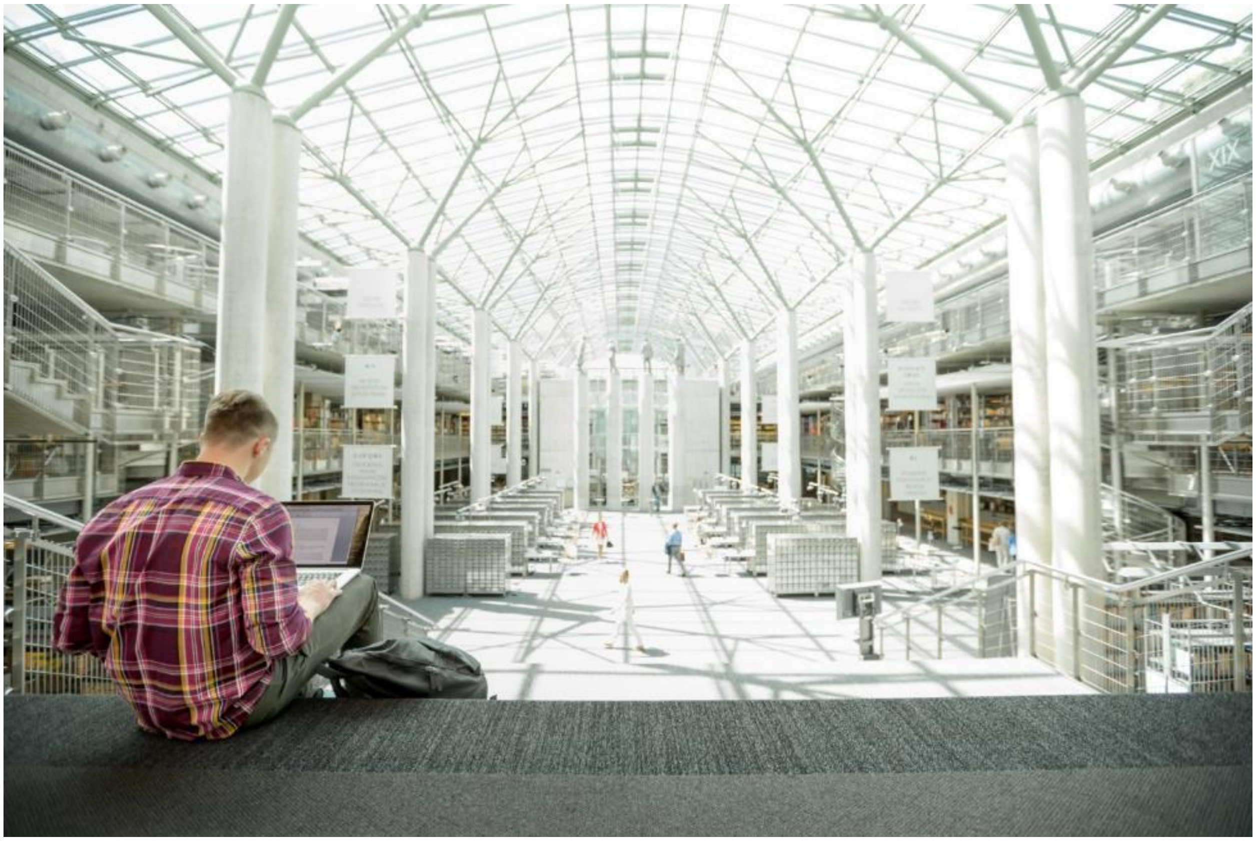

FlexPod Datacenter with Citrix XenDesktop/XenApp 7.15 and VMware vSphere 6.5 Update 1 for 6000 Seats

Cisco Validated Design for a 6000 Seat Virtual Desktop Infrastructure Built on Cisco UCS B200 M5, Cisco UCS Manager 3.2 with NetApp AFF A-Series using Citrix XenDesktop/XenApp 7.15, and

VMware vSphere ESXi 6.5 Update 1 Hypervisor Platform

Last Updated: June 26, 2018

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, visit:

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

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

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

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

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

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

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

This document provides a Reference Architecture for a virtual desktop and application design using Citrix XenApp/XenDesktop 7.15 built on Cisco UCS with a NetApp® All Flash FAS (AFF) A300 storage and the VMware vSphere ESXi 6.5 hypervisor platform.

The landscape of desktop and application virtualization is changing constantly. The new M5 high-performance Cisco UCS Blade Servers and Cisco UCS Unified Fabric combined as part of the FlexPod proven Infrastructure, with the latest generation NetApp AFF storage result in a more compact, more powerful, more reliable and more efficient platform.

This document provides the architecture and design of a virtual desktop infrastructure for up to 6000 mixed use-case users. The solution virtualized on fifth generation Cisco UCS B200 M5 blade servers, booting VMware vSphere 6.5 Update 1 through FC SAN from the AFF A300 storage array. The virtual desktops are powered using Citrix Provisioning Server 7.15 and Citrix XenApp/XenDesktop 7.15, with a mix of RDS hosted shared desktops (1900), pooled/non-persistent hosted virtual Windows 10 desktops (2050) and persistent hosted virtual Windows 10 desktops provisioned with Citrix Machine Creation Services (2050) to support the user population. Where applicable, the document provides best practice recommendations and sizing guidelines for customer deployments of this solution.

The solution is fully capable of supporting hardware accelerated graphicss 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 Citrix XenDesktop.

This solution provides an outstanding virtual desktop end-user experience as measured by the Login VSI 4.1.25.6 Knowledge Worker workload running in benchmark mode, along with the 6000-seat solution 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, NetApp 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 Citrix XenDesktop 7.15 mixed workload solution with NetApp AFF A300, Cisco UCS Blade Servers, Cisco Nexus 9000 series Ethernet switches and Cisco MDS 9000 series fibre channel switches.

What’s New?

This is the first Citrix XenDesktop desktop virtualization Cisco Validated Design with Cisco UCS 5th generation servers and a NetApp AFF A-Series system.

It incorporates the following features:

- Cisco UCS B200 M5 blade servers with Intel Xeon Scalable Family processors and 2666 MHz memory
- Validation of Cisco Nexus 9000 with NetApp AFF A300 system
- Validation of Cisco MDS 9000 with NetApp AFF A300 system
- Support for the Cisco UCS 3.2(3d) release and Cisco UCS B200-M5 servers
- Support for the latest release of NetApp AFF A300 hardware and NetApp ONTAP® 9.3
- VMware vSphere 6.5 U1 Hypervisor
- Citrix XenDesktop 7.15 Server 2016 RDS hosted shared virtual desktops
- Citrix XenDesktop 7.15 non-persistent hosted virtual Windows 10 desktops provisioned with Citrix Provisioning Services
- Citrix XenDesktop 7.15 persistent full clones hosted virtual Windows 10 desktops provisioned with Citrix Machine Creation Services

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

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

The use cases include:

- Enterprise Datacenter
- Service Provider Datacenter
- Large Commercial Datacenter

Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Citrix XenDesktop Microsoft Windows 10 virtual desktops and Citrix XenApp server desktop sessions based on Microsoft Server 2016.

The mixed workload solution includes NetApp AFF A300 storage, Cisco Nexus® and MDS networking, the Cisco Unified Computing System (Cisco UCS®), Citrix XenDesktop 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 access to shared storage using NFS mounts. The reference architecture reinforces the "wire-once" strategy because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

The combination of technologies from Cisco Systems, Inc., NetApp Inc., Citrix 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 the solution include the following:

- **More power, same size**. Cisco UCS B200 M5 half-width blade with dual 18-core 2.3 GHz Intel ® Xeon ® Gold (6140) processors and 768 GB of memory supports more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel 18-core 2.3 GHz Intel ® Xeon ® Gold (6140) 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 simulated login storms.** All 6000 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 datacenter.** The rack space required to support the system is a single 42U rack, conserving valuable data center floor space.
- **All Virtualized**: This CVD presents a validated design that is 100 percent virtualized on VMware ESXi 6.5. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, Citrix XenDesktop components, XenDesktop VDI desktops and XenApp servers were hosted as virtual machines. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the FlexPod converged infrastructure with stateless Cisco UCS Blade servers and NetApp FC storage.
- **Cisco maintains industry leadership** with the new Cisco UCS Manager 3.2(3d) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director ensure that customer environments are consistent locally, across Cisco UCS Domains and across the globe, 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 40G unified fabric story** gets additional validation on Cisco UCS 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.

- **NetApp AFF A300** array provides industry-leading storage solutions that efficiently handle the most demanding I/O bursts (for example, login storms), profile management, and user data management, deliver simple and flexible business continuance, and help reduce storage cost per desktop.
- **NetApp AFF A300** array provides a simple to understand storage architecture for hosting all user data components (VMs, profiles, user data) on the same storage array.
- **NetApp clustered Data ONTAP software** enables to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.
- **Citrix XenDesktop and XenApp Advantage**. XenApp and XenDesktop are virtualization solutions that give IT control of virtual machines, applications, licensing, and security while providing anywhere access for any device.
- XenApp and XenDesktop provides the following:
  - End users to run applications and desktops independently of the device's operating system and interface.
  - Administrators to manage the network and control access from selected devices or from all devices.
  - Administrators to manage an entire network from a single data center.

  XenApp and XenDesktop share a unified architecture called FlexCast Management Architecture (FMA). FMA's key features are the ability to run multiple versions of XenApp or XenDesktop from a single Site and integrated provisioning.
- **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 XenApp 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**. Citrix provides two core provisioning methods for XenDesktop and XenApp virtual machines: Citrix Provisioning Services for pooled virtual desktops and XenApp virtual servers and Citrix Machine Creation Services for pooled or persistent virtual desktops. This paper provides guidance on how to use each method and documents the performance of each technology.

Cisco Desktop Virtualization Solutions: Datacenter

The Evolving Workplace

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

This workforce is also increasingly mobile, conducting business in traditional offices, conference rooms across the enterprise campus, home offices, on the road, in hotels, and at the local coffee shop. This workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference. These trends are increasing pressure on IT to ensure the 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 2016.

Figure 1 Cisco Data Center Partner Collaboration

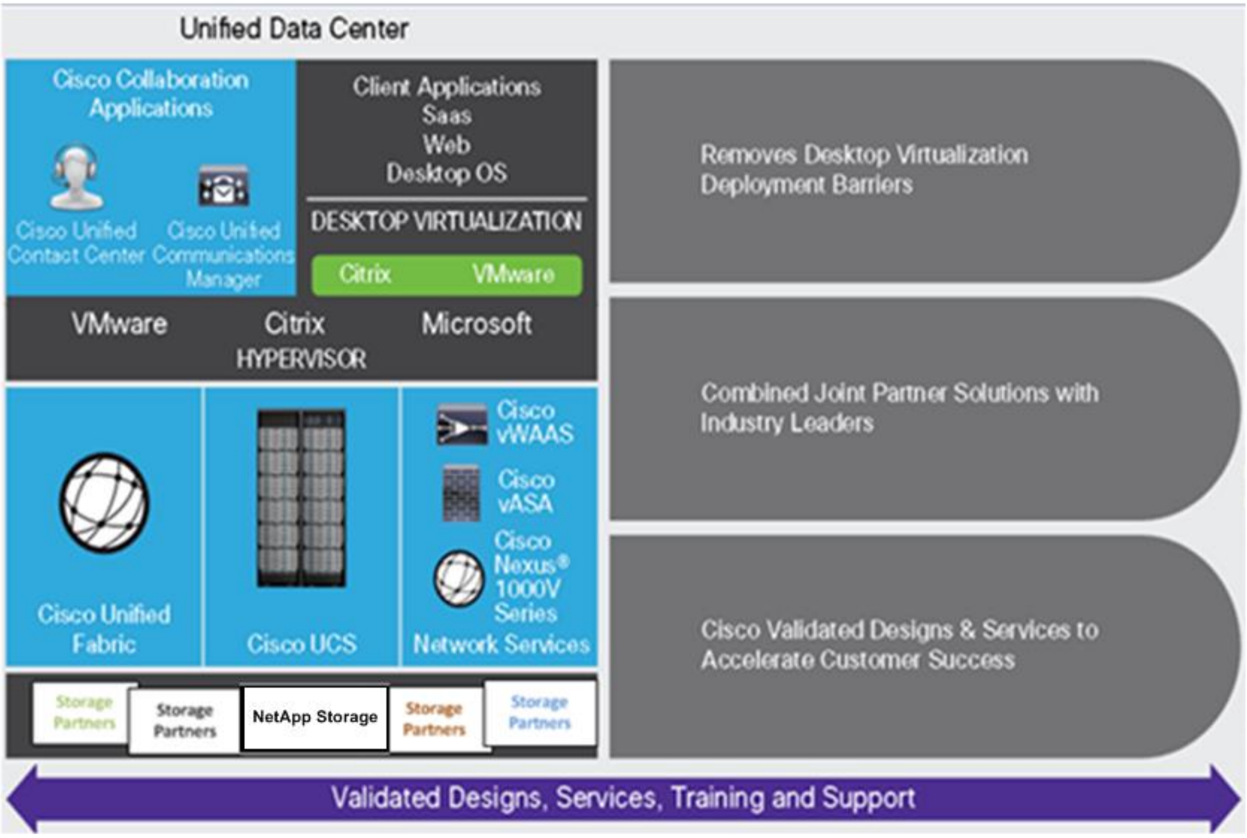
Unified Data Center
Cisco Collaboration Applications
Cisco Unified Contact Center
Cisco Unified Communications Manager
Client Applications
Saas
Web
Desktop OS
DESKTOP VIRTUALIZATION
Citrix
VMware
VMware
Citrix
Microsoft
HYPERVISOR
Cisco Unified Fabric
Cisco UCS
Cisco vWAAS
Cisco vASA
Cisco Nexus® 1000V Series
Network Services
Storage Partners
Storage Partners
NetApp 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 reprovision 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 (UCSM) 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.

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

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

Scalable

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

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 1 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partners NetApp help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs for End User Computing based on FlexPod solutions have demonstrated scalability and performance, with up to 6000 desktops up and running in less than 30 minutes.

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

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

Figure 2 illustrates the physical architecture.

Figure 2 Physical Architecture

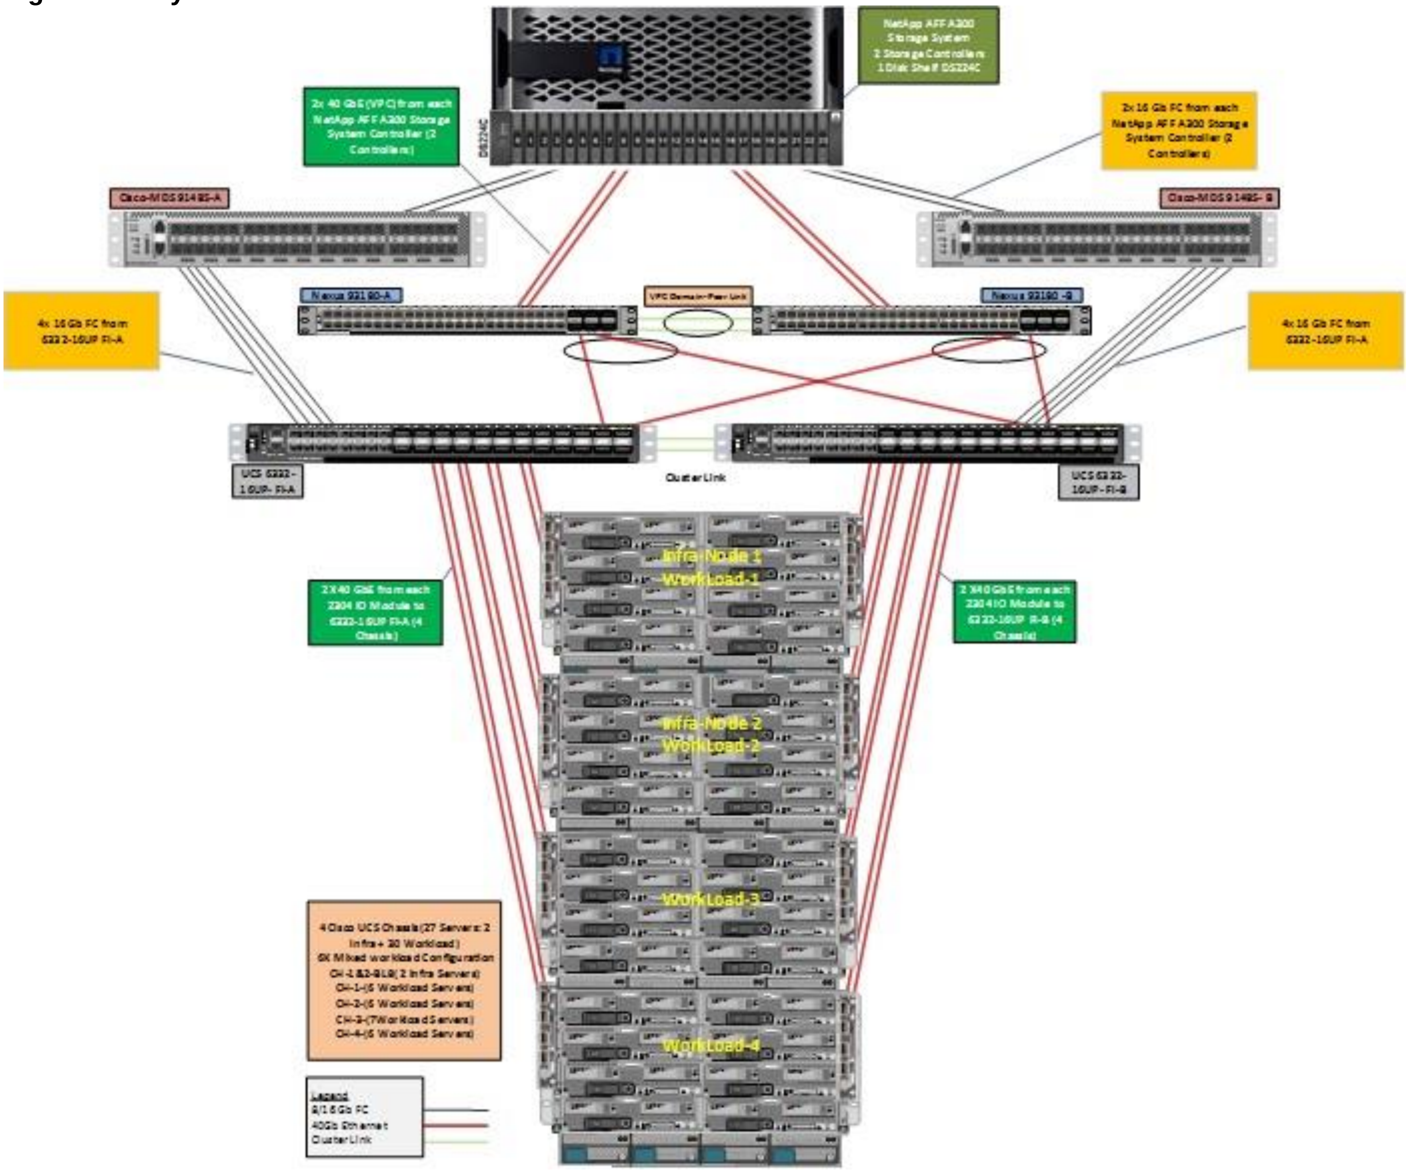
NetApp AFF A300 Storage System 2 Storage Controllers 1 Disk Shelf DS224C
DS224C
2x 40 GbE (VPC) from each NetApp AFF A300 Storage System Controller (2 Controllers)
2x 16 Gb FC from each NetApp AFF A300 Storage System Controller (2 Controllers)
Cisco-MDS 9148S-A
Cisco-MDS 9148S-B
Nexus 93180-A
VPC Domain-Peer Link
Nexus 93180-B
4x 16 Gb FC from 6332-16UP FI-A
4x 16 Gb FC from 6332-16UP FI-A
UCS 6332-16UP-FI-A
Cluster Link
UCS 6332-16UP-FI-B
Infra-Node 1
Workload-1
2X40 GbE from each 2304 IO Module to 6332-16UP FI-A (4 Chassis)
2X40 GbE from each 2304 IO Module to 6332-16UP FI-B (4 Chassis)
Infra-Node 2
Workload-2
Workload-3
Workload-4
4 Cisco UCS Chassis (27 Servers: 2 Infra + 30 Workload)
6K Mixed workload Configuration
CH-1 &2-BL8 (2 Infra Servers)
CH-1-(6 Workload Servers)
CH-2-(6 Workload Servers)
CH-3-(7 Workload Servers)
CH-4-(6 Workload Servers)
Legend
8/16 Gb FC
40Gb Ethernet
Cluster Link

The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco MDS 9148S 16GB Fibre Channel switches
- Two Cisco UCS 6332-16UP Fabric Interconnects
- Four Cisco UCS 5108 Blade Chassis
- Two Cisco UCS B200 M4 Blade Servers (2 Infra Server hosting Infrastructure VMs)
- 30 Cisco UCS B200 M5 Blade Servers (for workload)
- One NetApp AFF A300 Storage System

- One NetApp DS224C Disk Shelf

For desktop virtualization, the deployment includes Citrix XenDesktop 7.15 running on VMware vSphere 6.5.

The design is intended to provide a large-scale building block for XenDesktop mixed workloads consisting of RDS Windows Server 2016 hosted shared desktop sessions and Windows 10 non-persistent and persistent hosted desktops in the following ratio:

- 1900 Random Hosted Shared Windows 2016 user sessions with office 2016 (PVS)
- 2050 Random Pooled Windows 10 Desktops with office 2016 (PVS)
- 2050 Static Full Copy Windows 10 Desktops with office 2016 (MCS)

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

Configuration Guidelines

This Cisco Validated Design provides details for deploying a fully redundant, highly available 6000 seats mixed workload virtual desktop solution with VMware on a FlexPod Datacenter architecture. Configuration guidelines are provided that refer the reader to which redundant component is being configured with each step. For example, storage controller 01and storage controller 02 are used to identify the two AFF A300 storage controllers that are provisioned with this document, Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured and Cisco MDS A or Cisco MDS B identifies the pair of Cisco MDS switches that are configured.

The Cisco UCS 6332-16UP Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, VM-Host-RDSH-01, VM-Host-VDI-01 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

Solution Components

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

What is FlexPod?

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

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

Figure 3 FlexPod Component Families

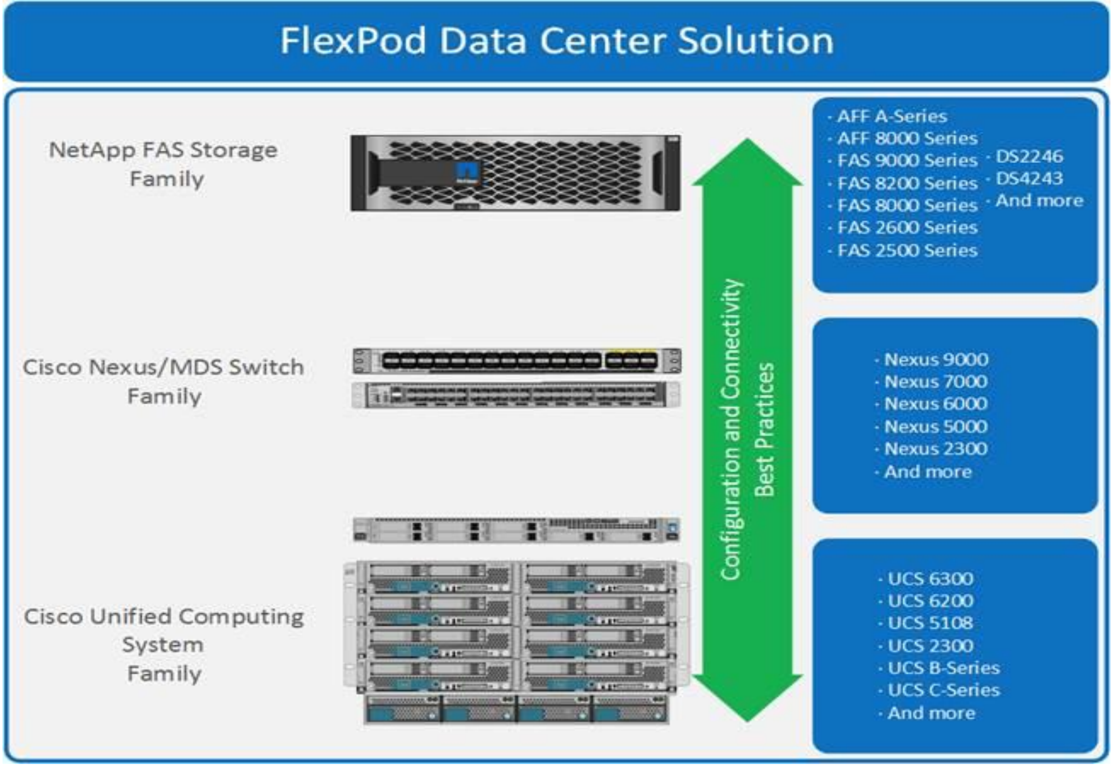

These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually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

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

Why FlexPod?

The following lists the benefits of FlexPod:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100 percent flash storage
  - Consolidate 100’s of enterprise-class applications in a single rack
  - Scales easily, without disruption
  - Continuous growth through multiple FlexPod CI deployments
- Operational Simplicity
  - Fully tested, validated, and documented for rapid deployment
  - Reduced management complexity
  - Auto-aligned 512B architecture removes storage alignment issues
  - No storage tuning or tiers necessary
- Lowest TCO
  - Dramatic savings in power, cooling, and space with 100 percent Flash
  - Industry leading data reduction
- Enterprise-Grade Resiliency
  - Highly available architecture with no single point of failure
  - Nondisruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication
  - Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases

Cisco Unified Computing System

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) through an intuitive GUI, a command-line interface (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® processor E5-2600/4600 v3 and E7-2800 v3 family CPUs.

- **Network**: The system is integrated on a low-latency, lossless, 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks 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 a choice for storage access and investment protection. In addition, server administrators can preassign 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. The manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

Figure 4 Cisco Data Center Overview

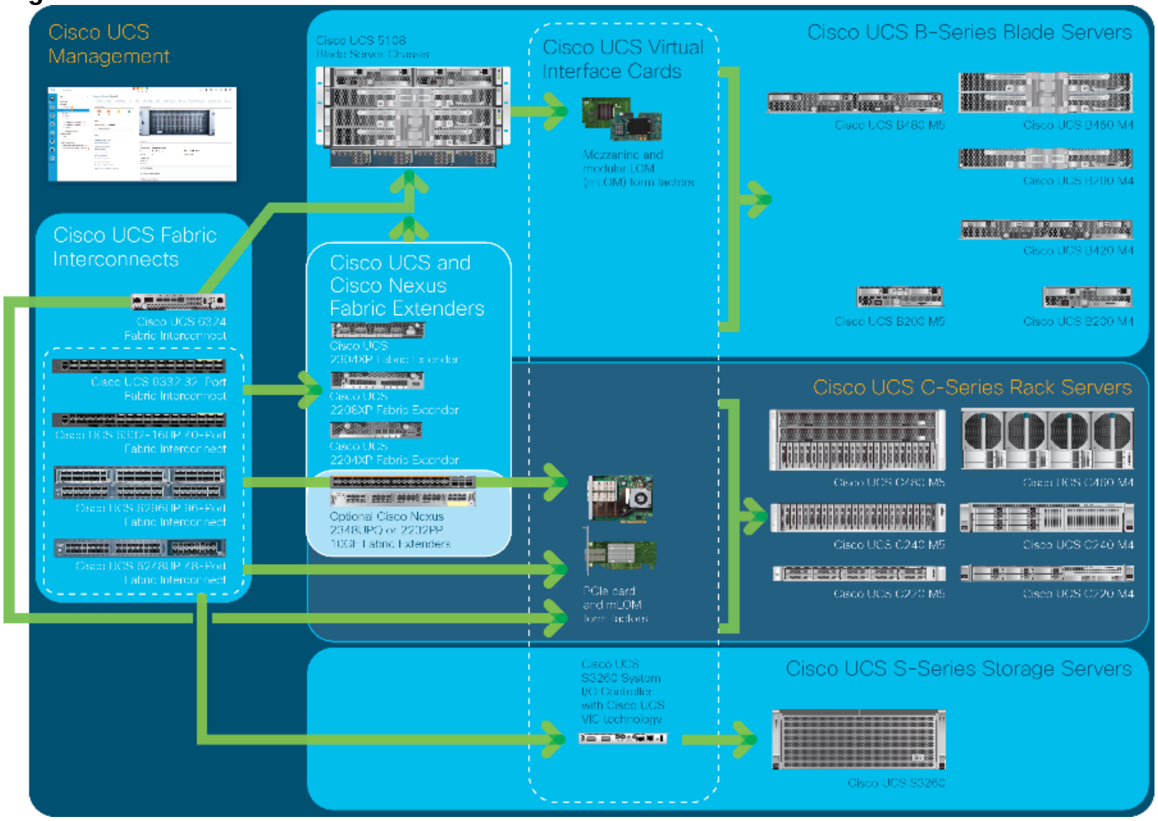
Cisco UCS Management
Cisco UCS 5108 Blade Server Chassis
Cisco UCS Virtual Interface Cards
Mezzanine and modular LOM (mLOM) form factors
Cisco UCS B-Series Blade Servers
Cisco UCS B480 M5
Cisco UCS B460 M4
Cisco UCS B260 M4
Cisco UCS B420 M4
Cisco UCS B200 M5
Cisco UCS B200 M4
Cisco UCS Fabric Interconnects
Cisco UCS 6324 Fabric Interconnect
Cisco UCS 6332 32-Port Fabric Interconnect
Cisco UCS 6332-16UP 40-Port Fabric Interconnect
Cisco UCS 6296UP 96-Port Fabric Interconnect
Cisco UCS 6248UP 48-Port Fabric Interconnect
Cisco UCS and Cisco Nexus Fabric Extenders
Cisco UCS 2304XP Fabric Extender
Cisco UCS 2208XP Fabric Extender
Cisco UCS 2204XP Fabric Extender
Optional Cisco Nexus 2348UPQ or 2232PP 10GE Fabric Extenders
PCIe card and mLOM form factors
Cisco UCS C-Series Rack Servers
Cisco UCS C480 M5
Cisco UCS C460 M4
Cisco UCS C240 M5
Cisco UCS C240 M4
Cisco UCS C220 M5
Cisco UCS C220 M4
Cisco UCS S3260 System I/O Controller with Cisco UCS VIC technology
Cisco UCS S-Series Storage Servers
Cisco UCS S3260

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 command-line interface (CLI), or an XML API for comprehensive access to all Cisco UCS Manager Functions.

Cisco UCS Fabric Interconnect

The Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 40 Gigabit Ethernet, FCoE, and Fibre Channel functions.

The fabric interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all blades in the domain.

For networking, the Cisco UCS 6300 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 40 Gigabit Ethernet on all ports, 2.4 plus terabit (Tb) switching capacity, and 320 Gbps of bandwidth per chassis IOM, independent of packet size and enabled services. The product series supports Cisco low-latency, lossless, 40 Gigabit Ethernet unified network fabric capabilities, increasing the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnects support multiple traffic classes over a lossless Ethernet fabric, from the blade server through the interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Figure 5 Cisco UCS 6300 Series Fabric Interconnect

Cisco UCS 5108 Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis. The Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+1 redundant, and grid redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS Fabric Extenders. A passive mid-plane provides up to 40 Gbps of I/O bandwidth per server slot from each Fabric Extender. The chassis is capable of supporting 40 Gigabit Ethernet standards.

Figure 6 Cisco UCS 5108 Blade Chassis Front and Rear Views

Cisco UCS B200 M5 Blade Server

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

Figure 7 Cisco UCS B200 M5 Front View

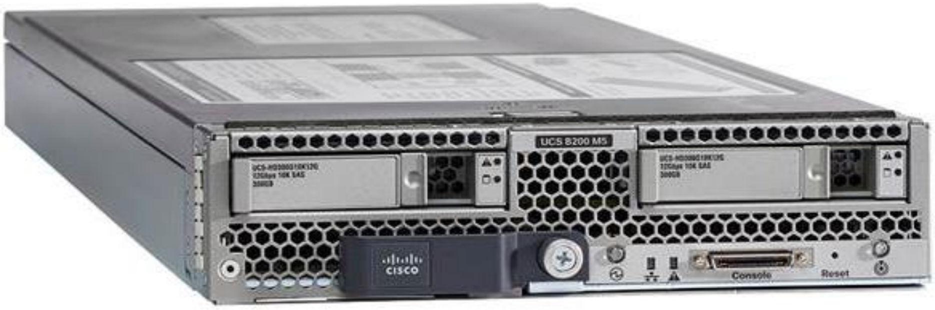

Figure 8 Cisco UCS B200 M5 Back View

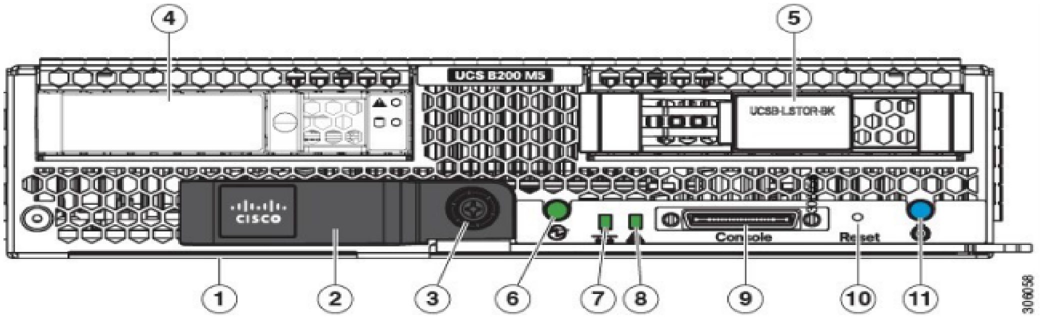

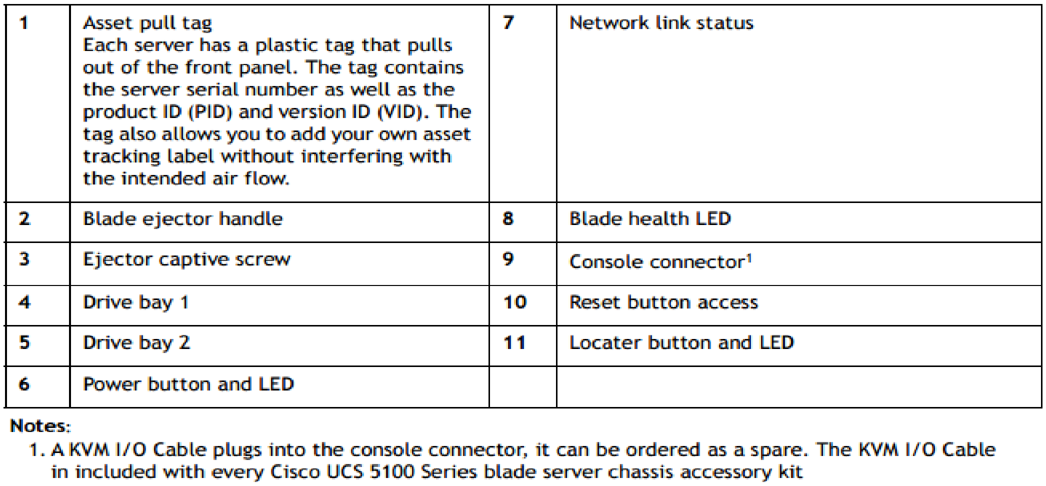

1	Asset pull tag Each server has a plastic tag that pulls out of the front panel. The tag contains the server serial number as well as the product ID (PID) and version ID (VID). The tag also allows you to add your own asset tracking label without interfering with the intended air flow.	7	Network link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

Notes:

1. A KVM I/O Cable plugs into the console connector, it can be ordered as a spare. The KVM I/O Cable 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 6140 Gold product family, it offers up to 3 TB of memory using 128GB DIMMs, up to two disk drives, and up to 320 Gbps of I/O throughput. The Cisco UCS B200 M5 offers exceptional levels of performance, flexibility, and I/O throughput to run your most demanding applications.

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

Product Overview

The Cisco UCS B200 M5 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA. The Cisco UCS B200 M5 server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager software and simplified server access through Cisco SingleConnect technology. It includes the following:

- Latest Intel® Xeon® Scalable processors with up to 28 cores per socket
- Up to 24 DDR4 DIMMs for improved performance
- Intel 3D XPoint-ready support, with built-in support for next-generation nonvolatile memory technology
- Two GPUs
- Two Small-Form-Factor (SFF) drives
- Two Secure Digital (SD) cards or M.2 SATA drives
- Up to 80 Gbps of I/O throughput

Main Features

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

The Cisco UCS B200 M5 provides these main features:

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

Table 1 **Product Specifications**

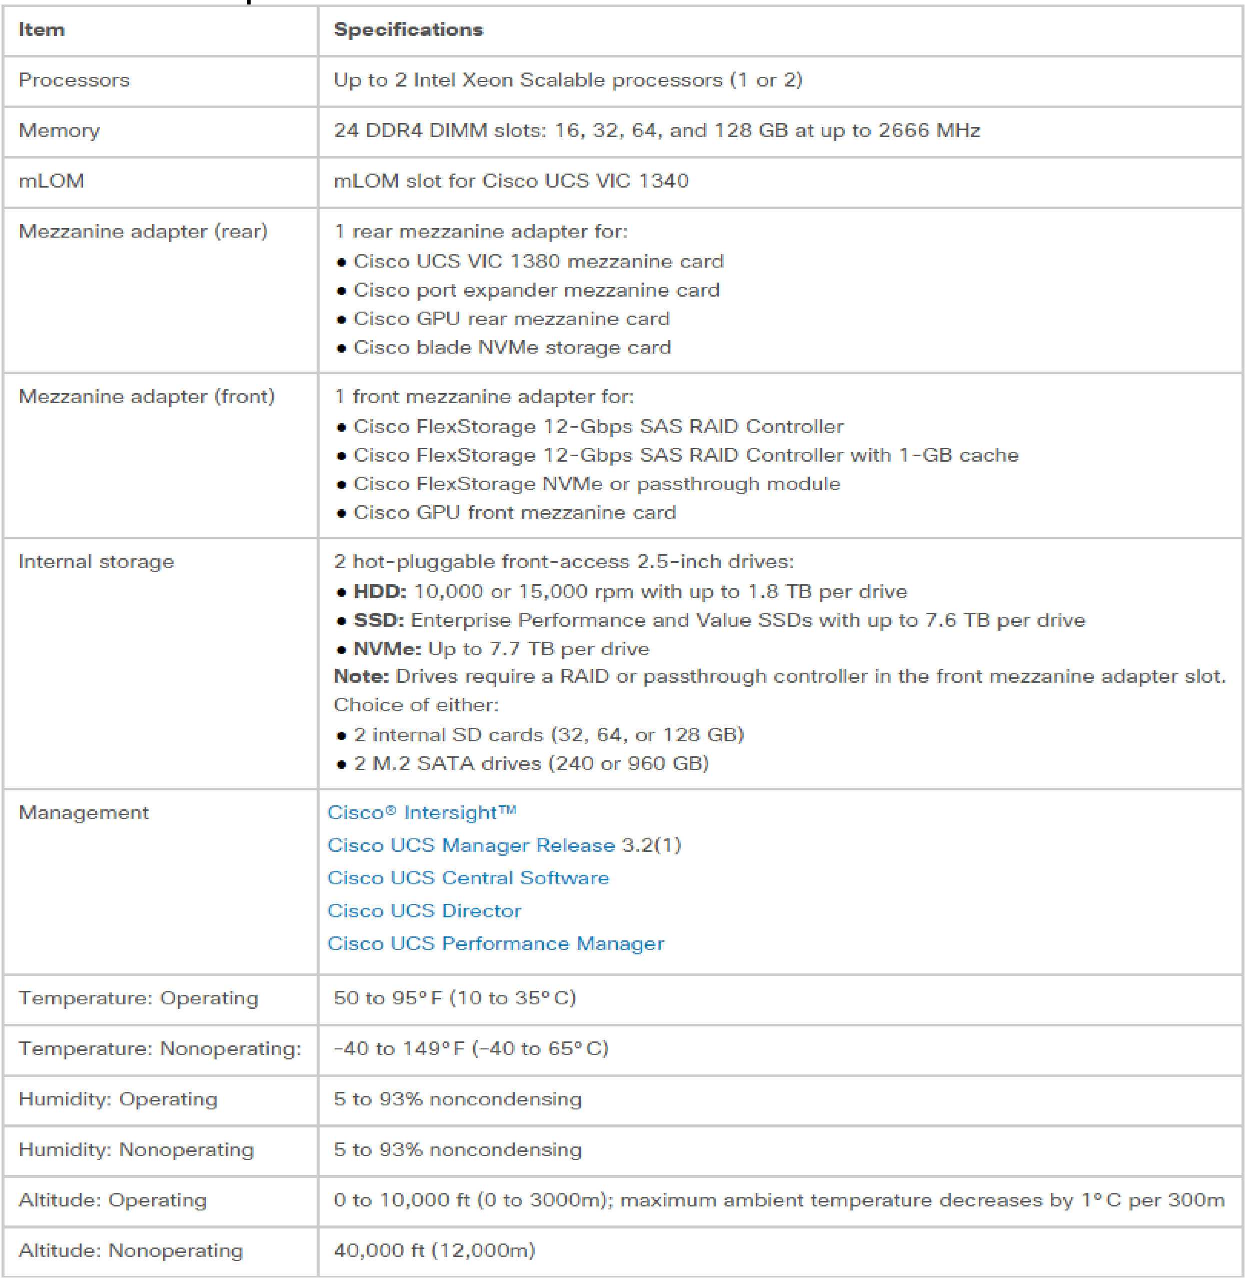

Item	Specifications
Processors	Up to 2 Intel Xeon Scalable processors (1 or 2)
Memory	24 DDR4 DIMM slots: 16, 32, 64, and 128 GB at up to 2666 MHz
mLOM	mLOM slot for Cisco UCS VIC 1340
Mezzanine adapter (rear)	1 rear mezzanine adapter for: • Cisco UCS VIC 1380 mezzanine card • Cisco port expander mezzanine card • Cisco GPU rear mezzanine card • Cisco blade NVMe storage card
Mezzanine adapter (front)	1 front mezzanine adapter for: • Cisco FlexStorage 12-Gbps SAS RAID Controller • Cisco FlexStorage 12-Gbps SAS RAID Controller with 1-GB cache • Cisco FlexStorage NVMe or passthrough module • Cisco GPU front mezzanine card
Internal storage	2 hot-pluggable front-access 2.5-inch drives: • **HDD:** 10,000 or 15,000 rpm with up to 1.8 TB per drive • **SSD:** Enterprise Performance and Value SSDs with up to 7.6 TB per drive • **NVMe:** Up to 7.7 TB per drive **Note:** Drives require a RAID or passthrough controller in the front mezzanine adapter slot. Choice of either: • 2 internal SD cards (32, 64, or 128 GB) • 2 M.2 SATA drives (240 or 960 GB)
Management	Cisco® Intersight™ Cisco UCS Manager Release 3.2(1) Cisco UCS Central Software Cisco UCS Director Cisco UCS Performance Manager
Temperature: Operating	50 to 95° F (10 to 35° C)
Temperature: Nonoperating:	-40 to 149° F (-40 to 65° C)
Humidity: Operating	5 to 93% noncondensing
Humidity: Nonoperating	5 to 93% noncondensing
Altitude: Operating	0 to 10,000 ft (0 to 3000m); maximum ambient temperature decreases by 1° C per 300m
Altitude: Nonoperating	40,000 ft (12,000m)

Table 2 **System Requirements**

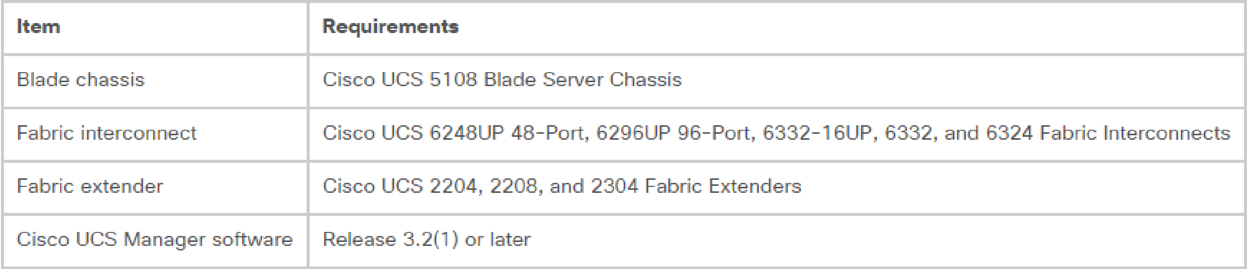

Item	Requirements
Blade chassis	Cisco UCS 5108 Blade Server Chassis
Fabric interconnect	Cisco UCS 6248UP 48-Port, 6296UP 96-Port, 6332-16UP, 6332, and 6324 Fabric Interconnects
Fabric extender	Cisco UCS 2204, 2208, and 2304 Fabric Extenders
Cisco UCS Manager software	Release 3.2(1) or later

Table 3 **Ordering Information**

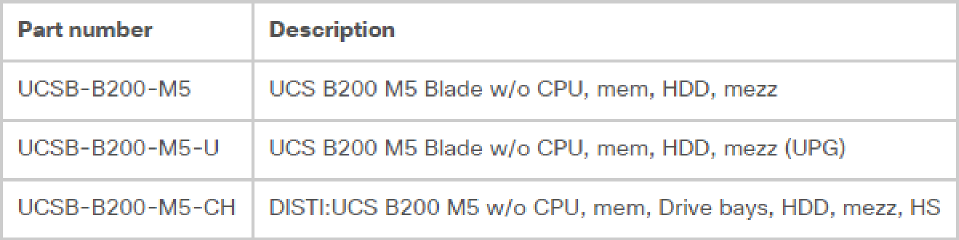

Part number	Description
UCSB-B200-M5	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz
UCSB-B200-M5-U	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz (UPG)
UCSB-B200-M5-CH	DISTI:UCS B200 M5 w/o CPU, mem, Drive bays, HDD, mezz, HS

Table 4 **Capabilities and Features**

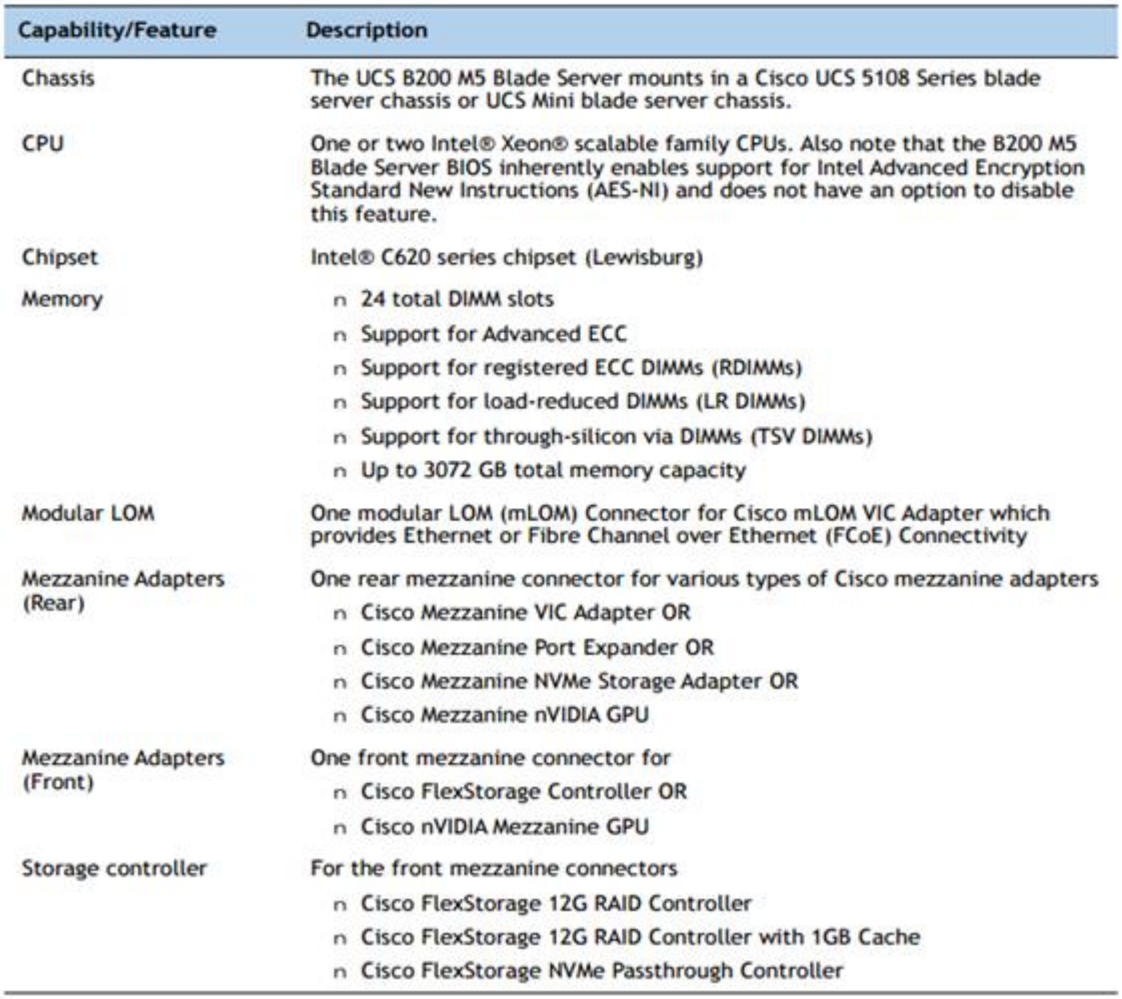

Capability/Feature	Description
Chassis	The UCS B200 M5 Blade Server mounts in a Cisco UCS 5108 Series blade server chassis or UCS Mini blade server chassis.
CPU	One or two Intel® Xeon® scalable family CPUs. Also note that the B200 M5 Blade Server BIOS inherently enables support for Intel Advanced Encryption Standard New Instructions (AES-NI) and does not have an option to disable this feature.
Chipset	Intel® C620 series chipset (Lewisburg)
Memory	n 24 total DIMM slots n Support for Advanced ECC n Support for registered ECC DIMMs (RDIMMs) n Support for load-reduced DIMMs (LR DIMMs) n Support for through-silicon via DIMMs (TSV DIMMs) n Up to 3072 GB total memory capacity
Modular LOM	One modular LOM (mLOM) Connector for Cisco mLOM VIC Adapter which provides Ethernet or Fibre Channel over Ethernet (FCoE) Connectivity
Mezzanine Adapters (Rear)	One rear mezzanine connector for various types of Cisco mezzanine adapters n Cisco Mezzanine VIC Adapter OR n Cisco Mezzanine Port Expander OR n Cisco Mezzanine NVMe Storage Adapter OR n Cisco Mezzanine nVIDIA GPU
Mezzanine Adapters (Front)	One front mezzanine connector for n Cisco FlexStorage Controller OR n Cisco nVIDIA Mezzanine GPU
Storage controller	For the front mezzanine connectors n Cisco FlexStorage 12G RAID Controller n Cisco FlexStorage 12G RAID Controller with 1GB Cache n Cisco FlexStorage NVMe Passthrough Controller

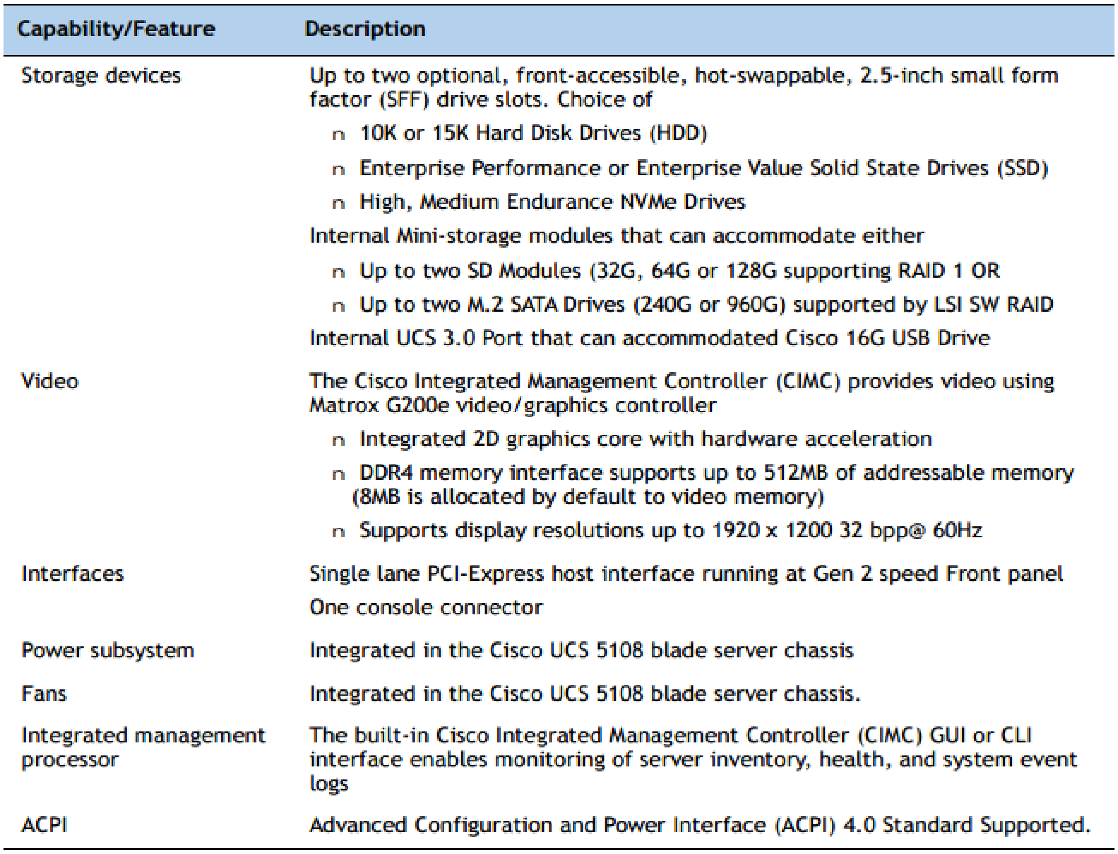

Capability/Feature	Description
Storage devices	Up to two optional, front-accessible, hot-swappable, 2.5-inch small form factor (SFF) drive slots. Choice of n 10K or 15K Hard Disk Drives (HDD) n Enterprise Performance or Enterprise Value Solid State Drives (SSD) n High, Medium Endurance NVMe Drives Internal Mini-storage modules that can accommodate either n Up to two SD Modules (32G, 64G or 128G supporting RAID 1 OR n Up to two M.2 SATA Drives (240G or 960G) supported by LSI SW RAID Internal UCS 3.0 Port that can accommodated Cisco 16G USB Drive
Video	The Cisco Integrated Management Controller (CIMC) provides video using Matrox G200e video/graphics controller n Integrated 2D graphics core with hardware acceleration n DDR4 memory interface supports up to 512MB of addressable memory (8MB is allocated by default to video memory) n Supports display resolutions up to 1920 x 1200 32 bpp@ 60Hz
Interfaces	Single lane PCI-Express host interface running at Gen 2 speed Front panel One console connector
Power subsystem	Integrated in the Cisco UCS 5108 blade server chassis
Fans	Integrated in the Cisco UCS 5108 blade server chassis.
Integrated management processor	The built-in Cisco Integrated Management Controller (CIMC) GUI or CLI interface enables monitoring of server inventory, health, and system event logs
ACPI	Advanced Configuration and Power Interface (ACPI) 4.0 Standard Supported.

For detailed information, refer to the Cisco UCS B200 M5 Blade Server Spec Sheet and the Cisco UCS B200 M5 Blade Server Data Sheet.

Cisco UCS VIC1340 Converged Network Adapter

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

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

Figure 9 Cisco UCS VIC 1340

Figure 10 illustrates the Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M5 Blade Servers.

Figure 10 Cisco UCS VIC 1340 Deployed in the Cisco UCS B200 M5

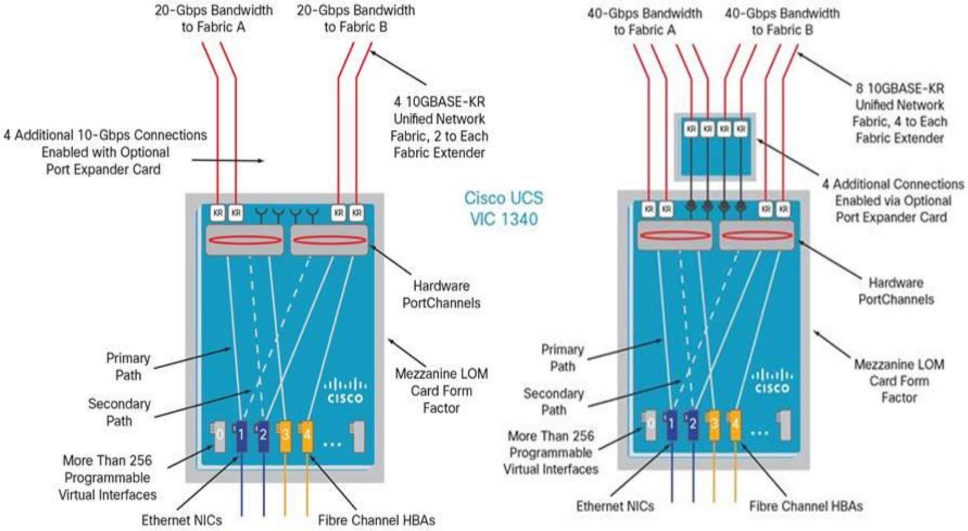
20-Gbps Bandwidth to Fabric A
20-Gbps Bandwidth to Fabric B
4 10GBASE-KR Unified Network Fabric, 2 to Each Fabric Extender
4 Additional 10-Gbps Connections Enabled with Optional Port Expander Card
KR
Cisco UCS VIC 1340
Hardware PortChannels
Primary Path
Secondary Path
Mezzanine LOM Card Form Factor
cisco
0
1
2
3
4
...
More Than 256 Programmable Virtual Interfaces
Ethernet NICs
Fibre Channel HBAs
40-Gbps Bandwidth to Fabric A
40-Gbps Bandwidth to Fabric B
8 10GBASE-KR Unified Network Fabric, 4 to Each Fabric Extender
4 Additional Connections Enabled via Optional Port Expander Card
Hardware PortChannels
Mezzanine LOM Card Form Factor
Primary Path
Secondary Path
More Than 256 Programmable Virtual Interfaces
Ethernet NICs
Fibre Channel HBAs

Cisco Switching

Cisco Nexus 93180YC EX Switches

The Cisco Nexus 93180YC-EX Switch has 48 1/10/25G-Gbps Small Form Pluggable Plus (SFP+) ports and 6 40/100-Gbps Quad SFP+ (QSFP+) uplink ports. All ports are line rate, delivering 3.6 Tbps of throughput in a 1-rack-unit (1RU) form factor.

Architectural Flexibility

- Includes top-of-rack, fabric extender aggregation, or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Includes leaf node support for Cisco ACI architecture
- 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

Real-Time Visibility and Telemetry

- Cisco Tetration Analytics Platform support with built-in hardware sensors for rich traffic flow telemetry and line-rate data collection
- Cisco Nexus Data Broker support for network traffic monitoring and analysis
- Real-time buffer utilization per port and per queue, for monitor traffic micro-bursts and application traffic patterns

Highly Available and Efficient Design

- High-performance, non-blocking architecture
- Easily deployed into either a hot-aisle and cold-aisle configuration
- Redundant, hot-swappable power supplies and fan trays

Simplified Operations

- Pre-boot execution environment (PXE) and Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
- Automate and configure switches with DevOps tools like Puppet, Chef, and Ansible
- An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
- Python scripting gives programmatic access to the switch command-line interface (CLI)
- Includes hot and cold patching and online diagnostics

Investment Protection

- A Cisco 40-Gb bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet
- Support for 10-Gb and 25-Gb access connectivity and 40-Gb and 100-Gb uplinks facilitate data centers migrating switching infrastructure to faster speeds
- 1.44 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10-Gbps SFP+ ports
- 6 fixed 40-Gbps QSFP+ for uplink connectivity that can be turned into 10 Gb ports through a QSFP to SFP or SFP+ Adapter (QSA)
- Latency of 1 to 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies

- Hot swappable 2+1 redundant fan tray

Figure 11 Cisco Nexus 93180YC-EX Switch

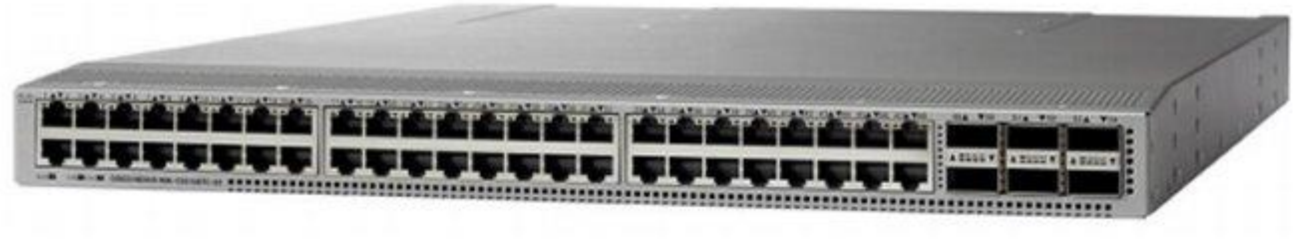

Cisco MDS 9148S Fiber Channel Switch

The Cisco MDS 9148S 16G Multilayer Fabric Switch is the next generation of the highly reliable Cisco MDS 9100 Series Switches. It includes up to 48 auto-sensing line-rate 16-Gbps Fibre Channel ports in a compact easy to deploy and manage 1-rack-unit (1RU) form factor. In all, the Cisco MDS 9148S is a powerful and flexible switch that delivers high performance and comprehensive Enterprise-class features at an affordable price.

MDS 9148S has a pay-as-you-grow model which helps you scale from a 12 port base license to a 48 port with an incremental 12 port license. This helps customers to pay and activate only the required ports.

MDS 9148S has a dual power supply and FAN trays to provide physical redundancy. The software features, like ISSU and ISSD, helps with upgrading and downgrading code without reloading the switch and without interrupting the live traffic.

Figure 12 Cisco 9148S MDS Fibre Channel Switch

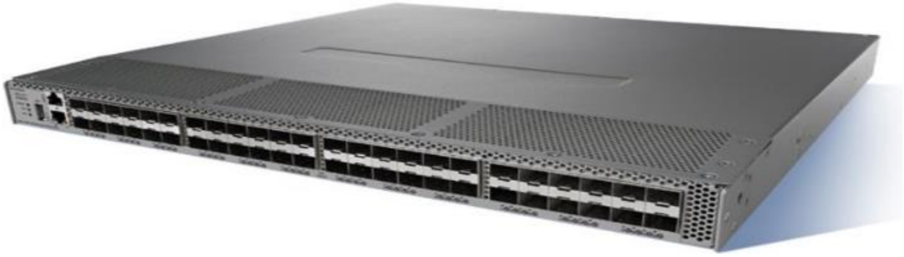

Features and Capabilities

Benefits

- Flexibility for growth and virtualization
- Easy deployment and management
- Optimized bandwidth utilization and reduced downtime
- Enterprise-class features and reliability at low cost

Features

- PowerOn Auto Provisioning and intelligent diagnostics
- In-Service Software Upgrade and dual redundant hot-swappable power supplies for high availability
- Role-based authentication, authorization, and accounting services to support regulatory requirements
- High-performance interswitch links with multipath load balancing
- Smart zoning and virtual output queuing
- Hardware-based slow port detection and recovery

Specifications at-a-Glance

Performance and Port Configuration

- 2/4/8/16-Gbps auto-sensing with 16 Gbps of dedicated bandwidth per port
- Up to 256 buffer credits per group of 4 ports (64 per port default, 253 maximum for a single port in the group)
- Supports configurations of 12, 24, 36, or 48 active ports, with pay-as-you-grow, on-demand licensing

Advanced Functions

- Virtual SAN (VSAN)
- Inter-VSAN Routing (IVR)
- PortChannel with multipath load balancing
- Flow-based and zone-based QoS

Hypervisor

This Cisco Validated Design includes VMware vSphere 6.5.

VMware vSphere 6.5

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

VMware vSphere 6.5 introduces many enhancements to vSphere Hypervisor, VMware virtual machines, vCenter Server, virtual storage, and virtual networking, further extending the core capabilities of the vSphere platform.

Today VMware announced vSphere 6.5, which is one of the most feature rich releases of vSphere in quite some time. The vCenter Server Appliance is taking charge in this release with several new features which we'll cover in this blog article. For starters, the installer has gotten an overhaul with a new modern look and feel. Users of both Linux and Mac will also be ecstatic since the installer is now supported on those platforms along with Microsoft Windows. If that wasn't enough, the vCenter Server Appliance now has features that are exclusive such as:

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

vSphere Client

With vSphere 6.5 I'm excited to say that we have a fully supported version of the HTML5-based vSphere Client that will run alongside the vSphere Web Client. The vSphere Client is built right into vCenter Server 6.5 (both Windows and Appliance) and is enabled by default. While the vSphere Client does not yet have full feature parity the team has prioritized many of the day to day tasks of administrators and continue to seek feedback on what is missing that will enable customers to use it full-time. The vSphere Web Client will continue to be accessible via "http://<vcenter_fqdn>/vsphere-client" while the vSphere Client will be reachable via "http://<vcenter_fqdn>/ui". VMware will also be periodically updating the vSphere Client outside of the normal vCenter Server release cycle. To make sure it is easy and simple for customers to stay up to date the vSphere Client will be able to be updated without any effects to the rest of vCenter Server.

The following are some of the benefits of the new vSphere Client:

- Clean, consistent UI built on VMware's new Clarity UI standards (to be adopted across our portfolio)
- Built on HTML5 so it is truly a cross-browser and cross-platform application
- No browser plugins to install/manage
- Integrated into vCenter Server for 6.5 and fully supported
- Fully supports Enhanced Linked Mode
- Users of the Fling have been extremely positive about its performance

VMware ESXi 6.5 Hypervisor

vSphere 6.5 introduces a number of new features in the hypervisor:

- Scalability Improvements

  ESXi 6.5 dramatically increases the scalability of the platform. With vSphere Hypervisor 6.0, clusters can scale to as many as 64 hosts, up from 32 in previous releases. With 64 hosts in a cluster, vSphere 6.0 can support 8000 virtual machines in a single cluster. This capability enables greater consolidation ratios, more efficient use of VMware vSphere Distributed Resource Scheduler (DRS), and fewer clusters that must be separately managed. Each vSphere Hypervisor 6.5 instance can support up to 480 logical CPUs, 12 terabytes (TB) of RAM, and 1024 virtual machines. By using the newest hardware advances, ESXi 6.5 enables the virtualization of applications that previously had been thought to be non-virtualizable.

- Security Enhancements

- Account management: ESXi 6.5 enables management of local accounts on the ESXi server using new ESXi CLI commands. The capability to add, list, remove, and modify accounts across all hosts in a cluster can be centrally managed using a vCenter Server system. Previously, the account and permission management functions for ESXi hosts were available only for direct host connections. The setup, removal, and listing of local permissions on ESXi servers can also be centrally managed.

- Account lockout: ESXi Host Advanced System Settings have two new options for the management of failed local account login attempts and account lockout duration. These parameters affect Secure Shell (SSH) and vSphere Web Services connections, but not ESXi direct console user interface (DCUI) or console shell access.

- Password complexity rules: In previous versions of ESXi, password complexity changes had to be made by manually editing the /etc/pam.d/passwd file on each ESXi host. In vSphere 6.0, an entry in Host Advanced System Settings enables changes to be centrally managed for all hosts in a cluster.

- Improved auditability of ESXi administrator actions: Prior to vSphere 6.0, actions at the vCenter Server level by a named user appeared in ESXi logs with the vpxuser username: for example, [user=vpxuser]. In vSphere 6.5, all actions at the vCenter Server level for an ESXi server appear in the ESXi logs with the vCenter Server username: for example, [user=vpxuser: DOMAIN\User]. This approach provides a better audit trail for actions run on a vCenter Server instance that conducted corresponding tasks on the ESXi hosts.

- Flexible lockdown modes: Prior to vSphere 6.5, only one lockdown mode was available. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.5, two lockdown modes are available:

  - In normal lockdown mode, DCUI access is not stopped, and users on the DCUI access list can access the DCUI.
  - In strict lockdown mode, the DCUI is stopped.
  - Exception users: vSphere 6.0 offers a new function called exception users. Exception users are local accounts or Microsoft Active Directory accounts with permissions defined locally on the host to which these users have host access. These exception users are not recommended for general user accounts, but they are recommended for use by third-party applications—for service accounts, for example—that need host access when either normal or strict lockdown mode is enabled. Permissions on these accounts should be set to the bare minimum required for the application to perform its task and with an account that needs only read-only permissions on the ESXi host.

- Smart card authentication to DCUI: This function is for U.S. federal customers only. It enables DCUI login access using a Common Access Card (CAC) and Personal Identity Verification (PIV). The ESXi host must be part of an Active Directory domain.

Citrix XenDestop 7.15

Enterprise IT organizations are tasked with the challenge of provisioning Microsoft Windows apps and desktops while managing cost, centralizing control, and enforcing corporate security policy. Deploying Windows apps to users in any location, regardless of the device type and available network bandwidth, enables a mobile workforce that can improve productivity. With Citrix XenDesktop 7.15, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

The XenDesktop 7.15 release offers these benefits:

- **Comprehensive virtual desktop delivery for any use case**. The XenDesktop 7.15 release incorporates the full power of XenApp, delivering full desktops or just applications to users. Administrators can deploy both XenApp published applications and desktops (to maximize IT control at low cost) or personalized VDI desktops (with simplified image management) from the same management console. Citrix XenDesktop 7.15 leverages common policies and cohesive tools to govern both infrastructure resources and user access.

- **Simplified support and choice of BYO (Bring Your Own) devices**. XenDesktop 7.15 brings thousands of corporate Microsoft Windows-based applications to mobile devices with a native-touch experience and optimized performance. HDX technologies create a “high definition” user experience, even for graphics intensive design and engineering applications.

- **Lower cost and complexity of application and desktop management**. XenDesktop 7.15 helps IT organizations take advantage of agile and cost-effective cloud offerings, allowing the virtualized infrastructure to flex and meet seasonal demands or the need for sudden capacity changes. IT organizations can deploy XenDesktop application and desktop workloads to private or public clouds.

- **Protection of sensitive information through centralization**. XenDesktop decreases the risk of corporate data loss, enabling access while securing intellectual property and centralizing applications since assets reside in the datacenter.

- **Virtual Delivery Agent improvements.** Universal print server and driver enhancements and support for the HDX 3D Pro graphics acceleration for Windows 10 are key additions in XenDesktop 7.15

- **Improved high-definition user experience.** XenDesktop 7.15 continues the evolutionary display protocol leadership with enhanced Thinwire display remoting protocol and Framehawk support for HDX 3D Pro.

Citrix XenApp and XenDesktop are application and desktop virtualization solutions built on a unified architecture so they're simple to manage and flexible enough to meet the needs of all your organization's users. XenApp and XenDesktop have a common set of management tools that simplify and automate IT tasks. You use the same architecture and management tools to manage public, private, and hybrid cloud deployments as you do for on premises deployments.

Citrix XenApp delivers:

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

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

- Virtual machine–hosted apps: These are applications hosted from machines running Windows desktop operating systems for applications that can’t be hosted in a server environment.

- Windows applications delivered with Microsoft App-V: These applications use the same management tools that you use for the rest of your XenApp deployment.

- Citrix XenDesktop: Includes significant enhancements to help customers deliver Windows apps and desktops as mobile services while addressing management complexity and associated costs. Enhancements in this release include:

- Unified product architecture for XenApp and XenDesktop: The FlexCast Management Architecture (FMA). This release supplies a single set of administrative interfaces to deliver both hosted-shared applications (RDS) and complete virtual desktops (VDI). Unlike earlier releases that separately provisioned Citrix XenApp and XenDesktop farms, the XenDesktop 7.15 release allows administrators to deploy a single infrastructure and use a consistent set of tools to manage mixed application and desktop workloads.

- Support for extending deployments to the cloud. This release provides the ability for hybrid cloud provisioning from Microsoft Azure, Amazon Web Services (AWS) or any Cloud Platform-powered public or private cloud. Cloud deployments are configured, managed, and monitored through the same administrative consoles as deployments on traditional on-premises infrastructure.

Citrix XenDesktop delivers:

- VDI desktops: These virtual desktops each run a Microsoft Windows desktop operating system rather than running in a shared, server-based environment. They can provide users with their own desktops that they can fully personalize.
- Hosted physical desktops: This solution is well suited for providing secure access to powerful physical machines, such as blade servers, from within your data center.
- Remote PC access: This solution allows users to log in to their physical Windows PC from anywhere over a secure XenDesktop connection.
- Server VDI: This solution is designed to provide hosted desktops in multitenant, cloud environments.
- Capabilities that allow users to continue to use their virtual desktops: These capabilities let users continue to work while not connected to your network.

This product release includes the following new and enhanced features:

Some XenDesktop editions include the features available in XenApp.

Zones

Deployments that span widely-dispersed locations connected by a WAN can face challenges due to network latency and reliability. Configuring zones can help users in remote regions connect to local resources without forcing connections to traverse large segments of the WAN. Using zones allows effective Site management from a single Citrix Studio console, Citrix Director, and the Site database. This saves the costs of deploying, staffing, licensing, and maintaining additional Sites containing separate databases in remote locations.

Zones can be helpful in deployments of all sizes. You can use zones to keep applications and desktops closer to end users, which improves performance.

For more information, see the Zones article.

Improved Database Flow and Configuration

When you configure the databases during Site creation, you can now specify separate locations for the Site, Logging, and Monitoring databases. Later, you can specify different locations for all three databases. In previous releases, all three databases were created at the same address, and you could not specify a different address for the Site database later.

You can now add more Delivery Controllers when you create a Site, as well as later. In previous releases, you could add more Controllers only after you created the Site.

For more information, see the Databases and Controllers articles.

Application Limits

Configure application limits to help manage application use. For example, you can use application limits to manage the number of users accessing an application simultaneously. Similarly, application limits can be used to manage the number of simultaneous instances of resource-intensive applications, this can help maintain server performance and prevent deterioration in service.

For more information, see the Manage applications article.

Multiple Notifications before Machine Updates or Scheduled Restarts

You can now choose to repeat a notification message that is sent to affected machines before the following types of actions begin:

- Updating machines in a Machine Catalog using a new master image
- Restarting machines in a Delivery Group according to a configured schedule

If you indicate that the first message should be sent to each affected machine 15 minutes before the update or restart begins, you can also specify that the message is repeated every five minutes until the update/restart begins.

For more information, see the Manage Machine Catalogs and Manage Delivery Groups articles.

API Support for Managing Session Roaming

By default, sessions roam between client devices with the user. When the user launches a session and then moves to another device, the same session is used and applications are available on both devices. The applications follow, regardless of the device or whether current sessions exist. Similarly, printers and other resources assigned to the application follow.

You can now use the PowerShell SDK to tailor session roaming. This was an experimental feature in the previous release.

For more information, see the Sessions article.

API Support for Provisioning VMs from Hypervisor Templates

When using the PowerShell SDK to create or update a Machine Catalog, you can now select a template from other hypervisor connections. This is in addition to the currently-available choices of VM images and snapshots.

Support for New and Additional Platforms

See the System requirements article for full support information. Information about support for third-party product versions is updated periodically.

By default, SQL Server 2014 SP2 Express is installed when installing the Controller, if an existing supported SQL Server installation is not detected.

You can install Studio or VDAs for Windows Desktop OS on machines running Windows 10.

You can create connections to Microsoft Azure virtualization resources.

Figure 13 Logical Architecture of Citrix XenDesktop

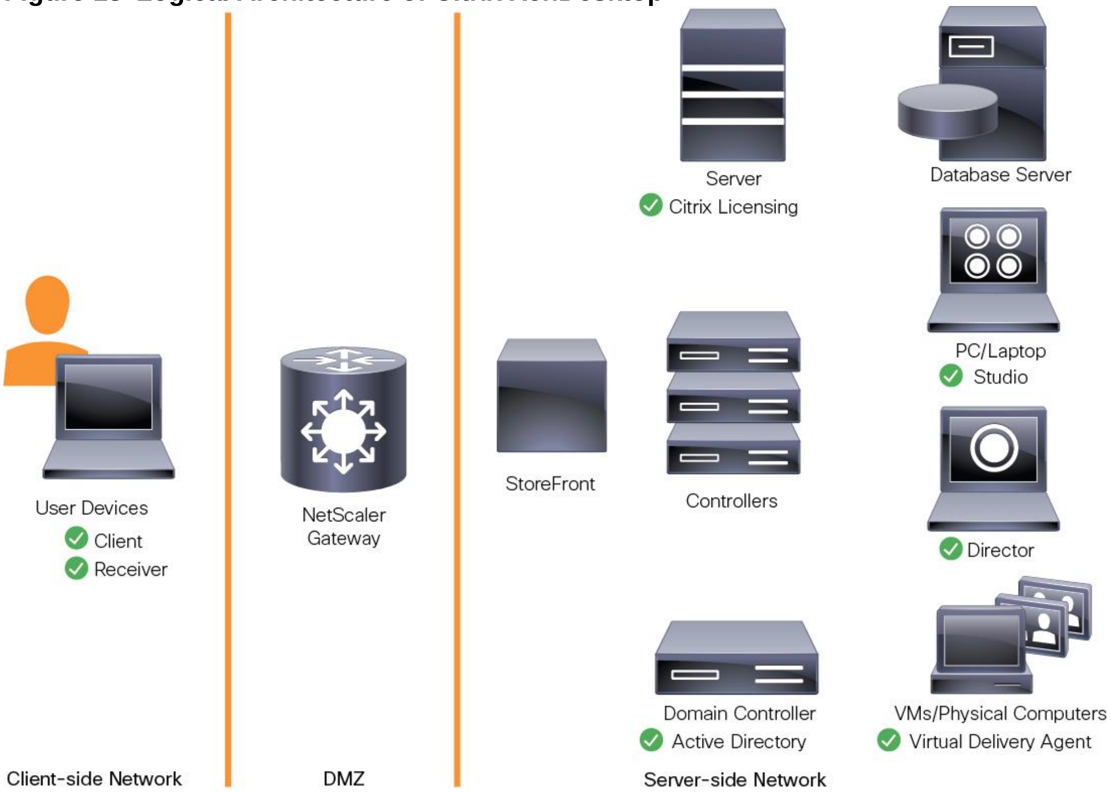

Citrix Provisioning Services 7.15

Most enterprises struggle to keep up with the proliferation and management of computers in their environments. Each computer, whether it is a desktop PC, a server in a data center, or a kiosk-type device, must be managed as an individual entity. The benefits of distributed processing come at the cost of distributed management. It costs time and money to set up, update, support, and ultimately decommission each computer. The initial cost of the machine is often dwarfed by operating costs.

Citrix PVS takes a very different approach from traditional imaging solutions by fundamentally changing the relationship between hardware and the software that runs on it. By streaming a single shared disk image (vDisk) rather than copying images to individual machines, PVS enables organizations to reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiency of centralized management and the benefits of distributed processing.

In addition, because machines are streaming disk data dynamically and in real time from a single shared image, machine image consistency is essentially ensured. At the same time, the configuration, applications, and even the OS of large pools of machines can be completed changed in the time it takes the machines to reboot.

Using PVS, any vDisk can be configured in standard-image mode. A vDisk in standard-image mode allows many computers to boot from it simultaneously, greatly reducing the number of images that must be maintained and the amount of storage that is required. The vDisk is in read-only format, and the image cannot be changed by target devices.

Benefits for Citrix XenApp and Other Server Farm Administrators

If you manage a pool of servers that work as a farm, such as Citrix XenApp servers or web servers, maintaining a uniform patch level on your servers can be difficult and time consuming. With traditional imaging solutions, you start with a clean golden master image, but as soon as a server is built with the master image, you must patch that individual server along with all the other individual servers. Rolling out patches to individual servers in your farm is not only inefficient, but the results can also be unreliable. Patches often fail on an individual server, and you may not realize you have a problem until users start complaining or the server has an outage. After that happens, getting the server resynchronized with the rest of the farm can be challenging, and sometimes a full reimaging of the machine is required.

With Citrix PVS, patch management for server farms is simple and reliable. You start by managing your golden image, and you continue to manage that single golden image. All patching is performed in one place and then streamed to your servers when they boot. Server build consistency is assured because all your servers use a single shared copy of the disk image. If a server becomes corrupted, simply reboot it, and it is instantly back to the known good state of your master image. Upgrades are extremely fast to implement. After you have your updated image ready for production, you simply assign the new image version to the servers and reboot them. You can deploy the new image to any number of servers in the time it takes them to reboot. Just as important, rollback can be performed in the same way, so problems with new images do not need to take your servers or your users out of commission for an extended period of time.

Benefits for Desktop Administrators

Because Citrix PVS is part of Citrix XenDesktop, desktop administrators can use PVS’s streaming technology to simplify, consolidate, and reduce the costs of both physical and virtual desktop delivery. Many organizations are beginning to explore desktop virtualization. Although virtualization addresses many of IT’s needs for consolidation and simplified management, deploying it also requires deployment of supporting infrastructure. Without PVS, storage costs can make desktop virtualization too costly for the IT budget. However, with PVS, IT can reduce the amount of storage required for VDI by as much as 90 percent. And with a single image to manage instead of hundreds or thousands of desktops, PVS significantly reduces the cost, effort, and complexity for desktop administration.

Different types of workers across the enterprise need different types of desktops. Some require simplicity and standardization, and others require high performance and personalization. XenDesktop can meet these requirements in a single solution using Citrix FlexCast delivery technology. With FlexCast, IT can deliver every type of virtual desktop, each specifically tailored to meet the performance, security, and flexibility requirements of each individual user.

Not all desktops applications can be supported by virtual desktops. For these scenarios, IT can still reap the benefits of consolidation and single-image management. Desktop images are stored and managed centrally in the data center and streamed to physical desktops on demand. This model works particularly well for standardized desktops such as those in lab and training environments and call centers and thin-client devices used to access virtual desktops.

Citrix Provisioning Services Solution

Citrix PVS streaming technology allows computers to be provisioned and re-provisioned in real time from a single shared disk image. With this approach, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is performed on the master image. The local hard drive of each system can be used for runtime data caching or, in some scenarios, removed from the system entirely, which reduces power use, system failure rate, and security risk.

The PVS solution’s infrastructure is based on software-streaming technology. After PVS components are installed and configured, a vDisk is created from a device’s hard drive by taking a snapshot of the OS and application image and then storing that image as a vDisk file on the network. A device used for this process is referred to as a master target device. The devices that use the vDisks are called target devices. vDisks can exist on a PVS, file share, or in larger deployments, on a storage system with which PVS can communicate (iSCSI, SAN, network-attached storage [NAS], and Common Internet File System [CIFS]). vDisks can be assigned to a single target device in private-image mode, or to multiple target devices in standard-image mode.

Citrix Provisioning Services Infrastructure

The Citrix PVS infrastructure design directly relates to administrative roles within a PVS farm. The PVS administrator role determines which components that administrator can manage or view in the console.

A PVS farm contains several components. Figure 14 provides a high-level view of a basic PVS infrastructure and shows how PVS components might appear within that implementation.

Figure 14 Logical Architecture of Citrix Provisioning Services

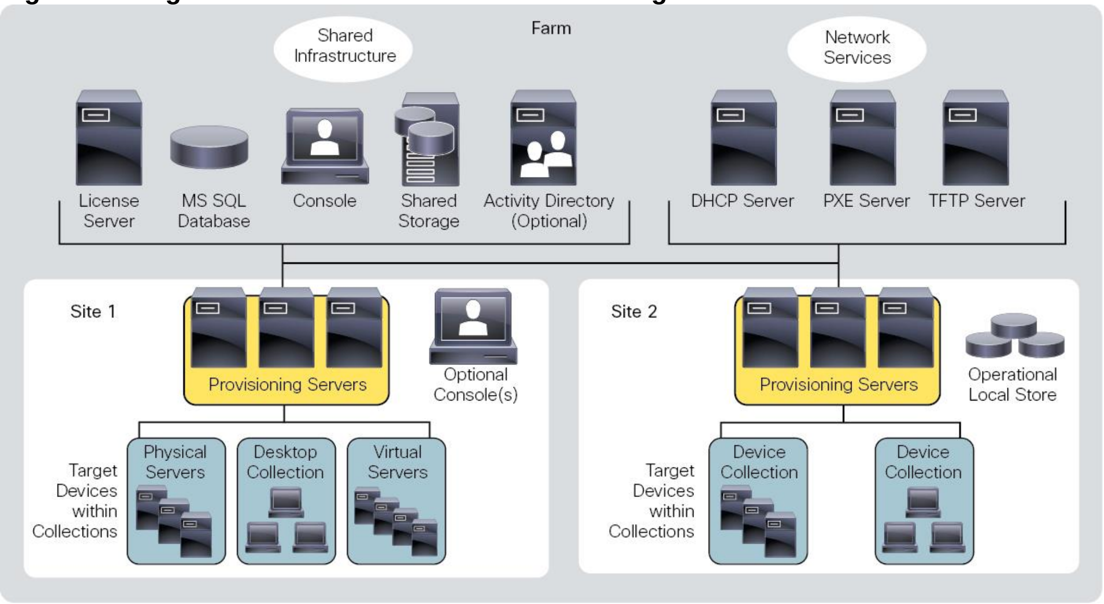

The following new features are available with Provisioning Services 7.15:

- Linux streaming
- XenServer proxy using PVS-Accelerator

NetApp A-Series All Flash FAS

With the new NetApp A-Series All Flash FAS (AFF) controller lineup, NetApp provides industry-leading performance while continuing to provide a full suite of enterprise-grade data management and data protection features. The A-Series lineup offers double the IOPS, while decreasing the latency. The AFF A-Series lineup includes the A200, A300, A700, and A700s. These controllers and their specifications listed in Table 5 . For more information about the A-Series AFF controllers, see:

- http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx
- https://hwu.netapp.com/Controller/Index

Table 5   NetApp A-Series Controller Specifications

	AFF A200	AFF A300	AFF A700	AFF A700s
NAS Scale-out	2-8 nodes	2-24 nodes	2-24 nodes	2-24 nodes
SAN Scale-out	2-8 nodes	2-12 nodes	2-12 nodes	2-12 nodes
Per HA Pair Specifications (Active-Active Dual Controller)				
Maximum SSDs	144	384	480	216
Maximum Raw Capacity	2.2PB	5.9PB	7.3PB	3.3PB
Effective Capacity	8.8PB	23.8PB	29.7PB	13PB

Chassis Form Factor	2U chassis with two HA controllers and 24 SSD slots	3U chassis with two HA controllers	8u chassis with two HA controllers	4u chassis with two HA controllers and 24 SSD slots

This solution utilizes the NetApp AFF A300, seen in Figure 15 and Figure 16. This controller provides the high-performance benefits of 40GbE and all flash SSDs, offering better performance than previous models, and occupying only 3U of rack space versus 6U with the AFF8040. When combined with the 2U disk shelf of 3.8TB disks, this solution can provide ample horsepower and over 90TB of raw capacity, all while occupying only 5U of valuable rack space. This makes it an ideal controller for a shared workload converged infrastructure. The A700s would be an ideal fit for situations where more performance is needed

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

For more information about the AFF A-Series product family, see: http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx

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

Figure 15 NetApp AFF A300 Front View

Figure 16 NetApp AFF A300 Rear View

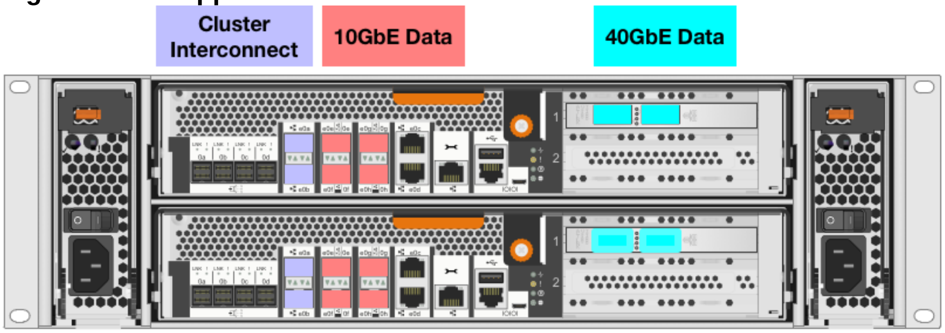

NetApp ONTAP 9.3

Storage Efficiency

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

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

Compaction, which is introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the NetApp WAFL® file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk-to-save space. This process is illustrated in Figure 17.

Figure 17 Storage Efficiency

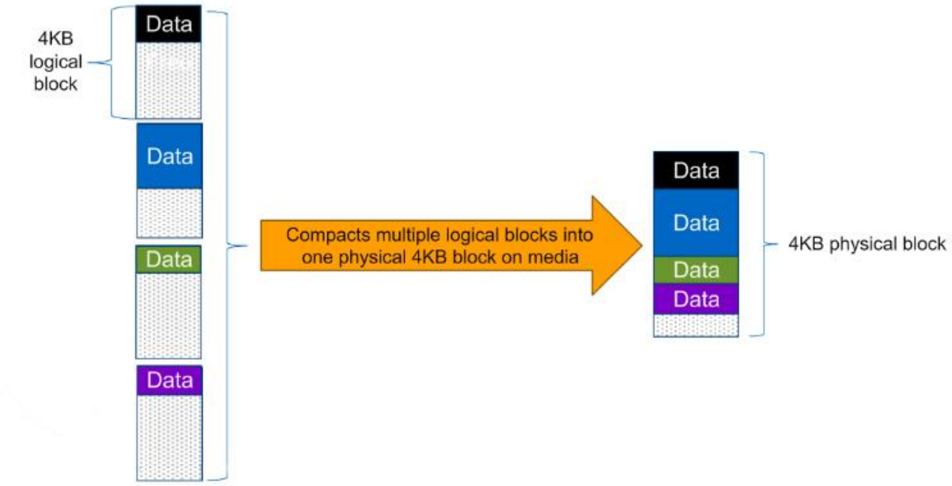
4KB logical block
Data
Data
Data
Data
Compacts multiple logical blocks into one physical 4KB block on media
Data
Data
Data
Data
4KB physical block

NetApp Storage Virtual Machine (SVM)

A cluster serves data through at least one and possibly multiple storage virtual machines (SVMs, formerly called Vservers). An SVM is a logical abstraction that represents the set of physical resources of the cluster. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node to another. For example, a flexible volume can be non-disruptively moved to a new node and aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware and it is not tied to any specific physical hardware.

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

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

SAN Boot

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

In FC SAN boot, each Cisco UCS server boots by connecting the NetApp All Flash FAS storage to the Cisco MDS switch. The 16G FC storage ports, in this example 0g and 0h, are connected to Cisco MDS switch. The FC LIFs are created on the physical ports and each FC LIF is uniquely identified by its target WWPN. The storage system target WWPNs can be zoned with the server initiator WWPNs in the Cisco MDS switches. The FC boot LUN is exposed to the servers through the FC LIF using the MDS switch; this enables only the authorized server to have access to the boot LUN. Refer to Figure 18 for the port and LIF layout

Figure 18 FC - SVM ports and LIF layout

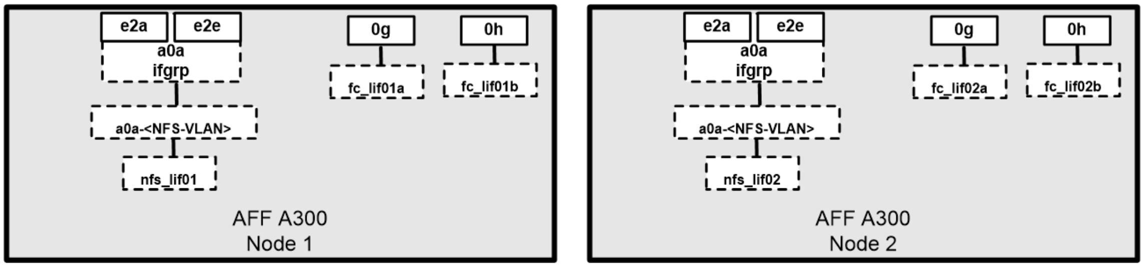

Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead if a network interface becomes unavailable, the host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp used to provide information about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

FlexGroups

ONTAP 9.3 brought an innovation in scale-out NAS file systems: NetApp FlexGroup volumes.

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

Figure 19 Illustration of FlexGroups

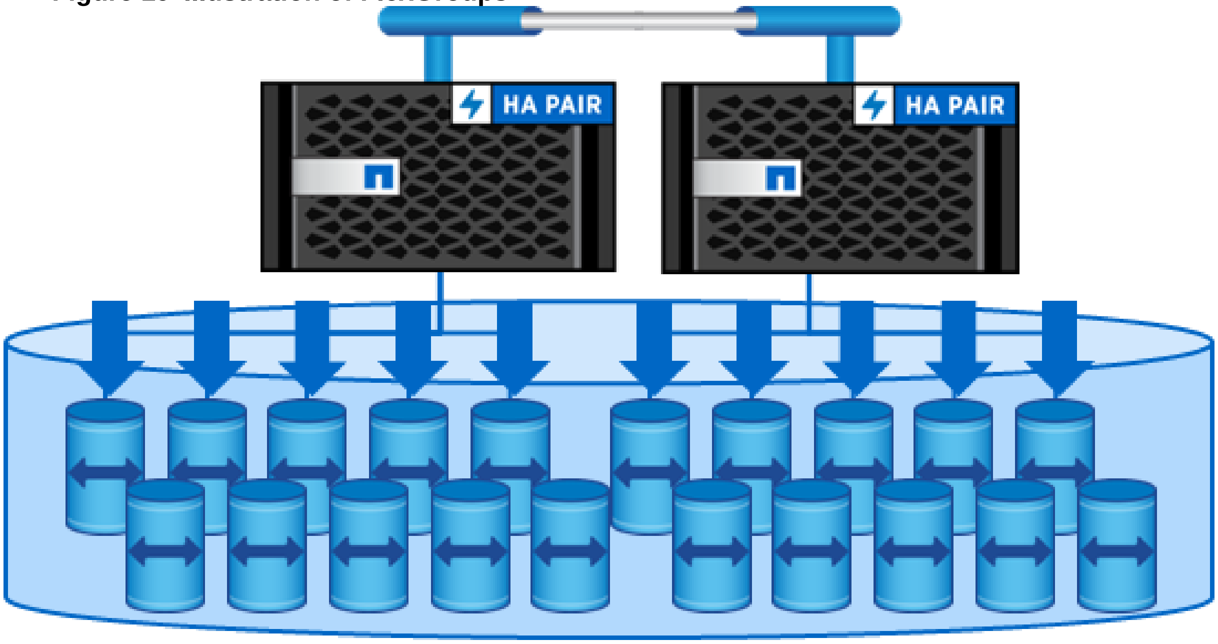

Architecture and Design Considerations for Desktop Virtualization

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: a physical device with a locally installed operating system.
- Hosted Shared Desktop: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2016, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remoted Desktop Server Hosted Server sessions: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead, the user interacts through a delivery protocol.
- Published Applications: Published applications run entirely on the Citrix XenApp 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 XenDesktop Virtual Desktops and XenApp Hosted Shared Desktop 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 of cloud applications, for example, SalesForce.com. This application and data analysis is beyond the scope of this Cisco Validated Design, but should not be omitted from the planning process. There are a variety of third-party tools available to assist organizations with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

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

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

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

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

- What is the desktop OS planned? Windows 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 OS planned for RDS Server Roles? Windows Server 2012 or Server 2016?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?

- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- What is the SQL server version for the database? SQL server 2012 or 2016?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere has been identified as the hypervisor for both HSD Sessions and HVD based desktops.

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

For this CVD, the hypervisor used was VMware ESXi 6.5 Update 1.

The server OS and desktop OS machines configured in this CVD support the Remote Desktop Server Hosted (RDSH) shared sessions and Hosted Virtual Desktops (both non-persistent and persistent).

Citrix XenDesktop Design Fundamentals

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

Citrix XenDesktop 7.15 integrates Hosted Shared and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, and manageable solution for delivering Windows applications and desktops as a service.

Users can select applications from an easy-to-use “store” that is accessible from tablets, smartphones, PCs, Macs, and thin clients. XenDesktop delivers a native touch-optimized experience with HDX high-definition performance, even over mobile networks.

Machine Catalogs

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Machine Catalog. In this CVD, VM provisioning relies on Citrix Provisioning Services to make sure that the machines in the catalog are consistent. In this CVD, machines in the Machine Catalog are configured to run either a Windows Server OS (for RDS hosted shared desktops) or a Windows Desktop OS (for hosted pooled VDI desktops).

Delivery Groups

To deliver desktops and applications to users, you create a Machine Catalog and then allocate machines from the catalog to users by creating Delivery Groups. Delivery Groups provide desktops, applications, or a combination of desktops and applications to users. Creating a Delivery Group is a flexible way of allocating machines and applications to users. In a Delivery Group, you can:

- Use machines from multiple catalogs
- Allocate a user to multiple machines
- Allocate multiple users to one machine

As part of the creation process, you specify the following Delivery Group properties:

- Users, groups, and applications allocated to Delivery Groups
- Desktop settings to match users' needs
- Desktop power management options

Figure 20 illustrates how users access desktops and applications through machine catalogs and delivery groups.

The Server OS and Desktop OS Machines configured in this CVD support the hosted shared desktops and hosted virtual desktops (both non-persistent and persistent).

Figure 20 Access Desktops and Applications through Machine Catalogs and Delivery Groups

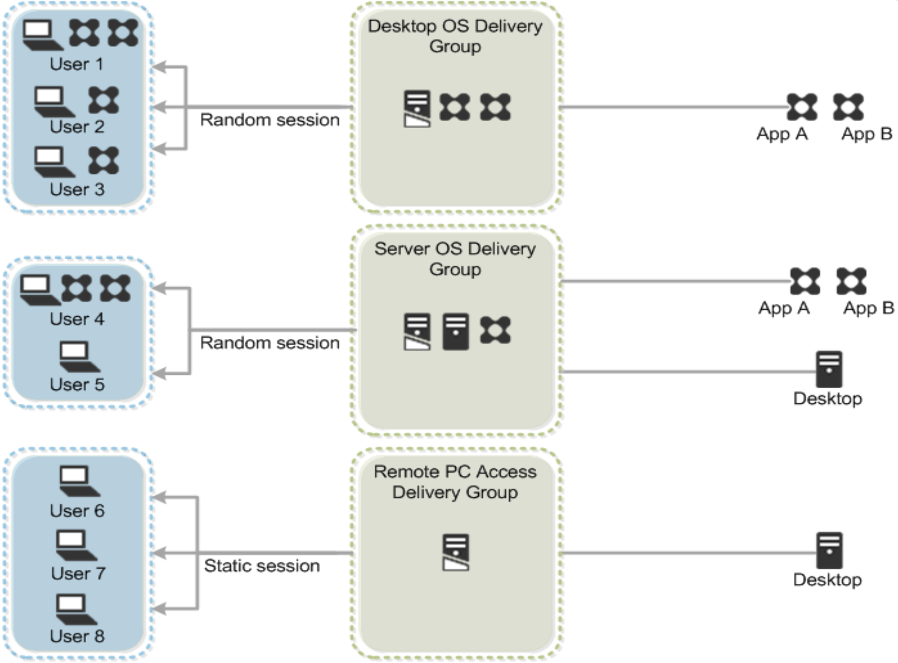

Citrix Provisioning Services

Citrix XenDesktop 7.15 can be deployed with or without Citrix Provisioning Services (PVS). The advantage of using Citrix PVS is that it allows virtual machines to be provisioned and re-provisioned in real-time from a single shared-disk image. In this way administrators can completely eliminate the need to manage and patch individual systems and reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiencies of a centralized management with the benefits of distributed processing.

The Provisioning Services solution's infrastructure is based on software-streaming technology. After installing and configuring Provisioning Services components, a single shared disk image (vDisk) is created from a device's hard drive by taking a snapshot of the OS and application image, and then storing that image as a vDisk file on the network. A device that is used during the vDisk creation process is the Master target device. Devices or virtual machines that use the created vDisks are called target devices.

When a target device is turned on, it is set to boot from the network and to communicate with a Provisioning Server. Unlike thin-client technology, processing takes place on the target device.

Figure 21 Citrix Provisioning Services Functionality

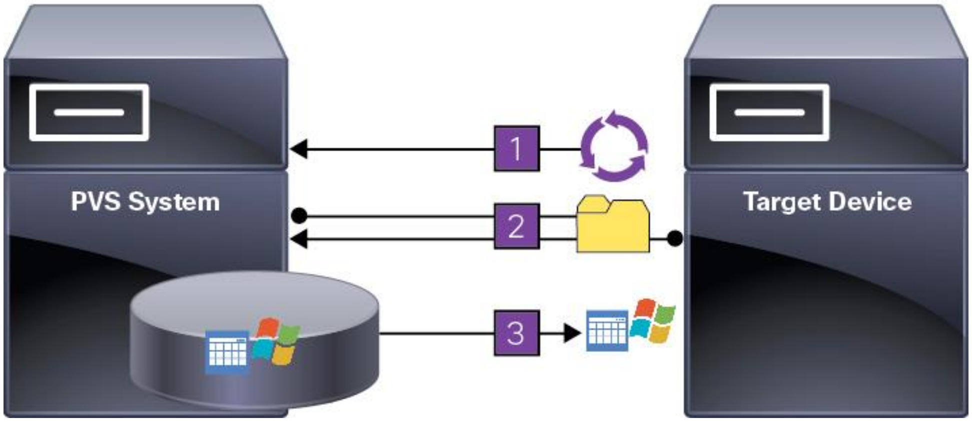

The target device downloads the boot file from a Provisioning Server (Step 2) and boots. Based on the boot configuration settings, the appropriate vDisk is mounted on the Provisioning Server (Step 3). The vDisk software is then streamed to the target device as needed, appearing as a regular hard drive to the system.

Instead of immediately pulling all the vDisk contents down to the target device (as with traditional imaging solutions), the data is brought across the network in real-time as needed. This approach allows a target device to get a completely new operating system and set of software in the time it takes to reboot. This approach dramatically decreases the amount of network bandwidth required and making it possible to support a larger number of target devices on a network without impacting performance

Citrix PVS can create desktops as Pooled or Private:

- Pooled Desktop: A pooled virtual desktop uses Citrix PVS to stream a standard desktop image to multiple desktop instances upon boot.
- Private Desktop: A private desktop is a single desktop assigned to one distinct user.

The alternative to Citrix Provisioning Services for pooled desktop deployments is Citrix Machine Creation Services (MCS), which is integrated with the XenDesktop Studio console.

Locating the PVS Write Cache

When considering a PVS deployment, there are some design decisions that need to be made regarding the write cache for the target devices that leverage provisioning services. The write cache is a cache of all data that the target device has written. If data is written to the PVS vDisk in a caching mode, the data is not written back to the base vDisk. Instead, it is written to a write cache file in one of the following locations:

- Cache on device hard drive. Write cache exists as a file in NTFS format, located on the target-device’s hard drive. This option frees up the Provisioning Server since it does not have to process write requests and does not have the finite limitation of RAM.
- Cache on device hard drive persisted. (Experimental Phase) This is the same as “Cache on device hard drive”, except that the cache persists. At this time, this method is an experimental feature only, and is only supported for NT6.1 or later (Windows 10 and Windows 2008 R2 and later). This method also requires a different bootstrap.
- Cache in device RAM. Write cache can exist as a temporary file in the target device’s RAM. This provides the fastest method of disk access since memory access is always faster than disk access.
- Cache in device RAM with overflow on hard disk. This method uses VHDX differencing format and is only available for Windows 10 and Server 2008 R2 and later. When RAM is zero, the target device write cache is only written to the local disk. When RAM is not zero, the target device write cache is written to RAM first. When RAM is full, the least recently used block of data is written to the local differencing disk to

accommodate newer data on RAM. The amount of RAM specified is the non-paged kernel memory that the target device will consume.

- Cache on a server. Write cache can exist as a temporary file on a Provisioning Server. In this configuration, all writes are handled by the Provisioning Server, which can increase disk I/O and network traffic. For additional security, the Provisioning Server can be configured to encrypt write cache files. Since the write-cache file persists on the hard drive between reboots, encrypted data provides data protection in the event a hard drive is stolen.
- Cache on server persisted. This cache option allows for the saved changes between reboots. Using this option, a rebooted target device is able to retrieve changes made from previous sessions that differ from the read only vDisk image. If a vDisk is set to this method of caching, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

In this CVD, Provisioning Server 7.15 was used to manage Pooled/Non-Persistent VDI Machines and XenApp RDS Machines with "Cache in device RAM with Overflow on Hard Disk" for each virtual machine. This design enables good scalability to many thousands of desktops. Provisioning Server 7.15 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

Example XenDesktop Deployments

Two examples of typical XenDesktop deployments are the following:

- A distributed components configuration
- A multiple site configuration

Since XenApp and XenDesktop 7.15 are based on a unified architecture, combined they can deliver a combination of Hosted Shared Desktops (HSDs, using a Server OS machine) and Hosted Virtual Desktops (HVDs, using a Desktop OS).

Distributed Components Configuration

You can distribute the components of your deployment among a greater number of servers, or provide greater scalability and failover by increasing the number of controllers in your site. You can install management consoles on separate computers to manage the deployment remotely. A distributed deployment is necessary for an infrastructure based on remote access through NetScaler Gateway (formerly called Access Gateway).

Figure 22 shows an example of a distributed components configuration. A simplified version of this configuration is often deployed for an initial proof-of-concept (POC) deployment. The CVD described in this document deploys Citrix XenDesktop in a configuration that resembles this distributed components configuration shown. Two Cisco UCS B200M4 blade servers host the required infrastructure services (AD, DNS, DHCP, License Server, SQL, Citrix XenDesktop management, and StoreFront servers).

Figure 22 Example of a Distributed Components Configuration

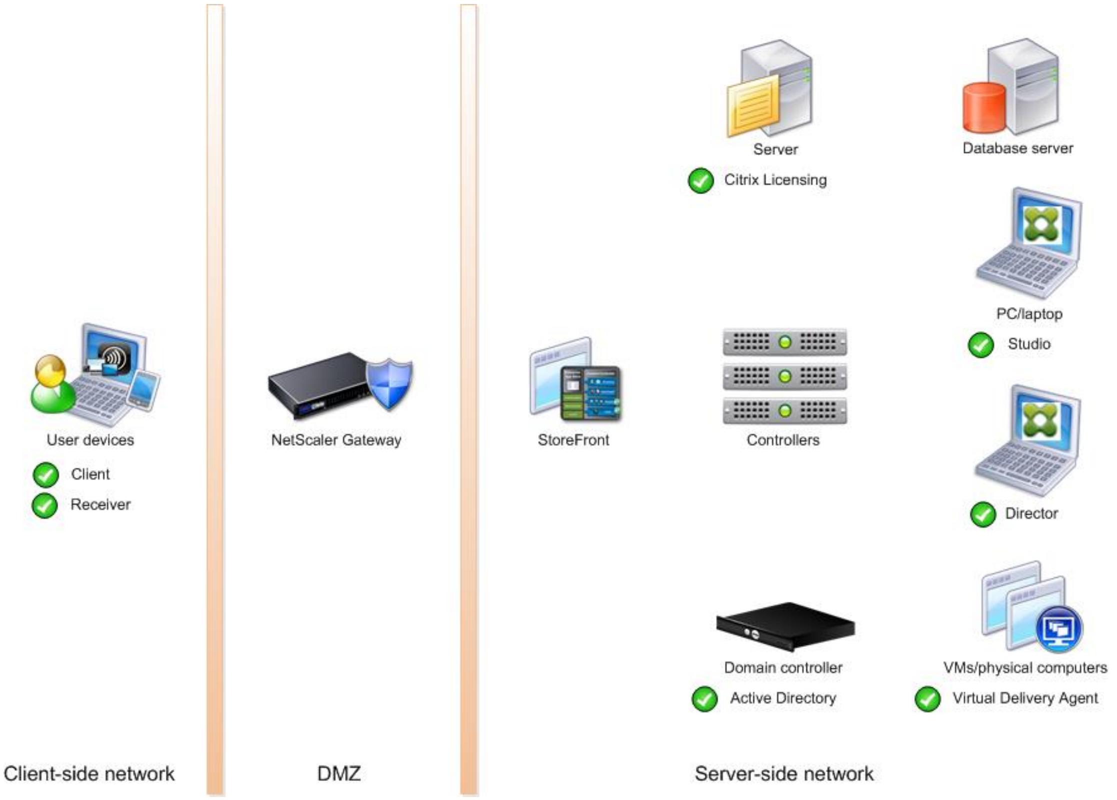

Multiple Site Configuration

If you have multiple regional sites, you can use Citrix NetScaler to direct user connections to the most appropriate site and StoreFront to deliver desktops and applications to users.

In Figure 23 depicting multiple sites, a site was created in two data centers. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic.

Figure 23 Multiple Sites

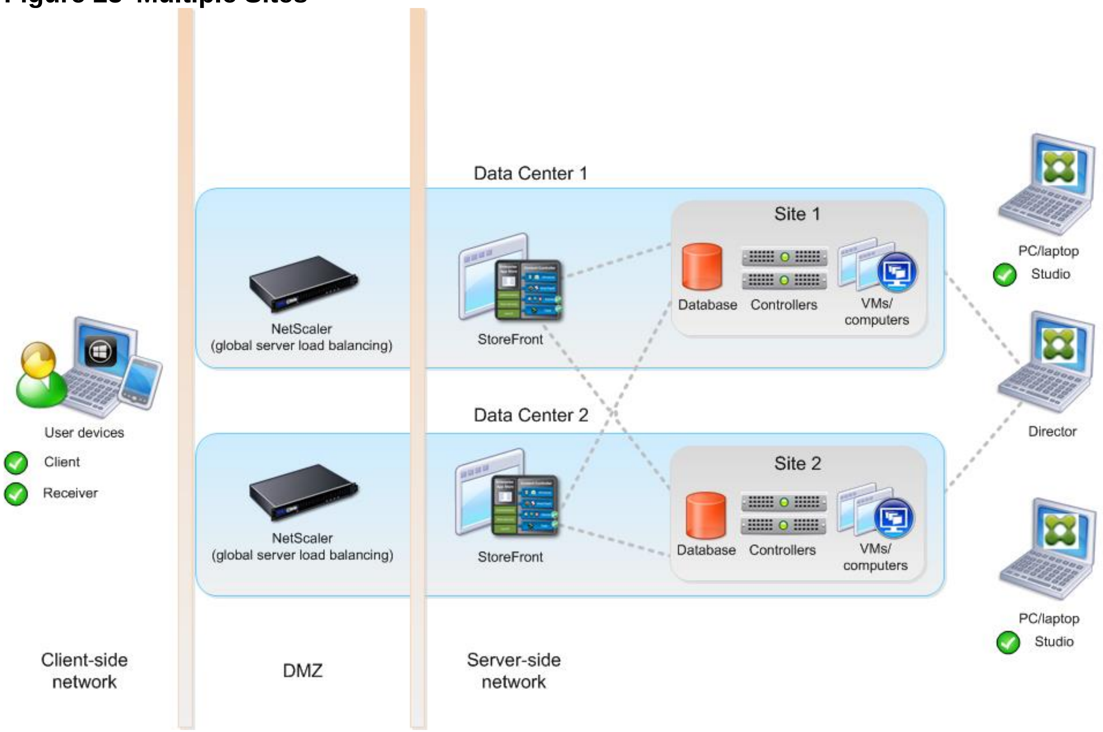

You can use StoreFront to aggregate resources from multiple sites to provide users with a single point of access with NetScaler. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

Citrix NetScaler accelerates application performance, load balances servers, increases security, and optimizes the user experience. In this example, two NetScalers are used to provide a high availability configuration. The NetScalers are configured for Global Server Load Balancing and positioned in the DMZ to provide a multi-site, fault-tolerant solution.

Citrix Cloud Services

Easily deliver the Citrix portfolio of products as a service. Citrix Cloud services simplify the delivery and management of Citrix technologies extending existing on-premises software deployments and creating hybrid workspace services.

- Fast: Deploy apps and desktops, or complete secure digital workspaces in hours, not weeks.
- Adaptable: Choose to deploy on any cloud or virtual infrastructure — or a hybrid of both.
- Secure: Keep all proprietary information for your apps, desktops, and data under your control.
- Simple: Implement a fully-integrated Citrix portfolio via a single-management plane to simplify administration

Designing a XenDesktop Environment for a Mixed Workload

With Citrix XenDesktop 7.15, 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.

Server OS	**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

machines	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 mix of Windows Server 2016 based Hosted Shared Desktop sessions (RDS) and Windows 10 Hosted Virtual desktops (Statically assigned Persistent and Random Pooled) were configured and tested.

Deployment Hardware and Software

Products Deployed

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

The Citrix solution includes Cisco networking, Cisco UCS, and NetApp AFF storage, which efficiently fits into a single data center rack, including the access layer network switches.

This validated design document details the deployment of the multiple configurations extending to 6000 users for a mixed XenDesktop workload featuring the following software:

- Citrix XenApp 7.15 Hosted Shared Virtual Desktops (HSD) with PVS write cache on NFS storage
- Citrix XenDesktop 7.15 Non-Persistent Hosted Virtual Desktops (HVD) with PVS write cache on NFS storage
- Citrix XenDesktop 7.15 Persistent Hosted Virtual Desktops (VDI) provisioned with MCS and stored on NFS storage
- Citrix Provisioning Server 7.15
- Citrix User Profile Manager 7.15
- Citrix StoreFront 7.15
- VMware vSphere ESXi 6.5 Update 1 Hypervisor
- Microsoft Windows Server 2016 and Windows 10 64-bit virtual machine Operating Systems
- Microsoft SQL Server 2016

Figure 24 Virtual Desktop and Application Workload Architecture

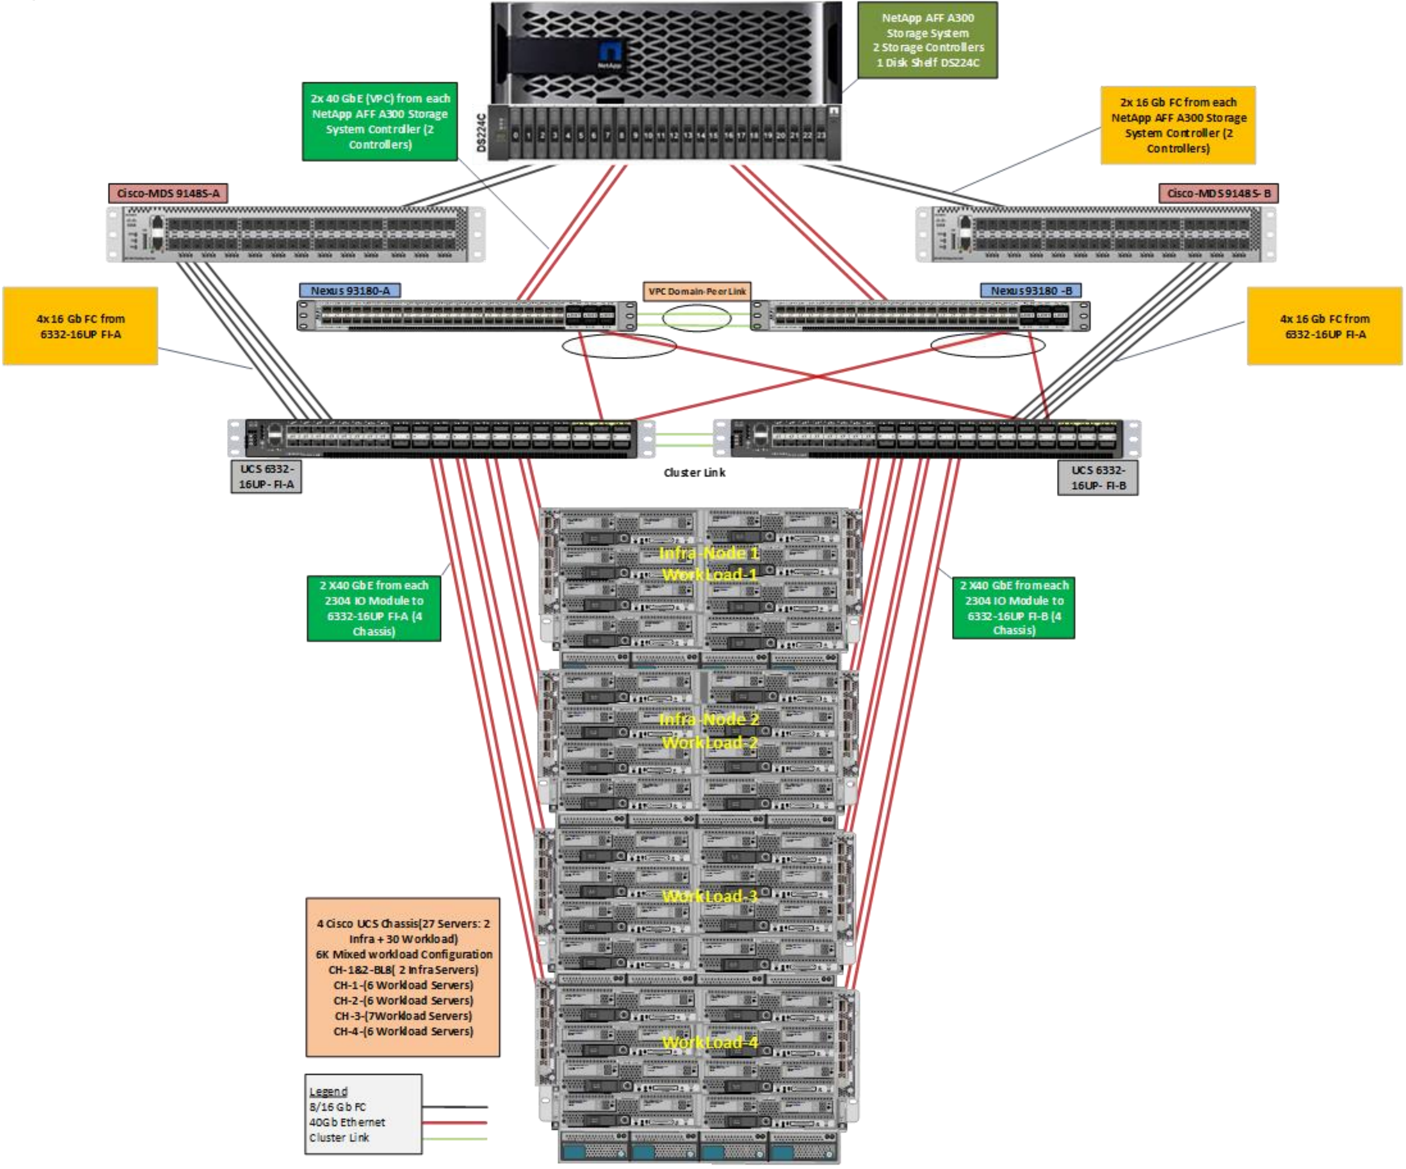
NetApp AFF A300
Storage System
2 Storage Controllers
1 Disk Shelf DS224C
DS224C
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
2x 40 GbE (VPC) from each NetApp AFF A300 Storage System Controller (2 Controllers)
2x 16 Gb FC from each NetApp AFF A300 Storage System Controller (2 Controllers)
Cisco-MDS 9148S-A
Cisco-MDS 9148S-B
Nexus 93180-A
VPC Domain-Peer Link
Nexus 93180 -B
4x 16 Gb FC from 6332-16UP FI-A
4x 16 Gb FC from 6332-16UP FI-A
UCS 6332-16UP - FI-A
Cluster Link
UCS 6332-16UP- FI-B
Infra-Node 1
WorkLoad-1
2 X40 GbE from each 2304 IO Module to 6332-16UP FI-A (4 Chassis)
2 X40 GbE from each 2304 IO Module to 6332-16UP FI-B (4 Chassis)
Infra-Node 2
WorkLoad-2
WorkLoad-3
WorkLoad-4
4 Cisco UCS Chassis(27 Servers: 2 Infra + 30 Workload)
6K Mixed workload Configuration
CH-1&2-BLB(2 Infra Servers)
CH-1-(6 Workload Servers)
CH-2-(6 Workload Servers)
CH-3-(7Workload Servers)
CH-4-(6 Workload Servers)
Legend
8/16 Gb FC
40Gb Ethernet
Cluster Link

The workload contains the following hardware as shown in Figure 24:

- Two Cisco Nexus 93180YC-FX Layer 2 Access Switches
- Four Cisco UCS 5108 Blade Server Chassis with two built-in UCS-IOM-2208XP IO Modules
- Two Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.60-GHz 10-core processors, 128GB 2133MHz RAM, and one Cisco VIC1340 mezzanine card for the hosted infrastructure, providing N+1 server fault tolerance
- Eight Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the Hosted Shared Desktop workload, providing N+1 server fault tolerance at the workload cluster level
- Eleven Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the Random Pooled desktops workload, providing N+1 server fault tolerance at the workload cluster level

- Nine Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the Static (Full Clones) desktops workload, providing N+1 server fault tolerance at the workload cluster level
- NetApp AFF A300 Storage System with dual redundant controllers, 1x DS224C disk shelf, and 24x 3.8 TB solid-state drives providing storage and FC/NFS/CIFS connectivity.

The LoginVSI Test infrastructure is not a part of this solution. The NetApp AFF300 configuration is detailed later in this document.

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 outlined in Figure 25.

Figure 25 Logical Architecture Overview

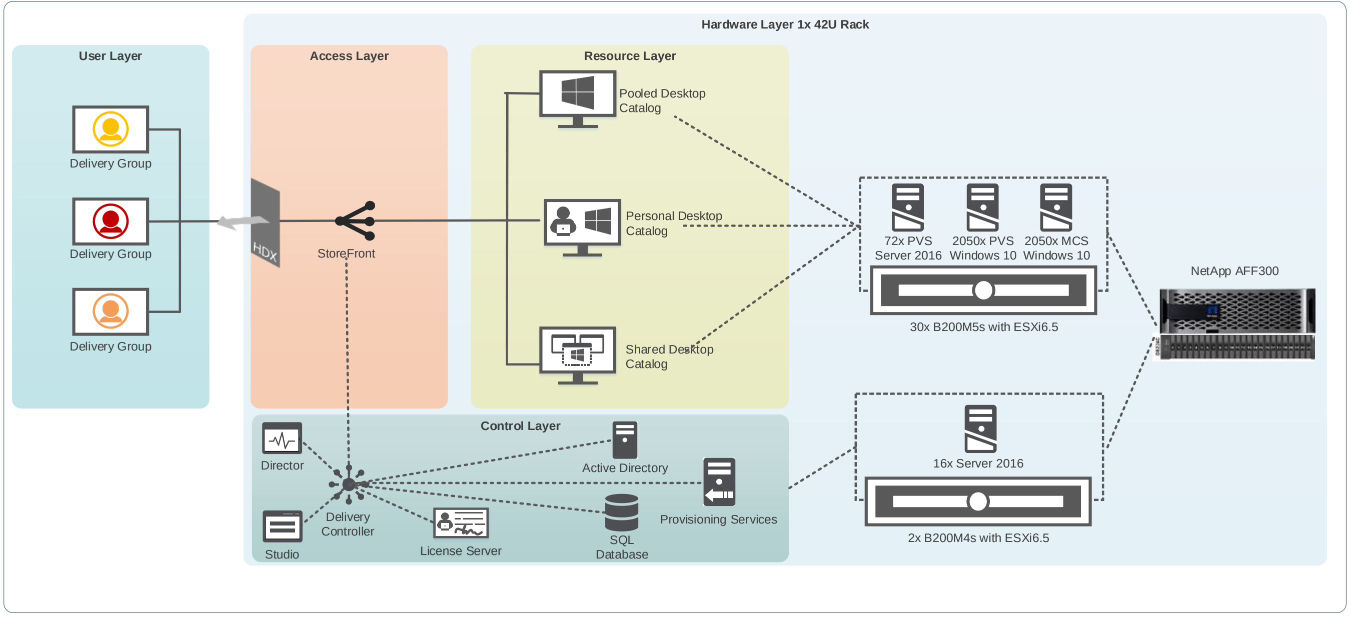

Software Revisions

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

Table 6 Software Revisions

Vendor	Product	Version
Cisco	UCS Component Firmware	3.2(3d) bundle release
Cisco	UCS Manager	3.2(3d) bundle release
Cisco	UCS B200 M4 Blades	3.2(3d) bundle release
Cisco	VIC 1340	4.1(1d)

Vendor	Product	Version
Cisco	UCS B200 M5 Blades	3.2(3d) bundle release
Cisco	UCS Performance Manager	2.0.3
Citrix	XenApp VDA	7.15
Citrix	XenDesktop VDA	7.15
Citrix	XenDesktop Controller	7.15
Citrix	Provisioning Services	7.15
Citrix	StoreFront Services	7.15
VMware	vCenter Server Appliance	6.5.0.5973321
VMware	vSphere ESXi 6.5 Update 1a	6.5.0.7967591
NetApp	Clustered Data ONTAP	9.3

Configuration Guidelines

The Citrix XenDesktop 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 Citrix XenDesktop 7.15 customer environment as a stand-alone solution.

VLANs

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

Table 7 VLAN Configuration

VLAN Name	VLAN ID	VLAN Purpose	VLAN Name
Default	1	Native VLAN	Default
In-Band-Mgmt	60	VLAN for in-band management	In-Band-Mgmt
Infra-Mgmt	61	VLAN for Virtual Infrastructure	Infra-Mgmt
CIFS	62	VLAN for CIFS traffic	CIFS

NFS	63	VLAN for Infrastructure NFS traffic	NFS
vMotion	66	VLAN for VMware vMotion	vMotion
VDI	102	VLAN for Desktop traffic	VDI
OB-Mgmt	164	VLAN for out-of-band management	OB-Mgmt

VMware Clusters

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

- VDI Cluster FlexPod Data Center with Cisco UCS
  - Infrastructure: Infra VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, XenDesktop Controllers, Provisioning Servers, and NetApp VSC, VSMs, etc.)
  - HSD-CLSTR: XenApp RDS VMs (Windows Server 2016 streamed with PVS)
  - HVDNonPersistent-CLSTR: XenDesktop VDI VMs (Windows 10 64-bit non-persistent virtual desktops streamed with PVS)
  - HVDPersistent-CLSTR: XenDesktop VDI VMs (Windows 10 64-bit persistent virtual desktops)
- VSI Launchersand Launcher Cluster
  - LVS-Launcher-CLSTR: Login VSI Cluster (The Login VSI launcher infrastructure was connected using the same set of switches and vCenter instance, but was hosted on separate storage and servers.)

Figure 26 vCenter Data Center and Clusters Deployed

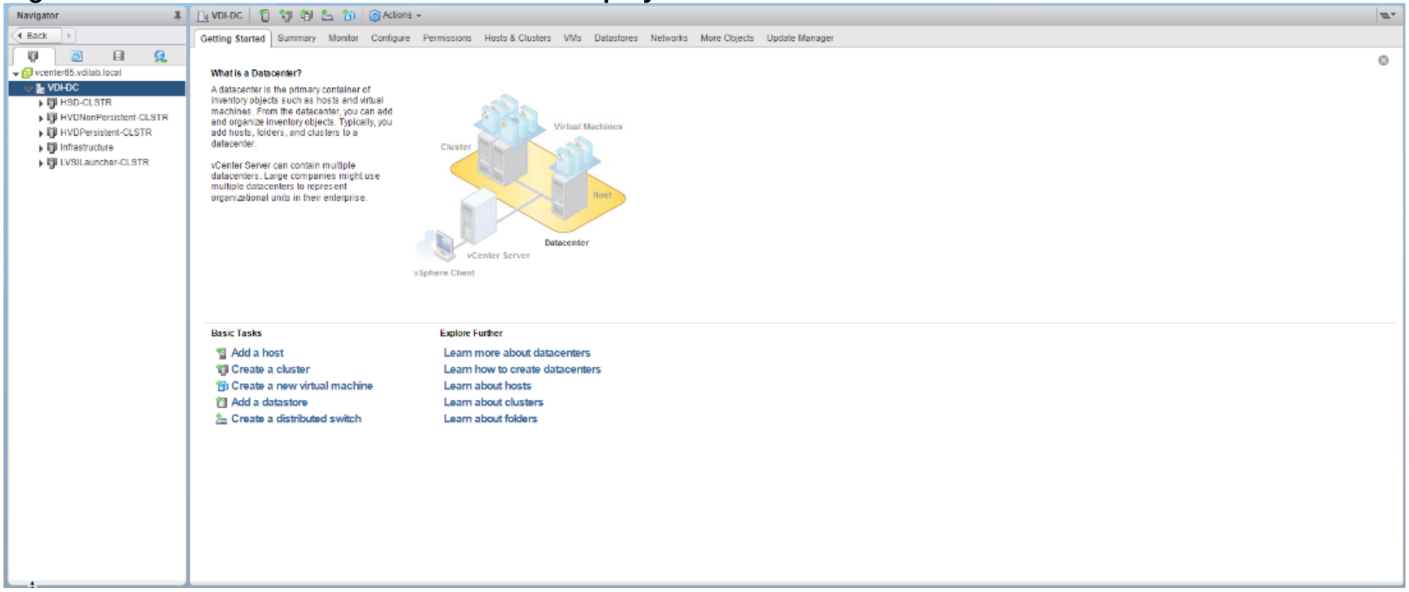

What is a Datacenter?

A datacenter is the primary container of inventory objects such as hosts and virtual machines. From the datacenter, you can add and organize inventory objects. Typically, you add hosts, folders, and clusters to a datacenter.

vCenter Server can contain multiple datacenters. Large companies might use multiple datacenters to represent organizational units in their enterprise.

Solution Configuration

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

Configuration Topology for a Scalable XenApp/XenDesktop 7.15 Mixed Workload Desktop Virtualization Solution

Figure 27 Component Layers for the FlexPod Data Center with Cisco UCS

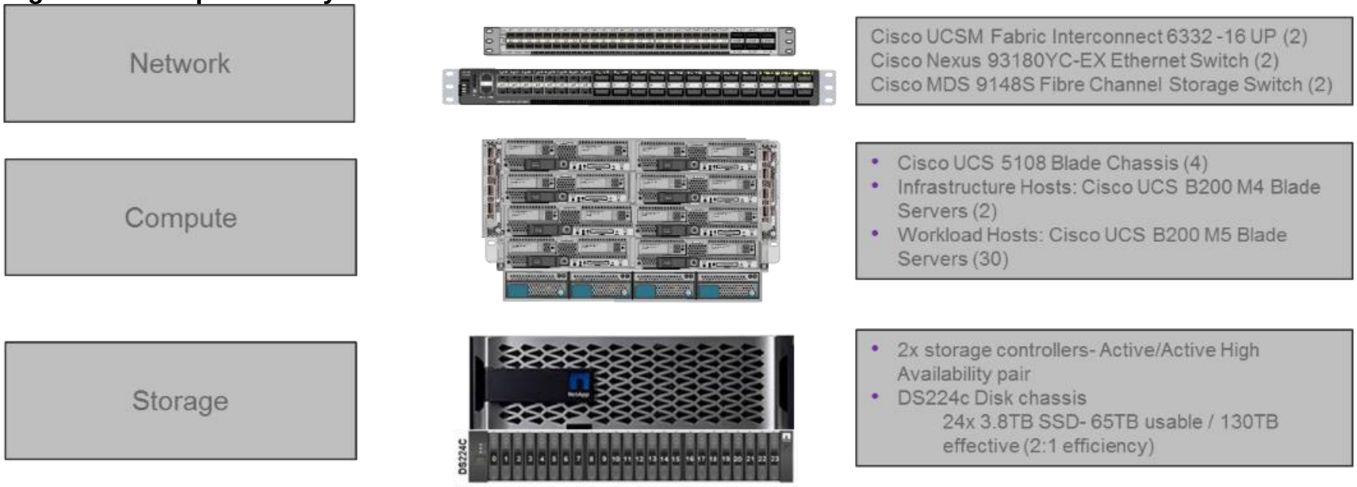

Network

Cisco UCSM Fabric Interconnect 6332 -16 UP (2)
Cisco Nexus 93180YC-EX Ethernet Switch (2)
Cisco MDS 9148S Fibre Channel Storage Switch (2)

Compute

- Cisco UCS 5108 Blade Chassis (4)
- Infrastructure Hosts: Cisco UCS B200 M4 Blade Servers (2)
- Workload Hosts: Cisco UCS B200 M5 Blade Servers (30)

Storage

- 2x storage controllers- Active/Active High Availability pair
- DS224c Disk chassis
  24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Figure 27 captures the architectural diagram for the purpose of this study. The architecture is divided into three distinct layers:

- Cisco UCS Compute Platform
- Network Access layer and LAN
- Storage Access to the NetApp AFF300

Figure 28 illustrates the physical connectivity configuration of the Citrix XenDesktop 7.15 environment.

Figure 28 Cabling Diagram of the FlexPod Datacenter with Cisco UCS

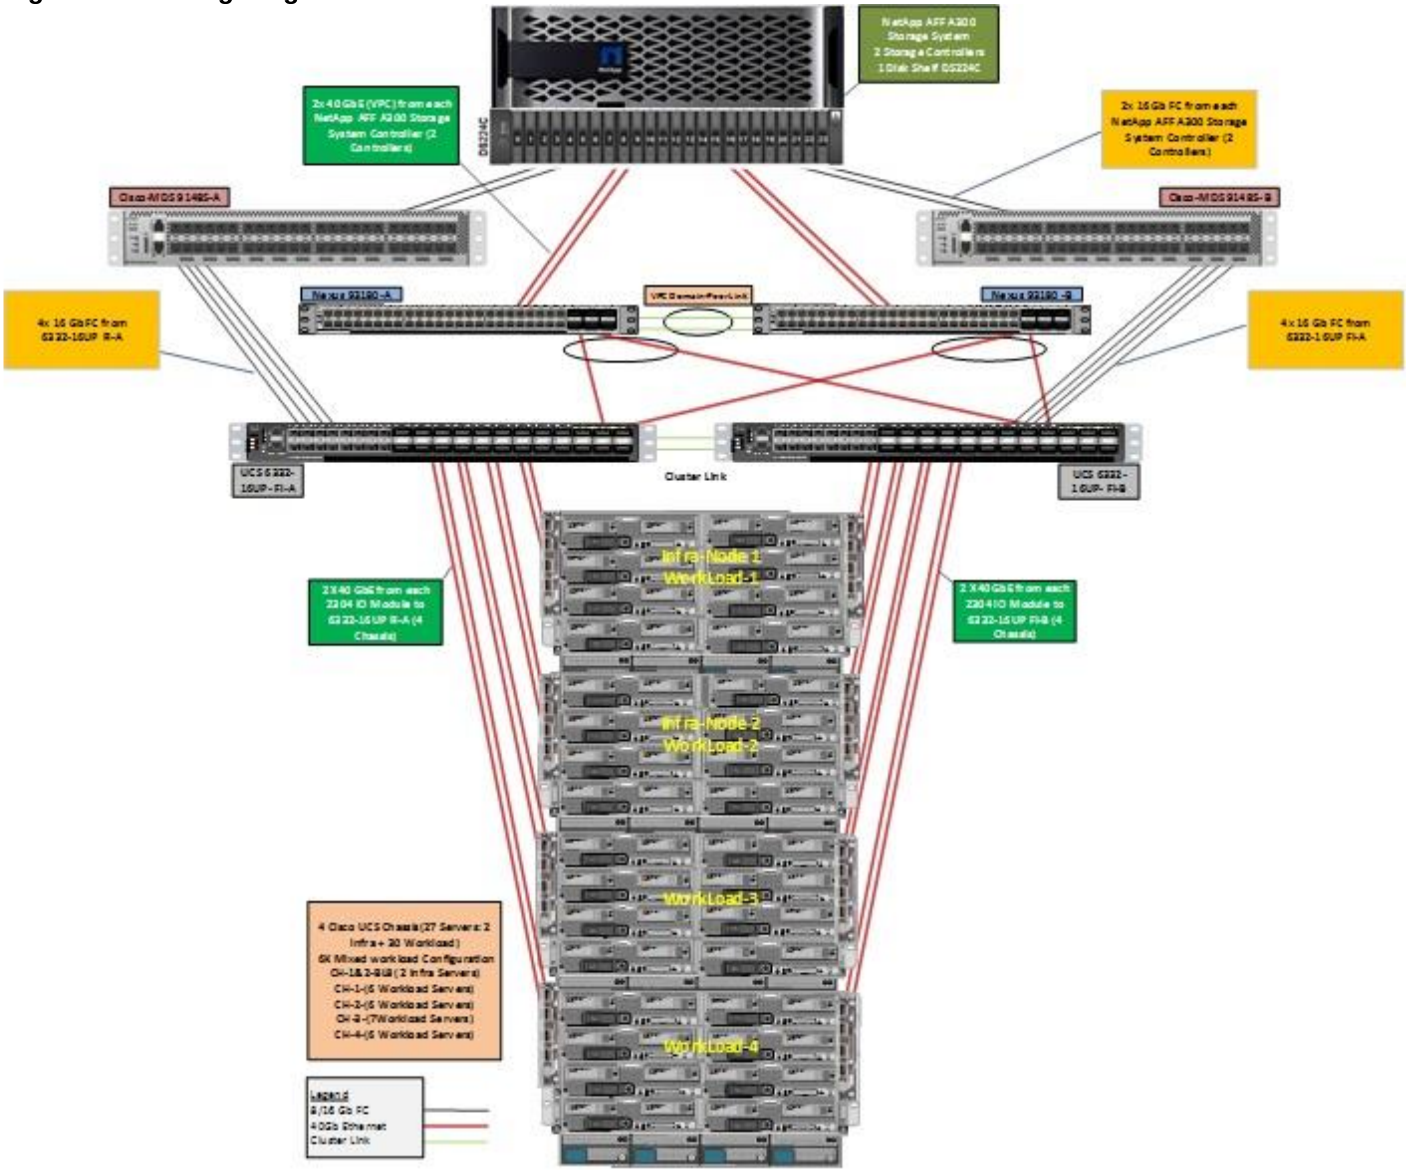
NetApp AFF A300 Storage System
2 Storage Controllers
1 Disk Shelf DS224C
DS224C
2x 40 GbE (VPC) from each NetApp AFF A300 Storage System Controller (2 Controllers)
2x 16 Gb FC from each NetApp AFF A300 Storage System Controller (2 Controllers)
Cisco-MDS 9148S-A
Cisco-MDS 9148S-B
4x 16 Gb FC from 6332-16UP FI-A
4x 16 Gb FC from 6332-16UP FI-A
Nexus 93180 -A
Nexus 93180 -B
VPC Domain-PeerLink
UCS 6332-16UP - FI-A
UCS 6332-16UP- FI-B
Cluster Link
2X40 GbE from each 2204 IO Module to 6332-16UP FI-A (4 Chassis)
2 X40 GbE from each 2204 IO Module to 6332-16UP FI-B (4 Chassis)
Infra-Node 1
WorkLoad-1
Infra-Node 2
WorkLoad-2
WorkLoad-3
WorkLoad-4
4 Cisco UCS Chassis (27 Servers: 2 Infra + 30 Workload)
6K Mixed workload Configuration
CH-1& 2-B8 (2 Infra Servers)
CH-1-(6 Workload Servers)
CH-2-(6 Workload Servers)
CH-3-(7Workload Servers)
CH-4-(6 Workload Servers)
Legend
8/16 Gb FC
40Gb Ethernet
Cluster Link

Table 8 through Table 13 list the details of all the connections in use.

Table 8   Cisco Nexus A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 A	Eth1/49	40GbE	Cisco Nexus B	Eth1/49
	Eth1/50	40GbE	Cisco Nexus B	Eth1/50
	Eth1/51	40GbE	Cisco UCS fabric interconnect A	Eth1/35
	Eth1/52	40GbE	Cisco UCS fabric interconnect B	Eth1/36
	Eth1/53	40GbE	NetApp Controller 1	e2a
	Eth1/54	40GbE	NetApp Controller 2	e2e

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

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

Table 9 Cisco Nexus B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 B	Eth1/49	40GbE	Cisco Nexus A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus A	Eth1/50
	Eth1/51	40GbE	Cisco UCS fabric interconnect B	Eth1/35
	Eth1/52	40GbE	Cisco UCS fabric interconnect A	Eth1/36
	Eth1/53	40GbE	NetApp Controller 1	e2e
	Eth1/54	40GbE	NetApp Controller 2	e2a
	MGMT0	GbE	GbE management switch	Any

Table 10 Cisco MDS A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148-A	FC1/37	16Gb	NetApp Controller 1	e0g
	FC1/38	16Gb	NetApp Controller 2	e0g
	FC1/43	16Gb	Cisco UCS fabric interconnect A	FC1/1
	FC1/44	16Gb	Cisco UCS fabric interconnect A	FC1/2
	FC1/45	16Gb	Cisco UCS fabric interconnect A	FC1/3
	FC1/46	16Gb	Cisco UCS fabric interconnect A	FC1/4

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

Table 11 Cisco MDS B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148-B	FC1/37	16Gb	NetApp Controller 1	e0h
	FC1/38	16Gb	NetApp Controller 2	e0h
	FC1/43	16Gb	Cisco UCS fabric interconnect B	FC1/1
	FC1/44	16Gb	Cisco UCS fabric interconnect B	FC1/2
	FC1/45	16Gb	Cisco UCS fabric interconnect B	FC1/3

Local Device	Local Port	Connection	Remote Device	Remote Port
	FC1/46	16Gb	Cisco UCS fabric interconnect B	FC1/4

Table 12 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6332-16 A	FC1/1	16Gb	Cisco MDS9148-A	FC1/43
	FC1/2	16Gb	Cisco MDS9148-A	FC1/44
	FC1/3	16Gb	Cisco MDS9148-A	FC1/45
	FC1/4	16Gb	Cisco MDS9148-A	FC1/46
	Eth1/17	40GbE	Cisco UCS Chassis1 FEX A	IOM 1/1
	Eth1/18	40GbE	Cisco UCS Chassis1 FEX A	IOM 1/2
	Eth1/19	40GbE	Cisco UCS Chassis2 FEX A	IOM 1/1
	Eth1/20	40GbE	Cisco UCS Chassis2 FEX A	IOM 1/2
	Eth1/21	40GbE	Cisco UCS Chassis3 FEX A	IOM 1/1
	Eth1/22	40GbE	Cisco UCS Chassis3 FEX A	IOM 1/2
	Eth1/23	40GbE	Cisco UCS Chassis4 FEX A	IOM 1/1
	Eth1/24	40GbE	Cisco UCS Chassis4 FEX A	IOM 1/2
	Eth1/35	40GbE	Cisco Nexus 93180 A	Eth1/51
	Eth1/36	40GbE	Cisco Nexus 93180 B	Eth1/52
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 13 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6332-16 B	FC1/1	16Gb	Cisco MDS9148-B	FC1/43
	FC1/2	16Gb	Cisco MDS9148-B	FC1/44
	FC1/3	16Gb	Cisco MDS9148-B	FC1/45
	FC1/4	16Gb	Cisco MDS9148-B	FC1/46
	Eth1/17	40GbE	Cisco UCS Chassis1 FEX B	IOM 1/1
	Eth1/18	40GbE	Cisco UCS Chassis1 FEX B	IOM 1/2
	Eth1/19	40GbE	Cisco UCS Chassis2 FEX B	IOM 1/1
	Eth1/20	40GbE	Cisco UCS Chassis2 FEX B	IOM 1/2
	Eth1/21	40GbE	Cisco UCS Chassis3 FEX B	IOM 1/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/22	40GbE	Cisco UCS Chassis3 FEX B	IOM 1/2
	Eth1/23	40GbE	Cisco UCS Chassis4 FEX B	IOM 1/1
	Eth1/24	40GbE	Cisco UCS Chassis4 FEX B	IOM 1/2
	Eth1/35	40GbE	Cisco Nexus 93180 B	Eth1/51
	Eth1/36	40GbE	Cisco Nexus 93180 A	Eth1/52
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Cisco Unified Computing System 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 Installation guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) and it is beyond the scope of this document.

