

FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 Seats

Deployment Guide for up to 6700 Seat Virtual Desktop Infrastructure built on Cisco UCS B200 M5 and Cisco UCS Manager 4.0 with NetApp AFF A-Series on VMware Horizon View 7.10 and VMware vSphere ESXi 6.7 Update 2 Hypervisor Platform

Published: March 2020



In partnership with:



About the Cisco Validated Design Program

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

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, 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. Lisa.

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)

© 2020 Cisco Systems, Inc. All rights reserved.

Contents

Executive Summary	4
Solution Overview	5
Solution Summary	7
Solution Components	16
Architecture and Design Considerations for Desktop Virtualization	47
Deployment Hardware and Software	52
Solution Configuration	58
Test Setup, Configuration, and Load Recommendation	222
Test Methodology and Success Criteria	234
Test Results	242
Summary	289
Appendix	290
About the Authors	385
Feedback	386

Executive Summary

Cisco Validated Designs (CVDs) include 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 customers. Cisco and NetApp have partnered to deliver this document, which serves as a specific step-by-step guide for implementing this solution. This Cisco Validated Design provides an efficient architectural design that is based on customer requirements. The solution that follows is a validated approach for deploying Cisco, NetApp, and VMware technologies as a shared, high performance, resilient, virtual desktop infrastructure.

This document provides a reference architecture and design guide for a 5000 to 6000 seat desktop workload end user computing environment on FlexPod Datacenter with Cisco UCS and NetApp® AFF A300 and NetApp ONTAP® data management software. The solution includes VMware Horizon server-based RDS Windows Server 2019 sessions, VMware Horizon persistent full clone Microsoft Windows 10 virtual desktops and VMware Horizon non-persistent instant-clone Microsoft Windows 10 virtual desktops on VMware vSphere 6.7U2.

The solution is a predesigned, best-practice datacenter architecture built on the FlexPod reference architecture. The FlexPod Datacenter used in this validation includes Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9100 family of Fibre Channel (FC) switches and NetApp AFF A300 system.

This solution is 100 percent virtualized on fifth generation Cisco UCS B200 M5 blade servers, booting VMware vSphere 6.7 Update 2 through FC SAN with multipathing from the NetApp AFF A300 storage array. The virtual desktops are powered using VMware Horizon 7.10, with a linked clones Windows Server 2019 RDS desktops, instant clone non-persistent virtual Windows 10 desktops, and full clones persistent virtual Windows 10 desktops to support the user population and provisioned on the NetApp AFF A300 storage array. Where applicable the document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

This solution delivers the design for the following user payload on fifth generation Cisco UCS Blade Servers making it more efficient and cost effective in the datacenter due to increased solution security and density:

- 6,700 users based on the Windows Server 2019 sessions
- 5,400 instant clone non-persistent virtual Windows 10 desktop users
- 5,400 full clones persistent virtual Windows 10 desktop users
- 5,800 mixed users

Further rack efficiencies were gained from a storage standpoint as all 6000 users were hosted on a single 3U NetApp AFF A300 storage array while previous large-scale FlexPod Cisco Validated Designs with VDI used a NetApp 3U base chassis along with a 2U expansion shelf.

The solution is fully capable of supporting hardware accelerated graphics workloads. The Cisco UCS B200 M5 server supports up to two NVIDIA P6 cards for high density, high-performance graphics workload support. See our Cisco Graphics White Paper for details about integrating NVIDIA GPU with VMware Horizon.

This solution provides an outstanding virtual desktop end-user experience as measured by the Login VSI 4.1.25.6 Knowledge Worker workload running in benchmark mode, along with the solution providing a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

Solution Overview

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

Audience

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

Purpose of this document

This document provides a step-by-step design, configuration and implementation guide for the Cisco Validated Design for a large-scale VMware Horizon 7.10 mixed workload solution with NetApp AFF A300 array, Cisco UCS Blade Servers, Cisco Nexus 9000 series Ethernet switches and Cisco MDS 9100 series Multilayer Fibre channel switches.

What's new in this release?

This is the first VMware Horizon desktop virtualization Cisco Validated Design with Cisco UCS 5th generation servers with 2nd Generation Intel® Xeon® Scalable Processors and a NetApp AFF A-Series system.

It incorporates the following features:

- Cisco UCS Fabric Interconnect 6454
- Cisco UCS B200 M5 Blade Servers with 2nd Generation Intel Xeon Scalable Family Processors and 2933 MHz memory
- Validation of Cisco Nexus 9000 with NetApp AFF A300 System
- Validation of Cisco MDS 9100 with NetApp AFF A300 System
- Support for the UCS 4.0(4e) release and Cisco UCS B200-M5 Servers
- Support for the latest release of NetApp ONTAP® 9.6
- A Fibre Channel storage design supporting SAN LUNs
- Cisco UCS Inband KVM Access
- Cisco UCS vMedia client for vSphere Installation
- Cisco UCS Firmware Auto Sync Server Policy
- VMware vSphere 6.7 U2 Hypervisor
- VMware Horizon 7.10 Linked Clones Server 2019 RDS Hosted Server Sessions
- VMware Horizon 7.10 non-persistent Instant Clone Windows 10 Virtual Machines
- VMware Horizon 7.10 persistent Full Clones Windows 10 Virtual Machines

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

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

The use cases include:

- Enterprise Datacenter
- Service Provider Datacenter
- Large Commercial Datacenter

Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both VMware Horizon RDSH server desktop sessions based on Microsoft Server 2019, VMware Horizon VDI persistent virtual machines and VMware Horizon VDI non-persistent virtual machines based on Windows 10 operating system.

The mixed workload solution includes NetApp AFF A300 storage, Cisco Nexus® and MDS networking, the Cisco Unified Computing System (Cisco UCS®), VMware Horizon and VMware vSphere software in a single package. The design is space optimized such that the network, compute, and storage required can be housed in one datacenter rack. Switch port density enables the networking components to accommodate multiple compute and storage configurations of this kind.

The infrastructure is deployed to provide Fibre Channel-booted hosts with block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

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

- More power, same size. Cisco UCS B200 M5 half-width blade with dual 20-core 2.1 GHz Intel® Xeon® Gold (6230) processors and 768 GB (2933MHz) of memory for VMware Horizon Desktop hosts supports more virtual desktop workloads than the previously released generation processors on the same hardware. The 20-core 2.1 GHz Intel® Xeon® Gold (6230) processors used in this study provided a balance between increased per-blade capacity and cost.
- Fewer servers. Because of the increased compute power in the Cisco UCS B200 M5 servers, we supported the 5000 seat design with 16% fewer servers compared to previous generation Cisco UCS B200 M4s.
- Fault-tolerance with high availability built into the design. The various designs are based on using one Cisco Unified Computing System chassis with multiple Cisco UCS B200 M5 blade servers for virtualized desktop and infrastructure workloads. The design provides N+1 server fault tolerance for virtual desktops, RDSH desktops, and infrastructure services.
- Stress-tested to the limits during aggressive boot scenario. The RDS hosted virtual sessions and VDI pooled shared desktop environment booted and registered with the VMware Horizon 7 Administrator in under 20 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- Stress-tested to the limits during simulated login storms. All simulated users logged in and started running work-loads up to steady 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 system is less than a single 42U rack, conserving valuable datacenter floor space.

- All Virtualized: This Cisco Validated Design (CVD) presents a validated design that is 100 percent virtualized on VMware ESXi 6.7 U2. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, Provisioning Servers, SQL Servers, VMware Horizon Connection Servers, VMware Horizon Composer Server, VMware Horizon Replica Servers, VMware Horizon Remote Desktop Server Hosted sessions and VDI virtual machine desktops. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the FlexPod converged infrastructure with stateless Cisco UCS Blade Servers and NetApp storage.
- Cisco maintains industry leadership with the new Cisco UCS Manager 4.0(4e) 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 insure that customer environments are consistent locally, across Cisco UCS Domains and across the globe, our software suite offers increasingly simplified operational and deployment management and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage and network.
- Our 25G unified fabric story gets additional validation on Cisco UCS 6400 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- NetApp AFF A300 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 A300 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.
- NetApp Virtual Storage Console (VSC) for VMware vSphere hypervisor has deep integrations with vSphere, providing easy-button automation for key storage tasks such as storage repository provisioning, storage resize, data deduplication, directly from vCenter.
- VMware Horizon 7. Latest and greatest virtual desktop and application product. VMware Horizon 7 follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new VMware 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 users increase. In addition, Horizon 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 to achieve the best possible performance and scale. For RDSH desktop sessions, the best performance was achieved when the number of vCPUs assigned to the VMware 7 RDS 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. Remote Desktop Server Hosted (RDSH) shared virtual machines and VMware Horizon 7, Microsoft Windows 10 virtual machines were created for this solution using VMware Instant and Composer pooled desktops.

Cisco desktop virtualization solutions: Datacenter

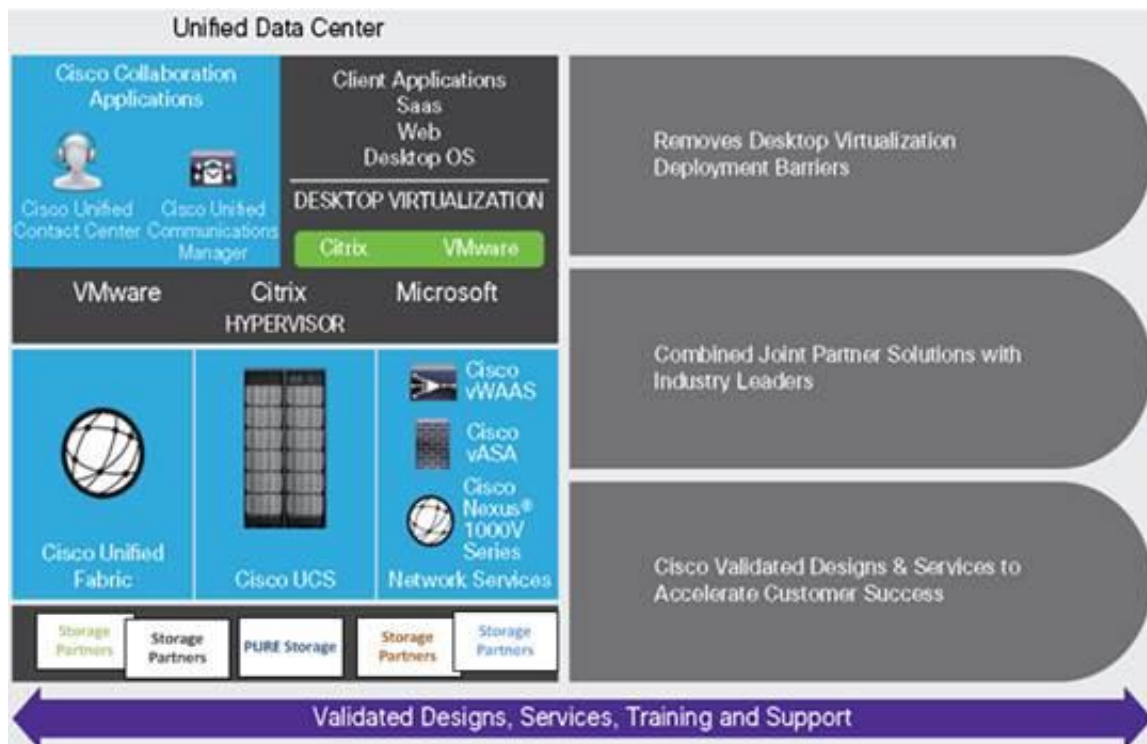
The evolving workplace

Today's IT departments are facing a rapidly evolving workplace environment. The workforce is becoming increasingly diverse and geographically dispersed, including offshore contractors, distributed call center operations, knowledge and task workers, partners, consultants, and executives connecting from locations around the world at all times.

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

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, specifically Microsoft Office 2016.

Figure 1. Cisco Datacenter Partner collaboration



Some of the key drivers for desktop virtualization are increased data security and reduced TCO through increased control and reduced management costs.

Cisco desktop virtualization focus

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

Simplified

Cisco UCS provides 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 and in the number of cables used per server, and the capability to rapidly deploy or re-provision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco UCS Manager Service Profiles and Cisco storage partners' storage-based cloning. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco UCS Manager automates many mundane, error-prone datacenter operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS B-Series Blade Servers and C-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 Technologies and NetApp have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlexPod. Cisco Desktop Virtualization Solutions have been tested with VMware vSphere, VMware Horizon.

Secure

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

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

Scalable

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 partner NetApp, helps 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 6700 desktops up and running in 20 minutes.

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

Savings and success

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

The simplified deployment of Cisco UCS 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.

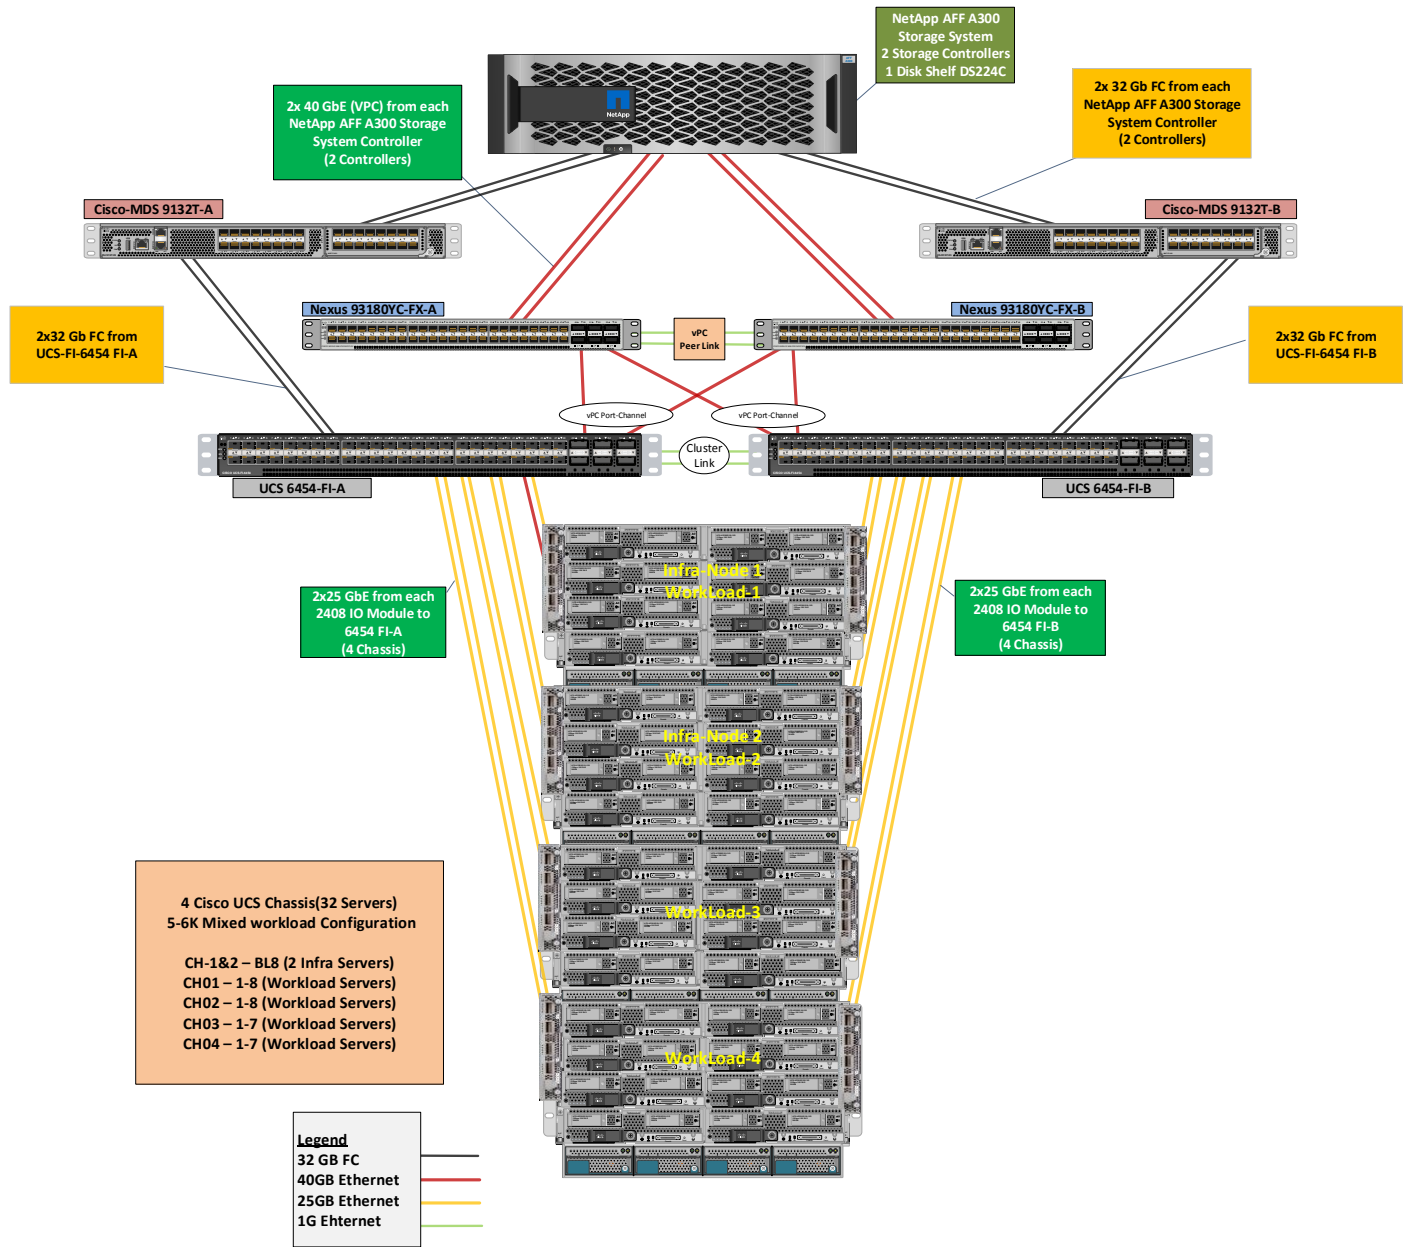
The ultimate measure of desktop virtualization for any organization is its efficiency and effectiveness in both the near term and the long term. The Cisco Desktop Virtualization Solutions are very efficient, allowing rapid deployment, requiring fewer devices and cables, and reducing costs. The solutions are also very 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 storage, 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.

Physical topology

High-scale mixed desktop workload solution reference architecture

Figure 2 illustrates the FlexPod System architecture used in this CVD to support very high scale mixed desktop user workload. It follows Cisco configuration requirements to deliver highly available and scalable architecture.

Figure 2. FlexPod solution reference architecture



The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco MDS 9132T 32-Gbps Fibre Channel switches
- Two Cisco UCS 6454 Fabric Interconnects
- Four Cisco UCS 5108 Blade Chassis
- Two Cisco UCS B200 M5 Blade Servers (for Infrastructure virtual machines)

-
- Thirty Cisco UCS B200 M5 Blade Servers (for workload)
 - One NetApp AFF A300 Storage System
 - One NetApp DS224C Disk Shelf

For desktop virtualization, the deployment includes VMware Horizon 7.10 running on VMware vSphere 6.7 Update 2.

The design is intended to provide a large-scale building block for VMware Horizon desktops in the following ratios:

- 6700 Random RDSH Windows 2019 (linked clone) user sessions with Office 2016
- 5400 Random Windows 10 Instant Clone Desktops with Office 2016
- 5400 Static Windows 10 Full Clone Desktops with Office 2016
- Mixed
 - 2440 Random RDSH Windows 2019 user sessions with Office 2016 (Linked Clones)
 - 1800 Random Windows 10 Instant Clone Desktops with Office 2016
 - 1800 Random Windows 10 Instant Clone Desktops with Office 2016

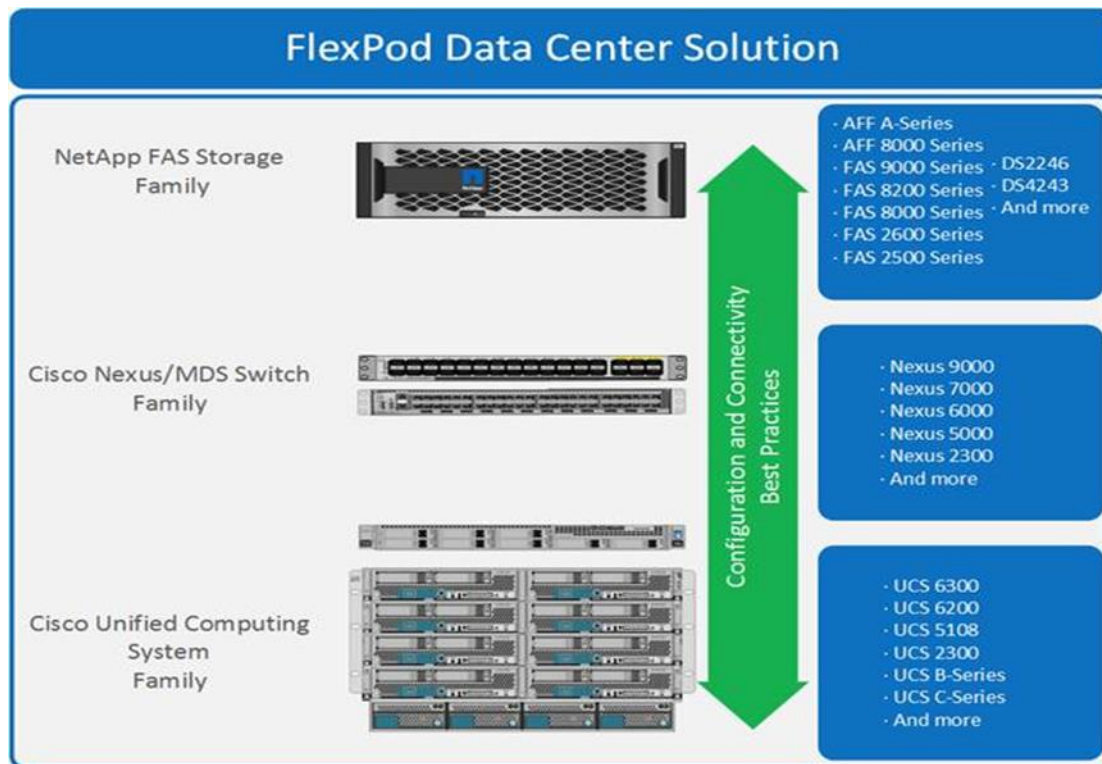
The data provided in this document will allow our customers to adjust the mix of the desktops to suit their environment. For example, additional blade servers and chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and 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 explains everything from physical cabling to network, compute and storage device configurations.

What is FlexPod?

FlexPod is a best practice datacenter architecture that includes the following components:

- Cisco Unified Computing System
- Cisco Nexus Switches
- Cisco MDS Switches
- NetApp AFF Systems
- FlexPod Component Families

Figure 3. FlexPod component families



These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed), or it can scale out for environments that require multiple consistent deployments (such as rolling out of additional FlexPod stacks). The reference architecture covered in this document leverages Cisco Nexus 9000 for the network switching element and pulls in the Cisco MDS 9100 for the SAN switching component.

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

Why FlexPod?

The following lists the benefits of FlexPod:

- Consistent Performance and Scalability
 - Consistent sub-millisecond latency with 100% flash storage
 - Consolidate 100's of enterprise-class applications in a single rack
 - Scales easily, without disruption
 - Continuous growth through multiple FlexPod CI deployments
- Operational Simplicity
 - Fully tested, validated, and documented for rapid deployment

-
- Reduced management complexity
 - Auto-aligned 512B architecture removes storage alignment issues
 - No storage tuning or tiers necessary
 - Lowest TCO
 - Dramatic savings in power, cooling, and space with 100 percent Flash
 - Industry leading data reduction
 - Enterprise-Grade Resiliency
 - Highly available architecture with no single point of failure
 - Nondisruptive operations with no downtime
 - Upgrade and expand without downtime or performance loss
 - Native data protection: snapshots and replication
 - Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases

Solution Components

Cisco Unified Computing System

Cisco UCS Manager (UCSM) provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) through an intuitive GUI, a CLI, and an XML API. The manager provides a unified management domain with centralized management capabilities and can control multiple chassis and thousands of virtual machines.

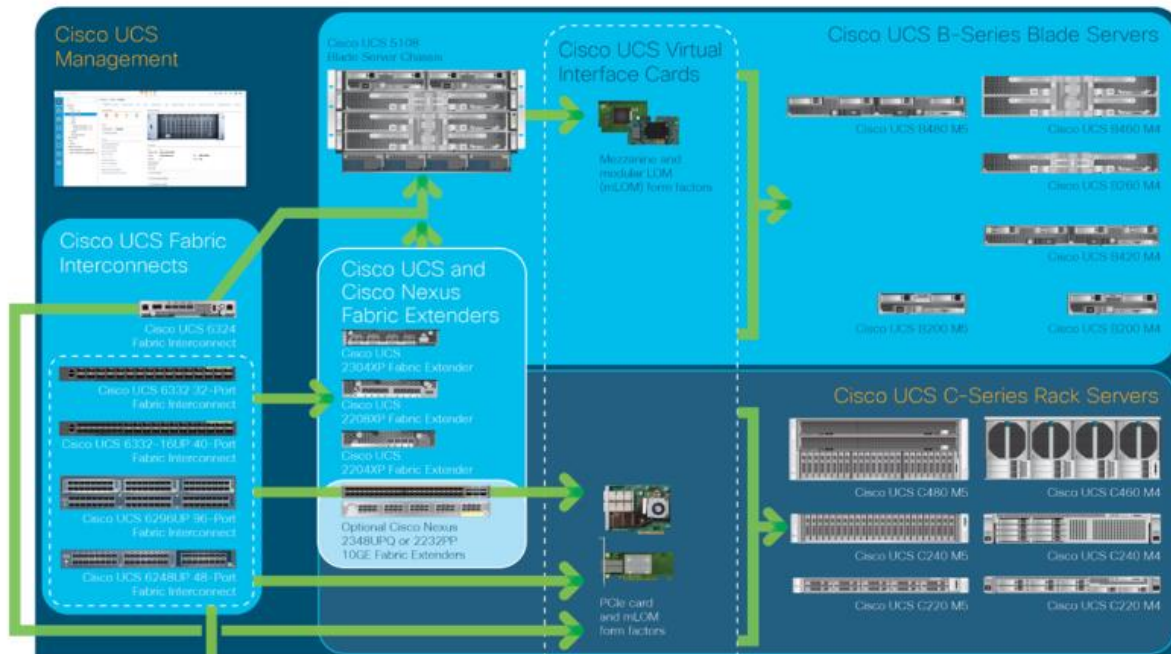
Cisco UCS is a next-generation data center platform that unites computing, networking, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 40 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Cisco Unified Computing System components

The main components of the Cisco UCS are:

- **Compute:** The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® Scalable Family processors.
- **Network:** The system is integrated on a low-latency, lossless, 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks to-day. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables needed, and by decreasing the power and cooling requirements.
- **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 virtualized environments to better support changing business and IT requirements.
- **Storage access:** The system provides consolidated access to local storage, SAN storage, and network-attached storage (NAS) over the unified fabric. With storage access unified, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Small Computer System Interface over IP (iSCSI) protocols. This capability provides customers with a choice for storage access and investment protection. In addition, server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management and helping increase productivity.
- **Management:** Cisco UCS uniquely integrates all system components, enabling the entire solution to be managed as a single entity by Cisco UCS Manager. Cisco UCS Manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

Figure 4. Cisco Datacenter overview



Cisco UCS is designed to deliver:

- Reduced TCO and increased business agility
- Increased IT staff productivity through just-in-time provisioning and mobility support
- A cohesive, integrated system that unifies the technology in the data center; the system is managed, serviced, and tested as a whole
- Scalability through design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand
- Industry standards supported by a partner ecosystem of industry leaders

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System across multiple chassis, rack servers, and thousands of virtual machines. Cisco UCS Manager manages Cisco UCS as a single entity through an intuitive GUI, a CLI, or an XML API for comprehensive access to all Cisco UCS Manager Functions.

Cisco UCS Fabric Interconnect

The Cisco UCS 6400 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6400 Series offers line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

The Cisco UCS 6400 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers, Cisco UCS 5108 B-Series Server Chassis, Cisco UCS Managed C-Series Rack Servers, and Cisco UCS S-Series Storage Servers. All servers attached to a Cisco UCS 6400 Series Fabric Interconnect

become part of a single, highly available management domain. In addition, by supporting a unified fabric, Cisco UCS 6400 Series Fabric Interconnect provides both the LAN and SAN connectivity for all servers within its domain.

From a networking perspective, the Cisco UCS 6400 Series use a cut-through architecture, supporting deterministic, low-latency, line-rate 10/25/40/100 Gigabit Ethernet ports, switching capacity of 3.82 Tbps for the 6454, 7.42 Tbps for the 64108, and 200 Gbps bandwidth between the Fabric Interconnect 6400 series and the IOM 2408 per 5108 blade chassis, independent of packet size and enabled services. 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 a lossless Ethernet fabric from the server through the fabric interconnect. 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.

Figure 5. Cisco UCS 6400 Series Fabric Interconnect - 6454 front view



Figure 6. Cisco UCS 6400 Series Fabric Interconnect - 6454 rear view




Cisco UCS B200 M5 Blade Server

The Cisco UCS B200 M5 Blade Server (Figure 7. and Figure 8.) is a density-optimized, half-width blade server that supports two CPU sockets for Intel Xeon processor 6140 Gold series CPUs and up to 24 DDR4 DIMMs. It supports one modular LAN-on-motherboard (LOM) dedicated slot for a Cisco virtual interface card (VIC) and one mezzanine adapter. In additions, the Cisco UCS B200 M5 supports an optional storage module that accommodates up to two SAS or SATA hard disk drives (HDDs) or solid-state disk (SSD) drives. You can install up to eight Cisco UCS B200 M5 servers in a chassis, mixing them with other models of Cisco UCS blade servers in the chassis if desired.

Figure 7. Cisco UCS B200 M5 front view



Figure 8. Cisco UCS B200 M5 back view



1	Asset pull tag Each server has a plastic tag that pulls out of the front panel. The tag contains the server serial number as well as the product ID (PID) and version ID (VID). The tag also allows you to add your own asset tracking label without interfering with the intended air flow.	7	Network link status
2	Blade ejector handle	8	Blade health LED
3	Ejector captive screw	9	Console connector ¹
4	Drive bay 1	10	Reset button access
5	Drive bay 2	11	Locater button and LED
6	Power button and LED		

Notes:

1. A KVM I/O Cable plugs into the console connector, it can be ordered as a spare. The KVM I/O Cable is included with every Cisco UCS 5100 Series blade server chassis accessory kit

Cisco UCS combines Cisco UCS B-Series Blade Servers and Cisco UCS C-Series Rack Servers with networking and storage access into a single converged system with simplified management, greater cost efficiency and agility, and increased visibility and control. The Cisco UCS B200 M5 Blade Server is one of the newest servers in the Cisco UCS portfolio.

The Cisco UCS B200 M5 delivers performance, flexibility, and optimization for data centers and remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads ranging from web infrastructure to distributed databases. The Cisco UCS B200 M5 can quickly deploy stateless physical and virtual workloads with the programmable ease of use of the Cisco UCS Manager software and simplified server access with Cisco® Single Connect technology. Based on the Intel Xeon processor 6140 Gold product family, it offers up to 3 TB of memory using 128GB DIMMs, up to two disk drives, and up to 320 Gbps of I/O throughput. The Cisco UCS B200 M5 offers exceptional levels of performance, flexibility, and I/O throughput to run your most demanding applications.

In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches, NICs, and HBAs in each blade server chassis. With a larger power budget per blade server, it provides uncompromised expandability and capabilities, as in the new Cisco UCS B200 M5 server with its leading memory-slot capacity and drive capacity.

The Cisco UCS B200 M5 provides:

- Latest Intel® Xeon® Scalable processors with up to 28 cores per socket

- Up to 24 DDR4 DIMMs for improved performance
- Intel 3D XPoint-ready support, with built-in support for next-generation nonvolatile memory technology
- Two GPUs
- Two Small-Form-Factor (SFF) drives
- Two Secure Digital (SD) cards or M.2 SATA drives
- Up to 80 Gbps of I/O throughput

Main features

The Cisco UCS B200 M5 server is a half-width blade. Up to eight servers can reside in the 6-Rack-Unit (6RU) Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit of blade chassis in the industry. You can configure the Cisco UCS B200 M5 to meet your local storage requirements without having to buy, power, and cool components that you do not need.

The Cisco UCS B200 M5 provides these main features:

- Up to two Intel Xeon Scalable CPUs with up to 28 cores per CPU
- 24 DIMM slots for industry-standard DDR4 memory at speeds up to 2666 MHz, with up to 3 TB of total memory when using 128-GB DIMMs
- Modular LAN On Motherboard (mLOM) card with Cisco UCS Virtual Interface Card (VIC) 1340, a 2-port, 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable mLOM mezzanine adapter
- Optional rear mezzanine VIC with two 40-Gbps unified I/O ports or two sets of 4 x 10-Gbps unified I/O ports, delivering 80 Gbps to the server; adapts to either 10- or 40-Gbps fabric connections
- Two optional, hot-pluggable, hard-disk drives (HDDs), solid-state drives (SSDs), or NVMe 2.5-inch drives with a choice of enterprise-class RAID or pass-through controllers
- Cisco FlexStorage local drive storage subsystem, which provides flexible boot and local storage capabilities and allows you to boot from dual, mirrored SD cards
- Support for up to two optional GPUs
- Support for up to one rear storage mezzanine card

For more information about Cisco UCS B200 M5, see the [Cisco UCS B200 M5 Blade Server Specs sheet](#)

Table 1. Ordering information

Part Number	Description
UCSB-B200-M5	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz
UCSB-B200-M5-U	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz (UPG)
UCSB-B200-M5-CH	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz,

Part Number	Description
	Drive bays, HS

Cisco UCS VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 9.) is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities are enabled for two ports of 40-Gbps Ethernet.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1340 supports Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

Figure 9. Cisco UCS VIC 1340




Figure 9. illustrates the Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M5 Blade Servers.

Cisco switching

Cisco Nexus 93180YC-FX switches

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

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

A Cisco 40 Gb bidirectional transceiver allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet Support for 1 Gb and 10 Gb 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-Gbps SFP+ ports
- 6 fixed 40/100-Gbps 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 10. Cisco Nexus 93180YC-EX Switch



Cisco MDS 9132T 32-Gbps Fiber Channel Switch

The next-generation Cisco® MDS 9132T 32-Gbps 32-Port Fibre Channel Switch (Figure 11.) 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-Gbps 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-Gbps) server rack gives you the option to upgrade to 32-Gbps server connectivity in the future using the 32-Gbps Host Bus Adapter (HBA) that are available today. The Cisco® MDS 9132T 32-Gbps 32-Port Fibre Channel switch also provides unmatched flexibility through a unique port expansion module (Figure 12.) 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 11. Cisco 9132T 32-Gbps MDS Fibre Channel Switch



Figure 12. Cisco MDS 9132T 32-Gbps 16-Port Fibre Channel Port Expansion Module



- Features
 - High performance: MDS 9132T architecture, with chip-integrated nonblocking arbitration, provides consistent 32-Gbps low-latency performance across all traffic conditions for every Fibre Channel port on the switch.
 - Capital Expenditure (CapEx) savings: The 32-Gbps ports allow users to deploy them on existing 16- or 8-Gbps transceivers, reducing initial CapEx with an option to upgrade to 32-Gbps 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 9100Family 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-Gbps 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-Gbps 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 allows 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-Gbps Cisco ASIC with an integrated network processor that powers the Cisco MDS 9700 48-Port 32-Gbps 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, smartzoning 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.

Hypervisor

This Cisco Validated Design includes VMware vSphere 6.7 Update2.

VMware vSphere 6.7

VMware provides virtualization software. VMware's enterprise software hypervisors for servers VMware vSphere ESX, vSphere ESXi, and vSphere—are bare-metal hypervisors that run directly on server hardware without requiring an additional underlying operating system. VMware vCenter Server for vSphere provides central management and complete control and visibility into clusters, hosts, virtual machines, storage, networking, and other critical elements of your virtual infrastructure.

VMware vSphere 6.7 introduces many enhancements to vSphere Hypervisor, VMware virtual machines, vCenter Server, virtual storage, and virtual networking, further extending the core capabilities of the vSphere platform.

Now VMware announced vSphere 6.7, which is one of the most feature rich releases of vSphere in quite some time. The vCenter Server Appliance is taking charge in this release with several new features which we'll cover in this blog article. For starters, the installer has gotten an overhaul with a new modern look and feel. Users of both Linux and Mac will also be ecstatic since the installer is now supported on those platforms along with Microsoft Windows. If that wasn't enough, the vCenter Server Appliance now has features that are exclusive such as:

- Migration
- Improved Appliance Management
- VMware Update Manager
- Native High Availability
- Built-in Backup / Restore

VMware vSphere Client

With VMware vSphere 6.7, a fully supported version of the HTML5-based vSphere Client that will run alongside the vSphere Web Client. The vSphere Client is built into vCenter Server 6.7 (both Windows and Appliance) and

is enabled by default. While the HTML-5 based vSphere Client does not have full feature parity, the team has prioritized many of the day-to-day tasks of administrators and continue to seek feedback on items that will enable customers to use it full time. The vSphere Web Client continues to be accessible through “http://<vcenter_fqdn>/vsphere-client” while the vSphere Client is reachable through “http://<vcenter_fqdn>/ui”. VMware is periodically updating the vSphere Client outside of the normal vCenter Server release cycle. To make sure it is easy and simple for customers to stay up to date the vSphere Client will be able to be updated without any effects to the rest of vCenter Server.

Some of the benefits of the new VMware vSphere Client:

- Clean, consistent UI built on VMware’s new Clarity UI standards (to be adopted across our portfolio)
- Built on HTML5 so it is truly a cross-browser and cross-platform application
- No browser plugins to install/manage
- Integrated into vCenter Server for 6.7 and fully supported
- Fully supports Enhanced Linked Mode
- Users of the Fling have been extremely positive about its performance

VMware ESXi 6.7 Hypervisor

VMware vSphere 6.7 introduces the following new features in the hypervisor:

- Scalability Improvements
 - ESXi 6.7 dramatically increases the scalability of the platform. With vSphere Hypervisor 6.0, clusters can scale to as many as 64 hosts, up from 32 in previous releases. With 64 hosts in a cluster, vSphere 6.0 can support 8000 virtual machines in a single cluster. This capability enables greater consolidation ratios, more efficient use of VMware vSphere Distributed Resource Scheduler (DRS), and fewer clusters that must be separately managed. Each vSphere Hypervisor 6.7 instance can support up to 480 logical CPUs, 12 terabytes (TB) of RAM, and 1024 virtual machines. By using the newest hardware advances, ESXi 6.7 enables the virtualization of applications that previously had been thought to be non-virtualizable.
- ESXi 6.7 security enhancements
 - Account management: ESXi 6.7 enables management of local accounts on the ESXi server using new ESXi CLI commands. The capability to add, list, remove, and modify accounts across all hosts in a cluster can be centrally managed using a vCenter Server system. Previously, the account and permission management functions for ESXi hosts were available only for direct host connections. The setup, removal, and listing of local permissions on ESXi servers can also be centrally managed.
 - Account lockout: ESXi Host Advanced System Settings have two new options for the management of failed local account login attempts and account lockout duration. These parameters affect Secure Shell (SSH) and vSphere Web Services connections, but not ESXi direct console user interface (DCUI) or console shell access.
 - Password complexity rules: In previous versions of ESXi, password complexity changes had to be made by manually editing the /etc/pam.d/passwd file on each ESXi host. In vSphere 6.0, an entry in Host Advanced System Settings enables changes to be centrally managed for all hosts in a cluster.

- Improved auditability of ESXi administrator actions: Prior to vSphere 6.0, actions at the vCenter Server level by a named user appeared in ESXi logs with the vpxuser username: for example, [user=vpxuser]. In vSphere 6.7, all actions at the vCenter Server level for an ESXi server appear in the ESXi logs with the vCenter Server username: for example, [user=vpxuser: DOMAIN\User]. This approach provides a better audit trail for actions run on a vCenter Server instance that conducted corresponding tasks on the ESXi hosts.
- Flexible lockdown modes: Prior to vSphere 6.7, only one lockdown mode was available. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.7, two lockdown modes are available:

In normal lockdown mode, DCUI access is not stopped, and users on the DCUI access list can access the DCUI.

In strict lockdown mode, the DCUI is stopped.
- Exception users: vSphere 6.0 offers a new function called exception users. Exception users are local accounts or Microsoft Active Directory accounts with permissions defined locally on the host to which these users have host access. These exception users are not recommended for general user accounts, but they are recommended for use by third-party applications—for service accounts, for example—that need host access when either normal or strict lockdown mode is enabled. Permissions on these accounts should be set to the bare minimum required for the application to perform its task and with an account that needs only read-only permissions on the ESXi host.
- Smart card authentication to DCUI: This function is for U.S. federal customers only. It enables DCUI login access using a Common Access Card (CAC) and Personal Identity Verification (PIV). The ESXi host must be part of an Active Directory domain.

Desktop Broker

This Cisco Validated Design includes VMware Horizon 7.10.

VMware Horizon Version 7

Enterprise IT organizations are tasked with the challenge of provisioning Microsoft Windows apps and desktops while managing cost, centralizing control, and enforcing corporate security policy. Deploying Windows apps to users in any location, regardless of the device type and available network bandwidth, enables a mobile workforce that can improve productivity. With VMware Horizon 7, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

VMware Horizon

VMware Horizon desktop virtualization solutions are built on a unified architecture so they are simple to manage and flexible enough to meet the needs of all your organization's users. You use the same architecture and management tools to manage public, private, and hybrid cloud deployments as you do for on-premises deployments.

VMware Horizon Virtual machines and RDSH known as server-based hosted sessions: These are applications hosted from Microsoft Windows servers to any type of device, including Windows PCs, Macs, smartphones, and tablets. Some VMware editions include technologies that further optimize the experience of using Windows applications on a mobile device by automatically translating native mobile-device display, navigation, and controls to Windows applications; enhancing performance over mobile networks; and enabling developers to optimize any custom Windows application for any mobile environment.

VMware Horizon RDSH session users also known as server-hosted desktops: These are inexpensive, locked-down Windows virtual desktops hosted from Windows server operating systems. They are well suited for users, such as call center employees, who perform a standard set of tasks.

Advantages of using VMware Horizon

VMware Horizon 7 provides the following new features and enhancements:

- Instant Clones
 - A new type of desktop virtual machines that can be provisioned significantly faster than the traditional Composer linked clones.
 - A fully functional desktop can be provisioned in two seconds or less.
 - Recreating a desktop pool with a new OS image can be accomplished in a fraction of the time it takes a Composer desktop pool because the parent image can be prepared well ahead of the scheduled time of pool recreation.
 - Clones are automatically rebalanced across available datastores.
 - View storage accelerator is automatically enabled.
- VMware Blast Extreme
 - VMware Blast Extreme is now fully supported on the Horizon platform.
 - Administrators can select the VMware Blast display protocol as the default or available protocol for pools, farms, and entitlements.
 - End users can select the VMware Blast display protocol when connecting to remote desktops and applications.
- VMware Blast Extreme features include:
 - TCP and UDP transport support.
 - H.264 support for the best performance across more devices.
 - Reduced device power consumption for longer battery life.
 - NVIDIA GRID acceleration for more graphical workloads per server, better performance, and a superior remote user experience.
- True SSO

For VMware Identity Manager integration, True SSO streamlines the end-to-end login experience. After users log in to VMware Identity Manager using a smart card or an RSA SecurID or RADIUS token, users are not required to also enter Active Directory credentials in order to use a remote desktop or application.

 - Uses a short-lived Horizon virtual certificate to enable a password-free Windows login.
 - Supports using either a native Horizon Client or HTML Access.
 - System health status for True SSO appears in the Horizon Administrator dashboard.
 - Can be used in a single domain, in a single forest with multiple domains, and in a multiple-forest, multiple-domain setup.
- Smart policies

-
- Control of the clipboard cut-and-paste, client drive redirection, USB redirection, and virtual printing desktop features through defined policies.
 - PCoIP session control through PCoIP profiles.
 - Conditional policies based on user location, desktop tagging, pool name, and Horizon Client registry values.
 - Configure the clipboard memory size for VMware Blast and PCoIP sessions.
 - Horizon administrators can configure the server clipboard memory size by setting GPOs for VMware Blast and PCoIP sessions. Horizon Client 4.1 users on Windows, Linux, and Mac OS X systems can configure the client clipboard memory size. The effective memory size is the lesser of the server and client clipboard memory size values.
 - VMware Blast network recovery enhancements
 - Network recovery is now supported for VMware Blast sessions initiated from iOS, Android, Mac OS X, Linux, and Chrome OS clients. Previously, network recovery was supported only for Windows client sessions. If you lose your network connection unexpectedly during a VMware Blast session, Horizon Client attempts to reconnect to the network, and you can continue to use your remote desktop or application. The network recovery feature also supports IP roaming, which means you can resume your VMware Blast session after switching to a WiFi network.
 - Configure Horizon Administrator to not remember the login name.
 - Horizon administrators can configure not to display the Remember user name check box and therefore not remember the administrator's login name.
 - Allow Mac OS X users to save credentials.
 - Horizon administrators can configure Connection Server to allow Horizon Client Mac OS X systems to remember a user's user name, password, and domain information. If users choose to have their credentials saved, the credentials are added to the login fields in Horizon Client on subsequent connections.
 - Microsoft Windows 10
 - Windows 10 is supported as a desktop guest operating system.
 - Horizon Client runs on Windows 10.
 - Smart card is supported on Windows 10.
 - The Horizon User Profile Migration tool migrates Windows 7, 8/8.1, Server 2008 R2, or Server 2012 R2 user profiles to Windows 10 user profiles.
 - RDS desktops and hosted apps
 - View Composer. View Composer and linked clones provide automated and efficient management of RDS server farms.
 - Graphics Support. Existing 3D vDGA and GRID vGPU graphics solutions on VDI desktops have been extended to RDS hosts, enabling graphics-intensive applications to run on RDS desktops and Hosted Apps.
 - Enhanced Load Balancing. A new capability provides load balancing of server farm applications based on memory and CPU resources.

-
- One-Way AD Trusts. One-way AD trust domains are now supported. This feature enables environments with limited trust relationships between domains without requiring Horizon Connection Server to be in an external domain.
 - Cloud Pod Architecture (CPA) enhancements
 - Hosted App Support. Support for application remoting allows applications to be launched using global entitlements across a pod federation.
 - HTML Access (Blast) Support. Users can use HTML Access to connect to remote desktops and applications in a Cloud Pod Architecture deployment.

- Access Point integration

Access Point is a hardened Linux-based virtual appliance that protects virtual desktop and application resources to allow secure remote access from the Internet. Access Point provides a new authenticating DMZ gateway to Horizon Connection Server. Smart card support on Access Point is available as a Tech Preview. Security server will continue to be available as an alternative configuration. For more information, see [Deploying and Configuring Access Point](#).

- FIPS

Install-time FIPS mode allows customers with high security requirements to deploy Horizon 6.

- Graphics enhancements

- AMD vDGA enables vDGA pass-through graphics for AMD graphics hardware.
- 4K resolution monitors (3840x2160) are supported.

- Horizon Administrator enhancements

- Horizon Administrator shows additional licensing information, including license key, named user and concurrent connection user count.
- Pool creation is streamlined by letting Horizon administrators to clone existing pools.

- Additional features

- Support for IPv6 with VMware Blast Extreme on security servers.
- Horizon Administrator security protection layer. See VMware Knowledge Base (KB) article 2144303 for more information:
https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2144303
- Protection against inadvertent pool deletion.
- RDS per-device licensing improvements.
- Support for Intel vDGA.
- Support for AMD Multiuser GPU Using vDGA.
- More resilient upgrades.
- Display scaling for Windows Horizon Clients.
- DPI scaling is supported if it is set at the system level and the scaling level is greater than 100.

What are VMware RDS hosted sessions?

An RDS host is a server computer that hosts applications and desktop sessions for remote access. An RDS host can be a virtual machine or a physical server.

An RDS host has the Microsoft Remote Desktop Services role, the Microsoft Remote Desktop Session Host service, and Horizon Agent installed. Remote Desktop Services was previously known as Terminal Services. The Remote Desktop Session Host service allows a server to host applications and remote desktop sessions. With Horizon Agent installed on an RDS host, users can connect to applications and desktop sessions by using the display protocol PCoIP or Blast Extreme. Both protocols provide an optimized user experience for the delivery of remote content, including images, audio and video.

The performance of an RDS host depends on many factors. For information on how to tune the performance of different versions of Windows Server, see <http://msdn.microsoft.com/library/windows/hardware/gg463392.aspx>.

Horizon 7 supports at most one desktop session and one application session per user on an RDS host.

When users submit print jobs concurrently from RDS desktops or applications that are hosted on the same RDS host, the ThinPrint server on the RDS host processes the print requests serially rather than in parallel. This can cause a delay for some users. Note that the print server does not wait for a print job to complete before processing the next one. Print jobs that are sent to different printers will print in parallel.

If a user launches an application and also an RDS desktop, and both are hosted on the same RDS host, they share the same user profile. If the user launches an application from the desktop, conflicts may result if both applications try to access or modify the same parts of the user profile, and one of the applications may fail to run properly.

The process of setting up applications or RDS desktops for remote access involves installing applications. If you plan to create application pools, you must install the applications on the RDS hosts. If you want Horizon 7 to automatically display the list of installed applications, you must install the applications so that they are available to all users from the Start menu. You can install an application at any time before you create the application pool. If you plan to manually specify an application, you can install the application at any time, either before or after creating an application pool.

IMPORTANT

When you install an application, you must install it on all the RDS hosts in a farm and in the same location on each RDS host. If you do not, a health warning will appear on the View Administrator dashboard. In such a situation, if you create an application pool, users might encounter an error when they try to run the application.

When you create an application pool, Horizon 7 automatically displays the applications that are available to all users rather than individual users from the Start menu on all of the RDS hosts in a farm. You can choose any applications from that list. In addition, you can manually specify an application that is not available to all users from the Start menu. There is no limit on the number of applications that you can install on an RDS host.

Farms, RDS Hosts, and Desktop and Application Pools

With VMware Horizon, you can create desktop and application pools to give users remote access to virtual machine-based desktops, session-based desktops, physical computers, and applications. Horizon takes advantage

of Microsoft Remote Desktop Services (RDS) and VMware PC-over-IP (PCoIP) technologies to provide high-quality remote access to users.

- RDS Hosts

RDS hosts are server computers that have Windows Remote Desktop Services and View Agent installed. These servers host applications and desktop sessions that users can access remotely. To use RDS desktop pools or applications, your end users must have access to Horizon Client 3.0 or later software.

- Desktop Pools

There are three types of desktop pools: automated, manual, and RDS. Automated desktop pools use a vCenter Server virtual machine template or snapshot to create a pool of identical virtual machines. Manual desktop pools are a collection of existing vCenter Server virtual machines, physical computers, or third-party virtual machines. In automated or manual pools, each machine is available for one user to access remotely at a time. RDS desktop pools are not a collection of machines, but instead, provide users with desktop sessions on RDS hosts. Multiple users can have desktop sessions on an RDS host simultaneously.

- Application Pools

Application pools let you deliver applications to many users. The applications in application pools run on a farm of RDS hosts.

- Farms

Farms are collections of RDS hosts and facilitate the management of those hosts. Farms can have a variable number of RDS hosts and provide a common set of applications or RDS desktops to users. When you create an RDS desktop pool or an application pool, you must specify a farm. The RDS hosts in the farm provide desktop and application sessions to users.

Some of the latest VMware Horizon features and enhancements are:

- Flash Redirection

You can compile a black list to ensure that the URLs specified in the list will not be able to redirect Flash content. You must enable the GPO setting `FlashMMRUrlListEnableType` to use either a white list or black list.

- Horizon Agent Policy Settings

The `VMwareAgentCIT` policy setting enables remote connections to Internet Explorer to use the Client's IP address instead of the IP address of the remote desktop machine.

The `FlashMMRUrlListEnableType` and `FlashMMRUrlList` policy settings specify and control the white list or black list that enables or disables the list of URLs from using Flash Redirection.

- Horizon PowerCLI

View PowerCLI is deprecated. Horizon PowerCLI replaces View PowerCLI and includes cmdlets that you can use with VMware PowerCLI.

For more information about Horizon PowerCLI cmdlets, read the [VMware PowerCLI Cmdlets Reference](#).

For information about the API specifications to create advanced functions and scripts to use with Horizon PowerCLI, see the API Reference at the [VMware Developer Center](#).

For more information about sample scripts that you can use to create your own Horizon PowerCLI scripts, see the [Horizon PowerCLI community on GitHub](#).

Horizon 7 for Linux desktops enhancements:

- UDP based Blast Extreme connectivity

User Datagram Protocol (UDP) is enabled by default in both the client and the agent. Note that Transmission Control Protocol (TCP) connectivity will have a better performance than UDP on the Local Area Network (LAN). UDP will have better performance than TCP over Wide Area Network (WAN). If you are on a LAN, disable the UDP feature to switch to using TCP to get better connectivity performance.

- KDE support

K Desktop Environment (KDE) support is now also available on CentOS 7, RHEL 7, Ubuntu 14.04, Ubuntu 16.04, and SLED 11 SP4 platforms.

- MATE support

MATE desktop environment is supported on Ubuntu 14.04 and 16.04 virtual machines.

- Hardware H.264 Encoder

The hardware H.264 encoder is now available and used when the vGPU is configured with the NVIDIA graphics card that has the NVIDIA driver 384 series or later installed on it.

- Additional platforms support

- RHEL 7.4 x64 and CentOS 7.4 x64 are now supported.
- Remote Desktop Operating System
- Windows 10 version 1607 Long-Term Servicing Branch (LTSB)
- Windows Server 2016

Horizon Agent

- HTML5 Multimedia Redirection

You can install the HTML5 Multimedia Redirection feature by selecting the HTML5 Multimedia Redirection custom setup option in the Horizon Agent installer. With HTML5 Multimedia Redirection, if an end user uses the Chrome browser, HTML5 multimedia content is sent from the remote desktop to the client system, reducing the load on the ESXi host. The client system plays the multimedia content and the user has a better audio and video experience.

- SHA-256 support

Horizon Agent has been updated to support the SHA-256 cryptographic hash algorithm. SHA-256 is also supported in Horizon Client 4.6 and Horizon 7 version 7.2 and later.

- Improved USB redirection with User Environment Manager

The default User Environment Manager timeout value has been increased. This change makes sure that the USB redirection smart policy takes effect even when the login process takes a long time. With Horizon Client 4.6, the User Environment Manager timeout value is configured only on the agent and is sent from the agent to the client.

You can now bypass User Environment Manager control of USB redirection by setting a registry key on the agent machine. This change helps ensure that smart card SSO works on Teradici zero clients.

- Composer

For enhanced security, you can enable the digest access authentication method for Composer.

- Persona Management

Persona Management supports guest operating systems that use the "v6" version of the user profile.

You can use the migration tool to migrate the "v2" and "v5" user profiles versions to the "v6" user profile version. The tool is installed with the Persona binary file.

Horizon Connection Server enhanced features

- Horizon Help Desk Tool

View application and process names and resource use within a virtual or published desktop to identify which applications and process are using up machine resources.

View event log information about the user's activities.

View updated metrics such as Horizon Client version and the Blast protocol.

View additional session metrics such as the VM information, CPU, or memory usage.

You can assign predefined administrator roles to Horizon Help Desk Tool administrators to delegate the troubleshooting tasks between administrator users. You can also create custom roles and add privileges based on the predefined administrator roles.

You can verify the product license key for Horizon Help Desk Tool and apply a valid license.

- Monitoring

If the event database shuts down, Horizon Administrator maintains an audit trail of the events that occur before and after the event database shutdown.

- Instant Clones

You can create dedicated instant-clone desktop pools.

Windows Server operating systems are supported for instant clones in this release. For an updated list of supported Windows Server operating systems, see the VMware Knowledge Base (KB) article [2150295](#).

You can copy, paste, or enter the path for the AD tree in the AD container field when you create an instant-clone desktop pool.

If there are no internal VMs in all four internal folders created in vSphere Web Client, these folders are unprotected and you can delete these folders.

You can use the enhanced instant-clone maintenance utility `lcUnprotect.cmd` to unprotect or delete template, replica, or parent VMs or folders from vSphere hosts.

Instant clones are compatible with Storage DRS (sDRS). Therefore, instant clones can reside in a datastore that is part of an sDRS cluster.

- Cloud Pod Architecture

The total session limit is increased to 140,000.

The site limit is increased to 7.

You can configure Windows Start menu shortcuts for global entitlements. When an entitled user connects to a Connection Server instance in the pod federation, Horizon Client for Windows places these shortcuts in the Start menu on the user's Windows client device.

- Published Desktops and Application Pools

You can restrict access to entitled desktop pools, application pools, global entitlements, and global application entitlements from certain client computers.

You can configure Windows start menu shortcuts for entitled desktop and application pools. When an entitled user connects to a Connection Server instance, Horizon Client for Windows places these shortcuts in the Start menu on the user's Windows client device.

- Virtual Desktops and Desktop Pools

Blast Extreme provides network continuity during momentary network loss on Windows clients.

Performance counters displayed using PerfMon on Windows agents for Blast session, imaging, audio, CDR, USB, and virtual printing provide an accurate representation of the current state of the system that also updates at a constant rate.

- Customer Experience Improvement Program

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the [Trust Assurance Center](#).

- Security

With the USB over Session Enhancement SDK feature, you do not need to open TCP port 32111 for USB traffic in a DMZ-based security server deployment. This feature is supported for both virtual desktops and published desktops on RDS hosts.

- Database Support

The Always On Availability Groups feature for Microsoft SQL Server 2014 is supported in this release of Horizon 7. For more information, refer to the [Release Notes](#).

Supported Windows Operating Systems

Horizon 7 version 7.10 supports the following Windows 10 operating systems:

- Win 10 1607 LTSB (Enterprise)
- Win 10 1703 CBB / Semi-Annual Channel (broad deployment) (Enterprise, Professional, Education)

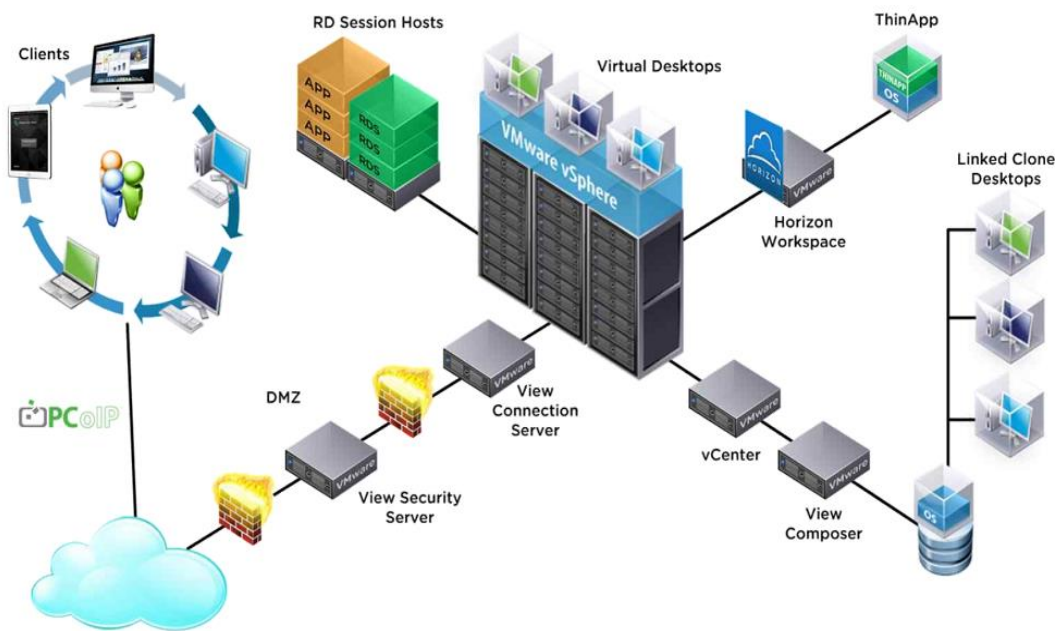
- Win 10 1709 Semi-Annual Channel (broad deployment) (Enterprise, Professional, Education)
- Windows 10 1809 LTSC (Enterprise)
- Windows 10 1903 SAC (Pro, Education, Enterprise)
- Windows 10 1909 SAC (Pro, Education, Enterprise)

For the complete list of supported Windows on Horizon including all VDI (Full Clones, Linked and Instant clones) click the following links: <https://kb.vmware.com/s/article/2149393> and https://kb.vmware.com/s/article/2150295?r=2&Quarterback.validateRoute=1&KM_Utility.getArticleData=1&KM_Utility.getUser=1&KM_Utility.getArticleLanguage=2&KM_Utility.getArticle=1

Note

Windows 10 version 1809 is used in this solution.

Figure 13. Logical architecture of VMware Horizon




VMware Horizon Composer

VMware Horizon Composer is a feature in Horizon that gives administrators the ability to manage virtual machine pools or the desktop pools that share a common [virtual disk](#). An administrator can update the [master image](#), then all desktops using [linked clones](#) of that master image can also be patched. Updating the master image will patch the cloned desktops of the users without touching their applications, data or settings.

The VMware View Composer pooled desktops solution’s infrastructure is based on software-streaming technology. After installing and configuring the composed pooled desktops, a single shared disk image (Master Image) is taken a snapshot of the OS and application image, and then storing that snapshot file accessible to host(s).

