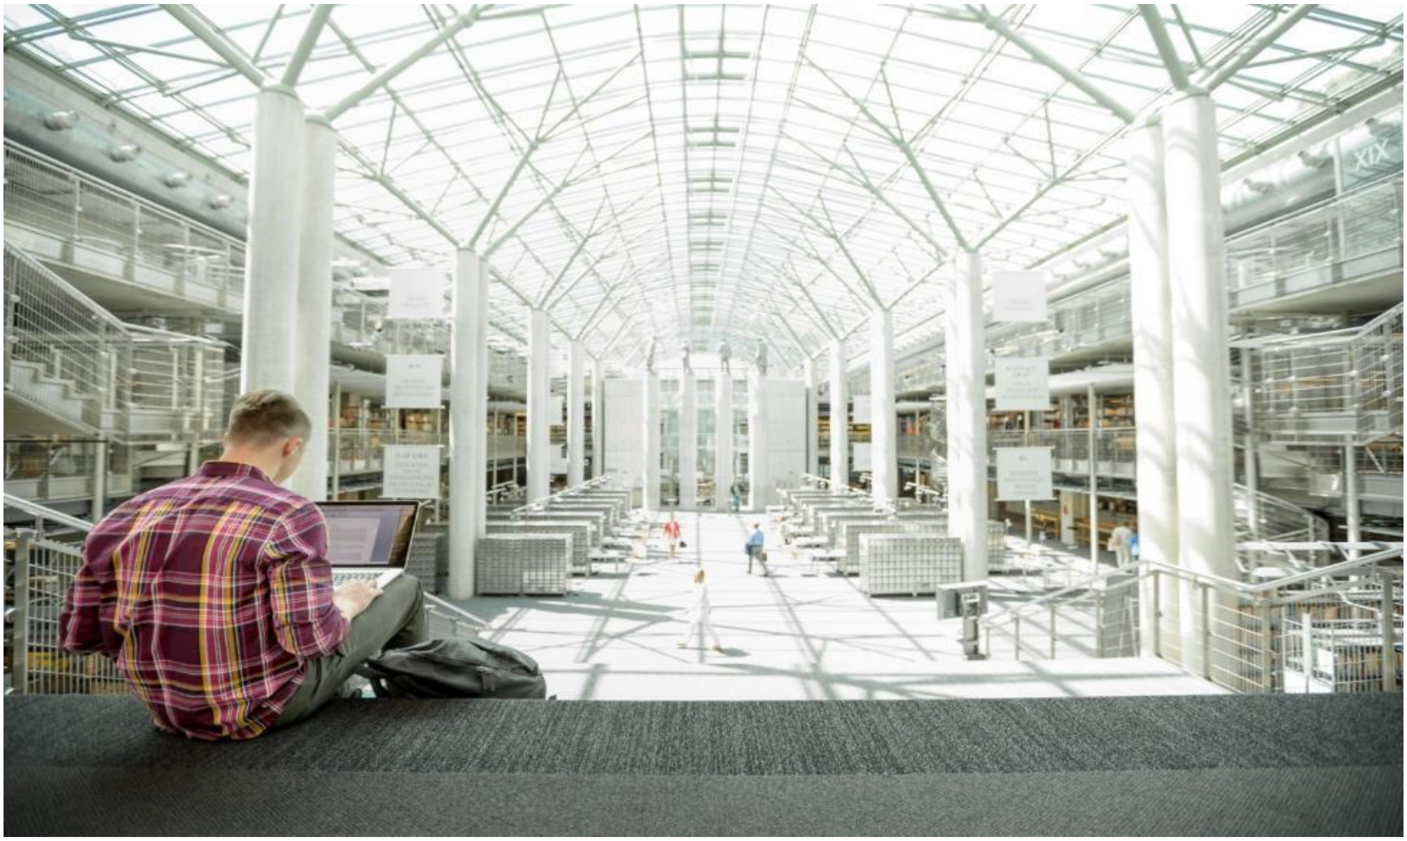

FlashStack Data Center with VMware Horizon 8 and VMware vSphere 7.0 with 4th Generation Cisco UCS

Deployment Guide for Virtual Desktop Infrastructure Built on Cisco UCS B200 M5 and Cisco UCS Manager 4.0 with Pure Storage FlashArray//X70 R3 Array, VMware Horizon 8, and VMware vSphere 7.0 Hypervisor Platform

Published: May 2021

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.

Executive Summary

Cisco Validated Designs 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, Pure and VMware 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, Pure Storage, 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 FlashStack Data Center with 4th Generation Cisco UCS and Pure Storage® FlashArray//X70 R3 with 100 percent DirectFlash Modules and DirectFlash Software. The solution includes VMware Horizon server-based RDS Windows Sever 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 ESXi 7.0 GA hypervisor.

The solution is a predesigned, best-practice data center architecture built on the FlashStack reference architecture. The FlashStack Data Center used in this validation includes Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel (FC) switches and Pure All-NVMe FlashArray//X system.

This solution is 100 percent virtualized on fifth generation Cisco UCS B200 M5 blade servers, booting VMware vSphere ESXi 7.0 GA through FC SAN from the FlashArray//X70 R3 storage array. Where applicable the document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

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 4.1.39.6 Knowledge Worker workload running in benchmark mode, providing a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility, and reducing costs. Cisco, Pure Storage, 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 8 desktop workload solution with Pure Storage FlashArray//X 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 VMware Horizon 8 Virtual Desktop Infrastructure (VDI) deployment Cisco Validated Design with Cisco UCS 5th generation servers and Pure Storage FlashArray//X Series system.

It incorporates the following features:

- Cisco UCS B200 M5 blade servers with Intel Xeon® Gold 6230 CPU
- 64GB DDR4-2933-MHz memory
- Support for the Cisco UCS 4.1(2a) release
- Support for the latest release of Pure Storage FlashArray//X70 R3 hardware and Purity//FA v6.0.3
- Introducing FA//File capabilities
- VMware vSphere 7.0 GA Hypervisor
- VMware Horizon 8 Server 2019 RDS hosted server sessions
- VMware Horizon 8 non-persistent Instant Clone Windows 10 virtual machines
- VMware Horizon 8 persistent Full Clone Windows 10 virtual machines

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

These factors have led to the need for 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 Data Center
- Service Provider Data Center
- Large Commercial Data Center

Solution Summary

FlashStack provides a jointly supported solution by Cisco and Pure Storage, providing a carefully validated architecture built on superior compute, world class networking, and the leading innovations in all flash storage.

The portfolio of validated offerings from FlashStack includes but is not limited to the following:

- **Consistent performance:** FlashStack provides higher, more consistent performance than disk-based solutions and delivers a converged infrastructure based on all-flash that provides non-disruptive upgrades and scalability.
- **Cost savings:** FlashStack uses less power, cooling, and data center space when compared to legacy disk/hybrid storage. It provides industry-leading storage data reduction and exceptional storage density.
- **Simplicity:** FlashStack requires low ongoing maintenance and reduces operational overhead. It also scales simply and smoothly in step with business requirements.
- **Deployment choices:** It is available as a custom-built single unit from FlashStack partners, but organizations can also deploy using equipment from multiple sources, including equipment they already own.
- **Unique business model:** The Pure Storage Evergreen Storage Model enables companies to keep their storage investments forever, which means no more forklift upgrades and no more downtime.
- **Mission-critical resiliency:** FlashStack offers best in class performance by providing active-active resiliency, no single point of failure, and non-disruptive operations, enabling organizations to maximize productivity.

- **Support choices:** Focused, high-quality single-number reach for FlashStack support is available from FlashStack Authorized Support Partners. Single-number support is also available directly from Cisco Systems as part of the Cisco Solution Support for Data Center offering. Support for FlashStack components is also available from Cisco, VMware, and Pure Storage individually and leverages TSANet for resolution of support queries between vendors.

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both VMware Horizon Microsoft Windows 10 virtual desktops and VMware Horizon RDS sessions based on Microsoft Windows Server 2019.

The mixed workload solution includes Pure Storage FlashArray//X®, 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 data center 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., Pure Storage Inc., and VMware Inc. produced a highly efficient, robust, and affordable desktop virtualization solution for a hosted virtual desktop and hosted shared desktop mixed deployment supporting different use cases. Key components of this solution include the following:

- **More power, same size.** Cisco UCS B200 M5 half-width blade with dual 20-core 2.1 GHz Intel ® Xeon ® Scalable Family Gold (6230) processors and 768 GB of memory for VMware Horizon hosts supports more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel 20-core 2.1 GHz Intel ® Xeon ® Gold Scalable Family (6230) processors used in this study provided a balance between increased per-blade capacity and cost.
- **Fault-tolerance with high availability built into the design.** The various designs are based on using one Unified Computing System chassis with multiple Cisco UCS B200 M5 blades for virtualized desktop and infrastructure workloads. The design provides N+1 server fault tolerance for hosted virtual desktops, hosted shared desktops and infrastructure services.
- **Stress-tested to the limits during aggressive boot scenario.** The servers hosting RDS sessions and VDI shared and statically assigned desktop environment booted and became available within very short time, 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 workloads 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 data center.** The rack space required to support the system is less than a single 42U rack, conserving valuable data center floor space.
- **All Virtualized:** This Cisco Validated Design (CVD) presents a validated design that is 100 percent virtualized on VMware ESXi 7.0 GA. All of the virtual desktops, user data, profiles, and supporting infrastructure

components, including Active Directory, SQL Servers, VMware Horizon components, VDI desktops and RDS servers were hosted as virtual machines. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the FlashStack converged infrastructure with stateless Cisco UCS Blade servers and Pure FC storage.

- **Cisco maintains industry leadership** with the new Cisco UCS Manager 4.1(2a) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager (UCSM), Cisco UCS Central, Cisco UCS Director and Cisco Intersight ensure 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.
- **Cisco SAN architectural benefit** of the next-generation 32-Gb fabric switches address the requirement for highly scalable, virtualized, intelligent SAN infrastructure in current-generation data center environments.
- **Pure All-NVMe FlashArray//X70 R3 storage 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.
- **Pure All-NVMe FlashArray//X70 R3 storage array** provides a simple to understand storage architecture for hosting all user data components (virtual machines, profiles, user data) on the same storage array.
- **Pure Storage** software enables to seamlessly add, upgrade, or remove capacity and/or controllers from the infrastructure to meet the needs of the virtual desktops transparently.
- **Pure Storage Management UI** for VMware vSphere hypervisor has deep integrations with vSphere, providing easy-button automation for key storage tasks such as storage repository provisioning, storage resize, directly from vCenter.
- **VMware Horizon 8.** VMware Horizon is a modern platform for secure delivery of virtual desktops and apps across the hybrid cloud. VMware's virtualization heritage provides Horizon unique benefits and best-in-class technologies that enable one-to-many provisioning and streamlined management of images, apps, profiles, and policies for an agile, lightweight, modern approach that speeds, simplifies and reduces costs. Horizon, powered by the Blast Extreme protocol, delivers an immersive, feature rich user experience for end users across devices, locations, media, and network connections. Enabled by enterprise-grade management capabilities and a deep VMware technology ecosystem, Horizon extends the digital workspace to all apps and secure productivity use cases
- **Optimized to achieve the best possible performance and scale**. For hosted shared 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 (RDS) shared virtual machines were setup as Manual farms and VMware Horizon 8, Microsoft Windows 10 virtual machines were created for this solution using VMware pooled desktops.

Cisco Desktop Virtualization Solutions: Data Center

The Evolving Workplace

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

This workforce is also increasingly mobile, conducting business in traditional offices, conference rooms across the enterprise campus, home offices, on the road, in hotels, and at the local coffee shop. This workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference.

These trends are increasing pressure on IT to ensure protection of corporate data and prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 1).

These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 10 and productivity tools, specifically Microsoft Office 2019.

Figure 1. Cisco Data Center Partner Collaboration

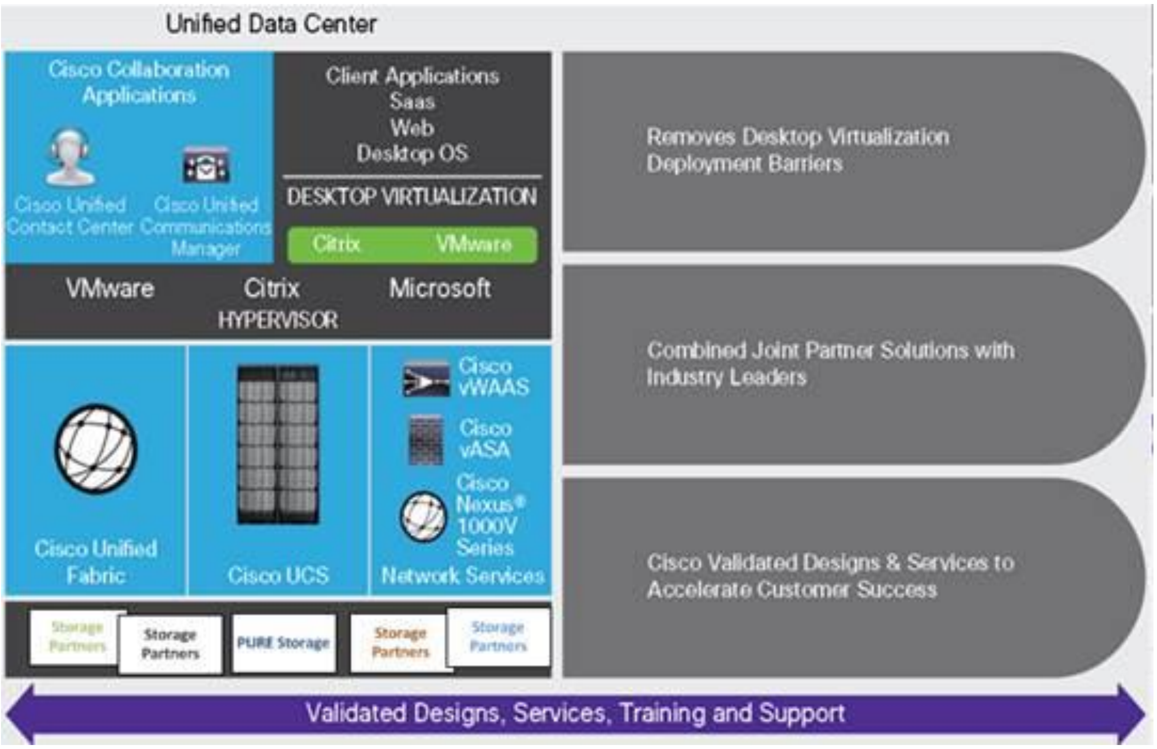
Unified Data Center
Cisco Collaboration Applications
Cisco Unified Contact Center
Cisco Unified Communications Manager
Client Applications
Saas
Web
Desktop OS
DESKTOP VIRTUALIZATION
Citrix
VMware
VMware
Citrix
Microsoft
HYPERVISOR
Cisco Unified Fabric
Cisco UCS
Cisco vWAAS
Cisco vASA
Cisco Nexus® 1000V Series
Network Services
Storage Partners
Storage Partners
PURE Storage
Storage Partners
Storage Partners
Removes Desktop Virtualization Deployment Barriers
Combined Joint Partner Solutions with Industry Leaders
Cisco Validated Designs & Services to Accelerate Customer Success
Validated Designs, Services, Training and Support

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 data center infrastructure: simplification, security, and scalability. The software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform.

Simplified

Cisco UCS 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 data center 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.

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

Simplification also leads to more successful desktop virtualization implementation. Cisco and its technology partners like VMware Technologies, and Pure Storage have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlashStack. 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 data center infrastructure (Cisco UCS and Cisco Nexus Family solutions) for desktop virtualization provides strong data center, network, and desktop security, with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, virtual machine–aware policies and administration, and network security across the LAN and WAN infrastructure.

Scalable

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 FlashStack Data Center infrastructure supports high virtual-desktop density (desktops per server), and additional servers and storage scale with near-linear performance. FlashStack Data Center 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 3 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 40 Gb per server, and the northbound Cisco UCS fabric interconnect can output 3.82 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 Pure, 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 FlashStack solutions have demonstrated scalability and performance.

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

Savings and Success

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

The simplified deployment of Cisco 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

Compute Connectivity

Each compute chassis in the design is redundantly connected to the managing fabric interconnects with at least two ports per IOM. Ethernet traffic from the upstream network and Fibre Channel frames coming from the FlashArray are converged within the fabric interconnect to be both Ethernet and Fibre Channel over Ethernet and transmitted to the UCS servers through the IOM. These IOM connections from the Cisco UCS Fabric Interconnects to the IOMs are automatically configured as port channels by specifying a Chassis/FEX Discovery Policy within UCSM.

Each rack server in the design is redundantly connected to the managing fabric interconnects with at least one port to each FI. Ethernet traffic from the upstream network and Fibre Channel frames coming from the FlashArray are converged within the fabric interconnect to be both Ethernet and Fibre Channel over Ethernet and transmitted to the UCS server.

These connections from the 4th Gen UCS 6454 Fabric Interconnect to the 2408 IOM hosted within the chassis are shown in Figure 2.

Figure 2. Compute Connectivity

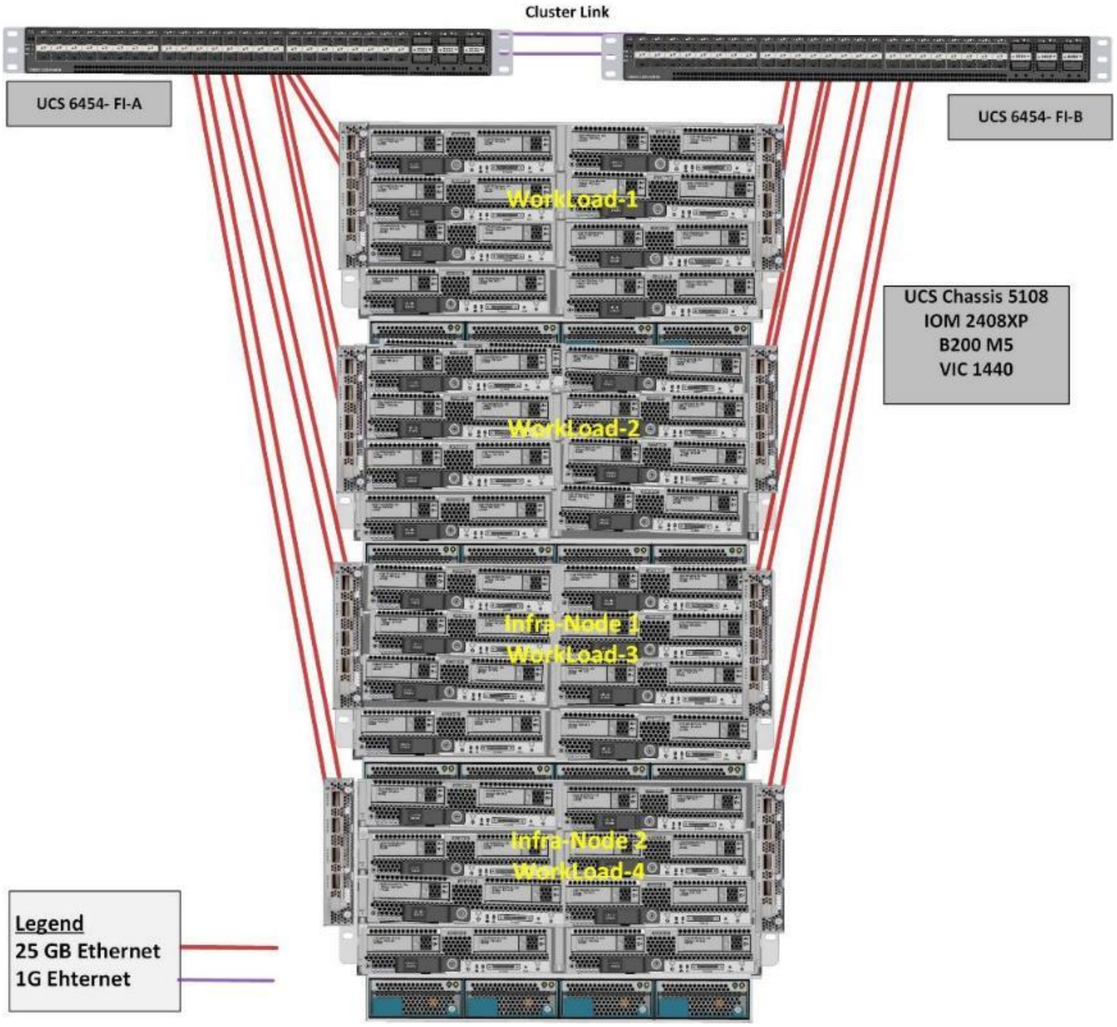

The 2408 IOM is shown with 2x25Gbe ports to delivers to the chassis, full population of the 2408 IOM can support 8x25Gbe ports, allowing for an aggregate of 200Gbe to the chassis.

Network Connectivity

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

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

The upstream network switches can connect to the Cisco UCS 6454 Fabric Interconnects using 10G, 25G, 40G, or 100G port speeds. In this design, the 100G ports from the 40/100G ports on the 6454 (1/49-54) were used for the virtual port channels.

Figure 3. **Network Connectivity**

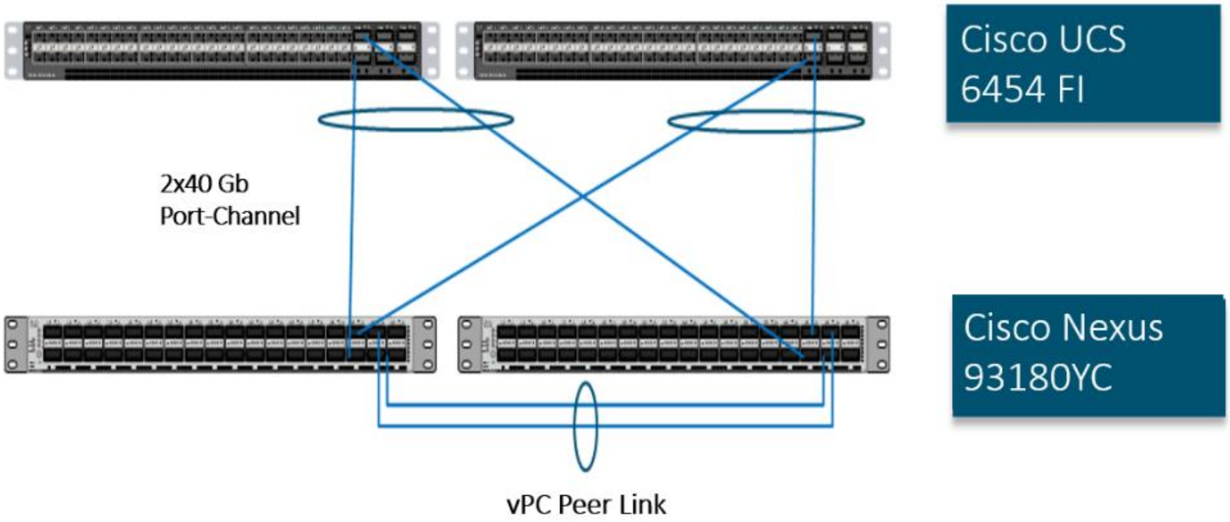

Fibre Channel Storage Connectivity

The Pure Storage FlashArray//X70 R3 platform is connected through both MDS 9132Ts to their respective Fabric Interconnects in a traditional air-gapped A/B fabric design. The Fabric Interconnects are configured in N-Port Virtualization (NPV) mode, known as FC end host mode in UCSM. The MDS has N-Port ID Virtualization (NPIV) enabled. This allows F-port channels to be used between the Fabric Interconnect and the MDS, providing the following benefits:

- Increased aggregate bandwidth between the fabric interconnect and the MDS
- Load balancing across the FC uplinks
- High availability in the event of a failure of one or more uplinks

Figure 4. **Fibre Channel Storage Connectivity**

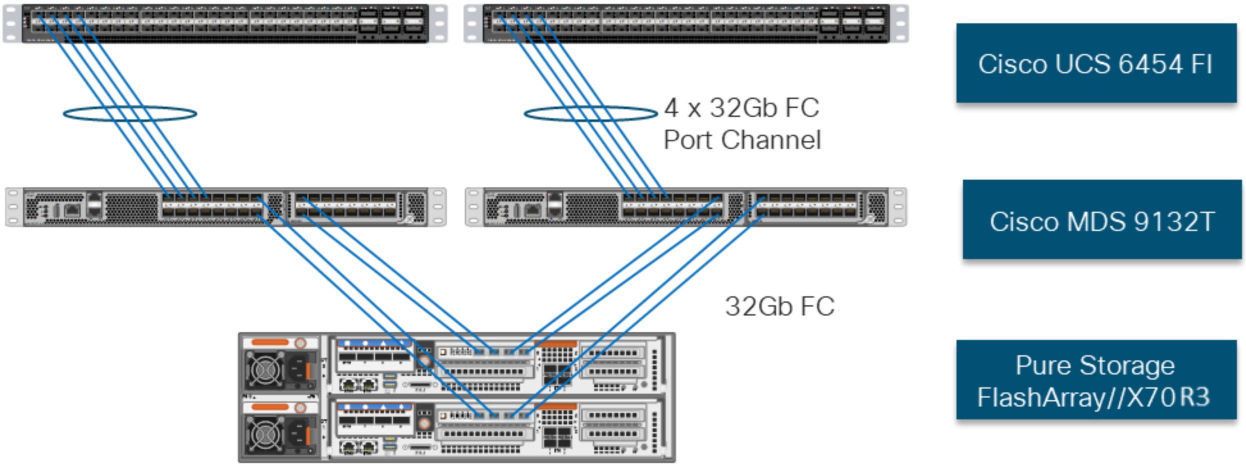

End-to-End Physical Connectivity

FC End-to-End Data Path

The FC end-to-end path in the design is a traditional air-gapped fabric with identical data path through each fabric as detailed below:

- Each Cisco UCS Server is equipped with a Cisco UCS VIC 1400 Series adapter
- In the Cisco B200 M5 server, a VIC 1440 provides 2x25Gbe to IOM A and 2x25Gbe to IOM B through the Cisco UCS Chassis 5108 chassis backplane
- Each IOM is connected to its respective Cisco UCS 6454 Fabric Interconnect using a port-channel for 4-8 links
- Each Cisco UCS 6454 FI connects to the MDS 9132T for the respective SAN fabric using an F-Port channel
- The Pure Storage FlashArray//X70 R3 is connected to both MDS 9132T switches to provide redundant paths through both fabrics

Figure 5. FC End-to-End Data Path

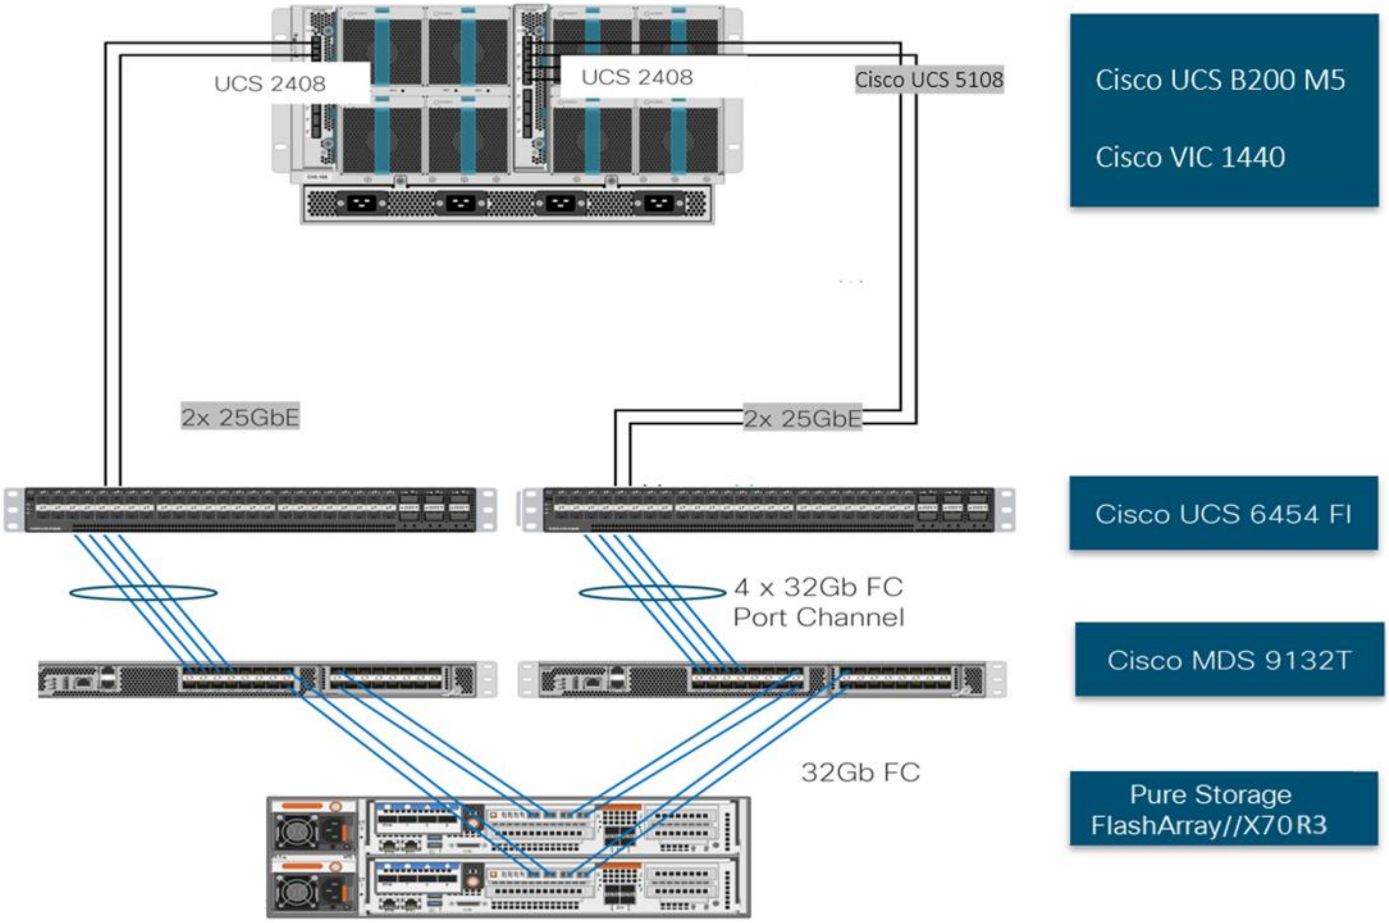

The components of this integrated architecture shown in Figure 5 are:

- Cisco Nexus 93180YC-FX – 10/25/40/100Gbe capable, LAN connectivity to the Cisco UCS compute resources
- Cisco UCS 6454 Fabric Interconnect – Unified management of Cisco UCS compute, and the compute's access to storage and networks
- Cisco UCS B200 M5 – High powered blade server, optimized for virtual computing
- Cisco MDS 9132T – 32Gb Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center
- Pure Storage FlashArray//X70 R3

High Scale RDS and VDI Workload Solution Reference Architecture

Figure 6 illustrates the FlashStack System architecture used in this Validated Design to support very high scale mixed desktop user workload. It follows Cisco configuration requirements to deliver highly available and scalable architecture.

Figure 6. FlashStack Solution Reference Architecture

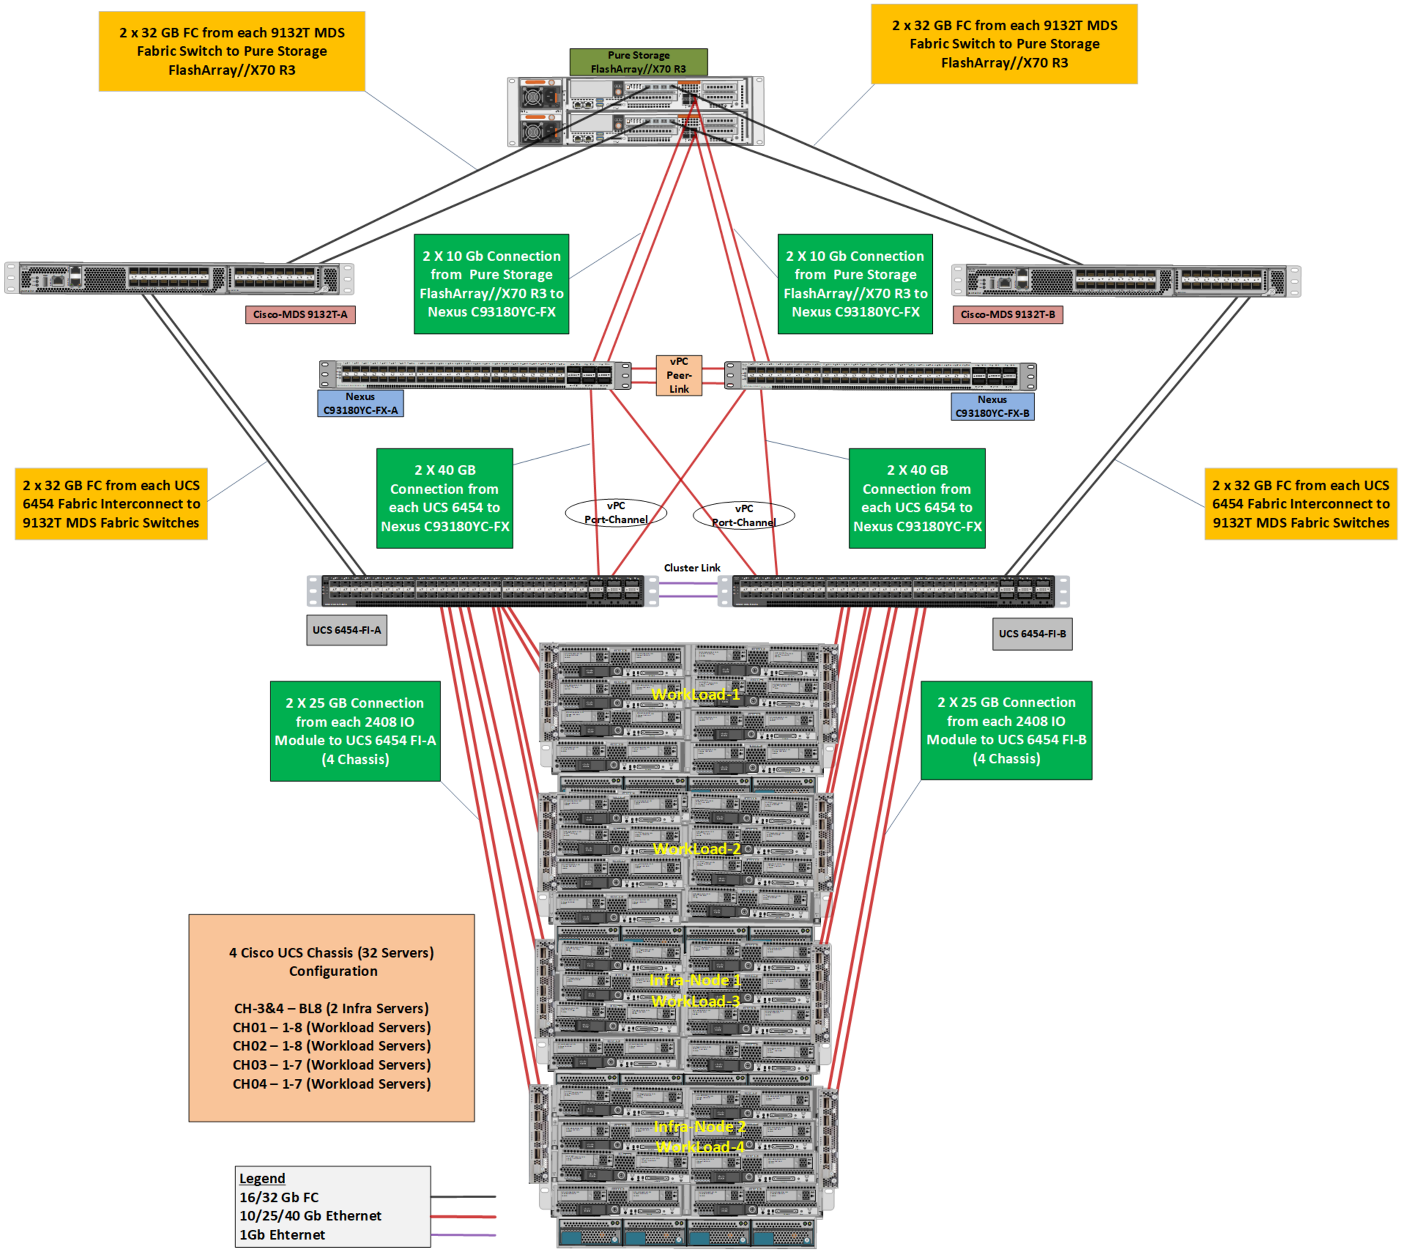
2 x 32 GB FC from each 9132T MDS Fabric Switch to Pure Storage FlashArray//X70 R3
Pure Storage FlashArray//X70 R3
2 x 32 GB FC from each 9132T MDS Fabric Switch to Pure Storage FlashArray//X70 R3
2 X 10 Gb Connection from Pure Storage FlashArray//X70 R3 to Nexus C93180YC-FX
2 X 10 Gb Connection from Pure Storage FlashArray//X70 R3 to Nexus C93180YC-FX
Cisco-MDS 9132T-A
Cisco-MDS 9132T-B
vPC Peer-Link
Nexus C93180YC-FX-A
Nexus C93180YC-FX-B
2 x 32 GB FC from each UCS 6454 Fabric Interconnect to 9132T MDS Fabric Switches
2 X 40 GB Connection from each UCS 6454 to Nexus C93180YC-FX
vPC Port-Channel
vPC Port-Channel
2 X 40 GB Connection from each UCS 6454 to Nexus C93180YC-FX
2 x 32 GB FC from each UCS 6454 Fabric Interconnect to 9132T MDS Fabric Switches
Cluster Link
UCS 6454-FI-A
UCS 6454-FI-B
2 X 25 GB Connection from each 2408 IO Module to UCS 6454 FI-A (4 Chassis)
WorkLoad-1
2 X 25 GB Connection from each 2408 IO Module to UCS 6454 FI-B (4 Chassis)
WorkLoad-2
4 Cisco UCS Chassis (32 Servers) Configuration
CH-3&4 – BL8 (2 Infra Servers)
CH01 – 1-8 (Workload Servers)
CH02 – 1-8 (Workload Servers)
CH03 – 1-7 (Workload Servers)
CH04 – 1-7 (Workload Servers)
Infra-Node 1
WorkLoad-3
Infra-Node 2
WorkLoad-4
Legend
16/32 Gb FC
10/25/40 Gb Ethernet
1Gb Ehternet

The reference hardware configuration includes:

- 2 Cisco Nexus 93180YC-FX switches
- 2 Cisco MDS 9132T 32-Gb Fibre Channel switches
- 2 Cisco UCS 6454 Fabric Interconnects
- 4 Cisco UCS 5108 Blade Chassis
- 2 Cisco UCS B200 M5 Blade Servers (2 Server hosting Infrastructure virtual machines)

- 30 Cisco UCS B200 M5 Blade Servers (for workload)
- 1 Pure Storage FlashArray//X70 R3 with All-NVMe DirectFlash Modules

For desktop virtualization, the deployment includes VMware Horizon 8 running on VMware vSphere ESXi 7.0 GA.

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

- 6000 Random RDS Windows 2019 user sessions with Office 2019 (Full Clones) on 30 UCS Hosts
- 5000 Random Windows 10 Instant Clone Desktops with Office 2019 on 30 UCS Hosts
- 5000 Random Windows 10 Full Clone Desktops with Office 2019 on 30 UCS Hosts

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 FlashStack?

The FlashStack platform, developed by Cisco and Pure Storage, is a flexible, integrated infrastructure solution that delivers pre-validated storage, networking, and server technologies. Cisco and Pure Storage have carefully validated and verified the FlashStack solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model.

FlashStack is a best practice data center architecture that includes the following components:

- Cisco Unified Computing System
- Cisco Nexus Switches
- Cisco MDS Switches
- Pure Storage FlashArray

Figure 7. **FlashStack Systems Components**

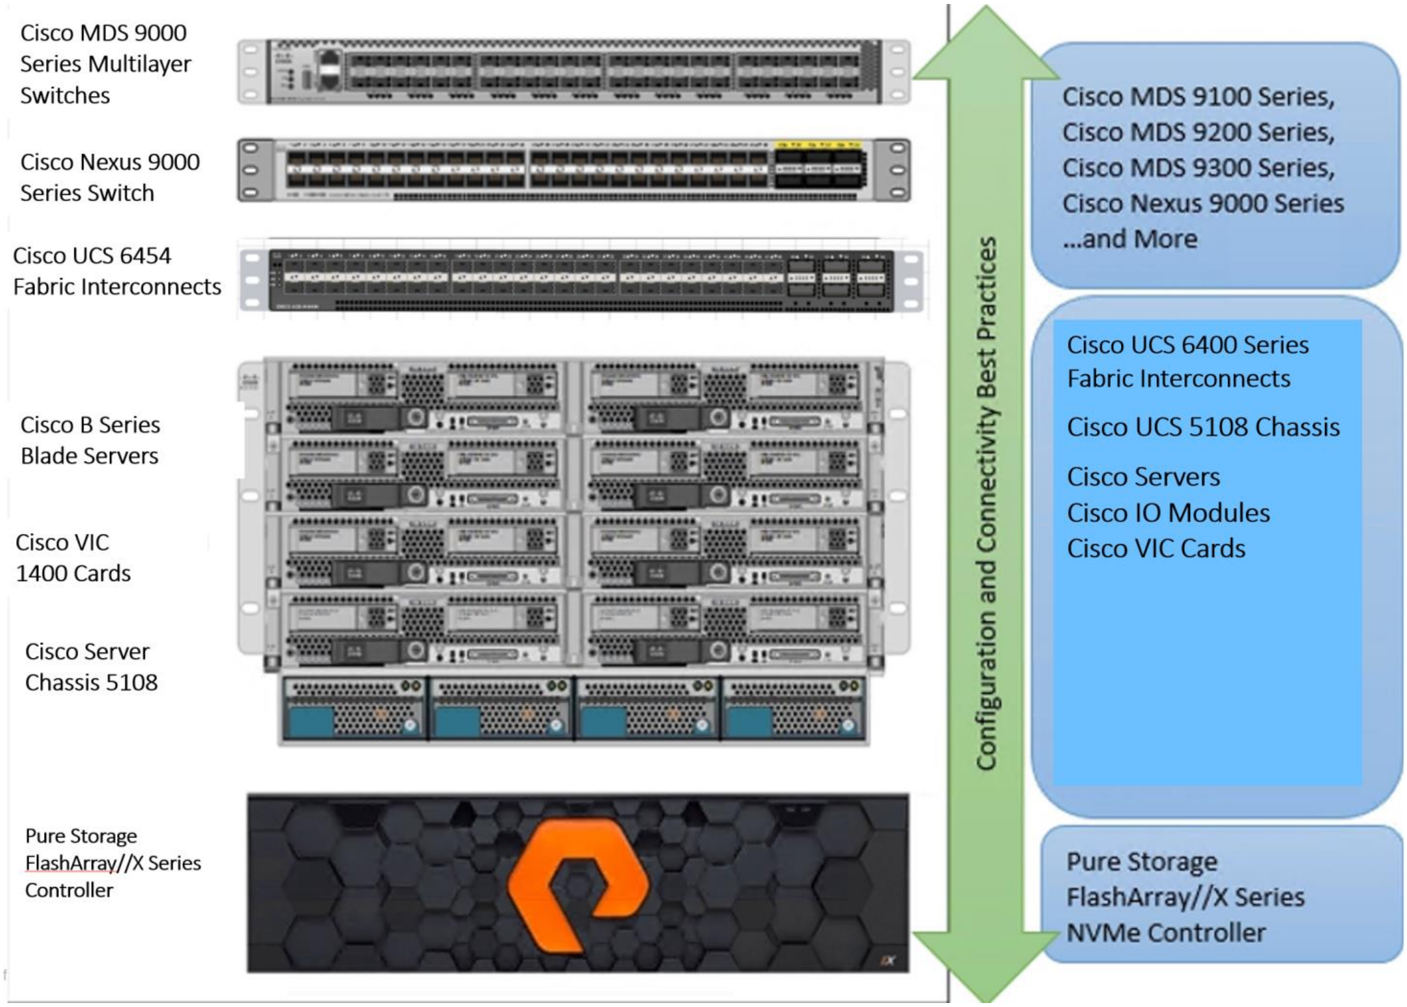
Cisco MDS 9000 Series Multilayer Switches
Cisco Nexus 9000 Series Switch
Cisco UCS 6454 Fabric Interconnects
Cisco B Series Blade Servers
Cisco VIC 1400 Cards
Cisco Server Chassis 5108
Pure Storage FlashArray//X Series Controller
Configuration and Connectivity Best Practices
Cisco MDS 9100 Series,
Cisco MDS 9200 Series,
Cisco MDS 9300 Series,
Cisco Nexus 9000 Series
...and More
Cisco UCS 6400 Series Fabric Interconnects
Cisco UCS 5108 Chassis
Cisco Servers
Cisco IO Modules
Cisco VIC Cards
Pure Storage FlashArray//X Series NVMe Controller

As shown in Figure 7, these components are connected and configured according to best practices of both Cisco and Pure Storage and provide the ideal platform for running a variety of enterprise database workloads with confidence. FlashStack 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.

The reference architecture covered in this document leverages the Pure Storage FlashArray//X70 R3 Controller with NVMe based DirectFlash modules for Storage, Cisco UCS B200 M5 Blade Server for Compute, Cisco Nexus 9000, and Cisco MDS 9100 Series for the switching element and Cisco Fabric Interconnects 6300 Series for System Management. As shown in Figure 7. , FlashStack Architecture can maintain consistency at scale. Each of the component families shown in (Cisco UCS, Cisco Nexus, Cisco MDS, Cisco FI and Pure Storage) 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 FlashStack.

FlashStack Solution Benefits

FlashStack provides a jointly supported solution by Cisco and Pure Storage. Bringing a carefully validated architecture built on superior compute, world-class networking, and the leading innovations in all flash storage. The portfolio of validated offerings from FlashStack includes but is not limited to the following:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100 percent NVMe enterprise flash storage
  - Consolidate hundreds of enterprise-class applications in a single rack
  - Scalability through a design for hundreds of discrete servers and thousands of virtual machines, and the capability to scale I/O bandwidth to match demand without disruption
  - Repeatable growth through multiple FlashStack CI deployments
- Operational Simplicity
  - Fully tested, validated, and documented for rapid deployment
  - Reduced management complexity
  - No storage tuning or tiers necessary
  - 3x better data reduction without any performance impact
- Lowest TCO
  - Dramatic savings in power, cooling and space with Cisco UCS and 100 percent Flash
  - Industry leading data reduction
  - Free FlashArray controller upgrades every three years with Forever Flash™
- Mission Critical and Enterprise Grade Resiliency
  - Highly available architecture with no single point of failure
  - Non-disruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication

Cisco and Pure Storage have also built a robust and experienced support team focused on FlashStack solutions, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance between Pure Storage and Cisco gives customers and channel services partners direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

What's New in this FlashStack Release

This CVD of the FlashStack release introduces new hardware with the Pure Storage FlashArray//X, that is 100 percent NVMe enterprise class all-flash array along with Cisco UCS B200 M5 Blade Servers featuring the Intel Xeon Scalable Family of CPUs. This is the second Oracle RAC Database deployment Cisco Validated Design with Pure Storage. It incorporates the following features:

- Pure Storage FlashArray//X70 R3 Purity//FA 6.0.3
- Cisco 4th Gen UCS 6454 with IOM 2408
- Cisco UCS Manager 4.1(2a)

- VMware vSphere ESXi 7.0 GA Hypervisor
- VMware Horizon 8
- VMware DEM Enterprise 10.0

Configuration Guidelines

This Cisco Validated Design provides the details to deploy a highly available 6000/5000/5000 seat RDS/VDI virtual desktop solution with VMware on a FlashStack Data Center architecture. Configuration guidelines are provided that refer the reader to which redundant component is being configured with each step.

Redundancy built-in the entire infrastructure is as follows:

- Storage Redundancy: FlashArray//X70 R3 Controller 0 and Controller 1
- Switching Redundancy: Cisco Nexus A and Cisco Nexus B
- SAN Switch redundancy: Cisco MDS A and Cisco MDS B
- Compute Redundancy: Cisco UCS 6454 FI- A and FI -B
- Compute Server redundancy: N+1
- Infrastructure Server redundancy: N+1

Additionally, this document explains the steps to provision multiple Cisco UCS hosts, and these are identified sequentially: Rack-Infra-01, Rack-Infra-02, Rack-WLHost-01, Rack-WLHost -02 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

Solution Components

This section describes the components used in the solution outlined in this solution.

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 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, 25-Gbe unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. 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 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 8. Cisco Data Center Overview

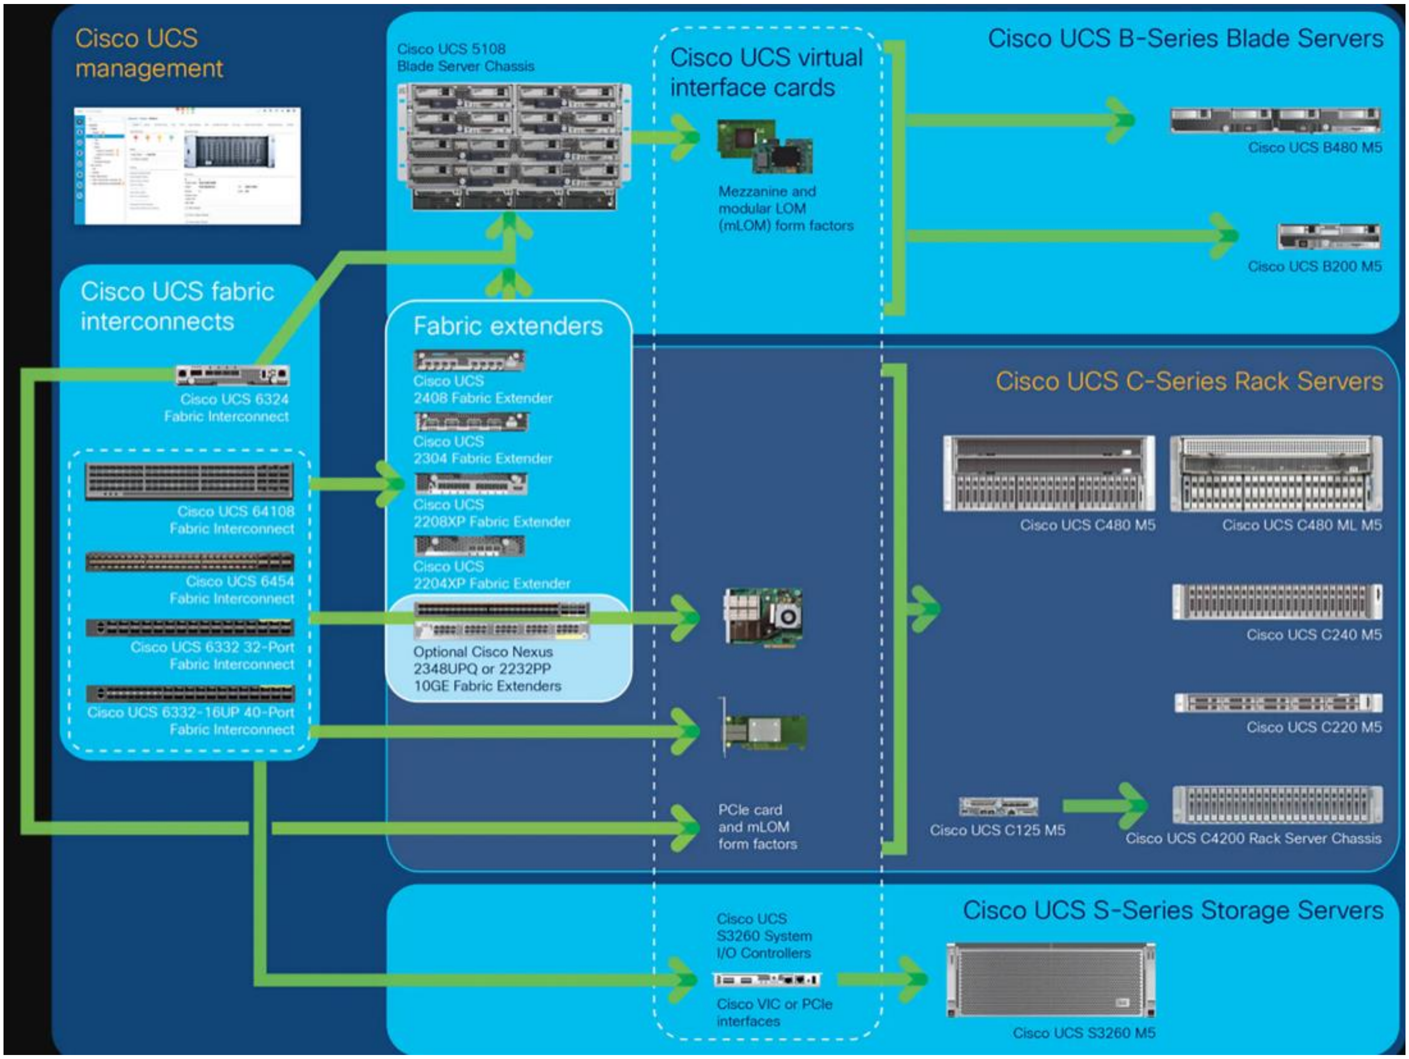
Cisco UCS management
Cisco UCS fabric interconnects
Cisco UCS 6324 Fabric Interconnect
Cisco UCS 64108 Fabric Interconnect
Cisco UCS 6454 Fabric Interconnect
Cisco UCS 6332 32-Port Fabric Interconnect
Cisco UCS 6332-16UP 40-Port Fabric Interconnect
Cisco UCS 5108 Blade Server Chassis
Fabric extenders
Cisco UCS 2408 Fabric Extender
Cisco UCS 2304 Fabric Extender
Cisco UCS 2208XP Fabric Extender
Cisco UCS 2204XP Fabric Extender
Optional Cisco Nexus 2348UPQ or 2232PP 10GE Fabric Extenders
Cisco UCS virtual interface cards
Mezzanine and modular LOM (mLOM) form factors
PCIe card and mLOM form factors
Cisco UCS S3260 System I/O Controllers
Cisco VIC or PCIe interfaces
Cisco UCS B-Series Blade Servers
Cisco UCS B480 M5
Cisco UCS B200 M5
Cisco UCS C-Series Rack Servers
Cisco UCS C480 M5
Cisco UCS C480 ML M5
Cisco UCS C240 M5
Cisco UCS C220 M5
Cisco UCS C125 M5
Cisco UCS C4200 Rack Server Chassis
Cisco UCS S-Series Storage Servers
Cisco UCS S3260 M5

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 a 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 offer 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 provide the management and communication backbone for the Cisco UCS B-Series Blade Servers, UCS 5108 B-Series Server Chassis, UCS Managed C-Series Rack Servers, and 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 Gbe 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 9. Cisco UCS 6400 Series Fabric Interconnect – 6454 Front View

Figure 10. Cisco UCS 6400 Series Fabric Interconnect – 6454 Rear View

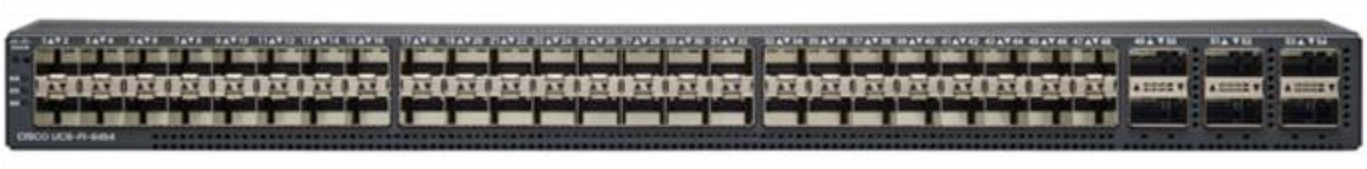

Cisco UCS B200 M5 Blade Server

The Cisco UCS B200 M5 Blade Server (Figure 11 and Figure 12) is a density-optimized, half-width blade server that supports two CPU sockets for Intel Xeon processor 6230 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 11. Cisco UCS B200 M5 Front View

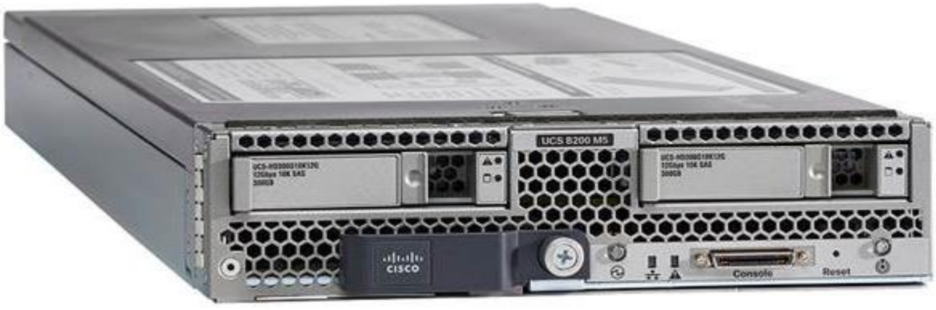

Figure 12. Cisco UCS B200 M5 Back View

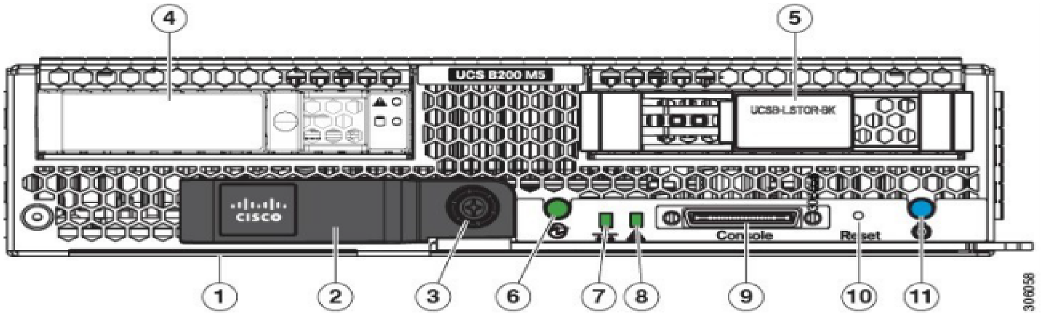

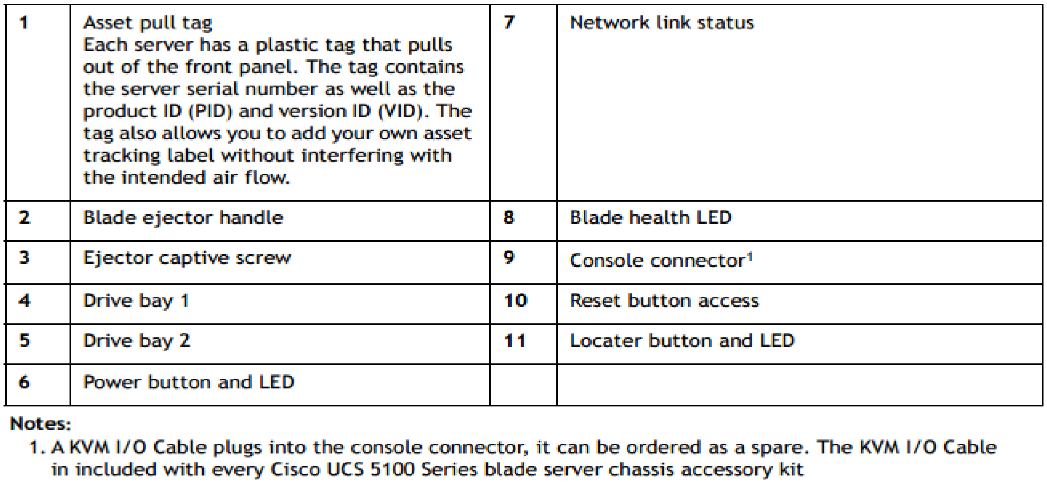

1	Asset pull tag Each server has a plastic tag that pulls out of the front panel. The tag contains the server serial number as well as the product ID (PID) and version ID (VID). The tag also allows you to add your own asset tracking label without interfering with the intended air flow.	7	Network link 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 in included with every Cisco UCS 5100 Series blade server chassis accessory kit

Cisco UCS combines Cisco UCS B-Series Blade Servers and 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 Gold 6230 product family, it offers up to 3 TB of memory using 128GB DIMMs, up to two disk drives, and up to 320

GB 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 Gbe 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) 1440 or 1340, a 2-port, 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable mLOM mezzanine adapter
- Optional rear mezzanine VIC with two 40-Gbe unified I/O ports or two sets of 4 x 10-Gbe unified I/O ports, delivering 80 Gbe to the server; adapts to either 10- or 40-Gbe 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
- Support for one 16-GB internal flash USB drive

For more information about Cisco UCS B200 M5, see the Cisco UCS B200 M5 Blade Server Specsheet.

Table 1. Ordering Information

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

Cisco UCS VIC1440 Converged Network Adapter

The Cisco UCS VIC 1440 (Figure 13) is a single-port 40-Gbe or 4x10-Gbe 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 1440 capabilities are enabled for two ports of 40-Gbe Ethernet. The Cisco UCS VIC 1440 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

Figure 13. Cisco UCS VIC 1440

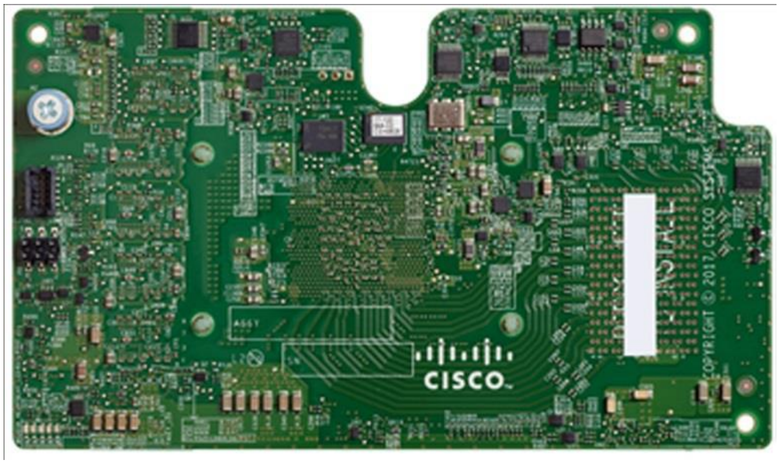

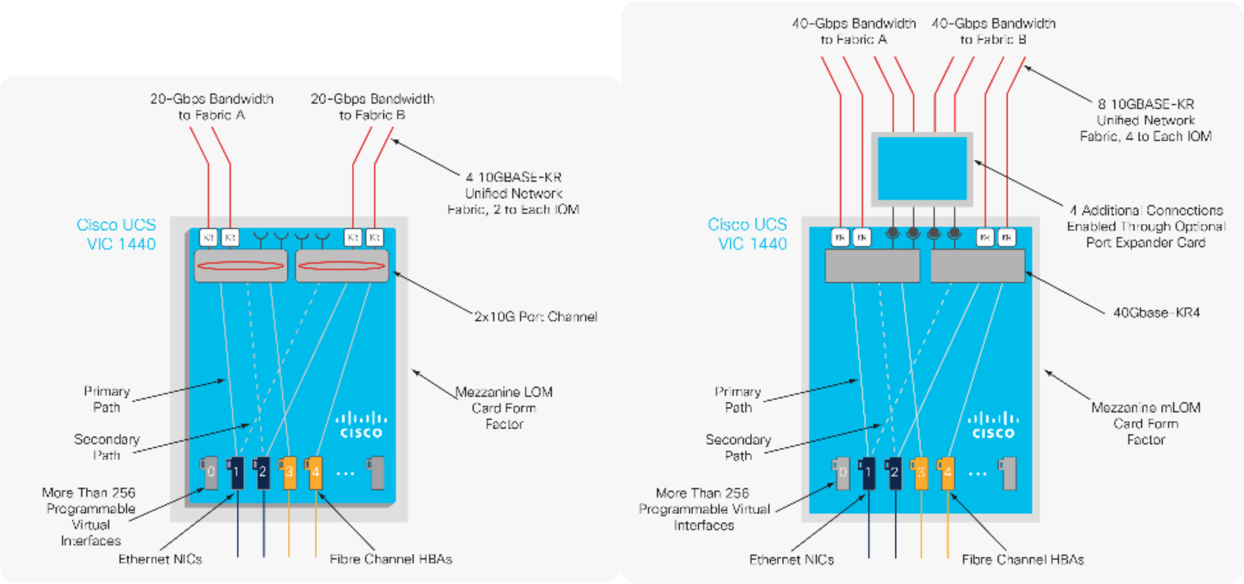
20-Gbps Bandwidth to Fabric A
20-Gbps Bandwidth to Fabric B
4 10GBASE-KR Unified Network Fabric, 2 to Each IOM
Cisco UCS VIC 1440
2x10G Port Channel
Mezzanine LOM Card Form Factor
Primary Path
Secondary Path
More Than 256 Programmable Virtual Interfaces
Ethernet NICs
Fibre Channel HBAs
0
1
2
3
4
...
cisco
40-Gbps Bandwidth to Fabric A
40-Gbps Bandwidth to Fabric B
8 10GBASE-KR Unified Network Fabric, 4 to Each IOM
Cisco UCS VIC 1440
4 Additional Connections Enabled Through Optional Port Expander Card
40Gbase-KR4
Mezzanine mLOM Card Form Factor
Primary Path
Secondary Path
More Than 256 Programmable Virtual Interfaces
Ethernet NICs
Fibre Channel HBAs
0
1
2
3
4
...
cisco

Figure 13 illustrates the Cisco UCS VIC 1440 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M5 Blade Servers.

Cisco Switching

Cisco Nexus 93180YC-FX Switches

The Cisco Nexus 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 data centers.

- Architectural Flexibility
  - Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
  - Leaf node support for Cisco ACI architecture is provided in the roadmap
  - Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support
- Feature Rich
  - Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
  - ACI-ready infrastructure helps users take advantage of automated policy-based systems management
  - Virtual Extensible LAN (VXLAN) routing provides network services
  - Rich traffic flow telemetry with line-rate data collection
  - Real-time buffer utilization per port and per queue, for monitoring traffic micro-bursts and application traffic patterns
- Highly Available and Efficient Design
  - High-density, non-blocking architecture

    - Easily deployed into either a hot-aisle and cold-aisle configuration
    - Redundant, hot-swappable power supplies and fan trays
- Simplified Operations
    - Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
    - An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
    - Python Scripting for programmatic access to the switch command-line interface (CLI)
    - Hot and cold patching, and online diagnostics
- Investment Protection

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

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

Figure 14. Cisco Nexus 93180YC-EX Switch

Cisco MDS 9132T 32-Gb Fiber Channel Switch

The next-generation Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switch (Figure 15) provides high-speed Fibre Channel connectivity from the server rack to the SAN core. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the dual benefits of greater bandwidth and consolidation.

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

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

Additionally, investing in this switch for the lower-speed (4- or 8- or 16-Gb) server rack gives you the option to upgrade to 32-Gb server connectivity in the future using the 32-Gb Host Bus Adapter (HBA) that are available today. The Cisco® MDS 9132T 32-Gb 32-Port Fibre Channel switch also provides unmatched flexibility through a unique port expansion module (Figure 16.) 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 15. Cisco 9132T 32-Gb MDS Fibre Channel Switch

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

- Features
  - ◦ High performance: MDS 9132T architecture, with chip-integrated nonblocking arbitration, provides consistent 32-Gb low-latency performance across all traffic conditions for every Fibre Channel port on the switch.
  - ◦ Capital Expenditure (CapEx) savings: The 32-Gb ports allow users to deploy them on existing 16- or 8-Gb transceivers, reducing initial CapEx with an option to upgrade to 32-Gb transceivers and adapters in the future.
  - ◦ High availability: MDS 9132T switches continue to provide the same outstanding availability and reliability as the previous-generation Cisco MDS 9000 Family switches by providing optional redundancy on all major components such as the power supply and fan. Dual power supplies also facilitate redundant power grids.
  - ◦ Pay-as-you-grow: The MDS 9132T Fibre Channel switch provides an option to deploy as few as eight 32-Gb Fibre Channel ports in the entry-level variant, which can grow by 8 ports to 16 ports, and thereafter with a port expansion module with sixteen 32-Gb ports, to up to 32 ports. This approach results in lower initial investment and power consumption for entry-level configurations of up to 16 ports compared to a fully loaded switch. Upgrading through an expansion module also reduces the overhead of

managing multiple instances of port activation licenses on the switch. This unique combination of port upgrade options allow four possible configurations of 8 ports, 16 ports, 24 ports and 32 ports.

- Next-generation Application-Specific Integrated Circuit (ASIC): The MDS 9132T Fibre Channel switch is powered by the same high-performance 32-Gb Cisco ASIC with an integrated network processor that powers the Cisco MDS 9700 48-Port 32-Gb Fibre Channel Switching Module. Among all the advanced features that this ASIC enables, one of the most notable is inspection of Fibre Channel and Small Computer System Interface (SCSI) headers at wire speed on every flow in the smallest form-factor Fibre Channel switch without the need for any external taps or appliances. The recorded flows can be analyzed on the switch and also exported using a dedicated 10/100/1000BASE-T port for telemetry and analytics purposes.
- Intelligent network services: Slow-drain detection and isolation, VSAN technology, Access Control Lists (ACLs) for hardware-based intelligent frame processing, 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.

Purity for FlashArray

The essential element of every FlashArray is the Purity Operating Environment software. Purity implements advanced data reduction, storage management, and flash management features, enabling organizations to enjoy Tier 1 data services for all workloads, proven 99.9999% availability over multiple years (inclusive of maintenance and generational upgrades), completely non-disruptive operations, 2X better data reduction versus alternative all-flash solutions, and – with FlashArray//X – the power and efficiency of DirectFlash™.

Moreover, Purity includes enterprise-grade data security, modern data protection options, and complete business continuity and global disaster recovery through ActiveCluster multi-site stretch cluster and ActiveDR* for continuous replication with near zero RPO. All these features are included with every array.

FlashArray File Services

Pure Storage acquired Compuverde last year, and they've been busy at work integrating this technology into the Purity//FA operating system. They emphasize the "integrating", because they didn't just take the existing product, drop it onto a FlashArray system, and run it on top of Purity. Instead, they incorporated key parts of it into Purity to give you the advantages of native files alongside blocks.

The SMB and NFS protocols bring consolidated storage to the Purity//FA operating system, complementing its block capabilities, while the file system offers features like directory snapshots and directory-level performance and space monitoring. For the purposes of this reference architecture, we will be focusing on using File Services for User Profile management.

Figure 17. FlashArray//X Specifications

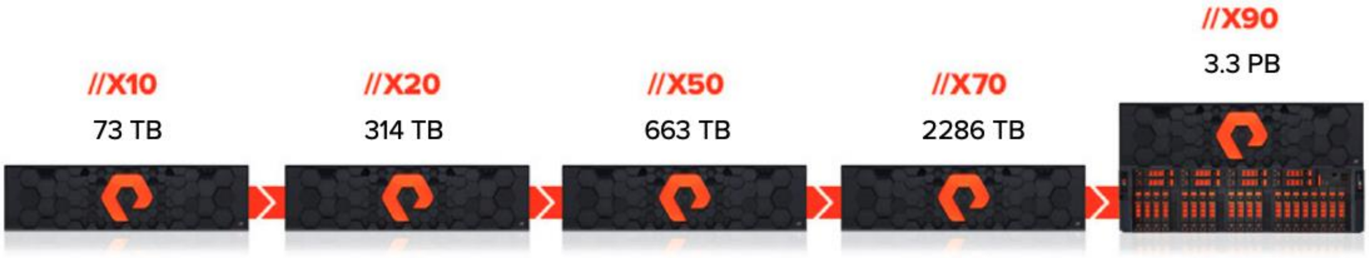

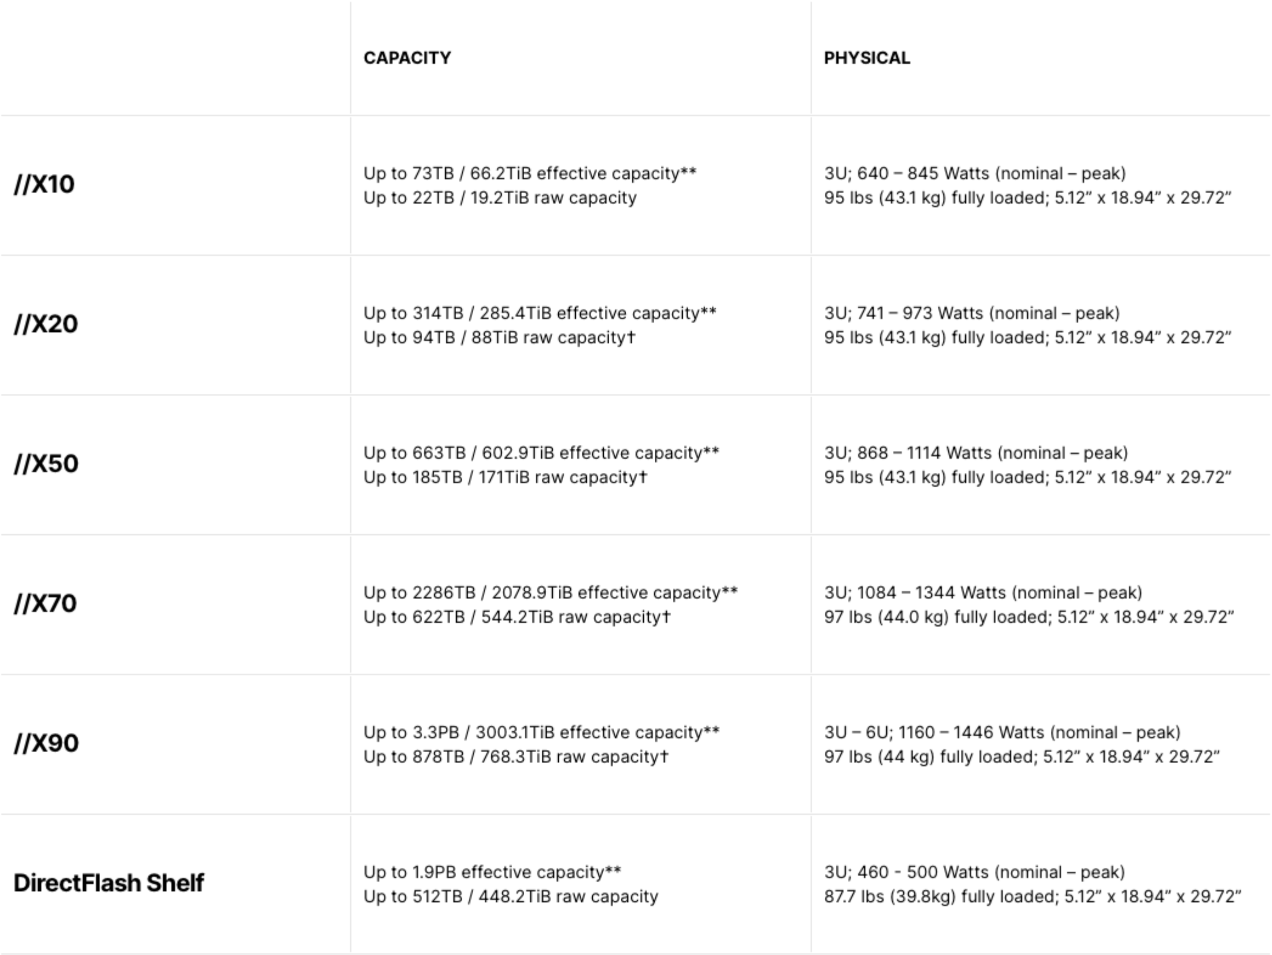

	CAPACITY	PHYSICAL
//X10	Up to 73TB / 66.2TiB effective capacity** Up to 22TB / 19.2TiB raw capacity	3U; 640 – 845 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12” x 18.94” x 29.72”
//X20	Up to 314TB / 285.4TiB effective capacity** Up to 94TB / 88TiB raw capacity†	3U; 741 – 973 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12” x 18.94” x 29.72”
//X50	Up to 663TB / 602.9TiB effective capacity** Up to 185TB / 171TiB raw capacity†	3U; 868 – 1114 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12” x 18.94” x 29.72”
//X70	Up to 2286TB / 2078.9TiB effective capacity** Up to 622TB / 544.2TiB raw capacity†	3U; 1084 – 1344 Watts (nominal – peak) 97 lbs (44.0 kg) fully loaded; 5.12” x 18.94” x 29.72”
//X90	Up to 3.3PB / 3003.1TiB effective capacity** Up to 878TB / 768.3TiB raw capacity†	3U – 6U; 1160 – 1446 Watts (nominal – peak) 97 lbs (44 kg) fully loaded; 5.12” x 18.94” x 29.72”
DirectFlash Shelf	Up to 1.9PB effective capacity** Up to 512TB / 448.2TiB raw capacity	3U; 460 - 500 Watts (nominal – peak) 87.7 lbs (39.8kg) fully loaded; 5.12” x 18.94” x 29.72”

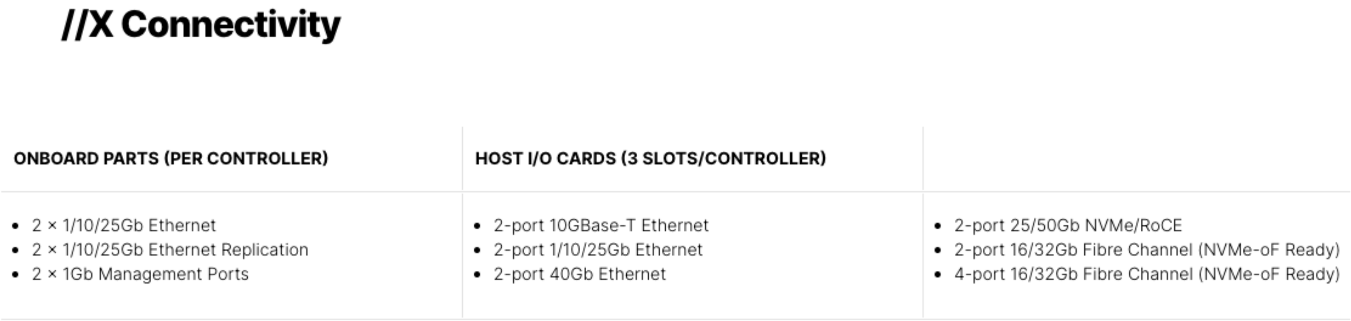

//X Connectivity

ONBOARD PARTS (PER CONTROLLER)	HOST I/O CARDS (3 SLOTS/CONTROLLER)	
• 2 × 1/10/25Gb Ethernet • 2 × 1/10/25Gb Ethernet Replication • 2 × 1Gb Management Ports	• 2-port 10GBase-T Ethernet • 2-port 1/10/25Gb Ethernet • 2-port 40Gb Ethernet	• 2-port 25/50Gb NVMe/RoCE • 2-port 16/32Gb Fibre Channel (NVMe-oF Ready) • 4-port 16/32Gb Fibre Channel (NVMe-oF Ready)

** Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1 and does not include thin provisioning or snapshots.

† Array accepts Pure Storage DirectFlash Shelf and/or Pure Storage SAS-based expansion shelf.

Evergreen™ Storage

Customers can deploy storage once and enjoy a subscription to continuous innovation through Pure's Evergreen Storage ownership model: expand and improve performance, capacity, density, and/or features for 10 years or more – all without downtime, performance impact, or data migrations. Pure has disrupted the industry's 3-5-year rip-and-replace cycle by engineering compatibility for future technologies right into its products, notably nondisruptive capability to upgrade from //M to //X with NVMe, DirectMemory, and NVMe-oF capability.

Pure1

Pure1®, our cloud-based management, analytics, and support platform, expands the self-managing, plug-n-play design of Pure all-flash arrays with the machine learning predictive analytics and continuous scanning of Pure1 Meta™ to enable an effortless, worry-free data platform.

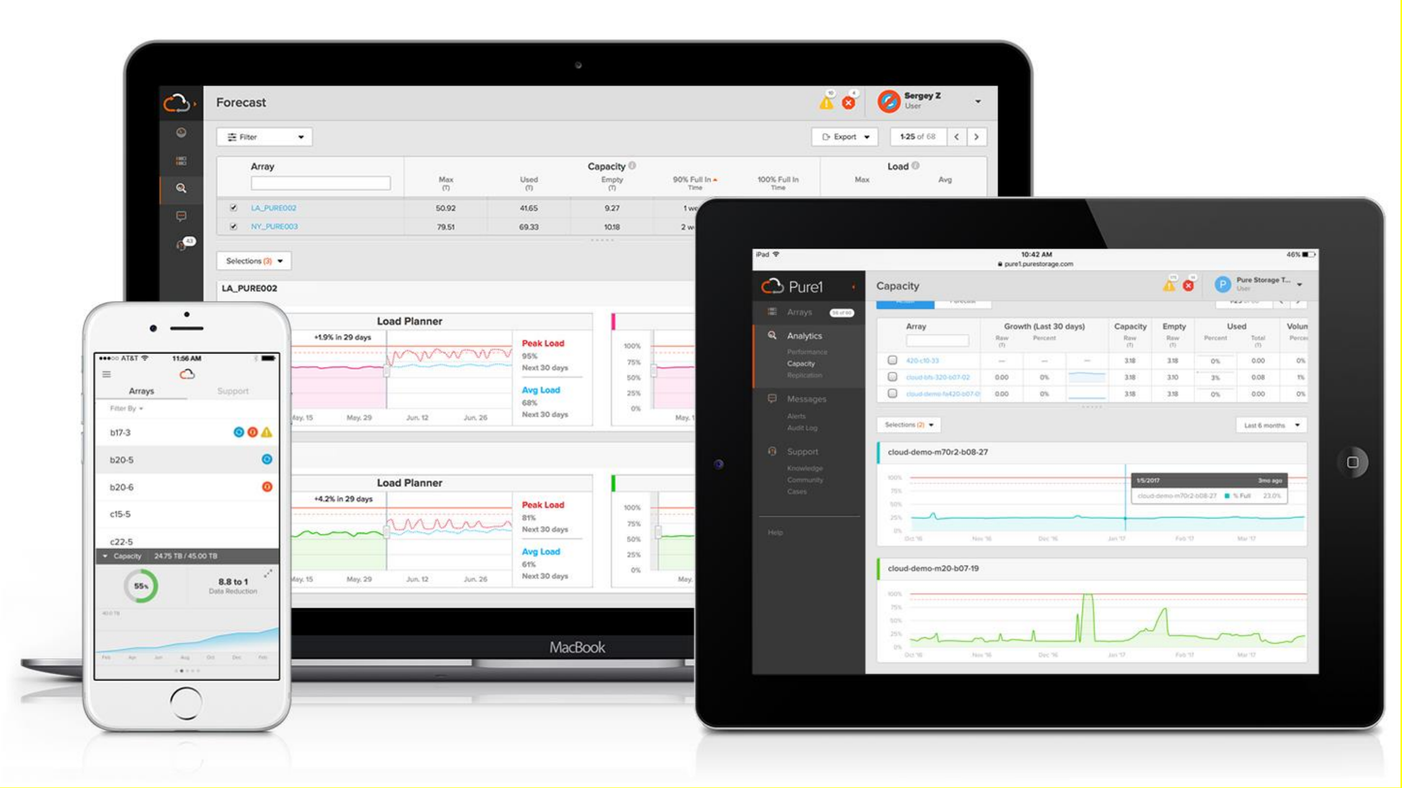

Pure1 Manage

In the Cloud IT operating model, installing, and deploying management software is an oxymoron: you simply log-in. Pure1 Manage is SaaS-based, allowing you to manage your array from any browser or from the Pure1 Mobile App – with nothing extra to purchase, deploy, or maintain. From a single dashboard you can manage all your arrays, with full visibility on the health and performance of your storage.

Pure1 Analyze

Pure1 Analyze delivers true performance forecasting – giving customers complete visibility into the performance and capacity needs of their arrays – now and in the future. Performance forecasting enables intelligent consolidation and unprecedented workload optimization.

Pure1 Support

Pure combines an ultra-proactive support team with the predictive intelligence of Pure1 Meta to deliver unrivaled support that's a key component in our proven FlashArray 99.9999% availability. Customers are often surprised and delighted when we fix issues they did not even know existed.

Pure1 META

The foundation of Pure1 services, Pure1 Meta is global intelligence built from a massive collection of storage array health and performance data. By continuously scanning call-home telemetry from Pure's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues and optimize workloads. The result is both a white glove customer support experience and breakthrough capabilities like accurate performance forecasting.

Meta is always expanding and refining what it knows about array performance and health, moving the Data Platform toward a future of self-driving storage.

Pure1 VM Analytics

Pure1 helps you narrow down the troubleshooting steps in your virtualized environment. VM Analytics provides you with a visual representation of the IO path from the VM all the way through to the FlashArray. Other tools and features guide you through identifying where an issue might be occurring in order to help eliminate potential candidates for a problem.

VM Analytics doesn't only help when there's a problem. The visualization allows you to identify which volumes and arrays particular applications are running on. This brings the whole environment into a more manageable domain.

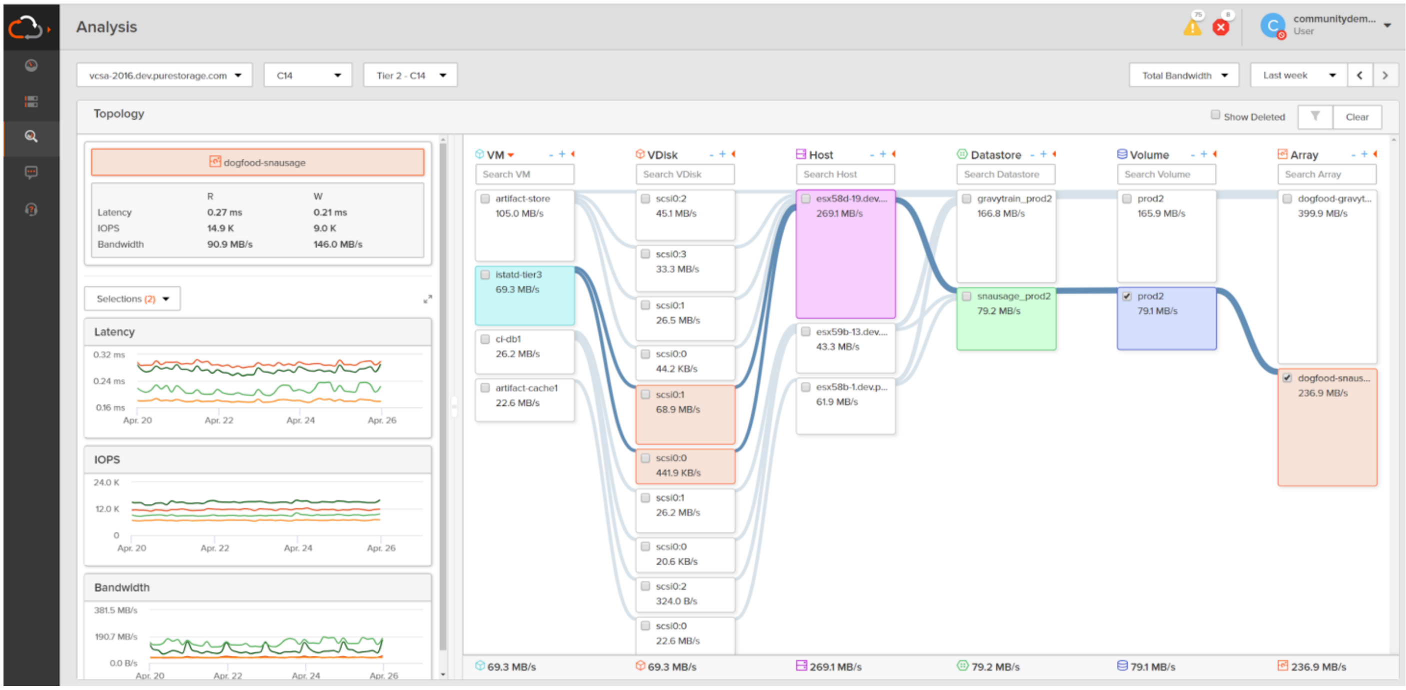

	R	W
Latency	0.27 ms	0.21 ms
IOPS	14.9 K	9.0 K
Bandwidth	90.9 MB/s	146.0 MB/s

Hypervisor

This Cisco Validated Design includes VMware vSphere ESXi 7.0 GA.

VMware vSphere 7.0

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.

vSphere 7 is the latest major vSphere release from VMware. vSphere 7 has been rearchitected with native Kubernetes to enable IT Admins to use vCenter Server® to operate Kubernetes clusters through namespaces. VMware vSphere with Tanzu allows IT Admins to operate with their existing skillset and deliver a self-service access to infrastructure for the Dev Ops teams; while providing observability and troubleshooting for Kubernetes workloads. vSphere 7 provides an enterprise platform for both traditional applications as well as modern applications - so customers and partners can deliver a developer-ready infrastructure, scale without compromise and simplify operations.

Deliver Developer-ready Infrastructure: IT teams can use existing vSphere environments to set up an Enterprise-grade Kubernetes infrastructure at a rapid pace (within one hour), while enabling enterprise-class governance, reliability, and security. After this one-time setup, vSphere with Tanzu enables a simple, fast, and self-service provisioning of Tanzu Kubernetes clusters within a few minutes1. Aligning DevOps teams and IT teams is critical to the success of modern application development; to bring efficiency, scale and security to Kubernetes deployments and operations. vSphere with Tanzu brings agile cloud operations to the IT admin to enable this transition into the role of Cloud Admin or SRE by delivering agility in day-to-day IT operations related to Kubernetes infrastructure.

Scale Without Compromise: vSphere can scale your infrastructure to meet the demands of high-performance applications and memory intensive databases including SAP HANA and Epic Caché Operational Database to name a few. With vSphere 7, a vSphere cluster can now support 50 percent more hosts compared to previous releases.

Simplify Operations: Simplified operations are delivered through key capabilities of vSphere 7 including elastic AI/ML infrastructure for sharing resources, simplified lifecycle management and intrinsic security across your hybrid cloud infrastructure.

The key features and capabilities are as follows:

- TKG Service2: Run the Tanzu Kubernetes Grid Service directly on vSphere to simplify operation of Kubernetes on-premises by putting cloud native constructs at the IT Admin’s fingertips. TKG allows IT admins to manage consistent, compliant, and conformant Kubernetes, while providing developers self-service access to infrastructure. vSphere with Tanzu enables a simple, fast, and self-service provisioning of Tanzu Kubernetes clusters within a few minutes1
- Drop-in to Existing Infrastructure2: Quickly deploy Kubernetes workloads on existing infrastructure with enterprise-grade governance, reliability, and security. Leverage existing networking infrastructure (or BYO networking) using vSphere Distributed Switch's (VDS) centralized interface to configure, monitor and administer switching access for VMs and Kubernetes workloads. Deploy existing block and file storage infrastructure (BYO storage) for containerized workloads. Choose your own L4 load balancing solution using HAProxy (commercial support offered directly by HAProxy) for Tanzu Kubernetes clusters.

- Application focused management2: Kubernetes makes vSphere better by providing DevOps teams (Platform Operators and SREs) with self-service access to infrastructure through Kubernetes APIs. vSphere makes Kubernetes better by empowering IT admins to use vCenter Server skills/tools to operate modern applications, alongside VMs, using namespaces as a unit of management. This is referred to as 'application focused management'. Using application focused management, IT admins can use vCenter Server to observe and troubleshoot Tanzu Kubernetes clusters alongside VMs, implement role-based access and allocate capacity to developer teams.
- Monster VMs: Deliver industry leading scale through Monster VMs designed for SAP HANA and Epic Cache Operational Database. Improve performance and scale for Monster VMs to support your large scale up environments. Scale up to 24TB memory and support up to 768 vCPUs through Monster VMs, leaving other hypervisor vendors far behind in the category. Speed-up the ESXi scheduler and co-scheduling logic for large VMs using selective latency sensitivity setting for workloads, removal of bottlenecks in vCPU sleep/wakeup paths and a reduced memory overhead.
- Cluster scale enhancements: Expand the number of hosts per cluster by 50% to support a total of 96 hosts per cluster, compared to previous releases.
- vLCM enhancements: Simplify software upgrades, patching, and firmware updates for vSphere, vSAN and NSX-T with a single tool. vLCM will also monitor for desired image compliance continuously and enable simple remediation in the event of any compliance drift. ▪ vSphere Ideas®: Submit feature requests right from the vSphere Client UI, track the status of the feature requests and look at all the other feature requests submitted by other users to vote for them, through the Ideas portal.
- vCenter connect®: Manage on-premises and off-premises (cloud providers) vCenter Servers in a single interface using the any to any vCenter connect capability

Additional information about VMware vSphere 7 can be found here.

Desktop Broker

This Cisco Validated Design includes VMware Horizon 8 Enterprise edition.

VMware Horizon

VMware Horizon is a modern platform for secure delivery of virtual desktops and apps across the hybrid cloud. Leveraging best-in-class management capabilities and deep integrations with the VMware technology ecosystem, the Horizon platform delivers a modern approach for desktop and app management that extends from on-premises to the hybrid and multi-cloud. The result is fast and simple virtual desktop and application delivery that extends the best digital workspace experience to all applications.

What's New in VMware Horizon Version 2012

VMware Horizon version 2012 provides the following new features and enhancements.

Horizon Connection Server

- Cloud Pod Architecture
  - When you create a global desktop entitlement, you can select Show Machine Alias Name to display the machine alias name set for the assigned users of the machine instead of the machine host name in Horizon Client. See Worksheet for Configuring a Global Entitlement. Additionally, you can configure the

Show Machine Alias Name option for instant-clone and manual desktop pools. See Worksheet for Creating an Instant-Clone Desktop Pool and Worksheet for Creating a Manual Desktop Pool.

  - The --displayMachineAlias option is added to the lmvutil --createGlobalEntitlement and --updateGlobalEntitlement commands. The --disableDisplayMachineAlias option is added to the lmvutil --updateGlobalEntitlement command. See Creating a Global Entitlement and Modifying a Global Entitlement.

- Published Desktops and Applications
  - You can enable or disable an application pool in Horizon Console. See Enable or Disable an Application Pool.
  - You can select VDS 7.0 as an ephemeral port when creating an automated instant-clone farm. See Worksheet for Creating an Automated Instant-Clone Farm.
- Virtual Desktops
  - You can set a remote machine power policy when creating a dedicated instant-clone desktop pool. See Power Policies for Desktop Pools.
  - You can select VDS 7.0 as an ephemeral port when creating instant clones. See Worksheet for Creating an Instant-Clone Desktop Pool.
  - VMware Update Manager can update ESXi hosts when performing maintenance on instant-clone hosts. See Perform Maintenance on Instant-Clone Hosts.
  - Horizon Administrators need minimum privileges to manage full clones and instant clones. See Minimum vCenter Server Privileges for Managing Full Clones and Instant Clones.
  - You can run virtual machines on Hyper-V hypervisor. See Running Virtual Machines on Hyper-V.
- Horizon Console
  - You can configure how long an idle Horizon Console session continues before the Connection Server session times out. See Global Settings for Client Sessions.
  - You can enable the setting Accept logon as current user to allow Connection Server to accept the user identity and credential information that is passed when users select Log in as current user. See Using the Log In as Current User Feature Available with Windows-Based Horizon Client.
  - You can collect log bundles for troubleshooting connection server, desktop pools, and farms in Horizon Console. See Collect Logs in Horizon Console.
  - The Monitor Events time period filter option All is removed from Horizon Console.
- Event Database
  - Additional columns are added to the event database. After a Connection Server upgrade, you can run DML update scripts to populate the data in these additional columns in the event database. See the VMware Knowledge Base article 80781.
- Horizon Agent
  - The drag and drop, file association, and file copy and paste features are no longer dependent on the client drive redirection feature being enabled. See Managing Access to Client Drive Redirection.
  - Installation of the serial port redirection and scanner redirection features is changed to improve version flexibility when Horizon Client for Windows and Horizon Agent are installed in the same virtual machine.

  - You can configure media optimization for Microsoft Teams, which supports SILK audio codec. See Configuring Media Optimization for Microsoft Teams.
  - You can set log levels and generate log files in a Data Collection Tool (DCT) bundle for remote desktop features. See Collecting Logs for Remote Desktop Features and Components.
  - Internationalization support is added to the VMware Horizon URL Content Redirection extension for all supported browsers.

> You can use the Microsoft Edge for Chromium browser with URL Content Redirection on a Mac. You must install a Web browser extension on the Mac client to use this feature. See Install the URL Content Redirection Helper Extension for Edge.

Horizon Agent for Linux

- Operating Systems

  Horizon Agent for Linux 2012 adds support for the following Linux distributions:

  - Ubuntu 20.04
  - Red Hat Enterprise Linux (RHEL) Workstation 7.9 and 8.3
  - Red Hat Enterprise Linux (RHEL) Server 7.8, 7.9, 8.2, and 8.3
  - CentOS 8.3
  - SUSE Linux Enterprise Desktop (SLED) 15 SP1 and 15 SP2
  - SUSE Linux Enterprise Server (SLES) 15 SP1 and 15 SP2

- Configurable X Display Numbers
  - Two new configuration options in the /etc/vmware/config file, Desktop.displayNumberMax and Desktop.displayNumberMin, let you define the range of X Windows System display numbers to allocate to user sessions. See Setting Options in Configuration Files on a Linux Desktop.

- Display Scaling
  - The Display Scaling feature allows Linux remote desktops and published applications to be displayed using a scale factor that matches the client system's display. This feature is turned off by default. You can enable it by configuring the rdeSvc.allowDisplayScaling option in the /etc/vmware/config file. See Setting Options in Configuration Files on a Linux Desktop.

- DPI Synchronization
  - The DPI Synchronization feature ensures that the DPI setting in a Linux remote session changes to match the DPI setting of the client system. This feature is enabled by default and configured using the DPISyncEnable option in the /etc/vmware/viewagent-custom.conf file. See Setting Options in Configuration Files on a Linux Desktop.

- Support for Unicode Input
  - The RemoteDisplay.allowVMWKeyEvent2Unicode configuration option in the /etc/vmware/config file allows Horizon Agent for Linux to process and display Unicode keyboard input from clients. This feature is enabled by default. See Setting Options in Configuration Files on a Linux Desktop.

- Enhancements to Session Collaboration

  - Horizon Agent for Linux can now remember the names of users invited to join a collaboration session. The next time a client user begins to type the name of an invitee into the Session Collaboration text box, an auto-completion menu appears with a list of selectable user names.
- Optimized Window Resizing for Published Applications
  - New performance enhancements allow client users to resize published application windows without the unwanted artifacts encountered in previous versions of Horizon Agent for Linux. This feature greatly improves the experience of users working in a published application session. Administrators can turn this feature on and off using the rdeSvc.enableOptimizedResize option in the /etc/vmware/config file. This feature is enabled by default. See Setting Options in Configuration Files on a Linux Desktop.
- Horizon GPO Bundle
  - The Do not redirect client printer(s) group policy setting stops client printers from being redirected. This setting is provided for the agent and the Windows client. For the agent setting, see VMware Integrated Printing Policy Settings. For the Windows client setting, see VMware Integrated Printing Settings for Client GPOs.
  - The Do not change default printer group policy setting stops the VMware Integrated Printing feature from changing the default printer in remote sessions. See VMware Integrated Printing Policy Settings.
  - The Printer Name for RDSH Agents group policy name is changed to Printer Name Schema. This setting now applies to virtual desktops as well as published desktops and published applications. See VMware Integrated Printing Policy Settings.
  - The Connect all ports automatically group policy setting connects all COM ports automatically, even if no individual group policy settings are enabled. See VMware View Agent Configuration ADMX Template Settings.
  - The Exclude Automatically Connection Device Family and Exclude Automatically Connection Vid/Pid Device group policy settings enable you to filter the USB devices that are forwarded automatically based on device family or vendor and product ID. For the client group policy settings, see USB Settings for GPOs. For the agent group policy settings, see USB Settings in the Horizon Agent Configuration ADMX Template.
  - You can optimize redirected USB HID devices with the Include HID Optimization Vid/Pid Device group policy setting. See USB Settings in the Horizon Agent Configuration ADMX Template.
  - Group policy settings are reorganized in the VMware View Agent Configuration folder Group Policy Management Editor. See VMware View Agent Configuration ADMX Template Settings.
  - The View Agent Direct-Connection Plug-in Configuration has a new GPO setting Allow NTLM Fallback for Log On As Current User.
  - In the Idle Time Until Disconnect (VDI) group policy setting, you can specify a minimum timeout value of 1 minute and a maximum timeout value of Never after which a desktop session will disconnect due to user inactivity.
  - The DPI Synchronization Per Monitor group policy setting adjusts the DPI settings in all monitors to match the client operating system's DPI setting during a remote session. The DPI Synchronization Per Connection group policy setting has been removed. See VMware View Agent Configuration ADMX Template Settings.

    - The Allow user to skip Horizon Client update, Automatically check for updates and Update message pop-up group policy settings enable you to customize the Horizon Client for Windows online update feature. See General Settings for Client GPOs.
    - For the session collaboration feature, you can enable the Include Outlook-formatted URL in clipboard text group policy setting to include a Microsoft Outlook-formatted invitation URL in clipboard invitation text. See VMware View Agent Configuration ADMX Template Settings.

- Horizon Client

  For information about new features in Horizon Client 2012, including HTML Access 2012, see the release notes on the VMware Horizon Client Documentation page.

- No Longer Supported Features in This Release

  The following features are no longer supported in this release:

    - View Composer - View Composer linked clones and persistent disks are no longer supported.

- Horizon Cloud Connector

  Applicable to VMware Horizon Universal License customers. The Horizon Cloud Connector virtual appliance is a required component for VMware Horizon to support the management of Horizon pods using Horizon Cloud Service.

- Horizon Deployed on VMware Cloud on AWS

  For a list of VMware Horizon features supported on VMware Cloud on AWS, see the VMware Knowledge Base article 58539.

- Horizon Deployed on Azure VMware Solution

  You can select Azure as an installation option to deploy Horizon on Azure VMware Solution (AVS).

Figure 18. Logical Architecture of VMware Horizon

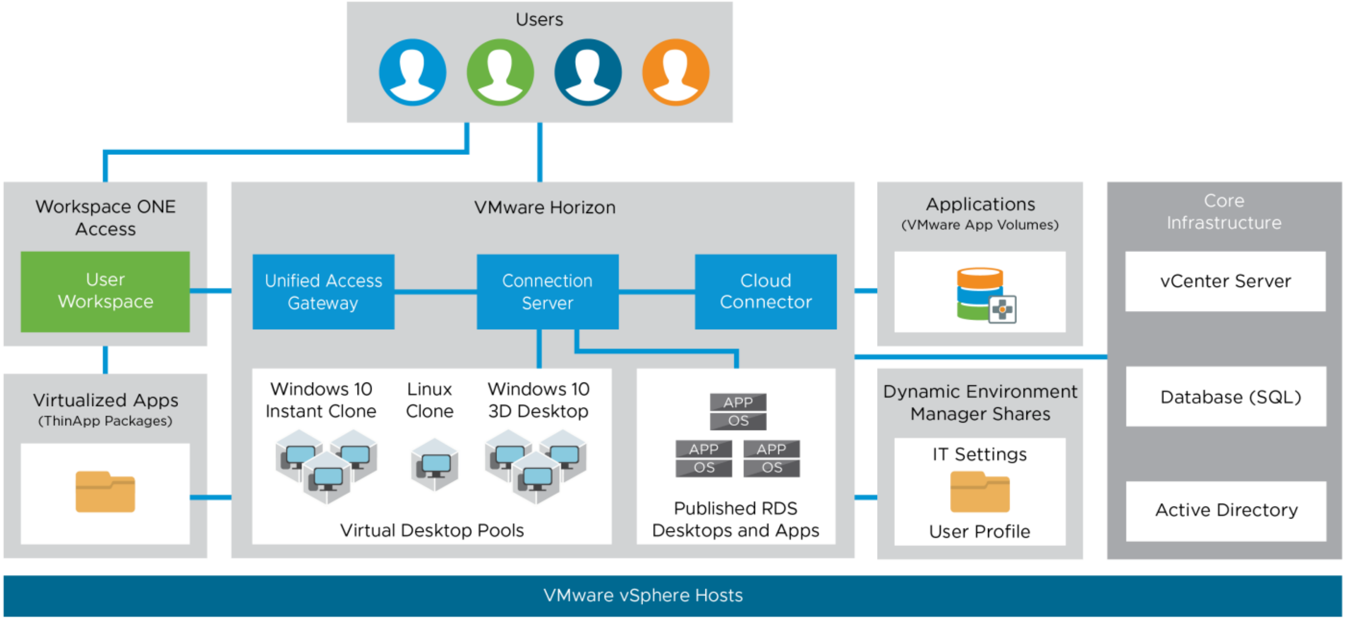

VMware Horizon Components and Features

- Connection Server

  The Horizon Connection Server securely brokers and connects users to the Horizon Agent that has been installed in the desktops and RDS Hosts. The Connection Server authenticates users through Active Directory and directs the request to the appropriate and entitled resource.

- Horizon Agent

  The Horizon Agent is installed on the guest OS of target VM or system. This agent allows the machine to be managed by Connection Servers and allows a Horizon Client to form a protocol session to the machine. Machines can be virtual desktops, Remote Desktop Session Hosts (RDS Host), physical desktops PCs.

- Horizon Client

  The Horizon Client is installed on a client device to access a Horizon-managed system that has the Horizon Agent installed. You can optionally use a web browser as an HTML client for devices on which installing client software is not possible.

- Unified Access Gateway

  VMware Unified Access Gateway is a virtual appliance that enables secure remote access from an external network to a variety of internal resources, including Horizon-managed resources. When providing access to internal resources, Unified Access Gateway can be deployed within the corporate DMZ or internal network, and acts as a reverse proxy host for connections to your company's resources. Unified Access Gateway directs authenticated requests to the appropriate resource and discards any unauthenticated requests. It also can perform the authentication itself, leveraging an additional layer of authentication when enabled.

- Horizon Console

  A web application that is part of the Connection Server, allowing administrators to configure the server, deploy and manage desktops, control user authentication, initiate and examine system and user events, carry out end-user support, and perform analytical activities.

- VMware Instant Clone Technology

  VMware technology that provides single-image management with automation capabilities. You can rapidly create automated pools or farms of instant-clone desktops or RDSH servers from a golden image VM. The technology reduces storage costs and streamlines desktop management by enabling easy updating and patching of hundreds or thousands of images from the golden image VM.

- RDSH servers

  Microsoft Windows Servers that provide published applications and session-based remote desktops to end users.

- Enrollment Server

  Server that delivers True SSO functionality by ensuring a user can single-sign-on to a Horizon resource when launched from Workspace ONE Access™, or through Unified Access Gateway, regardless of the authentication method.

- Horizon Cloud Connector

  The Horizon Cloud Connector is required to use with Horizon subscription licenses, services and management features hosted in the Horizon Cloud Service. The Horizon Cloud Connector is a virtual appliance that connects a Connection Server in a pod with the Horizon Cloud Service. You must have an active My VMware account to purchase a Horizon license from https://my.vmware.com.

- vSphere

  The vSphere product family includes VMware ESXi™ and VMware vCenter Server®, and it is designed for building and managing virtual infrastructures. The vCenter Server system provides key administrative and operational functions, such as provisioning, cloning, and VM management features, which are essential for VDI.

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 desktops environments for each user:

- Traditional PC: A traditional PC is what typically constitutes a desktop environment: physical device with a locally installed operating system.
- Remoted Desktop Server Hosted Sessions: 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 RDS server virtual machines and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.

- Streamed Applications: Streamed desktops and applications run entirely on the user‘s local client device and are sent from a server on demand. The user interacts with the application or desktop directly, but the resources may only 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 VDI desktops and VMware RDS server sessions were validated. Each of the sections provides some fundamental design decisions for this environment.

Understanding Applications and Data

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

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

Project Planning and Solution Sizing Sample Questions

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

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

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

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

- What is the desktop OS planned? Windows 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?

- Are there any applications installed? What application delivery methods will be used, Installed, Streamed, Layered, Hosted, or Local?
- What is the desktop OS planned for RDS Server Roles? Windows server 2016 or Server 2019?
- 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 there a 3rd party graphics component?
- Is anti-virus a part of the image?
- What is the SQL server version for database? SQL server 2016 or 2019?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere has been identified the hypervisor for both RDS 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 multipathing storage layer. More information on vSphere can be obtained at the VMware web site.

For this CVD, the hypervisor used was VMware vSphere 7.0 GA.

Server OS and Desktop OS Machines configured in this CVD to support Remoted Desktop Server Hosted (RDSH) shared sessions and Virtual Desktops (both non-persistent and persistent).

Storage Considerations

Boot from SAN

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

Pure Storage FlashArray Considerations

Make sure Each FlashArray Controller is connected to BOTH storage fabrics (A/B).

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

How big should a Volume be? With the Purity Operating Environment, we remove the complexities of aggregates, RAID groups, and so on. When managing storage, you just create a volume based on the size required, availability and performance are taken care of through RAID-HD and DirectFlash Software. As an administrator you can create 1 10TB volume or 10 1TB Volumes and their performance/availability will be the same, so instead of creating volumes for availability or performance you can think about recoverability, manageability, and administrative considerations. For example, what data do I want to present to this application or what data do I want to store together so I can replicate it to another site/system/cloud, and so on.

Port Connectivity

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

16/32Gb Fiber Channel support (N-2 support) – Pure Storage offer up to 32Gb FC support on the latest FlashArray//X series arrays. Always make sure the correct number of HBAs and the speed of SFPs are included in the original FlashArray BOM.

Oversubscription

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

Topology

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

VMware Virtual Volumes Considerations

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

Ensure that the Config vVol is either part of an existing FlashArray Protection Group, Storage Policy that includes snapshots, or manual snapshots of the Config vVol are taken. This will help with the VM recovery process if the VM is deleted.

There are some FlashArray limits on Volume Connections per Host, Volume Count, and Snapshot Count. For more information about FlashArray limits review the following: https://support.purestorage.com/FlashArray/PurityFA/General_Troubleshooting/Pure_Storage_FlashArray_Limits

When a Storage Policy is applied to a vVol VM, the volumes associated with that VM are added to the designated protection group when applying the policy to the VM. If replication is part of the policy, be mindful of the amount of VMs using that storage policy and replication group. A large amount of VMs with a high change rate could cause replication to miss its schedule due to increased replication bandwidth and time needed to complete the scheduled snapshot. Pure Storage recommends vVol VMs that have Storage Policies applied be balanced between protection groups.

Storage Best Practices

Pure Storage FlashArray Best Practices for VMware vSphere 7.0

The following Pure Storage best practices for VMware vSphere should be followed as part of a design:

- FlashArray Volumes are automatically presented to VMware vSphere using the Round Robin Path Selection Policy (PSP) and appropriate vendor Storage Array Type Plugin (SATP) for vSphere 7.0.
- vSphere 7.0 also uses the Latency SATP that was introduced in vSphere 6.7U1 (This replaces the I/O Operations Limit of 1 SATP, which was the default from vSphere 6.5U1).
- When using iSCSI connected FlashArray volumes, it is recommended to set DelayedAck to false (disabled) and LoginTimeout to 30 seconds. Jumbo Frames are optional when using iSCSI.
- In vSphere 6.x, if hosts have any VMFS-5 volumes, change EnableBlockDelete to enabled. If it is all VMFS-6, this change is not needed.
- For VMFS-5, Run UNMAP frequently.
- For VMFS-6, keep automatic UNMAP enabled.
- DataMover.HardwareAcceleratedMove, DataMover.HardwareAcceleratedInit, and VMFS3.HardwareAcceleratedLocking should all be enabled.
- Ensure all ESXi hosts are connected to both FlashArray controllers. A minimum of two paths to each. Aim for total redundancy.
- Install VMware tools or Open VM tools whenever possible.
- Queue depths should be left at the default. Changing queue depths on the ESXi host is a tweak and should only be examined if a performance problem (high latency) is observed.
- When mounting snapshots, use the ESXi resignature option and avoid force-mounting.
- Configure Host Groups on the FlashArray identically to clusters in vSphere. For example, if a cluster has four hosts in it, create a corresponding Host Group on the relevant FlashArray with exactly those four hosts—no more, no less.

- When possible, use Paravirtual SCSI adapters for virtual machines.
- Atomic Test and Set (ATS) is required on all Pure Storage volumes. This is a default configuration, and no changes should normally be needed.
- UseATSForHBOnVMFS5 should be enabled. This was introduced in vSphere 5.5 U2 and is enabled by default. It is NOT required though.

For more information about the VMware vSphere Pure Storage FlashArray Best Practices, see:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/001VMwareBestPractices/hhhWeb_Guide%3A_FlashArray_VMware_Best_Practices

Pure Storage FlashArray Best Practices for VMware Virtual Volumes (vVols)

Along with the Pure Storage Best Practices for VMware vSphere, the following should be considered as part of a design that includes the implementation of vVols as part of the solution:

- Create a Local FlashArray Array Admin user to register the storage provider with vs using the local pureuser account, vvols-admin for example.
- Use the Round Robin pathing policy (default) for the Protocol Endpoint.
- Use the Pure Storage Plugin for the vSphere Client to register the FlashArray storage provider and mount the vVols Datastore if possible.
- If manually registering the storage providers, Register both controllers' storage providers with CT0.ETH0 and CT1.ETH0. It is supported to use Eth1 if a customer certificate is used.
- If manually mounting the vVol datastore, you will need to connect the protocol endpoint.
- A single PE should be enough for the design utilizing the default device queue depth for the PE.
- Keep VM Templates on vVols when deploying new vVol VMs from a template.
- When resizing a VM's VMDK that resides on a vVol, complete the task from vSphere Client and not the FlashArray GUI.
- vCenter Server should not reside on a vVol
- All ESXi Hosts, vCenter Server and FlashArray should have the same NTP Server synchronization configuration and be configured to send their logs to a syslog target.

For more information about vVols Best practices, refer to the following quick reference:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/Quick_Reference_by_VMware_Product_and_Integration/Virtual_Volumes_Quick_Reference

VMware Horizon Design Fundamentals

Design a Vmware Horizon Environment for a Mixed Workload

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

Table 2. Designing a VMware Horizon Environment

Server OS machines	**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. **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. **Application types**: Any application.
Desktop OS machines	**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. **Your users**: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications. **Application types**: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines. Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.
Remote PC Access	**You want:** Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the datacenter. 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. Host: The same as Desktop OS machines. Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.

For the Cisco Validated Design described in this document, a Remote Desktop Server Hosted sessions (RDSH) using RDS based Server OS and VMware Horizon pooled Instant and Full Clone Virtual Machine Desktops using VDI based desktop OS machines were configured and tested.

The mixed workload test case consisted of a combination of all 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 Pure Storage FlashArrays).

The FlashStack Data Center solution includes Cisco networking, Cisco UCS and Pure Storage FlashArray//X70 R3, which efficiently fit into a single data center rack, including the access layer network switches.

This CVD details the deployment of up to 6000 users for RDSH and 5000 VDI users VMware Horizon desktop workload featuring the following software:

- VMware vSphere 7.0 GA Hypervisor
- Microsoft SQL Server 2019
- Microsoft Windows Server 2019 and Windows 10 64-bit virtual machine Operating Systems
- VMware Horizon 8 Remote Desktops (RDSH) provisioned as Full Clones
- VMware Horizon 8 Non-Persistent Virtual Desktops (VDI) provisioned as Instant Clones
- VMware Horizon 8 Persistent Virtual Desktops (VDI) provisioned as Full Clones
- VMware Dynamic Environment Manger 10.0

Figure 19. Virtual Desktop Workload Reference Architecture on FlashStack

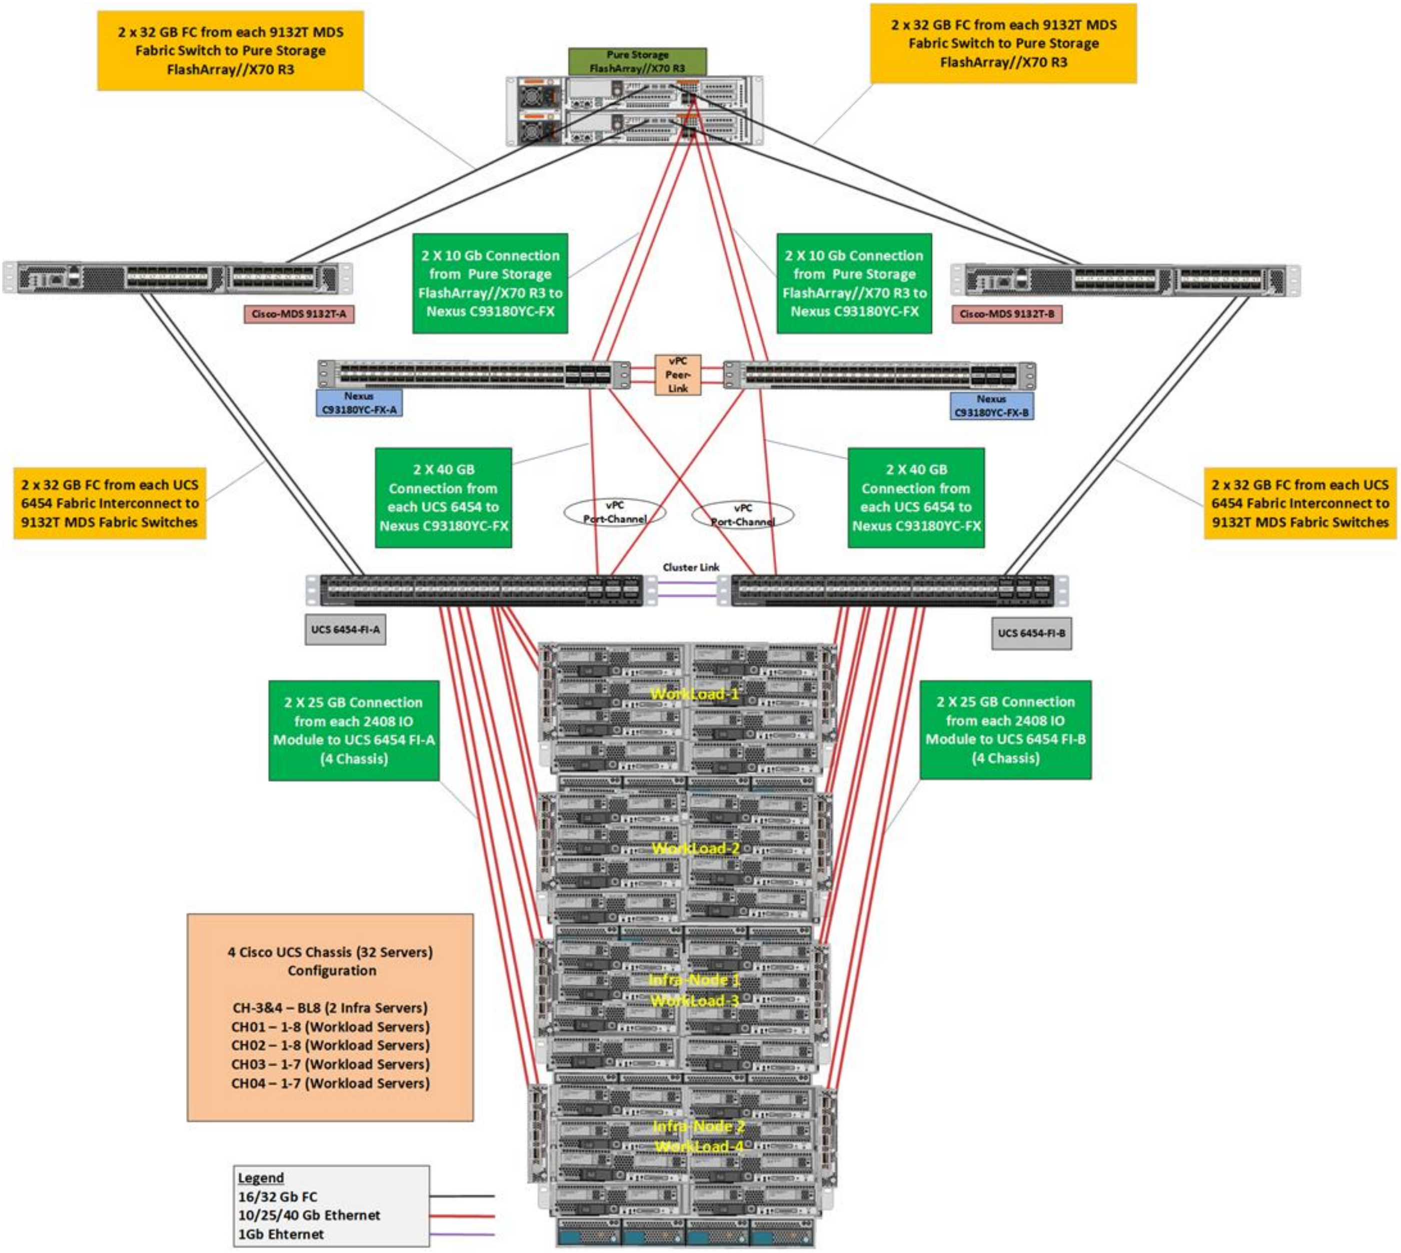
2 x 32 GB FC from each 9132T MDS Fabric Switch to Pure Storage FlashArray//X70 R3
Pure Storage FlashArray//X70 R3
2 x 32 GB FC from each 9132T MDS Fabric Switch to Pure Storage FlashArray//X70 R3
2 X 10 Gb Connection from Pure Storage FlashArray//X70 R3 to Nexus C93180YC-FX
2 X 10 Gb Connection from Pure Storage FlashArray//X70 R3 to Nexus C93180YC-FX
Cisco-MDS 9132T-A
Cisco-MDS 9132T-B
vPC Peer-Link
Nexus C93180YC-FX-A
Nexus C93180YC-FX-B
2 x 32 GB FC from each UCS 6454 Fabric Interconnect to 9132T MDS Fabric Switches
2 X 40 GB Connection from each UCS 6454 to Nexus C93180YC-FX
vPC Port-Channel
vPC Port-Channel
2 X 40 GB Connection from each UCS 6454 to Nexus C93180YC-FX
2 x 32 GB FC from each UCS 6454 Fabric Interconnect to 9132T MDS Fabric Switches
Cluster Link
UCS 6454-FI-A
UCS 6454-FI-B
2 X 25 GB Connection from each 2408 IO Module to UCS 6454 FI-A (4 Chassis)
Workload-1
2 X 25 GB Connection from each 2408 IO Module to UCS 6454 FI-B (4 Chassis)
Workload-2
4 Cisco UCS Chassis (32 Servers) Configuration
CH-3&4 – BL8 (2 Infra Servers)
CH01 – 1-8 (Workload Servers)
CH02 – 1-8 (Workload Servers)
CH03 – 1-7 (Workload Servers)
CH04 – 1-7 (Workload Servers)
Infra-Node 1
Workload-3
Infra-Node 2
Workload-4
Legend
16/32 Gb FC
10/25/40 Gb Ethernet
1Gb Ehternet

Figure 19 details the physical hardware and cabling deployed to enable this solution:

- Two Cisco Nexus 93180YC-FX Layer 2 Access Switches.
- Two Cisco MDS 9132T 32-Gb Fibre Channel Switches.
- Four Cisco UCS 5108 Blade Server Chassis with two Cisco UCS-IOM-2408 IO Modules.
- Two Cisco UCS B200 M5 Blade Servers with Intel® Xeon® Silver 4210 2.2-GHz 10-core processors, 384GB 2933MHz RAM, and one Cisco VIC1440 mezzanine card for the hosted infrastructure, providing N+1 server fault tolerance.

- Thirty Cisco UCS B200 M5 Blade Servers with Intel® Xeon® Gold 6230 2.1-GHz 20-core processors, 768GB 2933MHz RAM, and one Cisco VIC1440 mezzanine card, providing N+1 server fault tolerance.
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives.

The LoginVSI Test infrastructure is not a part of this solution. The Pure FlashArray//X70 R3 configuration is detailed later in this document.

Software Revisions

Table 3 lists the software versions of the primary products installed in the environment.

Table 3. Software and Firmware Versions

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	4.1(2a) bundle release
Cisco	UCS Manager	4.1(2a) bundle release
Cisco	UCS B200 M5 Blades	4.1(2a) bundle release
Cisco	VIC 1440	4.1(2a) bundle release
Pure Storage	FlashArray//X70 R3	Purity//FA v6.0.3
VMware	vCenter Server Appliance	7.0.0.10400
VMware	vSphere 7. 0 GA	7.0.0.15843807
VMware	Horizon Connection Server	8.1.0.17351278
VMware	Dynamic Environment Manger Enterprise	10.0.0.945
VMware	Horizon Agent	8.1.0
VMware	Tools	11.0.5.15389592

Logical Architecture

The logical architecture of the validated solution which is designed to support up to 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 20.

Figure 20. Logical Architecture Overview

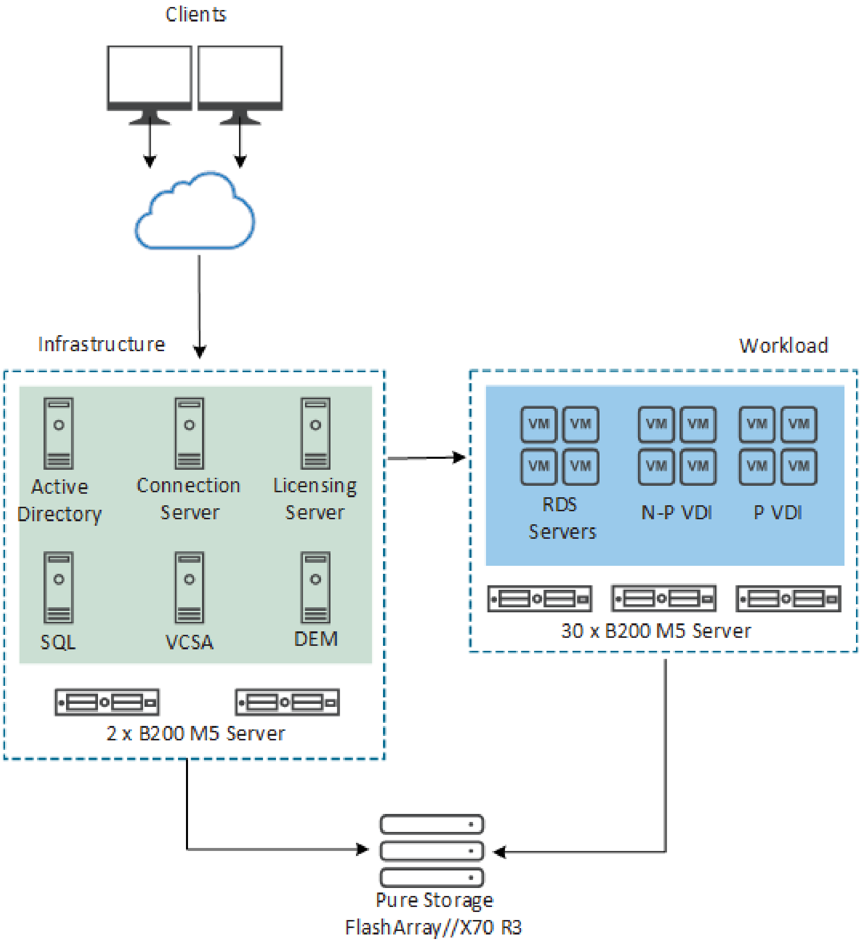

Configuration Guidelines

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

This document is intended to allow the reader to configure the VMware Horizon 8 customer environment as a stand-alone solution.

VLANs

The VLAN configuration recommended for the environment includes a total of six VLANs as outlined in Table 4.

Table 4. VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN

VLAN Name	VLAN ID	VLAN Purpose
In-Band-Mgmt	70	In-Band management interfaces
Infra-Mgmt	71	Infrastructure Virtual Machines
VCC/VM-Network	72	RDSH, VDI Persistent and Non-Persistent
vMotion	73	VMware vMotion
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 5.

Table 5. VSANs Configured in this Study

VSAN Name	VSAN ID	Purpose
VSAN 100	100	VSAN for Primary SAN communication
VSAN 101	101	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.

Solution Cabling

The following sections detail the physical connectivity configuration of the FlashStack 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 Pure Storage FlashArray//X70 R3 storage array to the Cisco 6454 Fabric Interconnects through Cisco MDS 9132T 32-Gb FC switches.

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.

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

Figure 21 shows a cabling diagram for a configuration using the Cisco Nexus 9000, Cisco MDS 9100 Series, and Pure Storage FlashArray//X70 R3 array.

Figure 21. FlashStack Solution Cabling Diagram

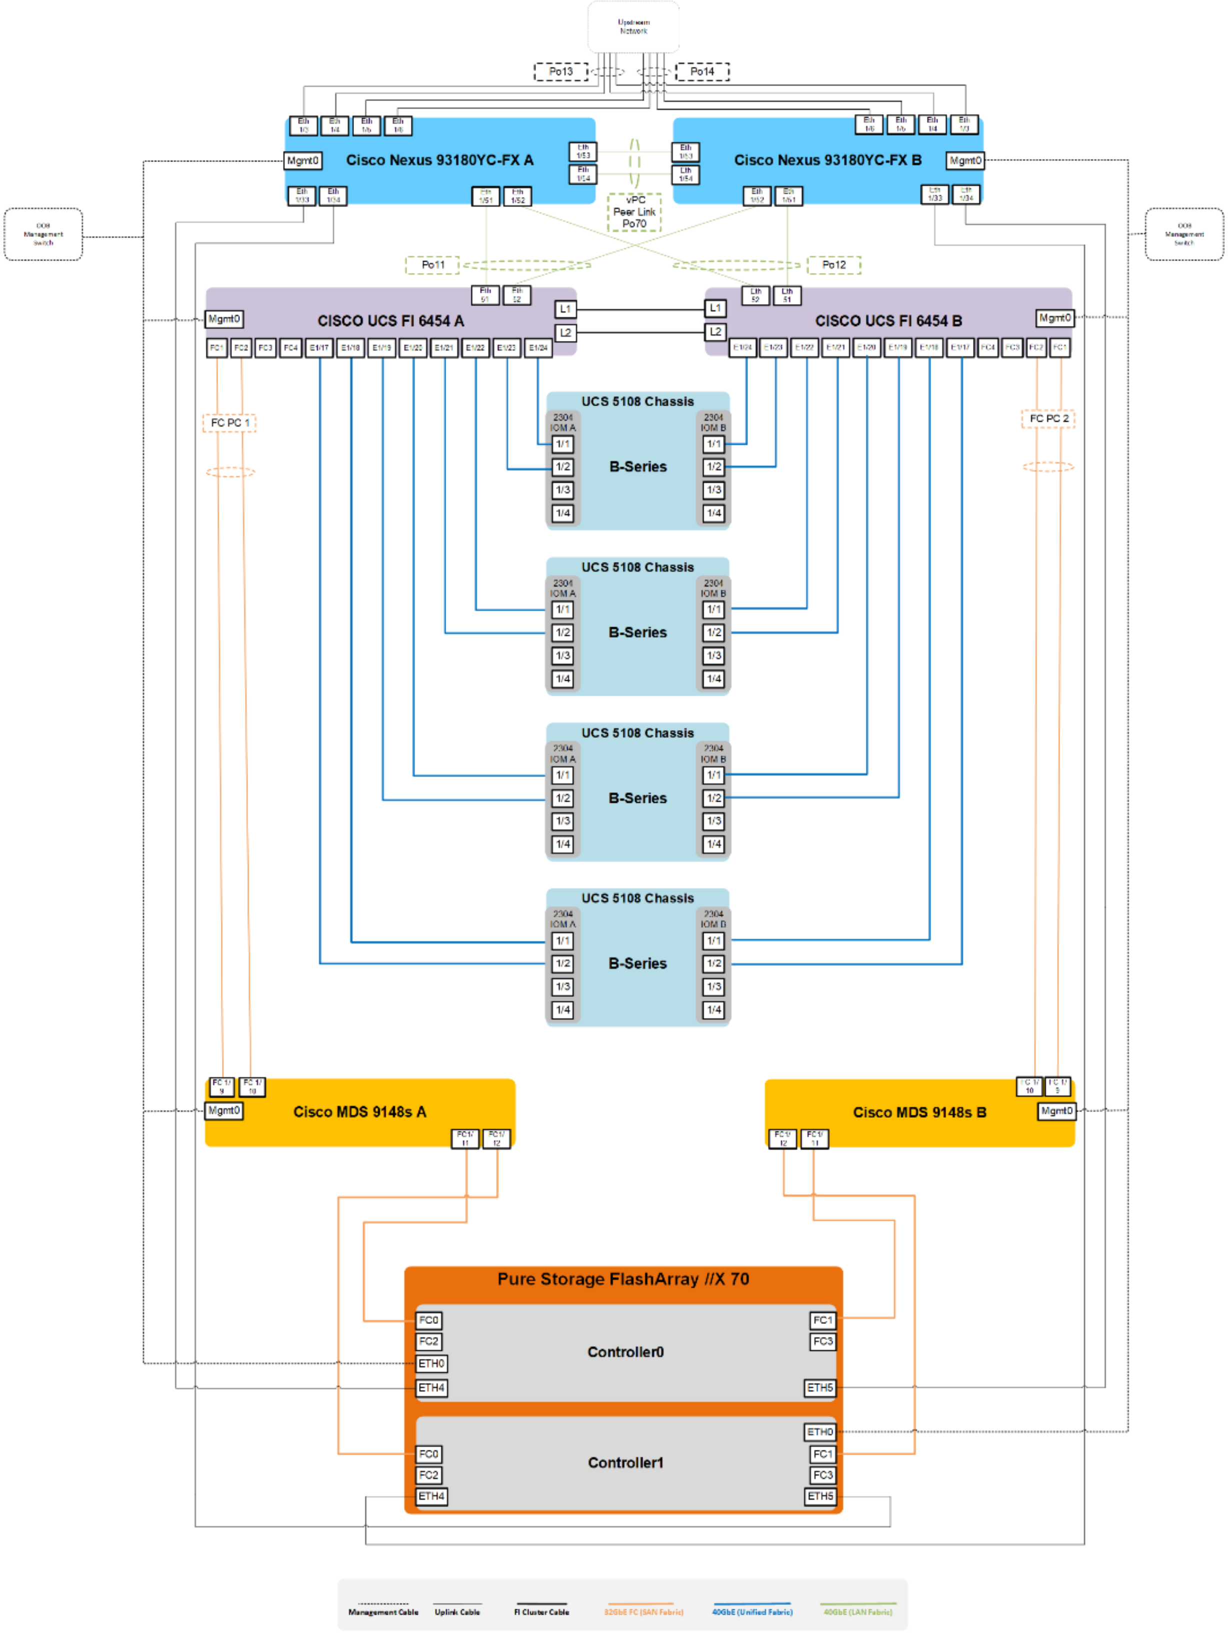

Cisco Unified Computing System Base Configuration

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 4.1(2a)

This document assumes you are using Cisco UCS Manager Software version 4.1(2a). To upgrade the Cisco UCS Manager software and the Cisco UCS 6454 Fabric Interconnect software to a higher version of the firmware,) refer to Cisco UCS Manager Install and Upgrade Guides.

Configure Fabric Interconnects at Console

To configure the fabric Interconnects, follow these steps:

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

2. If the fabric interconnect was previously deployed and you want to erase it to redeploy, follow these steps:

   a. Login with the existing user name and password.