For more information about each step, refer to the following documents: Cisco UCS Manager Configuration Guides – GUI and Command Line Interface (CLI) Cisco UCS Manager - Configuration Guides - Cisco

Cisco UCS Manager Software Version 3.2(3c)

This document assumes the use of Cisco UCS Manager Software version 3.2(3c). To upgrade the Cisco UCS Manager software and the Cisco UCS 6332-16UP 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 interconnect, complete the following steps:

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

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

   a. Login with the existing username and password
   b. Enter: connect local-mgmt
   c. Enter: erase config
   d. Enter: yes to confirm

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

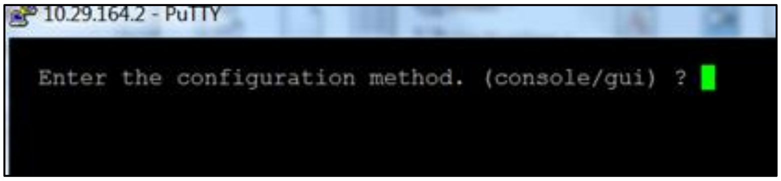

```
Enter the configuration method. (console/gui) ?
```

4. Type “setup” at the setup/restore prompt, then press Enter.

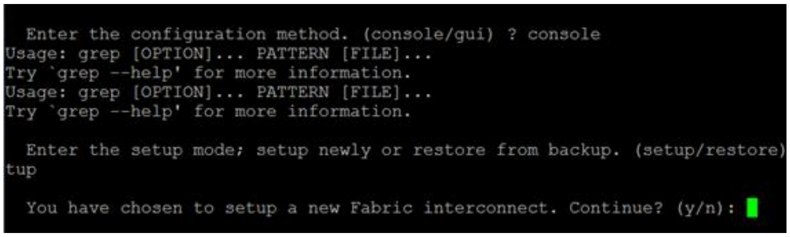

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  Enter the setup mode; setup newly or restore from backup. (setup/restore)
tup

  You have chosen to setup a new Fabric interconnect. Continue? (y/n):
```

5. Type “y” then press Enter to confirm the setup.

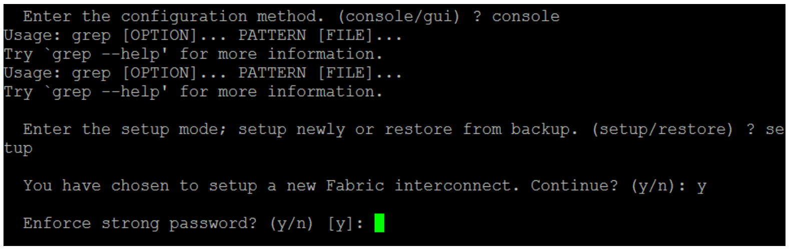

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  Enter the setup mode; setup newly or restore from backup. (setup/restore) ? se
tup

  You have chosen to setup a new Fabric interconnect. Continue? (y/n): y

  Enforce strong password? (y/n) [y]:
```

6. Type “y” or “n” depending on your organization’s security policies, then press Enter.

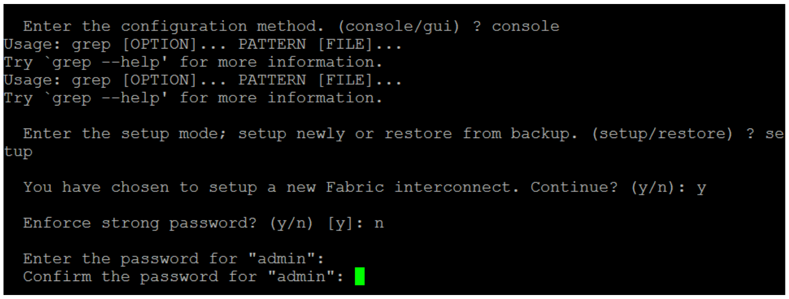

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  Enter the setup mode; setup newly or restore from backup. (setup/restore) ? se
tup

  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":
```

7. Enter and confirm the password and enter switch Fabric A.

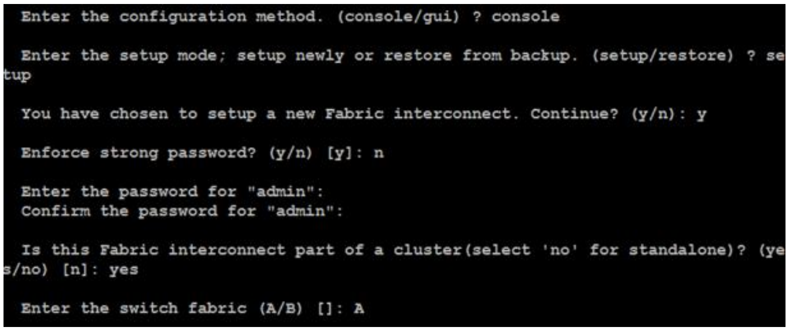

```
  Enter the configuration method. (console/gui) ? console

  Enter the setup mode; setup newly or restore from backup. (setup/restore) ? se
tup

  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)? (ye
s/no) [n]: yes

  Enter the switch fabric (A/B) []: A
```

8. Complete the setup dialog questions.

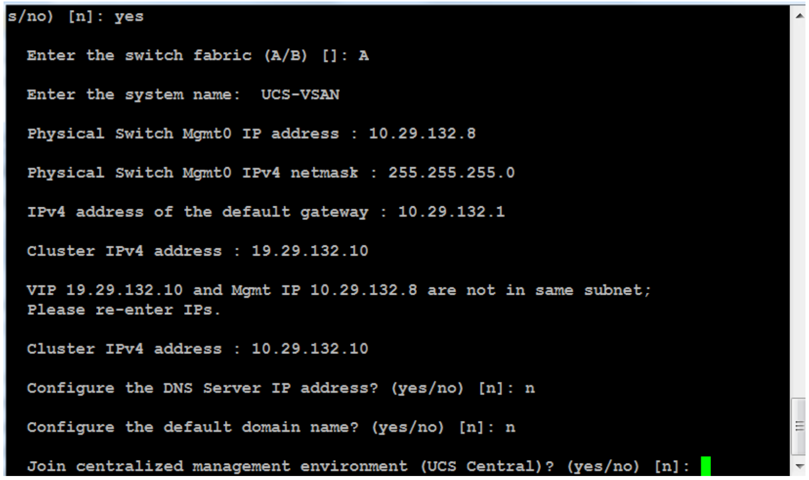

```
s/no) [n]: yes

  Enter the switch fabric (A/B) []: A

  Enter the system name:  UCS-VSAN

  Physical Switch Mgmt0 IP address : 10.29.132.8

  Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

  IPv4 address of the default gateway : 10.29.132.1

  Cluster IPv4 address : 19.29.132.10

  VIP 19.29.132.10 and Mgmt IP 10.29.132.8 are not in same subnet;
  Please re-enter IPs.

  Cluster IPv4 address : 10.29.132.10

  Configure the DNS Server IP address? (yes/no) [n]: n

  Configure the default domain name? (yes/no) [n]: n

  Join centralized management environment (UCS Central)? (yes/no) [n]:
```

9. Review the selections and type “yes”.

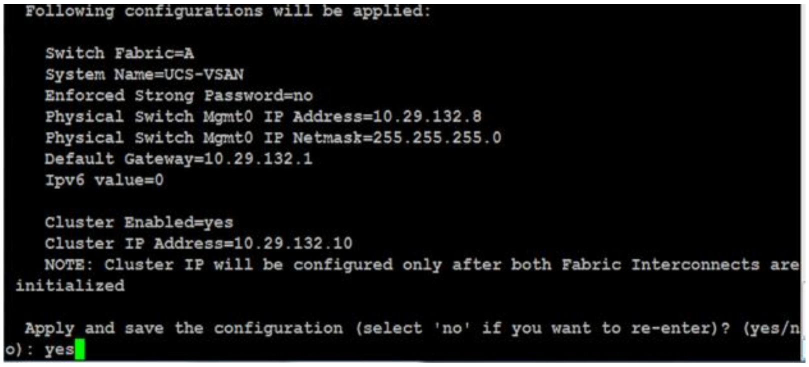

```
  Following configurations will be applied:

    Switch Fabric=A
    System Name=UCS-VSAN
    Enforced Strong Password=no
    Physical Switch Mgmt0 IP Address=10.29.132.8
    Physical Switch Mgmt0 IP Netmask=255.255.255.0
    Default Gateway=10.29.132.1
    Ipv6 value=0

    Cluster Enabled=yes
    Cluster IP Address=10.29.132.10
    NOTE: Cluster IP will be configured only after both Fabric Interconnects are
 initialized

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

10. Console on to the second fabric interconnect, select console as the configuration method, and provide the following inputs.

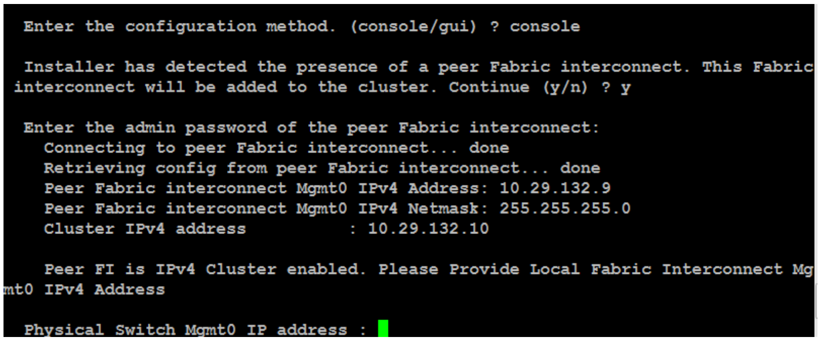

```
  Enter the configuration method. (console/gui) ? console

  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: 10.29.132.9
    Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
    Cluster IPv4 address          : 10.29.132.10

    Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mg
mt0 IPv4 Address

  Physical Switch Mgmt0 IP address :
```

11. Open a web browser and go to the Virtual IP address configured above.

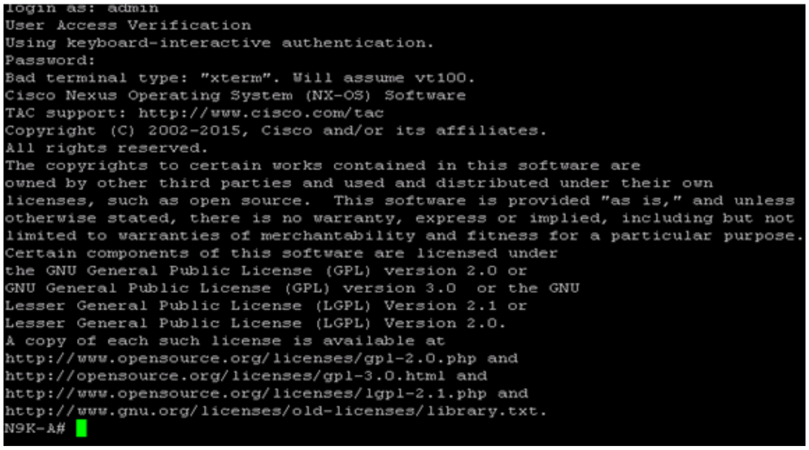

```
login as: admin
User Access Verification
Using keyboard-interactive authentication.
Password:
Bad terminal type: "xterm". Will assume vt100.
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2015, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source.  This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0  or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
N9K-A#
```

Base Cisco UCS System Configuration

To configure the Cisco Unified Computing System, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6332-16UP Fabric Interconnect cluster address.

2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.

3. If prompted to accept security certificates, accept as necessary.

4. When prompted, enter admin as the username and enter the administrative password.

5. To log in to Cisco UCS Manager, click Login.

Set Fabric Interconnects to Fibre Channel End Host Mode

To set the Fabric Interconnects to the Fibre Channel End Host Mode, complete the following steps:

1. On the Equipment tab, expand the Fabric Interconnects node and click Fabric Interconnect A.

2. On the General tab in the Actions pane, click Set FC End Host mode.

3. Follow the dialogs to complete the change.

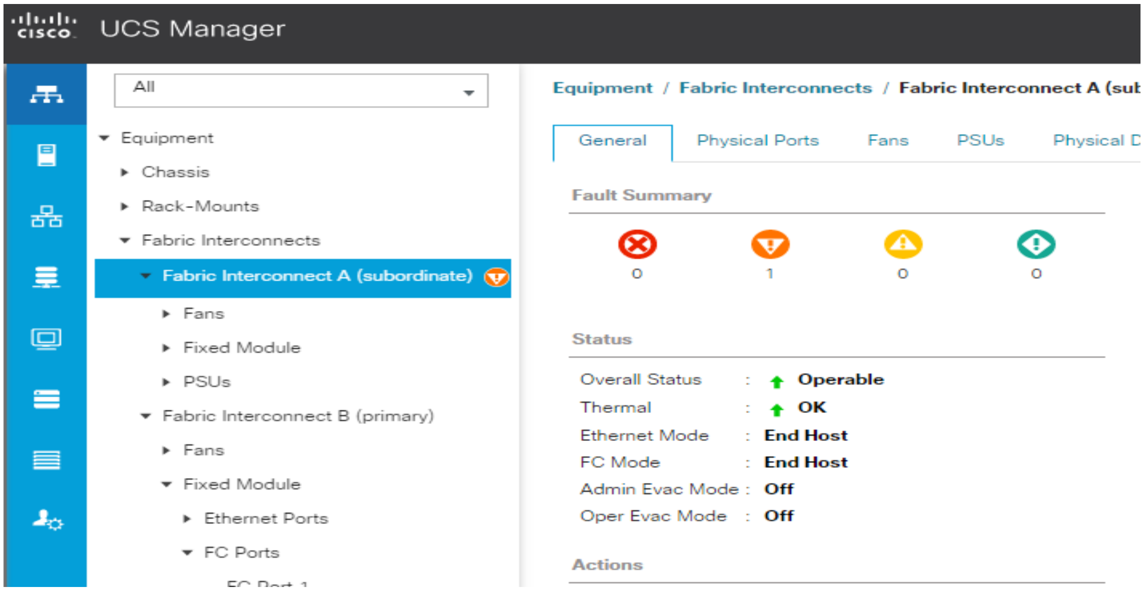

Overall Status : Operable
Thermal : OK
Ethernet Mode : End Host
FC Mode : End Host
Admin Evac Mode : Off
Oper Evac Mode : Off

Both Fabric Interconnects automatically reboot sequentially when you confirm you want to operate in this mode.

Configure Fibre Channel Uplink Ports

To configure the Fibre Channel Uplink Ports, complete the following steps:

1. After the restarts are complete, from the General tab, Actions pane, click Configure Unified ports.

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

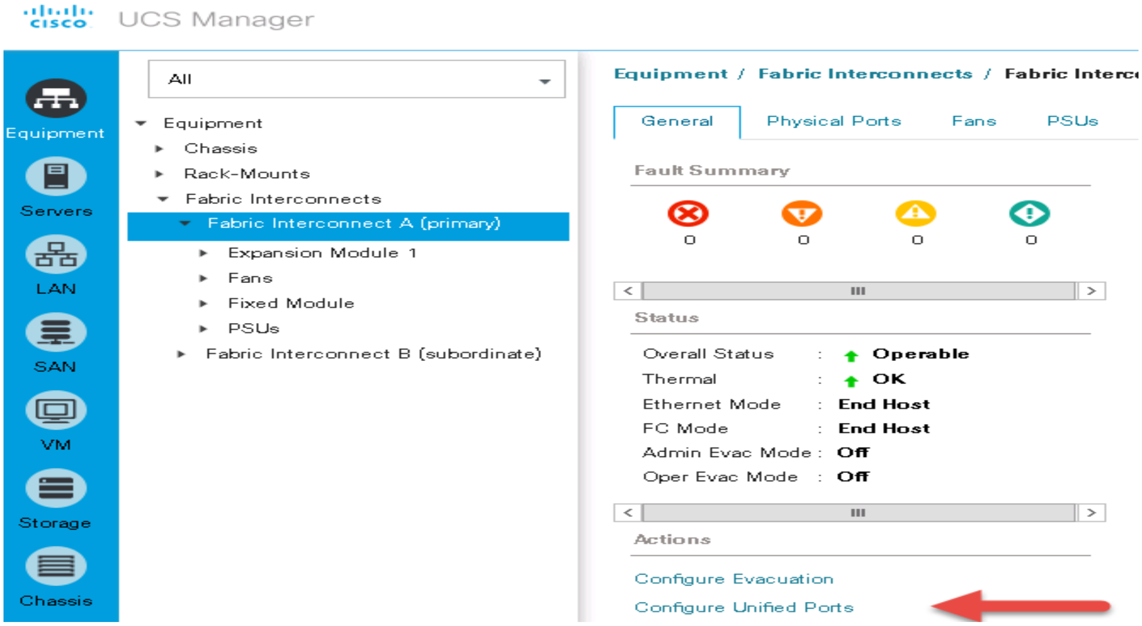

Overall Status : Operable
Thermal : OK
Ethernet Mode : End Host
FC Mode : End Host
Admin Evac Mode : Off
Oper Evac Mode : Off

3. Move the slider to the right to configure first 6 ports as FC ports

Ports to the left of the slider will become FC ports. For our study, we configured the 6 ports on the as FC ports.

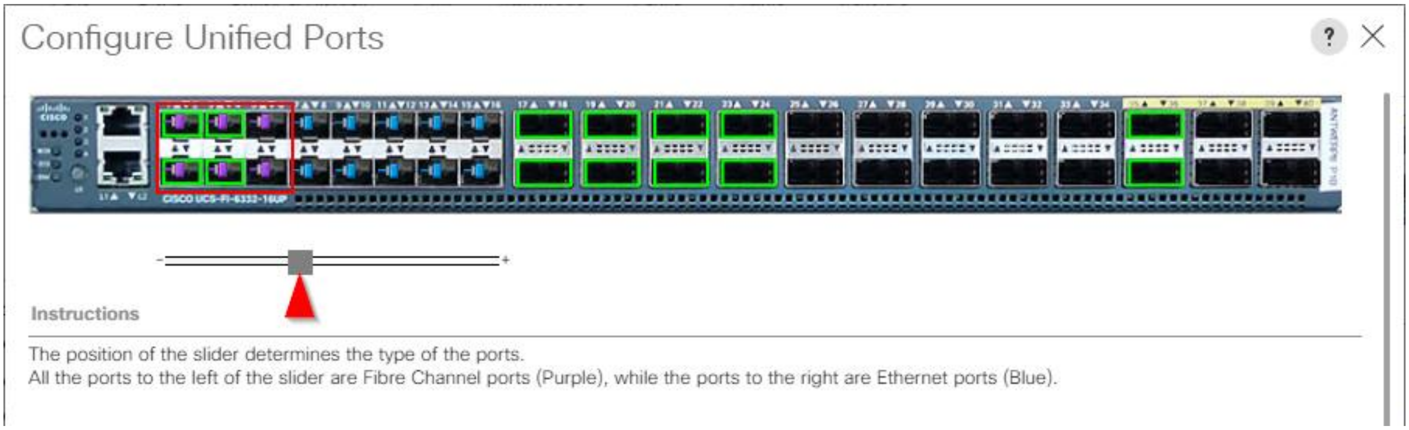

Configure Unified Ports

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).

4. Click OK, then click Yes to confirm. This action will cause a reboot of the Fabric Interconnect.

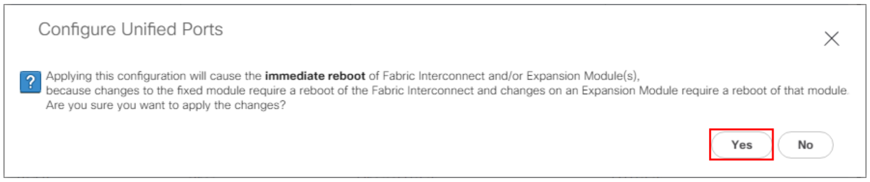

Configure Unified Ports

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?

After the reboot, your FC Ports configuration should look like the screenshot below:

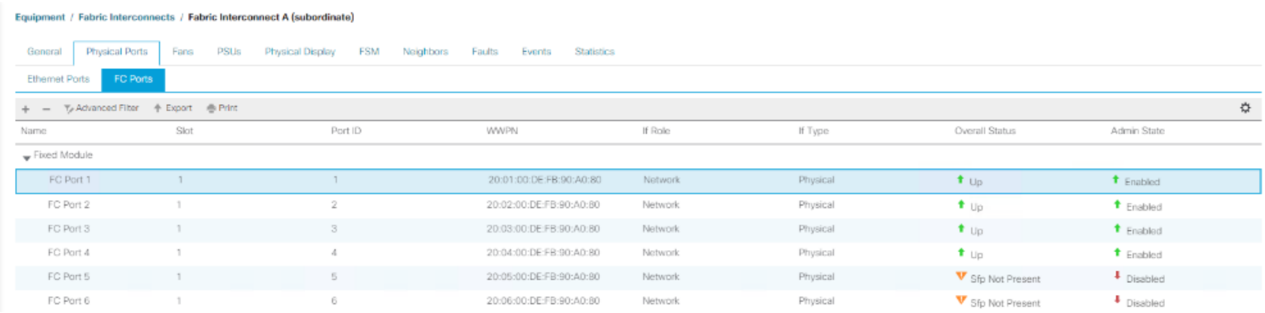

Equipment / Fabric Interconnects / **Fabric Interconnect A (subordinate)**

Name	Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
Fixed Module							
FC Port 1	1	1	20:01:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
FC Port 2	1	2	20:02:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
FC Port 3	1	3	20:03:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
FC Port 4	1	4	20:04:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
FC Port 5	1	5	20:05:00:DE:FB:90:A0:80	Network	Physical	▼ Sfp Not Present	↓ Disabled
FC Port 6	1	6	20:06:00:DE:FB:90:A0:80	Network	Physical	▼ Sfp Not Present	↓ Disabled

5. Repeat this procedure for Fabric Interconnect B.

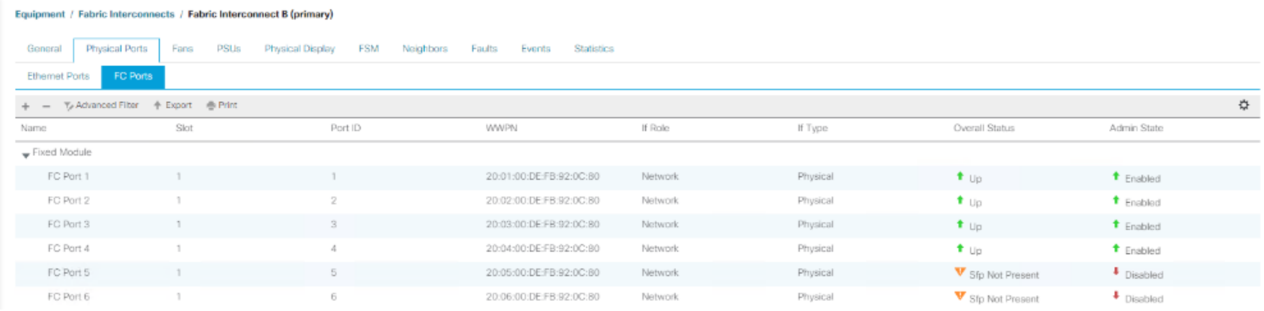

Equipment / Fabric Interconnects / **Fabric Interconnect B (primary)**

Name	Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
Fixed Module							
FC Port 1	1	1	20:01:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
FC Port 2	1	2	20:02:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
FC Port 3	1	3	20:03:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
FC Port 4	1	4	20:04:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
FC Port 5	1	5	20:05:00:DE:FB:92:0C:80	Network	Physical	▼ Sfp Not Present	↓ Disabled
FC Port 6	1	6	20:06:00:DE:FB:92:0C:80	Network	Physical	▼ Sfp Not Present	↓ Disabled

6. Insert Cisco SFP 16 Gbps FC (DS-SFP-FC16-SW) modules into ports 1 through 4 on both Fabric Interconnects and cable as prescribed later in this document.

Four FC Uplinks from Fabric A are connected to Cisco MDS 9148S switch A and four FC Uplinks from Fabric B to Cisco MDS 9148S switch B.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis.

To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment node and select Equipment in the list on the left.

2. In the right pane, click the Policies tab.

3. Under Global Policies, set the Chassis/FEX Discovery Policy to 2-link.

4. Set the Link Grouping Preference to Port Channel.

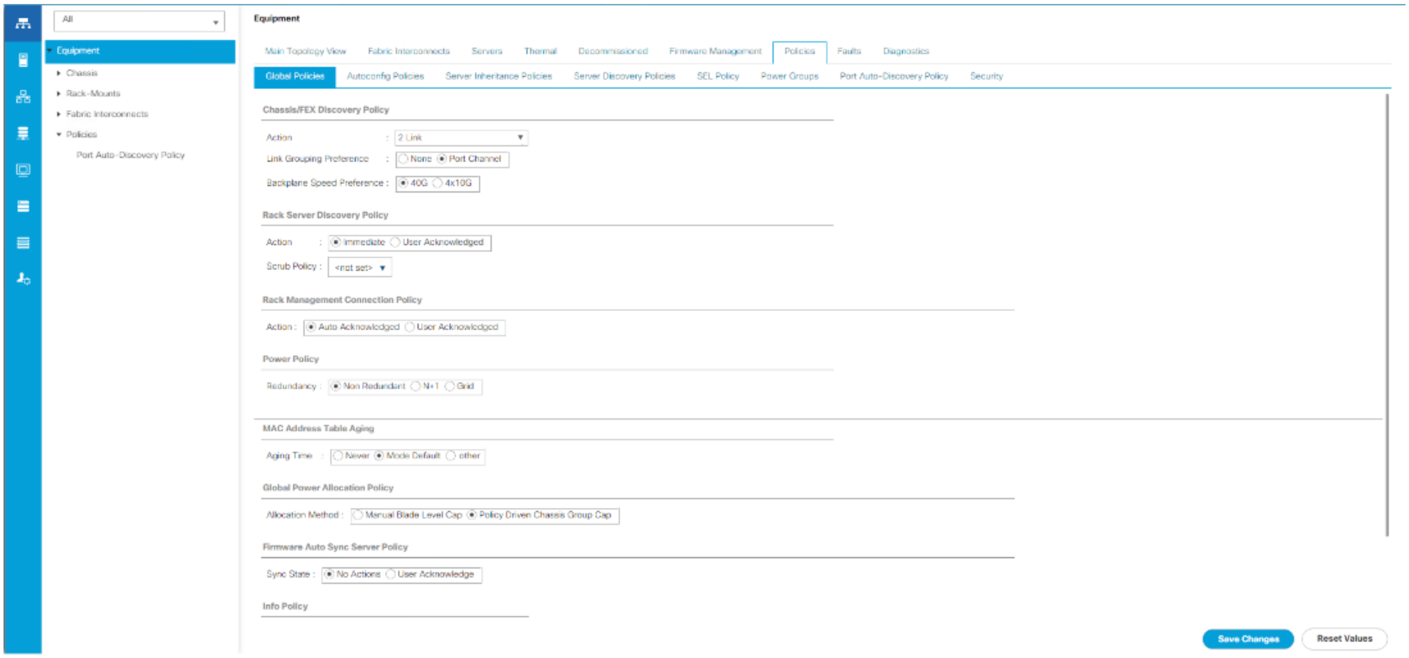

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : ○ None ● Port Channel

Backplane Speed Preference : ● 40G ○ 4x10G

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

Info Policy

5. Click Save Changes.

6. Click OK.

Acknowledge Cisco UCS Chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

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

2. Expand Chassis and select each chassis that is listed.

3. Right-click each chassis and select Acknowledge Chassis.

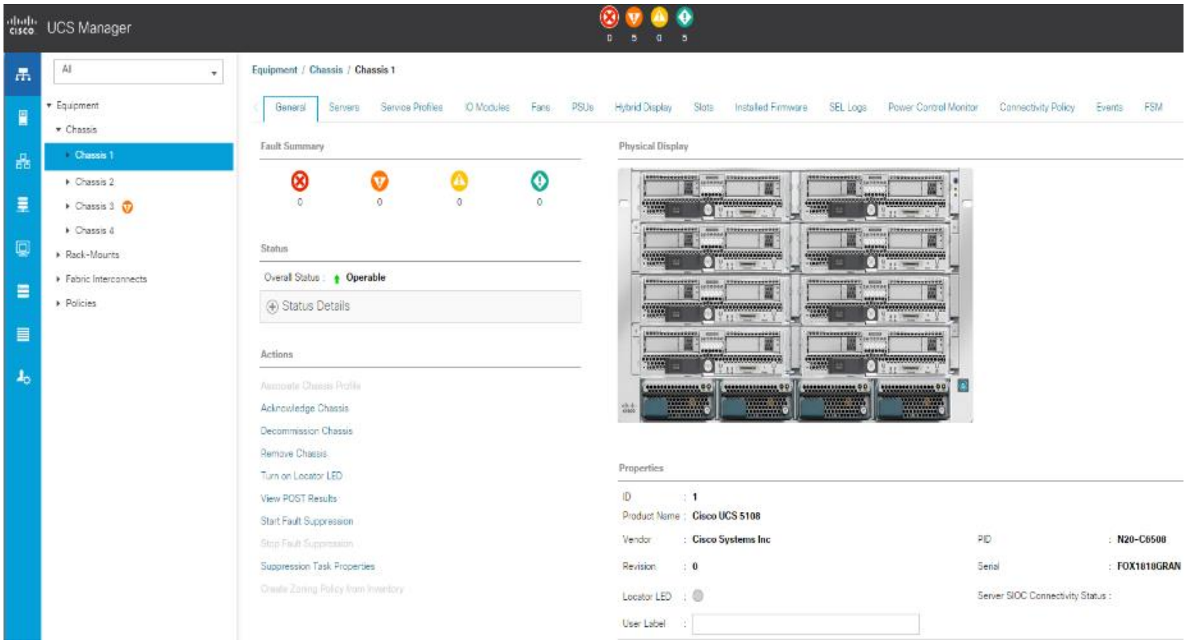

Equipment / Chassis / Chassis 1

Fault Summary

Status

Overall Status : Operable

Actions

Physical Display

Properties

ID : 1

Product Name : Cisco UCS 5108

Vendor : Cisco Systems Inc

PID : N20-C6508

Revision : 0

Serial : FOX1810GRAN

Locator LED :

Server SIOC Connectivity Status :

User Label :

4. Click Yes and then click OK to complete acknowledging the chassis.

5. Repeat for each of the remaining chassis.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

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

2. Select All > Timezone Management.

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

4. Click Save Changes and then click OK.

5. Click Add NTP Server.

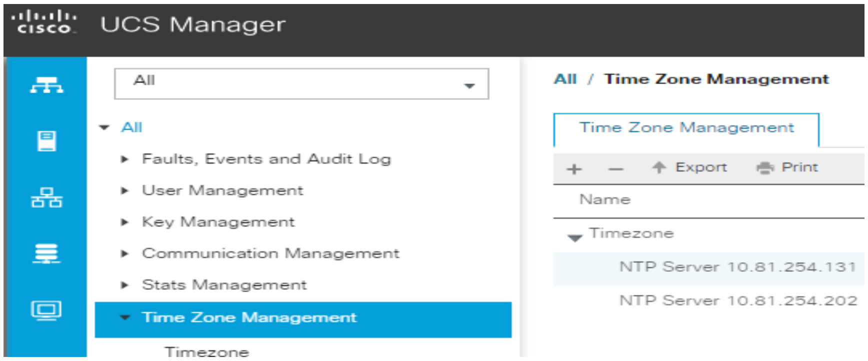

All / Time Zone Management

Name
Timezone
NTP Server 10.81.254.131
NTP Server 10.81.254.202

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

7. Click OK.

Enable Server and Ethernet Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Physical Ports > Ethernet Ports.
3. Expand Fixed Module.
4. Select ports 17 through 24 that are connected to the Cisco IO Modules of the four B-Series 5108 Chassis, right-click them and select Configure as Server Port.
5. Click Yes to confirm uplink ports and click OK.
6. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that ports have been configured correctly in the in the Role column.

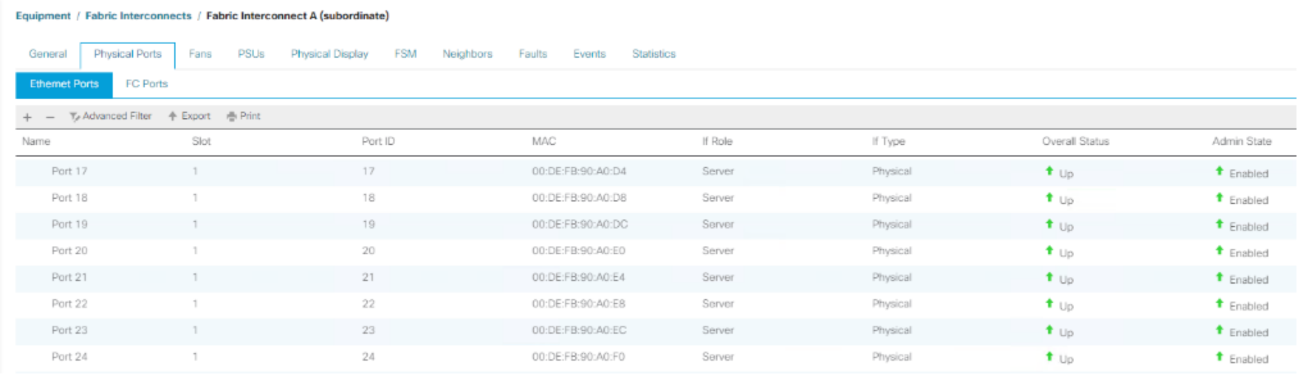

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 17	1	17	00:DE:FB:90:A0:D4	Server	Physical	Up	Enabled
Port 18	1	18	00:DE:FB:90:A0:D8	Server	Physical	Up	Enabled
Port 19	1	19	00:DE:FB:90:A0:DC	Server	Physical	Up	Enabled
Port 20	1	20	00:DE:FB:90:A0:E0	Server	Physical	Up	Enabled
Port 21	1	21	00:DE:FB:90:A0:E4	Server	Physical	Up	Enabled
Port 22	1	22	00:DE:FB:90:A0:E8	Server	Physical	Up	Enabled
Port 23	1	23	00:DE:FB:90:A0:EC	Server	Physical	Up	Enabled
Port 24	1	24	00:DE:FB:90:A0:F0	Server	Physical	Up	Enabled

7. Repeat the above steps for Fabric Interconnect B. The screenshot below shows the server ports for Fabric B.

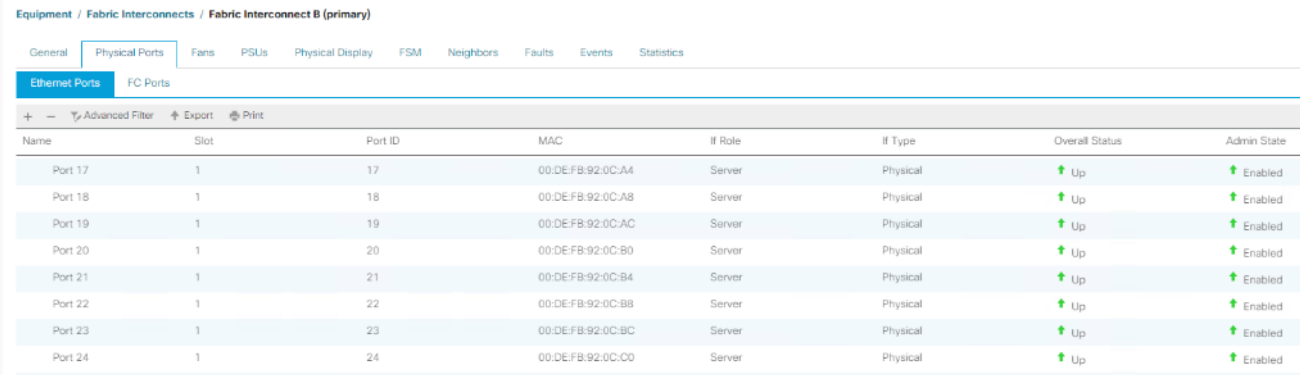

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 17	1	17	00:DE:FB:92:0C:A4	Server	Physical	Up	Enabled
Port 18	1	18	00:DE:FB:92:0C:A8	Server	Physical	Up	Enabled
Port 19	1	19	00:DE:FB:92:0C:AC	Server	Physical	Up	Enabled
Port 20	1	20	00:DE:FB:92:0C:B0	Server	Physical	Up	Enabled
Port 21	1	21	00:DE:FB:92:0C:B4	Server	Physical	Up	Enabled
Port 22	1	22	00:DE:FB:92:0C:B8	Server	Physical	Up	Enabled
Port 23	1	23	00:DE:FB:92:0C:BC	Server	Physical	Up	Enabled
Port 24	1	24	00:DE:FB:92:0C:C0	Server	Physical	Up	Enabled

To configure network ports used to uplink the Fabric Interconnects to the Cisco Nexus 9172PX switches, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Physical Ports > Ethernet Ports.
3. Expand Fixed Module.
4. Select ports 35 through 36 that are connected to the Nexus 93180YC-FX switches, right-click them and select Configure as Uplink Port.

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. Successful configuration should result in ports 35-36 configured as network ports as shown in the screenshot below:

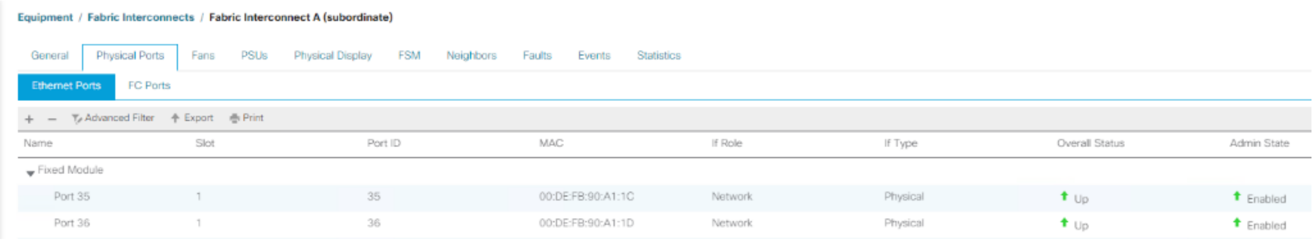

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
▾ Fixed Module							
Port 35	1	35	00:DE:FB:90:A1:1C	Network	Physical	↑ Up	↑ Enabled
Port 36	1	36	00:DE:FB:90:A1:1D	Network	Physical	↑ Up	↑ Enabled

8. Repeat the above steps for Fabric Interconnect B. The screenshot shows the network uplink ports for Fabric B.

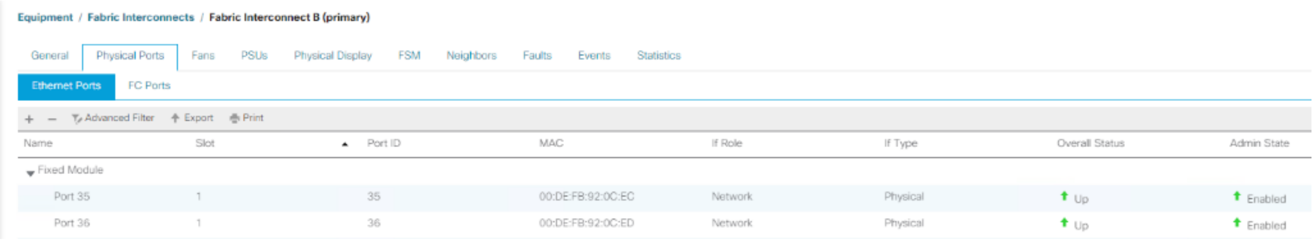

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
▾ Fixed Module							
Port 35	1	35	00:DE:FB:92:0C:EC	Network	Physical	↑ Up	↑ Enabled
Port 36	1	36	00:DE:FB:92:0C:ED	Network	Physical	↑ Up	↑ Enabled

Create Uplink Port Channels to Cisco Nexus 93180YC-FX Switches

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

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

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

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

3. Right-click Port Channels.

4. Select Create Port Channel.

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

6. Enter FI-A-Uplink as the name of the port channel.

7. Click Next.

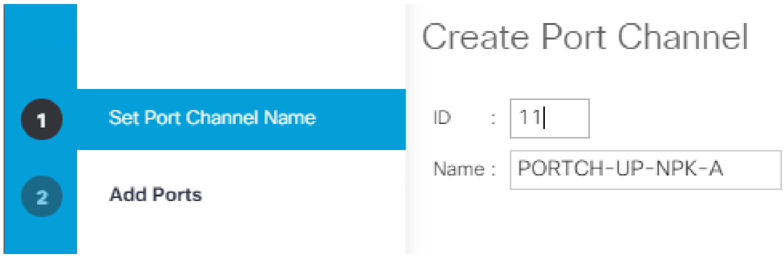

Create Port Channel

ID : 11

Name : PORTCH-UP-NPK-A

8. Select Ethernet ports 35-36 for the port channel.

9. Click Finish.

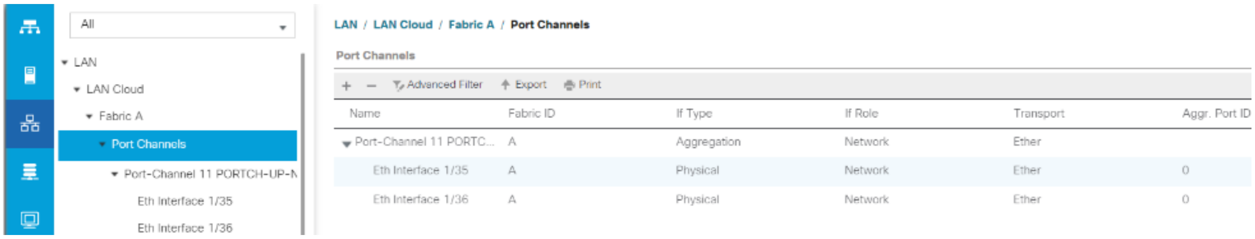

LAN / LAN Cloud / Fabric A / Port Channels

Port Channels

Name	Fabric ID	If Type	If Role	Transport	Aggr. Port ID
Port-Channel 11 PORTC...	A	Aggregation	Network	Ether	
Eth Interface 1/35	A	Physical	Network	Ether	0
Eth Interface 1/36	A	Physical	Network	Ether	0

Repeat steps 1-9 for Fabric Interconnect B, substituting 12 for the port channel number and FI-B-Uplink for the name. The configuration should look like the screenshot below:

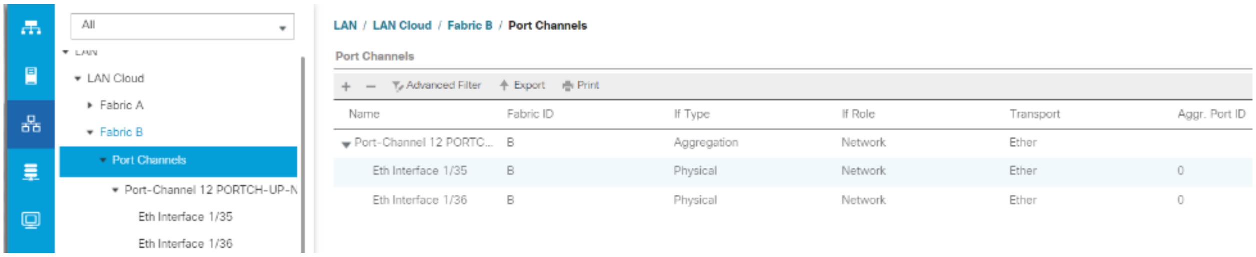

LAN / LAN Cloud / Fabric B / Port Channels

Port Channels

Name	Fabric ID	If Type	If Role	Transport	Aggr. Port ID
Port-Channel 12 PORTC...	B	Aggregation	Network	Ether	
Eth Interface 1/35	B	Physical	Network	Ether	0
Eth Interface 1/36	B	Physical	Network	Ether	0

Create Required Shared Resource Pools

This section details how to create the MAC address, iSCSI IQN, iSCSI IP, UUID suffix and server pools.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

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

2. Select Pools > root.

3. Right-click MAC Pools under the root organization.

4. Select Create MAC Pool to create the MAC address pool.

5. Enter MAC_Pool_A as the name for MAC pool.

6. Optional: Enter a description for the MAC pool.

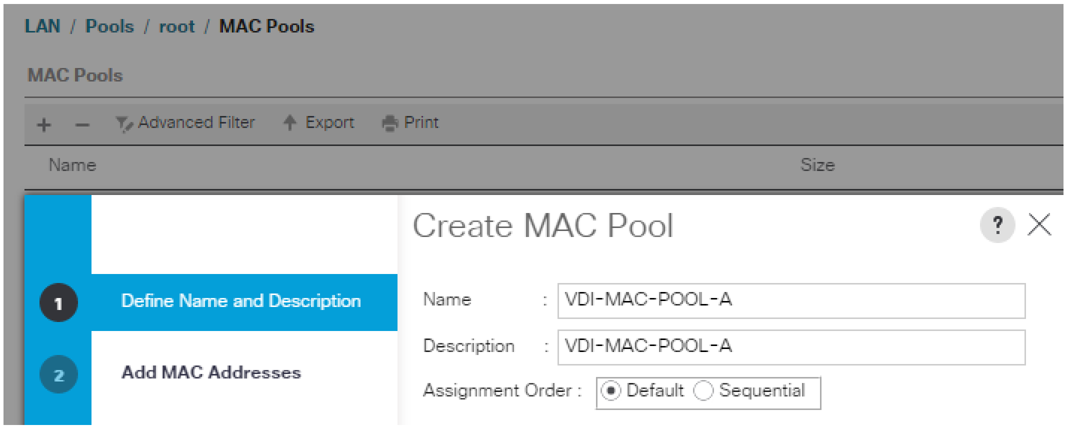

LAN / Pools / root / MAC Pools
MAC Pools
+ − Advanced Filter Export Print
Name
Size
Create MAC Pool
1 Define Name and Description
2 Add MAC Addresses

Name : VDI-MAC-POOL-A

Description : VDI-MAC-POOL-A

Assignment Order : ● Default ○ Sequential

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

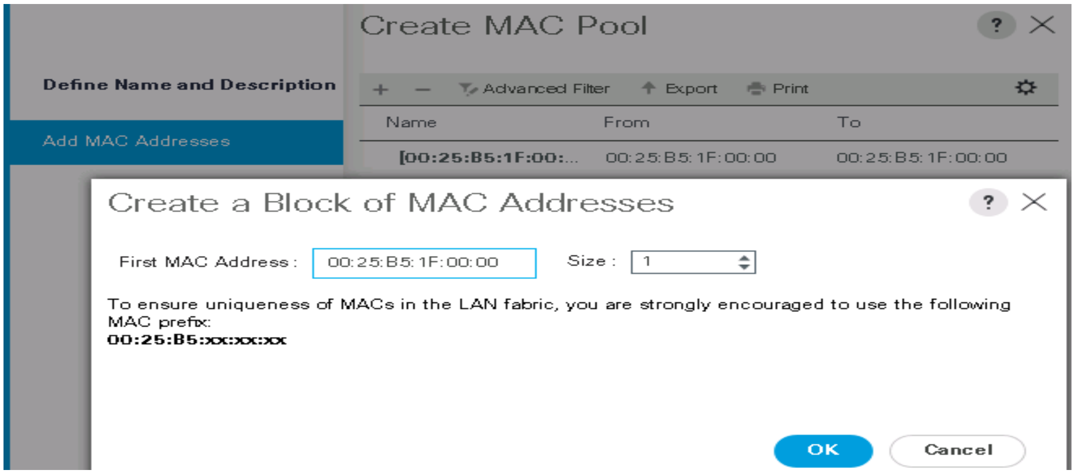

Create MAC Pool
Define Name and Description
Add MAC Addresses
+ − Advanced Filter Export Print
Name
From
To
Create a Block of MAC Addresses
OK
Cancel

Name	From	To
[00:25:B5:1F:00:...	00:25:B5:1F:00:00	00:25:B5:1F:00:00

First MAC Address : 00:25:B5:1F:00:00 Size : 1

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. Click OK, then click Finish.

9. In the confirmation message, click OK.

Create KVM IP Address Pool

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 the pool, complete the following steps:

1. Click the LAN tab in UCS Manager, expand the Pools node, expand the root node, right-click IP Pools, then click Create IP Pool.

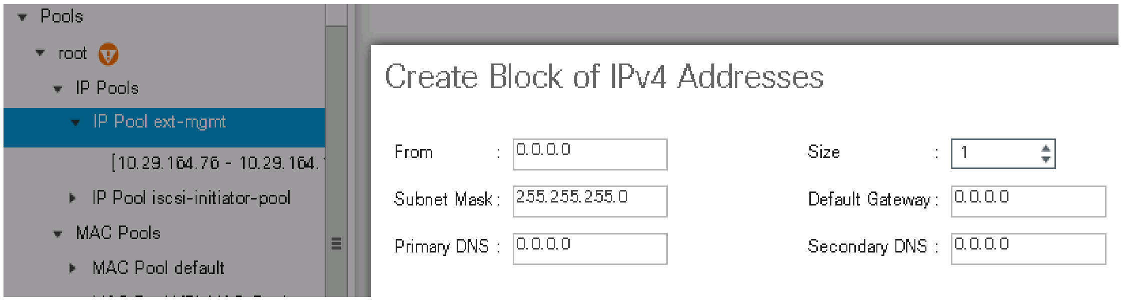

From : 0.0.0.0

Size : 1

Subnet Mask : 255.255.255.0

Default Gateway : 0.0.0.0

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

2. Provide a Name, choose Default or Sequential, and then click Next.

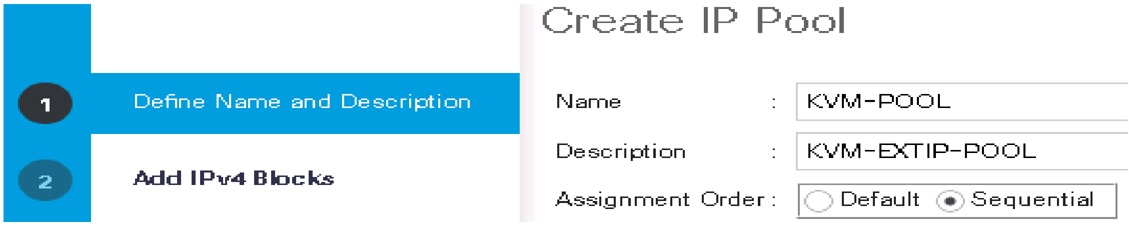

Create IP Pool

Name : KVM-POOL

Description : KVM-EXTIP-POOL

Assignment Order : ○ Default ● Sequential

3. Click the green + sign to add an IPv4 address block.

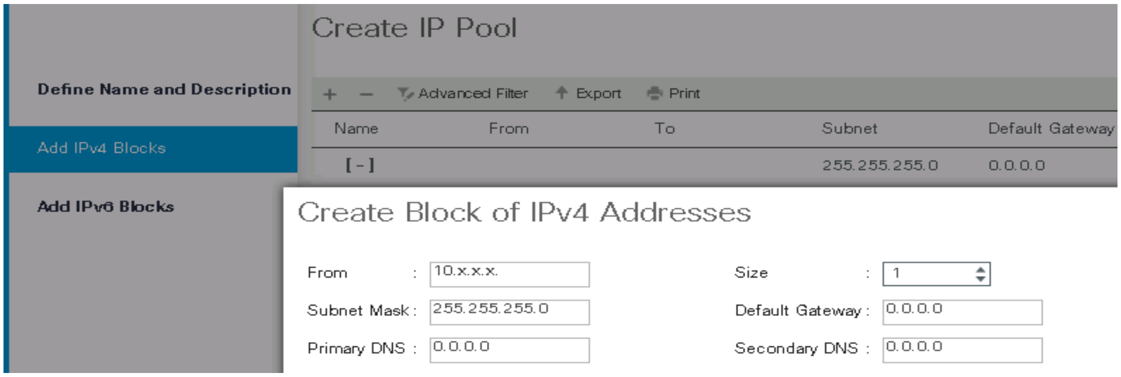

Create IP Pool

Name	From	To	Subnet	Default Gateway
[-]			255.255.255.0	0.0.0.0

Create Block of IPv4 Addresses

From : 10.x.x.x.

Size : 1

Subnet Mask : 255.255.255.0

Default Gateway : 0.0.0.0

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

4. Complete the starting IP address, size, subnet mask, default gateway, primary and secondary DNS values for your network, then click OK.

5. Click Finish.

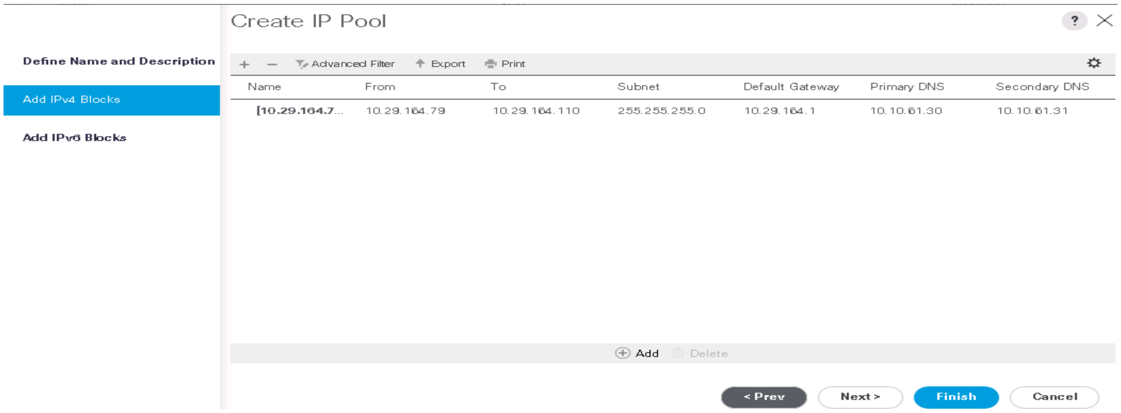

Create IP Pool

Name	From	To	Subnet	Default Gateway	Primary DNS	Secondary DNS
[10.29.164.7...	10.29.164.79	10.29.164.110	255.255.255.0	10.29.164.1	10.10.61.30	10.10.61.31

6. Click OK.

Create WWPN Pools

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

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

2. Select Pools > root.

3. Under WWPN Pools, right-click WWPN Pools and select Create WWPN Pool.

4. Assign a name and optional description.

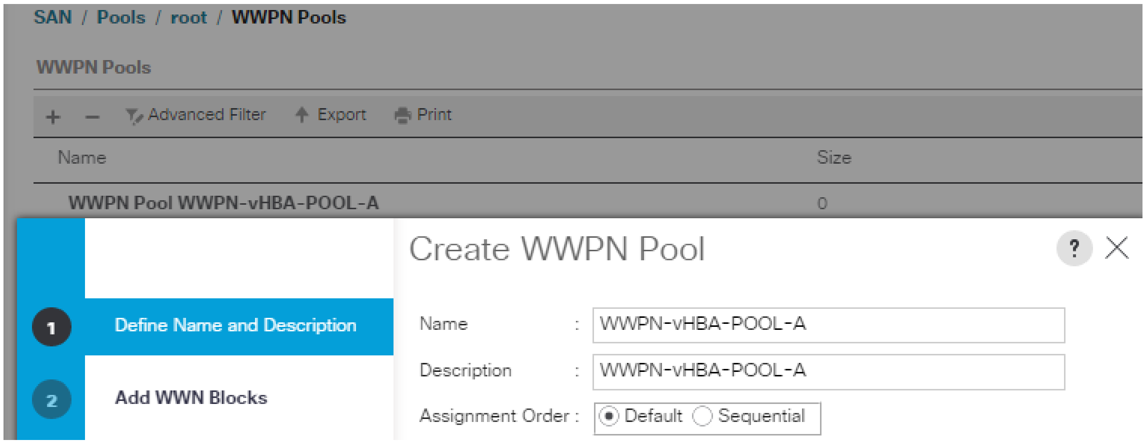
SAN / Pools / root / WWPN Pools
WWPN Pools
+ — Advanced Filter Export Print
Name
Size
WWPN Pool WWPN-vHBA-POOL-A
0
Create WWPN Pool
?
1 Define Name and Description
2 Add WWN Blocks

Name : WWPN-vHBA-POOL-A

Description : WWPN-vHBA-POOL-A

Assignment Order : ● Default ○ Sequential

5. Assignment order can remain Default.

6. Click Next.

7. Click Add to add a block of Ports.

8. Specify the size of a WWNN block sufficient enough to support 4 fully populated chassis.

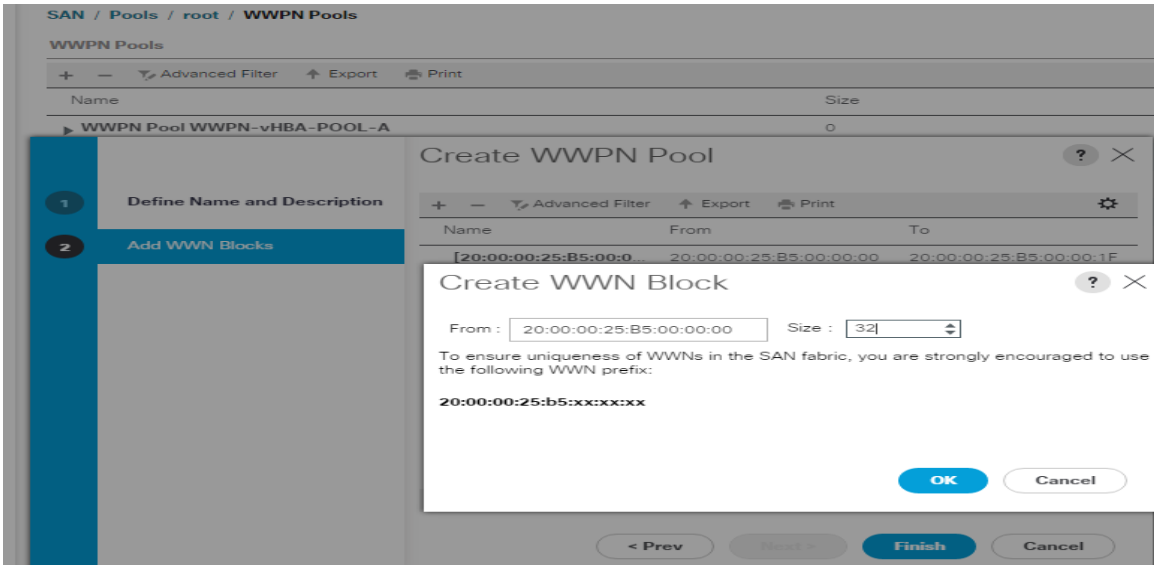
SAN / Pools / root / WWPN Pools
WWPN Pools
+ — Advanced Filter Export Print
Name
Size
WWPN Pool WWPN-vHBA-POOL-A
0
Create WWPN Pool
?
1 Define Name and Description
2 Add WWN Blocks
+ — Advanced Filter Export Print
Create WWN Block
?
OK
Cancel
< Prev
Next >
Finish
Cancel

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

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

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

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

9. Click Finish.

Create UUID Suffix Pool

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

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

2. Select Pools > root.

3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.

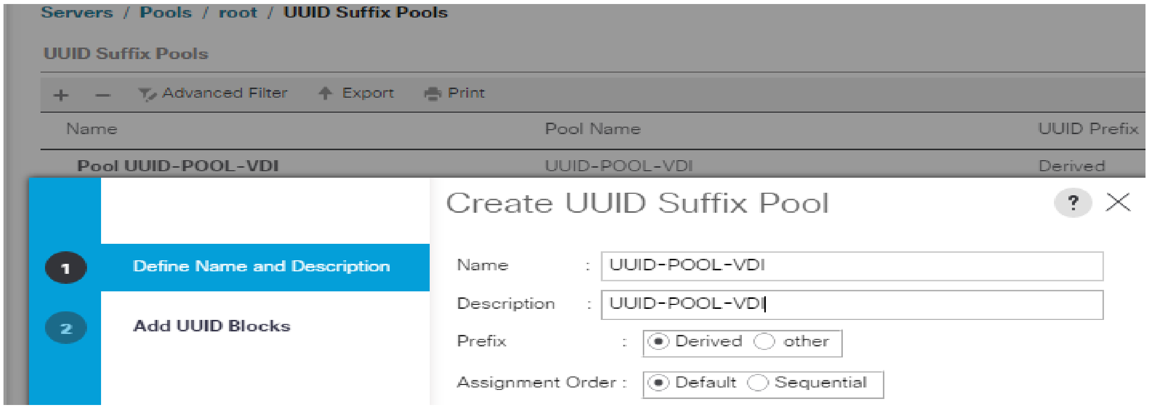
Servers / Pools / root / UUID Suffix Pools
UUID Suffix Pools
+ — Advanced Filter Export Print
Create UUID Suffix Pool
? ×
1 Define Name and Description
2 Add UUID Blocks

Name	Pool Name	UUID Prefix
Pool UUID-POOL-VDI	UUID-POOL-VDI	Derived

Name : UUID-POOL-VDI

Description : UUID-POOL-VDI

Prefix : ● Derived ○ other

Assignment Order : ● Default ○ Sequential

5. Enter UUID_Pool-VDI as the name of the UUID suffix pool.

6. Optional: Enter a description for the UUID suffix pool.

7. Keep the prefix at the derived option.

8. Click Next.

9. Click Add to add a block of UUIDs.

10. Create a starting point UUID seed for your environment.

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

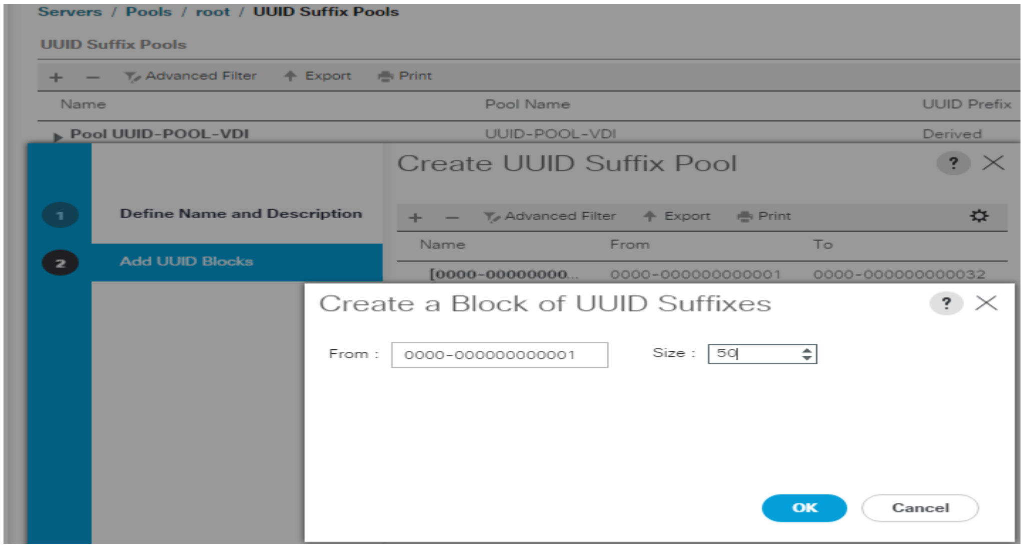
Servers / Pools / root / UUID Suffix Pools
UUID Suffix Pools
+ — Advanced Filter Export Print
Create UUID Suffix Pool
? ×
1 Define Name and Description
2 Add UUID Blocks
+ — Advanced Filter Export Print
Create a Block of UUID Suffixes
? ×
OK
Cancel

Name	Pool Name	UUID Prefix
Pool UUID-POOL-VDI	UUID-POOL-VDI	Derived

Name	From	To
[0000-00000000...	0000-000000000001	0000-000000000032

From : 0000-000000000001

Size : 50

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following 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. Select Pools > root.

3. Right-click Server Pools.

4. Select Create Server Pool.

5. Enter Infra_Pool as the name of the server pool.

6. Optional: Enter a description for the server pool.

7. Click Next.

8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra_Pool server pool.

9. Click Finish.

10. Click OK.

11. Create additional Server Pools for persistent, persistent, and RDS hosts

Create VLANs

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

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

In this procedure, eight unique VLANs are created. Refer to Table 14 .

Table 14 Created VLANs

VLAN Name	VLAN ID	VLAN Purpose	vNIC Assignment
Default	1	Native VLAN	vNIC-Template-A vNIC-Template-B
In-Band-Mgmt	60	VLAN for in-band management interfaces	vNIC-Template-A vNIC-Template-B
Infra-Mgmt	61	VLAN for Virtual Infrastructure	vNIC-Template-A vNIC-Template-B
NFS-Vlan	62	VLAN for NFS Share	vNIC-Template-A vNIC-Template-B
CIFS-Vlan	63	VLAN-CIFS Share User Profiles	vNIC-Template-A vNIC-Template-B
vMotion	66	VLAN for VMware vMotion	vNIC-Template-A vNIC-Template-B
VDI	102	Virtual Desktop traffic	vNIC-Template-A vNIC-Template-B

VLAN Name	VLAN ID	VLAN Purpose	vNIC Assignment
OB-Mgmt	164	VLAN for out-of-band management interfaces	vNIC-Template-A vNIC-Template-B

2. Select LAN > LAN Cloud.

3. Right-click VLANs.

4. Select Create VLANs

5. Enter MGMT as the name of the VLAN to be used for in-band management traffic.

6. Keep the Common/Global option selected for the scope of the VLAN.

7. Enter 60 as the ID of the management VLAN.

8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

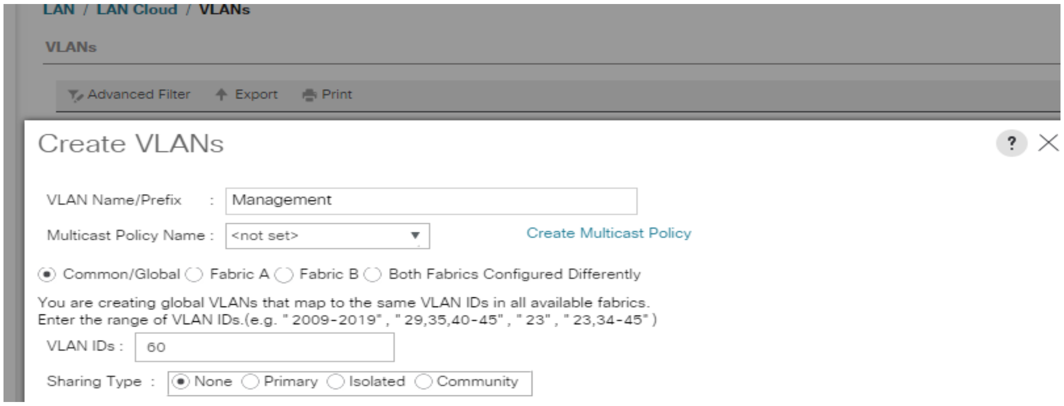

Create VLANs

VLAN Name/Prefix : Management

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 : 60

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

10. Repeat the above steps to create all VLANs and configure the Default VLAN as native.

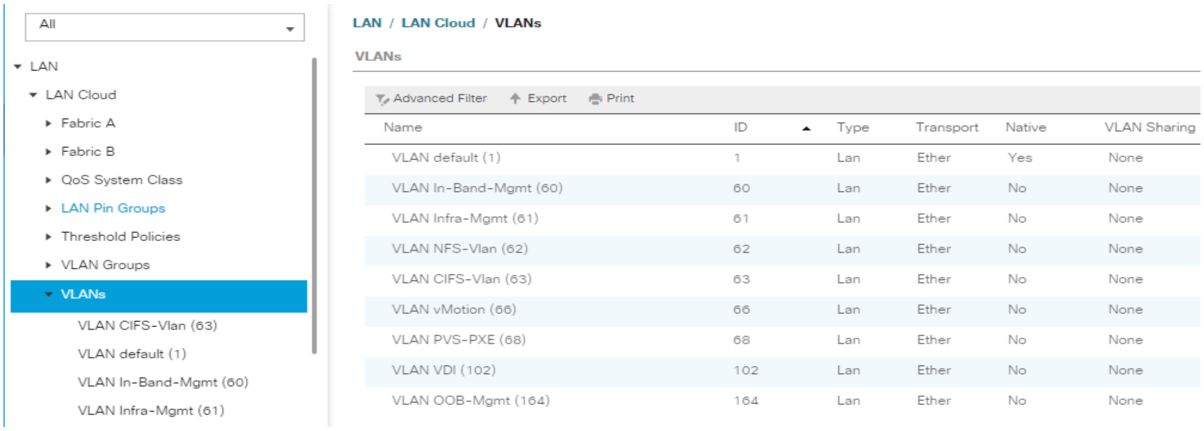

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN In-Band-Mgmt (60)	60	Lan	Ether	No	None
VLAN Infra-Mgmt (61)	61	Lan	Ether	No	None
VLAN NFS-Vlan (62)	62	Lan	Ether	No	None
VLAN CIFS-Vlan (63)	63	Lan	Ether	No	None
VLAN vMotion (66)	66	Lan	Ether	No	None
VLAN PVS-PXE (68)	68	Lan	Ether	No	None
VLAN VDI (102)	102	Lan	Ether	No	None
VLAN OOB-Mgmt (164)	164	Lan	Ether	No	None

Create VSANs

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

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

In this procedure, two VSANs are created. When these VSANs are created, be sure to add them to the uplink FC Interfaces created earlier.

2. Select SAN > SAN Cloud.
3. Under Fabric A, right-click VSANs.
4. Select Create VSANs.
5. Enter VSAN-400-A as the name of the VSAN to be used for in-band management traffic.
6. Select Fabric A for the scope of the VSAN.
7. Enter 400 as the ID of the VSAN.
8. Click OK and then click OK again.

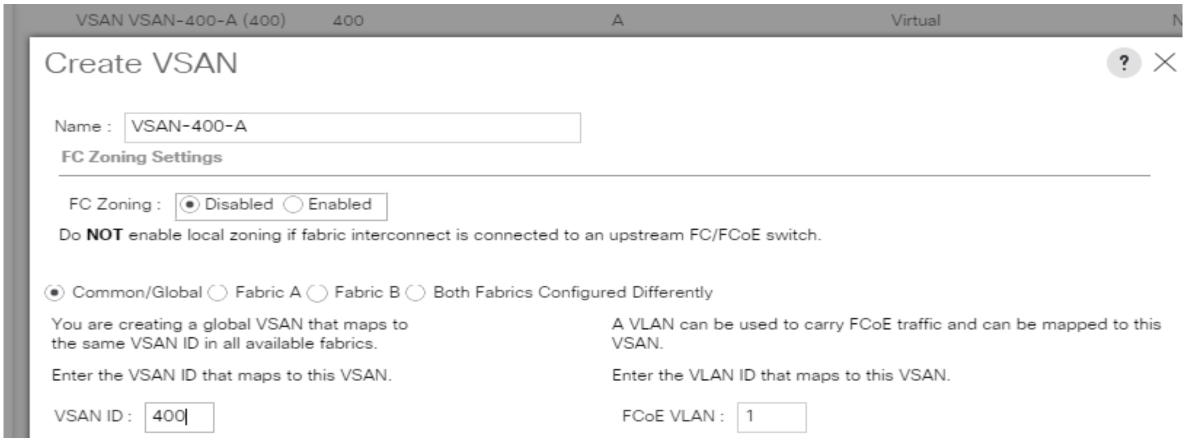

VSAN VSAN-400-A (400) 400 A Virtual

Create VSAN

Name : VSAN-400-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 global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 400

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 1

9. Repeat the above steps on Fabric B with VSAN-401-B to create the VSANs necessary for this solution.

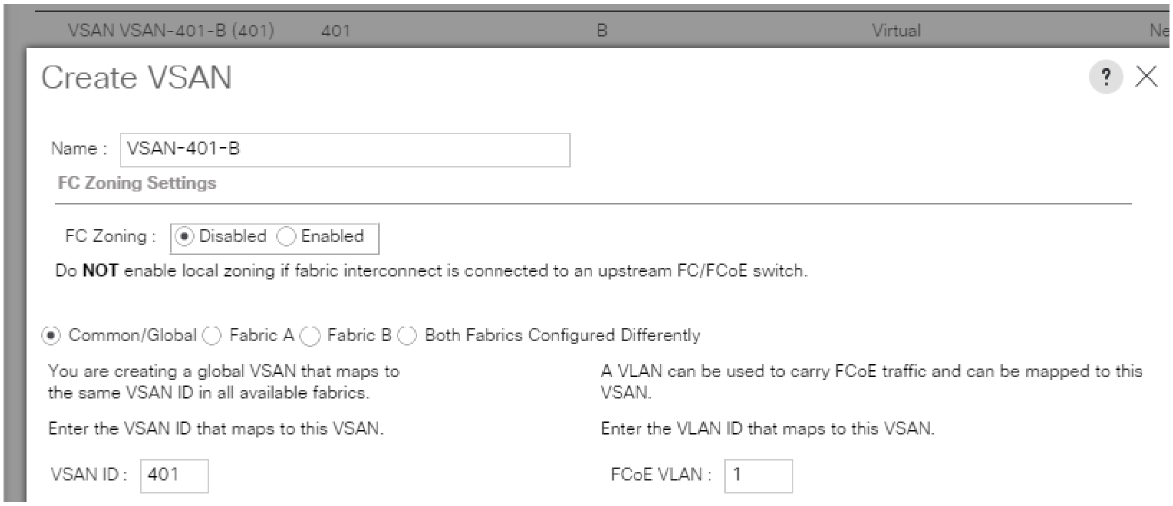

VSAN VSAN-401-B (401) 401 B Virtual

Create VSAN

Name : VSAN-401-B

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 global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 401

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 : 1

VSAN 400 and 401 are configured as shown below:

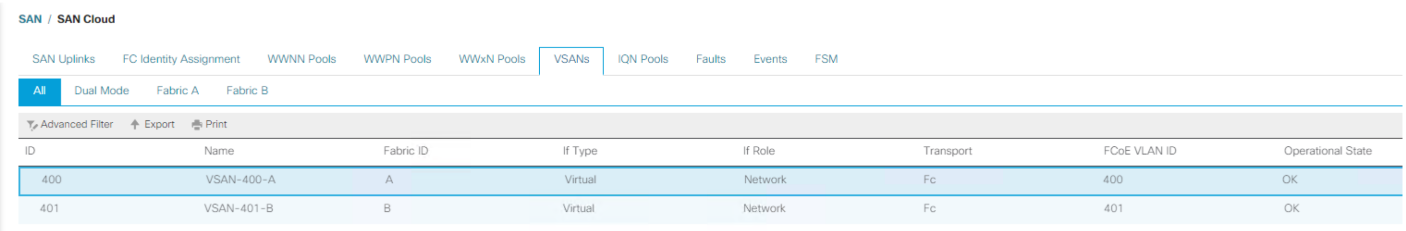

SAN / SAN Cloud

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

10. After configuring VSANs both sides, go into the port-channel created earlier in the section ‘Create uplinks for MDS 9148S and add the respective VSANs to their port channels. VSAN400 in this study is assigned to Fabric A and VSAN401 is assigned to Fabric B.

VSAN400 should only be on Fabric A and VSAN401 on Fabric B.

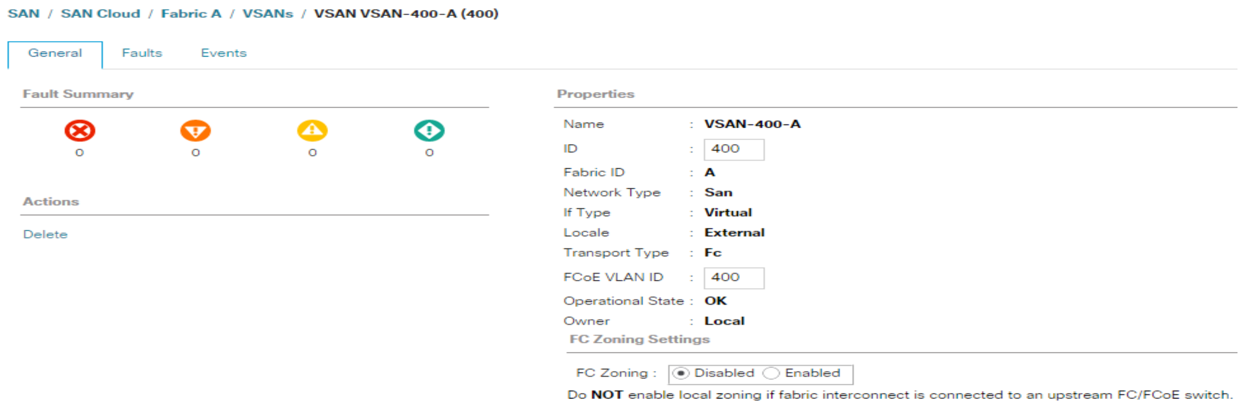

SAN / SAN Cloud / Fabric A / VSANs / VSAN VSAN-400-A (400)

Properties

Name : VSAN-400-A
ID : 400
Fabric ID : A
Network Type : San
If Type : Virtual
Locale : External
Transport Type : Fc
FCoE VLAN ID : 400
Operational State : OK
Owner : Local

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

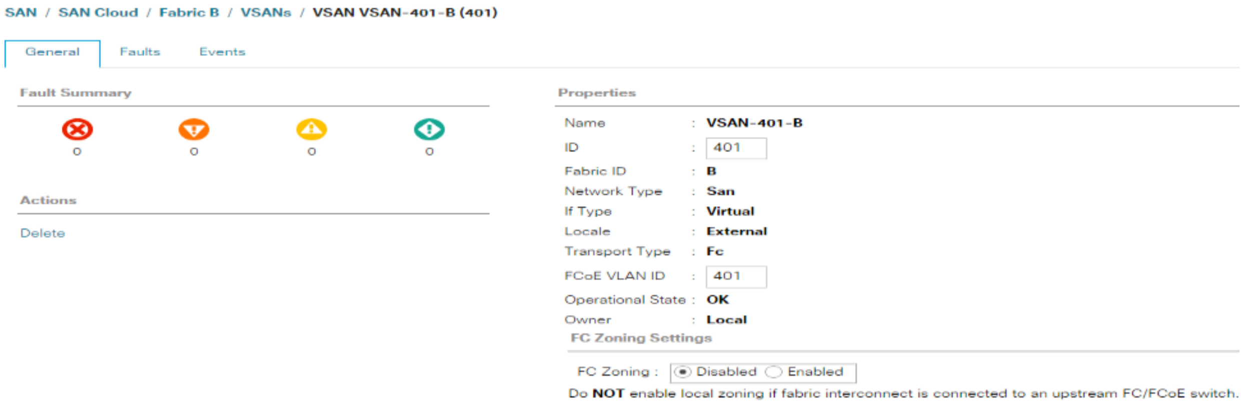

SAN / SAN Cloud / Fabric B / VSANs / VSAN VSAN-401-B (401)

Properties

Name : VSAN-401-B
ID : 401
Fabric ID : B
Network Type : San
If Type : Virtual
Locale : External
Transport Type : Fc
FCoE VLAN ID : 401
Operational State : OK
Owner : Local

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

11. Go to the Uplink FC interfaces for each Fabric and assign the VSAN appropriately to each FC Interface.

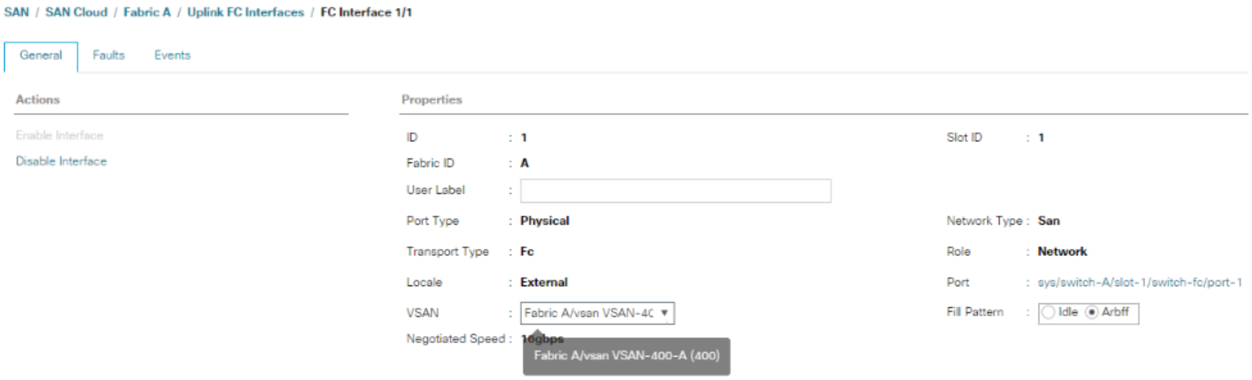

SAN / SAN Cloud / Fabric A / Uplink FC Interfaces / FC Interface 1/1

Properties

ID : 1
Fabric ID : A
User Label :
Port Type : Physical
Transport Type : Fc
Locale : External
VSAN : Fabric A/vsan VSAN-4C
Negotiated Speed : 16gbps

Fabric A/vsan VSAN-400-A (400)

Slot ID : 1
Network Type : San
Role : Network
Port : sys/switch-A/slot-1/switch-fc/port-1
Fill Pattern : ○ Idle ● Arbff

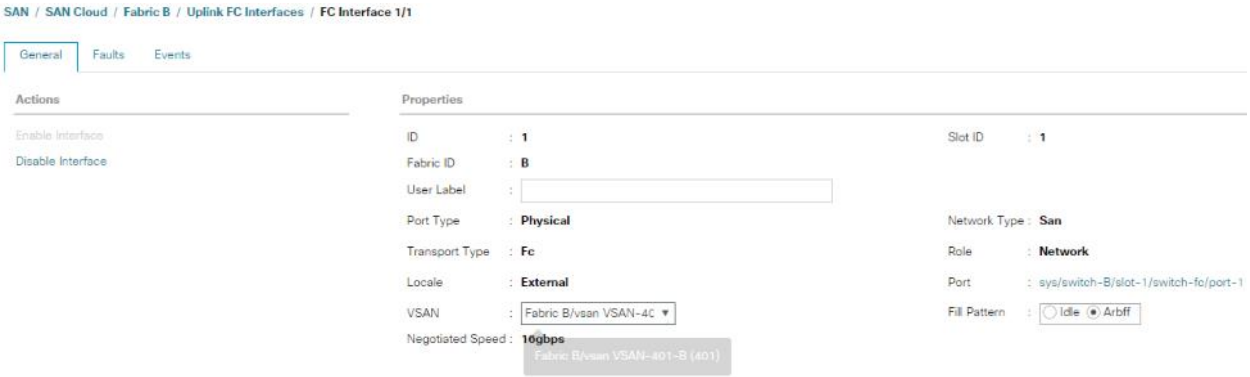

SAN / SAN Cloud / Fabric B / Uplink FC Interfaces / FC Interface 1/1

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID	: 1	Slot ID	: 1
Fabric ID	: B		
User Label	:		
Port Type	: Physical	Network Type	: San
Transport Type	: Fc	Role	: Network
Locale	: External	Port	: sys/switch-B/slot-1/switch-fc/port-1
VSAN	: Fabric B/vsan VSAN-40 ▾	Fill Pattern	: ○ Idle ● Arbff
Negotiated Speed	: 16gbps		

Fabric B/vsan VSAN-401-B (401)

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, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter VM-Host as the name of the host firmware package.
6. Leave Simple selected.
7. Select the version 3.2(3d) for the Blade Package
8. Click OK to create the host firmware package.

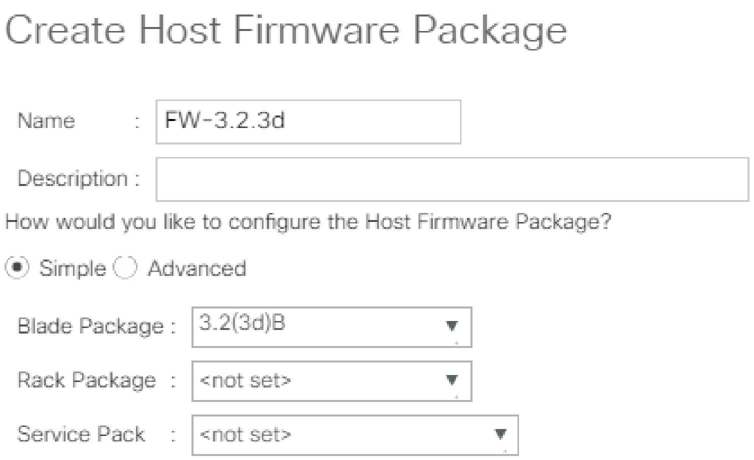

Create Host Firmware Package

Name : FW-3.2.3d

Description :

How would you like to configure the Host Firmware Package?

● Simple ○ Advanced

Blade Package : 3.2(3d)B ▾

Rack Package : <not set> ▾

Service Pack : <not set> ▾

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following 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 in the bottom of the window.
6. Click OK.

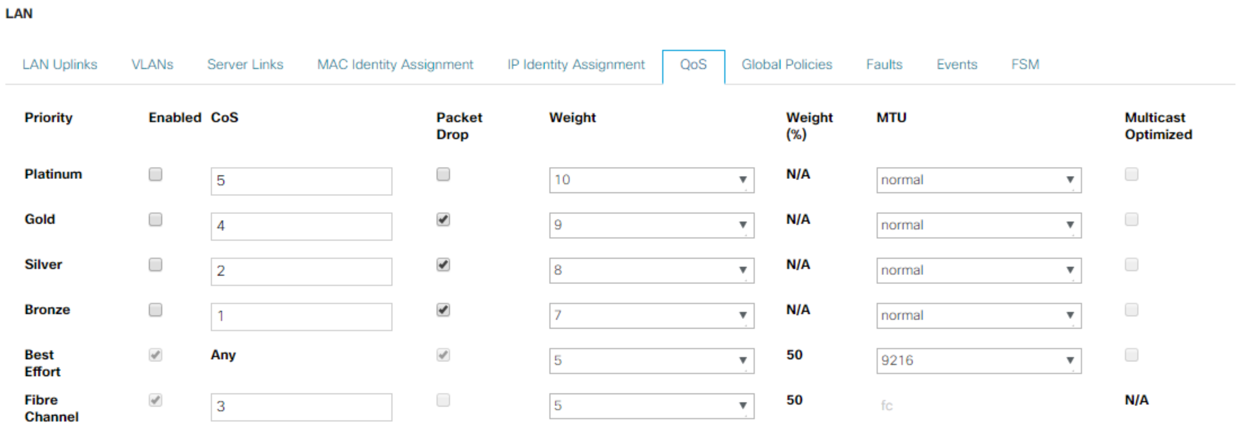

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	□	5	□	10	N/A	normal	□
Gold	□	4	☑	9	N/A	normal	□
Silver	□	2	☑	8	N/A	normal	□
Bronze	□	1	☑	7	N/A	normal	□
Best Effort	☑	Any	☑	5	50	9216	□
Fibre Channel	☑	3	□	5	50	fc	N/A

Create Network Control Policy for Cisco Discovery Protocol

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable_CDP as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

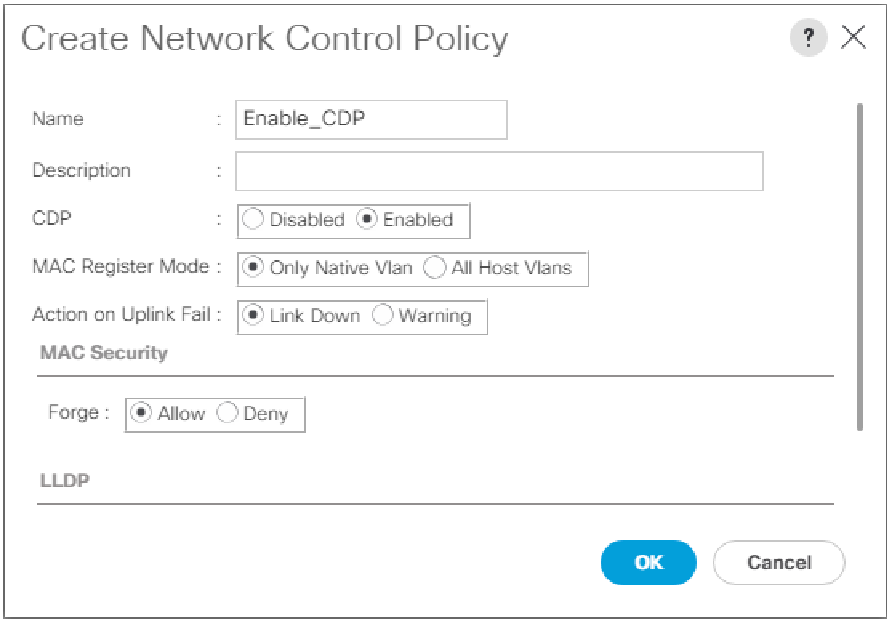

Create Network Control Policy

Name : Enable_CDP

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, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.

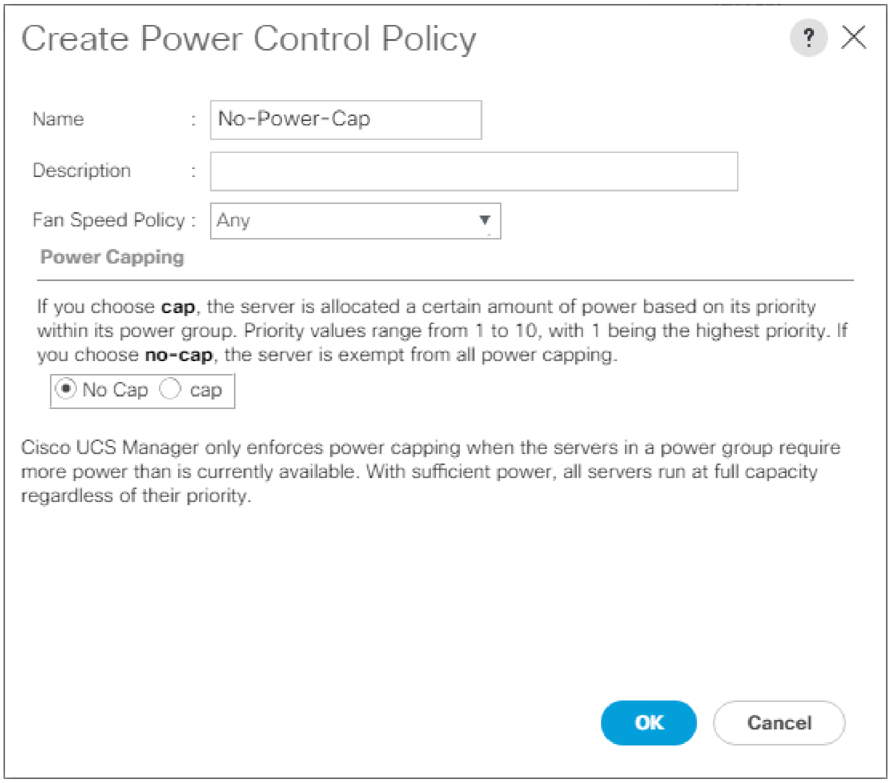

Create Power Control Policy

Name : No-Power-Cap

Description :

Fan Speed Policy : Any

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.

Cisco UCS System Configuration for Cisco UCS B-Series

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

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

2. Select Policies > root.

3. Right-click BIOS Policies.

4. Select Create BIOS Policy.

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

6. Configure the remaining BIOS policies as follows and click Finish.

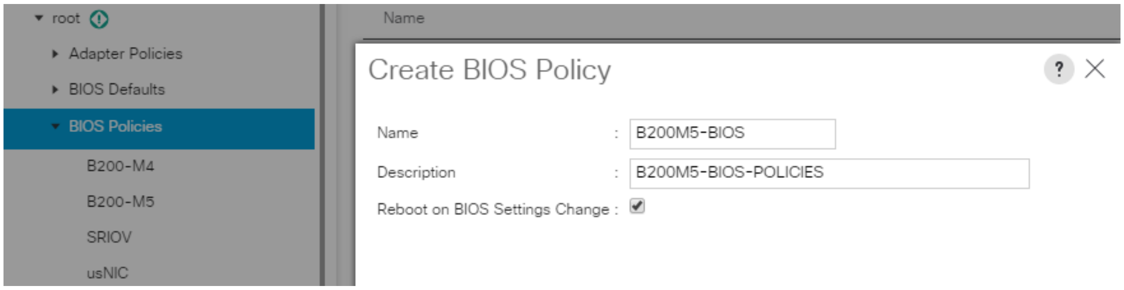

Create BIOS Policy

Name : B200M5-BIOS

Description : B200M5-BIOS-POLICIES

Reboot on BIOS Settings Change : ☑

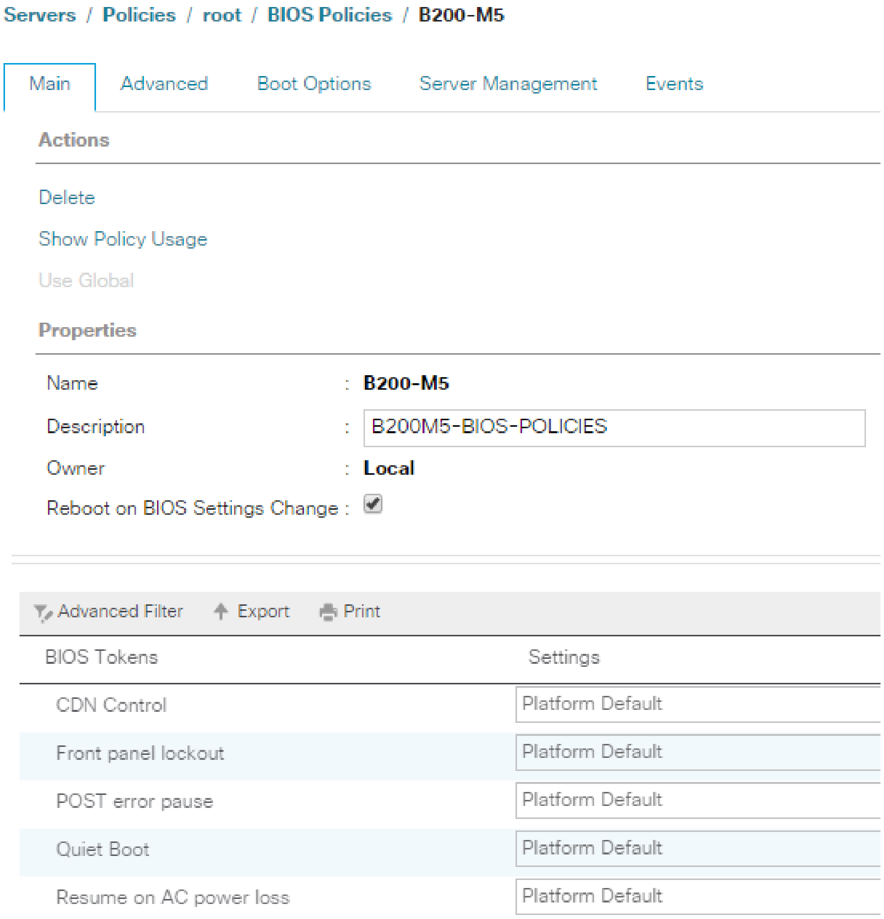

Servers / Policies / root / BIOS Policies / B200-M5

Properties

Name : **B200-M5**

Description : B200M5-BIOS-POLICIES

Owner : **Local**

Reboot on BIOS Settings Change : ☑

BIOS Tokens	Settings
CDN Control	Platform Default
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Platform Default
Resume on AC power loss	Platform Default

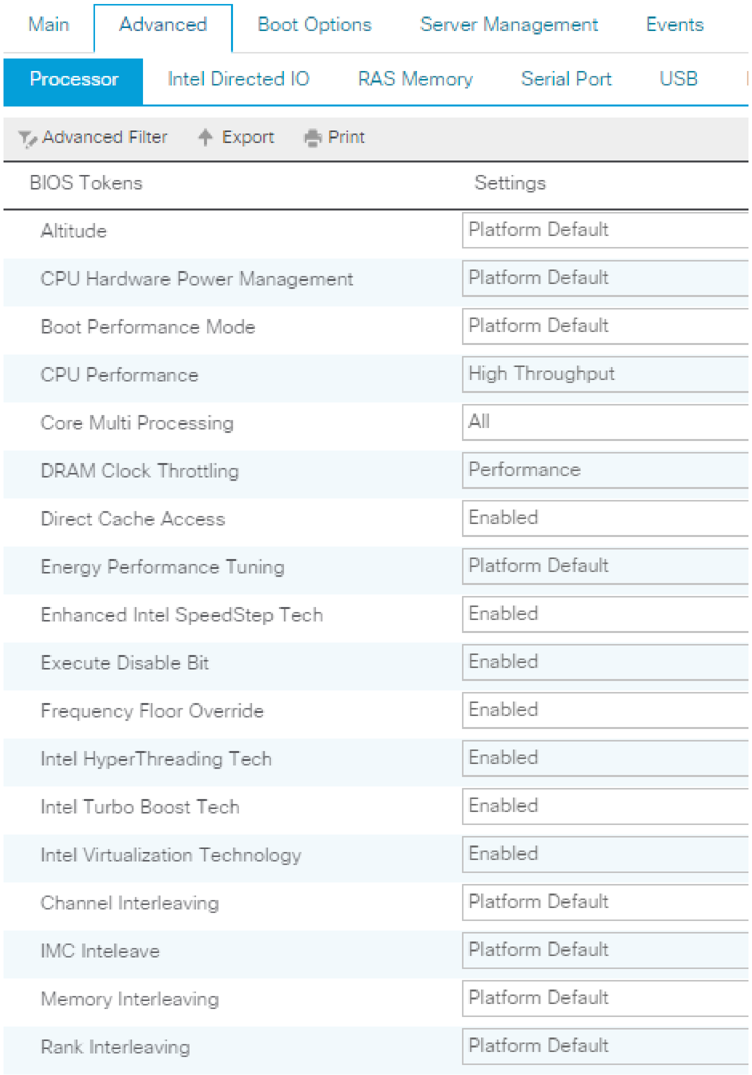

BIOS Tokens	Settings
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Platform Default
CPU Performance	High Throughput
Core Multi Processing	All
DRAM Clock Throttling	Performance
Direct Cache Access	Enabled
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Enabled
Execute Disable Bit	Enabled
Frequency Floor Override	Enabled
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	Enabled
Intel Virtualization Technology	Enabled
Channel Interleaving	Platform Default
IMC Inteleave	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default

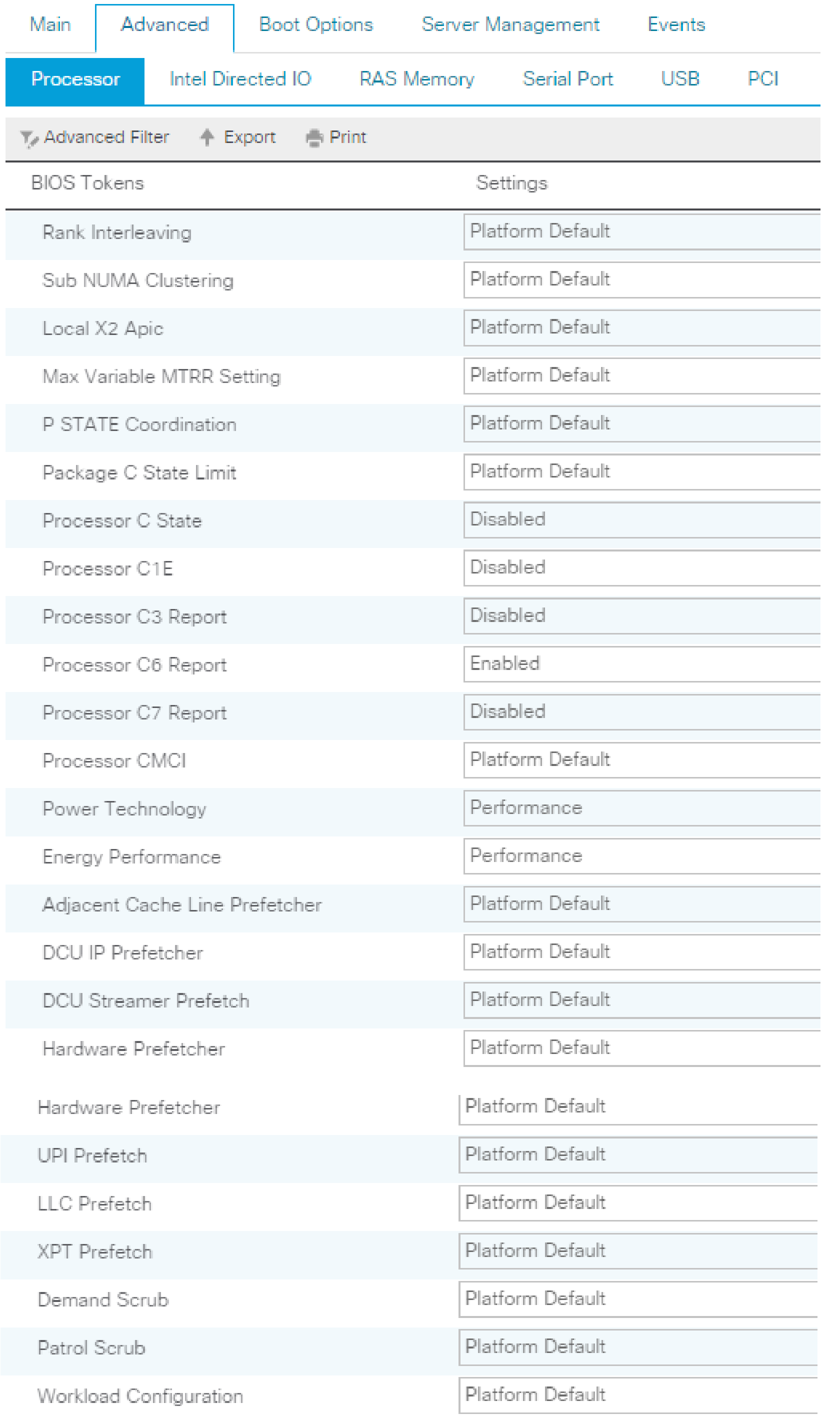

BIOS Tokens	Settings
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Enabled
Processor C7 Report	Disabled
Processor CMCI	Platform Default
Power Technology	Performance
Energy Performance	Performance
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	Platform Default

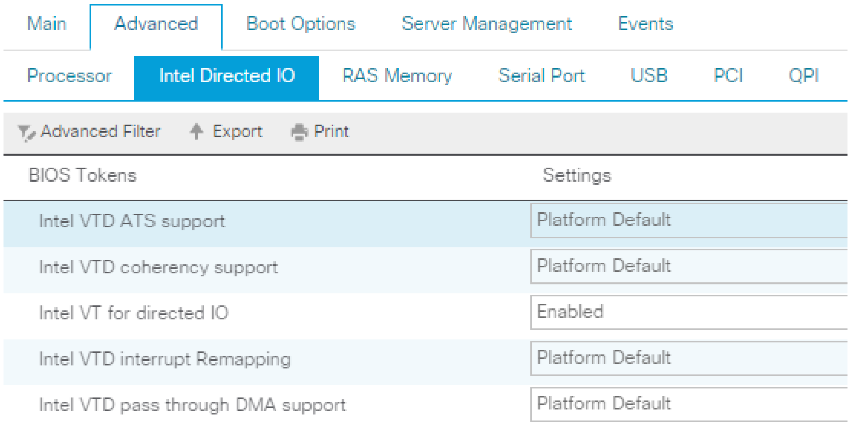

BIOS Tokens	Settings
Intel VTD ATS support	Platform Default
Intel VTD coherency support	Platform Default
Intel VT for directed IO	Enabled
Intel VTD interrupt Remapping	Platform Default
Intel VTD pass through DMA support	Platform Default

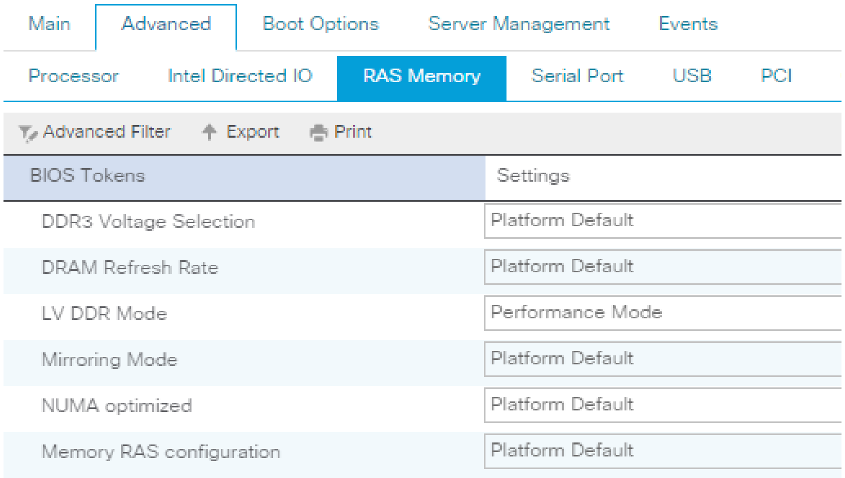

BIOS Tokens	Settings
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

7. Click Finish.

Configure Update Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Check **On Next Boot** check box
6. Click Save Changes.
7. Click OK to accept the change.

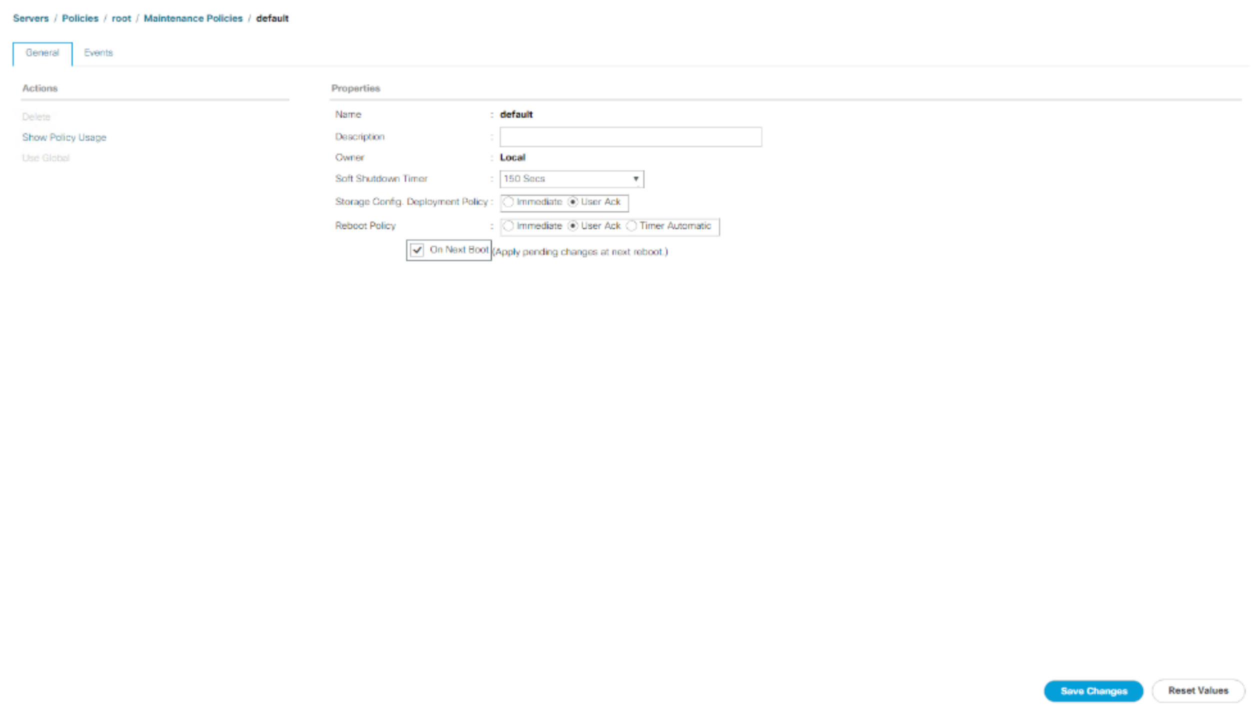

Name : default

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 for Cisco UCS B-Series

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for MGMT, Default, Infra, VDI, and vMotion.
11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select CDP_Enabled.

15. Click OK to create the vNIC template.

16. Click OK.

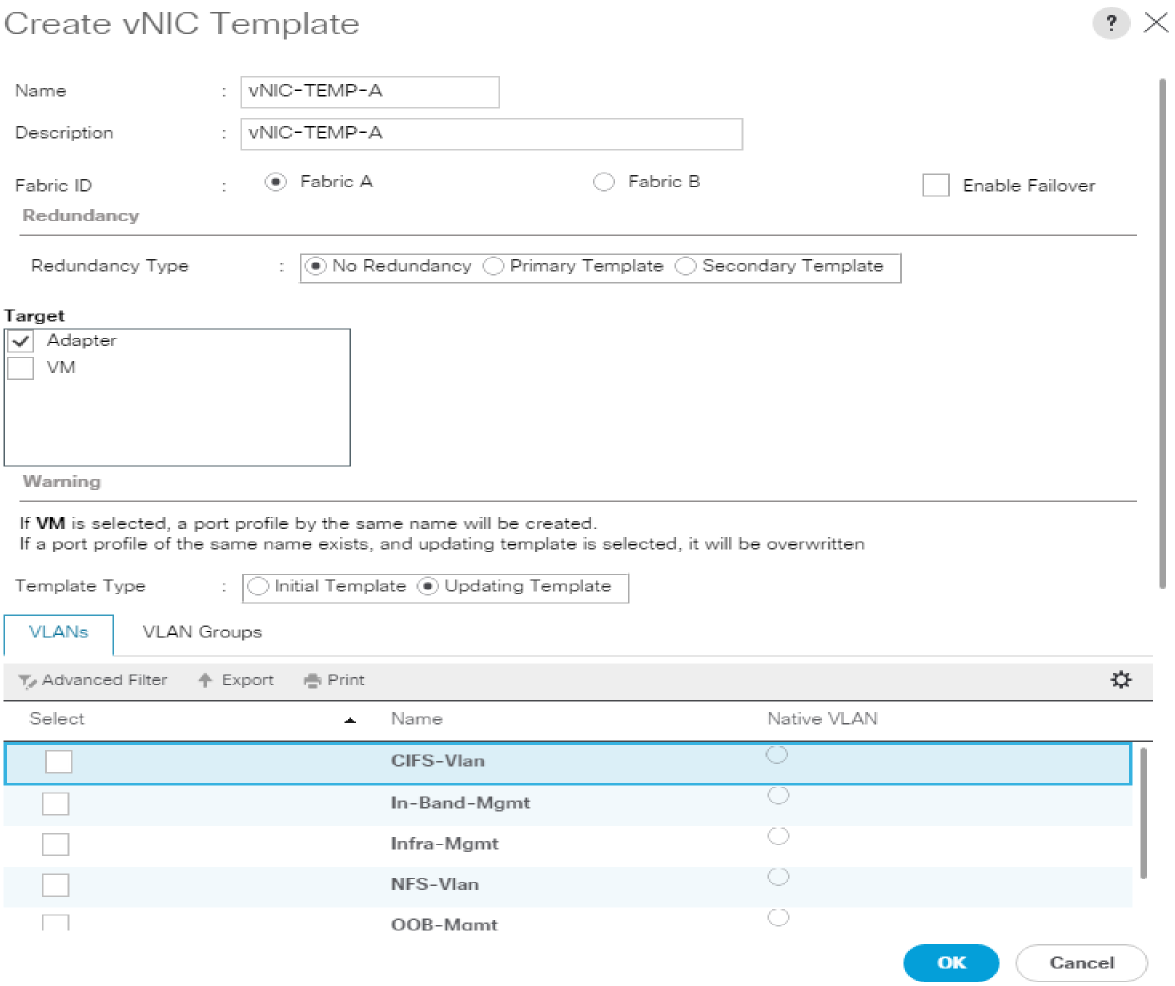

Create vNIC Template

Name : vNIC-TEMP-A

Description : vNIC-TEMP-A

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

Redundancy

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

Target

☑ Adapter

□ VM

Warning

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

Template Type : ○ Initial Template ● Updating Template

Select	Name	Native VLAN
□	CIFS-Vlan	○
□	In-Band-Mgmt	○
□	Infra-Mgmt	○
□	NFS-Vlan	○
□	OOB-Mgmt	○

17. In the navigation pane, select the LAN tab.

18. Select Policies > root.

19. Right-click vNIC Templates.

20. Select Create vNIC Template.

21. Enter vNIC_Template_B as the vNIC template name.

22. Select Fabric B.

23. Do not select the Enable Failover checkbox.

24. Under Target, make sure the VM checkbox is not selected.

25. Select Updating Template as the template type.

26. Under VLANs, select the checkboxes for MGMT, Default, VDI, Infra, and vMotion.

27. Set Native-VLAN as the native VLAN.
28. For MTU, enter 9000.
29. In the MAC Pool list, select MAC_Pool_B.
30. In the Network Control Policy list, select CDP_Enabled.
31. Click OK to create the vNIC template.
32. Click OK.

Create vHBA Templates for Cisco UCS B-Series

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter vHBA-FAB-A as the vHBA template name.
6. Keep Fabric A selected.
7. Select VSAN-400-A for Fabric A from the drop down.
8. Change to Updating Template.
9. For Max Data Field keep 2048.
10. Select VDI-WWPN (created earlier) for our WWPN Pool.
11. Leave the remaining as is.
12. Click OK.

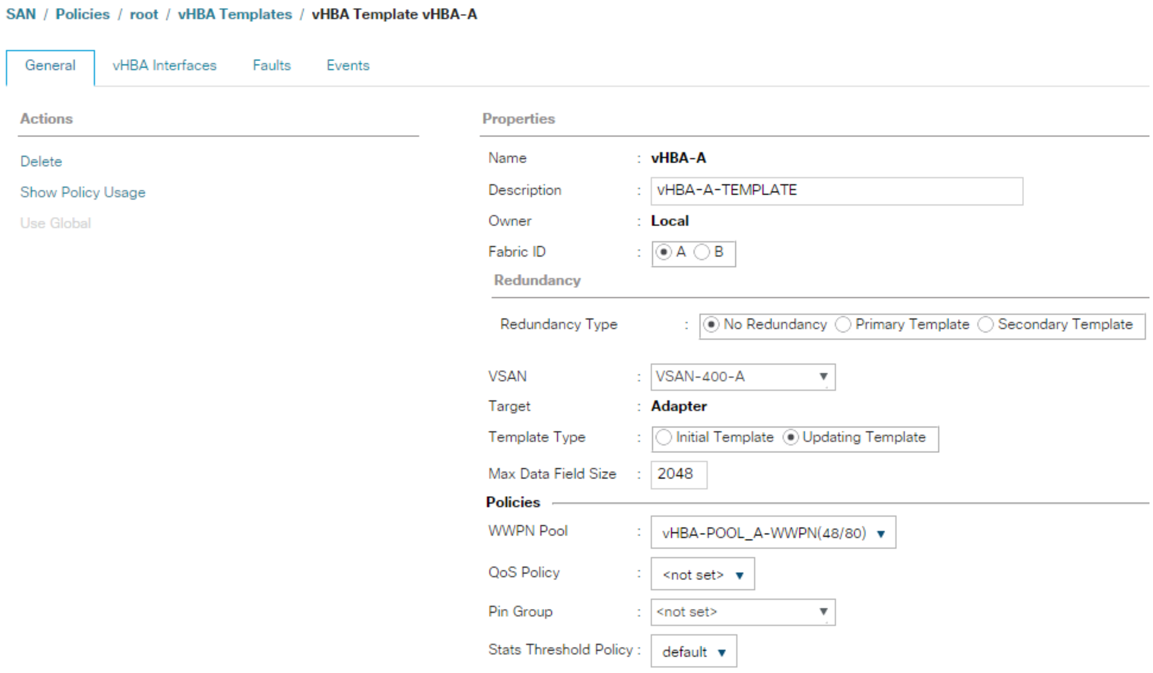

Properties

Name : vHBA-A

Description : vHBA-A-TEMPLATE

Owner : Local

Fabric ID : ● A ○ B

Redundancy

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

VSAN : VSAN-400-A ▾

Target : Adapter

Template Type : ○ Initial Template ● Updating Template

Max Data Field Size : 2048

Policies

WWPN Pool : vHBA-POOL_A-WWPN(48/80) ▾

QoS Policy : <not set> ▾

Pin Group : <not set> ▾

Stats Threshold Policy : default ▾

13. In the navigation pane, select the LAN tab.

14. Select Policies > root.

15. Right-click vHBA Templates.

16. Select Create vHBA Template.

17. Enter vHBA-FAB-B as the vHBA template name.

18. Select Fabric B.

19. Select VSAN-401-B for Fabric B from the drop down.

20. Change to Updating Template.

21. For Max Data Field keep 2048.

22. Select VDI-Pool-WWPN (created earlier) for our WWPN Pool.

23. Leave the remaining as is.

24. Click OK.

Configure Boot from SAN

All ESXi hosts were set to boot from SAN for the 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, name just a few.

To create a boot from SAN policy, complete the following steps:

1. Go to UCS Manager, right-click the ‘Boot Policies’ option shown below and select ‘Create Boot Policy.'

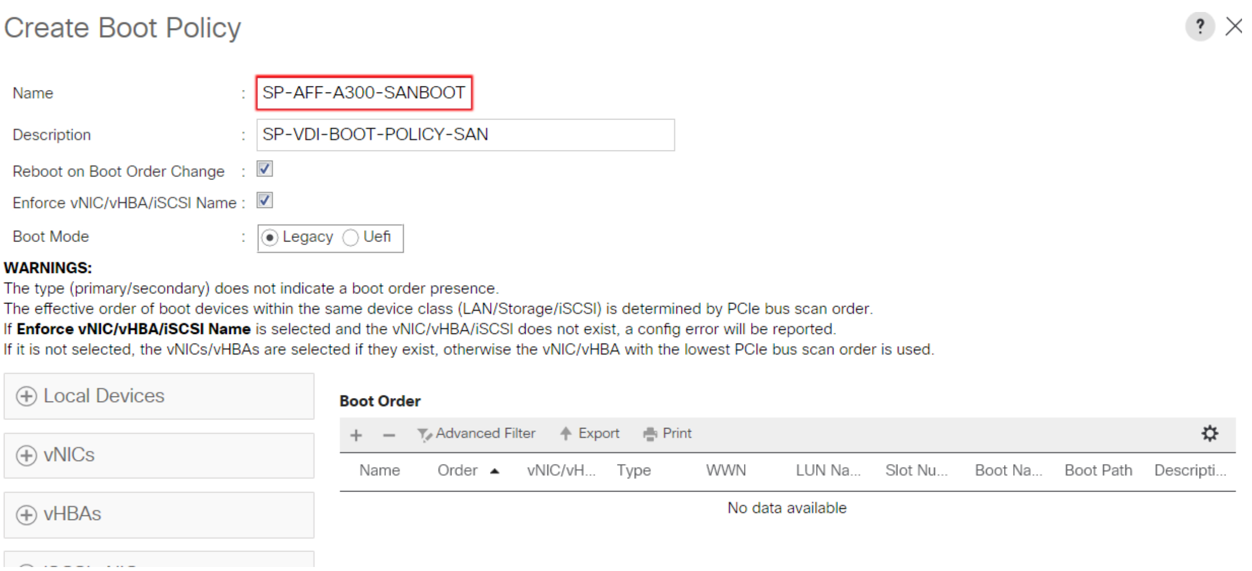

Create Boot Policy

Name : SP-AFF-A300-SANBOOT

Description : SP-VDI-BOOT-POLICY-SAN

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	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descripti...
				No data available					

2. Name the boot policy and expand the ‘vHBAs’ menu as shown below:

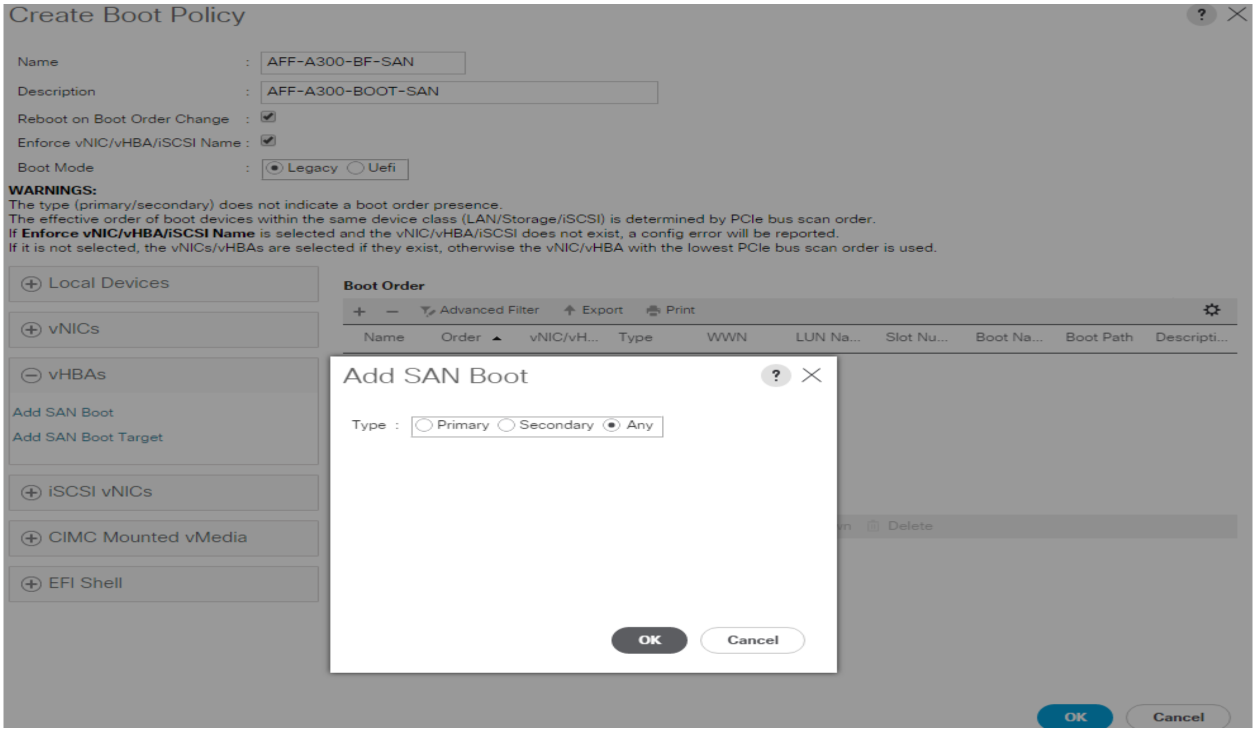

Create Boot Policy

Name : AFF-A300-BF-SAN

Description : AFF-A300-BOOT-SAN

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.

Add SAN Boot

Type : ○ Primary ○ Secondary ● Any

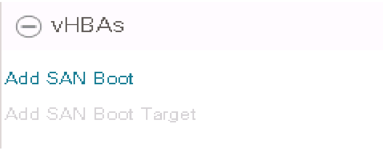

3. After selecting the ‘Add SAN Boot’ option, add the primary vHBA as shown below. Note that the vHBA name needs to match exactly. We will use the vHBA templates created in the previous step.

4. Repeat the steps to add a secondary SAN Boot option.

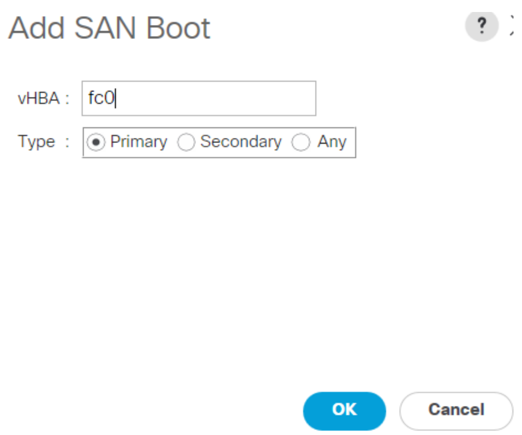

Add SAN Boot

vHBA : fc0

Type : ● Primary ○ Secondary ○ Any

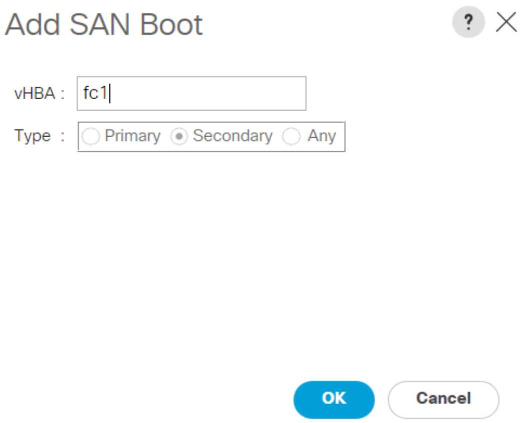

Add SAN Boot

vHBA : fc1

Type : ○ Primary ● Secondary ○ Any

5. Add the SAN Boot Targets to the primary and secondary. The SAN boot targets will also include primary and secondary options in order to maximize resiliency and number of paths.

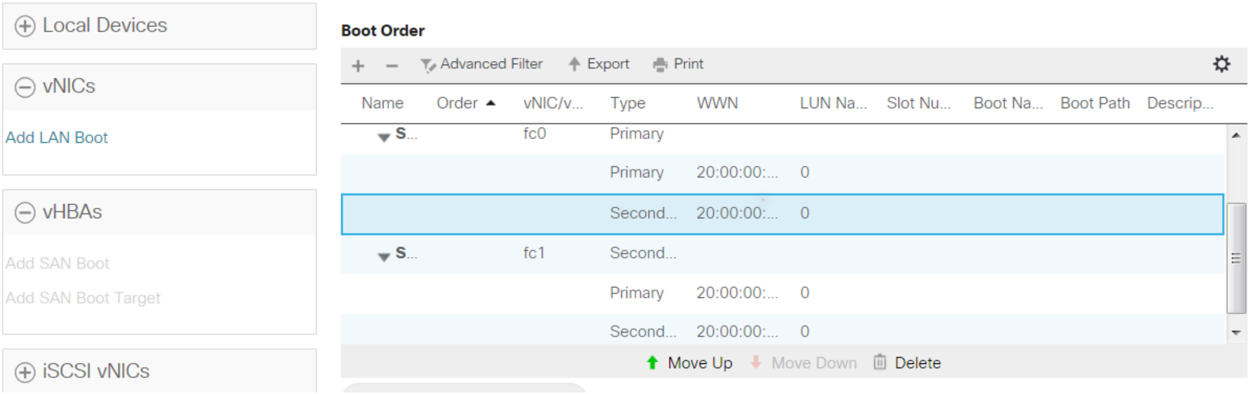

Boot Order

Name	Order	vNIC/v...	Type	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descrip...
S...		fc0	Primary						
			Primary	20:00:00:...	0				
			Second...	20:00:00:...	0				
S...		fc1	Second...						
			Primary	20:00:00:...	0				
			Second...	20:00:00:...	0				

6. Using the following command, find and record the WWPN for each FC LIF:

```
Network interface show -vserver <vserver> -data-protocol fcp
```

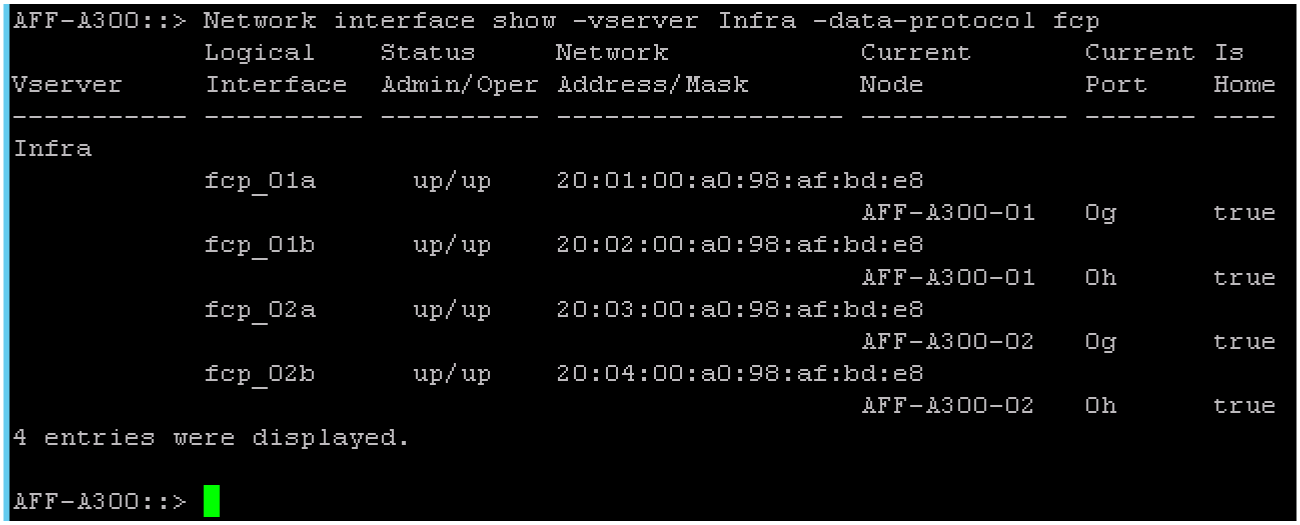

```
AFF-A300::> Network interface show -vserver Infra -data-protocol fcp
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
Infra
            fcp_01a      up/up    20:01:00:a0:98:af:bd:e8
                                                     AFF-A300-01   0g      true
            fcp_01b      up/up    20:02:00:a0:98:af:bd:e8
                                                     AFF-A300-01   0h      true
            fcp_02a      up/up    20:03:00:a0:98:af:bd:e8
                                                     AFF-A300-02   0g      true
            fcp_02b      up/up    20:04:00:a0:98:af:bd:e8
                                                     AFF-A300-02   0h      true
4 entries were displayed.

AFF-A300::> 
```

7. When the AFF A300 WWNs have been recorded, use fcp_01a for the first Boot Target WWPN:

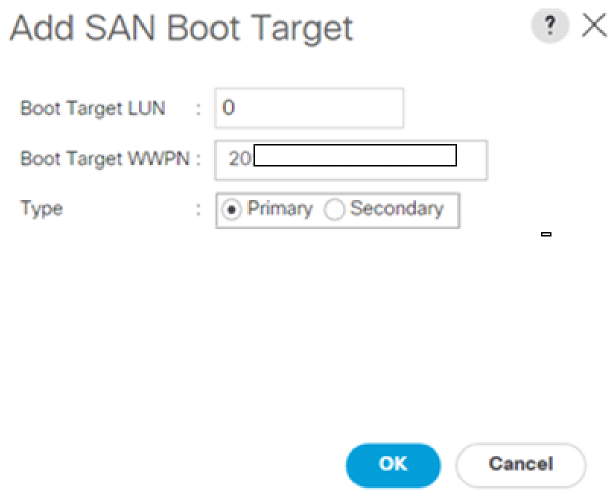

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20[redacted]

Type : ● Primary ○ Secondary

8. Add a secondary SAN Boot Target by clicking Add SAN Boot Target to SAN Primary while the primary SAN Boot option is highlighted. This time enter the AFF A300 WWPN for fcp_02a.

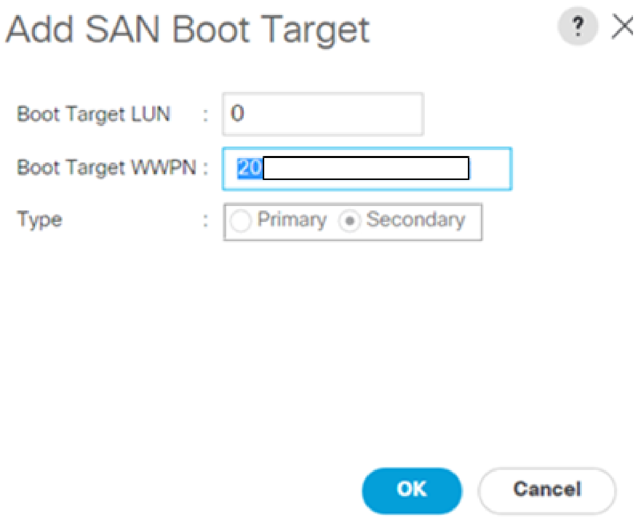

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20[redacted]

Type : ○ Primary ● Secondary

9. Repeat these steps for the secondary SAN boot target and use WWPN fcp_01b and fcp_02b in the primary and secondary SAN boot options.

10. For information about configuring boot and data LUNs on the NetApp A300 storage system, please refer to section NetApp A300 Storage System Configuration.

Create Service Profile Templates for Cisco UCS B-Series

To create service profile templates for the Cisco UCS B-Series environment, complete the following steps:

1. Under the Servers tab in UCSM Select Service Profile Templates.

2. Right-click and select Create Service Profile Template.

3. Name the template B-Series.

4. Select the UUID pool created earlier from the dropdown in the UUID Assignment dialog.

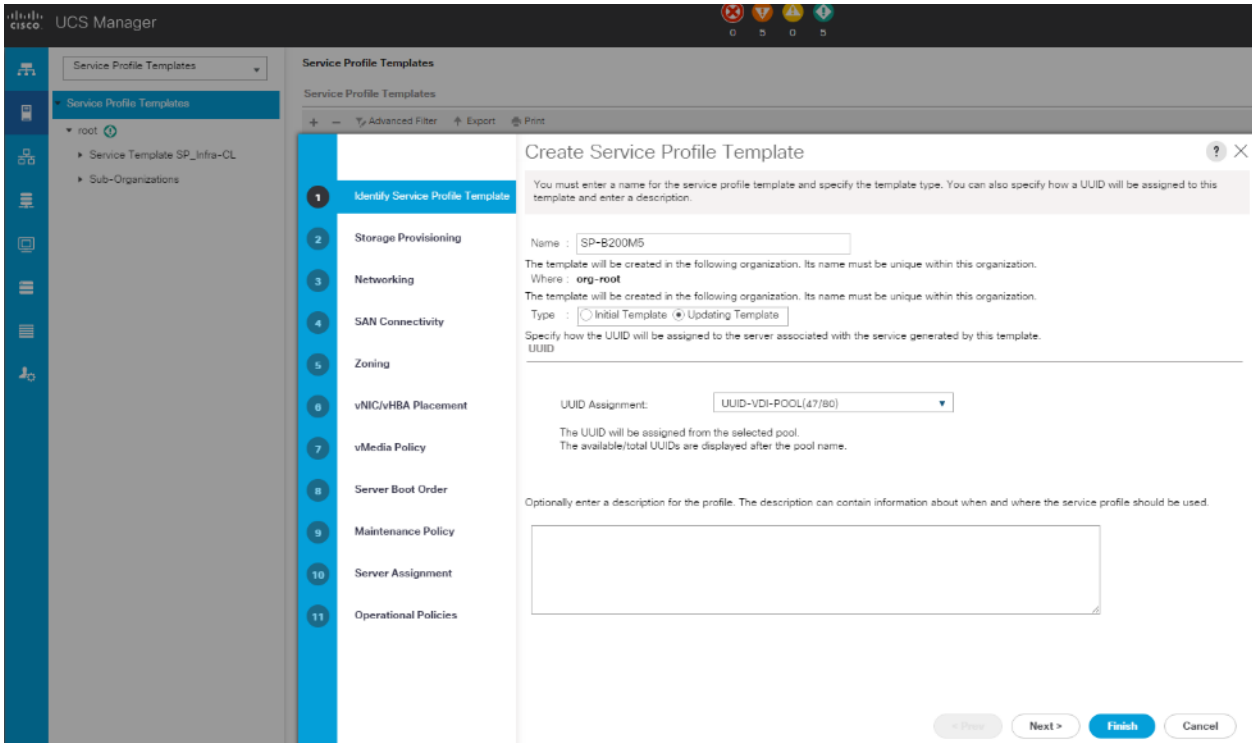

Create Service Profile Template

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 : SP-B200M5

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

Where : **org-root**

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: UUID-VDI-POOL(47/80)

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

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

5. Click Next.

6. Click Next through Storage Provisioning.

7. Under Networking, in the “How would you like to configure LAN connectivity?” dialogue, select the Expert radio button.

8. Click Add.

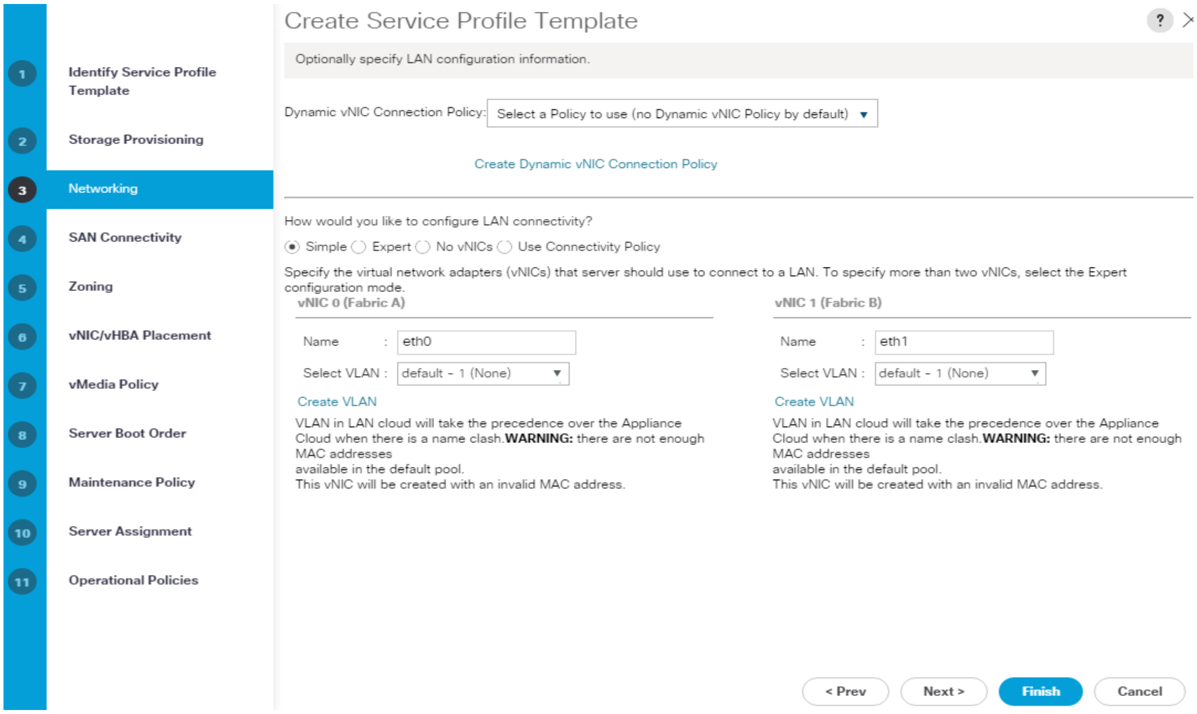

Create Service Profile Template

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

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) ▼

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

● Simple ○ Expert ○ No vNICs ○ Use Connectivity Policy

Specify the virtual network adapters (vNICs) that server should use to connect to a LAN. To specify more than two vNICs, select the Expert configuration mode.

vNIC 0 (Fabric A)

Name : eth0

Select VLAN : default - 1 (None) ▼

Create VLAN

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash. **WARNING:** there are not enough MAC addresses
available in the default pool.
This vNIC will be created with an invalid MAC address.

vNIC 1 (Fabric B)

Name : eth1

Select VLAN : default - 1 (None) ▼

Create VLAN

VLAN in LAN cloud will take the precedence over the Appliance Cloud when there is a name clash. **WARNING:** there are not enough MAC addresses
available in the default pool.
This vNIC will be created with an invalid MAC address.

< Prev | Next > | Finish | Cancel

9. Name it vNIC-A.

10. Select check box for Use vNIC Template.

11. Under vNIC template select the vNIC-A.

12. For Adapter Policy select VMware.

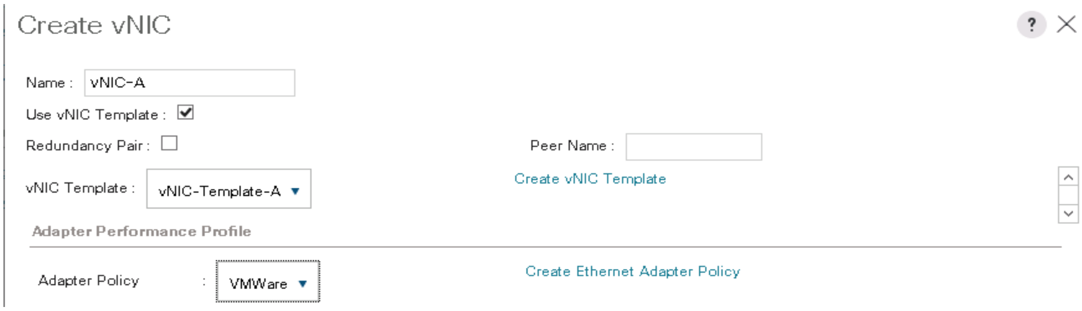

Create vNIC

Name : vNIC-A

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : vNIC-Template-A ▼

Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Ethernet Adapter Policy

13. Repeat networking steps for vNIC-B.

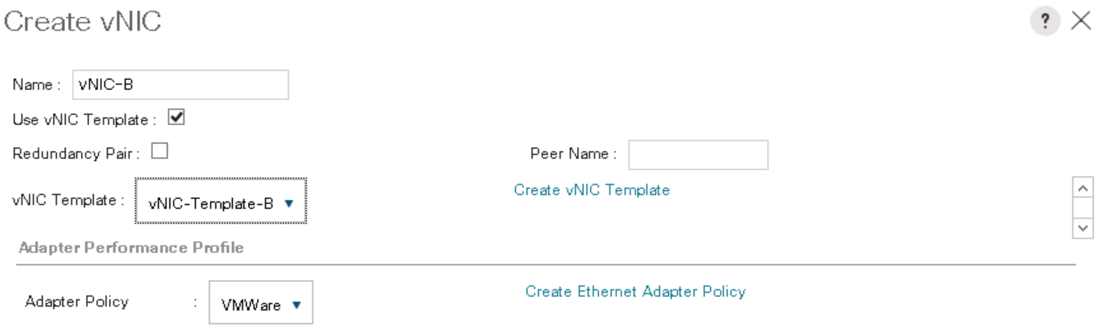

Create vNIC

Name : vNIC-B

Use vNIC Template : ☑

Redundancy Pair : □

Peer Name :

vNIC Template : vNIC-Template-B ▼

Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Ethernet Adapter Policy

14. Click Next.

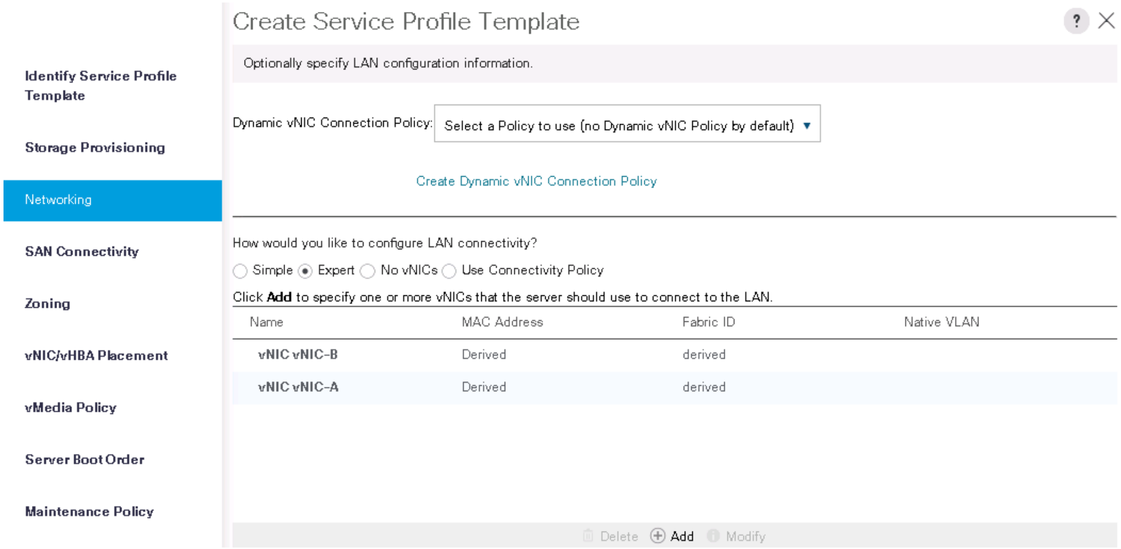

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) ▼

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

○ Simple ● Expert ○ No vNICs ○ Use Connectivity Policy

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

Name	MAC Address	Fabric ID	Native VLAN
vNIC vNIC-B	Derived	derived	
vNIC vNIC-A	Derived	derived	

15. Click Next.

16. Under SAN Connectivity, select the Expert button in the “How would you like to configure SAN Connectivity?

17. Select WWNN Assignment from the Pool created earlier.

18. Click Add.

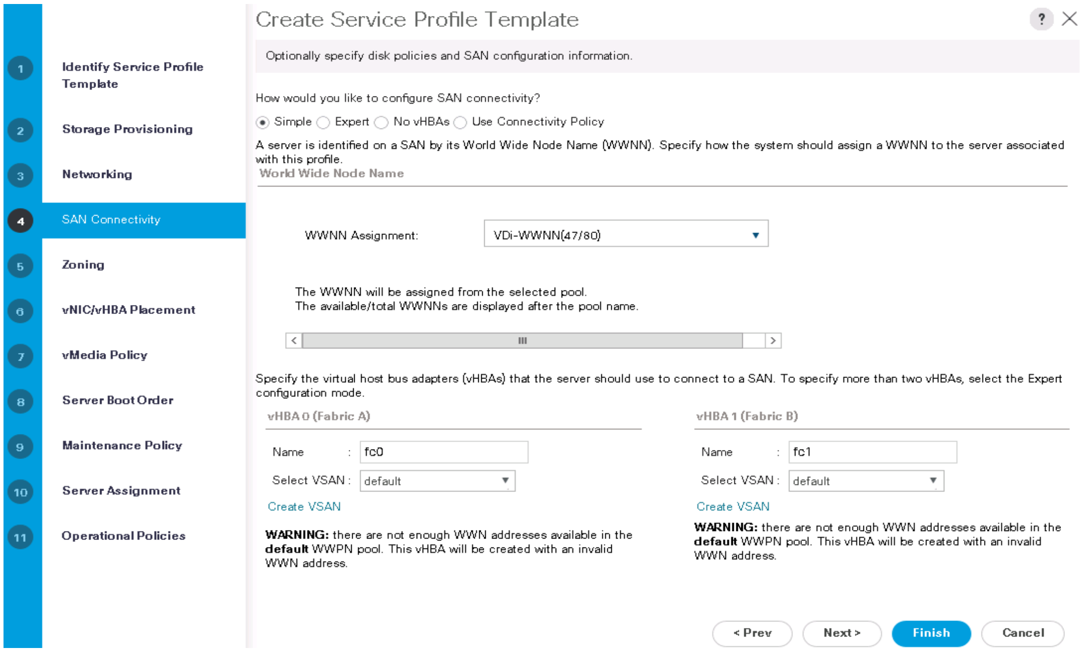

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

● Simple ○ Expert ○ No vHBAs ○ Use Connectivity Policy

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

World Wide Node Name

WWNN Assignment: VDi-WWNN(47/80) ▼

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

Specify the virtual host bus adapters (vHBAs) that the server should use to connect to a SAN. To specify more than two vHBAs, select the Expert configuration mode.

vHBA 0 (Fabric A)

Name : fc0

Select VSAN : default ▼

Create VSAN

WARNING: there are not enough WWN addresses available in the **default** WWPN pool. This vHBA will be created with an invalid WWN address.

vHBA 1 (Fabric B)

Name : fc1

Select VSAN : default ▼

Create VSAN

WARNING: there are not enough WWN addresses available in the **default** WWPN pool. This vHBA will be created with an invalid WWN address.