Figure 14. VMware Horizon Composer



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 (BYO) device to work programs are prime reasons for moving to a virtual desktop solution.

VMware Horizon design fundamentals

VMware Horizon 7 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 VDI Pool and RDSH Server pool

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 Composer aligning with VMware Horizon Connection Server with vCenter Server components. In this CVD, machines in the Pools are configured to run either a Windows Server 2019 OS (for RDS Hosted shared sessions) and a Windows 10 Desktop OS (for pooled VDI desktops).

Figure 15. VMware Horizon design overview


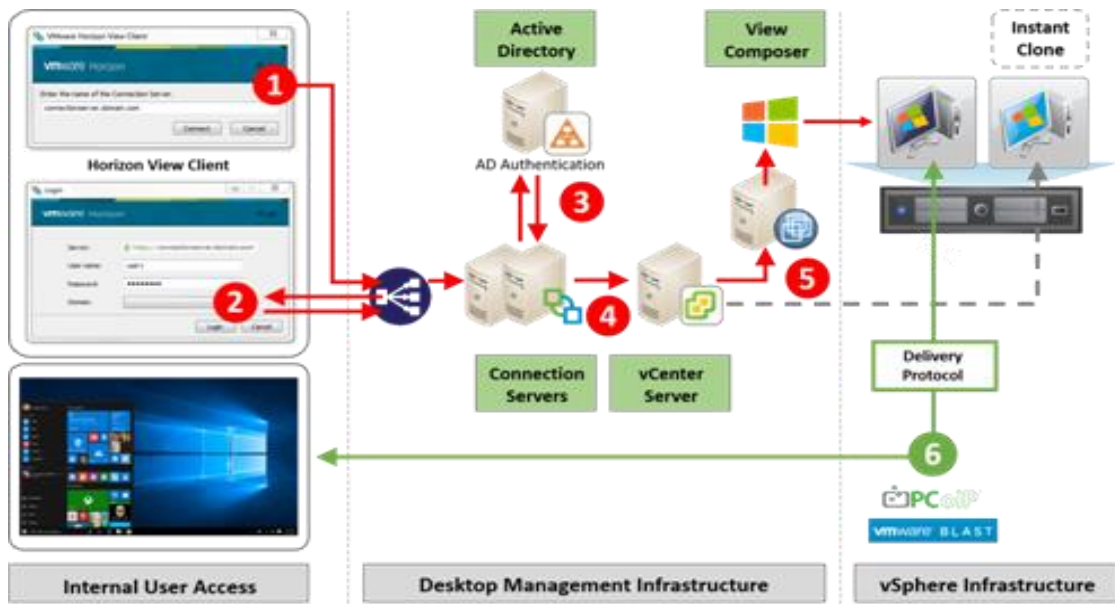


Figure 16. Horizon VDI and RDSH desktop delivery based on display protocol (PCoIP/Blast/RDP)



NetApp AFF A-Series Systems

With the new AFF A-Series controllers, NetApp provides industry-leading performance while continuing to provide a full suite of enterprise-grade data management and data protection features. The AFF A-Series systems offer double the IOPS, while decreasing the latency. The AFF A-Series lineup includes the A220, A300, A400, A700, and A800. These controllers and their specifications are listed in Table 2. For more information about the AFF A-Series controllers, see:

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

<https://hww.NetApp.com/Controller/Index>

Table 2. NetApp AFF A-series controller specifications

	AFF A220	AFF A300	AFF A400	AFF A700	AFF A800
NAS scale-out	2-24 nodes	2-24 nodes	2-24 nodes	2-24 nodes	2-24 nodes
SAN scale-out	2-12 nodes	2-12 nodes	2-12 nodes	2-12 nodes	2-12 nodes
Per HA Pair Specifications (Active-Active Dual Controller)					
Maximum SSDs	1728	4608	5760	5760	2880
Maximum raw capacity	4.3PB	11.47PB	14.34PB	14.34PB	6.45PB
Effective capacity	193.3PB	562.2PB	702.7PB	702.7PB	316.3PB
Chassis form factor	2U chassis with two HA controllers and 24 SSD slots	3U chassis with two HA controllers	4U chassis with two HA controllers	8U chassis with two HA controllers	4U chassis with two HA controllers and 24 SSD slots

This solution uses the NetApp AFF A300 controller, as shown in Figure 17. and Figure 18. This controller provides the high-performance benefits of 40GbE and all-flash SSDs, offering better performance than previous models, and occupying only 3U of rack space versus 6U with the AFF8040 systems. When combined with the 2U disk shelf of 3.48TB disks, this solution provides ample horsepower and over 83TB of raw capacity, all while occupying only 5U of valuable rack space. This configuration makes it an ideal controller for a shared workload converged infrastructure. The AFF A800 controller would be an ideal fit for situations in which more performance is required.

The FlexPod reference architecture supports a variety of NetApp FAS controllers such as FAS9000, FAS8000, FAS2600 and FAS2500; AFF A-Series platforms such as AFF8000; and legacy NetApp storage.

For more information about the AFF A-Series product family, see:

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

Note

The 40GbE cards are installed in the expansion slot 2 and the ports are e2a, e2e.

Figure 17. NetApp AFF A300 front view





Figure 18. NetApp AFF A300 rear view



NetApp ONTAP 9.6

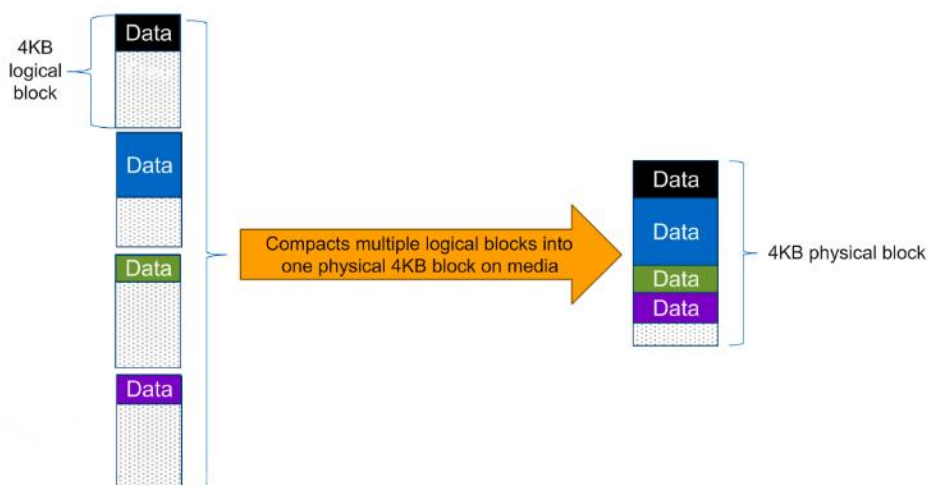
Storage efficiency

Storage efficiency has always been a primary architectural design point of the NetApp ONTAP data management software. 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, multitenancy, 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 was 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 19.

Figure 19. Storage efficiency



NetApp Storage Virtual Machines

A cluster serves data through at least one and possibly multiple storage virtual machines (SVMs, formerly called Vservers). An SVM is a logical abstraction that represents the set of physical resources of the cluster. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and may can on any node in the cluster to which the SVM has been given access. An SVM can own resources on multiple nodes concurrently, and those resources can be moved nondisruptively from one node to another. For example, a flexible volume can be nondisruptively 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 NFS and CIFS clients. 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.

Because it is a secure entity, an SVM is only aware of the resources that are assigned to it and has no knowledge of other SVMs and their respective resources. Each SVM operates as a separate and distinct entity with its own security domain. Tenants can manage the resources allocated to them through a delegated SVM administration account. Each SVM can connect to unique authentication zones such as Active Directory, LDAP, or NIS. A NetApp cluster can contain multiple SVMs. If you have multiple SVMs, you can delegate an SVM to a


specific application. This allows administrators of the application to access only the dedicated SVMs and associated storage, increasing manageability, and reducing risk.

SAN Boot

NetApp recommends implementing SAN boot for Cisco UCS servers in the FlexPod Datacenter solution. Doing so enables the operating system to be safely secured by the NetApp AFF 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 AFF storage to the Cisco MDS switch. The 16G 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. Refer to Figure 20. for the port and LIF layout


Figure 20. 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 host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp that is used to provide information about SCSI targets, which allows a host to identify the best path to the storage.

NetApp ONTAP FlexGroup

NetApp ONTAP FlexGroup volumes make it easier to expand the file share capacity or performance needs by adding more NetApp FlexVol® member volumes, which can span across multiple storage nodes. Files are not striped but instead are placed systematically into individual FlexVol member volumes that work together under a single access point.




ONTAP FlexGroup technology supports features such as SMB 3.x multichannel, antivirus scanning for SMB, volume autosize autogrow/autoshrink, and qtree (for quotas).

For more information, see [TR-4557: NetApp ONTAP FlexGroup Volumes](#).


Quality of Service

Storage Quality of Service (QoS) guarantees performance of critical workloads that are not degraded by competing workloads. A throughput ceiling on a competing workload to limit its impact on system resources, or set a throughput floor for a critical workload, ensuring that it meets minimum throughput targets, regardless of demand by competing workloads.

Before QoS



After QoS — Ceiling



Starting with ONTAP 9.4, a non-shared QoS policy group can be used to specify the defined throughput ceiling or floor that applies to each member workload individually. A Shared policy group with throughput ceilings, the total throughput for the workloads assigned to the shared policy cannot exceed the specified ceiling. The throughput floors shared policy group can be applied to single workload only.

Adaptive Quality of Service

Adaptive quality of service (AQoS) enables an IT organization to deliver predictable and agile services, aligning performance and cost to the needs of the applications. It does this by enabling a range of service levels aligned to application I/O requirements.

AQoS allows you to define IOPS per terabytes based on allocated or used space. AQoS defines the expected IOPS (floor) and peak IOPS (ceiling) based on given capacity. As the capacity size changes, the limit also varies. This is a to VDI environments where the volume size expands or shrinks based on the number of deployed desktops.

AQoS also allows you to define the static value for minimum IOPS (when the volume size is less than 1TB). The minimum limit is only available on AFF systems. AQoS policy groups are always nonshared. The defined ceiling or floor throughput applies to each member workload individually.

Table 3. AQoS workloads

Workload Support	Supported with ONTAP 9.6?
Volume	Yes
File	Yes

Workload Support	Supported with ONTAP 9.6?
LUN	Yes
SVM	Yes
FlexGroup volume	Yes
Multiple workloads per policy group	Yes
Nonshared policy groups	Yes

Table 4. Three default Adaptive QoS policy groups

Default Policy Group	Expected IOPS per TB (Floor)	Peak IOPS per TB (Ceiling)	Absolute Minimum IOPS
Extreme	6,144	12,288	1,000
Performance	2,048	4,096	500
Value	128	512	75

If a volume is assigned with performance AQoS, and if the allocated space is 10TB, it receives a minimum of 20,480 guaranteed IOPS. When the volume is 2TB, it receives a minimum of 4,096 guaranteed IOPS.

Active IQ Unified Manager

NetApp Active IQ® Unified Manager is a graphical management product that provides comprehensive monitoring and key management capabilities for ONTAP systems to help manage the availability, capacity, protection, and performance risks of your storage systems. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

Unified Manager allows you to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

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.

Unified Manager allows you to report different views of your network, providing actionable intelligence on capacity, health, performance, and protection data. You can customize your views by showing and hiding columns, rearranging columns, filtering data, sorting data, and searching the results. You can save custom views for re-use, download them as reports, and schedule them as recurring reports to distribute through email.

Unified Manager allows you to plan the storage requirements of your cluster objects by using the information provided in the capacity and health charts, for the respective cluster object.

Architecture and Design Considerations for Desktop Virtualization

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 as 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 desktop environments for each user:

- Traditional PC: A traditional PC is what typically constitutes a desktop environment: physical device with a locally installed operating system.
- Remote Desktop Hosted Session 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 2019, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remoted Desktop Server Hosted Server sessions: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- Published Applications: Published applications run entirely on the VMware Horizon RDS hosted 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 type 1 hypervisor and is synced 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 hosted virtual desktops and Remote Desktop Server Hosted sessions were validated. Each of the sections provides some fundamental de-sign 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 subgroup:

- What is the desktop OS planned? Windows 8 or Windows 10?
- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 8/10?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?

- Will VMware Horizon RDSH be used for Hosted Shared Server applications planned? Are there any applications in-stalled?
- What is the desktop OS planned for RDS Server Roles? Windows server 2012 or Server 2016?
- Will VMware Horizon Composer or Instant Clones or another method be used for virtual desktop deployment?
- 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 database? SQL server 2012 or 2016?
- 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 subgroup specific questions?

Hypervisor selection

VMware vSphere has been identified as the hypervisor for both RDS Hosted Sessions and VDI based 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 web site.

Note
For this CVD, the hypervisor used was VMware ESXi 6.7 Update 2.
Note
Server OS and Desktop OS Machines configured in this CVD to support Remote Desktop Server Hosted (RDSH) shared sessions and Virtual Desktops (both non-persistent and persistent).

Storage considerations

The datastores used by VMware Horizon desktop pool needs to be in the same storage DRS cluster. Storage vMotion is permitted only on the datastores consumed by the desktop pool. To reduce the IOPS for boot storm, the Storage Accelerator can be used. This uses a content-based read cache feature in vSphere to reduce the read IOPS. VMware Instant Clone uses Storage Accelerator by default.

Note: While using Horizon Storage Accelerator, NFS VAAI is not supported.

Many user profile solutions require SMB File Share. With ONTAP, NetApp AFF provides many feature of SMB3 and also secure multi-tenancy. FlexGroup Volume not only allows the volume to grow petabytes in size, but also effectively utilizes the resources across the storage nodes to provide better performance. If a need arises for unified global namespace across the countries, Global File Cache can be explored.

Vmware Horizon design fundamentals

Designing a Vmware Horizon environment for a mixed workload

With VMware Horizon 7 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.

Table 5. Designing a VMware Horizon Environment

Environment	Details
Server OS machines	<p>You want: Inexpensive server-based delivery 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-based 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 off-line 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.</p> <p>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</p>

Environment	Details
	<p>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, a mix of Remote Desktop Server Hosted sessions (RDSH) using RDS based Server OS and VMware Horizon pooled Linked Clone Virtual Machine Desktops using VDI based desktop OS machines were configured and tested.

The mix consisted of a combination of both use cases. The following sections discuss design decisions relative to the VMware Horizon deployment, including the CVD test environment.

Deployment Hardware and Software

Products deployed

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).

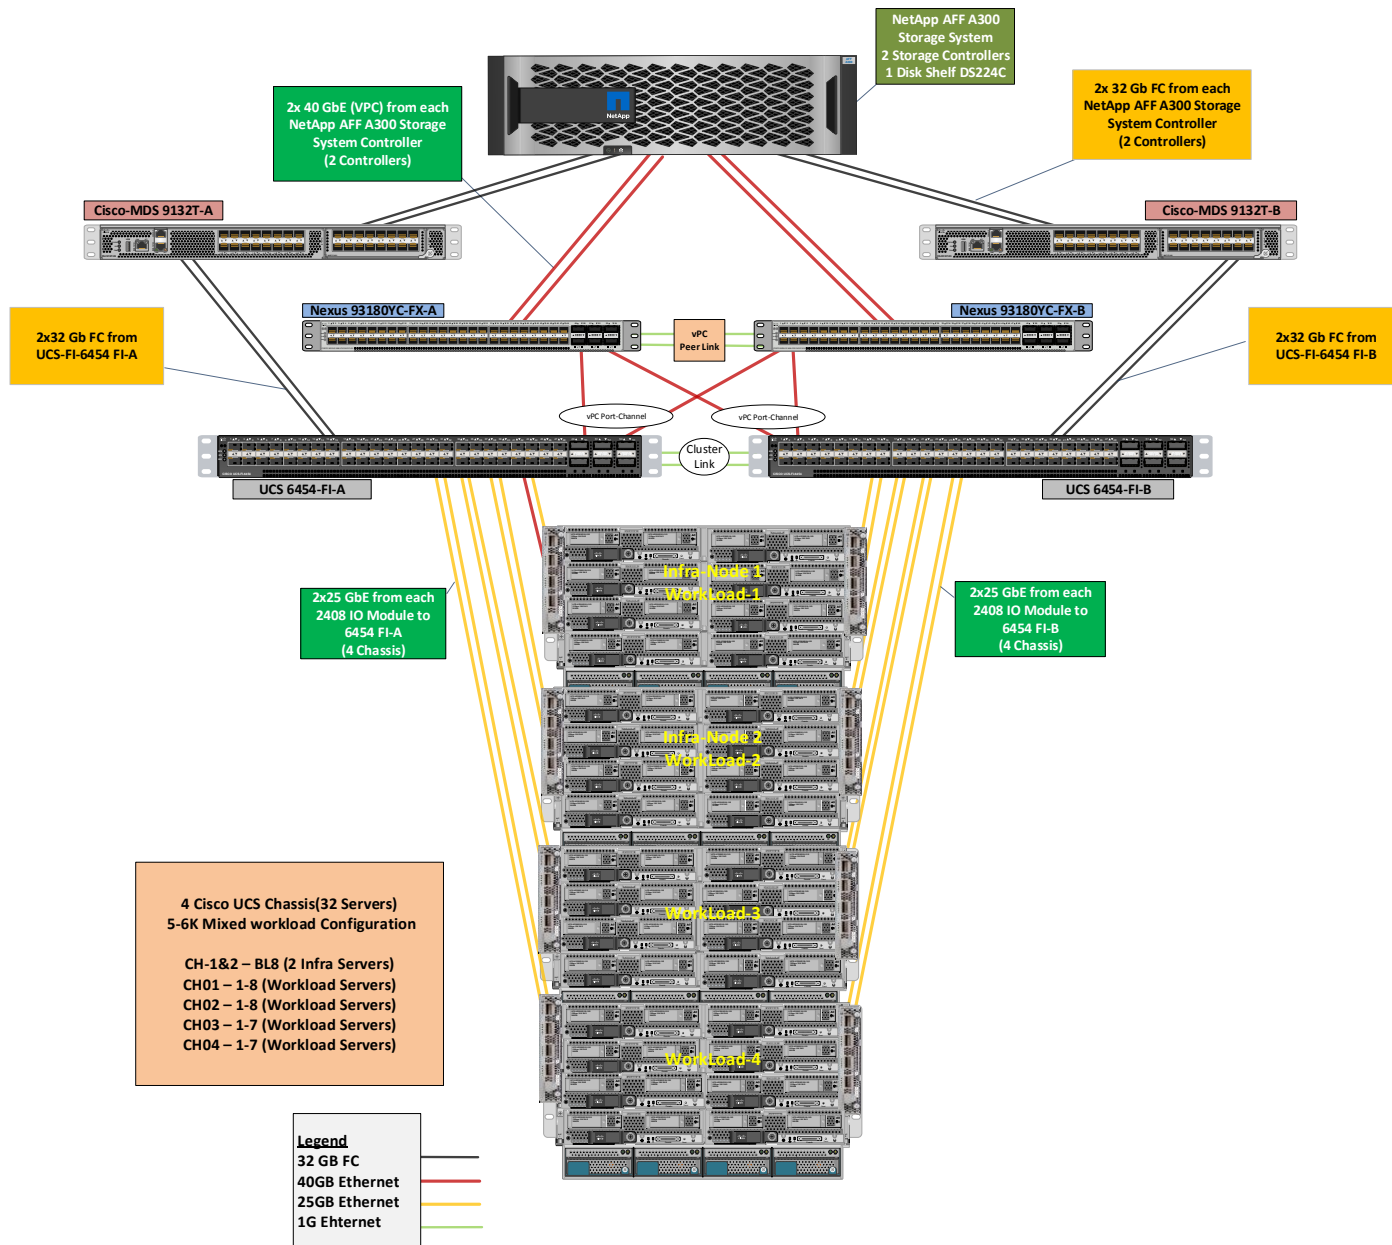
The FlexPod Datacenter solution includes Cisco networking, Cisco UCS and NetApp AFF A300, which efficiently fit into a single data center rack, including the access layer network switches.

This CVD details the deployment of the following software:

- VMware vSphere ESXi 6.7 Update 2 Hypervisor
- Microsoft SQL Server 2016 SP1
- Microsoft Windows Server 2019 and Windows 10 64-bit virtual machine Operating Systems
- VMware Horizon 7.10 Remote Desktops (RDSH) provisioned as Linked Clones and stored on the NFS storage
- VMware Horizon 7.10 Non-Persistent Virtual Desktops (VDI) provisioned as Instant Clones and stored on NFS storage
- VMware Horizon 7.10 Persistent Virtual Desktops (VDI) provisioned as Full Clones and stored on NFS storage
- NetApp Virtual Storage Console 9.6P1
- NetApp ONTAP 9.6P4

Figure 21. details the physical hardware and cabling deployed to enable this solution.

Figure 21. Virtual desktop workload reference architecture for the 5000-6000 Seat on Vmware Horizon 7.10 on FlexPod



The solution contains the following hardware:

- Two Cisco Nexus 93180YC-FX Layer 2 Access Switches.
- Two Cisco MDS 9132T 32-Gbps Fibre Channel Switches.
- Four Cisco UCS 5108 Blade Server Chassis with two Cisco UCS-IOM-2304 IO Modules

- Two Cisco UCS B200 M5 Blade servers with Intel Xeon Silver 4114 2.20-GHz 10-core processors, 192GB 2400MHz RAM, and one Cisco VIC1340 mezzanine card for the hosted infrastructure, providing N+1 server fault tolerance.
- Ten Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM, and one Cisco VIC1340 mezzanine card for the VDI/RDSH workload, providing N+1 server fault tolerance at the workload cluster level.
- Ten Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM, and one Cisco VIC1340 mezzanine card for the VDI/RDSH workload, providing N+1 server fault tolerance at the workload cluster level.
- Ten Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM, and one Cisco VIC1340 mezzanine card for the VDI/RDSH workload, providing N+1 server fault tolerance at the workload cluster level.
- NetApp AFF A300 with dual redundant controllers, with twenty-four 3.8TB SSD drives.

Note

The LoginVSI Test infrastructure is not a part of this solution. The NetApp AFF A300 configuration is detailed later in this document.

Software revisions

Table 6 lists the software versions of the primary products installed in the environment.

Table 6. Software and firmware versions

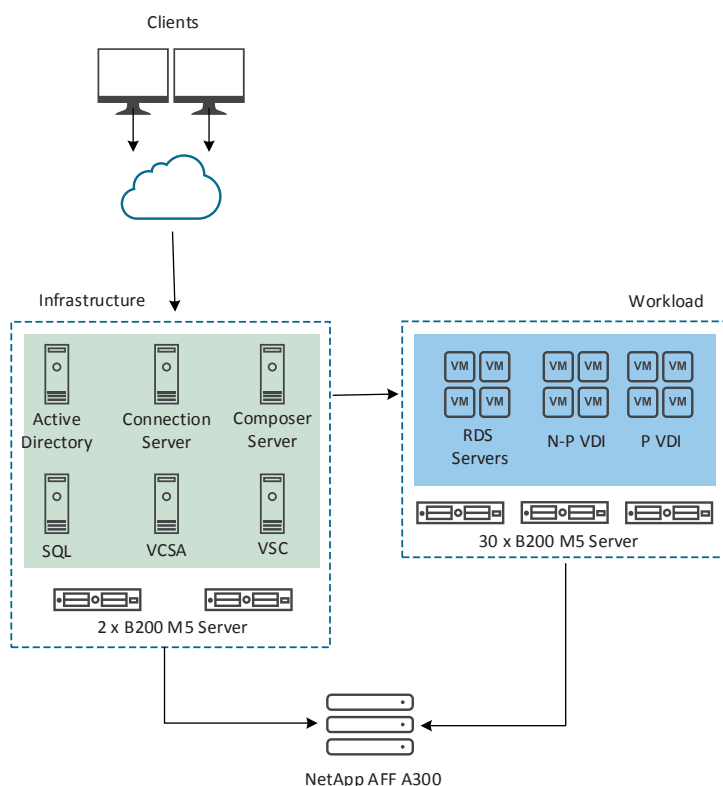
Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	4.0(4e) bundle release
Cisco	UCS Manager	4.0(4e) bundle release
Cisco	UCS B200 M5 Blades	4.0(4e) bundle release
Cisco	VIC 1340	4.3(3ba)
VMware	vCenter Server Appliance	6.7.0.32000
VMware	vSphere ESXi 6.7 Update 2	6.7.0.13006603
VMware	Horizon Connection Server	7.10.0.14584133
VMware	Horizon Composer Server	7.10.0.14535354
VMware	Horizon Agent	7.10.0

Vendor	Product / Component	Version / Build / Code
NetApp	VCS	9.6P1
NetApp	AFF A300	ONTAP 9.6P4

Logical architecture

The logical architecture of the validated solution which is designed to support 5-6000 users within a single 42u rack containing 32 blades in 4 chassis, with physical redundancy for the blade servers for each workload type is illustrated in Figure 22.

Figure 22. Logical architecture overview



Configuration guidelines

The VMware Horizon solution described in this document provides details for configuring a fully redundant, high-availability 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 7.10 customer environment as a stand-alone solution.

VLANS

The VLAN configuration recommended for the environment includes a total of eight VLANs as outlined in Table 7.

Table 7. VLANs configured in this CVD

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	60	In-Band management interfaces
Infra-Mgmt	61	Infrastructure Virtual Machines
VDI-Network	102	RDSH, Persistent and Non-Persistent
vMotion	66	VMware vMotion
NFS-Vlan	63	NFS storage access
CIFS-Vlan	62	CIFS storage access
OOB-Mgmt	164	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	Purpose
VSAN 400	400	VSAN for Primary SAN communication
VSAN 401	401	VSAN for Secondary SAN communication



Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution.

How to read deployment commands

The guide uses the following conventions for commands that you enter at the command-line interface (CLI).

Commands to enter at a CLI prompt:

```
configure terminal
```

Commands that specify a value for a variable (variable is in bold italics):

```
ntp server 10.4.0.1
```

Commands with variables that you must define (definition is bracketed in bold and italics):

```
class-map [highest class name]
```

Commands at a CLI or script prompt (entered commands are in bold):

```
Router# enable
```

Long command lines that wrap on a printed page (underlined text is entered as one command):

```
police rate 1000 pps burst 10000  
packets conform-action
```

Solution cabling

The following sections detail the physical connectivity configuration of the FlexPod 6000 seat mixed workload VMware Horizon 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 AFF A300 storage array to the Cisco 6454 Fabric Interconnects through Cisco MDS 9132T 32-Gbps 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.

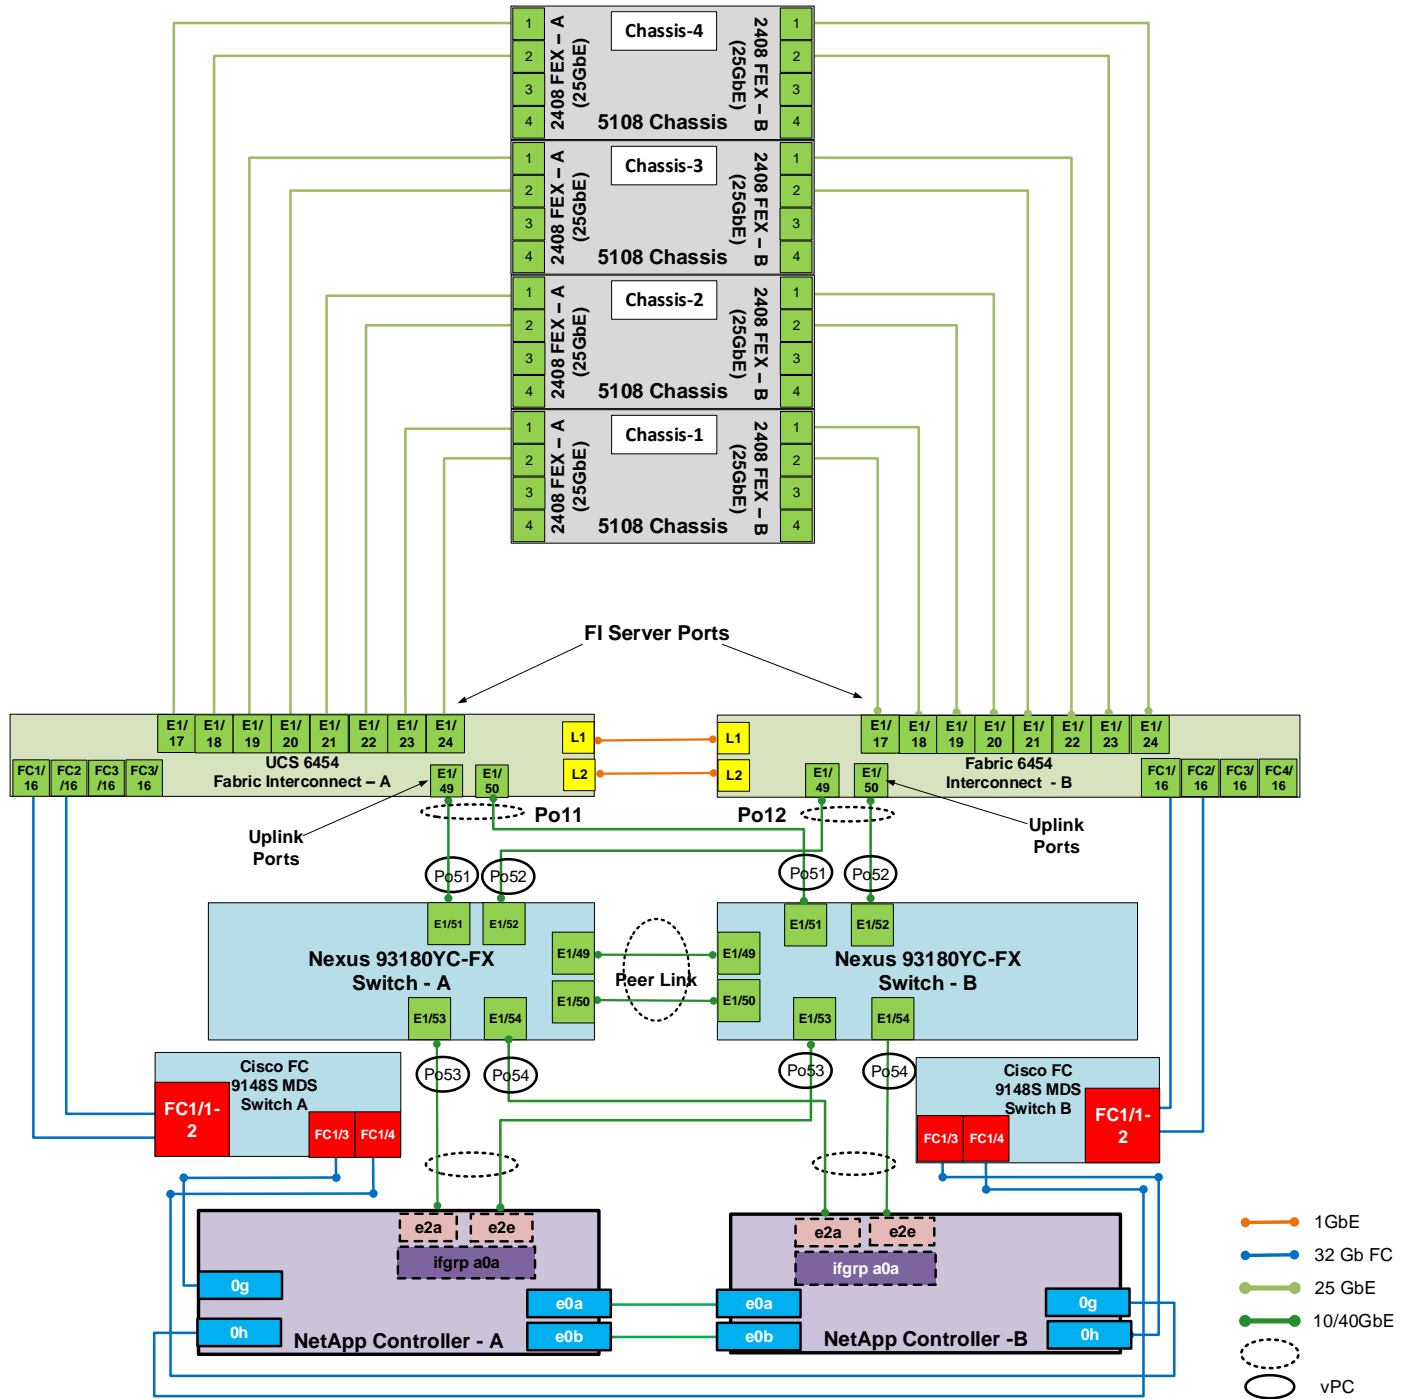
Note

Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment pro-

cedures that follow because specific port locations are mentioned.

Figure 23. shows a cabling diagram for a configuration using the Cisco Nexus 9000, Cisco MDS 9100 Series, and NetApp AFF A300 array.

Figure 23. FlexPod 5-6000 seat cabling diagram



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 4.0(4e)

This document assumes you are using Cisco UCS Manager Software version 4.0(4e). To upgrade the Cisco UCS Manager software and the Cisco UCS 6454 Fabric Interconnect software to a higher version of the firmware,) refer to the [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-box first time installation prompt appears, type **console** and press **Enter**.

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

Procedure 2. Configure fabric interconnects for a cluster setup

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

- The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
- The L1 ports on both fabric interconnects are directly connected to each other
- The L2 ports on both fabric interconnects are directly connected to each other

Step 2. Connect to the console port on the first Fabric Interconnect.

Step 3. Review the settings on the console. Answer **yes** to Apply and Save the configuration.

Step 4. Wait for the login prompt to make sure the configuration has been saved to Fabric Interconnect A.

Step 5. Connect the console port on the second Fabric Interconnect, configure secondary FI.

Figure 24. 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 25. 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.

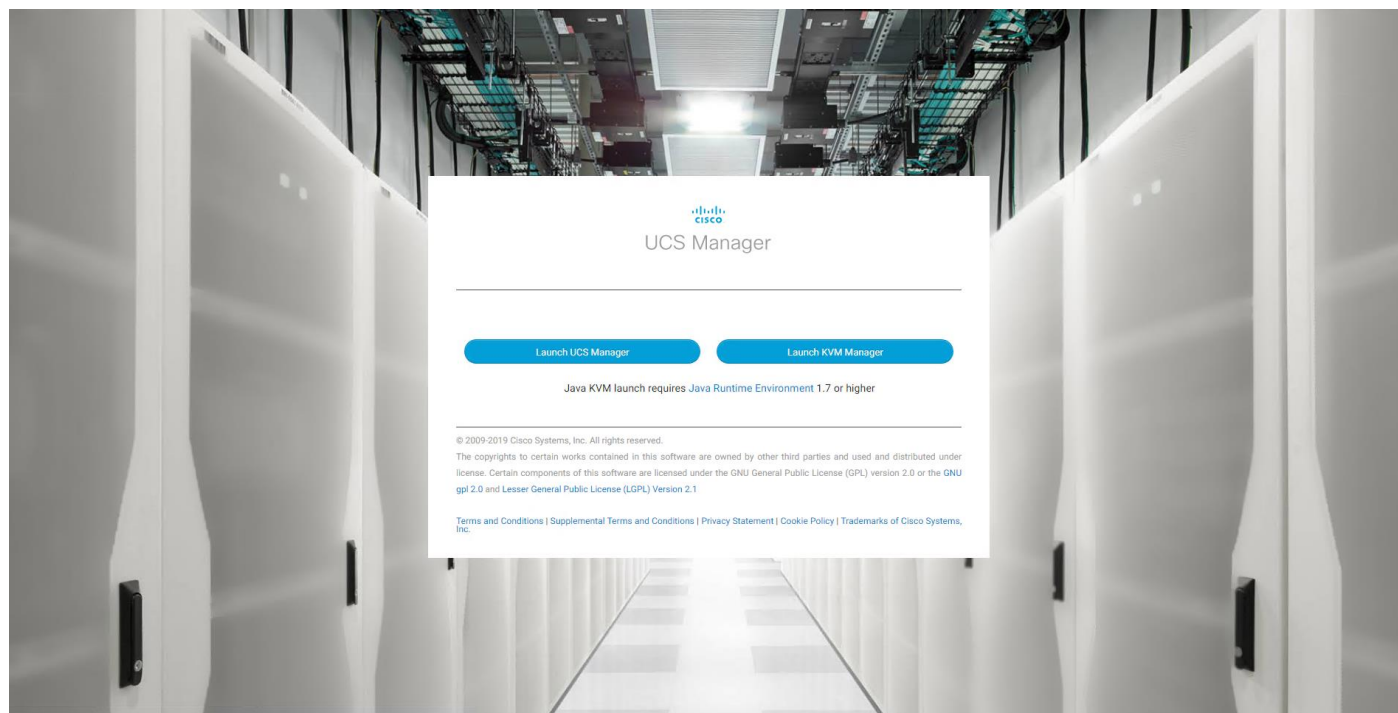
Fri Feb 16 18:53:15 UTC 2018
Configuration file - Ok

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

Step 6. Log into the Cisco Unified Computing System (Cisco UCS) environment by opening a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address configured above.

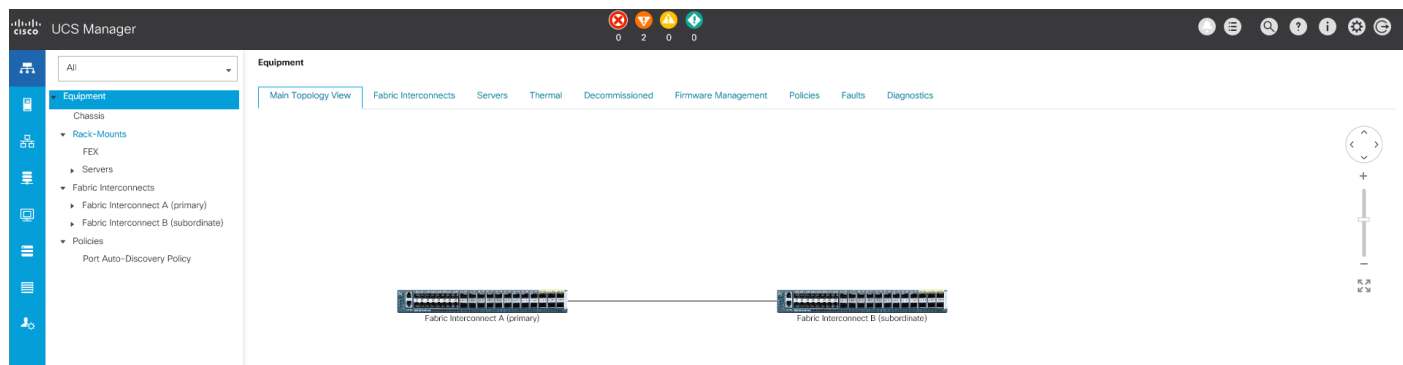
Step 7. Click the **Launch UCS Manager** link to download the Cisco UCS Manager software. If prompted, accept the security certificates.

Figure 26. Cisco UCS Manager web interface



Step 8. When prompted, enter the user name and password. Click **Log In** to log into Cisco UCS Manager.

Figure 27. Cisco UCS Manager web interface after login



Configure base Cisco Unified Computing System

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

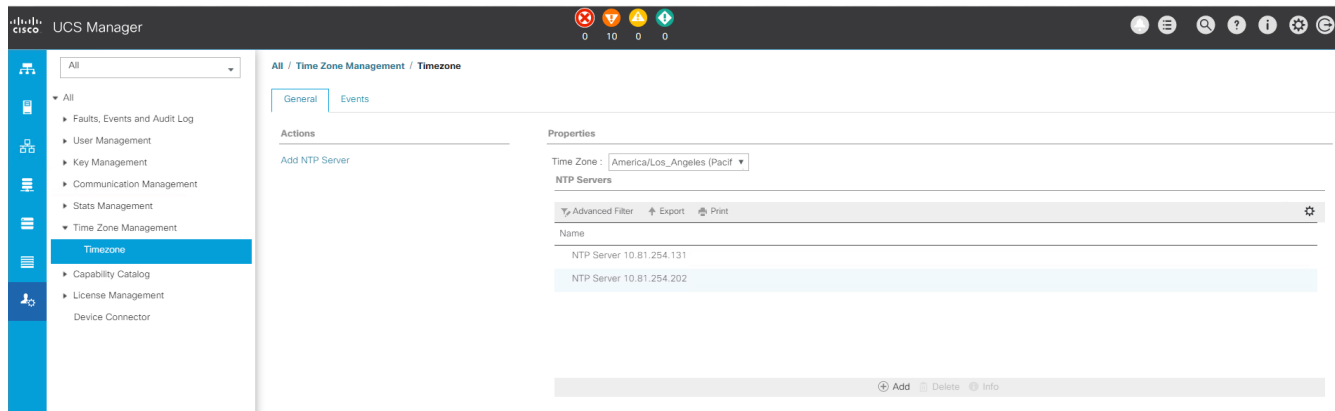
- Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

Procedure 1. Synchronize Cisco UCSM 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.** Click **Save Changes**.

Figure 28. Synchronize Cisco UCS Manager to NTP



- Step 9.** 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 2. Configure global policies

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 (right pane) > Global Policies > Chassis/FEX Discovery Policies**. As shown below, from the drop-down list select the **Action Platform Max** and select **Port Channel** for the Link Grouping Preference.
- Step 2.** Click **Save Changes**.
- Step 3.** Click **OK**.

Equipment

Main Topology View Fabric Interconnects Servers Thermal Decommissioned Firmware Management **Policies** Faults Diagnostics

Global Policies Autoconfig Policies Server Inheritance Policies Server Discovery Policies SEL Policy Power Groups Port Auto-Discovery Policy Security

Chassis/FEX Discovery Policy

Action : Platform Max

Link Grouping Preference : None Port Channel

Rack Server Discovery Policy

Action : Immediate User Acknowledged

Scrub Policy : <not set>

Rack Management Connection Policy

Action : Auto Acknowledged User Acknowledged

Power Policy

Redundancy : Non Redundant N+1 Grid

MAC Address Table Aging

Aging Time : Never Mode Default other

Global Power Allocation Policy

Allocation Method : Manual Blade Level Cap Policy Driven Chassis Group Cap

Firmware Auto Sync Server Policy

Sync State : No Actions User Acknowledge

[Save Changes](#) [Reset Values](#)

Procedure 3. Set fabric interconnects to fibre channel end host mode

Step 1. Configure the FC Uplink ports connected to Cisco UCS MDS 9132T 32-Gbps FC switch, set the Fabric Interconnects to the Fibre Channel End Host Mode.

Step 2. Verify that Fabric Interconnects are operating in **FC End-Host Mode**.

Fabric Interconnects

- Fabric Interconnects
 - Fabric Interconnect A (primary)
 - Fabric Interconnect B (subordinate)

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 on Locator LED

Fabric Interconnects / Fabric Interconnect A (primary)

General Physical Ports Fans PSUs Physical Di

Fault Summary

0 0 0

Operable

OK

End Host

End Host

ode : Off

de : Off

Actions

Note

Fabric Interconnect automatically reboot if switched operational mode; perform this task on one FI first, wait for FI to come up and follow the same on second FI.

Procedure 4. 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

Fault Summary

0 0 0 0

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

Physical Display

Legend: Up (Green), Admin Down (Yellow), Fail (Red), Link Down (Orange)

Properties

Name : **A**
Product Name : **Cisco UCS 6454**
Vendor : **Cisco Systems, Inc.** PID : **UCS-FI-6454**
Revision : **0** Serial : **FDO23320Q11**
Available Memory : **52.515 (GB)** Total Memory : **62.761 (GB)**
Locator LED :

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

Firmware

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

Configure Unified Ports

The Configure Unified Ports wizard allows you to change the port mode from Ethernet to Fibre Channel or FC to Ethernet. Changing the port mode on either module causes an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion module require a reboot of that module. Are you sure you want to launch this wizard and reboot the modules associated with any reconfigured ports?

Yes No

Step 3. Move the slider to the right.

Step 4. Click **OK**.


Note

Ports to the left of the slider will become FC ports. The unified ports of the Cisco UCS 6454 are configured in sets of four. For our CVD, we configured the first four ports on the FI as FC Uplink ports.

Note

Applying this configuration will cause the immediate reboot of Fabric Interconnect and/or Expansion Module(s).

Configure Unified Ports ? X



Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	

OK
Cancel

Step 5. Click **Yes** to apply the changes.

Step 6. After the FI reboot, your FC Ports configuration will look like Figure 29.

Step 7. Follow steps 1–6 on Fabric Interconnect B.

Figure 29. FC Uplink Ports on Fabric Interconnect A

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20:01:00:3A:9C:A4:FD:80	Network	Physical	↑ Up	↑ Enabled
1	2	20:02:00:3A:9C:A4:FD:80	Network	Physical	↑ Up	↑ Enabled
1	3	20:03:00:3A:9C:A4:FD:80	Network	Physical	↑ Up	↑ Enabled
1	4	20:04:00:3A:9C:A4:FD:80	Network	Physical	↑ Up	↑ Enabled

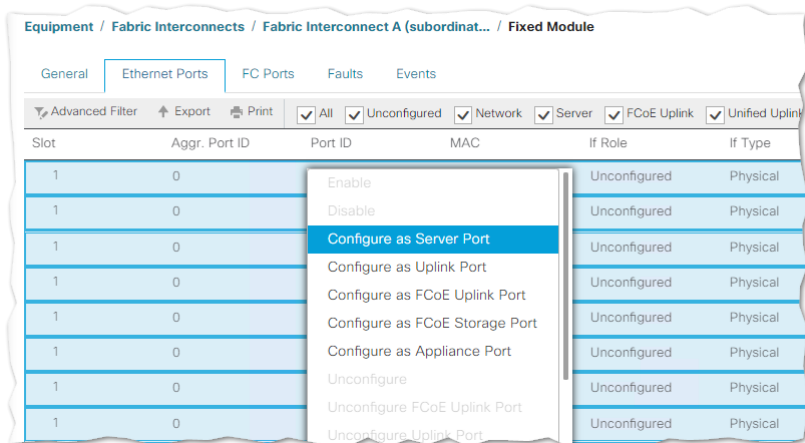
Procedure 5. Configure the server ports to initiate chassis and blade discovery

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 Cisco UCS B-Series 5108 Chassis.

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

Figure 30. Configure server port on Cisco UCS Manager Fabric Interconnect for chassis/server discovery



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 31. on both Fabric Interconnects.

Figure 31. Server ports on Fabric Interconnect A

The screenshot shows the Cisco UCS Manager interface displaying a list of configured server ports. The breadcrumb navigation is "Equipment / Fabric Interconnects / Fabric Interconnect A (subordinat... / Fixed Module / Ethernet Ports". The "Ethernet Ports" tab is selected. A table lists ports with columns for Slot, Aggr. Port ID, Port ID, MAC, If Role, If Type, Overall Status, Admin State, and Peer. All ports are in the "Up" state and "Enabled".

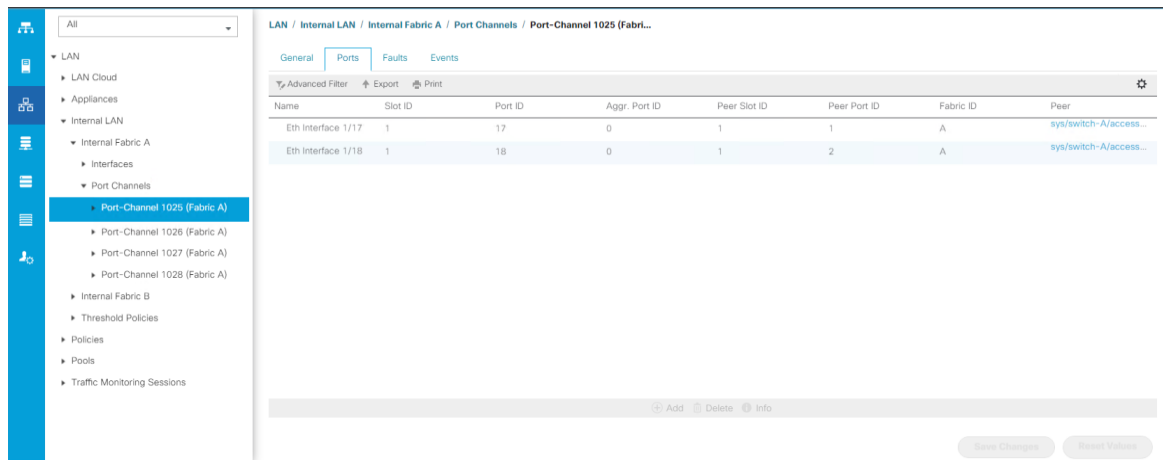
Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:A4:FD:98	Server	Physical	Up	Enabled	sys/chassis-1/slot-1/f...
1	0	18	00:3A:9C:A4:FD:99	Server	Physical	Up	Enabled	sys/chassis-1/slot-1/f...
1	0	19	00:3A:9C:A4:FD:9A	Server	Physical	Up	Enabled	sys/chassis-2/slot-1/f...
1	0	20	00:3A:9C:A4:FD:9B	Server	Physical	Up	Enabled	sys/chassis-2/slot-1/f...
1	0	21	00:3A:9C:A4:FD:9C	Server	Physical	Up	Enabled	sys/chassis-3/slot-1/f...
1	0	22	00:3A:9C:A4:FD:9D	Server	Physical	Up	Enabled	sys/chassis-3/slot-1/f...
1	0	23	00:3A:9C:A4:FD:9E	Server	Physical	Up	Enabled	sys/chassis-4/slot-1/f...
1	0	24	00:3A:9C:A4:FD:9F	Server	Physical	Up	Enabled	sys/chassis-4/slot-1/f...

Step 6. After configuring Server Ports, acknowledge the Chassis. Go to **Equipment > Chassis > Chassis 1 > General > Actions > Acknowledge Chassis**. Repeat this step to acknowledge chassis 2-4.

Step 7. After acknowledging the chassis, acknowledge all the servers placed in the chassis. Go to **Equipment > Chassis 1 > Servers > Server 1 > General > Actions > Server Maintenance > Re-acknowledge** and click **OK**. Repeat this step to acknowledge all eight servers.

Step 8. When the acknowledgement of the servers is completed, verify the port-channel of the Internal LAN. Go to the **LAN tab > Internal LAN > Internal Fabric A > Port Channels** as shown in Figure 32.

Figure 32. Internal LAN Port Channels

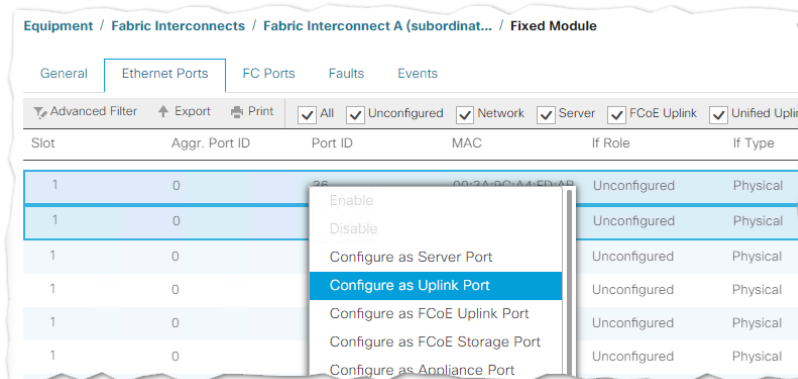


Procedure 6. Configure ethernet LAN uplink ports

Configure network ports that are used to uplink the fabric interconnects to the Cisco Nexus switches.

- 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 33. Network uplink port configuration on fabric interconnect configuration



- Step 5.** Click **Yes** to confirm ports and click **OK**.
- Step 6.** Verify the ports connected to the 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 34. Network uplink port on fabric interconnect

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	43	00:3A:9C:A4:FD:B2	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	44	00:3A:9C:A4:FD:B3	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	45	00:3A:9C:A4:FD:B4	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	46	00:3A:9C:A4:FD:B5	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	47	00:3A:9C:A4:FD:B6	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	48	00:3A:9C:A4:FD:B7	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	49	00:3A:9C:A4:FD:B8	Network	Physical	Up	Enabled	
1	0	50	00:3A:9C:A4:FD:BC	Network	Physical	Up	Enabled	
1	0	51	00:3A:9C:A4:FD:C0	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	52	00:3A:9C:A4:FD:C4	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	53	00:3A:9C:A4:FD:C8	Unconfigured	Physical	Sfp Not Present	Disabled	
1	0	54	00:3A:9C:A4:FD:CC	Unconfigured	Physical	Sfp Not Present	Disabled	

Now you’ve created two uplink ports on each Fabric Interconnect as shown above. These ports will be used to create a Virtual Port Channel.

Procedure 7. Create uplink port channels to Cisco Nexus switches

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. Under **LAN > LAN Cloud**, expand node Fabric A tree:

- Right-click **Port Channels**.
- Select **Create Port Channel**.
- Enter **11** as the unique ID of the port channel.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

ID : 11

Name :

Step 3. Enter the name of the port channel.

Create Port Channel

1 Set Port Channel Name

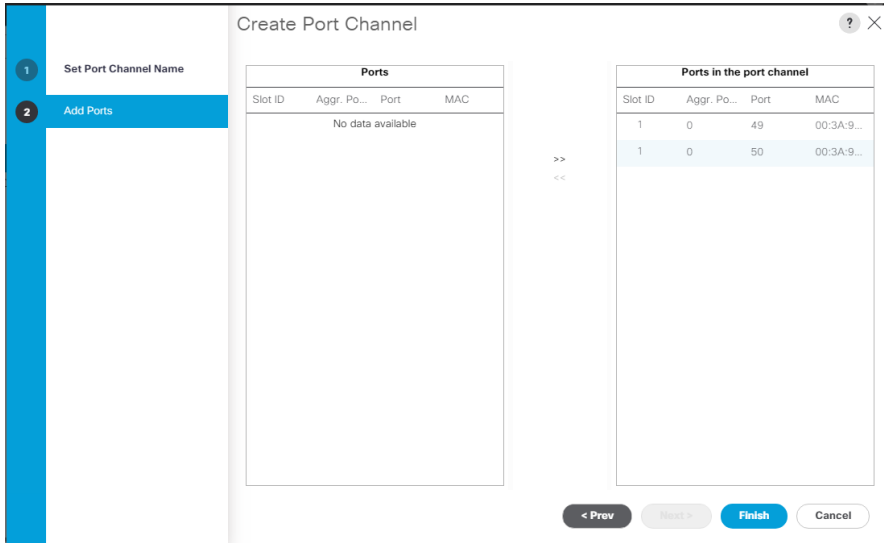
2 Add Ports

ID : 11

Name : PORTCH-UP-NPK-A

Step 4. Click **Next**.

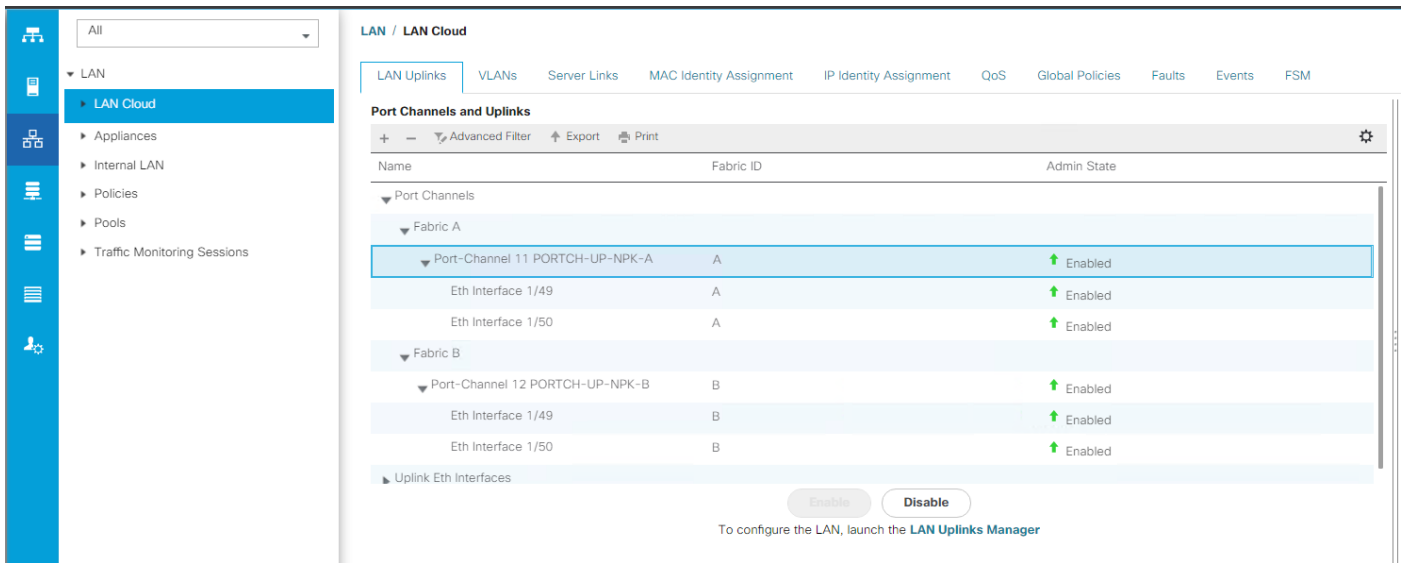
Step 5. Select Ethernet ports **49-50** for the port channel.



Step 6. Click **Finish**.

Step 7. Repeat steps 1-6 for the Port Channel configuration on FI-B making sure the port channel id is unique.

Figure 35. Port Channels configured for this solution



Procedure 8. Configure VLAN

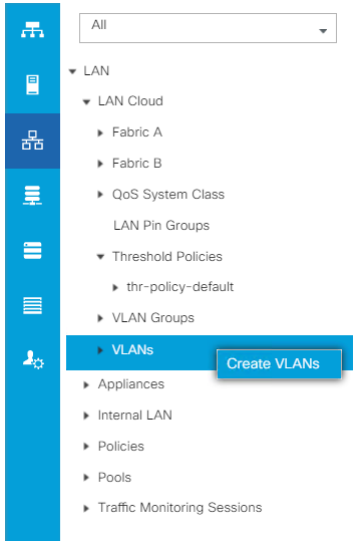
Configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment.

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

Step 2. Select **LAN > LAN Cloud**.

Step 3. Right-click **VLANs**.

Step 4. Select **Create VLANs**.



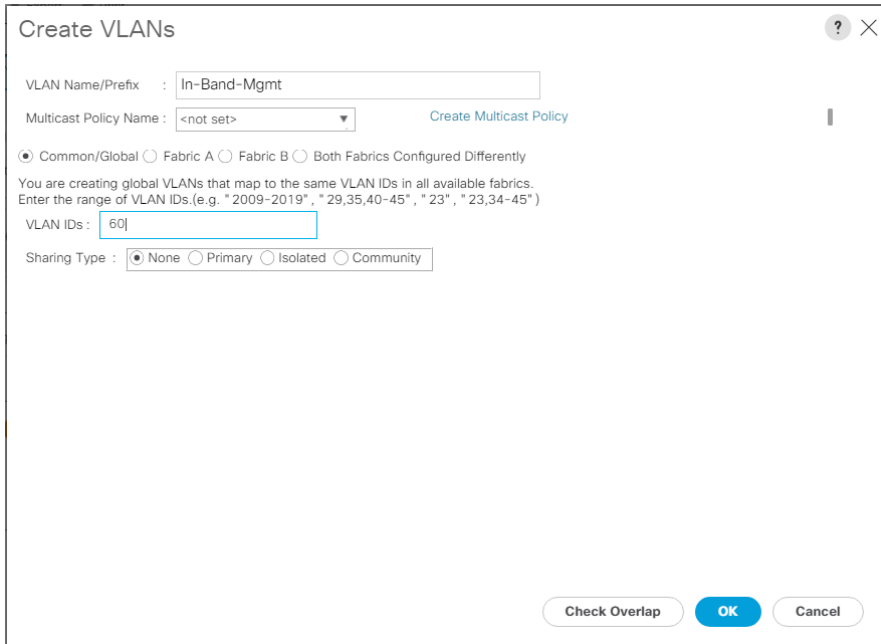
Step 5. Enter **In-Band-Mgmt** for the name of the VLAN to be used for Public Network Traffic.

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**.



Step 10. Repeat steps 1-9 to create the required VLANs. Figure 36. shows the VLANs configured for this solution.

Figure 36. VLANs configured for this solution

LAN / LAN Cloud / VLANs

VLANs

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
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 vMotion (66)	66	Lan	Ether	No	None
VLAN PVS-PXE (68)	68	Lan	Ether	No	None
VLAN VDI (102)	102	Lan	Ether	No	None

Buttons: + Add, - Delete, i Info

Tech tip

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

Procedure 9. Configure VSAN

Configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment.

- 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**.

SAN

- SAN Cloud
 - Fabric A
 - FC Port Channels
 - FCoE Port Channels
 - Uplink FC Interfaces
 - Uplink FCoE Interfaces
 - VSANs**
 - Create VSAN**
 - Fabric B
 - FC Port Channels

- Step 5.** Enter the name of the VSAN.

Note

In this solution, two VSANs are created; VSAN-A 400 and VSAN-B 401 for SAN Boot.

Step 6. Select **Fabric A** for the scope of the VSAN:

- For the ID of the VSAN and FCoE VLAN enter **400**.
- Click **OK** and then click **OK** again.

Create VSAN

Name :

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

Enter the VSAN ID that maps to this VSAN. Enter the VLAN ID that maps to this VSAN.

VSAN ID : FCoE VLAN :

Step 7. Repeat steps 1-6 to create the VSANs necessary for this solution.

VSAN 400 and 401 are configured as shown in Figure 37.

Figure 37. VSANs configured for this solution

The screenshot shows the SAN configuration interface with a left-hand navigation menu and a main content area. The main content area displays a table of VSANs configured for Fabric A and Fabric B.