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

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

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

Configure Fabric Interconnects for a Cluster Setup

To configure the Cisco UCS Fabric Interconnects, follow these steps:

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

   a. The management Ethernet port (mgmt0) is connected to an external hub, switch, or router

   b. The L1 ports on both fabric interconnects are directly connected to each other

   c. The L2 ports on both fabric interconnects are directly connected to each other

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

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

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

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

Figure 22. Initial Setup of Cisco UCS Manager on Primary Fabric Interconnect

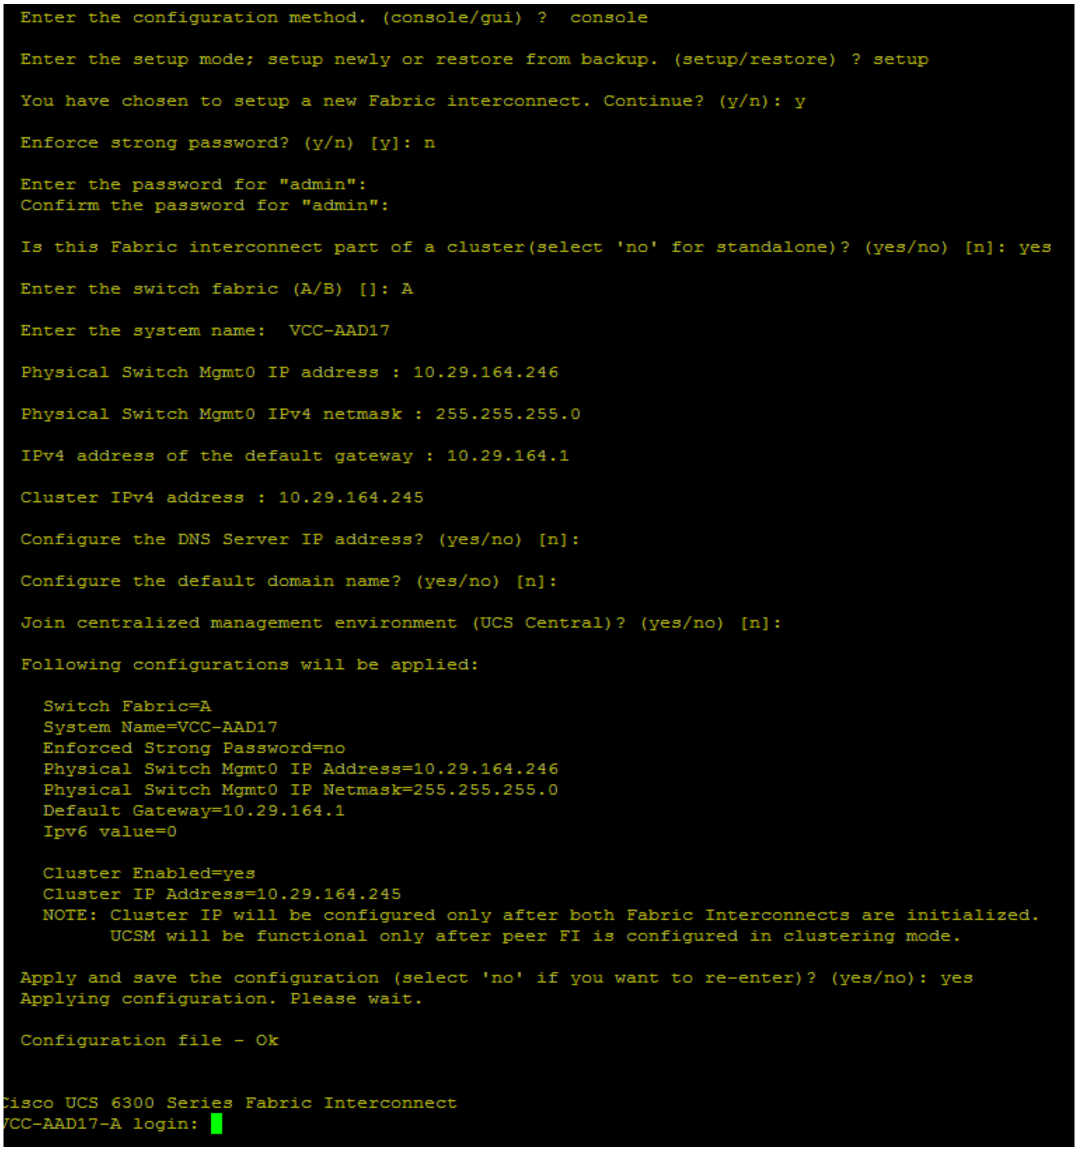

```
 Enter the configuration method. (console/gui) ?  console

 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 23. Initial Setup of Cisco UCS Manager on Secondary Fabric Interconnect

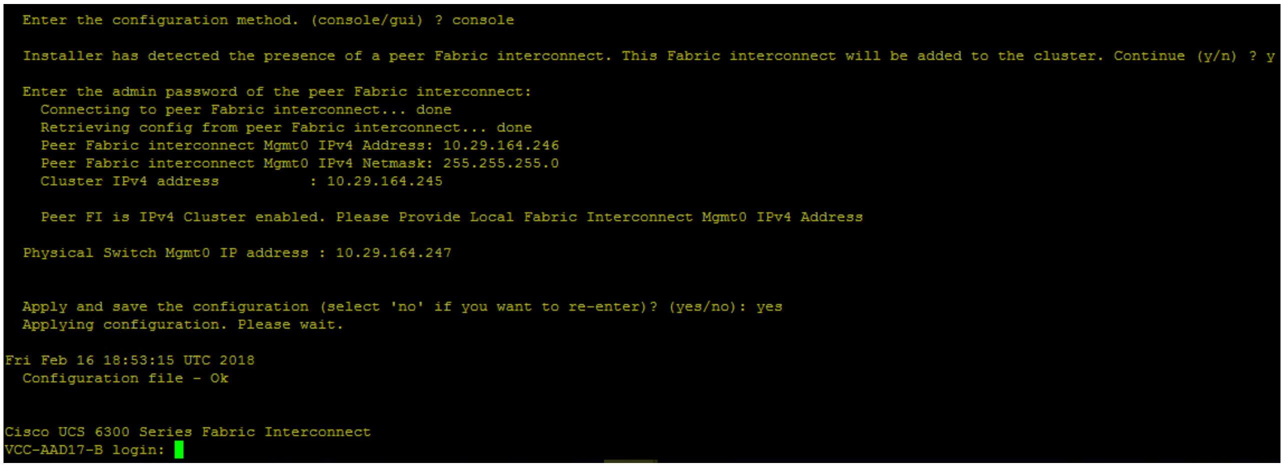

```
  Enter the configuration method. (console/gui) ? console

  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:
```

To log into the Cisco Unified Computing System (Cisco UCS) environment, follow these steps:

1. Open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address configured above.

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

Figure 24. Cisco UCS Manager Web Interface

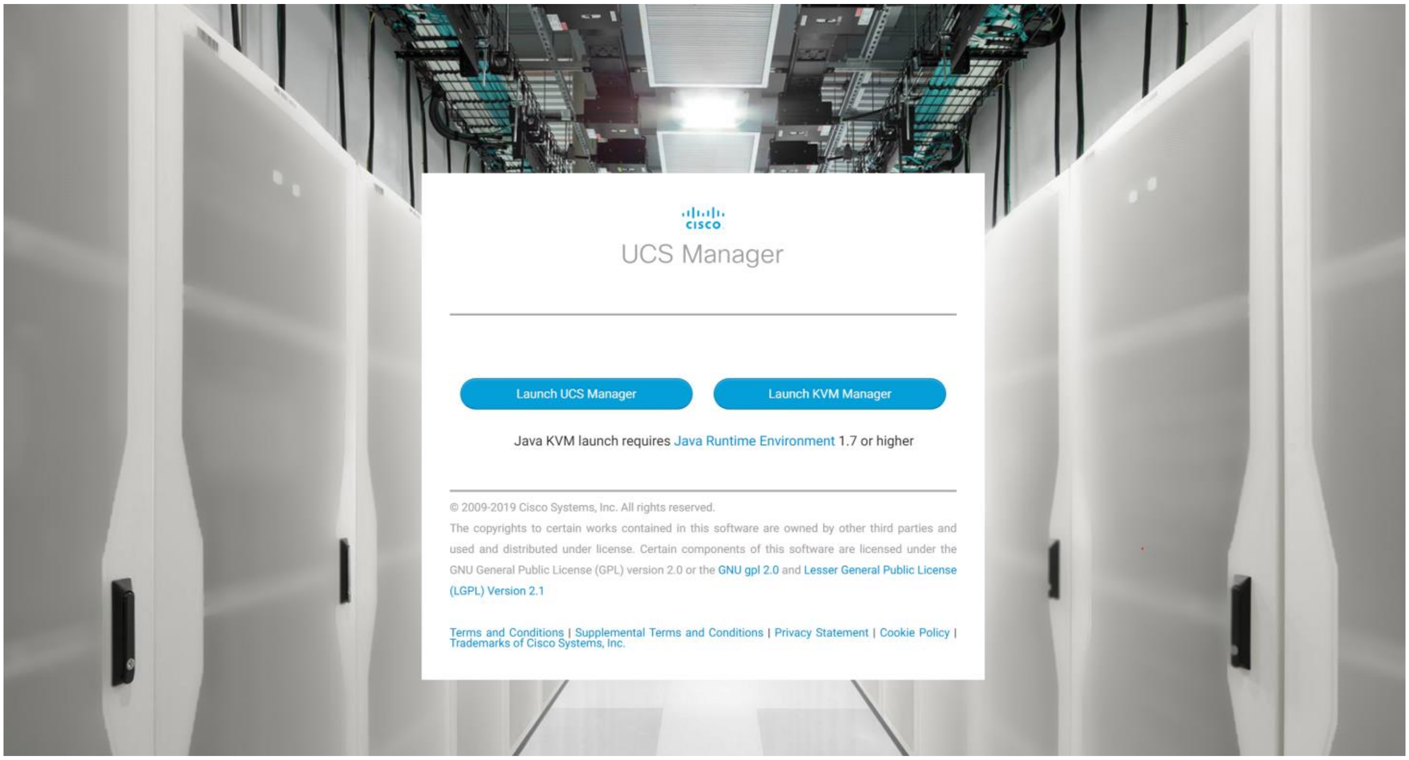

Java KVM launch requires Java Runtime Environment 1.7 or higher

© 2009-2019 Cisco Systems, Inc. All rights reserved.

The copyrights to certain works contained in this software are owned by other third parties and used and distributed under license. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or the GNU gpl 2.0 and Lesser General Public License (LGPL) Version 2.1

Terms and Conditions | Supplemental Terms and Conditions | Privacy Statement | Cookie Policy | Trademarks of Cisco Systems, Inc.

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

Figure 25. Cisco UCS Manager Web Interface after Login

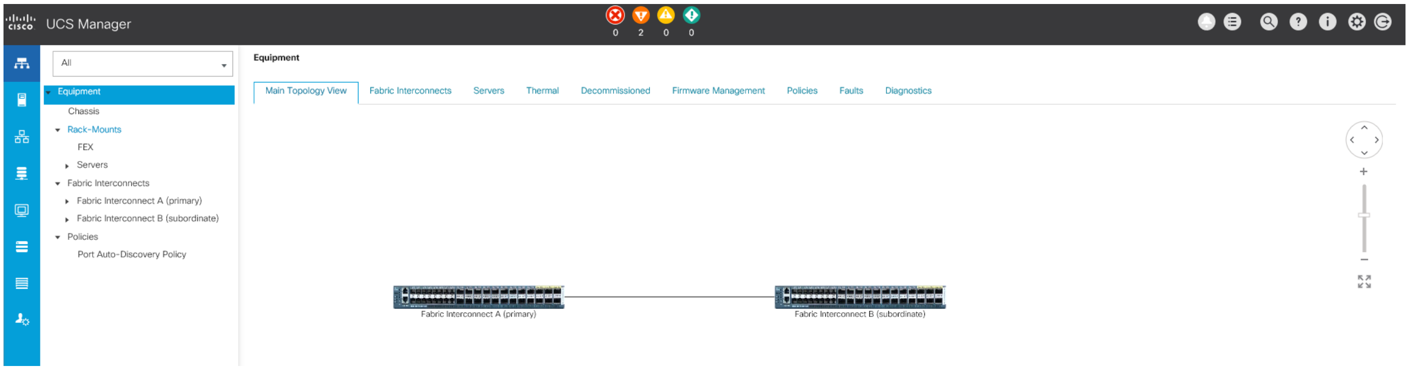

Configure Base Cisco Unified Computing System

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.

Synchronize Cisco UCSM to NTP

To synchronize the Cisco UCS environment to the NTP server, follow these steps:

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

2. Select All > Time zone Management.

3. In the Properties pane, select the appropriate time zone in the Time zone menu.

4. Click Save Changes and then click OK.

5. Click Add NTP Server.

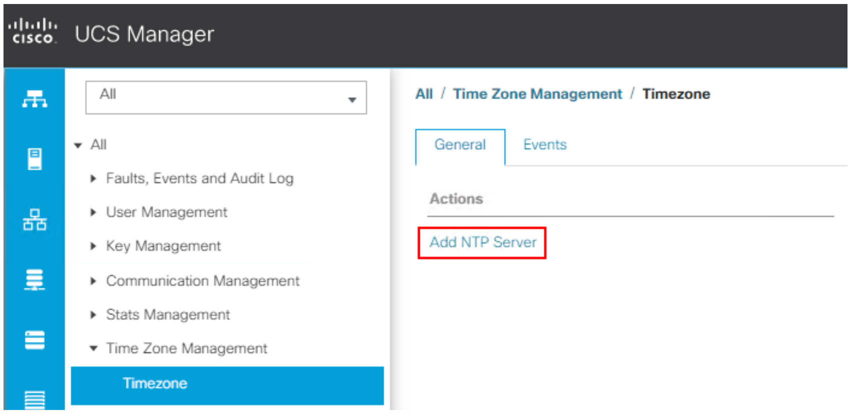

6. Enter the NTP server IP address and click OK.

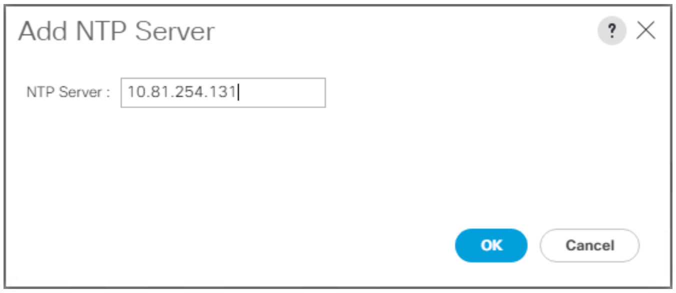

NTP Server : 10.81.254.131

7. Click OK to finish.

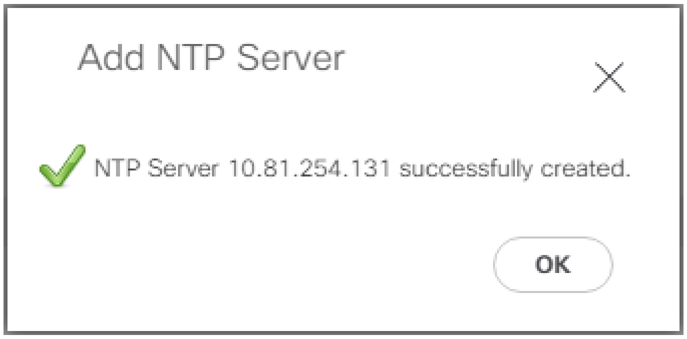

Add NTP Server

NTP Server 10.81.254.131 successfully created.

8. Repeat steps 1-7 to configure additional NTP servers.

9. Click Save Changes.

Figure 26. Synchronize Cisco UCS Manager to NTP

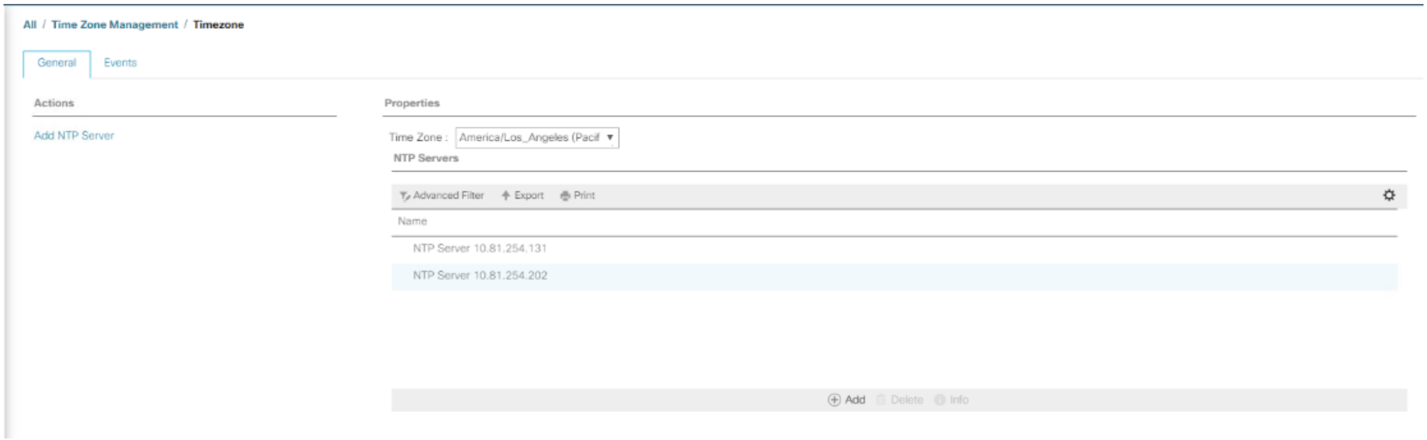

Time Zone : America/Los_Angeles (Pacif

NTP Servers

Name
NTP Server 10.81.254.131
NTP Server 10.81.254.202

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.

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.

To configure global policies, follow these steps:

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

2. Click Save Changes.

3. Click OK.

Figure 27. Cisco UCS Global Policy

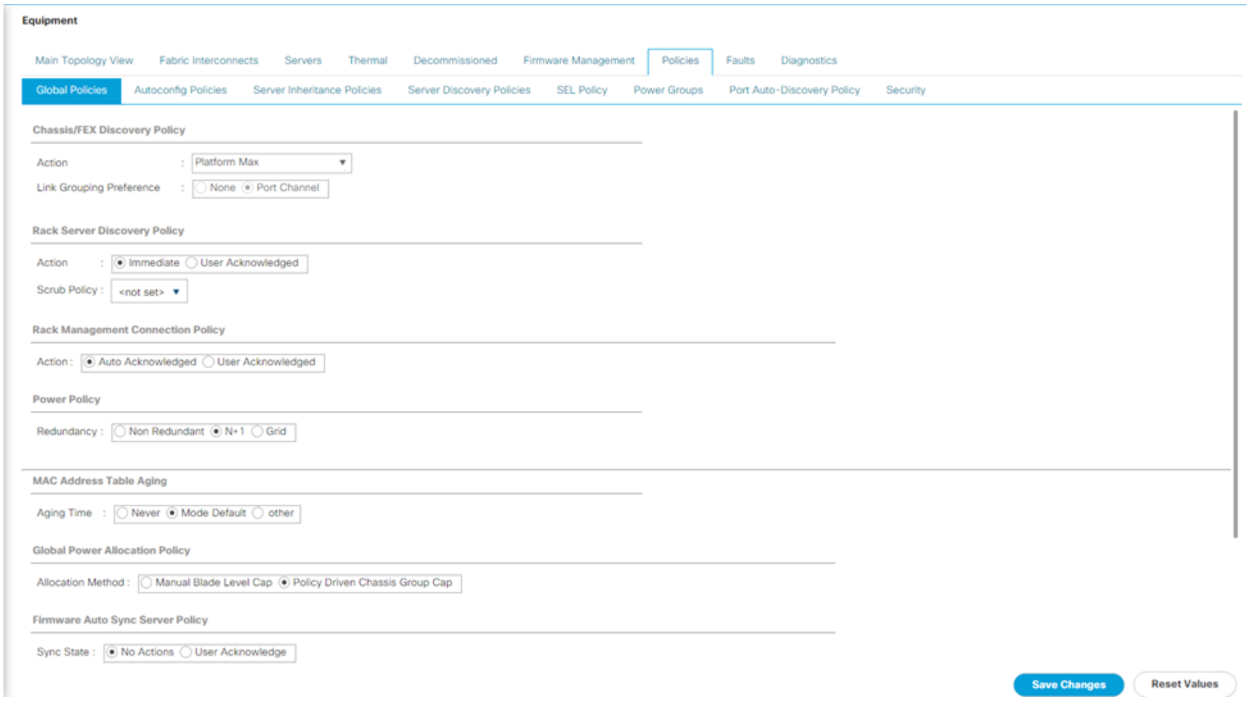

Chassis/FEX Discovery Policy

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

Fabric Ports: Discrete versus Port Channel Mode

Figure 28 illustrates the advantage of Discrete Vs Port-Channel mode in UCSM.

Figure 28. Port Channel versus Discrete Mode

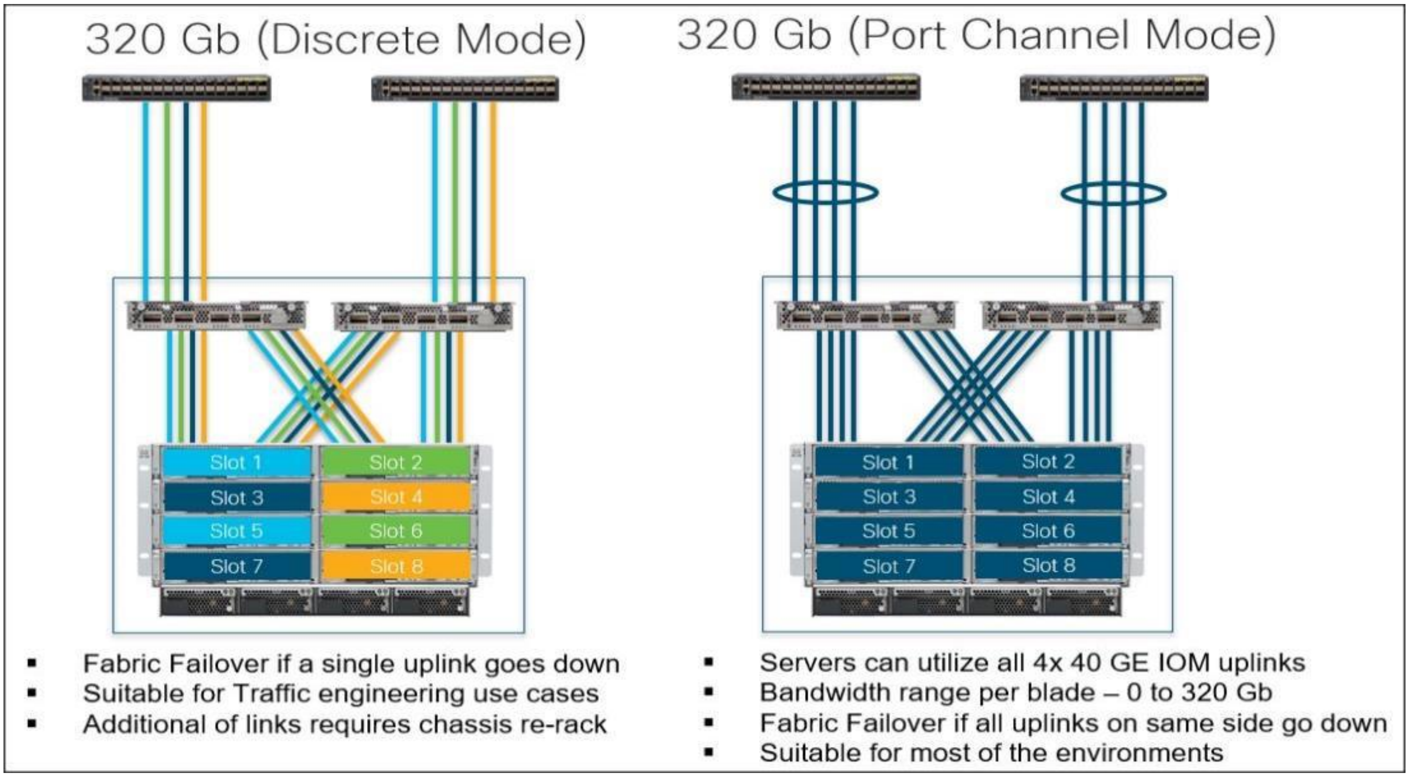

- Fabric Failover if a single uplink goes down
- Suitable for Traffic engineering use cases
- Additional of links requires chassis re-rack

- Servers can utilize all 4x 40 GE IOM uplinks
- Bandwidth range per blade – 0 to 320 Gb
- Fabric Failover if all uplinks on same side go down
- Suitable for most of the environments

Set Fabric Interconnects to Fibre Channel End Host Mode

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

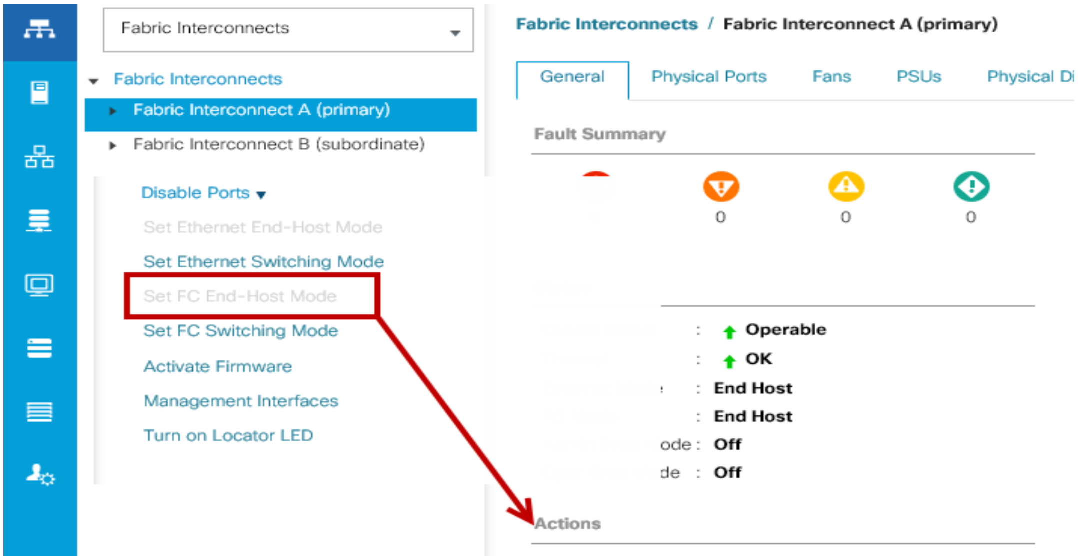

: Operable
: OK
: End Host
: End Host
ode : Off
de : Off

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

Configure FC SAN Uplink Ports

To configure Fibre Channel Uplink ports, follow these steps:

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

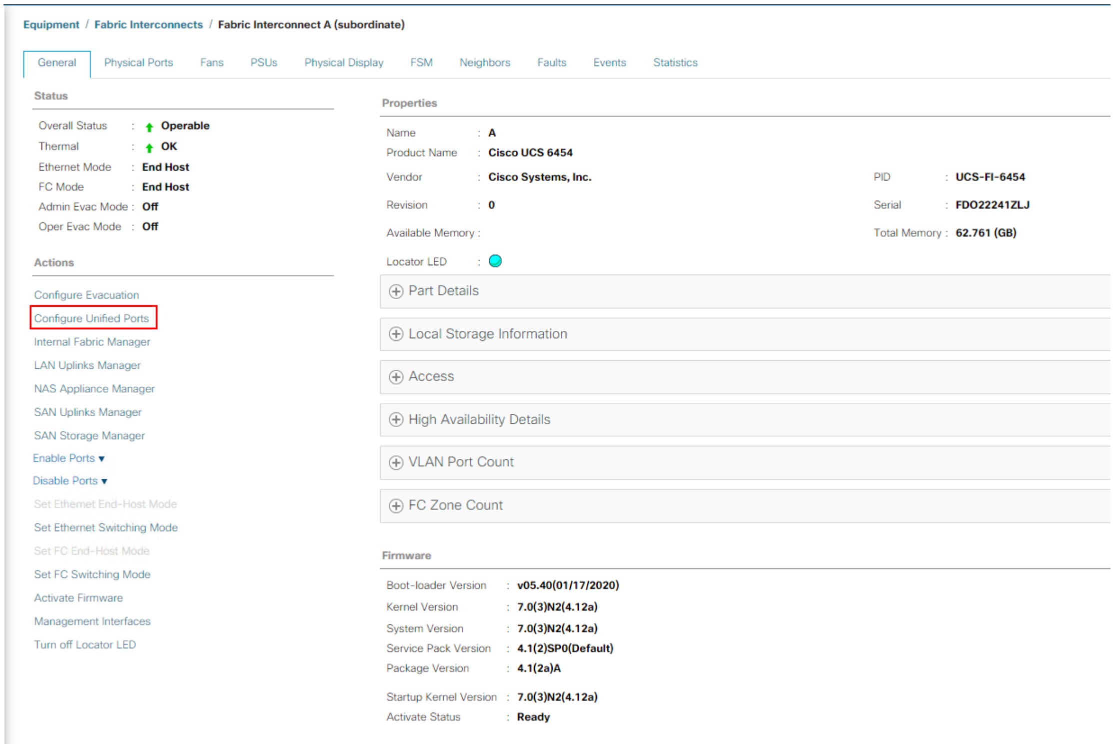

Overall Status : Operable
Thermal : OK
Ethernet Mode : End Host
FC Mode : End Host
Admin Evac Mode : Off
Oper Evac Mode : Off

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

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

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

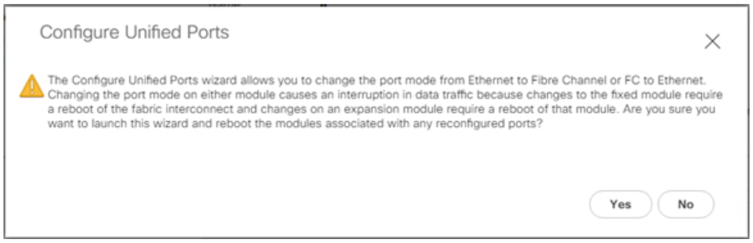

The Configure Unified Ports wizard allows you to change the port mode from Ethernet to Fibre Channel or FC to Ethernet. Changing the port mode on either module causes an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion 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?

3. Move the slider to the right.

4. Click OK.

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

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

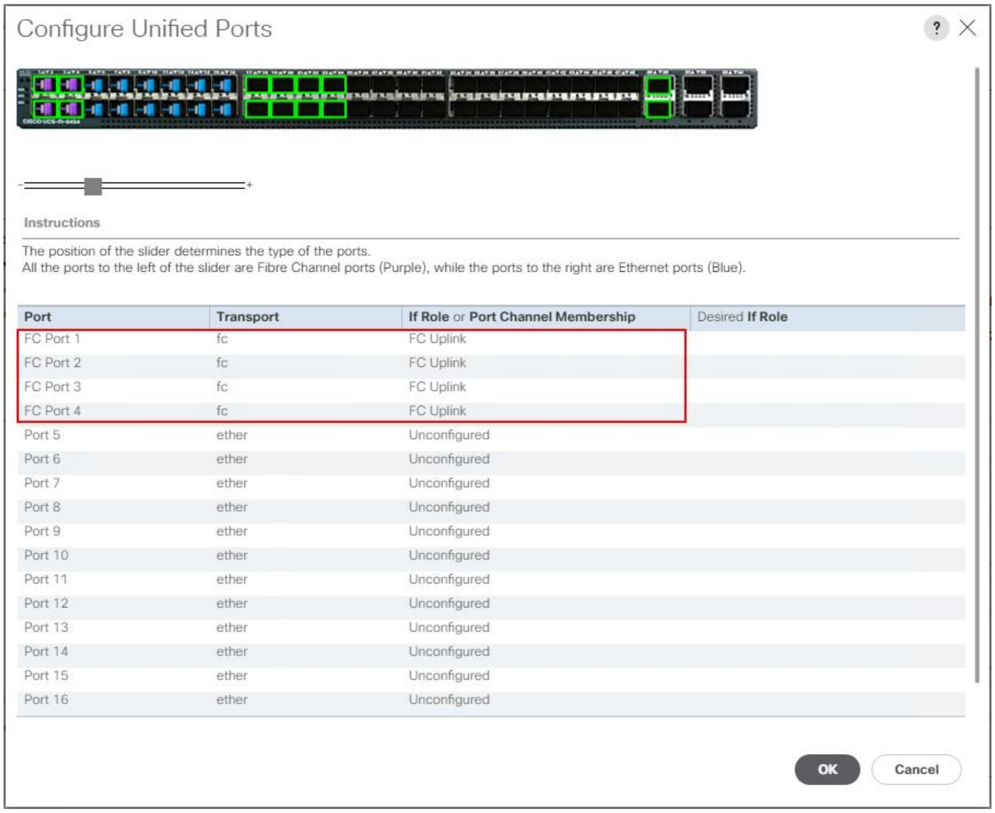

Configure Unified Ports

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
FC Port 1	fc	FC Uplink	
FC Port 2	fc	FC Uplink	
FC Port 3	fc	FC Uplink	
FC Port 4	fc	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	

5. Click Yes to apply the changes.

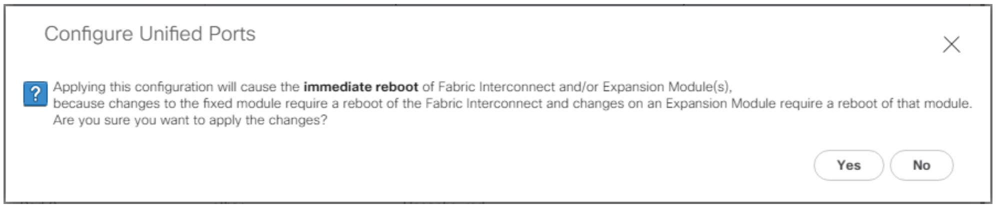

Configure Unified Ports

Applying this configuration will cause the **immediate reboot** of Fabric Interconnect and/or Expansion Module(s), 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 apply the changes?

6. Click OK to proceed.

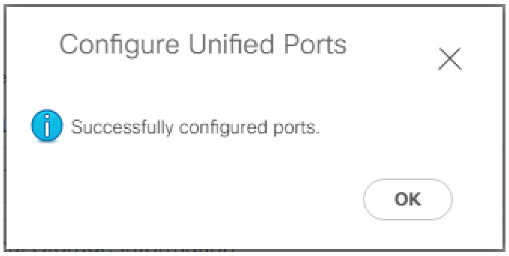

Configure Unified Ports

Successfully configured ports.

7. After the FI reboot, your FC Ports configuration will look like Figure 32.

8. Repeat steps 1-7 on Fabric Interconnect B.

Figure 29. FC Uplink Ports on Fabric Interconnect A

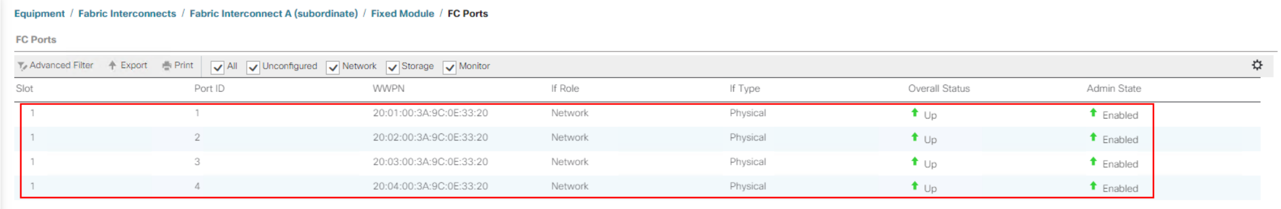

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module / FC Ports

FC Ports

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

Configure Server Ports

Configure the server ports to initiate chassis and blade discovery. To configure server ports, follow these steps:

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

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

3. Right-click and select "Configure as Server Port."

Figure 30. Configure Server Port on Cisco UCS Manager Fabric Interconnect for Chassis/Server Discovery

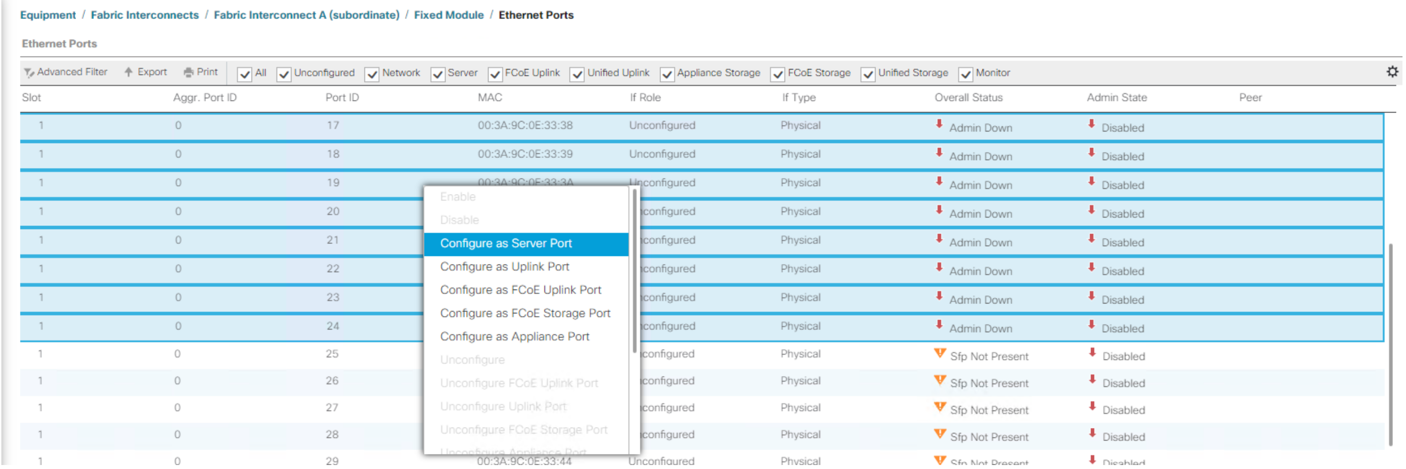

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module / Ethernet Ports

Ethernet Ports

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:0E:33:38	Unconfigured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	18	00:3A:9C:0E:33:39	Unconfigured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	19	00:3A:9C:0E:33:3A	Unconfigured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	20		configured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	21		configured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	22		configured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	23		configured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	24		configured	Physical	⬇ Admin Down	⬇ Disabled	
1	0	25		configured	Physical	⚠ Sfp Not Present	⬇ Disabled	
1	0	26		configured	Physical	⚠ Sfp Not Present	⬇ Disabled	
1	0	27		configured	Physical	⚠ Sfp Not Present	⬇ Disabled	
1	0	28		configured	Physical	⚠ Sfp Not Present	⬇ Disabled	
1	0	29	00:3A:9C:0E:33:44	Unconfigured	Physical	⚠ Sfp Not Present	⬇ Disabled	

4. Click Yes to confirm and click OK.

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

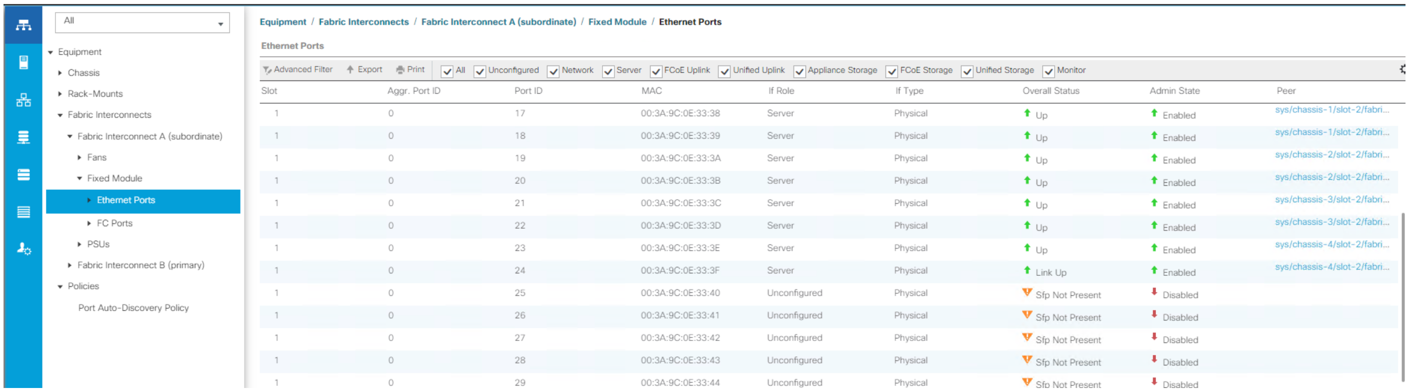
All
Equipment
Chassis
Rack-Mounts
Fabric Interconnects
Fabric Interconnect A (subordinate)
Fans
Fixed Module
Ethernet Ports
FC Ports
PSUs
Fabric Interconnect B (primary)
Policies
Port Auto-Discovery Policy
Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module / Ethernet Ports
Ethernet Ports
Advanced Filter
Export
Print
☑ All ☑ Unconfigured ☑ Network ☑ Server ☑ FCoE Uplink ☑ Unified Uplink ☑ Appliance Storage ☑ FCoE Storage ☑ Unified Storage ☑ Monitor

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

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

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

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

Figure 32. Internal LAN Port Channels

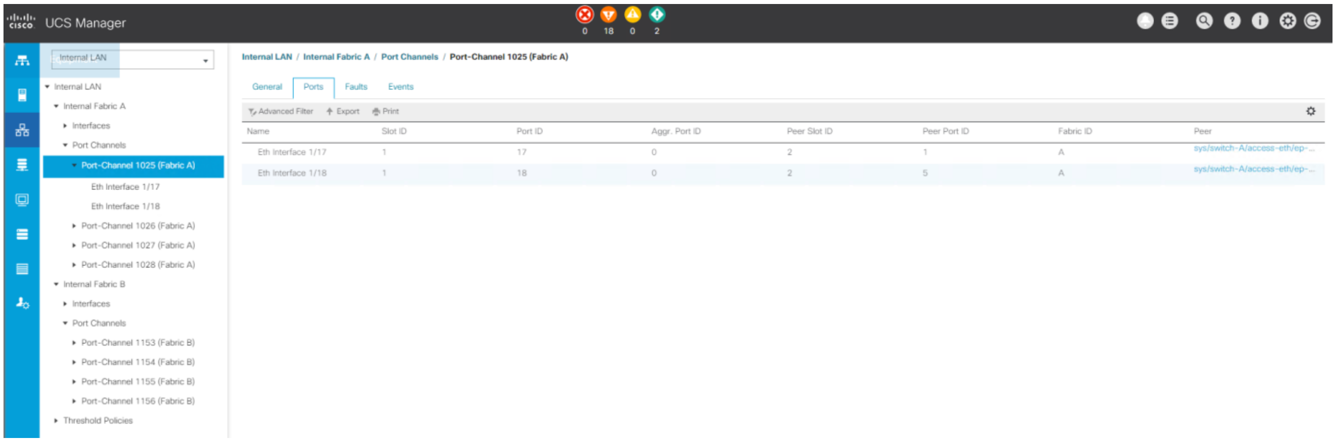
UCS Manager
0
18
0
2
Internal LAN
Internal LAN
Internal Fabric A
Interfaces
Port Channels
Port-Channel 1025 (Fabric A)
Eth Interface 1/17
Eth Interface 1/18
Port-Channel 1026 (Fabric A)
Port-Channel 1027 (Fabric A)
Port-Channel 1028 (Fabric A)
Internal Fabric B
Interfaces
Port Channels
Port-Channel 1153 (Fabric B)
Port-Channel 1154 (Fabric B)
Port-Channel 1155 (Fabric B)
Port-Channel 1156 (Fabric B)
Threshold Policies
Internal LAN / Internal Fabric A / Port Channels / Port-Channel 1025 (Fabric A)
General
Ports
Faults
Events
Advanced Filter
Export
Print

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

Configure Ethernet LAN Uplink Ports

To configure network ports that are used to uplink the Fabric Interconnects to the Cisco Nexus switches, follow these steps:

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

2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.

3. Expand Ethernet Ports.

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

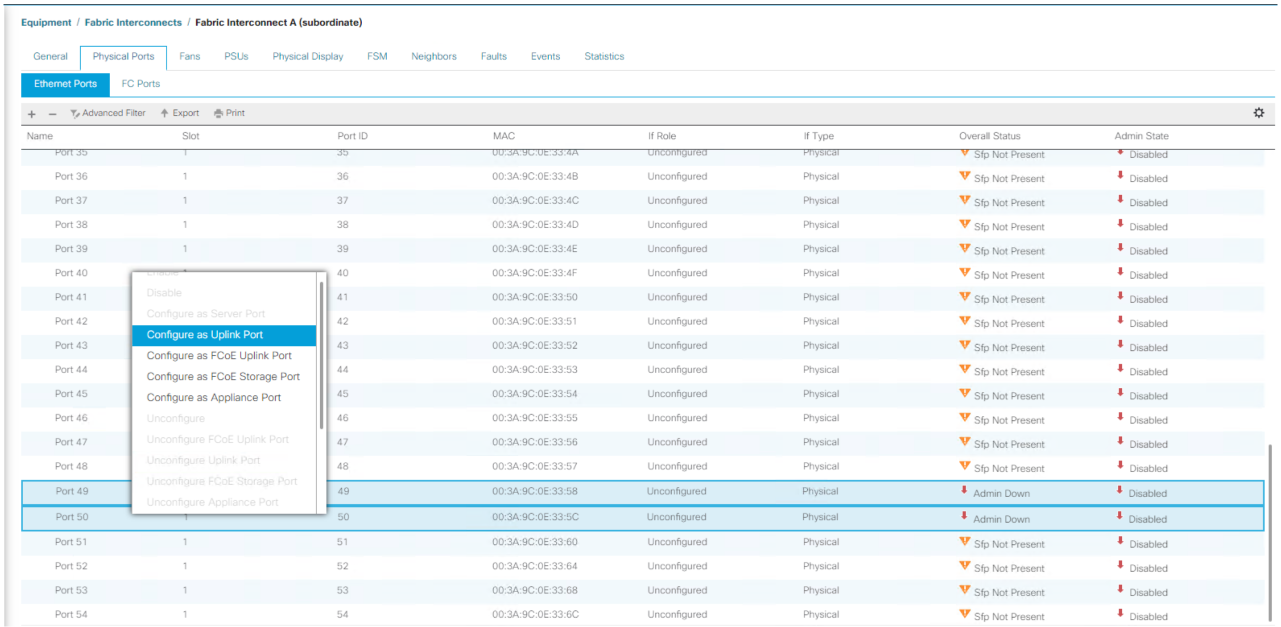
Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate)
General
Physical Ports
Fans
PSUs
Physical Display
FSM
Neighbors
Faults
Events
Statistics
Ethernet Ports
FC Ports
+ − Advanced Filter
Export
Print
Disable
Configure as Server Port
Configure as Uplink Port
Configure as FCoE Uplink Port
Configure as FCoE Storage Port
Configure as Appliance Port
Unconfigure
Unconfigure FCoE Uplink Port
Unconfigure Uplink Port
Unconfigure FCoE Storage Port
Unconfigure Appliance Port

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

5. Click Yes to confirm ports and click OK.

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

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

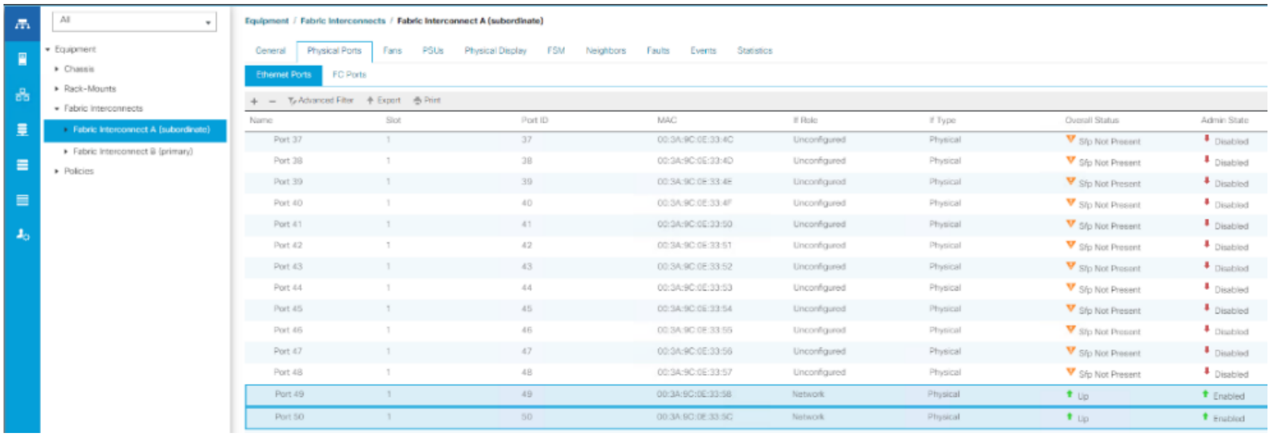

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 37	1	37	00:3A:9C:0E:33:4C	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	00:3A:9C:0E:33:4D	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	00:3A:9C:0E:33:4E	Unconfigured	Physical	Sfp Not Present	Disabled
Port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

You have now created two uplink ports on each Fabric Interconnect as shown above. These ports will be used to create Virtual Port Channel in the next section.

Create Uplink Port Channels to Cisco Nexus Switches

In this procedure, two port channels were created one from Fabric A to both Cisco Nexus 93180YC-FX switches and one from Fabric B to both Cisco Nexus 93180YC-FX switches. To configure the necessary port channels in the Cisco UCS environment, follow these steps:

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

2. Click LAN > LAN Cloud >Fabric A.

3. Right-click Port Channels.

4. Select Create Port Channel.

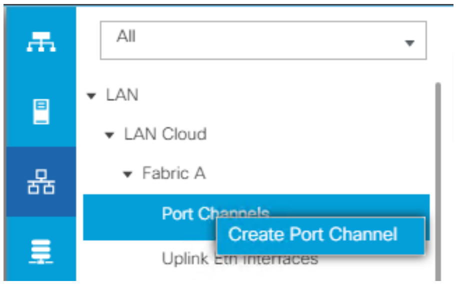

5. Enter 11 as the unique ID of the port channel and name of the port channel.

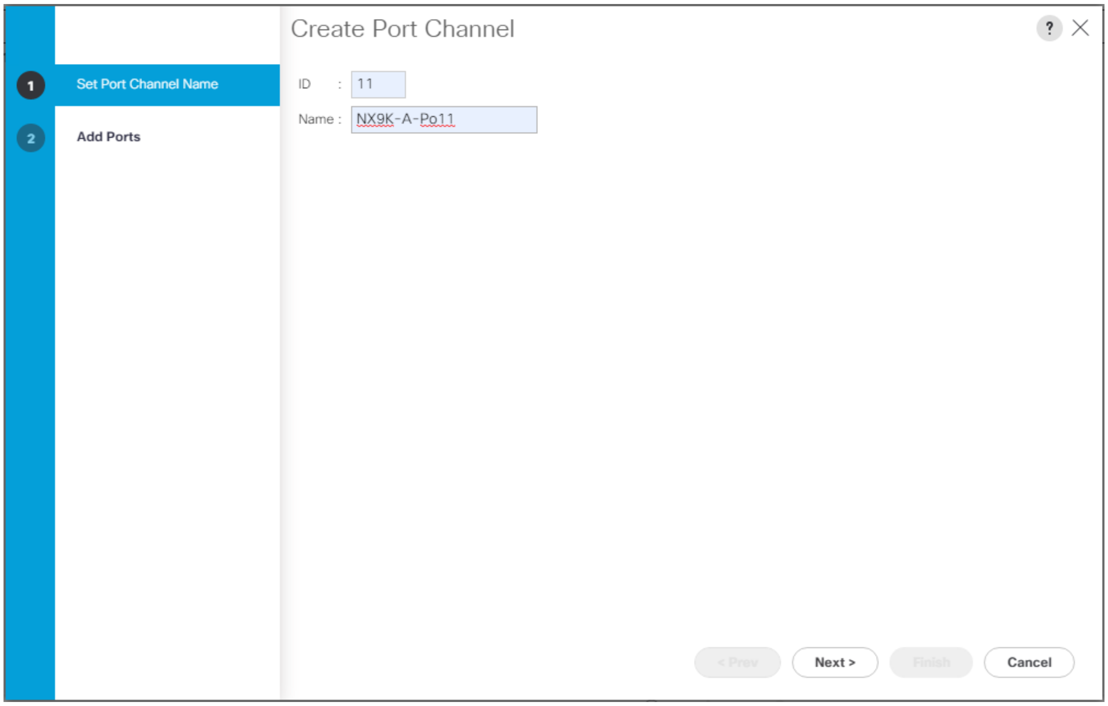
Create Port Channel

ID : 11

Name : NX9K-A-Po11

6. Click Next.

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

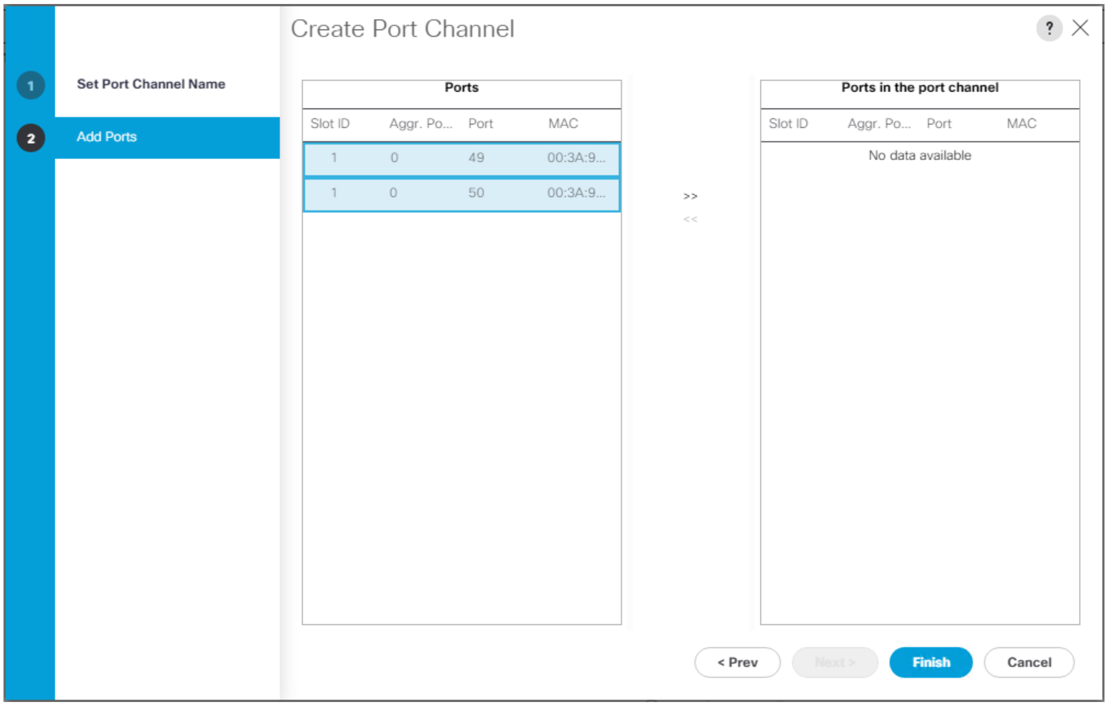

Ports

Slot ID	Aggr. Po...	Port	MAC
1	0	49	00:3A:9...
1	0	50	00:3A:9...

Ports in the port channel

Slot ID	Aggr. Po...	Port	MAC

No data available

8. Click Finish.

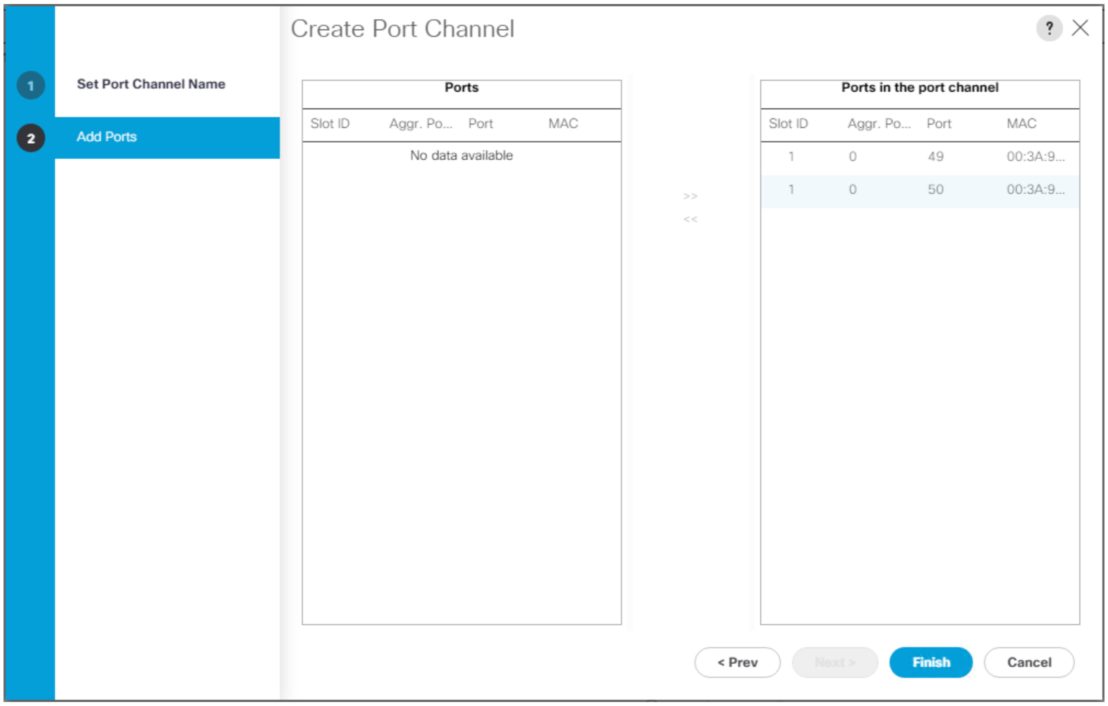

Ports

Slot ID	Aggr. Po...	Port	MAC

No data available

Ports in the port channel

Slot ID	Aggr. Po...	Port	MAC
1	0	49	00:3A:9...
1	0	50	00:3A:9...

9. Click OK.

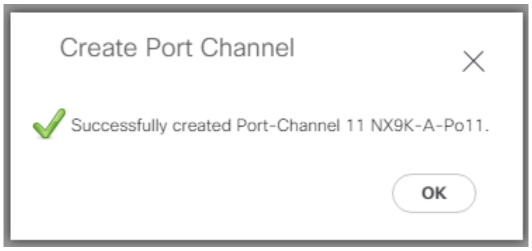

Successfully created Port-Channel 11 NX9K-A-Po11.

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

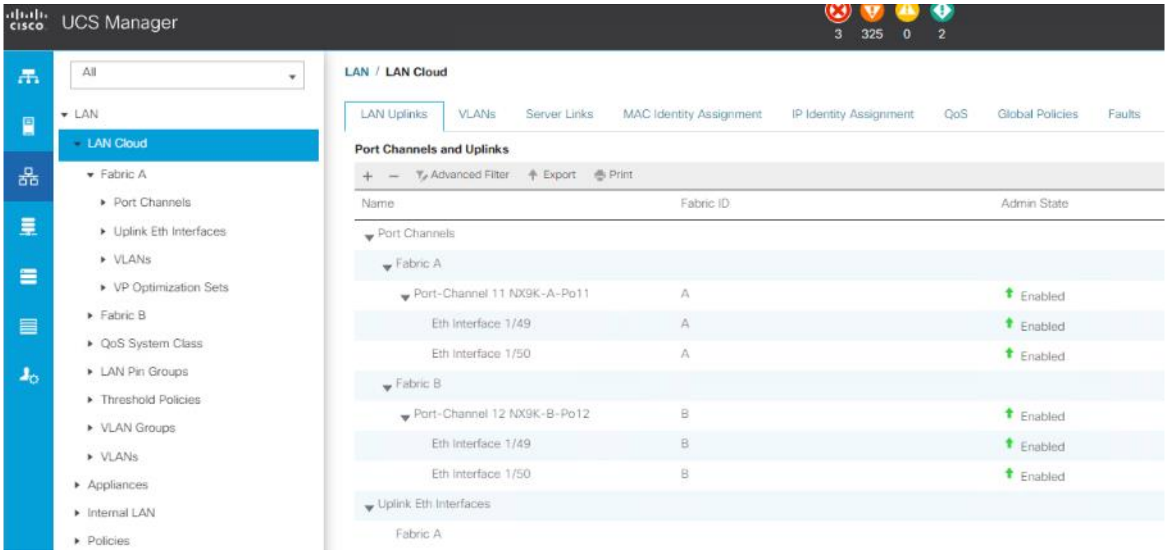

Name	Fabric ID	Admin State
Port Channels		
Fabric A		
Port-Channel 11 NX9K-A-Po11	A	Enabled
Eth Interface 1/49	A	Enabled
Eth Interface 1/50	A	Enabled
Fabric B		
Port-Channel 12 NX9K-B-Po12	B	Enabled
Eth Interface 1/49	B	Enabled
Eth Interface 1/50	B	Enabled
Uplink Eth Interfaces		
Fabric A		

Configure VLAN

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter InBand-Mgmt as the name of the VLAN to be used for Public Network Traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 70 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK.

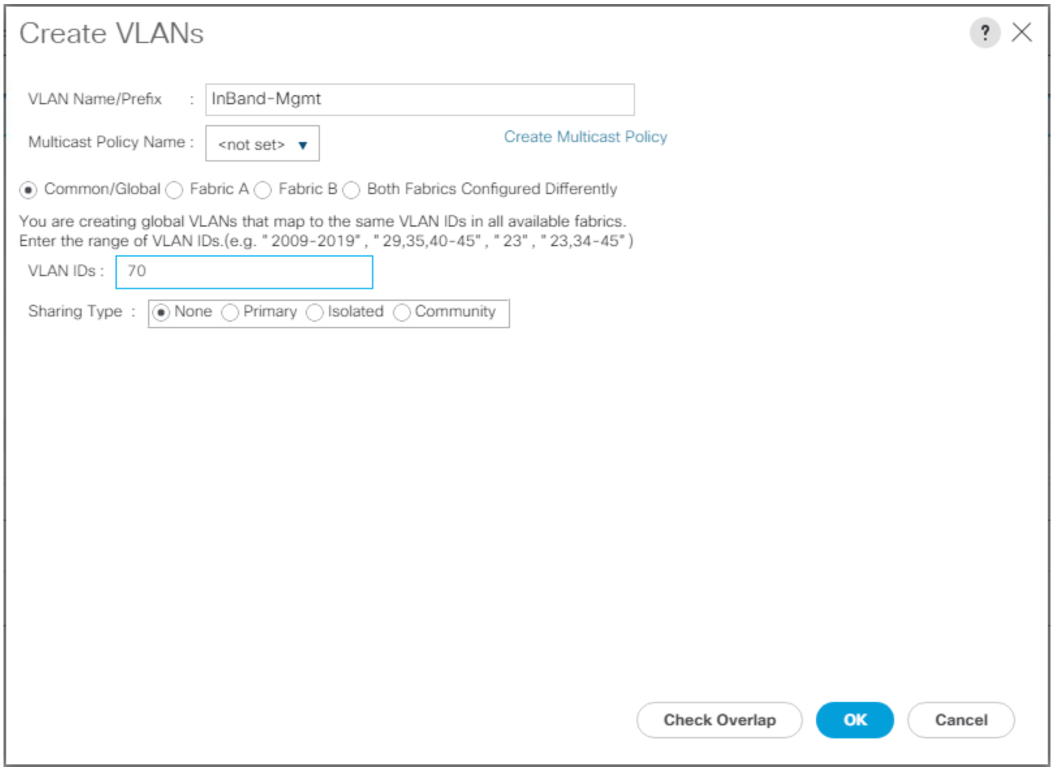

Create VLANs

VLAN Name/Prefix : InBand-Mgmt

Multicast Policy Name : <not set> ▾ Create Multicast Policy

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

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

VLAN IDs : 70

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

10. Repeat steps 1-9 to create required VLANs. Figure 35. shows the VLANs configured for this solution.

Figure 35. VLANs Configured for this Solution

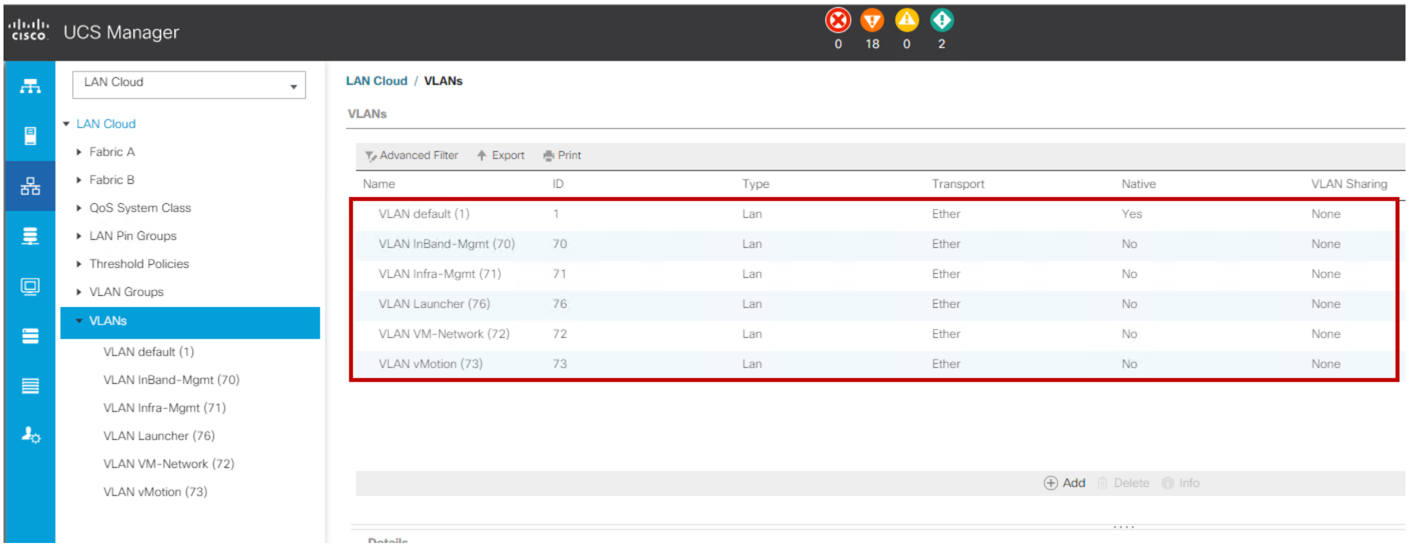

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN InBand-Mgmt (70)	70	Lan	Ether	No	None
VLAN Infra-Mgmt (71)	71	Lan	Ether	No	None
VLAN Launcher (76)	76	Lan	Ether	No	None
VLAN VM-Network (72)	72	Lan	Ether	No	None
VLAN vMotion (73)	73	Lan	Ether	No	None

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

Configure VSAN

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > SAN Cloud.
3. Under VSANs, right-click VSANs.
4. Select Create VSANs.

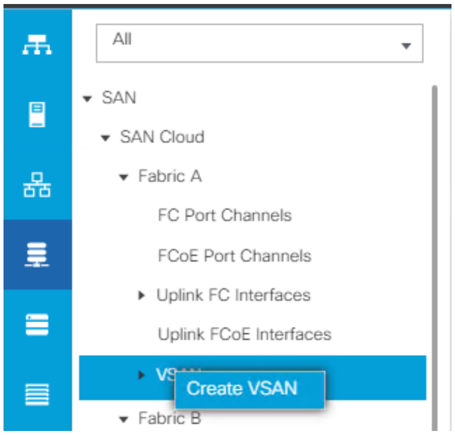

5. Enter the name of the VSAN, such as FlashStack-A.

In this solution, we created two VSANs; VSAN FlashStack-A 100 on the Cisco UCS Fabric A and VSAN FlashStack-B 101 on the Cisco UCS Fabric B for SAN Boot and Storage Access.

6. Select Disabled for FC Zoning

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

7. Select Fabric A for the scope of the VSAN:
   a. Enter 100 as VSAN ID and FCoE VLAN ID.
   b. Click OK.

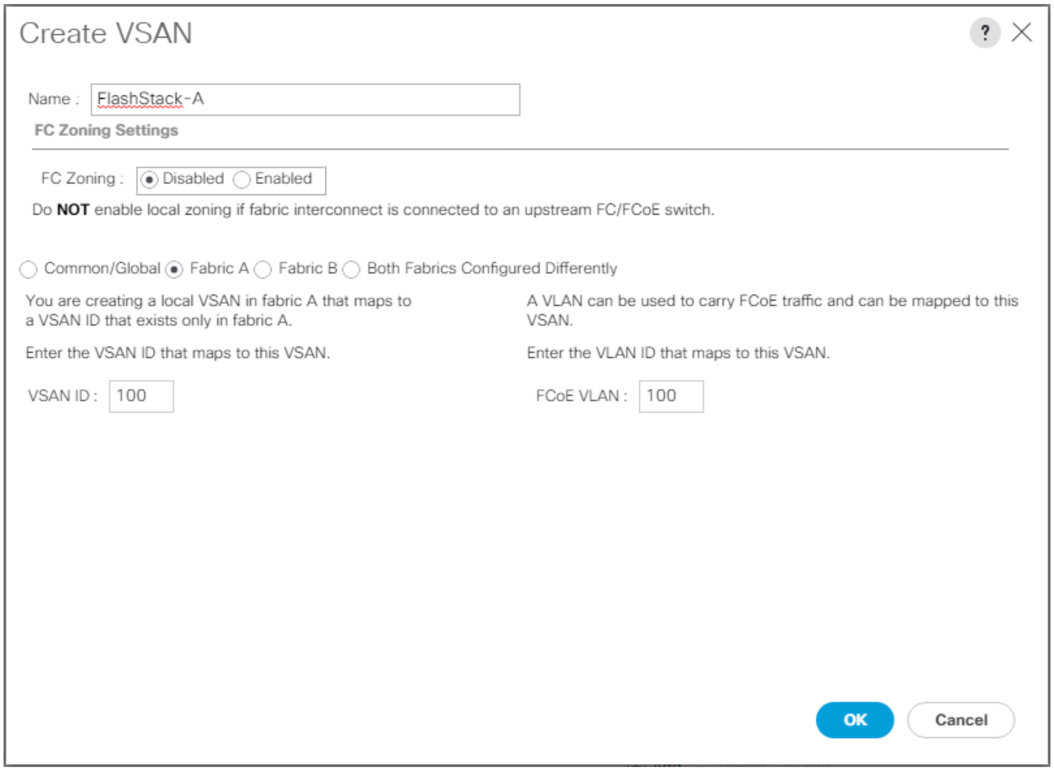

Create VSAN

Name : FlashStack-A

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

○ Common/Global ● Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

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

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 100

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 100

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

Figure 36 shows VSAN 100 and 101 configured for this solution.

Figure 36. VSANs Configured for this Solution

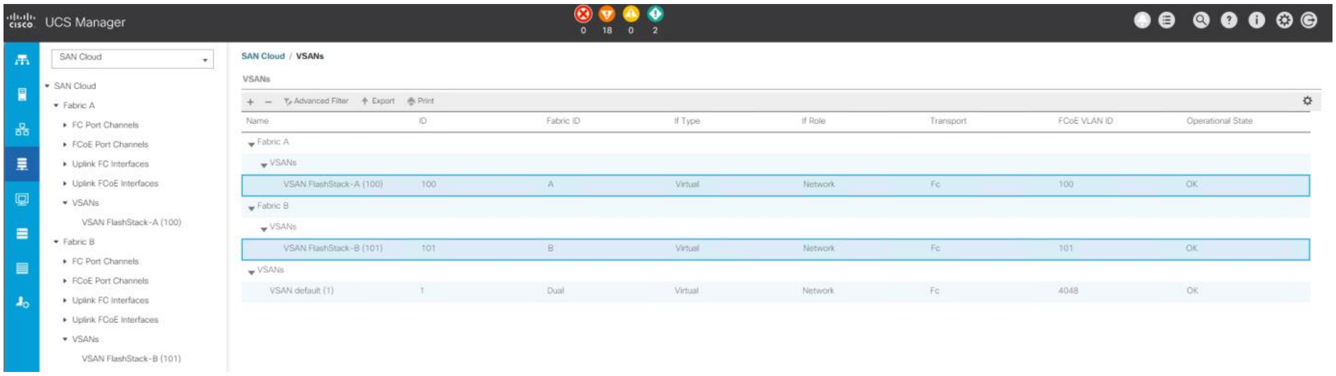

SAN Cloud / VSANs

VSANs

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational State
Fabric A							
VSANs							
VSAN FlashStack-A (100)	100	A	Virtual	Network	Fc	100	OK
Fabric B							
VSANs							
VSAN FlashStack-B (101)	101	B	Virtual	Network	Fc	101	OK
VSANs							
VSAN default (1)	1	Dual	Virtual	Network	Fc	4048	OK

Create New Sub-Organization

To configure the necessary Sub-Organization for the Cisco UCS environment, follow these steps:

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

2. Select root > Sub-Organization.

3. Right-click Sub-Organization.

4. Enter the name of the Sub-Organization.

5. Click OK.

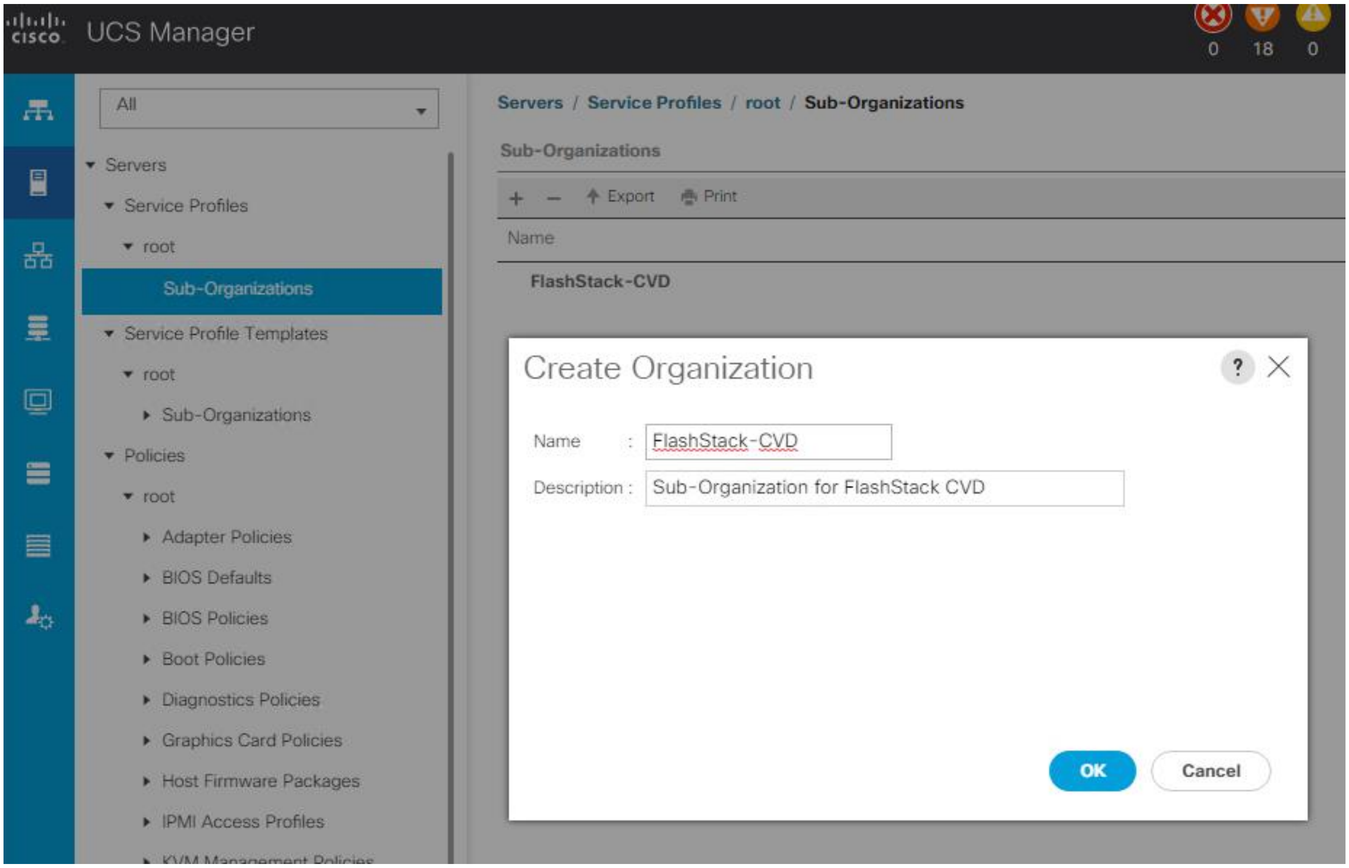

Name : FlashStack-CVD

Description : Sub-Organization for FlashStack CVD

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

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

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. To create a block of IP addresses for server KVM access in the Cisco UCS environment, follow these steps:

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

2. Click Pools > root > Sub-Organizations > FlashStack-CVD > IP Pools > click Create IP Pool.

3. Select the option Sequential to assign IP in sequential order then click Next.

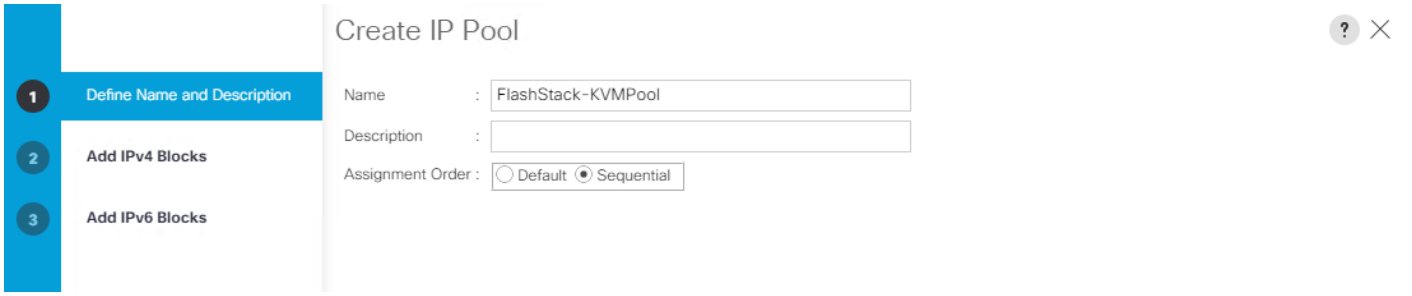

Name : FlashStack-KVMPool

Description :

Assignment Order : ○ Default ● Sequential

4. Click Add IPv4 Block.

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

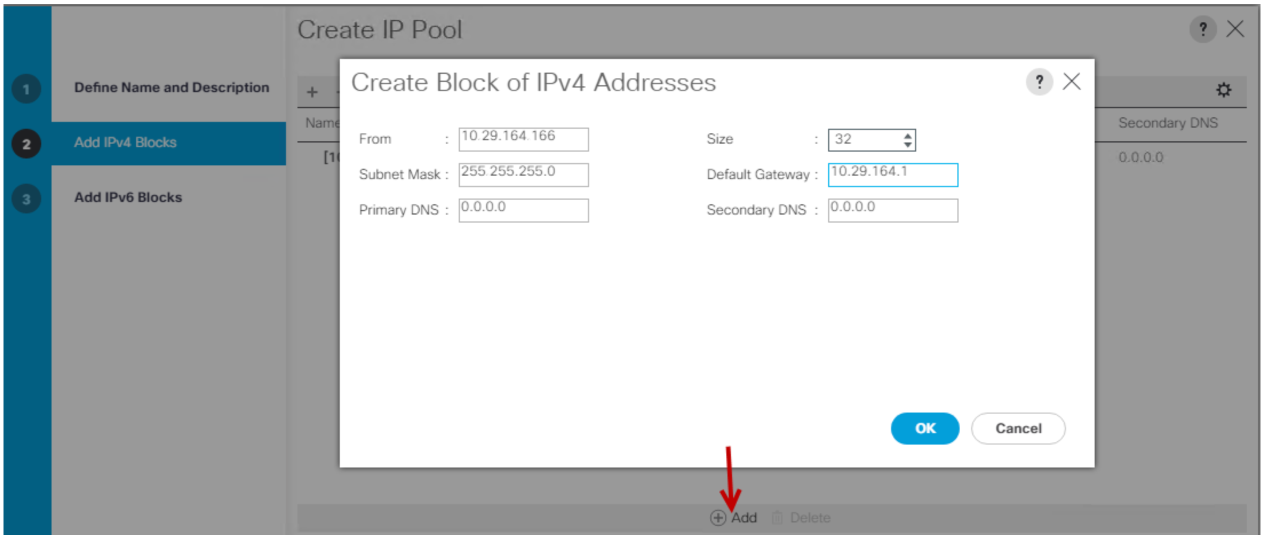

From : 10.29.164.166

Size : 32

Subnet Mask : 255.255.255.0

Default Gateway : 10.29.164.1

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

UUID Suffix Pool Creation

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these steps:

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

2. Click Pools > root > Sub-Organization > FlashStack-CVD.

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

4. Enter the name of the UUID name.

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

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

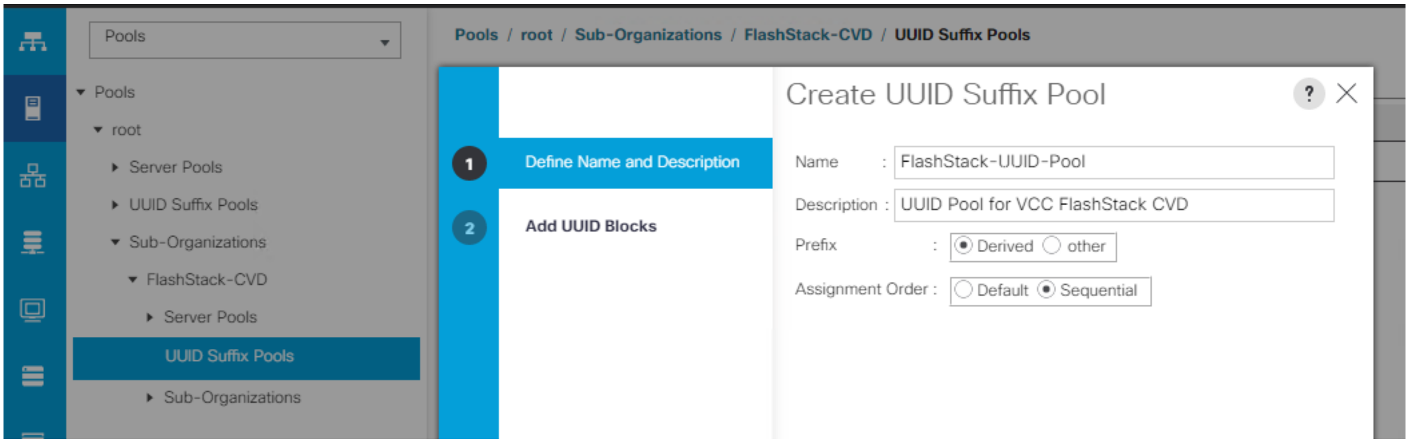

Create UUID Suffix Pool

Name : FlashStack-UUID-Pool

Description : UUID Pool for VCC FlashStack CVD

Prefix : ● Derived ○ other

Assignment Order : ○ Default ● Sequential

7. Click Add to add a block of UUIDs.

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

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

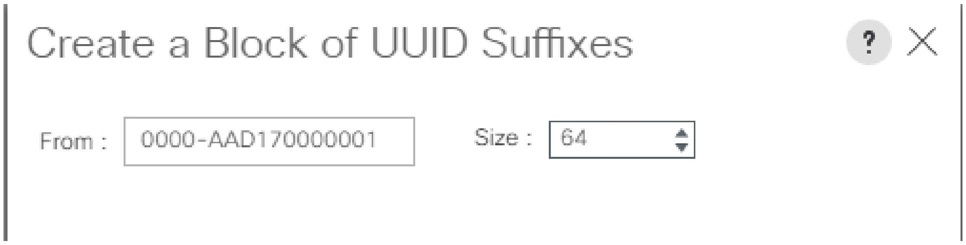

Create a Block of UUID Suffixes

From : 0000-AAD170000001

Size : 64

Server Pool Creation

To configure the necessary server pool for the Cisco UCS environment, follow these steps:

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

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

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

3. Enter name of the server pool.

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

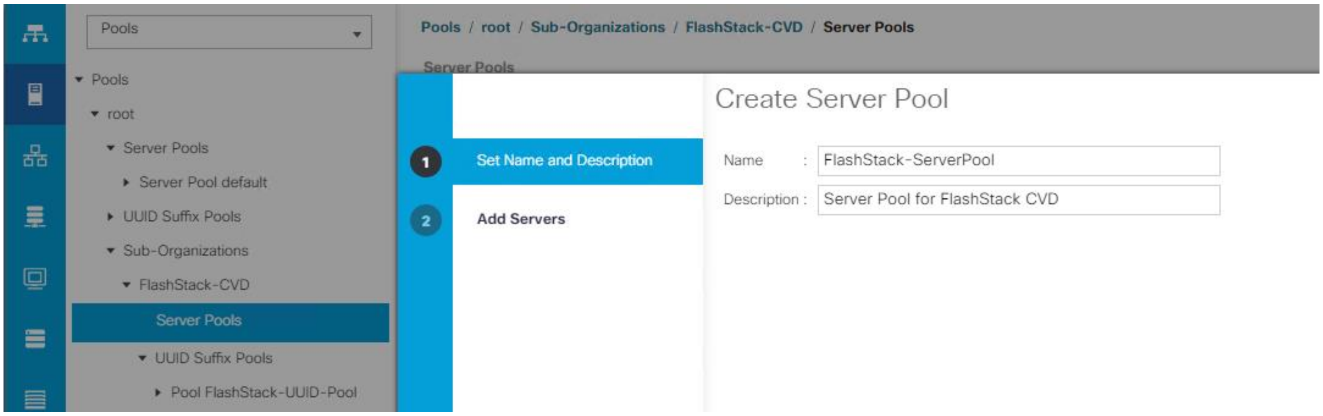

Create Server Pool

Name : FlashStack-ServerPool

Description : Server Pool for FlashStack CVD

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

6. Click Finish and then click OK.

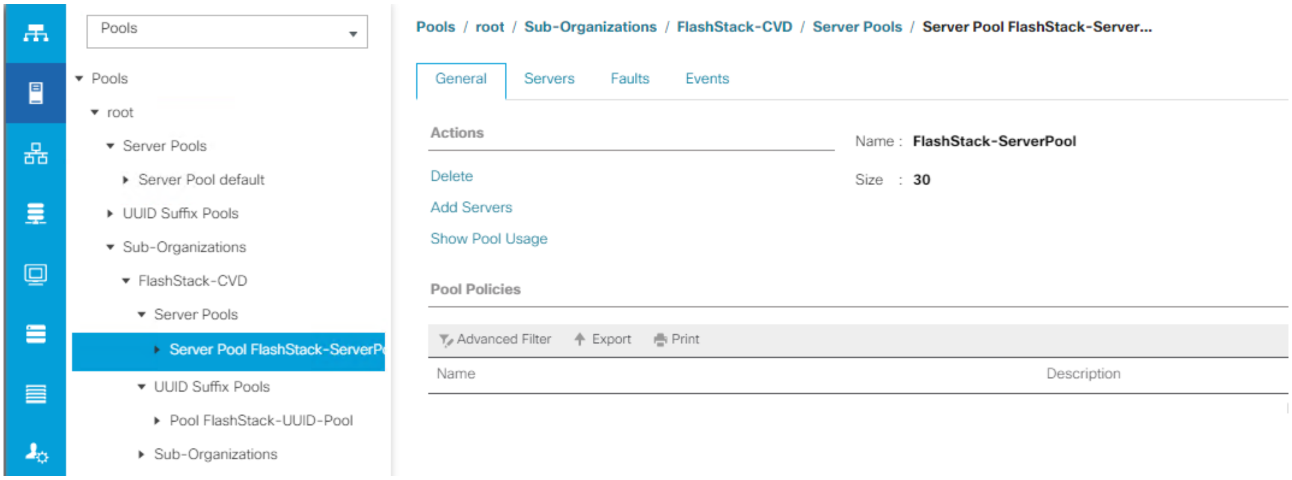

Name : FlashStack-ServerPool

Size : 30

MAC Pool Creation

To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

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

2. Click Pools > root > Sub-Organization > FlashStack > right-click MAC Pools under the root organization.

3. Click Create MAC Pool to create the MAC address pool.

4. Enter name for MAC pool. Select Assignment Order as “Sequential.”

5. Enter the seed MAC address and provide the number of MAC addresses to be provisioned.

6. Click OK and then click Finish.

7. In the confirmation message, click OK.

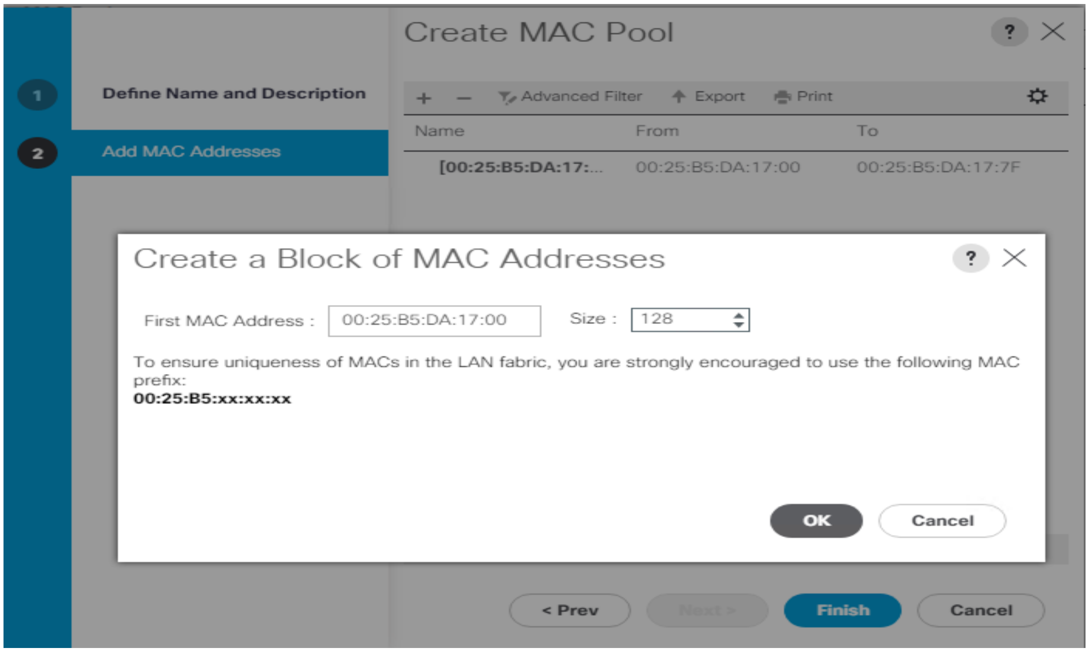

Create MAC Pool

1 Define Name and Description

2 Add MAC Addresses

Name	From	To
[00:25:B5:DA:17:...	00:25:B5:DA:17:00	00:25:B5:DA:17:7F

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:DA:17:00 Size : 128

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

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

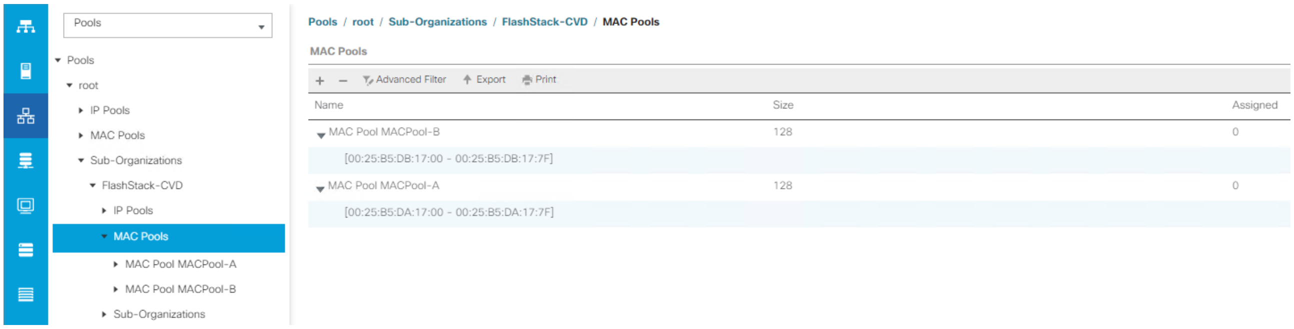

Pools / root / Sub-Organizations / FlashStack-CVD / **MAC Pools**

MAC Pools

Name	Size	Assigned
MAC Pool MACPool-B	128	0
[00:25:B5:DB:17:00 - 00:25:B5:DB:17:7F]		
MAC Pool MACPool-A	128	0
[00:25:B5:DA:17:00 - 00:25:B5:DA:17:7F]		

WWNN and WWPN Pool Creation

To configure the necessary WWNN pools for the Cisco UCS environment, follow these steps:

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

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

3. Assign name and Assignment Order as sequential.

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

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

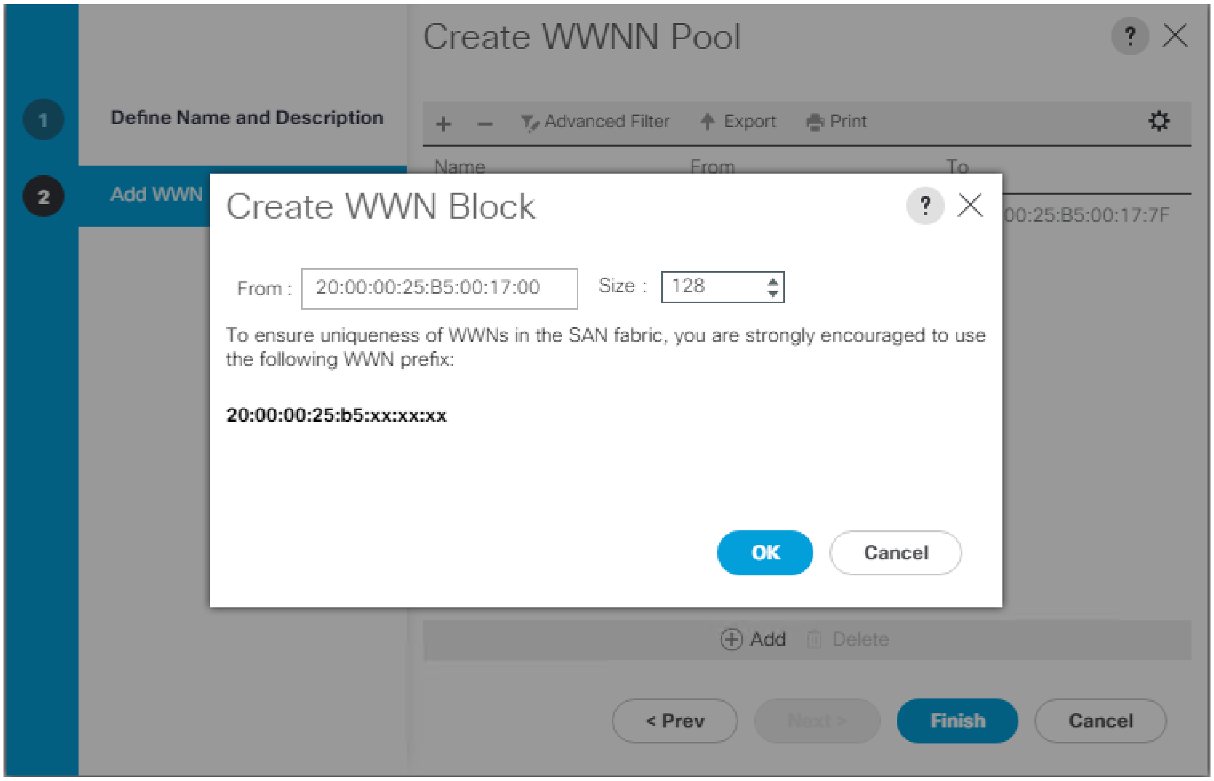
Create WWN Block

From : 20:00:00:25:B5:00:17:00 Size : 128

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

6. Click OK and then click Finish.

To configure the necessary WWPN pools for the Cisco UCS environment, follow these steps:

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > Root > WWPN Pools > right-click WWPN Pools > select Create WWPN Pool.
3. Assign name and Assignment Order as sequential.
4. Click Next and then click Add to add block of Ports.
5. Enter Block for WWN and size.
6. Click OK and then click Finish.

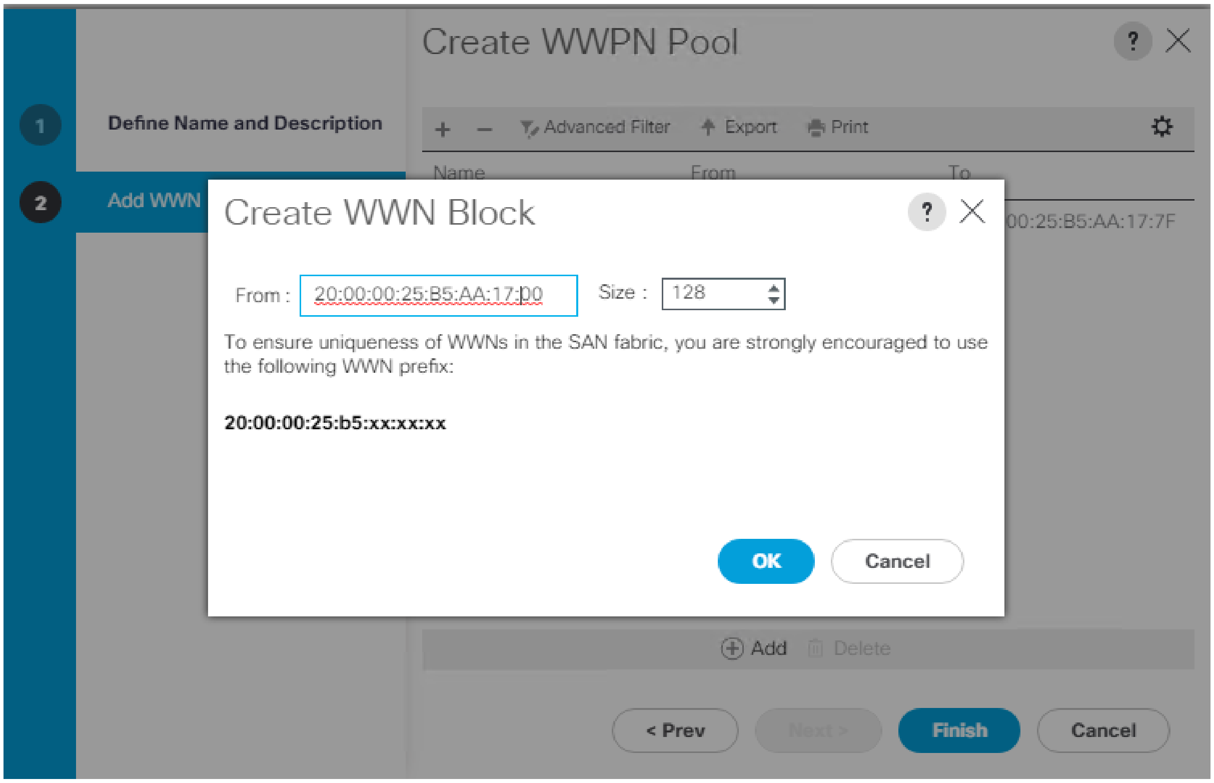

Create WWN Block

From : 20:00:00:25:B5:AA:17:00 Size : 128

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

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

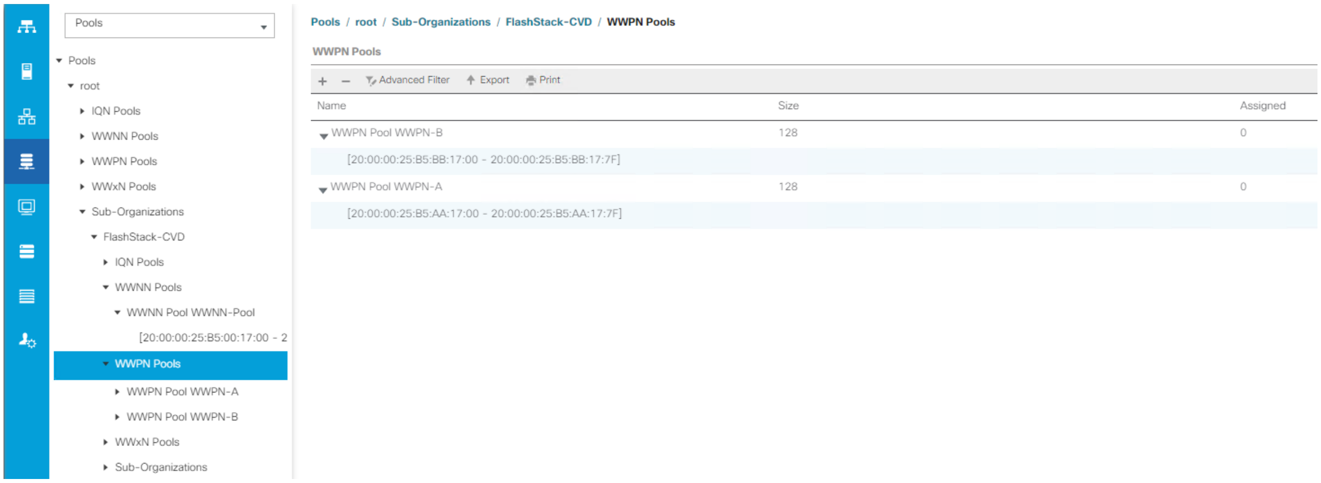

Pools / root / Sub-Organizations / FlashStack-CVD / **WWPN Pools**

WWPN Pools

Name	Size	Assigned
WWPN Pool WWPN-B	128	0
[20:00:00:25:B5:BB:17:00 - 20:00:00:25:B5:BB:17:7F]		
WWPN Pool WWPN-A	128	0
[20:00:00:25:B5:AA:17:00 - 20:00:00:25:B5:AA:17:7F]		

Set Jumbo Frames in both the Cisco Fabric Interconnect

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, follow these steps:

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

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

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

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

5. Click Save Changes.

6. Click OK.

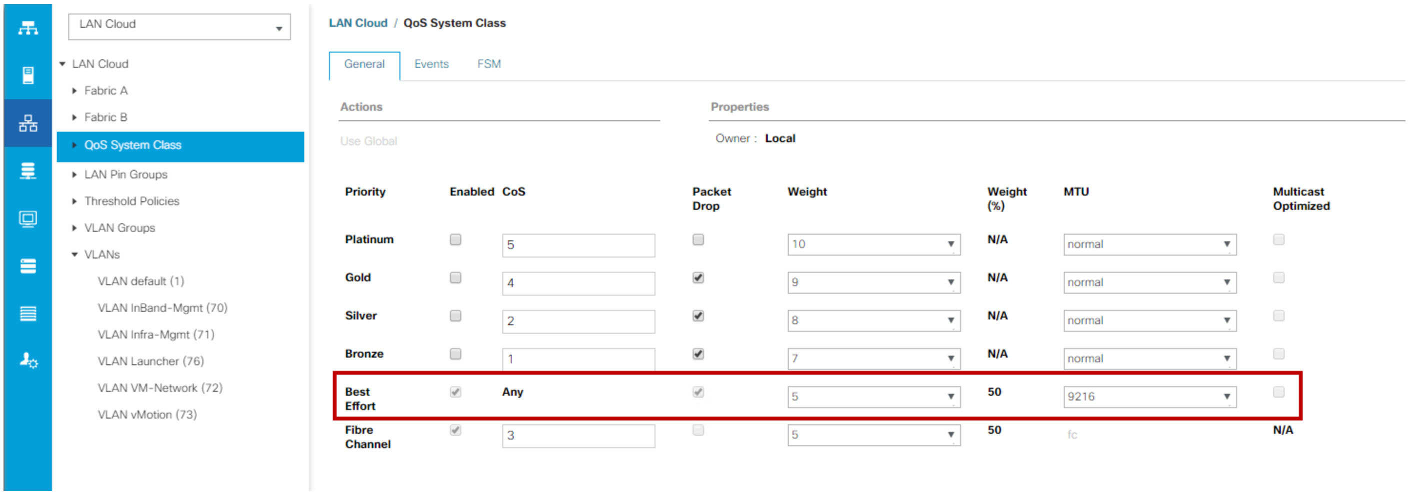

Owner : Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☐	5	☐	10	N/A	normal	☐
Gold	☐	4	☑	9	N/A	normal	☐
Silver	☐	2	☑	8	N/A	normal	☐
Bronze	☐	1	☑	7	N/A	normal	☐
Best Effort	☑	Any	☑	5	50	9216	☐
Fibre Channel	☑	3	☐	5	50	fc	N/A

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.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, follow these steps:

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

2. Select root > Sub-Organization > FlashStack-CVD > Host Firmware Packages.

3. Right-click Host Firmware Packages.

4. Select Create Host Firmware Package.

5. Enter name of the host firmware package.

6. Leave Simple selected.

7. Select the version 4.1(2a) for both the Blade Package.

8. Click OK to create the host firmware package.

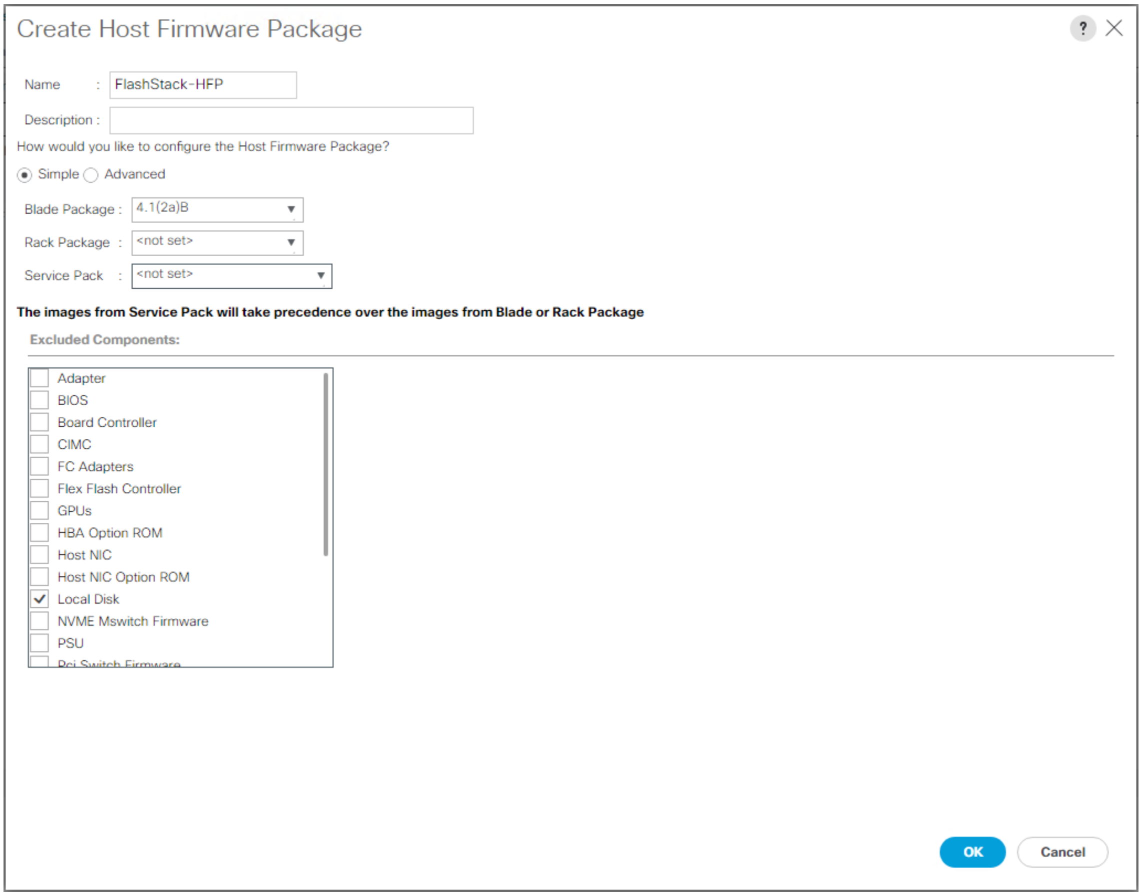
Create Host Firmware Package

Name : FlashStack-HFP

Description :

How would you like to configure the Host Firmware Package?

● Simple ○ Advanced

Blade Package : 4.1(2a)B

Rack Package : <not set>

Service Pack : <not set>

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

Create Server Pool Policy

Create Server Pools Policy

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

To create a Server Pools Policy, follow these steps:

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

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

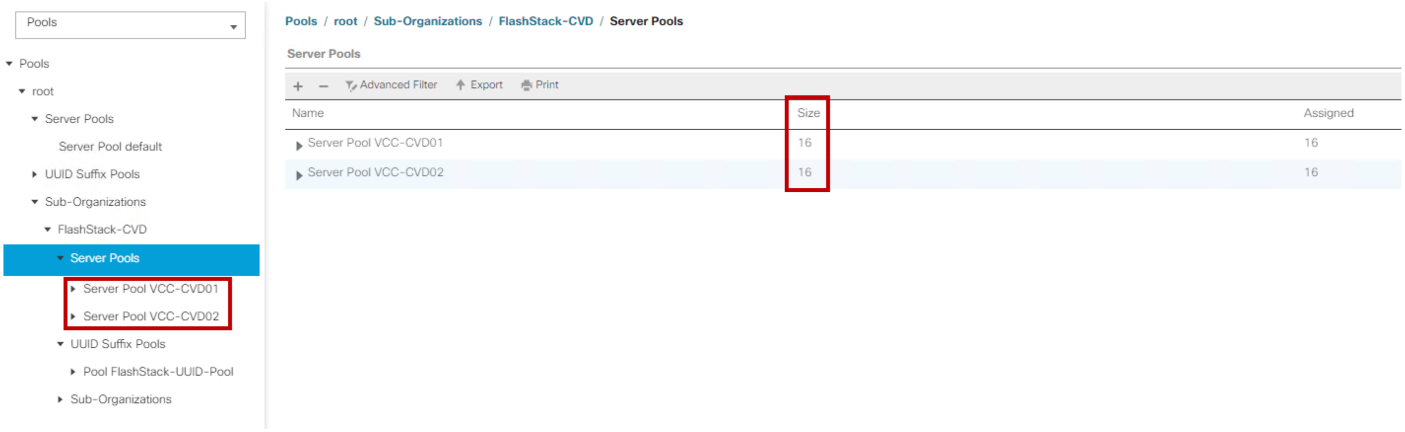

Pools / root / Sub-Organizations / FlashStack-CVD / Server Pools

Server Pools

Name	Size	Assigned
Server Pool VCC-CVD01	16	16
Server Pool VCC-CVD02	16	16

Create Server Pool Policy Qualifications

To create a Server Pool Policy Qualification Policy complete following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD > Server Pool Policy Qualification.
3. Right-click Server Pools Select Create Server Pool Policy Qualification; Enter Policy name.
4. Select Chassis/Server Qualification from left pane to add in Qualifications.
5. Click Add or OK to either Add more servers to existing policy to Finish creation of Policy.

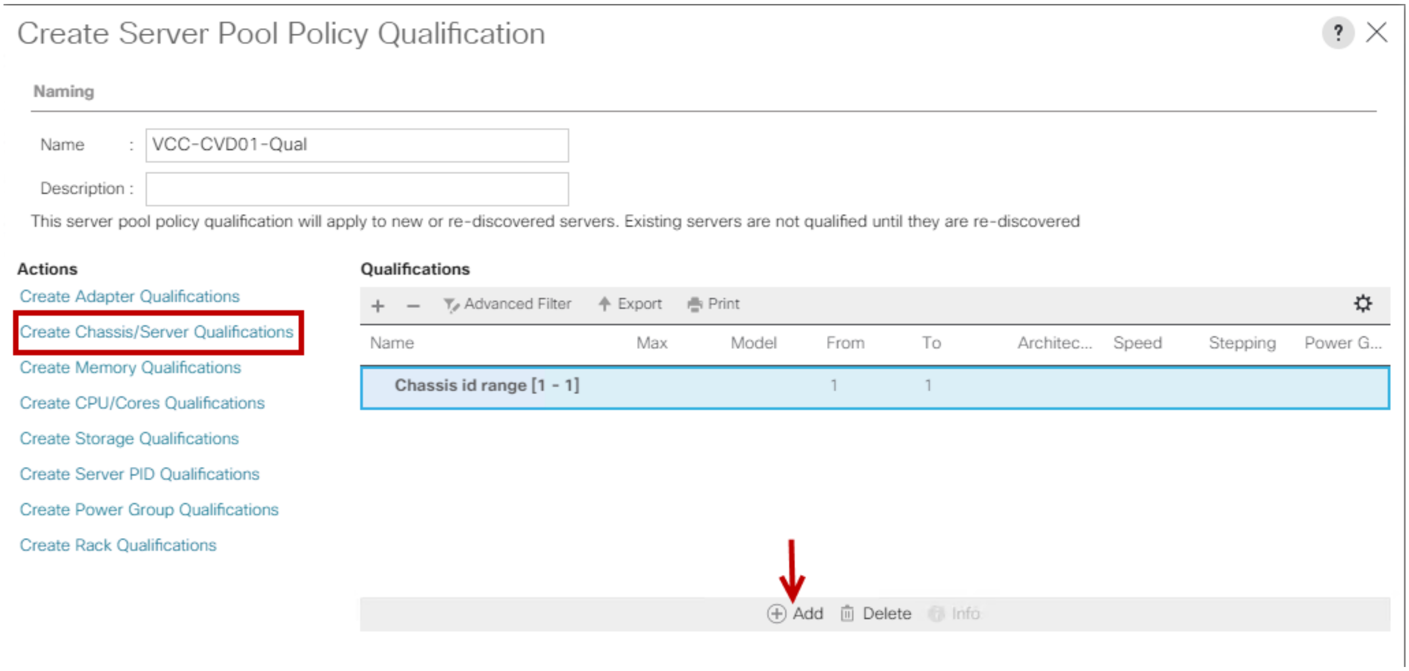
Actions
Create Adapter Qualifications
Create Chassis/Server Qualifications
Create Memory Qualifications
Create CPU/Cores Qualifications
Create Storage Qualifications
Create Server PID Qualifications
Create Power Group Qualifications
Create Rack Qualifications
+ − Advanced Filter Export Print
Add Delete Info

Create Server Pool Policy Qualification

Naming

Name : VCC-CVD01-Qual

Description :

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

Qualifications

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

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

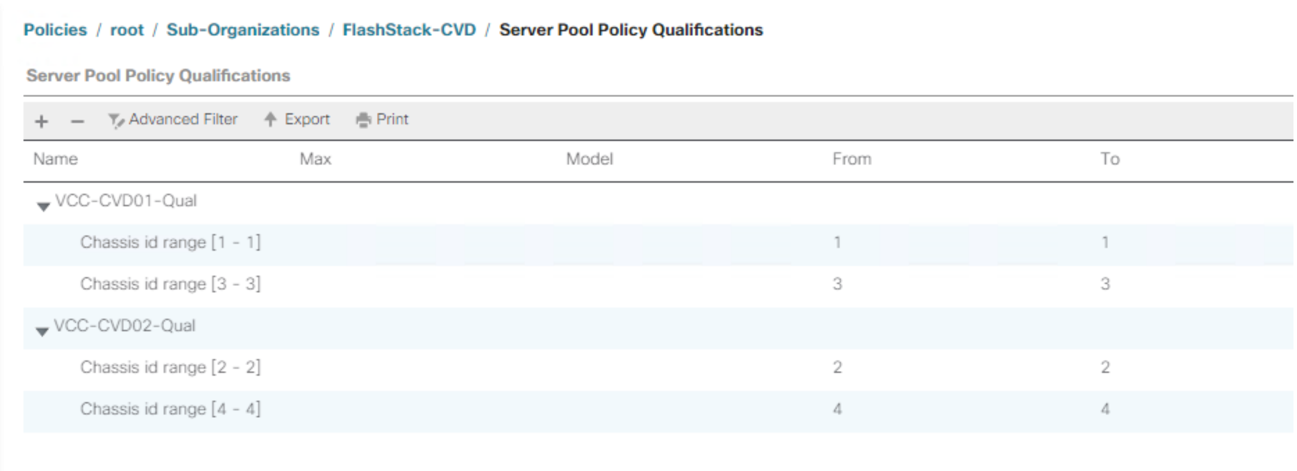

Policies / root / Sub-Organizations / FlashStack-CVD / Server Pool Policy Qualifications

Server Pool Policy Qualifications

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

To create a Server Pool Policy, follow these steps:

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

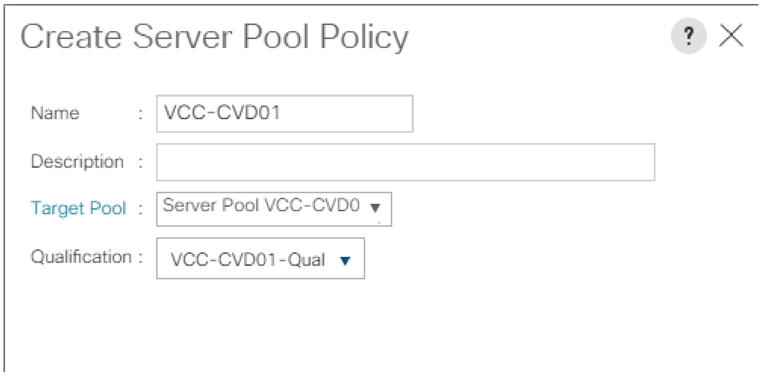

Create Server Pool Policy

Name : VCC-CVD01

Description :

Target Pool : Server Pool VCC-CVD0 ▾

Qualification : VCC-CVD01-Qual ▾

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

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, follow these steps:

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

2. Click Policies > root > Sub-Organization > FlashStack-CVD > Network Control Policies.

3. Right-click Network Control Policies.

4. Click Create Network Control Policy.

5. Enter policy name.

6. Select the Enabled option for "CDP."

7. Click OK to create the network control policy.

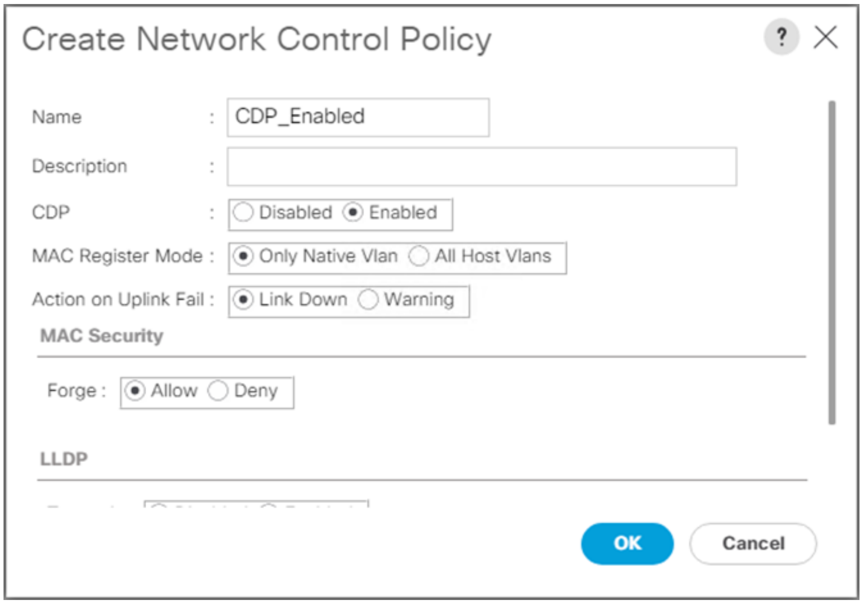

Create Network Control Policy

Name : CDP_Enabled

Description :

CDP : ○ Disabled ● Enabled

MAC Register Mode : ● Only Native Vlan ○ All Host Vlans

Action on Uplink Fail : ● Link Down ○ Warning

MAC Security

Forge : ● Allow ○ Deny

LLDP

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, follow these steps:

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

2. Click Policies > root > Sub-Organization > FlashStack-CVD > Power Control Policies.

3. Right-click Power Control Policies.

4. Click Create Power Control Policy.

5. Select Fan Speed Policy as "Max Power."

6. Enter NoPowerCap as the power control policy name.

7. Change the power capping setting to No Cap.

8. Click OK to create the power control policy.

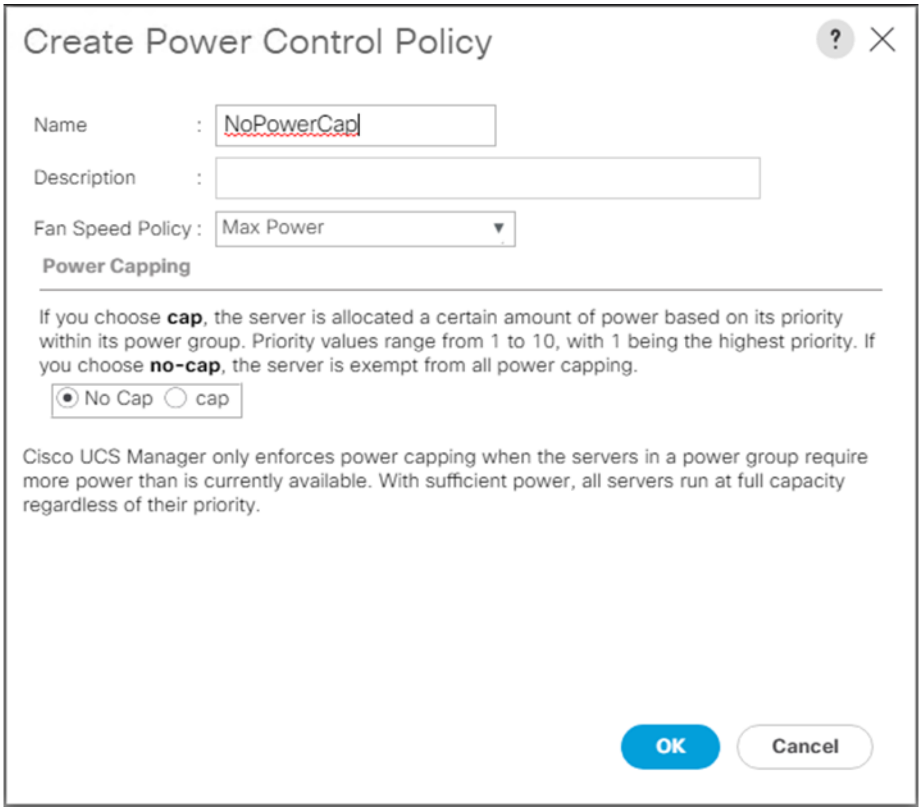

Create Power Control Policy

Name : NoPowerCap

Description :

Fan Speed Policy : Max Power

Power Capping

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

● No Cap ○ cap

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

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, follow these steps:

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

2. Click Policies > root > Sub-Organization > FlashStack-CVD > BIOS Policies.

3. Right-click BIOS Policies.

4. Click Create BIOS Policy.

5. Enter B200-M5-BIOS as the BIOS policy name.

6. Click OK to create policy.”

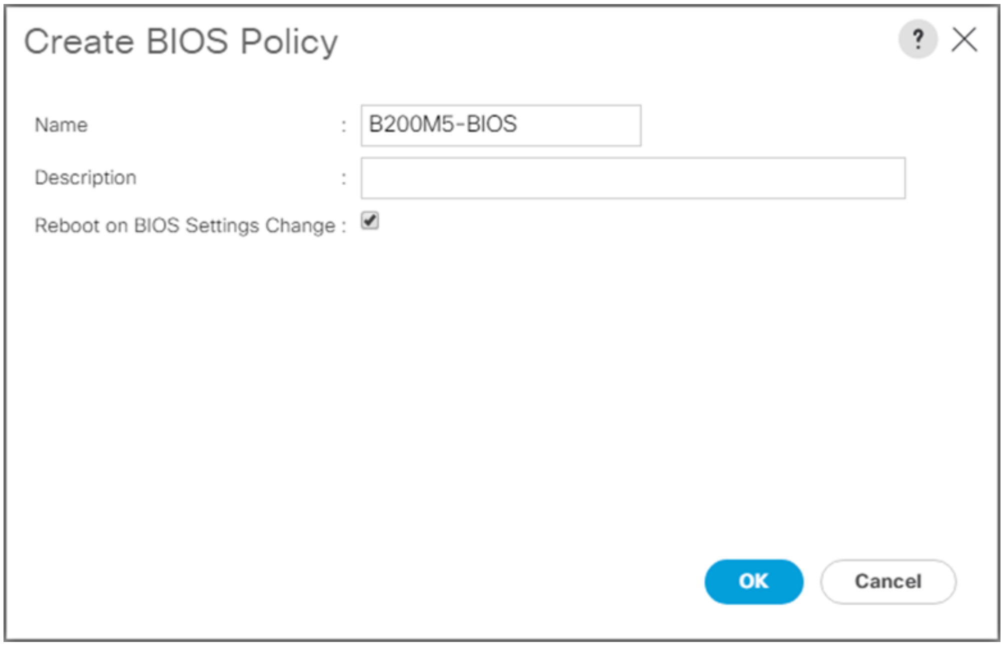

Create BIOS Policy

Name : B200M5-BIOS

Description :

Reboot on BIOS Settings Change : ☑

7. Leave all BIOS Settings as “Platform Default.”

Configure Maintenance Policy

To update the default Maintenance Policy, follow these steps:

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

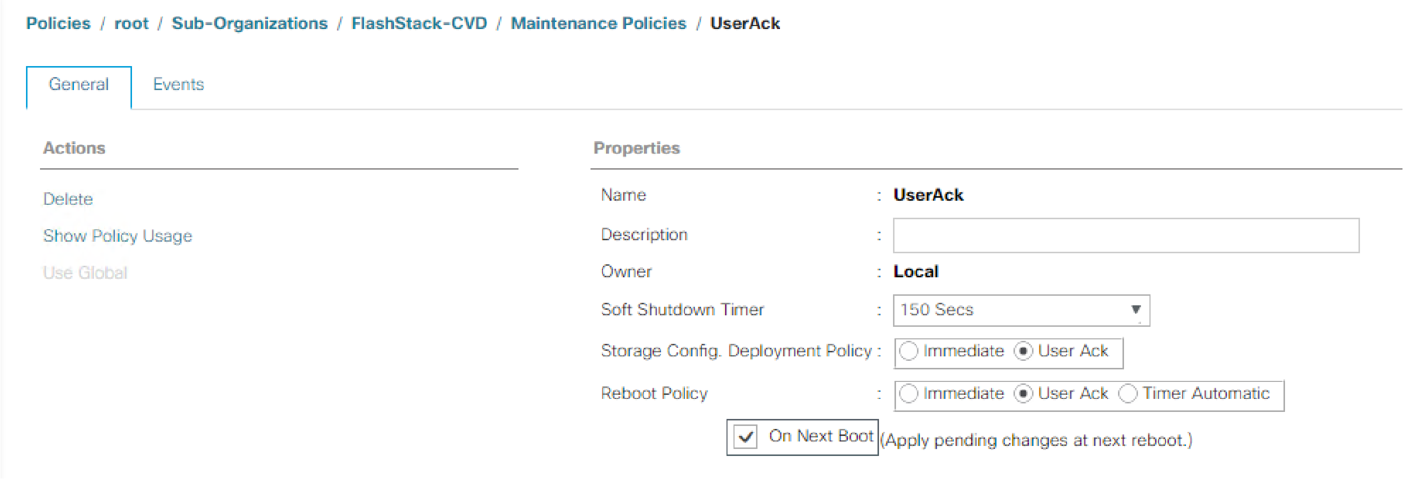

Policies / root / Sub-Organizations / FlashStack-CVD / Maintenance Policies / UserAck

Name : UserAck

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : ○ Immediate ● User Ack

Reboot Policy : ○ Immediate ● User Ack ○ Timer Automatic

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

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > vNIC Template.
3. Right-click vNIC Templates.
4. Click Create vNIC Template.
5. Enter name for vNIC template.
6. Keep Fabric A selected. Do not select the Enable Failover checkbox.
7. For Redundancy Type, Select “Primary Template.”
8. Select Updating Template as the Template Type.
9. Under VLANs, select the checkboxes for desired VLANs to add as part of the vNIC Template.
10. Set Native-VLAN as the native VLAN.
11. For MTU, enter 9000.
12. In the MAC Pool list, select MAC Pool configure for Fabric A.
13. In the Network Control Policy list, select CDP_Enabled.
14. Click OK to create the vNIC template.

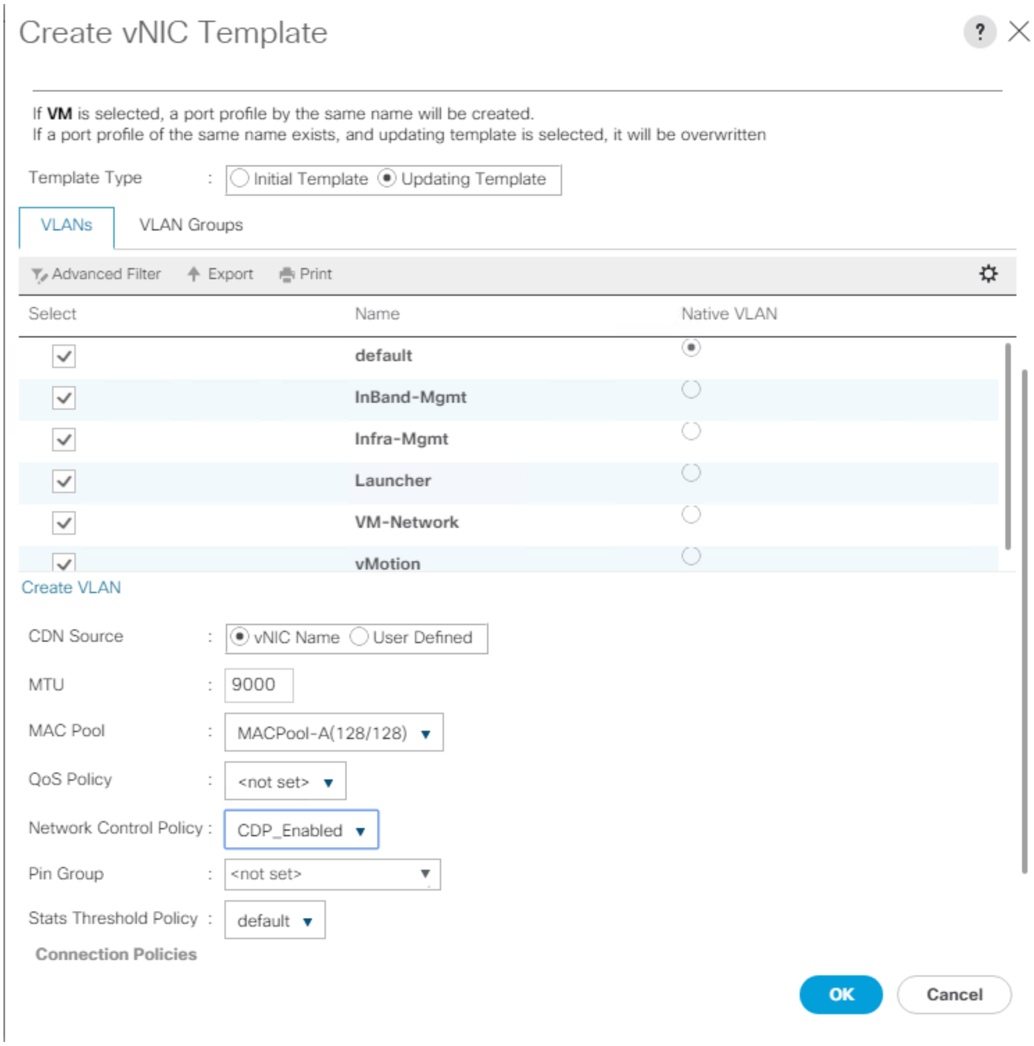

Create vNIC Template

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

Select	Name	Native VLAN
☑	default	●
☑	InBand-Mgmt	○
☑	Infra-Mgmt	○
☑	Launcher	○
☑	VM-Network	○
☑	vMotion	○

Create VLAN

CDN Source : ● vNIC Name ○ User Defined

MTU : 9000

MAC Pool : MACPool-A(128/128)

QoS Policy : <not set>

Network Control Policy : CDP_Enabled

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

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

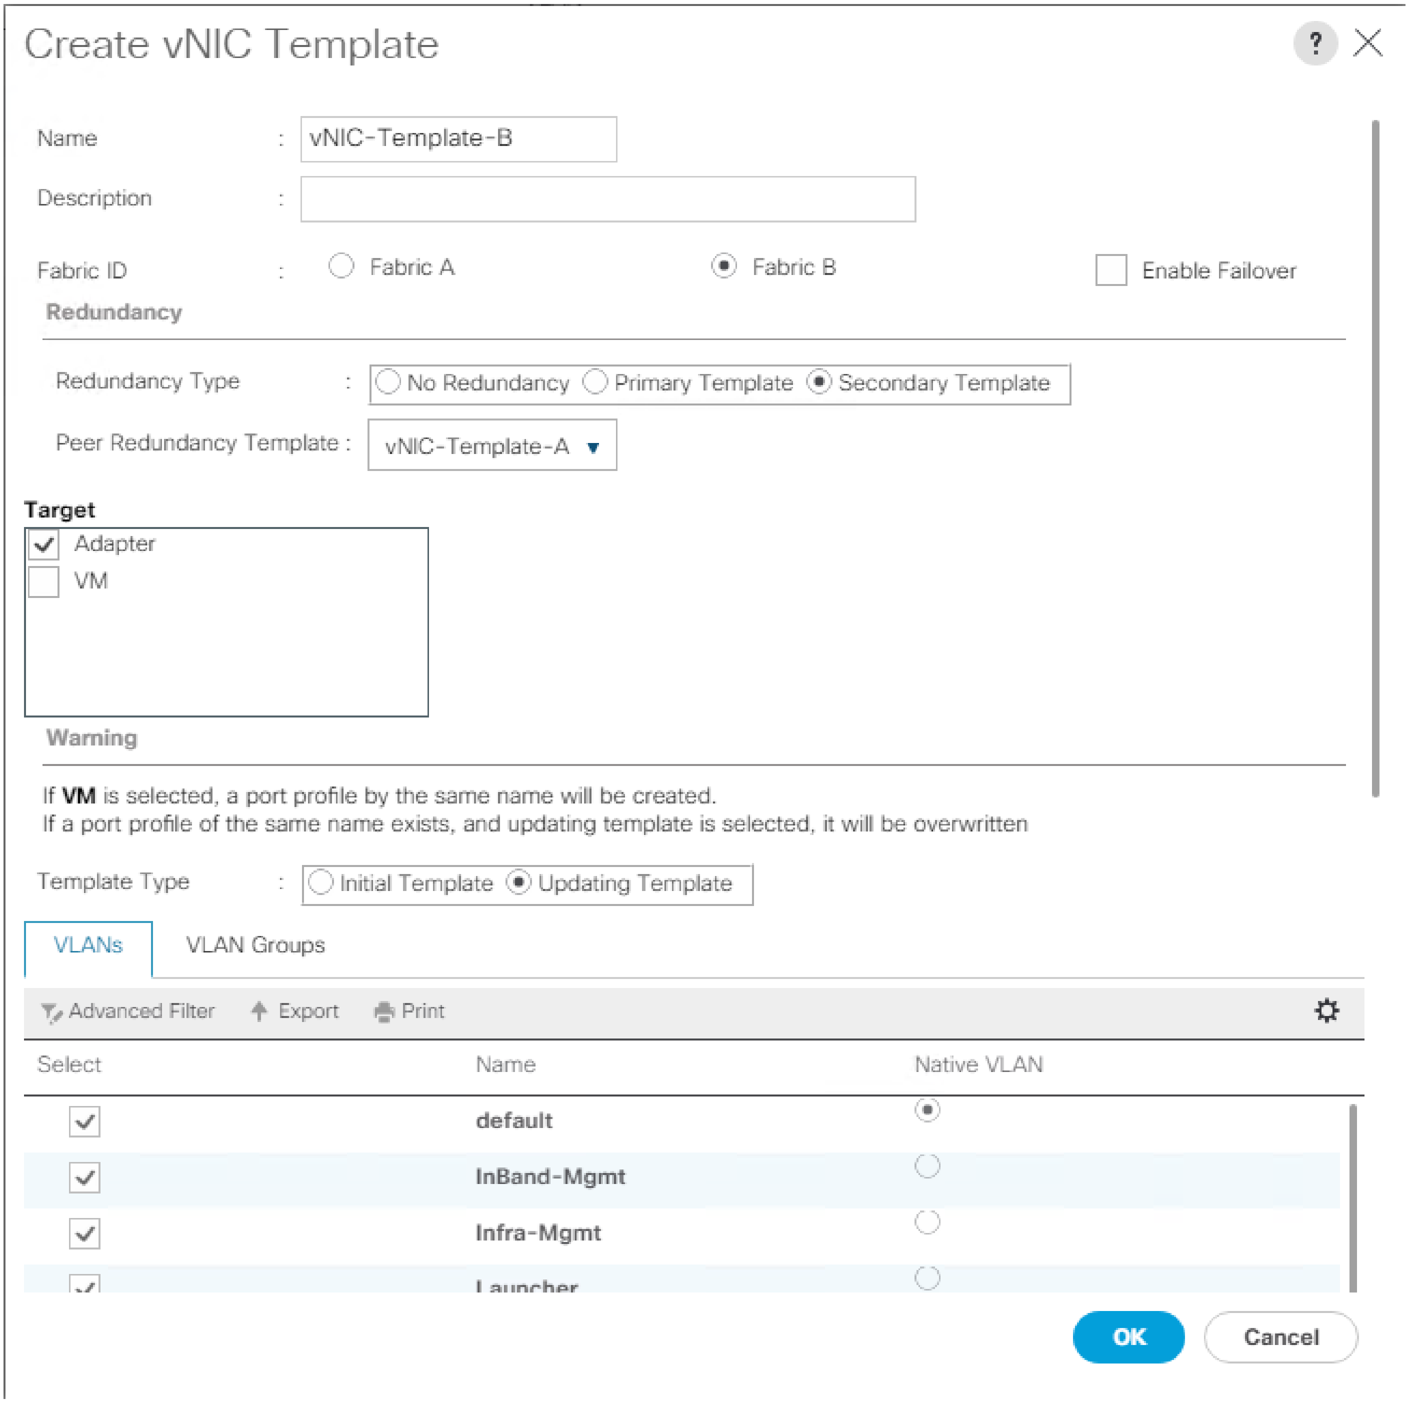
Create vNIC Template

Name : vNIC-Template-B

Description :

Fabric ID : ○ Fabric A ● Fabric B □ Enable Failover

Redundancy

Redundancy Type : ○ No Redundancy ○ Primary Template ● Secondary Template

Peer Redundancy Template : vNIC-Template-A

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

Select	Name	Native VLAN
☑	**default**	●
☑	**InBand-Mgmt**	○
☑	**Infra-Mgmt**	○
☑	**Launcher**	○

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

Create vHBA Templates

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, follow these steps:

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

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

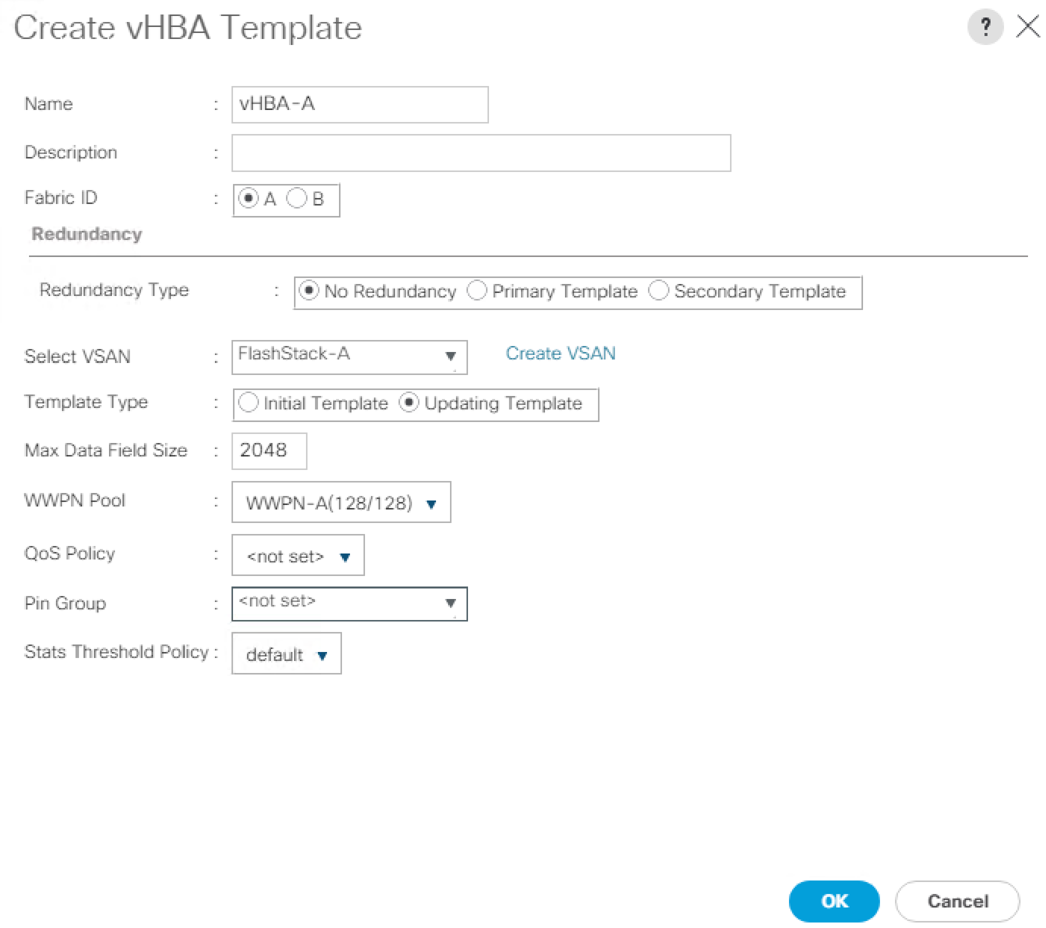

Create vHBA Template

Name : vHBA-A

Description :

Fabric ID : ● A ○ B

Redundancy

Redundancy Type : ● No Redundancy ○ Primary Template ○ Secondary Template

Select VSAN : FlashStack-A ▾ Create VSAN

Template Type : ○ Initial Template ● Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-A(128/128) ▾

QoS Policy : <not set> ▾

Pin Group : <not set> ▾

Stats Threshold Policy : default ▾

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

Create Server Boot Policy for SAN Boot

All Cisco UCS B200 M5 Blade Servers for the workload and the two Infrastructure servers were set to boot from SAN for this Cisco Validated Design as part of the Service Profile template. The benefits of booting from SAN are numerous; disaster recovery, lower cooling, and power requirements for each server since a local drive is not required, and better performance, to name just a few.

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.

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

To configure Local disk policy, follow these steps:

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

2. Click OK to create the policy.

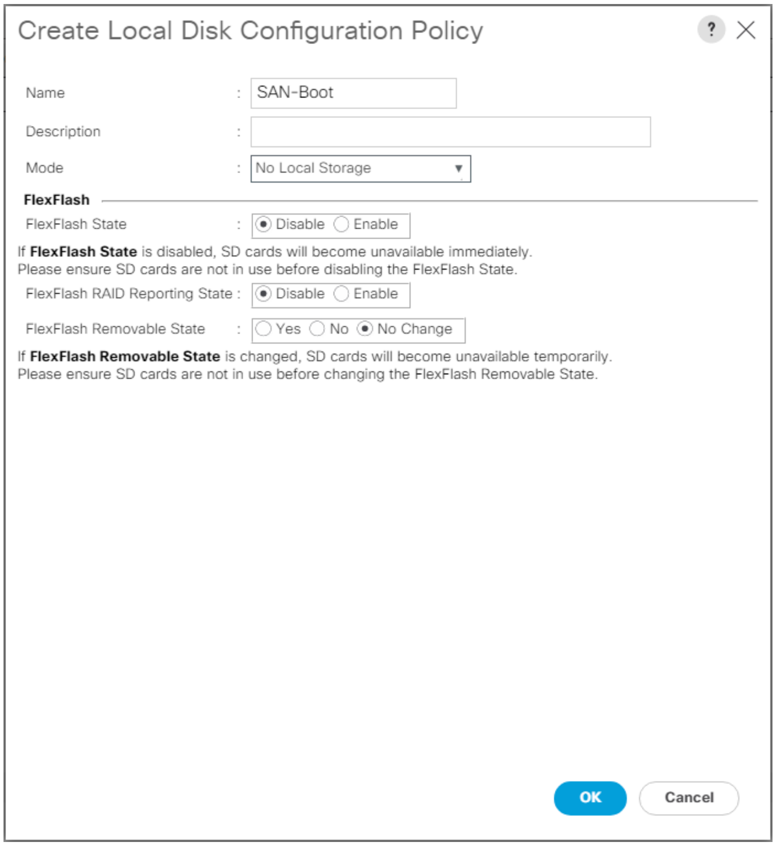

Create Local Disk Configuration Policy

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.

As shown below, the Pure Storage FlashArray have four active FC connections that pair with the Cisco MDS 9132T 32-Gb switches. Two FC ports are connected to Cisco MDS-A and the other Two FC ports are connected to Cisco MDS-B Switches. All FC ports are 32 Gb/s. The SAN Port CT0.FC0 of Pure Storage FlashArray Controller 0 is connected to Cisco MDS Switch A and SAN port CT0.FC2 is connected to MDS Switch B. The SAN Port CT1.FC0 of Pure Storage FlashArray Controller 1 is connected to Cisco MDS Switch A and SAN port CT1.FC2 connected to MDS Switch B.

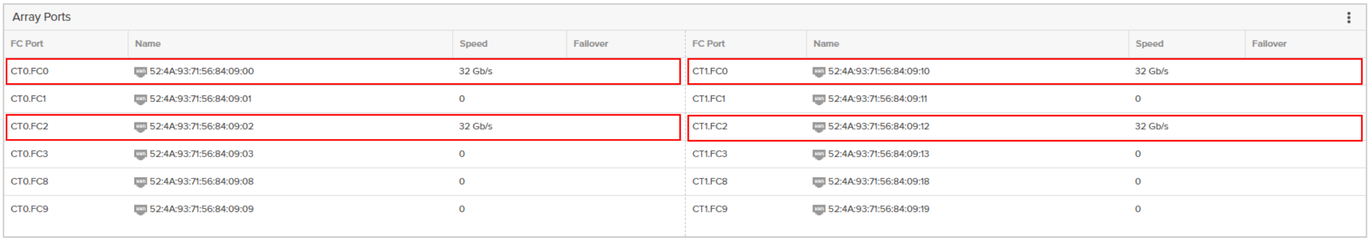

Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Create SAN Policy A

The SAN-A boot policy configures the SAN Primary's primary-target to be port CT0.FC0 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC0 on the Pure Storage cluster. Similarly, the SAN Secondary’s primary-target should be port CT1.FC2 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC2 on the Pure Storage cluster.

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

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

To create Boot Policies for the Cisco UCS environments, follow these steps:

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

2. Enter SAN-A as the name of the boot policy.

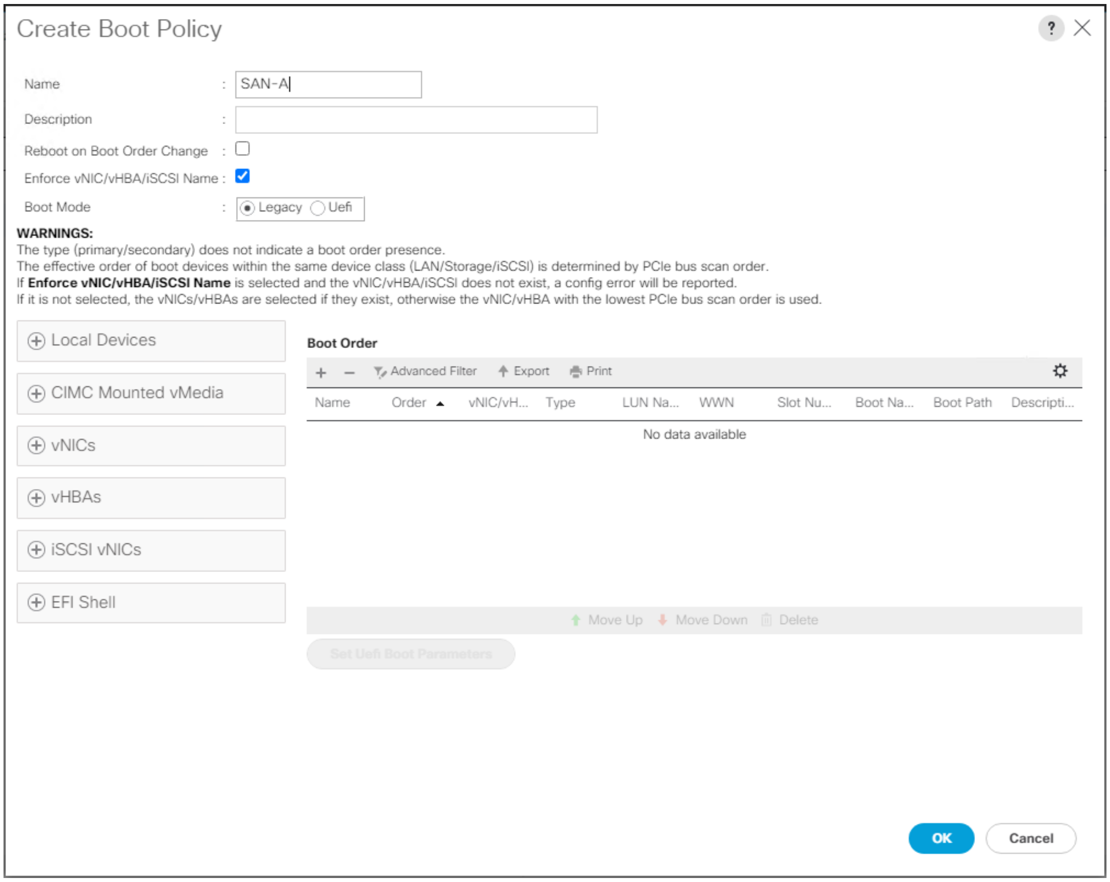

Create Boot Policy

Name : SAN-A
Description :
Reboot on Boot Order Change : □
Enforce vNIC/vHBA/iSCSI Name : ☑
Boot Mode : ● Legacy ○ Uefi

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

Name	Order ▲	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...

No data available

3. Expand the Local Devices drop-down list and Choose Add CD/DVD.

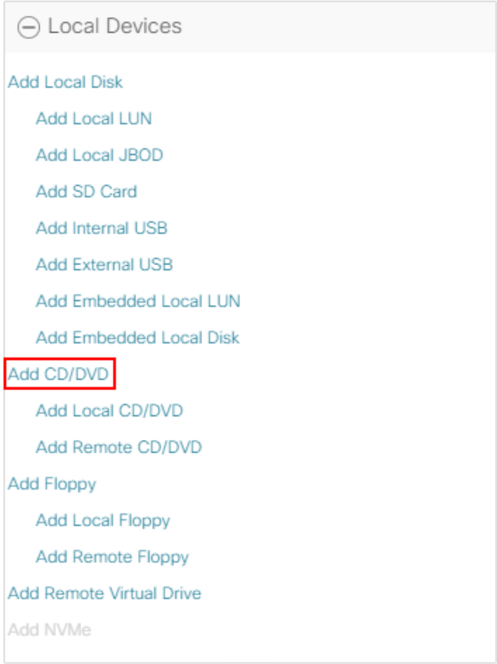

4. Expand the vHBAs drop-down list and Choose Add SAN Boot.

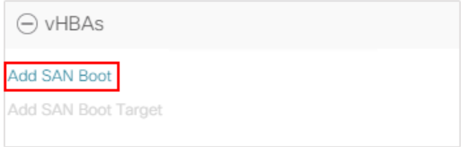

The SAN boot paths and targets will include primary and secondary options in order to maximize resiliency and number of paths.

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

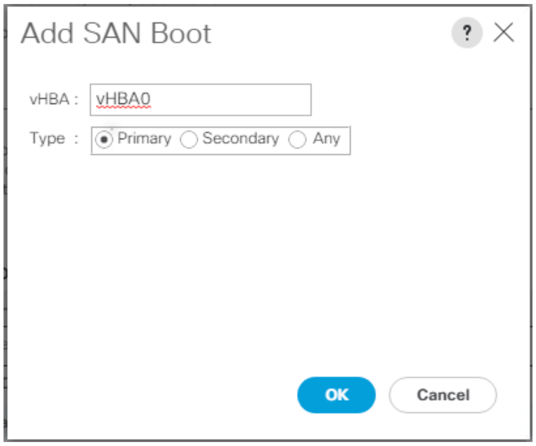

vHBA : vHBA0

Type : ● Primary ○ Secondary ○ Any

6. Select add SAN Boot Target.

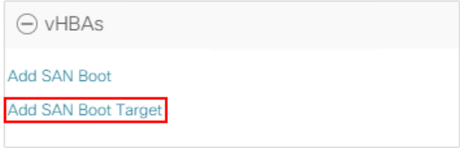

7. Keep **1** as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Primary Target.

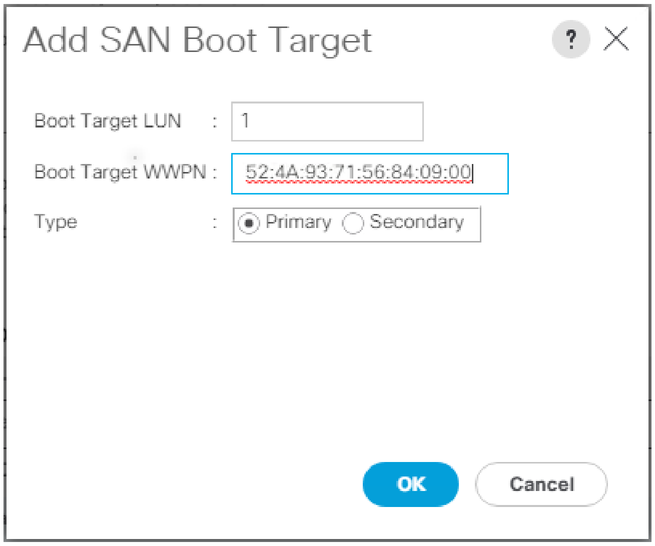

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:00

Type : ● Primary ○ Secondary

8. Add a secondary SAN Boot target into same hba0, enter the boot target LUN as **1** and WWPN for FC port CT1.FC0 of Pure Storage, and add SAN Boot Secondary Target.

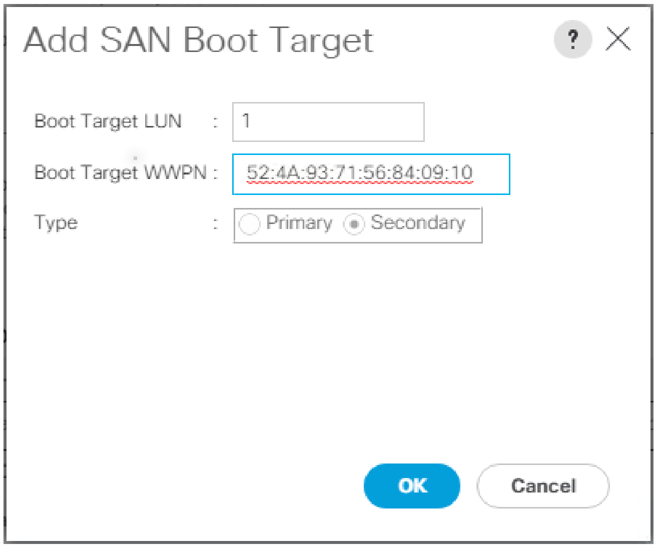

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:10

Type : ○ Primary ● Secondary

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

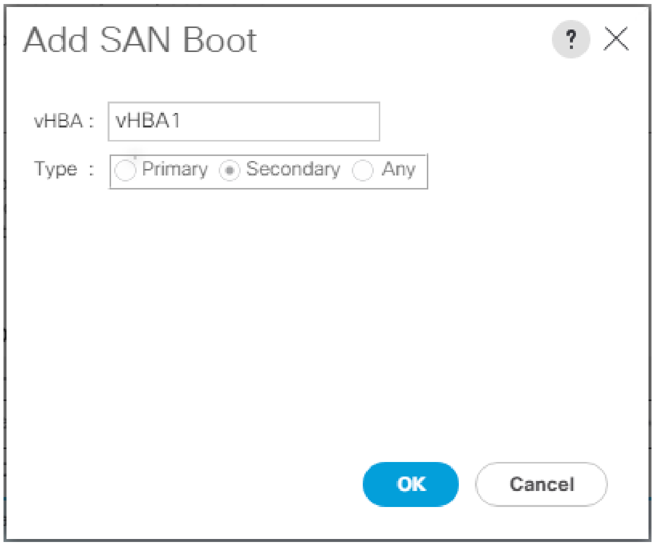

Add SAN Boot

vHBA : vHBA1

Type : ○ Primary ● Secondary ○ Any

10. Keep **1** as the value for the Boot Target LUN. Enter the WWPN for FC port CT1.FC2 of Pure Storage and add SAN Boot Primary Target.

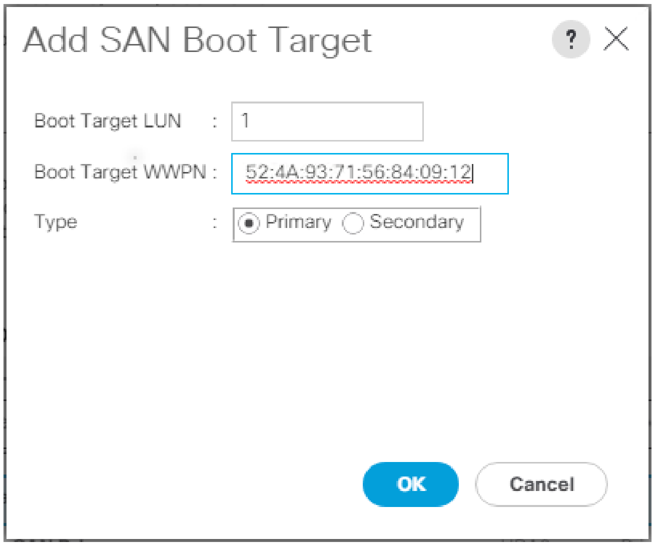

Add SAN Boot Target

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:12

Type : ● Primary ○ Secondary

11. Add a secondary SAN Boot target into same vhba1 and enter the boot target LUN as **1** and WWPN for FC port CT0.FC2 of Pure Storage and add SAN Boot Secondary Target.

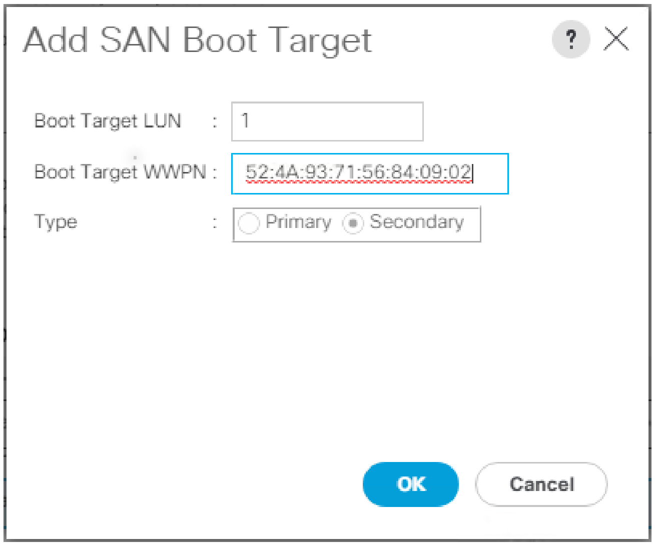

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:02

Type : ○ Primary ● Secondary

12. Click Save Changes.

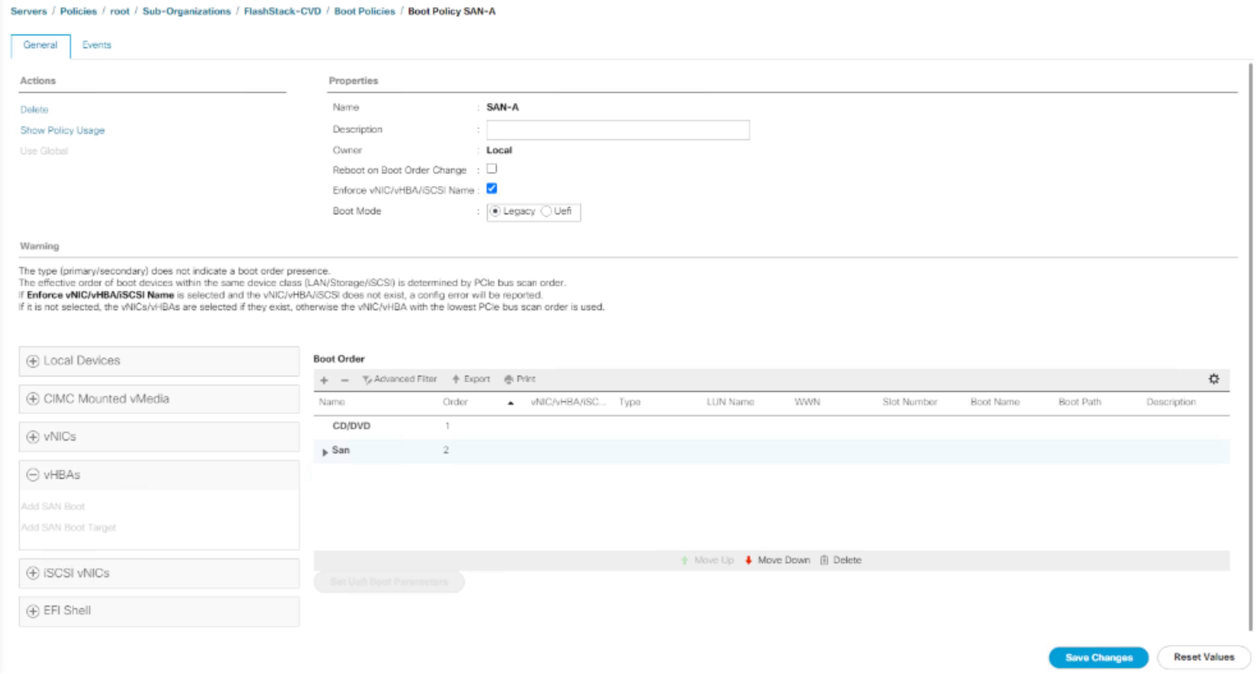

Name : SAN-A

Description :

Owner : Local

Reboot on Boot Order Change : ☐

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

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

Boot Order

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

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

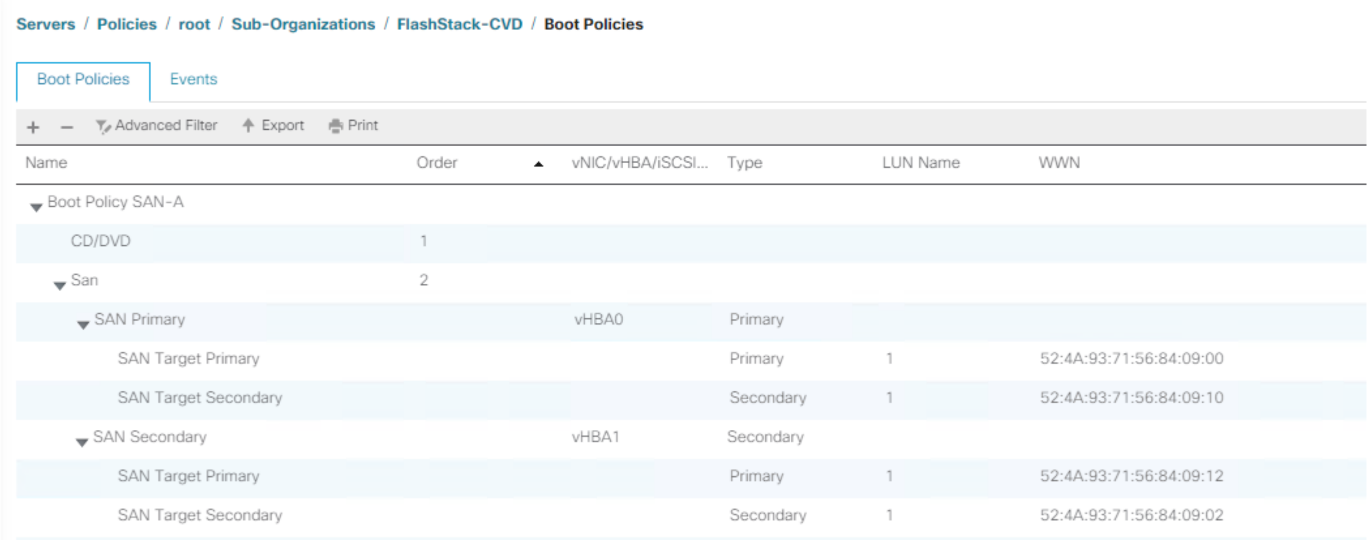
Servers / Policies / root / Sub-Organizations / FlashStack-CVD / Boot Policies
Boot Policies
Events
+ — Advanced Filter
Export
Print

Name	Order	vNIC/vHBA/iSCSI...	Type	LUN Name	WWN
Boot Policy SAN-A					
CD/DVD	1				
San	2				
SAN Primary		vHBA0	Primary		
SAN Target Primary			Primary	1	52:4A:93:71:56:84:09:00
SAN Target Secondary			Secondary	1	52:4A:93:71:56:84:09:10
SAN Secondary		vHBA1	Secondary		
SAN Target Primary			Primary	1	52:4A:93:71:56:84:09:12
SAN Target Secondary			Secondary	1	52:4A:93:71:56:84:09:02

Create SAN Policy B

The SAN-B boot policy configures the SAN Primary's primary-target to be port CT0.FC6 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC6 on the Pure Storage cluster. Similarly, the SAN Secondary’s primary-target should be port CT1.FC0 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC0 on the Pure Storage cluster.

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

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

To create boot policies for the Cisco UCS environments, follow these steps:

1. Go to Cisco UCS Manager and then go to tab Servers > Policies > root > Sub Organization > FlashStack-CVD > Boot Policies.

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

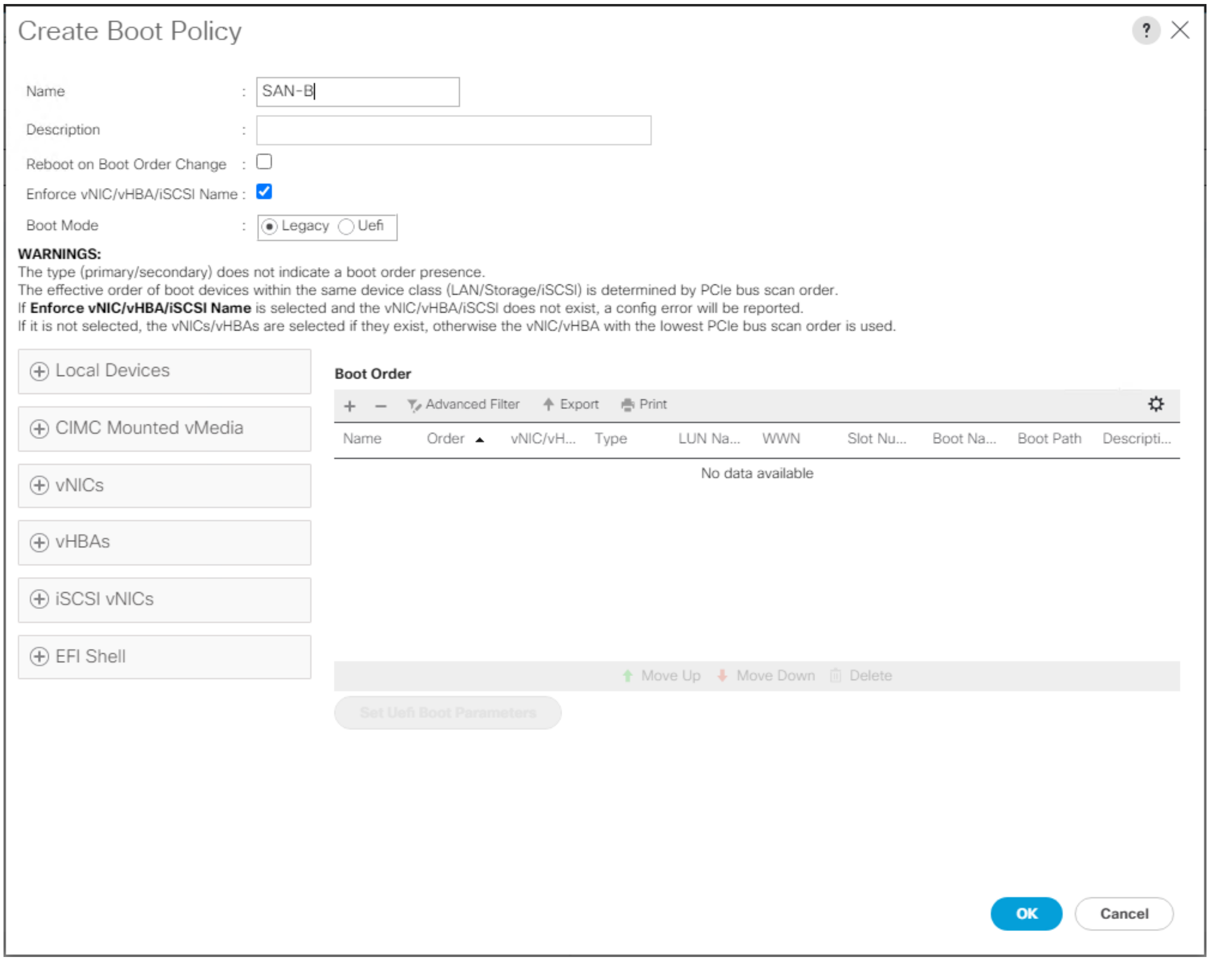

Create Boot Policy

Name : SAN-B

Description :

Reboot on Boot Order Change : □

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

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

Boot Order

Name	Order ▲	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...

No data available

3. Expand the Local Devices drop-down list and Choose Add CD/DVD. Expand the vHBAs drop-down list and choose Add SAN Boot.

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

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

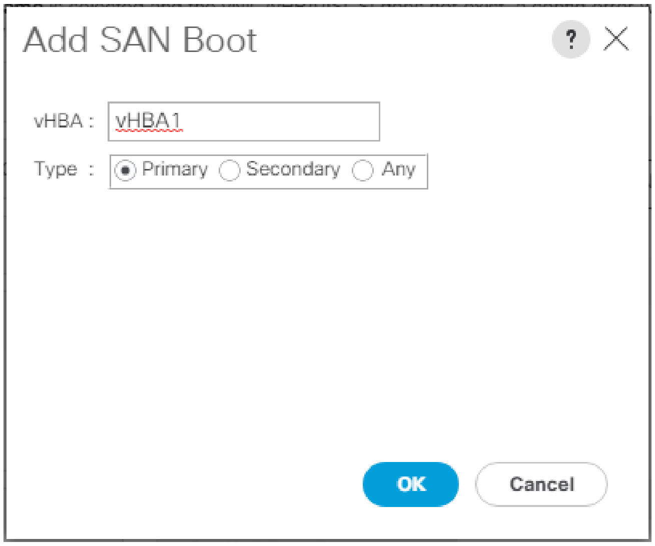

Add SAN Boot

vHBA : vHBA1

Type : ● Primary ○ Secondary ○ Any

5. Select Add SAN Boot Target to enter WWPN address of storage port. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC2 of Pure Storage and add SAN Boot Primary Target.

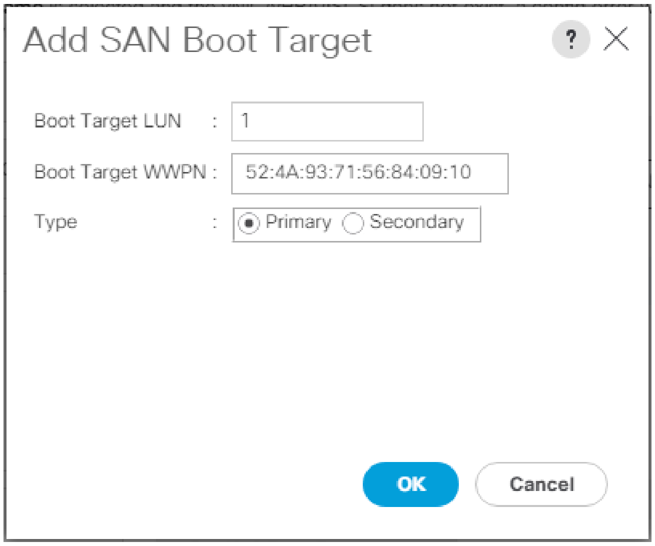

Add SAN Boot Target

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:10

Type : ● Primary ○ Secondary

6. Add the secondary SAN Boot target into the same hba0; enter boot target LUN as 1 and WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Secondary Target.

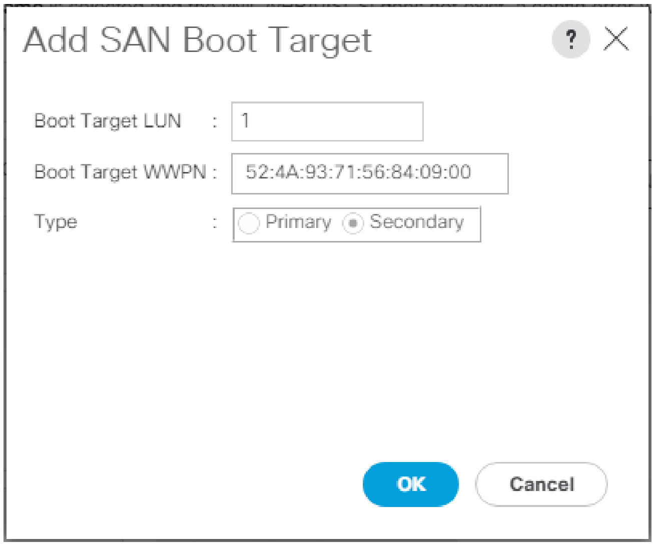

Add SAN Boot Target

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:00

Type : ○ Primary ● Secondary

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

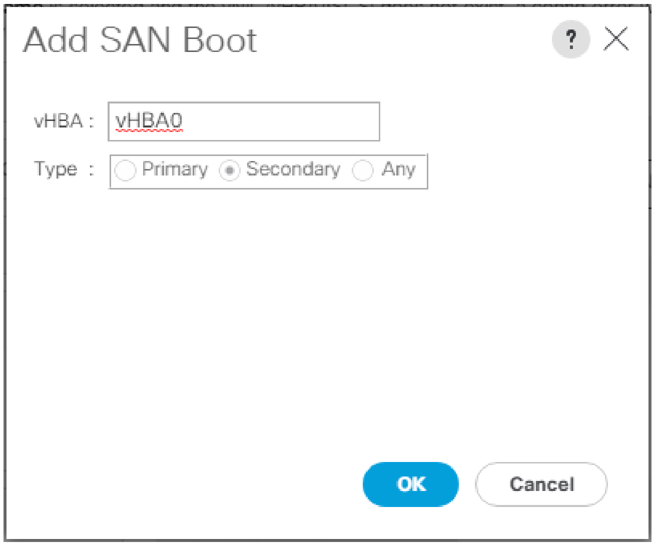

Add SAN Boot

vHBA : vHBA0

Type : ○ Primary ● Secondary ○ Any

8. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC1 of Pure Storage and Add SAN Boot Primary Target.

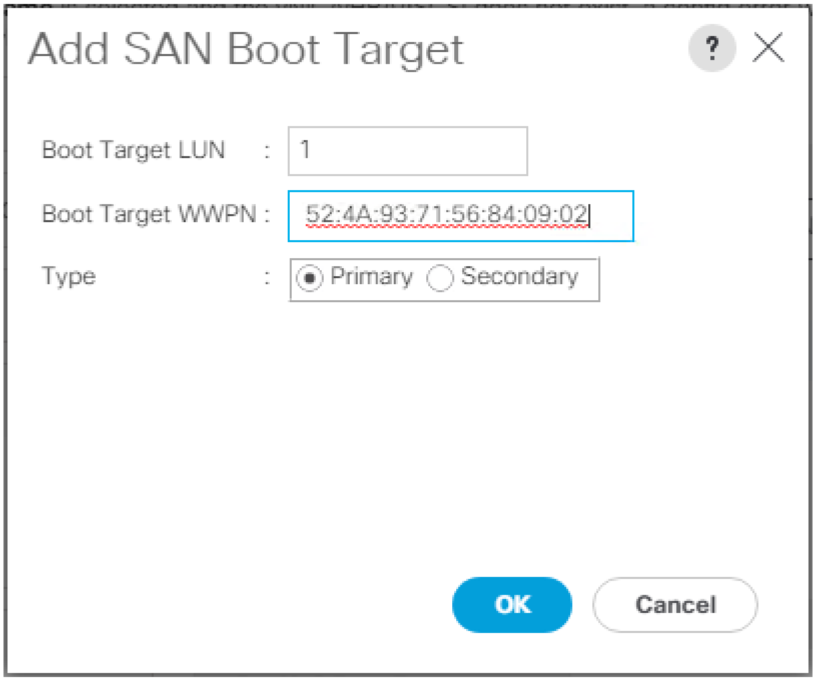

Add SAN Boot Target

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:02

Type : ● Primary ○ Secondary

9. Add secondary SAN Boot target into same hba1 and enter boot target LUN as 1 and WWPN for FC port CT1.FC1 of Pure Storage and add SAN Boot Secondary Target.

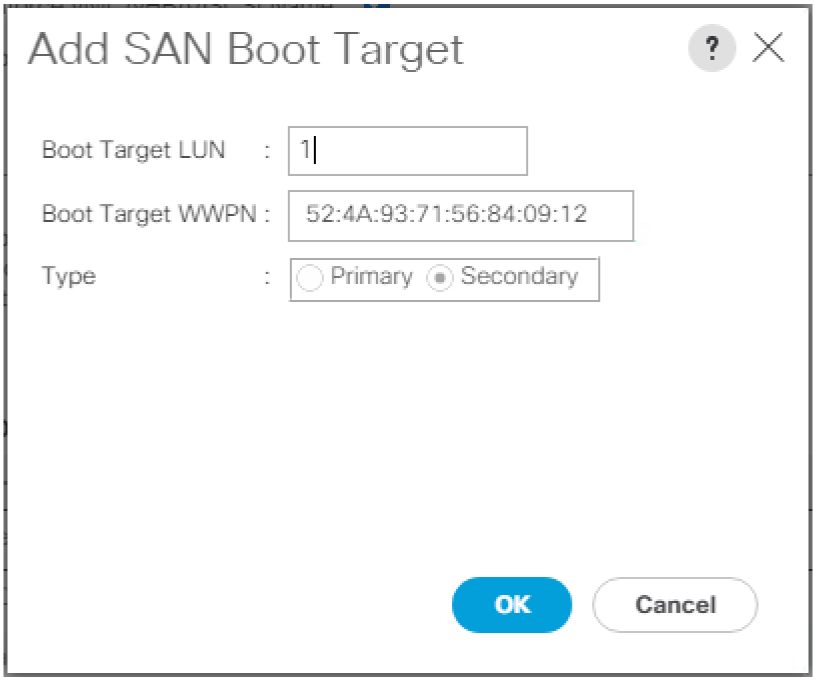

Add SAN Boot Target

Boot Target LUN : 1

Boot Target WWPN : 52:4A:93:71:56:84:09:12

Type : ○ Primary ● Secondary

10. Click OK.

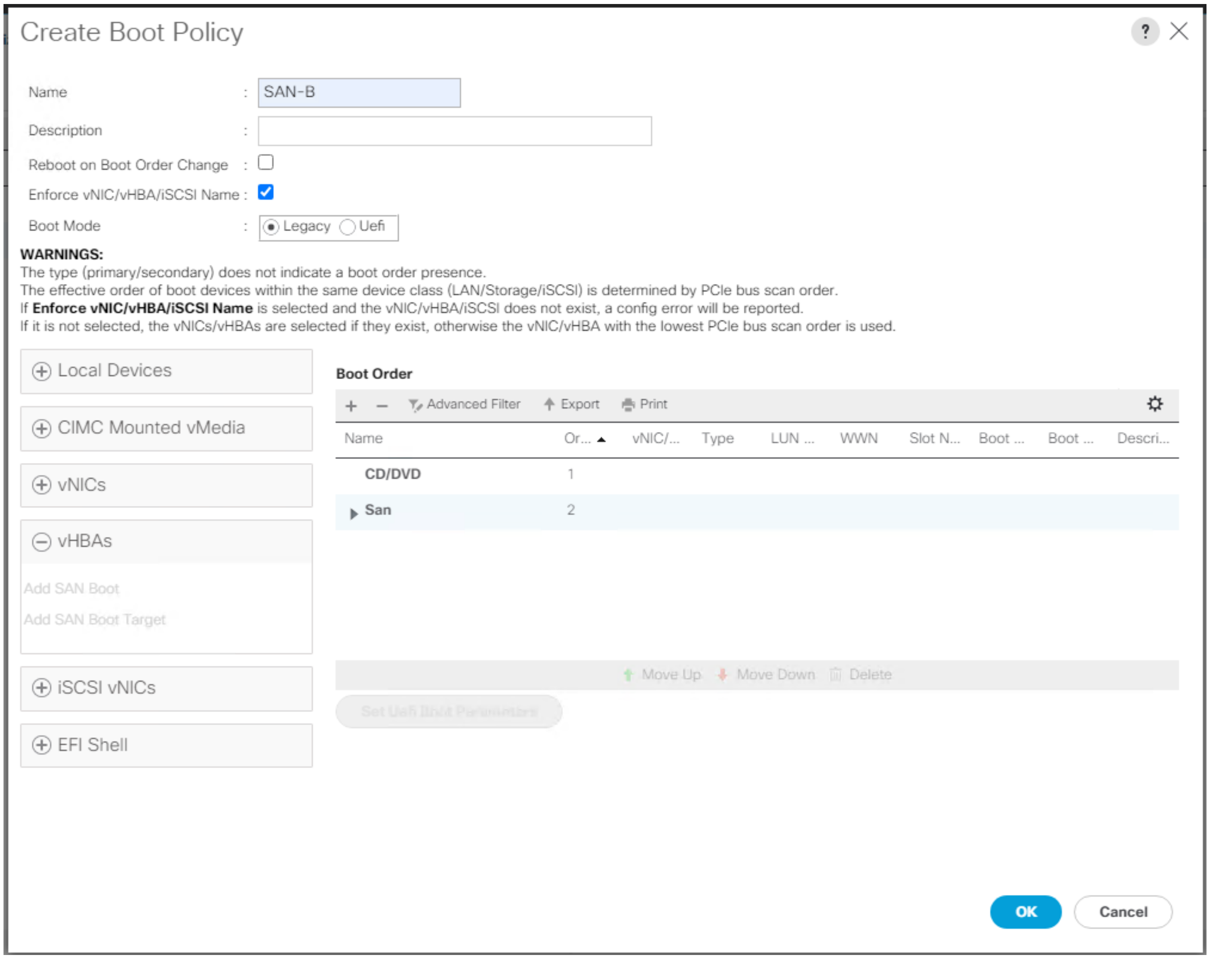

Create Boot Policy

Name : SAN-B

Description :

Reboot on Boot Order Change : □

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

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

Boot Order

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

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

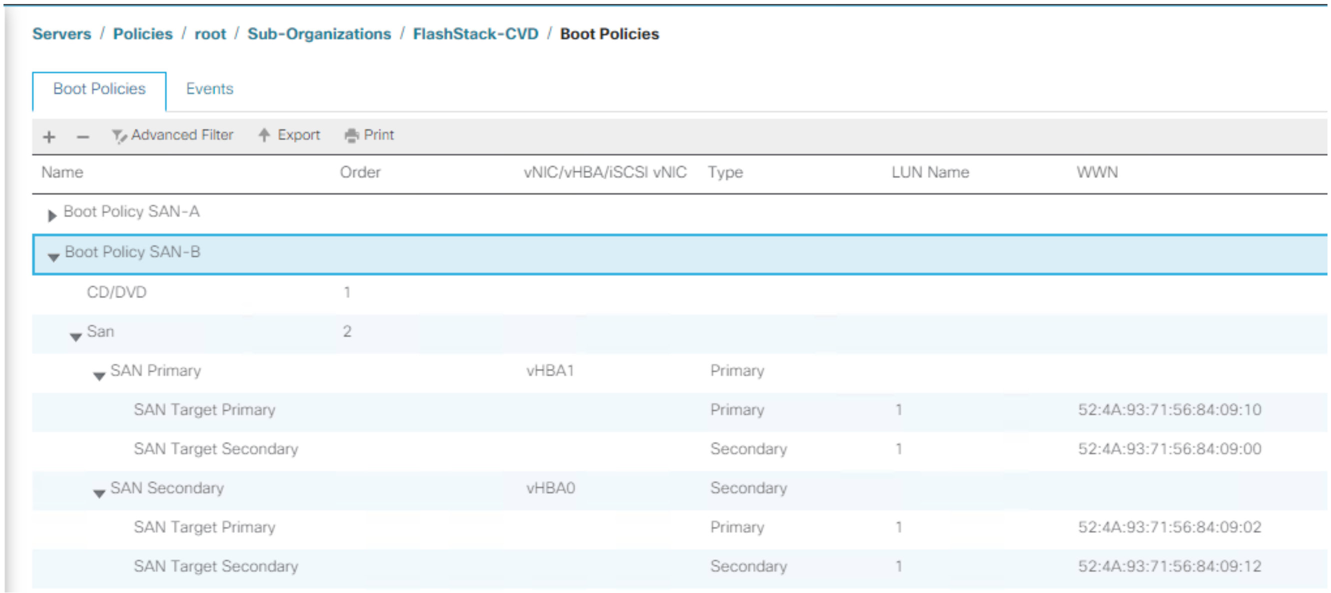
Servers / Policies / root / Sub-Organizations / FlashStack-CVD / Boot Policies
Boot Policies
Events
+ − Advanced Filter Export Print

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN
▸ Boot Policy SAN-A					
▾ Boot Policy SAN-B					
CD/DVD	1				
▾ San	2				
▾ SAN Primary		vHBA1	Primary		
SAN Target Primary			Primary	1	52:4A:93:71:56:84:09:10
SAN Target Secondary			Secondary	1	52:4A:93:71:56:84:09:00
▾ SAN Secondary		vHBA0	Secondary		
SAN Target Primary			Primary	1	52:4A:93:71:56:84:09:02
SAN Target Secondary			Secondary	1	52:4A:93:71:56:84:09:12

For this solution, we created two Boot Policy as “SAN-A” and “SAN-B”. For 32 Cisco UCS B200 M5 blade servers, you will assign the first 16 Service Profiles with SAN-A to the first 16 servers and the remaining 16 Service Profiles with SAN-B to the remaining 16 servers as explained in the following section.

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.

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

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

Create Service Profile Template

To create a service profile template, follow these steps:

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

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

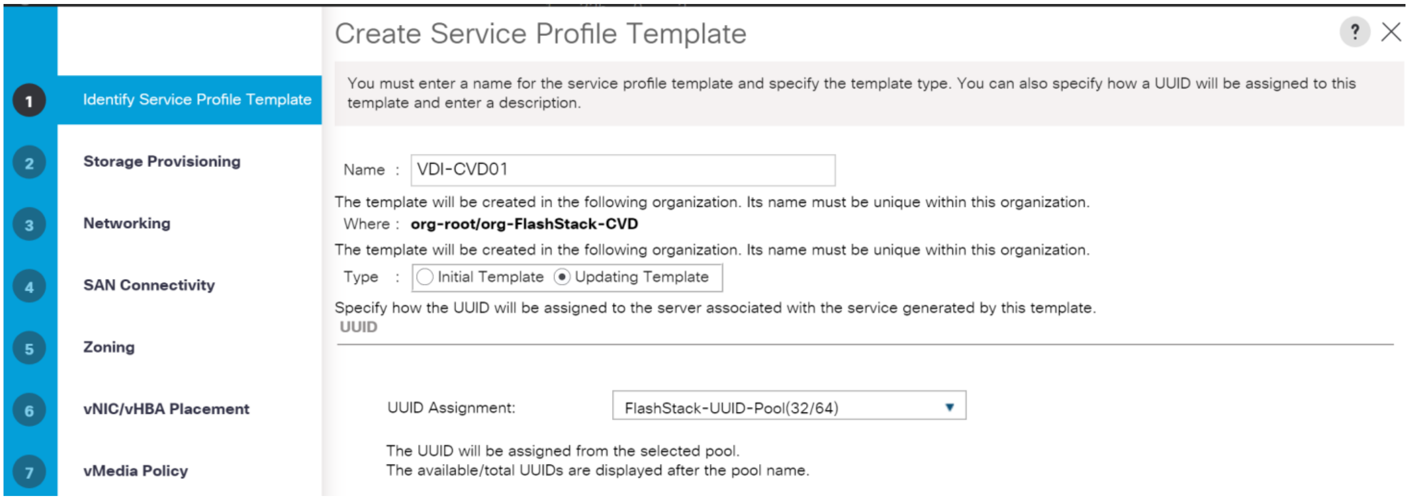

Create Service Profile Template

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 : VDI-CVD01

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-FlashStack-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: FlashStack-UUID-Pool(32/64)

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

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

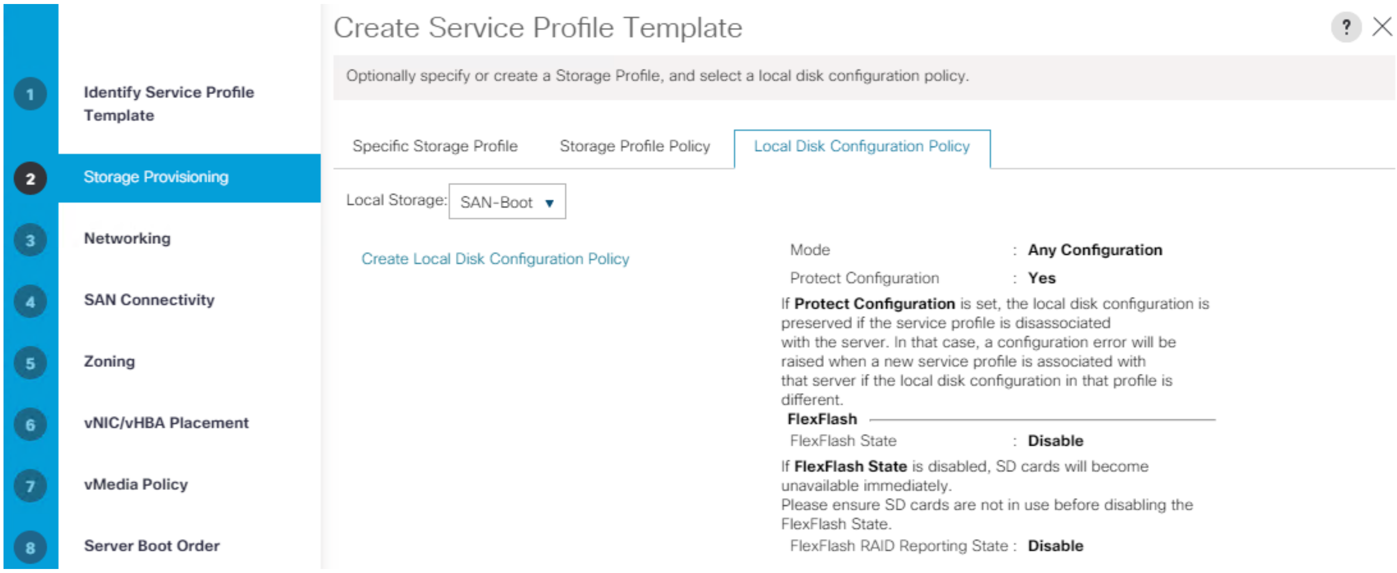

Create Service Profile Template

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

Local Storage: SAN-Boot

Mode : **Any Configuration**

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**

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.

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

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

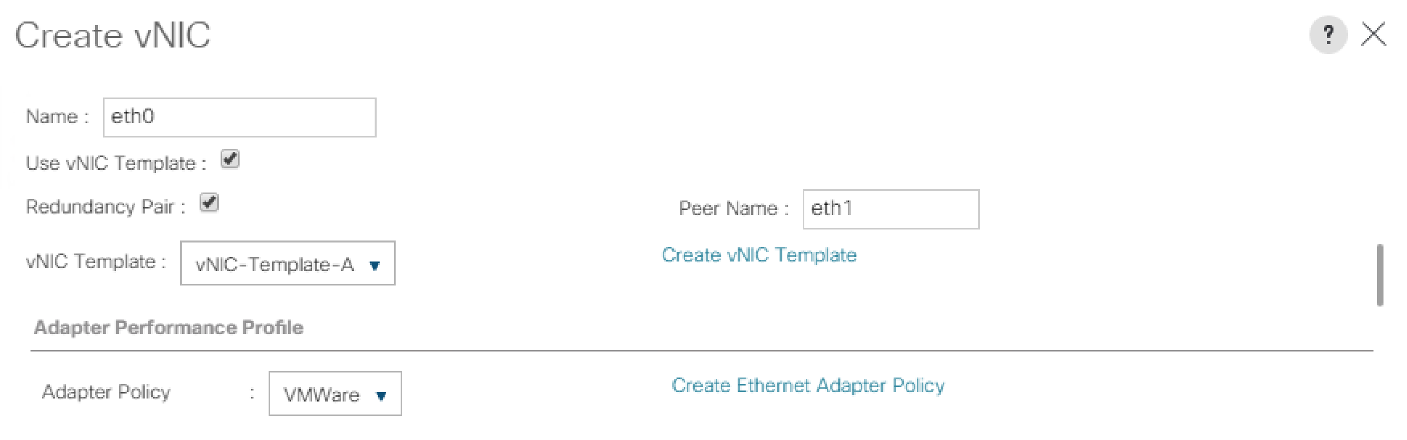

Create vNIC

Name : eth0

Use vNIC Template : ☑

Redundancy Pair : ☑

Peer Name : eth1

vNIC Template : vNIC-Template-A

Adapter Performance Profile

Adapter Policy : VMWare

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

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

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

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

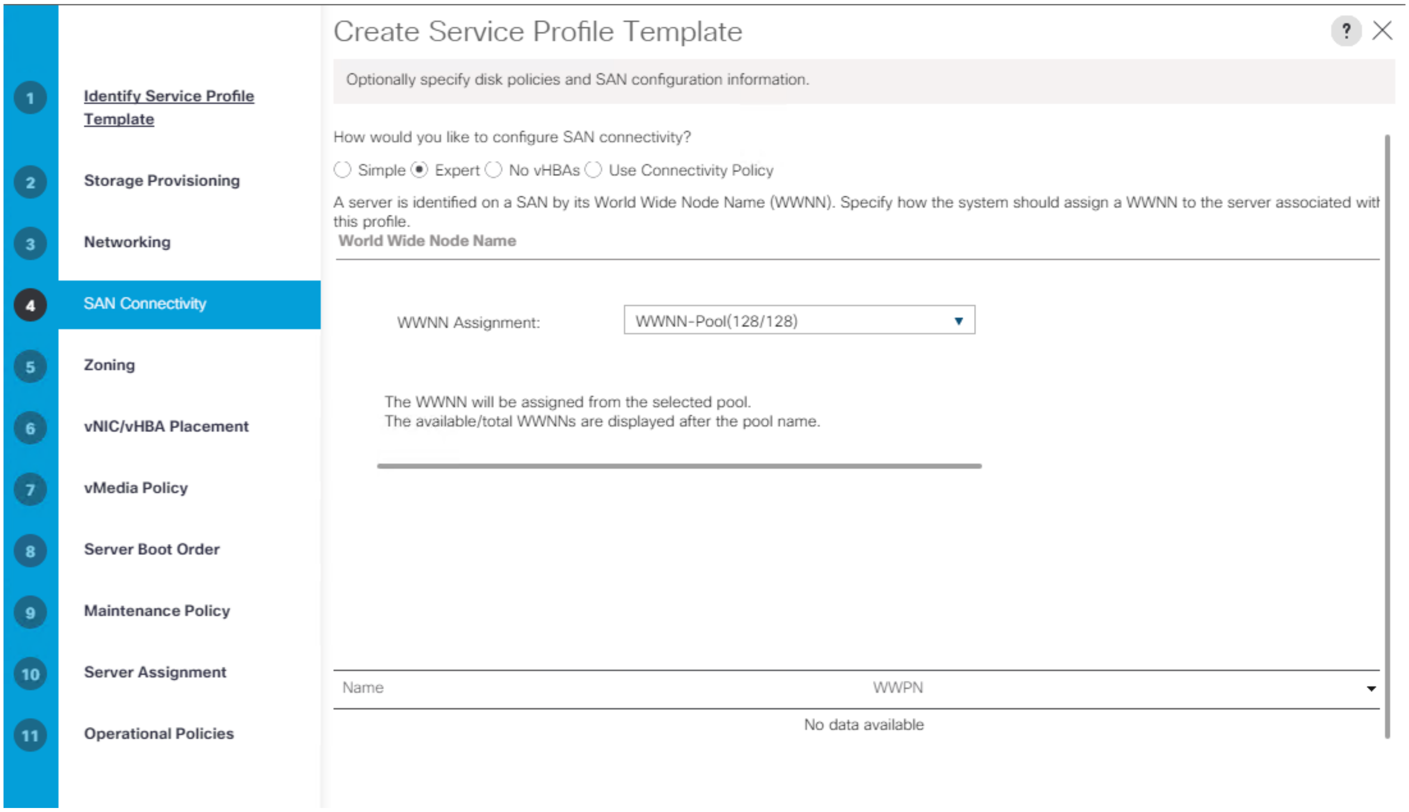

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

○ Simple ● Expert ○ No vHBAs ○ Use Connectivity Policy

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

World Wide Node Name

WWNN Assignment: WWNN-Pool(128/128)

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

Name	WWPN
No data available	

The following four HBAs were created:

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

Figure 37. **vHBA0**

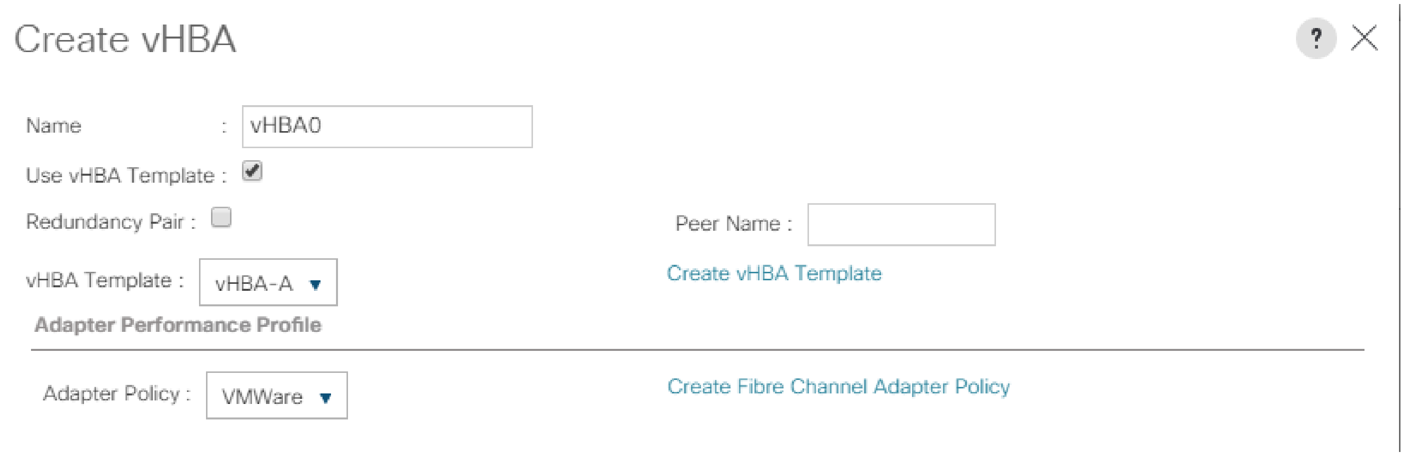

Create vHBA

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 38. **vHBA1**

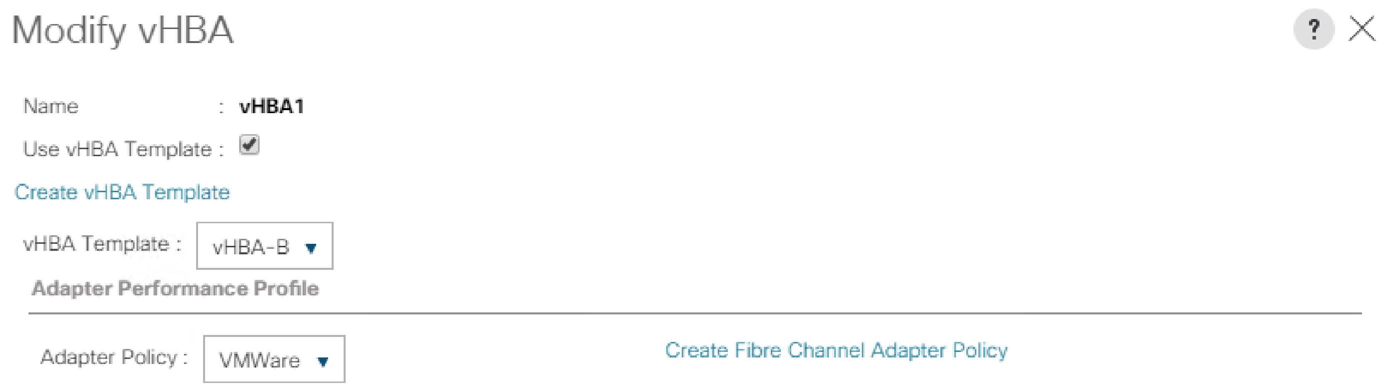

Modify vHBA

Name : **vHBA1**

Use vHBA Template : ☑

Create vHBA Template

vHBA Template : vHBA-B ▼

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Fibre Channel Adapter Policy

Figure 39. All vHBAs

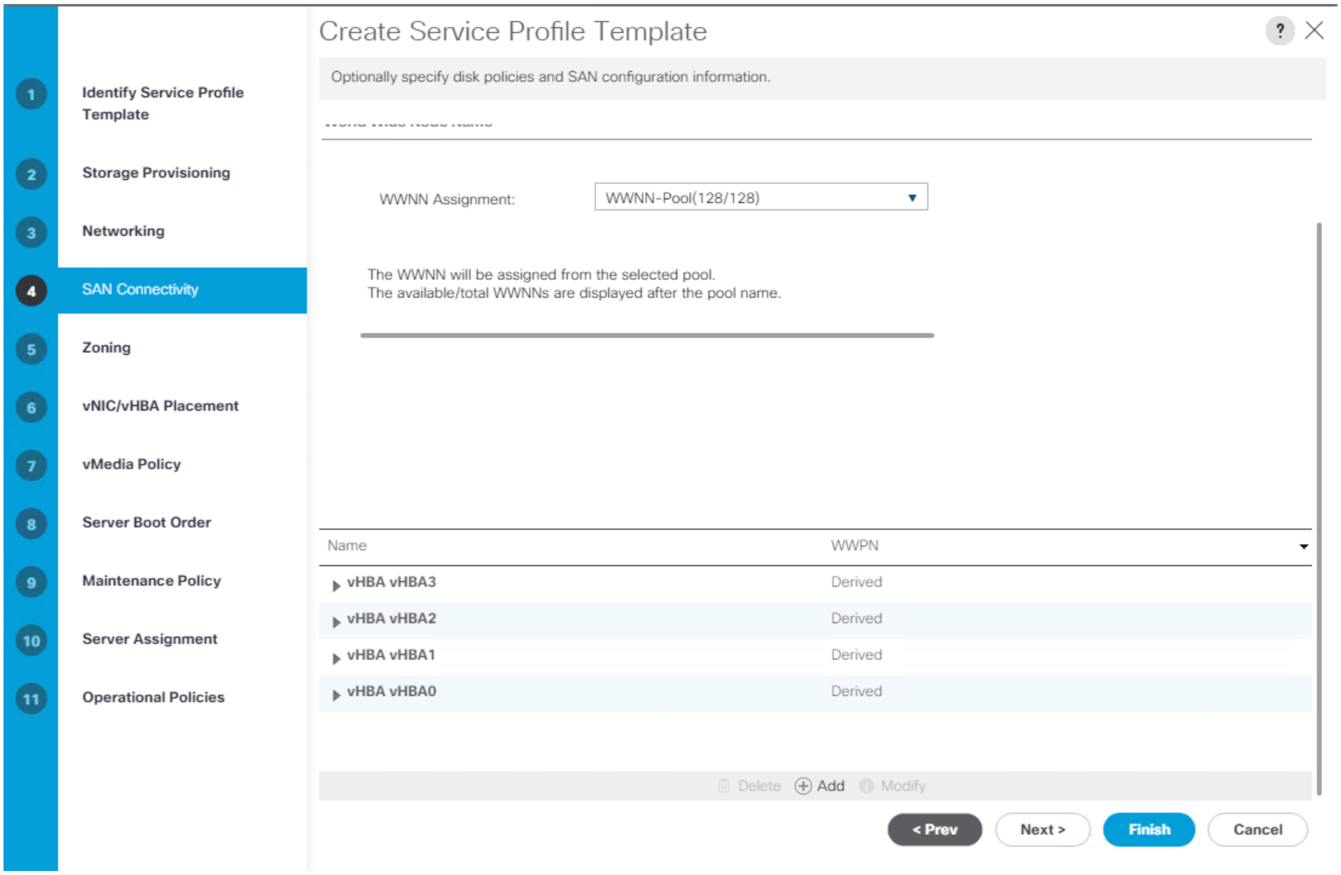

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

World Wide Node Name

WWNN Assignment: WWNN-Pool(128/128)

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

Name	WWPN
vHBA vHBA3	Derived
vHBA vHBA2	Derived
vHBA vHBA1	Derived
vHBA vHBA0	Derived

10. Skip zoning. For this FlashStack Configuration, the Cisco MDS 9132T 32-Gb is used for zoning.

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

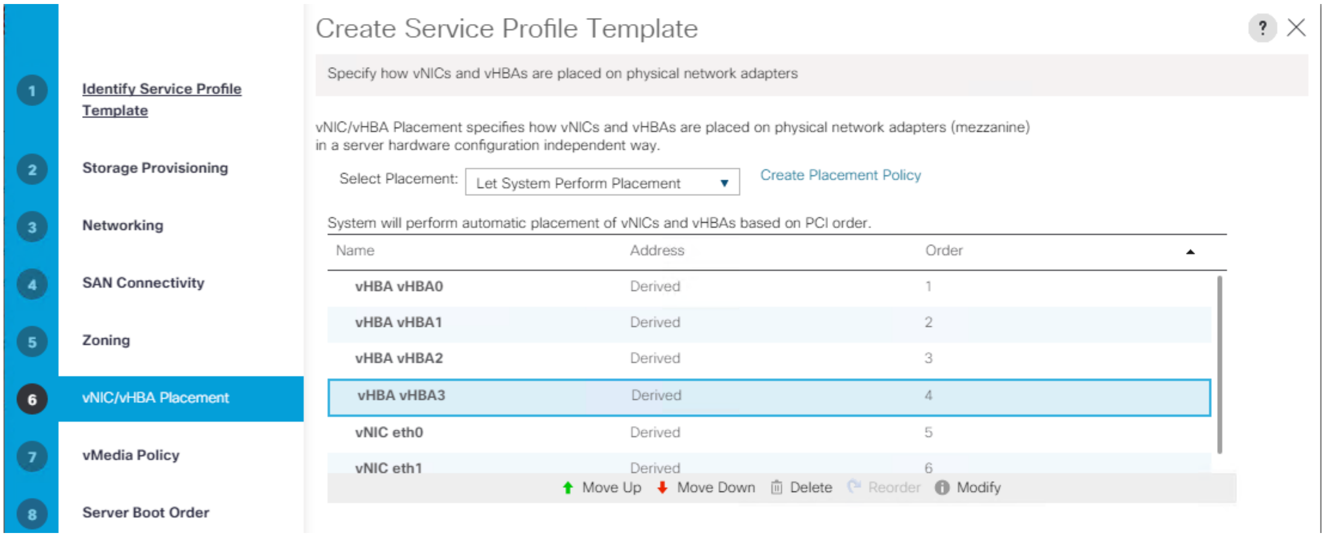

Create Service Profile Template

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: Let System Perform Placement

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

Name	Address	Order
vHBA vHBA0	Derived	1
vHBA vHBA1	Derived	2
vHBA vHBA2	Derived	3
vHBA vHBA3	Derived	4
vNIC eth0	Derived	5
vNIC eth1	Derived	6

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

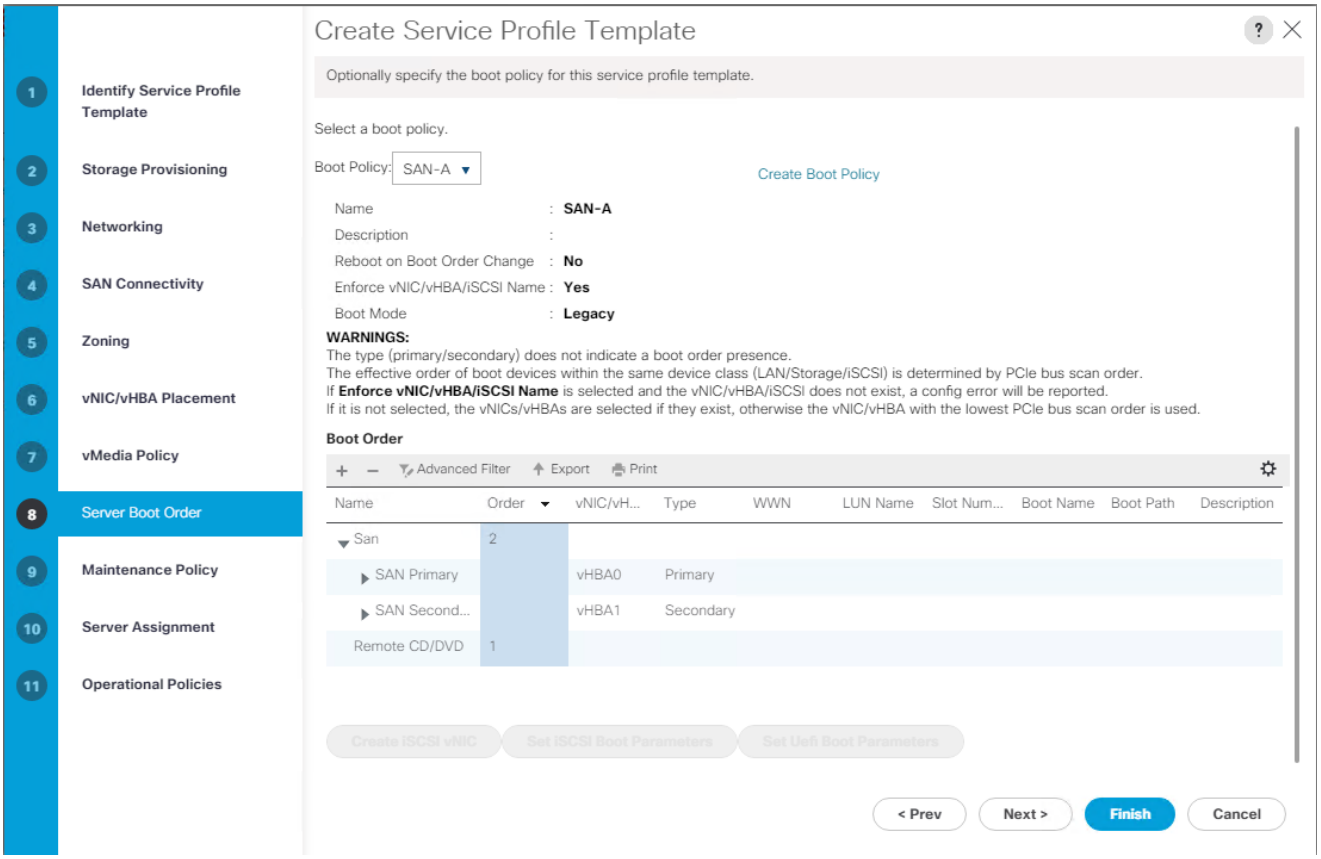

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: SAN-A ▾

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

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

Boot Order

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

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.

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

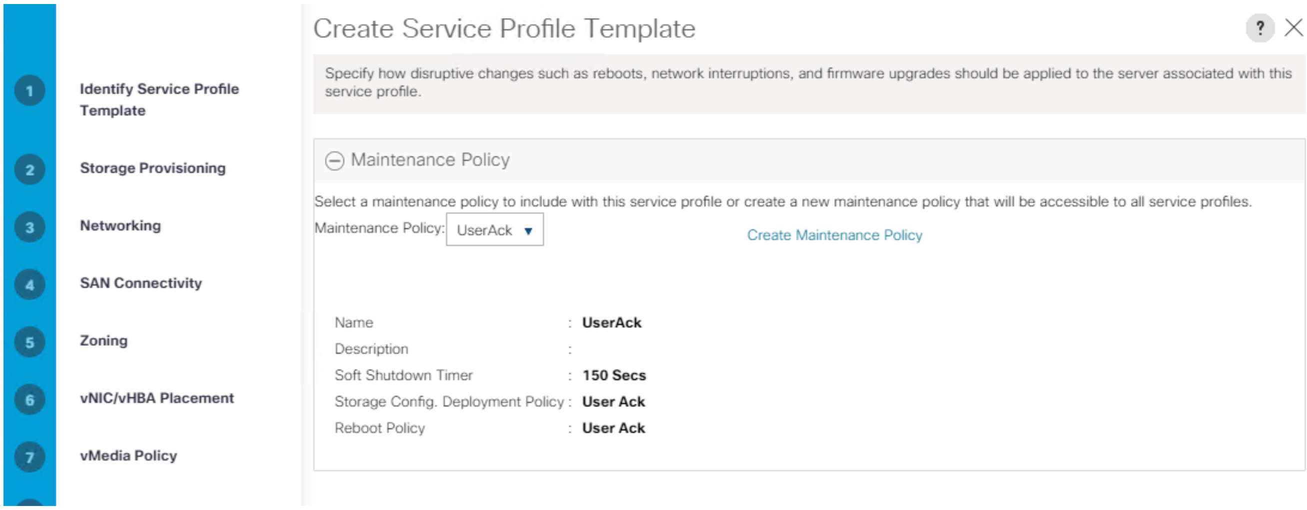

Create Service Profile Template

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: UserAck ▾

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

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

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

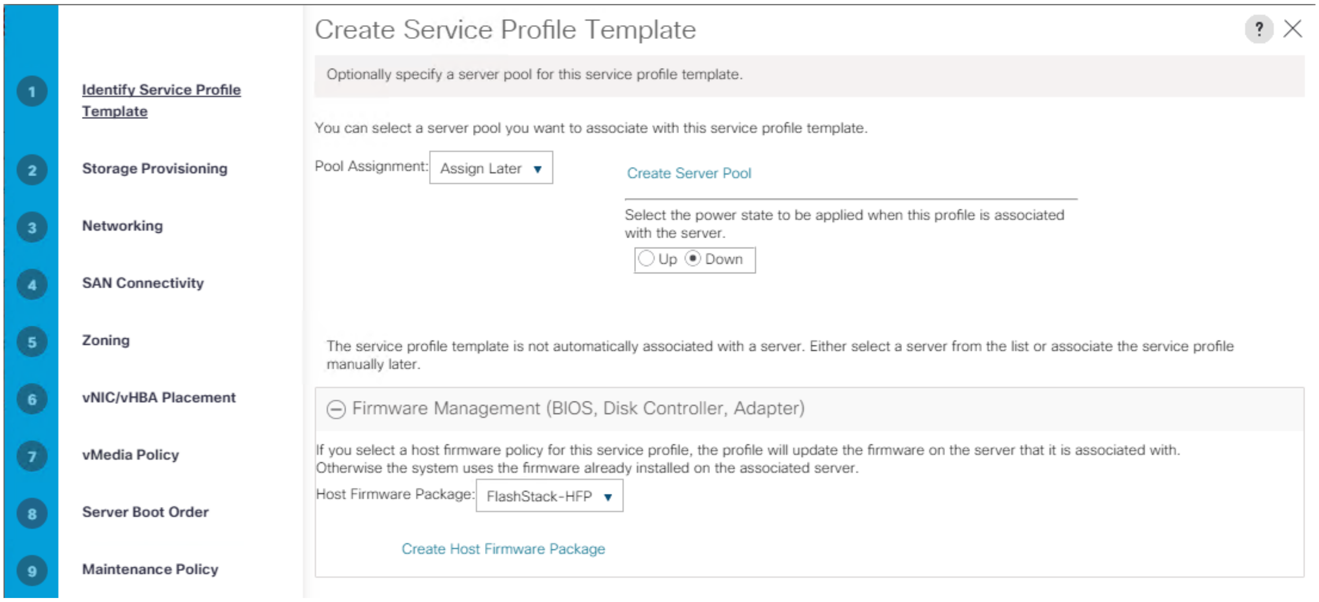

Create Service Profile Template

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: FlashStack-HFP ▾

Create Host Firmware Package

16. On the Operational Policy page, we configured BIOS policy for B200 M5 blade server, Power Control Policy with “NoPowerCap” for maximum performance and Graphics Card Policy for B200 M5 server configured with NVidia P6 GPU card.

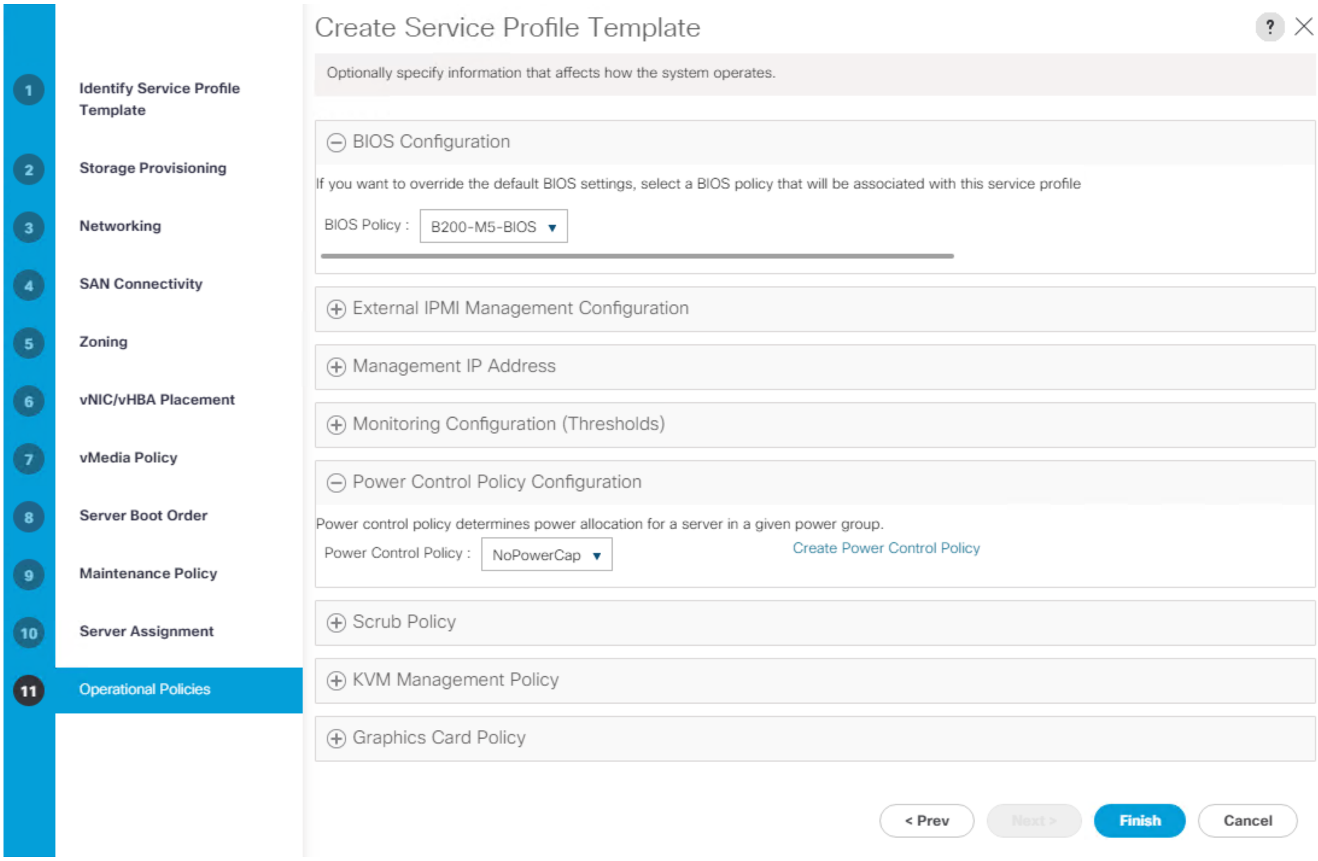

Create Service Profile Template

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy : B200-M5-BIOS ▾

⊕ External IPMI Management Configuration

⊕ Management IP Address

⊕ Monitoring Configuration (Thresholds)

⊖ Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : NoPowerCap ▾

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

Clone Service Profile Template

To clone the Service Profile template, follow these steps:

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

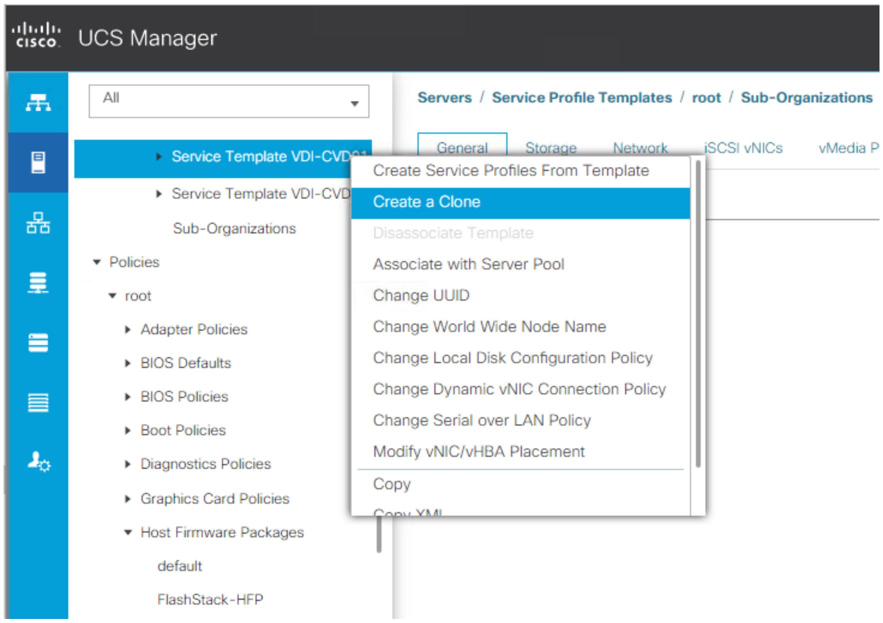
UCS Manager
All
Servers / Service Profile Templates / root / Sub-Organizations
General
Storage
Network
iSCSI vNICs
vMedia P
Service Template VDI-CVD01
Service Template VDI-CVD
Sub-Organizations
Policies
root
Adapter Policies
BIOS Defaults
BIOS Policies
Boot Policies
Diagnostics Policies
Graphics Card Policies
Host Firmware Packages
default
FlashStack-HFP
Create Service Profiles From Template
Create a Clone
Disassociate Template
Associate with Server Pool
Change UUID
Change World Wide Node Name
Change Local Disk Configuration Policy
Change Dynamic vNIC Connection Policy
Change Serial over LAN Policy
Modify vNIC/vHBA Placement
Copy
Copy XML

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

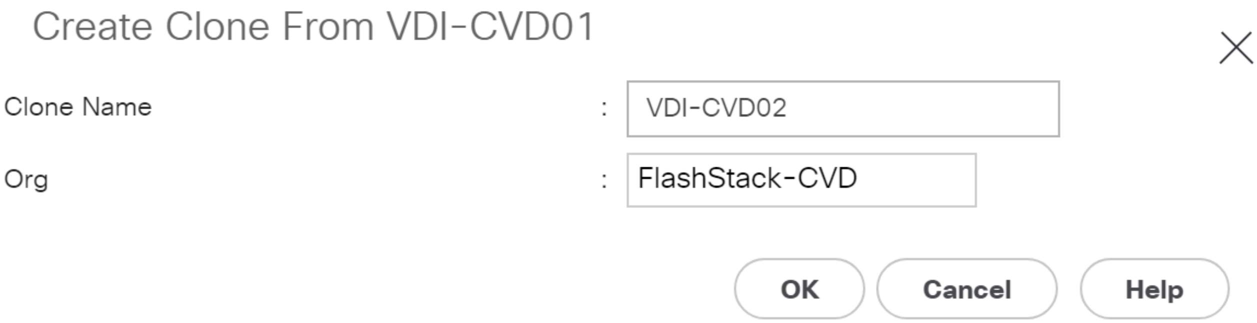
Create Clone From VDI-CVD01

Clone Name : VDI-CVD02

Org : FlashStack-CVD

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

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

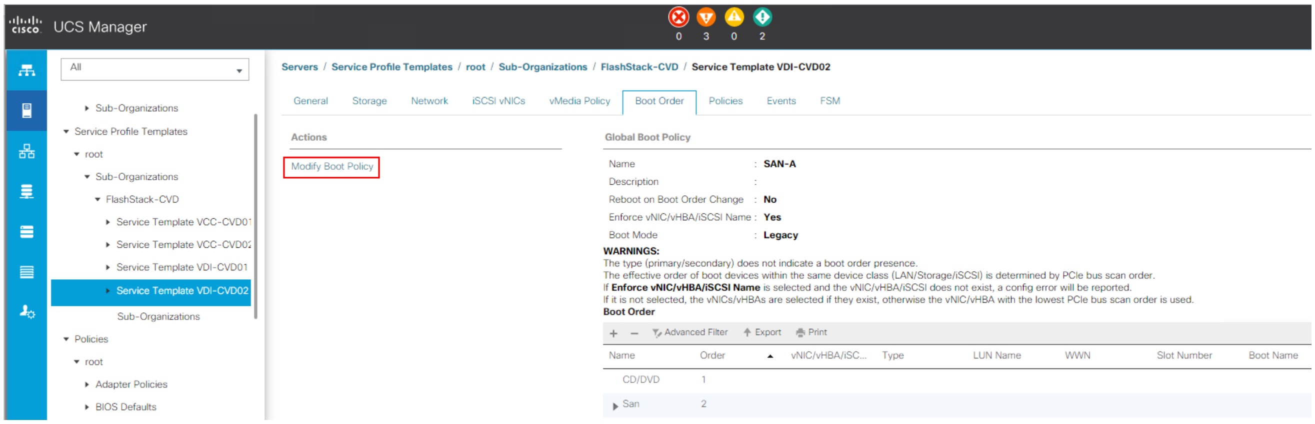

Servers / Service Profile Templates / root / Sub-Organizations / FlashStack-CVD / Service Template VDI-CVD02

Actions

Modify Boot Policy

Global Boot Policy

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

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

Boot Order

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

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

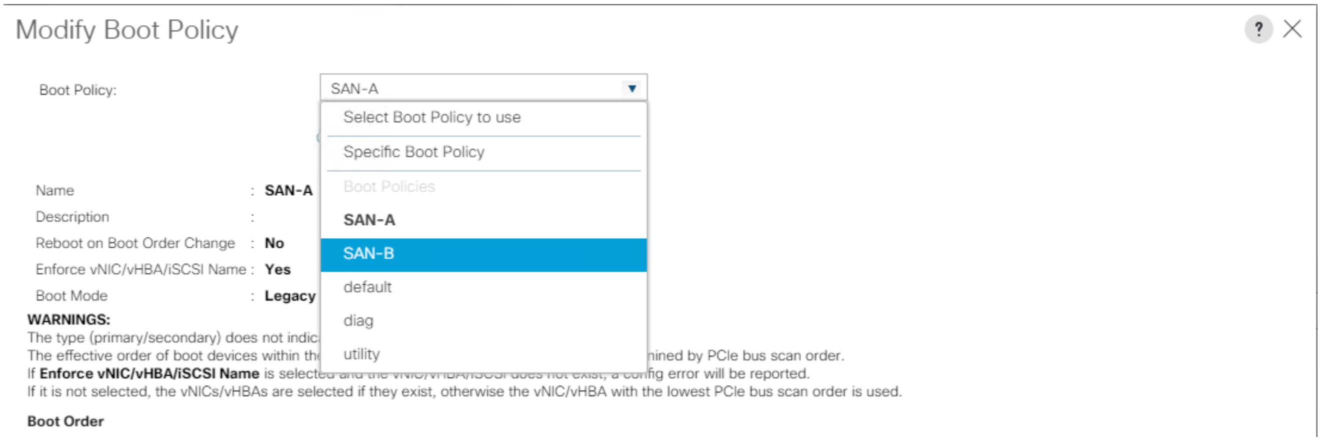

Boot Policy: SAN-A

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

WARNINGS:
The type (primary/secondary) does not indic
The effective order of boot devices within the ined 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

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

Create Service Profiles from Template and Associate to Servers

Create Service Profiles from Template

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

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

To create first four Service Profiles from Template, follow these steps:

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

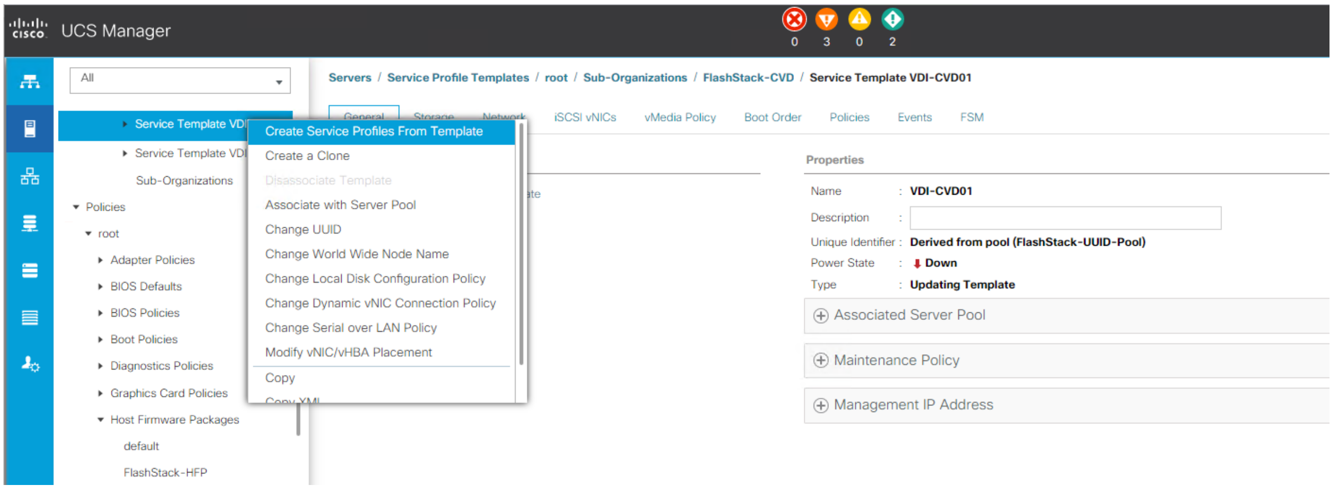

Properties

Name : VDI-CVD01

Description :

Unique Identifier : Derived from pool (FlashStack-UUID-Pool)

Power State : Down

Type : Updating Template

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

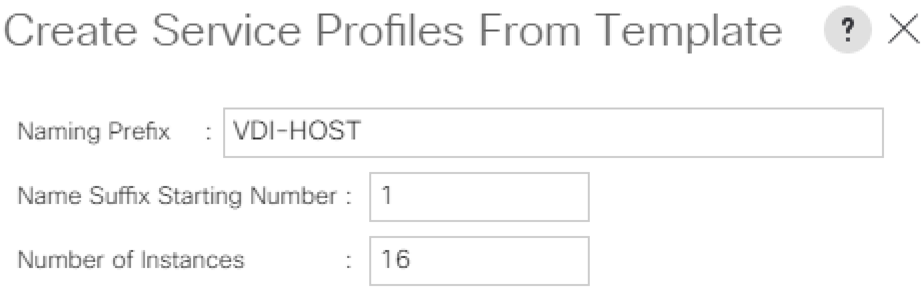
Create Service Profiles From Template

Naming Prefix : VDI-HOST

Name Suffix Starting Number : 1

Number of Instances : 16

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

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

4. Rename the Service Profiles on Chassis 3/8 as VDI-Infra01 and Service Profile on Chassis 4/8 as VDI-Infra02. Rename rest as necessary to have VDI-Host1to VDI-Host30.

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

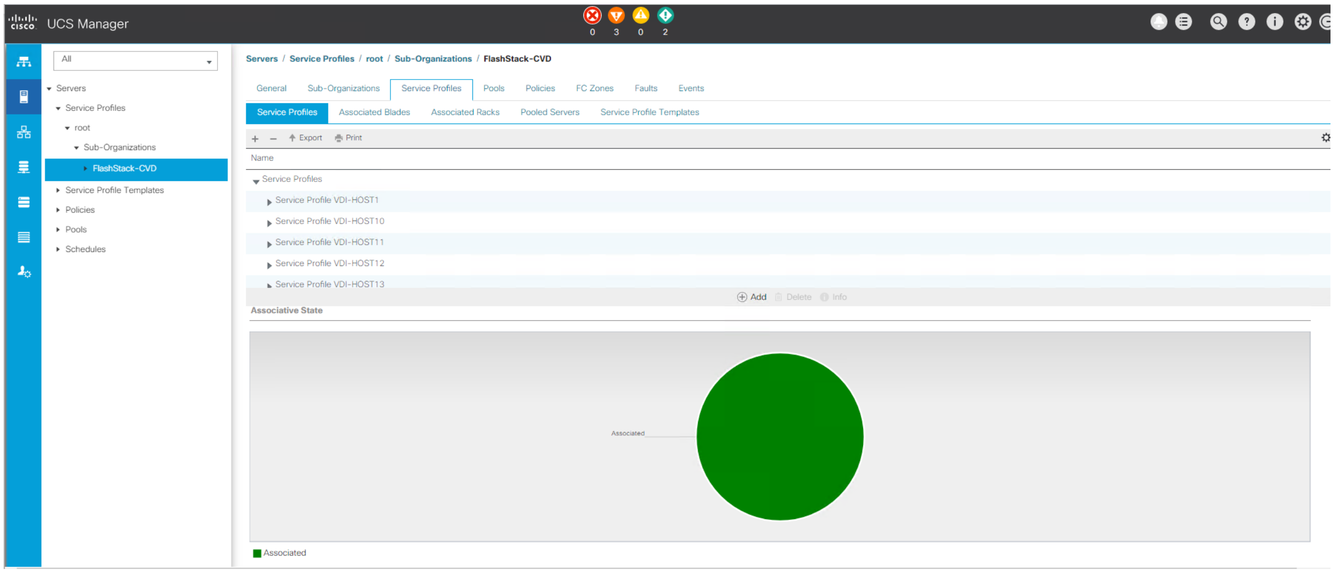

Configure Cisco Nexus 93180YC-FX Switches

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

Configure Global Settings for Cisco Nexus A and Cisco Nexus B

To set global configuration, follow these steps on both Nexus switches:

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 running-config startup-config
```