19. Name the adapter vHBA-A.

20. Click Use vHBA Template.

21. Select vHBA Template: vHBA-A.

22. Select Adapter Policy: VMWare.

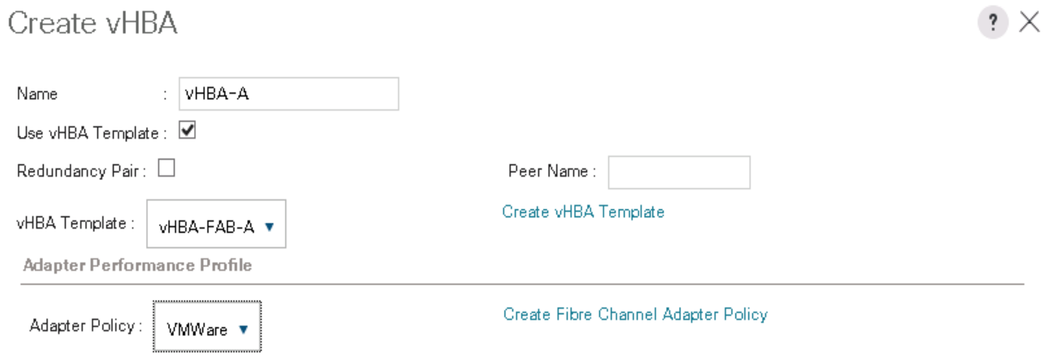

Create vHBA

Name : vHBA-A

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : vHBA-FAB-A ▾

Create vHBA Template

Adapter Performance Profile

Adapter Policy : VMWare ▾

Create Fibre Channel Adapter Policy

23. Repeat steps for vHBA-B on Fabric B.

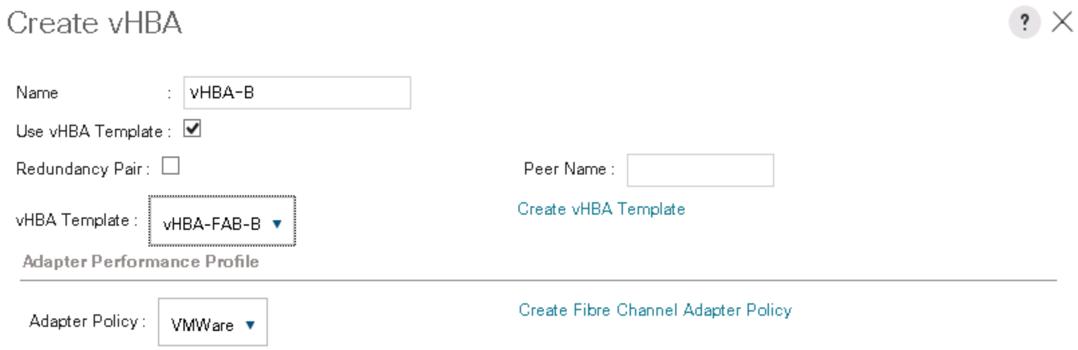

Create vHBA

Name : vHBA-B

Use vHBA Template : ☑

Redundancy Pair : □

Peer Name :

vHBA Template : vHBA-FAB-B ▾

Create vHBA Template

Adapter Performance Profile

Adapter Policy : VMWare ▾

Create Fibre Channel Adapter Policy

24. No Zoning will be used. Click Next.

25. Click Next through vNIC/vHBA Placement policy.

26. Click Next through vMedia Policy.

27. Use the Boot Policy drop down to select the Boot Policy created earlier, then click Finish.

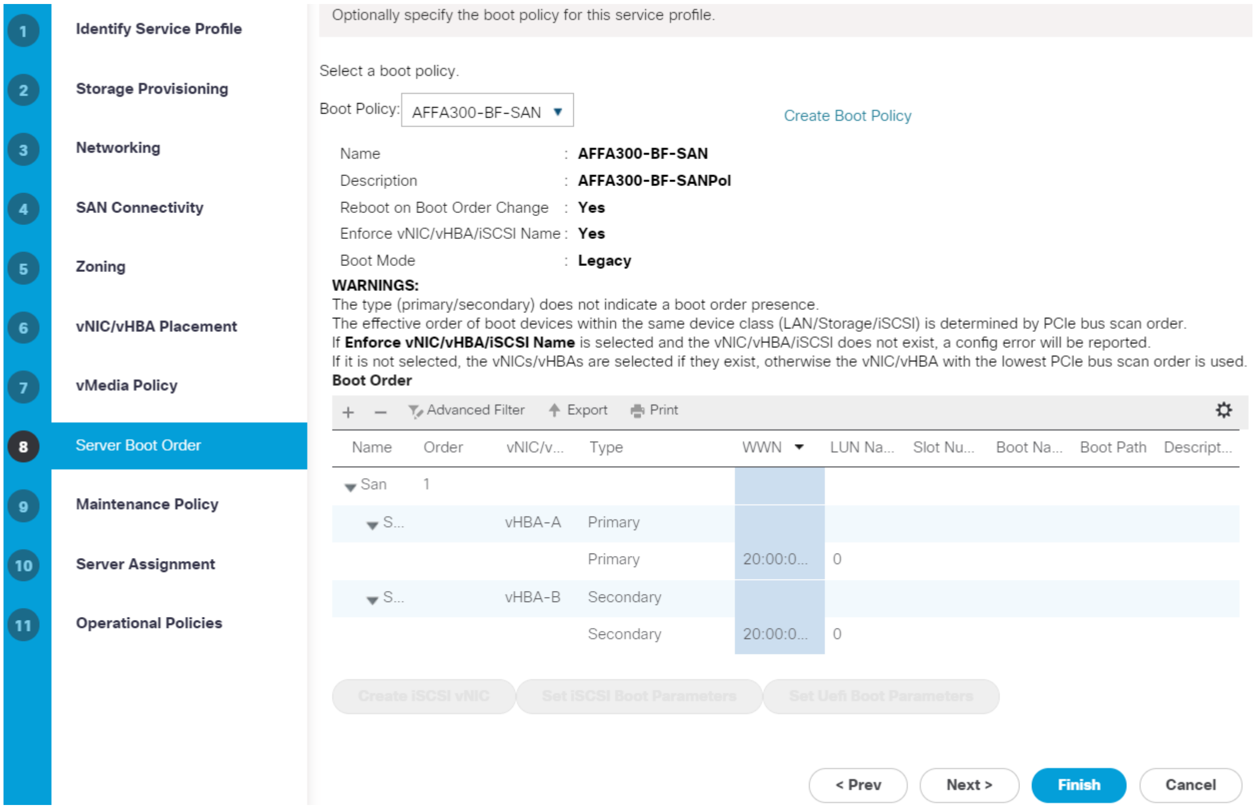

Optionally specify the boot policy for this service profile.

Select a boot policy.

Boot Policy: AFFA300-BF-SAN

Name	: **AFFA300-BF-SAN**
Description	: **AFFA300-BF-SANPol**
Reboot on Boot Order Change	: **Yes**
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/v...	Type	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descript...
San	1								
S...		vHBA-A	Primary						
			Primary	20:00:0...	0				
S...		vHBA-B	Secondary						
			Secondary	20:00:0...	0				

28. Select maintenance Policy and Server Assignment.

29. Click Finish and complete the Service Profile creation.

Create Service Profiles

To create service profiles for each of the blades in the NetApp solution, complete the following steps:

1. From the Servers tab in UCS Manager, under the Service Profile Templates node, right-click the Service Profile Template created in the previous step, then click Create Service Profiles from Template.

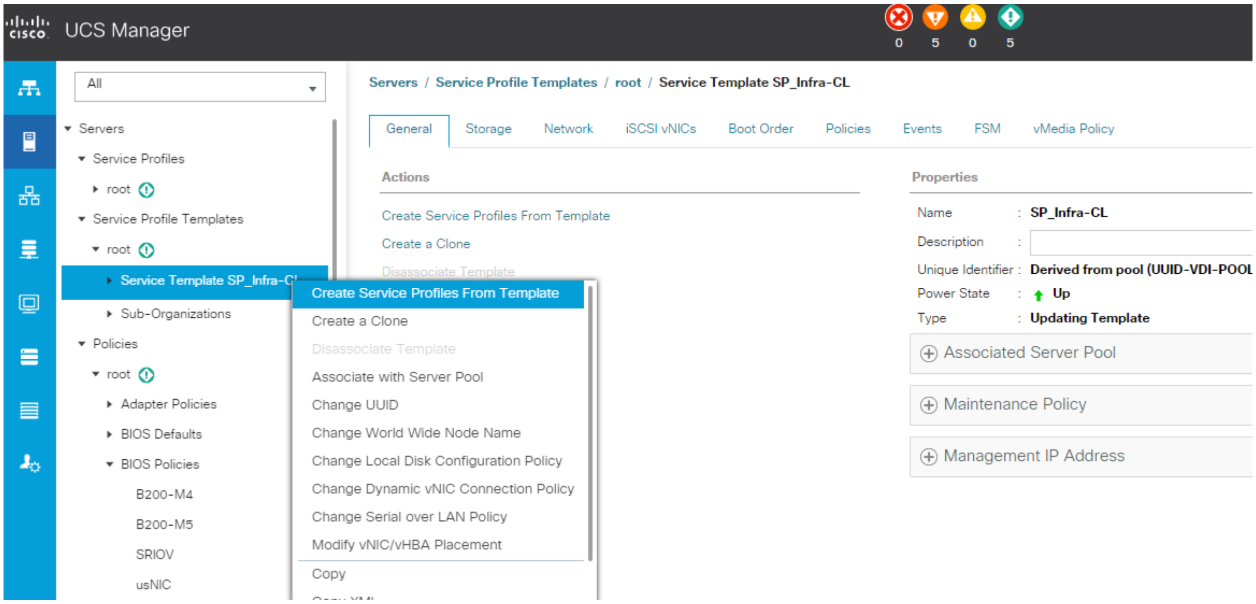
UCS Manager
0 5 0 5
All
Servers
Service Profiles
root
Service Profile Templates
root
Service Template SP_Infra-Cl
Sub-Organizations
Policies
root
Adapter Policies
BIOS Defaults
BIOS Policies
B200-M4
B200-M5
SRIOV
usNIC
Create Service Profiles From Template
Create a Clone
Disassociate Template
Associate 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
Servers / Service Profile Templates / root / Service Template SP_Infra-CL
General
Storage
Network
iSCSI vNICs
Boot Order
Policies
Events
FSM
vMedia Policy
Actions
Create Service Profiles From Template
Create a Clone
Disassociate Template
Properties
Associated Server Pool
Maintenance Policy
Management IP Address

Name	: **SP_Infra-CL**
Description	:
Unique Identifier	: **Derived from pool (UUID-VDI-POOL**
Power State	: **Up**
Type	: **Updating Template**

2. Provide a naming prefix, a starting number, and the number of services profiles to create, then click OK.

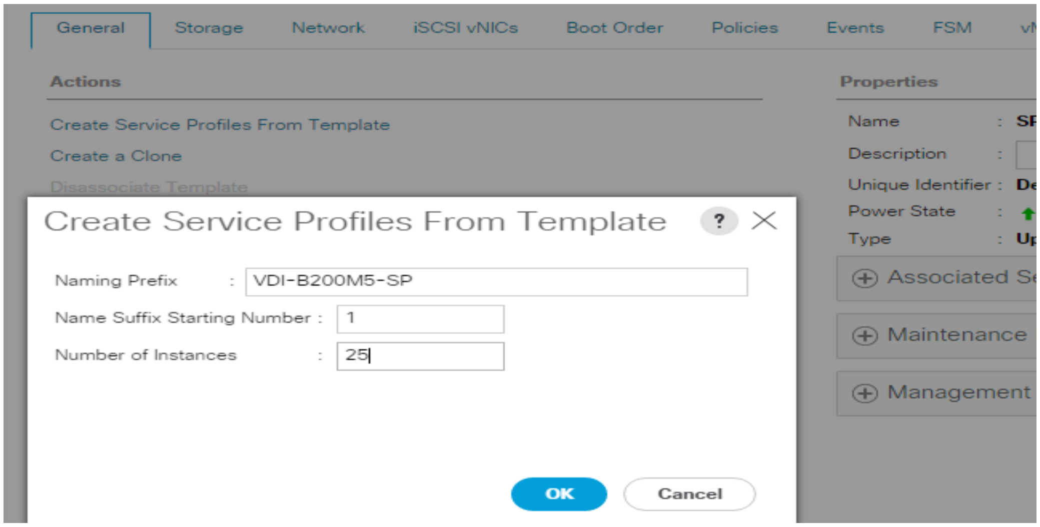

Create Service Profiles From Template

Naming Prefix : VDI-B200M5-SP

Name Suffix Starting Number : 1

Number of Instances : 25

The requested number of service profiles (for example, 25) are created in the Service Profiles root organization.

Configuration of AFF 300 with NetApp ONTAP 9.3

NetApp A300 Storage System Configuration

This design includes instructions on the steps necessary to perform initial setup and configuration of the NetApp A300 storage system. Specific details of the configuration as tested can be found in the NetApp A300 Storage Configuration and NetApp ONTAP9.3 sections below.

NetApp All Flash FAS A300 Controllers

See the following sections (NetApp Hardware Universe) for planning the physical location of the storage systems:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- AFF Systems

NetApp Hardware Universe

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

- Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the HWU application at the NetApp Support site.
- Access the HWU application to view the System Configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

Controllers

Follow the physical installation procedures for the controllers found in the AFF A300 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported by the AFF A300 is available at the NetApp Support site.

NetApp ONTAP 9.3

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the ONTAP 9.1 Software Setup Guide. You must have access to the NetApp Support site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the ONTAP 9.1 Software Setup Guide to learn about configuring ONTAP. Table 15 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 15 ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	`<node01-mgmt-ip>`
Cluster node 01 netmask	`<node01-mgmt-mask>`
Cluster node 01 gateway	`<node01-mgmt-gateway>`
Cluster node 02 IP address	`<node02-mgmt-ip>`
Cluster node 02 netmask	`<node02-mgmt-mask>`
Cluster node 02 gateway	`<node02-mgmt-gateway>`
ONTAP 9.1 URL	`<url-boot-software>`

Configure Node 01

To configure node 01, complete the following steps:

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

```
Starting AUTOBOOT press Ctrl-C to abort…
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

> If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select `e0M` for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for `e0M`.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

9. Enter the URL where the software can be found.

> This web server must be pingable.

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

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

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

13. Press Ctrl-C when the following message displays:

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

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

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

Configure Node 02

To configure node 02, complete the following steps:

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

```
Starting AUTOBOOT press Ctrl-C to abort…
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

> If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and y to reboot the node, then continue with step 14.

4. To install new software, select option 7.

5. Enter y to perform an upgrade.

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

7. Enter y to reboot now.

8. Enter the IP address, netmask, and default gateway for e0M.

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

9. Enter the URL where the software can be found.

> This web server must be pingable.

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

10. Press Enter for the user name, indicating no user name.

11. Enter y to set the newly installed software as the default to be used for subsequent reboots.

12. Enter y to reboot the node.

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

13. Press Ctrl-C when you see this message:

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

14. Select option 4 for Clean Configuration and Initialize All Disks.

15. Enter y to zero disks, reset config, and install a new file system.

16. Enter y to erase all the data on the disks.

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

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.1 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```
Welcome to node setup.

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

You can return to cluster setup at any time by typing “cluster setup”.
To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur
on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

Use your web browser to complete cluster setup by accesing https://<node01-mgmt-ip>

Otherwise press Enter to complete cluster setup using the command line interface:
```

2. To complete the cluster setup, open a web browser and navigate to https://<node01-mgmt-ip.

Table 16 Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>

Cluster Detail	Cluster Detail Value
Cluster node 01 gateway	`<node01-mgmt-gateway>`
Cluster node 02 IP address	`<node02-mgmt-ip>`
Cluster node 02 netmask	`<node02-mgmt-mask>`
Cluster node 02 gateway	`<node02-mgmt-gateway>`
Node 01 service processor IP address	`<node01-SP-ip>`
Node 02 service processor IP address	`<node02-SP-ip>`
DNS domain name	`<dns-domain-name>`
DNS server IP address	`<dns-ip>`
NTP server IP address	`<ntp-ip>`

Cluster setup can also be done using command line interface. This document describes the cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup on the Welcome screen.

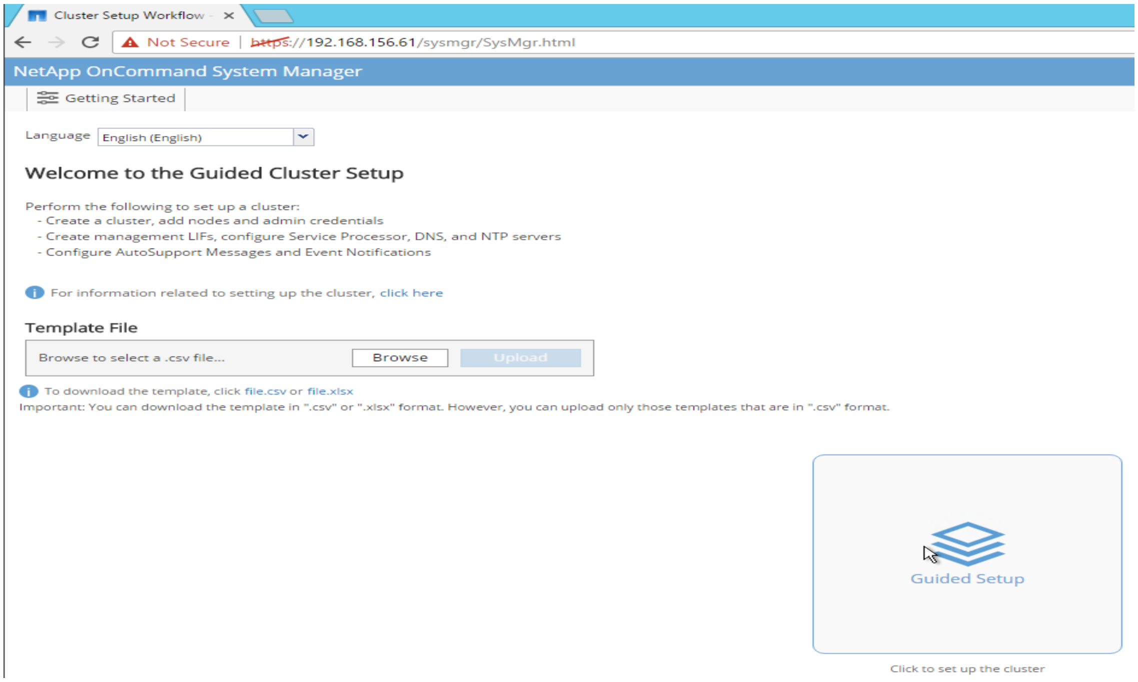
Cluster Setup Workflow - ×
Not Secure | ~~https~~://192.168.156.61/sysmgr/SysMgr.html
NetApp OnCommand System Manager
Getting Started
Language English (English)
Browse to select a .csv file...
Browse
Upload
Guided Setup
Click to set up the cluster

Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:
- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications

For information related to setting up the cluster, click here

Template File

To download the template, click file.csv or file.xlsx

Important: You can download the template in ".csv" or ".xlsx" format. However, you can upload only those templates that are in ".csv" format.

4. In the Cluster screen, do the following:

   a. Enter the cluster and node names.

   b. Select the cluster configuration.

c. Enter and confirm the password.

d. (Optional) Enter the cluster base and feature licenses.

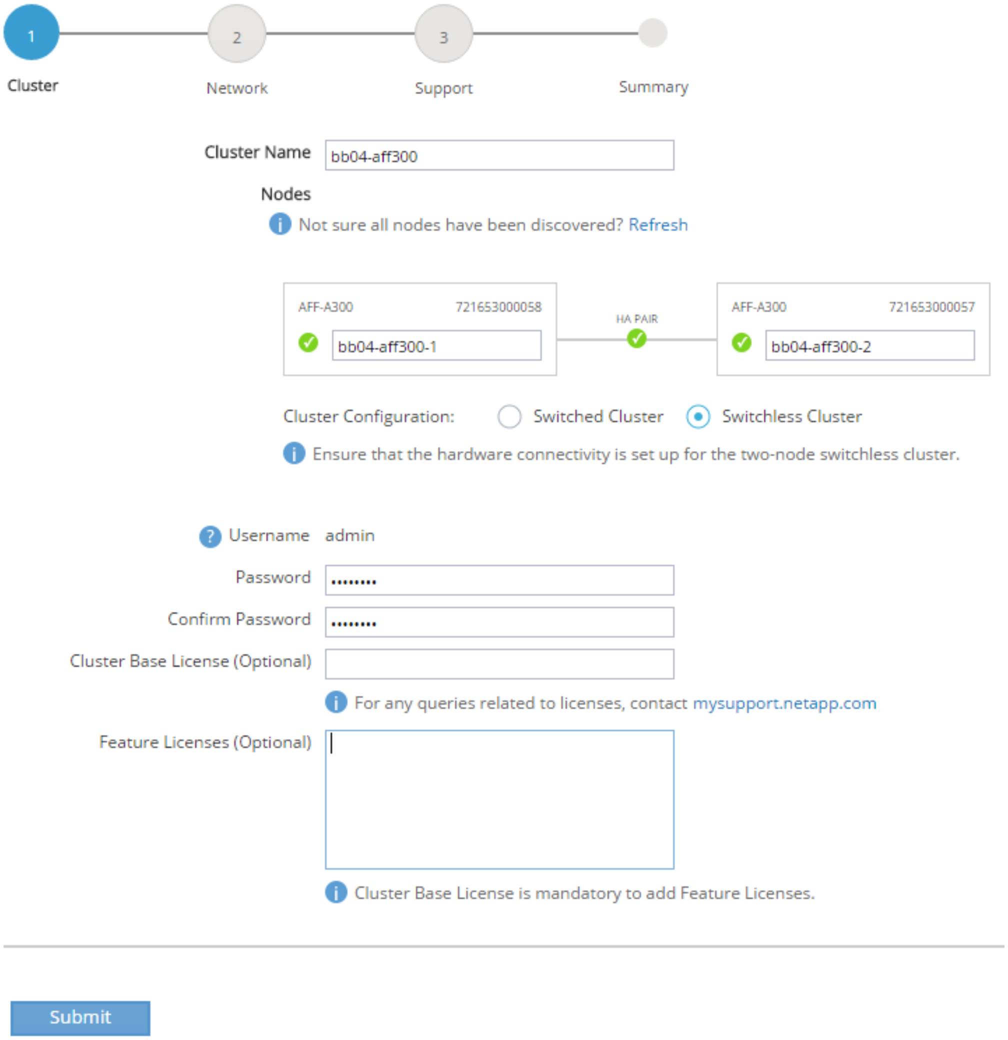

Cluster Name: bb04-aff300

Nodes

Not sure all nodes have been discovered? Refresh

AFF-A300	721653000058	HA PAIR	AFF-A300	721653000057
✓ bb04-aff300-1		✓	✓ bb04-aff300-2	

Cluster Configuration: ○ Switched Cluster ● Switchless Cluster

Ensure that the hardware connectivity is set up for the two-node switchless cluster.

Username: admin

Password: ••••••••

Confirm Password: ••••••••

Cluster Base License (Optional):

For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional):

Cluster Base License is mandatory to add Feature Licenses.

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

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

Cluster license and feature licenses can also be installed after completing the cluster creation.

5. Click Submit.

6. In the network page, complete the following sections:

- Cluster Management
  - Enter the IP address, netmask, gateway and port details.
- Node Management
  - Enter the node management IP addresses and port details for all the nodes.
- Service Processor Management
  - Enter the IP addresses for all the nodes.
- DNS Details
  - Enter the DNS domain names and server address.
- NTP Details
  - Enter the primary and alternate NTP server.

7. Click Submit.

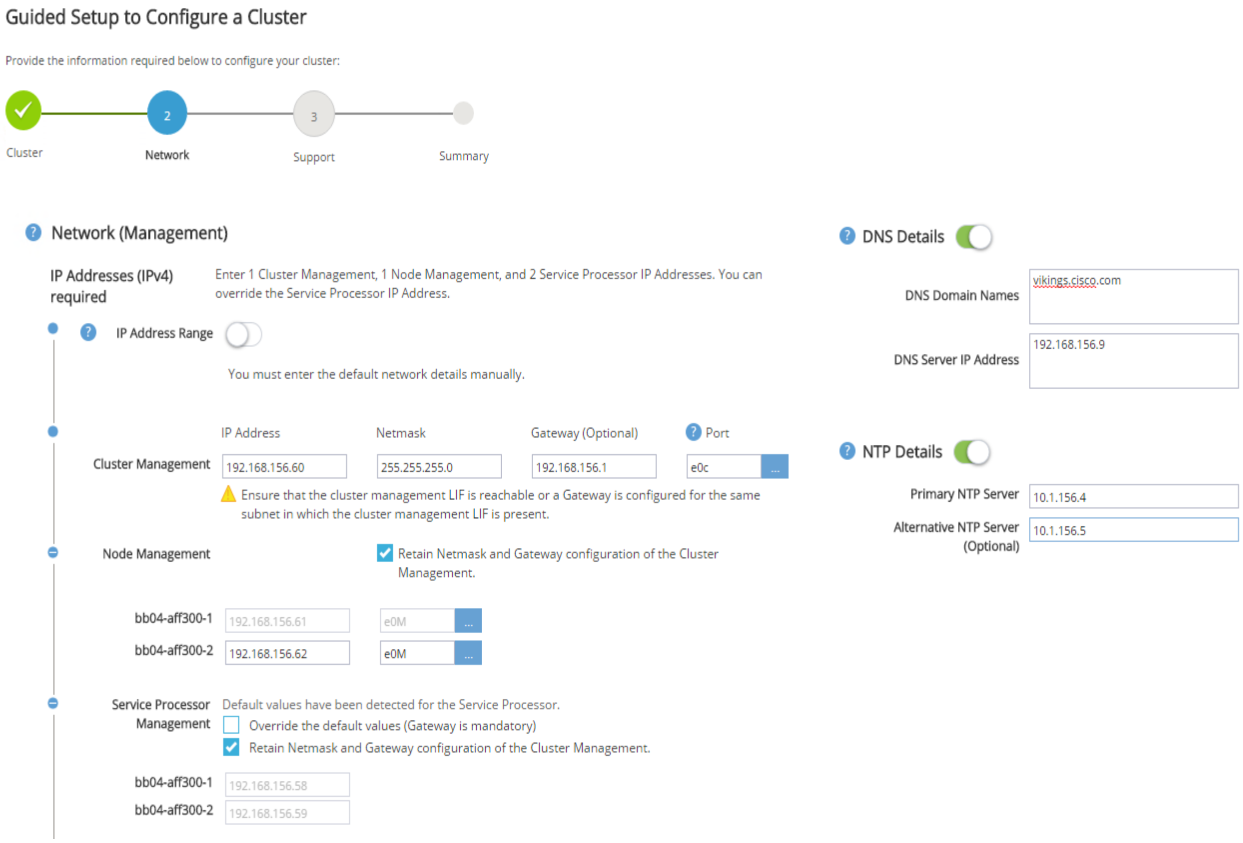

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Network (Management)

IP Addresses (IPv4) required

Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range: off

You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	Port
Cluster Management	192.168.156.60	255.255.255.0	192.168.156.1	e0c

Ensure that the cluster management LIF is reachable or a Gateway is configured for the same subnet in which the cluster management LIF is present.

Node Management: ☑ Retain Netmask and Gateway configuration of the Cluster Management.

bb04-aff300-1	192.168.156.61	e0M
bb04-aff300-2	192.168.156.62	e0M

Service Processor Management: Default values have been detected for the Service Processor.

☐ Override the default values (Gateway is mandatory)

☑ Retain Netmask and Gateway configuration of the Cluster Management.

bb04-aff300-1	192.168.156.58
bb04-aff300-2	192.168.156.59

DNS Details: on

DNS Domain Names: vikings.cisco.com

DNS Server IP Address: 192.168.156.9

NTP Details: on

Primary NTP Server: 10.1.156.4

Alternative NTP Server (Optional): 10.1.156.5

8. In the Support page, configure the AutoSupport and Event Notifications sections.

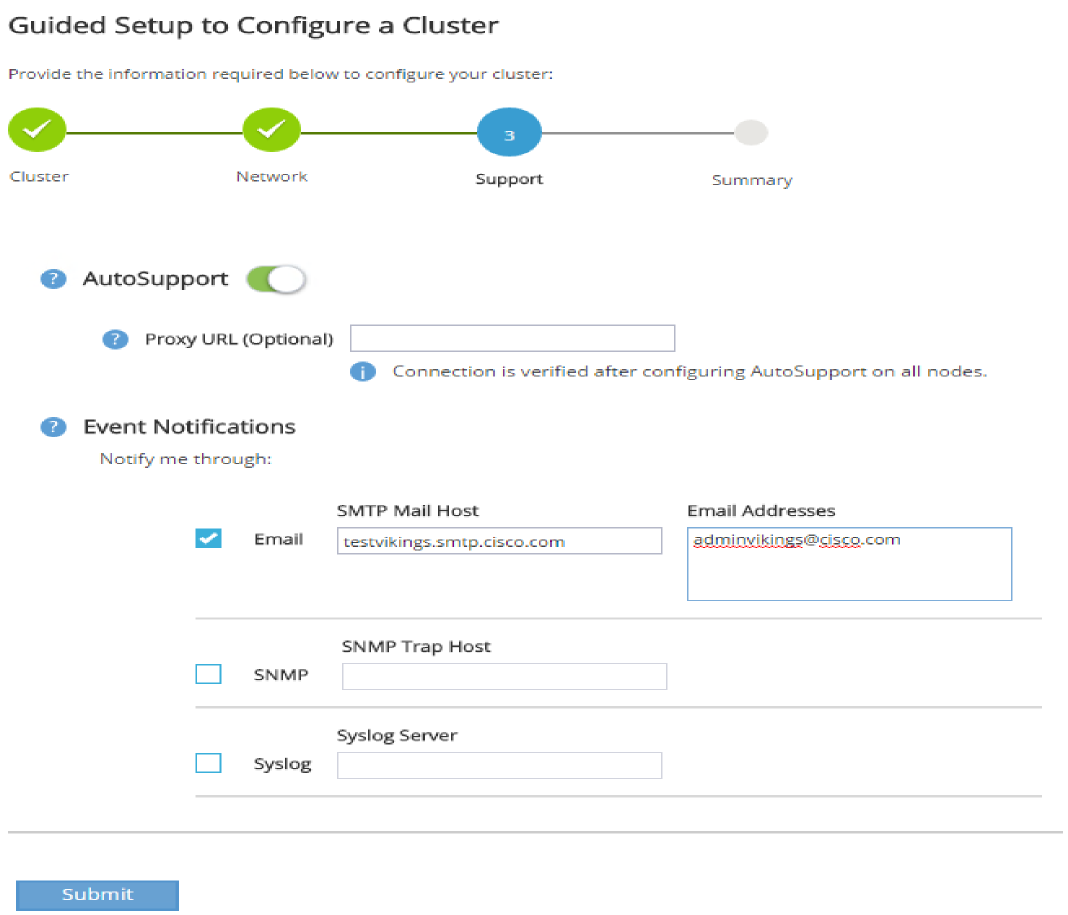

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

AutoSupport: on

Proxy URL (Optional):

Connection is verified after configuring AutoSupport on all nodes.

Event Notifications

Notify me through:

☑ Email — SMTP Mail Host: testvikings.smtp.cisco.com — Email Addresses: adminvikings@cisco.com

□ SNMP — SNMP Trap Host:

□ Syslog — Syslog Server:

9. Click Submit.
10. In the Summary page, review the configuration details if needed.

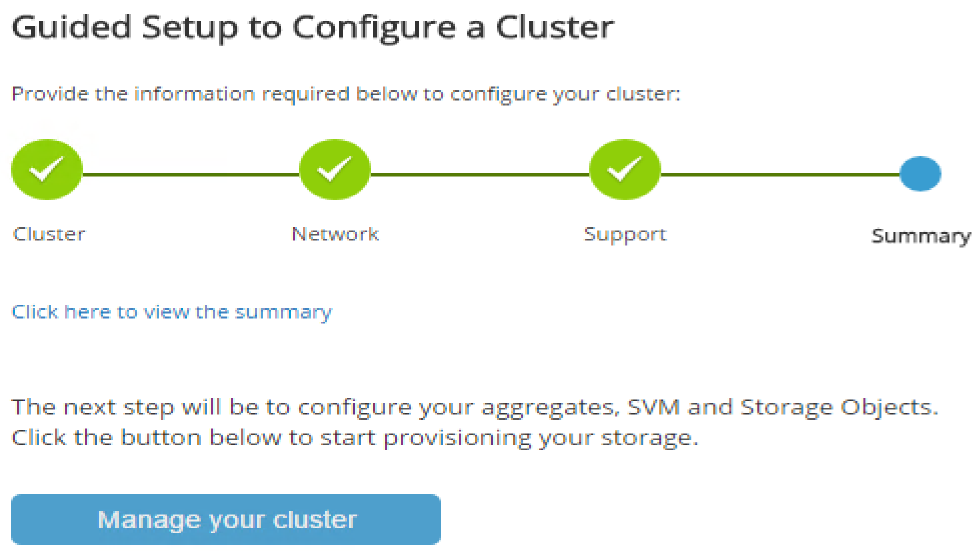

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

The next step will be to configure your aggregates, SVM and Storage Objects. Click the button below to start provisioning your storage.

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

Log into the Cluster

To log in to the cluster, complete the following steps:

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

2. Log in to the admin user with the password you provided earlier.

Zero All Spare Disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

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

Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard unified target adapter 2 (UTA2), complete the following steps:

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin show` command:

```
ucadmin show
                      Current  Current    Pending  Pending    Admin
Node          Adapter Mode     Type       Mode     Type       Status
------------  ------- -------  ---------  -------  ---------  -----------
<st-node01>
              0e      cna      target     -        -          online
<st-node01>
              0f      cna      target     -        -          online
<st-node01>
              0g      fc       target     -        -          online
<st-node01>
              0h      fc       target     -        -          online
<st-node02>
              0e      cna      target     -        -          online
<st-node02>
              0f      cna      target     -        -          online
<st-node02>
              0g      fc       target     -        -          online
<st-node02>
              0h      fc       target     -        -          online
8 entries were displayed.
```

2. Verify that the Current Mode and Current Type properties for all ports are set properly. Ports 0g and 0h are used for FC connectivity and should be set to mode fc if not already configured. The port type for all proto-cols should be set to target. Change the port personality by running the following command:

```
ucadmin modify -node <home-node-of-the-port> -adapter <port-name> -mode fc -type target
```

The ports must be offline to run this command. To take an adapter offline, run the `fcp adapter modify -node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, `0e` and `0f`).

After conversion, a reboot is required. After reboot, bring the ports online by running `fcp adapter modify -node <home-node-of-the-port> -adapter <port-name> -state up`.

Set Auto-Revert on Cluster Management

1. To set the `auto-revert` parameter on the cluster management interface, complete the following step:

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

2. Run the following command:

```
network interface modify -vserver <clustername> -lif cluster_mgmt -auto-revert true
```

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, `e0d`, `e2a`, and `e2e`) should be removed from the default broadcast domain, leaving just the management network ports (`e0c` and `e0M`). To perform this task, run the following commands:

```
broadcast-domain remove-ports  -broadcast-domain Default -ports <st-node01>:e0d, <st-node01>:e0e, <st-
node01>:e0e, <st-node01>:e2a, <st-node01>:e2e, <st-node02>:e0d, <st-node02>:e0e, <st-node02>:e0f, <st-
node02>:e2a, <st-node02>:e2e
broadcast-domain show
```

Set Up Service Processor Network Interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

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

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

The service processor IP addresses should be in the same subnet as the node management IP addresses.

Create Aggregates

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

This solution was validated using 1 data aggregate on each controller with 23 data partitions per aggregate. To create the required aggregates, run the following commands:

```
aggr create -aggregate aggr1_node01 -node <st-node01> -diskcount 23
aggr create -aggregate aggr1_node02 -node <st-node02> -diskcount 23
```

You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate.

For all flash aggregates, you should have a minimum of one hot spare disk or disk partition. For nonflash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For NetApp Flash Pool™ aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.

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

```
aggr show
aggr rename -aggregate aggr0 -newname <node01-rootaggrname>
```

Verify Storage Failover

To confirm that storage failover is enabled, run the following commands for a failover pair:

1. Verify the status of the storage failover.

```
storage failover show
```

Both `<st-node01>` and `<st-node02>` must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.

Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

5. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

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

Disable Flow Control on 10GbE and 40GbE Ports

NetApp recommends disabling flow control on all the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show -fields flowcontrol-admin
```

Disable Unused FCoE Capability on CNA Ports

If the UTA2 port is set to CNA mode and is only expected to handle Ethernet data traffic (for example NFS), then the unused FCoE capability of the port should be disabled by setting the corresponding FCP adapter to state down with the `fcp adapter modify` command. Here are some examples:

```
fcp adapter modify -node <st-node01> -adapter 0e -status-admin down
fcp adapter modify -node <st-node01> -adapter 0f -status-admin down
fcp adapter modify -node <st-node02> -adapter 0e -status-admin down
fcp adapter modify -node <st-node02> -adapter 0f -status-admin down
fcp adapter show -fields status-admin
```

Configure Network Time Protocol

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

```
timezone <timezone>
```

For example, in the eastern United States, the time zone is `America/New_York`.

2. Set the date for the cluster.

```
date <ccyymmddhhmm.ss>
```

The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]> (for example, 201703231549.30).

3. Configure the Network Time Protocol (NTP) servers for the cluster.

```
cluster time-service ntp server create -server <switch-a-ntp-ip>
cluster time-service ntp server create -server <switch-b-ntp-ip>
```

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

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

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

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

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

Configure SNMPv1 Access

To configure SNMPv1 access, set the shared, secret plain-text password (called a community):

```
snmp community add ro <snmp-community>
```

Configure AutoSupport

NetApp AutoSupport® sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

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

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command to enable CDP on ONTAP:

```
node run -node * options cdpd.enable on
```

To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

Create Jumbo Frame MTU Broadcast Domains in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands to create a broadcast domain for NFS on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
```

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2e

ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2e

ifgrp show
```

Create VLANs

To create VLANs, create NFS VLAN ports and add them to the NFS broadcast domain:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <st-node01>:a0a-<infra-nfs-vlan-id>, <st-
node02>:a0a-<infra-nfs-vlan-id>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-security-style
unix
```

3. Remove the unused data protocols from the SVM - CIFS, iSCSI, and NDMP.

```
vserver remove-protocols -vserver Infra-SVM -protocols iscsi,cifs,ndmp
```

4. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp Virtual Storage Console (VSC).

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

5. Enable and run the NFS protocol in the Infra-SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

If NFS license is not installed during the cluster configuration, make sure install the license for staring the NFS service.

6. Turn on the SVM vstorage parameter for the NetApp NFS VAAI plug-in.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled
vserver nfs show
```

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

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

```
volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m01 -type LS -
schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m02 -type LS -
schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol
snapmirror show
```

Create Block Protocol (FC) Service

Run the following command to create the FCP service on each SVM. This command also starts the FCP service and sets the WWN for the SVM.

```
fcp create -vserver Infra-SVM
fcp show
```

If FC license is not installed during the cluster configuration, make sure install the license for creating FC service

Configure HTTPS Access

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

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

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

```
security certificate show
```

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

```
security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -serial
<serial-number>
```

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

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

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

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

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

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

6. Disable HTTP cluster management access.

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

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

7. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set –privilege admin
vserver services web modify –name spi|ontapi|compat –vserver * -enabled true
```

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

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

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

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

```
volume modify –vserver Infra-SVM –volume rootvol –policy default
```

Create FlexVol Volumes

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

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

FlexVol volumes are created to house boot LUNs for ESXi servers, datastore NFS volumes for virtual desktops and RDS hosts, and for the Citrix PVS VDI desktops. For specific details about the volumes created during this validation, see the Storage Configuration section below.

To create a FlexVol volume, run the following command(s):

```
volume create -vserver Infra-SVM -volume vdi_nfs_01 -aggregate aggr1_AFF300_01 -size 10TB -state online -
policy default -space-guarantee none -percent-snapshot-space 0
```

Create FlexGroup Volumes

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

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

FlexGroup volumes are created to house the CIFS shares hosting user profile data and PVS vDisks. For specific details about the volumes created during this validation, see the Storage Configuration section below.

To create a FlexGroup volume, run the following command(s):

```
volume flexgroup deploy -vserver Infra -volume vdi_cifs -size 10TB -space-guarantee none -type RW
```

Create Boot LUNs

Boot LUNs are created for each ESXi host, and data LUNs are created to host virtual desktop and RDS host VMs. For specific details about the LUNs created during this validation, see the Storage Configuration section below. To create boot and data LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VDI-1 -size 10GB -ostype vmware -space-reserve
disabled

lun create -vserver Infra-SVM -volume esxi_boot -lun VDI-2 -size 10GB -ostype vmware -space-reserve
disabled
```

Create igroups

Igroups are created to map host initiators to the LUNs they are allowed to access. Igroups can be FCP protocol, iSCSI protocol, or both. An igroup is created for each ESXi host to map for access to a boot LUN. A separate igroup is created for the entire ESXi cluster to map all data LUNs to every node in the cluster.

1. To create igroups, run the following commands:

```
igroup create –vserver Infra-SVM –igroup VDI-1 –protocol fcp –ostype vmware –initiator <vm-host-VDI-1-
iqn-a>,<vm-host-VDI-1-iqn-b>
igroup create –vserver Infra-SVM –igroup VDI-2 –protocol fcp –ostype vmware –initiator <vm-host-VDI-2-
iqn-a>,<vm-host-VDI-2-iqn-b>
igroup create –vserver Infra-SVM –igroup VDI-3 –protocol fcp –ostype vmware –initiator <vm-host-VDI-3-
iqn-a>,<vm-host-VDI-3-iqn-b>
```

2. To view igroups, type igroup show.

Map Boot LUNs to igroups

To allow access to specific LUNs by specific hosts, map the LUN to the appropriate igroup. For specific details about the LUNs created during this validation, see the Storage Configuration section below. To map luns to igroups, run the following commands:

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-1 –igroup VDI-1 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-2 –igroup VDI-2 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-3 –igroup VDI-3 –lun-id 0
```

Schedule Deduplication

On NetApp All Flash FAS systems, deduplication is enabled by default. To schedule deduplication, complete the following step:

1. After the volumes are created, assign a once-a-day deduplication schedule to `esxi_boot, infra_datastore_1 and infra_datastore_2`:

```
efficiency modify –vserver Infra-SVM –volume esxi_boot –schedule sun-sat@0
efficiency modify –vserver Infra-SVM –volume infra_datastore_1 –schedule sun-sat@0
efficiency modify –vserver Infra-SVM –volume infra_datastore_2 –schedule sun-sat@0
```

Create FC LIFs

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

```
network interface create -vserver Infra-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <st-
node01> -home-port 0g –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <st-
node01> -home-port 0h –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node <st-
node02> -home-port 0g –status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node <st-
node02> -home-port 0h –status-admin up

network interface show
```

Create NFS LIF

To create an NFS LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_lif01 -role data -data-protocol nfs -home-node <st-
node01> -home-port a0a-<infra-nfs-vlan-id> –address <node01-nfs_lif01-ip> -netmask <node01-nfs_lif01-
mask> -status-admin up –failover-policy broadcast-domain-wide –firewall-policy data –auto-revert true

network interface create -vserver Infra-SVM -lif nfs_lif02 -role data -data-protocol nfs -home-node <st-
node02> -home-port a0a-<infra-nfs-vlan-id> –address <node02-nfs_lif02-ip> -netmask <node02-nfs_lif02-
mask>> -status-admin up –failover-policy broadcast-domain-wide –firewall-policy data –auto-revert true

network interface show
```

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, complete the following steps:

1. Run the following commands:

```
network interface create –vserver Infra-SVM –lif svm-mgmt –role data –data-protocol none –home-node <st-
node02> -home-port  e0c –address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up –failover-policy
broadcast-domain-wide –firewall-policy mgmt –auto-revert true
```

The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

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

network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

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

security login unlock -username vsadmin -vserver Infra-SVM
```

A cluster serves data through at least one and possibly several SVMs. We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.

NetApp A300 Storage Configuration

The storage components for this reference architecture are composed of one AFF A300 HA pair and one DS224C disk with 24x 3.8TB solid-state drives. This configuration delivers 65 TB of usable storage and over 200TB effective storage with deduplication, compression and compaction, and the potential for over 300,000 IOPs depending on the application workload.

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

Cluster Details

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

Table 17 lists the cluster details.

Table 17 Cluster Details

Cluster Name	ONTAP Version	Node Count	Data SVM Count	Cluster Raw Capacity
AFF A300	9.3P2	2	1	83.84TB

Storage Details

Table 18 lists the storage details for each HA pair.

Table 18 Storage Details

Node Names	Shelf Count	Disk Count	Disk Capacity	Raw Capacity
AFF-A300-01 AFF-A300-02	DS224-12: 1	SSD: 24	SSD: 83.84TB	83.84TB

Raw capacity is not the same as usable capacity.

Drive Allocation Details

Table 19 lists the drive allocation details for each node.

Table 19 Drive Allocation Details

Node Name	Total Disk Count	Allocated Disk Count	Disk Type	Raw Capacity	Spare Disk Count

Node Name	Total Disk Count	Allocated Disk Count	Disk Type	Raw Capacity	Spare Disk Count
AFF-A300-01	12	12	3.8TB_SSD	41.92TB	0
AFF-A300-02	12	12	3.8TB_SSD	41.92TB	0

Raw capacity is not the same as usable capacity.

Adapter Card Details

Table 20 lists the adapter cards present in each node.

Table 20 Adapter Card Details

Node Name	System Model	Slot Number	Part Number	Description
AFF-A300-01	AFF A300	1	X2069	PMC PM8072; PCI-E quad-port SAS (PM8072)
AFF-A300-01	AFF A300	2	X1144A	NIC,2x40GbE,QSFP
AFF-A300-02	AFF A300	1	X2069	PMC PM8072; PCI-E quad-port SAS (PM8072)
AFF-A300-02	AFF A300	2	X1144A	NIC,2x40GbE,QSFP

Firmware Details

Table 21 lists the relevant firmware details for each node.

Table 21 Firmware Details

Node Name	Node Firmware	Shelf Firmware	Drive Firmware	Remote Mgmt Firmware
AFF-A300-01	AFF-A300: 11.1	IOM12: A:0210, B:0210	X357_S163A3T8ATE: NA51	SP: 5.1
AFF-A300-02	AFF-A300: 11.1	IOM12: A:0210, B:0210	X357_S163A3T8ATE: NA51	SP: 5.0X21

Network Port Settings

You can modify the MTU, auto-negotiation, duplex, flow control, and speed settings of a physical network port or interface group.

Table 22 lists the network port settings.

Table 22 Network Port Settings for ONTAP

Node Name	Port Name	Link Status	Port Type	MTU Size	Flow Control (Admin/Oper)	IPspace Name	Broadcast Domain
AFF-A300-01	a0a	up	if_group	9000	full/-	Default	

Node Name	Port Name	Link Status	Port Type	MTU Size	Flow Control (Admin/Oper)	IPspace Name	Broadcast Domain
AFF-A300-01	a0a-61	up	vlan	1500	full/-	Default	IB
AFF-A300-01	a0a-62	up	vlan	1500	full/-	Default	cifs
AFF-A300-01	a0a-63	up	vlan	9000	full/-	Default	nfs
AFF-A300-01	e0a	up	physical	9000	none/none	Cluster	Cluster
AFF-A300-01	e0b	up	physical	9000	none/none	Cluster	Cluster
AFF-A300-01	e0c	down	physical	1500	none/none	Default	Default
AFF-A300-01	e0d	down	physical	1500	none/none	Default	
AFF-A300-01	e0e	up	physical	1500	none/none	Default	
AFF-A300-01	e0f	up	physical	1500	none/none	Default	
AFF-A300-01	e0M	up	physical	1500	full/full	Default	Default
AFF-A300-01	e2a	up	physical	9000	none/none	Default	
AFF-A300-01	e2e	up	physical	9000	none/none	Default	
AFF-A300-02	a0a	up	if_group	9000	full/-	Default	
AFF-A300-02	a0a-61	up	vlan	1500	full/-	Default	IB
AFF-A300-02	a0a-62	up	vlan	1500	full/-	Default	cifs
AFF-A300-02	a0a-63	up	vlan	9000	full/-	Default	nfs
AFF-A300-02	e0a	up	physical	9000	none/none	Cluster	Cluster
AFF-A300-02	e0b	up	physical	9000	none/none	Cluster	Cluster
AFF-A300-02	e0c	down	physical	1500	none/none	Default	Default

Node Name	Port Name	Link Status	Port Type	MTU Size	Flow Control (Admin/Oper)	IPspace Name	Broadcast Domain
AFF-A300-02	e0d	down	physical	1500	none/none	Default	
AFF-A300-02	e0e	up	physical	1500	none/none	Default	
AFF-A300-02	e0f	up	physical	1500	none/none	Default	
AFF-A300-02	e0M	up	physical	1500	full/full	Default	Default
AFF-A300-02	e2a	up	physical	9000	none/none	Default	
AFF-A300-02	e2e	up	physical	9000	none/none	Default	

Network Port Interface Group Settings

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

Table 23 lists the network port ifgrp settings.

Table 23 Network Port Ifgrp Settings

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

Network Routes

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

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

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

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

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

Error! Reference source not found.lists the network routes for Data ONTAP 8.3 or later.

Table 24 Network Routes

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

Network Port Broadcast Domains

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

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

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

Error! Reference source not found.lists the network port broadcast domains for Data ONTAP 8.3 or later.

Table 25 Network Port Broadcast Domains

Cluster Name	Broadcast Domain	IPspace Name	MTU Size	Subnet Names	Port List	Failover Group Names
AFF-A300	Cifs	Default	1500		AFF-A300-01:a0a-62 AFF-A300-02:a0a-62	cifs
AFF-A300	Cluster	Cluster	9000		AFF-A300-01:e0a AFF-A300-01:e0b AFF-A300-02:e0a AFF-A300-02:e0b	Cluster
AFF-A300	Default	Default	1500		AFF-A300-01:e0c AFF-A300-01:e0M AFF-A300-02:e0c AFF-A300-02:e0M	Default

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

Aggregate Configuration

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

- For business-critical applications that need the lowest possible latency and the highest possible performance, you might create an aggregate consisting entirely of SSDs.
- To tier data with different access patterns, you can create a hybrid aggregate, deploying flash as high-performance cache for a working data set, while using lower-cost HDDs or object storage for less frequently accessed data. A Flash Pool consists of both SSDs and HDDs. A Fabric Pool consists of an all-SSD aggregate with an attached object store.
- If you need to segregate archived data from active data for regulatory purposes, you can use an aggregate consisting of capacity HDDs, or a combination of performance and capacity HDDs.

Error! Reference source not found.contains all aggregate configuration information.

Table 26 Aggregate Configuration

Aggregate Name	Home Node Name	State	RAID Status	RAID Type	Disk Count (By Type)	RG Size (HDD / SSD)	HA Policy	Has Mroot	Mirrored	Size Nominal
aggr0_A300_01	AFF-A300-01	online	normal	raid_dp	11@3.8TB_SSD (Shared)	24	cfo	True	False	414.47GB
aggr0_A300_02	AFF-A300-02	online	normal	raid_dp	10@3.8TB_SSD (Shared)	24	cfo	True	False	368.42GB
aggr1_AFF300_01	AFF-A300-01	online	normal	raid_dp	23@3.8TB_SSD (Shared)	24	sfo	False	False	32.51TB
aggr1_AFF300_02	AFF-A300-02	online	normal	raid_dp	23@3.8TB_SSD (Shared)	24	sfo	False	False	32.51TB

Storage Virtual Machines

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

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

SVM Configuration

Table 27 lists the SVM configuration.

Table 27 SVM Configuration

Cluster Name	SVM Name	Type	Subtype	State	Allowed Protocols	Name Server Switch	Name Mapping Switch	Comment
AFF-A300	Infra	data	default	running	running	cifs, fcp	AFF-A300	

SVM Storage Configuration

Table 28 lists the SVM storage configuration.

Table 28 SVM Storage Configuration

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

Volume Configuration

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

Error! Reference source not found.the FlexVol configuration.

Table 29 FlexVol Configuration

Cluster Name	SVM Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
AFF-A300	Infra	esxi_boot	aggr1_AFF300_01	RW	default	default	unix	500.00GB
AFF-A300	Infra	home	aggr1_AFF300_02	RW	default	default	ntfs	1.00TB

Cluster Name	SVM Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
AFF-A300	Infra	infra_nfs_ds01	aggr1_AFF300_01	RW	default	default	unix	6.00TB
AFF-A300	Infra	vdi_cifs (FlexGroup)	aggr1_AFF300_01 & 02	RW	default	default	unix	1.00TB
AFF-A300	Infra	vdi_cifs_vDisk	aggr1_AFF300_02 & 02	RW	default	default	unix	500GB
AFF-A300	Infra	vdi_nfs_ds01	aggr1_AFF300_01	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds02	aggr1_AFF300_02	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds03	aggr1_AFF300_01	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds04	aggr1_AFF300_02	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds05	aggr1_AFF300_01	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds06	aggr1_AFF300_02	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds07	aggr1_AFF300_01	RW	default	default	unix	10.00TB
AFF-A300	Infra	vdi_nfs_ds08	aggr1_AFF300_02	RW	default	default	unix	10.00TB

Protocol Configuration

NAS

ONTAP can be accessed over Common Internet File System (CIFS), Server Message Block (SMB) and Network File System (NFS) capable clients.

This means clients can access all files on a SVM regardless of the protocol they are connecting with or the type of authentication they require.

Logical Interfaces

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

LIFs can be hosted on the following ports:

- Physical ports that are not part of interface groups
- Interface groups
- VLANs

- Physical ports or interface groups that host VLANs

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

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

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

NAS Logical Interface Settings

Error! Reference source not found.lists the NAS LIF settings.

Table 30 NAS LIF Settings

Cluster Name	SVM Name	Interface Name	Status (Admin/Oper)	IP Address	Current Node	Current Port	Is Home
AFF-A300	Infra	CIFS1-01	up/up	10.10.62.10/24	AFF-A300-01	a0a-62	False
AFF-A300	Infra	CIFS2-02	up/up	10.10.62.11/24	AFF-A300-02	a0a-62	True
AFF-A300	Infra	mgmt2	up/up	10.10.61.26/24	AFF-A300-01	a0a-61	True
AFF-A300	Infra	NFS1-01	up/up	10.10.63.10/24	AFF-A300-01	a0a-63	False
AFF-A300	Infra	NFS2-02	up/up	10.10.63.11/24	AFF-A300-02	a0a-63	True

Windows File Services

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

CIFS Servers

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

Error! Reference source not found.lists CIFS server configuration information.

Table 31 CIFS Servers

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

CIFS Options

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

Error! Reference source not found.lists CIFS options.

Table 32 **CIFS Options**

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

CIFS Local Users and Groups

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

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

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

Privileges do not provide ONTAP general administrative capabilities.

CIFS Shares

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

Error! Reference source not found.lists the CIFS shares.

Table 33 **CIFS Shares**

Cluster Name	SVM Name	Share Name	Path	Share Properties	Symlink Properties	Share ACL
AFF-A300	Infra	%w	%w	homedirectory	symlinks	Everyone:Full Control
AFF-A300	Infra	admin$	/	browsable		UTD
AFF-A300	Infra	c$	/	oplocks browsable changenotify show_previous_versions	symlinks	Administrators:Full Control

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

CIFS Home Directory Search Paths

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

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

Error! Reference source not found.lists the CIFS home directory search paths.

Table 34 CIFS Home Directory Search Paths

Cluster Name	SVM Name	Position	Path
AFF-A300	Infra	1	/home/LoginVSI

SAN

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

LUNs

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

Error! Reference source not found.lists the LUN details.

Table 35 LUN Configuration

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/esxi_boot/VDI-1	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-2	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-3	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-4	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-5	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-6	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-7	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-9	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-10	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-11	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-12	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-13	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-14	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-15	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-17	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-18	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-19	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-20	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-21	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-22	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-23	True	True	vmware	False	10.00GB

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF-A300	Infra	/vol/esxi_boot/VDI-24	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-25	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-26	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-27	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-28	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-29	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-30	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-31	True	True	vmware	False	10.00GB
AFF-A300	Infra	/vol/esxi_boot/VDI-32	True	True	vmware	False	10.00GB

Initiator Groups

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

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

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

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

Error! Reference source not found.lists the igroups that have been created.

Table 36 Initiator Groups

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF-A300	Infra	VDI-1	fcp	vmware	True	full
AFF-A300	Infra	VDI-2	fcp	vmware	True	full
AFF-A300	Infra	VDI-3	fcp	vmware	True	full
AFF-A300	Infra	VDI-4	fcp	vmware	True	full
AFF-A300	Infra	VDI-5	fcp	vmware	True	full

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF-A300	Infra	VDI-6	fcp	vmware	True	full
AFF-A300	Infra	VDI-7	fcp	vmware	True	full
AFF-A300	Infra	VDI-9	fcp	vmware	True	full
AFF-A300	Infra	VDI-10	fcp	vmware	True	full
AFF-A300	Infra	VDI-11	fcp	vmware	True	full
AFF-A300	Infra	VDI-12	fcp	vmware	True	full
AFF-A300	Infra	VDI-13	fcp	vmware	True	full
AFF-A300	Infra	VDI-14	fcp	vmware	True	full
AFF-A300	Infra	VDI-15	fcp	vmware	True	full
AFF-A300	Infra	VDI-17	fcp	vmware	True	full
AFF-A300	Infra	VDI-18	fcp	vmware	True	full
AFF-A300	Infra	VDI-19	fcp	vmware	True	full
AFF-A300	Infra	VDI-20	fcp	vmware	True	full
AFF-A300	Infra	VDI-21	fcp	vmware	True	full
AFF-A300	Infra	VDI-22	fcp	vmware	True	full
AFF-A300	Infra	VDI-23	fcp	vmware	True	full
AFF-A300	Infra	VDI-24	fcp	vmware	True	full
AFF-A300	Infra	VDI-25	fcp	vmware	True	full
AFF-A300	Infra	VDI-26	fcp	vmware	True	full
AFF-A300	Infra	VDI-27	fcp	vmware	True	full
AFF-A300	Infra	VDI-28	fcp	vmware	True	full
AFF-A300	Infra	VDI-29	fcp	vmware	True	full
AFF-A300	Infra	VDI-30	fcp	vmware	True	full
AFF-A300	Infra	VDI-31	fcp	vmware	True	full
AFF-A300	Infra	VDI-32	fcp	vmware	True	full

FCP Logical Interface Settings

Error! Reference source not found.lists the FCP LIF settings.

Table 37 FCP LIF Settings

Cluster Name	SVM Name	Interface Name	Status (Admin/Oper)	Port Name	Current Node	Current Port	Is Home
AFF-A300	Infra	fcp_01a	up/up	20:01:00:a0:98:af:bd:e8	AFF-A300-01	0g	True
AFF-A300	Infra	fcp_01b	up/up	20:02:00:a0:98:af:bd:e8	AFF-A300-01	0h	True
AFF-A300	Infra	fcp_02a	up/up	20:03:00:a0:98:af:bd:e8	AFF-A300-02	0g	True
AFF-A300	Infra	fcp_02b	up/up	20:04:00:a0:98:af:bd:e8	AFF-A300-02	0h	True

FCP / FCoE

FCP Service Configuration

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

Error! Reference source not found.lists the FCP service configuration details.

Table 38 FCP Service Configuration

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

FCP Adapter Configuration

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

Error! Reference source not found.lists the details of the storage controller target ports and the WWPN address of each.

Table 39 FCP Adapter Configuration

Node Name	Adapter Name	State	Data Link Rate	Media Type	Speed	Port Name
AFF-A300-01	0e	offlined by user/system	0	ptp	auto	50:0a:09:82:80:13:41:27
AFF-A300-01	0f	offlined by user/system	0	ptp	auto	50:0a:09:81:80:13:41:27

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

Storage Efficiency and Space Management

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

Volume Efficiency

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

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

Error! Reference source not found.lists the volume efficiency settings.

Table 40 Volume Efficiency Settings

Cluster Name	SVM Name	Volume Name	Space Guarantee	Dedupe	Schedule Or Policy Name	Compression	Inline Compression
AFF-A300	Infra	esxi_boot	none	True	sun-sat@1	True	True
AFF-A300	Infra	infra_nfs_ds01	none	True	inline-only	True	True
AFF-A300	Infra	svm_root	volume		-		
AFF-A300	Infra	vdi_cifs	none	True	inline-only	True	True
AFF-A300	Infra	vdi_cifs_vDisk	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds01	none	True	inline-only	True	True

Cluster Name	SVM Name	Volume Name	Space Guarantee	Dedupe	Schedule Or Policy Name	Compression	Inline Compression
AFF-A300	Infra	vdi_nfs_ds02	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds03	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds04	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds05	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds06	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds07	none	True	inline-only	True	True
AFF-A300	Infra	vdi_nfs_ds08	none	True	inline-only	True	True

LUN Efficiency

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

With thin provisioning, you can present more storage space to the hosts connecting to the SVM than what is actually available on the SVM. Storage provisioning with thinly provisioned LUNs enables storage administrators to provide actual storage that the LUN needs. As ONTAP writes blocks to the LUN, the LUN increases in size automatically.

Error! Reference source not found.lists the LUN efficiency settings.

Table 41 LUN Efficiency Settings

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF-A300	Infra	/vol/esxi_boot/VDI-1	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-2	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-3	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-4	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-5	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-6	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-7	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-9	False	False

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF-A300	Infra	/vol/esxi_boot/VDI-10	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-11	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-12	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-13	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-14	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-15	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-17	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-18	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-19	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-20	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-21	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-22	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-23	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-24	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-25	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-26	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-27	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-28	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-29	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-30	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-31	False	False
AFF-A300	Infra	/vol/esxi_boot/VDI-32	False	False

NetApp Storage Configuration for CIFS Shares

CIFS in Cluster Data ONTAP

NetApp is the leader in providing a fully functional CIFS storage server. NetApp has been providing CIFS server functions since SMB1 and NetApp provides support for SMB 2.0, 2.1 and 3.0. The benefit of using the integrated CIFS functionality within the storage array is that it removes need to have the IO processed twice. With a Windows File Server environment, the data is processed at the Windows File Server layer and then passed on to be processed by the storage array. With NetApp's CIFS functionality, the client maps the share on the NetApp storage cluster directly; therefore, the IO is only processed at the storage array level. In the NetApp CIFS model, the requirement for separate Windows file servers is removed, which then removes the overhead of having the data processed at the Windows File Server.
Windows File Services Features in Clustered Data ONTAP 9.3

Clustered Data ONTAP 9.3 contains the following new CIFS features:

- Microsoft Management Console (MMC) support for viewing and managing open files, open sessions, and shares
- NetBIOS aliases
- Storage-Level Access Guard (SLAG)
- Native file-access auditing for user logon and logoff
- Group Policy object (GPO) security policy support
- NetApp FPolicy pass-through read support
- Offloaded data transfer (ODX) enhancements
- Support for Microsoft Dynamic Access Control (DAC)

Table 42 presents a complete list of CIFS features.

Table 42 9.3 CIFS Features in Clustered Data ONTAP

CIFS Features
Support for Microsoft DAC (Dynamic Access Control)
AES 128/256 for CIFS Kerberos authentication
ODX direct-copy
MMC support for viewing and managing open files and sessions
NetBIOS aliases
SLAG
Native auditing for logon and logoff to shares
UNIX character mapping
GPO security policy support
FPolicy pass-through read support
CIFS restrict anonymous capability

CIFS Features
Control bypass traverse checking
CIFS home directory show user command
Control of CIFS home directory access for admins
Multidomain user mapping
LDAP over SSL (start-TLS)
Offbox antivirus
Separate AD licensing
SMB 3.0 , SMB 2.1, and SBM 2.0
Copy offload
SMB autolocation
BranchCache
Local users and groups
FSecurity
FPolicy
Roaming profiles and folder redirection
Access-based enumeration (ABE)
Offline folders
SMB signing
Remove VSS
File access auditing or file access monitoring

Best Practices
• Use CIFS shares on the NetApp storage cluster instead of a Windows File Server VM • Use CIFS shares on the NetApp storage cluster for VDI Home directories, VDI profiles, and other VDI CIFS data.

User Home Data

In this reference architecture, we did not put the user home directories on the NetApp storage but if you do not have a CIFS home directory storage array, you can put the Home Directories/Data on the NetApp storage too. With user home data, this data is the intellectual property of each company and is directly generated by the end user. In a virtual desktop environment, the home directory data located in a NetApp volume is shared through the CIFS protocol and is mapped as a drive letter in the virtual desktop. This data often requires backup and recovery as well as disaster recovery services. Using a CIFS home directory brings more efficiency in the management and protection of user data. End-user data files should be deduplicated and compressed to achieve storage efficiency and reduce the overall solution cost.

Best Practices
• Use deduplication and compression for end-user data files stored in home directories to achieve storage efficiency. NetApp strongly recommends storing user data on the CIFS home directory in the NetApp storage cluster. • Use Microsoft DFS to manage CIFS shares. NetApp supports client DFS to locate directories and files. • Use the NetApp home directory share feature to minimize the number of shares on the system. • Use SMB3 for home directories.

User Profile Data

The second type of user data is the user profile (personal data). This data allows the user to have a customized desktop environment when using a non-persistent virtual desktop. User profiles are typically stored in C:\Users on a Microsoft Windows physical machine and can be redirected to a CIFS share for use in a non-persistent, virtual desktop environment.

Many profile management solutions on the market simplify management, improve reliability, and reduce network requirements when compared with standard Windows folder redirection. A profile management solution speeds the migration process from a physical desktop or laptop by first virtualizing the profile and then virtualizing the desktop. This improves login times compared with using folder redirection alone and centralizes the end-user profile for better backup and recovery and disaster recovery of data.

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtual or physical Windows OS 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 can be customized 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 settings
- Internet Explorer favorites and homepage
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, these settings are stored within the user profile if folder redirection is not used.

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 for XenDesktop deployments.

Best Practices

For a faster login:

- NetApp recommends All Flash FAS (Solid State Drives) for profiles and PVS vDisk.
- NetApp recommends the use of Citrix User Profile Manager (UPM) software to eliminate unnecessary file copying during login and to allow users to personalize their desktops.
- NetApp recommends utilizing Folder redirection in Microsoft GPO's to eliminate an enlarged profile state, which slows down login time.
- NetApp recommends using SMB3 shares on the NetApp storage with NetApp FlexGroup volumes, which eliminates the CIFS I/O being processed by a Windows server and then by the storage array. The data is directly written to the storage array over multiple volumes of multiple clustered storage nodes.

Create Storage Volumes for PVS vDisks

Citrix recommends placing the PVS vDisks on a Microsoft SMB3 to centralize all Golden Master VDI templates images for all PVS servers in the cluster. This prevents the need to copy Golden Masters to each PVS server and potentially avoiding the issue of outdated or mixed version Golden Master templates. NetApp storage has the capability of placing the PVS vDisks on a SMB3 share residing on the NetApp storage for centralization, enterprise backups and the functional benefit of using enterprise storage. In this reference architecture, we placed the PVS vDisks on a NetApp SMB3 share on a FlexGroup volume. We successfully tested failover of the storage nodes while the PVS server was running and the PVS vDisks resided on the NetApp A300 storage array. FlexGroups do not support the Microsoft Continuous Availability (CA) feature of SMB3 (persistent handles). Therefore, there is no longer a need to enable NetApp's CA shares feature on the PVS vDisks SMB3 share when using a FlexGroup volume. This was a previous recommendation and is no longer required from ONTAP 9.3 and beyond.

NetApp OnCommand System Manager can be used to set up CIFS volumes, shares and LIFs for PVS vDisks. Although LIFs can be created and managed through the command line, this section focuses on the NetApp OnCommand System Manager GUI. Note that System Manager 3.0 or later is required to perform these steps.

CIFS Configuration

In clustered Data ONTAP 9.3, you can use Microsoft Management Console (MMC) to create shares or NetApp System Manager. In addition, you can use the NetApp System Manager tool to configure the CIFS server in a NetApp SVM and to configure the CIFS shares as well.

In this reference architecture, we used NetApp's System Manager to create and configure CIFS for the VDI environment, including the CIFS shares. Below, we show you how to create and configure CIFS shares with the NetApp System Manger GUI tool. To configure CIFS, complete the following steps:

1. To configure CIFS, sign into the System Manager Tool and go to the SVM menu.

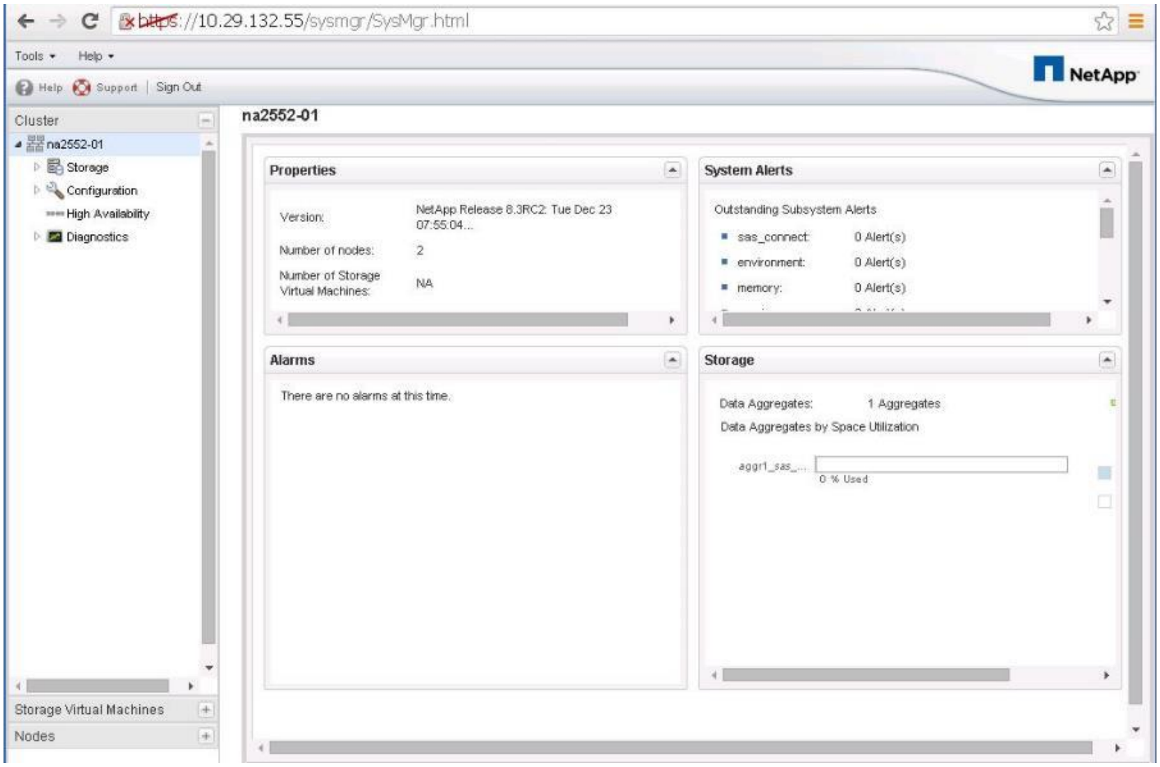

na2552-01

Properties

Version: NetApp Release 8.3RC2: Tue Dec 23 07:55:04...

Number of nodes: 2

Number of Storage Virtual Machines: NA

System Alerts

Outstanding Subsystem Alerts

- sas_connect: 0 Alert(s)
- environment: 0 Alert(s)
- memory: 0 Alert(s)

Alarms

There are no alarms at this time.

Storage

Data Aggregates: 1 Aggregates

Data Aggregates by Space Utilization

aggr1_sas_... 0 % Used

2. Click the SVM menu and then click the SVM that will contain the CIFS volumes and thus require the CIFS configuration.

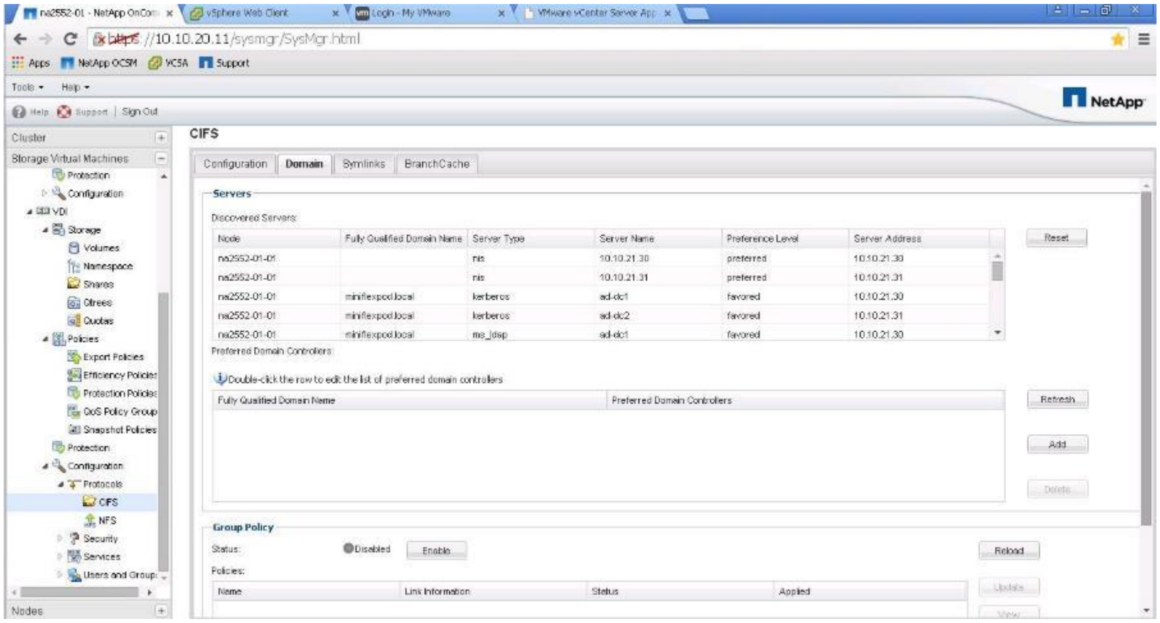

CIFS

Servers

Discovered Servers:

Node	Fully Qualified Domain Name	Server Type	Server Name	Preference Level	Server Address
na2552-01-01		nis	10.10.21.30	preferred	10.10.21.30
na2552-01-01		nis	10.10.21.31	preferred	10.10.21.31
na2552-01-01	miniflexpod.local	kerberos	ad-dc1	favored	10.10.21.30
na2552-01-01	miniflexpod.local	kerberos	ad-dc2	favored	10.10.21.31
na2552-01-01	miniflexpod.local	ms_ldap	ad-dc1	favored	10.10.21.30

Preferred Domain Controllers:

Double-click the row to edit the list of preferred domain controllers

Fully Qualified Domain Name	Preferred Domain Controllers

Group Policy

Status: Disabled

3. In the left pane, select Configuration > Protocols > CIFS.

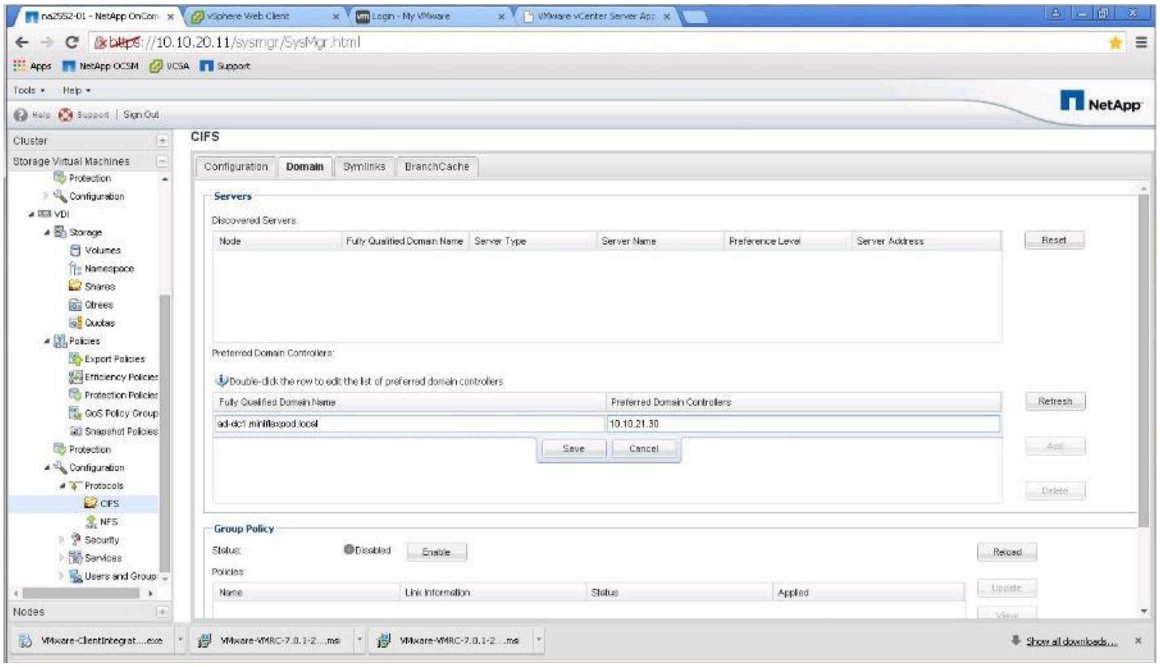

Discovered Servers:

Node	Fully Qualified Domain Name	Server Type	Server Name	Preference Level	Server Address

Preferred Domain Controllers:

Double-click the row to edit the list of preferred domain controllers

Fully Qualified Domain Name	Preferred Domain Controllers
ad-dc1.miniflexpod.local	10.10.21.30

Status: Disabled

Policies:

Name	Link Information	Status	Applied

4. In section “name of section”, we added CIFS licenses and enabled the CIFS service. To configure the Preferred Domain Controllers, click the line in the bottom window. Add the preferred DCs IP address and the FQDN and click save. Repeat this step for each DC that is local to your site and that you want to be on your preferred list.

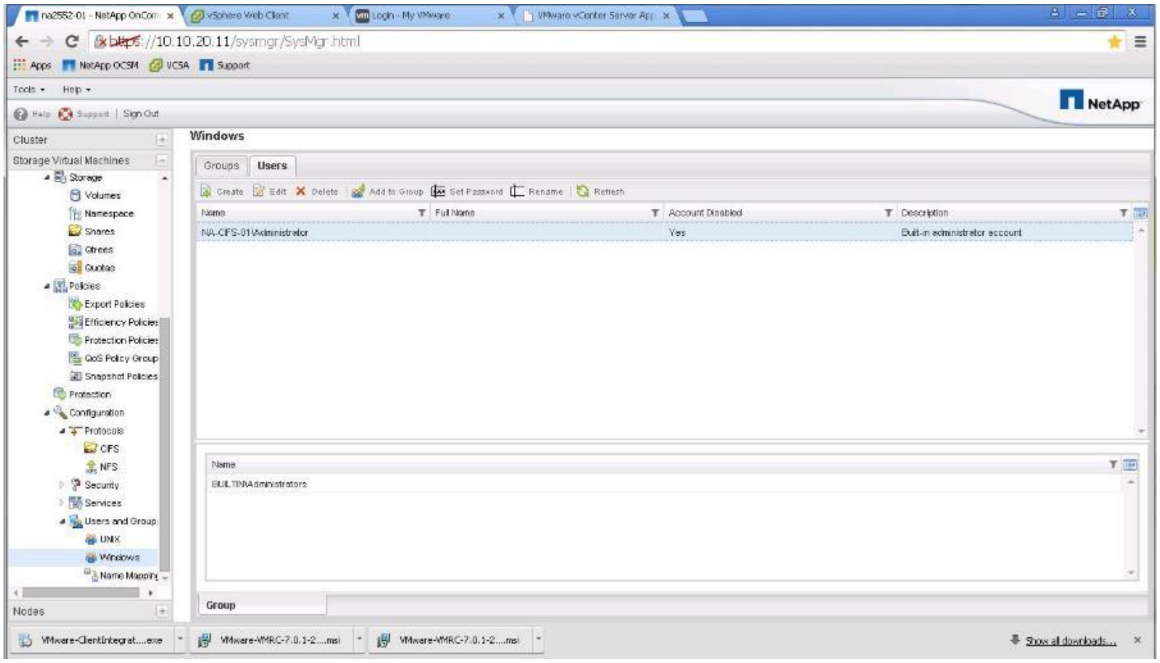

Name	Full Name	Account Disabled	Description
NA-CIFS-01\Administrator		Yes	Built-in administrator account

Name
BUILTIN\Administrators

5. Enable the built-in administrator account by selecting Users and Groups in the Configuration menu. Then click Windows. In the right pane, select the local administrator account and click Edit.

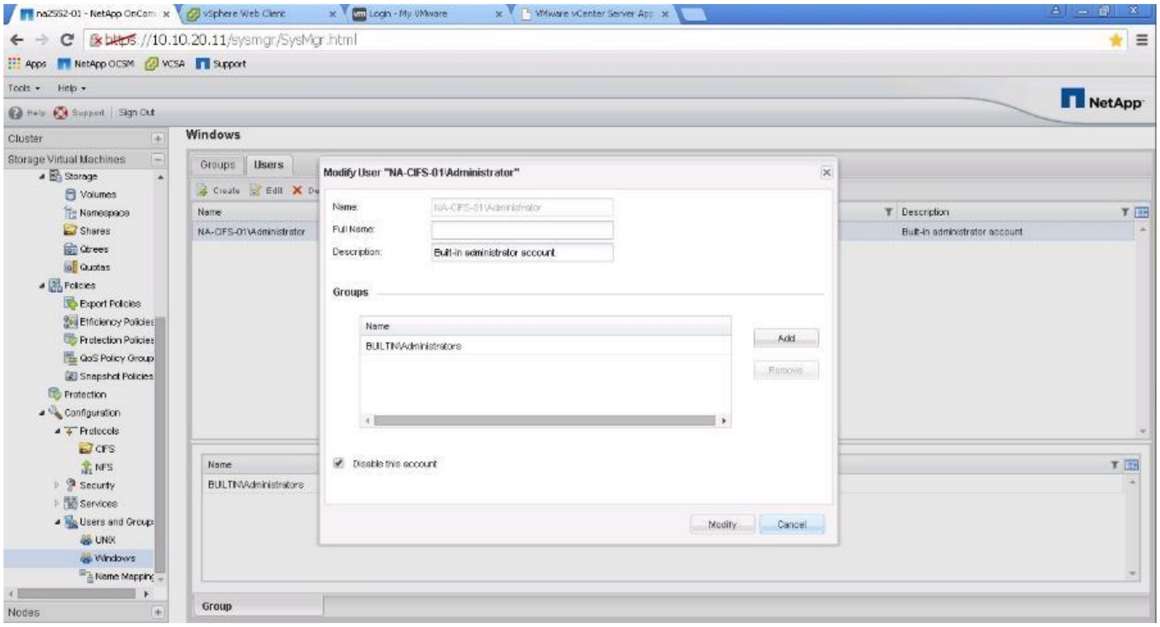

Modify User "NA-CIFS-01\Administrator"

Name: NA-CIFS-01\Administrator

Full Name:

Description: Built-in administrator account

Groups

Name
BUILTIN\Administrators

☑ Disable this account

6. Deselect Disable This Account and click Modify.

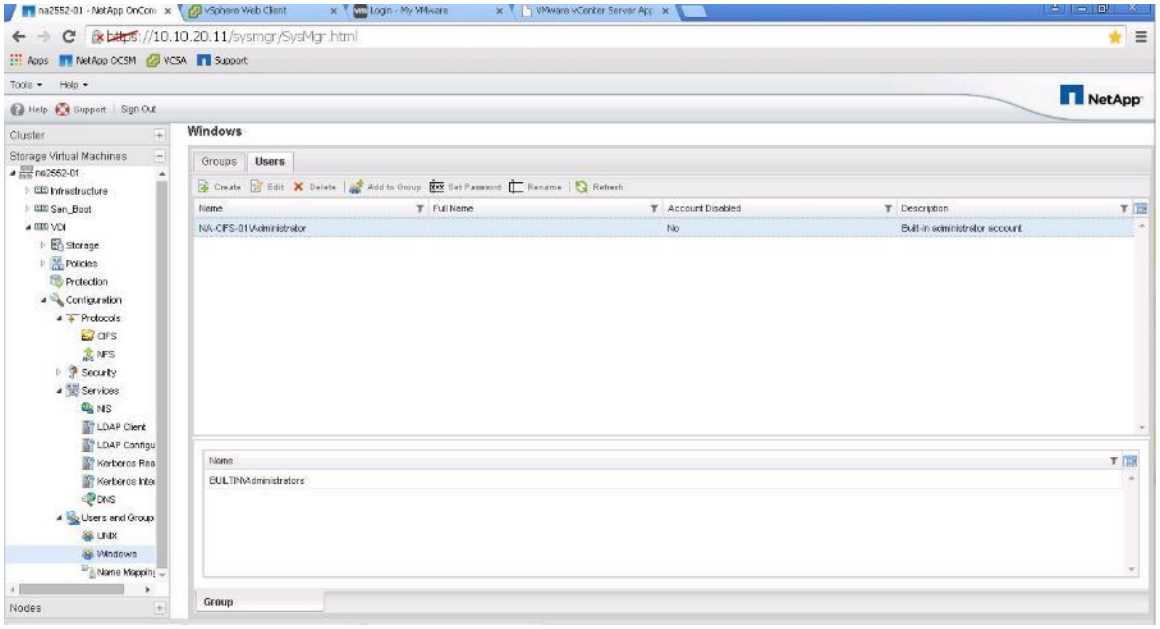

Name	Full Name	Account Disabled	Description
NA-CIFS-01\Administrator		No	Built-in administrator account

Name
BUILTIN\Administrators

The Account Disabled column should read No.

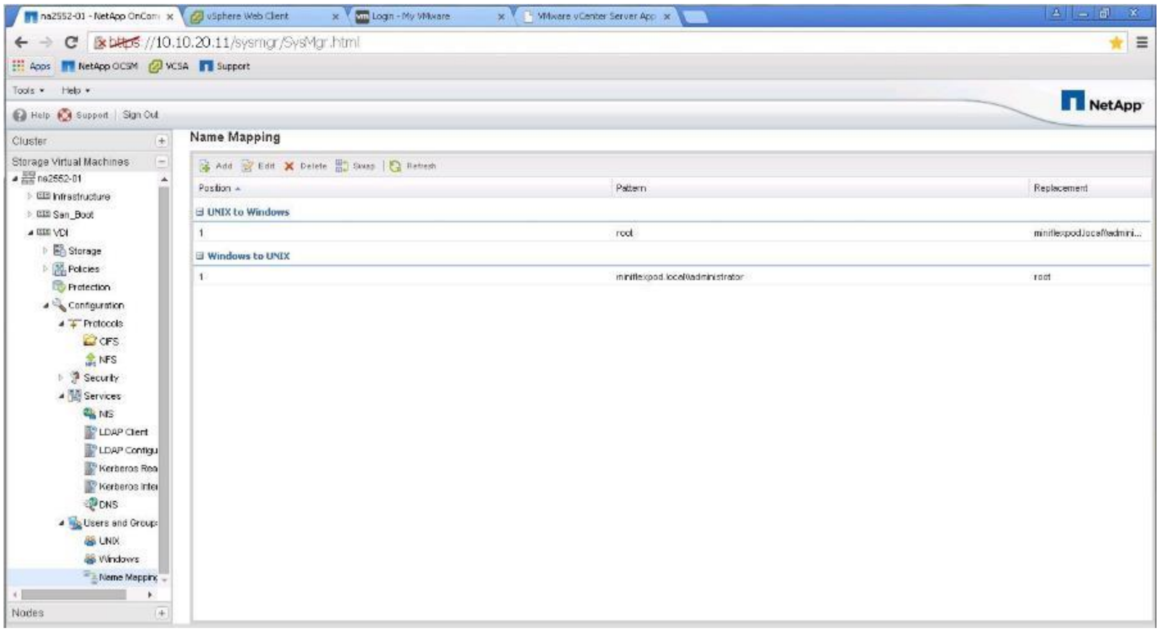

Name Mapping

Position	Pattern	Replacement
UNIX to Windows		
1	root	miniflexpod.local\admini...
Windows to UNIX		
1	miniflexpod.local\administrator	root

7. To configure Windows-to-Unix and Unix-to-Windows name mapping, select Name Mapping within the Users and Groups menu.

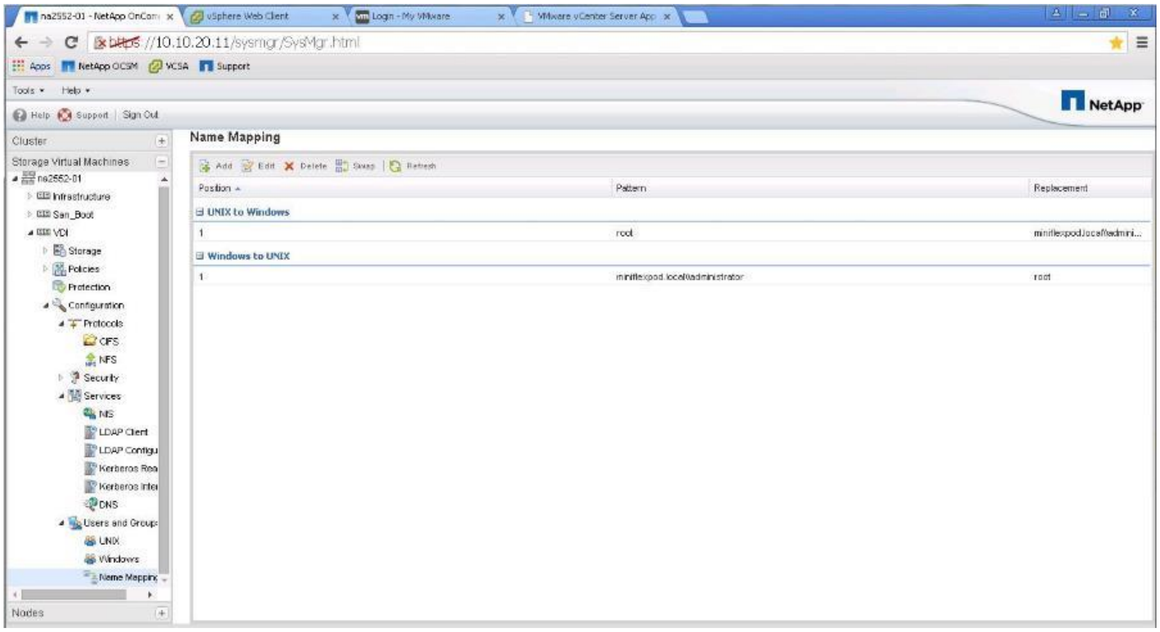

Name Mapping

Position	Pattern	Replacement
UNIX to Windows		
1	root	miniflexpod.local\admini...
Windows to UNIX		
1	miniflexpod.local\administrator	root

8. Click Add and then add the following:

    - Unix to Windows: ID=1, Pattern=root, Replacement=Domain administrator
    - Windows to Unix: ID=1, Pattern=Domain administrator, Replacement=root

Create CIFS Shares and Qtrees

Create Qtrees

As a part of the CIFS share design, we chose to utilize NetApp Qtrees to provide quotas at a lower level within the volume. A Qtree is a folder that is created within the NetApp volume and yet is maintained at the storage level, not the hypervisor level. The hypervisor has access to the Qtree, which appears as a normal mount point within the volume. The Qtree folder provides granular quota functions within the volume. A Qtree folder must be created prior to creating the CIFS share because we will export the Qtree folder as the CIFS share.

1. To create a Qtree, sign into the System Manager tool and go to the SVM menu. Expand the SVM menu and select Storage > Qtrees.

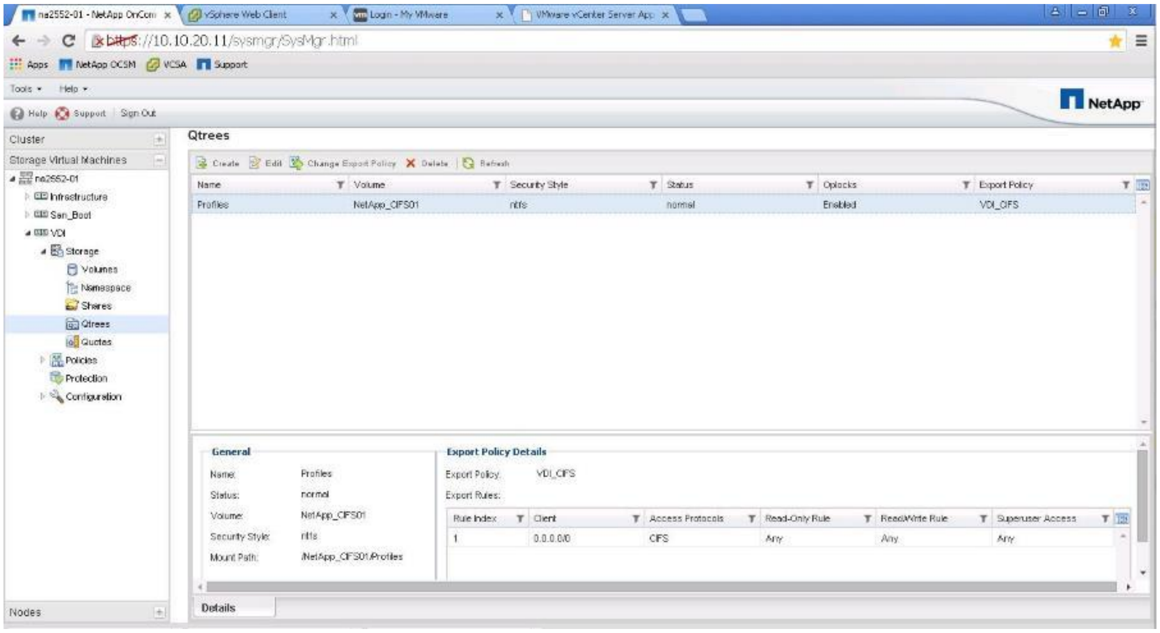

https://10.10.20.11/sysmgr/SysMgr.html

Qtrees

Name	Volume	Security Style	Status	Oplocks	Export Policy
Profiles	NetApp_CIFS01	ntfs	normal	Enabled	VDI_CIFS

General

Name: Profiles
Status: normal
Volume: NetApp_CIFS01
Security Style: ntfs
Mount Path: /NetApp_CIFS01/Profiles

Export Policy Details

Export Policy: VDI_CIFS
Export Rules:

Rule Index	Client	Access Protocols	Read-Only Rule	Read/Write Rule	Superuser Access
1	0.0.0.0/0	CIFS	Any	Any	Any

2. In the right pane, click Create to create a Qtree.

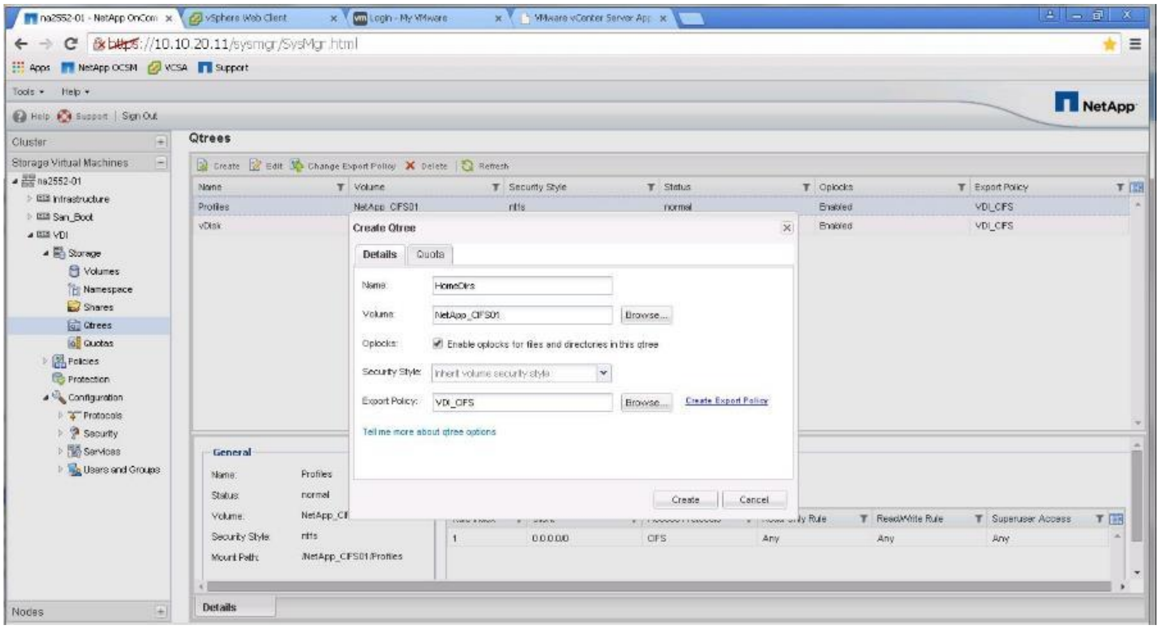

Name	Volume	Security Style	Status	Oplocks	Export Policy
Profiles	NetApp_CIFS01	ntfs	normal	Enabled	VDI_CIFS
vDisk				Enabled	VDI_CIFS

Create Qtree

Name: HomeDirs
Volume: NetApp_CIFS01
Oplocks: ☑ Enable oplocks for files and directories in this qtree
Security Style: Inherit volume security style
Export Policy: VDI_CIFS

3. Enter the Qtree folder name, chose the storage volume, select Enable Opslocks for Files and Directories in This Qtree, and enter the export policy. You can create the export policy prior to this step or by clicking Create Export Policy, then click the Quota tab.

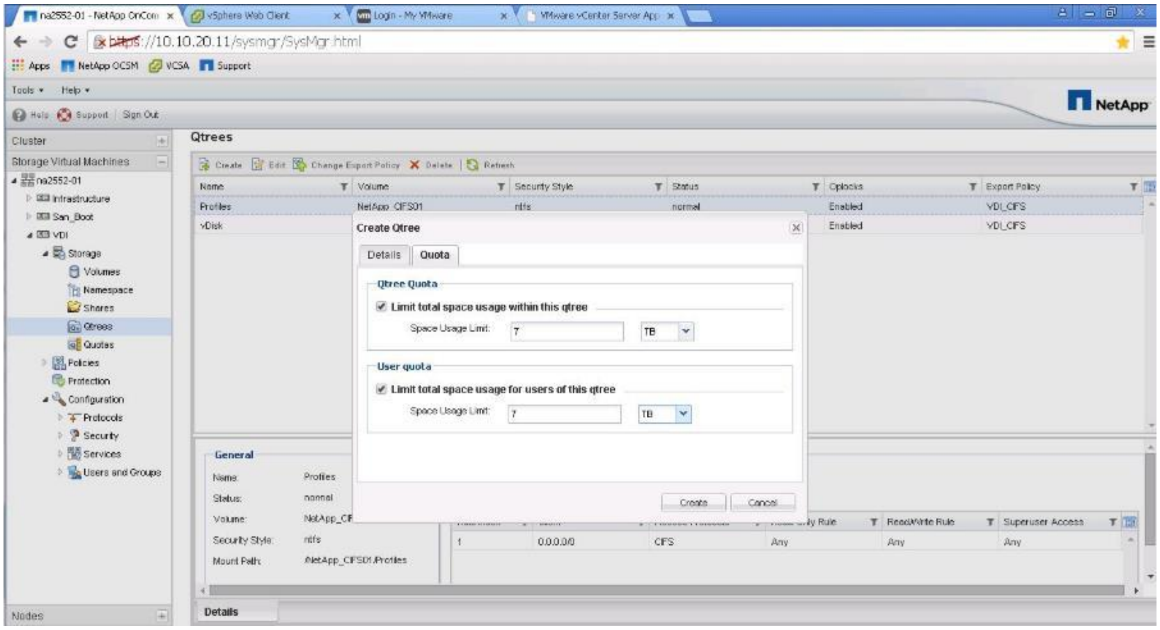

Name	Volume	Security Style	Status	Oplocks	Export Policy
Profiles	NetApp_CIFS01	ntfs	normal	Enabled	VDI_CIFS
vDisk				Enabled	VDI_CIFS

Qtree Quota

☑ Limit total space usage within this qtree

Space Usage Limit: 7 TB

User quota

☑ Limit total space usage for users of this qtree

Space Usage Limit: 7 TB

General

Name: Profiles

Status: normal

Volume: NetApp_CIF

Security Style: ntfs

Mount Path: /NetApp_CIFS01/Profiles

				...y Rule	Read/Write Rule	Superuser Access
1	0.0.0.0/0	CIFS		Any	Any	Any

4. Select Limit Total Space Usage Within This Qtree and enter the space usage limit in TB or GB. Then select the Limit Total Space Usage for Users of This Qtree and enter the space usage limit in TB or GB. Click Create.

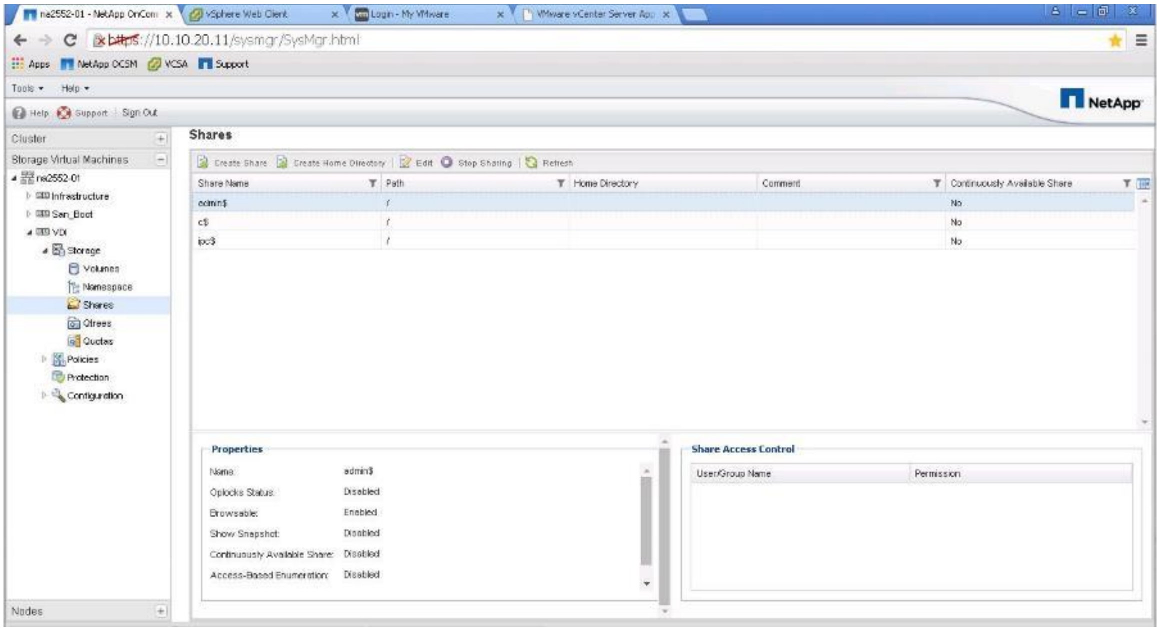

Share Name	Path	Home Directory	Comment	Continuously Available Share
admin$	/			No
c$	/			No
ipc$	/			No

Properties

Name: admin$

Oplocks Status: Disabled

Browsable: Enabled

Show Snapshot: Disabled

Continuously Available Share: Disabled

Access-Based Enumeration: Disabled

Share Access Control

User/Group Name	Permission

Create CIFS Shares

There are several tools that can create CIFS shares supported on NetApp storage. Some of the tools that can used to create CIFS shares on NetApp storage are:

- Microsoft Management Console as of cDOT 8.3
- The NetApp clustered Data ONTAP CLI
- NetApp System Manager

For this reference architecture, NetApp System Manager is used to take advantage of the NetApp User Home Directory Shares feature.

To create CIFS shares, complete the following steps:

1. Within System Manager, select SVM menu, expand the SVM, and select Storage > Shares in the left pane.

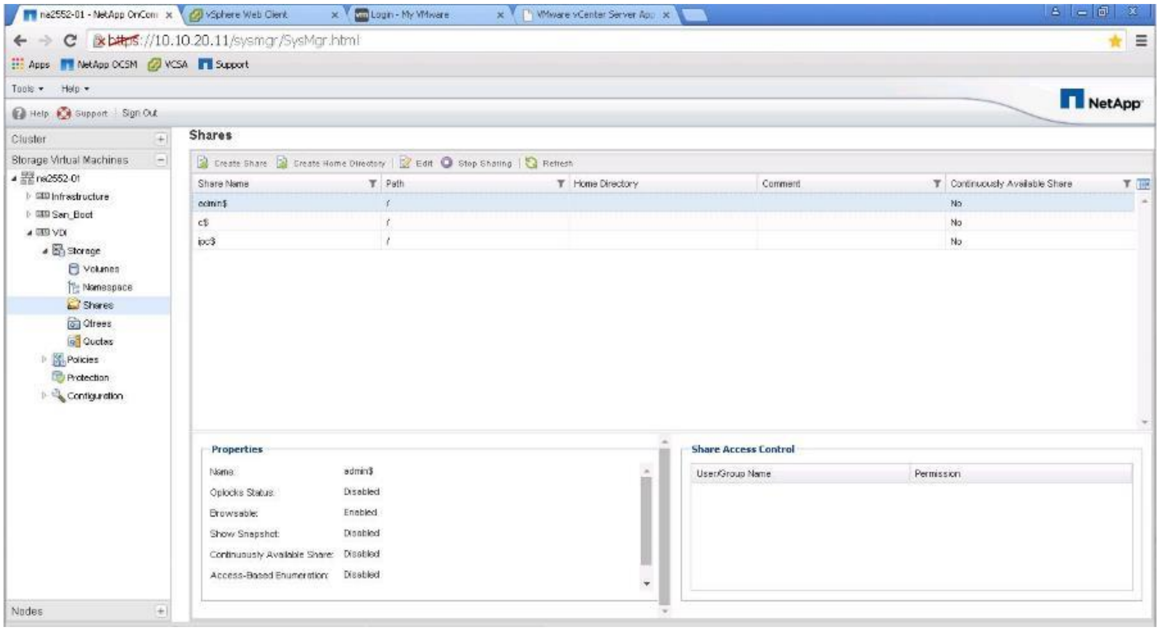

Shares

Share Name	Path	Home Directory	Comment	Continuously Available Share
admin$	/			No
c$	/			No
ipc$	/			No

Properties

Name:	admin$
Oplocks Status:	Disabled
Browsable:	Enabled
Show Snapshot:	Disabled
Continuously Available Share:	Disabled
Access-Based Enumeration:	Disabled

Share Access Control

User/Group Name	Permission

2. Click Create to create the CIFS share.

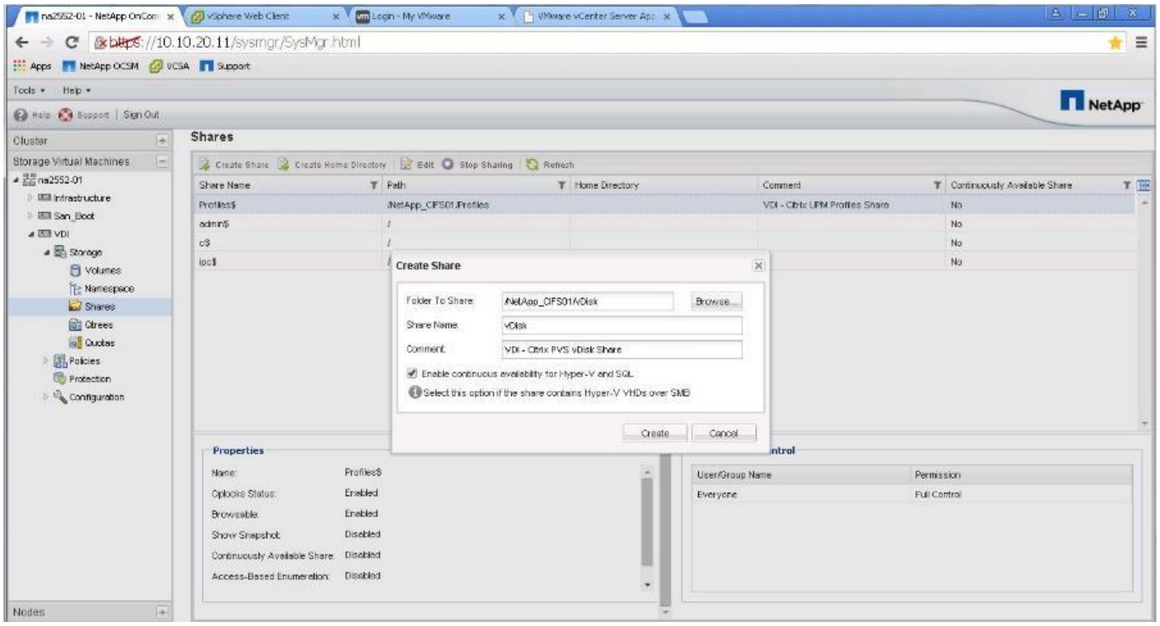

Shares

Share Name	Path	Home Directory	Comment	Continuously Available Share
Profiles$	/NetApp_CIFS01/Profiles		VDI - Citrix UPM Profiles Share	No
admin$	/			No
c$	/			No
ipc$	/			No

Create Share

Folder To Share: /NetApp_CIFS01/vDisk
Share Name: vDisk
Comment: VDI - Citrix PVS vDisk Share
☑ Enable continuous availability for Hyper-V and SQL
Select this option if the share contains Hyper-V VHDs over SMB

Properties

Name:	Profiles$
Oplocks Status:	Enabled
Browsable:	Enabled
Show Snapshot:	Disabled
Continuously Available Share:	Disabled
Access-Based Enumeration:	Disabled

User/Group Name	Permission
Everyone	Full Control

3. Enter the folder to share (the Qtree path). The CIFS share name is the advertised SMB share name mapped by the VDI clients. Enter a comment and, if needed, select Enable Continuous Availability for Hyper-V and SQL. Click Create.

Selecting Enable Continuous Availability for Hyper-V and SQL enables Microsoft Persistent Handles support on the NetApp SMB3 CIFS share. This feature is only available with SMB3 clients (Windows 2012 Servers) that map the NetApp CIFS share. Therefore, this feature is not available for Windows 2008 or 2003 servers.

Since it is Citrix PVS best practice to install the PVS vDisk on a CIFS share to eliminate the need for multiple copies of the Golden Templates, it is NetApp best practice to activate continuous availability (Persistent Handles) on the PVS vDisk share. If using FlexVol volumes for the PVS vDisks, not selecting the Continuous Availability (CA) option for the PVS vDisk share on the NetApp storage may result in a loss of PVS service if a storage controller failover event occurs (one storage node failing over its resources to another storage controller node). With FlexGroups, this is not the case.

In this reference architecture, we used FlexGroups for the CIFS requirements and FlexGroups do not support the Continuous Availability (CA) option. We conducted several successful tests of storage node failover while the PVS servers were running and did not experience any failures. Therefore, CA option is not needed for PVS servers if the PVS servers vDisks are placed on FlexGroup volumes.

Best Practices
• Use FlexGroup volumes for VDI CIFS requirements including the PVS vDisks. • If using FlexGroup volumes for the PVS vDisks, the Continuous Availability is not supported and not needed.

NetApp VDI Write-Cache Volume Considerations

Deduplication on Write Cache Volumes

Previously, it was NetApp best practice to disable deduplication on volumes that contain write-cache disks. With the advent of NetApp All Flash FAS (flash media) and the Citrix Ram Cache Plus Overflow feature starting in XenDesktop 7.15 and greater, NetApp recommends enabling deduplication on write-cache volumes. The two primary reasons for the change are the need to reduce the amount of stale data in the write cache disks and the need for capacity reduction.

The Citrix Ram Cache Plus Overflow feature reduces the majority of IOPS from centrally shared storage but still requires the full capacity of the write-cache disks. This creates a requirement for high capacity, low-IOPS on the write-cache disk volumes. This situation takes excellent advantage of the NetApp storage inline and post deduplication features. Storage deduplication is very beneficial in a low-IOPS, high-capacity environment.

Write-cache disks can build up stale data. The write-cache disk cache data is cleared when a VDI desktop is rebooted or deleted. However, in many VDI environments, customers have persistent VDI desktops or VDI desktops that do not get rebooted regularly. Therefore, stale data builds up in the write cache disks. Deduplication reduces the amount of stale data that resides on these disks.

Since capacity is an important requirement with an all flash disks environment, in this reference architecture, we enabled deduplication on the write-cache volumes to reduce the amount of capacity required. In addition, write-cache disks may vary in size from user to user. Therefore, it is uncertain how much capacity is needed for each user’s write-cache disk.

Thin-Provision the Write Cache Volumes

Another option for conserving capacity is to utilize NetApp volume thin provisioning. Thin provisioning on the write-cache volumes takes the guess work out of allocating capacity for each user’s write-cache disk file and prevents over provisioned in the environment. It is best practice to enable storage thin provisioning on the write-cache volumes.

Hypervisor Considerations with Write Cache Volumes

When creating the VMware vmdk disks, NetApp recommends provisioning the VMware vmdk disks as thick-eager zero in a Citrix VDI environment. The reasoning behind this option is to remove the extra vmdk formatting write I/O during production time. With thick-eager zero disks, the vmdk files are pre-formatted at time of creation with zeros. NetApp storage inline deduplication feature reduces those inline zeros to a couple of bytes. The VMware vmdk will display and 100Gb drive but it will only represent a couple of bytes on the NetApp storage when configured in this manner.

With thin-provisioning vmdk disks, the vmdk disks are formatted when a write I/O is requested during production; therefore, adding additional IO during production usage time to the storage. VDI environments are 90 percent writes and 10 percent reads so creating the VMware vmdk disks as thick-eager zero will reduce the amount of overall write I/O.

On the other hand, with NFS volumes, creating thick-eager zero vmdk disks does not allow a VMware vmdk to realize the space savings with storage deduplication. The NetApp storage realizes the space savings but the VMware vmdk will not see the space savings derived from the storage deduplication feature. If seeing the vmdk capacity savings within VMware is a higher priority than performance, you may want to create the VMware vmdk disks as thin-provisioned.

Best Practices
• Enable NetApp storage deduplication on the write-cache volumes • Enable thin provisioning on the write-cache volumes • Create VMware vmdk disks as Thick-eager zero disks for Citrix XenDesktop Environments

MDS 9100 Series Configuration

In this solution, we utilized the Cisco MDS 9148S Switches for Fiber Channel Switching. For racking, cable and initial setup of the MDS switches, please refer to the Quick Start Guide: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/hw/9148/quick/guide/MDS_9148_QSG.pdf

For this solution, the Fabric Interconnect A ports 1-4 connected to the MDS switch A ports 43-46, and MDS Switch A ports 37, 38 connected to Netapp A300 Controller A and Controller B ports 0g. Similarly, the Fabric Interconnect B ports 1 - 4 connected to the MDS switch B ports 43-46, and MDS Switch B ports 37, 38 connected to Netapp A300 Controller A and Controller B ports 0h. All ports carry 16 Gb/s FC Traffic.

For this design, two separate fabrics were created, each with their own unique VSAN. Fabric A side configured for VSAN400 while Fabric B side for VSAN401.

Figure 29 VSAN 400 Configured for Fabric A

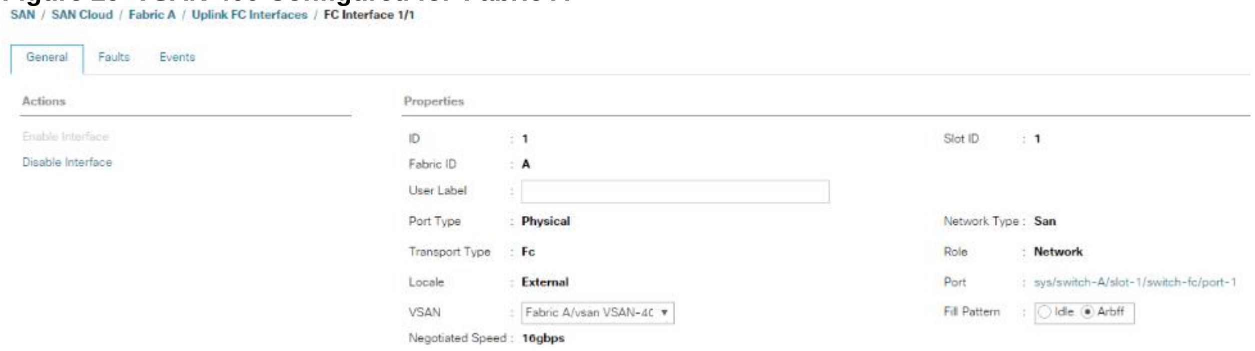

SAN / SAN Cloud / Fabric A / Uplink FC Interfaces / FC Interface 1/1

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID	: 1	Slot ID	: 1
Fabric ID	: A		
User Label	:		
Port Type	: Physical	Network Type	: San
Transport Type	: Fc	Role	: Network
Locale	: External	Port	: sys/switch-A/slot-1/switch-fc/port-1
VSAN	: Fabric A/vsan VSAN-4C ▾	Fill Pattern	: ○ Idle ● Arbff
Negotiated Speed	: 16gbps		

Figure 30 VSAN 401 Configured for Fabric B

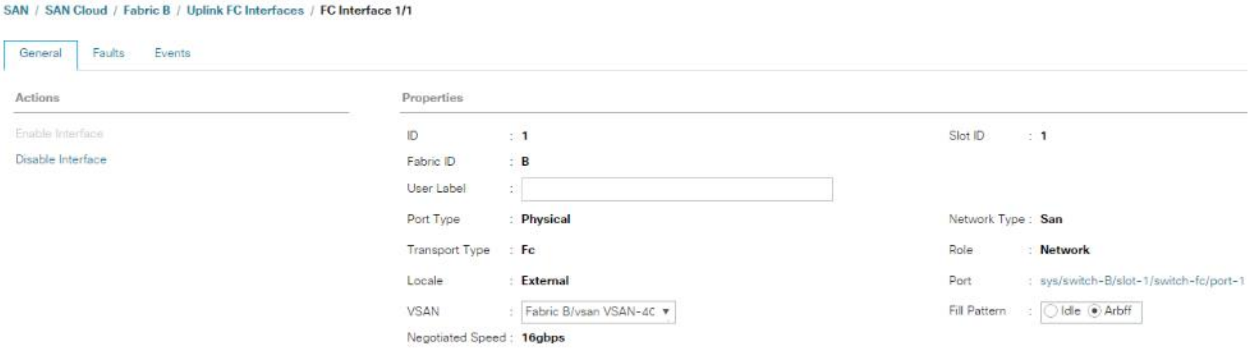

SAN / SAN Cloud / Fabric B / Uplink FC Interfaces / FC Interface 1/1

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID	: 1	Slot ID	: 1
Fabric ID	: B		
User Label	:		
Port Type	: Physical	Network Type	: San
Transport Type	: Fc	Role	: Network
Locale	: External	Port	: sys/switch-B/slot-1/switch-fc/port-1
VSAN	: Fabric B/vsan VSAN-4C ▾	Fill Pattern	: ○ Idle ● Arbff
Negotiated Speed	: 16gbps		

Figure 31 Fibre Channel Cable Connectivity from Netapp AFF A300 to Cisco MDS 9148S to Cisco 6332-16UP Fabric Interconnects

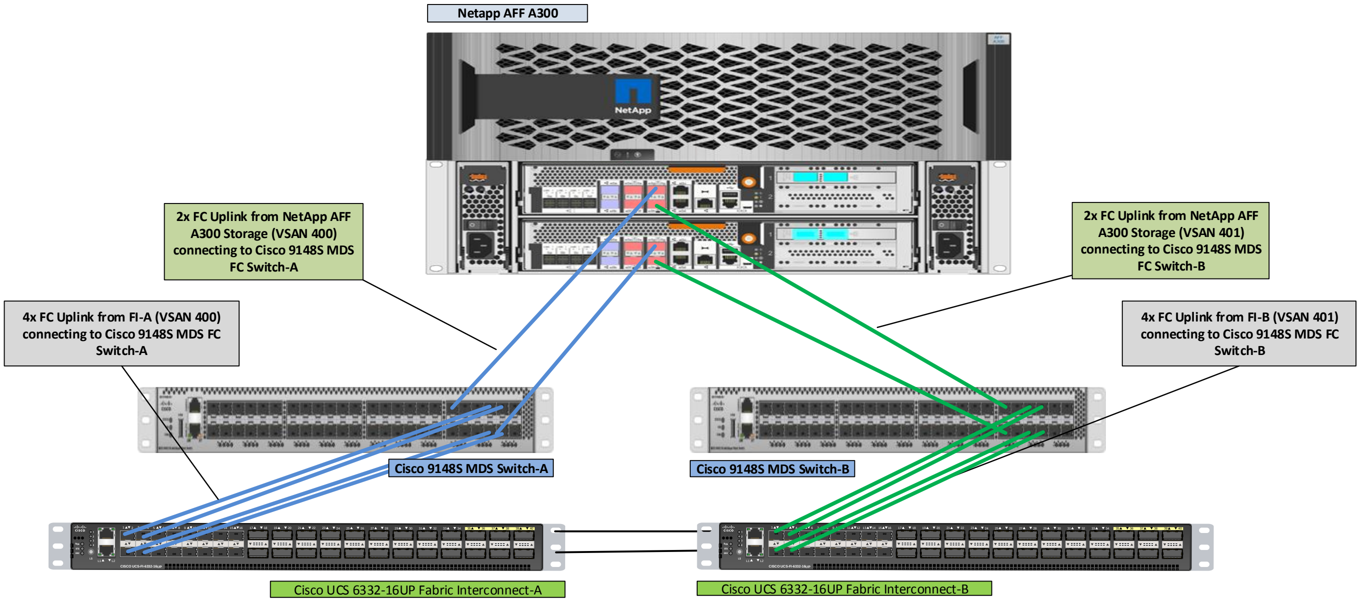
Netapp AFF A300
NetApp
2x FC Uplink from NetApp AFF A300 Storage (VSAN 400) connecting to Cisco 9148S MDS FC Switch-A
2x FC Uplink from NetApp AFF A300 Storage (VSAN 401) connecting to Cisco 9148S MDS FC Switch-B
4x FC Uplink from FI-A (VSAN 400) connecting to Cisco 9148S MDS FC Switch-A
4x FC Uplink from FI-B (VSAN 401) connecting to Cisco 9148S MDS FC Switch-B
Cisco 9148S MDS Switch-A
Cisco 9148S MDS Switch-B
Cisco UCS 6332-16UP Fabric Interconnect-A
Cisco UCS 6332-16UP Fabric Interconnect-B

All connections are 16Gb FC links.

Configure Feature for MDS Switch A and MDS Switch B

To set feature on MDS Switches, complete the following steps on both MDS switches:

1. Login as admin user into MDS Switch A.

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

2. Login 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

The next steps are to configure the VSANs, ports, and zones in the MDS switches. The commands listed below how to do it. The entire MDS 9148S FC switch configuration is included in Appendix A of this document.

MDS-A

```
config terminal
VSAN database
vsan 400
vsan 400 interface fc {interface 1/X}
exit

interface fc {interface 1/X} switchport trunk allowed vsan 400
switchport trunk mode off
port-license acquire
no shutdown
exit

Zoneset name AFF-A300_VDI vsan 400
Member {ESXi hostname-fc0}
exit

Zoneset activate name AFF-A300_VDI vsan 400
Zone commit vsan 400
exit

Copy running-config startup-config
```

MDS-B

```
config terminal
VSAN database
vsan 401
vsan 401 interface fc {interface 1/X}
exit

interface fc {interface 1/X} switchport trunk allowed vsan 401
switchport trunk mode off
port-license acquire
no shutdown
```

```
exit

Zoneset name AFF-A300_VDI vsan 401
Member {ESXi hostname-fc1}
exit

Zoneset activate name AFF-A300_VDI vsan 401
Zone commit vsan 401
exit

Copy running-config startup-config
```

Installing and Configuring VMware ESXi 6.5

This section provides detailed instructions for installing VMware ESXi 6.5 Update1 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6.5 Update 1

To download the Cisco Custom Image for ESXi 6.5 Update 1, complete the following steps:

1. Click the following link vmware login page.

2. Type your email or customer number and the password and then click Log in.

3. Click the following link:

https://my.vmware.com/web/vmware/details?productId=614&downloadGroup=ESXI65U1

4. Click Download Now.

5. Save it to your destination folder.

KVM Access to Hosts

To log in to the Cisco UCS environment, complete the following steps:

1. Log in to Cisco UCS Manager.

2. The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

3. Open a Web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.

4. Log in to Cisco UCS Manager by using the admin user name and password.

5. From the main menu, click the Servers tab.

6. Select Servers > Service Profiles > root > VM-Host-01.

7. Right-click VM-Host-01 and select KVM Console.

8. Repeat steps for 4-6 for all host servers.

Set Up VMware ESXi Installation

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click the Virtual Media tab.

2. Click Add Image.

3. Browse to the ESXi installer ISO image file and click Open.

4. Select the Mapped checkbox to map the newly added image.

5. Click the KVM tab to monitor the server boot.

6. Boot the server by selecting Boot Server and click OK, then click OK again.

Install ESXi

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that is displayed.

2. After the installer is finished loading, press Enter to continue with the installation.

3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

4. Select the AFF A300 boot LUN.

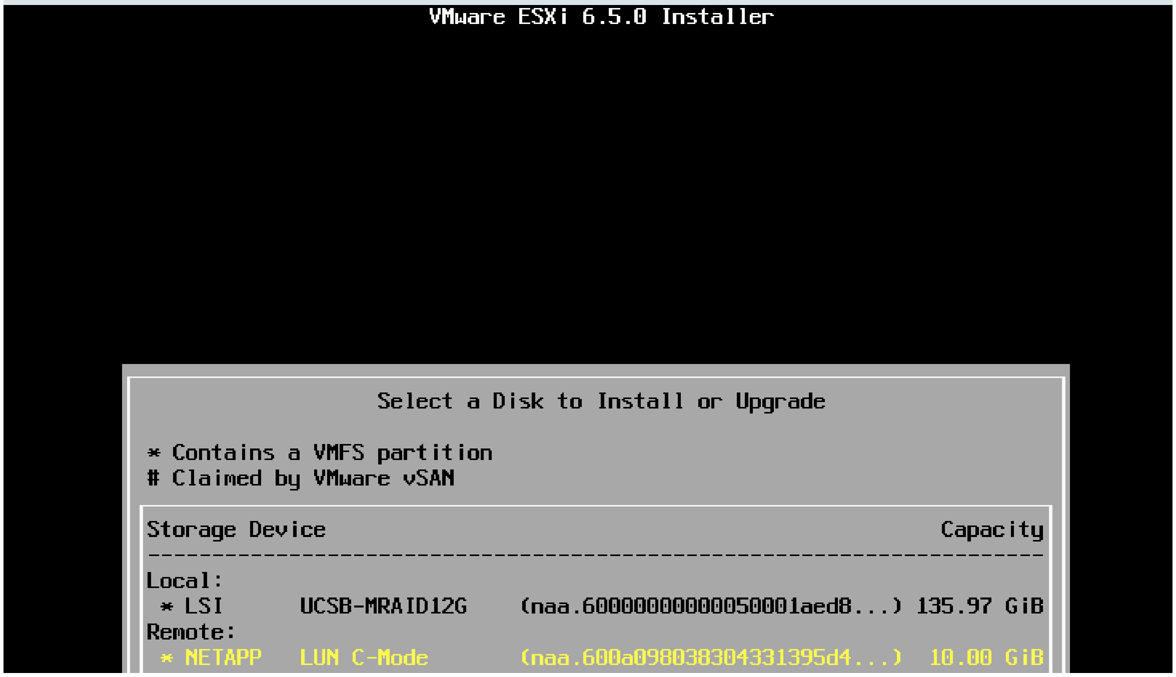

```
                         VMware ESXi 6.5.0 Installer

                    Select a Disk to Install or Upgrade

        * Contains a VMFS partition
        # Claimed by VMware vSAN

        Storage Device                                                Capacity
        ----------------------------------------------------------------------
        Local:
         * LSI      UCSB-MRAID12G    (naa.60000000000050001aed8...) 135.97 GiB
        Remote:
         * NETAPP   LUN C-Mode       (naa.600a098038304331395d4...)  10.00 GiB
```

5. (NetApp LUN C-Mode(naa.600a098038304331395d4) 10 GB that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

6. Select the appropriate keyboard layout and press Enter.

7. Enter and confirm the root password and press Enter.

8. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.

9. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

10. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, click Yes to unmap the image.

11. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host.

To configure the ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.

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.

5. Enter the VLAN in-band management ID and press Enter.

6. From the Configure Management Network menu, select IP Configuration and press Enter.

7. Select the Set Static IP Address and Network Configuration option by using the space bar.

8. Enter the IP address for managing the first ESXi host.

9. Enter the subnet mask for the first ESXi host.

10. Enter the default gateway for the first ESXi host.

11. Press Enter to accept the changes to the IP configuration.

12. Select the IPv6 Configuration option and press Enter.

13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.

14. Select the DNS Configuration option and press Enter.

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

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

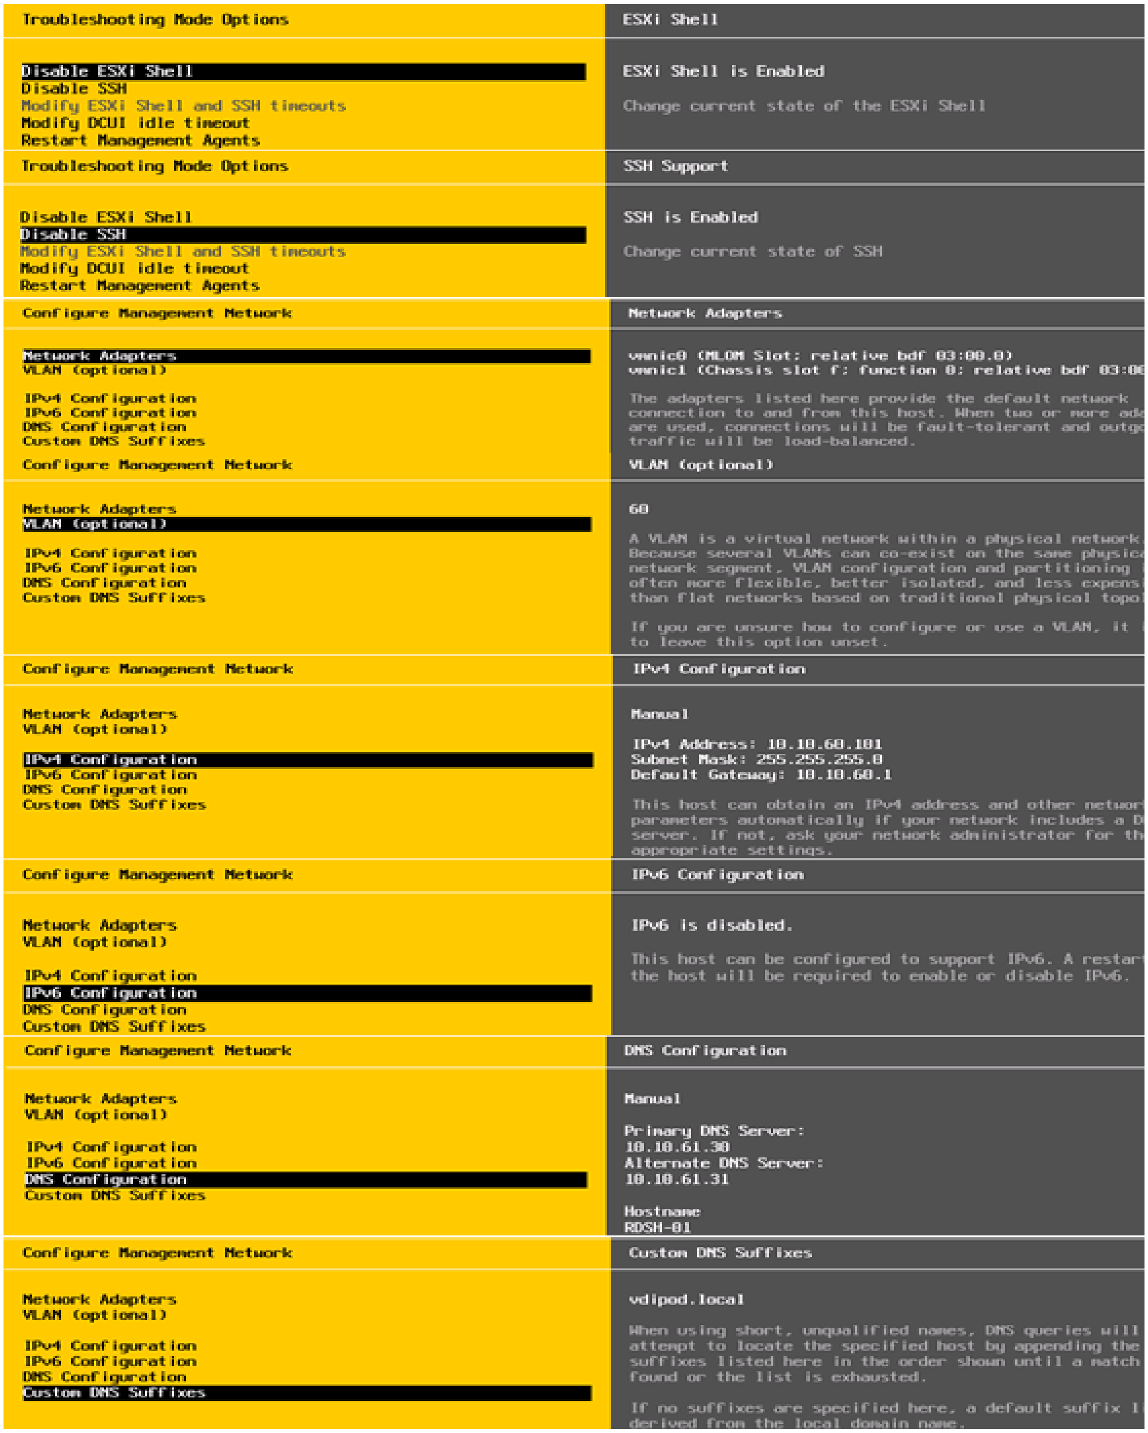

```
Troubleshooting Mode Options

Disable ESXi Shell
Disable SSH
Modify ESXi Shell and SSH timeouts
Modify DCUI idle timeout
Restart Management Agents
```

```
ESXi Shell

ESXi Shell is Enabled

Change current state of the ESXi Shell
```

```
Troubleshooting Mode Options

Disable ESXi Shell
Disable SSH
Modify ESXi Shell and SSH timeouts
Modify DCUI idle timeout
Restart Management Agents
```

```
SSH Support

SSH is Enabled

Change current state of SSH
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
Network Adapters

vmnic0 (MLOM Slot: relative bdf 03:00.0)
vmnic1 (Chassis slot f: function 0: relative bdf 03:00

The adapters listed here provide the default network
connection to and from this host. When two or more ada
are used, connections will be fault-tolerant and outgo
traffic will be load-balanced.
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
VLAN (optional)

60

A VLAN is a virtual network within a physical network.
Because several VLANs can co-exist on the same physica
network segment, VLAN configuration and partitioning i
often more flexible, better isolated, and less expensi
than flat networks based on traditional physical topol

If you are unsure how to configure or use a VLAN, it i
to leave this option unset.
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
IPv4 Configuration

Manual

IPv4 Address: 10.10.60.101
Subnet Mask: 255.255.255.0
Default Gateway: 10.10.60.1

This host can obtain an IPv4 address and other networ
parameters automatically if your network includes a D
server. If not, ask your network administrator for th
appropriate settings.
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
IPv6 Configuration

IPv6 is disabled.

This host can be configured to support IPv6. A restart
the host will be required to enable or disable IPv6.
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
DNS Configuration

Manual

Primary DNS Server:
10.10.61.30
Alternate DNS Server:
10.10.61.31

Hostname
RDSH-01
```

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
Custom DNS Suffixes

vdipod.local

When using short, unqualified names, DNS queries will
attempt to locate the specified host by appending the
suffixes listed here in the order shown until a match
found or the list is exhausted.

If no suffixes are specified here, a default suffix li
derived from the local domain name.
```

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-01 management IP address.

2. Download and install the vSphere Client.

This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 6

To download VMware vSphere CLI 6.5, complete the following steps:

1. Click the following link: https://my.vmware.com/en/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5
2. Select your OS and click Download.
3. Save it to your destination folder.
4. Run the VMware-vSphere-CLI.exe
5. Click Next.
6. Accept the terms for the license and click Next.
7. Click Next on the Destination Folder screen.
8. Click Install.
9. Click Finish.

Install VMware vSphere CLI 6.5 on the management workstation.

10. Log in to VMware ESXi Hosts by Using VMware vSphere Client.

Log in to VMware ESXi Hosts by using VMware vSphere Client

To log in to the `VM-Host-01` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-01` as the host you are trying to connect to: `<<var_vm_host_01_ip>>`.
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

Download and Install Updated Cisco VIC eNIC Drivers

The Cisco VIC drivers for VMware ESXi Hypervisor may require an update to match current Cisco Hardware and Software Interoperability Matrix.

Figure 32 Cisco UCS Hardware and Software Interoperability Matrix for vSphere 6.5 U1 and Cisco UCS B200 M5 on Cisco UCS Manager v3.2.3. Recommendation

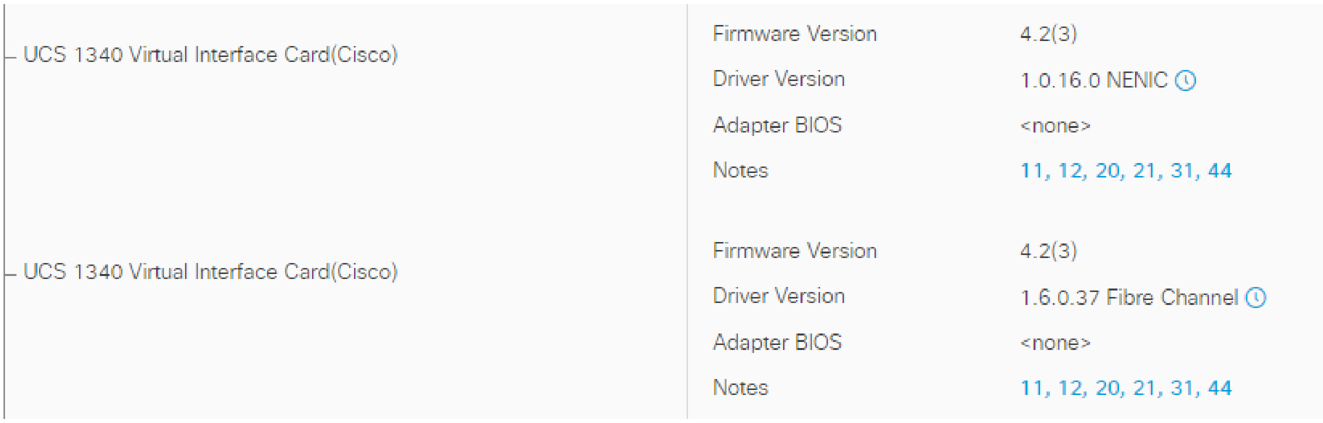

UCS 1340 Virtual Interface Card(Cisco)	Firmware Version	4.2(3)
	Driver Version	1.0.16.0 NENIC
	Adapter BIOS	<none>
	Notes	11, 12, 20, 21, 31, 44
UCS 1340 Virtual Interface Card(Cisco)	Firmware Version	4.2(3)
	Driver Version	1.6.0.37 Fibre Channel
	Adapter BIOS	<none>
	Notes	11, 12, 20, 21, 31, 44

1. Download the recommended Cisco Virtual Interface Card (VIC) eNIC and fNIC drivers:

The neNIC version is 1.0.16.0 and the fNIC version is 1.6.0.37 were used in this configuration

2. Open a Web browser on the management workstation and navigate to www.cisco.com.

3. Download the Cisco eNIC and fNIC driver bundle.

4. Open the neNIC driver bundle. This bundle includes the VMware driver bundle which will be uploaded to ESXi hosts.

5. Open the fNIC driver bundle. This bundle includes the VMware driver bundle which will be uploaded to ESXi hosts.

6. Save the location of these driver bundles for uploading to ESXi in the next section.

Go to www.cisco.com for the latest ISO images of Cisco UCS-related drivers.

Load Updated Cisco VIC neNIC and fNIC Drivers

To install VMware VIC Drivers on the ESXi host servers, complete the following steps:

1. From the vSphere Client, select the host in the inventory.

2. Click the Summary tab to view the environment summary.

3. From Resources > Storage, right-click datastore1 and select Browse Datastore.

4. Click the fourth button and select Upload File.

5. Navigate to the saved location for each downloaded VIC driver and select:

    a. VMware ESXi 6.5 NIC nenic 1.0.16.0 Driver for Cisco nenic

    b. VMware ESXi 6.5 fnic 1.6.0.37 FC Driver for Cisco

6. Click Open on each and click Yes to upload the file to datastore1.

7. Click the fourth button and select Upload File.

8. Repeat the process until the files have been uploaded to all ESXi hosts.

9. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.

10. At the command prompt, run the following commands to account for each host:

> To get the host thumbprint, type the command without the –-thumbprint option, then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/ VMW-ESX-6.5.0-
nenic-1.0.16.0-offline_bundle-7643104.zip
```

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/
fnic_driver_1.6.0.37-offline_bundle-7765239.zip
```

11. Back in the vSphere Client for each host, right click the host and select Reboot.

12. Click Yes and OK to reboot the host.

13. Log back into each host with vSphere Client.

> Verify the neNIC driver version installed by entering `vmkload_mod -s nenic` and `vmkload_mod -s fNIC` at the command prompt.

Figure 33 Verify the neNIC Driver Version

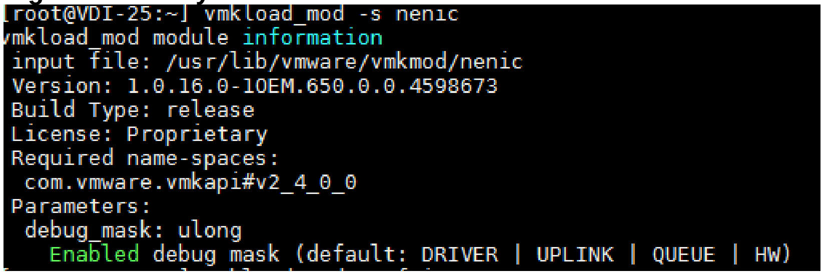

```
[root@VDI-25:~] vmkload_mod -s nenic
vmkload_mod module information
 input file: /usr/lib/vmware/vmkmod/nenic
 Version: 1.0.16.0-1OEM.650.0.0.4598673
 Build Type: release
 License: Proprietary
 Required name-spaces:
  com.vmware.vmkapi#v2_4_0_0
 Parameters:
  debug_mask: ulong
    Enabled debug mask (default: DRIVER | UPLINK | QUEUE | HW)
```

Figure 34 Verify the fNIC Driver Version

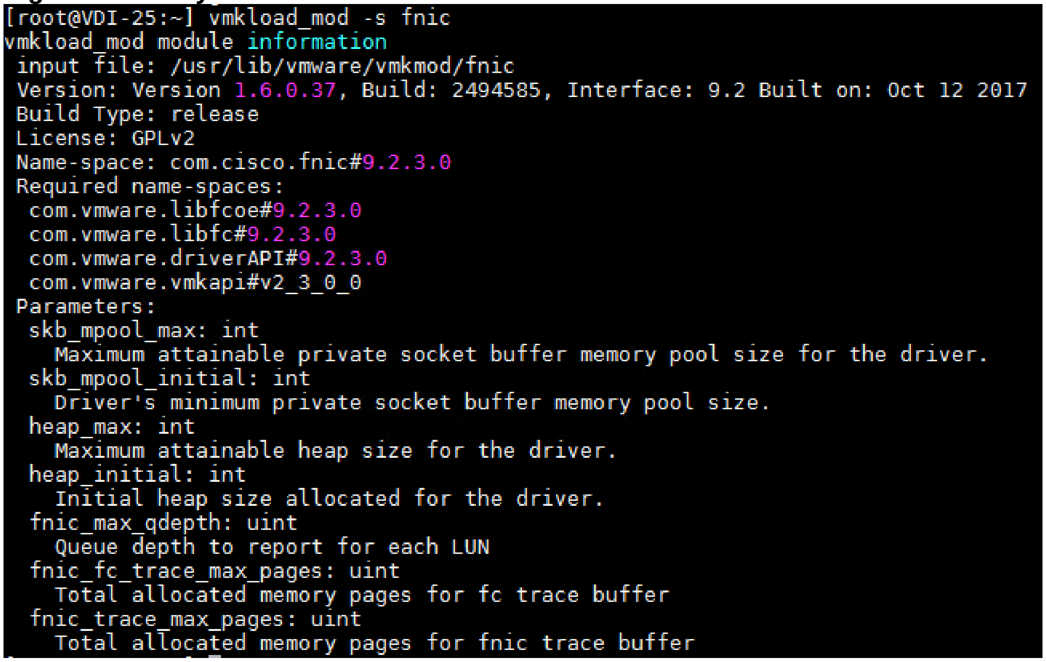

```
[root@VDI-25:~] vmkload_mod -s fnic
vmkload_mod module information
 input file: /usr/lib/vmware/vmkmod/fnic
 Version: Version 1.6.0.37, Build: 2494585, Interface: 9.2 Built on: Oct 12 2017
 Build Type: release
 License: GPLv2
 Name-space: com.cisco.fnic#9.2.3.0
 Required name-spaces:
  com.vmware.libfcoe#9.2.3.0
  com.vmware.libfc#9.2.3.0
  com.vmware.driverAPI#9.2.3.0
  com.vmware.vmkapi#v2_3_0_0
 Parameters:
  skb_mpool_max: int
    Maximum attainable private socket buffer memory pool size for the driver.
  skb_mpool_initial: int
    Driver's minimum private socket buffer memory pool size.
  heap_max: int
    Maximum attainable heap size for the driver.
  heap_initial: int
    Initial heap size allocated for the driver.
  fnic_max_qdepth: uint
    Queue depth to report for each LUN
  fnic_fc_trace_max_pages: uint
    Total allocated memory pages for fc trace buffer
  fnic_trace_max_pages: uint
    Total allocated memory pages for fnic trace buffer
```

Install and Configure VMware vCenter Appliance

Log in to the VM-Host-01 ESXi host by using the VMware vSphere Client and complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-01 as the host you are trying to connect to.

2. Enter root for the user name.

3. Enter the root password.

4. Click Login to connect.

To build the VMWare vCenter VM, complete the following steps:

1. From the vSphere 6 download page on the VMware Web site, download the vCenter ISO file for the vCenter Server appliance onto your system.

2. Mount the vSphere ISO file via Windows Explorer and navigate to the folder vcsa-ui-installer/win32 and click installer file to start the VCSA Appliance from Installer.