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

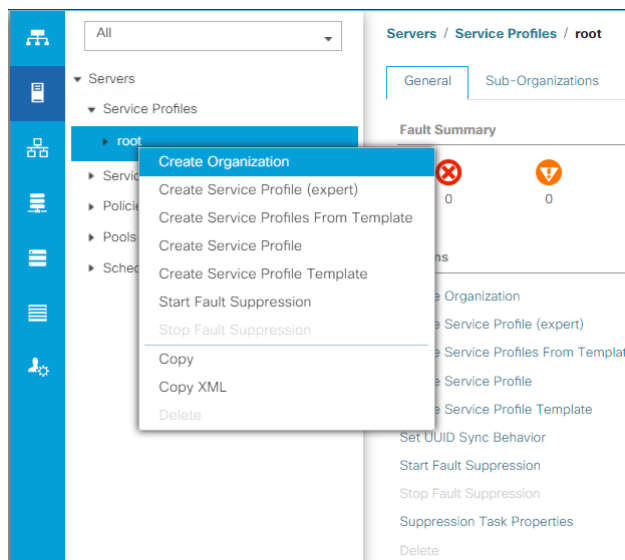
Procedure 10. Create new sub-organization

Configure the necessary Sub-Organization for the Cisco UCS environment.

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

Step 2. Select **Service Profiles > root**.

Step 3. Right-click **root > Create Organization**.



Step 4. Enter the name of the Sub-Organization.

Step 5. Click **OK**.

Create Organization

Name :

Description :

Note

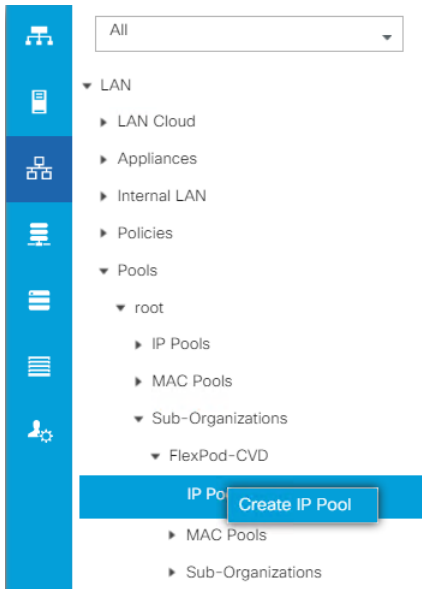
You will create pools and policies required for this solution under the newly created "FlexPod-CVD" sub-organization

Configure IP, UUID, Server, MAC, WWNN, and WWPN Pools

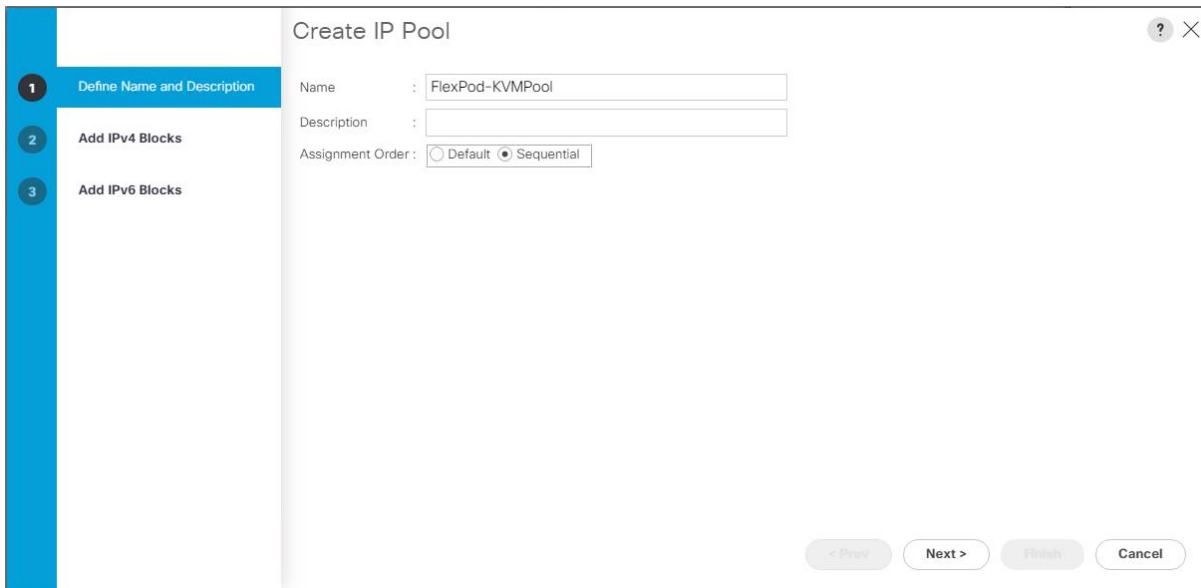
Procedure 1. IP pool creation

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.** Select **Pools > root > Sub-Organizations > FlexPod-CVD > IP Pools**.
- Step 3.** Right-click **IP Pools** and select **Create IP Pool**.



- Step 4.** For Assignment Order select **Sequential** to assign the IP in sequential order then click **Next**.

A screenshot of the 'Create IP Pool' dialog box. The dialog has a left sidebar with three steps: 1. Define Name and Description (selected), 2. Add IPv4 Blocks, and 3. Add IPv6 Blocks. The main area contains the following fields: Name: FlexPod-KVMPool, Description: (empty), and Assignment Order: Default Sequential. At the bottom right, there are four buttons: < Prev, Next >, Finish, and Cancel.

- Step 5.** In the navigation pane select **IPv4 Block** and click **Add**.

Step 6. Enter the starting IP address of the block and the number of IP addresses required and the subnet and gateway information as shown below then click **OK**.

Create IP Pool

1 Define Name and Description

2 **Add IPv4 Blocks**

3 Add IPv6 Blocks

Create Block of IPv4 Addresses

From : Size :

Subnet Mask : Default Gateway :

Primary DNS : Secondary DNS :

OK Cancel

Step 7. Click **Next**.

Create IP Pool

1 Define Name and Description

2 **Add IPv4 Blocks**

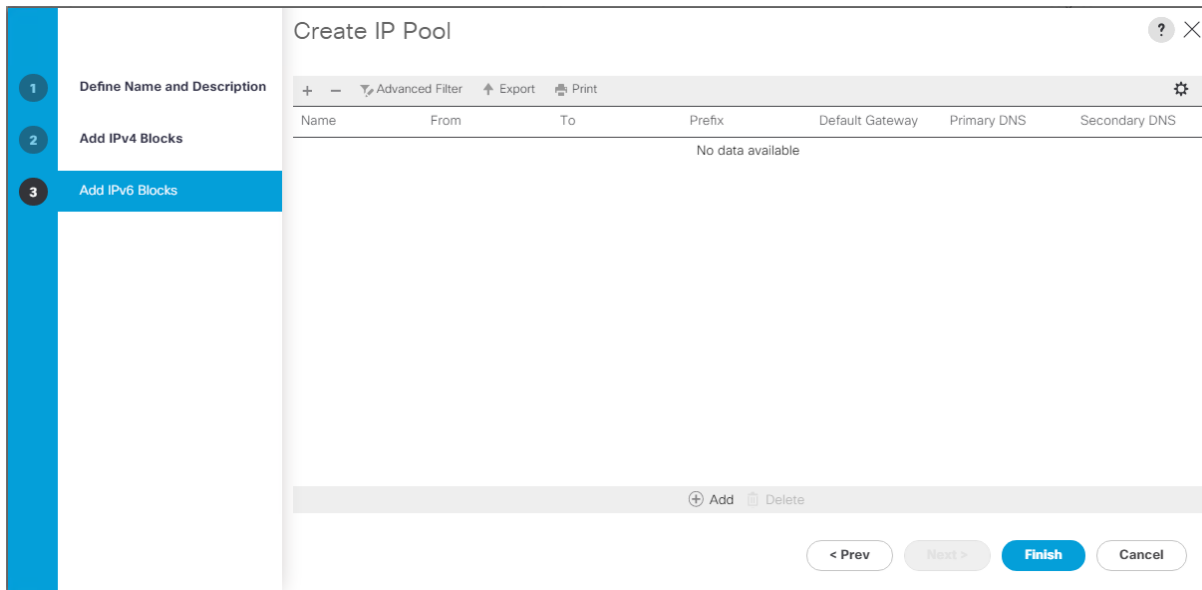
3 Add IPv6 Blocks

Name	From	To	Subnet	Default Gateway	Primary DNS	Secondary DNS
[10.29.164.8...	10.29.164.86	10.29.164.117	255.255.255.0	10.29.164.1	0.0.0.0	0.0.0.0

+ Add - Delete

< Prev **Next >** Finish Cancel

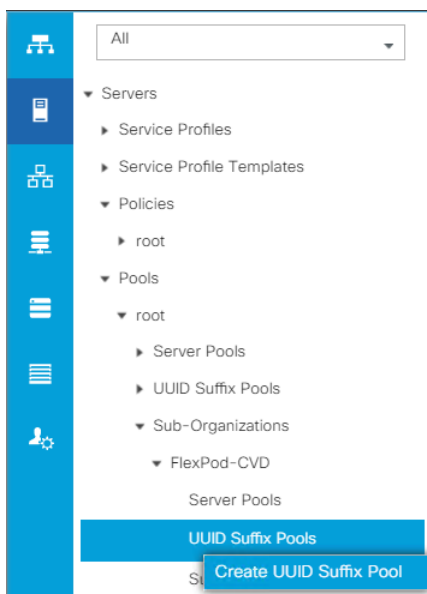
Step 8. Click **Finish** and then click **OK**.



Procedure 2. UUID Suffix Pool creation

Configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment.

- Step 1.** In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Step 2.** Select **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 for the Assignment Order select **Sequential** click **Next**.

1 Define Name and Description

2 Add UUID Blocks

Create UUID Suffix Pool

Name : FlexPod-UUIDPool

Description :

Prefix : Derived other

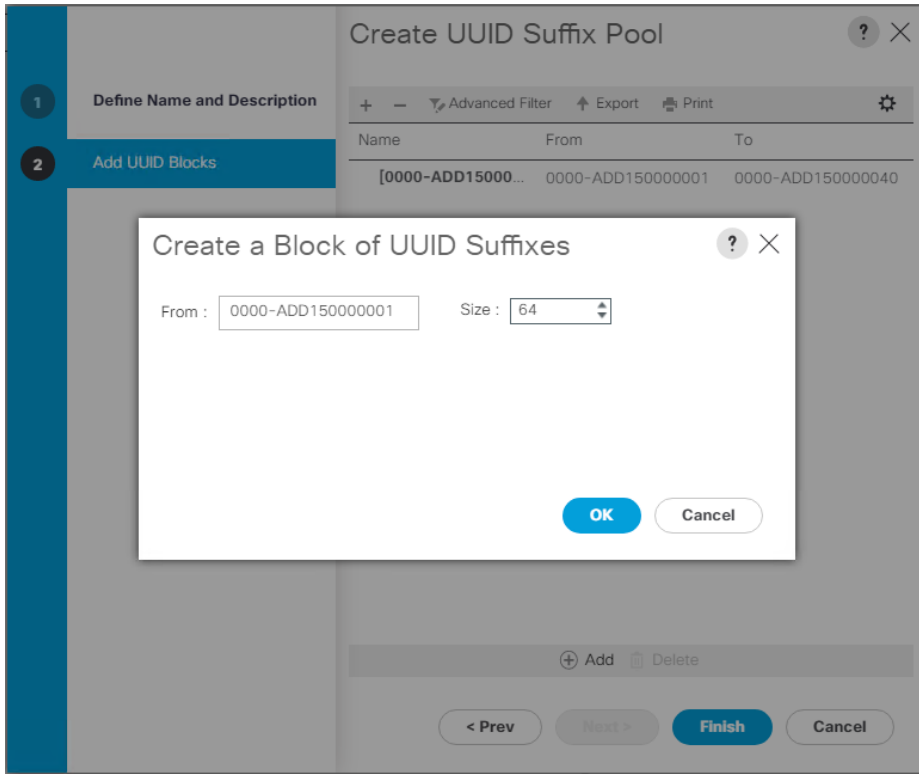
Assignment Order : Default Sequential

< Prev Next > Finish Cancel

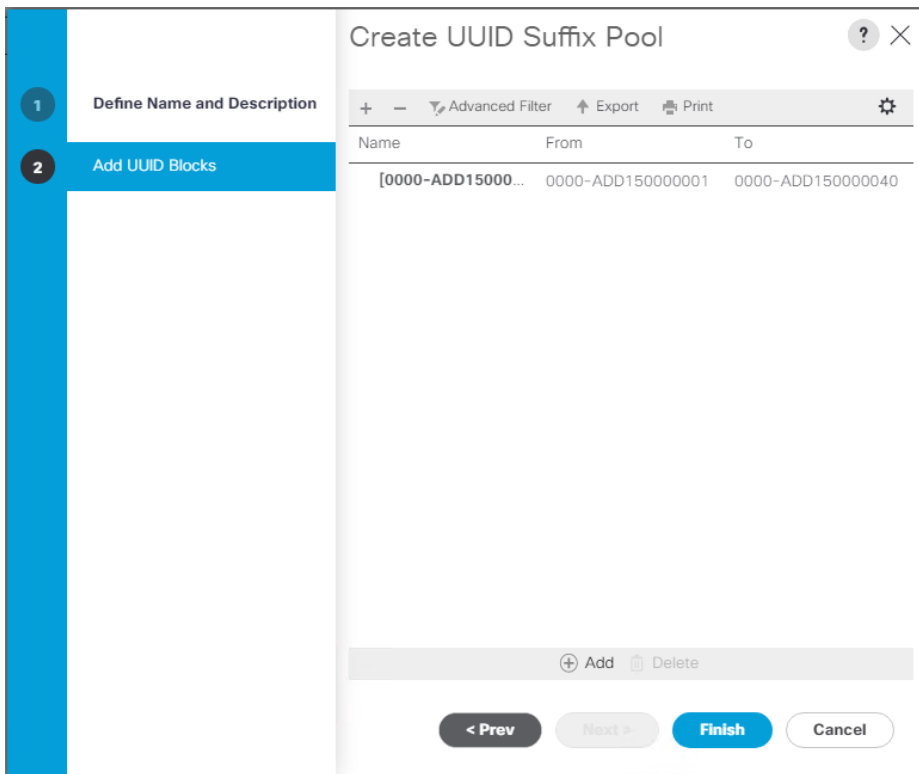
Step 7. Click **Add** to add a block of UUIDs.

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 then click **OK**.



Step 10. Click **Finish** and then click **OK**.



Procedure 3. Server Pool creation

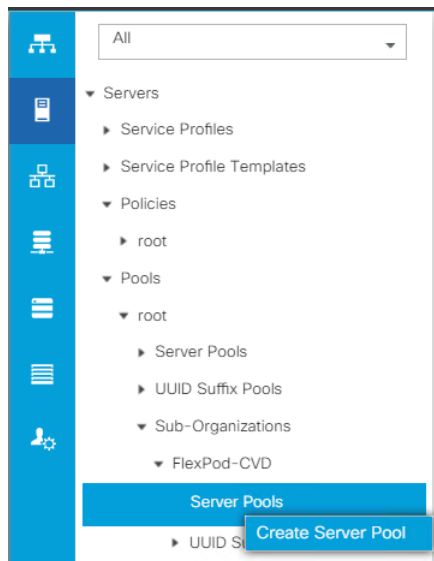
Configure the necessary server pools for the Cisco UCS environment.

Note

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. Select **Pools > root > Sub-Organization > FlexPod-CVD > right-click Server Pools > Select Create Server Pool.**



Step 3. Enter the name of the server pool.

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

1 Set Name and Description

2 Add Servers

Create Server Pool

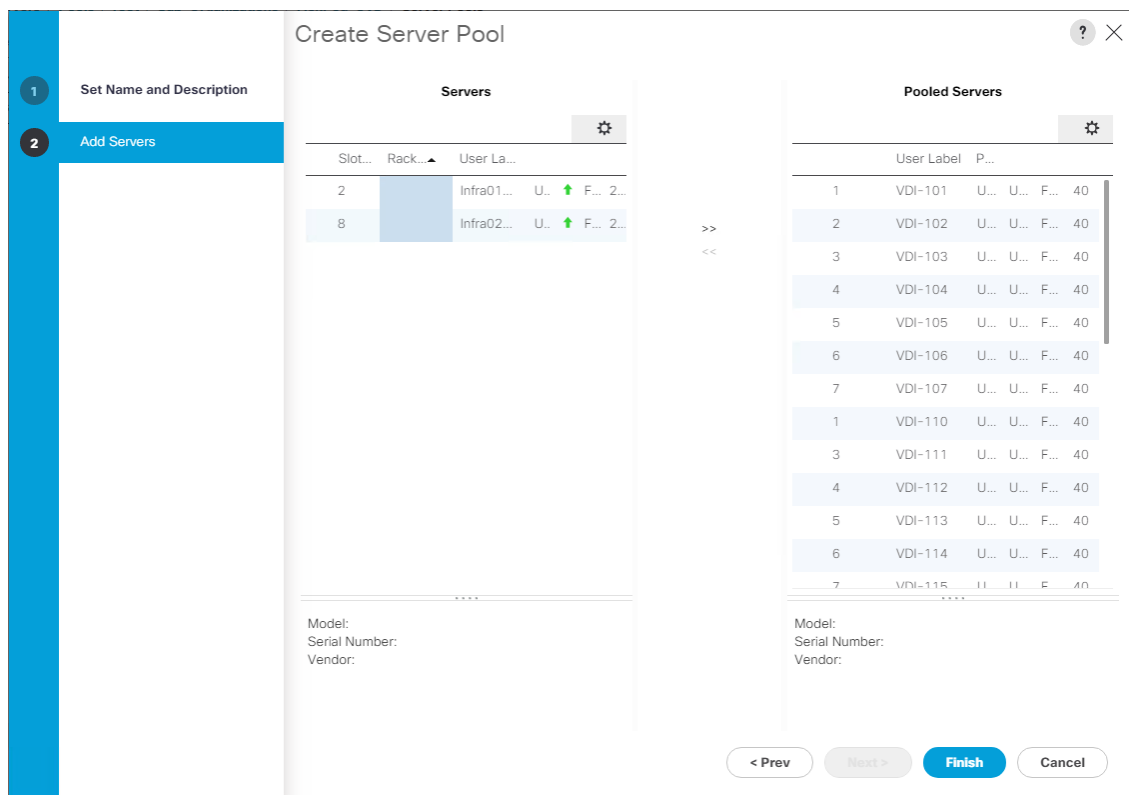
Name : FlexPod-ServerPool

Description :

< Prev Next > Finish Cancel

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

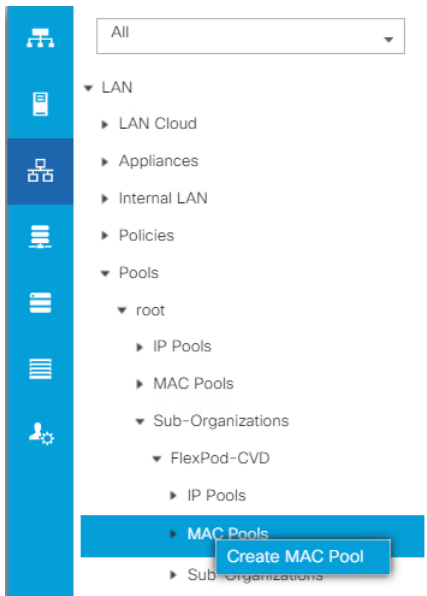
Step 6. Click **Finish** and then click **OK**.



Procedure 4. MAC Pool creation

Configure the necessary MAC address pools for the Cisco UCS environment.

- Step 1.** In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
- Step 2.** Select **Pools > root > Sub-Organization > FlexPod > right-click MAC Pools** under the root organization.
- Step 3.** Select **Create MAC Pool** to create the MAC address pool.



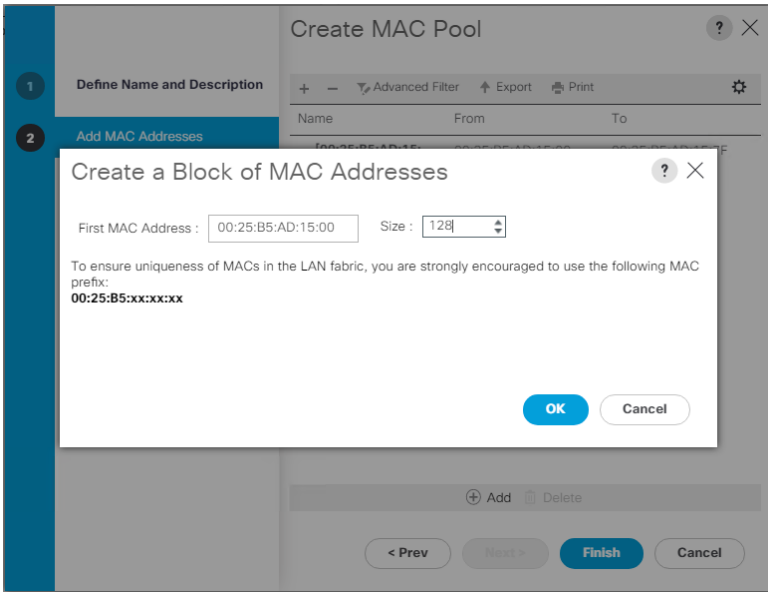
Step 4. Enter a name for MAC pool. For the Assignment Order select **Sequential**. Click **Next**.

A screenshot of the 'Create MAC Pool' dialog box. The dialog has a title bar with a question mark and a close button. On the left, there is a vertical navigation bar with two steps: '1 Define Name and Description' (highlighted) and '2 Add MAC Addresses'. The main area contains the following fields: 'Name : FlexPod-MACPool', 'Description :', and 'Assignment Order :'. The 'Assignment Order' has two radio buttons: 'Default' (unselected) and 'Sequential' (selected). At the bottom, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

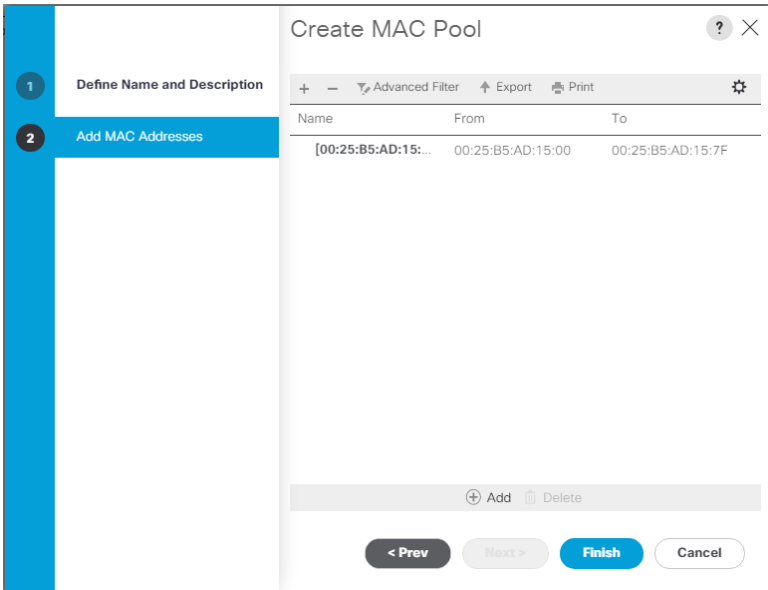
Step 5. Click **Add** to add a block of MAC addresses.

Step 6. Enter the seed MAC address and provide the appropriate number of MAC addresses to be provisioned.

Step 7. Click **OK**.



Step 8. Click **Finish** then click **OK**.



Step 9. Repeat steps 1-8 to create MAC Pool B and assign a unique MAC Addresses.

Figure 38. Configured MAC Pools

LAN / Pools / root / Sub-Organizations / FlexPod-CVD / MAC Pools

MAC Pools

+ - Advanced Filter Export Print

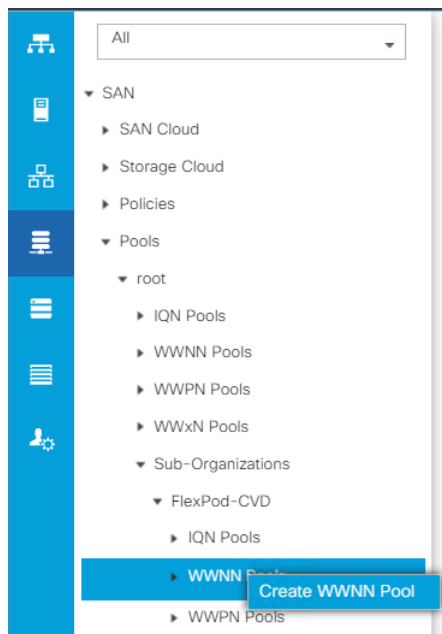
Name	Size	Assigned
MAC Pool FlexPod-MACPoolB	128	0
[00:25:B5:BD:15:00 - 00:25:B5:BD:15:7F]		
MAC Pool FlexPod-MACPool	128	0
[00:25:B5:AD:15:00 - 00:25:B5:AD:15:7F]		

Procedure 5. WWNN Pool creation

Configure the necessary WWNN (World Wide Node Name) pools for the Cisco UCS environment.

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

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



Step 3. Assign a name and for the Assignment Order select **Sequential** and click **Next**.

Create WWNN Pool

Name : FlexPod-WWNNPool

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

Step 4. Click **Add WWN Blocks** to add a block of Ports.

Step 5. Enter the Block for WWN and size of the WWNN Pool.

Create WWNN Pool

Define Name and Description

Add WWN Blocks

Create WWN Block

From : 20:00:00:25:B5:09:00:00 Size : 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

OK Cancel

Advanced Filter Export Print

From To

00:25:B5:09:00:1F

+ Add - Delete

< Prev Next > Finish Cancel

Step 6. Click **OK**.

Step 7. Click **Finish** then click **OK**.

Name	From	To
[20:00:00:25:B5:09:0...	20:00:00:25:B5:09:00:00	20:00:00:25:B5:09:00:1F

Procedure 6. WWPN Pool Creation

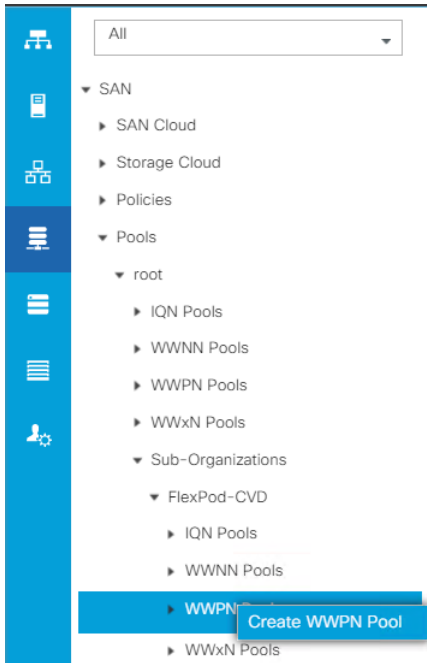
Configure the necessary WWPN pools for the Cisco UCS environment.

Note

We created two WWPN as WWPN-A Pool and WWPN-B and World Wide 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 > Create WWPN Pool**.



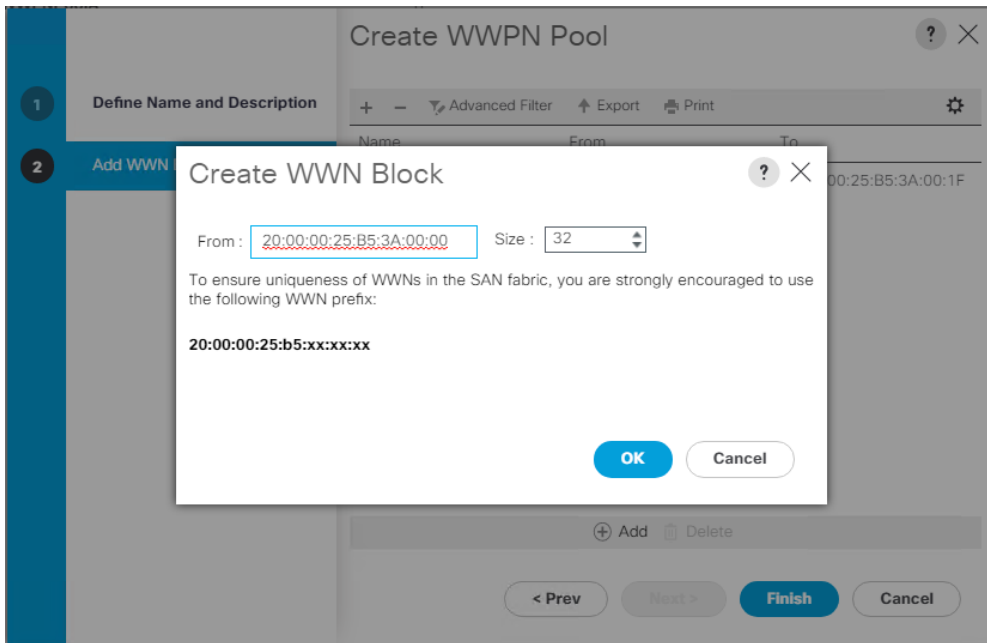
Step 3. Assign a name and for the Assignment Order select **Sequential** and click **Next**.

A screenshot of the 'Create WWPN Pool' dialog box. The dialog has a title bar with a question mark and a close button. On the left, there is a vertical sidebar with two steps: '1 Define Name and Description' (highlighted in blue) and '2 Add WWN Blocks'. The main area contains three input fields: 'Name' with the value 'FlexPod-WWPNPoolA', 'Description' (empty), and 'Assignment Order' with radio buttons for 'Default' and 'Sequential' (selected). At the bottom, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

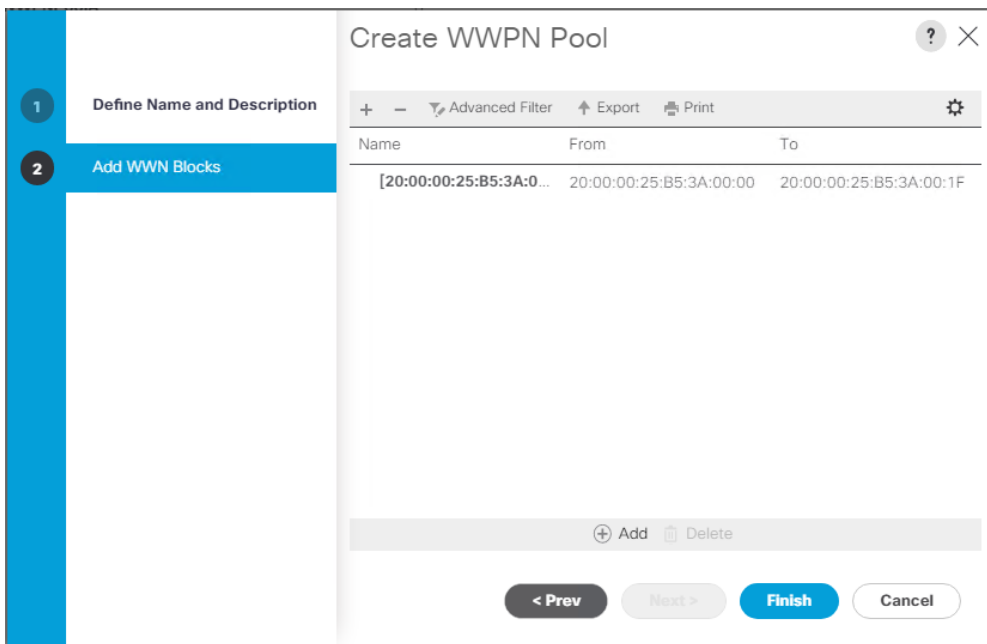
Step 4. Click **Add** to add a block of Ports.

Step 5. Enter a Block for WWN and size.

Step 6. Click **OK**.



Step 7. Click **Finish** and click **OK**.



Step 8. Repeat steps 1-7 to configure the WWPN Pool Block and assign the unique block IDs.

Figure 39. Configured WWPN Pools

SAN / Pools / root / Sub-Organizations / FlexPod-CVD / WWPN Pools

WWPN Pools

+ - Advanced Filter Export Print

Name	Size
WWPN Pool FlexPod-WWPNPoolA [20:00:00:25:B5:3A:00:00 - 20:00:00:25:B5:3A:00:1F]	32
WWPN Pool FlexPod-WWPNPoolB [20:00:00:25:D5:06:00:00 - 20:00:00:25:D5:06:00:1F]	32

Procedure 7. Set Jumbo Frames in Cisco Fabric Interconnect

Configure jumbo frames and enable quality of service in the Cisco UCS fabric.

- 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 under the MTU column, enter **9216** in the box.
- 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 8. Create host firmware package

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.

Create a firmware management policy for a given server configuration in the Cisco UCS environment,

- 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 the name of the host firmware package.
- Step 6.** Provide a Description.
- Step 7.** For “How would you like to configure the Host Firmware Package?” select **Simple**.
- Step 8.** For the Blade Package, select version **4.0(4d)**.
- Step 9.** Click **OK** to create the host firmware package.

Create Host Firmware Package
?
×

Name :

Description :

How would you like to configure the Host Firmware Package?

Simple Advanced

Blade Package :

Rack Package :

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

- Adapter
- BIOS
- Board Controller
- CIMC
- FC Adapters
- Flex Flash Controller
- GPUs
- HBA Option ROM
- Host NIC
- Host NIC Option ROM
- Local Disk
- NVME Mswitch Firmware
- PSU
- Pci Switch Firmware

Procedure 9. Create network control policy for Cisco Discovery Protocol

Create a network control policy that enables the Cisco Discovery Protocol (CDP) on virtual network ports.

- Step 1.** In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > Network Control Policies**.
- Step 3.** Right-click **Network Control Policies**.
- Step 4.** Select **Create Network Control Policy**.
- Step 5.** Enter a policy name.
- Step 6.** For CDP select **Enabled**.

Step 7. Click **OK** to create the network control policy.

Create Network Control Policy ? X

Name : CDP_Enabled

Description : FlexPod-Motomel

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 10. Create power control policy

Create a power control policy for the Cisco UCS environment.

- Step 1.** In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > Power Control Policies**.
- Step 3.** Right-click **Power Control Policies**.
- Step 4.** Select **Create Power Control Policy**.
- Step 5.** For the Name, enter **NoPowerCap** for the power control policy name.
- Step 6.** For the Description, enter **FlexPod-Motomel** or a description of your choice.
- Step 7.** For Power Capping, select **No Cap**.
- Step 8.** Click **OK** to create the power control policy.

Create Power Control Policy ? X

Name :

Description :

Fan Speed Policy :

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 11. Create server BIOS policy

Create a server BIOS policy for the Cisco UCS environment.

- Step 1.** In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > BIOS Policies**.
- Step 3.** Right-click **BIOS Policies**.
- Step 4.** Select **Create BIOS Policy**.
- Step 5.** For the BIOS policy name, enter **B200M5-BIOS**.
- Step 6.** Click **OK** to create the policy.

Create BIOS Policy ? X

Name : B200M5-BIOS

Description : FlexPod-Motomel

Reboot on BIOS Settings Change :

OK Cancel

Step 7. Leave all BIOS Settings set to **Platform Default**.

Procedure 12. Configure maintenance policy

Update the default maintenance policy.

- Step 1.** In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > Maintenance Policies**.
- Step 3.** Right-click **Maintenance Policies** to create a new policy.
- Step 4.** Enter the name for Maintenance Policy
- Step 5.** Change the Reboot Policy to **User Ack**.
- Step 6.** Check **On Next Boot** for applying pending changes at next reboot.
- Step 7.** Click **OK** to create the policy.

Create Maintenance Policy ? X

Name : FlexPodMaint

Description : FlexPod-Motomel

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.)

OK Cancel

Procedure 13. Create vNIC templates

Create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment.

- Step 1.** In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > vNIC Template**.
- Step 3.** Right-click **vNIC Templates**.
- Step 4.** Select **Create vNIC Template**.
- Step 5.** Enter the name for the 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.** For Template Type, select **Updating Template**.
- Step 9.** Under VLANs, select the checkboxes for the VLANs to add as part of the vNIC Template.
- Step 10.** Select **Native VLAN** for the native VLAN.
- Step 11.** For MTU, enter **9000**.
- Step 12.** In the MAC Pool drop-down list, select your MAC Pool configured 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

Name : vNIC-A

Description : FlexPod-Motomel

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : <not set>

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	VLAN ID
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>	1
<input checked="" type="checkbox"/>	In-Band-Mgmt	<input type="radio"/>	60
<input checked="" type="checkbox"/>	Infra-Mgmt	<input type="radio"/>	61
<input checked="" type="checkbox"/>	CIFS-VLAN	<input type="radio"/>	62
<input checked="" type="checkbox"/>	NFS-Vlan	<input type="radio"/>	63
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>	66

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : FlexPod-MACPool(128/128)

QoS Policy : <not set>

Network Control Policy : CDP_Enabled

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

usNIC Connection Policy : <not set>

OK Cancel

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

Create vNIC Template

Name : vNIC-B

Description : FlexPod-Motomel

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : vNIC-A ▼

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

Properties for: vNIC Template vNIC-A

General | VLANs | VLAN Groups | Faults | Events

Actions

- Modify VLANs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Use Global

Properties

Name : **vNIC-A**

Description : FlexPod-Motomel

Owner : **Local**

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : vNIC-B ▼ [Create vNIC Template](#)

Procedure 14. Create vHBA templates

Create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment

- Step 1.** In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
- Step 2.** Select **Policies > root > Sub-Organization > FlexPod-CVD > vHBA Template**.
- Step 3.** Right-click **vHBA Templates**.
- Step 4.** Select **Create vHBA Template**.
- Step 5.** For the name, enter **vHBA-A**.
- Step 6.** For Fabric ID, select **Fabric A**.
- Step 7.** For Select VSAN, choose the VSAN created for Fabric A from the drop-down list.
- Step 8.** For Template Type, select **Updating Template**.
- Step 9.** For Max Data Field Size, enter **2048**.
- Step 10.** For WWPN Pool, from the drop-down list select the WWPN Pool for Fabric A (created earlier).

Step 11. Leave the remaining fields as-is.

Step 12. Click **OK**.

Create vHBA Template

Name : vHBA-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : VSAN-400-A [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : FlexPod-WWPNPoolA(20/32)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

OK **Cancel**

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

Create vHBA Template ? X

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 :

Procedure 15. Create server boot policy for SAN boot

All Cisco UCS B200 M5 Blade Servers for 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.

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 the **Servers tab > Policies > root > Sub-Organization > FlexPod-CVD > right-click Local Disk Configuration Policy > enter SAN-Boot** for 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

Procedure 16. Create SAN boot policy

The SAN Boot policy configures the SAN Primary's primary-target to be port AFF300-01:0g on the NetApp cluster and SAN Primary's secondary-target to be port AFF300-02:0g on the NetApp cluster. Similarly, the SAN Secondary's primary-target should be port AFF300-01:0h on the NetApp cluster and SAN Secondary's secondary-target should be port AFF300-02:0h on the NetApp cluster.

Step 1. Log into the storage controller and verify that all port information is correct. This information can be found in the **NetApp ONTAP System Manager** under **Network > Network Interfaces**:

Network Interfaces

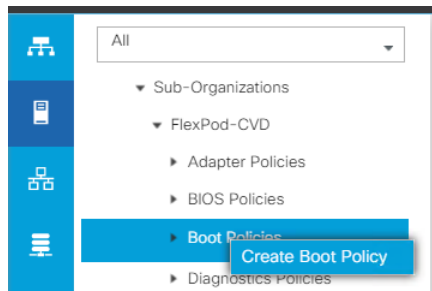
Interface Name	Storage Virtual Mac...	IP Address/WWPN	Current Port	Home Port	Data Protocol Access	Management Access	Subnet	Role	VIP LIF
fc_p_01a	Infra	20:01:00:a0:98:af:bd:e8	AFF-A300-01:0g	Yes	fc_p	No	-NA-	Data	No
fc_p_01b	Infra	20:02:00:a0:98:af:bd:e8	AFF-A300-01:0h	Yes	fc_p	No	-NA-	Data	No
fc_p_02a	Infra	20:03:00:a0:98:af:bd:e8	AFF-A300-02:0g	Yes	fc_p	No	-NA-	Data	No
fc_p_02b	Infra	20:04:00:a0:98:af:bd:e8	AFF-A300-02:0h	Yes	fc_p	No	-NA-	Data	No

Note

You have to create a SAN Primary (hba0) and a SAN Secondary (hba1) in SAN Boot Policy by entering WWPN of NetApp FC

Interfaces as explained in the following section.

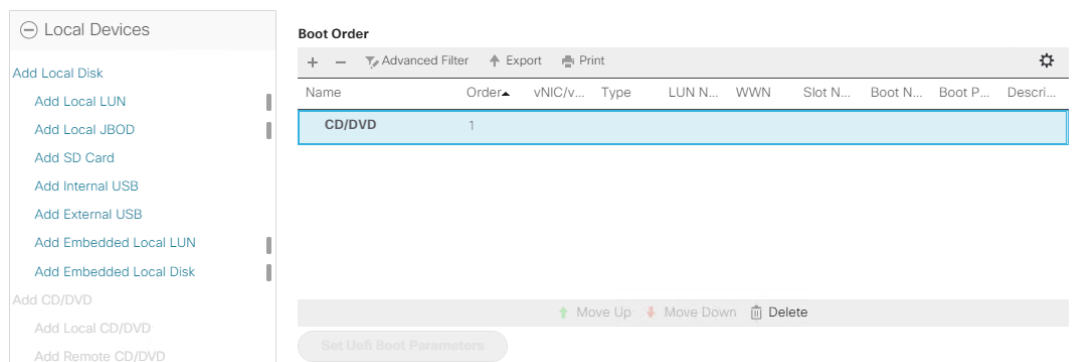
Step 2. From the Cisco UCS Manager and go to **Servers > Policies > root > Sub Organization > FlexPod-CVD > Boot Policies**. Right-click and select **Create Boot Policy**.



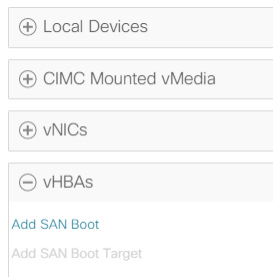
Step 3. For the name, enter **SAN** for the boot policy.

A screenshot of the 'Create Boot Policy' configuration form. The 'Name' field contains 'A300Boot'. The 'Description' field contains 'FlexPod-Motome1'. The 'Reboot on Boot Order Change' checkbox is unchecked. The 'Enforce vNIC/vHBA/iSCSI Name' checkbox is checked.

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



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

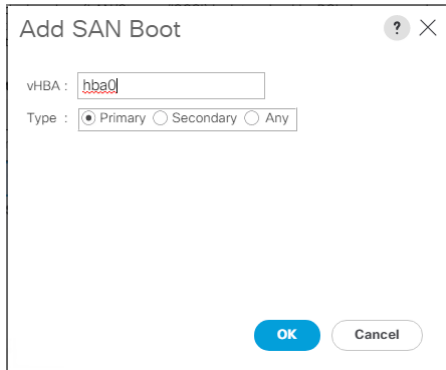


Note

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

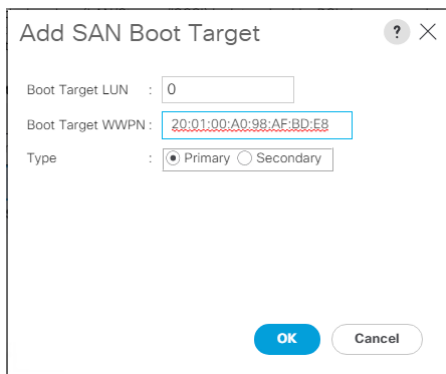
paths.

Step 6. In the Add SAN Boot dialog box, for vHBA enter **hba0**, and for Type select **Primary**. Click **OK** to add SAN Boot.



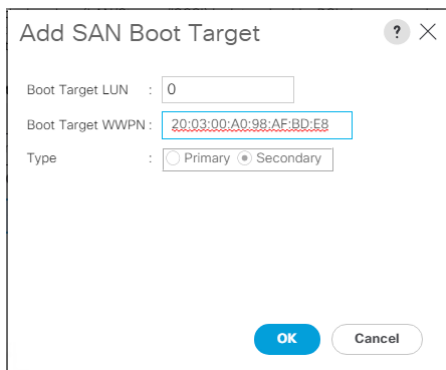
The screenshot shows a dialog box titled "Add SAN Boot" with a close button (X) and a help button (?). The "vHBA" field contains the text "hba0". Below it, the "Type" field has three radio buttons: "Primary" (which is selected), "Secondary", and "Any". At the bottom of the dialog are two buttons: "OK" and "Cancel".

Step 7. Select **Add SAN Boot Target** to enter the WWPN address of the storage port. For the Boot Target LUN enter **0** for the value. For the Boot Target WWPN, enter the WWPN for FC port AFF300-01:0g for NetApp and click **OK** to add the SAN Boot Primary Target.



The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) and a help button (?). The "Boot Target LUN" field contains "0". The "Boot Target WWPN" field contains "20:01:00:A0:98:AF:BD:E8". The "Type" field has two radio buttons: "Primary" (selected) and "Secondary". At the bottom are "OK" and "Cancel" buttons.

Step 8. Add a secondary SAN Boot target using the same hba0. For the Boot Target LUN, enter **0** and for the Boot Target WWPN, enter the WWPN for FC port AFF300-02:0g for NetApp, and for Type select **Secondary**.



The screenshot shows a dialog box titled "Add SAN Boot Target" with a close button (X) and a help button (?). The "Boot Target LUN" field contains "0". The "Boot Target WWPN" field contains "20:03:00:A0:98:AF:BD:E8". The "Type" field has two radio buttons: "Primary" and "Secondary" (which is selected). At the bottom are "OK" and "Cancel" buttons.

Step 9. Repeat steps 1-8 to add a secondary san boot for **hba1**.

Step 10. Click **OK** to create Boot Policy

Step 11. After creating the 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 A300Boot** to view the boot order in the right pane of the Cisco UCS Manager as shown below:

Boot Order

+ - Advanced Filter Export Print

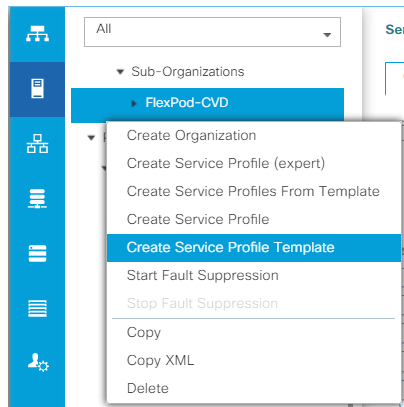
Name	vNIC/vH...	Type	WWN
▼ SAN Primary	hba0	Primary	...
SAN Target Primary	...	Primary	0. 20:01:00:A0:98:AF:BD:E8
SAN Target Secondary	...	Secondary	0. 20:03:00:A0:98:AF:BD:E8
▼ SAN Secondary	hba1	Secondary	...
SAN Target Primary	...	Primary	0. 20:02:00:A0:98:AF:BD:E8
SAN Target Secondary	...	Secondary	0. 20:04:00:A0:98:AF:BD:E8

Procedure 17. Configure and create a service profile template

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

Create two service profile templates; the first service profile template “VDI-FLEXPOD” for workload servers and the second service profile template “INFRA-FLEXPOD” for infrastructure servers.

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



Step 2. Enter a name for the Service Profile Template, for Type select **Updating Template**, for UUID Assignment select the UUID pool that was created earlier and click **Next**.

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-FlexPod-CVD**

The template will be created in the following organization. Its name must be unique within this organization.
Type : Initial Template Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.
UUID

UUID Assignment:

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

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

Step 3. For Local Storage, select **SAN-Boot** and for Mode, select **No Local Storage**.

Create Service Profile Template

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Specific Storage Profile Storage Profile Policy **Local Disk Configuration Policy**

Local Storage:

[Create Local Disk Configuration Policy](#)

Mode : **No Local Storage**

Protect Configuration : **Yes**

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : **Disable**

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**

FlexFlash Removable State : **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.

< Prev Next > **Finish** Cancel

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.

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
vNIC eth1	Derived	derived	
vNIC eth0	Derived	derived	

[Delete](#) [Add](#) [Modify](#)

[+ iSCSI vNICs](#)

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

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

Step 5. For the first vNIC Template name, enter **eth0**, for vNIC Template select **vNIC-A**, and for the Adapter Policy, select **VMware**:

Create vNIC

Name:

Use vNIC Template:

Redundancy Pair:

vNIC Template:

Peer Name:

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy:

[Create Ethernet Adapter Policy](#)

Step 6. For the second vNIC Template name, enter **eth1**, for the vNIC Template select **vNIC-B**, and for the Adapter Policy, select **VMware**.

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

Step 7. When the vNICs are created, click **Next**.

Step 8. In the SAN Connectivity menu, for SAN connectivity select **Expert**. For WWNN Assignment, select the WWNN pool that you created previously. Click **Add** to add vHBAs.

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

WWNN Assignment: FlexPod-WWNNPool(32/32)

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

Name	WWPN
	No data available

[Delete](#) [Add](#) [Modify](#)

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

The following two HBAs were created:

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

Figure 40. vHBA0

Create vHBA

Name : vHBA0

Use vHBA Template :

Redundancy Pair : Peer Name :

vHBA Template : vHBA-A

[Create vHBA Template](#)

Adapter Performance Profile

Adapter Policy : VMWare

[Create Fibre Channel Adapter Policy](#)

Figure 41. vHBA1

Modify vHBA



Name : vHBA1

Use vHBA Template :

Create vHBA Template

vHBA Template : vHBA-B ▼

Adapter Performance Profile

Adapter Policy : VMWare ▼

[Create Fibre Channel Adapter Policy](#)

Figure 42. All vHBAs

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

WWNN Assignment: FlexPod-WWNNPool(32/32) ▼

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

Name	WWPN
▶ vHBA vHBA1	Derived
▶ vHBA vHBA0	Derived

Delete Add Modify

< Prev Next > Finish Cancel

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

Step 10. For Select Placement, select **Let System Perform Placement**. Click **Next**.

Create Service Profile Template

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [Create Placement Policy](#)

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vHBA vHBA0	Derived	1
vHBA vHBA1	Derived	2
vNIC eth0	Derived	3
vNIC eth1	Derived	4

[Move Up](#)
[Move Down](#)
[Delete](#)
[Reorder](#)
[Modify](#)

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

Step 11. For the Boot Policy, select **A300Boot** for the Boot Policy which you created earlier.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **A300Boot**
 Description : **FlexPod-Motomel**
 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

+ - ▾ Advanced Filter ↑ Export 🖨 Print ⚙

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
CD/DVD	1								
▾ San	2								
▶ SAN Primary		hba0	Primary						
▶ SAN Secondary		hba1	Second...						

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

< Prev Next > **Finish** Cancel

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. For the Maintenance Policy, select **UserAck**, which requires user acknowledgement prior to re-booting the server when making changes to the policy or pool configuration tied to a service profile.

Create Service Profile Template
? X

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

⊖
Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: Create Maintenance Policy

Name	:	User-Ack
Description	:	User-Ackn_Policy
Soft Shutdown Timer	:	150 Secs
Storage Config. Deployment Policy	:	User Ack
Reboot Policy	:	User Ack

< Prev
Next >
Finish
Cancel

- 1 Identify Service Profile Template
- 2 Storage Provisioning
- 3 Networking
- 4 SAN Connectivity
- 5 Zoning
- 6 vNIC/vHBA Placement
- 7 vMedia Policy
- 8 Server Boot Order
- 9
- Maintenance Policy
- 10 Server Assignment
- 11 Operational Policies

Step 13. For the Pool Assignment, you can select a server pool policy to automatically assign a service profile to a server that meets the requirement for the server qualification based on the pool configuration or you can select **Assign Later** if you want manual assignment.

Step 14. On the same page you can configure the **Host Firmware Package Policy** that helps to keep the firmware in sync when associated to server.

Create Service Profile Template

?
×

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Assign Later ▼ [Create Server Pool](#)

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile template is not automatically associated with a server. Either select a server from the list or associate the service profile manually later.

⊖
Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: FlexPod-404 ▼

[Create Host Firmware Package](#)

< Prev
Next >
Finish
Cancel

1

 Identify Service Profile Template

2

 Storage Provisioning

3

 Networking

4

 SAN Connectivity

5

 Zoning

6

 vNIC/vHBA Placement

7

 vMedia Policy

8

 Server Boot Order

9

 Maintenance Policy

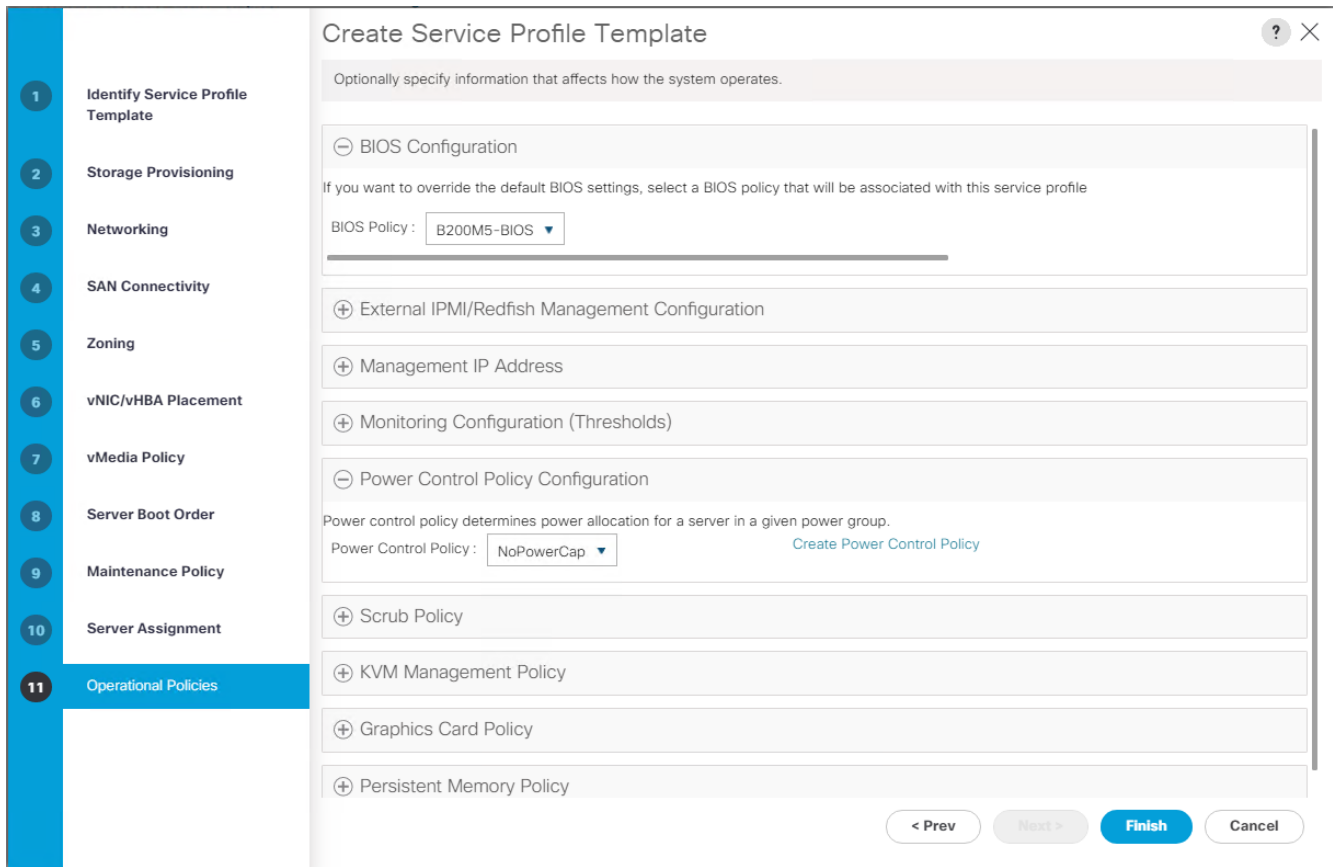
10

Server Assignment

11

 Operational Policies

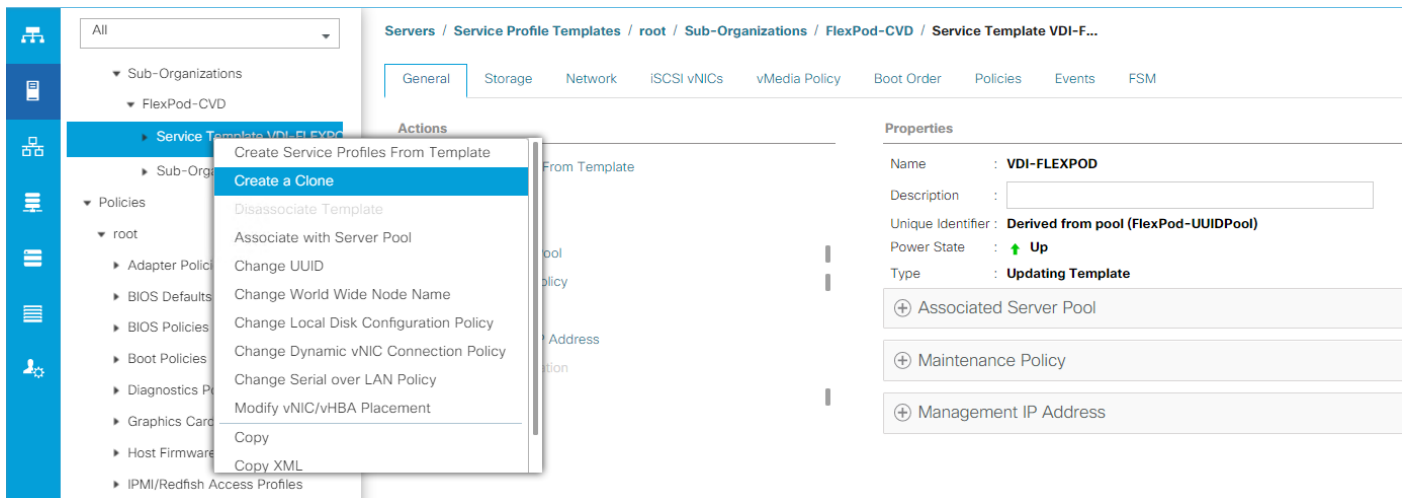
Step 15. For this solution we selected the following on the Operational Policy page (select what's appropriate for your solution): for the BIOS Policy, select **B200M5-BIOS**, for the Power Control Policy, select **NoPowerCap** for maximum performance, and for the Graphics Card Policy, select **B200M5** server configured with **Nvidia P6 GPU card**.



Step 16. Click **Next**, then click **Finish**, and click **OK** to create the “VDI-FLEXPOD” service profile template.

Procedure 18. 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-FLEXPOD** > right-click **Create a Clone** as shown below:



Step 2. For the Clone Name, enter INFRA-FLEXOD. Click **OK**.

Create Clone From VDI-FLEXPOD ✕

Clone Name :

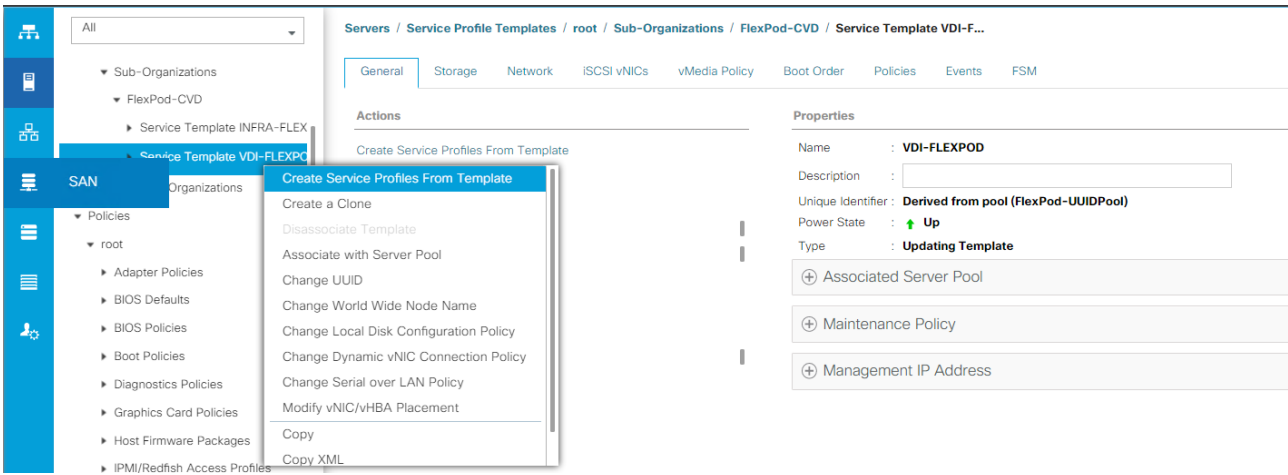
Org :

Procedure 19. Create Service Profiles from Template and associate to servers

Create thirty Service profiles from the VDI-FLEXPOD template and two Service profiles from the INFRA-FLEXPOD template as explained in the following sections.

Create the first four Service Profiles from Template.

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



Step 2. For the Service Profile Template, select **VDI-FLEXPOD** for the service profile template that you created earlier. For the Naming Prefix, **SP-VDIX**.

Step 3. To create thirty service profiles, for Number of Instances, enter **30**. This process will create service profiles “SP-VDI1”, “SP-VDI2”, and “SP-VDI30.”

Create Service Profiles From Template

Naming Prefix : SP-VDI

Name Suffix Starting Number : 1

Number of Instances : 30

OK Cancel

Step 4. Create the remaining two Service Profiles “SP_Infra1” and “SP_Infra2” from Template “INFRA-FLEXPOD.”