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

Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

To create the necessary virtual local area networks (VLANs), follow these steps on both Nexus switches. We created VLAN 70, 71, 72, 73 and 76.

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

2. Create VLAN 70:

```
config terminal
VLAN 70
name InBand-Mgmt
no shutdown
exit
copy running-config startup-config
```

3. Log in as admin user into the Nexus Switch B and create 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 HA, 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 is listed in Table 6.

Table 6. vPC Summary

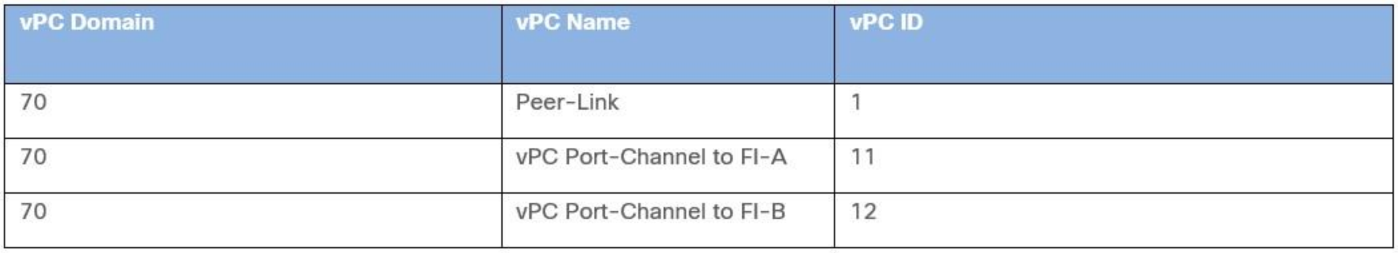

vPC Domain	vPC Name	vPC ID
70	Peer-Link	1
70	vPC Port-Channel to FI-A	11
70	vPC Port-Channel to FI-B	12

As listed in Table 6, a single vPC domain with Domain ID 70 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 3 vPCs:

- vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B.
- vPC IDs 11 and 12 are defined for traffic from Cisco UCS fabric interconnects.

Cisco Nexus 93180YC-FX Switch Cabling Details

The following tables list the cabling information.

Table 7. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch A	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/49
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/49
	Eth1/53	40Gbe	Cisco Nexus 93180YC-FX B	Eth1/53
	Eth1/54	40Gbe	Cisco Nexus 93180YC-FX B	Eth1/54
	MGMT0	1Gbe	Gbe management switch	Any

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

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch B	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/50
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/50
	Eth1/53	40Gbe	Cisco Nexus 93180YC-FX A	Eth1/53
	Eth1/54	40Gbe	Cisco Nexus 93180YC-FX A	Eth1/54
	MGMT0	Gbe	Gbe management switch	Any

Cisco UCS Fabric Interconnect 6454 Cabling

The following tables list the FI 6454 cabling information.

Table 9. Cisco UCS Fabric Interconnect (FI) A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-A	FC 1/1	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/13
	FC 1/2	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/14
	Eth1/17-24	40Gbe	UCS 5108 Chassis IOM-A 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

Local Device	Local Port	Connection	Remote Device	Remote Port
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Table 10. 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	32Gb FC	Cisco MDS 9132T 32-Gb-B	FC 1/13
	FC 1/2	32Gb FC	Cisco MDS 9132T 32-Gb-B	FC 1/14
	Eth1/17-24	40Gbe	UCS 5108 Chassis IOM-B 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
	L2	1Gbe	Cisco UCS FI - B	L2

Create vPC Peer-Link Between the Two Nexus Switches

To create the vPC Peer-Link, follow these steps:

1. Log in as “admin” user into the Nexus Switch A.

For vPC 1 as Peer-link, we used interfaces 53-54 for Peer-Link. You may choose the appropriate number of ports for your needs.

To create the necessary port channels between devices, follow these steps on both Nexus switches:

```
config terminal
feature vpc
feature lacp
vpc domain 1
peer-keepalive destination 10.29.164.234 source 10.29.164.233
exit
interface port-channel 70
```

```
description VPC peer-link
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type network
vpc peer-link
exit
interface Ethernet1/53
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
channel-group 70 mode active
no shutdown
exit
interface Ethernet1/54
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
channel-group 70 mode active
no shutdown
exit
copy running-config startup-config
```

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

Make sure to change the 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 11 and 12 for data network between the Nexus switches and fabric interconnects.

To create the necessary port channels between devices, follow these steps on both Nexus switches:

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

```
config terminal

interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 11
```

```
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

2. Log in as admin user into the Nexus Switch B and complete the following for the second switch configuration:

```
config Terminal
interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 11
no shutdown
```

```
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Verify All vPC Status is Up on Both Cisco Nexus Switches

Figure 40 shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 40. vPC Description for Cisco Nexus Switch A and B

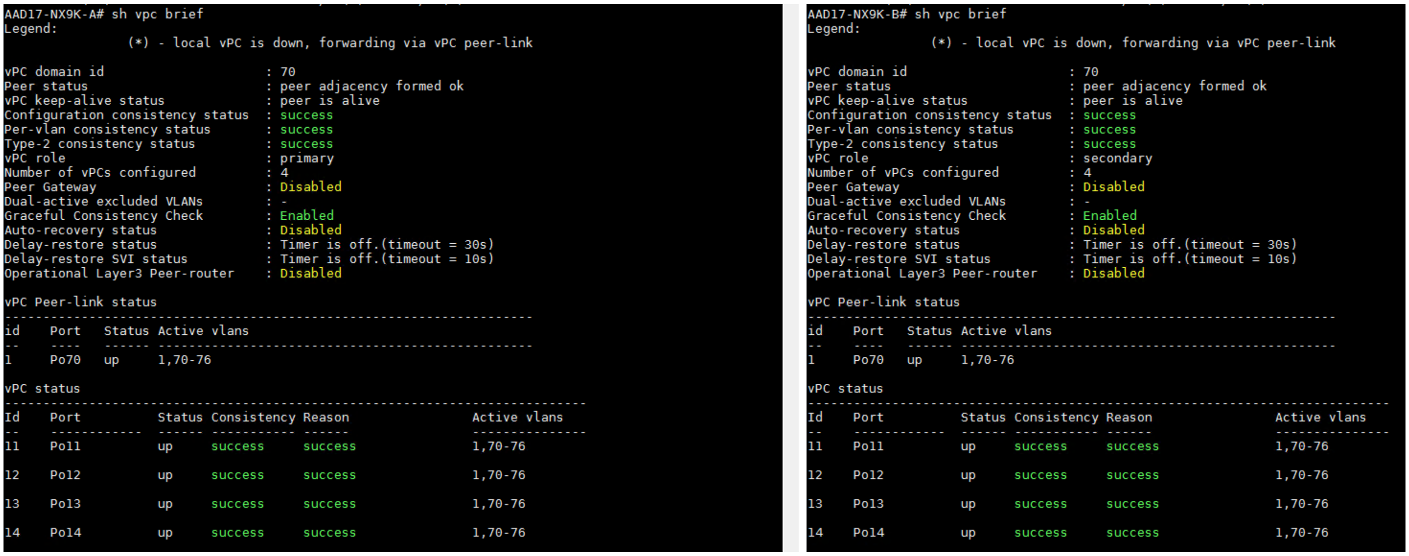

```
AAD17-NX9K-A# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 70
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Disabled
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po70   up     1,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,70-76

12    Po12          up     success     success               1,70-76

13    Po13          up     success     success               1,70-76

14    Po14          up     success     success               1,70-76
```

```
AAD17-NX9K-B# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 70
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Disabled
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po70   up     1,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,70-76

12    Po12          up     success     success               1,70-76

13    Po13          up     success     success               1,70-76

14    Po14          up     success     success               1,70-76
```

Cisco MDS 9132T 32-Gb FC Switch Configuration

Figure 21 illustrates the cable connectivity between the Cisco MDS 9132T 32-Gb switch and the Cisco 6454 Fabric Interconnects and Pure Storage FlashArray//X70 R3 storage.

We used two 32Gb FC connections from each fabric interconnect to each MDS switch and two 32Gb FC connections from each Pure Storage FlashArray//X70 R3 array controller to each MDS switch.

Table 11. Cisco MDS 9132T-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-A	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC0
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/1
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/2

Table 12. Cisco MDS 9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-B	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC2
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/2

Pure Storage FlashArray//X70 R3 to MDS SAN Fabric Connectivity

Pure Storage FlashArray//X70 R3 to MDS A and B Switches using VSAN 100 for Fabric A and VSAN 101 Configured for Fabric B

In this solution, two ports (ports FC1/9 and FC1/10) of MDS Switch A and two ports (ports FC1/9 abd FC1/10) of MDS Switch B are connected to Pure Storage System as listed in Table 13. All ports connected to the Pure Storage Array carry 32 Gb/s FC Traffic.

Table 13. MDS 9132T 32-Gb switch Port Connection to Pure Storage System

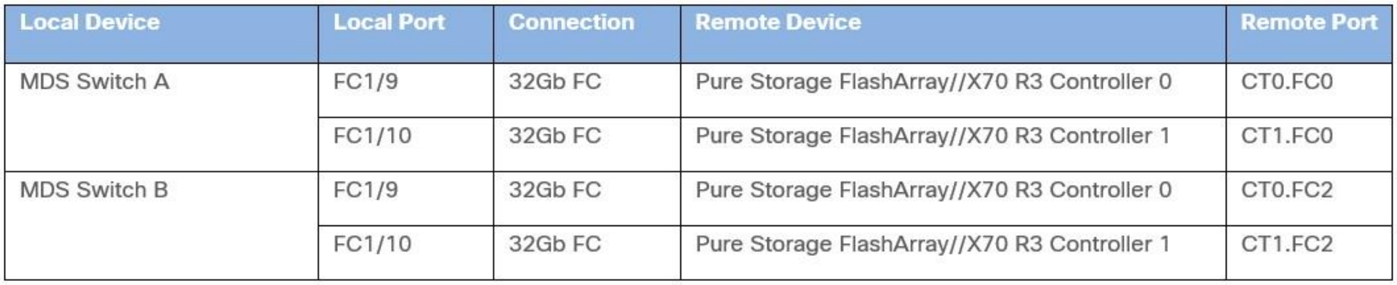

Local Device	Local Port	Connection	Remote Device	Remote Port
MDS Switch A	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC0
MDS Switch B	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC2

Configure Feature for MDS Switch A and MDS Switch B

To set feature on MDS Switches, follow these steps on both MDS switches:

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

```
config terminal
feature npiv
feature telnet
switchname FlashStack-MDS-A
copy running-config startup-config
```

2. Log in as admin user into MDS Switch B. Repeat the steps above on MDS Switch B.

Configure VSANs for MDS Switch A and MDS Switch B

To create VSANs, follow these steps:

1. Log in as admin user into MDS Switch A. Create VSAN 100 for Storage Traffic:

```
config terminal
VSAN database
vsan 100
vsan 100 interface fc 1/9-16
```

```
exit
interface fc 1/9-16
switchport trunk allowed vsan 100
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

2. Log in as admin user into MDS Switch B. Create VSAN 101 for Storage Traffic:

```
config terminal
VSAN database
vsan 101
vsan 101 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 101
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Add FC Uplink Ports to Corresponding VSAN on Fabric Interconnect

To add the FC Ports to the corresponding VSAN, follow these steps:

1. In Cisco UCS Manager, in the Equipment tab, select Fabric Interconnects > Fabric Interconnect A > Physical Ports > FC Ports.

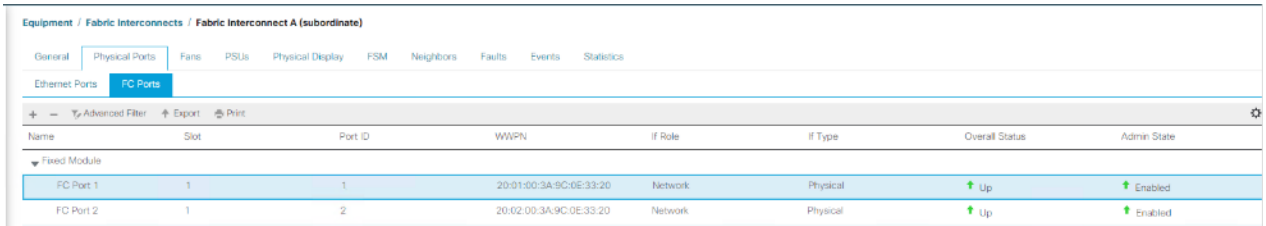

Name	Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
Fixed Module							
FC Port 1	1	1	20:01:00:3A:9C:0E:33:20	Network	Physical	↑ Up	↑ Enabled
FC Port 2	1	2	20:02:00:3A:9C:0E:33:20	Network	Physical	↑ Up	↑ Enabled

2. From the drop-down list double-click FC Port 1 and select VSAN 100.

Figure 41. VSAN Assignment on FC Uplink Ports to MDS Switch

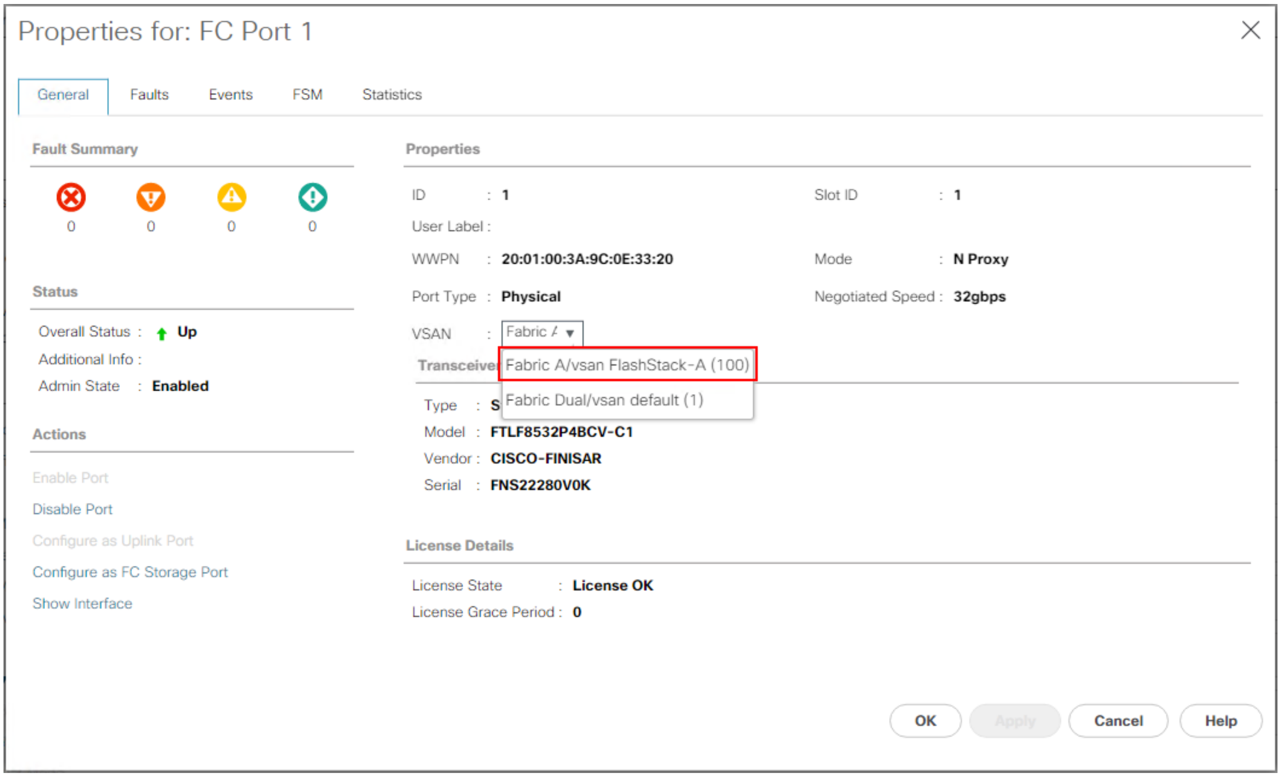

Status

Overall Status : Up
Additional Info :
Admin State : Enabled

Properties

ID : 1
Slot ID : 1
User Label :
WWPN : 20:01:00:3A:9C:0E:33:20
Mode : N Proxy
Port Type : Physical
Negotiated Speed : 32gbps
VSAN : Fabric A
Fabric A/vsan FlashStack-A (100)
Fabric Dual/vsan default (1)

Transceiver

Type : S
Model : FTLF8532P4BCV-C1
Vendor : CISCO-FINISAR
Serial : FNS22280V0K

License Details

License State : License OK
License Grace Period : 0

3. Repeat these steps to Add FC Port 1-4 to VSAN 100 on Fabric A and FC Port 1-4 to VSAN 101 on Fabric B.

Create and Configure Fiber Channel Zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9132T 32-Gb switches, the Cisco UCS Fabric Interconnects, and the Pure Storage FlashArray systems.

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 4 HBAs for each Server. Two HBAs (HBA0 and HBA2) are connected to MDS Switch-A and other two HBAs (HBA1 and HBA3) are connected to MDS Switch-B.

To create and configure the fiber channel zoning, follow these steps:

1. Log into the Cisco UCS Manager and go to Servers > Service Profiles > Sub-Organizations > FlashStack-CVD > VDI-HostX, then click the Storage tab and HBA's tab to get the WWPN of HBA's as shown in the screenshot below. Repeat for all the configured host profiles.

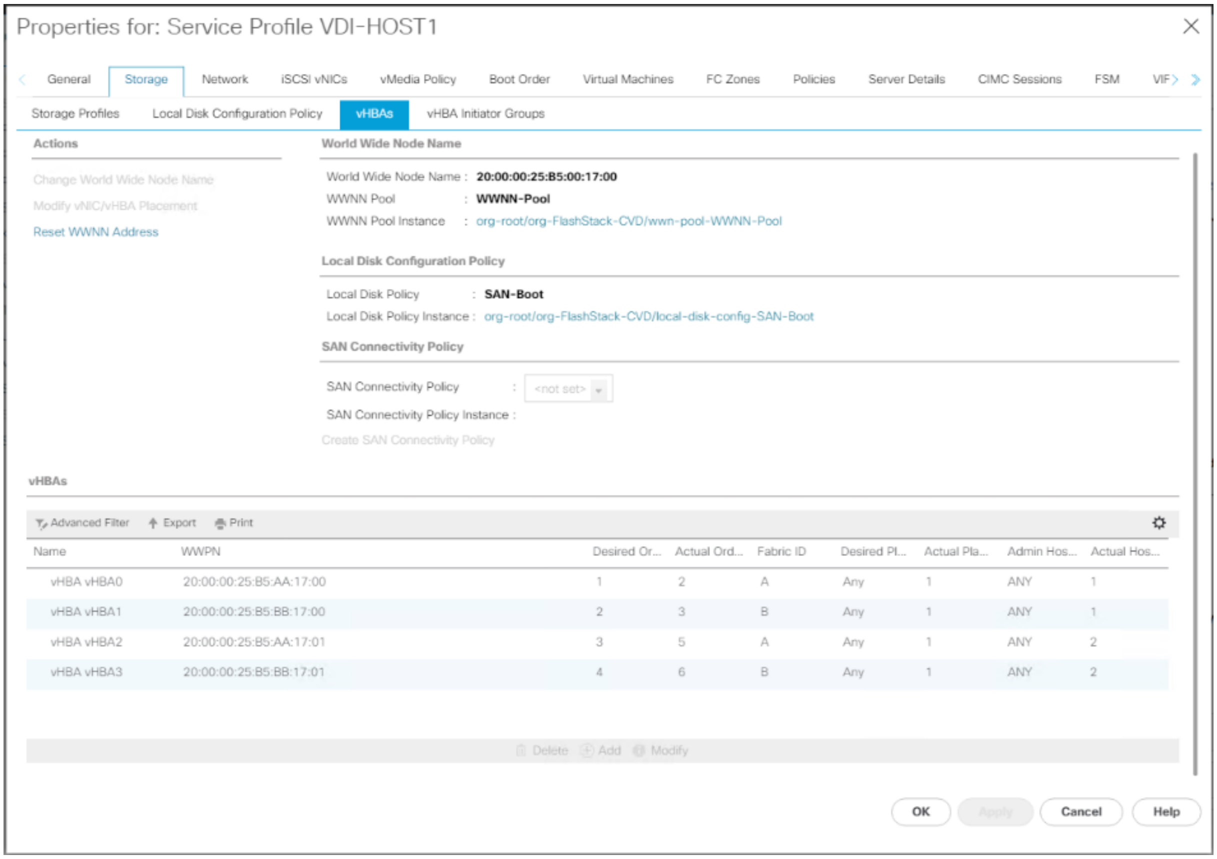

Properties for: Service Profile VDI-HOST1

World Wide Node Name

World Wide Node Name : **20:00:00:25:B5:00:17:00**
WWNN Pool : **WWNN-Pool**
WWNN Pool Instance : org-root/org-FlashStack-CVD/wwn-pool-WWNN-Pool

Local Disk Configuration Policy

Local Disk Policy : **SAN-Boot**
Local Disk Policy Instance : org-root/org-FlashStack-CVD/local-disk-config-SAN-Boot

SAN Connectivity Policy

SAN Connectivity Policy : <not set>
SAN Connectivity Policy Instance :
Create SAN Connectivity Policy

vHBAs

Name	WWPN	Desired Or...	Actual Ord...	Fabric ID	Desired Pl...	Actual Pla...	Admin Hos...	Actual Hos...
vHBA vHBA0	20:00:00:25:B5:AA:17:00	1	2	A	Any	1	ANY	1
vHBA vHBA1	20:00:00:25:B5:BB:17:00	2	3	B	Any	1	ANY	1
vHBA vHBA2	20:00:00:25:B5:AA:17:01	3	5	A	Any	1	ANY	2
vHBA vHBA3	20:00:00:25:B5:BB:17:01	4	6	B	Any	1	ANY	2

2. Connect to the Pure Storage System Health and go to the Connections tab and extract the WWPN of FC Ports connected to the Cisco MDS Switches from Array Ports section. We have connected 4 FC ports from Pure Storage System to Cisco MDS Switches. FC ports CT0.FC0, CT1.FC0 are connected to MDS Switch-A and similarly FC ports CT1.FC2, CT0.FC2 are connected to MDS Switch-B.

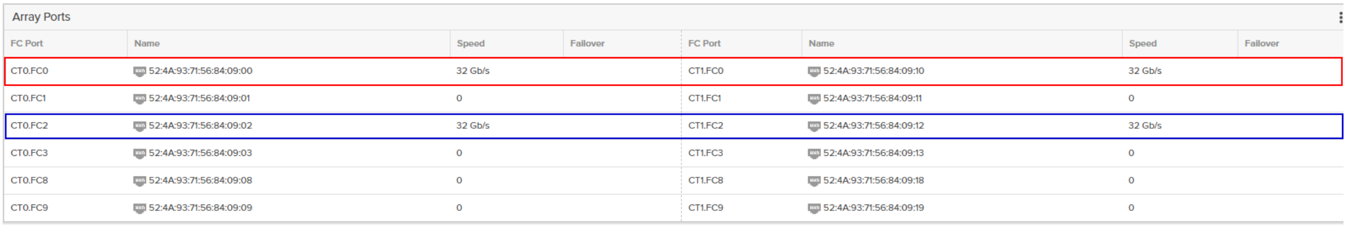

Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Create Device Aliases for Fiber Channel Zoning

Cisco MDS Switch A

To configure device aliases and zones for the SAN boot paths as well as the datapaths of MDS switch A, follow these steps:

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

```
configure terminal
device-alias database
```

```
device-alias name VDI-Host01-HBA0 pwwn 20:00:00:25:B5:AA:17:00
device-alias name VDI-Host01-HBA2 pwwn 20:00:00:25:B5:AA:17:01
device-alias name X70R3-CT0-FC0 pwwn 52:4A:93:71:56:84:09:00
device-alias name X70R3-CT1-FC0 pwwn 52:4A:93:71:56:84:09:10
exit
device-alias commit
```

Cisco MDS Switch B

To configure device aliases and zones for the SAN boot paths as well as datapaths of MDS switch B, follow this step:

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

```
configure terminal
device-alias database
device-alias name VDI-Host01-HBA1 pwwn 20:00:00:25:B5:AA:17:00
device-alias name VDI-Host01-HBA3 pwwn 20:00:00:25:B5:AA:17:01
device-alias name X70R3-CT0-FC2 pwwn 52:4A:93:71:56:84:09:02
device-alias name X70R3-CT1-FC2 pwwn 52:4A:93:71:56:84:09:12
exit
device-alias commit
```

Create Fiber Channel Zoning

Cisco MDS Switch A

To configure zones for the MDS switch A, follow these steps to create a zone for each server service profile:

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

```
configure terminal
zone name FlaskStack-VCC-CVD-WLHost01 vsan 100
    member pwwn 52:4A:93:71:56:84:09:00
    member pwwn 52:4A:93:71:56:84:09:10
    member pwwn 20:00:00:25:b5:aa:17:00
    member pwwn 20:00:00:25:b5:aa:17:01
```

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

```
configure terminal
zoneset name FlashStack-VDI-CVD vsan 100
member FlaskStack-VDI-CVD-Host01
```

3. Activate the zone set by running following commands:

```
zoneset activate name FlashStack-VDI-CVD vsan 100
exit
copy running-config startup-config
```

Cisco MDS Switch B

To configure zones for the MDS switch B, follow these steps to create a zone for each server service profile:

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

```
configure terminal
zone name FlaskStack-VCC-CVD-WLHost01 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:00
    member pwwn 20:00:00:25:b5:bb:17:01
    member pwwn 52:4a:93:71:56:84:09:02
    member pwwn 52:4a:93:71:56:84:09:12
```

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

```
zoneset name FlashStack-VDI-CVD vsan 101
member FlaskStack-VDI-CVD-Host01
```

3. Activate the zone set by running following commands:

```
zoneset activate name FlashStack-VDI-CVD vsan 101
exit
copy running-config startup-config
```

Configure Pure Storage FlashArray//X70 R3

The design goal of the reference architecture is to best represent a real-world environment as closely as possible. The approach included the features of Cisco UCS to rapidly deploy stateless servers and use Pure Storage FlashArray's boot LUNs to provision the ESXi on top of Cisco UCS. Zoning was performed on the Cisco MDS 9132T 32-Gb switches to enable the initiators discover the targets during boot process.

A Service Profile was created within Cisco UCS Manager to deploy the thirty-two servers quickly with a standard configuration. SAN boot volumes for these servers were hosted on the same Pure Storage FlashArray//X70 R3. Once the stateless servers were provisioned, following process was performed to enable rapid deployment of thirty-two Blade Servers.

Each Blade Server has dedicated single LUN to install operating system and all the thirty-two Blade Servers configured to boot from SAN. For this solution, we have installed vSphere ESXi 7.0 GA Cisco Custom ISO on this LUNs to create solution.

Using logical servers that are disassociated from the physical hardware removes many limiting constraints around how servers are provisioned. Cisco UCS Service Profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and HA information. The service profiles represent all the attributes of a logical server in Cisco UCS model. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. Furthermore, Cisco is the only hardware provider to offer a truly unified management platform, with Cisco UCS Service Profiles and hardware abstraction capabilities extending to both blade and rack servers.

In addition to the service profiles, the use of Pure Storage's FlashArray's with SAN boot policy provides the following benefits:

- Scalability - Rapid deployment of new servers to the environment in a very few steps.
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a huge benefit in comparison to another appliance model like Exadata.
- Flexibility - Easy to repurpose physical servers for different applications and services as needed.
- Availability - Hardware failures are not impactful and critical. In rare case of a server failure, it is easier to associate the logical service profile to another healthy physical server to reduce the impact.

Configure Host, WWNs, and Volume Connectivity with FlashArray Management Tools

Configure Host

Before using a boot volume (LUN) by a Cisco UCS Blade Server, a host representing this blade server must be defined on Pure Storage FlashArray. To set up a host, follow these steps:

1. Log into Pure Storage FlashArray Management interface.
2. Click the Storage tab.
3. Click the + sign in the Hosts section and select Create Host.

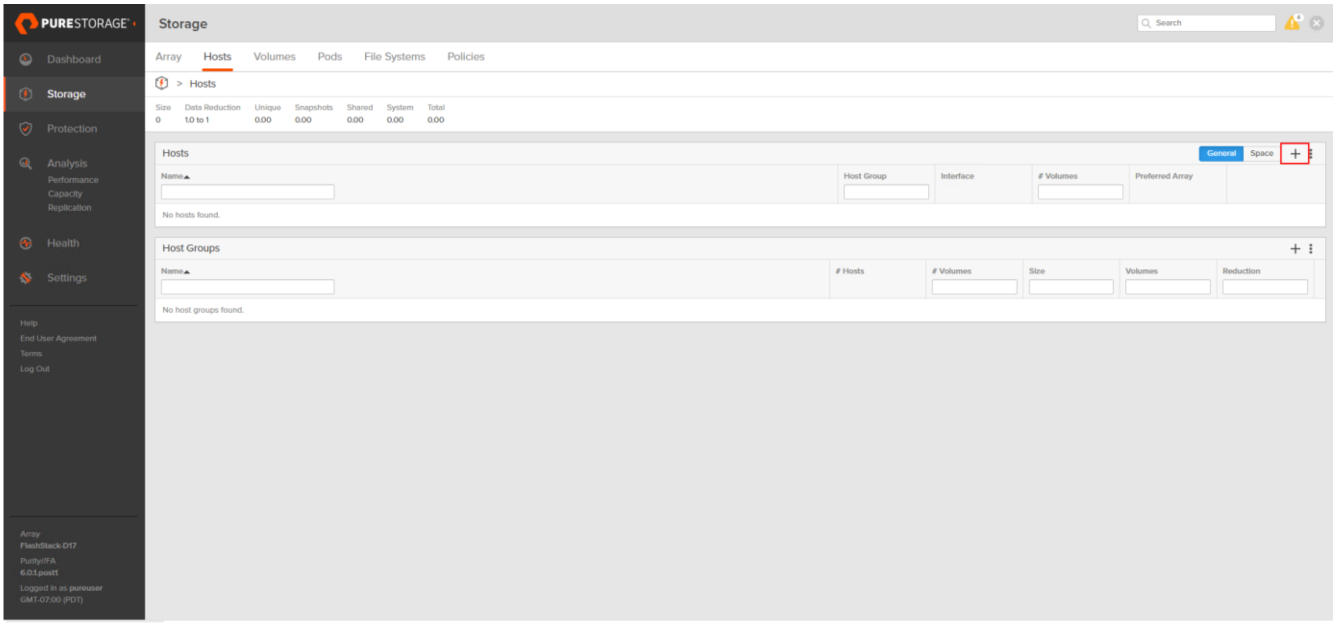

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	0.00	0.00	0.00

Hosts

Name▲	Host Group	Interface	# Volumes	Preferred Array
No hosts found.				

Host Groups

Name▲	# Hosts	# Volumes	Size	Volumes	Reduction
No host groups found.					

Array
FlashStack-D17
Purity//FA
6.0.1.post1
Logged in as pureuser
GMT-07:00 (PDT)

4. Select Create Multiple to create a Host entries under the Hosts category.

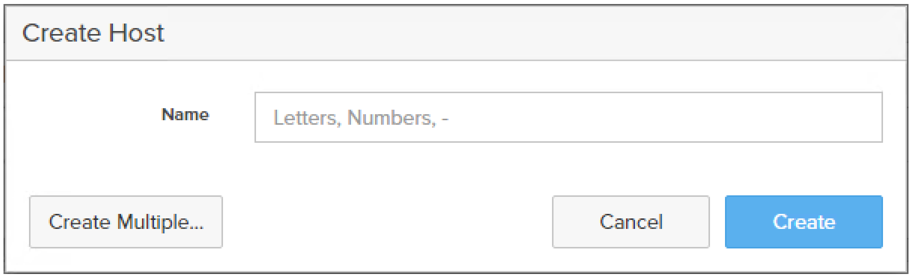

Create Host

Name: Letters, Numbers, -

5. Enter the required information and click Create.

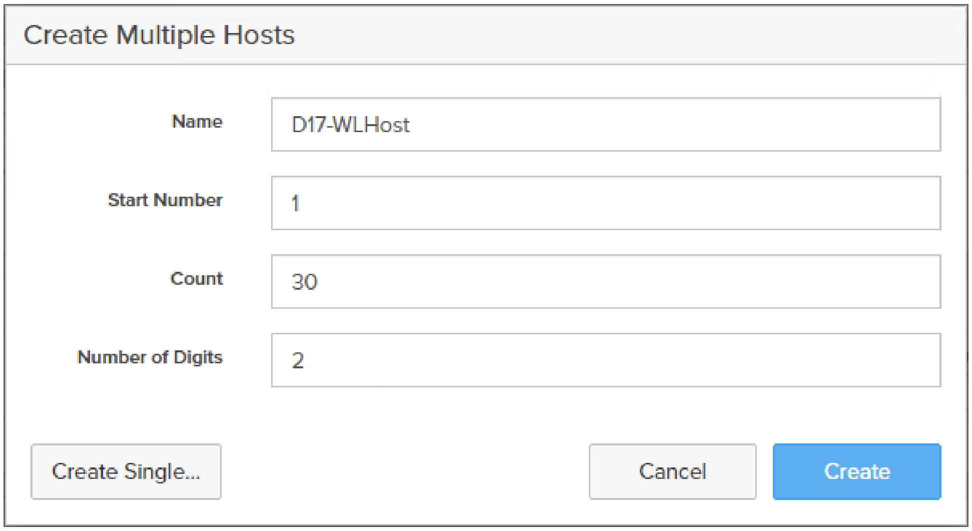

Create Multiple Hosts

Name: D17-WLHost

Start Number: 1

Count: 30

Number of Digits: 2

6. Select one of the newly created hosts, in Host Ports section from the drop-down list select “Configure WWNs.”

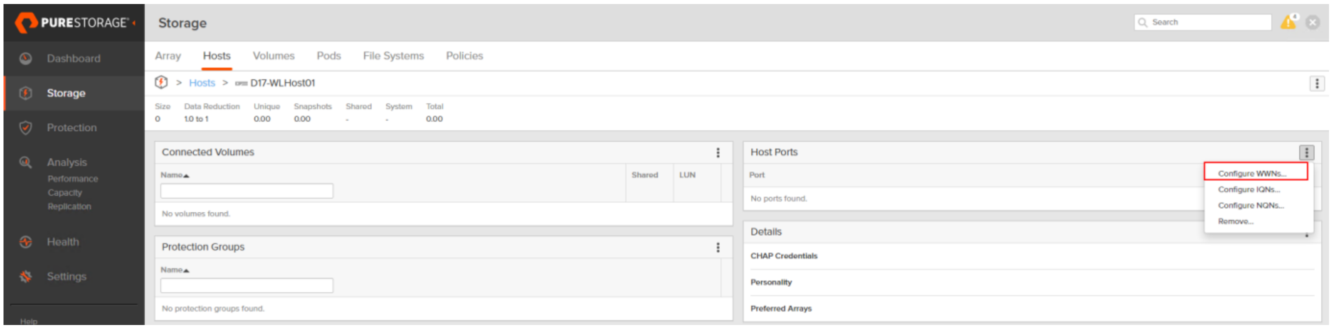

7. Select the list of WWNs that belongs to the host in the next window and click Add.

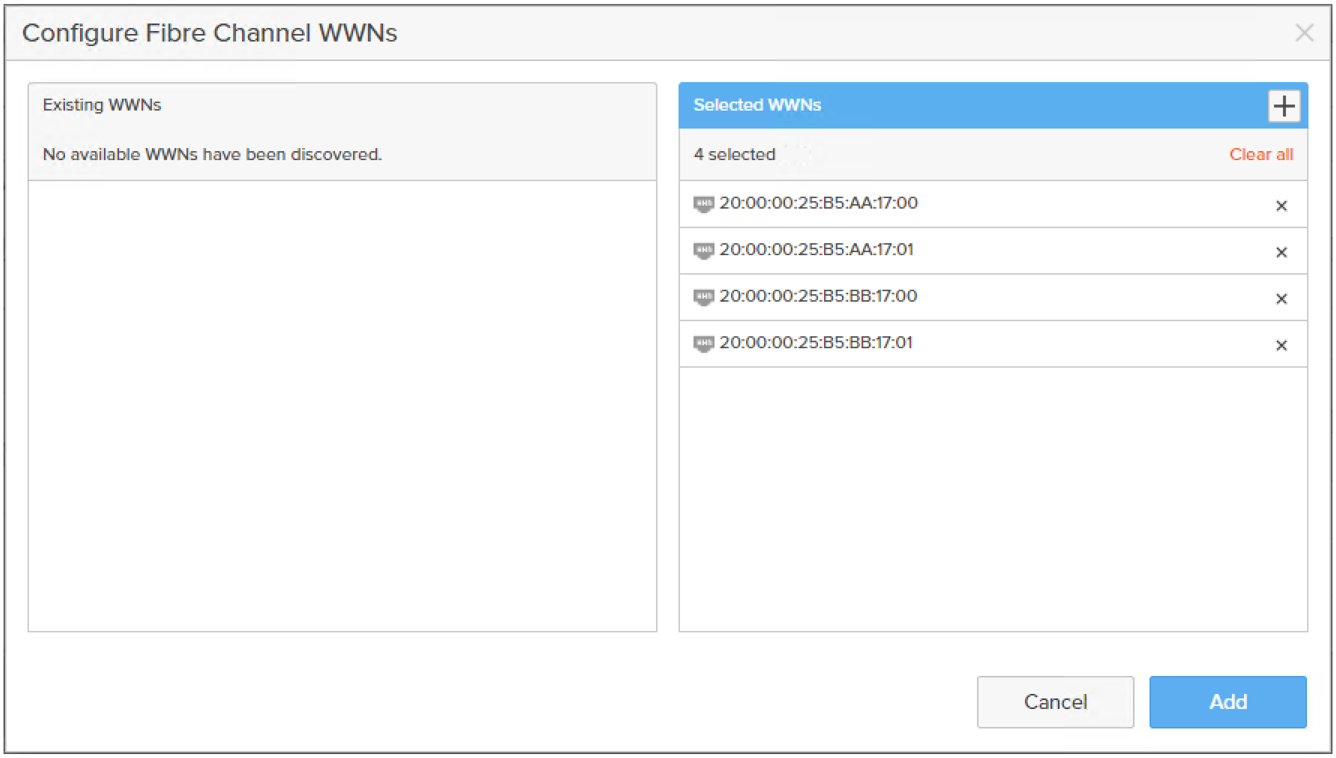

Configure Fibre Channel WWNs

Existing WWNs

No available WWNs have been discovered.

Selected WWNs

4 selected

20:00:00:25:B5:AA:17:00

20:00:00:25:B5:AA:17:01

20:00:00:25:B5:BB:17:00

20:00:00:25:B5:BB:17:01

Make sure the zoning has been setup to include the WWNs details of the initiators along with the target, without which the SAN boot will not work.

WWNs will appear only if the appropriate FC connections were made, and the zones were setup on the underlying FC switch.

Alternatively, the WWN can be added manually by clicking the + in the Selected WWNs section and manually inputting the blade’s WWNs.

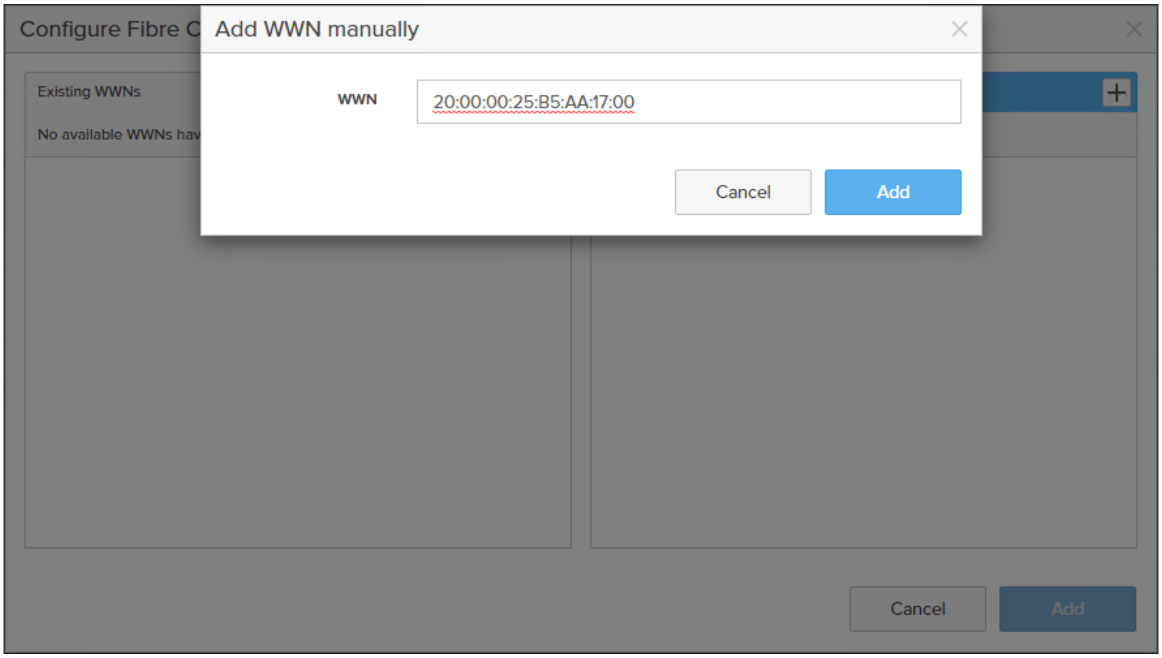

WWN: 20:00:00:25:B5:AA:17:00

Configure Volume Connectivity

To configure a volume and volume connectivity, follow these steps:

1. Click the Storage tab.

2. Click the + sign in the Volumes section and click Create Volume.

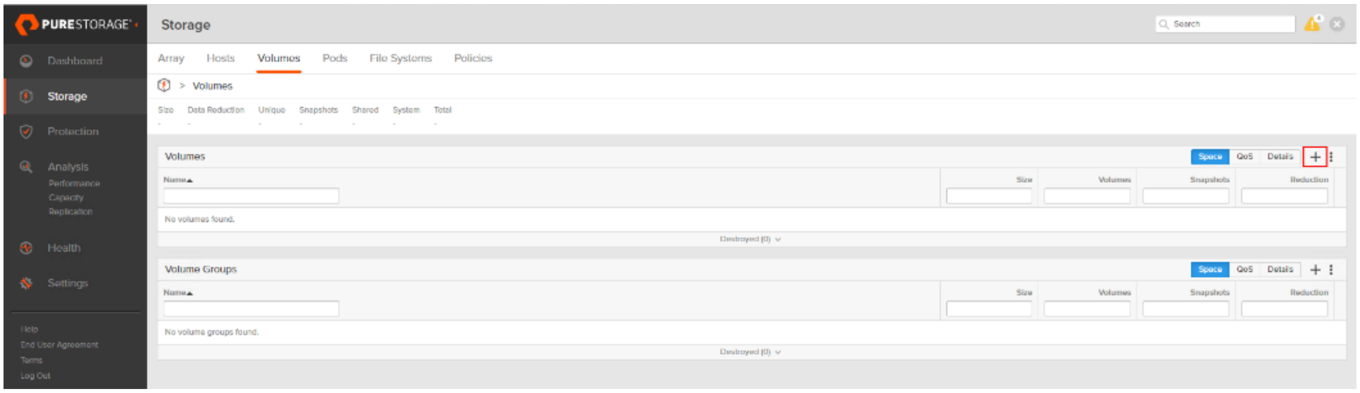

3. Select Create Multiple to open Create Multiple Volumes wizard.

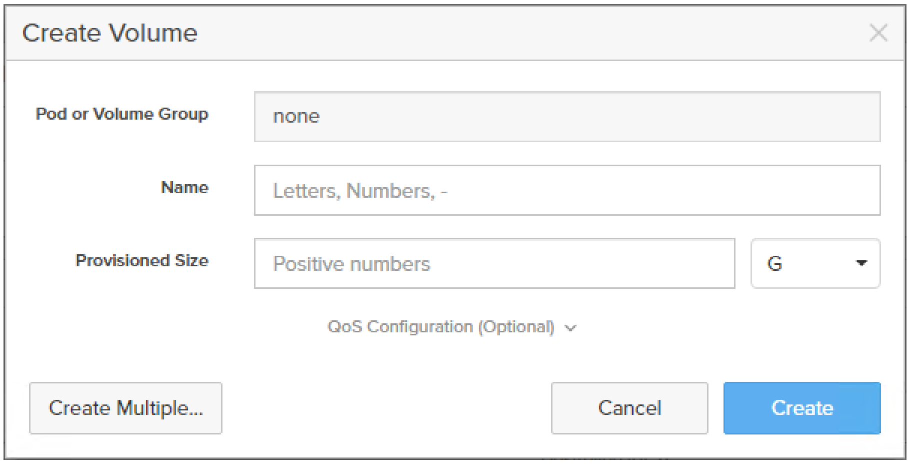

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

4. Provide the common name of the volume, size, choose the size type (KB, MB, GB, TB, PB) and click Create to create volumes.

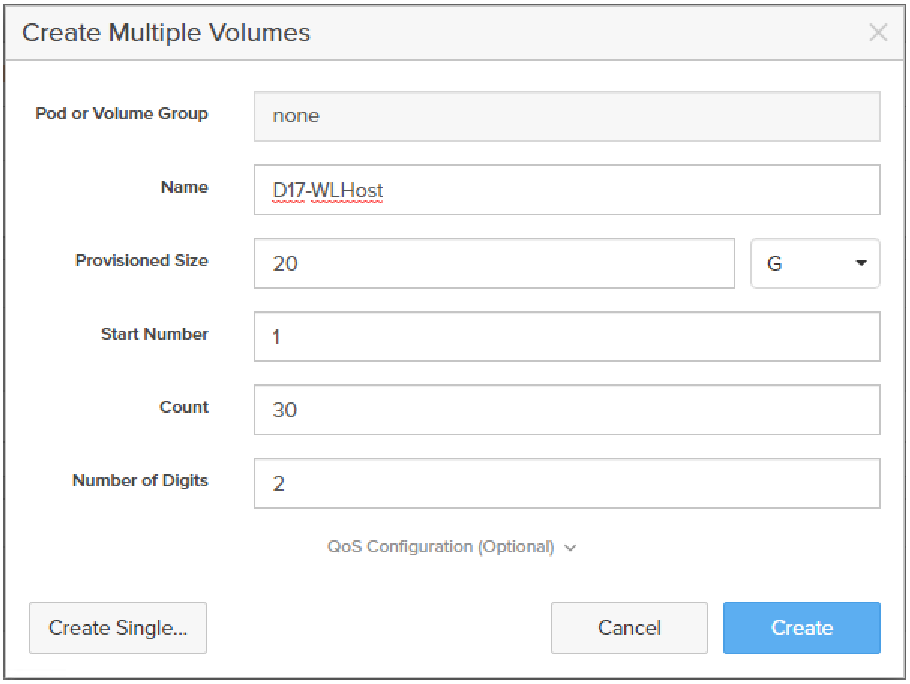

Create Multiple Volumes

Pod or Volume Group: none

Name: D17-WLHost

Provisioned Size: 20 G

Start Number: 1

Count: 30

Number of Digits: 2

5. Select one of the hosts and in Connected Volumes section from the drop-down list select Connect.

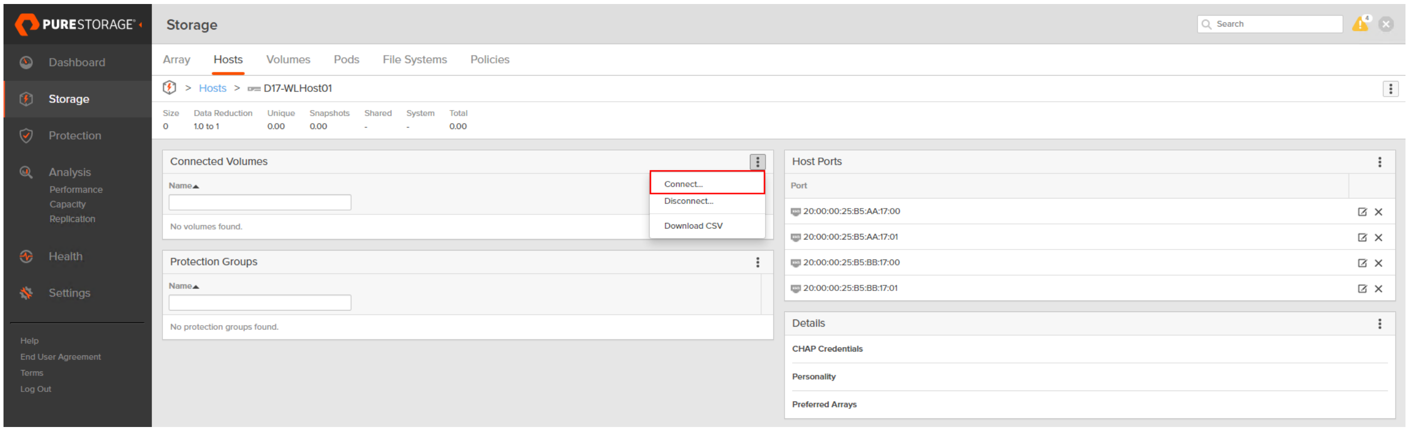

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

No volumes found.

No protection groups found.

Port
20:00:00:25:B5:AA:17:00
20:00:00:25:B5:AA:17:01
20:00:00:25:B5:BB:17:00
20:00:00:25:B5:BB:17:01

6. In the Connect Volumes to Host wizard select the volume configured for ESXi installation, click Connect.

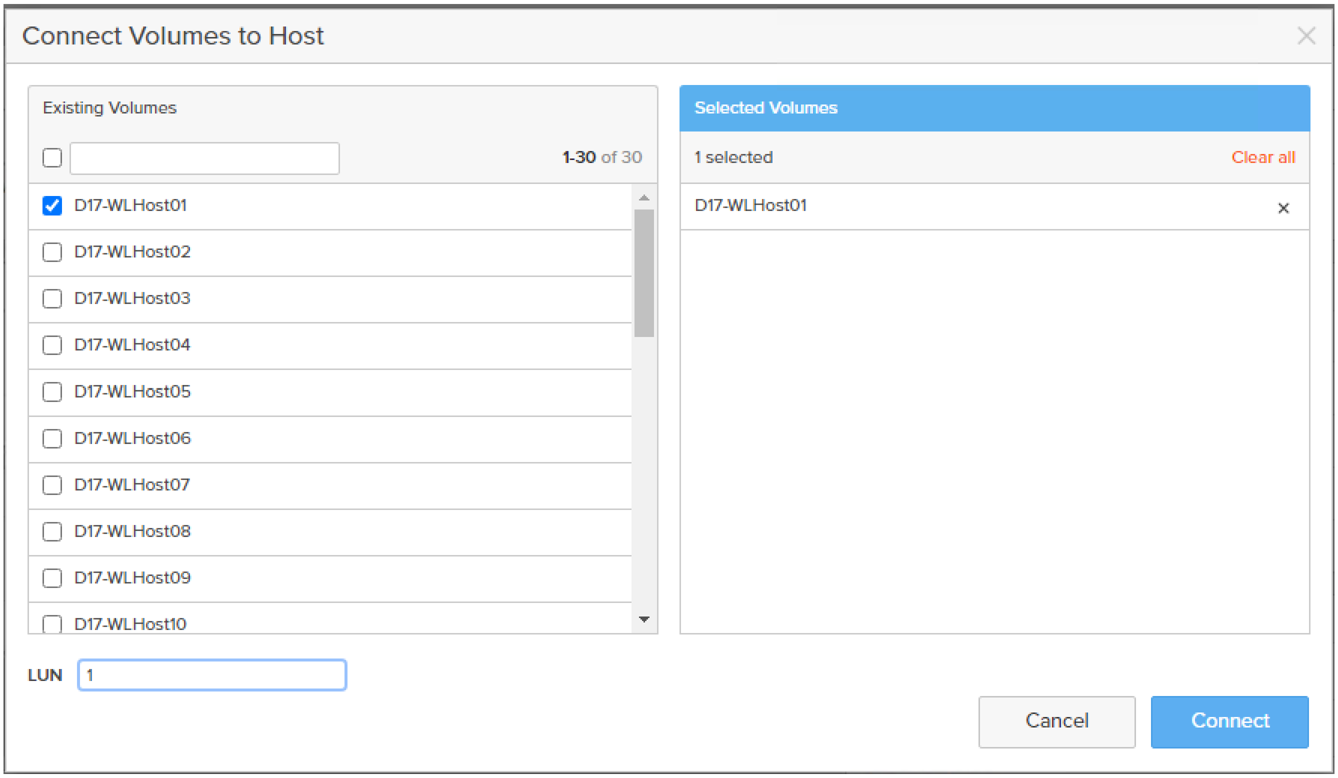

Existing Volumes — 1-30 of 30

- ☑ D17-WLHost01
- ☐ D17-WLHost02
- ☐ D17-WLHost03
- ☐ D17-WLHost04
- ☐ D17-WLHost05
- ☐ D17-WLHost06
- ☐ D17-WLHost07
- ☐ D17-WLHost08
- ☐ D17-WLHost09
- ☐ D17-WLHost10

Selected Volumes — 1 selected

D17-WLHost01 ×

LUN: 1

Make sure the SAN Boot Volumes has the LUN ID “1” since this is important while configuring Boot from SAN. You will also configure the LUN ID as “1” when configuring Boot from SAN policy in Cisco UCS Manager.