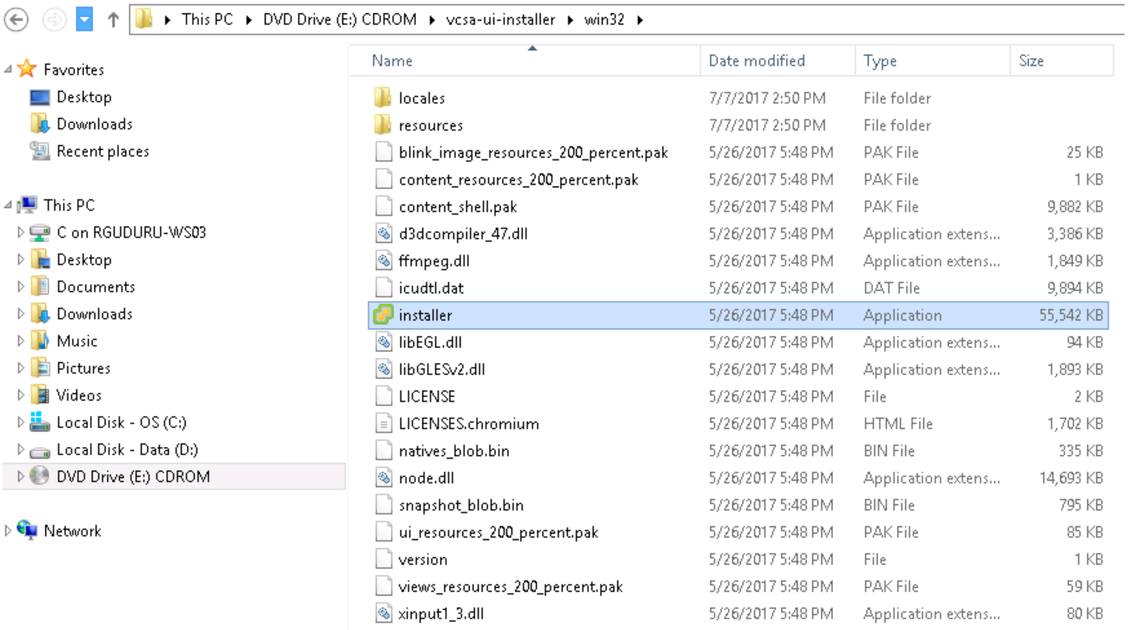

Name	Date modified	Type	Size
locales	7/7/2017 2:50 PM	File folder	
resources	7/7/2017 2:50 PM	File folder	
blink_image_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	25 KB
content_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	1 KB
content_shell.pak	5/26/2017 5:48 PM	PAK File	9,882 KB
d3dcompiler_47.dll	5/26/2017 5:48 PM	Application extens...	3,386 KB
ffmpeg.dll	5/26/2017 5:48 PM	Application extens...	1,849 KB
icudtl.dat	5/26/2017 5:48 PM	DAT File	9,894 KB
installer	5/26/2017 5:48 PM	Application	55,542 KB
libEGL.dll	5/26/2017 5:48 PM	Application extens...	94 KB
libGLESv2.dll	5/26/2017 5:48 PM	Application extens...	1,893 KB
LICENSE	5/26/2017 5:48 PM	File	2 KB
LICENSES.chromium	5/26/2017 5:48 PM	HTML File	1,702 KB
natives_blob.bin	5/26/2017 5:48 PM	BIN File	335 KB
node.dll	5/26/2017 5:48 PM	Application extens...	14,693 KB
snapshot_blob.bin	5/26/2017 5:48 PM	BIN File	795 KB
ui_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	85 KB
version	5/26/2017 5:48 PM	File	1 KB
views_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	59 KB
xinput1_3.dll	5/26/2017 5:48 PM	Application extens...	80 KB

3. Click Install.

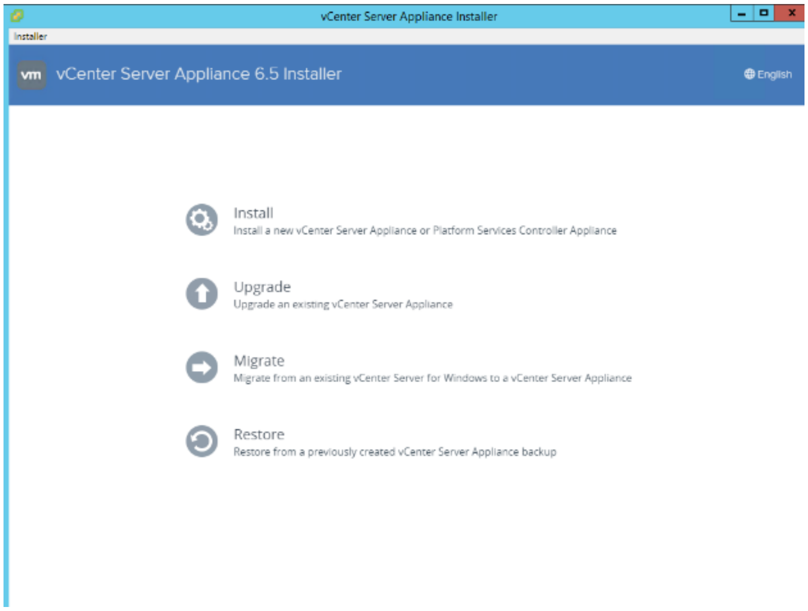

vCenter Server Appliance 6.5 Installer

Install
Install a new vCenter Server Appliance or Platform Services Controller Appliance

Upgrade
Upgrade an existing vCenter Server Appliance

Migrate
Migrate from an existing vCenter Server for Windows to a vCenter Server Appliance

Restore
Restore from a previously created vCenter Server Appliance backup

4. Click Next.

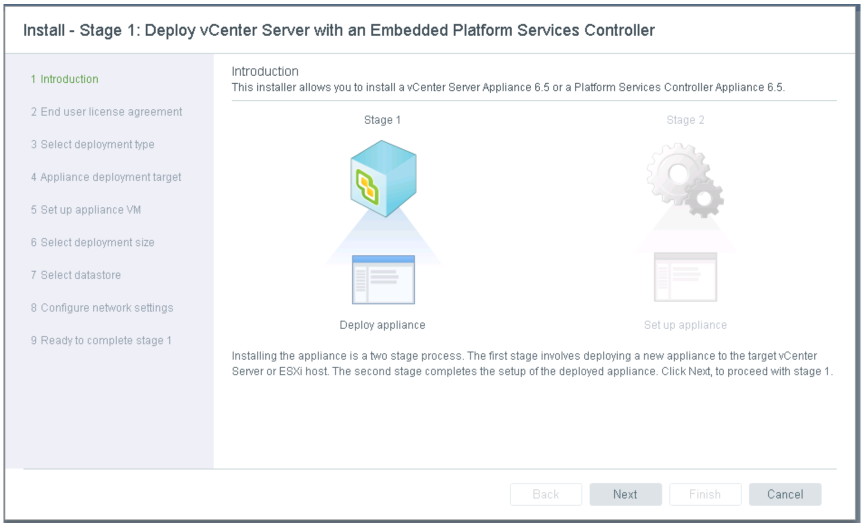

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

Introduction
This installer allows you to install a vCenter Server Appliance 6.5 or a Platform Services Controller Appliance 6.5.

Installing the appliance is a two stage process. The first stage involves deploying a new appliance to the target vCenter Server or ESXi host. The second stage completes the setup of the deployed appliance. Click Next, to proceed with stage 1.

5. Follow the onscreen prompts. Accept EULA.

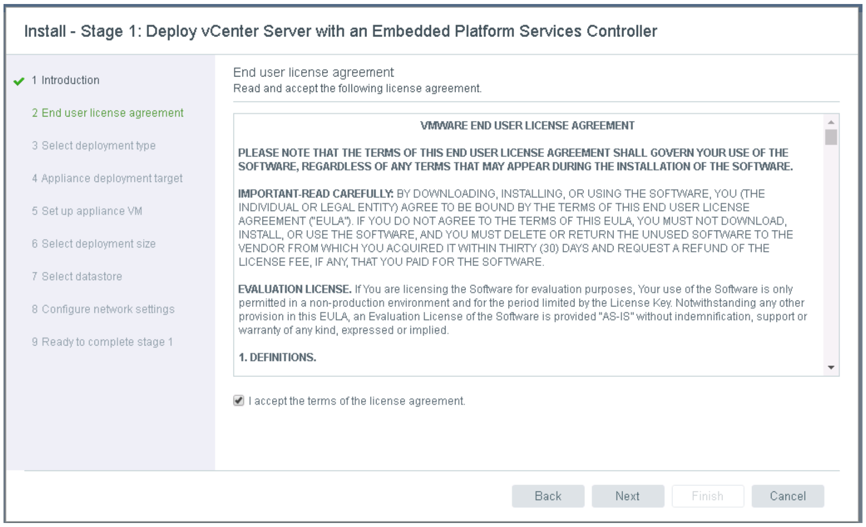

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

✓ 1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1

End user license agreement
Read and accept the following license agreement.

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 USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE VENDOR FROM WHICH YOU ACQUIRED IT WITHIN THIRTY (30) DAYS AND REQUEST A REFUND OF THE LICENSE FEE, IF ANY, THAT YOU PAID FOR THE SOFTWARE.

EVALUATION LICENSE. If You are licensing the Software for evaluation purposes, Your use of the Software is only permitted in a non-production environment and for the period limited by the License Key. Notwithstanding any other provision in this EULA, an Evaluation License of the Software is provided "AS-IS" without indemnification, support or warranty of any kind, expressed or implied.

1. DEFINITIONS.

☑ I accept the terms of the license agreement.

6. Select Install vCenter Server with and Embedded Platform Services Controller (unless your environment already has a PSC).

7. Click Next.

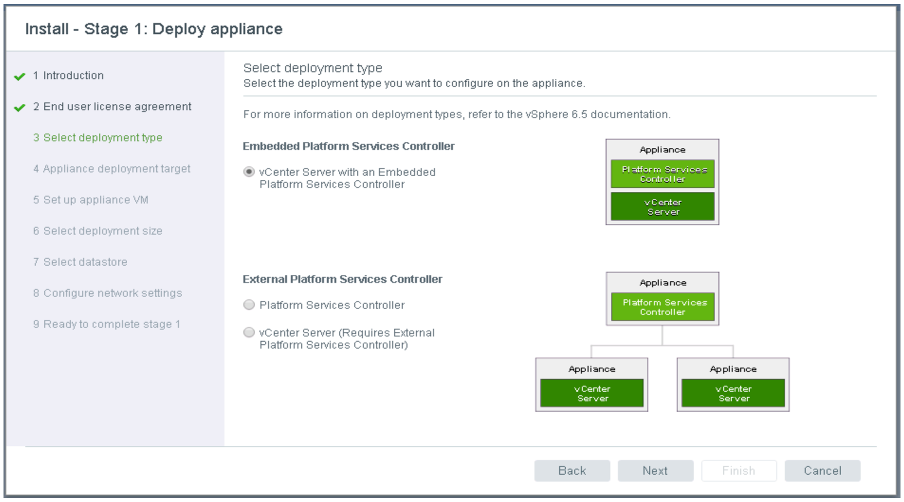
Install - Stage 1: Deploy appliance
✓ 1 Introduction
✓ 2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Appliance
Platform Services Controller
vCenter Server
Appliance
Platform Services Controller
Appliance
vCenter Server
Appliance
vCenter Server
Back
Next
Finish
Cancel

Select deployment type

Select the deployment type you want to configure on the appliance.

For more information on deployment types, refer to the vSphere 6.5 documentation.

Embedded Platform Services Controller

● vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Platform Services Controller

○ vCenter Server (Requires External Platform Services Controller)

8. Provide Host IP or FQDN and User Name, Password Credentials of the Host to connect.

9. Click Next to continue.

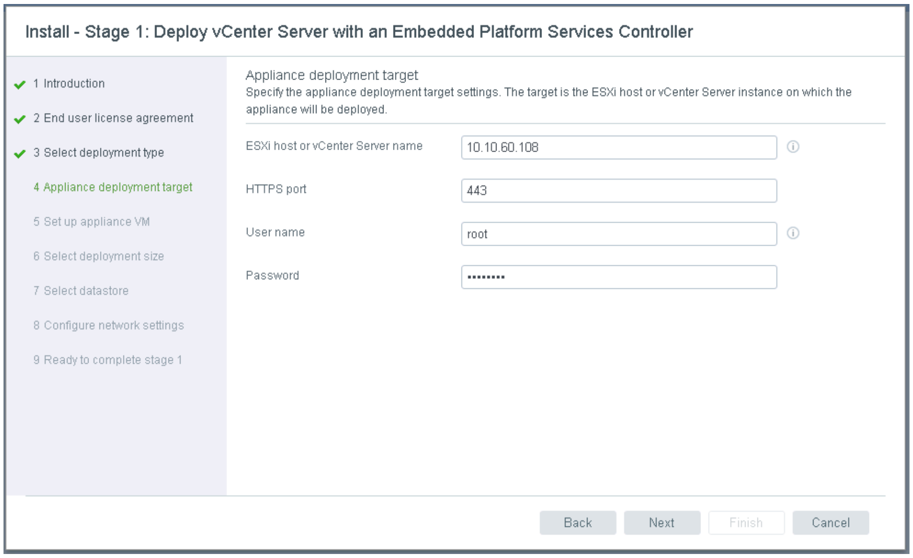
Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller
✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

ESXi host or vCenter Server name	10.10.60.108
HTTPS port	443
User name	root
Password	••••••••

10. Click Yes to accept Certificate Warning.

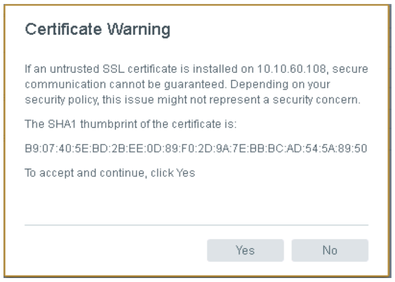

Certificate Warning

If an untrusted SSL certificate is installed on 10.10.60.108, secure communication cannot be guaranteed. Depending on your security policy, this issue might not represent a security concern.

The SHA1 thumbprint of the certificate is:

B9:07:40:5E:BD:2B:EE:0D:89:F0:2D:9A:7E:BB:BC:AD:54:5A:89:50

To accept and continue, click Yes

11. Provide a VM name and a root password for the vCenter appliance.

12. Click Next to continue.

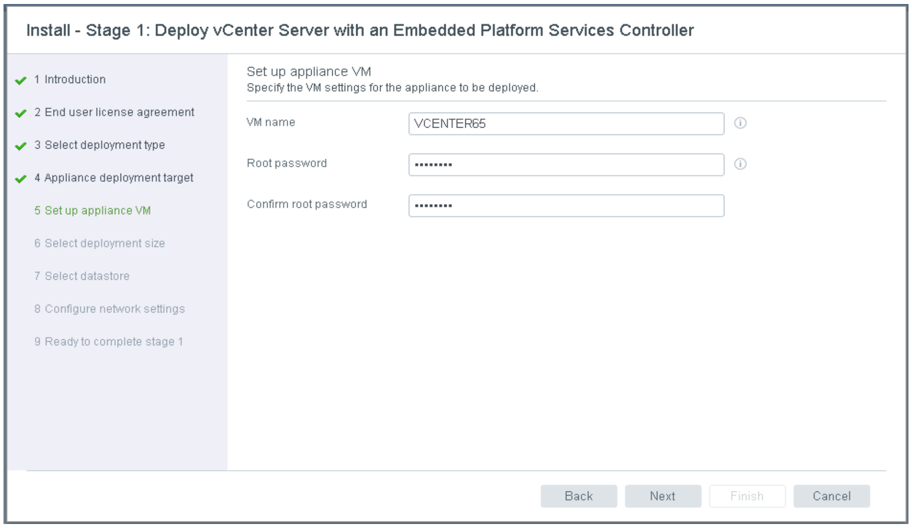

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name: VCENTER65

Root password: ••••••••

Confirm root password: ••••••••

13. Select the proper appliance size for your deployment. In our study, Large was selected.

14. Click Next to continue

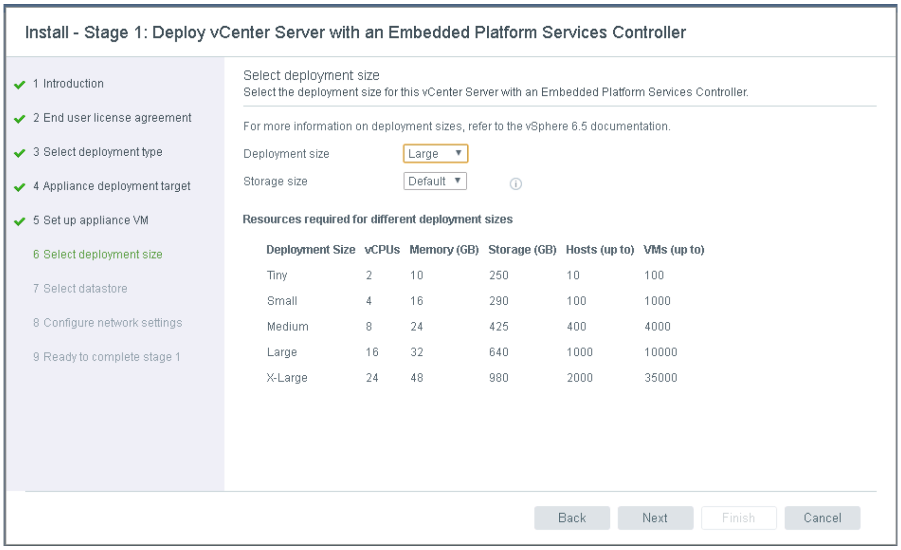

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Select deployment size

Select the deployment size for this vCenter Server with an Embedded Platform Services Controller.

For more information on deployment sizes, refer to the vSphere 6.5 documentation.

Deployment size: Large

Storage size: Default

Resources required for different deployment sizes

Deployment Size	vCPUs	Memory (GB)	Storage (GB)	Hosts (up to)	VMs (up to)
Tiny	2	10	250	10	100
Small	4	16	290	100	1000
Medium	8	24	425	400	4000
Large	16	32	640	1000	10000
X-Large	24	48	980	2000	35000

15. Select appropriate Data store. Check Enable Thin Disk Mode.

16. Click Next to continue.

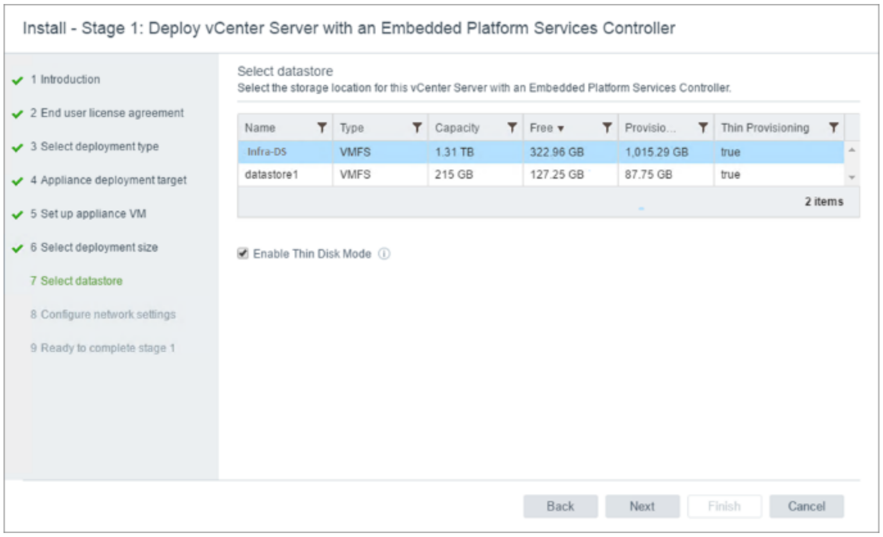
✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
✓ 5 Set up appliance VM
✓ 6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Select datastore

Select the storage location for this vCenter Server with an Embedded Platform Services Controller.

Name	Type	Capacity	Free ▾	Provisio...	Thin Provisioning
Infra-DS	VMFS	1.31 TB	322.96 GB	1,015.29 GB	true
datastore1	VMFS	215 GB	127.25 GB	87.75 GB	true

2 items

☑ Enable Thin Disk Mode

17. Provide Network Settings for the appliance.

It is important to note at this step that you should create a DNS A record for your appliance prior to running the install. The services will fail to startup and your install will fail if it cannot resolve properly.

18. Click Next.

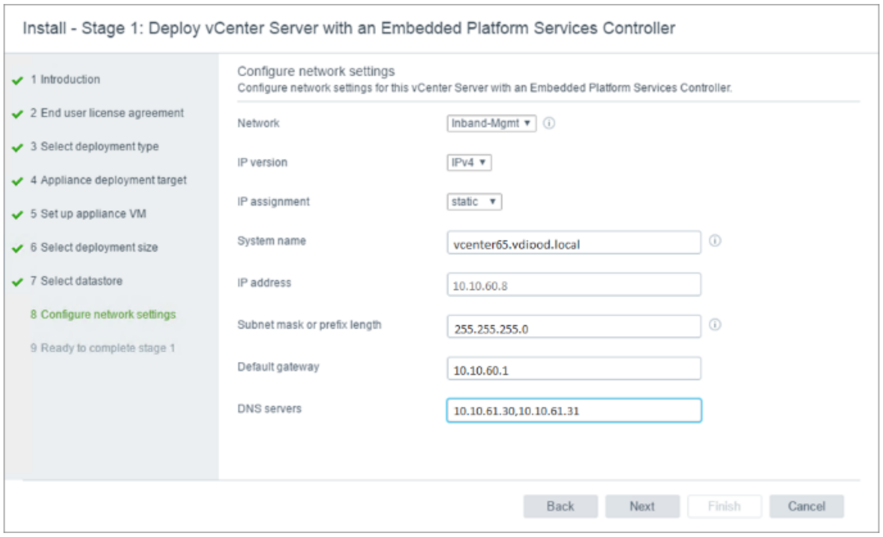
✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
✓ 5 Set up appliance VM
✓ 6 Select deployment size
✓ 7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Configure network settings

Configure network settings for this vCenter Server with an Embedded Platform Services Controller.

Network	Inband-Mgmt ▾
IP version	IPv4 ▾
IP assignment	static ▾
System name	vcenter65.vdipod.local
IP address	10.10.60.8
Subnet mask or prefix length	255.255.255.0
Default gateway	10.10.60.1
DNS servers	10.10.61.30,10.10.61.31

19. Review Settings and click Finish.

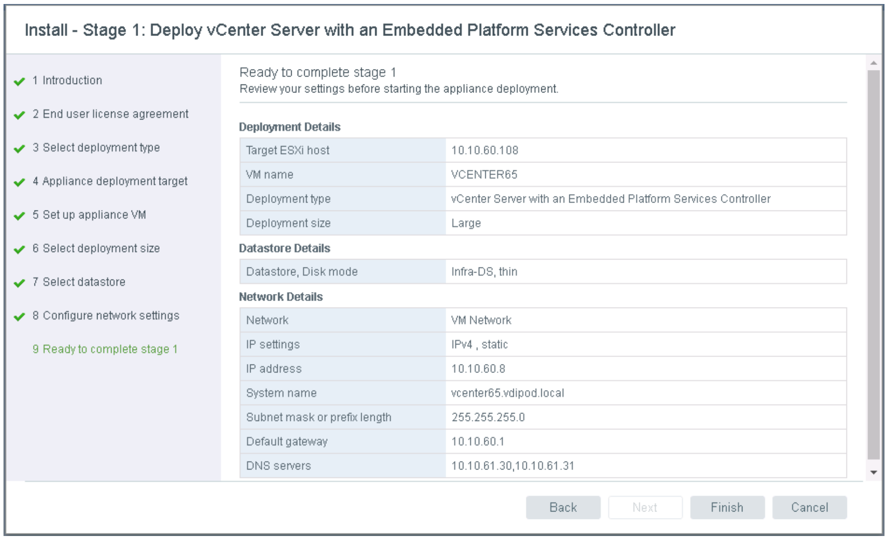

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

✓ 1 Introduction
✓ 2 End user license agreement
✓ 3 Select deployment type
✓ 4 Appliance deployment target
✓ 5 Set up appliance VM
✓ 6 Select deployment size
✓ 7 Select datastore
✓ 8 Configure network settings
9 Ready to complete stage 1

Ready to complete stage 1
Review your settings before starting the appliance deployment.

Deployment Details

Target ESXi host	10.10.60.108
VM name	VCENTER65
Deployment type	vCenter Server with an Embedded Platform Services Controller
Deployment size	Large

Datastore Details

Datastore, Disk mode	Infra-DS, thin

Network Details

Network	VM Network
IP settings	IPv4 , static
IP address	10.10.60.8
System name	vcenter65.vdipod.local
Subnet mask or prefix length	255.255.255.0
Default gateway	10.10.60.1
DNS servers	10.10.61.30,10.10.61.31

20. Once deployment completed click Continue to proceed with set up.

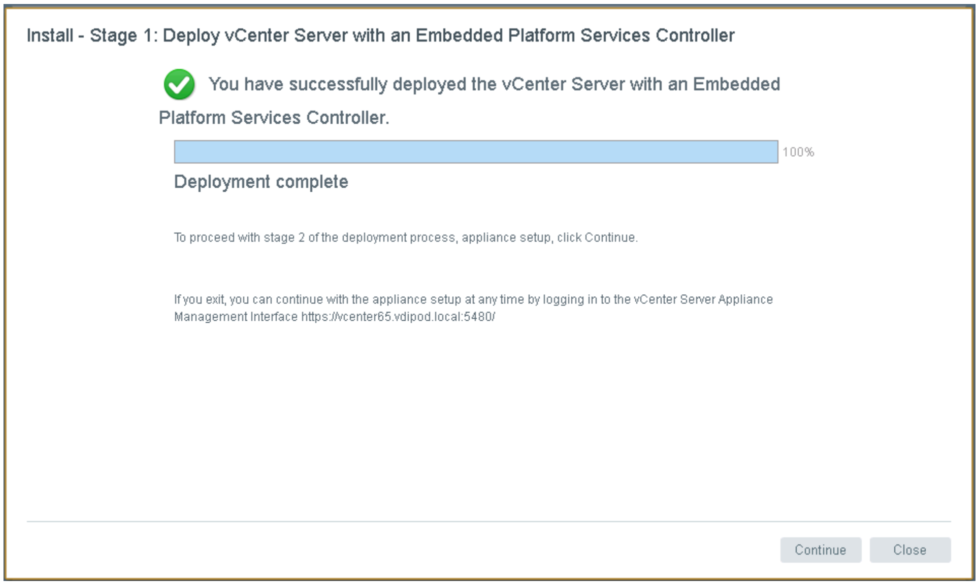

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

You have successfully deployed the vCenter Server with an Embedded Platform Services Controller.

Deployment complete

To proceed with stage 2 of the deployment process, appliance setup, click Continue.

If you exit, you can continue with the appliance setup at any time by logging in to the vCenter Server Appliance Management Interface https://vcenter65.vdipod.local:5480/

21. To start Set Up vCenter Server Appliance with an Embedded PSC click Next.

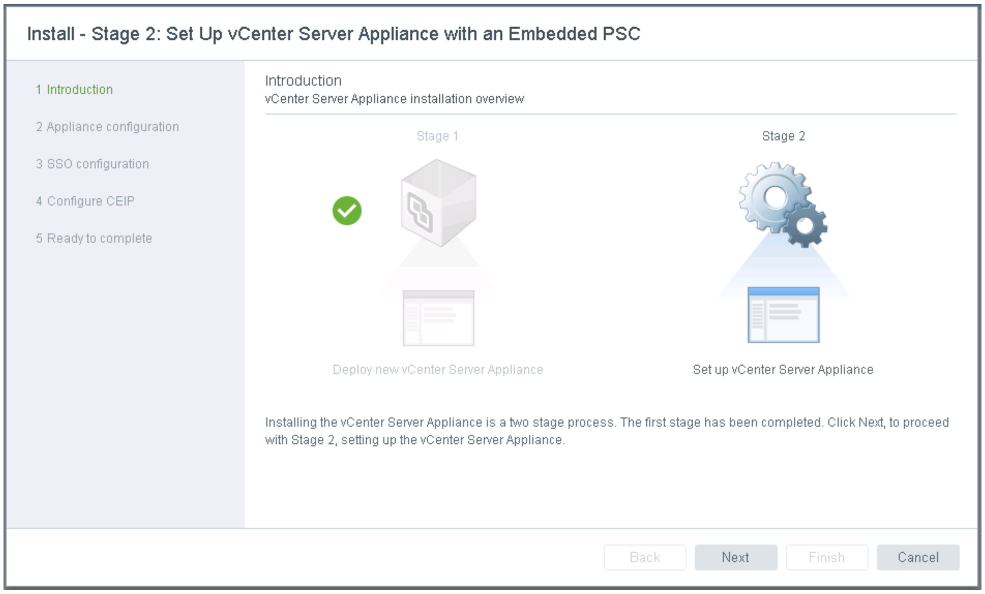

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Introduction
vCenter Server Appliance installation overview

Installing the vCenter Server Appliance is a two stage process. The first stage has been completed. Click Next, to proceed with Stage 2, setting up the vCenter Server Appliance.

22. Configure NTP time synchronization for the appliance.

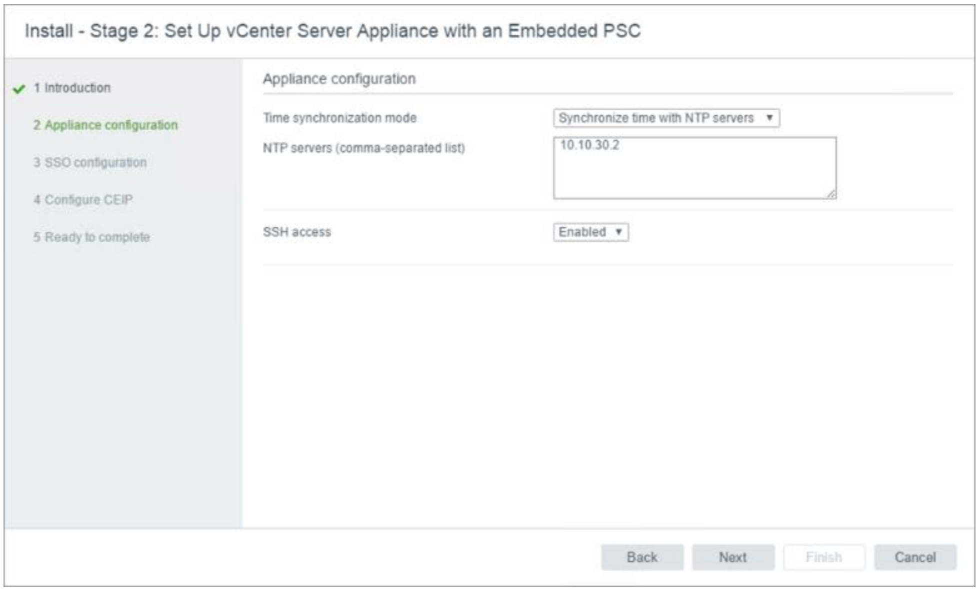

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Appliance configuration

Time synchronization mode: Synchronize time with NTP servers

NTP servers (comma-separated list): 10.10.30.2

SSH access: Enabled

23. Create a new SSO domain (unless your environment already has and SSO domain. Multiple SSO domains can co-exist).

24. Provide Single Sign On Password and Site Name Credentials.

25. Click Next.

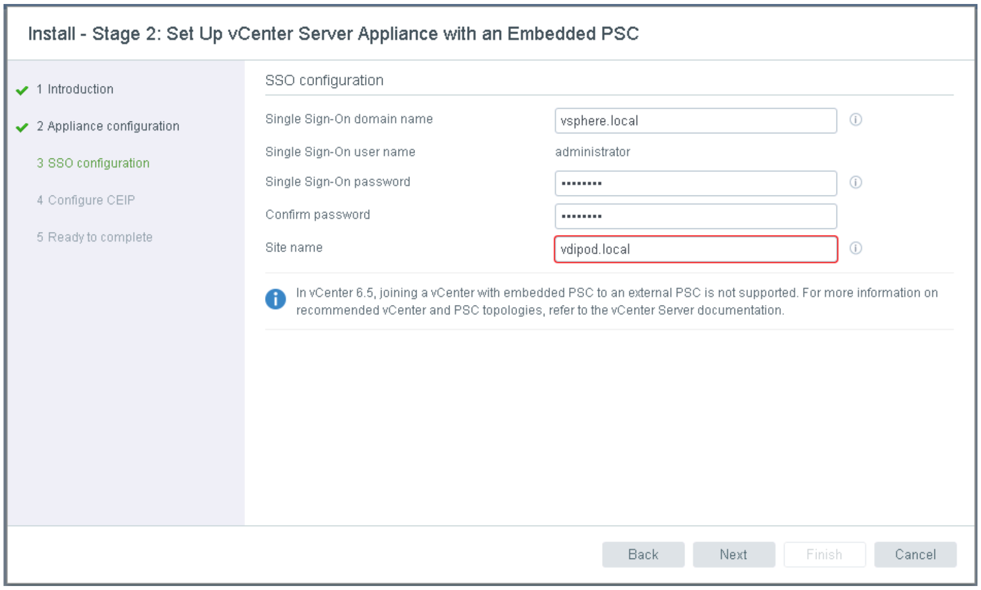

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

SSO configuration

Single Sign-On domain name	vsphere.local
Single Sign-On user name	administrator
Single Sign-On password	••••••••
Confirm password	••••••••
Site name	vdipod.local

In vCenter 6.5, joining a vCenter with embedded PSC to an external PSC is not supported. For more information on recommended vCenter and PSC topologies, refer to the vCenter Server documentation.

26. Configure participation in CEICP program.

27. Click Next.

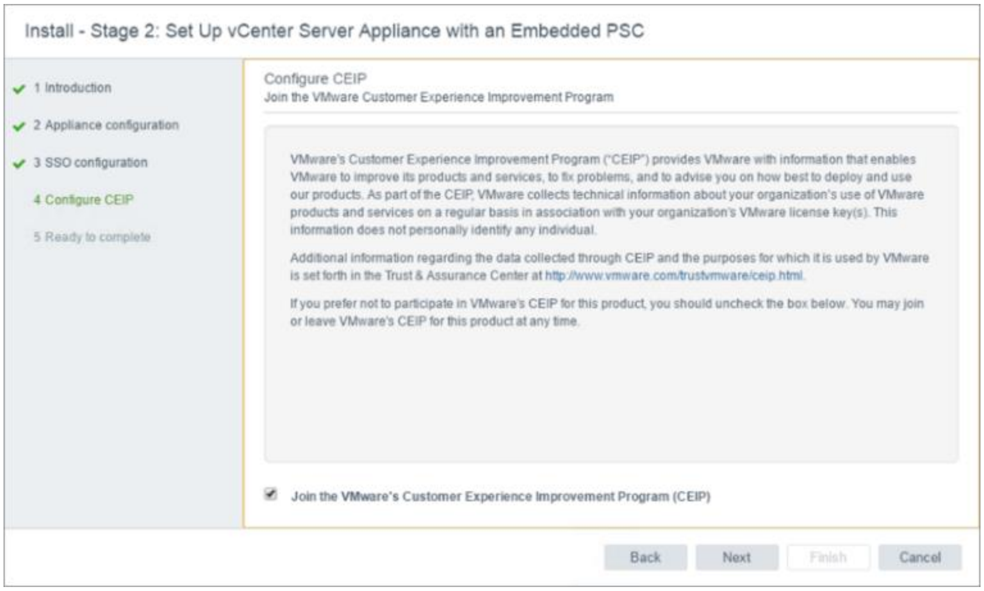

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Configure CEIP

Join the VMware Customer Experience Improvement Program

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. As part of the CEIP, VMware collects technical information about your organization's use of VMware products and services on a regular basis in association with your organization's VMware license key(s). This information does not personally identify any individual.

Additional information regarding the data collected through CEIP and the purposes for which it is used by VMware is set forth in the Trust & Assurance Center at http://www.vmware.com/trustvmware/ceip.html.

If you prefer not to participate in VMware's CEIP for this product, you should uncheck the box below. You may join or leave VMware's CEIP for this product at any time.

☑ Join the VMware's Customer Experience Improvement Program (CEIP)

28. Review Settings and click Finish.

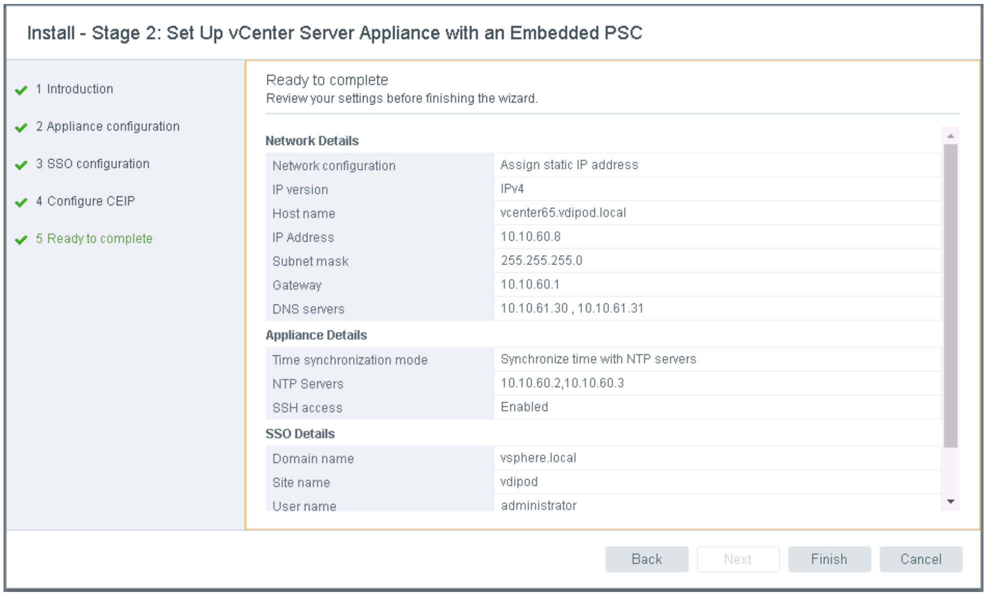

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Ready to complete
Review your settings before finishing the wizard.

Network Details

Network configuration	Assign static IP address
IP version	IPv4
Host name	vcenter65.vdipod.local
IP Address	10.10.60.8
Subnet mask	255.255.255.0
Gateway	10.10.60.1
DNS servers	10.10.61.30 , 10.10.61.31

Appliance Details

Time synchronization mode	Synchronize time with NTP servers
NTP Servers	10.10.60.2,10.10.60.3
SSH access	Enabled

SSO Details

Domain name	vsphere.local
Site name	vdipod
User name	administrator

29. Click Close upon a successful set up.

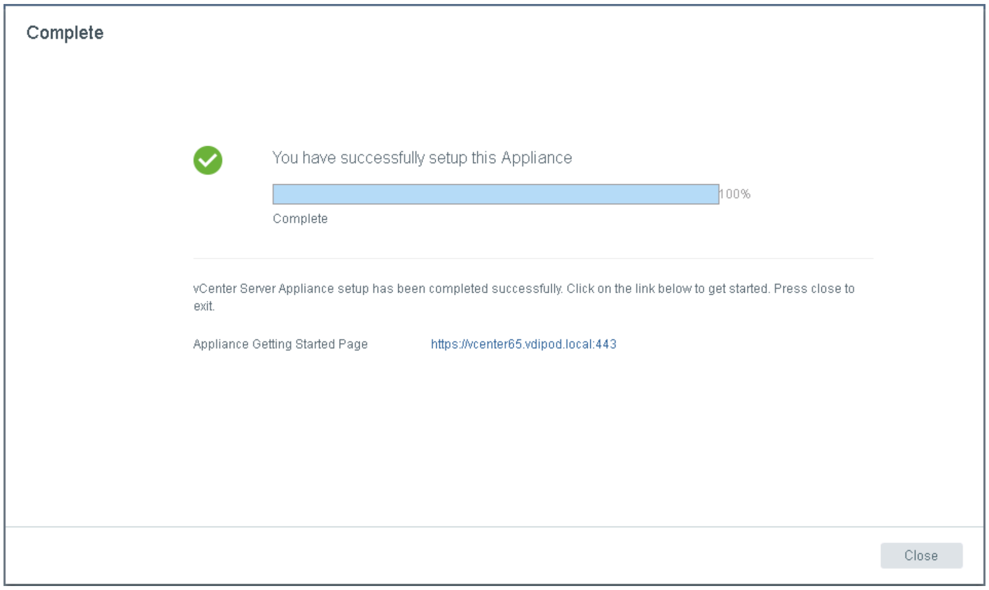

Complete

You have successfully setup this Appliance

100%

Complete

vCenter Server Appliance setup has been completed successfully. Click on the link below to get started. Press close to exit.

Appliance Getting Started Page https://vcenter65.vdipod.local:443

30. Log in using the Single Sign-On username and password created during the vCenter installation into the vSphere Web Client (ie. https://vcenter65.vdilab.local/vsphere-client).

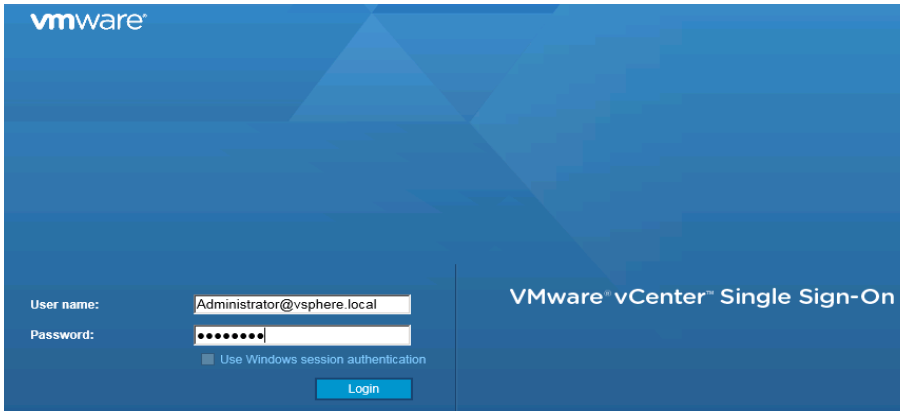

User name: Administrator@vsphere.local

Password: ••••••••

□ Use Windows session authentication

31. Click Create Datacenter in the center pane.

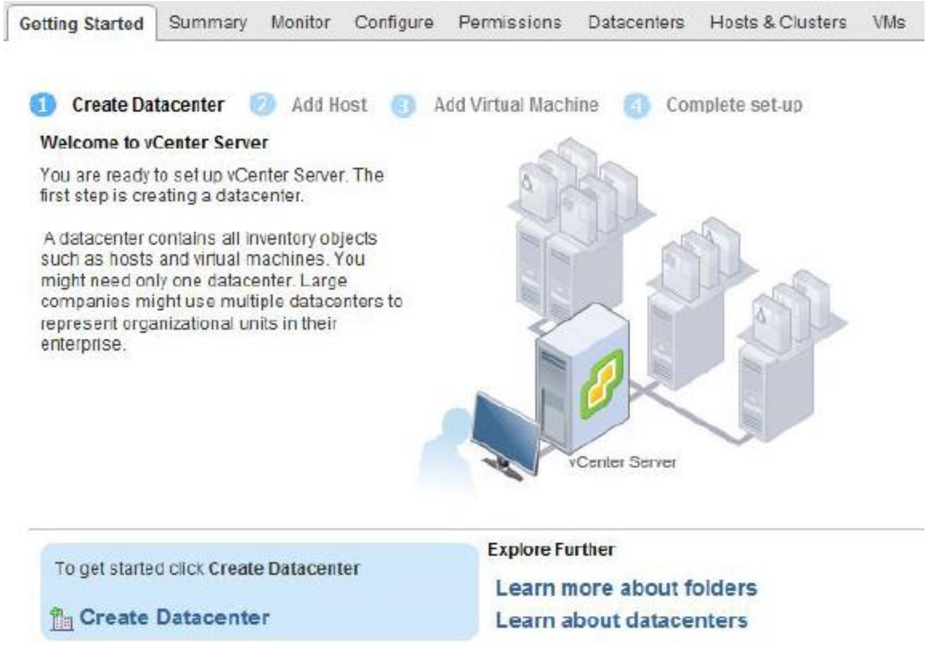

Welcome to vCenter Server

You are ready to set up vCenter Server. The first step is creating a datacenter.

A datacenter contains all inventory objects such as hosts and virtual machines. You might need only one datacenter. Large companies might use multiple datacenters to represent organizational units in their enterprise.

To get started click Create Datacenter

Explore Further

32. Type VDI-DC as the Datacenter name.

33. Click OK to continue.

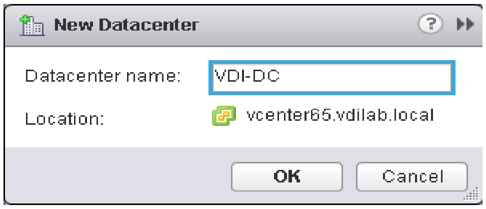

Datacenter name: VDI-DC

Location: vcenter65.vdilab.local

34. Right-click Datacenters > VDI-DC in the list in the center pane, then click New Cluster.

35. Name the cluster Infrastructure.

36. Check the box to turn on DRS. Leave the default values.

Set DRS to Manual for Clusters hosting non persistent desktops.

37. Check the box to turn on vSphere HA. Leave the default values.

38. Click OK to create the new cluster.

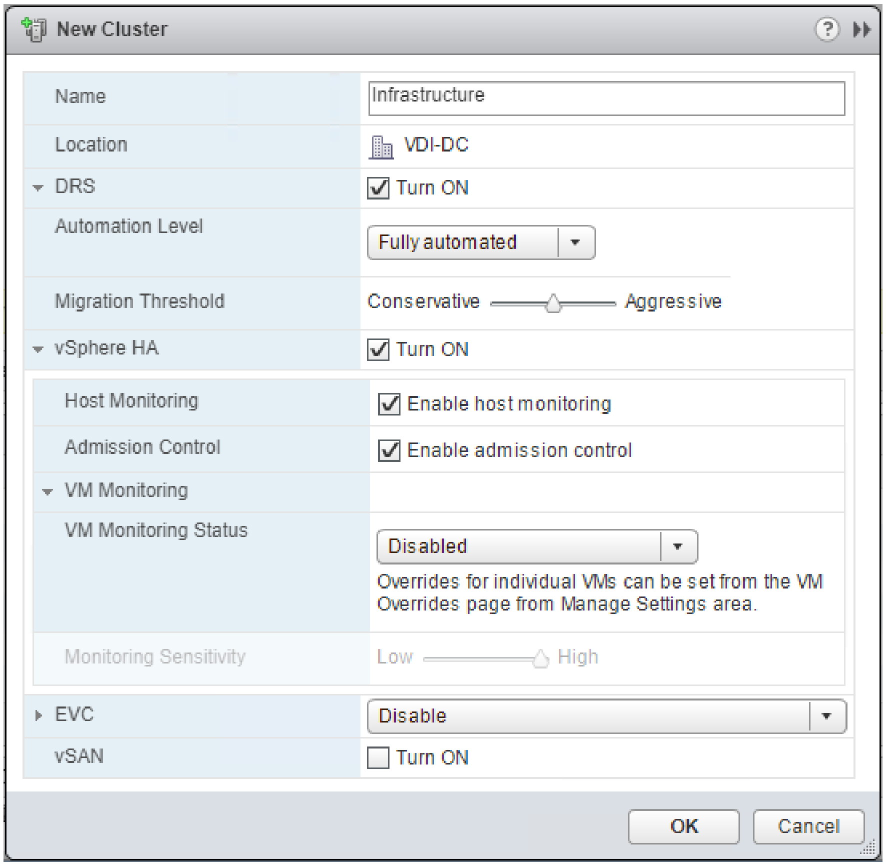

New Cluster

Name	Infrastructure
Location	VDI-DC
DRS	☑ Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative — Aggressive
vSphere HA	☑ Turn ON
Host Monitoring	☑ Enable host monitoring
Admission Control	☑ Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low — High
EVC	Disable
vSAN	□ Turn ON

If mixing Cisco UCS B200 M5 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

39. Right-click Infrastructure in the left pane.

40. Select Add Host.

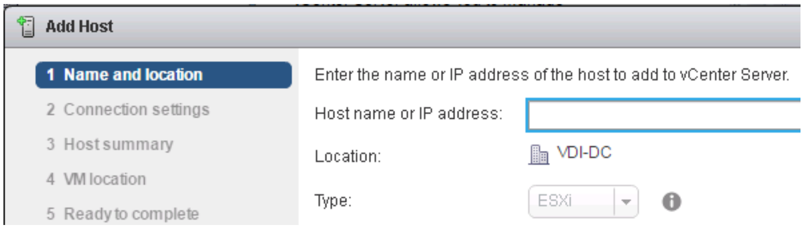

Add Host

Enter the name or IP address of the host to add to vCenter Server.

Host name or IP address:

Location: VDI-DC

Type: ESXi

41. Type the host IP address and click Next.

42. Type root as the user name and root password as the password. Click Next to Continue.

43. Click Yes to accept the certificate.

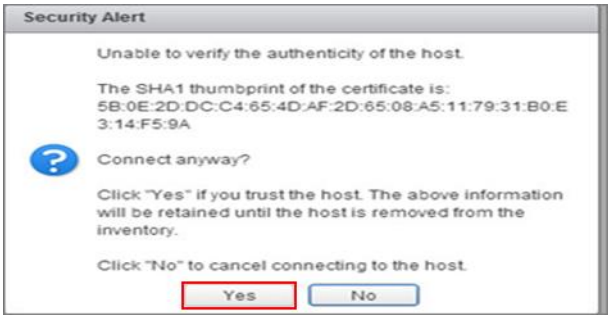

Security Alert

Unable to verify the authenticity of the host.

The SHA1 thumbprint of the certificate is:
5B:0E:2D:DC:C4:65:4D:AF:2D:65:08:A5:11:79:31:B0:E
3:14:F5:9A

Connect anyway?

Click "Yes" if you trust the host. The above information will be retained until the host is removed from the inventory.

Click "No" to cancel connecting to the host.

44. Review the host details and click Next to continue.

45. Assign a license and click Next to continue.

46. Click Next to continue.

47. Click Next to continue.

48. Review the configuration parameters then click Finish to add the host.

49. Repeat this for the other hosts and clusters.

50. When completed, the vCenter cluster configuration is comprised of the following clusters, including a cluster to manage the workload launcher hosts:

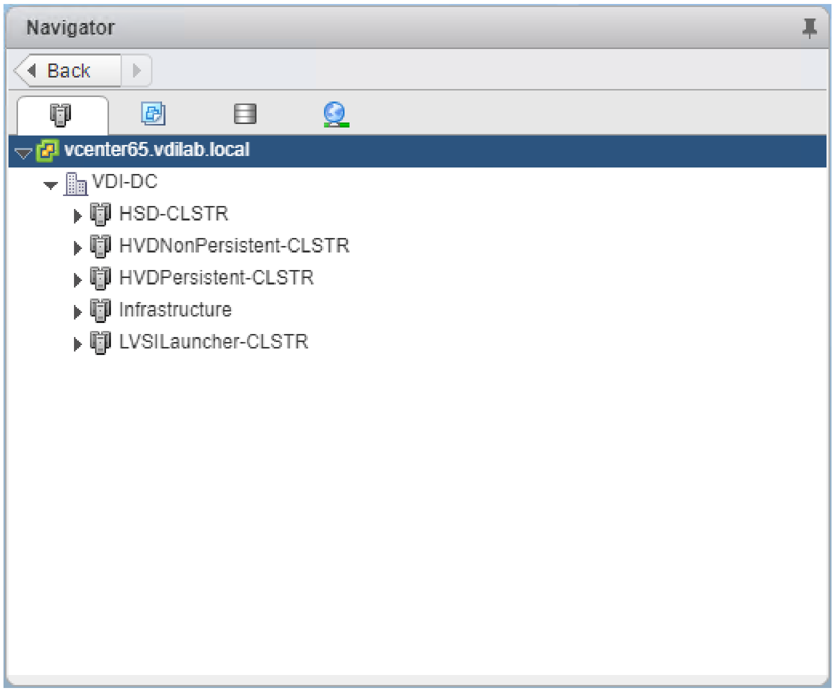

- vcenter65.vdilab.local
  - VDI-DC
    - HSD-CLSTR
    - HVDNonPersistent-CLSTR
    - HVDPersistent-CLSTR
    - Infrastructure
    - LVSILauncher-CLSTR

Configure NTP on ESXi Hosts

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From each vSphere Client, select the host in the inventory.

2. Click the Configuration tab to enable configurations.
3. Click Time Configuration in the Software pane.
4. Click Properties at the upper right side of the window.
5. At the bottom of the Time Configuration dialog box, click Options.
6. In the NTP Daemon Options dialog box, complete the following steps:
    - Click General in the left pane and select Start and stop with host.
    - Click NTP Settings in the left pane and click Add.
7. In the Add NTP Server dialog box, enter <<var_global_ntp_server_ip>> as the IP address of the NTP server and click OK.
8. In the NTP Daemon Options dialog box, select the Restart NTP Service to Apply Changes checkbox and click OK.
9. In the Time Configuration dialog box, complete the following steps:
    - Select the NTP Client Enabled checkbox and click OK.
    - Verify that the clock is now set to approximately the correct time.

The NTP server time may vary slightly from the host time.

ESXi Dump Collector Setup for SAN-Booted Hosts

ESXi hosts booted from SAN need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, complete the following steps:

1. In the vSphere web client, select Home > Administration.
2. In the left hand pane, click System Configuration.
3. In the left hand pane, click Services > VMware vSphere ESXi Dump Collector.
4. In the Actions menu, choose Start.
5. In the Actions menu, click Edit Startup Type.
6. Select Automatic.
7. Click OK.

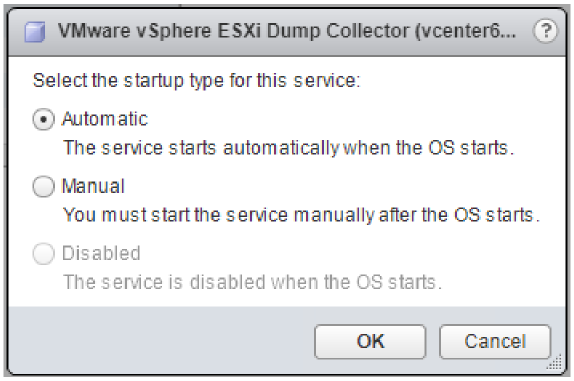
VMware vSphere ESXi Dump Collector (vcenter6...

Select the startup type for this service:

● Automatic
The service starts automatically when the OS starts.

○ Manual
You must start the service manually after the OS starts.

○ Disabled
The service is disabled when the OS starts.

8. On the Management Workstation, open the VMware vSphere CLI command prompt.

9. Set each SAN-booted ESXi Host to coredump to the ESXi Dump Collector by running the following commands:

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> system coredump network set --interface-name vmk0 --server-
ipv4 <<var_vcenter_server_ip> --server-port 6500
```

To get the host thumbprint, type the command without the –-thumbprint option, then copy and paste the thumbprint into the command.

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> system coredump network set --enable true
```