Service Profiles Association

When the service profiles are created, the association of the Service Profile automatically starts for servers based on the Server Pool Policies or can be manually assigned if none are configured. To associate the Service Profile manually, go to **Equipment > Chassis > right-click the Server > select Associate Service Profile.**

Configure Cisco Nexus 93180YC-FX Switches

The following section details the steps for the Nexus 93180YC-FX switch configuration.

Procedure 1. Configure Global Settings for Cisco Nexus A and Cisco Nexus B

Step 1. Log in as **admin** user into the Nexus Switch A and run the following commands to set global configurations and jumbo frames in QoS:

```
conf terminal
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9216
exit
class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy run start
```

Step 2. Log in as **admin** user into the Nexus Switch B and run the same commands (above) to set global configurations and jumbo frames in QoS.

Procedure 2. Configure VLANs for Cisco Nexus A and Cisco Nexus B switches

Create the necessary virtual local area networks (VLANs) on both Nexus switches. We created VLAN 60, 61, 62, 63, 66, 68 and 102.

Step 1. Log in as **admin** user into the Nexus Switch A.

Step 2. Create VLAN 60:

```
config terminal
VLAN 60
name In-Band-Mgmt
no shutdown
exit
copy running-config startup-config
exit
```

Step 3. Log in as **admin** user into the Nexus Switch B and create the same VLANs.

Virtual Port Channel (vPC) summary for data and storage network

In the Cisco Nexus 93180YC-FX switch topology, a single vPC feature is enabled to provide high availability, faster convergence in the event of a failure, and greater throughput. Cisco Nexus 93180YC-FX vPC configurations with the vPC domains and corresponding vPC names and IDs for Oracle Database Servers are listed in Table 9.

Table 9. vPC Summary

vPC Domain	vPC Name	vPC ID
10	Peer-Link	1
10	vPC Port-Channel to FI-A	51
10	vPC Port-Channel to FI-B	52
10	vPC Port-Channel to AFF-A300-1	53
10	vPC Port-Channel to AFF-A300-2	54

As listed in Table 9, a single vPC domain with Domain ID 10 is created across two Cisco Nexus 93180YC-FX member switches to define vPC members to carry specific VLAN network traffic. In this topology, we defined a total number of 4 vPCs:

- vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B.
- vPC IDs 51 and 52 are defined for traffic from Cisco UCS fabric interconnects.
- vPC IDs 53 and 54 are defined for traffic from NetApp AFF A300 Controllers.

Cisco Nexus 93180YC-FX switch cabling details

The following tables list the cabling information.

Table 10. Cisco Nexus 93180YC-FX-A cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch A	Eth1/49	40GbE	Cisco Nexus 93180YC-FX B	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 93180YC-FX B	Eth1/50
	Eth1/51	40GbE	Cisco UCS fabric interconnect A	Eth1/49
	Eth1/52	40GbE	Cisco UCS fabric interconnect B	Eth1/50
	Eth1/53	40GbE	NetApp AFF-A300-1	300-01:0g
	Eth1/54	40GbE	NetApp AFF-A300-2	300-02:0g
	Eth1/51	40GbE	Cisco UCS fabric interconnect A	/49

Table 11. Cisco Nexus 93180YC-FX-B cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch B	Eth1/49	40GbE	Cisco Nexus 93180YC-FX A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 93180YC-FX A	Eth1/50

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/51	40GbE	Cisco UCS fabric interconnect B	Eth1/49
	Eth1/52	40GbE	Cisco UCS fabric interconnect A	Eth1/50
	Eth1/53	40GbE	NetApp AFF-A300-1	AFF300-01:0h
	Eth1/54	40GbE	NetApp AFF-A300-2	AFF300-02:0h
	Eth1/49	40GbE	Cisco Nexus 93180YC-FX A	Eth1/49

Cisco UCS Fabric Interconnect 6454 cabling

The following tables list the FI 6454 cabling information.

Table 12. Cisco UCS Fabric Interconnect (FI) A cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-A Chassis 1-4	FC 1/1	32G FC	Cisco MDS 9132T 32-Gbps-A	FC 1/1
	FC 1/2	32G FC	Cisco MDS 9132T 32-Gbps-A	FC 1/2
	Eth1/17-24	25GbE	UCS 5108 Chassis IOM-A Chassis 1-4	IO Module Port1-2
	Eth1/49	40GbE	Cisco Nexus 93180YC-FX Switch A	Eth1/51
	Eth1/50	40GbE	Cisco Nexus 93180YC-FX Switch B	Eth1/51
	Mgmt 0	1GbE	Management Switch	Any
	L1	1GbE	Cisco UCS FI - A	L1

Local Device	Local Port	Connection	Remote Device	Remote Port
	L2	1GbE	Cisco UCS FI - B	L2

Table 13. Cisco UCS Fabric Interconnect (FI) B cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-B	FC 1/1	32G FC	Cisco MDS 9132T 32-Gbps-B	FC 1/1
	FC 1/2	32G FC	Cisco MDS 9132T 32-Gbps-B	FC 1/2
	Eth1/17-24	25GbE	UCS 5108 Chassis IOM-B	
	Chassis 1-4	IO Module Port1-2		
	Eth1/49	40GbE	Cisco Nexus 93180YC-FX Switch A	Eth1/52
	Eth1/50	40GbE	Cisco Nexus 93180YC-FX Switch B	Eth1/52
	Mgmt 0	1GbE	Management Switch	Any
	L1	1GbE	Cisco UCS FI - A	L1

Procedure 3. Create vPC peer-link between the two Nexus switches

Step 1. Log in as **admin** user into the Nexus Switch A.

Note

For vPC 10 as Peer-link, we used interfaces 49-50 for Peer-Link. You may choose the appropriate number of ports for your needs.

Step 2. Create the necessary port channels between devices on both Nexus Switches:

```
config terminal
feature vpc
feature lacp
vpc domain 10
```



```
peer-keepalive destination 10.29.164.66 source 10.29.164.65
exit
interface port-channel10
description VPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1-2,60-70,102
spanning-tree port type network
vpc peer-link
exit
interface Ethernet1/53
description vPC PeerLink between 9ks
switchport mode trunk
switchport trunk allowed VLAN 1-2,60-70,102
channel-group 10 mode active
no shutdown
exit
interface Ethernet1/54
description vPC PeerLink between 9ks
switchport mode trunk
switchport trunk allowed VLAN 1-2,60-70,102
channel-group 10 mode active
no shutdown
exit
```

Step 3. Log in as **admin** user into the Nexus Switch B and repeat the above steps to configure second nexus switch.

Make sure to change peer-keepalive destination and source IP address appropriately for Nexus Switch B.

Create vPC configuration between Nexus 93180YC-FX and fabric interconnects

Create and configure vPC 51 and 52 for the data network between the Nexus switches and fabric interconnects.

Procedure 1. Create the necessary port channels between devices on both Nexus switches

Step 1. Log in as **admin** user into Nexus Switch A and enter the following:

```
config Terminal
interface port-channel51
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
```

```
vpc 51
no shutdown
exit
interface port-channel52
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
vpc 52
no shutdown
exit
interface Ethernet1/51
description FI-A-N9K-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
mtu 9216
speed 40000
no negotiate auto
channel-group 51 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-N9K-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
channel-group 52 mode active
no shutdown
exit
copy running-config startup-config
```

Step 2. Log in as **admin** user into the Nexus Switch B and repeat Step 1 (above) to configure the second Nexus switch.

Procedure 2. Create vPC configuration between Nexus 93180YC-FX and NetApp AFF A300

Create and configure vPC 53 and 54 for the data network between the Nexus switches NetApp AFF A300 controllers.

Step 1. Log in as **admin** user into Nexus Switch A and enter the following:

```
config Terminal
interface port-channel53
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  spanning-tree port type edge trunk
  mtu 9216
  vpc 53
  no shutdown
exit
interface port-channel54
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  spanning-tree port type edge trunk
  mtu 9216
  vpc 54
  no shutdown
exit
interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  mtu 9216
  channel-group 53 mode active
  no shutdown
  exit
interface Ethernet1/54
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  mtu 9216
  channel-group 54 mode active
  no shutdown
  exit
copy running-config startup-config
```

Step 2. Log in as **admin** user into the Nexus Switch B and complete steps above to configure the second Nexus switch.

Verify all vPC status is up on both Cisco Nexus Switches

Figure 43. shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 43. vPC Description for Cisco Nexus Switch A and B

```

Switch A:
vPC domain id      : 10
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
P2P-vlan consistency status : success
Type-2 consistency status : success
vPC role           : primary
Number of vPCs configured : 4
Peer Gateway       : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled, timer is off. (timeout = 240s)
Auto-recovery status : Enabled, timer is off. (timeout = 150s)
Delay-restore status : Timer is off. (timeout = 10s)
Operational Layer3 Peer-router : Disabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  ---  -
0   Po10  up    1-2,60-70,102

vPC status
-----
id  Port  Status Consistency Reason  Active vlans
--  ---  -
10  Po10  up    success success          1-2,60-70,102
11  Po11  up    success success          1-2,60-70,102
12  Po12  up    success success          1-2,60-70,102
13  Po13  up    success success          1-2,60-70,102
14  Po14  up    success success          1-2,60-70,102

Switch B:
vPC domain id      : 10
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
P2P-vlan consistency status : success
Type-2 consistency status : success
vPC role           : secondary
Number of vPCs configured : 4
Peer Gateway       : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled, timer is off. (timeout = 240s)
Auto-recovery status : Enabled, timer is off. (timeout = 150s)
Delay-restore status : Timer is off. (timeout = 10s)
Operational Layer3 Peer-router : Disabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  ---  -
1   Po10  up    1-2,60-70,102

vPC status
-----
id  Port  Status Consistency Reason  Active vlans
--  ---  -
15  Po15  up    success success          1-2,60-70,102
16  Po16  up    success success          1-2,60-70,102
17  Po17  up    success success          1-2,60-70,102
18  Po18  up    success success          1-2,60-70,102
19  Po19  up    success success          1-2,60-70,102
20  Po20  up    success success          1-2,60-70,102

```

Cisco MDS 9132T 32-Gbps FC switch configuration

Figure 23. illustrates the cable connectivity between the Cisco MDS 9132T 32-Gbps switch and the Cisco 6332 Fabric Interconnects and NetApp AFF A300 storage.

Note

We used two 32Gb connections from each Fabric Interconnect to each MDS switch and two 32Gb FC connections from each NetApp AFF A300 array controller to each MDS switch.

Table 14. Cisco MDS 9132T-A cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-A	FC1/1	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/2
	FC1/2	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/2
	FC1/3	32Gb FC	NetApp AFF300-01 Controller	0g
	FC1/4	32Gb FC	NetApp AFF300-02 Controller	0g

Table 15. Cisco MDS 9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-B	FC1/1	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/2
	FC1/2	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/2
	FC1/3	32Gb FC	NetApp AFF300-01 Controller	0h
	FC1/4	32Gb FC	NetApp AFF300-02 Controller	0h

NetApp AFF A300 to MDS SAN fabric connectivity

NetApp AFF A300 to MDS A and B switches using VSAN 400 for Fabric A and VSAN 401 configured for Fabric B

In this solution, two ports (ports FC1/3 and FC1/4) of MDS Switch A and two ports (ports FC1/3 and FC1/4) of MDS Switch B connected to NetApp AFF A300 Storage System as shown in Table 16. All ports are connected to the NetApp Array carry 32 Gb/s FC Traffic.

Table 16. MDS 9132T 32-Gbps switch port connection to NetApp system

Local Device	Local Port	Connection	Remote Device	Remote Port
MDS Switch A	FC1/3	32Gb FC	NetApp AFF300-01 Controller	0g
	FC1/4	32Gb FC	NetApp AFF300-02 Controller	0g
MDS Switch B	FC1/3	32Gb FC	NetApp AFF300-01 Controller	0h
	FC1/4	32Gb FC	NetApp AFF300-02 Controller	0h

Procedure 1. Configure feature for MDS switch A and MDS switch B

Set feature on MDS Switches on both MDS switches:

Step 1. Log in as **admin** user into MDS Switch A:

```
config terminal
feature npiv
```

```
feature telnet
switchname MDS-A
copy running-config startup-config
```

Step 2. Log in as **admin** user into MDS Switch B. Repeat the Step 1 (above) on MDS Switch B.

Procedure 2. Configure VSANs for MDS switch A and MDS switch B

Create VSANs on both MDS switches:

Step 1. Log in as **admin** user into MDS Switch A. Create VSAN 400 for Storage Traffic:

```
config terminal
VSAN database
vsan 400
vsan 400 interface fc 1/1-4
exit
interface fc 1/1-4
    switchport trunk allowed vsan 400
    switchport trunk mode off
    port-license acquire
    no shutdown
    exit
copy running-config startup-config
```

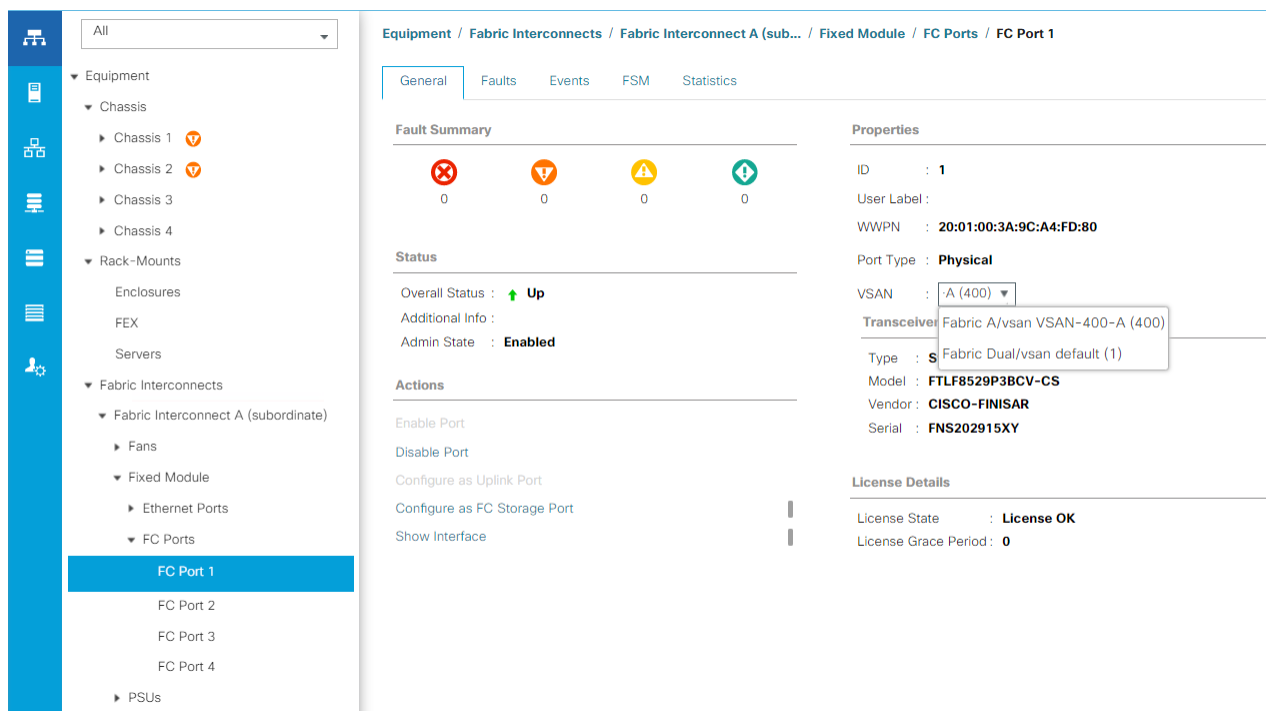
Step 2. Log in as **admin** user into MDS Switch B. Create VSAN 401 for Storage Traffic:

```
config terminal
VSAN database
vsan 401
vsan 401 interface fc 1/1-4
exit
interface fc 1/1-4
    switchport trunk allowed vsan 401
    switchport trunk mode off
    port-license acquire
    no shutdown
    exit
copy running-config startup-config
```

Procedure 3. Add FC uplink ports to corresponding VSAN on fabric interconnect

Step 1. In Cisco UCS Manager in the Equipment tab, select **Fabric Interconnects > Fabric Interconnect A > Fixed Module > FC Ports**.

Step 2. Select **FC Port 1** and from the VSAN drop-down list select **VSAN 400**.



Step 3. Repeat Steps 1 and 2 to add the FC Port 1-2 to VSAN 400 on Fabric A and FC Port 1-2 to VSAN 401 on Fabric B.

Procedure 4. Create and configure fiber channel zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9132T 32-Gbps switches, the Cisco UCS Fabric Interconnects, and the NetApp AFF system.

Before you configure the zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs from each server. We used 2 HBAs for each Server. HBA0 is connected to MDS Switch-A and is connected to MDS Switch-B.

Step 1. Log into the **Cisco UCS Manager > Equipment > Chassis > Servers** and select the desired server.

Step 2. Click the **Inventory** tab and then click the **HBAs** tab to get the WWPN of HBAs as shown in the screenshot below:

Name	vHBA	Vendor	Model	Operability	WWPN	Original WWPN
HBA 1	vHBA-A	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:85:3A:00:3F	00:00:00:00:00:00:00
HBA 2	vHBA-B	Cisco Systems Inc	UCSB-MLOM-40G-03	Operable	20:00:00:25:05:06:00:3F	00:00:00:00:00:00:00

Step 3. Connect to the **NetApp ONTAP System Manager** and extract the WWPN of FC Ports connected to the Cisco MDS Switches.

Network Interfaces

Interface Name	Storage Virtual Mac...	IP Address/WWN	Current Port	Home Port	Data Protocol Access	Management Access	Subnet	Role	VIP LIF
✓ fcp_01a	Infra	20:01:00:a0:98:af:bd:e8	AFF-A300-01:0g	Yes	fcp	No	-NA-	Data	No
✓ fcp_01b	Infra	20:02:00:a0:98:af:bd:e8	AFF-A300-01:0h	Yes	fcp	No	-NA-	Data	No
✓ fcp_02a	Infra	20:03:00:a0:98:af:bd:e8	AFF-A300-02:0g	Yes	fcp	No	-NA-	Data	No
✓ fcp_02b	Infra	20:04:00:a0:98:af:bd:e8	AFF-A300-02:0h	Yes	fcp	No	-NA-	Data	No

Procedure 5. Create device aliases for fiber channel zoning

Cisco MDS Switch A:

Step 1. Log in as **admin** user and run the following commands:

```
conf t
device-alias database
device-alias name VDI-1-HBA1 pwwn 20:00:00:25:b5:3a:00:3f
device-alias name a300-01-0g pwwn 20:01:00:a0:98:af:bd:e8
device-alias name a300-02-0g pwwn 20:03:00:a0:98:af:bd:e8
```

Cisco MDS Switch B:

Step 2. Log in as **admin** user and run the following commands:

```
conf t
device-alias database
device-alias name VDI-1-HBA2 pwwn 20:00:00:25:d5:06:00:3f
device-alias name a300-01-0h pwwn 20:02:00:a0:98:af:bd:e8
device-alias name a300-02-0h pwwn 20:04:00:a0:98:af:bd:e8
```

Procedure 6. Create zoning for Cisco MDS switch A

Configure zones for the MDS switch A for each server service profile.

Step 1. Log in as **admin** user and create the zone:

```
conf t
zone name a300_VDI-1-HBA1 vsan 400
  pwwn 20:00:00:25:b5:3a:00:3f [VDI-1-HBA1]
  pwwn 20:01:00:a0:98:af:bd:e8 [a300-01-0g]
  pwwn 20:03:00:a0:98:af:bd:e8 [a300-02-0g]
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members:

```
conf t
zoneset name FlexPod_FabricA vsan 400
member a300_VDI-1-HBA1
Activate the zone set by running following commands:
zoneset activate name FlexPod_FabricA vsan 400
exit
```



```
copy running-config startup-config
```

Procedure 7. Create zoning for Cisco MDS switch B

Configure zones for the MDS switch A for each server service profile.

Step 1. Log in as admin user and create the zone:

```
conf t
zone name a300_VDI-1-HBA2 vsan 401
  pwwn 20:00:00:25:d5:06:00:3f [VDI-1-HBA2]
  pwwn 20:02:00:a0:98:af:bd:e8 [a300-01-0h]
  pwwn 20:04:00:a0:98:af:bd:e8 [a300-02-0h]
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members:

```
conf t
zoneset name FlexPod_FabricB vsan 401
member a300_VDI-1-HBA2
Activate the zone set by running following commands:
zoneset activate name FlexPod_FabricB vsan 401
exit
copy running-config startup-config
```

Configure NetApp AFF A300 storage

The storage components for this reference architecture are composed of one AFF A300 high availability pair and one DS224C disk with 24x 3.49TB SSDs. This configuration delivers 65TB of usable storage and over 200TB effective storage with deduplication, compression and compaction, and the potential for over 300,000 IOPs, depending on the application workload.

This section details the specific storage system configuration used in this solution validation. This section does not include all the possible configuration options, only those necessary to support this solution.

Cluster details

A cluster consists of one or more nodes grouped as high availability pairs to form a scalable cluster. Creating a cluster enables the nodes to pool their resources and distribute work across the cluster, while presenting administrators with a single entity to manage. Clustering also enables continuous service to end users if individual nodes go offline.

Table 17 lists the cluster details.

Table 17. Cluster details

Cluster Name	ONTAP Version	Node Count	Data SVM Count	Cluster Raw Capacity
AFF A300	9.6P4	2	1	83.76TB

Storage details

Table 18 lists the storage details for each HA pair.

Table 18. Storage details

Node Names	Shelf Count	Disk Count	Disk Capacity	Raw Capacity
AFF A3 AFF-A300-01	DS224-12: 1	SSD: 24	SSD: 83.76TB	83.76TB
AFF-A300-0200				

Raw capacity is not the same as usable capacity.

Drive allocation details

Table 19 lists the drive allocation details for each node.

Table 19. Drive allocation details

Node Name	Total Disk Count	Allocated Disk Count	Disk Type	Raw Capacity	Spare Disk Count
AFF-A300-01	12	12	3.49TB_SSD	41.92TB	0
AFF-A300-02	12	12	3.49TB_SSD	41.92TB	0

Raw capacity is not the same as usable capacity.

Adapter card details

Table 20 lists the adapter cards present in each node.

Table 20. Adapter card details

Node Name	System Model	Slot Number	Part Number	Description
AFF-A300-01	AFF A300	1	110-00401	PMC PM8072; PCI-E quad-port SAS (PM8072)
AFF-A300-01	AFF A300	2	XL710	NIC,2x 10/40GbE,QSFP+
AFF-A300-02	AFF A300	1	110-00401	PMC PM8072; PCI-E quad-port SAS (PM8072)
AFF-A300-02	AFF A300	2	XL710	NIC,2x 10/40GbE,QSFP+

Firmware details

Table 21 lists the relevant firmware details for each node.

Table 21. Firmware details

Node Name	Node Firmware	Shelf Firmware	Drive Firmware	Remote Mgmt Firmware
AFF-A300-01	AFF-A300: 11.5	IOM12: A:0240, B:0240	X357_S163A3T8ATE: NA51	SP: 5.6P2
AFF-A300-02	AFF-A300: 11.5	IOM12: A:0240, B:0240	X357_S163A3T8ATE: NA51	SP: 5.6P2

Network port settings

You can modify the MTU, autonegotiation, duplex, flow control, and speed settings of a physical network port or interface group (ifgrp).

Table 22 lists the network port settings.

Table 22. Network port settings for ONTAP

Node Name	Port Name	Link Status	Port Type	MTU Size	Speed (Mbps)	Node Name	Port Name
AFF-A300-01	a0a	Up	if_group	9000	Auto/-	Default	-
AFF-A300-01	a0a-61	Up	VLAN	1500	Auto/-	Default	IB
AFF-A300-01	a0a-62	Up	VLAN	1500	Auto/-	Default	CIFS
AFF-A300-01	a0a-63	Up	VLAN	9000	Auto/-	Default	NFS
AFF-A300-01	e0a	Up	Physical	9000	Auto/10000	Cluster	Cluster
AFF-A300-01	e0b	Up	Physical	9000	Auto/10000	Cluster	Cluster
AFF-A300-	e0c	Dow	Physical	1500	1000/-	Default	Default

Node Name	Port Name	Link Status	Port Type	MTU Size	Speed (Mbps)	Node Name	Port Name
01		n					
AFF-A300-01	e0d	Down	Physical	1500	1000/-	Default	-
AFF-A300-01	e0e	Down	Physical	1500	1000/-	Default	-
AFF-A300-01	e0f	Down	Physical	1500	1000/-	Default	-
AFF-A300-01	e0M	Up	Physical	1500	Auto/1000	Default	Default
AFF-A300-01	e2a	Up	Physical	9000	1000/40000	Default	-
AFF-A300-01	e2e	Up	Physical	9000	1000/40000	Default	-
AFF-A300-02	a0a	Up	if_group	9000	Auto/-	Default	-
AFF-A300-02	a0a-61	Up	VLAN	1500	Auto/-	Default	IB
AFF-A300-02	a0a-62	Up	VLAN	1500	Auto/-	Default	CIFS
AFF-A300-02	a0a-63	up	VLAN	9000	Auto/-	Default	NFS
AFF-A300-02	e0a	up	Physical	9000	1000/10000	Cluster	Cluster
AFF-A300-02	e0b	up	Physical	9000	1000/10000	Cluster	Cluster
AFF-A300-02	e0c	Down	Physical	1500	1000/-	Default	Default

Node Name	Port Name	Link Status	Port Type	MTU Size	Speed (Mbps)	Node Name	Port Name
AFF-A300-02	e0d	Down	Physical	1500	1000/-	Default	-
AFF-A300-02	e0e	Down	Physical	1500	1000/-	Default	-
AFF-A300-02	e0f	Down	Physical	1500	1000/-	Default	-
AFF-A300-02	e0M	Up	Physical	1500	Auto/1000	Default	Default
AFF-A300-02	e2a	Up	Physical	9000	1000/40000	Default	-
AFF-A300-02	e2e	Up	Physical	9000	1000/40000	Default	-

Network port interface group settings

An ifgrp is a port aggregate containing two or more physical ports that act as a single trunk port. Expanded capabilities include increased resiliency, increased availability, and load distribution. You can create three different types of ifgrps on your storage system: single-mode, static multimode, and dynamic multimode. Each interface group provides different levels of fault tolerance. Multimode ifgrps provide methods for load balancing network traffic.

Table 23 lists the network port ifgrp settings.

Table 23. Network port Ifgrp settings

Node Name	Ifgrp Name	Mode	Distribution Function	Ports
AFF-A300-01	a0a	multimode_lacp	port	e2a, e2e
AFF-A300-02	a0a	multimode_lacp	port	e2a, e2e

Network routes

You control how LIFs in an SVM use your network for outbound traffic by configuring routing tables and static routes.

- Routing tables. Routes are configured for each SVM and identify the SVM, subnet, and destination. Because routing tables are for each SVM, routing changes to one SVM do not alter the route table of another SVM.

Routes are created in an SVM when a service or application is configured for the SVM. Like data SVMs, the admin SVM of each IPspace has its own routing table because LIFs can be owned by admin SVMs and might need route configurations different from those on data SVMs.

If you define a default gateway when creating a subnet, a default route to that gateway is added automatically to the SVM that uses a LIF from that subnet.

- Static route. This route is a defined route between a LIF and a specific destination IP address. The route can use a gateway IP address.

Table 24 lists the network routes for Data ONTAP 8.3 or later.

Table 24. Network routes

Cluster Name	SVM Name	Destination Address	Gateway Address	Metric	LIF Names
AFF-A300	AFF-A300	0.0.0.0/0	10.29.164.1	20	AFF-A300-01_mgmt1 AFF-A300-02_mgmt1 cluster_mgmt
AFF-A300	Infra	0.0.0.0/0	10.10.62.1	20	CIFS1-01 CIFS2-02

Network port broadcast domains

Broadcast domains enable you to group network ports that belong to the same layer 2 network. The ports in the group can then be used by an SVM for data or management traffic. A broadcast domain resides in an IPspace.

During cluster initialization, the system creates two default broadcast domains:

- The default broadcast domain contains ports that are in the default IPspace. These ports are used primarily to serve data. Cluster management and node management ports are also in this broadcast domain.
- The cluster broadcast domain contains ports that are in the cluster IPspace. These ports are used for cluster communication and include all cluster ports from all nodes in the cluster.

Table 25 lists the network port broadcast domains for Data ONTAP 8.3 or later.

Table 25. Network port broadcast domains

Cluster Name	Broadcast Domain	IPspace Name	MTU Size	Subnet Names	Port List	Failover Group Names
AFF-A300	CIFS	Default	1500	-	AFF- A300- 01:a0a- -62 AFF- A300- 02:a0a- -62	CIFS
AFF-A300	Infra	0.0.0.0/0	10.10.62.1	20	AFF- A300- 01:e0a AFF- A300- 01:e0 b AFF- A300- 02:e0a AFF- A300- 02:e0 b	Cluster
AFF-A300	Default	Default	1500	-	AFF- A300- 01:e0c AFF- A300- 01:e0 M AFF- A300- 02:e0c AFF- A300- 02:e0 M	Default

Cluster Name	Broadcast Domain	IPspace Name	MTU Size	Subnet Names	Port List	Failover Group Names
AFF-A300	IB	Default	1500	-	AFF-A300-01:a0a-61 AFF-A300-02:a0a-61	IB
AFF-A300	NFS	Default	9000	-	AFF-A300-01:a0a-63 AFF-A300-02:a0a-63	NFS

Aggregate configuration

Aggregates are containers for the disks managed by a node. You can use aggregates to isolate workloads with different performance demands, to tier data with different access patterns, or to segregate data for regulatory purposes.

For business-critical applications that need the lowest possible latency and the highest possible performance, you might create an aggregate consisting entirely of SSDs.

To tier data with different access patterns, you can create a hybrid aggregate, deploying flash as high-performance cache for a working data set, while using lower-cost HDDs or object storage for less frequently accessed data. A Flash Pool consists of both SSDs and HDDs. A Fabric Pool consists of an all-SSD aggregate with an attached object store.

If you need to segregate archived data from active data for regulatory purposes, you can use an aggregate consisting of capacity HDDs, or a combination of performance and capacity HDDs.

Table 26 lists the aggregate configuration information.

Table 26. Aggregate configuration

Aggregate Name	Home Node Name	State	RAID Status	RAID Type	Disk Count (By Type)	RG Size (HDD/SSD)	HA Policy	Has Mirror	Mirrored	Size Nominal
aggr0_A300_01	AFF-300-01	Online	Normal	raid	11@3.49TB_SSD (shared)	24	CO	True	False	414.56GB
aggr0_A300_02	AFF-300-02	Online	Normal	raid	10@3.49TB_SSD (shared)	24	CO	True	False	368.4GB
aggr1_AFF300_01	AFF-300-01	Online	Normal	raid	23@3.49TB_SSD (Shared)	24	SFO	False	False	32.51TB
aggr1_AFF300_02	AFF-300-02	Online	Normal	raid	23@3.49TB_SSD (shared)	24	SFO	False	False	32.51TB

Storage virtual machines configuration

An SVM is a secure virtual storage server that contains data volumes and one or more LIFs through which it serves data to the clients. An SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by its SVM administrator.

In a cluster, an SVM facilitates data access. A cluster must have at least one SVM to serve data. SVMs use the storage and network resources of the cluster. However, the volumes and LIFs are exclusive to the SVM. Multiple SVMs can coexist in a single cluster without being bound to any node in a cluster. However, they are bound to the physical cluster on which they exist.

Table 27 lists the SVM configuration.

Table 27. SVM configuration

Cluster Name	SVM Name	Type	Sub-type	State	Allowed Protocols	Comment
AFF-A300	Infra	Data	Default	Running	NFS, CIFS,FCP	

Table 28 lists the SVM storage configuration.

Table 28. SVM storage configuration

Cluster Name	SVM Name	Root Volume Security Style	Language	Root Volume	Root Aggregate	Aggregate List
AFF-A300	Infra	UNIX	c.utf_8	svm_root	aggr1_AFF300_01	aggr1_AFF300_01, aggr1_AFF300_02

Volume configuration

A FlexVol volume is a data container associated with a SVM with FlexVol volumes. It gets its storage from a single associated aggregate, which it might share with other FlexVol volumes or infinite volumes. It can be used to contain files in a NAS environment, or LUNs in a SAN environment.

Table 29 lists the FlexVol configuration.

Table 29. FlexVol configuration

Cluster Name	SV M Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
AFF-A300	In- fra	esxi_bo ot	aggr1_AFF300_0 1	R W	Default	Default	UNIX	500.00GB
AFF-A300	In- fra	vdi_cifs	aggr1_AFF300_0 1 aggr1_AFF300_0 2	R W	Default	Default	UNIX	1.00TB
AFF-A300	In- fra	in- fra _nf s_ ds 01	aggr1_AFF300_0 1	R W	Default	Default	UNIX	6.00TB
AFF-A300	In- fra	vdi_nfs _d s0 1	aggr1_AFF300_0 1	R W	Default	Default	UNIX	15.00TB
AFF-A300	In- fra	vdi_nfs _d s0 2	aggr1_AFF300_0 2	R W	Default	Default	UNIX	10.00TB
AFF-A300	In- fra	vdi_nfs _d s0 3	aggr1_AFF300_0 1	R W	Default	Default	UNIX	10.00TB
AFF-A300	In- fra	vdi_nfs _d s0 4	aggr1_AFF300_0 2	R W	Default	Default	UNIX	10.00TB
AFF-A300	In- fra	vdi_nfs _d s0	aggr1_AFF300_0 1	R W	Default	Default	UNIX	10.00TB

Cluster Name	SVM Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
	ra	5						
AFF-A300	Infra	vdi_nfs_d s0 6	aggr1_AFF300_0 2	RW	Default	Default	UNIX	10.00TB
AFF-A300	Infra	vdi_nfs_d s0 7	aggr1_AFF300_0 1	RW	Default	Default	UNIX	10.00TB
AFF-A300	Infra	vdi_nfs_d s0 8	aggr1_AFF300_0 2	RW	Default	Default	UNIX	10.00TB
AFF-A300	Infra	svm_root	aggr1_AFF300_0 1	RW	Default	Default	UNIX	1GB

Protocol configuration

NAS

ONTAP can be accessed over CIFS, SMB and NFS capable clients. This means that clients can access all files on an SVM regardless of the protocol they are connecting with or the type of authentication they require.

Logical interfaces

A LIF is an IP address with associated characteristics, such as a role, a home port, a home node, a routing group, a list of ports to fail over to, and a firewall policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network.

LIFs can be hosted on the following ports:

- Physical ports that are not part of ifgrps
- Ifgrps
- VLANs
- Physical ports or ifgrps that host VLANs

While configuring SAN protocols such as FC on a LIF, a LIF role determines the kind of traffic that is supported over the LIF, along with the failover rules that apply and the firewall restrictions that are in place.

LIF failover refers to the automatic migration of a LIF in response to a link failure on the LIF's current network port. When such a port failure is detected, the LIF is migrated to a working port.

A failover group contains a set of network ports (physical, VLANs, and ifgrps) on one or more nodes. A LIF can subscribe to a failover group. The network ports that are present in the failover group define the failover targets for the LIF.

NAS logical interface settings

Table 30 lists the NAS LIF settings.

Table 30. NAS LIF settings

Cluster Name	SVM Name	Interface Name	Status (Admin/Oper)	IP Address	Current Node	Current Port	Is Home
AFF-A300	Infra	CIFS1-01	Up/up	10.10.6 2.1 0/2 4	AFF- A 3 0 0- 0 1	a0a-62	True
AFF-A300	Infra	CIFS2-02	Up/up	10.10.6 2.1 1/2 4	AFF- A 3 0 0- 0 2	a0a-62	True
AFF-A300	Infra	Mgmt2	Up/up	10.10.6 1.2 6/2 4	AFF- A 3 0 0- 0 1	a0a-61	True
AFF-A300	Infra	NFS1-01	Up/up	10.10.6 3.1 0/2	AFF- A 3 0	a0a-63	True

Cluster Name	SVM Name	Interface Name	Status (Admin/Oper)	IP Address	Current Node	Current Port	Is Home
				4	0-0-1		
AFF-A300	Infra	NFS2-02	Up/up	10.10.6 3.1 1/2 4	AFF- A 3 0 0- 0 2	a0a-63	True

Windows file services

You can enable and configure a CIFS SVM to let SMB clients access files on your SVM. Each data SVM in the cluster can be bound to only one Active Directory domain; however, the data SVMs do not need to be bound to the same domain. Each SVM can be bound to a unique Active Directory domain. Additionally, a CIFS SVM can be used to tunnel cluster administration authentication, which can be bound to only one Active Directory domain.

CIFS servers

CIFS clients can access files on an SVM by using the CIFS protocol, provided ONTAP can properly authenticate the user.

Table 31 lists the CIFS server configuration information.

Table 31. CIFS servers

Cluster Name	SVM Name	CIFS Server	Domain	Domain Net-BIOS Name	WINS Servers	Preferred DC
AFF-A300	Infra	Infra	VDILAB.LOCAL	VDILAB	-	-

CIFS options

Most of these options are only available starting with Data ONTAP 8.2.

Table 32 lists the CIFS options.

Table 32. CIFS options

Cluster Name	SVM Name	SMB v2 Enabled	SMB v3 Enabled	Export Policy Enabled	Copy Offload Enabled	Local Users and Groups Enabled	Referral Enabled	Shadow Copy Enabled
AFF-A300	Infra	True	True	False	True	True	False	True

CIFS local users and groups

You can create local users and groups on the SVM. The CIFS server can use local users for CIFS authentication and can use both local users and groups for authorization when determining both share and file and directory access rights.

Local group members can be local users, domain users and groups, and domain machine accounts.

Local users and groups can also be assigned privileges. Privileges control access to SVM resources and can override the permissions that are set on objects. A user or member of a group that is assigned a privilege is granted the specific rights that the privilege allows.

Note

Privileges do not provide ONTAP general administrative capabilities.

CIFS shares

A CIFS share is a named access point in a volume and/or namespace that enables CIFS clients to view, browse, and manipulate files on an SVM.

Table 33 lists the CIFS shares.

Table 33. CIFS shares

Cluster Name	SVM Name	Share Name	Path	Share Properties	Symlink Properties	Share ACL
AFF-A300	Infra	%w	%w	homedirecto- ry	symlinks	Everyone:Full Control

Cluster Name	SVM Name	Share Name	Path	Share Properties	Symlink Properties	Share ACL
AFF-A300	Infra	admin\$	/	browsable	-	UTD
AFF-A300	Infra	c\$	/	oplocks browsable changenotify show_previous_versions	symlinks	Administrators:Full Control
AFF-A300	Infra	ipc\$	/	browsable	-	UTD
AFF-A300	Infra	Profile\$	/vdi_cifs/Profiles	oplocks browsable changenotify show_previous_versions	symlinks	Everyone:Full Control

CIFS home directory search paths

ONTAP home directories enable you to configure an SMB share that maps to different directories based on the user that connects to it and a set of variables. Instead of having to create separate shares for each user, you can configure a single share with a few home directory parameters to define a user's relationship between an entry point (the share) and their home directory (a directory on the SVM).

The home directory search paths are a set of absolute paths from the root of the SVM that you specify that directs the ONTAP search for home directories. You specify one or more search paths by using the `vserver cifs home-directory search-path add` command. If you specify multiple search paths, ONTAP tries them in the order specified until it finds a valid path.

SAN

SAN is a term used to describe a purpose-built storage controller that provides block-based data access. ONTAP supports traditional FC (as well as iSCSI and FCoE) within a unified architecture.

LUNs

LUNs are created and exist within a given FlexVol volume and are used to store data, which is presented to servers or clients. LUNs provide storage for block-based protocols such as FC or iSCSI.

Table 34 lists the LUN details.

Table 34. LUN configuration

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/esxi_boot/infra_host_01	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/infra_host_02	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-1	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-2	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-3	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-4	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-5	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-6	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-7	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-9	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-10	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-11	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-12	True	True	VMware	False	10.00GB

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/esxi_boot/VDI-13	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-14	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-15	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-17	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-18	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-19	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-20	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-21	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-22	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-23	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-24	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-25	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-26	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-27	True	True	VMware	False	10.00GB

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/esxi_boot/VDI-28	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-29	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-30	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-31	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-32	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/vGPU-Boot01	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/vGPU-Boot02	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/vGPU-Boot03	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/vGPU-Boot04	True	True	VMware	False	10.00GB
AFF-A300	Infra	/vol/vdi_fc_ds01/vdi_fc_ds01	True	True	VMware	False	10.00TB
AFF-A300	Infra	/vol/vdi_fc_ds02/vdi_fc_ds02	True	True	VMware	False	10.00TB
AFF-A300	Infra	/vol/vdi_fc_ds03/vdi_fc_ds03	True	True	VMware	False	10.00TB
AFF-A300	Infra	/vol/vdi_fc_ds04/vdi_fc_ds04	True	True	VMware	False	10.00TB
AFF-A300	Infra	/vol/vdi_fc_ds05/vdi_fc_ds05	True	True	VMware	False	10.00TB

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/vdi_fc_ds06/vdi_fc_ds06	True	True	VMware	False	10.00TB

Initiator groups

Initiator groups (igroups) are tables of FC protocol host WWPNs or iSCSI host node names. You can define igroups and map them to LUNs to control which initiators have access to LUNs.

Typically, you want all of the host's initiator ports or software initiators to have access to a LUN. If you are using multipathing software or have clustered hosts, each initiator port or software initiator of each clustered host needs redundant paths to the same LUN.

You can create igroups that specify which initiators have access to the LUNs either before or after you create LUNs, but you must create igroups before you can map a LUN to an igroup.

Igroups can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator. An initiator cannot be a member of igroups of differing OS types.

Table 35 lists the igroups that have been created.

Table 35. Igroups

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF-A300	Infra	SP_Infra1	FCP	VMware	True	Full
AFF-A300	Infra	SP_Infra2	FCP	VMware	True	Full
AFF-A300	Infra	VDI-1	FCP	VMware	True	Full
AFF-A300	Infra	VDI-2	FCP	VMware	True	Full
AFF-A300	Infra	VDI-3	FCP	VMware	True	Full
AFF-A300	Infra	VDI-4	FCP	VMware	True	Full
AFF-A300	Infra	VDI-5	FCP	VMware	True	Full
AFF-A300	Infra	VDI-6	FCP	VMware	True	Full

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF-A300	Infra	VDI-7	FCP	VMware	True	Full
AFF-A300	Infra	VDI-9	FCP	VMware	True	Full
AFF-A300	Infra	VDI-10	FCP	VMware	True	Full
AFF-A300	Infra	VDI-11	FCP	VMware	True	Full
AFF-A300	Infra	VDI-12	FCP	VMware	True	Full
AFF-A300	Infra	VDI-13	FCP	VMware	True	Full
AFF-A300	Infra	VDI-14	FCP	VMware	True	Full
AFF-A300	Infra	VDI-15	FCP	VMware	True	Full
AFF-A300	Infra	VDI-17	FCP	VMware	True	Full
AFF-A300	Infra	VDI-18	FCP	VMware	True	Full
AFF-A300	Infra	VDI-19	FCP	VMware	True	Full
AFF-A300	Infra	VDI-20	FCP	VMware	True	Full
AFF-A300	Infra	VDI-21	FCP	VMware	True	Full
AFF-A300	Infra	VDI-22	FCP	VMware	True	Full
AFF-A300	Infra	VDI-23	FCP	VMware	True	Full
AFF-A300	Infra	VDI-24	FCP	VMware	True	Full
AFF-A300	Infra	VDI-25	FCP	VMware	True	Full
AFF-A300	Infra	VDI-26	FCP	VMware	True	Full
AFF-A300	Infra	VDI-27	FCP	VMware	True	Full
AFF-A300	Infra	VDI-28	FCP	VMware	True	Full

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF-A300	Infra	VDI-29	FCP	VMware	True	Full
AFF-A300	Infra	VDI-30	FCP	VMware	True	Full
AFF-A300	Infra	VDI-31	FCP	VMware	True	Full
AFF-A300	Infra	VDI-32	FCP	VMware	True	Full
AFF-A300	Infra	VDI_cluster	FCP	VMware	True	Full
AFF-A300	Infra	vGPU01	FCP	VMware	True	Full
AFF-A300	Infra	vGPU02	FCP	VMware	True	Full
AFF-A300	Infra	vGPU03	FCP	VMware	True	Full
AFF-A300	Infra	vGPU04	FCP	VMware	True	Full

FCP logical interface settings

Table 36 lists the FCP LIF settings.

Table 36. FCP LIF settings

Cluster Name	SVM Name	Interface Name	Status (Admin/Operator)	Port Name	Current Node	Current Port	Is Home
AFF-A300	Infra	fcp_01a	Up/up	20:01:00:a0:98:af:bd:e8	AFF-A300-01	0g	True
AFF-A300	Infra	fcp_01b	Up/up	20:02:00:a0:98:af:bd:e8	AFF-A300-01	0h	True
AFF-A3	Infra	fcp_02a	Up/up	20:03:00:a0:98:af:bd:e8	AFF-A3	0g	True

Cluster Name	SVM Name	Interface Name	Status (Admin/Operator)	Port Name	Current Node	Current Port	Is Home
00					00-02		
AFF-A300	Infra	fcp_02b	Up/up	20:04:00:a0:98:af:bd:e8	AFF-A300-02	0h	True

FCP / FCoE

FCP service configuration

FCP is a licensed service on the storage system that enables you to export LUNs and transfer block data to hosts using the SCSI protocol over an FC fabric.

Table 37 lists the FCP service configuration details.

Table 37. FCP service configuration

Cluster Name	SVM Name	Node Name	Available
AFF-A300	Infra	20:00:00:a0:98:af:bd:e8	True

FCP adapter configuration

You can use storage controller onboard FC ports as both initiators and targets. You can also add storage controller FC ports on expansion adapters and use them as initiators or targets, depending on the type of expansion adapter installed.

Table 38 lists the details of the storage controller target ports and the WWPN address of each.

Table 38. FCP adapter configuration

Node Name	Adapter Name	State	Data Link Rate	Media Type	Speed	Port Name
AFF-A300-01	0e	Offlined by user/system	0	PTP	Auto	50:0a:09:82:80:13:41:27

Node Name	Adapter Name	State	Data Link Rate	Media Type	Speed	Port Name
AFF-A300-01	0f	Offlined by user/system	0	PTP	Auto	50:0a:09:81:80:13:41:27
AFF-A300-01	0g	Online	8	PTP	Auto	50:0a:09:84:80:13:41:27
AFF-A300-01	0h	Online	8	PTP	Auto	50:0a:09:83:80:13:41:27
AFF-A300-02	0e	Offlined by user/system	0	PTP	Auto	50:0a:09:82:80:d3:67:d3
AFF-A300-02	0f	Offlined by user/system	0	PTP	Auto	50:0a:09:81:80:d3:67:d3
AFF-A300-02	0g	Online	8	PTP	Auto	50:0a:09:84:80:d3:67:d3
AFF-A300-02	0h	Online	8	PTP	Auto	50:0a:09:83:80:d3:67:d3

Storage efficiency and space management

ONTAP offers a wide range of storage-efficiency technologies in addition to Snapshot technology. Key technologies include thin provisioning, deduplication, compression, and FlexClone volumes, files, and LUNs. As with Snapshot technology, all are built on ONTAP WAFL.

Volume efficiency

You can run deduplication, data compression, and data compaction together or independently on a FlexVol volume or an infinite volume to achieve optimal space savings. Deduplication eliminates duplicate data blocks and data compression compresses the data blocks to reduce the amount of physical storage that is required. Data compaction stores more data in less space to increase storage efficiency.

Beginning with ONTAP 9.2, all inline storage-efficiency features, such as inline deduplication and inline compression, are enabled by default on AFF volumes.

Table 39 lists the volume efficiency settings.

Table 39. Volume efficiency settings

Cluster Name	SVM Name	Volume Name	Space Guarantee	Dedupe	Schedule or Policy Name	Compression	Inline Compression
AFF-A300	Infra	esxi_boot	None	True	Inline-only	True	True
AFF-A300	Infra	esxi_boot	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_cifs	None	True	Inline-only	True	True
AFF-A300	Infra	infra_nfs_ds01	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds01	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds02	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds03	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds04	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds05	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds06	None	True	Inline-only	True	True

Cluster Name	SVM Name	Volume Name	Space Guarantee	Dedupe	Schedule or Policy Name	Compression	Inline Compression
AFF-A300	Infra	vdi_nfs_ds07	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds08	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds01	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds02	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds03	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds04	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds05	None	True	Inline-only	True	True
AFF-A300	Infra	vdi_fc_ds06	None	True	Inline-only	True	True

LUN efficiency

Thin provisioning enables storage administrators to provision more storage on a LUN than is physically present on the volume. By overprovisioning the volume, storage administrators can increase the capacity utilization of that volume. As the blocks are written to the LUN, ONTAP adds more space to the LUN from available space on the volume.

With thin provisioning, you can present more storage space to the hosts connecting to the SVM than what is actually available on the SVM. Storage provisioning with thinly provisioned LUNs enables storage administrators to

pro-
vide actual storage that the LUN needs. As ONTAP writes blocks to the LUN, the LUN increases in size automatically.

Table 40 lists the LUN efficiency settings.

Table 40. LUN efficiency settings

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF-A300	Infra	/vol/esxi_boot/infra_host_01	False	False
AFF-A300	Infra	/vol/esxi_boot/infra_host_02	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-1	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-2	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-3	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-4	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-5	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-6	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-7	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-9	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-10	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-11	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-12	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-13	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-14	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-15	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-17	False	False

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF-A300	Infra	/vol/esxi_boot/VDI-18	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-19	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-20	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-21	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-22	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-23	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-24	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-25	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-26	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-27	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-28	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-29	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-25	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-30	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-31	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-32	False	False

Install and configure VMware ESXi 6.7

This section explains how to install VMware ESXi 6.7 Update 2 in an environment.

There are several methods to install 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 install ESXi on boot logical unit number (LUN). Upon completion of steps outlined here, ESXi hosts will be booted from their corresponding SAN Boot LUNs.

Procedure 1. Download Cisco Custom image for ESXi 6.7 Update 2

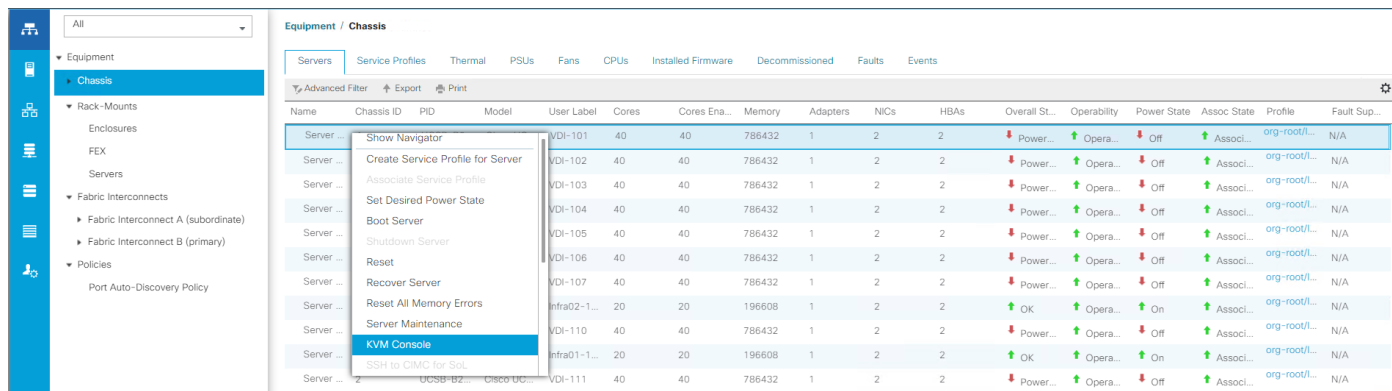
Step 1. From the [VMware vSphere Hypervisor 6.7 U2](#) page click the **Custom ISOs** tab and download the image.

Procedure 2. Install VMware vSphere ESXi 6.7

Step 1. In the Cisco UCS Manager navigation pane, click the **Equipment** tab.

Step 2. Go to **Equipment > Chassis > Chassis 1 > Server 1**.

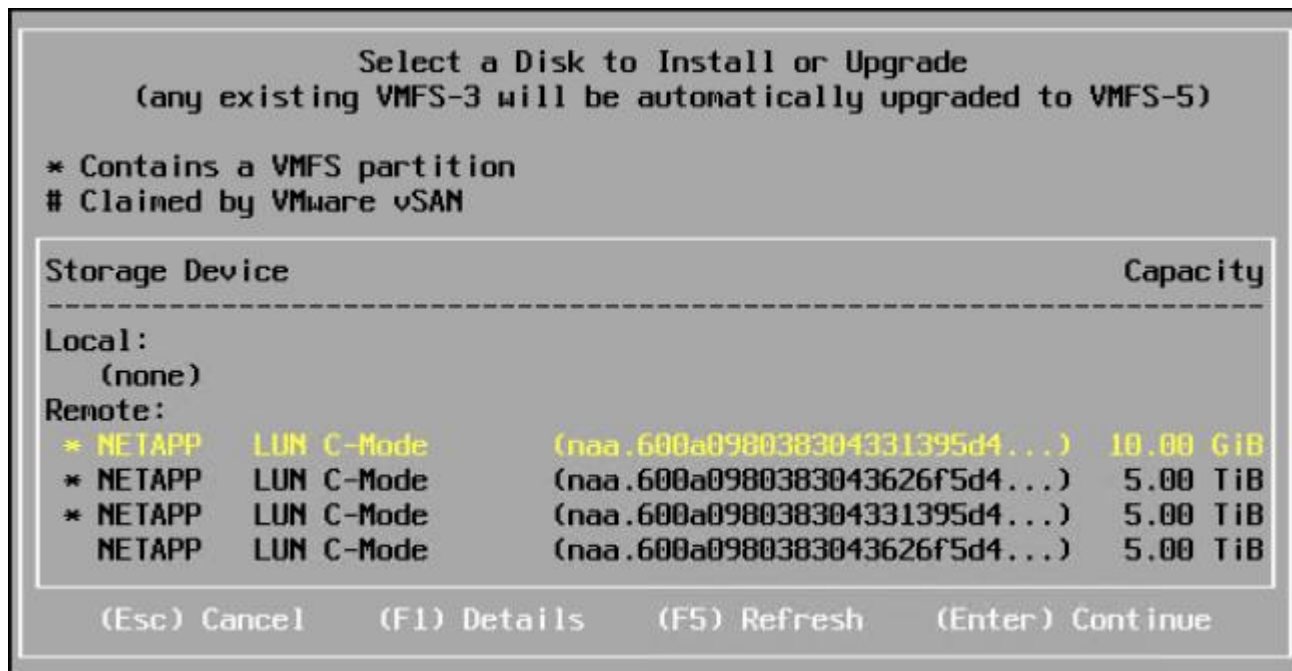
Step 3. Right-click **Server 1** and select **KVM Console**.



Step 4. Click **Activate Virtual Devices** and mount the ESXi ISO image.

Step 5. Follow the prompts to complete installing VMware vSphere ESXi hypervisor.

Step 6. When selecting a storage device to install ESXi, select the Remote LUN that was provisioned through the ONTAP System Manager and access it through the FC connection.



Procedure 3. Set up management networking for ESXi hosts

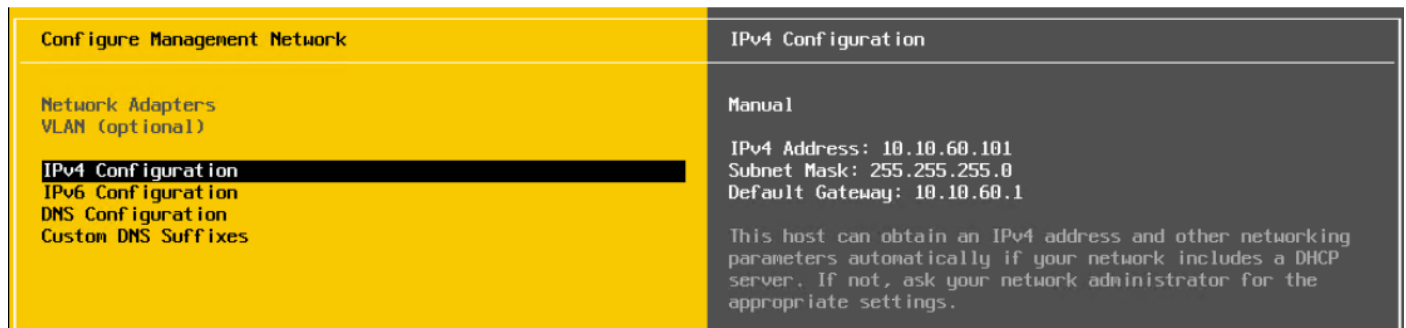
Adding a management network for each VMware host is necessary for managing the host and connection to vCenter Server. Please select the IP address that can communicate with existing or new vCenter Server.

- Step 1.** After the server has finished rebooting, press **F2** to enter the configuration wizard for ESXi Hypervisor.
- Step 2.** Log in as **root** and enter the corresponding password.
- Step 3.** Select **Configure the Management Network** option and press **Enter**.
- Step 4.** Select **VLAN (Optional)** and press **Enter**. Enter the VLAN In-Band management ID and press **Enter**.
- Step 5.** From the Configure Management Network menu, select **IPv4 Configuration** and press **Enter**.
- Step 6.** Select **Set Static IP Address and Network Configuration** 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 7.** IPv6 Configuration is set to **automatic**.
- Step 8.** Select **DNS Configuration** and press **Enter**.
- Step 9.** Enter the IP address of the primary and secondary DNS server. Enter the Hostname.
- Step 10.** Enter the DNS Suffixes.

Since the IP address is assigned manually, the DNS information must also be entered manually.

The steps provided varies based on the configuration. Please make the necessary changes according to your configuration.

Figure 44. Sample ESXi Configure Management Network



Procedure 4. Update Cisco VIC drivers for ESXi

When ESXi is installed from Cisco Custom ISO, you might have to update 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: VMW-ESX-6.7.0-nenic-1.0.29.0 and VMW-ESX-6.7.0-nfnic-4.0.0.40.

- Step 1.** Log into your VMware Account to download the 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 following VMware Clusters were configured to support the solution and testing environment:

- FlexPod-Datacenter: NetApp AFF A300 with Cisco UCS
 - Infrastructure Cluster: Infrastructure virtual machines (vCenter, Active Directory, DNS, DHCP, SQL Server, XenDesktop Controllers, Provisioning Servers, and other common services).
 - VDI: Virtual Desktop or RDS Server workload.
 - VDI1: Virtual Desktop or RDS Server workload.
 - VDI2: Virtual Desktop or RDS Server workload.

Login VSI Cluster: The Login VSI launcher infrastructure was connected using the same set of switches but hosted on separate SAN storage and servers.

Figure 45. VMware vSphere WebUI reporting cluster configuration for this CVD



Procedure 1. Install NetApp Virtual Storage Console (VSC) Plug-in

Install the NetApp Virtual Storage Console plug-in to simplify the system management.

- Step 1.** Download the **.ova file** from the [NetApp Support Site](#) to a vSphere Client system to deploy the virtual appliance for VSC.
- Step 2.** Log into the **vSphere Web Client**, select **Home > Host & Clusters**.
- Step 3.** Right-click the required data center and click **Deploy OVA template**.
- Step 4.** Browse to the folder where the .ova file is saved and click **Next**.
- Step 5.** Enter the required details to complete the deployment.

You can view the progress of the deployment from the Tasks tab, and wait for deployment to complete.

- Step 6.** Right-click the deployed virtual appliance for VSC and click **Install VMware tools**.
- Step 7.** Complete the VSC registration from the `<VSC_IPADDRESS:8143/Register.html>` page.