More LUNs can be connected by adding a connection to existing or new volume(s) to an existing node.

Configure File Services

FA File services can be activated by Pure Storage Technical Services (Support). Please refer to [FA File Services Support Matrix]() to verify that your hardware offers support for running File Services.

Currently all FA File services activations require Pure Storage Product Management approval. Customers can work with their local account representatives to obtain approval to activate File Services.

Create Virtual Interface(s)

The VIF provides high-availability network access across 2 physical Ethernet ports per array controller. Each VIF requires 2 physical ports per controller. Any physical ethernet port can be used with the restriction that any port that is in use by management services, a bond, or subnet configuration cannot be part of a VIF. For the maximum number of VIFs supported, please see the FA File Services Limits KB.

VIFs created by CLI over SSH, configured and enabled via Management Console. Account with administrator privileges is required.

To create File Virtual Interface, follow these steps:

1. Connect to the array via SSH.

2. Run the following syntax to create the VIF on the array:

```
purenetwork create vif --subinterfacelist ct0.ethX,ct1.ethX,ct0.ethY,ct1.ethY <name of interface>
```

Configure and Enable the Virtual Interface for File Services

To configure and enable the virtual interface, follow these steps:

1. Connect to the array GUI.

2. Navigate to Settings > Network.

3. Locate the File VIF in the interface list and click the edit icon.

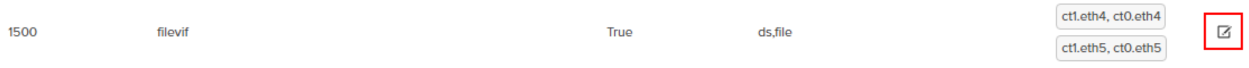

1500	filevif	True	ds,file	ct1.eth4, ct0.eth4 ct1.eth5, ct0.eth5	☐

4. In the Edit Interface dialog turn on the Enabled option, provide the IP Address, Netmask, and Gateway used by the interface. Click Save.

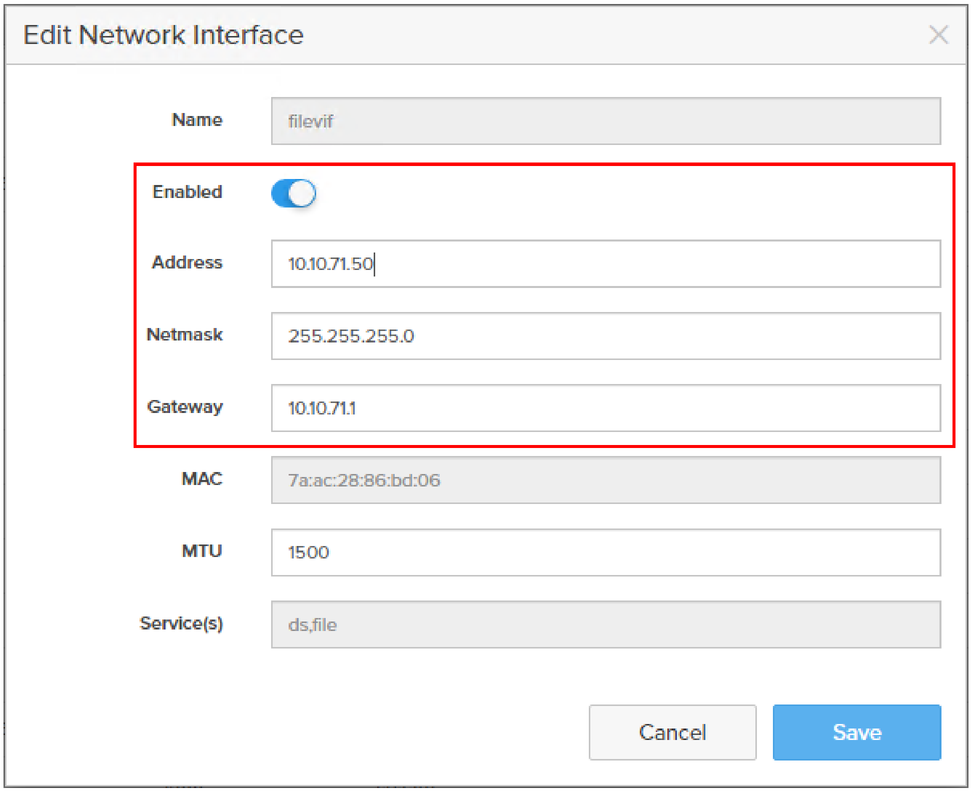

Edit Network Interface

Name	filevif
Enabled	●
Address	10.10.71.50
Netmask	255.255.255.0
Gateway	10.10.71.1
MAC	7a:ac:28:86:bd:06
MTU	1500
Service(s)	ds,file

5. Scroll to the bottom of the Network tab and click the edit icon for DNS Settings.

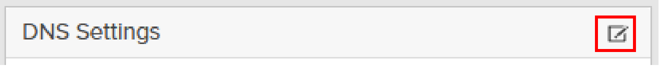

DNS Settings

6. In the Edit DNS Settings dialog, enter desired values for Domain and DNS server IPs. Click Save.

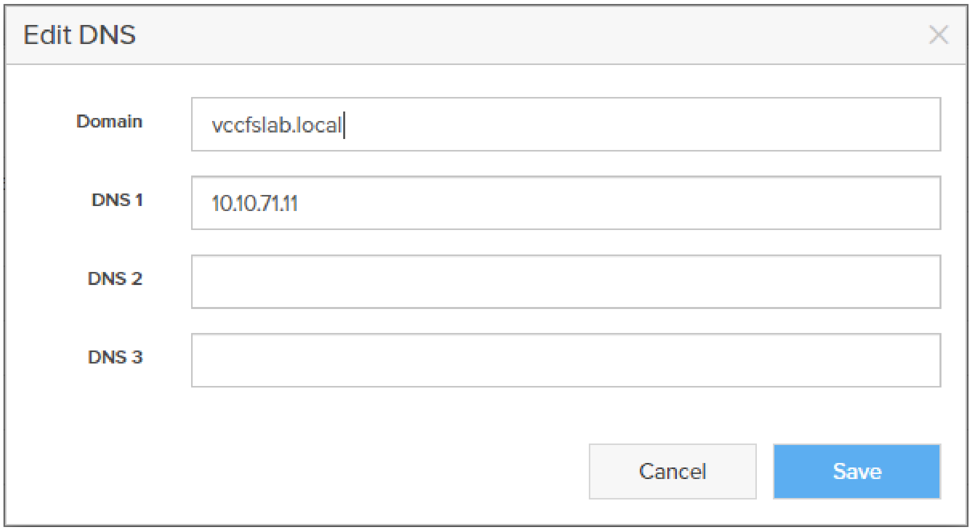

Edit DNS

Domain	vccfslab.local
DNS 1	10.10.71.11
DNS 2	
DNS 3	

More than one DNS server can be configured with the caveat that all DNS servers must have a record for Directory Service servers such as LDAP or Microsoft Active Directory.

Create Active Directory Account for the Array

To create the Active Directory Account, follow these steps:

1. Navigate to Settings > Access > Active Directory Accounts.
2. To open the Create Dialog, click the + icon.

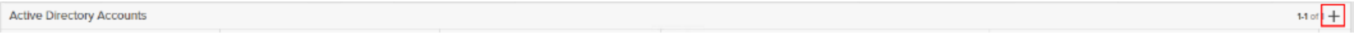

3. Enter the following information:
    a. Name = Array management name for this AD account
    b. Domain = AD domain name
    c. Computer Name = Computer Object name within AD
    d. User = Domain user that can create computer objects and join to the domain.
    e. Password = Users password for the above domain user
4. Click Create to finalize AD account creation.

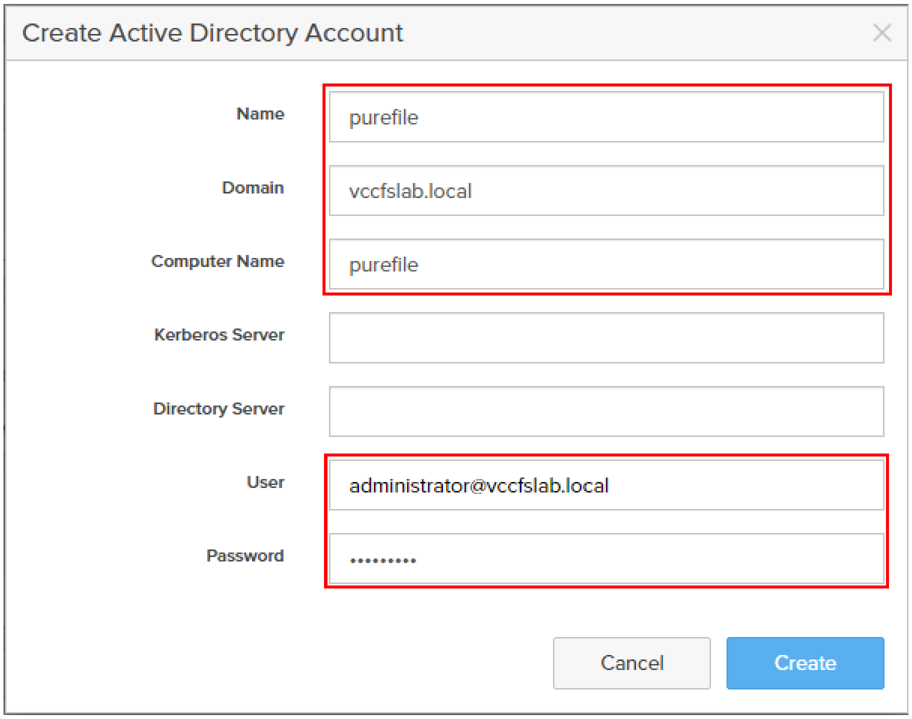

Create Active Directory Account

Name: purefile
Domain: vccfslab.local
Computer Name: purefile
Kerberos Server:
Directory Server:
User: administrator@vccfslab.local
Password: •••••••••

Create a File System and Shared Directory

To create a file system and shared directory, follow these steps:

1. Navigate to Storage > File Systems.

2. Click the + icon.

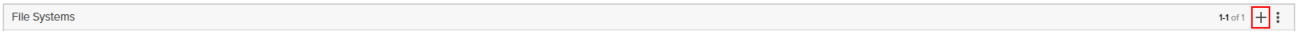

3. In Create File System enter a file system name and click Create.

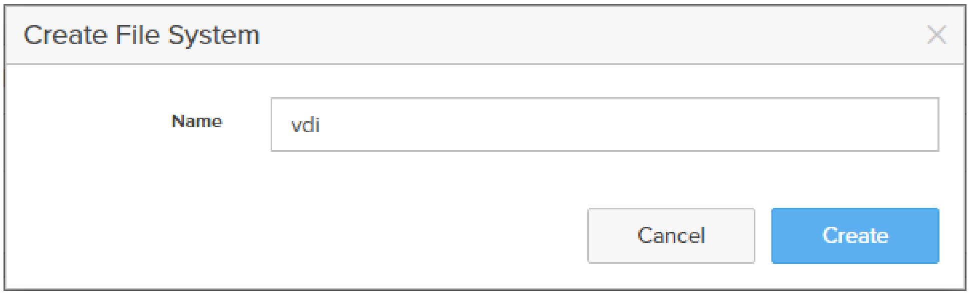
Create File System

Name: vdi

4. Navigate to Storage > File Systems > Directories.

5. Click the + icon.

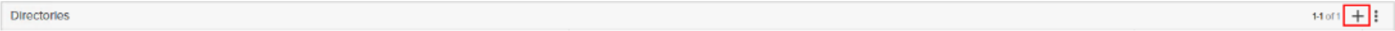

6. In Create Directory pop-up dialog enter Select a file system from the drop-down list, enter the desired management name of the directory, and enter the directory path in the file system. (for example, dir or /dir, for sub-level directories /dir/subdir or /dir/subdir/subdir1 can be used). Click Create.

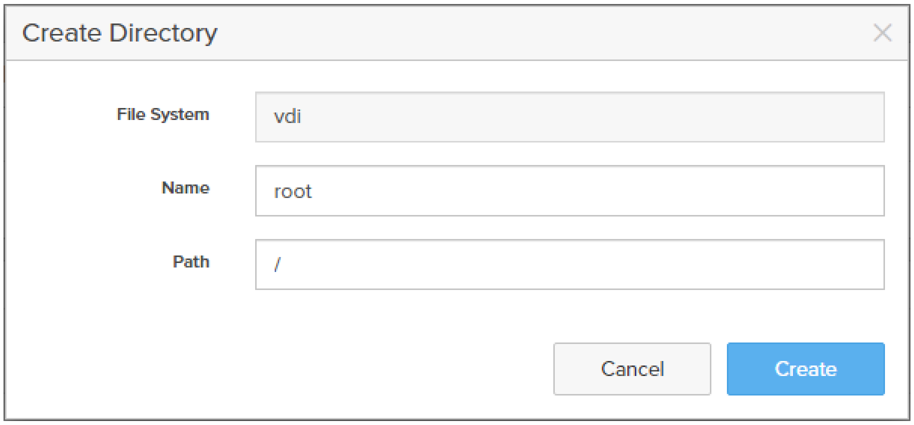
Create Directory

File System: vdi

Name: root

Path: /

Polices for exports/shares/snapshots can only be attached to managed directories at the file system root or 1 level deep (/ and /dir in the example above). Space and performance metrics can be seen at all levels of managed directories.

7. Navigate to Storage > Policies.

8. Click the + icon.

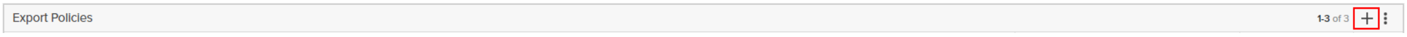

9. In the Create Export Policy pop-up choose SMB from the Type drop-down and enter a name for the policy. Click Create.

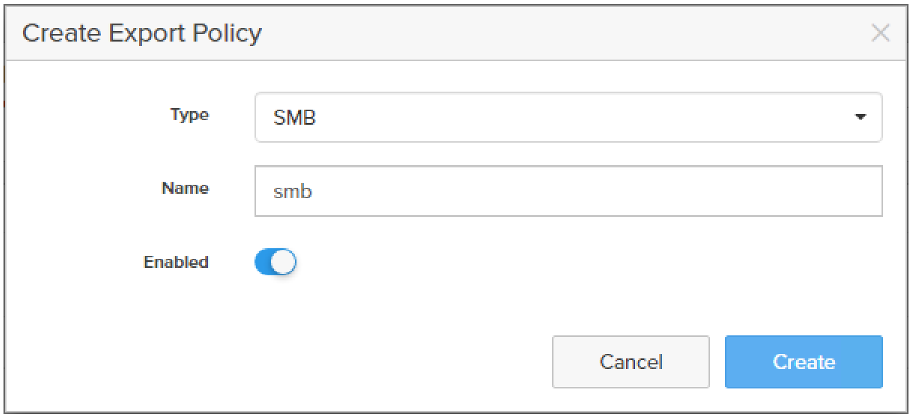

Create Export Policy

Type: SMB

Name: smb

Enabled: ●

10. Select Created Policy and click the + icon.

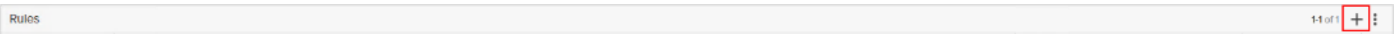

11. Complete the Client filter for read-write access and click Add to complete the rule creation.

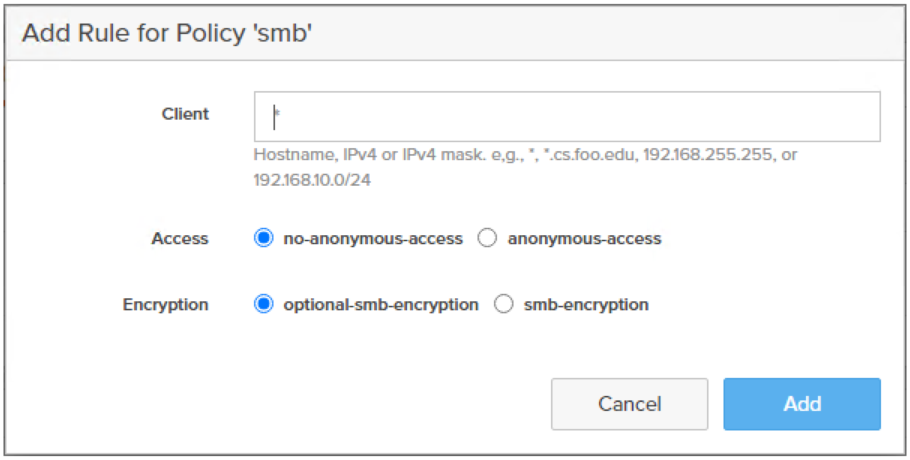

Add Rule for Policy 'smb'

Client: *

Hostname, IPv4 or IPv4 mask. e,g., *, *.cs.foo.edu, 192.168.255.255, or 192.168.10.0/24

Access: ● no-anonymous-access ○ anonymous-access

Encryption: ● optional-smb-encryption ○ smb-encryption

12. Attach the export policy(s) to a managed directory. Click the + icon.

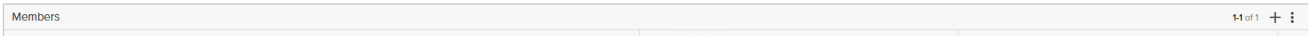

13. Select a managed directory from the drop-down list, enter a share/export name, and click Create.

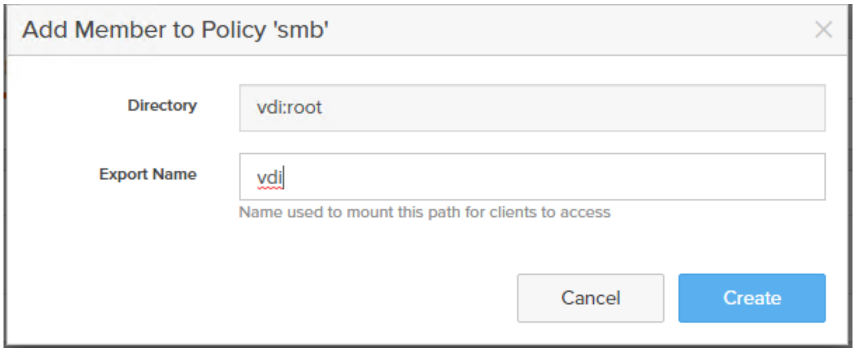

Add Member to Policy 'smb'

Directory: vdi:root

Export Name: vdi

Name used to mount this path for clients to access

14. Verify access to the created share from the Windows client.

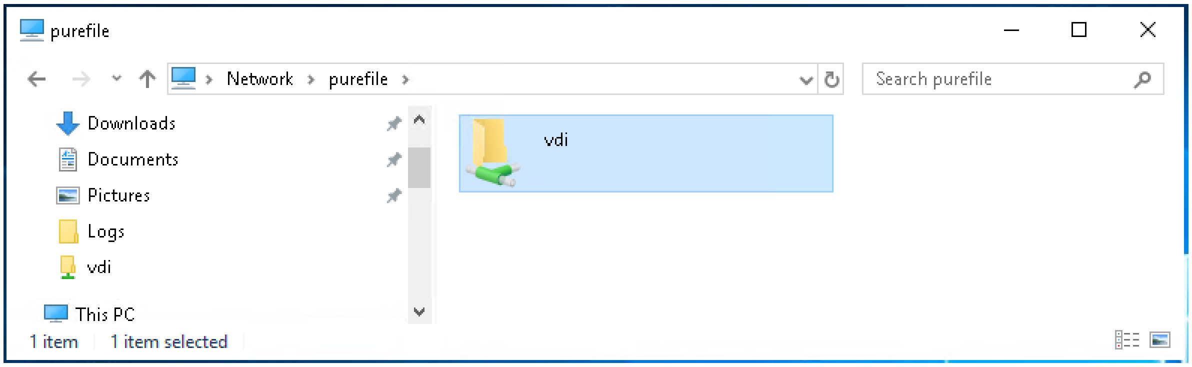

Install and Configure VMware ESXi 7.0

This section explains how to install VMware ESXi 7.0 GA 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.

Download Cisco Custom Image for VMware vSphere ESXi 7.0

To download the Cisco Custom Image for VMware ESXi 7.0 GA, from the VMware vSphere Hypervisor 7.0 GA page click the “Custom ISOs” tab.

Install VMware vSphere ESXi 7.0

To install VMware vSphere ESXi hypervisor on Cisco UCS Server, follow these steps:

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

2. Under Servers > Service Profiles> VDI-Host1

3. Right-click on VDI-Host1 and select KVM Console.

4. Click Activate Virtual Devices and then select CD/DVD.

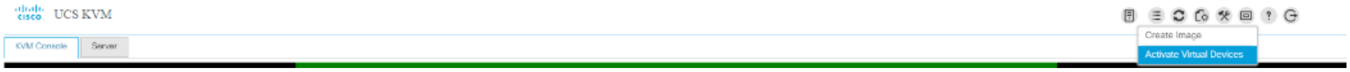

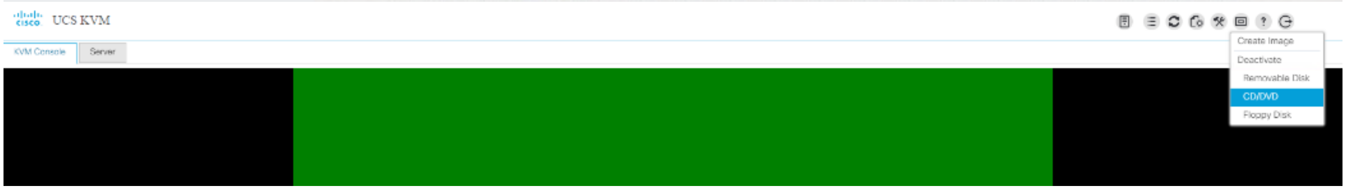

5. Mount the ESXi ISO image.

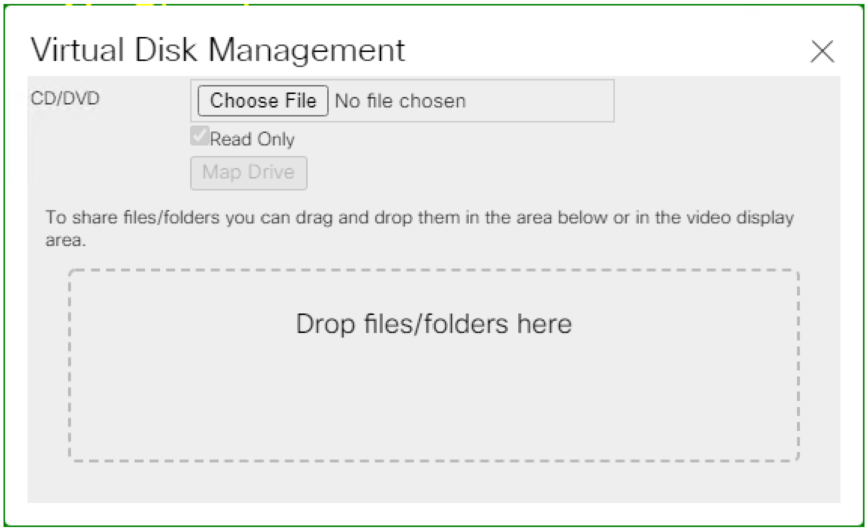
Virtual Disk Management

CD/DVD: No file chosen

☑ Read Only

To share files/folders you can drag and drop them in the area below or in the video display area.

6. Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.

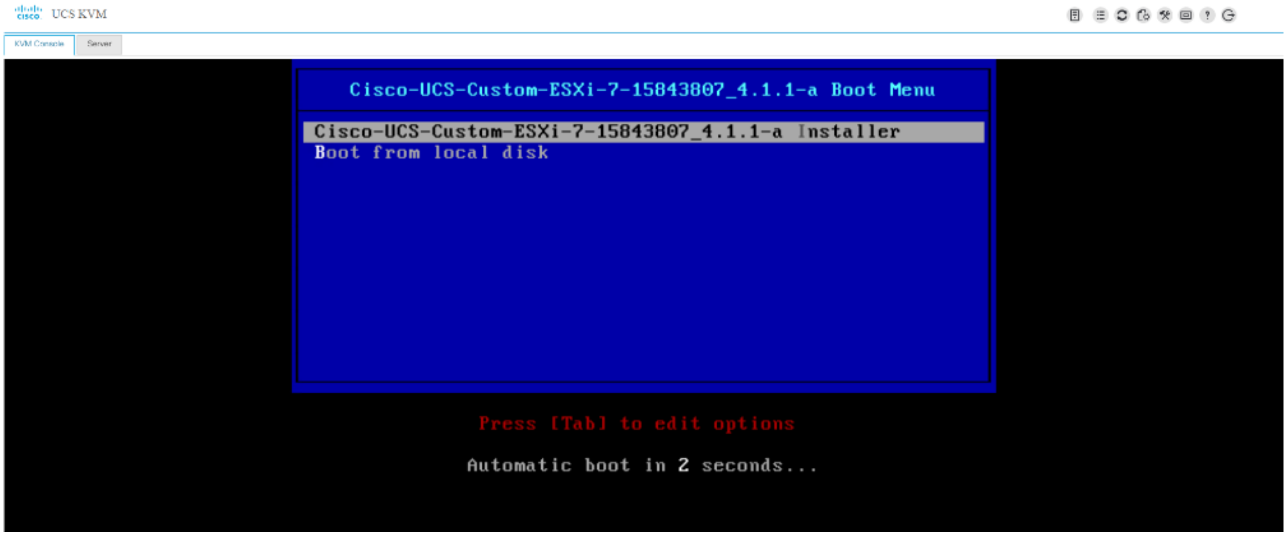

```
Cisco-UCS-Custom-ESXi-7-15843807_4.1.1-a Boot Menu

Cisco-UCS-Custom-ESXi-7-15843807_4.1.1-a Installer
Boot from local disk

Press [Tab] to edit options

Automatic boot in 2 seconds...
```

7. When selecting a storage device to install ESXi, select Remote LUN provisioned through Pure Storage Administrative console and access through FC connection.

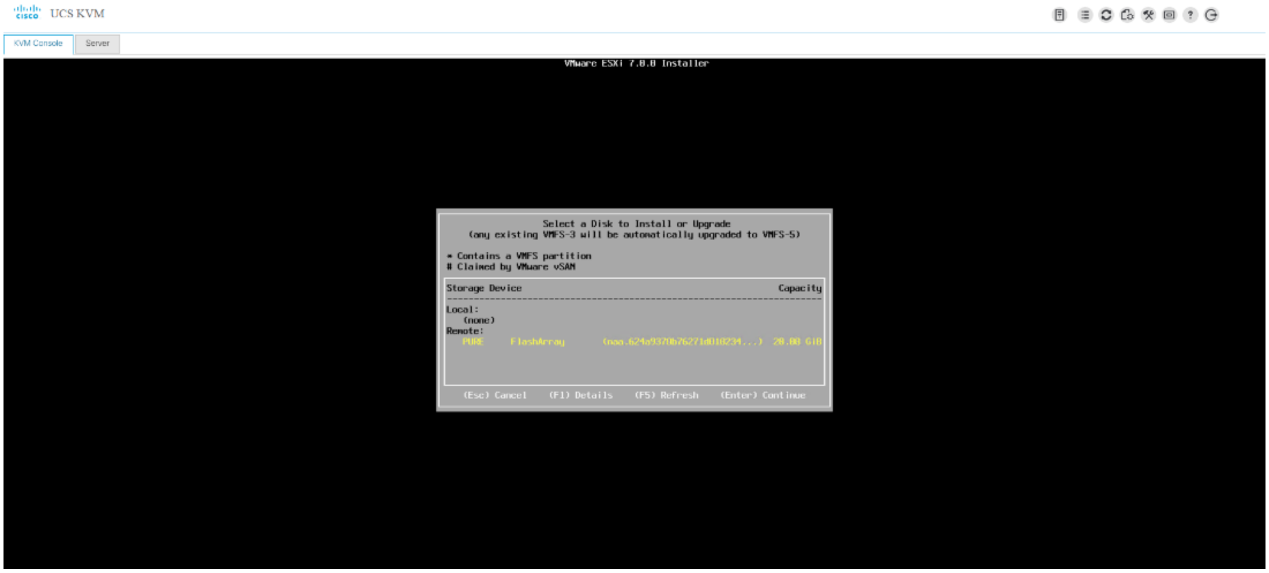

```
VMware ESXi 7.0.0 Installer

                 Select a Disk to Install or Upgrade
  (any existing VMFS-3 will be automatically upgraded to VMFS-5)

* Contains a VMFS partition
# Claimed by VMware vSAN

Storage Device                                              Capacity
--------------------------------------------------------------------
Local:
   (none)
Remote:
   PURE     FlashArray      (naa.624a9370b76271d018234...)  20.00 GiB

  (Esc) Cancel    (F1) Details    (F5) Refresh    (Enter) Continue
```

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.

To configure the ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to enter in to configuration wizard for ESXi Hypervisor.

2. Log in as root and enter the corresponding password.

3. Select the Configure the Management Network option and press Enter.

4. Select the VLAN (Optional) option and press Enter. Enter the VLAN In-Band management ID and press Enter.

5. From the Configure Management Network menu, select “IP Configuration” and press Enter.

6. Select “Set Static IP Address and Network Configuration” option by using the space bar. Enter the IP address to manage the first ESXi host. Enter the subnet mask for the first ESXi host. Enter the default gateway for the first ESXi host. Press Enter to accept the changes to the IP configuration.

7. IPv6 Configuration is set to automatic.

8. Select the DNS Configuration option and press Enter.

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

10. Enter DNS Suffixes.

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

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

Figure 42. Sample ESXi Configure Management Network

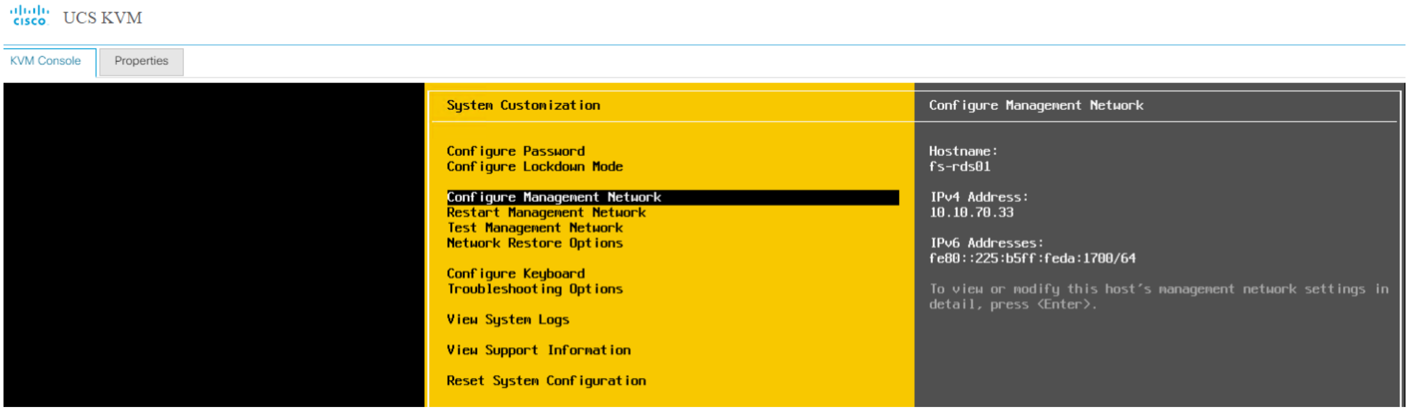

```
System Customization                    Configure Management Network

Configure Password                      Hostname:
Configure Lockdown Mode                 fs-rds01

Configure Management Network            IPv4 Address:
Restart Management Network              10.10.70.33
Test Management Network
Network Restore Options                 IPv6 Addresses:
                                        fe80::225:b5ff:feda:1700/64
Configure Keyboard
Troubleshooting Options                 To view or modify this host's management network settings in
                                        detail, press <Enter>.
View System Logs

View Support Information

Reset System Configuration
```

Update Cisco VIC Drivers for ESXi

When ESXi is installed from Cisco Custom ISO, you might have to update the Cisco VIC drivers for VMware ESXi Hypervisor to match the current Cisco Hardware and Software Interoperability Matrix.

In this Validated Design the following drivers were used:
- Cisco-nenic- 1.0.33.0
- Cisco-nfnic- 4.0.0.56

To update the Cisco VIC drivers for ESXi, follow these steps:

1. Log into your VMware Account to download required drivers for FNIC and NENIC as per the recommendation.

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:

- FlashStack-Datacenter: Pure Storage FlashArray//X70 R3 with Cisco UCS
- Infrastructure Cluster: Infrastructure virtual machines (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware Connection 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 43. VMware vSphere WebUI Reporting Cluster Configuration for this Validated Design

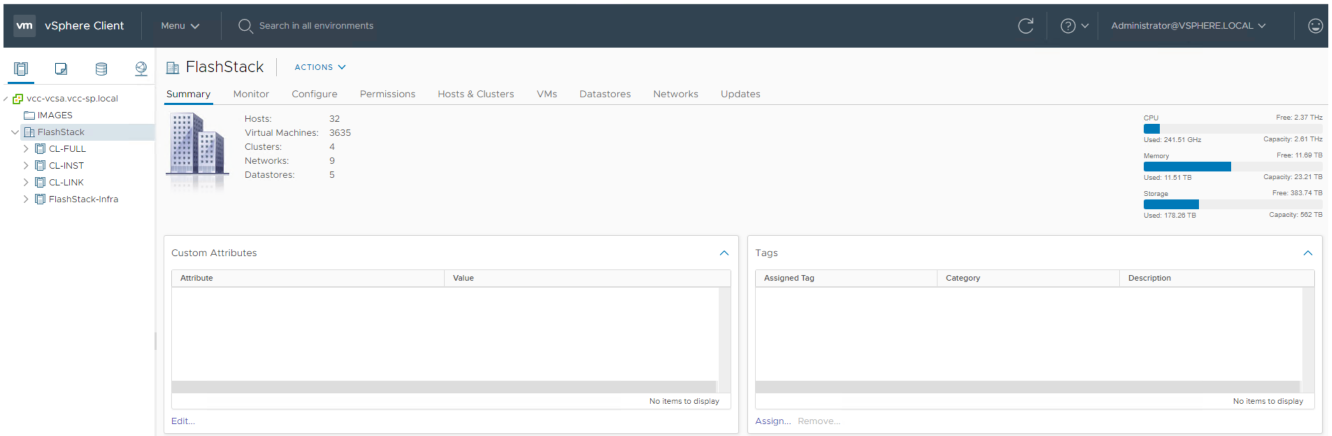

Hosts: 32
Virtual Machines: 3635
Clusters: 4
Networks: 9
Datastores: 5

CPU Free: 2.37 THz
Used: 241.51 GHz Capacity: 2.61 THz
Memory Free: 11.69 TB
Used: 11.51 TB Capacity: 23.21 TB
Storage Free: 383.74 TB
Used: 178.26 TB Capacity: 562 TB

Build the Virtual Machines and Environment for Workload Testing

Prerequisites

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

Figure 44. Example of the DHCP Scopes used in this CVD

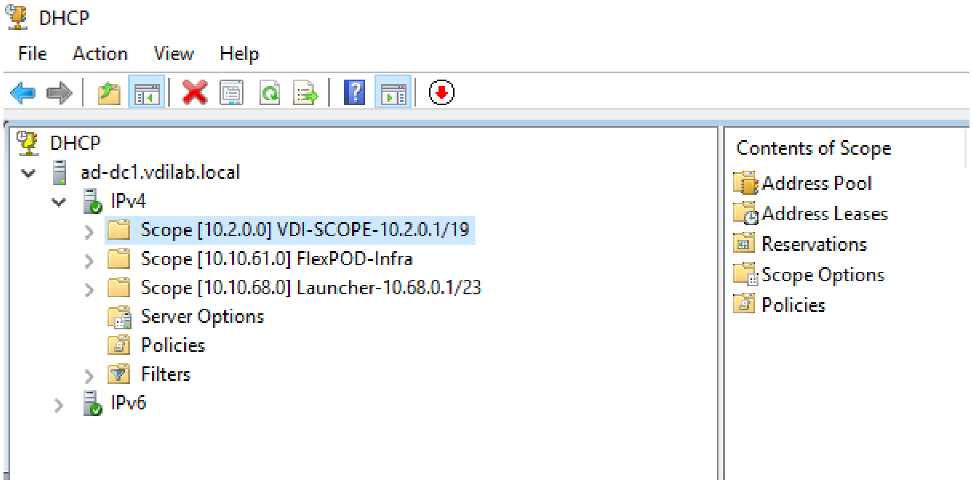

Software Infrastructure Configuration

This section explains how to configure the software infrastructure components that comprise this solution.

Install and configure the infrastructure virtual machines by following the process listed in Table 14.

Table 14. Test Infrastructure Virtual Machine Configuration

Configuration	Microsoft Active Directory DCs	Configuration

Configuration	Microsoft Active Directory DCs	Configuration
Operating system	Microsoft Windows Server 2019	VCSA – SUSE Linux
Virtual CPU amount	2	24
Memory amount	8 GB	48 GB
Network	VMXNET3 Infra	VMXNET3 OOB-Mgmt
Disk size	40 GB	2 TB (across 13 VMDKs)

Configuration	Microsoft SQL Server	
Virtual Machine	VMware Connection Servers	
Virtual Machines		
Operating system	Microsoft Windows Server 2019	
Microsoft SQL Server 2016 SP1	Microsoft Windows Server 2019	
Virtual CPU amount	4	10
Memory amount	16GB	16 GB

Prepare 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 2019 were installed.

To prepare the master virtual machines, there are three major steps: installing 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.

The master target VDI and RDS virtual machines were configured as listed in Table 15.

Table 15. VDI and RDS Virtual Machines Configurations

Configuration	VDI Virtual Machines	RDS Virtual Machines
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2019
Virtual CPU amount	2	10
Memory amount	3.5 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 2019 Login VSI 4.1.25 (Knowledge Worker Workload)	Microsoft Office 2019 Login VSI 4.1.25 (Knowledge Worker Workload)

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

VMware Horizon Agent Installation

To install the VMware Horizon Agent, follow these steps:

1. Download VMware-Horizon-Agent-x86_64-2012-8.1.0-17352461 version.

2. Click the VMware Horizon Agent installer.

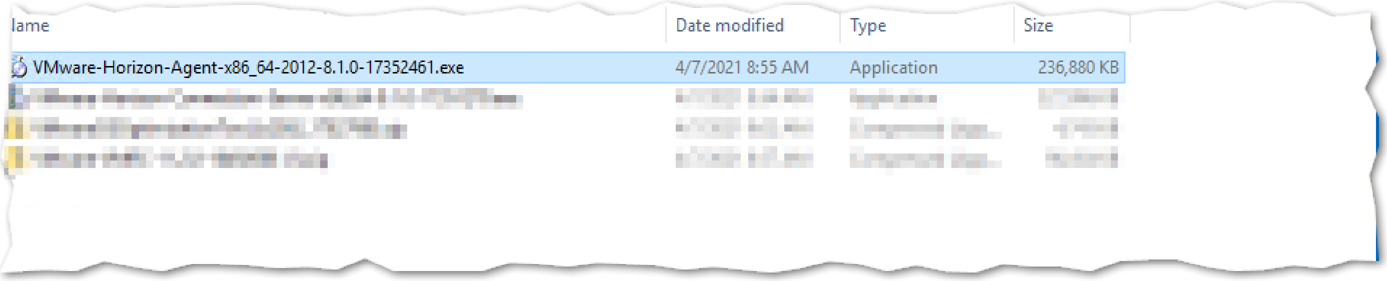

lame	Date modified	Type	Size
VMware-Horizon-Agent-x86_64-2012-8.1.0-17352461.exe	4/7/2021 8:55 AM	Application	236,880 KB

3. Click Next.

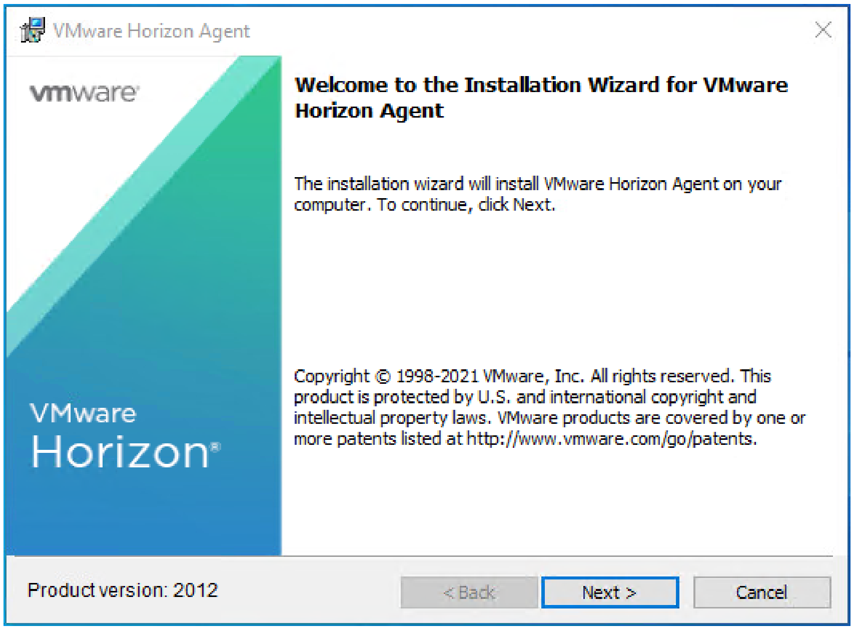

VMware Horizon Agent

Welcome to the Installation Wizard for VMware Horizon Agent

The installation wizard will install VMware Horizon Agent on your computer. To continue, click Next.

Copyright © 1998-2021 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 2012

4. Accept the license agreement and click Next.

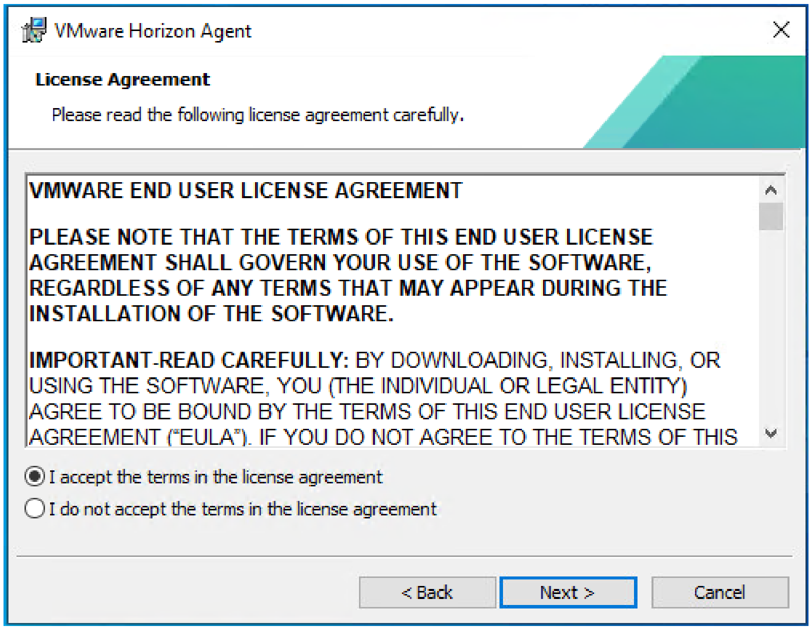

VMware Horizon Agent

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

5. Chose Desktop OS Configuration:

6. During the installation on the Windows 10 select Desktop Mode for the agent installation.

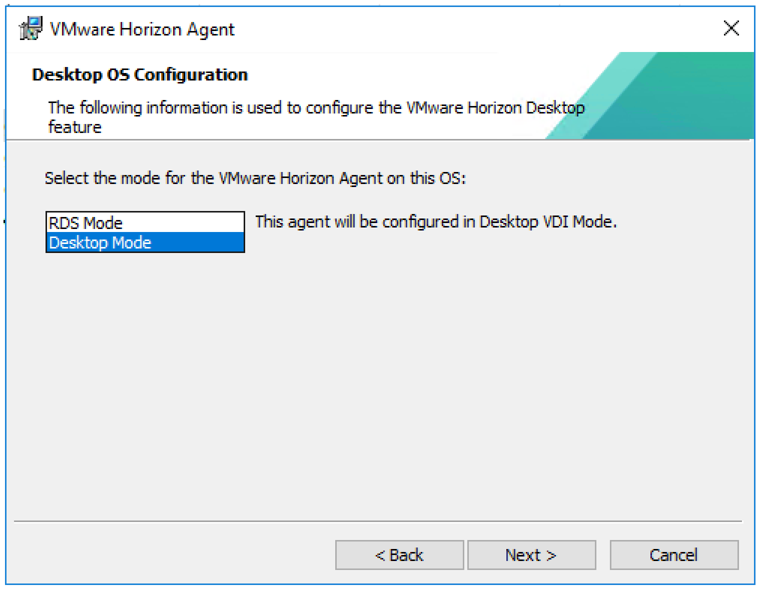

VMware Horizon Agent

Desktop OS Configuration

The following information is used to configure the VMware Horizon Desktop feature

Select the mode for the VMware Horizon Agent on this OS:

RDS Mode
[✓ Desktop Mode]

This agent will be configured in Desktop VDI Mode.

7. During the installation on the Windows 2019 Server select RDS Mode for the agent installation.

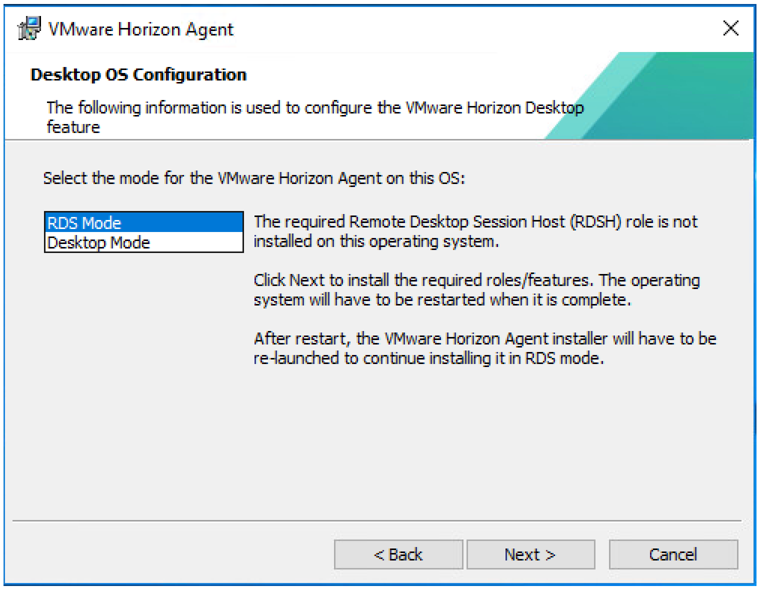

VMware Horizon Agent

Desktop OS Configuration

The following information is used to configure the VMware Horizon Desktop feature

Select the mode for the VMware Horizon Agent on this OS:

[✓ RDS Mode]
Desktop Mode

The required Remote Desktop Session Host (RDSH) role is not installed on this operating system.

Click Next to install the required roles/features. The operating system will have to be restarted when it is complete.

After restart, the VMware Horizon Agent installer will have to be re-launched to continue installing it in RDS mode.

8. Select the default IPV4 and click Next.

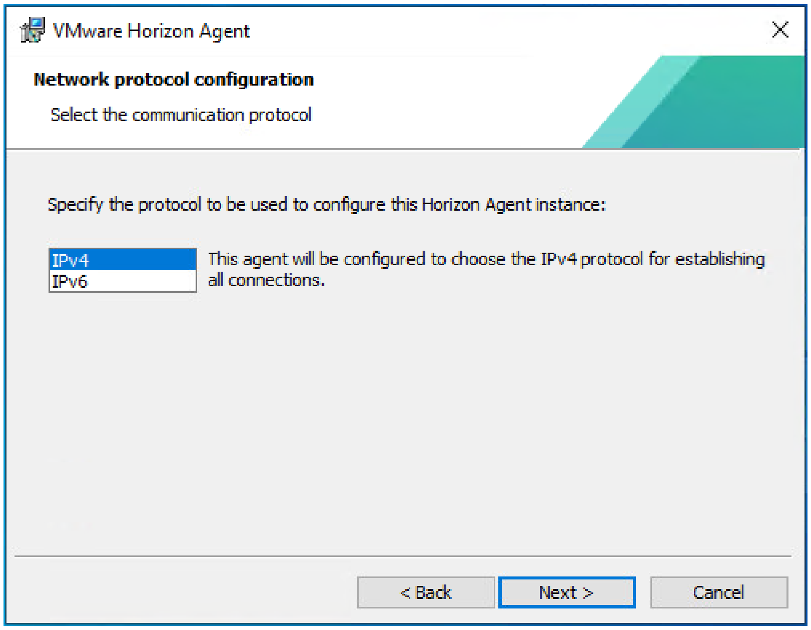

Network protocol configuration

Select the communication protocol

Specify the protocol to be used to configure this Horizon Agent instance:

IPv4
IPv6

This agent will be configured to choose the IPv4 protocol for establishing all connections.

9. Select the features to install. Click Next to continue.

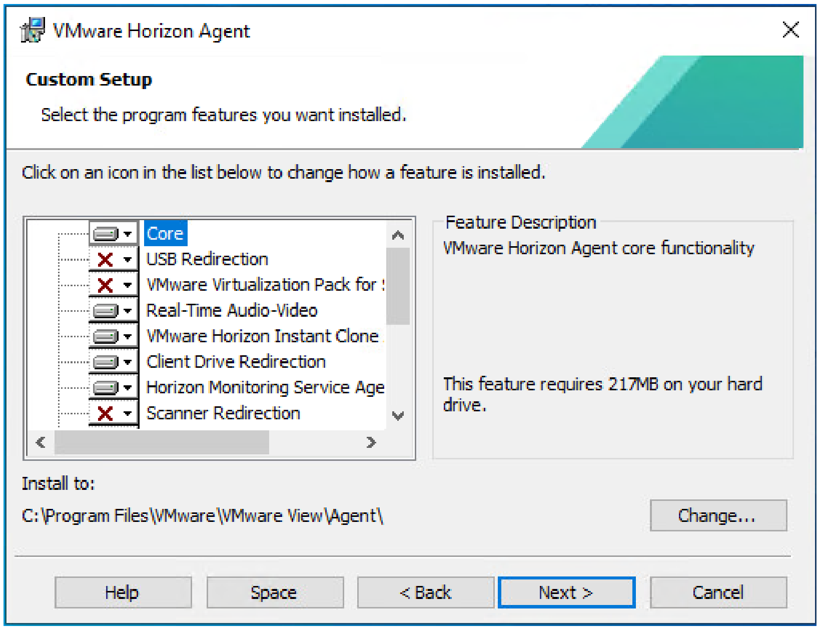

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- Core
- ☒ USB Redirection
- ☒ VMware Virtualization Pack for :
- Real-Time Audio-Video
- VMware Horizon Instant Clone
- Client Drive Redirection
- Horizon Monitoring Service Age
- ☒ Scanner Redirection

Feature Description

VMware Horizon Agent core functionality

This feature requires 217MB on your hard drive.

Install to:

C:\Program Files\VMware\VMware View\Agent\

10. Enable RDP and click Next to continue.

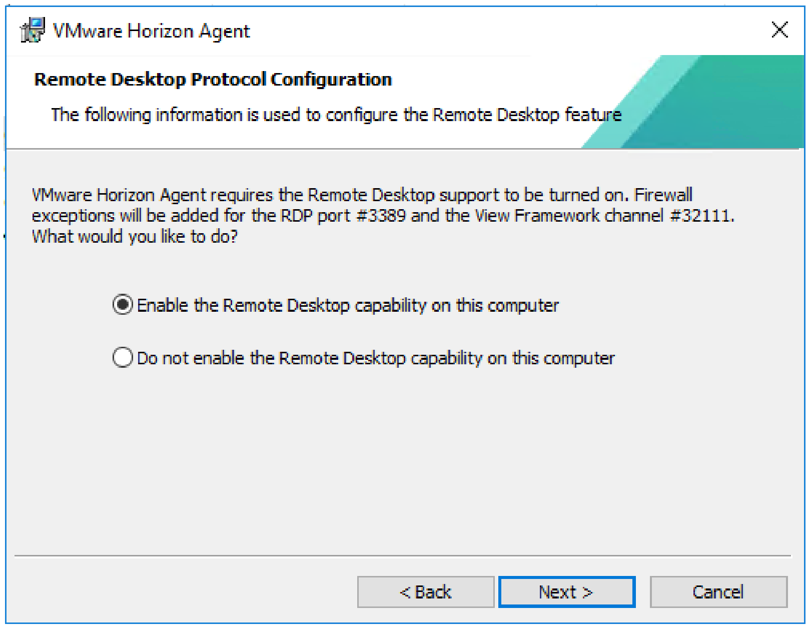

Remote Desktop Protocol Configuration

The following information is used to configure the Remote Desktop feature

VMware Horizon Agent requires the Remote Desktop support to be turned on. Firewall exceptions will be added for the RDP port #3389 and the View Framework channel #32111. What would you like to do?

● Enable the Remote Desktop capability on this computer

○ Do not enable the Remote Desktop capability on this computer

11. Click Next to begin Horizon Agent installation on the Master image.

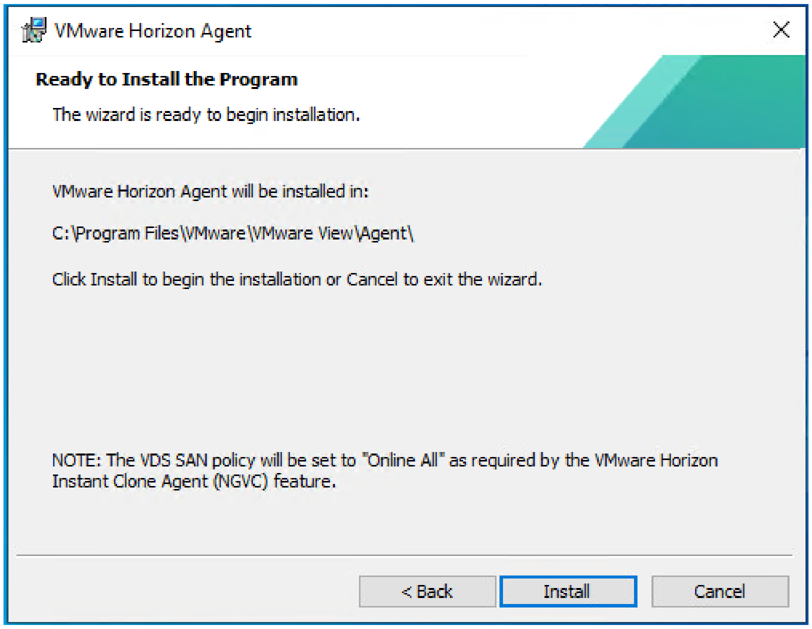

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Agent will be installed in:

C:\Program Files\VMware\VMware View\Agent\

Click Install to begin the installation or Cancel to exit the wizard.

NOTE: The VDS SAN policy will be set to "Online All" as required by the VMware Horizon Instant Clone Agent (NGVC) feature.

12. Click Finish to complete the Horizon Agent installation on the Master image.

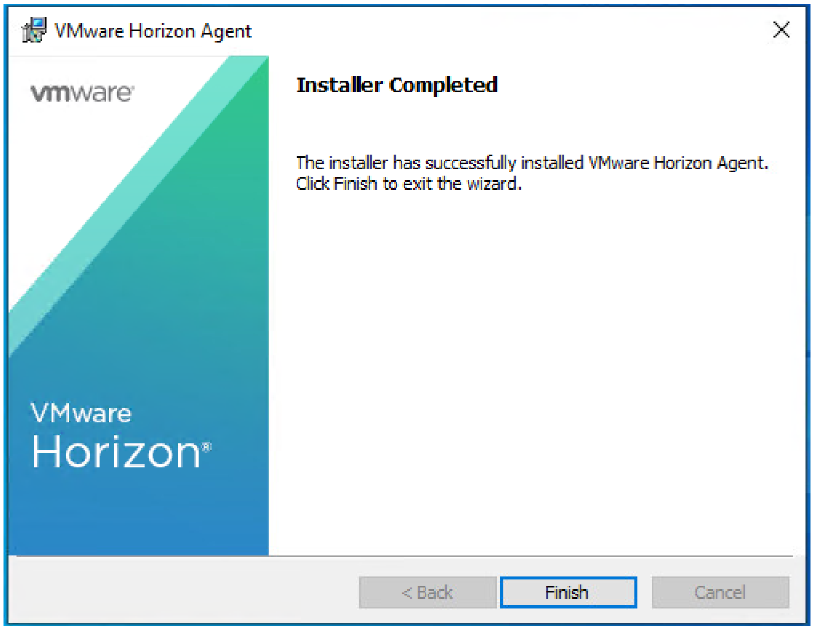

VMware Horizon Agent

Installer Completed

The installer has successfully installed VMware Horizon Agent. Click Finish to exit the wizard.

VMware Dynamic Environment Manager Enterprise (DEM) Setup

To install the VMware DEM on the master image, follow these steps:

1. Download VMware Dynamic Environment Manager Enterprise 10.0-GA version.
2. Click the VMware Dynamic Environment Manager installer for your OS architecture.

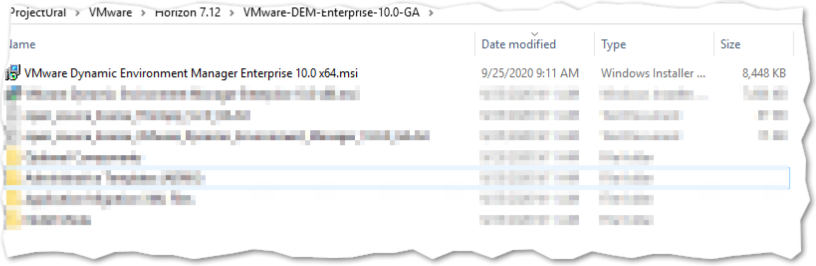

ProjectUral › VMware › Horizon 7.12 › VMware-DEM-Enterprise-10.0-GA ›

Name	Date modified	Type	Size
VMware Dynamic Environment Manager Enterprise 10.0 x64.msi	9/25/2020 9:11 AM	Windows Installer ...	8,448 KB

3. Click Next.

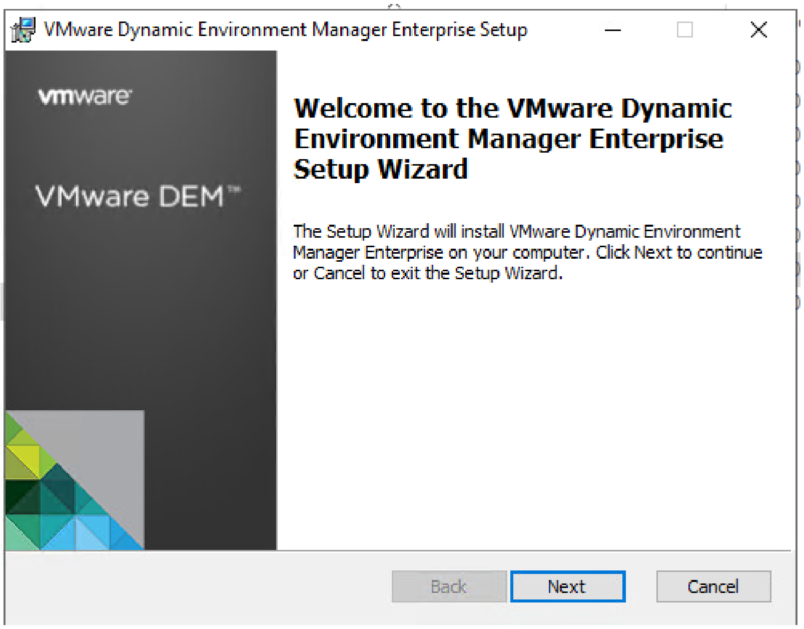

Welcome to the VMware Dynamic Environment Manager Enterprise Setup Wizard

The Setup Wizard will install VMware Dynamic Environment Manager Enterprise on your computer. Click Next to continue or Cancel to exit the Setup Wizard.

4. Read and accept the End User License Agreement and click Next.

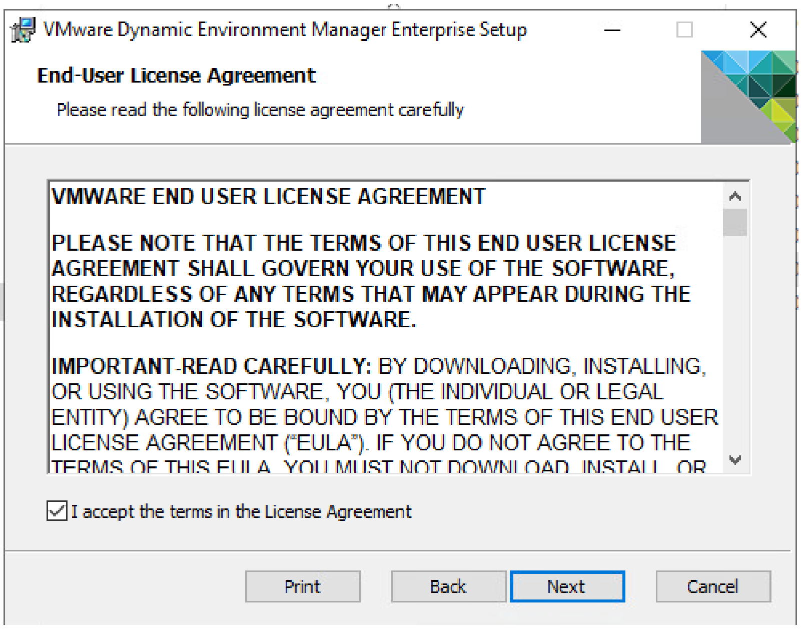

End-User License Agreement

Please read the following license agreement carefully

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR

☑ I accept the terms in the License Agreement

5. Select the destination folder where you want to install the application and click Next.

As a best practice, install VMware Dynamic Environment Manager in the default folder.

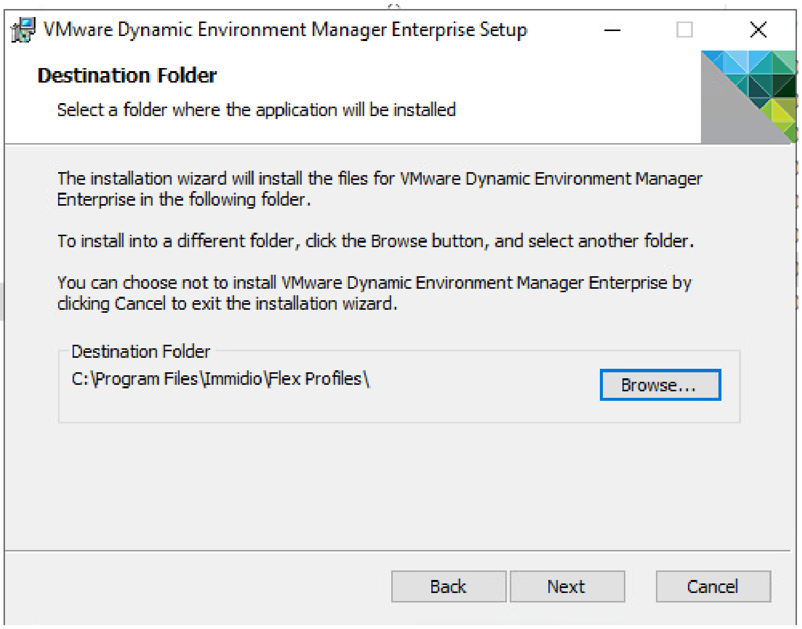

VMware Dynamic Environment Manager Enterprise Setup

Destination Folder

Select a folder where the application will be installed

The installation wizard will install the files for VMware Dynamic Environment Manager Enterprise in the following folder.

To install into a different folder, click the Browse button, and select another folder.

You can choose not to install VMware Dynamic Environment Manager Enterprise by clicking Cancel to exit the installation wizard.

Destination Folder
C:\Program Files\Immidio\Flex Profiles\

6. Select the Custom installation option for VMware Dynamic Environment Manager.

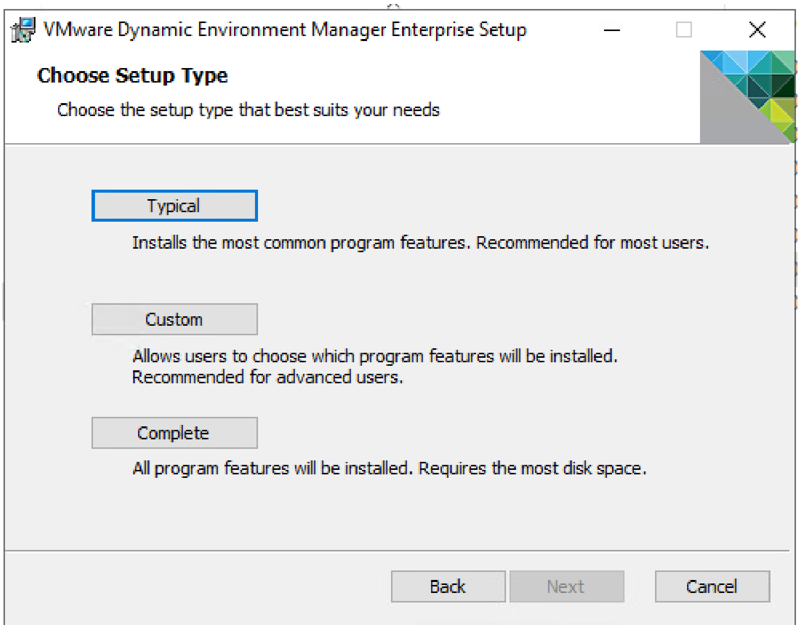

VMware Dynamic Environment Manager Enterprise Setup

Choose Setup Type

Choose the setup type that best suits your needs

Installs the most common program features. Recommended for most users.

Allows users to choose which program features will be installed. Recommended for advanced users.

All program features will be installed. Requires the most disk space.

7. Manually select components to install.

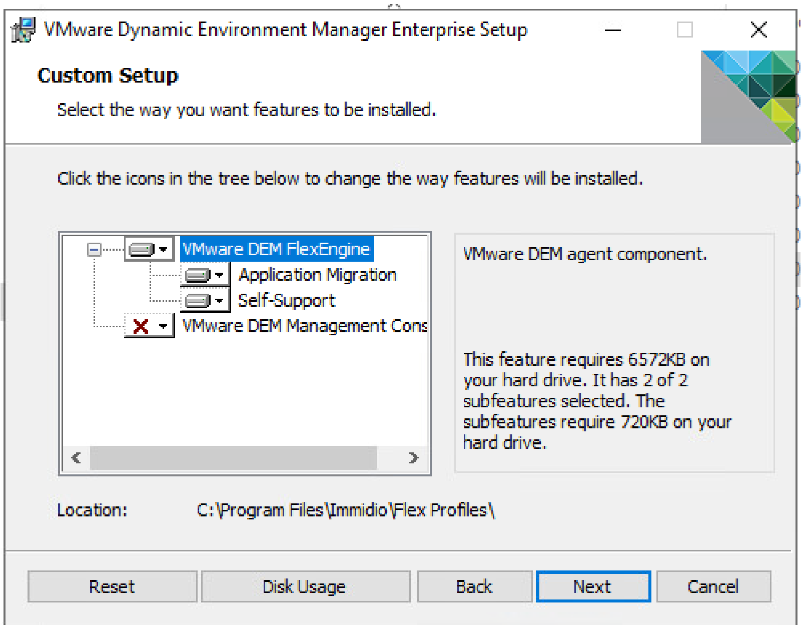

VMware Dynamic Environment Manager Enterprise Setup

Custom Setup

Select the way you want features to be installed.

Click the icons in the tree below to change the way features will be installed.

VMware DEM FlexEngine
- Application Migration
- Self-Support

VMware DEM Management Cons

VMware DEM agent component.

This feature requires 6572KB on your hard drive. It has 2 of 2 subfeatures selected. The subfeatures require 720KB on your hard drive.

Location: C:\Program Files\Immidio\Flex Profiles\

8. Click Next.

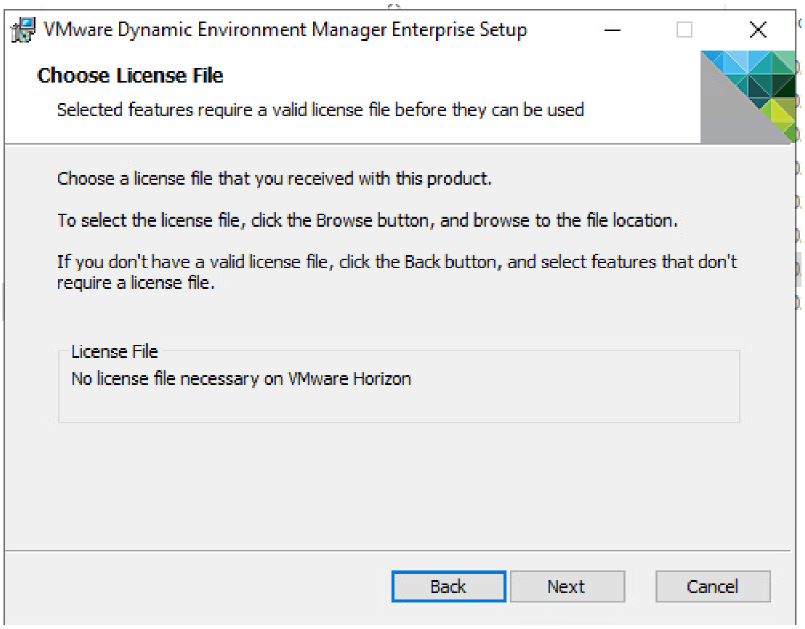

VMware Dynamic Environment Manager Enterprise Setup

Choose License File

Selected features require a valid license file before they can be used

Choose a license file that you received with this product.

To select the license file, click the Browse button, and browse to the file location.

If you don't have a valid license file, click the Back button, and select features that don't require a license file.

License File

No license file necessary on VMware Horizon

9. Click Install.

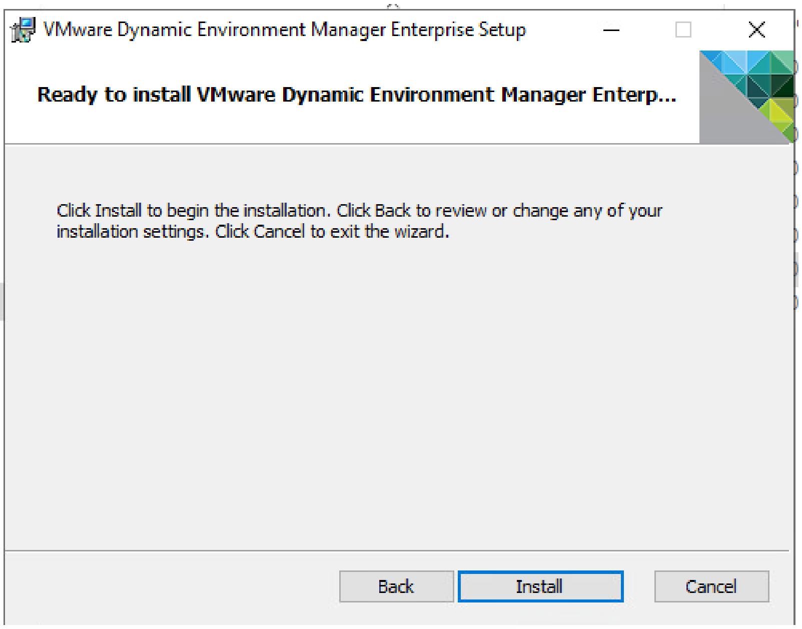
VMware Dynamic Environment Manager Enterprise Setup

Ready to install VMware Dynamic Environment Manager Enterp...

Click Install to begin the installation. Click Back to review or change any of your installation settings. Click Cancel to exit the wizard.

10. After the installation is complete, click Finish.

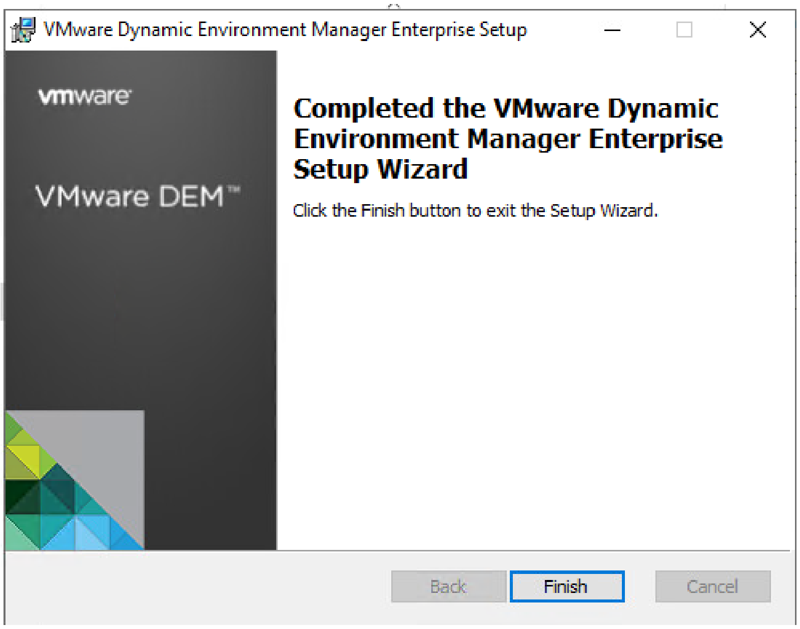
VMware Dynamic Environment Manager Enterprise Setup

Completed the VMware Dynamic Environment Manager Enterprise Setup Wizard

Click the Finish button to exit the Setup Wizard.

Install and Configure VMware Horizon Components

This section details the installation of the VMware Horizon core components. This CVD installs 4 VMware Horizon Connection 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 2012 Configuration Limits.

VMware Horizon Connection Server Configuration

To configure the VMware View Connection Server, follow these steps:

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 8.1.0 Version 2012. For the download, see Download VMware Horizon 8.

2. Click the Connection Server installer based on your Operating System.

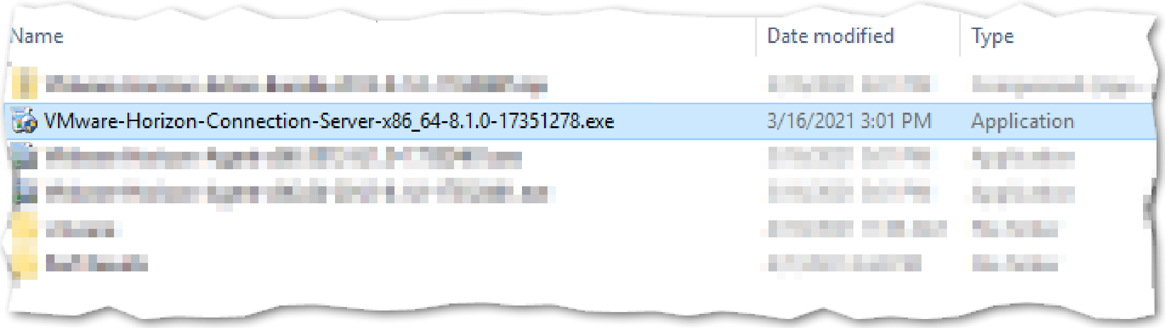

Name	Date modified	Type
[illegible]	[illegible]	[illegible]
VMware-Horizon-Connection-Server-x86_64-8.1.0-17351278.exe	3/16/2021 3:01 PM	Application
[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]

3. Click Next.

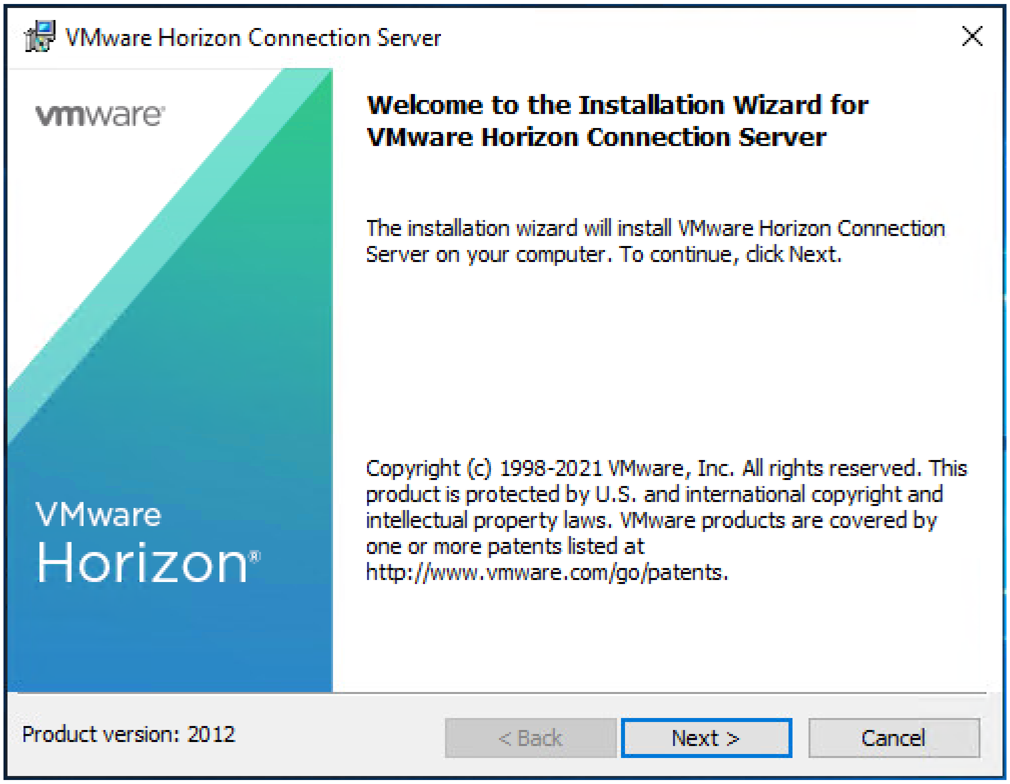

VMware Horizon Connection Server

Welcome to the Installation Wizard for VMware Horizon Connection Server

The installation wizard will install VMware Horizon Connection Server on your computer. To continue, click Next.

Copyright (c) 1998-2021 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 2012

4. Read and accept the End User License Agreement and click Next.

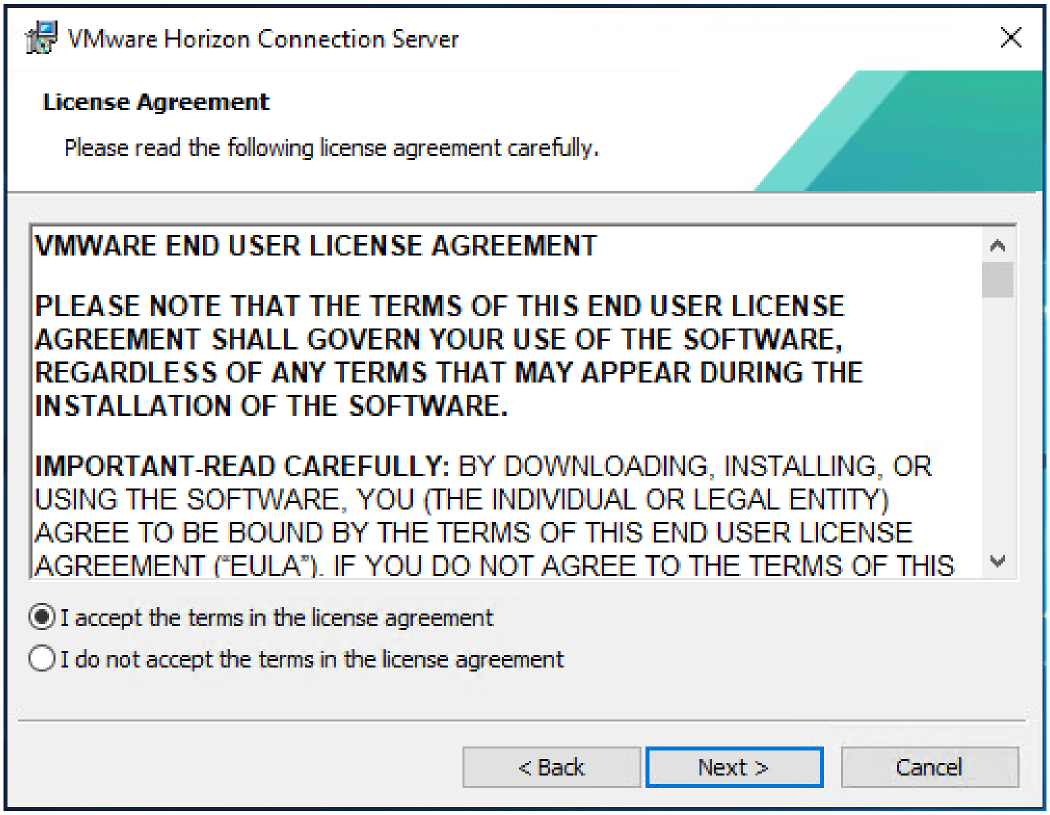

VMware Horizon Connection Server

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

5. Select the destination folder where you want to install the application and click Next.

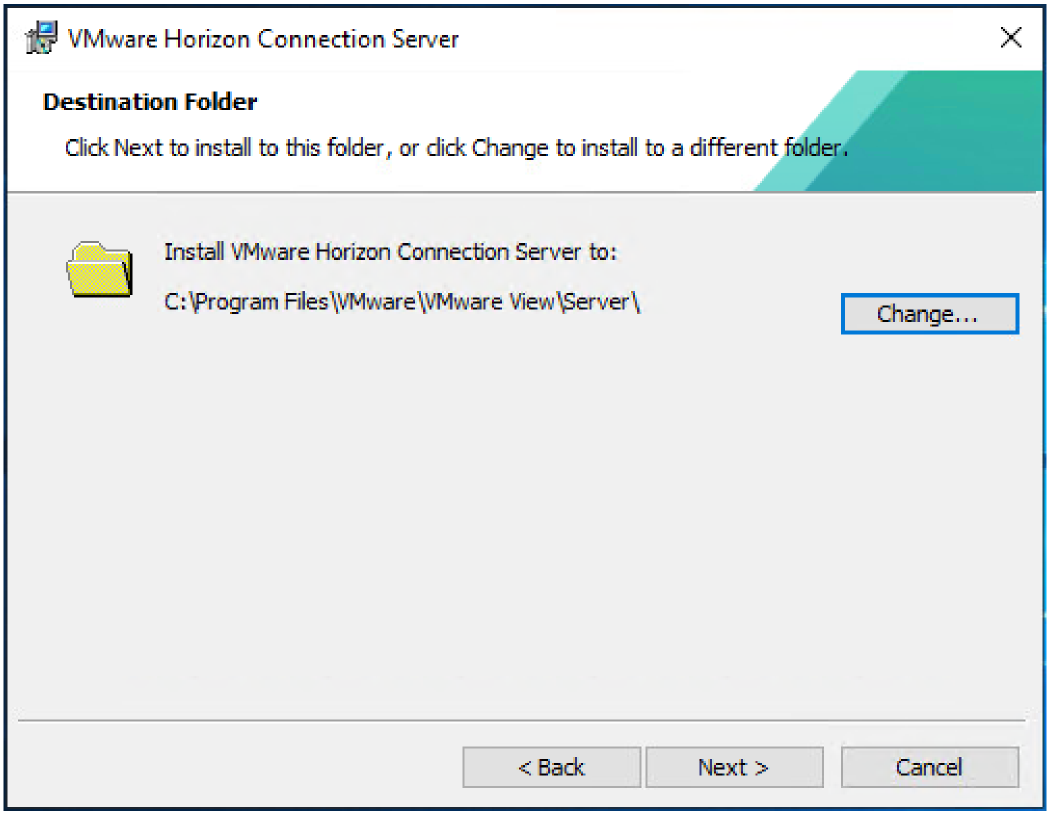

VMware Horizon Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware Horizon Connection Server to:

C:\Program Files\VMware\VMware View\Server\

6. Select the Standard Server and IPv4 for the IP protocol version.

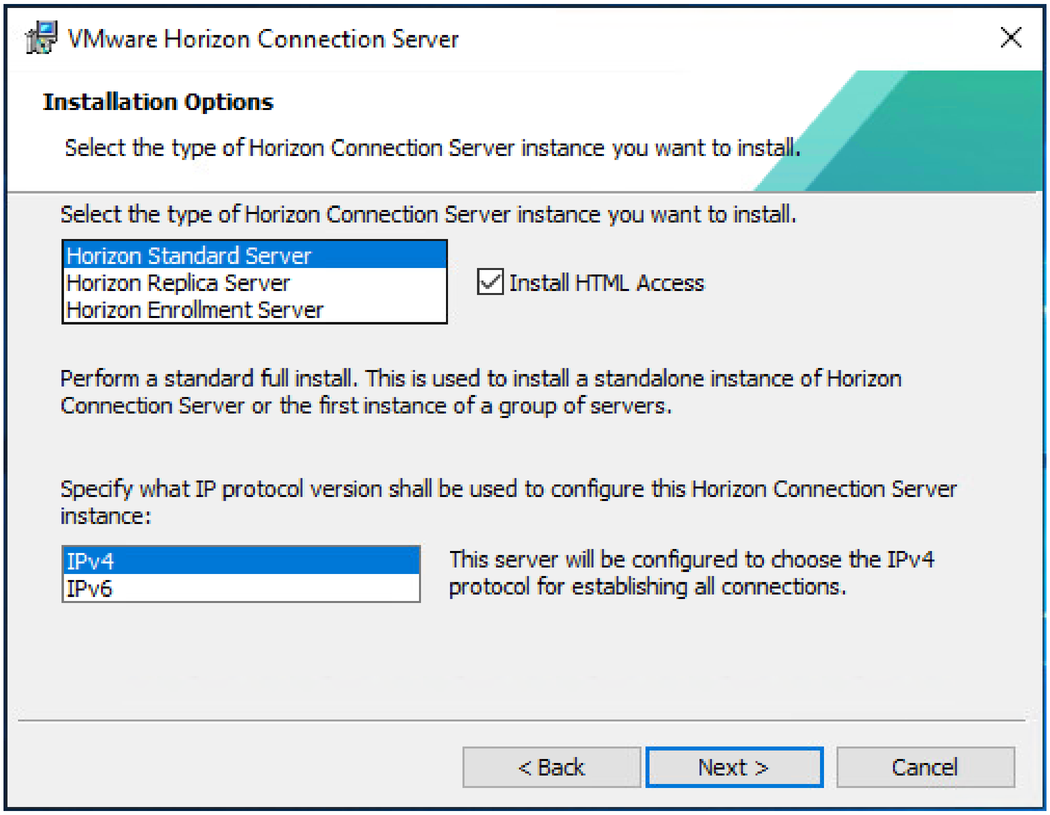

Installation Options

Select the type of Horizon Connection Server instance you want to install.

Select the type of Horizon Connection Server instance you want to install.

- Horizon Standard Server
- Horizon Replica Server
- Horizon Enrollment Server

☑ Install HTML Access

Perform a standard full install. This is used to install a standalone instance of Horizon Connection Server or the first instance of a group of servers.

Specify what IP protocol version shall be used to configure this Horizon Connection Server instance:

- IPv4
- IPv6

This server will be configured to choose the IPv4 protocol for establishing all connections.

7. Provide the Data Recovery password. Click Next.

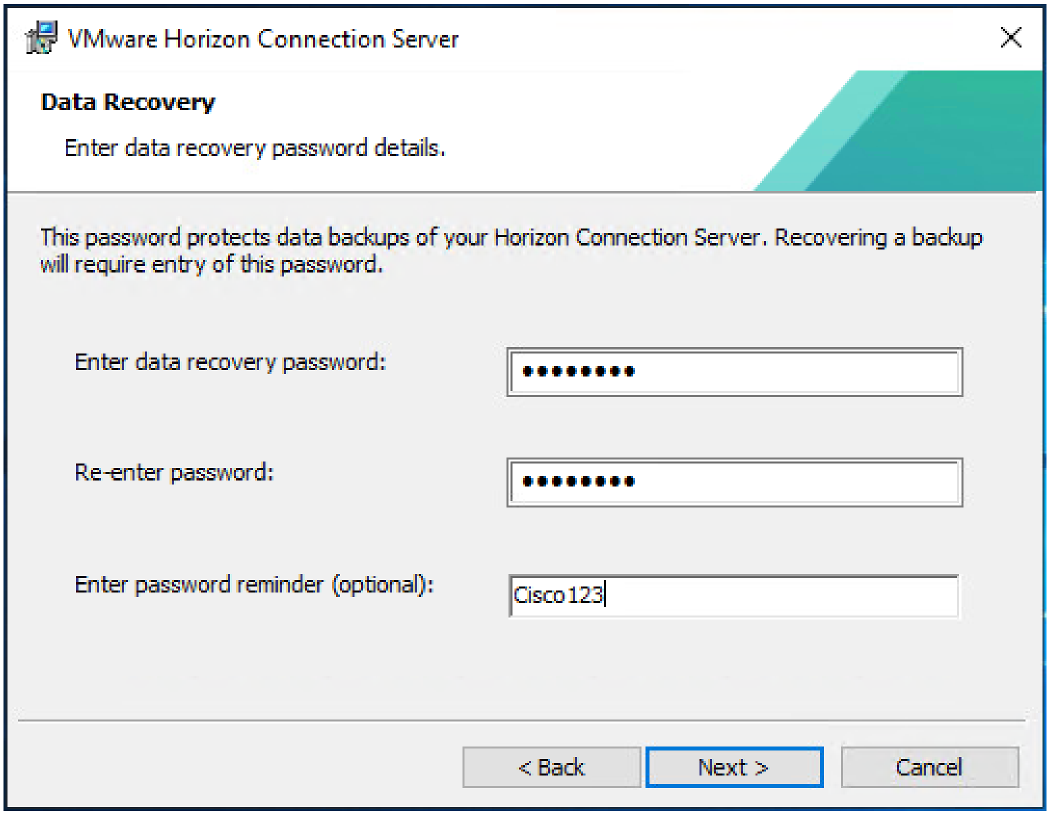

VMware Horizon Connection Server

Data Recovery

Enter data recovery password details.

This password protects data backups of your Horizon Connection Server. Recovering a backup will require entry of this password.

Enter data recovery password: ●●●●●●●●●

Re-enter password: ●●●●●●●●●

Enter password reminder (optional): Cisco123

8. Select Configure Windows Firewall automatically. Click Next.

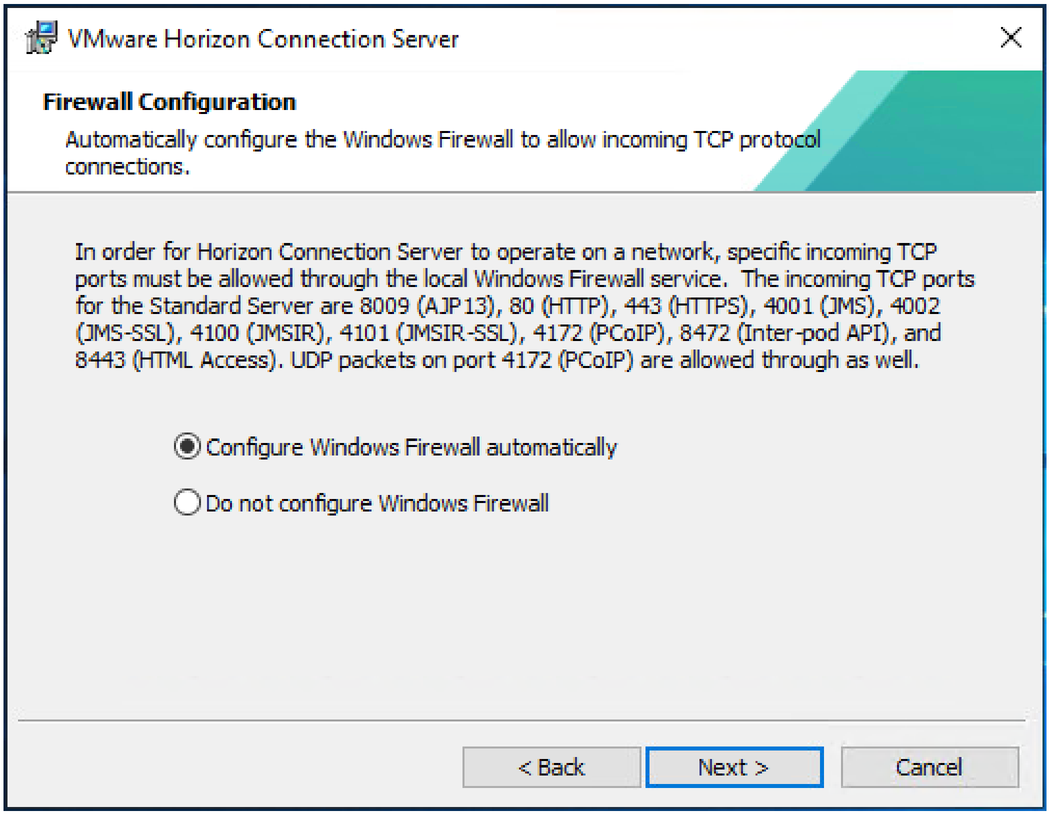

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for Horizon Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Standard Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), 4002 (JMS-SSL), 4100 (JMSIR), 4101 (JMSIR-SSL), 4172 (PCoIP), 8472 (Inter-pod API), and 8443 (HTML Access). UDP packets on port 4172 (PCoIP) are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

9. Authorize Domain Admins to be VMware Horizon administrators.

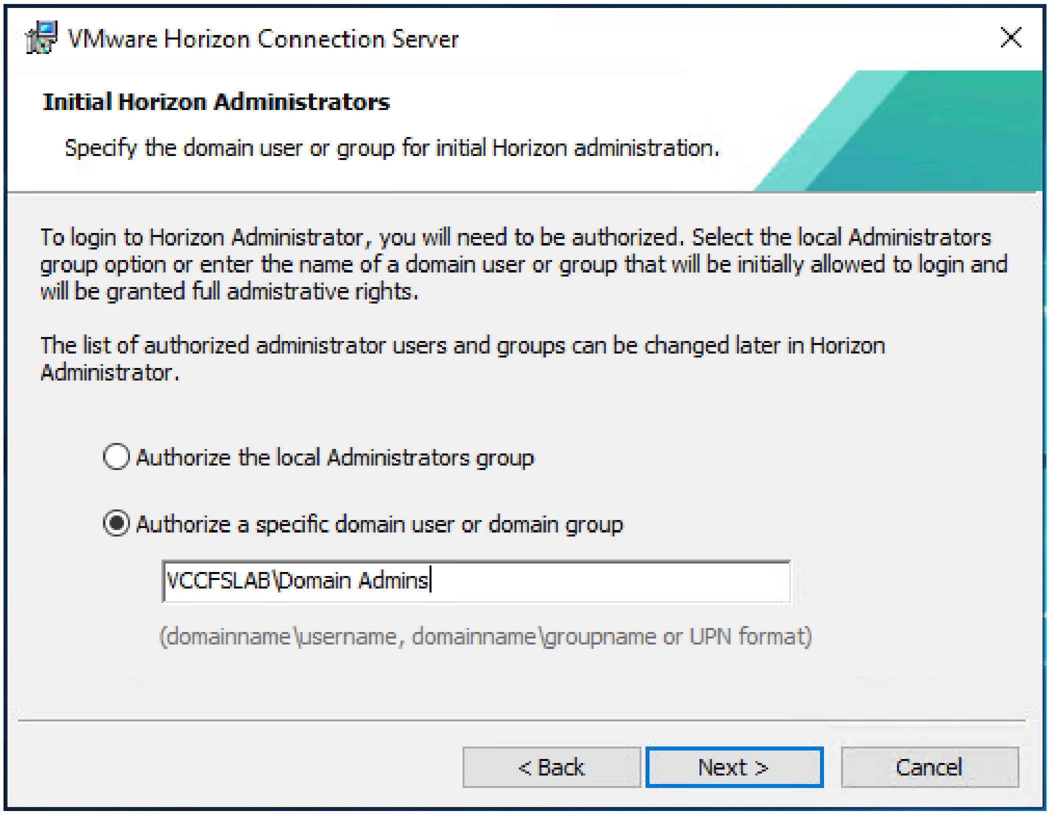

VMware Horizon Connection Server

Initial Horizon Administrators

Specify the domain user or group for initial Horizon administration.

To login to Horizon Administrator, you will need to be authorized. Select the local Administrators group option or enter the name of a domain user or group that will be initially allowed to login and will be granted full admistrative rights.

The list of authorized administrator users and groups can be changed later in Horizon Administrator.

○ Authorize the local Administrators group

● Authorize a specific domain user or domain group

VCCFSLAB\Domain Admins

(domainname\username, domainname\groupname or UPN format)

10. (Optional) Join Customer Experience Program.

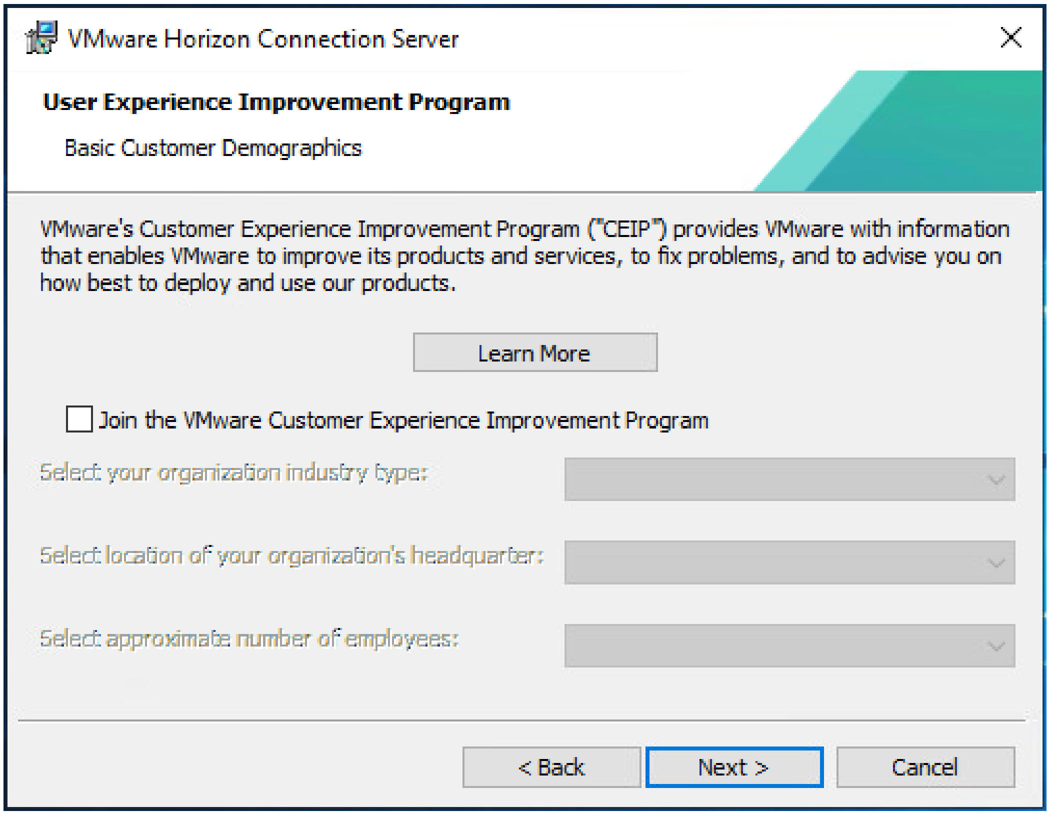

VMware Horizon Connection Server

User Experience Improvement Program

Basic Customer Demographics

VMware's Customer Experience Improvement Program ("CEIP") provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products.

□ Join the VMware Customer Experience Improvement Program

Select your organization industry type:

Select location of your organization's headquarter:

Select approximate number of employees:

11. Select General from drop-down list. Click Install to begin installation.

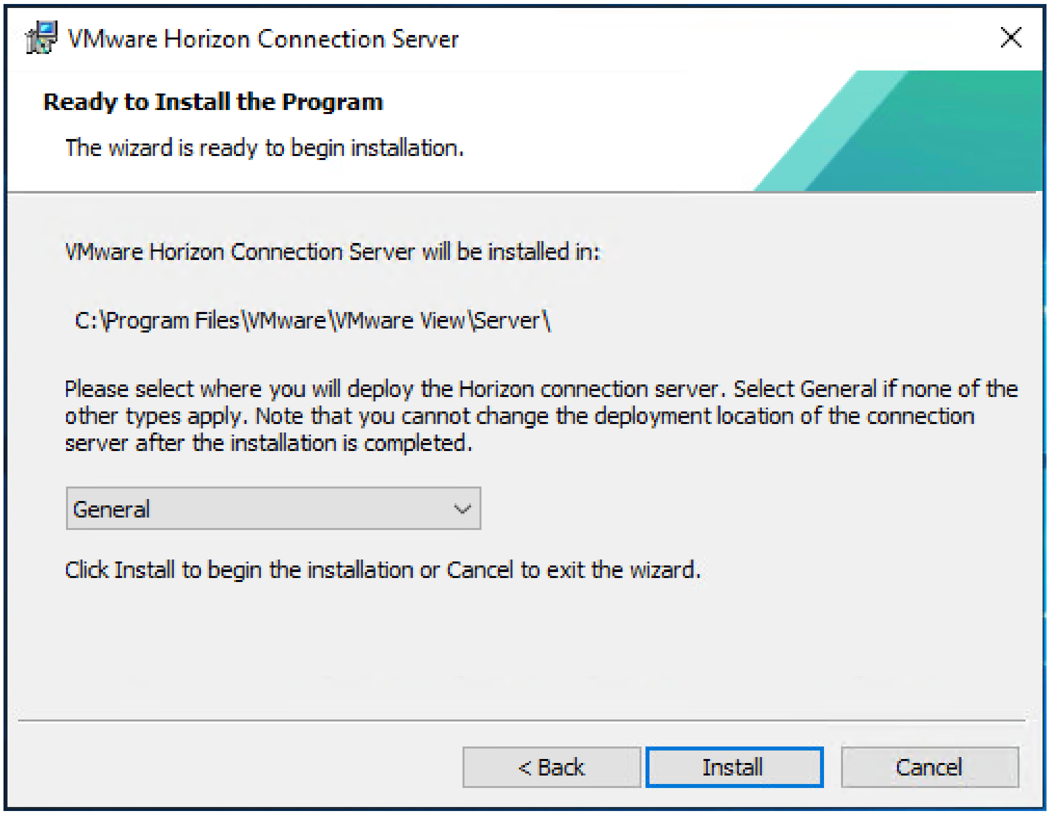

VMware Horizon Connection Server

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Connection Server will be installed in:

C:\Program Files\VMware\VMware View\Server\

Please select where you will deploy the Horizon connection server. Select General if none of the other types apply. Note that you cannot change the deployment location of the connection server after the installation is completed.

General

Click Install to begin the installation or Cancel to exit the wizard.

12. After Horizon Connection Server installation is complete, click Finish.

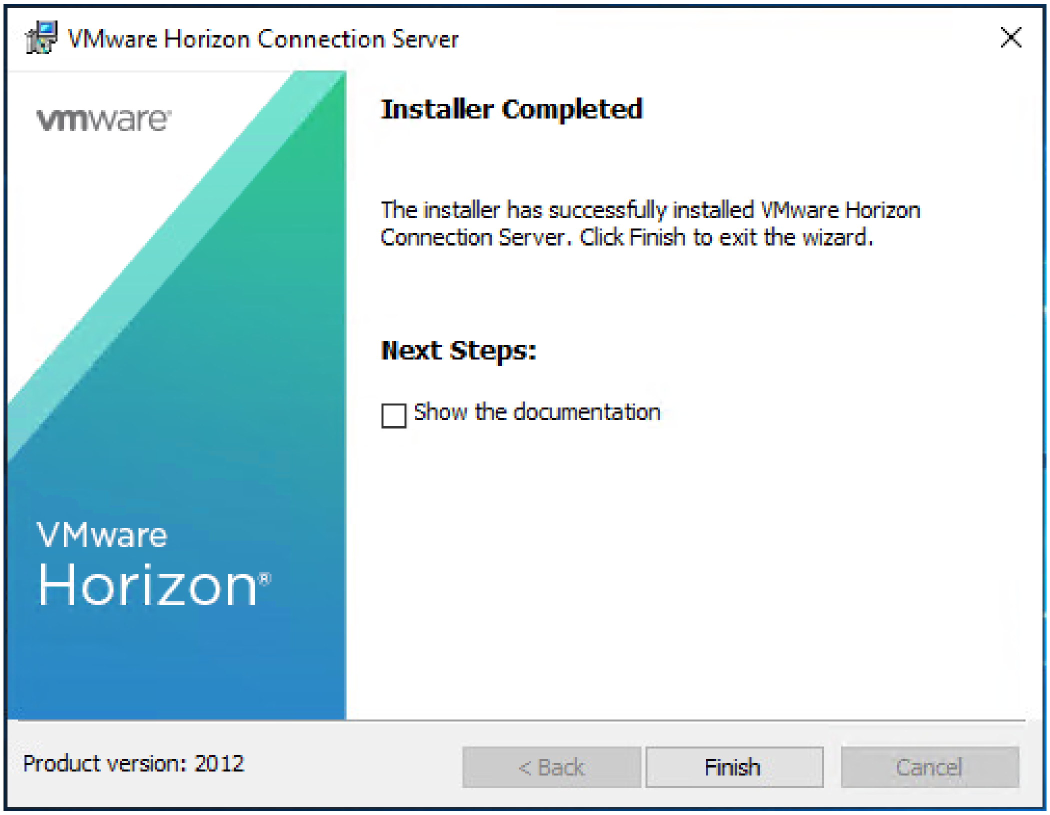

VMware Horizon Connection Server

Installer Completed

The installer has successfully installed VMware Horizon Connection Server. Click Finish to exit the wizard.

Next Steps:

□ Show the documentation

Product version: 2012

Install VMware Horizon Replica Server

To install the VMware Horizon Replica Server, follow these steps:

1. Click the Connection Server installer based on your Operating System.

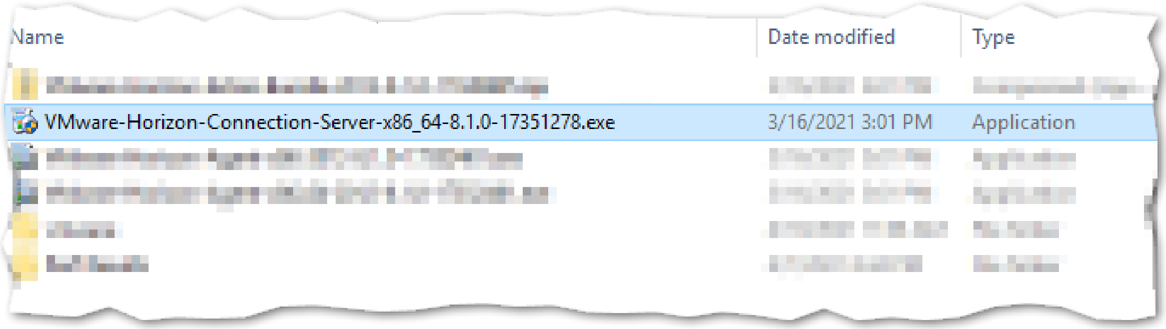

Name	Date modified	Type
VMware-Horizon-Connection-Server-x86_64-8.1.0-17351278.exe	3/16/2021 3:01 PM	Application

2. Click Next.

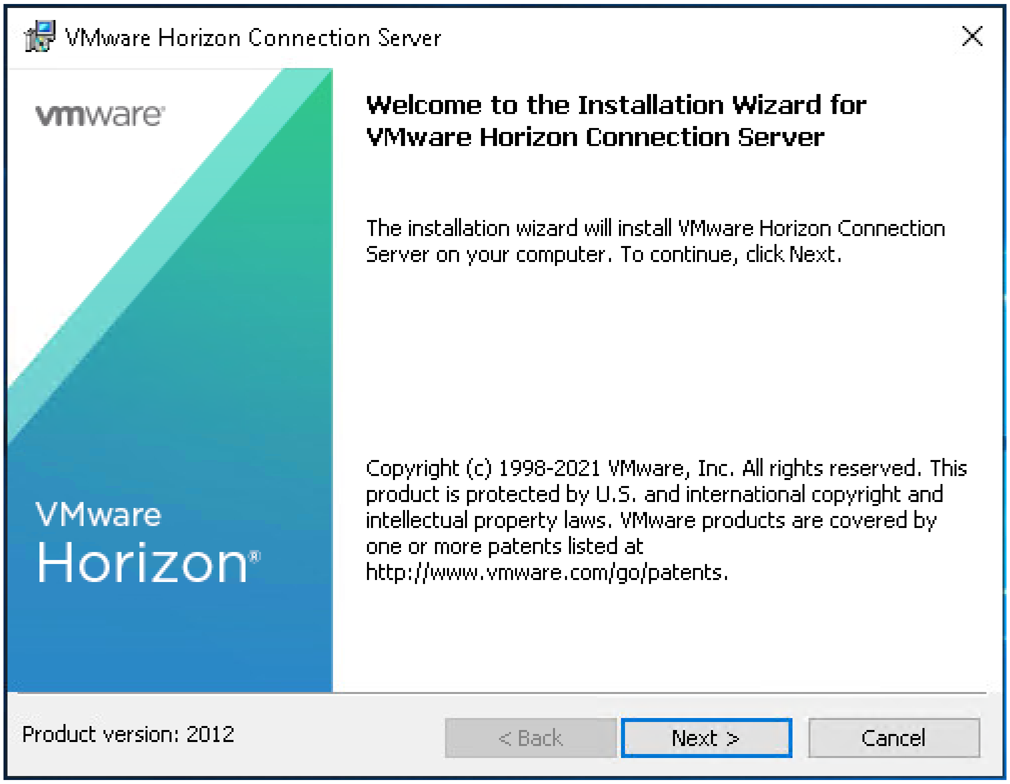

VMware Horizon Connection Server

Welcome to the Installation Wizard for VMware Horizon Connection Server

The installation wizard will install VMware Horizon Connection Server on your computer. To continue, click Next.

Copyright (c) 1998-2021 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 2012

3. Read and accept the End User License Agreement and click Next.

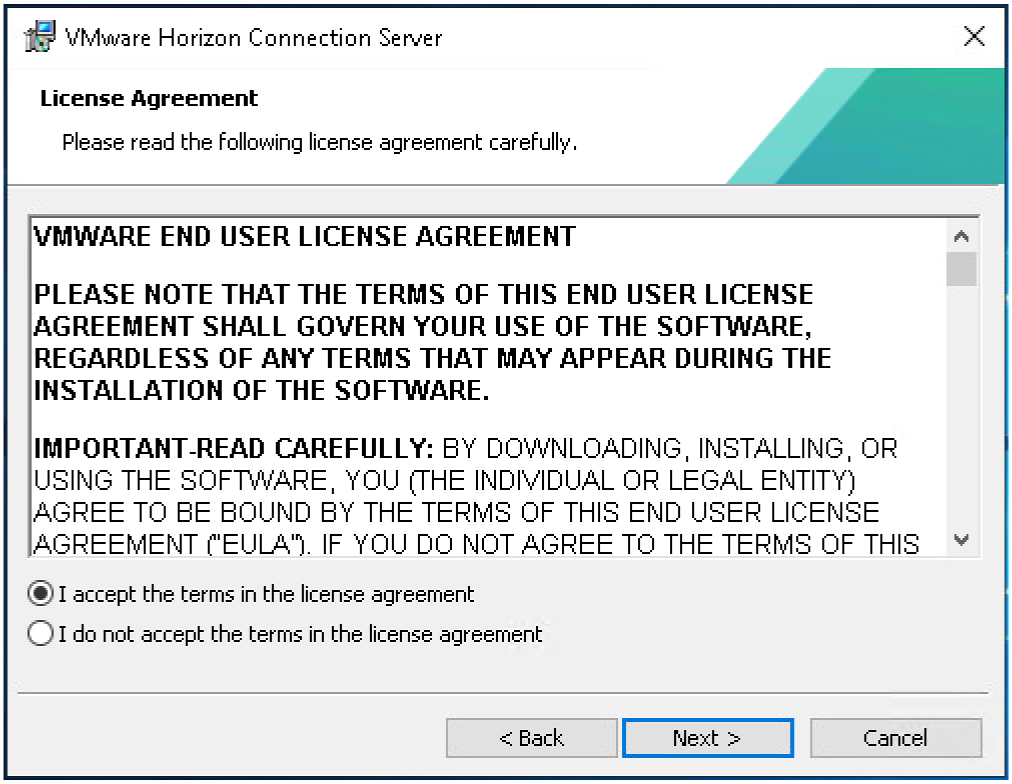

VMware Horizon Connection Server

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

PLEASE NOTE THAT THE TERMS OF THIS END USER LICENSE AGREEMENT SHALL GOVERN YOUR USE OF THE SOFTWARE, REGARDLESS OF ANY TERMS THAT MAY APPEAR DURING THE INSTALLATION OF THE SOFTWARE.

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

4. Select the destination folder where you want to install the application and click Next.

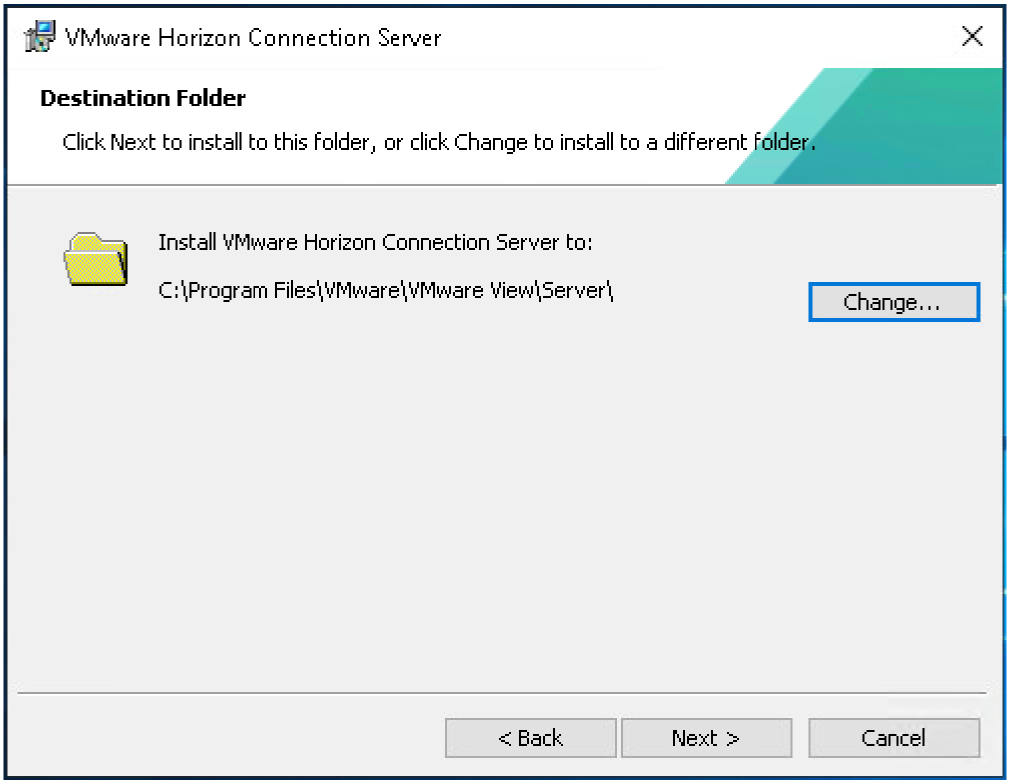

VMware Horizon Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware Horizon Connection Server to:

C:\Program Files\VMware\VMware View\Server\

5. Select the Replica Server and IPv4 for the IP protocol version.

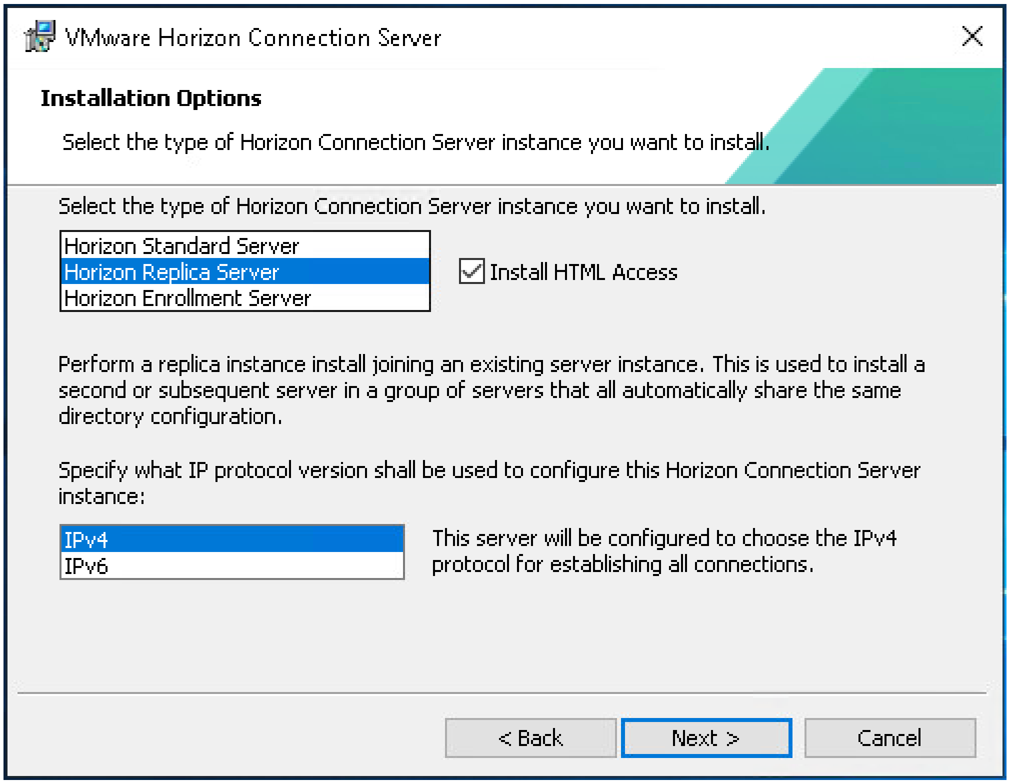

VMware Horizon Connection Server

Installation Options

Select the type of Horizon Connection Server instance you want to install.

Select the type of Horizon Connection Server instance you want to install.

- Horizon Standard Server
- Horizon Replica Server
- Horizon Enrollment Server

☑ Install HTML Access

Perform a replica instance install joining an existing server instance. This is used to install a second or subsequent server in a group of servers that all automatically share the same directory configuration.

Specify what IP protocol version shall be used to configure this Horizon Connection Server instance:

- IPv4
- IPv6

This server will be configured to choose the IPv4 protocol for establishing all connections.

6. Provide existing Standard View Connection Server's FQDN or IP address and click Next.

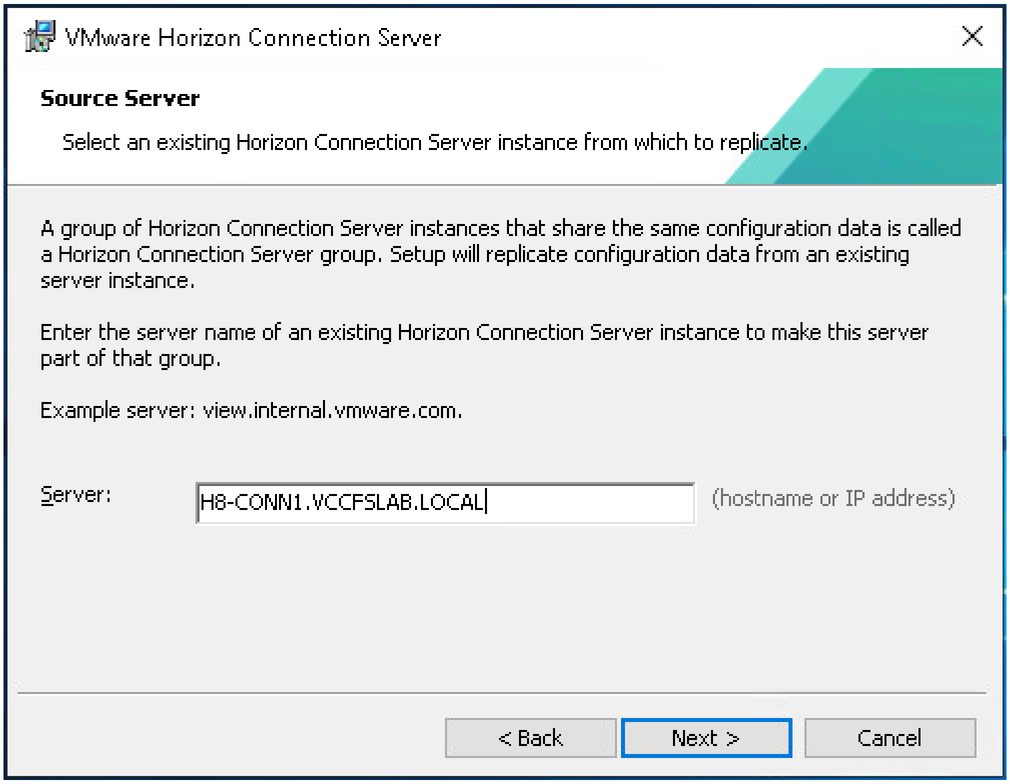

VMware Horizon Connection Server

Source Server

Select an existing Horizon Connection Server instance from which to replicate.

A group of Horizon Connection Server instances that share the same configuration data is called a Horizon Connection Server group. Setup will replicate configuration data from an existing server instance.

Enter the server name of an existing Horizon Connection Server instance to make this server part of that group.

Example server: view.internal.vmware.com.

Server: H8-CONN1.VCCFSLAB.LOCAL (hostname or IP address)

7. Select Configure the Windows Firewall automatically.

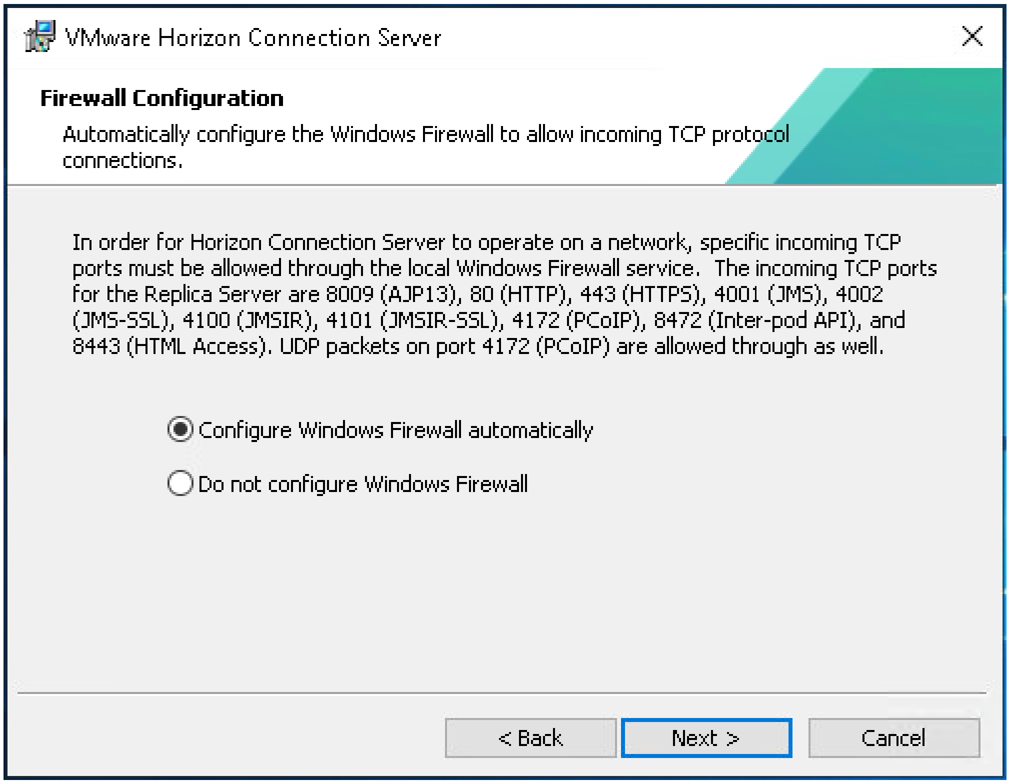

VMware Horizon Connection Server

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for Horizon Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Replica Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), 4002 (JMS-SSL), 4100 (JMSIR), 4101 (JMSIR-SSL), 4172 (PCoIP), 8472 (Inter-pod API), and 8443 (HTML Access). UDP packets on port 4172 (PCoIP) are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

8. Click Install to begin the installation process.

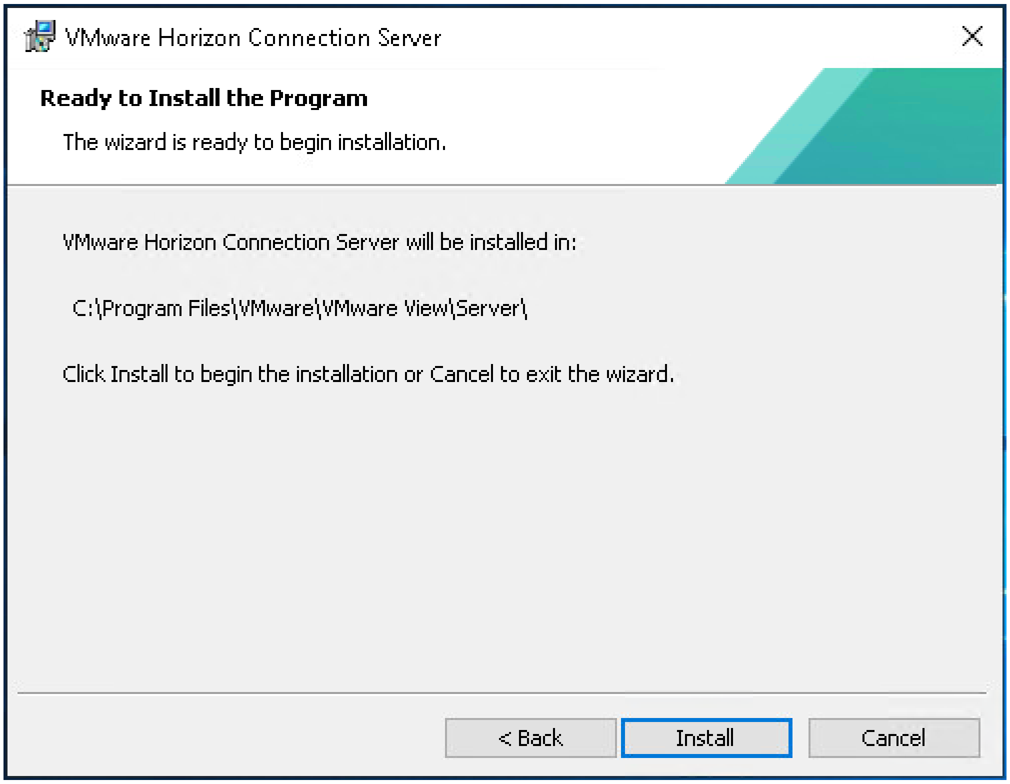
VMware Horizon Connection Server

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Connection Server will be installed in:

C:\Program Files\VMware\VMware View\Server\

Click Install to begin the installation or Cancel to exit the wizard.

9. After installation is complete, click Finish.

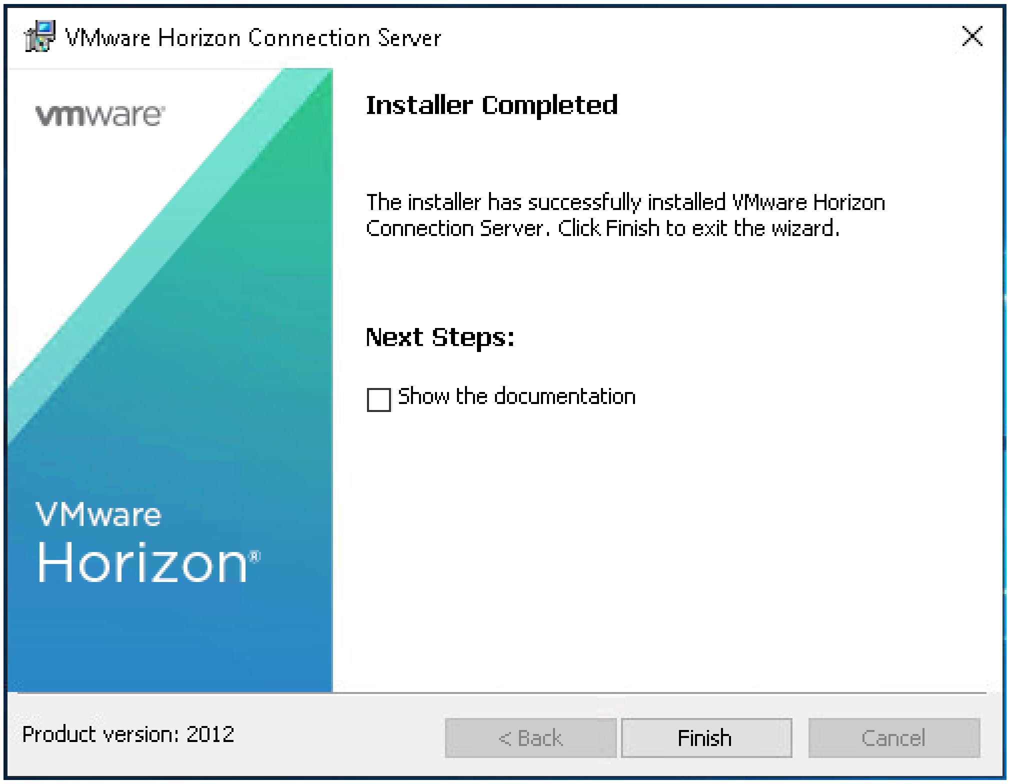
VMware Horizon Connection Server

Installer Completed

The installer has successfully installed VMware Horizon Connection Server. Click Finish to exit the wizard.

Next Steps:

□ Show the documentation

Product version: 2012

VMware Horizon Desktop Configuration

Management of the desktops, application pools and farms is accomplished in VMware Horizon Console (HTML5) or Horizon Administrator (Flex). We used Horizon Console to administer VMware Horizon environment in this validated design.

> VMware recommends using Horizon Console, an HTML5 based interface with enhanced security, capabilities, and performance.

To create the VMware Horizon Desktop Initial configuration, follow these steps:

1. Log into Horizon Console via a web browser using Address or FQDN>/admin/#/login.

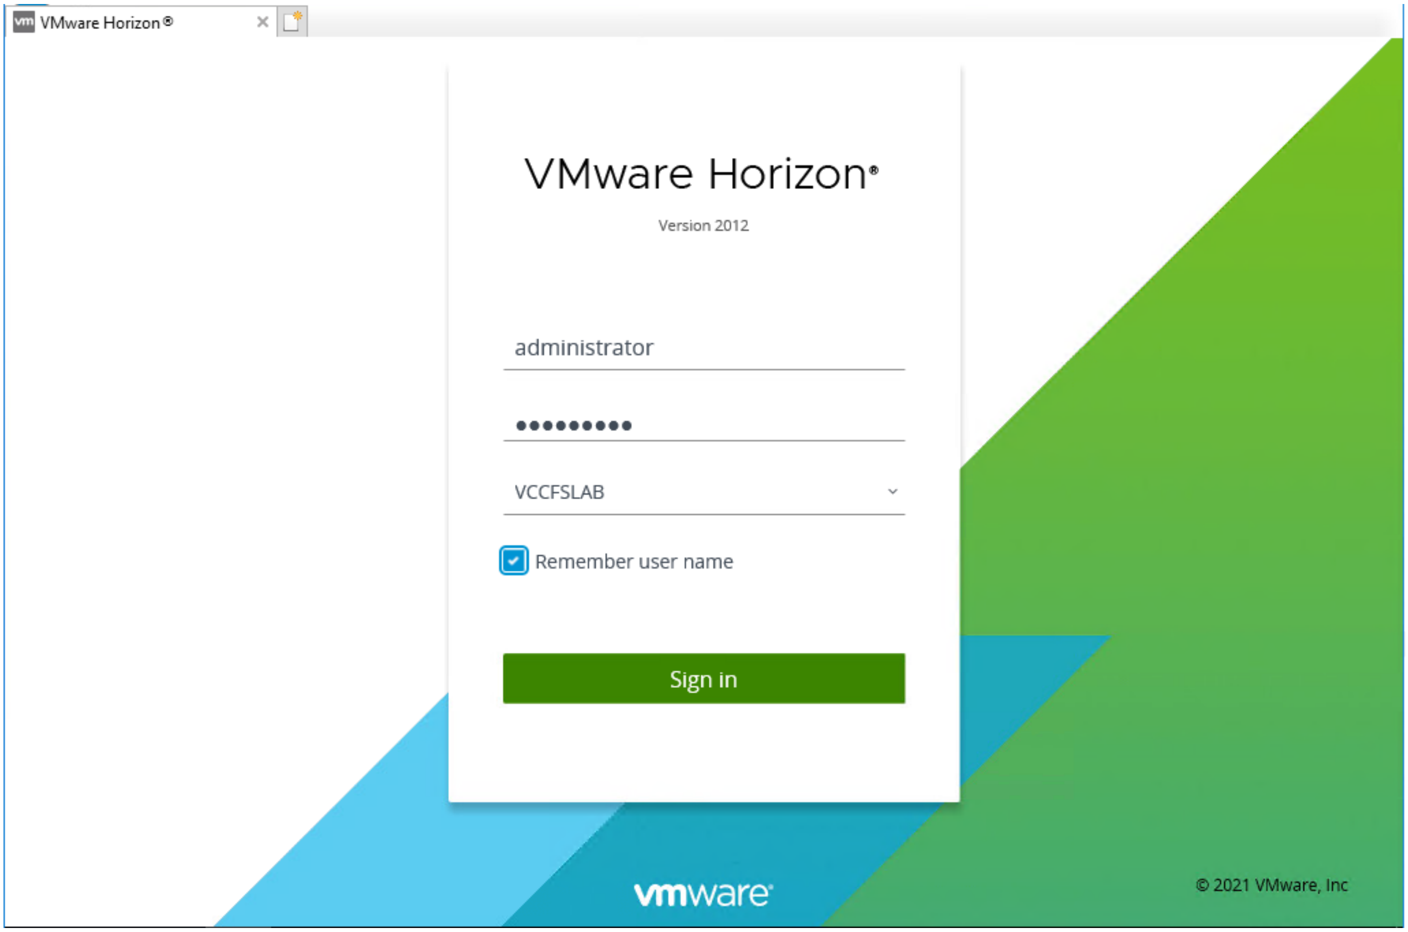

VMware Horizon®

Version 2012

administrator

●●●●●●●●●●

VCCFSLAB

☑ Remember user name

2. In Horizon Console, expand Settings and click Servers.

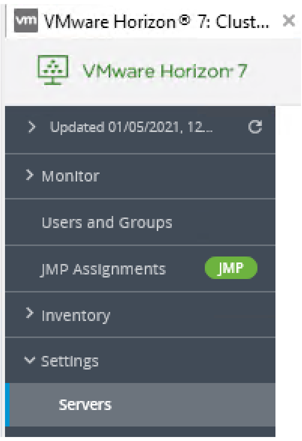

3. Select the vCenter Servers tab and click Add.

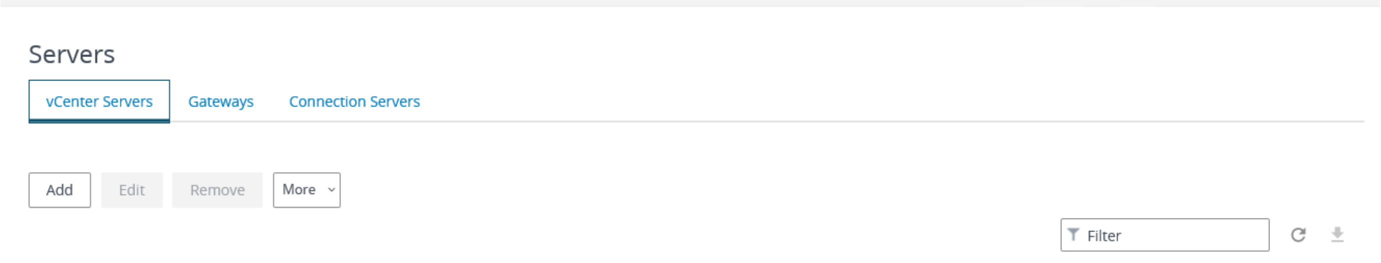

4. Provide Server Address (IP or FQDN) and credentials that Horizon will use to login to vCenter, then click Next.

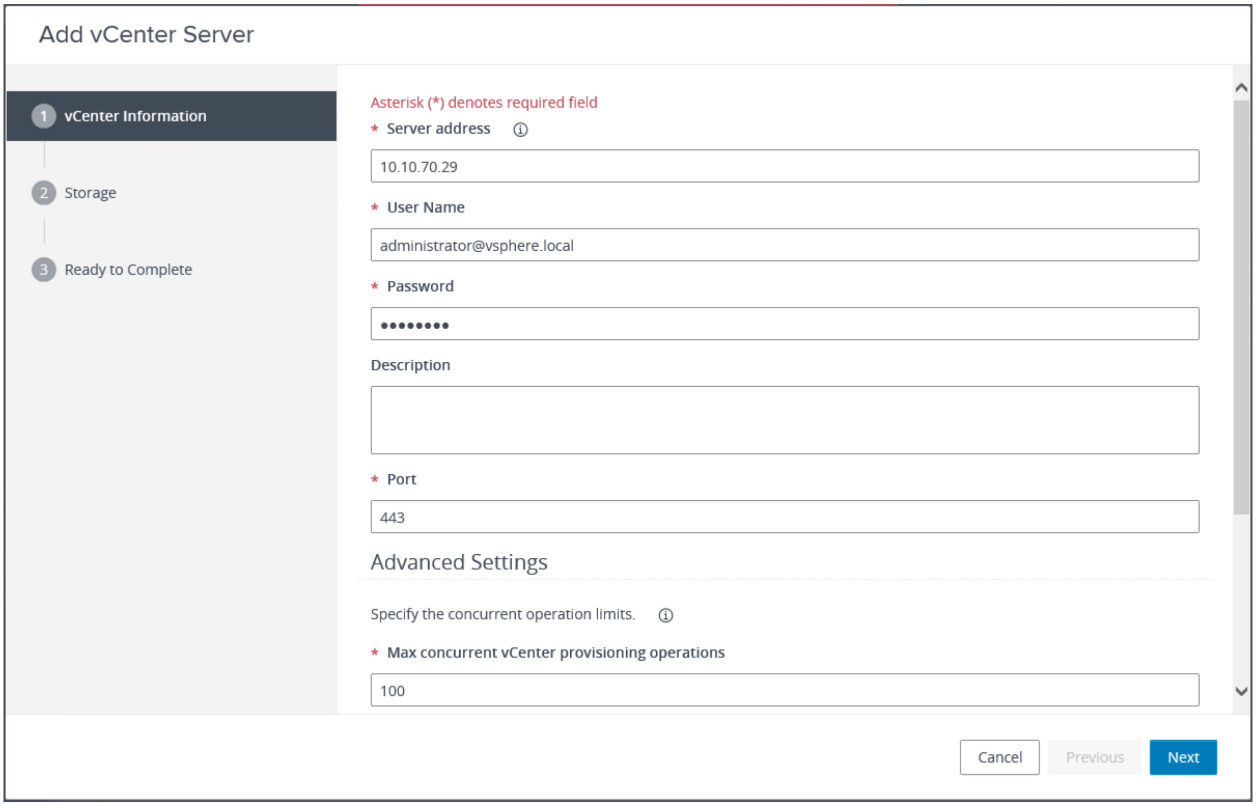

Add vCenter Server

Asterisk (*) denotes required field

* Server address

10.10.70.29

* User Name

administrator@vsphere.local

* Password

••••••••

Description

* Port

443

Advanced Settings

Specify the concurrent operation limits.