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> system coredump network check
```

FlexPod VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for installing the VMware vDS on the FlexPod ESXi Desktop Workload Hosts.

In the Cisco UCS Configuration section of this document one set of vNICs (A and B) was setup. The vmnic ports associated with the A and B vNICs will be migrated to VMware vDS in this procedure. The critical infrastructure VLAN interfaces and vMotion interfaces will be placed on the vDS.

Configure the VMware vDS in vCenter

To configure the vDS from VMware vSphere Web Client, complete the following steps:

1. After logging into the VMware vSphere Web Client, select Networking under the Home tab.

2. Right-click the VDI-DC datacenter and select Distributed Switch > New Distributed Switch.

3. Give the Distributed Switch a name VDI-DVS and click Next.

4. Make sure Distributed switch: 6.5.0 is selected and click Next.

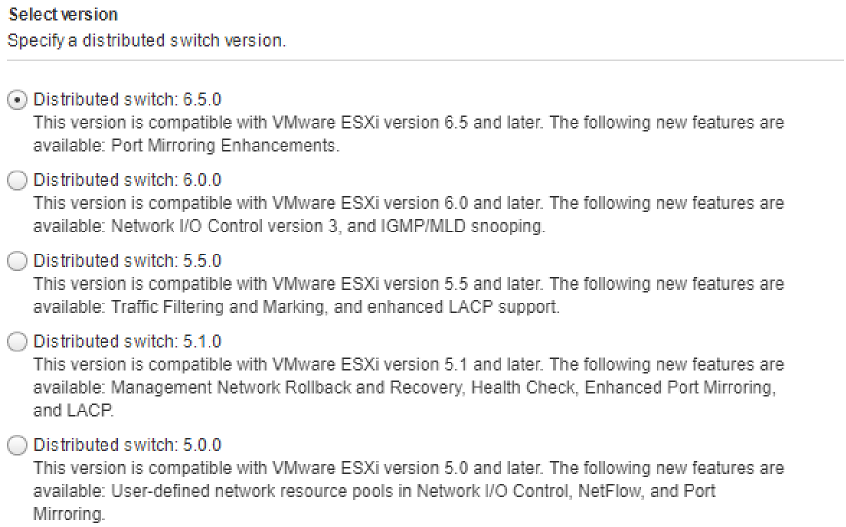

Select version

Specify a distributed switch version.

● Distributed switch: 6.5.0
This version is compatible with VMware ESXi version 6.5 and later. The following new features are available: Port Mirroring Enhancements.

○ Distributed switch: 6.0.0
This version is compatible with VMware ESXi version 6.0 and later. The following new features are available: Network I/O Control version 3, and IGMP/MLD snooping.

○ Distributed switch: 5.5.0
This version is compatible with VMware ESXi version 5.5 and later. The following new features are available: Traffic Filtering and Marking, and enhanced LACP support.

○ Distributed switch: 5.1.0
This version is compatible with VMware ESXi version 5.1 and later. The following new features are available: Management Network Rollback and Recovery, Health Check, Enhanced Port Mirroring, and LACP.

○ Distributed switch: 5.0.0
This version is compatible with VMware ESXi version 5.0 and later. The following new features are available: User-defined network resource pools in Network I/O Control, NetFlow, and Port Mirroring.

5. Change the Number of uplinks to 2. Leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Do not check Create default Port group name. Click Next.

6. Review the information and click Finish to complete creating the vDS.

7. On the left, expand the VDI-DC datacenter and the newly created vDS. Select the newly created vDS.

8. In the center pane, select the New Distributed Port Group icon. Configure the following port groups:

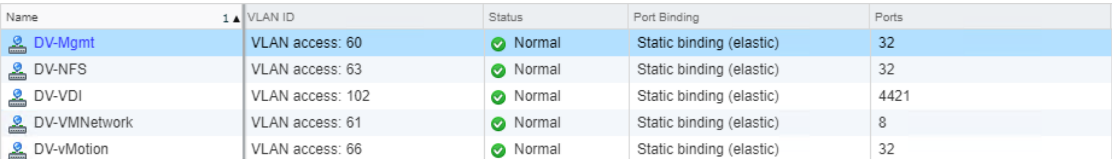

Name 1 ▲	VLAN ID	Status	Port Binding	Ports
DV-Mgmt	VLAN access: 60	✓ Normal	Static binding (elastic)	32
DV-NFS	VLAN access: 63	✓ Normal	Static binding (elastic)	32
DV-VDI	VLAN access: 102	✓ Normal	Static binding (elastic)	4421
DV-VMNetwork	VLAN access: 61	✓ Normal	Static binding (elastic)	8
DV-vMotion	VLAN access: 66	✓ Normal	Static binding (elastic)	32

9. Edit newly created Port Groups. Go to teaming and failover and using arrows place Uplink 1 and Uplink 2 on the list of Active uplinks.

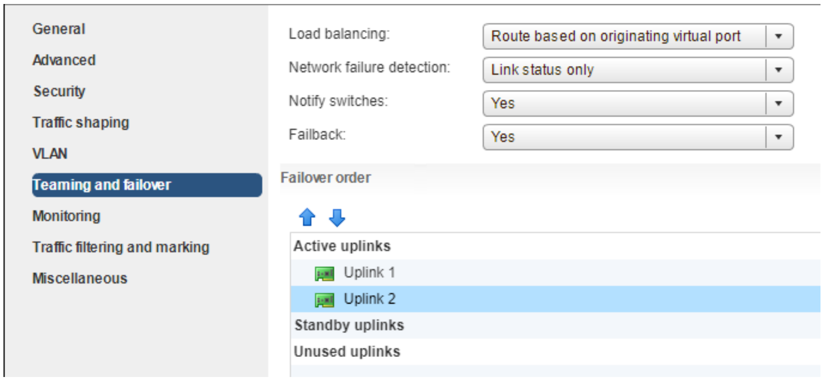

Load balancing: Route based on originating virtual port

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active uplinks
Uplink 1
Uplink 2
Standby uplinks
Unused uplinks

10. Select the vDS on the left. Click the Edit distributed switch settings icon on the right.

11. On the left in the Edit Settings window, select Advanced.

12. Change the MTU to 9000. Click OK.

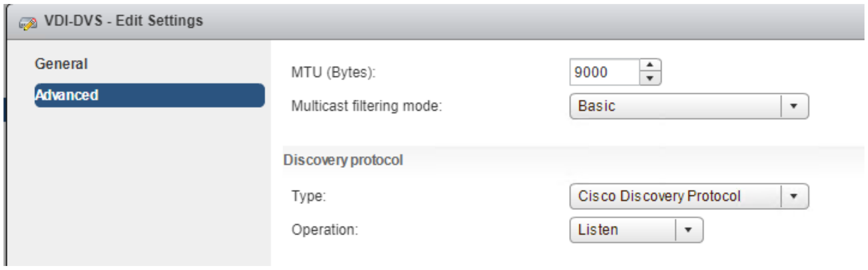

VDI-DVS - Edit Settings

General

Advanced

MTU (Bytes): 9000

Multicast filtering mode: Basic

Discovery protocol

Type: Cisco Discovery Protocol

Operation: Listen

Add Hosts to a vSphere Distributed Switch in the vSphere Web Client

After a vSphere VDI-DVS distributed switch is created, add hosts from each vSphere cluster and physical adapters to create FlexPod virtual network using these steps:

1. Right-click VDI-DVS distributed switch in the vSphere Web Client and select Add and Manage Hosts.

2. On the Select tasks page, select Add hosts.

3. Click Next.

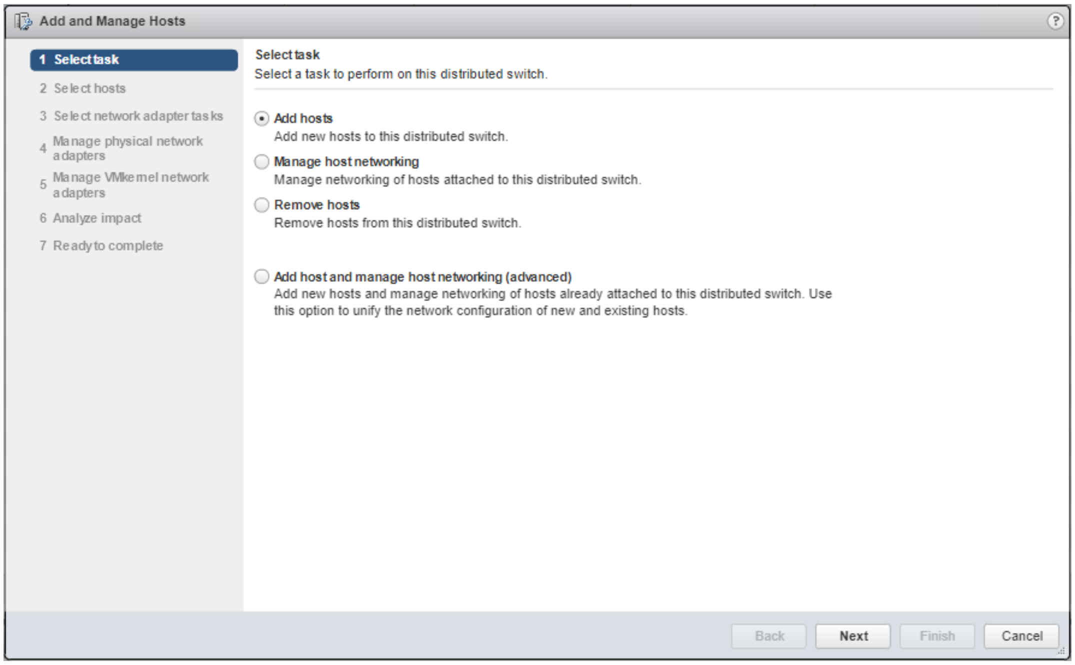

Add and Manage Hosts

1 Select task
2 Select hosts
3 Select network adapter tasks
4 Manage physical network adapters
5 Manage VMkernel network adapters
6 Analyze impact
7 Ready to complete

Select task
Select a task to perform on this distributed switch.

● Add hosts
Add new hosts to this distributed switch.

○ Manage host networking
Manage networking of hosts attached to this distributed switch.

○ Remove hosts
Remove hosts from this distributed switch.

○ Add host and manage host networking (advanced)
Add new hosts and manage networking of hosts already attached to this distributed switch. Use this option to unify the network configuration of new and existing hosts.

4. On the Select hosts page, click Add New hosts. The select new hosts dialog box opens. Select a host from the list and click OK.

5. Click Next.

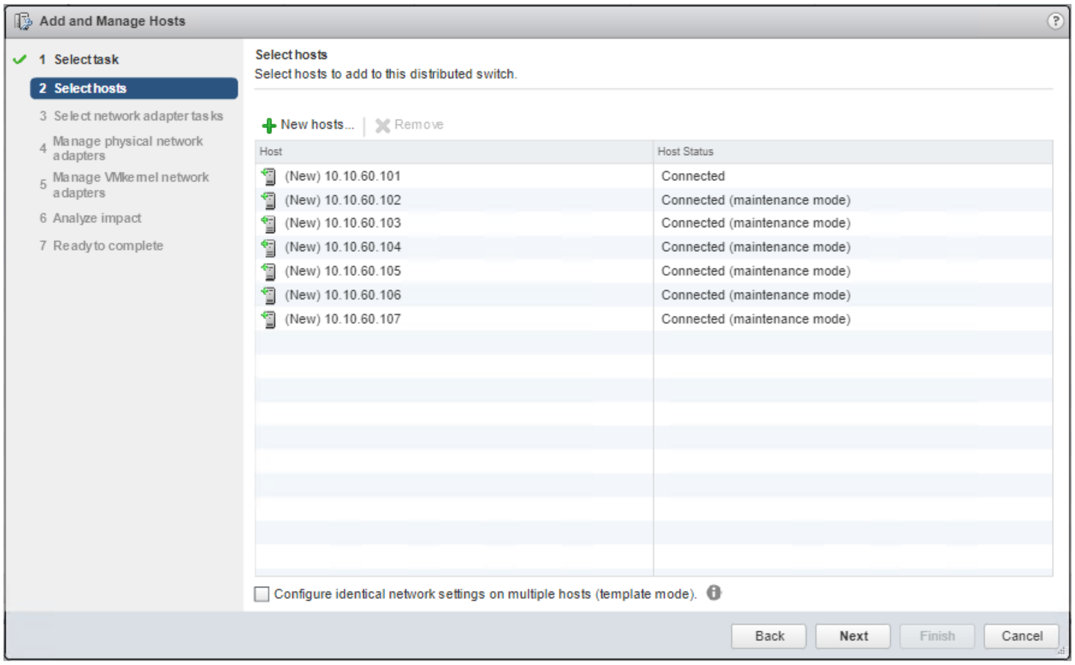

Add and Manage Hosts

✓ 1 Select task
2 Select hosts
3 Select network adapter tasks
4 Manage physical network adapters
5 Manage VMkernel network adapters
6 Analyze impact
7 Ready to complete

Select hosts
Select hosts to add to this distributed switch.

New hosts... | Remove

Host	Host Status
(New) 10.10.60.101	Connected
(New) 10.10.60.102	Connected (maintenance mode)
(New) 10.10.60.103	Connected (maintenance mode)
(New) 10.10.60.104	Connected (maintenance mode)
(New) 10.10.60.105	Connected (maintenance mode)
(New) 10.10.60.106	Connected (maintenance mode)
(New) 10.10.60.107	Connected (maintenance mode)

☐ Configure identical network settings on multiple hosts (template mode).

Back | Next | Finish | Cancel

6. Select Manage physical network adapters and Manage VMkernel adapters.

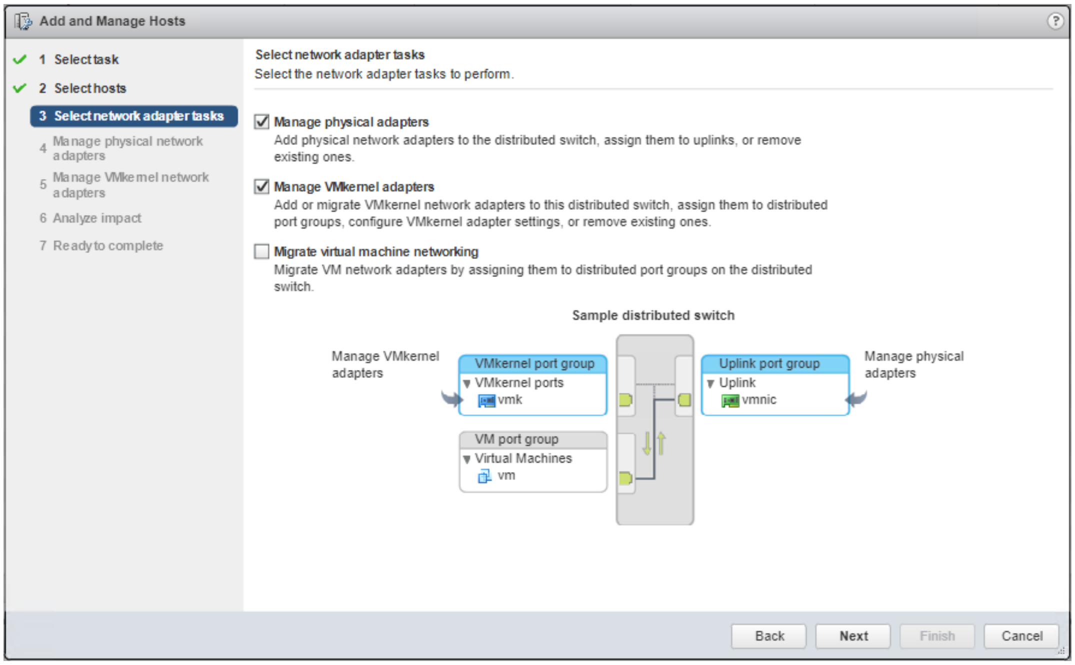

Add and Manage Hosts

✓ 1 Select task
✓ 2 Select hosts
3 Select network adapter tasks
4 Manage physical network adapters
5 Manage VMkernel network adapters
6 Analyze impact
7 Ready to complete

Select network adapter tasks
Select the network adapter tasks to perform.

☑ Manage physical adapters
Add physical network adapters to the distributed switch, assign them to uplinks, or remove existing ones.

☑ Manage VMkernel adapters
Add or migrate VMkernel network adapters to this distributed switch, assign them to distributed port groups, configure VMkernel adapter settings, or remove existing ones.

☐ Migrate virtual machine networking
Migrate VM network adapters by assigning them to distributed port groups on the distributed switch.

Sample distributed switch

Back | Next | Finish | Cancel

7. On the Manage physical network adapters page, configure physical NICs on the distributed switch.

   – From the On other switches/unclaimed list, select a physical NIC.

> If you select physical NICs that are already connected to other switches, they are migrated to the current distributed switch.

   – Click Assign uplink.

   – Select an uplink and click OK.

8. Click Next.

9. On the Manage VMkernel network adapters page, configure VMkernel adapters.

   – Select Host management VMkernel adapter and click Assign port group.

   – Select a DV-Mgmt distributed port group and click OK.

10. Repeat the process for the vMotion VMkernel adapter. Use appropriate distributed port group (DV-vMotion).

11. Click Next.

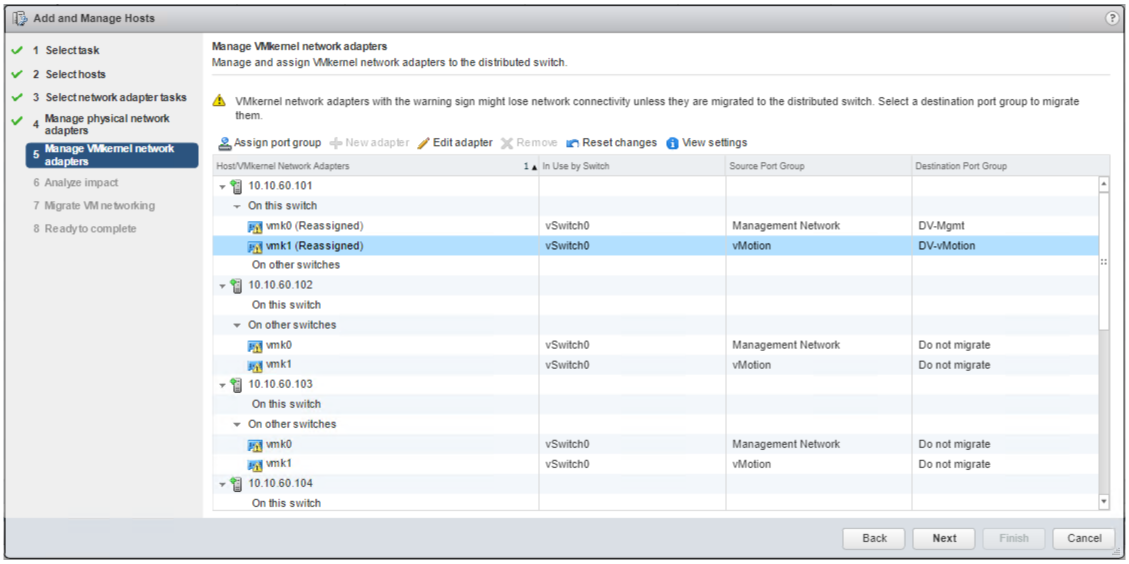

Add and Management Hosts

✓ 1 Select task
✓ 2 Select hosts
✓ 3 Select network adapter tasks
✓ 4 Manage physical network adapters
5 Manage VMkernel network adapters
6 Analyze impact
7 Migrate VM networking
8 Ready to complete

Manage VMkernel network adapters
Manage and assign VMkernel network adapters to the distributed switch.

VMkernel network adapters with the warning sign might lose network connectivity unless they are migrated to the distributed switch. Select a destination port group to migrate them.

Host/VMkernel Network Adapters	In Use by Switch	Source Port Group	Destination Port Group
10.10.60.101			
On this switch			
vmk0 (Reassigned)	vSwitch0	Management Network	DV-Mgmt
vmk1 (Reassigned)	vSwitch0	vMotion	DV-vMotion
On other switches			
10.10.60.102			
On this switch			
On other switches			
vmk0	vSwitch0	Management Network	Do not migrate
vmk1	vSwitch0	vMotion	Do not migrate
10.10.60.103			
On this switch			
On other switches			
vmk0	vSwitch0	Management Network	Do not migrate
vmk1	vSwitch0	vMotion	Do not migrate
10.10.60.104			
On this switch			

12. Review the impacted services as well as the level of impact.

13. Click Next.

14. On the Ready to complete page, review the settings you selected and click Finish.

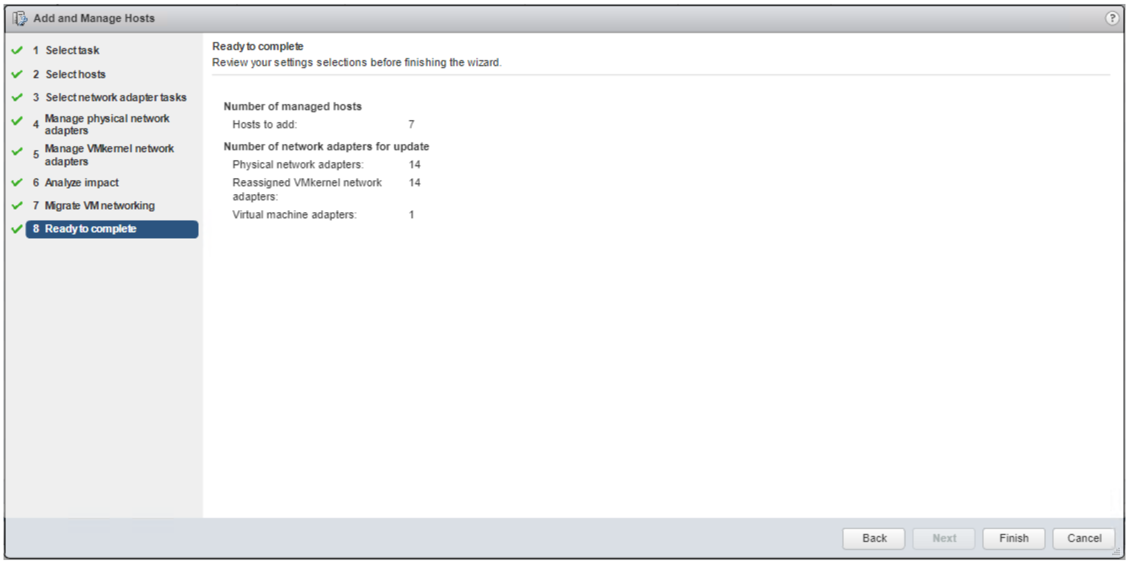

Add and Manage Hosts

✓ 1 Select task
✓ 2 Select hosts
✓ 3 Select network adapter tasks
✓ 4 Manage physical network adapters
✓ 5 Manage VMkernel network adapters
✓ 6 Analyze impact
✓ 7 Migrate VM networking
✓ 8 Ready to complete

Ready to complete
Review your settings selections before finishing the wizard.

Number of managed hosts

Hosts to add:	7

Number of network adapters for update

Physical network adapters:	14
Reassigned VMkernel network adapters:	14
Virtual machine adapters:	1

Building the Virtual Machines and Environment

Software Infrastructure Configuration

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

Install and configure the infrastructure virtual machines by following the process provided in Table 43 .

Table 43 Test Infrastructure Virtual Machine Configuration

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Operating system	Microsoft Windows Server 2016	Microsoft Windows Server 2016
Virtual CPU amount	6	8
Memory amount	8 GB	16 GB
Network	VMXNET3 Infra-Mgmt	VMXNET3 VDI
Disk-1 (OS) size	40 GB	40 GB
Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2016	VCSA – SUSE Linux
Virtual CPU amount	2	16
Memory amount	4 GB	32 GB
Network	VMXNET3 Infra-Mgmt	VMXNET3 In-Band-Mgmt
Disk size	40 GB	599 GB (across 12 VMDKs)

Configuration	Microsoft SQL Server Virtual Machine	Citrix StoreFront Controller Virtual Machine
Operating system	Microsoft Windows Server 2016 Microsoft SQL Server 2012 SP1	Microsoft Windows Server 2016
Virtual CPU amount	6	4

Memory amount	24GB	8 GB
Network	VMXNET3 Infra-Mgmt	VMXNET3 Infra-Mgmt
Disk-1 (OS) size	40 GB	40 GB
Disk-2 size	100 GB SQL Databases\Logs	-

Preparing the Master Targets

This section provides guidance around creating the golden (or master) images for the environment. VMs for the master targets must first be installed with the software components needed to build the golden images. For this CVD, the images contain the basics needed to run the Login VSI workload.

To prepare the master VMs for the Hosted Virtual Desktops (HVDs) and Hosted Shared Desktops (HSDs), there are three major steps: installing the PVS Target Device x64 software, installing the Virtual Delivery Agents (VDAs), and installing application software.

The master target Hosted Virtual Desktop (HVD) and Hosted Shared Desktop (HSD) VMs were configured as detailed in Table 44 :

Table 44 VDI and RDS Configurations

Configuration	HVD Virtual Machines	HSD Virtual Machines
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2016
Virtual CPU amount	2	9
Memory amount	2 GB reserve for all guest memory	24 GB reserve for all guest memory
Network	VMXNET3 DV-VDI	VMXNET3 DV-VDI
Citrix PVS vDisk size Full Clone Disk Size	24 GB (dynamic) 100 GB	40 GB (dynamic)
Citrix PVS write cache Disk size	6 GB	30 GB
Citrix PVS write cache RAM cache size	64 MB	1024 MB
Additional software used for testing	Microsoft Office 2016 Login VSI 4.1.25 (Knowledge Worker Workload)	Microsoft Office 2016 Login VSI 4.1.25 (Knowledge Worker Workload)

Installing and Configuring XenDesktop and XenApp

This section details the installation of the core components of the XenDesktop/XenApp 7.15 system. This CVD installs two XenDesktop Delivery Controllers to support both hosted shared desktops (HSD), non-persistent hosted virtual desktops (HVD), and persistent hosted virtual desktops (HVD).

Prerequisites

Citrix recommends that you use Secure HTTP (HTTPS) and a digital certificate to protect vSphere communications. Citrix recommends that you use a digital certificate issued by a certificate authority (CA) according to your organization's security policy. Otherwise, if security policy allows, use the VMware-installed self-signed certificate.

To install vCenter Server self-signed Certificate, complete the following steps:

1. Add the FQDN of the computer running vCenter Server to the hosts file on that server, located at SystemRoot/ WINDOWS/system32/Drivers/etc/. This step is required only if the FQDN of the computer running vCenter Server is not already present in DNS.
2. Open Internet Explorer and enter the address of the computer running vCenter Server (for example, https://FQDN as the URL).
3. Accept the security warnings.
4. Click the Certificate Error in the Security Status bar and select **View certificates**.
5. Click Install certificate, select Local Machine, and then click **Next**.
6. Select Place all certificates in the following store and then click **Browse.**
7. Select Show physical stores.
8. Select Trusted People.

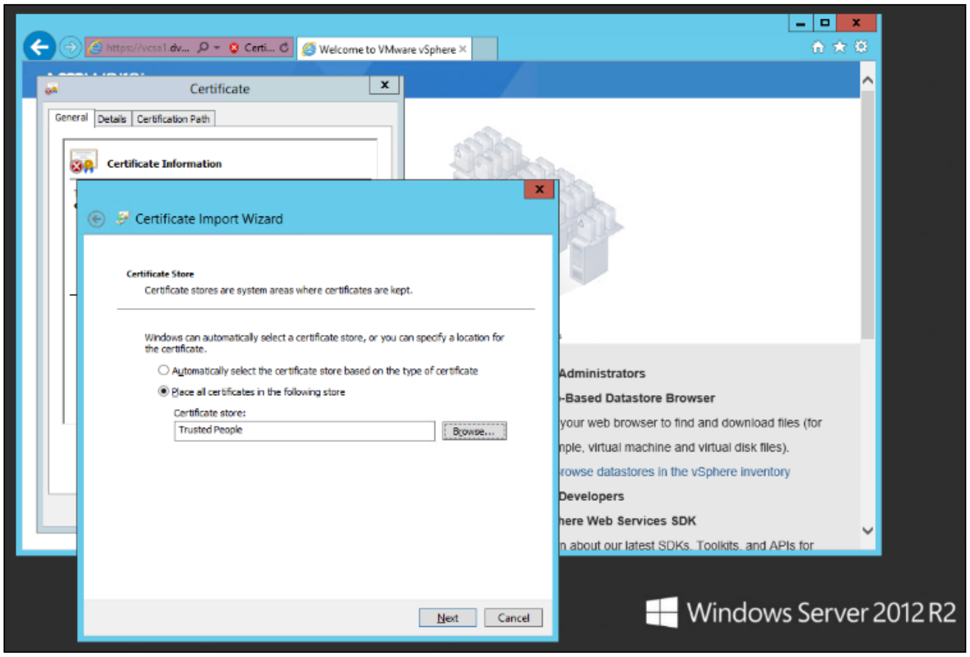

https://vcsa1.dv...
Certi...
Welcome to VMware vSphere
Certificate
General
Details
Certification Path
Certificate Information
Certificate Import Wizard

Certificate Store

Certificate stores are system areas where certificates are kept.

Windows can automatically select a certificate store, or you can specify a location for the certificate.

○ Automatically select the certificate store based on the type of certificate

● Place all certificates in the following store

Certificate store: Trusted People

Browse...
Next
Cancel
Administrators
-Based Datastore Browser
your web browser to find and download files (for
nple, virtual machine and virtual disk files).
rowse datastores in the vSphere inventory
Developers
here Web Services SDK
n about our latest SDKs, Toolkits, and APIs for
Windows Server 2012 R2

9. Click **Next** and then click **Finish**.

10. Perform the above steps on all Delivery Controllers and Provisioning Servers.

Install XenDesktop Delivery Controller, Citrix Licensing, and StoreFront

The process of installing the XenDesktop Delivery Controller also installs other key XenDesktop software components, including Studio, which is used to create and manage infrastructure components, and Director, which is used to monitor performance and troubleshoot problems.

Dedicated StoreFront and License servers should be implemented for large scale deployments.

Installing Citrix License Server

To install the Citrix License Server, complete the following steps:

1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix XenDesktop 7.15 ISO.

2. Click Start.

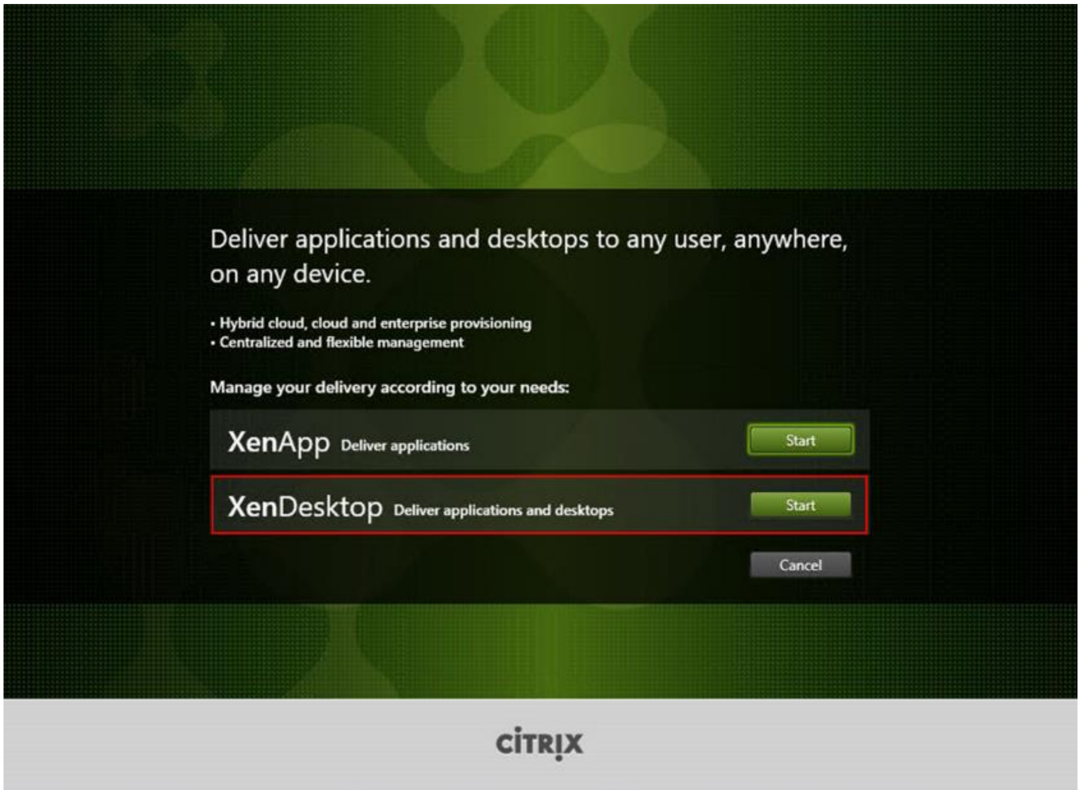

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications
Start
XenDesktop Deliver applications and desktops
Start
Cancel
CITRIX

3. Click “Extend Deployment – Citrix License Server.”

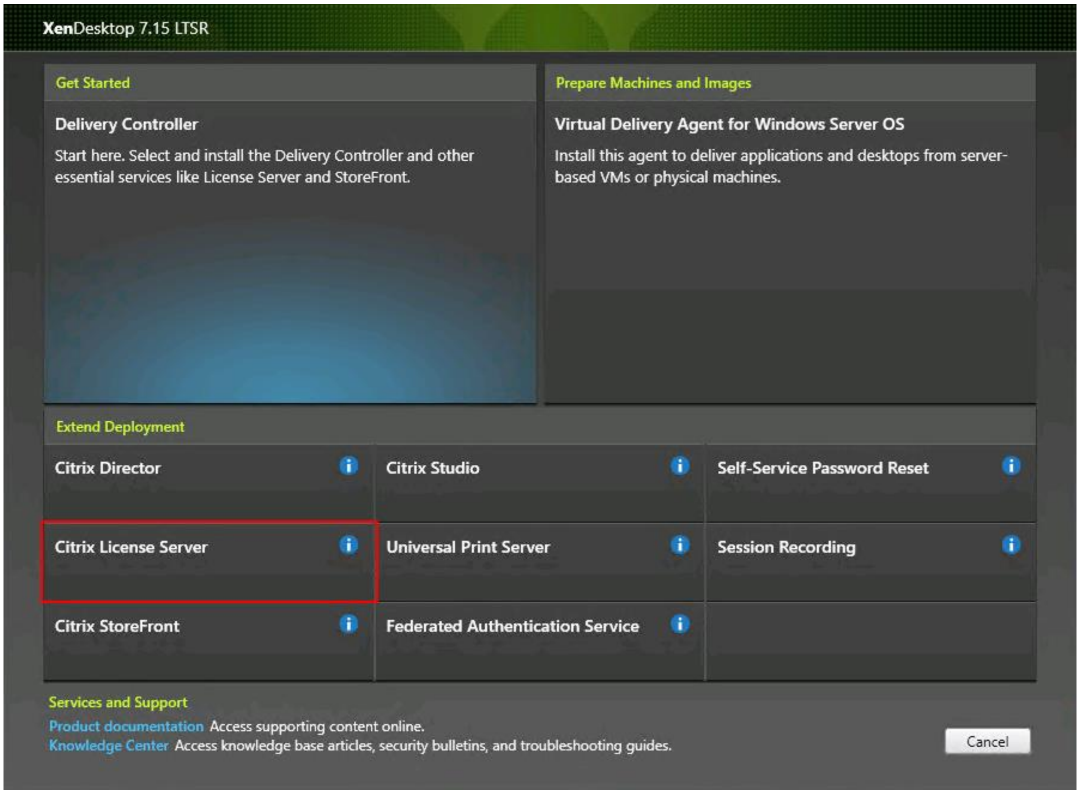

XenDesktop 7.15 LTSR

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Extend Deployment

Services and Support

Product documentation Access supporting content online.

Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. Read the Citrix License Agreement.

5. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

6. Click Next.

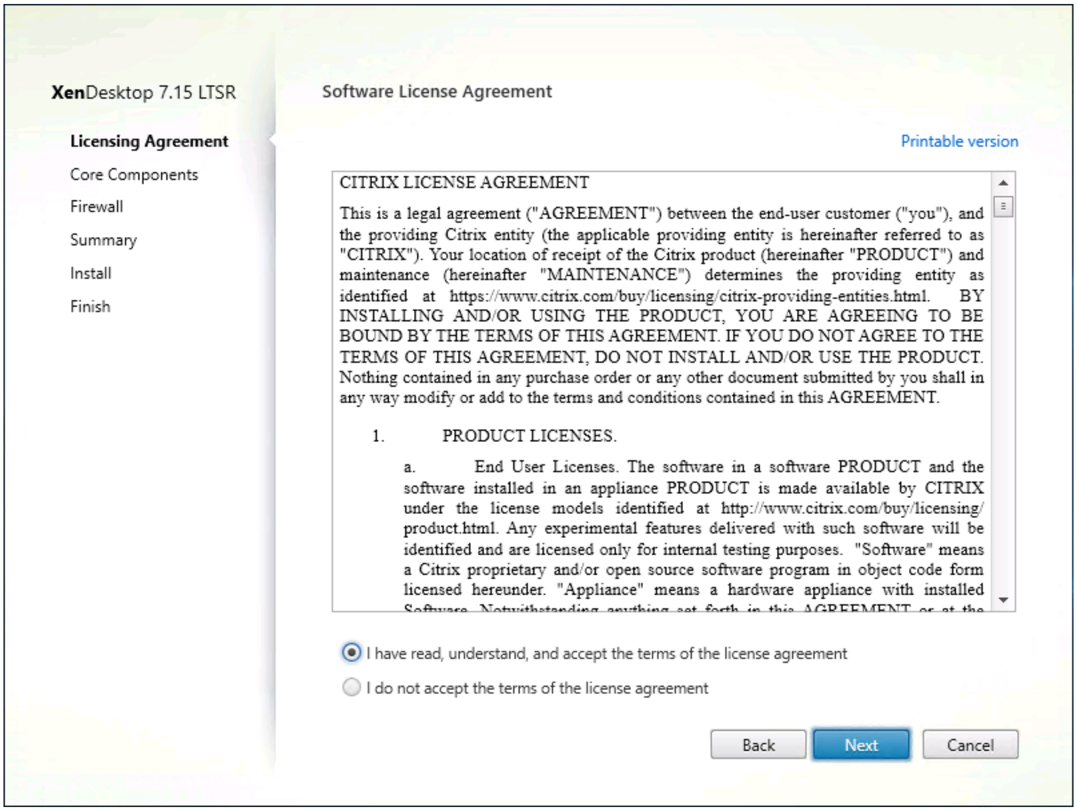

Software License Agreement

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY INSTALLING AND/OR USING THE PRODUCT, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT. Nothing contained in any purchase order or any other document submitted by you shall in any way modify or add to the terms and conditions contained in this AGREEMENT.

1. PRODUCT LICENSES.

a. End User Licenses. The software in a software PRODUCT and the software installed in an appliance PRODUCT is made available by CITRIX under the license models identified at http://www.citrix.com/buy/licensing/product.html. Any experimental features delivered with such software will be identified and are licensed only for internal testing purposes. "Software" means a Citrix proprietary and/or open source software program in object code form licensed hereunder. "Appliance" means a hardware appliance with installed Software. Notwithstanding anything set forth in this AGREEMENT or at the

● I have read, understand, and accept the terms of the license agreement

○ I do not accept the terms of the license agreement

7. Click Next.

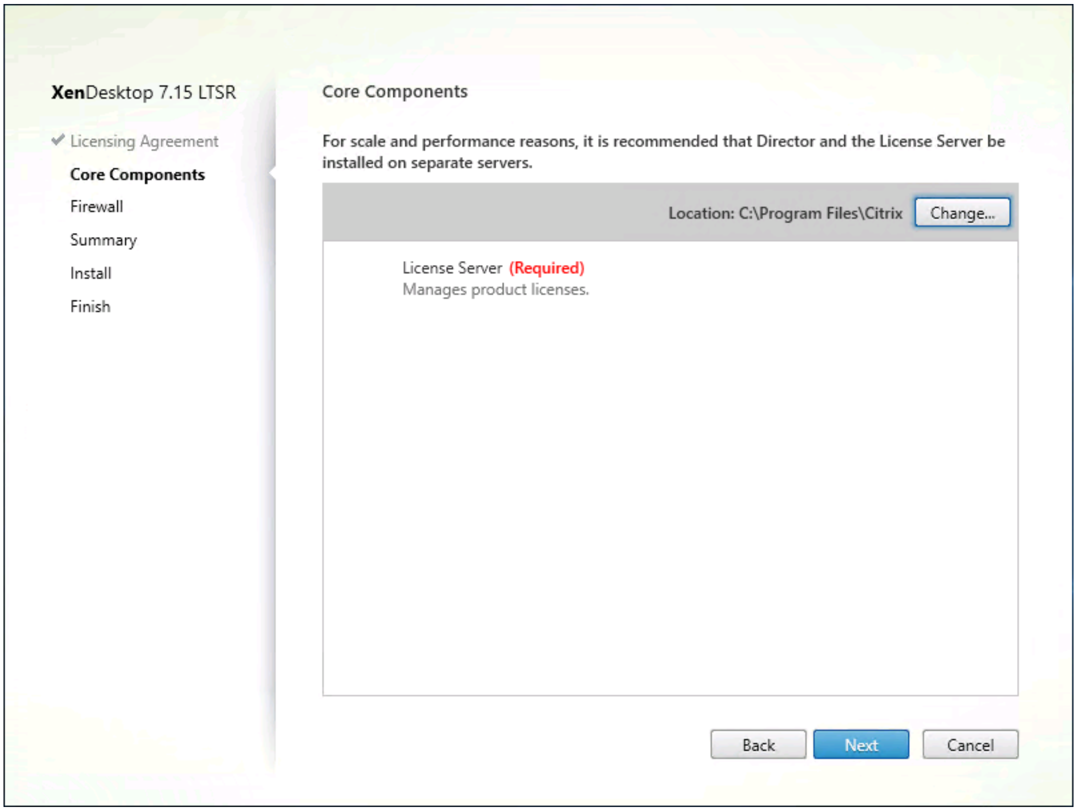

Core Components

For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.

Location: C:\Program Files\Citrix

License Server (Required)
Manages product licenses.

8. Select the default ports and automatically configured firewall rules.

9. Click Next.

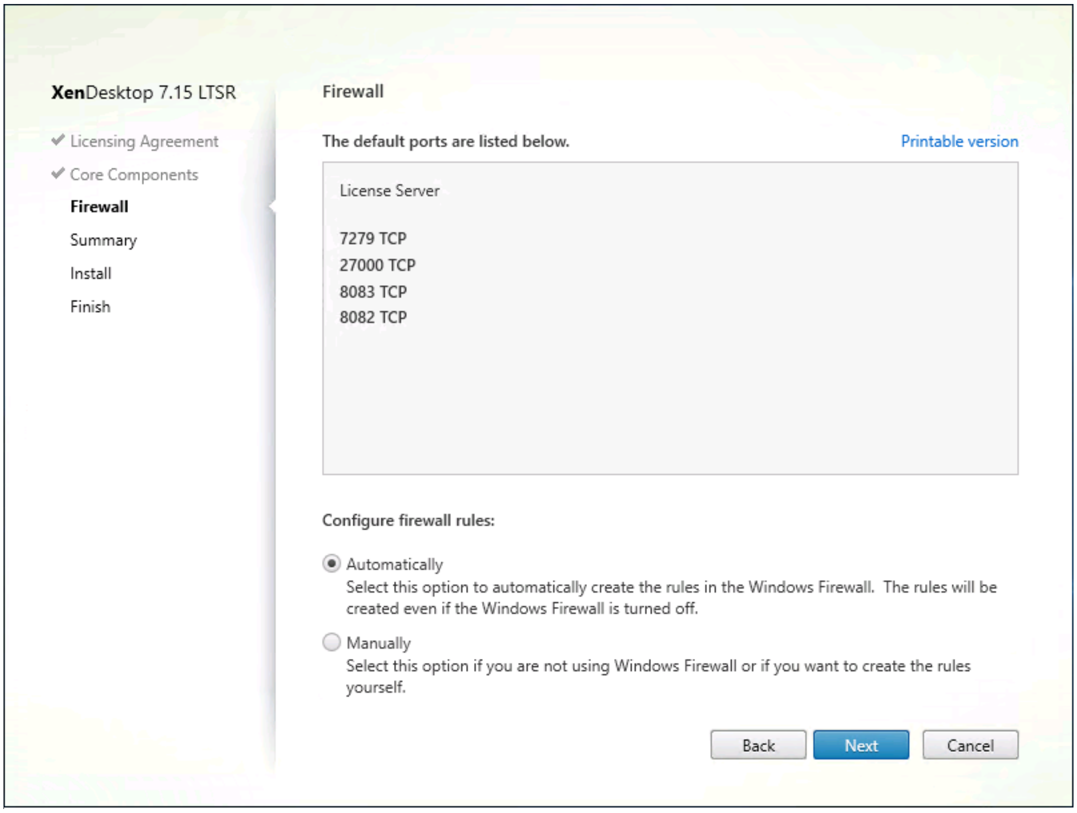

Firewall

The default ports are listed below.

License Server

7279 TCP
27000 TCP
8083 TCP
8082 TCP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

10. Click Install.

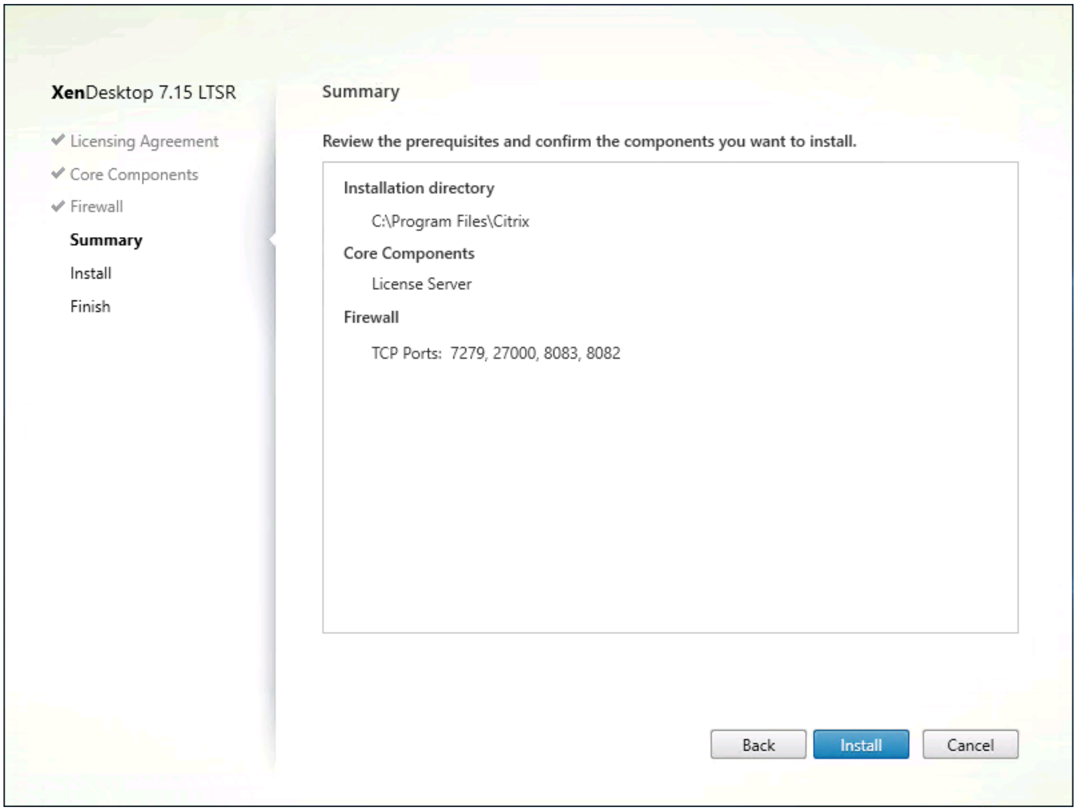

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory

C:\Program Files\Citrix

Core Components

License Server

Firewall

TCP Ports: 7279, 27000, 8083, 8082

11. Click Finish to complete installation.

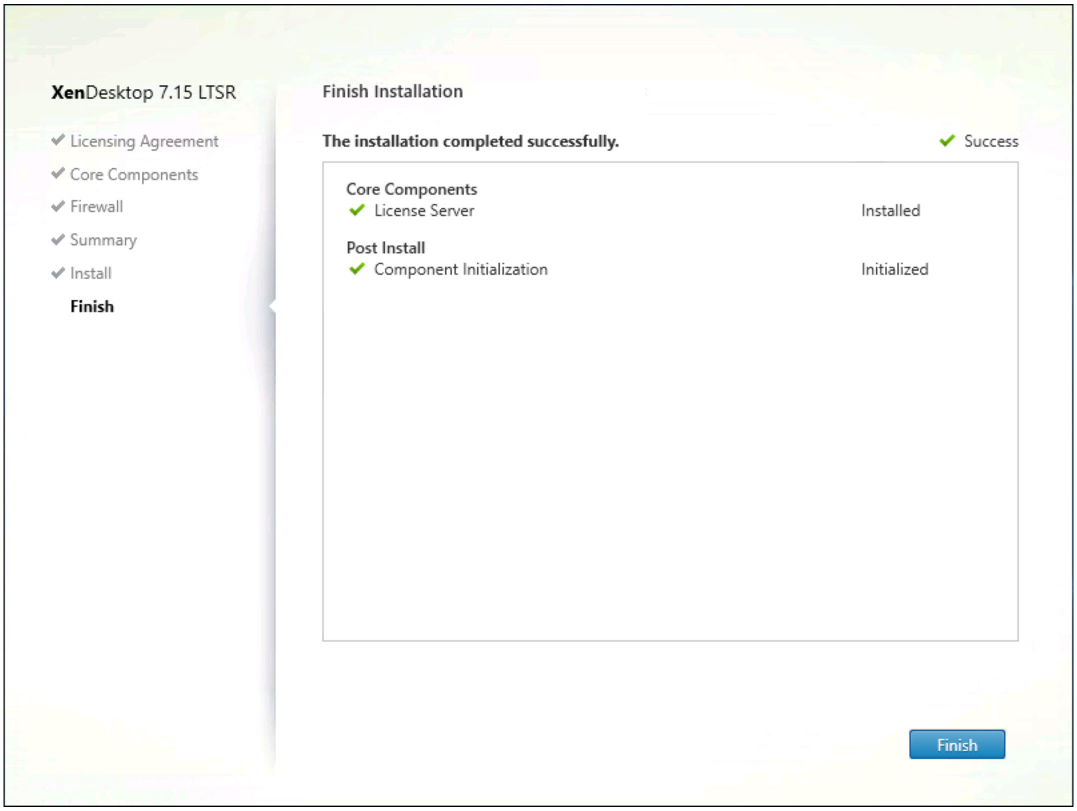

Finish Installation

The installation completed successfully. ✓ Success

Core Components	
✓ License Server	Installed
Post Install	
✓ Component Initialization	Initialized

Installing Citrix Licenses

To install the Citrix Licenses, complete the following steps:

1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\ MyFiles) on the license server.

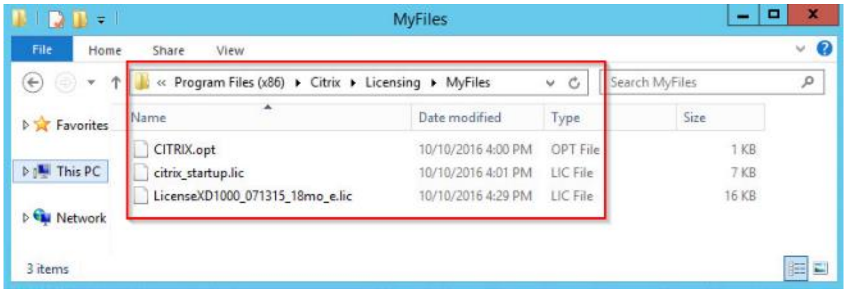

Name	Date modified	Type	Size
CITRIX.opt	10/10/2016 4:00 PM	OPT File	1 KB
citrix_startup.lic	10/10/2016 4:01 PM	LIC File	7 KB
LicenseXD1000_071315_18mo_e.lic	10/10/2016 4:29 PM	LIC File	16 KB

2. Restart the server or Citrix licensing services so that the licenses are activated.

3. Run the application Citrix License Administration Console.

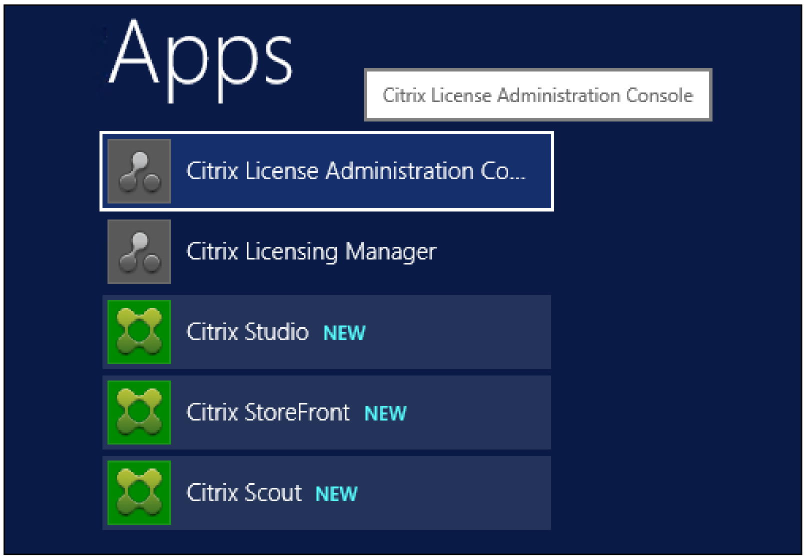

4. Confirm that the license files have been read and enabled correctly.

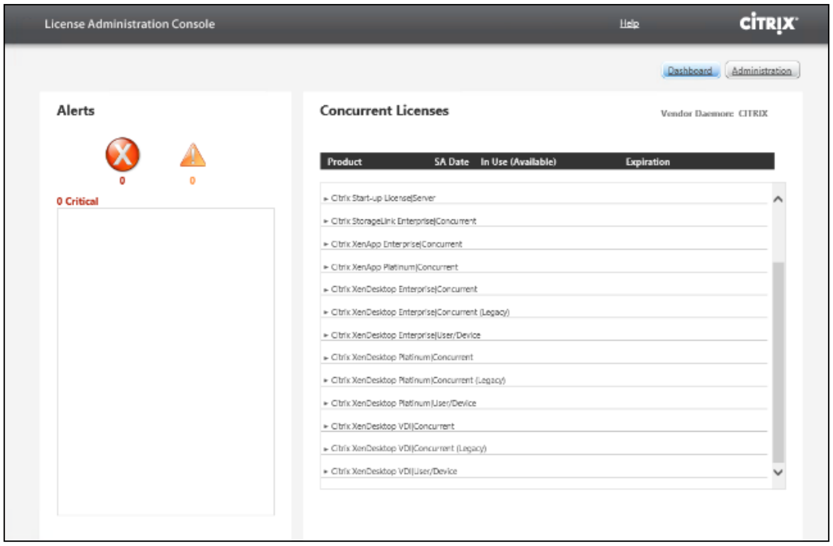

Alerts

0 0

0 Critical

Concurrent Licenses

Vendor Daemon: CITRIX

Product	SA Date	In Use (Available)	Expiration
▸ Citrix Start-up License|Server			
▸ Citrix StorageLink Enterprise|Concurrent			
▸ Citrix XenApp Enterprise|Concurrent			
▸ Citrix XenApp Platinum|Concurrent			
▸ Citrix XenDesktop Enterprise|Concurrent			
▸ Citrix XenDesktop Enterprise|Concurrent (Legacy)			
▸ Citrix XenDesktop Enterprise|User/Device			
▸ Citrix XenDesktop Platinum|Concurrent			
▸ Citrix XenDesktop Platinum|Concurrent (Legacy)			
▸ Citrix XenDesktop Platinum|User/Device			
▸ Citrix XenDesktop VDI|Concurrent			
▸ Citrix XenDesktop VDI|Concurrent (Legacy)			
▸ Citrix XenDesktop VDI|User/Device			

Install the XenDesktop

1. To begin the installation, connect to the first XenDesktop server and launch the installer from the Citrix XenDesktop 7.15ISO.

2. Click Start.

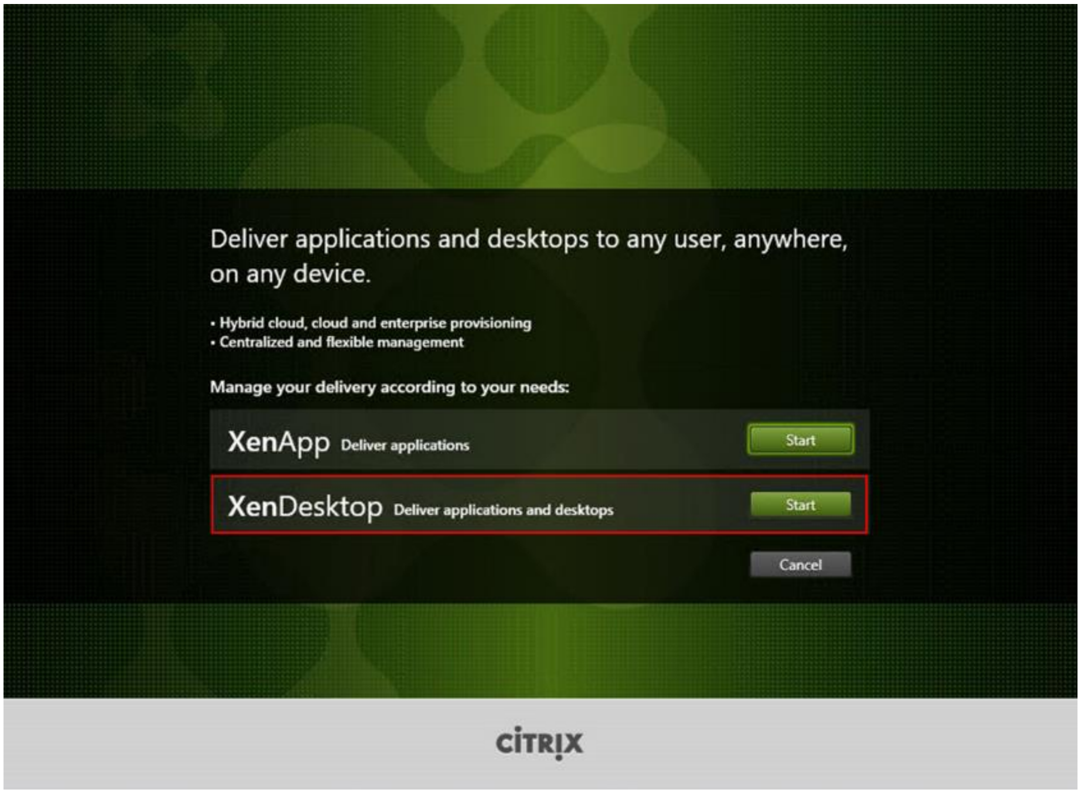

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications

XenDesktop Deliver applications and desktops

The installation wizard presents a menu with three subsections.

3. Click “Get Started - Delivery Controller.”

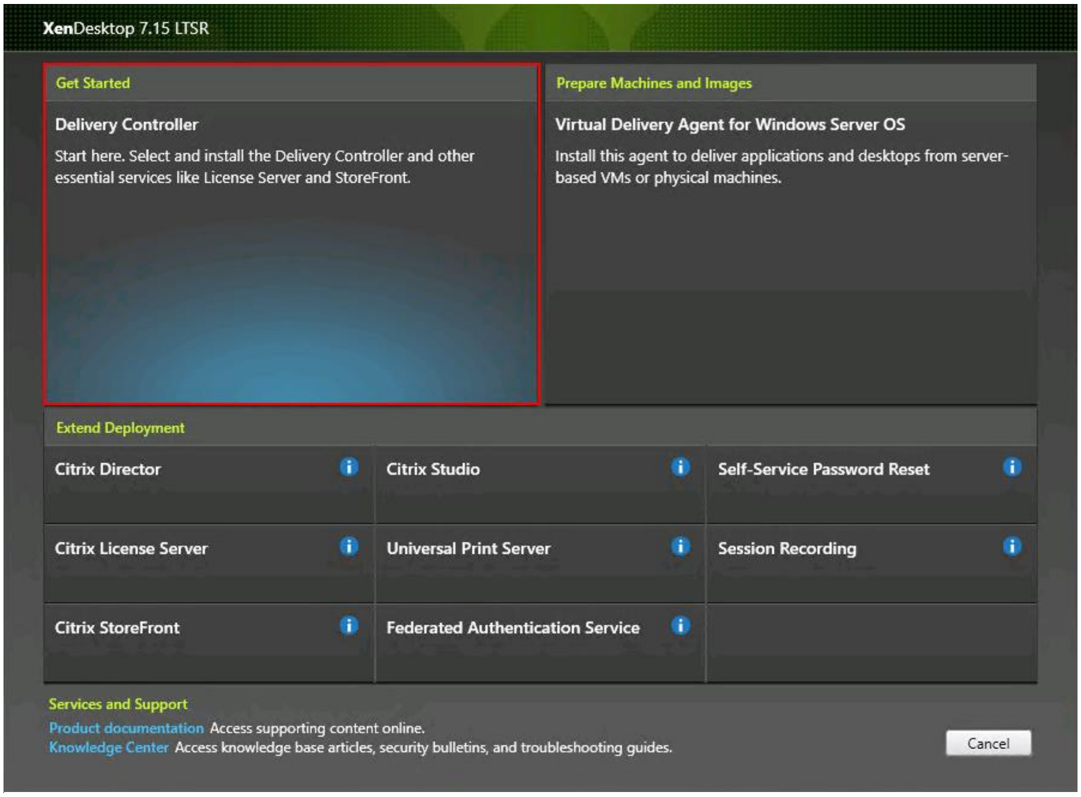

XenDesktop 7.15 LTSR

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Extend Deployment

Citrix Director	Citrix Studio	Self-Service Password Reset
Citrix License Server	Universal Print Server	Session Recording
Citrix StoreFront	Federated Authentication Service	

Services and Support

Product documentation Access supporting content online.

Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. Read the Citrix License Agreement.

5. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

6. Click Next.

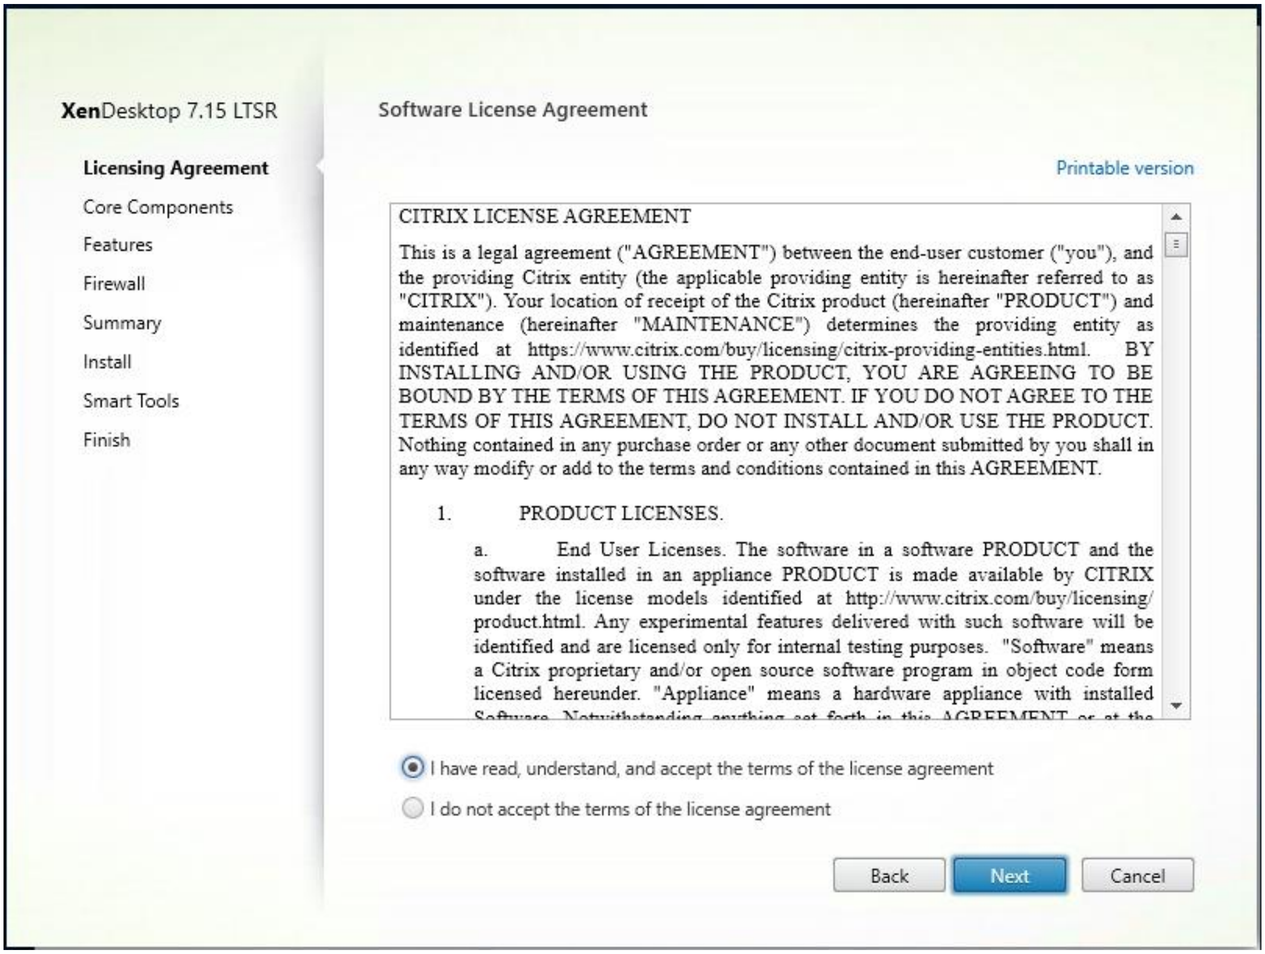

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY INSTALLING AND/OR USING THE PRODUCT, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT. Nothing contained in any purchase order or any other document submitted by you shall in any way modify or add to the terms and conditions contained in this AGREEMENT.

1. PRODUCT LICENSES.

a. End User Licenses. The software in a software PRODUCT and the software installed in an appliance PRODUCT is made available by CITRIX under the license models identified at http://www.citrix.com/buy/licensing/product.html. Any experimental features delivered with such software will be identified and are licensed only for internal testing purposes. "Software" means a Citrix proprietary and/or open source software program in object code form licensed hereunder. "Appliance" means a hardware appliance with installed Software. Notwithstanding anything set forth in this AGREEMENT or at the

● I have read, understand, and accept the terms of the license agreement

○ I do not accept the terms of the license agreement

7. Select the components to be installed on the first Delivery Controller Server:

   - Delivery Controller
   - Studio
   - Director

8. Click Next.

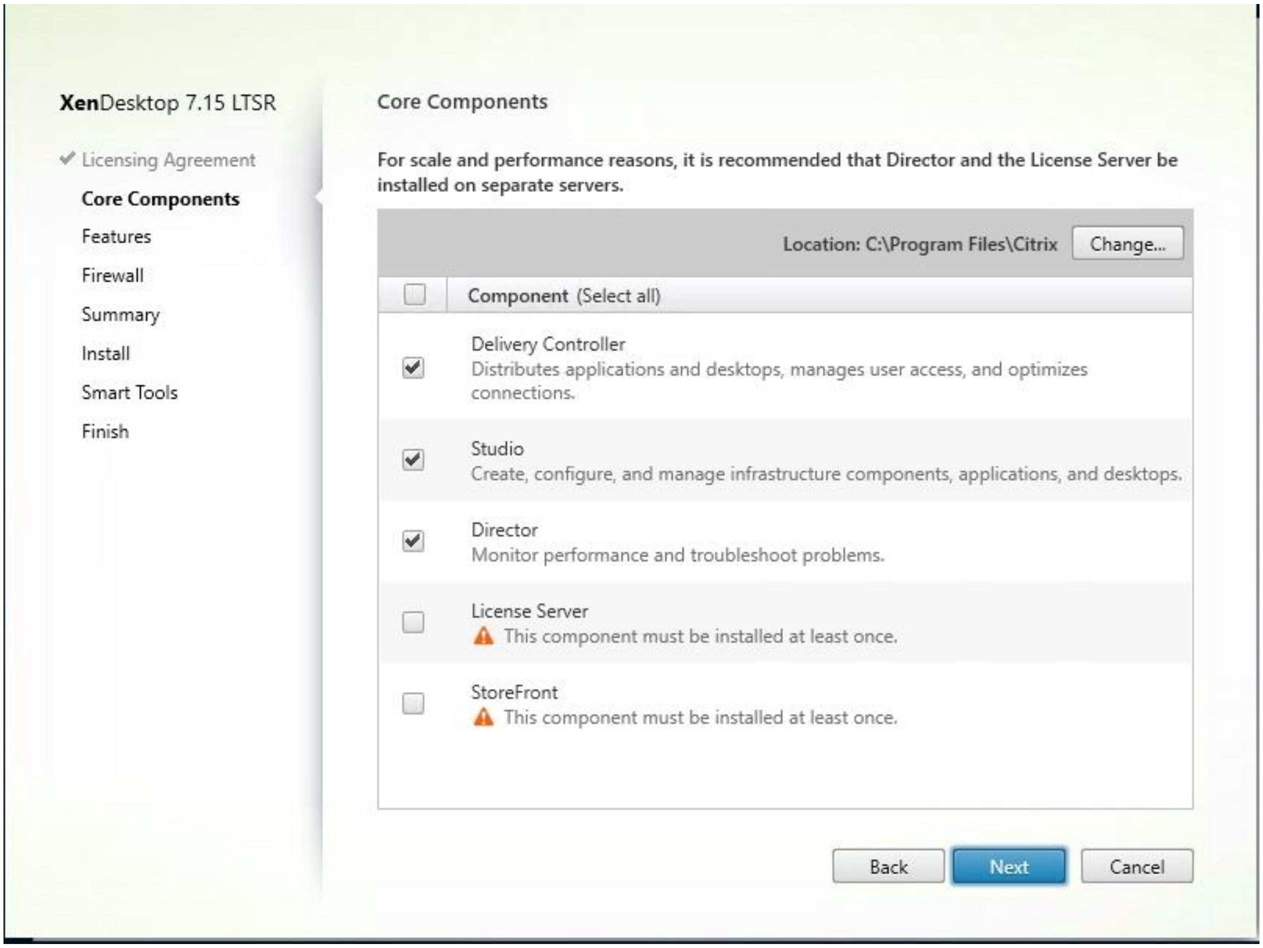

XenDesktop 7.15 LTSR

✓ Licensing Agreement
Core Components
Features
Firewall
Summary
Install
Smart Tools
Finish

Core Components

For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.

Location: C:\Program Files\Citrix

□	Component (Select all)
☑	Delivery Controller Distributes applications and desktops, manages user access, and optimizes connections.
☑	Studio Create, configure, and manage infrastructure components, applications, and desktops.
☑	Director Monitor performance and troubleshoot problems.
□	License Server ⚠ This component must be installed at least once.
□	StoreFront ⚠ This component must be installed at least once.

9. Since a dedicated SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2012 SP1 Express” unchecked.

10. Click Next.

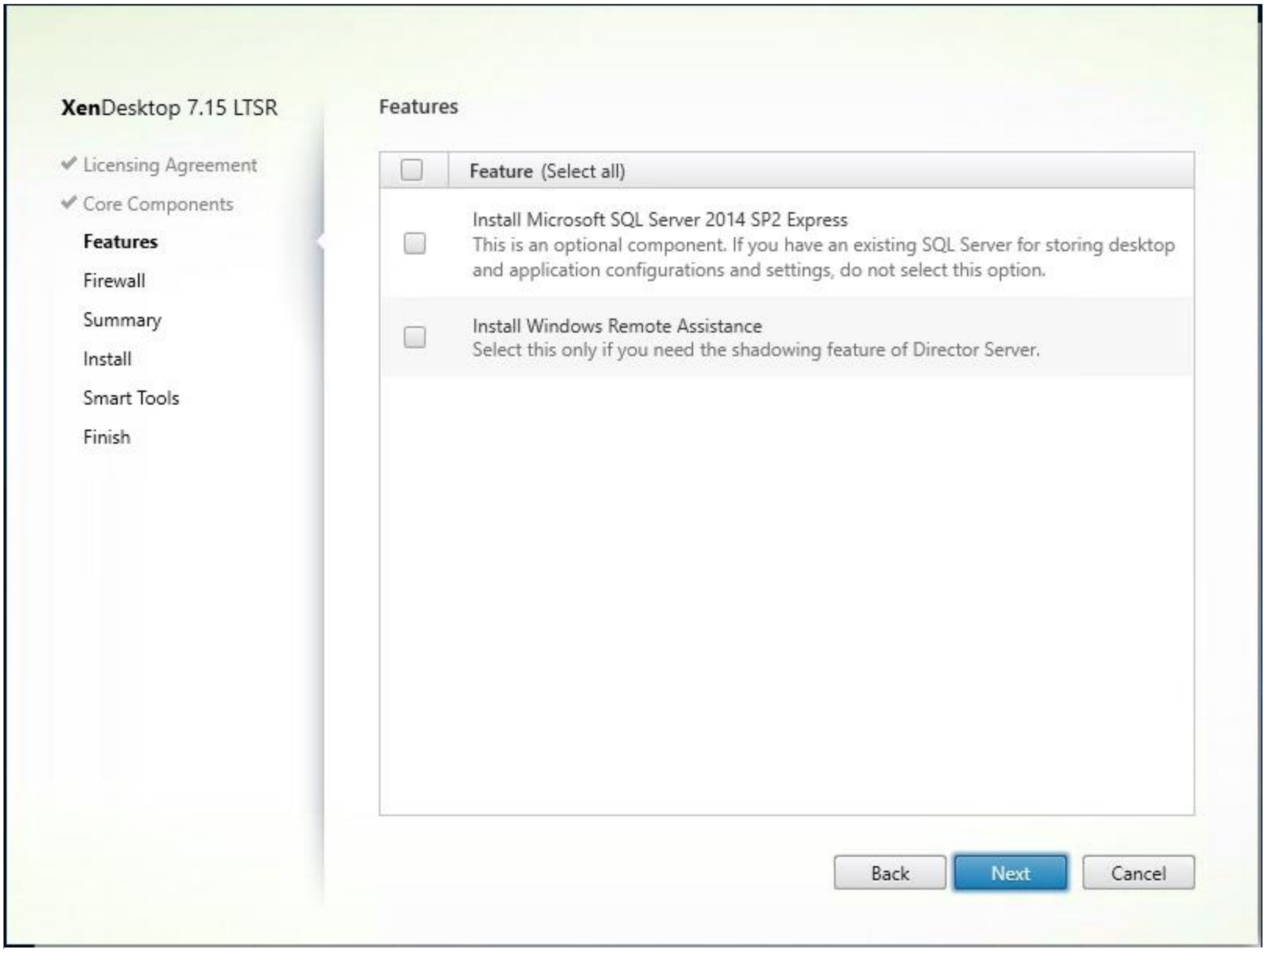

XenDesktop 7.15 LTSR

Features

□	Feature (Select all)
□	Install Microsoft SQL Server 2014 SP2 Express This is an optional component. If you have an existing SQL Server for storing desktop and application configurations and settings, do not select this option.
□	Install Windows Remote Assistance Select this only if you need the shadowing feature of Director Server.

11. Select the default ports and automatically configured firewall rules.

12. Click Next.

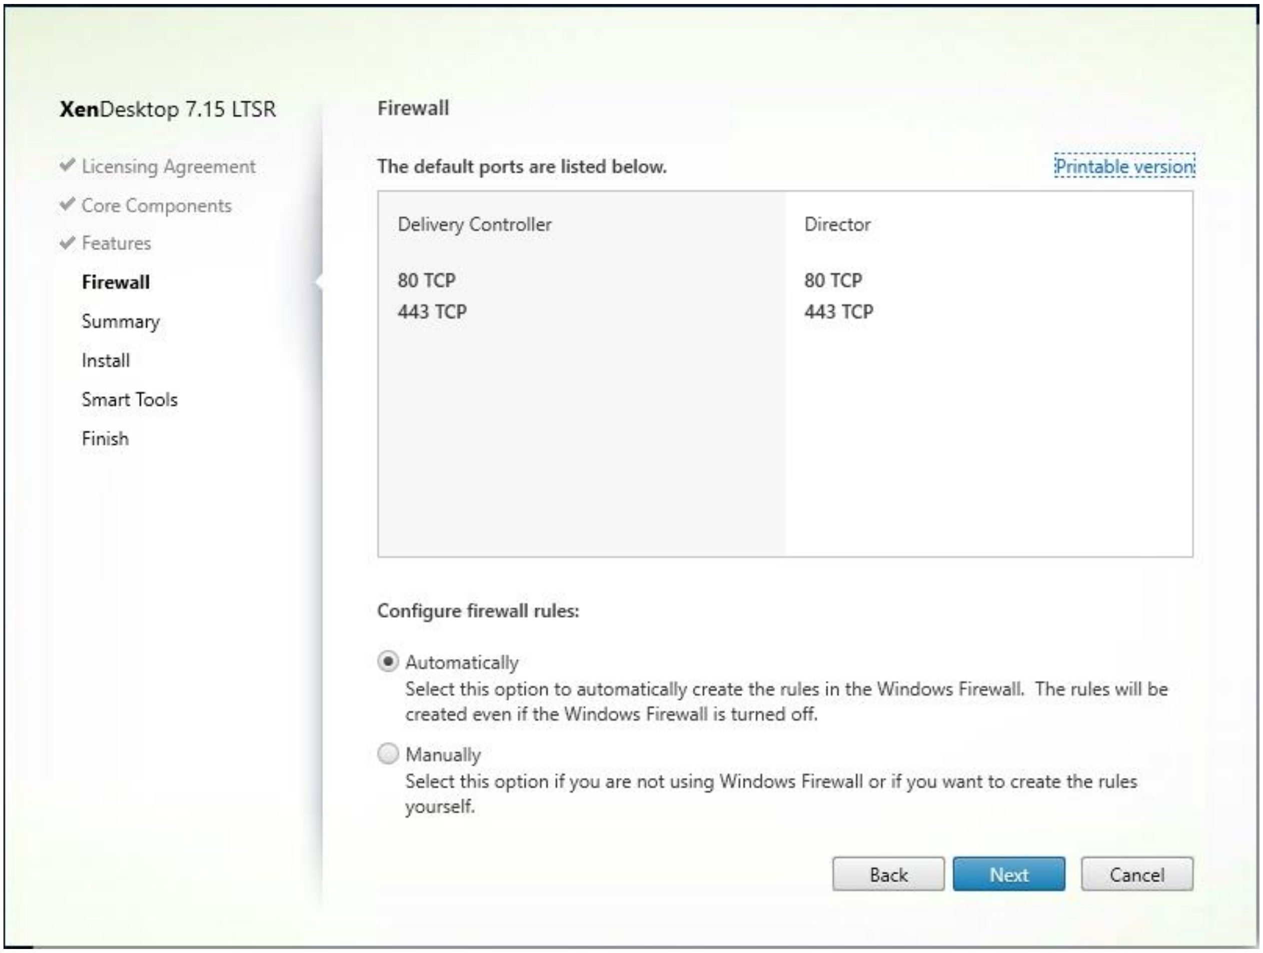

Firewall

The default ports are listed below.

Delivery Controller	Director
80 TCP 443 TCP	80 TCP 443 TCP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

13. Click Install to begin the installation.

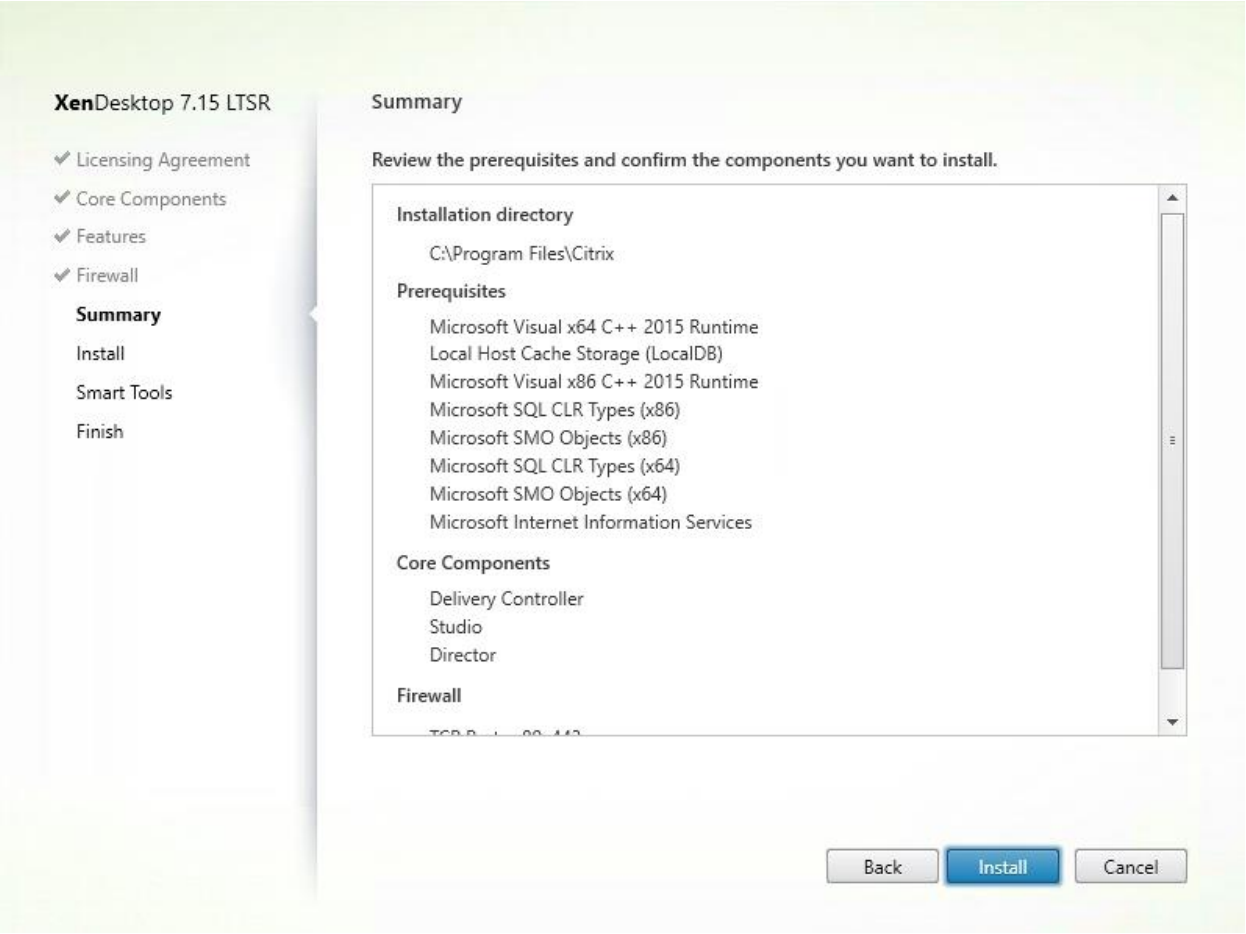

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory

C:\Program Files\Citrix

Prerequisites

Microsoft Visual x64 C++ 2015 Runtime
Local Host Cache Storage (LocalDB)
Microsoft Visual x86 C++ 2015 Runtime
Microsoft SQL CLR Types (x86)
Microsoft SMO Objects (x86)
Microsoft SQL CLR Types (x64)
Microsoft SMO Objects (x64)
Microsoft Internet Information Services

Core Components

Delivery Controller
Studio
Director

Firewall

14. (Optional) Configure Smart Tools/Call Home participation.

15. Click Next.

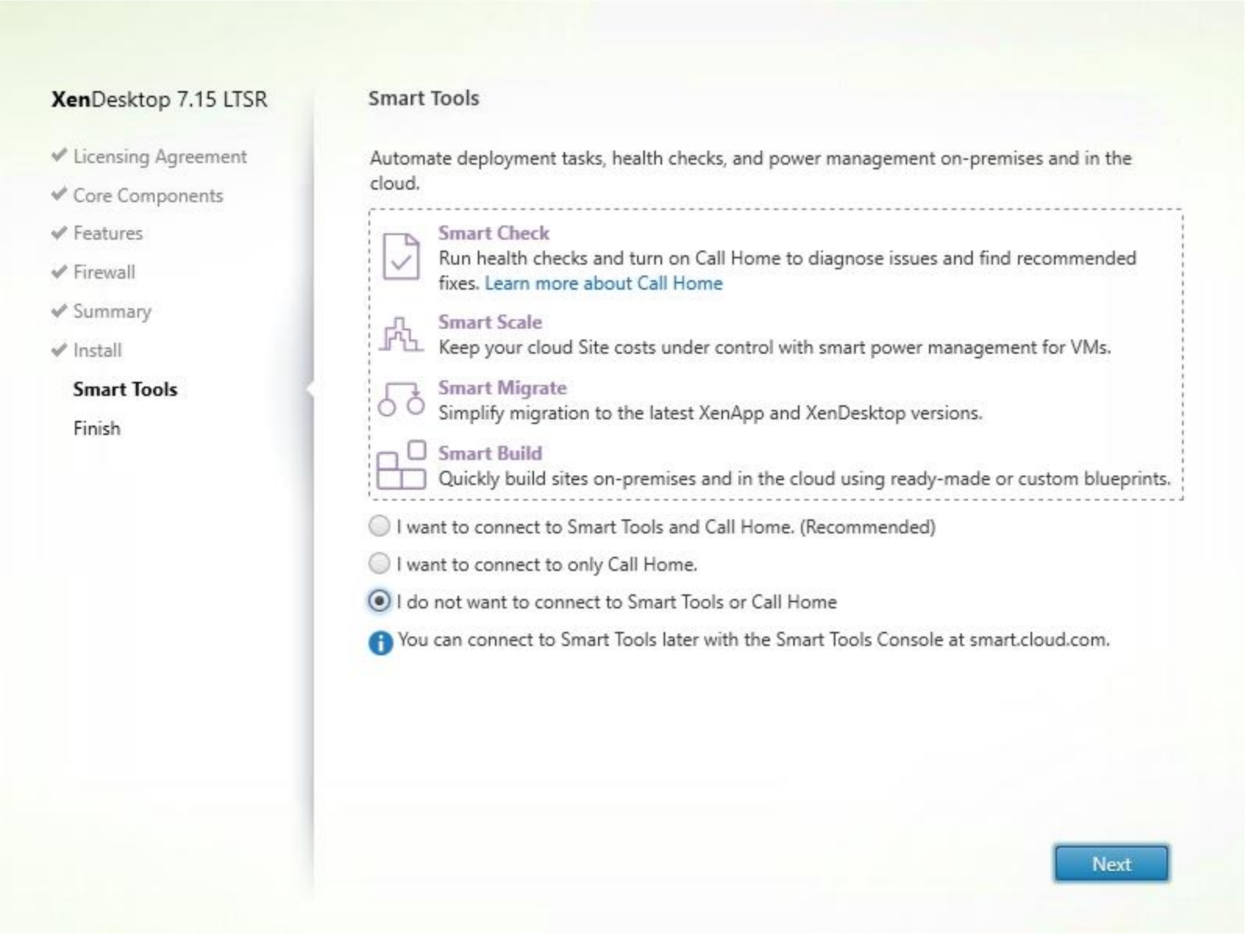

Smart Tools

Automate deployment tasks, health checks, and power management on-premises and in the cloud.

Smart Check Run health checks and turn on Call Home to diagnose issues and find recommended fixes. Learn more about Call Home

Smart Scale Keep your cloud Site costs under control with smart power management for VMs.

Smart Migrate Simplify migration to the latest XenApp and XenDesktop versions.

Smart Build Quickly build sites on-premises and in the cloud using ready-made or custom blueprints.

○ I want to connect to Smart Tools and Call Home. (Recommended)

○ I want to connect to only Call Home.

● I do not want to connect to Smart Tools or Call Home

You can connect to Smart Tools later with the Smart Tools Console at smart.cloud.com.

16. Click Finish to complete the installation.

17. (Optional) Check Launch Studio to launch Citrix Studio Console.

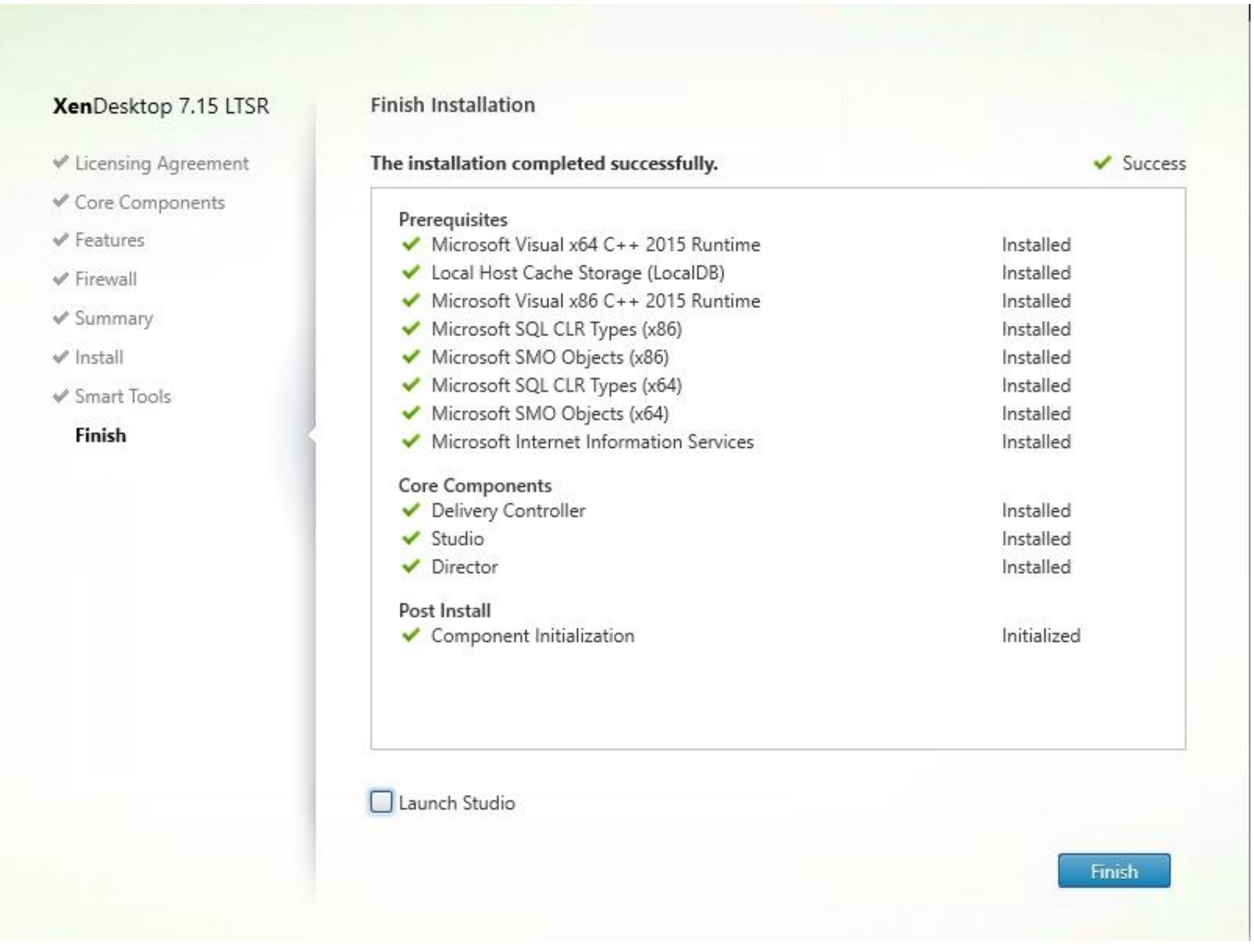

XenDesktop 7.15 LTSR

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites

✓ Microsoft Visual x64 C++ 2015 Runtime	Installed
✓ Local Host Cache Storage (LocalDB)	Installed
✓ Microsoft Visual x86 C++ 2015 Runtime	Installed
✓ Microsoft SQL CLR Types (x86)	Installed
✓ Microsoft SMO Objects (x86)	Installed
✓ Microsoft SQL CLR Types (x64)	Installed
✓ Microsoft SMO Objects (x64)	Installed
✓ Microsoft Internet Information Services	Installed

Core Components

✓ Delivery Controller	Installed
✓ Studio	Installed
✓ Director	Installed

Post Install

✓ Component Initialization	Initialized

☐ Launch Studio

Additional XenDesktop Controller Configuration

After the first controller is completely configured and the Site is operational, you can add additional controllers. In this CVD, we created two Delivery Controllers.

To configure additional XenDesktop controllers, complete the following steps:

1. To begin the installation of the second Delivery Controller, connect to the second XenDesktop server and launch the installer from the Citrix XenDesktop 7.15ISO.

2. Click Start.

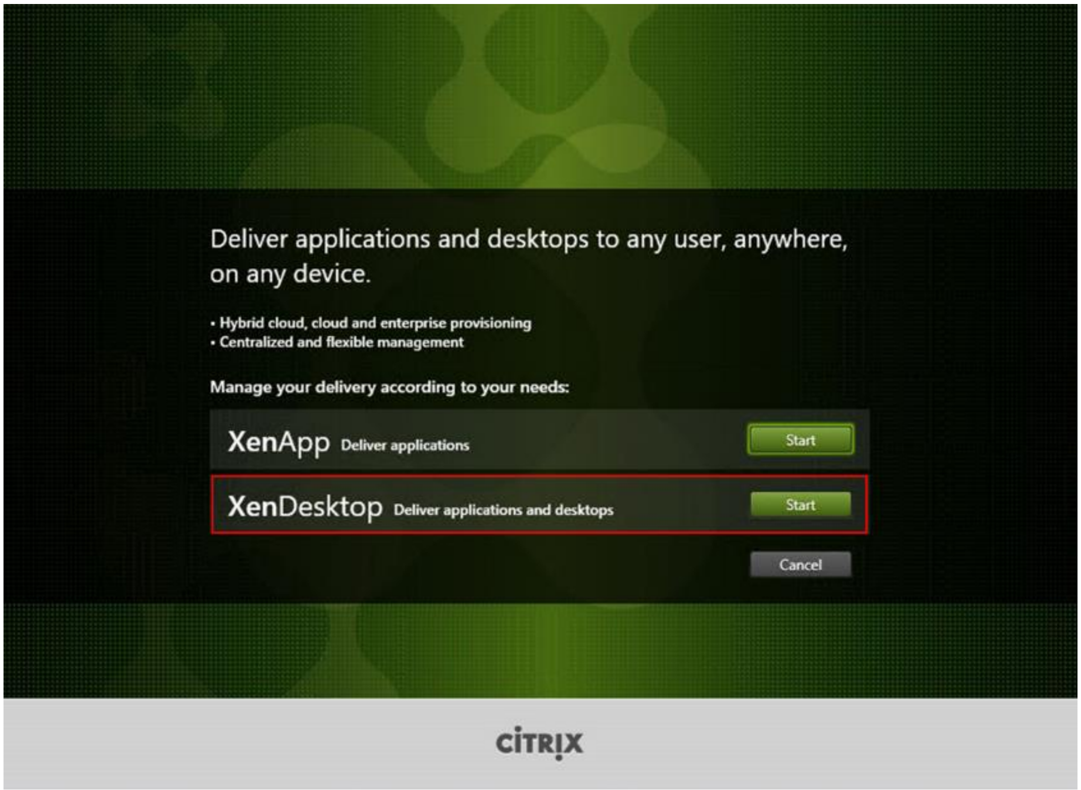

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications
Start
XenDesktop Deliver applications and desktops
Start
Cancel
CiTRIX

3. Click Delivery Controller.

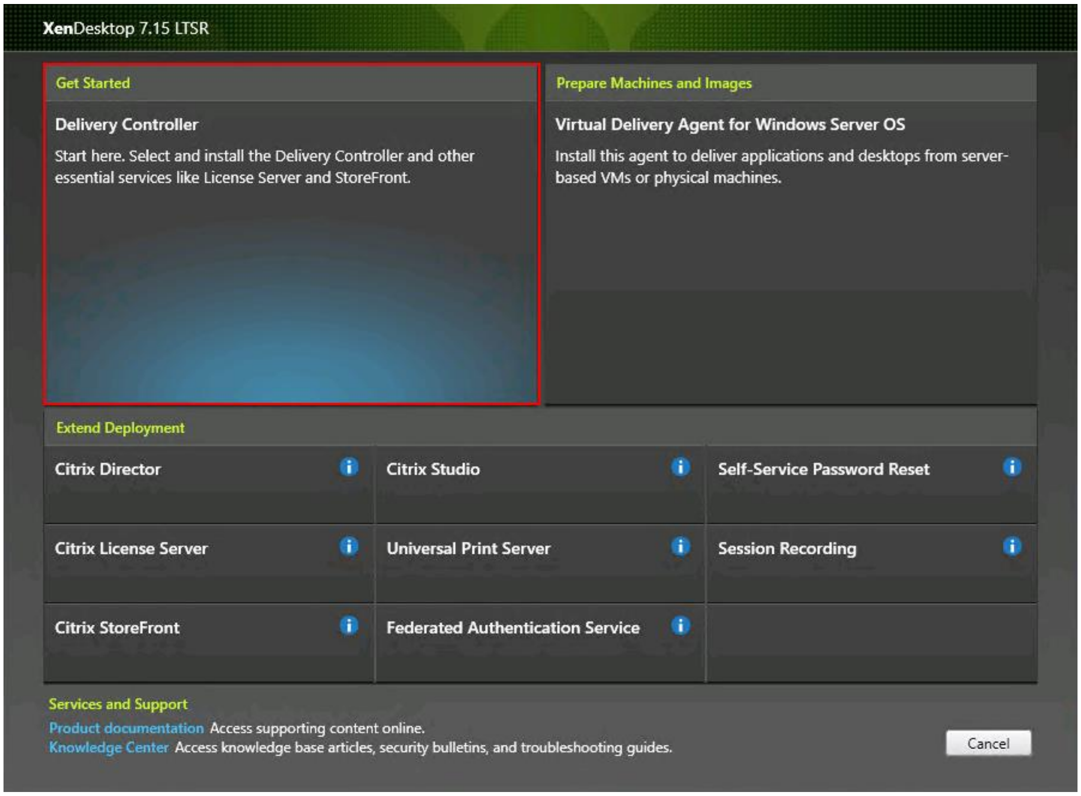

XenDesktop 7.15 LTSR

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Extend Deployment

Services and Support

Product documentation Access supporting content online.

Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. Repeat the same steps used to install the first Delivery Controller, including the step of importing an SSL certificate for HTTPS between the controller and vSphere.

5. Review the Summary configuration.

6. Click Install.

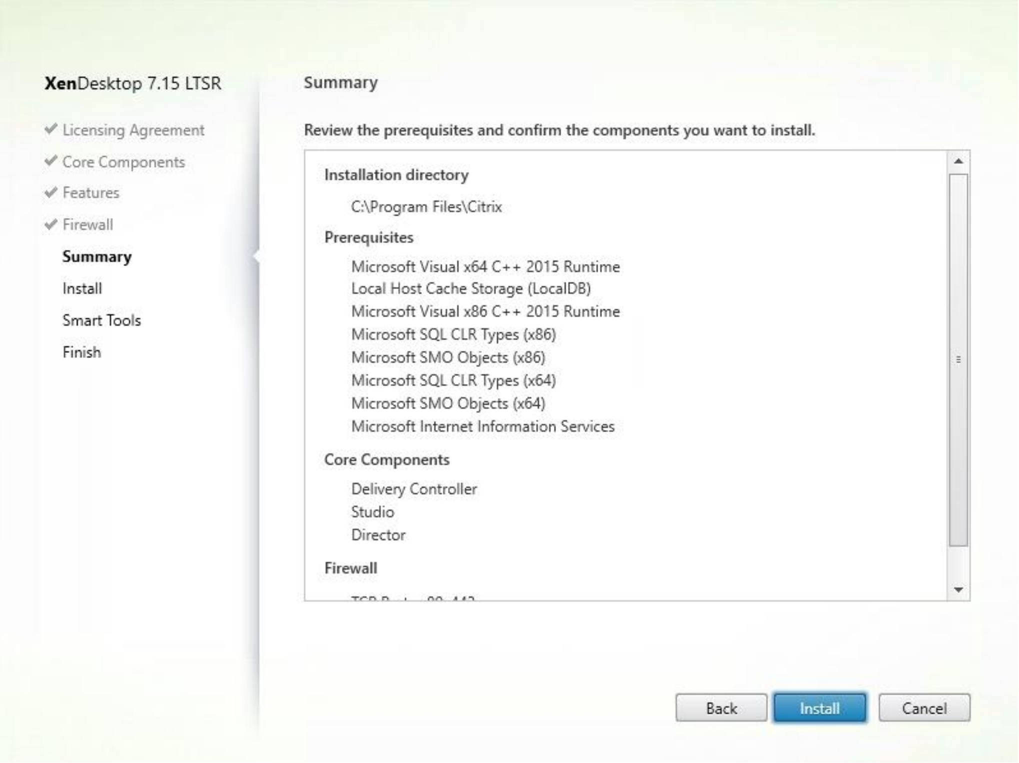

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory

C:\Program Files\Citrix

Prerequisites

Microsoft Visual x64 C++ 2015 Runtime
Local Host Cache Storage (LocalDB)
Microsoft Visual x86 C++ 2015 Runtime
Microsoft SQL CLR Types (x86)
Microsoft SMO Objects (x86)
Microsoft SQL CLR Types (x64)
Microsoft SMO Objects (x64)
Microsoft Internet Information Services

Core Components

Delivery Controller
Studio
Director

Firewall

TCP Ports: 80, 443

7. (Optional) Configure Smart Tools/Call Home participation.

8. Click Next.

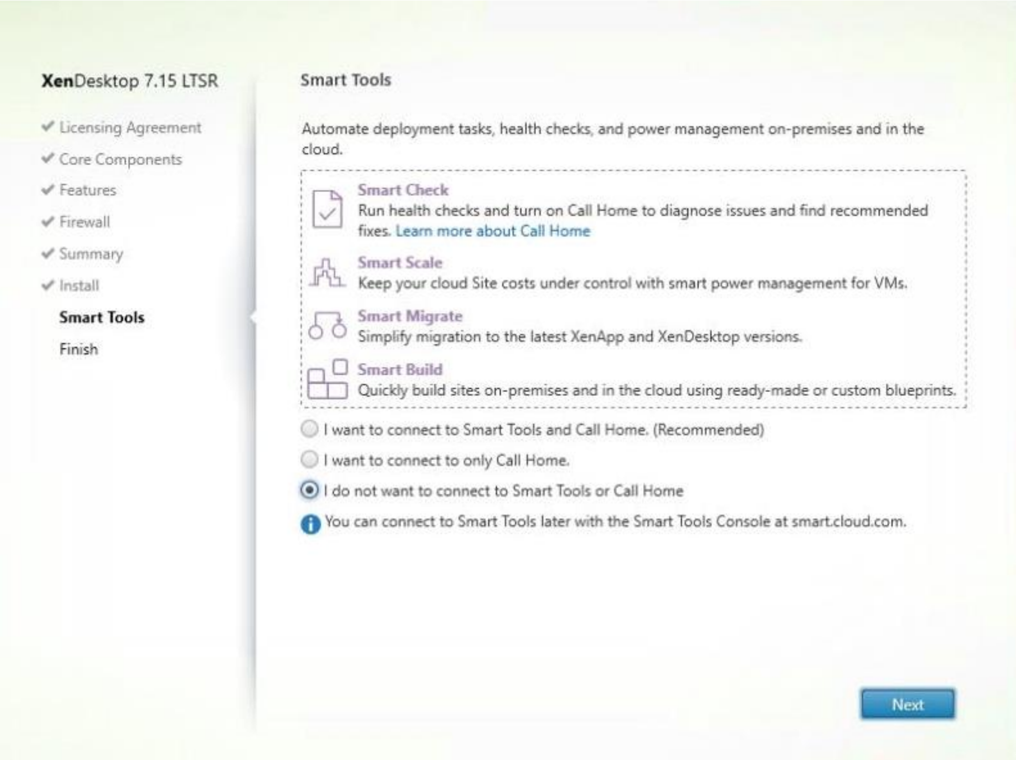

Smart Tools

Automate deployment tasks, health checks, and power management on-premises and in the cloud.

Smart Check
Run health checks and turn on Call Home to diagnose issues and find recommended fixes. Learn more about Call Home

Smart Scale
Keep your cloud Site costs under control with smart power management for VMs.

Smart Migrate
Simplify migration to the latest XenApp and XenDesktop versions.

Smart Build
Quickly build sites on-premises and in the cloud using ready-made or custom blueprints.

○ I want to connect to Smart Tools and Call Home. (Recommended)

○ I want to connect to only Call Home.

● I do not want to connect to Smart Tools or Call Home

You can connect to Smart Tools later with the Smart Tools Console at smart.cloud.com.

9. Verify the components installed successfully.

10. Click Finish.

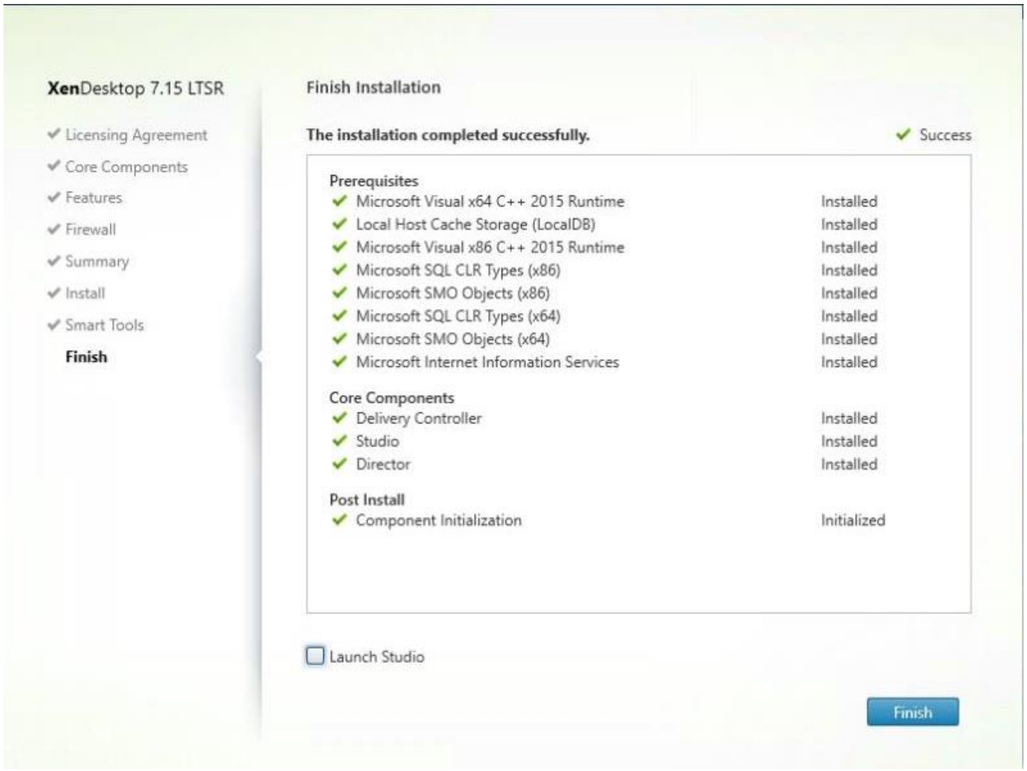

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites

✓ Microsoft Visual x64 C++ 2015 Runtime	Installed
✓ Local Host Cache Storage (LocalDB)	Installed
✓ Microsoft Visual x86 C++ 2015 Runtime	Installed
✓ Microsoft SQL CLR Types (x86)	Installed
✓ Microsoft SMO Objects (x86)	Installed
✓ Microsoft SQL CLR Types (x64)	Installed
✓ Microsoft SMO Objects (x64)	Installed
✓ Microsoft Internet Information Services	Installed

Core Components

✓ Delivery Controller	Installed
✓ Studio	Installed
✓ Director	Installed

Post Install

✓ Component Initialization	Initialized

□ Launch Studio

Configure the XenDesktop Site

Citrix Studio is a management console that allows you to create and manage infrastructure and resources to deliver desktops and applications. Replacing Desktop Studio from earlier releases, it provides wizards to set up your environment, create workloads to host applications and desktops, and assign applications and desktops to users.

Citrix Studio launches automatically after the XenDesktop Delivery Controller installation, or if necessary, it can be launched manually. Studio is used to create a Site, which is the core XenDesktop 7.15environment consisting of the Delivery Controller and the Database.

To configure XenDesktop, complete the following steps:

1. From Citrix Studio, click the Deliver applications and desktops to your users button.

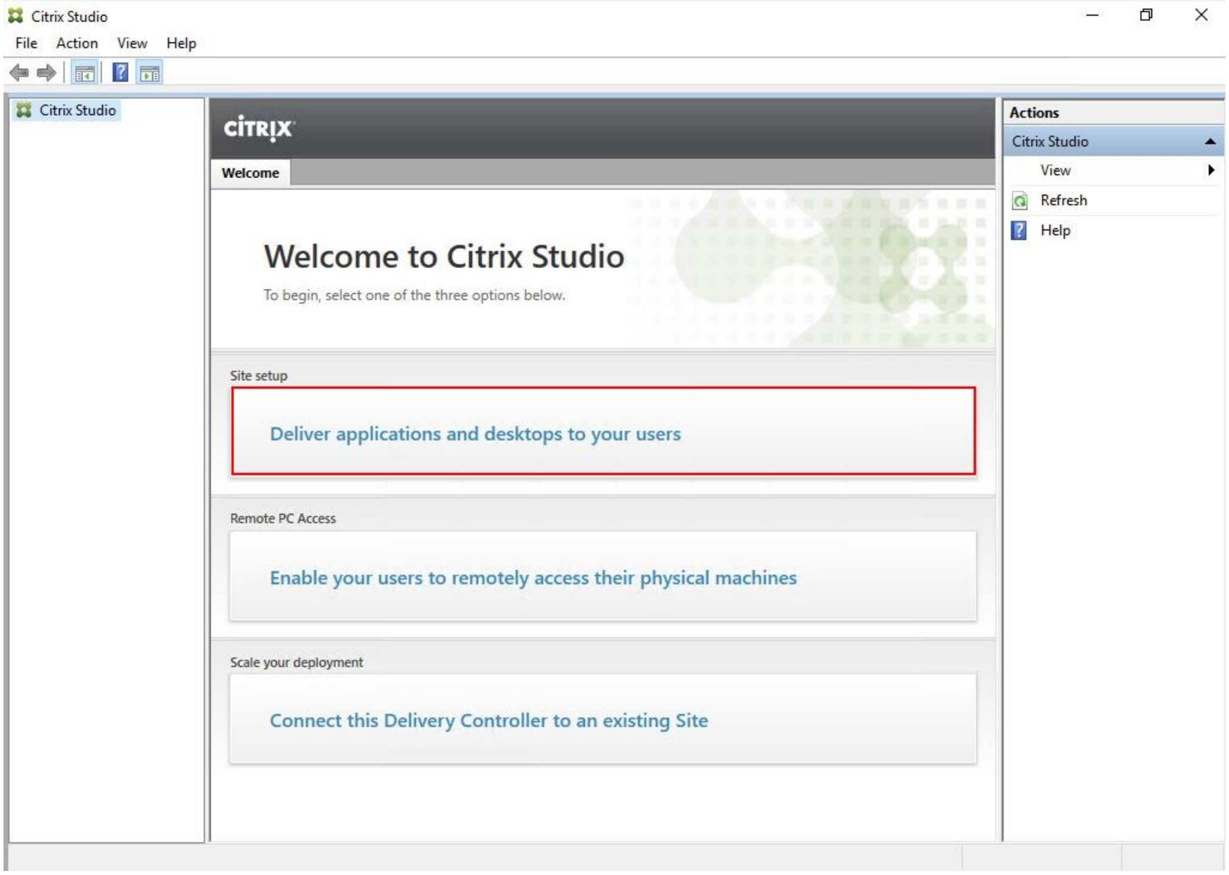
Citrix Studio
File
Action
View
Help
Citrix Studio
CITRIX
Welcome
Welcome to Citrix Studio
To begin, select one of the three options below.
Site setup
Deliver applications and desktops to your users
Remote PC Access
Enable your users to remotely access their physical machines
Scale your deployment
Connect this Delivery Controller to an existing Site
Actions
Citrix Studio
View
Refresh
Help

2. Select the “An empty, unconfigured Site” radio button.

3. Enter a site name.

4. Click Next.

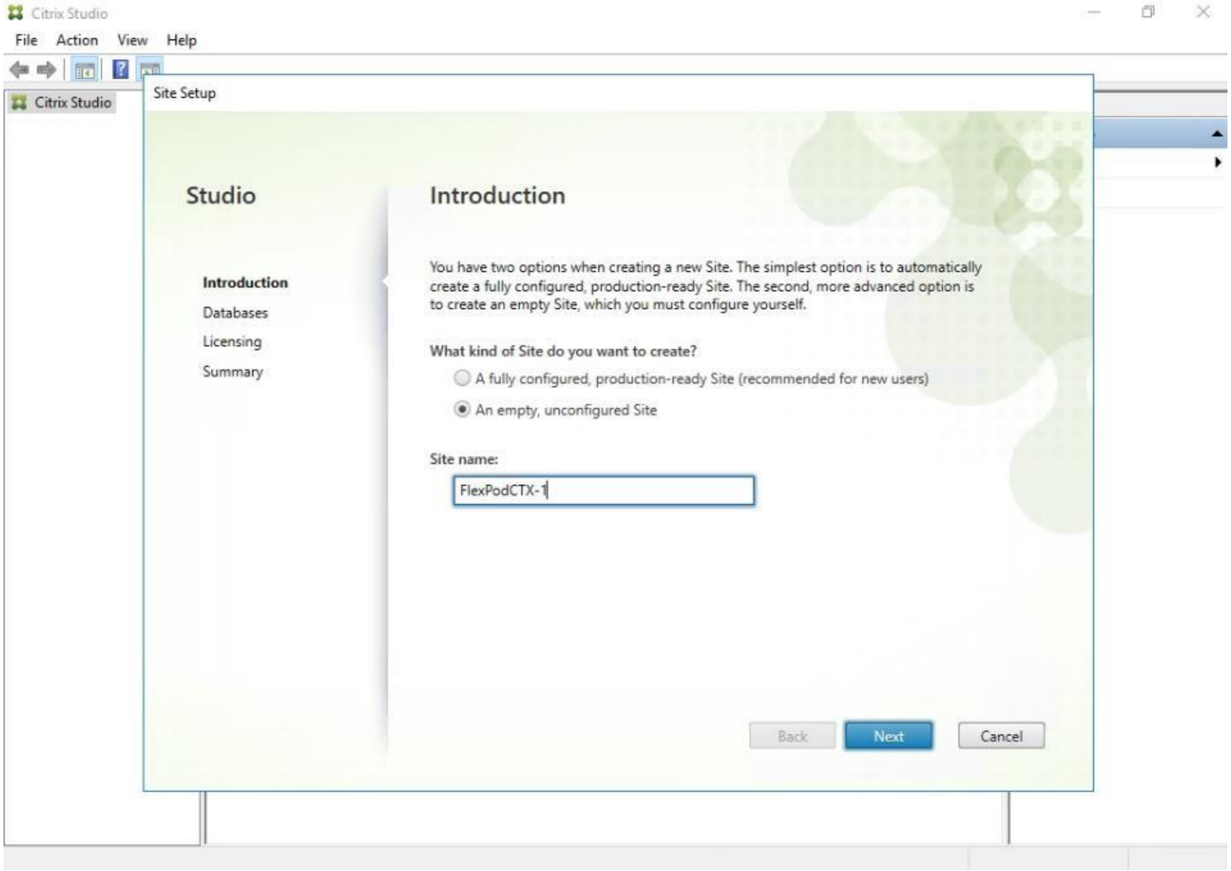

Introduction

You have two options when creating a new Site. The simplest option is to automatically create a fully configured, production-ready Site. The second, more advanced option is to create an empty Site, which you must configure yourself.

What kind of Site do you want to create?

○ A fully configured, production-ready Site (recommended for new users)

● An empty, unconfigured Site

Site name: FlexPodCTX-1

5. Provide the Database Server Locations for each data type and click Next.

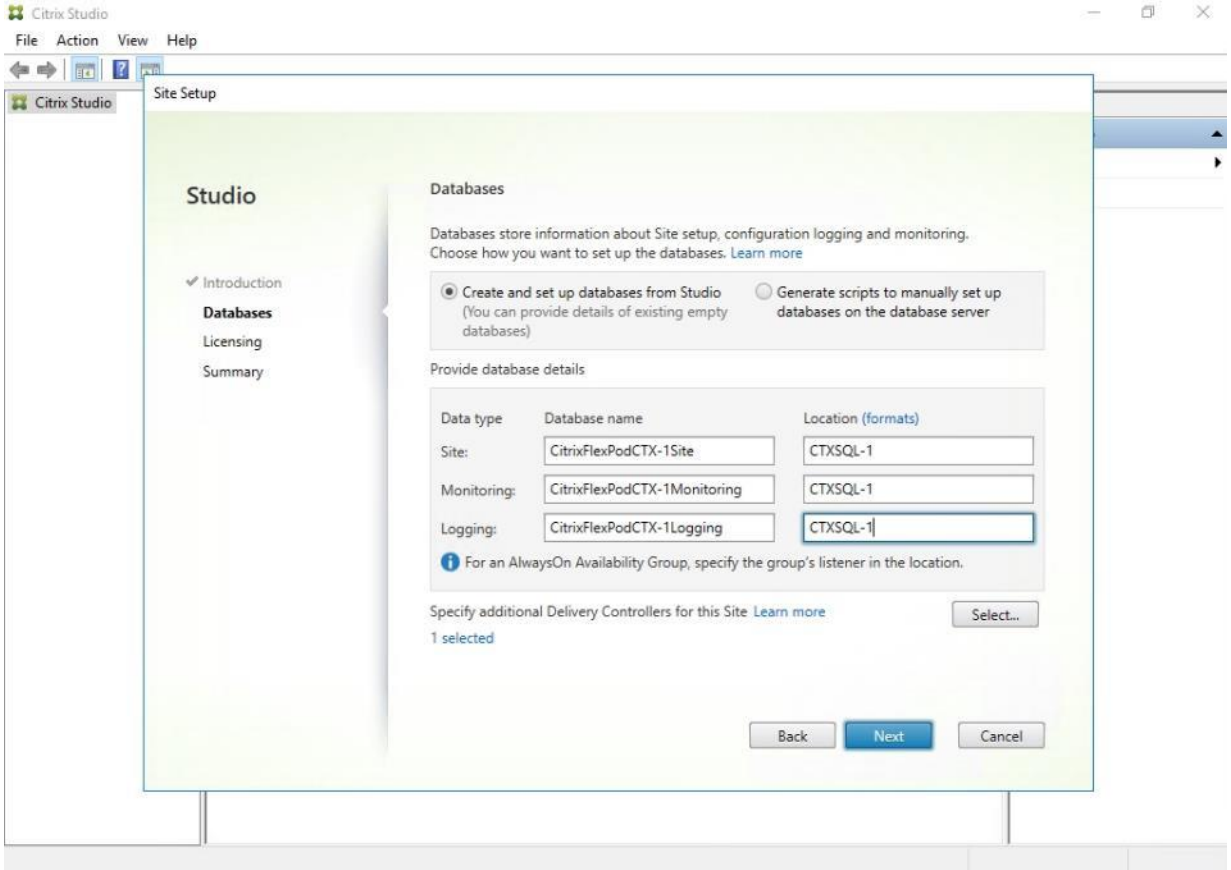

Databases

Databases store information about Site setup, configuration logging and monitoring. Choose how you want to set up the databases. Learn more

● Create and set up databases from Studio (You can provide details of existing empty databases)

○ Generate scripts to manually set up databases on the database server

Provide database details

Data type	Database name	Location (formats)
Site:	CitrixFlexPodCTX-1Site	CTXSQL-1
Monitoring:	CitrixFlexPodCTX-1Monitoring	CTXSQL-1
Logging:	CitrixFlexPodCTX-1Logging	CTXSQL-1

For an AlwaysOn Availability Group, specify the group's listener in the location.

Specify additional Delivery Controllers for this Site Learn more

1 selected

For an AlwaysOn Availability Group, use the group’s listener DNS name.

6. Click Select to specify additional controllers.

7. Click Add, specify controller FQDN, and click OK.

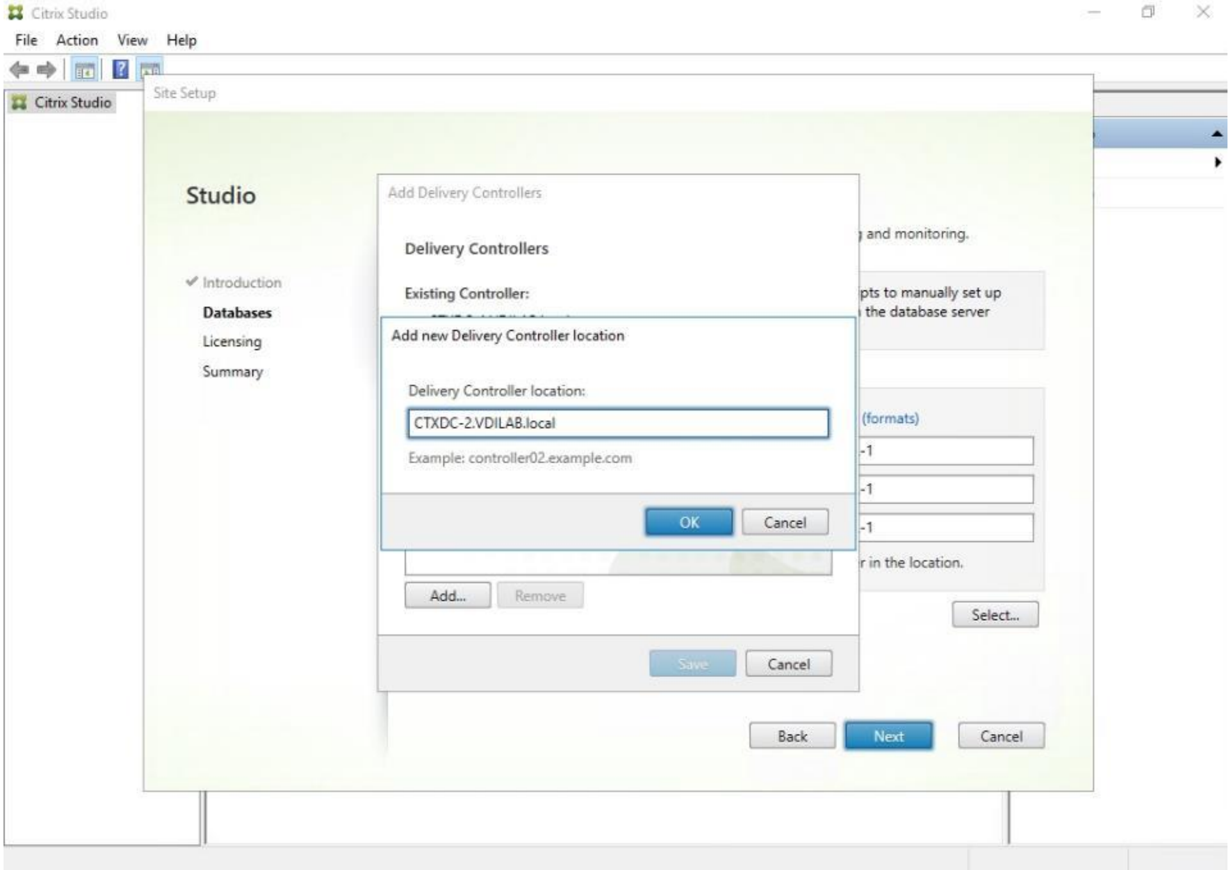

Delivery Controllers

Existing Controller:

Add new Delivery Controller location

Delivery Controller location:

CTXDC-2.VDILAB.local

Example: controller02.example.com

8. Click Save.

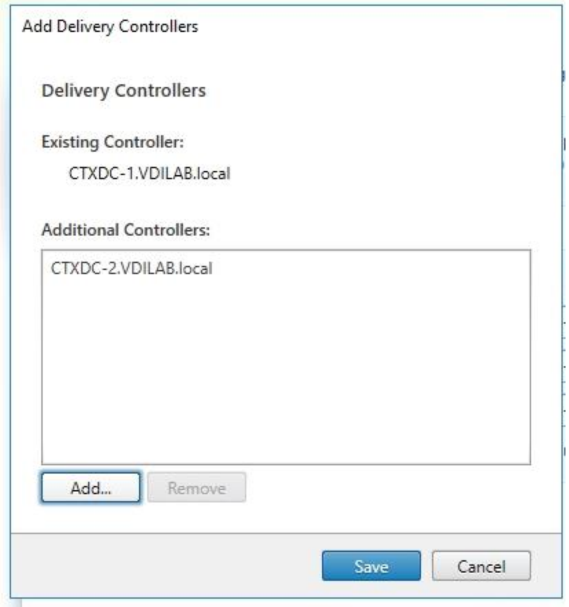

Delivery Controllers

Existing Controller:

CTXDC-1.VDILAB.local

Additional Controllers:

CTXDC-2.VDILAB.local

9. Click Next.

10. Provide the FQDN of the license server.

11. Click Connect to validate and retrieve any licenses from the server.

If no licenses are available, you can use the 30-day free trial or activate a license file.

12. Select the appropriate product edition using the license radio button.

13. Click Next.

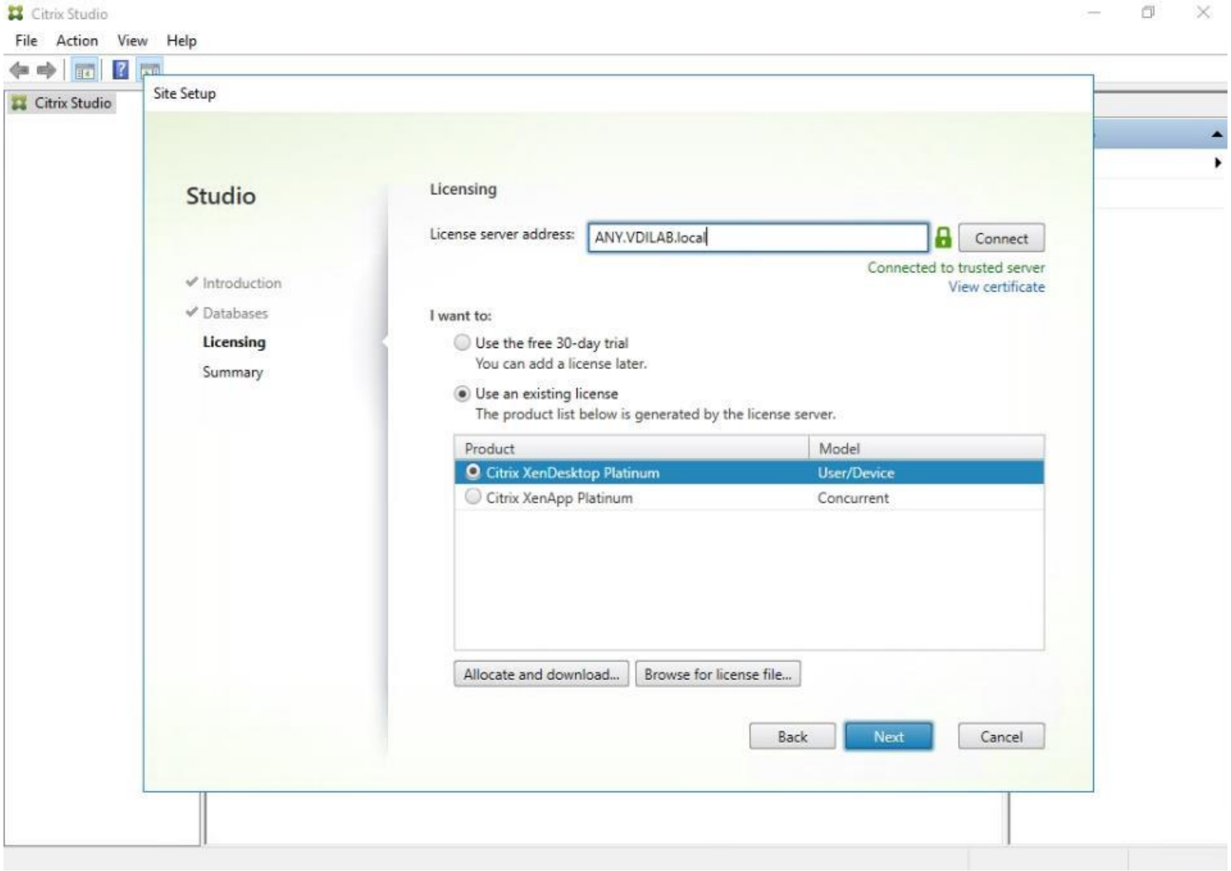

Licensing

License server address: ANY.VDILAB.local

Connected to trusted server
View certificate

I want to:

○ Use the free 30-day trial
You can add a license later.

● Use an existing license
The product list below is generated by the license server.

Product	Model
● Citrix XenDesktop Platinum	User/Device
○ Citrix XenApp Platinum	Concurrent

14. Verify information on the Summary page.

15. Click Finish.

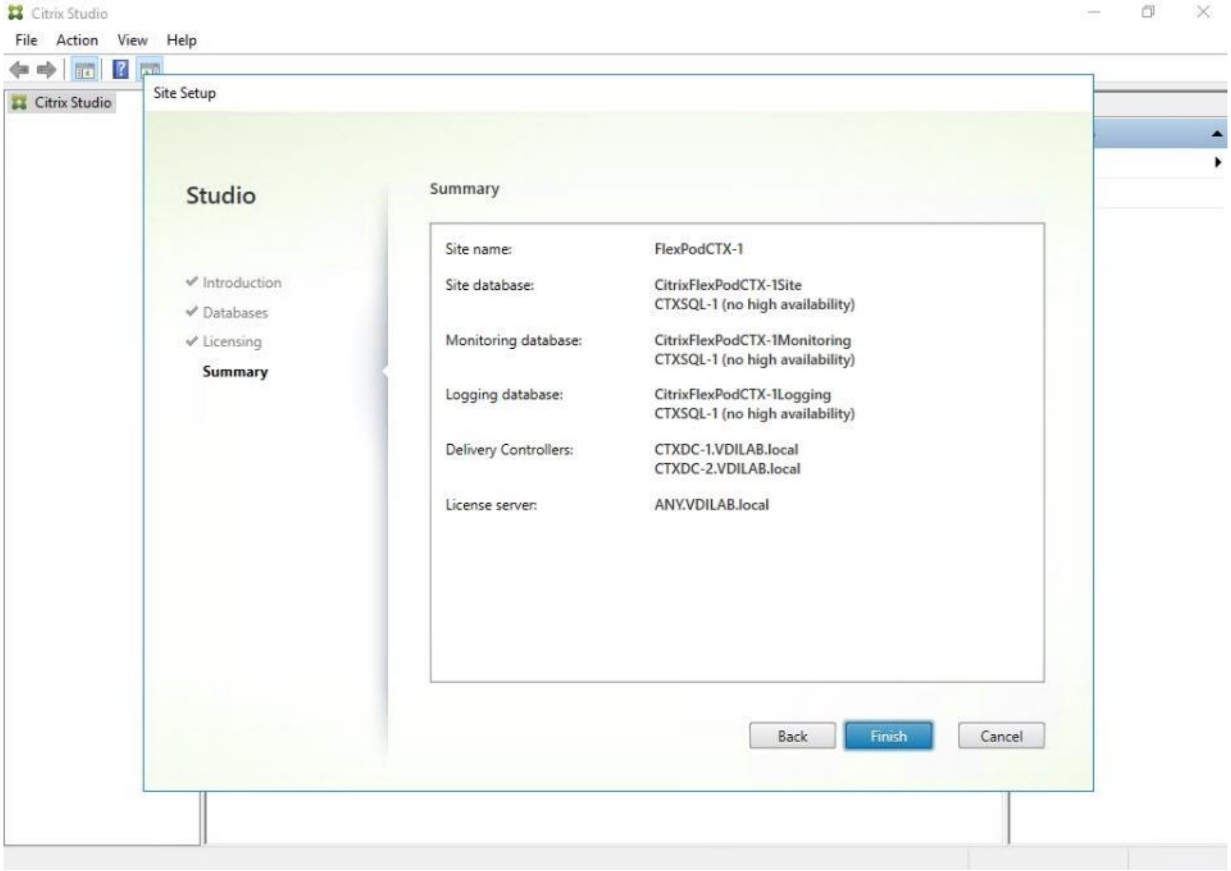

Summary

Site name:	FlexPodCTX-1
Site database:	CitrixFlexPodCTX-1Site CTXSQL-1 (no high availability)
Monitoring database:	CitrixFlexPodCTX-1Monitoring CTXSQL-1 (no high availability)
Logging database:	CitrixFlexPodCTX-1Logging CTXSQL-1 (no high availability)
Delivery Controllers:	CTXDC-1.VDILAB.local CTXDC-2.VDILAB.local
License server:	ANY.VDILAB.local

Configure the XenDesktop Site Hosting Connection

1. From Configuration > Hosting in Studio click Add Connection and Resources in the right pane.

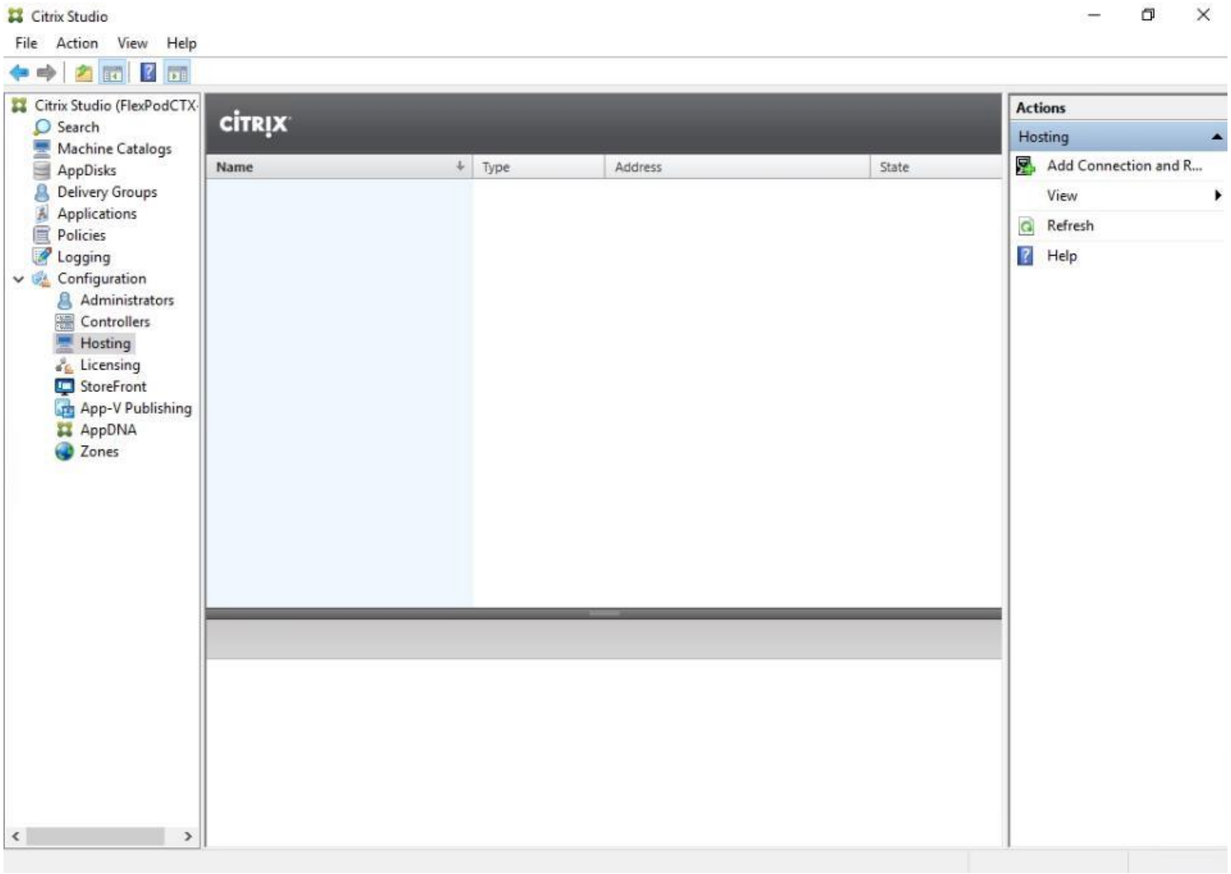

2. Select the Connection type of VMware vSphere®.

3. Enter the FQDN of the vCenter server (in Server_FQDN/sdk format).

4. Enter the username (in domain\username format) for the vSphere account.

5. Provide the password for the vSphere account.

6. Provide a connection name.

7. Check Studio Tools radio button required to support desktop provisioning task by this connection.

8. Click Next.

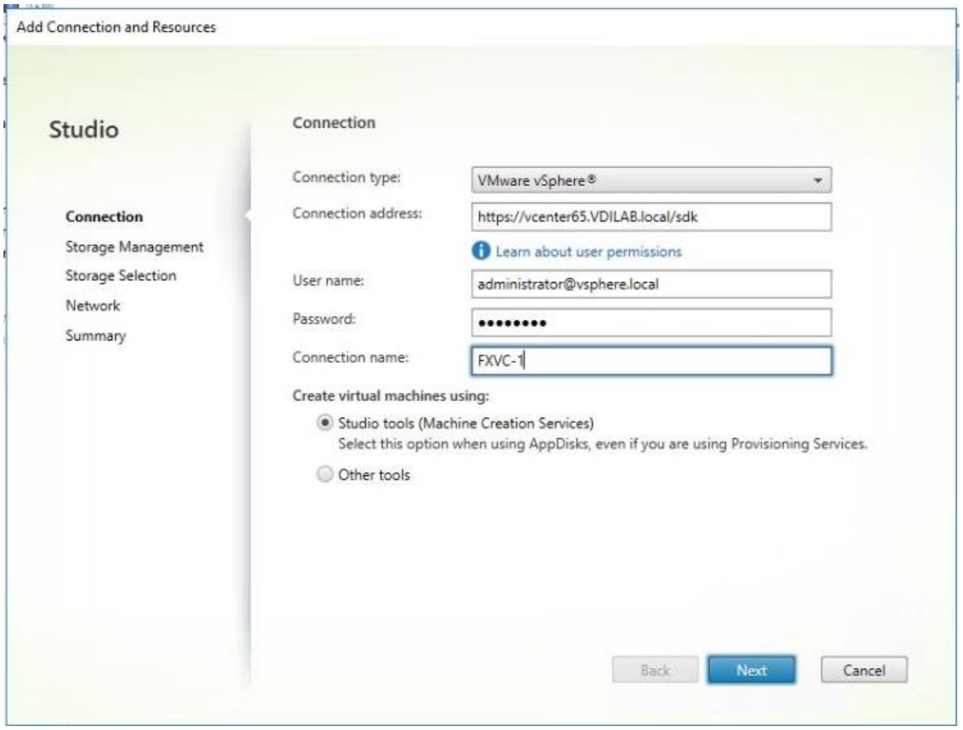

Add Connection and Resources

Studio

Connection

Connection type: VMware vSphere®

Connection address: https://vcenter65.VDILAB.local/sdk

Learn about user permissions

User name: administrator@vsphere.local

Password: ••••••••

Connection name: FXVC-1

Create virtual machines using:

● Studio tools (Machine Creation Services)
Select this option when using AppDisks, even if you are using Provisioning Services.

○ Other tools

9. Accept the certificate and click OK to trust the hypervisor connection.

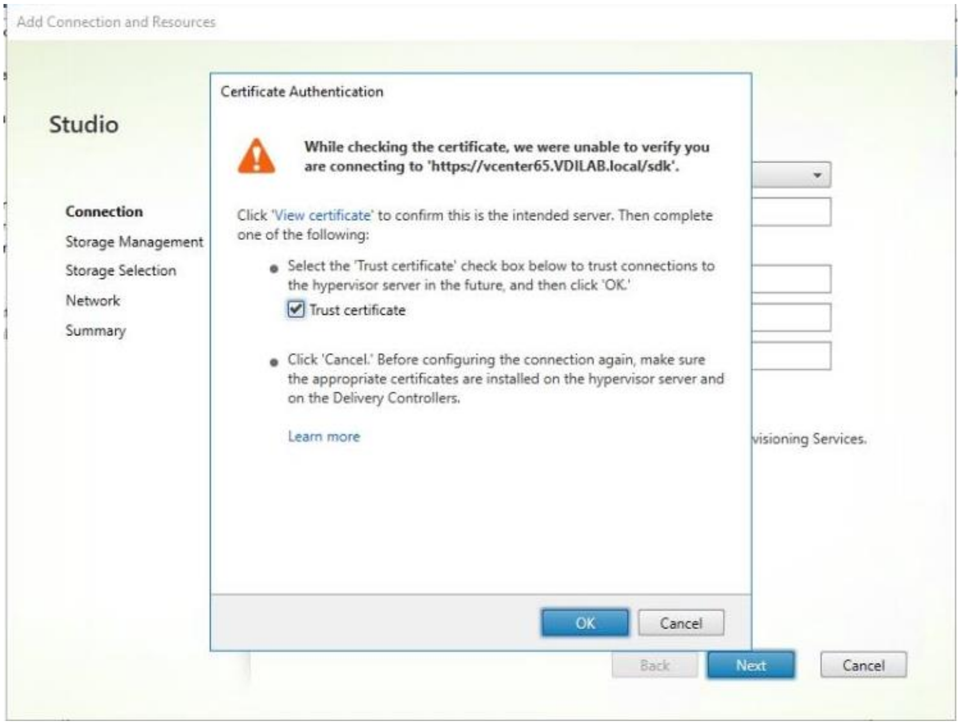

Certificate Authentication

While checking the certificate, we were unable to verify you are connecting to 'https://vcenter65.VDILAB.local/sdk'.

Click 'View certificate' to confirm this is the intended server. Then complete one of the following:

- Select the 'Trust certificate' check box below to trust connections to the hypervisor server in the future, and then click 'OK.'
  ☑ Trust certificate
- Click 'Cancel.' Before configuring the connection again, make sure the appropriate certificates are installed on the hypervisor server and on the Delivery Controllers.

Learn more

10. Select Cluster that will be used by this connection.

11. Check Use storage shared by hypervisors radio button.

12. Click Next.

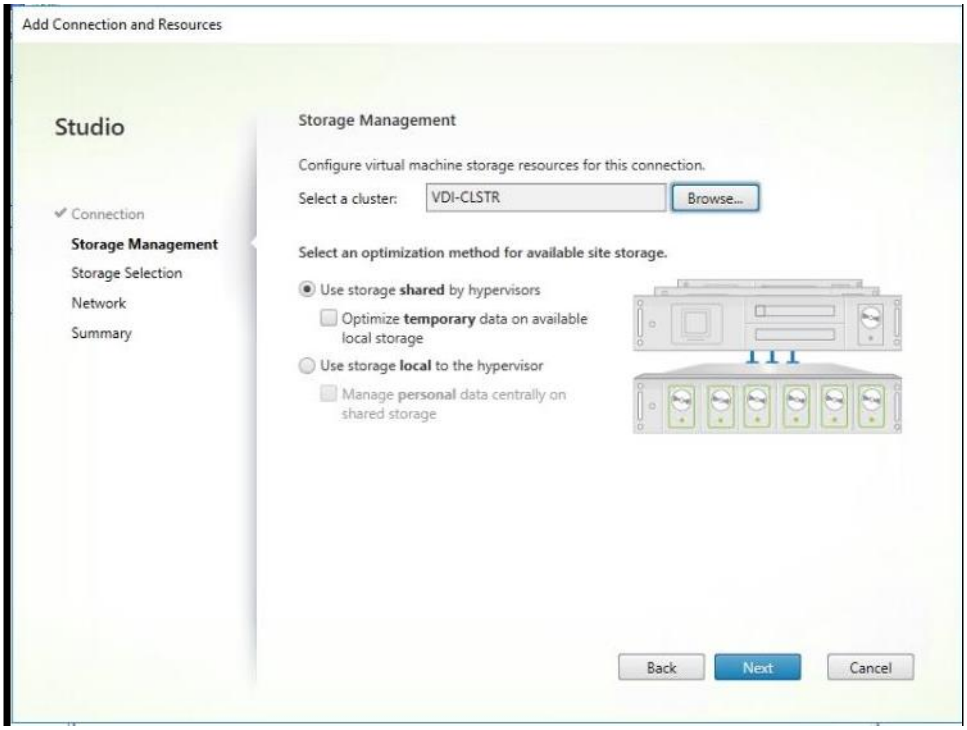

Add Connection and Resources

Storage Management

Configure virtual machine storage resources for this connection.

Select a cluster: VDI-CLSTR

Select an optimization method for available site storage.

● Use storage **shared** by hypervisors

□ Optimize **temporary** data on available local storage

○ Use storage **local** to the hypervisor

□ Manage **personal** data centrally on shared storage

13. Make Storage selection to be used by this connection, use all provisioned for desktops NFS datastores.

14. Click Next.

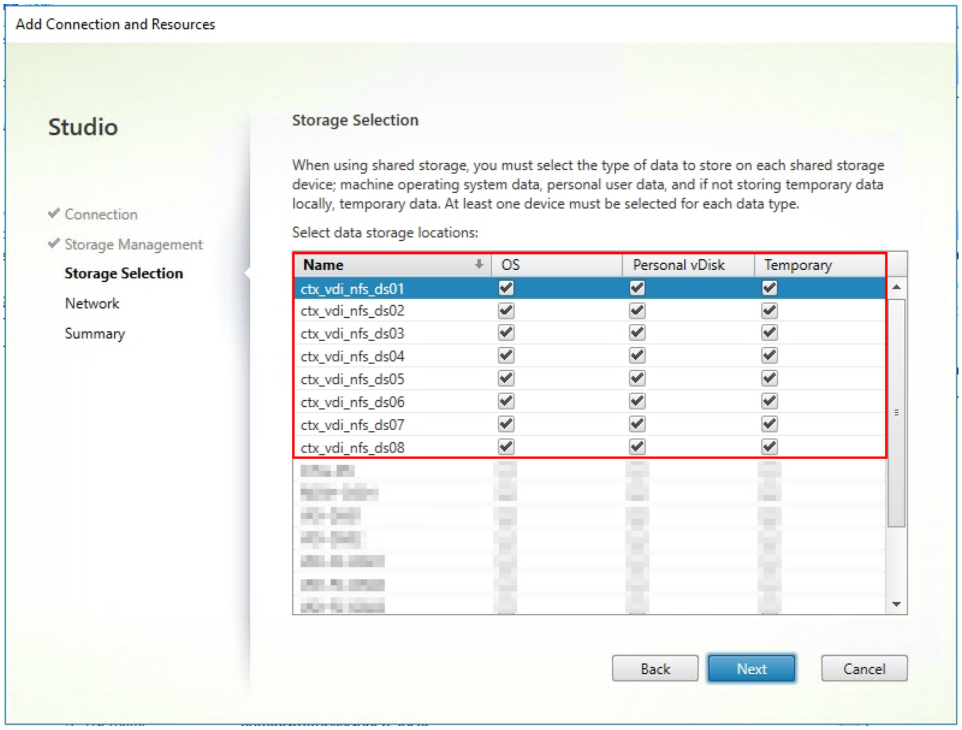

Add Connection and Resources

Storage Selection

When using shared storage, you must select the type of data to store on each shared storage device; machine operating system data, personal user data, and if not storing temporary data locally, temporary data. At least one device must be selected for each data type.

Select data storage locations:

Name	OS	Personal vDisk	Temporary
ctx_vdi_nfs_ds01	☑	☑	☑
ctx_vdi_nfs_ds02	☑	☑	☑
ctx_vdi_nfs_ds03	☑	☑	☑
ctx_vdi_nfs_ds04	☑	☑	☑
ctx_vdi_nfs_ds05	☑	☑	☑
ctx_vdi_nfs_ds06	☑	☑	☑
ctx_vdi_nfs_ds07	☑	☑	☑
ctx_vdi_nfs_ds08	☑	☑	☑

15. Make Network selection to be used by this connection.

16. Click Next.

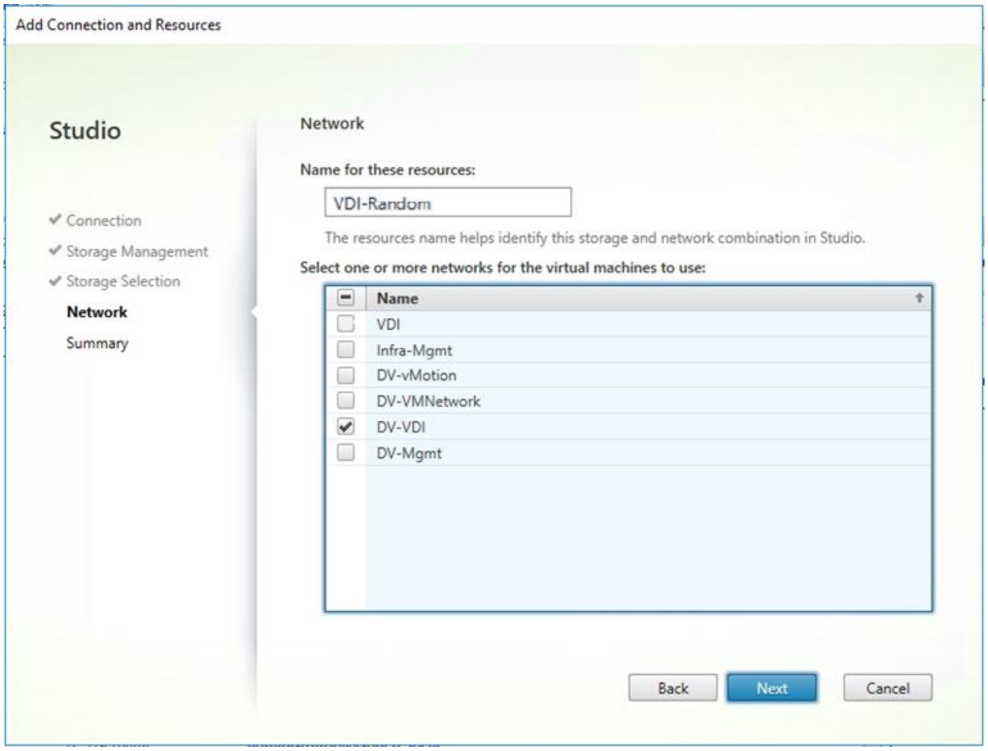

Add Connection and Resources

Network

Name for these resources: VDI-Random

The resources name helps identify this storage and network combination in Studio.

Select one or more networks for the virtual machines to use:

☐	Name
☐	VDI
☐	Infra-Mgmt
☐	DV-vMotion
☐	DV-VMNetwork
☑	DV-VDI
☐	DV-Mgmt

17. Review Site configuration Summary and click Finish.

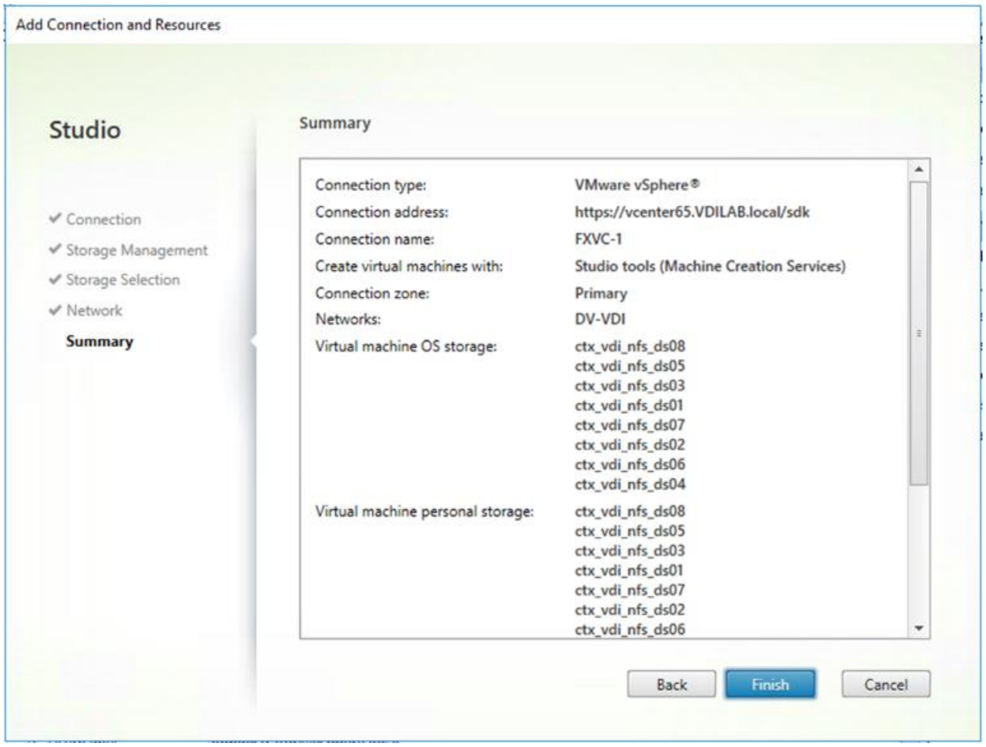

Add Connection and Resources

Summary

Connection type:	VMware vSphere®
Connection address:	https://vcenter65.VDILAB.local/sdk
Connection name:	FXVC-1
Create virtual machines with:	Studio tools (Machine Creation Services)
Connection zone:	Primary
Networks:	DV-VDI
Virtual machine OS storage:	ctx_vdi_nfs_ds08 ctx_vdi_nfs_ds05 ctx_vdi_nfs_ds03 ctx_vdi_nfs_ds01 ctx_vdi_nfs_ds07 ctx_vdi_nfs_ds02 ctx_vdi_nfs_ds06 ctx_vdi_nfs_ds04
Virtual machine personal storage:	ctx_vdi_nfs_ds08 ctx_vdi_nfs_ds05 ctx_vdi_nfs_ds03 ctx_vdi_nfs_ds01 ctx_vdi_nfs_ds07 ctx_vdi_nfs_ds02 ctx_vdi_nfs_ds06

Adding Resources to the Site Hosting Connection

To add resources to the additional vcenter clusters, complete the following steps:

1. From Configuration > Hosting in Studio click Add Connection and Resources in the right pane.

2. Select Use an existing Connection, use connection previously created for FlexPod environment.

3. Click Next.

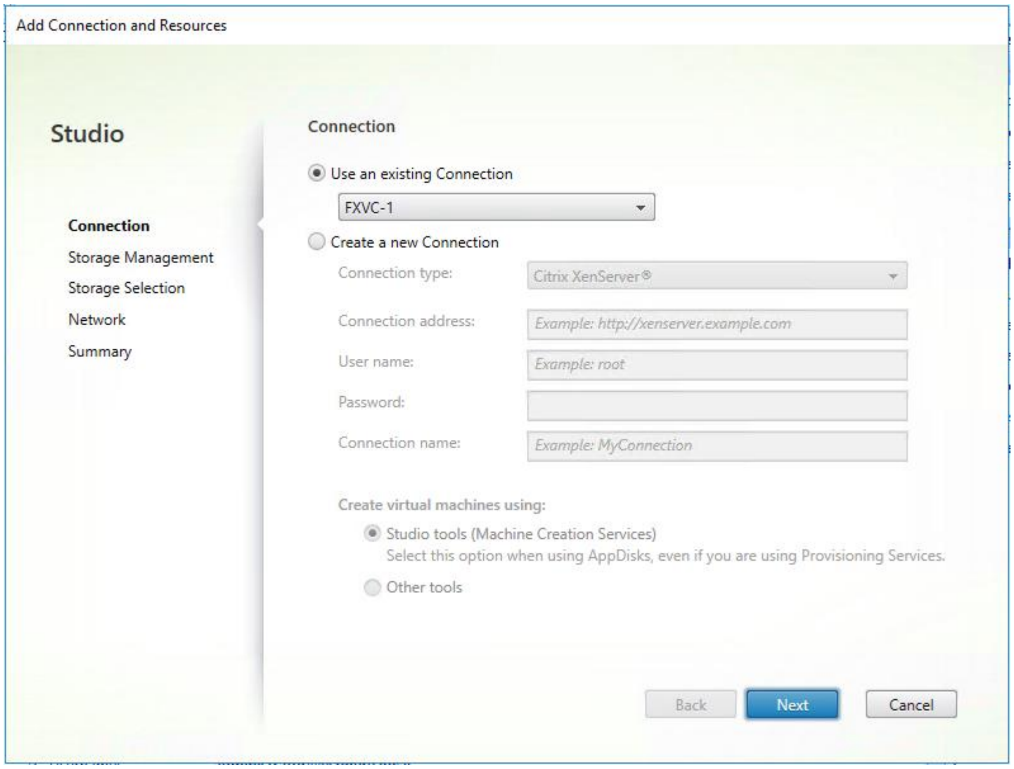

Add Connection and Resources

Studio

Connection

● Use an existing Connection

FXVC-1

○ Create a new Connection

Connection type: Citrix XenServer®

Connection address: Example: http://xenserver.example.com

User name: Example: root

Password:

Connection name: Example: MyConnection

Create virtual machines using:

● Studio tools (Machine Creation Services)
Select this option when using AppDisks, even if you are using Provisioning Services.

○ Other tools

4. Select Cluster you adding to this connection.

5. Check Use storage shared by hypervisors radio button.

6. Click Next.

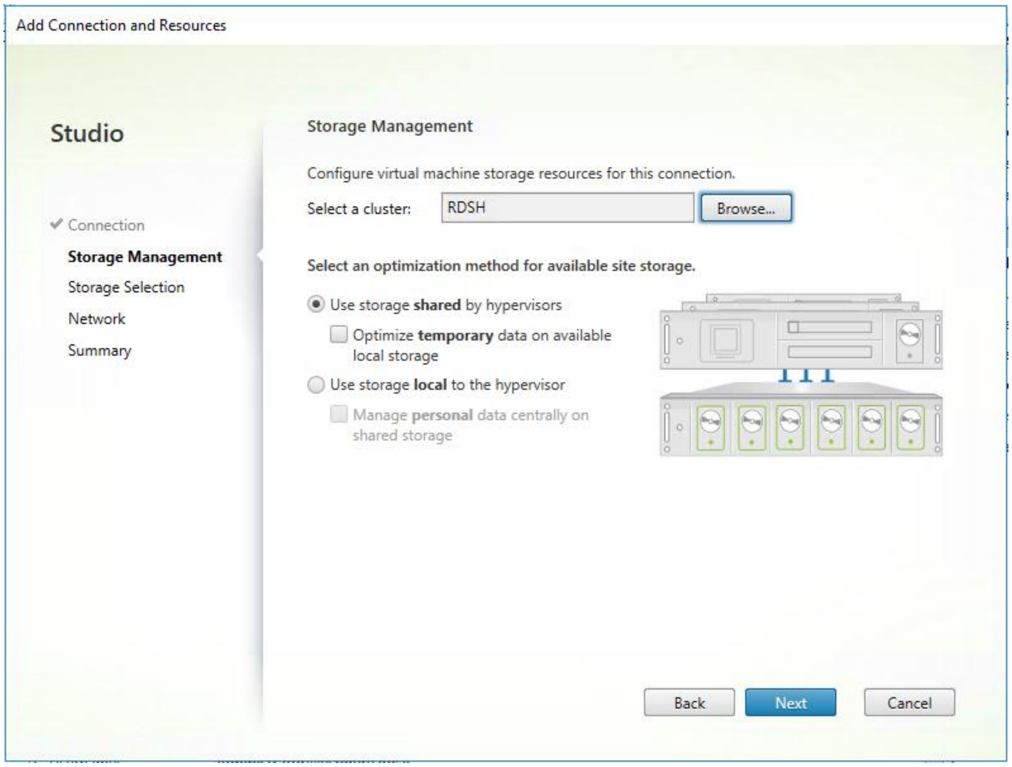

Add Connection and Resources

Studio

- ✓ Connection
- **Storage Management**
- Storage Selection
- Network
- Summary

Storage Management

Configure virtual machine storage resources for this connection.

Select a cluster: RDSH [Browse...]

Select an optimization method for available site storage.

● Use storage **shared** by hypervisors

□ Optimize **temporary** data on available local storage

○ Use storage **local** to the hypervisor

□ Manage **personal** data centrally on shared storage

7. Make Storage selection to be used by this connection, use all provisioned for desktops NFS datastores.

8. Click Next.

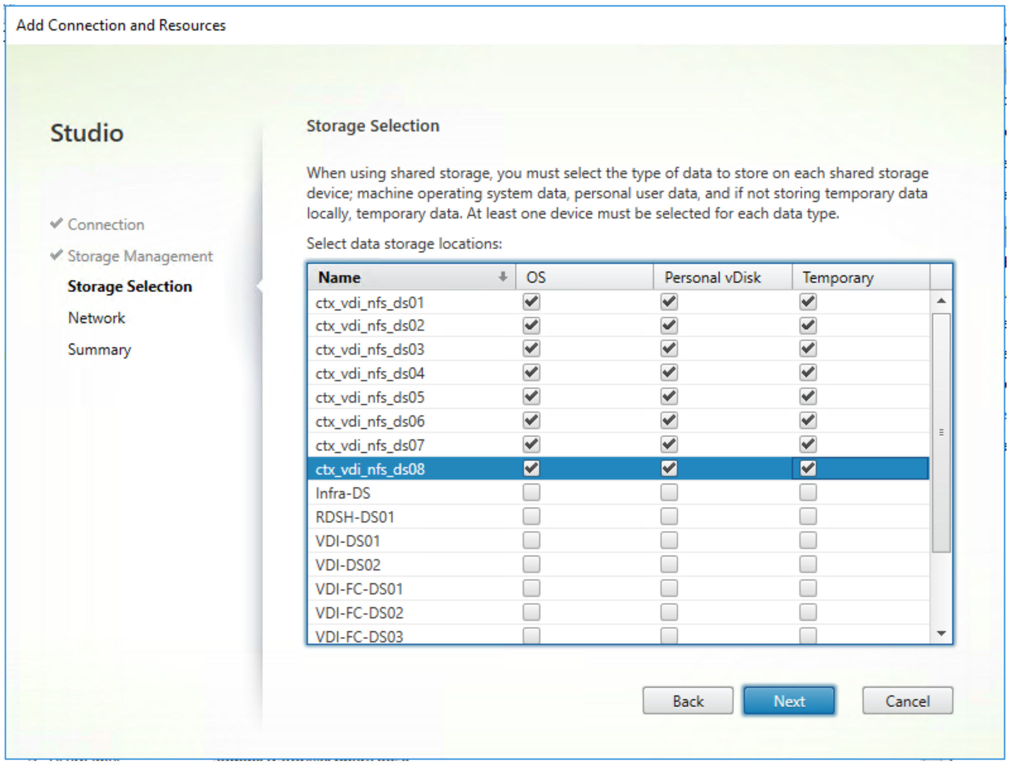

Add Connection and Resources

Studio

- ✓ Connection
- ✓ Storage Management
- **Storage Selection**
- Network
- Summary

Storage Selection

When using shared storage, you must select the type of data to store on each shared storage device; machine operating system data, personal user data, and if not storing temporary data locally, temporary data. At least one device must be selected for each data type.

Select data storage locations:

Name	OS	Personal vDisk	Temporary
ctx_vdi_nfs_ds01	☑	☑	☑
ctx_vdi_nfs_ds02	☑	☑	☑
ctx_vdi_nfs_ds03	☑	☑	☑
ctx_vdi_nfs_ds04	☑	☑	☑
ctx_vdi_nfs_ds05	☑	☑	☑
ctx_vdi_nfs_ds06	☑	☑	☑
ctx_vdi_nfs_ds07	☑	☑	☑
ctx_vdi_nfs_ds08	☑	☑	☑
Infra-DS	□	□	□
RDSH-DS01	□	□	□
VDI-DS01	□	□	□
VDI-DS02	□	□	□
VDI-FC-DS01	□	□	□
VDI-FC-DS02	□	□	□
VDI-FC-DS03	□	□	□

9. Make Network selection to be used by this connection.

10. Click Next.

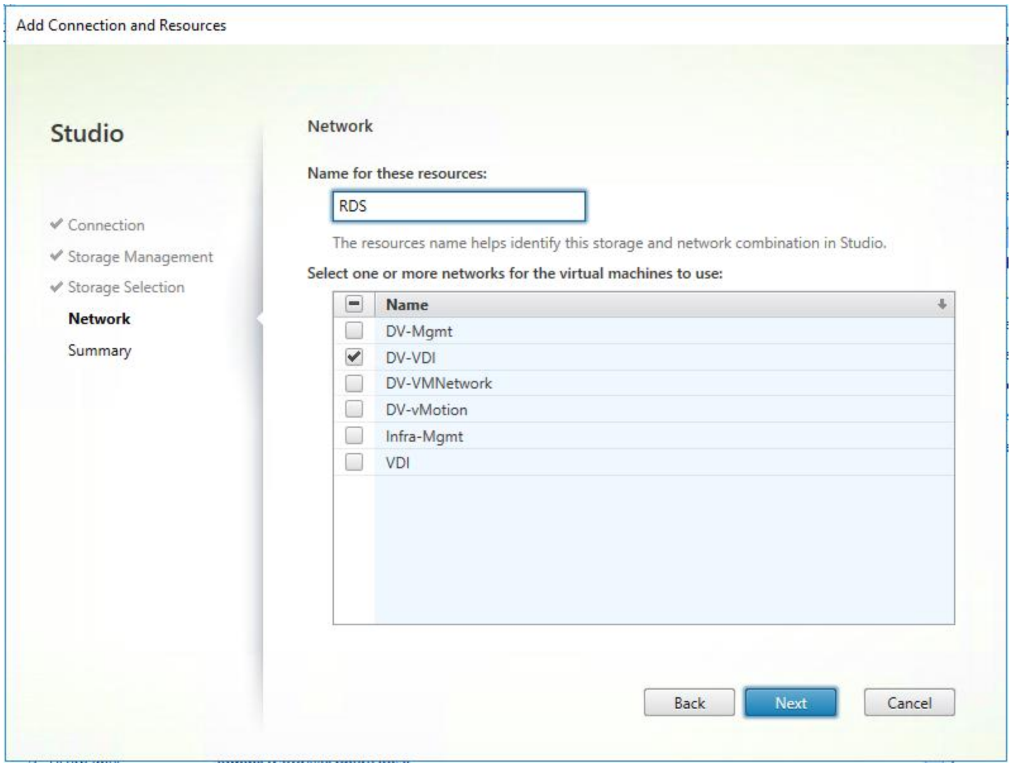

Add Connection and Resources

Network

Name for these resources: RDS

The resources name helps identify this storage and network combination in Studio.

Select one or more networks for the virtual machines to use:

	Name
□	DV-Mgmt
☑	DV-VDI
□	DV-VMNetwork
□	DV-vMotion
□	Infra-Mgmt
□	VDI

11. Review the Site configuration Summary and click Finish.

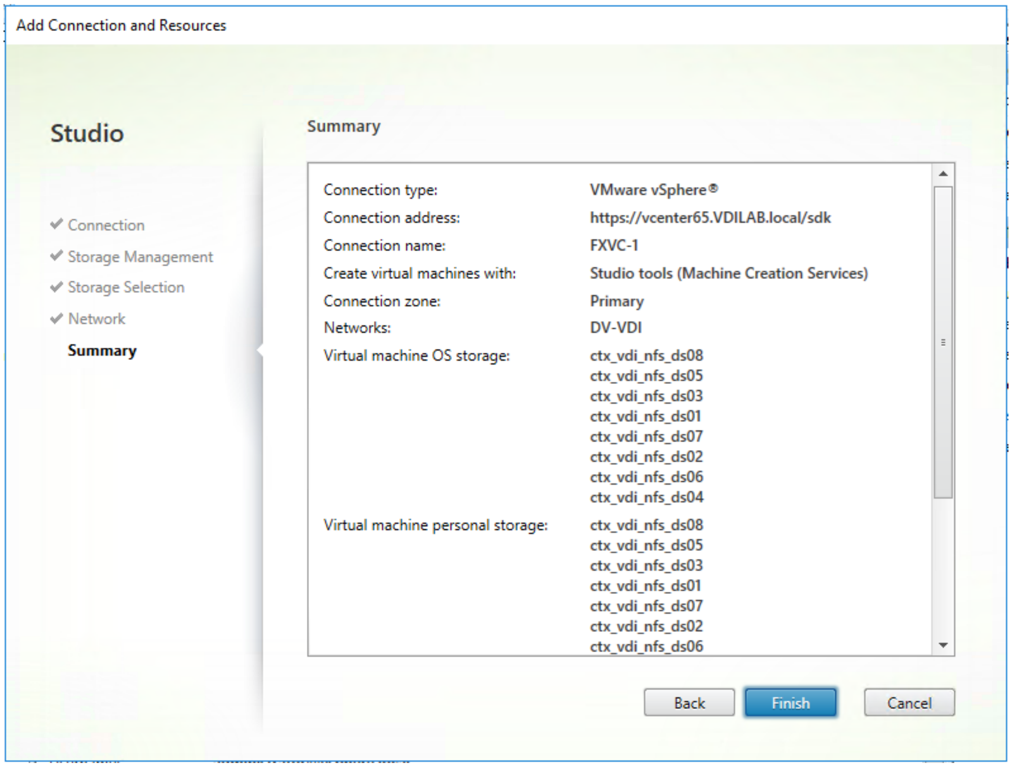

Add Connection and Resources

Summary

Connection type:	VMware vSphere®
Connection address:	https://vcenter65.VDILAB.local/sdk
Connection name:	FXVC-1
Create virtual machines with:	Studio tools (Machine Creation Services)
Connection zone:	Primary
Networks:	DV-VDI
Virtual machine OS storage:	ctx_vdi_nfs_ds08 ctx_vdi_nfs_ds05 ctx_vdi_nfs_ds03 ctx_vdi_nfs_ds01 ctx_vdi_nfs_ds07 ctx_vdi_nfs_ds02 ctx_vdi_nfs_ds06 ctx_vdi_nfs_ds04
Virtual machine personal storage:	ctx_vdi_nfs_ds08 ctx_vdi_nfs_ds05 ctx_vdi_nfs_ds03 ctx_vdi_nfs_ds01 ctx_vdi_nfs_ds07 ctx_vdi_nfs_ds02 ctx_vdi_nfs_ds06

12. Repeat these steps to add all additional clusters (Figure 35).

Figure 35 FlexPod Hosting Connection in Studio with Three Clusters

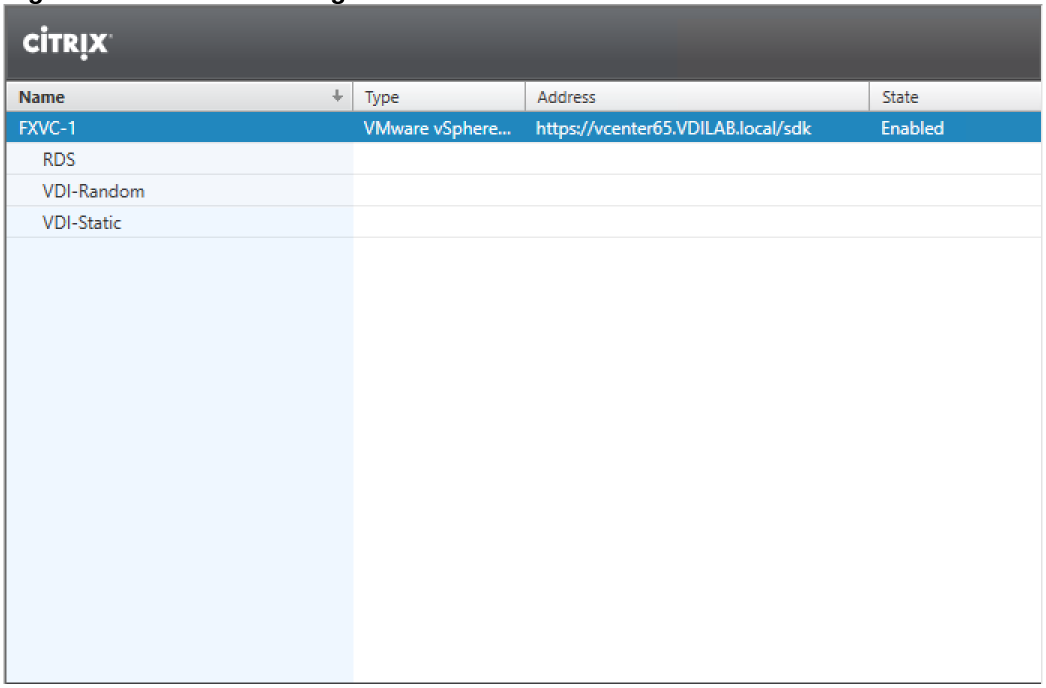

Name	Type	Address	State
FXVC-1	VMware vSphere...	https://vcenter65.VDILAB.local/sdk	Enabled
RDS			
VDI-Random			
VDI-Static			

Configure the XenDesktop Site Administrators

1. Connect to the XenDesktop server and open Citrix Studio Management console.

2. From the Configuration menu, right-click Administrator and select Create Administrator from the drop-down list.

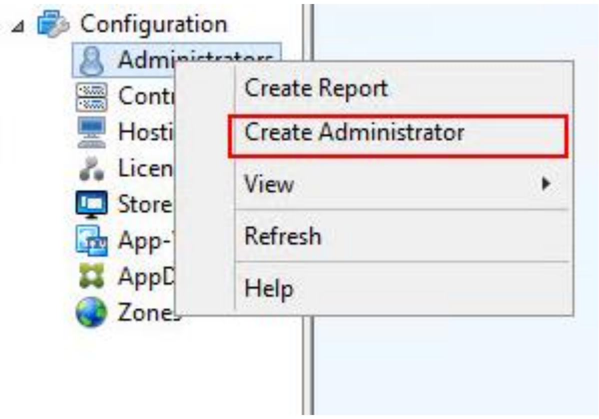

3. Select/Create appropriate scope and click Next.

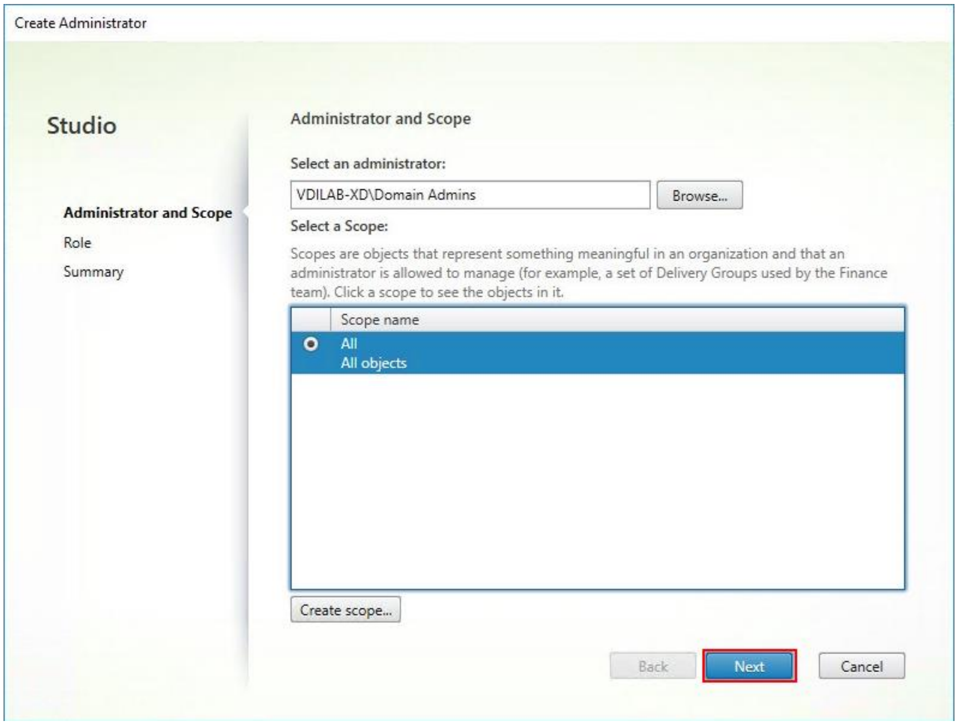

Create Administrator

Studio

Administrator and Scope

Role

Summary

Administrator and Scope

Select an administrator:

VDILAB-XD\Domain Admins

Browse...

Select a Scope:

Scopes are objects that represent something meaningful in an organization and that an administrator is allowed to manage (for example, a set of Delivery Groups used by the Finance team). Click a scope to see the objects in it.

	Scope name
●	All All objects

Create scope...

Back Next Cancel

4. Choose an appropriate Role.

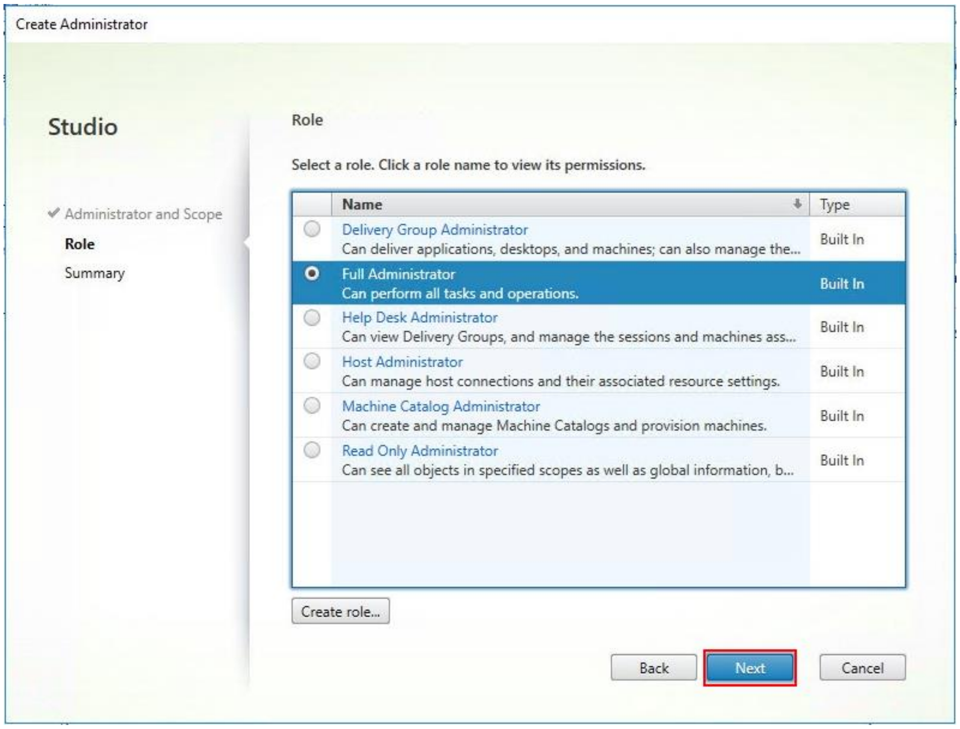

Create Administrator

Studio

✔ Administrator and Scope

Role

Summary

Role

Select a role. Click a role name to view its permissions.

	Name	Type
○	Delivery Group Administrator Can deliver applications, desktops, and machines; can also manage the...	Built In
●	Full Administrator Can perform all tasks and operations.	Built In
○	Help Desk Administrator Can view Delivery Groups, and manage the sessions and machines ass...	Built In
○	Host Administrator Can manage host connections and their associated resource settings.	Built In
○	Machine Catalog Administrator Can create and manage Machine Catalogs and provision machines.	Built In
○	Read Only Administrator Can see all objects in specified scopes as well as global information, b...	Built In

Create role...

Back Next Cancel

5. Review the Summary, check Enable administrator and click Finish.

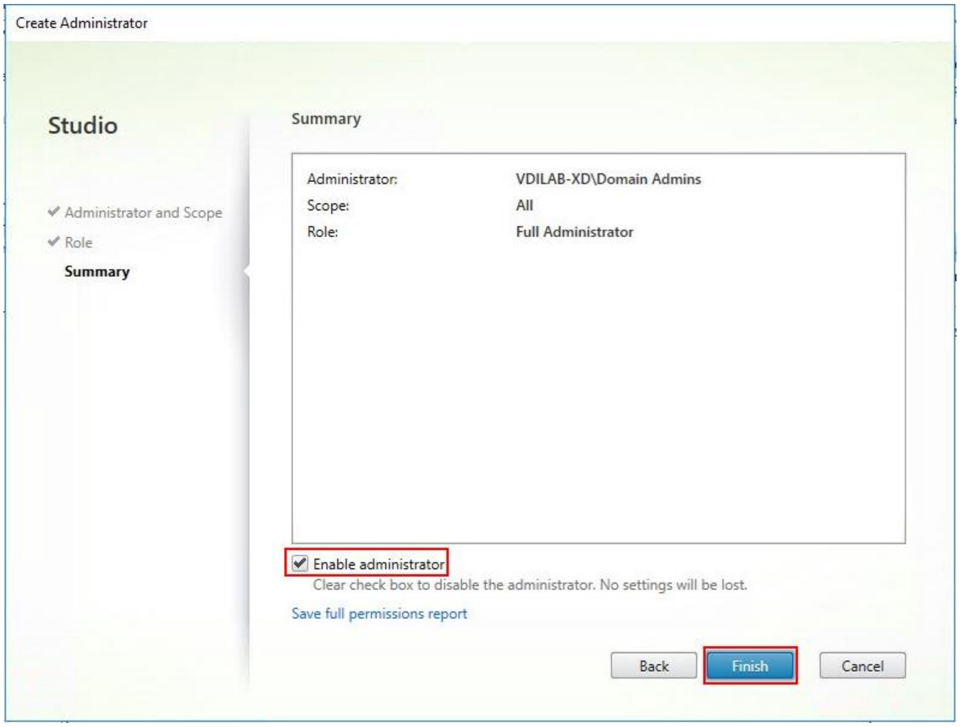

Summary

Administrator:	VDILAB-XD\Domain Admins
Scope:	All
Role:	Full Administrator

☑ Enable administrator

Clear check box to disable the administrator. No settings will be lost.

Save full permissions report

Installing and configuring StoreFront

Citrix StoreFront stores aggregate desktops and applications from XenDesktop sites, making resources readily available to users. In this CVD, we created two StoreFront servers on dedicated virtual machines.

1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix XenDesktop 7.15 ISO.

2. Click Start.

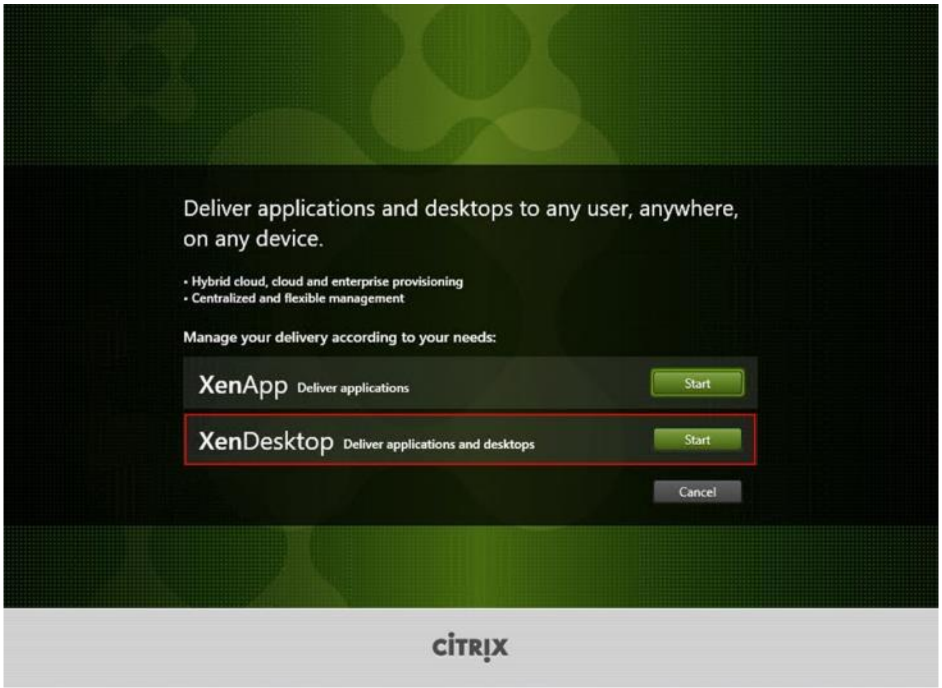

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications
Start
XenDesktop Deliver applications and desktops
Start
Cancel
CITRIX

3. Click Extend Deployment Citrix StoreFront.

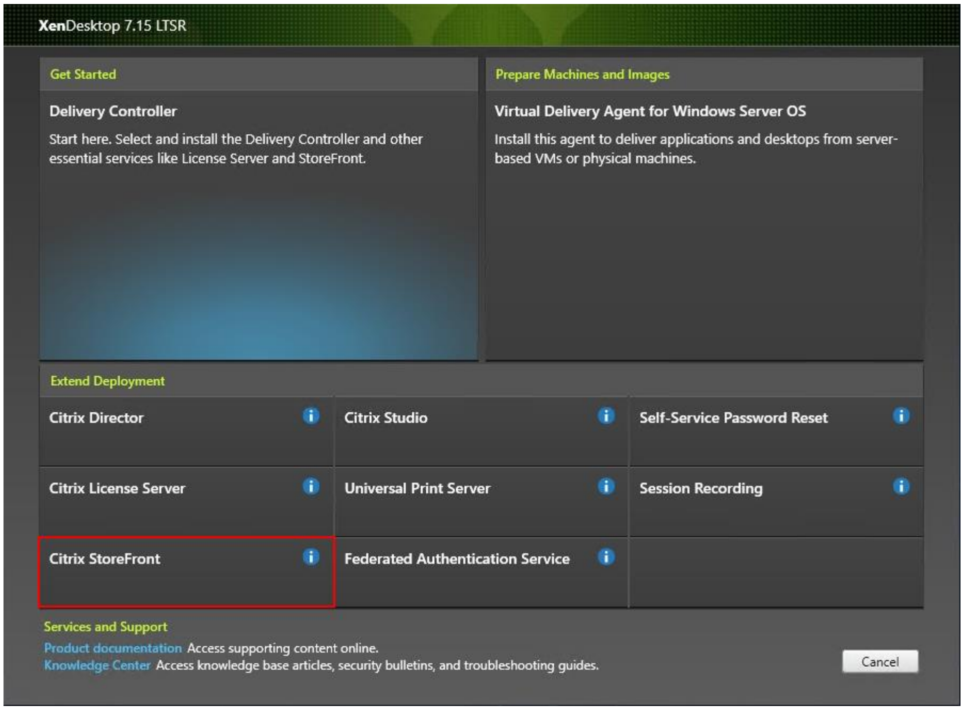

XenDesktop 7.15 LTSR

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Extend Deployment

Services and Support

Product documentation Access supporting content online.

Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

5. Click Next.

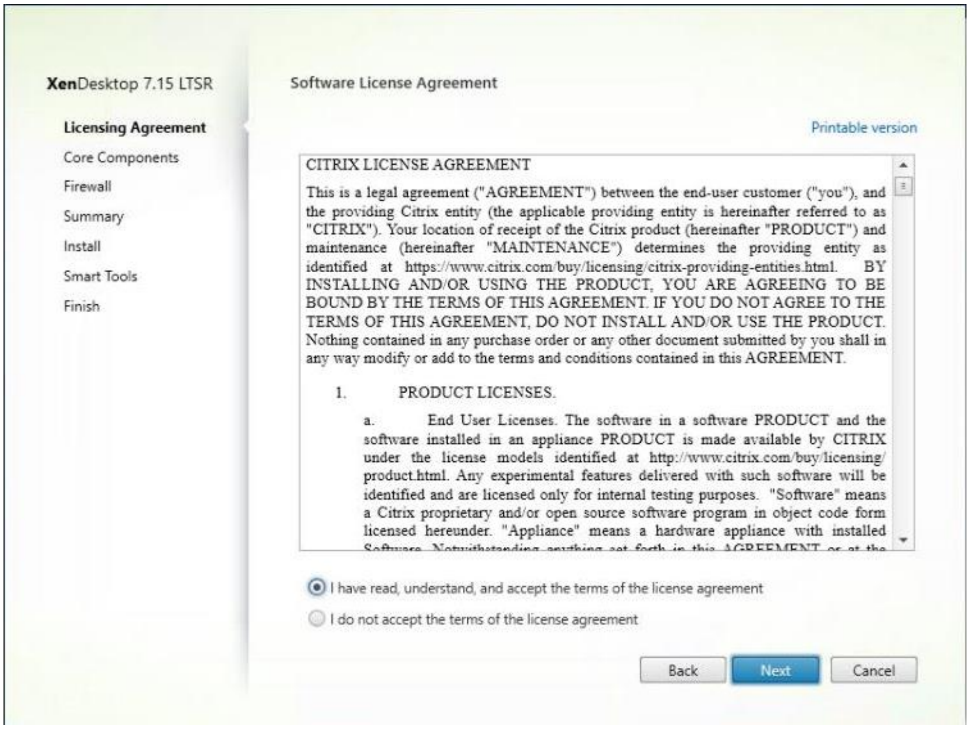

Software License Agreement

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY INSTALLING AND/OR USING THE PRODUCT, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT. Nothing contained in any purchase order or any other document submitted by you shall in any way modify or add to the terms and conditions contained in this AGREEMENT.

1. PRODUCT LICENSES.

a. End User Licenses. The software in a software PRODUCT and the software installed in an appliance PRODUCT is made available by CITRIX under the license models identified at http://www.citrix.com/buy/licensing/product.html. Any experimental features delivered with such software will be identified and are licensed only for internal testing purposes. "Software" means a Citrix proprietary and/or open source software program in object code form licensed hereunder. "Appliance" means a hardware appliance with installed Software. Notwithstanding anything set forth in this AGREEMENT or at the

● I have read, understand, and accept the terms of the license agreement

○ I do not accept the terms of the license agreement

6. Click Next.

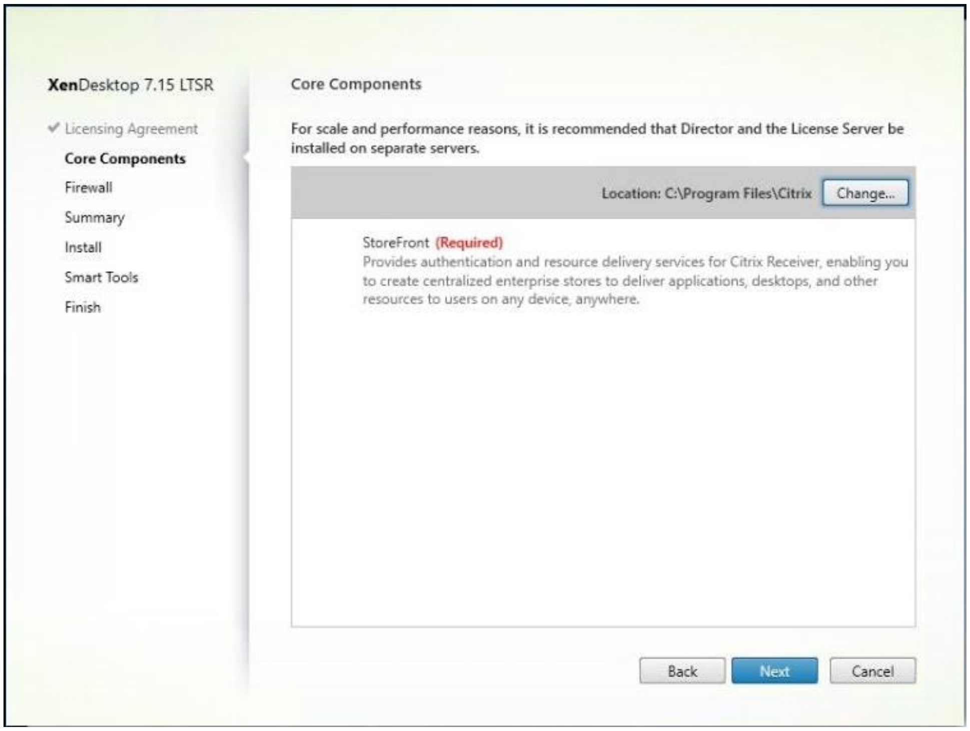

Core Components

For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.

Location: C:\Program Files\Citrix

StoreFront (Required)
Provides authentication and resource delivery services for Citrix Receiver, enabling you to create centralized enterprise stores to deliver applications, desktops, and other resources to users on any device, anywhere.

7. Select the default ports and automatically configured firewall rules.

8. Click Next.

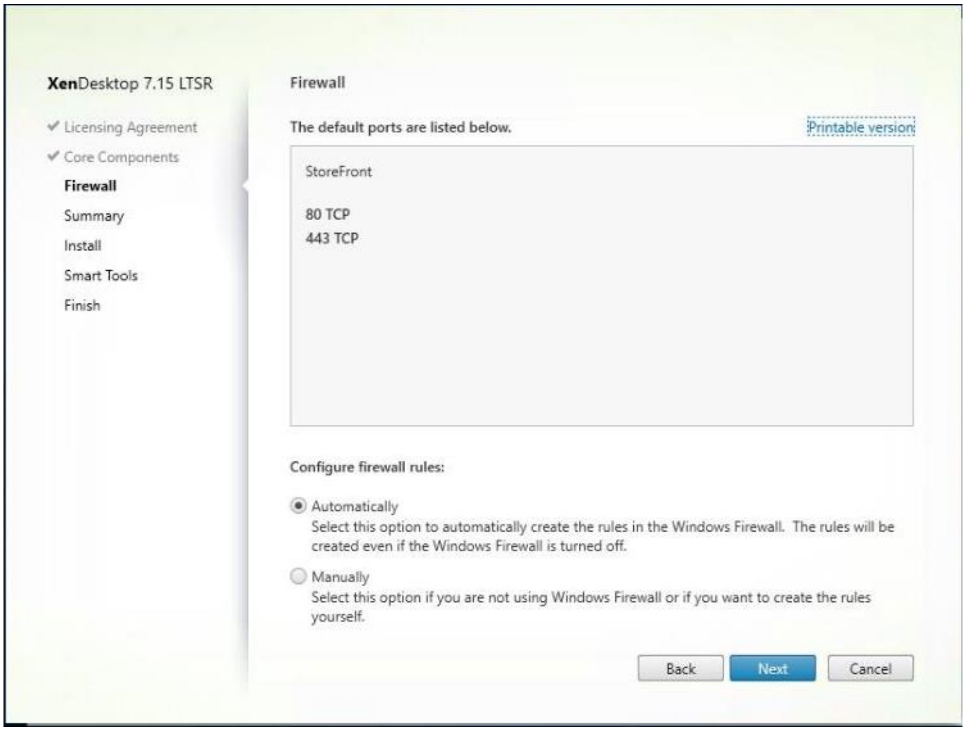

Firewall

The default ports are listed below.

StoreFront

80 TCP
443 TCP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

9. Click Install.

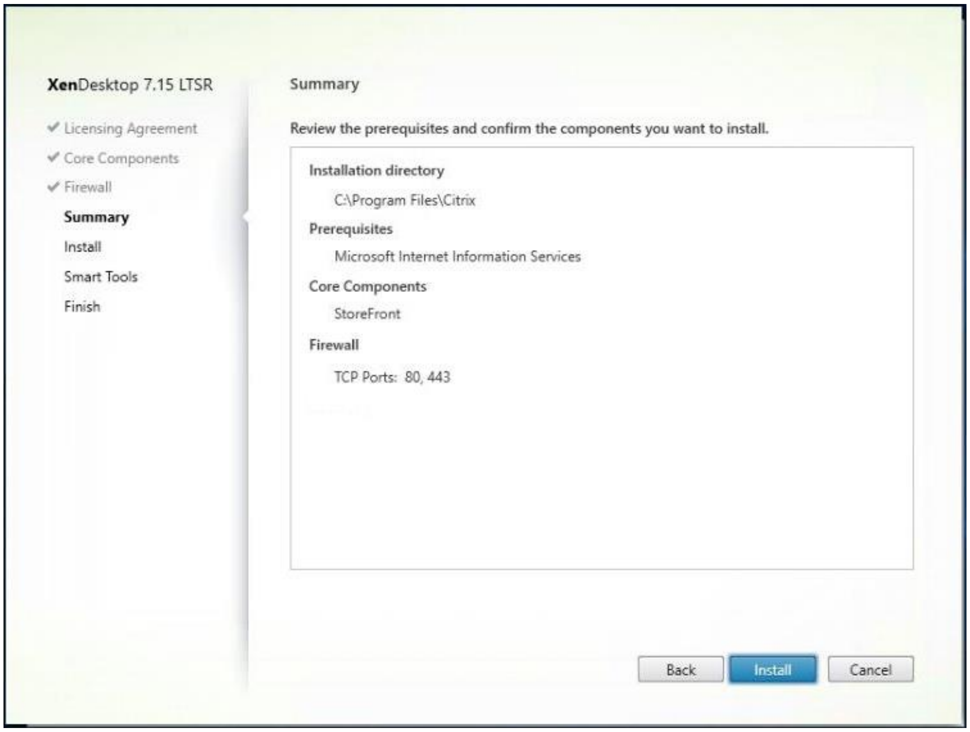

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory
C:\Program Files\Citrix

Prerequisites
Microsoft Internet Information Services

Core Components
StoreFront

Firewall
TCP Ports: 80, 443

10. (Optional) Click “I want to participate in Call Home.”

11. Click Next.

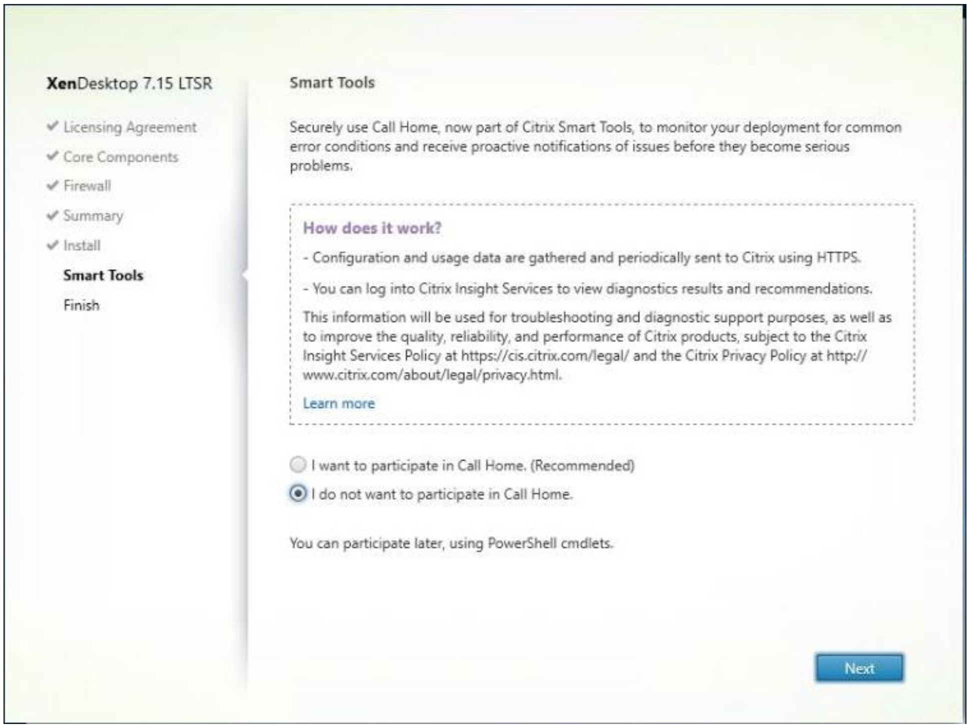

Smart Tools

Securely use Call Home, now part of Citrix Smart Tools, to monitor your deployment for common error conditions and receive proactive notifications of issues before they become serious problems.

How does it work?

- Configuration and usage data are gathered and periodically sent to Citrix using HTTPS.

- You can log into Citrix Insight Services to view diagnostics results and recommendations.

This information will be used for troubleshooting and diagnostic support purposes, as well as to improve the quality, reliability, and performance of Citrix products, subject to the Citrix Insight Services Policy at https://cis.citrix.com/legal/ and the Citrix Privacy Policy at http://www.citrix.com/about/legal/privacy.html.

○ I want to participate in Call Home. (Recommended)

● I do not want to participate in Call Home.

You can participate later, using PowerShell cmdlets.

12. Check “Open the StoreFront Management Console."

13. Click Finish.

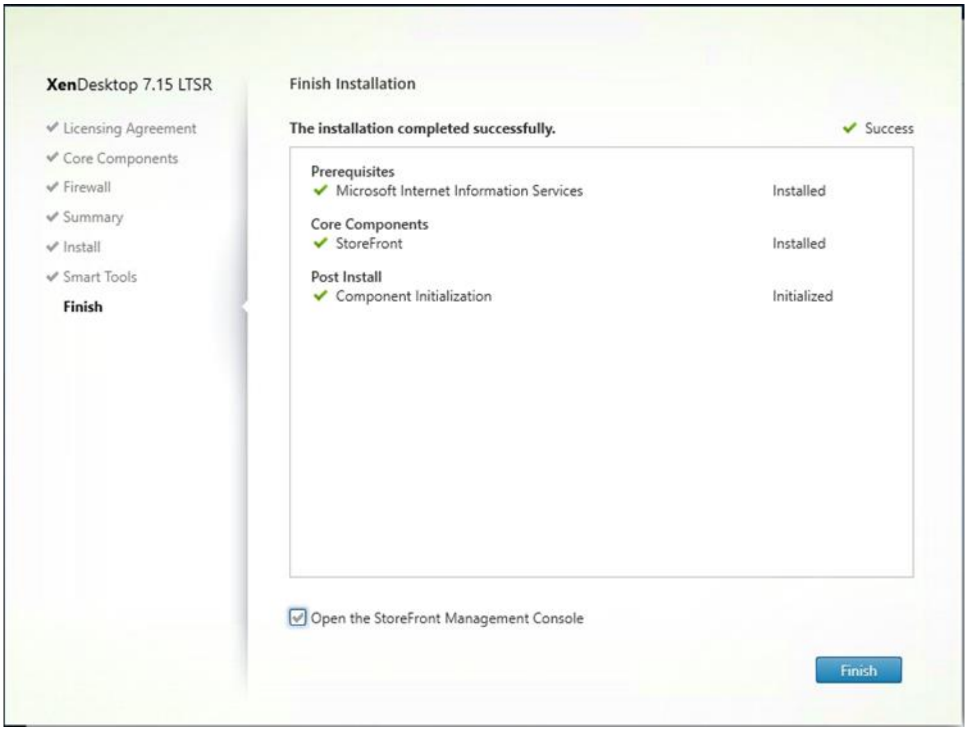

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites
✓ Microsoft Internet Information Services — Installed

Core Components
✓ StoreFront — Installed

Post Install
✓ Component Initialization — Initialized

☑ Open the StoreFront Management Console

14. Click Create a new deployment.

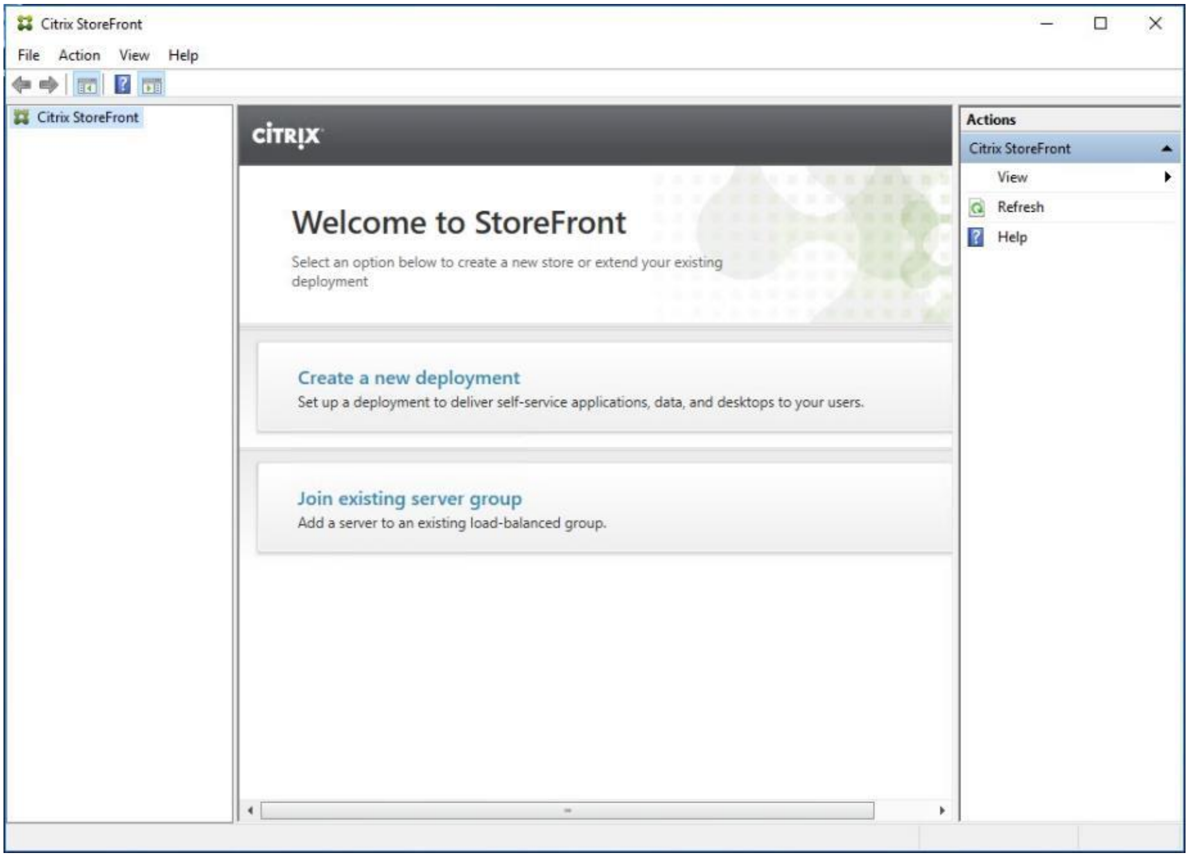

Welcome to StoreFront

Select an option below to create a new store or extend your existing deployment

Create a new deployment
Set up a deployment to deliver self-service applications, data, and desktops to your users.

Join existing server group
Add a server to an existing load-balanced group.

15. Specify the URL of the StoreFront server.

For a multiple server deployment use the load balancing environment in the Base URL box.

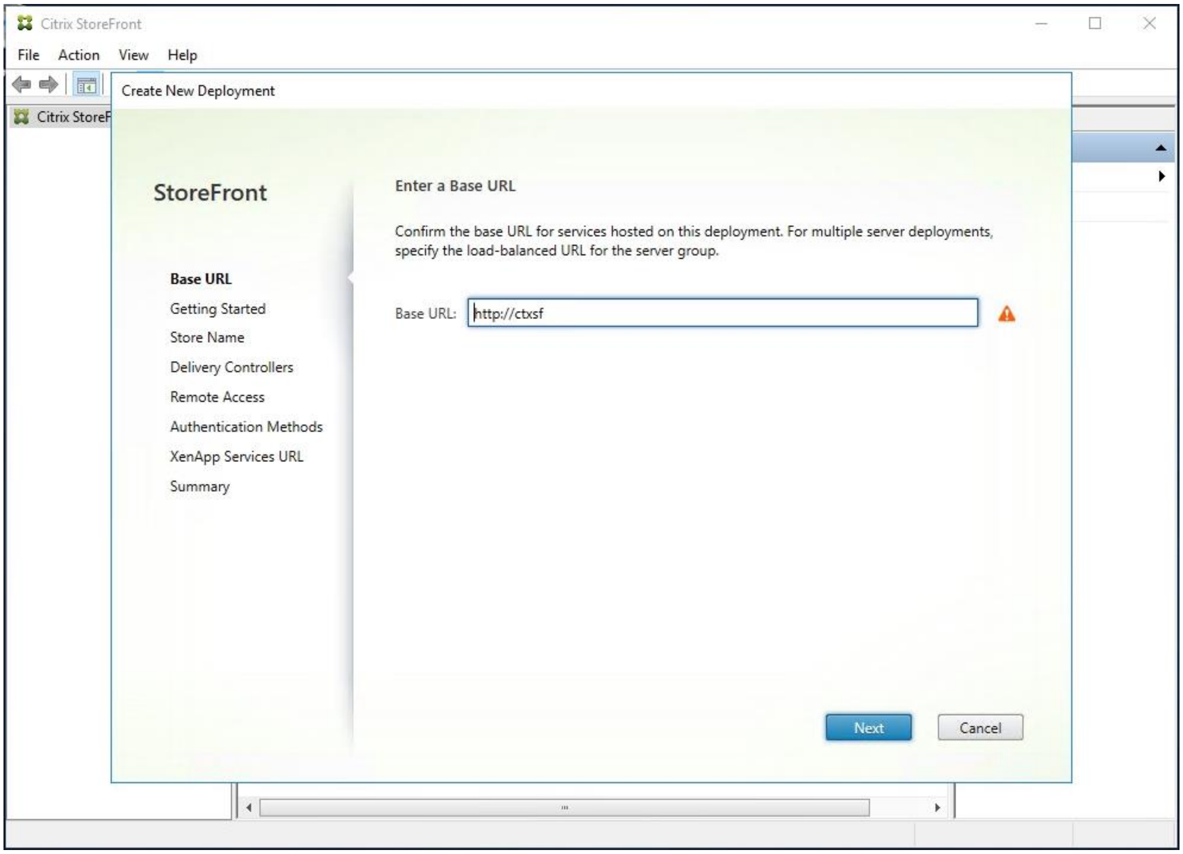

Enter a Base URL

Confirm the base URL for services hosted on this deployment. For multiple server deployments, specify the load-balanced URL for the server group.

Base URL: http://ctxsf

16. Click Next.

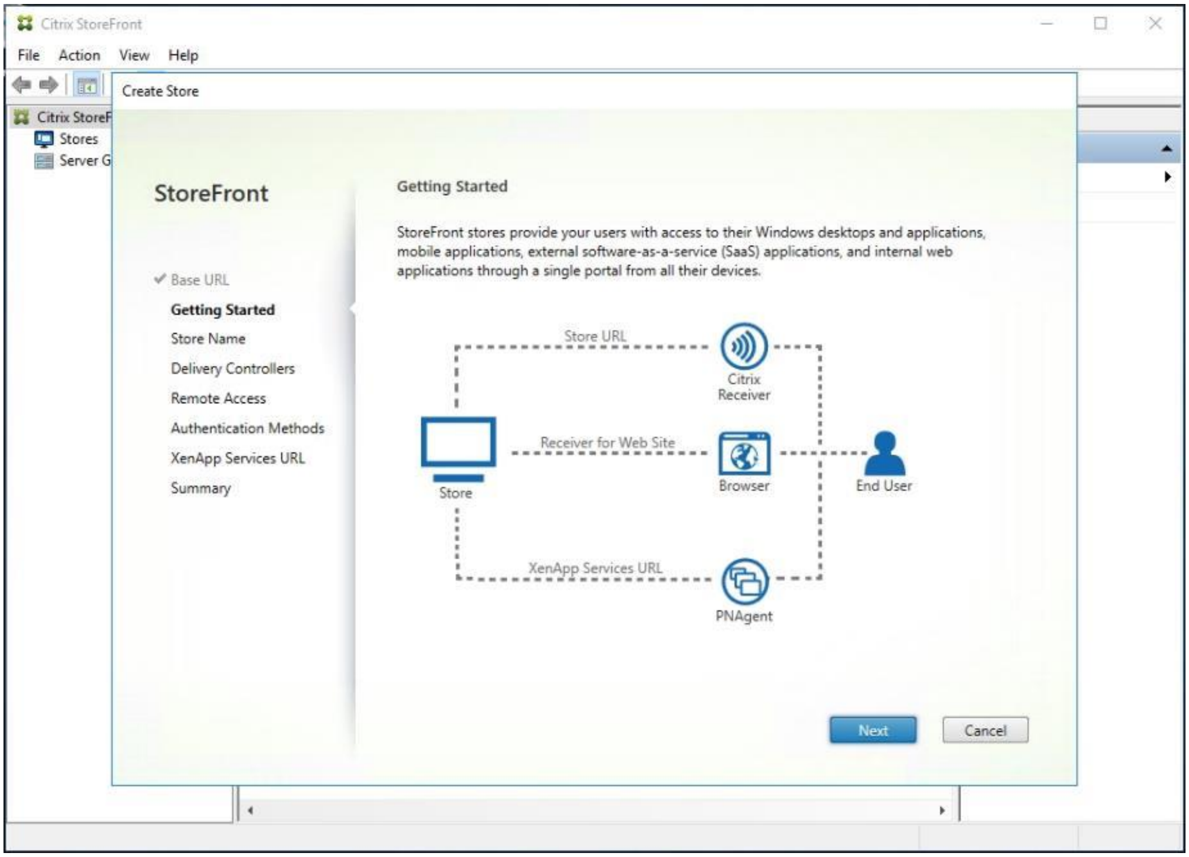

Getting Started

StoreFront stores provide your users with access to their Windows desktops and applications, mobile applications, external software-as-a-service (SaaS) applications, and internal web applications through a single portal from all their devices.

17. Specify a name for your store.

18. Click Next.

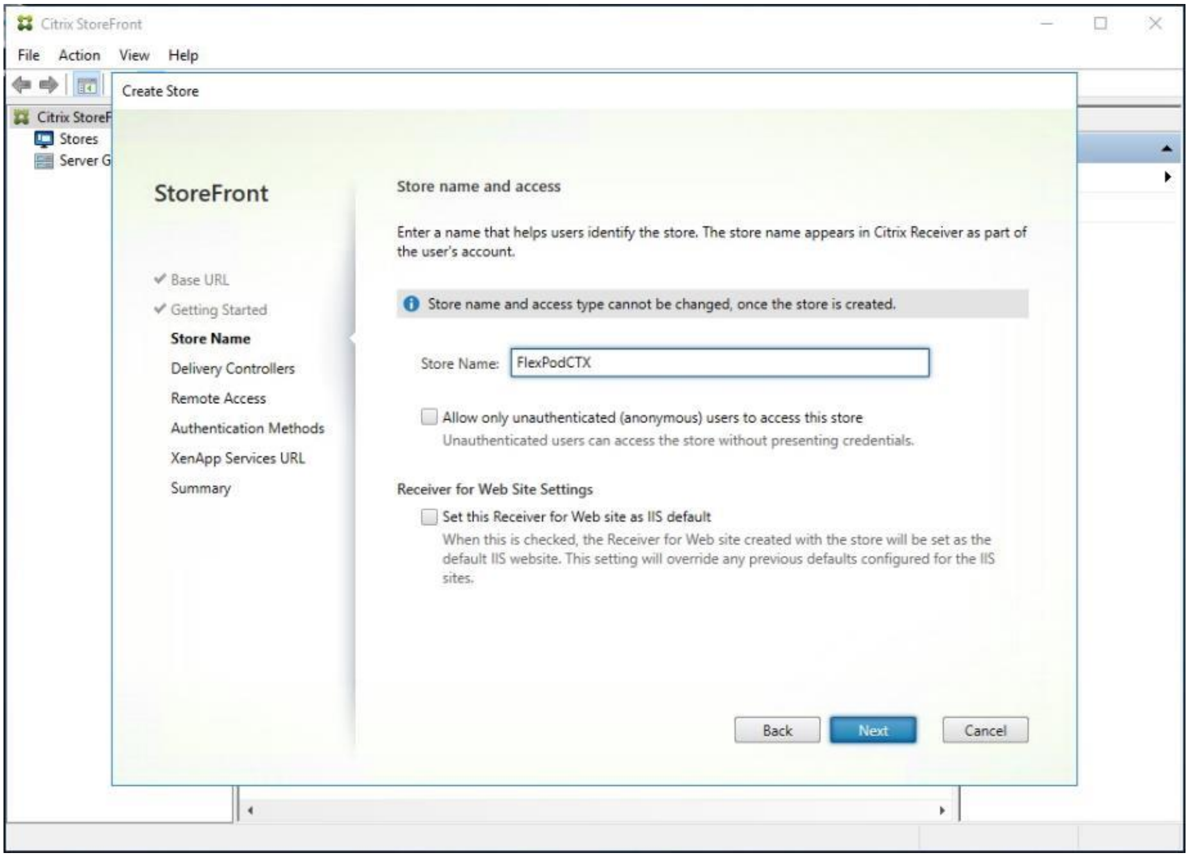

Store name and access

Enter a name that helps users identify the store. The store name appears in Citrix Receiver as part of the user's account.

Store name and access type cannot be changed, once the store is created.

Store Name: FlexPodCTX

□ Allow only unauthenticated (anonymous) users to access this store
Unauthenticated users can access the store without presenting credentials.

Receiver for Web Site Settings

□ Set this Receiver for Web site as IIS default
When this is checked, the Receiver for Web site created with the store will be set as the default IIS website. This setting will override any previous defaults configured for the IIS sites.

19. Add the required Delivery Controllers to the store.

20. Click Next.

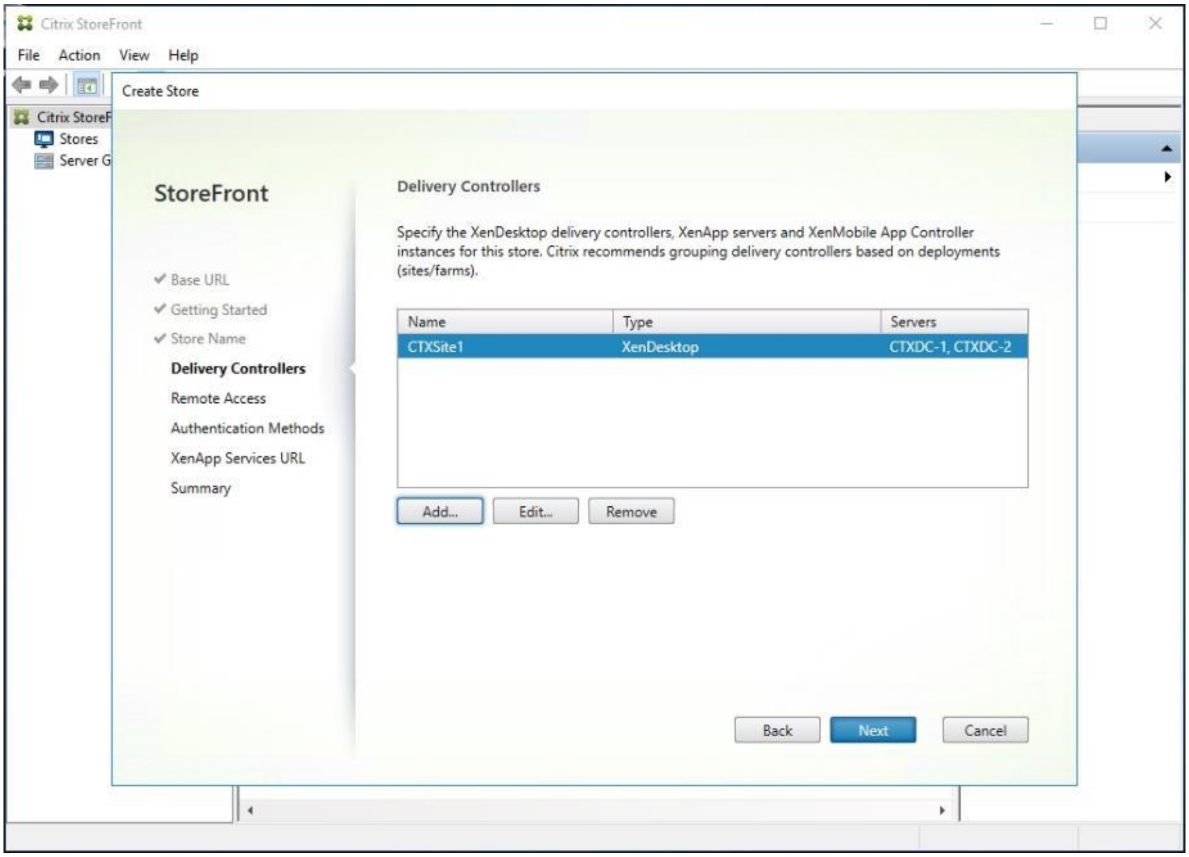

Delivery Controllers

Specify the XenDesktop delivery controllers, XenApp servers and XenMobile App Controller instances for this store. Citrix recommends grouping delivery controllers based on deployments (sites/farms).

Name	Type	Servers
CTXSite1	XenDesktop	CTXDC-1, CTXDC-2

21. Specify how connecting users can access the resources, in this environment only local users on the internal network are able to access the store.

22. Click Next.

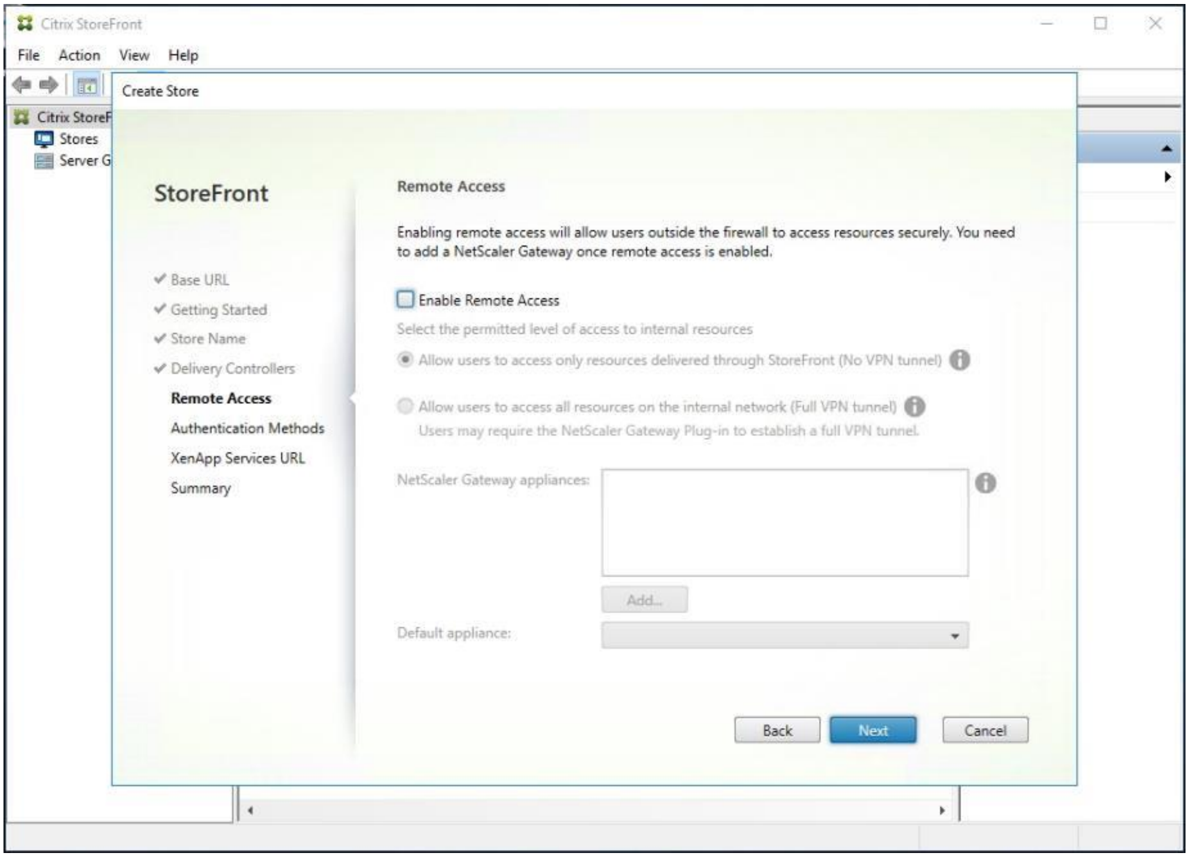

Remote Access

Enabling remote access will allow users outside the firewall to access resources securely. You need to add a NetScaler Gateway once remote access is enabled.

☐ Enable Remote Access

Select the permitted level of access to internal resources

● Allow users to access only resources delivered through StoreFront (No VPN tunnel)

○ Allow users to access all resources on the internal network (Full VPN tunnel)
Users may require the NetScaler Gateway Plug-in to establish a full VPN tunnel.

NetScaler Gateway appliances:

Default appliance:

23. On the ”Authentication Methods” page, select the methods your users will use to authenticate to the store. The following methods were configured in this deployment:

    – Username and password: Users enter their credentials and are authenticated when they access their stores.

    – Domain passthrough: Users authenticate to their domain-joined Windows computers and their credentials are used to log them on automatically when they access their stores.

24. Click Next.

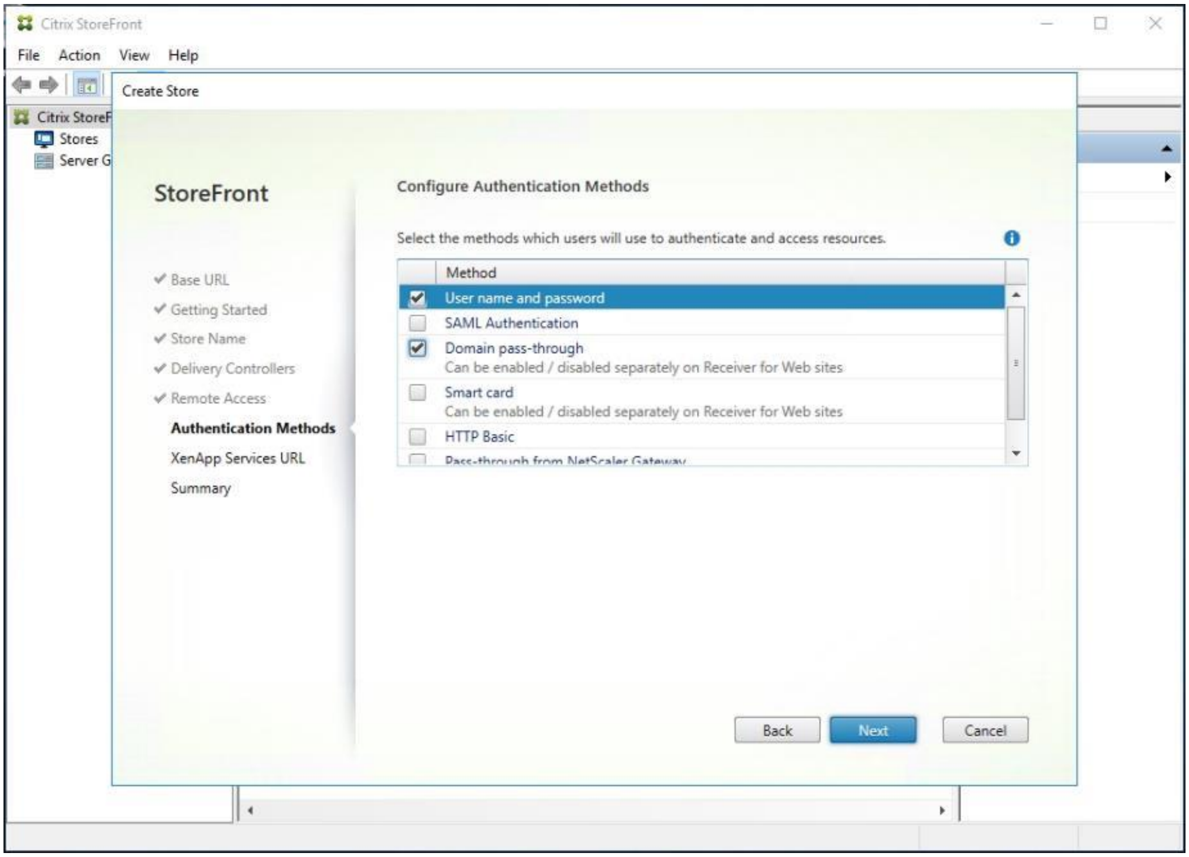

Configure Authentication Methods

Select the methods which users will use to authenticate and access resources.

	Method
☑	User name and password
□	SAML Authentication
☑	Domain pass-through Can be enabled / disabled separately on Receiver for Web sites
□	Smart card Can be enabled / disabled separately on Receiver for Web sites
□	HTTP Basic
□	Pass-through from NetScaler Gateway

25. Configure the XenApp Service URL for users who use PNAgent to access the applications and desktops.

26. Click Create.

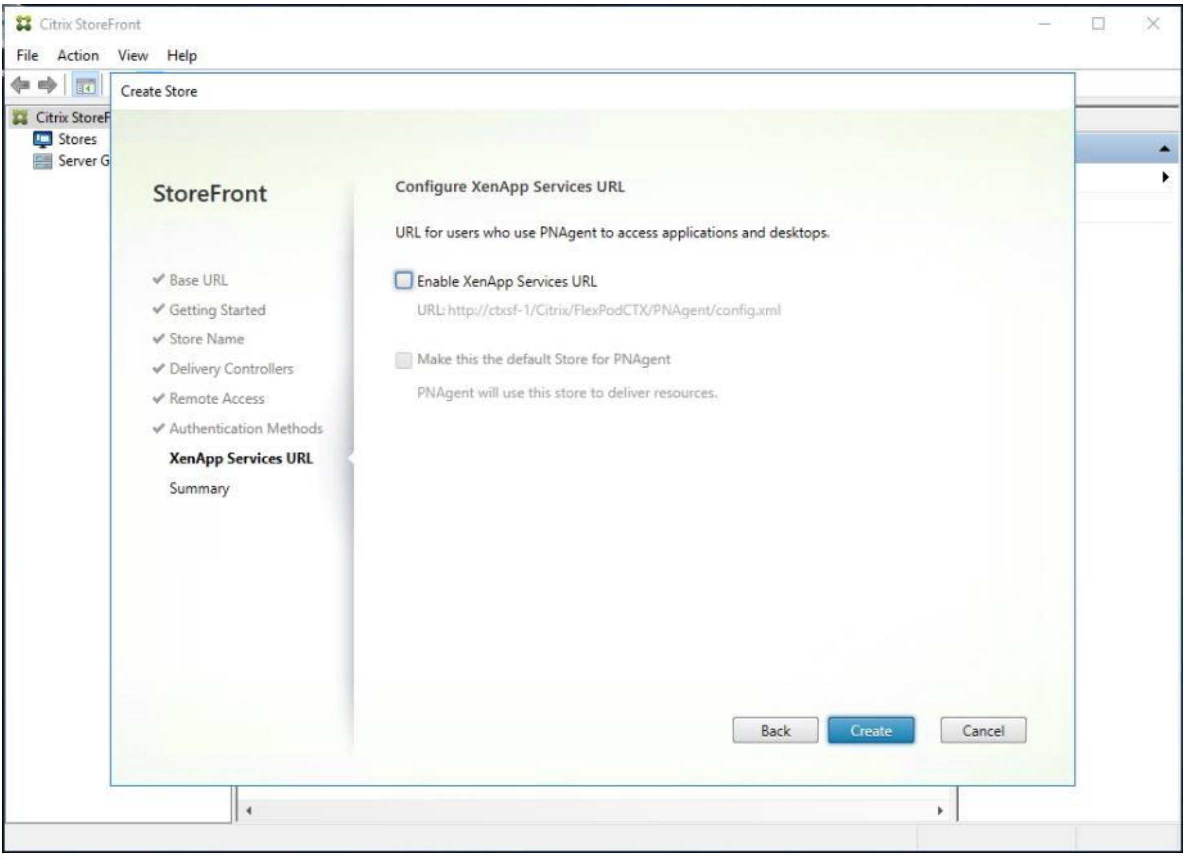

Configure XenApp Services URL

URL for users who use PNAgent to access applications and desktops.

□ Enable XenApp Services URL

URL: http://ctxsf-1/Citrix/FlexPodCTX/PNAgent/config.xml

□ Make this the default Store for PNAgent

PNAgent will use this store to deliver resources.

.

27. After creating the store click Finish.

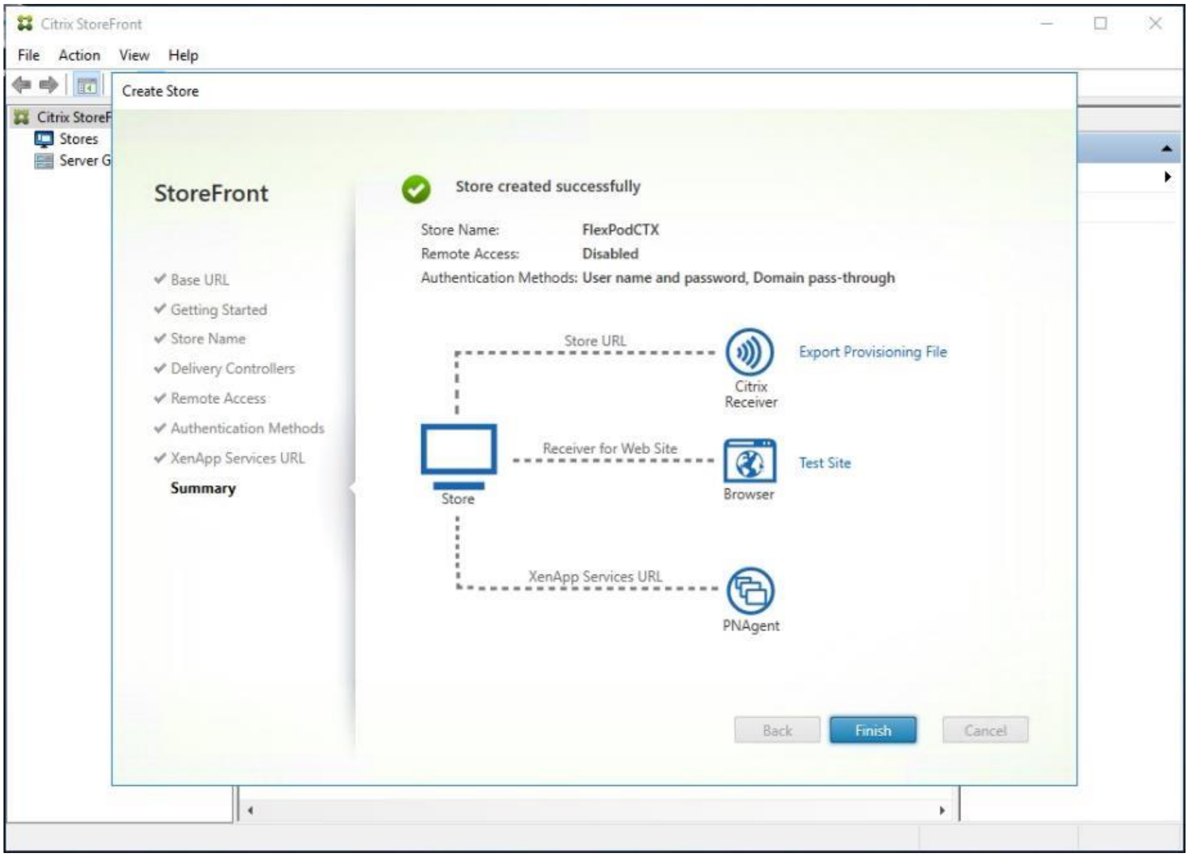

Store created successfully

Store Name: FlexPodCTX
Remote Access: Disabled
Authentication Methods: User name and password, Domain pass-through

Additional StoreFront Configuration

After the first StoreFront server is completely configured and the Store is operational, you can add additional servers.

To configure additional StoreFront server, complete the following steps:

1. To begin the installation of the second StoreFront, connect to the second StoreFront server and launch the installer from the Citrix XenDesktop 7.15 ISO.

2. Click Start.

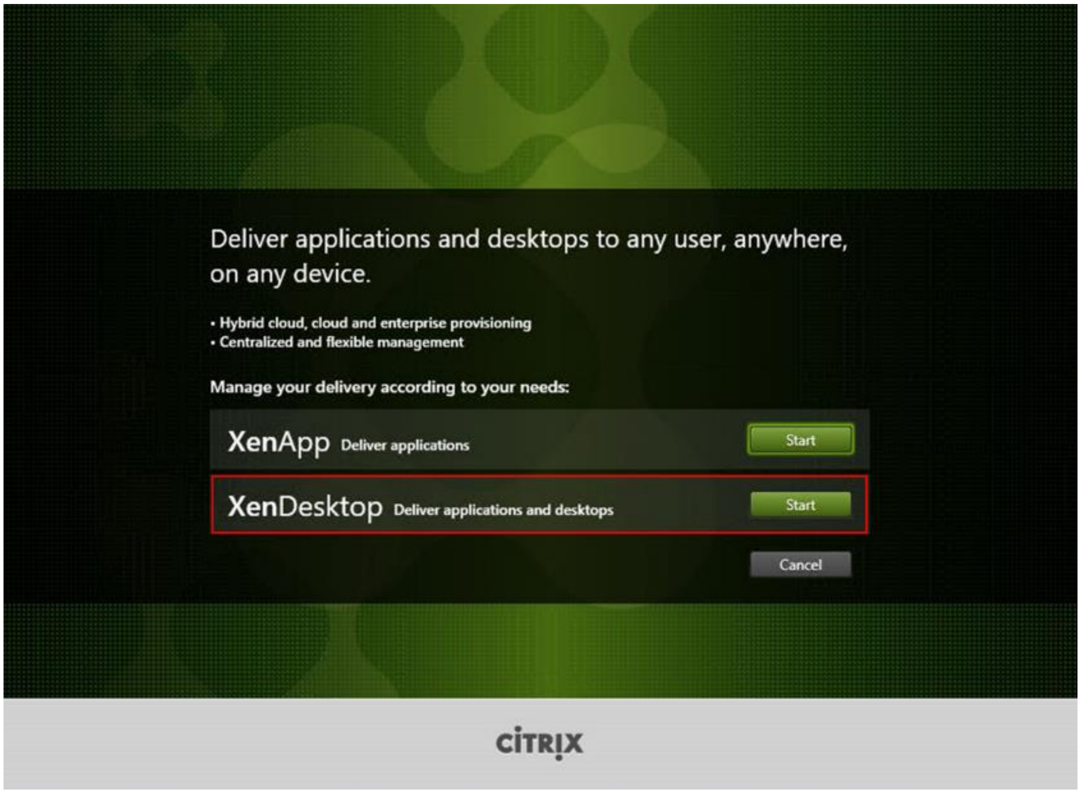

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications

XenDesktop Deliver applications and desktops

3. Click Extended Deployment Citrix StoreFront.

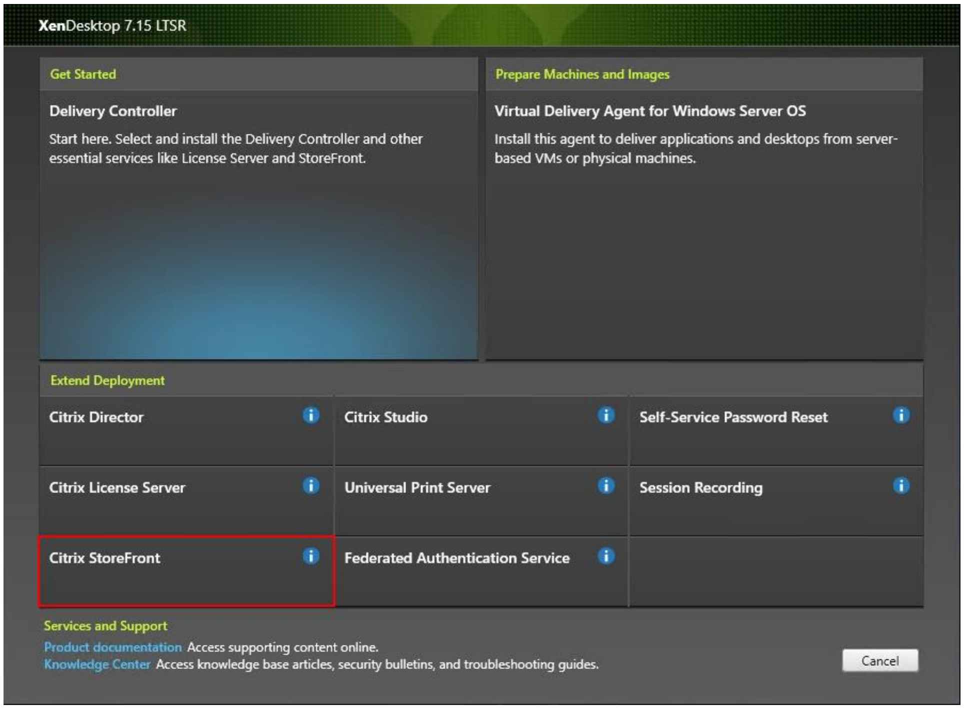

XenDesktop 7.15 LTSR

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Extend Deployment

Services and Support

Product documentation Access supporting content online.

Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. Repeat the same steps used to install the first StoreFront.

5. Review the Summary configuration.

6. Click Install.

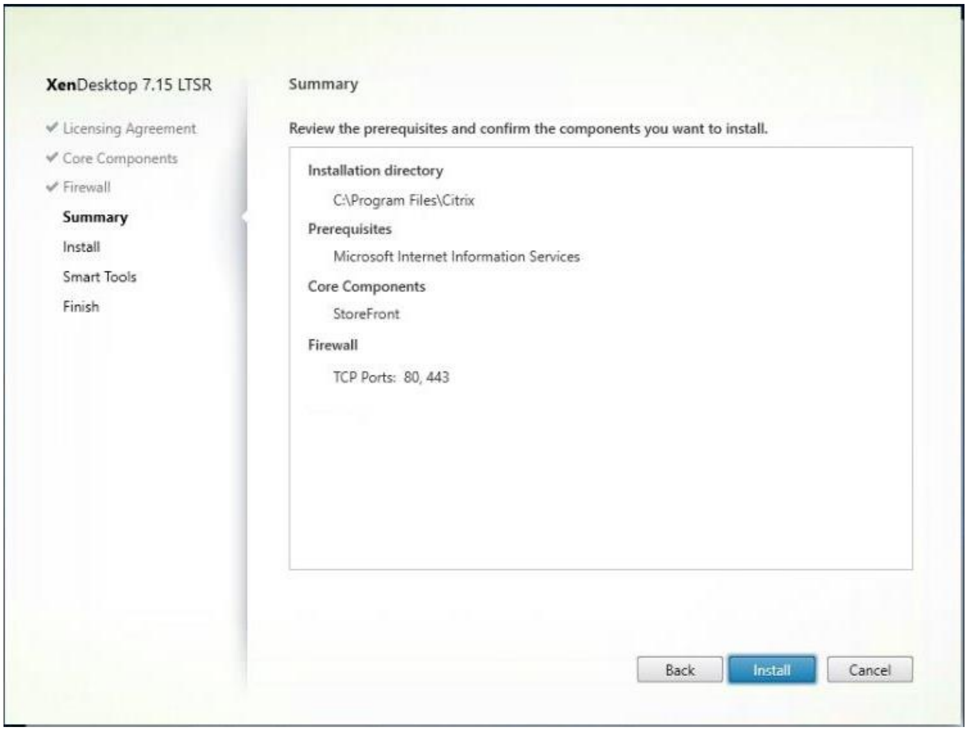

XenDesktop 7.15 LTSR

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory

C:\Program Files\Citrix

Prerequisites

Microsoft Internet Information Services

Core Components

StoreFront

Firewall

TCP Ports: 80, 443

7. (Optional) Click “I want to participate in Call Home.”

8. Click Next.

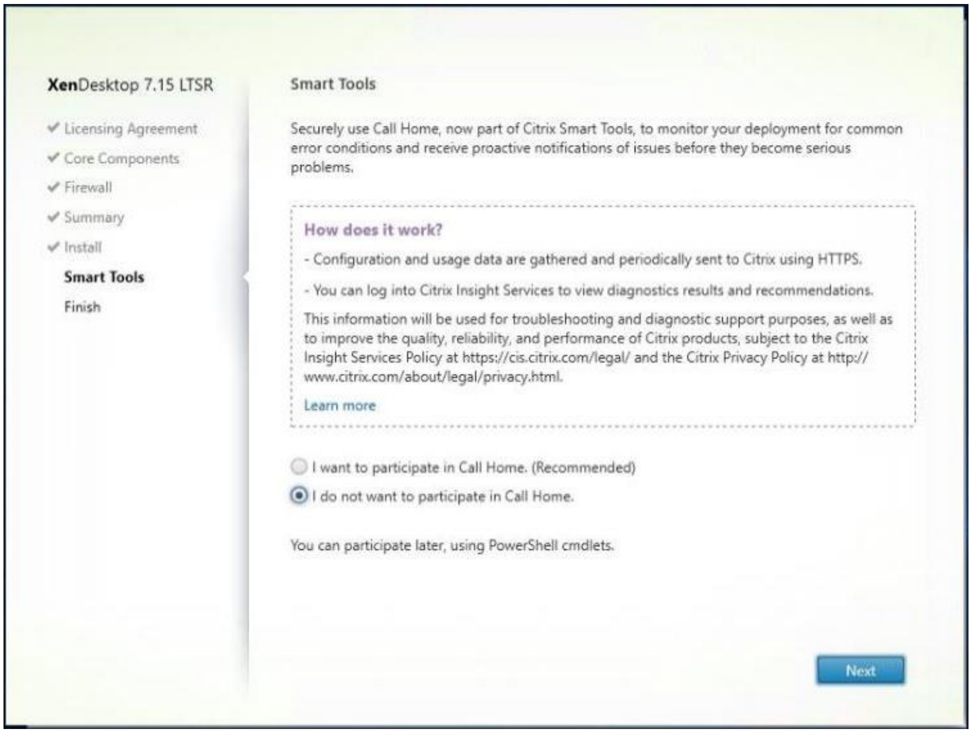

XenDesktop 7.15 LTSR

Smart Tools

Securely use Call Home, now part of Citrix Smart Tools, to monitor your deployment for common error conditions and receive proactive notifications of issues before they become serious problems.

How does it work?

- Configuration and usage data are gathered and periodically sent to Citrix using HTTPS.

- You can log into Citrix Insight Services to view diagnostics results and recommendations.

This information will be used for troubleshooting and diagnostic support purposes, as well as to improve the quality, reliability, and performance of Citrix products, subject to the Citrix Insight Services Policy at https://cis.citrix.com/legal/ and the Citrix Privacy Policy at http://www.citrix.com/about/legal/privacy.html.

Learn more

○ I want to participate in Call Home. (Recommended)

● I do not want to participate in Call Home.

You can participate later, using PowerShell cmdlets.

9. (Optional) check “Open the StoreFront Management Console."

10. Click Finish.

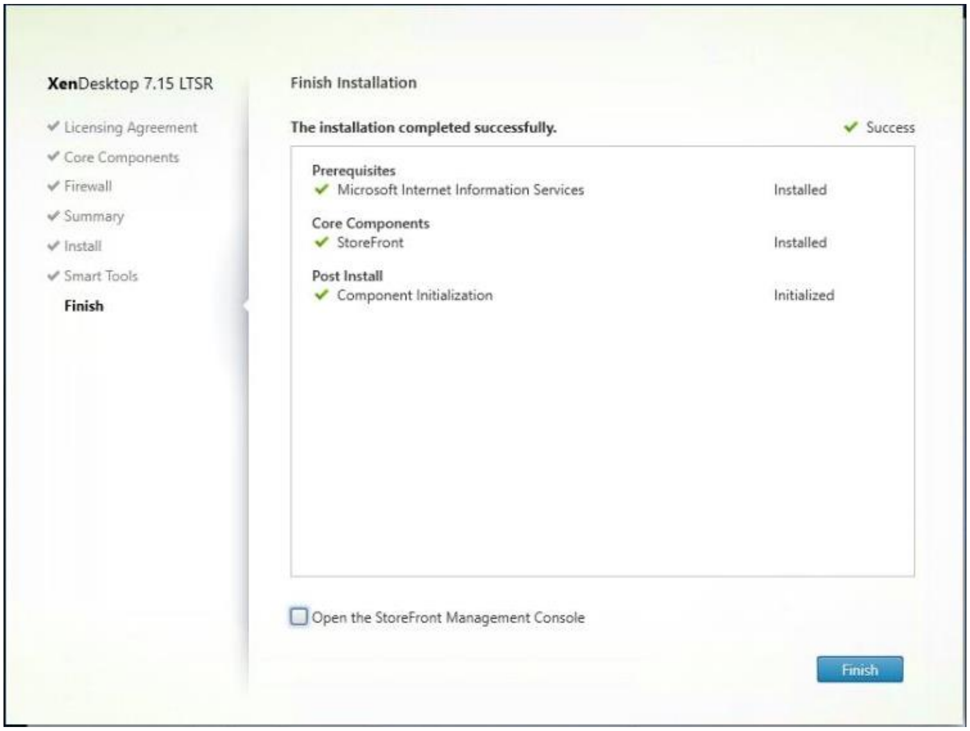

The installation completed successfully. ✓ Success

Prerequisites
✓ Microsoft Internet Information Services Installed

Core Components
✓ StoreFront Installed

Post Install
✓ Component Initialization Initialized

□ Open the StoreFront Management Console

To configure second StoreFront if used, complete the following steps:

1. From the StoreFront Console on the second server select “Join existing server group.”

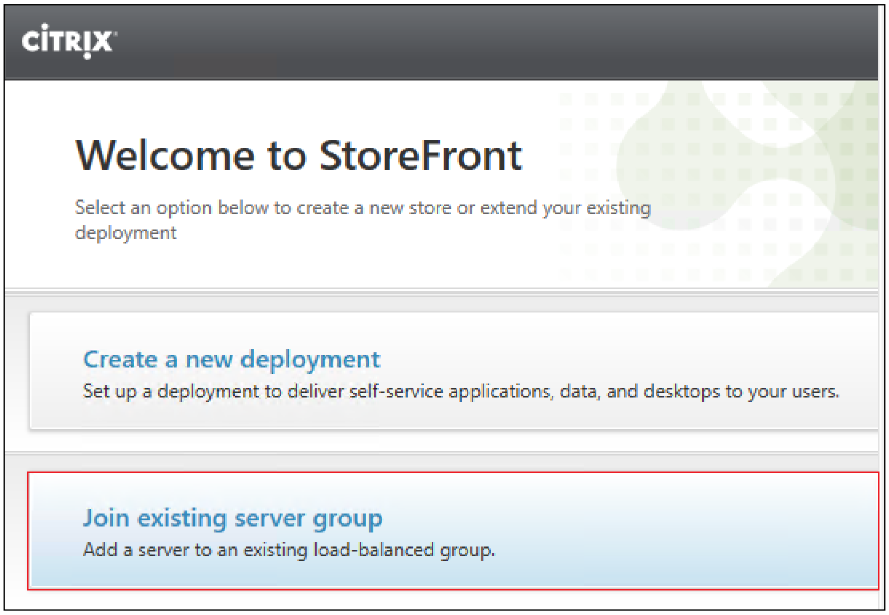

Welcome to StoreFront

Select an option below to create a new store or extend your existing deployment

Create a new deployment
Set up a deployment to deliver self-service applications, data, and desktops to your users.

Join existing server group
Add a server to an existing load-balanced group.

2. In the Join Server Group dialog, enter the name of the first Storefront server.

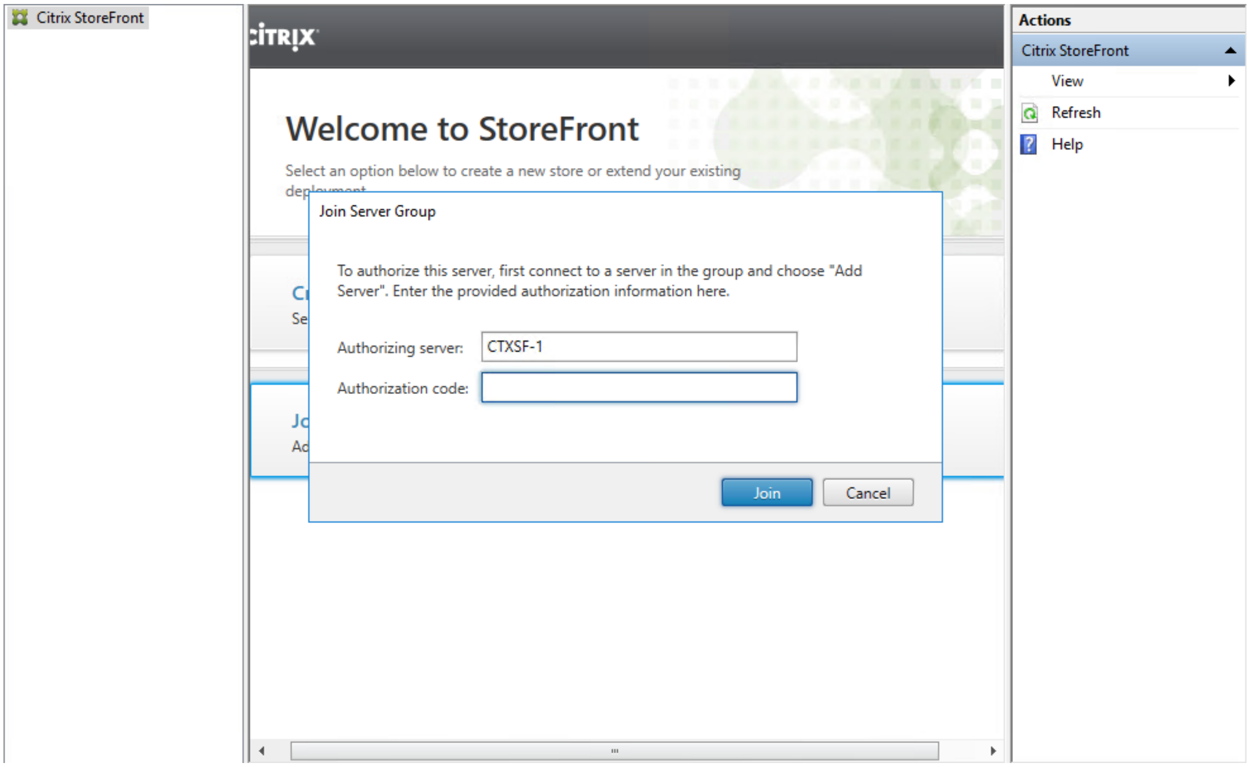

To authorize this server, first connect to a server in the group and choose "Add Server". Enter the provided authorization information here.

Authorizing server: CTXSF-1

Authorization code:

3. Before the additional StoreFront server can join the server group, you must connect to the first Storefront server, add the second server, and obtain the required authorization information.

4. Connect to the first StoreFront server.

5. Using the StoreFront menu on the left, you can scroll through the StoreFront management options.

6. Select Server Group from the menu.

7. To add the second server and generate the authorization information that allows the additional StoreFront server to join the server group, select Add Server from Actions pane.

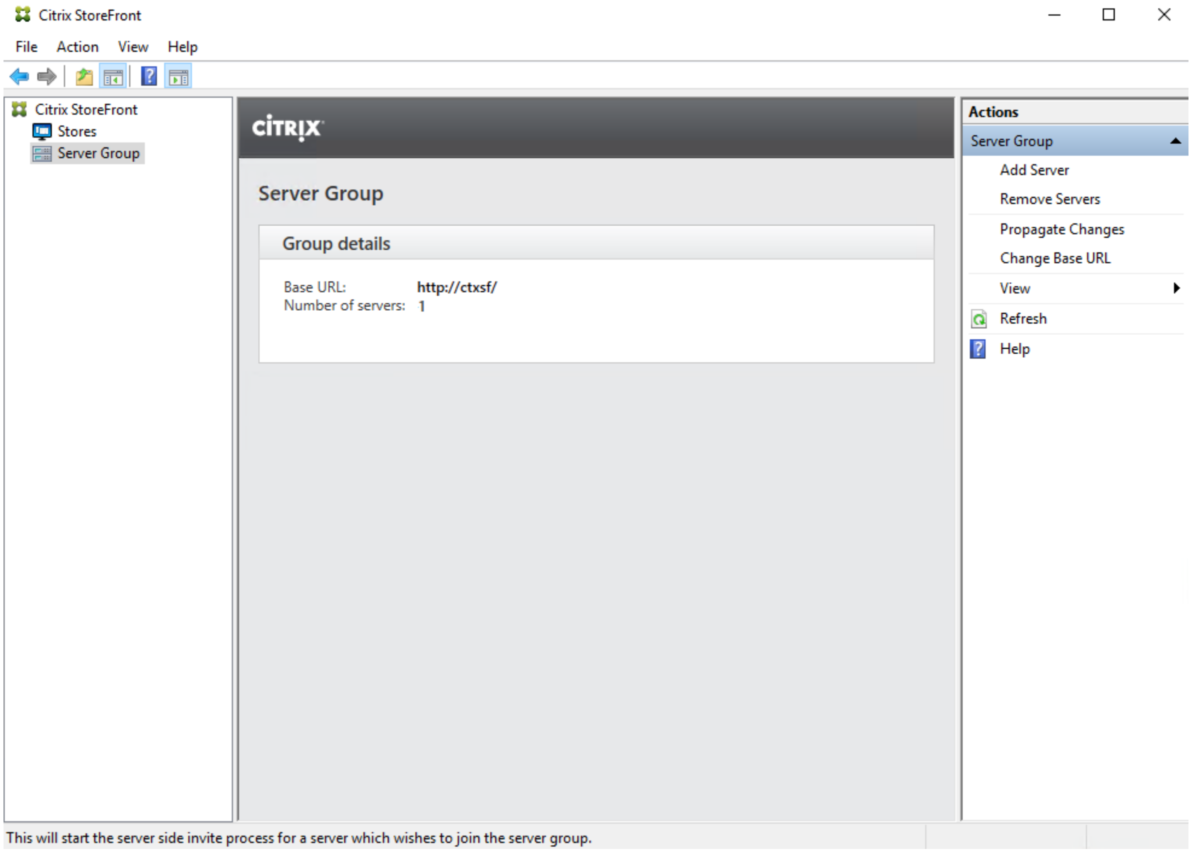

Server Group

Group details

Base URL: http://ctxsf/
Number of servers: 1

This will start the server side invite process for a server which wishes to join the server group.

8. Copy the Authorization code from the Add Server dialog.

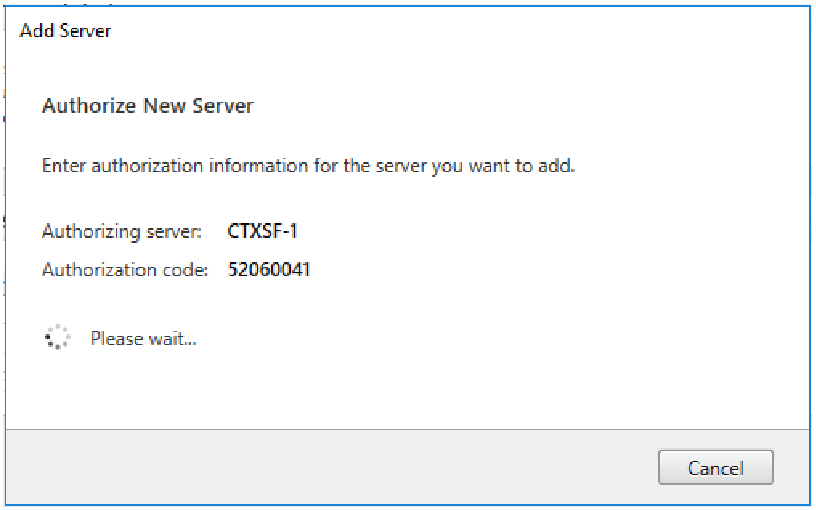

Authorize New Server

Enter authorization information for the server you want to add.

Authorizing server: CTXSF-1
Authorization code: 52060041

Please wait...

9. Connect to the second Storefront server and paste the Authorization code into the Join Server Group dialog.

10. Click Join.

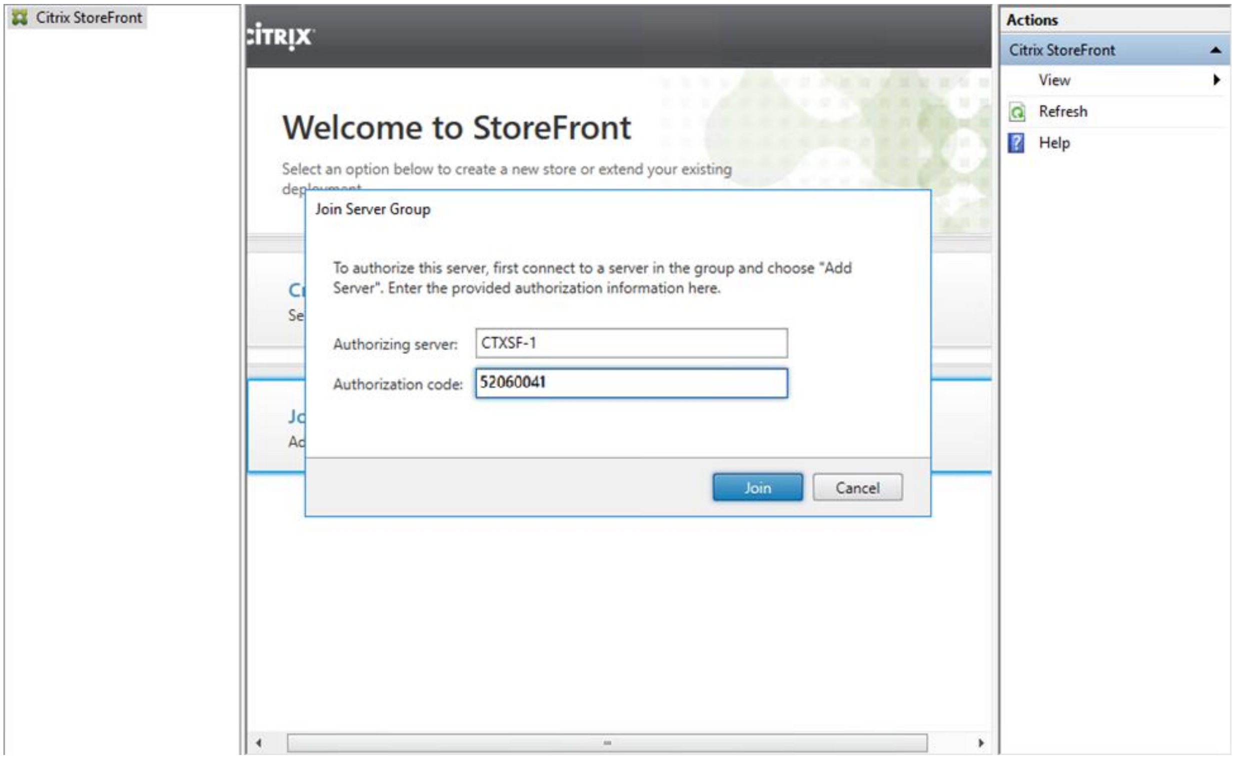

Welcome to StoreFront

Select an option below to create a new store or extend your existing deployment

Join Server Group

To authorize this server, first connect to a server in the group and choose "Add Server". Enter the provided authorization information here.

Authorizing server: CTXSF-1

Authorization code: 52060041

11. A message appears when the second server has joined successfully.

12. Click OK.

13. Verify Server Group status on the first StoreFront Server.

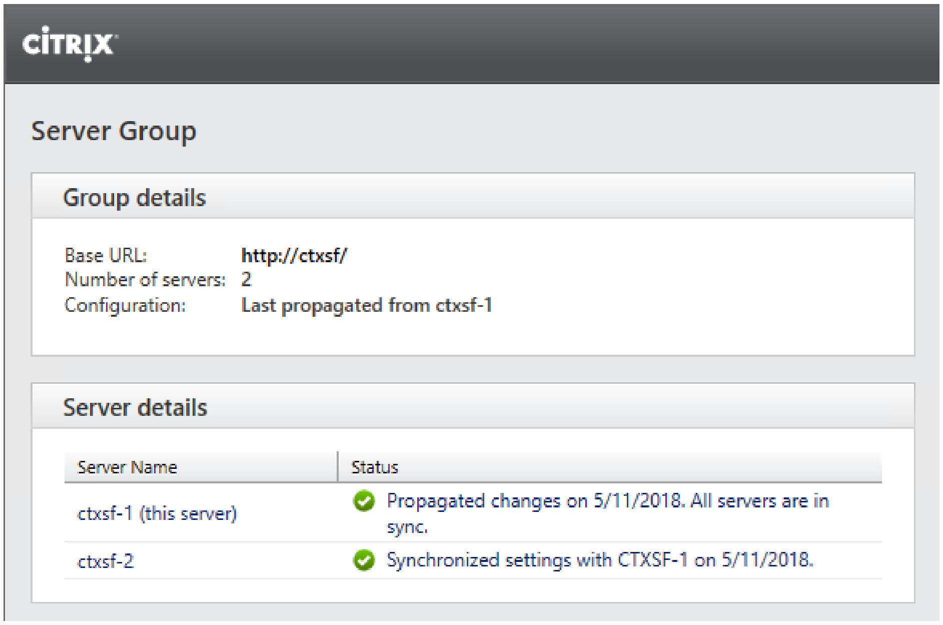

Server Group

Group details

Base URL: http://ctxsf/
Number of servers: 2
Configuration: Last propagated from ctxsf-1

Server details

Server Name	Status
ctxsf-1 (this server)	✓ Propagated changes on 5/11/2018. All servers are in sync.
ctxsf-2	✓ Synchronized settings with CTXSF-1 on 5/11/2018.

The second StoreFront is now in the Server Group.

Installing and Configuring Citrix Provisioning Server 7.15

In most implementations, there is a single vDisk providing the standard image for multiple target devices. Thousands of target devices can use a single vDisk shared across multiple Provisioning Services (PVS) servers in the same farm, simplifying virtual desktop management. This section describes the installation and configuration tasks required to create a PVS implementation.

The PVS server can have many stored vDisks, and each vDisk can be several gigabytes in size. Your streaming performance and manageability can be improved using a RAID array, SAN, or NAS. PVS software and hardware requirements are available in the [Provisioning Services 7.15 document.]()

Prerequisites

Set the following Scope Options on the DHCP server hosting the PVS target machines (for example, VDI, RDS).

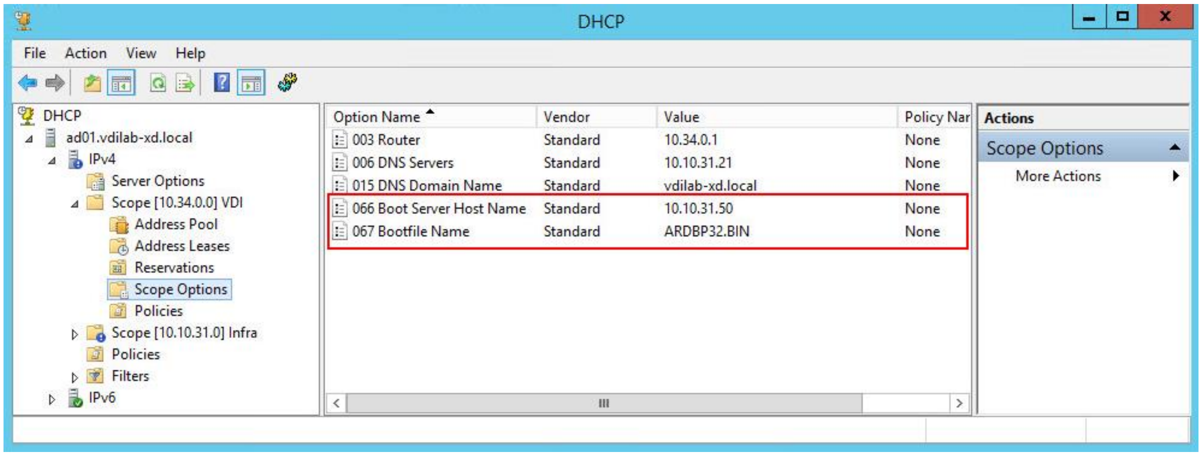

Option Name	Vendor	Value	Policy Nar
003 Router	Standard	10.34.0.1	None
006 DNS Servers	Standard	10.10.31.21	None
015 DNS Domain Name	Standard	vdilab-xd.local	None
066 Boot Server Host Name	Standard	10.10.31.50	None
067 Bootfile Name	Standard	ARDBP32.BIN	None

The Boot Server IP was configured for Load Balancing by NetScaler VPX to support High Availability of TFTP service.

To Configure TFTP Load Balancing, complete the following steps:

1. Create Virtual IP for TFTP Load Balancing.

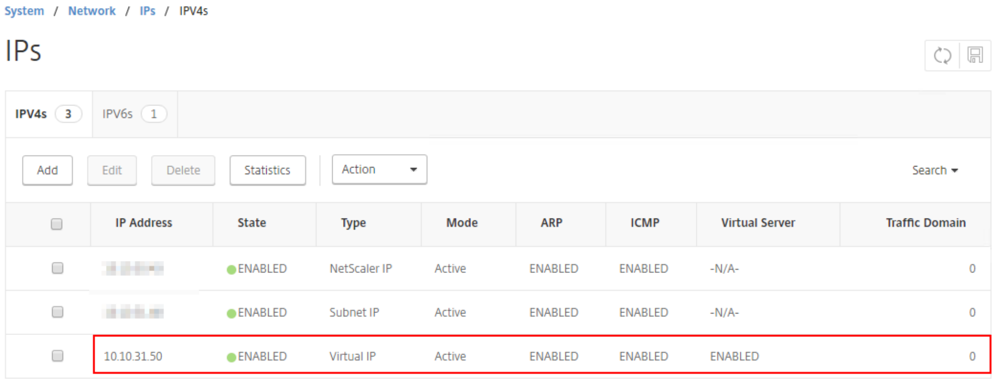

□	IP Address	State	Type	Mode	ARP	ICMP	Virtual Server	Traffic Domain
□	[redacted]	● ENABLED	NetScaler IP	Active	ENABLED	ENABLED	-N/A-	0
□	[redacted]	● ENABLED	Subnet IP	Active	ENABLED	ENABLED	-N/A-	0
□	10.10.31.50	● ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED	0

2. Configure servers that are running TFTP (your Provisioning Servers).

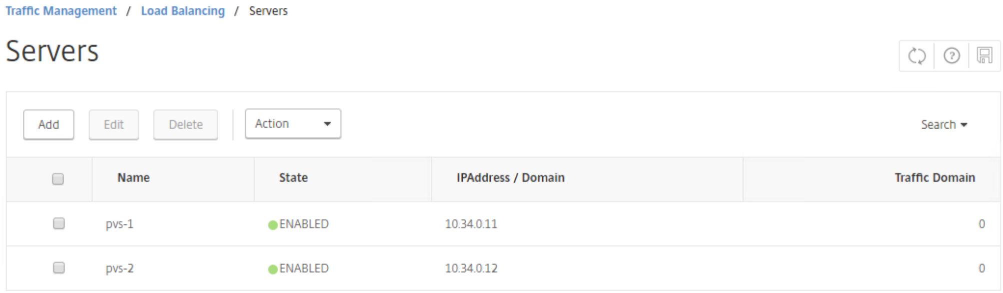

Traffic Management / Load Balancing / Servers

Servers

	Name	State	IPAddress / Domain	Traffic Domain
□	pvs-1	● ENABLED	10.34.0.11	0
□	pvs-2	● ENABLED	10.34.0.12	0

3. Define TFTP service for the servers (Monitor used: **udp-ecv**).

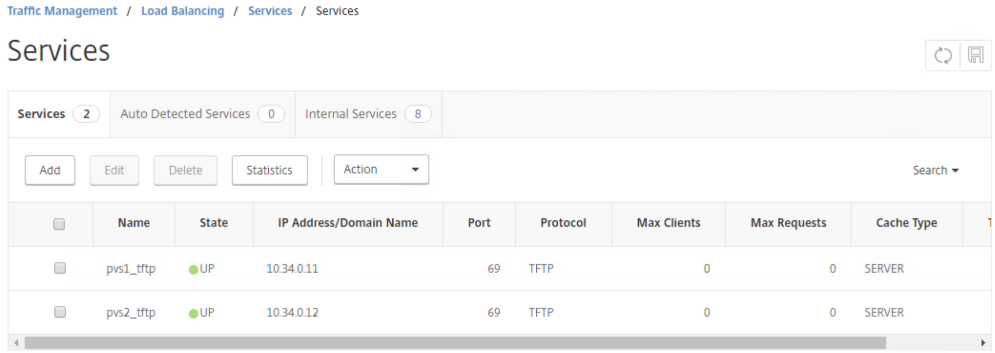

Traffic Management / Load Balancing / Services / Services

Services

Services 2 | Auto Detected Services 0 | Internal Services 8

	Name	State	IP Address/Domain Name	Port	Protocol	Max Clients	Max Requests	Cache Type	T
□	pvs1_tftp	● UP	10.34.0.11	69	TFTP	0	0	SERVER	
□	pvs2_tftp	● UP	10.34.0.12	69	TFTP	0	0	SERVER	

4. Configure TFTP for load balancing.

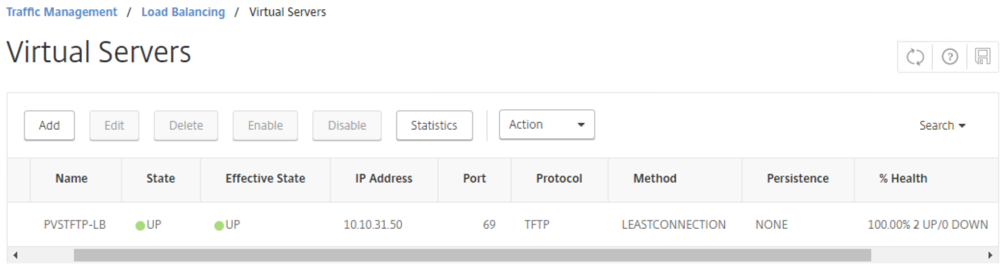

Traffic Management / Load Balancing / Virtual Servers

Virtual Servers

Name	State	Effective State	IP Address	Port	Protocol	Method	Persistence	% Health
PVSTFTP-LB	● UP	● UP	10.10.31.50	69	TFTP	LEASTCONNECTION	NONE	100.00% 2 UP/0 DOWN

As a Citrix best practice cited in this CTX article, apply the following registry setting both the PVS servers and target machines:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TCPIP\Parameters\
Key: "DisableTaskOffload" (dword)
Value: "1"

Only one MS SQL database is associated with a farm. You can choose to install the Provisioning Services database software on an existing SQL database, if that machine can communicate with all Provisioning Servers within the farm, or with a new SQL Express database machine, created using the SQL Express software that is free from Microsoft.

The following databases are supported: Microsoft SQL Server 2008 SP3 through 2016 (x86, x64, and Express editions). Microsoft SQL 2016 was installed separately for this CVD.

High availability will be available for the databases once added to the SQL AlwaysOn Availability Group CTX201203

To install and configure Citrix Provisioning Service 7.15, complete the following steps:

1. Insert the Citrix Provisioning Services 7.15 ISO and let AutoRun launch the installer.

2. Click the Console Installation button.

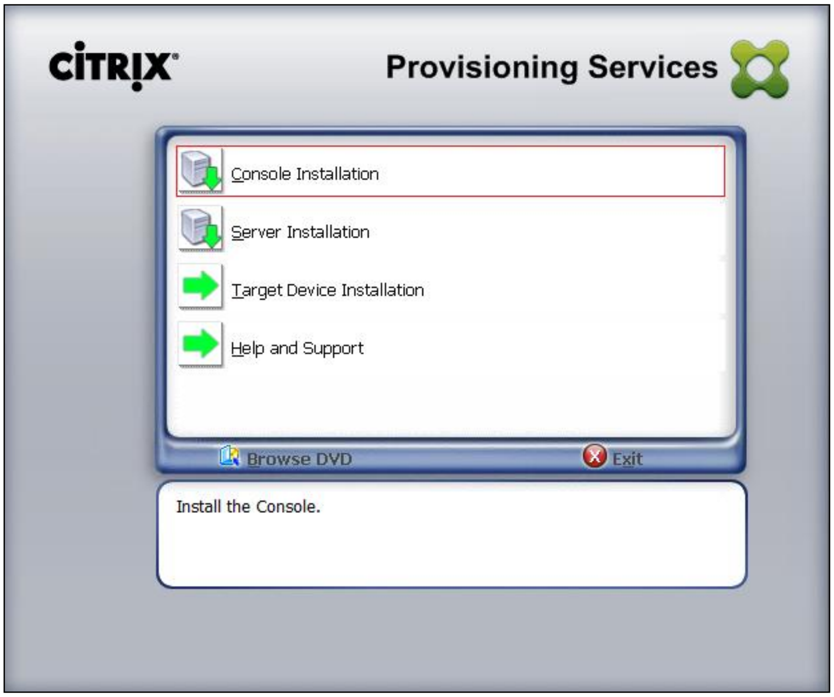

Install the Console.

3. Click Install to install the required prerequisites.

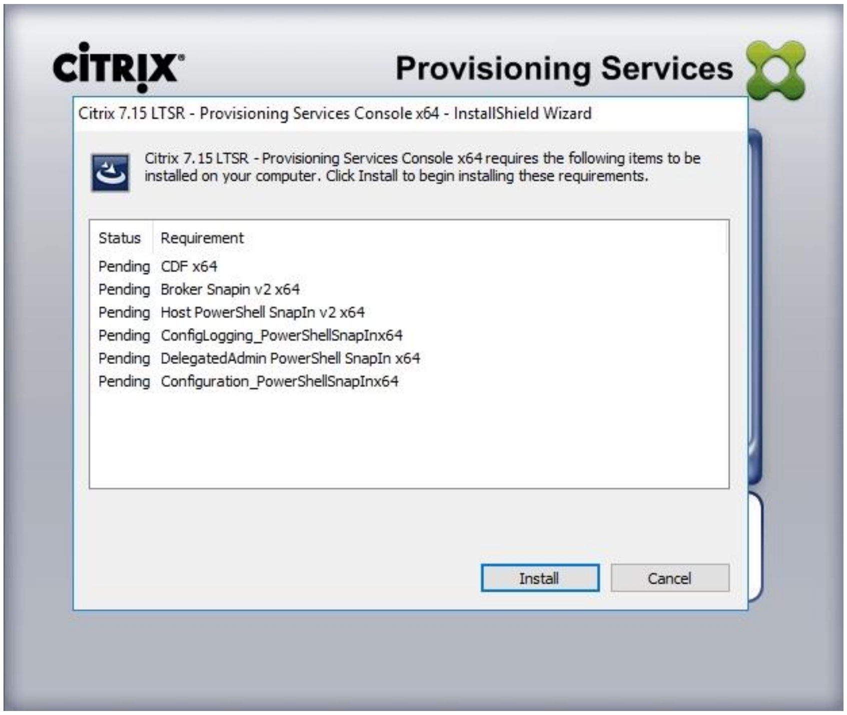

Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard

Citrix 7.15 LTSR - Provisioning Services Console x64 requires the following items to be installed on your computer. Click Install to begin installing these requirements.

Status	Requirement
Pending	CDF x64
Pending	Broker Snapin v2 x64
Pending	Host PowerShell SnapIn v2 x64
Pending	ConfigLogging_PowerShellSnapInx64
Pending	DelegatedAdmin PowerShell SnapIn x64
Pending	Configuration_PowerShellSnapInx64

4. Click Next to start console installation.

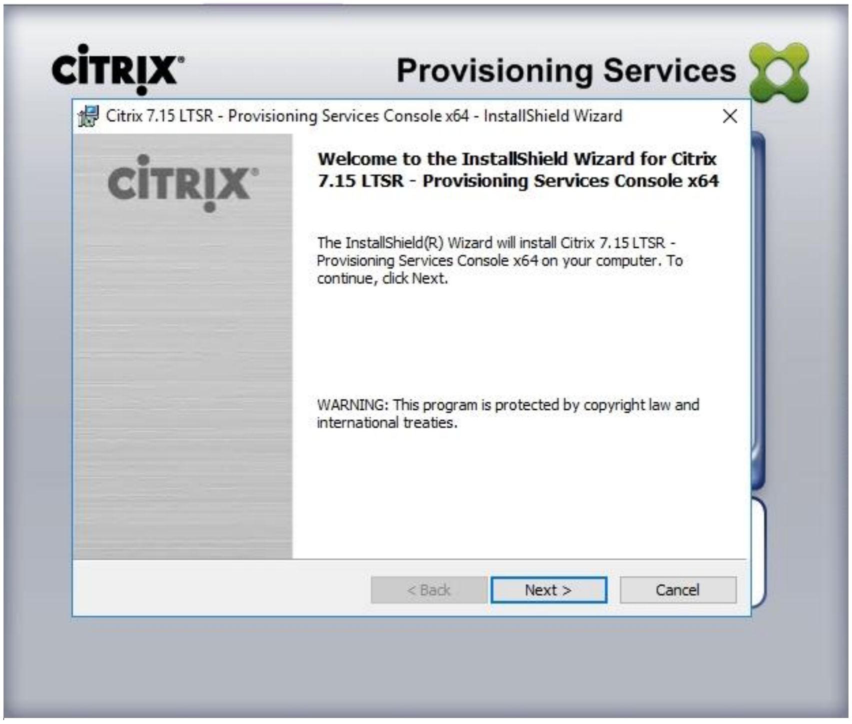

Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard

Welcome to the InstallShield Wizard for Citrix 7.15 LTSR - Provisioning Services Console x64

The InstallShield(R) Wizard will install Citrix 7.15 LTSR - Provisioning Services Console x64 on your computer. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

5. Read the Citrix License Agreement.

6. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”

7. Click Next.

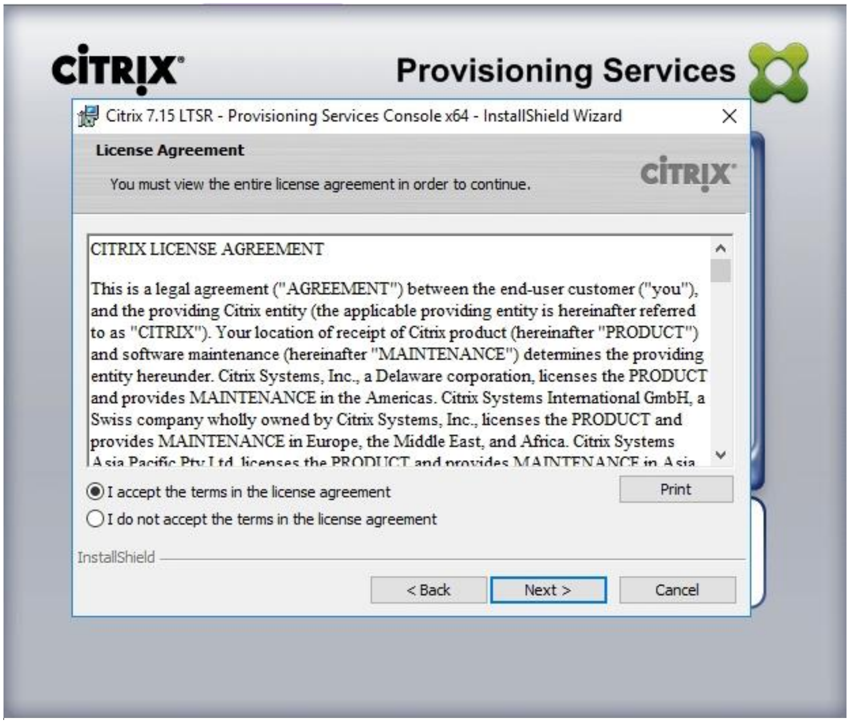

Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard

License Agreement

You must view the entire license agreement in order to continue.

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of Citrix product (hereinafter "PRODUCT") and software maintenance (hereinafter "MAINTENANCE") determines the providing entity hereunder. Citrix Systems, Inc., a Delaware corporation, licenses the PRODUCT and provides MAINTENANCE in the Americas. Citrix Systems International GmbH, a Swiss company wholly owned by Citrix Systems, Inc., licenses the PRODUCT and provides MAINTENANCE in Europe, the Middle East, and Africa. Citrix Systems Asia Pacific Pty Ltd. licenses the PRODUCT and provides MAINTENANCE in Asia

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

8. Optionally provide User Name and Organization.

9. Click Next.

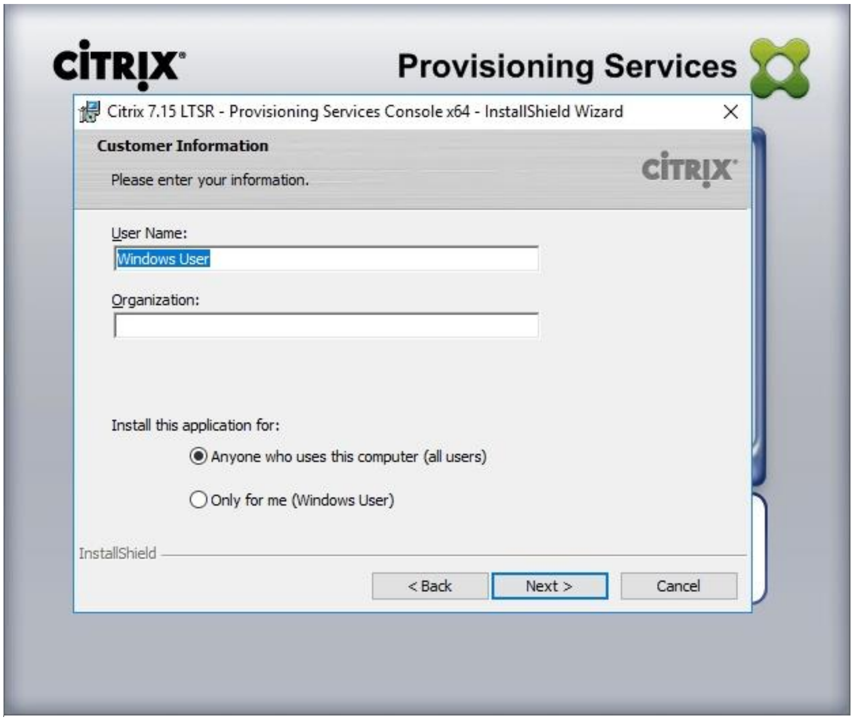

Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard

Customer Information

Please enter your information.

User Name: Windows User

Organization:

Install this application for:

● Anyone who uses this computer (all users)

○ Only for me (Windows User)

10. Accept the default path.

11. Click Next.

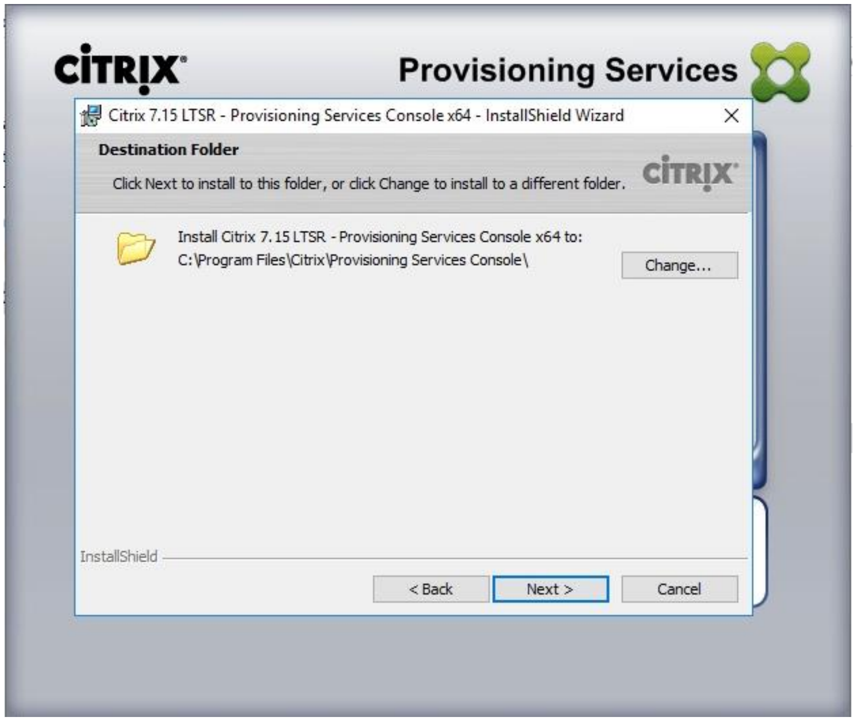

Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard

Destination Folder

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

Install Citrix 7.15 LTSR - Provisioning Services Console x64 to:
C:\Program Files\Citrix\Provisioning Services Console\

12. Click Install to start the console installation.

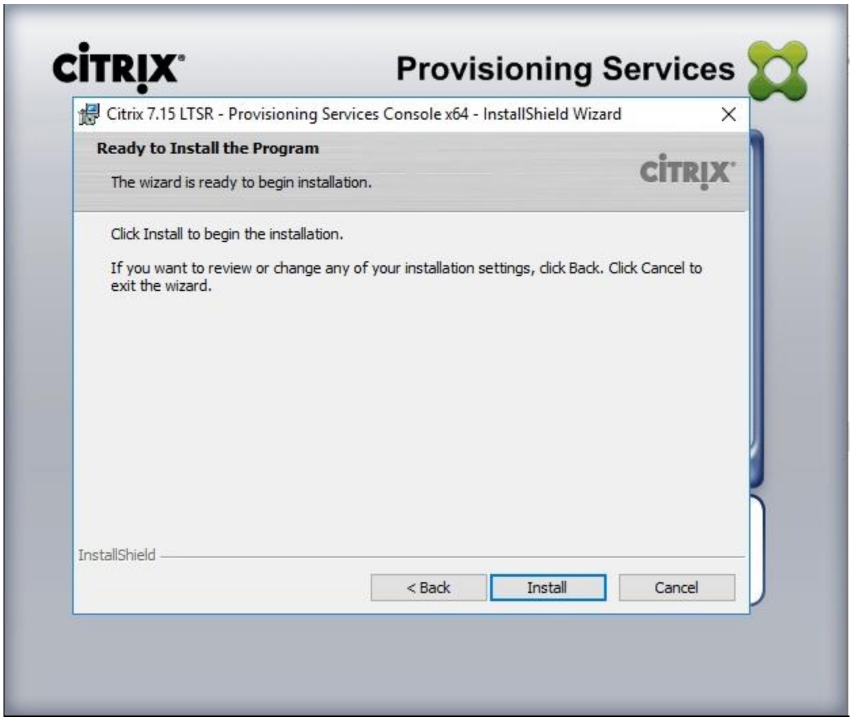
CITRIX
Provisioning Services
Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard
CITRIX
InstallShield
< Back
Install
Cancel

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

13. Click Finish after successful installation.

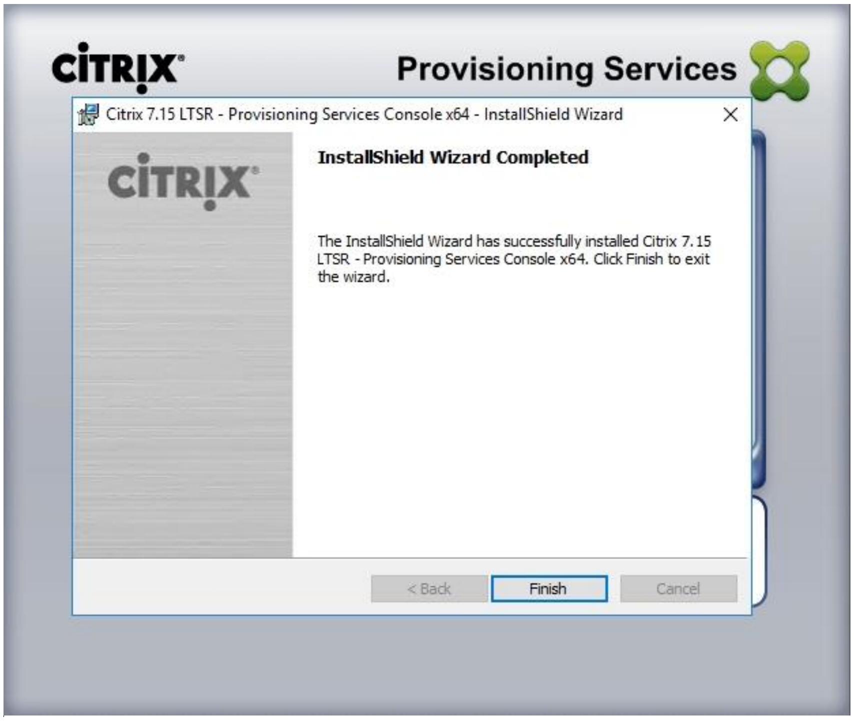
CITRIX
Provisioning Services
Citrix 7.15 LTSR - Provisioning Services Console x64 - InstallShield Wizard
CITRIX
< Back
Finish
Cancel

InstallShield Wizard Completed

The InstallShield Wizard has successfully installed Citrix 7.15 LTSR - Provisioning Services Console x64. Click Finish to exit the wizard.

14. From the main installation screen, select Server Installation.

15. The installation wizard will check to resolve dependencies and then begin the PVS server installation process.

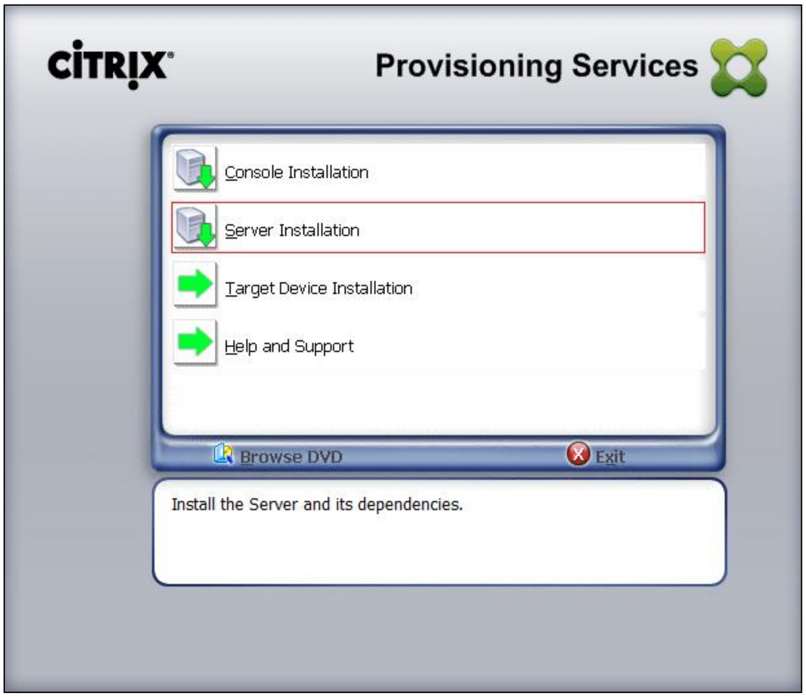

Install the Server and its dependencies.

16. Click Install on the prerequisites dialog.

17. Click Yes when prompted to install the SQL Native Client.

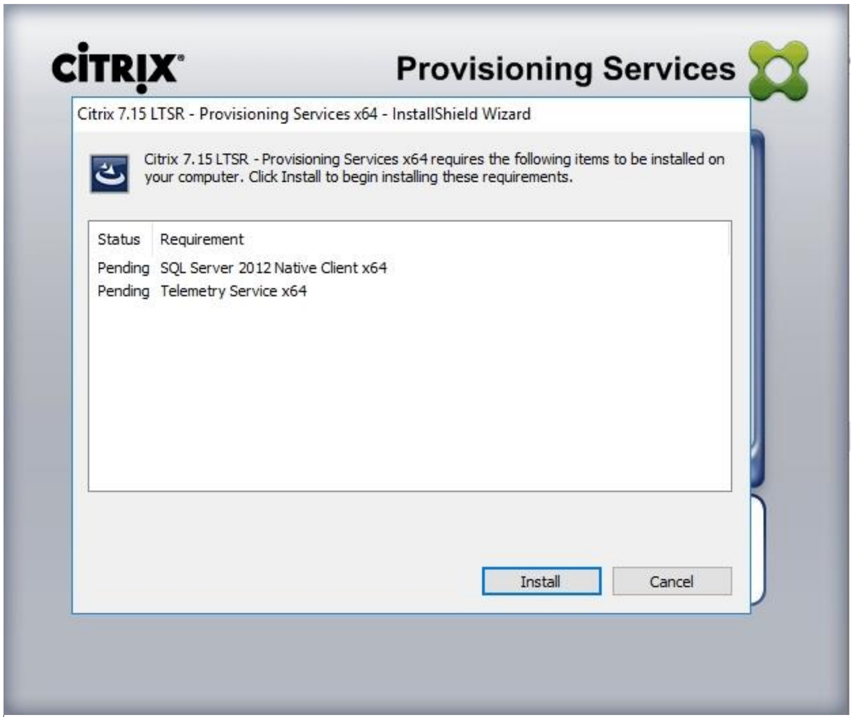

Citrix 7.15 LTSR - Provisioning Services x64 - InstallShield Wizard

Citrix 7.15 LTSR - Provisioning Services x64 requires the following items to be installed on your computer. Click Install to begin installing these requirements.

Status	Requirement
Pending	SQL Server 2012 Native Client x64
Pending	Telemetry Service x64

18. Click Next when the Installation wizard starts.

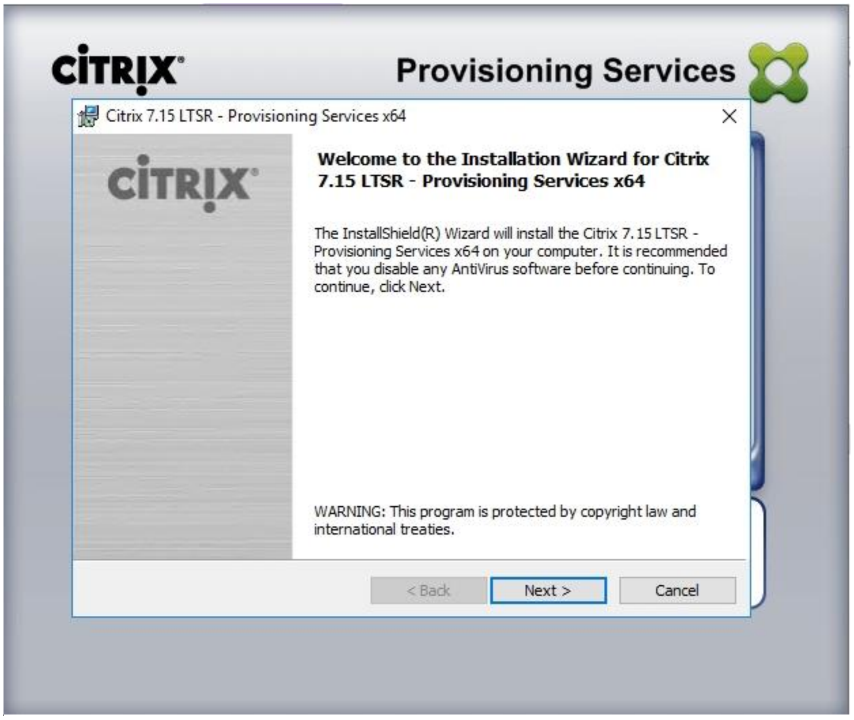

Citrix 7.15 LTSR - Provisioning Services x64

Welcome to the Installation Wizard for Citrix 7.15 LTSR - Provisioning Services x64

The InstallShield(R) Wizard will install the Citrix 7.15 LTSR - Provisioning Services x64 on your computer. It is recommended that you disable any AntiVirus software before continuing. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

19. Review the license agreement terms.

20. If acceptable, select the radio button labeled "I accept the terms in the license agreement."

21. Click Next.

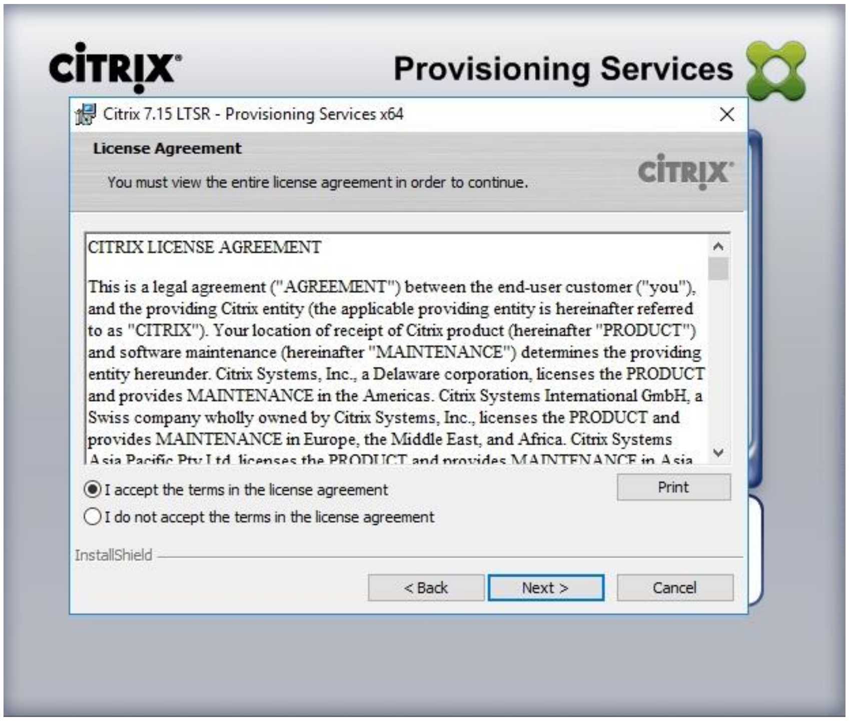

Citrix 7.15 LTSR - Provisioning Services x64

License Agreement

You must view the entire license agreement in order to continue.

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of Citrix product (hereinafter "PRODUCT") and software maintenance (hereinafter "MAINTENANCE") determines the providing entity hereunder. Citrix Systems, Inc., a Delaware corporation, licenses the PRODUCT and provides MAINTENANCE in the Americas. Citrix Systems International GmbH, a Swiss company wholly owned by Citrix Systems, Inc., licenses the PRODUCT and provides MAINTENANCE in Europe, the Middle East, and Africa. Citrix Systems Asia Pacific Pty Ltd. licenses the PRODUCT and provides MAINTENANCE in Asia

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

22. Provide User Name, and Organization information. Select who will see the application.

23. Click Next.

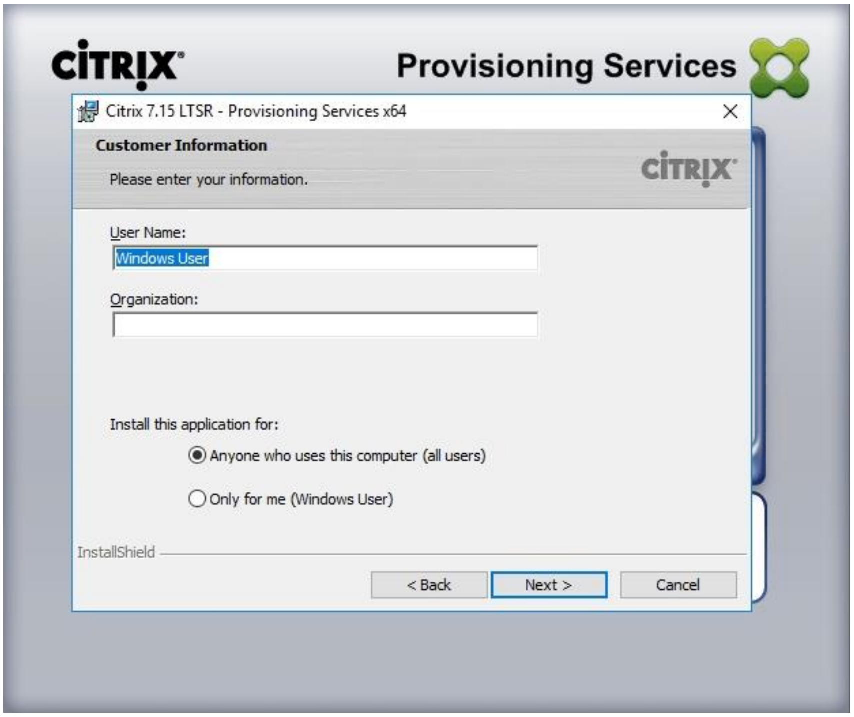

Customer Information

Please enter your information.

User Name:
Windows User

Organization:

Install this application for:

● Anyone who uses this computer (all users)

○ Only for me (Windows User)

24. Accept the default installation location.

25. Click Next.

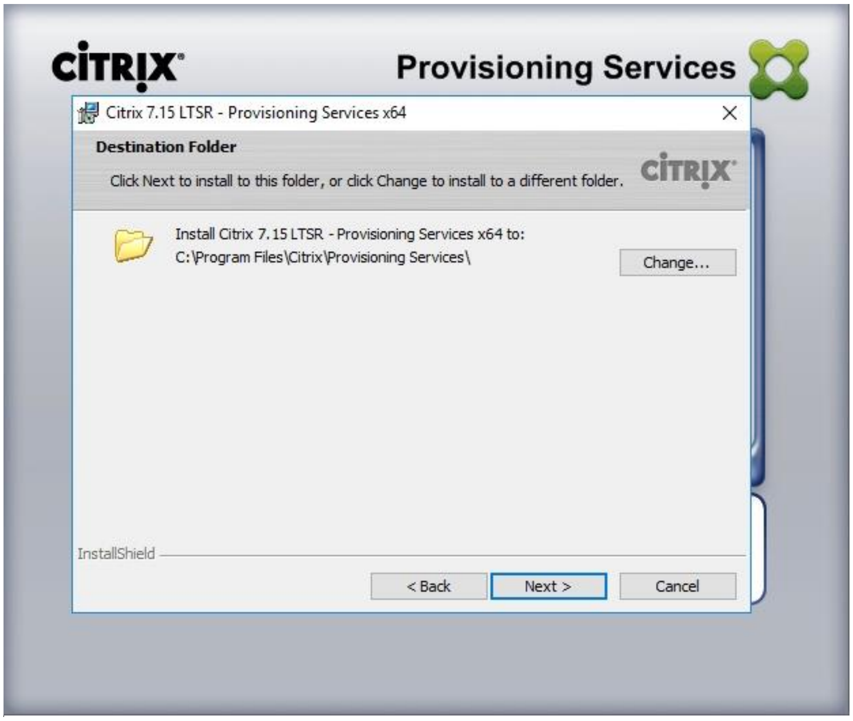

Destination Folder

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

Install Citrix 7.15 LTSR - Provisioning Services x64 to:
C:\Program Files\Citrix\Provisioning Services\

26. Click Install to begin the installation.

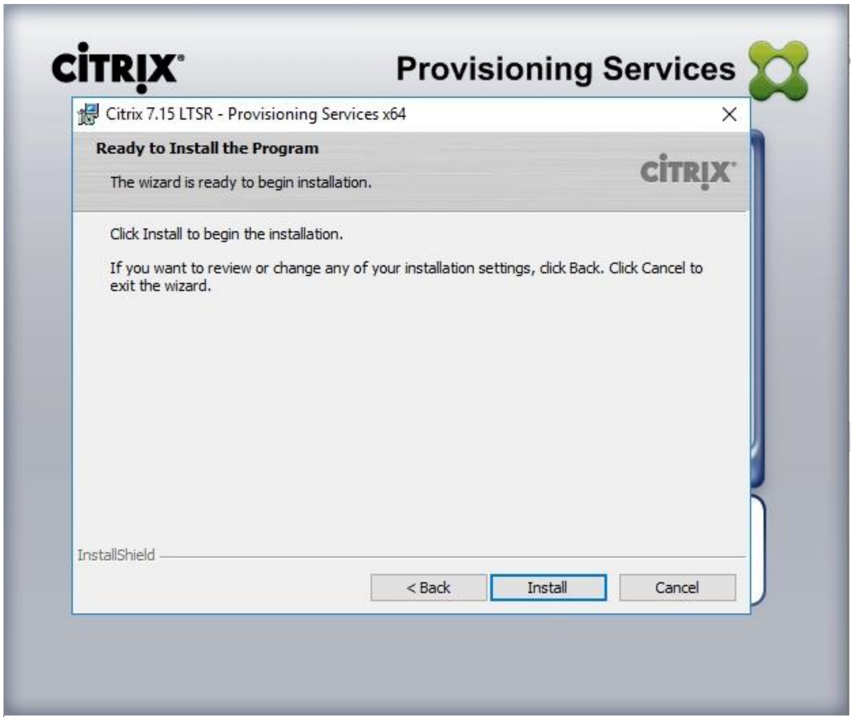

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

27. Click Finish when the install is complete.

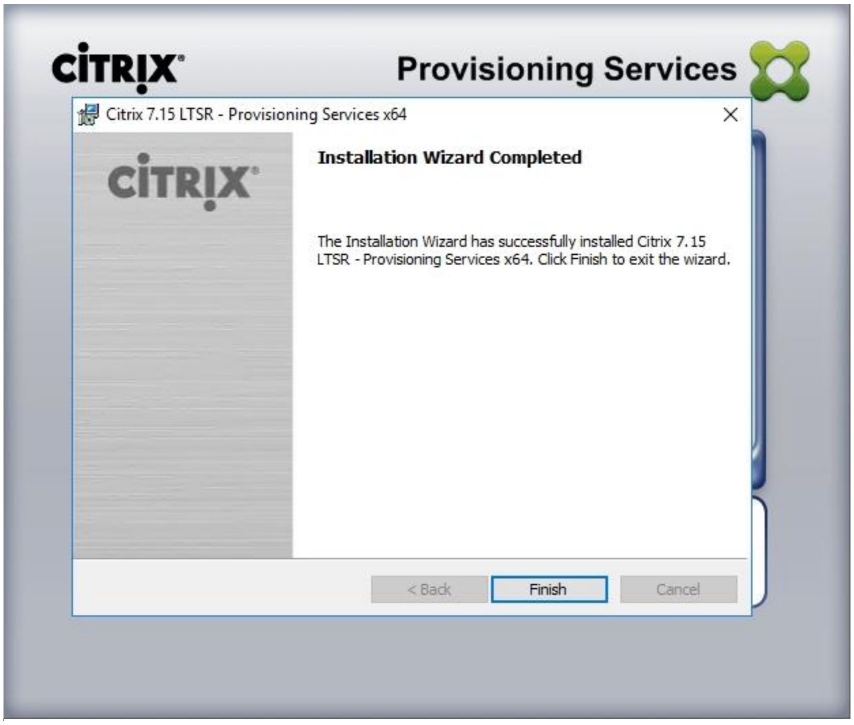

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix 7.15 LTSR - Provisioning Services x64. Click Finish to exit the wizard.

28. The PVS Configuration Wizard starts automatically.

29. Click Next.

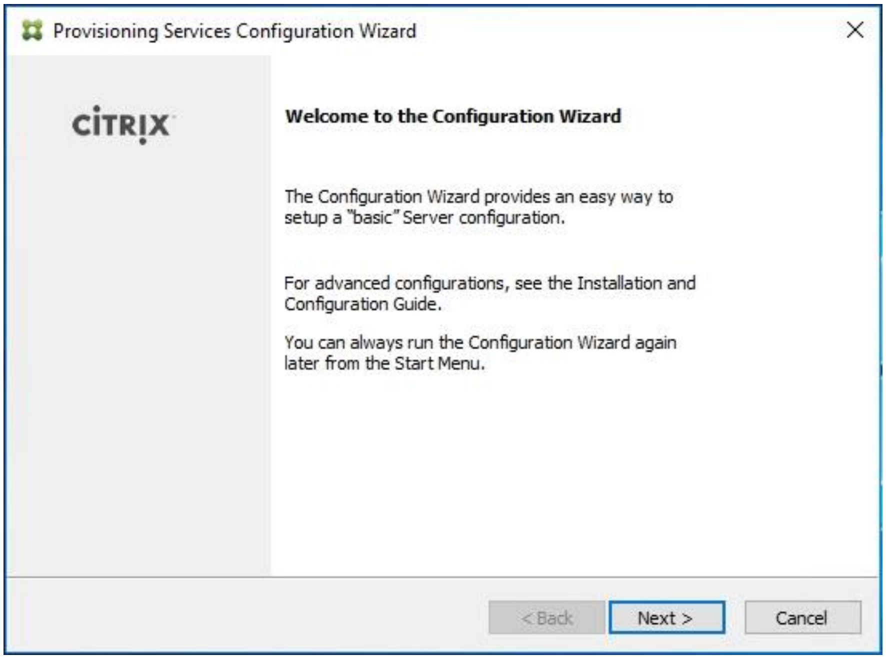

Provisioning Services Configuration Wizard

Welcome to the Configuration Wizard

The Configuration Wizard provides an easy way to setup a "basic" Server configuration.

For advanced configurations, see the Installation and Configuration Guide.

You can always run the Configuration Wizard again later from the Start Menu.

30. Since the PVS server is not the DHCP server for the environment, select the radio button labeled, "The service that runs on another computer."

31. Click Next.

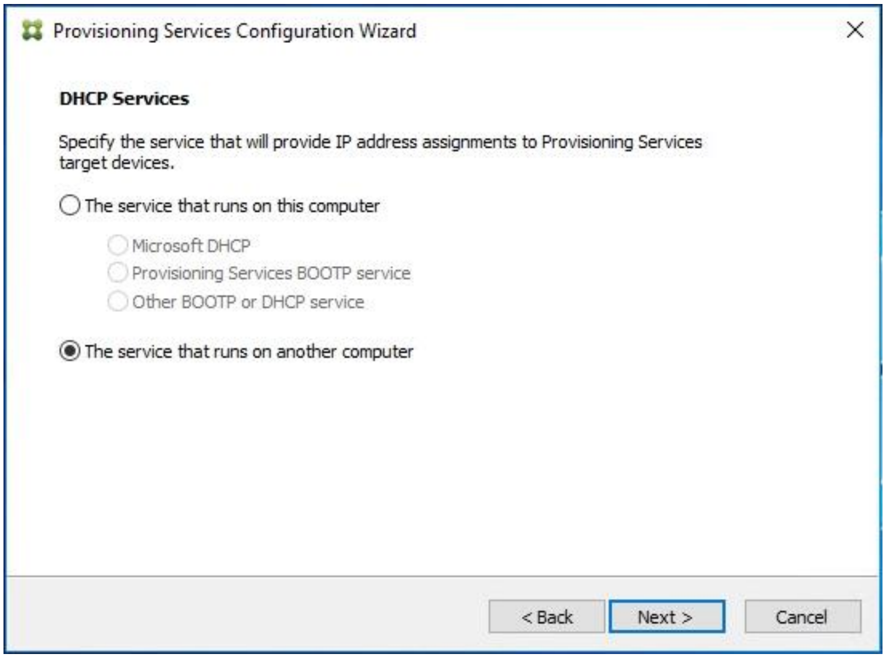

Provisioning Services Configuration Wizard

DHCP Services

Specify the service that will provide IP address assignments to Provisioning Services target devices.

○ The service that runs on this computer

- ○ Microsoft DHCP
- ○ Provisioning Services BOOTP service
- ○ Other BOOTP or DHCP service

● The service that runs on another computer

32. Since DHCP boot options 66 and 67 are used for TFTP services, select the radio button labeled, "The service that runs on another computer."

33. Click Next.

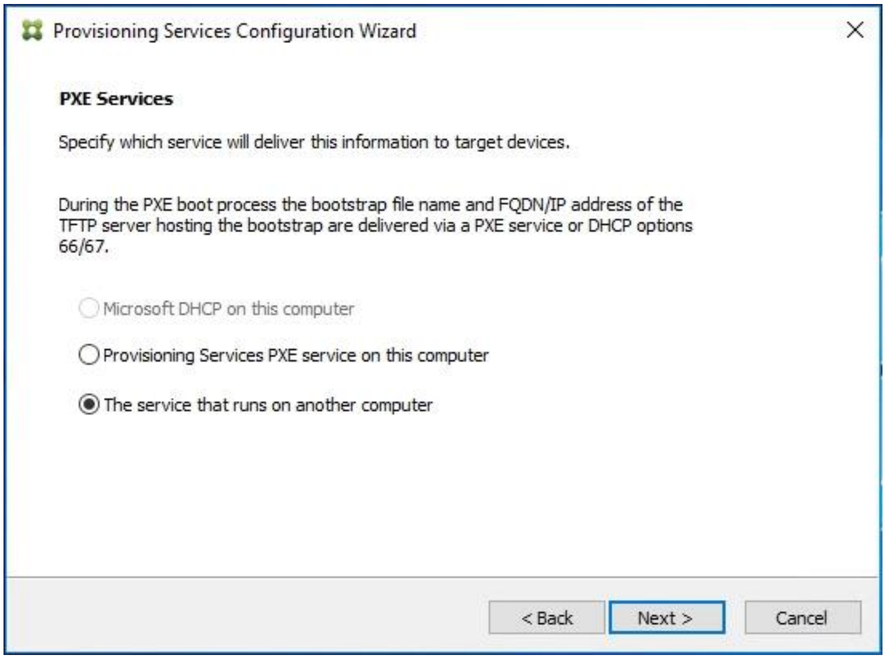

Provisioning Services Configuration Wizard

PXE Services

Specify which service will deliver this information to target devices.

During the PXE boot process the bootstrap file name and FQDN/IP address of the TFTP server hosting the bootstrap are delivered via a PXE service or DHCP options 66/67.

○ Microsoft DHCP on this computer

○ Provisioning Services PXE service on this computer

● The service that runs on another computer

34. Since this is the first server in the farm, select the radio button labeled, “Create farm.”

35. Click Next.

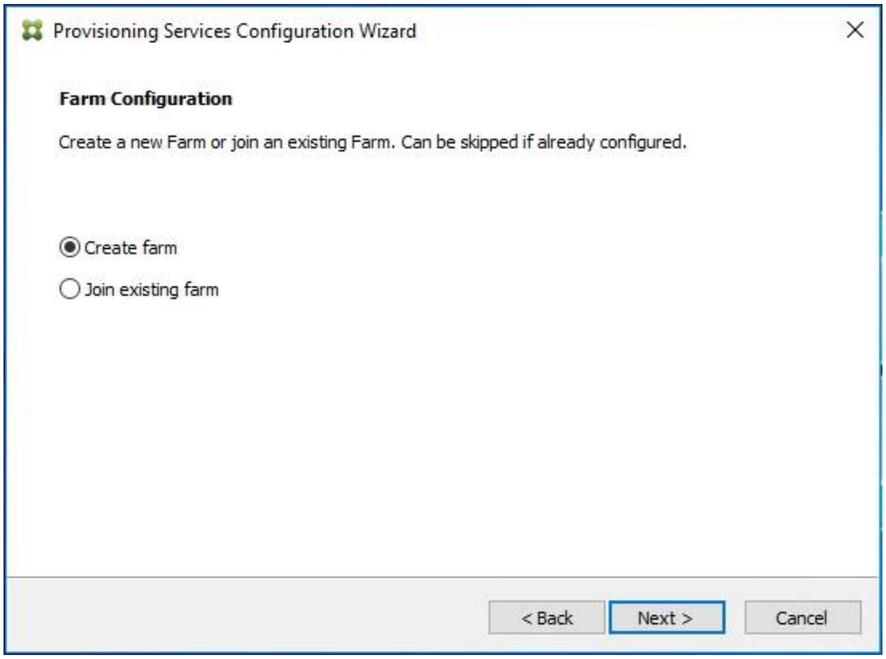

Provisioning Services Configuration Wizard

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

● Create farm

○ Join existing farm

36. Enter the FQDN of the SQL server.

37. Click Next.

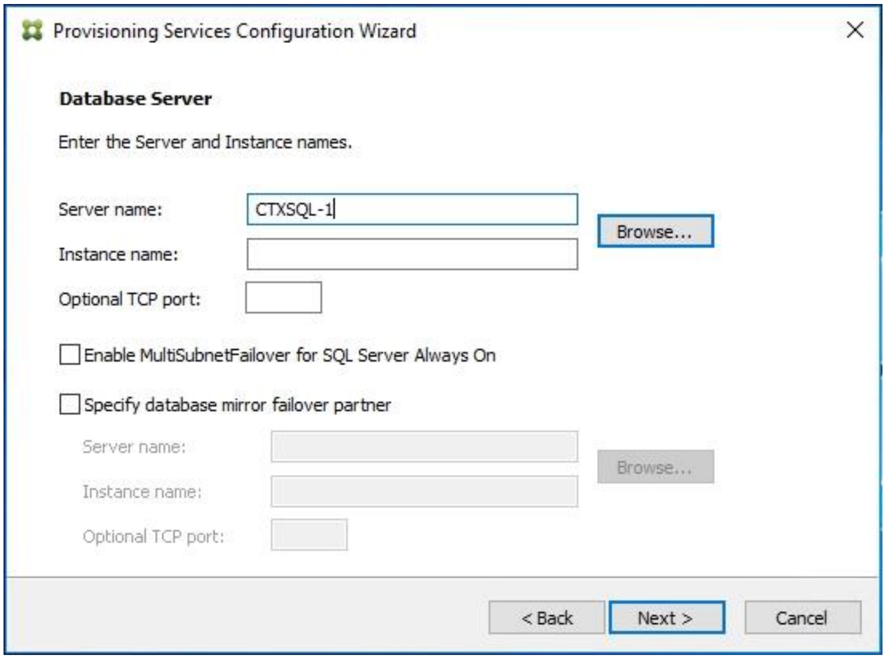

Provisioning Services Configuration Wizard

Database Server

Enter the Server and Instance names.

Server name: CTXSQL-1

Instance name:

Optional TCP port:

□ Enable MultiSubnetFailover for SQL Server Always On

□ Specify database mirror failover partner

Server name:

Instance name:

Optional TCP port:

38. Provide the Database, Farm, Site, and Collection names.

39. Click Next.

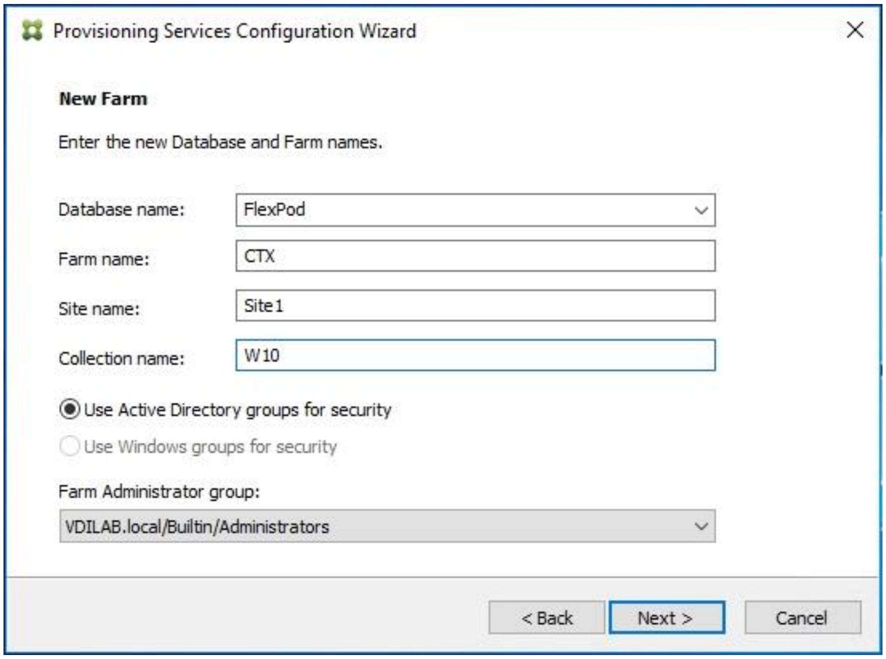

Provisioning Services Configuration Wizard

New Farm

Enter the new Database and Farm names.

Database name: FlexPod

Farm name: CTX

Site name: Site1

Collection name: W10

● Use Active Directory groups for security

○ Use Windows groups for security

Farm Administrator group:

VDILAB.local/Builtin/Administrators

40. Provide the vDisk Store details.

41. Click Next.

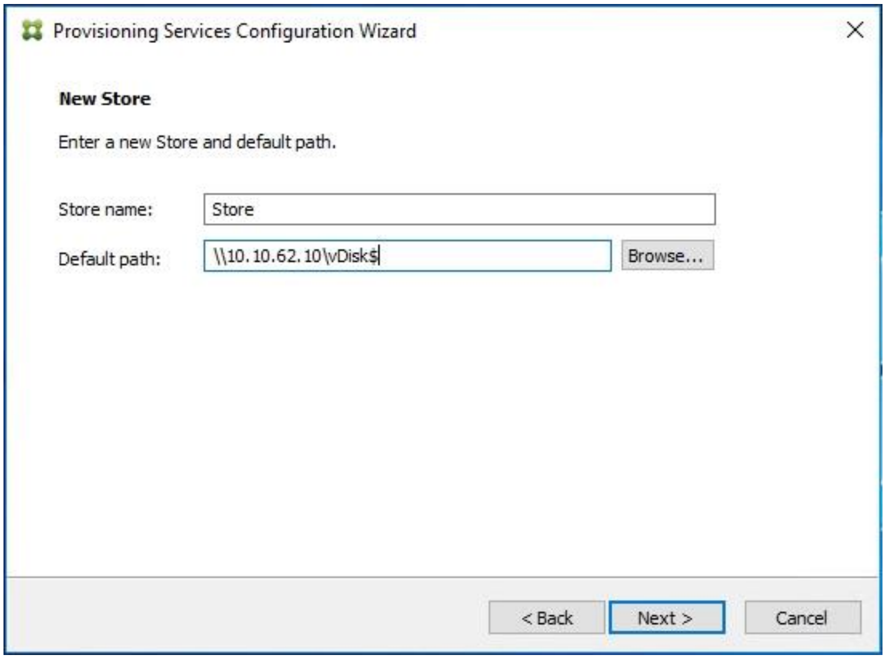

Provisioning Services Configuration Wizard

New Store

Enter a new Store and default path.

Store name: Store

Default path: \\10.10.62.10\vDisk$

For large scale PVS environment it is recommended to create the share using support for CIFS/SMB3 on an enterprise ready File Server.

42. Provide the FQDN of the license server.

43. Optionally, provide a port number if changed on the license server.

44. Click Next.

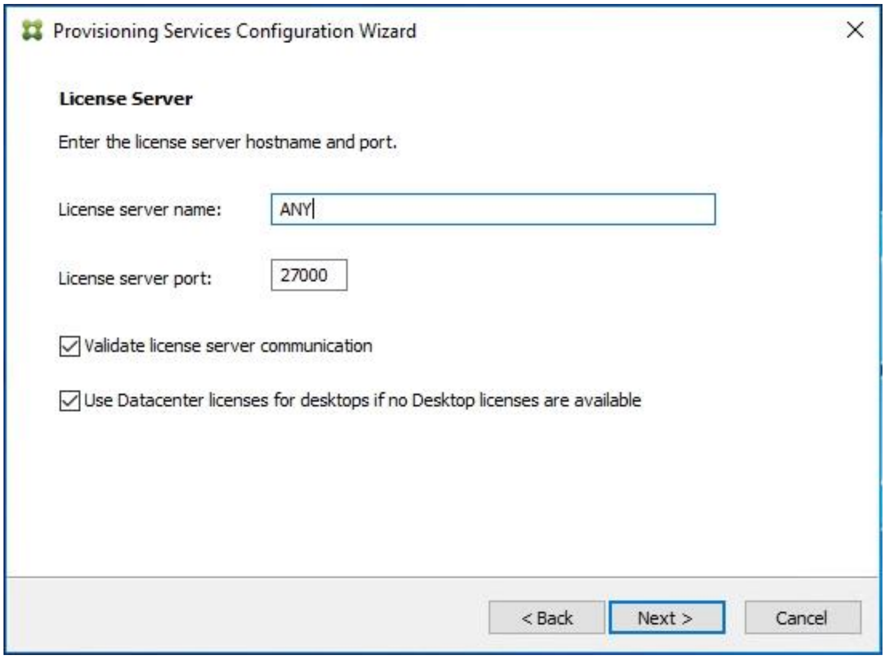

Provisioning Services Configuration Wizard

License Server

Enter the license server hostname and port.

License server name: ANY

License server port: 27000

☑ Validate license server communication

☑ Use Datacenter licenses for desktops if no Desktop licenses are available

If an Active Directory service account is not already setup for the PVS servers, create that account prior to clicking Next on this dialog.

45. Select the Specified user account radio button.

46. Complete the User name, Domain, Password, and Confirm password fields, using the PVS account information created earlier.

47. Click Next.

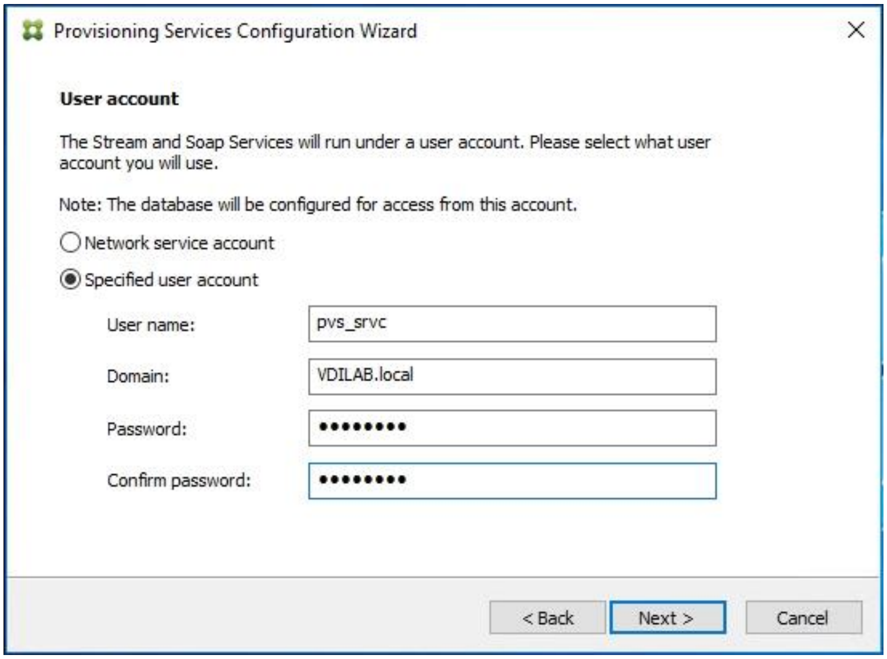

Provisioning Services Configuration Wizard

User account

The Stream and Soap Services will run under a user account. Please select what user account you will use.

Note: The database will be configured for access from this account.

○ Network service account

● Specified user account

User name: pvs_srvc

Domain: VDILAB.local

Password: ••••••••

Confirm password: ••••••••

48. Set the Days between password updates to 7.

This will vary per environment. “7 days” for the configuration was appropriate for testing purposes.

49. Click Next.

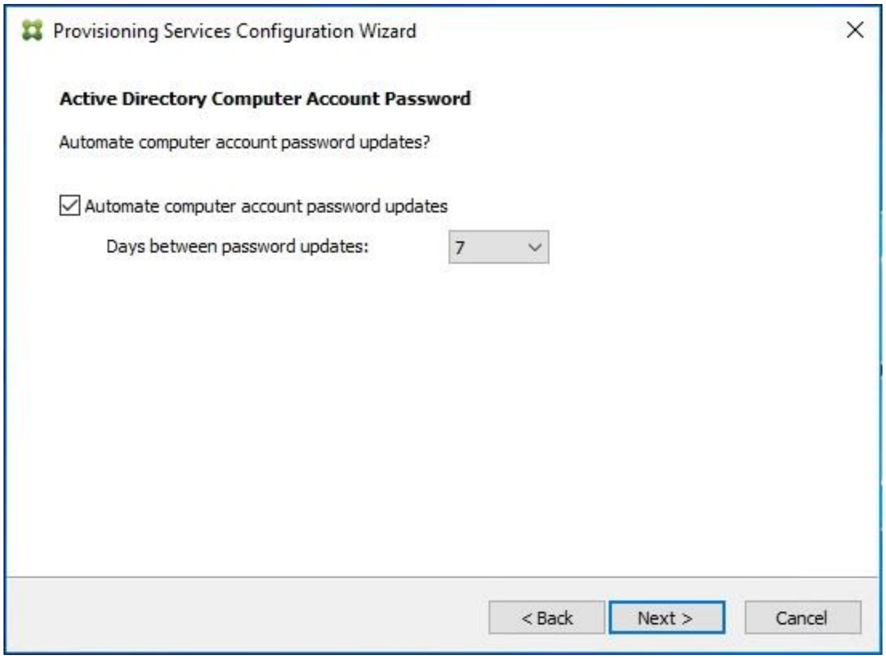
Provisioning Services Configuration Wizard

Active Directory Computer Account Password

Automate computer account password updates?

☑ Automate computer account password updates

Days between password updates: 7

50. Keep the defaults for the network cards.

51. Click Next.

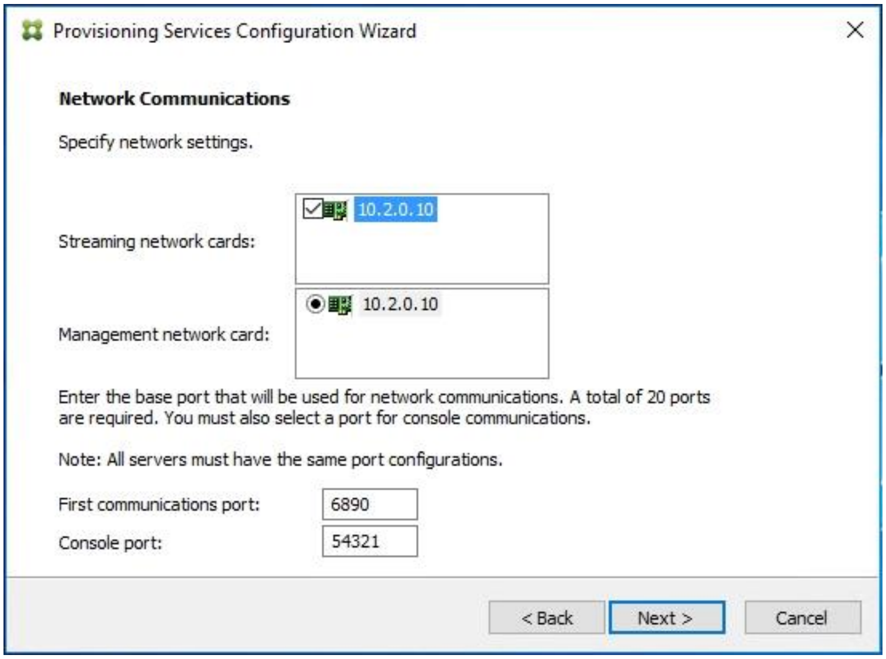
Provisioning Services Configuration Wizard

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.2.0.10

Management network card: ● 10.2.0.10

Enter the base port that will be used for network communications. A total of 20 ports are required. You must also select a port for console communications.

Note: All servers must have the same port configurations.

First communications port: 6890

Console port: 54321

52. Select Use the Provisioning Services TFTP service checkbox.

53. Click Next.

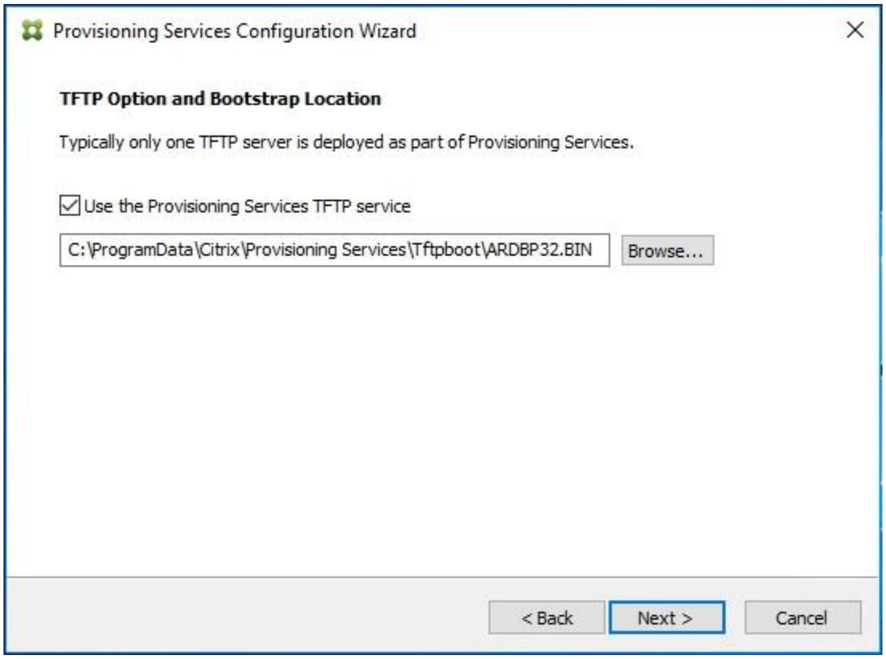

Provisioning Services Configuration Wizard

TFTP Option and Bootstrap Location

Typically only one TFTP server is deployed as part of Provisioning Services.

☑ Use the Provisioning Services TFTP service

C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN

54. Make sure that the IP Addresses for all PVS servers are listed in the Stream Servers Boot List.

55. Click Next.

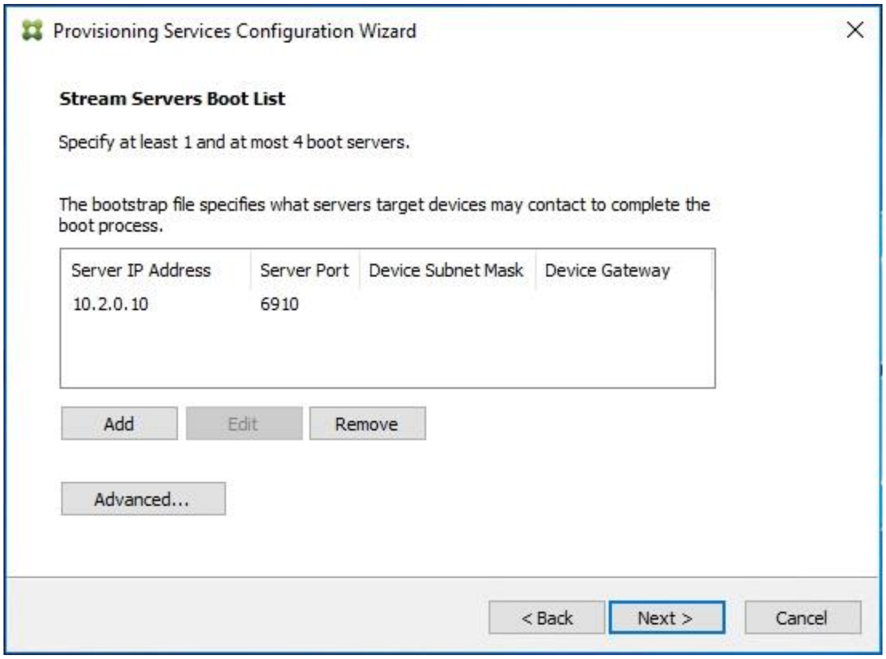

Provisioning Services Configuration Wizard

Stream Servers Boot List

Specify at least 1 and at most 4 boot servers.

The bootstrap file specifies what servers target devices may contact to complete the boot process.

Server IP Address	Server Port	Device Subnet Mask	Device Gateway
10.2.0.10	6910		

56. If Soap Server is used, provide details.

57. Click Next.

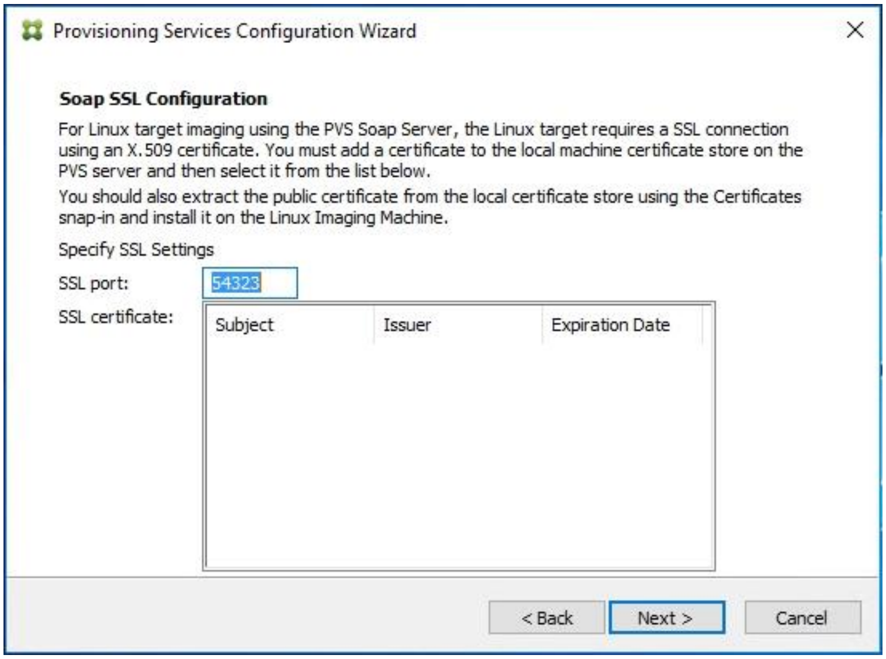

Soap SSL Configuration

For Linux target imaging using the PVS Soap Server, the Linux target requires a SSL connection using an X.509 certificate. You must add a certificate to the local machine certificate store on the PVS server and then select it from the list below.

You should also extract the public certificate from the local certificate store using the Certificates snap-in and install it on the Linux Imaging Machine.

Specify SSL Settings

SSL port: 54323

SSL certificate:

Subject	Issuer	Expiration Date

58. If desired fill in Problem Report Configuration.

59. Click Next.

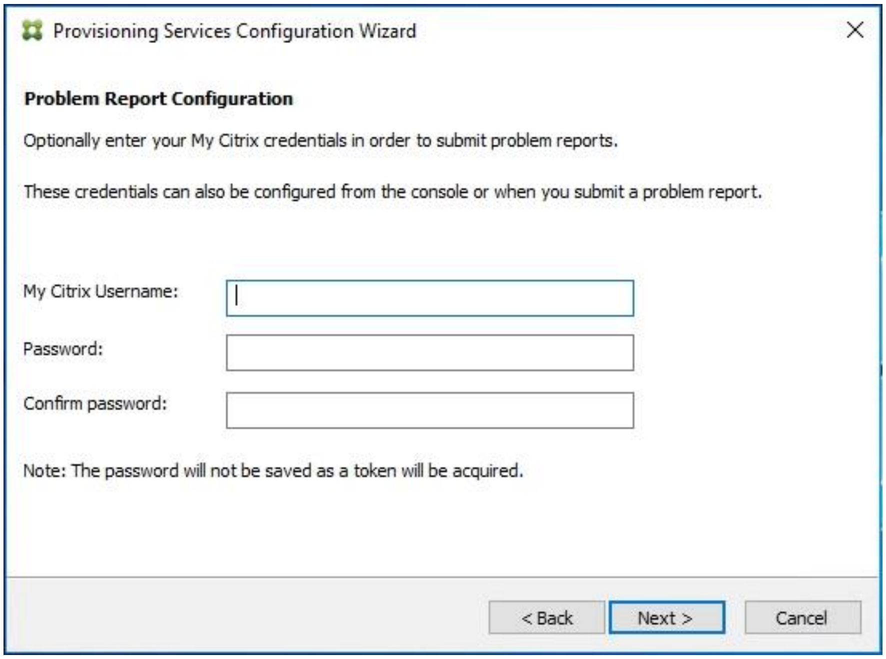

Problem Report Configuration

Optionally enter your My Citrix credentials in order to submit problem reports.

These credentials can also be configured from the console or when you submit a problem report.

My Citrix Username:

Password:

Confirm password:

Note: The password will not be saved as a token will be acquired.

60. Click Finish to start installation.

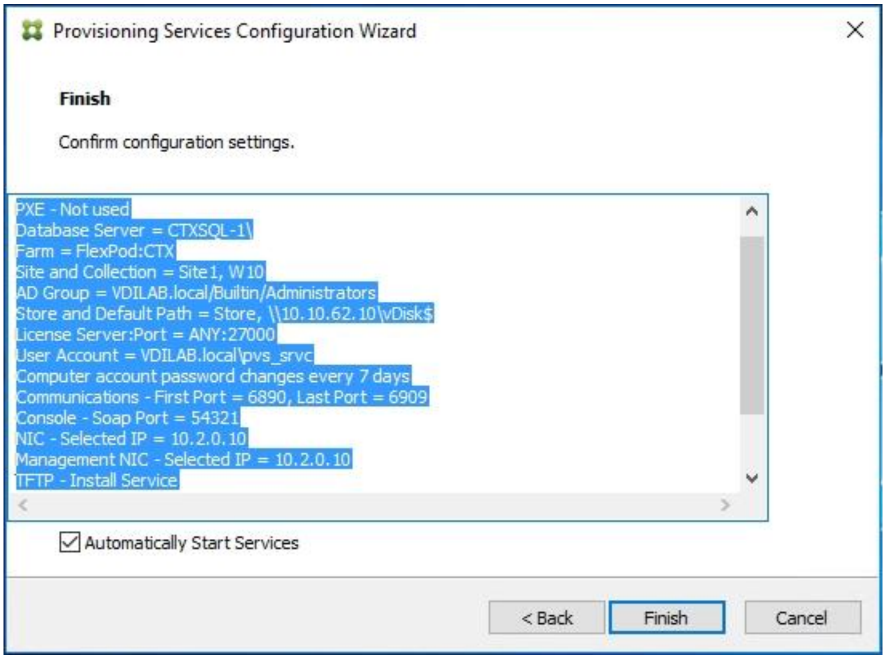

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = CTXSQL-1\
Farm = FlexPod:CTX
Site and Collection = Site1, W10
AD Group = VDILAB.local/Builtin/Administrators
Store and Default Path = Store, \\10.10.62.10\vDisk$
License Server:Port = ANY:27000
User Account = VDILAB.local\pvs_srvc
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.2.0.10
Management NIC - Selected IP = 10.2.0.10
TFTP - Install Service
```