Figure 46. VCS registration

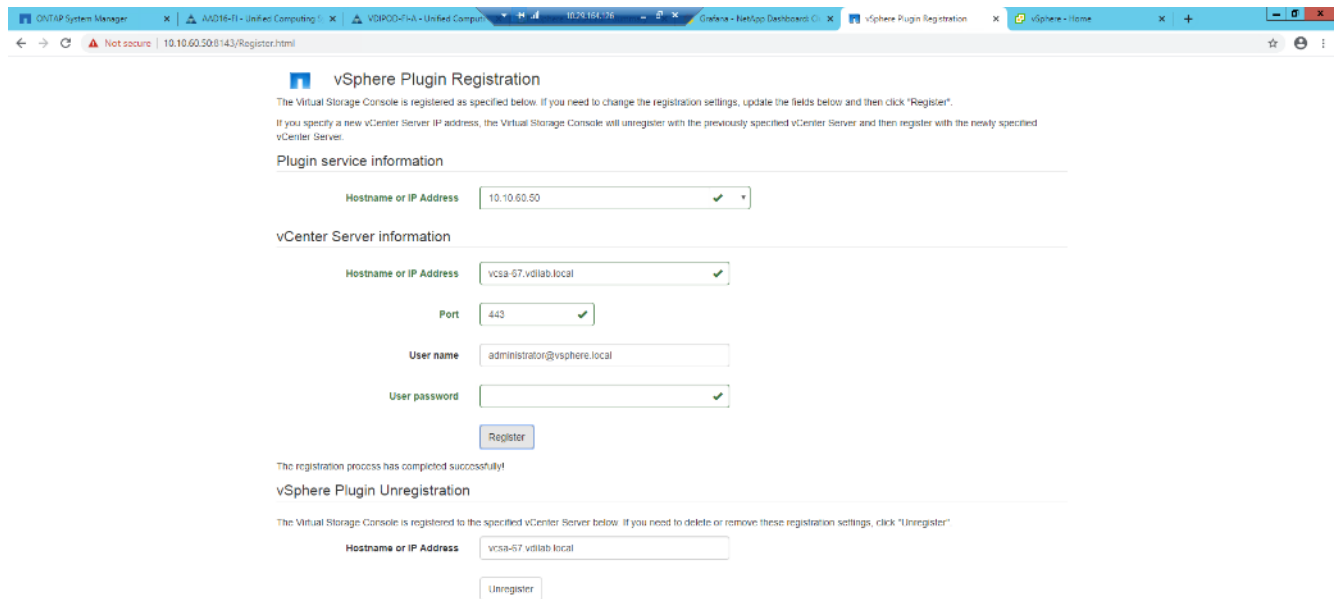
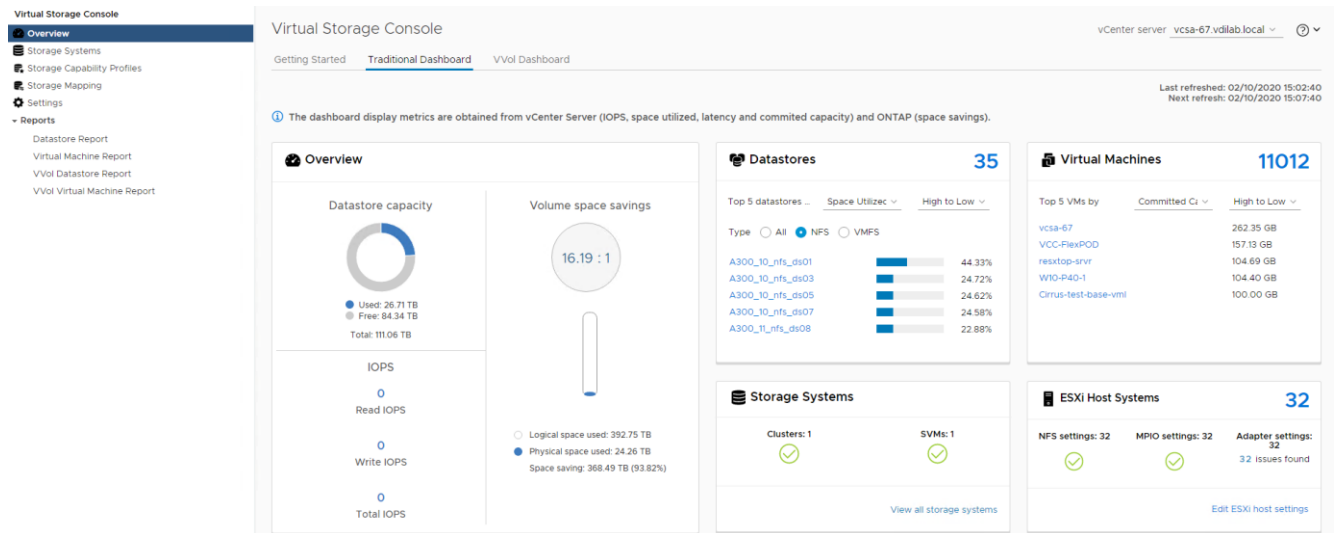


Figure 47. VSC dashboard in vCenter




Building the virtual machines and environment for workload testing

Prerequisites

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

Figure 48. Example of the DHCP scopes used in this CVD



Procedure 1. 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 provided in Table 41.

Table 41. Test infrastructure virtual machine configuration

Configuration	Vmware Connection Servers Virtual Machines	Vmware Composer Servers Virtual Machines
Operating system	Microsoft Windows Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	10	2
Memory amount	16 GB	16 GB
Network	VMXNET3	
Infra	VMXNET3	
Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2019	VCSA – SUSE Linux
Virtual CPU amount	2	24
Memory amount	8 GB	48 GB
Network	VMXNET3	VMXNET3

Configuration	Vmware Connection Servers Virtual Machines	Vmware Composer Servers Virtual Machines
	Infra	OOB-Mgmt
Disk size	40 GB	2 TB (across 13 VMDKs)
Configuration	Microsoft SQL Server Virtual Machine	NetApp VSC-9.6P1 Virtual Machine
Operating system	Microsoft Windows Server 2019 Microsoft SQL Server 2016 SP1	Linux Server (64-bit)
Virtual CPU amount	4	2
Memory amount	16GB	12 GB
Network	VMXNET3	VMXNET3 Infra
Disk-1 (OS) size	Infra	53 GB (across 4 VMDKs)
Disk-2 size	100 GB SQL Databases\Logs	-

Procedure 2. Preparing the master targets

This section provides guidance regarding creating the golden (or master) images for the environment. Virtual machines for the master targets must first be installed with the software components needed to build the golden images. Additionally, all available patches as of August 30, 2019 for the Microsoft operating systems, SQL server and Microsoft Office 2016 were installed.

To prepare the master virtual machines, there are three major steps: installing the Operating System and VMware tools, installing application software, and installing the VMware Horizon Agent.

For this CVD, the images contain the basics needed to run the Login VSI workload.

To configure the master target VDI and RDS virtual machines, see Table 42.

Table 42. VDI and RDS virtual machines configurations

Configuration	VDI Virtual Machines	RDS Virtual Machines
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2019


Configuration	VDI Virtual Machines	RDS Virtual Machines
Virtual CPU amount	2	10
Memory amount	3 GB reserve for all guest memory	32 GB reserve for all guest memory
Network	VMXNET3 VDI	VMXNET3 VDI
vDisk size	32 GB	40 GB
Additional software used for testing	Microsoft Office 2016 Login VSI 4.1.25 (Knowledge Worker Workload)	Microsoft Office 2016 Login VSI 4.1.25 (Knowledge Worker Workload)

RDS Server Roles need to be deployed on the RDS Master image.

Procedure 3. VMware Horizon Agent installation

Step 1. Download **VMware-viewagent-x86_64-7.10.0-14590940** version.

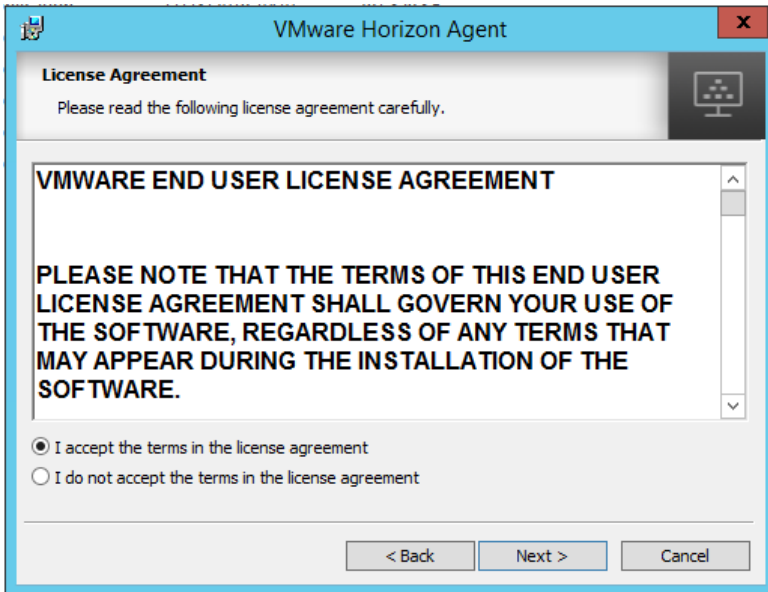
Step 2. Click the **VMware Horizon Agent installer**.



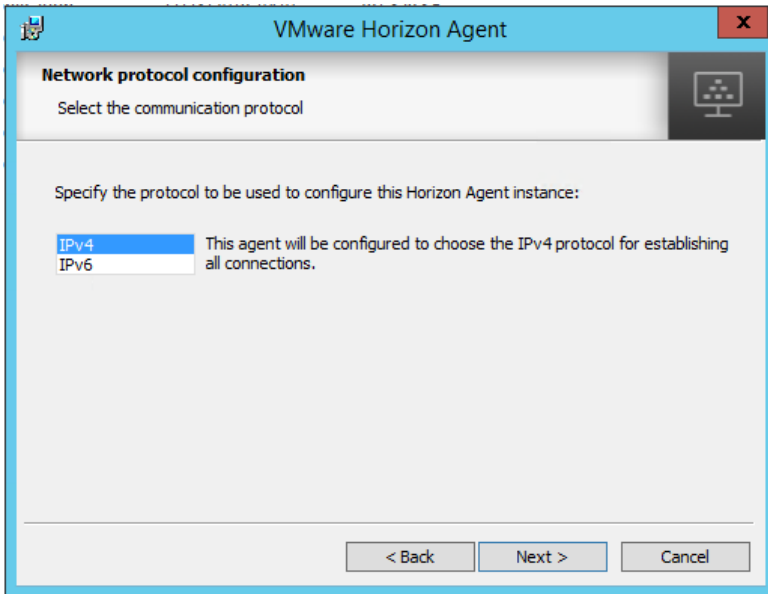
Step 3. Click **Next**.



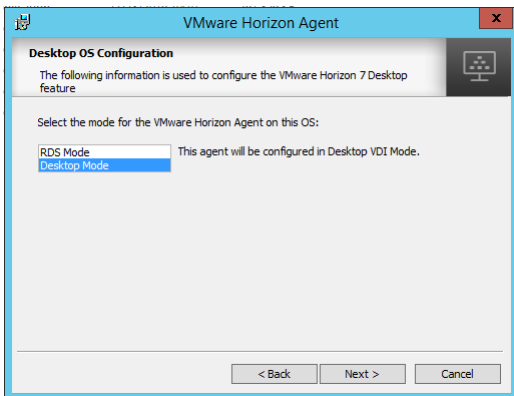
Step 4. Accept the License Agreement and click **Next**.



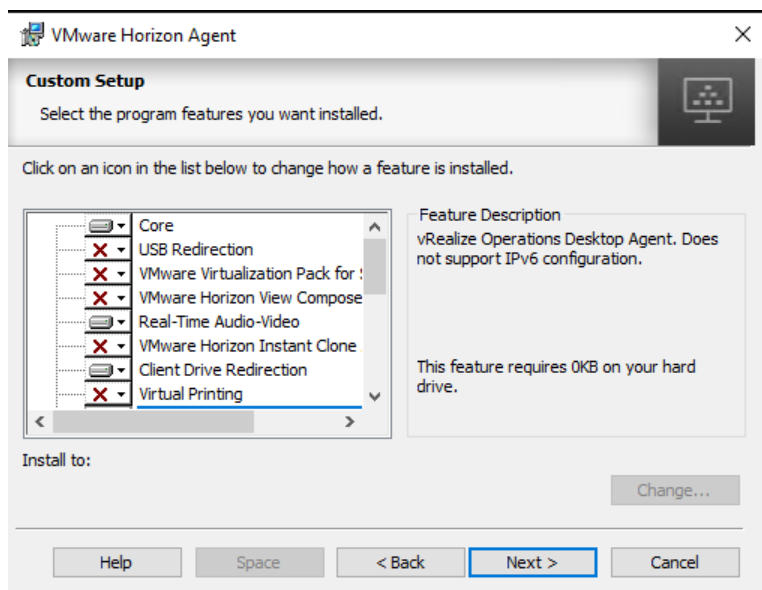
Step 5. Select the default **IPV4** and click **Next**.



Step 6. During the installation on the Windows 2019 Server select **RDS Mode** for the agent installation.

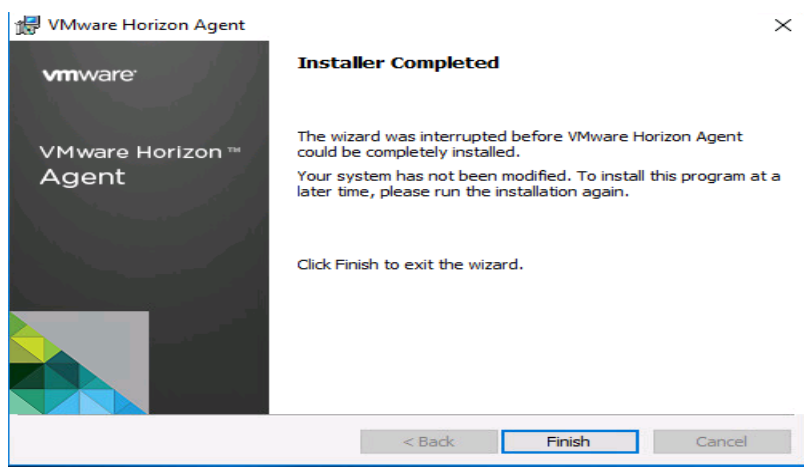


Step 7. Select the features to install.



Step 8. Click **Install**.

Step 9. Click **Finish** to complete the Horizon Agent installation on the Master image.




Install and configure VMware Horizon components

This section details the installation of the VMware core components of the Horizon Connection Server and Replica Servers. This CVD installs 1 VMware Horizon Replica server, 1 VMware Horizon Connection server and 3 VMware Horizon Replica Servers to support Remote Desktop Server Hosted Sessions (RDSH), non-persistent virtual desktops (VDI) instant clones, and persistent virtual desktops (VDI) full clones based on the best practices from VMware. For information about sizing limits, see [VMware Horizon View 7 sizing limits and recommendations](#).

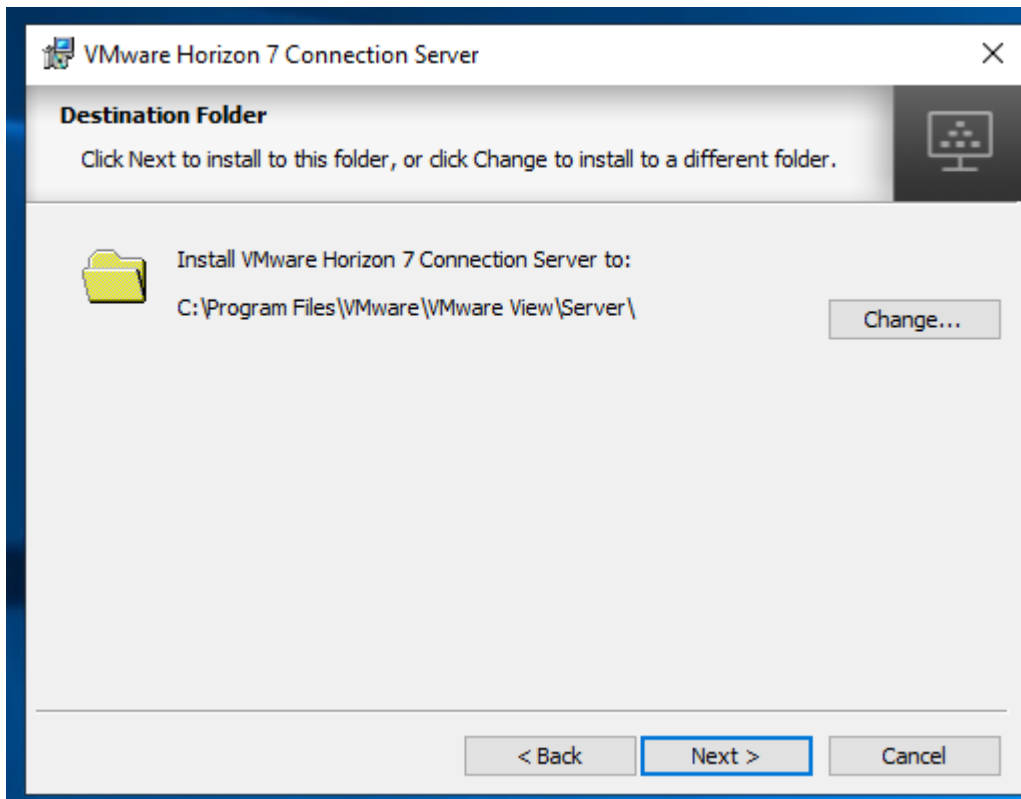
Procedure 1. VMware Horizon Connection Server configuration

Step 1. Download the **Horizon Connection Server Installer** from VMware and click **Install** on the Connection Server Windows Server Image. In this study, we used version Connection Server 7.10.0 build14590940. For the download, see [Download VMware Horizon 7.10.0](#).

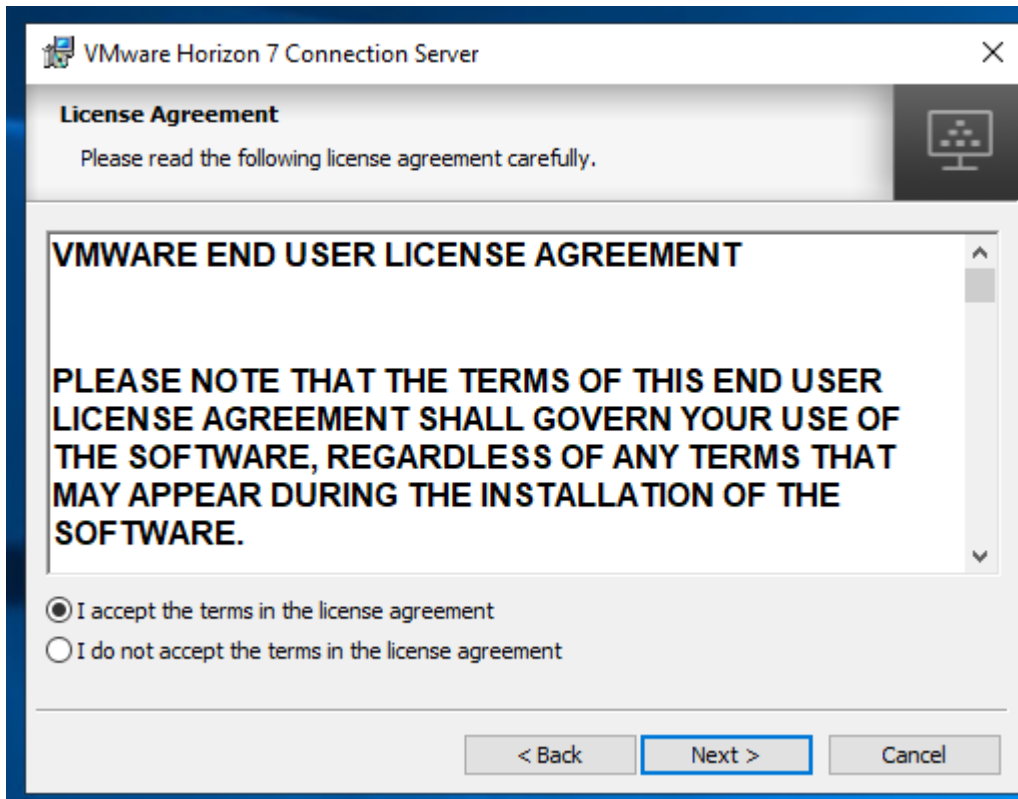
Step 2. Click the **Connection Server installer.**



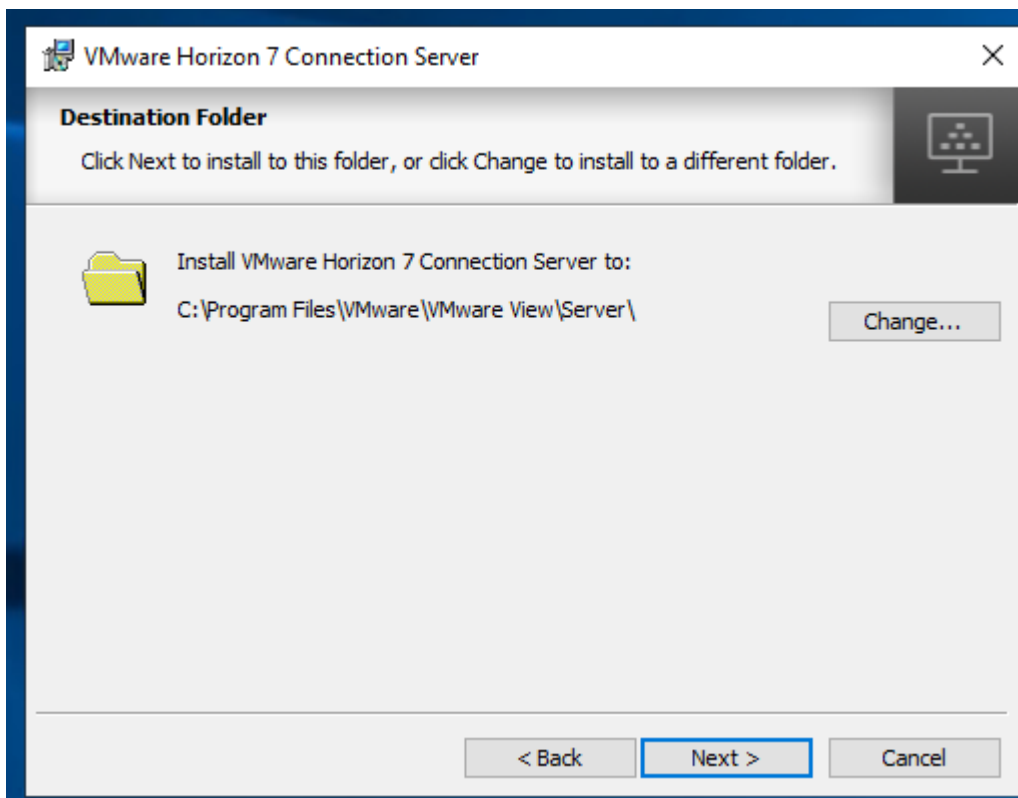
Step 3. Click **Next.**



Step 4. Accept the terms in the License Agreement. Click **Next.**

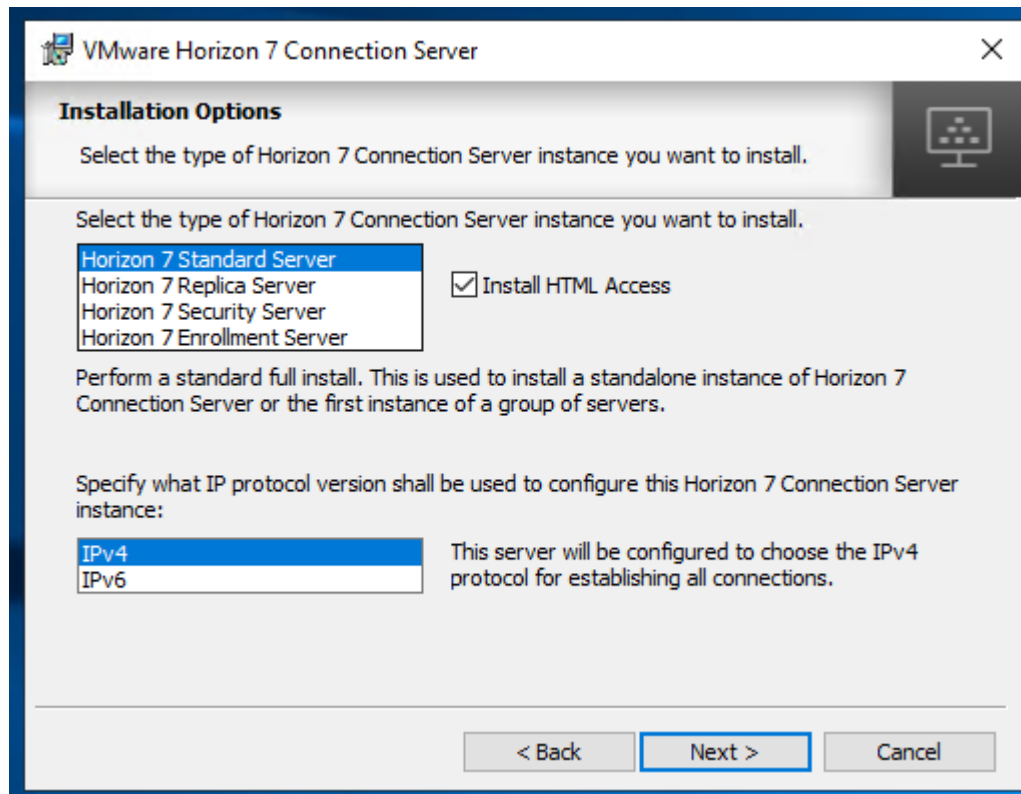


Step 5. Click Next.

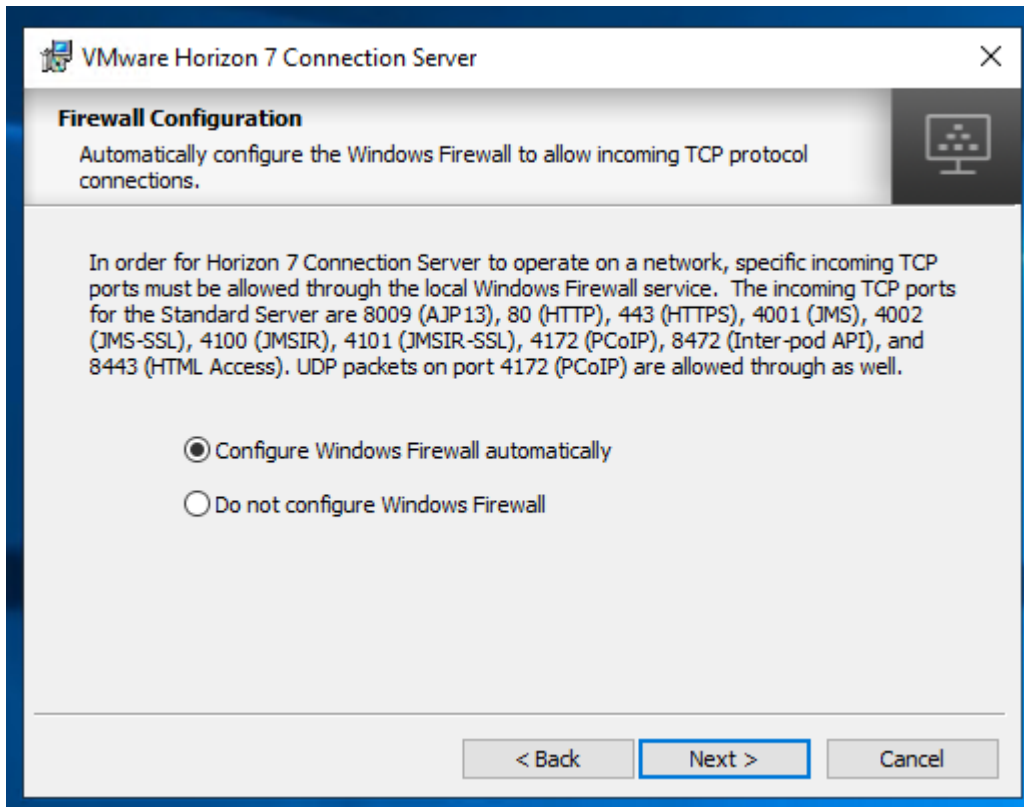


Step 6. Select the **Standard Server**. Click **Next**.

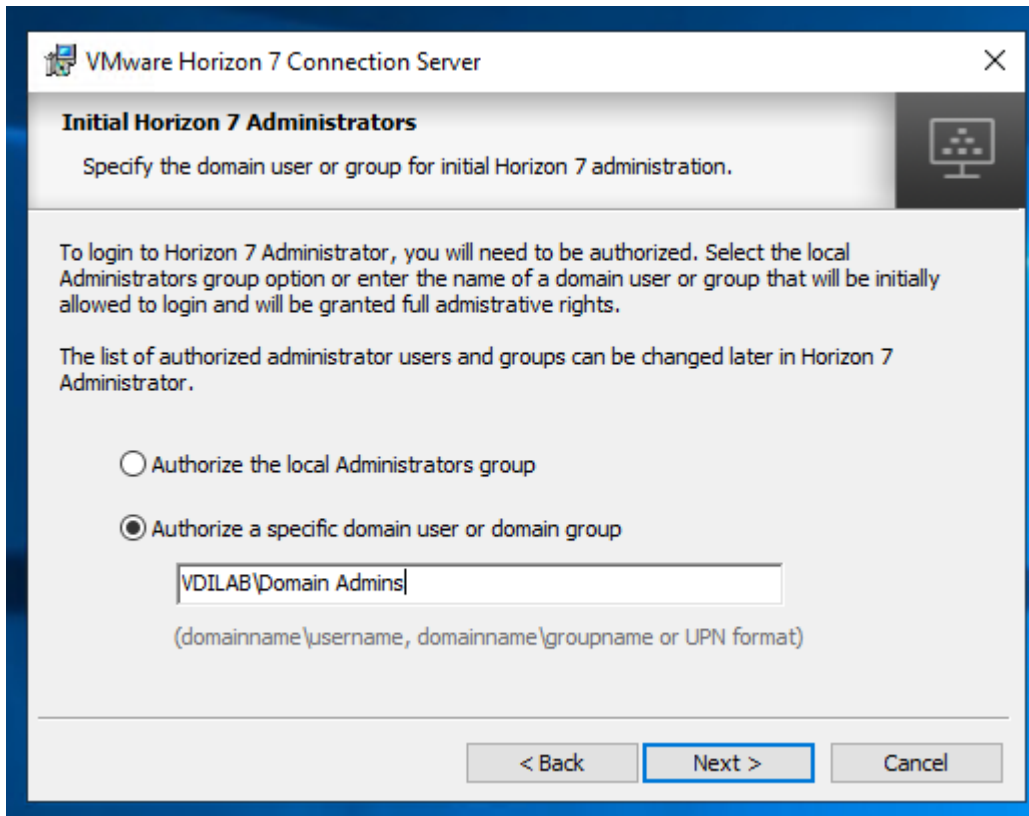
Step 7. To install additional VMware Horizon Replica servers, during the installation, select **Horizon 7 Connection Server** Option to sync the Replica Servers with the existing Standard server by providing View Connection Server's FQDN or IP address.



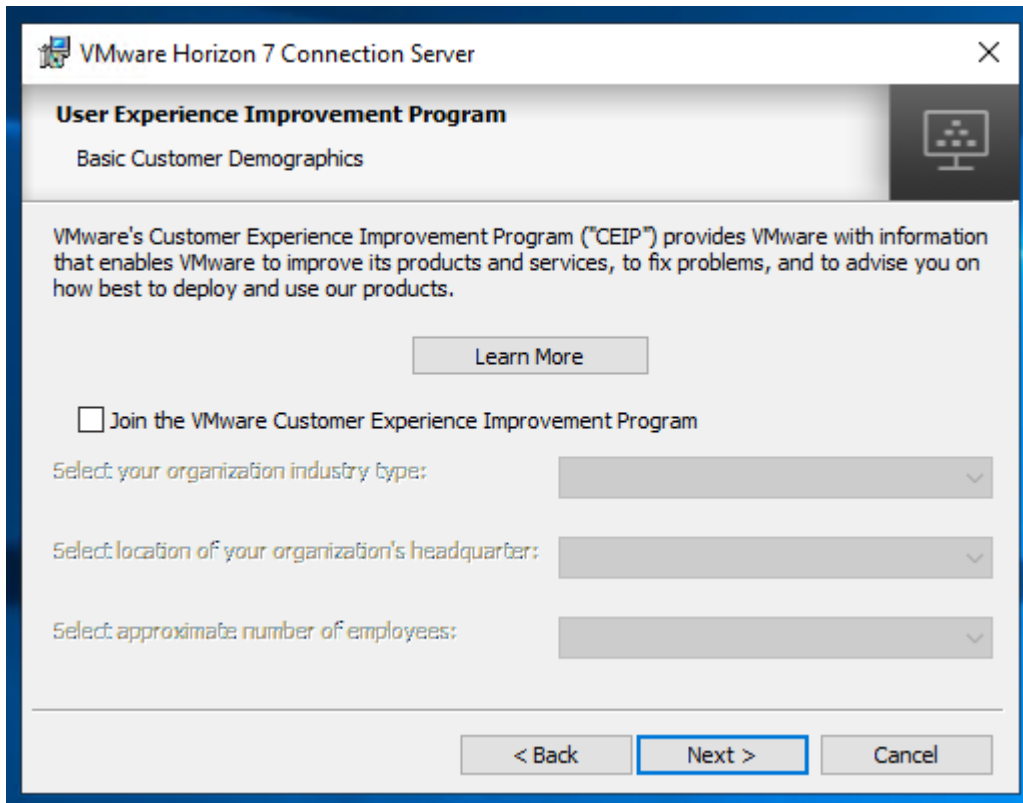
Step 8. Select **Configure Windows firewall automatically**. Click **Next**.



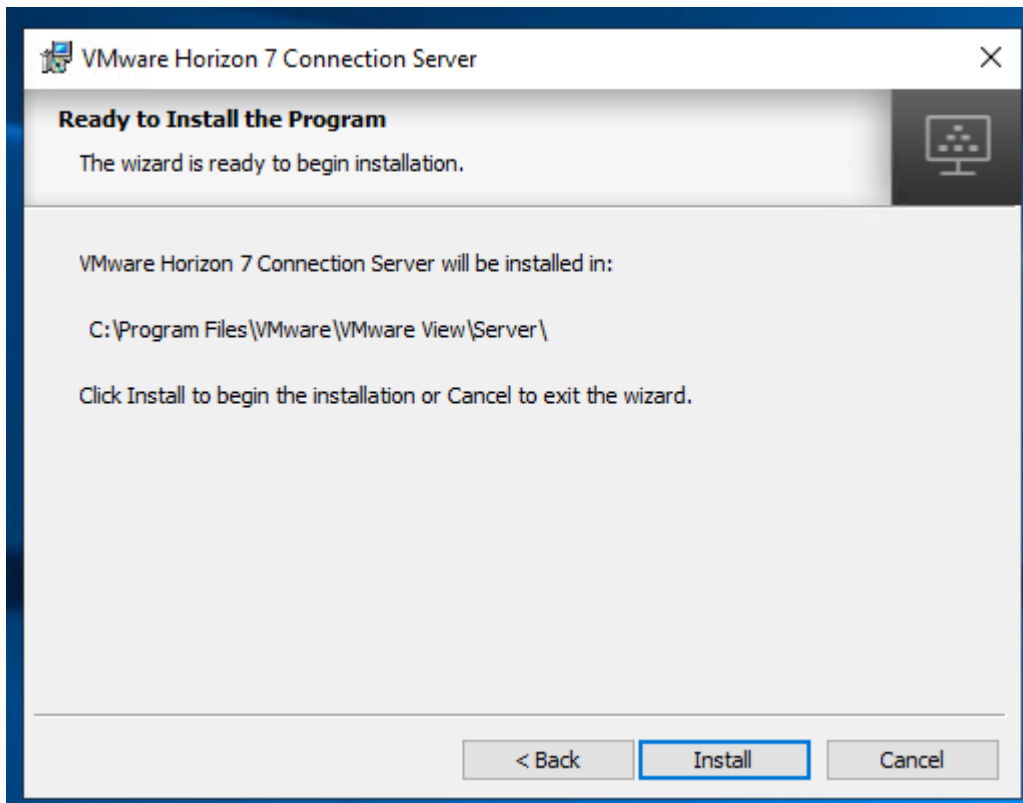
Step 9. Choose Horizon administrators. Click **Next**.

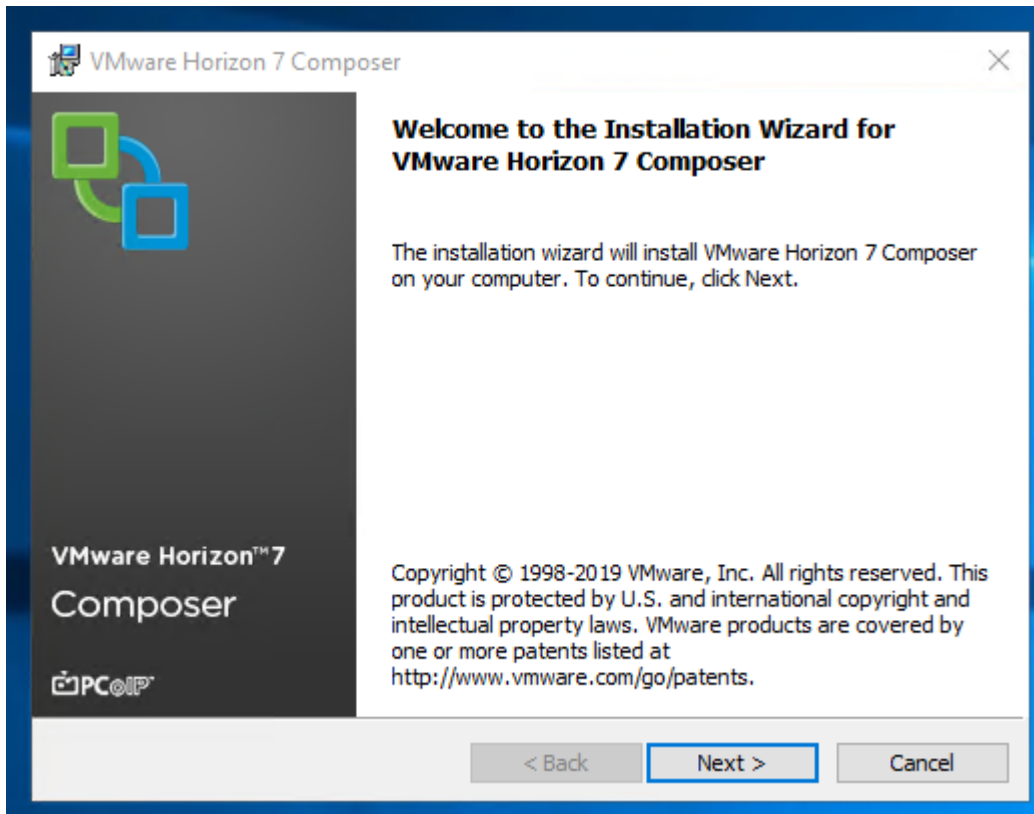


Step 10. (Optional) Select **Join Customer Experience Program**. Click **Next**.

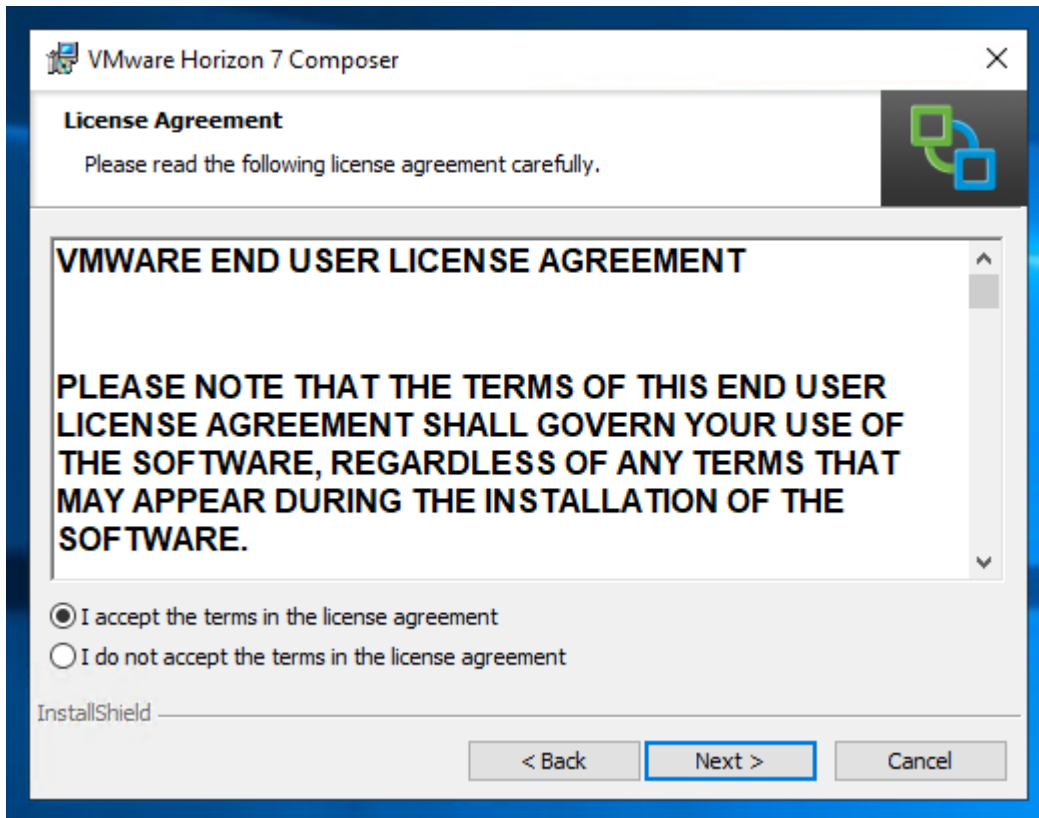


Step 11. Click **Install** to begin installation.

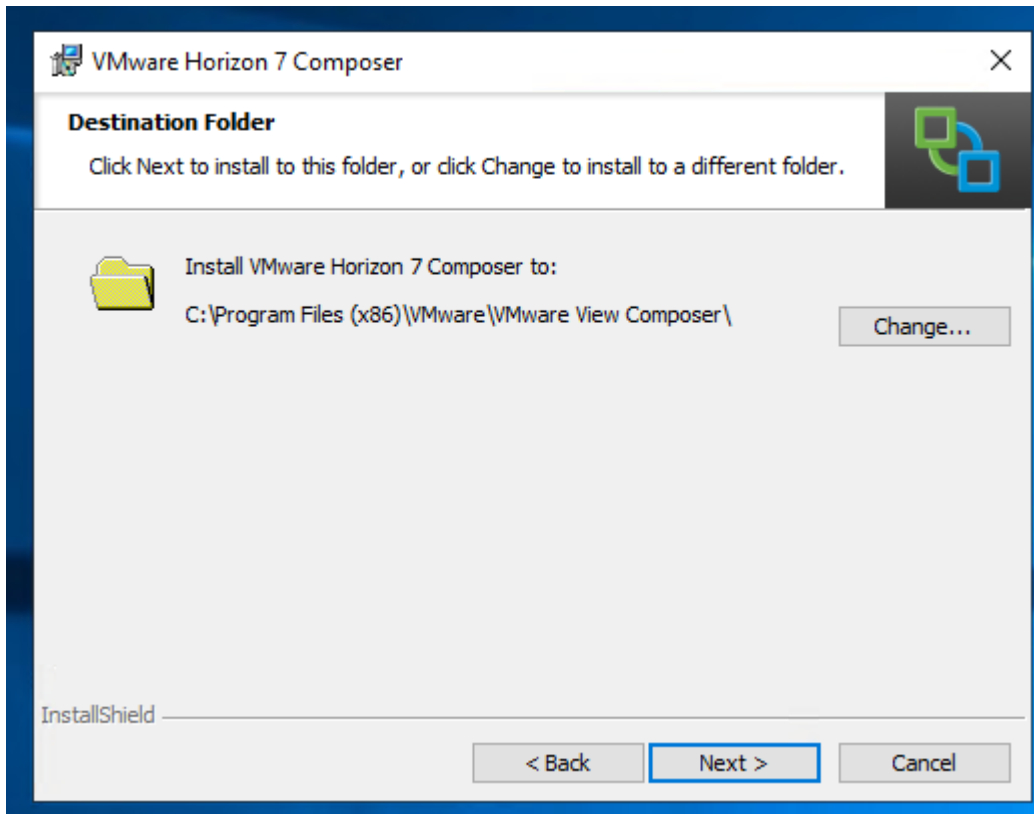




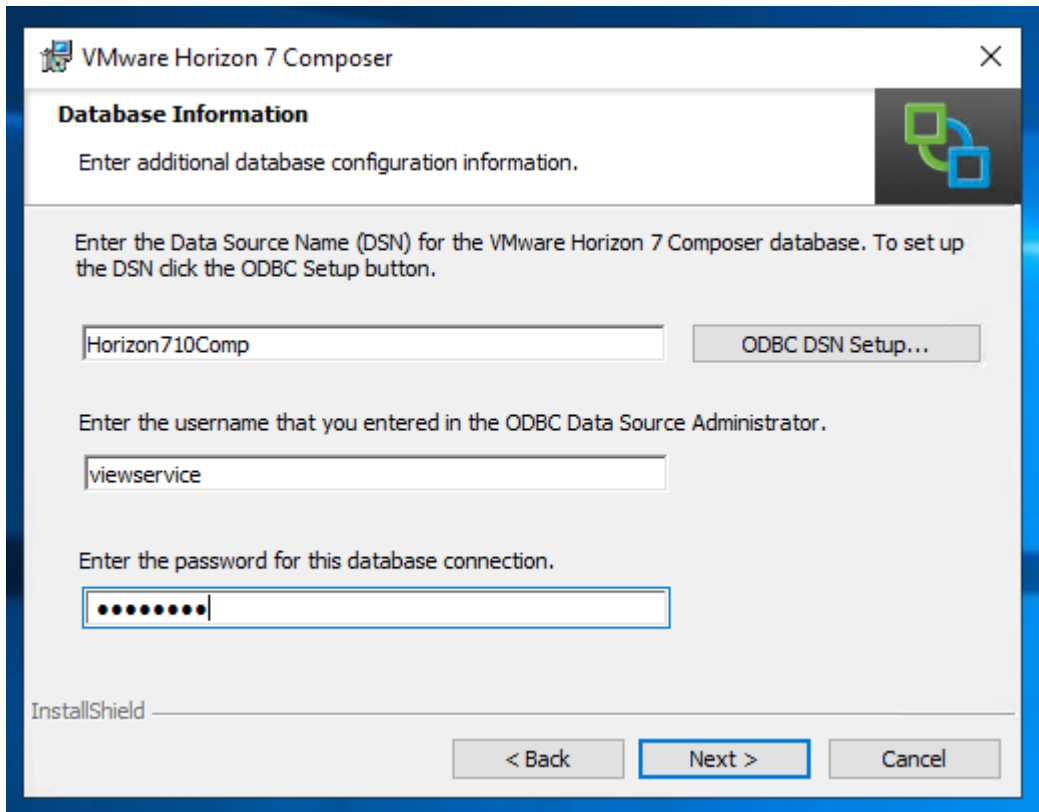
Step 3. Accept the License Agreement and click Next.



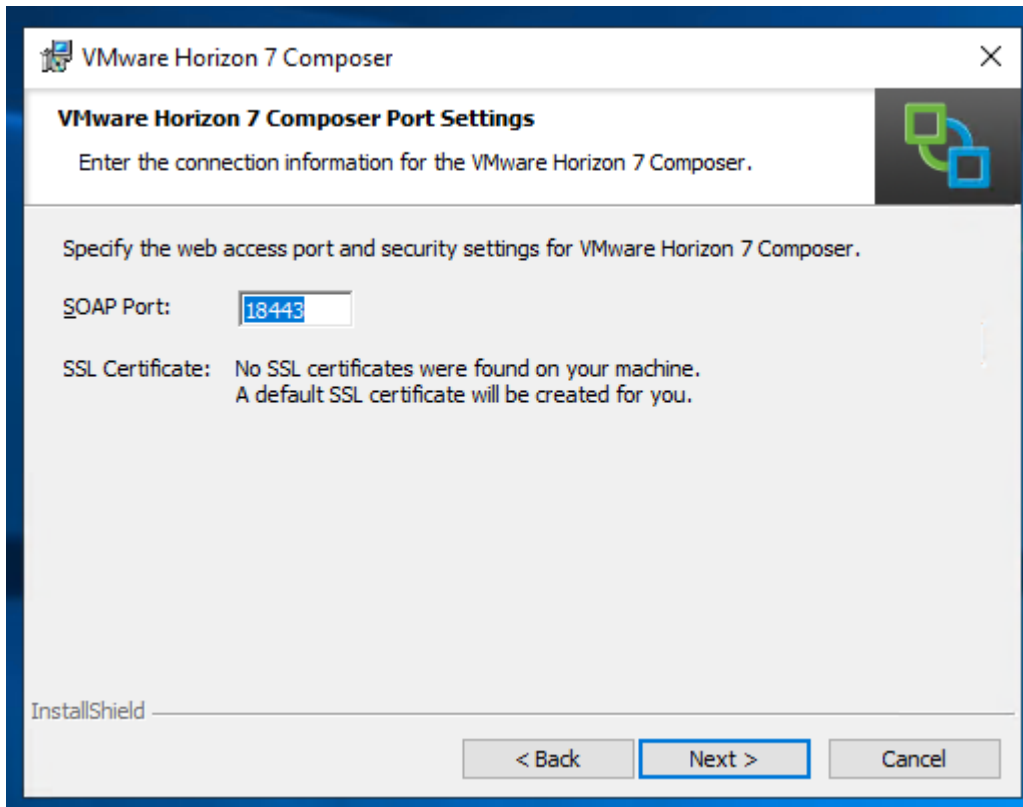
Step 4. Click **Next**.



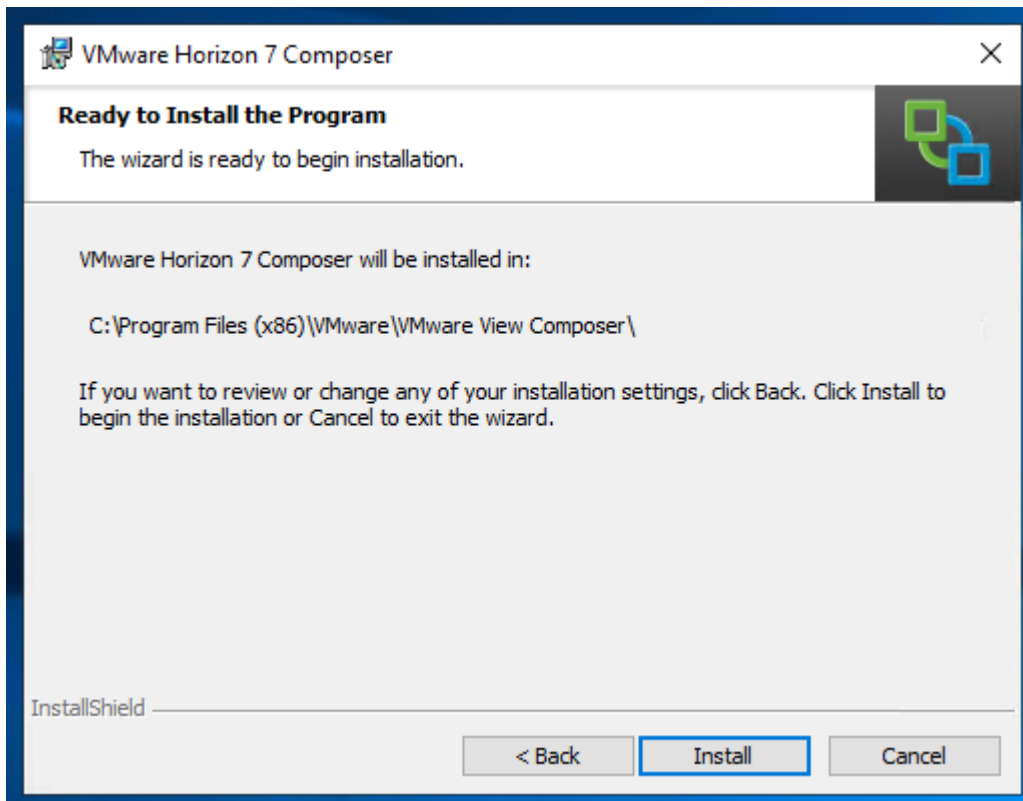
Step 5. Provide ODBC database connection details and click **Next**.



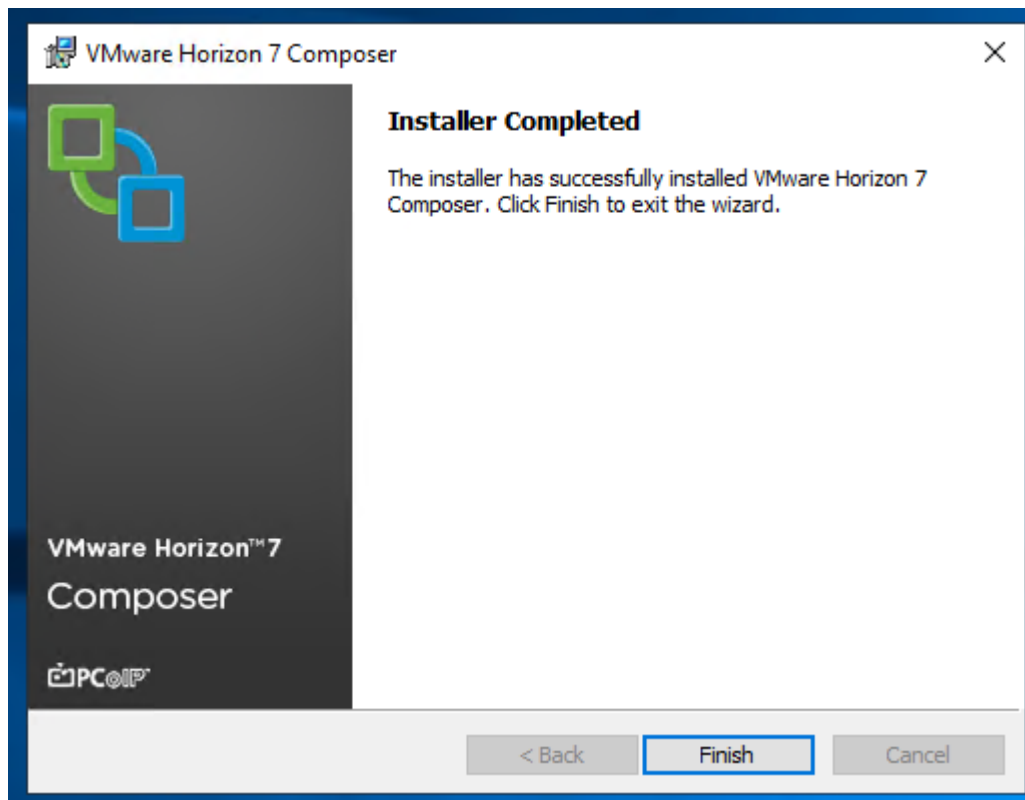
Step 6. Provide access port, security setting details and click **Next**.



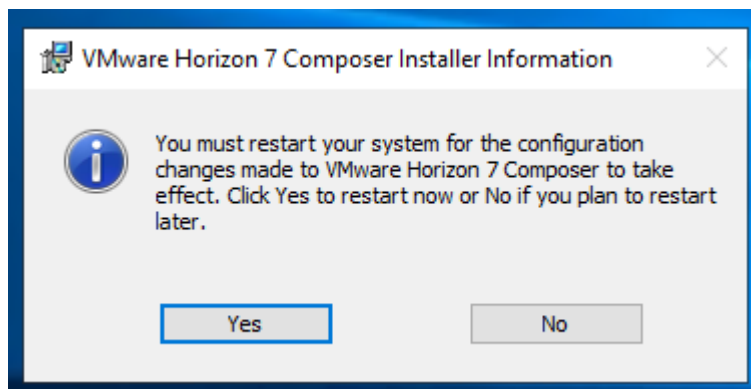
Step 7. Click **Install** to begin the installation.



Step 8. Click **Finish** to complete installation process.

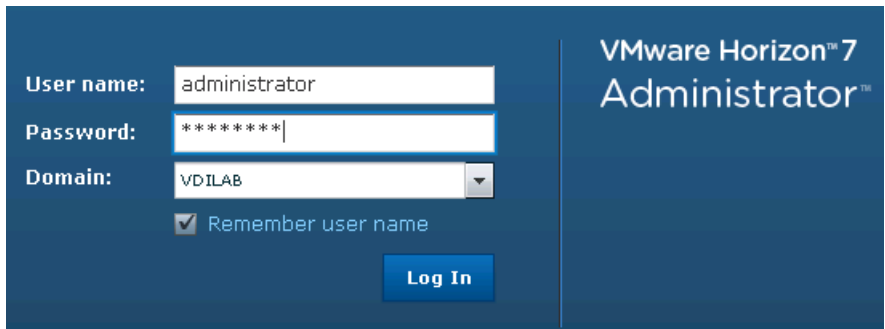


Step 9. If required, click **Yes** to reboot the Composer virtual machine.



Procedure 3. Create VMware Horizon Desktop Pool

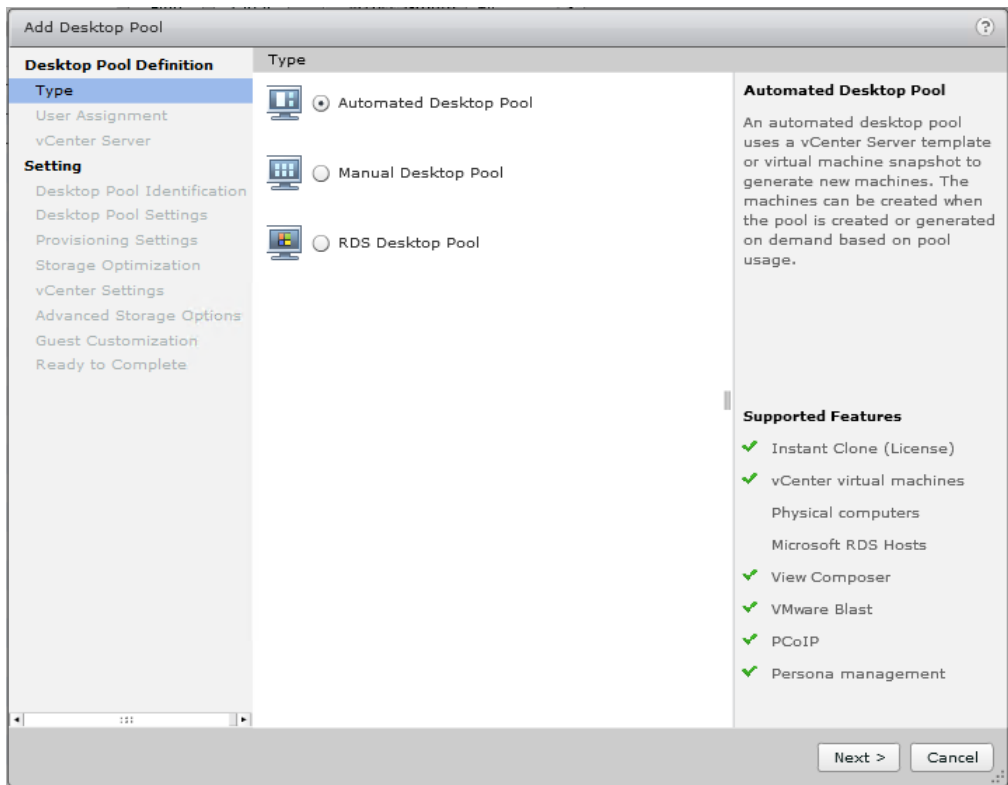
Step 1. Log into **Horizon 7 Administrator Console** via a web browser using Address or FQDN>/admin.



VMware Horizon™ 7
Administrator™

User name: administrator
Password: *****
Domain: VDILAB
 Remember user name
Log In

Step 2. Select the Type of Desktop Pool; **Automated**, **Manual**, or **RDS**. Click **Next**.



Add Desktop Pool

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Type

- Automated Desktop Pool
- Manual Desktop Pool
- RDS Desktop Pool

Automated Desktop Pool

An automated desktop pool uses a vCenter Server template or virtual machine snapshot to generate new machines. The machines can be created when the pool is created or generated on demand based on pool usage.

Supported Features

- ✓ Instant Clone (License)
- ✓ vCenter virtual machines
 - Physical computers
 - Microsoft RDS Hosts
- ✓ View Composer
- ✓ VMware Blast
- ✓ PCoIP
- ✓ Persona management

Next > Cancel

Step 3. Select the User assignment to be used by the desktop pool; **Dedicated** or **Floating**. Click **Next**.

Step 5. Provide Desktop Pool ID and Display name. Click Next.

Add Desktop Pool - W10-IC-VDI2

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification**
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings
- Guest Customization
- Ready to Complete

Desktop Pool Identification

ID: W10-IC-VDI2

Display name: W10-IC-VDI2

Access group: /

Description:

ID

The desktop pool ID is the unique name used to identify this desktop pool.

Display Name

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

Access Group

Access groups can organize the desktop pools in your organization. They can also be used for delegated administration.

Description

This description is only shown on the Settings tab for a desktop pool within View Administrator.

< Back Next > Cancel

Step 6. Select the Desktop Pool Settings. Click Next.

Add Desktop Pool - W10-IC-VDI2

Desktop Pool Definition

Type

User Assignment

vCenter Server

Setting

Desktop Pool Identification

Desktop Pool Settings

Provisioning Settings

Storage Optimization

vCenter Settings

Guest Customization

Ready to Complete

Desktop Pool Settings

General

State: ▾

Connection Server restrictions: None

Category Folder: None

Session Types: ▾ ?

Remote Settings

Automatically logoff after disconnect: ▾

Allow users to reset/restart their machines: ▾

Allow user to initiate separate desktop sessions from different client devices (desktops only): ▾ ?

Remote Display Protocol

Default display protocol: ▾

Allow users to choose protocol: ▾

3D Renderer: ▾ ?

HTML Access: Enabled ?

Requires installation of HTML Access.

Step 7. Provide the **Naming Pattern** and the **Max number of machines** to be provisioned (VDI pools consist of 1800 desktops). Click **Next**.

Add Desktop Pool - W10-IC-VDI2

Desktop Pool Definition

Type

User Assignment

vCenter Server

Setting

Desktop Pool Identification

Desktop Pool Settings

Provisioning Settings

Storage Optimization

vCenter Settings

Guest Customization

Ready to Complete

Provisioning Settings

Basic

Enable provisioning

Stop provisioning on error

Virtual Machine Naming

Use a naming pattern

Naming Pattern:

Desktop Pool Sizing

Max number of machines:

Number of spare (powered on) machines:

Provisioning Timing

Provision machines on demand

Min number of machines:

Provision all machines up-front

Virtual Device

Add a Trusted Platform Module (vTPM) device to the VMs

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.)

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales).

See the help for more naming pattern syntax options.

Step 8. For Storage Policy Management, select the appropriate option. Click **Next**.