* Max concurrent vCenter provisioning operations

100

5. If you see a message regarding an invalid certificate, click View Certificate.

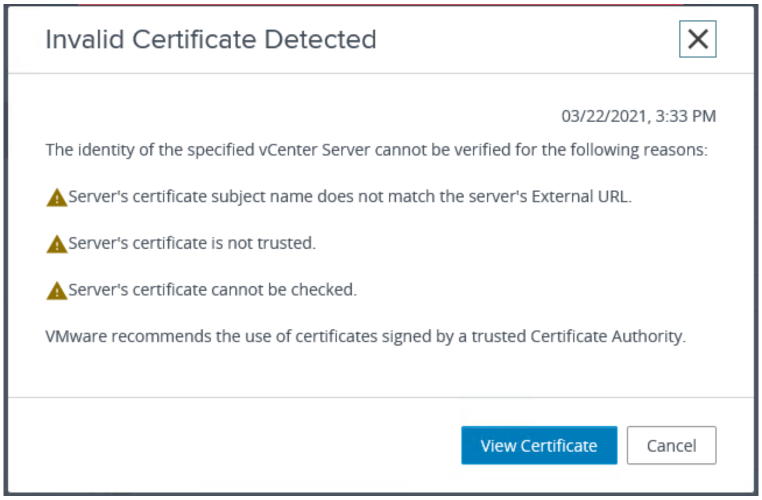

Invalid Certificate Detected

03/22/2021, 3:33 PM

The identity of the specified vCenter Server cannot be verified for the following reasons:

⚠Server's certificate subject name does not match the server's External URL.

⚠Server's certificate is not trusted.

⚠Server's certificate cannot be checked.

VMware recommends the use of certificates signed by a trusted Certificate Authority.

6. Click Accept.

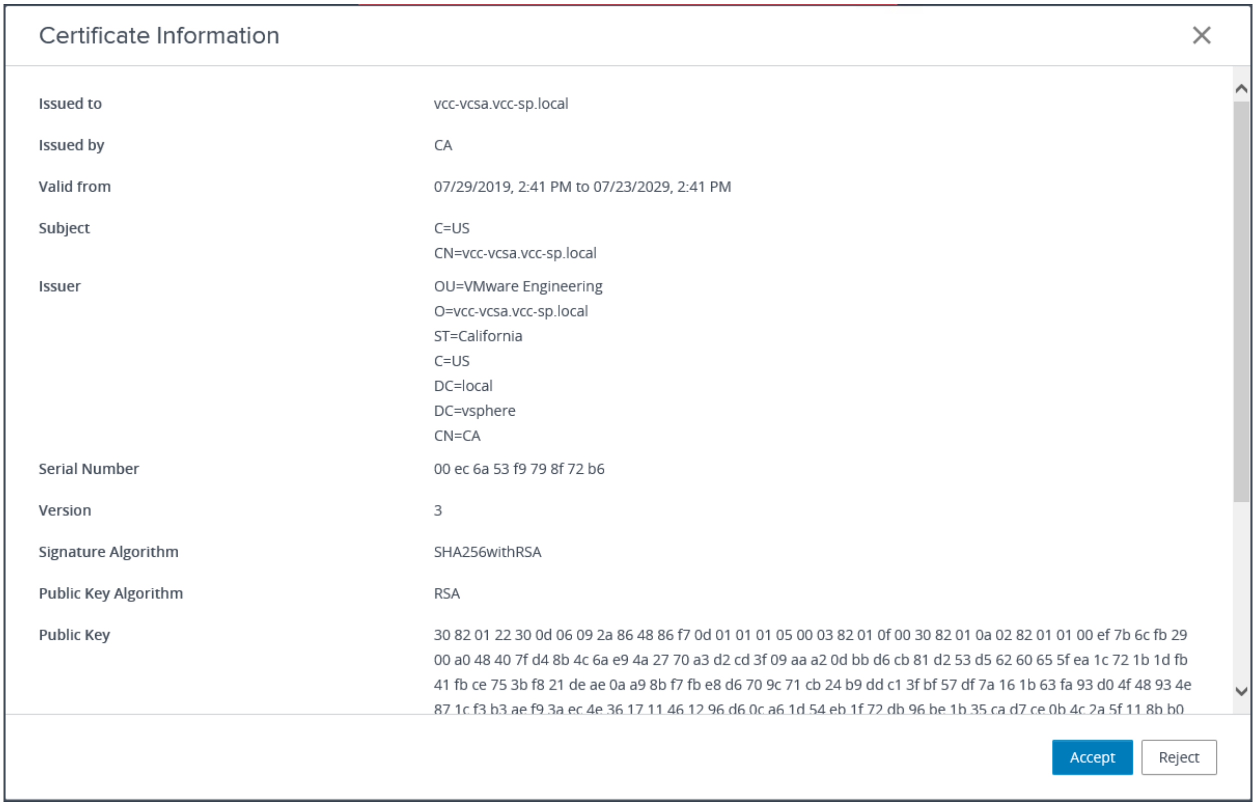

Certificate Information

Issued to	vcc-vcsa.vcc-sp.local
Issued by	CA
Valid from	07/29/2019, 2:41 PM to 07/23/2029, 2:41 PM

Subject	C=US
	CN=vcc-vcsa.vcc-sp.local

Issuer	OU=VMware Engineering
	O=vcc-vcsa.vcc-sp.local
	ST=California
	C=US
	DC=local
	DC=vsphere
	CN=CA

Serial Number	00 ec 6a 53 f9 79 8f 72 b6
Version	3
Signature Algorithm	SHA256withRSA
Public Key Algorithm	RSA
Public Key	30 82 01 22 30 0d 06 09 2a 86 48 86 f7 0d 01 01 01 05 00 03 82 01 0f 00 30 82 01 0a 02 82 01 01 00 ef 7b 6c fb 29 00 a0 48 40 7f d4 8b 4c 6a e9 4a 27 70 a3 d2 cd 3f 09 aa a2 0d bb d6 cb 81 d2 53 d5 62 60 65 5f ea 1c 72 1b 1d fb 41 fb ce 75 3b f8 21 de ae 0a a9 8b f7 fb e8 d6 70 9c 71 cb 24 b9 dd c1 3f bf 57 df 7a 16 1b 63 fa 93 d0 4f 48 93 4e 87 1c f3 b3 ae f9 3a ec 4e 36 17 11 46 12 96 d6 0c a6 1d 54 eb 1f 72 db 96 be 1b 35 ca d7 ce 0b 4c 2a 5f 11 8b b0

7. Select Reclaim VM disk space and Enable Horizon Storage Accelerator with cache size of 1024MB.

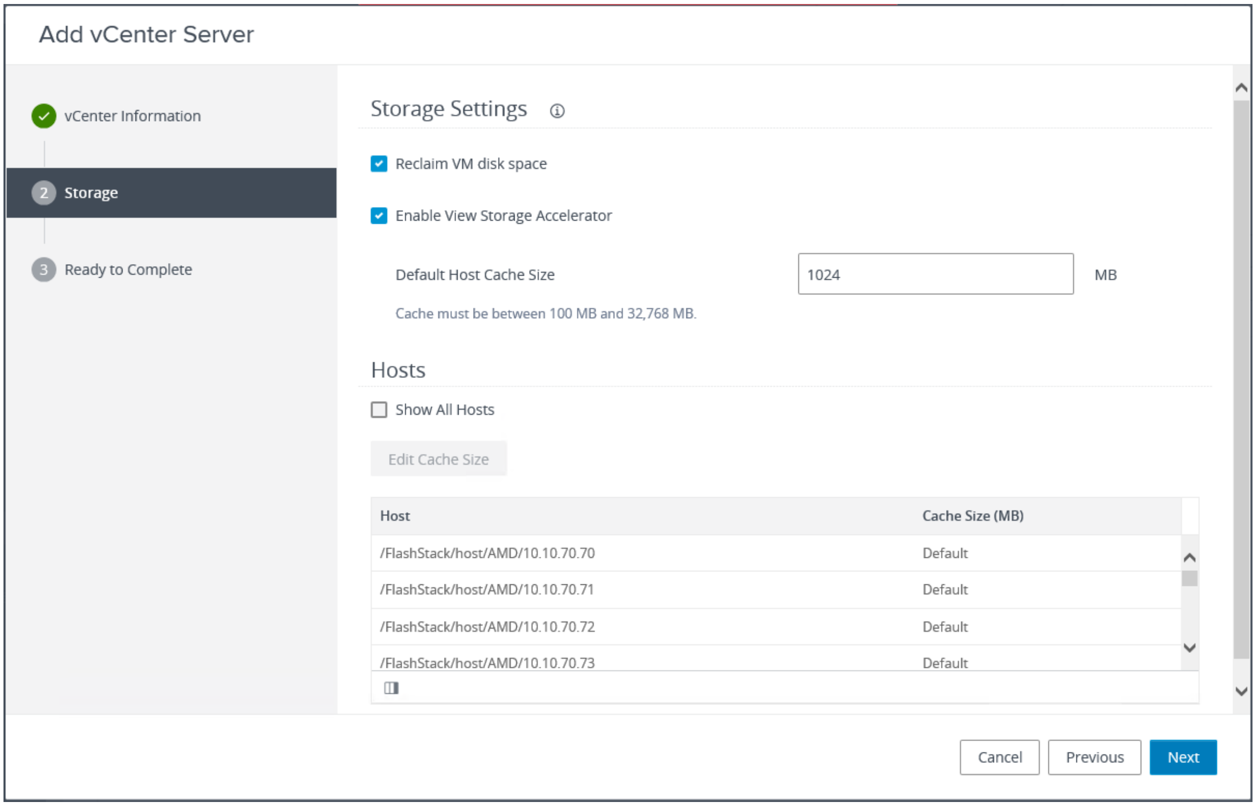

Add vCenter Server

Storage Settings

☑ Reclaim VM disk space

☑ Enable View Storage Accelerator

Default Host Cache Size: 1024 MB

Cache must be between 100 MB and 32,768 MB.

Hosts

☐ Show All Hosts

Host	Cache Size (MB)
/FlashStack/host/AMD/10.10.70.70	Default
/FlashStack/host/AMD/10.10.70.71	Default
/FlashStack/host/AMD/10.10.70.72	Default
/FlashStack/host/AMD/10.10.70.73	Default

8. Review the information you provided and click Submit.

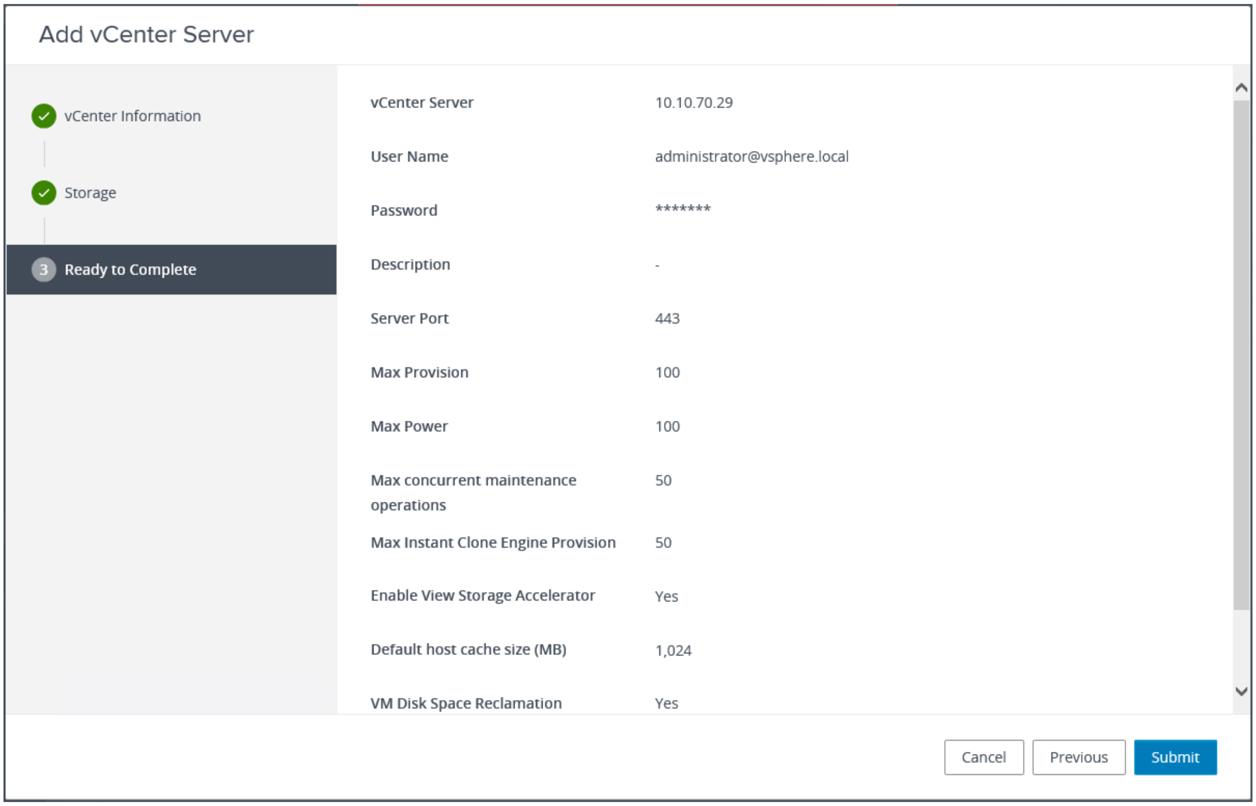

Add vCenter Server

vCenter Server	10.10.70.29
User Name	administrator@vsphere.local
Password	*******
Description	-
Server Port	443
Max Provision	100
Max Power	100
Max concurrent maintenance operations	50
Max Instant Clone Engine Provision	50
Enable View Storage Accelerator	Yes
Default host cache size (MB)	1,024
VM Disk Space Reclamation	Yes

Create VDI Instant Clone Desktop Pool

To create a VDI Instant Clone Desktop Pool, follow these steps:

1. In Horizon Console on the left plane, expand Inventory, select Desktops. Click Add.

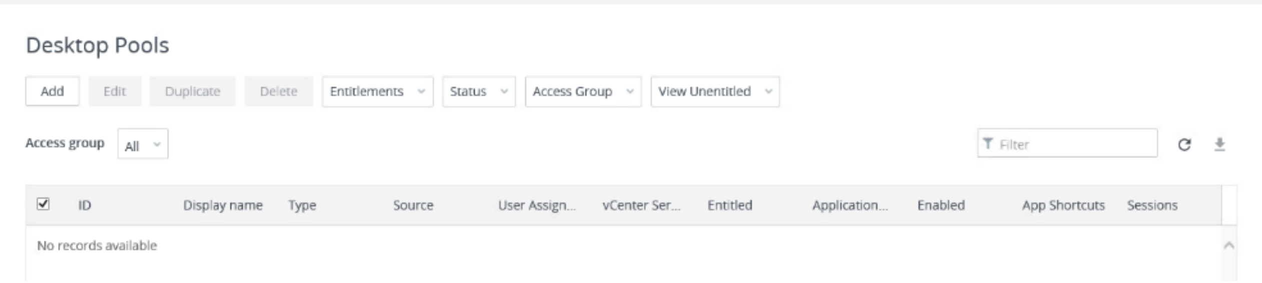

Desktop Pools

☑	ID	Display name	Type	Source	User Assign...	vCenter Ser...	Entitled	Application...	Enabled	App Shortcuts	Sessions
No records available											

2. Select Type of Desktop pool to be created. Click Next

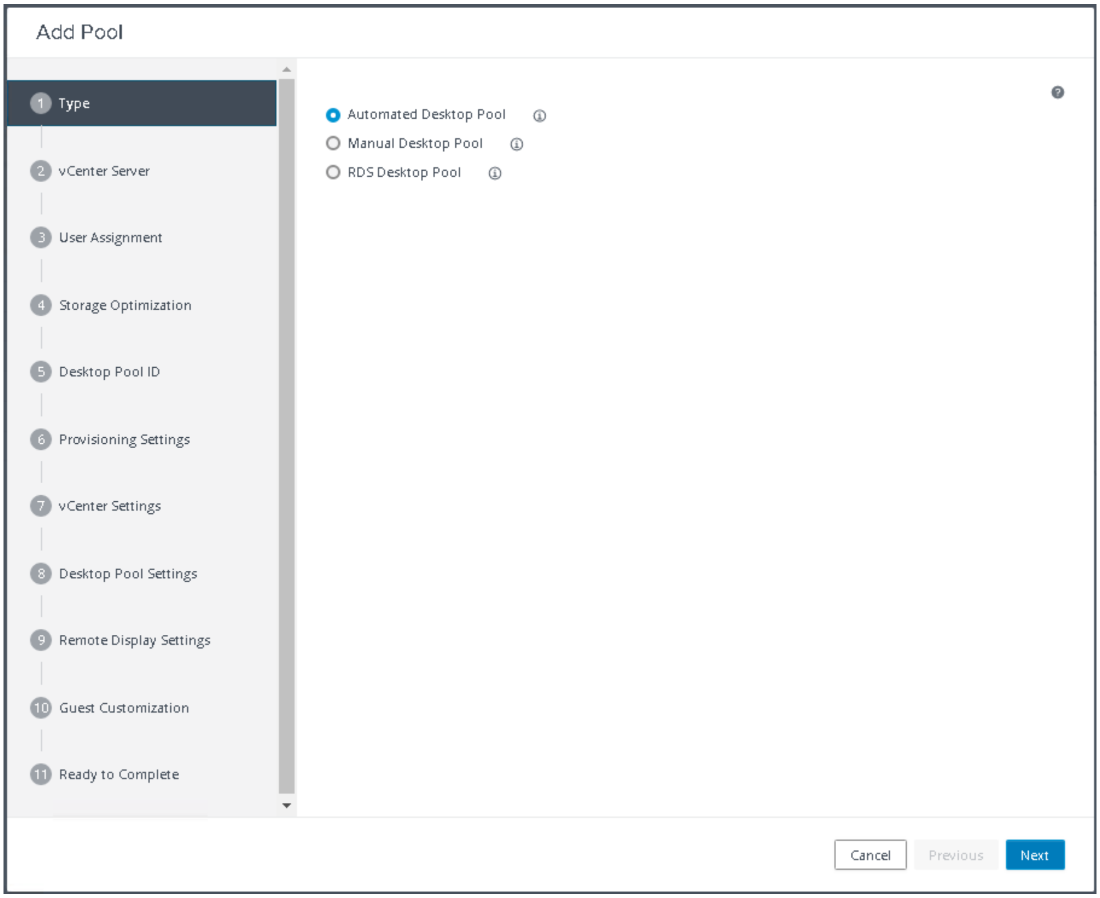

Add Pool

1 Type
2 vCenter Server
3 User Assignment
4 Storage Optimization
5 Desktop Pool ID
6 Provisioning Settings
7 vCenter Settings
8 Desktop Pool Settings
9 Remote Display Settings
10 Guest Customization
11 Ready to Complete

● Automated Desktop Pool
○ Manual Desktop Pool
○ RDS Desktop Pool

3. Choose provisioning type for the desktops in the pool (we created Instant Clones and Full Virtual Machines pools in this design). Click Next.

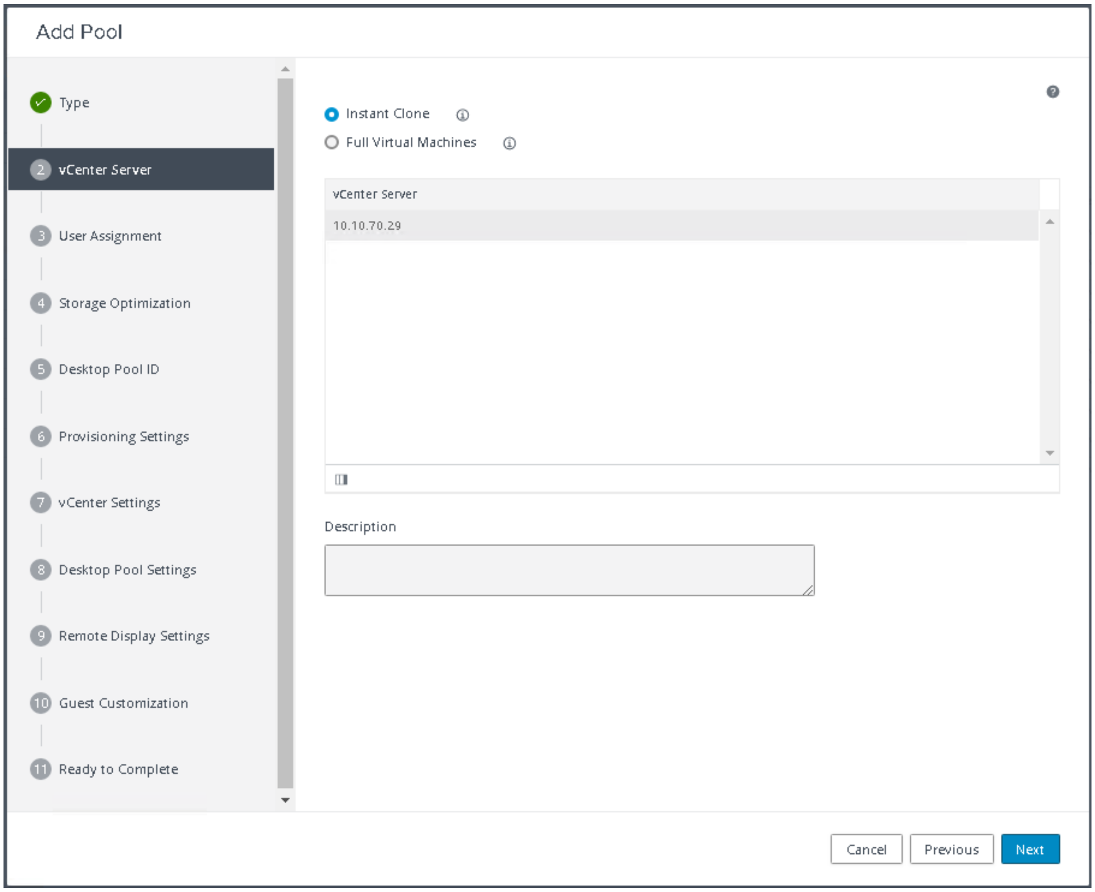
Add Pool

● Instant Clone

○ Full Virtual Machines

vCenter Server
10.10.70.29

Description

4. Select the User assignment to be used by the desktop pool. Click Next.

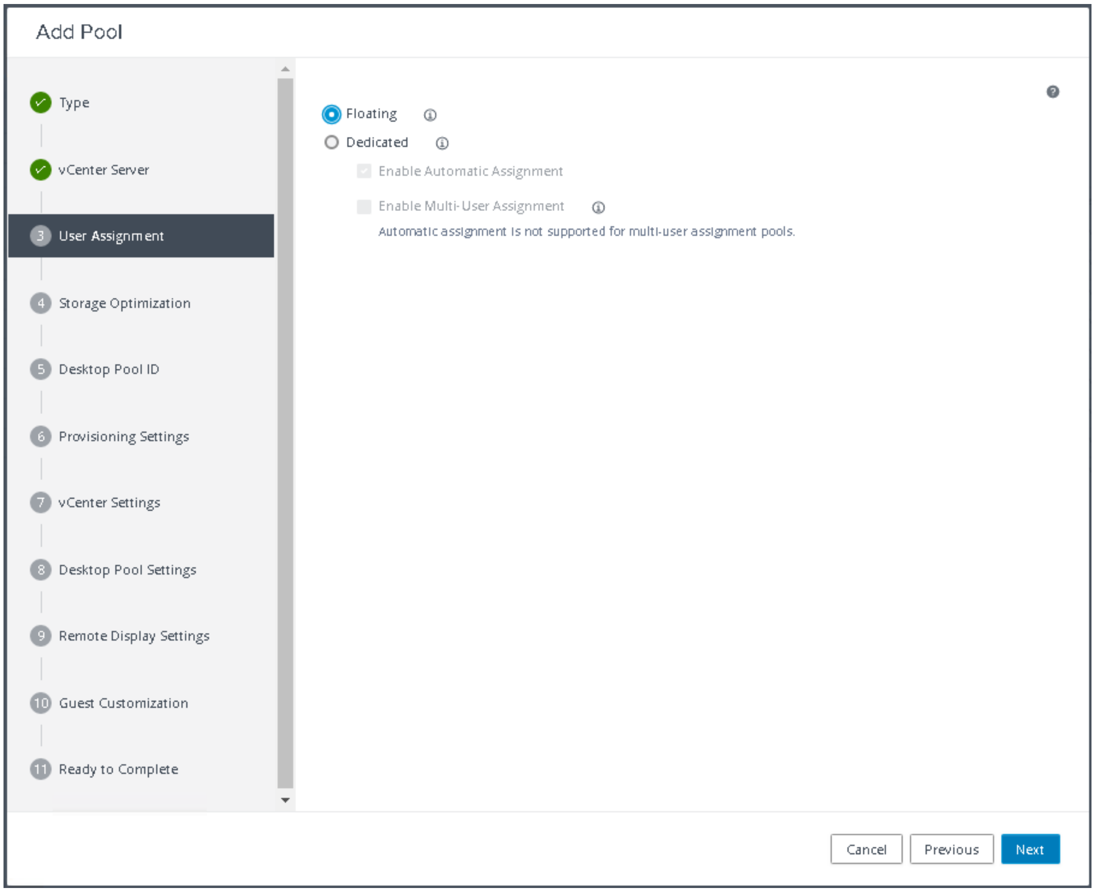

● Floating

○ Dedicated

☑ Enable Automatic Assignment

□ Enable Multi-User Assignment

Automatic assignment is not supported for multi-user assignment pools.

5. Select the required option for Storage Policy Management. Click Next.

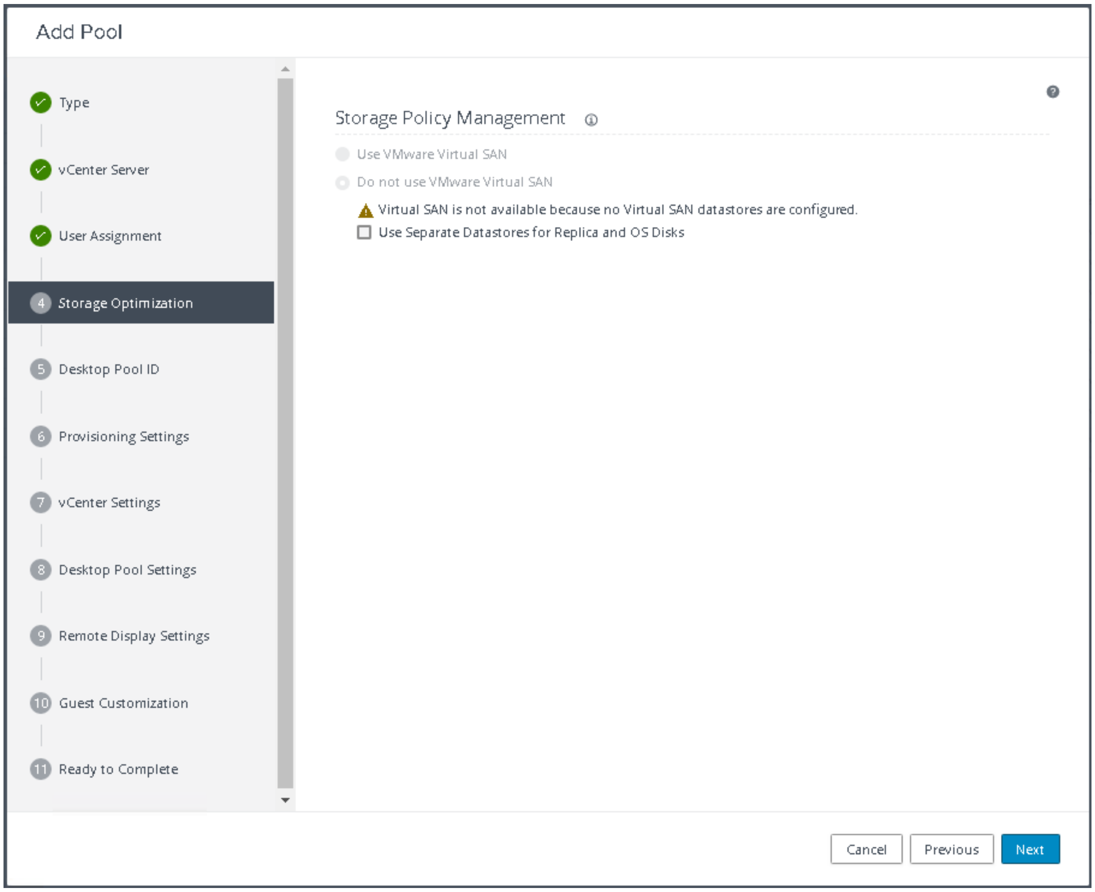

Storage Policy Management

○ Use VMware Virtual SAN

○ Do not use VMware Virtual SAN

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

□ Use Separate Datastores for Replica and OS Disks

6. Provide Desktop Pool ID and virtual display name. Click Next.

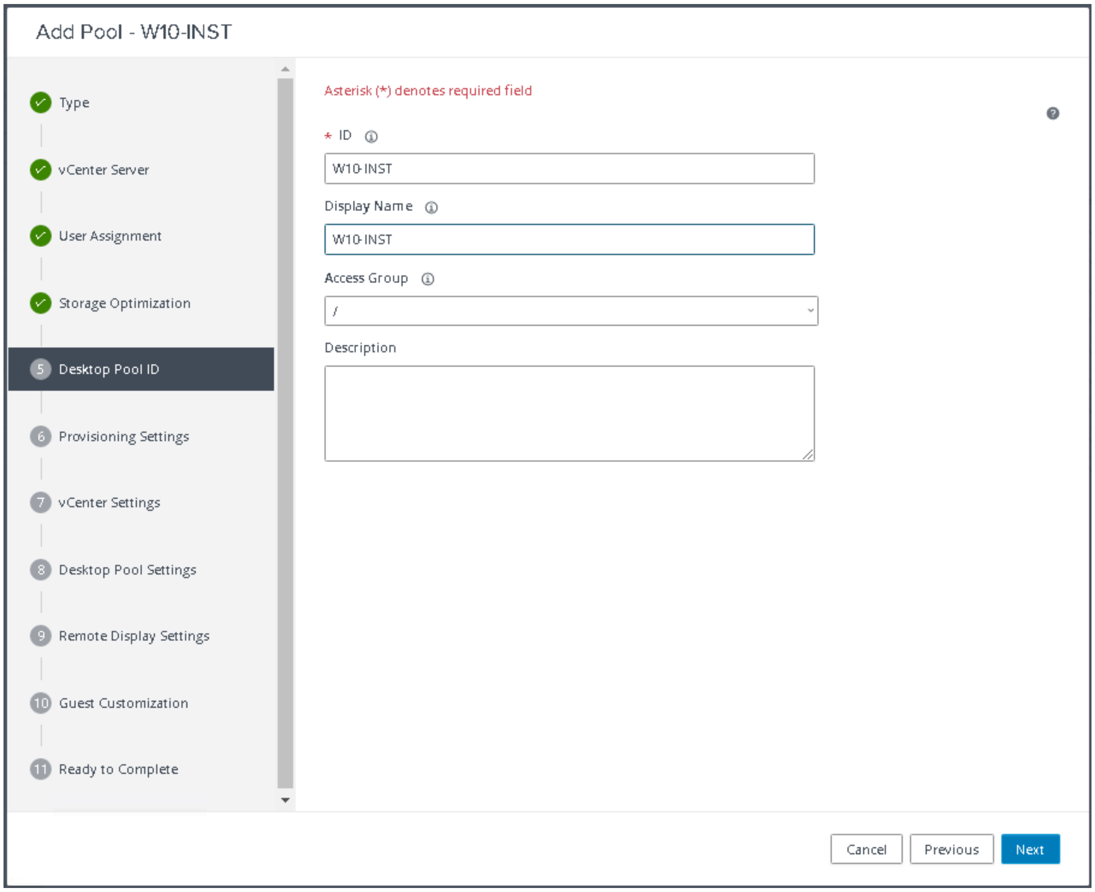
Add Pool - W10-INST

Asterisk (*) denotes required field

* ID: W10-INST

Display Name: W10-INST

Access Group: /

Description:

7. Provide the naming pattern and the number of desktops to be provisioned. Click Next.

In this validated design we used:
Single Server pool – 186
Cluster pool – 1674

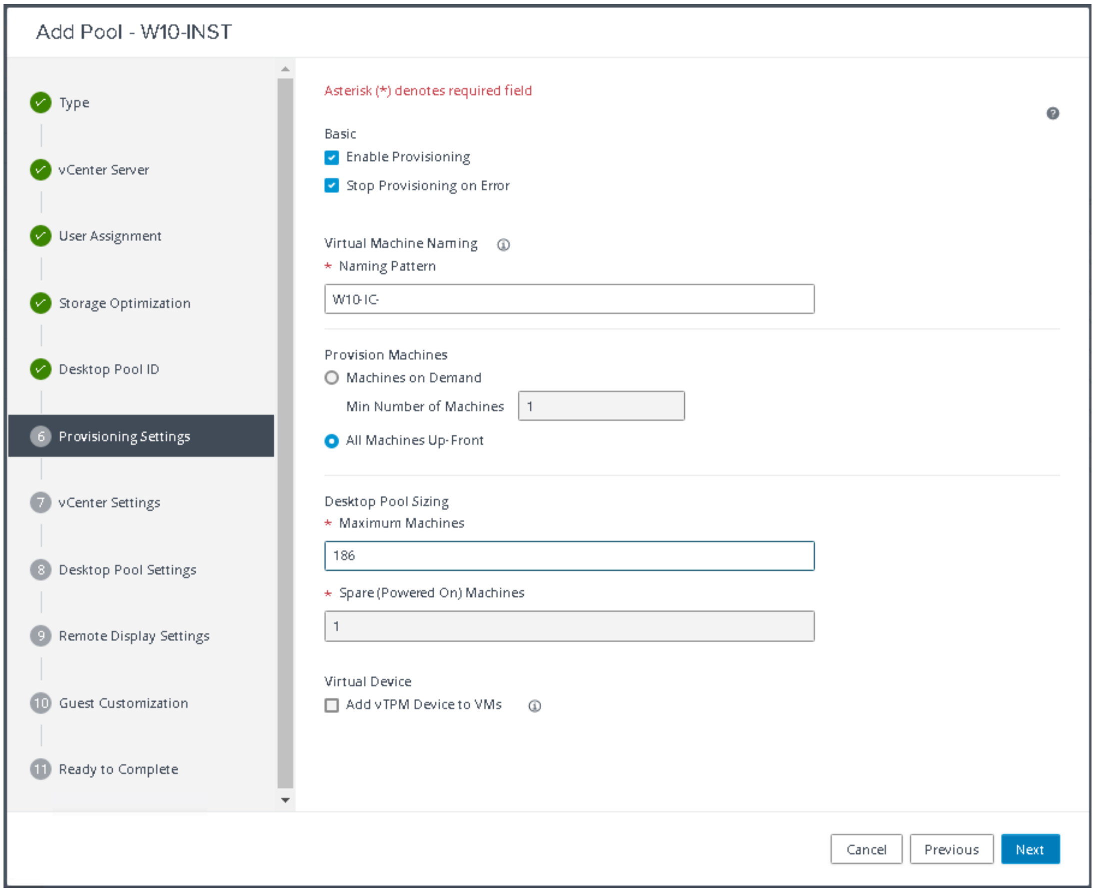

Add Pool - W10-INST

Asterisk (*) denotes required field

Basic

☑ Enable Provisioning

☑ Stop Provisioning on Error

Virtual Machine Naming

* Naming Pattern: W10-IC-

Provision Machines

○ Machines on Demand

Min Number of Machines: 1

● All Machines Up-Front

Desktop Pool Sizing

* Maximum Machines: 186

* Spare (Powered On) Machines: 1

Virtual Device

□ Add vTPM Device to VMs

8. Provide the parent VM, snapshot and host/cluster info, and data store information for the virtual machines to create. Click Next.

A single datastore was used per 10 host pools.

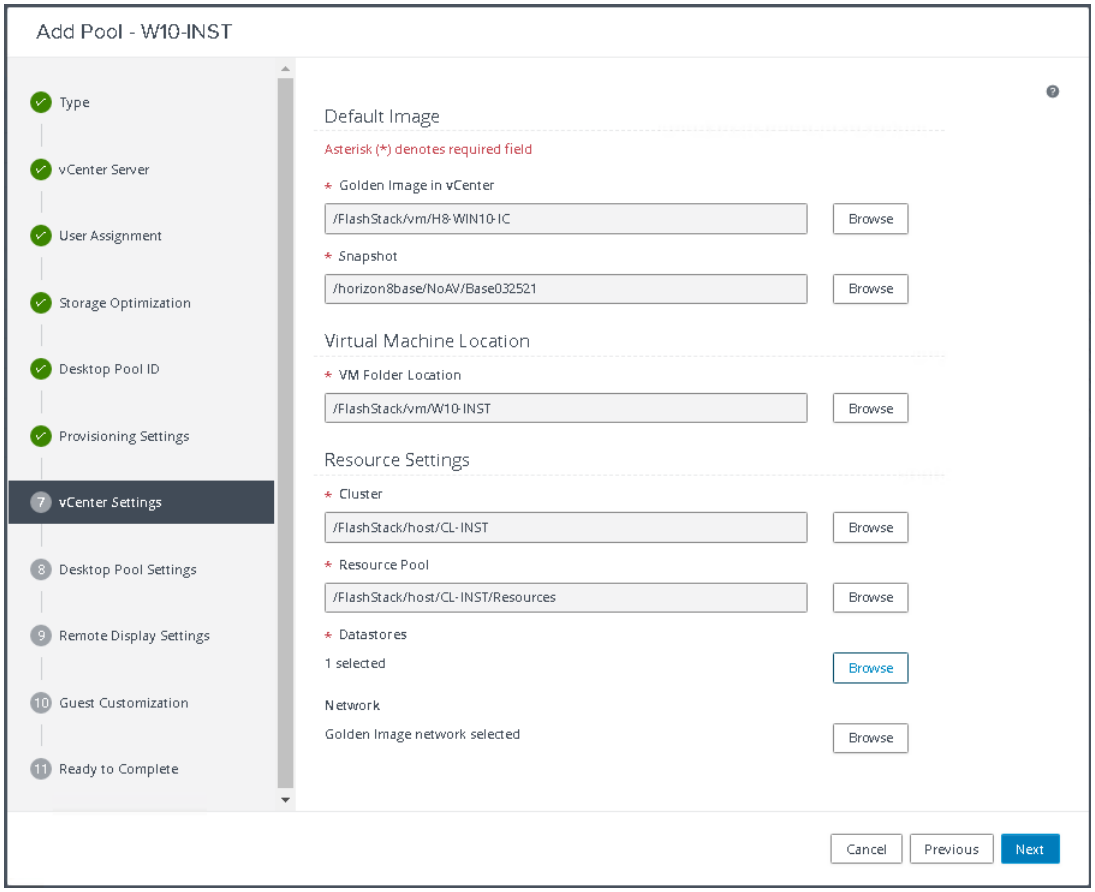

Add Pool - W10-INST

Default Image

Asterisk (*) denotes required field

* Golden Image in vCenter: /FlashStack/vm/H8-WIN10-IC

* Snapshot: /horizon8base/NoAV/Base032521

Virtual Machine Location

* VM Folder Location: /FlashStack/vm/W10-INST

Resource Settings

* Cluster: /FlashStack/host/CL-INST

* Resource Pool: /FlashStack/host/CL-INST/Resources

* Datastores: 1 selected

Network: Golden Image network selected

9. Configure the State and Session Type for Desktop Pool Settings. Click Next.

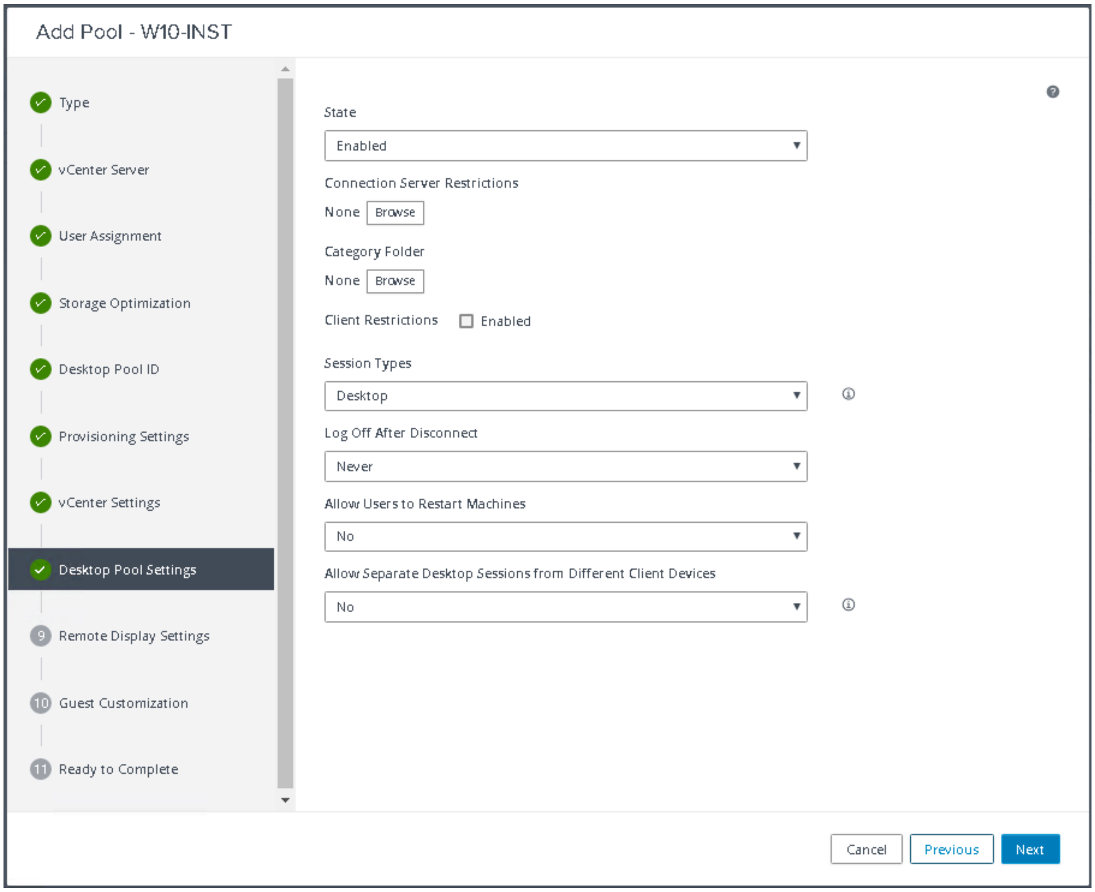

Add Pool - W10-INST

1. Type
2. vCenter Server
3. User Assignment
4. Storage Optimization
5. Desktop Pool ID
6. Provisioning Settings
7. vCenter Settings
8. Desktop Pool Settings
9. Remote Display Settings
10. Guest Customization
11. Ready to Complete

State: Enabled

Connection Server Restrictions: None Browse

Category Folder: None Browse

Client Restrictions ☐ Enabled

Session Types: Desktop

Log Off After Disconnect: Never

Allow Users to Restart Machines: No

Allow Separate Desktop Sessions from Different Client Devices: No

10. Provide the customizations to be used during the desktop deployment. Click Next.

PCoIP Display protocol was used for the VDI pools in this validated design.

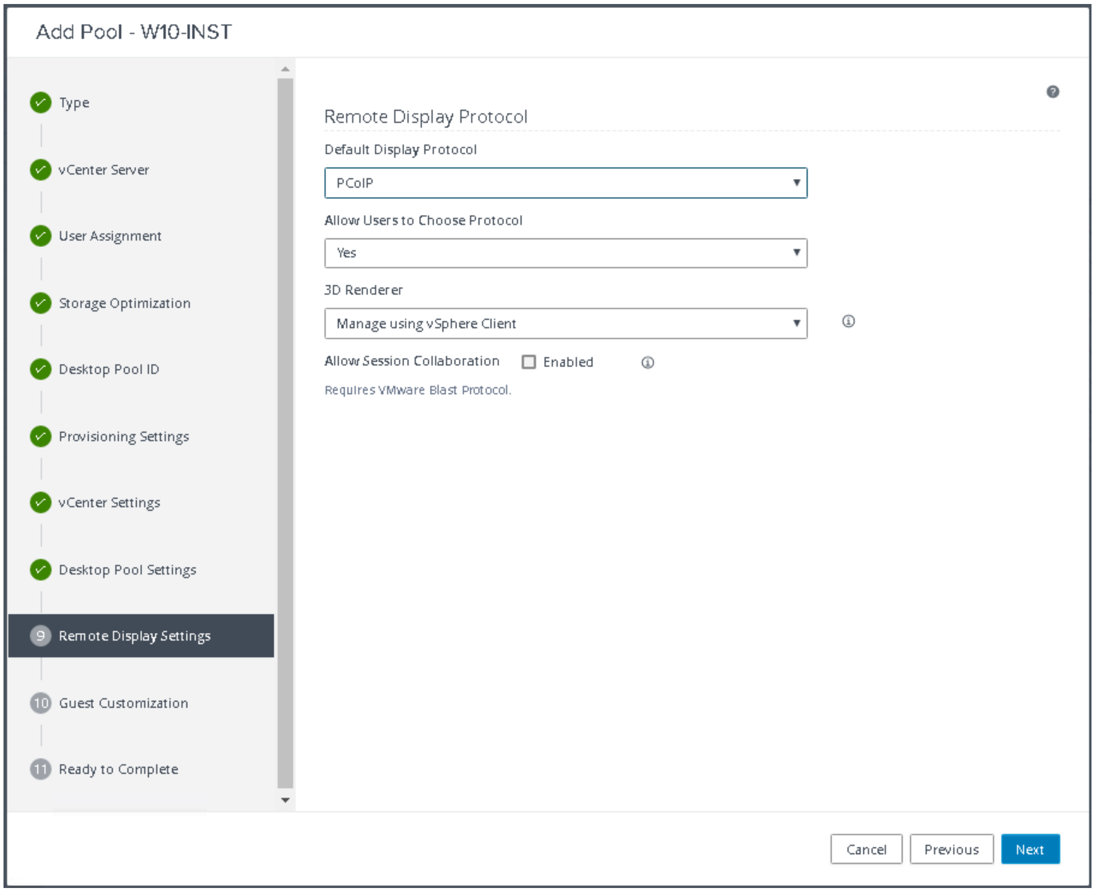

Add Pool - W10-INST

Remote Display Protocol

Default Display Protocol: PCoIP

Allow Users to Choose Protocol: Yes

3D Renderer: Manage using vSphere Client

Allow Session Collaboration ☐ Enabled

Requires VMware Blast Protocol.

11. Select the AD Container for desktops to place in a Domain Controller computer location.

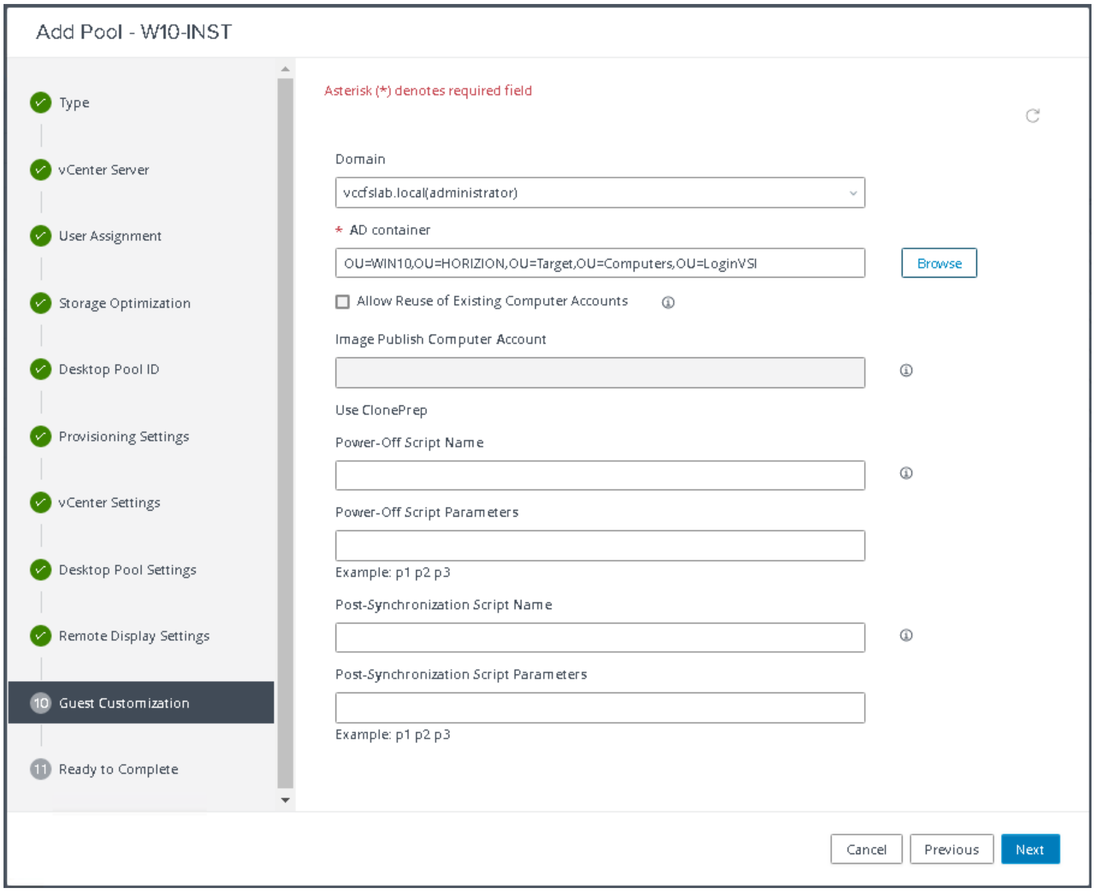

Add Pool - W10-INST

Asterisk (*) denotes required field

Domain

vccfslab.local(administrator)

* AD container

OU=WIN10,OU=HORIZION,OU=Target,OU=Computers,OU=LoginVSI

□ Allow Reuse of Existing Computer Accounts

Image Publish Computer Account

Use ClonePrep

Power-Off Script Name

Power-Off Script Parameters

Example: p1 p2 p3

Post-Synchronization Script Name

Post-Synchronization Script Parameters

Example: p1 p2 p3

12. Review all the deployment specifications and click Submit to complete the deployment.

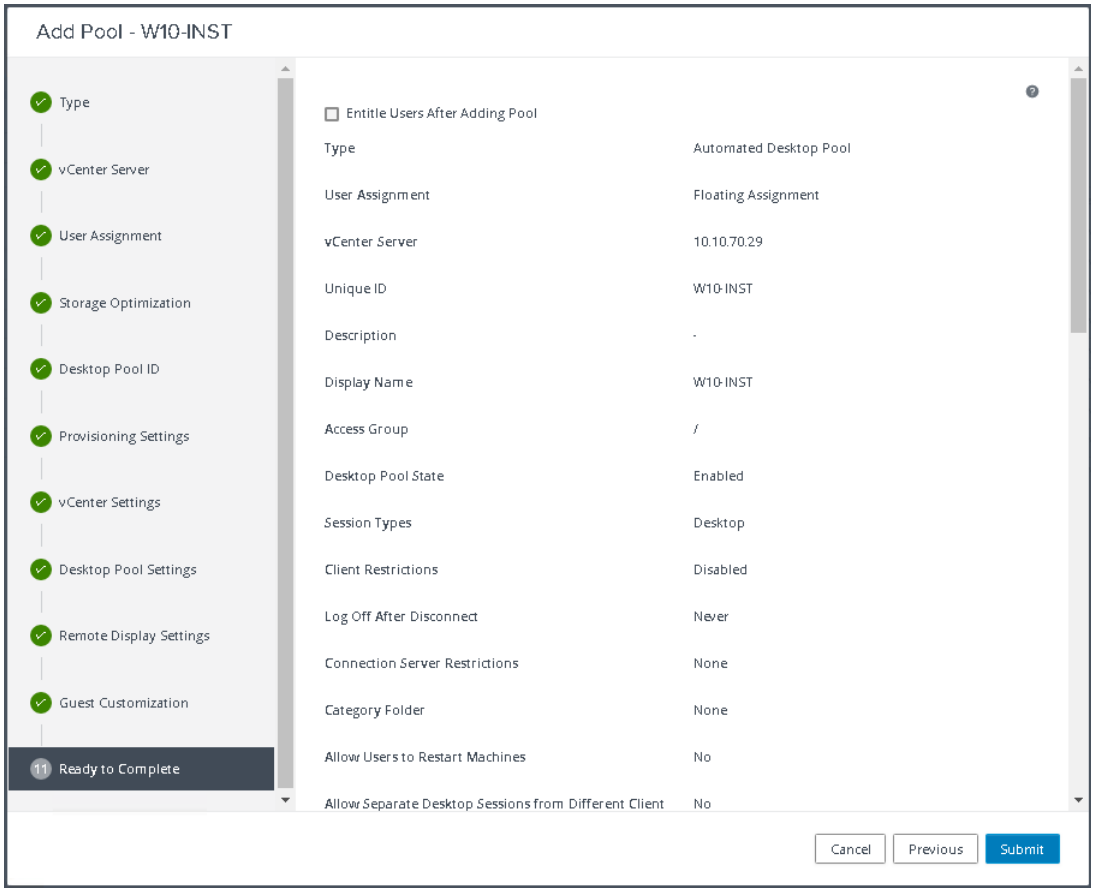

Add Pool - W10-INST

□ Entitle Users After Adding Pool

Type	Automated Desktop Pool
User Assignment	Floating Assignment
vCenter Server	10.10.70.29
Unique ID	W10-INST
Description	-
Display Name	W10-INST
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Client Restrictions	Disabled
Log Off After Disconnect	Never
Connection Server Restrictions	None
Category Folder	None
Allow Users to Restart Machines	No
Allow Separate Desktop Sessions from Different Client	No

13. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

Create VDI Full Clone Desktop Pool

To create a VDI Full Clone Desktop Pool, follow these steps:

1. Select Type of Desktop pool to be created. Click Next.

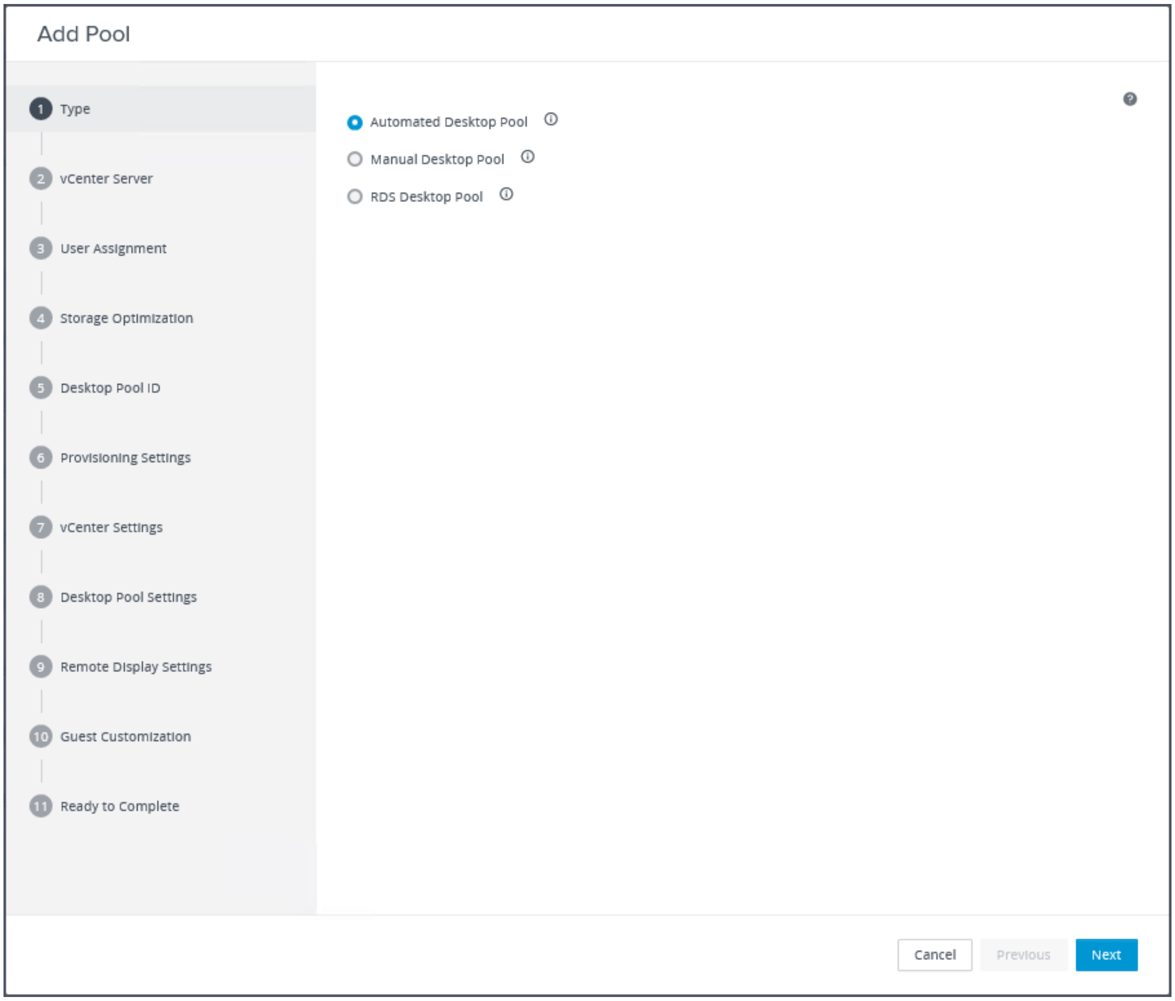

● Automated Desktop Pool

○ Manual Desktop Pool

○ RDS Desktop Pool

2. Choose the provisioning type for the desktops in the pool (we created Instant Clones and Full Virtual Machines pools in this design). Click Next.

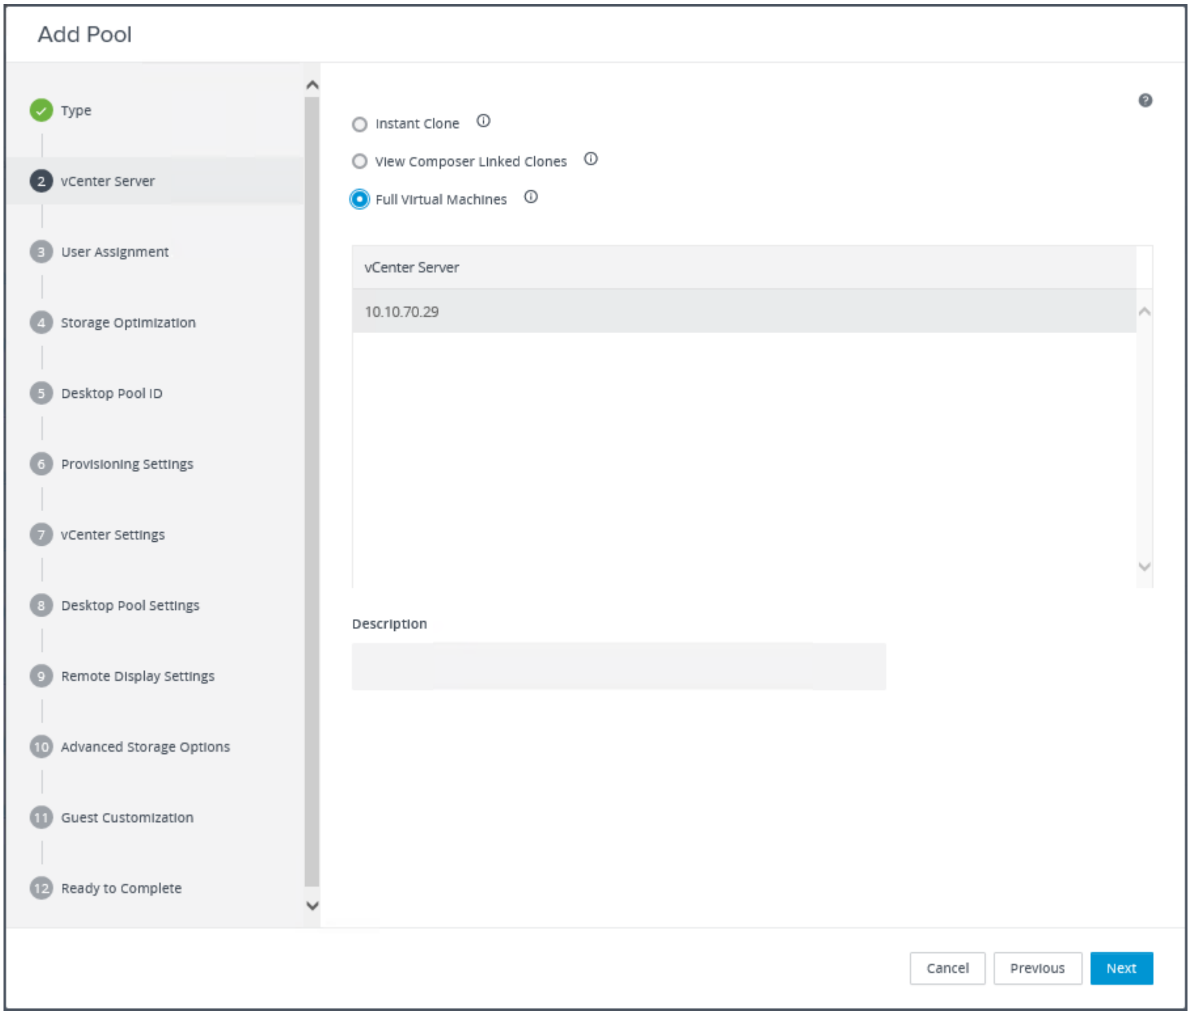
Add Pool

○ Instant Clone
○ View Composer Linked Clones
● Full Virtual Machines

vCenter Server
10.10.70.29

Description

3. Select the User assignment to be used by the desktop pool. Click Next.

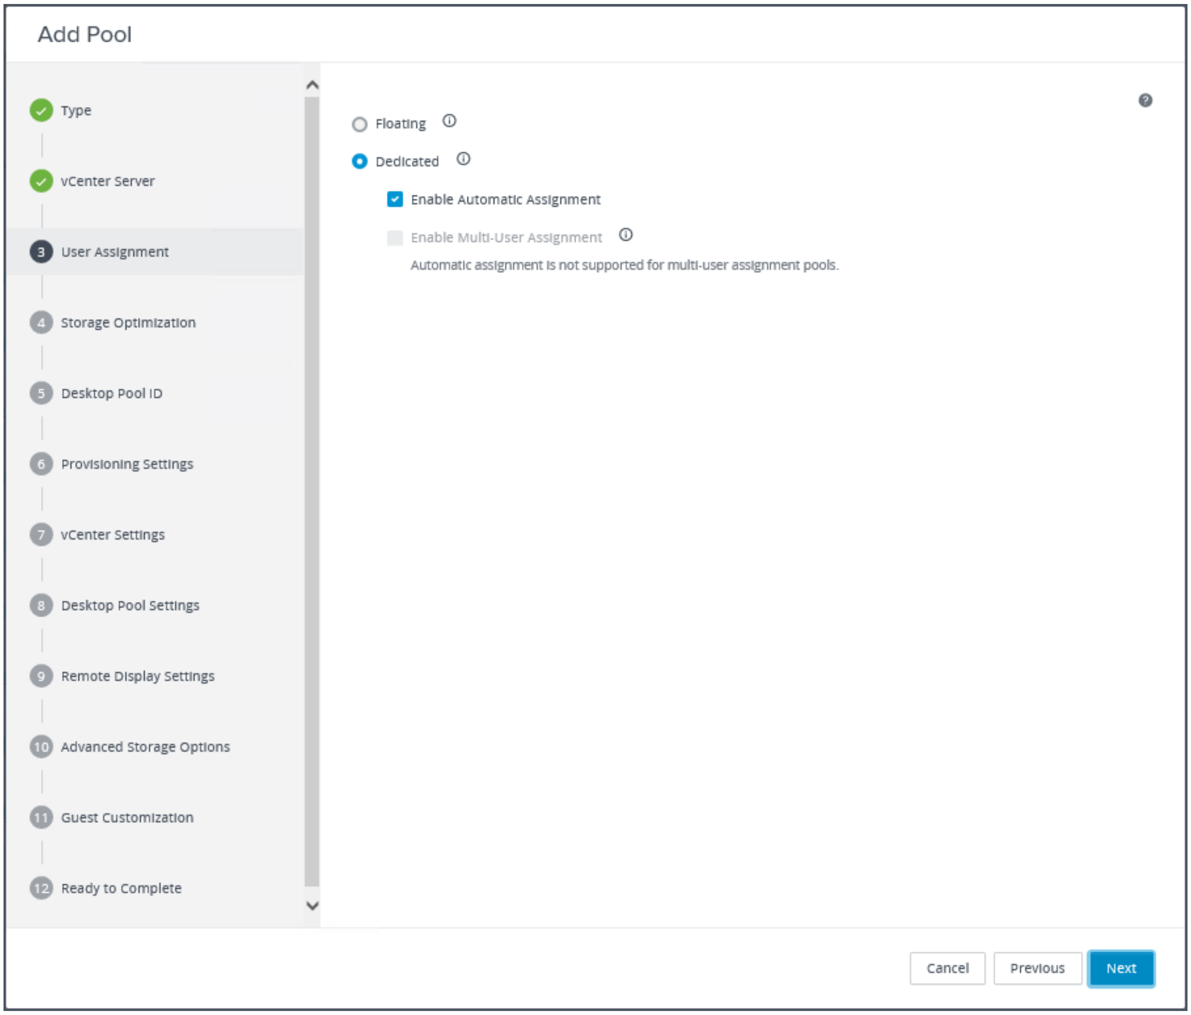

Add Pool

○ Floating

● Dedicated

☑ Enable Automatic Assignment

□ Enable Multi-User Assignment

Automatic assignment is not supported for multi-user assignment pools.

4. On Storage Optimization screen click Next.

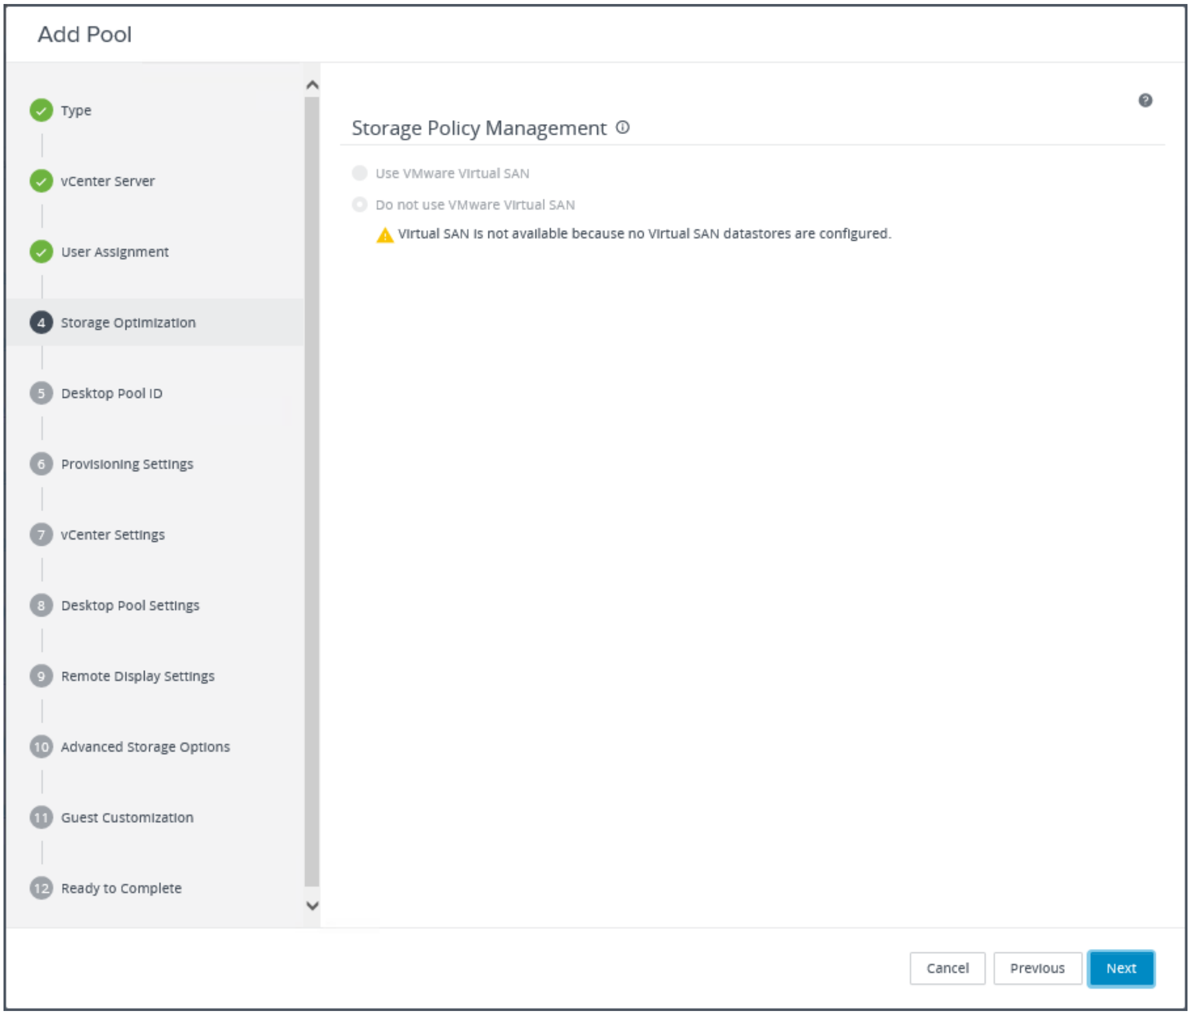

Storage Policy Management

○ Use VMware Virtual SAN

○ Do not use VMware Virtual SAN

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

5. Provide the Desktop Pool ID and Display Name. Click Next.

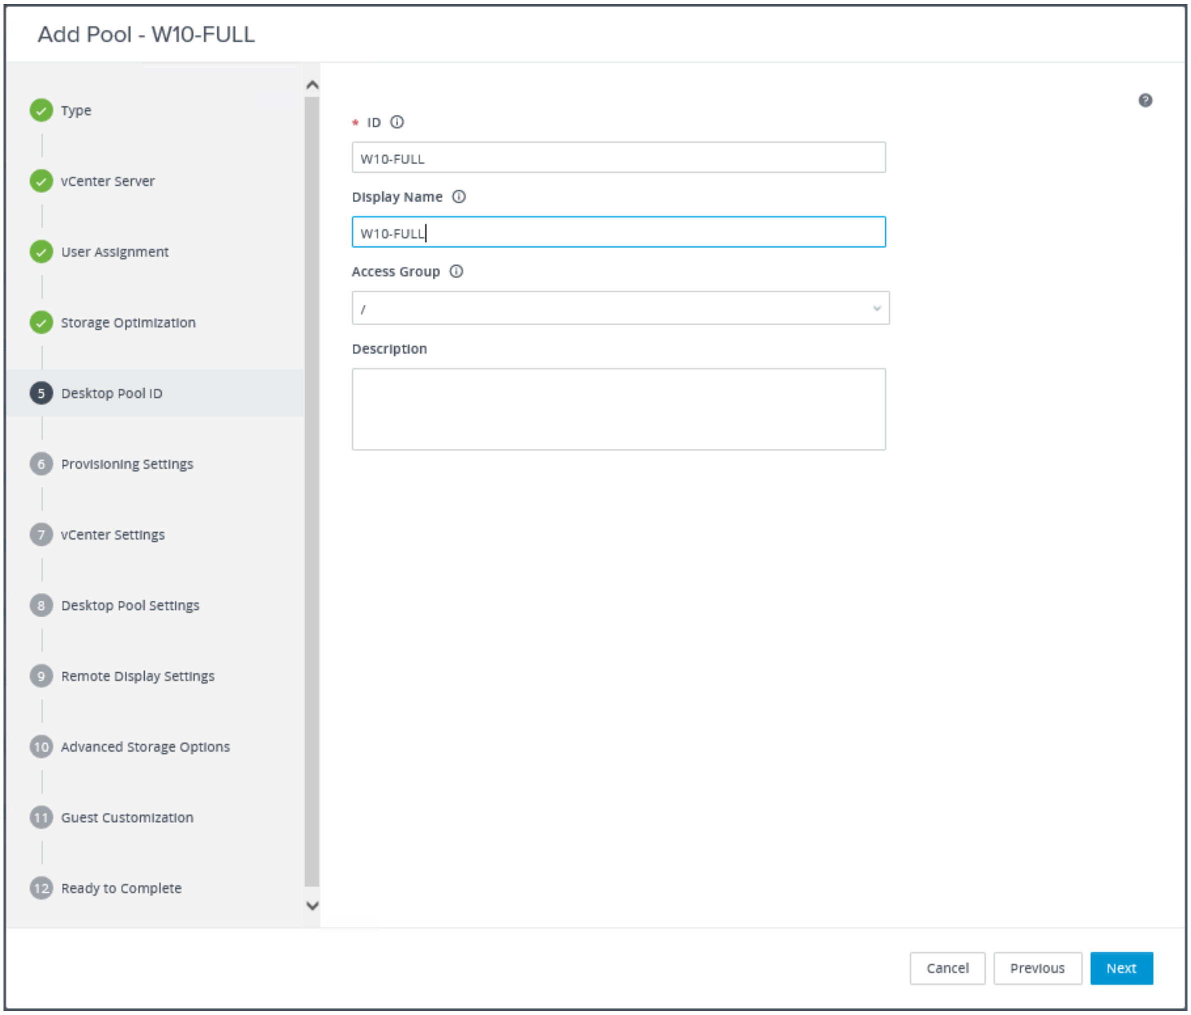
Add Pool - W10-FULL

* ID: W10-FULL

Display Name: W10-FULL

Access Group: /

Description:

6. Provide the naming pattern and the number of desktops to be provisioned. Click Next.

In this validated design for VDI pools we used:
Single Server pool – 186
Cluster pool – 1674

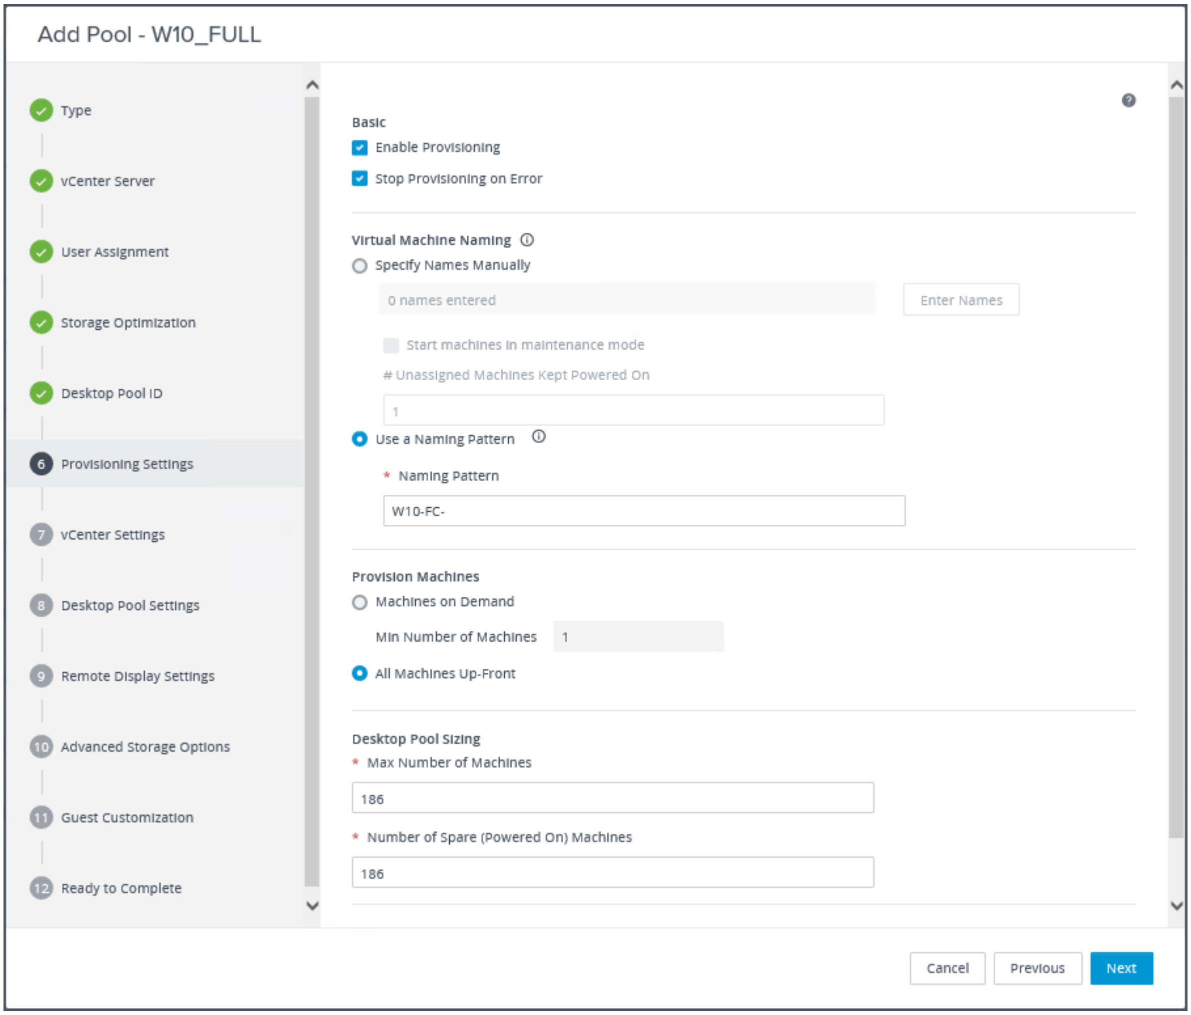

Add Pool - W10_FULL

Basic

☑ Enable Provisioning

☑ Stop Provisioning on Error

Virtual Machine Naming

○ Specify Names Manually

0 names entered

Enter Names

□ Start machines in maintenance mode

Unassigned Machines Kept Powered On

1

● Use a Naming Pattern

* Naming Pattern

W10-FC-

Provision Machines

○ Machines on Demand

Min Number of Machines 1

● All Machines Up-Front

Desktop Pool Sizing

* Max Number of Machines

186

* Number of Spare (Powered On) Machines

186

7. Provide the parent VM, snapshot and host/cluster info, data store information for the virtual machines to create.

A single datastore was used per 10 host pools.

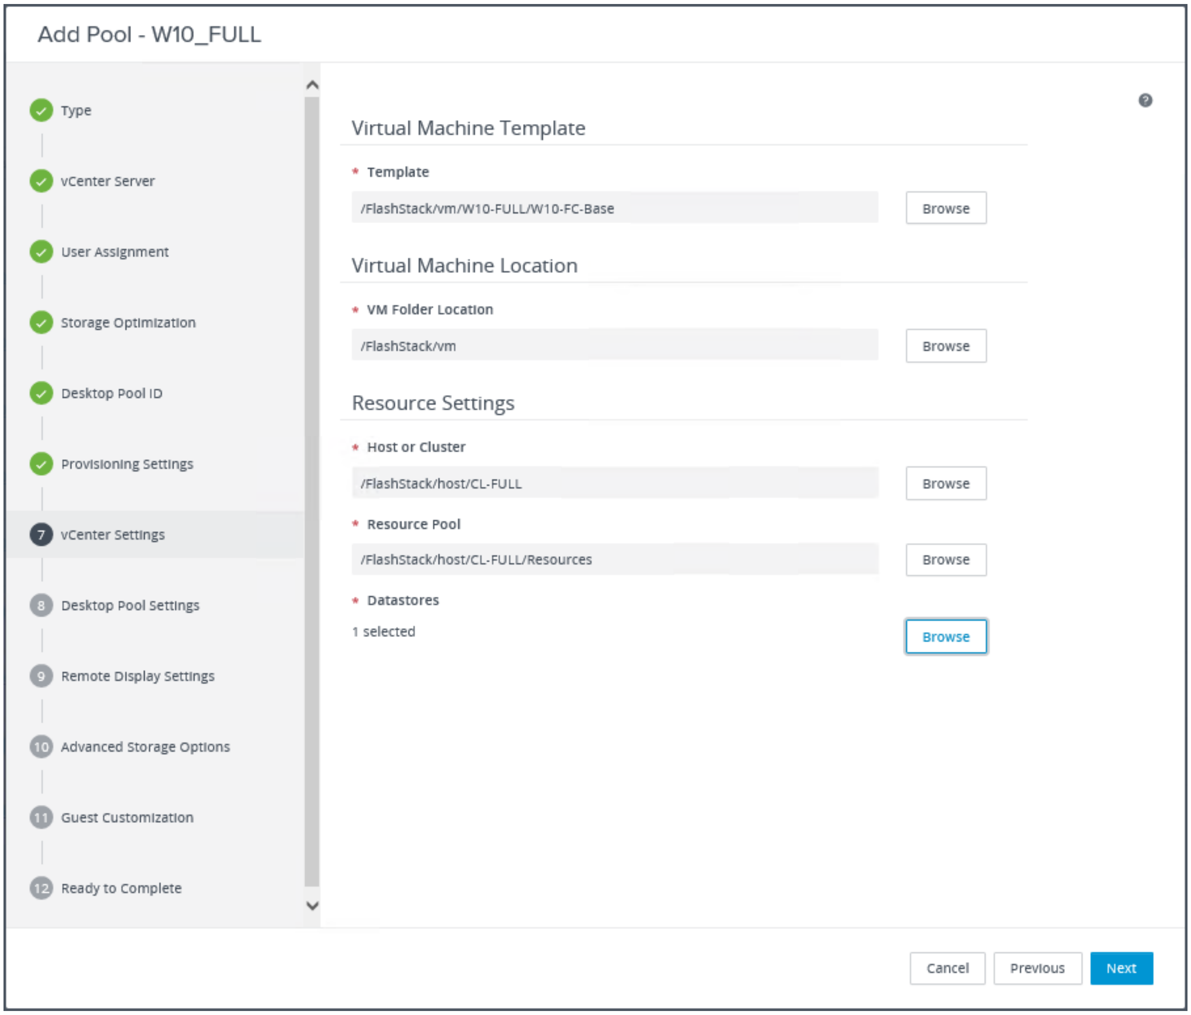

Add Pool - W10_FULL

Virtual Machine Template

* Template

/FlashStack/vm/W10-FULL/W10-FC-Base

Virtual Machine Location

* VM Folder Location

/FlashStack/vm

Resource Settings

* Host or Cluster

/FlashStack/host/CL-FULL

* Resource Pool

/FlashStack/host/CL-FULL/Resources

* Datastores

1 selected

8. Configure Desktop Pool settings.

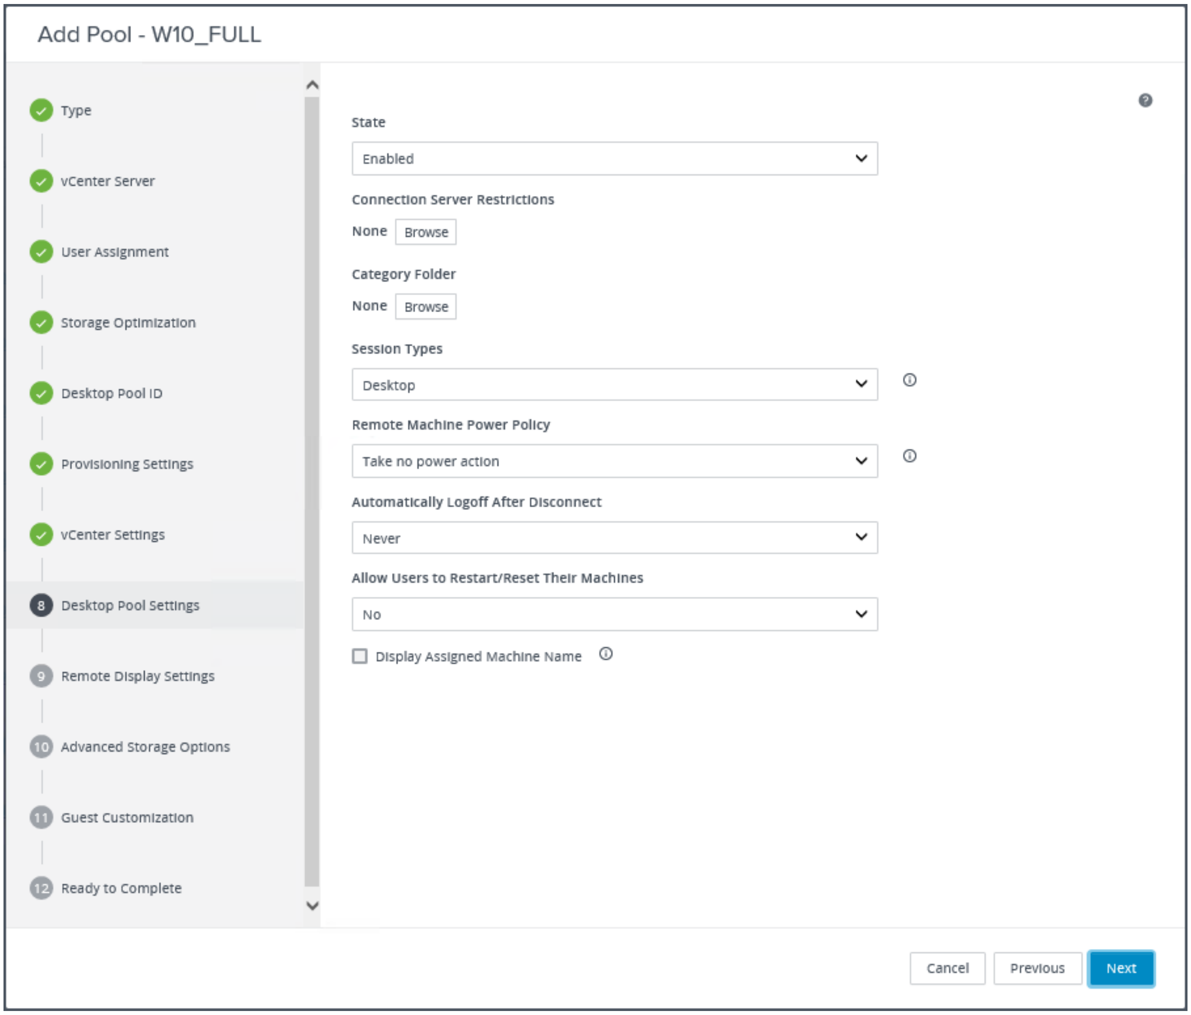

Add Pool - W10_FULL

State: Enabled

Connection Server Restrictions: None Browse

Category Folder: None Browse

Session Types: Desktop

Remote Machine Power Policy: Take no power action

Automatically Logoff After Disconnect: Never

Allow Users to Restart/Reset Their Machines: No

☐ Display Assigned Machine Name

9. Provide the customizations to be used during the desktop deployment.

PCoIP Display protocol was used for the VDI pools in this validated design.

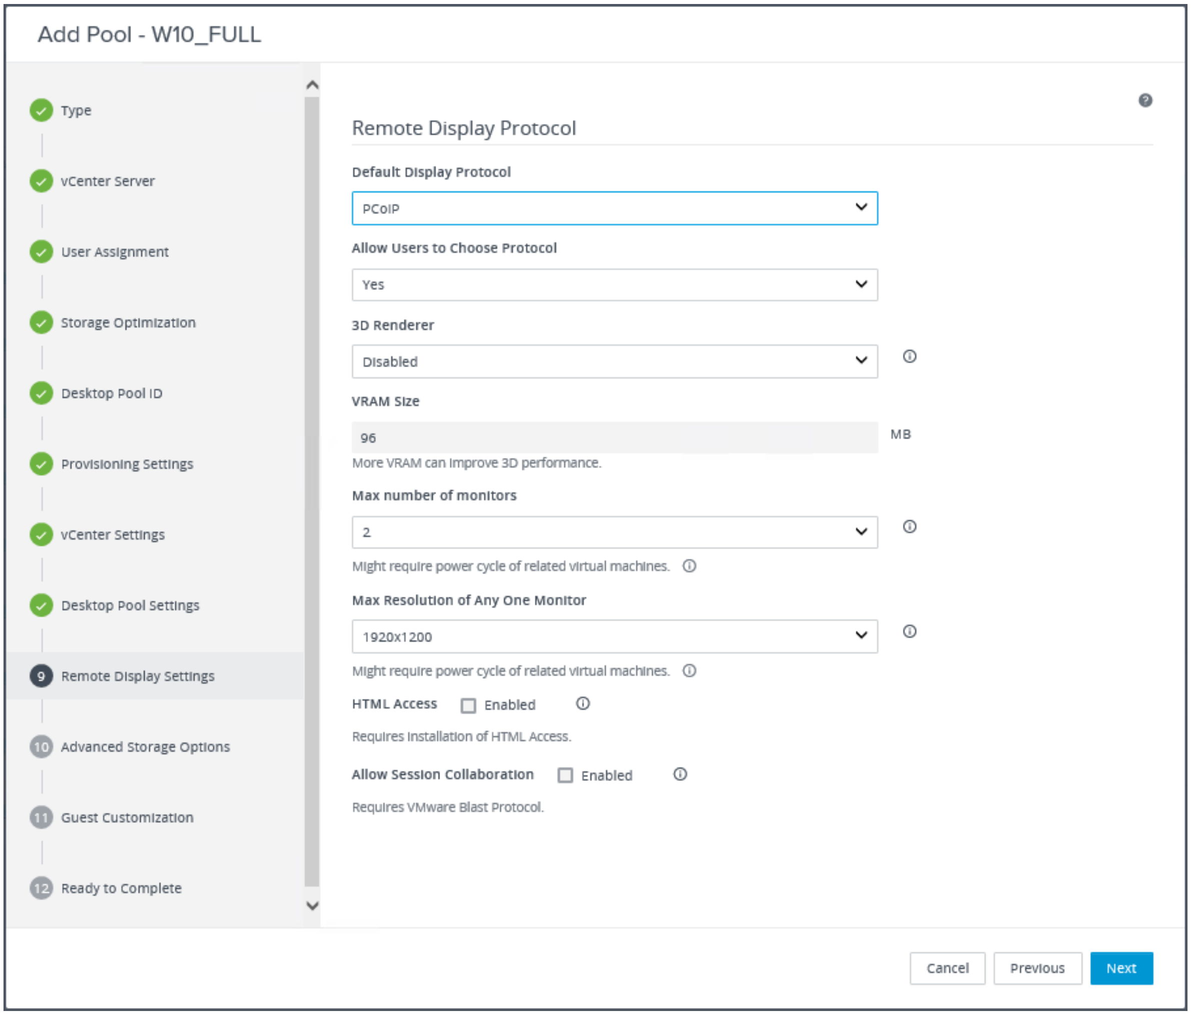
Add Pool - W10_FULL

Remote Display Protocol

Default Display Protocol: PCoIP

Allow Users to Choose Protocol: Yes

3D Renderer: Disabled

VRAM Size: 96 MB

More VRAM can improve 3D performance.

Max number of monitors: 2

Might require power cycle of related virtual machines.

Max Resolution of Any One Monitor: 1920x1200

Might require power cycle of related virtual machines.

HTML Access ☐ Enabled

Requires installation of HTML Access.

Allow Session Collaboration ☐ Enabled

Requires VMware Blast Protocol.

10. Configure Advanced Storage Options. Click Next.

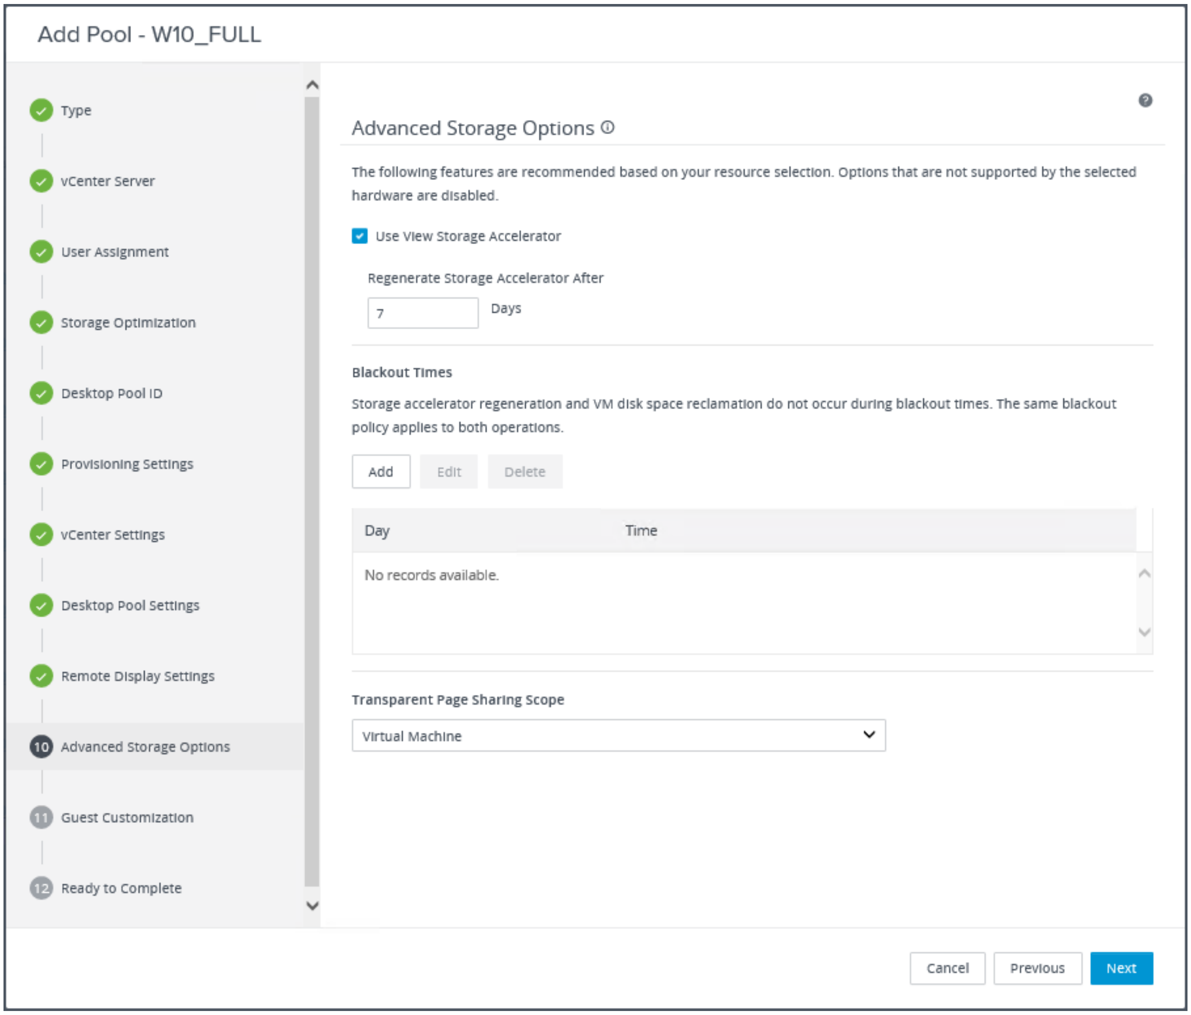

Add Pool - W10_FULL

Advanced Storage Options

The following features are recommended based on your resource selection. Options that are not supported by the selected hardware are disabled.

☑ Use View Storage Accelerator

Regenerate Storage Accelerator After 7 Days

Blackout Times

Storage accelerator regeneration and VM disk space reclamation do not occur during blackout times. The same blackout policy applies to both operations.

Day	Time
No records available.	

Transparent Page Sharing Scope

Virtual Machine

11. Select the VM Customization Specification to be used during deployment. Click Next.

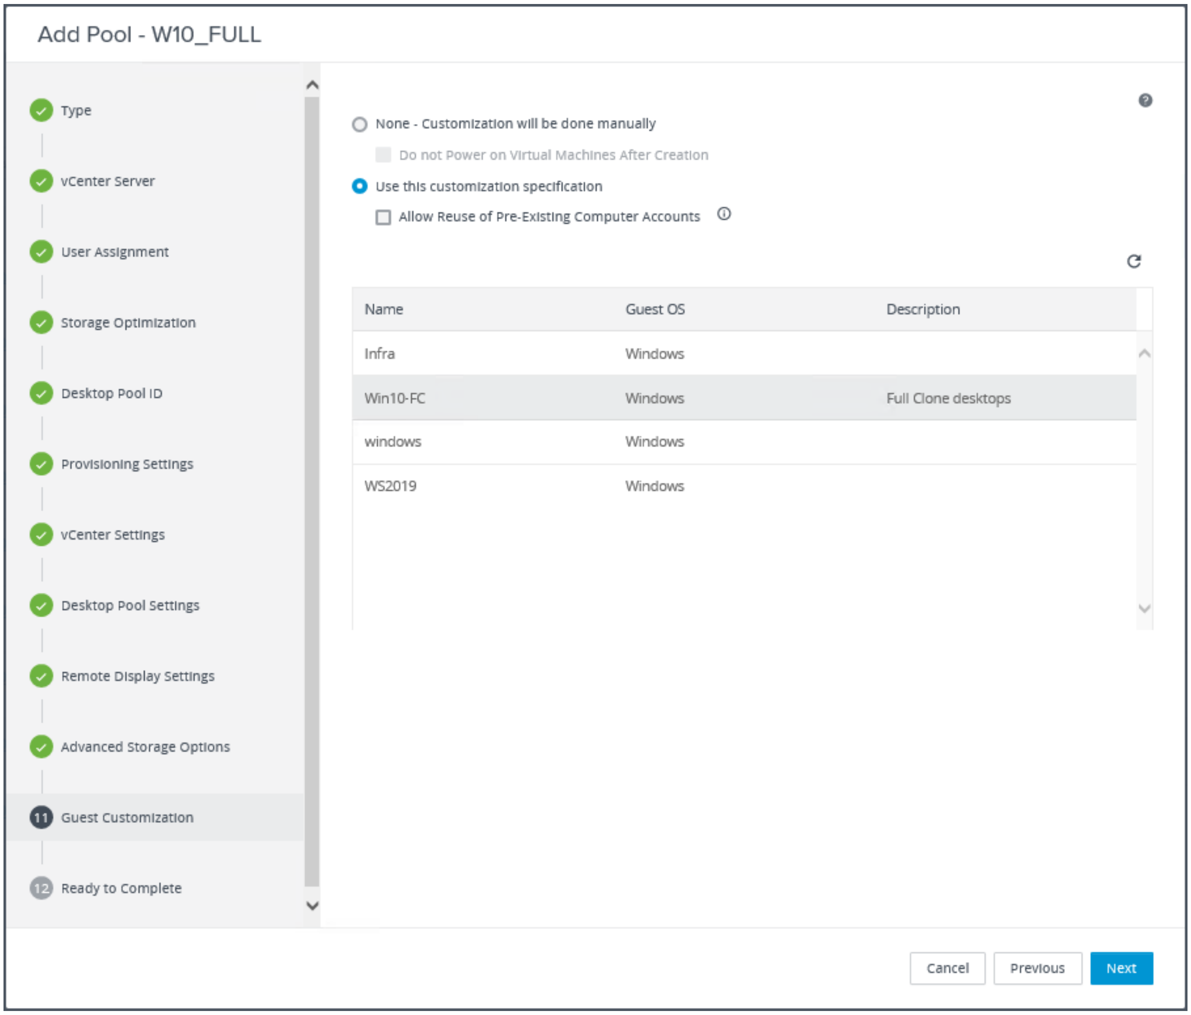

Add Pool - W10_FULL

○ None - Customization will be done manually

□ Do not Power on Virtual Machines After Creation

● Use this customization specification

□ Allow Reuse of Pre-Existing Computer Accounts

Name	Guest OS	Description
Infra	Windows	
Win10-FC	Windows	Full Clone desktops
windows	Windows	
WS2019	Windows	

12. Review all the deployment specifications and click Submit to complete the deployment.

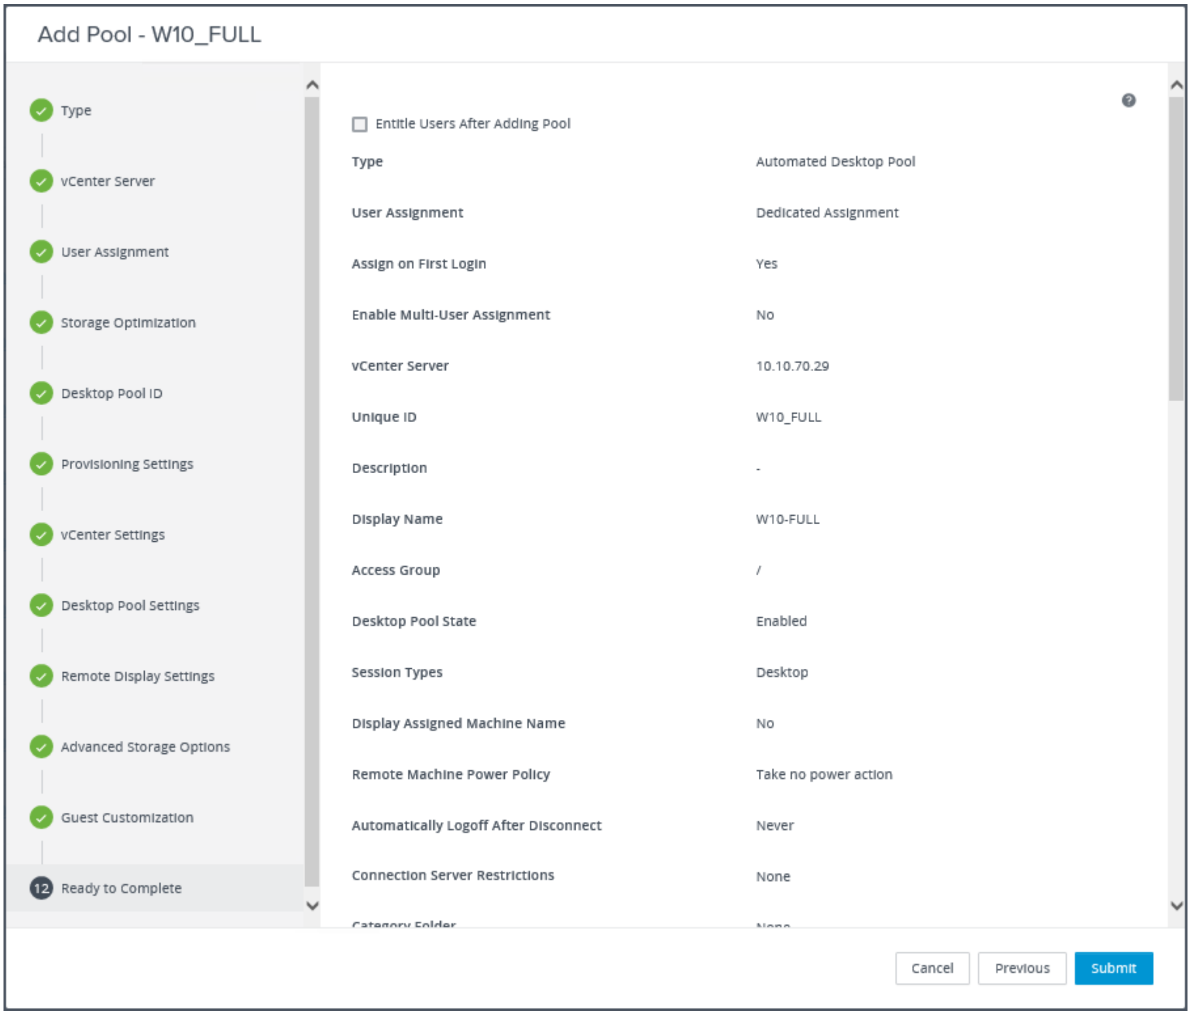

☐ Entitle Users After Adding Pool

Type	Automated Desktop Pool
User Assignment	Dedicated Assignment
Assign on First Login	Yes
Enable Multi-User Assignment	No
vCenter Server	10.10.70.29
Unique ID	W10_FULL
Description	-
Display Name	W10-FULL
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Display Assigned Machine Name	No
Remote Machine Power Policy	Take no power action
Automatically Logoff After Disconnect	Never
Connection Server Restrictions	None
Category Folder	None

Create RDSH Farm and Pool

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

> Full Clones provisioned with PowerCLI were used for RDS Virtual Machines.

Create RDS Farm

To create the RDSH Farm and Pool, follow these steps:

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.

2. Select Type of the Farm. We used Manual Farm of the previously provisioned virtual machines for the RDS desktops in this design. Click Next.

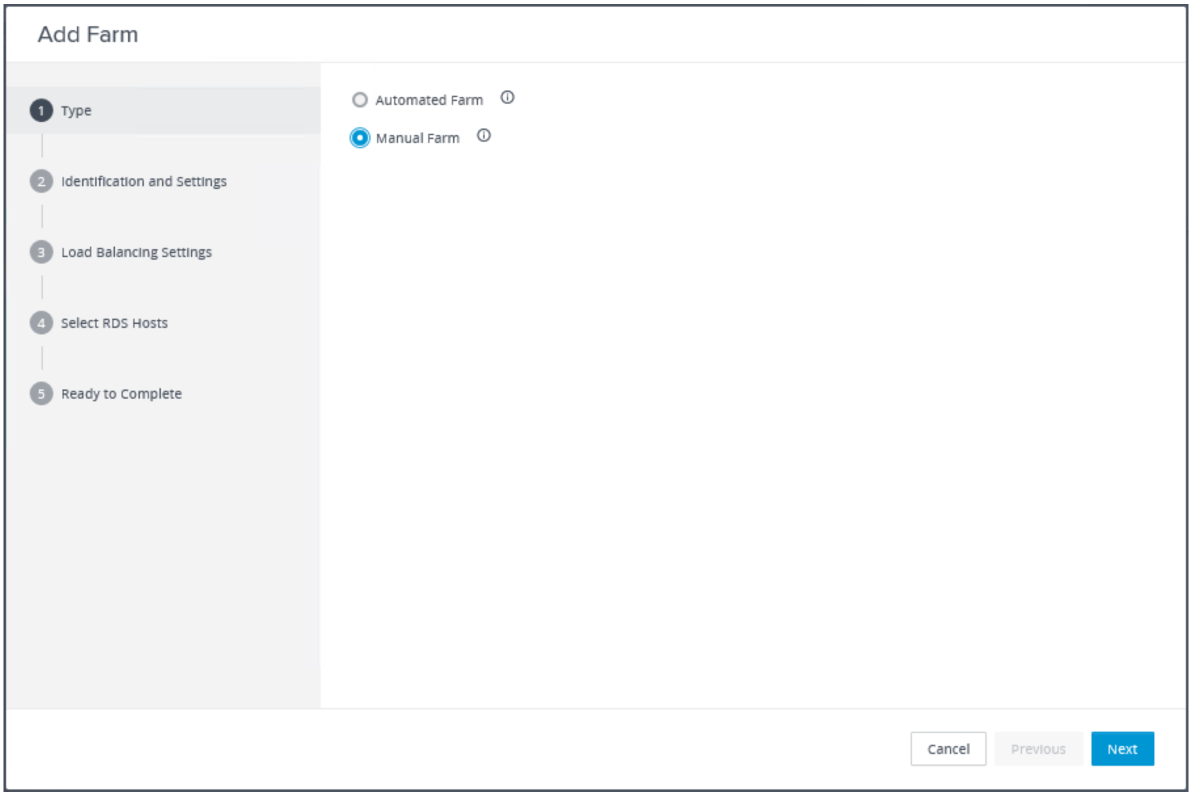

Add Farm

○ Automated Farm

● Manual Farm

3. Provide ID and Description for RDS FARM. Select the Display Protocol which is required for users to connect to the RDS Sessions. Click Next.

RDP Display protocol was used for the RDSH farms in this validated design.

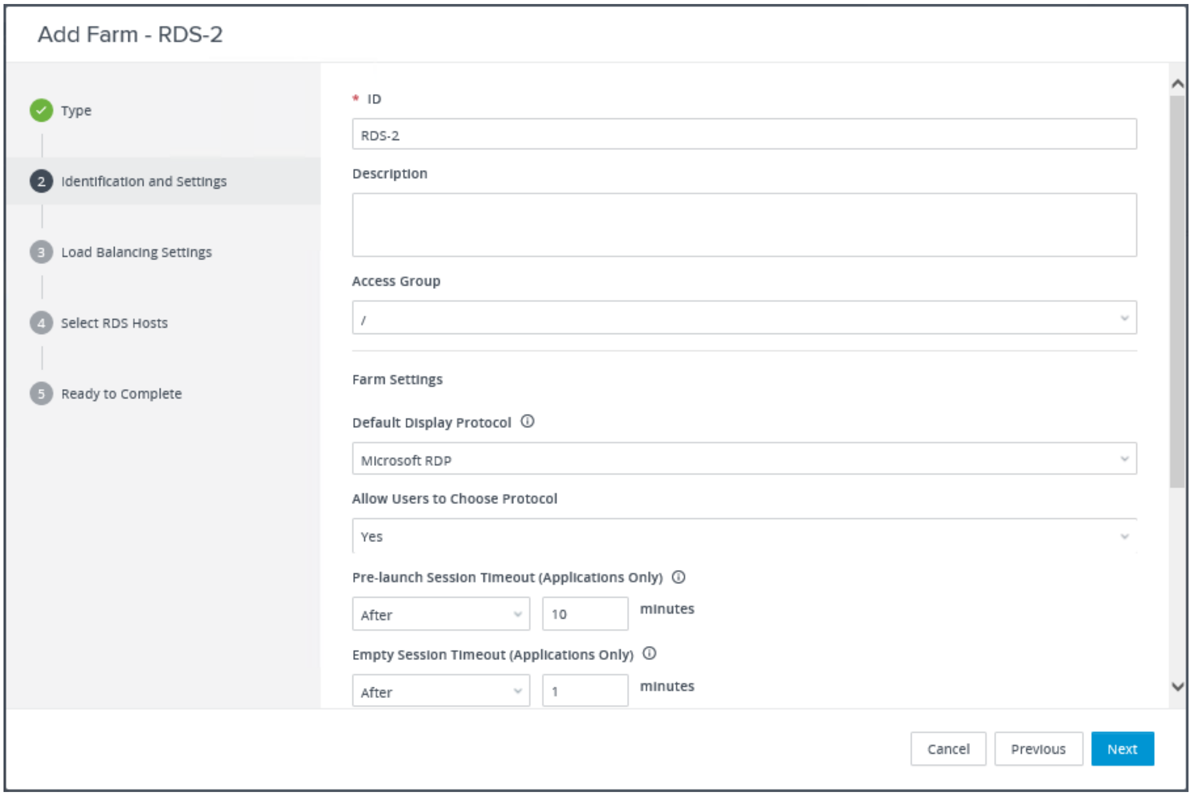

Add Farm - RDS-2

* ID: RDS-2

Description:

Access Group: /

Farm Settings

Default Display Protocol: Microsoft RDP

Allow Users to Choose Protocol: Yes

Pre-launch Session Timeout (Applications Only): After 10 minutes

Empty Session Timeout (Applications Only): After 1 minutes

4. Select Load Balancing Settings. Click Next.

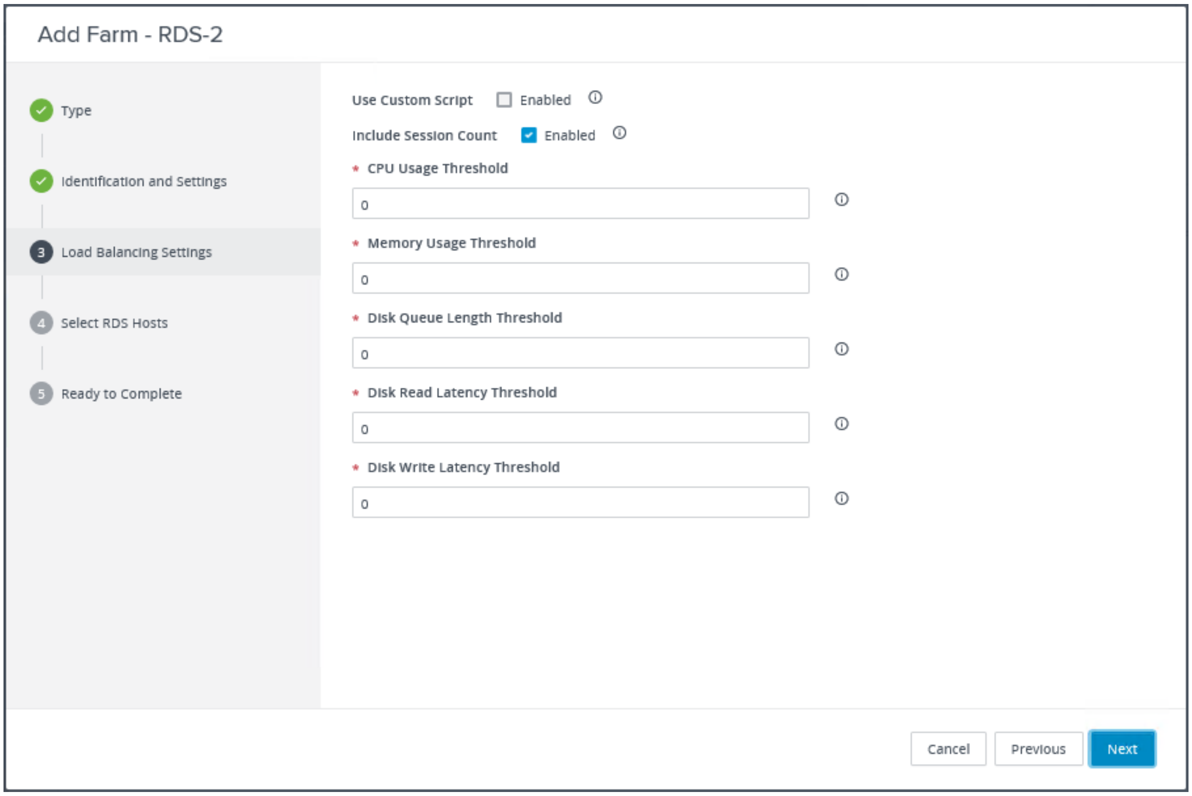

Add Farm - RDS-2

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

5. Select previously created virtual machine to be used as RDS host. Click Next.

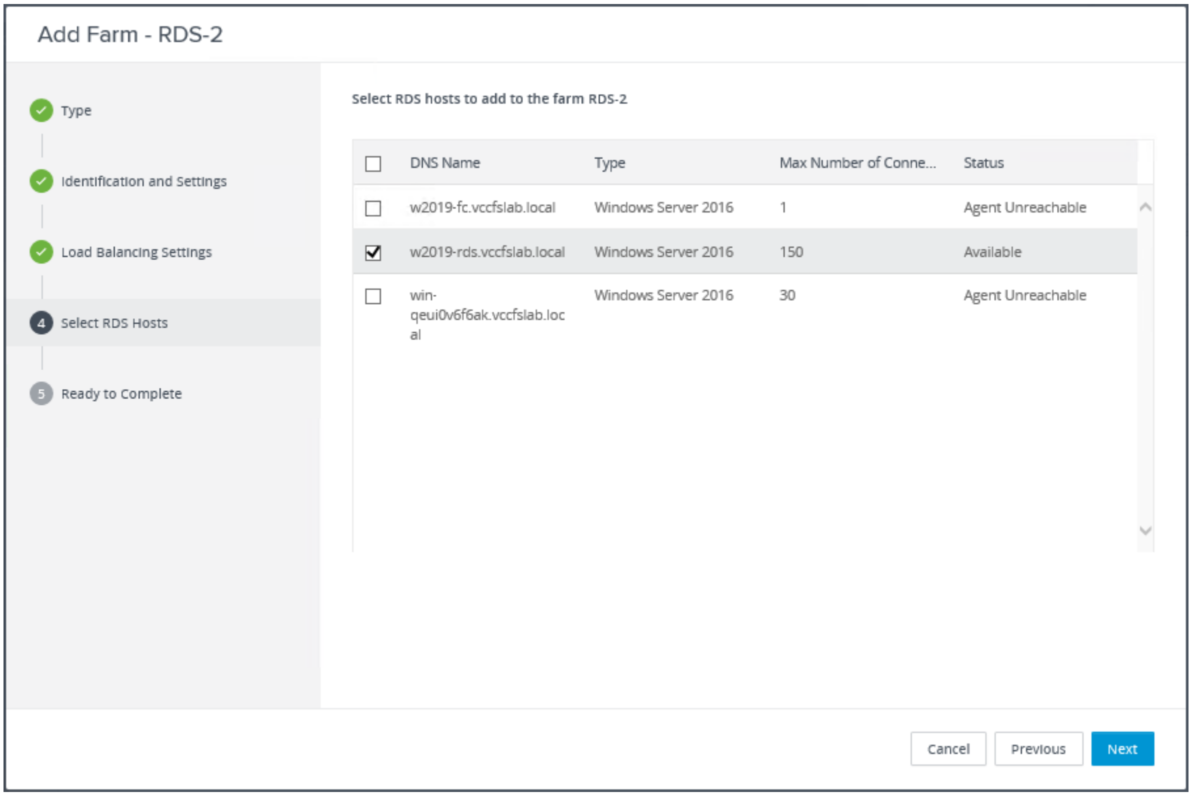

Add Farm - RDS-2

Select RDS hosts to add to the farm RDS-2

□	DNS Name	Type	Max Number of Conne...	Status
□	w2019-fc.vccfslab.local	Windows Server 2016	1	Agent Unreachable
☑	w2019-rds.vccfslab.local	Windows Server 2016	150	Available
□	win-qeui0v6f6ak.vccfslab.local	Windows Server 2016	30	Agent Unreachable

6. Review Farm information and click Submit to complete the RDS Farm creation.

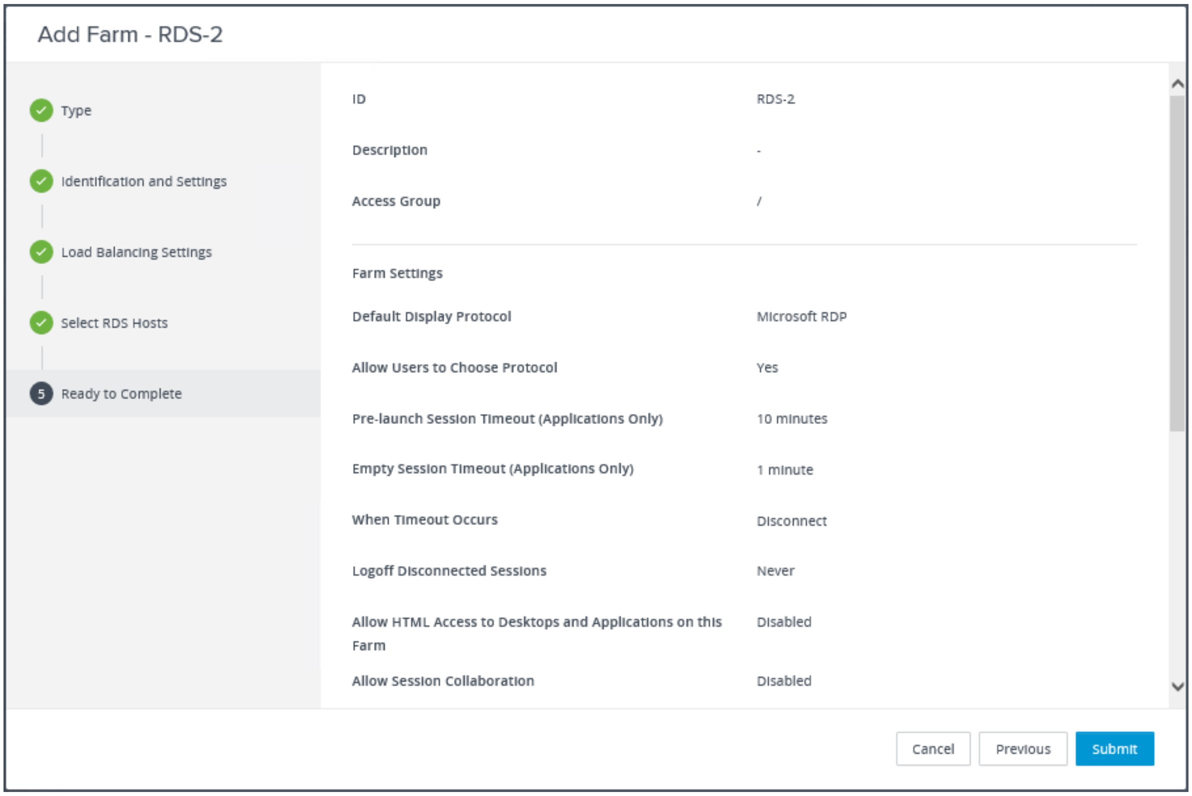

Add Farm - RDS-2

ID	RDS-2
Description	-
Access Group	/
Farm Settings	
Default Display Protocol	Microsoft RDP
Allow Users to Choose Protocol	Yes
Pre-launch Session Timeout (Applications Only)	10 minutes
Empty Session Timeout (Applications Only)	1 minute
When Timeout Occurs	Disconnect
Logoff Disconnected Sessions	Never
Allow HTML Access to Desktops and Applications on this Farm	Disabled
Allow Session Collaboration	Disabled

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. To create c RDS pool, follow these steps:

1. Select type as RDS Desktop Pool.

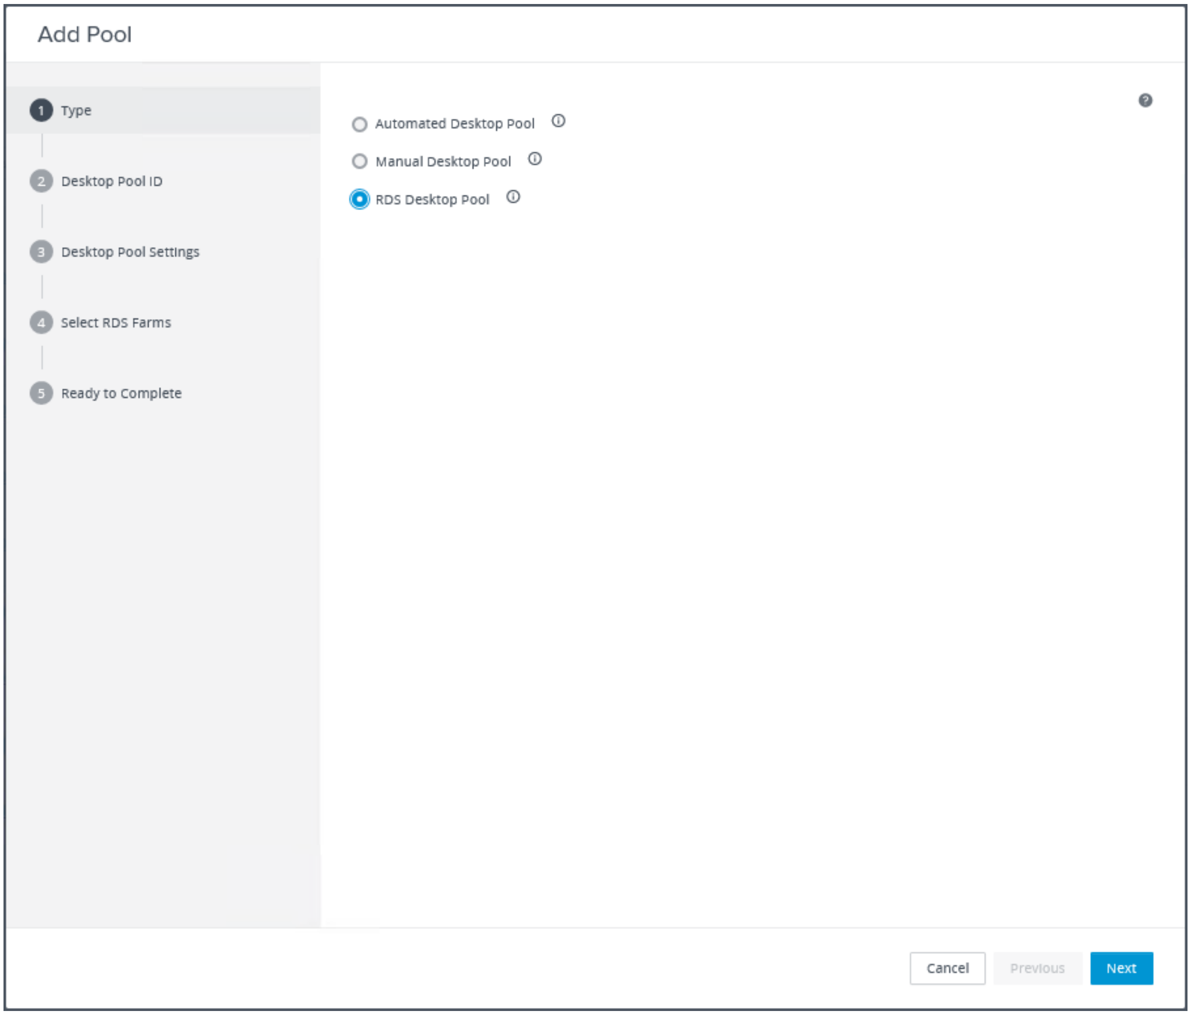
Add Pool

○ Automated Desktop Pool

○ Manual Desktop Pool

● RDS Desktop Pool

2. Provide an ID and Display Name for the Pool. Click Next.

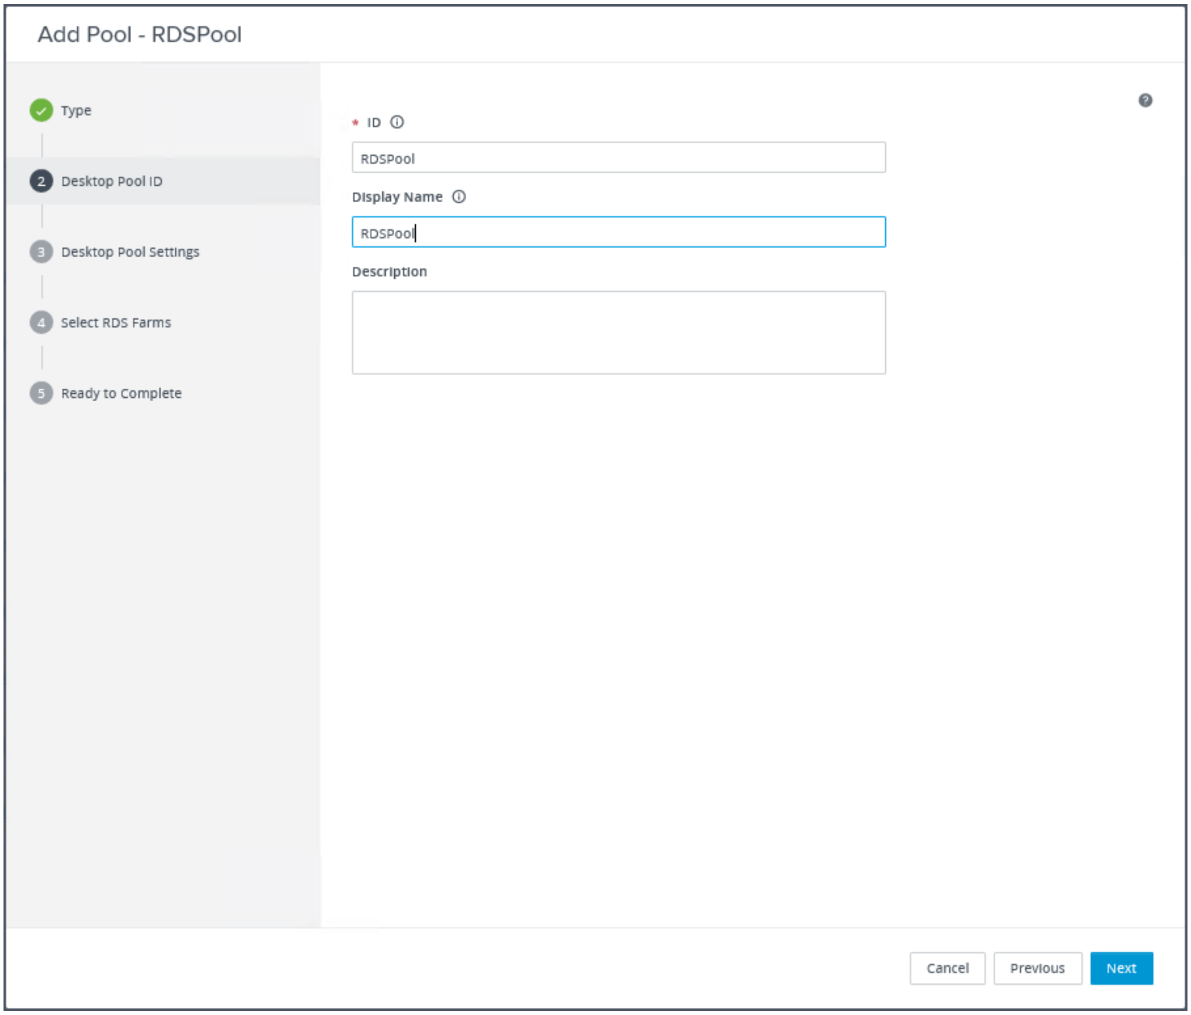

Add Pool - RDSPool

* ID: RDSPool

Display Name: RDSPool

Description:

3. Leave the default settings for the Desktop Pool Settings. Click Next.

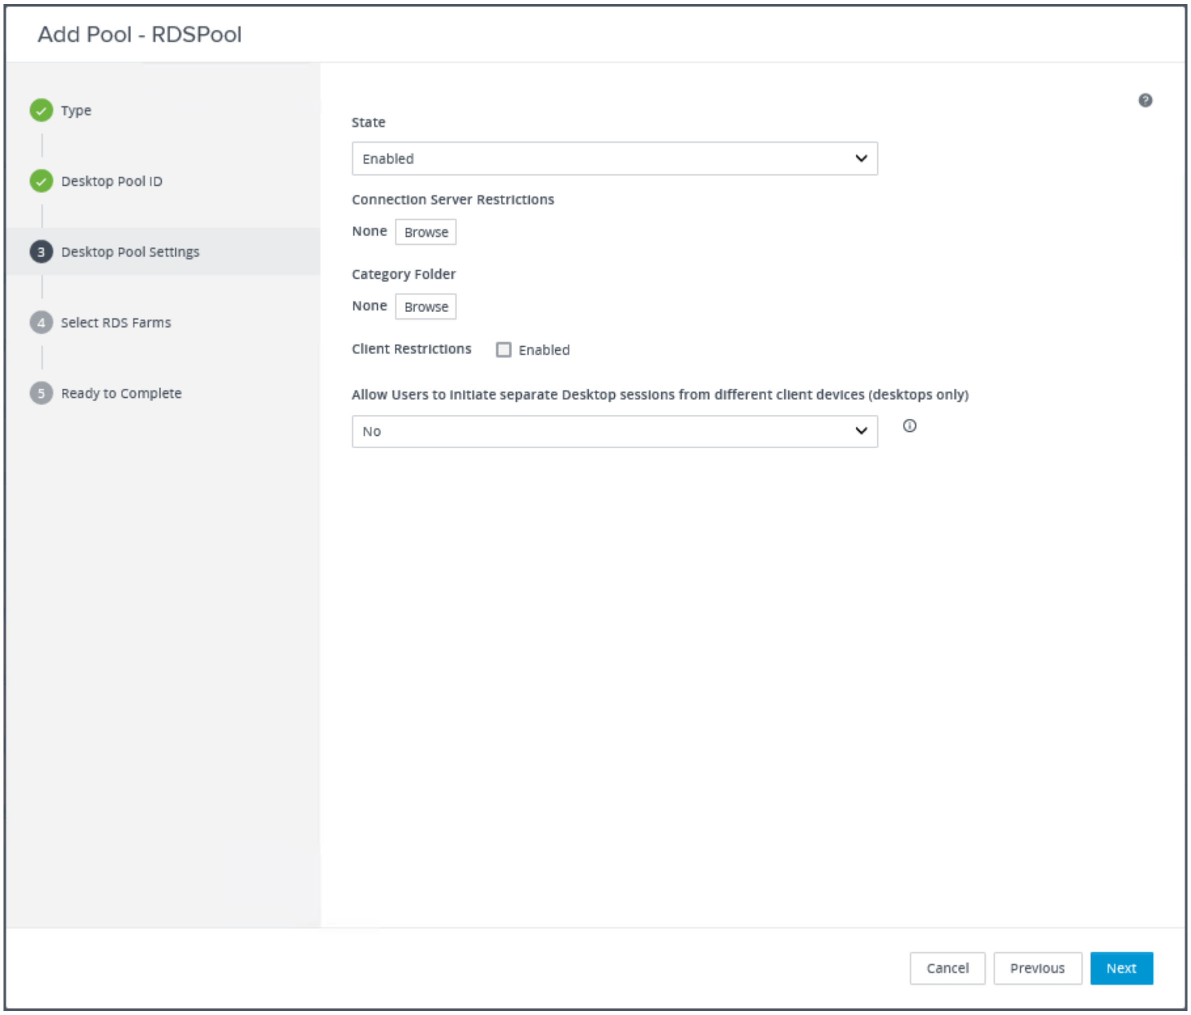

Add Pool - RDSPool

State

Enabled

Connection Server Restrictions

None Browse

Category Folder

None Browse

Client Restrictions □ Enabled

Allow Users to initiate separate Desktop sessions from different client devices (desktops only)

No

4. Select the RDS Farm. Select the farm which was already created for this desktop pool. Click Next.

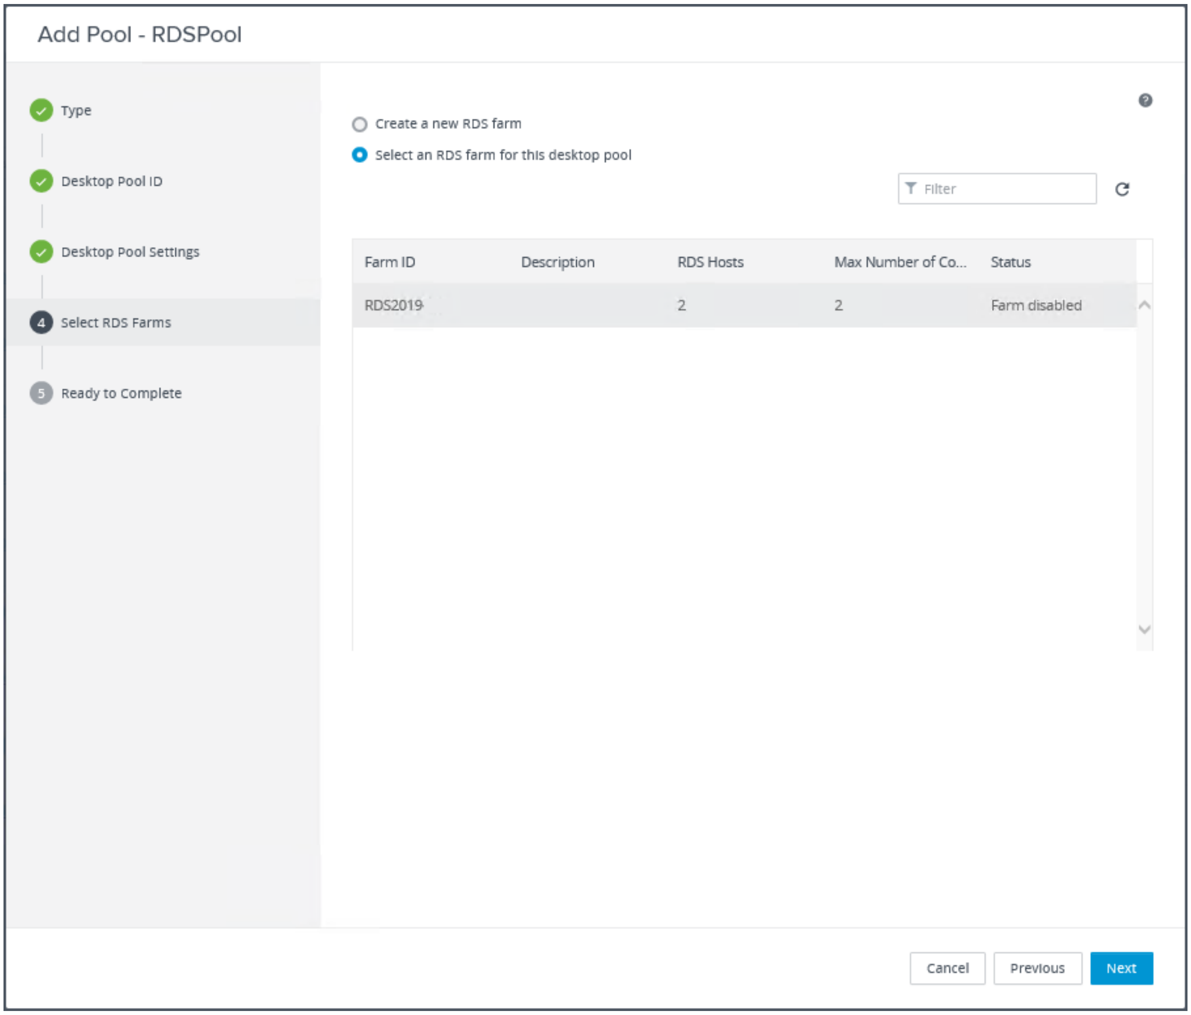

Add Pool - RDSPool

○ Create a new RDS farm

● Select an RDS farm for this desktop pool

Farm ID	Description	RDS Hosts	Max Number of Co...	Status
RDS2019		2	2	Farm disabled

5. Review the RDS Pool deployment specifications and click Finish to complete the RDS pool deployment.

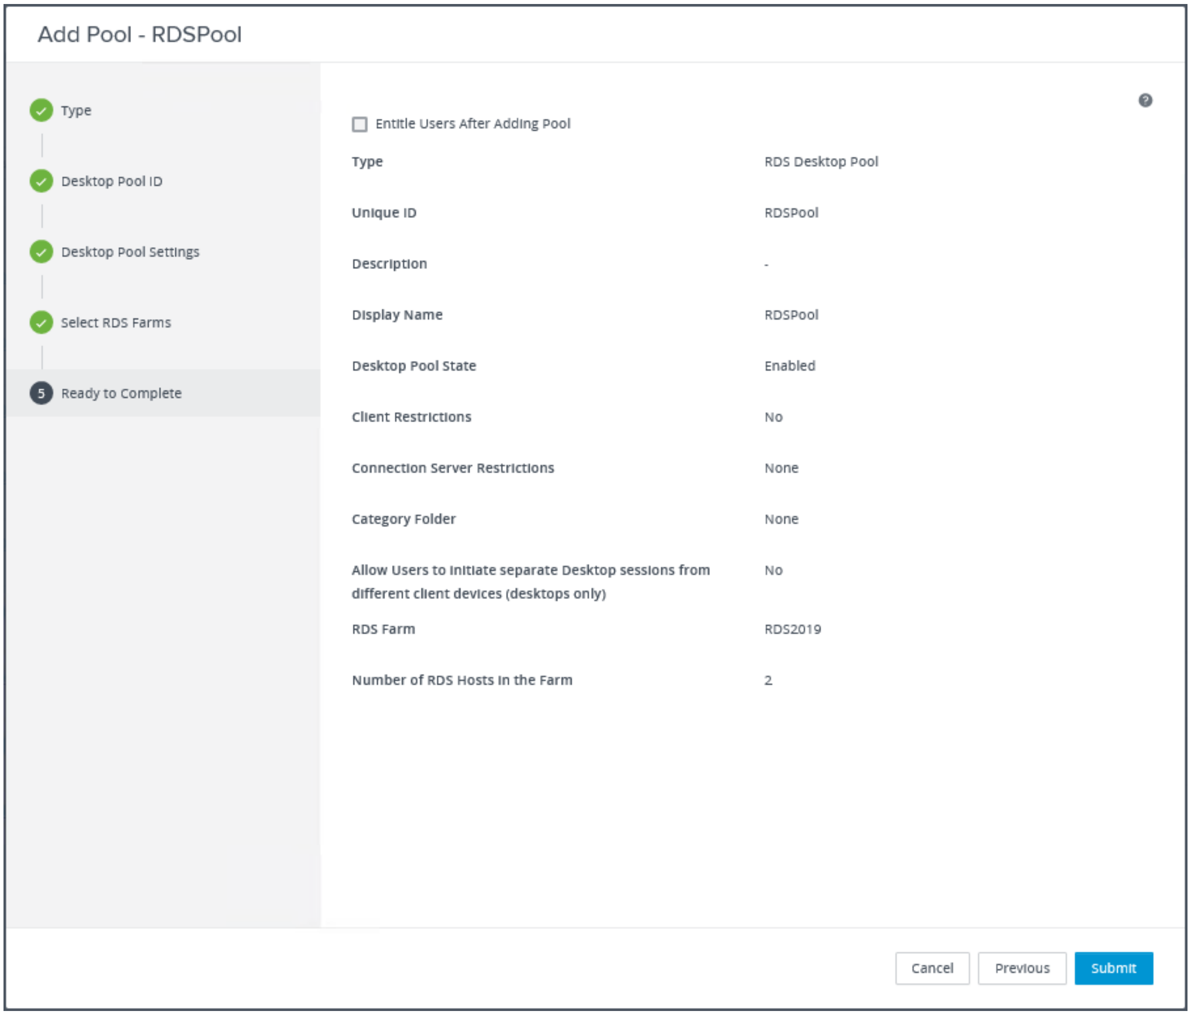

Add Pool - RDSPool

□ Entitle Users After Adding Pool

Type	RDS Desktop Pool
Unique ID	RDSPool
Description	-
Display Name	RDSPool
Desktop Pool State	Enabled
Client Restrictions	No
Connection Server Restrictions	None
Category Folder	None
Allow Users to initiate separate Desktop sessions from different client devices (desktops only)	No
RDS Farm	RDS2019
Number of RDS Hosts in the Farm	2

6. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

Configure 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.