☑ Automatically Start Services

61. When the installation is completed, click Done.

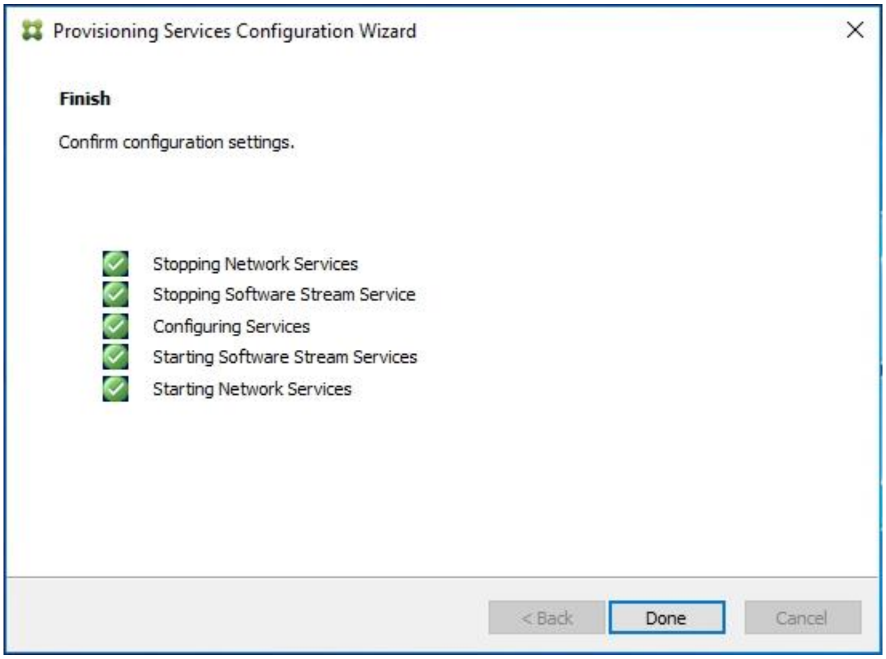

Finish

Confirm configuration settings.

- ✓ Stopping Network Services
- ✓ Stopping Software Stream Service
- ✓ Configuring Services
- ✓ Starting Software Stream Services
- ✓ Starting Network Services

Install Additional PVS Servers

Complete the installation steps on the additional PVS servers up to the configuration step where it asks to Create or Join a farm. In this CVD, we repeated the procedure to add a total of three PVS servers. To install additional PVS server, complete the following steps:

1. On the Farm Configuration dialog, select “Join existing farm.”

2. Click Next.

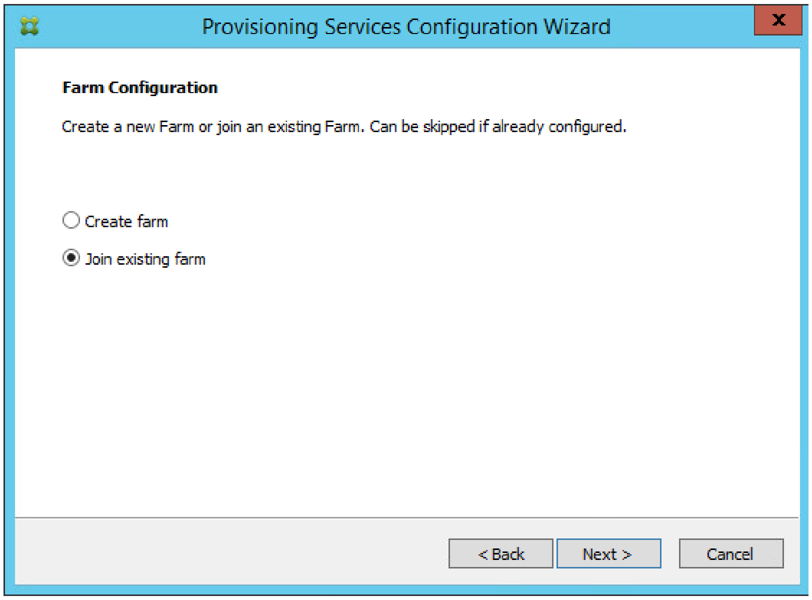

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

○ Create farm

● Join existing farm

3. Provide the FQDN of the SQL Server.

4. Click Next.

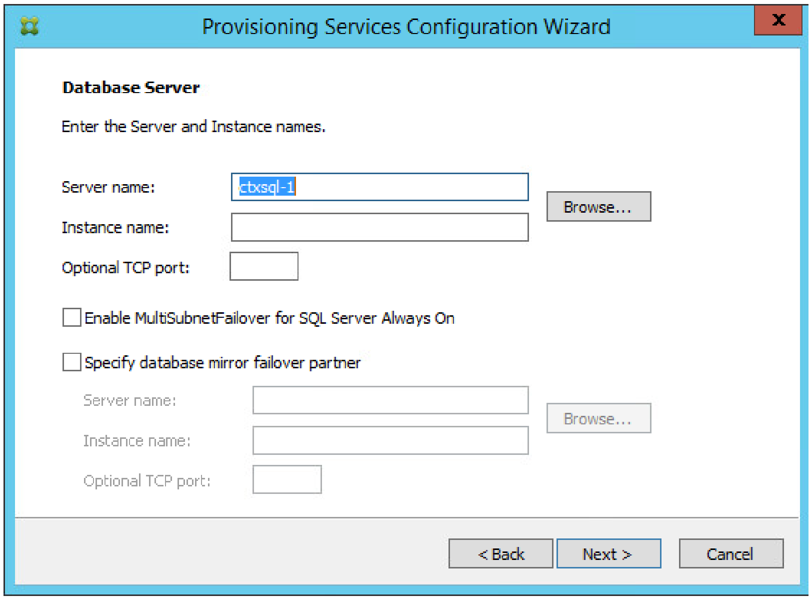

Database Server

Enter the Server and Instance names.

Server name: ctxsql-1

Instance name:

Optional TCP port:

□ Enable MultiSubnetFailover for SQL Server Always On

□ Specify database mirror failover partner

Server name:

Instance name:

Optional TCP port:

5. Accept the Farm Name.

6. Click Next.

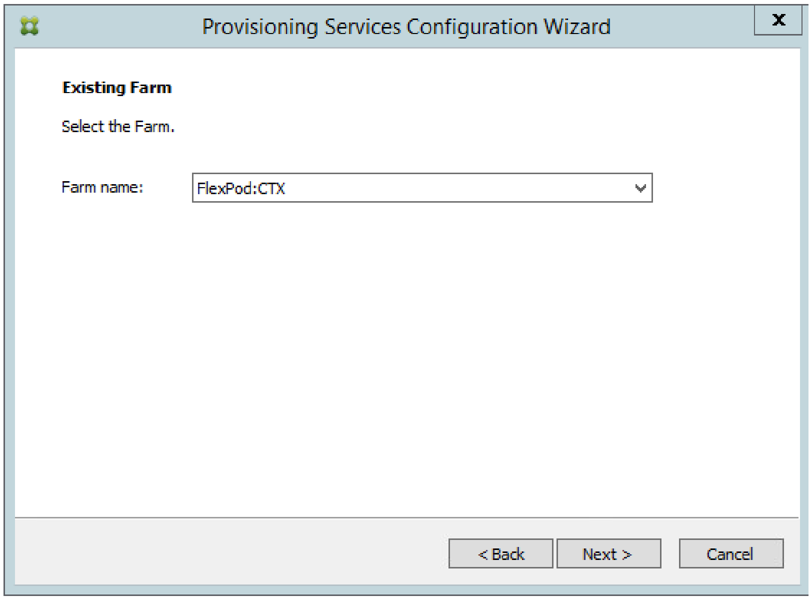

Existing Farm

Select the Farm.

Farm name: FlexPod:CTX

7. Accept the Existing Site.

8. Click Next.

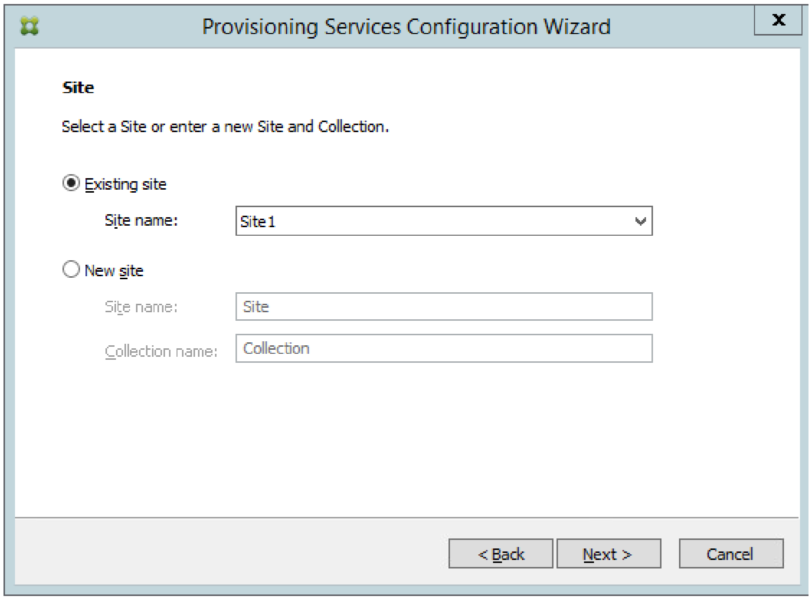

Site

Select a Site or enter a new Site and Collection.

● Existing site

Site name: Site1

○ New site

Site name: Site

Collection name: Collection

9. Accept the existing vDisk store.

10. Click Next.

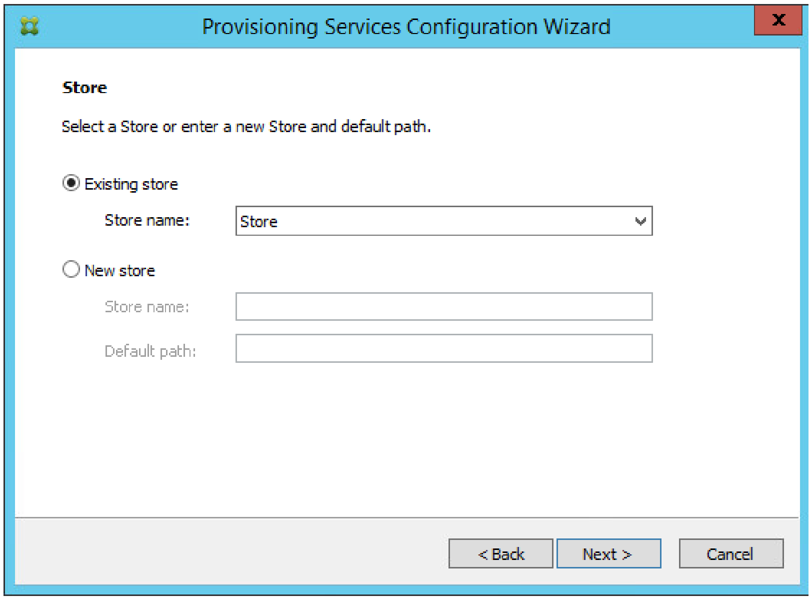

Store

Select a Store or enter a new Store and default path.

● Existing store

Store name: Store

○ New store

Store name:

Default path:

11. Provide the PVS service account information.

12. Click Next.

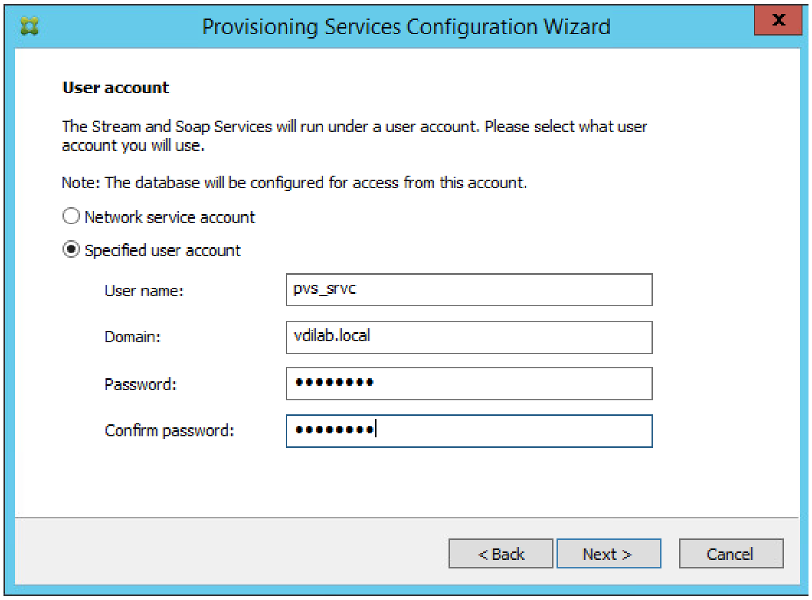

User account

The Stream and Soap Services will run under a user account. Please select what user account you will use.

Note: The database will be configured for access from this account.

○ Network service account

● Specified user account

User name: pvs_srvc

Domain: vdilab.local

Password: ••••••••

Confirm password: ••••••••

13. Set the Days between password updates to 7.

14. Click Next.

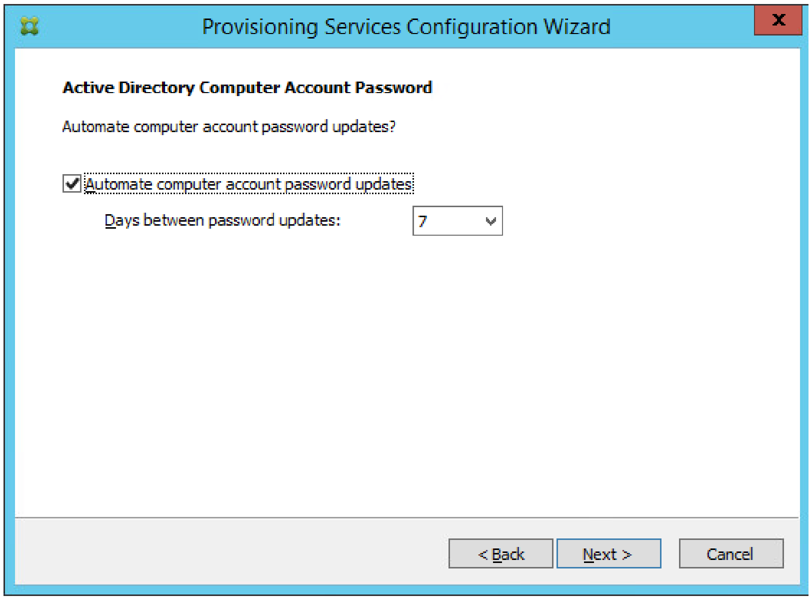

Provisioning Services Configuration Wizard

Active Directory Computer Account Password

Automate computer account password updates?

☑ Automate computer account password updates

Days between password updates: 7

15. Accept the network card settings.

16. Click Next.

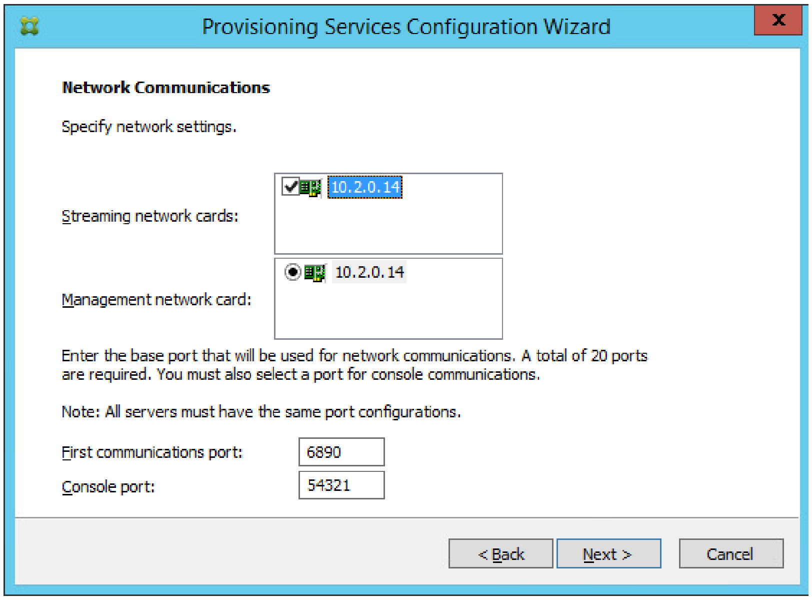

Provisioning Services Configuration Wizard

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.2.0.14

Management network card: ● 10.2.0.14

Enter the base port that will be used for network communications. A total of 20 ports are required. You must also select a port for console communications.

Note: All servers must have the same port configurations.

First communications port: 6890

Console port: 54321

17. Select Use the Provisioning Services TFTP service checkbox.

18. Click Next.

19. Make sure that the IP Addresses for all PVS servers are listed in the Stream Servers Boot List.

20. Click Next.

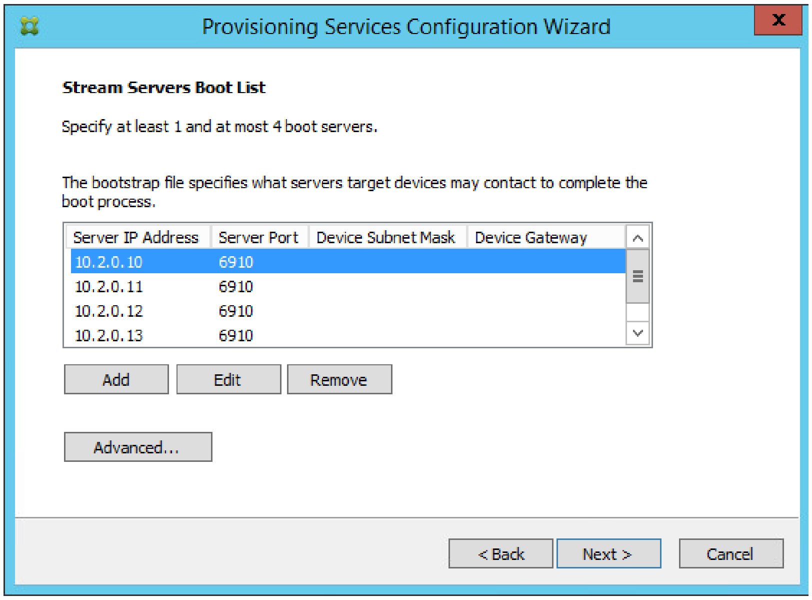

Provisioning Services Configuration Wizard

Stream Servers Boot List

Specify at least 1 and at most 4 boot servers.

The bootstrap file specifies what servers target devices may contact to complete the boot process.

Server IP Address	Server Port	Device Subnet Mask	Device Gateway
10.2.0.10	6910		
10.2.0.11	6910		
10.2.0.12	6910		
10.2.0.13	6910		

21. If Soap Server is used, provide details.

22. Click Next.

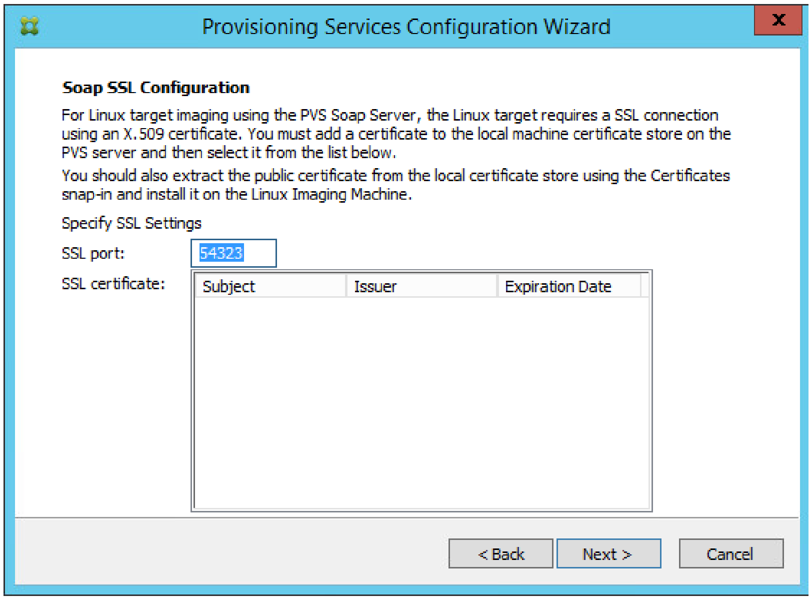

Provisioning Services Configuration Wizard

Soap SSL Configuration

For Linux target imaging using the PVS Soap Server, the Linux target requires a SSL connection using an X.509 certificate. You must add a certificate to the local machine certificate store on the PVS server and then select it from the list below.

You should also extract the public certificate from the local certificate store using the Certificates snap-in and install it on the Linux Imaging Machine.

Specify SSL Settings

SSL port: 54323

SSL certificate:

Subject	Issuer	Expiration Date

23. If desired fill in Problem Report Configuration.

24. Click Next.

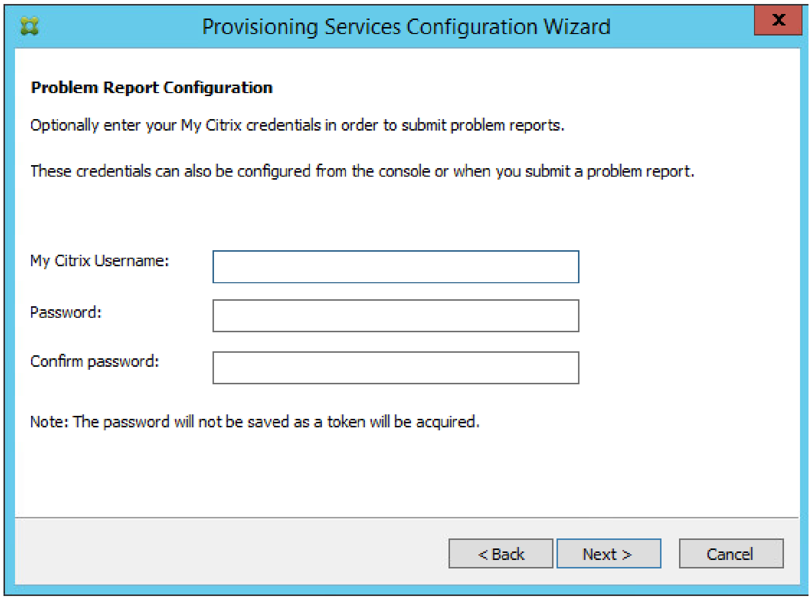

Problem Report Configuration

Optionally enter your My Citrix credentials in order to submit problem reports.

These credentials can also be configured from the console or when you submit a problem report.

My Citrix Username:

Password:

Confirm password:

Note: The password will not be saved as a token will be acquired.