Add Desktop Pool - W10-IC-VDI2

Desktop Pool Definition

Type

User Assignment

vCenter Server

Setting

Desktop Pool Identification

Desktop Pool Settings

Provisioning Settings

Storage Optimization

vCenter Settings

Guest Customization

Ready to Complete

Storage Optimization

Storage Policy Management

Use VMware Virtual SAN

Do not use VMware Virtual SAN

Virtual SAN is not available because no Virtual SAN datastores are configured.

Select separate datastores for replica and OS disks

Storage Optimization

Storage can be optimized by storing different kinds of data separately.

Step 9. Provide the **Parent VM in vCenter, Snapshot, VM folder location, Cluster, Resource pool,** and the **Datastores** information for the virtual machines.

The desktops are distributed across eight NFS datastores.

Add Desktop Pool - W10-IC-VDI2

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings**
- Guest Customization
- Ready to Complete

vCenter Settings

Default Image

- 1 Parent VM in vCenter: /FlexPod-Datacenter/vm/w10-IC-Maste
- 2 Snapshot: /VM Snapshot 11%252f15%252f2019,

Virtual Machine Location

- 3 VM folder location: /FlexPod-Datacenter/vm/Discovered vni

Resource Settings

- 4 Cluster: /FlexPod-Datacenter/host/VDI2
- 5 Resource pool: /FlexPod-Datacenter/host/VDI2/Resou
- 6 Datastores: 8 selected
- 7 Networks: Parent VM network selected

< Back Next > Cancel

Step 10. Select the **AD container** for the desktops to place in a Domain Controller computer location. Click **Next**.

The screenshot shows the 'Add Desktop Pool - W10-IC-VDI2' wizard window. The left sidebar is titled 'Desktop Pool Definition' and includes sections for 'Type', 'User Assignment', 'vCenter Server', and 'Setting'. Under 'Setting', there are links for 'Desktop Pool Identification', 'Desktop Pool Settings', 'Provisioning Settings', 'Storage Optimization', 'vCenter Settings', and 'Guest Customization' (which is currently selected and highlighted in blue). Below 'Guest Customization' is the text 'Ready to Complete'. The main area of the wizard is titled 'Guest Customization' and contains the following fields and options:

- Domain:** A dropdown menu showing 'VDILAB.local(administrator)'.
- AD container:** A text box containing 'OU=Target,OU=Computers,OU=LoginVSI' and a 'Browse...' button.
- Allow reuse of pre-existing computer accounts** (with a help icon).
- Image Publish Computer Account:** An empty text box with a help icon.
- Use ClonePrep** (checkbox).
- Power-off script name:** An empty text box with a help icon.
- Power-off script parameters:** An empty text box with an example: 'Example: p1 p2 p3'.
- Post-synchronization script name:** An empty text box with a help icon.
- Post-synchronization script parameters:** An empty text box with an example: 'Example: p1 p2 p3'.

At the bottom of the wizard window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

- Step 11.** Select **Entitle users after this wizard finishes** to enable desktop user group/users to access this pool.
- Step 12.** Review all the deployment specifications and click **Finish** to complete the deployment.

Add Desktop Pool - W10-IC-VDI2

Ready to Complete

Entitle users after this wizard finishes

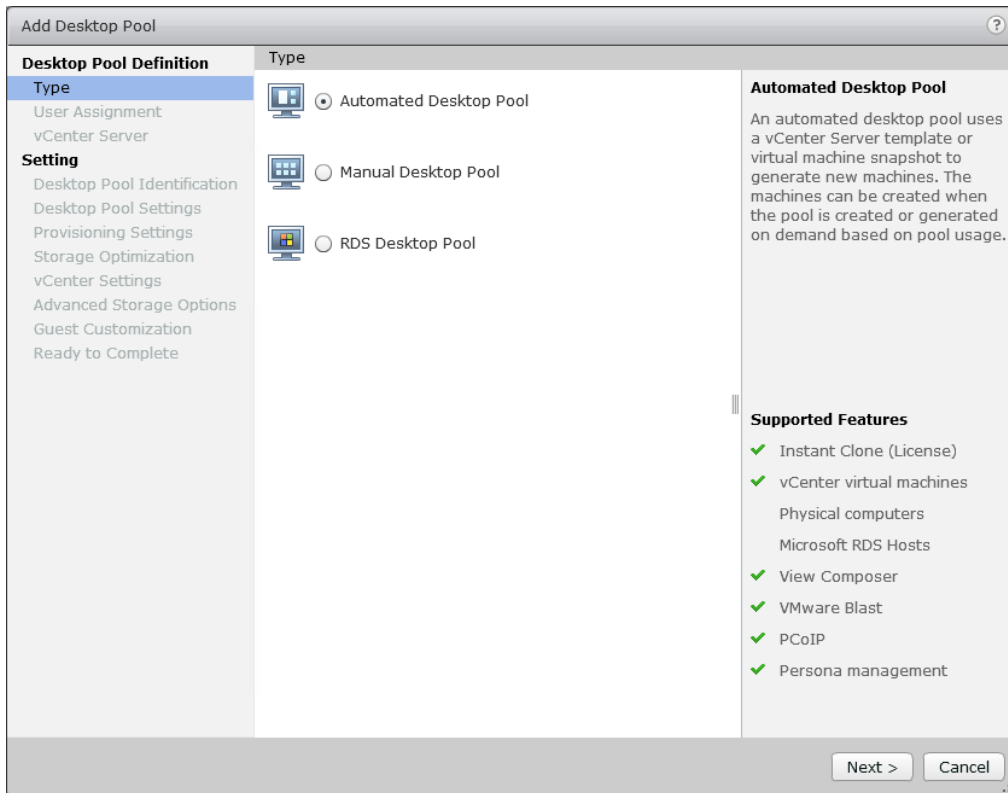
Type:	Automated
User assignment:	Floating assignment
vCenter Server:	vcsa-67.vdilab.local(administrator@vsphere.local)
Use View Composer:	No
Unique ID:	W10-IC-VDI2
Description:	
Display name:	W10-IC-VDI2
Access Group:	/
Desktop pool state:	Enabled
Session Types:	Desktop
Automatic logoff after disconnect:	Never
Connection Server restrictions:	None
Category Folder:	None
Allow users to reset/restart their machine:	No
Allow user to initiate separate desktop sessions from different client devices (desktops only):	No
Default display protocol:	PCoIP
Allow users to choose protocol:	Yes
3D Renderer:	Manage using vSphere Client
VRAM Size:	8 MB

< Back Finish Cancel

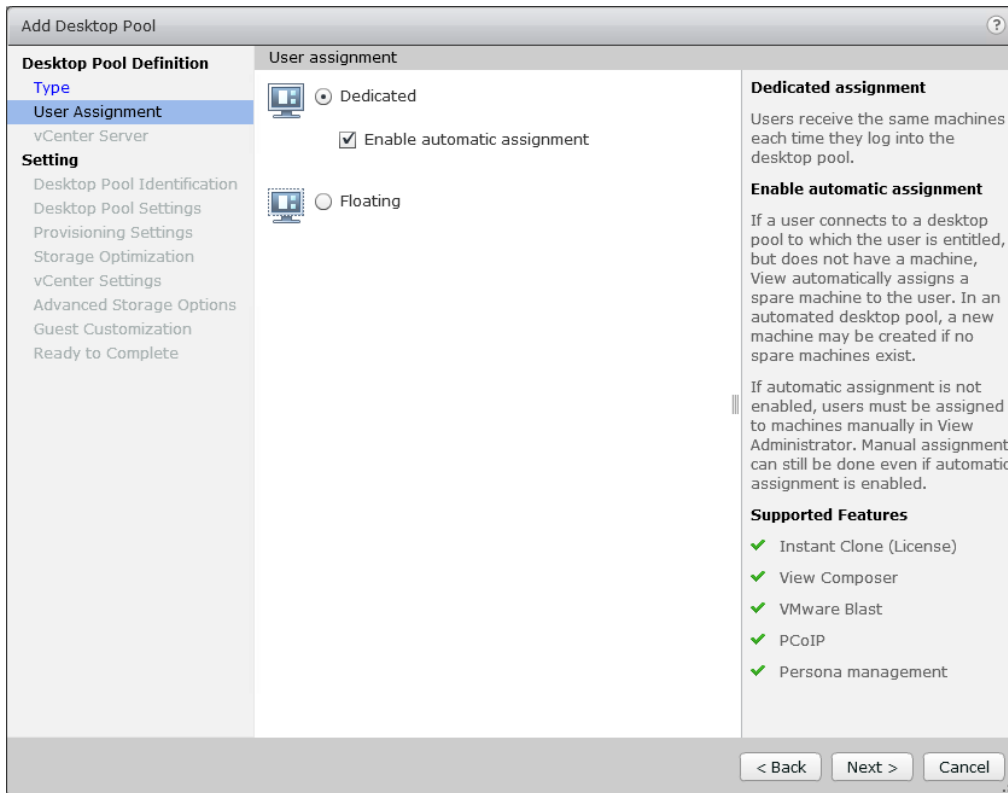
Procedure 4. Create VDI Full Clone Desktop Pool

This process describes the deployment using the Horizon Administrator Console.

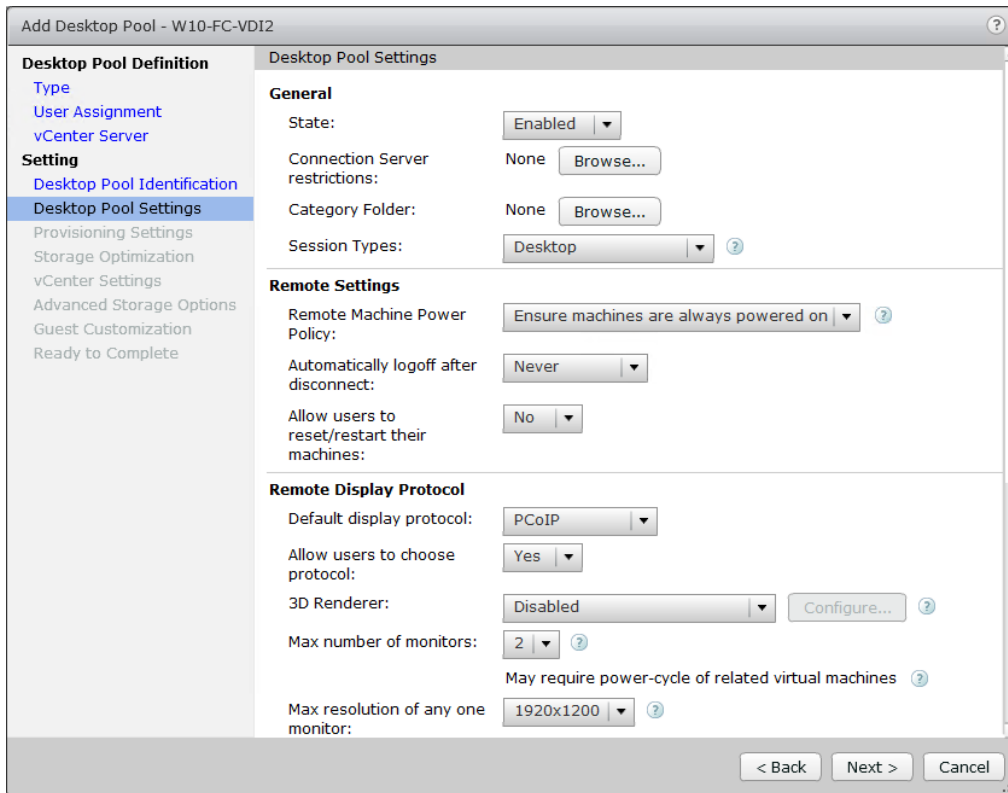
Step 1. Select the Type of Desktop Pool; **Automated**, **Manual**, or **RDS**. Click **Next**.



Step 2. Select the User assignment to be used by the desktop pool; **Dedicated** (select **Enable automated assignment**) or **Floating**. Click **Next**.



Step 3. Select the vCenter Server; **Instant clones**, **View Composer linked clones**, or **Full virtual machines**, and the type of Desktop deployment. (We created Horizon Instant Clones). Click **Next**.



Step 6. Provide the **Naming Pattern** and the **Max number of machines** to be provisioned (VDI pools consist of 1800 desktops). Click **Next**.

Add Desktop Pool - W10-FC-VDI2

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Provisioning Settings

Basic

Enable provisioning

Stop provisioning on error

Virtual Machine Naming

Specify names manually

0 names entered

Start machines in maintenance mode

Unassigned machines kept powered on:

Use a naming pattern

Naming Pattern:

Desktop Pool Sizing

Max number of machines:

Number of spare (powered on) machines:

Provisioning Timing

Provision machines on demand

Min number of machines:

Provision all machines up-front

Virtual Device

Add a Trusted Platform Module (vTPM) device to the VMs

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.)

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales).

See the help for more naming pattern syntax options.

< Back
Next >
Cancel

Step 7. For Storage Policy Management, select the appropriate option. Click **Next**.

Add Desktop Pool - W10-FC-VDI2

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

- Desktop Pool Identification
- Desktop Pool Settings
- Provisioning Settings
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Storage Optimization

Storage Policy Management

Use VMware Virtual SAN

Do not use VMware Virtual SAN

Virtual SAN is not available because no Virtual SAN datastores are configured.

Storage Optimization

Storage can be optimized by storing different kinds of data separately.

< Back
Next >
Cancel

Step 8. Provide the **Template**, **VM folder location**, **Host or cluster**, **Resource pool**, and the **Datastores** information for the virtual machines. Click **Next**.

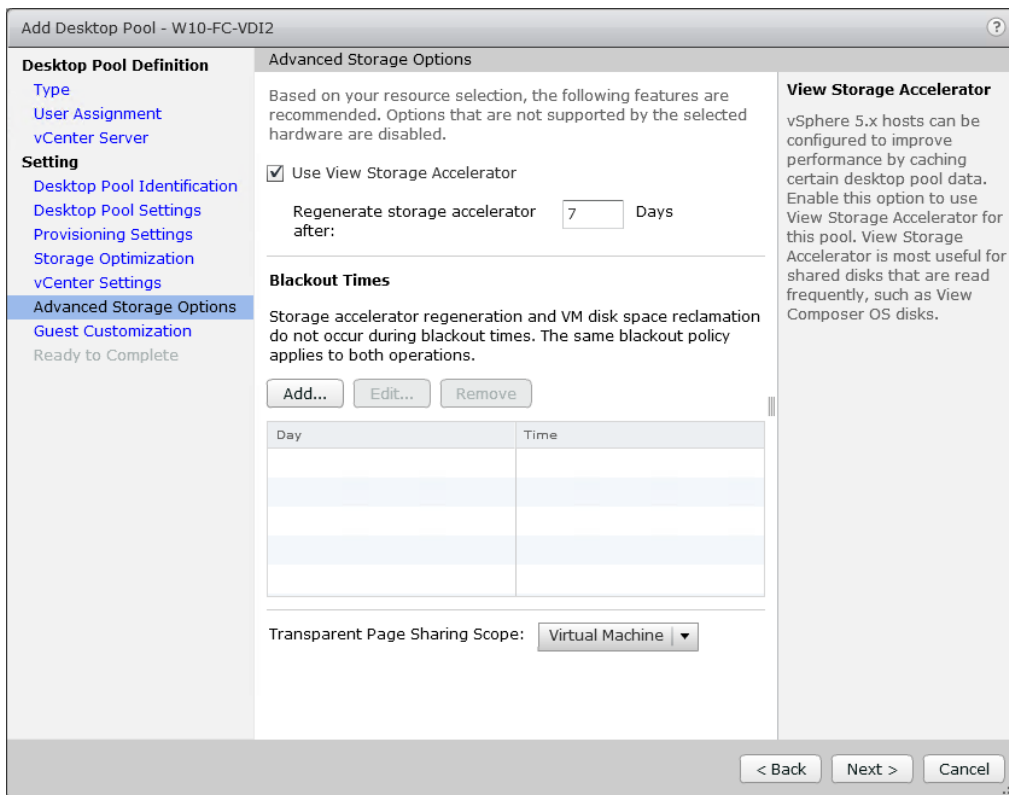
The desktops are distributed across eight NFS datastores.

The screenshot shows the 'Add Desktop Pool - W10-FC-VDI2' wizard window. The left sidebar contains a navigation menu with the following items: Desktop Pool Definition (Type, User Assignment, vCenter Server), Setting (Desktop Pool Identification, Desktop Pool Settings, Provisioning Settings, Storage Optimization), and vCenter Settings (Advanced Storage Options, Guest Customization, Ready to Complete). The main area is titled 'vCenter Settings' and contains the following sections:

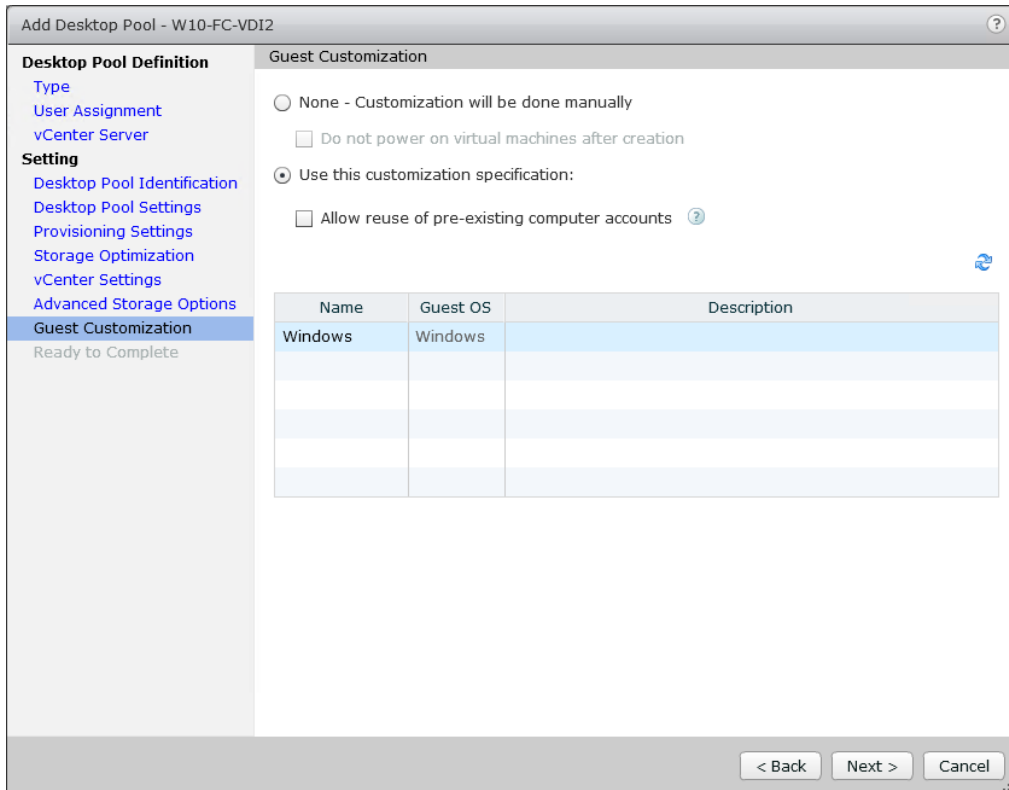
- Virtual Machine Template**: 1 Template: /FlexPod-Datacenter/vm/w10-1809-fc- [Browse...]
- Virtual Machine Location**: 2 VM folder location: /FlexPod-Datacenter/vm [Browse...]
- Resource Settings**: 3 Host or cluster: /FlexPod-Datacenter/host/VDI2 [Browse...]; 4 Resource pool: /FlexPod-Datacenter/host/VDI2/Resou [Browse...]; 5 Datastores: 8 selected [Browse...]

At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Step 9. Provide **Advanced Storage Options**.

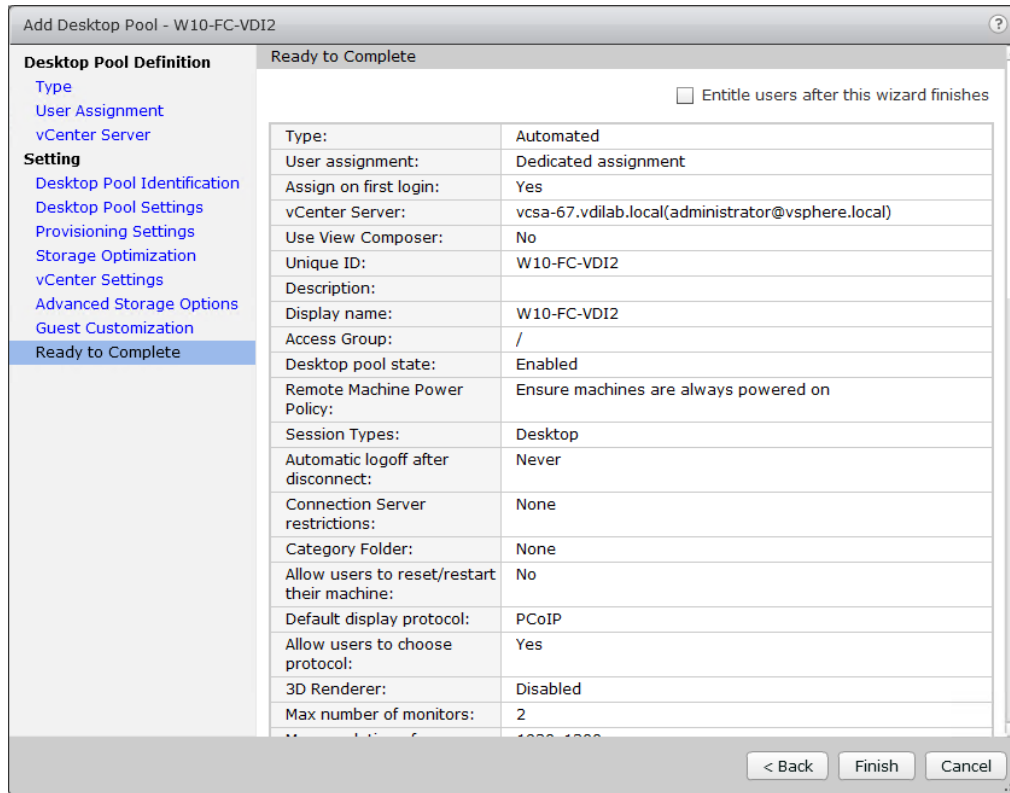


Step 10. Provide **Guest Customization** to be used during desktop deployment.



Step 11. Select **Entitle users after this wizard finishes** to enable desktop user group/users to access this pool.

Step 12. Review all the deployment specifications and click **Finish** to complete the deployment.



Procedure 5. Create RDS Farm

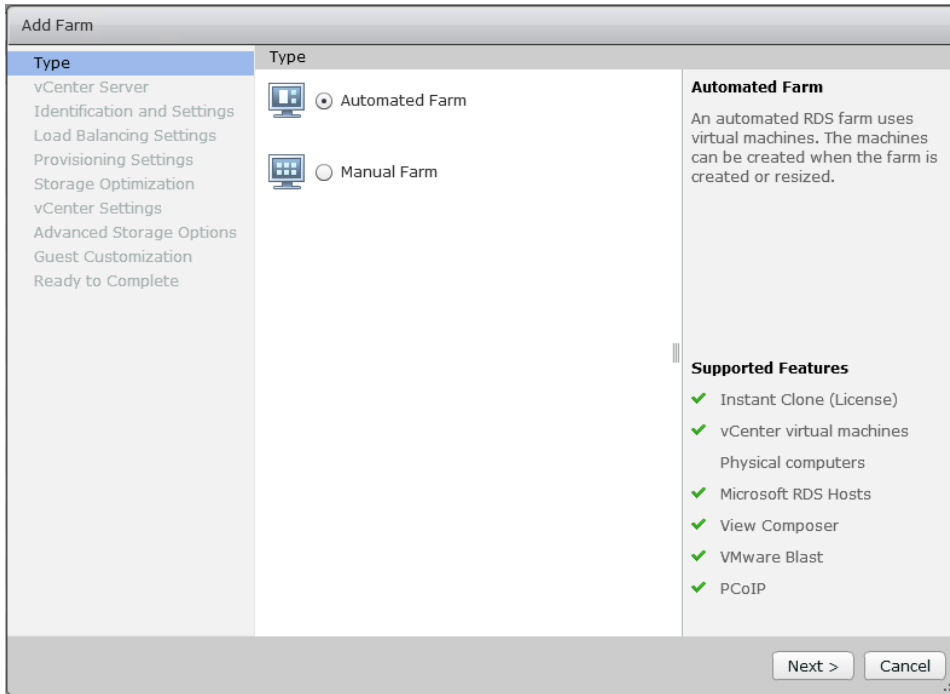
Tech tip

It is recommended to create a RDS Farm first with specifications set for RDS Server VMs and deploying a number of RDS servers required for users.

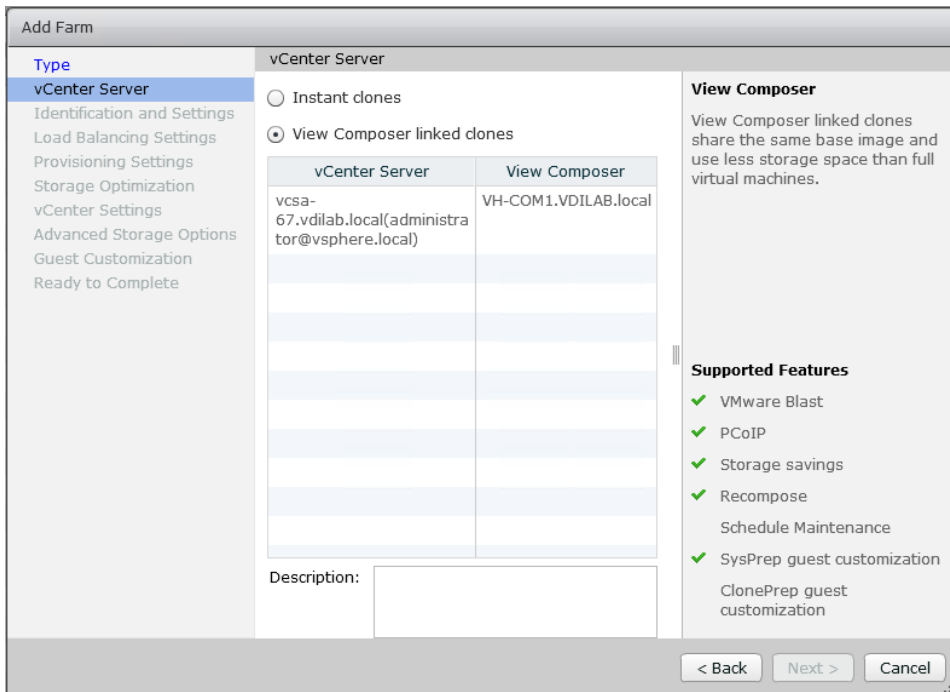
Step 1. Select the **FARM** when creating the RDS Pool.

You can entitle the user access at the RDS Pool level for RDS users/groups who can access the RDS VMs.

Step 2. From **Add Farm**, select the Type of Desktop Pool; **Automated Farm** or **Manual Farm**. We used Automated Farm for the RDS desktops in this design. Click **Next**.



Step 3. Select the vCenter Server; **Instant clones** or **View Composer linked clones** and the type of Desktop deployment. (We created Horizon Instant Clones). Click **Next**.



Step 4. Provide **ID** and **Description** for the RDS FARM. In Farm Settings, select the **Default display protocol** which is required for users to connect to the RDS Sessions. Click **Next**.

Add Farm - W2019-RDSH

Type
vCenter Server
Identification and Settings
Load Balancing Settings
Provisioning Settings
Storage Optimization
vCenter Settings
Advanced Storage Options
Guest Customization
Ready to Complete

Identification and Settings

General

ID: W2019-RDSH

Description:

Access group: /

Farm Settings

Default display protocol: PCoIP

Allow users to choose protocol: Yes

Pre-launch session timeout (applications only): After... 10 Minutes

Empty session timeout (applications only): After... 1 Minutes

When timeout occurs: Disconnect

Log off disconnected sessions: Never

< Back Next > Cancel

Step 5. Select the **Load Balancing Settings**. Click **Next**.

Add Farm - W2019-RDSH

Type
vCenter Server
Identification and Settings
Load Balancing Settings
Provisioning Settings
Storage Optimization
vCenter Settings
Advanced Storage Options
Guest Customization
Ready to Complete

Load Balancing Settings

Use custom script: Enabled

Include session count: Enabled

CPU usage threshold: 0

Memory usage threshold: 0

Disk queue length threshold: 0

Disk read latency threshold: 0

Disk write latency threshold: 0

< Back Next > Cancel

Step 6. In Provisioning Settings, enter the **Naming Pattern** for the RDS Desktop VMs you want to create and enter the **Max number of machines** you want to create on the RDS host or RDS cluster. Click **Next**.

Add Farm - W2019-RDSH

Type

- vCenter Server
- Identification and Settings
- Load Balancing Settings
- Provisioning Settings**
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Provisioning Settings

Basic

- Enable provisioning
- Stop provisioning on error

Virtual Machine Naming

Naming Pattern:

Farm Sizing

Max number of machines:

Minimum number of ready(provisioned) machines during View Composer maintenance operations:

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.).

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales.).

See the help for more naming pattern syntax options.

< Back Next > Cancel

Step 7. Complete the Storage Optimization settings as required. Click **Next**.

Add Farm - W2019-RDSH

Type

- vCenter Server
- Identification and Settings
- Load Balancing Settings
- Provisioning Settings
- Storage Optimization**
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Storage Optimization

Storage Policy Management

- Use VMware Virtual SAN
- Do not use VMware Virtual SAN

Virtual SAN is not available because no Virtual SAN datastores are configured.

- Select separate datastores for replica and OS disks

Virtual Volumes(VVOL) and fast NFS clones (VAAI) will be unavailable if the replica disks and OS disks are stored on separate datastores.

Storage can be optimized by storing different kinds of data separately.

Replica disks

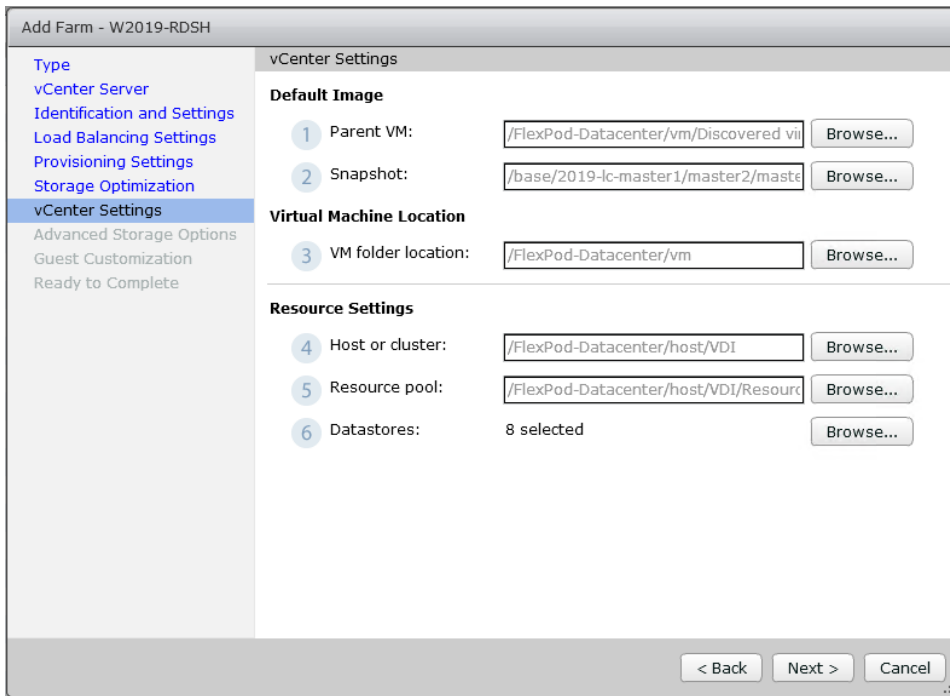
This option enables control over the placement of the replica that linked clones use as their base image.

It is recommended that a high performance datastore be chosen for these images. Depending on your hardware configuration, storing replicas on a separate datastore might create a single point of failure.

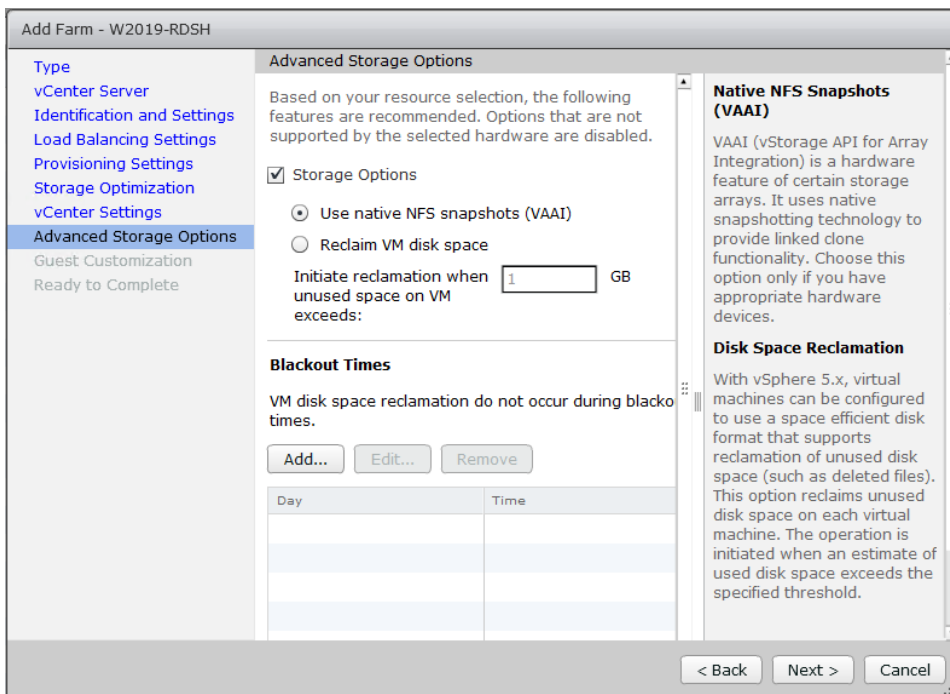
< Back Next > Cancel

Step 8. Provide the **Parent VM, Snapshot, VM folder location, Host or cluster, Resource pool**, and the **Datastores** information for the virtual machines. Click **Next**.

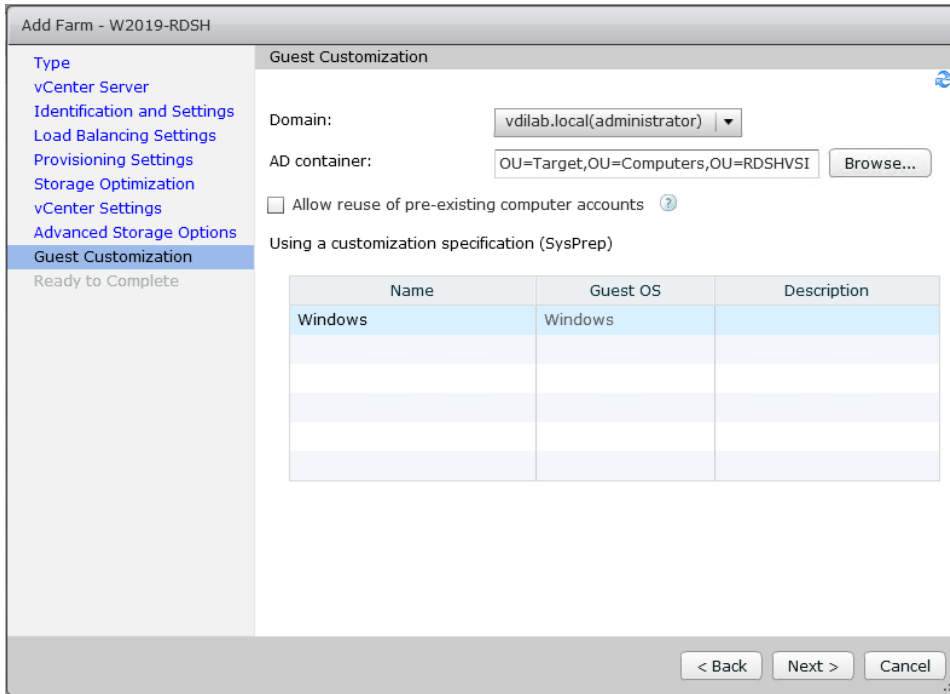
The desktops are distributed across eight NFS datastores.



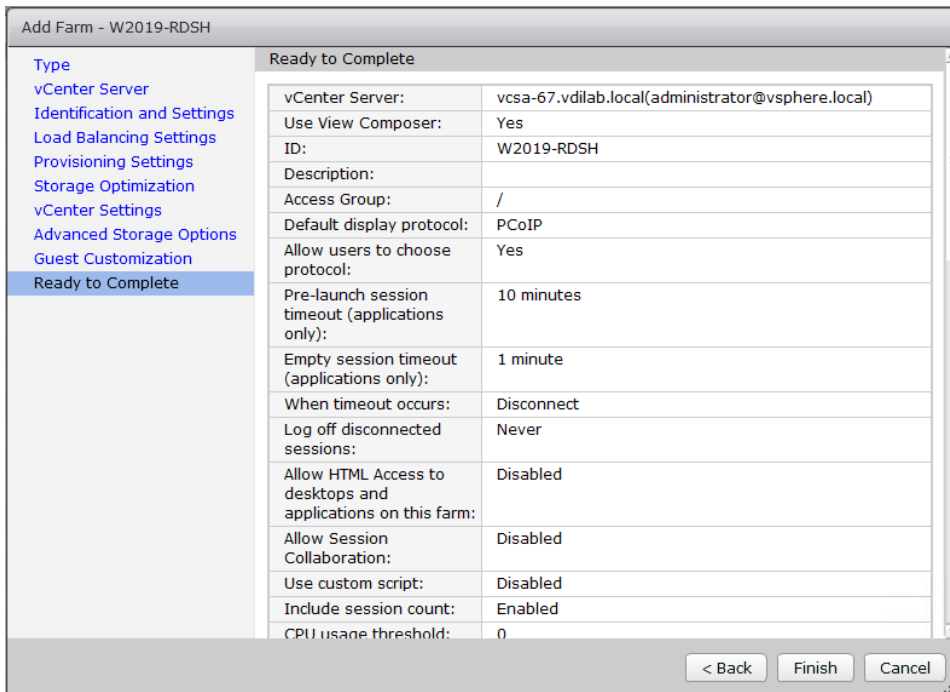
Step 9. Select the required **Advanced Storage Options**. Since VSC was used to configure VAAI support, native NFS snapshots option is used in this CVD. Click **Next**.



Step 10. In Guest Customization, select the **AD container** (VMs to be stored on the separate Computer VM (RDSHVS1) container in the Domain Controller) intended for storing RDS VMs and select the sysprep customization specs for creating VMware Composer provisioned RDS VMs. Click **Next**.



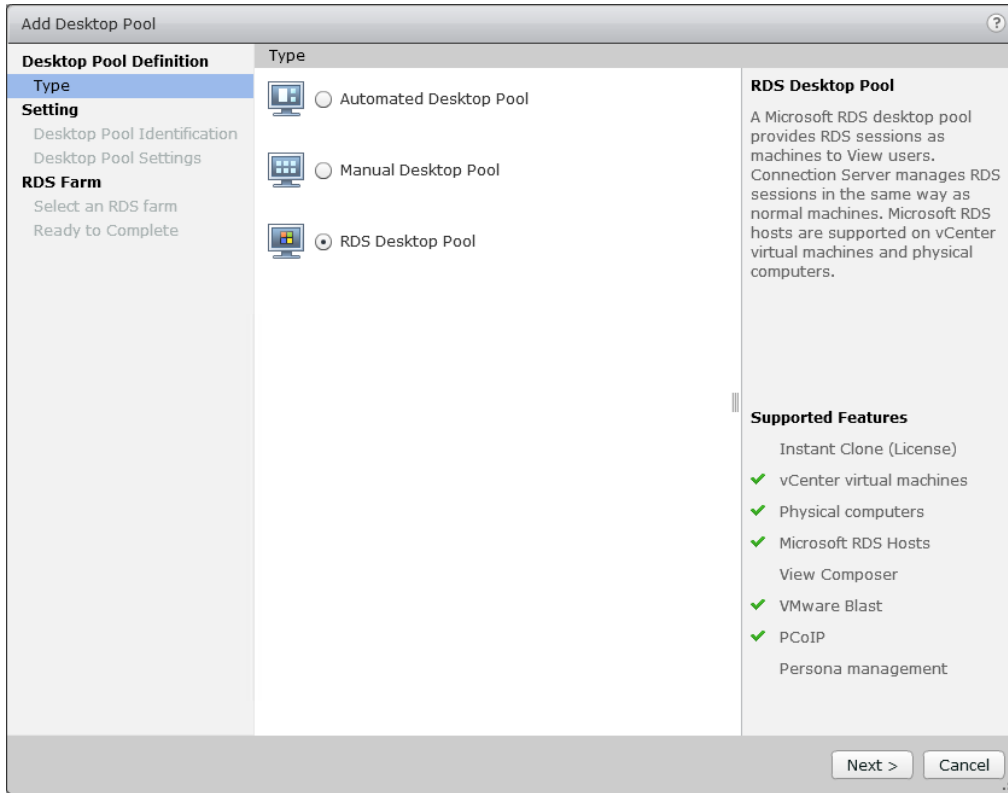
Step 11. Review the RDS Farm automatic deployment specifications and click **Finish** to complete the RDS pool deployment.



Procedure 6. Create RDS Pool

When the RDS Farm is created, you need to create an RDS Pool to absorb the RDS VMS FARM into the Pool for further managing the RDS pool.

Step 1. From the **Horizon Administrator Console**, click **Add Desktop Pool > Type >** and select; Auto-mated Desktop Pool, Manual Desktop Pool, or RDS Desktop Pool. The default choice is **RDS Desktop Pool**. Click **Next**.



Step 2. Provide Desktop Pool **ID** and **Display name**. Click **Next**.

Add Desktop Pool - W2019-LC-RDS

Desktop Pool Definition

Type

Setting

Desktop Pool Identification

Desktop Pool Settings

RDS Farm

Select an RDS farm

Ready to Complete

Desktop Pool Identification

ID:

Display name:

Description:

ID

The desktop pool ID is the unique name used to identify this desktop pool.

Display Name

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

Access groups can organize the desktop pools in your organization. They can also be used for delegated administration.

Description

This description is only shown on the Settings tab for a desktop pool within View Administrator.

Step 3. For the Desktop Pool Settings, keep the default settings. Click **Next**.

Add Desktop Pool - W2019-LC-RDS

Desktop Pool Definition

Type

Setting

Desktop Pool Identification

Desktop Pool Settings

RDS Farm

Select an RDS farm

Ready to Complete

Desktop Pool Settings

General

State: ▾

Connection Server restrictions: None

Category Folder: None

Client Restrictions: Enabled

Allow user to initiate separate desktop sessions from different client devices (desktops only): ▾ ?

Adobe Flash Settings for Sessions

Adobe Flash quality: ▾ ?

Adobe Flash throttling: ▾ ?

Add Desktop Pool - W2019-LC-RDS

Ready to Complete

Entitle users after this wizard finishes

Type:	RDS Desktop Pool
Unique ID:	W2019-LC-RDS
Description:	
Display name:	W2019-LC-RDS
Desktop pool state:	Enabled
Connection Server restrictions:	None
Category Folder:	None
Client Restrictions:	Disabled
Allow user to initiate separate desktop sessions from different client devices (desktops only):	No
Adobe Flash quality:	Do not control
Adobe Flash throttling:	Disabled
RDS Farm:	W2019-RDSH
Number of RDS hosts in the farm:	80

< Back Finish Cancel

Configuring user profile management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. 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. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting
- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for VMware Horizon desktop deployments. Screenshots of the User Profile Management interfaces that establish policies for this CVD's RDS and VDI users (for testing purposes) are shown below. Basic profile management policy settings are documented here:

Install and configure NVIDIA P6 card

This section focuses on installing and configuring the NVIDIA P6 cards with the Cisco UCS B200 M5 servers to deploy vGPU enabled virtual desktops.

Physical installation of P6 card into Cisco UCS B200 M5 server

The NVIDIA P6 graphics processing unit (GPU) card provides graphics and computing capabilities to the server. There are two supported versions of the NVIDIA P6 GPU card:

- UCSB-GPU-P6-F can be installed only in the front mezzanine slot of the server


Note

No front mezzanine cards can be installed when the server has CPUs greater than 165 W.

- UCSB-GPU-P6-R can be installed only in the rear mezzanine slot (slot 2) of the server.

Figure 49. illustrates the installed NVIDIA P6 GPU in the front and rear mezzanine slots.


Figure 49. NVIDIA GPU installed in the front and rear mezzanine slots



Procedure 1. Installing an NVIDIA GPU card in the front of the server


Figure 50. illustrates the front NVIDIA P6 GPU (UCSB-GPU-P6-F).

Figure 50. NVIDIA P6 GPU that installs in the front of the server



1	Leg with thumb screw that attaches to the server motherboard at the front	2	Handle to press down on when installing the GPU
---	---	---	---

Figure 51. Top-down view of the NVIDIA P6 GPU for the front of the server




1	Leg with thumb screw that attaches to the server motherboard	2	Thumb screw that attaches to a standoff below
---	--	---	---

Tech tip

Before installing the NVIDIA P6 GPU (UCSB-GPU-P6-F) in the front mezzanine slot, you need to upgrade the Cisco UCS domain that the GPU will be installed into to a version of Cisco UCS Manager that supports this card. Refer to the latest version of the release notes for Cisco UCS Software regarding supported hardware: <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html>. Remove the front mezzanine storage module if it is present. You cannot use the storage module in the front mezzanine slot when the NVIDIA P6 GPU is installed in the front of the server.

- Step 1.** Position the GPU in the correct orientation to the front of the server (callout 1) as shown in Figure 52.
- Step 2.** Install the GPU into the server. Press down on the handles (callout 5) to firmly secure the GPU.
- Step 3.** Tighten the thumb screws (callout 3) at the back of the GPU with the standoffs (callout 4) on the motherboard.
- Step 4.** Tighten the thumb screws on the legs (callout 2) to the motherboard.
- Step 5.** Install the drive blanking panels.

Figure 52. Installing the NVIDIA GPU in the front of the server




3006157

1	Front of the server	2	Leg with thumb screw that attaches to the motherboard
3	Thumbscrew to attach to standoff below	4	Standoff on the motherboard
5	Handle to press down on to firmly install the GPU	-	

Procedure 2. Installing an NVIDIA GPU card in the rear of the server

If you are installing the UCSB-GPU-P6-R to a server in the field, the option kit comes with the GPU itself (CPU and heatsink), a T-shaped installation wrench, and a custom standoff to support and attach the GPU to the mother-board. Figure 53. shows the three components of the option kit.

Figure 53. NVIDIA P6 GPU (UCSB-GPU-P6-R) option kit



1	NVIDIA P6 GPU (CPU and heatsink)	2	T-shaped wrench
3	Custom standoff	-	


Tech tip

Before installing the NVIDIA P6 GPU (UCSB-GPU-P6-R) in the rear mezzanine slot, you need to upgrade the Cisco UCS domain that the GPU will be installed into to a version of Cisco UCS Manager that supports this card. Refer to the latest version of the release notes for Cisco UCS Software regarding supported hardware: <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html>. Remove any other card, such as a VIC 1480, VIC 1380, or VIC port expander card from the rear mezzanine slot. You cannot use any other card in the rear mezzanine slot when the NVIDIA P6 GPU is installed.

Step 1. Use the T-shaped wrench that comes with the GPU to remove the existing standoff at the back end of the mother-board.

- Step 2.** Install the custom standoff in the same location at the back end of the motherboard.
- Step 3.** Position the GPU over the connector on the motherboard and align all the captive screws to the standoff posts (callout 1).
- Step 4.** Tighten the captive screws (callout 2).

Figure 54. Installing the NVIDIA P6 GPU in the rear mezzanine slot



Procedure 3. Install the NVIDIA VMware VIB driver

- Step 1.** From the **Cisco UCS Manager**, verify the GPU card has been properly installed.

Equipment / Chassis / Chassis 1 / Servers / Server 3

General | Inventory | Virtual Machines | Installed Firmware | CIMC Sessions | SEL Logs | VIF Paths | Health | Diagnostics | Faults | Events | FSM | Statistics | Temperatures | Power

Motherboard | CIMC | CPUs | **GPUs** | Memory | Adapters | HBAs | NICs | iSCSI vNICs | Security | Storage

Advanced Filter | Export | Print

Name	ID	Model	Serial	Mode
Graphics Card 2	2	UCSB-GPU-P6-R	FCH212373US	NA
Graphics Card 3	3	UCSB-GPU-P6-F	FCH21237472	NA

Step 2. Download the **NVIDIA GRID GPU driver pack** for VMware vSphere ESXi 6.7.

Step 3. Upload the NVIDIA driver (vSphere Installation Bundle [VIB] file) to the /tmp directory on the ESXi host using a tool such as WinSCP. (Shared storage is preferred if you are installing drivers on multiple servers or using the VMware Update Manager.)

Step 4. Log in as **root** to the vSphere console through SSH using a tool such as Putty.

Step 5. The ESXi host must be in maintenance mode for you to install the VIB module. To place the host in maintenance mode, use the command `esxcli system maintenanceMode set -enable true`.

Step 6. Enter the following command to install the NVIDIA vGPU drivers:

```
esxcli software vib install --no-sig-check -v /<path>/<filename>.VIB
```

The command should return output similar to that shown here:

```
# esxcli software vib install --no-sig-check -v /tmp/NVIDIA-
VMware_ESXi_6.7_Host_Driver_384.99-1OEM.650.0.0.4598673.vib
Installation Result
    Message: Operation finished successfully.
    Reboot Required: false
    VIBs Installed: NVIDIA_bootbank_NVIDIA-VMware_ESXi_6.7_Host_Driver_384.99-
1OEM.650.0.0.4598673
    VIBs Removed:
    VIBs Skipped:
```

Note

Although the display shows “Reboot Required: false,” a reboot is necessary for the VIB file to load and for xorg to start.

Step 7. Exit the ESXi host from maintenance mode and reboot the host by using the vSphere Web Client or by entering the following commands:

```
#esxcli system maintenanceMode set -e false
#reboot
```

Step 8. After the host reboots successfully, verify that the kernel module has loaded successfully using the following command:

```
#esxcli software vib list | grep -i nvidia
```

The command should return output similar to that shown here:

```
# esxcli software vib list | grep -i nvidia
NVIDIA-VMware_ESXi_6.7_Host_Driver  384.99-1OEM.650.0.0.4598673      NVIDIA
VMwareAccepted      2017-11-27
```

See the VMware knowledge base article for information about removing any existing NVIDIA drivers before installing new drivers:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2033434

Step 9. Confirm GRID GPU detection on the ESXi host. To determine the status of the GPU card’s CPU, the card’s memory, and the amount of disk space remaining on the card, enter the following command:

```
#nvidia-smi
```

The command should return output similar to that shown in Figure 55. depending on the card used in your environment.

Figure 55. VMware ESX SSH Console report for GPU P6 card detection on Cisco UCS B200 M5 Blade Server

```
-sh: nvidia-smi: not found
[root@M5:~] nvidia-smi
Wed Sep  6 00:43:04 2017
-----+-----
| NVIDIA-SMI 384.73                Driver Version: 384.73          |
|-----+-----|
| GPU Name               Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|-----+-----|
| 0 Tesla P6             On         | 00000000:18:00.0 Off  |    Off                |
| N/A   21C    P8      9W /  90W |  41MiB / 16383MiB |      0%      Default  |
|-----+-----|
| 1 Tesla P6             On         | 00000000:D8:00.0 Off  |    Off                |
| N/A   35C    P8     10W /  90W |  41MiB / 16383MiB |      0%      Default  |
|-----+-----|
|
| Processes:
| GPU      PID   Type   Process name                      GPU Memory
|-----+-----|
|          |          |       |                                 Usage
|-----+-----|
| No running processes found
|-----+-----|
[root@M5:~] █
```

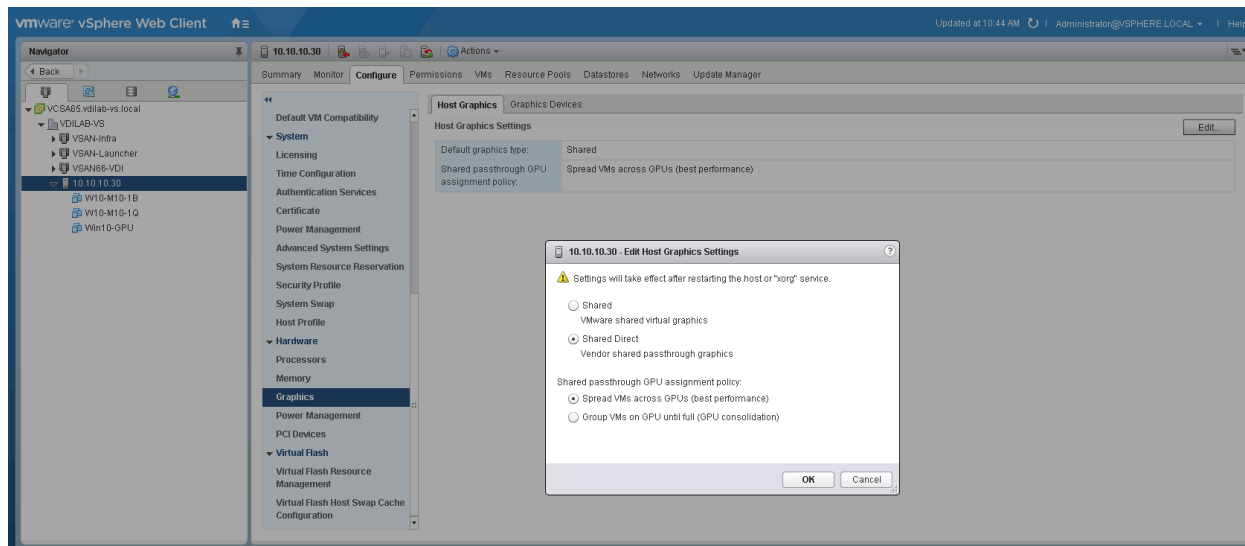
The NVIDIA system management interface (SMI) also allows GPU monitoring using the following command: `nvidia-smi -l` (this command adds a loop, automatically refreshing the display).

Procedure 4. Configure a virtual machine with a vGPU

Create the virtual machine that you will use as the VDI base image.

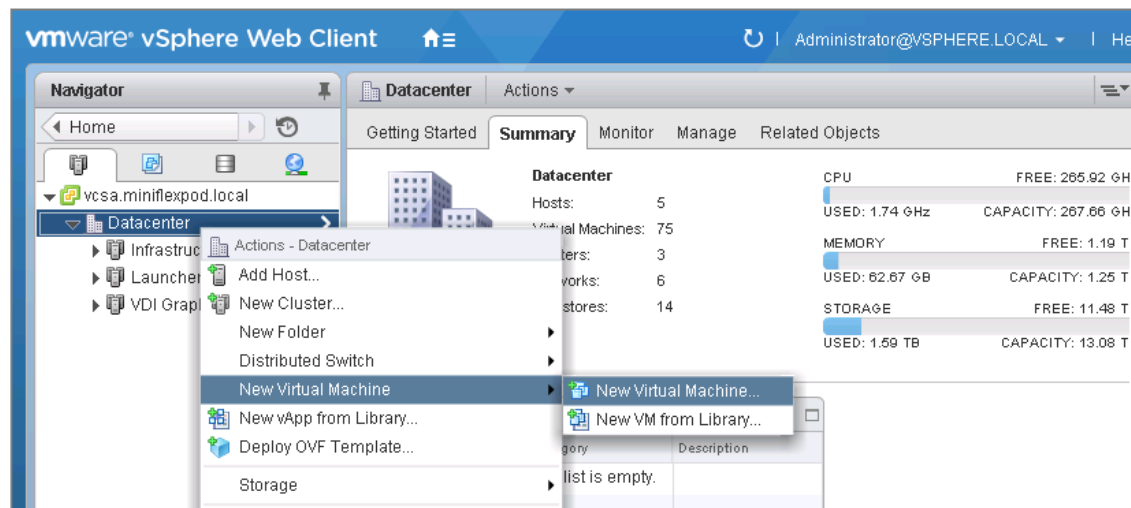
Step 1. Select the **ESXi host** and click the **Configure** tab. From the list of options, choose **Graphics > Edit Host Graphics Settings**. Select **Shared Direct Vendor shared passthrough graphics**. Reboot the system to make the changes effective.

Figure 56. Edit Host Graphics Settings



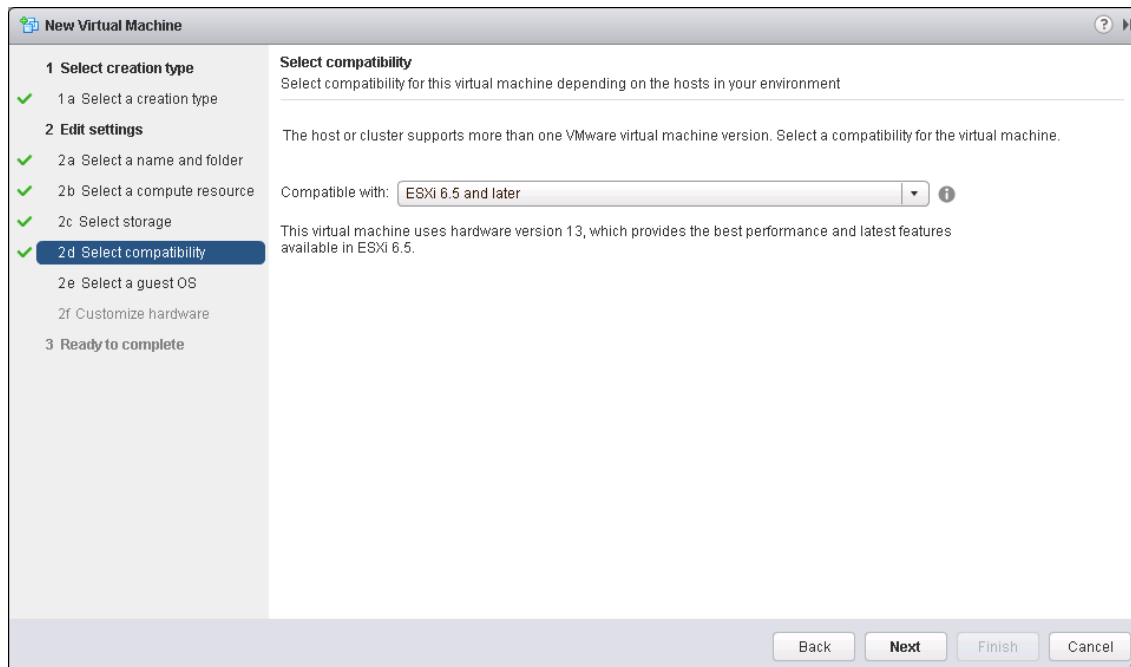
Step 2. Using the **vSphere Web Client**, create a new virtual machine. To do this, right-click a host or cluster and choose **New Virtual Machine**. Work through the New Virtual Machine wizard. Unless another configuration is specified, select the configuration settings appropriate for your environment

Figure 57. Creating a New Virtual Machine in VMware vSphere Web Client



Step 3. From **Select compatibility**, from the **Compatible with** drop-down list, choose **ESXi 6.0 and later** to use the latest features, including the mapping of shared PCI devices, which is required for the vGPU feature. This solution uses ESXi 6.7 and later which provides the latest features available in ESXi 6.7 and virtual machine hardware Version 13.

Figure 58. Selecting virtual machine hardware version 11 or later




Step 4. To customize the hardware of the new virtual machine, from the drop-down list add a new shared PCI device, select the appropriate GPU profile, and reserve all virtual machine memory. Click **Add**.

Note

If you are creating a new virtual machine and using the vSphere Web Client's virtual machine console functions, the mouse will not be usable in the virtual machine until after both the operating system and VMware Tools have been installed. If you cannot use the traditional vSphere Web Client to connect to the virtual machine, do not enable the NVIDIA GRID vGPU at this time.

Figure 59. Adding a Shared PCI Device to the virtual machine to attach the GPU profile



Step 5. A virtual machine with a vGPU assigned will not start if ECC is enabled. If this is the case, as a workaround, disable ECC by entering the following commands:

```
# nvidia-smi -i 0 -e 0
# nvidia-smi -i 1 -e 0
```

Step 6. Use `-i` to target a specific GPU. If two cards are installed in a server, run the command twice as shown in the example below, where 0 and 1 each specify a GPU card.

Figure 60. Disabling ECC

```

-sh: nvidia-smi: not found
[root@M5:~] nvidia-smi
Wed Sep  6 00:43:04 2017
+-----+
| NVIDIA-SMI 384.73                 Driver Version: 384.73          |
+-----+

+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|  0   Tesla P6             0n      | 0000:18:00.0  Off   |    0%      Default    |
| N/A   22C   P8      9W /  90W |  39MiB / 15359MiB |             |
+-----+-----+
|  1   Tesla P6             0n      | 0000:D8:00.0  Off   |    0%      Default    |
| N/A   37C   P8     10W /  90W |  39MiB / 15359MiB |             |
+-----+-----+

+-----+
| Processes:                        GPU Memory |
| GPU       PID  Type  Process name                        Usage    |
+-----+-----+
| No running processes found         |
+-----+

[root@M5:~] esxtop -a -b -d 10 -n 600 > /vmfs/volumes/594d8376-1531284a-003b-0025b5000a2f/215U-003.csv
[root@M5:~] nvidia-smi -i 0 -e 0
-sh: nvidia-smi: not found
[root@M5:~] nvidia-smi -i 0 -e 0
Disabled ECC support for GPU 0000:18:00.0.
All done.
Reboot required.
[root@M5:~] nvidia-smi -i 1 -e 0
Disabled ECC support for GPU 0000:D8:00.0.
All done.
Reboot required.
[root@M5:~]

```

Procedure 5. Install and configure Microsoft Windows on the virtual machine

- Step 1.** Configure the virtual machine with the appropriate amount of vCPU and RAM according to the GPU profile selected.
- Step 2.** Install **VMware Tools**.
- Step 3.** Join the virtual machine to the Microsoft Active Directory domain.
- Step 4.** From the Windows System Properties menu, select **Allow remote connections to this computer**.
- Step 5.** Install **VMware Horizon Agent** with appropriate settings. Enable the remote desktop capability if prompted to do so.
- Step 6.** Install **Horizon Direct Connection agent**.
- Step 7.** Optimize the Windows OS. [VMware OSOT](#), the Operating System Optimization Tool, includes customizable templates to enable or disable Windows system services and features using VMware recommendations and best practices across multiple systems. Since most Windows system services are enabled by default, the optimization tool can be used to easily disable unnecessary services and features to improve performance.
- Step 8.** Restart the **Windows OS** when prompted to do so.