Initial Configuration (Easy Start)

To perform an initial configuration of Dynamic Environment Manager, follow these steps:

1. Launch the DEM Management Console from the Start Menu.

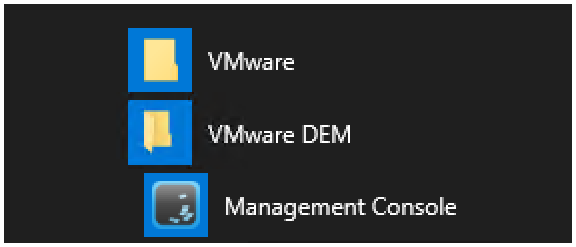

2. Enter the path to the DEM configuration share and click OK.

In this validated design the DEM configuration and user profiles were stored on the share hosted by Pure Storage FlashArray//X70 R3.

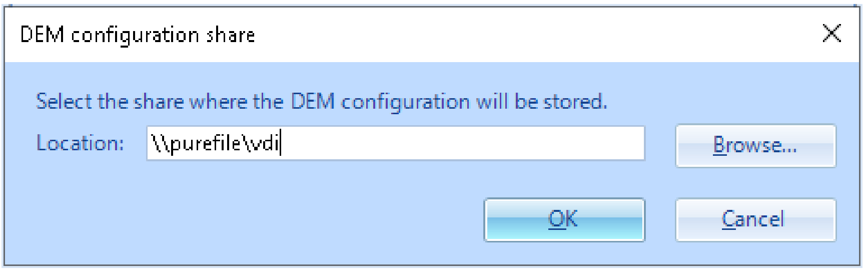
DEM configuration share

Select the share where the DEM configuration will be stored.

Location: \\purefile\vdi

3. Leave the defaults and click OK.

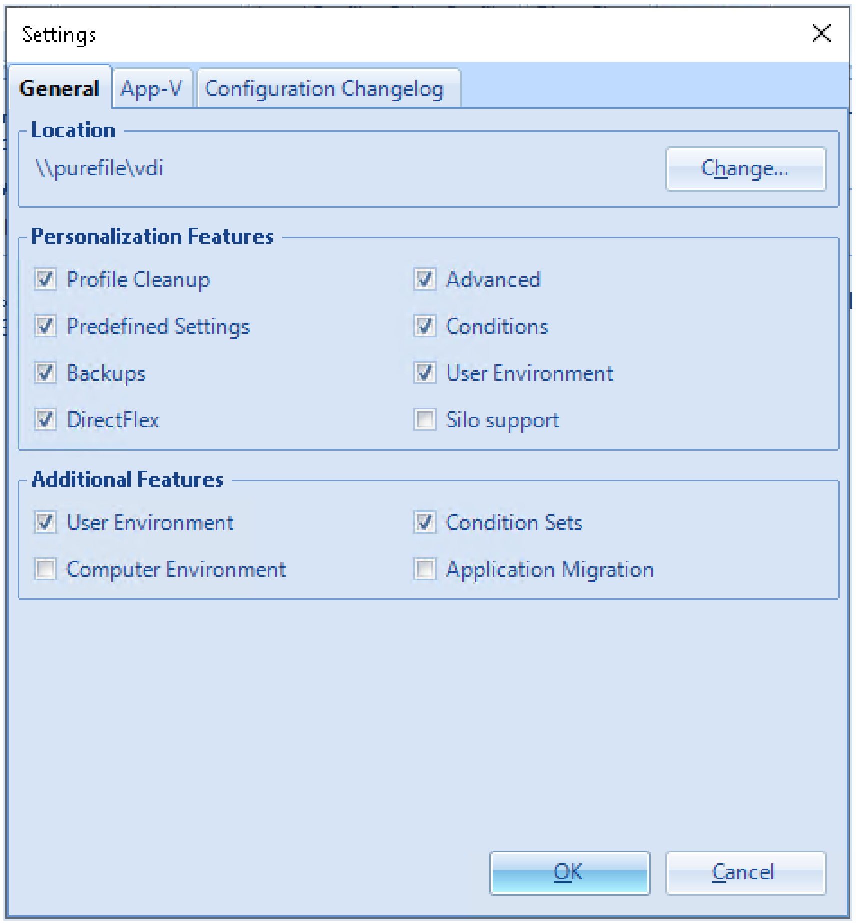

Settings

General | App-V | Configuration Changelog

Location

\\purefile\vdi

Personalization Features

☑ Profile Cleanup
☑ Advanced
☑ Predefined Settings
☑ Conditions
☑ Backups
☑ User Environment
☑ DirectFlex
□ Silo support

Additional Features

☑ User Environment
☑ Condition Sets
□ Computer Environment
□ Application Migration

4. On Personalization tab of the Management Console, click Easy Start.

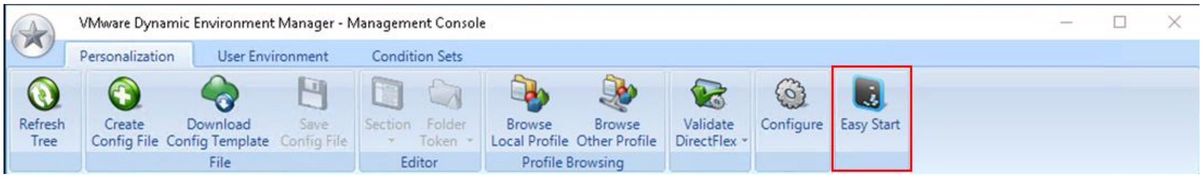

5. Select your Office version and click OK.

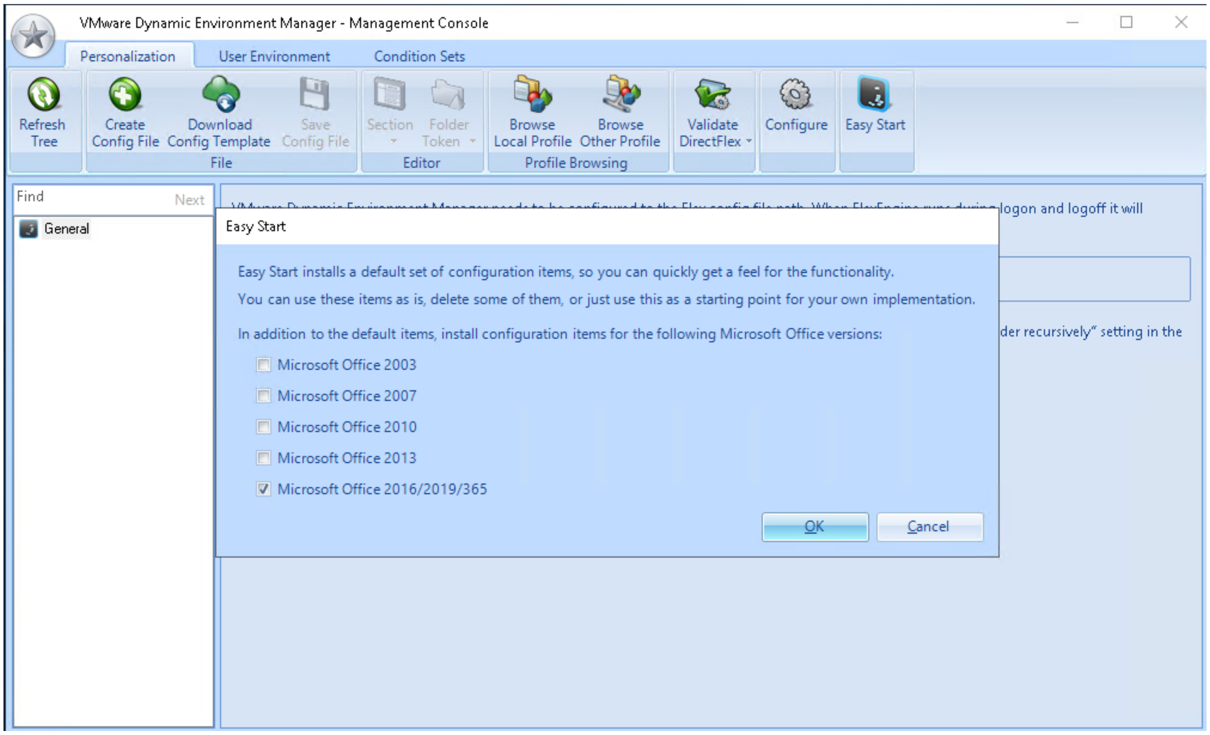

Easy Start

Easy Start installs a default set of configuration items, so you can quickly get a feel for the functionality.
You can use these items as is, delete some of them, or just use this as a starting point for your own implementation.

In addition to the default items, install configuration items for the following Microsoft Office versions:

☐ Microsoft Office 2003
☐ Microsoft Office 2007
☐ Microsoft Office 2010
☐ Microsoft Office 2013
☑ Microsoft Office 2016/2019/365

6. When configuration items successfully installed click OK.

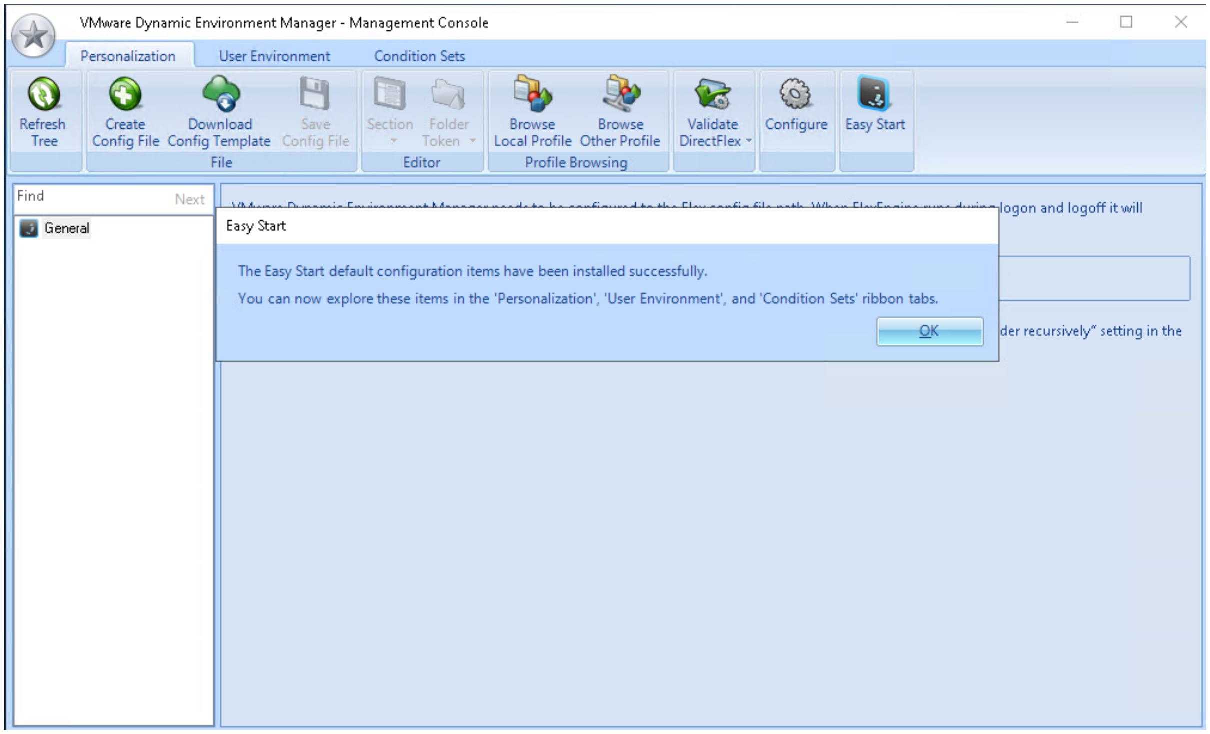

Easy Start

The Easy Start default configuration items have been installed successfully.
You can now explore these items in the 'Personalization', 'User Environment', and 'Condition Sets' ribbon tabs.

7. Click Personalization tab, click + Create Config File, select Use a Windows Common Setting and click Next.

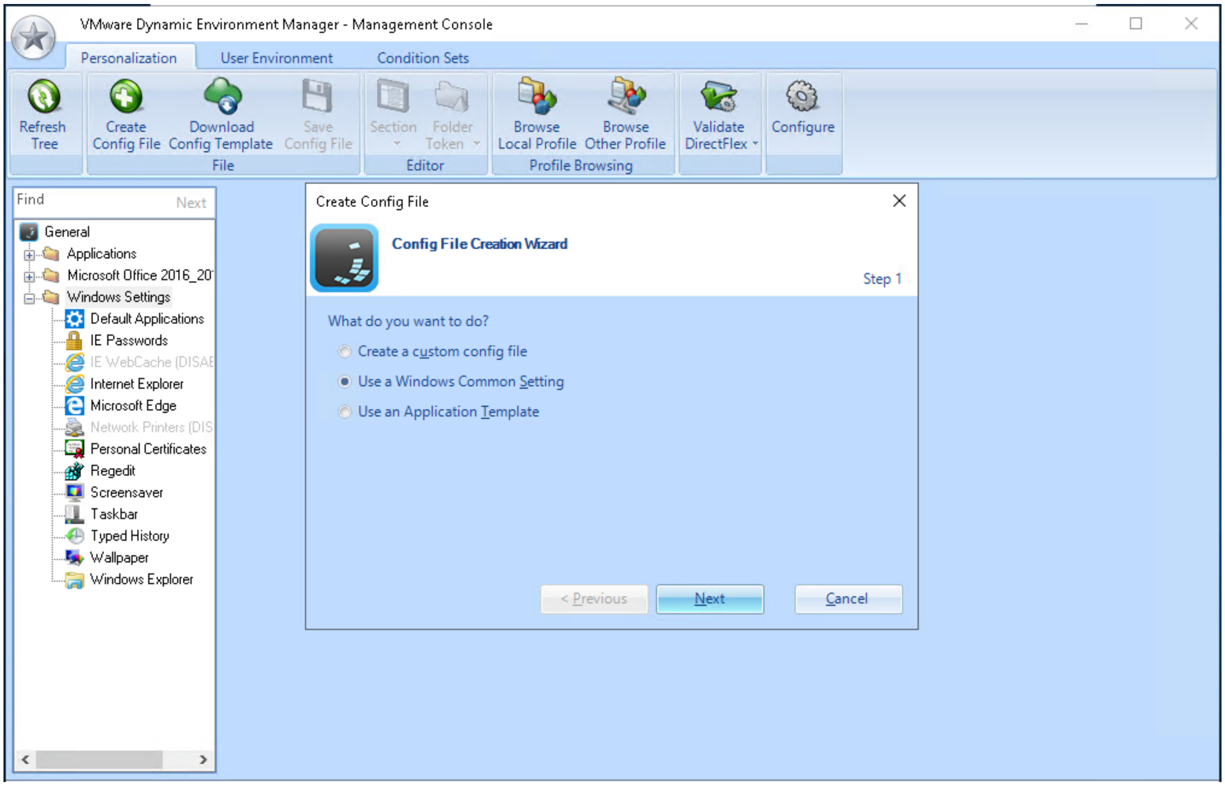

Create Config File

Config File Creation Wizard

Step 1

What do you want to do?

○ Create a custom config file

● Use a Windows Common Setting

○ Use an Application Template

8. Select Windows 10 Start Menu – Windows 10 Version 1703 and higher and click Next.

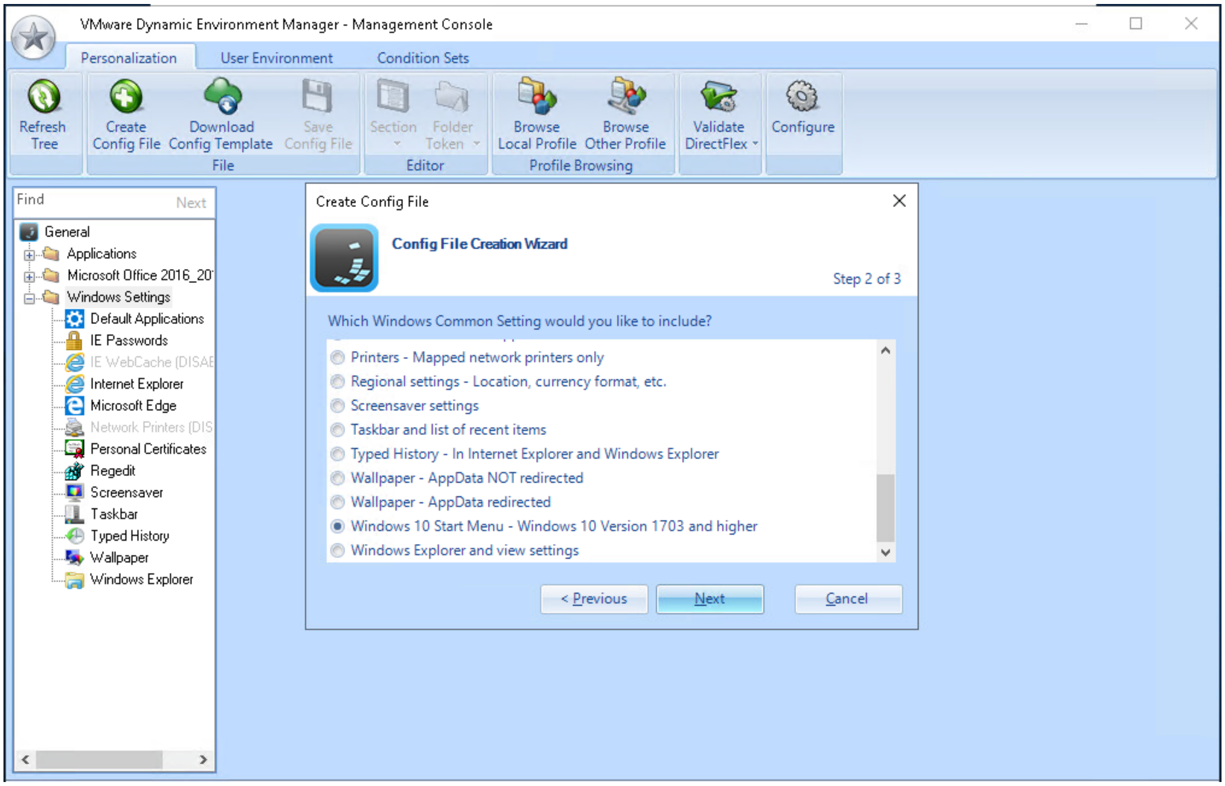

Create Config File

Config File Creation Wizard

Step 2 of 3

Which Windows Common Setting would you like to include?

○ Printers - Mapped network printers only

○ Regional settings - Location, currency format, etc.

○ Screensaver settings

○ Taskbar and list of recent items

○ Typed History - In Internet Explorer and Windows Explorer

○ Wallpaper - AppData NOT redirected

○ Wallpaper - AppData redirected

● Windows 10 Start Menu - Windows 10 Version 1703 and higher

○ Windows Explorer and view settings

9. Enter a file name and click Finish to create configuration.

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

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 45 illustrates the installed NVIDIA P6 GPU in the front and rear mezzanine slots.

Figure 45. NVIDIA GPU Installed in the Front and Rear Mezzanine Slots

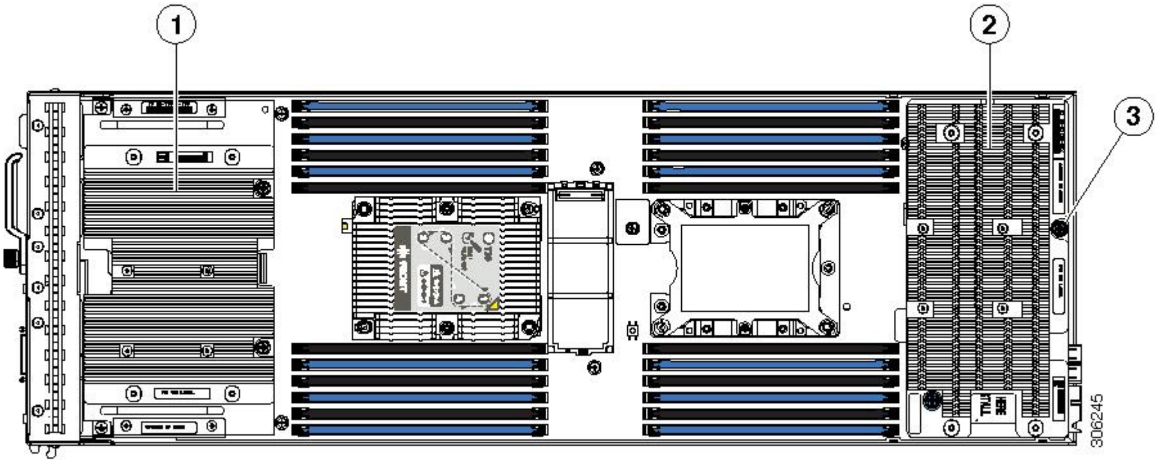

1	Front GPU	2	Rear GPU
3	Custom standoff screw	-	

Install an NVIDIA GPU Card in the Front of the Server

Figure 46 illustrates the front NVIDIA P6 GPU (UCSB-GPU-P6-F).

Figure 46. NVIDIA P6 GPU That Installs in the Front of the Server

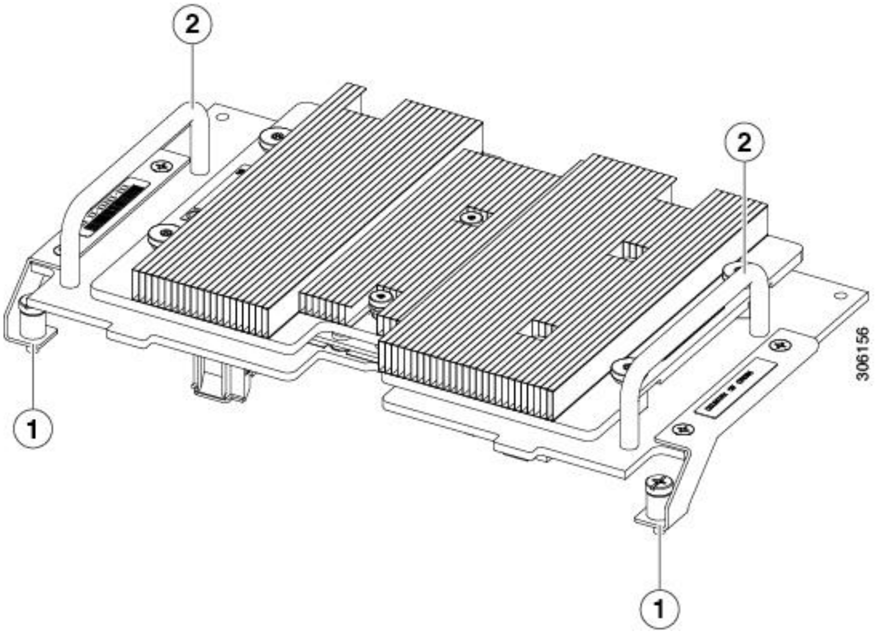

1	Leg with thumb screw that attaches to the server motherboard at the front	2	Handle to press down on when installing the GPU

Figure 47. Top-Down View of the NVIDIA P6 GPU for the Front of the Server

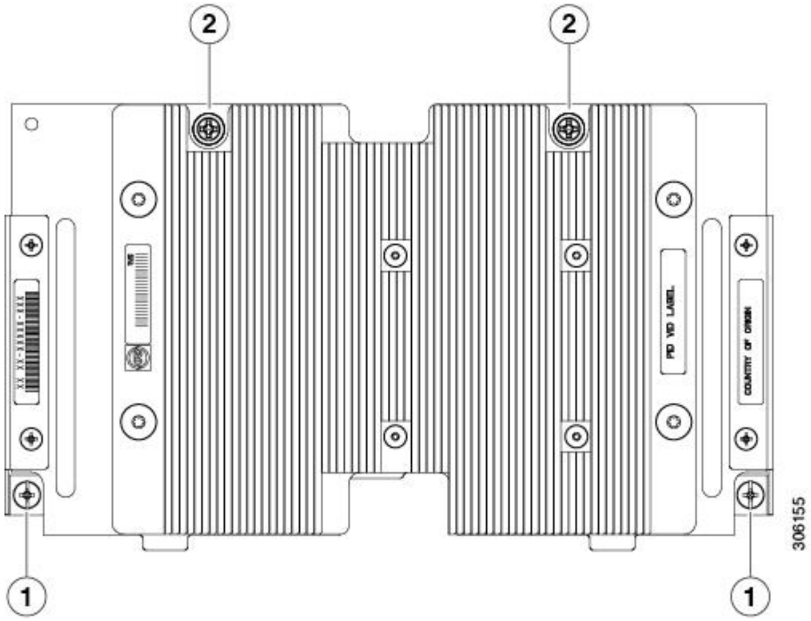

1	Leg with thumb screw that attaches to the server motherboard	2	Thumb screw that attaches to a standoff below

To install the NVIDIA GPU, follow these steps:

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 at the following URL for information about 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.

1. Position the GPU in the correct orientation to the front of the server (callout 1) as shown in Figure 48.

2. Install the GPU into the server. Press down on the handles (callout 5) to firmly secure the GPU.

3. Tighten the thumb screws (callout 3) at the back of the GPU with the standoffs (callout 4) on the motherboard.

4. Tighten the thumb screws on the legs (callout 2) to the motherboard.

5. Install the drive blanking panels.

Figure 48. Installing the NVIDIA GPU in the Front of the Server

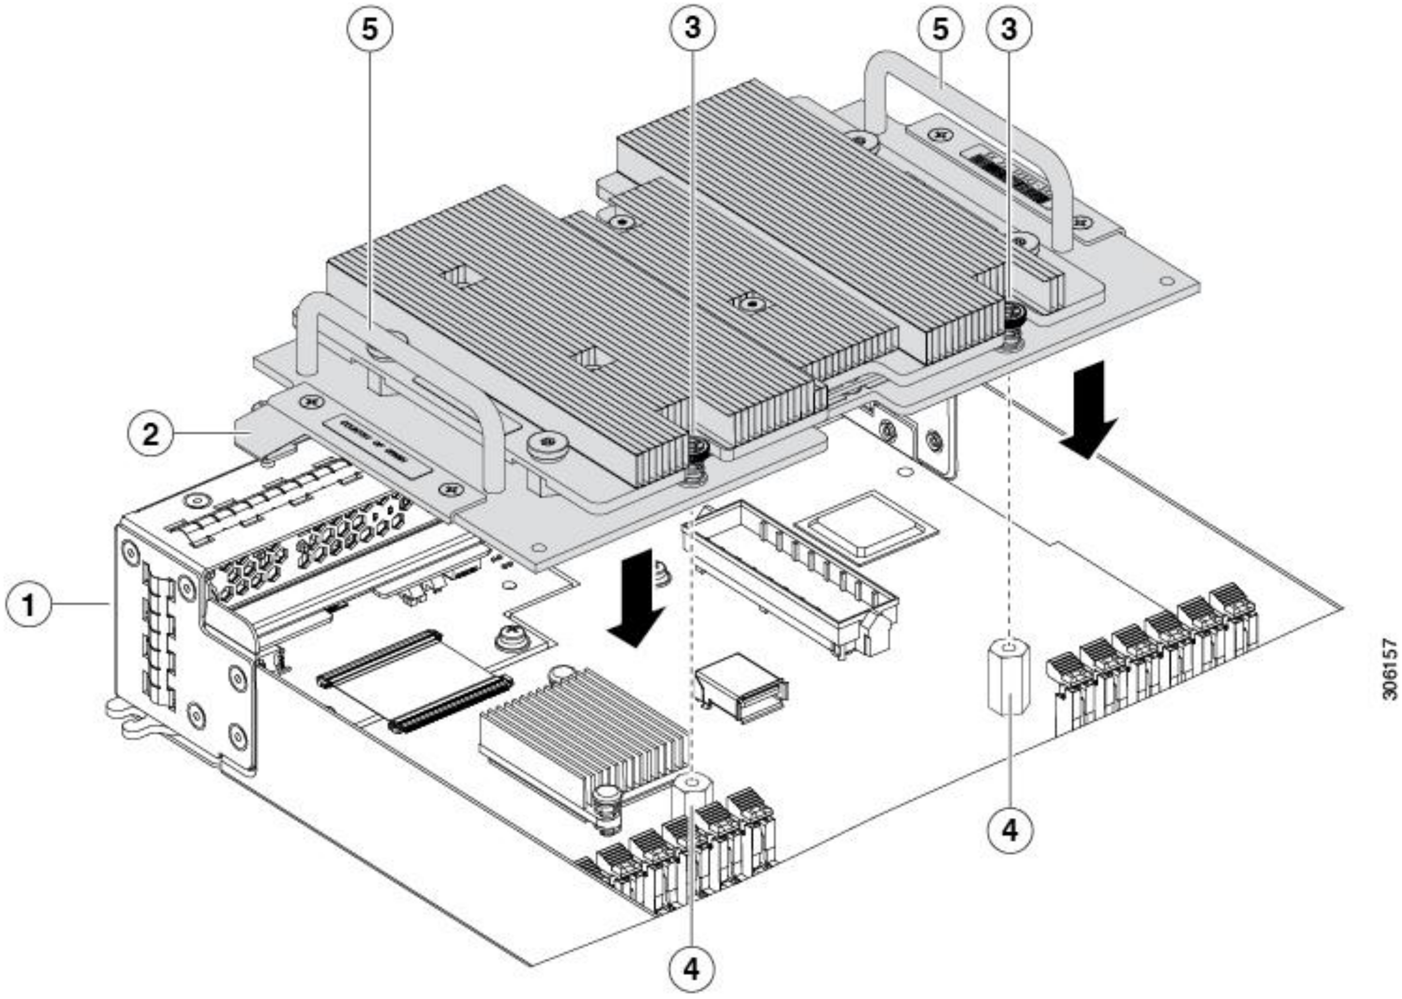

1	Front of the server	2	Leg with thumb screw that attaches to the motherboard
3	Thumbscrew to attach to standoff below	4	Standoff on the motherboard
5	Handle to press down on to firmly install the GPU	–	

Install 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 motherboard. Figure 49 shows the three components of the option kit.

Figure 49. NVIDIA P6 GPU (UCSB-GPU-P6-R) Option Kit

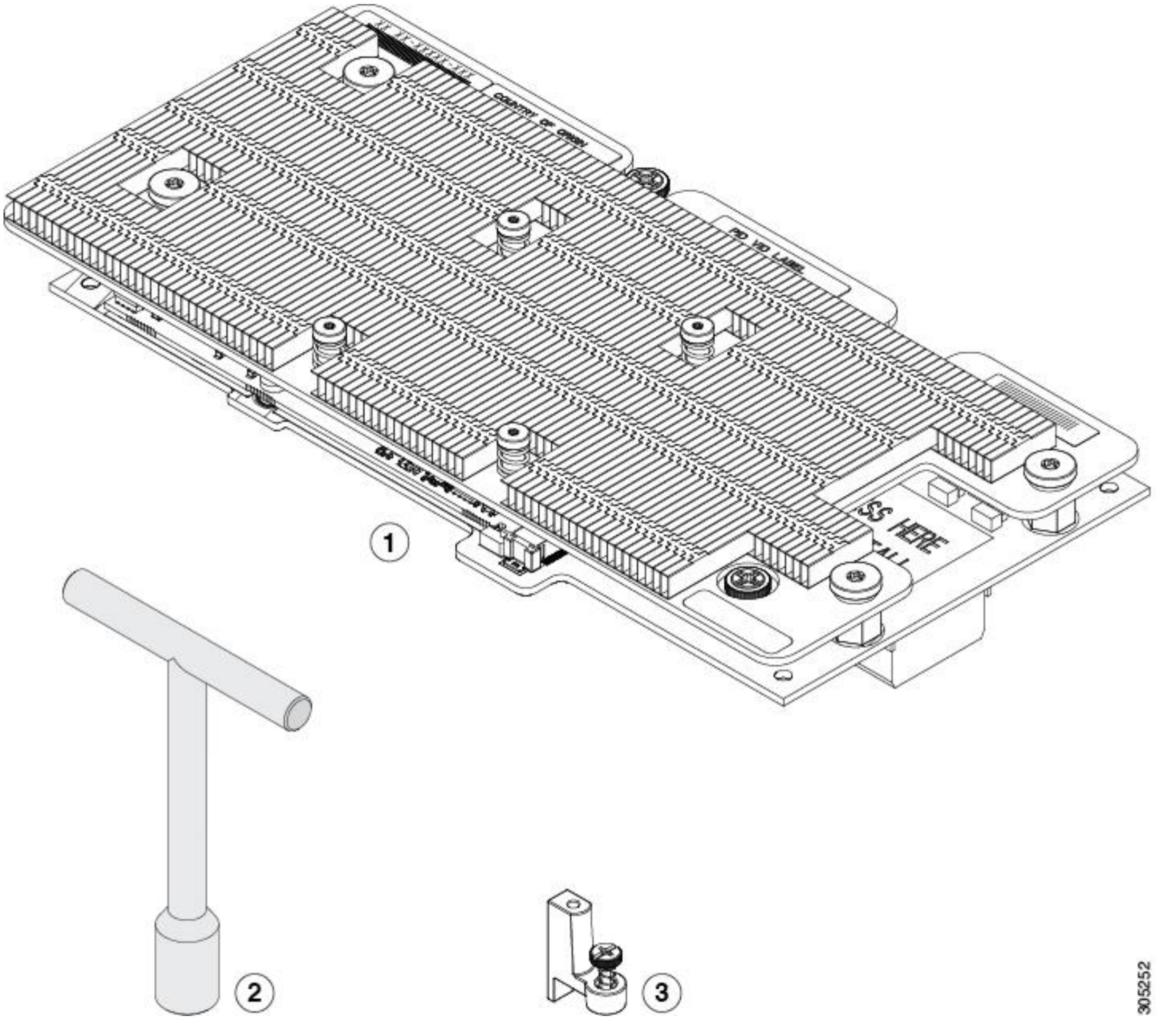

1	NVIDIA P6 GPU (CPU and heatsink)	2	T-shaped wrench
3	Custom standoff	-	

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* at the following URL for information about 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.

To install an NVIDIA GPU Card in the rear of the server, follow these steps:

1. Use the T-shaped wrench that comes with the GPU to remove the existing standoff at the back end of the motherboard.

2. Install the custom standoff in the same location at the back end of the motherboard.

3. Position the GPU over the connector on the motherboard and align all the captive screws to the standoff posts (callout 1).

4. Tighten the captive screws (callout 2).

Figure 50. Installing the NVIDIA P6 GPU in the Rear Mezzanine Slot

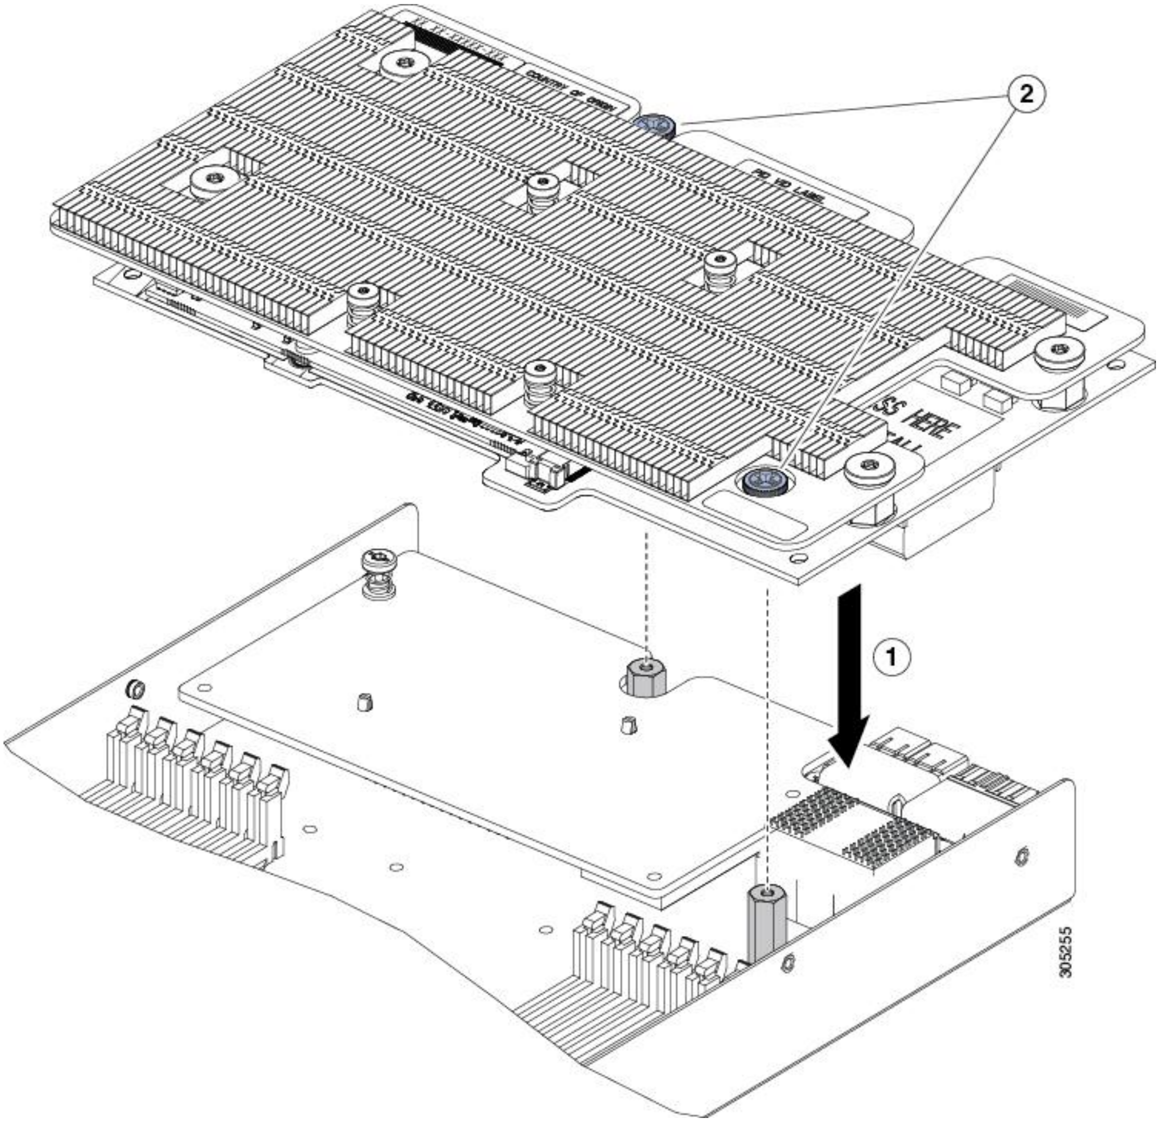

Install the NVIDIA VMware VIB Driver

To install the NVIDIA VMware VIB driver, follow these steps:

1. From the Cisco UCS Manager, verify the GPU card has been properly installed.

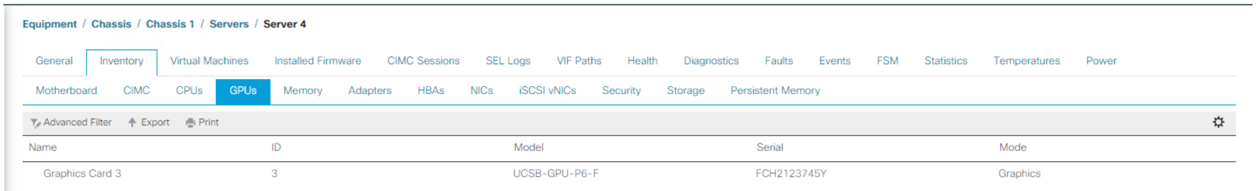

Name	ID	Model	Serial	Mode
Graphics Card 3	3	UCSB-GPU-P6-F	FCH2123745Y	Graphics

2. Download the NVIDIA GRID GPU driver pack for VMware vSphere ESXi 7.0.

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.)

4. Log in as root to the vSphere console through SSH using a tool such as Putty.

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.

5. 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_bootbank_NVIDIA-
VMware_ESXi_7.0_Host_Driver_450.102-1OEM.700.0.0.15525992.vib

Installation Result

   Message: Operation finished successfully.

   Reboot Required: false

   VIBs Installed: NVIDIA_bootbank_NVIDIA-VMware_ESXi_7.0_Host_Driver_450.102-1OEM.700.0.0.15525992

   VIBs Removed:

   VIBs Skipped:
```

VIBs Skipped: although the display shows “Reboot Required: false,” a reboot is necessary for the VIB file to load and for xorg to start.

6. 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
```

7. 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_7.0_Host_Driver  450.102-1OEM.700.0.0.15525992      NVIDIA  VMwareAccepted
   2021-01-14
```

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.

8. 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 51, depending on the card used in your environment.

Figure 51. VMware ESX SSH Console Report for GPU P6 Card Detection on Cisco UCS B200 M5 Blade Server

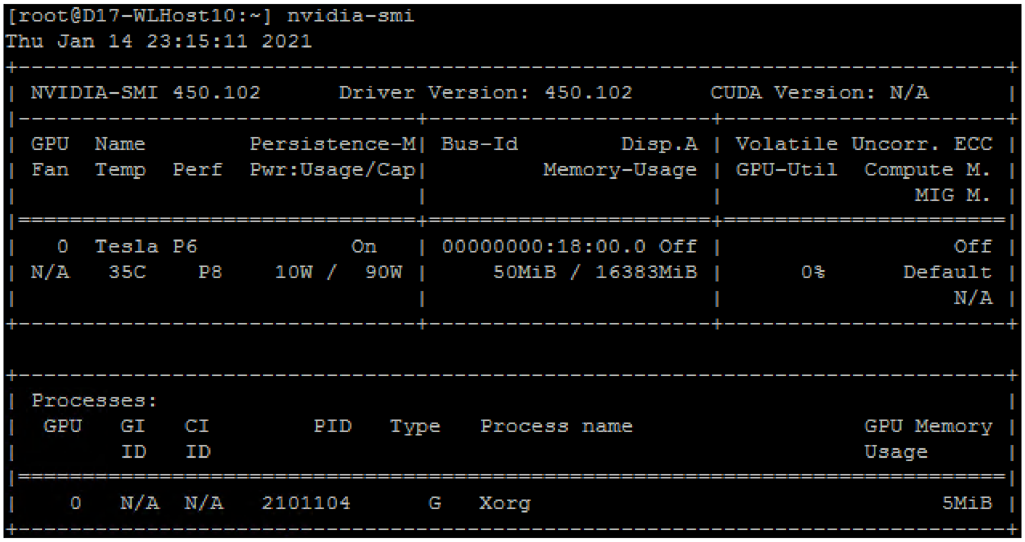

```
[root@D17-WLHost10:~] nvidia-smi
Thu Jan 14 23:15:11 2021
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 450.102      Driver Version: 450.102      CUDA Version: N/A      |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla P6            On   | 00000000:18:00.0 Off |                  Off |
| N/A   35C    P8    10W /  90W |     50MiB / 16383MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                                  |
|  GPU   GI   CI        PID   Type   Process name                  GPU Memory |
|        ID   ID                                                   Usage      |
|=============================================================================|
|    0   N/A  N/A   2101104      G   Xorg                                5MiB |
+-----------------------------------------------------------------------------+
```

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).

Configure a Virtual Machine with a vGPU

To create the virtual machine that you will use as the VDI base image, follow these steps:

1. Using the vSphere Client, select the ESXi host and click the Configure tab. From the list of options at the left, choose Hardware > Graphics > Host Graphics. Click Edit. Select Shared Direct "Vendor shared passthrough graphics" and click OK. Reboot the system to make the changes effective.

Figure 52. **Edit Host Graphics Settings**

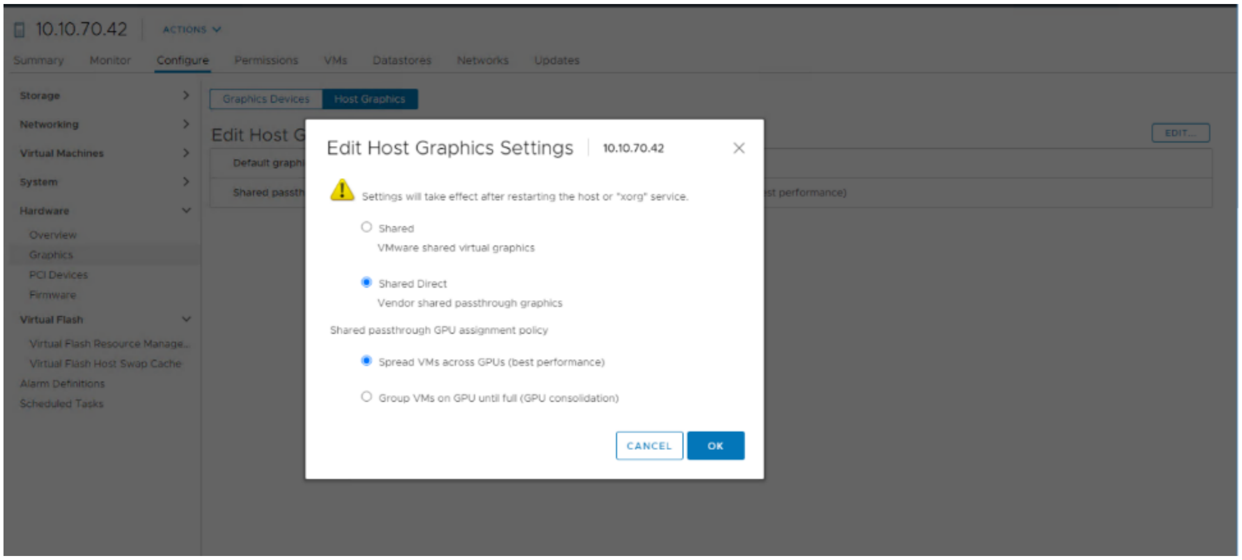

Edit Host Graphics Settings | 10.10.70.42

Settings will take effect after restarting the host or "xorg" service.

○ Shared
VMware shared virtual graphics

● Shared Direct
Vendor shared passthrough graphics

Shared passthrough GPU assignment policy

● Spread VMs across GPUs (best performance)

○ Group VMs on GPU until full (GPU consolidation)

2. Using the vSphere 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 53. **New Virtual Machine Wizard in VMware vSphere Client**

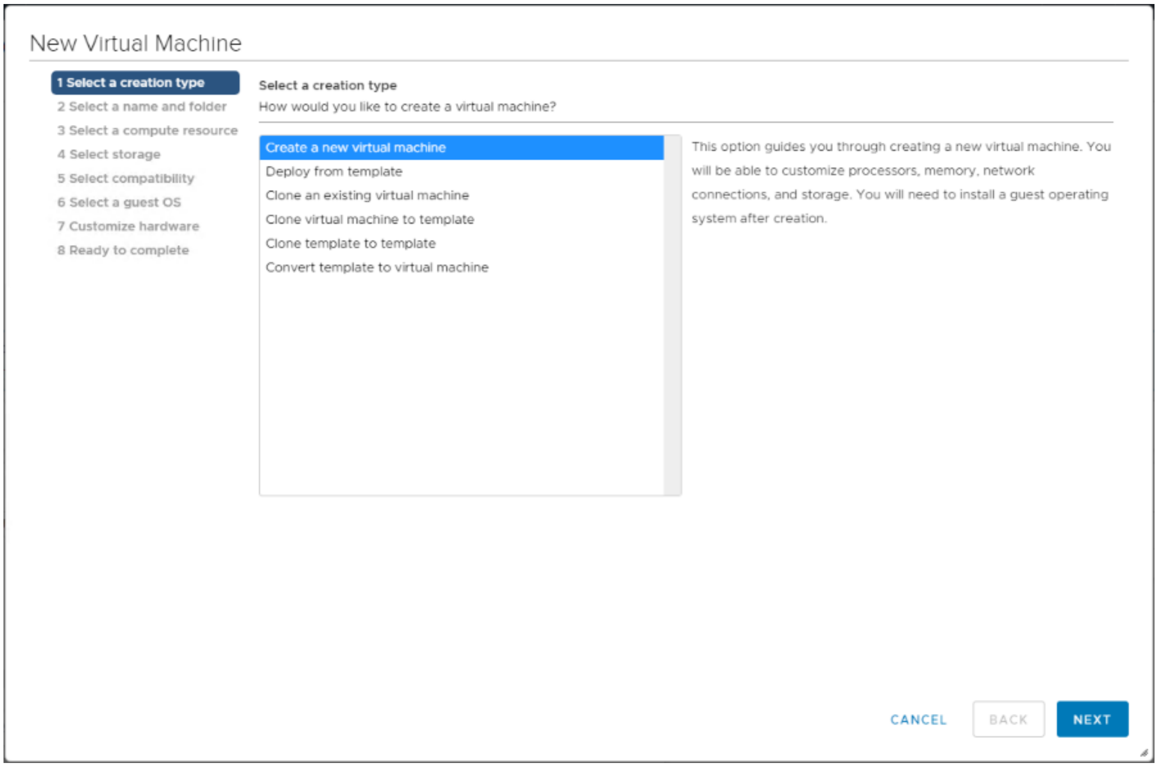

New Virtual Machine

1 Select a creation type
2 Select a name and folder
3 Select a compute resource
4 Select storage
5 Select compatibility
6 Select a guest OS
7 Customize hardware
8 Ready to complete

Select a creation type

How would you like to create a virtual machine?

Create a new virtual machine
Deploy from template
Clone an existing virtual machine
Clone virtual machine to template
Clone template to template
Convert template to virtual machine

This option guides you through creating a new virtual machine. You will be able to customize processors, memory, network connections, and storage. You will need to install a guest operating system after creation.

3. Choose “ESXi 6.0 and later” from the “Compatible with” drop-down list to use the latest features, including the mapping of shared PCI devices, which is required for the vGPU feature.

This solution uses “ESXi 7.0 and later,” which provides the latest features available in ESXi 7.0 and virtual machine hardware Version 17.

Figure 54. Selecting Virtual Machine Hardware Version 17

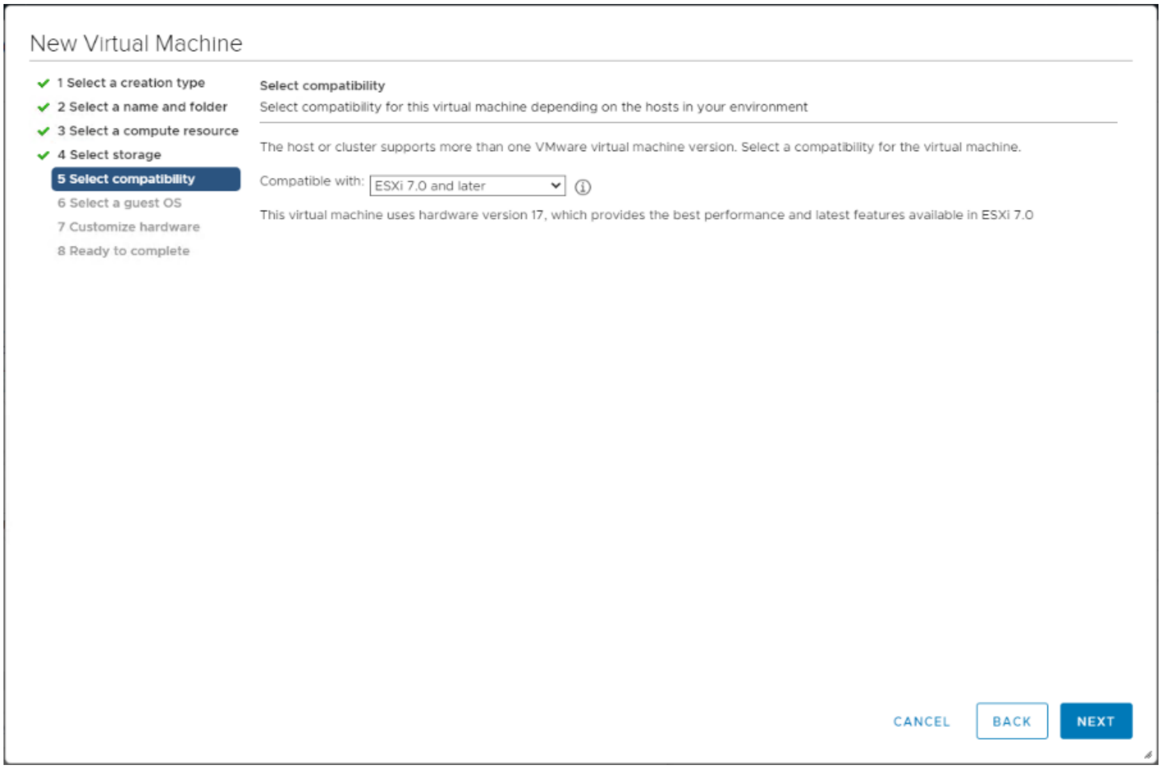

New Virtual Machine

Select compatibility

Select compatibility for this virtual machine depending on the hosts in your environment

The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.

Compatible with: ESXi 7.0 and later

This virtual machine uses hardware version 17, which provides the best performance and latest features available in ESXi 7.0

4. To customize the hardware of the new virtual machine, add a new PCI device, select the appropriate GPU profile, and reserve all virtual machine memory.

If you are creating a new virtual machine and using the vSphere 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 Client to connect to the virtual machine, do not enable the NVIDIA GRID vGPU at this time.

Figure 55. Adding a PCI Device to the Virtual Machine to attach the GPU Profile

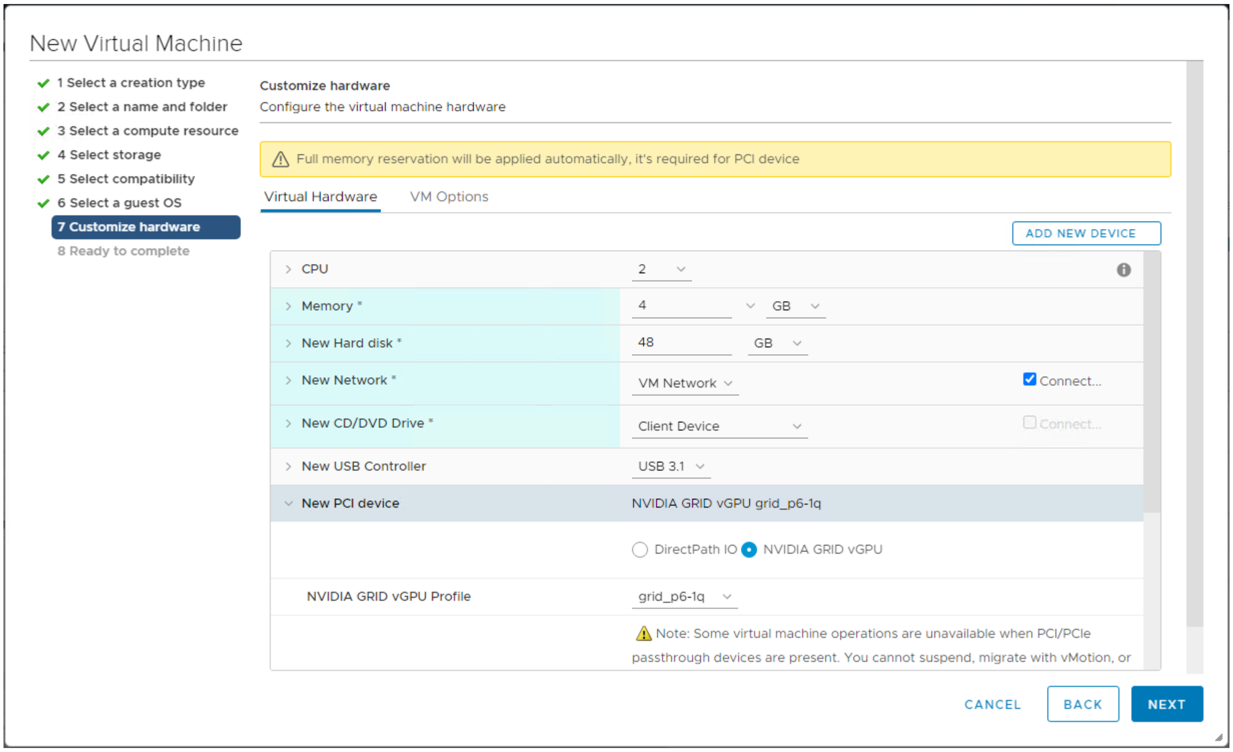

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- ✓ 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Customize hardware

Configure the virtual machine hardware

⚠ Full memory reservation will be applied automatically, it's required for PCI device

Virtual Hardware | VM Options

ADD NEW DEVICE

CPU	2
Memory *	4 GB
New Hard disk *	48 GB
New Network *	VM Network ☑ Connect...
New CD/DVD Drive *	Client Device □ Connect...
New USB Controller	USB 3.1
New PCI device	NVIDIA GRID vGPU grid_p6-1q
	○ DirectPath IO ● NVIDIA GRID vGPU
NVIDIA GRID vGPU Profile	grid_p6-1q
	⚠ Note: Some virtual machine operations are unavailable when PCI/PCIe passthrough devices are present. You cannot suspend, migrate with vMotion, or

CANCEL | BACK | NEXT

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 -e 0
```

Use **-i** to target a specific GPU. If two cards are installed in a server, run the command twice as shown in the example here, where **0** and **1** each specify a GPU card.

6. Install and configure Microsoft Windows on the virtual machine:

   a. Configure the virtual machine with the appropriate amount of vCPU and RAM according to the GPU profile selected.

   b. Install VMware Tools.

   c. Join the virtual machine to the Microsoft Active Directory domain.

   d. Choose "Allow remote connections to this computer" on the Windows System Properties menu.

   e. Install VMware Horizon Agent with appropriate settings. Enable the remote desktop capability if prompted to do so.

   f. Install Horizon Direct Connection agent.

   g. Optimize the Windows OS. VMware OSOT, the optimization tool, includes customizable templates to enable or disable Windows system services and features using VMware recommendations and best practices across multiple systems. Because most Windows system services are enabled by default, the optimization tool can be used to easily disable unnecessary services and features to improve performance.

h. Restart the Windows OS when prompted to do so.

Install the GPU Drivers Inside Windows Virtual Machine

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 vib, you must do the same with the NVIDIA driver in your Windows master image.

In this design we used ESXi Host Driver version 450.102 and 452.77 for the Windows VDI image. These drivers come in the same download package from NVIDIA.

To install the GPU drivers, follow these steps:

1. Copy the Microsoft Windows drivers from the NVIDIA GRID vGPU driver pack downloaded earlier to the master virtual machine.

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 56. NVIDIA Driver Pack

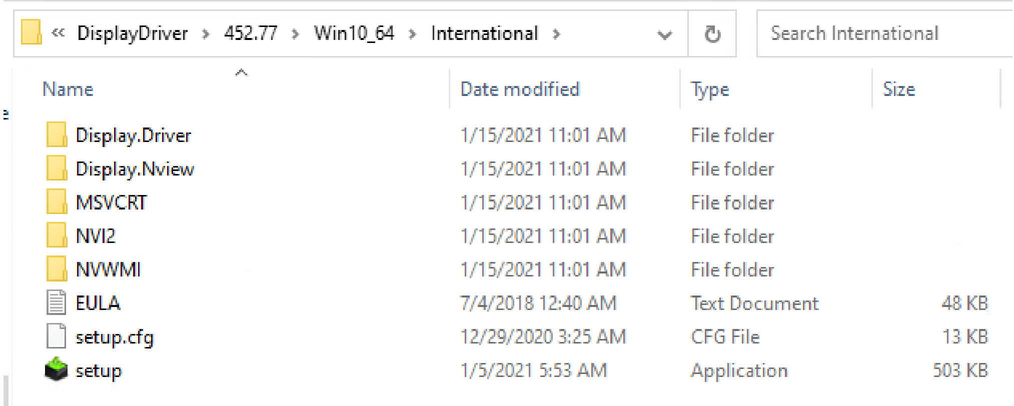

« DisplayDriver › 452.77 › Win10_64 › International ›

Name	Date modified	Type	Size
Display.Driver	1/15/2021 11:01 AM	File folder	
Display.Nview	1/15/2021 11:01 AM	File folder	
MSVCRT	1/15/2021 11:01 AM	File folder	
NVI2	1/15/2021 11:01 AM	File folder	
NVWMI	1/15/2021 11:01 AM	File folder	
EULA	7/4/2018 12:40 AM	Text Document	48 KB
setup.cfg	12/29/2020 3:25 AM	CFG File	13 KB
setup	1/5/2021 5:53 AM	Application	503 KB

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.

3. Agree to the NVIDIA software license.

Figure 57. Agreeing to the NVIDIA Software License

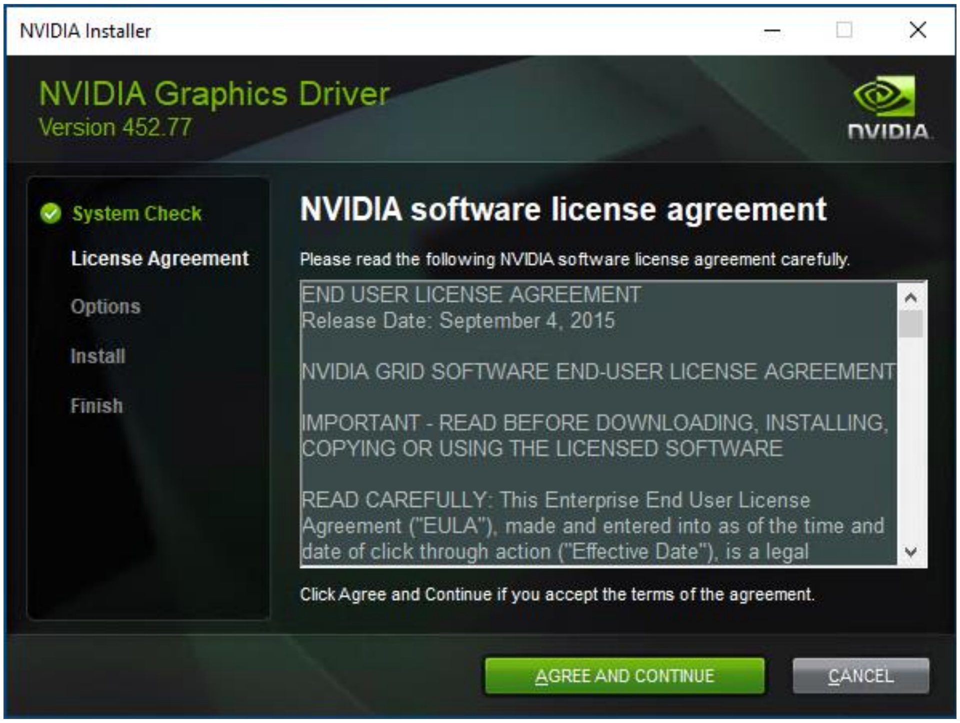

NVIDIA software license agreement

Please read the following NVIDIA software license agreement carefully.

END USER LICENSE AGREEMENT
Release Date: September 4, 2015

NVIDIA GRID SOFTWARE END-USER LICENSE AGREEMENT

IMPORTANT - READ BEFORE DOWNLOADING, INSTALLING, COPYING OR USING THE LICENSED SOFTWARE

READ CAREFULLY: This Enterprise End User License Agreement ("EULA"), made and entered into as of the time and date of click through action ("Effective Date"), is a legal

Click Agree and Continue if you accept the terms of the agreement.

4. Install the graphics drivers using the Express or Custom option. After the installation has completed successfully, restart the virtual machine.

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 58. Selecting the Express or Custom Installation Option

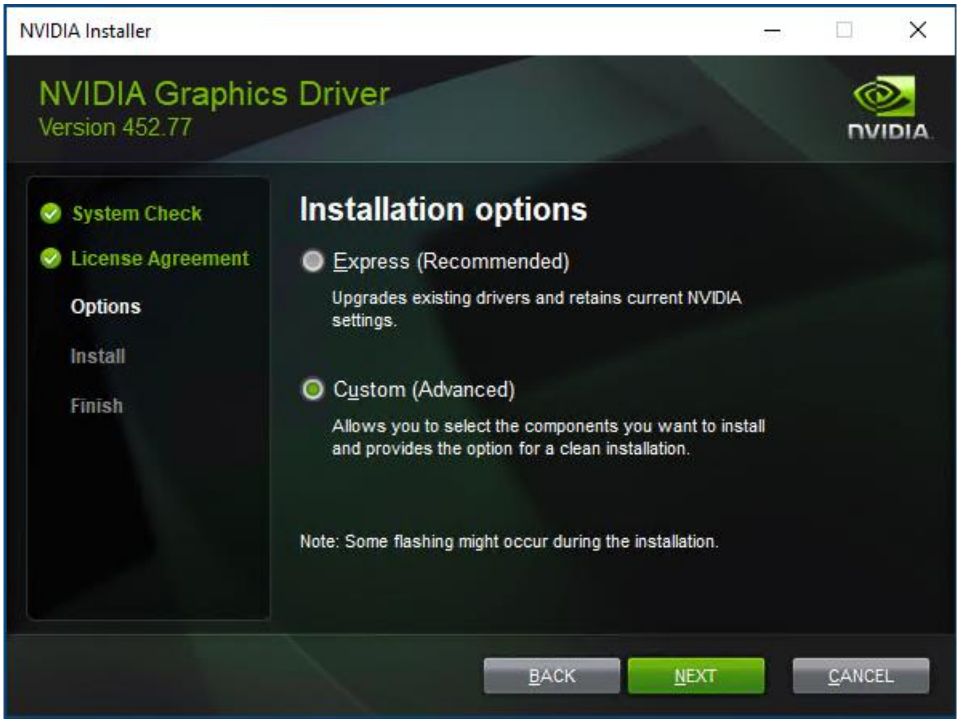

Installation options

○ Express (Recommended)

Upgrades existing drivers and retains current NVIDIA settings.

● Custom (Advanced)

Allows you to select the components you want to install and provides the option for a clean installation.

Note: Some flashing might occur during the installation.

Figure 59. Components Installed During NVIDIA Graphics Driver Custom Installation Process

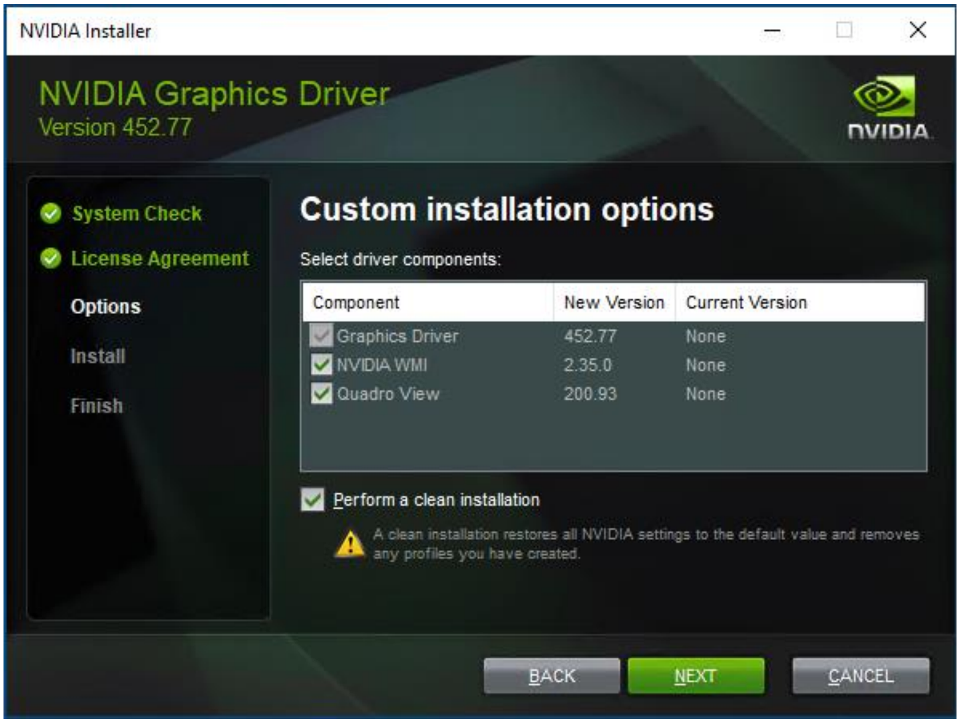

Custom installation options

Select driver components:

Component	New Version	Current Version
☑ Graphics Driver	452.77	None
☑ NVIDIA WMI	2.35.0	None
☑ Quadro View	200.93	None

☑ Perform a clean installation

A clean installation restores all NVIDIA settings to the default value and removes any profiles you have created.

Figure 60. Restarting the Virtual Machine

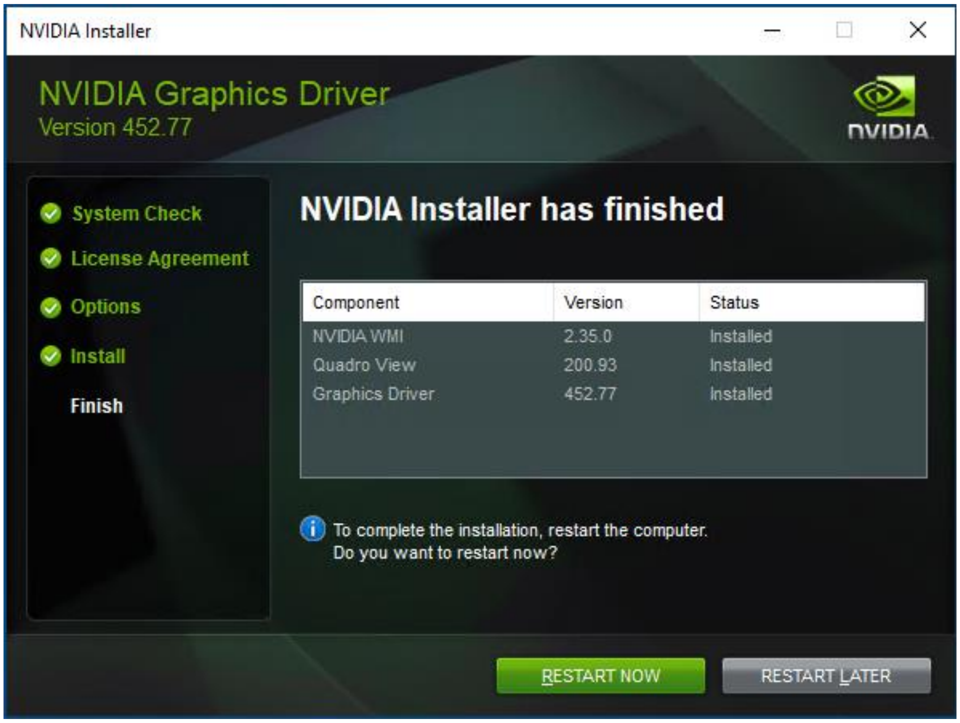

NVIDIA Installer has finished

Component	Version	Status
NVIDIA WMI	2.35.0	Installed
Quadro View	200.93	Installed
Graphics Driver	452.77	Installed

To complete the installation, restart the computer.
Do you want to restart now?

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. To do so, follow these steps:

1. In the Windows Control Panel, double-click the NVidia Control Panel.

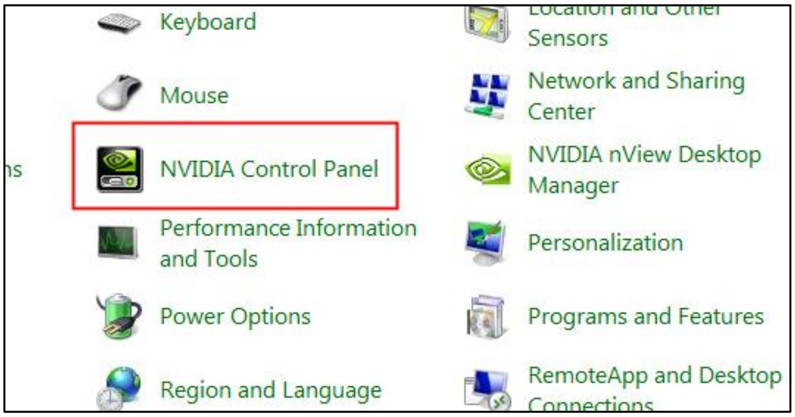
Keyboard
Location and Other Sensors
Mouse
Network and Sharing Center
ns
NVIDIA Control Panel
NVIDIA nView Desktop Manager
Performance Information and Tools
Personalization
Power Options
Programs and Features
Region and Language
RemoteApp and Desktop Connections

2. In the Control Panel, enter the IP or FQDN of the Grid License Server. You will receive a result similar to the one shown below.

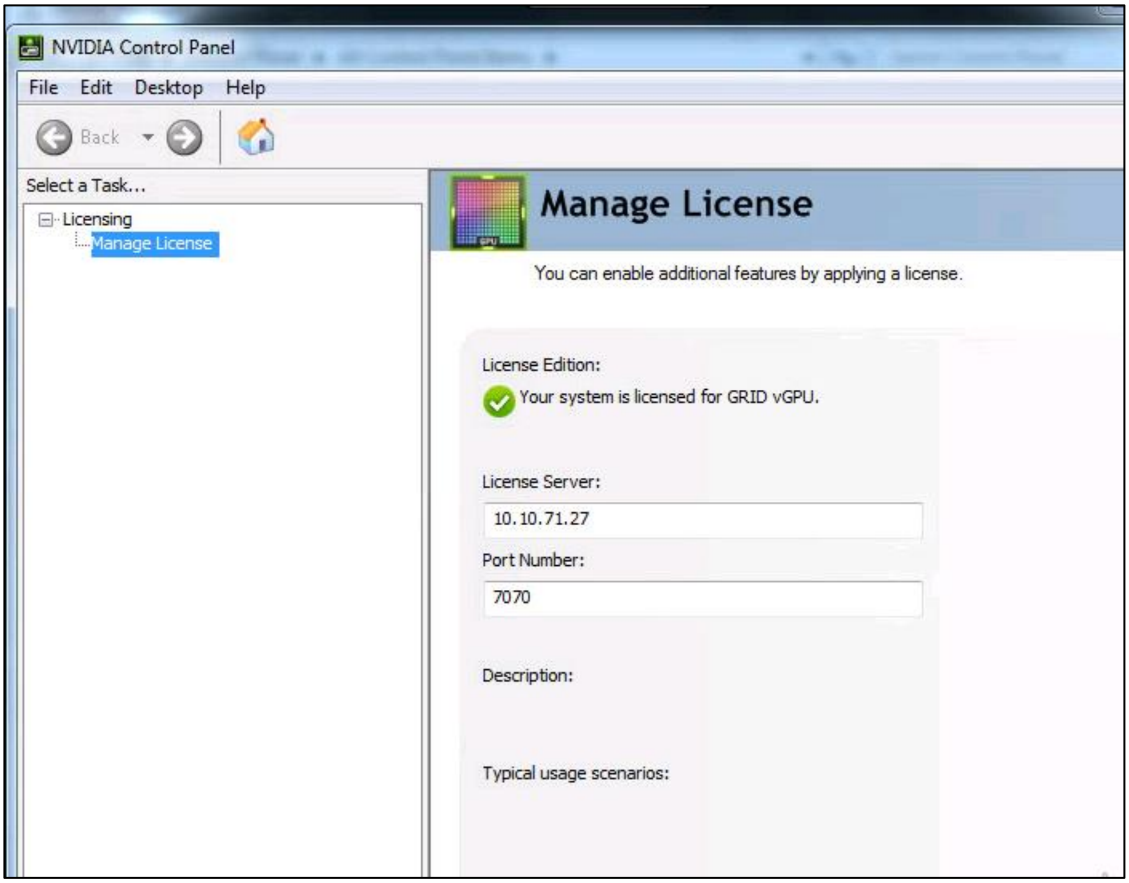

You can enable additional features by applying a license.

License Edition:

✓ Your system is licensed for GRID vGPU.

License Server:

10.10.71.27

Port Number:

7070

Description:

Typical usage scenarios:

Cisco Intersight Cloud Based Management

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

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

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

Cisco Intersight License

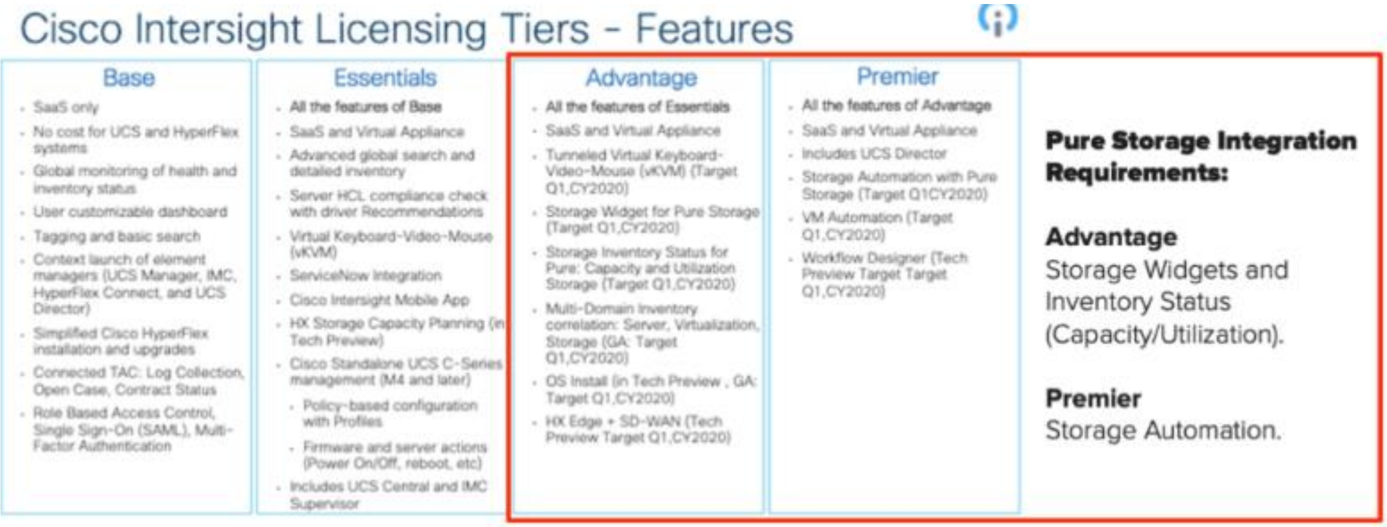

Cisco Intersight Licensing Tiers - Features

Base

- SaaS only
- No cost for UCS and HyperFlex systems
- Global monitoring of health and inventory status
- User customizable dashboard
- Tagging and basic search
- Context launch of element managers (UCS Manager, IMC, HyperFlex Connect, and UCS Director)
- Simplified Cisco HyperFlex installation and upgrades
- Connected TAC: Log Collection, Open Case, Contract Status
- Role Based Access Control, Single Sign-On (SAML), Multi-Factor Authentication

Essentials

- All the features of Base
- SaaS and Virtual Appliance
- Advanced global search and detailed inventory
- Server HCL compliance check with driver Recommendations
- Virtual Keyboard-Video-Mouse (vKVM)
- ServiceNow Integration
- Cisco Intersight Mobile App
- HX Storage Capacity Planning (in Tech Preview)
- Cisco Standalone UCS C-Series management (M4 and later)
  - Policy-based configuration with Profiles
  - Firmware and server actions (Power On/Off, reboot, etc)
- Includes UCS Central and IMC Supervisor

Advantage

- All the features of Essentials
- SaaS and Virtual Appliance
- Tunneled Virtual Keyboard-Video-Mouse (vKVM) (Target Q1,CY2020)
- Storage Widget for Pure Storage (Target Q1,CY2020)
- Storage Inventory Status for Pure: Capacity and Utilization Storage (Target Q1,CY2020)
- Multi-Domain Inventory correlation: Server, Virtualization, Storage (GA: Target Q1,CY2020)
- OS Install (in Tech Preview , GA: Target Q1,CY2020)
- HX Edge + SD-WAN (Tech Preview Target Q1,CY2020)

Premier

- All the features of Advantage
- SaaS and Virtual Appliance
- Includes UCS Director
- Storage Automation with Pure Storage (Target Q1CY2020)
- VM Automation (Target Q1,CY2020)
- Workflow Designer (Tech Preview Target Target Q1,CY2020)

Pure Storage Integration Requirements:

Advantage
Storage Widgets and Inventory Status (Capacity/Utilization).

Premier
Storage Automation.

Getting Started with the Cisco Intersight Platform

The Cisco Intersight platform provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with the Cisco Intersight platform is quick and easy.

To configure the Cisco Intersight platform, follow these steps:

1. If you do not already have a Cisco Intersight account, to claim your Cisco UCS system in a new account on Cisco Intersight, connect to https://intersight.com. If you have an existing Cisco Intersight account, connect to https://intersight.com and sign in with your Cisco ID, select the appropriate account, and skip to step 6.

2. Click Create an account.

3. Sign in with your Cisco ID.

4. Read, scroll through, and accept the End User License Agreement and click Next.

5. Enter an account name and click Create.

6. Choose ADMIN > Targets. Click Claim a New Target.

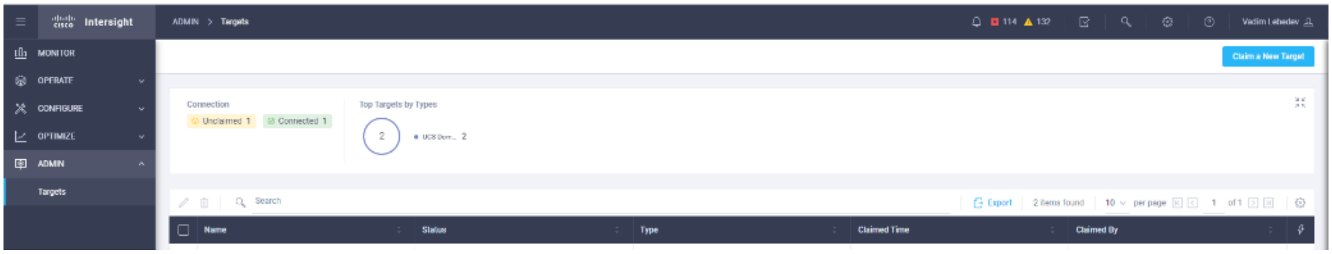

7. Select Cisco UCS Domain (UCSM Managed) and click Start.

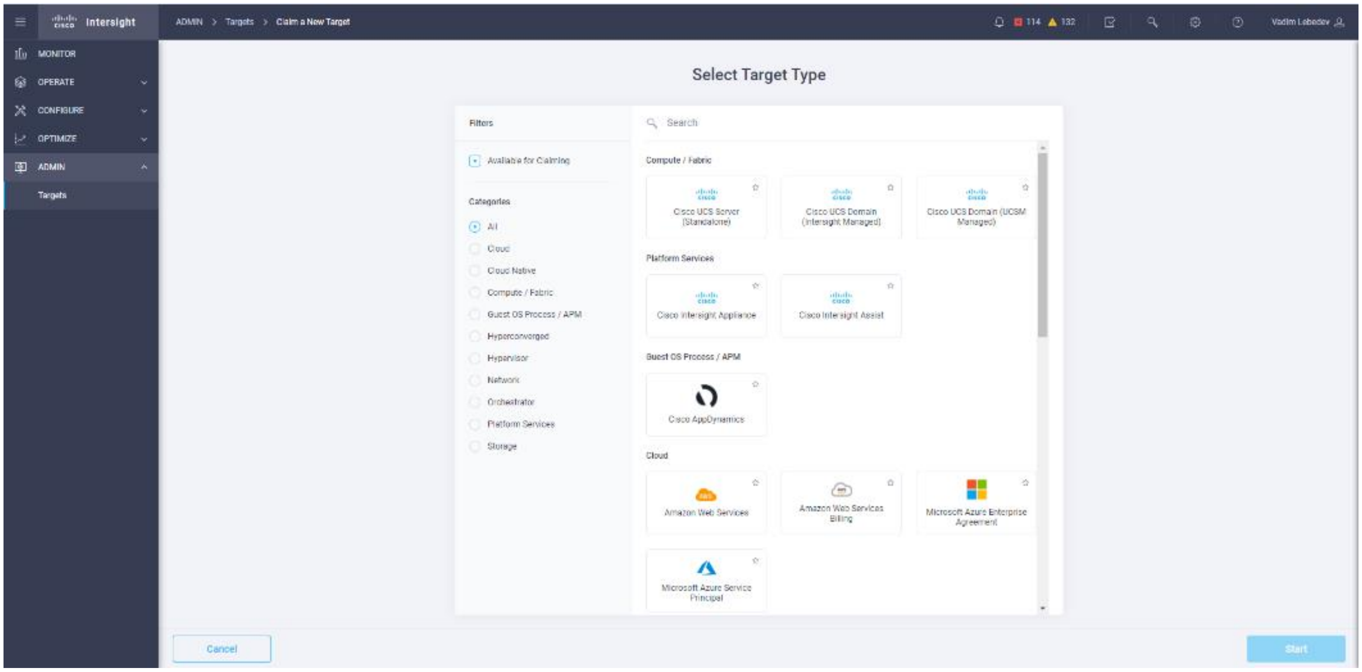

8. Fill in the device ID and claim code and click Claim.

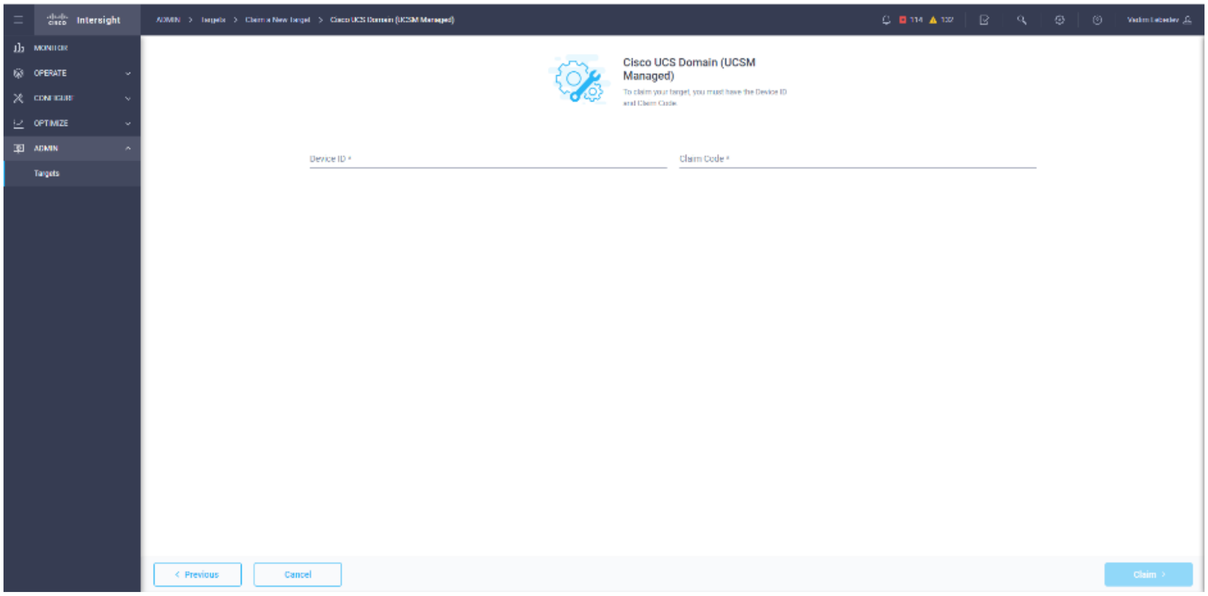
Intersight
ADMIN > Targets > Claim a New Target > Cisco UCS Domain (UCSM Managed)
114 132
Vadim Lebedev
MONITOR
OPERATE
CONFIGURE
OPTIMIZE
ADMIN
Targets
Cisco UCS Domain (UCSM Managed)
To claim your target, you must have the Device ID and Claim Code
Device ID *
Claim Code *
< Previous
Cancel
Claim >

The device ID and claim code can be obtained by connecting to Cisco UCS Manager and choosing Admin > All > Device Connector. The device ID and claim code are on the right.

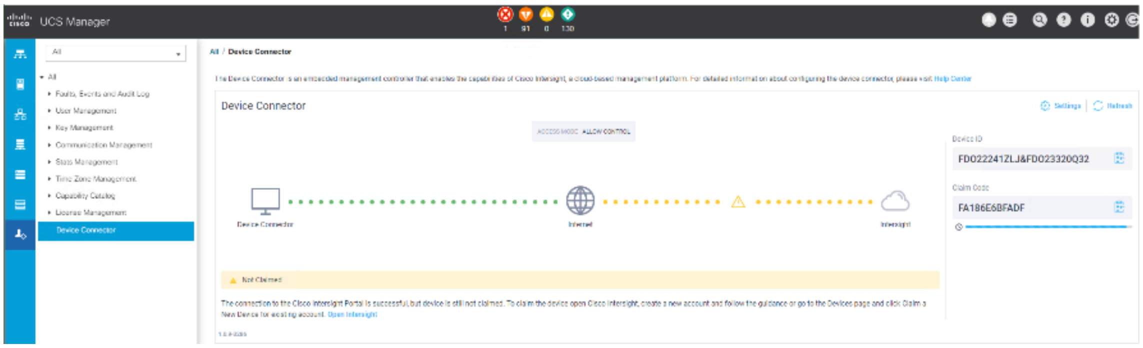

UCS Manager

All / Device Connector

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device Connector

ACCESS MODE ALLOW CONTROL

Device ID: FDO22241ZLJ&FDO23320Q32

Claim Code: FA186E6BFADF

Not Claimed

The connection to the Cisco Intersight Portal is successful, but device is still not claimed. To claim the device open Cisco Intersight, create a new account and follow the guidance or go to the Devices page and click Claim a New Device for existing account. Open Intersight

9. The target will be visible in the list of your available targets.

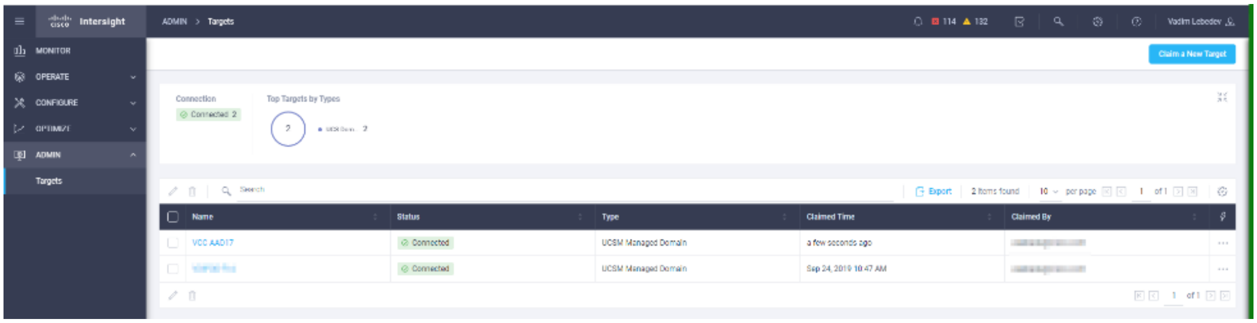

Intersight — ADMIN > Targets

Connection: Connected 2

Top Targets by Types: 2

2 items found

Name	Status	Type	Claimed Time	Claimed By
VCC AAD17	Connected	UCSM Managed Domain	a few seconds ago	[redacted]
[redacted]	Connected	UCSM Managed Domain	Sep 24, 2019 10 47 AM	[redacted]

10. From the Cisco Intersight window, click the gear icon (⚙)and then click Licensing. If this is a new account, all servers connected to the Cisco UCS domain will appear under the Base license tier. If you have purchased Cisco Intersight licenses and have them in your Cisco Smart Account, click Register and follow the prompts to register this Cisco Intersight account in your Cisco Smart Account. Cisco Intersight also offers a one-time 90-day trial of Premier licensing for new accounts. Click Start Trial and then Start to begin this evaluation. The remainder of this section assumes that you are using Premier licensing.

Test Setup, Configuration, and Load Recommendation

In this solution, we tested a single Cisco UCS B200 M5 blade to validate against the performance of one blade and thirty B200 M5 blades 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 Recommended Maximum Workload per host server using VMware Horizon 8 with 224 RDS sessions, 186 VDI Non-Persistent sessions, and 186 VDI Persistent sessions.

Figure 61. Test Configuration for Single Server Scalability VMware Horizon 8 VDI (Persistent) Full Clones

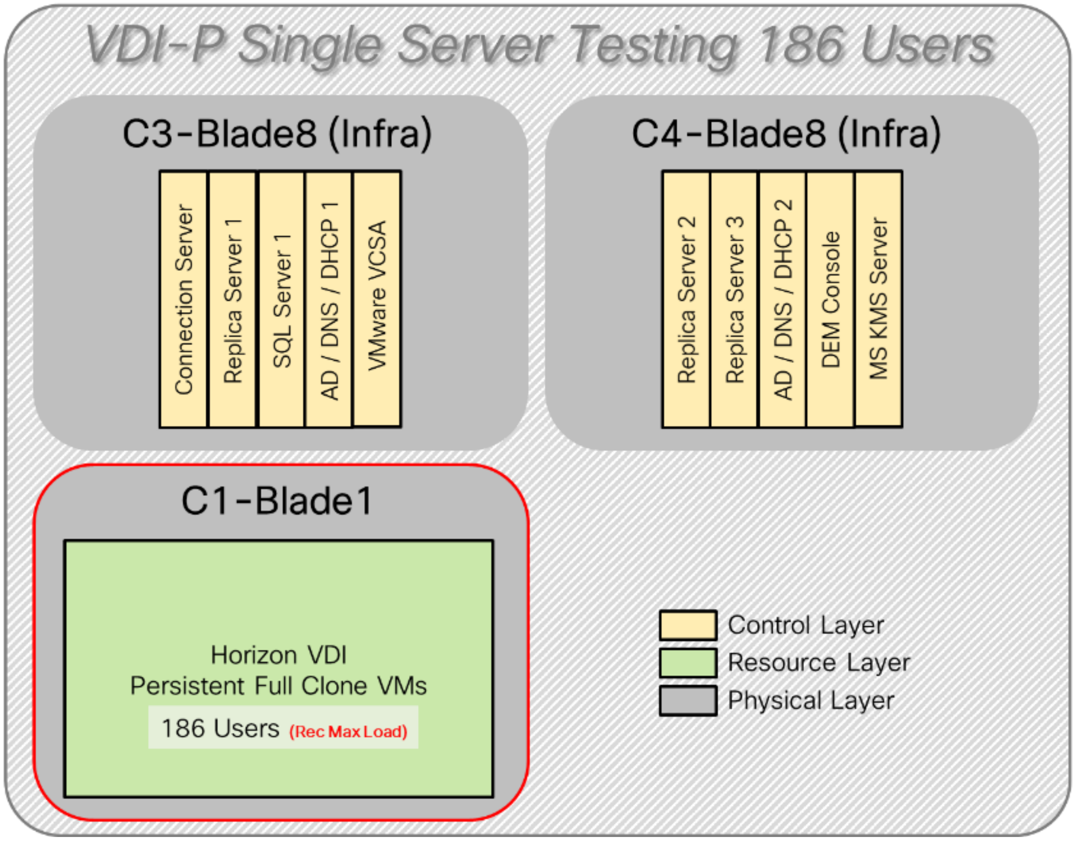

Figure 62. Test Configuration for Single Server Scalability VMware Horizon 8 VDI (Non-Persistent) Instant Clones

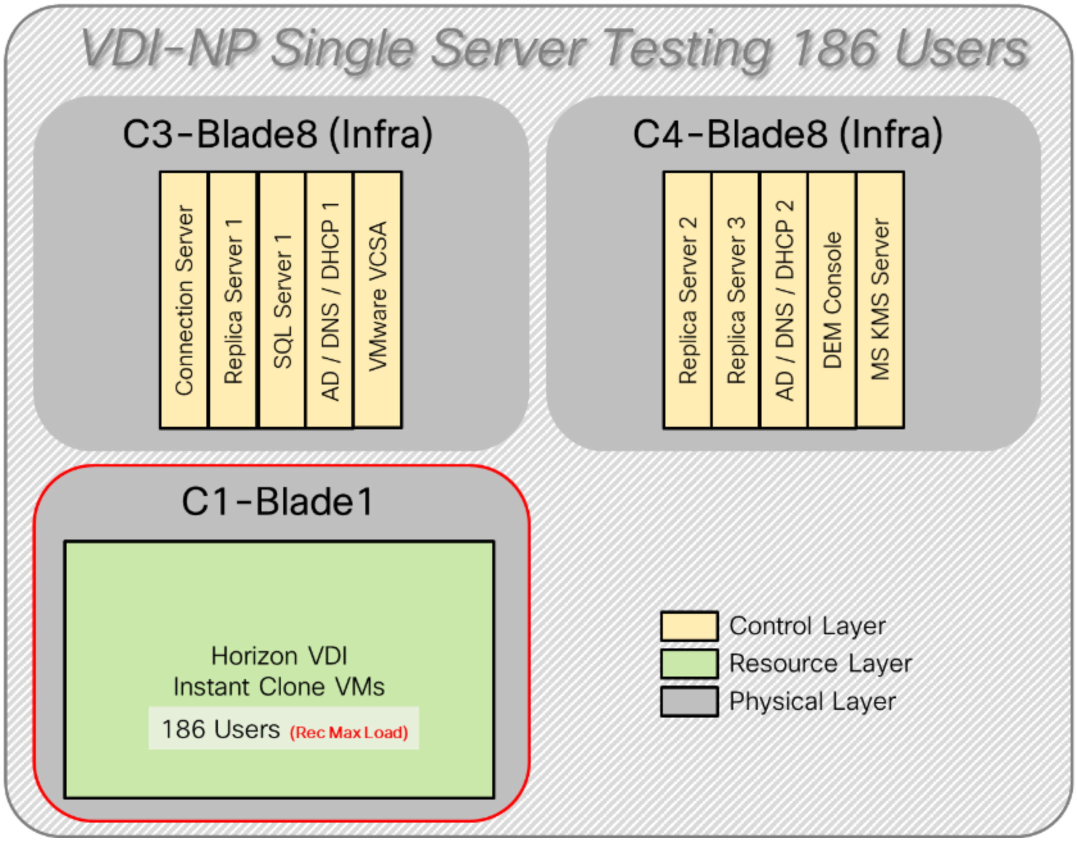

Figure 63. Test configuration for Single Server Scalability VMware Horizon 8 RDS Full Clones

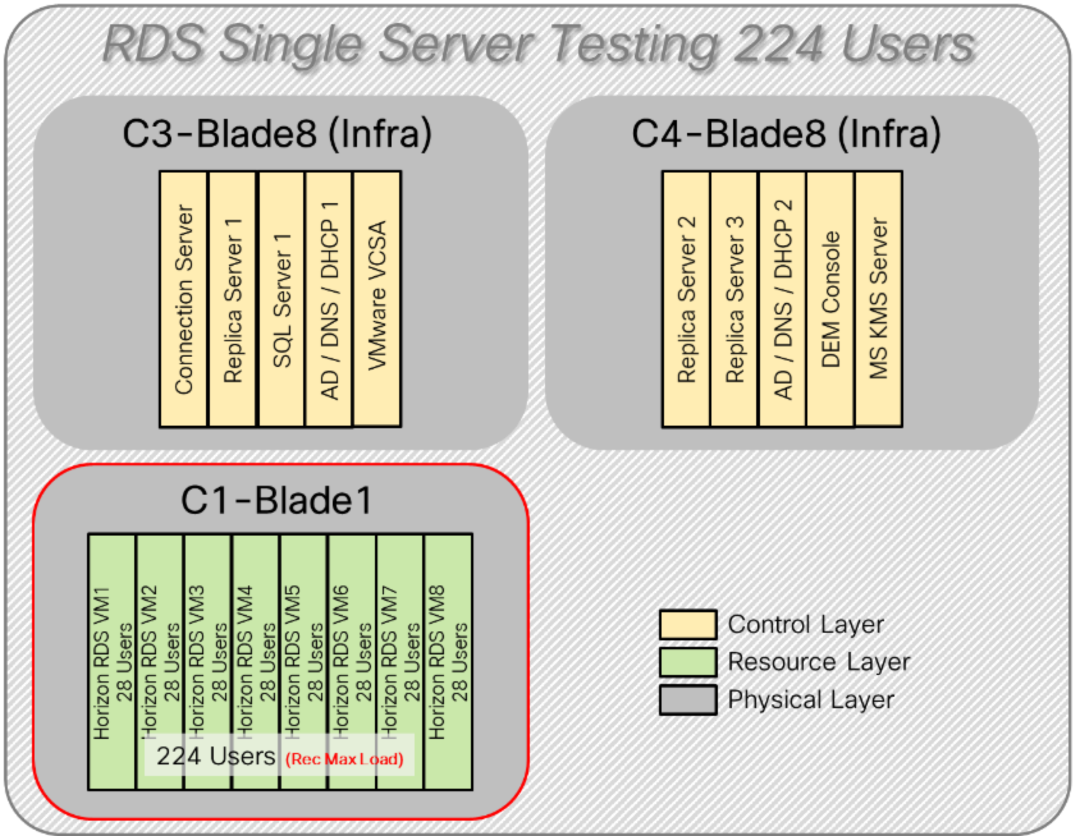

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M5 Blade servers with Intel Xeon Silver 4210 2.20-GHz 10-core processors, 384GB 2933MHz RAM for all host blades
- 1 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.1-GHz 20-core processors, 768GB 2933MHz RAM for all host blades
- Cisco VIC 1440 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switches
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives

Software components:

- Cisco UCS firmware 4.1(2a)
- Pure Storage Purity//FA 6.0.3

- VMware ESXi 7.0 GA for host blades
- VMware Horizon 8 VDI Desktops and RDSH Desktops
- VMware DEM Enterprise 10
- Microsoft SQL Server 2019
- Microsoft Windows 10 64 bit (1909), 2vCPU, 3 GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 GB vDisk (master)
- Microsoft Office 2019
- Login VSI 4.1.39 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Test Configuration for Cluster Scale Testing

This test case validates three workload clusters of ten blades using VMware Horizon 8 with 2016 RDS sessions, 1674 VDI Non-Persistent sessions, and 1674 VDI Persistent sessions. Server N+1 fault tolerance is factored into this test scenario for each workload and infrastructure cluster.

Figure 64. Test Configuration for Cluster Scalability VMware Horizon 8 VDI (Persistent) Full Clones

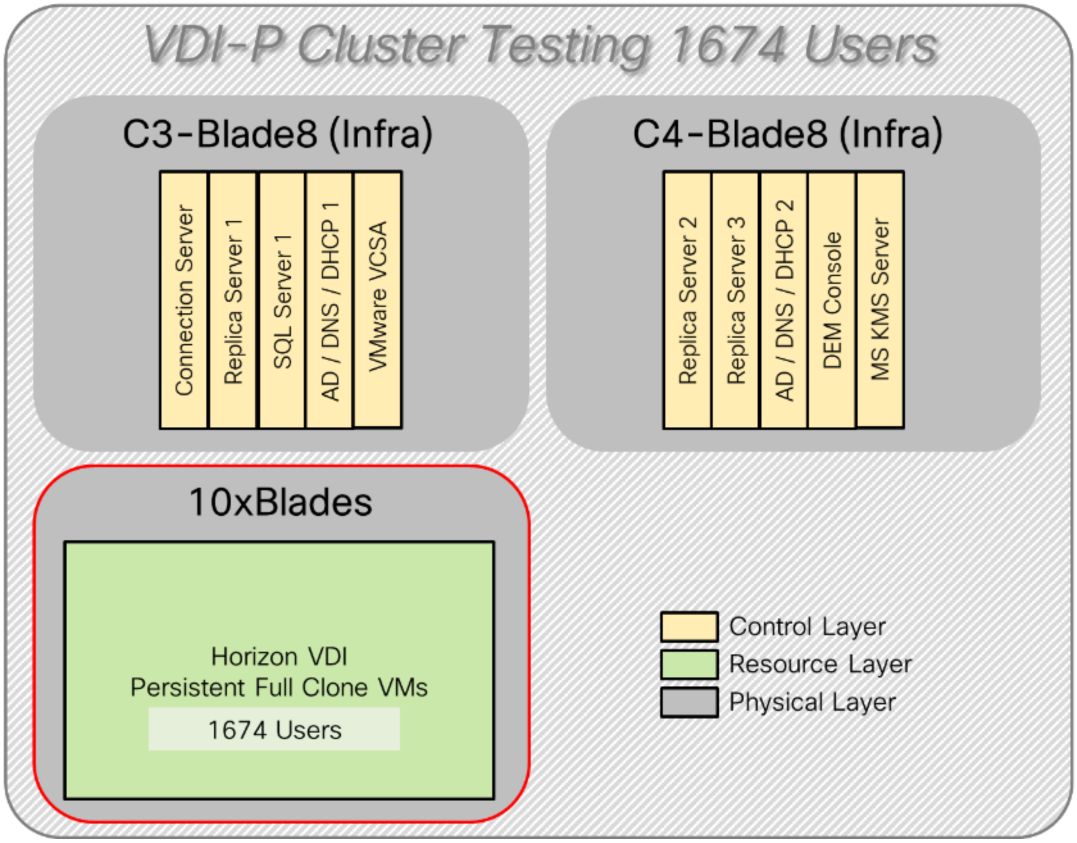

Figure 65. Test Configuration for Cluster Scalability VMware Horizon 8 VDI (Non-Persistent) Instant Clones

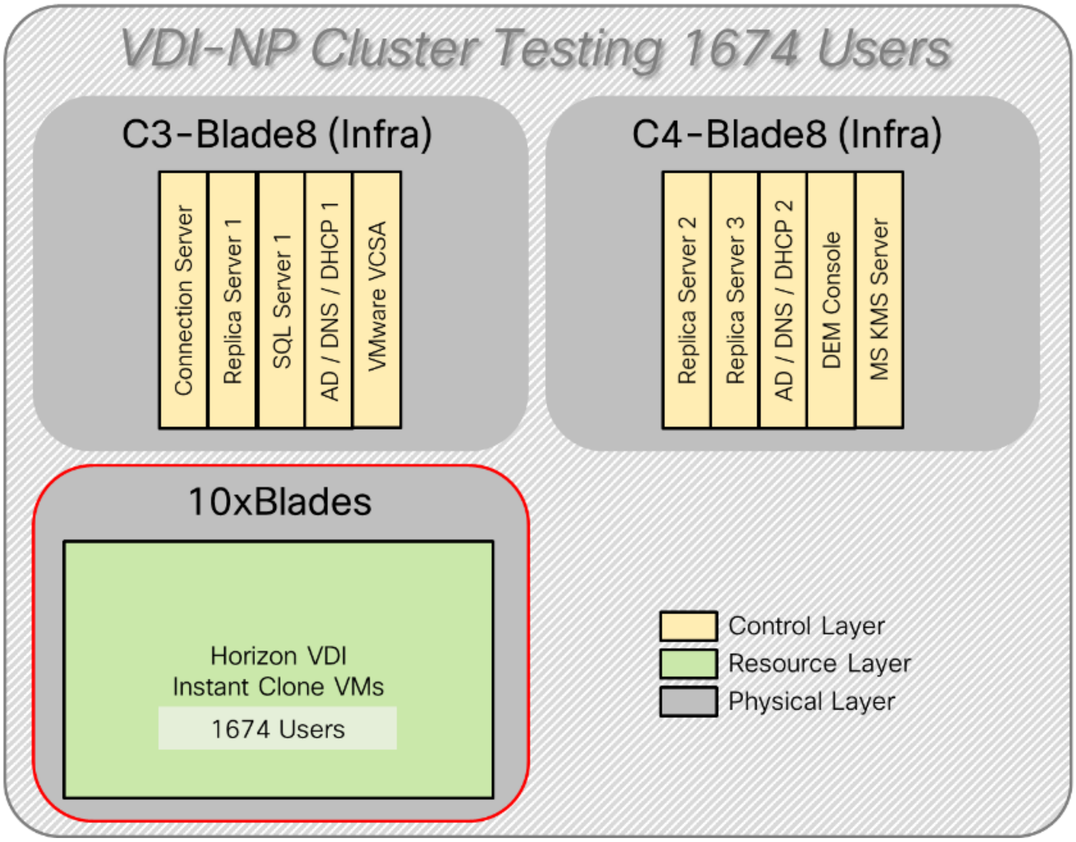

Figure 66. Test Configuration for Single Server Scalability VMware Horizon 8 RDS Full Clones

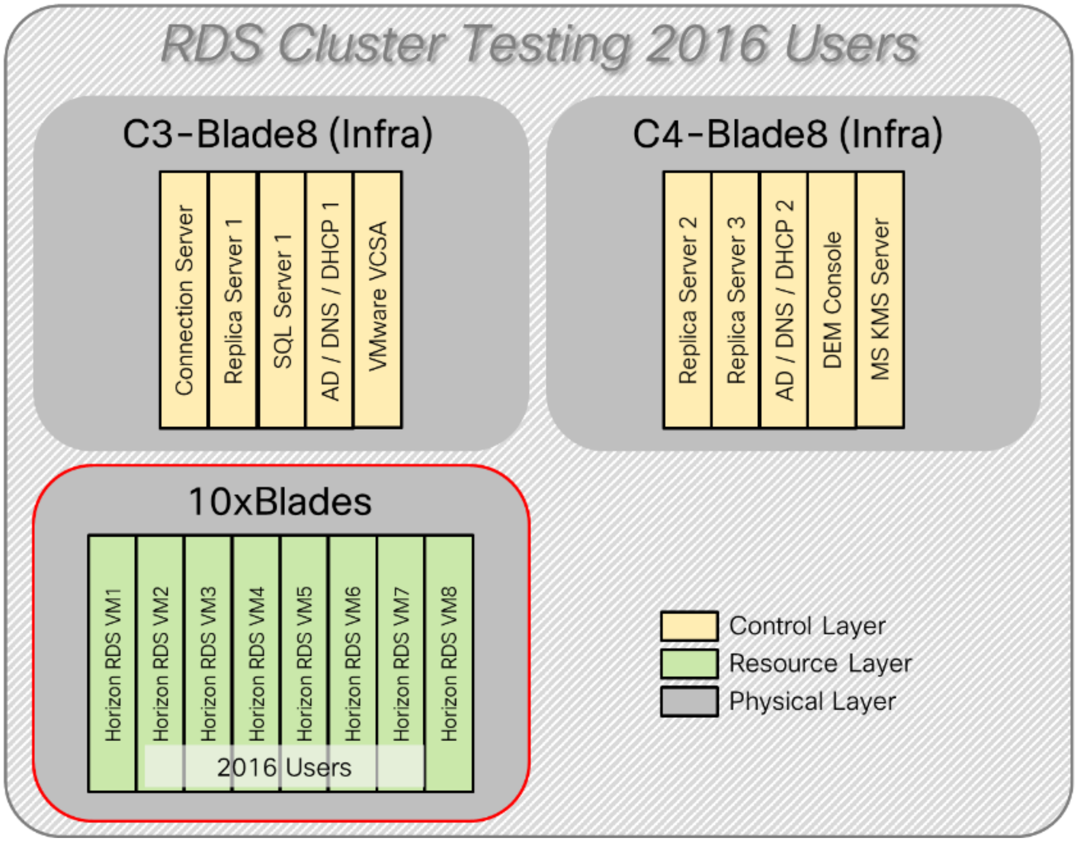

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M5 Blade servers with Intel Xeon Silver 4210 2.20-GHz 10-core processors, 384GB 2933MHz RAM for all host blades
- 10 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6230 2.1-GHz 20-core processors, 768GB 2933MHz RAM for all host blades
- Cisco VIC 1440 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switches
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives

Software components:

- Cisco UCS firmware 4.1(2a)
- Pure Storage Purity//FA 6.0.3

- VMware ESXi 7.0 GA for host blades
- VMware Horizon 8 VDI Desktops and RDSH Desktops
- VMware DEM Enterprise 10
- Microsoft SQL Server 2019
- Microsoft Windows 10 64 bit (1909), 2vCPU, 3 GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 GB vDisk (master)
- Microsoft Office 2019
- Login VSI 4.1.39 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Test Configuration for Full Scale Testing

These test cases validate thirty blades in three clusters of three distinct workloads using VMware Horizon 8 with:

- 6048 Non-Persistent RDS sessions (Full clones).
- 5022 Persistent VDI sessions (Full clones).
- 5022 Non-Persistent VDI sessions (Instant clones).

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

Figure 67. Test Configuration for Full Scale VMware Horizon 8 VDI (Persistent) Full Clones

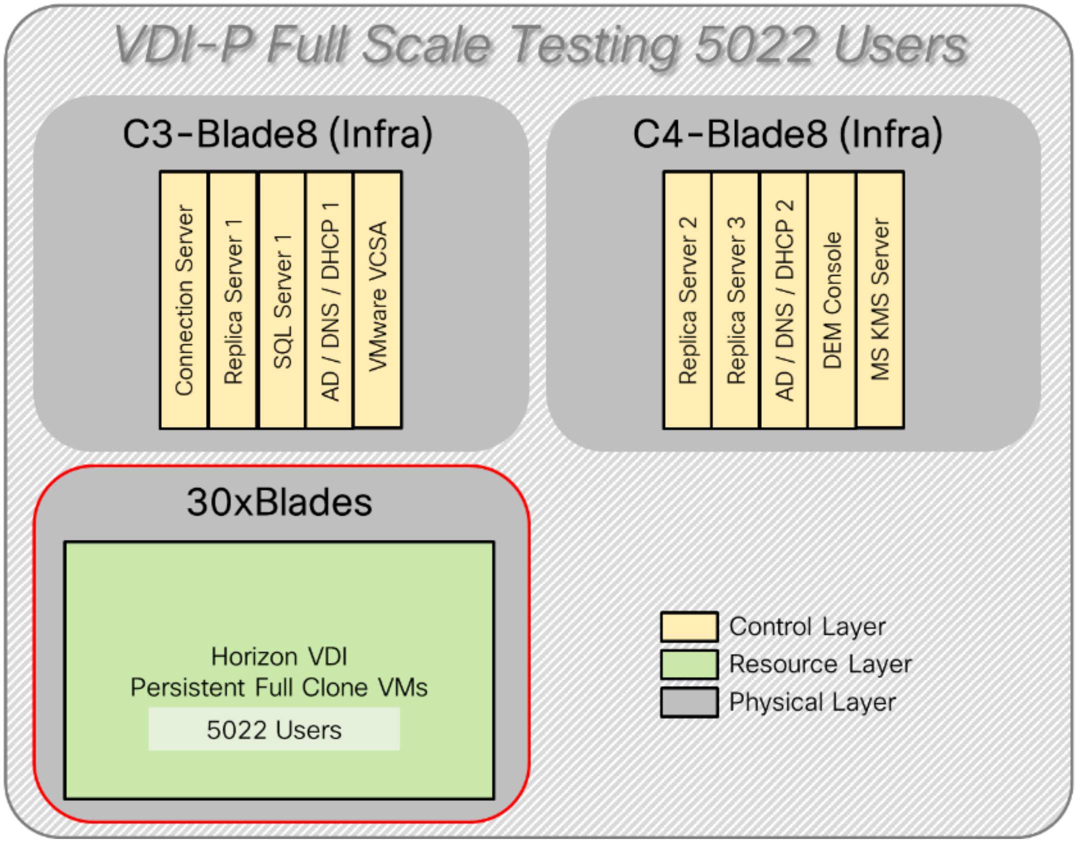

Figure 68. Test Configuration for Full Scale VMware Horizon 8 VDI (Non-Persistent) Instant Clones

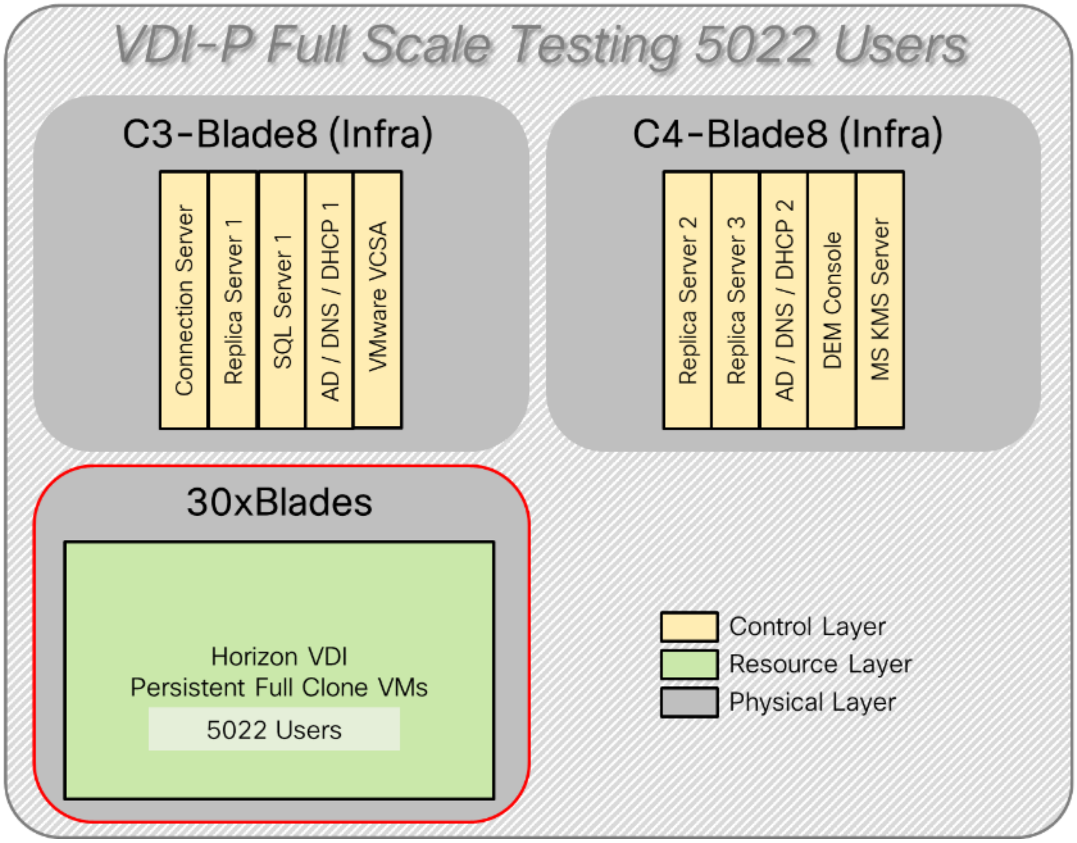

Figure 69. Test Configuration for Full Scale VMware Horizon 8 RDS Full Clones

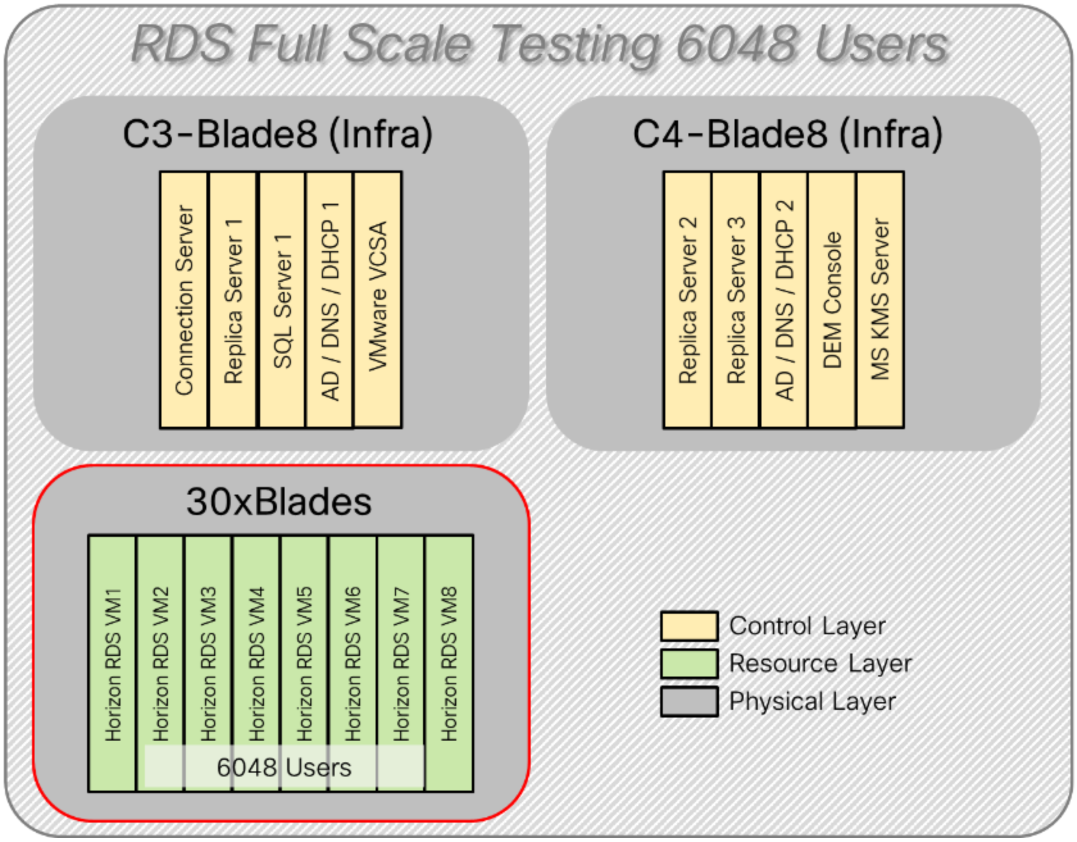

Test Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the RDSH/VDI Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from http://www.loginvsi.com

Test Procedure

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 Console and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

All VMware ESXi VDI host blades to be tested were restarted prior to each test cycle.

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. To do so, follow these steps:

1. Time 0:00:00 Start PerfMon/Esxtop Logging on the following systems:
    a. Infrastructure and VDI Host Blades used in the test run
    b. vCenter used in the test run
    c. All Infrastructure virtual machines used in test run (AD, SQL, brokers, image mgmt., etc.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using View Connection server.

The boot rate should be around 10-12 virtual machines per minute per server.

4. Time 0:06 First machines boot.

5. Time 0:30 Single Server or Scale target number of desktop virtual machines booted on 1 or more blades.

No more than 30 minutes for boot up of all virtual desktops is allowed.

6. Time 0:35 Single Server or Scale target number of desktop virtual machines desktops available on View Connection Server.

7. Virtual machine settling time.

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 30-45-minute rest period is sufficient.

8. Time 1:35 Start Login VSI 4.1.x Office Worker Benchmark Mode Test, setting auto-logoff time at 15 minutes, with Single Server or Scale target number of desktop virtual machines utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).

9. Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48 minute benchmark launch rate).

10. Time 2:25 All launched sessions must become active. id test run within this window.

11. Time 2:40 Login VSI Test Ends (based on Auto Logoff 15 minutes period designated above).

    a. Time 2:55 All active sessions logged off.
    b. Time 2:57 All logging terminated; Test complete.
    c. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines.
    d. Time 3:30 Reboot all hypervisor hosts.
    e. Time 3:45 Ready for the new test sequence.

Success Criteria

Our pass criteria for this testing is as follows:

Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1.x Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The VMware Horizon Console 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 above and will note that we did not reach a VSImax dynamic in our testing. FlashStack Data Center with Cisco UCS and VMware Horizon 8 on VMware ESXi 7.0 GA Test Results.

The purpose of this testing is to provide the data needed to validate VMware Horizon Remote Desktop Sessions (RDS) and VMware Horizon Virtual Desktop (VDI) instant-clones and VMware Horizon Virtual Desktop (VDI) full-clones models using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2019 sessions on Cisco UCS B200 M5 Blade Servers using the Pure Storage FlashArray//X70 R3 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, linear scalability.

VSImax 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 RDSH or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer, and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

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 coming closer to its saturation point, response times will rise. When reviewing the average response time, it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system and are initiated at logon within the simulated user's desktop session context.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI, the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solution, for a benchmark like Login VSI.

Calculating VSImax v4.1.x

The simulated desktop workload is scripted in a 48 minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop, the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The 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 do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, and so on. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

Figure 70. Sample of a VSI Max Response Time Graph, Representing a Normal Test

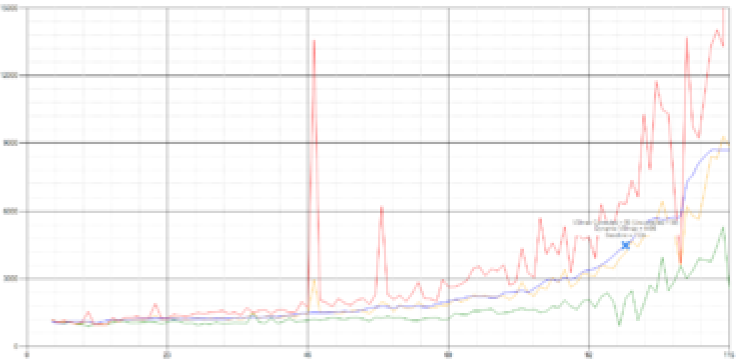

Figure 71. Sample of a VSI Test Response Time Graph with a Performance Issue

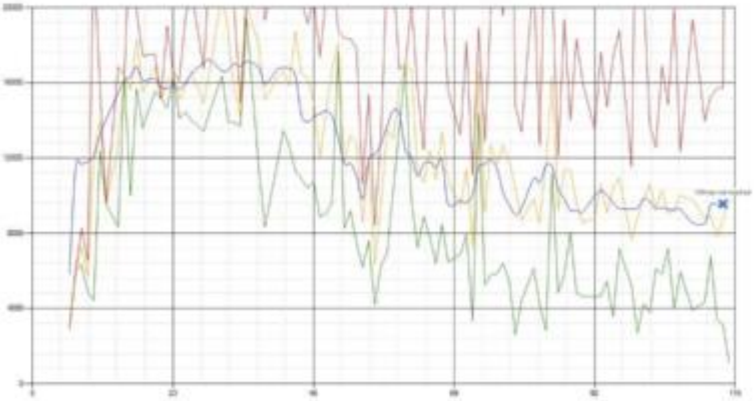

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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 times of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSImax models, this weighting much better represents system performance. All actions have 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, and so on) enabling 'basephase' for a more reliable baseline has become obsolete. The calculation is explained below. In total the 15 lowest VSI response time samples are taken from the entire test; the lowest 2 samples are removed. and the 13 remaining samples are averaged. The result is the Baseline. To summarize:

- Take the lowest 15 samples of the complete test
- From those 15 samples remove the lowest 2
- Average the 13 results that are left is the baseline

The VSImax average response time in Login VSI 4.1.x is calculated on the number of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40 percent of the number of “active” sessions. For example, if the active sessions are 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement is used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5 percent, and bottom 5 percent of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples, the top 2 samples are removed, and the lowest 2 results are removed (5 percent of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

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.

Lastly, 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 and 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 insight into system performance and scales to extremely large systems.

Single-Server Recommended Maximum Workload

For both the VMware Horizon 8 Virtual Desktop and VMware Horizon 8 Remote Desktop Service Hosts (RDSH) use cases, a recommended maximum workload was determined by the Login VSI Knowledge Worker Workload in VSI Benchmark Mode end user experience measurements and blade server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to ensure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.

Memory should never be oversubscribed for Desktop Virtualization workloads.

Table 16. Phases of Test Runs

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time
Idle	The rest time after the last desktop is registered on the XD Studio. (typically, a 30-45 minute, <60 min)
Logon	The Login VSI phase of the test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files (typically for the 15-minute duration)
Logoff	Sessions finish executing the Login VSI workload and logoff

Test Results

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 following three tests:

- 186 VDI Persistent sessions.
- 186 VDI Non-Persistent sessions
- 224 RDSH 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.

LoginVSI data is as follows:

Figure 72. Single Server Recommended Maximum Workload | Horizon 8 RDS | VSI Score

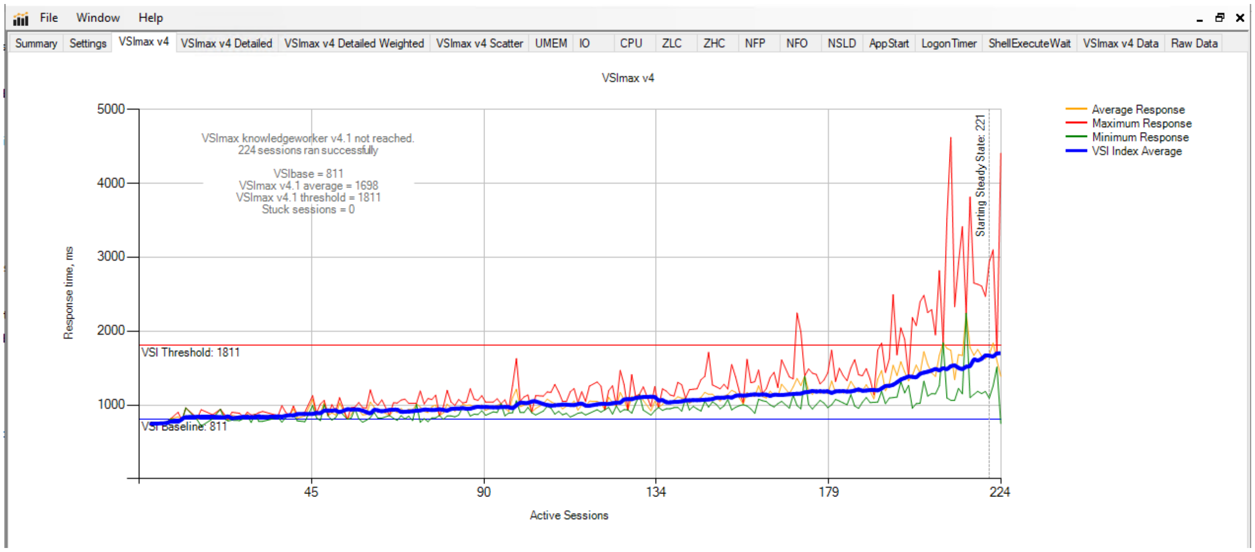

VSImax knowledgeworker v4.1 not reached.
224 sessions ran successfully

VSIbase = 811
VSImax v4.1 average = 1698
VSImax v4.1 threshold = 1811
Stuck sessions = 0

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

Figure 73. Single Server Recommended Maximum Workload | Horizon 8 RDS | Host CPU Utilization

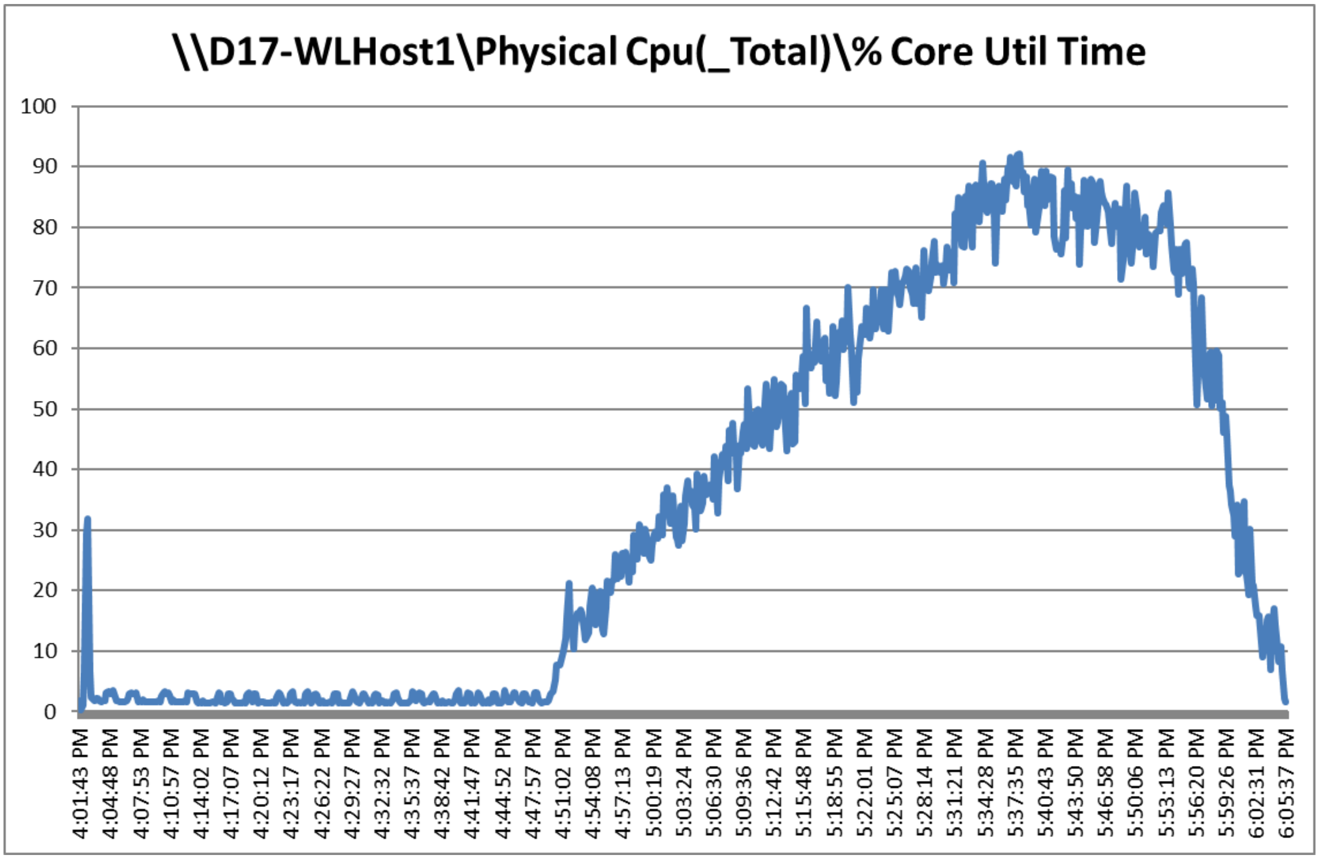

Figure 74. Single Server Recommended Maximum Workload | Horizon 8 RDS | Host Memory Utilization

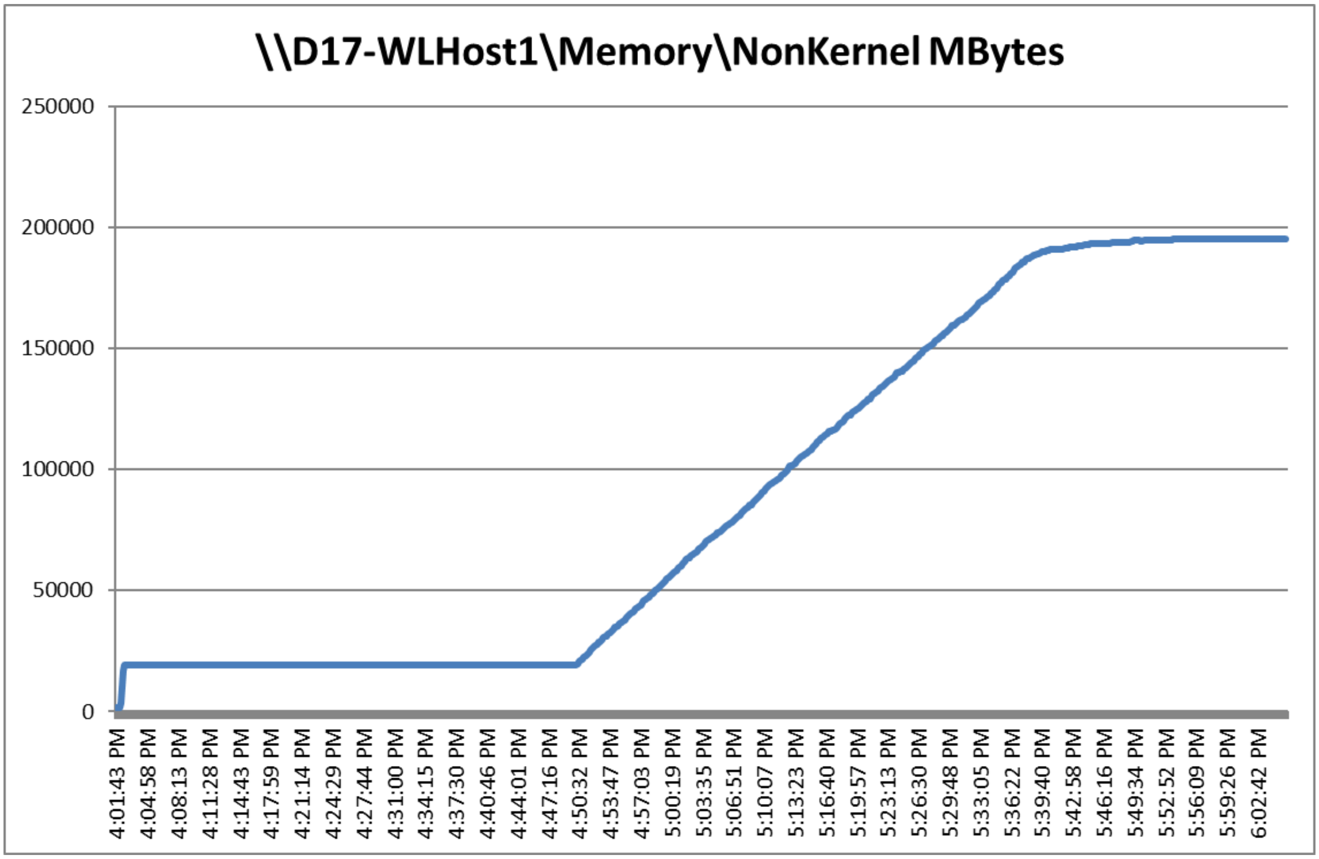

Figure 75. Single Server | Horizon 8 RDS | Host Network Utilization

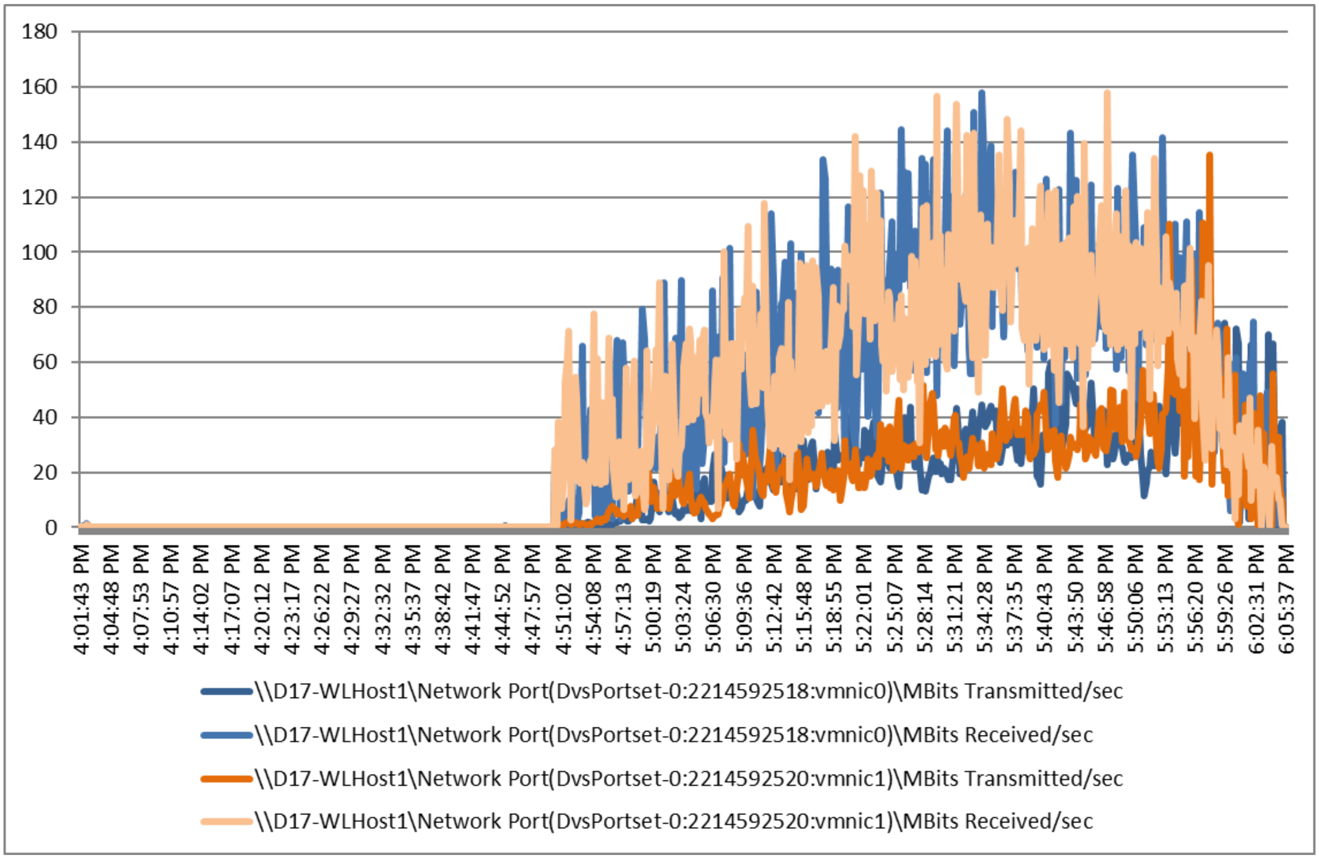

Performance data for the RDS Virtual Machine running the workload is as follows:

Figure 76. Single Server Recommended Maximum Workload | Horizon 8 RDS | Virtual Machine CPU Utilization

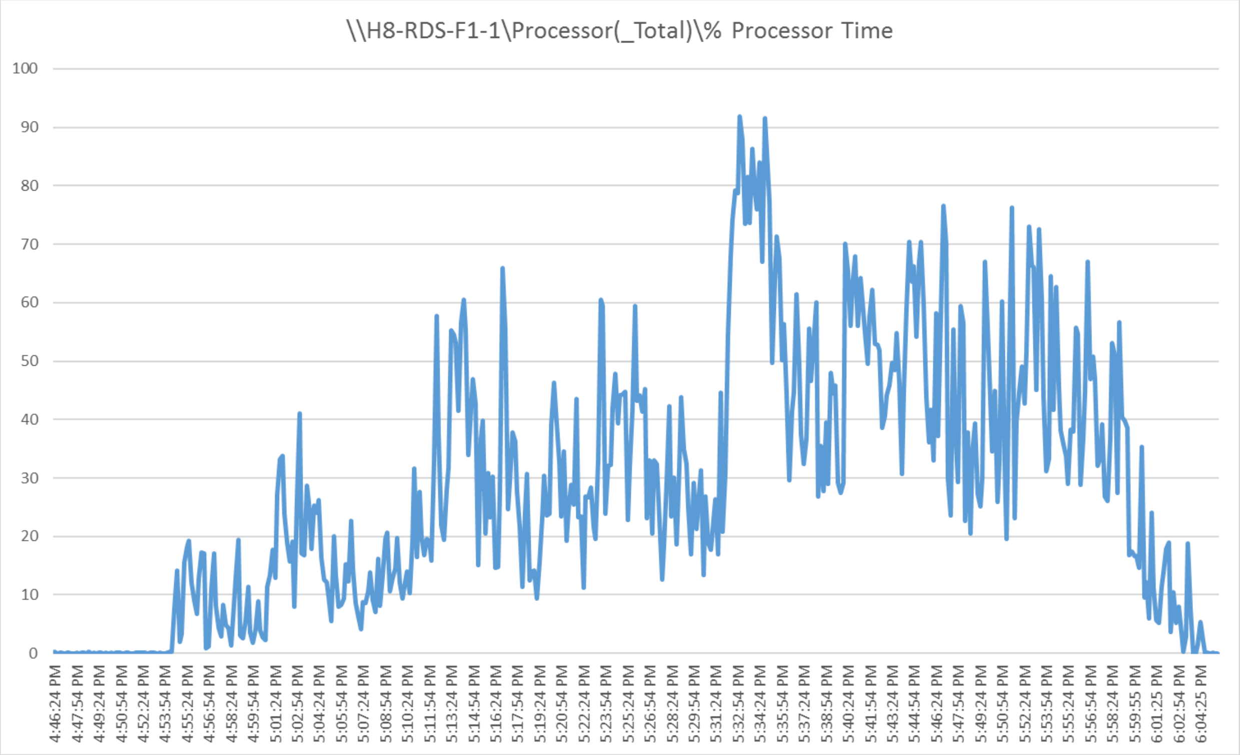

Figure 77. Single Server Recommended Maximum Workload | Horizon 8 RDS | Virtual Machine Memory Utilization

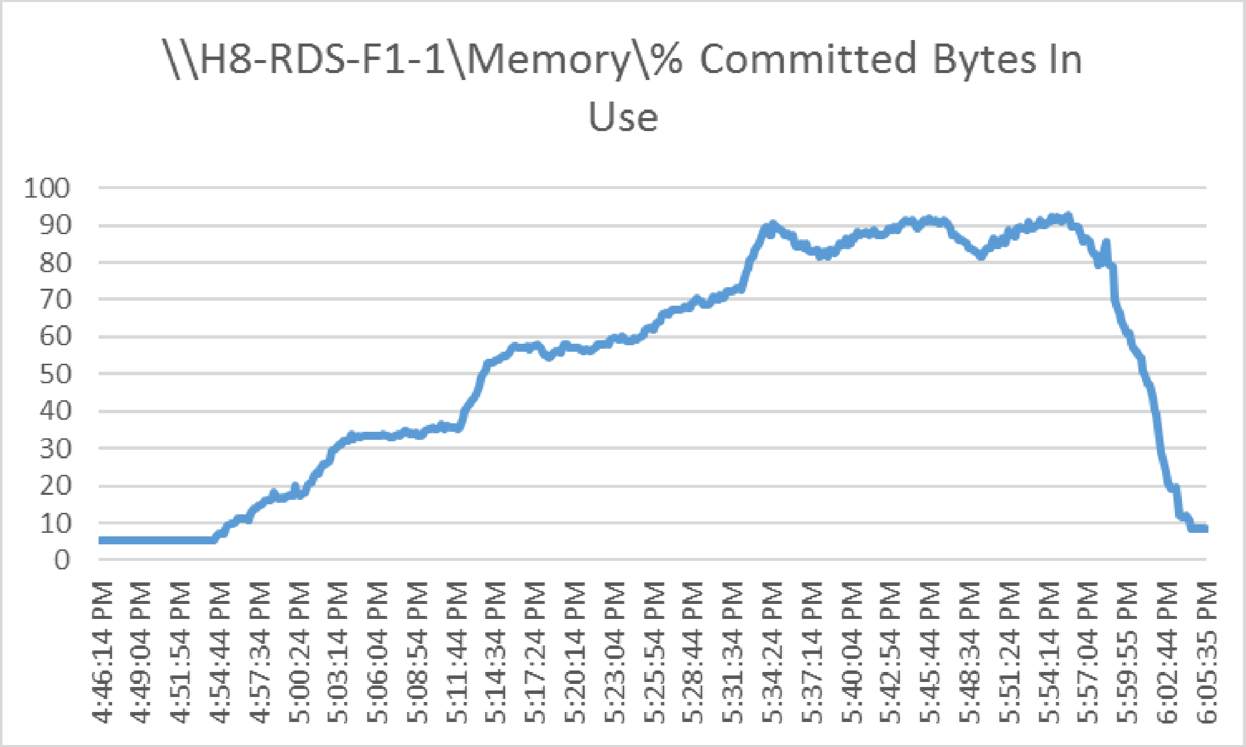

Single-Server Recommended Maximum Workload for VDI Non-Persistent with 186 Users

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 186 Windows 10 64-bit VDI non-persistent instant clone virtual machines with 2 vCPU and 3.5 GB RAM.