25. Click Finish to start the installation process.

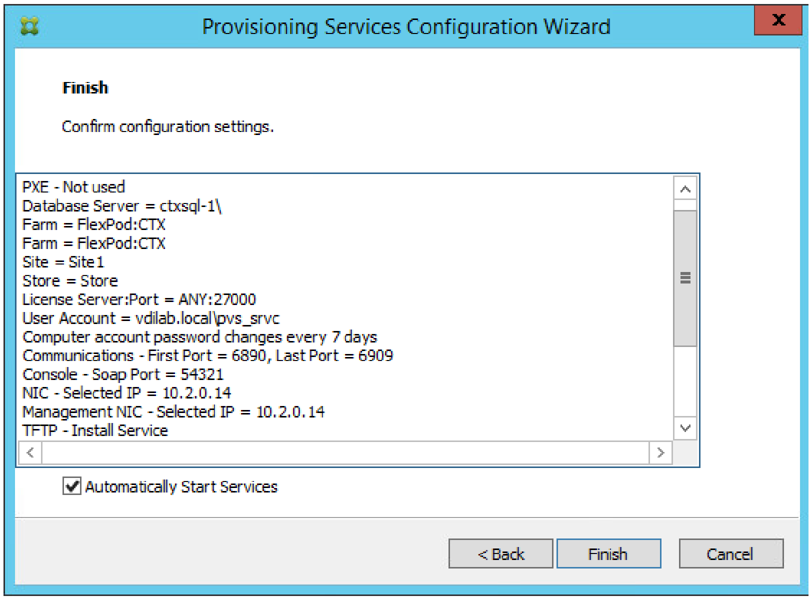

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = ctxsql-1\
Farm = FlexPod:CTX
Farm = FlexPod:CTX
Site = Site1
Store = Store
License Server:Port = ANY:27000
User Account = vdilab.local\pvs_srvc
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.2.0.14
Management NIC - Selected IP = 10.2.0.14
TFTP - Install Service
```

☑ Automatically Start Services

26. Click Done when the installation finishes.

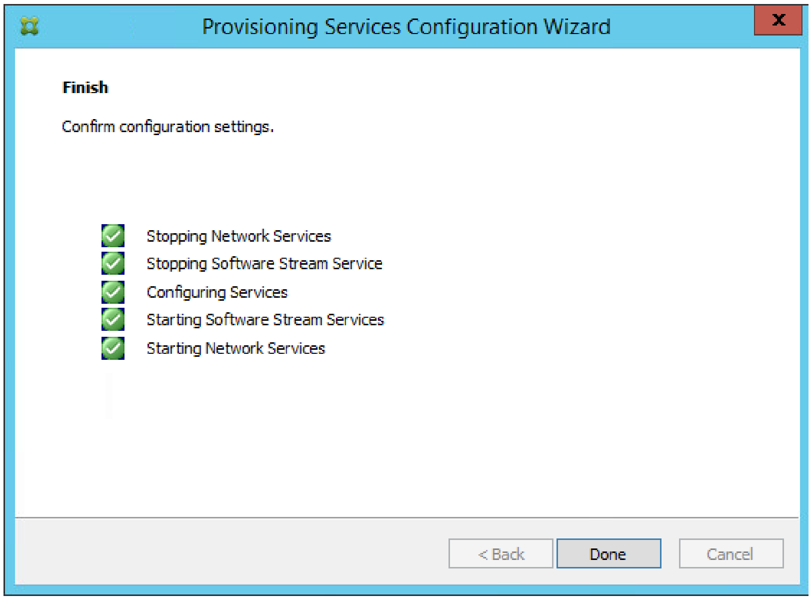

Finish

Confirm configuration settings.

✓ Stopping Network Services
✓ Stopping Software Stream Service
✓ Configuring Services
✓ Starting Software Stream Services
✓ Starting Network Services

You can optionally install the Provisioning Services console on the second PVS server following the procedure in the section Installing Provisioning Services.

After completing the steps to install the three additional PVS servers, launch the Provisioning Services Console to verify that the PVS Servers and Stores are configured and that DHCP boot options are defined.

27. Launch Provisioning Services Console and select Connect to Farm.

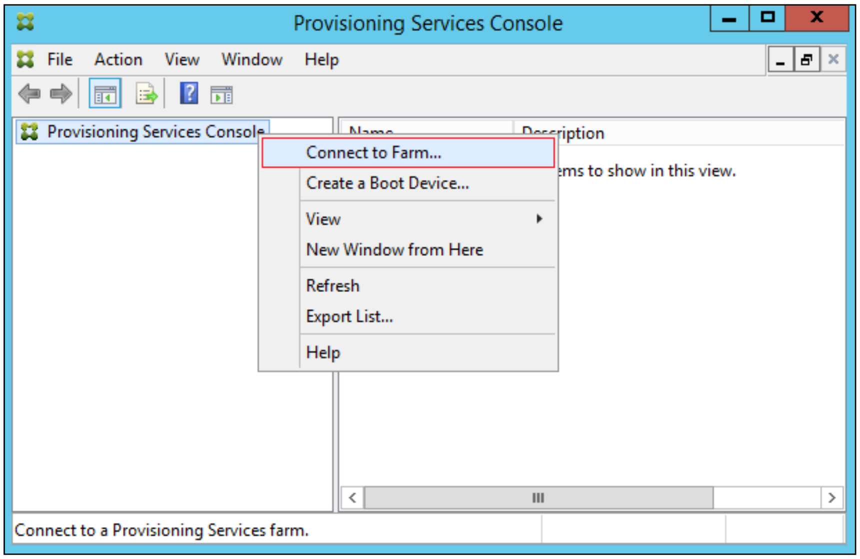

There are no items to show in this view.

Connect to a Provisioning Services farm.

28. Enter localhost for the PVS1 server.

29. Click Connect.

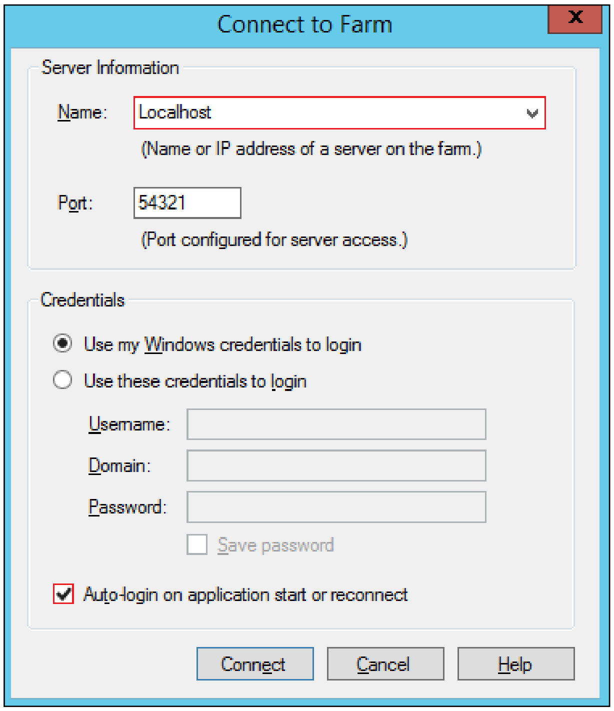

Connect to Farm

Server Information

Name: Localhost

(Name or IP address of a server on the farm.)

Port: 54321

(Port configured for server access.)

Credentials

● Use my Windows credentials to login

○ Use these credentials to login

Username:

Domain:

Password:

□ Save password

☑ Auto-login on application start or reconnect

30. Select Store Properties from the drop-down list.

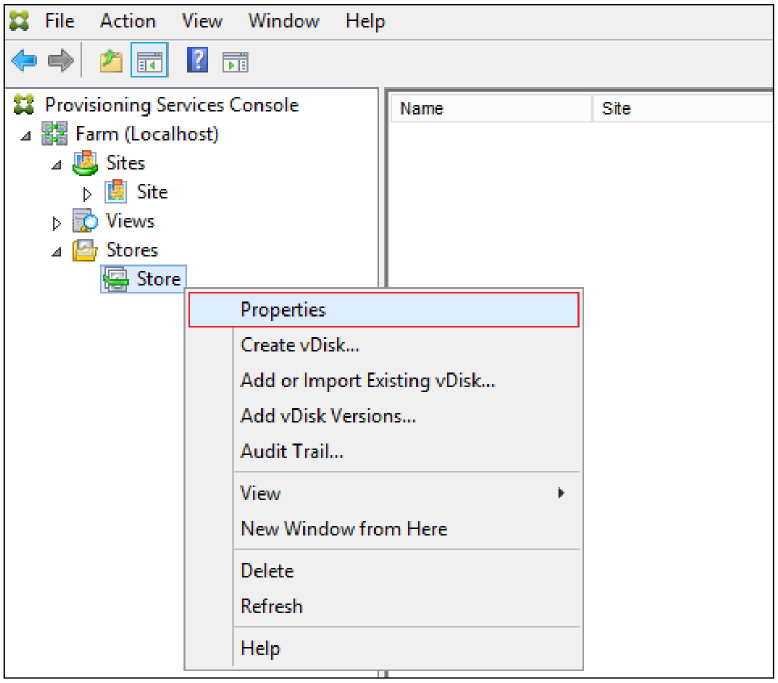

31. In the Store Properties dialog, add the Default store path to the list of Default write cache paths.

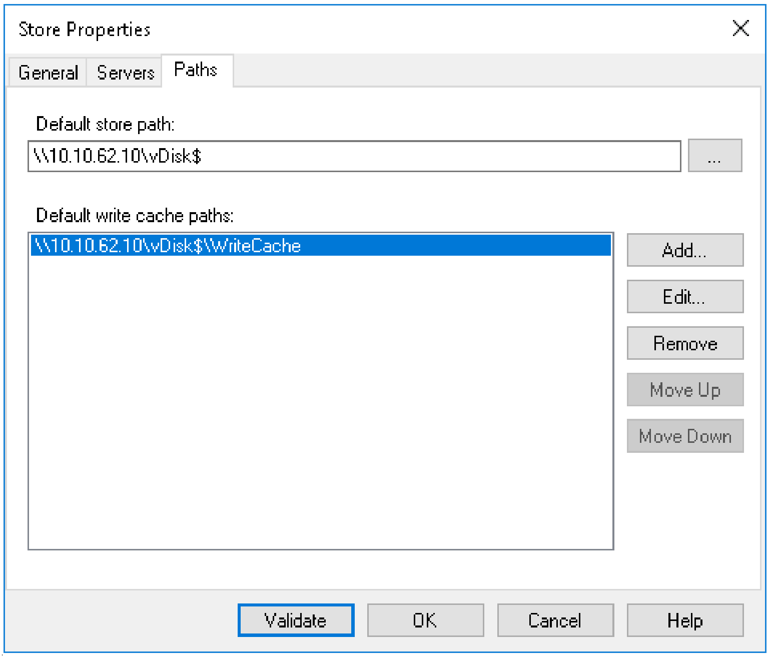

Default store path: \\10.10.62.10\vDisk$

Default write cache paths:
\\10.10.62.10\vDisk$\WriteCache

32. Click Validate. If the validation is successful, click Close and OK to continue.

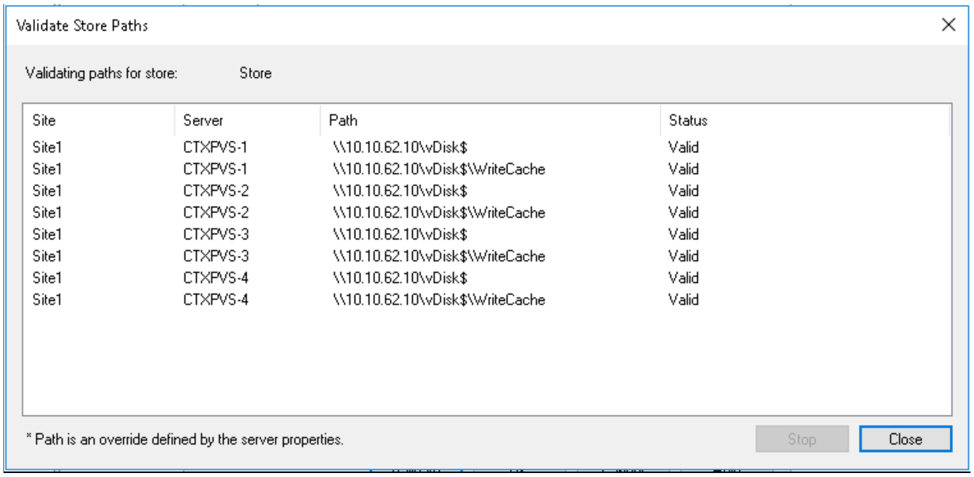

Validating paths for store: Store

Site	Server	Path	Status
Site1	CTXPVS-1	\\10.10.62.10\vDisk$	Valid
Site1	CTXPVS-1	\\10.10.62.10\vDisk$\WriteCache	Valid
Site1	CTXPVS-2	\\10.10.62.10\vDisk$	Valid
Site1	CTXPVS-2	\\10.10.62.10\vDisk$\WriteCache	Valid
Site1	CTXPVS-3	\\10.10.62.10\vDisk$	Valid
Site1	CTXPVS-3	\\10.10.62.10\vDisk$\WriteCache	Valid
Site1	CTXPVS-4	\\10.10.62.10\vDisk$	Valid
Site1	CTXPVS-4	\\10.10.62.10\vDisk$\WriteCache	Valid

* Path is an override defined by the server properties.

Install XenDesktop Virtual Desktop Agents

Virtual Delivery Agents (VDAs) are installed on the server and workstation operating systems, and enable connections for desktops and apps. The following procedure was used to install VDAs for both HVD and HSD environments.

By default, when you install the Virtual Delivery Agent, Citrix User Profile Management is installed silently on master images. (Using profile management as a profile solution is optional but was used for this CVD, and is described in a later section.)

To install XenDesktop Virtual Desktop Agents, complete the following steps:

1. Launch the XenDesktop installer from the XenDesktop 7.15 ISO.

2. Click **Start** on the Welcome Screen.

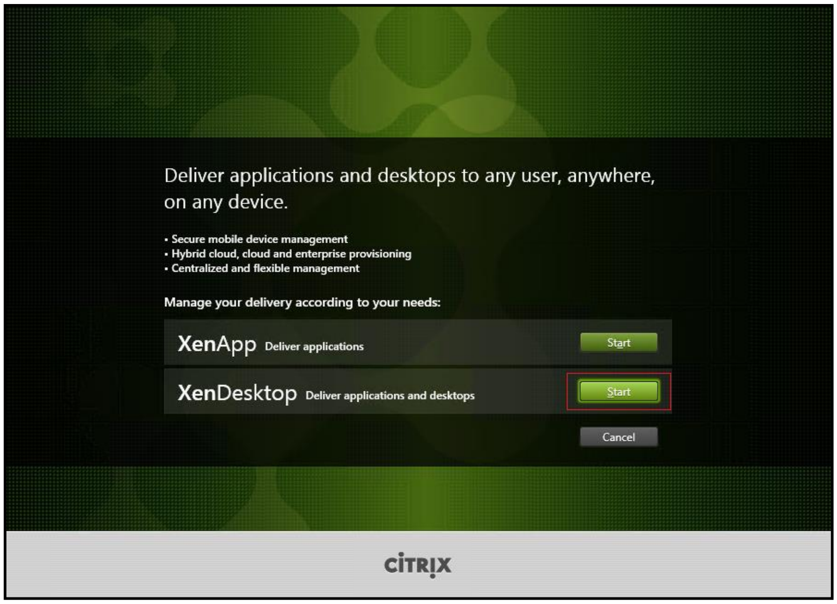

Deliver applications and desktops to any user, anywhere, on any device.

- Secure mobile device management
- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

XenApp Deliver applications
Start
XenDesktop Deliver applications and desktops
Start
Cancel
CITRIX

3. To install the VDA for the Hosted Virtual Desktops (VDI), select **Virtual Delivery Agent for Windows Desktop OS**.

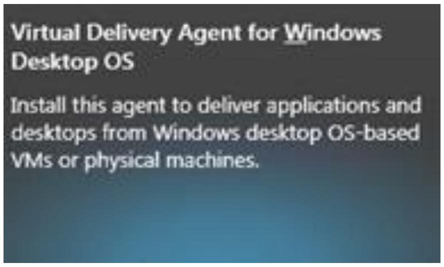

Virtual Delivery Agent for Windows Desktop OS

Install this agent to deliver applications and desktops from Windows desktop OS-based VMs or physical machines.

4. After the VDA is installed for Hosted Virtual Desktops, repeat the procedure to install the VDA for Hosted Shared Desktops (RDS). In this case, select **Virtual Delivery Agent for Windows Server OS** and follow the same basic steps.

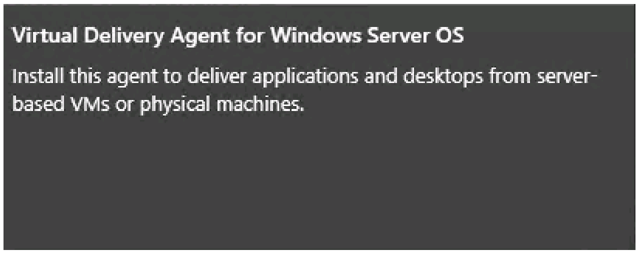

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

5. Select “Create a Master Image.”

6. Click Next.

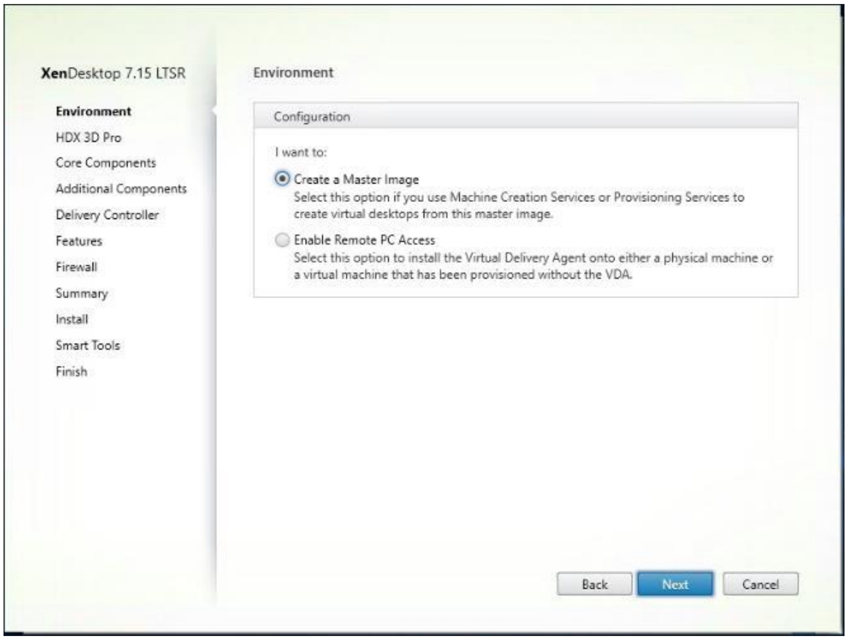

XenDesktop 7.15 LTSR

Environment

Configuration

I want to:

● Create a Master Image
Select this option if you use Machine Creation Services or Provisioning Services to create virtual desktops from this master image.

○ Enable Remote PC Access
Select this option to install the Virtual Delivery Agent onto either a physical machine or a virtual machine that has been provisioned without the VDA.

7. For the VDI vDisk, select "No, install the standard VDA."

Select Yes, install in HDX 3D Pro Mode if VM will be used with vGPU. For more information, go to section Configure VM with vGPU.

8. Click Next.

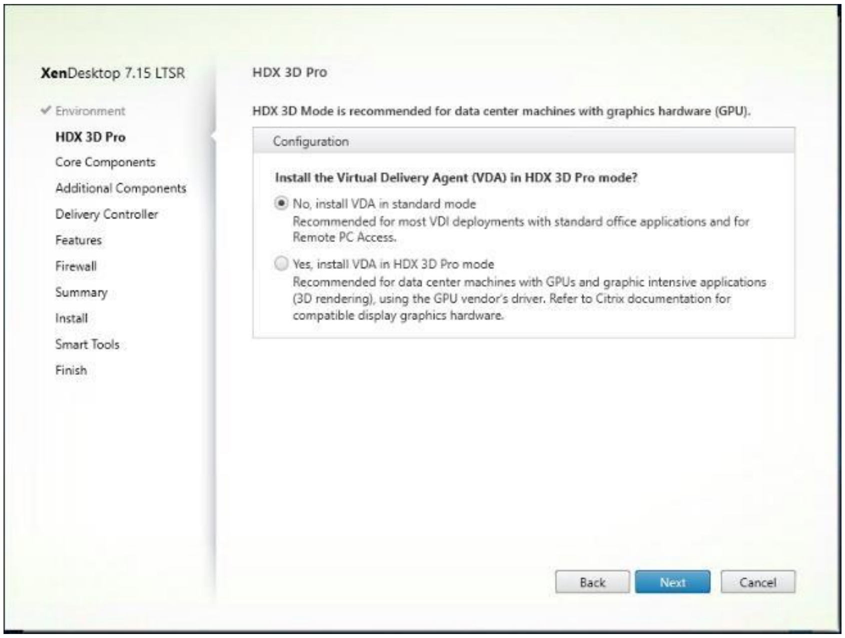

XenDesktop 7.15 LTSR

HDX 3D Pro

HDX 3D Mode is recommended for data center machines with graphics hardware (GPU).

Configuration

Install the Virtual Delivery Agent (VDA) in HDX 3D Pro mode?

● No, install VDA in standard mode
Recommended for most VDI deployments with standard office applications and for Remote PC Access.

○ Yes, install VDA in HDX 3D Pro mode
Recommended for data center machines with GPUs and graphic intensive applications (3D rendering), using the GPU vendor's driver. Refer to Citrix documentation for compatible display graphics hardware.

9. Optional: Do not select Citrix Receiver**.**

10. Click Next.

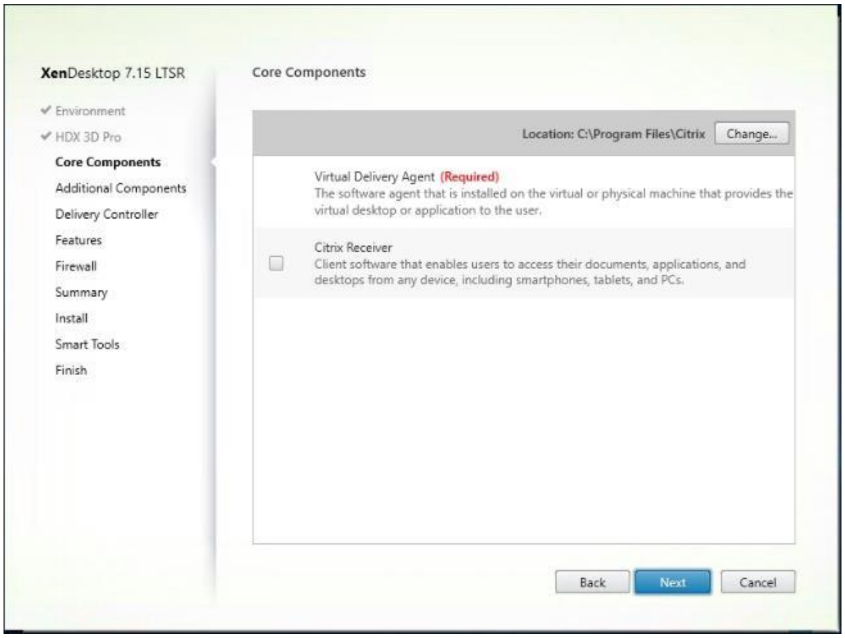

XenDesktop 7.15 LTSR

✓ Environment
✓ HDX 3D Pro
Core Components
Additional Components
Delivery Controller
Features
Firewall
Summary
Install
Smart Tools
Finish

Core Components

Location: C:\Program Files\Citrix

Virtual Delivery Agent (Required)
The software agent that is installed on the virtual or physical machine that provides the virtual desktop or application to the user.

□ Citrix Receiver
Client software that enables users to access their documents, applications, and desktops from any device, including smartphones, tablets, and PCs.

11. Select additional components required for your image. In this design, only UPM and MCS components were installed on the image.

> Deselect Citrix Machine Identity Service when building a master image for use with Citrix Provisioning Services.

12. Click **Next**

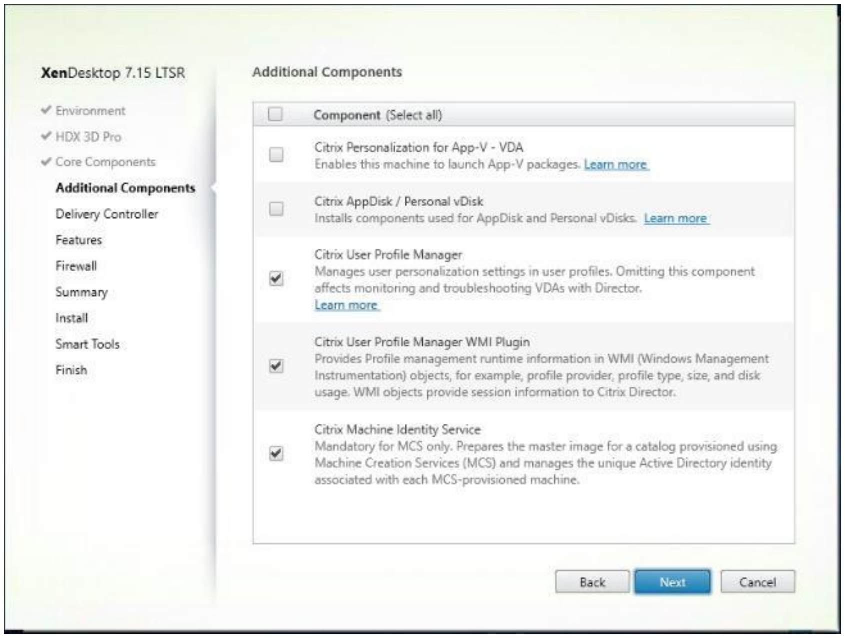

XenDesktop 7.15 LTSR

✓ Environment
✓ HDX 3D Pro
✓ Core Components
Additional Components
Delivery Controller
Features
Firewall
Summary
Install
Smart Tools
Finish

Additional Components

□ Component (Select all)

□ Citrix Personalization for App-V - VDA
Enables this machine to launch App-V packages. Learn more

□ Citrix AppDisk / Personal vDisk
Installs components used for AppDisk and Personal vDisks. Learn more

☑ Citrix User Profile Manager
Manages user personalization settings in user profiles. Omitting this component affects monitoring and troubleshooting VDAs with Director.
Learn more

☑ Citrix User Profile Manager WMI Plugin
Provides Profile management runtime information in WMI (Windows Management Instrumentation) objects, for example, profile provider, profile type, size, and disk usage. WMI objects provide session information to Citrix Director.

☑ Citrix Machine Identity Service
Mandatory for MCS only. Prepares the master image for a catalog provisioned using Machine Creation Services (MCS) and manages the unique Active Directory identity associated with each MCS-provisioned machine.

13. Do not configure Delivery Controllers at this time.

14. Click Next.

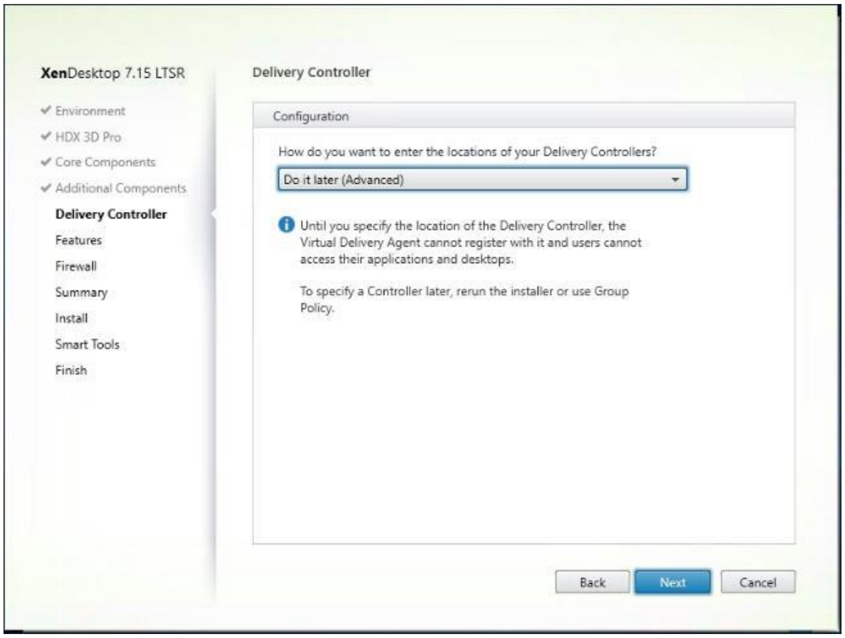

Delivery Controller

Configuration

How do you want to enter the locations of your Delivery Controllers?

Do it later (Advanced)

Until you specify the location of the Delivery Controller, the Virtual Delivery Agent cannot register with it and users cannot access their applications and desktops.

To specify a Controller later, rerun the installer or use Group Policy.

15. Accept the default features**.**

16. Click Next.

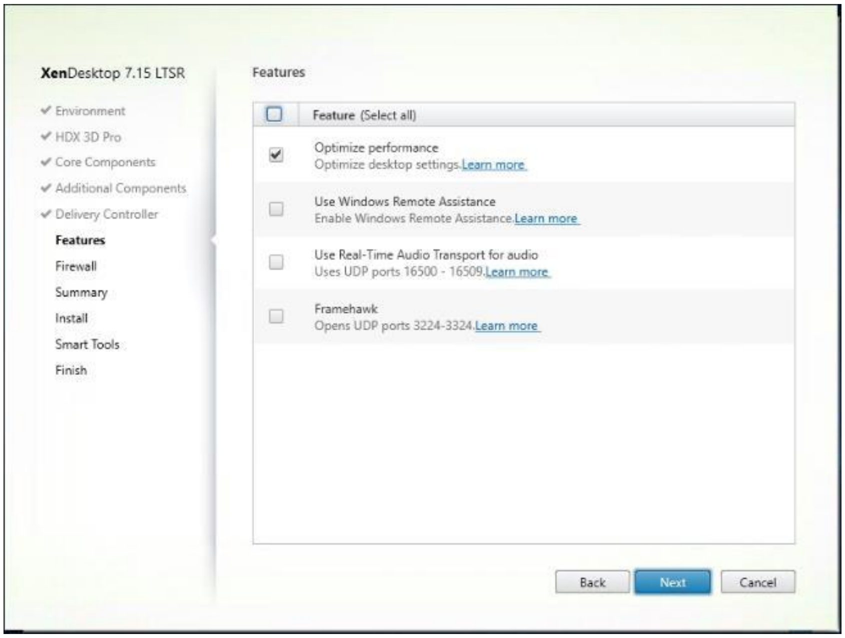

Features

□	Feature (Select all)
☑	Optimize performance Optimize desktop settings.Learn more
□	Use Windows Remote Assistance Enable Windows Remote Assistance.Learn more
□	Use Real-Time Audio Transport for audio Uses UDP ports 16500 - 16509.Learn more
□	Framehawk Opens UDP ports 3224-3324.Learn more

17. Allow the firewall rules to be configured Automatically.

18. Click Next.

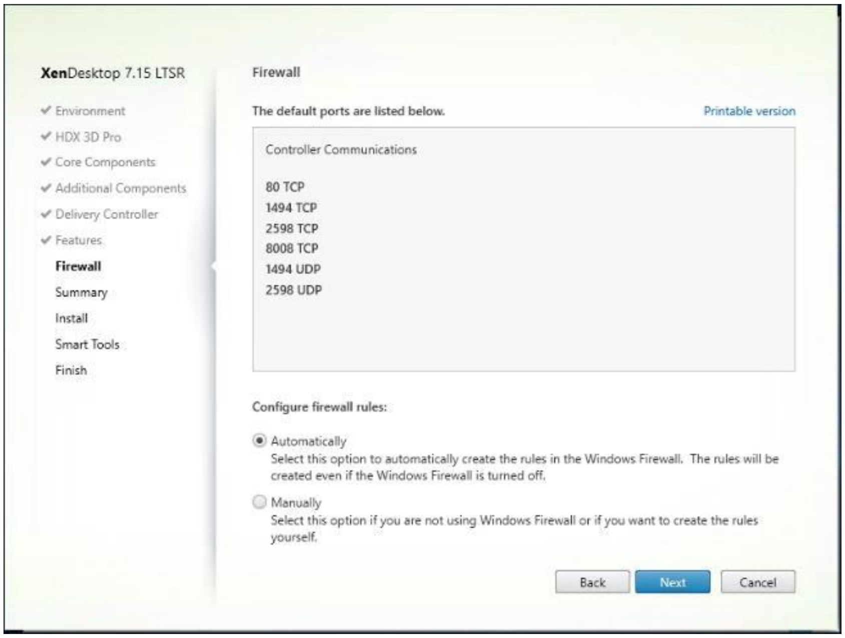

Firewall

The default ports are listed below.

Controller Communications

80 TCP
1494 TCP
2598 TCP
8008 TCP
1494 UDP
2598 UDP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

19. Verify the Summary and click Install**.**

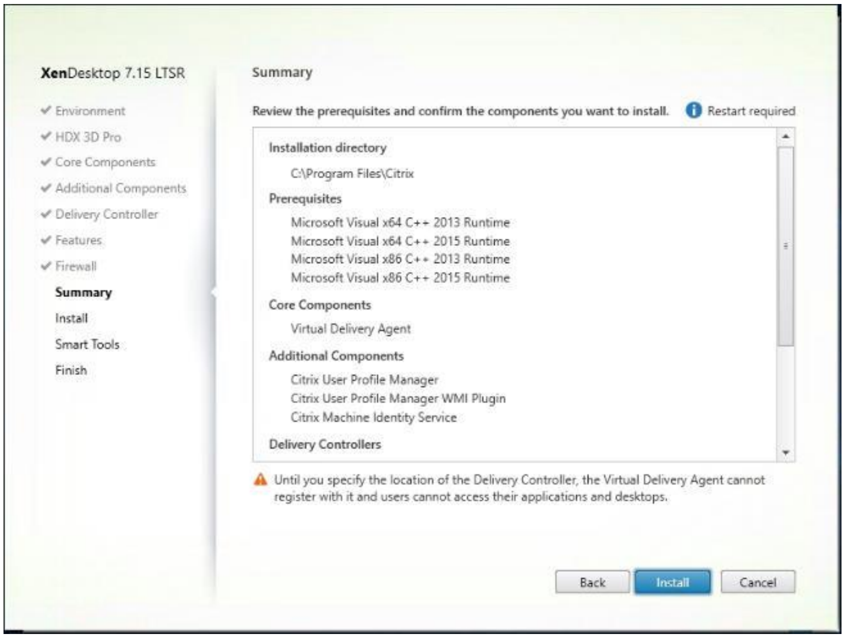

Summary

Review the prerequisites and confirm the components you want to install. Restart required

Installation directory
C:\Program Files\Citrix

Prerequisites
Microsoft Visual x64 C++ 2013 Runtime
Microsoft Visual x64 C++ 2015 Runtime
Microsoft Visual x86 C++ 2013 Runtime
Microsoft Visual x86 C++ 2015 Runtime

Core Components
Virtual Delivery Agent

Additional Components
Citrix User Profile Manager
Citrix User Profile Manager WMI Plugin
Citrix Machine Identity Service

Delivery Controllers

Until you specify the location of the Delivery Controller, the Virtual Delivery Agent cannot register with it and users cannot access their applications and desktops.

The machine will reboot automatically during installation.

20. (Optional) Configure Smart Tools/Call Home participation.

21. Click Next.

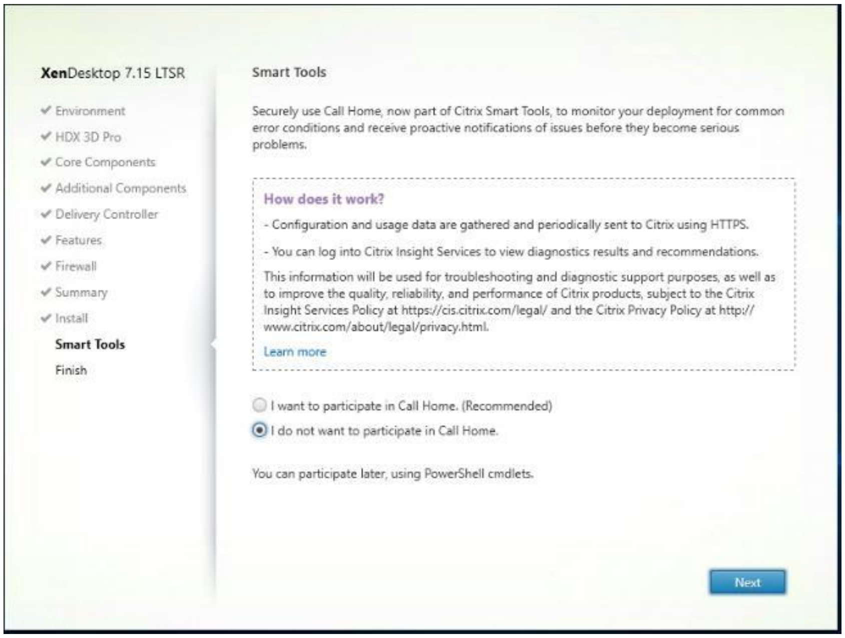

Smart Tools

Securely use Call Home, now part of Citrix Smart Tools, to monitor your deployment for common error conditions and receive proactive notifications of issues before they become serious problems.

How does it work?

- Configuration and usage data are gathered and periodically sent to Citrix using HTTPS.

- You can log into Citrix Insight Services to view diagnostics results and recommendations.

This information will be used for troubleshooting and diagnostic support purposes, as well as to improve the quality, reliability, and performance of Citrix products, subject to the Citrix Insight Services Policy at https://cis.citrix.com/legal/ and the Citrix Privacy Policy at http://www.citrix.com/about/legal/privacy.html.

Learn more

○ I want to participate in Call Home. (Recommended)

● I do not want to participate in Call Home.

You can participate later, using PowerShell cmdlets.

22. Check "Restart Machine."

23. Click **Finish** and the machine will reboot automatically.

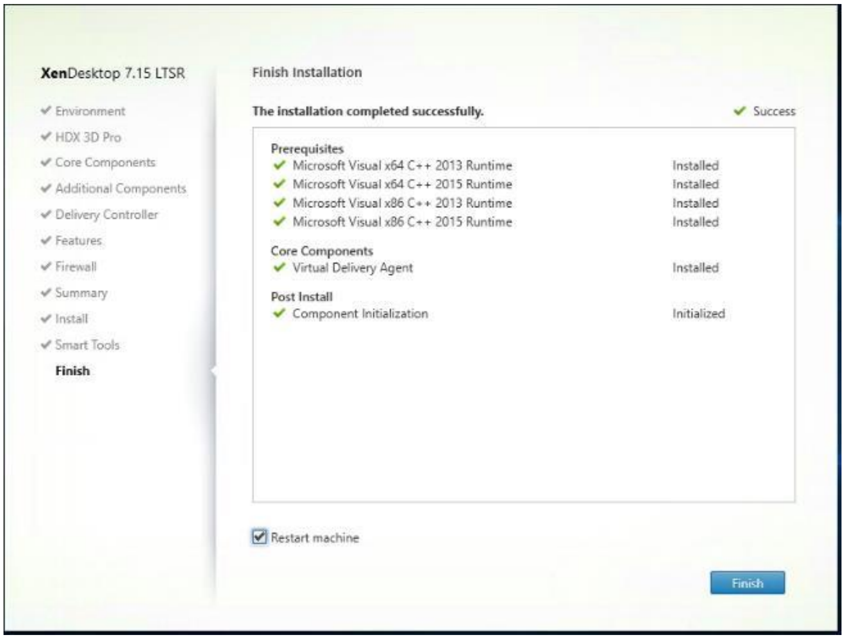

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites	
✓ Microsoft Visual x64 C++ 2013 Runtime	Installed
✓ Microsoft Visual x64 C++ 2015 Runtime	Installed
✓ Microsoft Visual x86 C++ 2013 Runtime	Installed
✓ Microsoft Visual x86 C++ 2015 Runtime	Installed
Core Components	
✓ Virtual Delivery Agent	Installed
Post Install	
✓ Component Initialization	Initialized

☑ Restart machine

Install the Citrix Provisioning Server Target Device Software

The Master Target Device refers to the target device from which a hard disk image is built and stored on a vDisk. Provisioning Services then streams the contents of the vDisk created to other target devices. This procedure installs the PVS Target Device software that is used to build the RDS and VDI golden images.

To install the Citrix Provisioning Server Target Device software, complete the following steps:

The instructions below outline the installation procedure to configure a vDisk for VDI desktops. When you have completed these installation steps, repeat the procedure to configure a vDisk for RDS.

1. Launch the PVS installer from the Provisioning Services 7.15 ISO.

2. Click the Target Device Installation button.

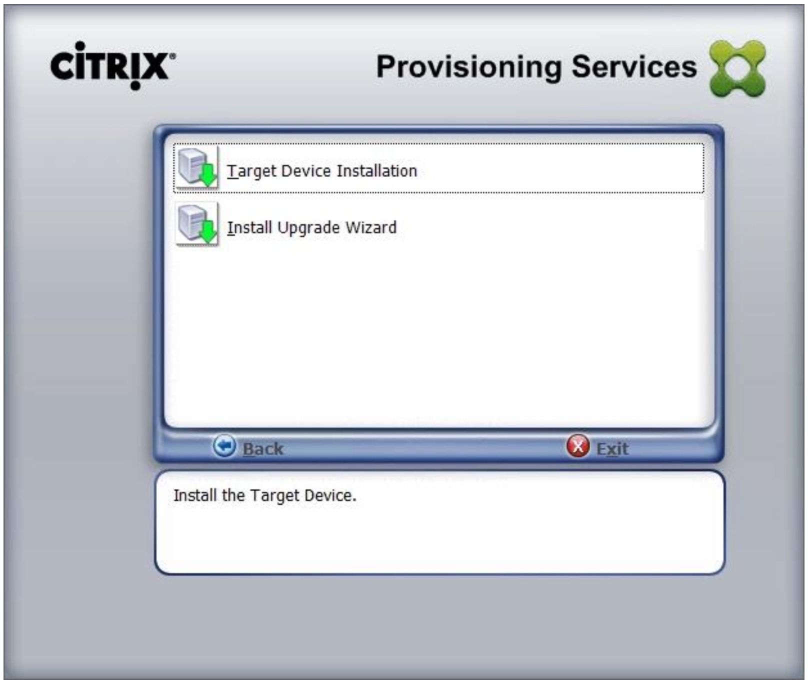

Install the Target Device.

The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

3. Click Next.

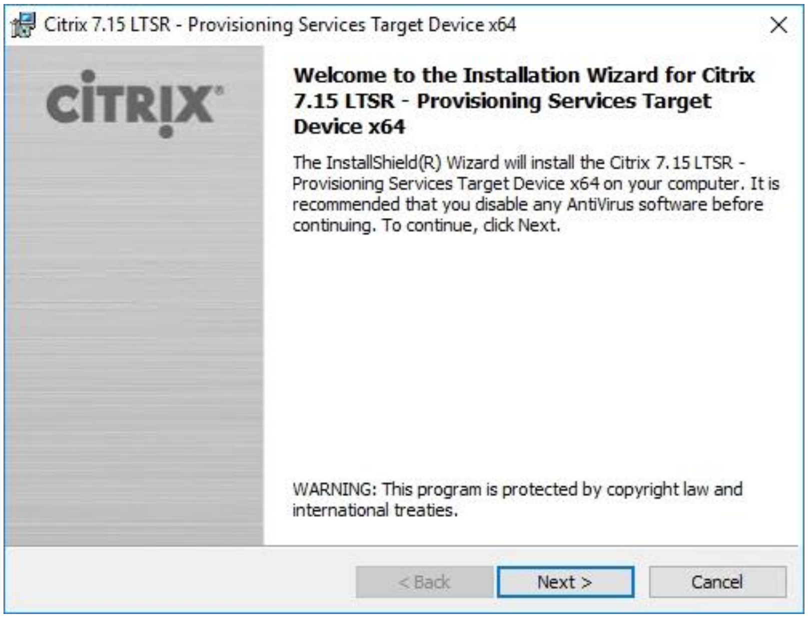

Welcome to the Installation Wizard for Citrix 7.15 LTSR - Provisioning Services Target Device x64

The InstallShield(R) Wizard will install the Citrix 7.15 LTSR - Provisioning Services Target Device x64 on your computer. It is recommended that you disable any AntiVirus software before continuing. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

4. Indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

5. Click Next.

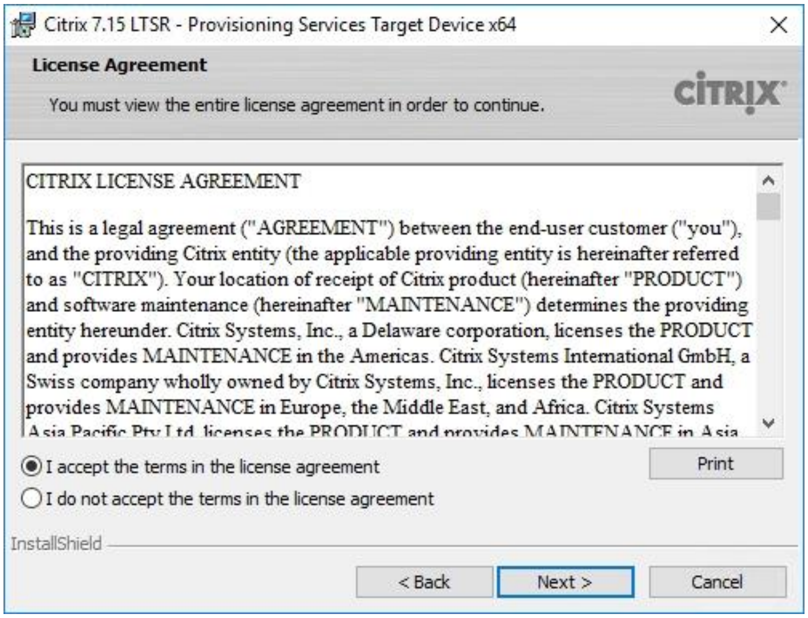

Citrix 7.15 LTSR - Provisioning Services Target Device x64

License Agreement

You must view the entire license agreement in order to continue.

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). Your location of receipt of Citrix product (hereinafter "PRODUCT") and software maintenance (hereinafter "MAINTENANCE") determines the providing entity hereunder. Citrix Systems, Inc., a Delaware corporation, licenses the PRODUCT and provides MAINTENANCE in the Americas. Citrix Systems International GmbH, a Swiss company wholly owned by Citrix Systems, Inc., licenses the PRODUCT and provides MAINTENANCE in Europe, the Middle East, and Africa. Citrix Systems Asia Pacific Pty Ltd licenses the PRODUCT and provides MAINTENANCE in Asia

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

6. Optionally. provide the Customer information.

7. Click Next.

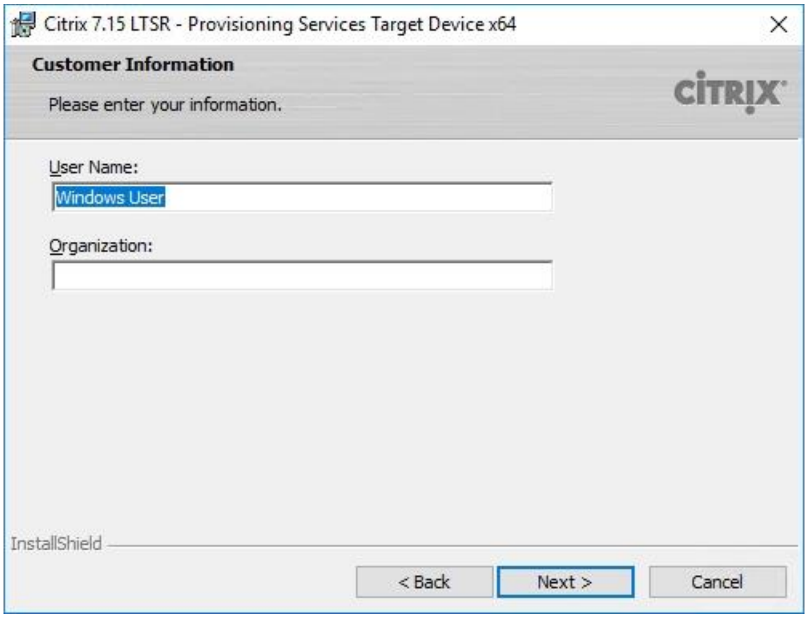

Citrix 7.15 LTSR - Provisioning Services Target Device x64

Customer Information

Please enter your information.

User Name: Windows User

Organization:

8. Accept default installation path.

9. Click Next.

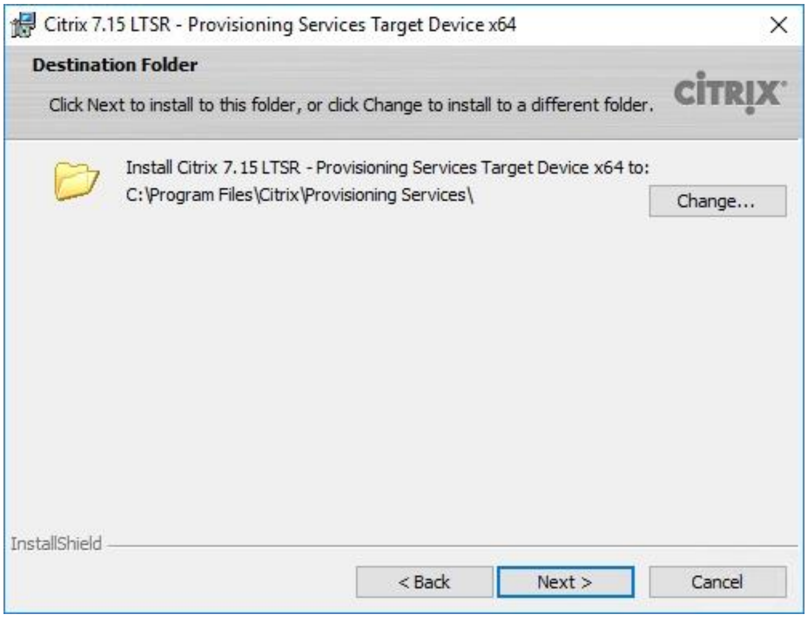

Citrix 7.15 LTSR - Provisioning Services Target Device x64

Destination Folder

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

Install Citrix 7.15 LTSR - Provisioning Services Target Device x64 to:
C:\Program Files\Citrix\Provisioning Services\

10. Click Install.

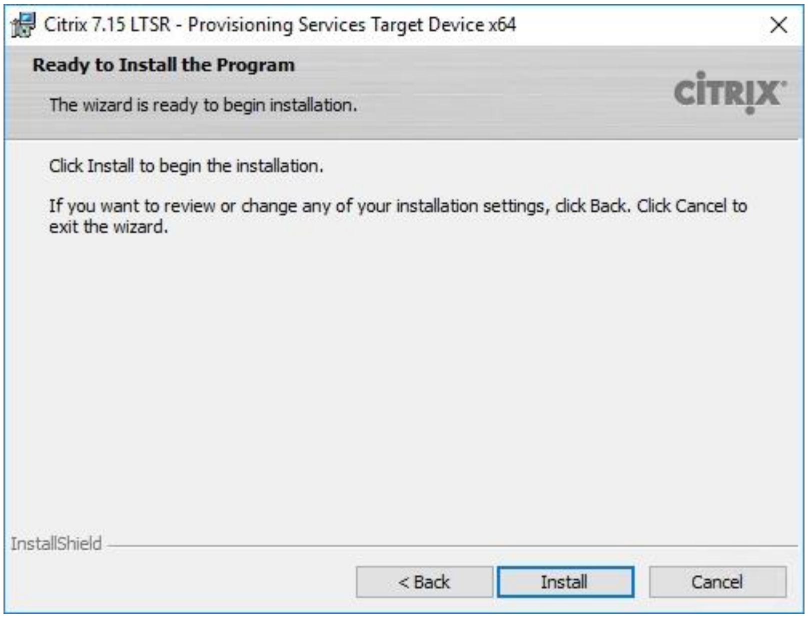

Citrix 7.15 LTSR - Provisioning Services Target Device x64

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

11. Deselect the checkbox to launch the Imaging Wizard and click Finish.

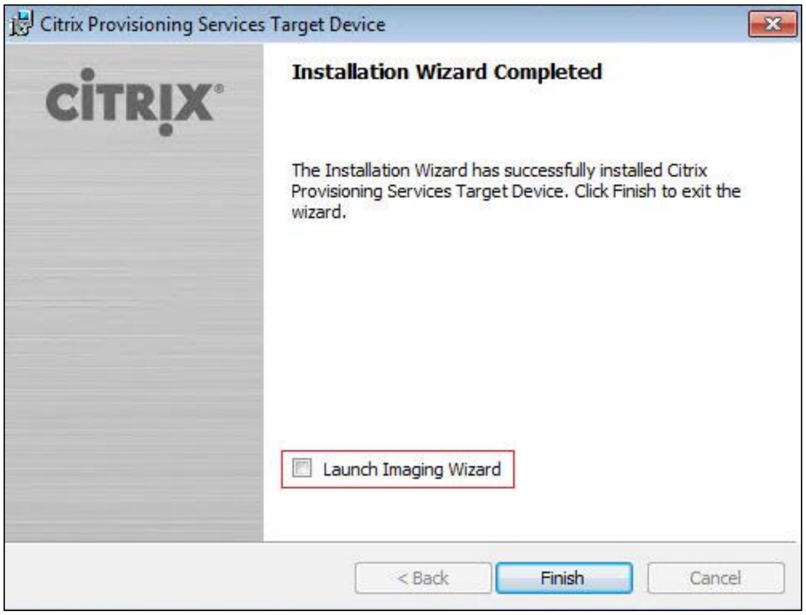
Citrix Provisioning Services Target Device

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix Provisioning Services Target Device. Click Finish to exit the wizard.

☐ Launch Imaging Wizard

12. Click Yes to reboot the machine.

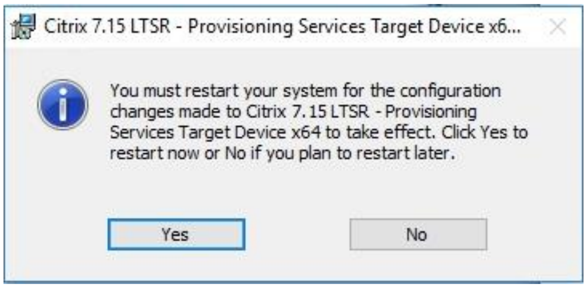
Citrix 7.15 LTSR - Provisioning Services Target Device x6...

You must restart your system for the configuration changes made to Citrix 7.15 LTSR - Provisioning Services Target Device x64 to take effect. Click Yes to restart now or No if you plan to restart later.

Create Citrix Provisioning Server vDisks

The PVS Imaging Wizard automatically creates a base vDisk image from the master target device. To create the Citrix Provisioning Server vDisks, complete the following steps:

The instructions below describe the process of creating a vDisk for VDI desktops. When you have completed these steps, repeat the procedure to build a vDisk for RDS.

1. The PVS Imaging Wizard's Welcome page appears.

2. Click **Next**.

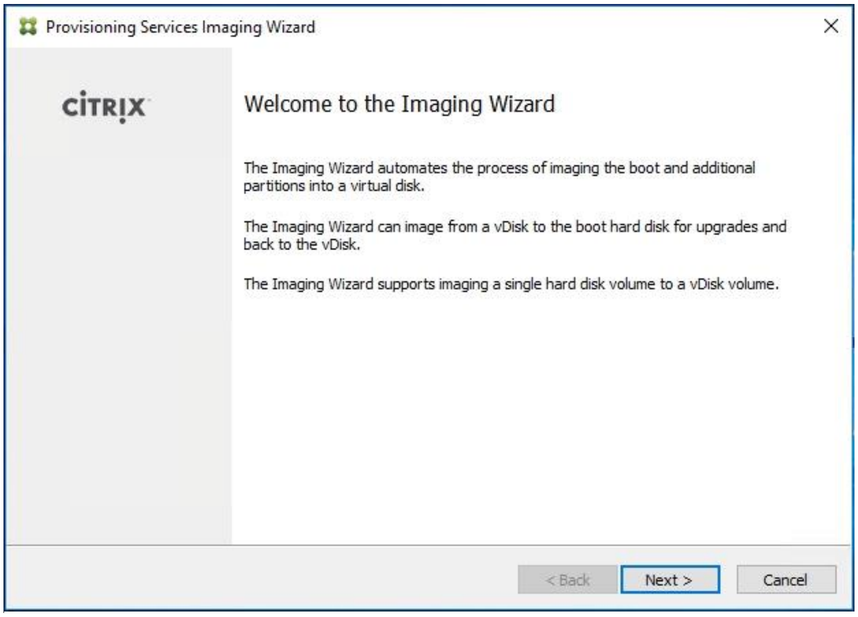

Welcome to the Imaging Wizard

The Imaging Wizard automates the process of imaging the boot and additional partitions into a virtual disk.

The Imaging Wizard can image from a vDisk to the boot hard disk for upgrades and back to the vDisk.

The Imaging Wizard supports imaging a single hard disk volume to a vDisk volume.

3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Server within the farm to connect to and the port to use to make that connection.

4. Use the Windows credentials (default) or enter different credentials.

5. Click Next.

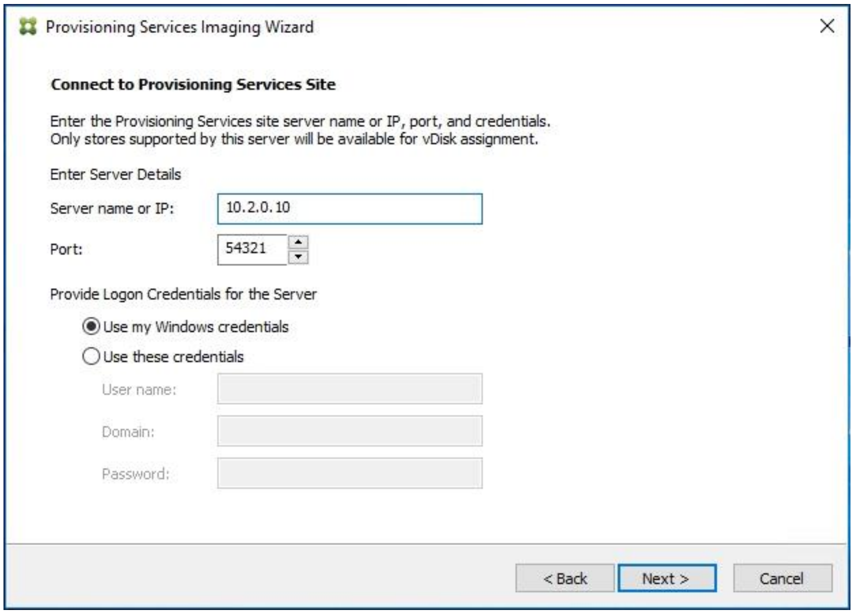

Connect to Provisioning Services Site

Enter the Provisioning Services site server name or IP, port, and credentials.
Only stores supported by this server will be available for vDisk assignment.

Enter Server Details

Server name or IP: 10.2.0.10

Port: 54321

Provide Logon Credentials for the Server

● Use my Windows credentials

○ Use these credentials

User name:

Domain:

Password:

6. Select Create new vDisk.

7. Click Next.

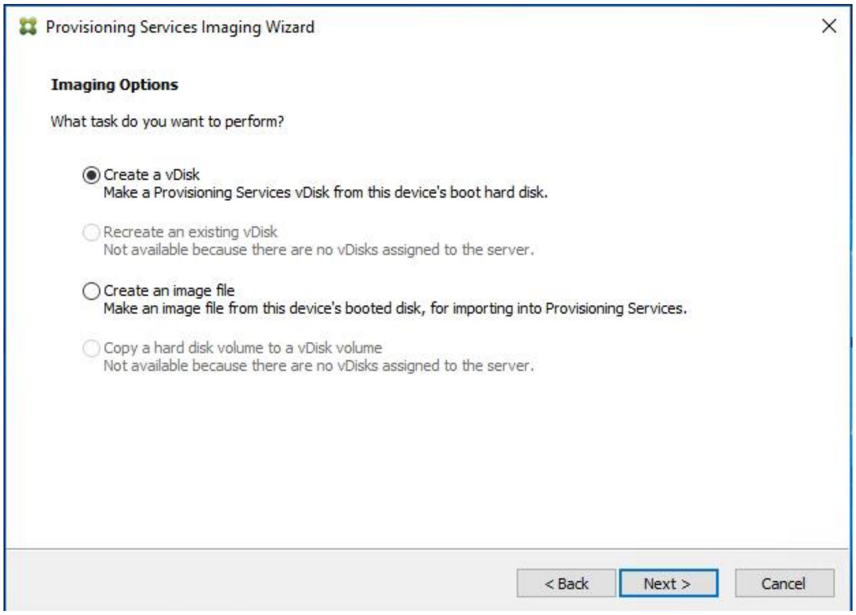

Provisioning Services Imaging Wizard

Imaging Options

What task do you want to perform?

● Create a vDisk
Make a Provisioning Services vDisk from this device's boot hard disk.

○ Recreate an existing vDisk
Not available because there are no vDisks assigned to the server.

○ Create an image file
Make an image file from this device's booted disk, for importing into Provisioning Services.

○ Copy a hard disk volume to a vDisk volume
Not available because there are no vDisks assigned to the server.

8. The Add Target Device page appears.

9. Select the Target Device Name, the MAC address associated with one of the NICs that was selected when the target device software was installed on the master target device, and the Collection to which you are adding the device.

10. Click Next.

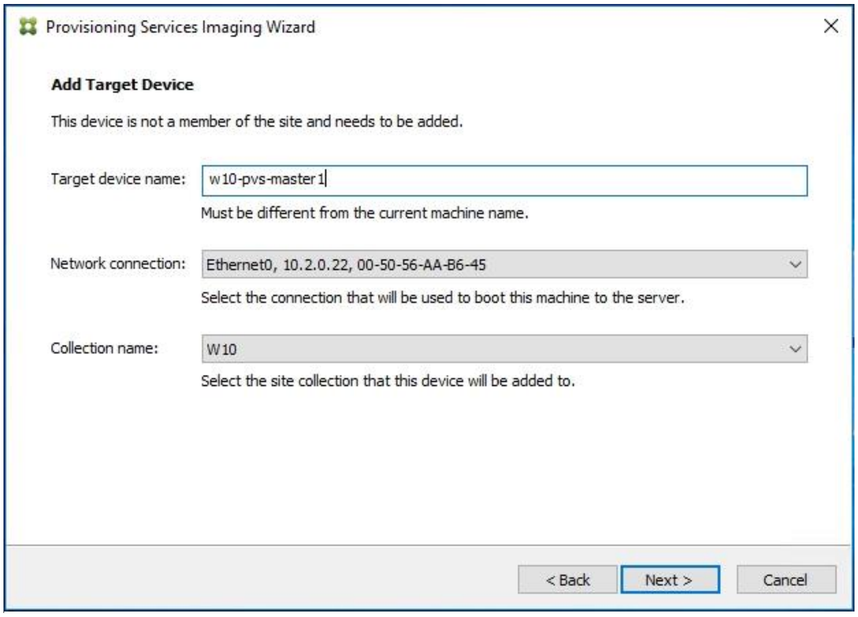

Provisioning Services Imaging Wizard

Add Target Device

This device is not a member of the site and needs to be added.

Target device name: w10-pvs-master1
Must be different from the current machine name.

Network connection: Ethernet0, 10.2.0.22, 00-50-56-AA-B6-45
Select the connection that will be used to boot this machine to the server.

Collection name: W10
Select the site collection that this device will be added to.

11. The New vDisk dialog displays. Enter the name of the vDisk.

12. Select the Store where the vDisk will reside. Select the vDisk type, either Fixed or Dynamic, from the drop-down list.

This CVD used Dynamic rather than Fixed vDisks.

13. Click Next.

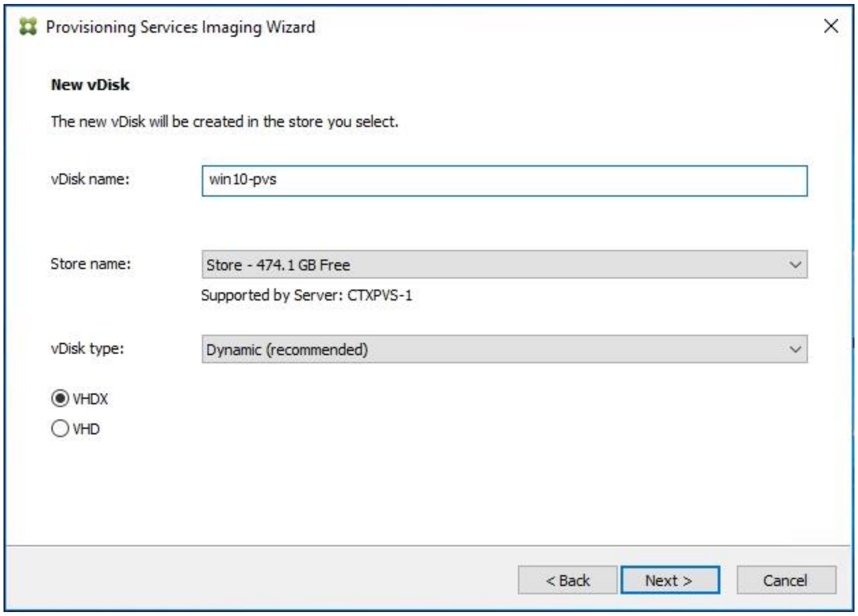

Provisioning Services Imaging Wizard

New vDisk

The new vDisk will be created in the store you select.

vDisk name: win10-pvs

Store name: Store - 474.1 GB Free

Supported by Server: CTXPVS-1

vDisk type: Dynamic (recommended)

● VHDX

○ VHD

14. On the Microsoft Volume Licensing page, select the volume license option to use for target devices. For this CVD, volume licensing is not used, so the None button is selected.

15. Click Next.

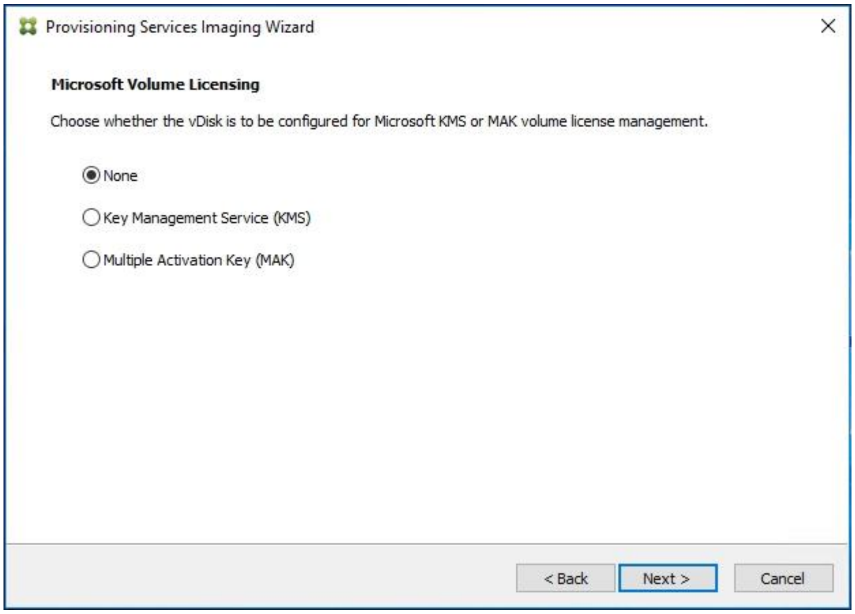

Provisioning Services Imaging Wizard

Microsoft Volume Licensing

Choose whether the vDisk is to be configured for Microsoft KMS or MAK volume license management.

● None

○ Key Management Service (KMS)

○ Multiple Activation Key (MAK)

16. Select Image entire boot disk on the Configure Image Volumes page.

17. Click Next.

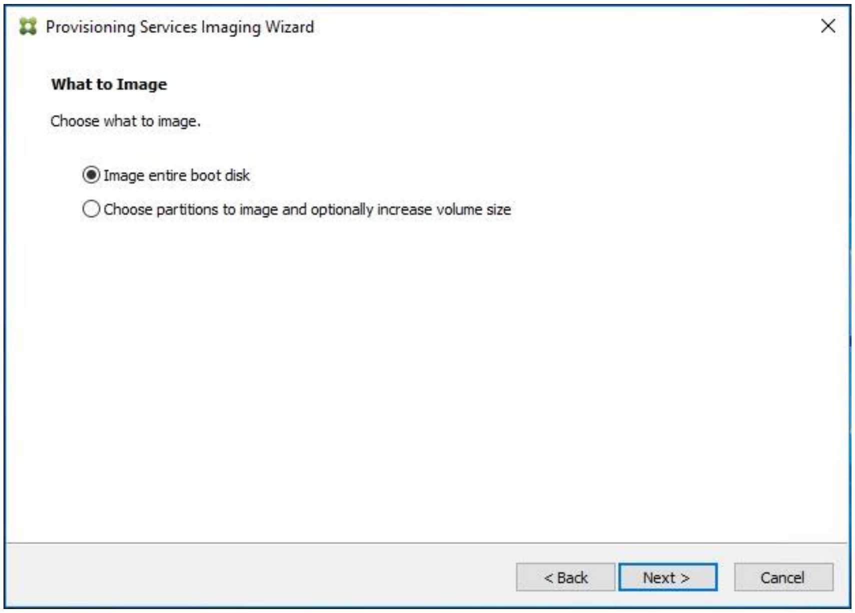

What to Image

Choose what to image.

● Image entire boot disk

○ Choose partitions to image and optionally increase volume size

18. Select Optimize for hard disk again for Provisioning Services before imaging on the Optimize Hard Disk for Provisioning Services.

19. Click **Next**.

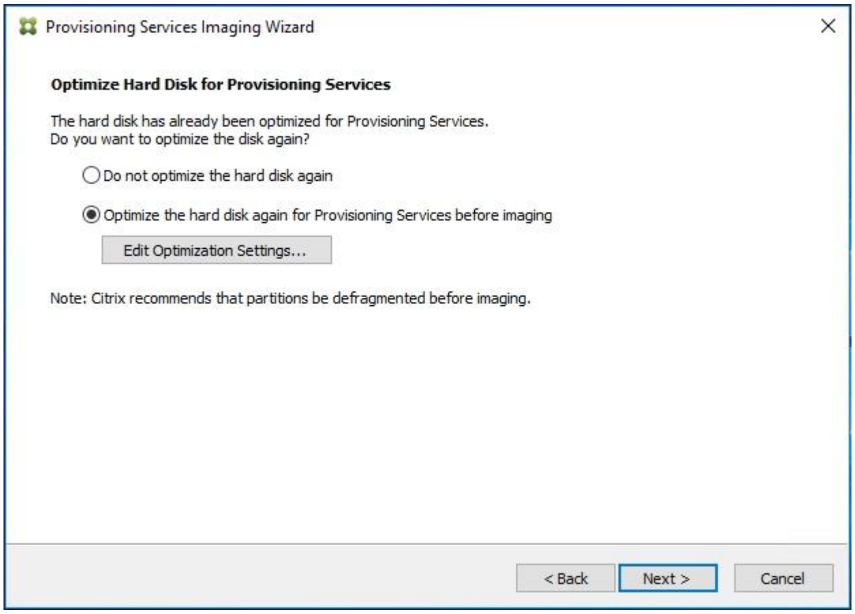
Provisioning Services Imaging Wizard
Edit Optimization Settings...
< Back
Next >
Cancel

Optimize Hard Disk for Provisioning Services

The hard disk has already been optimized for Provisioning Services.
Do you want to optimize the disk again?

○ Do not optimize the hard disk again

● Optimize the hard disk again for Provisioning Services before imaging

Note: Citrix recommends that partitions be defragmented before imaging.

20. Select Create on the Summary page.

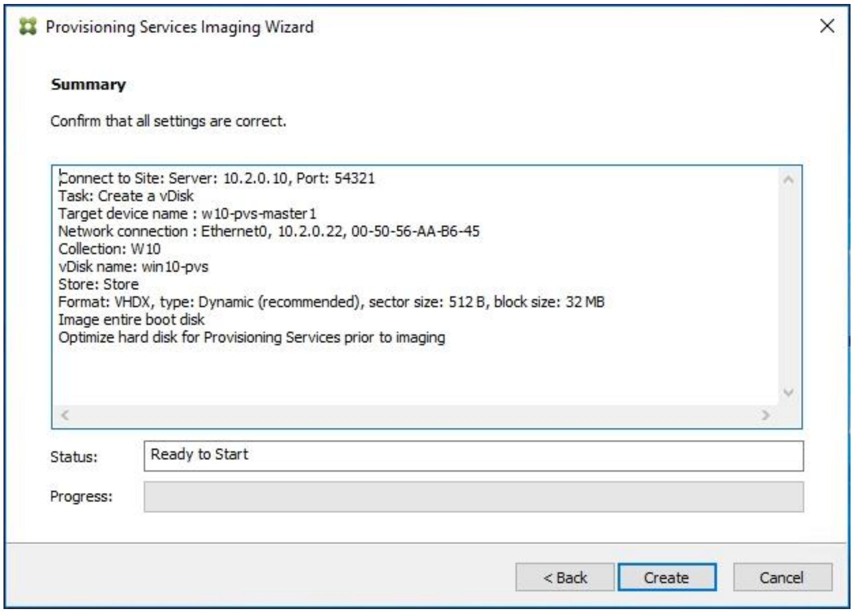

Summary

Confirm that all settings are correct.

```
Connect to Site: Server: 10.2.0.10, Port: 54321
Task: Create a vDisk
Target device name : w10-pvs-master1
Network connection : Ethernet0, 10.2.0.22, 00-50-56-AA-B6-45
Collection: W10
vDisk name: win10-pvs
Store: Store
Format: VHDX, type: Dynamic (recommended), sector size: 512 B, block size: 32 MB
Image entire boot disk
Optimize hard disk for Provisioning Services prior to imaging
```

Status: Ready to Start

Progress:

21. Review the configuration and click Continue.

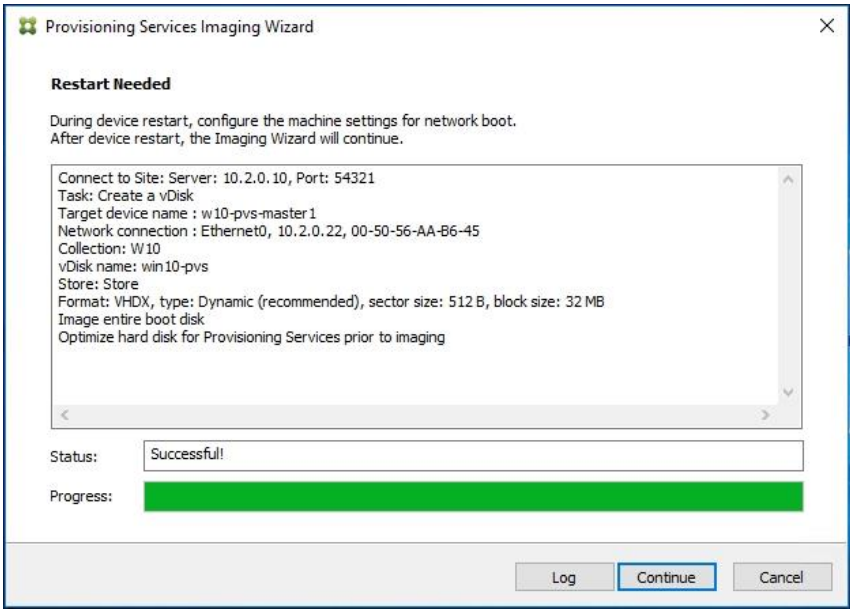

Restart Needed

During device restart, configure the machine settings for network boot.
After device restart, the Imaging Wizard will continue.

```
Connect to Site: Server: 10.2.0.10, Port: 54321
Task: Create a vDisk
Target device name : w10-pvs-master1
Network connection : Ethernet0, 10.2.0.22, 00-50-56-AA-B6-45
Collection: W10
vDisk name: win10-pvs
Store: Store
Format: VHDX, type: Dynamic (recommended), sector size: 512 B, block size: 32 MB
Image entire boot disk
Optimize hard disk for Provisioning Services prior to imaging
```

Status: Successful!

Progress:

22. When prompted, click **No** to shut down the machine.

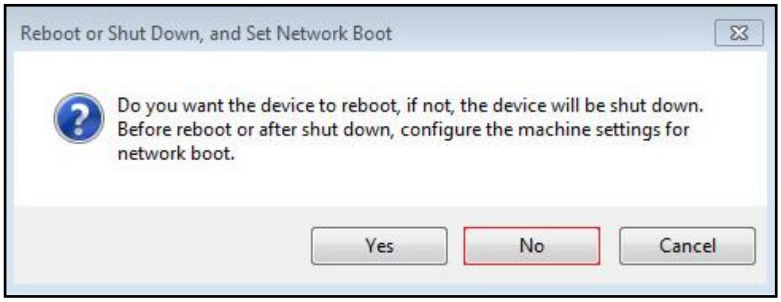

Do you want the device to reboot, if not, the device will be shut down. Before reboot or after shut down, configure the machine settings for network boot.

23. Edit the VM settings and select Force BIOS Setup under Boot Options.

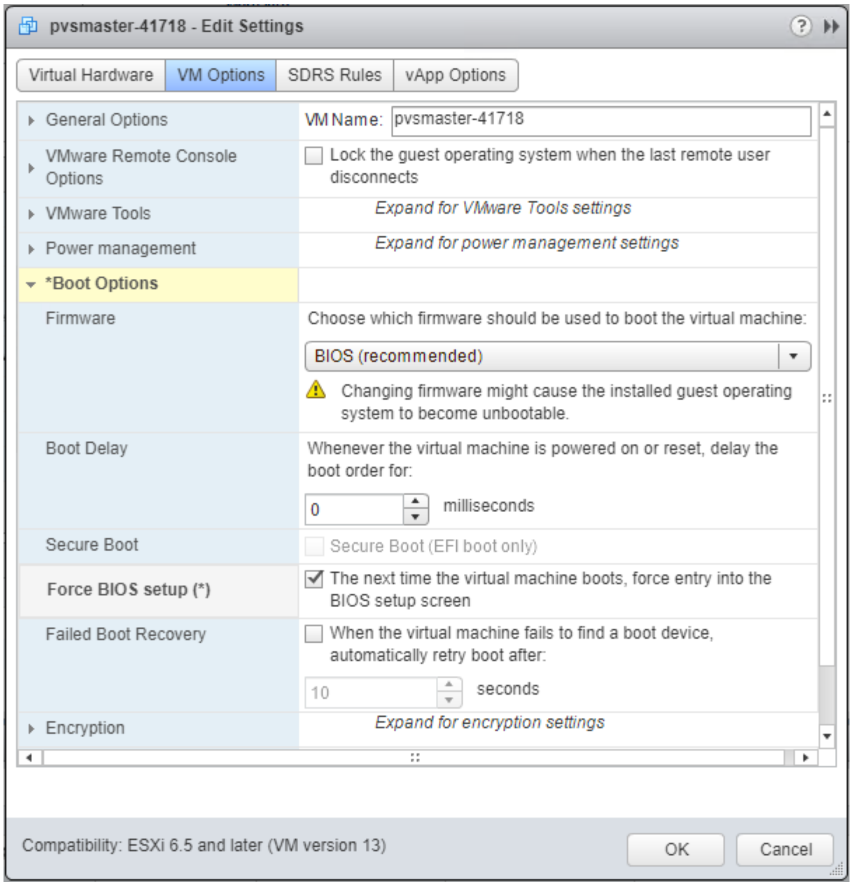

pvsmaster-41718 - Edit Settings

- General Options — VM Name: pvsmaster-41718
- VMware Remote Console Options — □ Lock the guest operating system when the last remote user disconnects
- VMware Tools — *Expand for VMware Tools settings*
- Power management — *Expand for power management settings*
- *Boot Options
  - Firmware — Choose which firmware should be used to boot the virtual machine: BIOS (recommended)
    ⚠ Changing firmware might cause the installed guest operating system to become unbootable.
  - Boot Delay — Whenever the virtual machine is powered on or reset, delay the boot order for: 0 milliseconds
  - Secure Boot — □ Secure Boot (EFI boot only)
  - Force BIOS setup (*) — ☑ The next time the virtual machine boots, force entry into the BIOS setup screen
  - Failed Boot Recovery — □ When the virtual machine fails to find a boot device, automatically retry boot after: 10 seconds
- Encryption — *Expand for encryption settings*

Compatibility: ESXi 6.5 and later (VM version 13)

24. Configure the BIOS/VM settings for PXE/network boot, putting Network boot from VMware VMXNET3 at the top of the boot device list.

25. Select Exit Saving Changes.

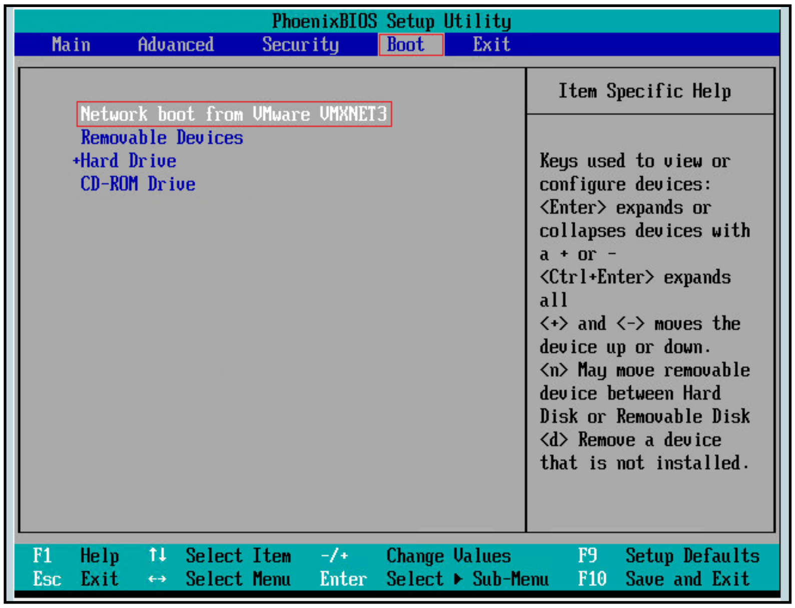

```
PhoenixBIOS Setup Utility
  Main    Advanced    Security   Boot   Exit

    Network boot from VMware VMXNET3        Item Specific Help
    Removable Devices
   +Hard Drive                              Keys used to view or
    CD-ROM Drive                            configure devices:
                                            <Enter> expands or
                                            collapses devices with
                                            a + or -
                                            <Ctrl+Enter> expands
                                            all
                                            <+> and <-> moves the
                                            device up or down.
                                            <n> May move removable
                                            device between Hard
                                            Disk or Removable Disk
                                            <d> Remove a device
                                            that is not installed.

F1   Help  ↑↓  Select Item   -/+    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu   Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

After restarting the VM, log into the HVD or HSD master target. The PVS imaging process begins, copying the contents of the C: drive to the PVS vDisk located on the server.

26. If prompted to Restart select Restart Later.

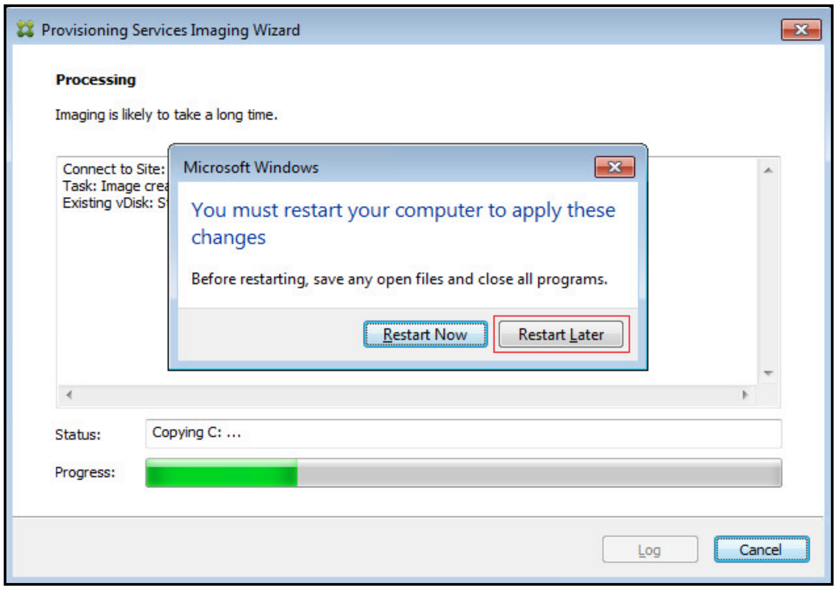

Processing

Imaging is likely to take a long time.

```
Connect to Site:
Task: Image crea
Existing vDisk: S
```

You must restart your computer to apply these changes

Before restarting, save any open files and close all programs.

Status: Copying C: ...

Progress:

27. A message is displayed when the conversion is complete, click **Done**.

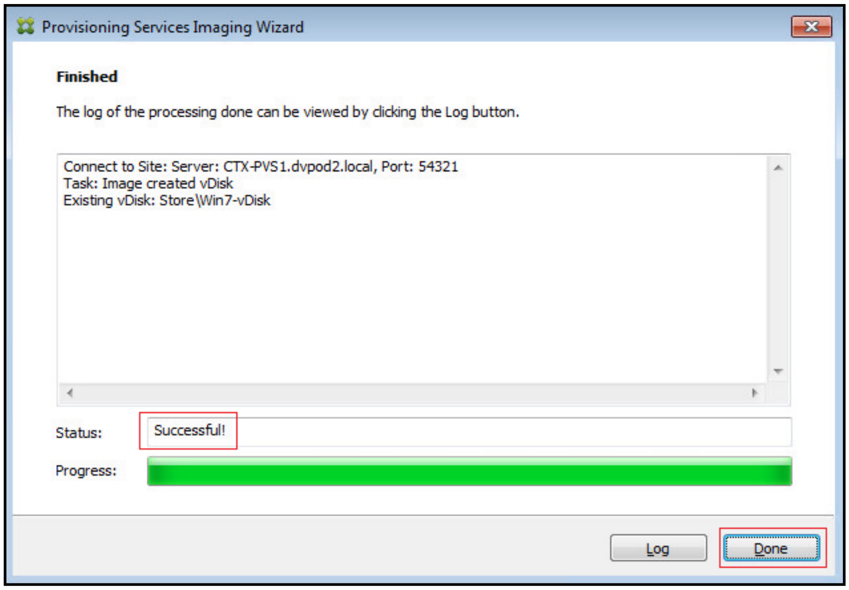

Finished

The log of the processing done can be viewed by clicking the Log button.

```
Connect to Site: Server: CTX-PVS1.dvpod2.local, Port: 54321
Task: Image created vDisk
Existing vDisk: Store\Win7-vDisk
```

Status: Successful!

Progress:

28. Shutdown the VM used as the VDI or RDS master target.

29. Connect to the PVS server and validate that the vDisk image is available in the Store.

30. Right-click the newly created vDisk and select **Properties**.

31. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access).”

32. Set the Cache Type to “Cache in device RAM with overflow on hard disk.”

33. Set Maximum RAM size (MBs): 128 for HVD and set 2048 MB for HSD vDisk.

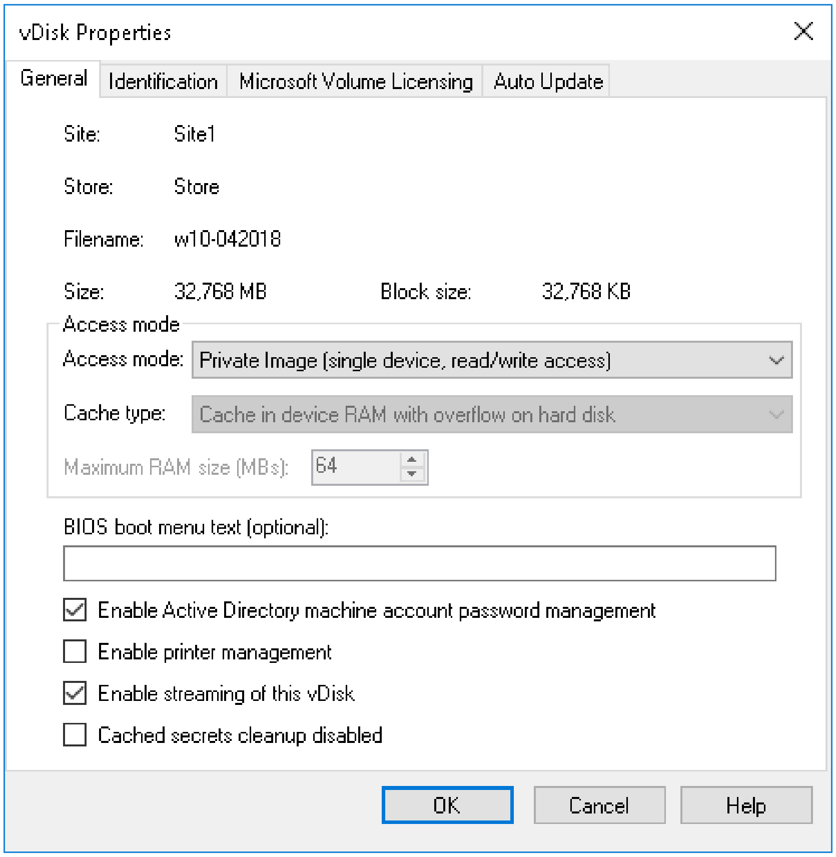

vDisk Properties

Site: Site1

Store: Store

Filename: w10-042018

Size: 32,768 MB Block size: 32,768 KB

Access mode

Access mode: Private Image (single device, read/write access)

Cache type: Cache in device RAM with overflow on hard disk

Maximum RAM size (MBs): 64

BIOS boot menu text (optional):

☑ Enable Active Directory machine account password management

□ Enable printer management

☑ Enable streaming of this vDisk

□ Cached secrets cleanup disabled

34. Click **OK**

> Repeat this procedure to create vDisks for both the Hosted VDI Desktops (using the Windows 10 OS image) and the Hosted Shared Desktops (using the Windows Server 2016 image).

Provision Virtual Desktop Machines

Citrix Provisioning Services XenDesktop Setup Wizard

To create virtual desktop machines, complete the following steps:

1. Create diskless Master Target Virtual Machine (VM).

> Hard disk is not required to provision desktop machines as the XenDesktop Setup Wizard dynamically creates the write cache disk.

2. Select the Master Target VM from the vSphere Client.

3. Right-click the VM and select Clone to Template...

4. Name the cloned VM Desktop-Template.

5. Select the cluster and datastore where the first phase of provisioning will occur.

6. Click Finish

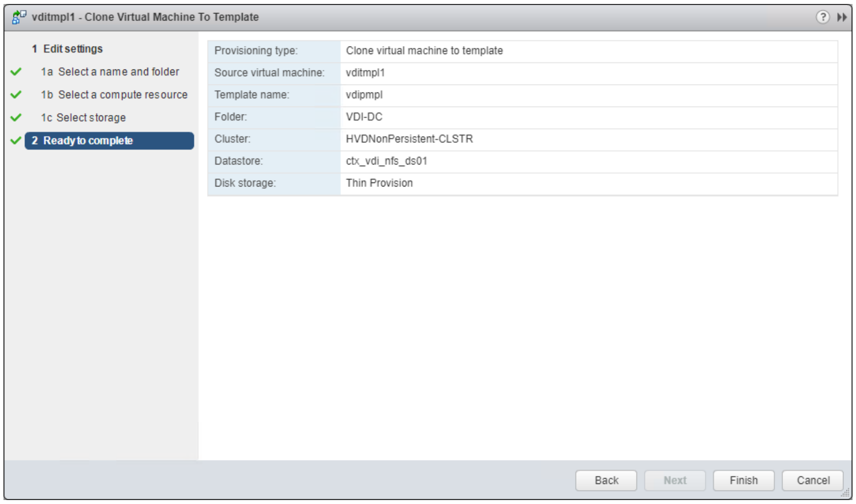
vditmpl1 - Clone Virtual Machine To Template
1 Edit settings
1a Select a name and folder
1b Select a compute resource
1c Select storage
2 Ready to complete
Back
Next
Finish
Cancel

Provisioning type:	Clone virtual machine to template
Source virtual machine:	vditmpl1
Template name:	vdipmpl
Folder:	VDI-DC
Cluster:	HVDNonPersistent-CLSTR
Datastore:	ctx_vdi_nfs_ds01
Disk storage:	Thin Provision

7. Start the XenDesktop Setup Wizard from the Provisioning Services Console.

8. Right-click the Site.

9. Choose XenDesktop Setup Wizard… from the context menu.

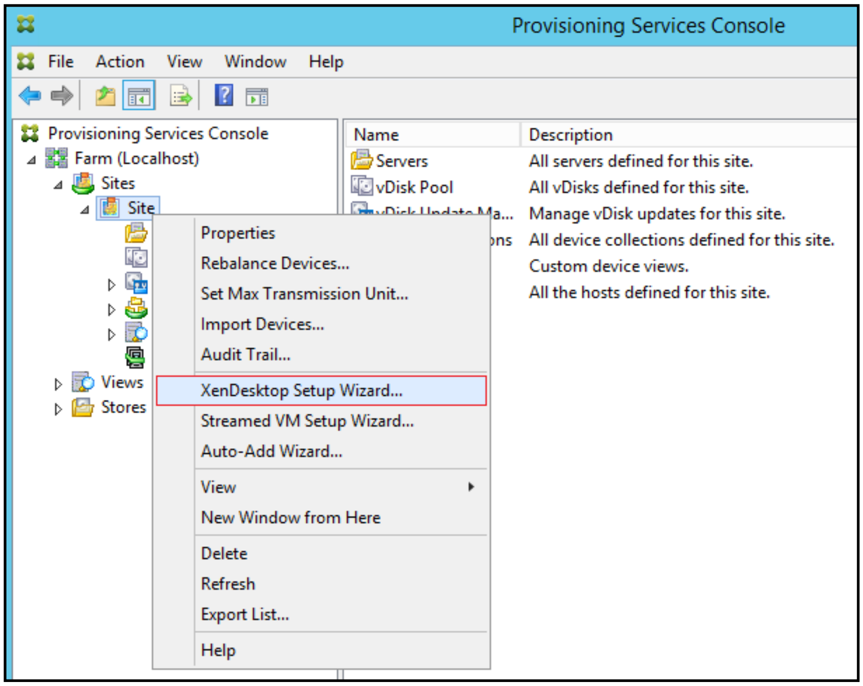

Name	Description
Servers	All servers defined for this site.
vDisk Pool	All vDisks defined for this site.
vDisk Update Ma...	Manage vDisk updates for this site.
ons	All device collections defined for this site.
	Custom device views.
	All the hosts defined for this site.

10. Click **Next**.

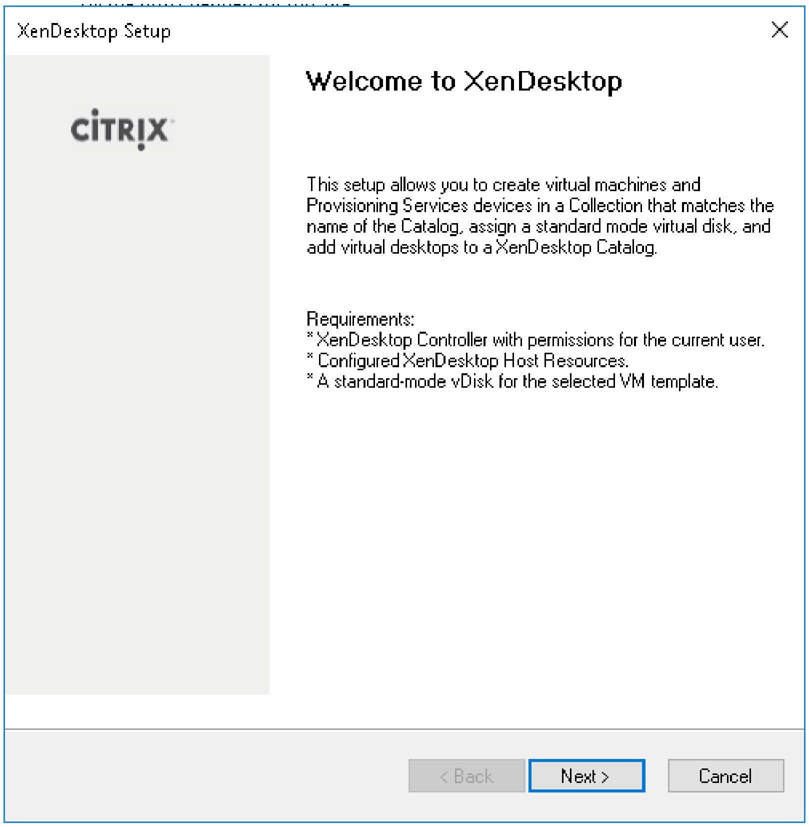

Welcome to XenDesktop

This setup allows you to create virtual machines and Provisioning Services devices in a Collection that matches the name of the Catalog, assign a standard mode virtual disk, and add virtual desktops to a XenDesktop Catalog.

Requirements:
* XenDesktop Controller with permissions for the current user.
* Configured XenDesktop Host Resources.
* A standard-mode vDisk for the selected VM template.

11. Enter the XenDesktop Controller address that will be used for the wizard operations.

12. Click Next.

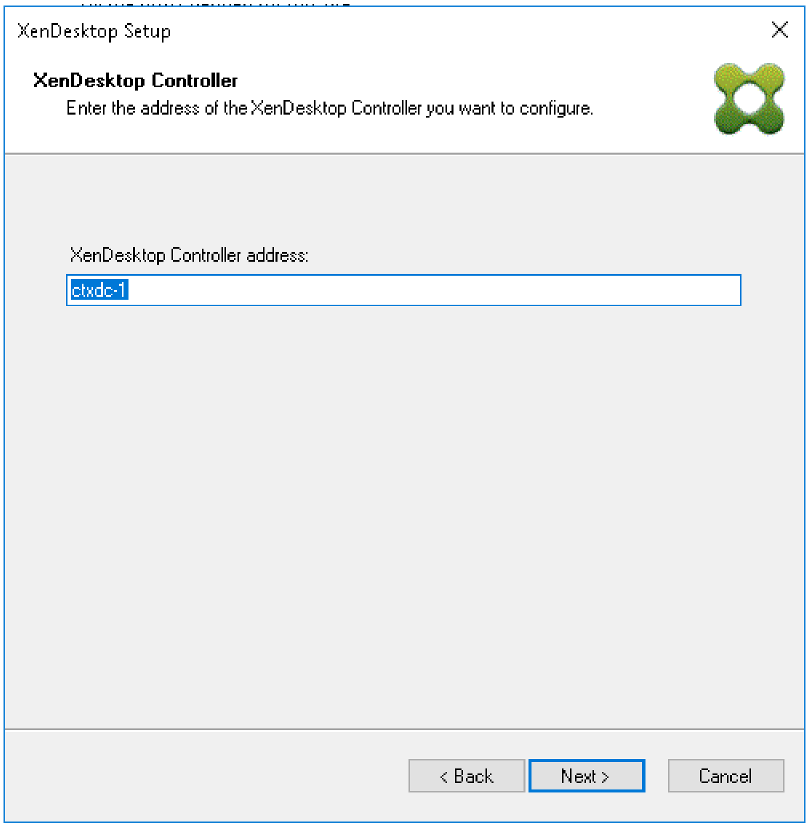

XenDesktop Setup

XenDesktop Controller

Enter the address of the XenDesktop Controller you want to configure.

XenDesktop Controller address: ctxdc-1

13. Select the Host Resources on which the virtual machines will be created.

14. Click Next.

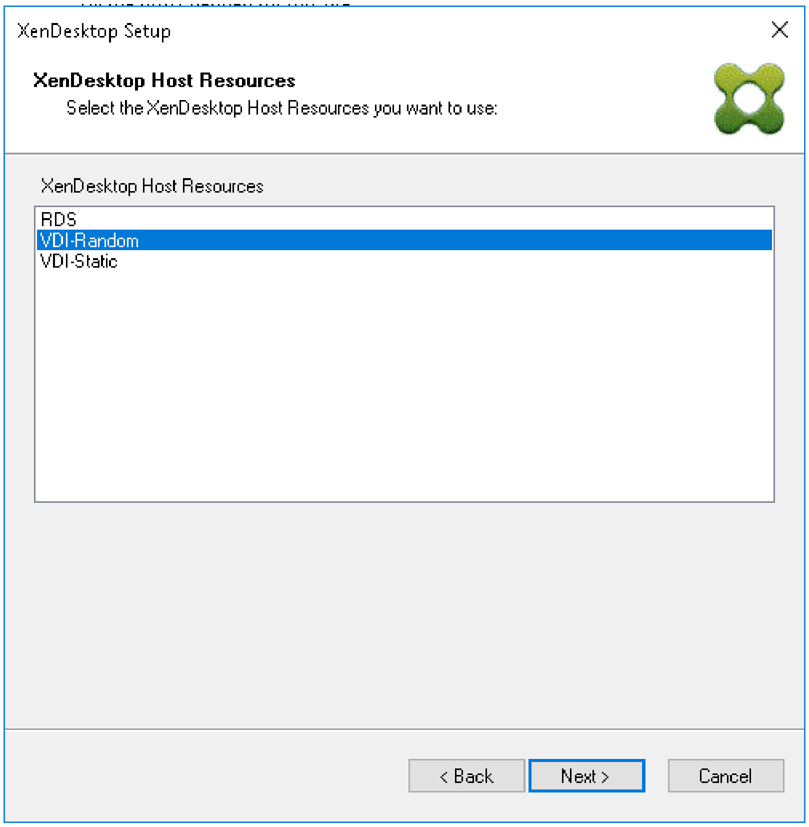

XenDesktop Setup

XenDesktop Host Resources

Select the XenDesktop Host Resources you want to use:

XenDesktop Host Resources

- RDS
- VDI-Random
- VDI-Static

15. Provide the Host Resources Credentials (Username and Password) to the XenDesktop controller when prompted.

16. Click OK.

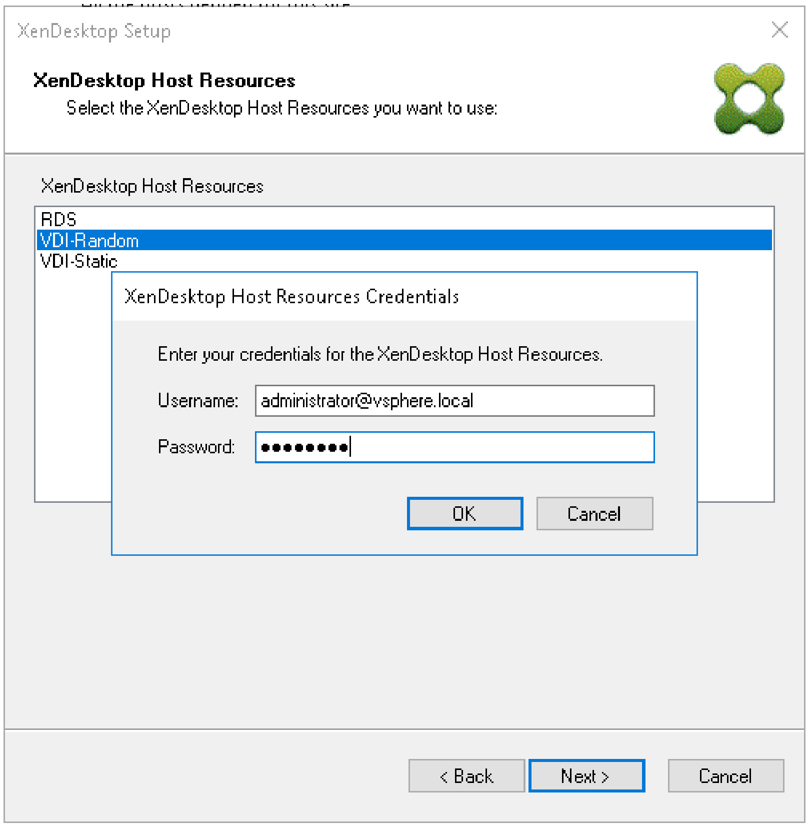

XenDesktop Setup

XenDesktop Host Resources

Select the XenDesktop Host Resources you want to use:

XenDesktop Host Resources

RDS
VDI-Random
VDI-Static

XenDesktop Host Resources Credentials

Enter your credentials for the XenDesktop Host Resources.

Username: administrator@vsphere.local

Password: ••••••••

17. Select the Template created earlier.

18. Click Next.

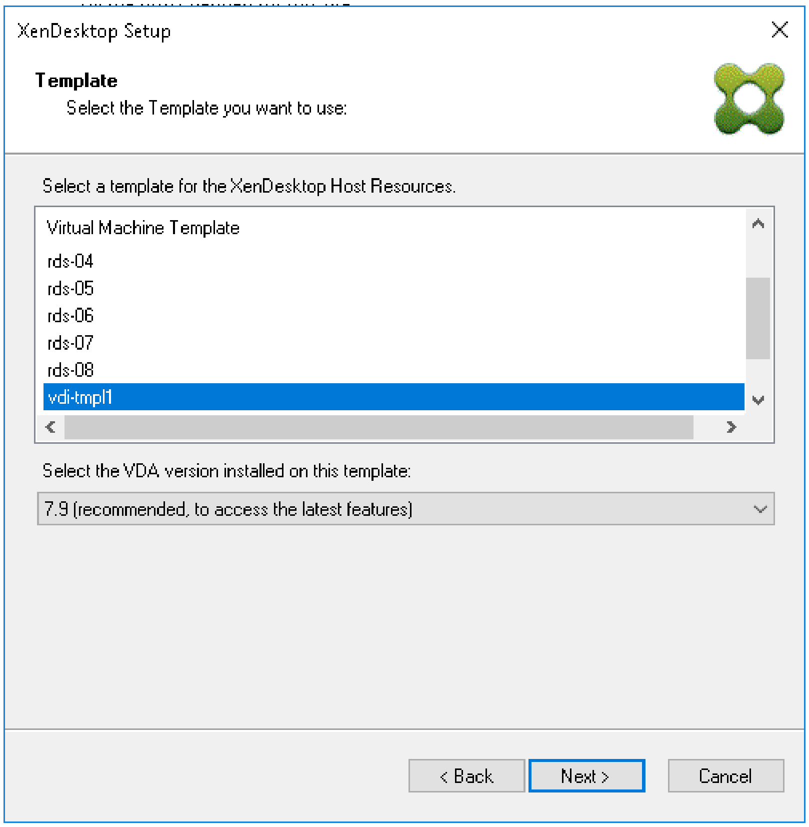

XenDesktop Setup

Template

Select the Template you want to use:

Select a template for the XenDesktop Host Resources.

Virtual Machine Template
- rds-04
- rds-05
- rds-06
- rds-07
- rds-08
- vdi-tmpl1

Select the VDA version installed on this template:

7.9 (recommended, to access the latest features)

19. Select the virtual disk (vDisk) that will be used to stream the provisioned virtual machines.

20. Click Next.

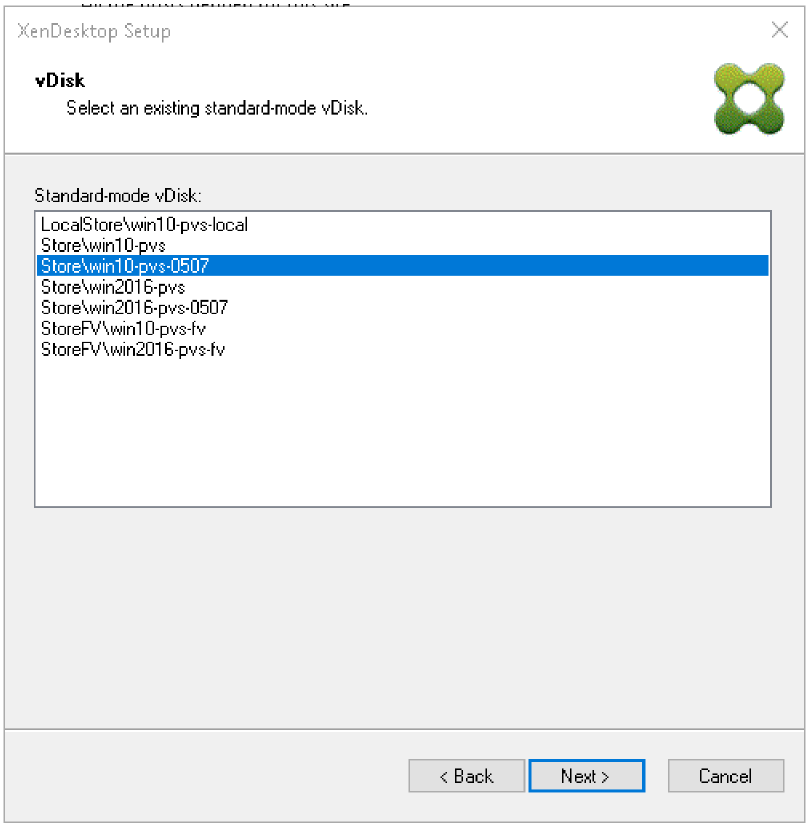

vDisk

Select an existing standard-mode vDisk.

Standard-mode vDisk:

```
LocalStore\win10-pvs-local
Store\win10-pvs
Store\win10-pvs-0507
Store\win2016-pvs
Store\win2016-pvs-0507
StoreFV\win10-pvs-fv
StoreFV\win2016-pvs-fv
```

21. Select “Create a new catalog.”

22. Type in Catalog name

> The catalog name is also used as the collection name in the PVS site.

23. Click Next.

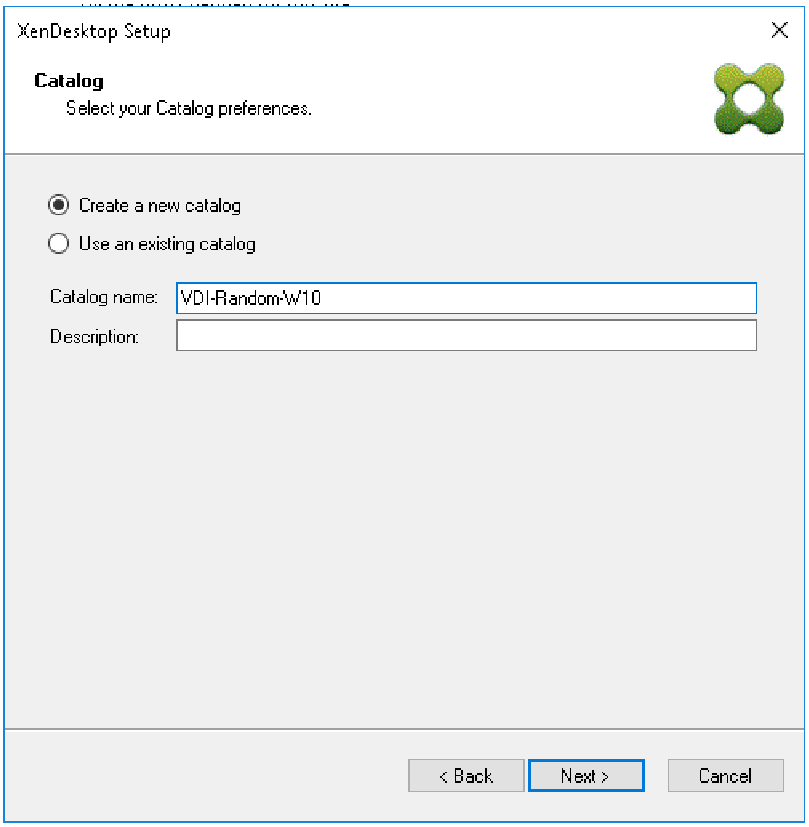

XenDesktop Setup

Catalog

Select your Catalog preferences.

● Create a new catalog

○ Use an existing catalog

Catalog name: VDI-Random-W10

Description:

24. On the Operating System dialog, specify the operating system for the catalog. Specify Windows Desktop Operating System for VDI and Windows Server Operating System for RDS.

25. Click Next.

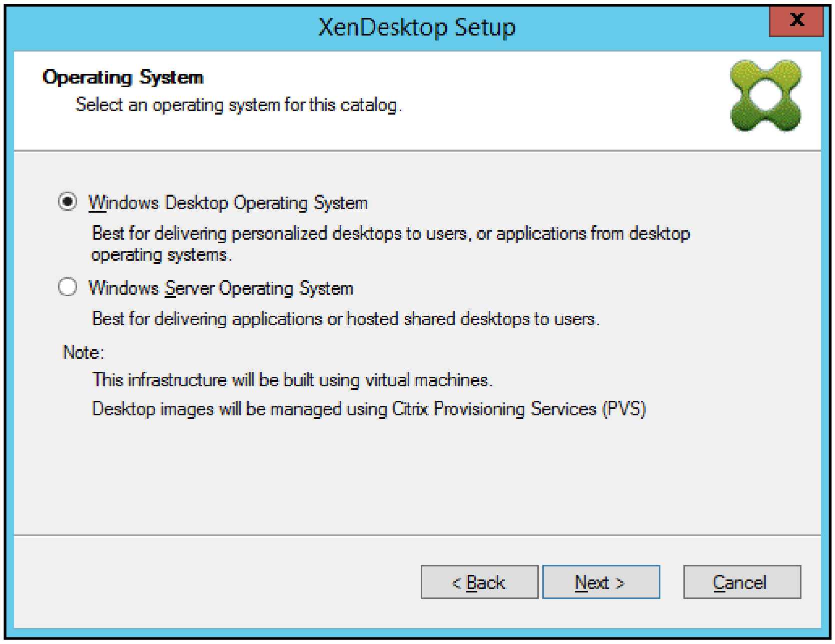

XenDesktop Setup

Operating System

Select an operating system for this catalog.

● Windows Desktop Operating System

Best for delivering personalized desktops to users, or applications from desktop operating systems.

○ Windows Server Operating System

Best for delivering applications or hosted shared desktops to users.

Note:

This infrastructure will be built using virtual machines.

Desktop images will be managed using Citrix Provisioning Services (PVS)

26. If you specified a Windows Desktop OS for VDIs, a **User Experience** dialog appears. Specify that the user will connect to “**A fresh new (random) desktop each time**.”

27. Click **Next.**

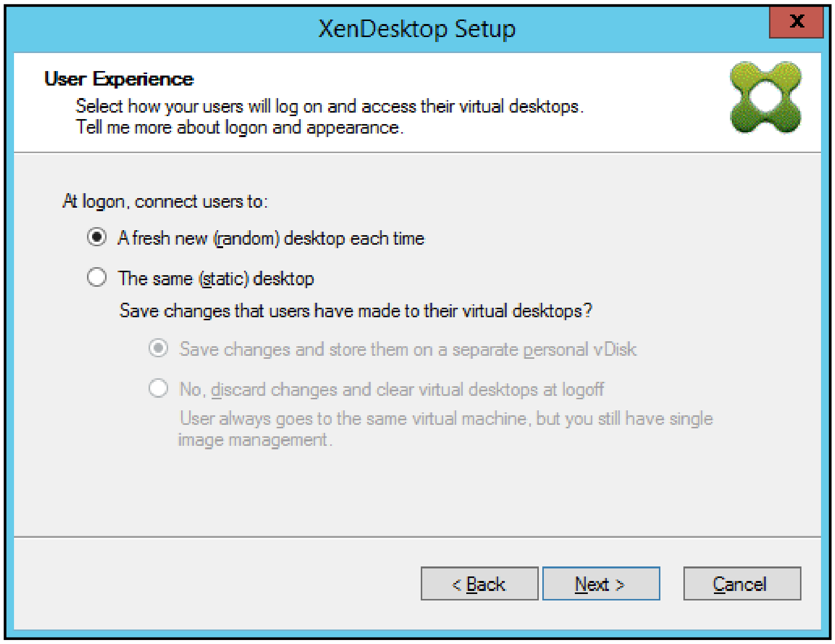

User Experience

Select how your users will log on and access their virtual desktops.
Tell me more about logon and appearance.

At logon, connect users to:

● A fresh new (random) desktop each time

○ The same (static) desktop

Save changes that users have made to their virtual desktops?

● Save changes and store them on a separate personal vDisk

○ No, discard changes and clear virtual desktops at logoff

User always goes to the same virtual machine, but you still have single image management.

28. On the Virtual machines dialog, specify:

   – The number of VMs to create. (Note that it is recommended to create 200 or less per provisioning run. Create a single VM at first to verify the procedure.)

   – Number of vCPUs for the VM (2 for HVD, 9 for HSD)

   – The amount of memory for the VM (2GB for HVD, 24GB for HSD)

   – The write-cache disk size (6GB for HVD, 30GB for HSD)

   – PXE boot as the Boot Mode

29. Click Next.

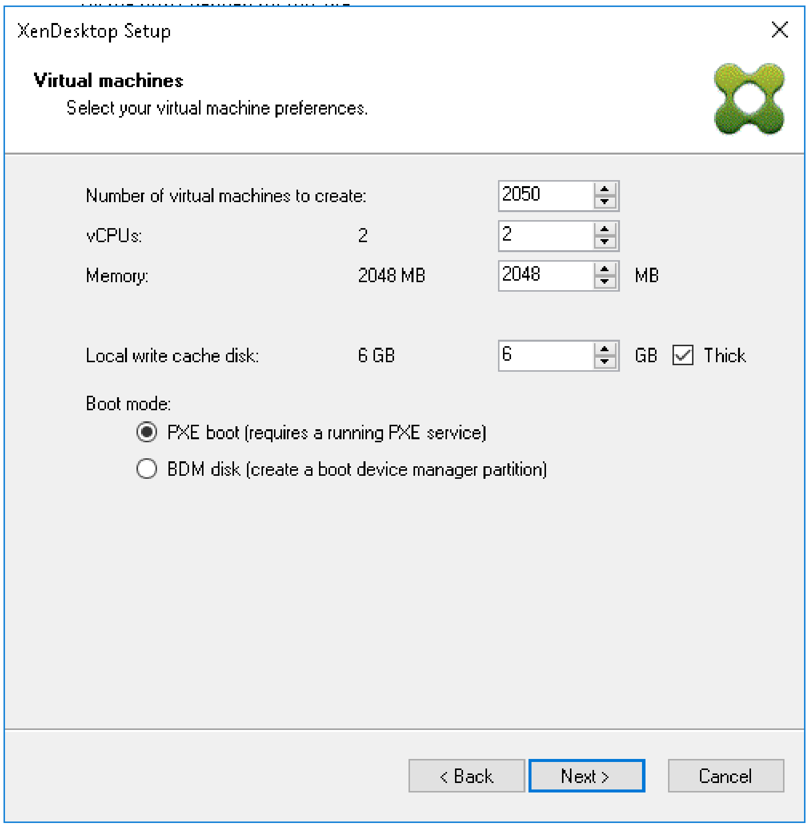
XenDesktop Setup

Virtual machines

Select your virtual machine preferences.

Number of virtual machines to create:		2050	
vCPUs:	2	2	
Memory:	2048 MB	2048	MB
Local write cache disk:	6 GB	6	GB ☑ Thick

Boot mode:

● PXE boot (requires a running PXE service)

○ BDM disk (create a boot device manager partition)

30. Select the **Create new accounts** radio button.

31. Click **Next**.

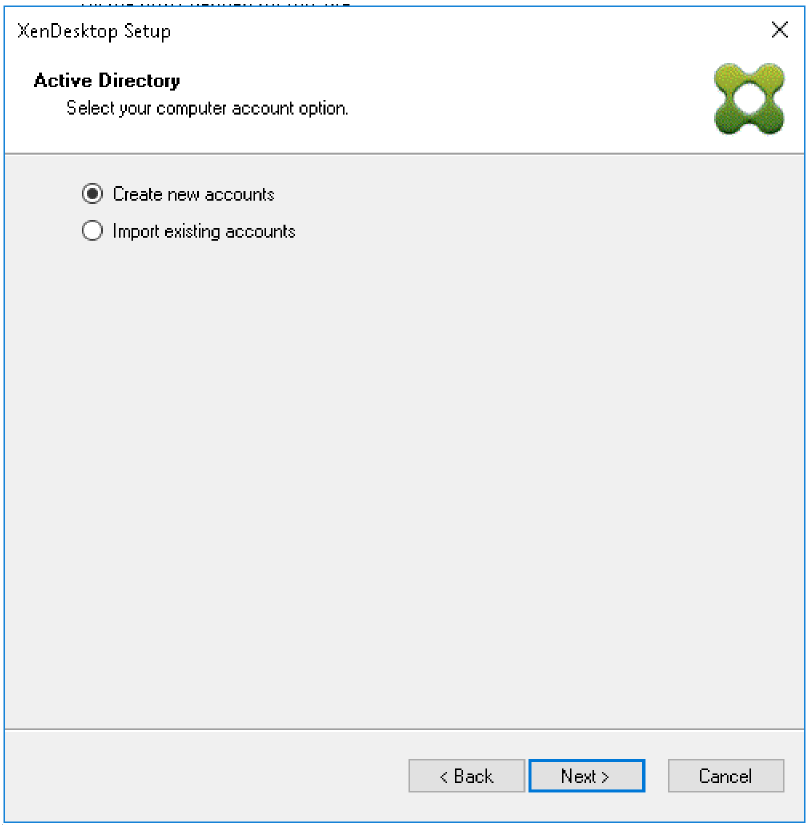

XenDesktop Setup

Active Directory

Select your computer account option.

● Create new accounts

○ Import existing accounts

32. Specify the Active Directory Accounts and Location. This is where the wizard should create the computer accounts.

33. Provide the Account naming scheme. An example name is shown in the text box below the name scheme selection location.

34. Click Next.

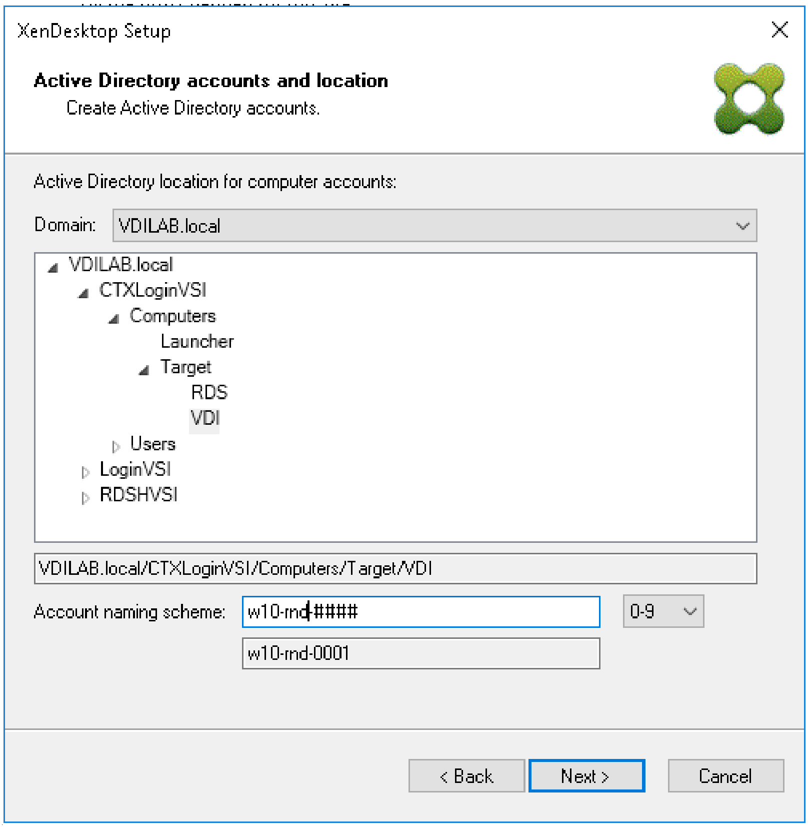

XenDesktop Setup

Active Directory accounts and location

Create Active Directory accounts.

Active Directory location for computer accounts:

Domain: VDILAB.local

```
◢ VDILAB.local
    ◢ CTXLoginVSI
        ◢ Computers
              Launcher
            ◢ Target
                  RDS
                  VDI
        ▷ Users
    ▷ LoginVSI
    ▷ RDSHVSI
```

VDILAB.local/CTXLoginVSI/Computers/Target/VDI

Account naming scheme: w10-rnd-#### | 0-9

w10-rnd-0001

35. Click Finish to begin the virtual machine creation.

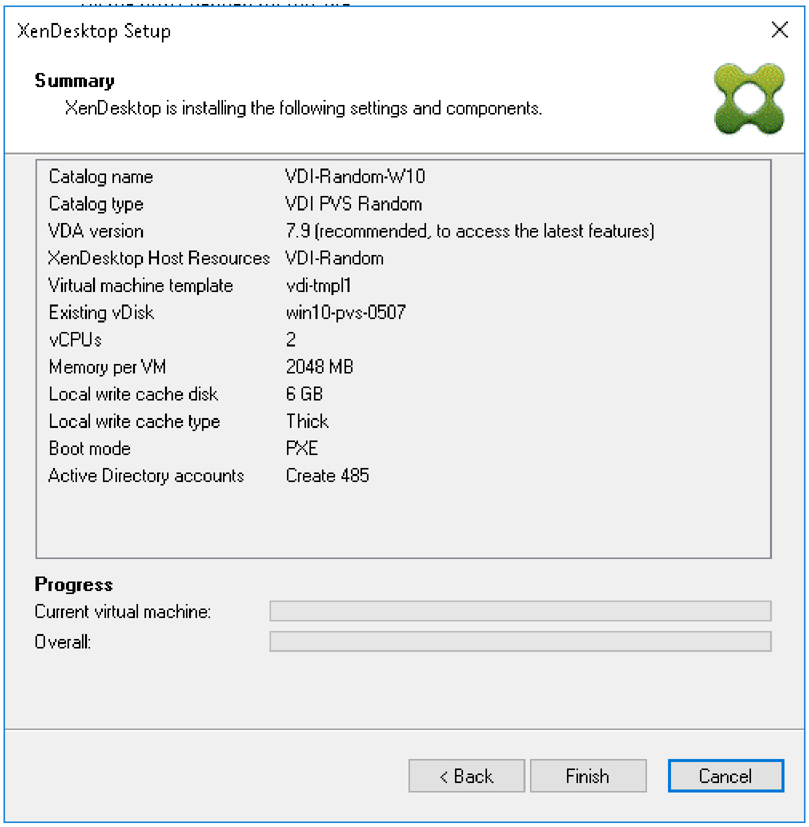

XenDesktop Setup

Summary

XenDesktop is installing the following settings and components.

Catalog name	VDI-Random-W10
Catalog type	VDI PVS Random
VDA version	7.9 (recommended, to access the latest features)
XenDesktop Host Resources	VDI-Random
Virtual machine template	vdi-tmpl1
Existing vDisk	win10-pvs-0507
vCPUs	2
Memory per VM	2048 MB
Local write cache disk	6 GB
Local write cache type	Thick
Boot mode	PXE
Active Directory accounts	Create 485

Progress

Current virtual machine:

Overall:

36. When the wizard is done provisioning the virtual machines, click Done.

Citrix Machine Creation Services

1. Connect to a XenDesktop server and launch Citrix Studio.
2. Choose Create Machine Catalog from the Actions pane.
3. Click Next.

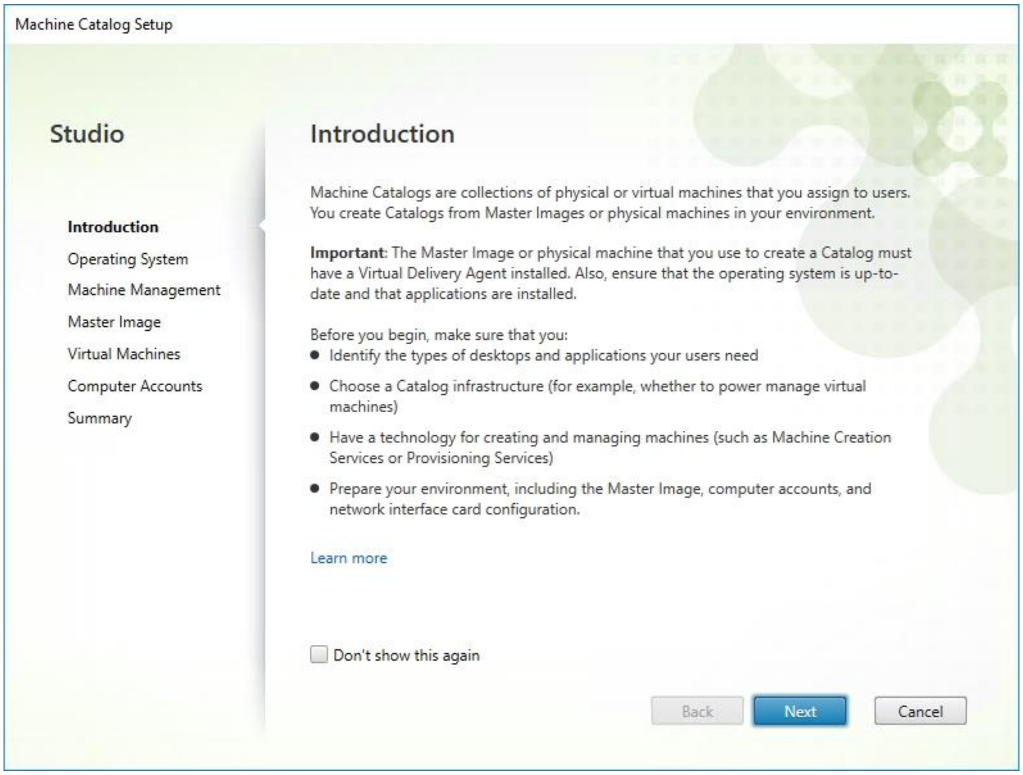

Machine Catalog Setup

Introduction

Machine Catalogs are collections of physical or virtual machines that you assign to users. You create Catalogs from Master Images or physical machines in your environment.

Important: The Master Image or physical machine that you use to create a Catalog must have a Virtual Delivery Agent installed. Also, ensure that the operating system is up-to-date and that applications are installed.

Before you begin, make sure that you:

- Identify the types of desktops and applications your users need
- Choose a Catalog infrastructure (for example, whether to power manage virtual machines)
- Have a technology for creating and managing machines (such as Machine Creation Services or Provisioning Services)
- Prepare your environment, including the Master Image, computer accounts, and network interface card configuration.

Learn more

□ Don't show this again

4. Select Desktop OS.
5. Click Next.

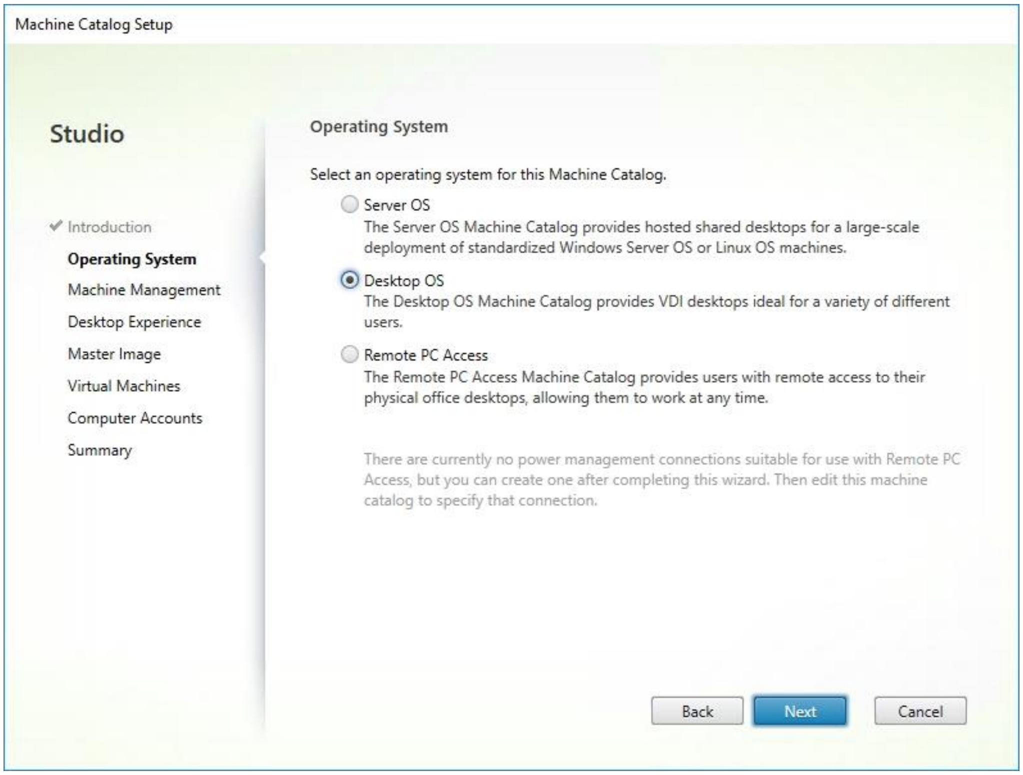

Machine Catalog Setup

Operating System

Select an operating system for this Machine Catalog.

○ Server OS
The Server OS Machine Catalog provides hosted shared desktops for a large-scale deployment of standardized Windows Server OS or Linux OS machines.

● Desktop OS
The Desktop OS Machine Catalog provides VDI desktops ideal for a variety of different users.

○ Remote PC Access
The Remote PC Access Machine Catalog provides users with remote access to their physical office desktops, allowing them to work at any time.

There are currently no power management connections suitable for use with Remote PC Access, but you can create one after completing this wizard. Then edit this machine catalog to specify that connection.

6. Select appropriate machine management.

7. Click Next.

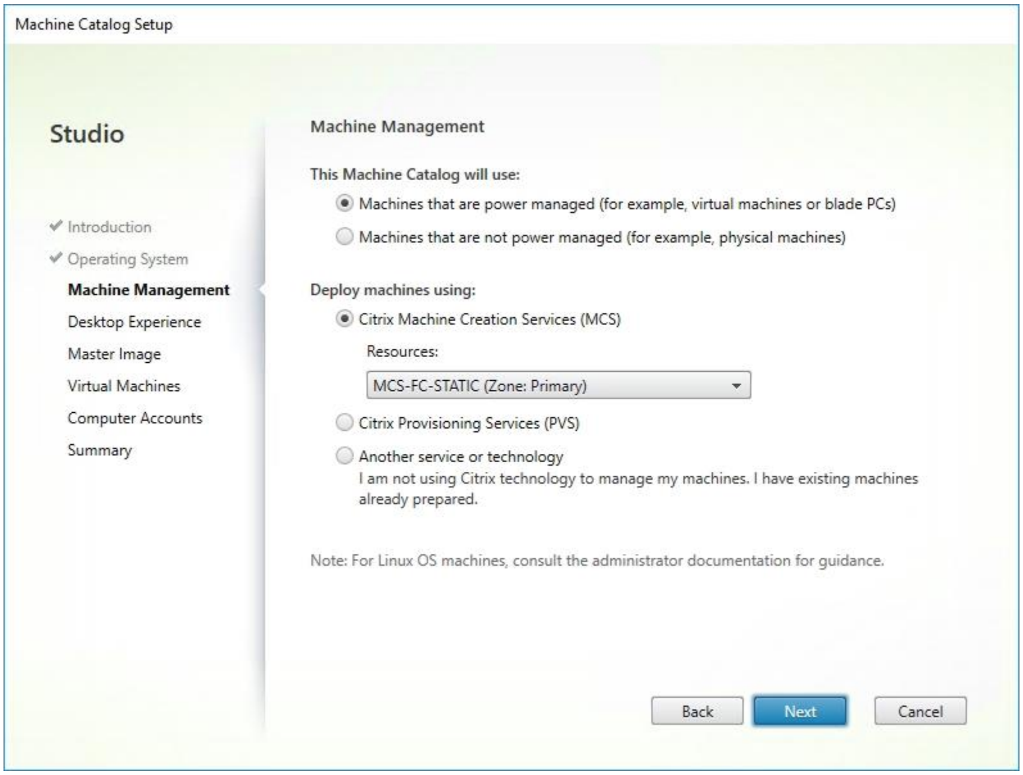

Machine Catalog Setup

Machine Management

This Machine Catalog will use:

● Machines that are power managed (for example, virtual machines or blade PCs)

○ Machines that are not power managed (for example, physical machines)

Deploy machines using:

● Citrix Machine Creation Services (MCS)

Resources: MCS-FC-STATIC (Zone: Primary)

○ Citrix Provisioning Services (PVS)

○ Another service or technology
I am not using Citrix technology to manage my machines. I have existing machines already prepared.

Note: For Linux OS machines, consult the administrator documentation for guidance.

8. Select Static, Dedicated Virtual Machine for Desktop Experience.

9. Click Next.

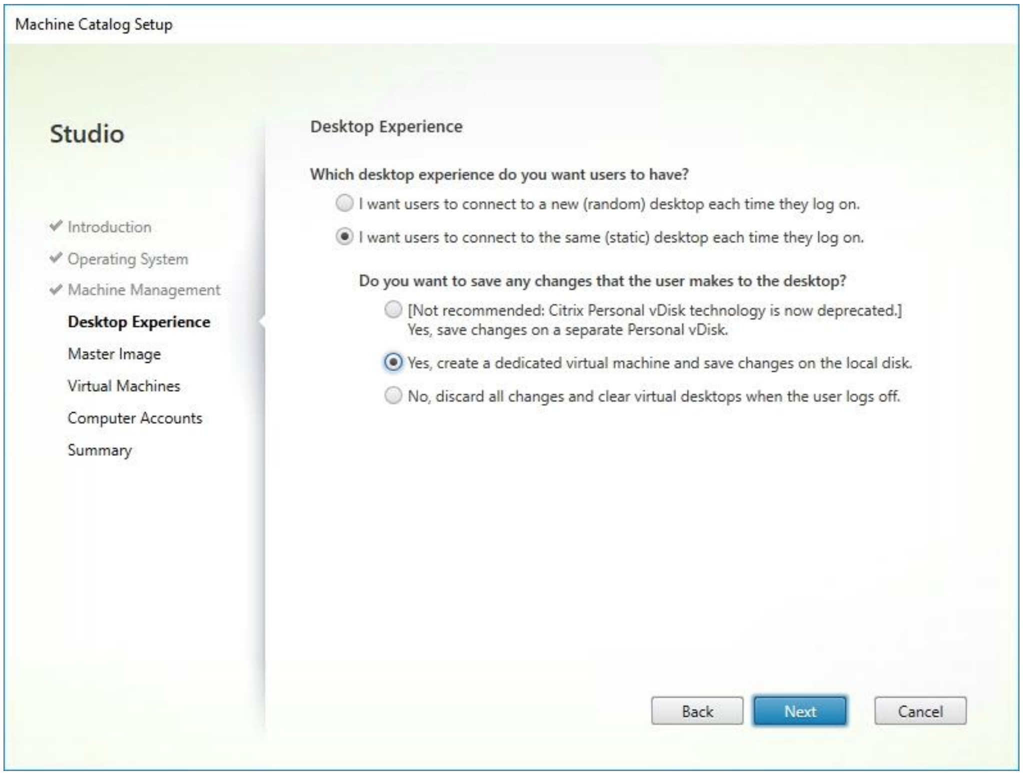

Machine Catalog Setup

Desktop Experience

Which desktop experience do you want users to have?

○ I want users to connect to a new (random) desktop each time they log on.

● I want users to connect to the same (static) desktop each time they log on.

Do you want to save any changes that the user makes to the desktop?

○ [Not recommended: Citrix Personal vDisk technology is now deprecated.] Yes, save changes on a separate Personal vDisk.

● Yes, create a dedicated virtual machine and save changes on the local disk.

○ No, discard all changes and clear virtual desktops when the user logs off.

10. Select a Virtual Machine to be used for Catalog Master Image.

11. Click Next.

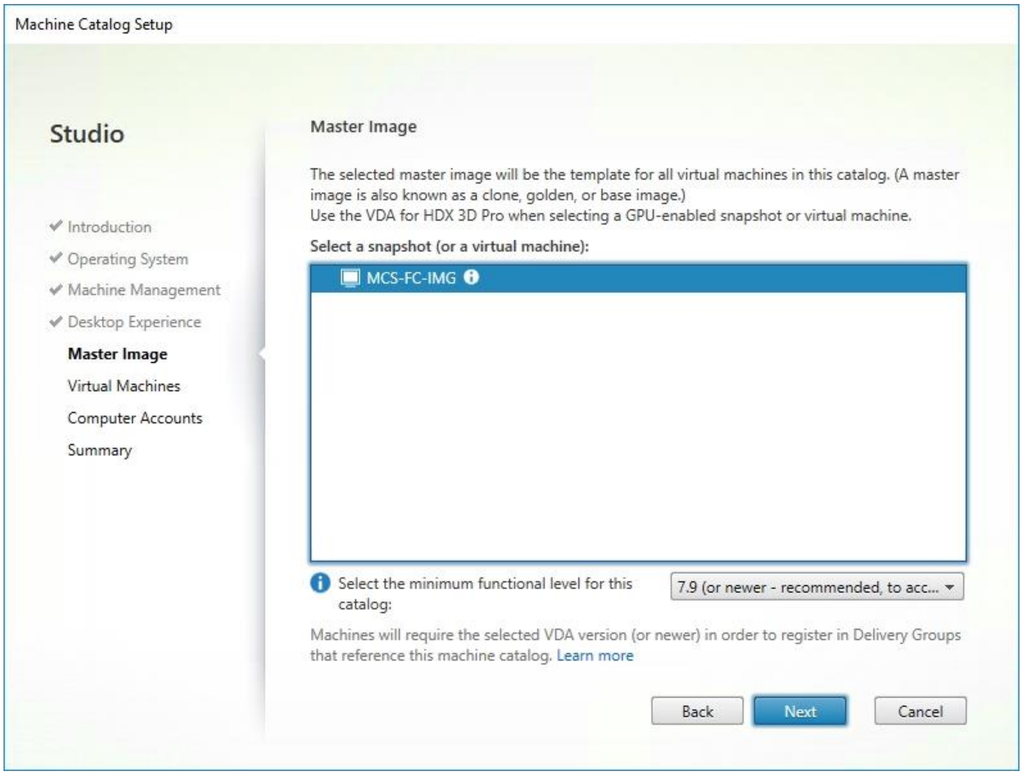

Machine Catalog Setup

Master Image

The selected master image will be the template for all virtual machines in this catalog. (A master image is also known as a clone, golden, or base image.)
Use the VDA for HDX 3D Pro when selecting a GPU-enabled snapshot or virtual machine.

Select a snapshot (or a virtual machine):

MCS-FC-IMG

Select the minimum functional level for this catalog: 7.9 (or newer - recommended, to acc...

Machines will require the selected VDA version (or newer) in order to register in Delivery Groups that reference this machine catalog. Learn more

12. Specify the number of desktops to create and machine configuration.

13. Set amount of memory (MB) to be used by virtual desktops.

14. Select Full Copy for machine copy mode.

15. Click Next.

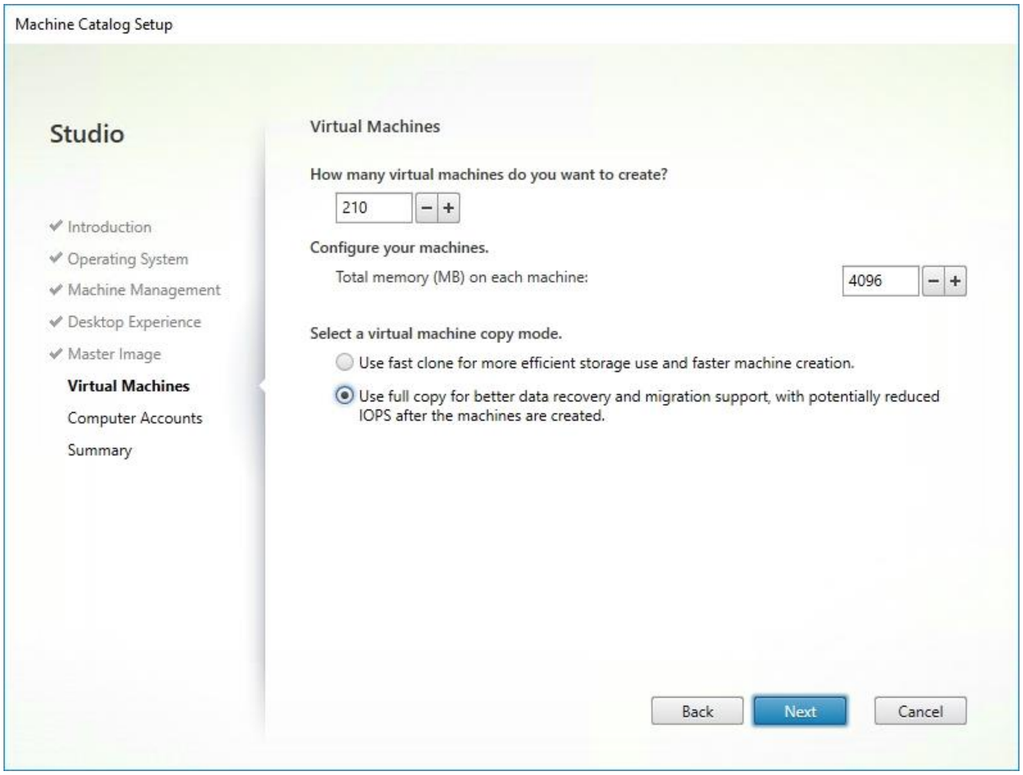

Machine Catalog Setup

Studio

Virtual Machines

How many virtual machines do you want to create?

210

Configure your machines.

Total memory (MB) on each machine: 4096

Select a virtual machine copy mode.

○ Use fast clone for more efficient storage use and faster machine creation.

● Use full copy for better data recovery and migration support, with potentially reduced IOPS after the machines are created.

16. Specify AD account naming scheme and OU where accounts will be created.

17. Click Next.

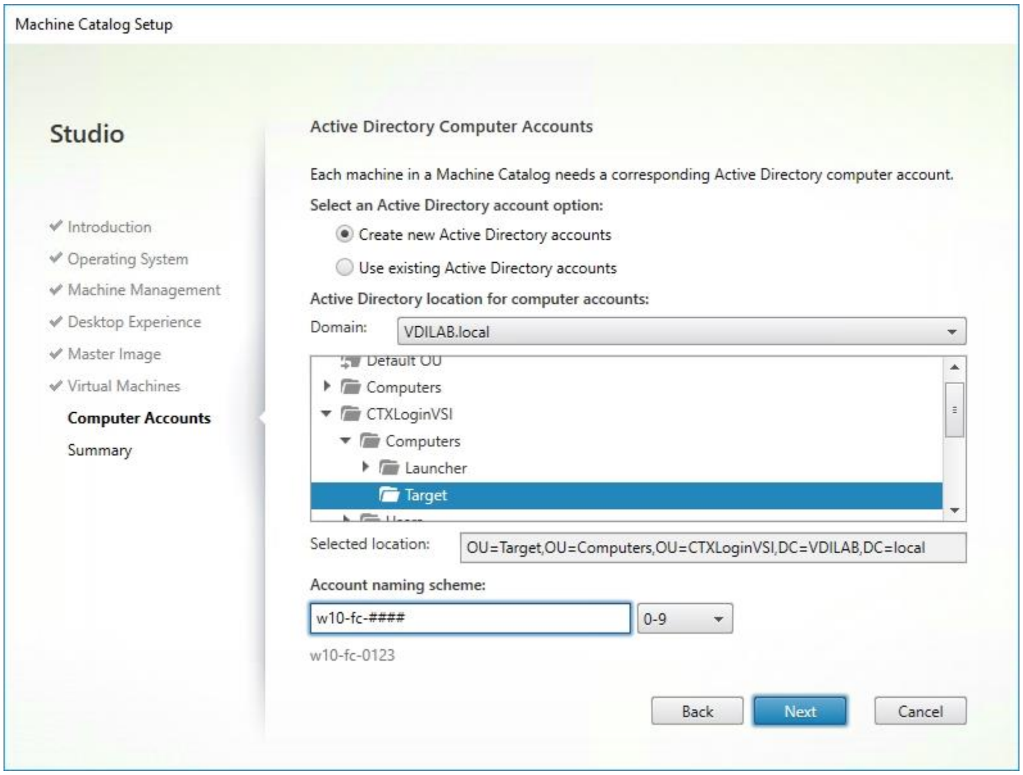

Machine Catalog Setup

Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Virtual Machines
- **Computer Accounts**
- Summary

Active Directory Computer Accounts

Each machine in a Machine Catalog needs a corresponding Active Directory computer account.

Select an Active Directory account option:

● Create new Active Directory accounts

○ Use existing Active Directory accounts

Active Directory location for computer accounts:

Domain: VDILAB.local

Default OU
Computers
CTXLoginVSI
Computers
Launcher
Target

Selected location: OU=Target,OU=Computers,OU=CTXLoginVSI,DC=VDILAB,DC=local

Account naming scheme:

w10-fc-#### 0-9

w10-fc-0123

18. On Summary page specify Catalog name and click Finish to start deployment.

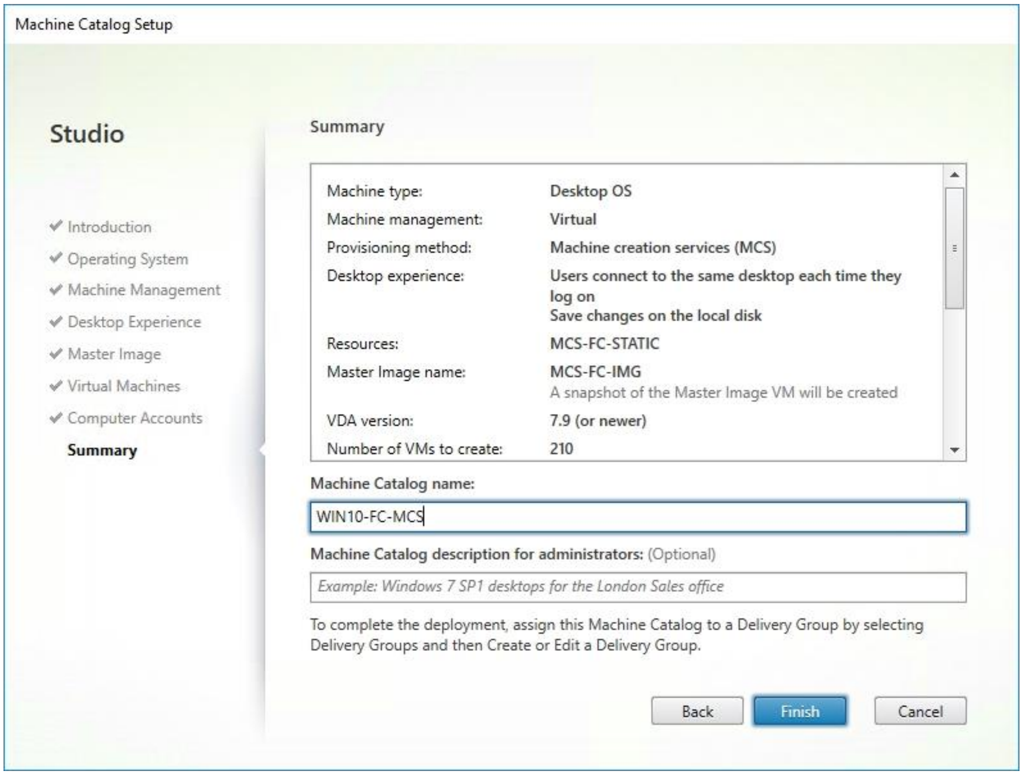

Machine Catalog Setup

Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Virtual Machines
- ✓ Computer Accounts
- **Summary**

Summary

Machine type:	Desktop OS
Machine management:	Virtual
Provisioning method:	Machine creation services (MCS)
Desktop experience:	Users connect to the same desktop each time they log on Save changes on the local disk
Resources:	MCS-FC-STATIC
Master Image name:	MCS-FC-IMG A snapshot of the Master Image VM will be created
VDA version:	7.9 (or newer)
Number of VMs to create:	210

Machine Catalog name:

WIN10-FC-MCS

Machine Catalog description for administrators: (Optional)

Example: Windows 7 SP1 desktops for the London Sales office

To complete the deployment, assign this Machine Catalog to a Delivery Group by selecting Delivery Groups and then Create or Edit a Delivery Group.

Create Delivery Groups

Delivery Groups are collections of machines that control access to desktops and applications. With Delivery Groups, you can specify which users and groups can access which desktops and applications.

To create delivery groups, complete the following steps:

The instructions below outline the procedure to create a Delivery Group for VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for RDS desktops.

1. Connect to a XenDesktop server and launch Citrix Studio.

2. Choose **Create Delivery Group** from the drop-down list.

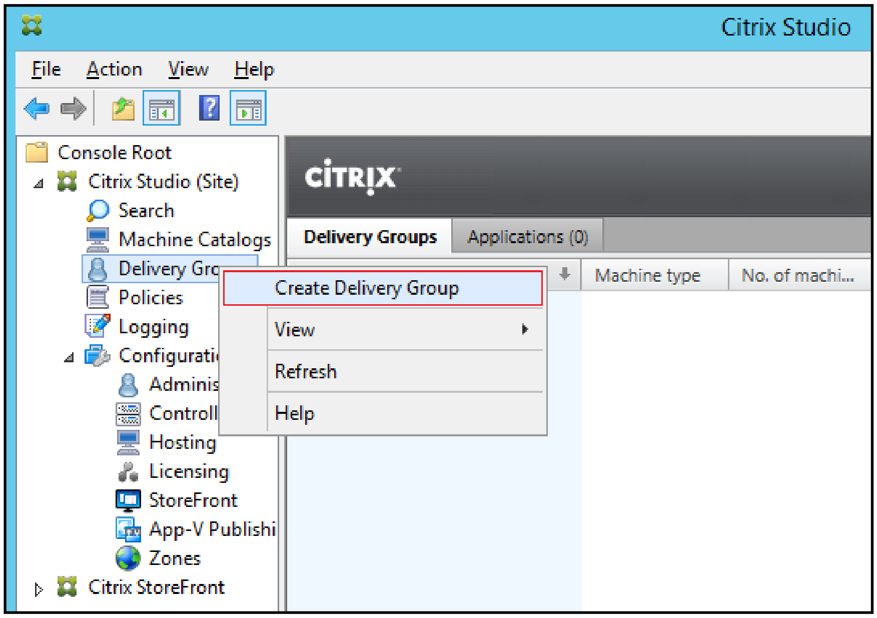

3. Click Next.

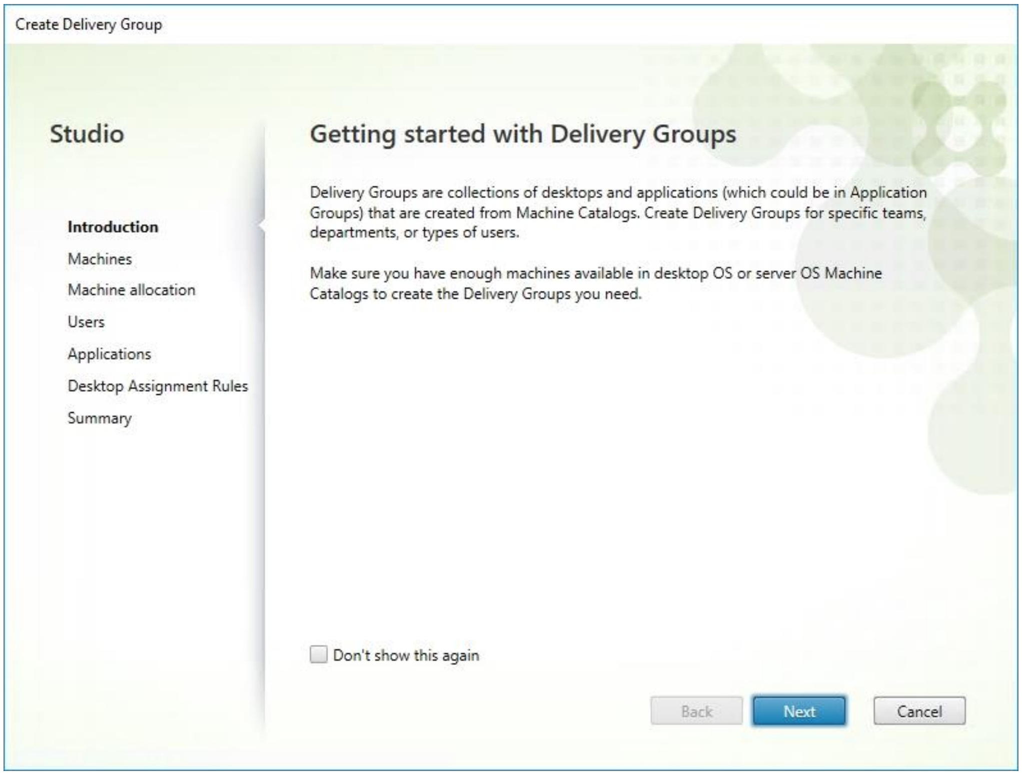

Create Delivery Group

Studio

Getting started with Delivery Groups

Delivery Groups are collections of desktops and applications (which could be in Application Groups) that are created from Machine Catalogs. Create Delivery Groups for specific teams, departments, or types of users.

Make sure you have enough machines available in desktop OS or server OS Machine Catalogs to create the Delivery Groups you need.

□ Don't show this again

4. Specify the Machine Catalog and increment the number of machines to add.

5. Click Next.

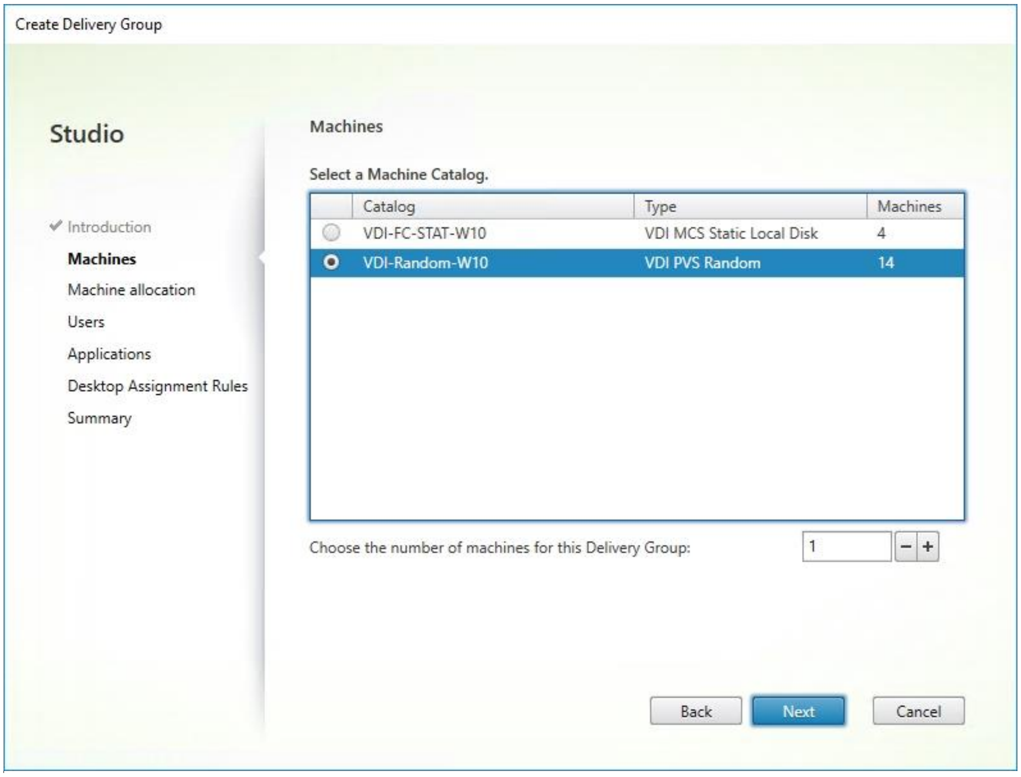

Create Delivery Group

Studio

Machines

Select a Machine Catalog.

	Catalog	Type	Machines
○	VDI-FC-STAT-W10	VDI MCS Static Local Disk	4
●	VDI-Random-W10	VDI PVS Random	14

Choose the number of machines for this Delivery Group: 1

6. Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.

7. Select Desktops.

8. Click Next.

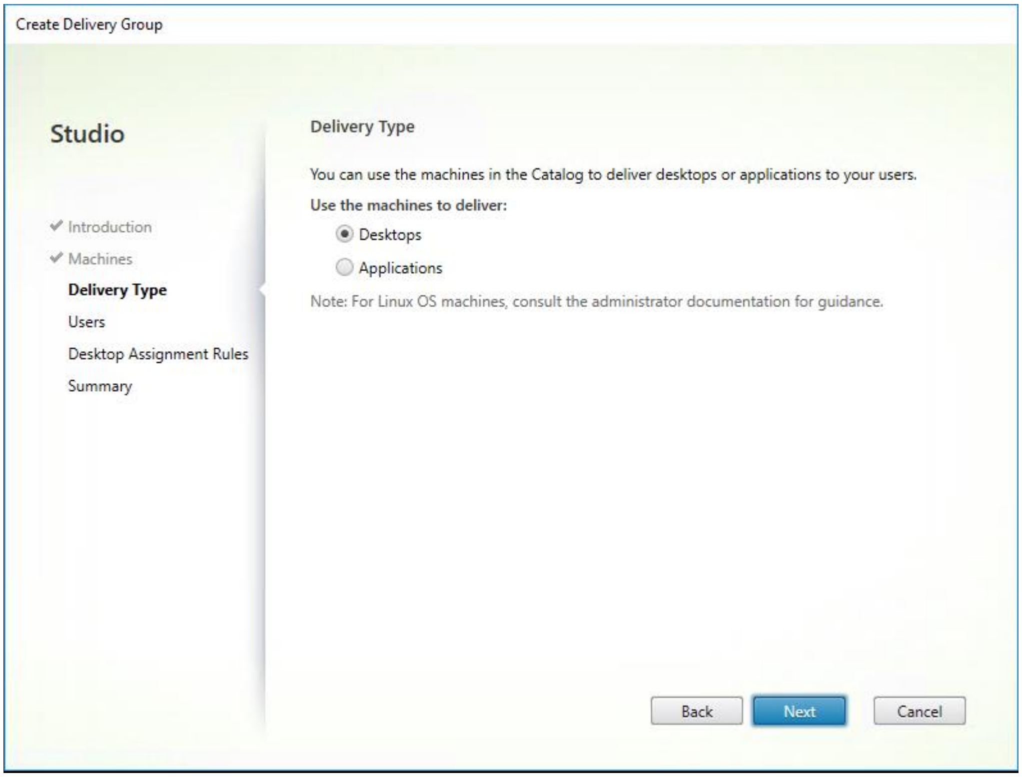

Create Delivery Group

Studio

Delivery Type

You can use the machines in the Catalog to deliver desktops or applications to your users.

Use the machines to deliver:

● Desktops

○ Applications

Note: For Linux OS machines, consult the administrator documentation for guidance.

9. To make the Delivery Group accessible, you must add users, select Allow any authenticated users to use this Delivery Group.

> User assignment can be updated any time after Delivery group creation by accessing Delivery group properties in Desktop Studio.

10. Click Next.

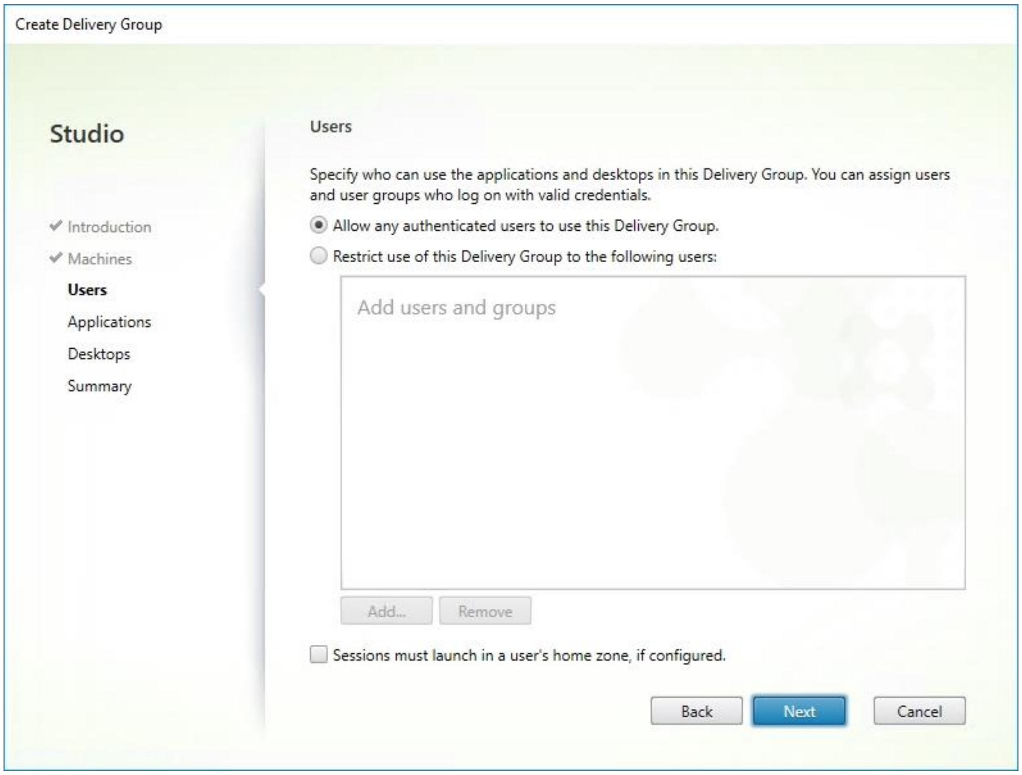

Create Delivery Group

Studio

Users

Specify who can use the applications and desktops in this Delivery Group. You can assign users and user groups who log on with valid credentials.

● Allow any authenticated users to use this Delivery Group.

○ Restrict use of this Delivery Group to the following users:

□ Sessions must launch in a user's home zone, if configured.

11. Click Next (no applications used in this design).

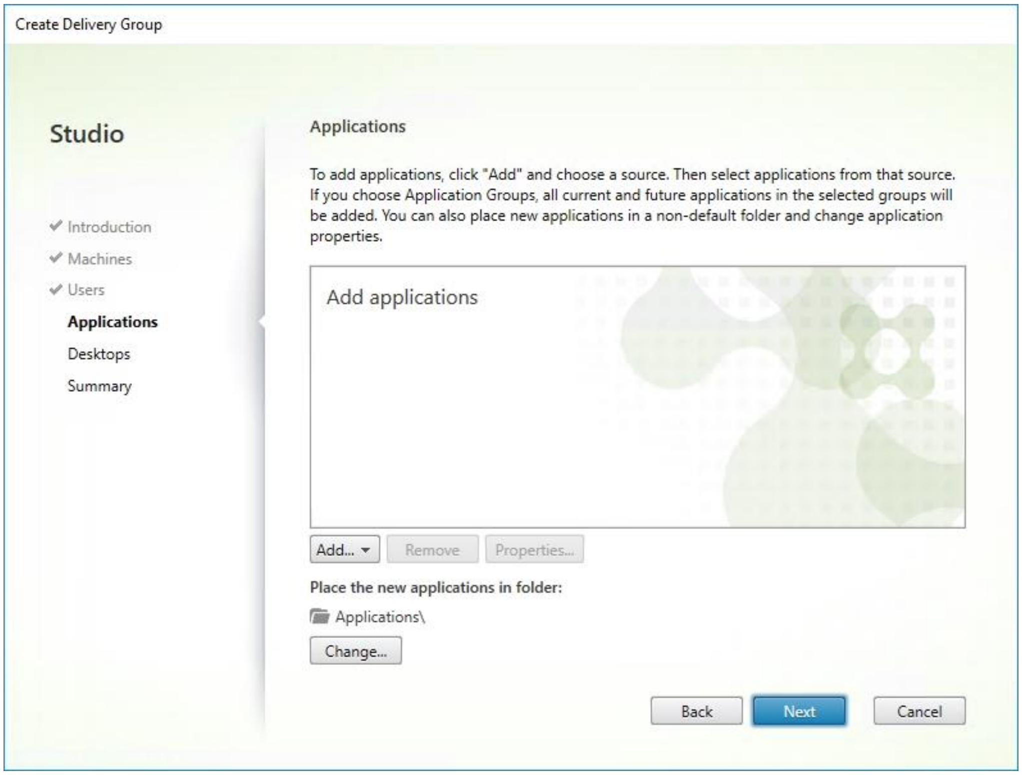

Create Delivery Group

Studio

Applications

To add applications, click "Add" and choose a source. Then select applications from that source. If you choose Application Groups, all current and future applications in the selected groups will be added. You can also place new applications in a non-default folder and change application properties.