Procedure 6. Install the GPU drivers inside Windows Virtual Machine

Tech tip

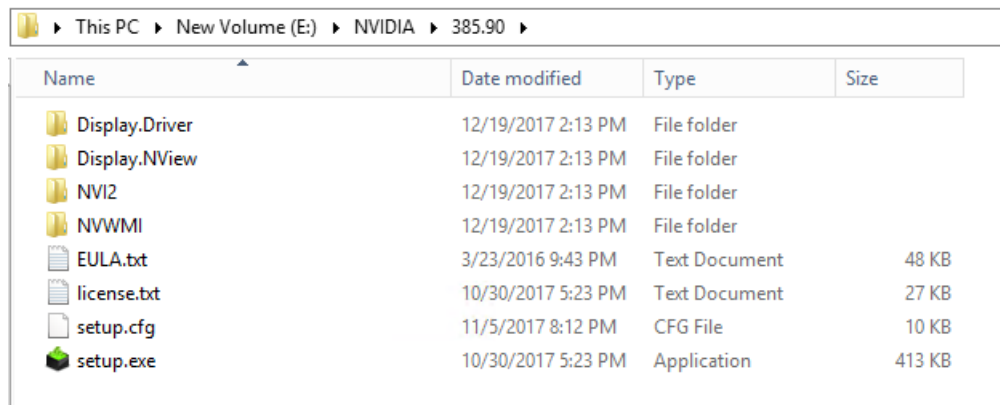
It is important to note that the drivers installed with the Windows VDI desktop must match the version that accompanies the driver for the ESXi host. So, if you downgrade or upgrade the ESXi host vtb, you must do the same with the NVIDIA driver in your Windows master image.

In this study, we used ESXi Host Driver version 352.83 and 354.80 for the Windows VDI image. These drivers come in the same download package from NVIDIA.

Step 1. Copy the **Microsoft Windows drivers** from the NVIDIA GRID vGPU driver pack downloaded earlier to the master virtual machine.

Step 2. Copy the **32- or 64-bit NVIDIA Windows driver** from the vGPU driver pack to the desktop virtual machine and run setup.exe.

Figure 61. NVIDIA Driver Pack



The screenshot shows a Windows File Explorer window with the address bar set to "This PC > New Volume (E:) > NVIDIA > 385.90". The main area displays a list of files and folders with columns for Name, Date modified, Type, and Size.

Name	Date modified	Type	Size
Display.Driver	12/19/2017 2:13 PM	File folder	
Display.NView	12/19/2017 2:13 PM	File folder	
NVI2	12/19/2017 2:13 PM	File folder	
NVWMI	12/19/2017 2:13 PM	File folder	
EULA.txt	3/23/2016 9:43 PM	Text Document	48 KB
license.txt	10/30/2017 5:23 PM	Text Document	27 KB
setup.cfg	11/5/2017 8:12 PM	CFG File	10 KB
setup.exe	10/30/2017 5:23 PM	Application	413 KB

Note

The vGPU host driver and guest driver versions need to match. Do not attempt to use a newer guest driver with an older vGPU host driver or an older guest driver with a newer vGPU host driver. In addition, the vGPU driver from NVIDIA is a different driver than the GPU pass-through driver.

Step 3. To accept the terms of the NVIDIA software license agreement, click **Agree and Continue**.

Figure 62. NVIDIA Software License



Step 4. From the Installation options, select **Express (Recommended)** or **Custom (Advanced)**. Click **Next**. After the installation has completed successfully, **restart** the **virtual machine**.

Note

Make sure that remote desktop connections are enabled. After this step, console access may not be available for the virtual machine when you connect from a vSphere Client.

Figure 63. Express or Custom Installation options



Figure 64. Components installed during NVIDIA Graphics Driver custom installation process



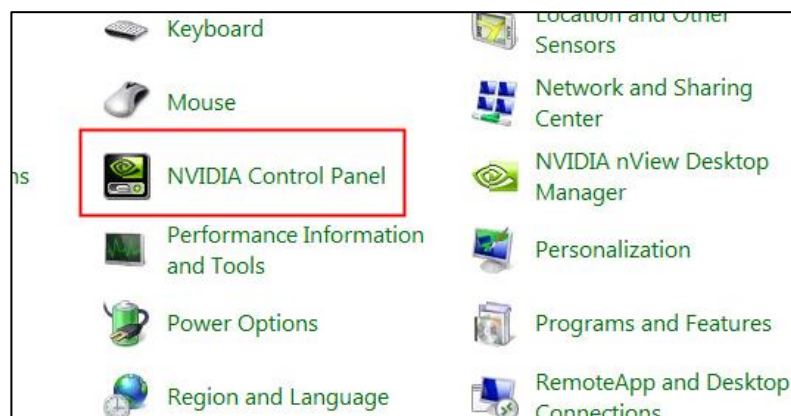
Figure 65. Restarting the virtual machine




Procedure 7. Configure NVIDIA Grid License Server on virtual machine

When the license server is properly installed, you must point the master image to the license server so the virtual machines with vGPUs can obtain a license.

Step 1. In the Windows Control Panel, double-click **NVidia Control Panel**.



Step 2. In the **NVIDIA Control Panel**, enter the IP or FQDN of the Grid License Server. You will receive a result similar to the one shown below.



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 de-liver 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 use 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 the 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 66. Example of Cisco Intersight Dashboard for FlexPod UCS Domain

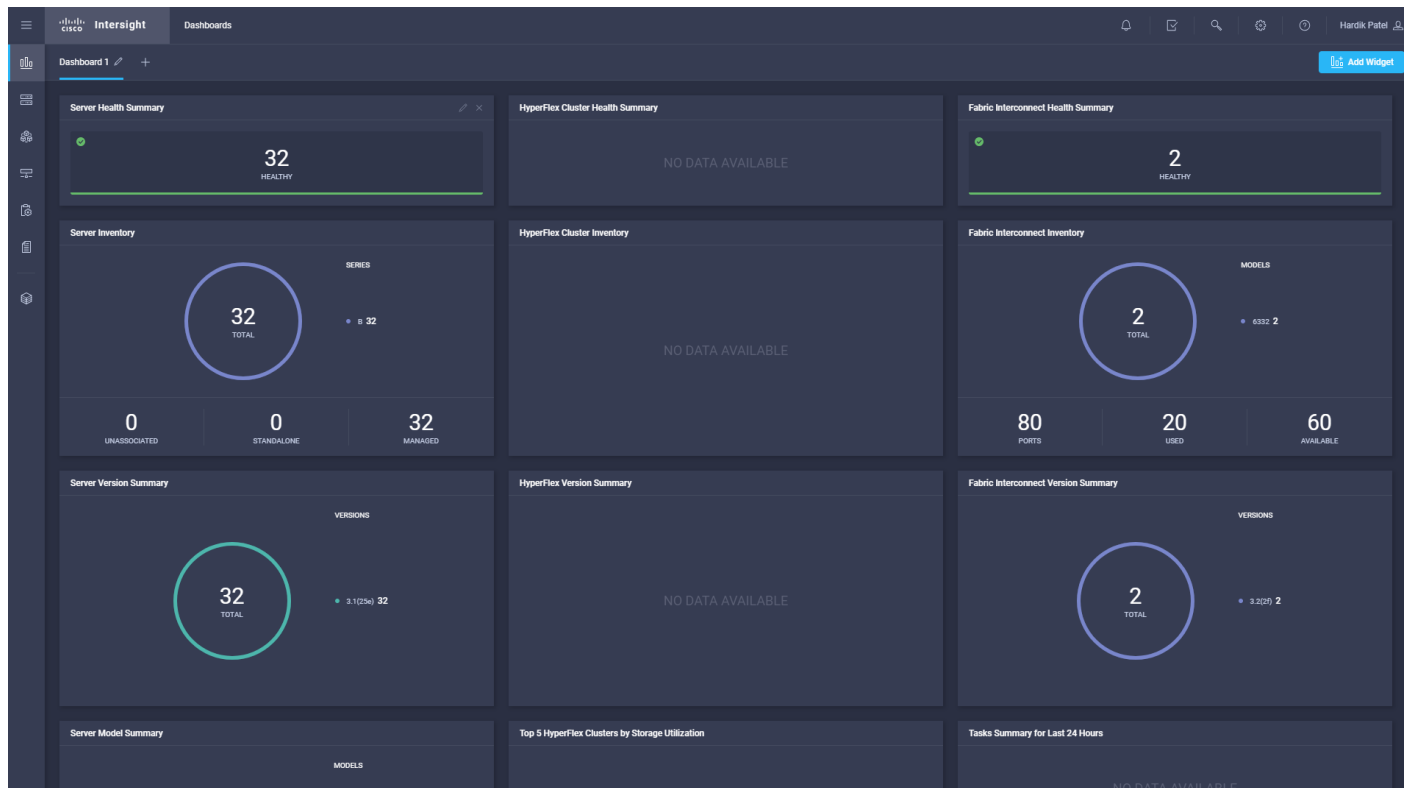


Figure 67. Cisco UCS Manager Device Connector example



Test Setup, Configuration, and Load Recommendation

In this solution, we tested a single Cisco UCS B200 M5 blade server to validate against the performance of one blade and thirty Cisco UCS B200 M5 blade servers across four chassis to illustrate linear scalability for each workload use case studied.

Cisco UCS test configuration for single blade scalability

This test case validates each workload on a single blade to determine the Recommended Maximum Workload per host server using Vmware Horizon 7.10 with 224 RDS sessions, 180 VDI Non-Persistent sessions, and 180 VDI Persistent sessions.

Figure 68. Cisco UCS B200 M5 Blade Server for Single Server Scalability Vmware Horizon 7.10 RDS Linked Clones

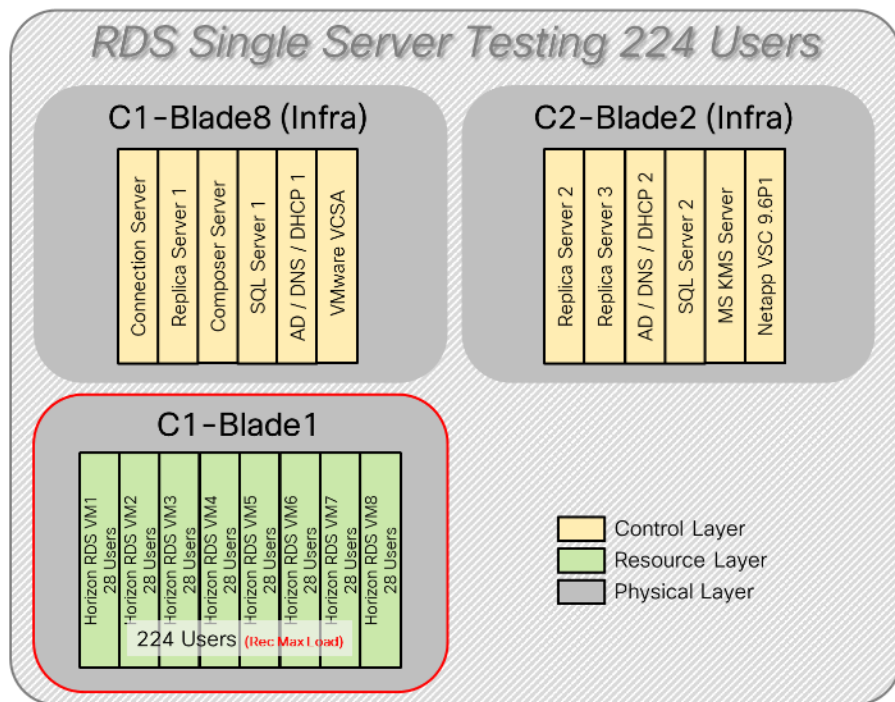


Figure 69. Cisco UCS B200 M5 Blade Server for Single Server Scalability Vmware Horizon 7.10 VDI (Non-Persistent) Instant Clones

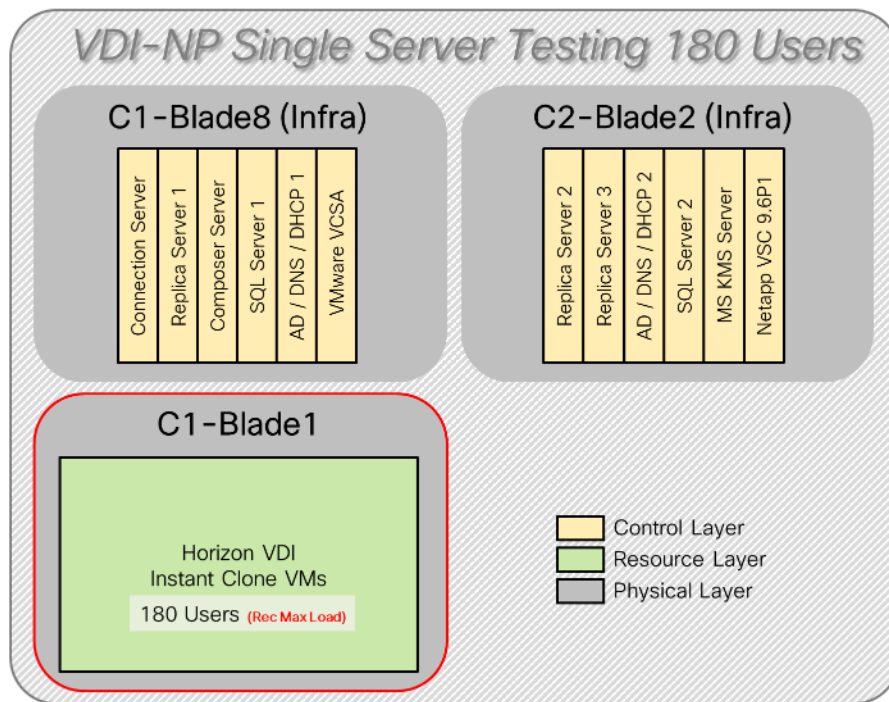
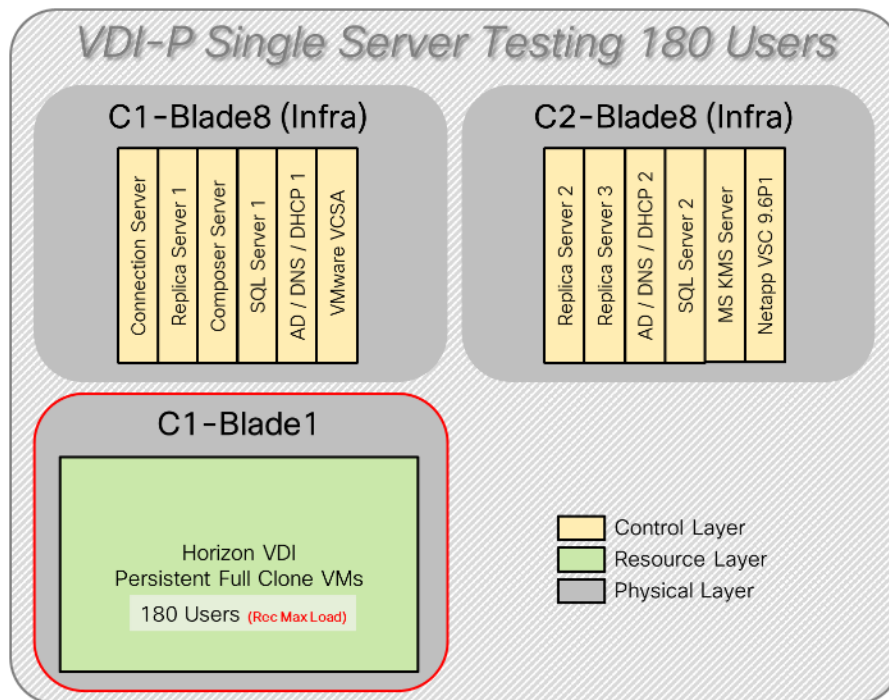


Figure 70. Cisco UCS B200 M5 Blade Server for Single Server Scalability Vmware Horizon 7.10 VDI (Persistent) Full Clones



Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M5 Blade Servers with Intel Xeon Silver 4114 2.20-GHz 10-core processors, 192GB 2400MHz RAM for all host blades
- 1 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM for all host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switches
- 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

- Cisco UCS firmware 4.0(4e)
- NetApp ONTAP 9.6P4
- VMware ESXi 6.7 Update 2 for host blades
- VMware Horizon 7.10 VDI Desktops and RDSH Desktops
- Microsoft SQL Server 2016 SP1
- Microsoft Windows 10 64 bit (1809), 2vCPU, 3 GB RAM, 32 GB vDisk (master)
- Microsoft Windows Server 2019 (1809), 10vCPU, 32GB RAM, 40 GB vDisk (master)
- Microsoft Office 2016
- Login VSI 4.1.25 Knowledge Worker Workload (Benchmark Mode)
- Cisco UCS Configuration for Cluster Testing

Cisco UCS test configuration for cluster scale testing

This test case validates three workload clusters of ten blades using VMware Horizon 7.10 with 2240 RDS sessions, 1800 VDI non-persistent sessions, and 1800 VDI persistent sessions. Server N+1 fault tolerance is factored into this test scenario for each workload and infrastructure cluster.

Figure 71. RDS Cluster Test Configuration with 10 blades

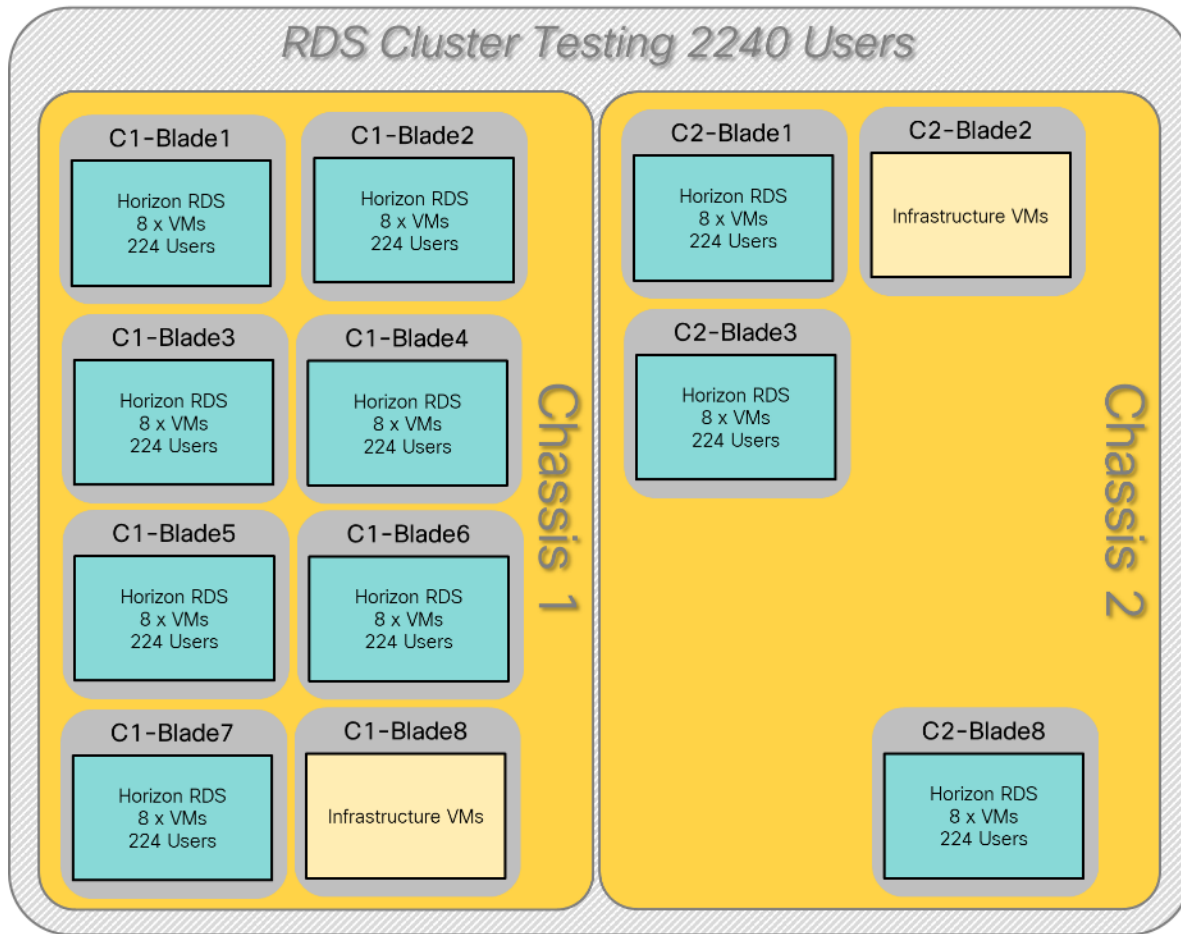


Figure 72. VDI Persistent Cluster Test Configuration with 10 blades

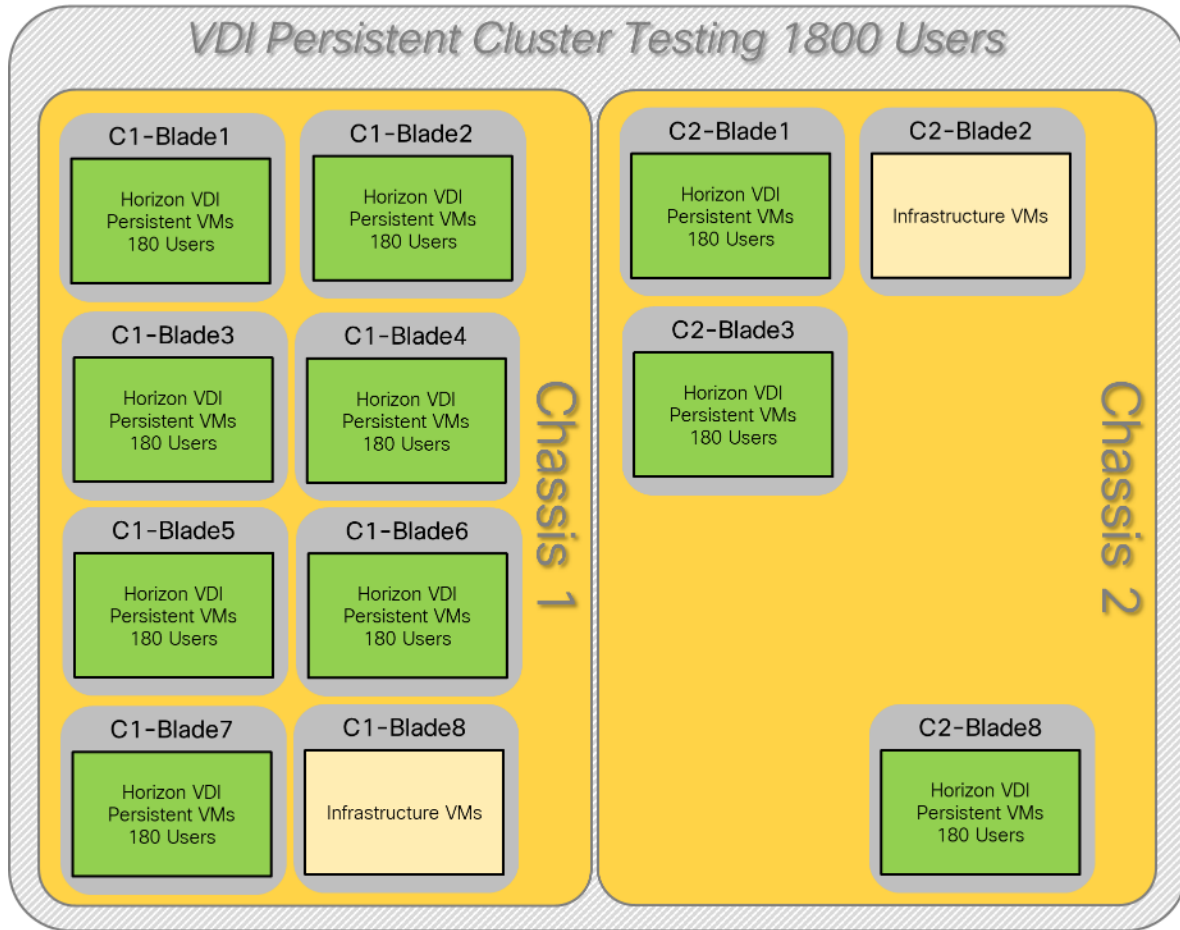
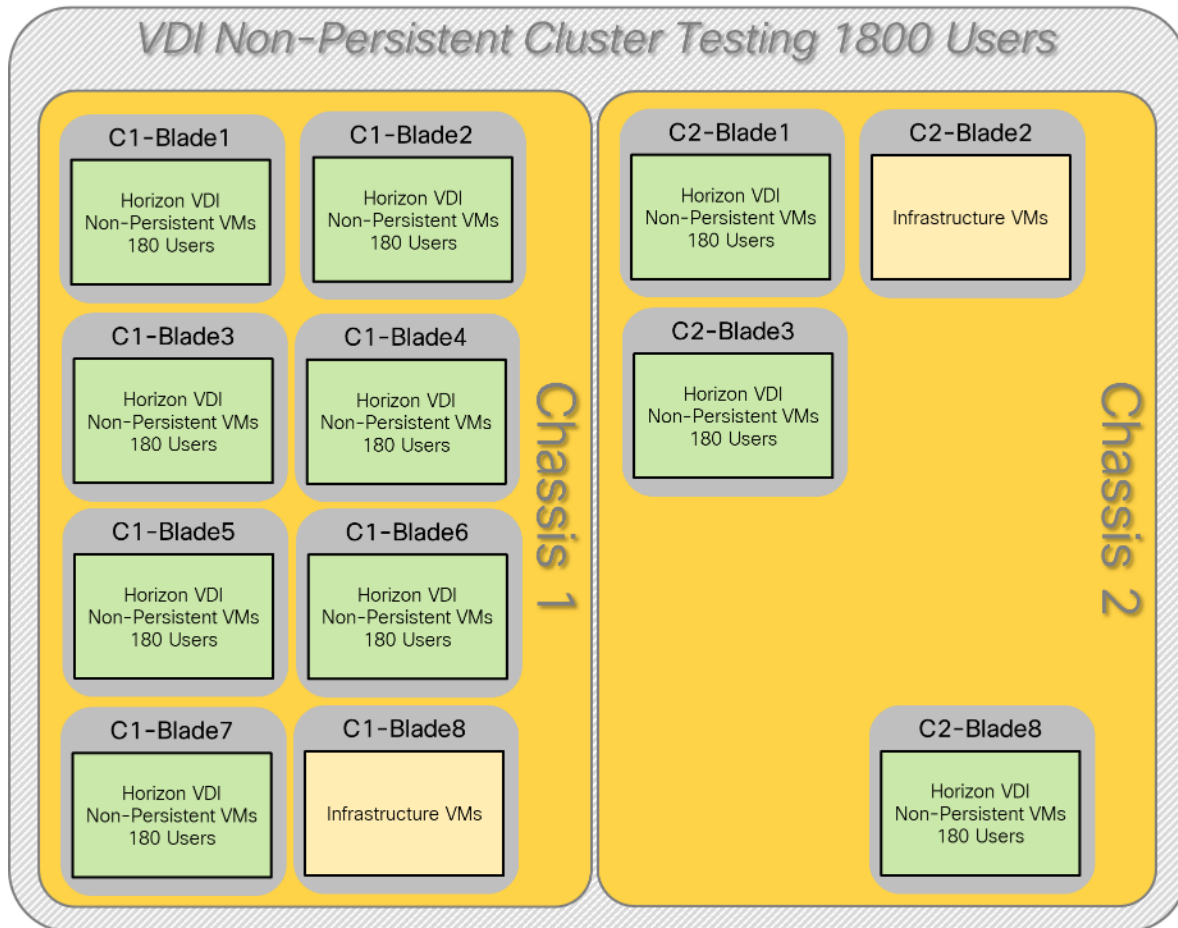


Figure 73. VDI Non-Persistent Cluster Test Configuration with 10 blades



Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M5 Blade Servers with Intel Xeon Silver 4114 2.20-GHz 10-core processors, 192GB 2400MHz RAM for all host blades
- 10 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM for all host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switches
- 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

-
- Cisco UCS firmware 4.0(4e)
 - NetApp ONTAP 9.6P4
 - VMware ESXi 6.7 Update 2 for host blades
 - VMware Horizon 7.10 VDI Desktops and RDSH Desktops
 - Microsoft SQL Server 2016 SP1
 - Microsoft Windows 10 64 bit (1809), 2vCPU, 3 GB RAM, 32 GB vDisk (master)
 - Microsoft Windows Server 2019 (1809), 10vCPU, 32GB RAM, 40 GB vDisk (master)
 - Microsoft Office 2016
 - Login VSI 4.1.25 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS configuration for full-scale testing

These test cases validate 30 blades of three distinct workloads and a mixed workload using VMware Horizon 7.10 with:

- 6700 Non-Persistent RDS sessions (Linked clones).
- 5400 Persistent VDI sessions (Full clones).
- 5400 Non-Persistent VDI sessions (Instant clones).
- 2240 RDS sessions, 1800 VDI Non-Persistent sessions, and 1800 VDI Persistent sessions for a total sum of 5800 users

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

Figure 74. Full-scale Test Configuration 6700 Non-Persistent RDS Sessions with 30 blades

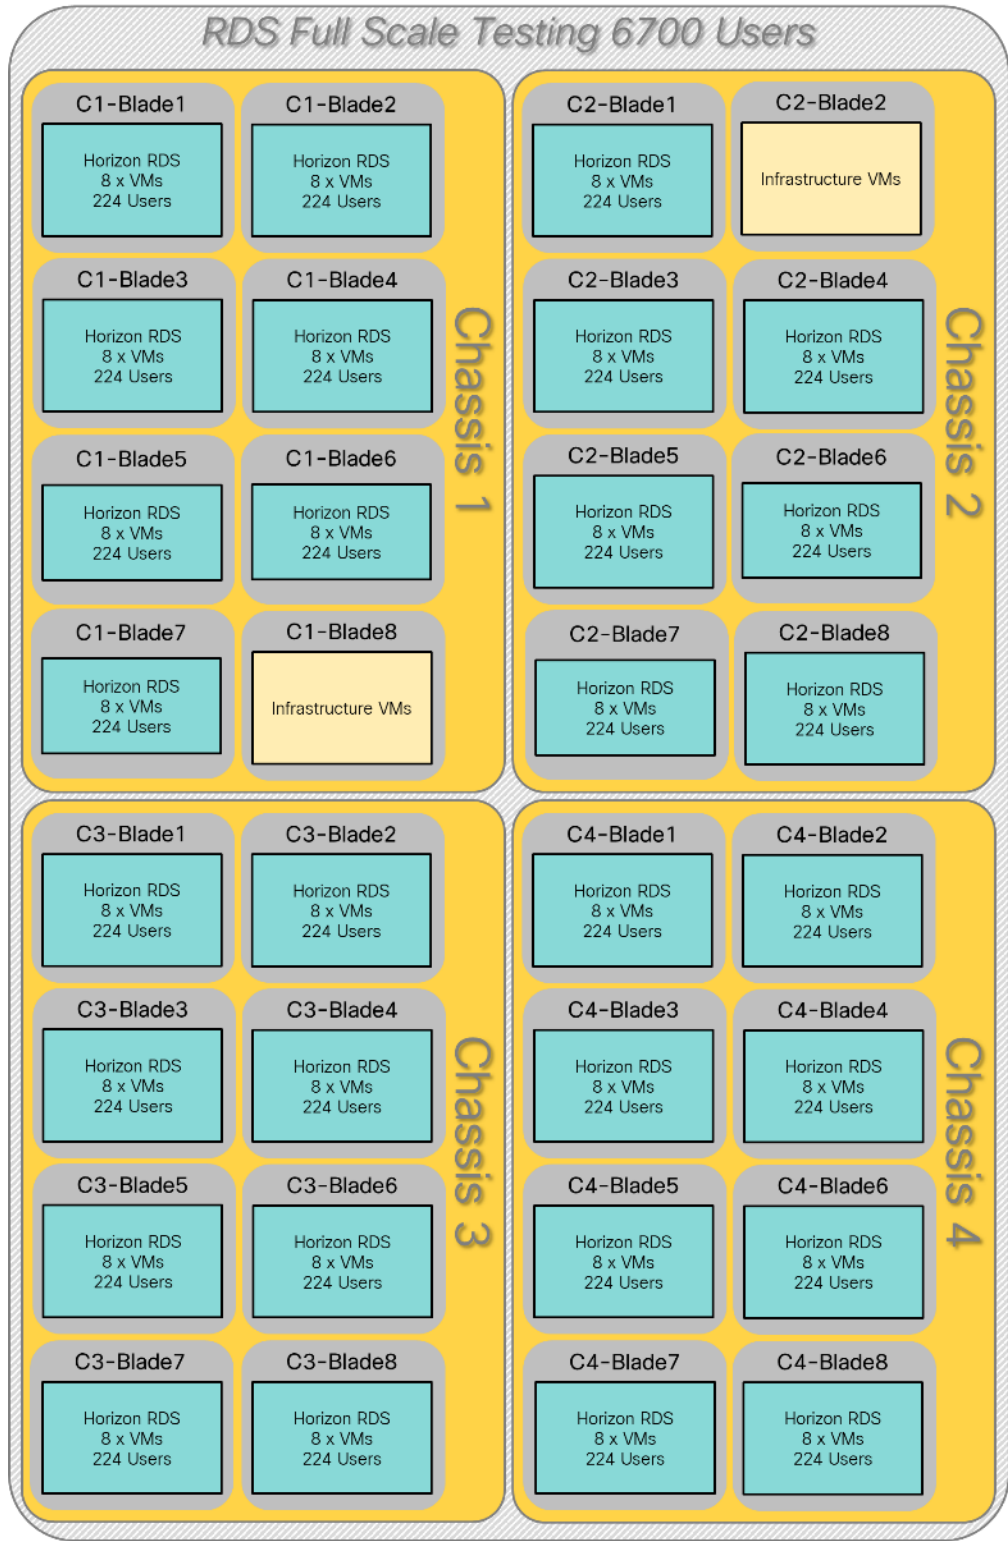


Figure 75. Full-scale Test Configuration 5400 Persistent VDI Sessions with 30 blades

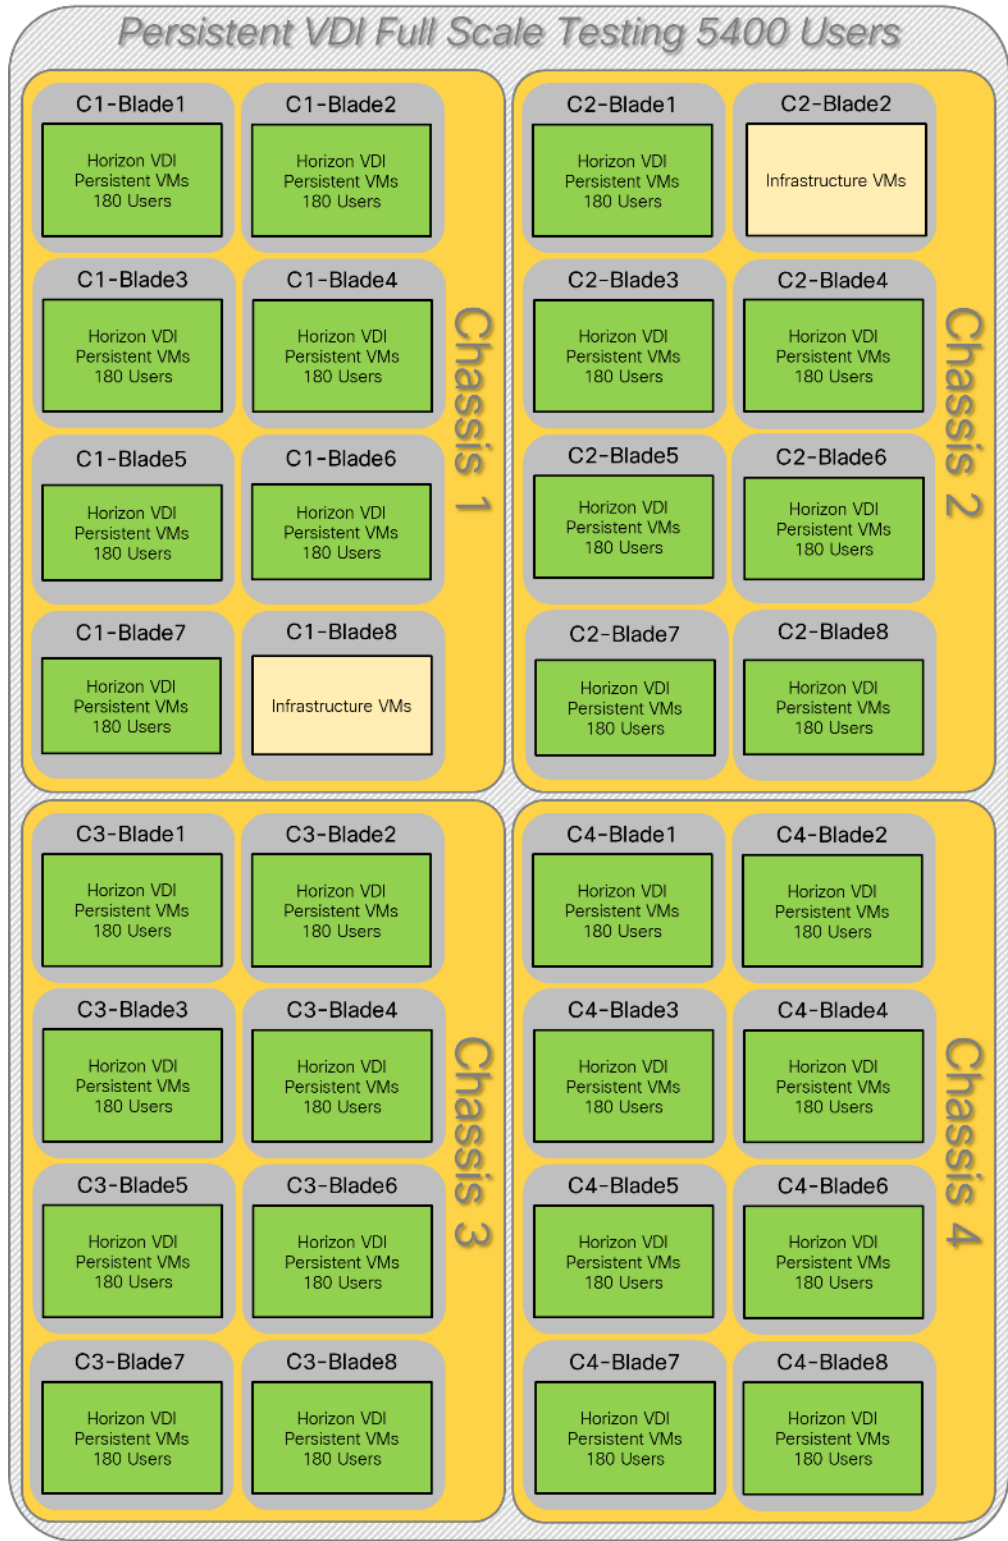


Figure 76. Full-scale Test Configuration 5400 Non-Persistent VDI Sessions with 30 blades

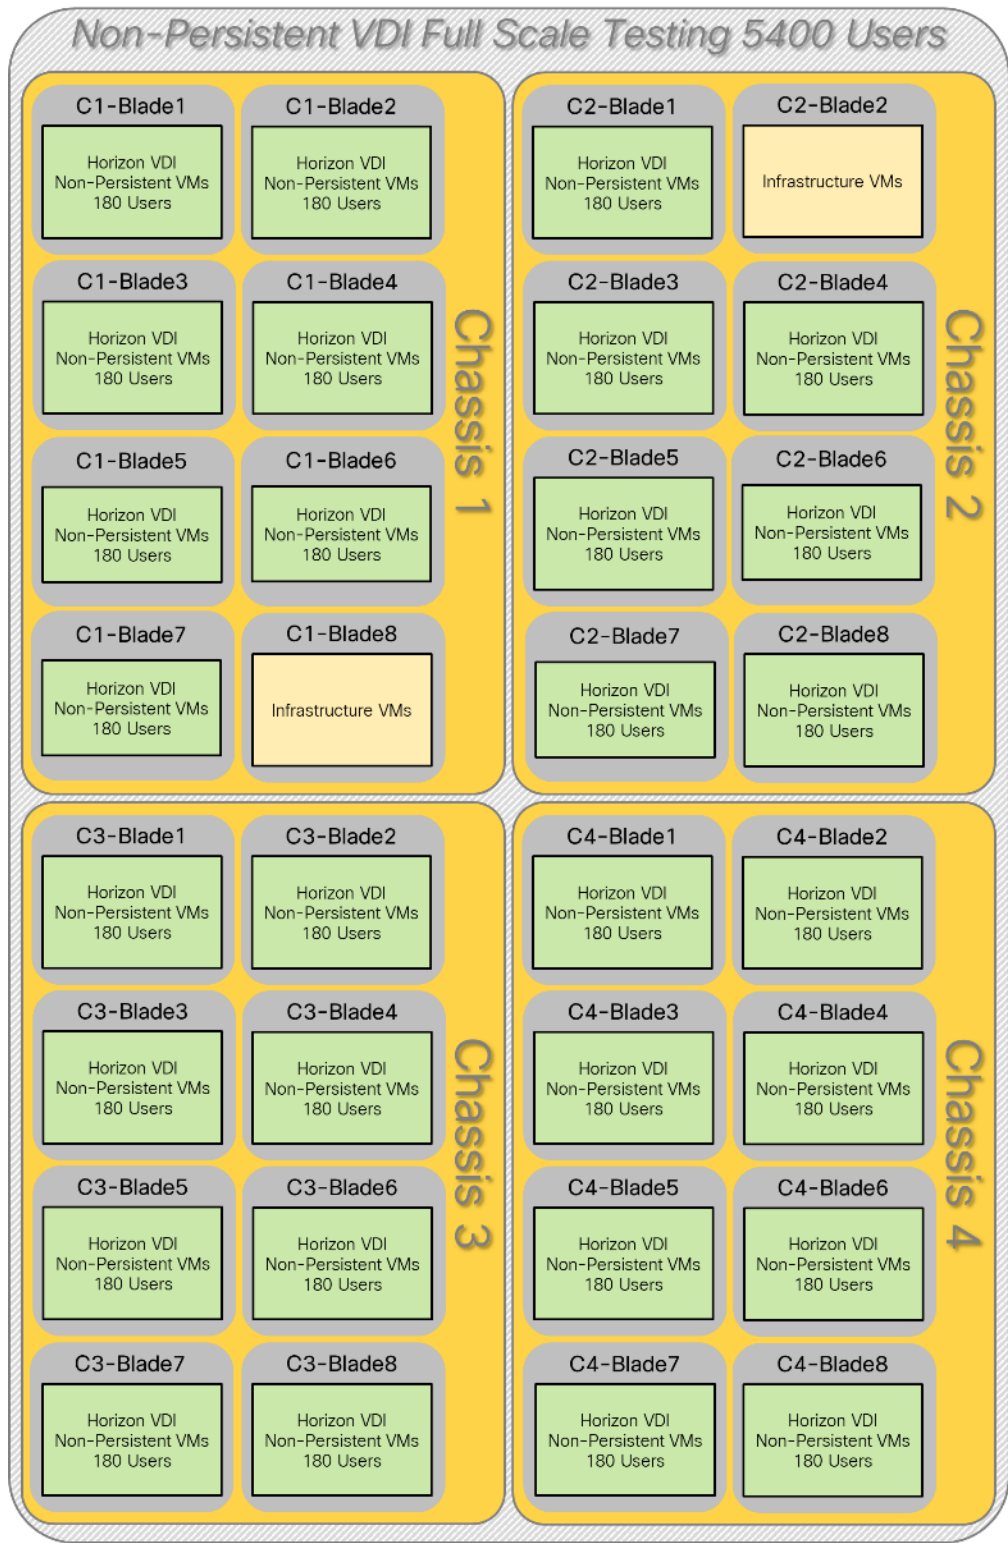
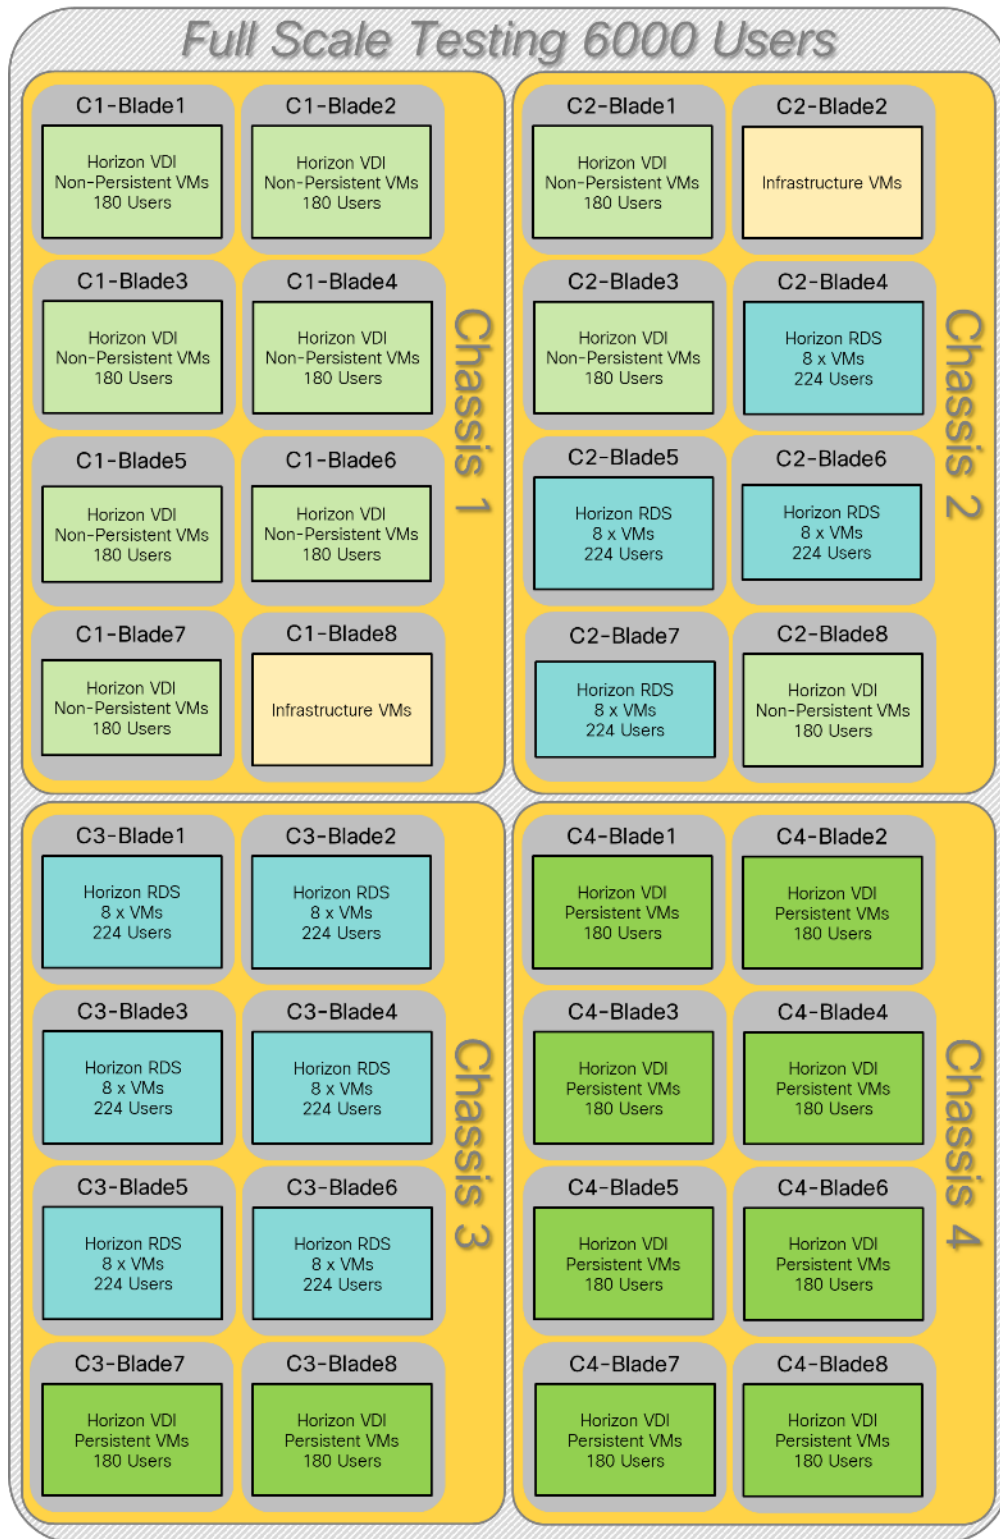


Figure 77. Full-scale Test Configuration Mixed with 30 blades



Hardware components:

-
- Cisco UCS 5108 Blade Server Chassis
 - 2 Cisco UCS 6454 Fabric Interconnects
 - 2 (Infrastructure Hosts) Cisco UCS B200 M5 Blade Servers with Intel Xeon Silver 4114 2.20-GHz 10-core processors, 192GB 2400MHz RAM for all host blades
 - 30 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM for all host blades
 - Cisco VIC 1340 CNA (1 per blade)
 - 2 Cisco Nexus 93180YC-FX Access Switches
 - 2 Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switches
 - 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

- Cisco UCS firmware 4.0(4e)
- NetApp ONTAP 9.6P4
- VMware ESXi 6.7 Update 2 for host blades
- VMware Horizon 7.10 VDI Desktops and RDSH Desktops
- Microsoft SQL Server 2016 SP1
- Microsoft Windows 10 64 bit (1809), 2vCPU, 3 GB RAM, 32 GB vDisk (master)
- Microsoft Windows Server 2019 (1809), 10vCPU, 32GB RAM, 40 GB vDisk (master)
- Microsoft Office 2016
- Login VSI 4.1.25 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 Servers 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

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 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.

All VMware ESXi VDI host blades to be tested were restarted prior to each test cycle.

Procedure 1. 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.

In addition, Cisco requires that the Login VSI Benchmark method is 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.

Step 1. Time 0:00:00 Start PerfMon/Esxstop/XenServer Logging on the following systems:

- Infrastructure and VDI Host Blades used in the test run
- SCVMM/vCenter used in the test run
- All Infrastructure virtual machines used in test run (AD, SQL, brokers, image mgmt., etc.)

Step 2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.

Step 3. Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using XenDesktop Studio or View Connection server.

Note

The boot rate should be around 10-12 virtual machines per minute per server.

Step 4. Time 0:06 First machines boot.

Step 5. Time 0:30 Single Server or Scale target number of desktop virtual machines booted on 1 or more blades.

Note

No more than 30 minutes for boot up of all virtual desktops is allowed.

Step 6. Time 0:35 Single Server or Scale target number of desktop virtual machines desktops registered on XD Studio or available on View Connection Server.

Step 7. Virtual machine settling time.

Note

No more than 60 Minutes of rest time is allowed after the last desktop is registered on the XD Studio or available in View Connection Server dashboard. Typically, a 30-40 minute rest period is sufficient.

Step 8. Time 1:35 Start Login VSI 4.1.x Office Worker Benchmark Mode Test, setting auto-logout time at 900 seconds, with Single Server or Scale target number of desktop virtual machines utilizing a sufficient number of Launchers (at 20-25 sessions/Launcher).

Step 9. Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48 minute bench-mark launch rate).

Step 10. Time 2:25 All launched sessions must become active.

Note

All sessions launched must become active for a valid test run within this window.

Step 11. Time 2:40 Login VSI Test Ends (based on Auto Logout 900 Second period designated above).

Step 12. Time 2:55 All active sessions logged off.

Step 13. Time 2:57 All logging terminated; Test complete.

Step 14. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines.

Step 15. Time 3:30 Reboot all hypervisor hosts.

Step 16. 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 the 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 Citrix Desktop Studio be monitored throughout the steady state to 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 described above and will note that we did not reach a VSImax dynamic in our testing.

The purpose of this testing is to provide the data needed to validate VMware Horizon 7.10 Remote Desktop Service Hosts (RDSH) linked clones and VMware Horizon 7.10 Virtual Desktop (VDI) randomly assigned, non-persistent instant clones and VMware Horizon 7.10 Virtual Desktop (VDI) statically assigned, persistent full clones models using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2019 sessions on Cisco UCS B200 M5 Blade Servers using a NetApp AFF300 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 Citrix products with VMware vSphere.

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 SBC 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 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 VSImax. When the system is reaching its saturation point, re-

response times will rise. When reviewing the average response time, you will see 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, it was decided to execute the scripts completely on the server side. This is the only practical and platform-independent solution for a benchmark like Login VSI.

Calculating VSI_{max} 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 VSI_{max}.

The five 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 effect considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS, the application, print, GDI, and so on. These operations are specifically short by nature. When such operations become consistently long, the system is saturated due to excessive queuing on the 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 78. Sample of a VSI max response time graph, representing a normal test



Figure 79. 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 very 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 time 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 very 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, etc.) 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:

- 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 amount of “active” sessions. For ex-ample, 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.

VSImax v4.1.x is reached when the VSIBase + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSImax 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 VSImax 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 VSImax v4.1.x, the performance of the system is not decided by the total average response time, but by the latency is 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, VSImax is not hit and the number of sessions ran successfully. This approach is fundamentally different in comparison to previous VSImax methods, as it was always required to saturate the system beyond VSImax threshold.

VSImax v4.1.x is now always reported with the average baseline VSI response time result. For example: “The VSImax 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. VSImax indicates what the total user capacity is for the system. These two are not automatically connected and related.

When a server with a very 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 VSImax v4.1.x, and the higher VSImax is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSImax method is introduced: VSImax v4.1.x. This methodology gives much better in-sight into system performance and scales to extremely large systems.

Single-server recommended maximum workload

For both the Vmware Horizon 7.10 Remote Desktop Service Hosts (RDSH) and Vmware Horizon 7.10 Virtual Desk-top (VDI) 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%.

Note
Memory should never be oversubscribed for Desktop Virtualization workloads.

Figure 80. Phases of test runs

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time

Test Phase	Description
Idle	The rest time after the last desktop is registered on the Connection Server. (typically, a 30-40 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


Single-server recommended maximum workload testing

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 three tests: 224 RDS sessions, 180 VDI non-persistent sessions, and 180 VDI persistent sessions.

Single-server recommended maximum workload for RDS with 224 users

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM is 224 RDS Windows Server 2019 sessions. Each dedicated blade server ran 8 Windows Server 2019 Virtual Machines. Each virtual server was configured with 10 vCPUs and 32GB RAM.

Figure 81. Single-server recommended maximum workload | Horizon 7.10 RDS | VSI score



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

Figure 82. Single-server recommended maximum workload | Horizon 7.10 RDS | Host CPU utilization




Figure 83. Single-server recommended maximum workload | Horizon 7.10 RDS | Host memory utilization




Figure 84. Single-server | Horizon 7.10 RDS | Host network utilization





Figure 85. Single-server recommended maximum workload | Horizon 7.10 RDS | Virtual machine CPU utilization




Single-server recommended maximum workload for VDI non-persistent with 180 users

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM is 180 Windows 10 64-bit VDI non-persistent instant clone virtual machines with 2 vCPU and 3GB RAM.

Login VSI performance data is as follows:

Figure 86. Single-server | VMware Horizon 7.10 VDI-NP | VSI score



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

Figure 87. Single-server | VMware Horizon 7.10 VDI-NP | Host CPU utilization




Figure 88. Single-server | Vmware Horizon 7.10 VDI-NP | Host memory utilization





Figure 89. Single-server | Vmware Horizon 7.10 VDI-NP | Host network utilization




Single-server recommended maximum workload for VDI persistent with 180 Users

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6230 2.10-GHz 20-core processors, 768GB 2933MHz RAM is 180 Windows 10 64-bit VDI persistent virtual machines with 2 vCPU and 3GB RAM.

Login VSI performance data is as follows:

Figure 90. Single-server | VMware Horizon 7.10 VDI-P | VSI score



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

Figure 91. Single-server | Vmware Horizon 7.10 VDI-P | Host CPU utilization




Figure 92. Single-server | Vmware Horizon 7.10 VDI-P | Host memory utilization





Figure 93. Single-server | Vmware Horizon 7.10 VDI-P | Host network utilization



Cluster recommended maximum workload testing

This section details the key performance metrics captured on the Cisco UCS host blades during the cluster testing to determine the per host server workload in the N+1 environment. The cluster testing comprised of three tests: 2240 RDS sessions, 1800 VDI non-persistent sessions, and 1800 VDI persistent sessions.

Cluster workload testing with 2240 RDS users

This section details the key performance metrics captured on the Cisco UCS, NetApp array, and Infra-structure virtual machines during the non-persistent desktop testing. The cluster testing was comprised of 2240 RDS sessions using 10 workload blades.

The workload for the test is 2240 RDS users. To achieve the target, sessions were launched against all workload clusters 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 94. Eight node cluster | 2240 RDS users | VSI score




Figure 95. Eight node cluster | 2240 RDS users | VSI repeatability




Figure 96. Cluster | 2240 RDS users | 8 RDS hosts | Host CPU utilization




Figure 97. Cluster | 2240 RDS users | RDS hosts | Host memory utilization




Figure 98. Cluster | 2240 RDS users | RDS hosts | Host system uplink network utilization




Figure 99. Cluster | 2240 RDS users | RDS hosts | NetApp AF A300 utilization | Total throughput




Figure 100. Cluster | 2240 RDS users | RDS hosts | NetApp AFF A300 utilization | Total IOPs





Figure 101. Cluster | 2240 RDS users | RDS hosts | NetApp AF A300 utilization | Latency



Cluster workload testing with 1800 non-persistent desktop users

This section details the key performance metrics that were captured on the Cisco UCS, NetApp array, and Infrastructure virtual machines during the non-persistent desktop testing. The cluster testing comprised of 1800 VDI non-persistent desktop sessions using 10 workload blades.

The workload for the test is 1800 VDI non-persistent desktop users. To achieve the target, sessions were launched against all workload clusters 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 102. Cluster | 1800 VDI-NP users | VSI score




Figure 103. Cluster | 1800 VDI-NP users | VSI repeatability




Figure 104. Cluster | 1800 VDI-NP users | Non-persistent hosts | Host CPU utilization




Figure 105. Cluster | 1800 VDI-NP users | Non-persistent hosts | Host memory utilization




Figure 106. Cluster | 1800 VDI-NP users | Non-persistent hosts | Host network utilization




Figure 107. Cluster | 1800 VDI-NP users | Non-persistent hosts | NetApp AFF A300 utilization | Total throughput




Figure 108. Cluster | 1800 VDI-NP users | Non-persistent hosts | NetApp AFF A300 utilization | Total IOPs





Figure 109. Cluster | 1800 VDI-NP users | Non-persistent hosts | NetApp AFF A300 utilization | Latency



Cluster workload testing with 1800 persistent desktop users

This section details the key performance metrics that were captured on the Cisco UCS, NetApp array, and Infrastructure virtual machines during the persistent desktop testing. The cluster testing with comprised of 1800 VDI Persistent desktop sessions using 10 workload blades.

The workload for the test is 1800 VDI persistent desktop users. To achieve the target, sessions were launched against all workload clusters 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 110. Cluster | 1800 VDI-P users | VSI score




Figure 111. Cluster | 1800 VDI-P users | VSI repeatability

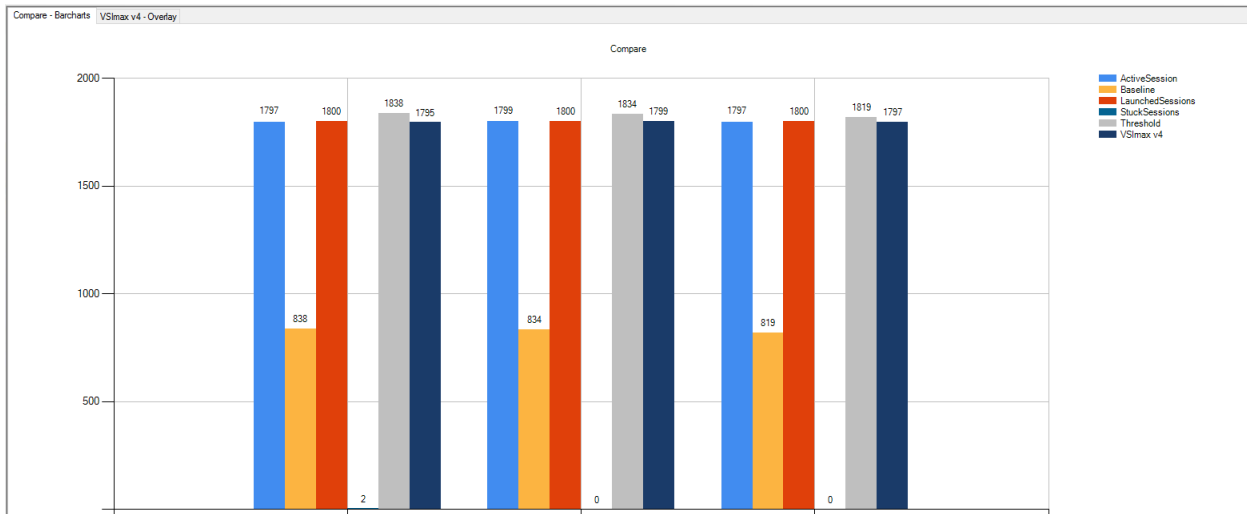


Figure 112. Cluster | 1800 VDI-P users | Persistent hosts | Host CPU utilization




Figure 113. Cluster | 1800 VDI-P users | Persistent hosts | Host memory utilization




Figure 114. Cluster | 1800 VDI-P users | Persistent hosts | Host network utilization




Figure 115. Cluster | 1800 VDI-P users | Persistent hosts | NetApp AFF A300 utilization | Total throughput




Figure 116. Cluster | 1800 VDI-P users | Persistent hosts | NetApp AFF A300 utilization | Total IOPs





Figure 117. Cluster | 1800 VDI-P users | Persistent hosts | NetApp AFF A300 utilization | Latency



Full-scale workload testing

Full-scale workload testing with 6700 RDS users

This section details the key performance metrics that were captured on the Cisco UCS, during the RDS full-scale testing with 6700 desktop sessions using 30 blades.