Login VSI performance data is as follows:

Figure 78. Single Server | Horizon 8 VDI-NP | VSI Score

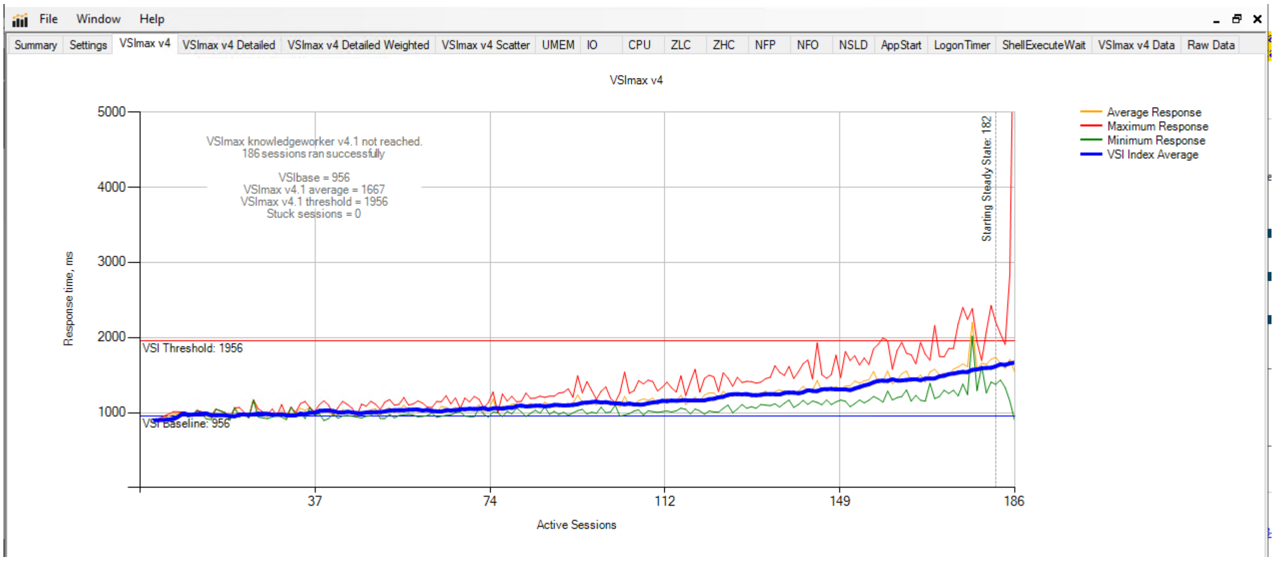

VSImax knowledgeworker v4.1 not reached.
186 sessions ran successfully

VSIbase = 956
VSImax v4.1 average = 1667
VSImax v4.1 threshold = 1956
Stuck sessions = 0

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

Figure 79. Single Server | Horizon 8 VDI-NP | Host CPU Utilization

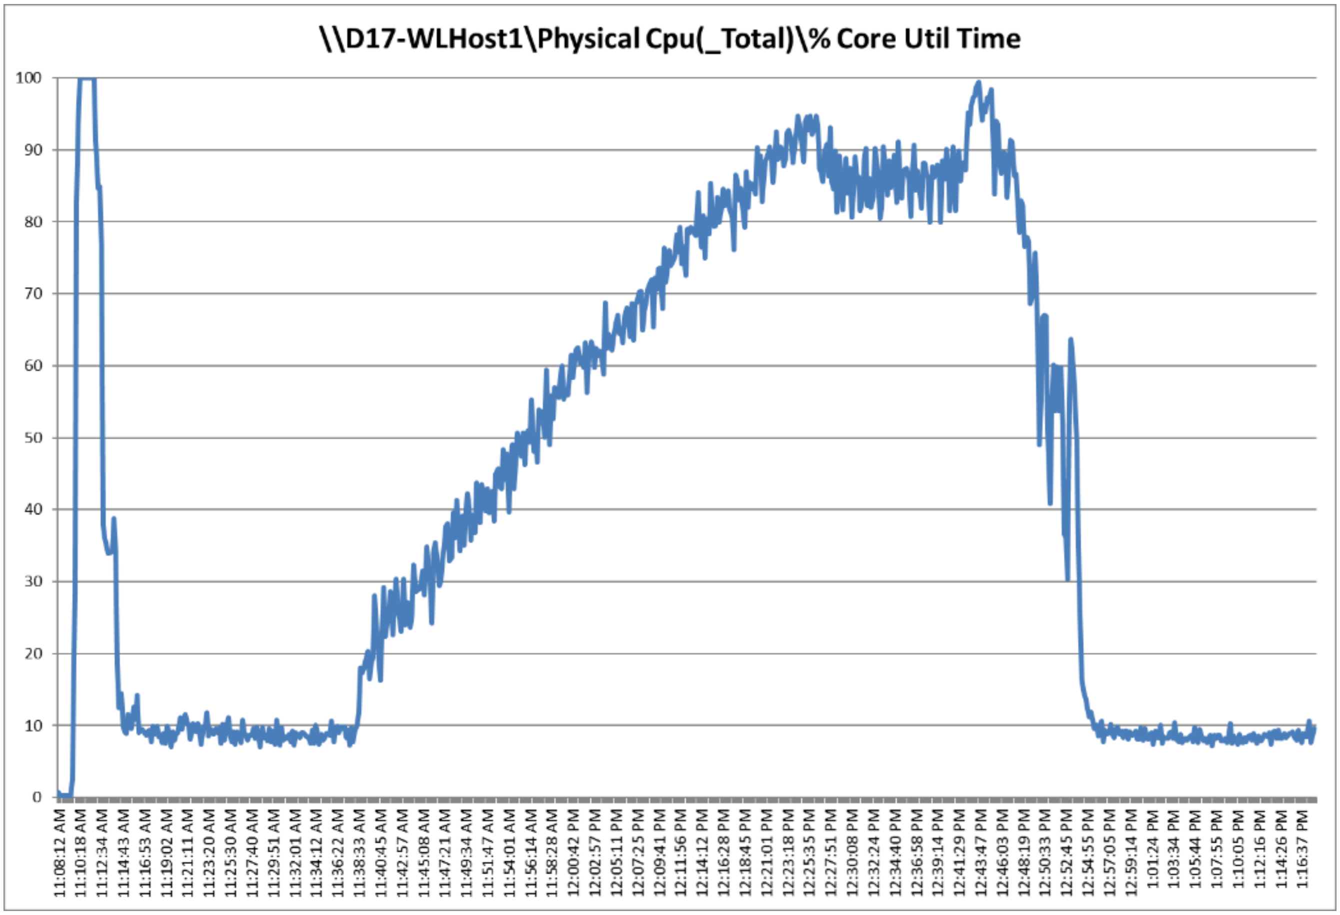

Figure 80. Single Server | Horizon 8 VDI-NP | Host Memory Utilization

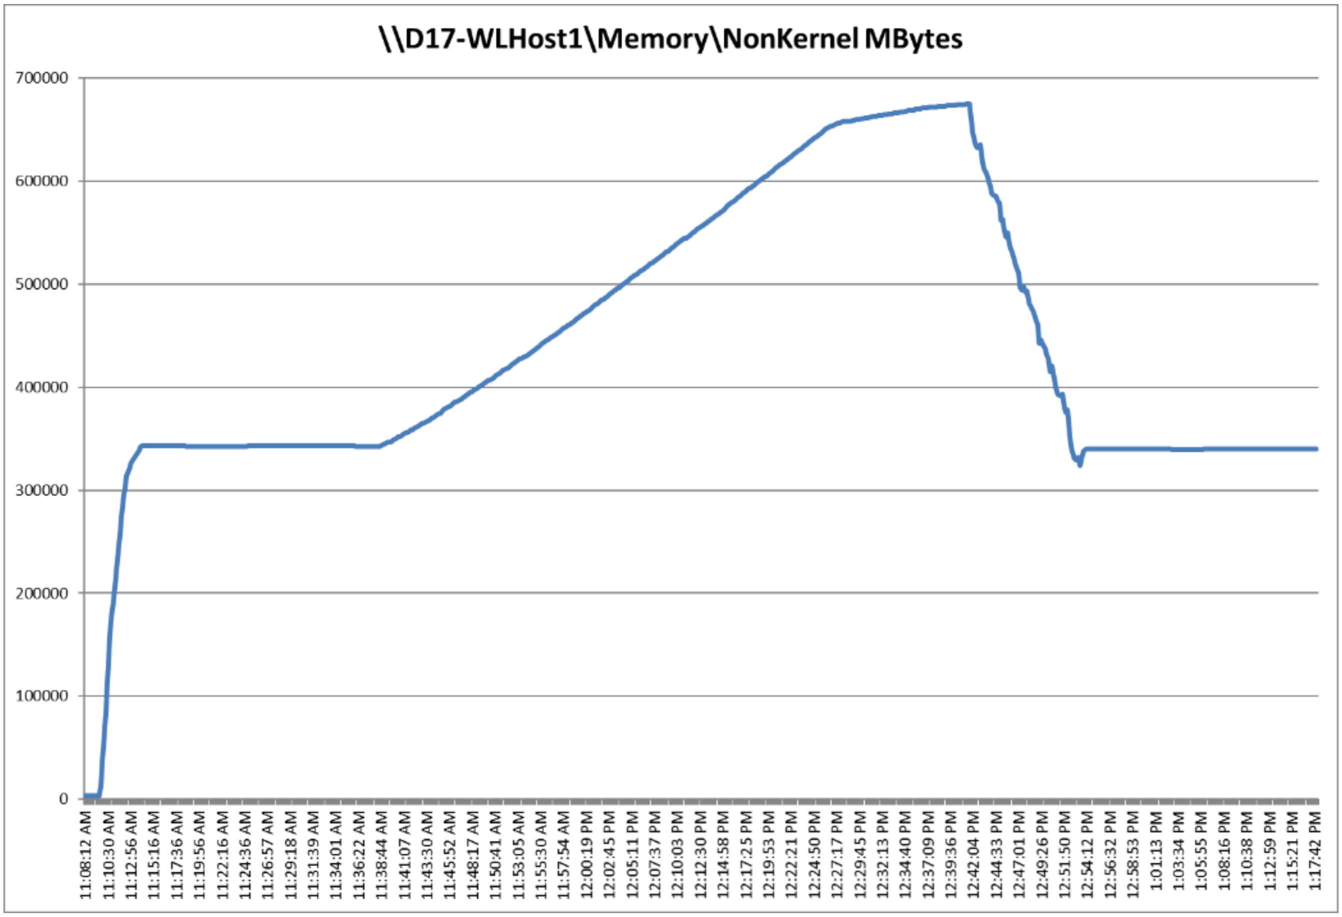

Figure 81. Single Server | Horizon 8 VDI-NP | Host Network Utilization

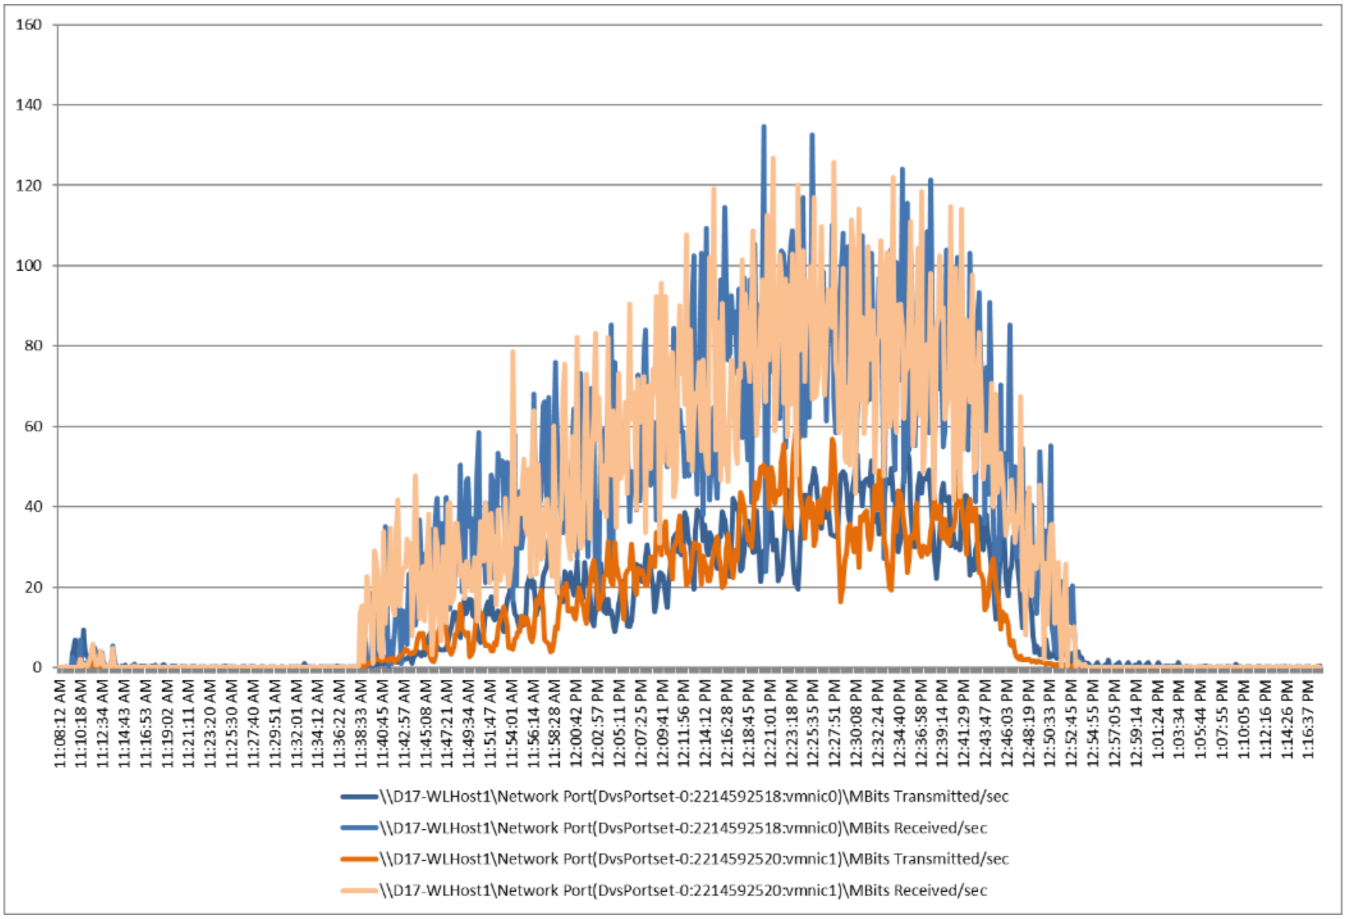

Single-Server Recommended Maximum Workload for VDI Persistent with 186 Users

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 186 Windows 10 64-bit VDI persistent virtual machines with 2 vCPU and 3.5GB RAM.

Login VSI performance data is as follows:

Figure 82. Single Server | Horizon 8 VDI-P | VSI Score

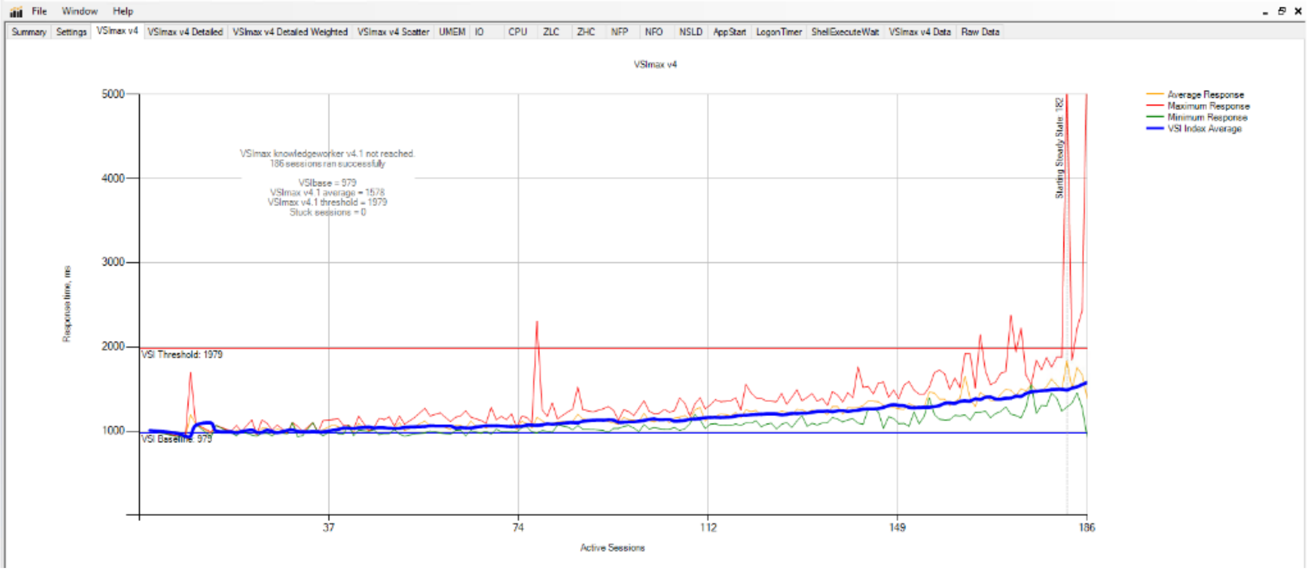

VSImax knowledgeworker v4.1 not reached.
186 sessions ran successfully

VSIbase = 979
VSImax v4.1 average = 1578
VSImax v4.1 threshold = 1979
Stuck sessions = 0

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

Figure 83. Single Server Recommended Maximum Workload | Horizon 8 VDI-P | Host CPU Utilization

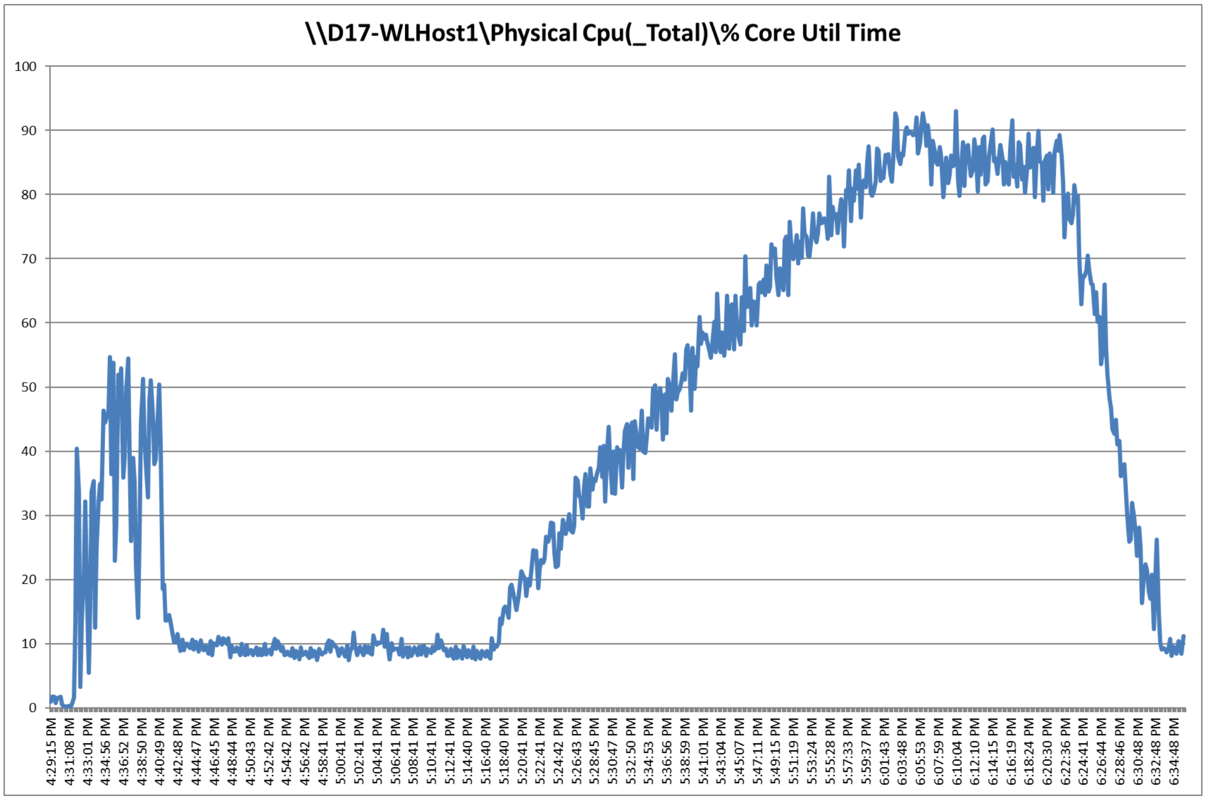

Figure 84. Single Server Recommended Maximum Workload | Horizon 8 VDI-P | Host Memory Utilization

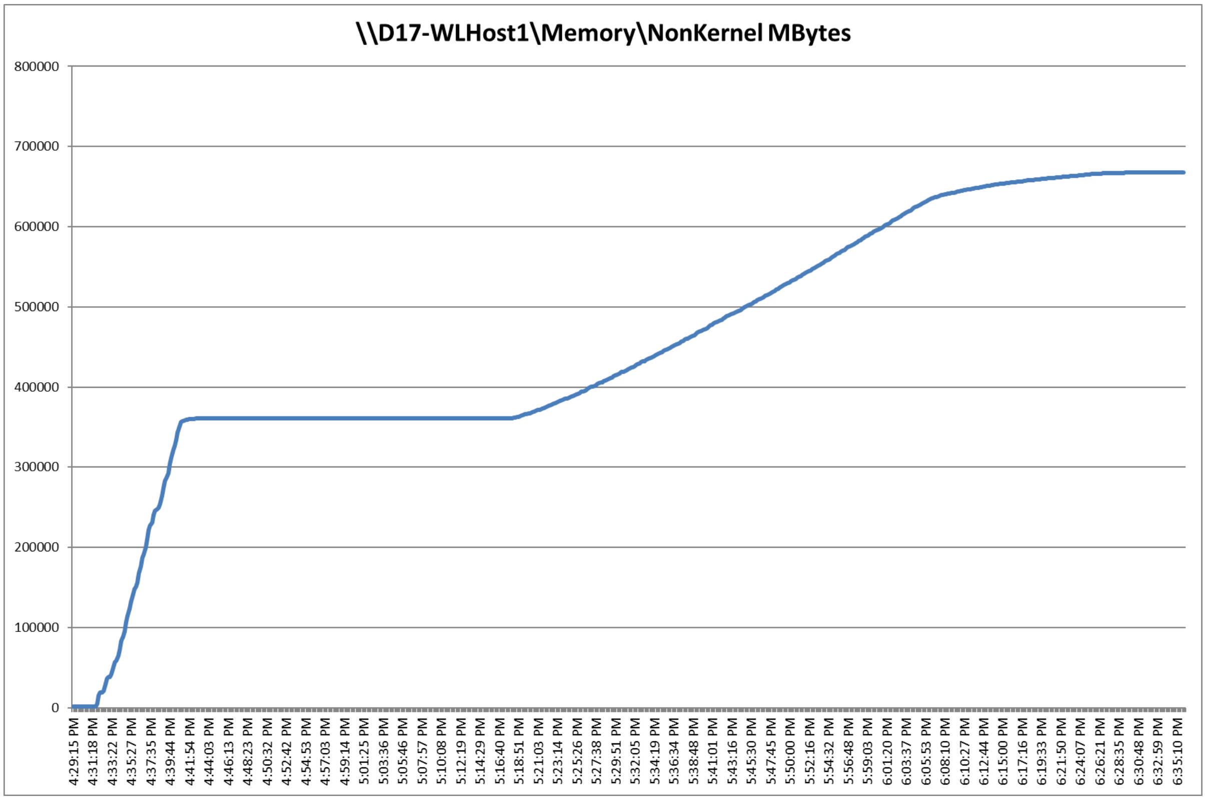

Figure 85. Single Server | Horizon 8 VDI-P | Host Network Utilization

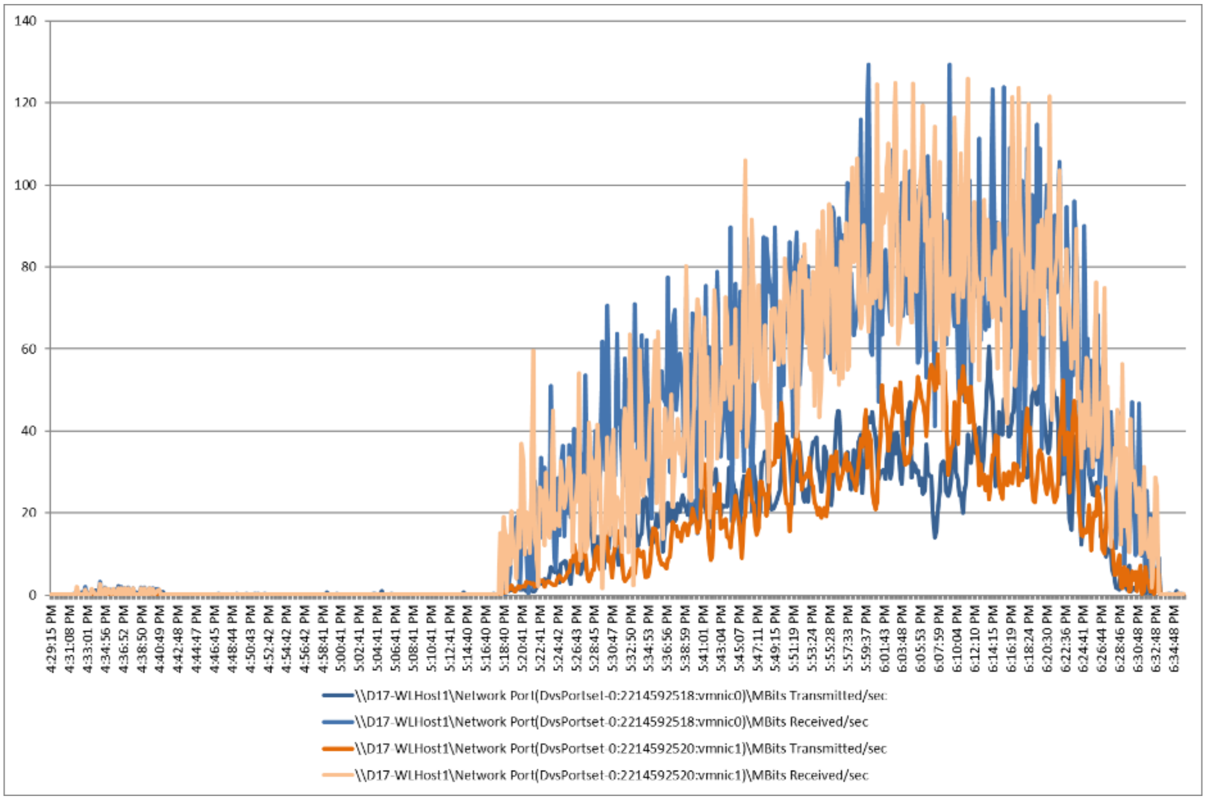

Cluster Recommended Maximum Workload Testing

This section shows the key performance metrics that were 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: 2016 RDS sessions, 1674 VDI Non-Persistent sessions, and 1674 VDI Persistent sessions.

Cluster Workload Testing with 2016 RDS Users

The cluster testing was comprised of 2016 RDS sessions using 10 workload blades.

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 86. **Ten Node Cluster | 2016 RDS Users | VSI Score**

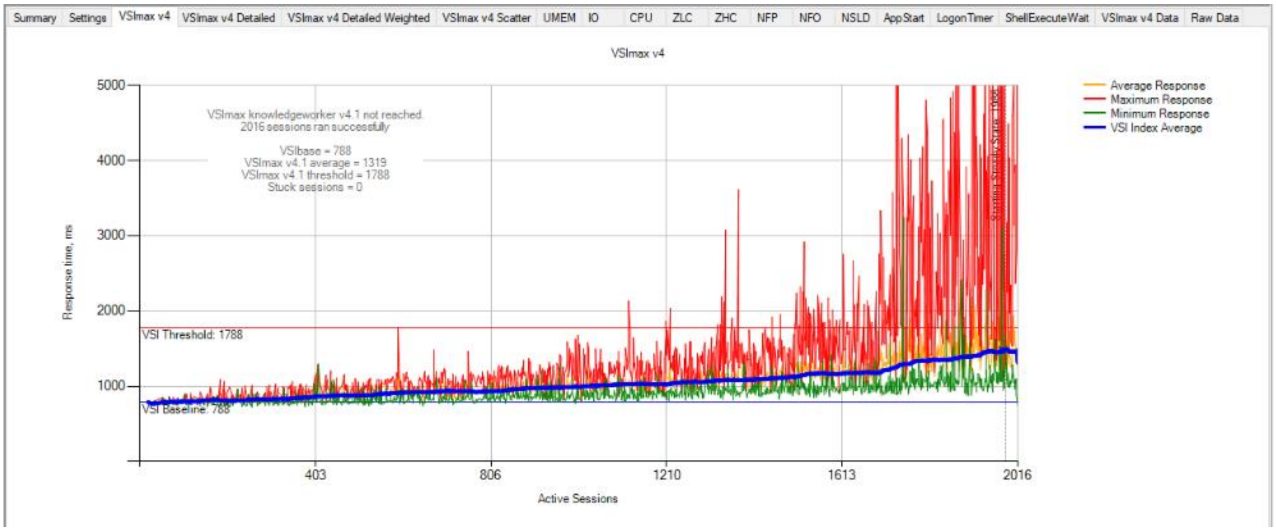

VSImax knowledgeworker v4.1 not reached.
2016 sessions ran successfully

VSIbase = 788
VSImax v4.1 average = 1319
VSImax v4.1 threshold = 1788
Stuck sessions = 0

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

Figure 87. **Cluster | 2016 RDS Users | 10 RDS Hosts | Host CPU Utilization**

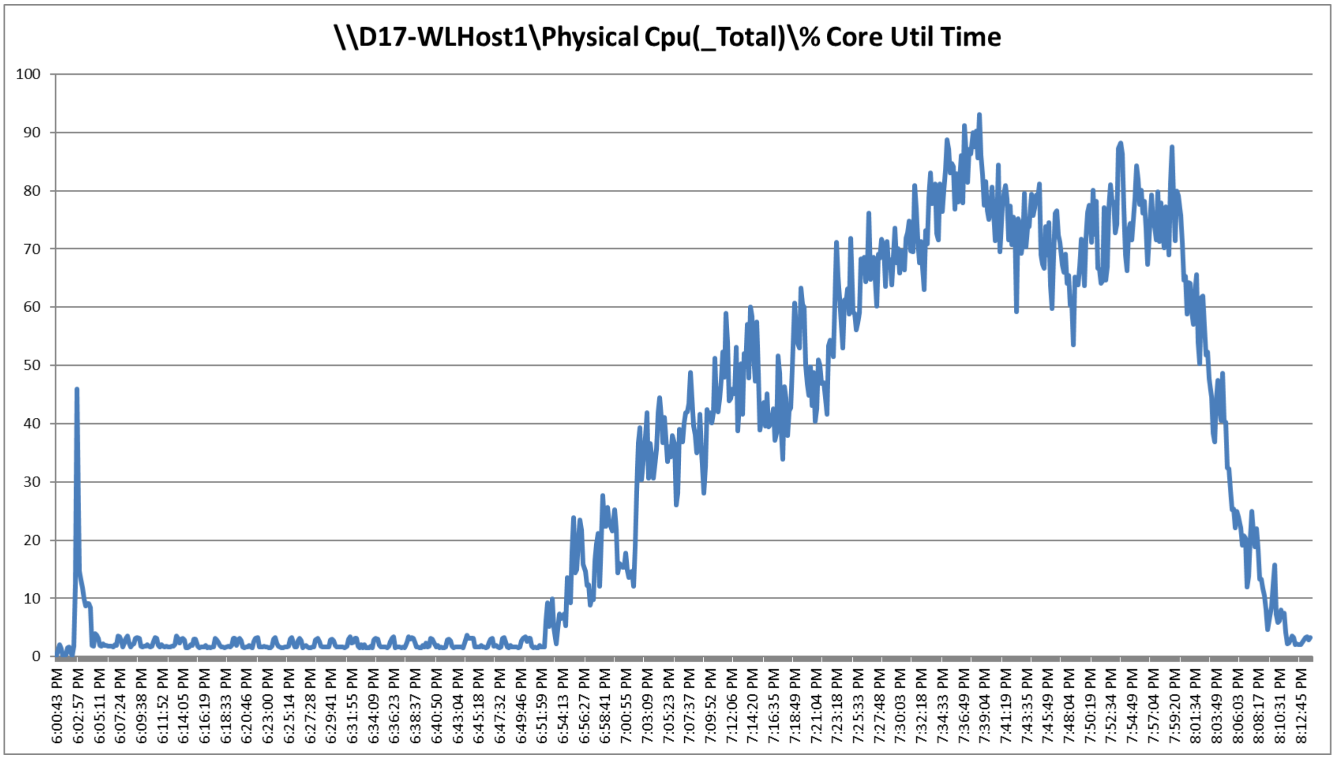

Figure 88. Cluster | 2016 RDS Users | 10 RDS Hosts | Host Memory Utilization

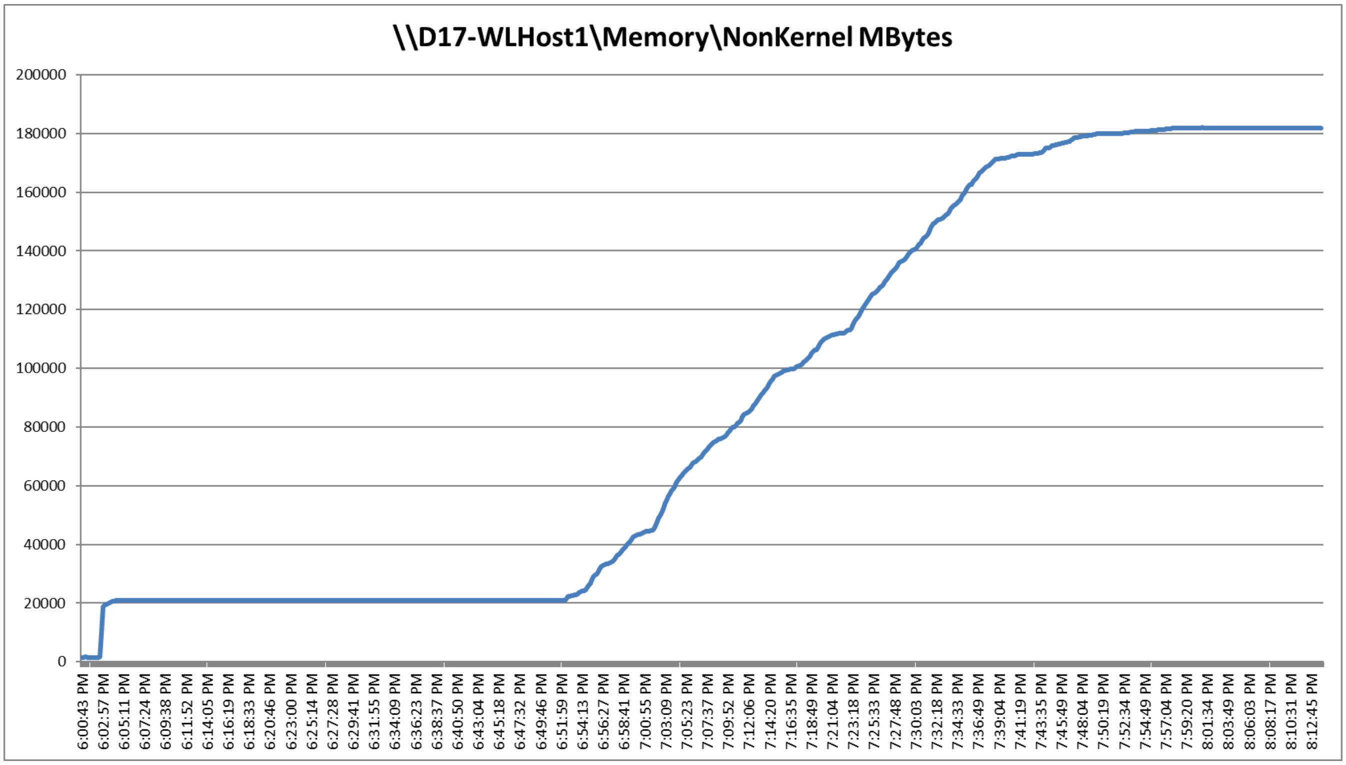

Figure 89. Cluster | 2016 RDS Users | 10 RDS Hosts | Host System Uplink Network Utilization

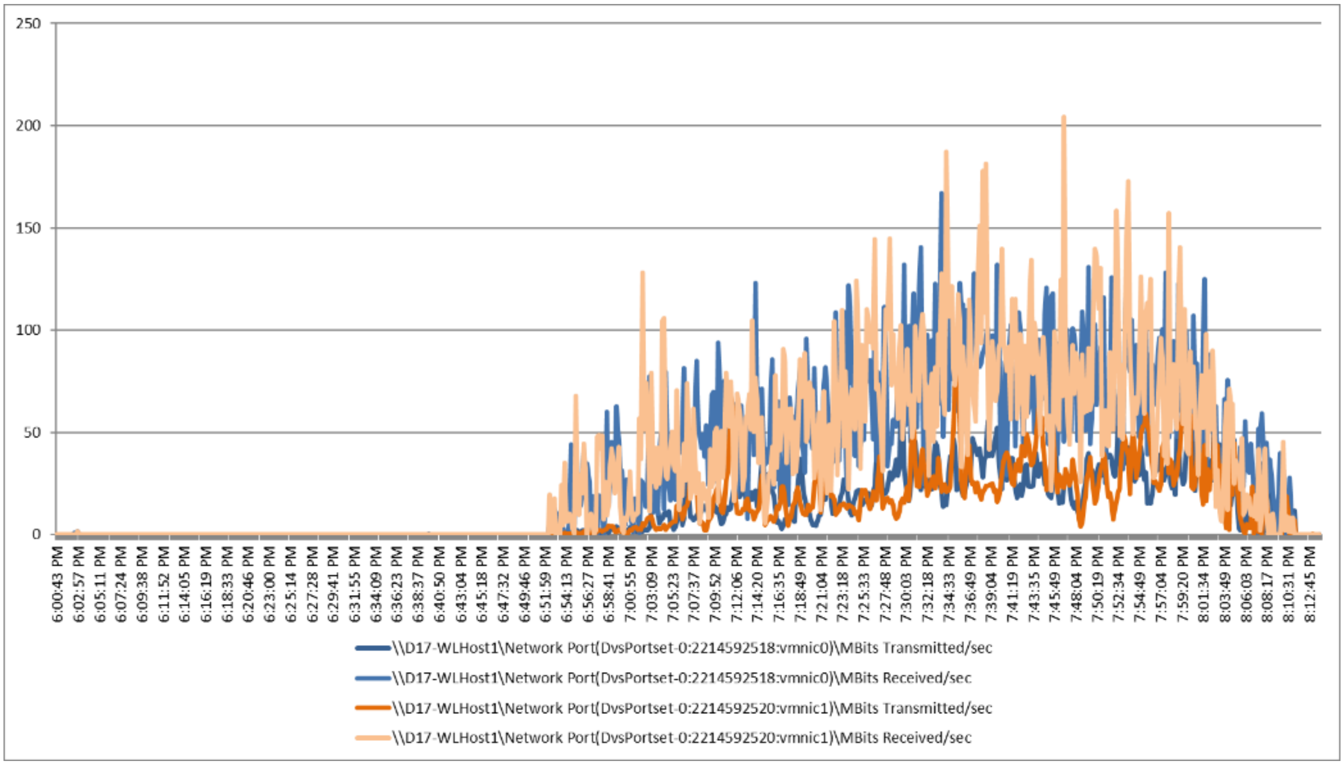

Cluster Workload Testing with 1674 Non-Persistent Desktop Users

The cluster testing comprised of 1674 VDI non-persistent desktop sessions using 10 workload blades.

The workload for the test is 1674 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 90. **Cluster | 1674 VDI-NP Users | VSI Score**

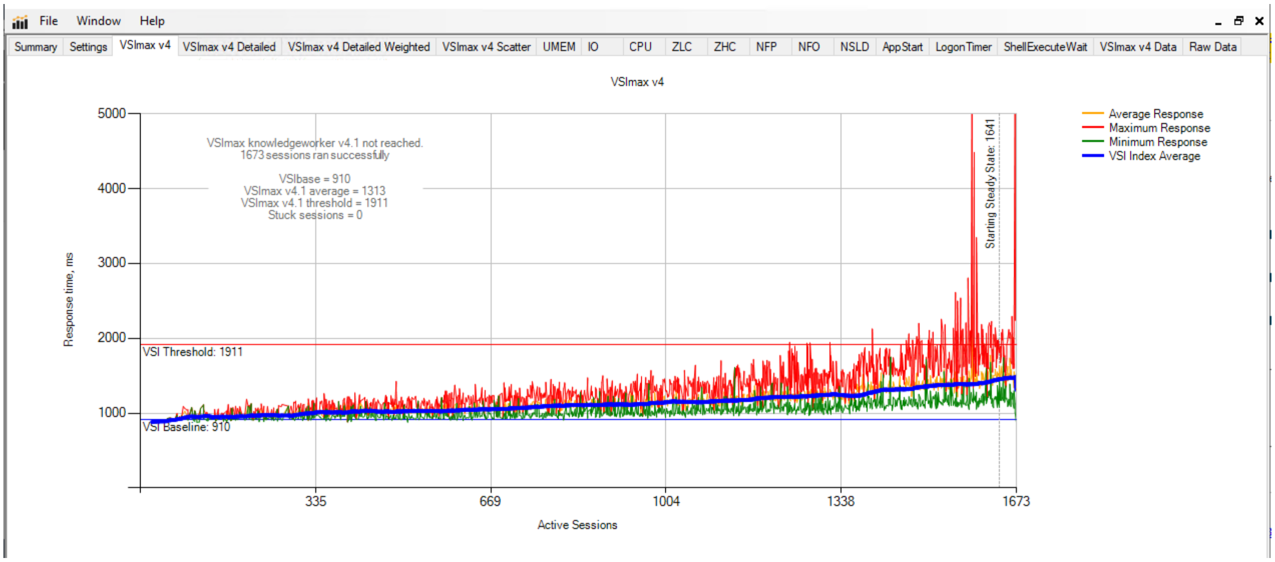

VSImax knowledgeworker v4.1 not reached.
1673 sessions ran successfully

VSIbase = 910
VSImax v4.1 average = 1313
VSImax v4.1 threshold = 1911
Stuck sessions = 0

Figure 91. Cluster | 1674 VDI-NP Users | Non-Persistent Hosts | Host CPU Utilization

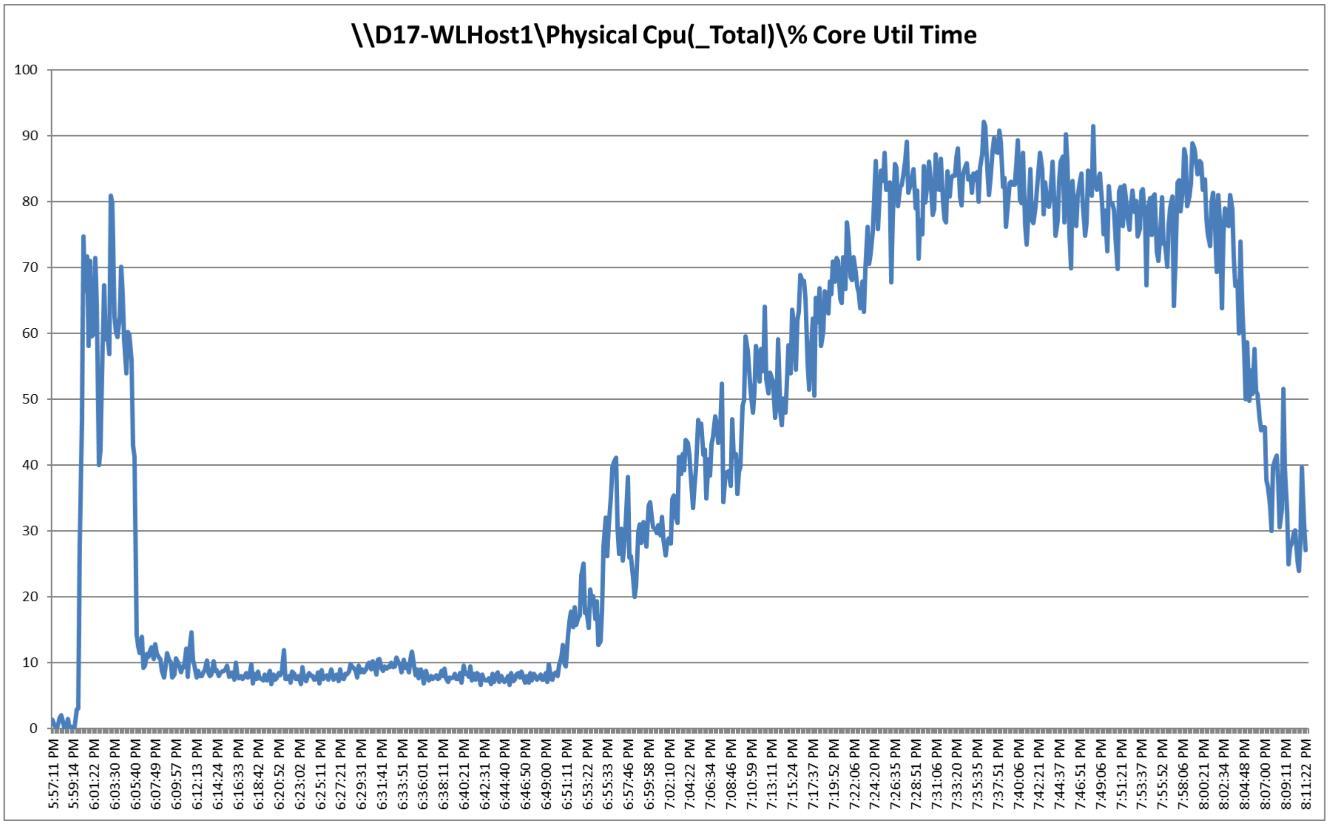

Figure 92. Cluster | 1674 VDI-NP Users | Non-Persistent Hosts | Host Memory Utilization

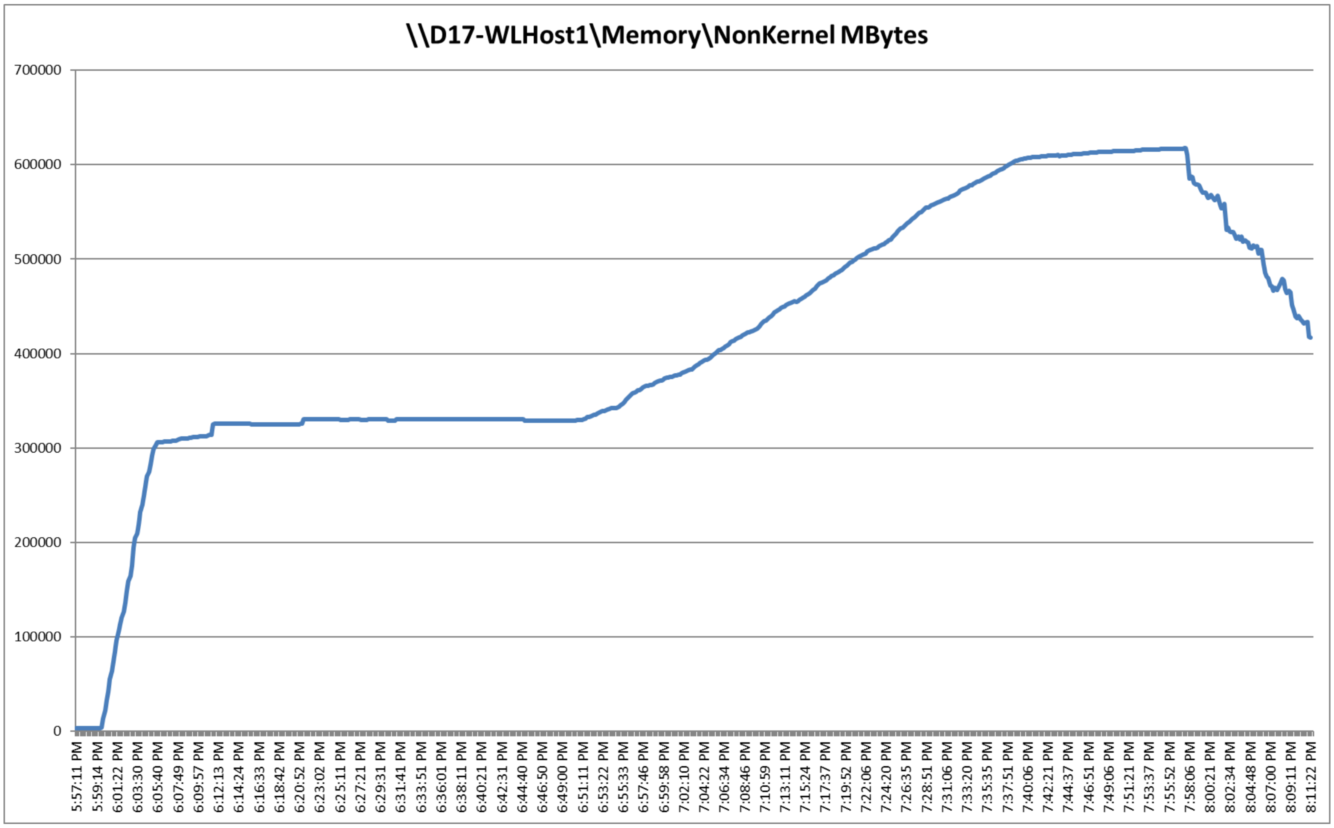

Figure 93. Cluster | 1674 VDI-NP Users | Non-Persistent Hosts | Host Network Utilization

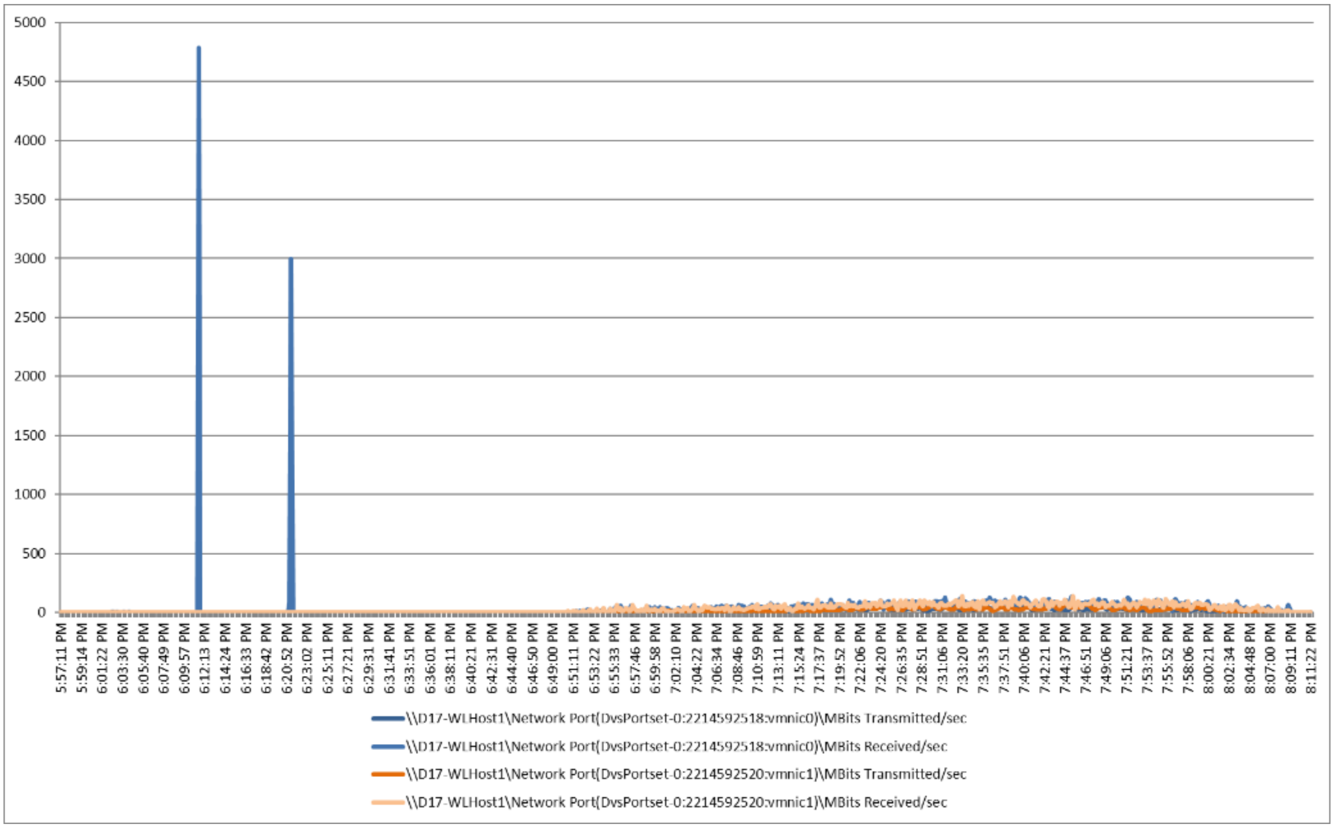

Cluster Workload Testing with 1674 Persistent Desktop Users

This section describes the key performance metrics that were captured on the Cisco UCS, Pure Storage FlashArray array, and Infrastructure virtual machines during the persistent desktop testing. The cluster testing with comprised of 1674 VDI Persistent desktop sessions using 10 workload blades.

The workload for the test is 1674 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 94. Cluster | 1674 VDI-P Users | VSI Score

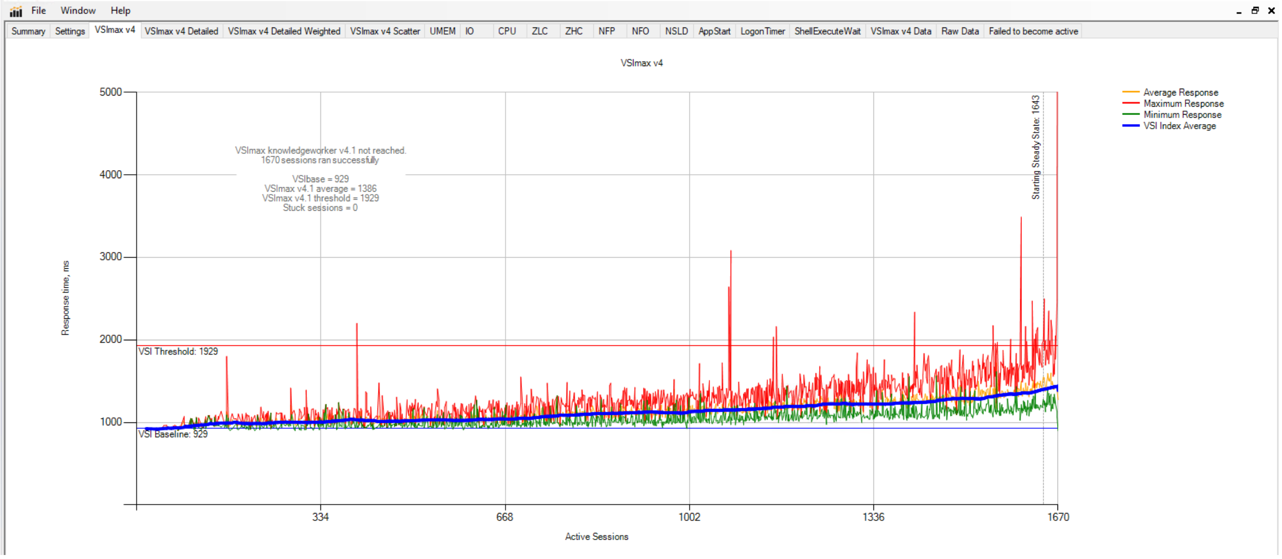

VSImax knowledgeworker v4.1 not reached.
1670 sessions ran successfully

VSIbase = 929
VSImax v4.1 average = 1386
VSImax v4.1 threshold = 1929
Stuck sessions = 0

Figure 95. Cluster | 1674 VDI-P Users | Persistent Hosts | Host CPU Utilization

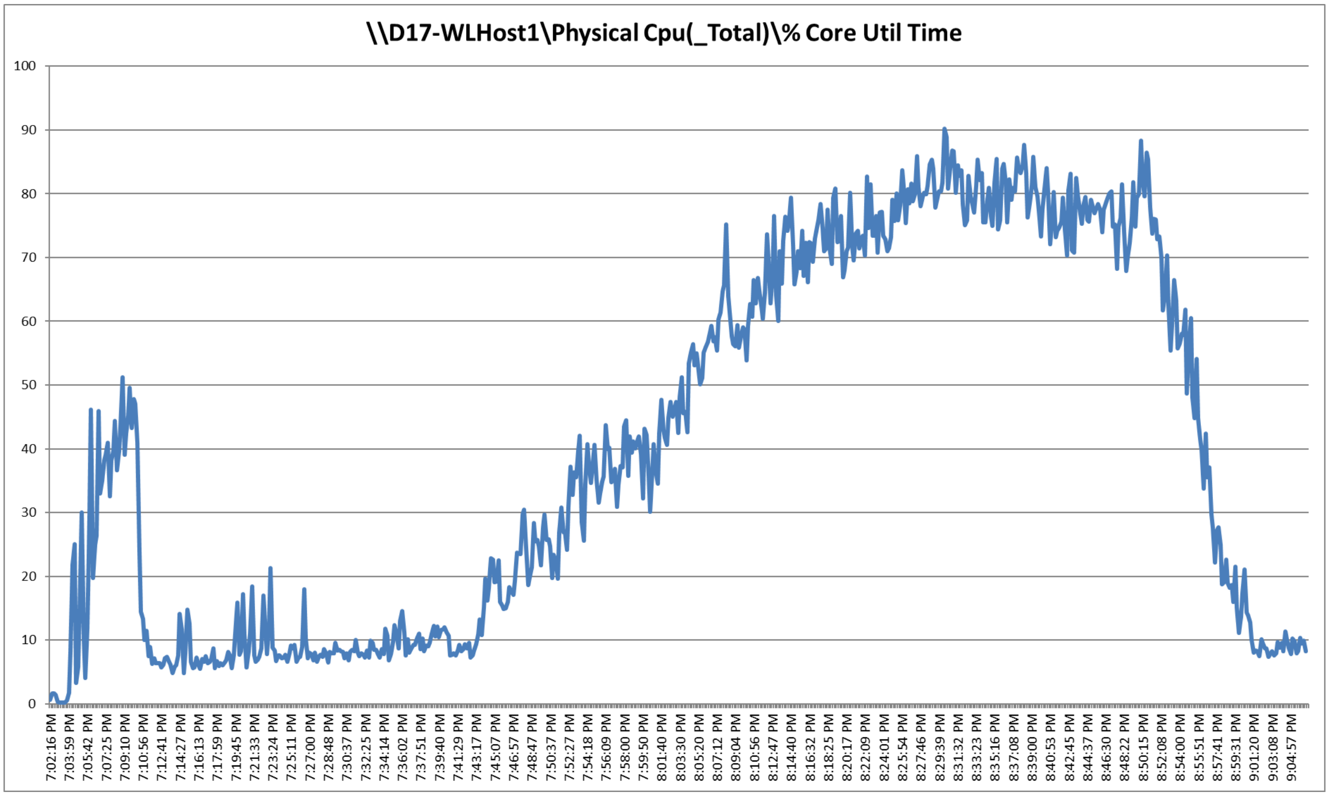

Figure 96. Cluster | 1674 VDI-P Users | Persistent Hosts | Host Memory Utilization

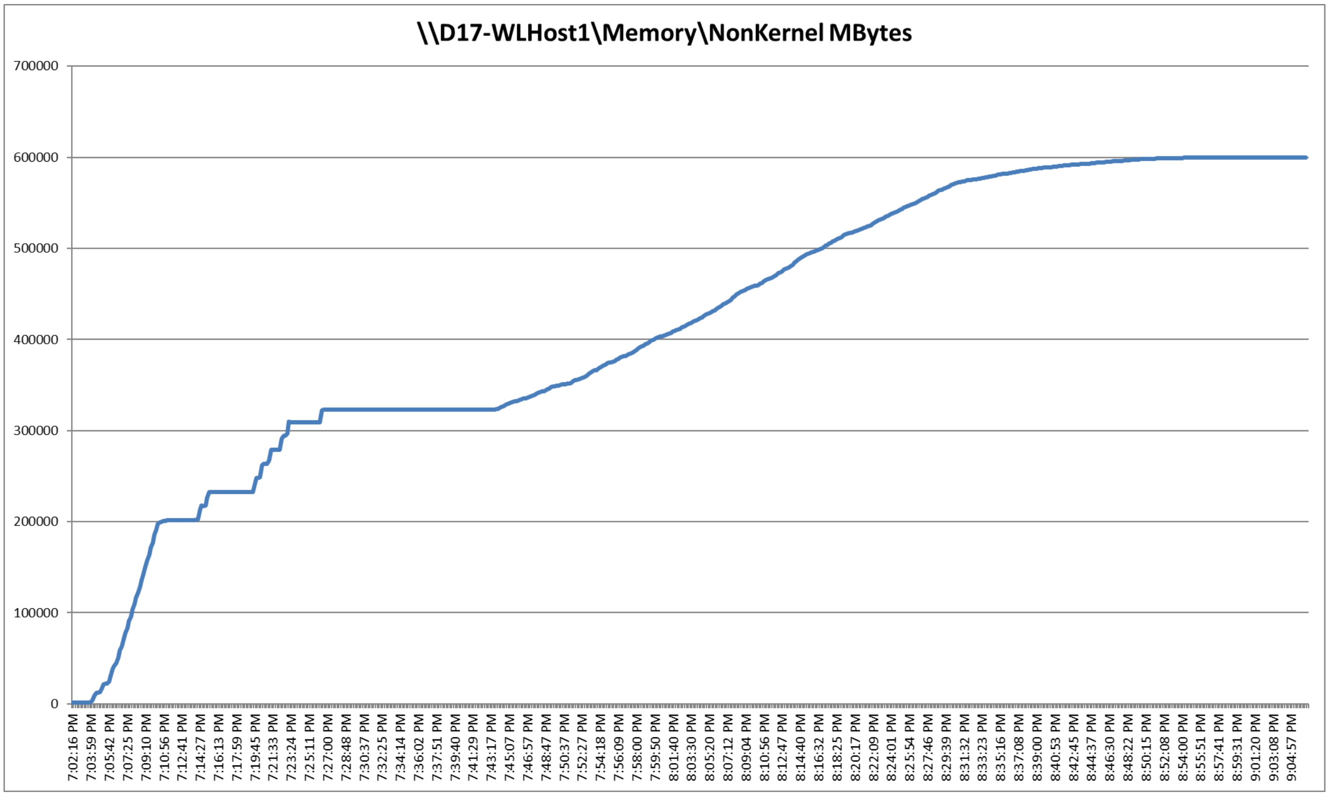

Figure 97. **Cluster | 1800 VDI-P Users | Persistent Hosts | Host Network Utilization**

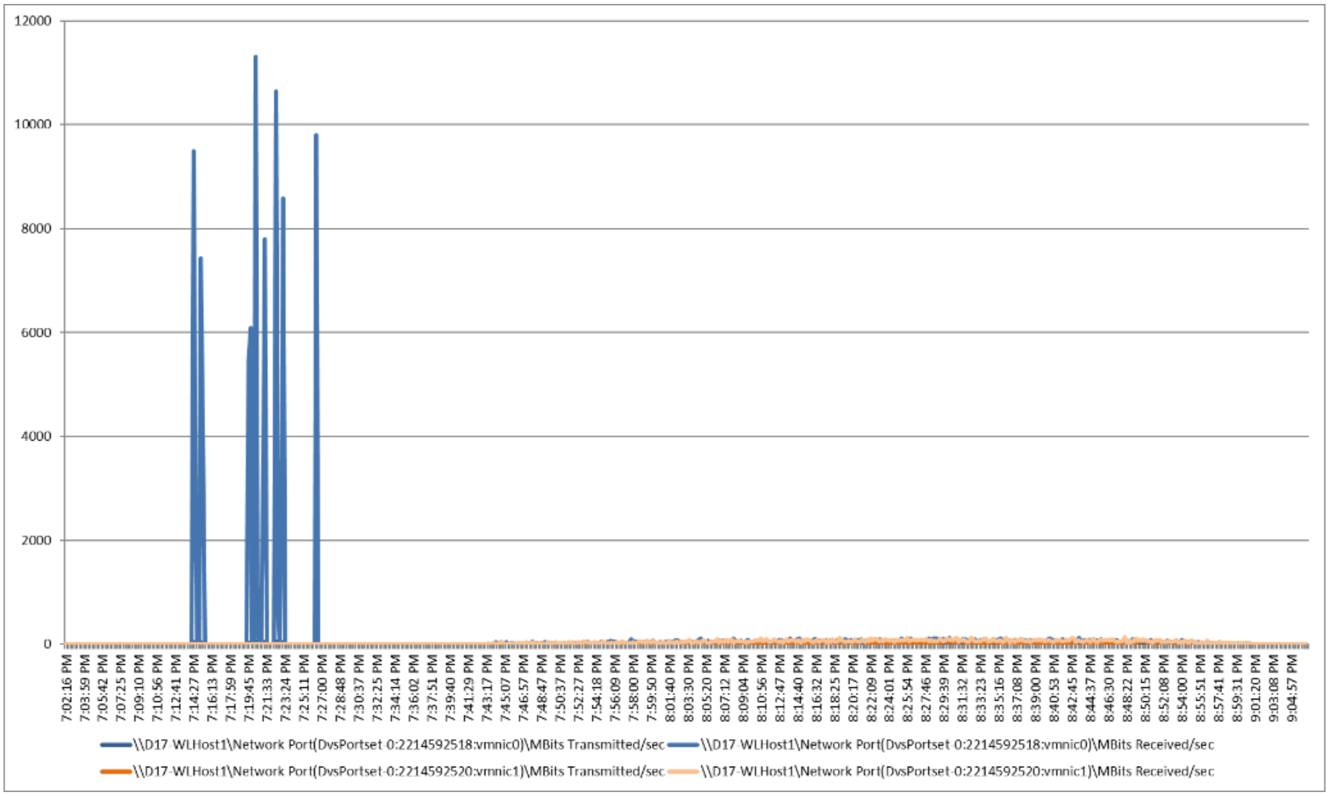

Full Scale Workload Testing

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. Full Scale testing Was done with following Workloads using 30 Hosts (configured in three 10 Host Pools) for workload and 2 hosts for Infrastructure VMs:

- RDS Test -6048 sessions
- VDI Non-Persistent Desktop test- 5022 users
- VDI Persistent Desktop test- 5022 users

To achieve the target, sessions were launched against each workload set at a time. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

Full Scale Workload Testing with 6048 RDS Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray//X70 R3 array, during the RDS full-scale testing with 6048 Desktop Sessions using 30 blades (configured in three 10 Host Pools).

The RDS workload for the solution is 6048 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.

Figure 98. Full Scale 6048 RDS Sessions Test - Workload Distribution

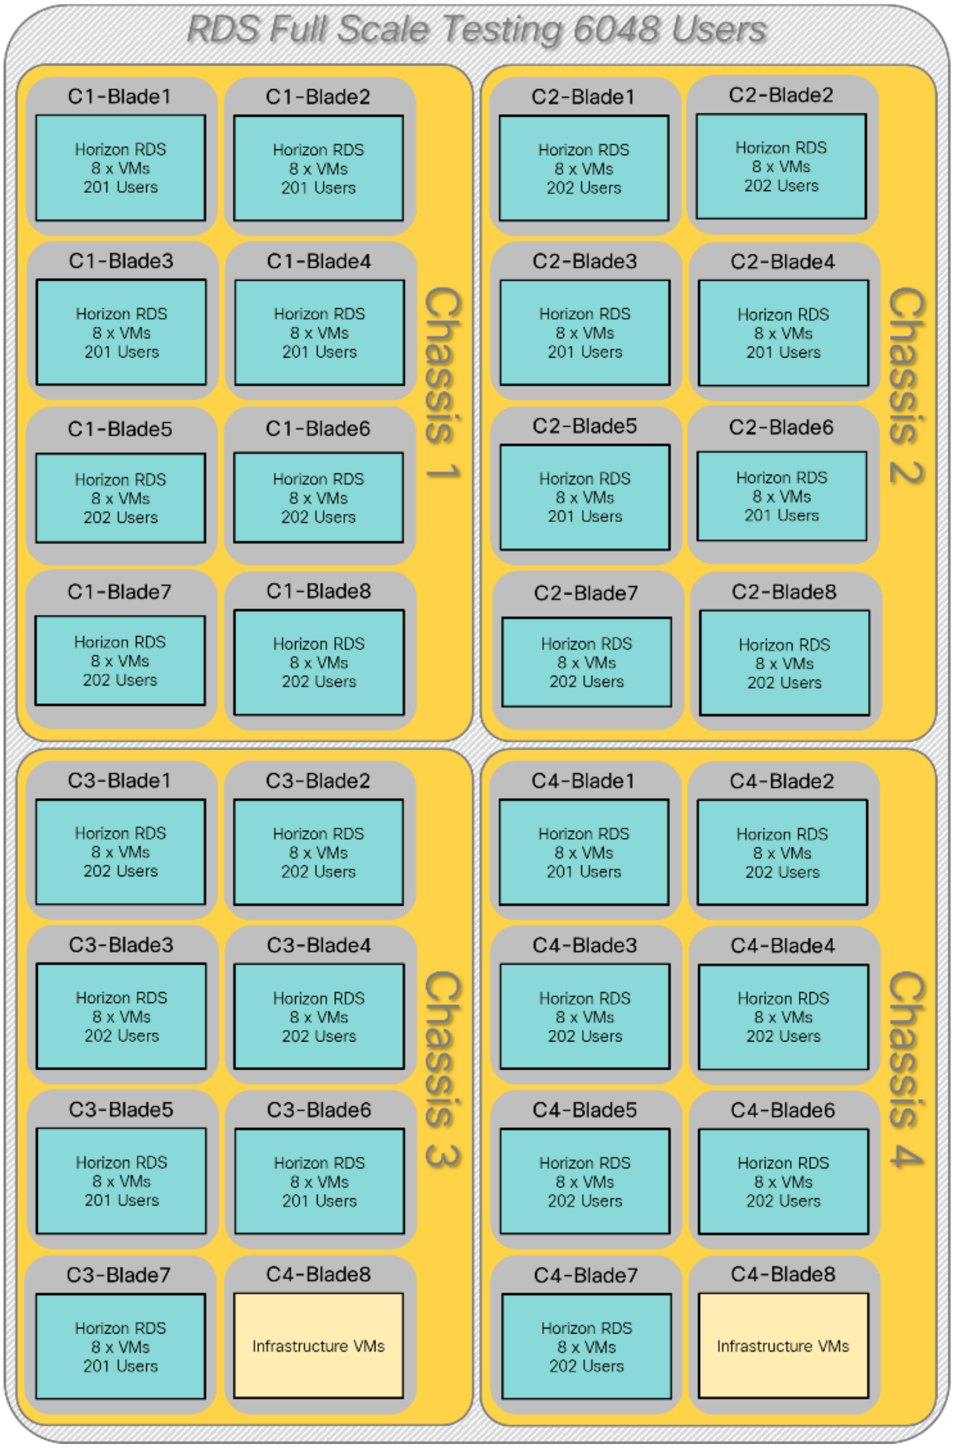

RDS Full Scale Testing 6048 Users

Chassis 1

C1-Blade1
Horizon RDS
8 x VMs
201 Users

C1-Blade2
Horizon RDS
8 x VMs
201 Users

C1-Blade3
Horizon RDS
8 x VMs
201 Users

C1-Blade4
Horizon RDS
8 x VMs
201 Users

C1-Blade5
Horizon RDS
8 x VMs
202 Users

C1-Blade6
Horizon RDS
8 x VMs
202 Users

C1-Blade7
Horizon RDS
8 x VMs
202 Users

C1-Blade8
Horizon RDS
8 x VMs
202 Users

Chassis 2

C2-Blade1
Horizon RDS
8 x VMs
202 Users

C2-Blade2
Horizon RDS
8 x VMs
202 Users

C2-Blade3
Horizon RDS
8 x VMs
201 Users

C2-Blade4
Horizon RDS
8 x VMs
201 Users

C2-Blade5
Horizon RDS
8 x VMs
201 Users

C2-Blade6
Horizon RDS
8 x VMs
201 Users

C2-Blade7
Horizon RDS
8 x VMs
202 Users

C2-Blade8
Horizon RDS
8 x VMs
202 Users

Chassis 3

C3-Blade1
Horizon RDS
8 x VMs
202 Users

C3-Blade2
Horizon RDS
8 x VMs
202 Users

C3-Blade3
Horizon RDS
8 x VMs
202 Users

C3-Blade4
Horizon RDS
8 x VMs
202 Users

C3-Blade5
Horizon RDS
8 x VMs
201 Users

C3-Blade6
Horizon RDS
8 x VMs
201 Users

C3-Blade7
Horizon RDS
8 x VMs
201 Users

C4-Blade8
Infrastructure VMs

Chassis 4

C4-Blade1
Horizon RDS
8 x VMs
201 Users

C4-Blade2
Horizon RDS
8 x VMs
202 Users

C4-Blade3
Horizon RDS
8 x VMs
202 Users

C4-Blade4
Horizon RDS
8 x VMs
202 Users

C4-Blade5
Horizon RDS
8 x VMs
202 Users

C4-Blade6
Horizon RDS
8 x VMs
202 Users

C4-Blade7
Horizon RDS
8 x VMs
202 Users

C4-Blade8
Infrastructure VMs

The configured system efficiently and effectively delivered the following results:

Figure 99. Full Scale | 6048 Users | Win 2019 RDSH | VSI Score

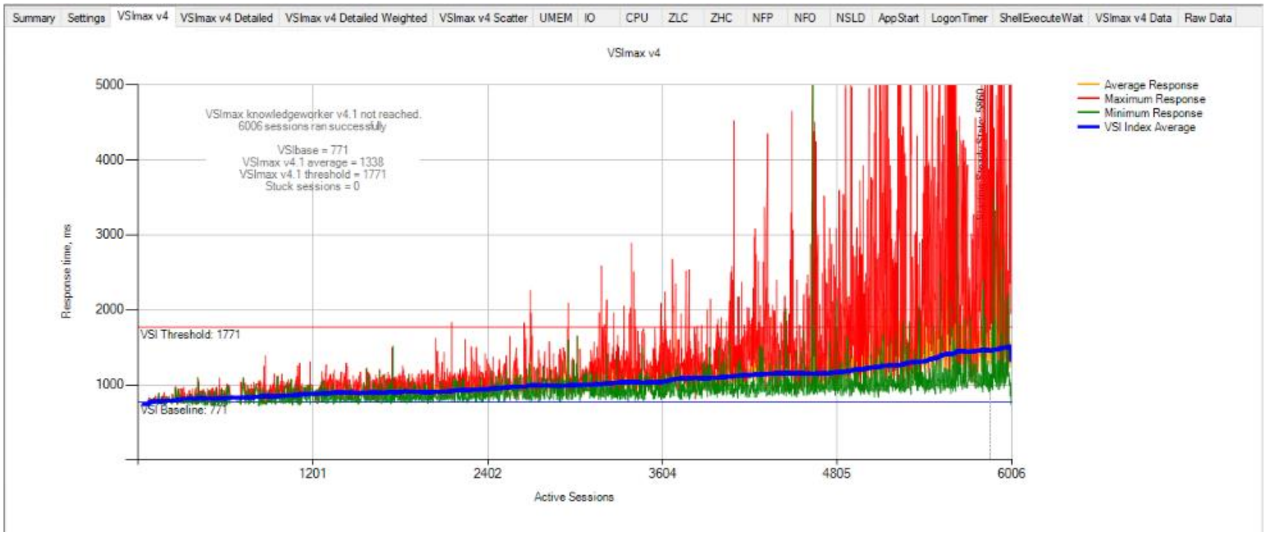

VSImax knowledgeworker v4.1 not reached.
6006 sessions ran successfully

VSIbase = 771
VSImax v4.1 average = 1338
VSImax v4.1 threshold = 1771
Stuck sessions = 0

Figure 100. Full Scale | 6048 Users | Win 2019 RDS Hosts| Host CPU Utilization

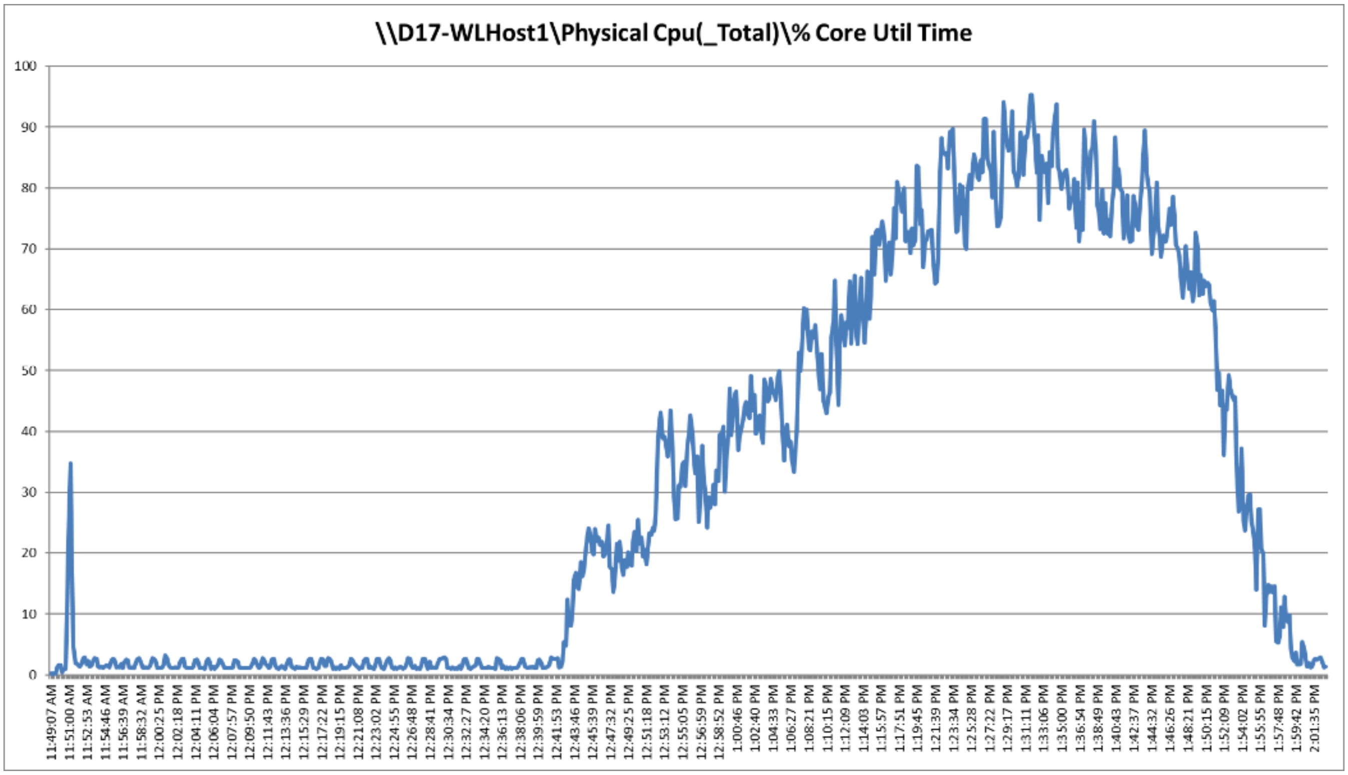

Figure 101. Full Scale | 6048 Users | Win 2019 RDS Hosts | Host Memory Utilization

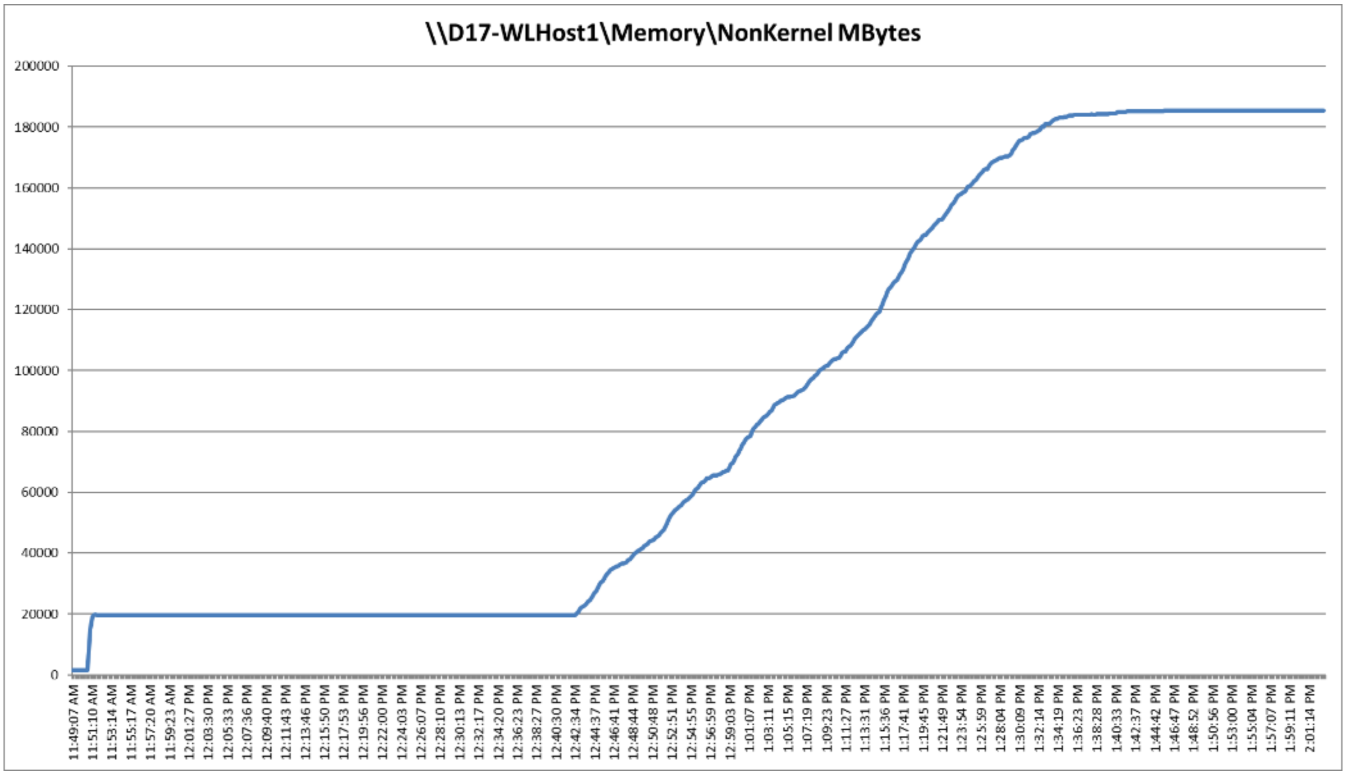

Figure 102. Full Scale | 6048 Users | Win 2019 RDS Hosts | Host Network Utilization

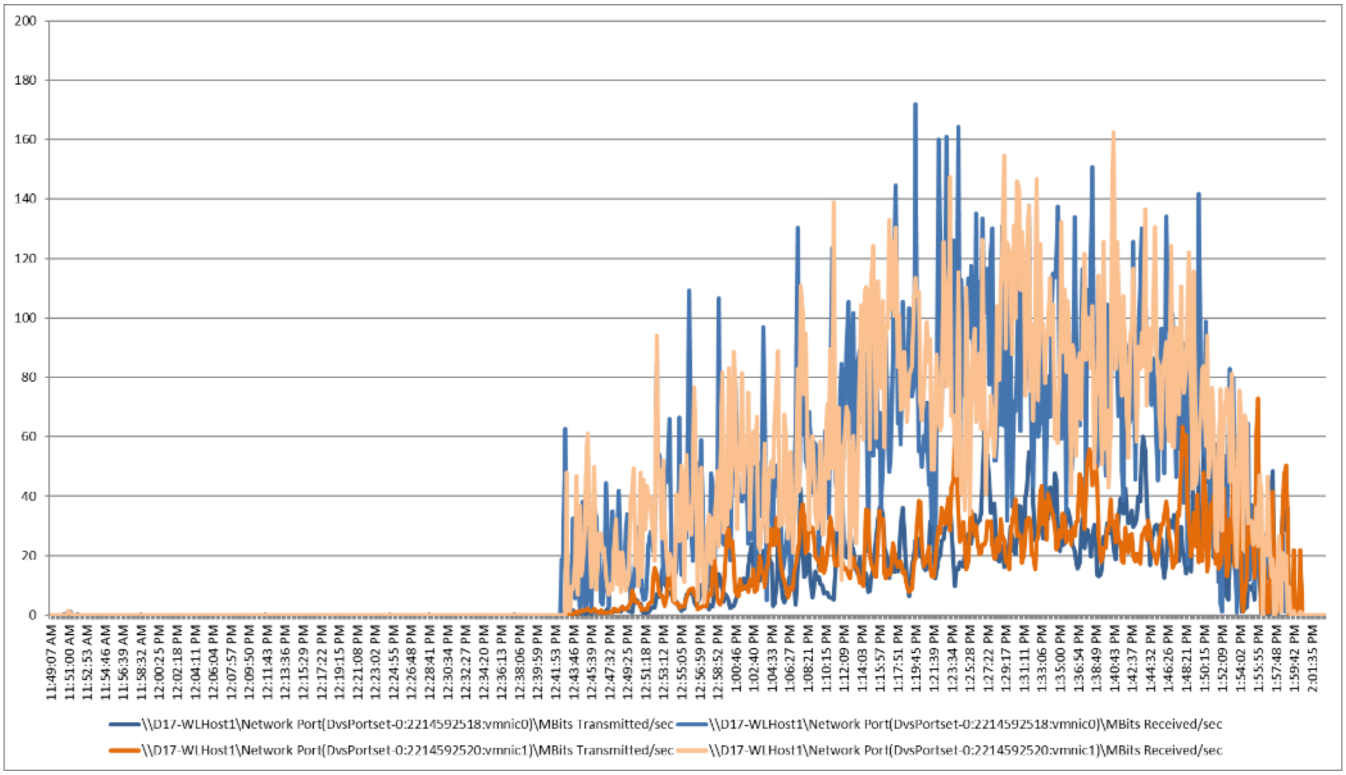

Pure Storage FlashArray//X70 R3 Storage System Graph for 6048 Users RDS Workload Test

Figure 103. Full Scale | 6048 Users | Win 2019 RDS Hosts | Pure Storage FlashArray//X70 R3 System Latency Chart

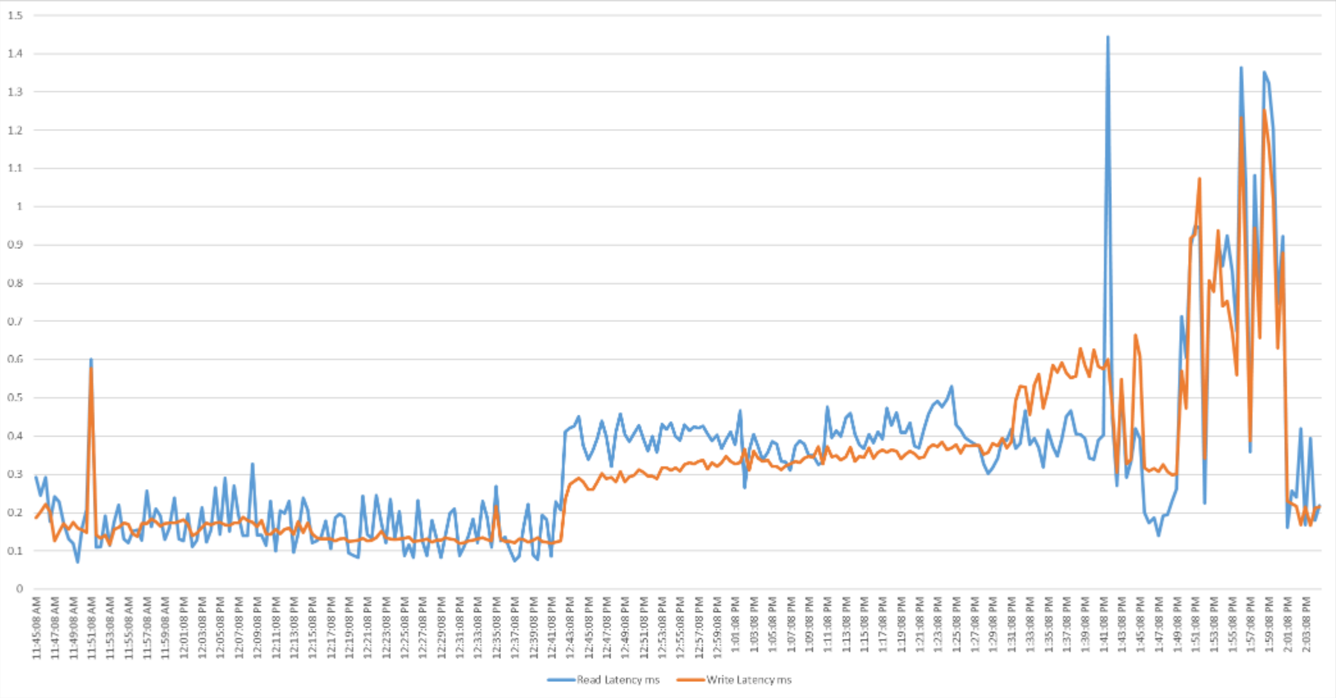

Figure 104. Full Scale | 6048 Users | Win 2019 RDS Hosts | Pure Storage FlashArray//X70 R3 System IOPS Chart

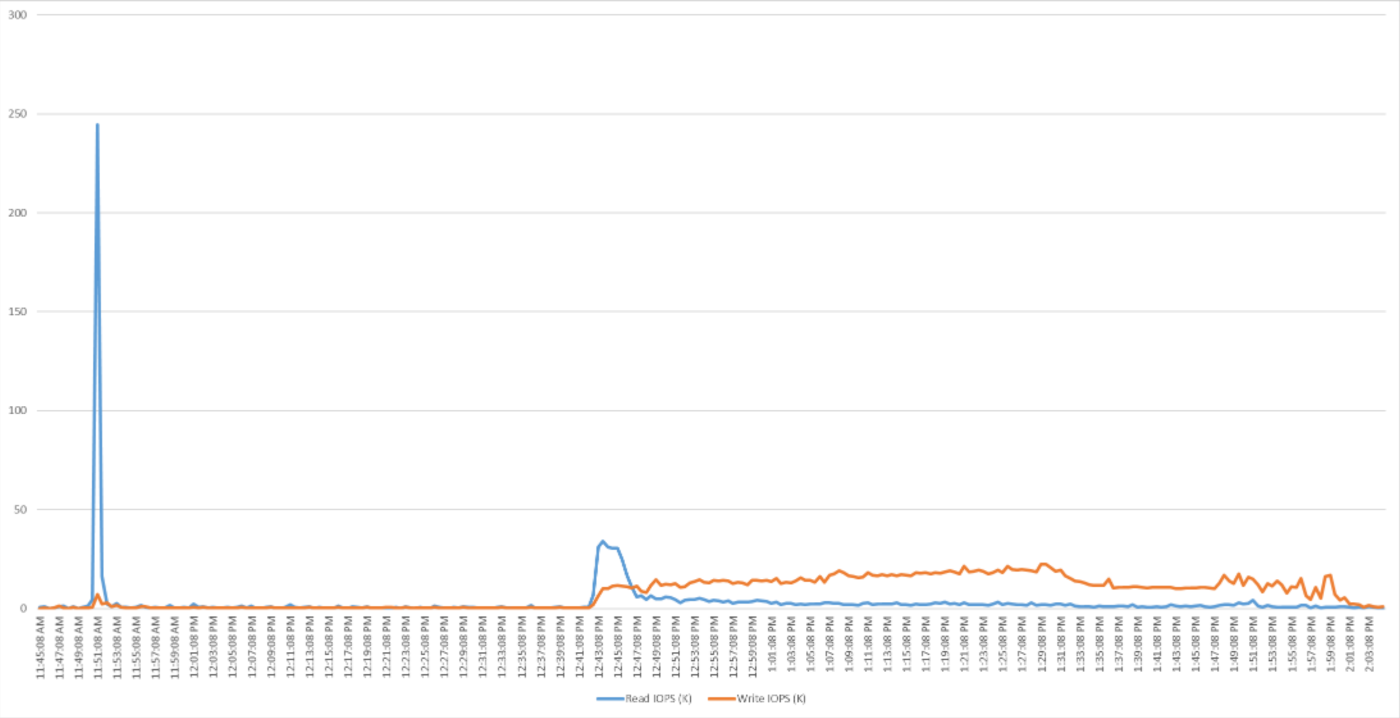

Figure 105. Full Scale | 6048 Users | Win 2019 RDS Hosts | Pure Storage FlashArray//X70 R3 System Bandwidth Chart

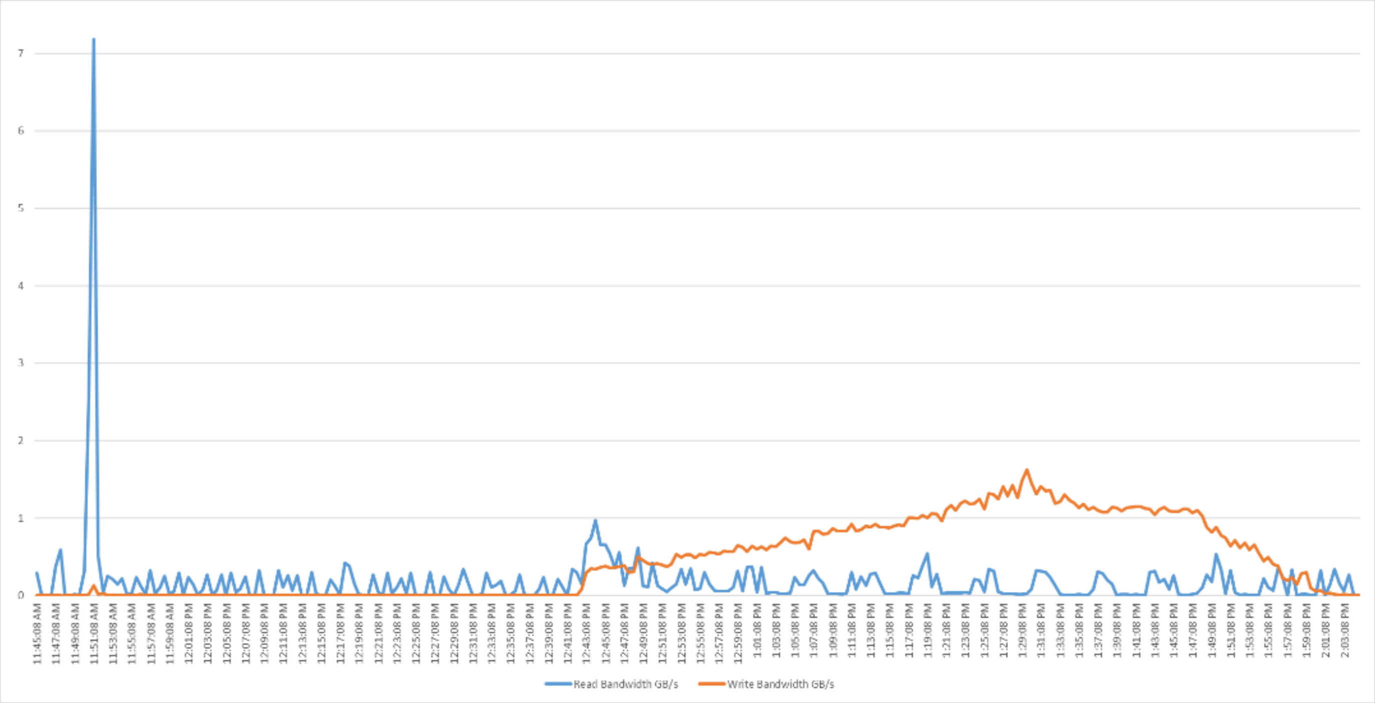

Figure 106. Full Scale | 6048 Users | Win 2019 RDS Hosts | FlashArray//X70 R3 System Performance Chart

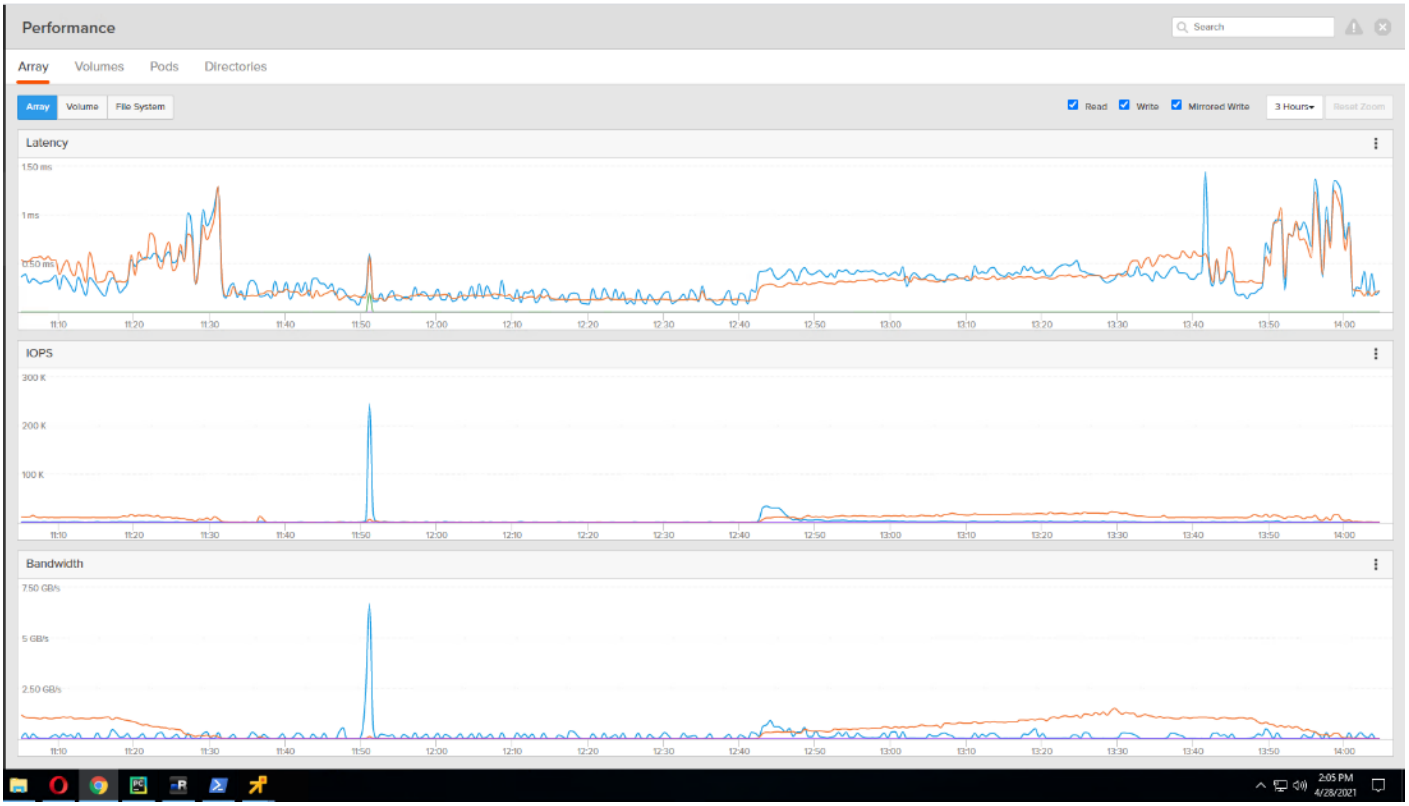

Full Scale Workload Testing with 5022 VDI Non-Persistent Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray//X70 R3 array during the persistent desktop full-scale testing with 5022 VDI Non-Persistent desktops using 30 blades (configured in three 10 Host Pools).

The workload for the test is 5022 Non-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.

Figure 107. Full Scale 5022 User Non-Persistent VDI Test – Workload Distribution

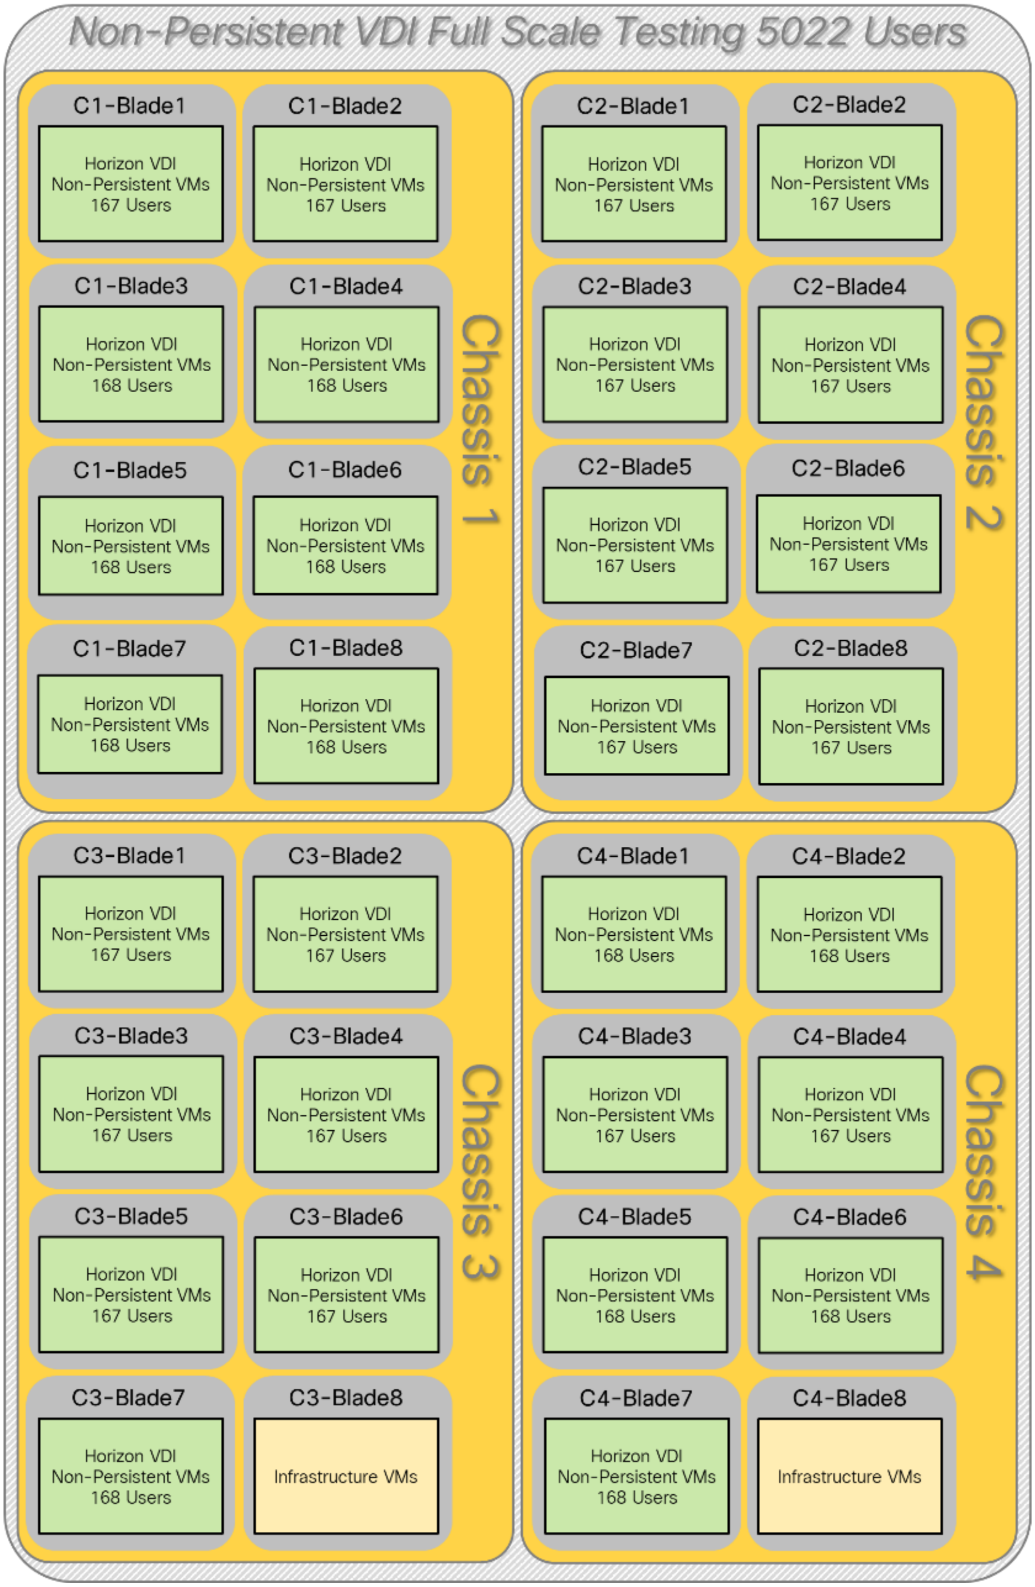
Non-Persistent VDI Full Scale Testing 5022 Users
Chassis 1
C1-Blade1 Horizon VDI Non-Persistent VMs 167 Users
C1-Blade2 Horizon VDI Non-Persistent VMs 167 Users
C1-Blade3 Horizon VDI Non-Persistent VMs 168 Users
C1-Blade4 Horizon VDI Non-Persistent VMs 168 Users
C1-Blade5 Horizon VDI Non-Persistent VMs 168 Users
C1-Blade6 Horizon VDI Non-Persistent VMs 168 Users
C1-Blade7 Horizon VDI Non-Persistent VMs 168 Users
C1-Blade8 Horizon VDI Non-Persistent VMs 168 Users
Chassis 2
C2-Blade1 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade2 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade3 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade4 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade5 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade6 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade7 Horizon VDI Non-Persistent VMs 167 Users
C2-Blade8 Horizon VDI Non-Persistent VMs 167 Users
Chassis 3
C3-Blade1 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade2 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade3 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade4 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade5 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade6 Horizon VDI Non-Persistent VMs 167 Users
C3-Blade7 Horizon VDI Non-Persistent VMs 168 Users
C3-Blade8 Infrastructure VMs
Chassis 4
C4-Blade1 Horizon VDI Non-Persistent VMs 168 Users
C4-Blade2 Horizon VDI Non-Persistent VMs 168 Users
C4-Blade3 Horizon VDI Non-Persistent VMs 167 Users
C4-Blade4 Horizon VDI Non-Persistent VMs 167 Users
C4-Blade5 Horizon VDI Non-Persistent VMs 168 Users
C4-Blade6 Horizon VDI Non-Persistent VMs 168 Users
C4-Blade7 Horizon VDI Non-Persistent VMs 168 Users
C4-Blade8 Infrastructure VMs

The configured system efficiently and effectively delivered the following results:

Figure 108. Full Scale | 5022 Users | Win10 Non-Persistent Desktop | VSI Score

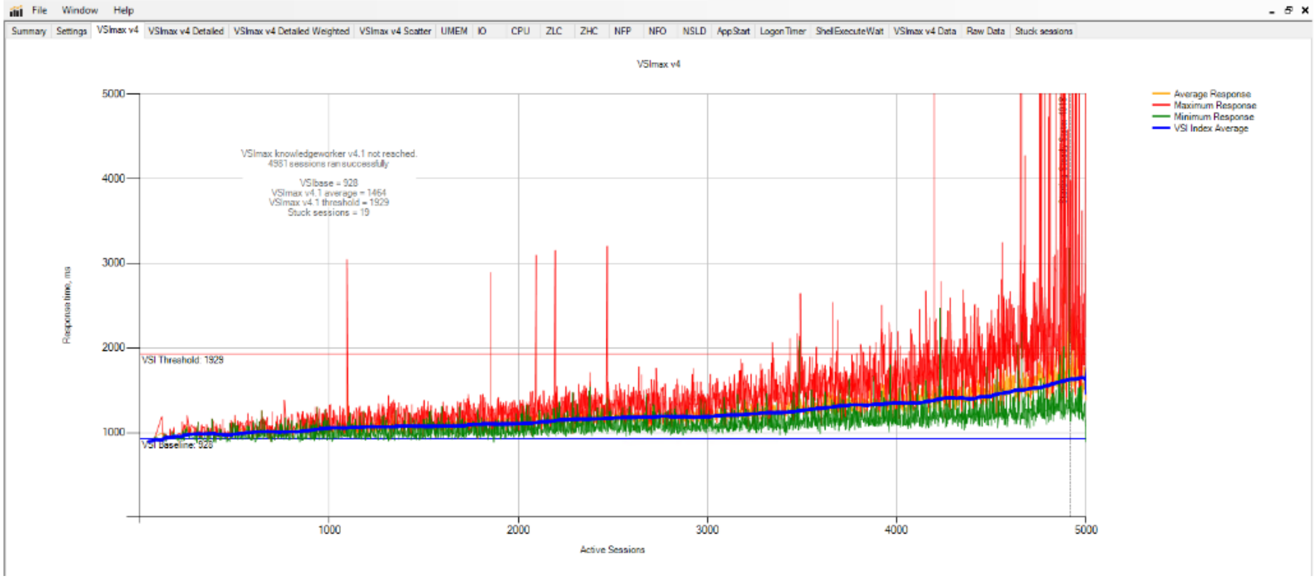

VSImax knowledgeworker v4.1 not reached.
4981 sessions ran successfully

VSIbase = 928
VSImax v4.1 average = 1464
VSImax v4.1 threshold = 1929
Stuck sessions = 19

Figure 109. Full Scale | 5022 Users | VDI Non-Persistent VM Hosts | Host CPU Utilization

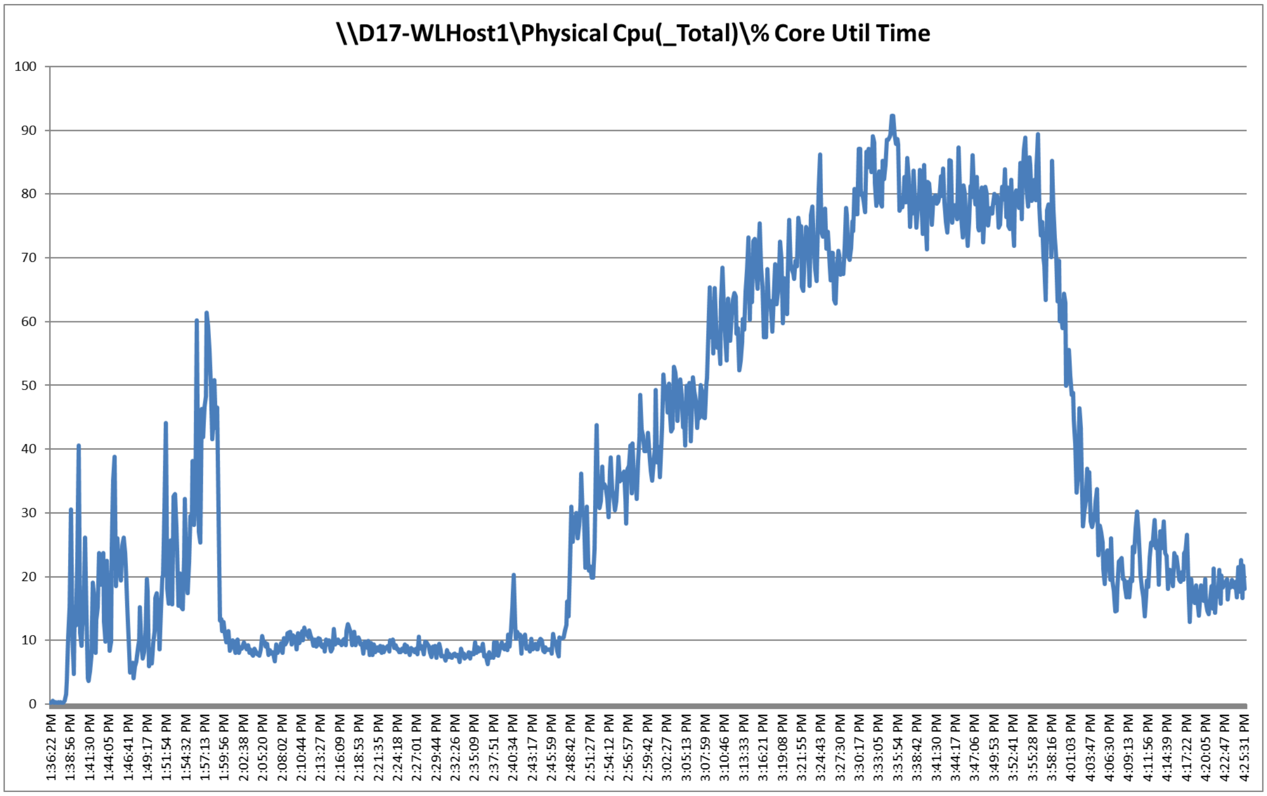

Figure 110. Full Scale | 5022 Users | VDI Non-Persistent Hosts | Host Memory Utilization

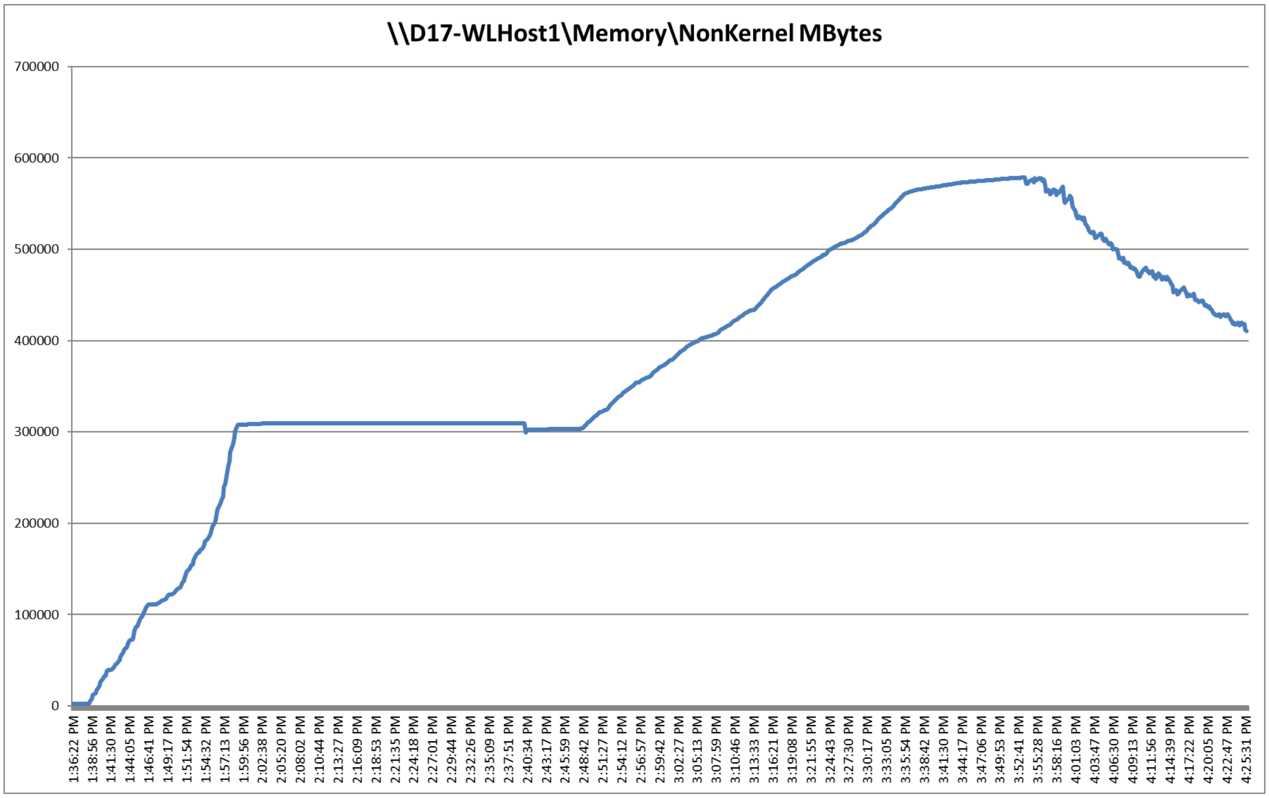

Figure 111. Full Scale | 5022 Users | VDI Non-Persistent Hosts | Host Network Utilization

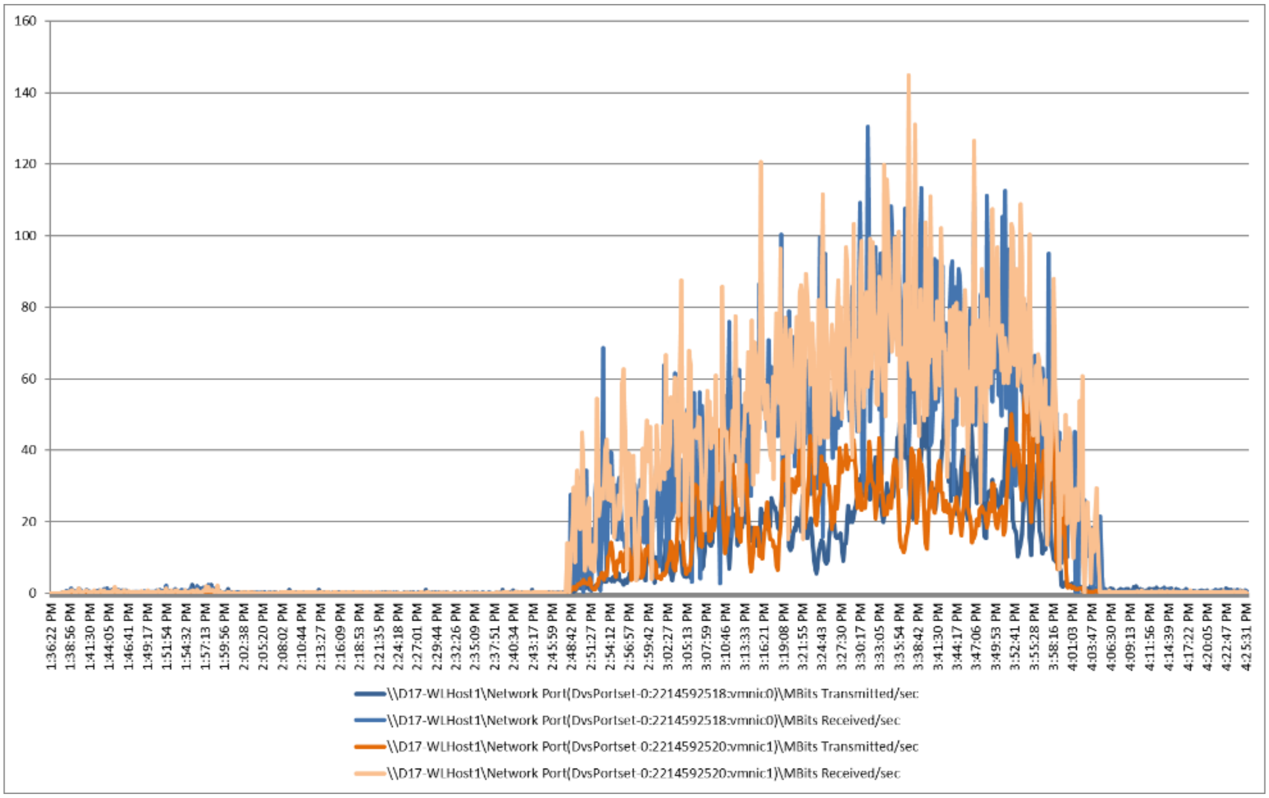

Pure Storage FlashArray//X70 R3 Storage System Graph for 5500 Non-Persistent Workload Test

Figure 112. Full Scale | 5022 Users | VDI Non-Persistent VM Hosts | Pure Storage FlashArray//X70 R3 System Latency Chart

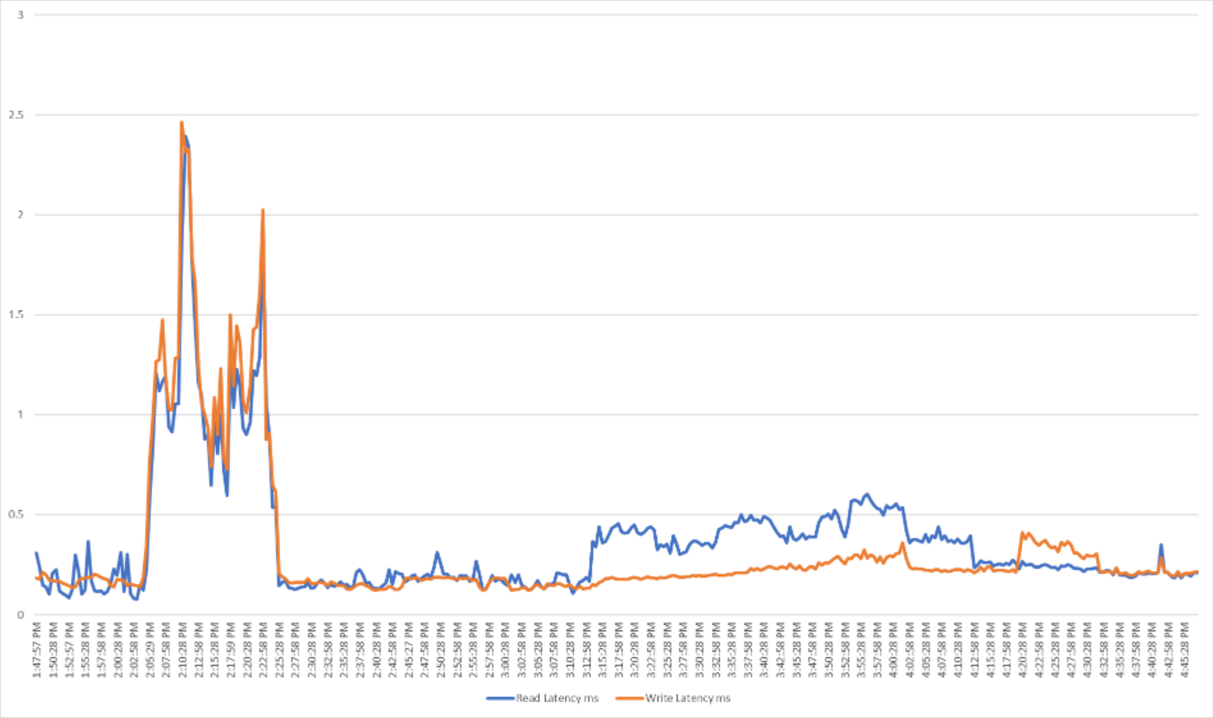

Figure 113. Full Scale | 5022 Users | VDI Non-Persistent VM Hosts | FlashArray//X70 R3 System IOPS Chart

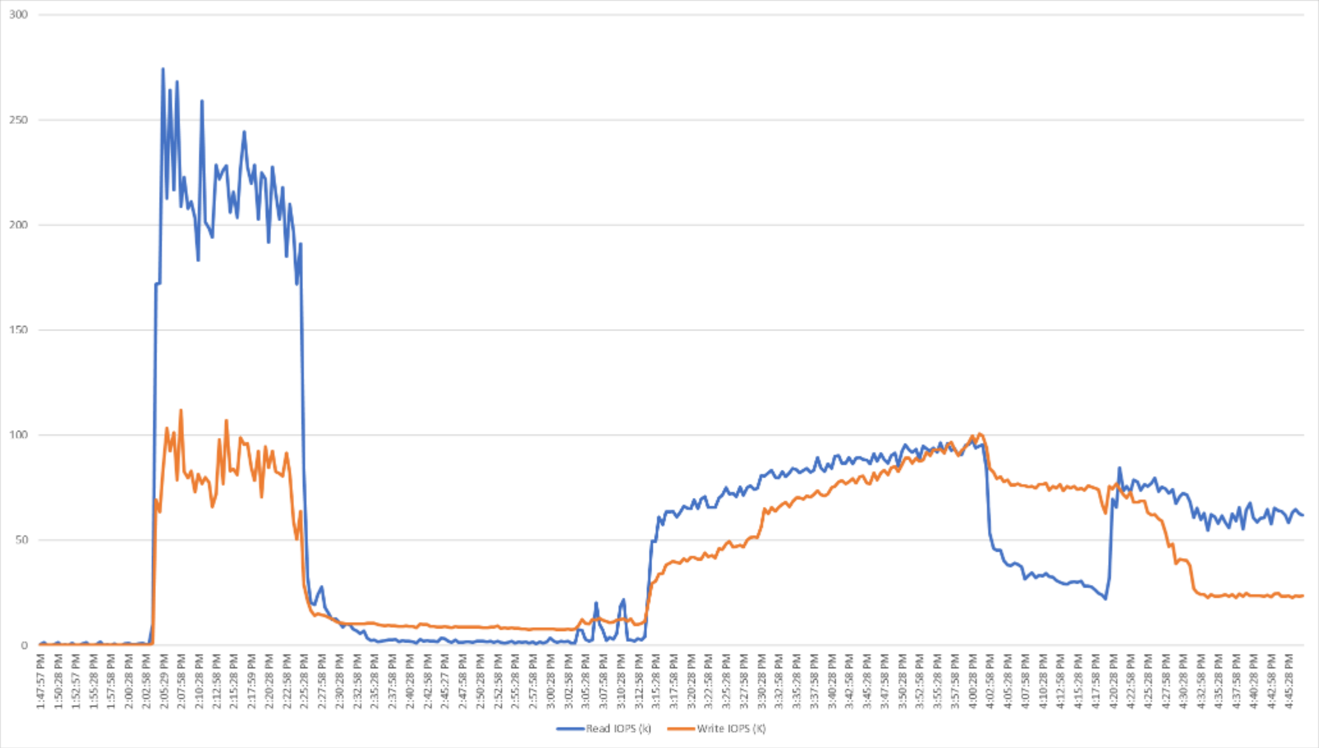

Figure 114. Full Scale | 5022 Users | VDI Non-Persistent VM Hosts | FlashArray//X70 R3 System Bandwidth Chart

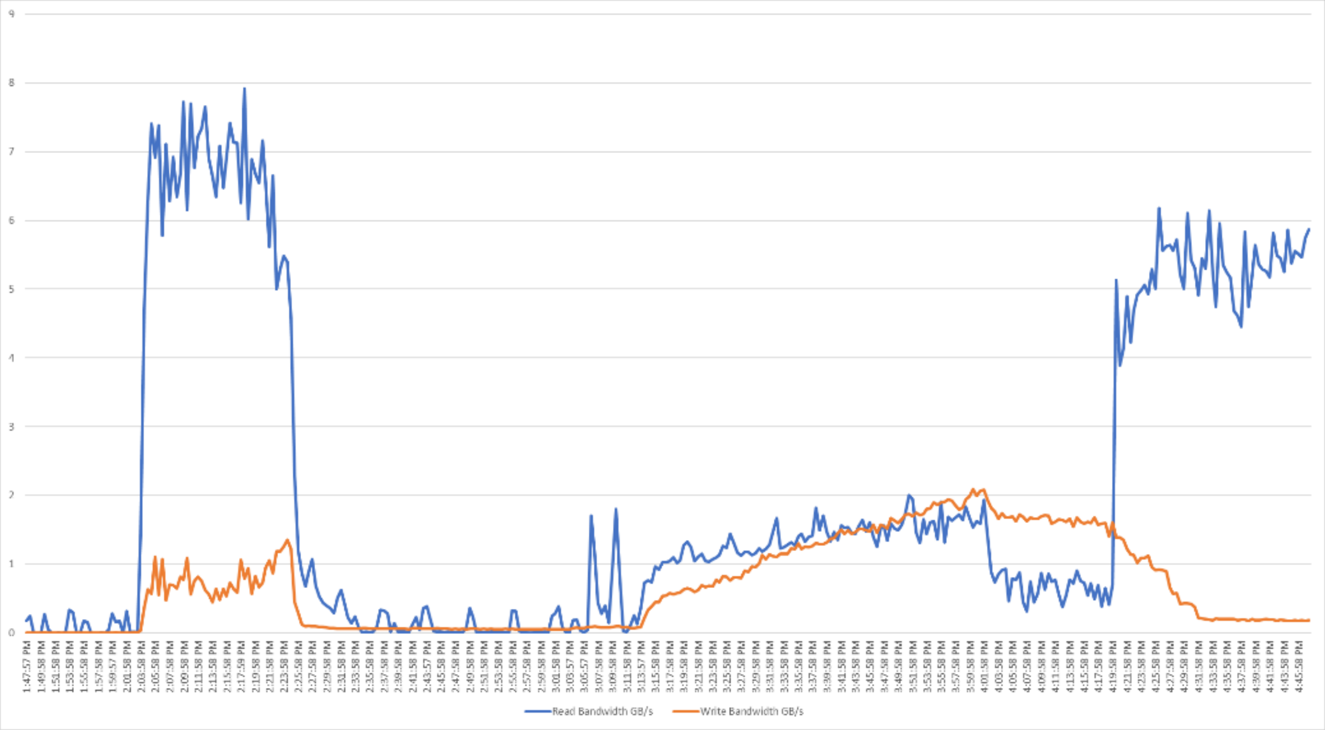

Figure 115. Full Scale | 5022 Users | VDI Non-Persistent VM Hosts | FlashArray//X70 R3 Performance Chart

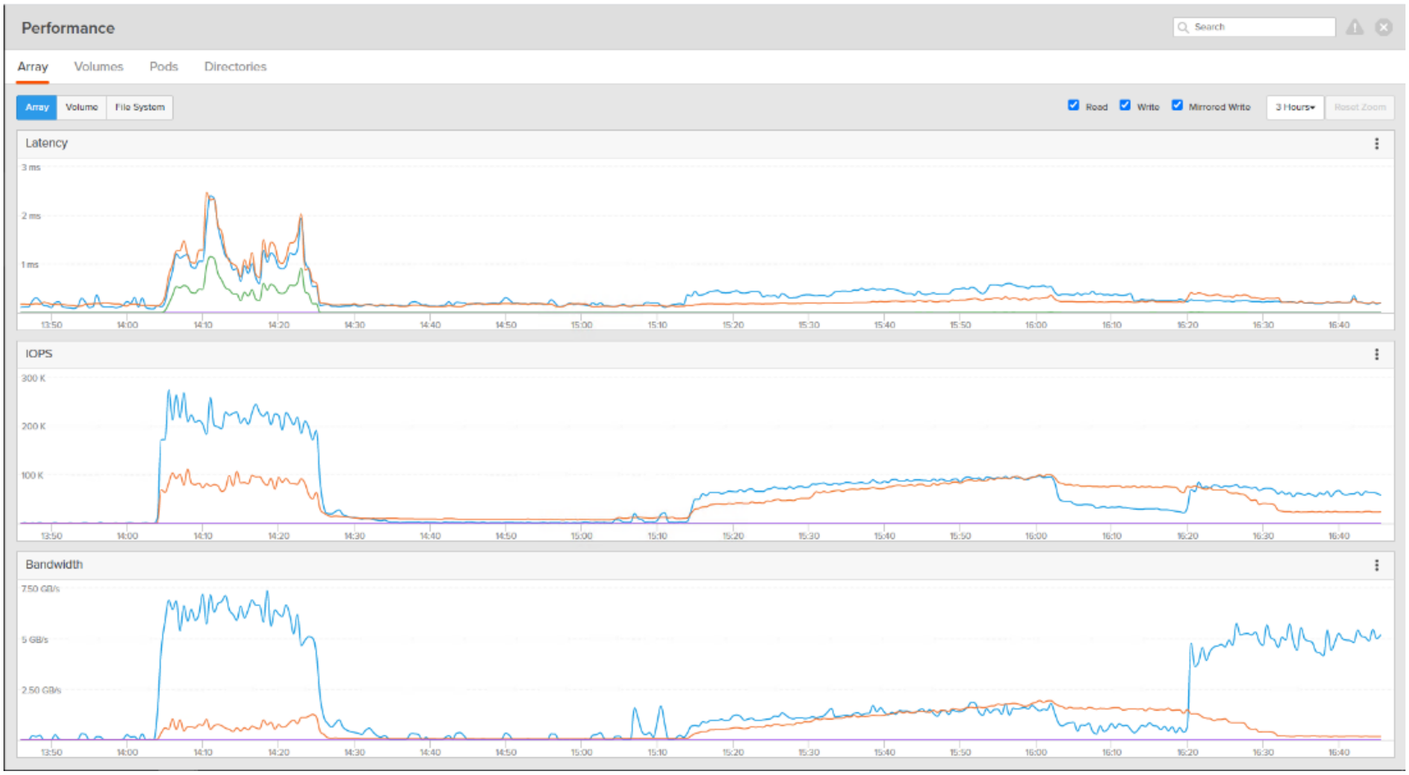

Full Scale Workload Testing with 5022 VDI Persistent Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray during the persistent desktop full-scale testing with 5022 VDI Persistent desktops using 30 blades (configured in three 10 Host Pools).

The workload for the test is 5022 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.