The RDS workload for the solution is 6700 users. To achieve the target, sessions were launched against all workload clusters 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 118. Full-scale | 6700 RDS users | VSI score




Figure 119. Full-scale | 6700 RDS users | VSI repeatability




Figure 120. Full-scale | 6700 RDS users | RDS hosts | Host CPU utilization




Figure 121. Full-scale | 6700 RDS users | RDS hosts | Host memory utilization




Figure 122. Full-scale | 6700 RDS users | RDS hosts | Host network utilization




Figure 123. Full-scale | 6700 RDS users | Persistent hosts | NetApp AFF A300 utilization | Total throughput




Figure 124. Full-scale | 6700 RDS users | Persistent hosts | NetApp AFF A300 utilization | Total IOPs





Figure 125. Full-scale | 6700 RDS users | Persistent hosts | NetApp AFF A300 utilization | Latency



Full-scale workload testing with 5400 VDI persistent users

This section details the key performance metrics that were captured on the Cisco UCS and NetApp array during the persistent desktop full-scale testing with 5400 VDI persistent desktops using 30 blades.

The workload for the test is 5400 persistent VDI users. To achieve the target, sessions were launched against all workload clusters 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 126. Full-scale | 5400 VDI-P users | VSI score




Figure 127. Full-scale | 5400 VDI-P users | VSI repeatability




Figure 128. Full-scale | 5400 VDI-P users | Host CPU utilization




Figure 129. Full-scale | 5400 VDI-P users | Host memory utilization




Figure 130. Full-scale | 5400 VDI-P users | Host network utilization




Figure 131. Full-scale | 5400 VDI-P users | NetApp AFF A300 | Throughput




Figure 132. Full-scale | 5400 VDI-P users | NetApp AFF A300 | IOPs

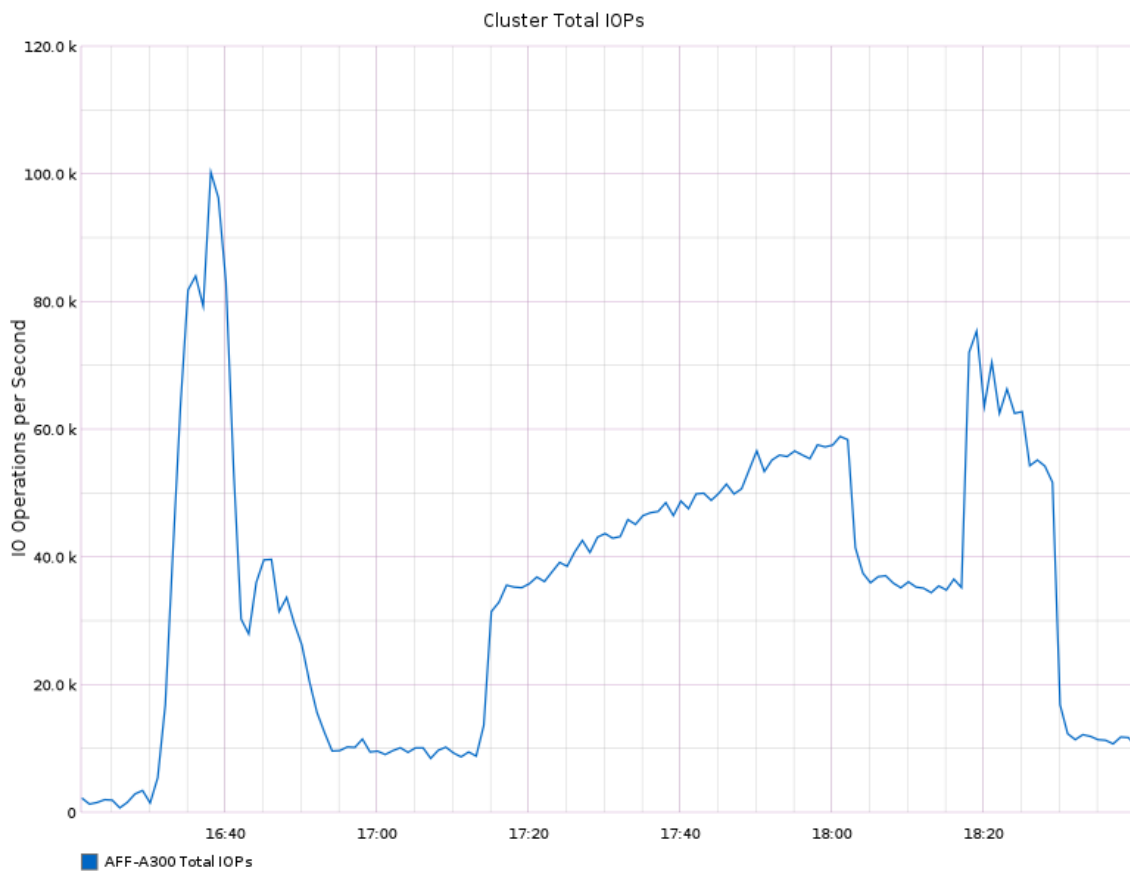



Figure 133. Full-scale | 5400 VDI-P users | NetApp AFF A300 | Latency



Full-scale workload testing with 5400 VDI non-persistent desktop users

This section describes the key performance metrics that were captured on the Cisco UCS, NetApp array, and Infrastructure virtual machines during the non-persistent desktop testing with 5400 VDI non-persistent desktops using 30 workload blades.

The workload for the test is 2400 VDI non-persistent desktop users. To achieve the target, sessions were launched against all workload clusters 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 134. Full-scale | 5400 VDI-NP users | VSI score




Figure 135. Full-scale | 5400 VDI-NP users | VSI repeatability




Figure 136. Full-scale | 5400 VDI-NP users | Non-persistent hosts | Host CPU utilization




Figure 137. Full-scale | 5400 VDI-NP users | Non-Persistent hosts | Host memory utilization




Figure 138. Full-scale | 5400 VDI-NP users | Non-persistent hosts | Host network utilization




Figure 139. Full-scale | 5400 VDI-NP users | Non-persistent hosts | NetApp AFF A300




Figure 140. Full-scale | 5400 VDI-NP users | Non-persistent hosts | NetApp AFF A300





Figure 141. Full-scale | 5400 VDI-NP users | Non-persistent hosts | NetApp AFF A300



Full-scale workload testing with 5800 mixed desktop users

This section describes the key performance metrics that were captured on the Cisco UCS and NetApp array during the mixed desktop users full-scale testing. The full-scale testing with 5800 users comprised of: 2240 RDS desktop sessions using 10 blades, 1800 VDI instant clones non-persistent sessions using 10 blades, and 1800 VDI full clones persistent sessions using 10 blades.

The combined mixed workload for the solution is 5800 users. To achieve the target, sessions were launched against all workload clusters 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 142. Full-scale | 5800 mixed users | VSI score




Figure 143. Full-scale | 5800 mixed users | VSI repeatability




Figure 144. Full-scale | 5800 mixed users | Non-persistent hosts | Host CPU utilization




Figure 145. Full-scale | 5800 mixed users | Non-persistent hosts | Host memory utilization




Figure 146. Full-scale | 5800 mixed users | Non-persistent hosts | Host network utilization




Figure 147. Full-scale | 5800 mixed users | RDS hosts | Host CPU utilization




Figure 148. Full-scale | 5800 mixed users | RDS hosts | Host memory utilization




Figure 149. Full-scale | 5800 mixed users | RDS hosts | Host network utilization




Figure 150. Full-scale | 5800 mixed users | Persistent hosts | Host CPU utilization




Figure 151. Full-scale | 5800 mixed users | Persistent hosts | Host memory utilization




Figure 152. Full-scale | 5800 mixed users | Persistent hosts | Host network utilization




Figure 153. Full-scale | 5800 mixed users | Mixed hosts | NetApp AFF A300 | Throughput




Figure 154. Full-scale | 5800 mixed users | Mixed hosts | NetApp AFF A300 | IOPs





Figure 155. Full-scale | 5800 mixed users | Mixed hosts | NetApp AFF A300 | Latency



Adaptive Quality of Service validation

AQoS adjusts the QoS ceiling based on allocated or used space, as defined in the policy. This feature is useful to virtual desktop environments because as the number of virtual desktops on the given datastore increases, the throughput adjusts correctly. To validate this feature, we created the following AQoS policy by using this CLI:

```

gos adaptive-policy-group create -policy-group VDI -expected-iops 585/TB -peak-iops
650/TB -expected-iops-allocation used-space -peak-iops-allocation used-space -absolute-
min-iops 100
    
```


Policy Group	Expected IOPS per TB (Floor)	Peak IOPS per TB (Ceiling)	Absolute Minimum IOPS	Allocation for Expected IOPS and Peak IOPS
VDI	585	650	100	used-space

We assigned the AQoS created to volumes `vdi_nfs_ds01` to `vdi_nfs_ds08`. We used IOMeter to generate IOPS on the target volumes.

We targeted two volumes: `vdi_nfs_ds01` and `vdi_nfs_ds04`.

Volume	Used Space	Expected IOPS	Peak IOPS
Vdi_nfs_ds01	22.1	12,929	14,365
Vdi_nfs_ds04	4.01	2,346	2,607

The same workload was used to generate the load against those targets. We can clearly see that the volume with more used space has a higher ceiling.



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 AFF A-Series storage systems. FlexPod is designed and validated by 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 wanting to deploy an enterprise-class VDI solution for 5,000 to 6,000 users at one time.

Get more business value with services

Whether you are planning your next-generation environment, need specialized know-how for a major deployment, or want to get the most from your current storage, Cisco Advanced Services, NetApp AFF A300 storage system, and our certified partners can help. We collaborate with you to enhance your IT capabilities through a full portfolio of services that covers your IT lifecycle with:

- Strategy services to align IT with your business goals
- Design services to architect your best storage environment
- Deploy and transition services to implement validated architectures and prepare your storage environment
- Operations services to deliver continuous operations while driving operational excellence and efficiency

In addition, Cisco Advanced Services and NetApp Support provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

Appendix

Ethernet network configuration

The following section provides a detailed procedure for configuring the Cisco Nexus 9000 Switches used in this study.

Cisco Nexus 93180YC-A configuration

```
version 7.0(3)I7(2)
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
feature nxapi
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
```

```
feature dhcp
feature vpc
feature lldp
clock timezone PST -7 0

no password strength-check
username admin password 5 $1$tYYajkfc$7P7nLjWYvfTWAlvFDnwJZ. role network-admin
ip domain-lookup
system default switchport
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 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 0xa075be936e36177e1912888e7aed3223 priv
0xa075be936e36177e1912888e7aed3223 localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 10.10.60.2 use-vrf default
ntp peer 10.10.60.3 use-vrf default
ntp server 10.10.160.2 use-vrf default
ntp peer 10.10.160.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8

vlan 1-2,60-70,102
vlan 60
  name In-Band-Mgmt
```

```
vlan 61
  name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name NFS
vlan 66
  name vMotion
vlan 68
  name LauncherPXE
vlan 69
  name Launcher81
vlan 102
  name VDI

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 14port
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 shutdown
  no ip redirects
  ip address 10.29.164.2/24
  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
  preempt
  ip 10.10.61.1

interface Vlan62
  description CIFS vlan 62
  no shutdown
  no ip redirects
  ip address 10.10.62.2/24
  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 10.10.63.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 63
    preempt
    ip 10.10.63.1
```

```
interface Vlan66
  description vMotion network vlan 66
  no shutdown
  no ip redirects
  ip address 10.10.66.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1
```

```
interface Vlan68
  description LoginVSI Launchers vlan 68
  no shutdown
  no ip redirects
  ip address 10.10.68.2/23
  no ipv6 redirects
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
  ip dhcp relay address 10.10.61.30
```

```
interface Vlan69
  description LoginVSI Launchers 10.10.81-network vlan 69
  no shutdown
  no ip redirects
```

```
ip address 10.10.81.2/24
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 10.2.0.2/19
    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
    ip dhcp relay address 10.10.61.31

interface port-channel10
    description VPC-PeerLink
    switchport mode trunk
    switchport trunk allowed vlan 1-2,60-70,102
    spanning-tree port type network
    vpc peer-link

interface port-channel15
    description FI-A_6k_Launchers-Uplink
    switchport mode trunk
    switchport trunk allowed vlan 1-2,60-70,102
    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
spanning-tree port type edge trunk
mtu 9216
vpc 16
```

```
interface port-channel51
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
vpc 51
```

```
interface port-channel52
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
vpc 52
```

```
interface port-channel53
switchport mode trunk
switchport trunk allowed vlan 1-163,165-263,265-4094
spanning-tree port type edge trunk
mtu 9216
vpc 53
```

```
interface port-channel54
switchport mode trunk
switchport trunk allowed vlan 1-163,165-263,265-4094
spanning-tree port type edge trunk
mtu 9216
vpc 54
```

```
interface Ethernet1/1
```

```
interface Ethernet1/2
```

```
interface Ethernet1/3
```

```
interface Ethernet1/4
```

```
interface Ethernet1/5
```

```
interface Ethernet1/6
```

```
interface Ethernet1/7
```

```
interface Ethernet1/8
```

```
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
```

```
interface Ethernet1/18
```

```
interface Ethernet1/19
```

```
interface Ethernet1/20
```

```
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
  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
  mtu 9216
  channel-group 16 mode active

interface Ethernet1/47
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  mtu 9216
  channel-group 15 mode active

interface Ethernet1/48
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  mtu 9216
  channel-group 16 mode active

interface Ethernet1/49
  description VPC Peer Link between 9ks
```

```
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
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
channel-group 10 mode active
```

```
interface Ethernet1/51
description FI-A-N9K-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
mtu 9216
speed 40000
no negotiate auto
channel-group 51 mode active
```

```
interface Ethernet1/52
description FI-B-N9K-Uplink
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
channel-group 52 mode active
```

```
interface Ethernet1/53
switchport mode trunk
switchport trunk allowed vlan 1-163,165-263,265-4094
mtu 9216
channel-group 53 mode active
```

```
interface Ethernet1/54
switchport mode trunk
switchport trunk allowed vlan 1-163,165-263,265-4094
mtu 9216
```

```
channel-group 54 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.65/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
no system default switchport shutdown
```

Cisco Nexus 93180YC -B configuration

```
version 7.0(3)I7(2)
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
feature nxapi
cfs ipv4 distribute
cfs eth distribute
feature udd
```

```
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock timezone PST -7 0

no password strength-check
username admin password 5 $1$fp3LrGLC$PF8eML85qkPBgdH/bZAKK/ role network-admin
ip domain-lookup
system default switchport
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 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 0x142e177306873a75257c9a8388b47fb7 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 peer 10.10.60.2 use-vrf default
ntp server 10.10.60.3 use-vrf default
ntp peer 10.10.160.2 use-vrf default
ntp server 10.10.160.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8
```

```
vlan 1-2,60-70,102
```

```
vlan 60
```

```
    name In-Band-Mgmt
```

```
vlan 61
```

```
    name Infra-Mgmt
```

```
vlan 62
```

```
    name CIFS
```

```
vlan 63
```

```
    name NFS
```

```
vlan 66
```

```
    name vMotion
```

```
vlan 68
```

```
    name LauncherPXE
```

```
vlan 69
```

```
    name Launcher81
```

```
vlan 70
```

```
    name other-3
```

```
vlan 102
```

```
    name VDI
```

```
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 shutdown
no ip redirects
ip address 10.29.164.3/24
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
  mtu 9216
  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 Vlan66
  description vMotion network vlan 66
  no shutdown
  no ip redirects
  ip address 10.10.66.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1
```

```
interface Vlan68
  description LoginVSI Launchers vlan 68
  no shutdown
  no ip redirects
  ip address 10.10.68.3/23
  no ipv6 redirects
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
```

```
ip dhcp relay address 10.10.61.30
```

```
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
  ip dhcp relay address 10.10.61.31
```

```
interface port-channel10
  description VPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  spanning-tree port type network
  vpc peer-link
```

```
interface port-channel15
  description FI-A_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
```

```
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
spanning-tree port type edge trunk
mtu 9216
vpc 16
```

```
interface port-channel51
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
vpc 51
```

```
interface port-channel52
switchport mode trunk
switchport trunk allowed vlan 1-2,60-70,102
spanning-tree port type edge trunk
mtu 9216
speed 40000
no negotiate auto
vpc 52
```

```
interface port-channel53
switchport mode trunk
switchport trunk allowed vlan 1-163,165-263,265-4094
spanning-tree port type edge trunk
mtu 9216
vpc 53
```

```
interface port-channel54
switchport mode trunk
```

```
switchport trunk allowed vlan 1-163,165-263,265-4094
spanning-tree port type edge trunk
mtu 9216
vpc 54
```

```
interface Ethernet1/1
```

```
interface Ethernet1/2
```

```
interface Ethernet1/3
```

```
interface Ethernet1/4
```

```
interface Ethernet1/5
```

```
interface Ethernet1/6
```

```
interface Ethernet1/7
```

```
interface Ethernet1/8
```

```
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

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

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
  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
  mtu 9216
  channel-group 16 mode active

interface Ethernet1/47
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  mtu 9216
  channel-group 15 mode active
```

```
interface Ethernet1/48
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  mtu 9216
  channel-group 16 mode active
```

```
interface Ethernet1/49
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  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
  channel-group 10 mode active
```

```
interface Ethernet1/51
  description FI-B-N9KUPLink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  mtu 9216
  speed 40000
  no negotiate auto
  channel-group 51 mode active
```

```
interface Ethernet1/52
  description FI-B-N9KUPLink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102
  spanning-tree port type edge trunk
  mtu 9216
  speed 40000
  no negotiate auto
  channel-group 52 mode active
```

```
interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  mtu 9216
  channel-group 53 mode active

interface Ethernet1/54
  switchport mode trunk
  switchport trunk allowed vlan 1-163,165-263,265-4094
  mtu 9216
  channel-group 54 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.66/24
  line console
  line vty
  boot nxos bootflash:/nxos.7.0.3.I7.2.bin
  no system default switchport shutdown
```

Fibre channel network configuration

The following section provides a detailed procedure for configuring the Cisco MDS 9100 Switches used in this study.

Cisco MDS 9132T-A configuration

```
version 8.3(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 $5$Pfdlzb5$OMcmDntTwIALadvV4BEPNQH0CfetmM9GCfoHpWeD905 role
network-admin
ip domain-lookup
```



```
ip host ADD16-MDS-A 10.29.164.238
```

```
aaa group server radius radius
```

```
snmp-server user admin network-admin auth md5 0x7dbf6d4613d0c91e60cc0f12b12ca4fd priv  
0x7dbf6d4613d0c91e60cc0f12b12ca4fd 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 10.81.254.131
```

```
vsan database
```

```
vsan 400 name "FlexPod-A"
```

```
device-alias database
```

```
device-alias name VDI-1-HBA1 pwn 20:00:00:25:b5:3a:00:3f
```

```
device-alias name VDI-2-HBA1 pwn 20:00:00:25:b5:3a:00:0f
```

```
device-alias name VDI-3-HBA1 pwn 20:00:00:25:b5:3a:00:1f
```

```
device-alias name VDI-4-HBA1 pwn 20:00:00:25:b5:3a:00:4e
```

```
device-alias name VDI-5-HBA1 pwn 20:00:00:25:b5:3a:00:2e
```

```
device-alias name VDI-6-HBA1 pwn 20:00:00:25:b5:3a:00:3e
```

```
device-alias name VDI-7-HBA1 pwn 20:00:00:25:b5:3a:00:0e
```

```
device-alias name VDI-9-HBA1 pwn 20:00:00:25:b5:3a:00:4d
```

```
device-alias name a300-01-0g pwn 20:01:00:a0:98:af:bd:e8
```

```
device-alias name a300-02-0g pwn 20:03:00:a0:98:af:bd:e8
```

```
device-alias name VDI-10-HBA1 pwn 20:00:00:25:b5:3a:00:2d
```

```
device-alias name VDI-11-HBA1 pwn 20:00:00:25:b5:3a:00:3d
```

```
device-alias name VDI-12-HBA1 pwn 20:00:00:25:b5:3a:00:0d
```

```
device-alias name VDI-13-HBA1 pwn 20:00:00:25:b5:3a:00:1d
```

```
device-alias name VDI-14-HBA1 pwn 20:00:00:25:b5:3a:00:4c
```

```
device-alias name VDI-15-HBA1 pwn 20:00:00:25:b5:3a:00:2c
```

```
device-alias name VDI-17-HBA1 pwn 20:00:00:25:b5:3a:00:0c
```

```
device-alias name VDI-18-HBA1 pwn 20:00:00:25:b5:3a:00:1c
```

```
device-alias name VDI-19-HBA1 pwn 20:00:00:25:b5:3a:00:4b
```

```
device-alias name VDI-20-HBA1 pwn 20:00:00:25:b5:3a:00:2b
```

```
device-alias name VDI-21-HBA1 pwwn 20:00:00:25:b5:3a:00:3b
device-alias name VDI-22-HBA1 pwwn 20:00:00:25:b5:3a:00:0b
device-alias name VDI-23-HBA1 pwwn 20:00:00:25:b5:3a:00:1b
device-alias name VDI-24-HBA1 pwwn 20:00:00:25:b5:3a:00:4a
device-alias name VDI-25-HBA1 pwwn 20:00:00:25:b5:3a:00:2a
device-alias name VDI-26-HBA1 pwwn 20:00:00:25:b5:3a:00:3a
device-alias name VDI-27-HBA1 pwwn 20:00:00:25:b5:3a:00:0a
device-alias name VDI-28-HBA1 pwwn 20:00:00:25:b5:3a:00:1a
device-alias name VDI-29-HBA1 pwwn 20:00:00:25:b5:3a:00:49
device-alias name VDI-30-HBA1 pwwn 20:00:00:25:b5:3a:00:39
device-alias name VDI-31-HBA1 pwwn 20:00:00:25:b5:3a:00:1e
device-alias name VDI-32-HBA1 pwwn 20:00:00:25:b5:3a:00:3c
device-alias name Infra01-8-HBA1 pwwn 20:00:00:25:b5:3a:00:4f
device-alias name Infra02-16-HBA1 pwwn 20:00:00:25:b5:3a:00:2f
device-alias commit
```

```
fcdomain fcid database
```

```
vsan 400 wwn 50:0a:09:84:80:13:41:27 fcid 0x680000 dynamic
vsan 400 wwn 50:0a:09:84:80:d3:67:d3 fcid 0x680100 dynamic
vsan 400 wwn 20:01:00:de:fb:90:a0:80 fcid 0x680200 dynamic
vsan 400 wwn 20:03:00:de:fb:90:a0:80 fcid 0x680300 dynamic
vsan 400 wwn 20:02:00:de:fb:90:a0:80 fcid 0x680400 dynamic
vsan 400 wwn 20:04:00:de:fb:90:a0:80 fcid 0x680500 dynamic
vsan 400 wwn 20:03:00:a0:98:af:bd:e8 fcid 0x680101 dynamic
!
    [a300-02-0g]
vsan 400 wwn 20:01:00:a0:98:af:bd:e8 fcid 0x680001 dynamic
!
    [a300-01-0g]
vsan 400 wwn 20:00:00:25:b5:3a:00:1a fcid 0x680402 dynamic
!
    [VDI-28-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1b fcid 0x680403 dynamic
!
    [VDI-23-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:39 fcid 0x680301 dynamic
!
    [VDI-30-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3a fcid 0x680501 dynamic
!
    [VDI-26-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3b fcid 0x680502 dynamic
!
    [VDI-21-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0a fcid 0x680202 dynamic
!
    [VDI-27-HBA1]
```

```
vsan 400 wwn 20:00:00:25:b5:3a:00:1e fcid 0x680302 dynamic
!           [VDI-31-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2b fcid 0x680303 dynamic
!           [VDI-20-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4b fcid 0x680404 dynamic
!           [VDI-19-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1c fcid 0x680203 dynamic
!           [VDI-18-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0b fcid 0x680304 dynamic
!           [VDI-22-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3e fcid 0x680204 dynamic
!           [VDI-6-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0c fcid 0x680503 dynamic
!           [VDI-17-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:49 fcid 0x680405 dynamic
!           [VDI-29-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3c fcid 0x680504 dynamic
!           [VDI-32-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4a fcid 0x680205 dynamic
!           [VDI-24-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1d fcid 0x680505 dynamic
!           [VDI-13-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2c fcid 0x680506 dynamic
!           [VDI-15-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4c fcid 0x680206 dynamic
!           [VDI-14-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:00:4f fcid 0x680406 dynamic
!           [Infra01-8-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2f fcid 0x680507 dynamic
!           [Infra02-16-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2a fcid 0x68020a dynamic
!           [VDI-25-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2e fcid 0x680306 dynamic
!           [VDI-5-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0e fcid 0x680407 dynamic
!           [VDI-7-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3f fcid 0x680307 dynamic
```

```
!           [VDI-1-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0f fcid 0x680208 dynamic
!           [VDI-2-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1f fcid 0x680508 dynamic
!           [VDI-3-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4e fcid 0x680408 dynamic
!           [VDI-4-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4d fcid 0x680309 dynamic
!           [VDI-9-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0d fcid 0x680409 dynamic
!           [VDI-12-HBA1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3d fcid 0x680209 dynamic
!           [VDI-11-HBA1]
```

```
!Active Zone Database Section for vsan 400
```

```
zone name a300_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
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
```

```
zone name a300_VDI-2-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0f
!           [VDI-2-HBA1]
```

```
zone name a300_VDI-3-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1f
!           [VDI-3-HBA1]
```

```
zone name a300_VDI-4-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:4e
!           [VDI-4-HBA1]
```

```
zone name a300_VDI-5-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:2e
!           [VDI-5-HBA1]
```

```
zone name a300_VDI-6-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:3e
!           [VDI-6-HBA1]
```

```
zone name a300_VDI-7-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:0e
!           [VDI-7-HBA1]
```

```
zone name a300_Infra01-8-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:4f
!           [Infra01-8-HBA1]
```

```
zone name a300_VDI-9-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4d
!           [VDI-9-HBA1]
```

```
zone name a300_VDI-10-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2d
!           [VDI-10-HBA1]
```

```
zone name a300_VDI-11-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3d
!           [VDI-11-HBA1]
```

```
zone name a300_VDI-12-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0d
!           [VDI-12-HBA1]
```

```
zone name a300_VDI-13-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1d
```

```
! [VDI-13-HBA1]

zone name a300_VDI-14-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
! [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
! [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:4c
! [VDI-14-HBA1]
```

```
zone name a300_VDI-15-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
! [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
! [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:2c
! [VDI-15-HBA1]
```

```
zone name a300_Infra02-16-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
! [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
! [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:2f
! [Infra02-16-HBA1]
```

```
zone name a300_VDI-17-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
! [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
! [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:0c
! [VDI-17-HBA1]
```

```
zone name a300_VDI-18-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
! [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
! [a300-02-0g]
```

```
member pwnn 20:00:00:25:b5:3a:00:1c
!           [VDI-18-HBA1]

zone name a300_VDI-19-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4b
!           [VDI-19-HBA1]

zone name a300_VDI-20-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2b
!           [VDI-20-HBA1]

zone name a300_VDI-21-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:3b
!           [VDI-21-HBA1]

zone name a300_VDI-22-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0b
!           [VDI-22-HBA1]

zone name a300_VDI-23-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
```



```
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1b
!           [VDI-23-HBA1]

zone name a300_VDI-24-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4a
!           [VDI-24-HBA1]

zone name a300_VDI-25-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2a
!           [VDI-25-HBA1]

zone name a300_VDI-26-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3a
!           [VDI-26-HBA1]

zone name a300_VDI-27-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0a
!           [VDI-27-HBA1]

zone name a300_VDI-28-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
```

```
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1a
!           [VDI-28-HBA1]

zone name a300_VDI-29-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:49
!           [VDI-29-HBA1]

zone name a300_VDI-30-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:39
!           [VDI-30-HBA1]

zone name a300_VDI-31-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1e
!           [VDI-31-HBA1]

zone name a300_VDI-32-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:3c
!           [VDI-32-HBA1]

zoneset name FlexPod_FabricA vsan 400
member a300_VDI-1-HBA1
```

```
member a300_VDI-2-HBA1
member a300_VDI-3-HBA1
member a300_VDI-4-HBA1
member a300_VDI-5-HBA1
member a300_VDI-6-HBA1
member a300_VDI-7-HBA1
member a300_Infra01-8-HBA1
member a300_VDI-9-HBA1
member a300_VDI-10-HBA1
member a300_VDI-11-HBA1
member a300_VDI-12-HBA1
member a300_VDI-13-HBA1
member a300_VDI-14-HBA1
member a300_VDI-15-HBA1
member a300_Infra02-16-HBA1
member a300_VDI-17-HBA1
member a300_VDI-18-HBA1
member a300_VDI-19-HBA1
member a300_VDI-20-HBA1
member a300_VDI-21-HBA1
member a300_VDI-22-HBA1
member a300_VDI-23-HBA1
member a300_VDI-24-HBA1
member a300_VDI-25-HBA1
member a300_VDI-26-HBA1
member a300_VDI-27-HBA1
member a300_VDI-28-HBA1
member a300_VDI-29-HBA1
member a300_VDI-30-HBA1
member a300_VDI-31-HBA1
member a300_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 a300_VDI-1-HBA1 vsan 400
    member pwnn 20:00:00:25:b5:3a:00:3f
!
    [VDI-1-HBA1]
    member pwnn 20:01:00:a0:98:af:bd:e8
```

```
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]

zone name a300_VDI-2-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0f
!           [VDI-2-HBA1]

zone name a300_VDI-3-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1f
!           [VDI-3-HBA1]

zone name a300_VDI-4-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4e
!           [VDI-4-HBA1]

zone name a300_VDI-5-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2e
!           [VDI-5-HBA1]

zone name a300_VDI-6-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
```

```
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:3e
!           [VDI-6-HBA1]

zone name a300_VDI-7-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0e
!           [VDI-7-HBA1]

zone name a300_Infra01-8-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4f
!           [Infra01-8-HBA1]

zone name a300_VDI-9-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4d
!           [VDI-9-HBA1]

zone name a300_VDI-10-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2d
!           [VDI-10-HBA1]

zone name a300_VDI-11-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
```

```
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3d
!           [VDI-11-HBA1]

zone name a300_VDI-12-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0d
!           [VDI-12-HBA1]

zone name a300_VDI-13-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1d
!           [VDI-13-HBA1]

zone name a300_VDI-14-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:4c
!           [VDI-14-HBA1]

zone name a300_VDI-15-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:2c
!           [VDI-15-HBA1]

zone name a300_Infra02-16-HBA1 vsan 400
```

```
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2f
!           [Infra02-16-HBA1]

zone name a300_VDI-17-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:0c
!           [VDI-17-HBA1]

zone name a300_VDI-18-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:1c
!           [VDI-18-HBA1]

zone name a300_VDI-19-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:4b
!           [VDI-19-HBA1]

zone name a300_VDI-20-HBA1 vsan 400
member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
member pwnn 20:00:00:25:b5:3a:00:2b
!           [VDI-20-HBA1]
```

```
zone name a300_VDI-21-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:3b
!           [VDI-21-HBA1]
```

```
zone name a300_VDI-22-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:0b
!           [VDI-22-HBA1]
```

```
zone name a300_VDI-23-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:1b
!           [VDI-23-HBA1]
```

```
zone name a300_VDI-24-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:4a
!           [VDI-24-HBA1]
```

```
zone name a300_VDI-25-HBA1 vsan 400
  member pwnn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwnn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwnn 20:00:00:25:b5:3a:00:2a
!           [VDI-25-HBA1]
```

```
zone name a300_VDI-26-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:3a
!           [VDI-26-HBA1]
```

```
zone name a300_VDI-27-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:0a
!           [VDI-27-HBA1]
```

```
zone name a300_VDI-28-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:1a
!           [VDI-28-HBA1]
```

```
zone name a300_VDI-29-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:49
!           [VDI-29-HBA1]
```

```
zone name a300_VDI-30-HBA1 vsan 400
  member pwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwn 20:00:00:25:b5:3a:00:39
```

```
!                               [VDI-30-HBA1]

zone name a300_VDI-31-HBA1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!                               [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!                               [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:1e
!                               [VDI-31-HBA1]

zone name a300_VDI-32-HBA1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!                               [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!                               [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:3c
!                               [VDI-32-HBA1]

zoneset name FlexPod_FabricA vsan 400
  member a300_VDI-1-HBA1
  member a300_VDI-2-HBA1
  member a300_VDI-3-HBA1
  member a300_VDI-4-HBA1
  member a300_VDI-5-HBA1
  member a300_VDI-6-HBA1
  member a300_VDI-7-HBA1
  member a300_Infra01-8-HBA1
  member a300_VDI-9-HBA1
  member a300_VDI-10-HBA1
  member a300_VDI-11-HBA1
  member a300_VDI-12-HBA1
  member a300_VDI-13-HBA1
  member a300_VDI-14-HBA1
  member a300_VDI-15-HBA1
  member a300_Infra02-16-HBA1
  member a300_VDI-17-HBA1
  member a300_VDI-18-HBA1
  member a300_VDI-19-HBA1
  member a300_VDI-20-HBA1
```

```
member a300_VDI-21-HBA1
member a300_VDI-22-HBA1
member a300_VDI-23-HBA1
member a300_VDI-24-HBA1
member a300_VDI-25-HBA1
member a300_VDI-26-HBA1
member a300_VDI-27-HBA1
member a300_VDI-28-HBA1
member a300_VDI-29-HBA1
member a300_VDI-30-HBA1
member a300_VDI-31-HBA1
member a300_VDI-32-HBA1
```

```
interface mgmt0
  ip address 10.29.164.238 255.255.255.0
clock timezone PST 0 0
vsan database
  vsan 400 interface fc1/1
  vsan 400 interface fc1/2
  vsan 400 interface fc1/3
  vsan 400 interface fc1/4
  vsan 400 interface fc1/5
  vsan 400 interface fc1/6
  vsan 400 interface fc1/7
  vsan 400 interface fc1/8
```

```
switchname ADD16-MDS-A
cli alias name autozone source sys/autozone.py
```

```
line console
line vty
boot kickstart bootflash:/m9100-s6ek9-kickstart-mz.8.3.1.bin
boot system bootflash:/m9100-s6ek9-mz.8.3.1.bin
```

```
interface fc1/1
interface fc1/2
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 fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
```

```
interface fc1/1
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/2
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/3
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/4
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/5
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/6
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/7
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/8
  switchport trunk allowed vsan 400 no-warning
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/9
  no port-license acquire
```

```
interface fcl/10
  no port-license acquire
```

```
interface fcl/11
  no port-license acquire
```

```
interface fcl/12
  no port-license acquire
```

```
interface fcl/13
  no port-license acquire
```

```
interface fcl/14
  no port-license acquire

interface fcl/15
  no port-license acquire

interface fcl/16
  no port-license acquire

ip default-gateway 10.29.164.1
```

Cisco MDS 9132T-B configuration

```
version 8.3(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 $5$1qs42bIH$hp2kMO3FA/4Zzg6EekVHWpA81A7Mc/kBsFZVU8q1uU7 role
network-admin
ip domain-lookup
ip host ADD16-MDS-B 10.29.164.239
aaa group server radius radius
snmp-server user admin network-admin auth md5 0x6fa97f514b0cdf3638e31dfd0bd19c71 priv
0x6fa97f514b0cdf3638e31dfd0bd19c71 localizedkey
snmp-server host 10.155.160.97 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
ntp server 10.81.254.131
ntp server 10.81.254.202
vsan database
```

```
vsan 401 name "FlexPod-B"
```

```
device-alias database
```

```
device-alias name VDI-1-HBA2 pwwn 20:00:00:25:d5:06:00:3f
device-alias name VDI-2-HBA2 pwwn 20:00:00:25:d5:06:00:0f
device-alias name VDI-3-HBA2 pwwn 20:00:00:25:d5:06:00:1f
device-alias name VDI-4-HBA2 pwwn 20:00:00:25:d5:06:00:4e
device-alias name VDI-5-HBA2 pwwn 20:00:00:25:d5:06:00:2e
device-alias name VDI-6-HBA2 pwwn 20:00:00:25:d5:06:00:3e
device-alias name VDI-7-HBA2 pwwn 20:00:00:25:d5:06:00:0e
device-alias name VDI-9-HBA2 pwwn 20:00:00:25:d5:06:00:4d
device-alias name a300-01-0h pwwn 20:02:00:a0:98:af:bd:e8
device-alias name a300-02-0h pwwn 20:04:00:a0:98:af:bd:e8
device-alias name VDI-10-HBA2 pwwn 20:00:00:25:d5:06:00:2d
device-alias name VDI-11-HBA2 pwwn 20:00:00:25:d5:06:00:3d
device-alias name VDI-12-HBA2 pwwn 20:00:00:25:d5:06:00:0d
device-alias name VDI-13-HBA2 pwwn 20:00:00:25:d5:06:00:1d
device-alias name VDI-14-HBA2 pwwn 20:00:00:25:d5:06:00:4c
device-alias name VDI-15-HBA2 pwwn 20:00:00:25:d5:06:00:2c
device-alias name VDI-17-HBA2 pwwn 20:00:00:25:d5:06:00:0c
device-alias name VDI-18-HBA2 pwwn 20:00:00:25:d5:06:00:1c
device-alias name VDI-19-HBA2 pwwn 20:00:00:25:d5:06:00:4b
device-alias name VDI-20-HBA2 pwwn 20:00:00:25:d5:06:00:2b
device-alias name VDI-21-HBA2 pwwn 20:00:00:25:d5:06:00:3b
device-alias name VDI-22-HBA2 pwwn 20:00:00:25:d5:06:00:6b
device-alias name VDI-23-HBA2 pwwn 20:00:00:25:d5:06:00:1b
device-alias name VDI-24-HBA2 pwwn 20:00:00:25:d5:06:00:4a
device-alias name VDI-25-HBA2 pwwn 20:00:00:25:d5:06:00:2a
device-alias name VDI-26-HBA2 pwwn 20:00:00:25:d5:06:00:3a
device-alias name VDI-27-HBA2 pwwn 20:00:00:25:d5:06:00:0a
device-alias name VDI-28-HBA2 pwwn 20:00:00:25:d5:06:00:1a
device-alias name VDI-29-HBA2 pwwn 20:00:00:25:d5:06:00:49
device-alias name VDI-30-HBA2 pwwn 20:00:00:25:d5:06:00:39
device-alias name VDI-31-HBA2 pwwn 20:00:00:25:d5:06:00:1e
device-alias name VDI-32-HBA2 pwwn 20:00:00:25:d5:06:00:3c
```

```
device-alias name Infra01-8-HBA2 pwwn 20:00:00:25:d5:06:00:4f
device-alias name Infra02-16-HBA2 pwwn 20:00:00:25:d5:06:00:2f
device-alias commit
```

```
fcdomain fcid database
```

```
vsan 401 wwn 50:0a:09:84:80:01:c7:87 fcid 0x5b0000 dynamic
vsan 401 wwn 20:00:00:25:d5:06:00:2e fcid 0x870509 dynamic
!
    [VDI-5-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0e fcid 0x870508 dynamic
!
    [VDI-7-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1d fcid 0x870203 dynamic
!
    [VDI-13-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:2c fcid 0x870507 dynamic
!
    [VDI-15-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:4c fcid 0x870302 dynamic
!
    [VDI-14-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:2f fcid 0x870207 dynamic
!
    [Infra02-16-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0c fcid 0x870505 dynamic
!
    [VDI-17-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1b fcid 0x870406 dynamic
!
    [VDI-23-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:39 fcid 0x870504 dynamic
!
    [VDI-30-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3f fcid 0x870403 dynamic
!
    [VDI-1-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1e fcid 0x870501 dynamic
!
    [VDI-31-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1c fcid 0x87020a dynamic
!
    [VDI-18-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:2a fcid 0x870301 dynamic
!
    [VDI-25-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1f fcid 0x870402 dynamic
!
    [VDI-3-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:4e fcid 0x870303 dynamic
!
    [VDI-4-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:1a fcid 0x87030a dynamic
!
    [VDI-28-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3e fcid 0x870405 dynamic
```



```
! [VDI-6-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0d fcid 0x870201 dynamic
! [VDI-12-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3d fcid 0x870306 dynamic
! [VDI-11-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3a fcid 0x87040a dynamic
! [VDI-26-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0a fcid 0x870204 dynamic
! [VDI-27-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:4b fcid 0x870503 dynamic
! [VDI-19-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0f fcid 0x870202 dynamic
! [VDI-2-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:49 fcid 0x870304 dynamic
! [VDI-29-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3b fcid 0x870408 dynamic
! [VDI-21-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:0b fcid 0x870307 dynamic
vsan 401 wwn 20:00:00:25:d5:06:00:4f fcid 0x870401 dynamic
! [Infra01-8-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:4d fcid 0x870506 dynamic
! [VDI-9-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:2d fcid 0x870407 dynamic
! [VDI-10-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:3c fcid 0x870206 dynamic
! [VDI-32-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:2b fcid 0x870308 dynamic
! [VDI-20-hba2]
vsan 401 wwn 20:00:00:25:d5:06:00:4a fcid 0x87020b dynamic
! [VDI-24-hba2]
vsan 401 wwn 50:0a:09:83:80:d3:67:d3 fcid 0x870000 dynamic
vsan 401 wwn 20:04:00:a0:98:af:bd:e8 fcid 0x870001 dynamic
! [a300-02-0h]
vsan 401 wwn 50:0a:09:83:80:13:41:27 fcid 0x870100 dynamic
vsan 401 wwn 20:02:00:a0:98:af:bd:e8 fcid 0x870101 dynamic
! [a300-01-0h]
vsan 401 wwn 20:01:00:de:fb:92:0c:80 fcid 0x870200 dynamic
vsan 401 wwn 20:02:00:de:fb:92:0c:80 fcid 0x870300 dynamic
vsan 401 wwn 20:03:00:de:fb:92:0c:80 fcid 0x870400 dynamic
```

vsan 401 wwn 20:04:00:de:fb:92:0c:80 fcid 0x870500 dynamic

!Active Zone Database Section for vsan 401

zone name a300_VDI-1-HBA2 vsan 401

member pwnn 20:00:00:25:d5:06:00:3f

! [VDI-1-HBA2]

member pwnn 20:02:00:a0:98:af:bd:e8

! [a300-01-0h]

member pwnn 20:04:00:a0:98:af:bd:e8

! [a300-02-0h]

zone name a300_VDI-2-HBA2 vsan 401

member pwnn 20:00:00:25:d5:06:00:0f

! [VDI-2-HBA2]

member pwnn 20:02:00:a0:98:af:bd:e8

! [a300-01-0h]

member pwnn 20:04:00:a0:98:af:bd:e8

! [a300-02-0h]

zone name a300_VDI-3-HBA2 vsan 401

member pwnn 20:00:00:25:d5:06:00:1f

! [VDI-3-HBA2]

member pwnn 20:02:00:a0:98:af:bd:e8

! [a300-01-0h]

member pwnn 20:04:00:a0:98:af:bd:e8

! [a300-02-0h]

zone name a300_VDI-4-HBA2 vsan 401

member pwnn 20:00:00:25:d5:06:00:4e

! [VDI-4-HBA2]

member pwnn 20:02:00:a0:98:af:bd:e8

! [a300-01-0h]

member pwnn 20:04:00:a0:98:af:bd:e8

! [a300-02-0h]

zone name a300_VDI-5-HBA2 vsan 401

member pwnn 20:00:00:25:d5:06:00:2e

! [VDI-5-HBA2]

member pwnn 20:02:00:a0:98:af:bd:e8

```
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-6-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:3e
!           [VDI-6-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-7-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:0e
!           [VDI-7-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_Infra01-8-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:4f
!           [Infra01-8-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-9-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:4d
!           [VDI-9-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-10-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:2d
!           [VDI-10-HBA2]
```

```
member pwnn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-11-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:3d
!           [VDI-11-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-12-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:0d
!           [VDI-12-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-13-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:1d
!           [VDI-13-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-14-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:4c
!           [VDI-14-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-15-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:2c
```

```
!           [VDI-15-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_Infra02-16-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:2f
!           [Infra02-16-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-17-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:0c
!           [VDI-17-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-18-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:1c
!           [VDI-18-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-19-HBA2 vsan 401
member pwn 20:00:00:25:d5:06:00:4b
!           [VDI-19-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-20-HBA2 vsan 401
```

```
member pwnn 20:00:00:25:d5:06:00:2b
!          [VDI-20-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-21-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:3b
!          [VDI-21-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-22-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:6b
!          [VDI-22-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-23-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:1b
!          [VDI-23-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-24-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:4a
!          [VDI-24-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-25-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:2a
  !          [VDI-25-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-26-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:3a
  !          [VDI-26-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-27-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0a
  !          [VDI-27-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-28-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:1a
  !          [VDI-28-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-29-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:49
  !          [VDI-29-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-30-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:39
  !          [VDI-30-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-31-HBA2 vsan 401
  member pwnn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
  member pwnn 20:00:00:25:d5:06:00:1e
  !          [VDI-31-HBA2]
```

```
zone name a300_VDI-32-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:3c
  !          [VDI-32-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zoneset name FlexPod_FabricB vsan 401
  member a300_VDI-1-HBA2
  member a300_VDI-2-HBA2
  member a300_VDI-3-HBA2
  member a300_VDI-4-HBA2
  member a300_VDI-5-HBA2
  member a300_VDI-6-HBA2
  member a300_VDI-7-HBA2
  member a300_Infra01-8-HBA2
  member a300_VDI-9-HBA2
  member a300_VDI-10-HBA2
  member a300_VDI-11-HBA2
  member a300_VDI-12-HBA2
  member a300_VDI-13-HBA2
```



```
member a300_VDI-14-HBA2
member a300_VDI-15-HBA2
member a300_Infra02-16-HBA2
member a300_VDI-17-HBA2
member a300_VDI-18-HBA2
member a300_VDI-19-HBA2
member a300_VDI-20-HBA2
member a300_VDI-21-HBA2
member a300_VDI-22-HBA2
member a300_VDI-23-HBA2
member a300_VDI-24-HBA2
member a300_VDI-25-HBA2
member a300_VDI-26-HBA2
member a300_VDI-27-HBA2
member a300_VDI-28-HBA2
member a300_VDI-29-HBA2
member a300_VDI-30-HBA2
member a300_VDI-31-HBA2
member a300_VDI-32-HBA2
```

```
zoneset activate name FlexPod_FabricB vsan 401
```

```
do clear zone database vsan 401
```

```
!Full Zone Database Section for vsan 401
```

```
zone name a300_VDI-1-HBA2 vsan 401
```

```
member pwn 20:00:00:25:d5:06:00:3f
```

```
! [VDI-1-HBA2]
```

```
member pwn 20:02:00:a0:98:af:bd:e8
```

```
! [a300-01-0h]
```

```
member pwn 20:04:00:a0:98:af:bd:e8
```

```
! [a300-02-0h]
```

```
zone name a300_VDI-2-HBA2 vsan 401
```

```
member pwn 20:00:00:25:d5:06:00:0f
```

```
! [VDI-2-HBA2]
```

```
member pwn 20:02:00:a0:98:af:bd:e8
```

```
! [a300-01-0h]
```

```
member pwn 20:04:00:a0:98:af:bd:e8
```

```
! [a300-02-0h]
```

```
zone name a300_VDI-3-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:1f
  !          [VDI-3-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-4-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:4e
  !          [VDI-4-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-5-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:2e
  !          [VDI-5-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-6-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:3e
  !          [VDI-6-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_VDI-7-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0e
  !          [VDI-7-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  !          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  !          [a300-02-0h]
```

```
zone name a300_Infra01-8-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:4f
    !          [Infra01-8-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
    !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
    !          [a300-02-0h]
```

```
zone name a300_VDI-9-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:4d
    !          [VDI-9-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
    !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
    !          [a300-02-0h]
```

```
zone name a300_VDI-10-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:2d
    !          [VDI-10-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
    !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
    !          [a300-02-0h]
```

```
zone name a300_VDI-11-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:3d
    !          [VDI-11-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
    !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
    !          [a300-02-0h]
```

```
zone name a300_VDI-12-HBA2 vsan 401
  member pwnn 20:00:00:25:d5:06:00:0d
    !          [VDI-12-HBA2]
  member pwnn 20:02:00:a0:98:af:bd:e8
    !          [a300-01-0h]
  member pwnn 20:04:00:a0:98:af:bd:e8
```

```
! [a300-02-0h]

zone name a300_VDI-13-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:1d
  ! [VDI-13-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  ! [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  ! [a300-02-0h]

zone name a300_VDI-14-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:4c
  ! [VDI-14-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  ! [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  ! [a300-02-0h]

zone name a300_VDI-15-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:2c
  ! [VDI-15-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  ! [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  ! [a300-02-0h]

zone name a300_Infra02-16-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:2f
  ! [Infra02-16-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  ! [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
  ! [a300-02-0h]

zone name a300_VDI-17-HBA2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0c
  ! [VDI-17-HBA2]
  member pwwn 20:02:00:a0:98:af:bd:e8
  ! [a300-01-0h]
```

```
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-18-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:1c
!          [VDI-18-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-19-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:4b
!          [VDI-19-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-20-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:2b
!          [VDI-20-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-21-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:3b
!          [VDI-21-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-22-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:6b
!          [VDI-22-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
```

```
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-23-HBA2 vsan 401
  member pwn 20:00:00:25:d5:06:00:1b
  !           [VDI-23-HBA2]
  member pwn 20:02:00:a0:98:af:bd:e8
  !           [a300-01-0h]
  member pwn 20:04:00:a0:98:af:bd:e8
  !           [a300-02-0h]
```

```
zone name a300_VDI-24-HBA2 vsan 401
  member pwn 20:00:00:25:d5:06:00:4a
  !           [VDI-24-HBA2]
  member pwn 20:02:00:a0:98:af:bd:e8
  !           [a300-01-0h]
  member pwn 20:04:00:a0:98:af:bd:e8
  !           [a300-02-0h]
```

```
zone name a300_VDI-25-HBA2 vsan 401
  member pwn 20:00:00:25:d5:06:00:2a
  !           [VDI-25-HBA2]
  member pwn 20:02:00:a0:98:af:bd:e8
  !           [a300-01-0h]
  member pwn 20:04:00:a0:98:af:bd:e8
  !           [a300-02-0h]
```

```
zone name a300_VDI-26-HBA2 vsan 401
  member pwn 20:00:00:25:d5:06:00:3a
  !           [VDI-26-HBA2]
  member pwn 20:02:00:a0:98:af:bd:e8
  !           [a300-01-0h]
  member pwn 20:04:00:a0:98:af:bd:e8
  !           [a300-02-0h]
```

```
zone name a300_VDI-27-HBA2 vsan 401
  member pwn 20:00:00:25:d5:06:00:0a
  !           [VDI-27-HBA2]
```

```
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-28-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:1a
!          [VDI-28-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-29-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:49
!          [VDI-29-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-30-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:39
!          [VDI-30-HBA2]
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-31-HBA2 vsan 401
member pwnn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
member pwnn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
member pwnn 20:00:00:25:d5:06:00:1e
!          [VDI-31-HBA2]
```

```
zone name a300_VDI-32-HBA2 vsan 401
member pwnn 20:00:00:25:d5:06:00:3c
```

```
!           [VDI-32-HBA2]
member pwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
member pwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zoneset name FlexPod_FabricB vsan 401
```

```
member a300_VDI-1-HBA2
member a300_VDI-2-HBA2
member a300_VDI-3-HBA2
member a300_VDI-4-HBA2
member a300_VDI-5-HBA2
member a300_VDI-6-HBA2
member a300_VDI-7-HBA2
member a300_Infra01-8-HBA2
member a300_VDI-9-HBA2
member a300_VDI-10-HBA2
member a300_VDI-11-HBA2
member a300_VDI-12-HBA2
member a300_VDI-13-HBA2
member a300_VDI-14-HBA2
member a300_VDI-15-HBA2
member a300_Infra02-16-HBA2
member a300_VDI-17-HBA2
member a300_VDI-18-HBA2
member a300_VDI-19-HBA2
member a300_VDI-20-HBA2
member a300_VDI-21-HBA2
member a300_VDI-22-HBA2
member a300_VDI-23-HBA2
member a300_VDI-24-HBA2
member a300_VDI-25-HBA2
member a300_VDI-26-HBA2
member a300_VDI-27-HBA2
member a300_VDI-28-HBA2
member a300_VDI-29-HBA2
member a300_VDI-30-HBA2
member a300_VDI-31-HBA2
member a300_VDI-32-HBA2
```



```
interface mgmt0
  ip address 10.29.164.239 255.255.255.0

vsan database
  vsan 401 interface fc1/1
  vsan 401 interface fc1/2
  vsan 401 interface fc1/3
  vsan 400 interface fc1/4
  vsan 401 interface fc1/5
  vsan 401 interface fc1/6
  vsan 401 interface fc1/7
  vsan 401 interface fc1/8

clock timezone PST 0 0
clock summer-time PDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
switchname ADD16-MDS-B
cli alias name autozone source sys/autozone.py
line console
line vty
boot kickstart bootflash:/m9100-s6ek9-kickstart-mz.8.3.1.bin
boot system bootflash:/m9100-s6ek9-mz.8.3.1.bin

interface fc1/1
interface fc1/2
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 fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
```

```
interface fcl/16
```

```
interface fcl/1
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/2
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/3
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/4
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/5
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/6
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/7
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/8
  switchport trunk allowed vsan 401
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fcl/9
  no port-license acquire
```

```
interface fcl/10
  no port-license acquire
```

```
interface fcl/11
  no port-license acquire
```

```
interface fcl/12
  no port-license acquire
```

```
interface fcl/13
  no port-license acquire
```

```
interface fcl/14
  no port-license acquire
```

```
interface fcl/15
  no port-license acquire
```

```
interface fcl/16
  no port-license acquire
```

```
ip default-gateway 10.29.164.1
```



Full-scale server performance chart with boot and loginvsi knowledge worker workload test

This section provides a detailed performance chart for ESXi 6.7 U2 installed on Cisco UCS B200 M5 Blade Server part of the workload test with VMware Horizon 7 deployed on NetApp AFF A300 system running LoginVSI v4.1.25 based knowledge worker workload part of the FlexPod Datacenter reference architecture defined here.

The charts below are defined in the set of 10 hosts in the single performance chart.

VDI server performance monitor data for one sample VDI server: 5400 users VDI non-persistent scale testing

Figure 156. Full-scale | 5400 Non-persistent users| VDI host | Host CPU utilization











Figure 159. Full-scale | 5400 Non-persistent users| VDI host | Host network utilization | Transmitted





VDI server performance monitor data for one sample VDI server: 5400 users VDI persistent scale testing

Figure 160. Full-scale | 5400 Persistent users| VDI host | Host CPU utilization






Figure 161. Full-scale | 5400 Persistent users| VDI host | Host memory utilization






Figure 162. Full-scale | 5400 Persistent users| VDI host | Host network utilization | Received











Figure 163. Full-scale | 5400 Persistent users| VDI host | Host network utilization | Transmitted















Figure 167. Full-scale | 6700 RDS users | VDI host | Host network utilization | Transmitted





VDI server performance monitor data for one sample VDI server: 5800 users mixed scale testing

Figure 168. Full-scale | 5800 Mixed users | VDI host | Host CPU utilization






Figure 169. Full-scale | 5800 Mixed users | VDI host | Host memory utilization










Figure 171. Full-scale | 5800 Mixed users | VDI host | Host network utilization | Transmitted







PowerShell example to clones persistent desktops from a template using VAAI

```
# Name: Clone Desktops Powershell script
# Date: 1/02/2018
# Description: Clones persistent desktops from a template and utilizes VAAI
#
# Author: Dave Arnette, NetApp, Inc.
#
# Revisions:
#
# Starting script
[CmdletBinding()]
Param(

    [Parameter()]
    [switch]$deploy = $false,

    [Parameter()]
    [switch]$changeNetwork = $false,

    [Parameter()]
    [switch]$test = $false

)

$TargetCluster = Get-Cluster -Name "TMEonly"
```

```
$targetFolder = "Launchers"
$vmHost = "192.168.201.101"
$SourceVMTemplate = Get-Template -Name "W2016_template_0510"
$SourceCustomSpec = Get-OSCustomizationSpec -Name "desktop-v1"
$datastore = get-datastore -name "TMEinfraONLY"
$domaincred = (New-Object System.Management.Automation.PSCredential
"vdi\administrator", (get-content c:\vdi_domaincred.txt | convertto-securestring) )

$servers = (1,3,5,7,8)
$basename = "launcher"

$vmList = @()

$servers | foreach-object {

    $vm = New-Object System.Object

    if ( $_ -le 9) {
        $name = "${basename}-00${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.10${_}"

    }

    elseif ( $_ -ge 10 -and $_ -le 99 ) {
        $name = "${basename}-0${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.1${_}"

    }

    else {
        $name = "${basename}-${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.${_}"

    }

    $vmList = $vmList + $vm
}
}
```

```
function change_dns {
    Write-host "Waiting for VMware Tools to start on $vmname"
    wait-tools -vm $vmname -TimeoutSeconds 300

    write-host " Changing DNS server on $vmname"

    $dnsChangeCmd = { get-DNSClientServerAddress -interfaceAlias "Ethernet0" | set-
DNSClientServerAddress -serveraddresses "192.168.201.41" }
    invoke-command -computername $vmname -credential $domaincred -scriptblock $dnsChan-
geCmd

}

function change_ip {

    write-host " Changing IP address on $vmname"

    $ipChangeCmd = { new-netIPAddress -interfaceAlias "Ethernet0" -ipaddress $Using:ip
-PrefixLength "16" -DefaultGateway "172.18.255.254" }

    invoke-command -computername $vmname -credential $domaincred -scriptblock $ipChan-
geCmd -AsJob

}

if ($deploy) {
    if ($test) {
        $vmList| foreach-object {
            $vmname = $_.name

            write-host " Creating VM $vmname from template $sourceVMtemplate " -
ForegroundColor Green
        }
    }
    else {

        #create the VMs asynchronously

        $createJobList = @()
```

```
$vmList| foreach-object {
    # $vmname = ""
    $vmname = $_.Name

    write-host " Creating VM $vmname from template $sourceVMtemplate " -
    ForegroundColor Green

    new-vm -name $vmname -resourcepool $targetCluster -vmhost $vmhost -location
    $targetFolder -template $sourceVMtemplate -OSCustomizationSpec $SourceCustomSpec -
    datastore $datastore -runasync -OutVariable createJob

    $createJob | add-member -MemberType NoteProperty -name "VM" -value
    "$vmname"

    $createJobList = $createJobList + $createJob
}

# Wait for each clone job to finish, then start the VM

$createJobList| foreach-object {
    $vmname = $_.VM

    write-host " Waiting for VM $vmname clone process to finish " -
    ForegroundColor Green

    do { start-sleep -seconds 5 }
    until ( (get-task -id $_.ID).state -eq "Success" )

    write-host " Starting VM $vmname " -ForegroundColor Green
    start-vm -vm $vmname -runAsync
}

}

}

if ($changeNetwork) {
```

```
if ($test) {
    $vmList| foreach-object {
        $vmname = $_.Name
        $ip = $_.IP

        write-host " Changing DNS on $vmname" -ForegroundColor Green
        Write-host " Changing IP on $vmname to $ip"
    }
}
else {

    $vmList | foreach-object {
        $vmname = $_.name
        $ip = $_.IP

        change_dns
        change_ip

    }
}
}
```

About the Authors

- Vadim Lebedev, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Vadim Lebedev is a member of the Cisco's Computing Systems Product Group team focusing on design, testing, and solutions validation, technical content creation, and performance testing/benchmarking. He has years of experience in server and desktop virtualization. Vadim is a subject matter expert on Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA Graphics.

- Suresh Thoppay, Sr. Technical Marketing Engineer, Digital Workspace Solutions, NetApp

Suresh Thoppay is part of the NetApp Converged Infrastructure Engineering Solutions team focusing on design, testing and solutions validation, technical content creation, and performance testing/benchmarking. He is a member of Login VSI Technology Advocates and NVIDIA GRID Community Advisors. Suresh has obtained various IT vendor certifications over the last 25 years.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the following for their contribution and expertise that resulted in developing this document:

- Mike Brennan, Product Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.
- Ketan Mota, Sr. Product Manager, FlexPod Solutions, NetApp

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>.

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)