Figure 116. Full Scale VDI Persistent Test – Workload Distribution

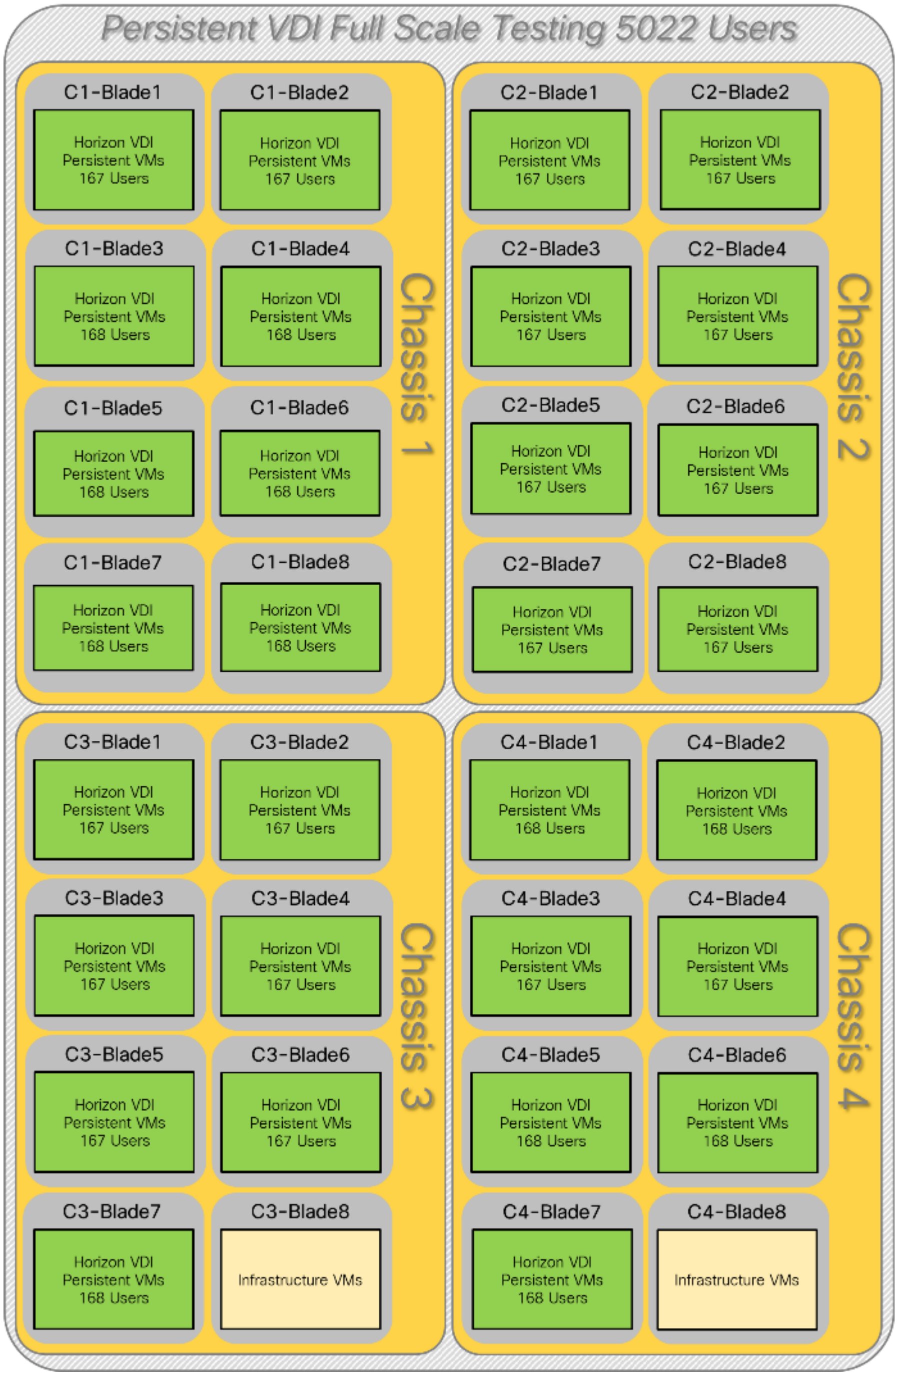
Persistent VDI Full Scale Testing 5022 Users
Chassis 1
C1-Blade1 Horizon VDI Persistent VMs 167 Users
C1-Blade2 Horizon VDI Persistent VMs 167 Users
C1-Blade3 Horizon VDI Persistent VMs 168 Users
C1-Blade4 Horizon VDI Persistent VMs 168 Users
C1-Blade5 Horizon VDI Persistent VMs 168 Users
C1-Blade6 Horizon VDI Persistent VMs 168 Users
C1-Blade7 Horizon VDI Persistent VMs 168 Users
C1-Blade8 Horizon VDI Persistent VMs 168 Users
Chassis 2
C2-Blade1 Horizon VDI Persistent VMs 167 Users
C2-Blade2 Horizon VDI Persistent VMs 167 Users
C2-Blade3 Horizon VDI Persistent VMs 167 Users
C2-Blade4 Horizon VDI Persistent VMs 167 Users
C2-Blade5 Horizon VDI Persistent VMs 167 Users
C2-Blade6 Horizon VDI Persistent VMs 167 Users
C2-Blade7 Horizon VDI Persistent VMs 167 Users
C2-Blade8 Horizon VDI Persistent VMs 167 Users
Chassis 3
C3-Blade1 Horizon VDI Persistent VMs 167 Users
C3-Blade2 Horizon VDI Persistent VMs 167 Users
C3-Blade3 Horizon VDI Persistent VMs 167 Users
C3-Blade4 Horizon VDI Persistent VMs 167 Users
C3-Blade5 Horizon VDI Persistent VMs 167 Users
C3-Blade6 Horizon VDI Persistent VMs 167 Users
C3-Blade7 Horizon VDI Persistent VMs 168 Users
C3-Blade8 Infrastructure VMs
Chassis 4
C4-Blade1 Horizon VDI Persistent VMs 168 Users
C4-Blade2 Horizon VDI Persistent VMs 168 Users
C4-Blade3 Horizon VDI Persistent VMs 167 Users
C4-Blade4 Horizon VDI Persistent VMs 167 Users
C4-Blade5 Horizon VDI Persistent VMs 168 Users
C4-Blade6 Horizon VDI Persistent VMs 168 Users
C4-Blade7 Horizon VDI Persistent VMs 168 Users
C4-Blade8 Infrastructure VMs

The configured system efficiently and effectively delivered the following results:

Figure 117. Full Scale | 5022 Users | Win10 Persistent Desktop | VSI Score

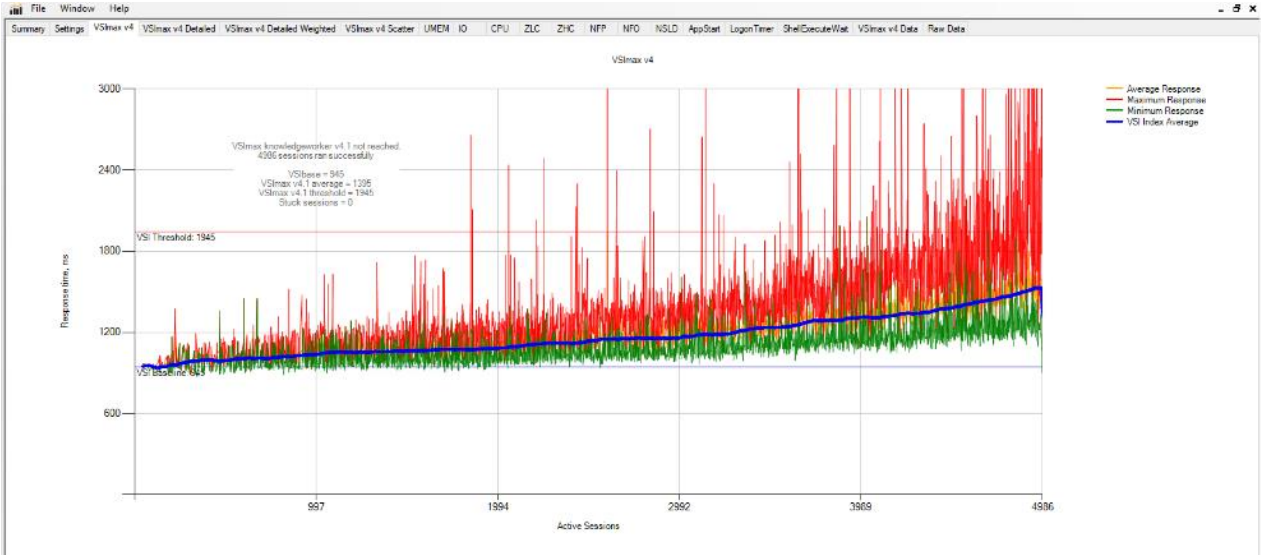

VSImax knowledgeworker v4.1 not reached.
4986 sessions ran successfully

VSIbase = 945
VSImax v4.1 average = 1395
VSImax v4.1 threshold = 1945
Stuck sessions = 0

Figure 118. Full Scale | 5022 Users | VDI Persistent VM Host | Host CPU Utilization

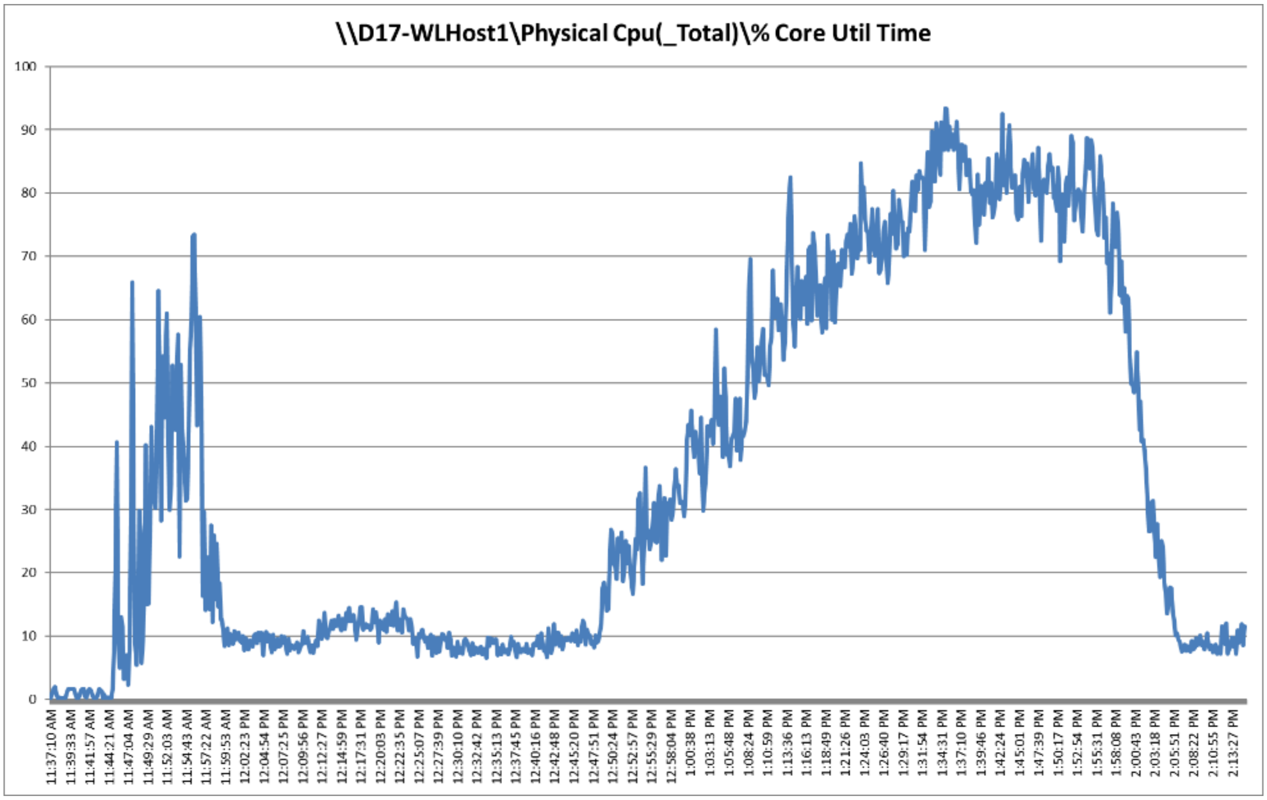

Figure 119. Full Scale | 5022 Users | VDI Persistent VM Hosts | Host Memory Utilization

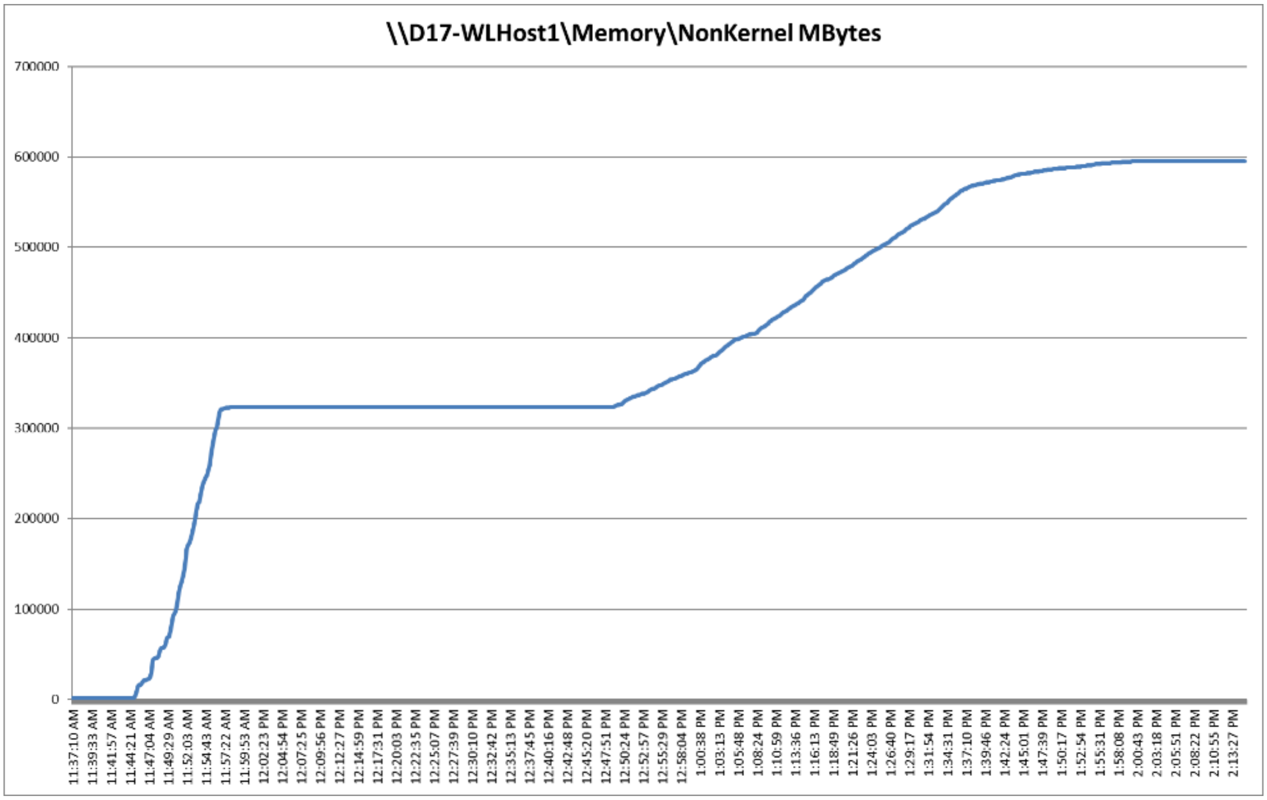

Figure 120. Full Scale | 5022 Users | VDI Persistent VM Hosts | Host Network Utilization

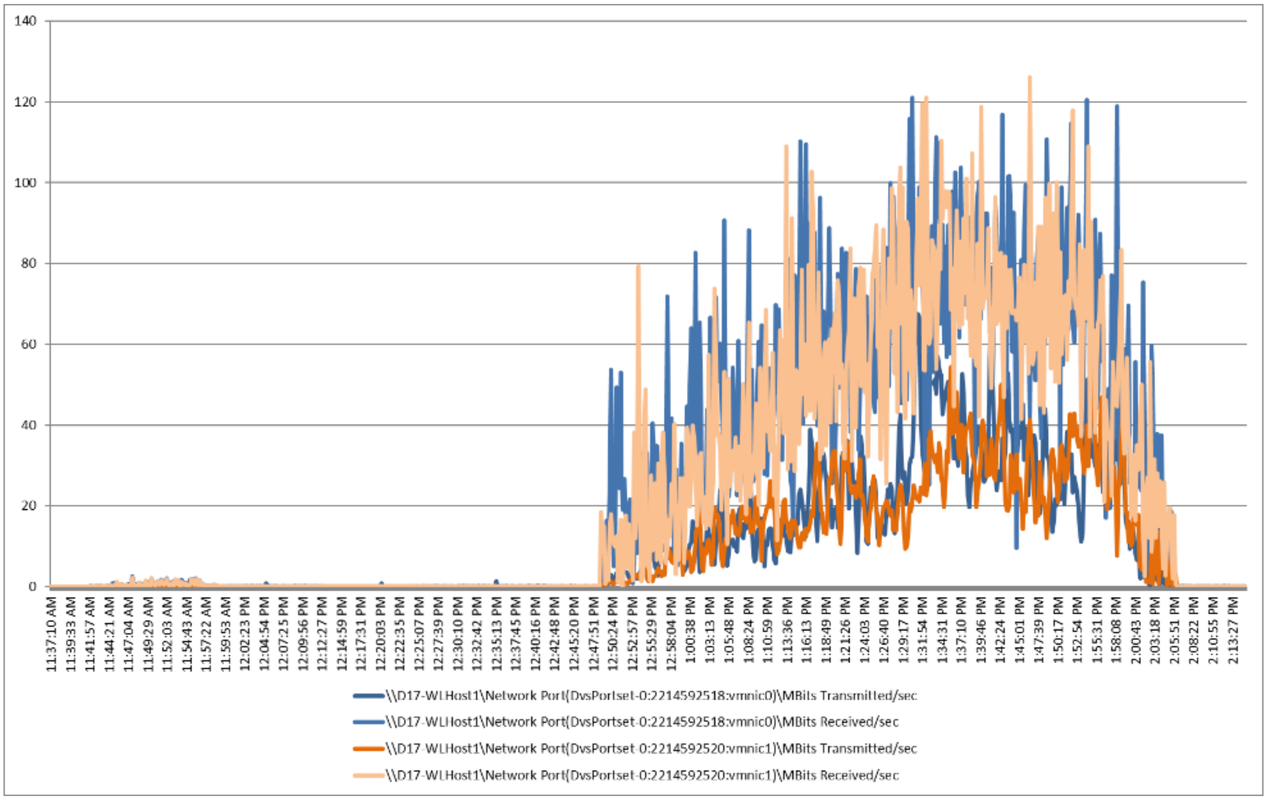

Pure Storage FlashArray//X70 R3 Storage System Graph for 5022 Users Persistent Desktop Workload Test

Figure 121. Full Scale | 5022 Users | VDI-Persistent VM Hosts | FlashArray//X70 R3 System Latency Chart

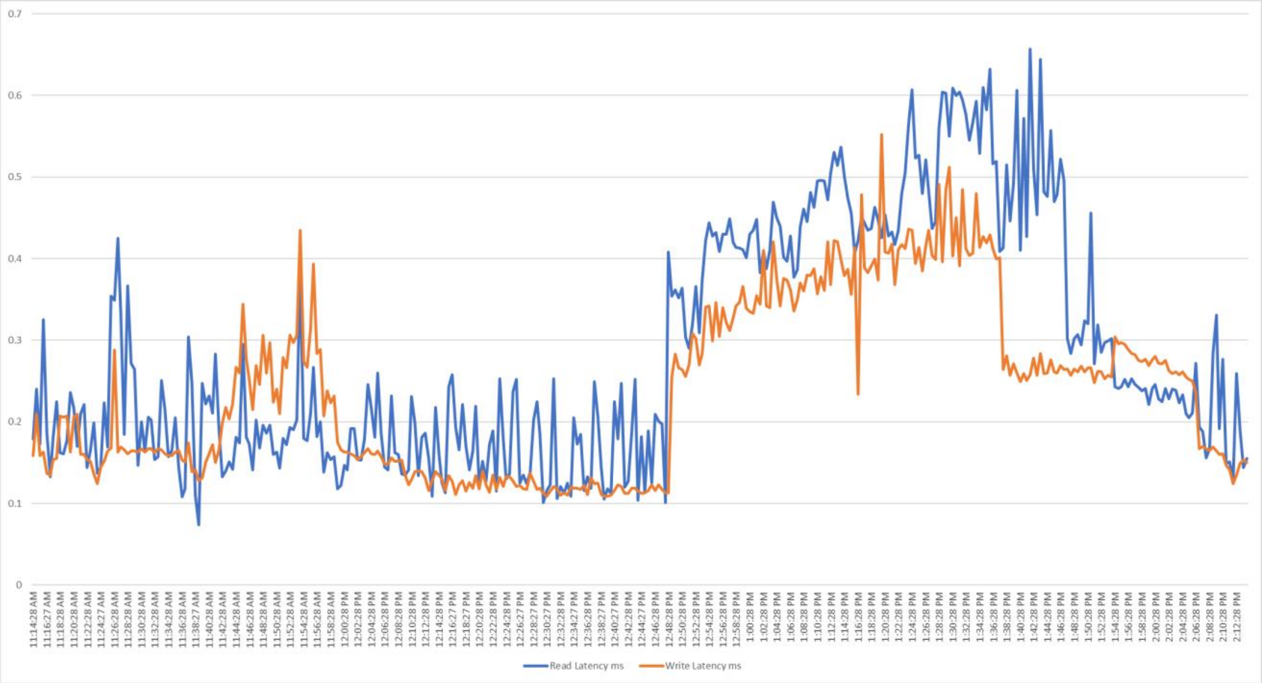

Figure 122. Full Scale | 5022 Users | VDI-Persistent VM Hosts | FlashArray//X70 R3 System IOPS Chart

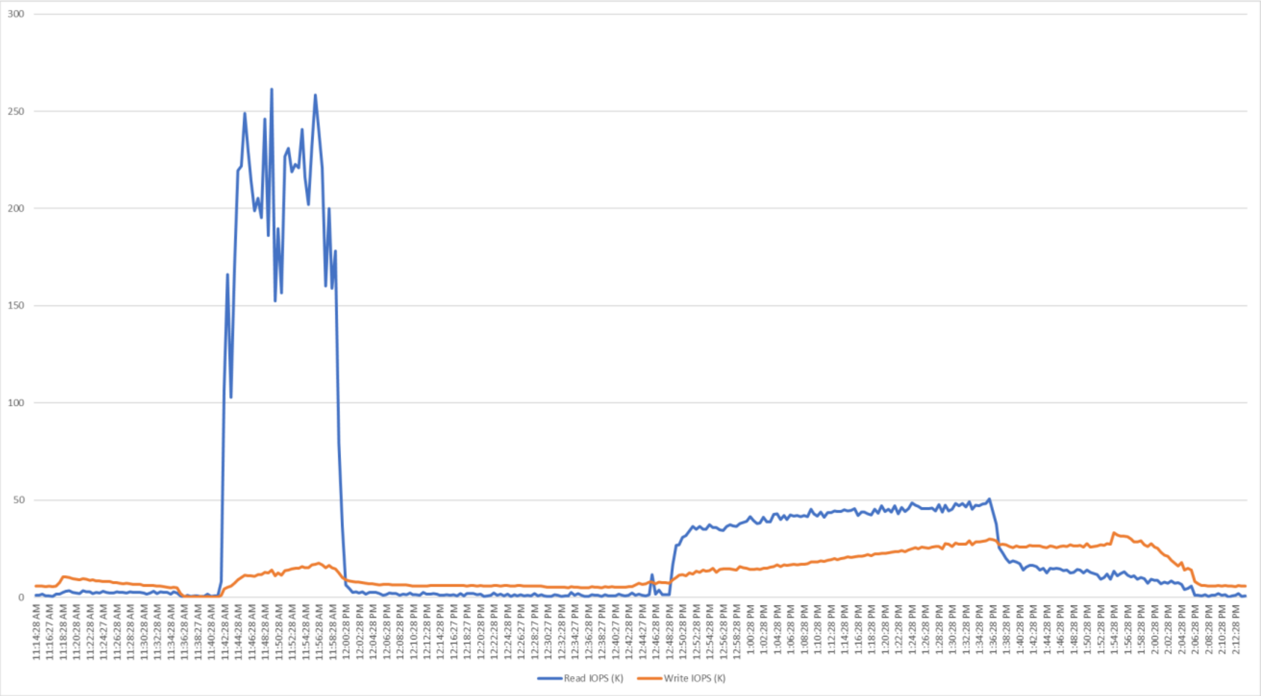

Figure 123. Full Scale | 5022 Users | VDI-Persistent VM Hosts | FlashArray//X70 R3 System Bandwidth Chart

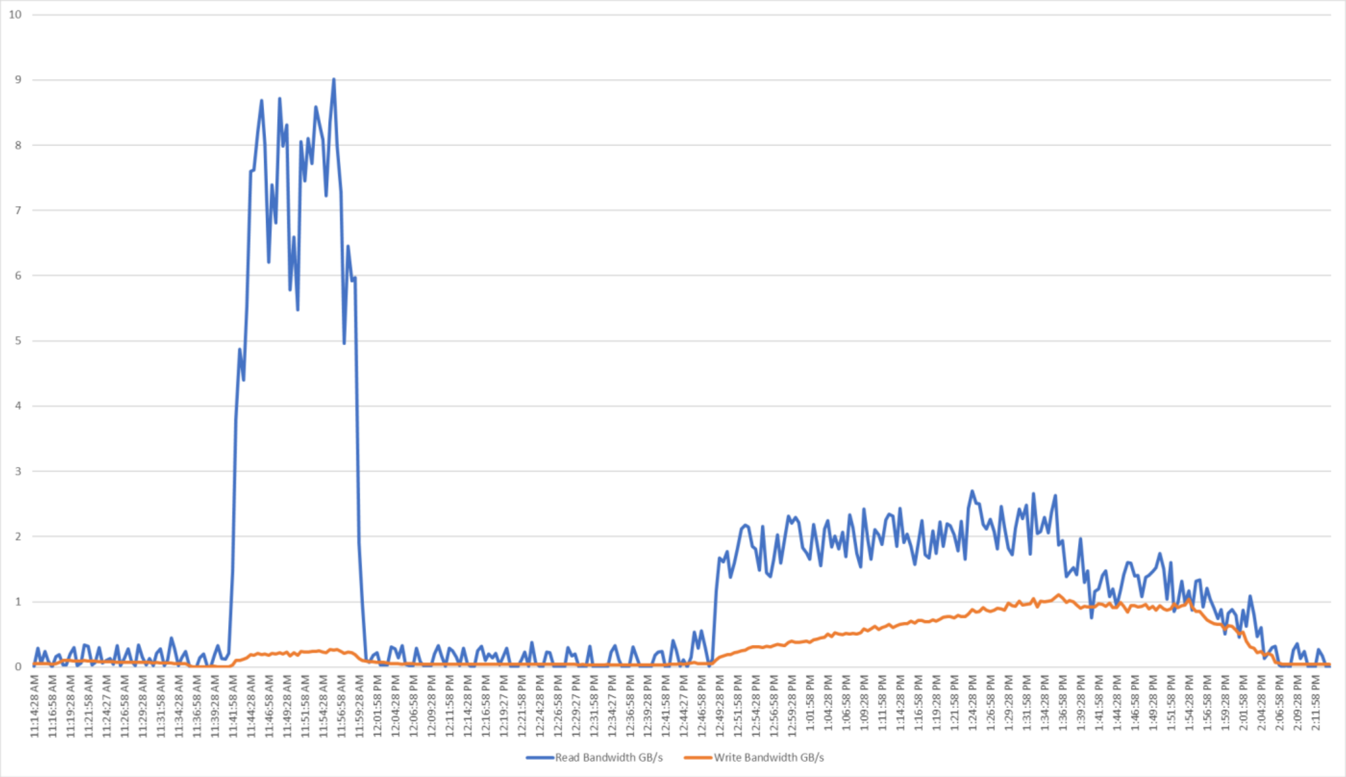

Figure 124. Full Scale | 5022 Users | VDI-Persistent VM Hosts | FlashArray//X70 R3 System Performance Chart

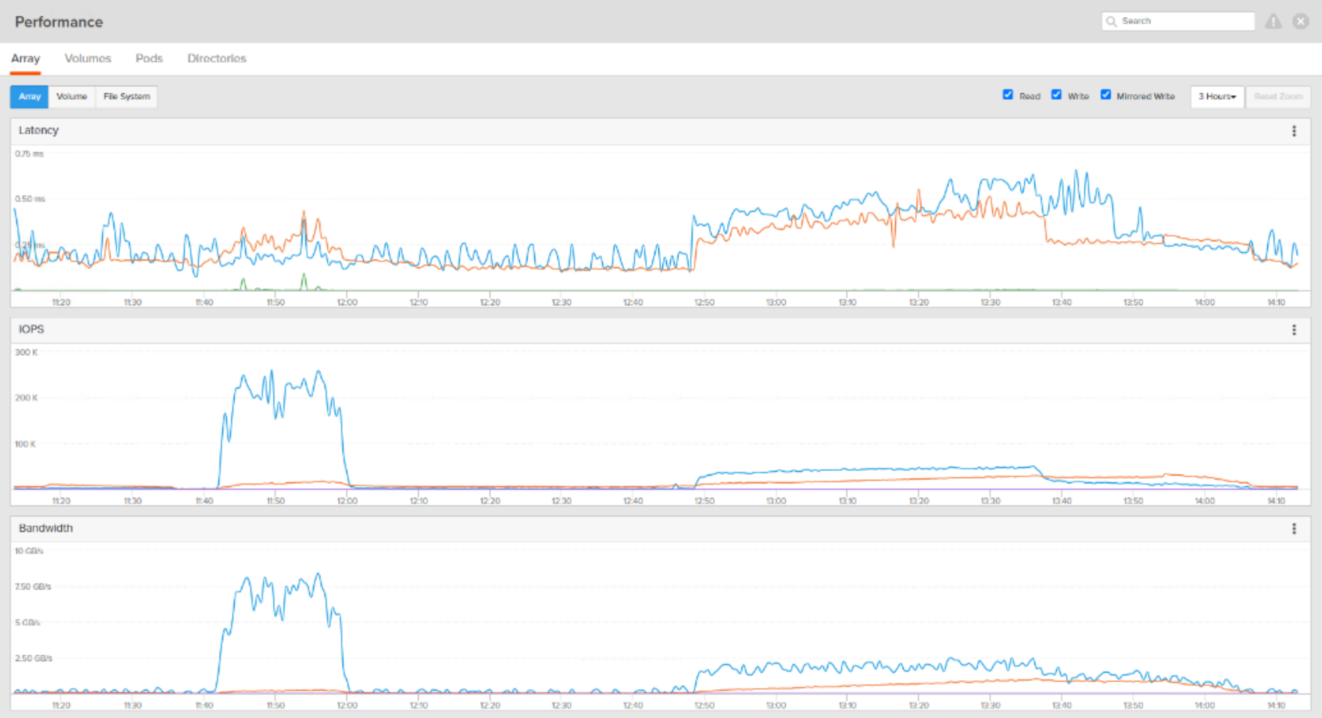

☑ Read ☑ Write ☑ Mirrored Write 3 Hours▾

Summary

FlashStack 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 Pure Storage FlashArray//X70 R3 Storage Array. FlashStack is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlashStack provides to customers wishing to deploy enterprise-class VDI and RDS.

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, Pure Storage FlashArray//X70 R3 storage and our certified partners can help. We collaborate with you to enhance your IT capabilities through a full portfolio of services for 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.

Additionally, Cisco Advanced Services and Pure Storage Support provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Author

Vadim Lebedev, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Vadim Lebedev for the last 5 years 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 Desktop/Server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA Graphics.

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:

- Keith Hageman, Technical Marketing Performance Engineer, Pure Storage, Inc.
- Craig Waters, Technical Director Solutions Architecture/Product Management, Pure Storage, Inc.

References

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

Cisco UCS B-Series Servers

- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html
- https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b200m5-specsheet.pdf
- https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/datasheet-listing.html
- https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-b200-m5-blade-server/model.html
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/blade-servers/B200M5.pdf

Cisco UCS Manager Configuration Guides

- http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
- https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html

Cisco UCS Virtual Interface Cards

- https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.html

Cisco Nexus Switching References

- http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html
- https://www.cisco.com/c/en/us/products/switches/nexus-93180yc-fx-switch/index.html

Cisco MDS 9000 Service Switch References

- http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html
- http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html
- https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9132T_32-Gb-16g-multilayer-fabric-switch/datasheet-c78-731523.html

VMware References

- https://docs.vmware.com/en/VMware-vSphere/index.html

- https://docs.vmware.com/en/VMware-Horizon/index.html

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

Pure Storage Reference Documents

- https://www.flashstack.com/
- https://www.purestorage.com/content/dam/purestorage/pdf/datasheets/ps_ds_flasharray_03.pdf
- https://www.purestorage.com
- https://www.purestorage.com/products/evergreen-subscriptions.html
- https://www.purestorage.com/solutions/infrastructure/vdi.html
- https://www.purestorage.com/solutions/infrastructure/vdi-calculator.html

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 AAD17-NX9K-A
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-all-flood
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-fcoe
    mtu 2158
  class type network-qos class-default
    mtu 9216
install feature-set fcoe-npv
vdc AAD17-NX9K-A id 1
  allow feature-set fcoe-npv
  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-set fcoe-npv

feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp

no password strength-check
username admin password 5 $5$d3vc8gvD$hmf.YoRRPcqZ2dDGV2IaVKYZsPSPls8E9bpUzMciMZ0  role network-admin
ip domain-lookup
system default switchport
class-map type qos match-all class-fcoe
policy-map type qos jumbo
  class class-default
    set qos-group 0
system qos
  service-policy type network-qos jumbo
copp profile lenient
snmp-server user admin network-admin auth md5 0xc9a73d344387b8db2dc0f3fc624240ac priv 0xc9a73d344387b8db2dc0f3fc624240ac localizedkey
snmp-server host 10.24.66.169 traps version 2c public udp-port 1165
snmp-server host 10.24.72.119 traps version 2c public udp-port 1163
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 10.10.70.2 use-vrf default
ntp peer 10.10.70.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8

vlan 1,70-76
vlan 70
  name InBand-Mgmt-SP
vlan 71
  name Infra-Mgmt-SP
vlan 72
  name VM-Network-SP
vlan 73
  name vMotion-SP
vlan 74
  name Storage_A-SP
```

```
vlan 75
  name Storage_B-SP
vlan 76
  name Launcher-SP

service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-redirect 256
vpc domain 70
  role priority 1000
  peer-keepalive destination 10.29.164.234 source 10.29.164.233


interface Vlan1
  no shutdown
  ip address 10.29.164.241/24

interface Vlan70
  no shutdown
  ip address 10.10.70.2/24
  hsrp version 2
  hsrp 70
    preempt
    priority 110
    ip 10.10.70.1

interface Vlan71
  no shutdown
  ip address 10.10.71.2/24
  hsrp version 2
  hsrp 71
    preempt
    priority 110
    ip 10.10.71.1
```

```
interface Vlan72
  no shutdown
  ip address 10.72.0.2/19
  hsrp version 2
  hsrp 72
    preempt
    priority 110
    ip 10.72.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan73
  no shutdown
  ip address 10.10.73.2/24
  hsrp version 2
  hsrp 73
    preempt
    priority 110
    ip 10.10.73.1

interface Vlan74
  no shutdown
  ip address 10.10.74.2/24
  hsrp version 2
  hsrp 74
    preempt
    priority 110
    ip 10.10.74.1

interface Vlan75
  no shutdown
  ip address 10.10.75.2/24
  hsrp version 2
  hsrp 75
    preempt
    priority 110
    ip 10.10.75.1
```

```
interface Vlan76
  no shutdown
  ip address 10.10.76.2/23
  hsrp version 2
  hsrp 76
    preempt
    priority 110
    ip 10.10.76.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface port-channel10

interface port-channel11
  description FI-Uplink-D17
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 11

interface port-channel12
  description FI-Uplink-D17
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 12

interface port-channel13
  description FI-Uplink-D16
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 13
```

```
interface port-channel14
  description FI-Uplink-D16
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 14

interface port-channel70
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type network
  service-policy type qos input jumbo
  vpc peer-link

interface port-channel101
  description to PureStorage ethernet port eth2
  shutdown
  switchport access vlan 72
  spanning-tree port type edge
  mtu 9216
  service-policy type qos input jumbo
  vpc 101

interface Ethernet1/1

interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76

interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 13 mode active
```

```
interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 13 mode active

interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 14 mode active

interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 14 mode active

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
  switchport access vlan 71
  spanning-tree port type edge

interface Ethernet1/34
  switchport access vlan 71
  spanning-tree port type edge
```

```
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

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/52
```

```
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active

interface Ethernet1/54
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.233/24
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 AAD17-NX9K-B
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-all-flood
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-fcoe
    mtu 2158
  class type network-qos class-default
    mtu 9216
install feature-set fcoe-npv
vdc AAD17-NX9K-B id 1
  allow feature-set fcoe-npv
  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-set fcoe-npv

feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp

no password strength-check
username admin password 5 $5$/48.OHa8$g6pOMLIwrzqxJesMYoP5CNphujBksPPRjn4I3iFfOp.  role network-admin
ip domain-lookup
system default switchport
class-map type qos match-all class-fcoe
policy-map type qos jumbo
  class class-default
    set qos-group 0
system qos
  service-policy type network-qos jumbo
copp profile lenient
snmp-server user admin network-admin auth md5 0x6d450e3d5a3927ddee1dadd30e5f616f priv 0x6d450e3d5a3927ddee1dadd30e5f616f localizedkey
snmp-server host 10.24.66.169 traps version 2c public udp-port 1166
snmp-server host 10.24.72.119 traps version 2c public udp-port 1164
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 peer 10.10.70.2 use-vrf default
ntp server 10.10.70.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8

vlan 1,70-76
vlan 70
  name InBand-Mgmt-SP
vlan 71
  name Infra-Mgmt-SP
vlan 72
  name VM-Network-SP
vlan 73
  name vMotion-SP
vlan 74
  name Storage_A-SP
vlan 75
  name Storage_B-SP
vlan 76
  name Launcher-SP

service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-redirect 256
vpc domain 70
  role priority 2000
  peer-keepalive destination 10.29.164.233 source 10.29.164.234


interface Vlan1
  no shutdown
  ip address 10.29.164.240/24

interface Vlan70
```

```
  no shutdown
  ip address 10.10.70.3/24
  hsrp version 2
  hsrp 70
    preempt
    priority 110
    ip 10.10.70.1

interface Vlan71
  no shutdown
  ip address 10.10.71.3/24
  hsrp version 2
  hsrp 71
    preempt
    priority 110
    ip 10.10.71.1

interface Vlan72
  no shutdown
  ip address 10.72.0.2/19
  hsrp version 2
  hsrp 72
    preempt
    priority 110
    ip 10.72.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan73
  no shutdown
  ip address 10.10.73.3/24
  hsrp version 2
  hsrp 73
    preempt
    priority 110
    ip 10.10.73.1

interface Vlan74
  no shutdown
  ip address 10.10.74.3/24
  hsrp version 2
  hsrp 74
    preempt
    priority 110
    ip 10.10.74.1

interface Vlan75
  no shutdown
  ip address 10.10.75.3/24
  hsrp version 2
  hsrp 75
    preempt
    priority 110
    ip 10.10.75.1

interface Vlan76
  no shutdown
  ip address 10.10.76.3/23
  hsrp version 2
  hsrp 76
    preempt
    priority 110
    ip 10.10.76.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface port-channel10

interface port-channel11
  description FI-Uplink-D17
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 11

interface port-channel12
```

```
  description FI-Uplink-D17
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 12

interface port-channel13
  description FI-Uplink-D16
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 13

interface port-channel14
  description FI-Uplink-D16
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 14

interface port-channel70
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type network
  service-policy type qos input jumbo
  vpc peer-link

interface port-channel101
  description to PureStorage ethernet port eth2
  shutdown
  switchport access vlan 72
  mtu 9216
  service-policy type qos input jumbo
  vpc 101

interface Ethernet1/1
  switchport access vlan 70
  speed 1000

interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76

interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 13 mode active

interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 13 mode active

interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 14 mode active

interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 14 mode active

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
  switchport access vlan 71
  spanning-tree port type edge

interface Ethernet1/34
  switchport access vlan 71
  spanning-tree port type edge

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
interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/52
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active

interface Ethernet1/54
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.234/24
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 $5$Dcs72Ao/$8lHyVrotTm4skqb/84BC793tgdly/yWf9IoMx2OEg6C  role network-admin
ip domain-lookup
ip name-server 10.10.61.30
ip host ADD16-MDS-A  10.29.164.238
aaa group server radius radius
snmp-server user admin network-admin auth md5 0x616758aed4f07bab2d24f3d594ebd649 priv 0x616758aed4f07bab2d24f3d594ebd649 localizedkey
snmp-server host 10.24.30.91 traps version 2c public udp-port 1163
snmp-server host 10.24.46.67 traps version 2c public udp-port 1163
snmp-server host 10.24.66.169 traps version 2c public udp-port 1163
snmp-server host 10.24.72.119 traps version 2c public udp-port 1165
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 100 name "FlashStack-VCC-CVD-Fabric-A"
device-alias database
  device-alias name X70R3-CT0-FC0 pwwn 52:4a:93:71:56:84:09:00
```

```
device-alias name X70R3-CT1-FC0 pwwn 52:4a:93:71:56:84:09:10
device-alias name VCC-Infra01-HBA0 pwwn 20:00:00:25:b5:aa:17:1e
device-alias name VCC-Infra01-HBA2 pwwn 20:00:00:25:b5:aa:17:1f
device-alias name VCC-Infra02-HBA0 pwwn 20:00:00:25:b5:aa:17:3e
device-alias name VCC-Infra02-HBA2 pwwn 20:00:00:25:b5:aa:17:3f
device-alias name VCC-WLHost01-HBA0 pwwn 20:00:00:25:b5:aa:17:00
device-alias name VCC-WLHost01-HBA2 pwwn 20:00:00:25:b5:aa:17:01
device-alias name VCC-WLHost02-HBA0 pwwn 20:00:00:25:b5:aa:17:02
device-alias name VCC-WLHost02-HBA2 pwwn 20:00:00:25:b5:aa:17:03
device-alias name VCC-WLHost03-HBA0 pwwn 20:00:00:25:b5:aa:17:04
device-alias name VCC-WLHost03-HBA2 pwwn 20:00:00:25:b5:aa:17:05
device-alias name VCC-WLHost04-HBA0 pwwn 20:00:00:25:b5:aa:17:06
device-alias name VCC-WLHost04-HBA2 pwwn 20:00:00:25:b5:aa:17:07
device-alias name VCC-WLHost05-HBA0 pwwn 20:00:00:25:b5:aa:17:08
device-alias name VCC-WLHost05-HBA2 pwwn 20:00:00:25:b5:aa:17:09
device-alias name VCC-WLHost06-HBA0 pwwn 20:00:00:25:b5:aa:17:0a
device-alias name VCC-WLHost06-HBA2 pwwn 20:00:00:25:b5:aa:17:0b
device-alias name VCC-WLHost07-HBA0 pwwn 20:00:00:25:b5:aa:17:0c
device-alias name VCC-WLHost07-HBA2 pwwn 20:00:00:25:b5:aa:17:0d
device-alias name VCC-WLHost08-HBA0 pwwn 20:00:00:25:b5:aa:17:0e
device-alias name VCC-WLHost08-HBA2 pwwn 20:00:00:25:b5:aa:17:0f
device-alias name VCC-WLHost09-HBA0 pwwn 20:00:00:25:b5:aa:17:10
device-alias name VCC-WLHost09-HBA2 pwwn 20:00:00:25:b5:aa:17:11
device-alias name VCC-WLHost10-HBA0 pwwn 20:00:00:25:b5:aa:17:12
device-alias name VCC-WLHost10-HBA2 pwwn 20:00:00:25:b5:aa:17:13
device-alias name VCC-WLHost11-HBA0 pwwn 20:00:00:25:b5:aa:17:14
device-alias name VCC-WLHost11-HBA2 pwwn 20:00:00:25:b5:aa:17:15
device-alias name VCC-WLHost12-HBA0 pwwn 20:00:00:25:b5:aa:17:16
device-alias name VCC-WLHost12-HBA2 pwwn 20:00:00:25:b5:aa:17:17
device-alias name VCC-WLHost13-HBA0 pwwn 20:00:00:25:b5:aa:17:18
device-alias name VCC-WLHost13-HBA2 pwwn 20:00:00:25:b5:aa:17:19
device-alias name VCC-WLHost14-HBA0 pwwn 20:00:00:25:b5:aa:17:1a
device-alias name VCC-WLHost14-HBA2 pwwn 20:00:00:25:b5:aa:17:1b
device-alias name VCC-WLHost15-HBA0 pwwn 20:00:00:25:b5:aa:17:1c
device-alias name VCC-WLHost15-HBA2 pwwn 20:00:00:25:b5:aa:17:1d
device-alias name VCC-WLHost16-HBA0 pwwn 20:00:00:25:b5:aa:17:20
device-alias name VCC-WLHost16-HBA2 pwwn 20:00:00:25:b5:aa:17:21
device-alias name VCC-WLHost17-HBA0 pwwn 20:00:00:25:b5:aa:17:22
device-alias name VCC-WLHost17-HBA2 pwwn 20:00:00:25:b5:aa:17:23
```

```
  device-alias name VCC-WLHost18-HBA0 pwwn 20:00:00:25:b5:aa:17:24
  device-alias name VCC-WLHost18-HBA2 pwwn 20:00:00:25:b5:aa:17:25
  device-alias name VCC-WLHost19-HBA0 pwwn 20:00:00:25:b5:aa:17:26
  device-alias name VCC-WLHost19-HBA2 pwwn 20:00:00:25:b5:aa:17:27
  device-alias name VCC-WLHost20-HBA0 pwwn 20:00:00:25:b5:aa:17:28
  device-alias name VCC-WLHost20-HBA2 pwwn 20:00:00:25:b5:aa:17:29
  device-alias name VCC-WLHost21-HBA0 pwwn 20:00:00:25:b5:aa:17:2a
  device-alias name VCC-WLHost21-HBA2 pwwn 20:00:00:25:b5:aa:17:2b
  device-alias name VCC-WLHost22-HBA0 pwwn 20:00:00:25:b5:aa:17:2c
  device-alias name VCC-WLHost22-HBA2 pwwn 20:00:00:25:b5:aa:17:2d
  device-alias name VCC-WLHost23-HBA0 pwwn 20:00:00:25:b5:aa:17:2e
  device-alias name VCC-WLHost23-HBA2 pwwn 20:00:00:25:b5:aa:17:2f
  device-alias name VCC-WLHost24-HBA0 pwwn 20:00:00:25:b5:aa:17:30
  device-alias name VCC-WLHost24-HBA2 pwwn 20:00:00:25:b5:aa:17:31
  device-alias name VCC-WLHost25-HBA0 pwwn 20:00:00:25:b5:aa:17:32
  device-alias name VCC-WLHost25-HBA2 pwwn 20:00:00:25:b5:aa:17:33
  device-alias name VCC-WLHost26-HBA0 pwwn 20:00:00:25:b5:aa:17:34
  device-alias name VCC-WLHost26-HBA2 pwwn 20:00:00:25:b5:aa:17:35
  device-alias name VCC-WLHost27-HBA0 pwwn 20:00:00:25:b5:aa:17:36
  device-alias name VCC-WLHost27-HBA2 pwwn 20:00:00:25:b5:aa:17:37
  device-alias name VCC-WLHost28-HBA0 pwwn 20:00:00:25:b5:aa:17:38
  device-alias name VCC-WLHost28-HBA2 pwwn 20:00:00:25:b5:aa:17:39
  device-alias name VCC-WLHost29-HBA0 pwwn 20:00:00:25:b5:aa:17:3a
  device-alias name VCC-WLHost29-HBA2 pwwn 20:00:00:25:b5:aa:17:3b
  device-alias name VCC-WLHost30-HBA0 pwwn 20:00:00:25:b5:aa:17:3c
  device-alias name VCC-WLHost30-HBA2 pwwn 20:00:00:25:b5:aa:17:3d

device-alias commit

fcdomain fcid database
  vsan 100 wwn 20:03:00:de:fb:92:8d:00 fcid 0x300000 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:02 fcid 0x300020 dynamic
    !          [X70-CT0-FC2]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:17 fcid 0x300040 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:06 fcid 0x300041 dynamic
    !          [X70-CT0-FC8]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:07 fcid 0x300042 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:16 fcid 0x300043 dynamic
    !          [X70-CT1-FC8]
```

```
vsan 100 wwn 20:00:00:25:b5:aa:17:3e fcid 0x300060 dynamic
  !          [VCC-Infra02-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:07 fcid 0x300061 dynamic
  !          [VCC-WLHost04-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:06 fcid 0x300062 dynamic
  !          [VCC-WLHost04-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:3a fcid 0x300063 dynamic
  !          [VCC-WLHost29-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:29 fcid 0x300064 dynamic
  !          [VCC-WLHost20-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:13 fcid 0x300065 dynamic
  !          [VCC-WLHost10-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:1c fcid 0x300066 dynamic
  !          [VCC-WLHost15-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:32 fcid 0x300067 dynamic
  !          [VCC-WLHost25-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:17 fcid 0x300068 dynamic
  !          [VCC-WLHost12-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:2e fcid 0x300069 dynamic
  !          [VCC-WLHost23-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:1f fcid 0x30006a dynamic
  !          [VCC-Infra01-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:1b fcid 0x30006b dynamic
  !          [VCC-WLHost14-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:1a fcid 0x30006c dynamic
  !          [VCC-WLHost14-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:0a fcid 0x30006d dynamic
  !          [VCC-WLHost06-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:34 fcid 0x30006e dynamic
  !          [VCC-WLHost26-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:19 fcid 0x30006f dynamic
  !          [VCC-WLHost13-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:36 fcid 0x300070 dynamic
  !          [VCC-WLHost27-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:01 fcid 0x300071 dynamic
  !          [VCC-WLHost01-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:12 fcid 0x300072 dynamic
  !          [VCC-WLHost10-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:16 fcid 0x300073 dynamic
```

```
  !          [VCC-WLHost12-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:2b fcid 0x300074 dynamic
  !          [VCC-WLHost21-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:25 fcid 0x300075 dynamic
  !          [VCC-WLHost18-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:27 fcid 0x300076 dynamic
  !          [VCC-WLHost19-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:3d fcid 0x300077 dynamic
  !          [VCC-WLHost30-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:15 fcid 0x300078 dynamic
  !          [VCC-WLHost11-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:38 fcid 0x300079 dynamic
  !          [VCC-WLHost28-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:23 fcid 0x30007a dynamic
  !          [VCC-WLHost17-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:00 fcid 0x30007b dynamic
  !          [VCC-WLHost01-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:04 fcid 0x30007c dynamic
  !          [VCC-WLHost03-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:03 fcid 0x30007d dynamic
  !          [VCC-WLHost02-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:0f fcid 0x30007e dynamic
  !          [VCC-WLHost08-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:1d fcid 0x30007f dynamic
  !          [VCC-WLHost15-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:31 fcid 0x300080 dynamic
  !          [VCC-WLHost24-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:30 fcid 0x300081 dynamic
  !          [VCC-WLHost24-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:02 fcid 0x300082 dynamic
  !          [VCC-WLHost02-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:08 fcid 0x300083 dynamic
  !          [VCC-WLHost05-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:26 fcid 0x300084 dynamic
  !          [VCC-WLHost19-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:22 fcid 0x300085 dynamic
  !          [VCC-WLHost17-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:2c fcid 0x300086 dynamic
  !          [VCC-WLHost22-HBA0]
```

```
vsan 100 wwn 20:00:00:25:b5:aa:17:33 fcid 0x300087 dynamic
  !         [VCC-WLHost25-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:21 fcid 0x300088 dynamic
  !         [VCC-WLHost16-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:2d fcid 0x300089 dynamic
  !         [VCC-WLHost22-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:24 fcid 0x30008a dynamic
  !         [VCC-WLHost18-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:3f fcid 0x30008b dynamic
  !         [VCC-Infra02-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:39 fcid 0x30008c dynamic
  !         [VCC-WLHost28-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:3c fcid 0x30008d dynamic
  !         [VCC-WLHost30-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:14 fcid 0x30008e dynamic
  !         [VCC-WLHost11-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:11 fcid 0x30008f dynamic
  !         [VCC-WLHost09-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:10 fcid 0x300090 dynamic
  !         [VCC-WLHost09-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:05 fcid 0x300091 dynamic
  !         [VCC-WLHost03-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:0e fcid 0x300092 dynamic
  !         [VCC-WLHost08-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:0d fcid 0x300093 dynamic
  !         [VCC-WLHost07-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:0c fcid 0x300094 dynamic
  !         [VCC-WLHost07-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:1e fcid 0x300095 dynamic
  !         [VCC-Infra01-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:0b fcid 0x300096 dynamic
  !         [VCC-WLHost06-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:28 fcid 0x300097 dynamic
  !         [VCC-WLHost20-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:37 fcid 0x300098 dynamic
  !         [VCC-WLHost27-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:3b fcid 0x300099 dynamic
  !         [VCC-WLHost29-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:09 fcid 0x30009a dynamic
```

```
  !          [VCC-WLHost05-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:2a fcid 0x30009b dynamic
  !          [VCC-WLHost21-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:2f fcid 0x30009c dynamic
  !          [VCC-WLHost23-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:20 fcid 0x30009d dynamic
  !          [VCC-WLHost16-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:35 fcid 0x30009e dynamic
  !          [VCC-WLHost26-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:18 fcid 0x30009f dynamic
  !          [VCC-WLHost13-HBA0]
vsan 100 wwn 20:02:00:de:fb:92:8d:00 fcid 0x3000a0 dynamic
vsan 100 wwn 20:04:00:de:fb:92:8d:00 fcid 0x3000c0 dynamic
vsan 100 wwn 20:01:00:de:fb:92:8d:00 fcid 0x3000e0 dynamic
vsan 100 wwn 52:4a:93:75:dd:91:0a:00 fcid 0x300044 dynamic
  !          [X70-CT0-FC0]
vsan 100 wwn 20:01:00:3a:9c:0e:33:20 fcid 0x3000e1 dynamic
vsan 100 wwn 20:02:00:3a:9c:0e:33:20 fcid 0x3000a1 dynamic
vsan 100 wwn 20:04:00:3a:9c:0e:33:20 fcid 0x3000c1 dynamic
vsan 100 wwn 20:03:00:3a:9c:0e:33:20 fcid 0x300100 dynamic
vsan 100 wwn 52:4a:93:75:dd:91:0a:10 fcid 0x300021 dynamic
  !          [X70-CT1-FC0]
vsan 100 wwn 52:4a:93:71:56:84:09:12 fcid 0x300022 dynamic
vsan 100 wwn 52:4a:93:71:56:84:09:10 fcid 0x300045 dynamic
  !          [X70R3-CT1-FC0]
vsan 100 wwn 52:4a:93:71:56:84:09:02 fcid 0x300046 dynamic
vsan 100 wwn 52:4a:93:71:56:84:09:00 fcid 0x300023 dynamic
  !          [X70R3-CT0-FC0]
vsan 100 wwn 20:00:00:25:b5:aa:17:40 fcid 0x3000e2 dynamic
  !          [AMD-VMHost70-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:41 fcid 0x3000a2 dynamic
  !          [AMD-VMHost70-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:44 fcid 0x3000e3 dynamic
  !          [AMD-VMHost72-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:45 fcid 0x3000a3 dynamic
  !          [AMD-VMHost72-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:4e fcid 0x3000e4 dynamic
  !          [AMD-VMHost73-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:4f fcid 0x3000a4 dynamic
```

```
    !          [AMD-VMHost73-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:42 fcid 0x3000e5 dynamic
    !          [AMD-VMHost71-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:43 fcid 0x3000a5 dynamic
    !          [AMD-VMHost71-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:46 fcid 0x3000e6 dynamic
    !          [AMD-VMHost74-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:47 fcid 0x3000a6 dynamic
    !          [AMD-VMHost74-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:48 fcid 0x3000e7 dynamic
    !          [AMD-VMHost75-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:49 fcid 0x3000a7 dynamic
    !          [AMD-VMHost75-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:4a fcid 0x3000e8 dynamic
    !          [AMD-VMHost76-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:4b fcid 0x3000a8 dynamic
    !          [AMD-VMHost76-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:4c fcid 0x3000e9 dynamic
    !          [AMD-VMHost77-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:4d fcid 0x3000a9 dynamic
    !          [AMD-VMHost77-HBA2]

!Active Zone Database Section for vsan 100
zone name FlaskStack-VCC-CVD-WLHost01 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:00
    !           [VCC-WLHost01-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:01
    !           [VCC-WLHost01-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:02
    !           [VCC-WLHost02-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:03
    !           [VCC-WLHost02-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:04
    !           [VCC-WLHost03-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:05
    !           [VCC-WLHost03-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:06
    !           [VCC-WLHost04-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:07
    !           [VCC-WLHost04-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:08
    !           [VCC-WLHost05-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:09
    !           [VCC-WLHost05-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0a
    !           [VCC-WLHost06-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0b
    !           [VCC-WLHost06-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0c
    !           [VCC-WLHost07-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0d
    !           [VCC-WLHost07-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0e
    !           [VCC-WLHost08-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0f
    !           [VCC-WLHost08-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:10
    !           [VCC-WLHost09-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:11
    !           [VCC-WLHost09-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:12
    !           [VCC-WLHost10-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:13
    !           [VCC-WLHost10-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:14
    !           [VCC-WLHost11-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:15
    !           [VCC-WLHost11-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:16
    !           [VCC-WLHost12-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:17
    !           [VCC-WLHost12-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:18
    !           [VCC-WLHost13-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:19
    !           [VCC-WLHost13-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1a
    !           [VCC-WLHost14-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1b
    !           [VCC-WLHost14-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1c
    !           [VCC-WLHost15-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1d
    !           [VCC-WLHost15-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1e
    !           [VCC-Infra01-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1f
    !           [VCC-Infra01-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:20
    !           [VCC-WLHost16-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:21
    !           [VCC-WLHost16-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:22
    !            [VCC-WLHost17-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:23
    !            [VCC-WLHost17-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:24
    !            [VCC-WLHost18-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:25
    !            [VCC-WLHost18-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:26
    !            [VCC-WLHost19-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:27
    !            [VCC-WLHost19-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:28
    !            [VCC-WLHost20-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:29
    !            [VCC-WLHost20-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100
```

```
    member pwwn 20:00:00:25:b5:aa:17:2a
    !           [VCC-WLHost21-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2b
    !           [VCC-WLHost21-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:2c
    !           [VCC-WLHost22-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2d
    !           [VCC-WLHost22-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:2e
    !           [VCC-WLHost23-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2f
    !           [VCC-WLHost23-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:30
    !           [VCC-WLHost24-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:31
    !           [VCC-WLHost24-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]
```

```
zone name FlaskStack-VCC-CVD-WLHost25 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:32
    !           [VCC-WLHost25-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:33
    !           [VCC-WLHost25-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:34
    !           [VCC-WLHost26-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:35
    !           [VCC-WLHost26-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:36
    !           [VCC-WLHost27-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:37
    !           [VCC-WLHost27-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:38
    !           [VCC-WLHost28-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:39
    !           [VCC-WLHost28-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]
```

```
zone name FlaskStack-VCC-CVD-WLHost29 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3a
    !           [VCC-WLHost29-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3b
    !           [VCC-WLHost29-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3c
    !           [VCC-WLHost30-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3d
    !           [VCC-WLHost30-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3e
    !           [VCC-Infra02-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3f
    !           [VCC-Infra02-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost70 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:40
    !           [AMD-VMHost70-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:41
    !           [AMD-VMHost70-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost71 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:42
    !           [AMD-VMHost71-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:43
    !           [AMD-VMHost71-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost72 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:44
    !           [AMD-VMHost72-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:45
    !           [AMD-VMHost72-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost73 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:4e
    !           [AMD-VMHost73-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4f
    !           [AMD-VMHost73-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost74 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:46
    !           [AMD-VMHost74-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:47
    !           [AMD-VMHost74-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost75 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:48
    !           [AMD-VMHost75-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:49
    !           [AMD-VMHost75-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost76 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:4a
    !           [AMD-VMHost76-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4b
    !           [AMD-VMHost76-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost77 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:4c
    !           [AMD-VMHost77-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4d
    !           [AMD-VMHost77-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zoneset name FlashStack-VCC-CVD vsan 100
    member FlaskStack-VCC-CVD-WLHost01
    member FlaskStack-VCC-CVD-WLHost02
    member FlaskStack-VCC-CVD-WLHost03
    member FlaskStack-VCC-CVD-WLHost04
    member FlaskStack-VCC-CVD-WLHost05
    member FlaskStack-VCC-CVD-WLHost06
    member FlaskStack-VCC-CVD-WLHost07
    member FlaskStack-VCC-CVD-WLHost08
    member FlaskStack-VCC-CVD-WLHost09
    member FlaskStack-VCC-CVD-WLHost10
    member FlaskStack-VCC-CVD-WLHost11
    member FlaskStack-VCC-CVD-WLHost12
    member FlaskStack-VCC-CVD-WLHost13
    member FlaskStack-VCC-CVD-WLHost14
    member FlaskStack-VCC-CVD-WLHost15
    member FlaskStack-VCC-CVD-Infra01
    member FlaskStack-VCC-CVD-WLHost16
    member FlaskStack-VCC-CVD-WLHost17
    member FlaskStack-VCC-CVD-WLHost18
    member FlaskStack-VCC-CVD-WLHost19
    member FlaskStack-VCC-CVD-WLHost20
    member FlaskStack-VCC-CVD-WLHost21
    member FlaskStack-VCC-CVD-WLHost22
    member FlaskStack-VCC-CVD-WLHost23
    member FlaskStack-VCC-CVD-WLHost24
    member FlaskStack-VCC-CVD-WLHost25
    member FlaskStack-VCC-CVD-WLHost26
    member FlaskStack-VCC-CVD-WLHost27
    member FlaskStack-VCC-CVD-WLHost28
    member FlaskStack-VCC-CVD-WLHost29
    member FlaskStack-VCC-CVD-WLHost30
    member FlaskStack-VCC-CVD-Infra02
    member FlaskStack-AMD-VMHost70
    member FlaskStack-AMD-VMHost71
    member FlaskStack-AMD-VMHost72
    member FlaskStack-AMD-VMHost73
    member FlaskStack-AMD-VMHost74
    member FlaskStack-AMD-VMHost75
    member FlaskStack-AMD-VMHost76
    member FlaskStack-AMD-VMHost77

zoneset activate name FlashStack-VCC-CVD vsan 100
do clear zone database vsan 100
!Full Zone Database Section for vsan 100
```

```
zone name FlaskStack-VCC-CVD-WLHost01 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:00
    !           [VCC-WLHost01-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:01
    !           [VCC-WLHost01-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:02
    !           [VCC-WLHost02-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:03
    !           [VCC-WLHost02-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:04
    !           [VCC-WLHost03-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:05
    !           [VCC-WLHost03-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:06
    !           [VCC-WLHost04-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:07
    !           [VCC-WLHost04-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]
```

```
zone name FlaskStack-VCC-CVD-WLHost05 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:08
    !           [VCC-WLHost05-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:09
    !           [VCC-WLHost05-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0a
    !           [VCC-WLHost06-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0b
    !           [VCC-WLHost06-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0c
    !           [VCC-WLHost07-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0d
    !           [VCC-WLHost07-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:0e
    !           [VCC-WLHost08-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:0f
    !           [VCC-WLHost08-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:10
    !           [VCC-WLHost09-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:11
    !           [VCC-WLHost09-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:12
    !           [VCC-WLHost10-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:13
    !           [VCC-WLHost10-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:14
    !           [VCC-WLHost11-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:15
    !           [VCC-WLHost11-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:16
    !           [VCC-WLHost12-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:17
    !           [VCC-WLHost12-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:18
    !           [VCC-WLHost13-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:19
    !           [VCC-WLHost13-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1a
    !           [VCC-WLHost14-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1b
    !           [VCC-WLHost14-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1c
    !           [VCC-WLHost15-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1d
    !           [VCC-WLHost15-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:1e
    !           [VCC-Infra01-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:1f
    !           [VCC-Infra01-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:20
    !           [VCC-WLHost16-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:21
    !           [VCC-WLHost16-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:22
    !           [VCC-WLHost17-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:23
    !           [VCC-WLHost17-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:24
    !           [VCC-WLHost18-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:25
    !           [VCC-WLHost18-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:26
    !           [VCC-WLHost19-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:27
    !           [VCC-WLHost19-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:28
    !           [VCC-WLHost20-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:29
    !           [VCC-WLHost20-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:2a
    !           [VCC-WLHost21-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2b
    !           [VCC-WLHost21-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:2c
    !           [VCC-WLHost22-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2d
    !           [VCC-WLHost22-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:2e
    !           [VCC-WLHost23-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:2f
    !           [VCC-WLHost23-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:30
    !           [VCC-WLHost24-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:31
    !           [VCC-WLHost24-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:32
    !           [VCC-WLHost25-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:33
    !           [VCC-WLHost25-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:34
    !           [VCC-WLHost26-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:35
    !           [VCC-WLHost26-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:36
    !           [VCC-WLHost27-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:37
    !            [VCC-WLHost27-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:38
    !            [VCC-WLHost28-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:39
    !            [VCC-WLHost28-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3a
    !            [VCC-WLHost29-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3b
    !            [VCC-WLHost29-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3c
    !            [VCC-WLHost30-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3d
    !            [VCC-WLHost30-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !            [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !            [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:3e
    !           [VCC-Infra02-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:3f
    !           [VCC-Infra02-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost70 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:40
    !           [AMD-VMHost70-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:41
    !           [AMD-VMHost70-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost71 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:42
    !           [AMD-VMHost71-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:43
    !           [AMD-VMHost71-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost72 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:44
    !           [AMD-VMHost72-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:45
    !           [AMD-VMHost72-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost73 vsan 100
```

```
    member pwwn 20:00:00:25:b5:aa:17:4e
    !           [AMD-VMHost73-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4f
    !           [AMD-VMHost73-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost74 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:46
    !           [AMD-VMHost74-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:47
    !           [AMD-VMHost74-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost75 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:48
    !           [AMD-VMHost75-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:49
    !           [AMD-VMHost75-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost76 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:4a
    !           [AMD-VMHost76-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4b
    !           [AMD-VMHost76-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]
```

```
zone name FlaskStack-AMD-VMHost77 vsan 100
    member pwwn 20:00:00:25:b5:aa:17:4c
    !           [AMD-VMHost77-HBA0]
    member pwwn 20:00:00:25:b5:aa:17:4d
    !           [AMD-VMHost77-HBA2]
    member pwwn 52:4a:93:71:56:84:09:00
    !           [X70R3-CT0-FC0]
    member pwwn 52:4a:93:71:56:84:09:10
    !           [X70R3-CT1-FC0]

zoneset name FlashStack-VCC-CVD vsan 100
    member FlaskStack-VCC-CVD-WLHost01
    member FlaskStack-VCC-CVD-WLHost02
    member FlaskStack-VCC-CVD-WLHost03
    member FlaskStack-VCC-CVD-WLHost04
    member FlaskStack-VCC-CVD-WLHost05
    member FlaskStack-VCC-CVD-WLHost06
    member FlaskStack-VCC-CVD-WLHost07
    member FlaskStack-VCC-CVD-WLHost08
    member FlaskStack-VCC-CVD-WLHost09
    member FlaskStack-VCC-CVD-WLHost10
    member FlaskStack-VCC-CVD-WLHost11
    member FlaskStack-VCC-CVD-WLHost12
    member FlaskStack-VCC-CVD-WLHost13
    member FlaskStack-VCC-CVD-WLHost14
    member FlaskStack-VCC-CVD-WLHost15
    member FlaskStack-VCC-CVD-Infra01
    member FlaskStack-VCC-CVD-WLHost16
    member FlaskStack-VCC-CVD-WLHost17
    member FlaskStack-VCC-CVD-WLHost18
    member FlaskStack-VCC-CVD-WLHost19
    member FlaskStack-VCC-CVD-WLHost20
    member FlaskStack-VCC-CVD-WLHost21
    member FlaskStack-VCC-CVD-WLHost22
    member FlaskStack-VCC-CVD-WLHost23
    member FlaskStack-VCC-CVD-WLHost24
    member FlaskStack-VCC-CVD-WLHost25
    member FlaskStack-VCC-CVD-WLHost26
    member FlaskStack-VCC-CVD-WLHost27
    member FlaskStack-VCC-CVD-WLHost28
    member FlaskStack-VCC-CVD-WLHost29
    member FlaskStack-VCC-CVD-WLHost30
    member FlaskStack-VCC-CVD-Infra02
    member FlaskStack-AMD-VMHost70
    member FlaskStack-AMD-VMHost71
    member FlaskStack-AMD-VMHost72
    member FlaskStack-AMD-VMHost73
    member FlaskStack-AMD-VMHost74
    member FlaskStack-AMD-VMHost75
    member FlaskStack-AMD-VMHost76
    member FlaskStack-AMD-VMHost77

interface mgmt0
  ip address 10.29.164.238 255.255.255.0
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
  vsan 100 interface fc1/9
  vsan 100 interface fc1/10
  vsan 100 interface fc1/11
  vsan 100 interface fc1/12
  vsan 100 interface fc1/13
  vsan 100 interface fc1/14
  vsan 100 interface fc1/15
  vsan 100 interface fc1/16
clock timezone PST 0 0
clock summer-time PDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
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 fc1/4
  switchport speed auto
interface fc1/1
interface fc1/2
interface fc1/3
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 fc1/4

interface fc1/1
  port-license acquire
  no shutdown

interface fc1/2
  port-license acquire
  no shutdown

interface fc1/3
  port-license acquire
  no shutdown

interface fc1/4
  port-license acquire
  no shutdown

interface fc1/5
  no port-license
```

```
interface fc1/6
  no port-license

interface fc1/7
  no port-license

interface fc1/8
  no port-license

interface fc1/9
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/10
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/11
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/12
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/13
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/14
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/15
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/16
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown
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/4Zzg6EekVHWpA8lA7Mc/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
```

```
snmp-server host 10.24.66.169 traps version 2c public udp-port 1164
snmp-server host 10.24.72.119 traps version 2c public udp-port 1166
snmp-server host 10.29.164.250 traps version 2c public udp-port 1163
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 101 name "FlashStack-VCC-CVD-Fabric-B"
device-alias database
  device-alias name X70R3-CT0-FC2 pwwn 52:4a:93:71:56:84:09:02
  device-alias name X70R3-CT1-FC2 pwwn 52:4a:93:71:56:84:09:12
  device-alias name VCC-Infra01-HBA1 pwwn 20:00:00:25:b5:bb:17:1e
  device-alias name VCC-Infra01-HBA3 pwwn 20:00:00:25:b5:bb:17:1f
  device-alias name VCC-Infra02-HBA1 pwwn 20:00:00:25:b5:bb:17:3e
  device-alias name VCC-Infra02-HBA3 pwwn 20:00:00:25:b5:bb:17:3f
  device-alias name VCC-WLHost01-HBA1 pwwn 20:00:00:25:b5:bb:17:00
  device-alias name VCC-WLHost01-HBA3 pwwn 20:00:00:25:b5:bb:17:01
  device-alias name VCC-WLHost02-HBA1 pwwn 20:00:00:25:b5:bb:17:02
  device-alias name VCC-WLHost02-HBA3 pwwn 20:00:00:25:b5:bb:17:03
  device-alias name VCC-WLHost03-HBA1 pwwn 20:00:00:25:b5:bb:17:04
  device-alias name VCC-WLHost03-HBA3 pwwn 20:00:00:25:b5:bb:17:05
  device-alias name VCC-WLHost04-HBA1 pwwn 20:00:00:25:b5:bb:17:06
  device-alias name VCC-WLHost04-HBA3 pwwn 20:00:00:25:b5:bb:17:07
  device-alias name VCC-WLHost05-HBA1 pwwn 20:00:00:25:b5:bb:17:08
  device-alias name VCC-WLHost05-HBA3 pwwn 20:00:00:25:b5:bb:17:09
  device-alias name VCC-WLHost06-HBA1 pwwn 20:00:00:25:b5:bb:17:0a
  device-alias name VCC-WLHost06-HBA3 pwwn 20:00:00:25:b5:bb:17:0b
  device-alias name VCC-WLHost07-HBA1 pwwn 20:00:00:25:b5:bb:17:0c
  device-alias name VCC-WLHost07-HBA3 pwwn 20:00:00:25:b5:bb:17:0d
  device-alias name VCC-WLHost08-HBA1 pwwn 20:00:00:25:b5:bb:17:0e
  device-alias name VCC-WLHost08-HBA3 pwwn 20:00:00:25:b5:bb:17:0f
  device-alias name VCC-WLHost09-HBA1 pwwn 20:00:00:25:b5:bb:17:10
  device-alias name VCC-WLHost09-HBA3 pwwn 20:00:00:25:b5:bb:17:11
  device-alias name VCC-WLHost10-HBA1 pwwn 20:00:00:25:b5:bb:17:12
  device-alias name VCC-WLHost10-HBA3 pwwn 20:00:00:25:b5:bb:17:13
```

```
device-alias name VCC-WLHost11-HBA1 pwwn 20:00:00:25:b5:bb:17:14
device-alias name VCC-WLHost11-HBA3 pwwn 20:00:00:25:b5:bb:17:15
device-alias name VCC-WLHost12-HBA1 pwwn 20:00:00:25:b5:bb:17:16
device-alias name VCC-WLHost12-HBA3 pwwn 20:00:00:25:b5:bb:17:17
device-alias name VCC-WLHost13-HBA1 pwwn 20:00:00:25:b5:bb:17:18
device-alias name VCC-WLHost13-HBA3 pwwn 20:00:00:25:b5:bb:17:19
device-alias name VCC-WLHost14-HBA1 pwwn 20:00:00:25:b5:bb:17:1a
device-alias name VCC-WLHost14-HBA3 pwwn 20:00:00:25:b5:bb:17:1b
device-alias name VCC-WLHost15-HBA1 pwwn 20:00:00:25:b5:bb:17:1c
device-alias name VCC-WLHost15-HBA3 pwwn 20:00:00:25:b5:bb:17:1d
device-alias name VCC-WLHost16-HBA1 pwwn 20:00:00:25:b5:bb:17:20
device-alias name VCC-WLHost16-HBA3 pwwn 20:00:00:25:b5:bb:17:21
device-alias name VCC-WLHost17-HBA1 pwwn 20:00:00:25:b5:bb:17:22
device-alias name VCC-WLHost17-HBA3 pwwn 20:00:00:25:b5:bb:17:23
device-alias name VCC-WLHost18-HBA1 pwwn 20:00:00:25:b5:bb:17:24
device-alias name VCC-WLHost18-HBA3 pwwn 20:00:00:25:b5:bb:17:25
device-alias name VCC-WLHost19-HBA1 pwwn 20:00:00:25:b5:bb:17:26
device-alias name VCC-WLHost19-HBA3 pwwn 20:00:00:25:b5:bb:17:27
device-alias name VCC-WLHost20-HBA1 pwwn 20:00:00:25:b5:bb:17:28
device-alias name VCC-WLHost20-HBA3 pwwn 20:00:00:25:b5:bb:17:29
device-alias name VCC-WLHost21-HBA1 pwwn 20:00:00:25:b5:bb:17:2a
device-alias name VCC-WLHost21-HBA3 pwwn 20:00:00:25:b5:bb:17:2b
device-alias name VCC-WLHost22-HBA1 pwwn 20:00:00:25:b5:bb:17:2c
device-alias name VCC-WLHost22-HBA3 pwwn 20:00:00:25:b5:bb:17:2d
device-alias name VCC-WLHost23-HBA1 pwwn 20:00:00:25:b5:bb:17:2e
device-alias name VCC-WLHost23-HBA3 pwwn 20:00:00:25:b5:bb:17:2f
device-alias name VCC-WLHost24-HBA1 pwwn 20:00:00:25:b5:bb:17:30
device-alias name VCC-WLHost24-HBA3 pwwn 20:00:00:25:b5:bb:17:31
device-alias name VCC-WLHost25-HBA1 pwwn 20:00:00:25:b5:bb:17:32
device-alias name VCC-WLHost25-HBA3 pwwn 20:00:00:25:b5:bb:17:33
device-alias name VCC-WLHost26-HBA1 pwwn 20:00:00:25:b5:bb:17:34
device-alias name VCC-WLHost26-HBA3 pwwn 20:00:00:25:b5:bb:17:35
device-alias name VCC-WLHost27-HBA1 pwwn 20:00:00:25:b5:bb:17:36
device-alias name VCC-WLHost27-HBA3 pwwn 20:00:00:25:b5:bb:17:37
device-alias name VCC-WLHost28-HBA1 pwwn 20:00:00:25:b5:bb:17:38
device-alias name VCC-WLHost28-HBA3 pwwn 20:00:00:25:b5:bb:17:39
device-alias name VCC-WLHost29-HBA1 pwwn 20:00:00:25:b5:bb:17:3a
device-alias name VCC-WLHost29-HBA3 pwwn 20:00:00:25:b5:bb:17:3b
device-alias name VCC-WLHost30-HBA1 pwwn 20:00:00:25:b5:bb:17:3c
```

```
  device-alias name VCC-WLHost30-HBA3 pwwn 20:00:00:25:b5:bb:17:3d

device-alias commit

fcdomain fcid database
  vsan 101 wwn 20:03:00:de:fb:90:a4:40 fcid 0xc40000 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:17 fcid 0xc40020 dynamic
    !          [X70-CT1-FC9]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:07 fcid 0xc40040 dynamic
    !          [X70-CT0-FC9]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:16 fcid 0xc40021 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:13 fcid 0xc40041 dynamic
    !          [X70-CT1-FC3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3e fcid 0xc40060 dynamic
    !          [VCC-Infra02-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:07 fcid 0xc40061 dynamic
    !          [VCC-WLHost04-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3c fcid 0xc40062 dynamic
    !          [VCC-WLHost30-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:11 fcid 0xc40063 dynamic
    !          [VCC-WLHost09-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:01 fcid 0xc40064 dynamic
    !          [VCC-WLHost01-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:00 fcid 0xc40065 dynamic
    !          [VCC-WLHost01-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:13 fcid 0xc40066 dynamic
    !          [VCC-WLHost10-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:04 fcid 0xc40067 dynamic
    !          [VCC-WLHost03-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:17 fcid 0xc40068 dynamic
    !          [VCC-WLHost12-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:16 fcid 0xc40069 dynamic
    !          [VCC-WLHost12-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:30 fcid 0xc4006a dynamic
    !          [VCC-WLHost24-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:21 fcid 0xc4006b dynamic
    !          [VCC-WLHost16-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:1f fcid 0xc4006c dynamic
    !          [VCC-Infra01-HBA3]
```

```
vsan 101 wwn 20:00:00:25:b5:bb:17:1a fcid 0xc4006d dynamic
  !          [VCC-WLHost14-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:3f fcid 0xc4006e dynamic
  !          [VCC-Infra02-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0a fcid 0xc4006f dynamic
  !          [VCC-WLHost06-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:38 fcid 0xc40070 dynamic
  !          [VCC-WLHost28-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:19 fcid 0xc40071 dynamic
  !          [VCC-WLHost13-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:22 fcid 0xc40072 dynamic
  !          [VCC-WLHost17-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2f fcid 0xc40073 dynamic
  !          [VCC-WLHost23-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1b fcid 0xc40074 dynamic
  !          [VCC-WLHost14-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3b fcid 0xc40075 dynamic
  !          [VCC-WLHost29-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2a fcid 0xc40076 dynamic
  !          [VCC-WLHost21-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:29 fcid 0xc40077 dynamic
  !          [VCC-WLHost20-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1c fcid 0xc40078 dynamic
  !          [VCC-WLHost15-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0b fcid 0xc40079 dynamic
  !          [VCC-WLHost06-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0d fcid 0xc4007a dynamic
  !          [VCC-WLHost07-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:37 fcid 0xc4007b dynamic
  !          [VCC-WLHost27-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:31 fcid 0xc4007c dynamic
  !          [VCC-WLHost24-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:08 fcid 0xc4007d dynamic
  !          [VCC-WLHost05-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:10 fcid 0xc4007e dynamic
  !          [VCC-WLHost09-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:34 fcid 0xc4007f dynamic
  !          [VCC-WLHost26-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:25 fcid 0xc40080 dynamic
```

```
  !          [VCC-WLHost18-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3d fcid 0xc40081 dynamic
  !          [VCC-WLHost30-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:15 fcid 0xc40082 dynamic
  !          [VCC-WLHost11-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:23 fcid 0xc40083 dynamic
  !          [VCC-WLHost17-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3a fcid 0xc40084 dynamic
  !          [VCC-WLHost29-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:28 fcid 0xc40085 dynamic
  !          [VCC-WLHost20-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:32 fcid 0xc40086 dynamic
  !          [VCC-WLHost25-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0f fcid 0xc40087 dynamic
  !          [VCC-WLHost08-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0c fcid 0xc40088 dynamic
  !          [VCC-WLHost07-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2e fcid 0xc40089 dynamic
  !          [VCC-WLHost23-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:03 fcid 0xc4008a dynamic
  !          [VCC-WLHost02-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:02 fcid 0xc4008b dynamic
  !          [VCC-WLHost02-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2b fcid 0xc4008c dynamic
  !          [VCC-WLHost21-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:35 fcid 0xc4008d dynamic
  !          [VCC-WLHost26-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2c fcid 0xc4008e dynamic
  !          [VCC-WLHost22-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:27 fcid 0xc4008f dynamic
  !          [VCC-WLHost19-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:18 fcid 0xc40090 dynamic
  !          [VCC-WLHost13-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:14 fcid 0xc40091 dynamic
  !          [VCC-WLHost11-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0e fcid 0xc40092 dynamic
  !          [VCC-WLHost08-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:1e fcid 0xc40093 dynamic
  !          [VCC-Infra01-HBA1]
```

```
vsan 101 wwn 20:00:00:25:b5:bb:17:06 fcid 0xc40094 dynamic
  !          [VCC-WLHost04-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:09 fcid 0xc40095 dynamic
  !          [VCC-WLHost05-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:26 fcid 0xc40096 dynamic
  !          [VCC-WLHost19-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:24 fcid 0xc40097 dynamic
  !          [VCC-WLHost18-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:20 fcid 0xc40098 dynamic
  !          [VCC-WLHost16-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:1d fcid 0xc40099 dynamic
  !          [VCC-WLHost15-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:33 fcid 0xc4009a dynamic
  !          [VCC-WLHost25-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:36 fcid 0xc4009b dynamic
  !          [VCC-WLHost27-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:39 fcid 0xc4009c dynamic
  !          [VCC-WLHost28-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2d fcid 0xc4009d dynamic
  !          [VCC-WLHost22-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:12 fcid 0xc4009e dynamic
  !          [VCC-WLHost10-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:05 fcid 0xc4009f dynamic
  !          [VCC-WLHost03-HBA3]
vsan 101 wwn 20:02:00:de:fb:90:a4:40 fcid 0xc400a0 dynamic
vsan 101 wwn 20:01:00:de:fb:90:a4:40 fcid 0xc400c0 dynamic
vsan 101 wwn 20:04:00:de:fb:90:a4:40 fcid 0xc400e0 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:00 fcid 0xc40022 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:12 fcid 0xc40042 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:11 fcid 0xc40023 dynamic
  !          [X70-CT1-FC1]
vsan 101 wwn 20:01:00:3a:9c:a4:fd:20 fcid 0xc400c1 dynamic
vsan 101 wwn 20:02:00:3a:9c:a4:fd:20 fcid 0xc400a1 dynamic
vsan 101 wwn 20:03:00:3a:9c:a4:fd:20 fcid 0xc40100 dynamic
vsan 101 wwn 20:04:00:3a:9c:a4:fd:20 fcid 0xc400e1 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:01 fcid 0xc40043 dynamic
  !          [X70-CT0-FC1]
vsan 101 wwn 52:4a:93:71:56:84:09:02 fcid 0xc40044 dynamic
  !          [X70R3-CT0-FC2]
  vsan 101 wwn 52:4a:93:71:56:84:09:00 fcid 0xc40024 dynamic
  vsan 101 wwn 52:4a:93:71:56:84:09:12 fcid 0xc40045 dynamic
    !          [X70R3-CT1-FC2]
  vsan 101 wwn 20:00:00:25:b5:bb:17:40 fcid 0xc400c2 dynamic
    !          [AMD-VMHost70-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:41 fcid 0xc400a2 dynamic
    !          [AMD-VMHost70-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:44 fcid 0xc400c3 dynamic
    !          [AMD-VMHost72-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:45 fcid 0xc400a3 dynamic
    !          [AMD-VMHost72-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4e fcid 0xc400c4 dynamic
    !          [AMD-VMHost73-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4f fcid 0xc400a4 dynamic
    !          [AMD-VMHost73-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:42 fcid 0xc400c5 dynamic
    !          [AMD-VMHost71-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:43 fcid 0xc400a5 dynamic
    !          [AMD-VMHost71-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:46 fcid 0xc400c6 dynamic
    !          [AMD-VMHost74-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:47 fcid 0xc400a6 dynamic
    !          [AMD-VMHost74-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:48 fcid 0xc400c7 dynamic
    !          [AMD-VMHost75-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:49 fcid 0xc400a7 dynamic
    !          [AMD-VMHost75-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4a fcid 0xc400c8 dynamic
    !          [AMD-VMHost76-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4b fcid 0xc400a8 dynamic
    !          [AMD-VMHost76-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4c fcid 0xc400c9 dynamic
    !          [AMD-VMHost77-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:4d fcid 0xc400a9 dynamic
    !          [AMD-VMHost77-HBA3]

!Active Zone Database Section for vsan 101
zone name FlaskStack-VCC-CVD-WLHost01 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:00
    !           [VCC-WLHost01-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:01
    !           [VCC-WLHost01-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:02
    !           [VCC-WLHost02-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:03
    !           [VCC-WLHost02-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:04
    !           [VCC-WLHost03-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:05
    !           [VCC-WLHost03-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:06
    !           [VCC-WLHost04-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:07
    !           [VCC-WLHost04-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 101
```

```
    member pwwn 20:00:00:25:b5:bb:17:08
    !           [VCC-WLHost05-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:09
    !           [VCC-WLHost05-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0a
    !           [VCC-WLHost06-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0b
    !           [VCC-WLHost06-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0c
    !           [VCC-WLHost07-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0d
    !           [VCC-WLHost07-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0e
    !           [VCC-WLHost08-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0f
    !           [VCC-WLHost08-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]
```

```
zone name FlaskStack-VCC-CVD-WLHost09 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:10
    !           [VCC-WLHost09-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:11
    !           [VCC-WLHost09-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:12
    !           [VCC-WLHost10-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:13
    !           [VCC-WLHost10-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:14
    !           [VCC-WLHost11-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:15
    !           [VCC-WLHost11-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:16
    !           [VCC-WLHost12-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:17
    !           [VCC-WLHost12-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]
```

```
zone name FlaskStack-VCC-CVD-WLHost13 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:18
    !           [VCC-WLHost13-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:19
    !           [VCC-WLHost13-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1a
    !           [VCC-WLHost14-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1b
    !           [VCC-WLHost14-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1c
    !           [VCC-WLHost15-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1d
    !           [VCC-WLHost15-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1e
    !           [VCC-Infra01-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1f
    !           [VCC-Infra01-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:20
    !           [VCC-WLHost16-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:21
    !           [VCC-WLHost16-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:22
    !           [VCC-WLHost17-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:23
    !           [VCC-WLHost17-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:24
    !           [VCC-WLHost18-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:25
    !           [VCC-WLHost18-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:26
    !           [VCC-WLHost19-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:27
    !           [VCC-WLHost19-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:28
    !           [VCC-WLHost20-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:29
    !           [VCC-WLHost20-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2a
    !           [VCC-WLHost21-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2b
    !           [VCC-WLHost21-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2c
    !           [VCC-WLHost22-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2d
    !           [VCC-WLHost22-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2e
    !           [VCC-WLHost23-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2f
    !           [VCC-WLHost23-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:30
    !           [VCC-WLHost24-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:31
    !           [VCC-WLHost24-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:32
    !           [VCC-WLHost25-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:33
    !           [VCC-WLHost25-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:34
    !           [VCC-WLHost26-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:35
    !           [VCC-WLHost26-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:36
    !           [VCC-WLHost27-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:37
    !           [VCC-WLHost27-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:38
    !           [VCC-WLHost28-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:39
    !           [VCC-WLHost28-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3a
    !           [VCC-WLHost29-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3b
    !           [VCC-WLHost29-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3c
    !           [VCC-WLHost30-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3d
    !           [VCC-WLHost30-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3e
    !           [VCC-Infra02-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3f
    !           [VCC-Infra02-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost70 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:40
    !           [AMD-VMHost70-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:41
    !           [AMD-VMHost70-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost71 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:42
    !           [AMD-VMHost71-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:43
    !           [AMD-VMHost71-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost72 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:44
    !           [AMD-VMHost72-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:45
    !           [AMD-VMHost72-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost73 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4e
    !           [AMD-VMHost73-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4f
    !           [AMD-VMHost73-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost74 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:46
    !           [AMD-VMHost74-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:47
    !           [AMD-VMHost74-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost75 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:48
    !           [AMD-VMHost75-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:49
    !           [AMD-VMHost75-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost76 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4a
    !           [AMD-VMHost76-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4b
    !           [AMD-VMHost76-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost77 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4c
    !           [AMD-VMHost77-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4d
    !           [AMD-VMHost77-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zoneset name FlashStack-VCC-CVD vsan 101
    member FlaskStack-VCC-CVD-WLHost01
    member FlaskStack-VCC-CVD-WLHost02
    member FlaskStack-VCC-CVD-WLHost03
    member FlaskStack-VCC-CVD-WLHost04
    member FlaskStack-VCC-CVD-WLHost05
    member FlaskStack-VCC-CVD-WLHost06
    member FlaskStack-VCC-CVD-WLHost07
    member FlaskStack-VCC-CVD-WLHost08
    member FlaskStack-VCC-CVD-WLHost09
    member FlaskStack-VCC-CVD-WLHost10
    member FlaskStack-VCC-CVD-WLHost11
    member FlaskStack-VCC-CVD-WLHost12
    member FlaskStack-VCC-CVD-WLHost13
    member FlaskStack-VCC-CVD-WLHost14
    member FlaskStack-VCC-CVD-WLHost15
    member FlaskStack-VCC-CVD-Infra01
    member FlaskStack-VCC-CVD-WLHost16
    member FlaskStack-VCC-CVD-WLHost17
    member FlaskStack-VCC-CVD-WLHost18
    member FlaskStack-VCC-CVD-WLHost19
    member FlaskStack-VCC-CVD-WLHost20
    member FlaskStack-VCC-CVD-WLHost21
    member FlaskStack-VCC-CVD-WLHost22
    member FlaskStack-VCC-CVD-WLHost23
    member FlaskStack-VCC-CVD-WLHost24
    member FlaskStack-VCC-CVD-WLHost25
    member FlaskStack-VCC-CVD-WLHost26
    member FlaskStack-VCC-CVD-WLHost27
    member FlaskStack-VCC-CVD-WLHost28
    member FlaskStack-VCC-CVD-WLHost29
    member FlaskStack-VCC-CVD-WLHost30
    member FlaskStack-VCC-CVD-Infra02
    member FlaskStack-AMD-VMHost70
    member FlaskStack-AMD-VMHost71
    member FlaskStack-AMD-VMHost72
    member FlaskStack-AMD-VMHost73
    member FlaskStack-AMD-VMHost74
    member FlaskStack-AMD-VMHost75
    member FlaskStack-AMD-VMHost76
    member FlaskStack-AMD-VMHost77

zoneset activate name FlashStack-VCC-CVD vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zone name FlaskStack-VCC-CVD-WLHost01 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:00
    !           [VCC-WLHost01-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:01
    !           [VCC-WLHost01-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:02
    !           [VCC-WLHost02-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:03
    !           [VCC-WLHost02-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:04
    !           [VCC-WLHost03-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:05
    !           [VCC-WLHost03-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:06
    !           [VCC-WLHost04-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:07
    !           [VCC-WLHost04-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:08
    !           [VCC-WLHost05-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:09
    !           [VCC-WLHost05-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0a
    !           [VCC-WLHost06-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0b
    !           [VCC-WLHost06-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0c
    !           [VCC-WLHost07-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0d
    !           [VCC-WLHost07-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:0e
    !           [VCC-WLHost08-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:0f
    !           [VCC-WLHost08-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:10
    !           [VCC-WLHost09-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:11
    !           [VCC-WLHost09-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:12
    !           [VCC-WLHost10-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:13
    !           [VCC-WLHost10-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:14
    !           [VCC-WLHost11-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:15
    !           [VCC-WLHost11-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:16
    !           [VCC-WLHost12-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:17
    !           [VCC-WLHost12-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:18
    !           [VCC-WLHost13-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:19
    !           [VCC-WLHost13-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1a
    !           [VCC-WLHost14-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1b
    !           [VCC-WLHost14-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1c
    !           [VCC-WLHost15-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1d
    !           [VCC-WLHost15-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:1e
    !           [VCC-Infra01-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:1f
    !           [VCC-Infra01-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:20
    !           [VCC-WLHost16-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:21
    !           [VCC-WLHost16-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:22
    !           [VCC-WLHost17-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:23
    !           [VCC-WLHost17-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101
```

```
    member pwwn 20:00:00:25:b5:bb:17:24
    !           [VCC-WLHost18-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:25
    !           [VCC-WLHost18-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:26
    !           [VCC-WLHost19-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:27
    !           [VCC-WLHost19-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:28
    !           [VCC-WLHost20-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:29
    !           [VCC-WLHost20-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2a
    !           [VCC-WLHost21-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2b
    !           [VCC-WLHost21-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]
```

```
zone name FlaskStack-VCC-CVD-WLHost22 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2c
    !           [VCC-WLHost22-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2d
    !           [VCC-WLHost22-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:2e
    !           [VCC-WLHost23-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:2f
    !           [VCC-WLHost23-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:30
    !           [VCC-WLHost24-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:31
    !           [VCC-WLHost24-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:32
    !           [VCC-WLHost25-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:33
    !           [VCC-WLHost25-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]
```

```
zone name FlaskStack-VCC-CVD-WLHost26 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:34
    !           [VCC-WLHost26-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:35
    !           [VCC-WLHost26-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:36
    !           [VCC-WLHost27-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:37
    !           [VCC-WLHost27-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:38
    !           [VCC-WLHost28-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:39
    !           [VCC-WLHost28-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3a
    !           [VCC-WLHost29-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3b
    !           [VCC-WLHost29-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3c
    !           [VCC-WLHost30-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3d
    !           [VCC-WLHost30-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:3e
    !           [VCC-Infra02-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:3f
    !           [VCC-Infra02-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost70 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:40
    !           [AMD-VMHost70-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:41
    !           [AMD-VMHost70-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost71 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:42
    !           [AMD-VMHost71-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:43
    !           [AMD-VMHost71-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost72 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:44
    !           [AMD-VMHost72-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:45
    !           [AMD-VMHost72-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost73 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4e
    !           [AMD-VMHost73-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4f
    !           [AMD-VMHost73-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost74 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:46
    !           [AMD-VMHost74-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:47
    !           [AMD-VMHost74-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost75 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:48
    !           [AMD-VMHost75-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:49
    !           [AMD-VMHost75-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost76 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4a
    !           [AMD-VMHost76-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4b
    !           [AMD-VMHost76-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost77 vsan 101
    member pwwn 20:00:00:25:b5:bb:17:4c
    !           [AMD-VMHost77-HBA1]
    member pwwn 20:00:00:25:b5:bb:17:4d
    !           [AMD-VMHost77-HBA3]
    member pwwn 52:4a:93:71:56:84:09:02
    !           [X70R3-CT0-FC2]
    member pwwn 52:4a:93:71:56:84:09:12
    !           [X70R3-CT1-FC2]

zoneset name FlashStack-VCC-CVD vsan 101
    member FlaskStack-VCC-CVD-WLHost01
    member FlaskStack-VCC-CVD-WLHost02
    member FlaskStack-VCC-CVD-WLHost03
    member FlaskStack-VCC-CVD-WLHost04
    member FlaskStack-VCC-CVD-WLHost05
    member FlaskStack-VCC-CVD-WLHost06
    member FlaskStack-VCC-CVD-WLHost07
    member FlaskStack-VCC-CVD-WLHost08
    member FlaskStack-VCC-CVD-WLHost09
    member FlaskStack-VCC-CVD-WLHost10
    member FlaskStack-VCC-CVD-WLHost11
    member FlaskStack-VCC-CVD-WLHost12
    member FlaskStack-VCC-CVD-WLHost13
    member FlaskStack-VCC-CVD-WLHost14
    member FlaskStack-VCC-CVD-WLHost15
    member FlaskStack-VCC-CVD-Infra01
    member FlaskStack-VCC-CVD-WLHost16
    member FlaskStack-VCC-CVD-WLHost17
    member FlaskStack-VCC-CVD-WLHost18
    member FlaskStack-VCC-CVD-WLHost19
    member FlaskStack-VCC-CVD-WLHost20
    member FlaskStack-VCC-CVD-WLHost21
    member FlaskStack-VCC-CVD-WLHost22
    member FlaskStack-VCC-CVD-WLHost23
    member FlaskStack-VCC-CVD-WLHost24
    member FlaskStack-VCC-CVD-WLHost25
    member FlaskStack-VCC-CVD-WLHost26
    member FlaskStack-VCC-CVD-WLHost27
    member FlaskStack-VCC-CVD-WLHost28
    member FlaskStack-VCC-CVD-WLHost29
    member FlaskStack-VCC-CVD-WLHost30
    member FlaskStack-VCC-CVD-Infra02
    member FlaskStack-AMD-VMHost70
    member FlaskStack-AMD-VMHost71
    member FlaskStack-AMD-VMHost72
    member FlaskStack-AMD-VMHost73
    member FlaskStack-AMD-VMHost74
    member FlaskStack-AMD-VMHost75
    member FlaskStack-AMD-VMHost76
    member FlaskStack-AMD-VMHost77



interface mgmt0
  ip address 10.29.164.239 255.255.255.0
vsan database
  vsan 101 interface fc1/9
  vsan 101 interface fc1/10
  vsan 101 interface fc1/11
  vsan 101 interface fc1/12
  vsan 101 interface fc1/13
  vsan 101 interface fc1/14
  vsan 101 interface fc1/15
  vsan 101 interface fc1/16
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 fc1/16

interface fc1/1
  no port-license

interface fc1/2
  no port-license

interface fc1/3
  no port-license

interface fc1/4
 no port-license

interface fc1/5
```

```
  no port-license

interface fc1/6
  no port-license

interface fc1/7
  no port-license

interface fc1/8
  no port-license

interface fc1/9
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/10
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/11
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/12
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/13
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
```

```
    no shutdown

  interface fc1/14
    switchport trunk allowed vsan 101
    switchport trunk mode off
    port-license acquire
    no shutdown

  interface fc1/15
    switchport trunk allowed vsan 101
    switchport trunk mode off
    port-license acquire
    no shutdown

  interface fc1/16
    switchport trunk allowed vsan 101
    switchport trunk mode off
    port-license acquire
    no shutdown
  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 7.0 GA installed on Cisco UCS B200 M5 Blade Server as part of the workload test with VMware Horizon 8 deployed on Pure Storage FlashArray //70 R3 system running LoginVSI v4.1.39 based knowledge worker workload part of the FlashStack reference architecture defined here.

The charts below are defined in the set of 30 hosts in the single performance chart.

VDI Server Performance Monitor Data for One Sample Test: 5022 Users VDI Non-Persistent (Instant Clones) Scale Testing

Figure 125. Full Scale | 5022 Non-Persistent Users| VDI Host | Host CPU Utilization

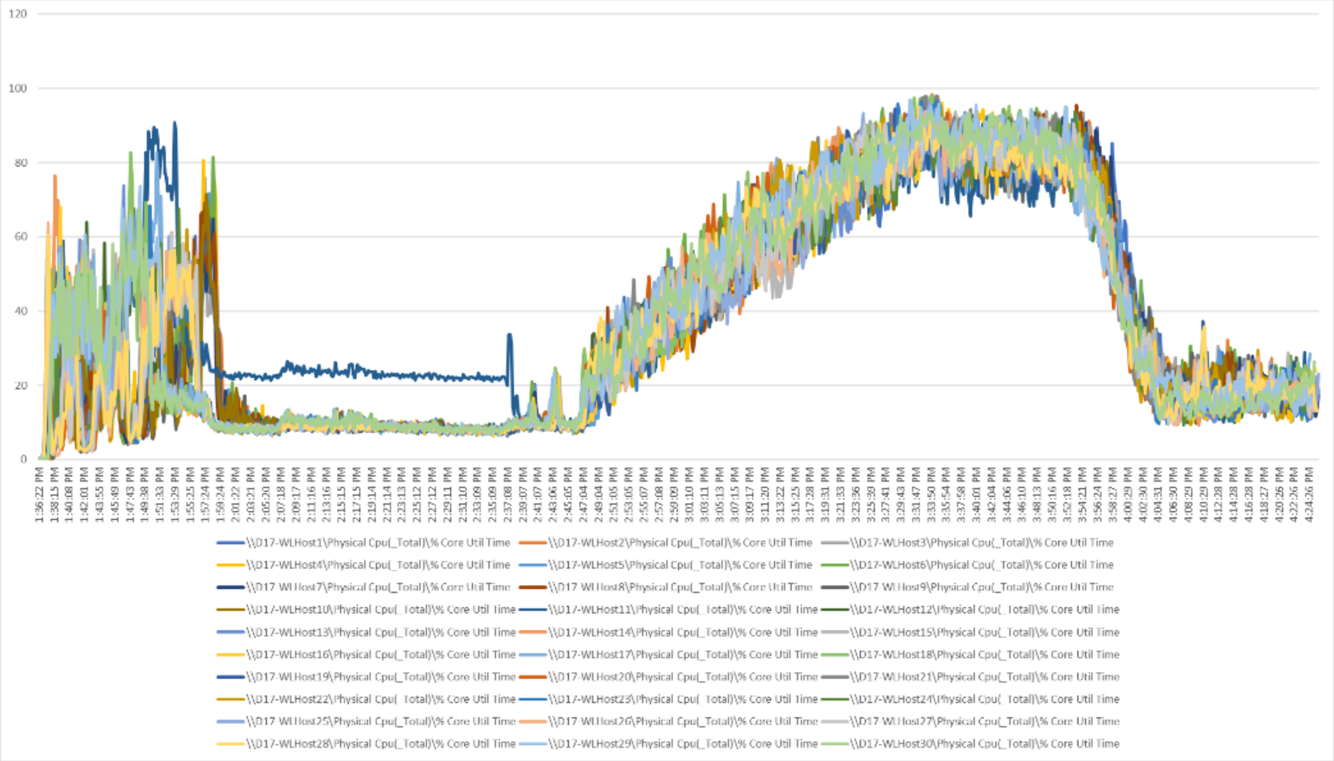

Figure 126. Full Scale | 5022 Non-Persistent Users| VDI Host | Host Memory Utilization

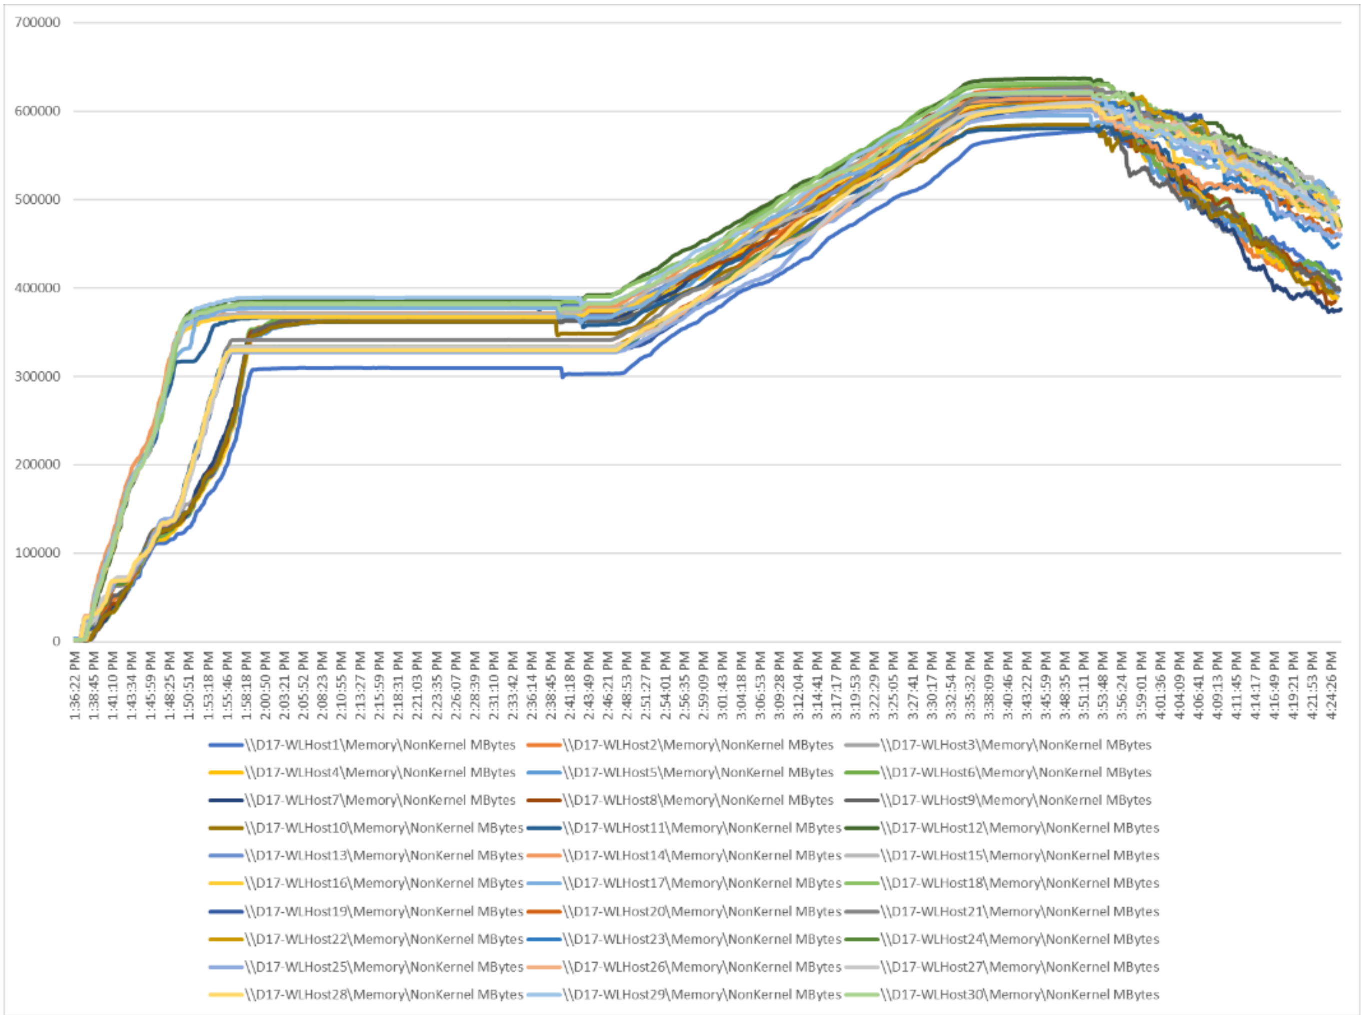

Figure 127. Full Scale | 5022 Non-Persistent Users| VDI Host | Host Network Utilization | Received

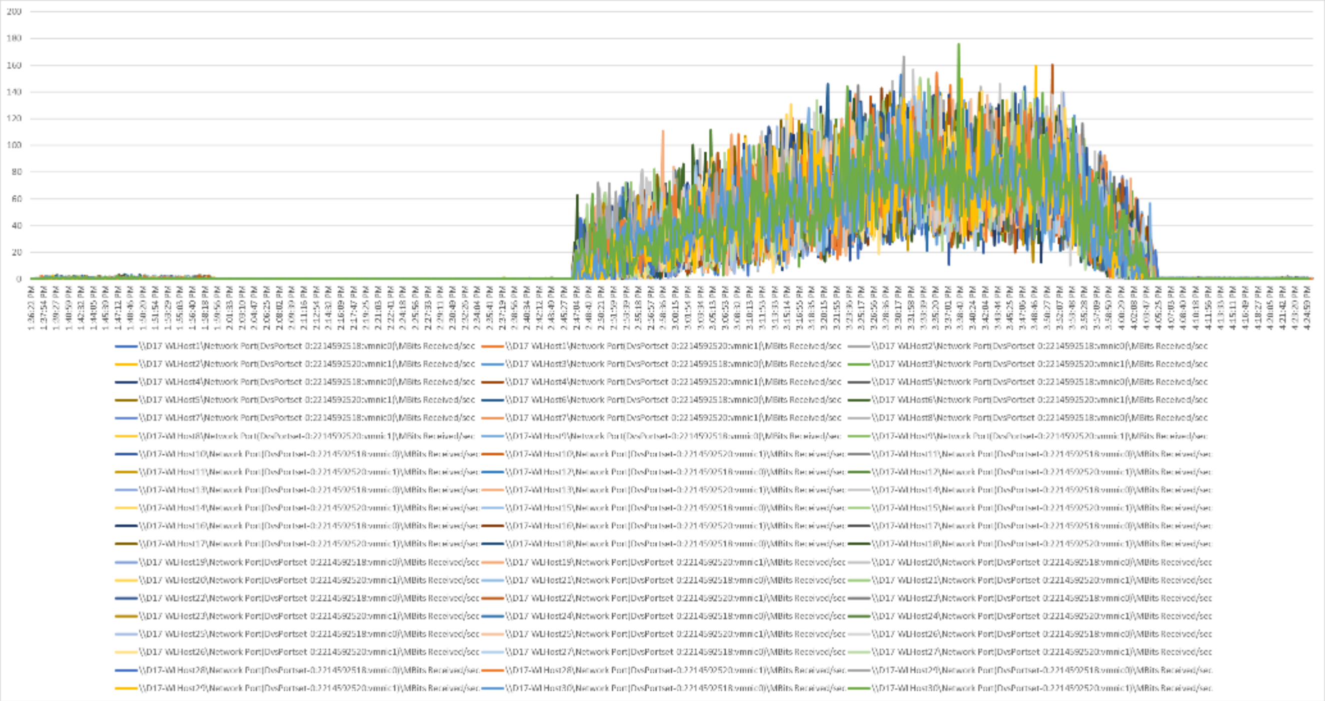

Figure 128. Full Scale | 5022 Non-Persistent Users| VDI Host | Host Network Utilization | Transmitted

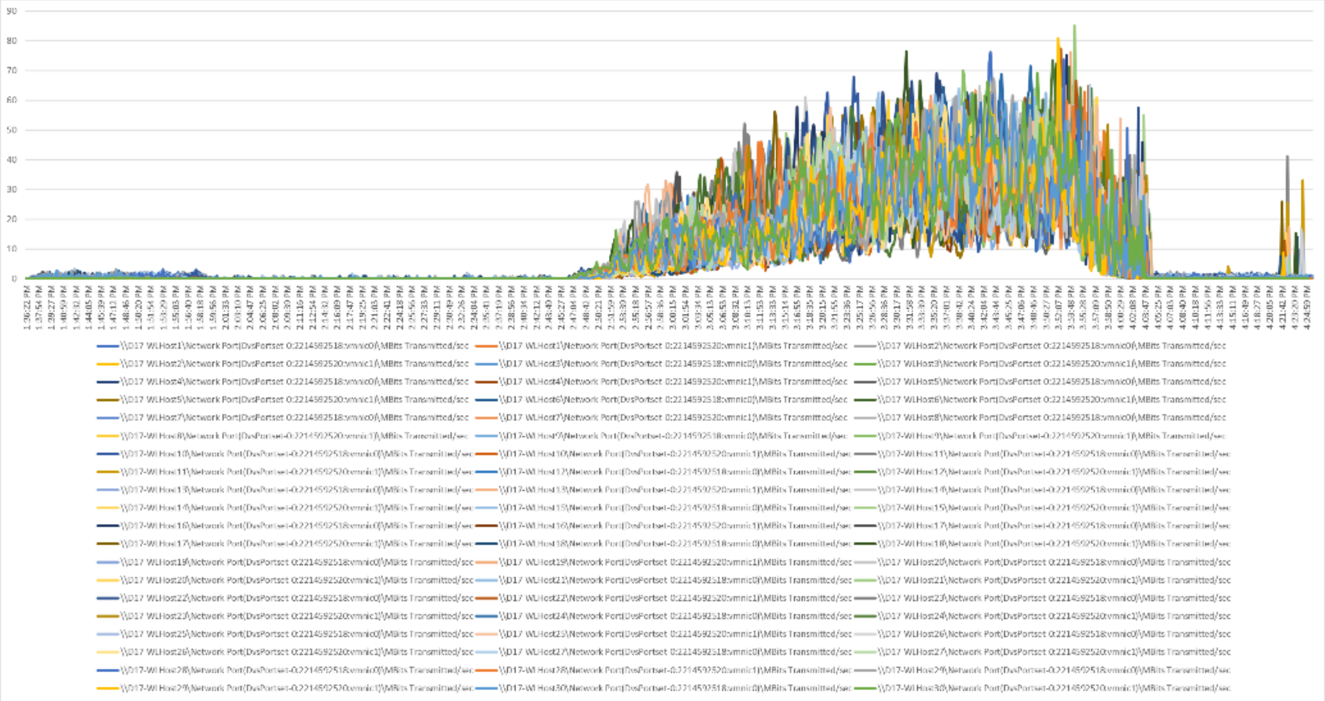

VDI Server Performance Monitor Data for One Sample Test: 5022 Users VDI Persistent (Full Clones) Scale Testing

Figure 129. Full Scale | 5022 Persistent Users| VDI Host | Host CPU Utilization

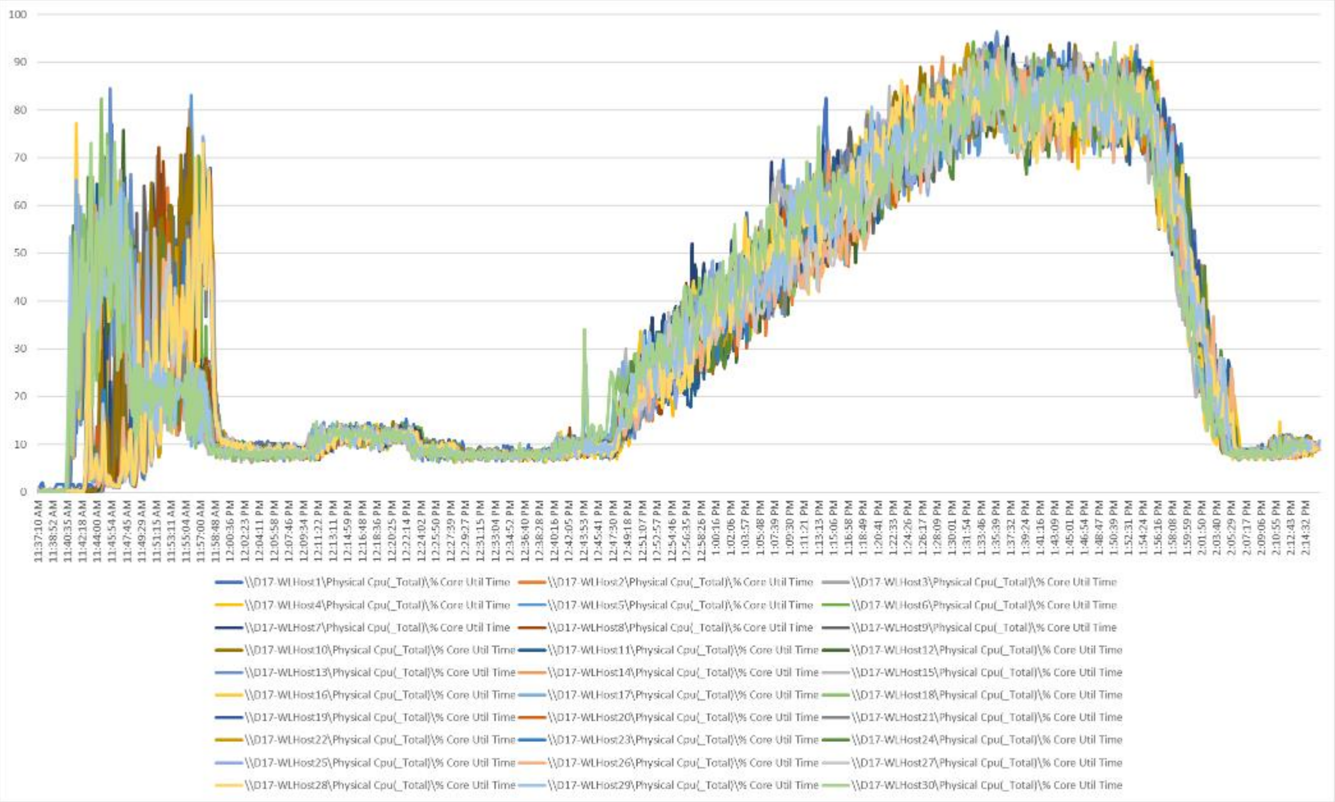

Figure 130. **Full Scale | 5022 Persistent Users| VDI Host | Host Memory Utilization**

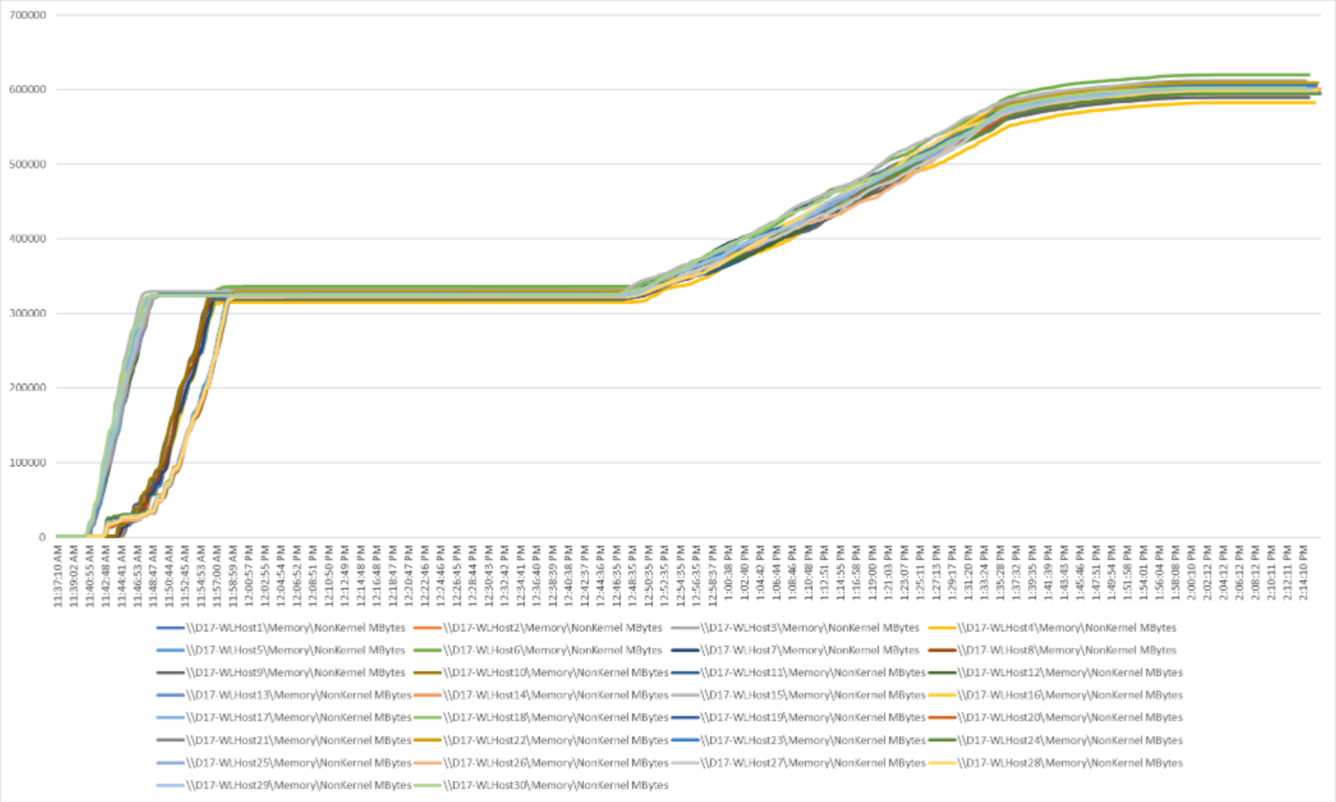

Figure 131. Full Scale | 5022 Persistent Users| VDI Host | Host Network Utilization | Received

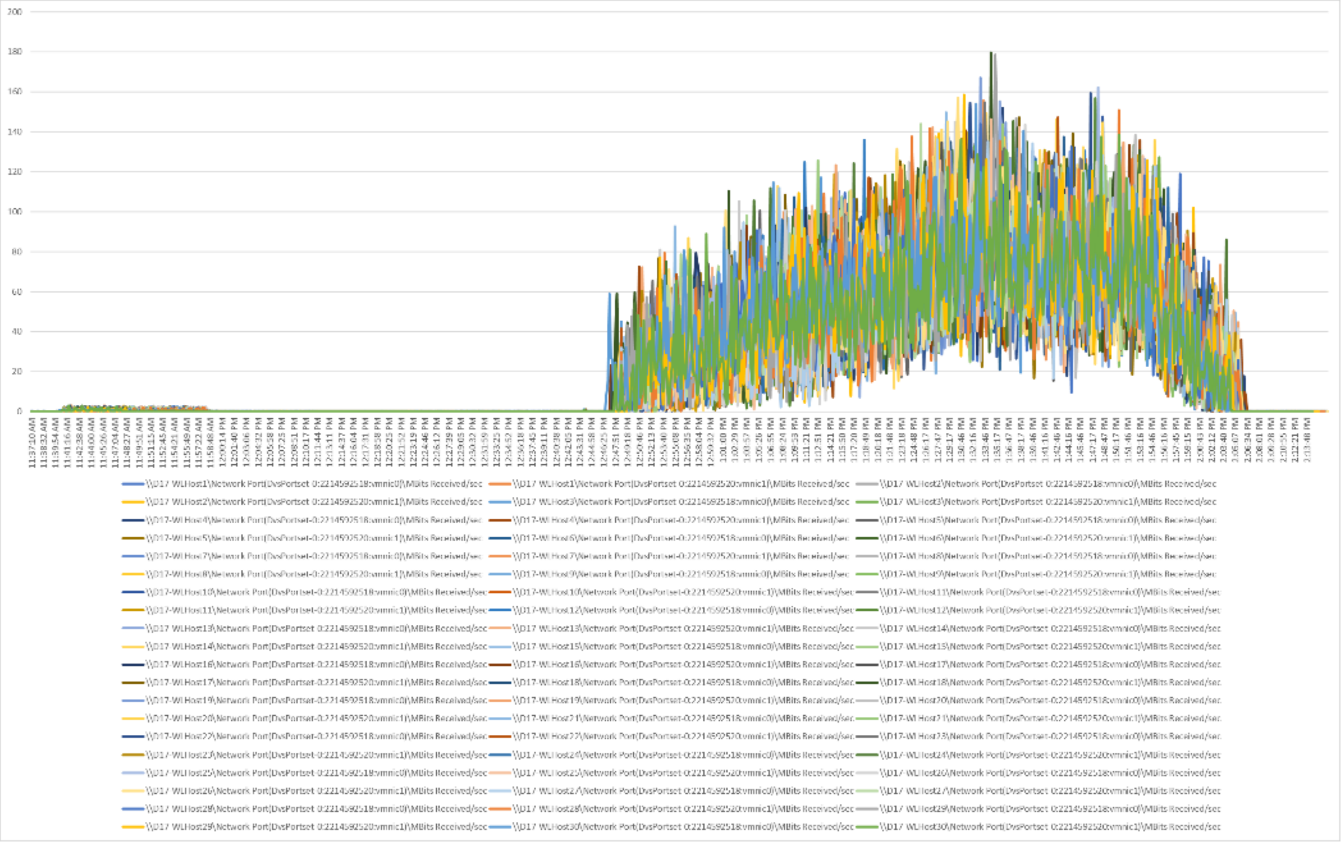

Figure 132. **Full Scale | 5022 Persistent Users| VDI Host | Host Network Utilization | Transmitted**

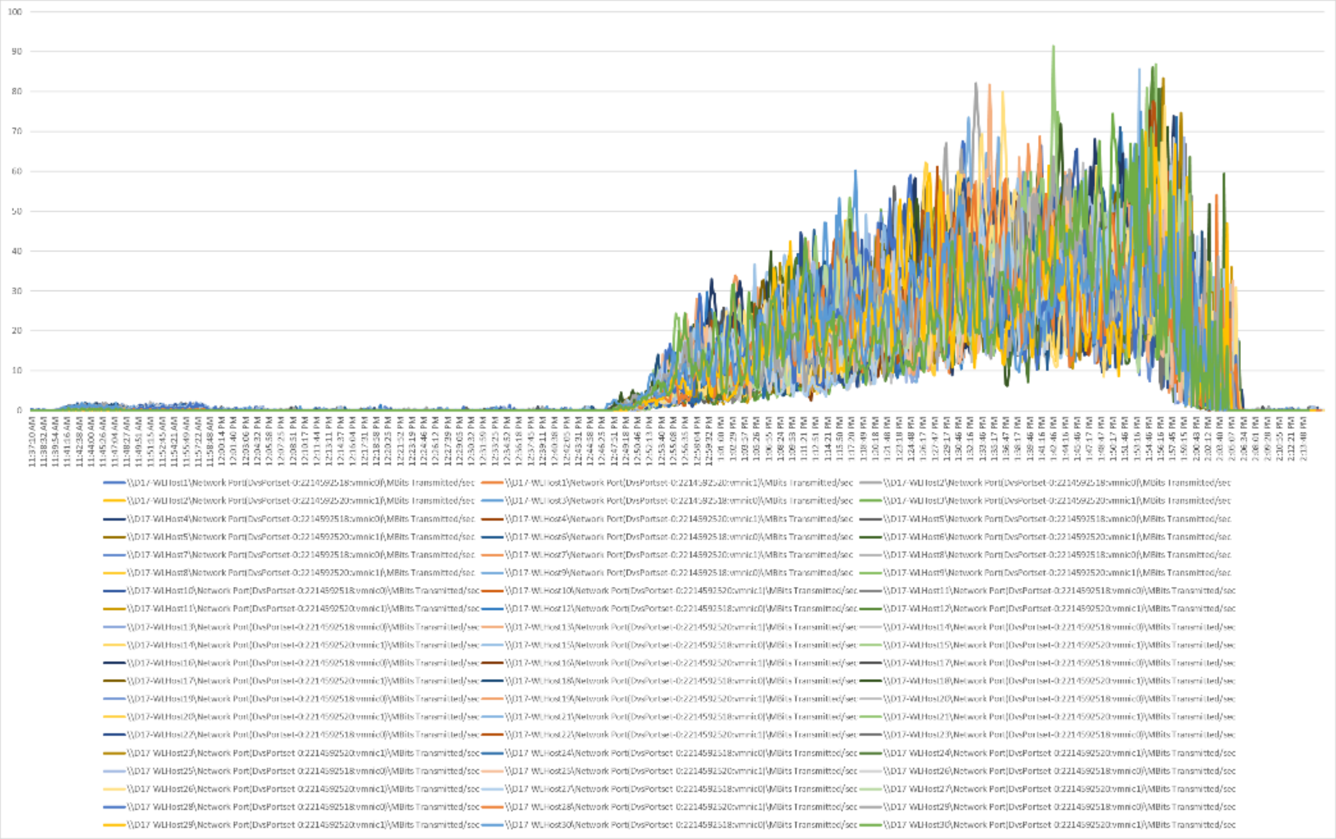

VDI Server Performance Monitor Data for One Sample Test: 6048 Users RDS Scale Testing

Figure 133. Full Scale | 6048 RDS Users| VDI Host | Host CPU Utilization

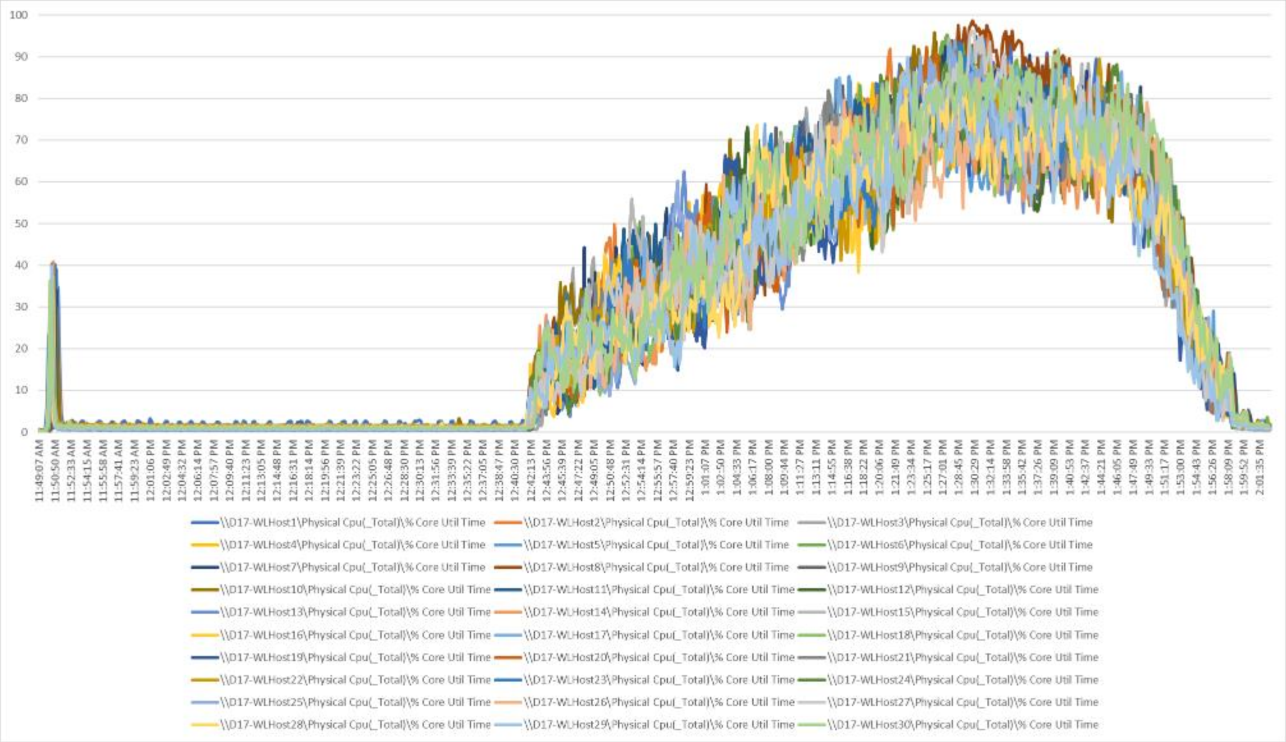

Figure 134. Full Scale | 6048 RDS Users| VDI Host | Host Memory Utilization

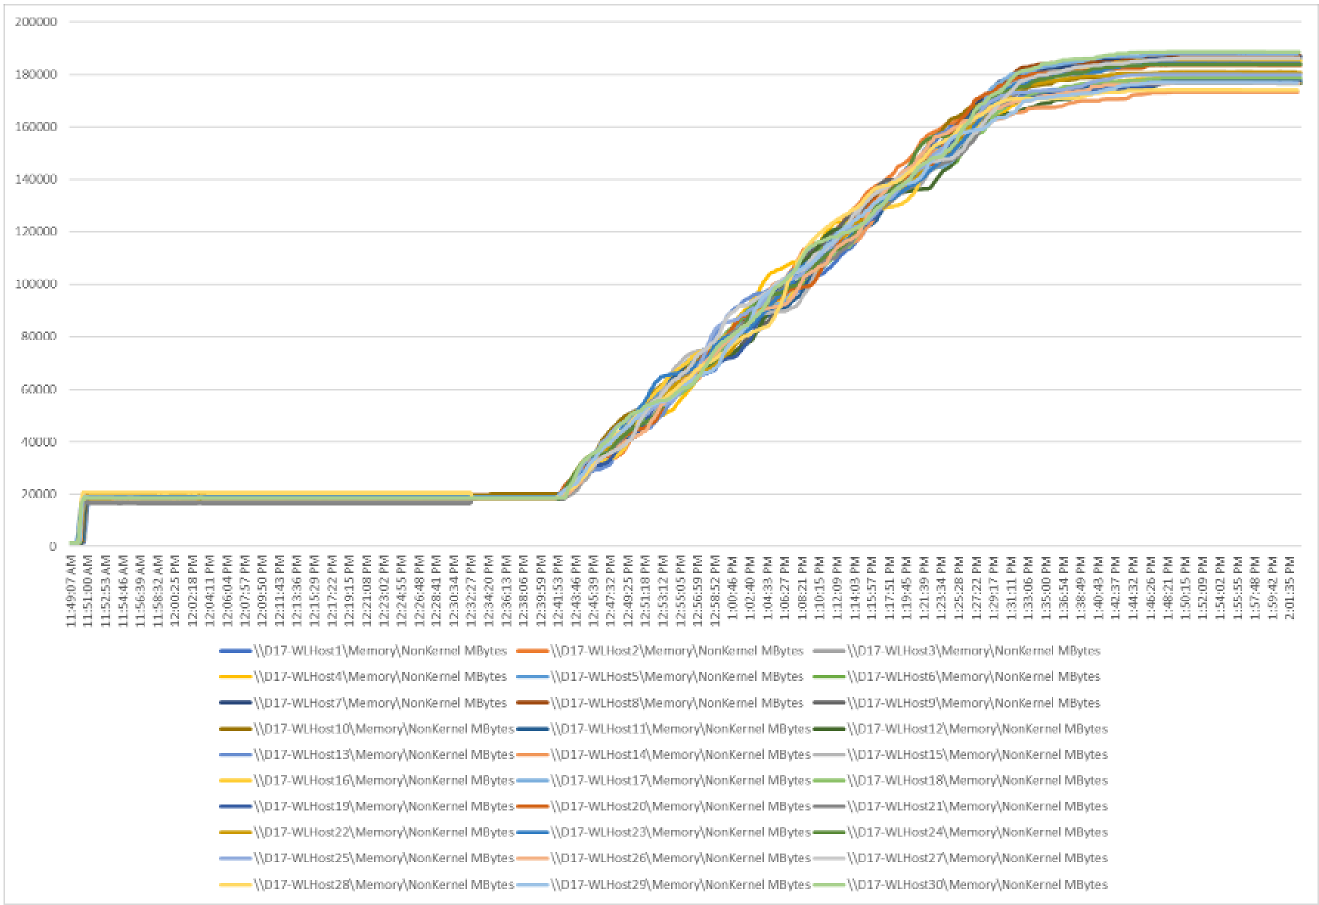

Figure 135. Full Scale | 6048 RDS Users| VDI Host | Host Network Utilization | Received

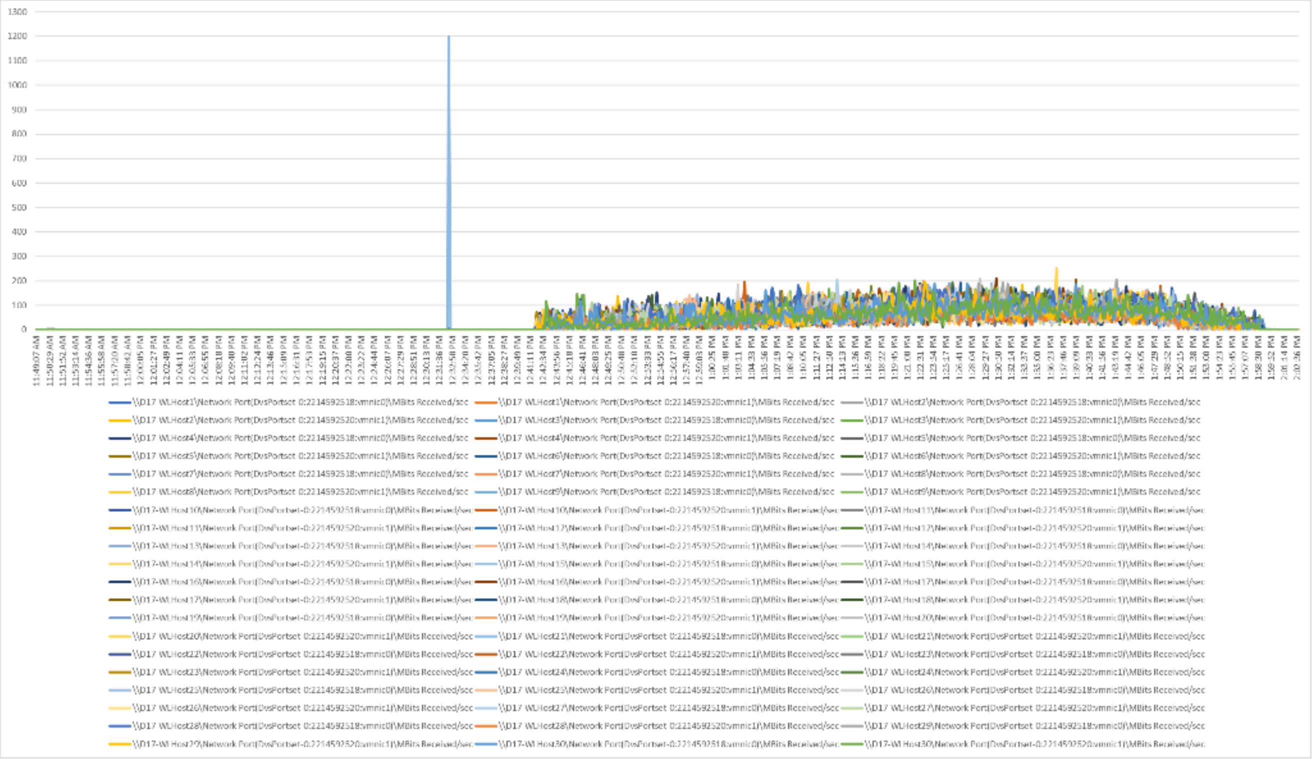

Figure 136. Full Scale | 6048 RDS Users| VDI Host | Host Network Utilization | Transmitted

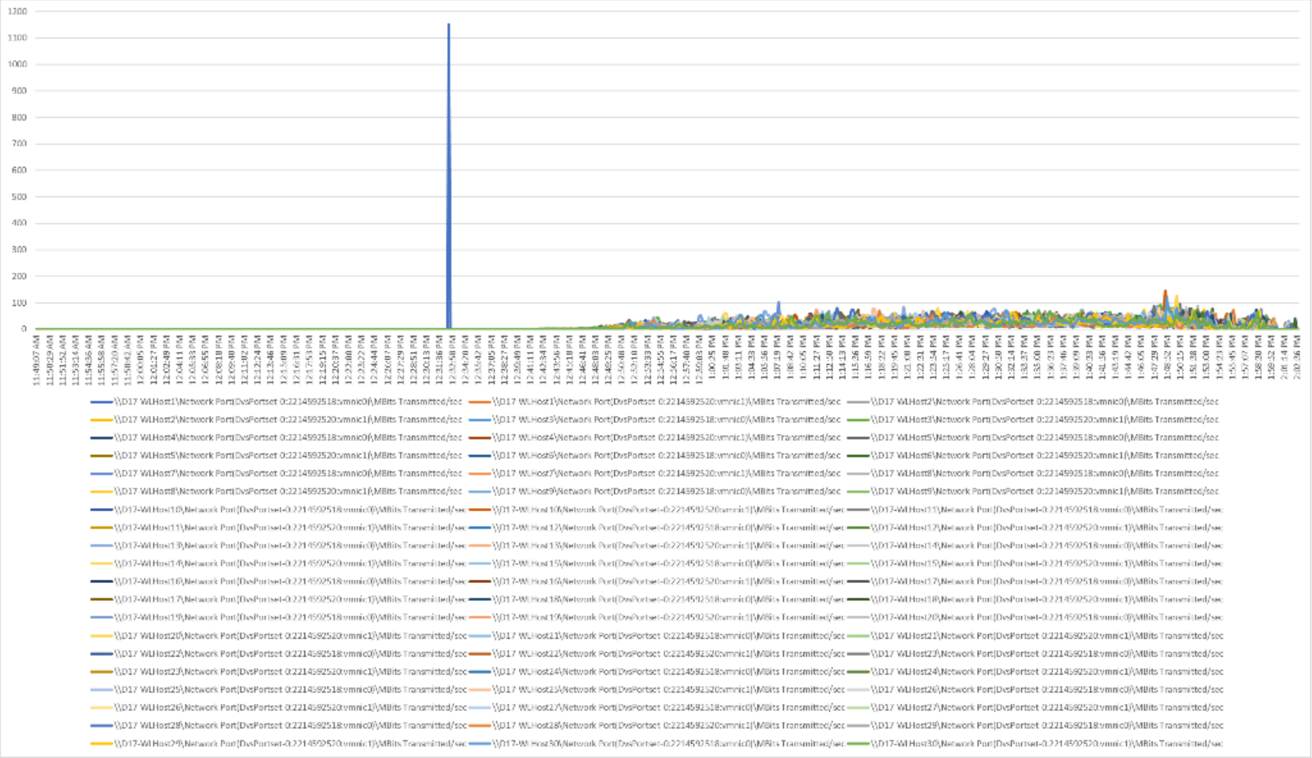

Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community 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**.