Place the new applications in folder:

Applications\

12. Enable Users to access the desktops.

13. Click Next.

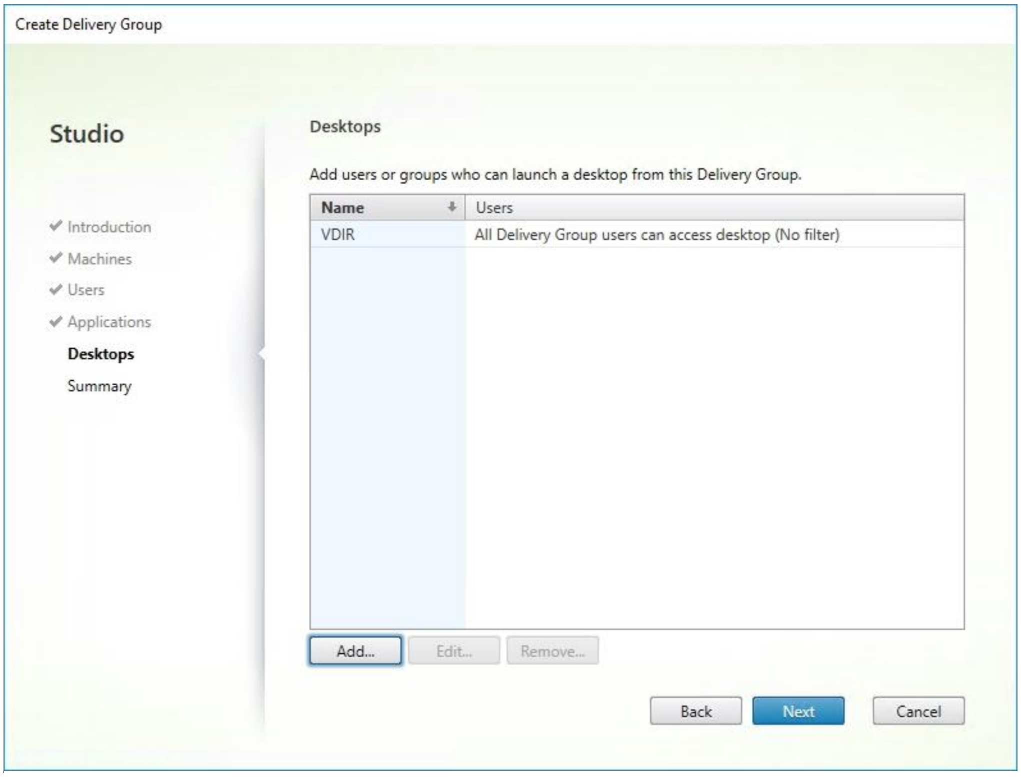

Desktops

Add users or groups who can launch a desktop from this Delivery Group.

Name	Users
VDIR	All Delivery Group users can access desktop (No filter)

14. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Description (Optional).

15. Click Finish.

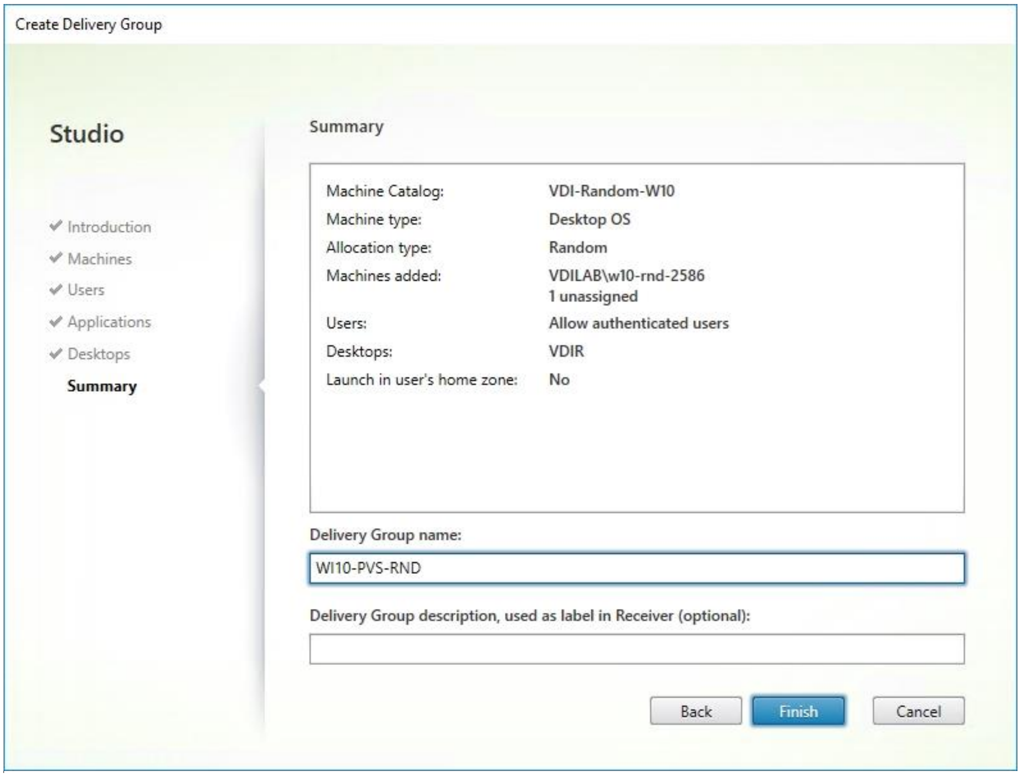

Summary

Machine Catalog:	VDI-Random-W10
Machine type:	Desktop OS
Allocation type:	Random
Machines added:	VDILAB\w10-rnd-2586 1 unassigned
Users:	Allow authenticated users
Desktops:	VDIR
Launch in user's home zone:	No

Delivery Group name: WI10-PVS-RND

Delivery Group description, used as label in Receiver (optional):

16. Citrix Studio lists the created Delivery Groups as well as the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab.

17. On the drop-down list, select "Turn on Maintenance Mode."

Citrix XenDesktop Policies and Profile Management

Policies and profiles allow the Citrix XenDesktop environment to be easily and efficiently customized.

Configure Citrix XenDesktop Policies

Citrix XenDesktop policies control user access and session environments, and are the most efficient method of controlling connection, security, and bandwidth settings. You can create policies for specific groups of users, devices, or connection types with each policy. Policies can contain multiple settings and are typically defined through Citrix Studio. (The Windows Group Policy Management Console can also be used if the network environment includes Microsoft Active Directory and permissions are set for managing Group Policy Objects). The screenshot below shows policies for Login VSI testing in this CVD.

Figure 36 XenDesktop Policy

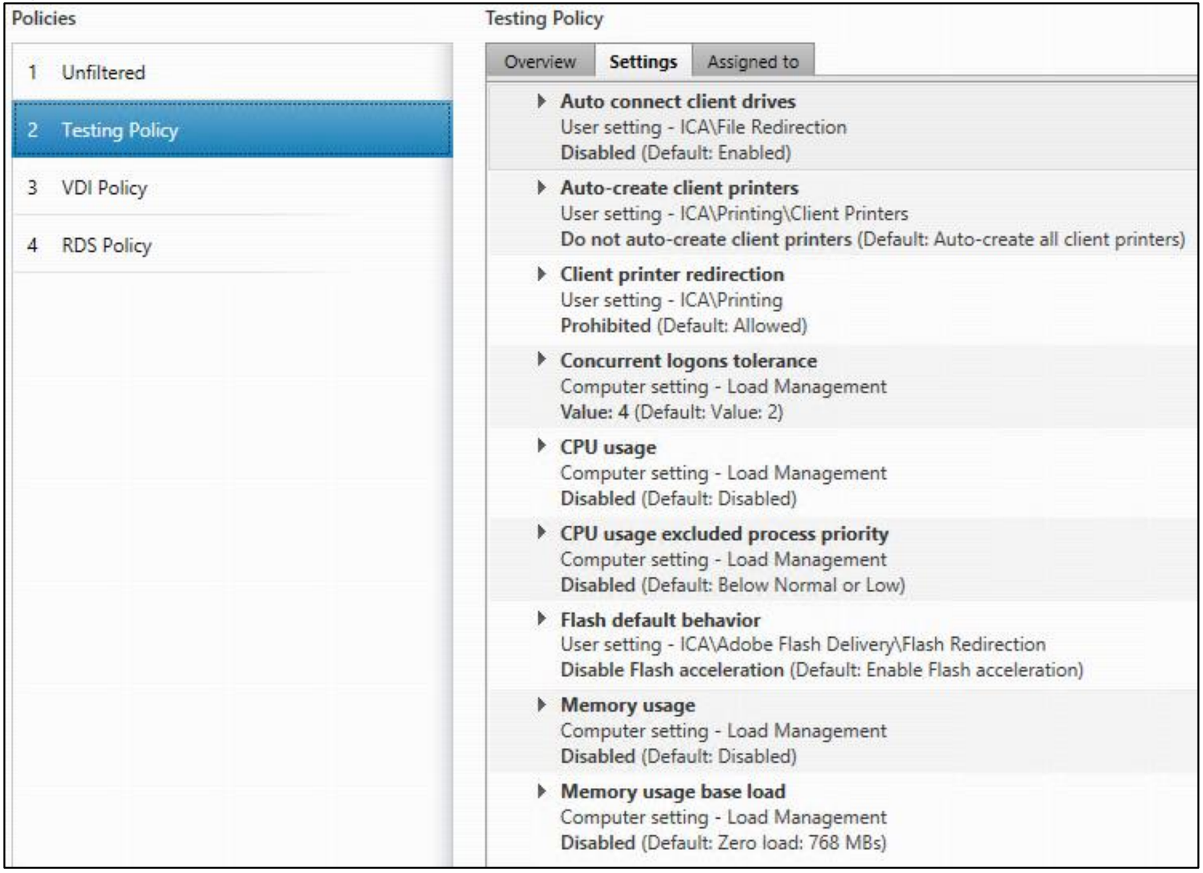

Auto connect client drives
User setting - ICA\File Redirection
Disabled (Default: Enabled)

Auto-create client printers
User setting - ICA\Printing\Client Printers
Do not auto-create client printers (Default: Auto-create all client printers)

Client printer redirection
User setting - ICA\Printing
Prohibited (Default: Allowed)

Concurrent logons tolerance
Computer setting - Load Management
Value: 4 (Default: Value: 2)

CPU usage
Computer setting - Load Management
Disabled (Default: Disabled)

CPU usage excluded process priority
Computer setting - Load Management
Disabled (Default: Below Normal or Low)

Flash default behavior
User setting - ICA\Adobe Flash Delivery\Flash Redirection
Disable Flash acceleration (Default: Enable Flash acceleration)

Memory usage
Computer setting - Load Management
Disabled (Default: Disabled)

Memory usage base load
Computer setting - Load Management
Disabled (Default: Zero load: 768 MBs)

Figure 37   Delivery Controllers Policy

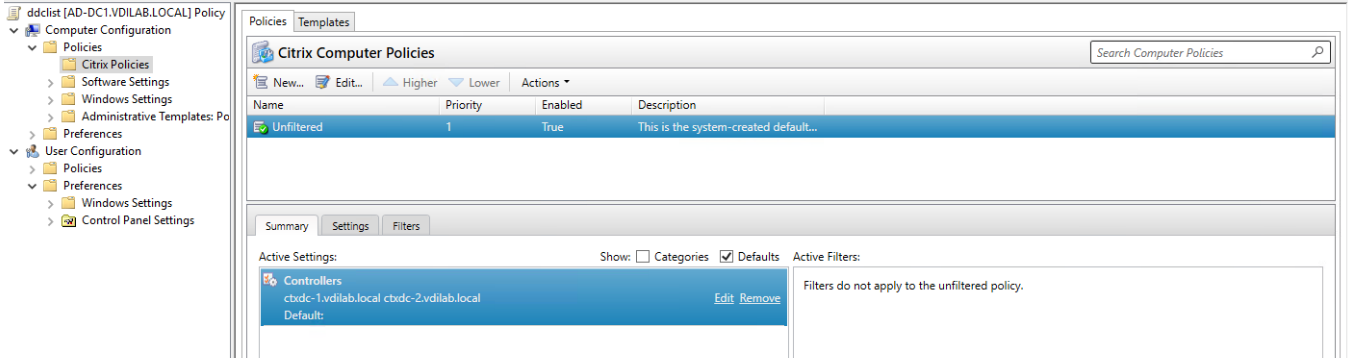

Name	Priority	Enabled	Description
Unfiltered	1	True	This is the system-created default...

Active Settings: Show: □ Categories ☑ Defaults

Controllers
ctxdc-1.vdilab.local ctxdc-2.vdilab.local   Edit Remove
Default:

Active Filters:

Filters do not apply to the unfiltered policy.

Configuring User Profile Management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting

- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for XenDesktop 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:

https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-15-ltsr.html

Figure 38 VDI User Profile Manager Policy

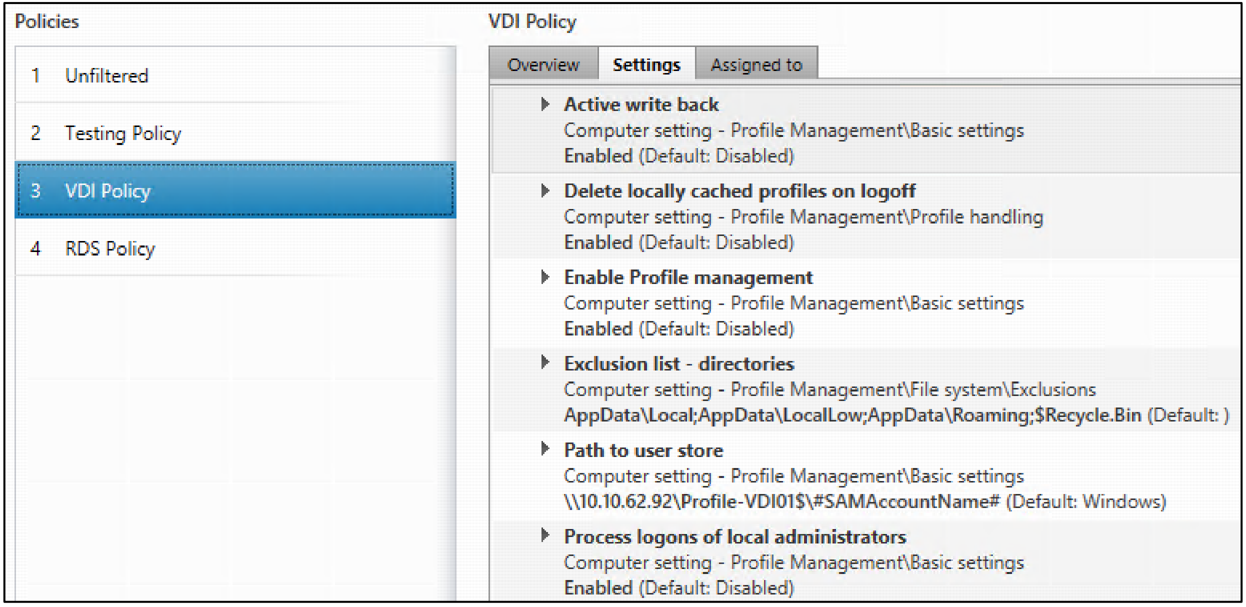

Active write back
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

Delete locally cached profiles on logoff
Computer setting - Profile Management\Profile handling
Enabled (Default: Disabled)

Enable Profile management
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

Exclusion list - directories
Computer setting - Profile Management\File system\Exclusions
AppData\Local;AppData\LocalLow;AppData\Roaming;$Recycle.Bin (Default:)

Path to user store
Computer setting - Profile Management\Basic settings
\\10.10.62.92\Profile-VDI01$\#SAMAccountName# (Default: Windows)

Process logons of local administrators
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

Figure 39 RDS User Profile Manager Policy

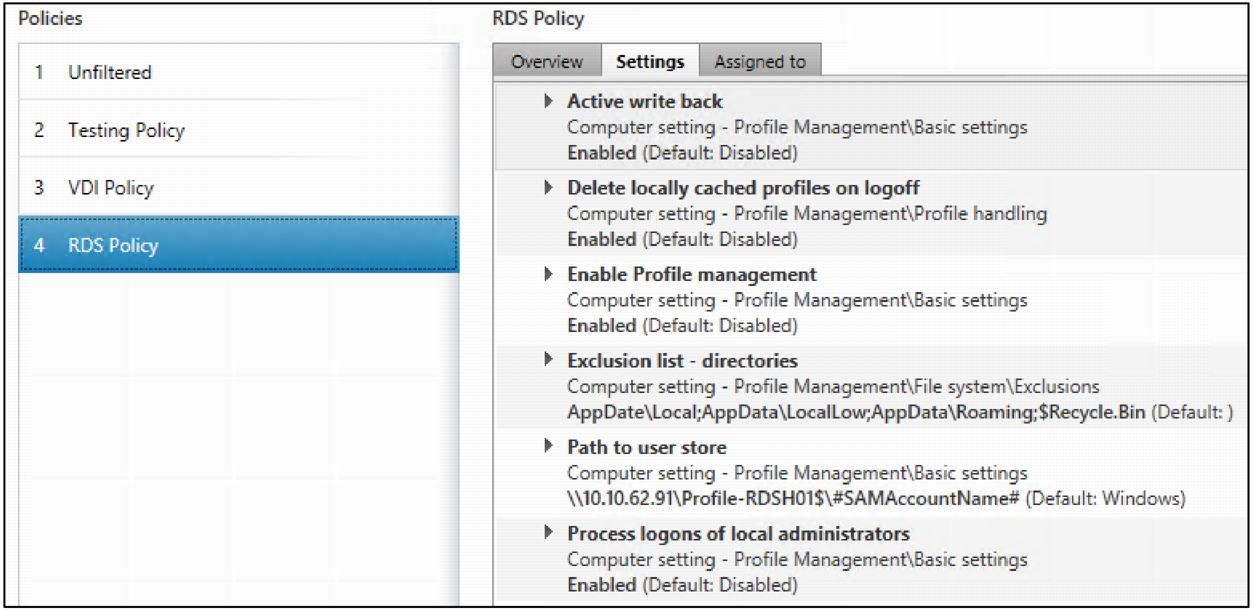

Active write back
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

Delete locally cached profiles on logoff
Computer setting - Profile Management\Profile handling
Enabled (Default: Disabled)

Enable Profile management
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

Exclusion list - directories
Computer setting - Profile Management\File system\Exclusions
AppDate\Local;AppData\LocalLow;AppData\Roaming;$Recycle.Bin (Default:)

Path to user store
Computer setting - Profile Management\Basic settings
\\10.10.62.91\Profile-RDSH01$\#SAMAccountName# (Default: Windows)

Process logons of local administrators
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

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 Install of P6 Card into 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
- UCSB-GPU-P6-R can be installed only in the rear mezzanine slot (slot 2) of the server.

No front mezzanine cards can be installed when the server has CPUs greater than 165 W.

The following figure shows the installed NVIDIA P6 GPU in the front and rear mezzanine slots.

Figure 40 NVIDIA GPU Installed in the Front and Rear Mezzanine Slots

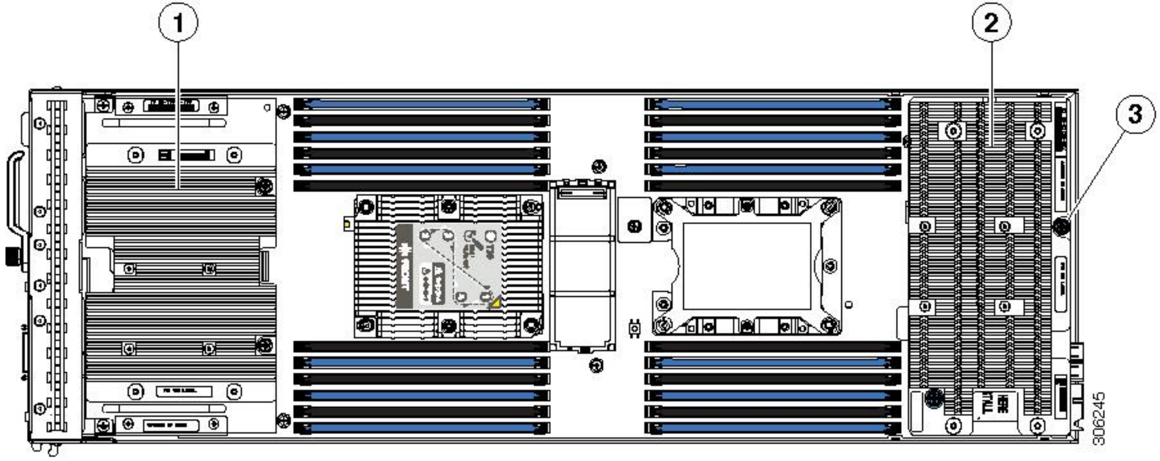

1	Front GPU	2	Rear GPU
3	Custom standoff screw	-	

Installing an NVIDIA GPU Card in the Front of the Server

The following figure shows the front NVIDIA P6 GPU (UCSB-GPU-P6-F).

Figure 41 NVIDIA P6 GPU That Installs in the Front of the Server

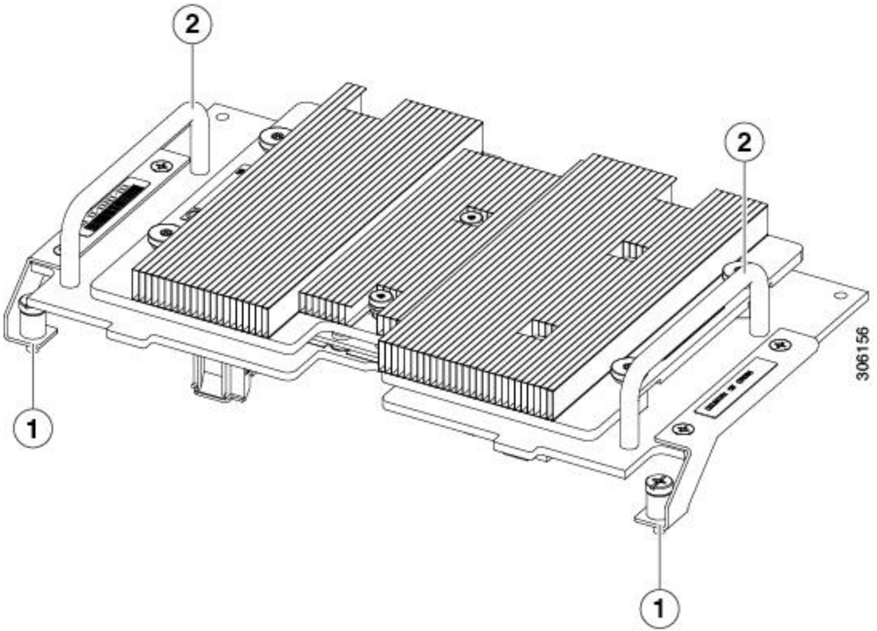

1	Leg with thumb screw that attaches to the server motherboard at the front	2	Handle to press down on when installing the GPU

Figure 42 Top Down View of the NVIDIA P6 GPU for the Front of the Server

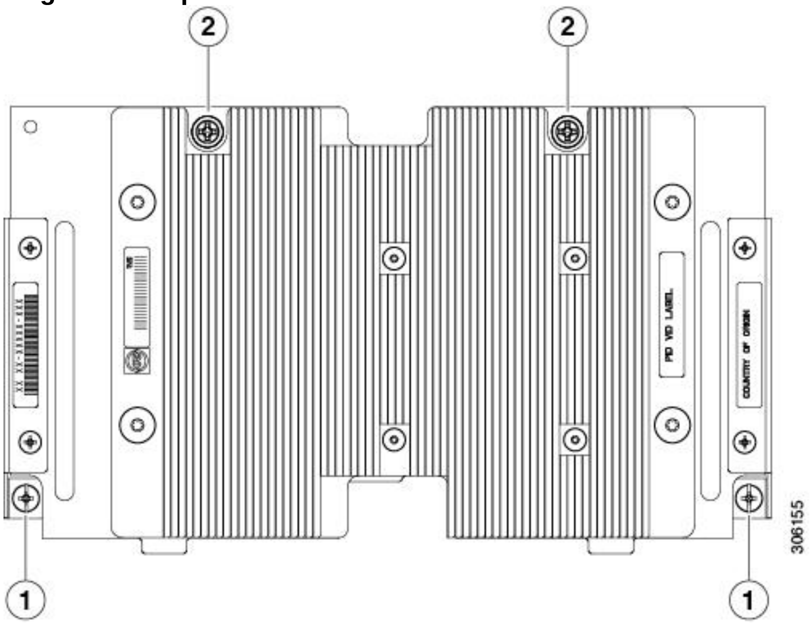

1	Leg with thumb screw that attaches to the server motherboard	2	Thumb screw that attaches to a standoff below

To install the NVIDIA GPU, follow these steps:

Before installing the NVIDIA P6 GPU (UCSB-GPU-P6-F) in the front mezzanine slot: 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 the following figure.

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 43 Installing the NVIDIA GPU in the Front of the Server

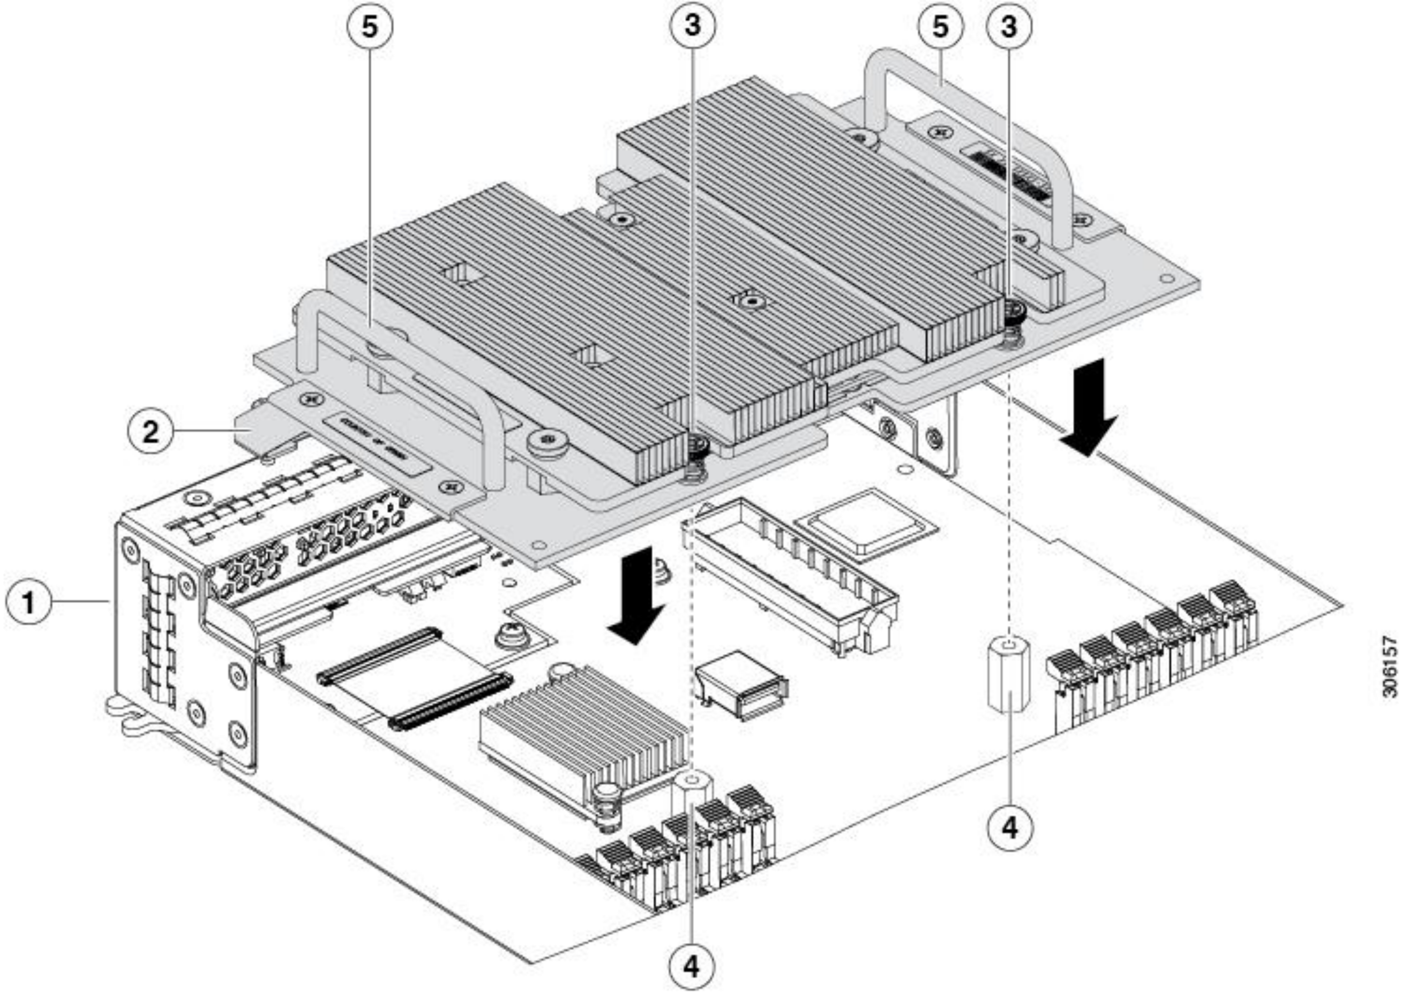

1	Front of the server	2	Leg with thumb screw that attaches to the motherboard
3	Thumbscrew to attach to standoff below	4	Standoff on the motherboard
5	Handle to press down on to firmly install the GPU	–	

Installing an NVIDIA GPU Card in the Rear of the Server

If you are installing the UCSB-GPU-P6-R to a server in the field, the option kit comes with the GPU itself (CPU and heatsink), a T-shaped installation wrench, and a custom standoff to support and attach the GPU to the motherboard. The following figure shows the three components of the option kit.

Figure 44 NVIDIA P6 GPU (UCSB-GPU-P6-R) Option Kit

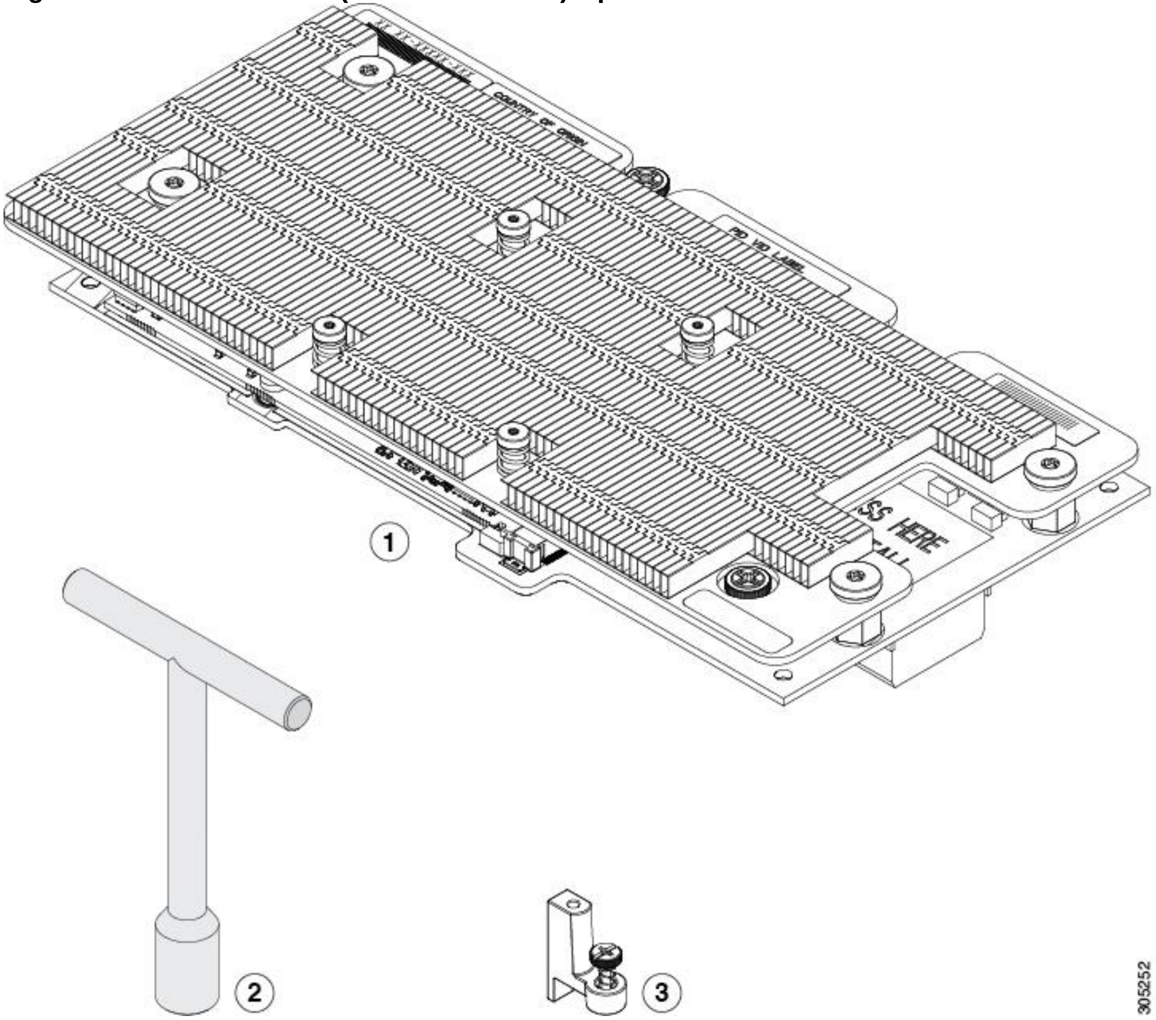

1	NVIDIA P6 GPU (CPU and heatsink)	2	T-shaped wrench
3	Custom standoff	-	

Before installing the NVIDIA P6 GPU (UCSB-GPU-P6-R) in the rear mezzanine slot: 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.

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 45 Installing the NVIDIA P6 GPU in the Rear Mezzanine Slot

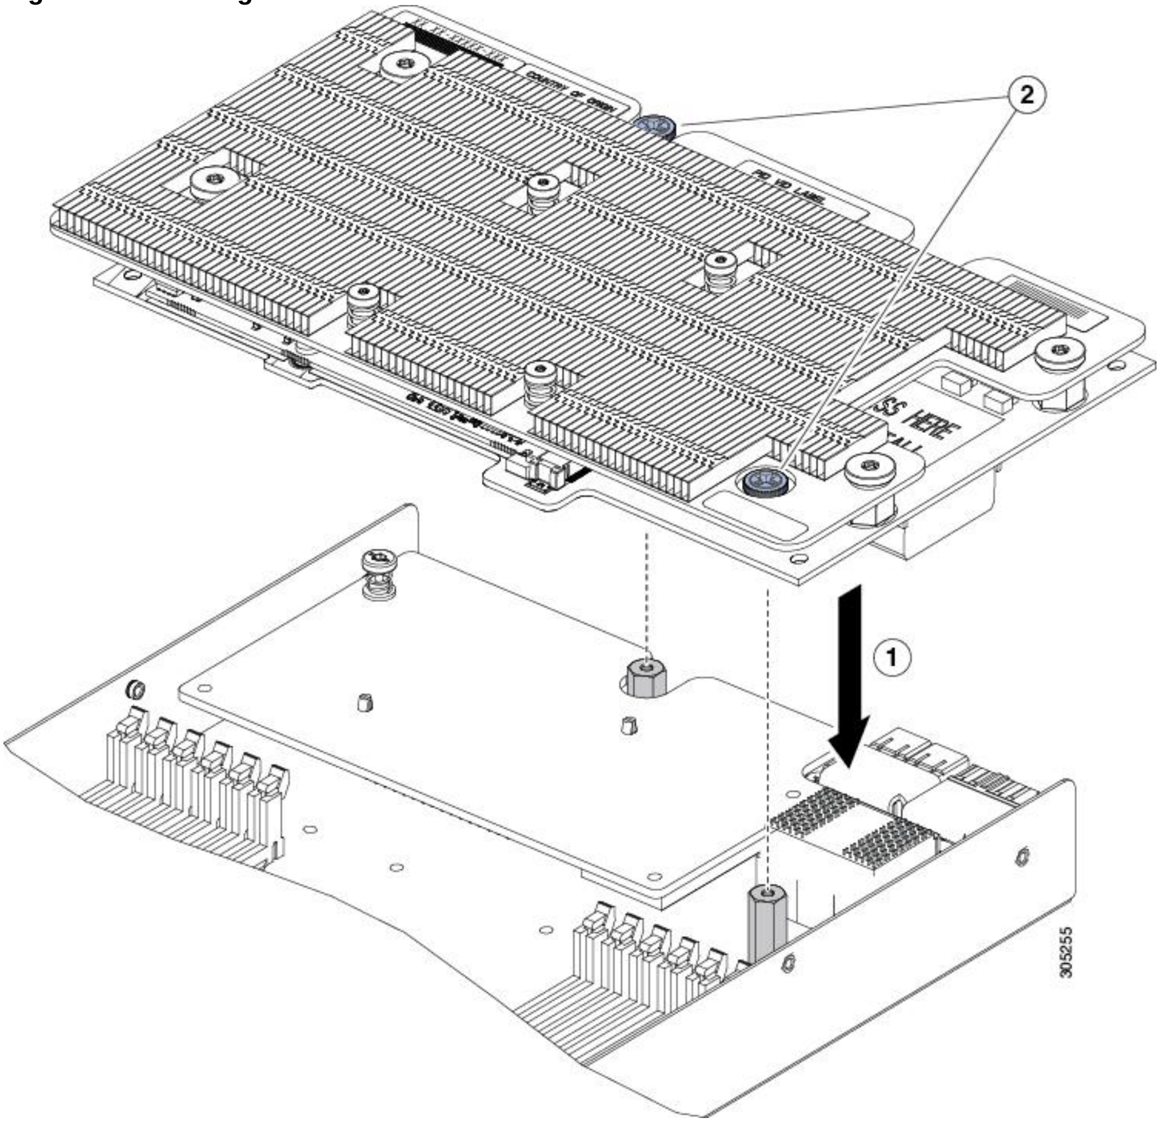

Install the NVIDIA VMware VIB Driver

To install the NVIDIA VMware VIB driver, complete the following steps:

1. From UCS Manager, verify the GPU card has been properly installed.

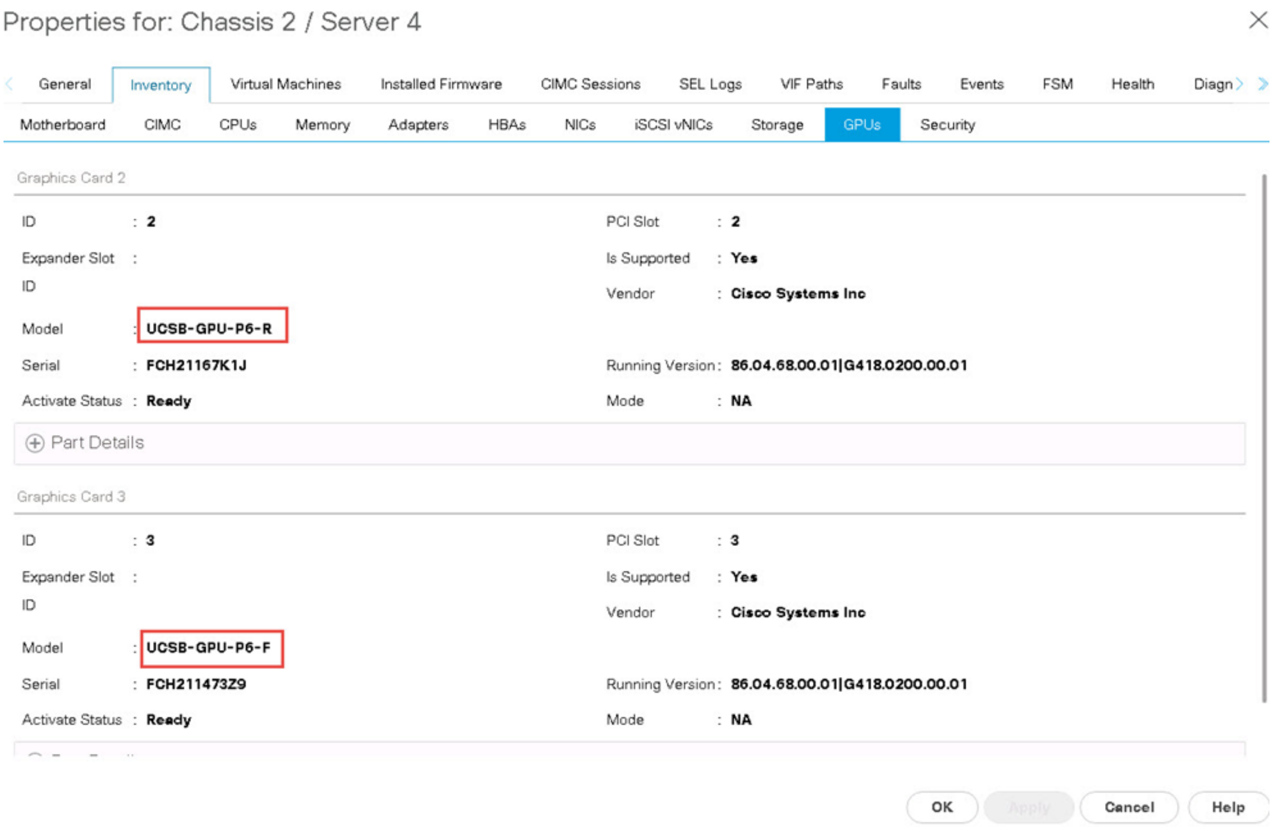

Graphics Card 2

ID : 2
PCI Slot : 2
Expander Slot ID :
Is Supported : Yes
Vendor : Cisco Systems Inc
Model : UCSB-GPU-P6-R
Serial : FCH21167K1J
Running Version : 86.04.68.00.01|G418.0200.00.01
Activate Status : Ready
Mode : NA

Graphics Card 3

ID : 3
PCI Slot : 3
Expander Slot ID :
Is Supported : Yes
Vendor : Cisco Systems Inc
Model : UCSB-GPU-P6-F
Serial : FCH211473Z9
Running Version : 86.04.68.00.01|G418.0200.00.01
Activate Status : Ready
Mode : NA

2. Download the NVIDIA GRID GPU driver pack for VMware vSphere ESXi 6.5.

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-VMware_ESXi_6.5_Host_Driver_384.99-
1OEM.650.0.0.4598673.vib
Installation Result
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: NVIDIA_bootbank_NVIDIA-VMware_ESXi_6.5_Host_Driver_384.99-1OEM.650.0.0.4598673
   VIBs Removed:
   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_6.5_Host_Driver  384.99-1OEM.650.0.0.4598673           NVIDIA
VMwareAccepted    2017-11-27
```

See the VMware knowledge base article for information about removing any existing NVIDIA drivers before installing new drivers: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2033434.

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 46, depending on the card used in your environment.

Figure 46 VMware ESX SSH Console Report for GPU P6 Card Detection on Cisco UCs B200 M5 Blade Server

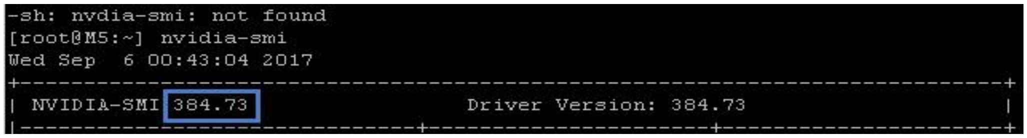

```
-sh: nvdia-smi: not found
[root@M5:~] nvidia-smi
Wed Sep  6 00:43:04 2017
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 384.73                 Driver Version: 384.73                    |
|-------------------------------+----------------------+----------------------+
```

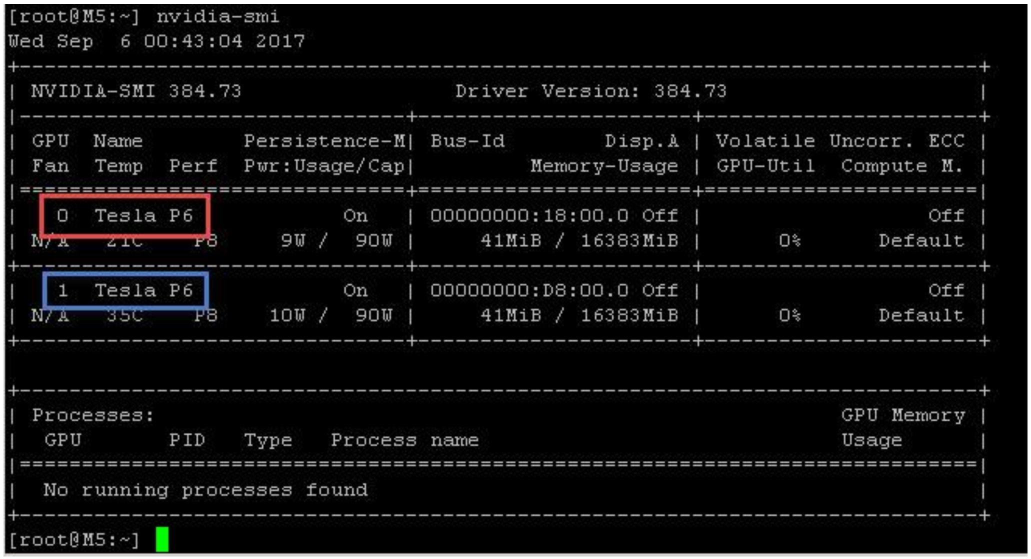

```
[root@M5:~] nvidia-smi
Wed Sep  6 00:43:04 2017
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 384.73                 Driver Version: 384.73                    |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P6            On   | 00000000:18:00.0 Off |                  Off |
| N/A   21C    P8     9W /  90W |     41MiB / 16383MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P6            On   | 00000000:D8:00.0 Off |                  Off |
| N/A   35C    P8    10W /  90W |     41MiB / 16383MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
[root@M5:~]
```

The NVIDIA system management interface (SMI) also allows GPU monitoring using the following command: `nvidia-smi –l` (this command adds a loop, automatically refreshing the display).

Configure a VM with a vGPU

To create the virtual machine that will later be used as the VDI base image, complete the following steps:

1. Select the ESXi host and click the Configure tab. From the list of options at the left, choose Graphics > Edit Host Graphics Settings. Select Shared Direct "Vendor shared passthrough graphics." Reboot the system to make the changes effective.

Figure 47 Edit Host Graphics Settings Window

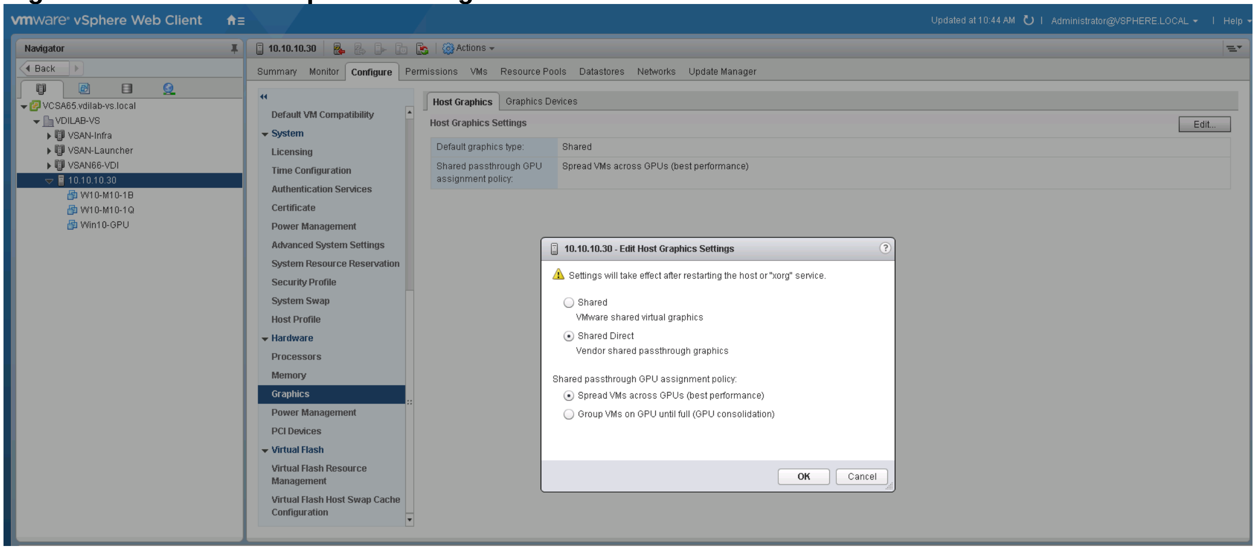

Default graphics type:	Shared
Shared passthrough GPU assignment policy:	Spread VMs across GPUs (best performance)

10.10.10.30 - Edit Host Graphics Settings

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 Web Client, create a new virtual machine. To do this, right-click a host or cluster and choose New Virtual Machine. Work through the New Virtual Machine wizard. Unless another configuration is specified, select the configuration settings appropriate for your environment.

Figure 48 Creating a New Virtual Machine in VMware vSphere Web Client

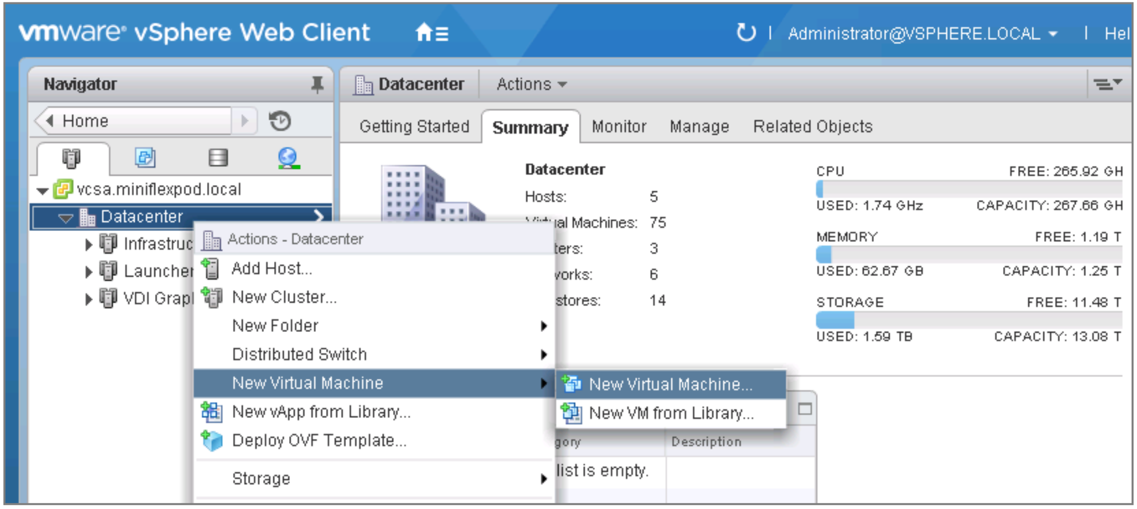

Datacenter
Hosts: 5
Virtual Machines: 75
ters: 3
vorks: 6
stores: 14

CPU FREE: 265.92 GH
USED: 1.74 GHz CAPACITY: 267.66 GH
MEMORY FREE: 1.19 T
USED: 62.67 GB CAPACITY: 1.25 T
STORAGE FREE: 11.48 T
USED: 1.59 TB CAPACITY: 13.08 T

3. Choose "ESXi 6.0 and later" from the "Compatible with" drop-down menu to use the latest features, including the mapping of shared PCI devices, which is required for the vGPU feature. This document uses "ESXi 6.5 and later," which provides the latest features available in ESXi 6.5 and virtual machine hardware Version 13.

Figure 49 Selecting Virtual Machine Hardware Version 11 or Later

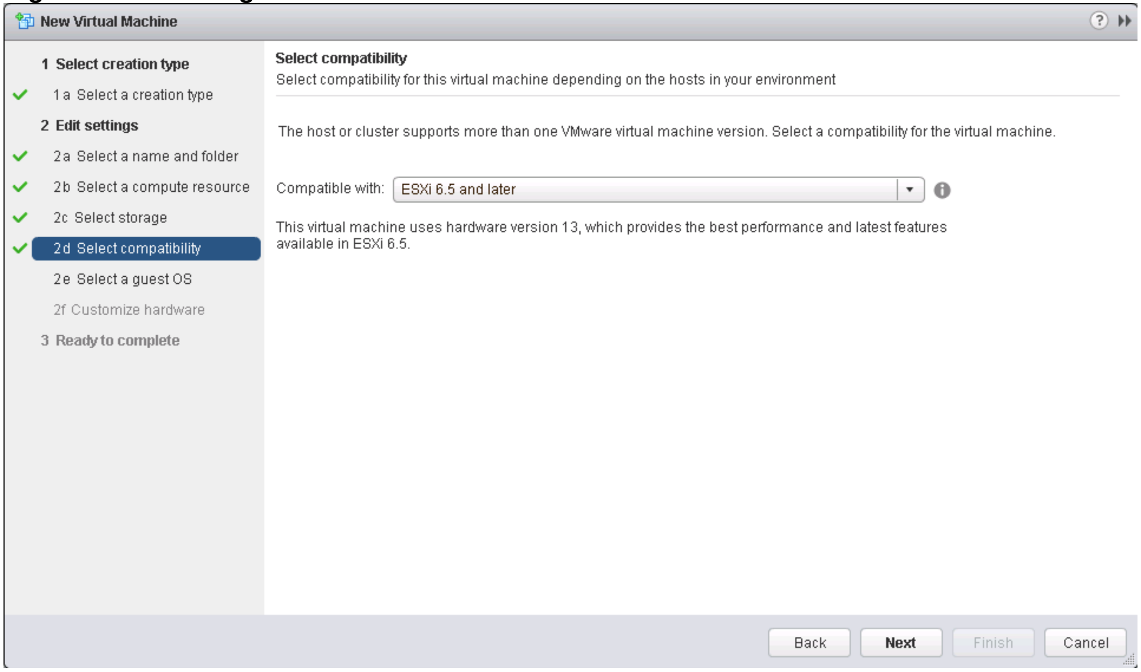

New Virtual Machine

1 Select creation type
✓ 1a Select a creation type
2 Edit settings
✓ 2a Select a name and folder
✓ 2b Select a compute resource
✓ 2c Select storage
✓ 2d Select compatibility
2e Select a guest OS
2f Customize hardware
3 Ready to complete

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 6.5 and later

This virtual machine uses hardware version 13, which provides the best performance and latest features available in ESXi 6.5.

Back | **Next** | Finish | Cancel

4. To customize the hardware of the new virtual machine, add a new shared 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 Web Client's virtual machine console functions, the mouse will not be usable in the virtual machine until after both the operating system and VMware Tools have been installed. If you cannot use the traditional vSphere Web Client to connect to the virtual machine, do not enable the NVIDIA GRID vGPU at this time.

Figure 50 Adding a Shared PCI Device to the Virtual Machine to Attach the GPU Profile

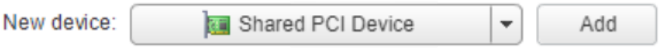

New device: Shared PCI Device | Add

Figure 51 Attaching the GPU Profile to a Shared PCI Device

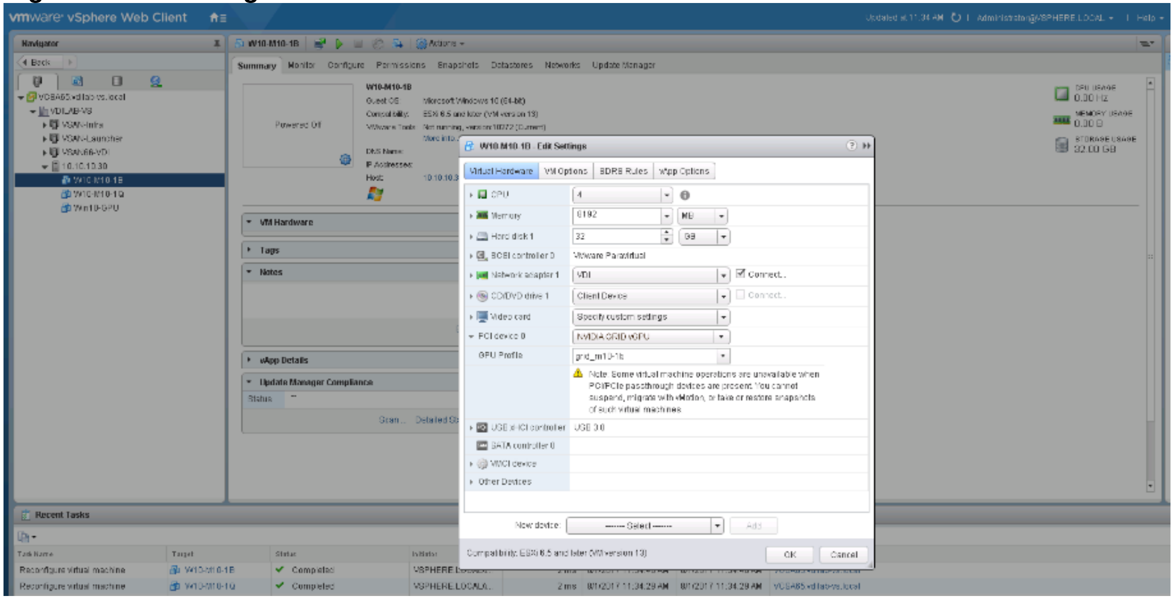

W10 M10 1B - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

CPU	4
Memory	8192 MB
Hard disk 1	32 GB
SCSI controller 0	VMware Paravirtual
Network adapter 1	VDI ☑ Connect...
CD/DVD drive 1	Client Device □ Connect...
Video card	Specify custom settings
PCI device 0	NVIDIA GRID vGPU
GPU Profile	grid_m10-1b
	Note: Some virtual machine operations are unavailable when PCI/PCIe passthrough devices are present. You cannot suspend, migrate with vMotion, or take or restore snapshots of such virtual machines.
USB xHCI controller	USB 3.0
SATA controller 0	
VMCI device	
Other Devices	

New device: ------- Select ------- | Add

Compatibility: ESXi 6.5 and later (VM version 13)

OK | Cancel

5. A virtual machine with a vGPU assigned will not start if ECC is enabled. If this is the case, as a workaround disable ECC by entering the following commands:

```
# nvidia-smi -i 0 -e 0

# nvidia-smi -i 1 -e 0
```

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.

Figure 52 Disabling ECC

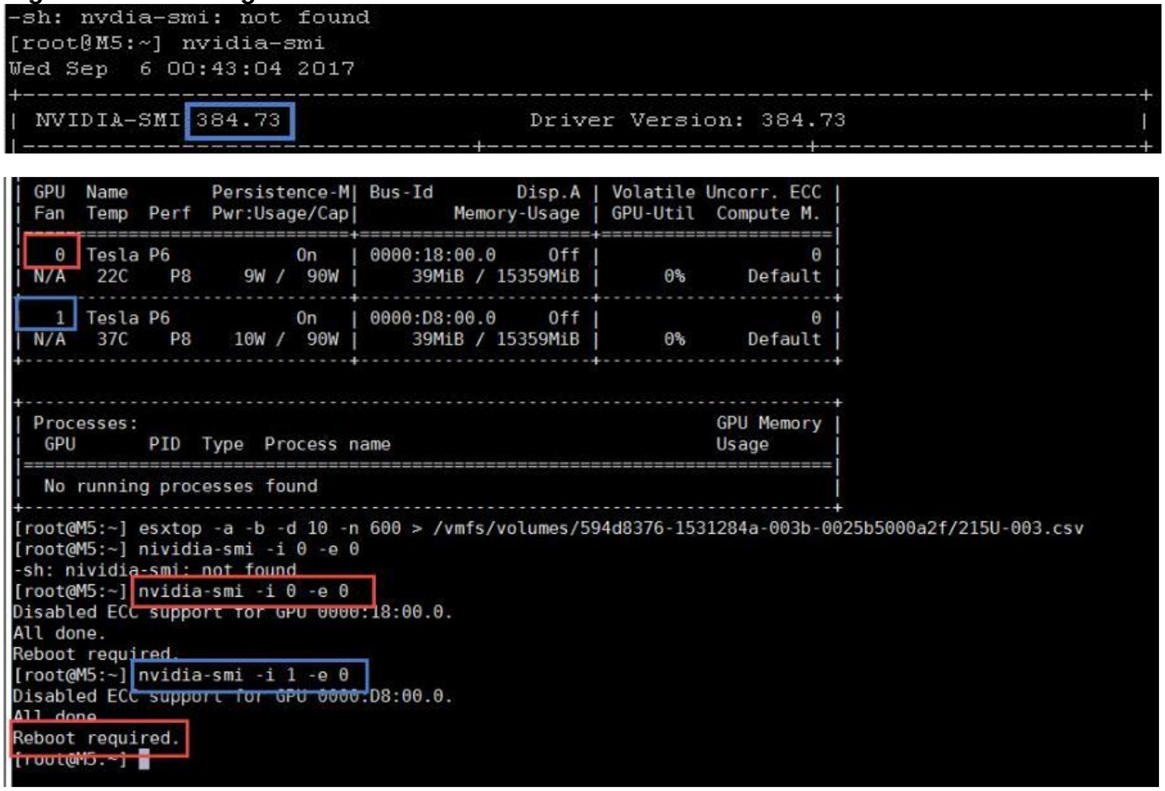

```
-sh: nvdia-smi: not found
[root@M5:~] nvidia-smi
Wed Sep  6 00:43:04 2017
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 384.73                 Driver Version: 384.73                    |
|-------------------------------+----------------------+----------------------+

| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla P6            On   | 0000:18:00.0     Off |                    0 |
| N/A   22C    P8     9W /  90W |     39MiB / 15359MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+
|   1  Tesla P6            On   | 0000:D8:00.0     Off |                    0 |
| N/A   37C    P8    10W /  90W |     39MiB / 15359MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID  Type  Process name                               Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
[root@M5:~] esxtop -a -b -d 10 -n 600 > /vmfs/volumes/594d8376-1531284a-003b-0025b5000a2f/215U-003.csv
[root@M5:~] nividia-smi -i 0 -e 0
-sh: nividia-smi: not found
[root@M5:~] nvidia-smi -i 0 -e 0
Disabled ECC support for GPU 0000:18:00.0.
All done.
Reboot required.
[root@M5:~] nvidia-smi -i 1 -e 0
Disabled ECC support for GPU 0000:D8:00.0.
All done.
Reboot required.
[root@M5:~]
```

6. Install and configure Microsoft Windows on the virtual machine:

   – Configure the virtual machine with the appropriate amount of vCPU and RAM according to the GPU profile selected.

   – Install VMware Tools.

   – Join the virtual machine to the Microsoft Active Directory domain.

   – Install or upgrade Citrix HDX 3D Pro Virtual Desktop Agent.

Figure 53 Selecting HDX 3D Pro During VDA Installation

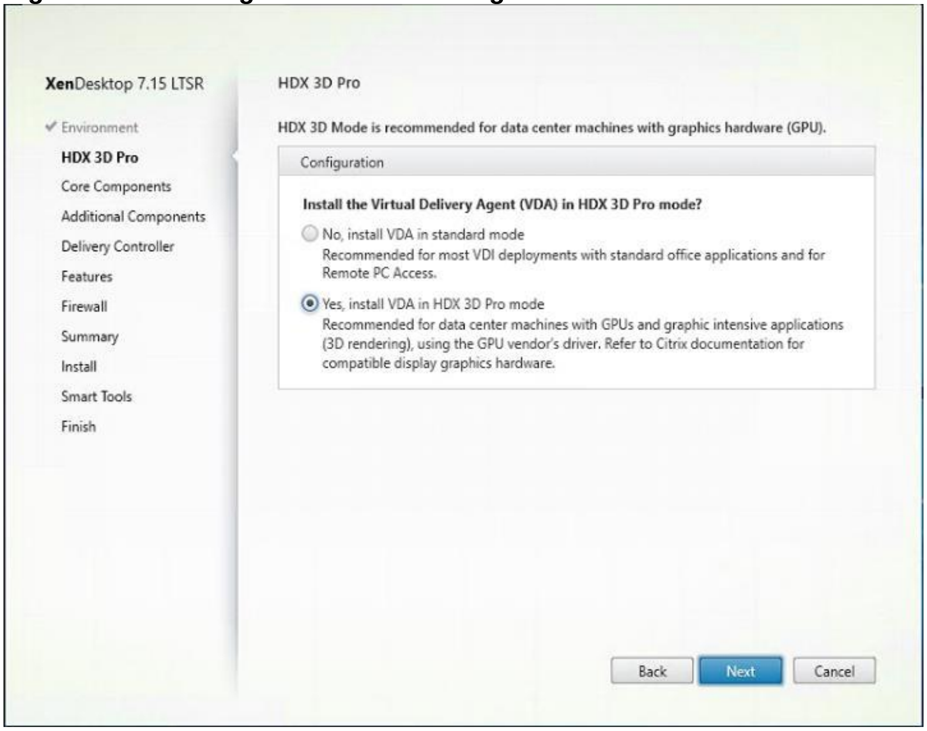

HDX 3D Pro

HDX 3D Mode is recommended for data center machines with graphics hardware (GPU).

Configuration

Install the Virtual Delivery Agent (VDA) in HDX 3D Pro mode?

○ No, install VDA in standard mode
Recommended for most VDI deployments with standard office applications and for Remote PC Access.

● Yes, install VDA in HDX 3D Pro mode
Recommended for data center machines with GPUs and graphic intensive applications (3D rendering), using the GPU vendor's driver. Refer to Citrix documentation for compatible display graphics hardware.

When you use the command-line interface (CLI) to install the VDA, include the **/enable_hdx_3d_pro** option with the XenDesktop VdaSetup.exe command.

To upgrade HDX 3D Pro, uninstall both the separate HDX 3D for Professional Graphics component and the VDA before installing the VDA for HDX 3D Pro. Similarly, to switch from the standard VDA for a Windows desktop to the HDX 3D Pro VDA, uninstall the standard VDA and then install the VDA for HDX 3D Pro.

Install the GPU Drivers inside your Windows VM

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 study, we used ESXi Host Driver version 352.83 and 354.80 for the Windows VDI image. These drivers come in the same download package from NVIDIA.

To install the GPU drivers, complete the following 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 54 NVIDIA Driver Pack

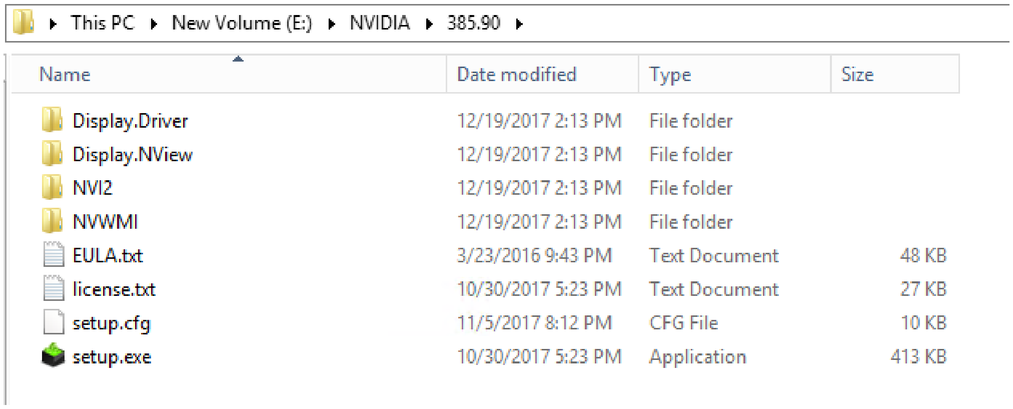

This PC ▸ New Volume (E:) ▸ NVIDIA ▸ 385.90 ▸

Name	Date modified	Type	Size
Display.Driver	12/19/2017 2:13 PM	File folder	
Display.NView	12/19/2017 2:13 PM	File folder	
NVI2	12/19/2017 2:13 PM	File folder	
NVWMI	12/19/2017 2:13 PM	File folder	
EULA.txt	3/23/2016 9:43 PM	Text Document	48 KB
license.txt	10/30/2017 5:23 PM	Text Document	27 KB
setup.cfg	11/5/2017 8:12 PM	CFG File	10 KB
setup.exe	10/30/2017 5:23 PM	Application	413 KB

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 55 Agreeing to the NVIDIA Software License

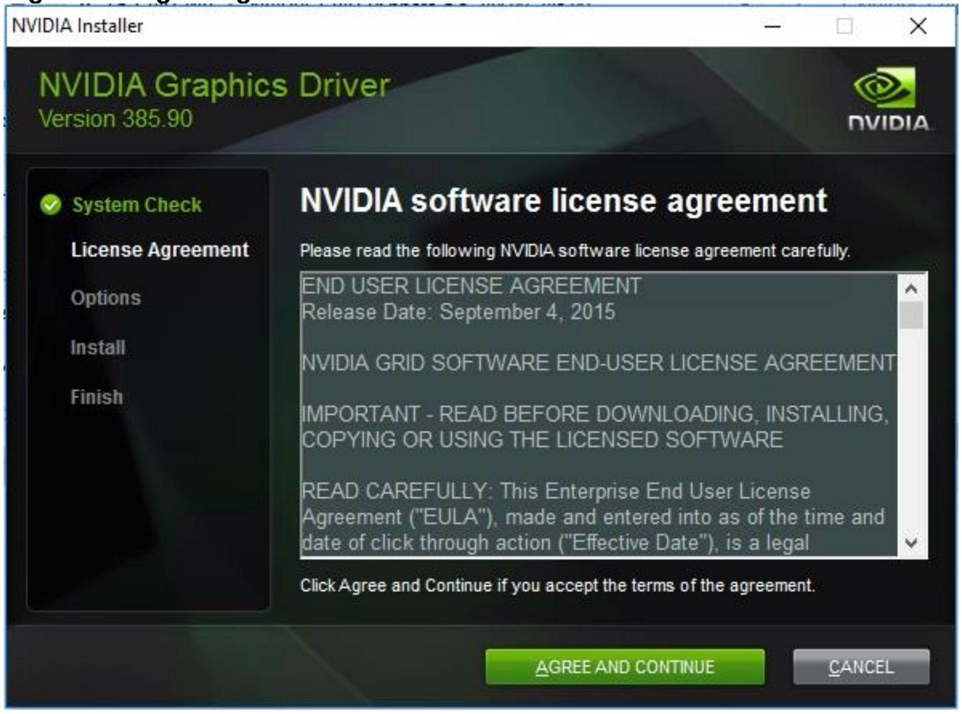

NVIDIA software license agreement

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 (Figures 51 and 52). After the installation has completed successfully, restart the virtual machine.

Be 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 56 Selecting the Express or Custom Installation Option

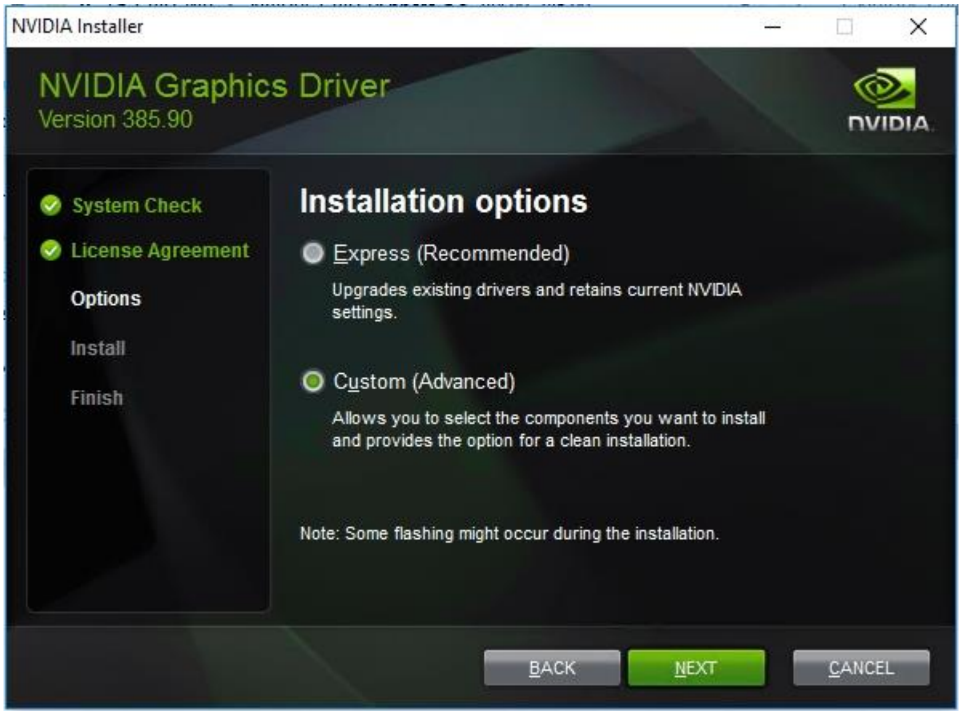

NVIDIA Installer

NVIDIA Graphics Driver
Version 385.90

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 57 Components Installed During NVIDIA Graphics Driver Custom Installation Process

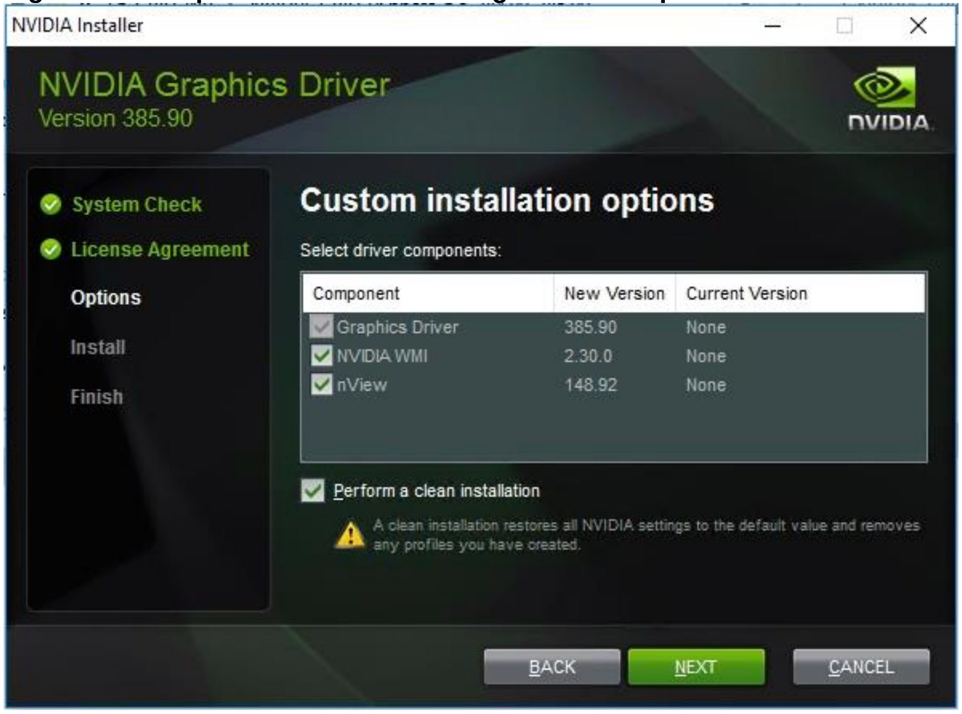

NVIDIA Installer

NVIDIA Graphics Driver
Version 385.90

Custom installation options

Select driver components:

Component	New Version	Current Version
☑ Graphics Driver	385.90	None
☑ NVIDIA WMI	2.30.0	None
☑ nView	148.92	None

☑ Perform a clean installation

A clean installation restores all NVIDIA settings to the default value and removes any profiles you have created.

Figure 58 Resarting the Virtual Machine

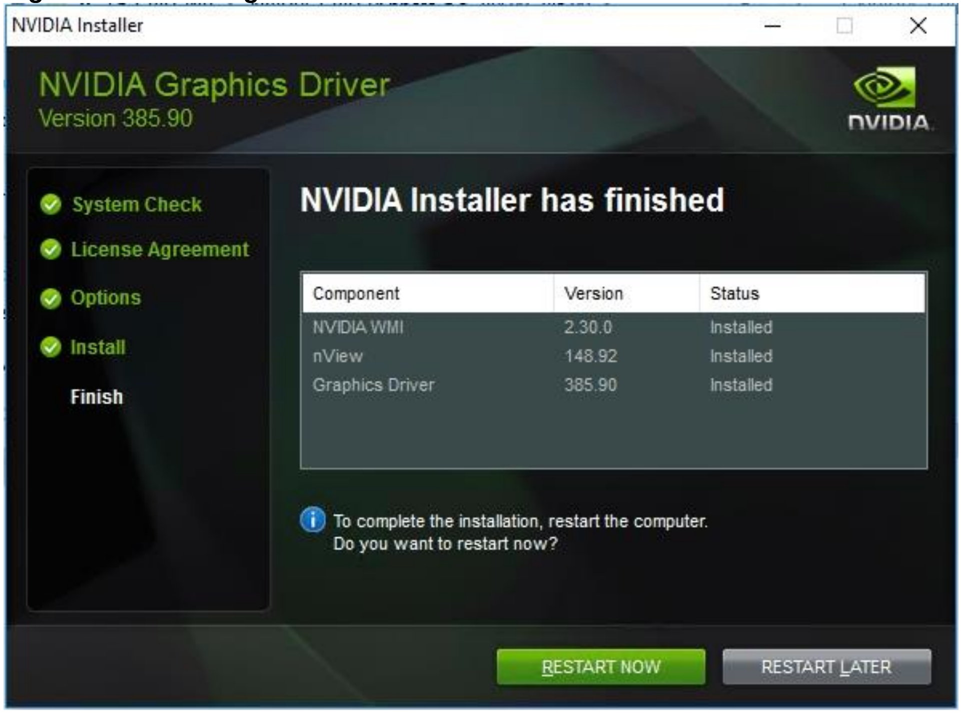

NVIDIA Installer has finished

Component	Version	Status
NVIDIA WMI	2.30.0	Installed
nView	148.92	Installed
Graphics Driver	385.90	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, we must point our master image to the license server so the VMs with vGPUs can obtain a license.

1. In Windows – Control Panel, double click the NVidia Control Panel.

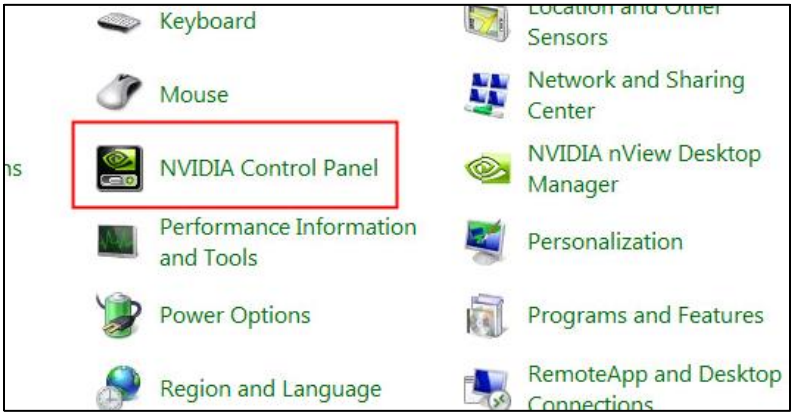

Keyboard
Location and Other Sensors
Mouse
Network and Sharing Center
ns
NVIDIA Control Panel
NVIDIA nView Desktop Manager
Performance Information and Tools
Personalization
Power Options
Programs and Features
Region and Language
RemoteApp and Desktop Connections

2. In the Control Panel, enter the IP or FQDN of the Grid License Server. You should receive a result similar to the image below.

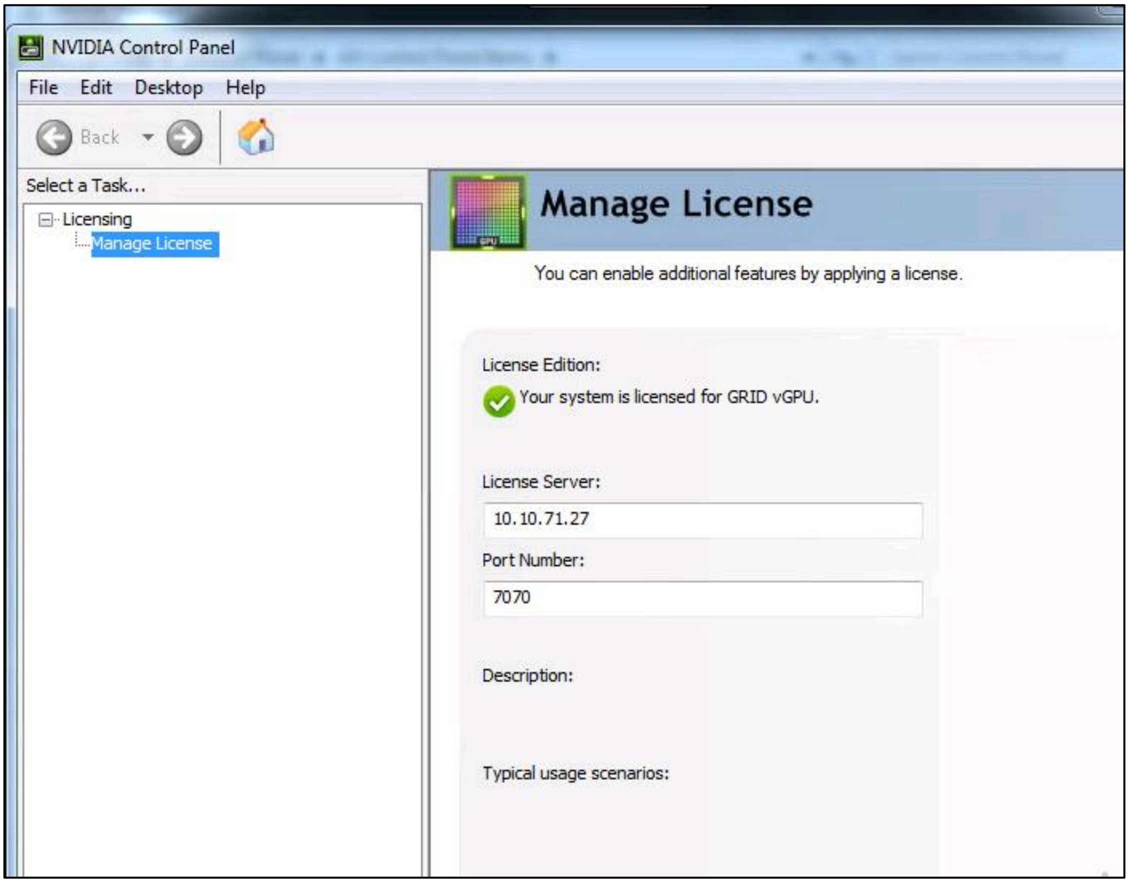

You can enable additional features by applying a license.

License Edition:
Your system is licensed for GRID vGPU.

License Server:
10.10.71.27

Port Number:
7070

Description:

Typical usage scenarios:

Cisco UCS Performance Manager

Cisco UCS Performance Manager provides visibility from a single console into Cisco UCS components for performance monitoring and capacity planning. It provides data center assurance of integrated infrastructures and ties application performance to physical and virtual infrastructure performance. This allows you to optimize resources and deliver better service levels to your customers.

The release used in this solution features an additional component, Control Center, which is an open-source, application service orchestrator based on Docker.

Control Center greatly simplifies the installation, deployment, and management of Cisco UCS Performance Manager.

This section provides a brief introduction to Control Center, and describes how it affects Cisco UCS Performance Manager deployments.

Installing Cisco UCS Performance Manager

Installing the Control Center Master Host

To install a Cisco UCS Performance Manager appliance package as a Control Center master host, using VMware vSphere, complete the following steps:

1. Download the Cisco UCS Performance Manager OVA file from the Cisco UCS Performance Manager site to your workstation.

2. Use the VMware vSphere Client to log in to vCenter as root, or as a user with superuser privileges, and then display the Home view.

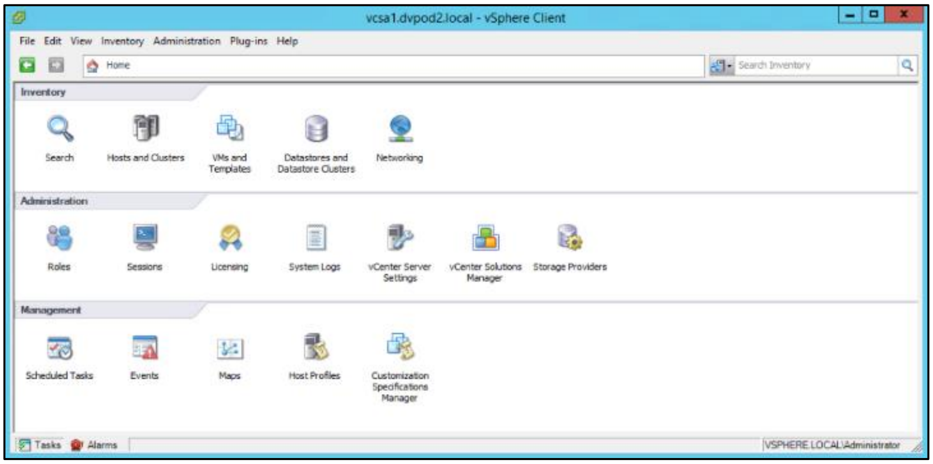

3. From the File menu, select Deploy OVF Template....

4. In the Source panel, specify the path of the Cisco UCS Performance Manager package, and then click Next >.

5. In the OVF Template Details panel, click Next.

6. In the Name and Location panel, provide a name and a location for the server:

   a. In the Name field, enter a new name or use the default.
   b. In the Inventory Location area, select a data center for the virtual machine.
   c. Click Next.

7. In the Host / Cluster panel, select a host system, and then click Next.

8. In the Storage panel, select a storage system with sufficient space for your Cisco system, and then click Next.

9. In the Disk Format panel, select Thin Provision, and then click Next.

10. In the Ready to Complete panel, review the deployment settings, and then click Finish. Please do not check the check box labeled Power on after deployment.

11. Navigate to the new virtual machine's Getting Started tab, and then click the Edit virtual machine settings link.

12. In the Virtual Machine Properties dialog, select Memory in the Hardware table.

13. In the Memory Configuration area, set the Memory Size field to 64GB, and then click the OK button.

14. On the new virtual machine's Getting Started tab, click the Power on virtual machine link.

Configure the Control Center Host Mode

Perform this procedure immediately after creating and starting a Control Center host. All Control Center deployments must include one system configured as the master host.

To configure the Control Center host mode, complete the following steps:

1. Gain access to the console interface of the Control Center host through your hypervisor console interface.

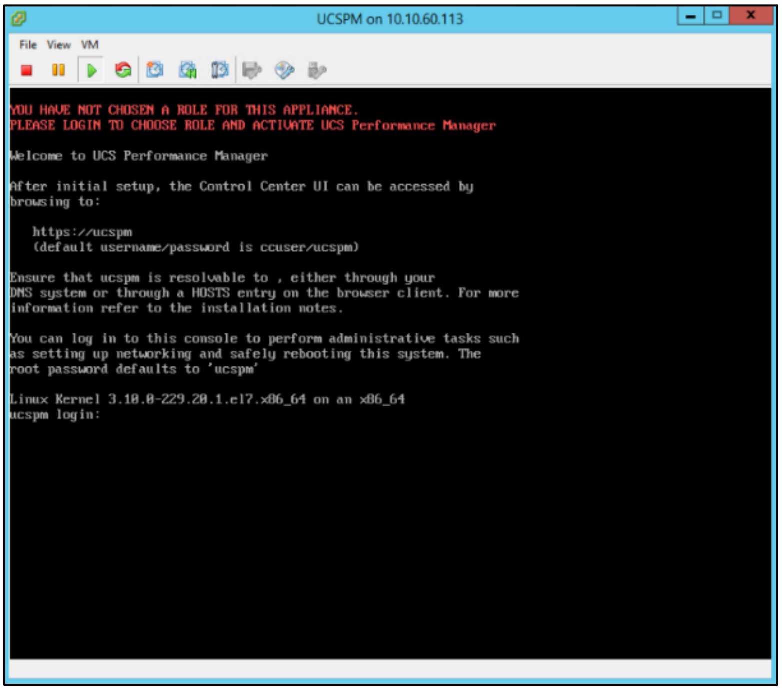

```
YOU HAVE NOT CHOSEN A ROLE FOR THIS APPLIANCE.
PLEASE LOGIN TO CHOOSE ROLE AND ACTIVATE UCS Performance Manager

Welcome to UCS Performance Manager

After initial setup, the Control Center UI can be accessed by
browsing to:

   https://ucspm
   (default username/password is ccuser/ucspm)

Ensure that ucspm is resolvable to , either through your
DNS system or through a HOSTS entry on the browser client. For more
information refer to the installation notes.

You can log in to this console to perform administrative tasks such
as setting up networking and safely rebooting this system. The
root password defaults to 'ucspm'

Linux Kernel 3.10.0-229.20.1.el7.x86_64 on an x86_64
ucspm login:
```

2. Log in as the root user.

3. The initial password is ucspm.

4. The system prompts you to enter a new password for root.

Passwords must include a minimum of eight characters, with at least one character from three of the following character classes: uppercase letter, lowercase letter, digit, and special.

5. The system prompts you to enter a new password for ccuser. The ccuser account is the default account for gaining access to the Control Center browser interface.

6. Select the master role for the host.

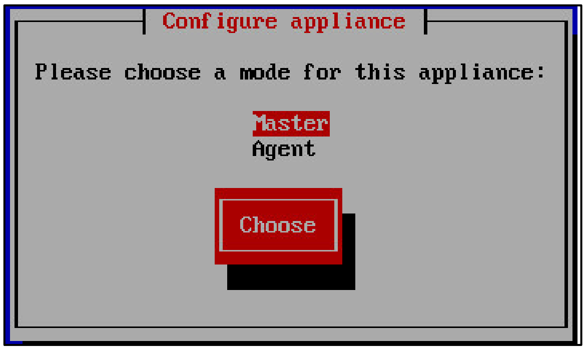

```
Configure appliance

Please choose a mode for this appliance:

                Master
                Agent

               Choose
```

7. In the Configure appliance menu, press the Tab key to select the Choose button.

8. Press the Enter key.

The system will now restart.

Edit a Connection

The default configuration for network connections is DHCP. To configure static IPv4 addressing, complete the following steps:

1. After the systems restarts, login as the root user.

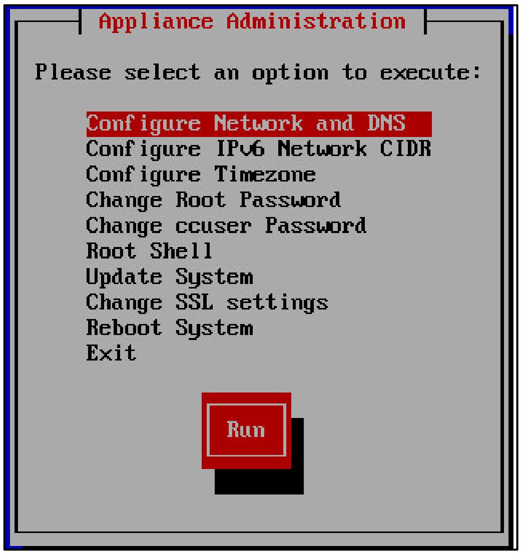

2. Select the NetworkManager TUI menu.

   a. In the Appliance Administration menu, select the Configure Network and DNS option.

   b. Press the Tab key to select the Run button.

   c. Press the Enter key.

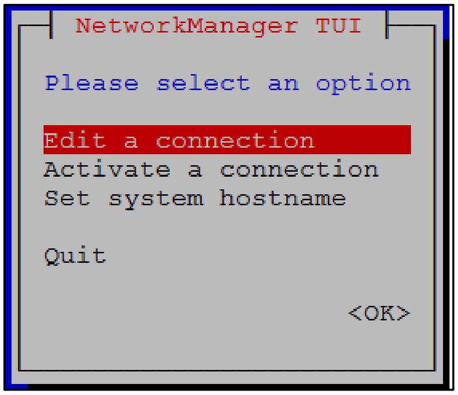

3. On the NetworkManager TUI menu, select Edit a connection, and then press the Return key.

The TUI displays the connections that are available on this host.

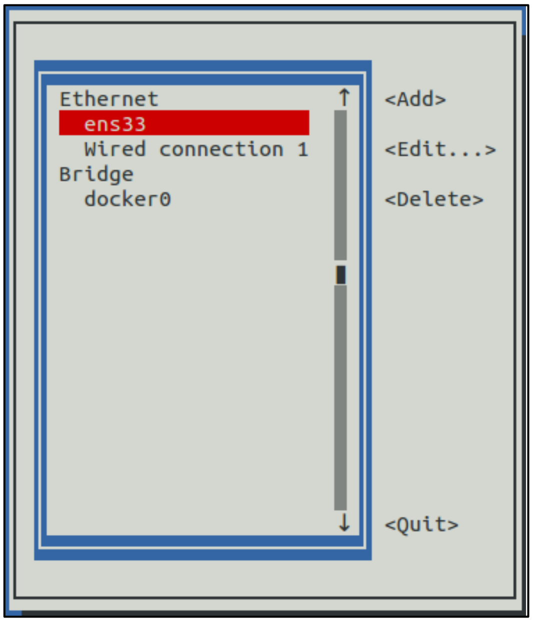

```
Ethernet                     <Add>
  ens33
  Wired connection 1         <Edit...>
Bridge
  docker0                    <Delete>



                             <Quit>
```

4. Use the down-arrow key to select Wired Connection 1, and then press the Return key.

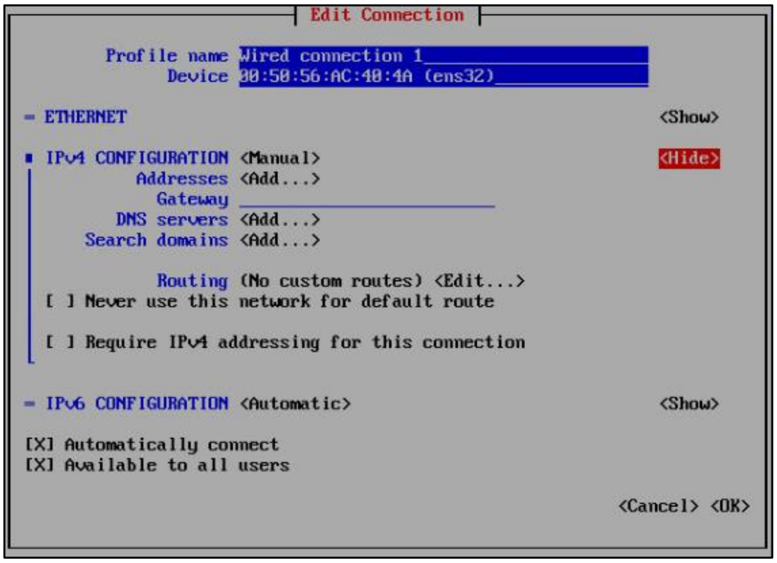

```
                         Edit Connection

       Profile name Wired connection 1_______________
             Device 00:50:56:AC:40:4A (ens32)________

= ETHERNET                                            <Show>

╤ IPv4 CONFIGURATION <Manual>                         <Hide>
│          Addresses <Add...>
│            Gateway _________________________
│        DNS servers <Add...>
│     Search domains <Add...>
│
│            Routing (No custom routes) <Edit...>
│ [ ] Never use this network for default route
│
│ [ ] Require IPv4 addressing for this connection
└

= IPv6 CONFIGURATION <Automatic>                      <Show>

[X] Automatically connect
[X] Available to all users

                                               <Cancel> <OK>
```

5. Use the Tab key and the arrow keys to navigate among options in the Edit Connection screen, and use the Return key to toggle an option or to display a menu of options.

6. Optional: If the IPv4 CONFIGURATION area is not visible, select its display option (<Show>), and then press the Return key.

7. In the IPv4 CONFIGURATION area, select <Automatic>, and then press the Return key.

8. Configure static IPv4 networking:

   a. Use the down arrow key to select Manual, and then press the Return key.

   b. Use the Tab key or the down arrow key to select the <Add...> option next to Addresses, and then press

   c. the Return key.

   d. In the Addresses field, enter an IPv4 address for the virtual machine, and then press the Return key.

   e. Repeat the preceding two steps for the Gateway and DNS servers fields.

9. Use the Tab key or the down arrow key to select the <OK> option at the bottom of the Edit Connection screen, and then press the Return key.

10. In the available connections screen, use the Tab key to select the <Quit> option, and then press the Return key.

11. Reboot the operating system:

    a. In the Appliance Administration menu, use the down-arrow key to select the Reboot System option.

    b. Press the Tab key to select the Run button.

    c. Press the Enter key.

Enabling Access to Browser Interfaces

Control Center and Cisco UCS Performance Manager have independent browser interfaces served by independent web servers.

- The Control Center web server listens at HostnameOrIP:443. So, for a Control Center master host named cc-master.example.com, the hostname-based URL to use is https://cc-master.

- The Cisco UCS Performance Manager web server listens at a virtual hostname, ucspm.HostnameOrIP:443. For a Control Center master host named cc-master.example.com, the hostname-based URL to use is https://ucspm.cc-master.

To enable access to the browser interfaces by hostname, add name resolution entries to the DNS servers in your environment, or to the hosts files of individual client systems.

- On Windows client systems, the file is C:\Windows\System32\drivers\etc\hosts.

- Linux and OS/X client systems, the file is /etc/hosts.

The following line shows the syntax of the entry to add to a name resolution file:

```
IP-Address FQDN Hostname ucspm.Hostname
```

For example, the following entry identifies a Control Center master host at IP address 10.24.164.120, hostname cc-master, in the example.com domain.

```
10.24.164.120 cc-master.example.com cc-master ucspm.cc-master
```

Deploy Cisco UCS Performance Manager

To log into Control Center for the first time, complete the following steps:

1. Display the login page of the Control Center browser interface.

2. Replace Hostname with the name of the Cisco UCS Performance Manager virtual machine.

   https://cc-master.dvpod2.local

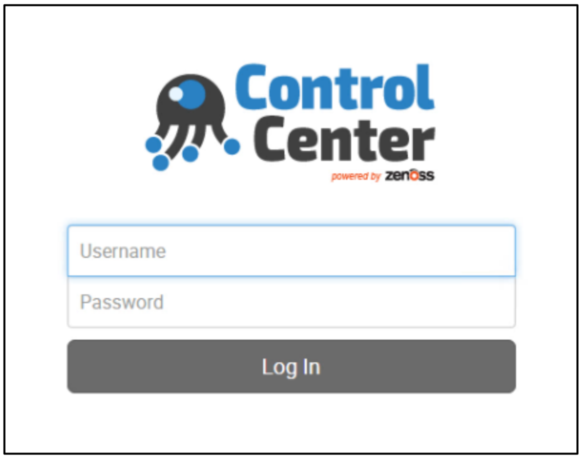

3. At the login page, enter ccuser and its password.

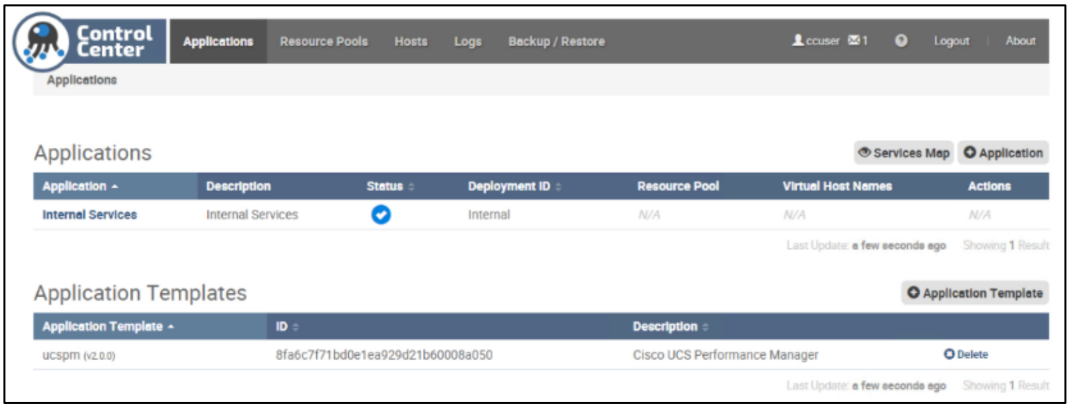

Applications

Application	Description	Status	Deployment ID	Resource Pool	Virtual Host Names	Actions
Internal Services	Internal Services	✓	Internal	N/A	N/A	N/A

Last Update: a few seconds ago Showing 1 Result

Application Templates

Application Template	ID	Description	
ucspm (v2.0.0)	8fa6c7f71bd0e1ea929d21b60008a050	Cisco UCS Performance Manager	Delete

Last Update: a few seconds ago Showing 1 Result

4. On the Applications page, click the + Application button, located at the right side of the page.

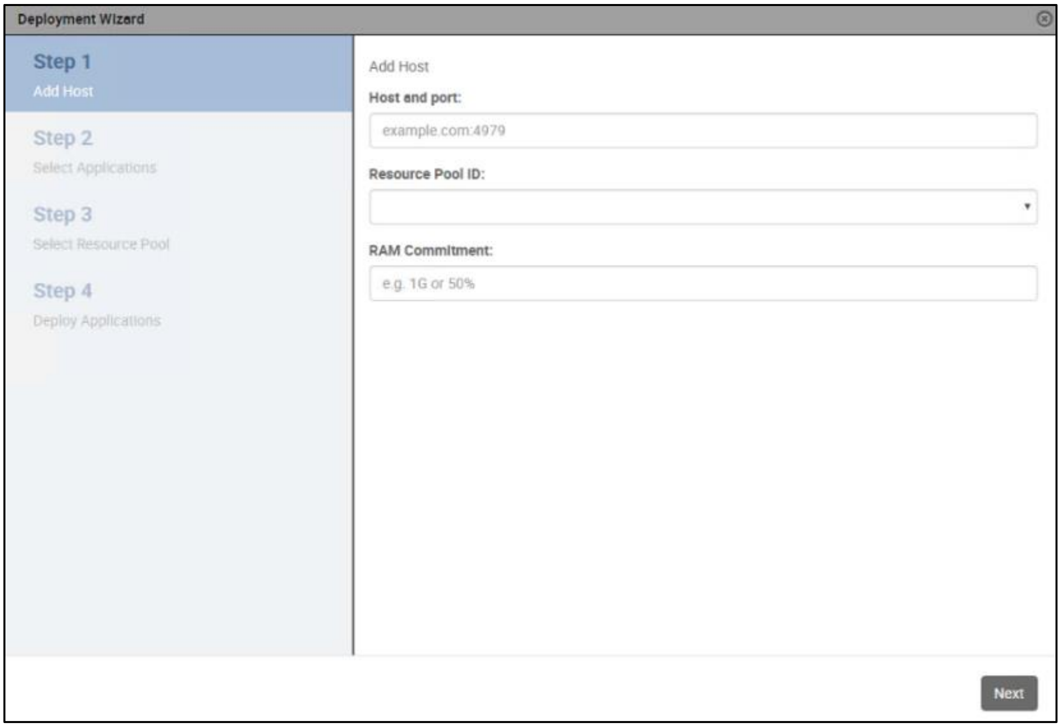

Add Host

Host and port: example.com:4979

Resource Pool ID:

RAM Commitment: e.g. 1G or 50%

5. In the Deployment Wizard, add the master host to the default resource pool. The host to add is the Control Center master host:

   a. In the Host and Port field, enter the hostname or IP address of the Control Center master host, followed by a colon character (:), and then 4979.

   b. If you enter a hostname, all hosts in your Control Center cluster must be able to resolve the name, either through an entry in /etc/hosts, or through a nameserver on your network.

   c. In the Resource Pool ID field, select default from the list, and then click Next.

   d. In the RAM Commitment field, enter the percentage of master host RAM to devote to Control Center and Cisco UCS Performance Manager.

   e. The amount of RAM required for the operating system is not included in this value. Cisco recommends entering 100 in the field.

   f. At the bottom of the Deployment Wizard, click Next.

6. Select the application to deploy.

   a. Select ucspm.

   b. At the bottom of the Deployment Wizard, click Next.

7. Select the resource pool for the application.

   a. Select default.

   b. At the bottom of the Deployment Wizard, click Next.

8. Choose a deployment ID and deploy Cisco UCS Performance Manager.

   a. In the Deployment ID field, enter a name for this deployment of Cisco UCS Performance Manager.

   b. At the bottom of the Deployment Wizard, click Deploy.

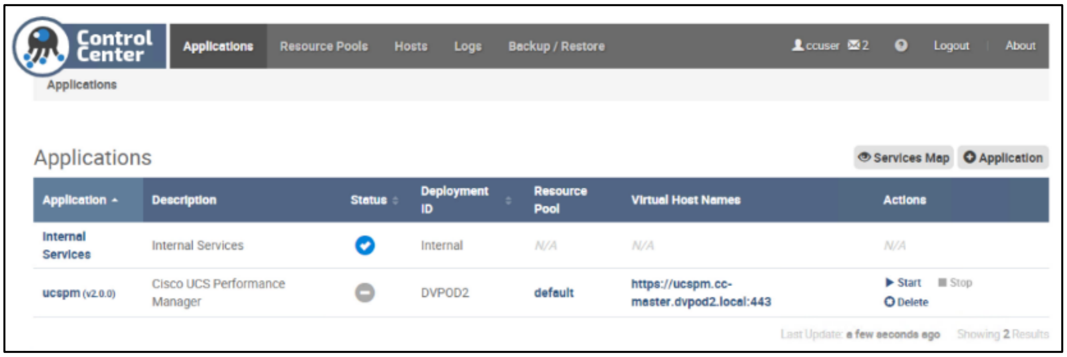

Applications

Application	Description	Status	Deployment ID	Resource Pool	Virtual Host Names	Actions
Internal Services	Internal Services	✓	Internal	N/A	N/A	N/A
ucspm (v2.0.0)	Cisco UCS Performance Manager	⊖	DVPOD2	default	https://ucspm.cc-master.dvpod2.local:443	▶ Start ■ Stop ○ Delete

Last Update: a few seconds ago Showing 2 Results

9. At the top of the page, click Logout. The control is located at the right side of the page.

10. In the Actions column of the Applications table, click the Start control of the ucspm row.

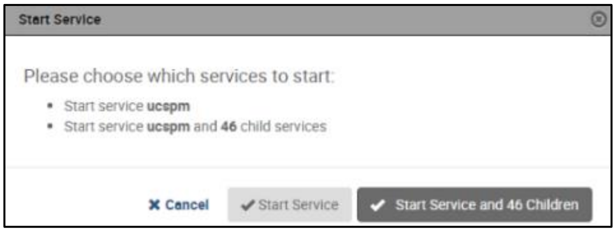

Start Service

Please choose which services to start:

- Start service **ucspm**
- Start service **ucspm** and **46** child services

11. In the Start Service dialog, click Start Service and 46 Children button.

12. In the Application column of the Applications table, click ucspm in the ucspm row.

13. Scroll down to watch child services starting.

> Typically, child services take 4-5 minutes to start. When no child service shows a red exclamation point icon, Cisco UCS Performance Manager is running.

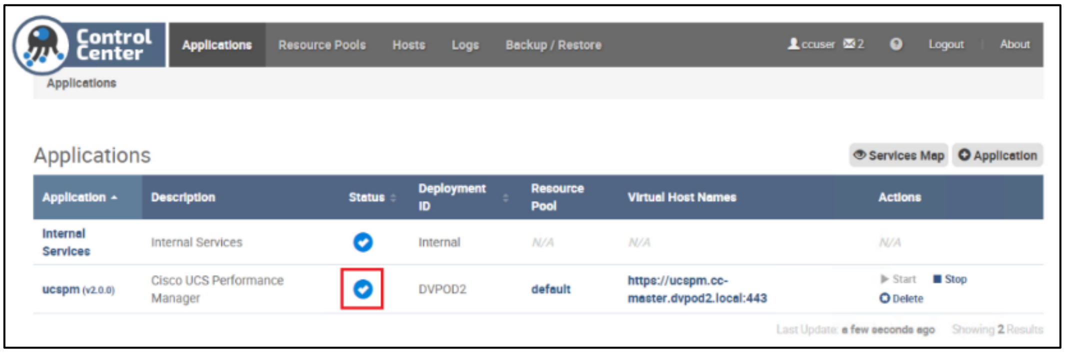

Applications

Application	Description	Status	Deployment ID	Resource Pool	Virtual Host Names	Actions
Internal Services	Internal Services	✓	Internal	N/A	N/A	N/A
ucspm (v2.0.0)	Cisco UCS Performance Manager	✓	DVPOD2	default	https://ucspm.cc-master.dvpod2.local:443	▶ Start ■ Stop ○ Delete

Last Update: a few seconds ago Showing 2 Results

Setting up Cisco UCS Performance Manager

This section describes how to use the Cisco UCS Performance Manager Setup Wizard to accept the end-user license agreement, to provide your license key, define users and passwords, to set up UCS Domains, and to add additional infrastructure.

Initial Setup

After installing Cisco UCS Performance Manager on a virtual machine, and starting it in Control Center, complete the following steps:

1. In a web browser, navigate to the login page of the Cisco UCS Performance Manager interface. Cisco UCS Performance Manager redirects the first login attempt to the Setup page, which includes the End User License Agreement (EULA) dialog.

2. Read through the agreement. At the bottom of the EULA dialog, check the checkbox on the left side, and then click the Accept License button on the right side.

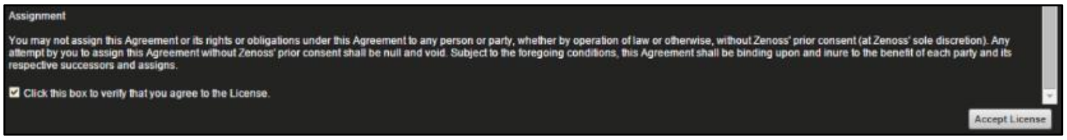

Assignment

You may not assign this Agreement or its rights or obligations under this Agreement to any person or party, whether by operation of law or otherwise, without Zenoss' prior consent (at Zenoss' sole discretion). Any attempt by you to assign this Agreement without Zenoss' prior consent shall be null and void. Subject to the foregoing conditions, this Agreement shall be binding upon and inure to the benefit of each party and its respective successors and assigns.

☑ Click this box to verify that you agree to the License.

3. On the Cisco UCS Performance Manager Setup page, click Get Started!

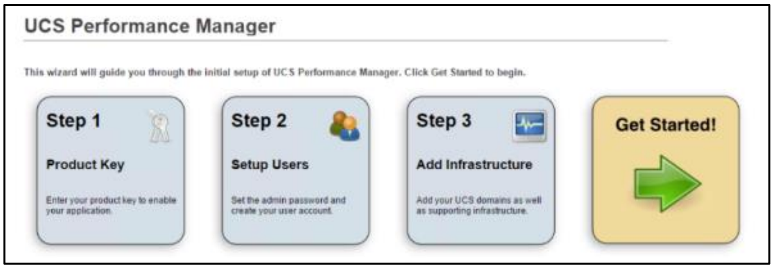

UCS Performance Manager

This wizard will guide you through the initial setup of UCS Performance Manager. Click Get Started to begin.

Step 1 Product Key Enter your product key to enable your application.

Step 2 Setup Users Set the admin password and create your user account.

Step 3 Add Infrastructure Add your UCS domains as well as supporting infrastructure.

4. On the Add Licenses page, click the Add License File button.

> If you do not have your license file yet, you can use the trial version for up to 30 days. You can enter your license file at a later date through the user interface. See the "Product Licensing" section of the Cisco UCS Performance Manager Administration Guide.

5. In the Open dialog, select your license file, and then click Open.

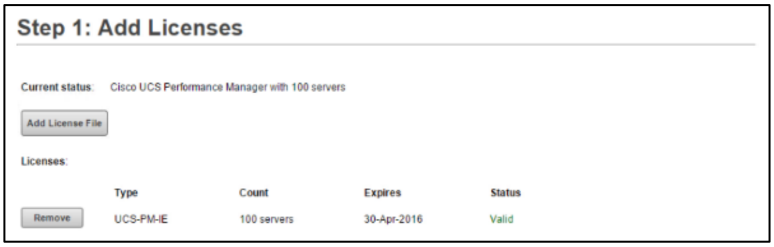

Step 1: Add Licenses

Current status: Cisco UCS Performance Manager with 100 servers

Licenses:

	Type	Count	Expires	Status
Remove	UCS-PM-IE	100 servers	30-Apr-2016	Valid

6. Proceed to the next task or repeat the preceding step.

7. In the Set admin password area, enter and confirm a password for the admin user account.

> Passwords must contain a minimum of 8 characters, including one capital letter and one digit.

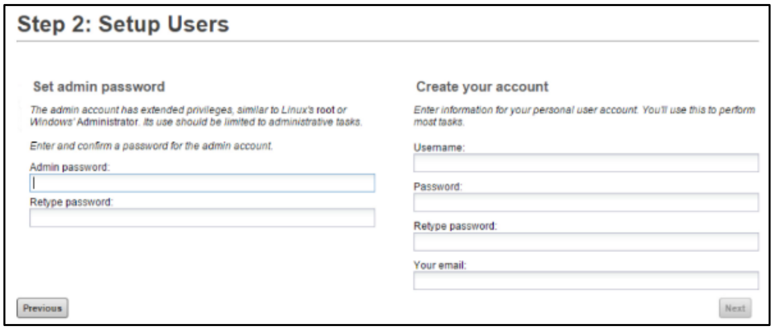

Step 2: Setup Users

Set admin password

The admin account has extended privileges, similar to Linux's root *or Windows'* Administrator*. Its use should be limited to administrative tasks.*

Enter and confirm a password for the admin account.

Admin password:

Retype password:

Create your account

Enter information for your personal user account. You'll use this to perform most tasks.

Username:

Password:

Retype password:

Your email:

8. In the Create your account area, create one additional administrative user account name and password.

9. Click Next.

Add Cisco UCS Domains

To add the Cisco UCS Domain to Cisco UCS Performance Manager after completing the initial setup configuration, complete the following steps:

1. On the Add UCS Domains page, provide connection credentials for one or more Cisco UCS domains.

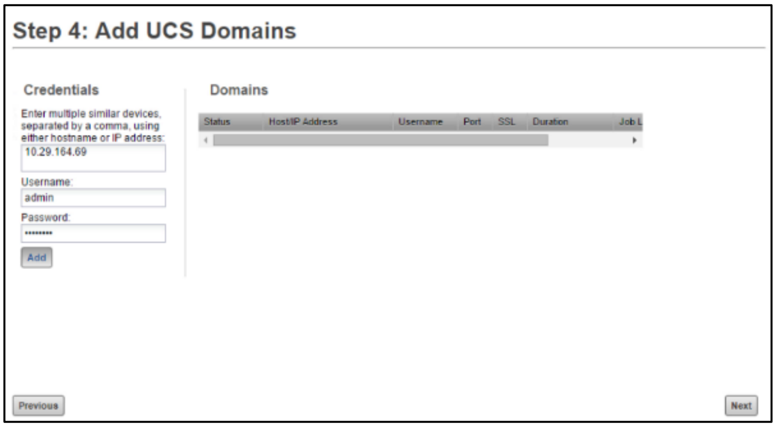

Step 4: Add UCS Domains

Credentials

Enter multiple similar devices, separated by a comma, using either hostname or IP address: 10.29.164.69

Username: admin

Password: ••••••••

Status	Host/IP Address	Username	Port	SSL	Duration	Job L

a. In the Enter multiple similar devices, separated by a comma, using either hostname or IP address field, enter the fully-qualified domain name or IP address of a UCS domain server.

b. In the Username field, enter the name of a user account in the UCS domain that is authorized for read access to the resources you plan to monitor.

c. In the Password field, enter the password of the user account specified in the preceding step.

d. Click Add.

2. Review the information in the Status column of the Domains table, and then remove a domain, add a domain, or continue.

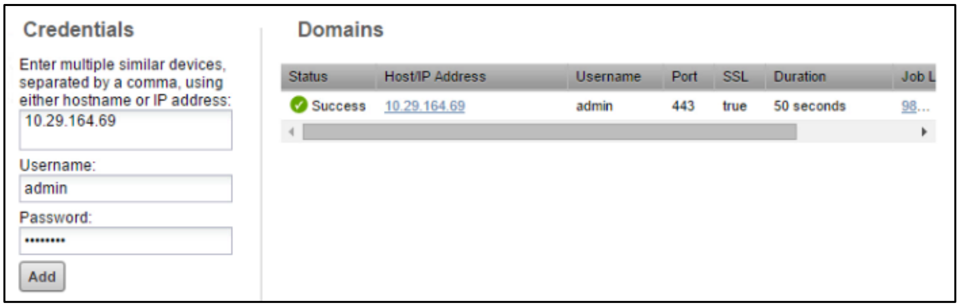

Credentials

Enter multiple similar devices, separated by a comma, using either hostname or IP address: 10.29.164.69

Username: admin

Password: ••••••••

Status	Host/IP Address	Username	Port	SSL	Duration	Job L
✓ Success	10.29.164.69	admin	443	true	50 seconds	98...

If the final message in the Status column is Failure, click the button in the Remove column, and then try again to add a domain.

If the final message in the Status column is Success, you may add another domain or continue to the next page.

3. Click Next to continue to the Add Infrastructure step.

Adding Infrastructure Devices

To add the Infrastructure Devices to Cisco UCS Performance Manager after completing the initial setup configuration, complete the following steps:

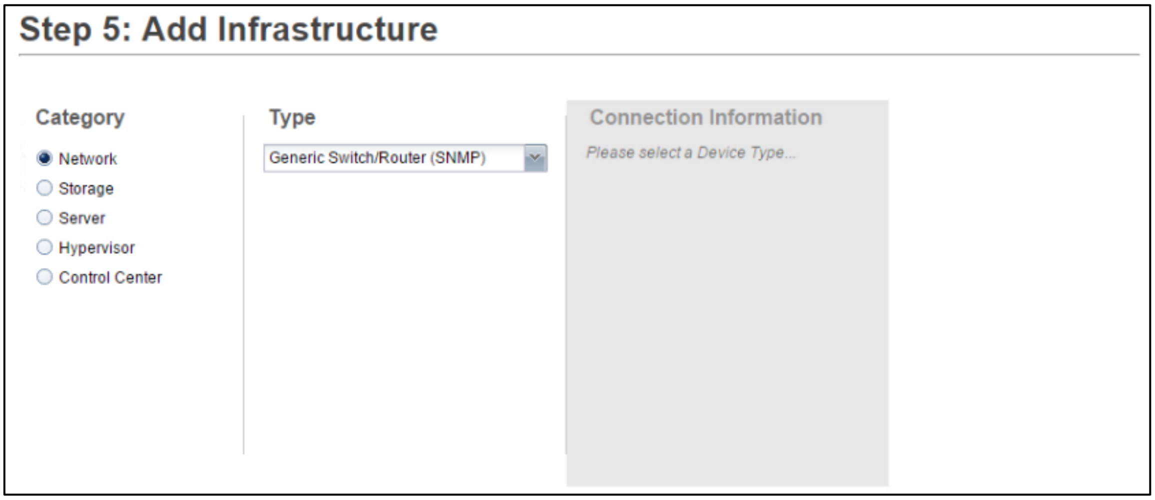

Step 5: Add Infrastructure

Category

- ● Network
- ○ Storage
- ○ Server
- ○ Hypervisor
- ○ Control Center

Type: Generic Switch/Router (SNMP)

Connection Information

Please select a Device Type...

1. This step is optional. Click Finish to exit the Setup Wizard. You will then be taken to the Dashboard.

2. The Setup Wizard times out after 20 minutes if you have not completed it. You may restart Setup Wizard by closing its browser window or tab, and then logging in again. Also, you may add devices through the Add Infrastructure page at any time.

3. As it relates to this solution, other infrastructure devices that can be added include the Cisco Nexus 1000V, NetApp storage using Data ONTAP API (ZAPI), ESXi hosts using SOAP, and Windows Servers using SNMP or WinRM.

Add Nexus 9000 Series Switches

Perform this procedure to add the Infrastructure Devices to Cisco UCS Performance Manager after completing the initial setup configuration.

In order to monitor Cisco Nexus 9000 Series devices, you must first enable NX-API with the feature manager CLI command on the device. For detailed instructions on performing this task, see the following Cisco documentation: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/programmability/guide/b_Cisco_Nexus_9000_Series_NXOS_Programmability_Guide/b_Cisco_Nexus_9000_Series_NXOS_Programmability_Configuration_Guide_chapter_0101.html#concept_BCCB1EFF9C4A4138BECE9ECC0C4E38DF

1. In the Category area, select Network.

2. In the Type list, select Cisco Nexus 9000 (SNMP + Netconf).

The protocol used to gather data from the device is included in the list, in parentheses.

3. In the Connection Information area, specify the two 93180 switches to add.

   a. In the Enter multiple similar devices, separated by a comma, using either hostname or IP Address field, enter the hostname or IP address of one or more switch or router devices on your network.

   b. In the Username or Netconf Username field, enter the name of a user account on the device.

   c. In the Password or Netconf Password field, enter the password of the user account specified in the previous field.

   d. Click Add.

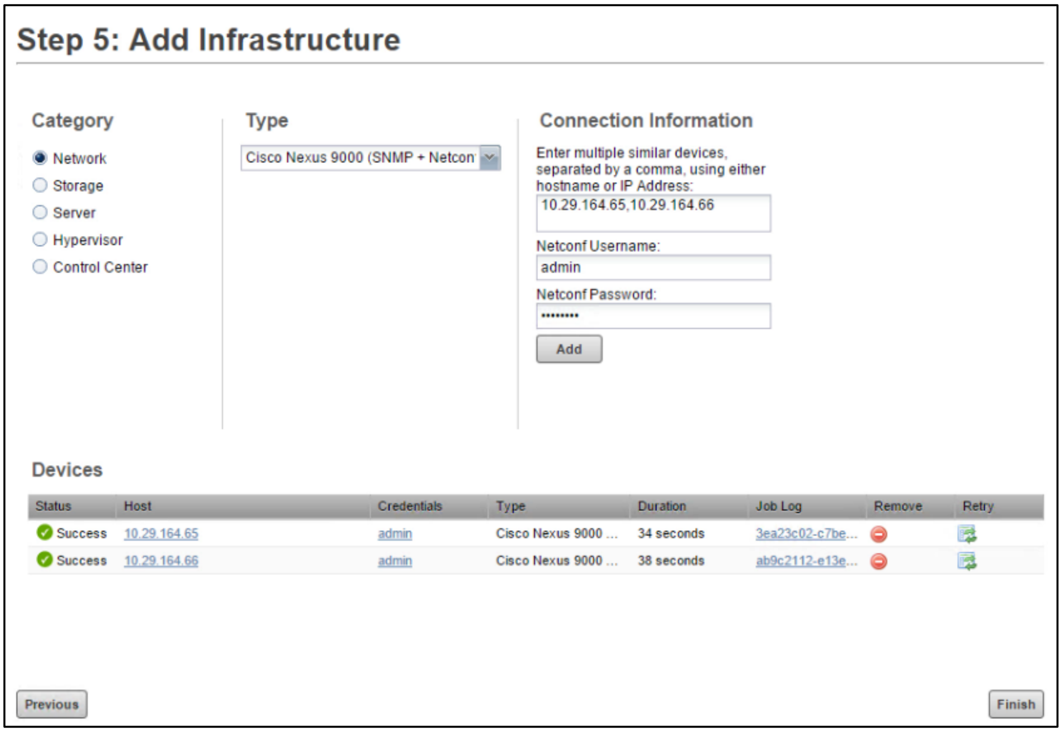

Step 5: Add Infrastructure

Category

● Network
○ Storage
○ Server
○ Hypervisor
○ Control Center

Type

Cisco Nexus 9000 (SNMP + Netcon

Connection Information

Enter multiple similar devices, separated by a comma, using either hostname or IP Address:
10.29.164.65,10.29.164.66

Netconf Username:
admin

Netconf Password:
••••••••

Devices

Status	Host	Credentials	Type	Duration	Job Log	Remove	Retry
✓ Success	10.29.164.65	admin	Cisco Nexus 9000 ...	34 seconds	3ea23c02-c7be...	⊖	↻
✓ Success	10.29.164.66	admin	Cisco Nexus 9000 ...	38 seconds	ab9c2112-e13e...	⊖	↻

4. When finished adding network devices, click Next.

Cisco UCS Performance manager, in addition to monitoring Cisco hardware operations, is able to monitor vSphere environment.

The following operational reports are available for vSphere.

■ Clusters - Shows all clusters, with the count of VMs (total and powered on), hosts, and CPU/Memory utilization within each cluster.

■ Datastores - Shows all datastores, with the number of connected VMs (total and powered on) and the disk space available and consumed on each datastore.

■ Hosts - Shows all hosts, with the count of VMs (total and powered on), hosts, CPU/Memory reservation and utilization on each host

■ VMs - Shows all VMs, their operating system, CPU/Memory utilization, and which host/cluster they reside within.

■ VMware Utilization - Provides a summary of VMs, CPU, memory, and disk utilization over a specified time interval, broken down by host.

Cisco UCS Performance Manager Sample Test Data

The following samples represent just some of the useful data that can be obtained using UCS Performance Manager.

The chart shows the network usage from a Fabric Interconnect (Fabric A) during the boot storm and 6000 user mixed workload test.

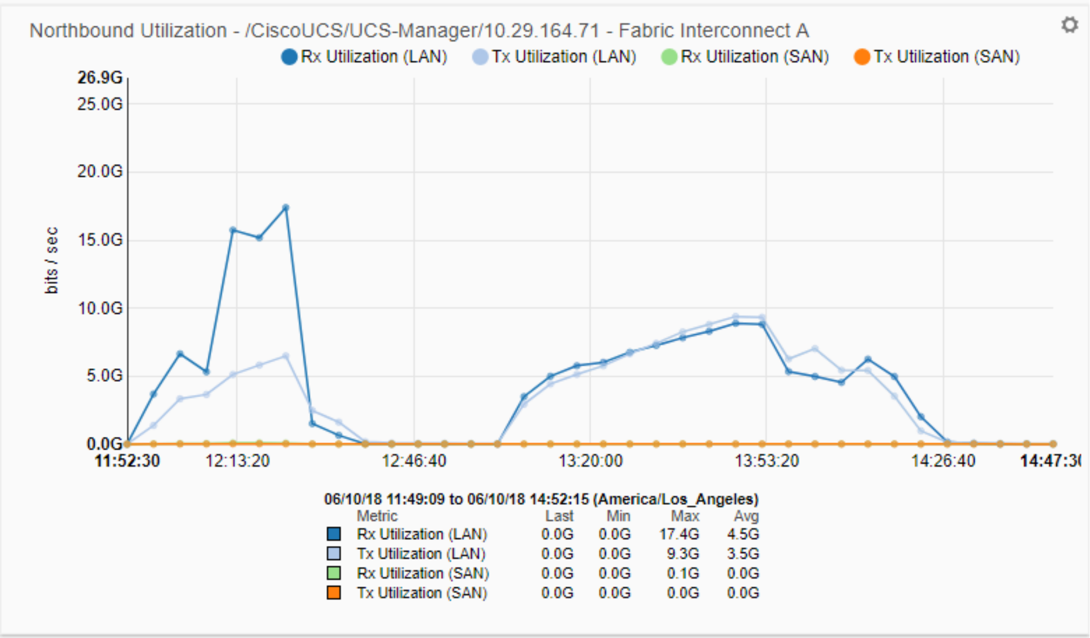

06/10/18 11:49:09 to 06/10/18 14:52:15 (America/Los_Angeles)

	Metric	Last	Min	Max	Avg
🟦	Rx Utilization (LAN)	0.0G	0.0G	17.4G	4.5G
🟦	Tx Utilization (LAN)	0.0G	0.0G	9.3G	3.5G
🟩	Rx Utilization (SAN)	0.0G	0.0G	0.1G	0.0G
🟧	Tx Utilization (SAN)	0.0G	0.0G	0.0G	0.0G

The chart shows the throughput from Nexus switch (Fabric A) during the 6000 user mixed workload test.

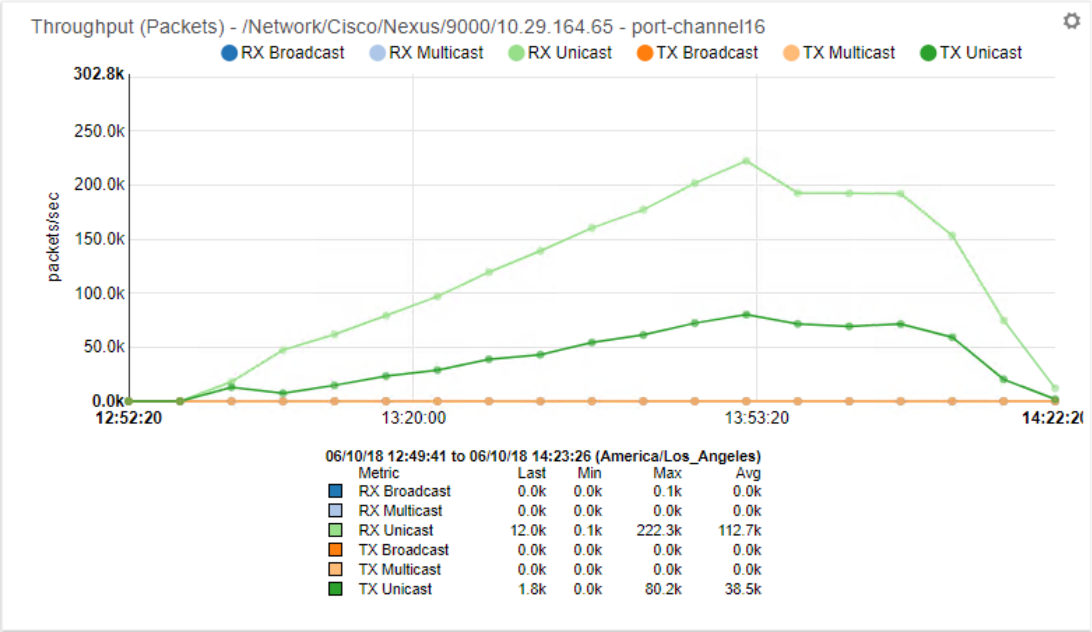

06/10/18 12:49:41 to 06/10/18 14:23:26 (America/Los_Angeles)

	Metric	Last	Min	Max	Avg
🟦	RX Broadcast	0.0k	0.0k	0.1k	0.0k
🟦	RX Multicast	0.0k	0.0k	0.0k	0.0k
🟩	RX Unicast	12.0k	0.1k	222.3k	112.7k
🟧	TX Broadcast	0.0k	0.0k	0.0k	0.0k
🟧	TX Multicast	0.0k	0.0k	0.0k	0.0k
🟩	TX Unicast	1.8k	0.0k	80.2k	38.5k

Test Setup and Configurations

In this solution, we tested a single Cisco UCS B200 M5 blade server to validate against the performance of one blade and thirty Cisco UCS B200 M5 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 each workload on a single blade to determine the Recommended Maximum Workload per host server using XenApp/XenDesktop 7.15 with 270 RDS sessions, 205 VDI Non-Persistent sessions, and 205 VDI Persistent sessions.

Figure 59 Cisco UCS B200 M5 Blade Server for Single Server Scalability XenApp 7.15 RDS with PVS 7.15

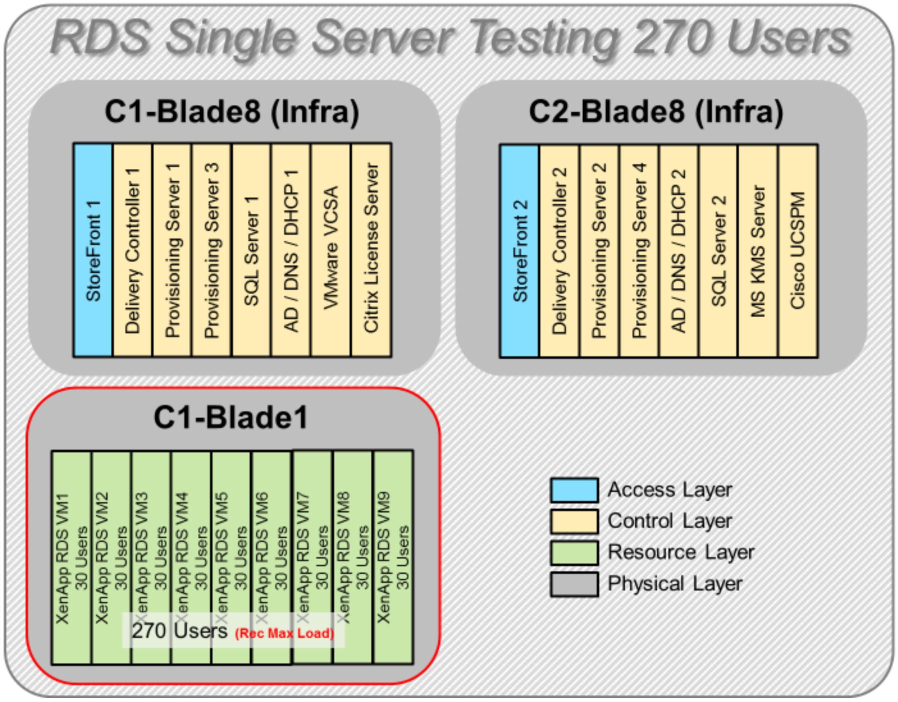

Figure 60 Cisco UCS B200 M5 Blade Server for Single Server Scalability XenDesktop 7.15 VDI (Non-Persistent) with PVS 7.15

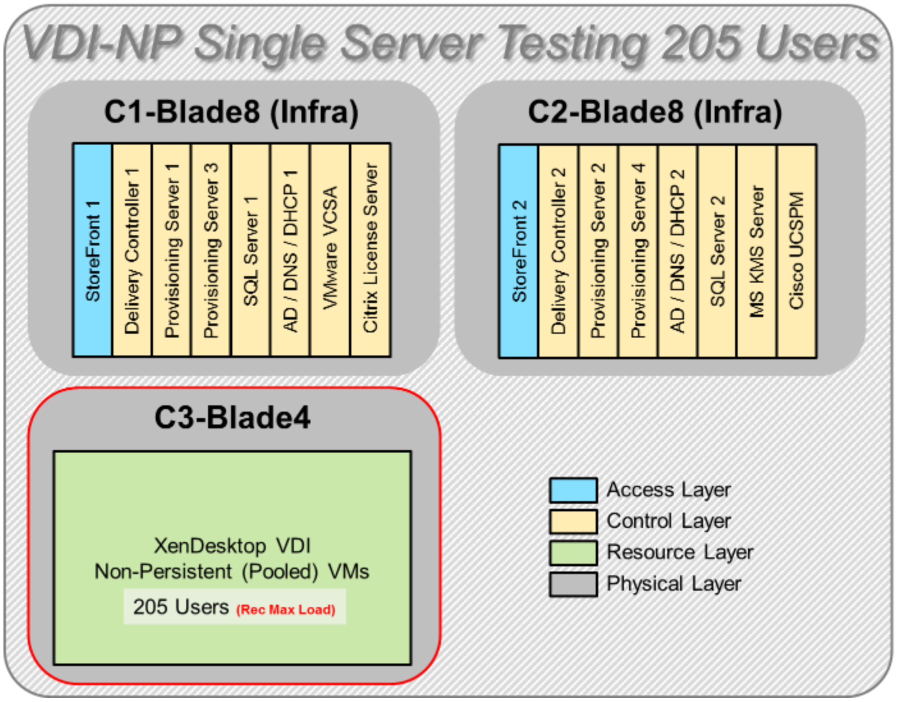

Figure 61 Cisco UCS B200 M5 Blade Server for Single Server Scalability XenDesktop 7.15 VDI (Persistent) Full Clones

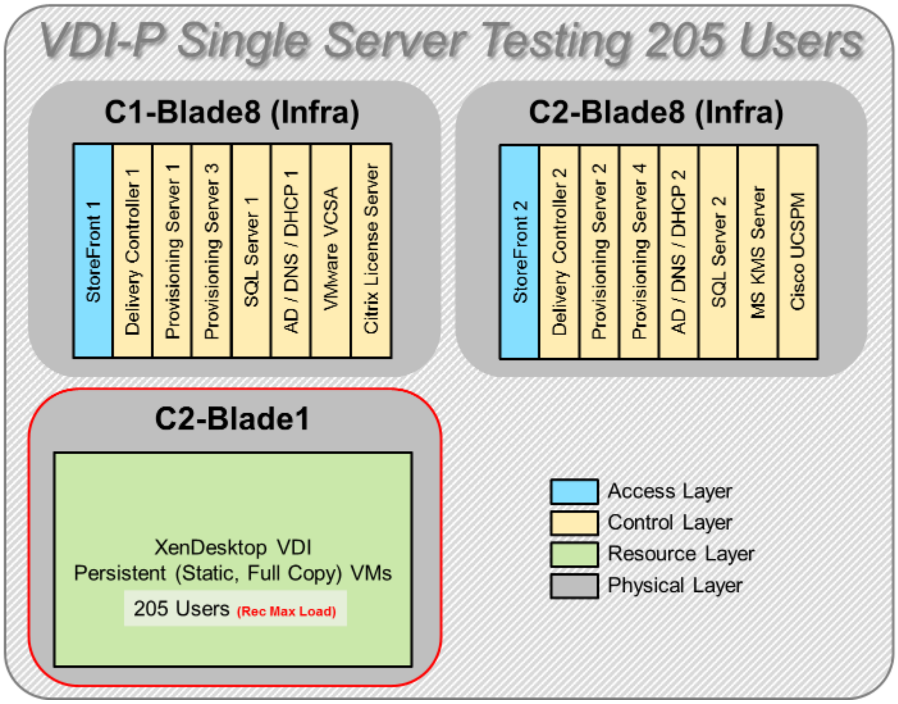

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332 - 16 UP Fabric Interconnects

- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.60-GHz 10-core processors, 128GB 2133MHz RAM for all host blades
- 1 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM for all host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Storage Switches
- 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

- Cisco UCS firmware 3.2(3d)
- VMware ESXi 6.5 Update 1 for host blades
- Citrix XenApp/XenDesktop 7.15 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.15
- Citrix User Profile Manager
- Microsoft SQL Server 2016
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB vdisk (master)
- Microsoft Windows Server 2016, 9vCPU, 24GB RAM, 40 GB vdisk (master)
- Microsoft Office 2016
- Login VSI 4.1.25 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Configuration for Cluster Testing

This test case validates three workload clusters using XenApp/XenDesktop 7.15 with 1900 RDS sessions, 2050 VDI Non-Persistent sessions, and 2050 VDI Persistent sessions. Server N+1 fault tolerance is factored into this test scenario for each workload and infrastructure cluster.

Figure 62 RDS Cluster Test Configuration with Eight Blades

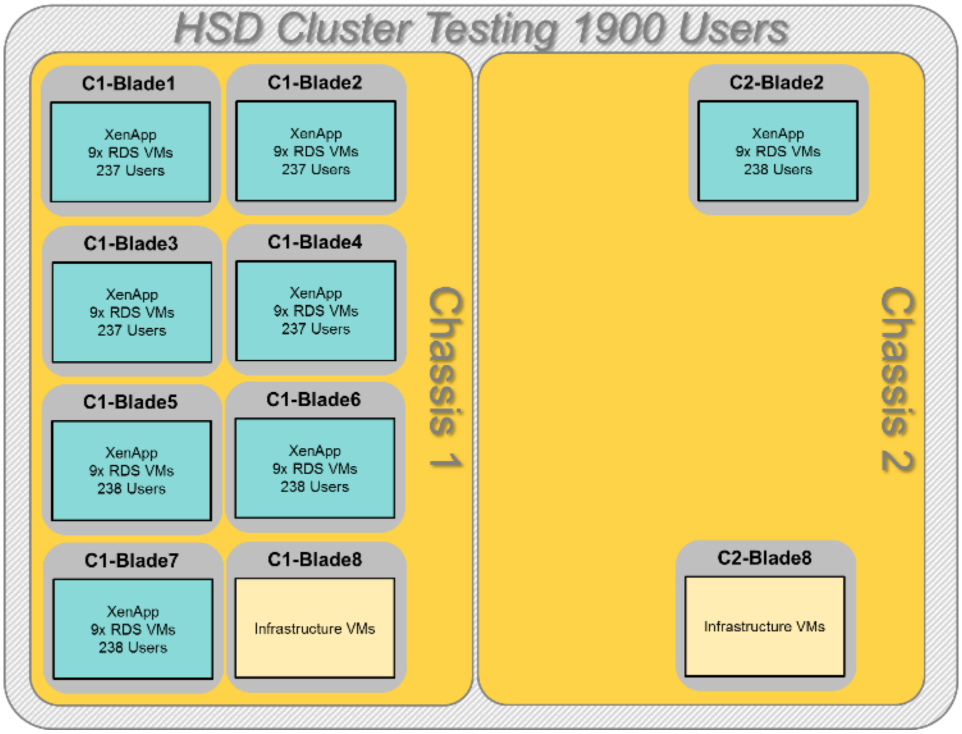

Figure 63 VDI Persistent Cluster Test Configuration with Eleven

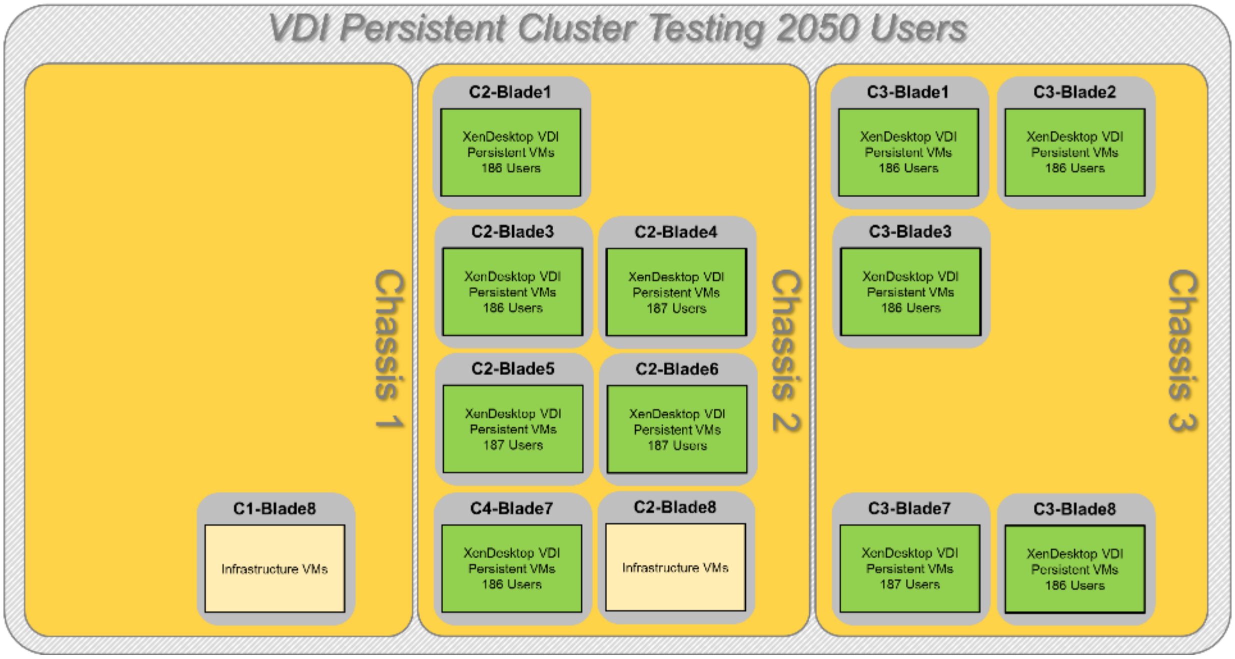

Figure 64 Blades VDI Non-Persistent Cluster Test Configuration with Eleven Blades

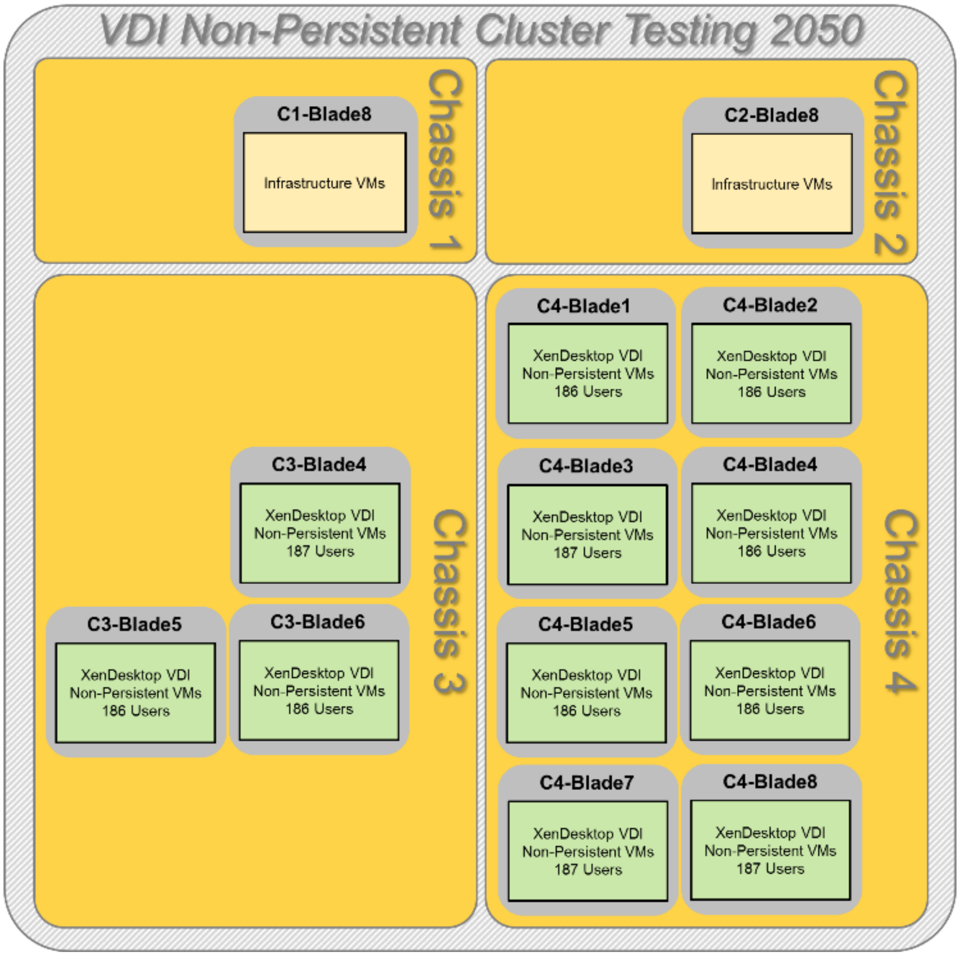
VDI Non-Persistent Cluster Testing 2050
Chassis 1
C1-Blade8
Infrastructure VMs
Chassis 2
C2-Blade8
Infrastructure VMs
Chassis 3
C3-Blade4
XenDesktop VDI Non-Persistent VMs 187 Users
C3-Blade5
XenDesktop VDI Non-Persistent VMs 186 Users
C3-Blade6
XenDesktop VDI Non-Persistent VMs 186 Users
Chassis 4
C4-Blade1
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade2
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade3
XenDesktop VDI Non-Persistent VMs 187 Users
C4-Blade4
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade5
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade6
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade7
XenDesktop VDI Non-Persistent VMs 187 Users
C4-Blade8
XenDesktop VDI Non-Persistent VMs 187 Users

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332 - 16 UP Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.60-GHz 10-core processors, 128GB 2133MHz RAM for all host blades
- 8 (RDS Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM for all host blades
- 11 (VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM for all host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Storage Switches
- 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

- Cisco UCS firmware 3.2(3d)
- VMware ESXi 6.5 Update 1 for host blades
- Citrix XenApp/XenDesktop 7.15 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.15
- Citrix User Profile Manager
- Microsoft SQL Server 2016
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB vdisk (master)
- Microsoft Windows Server 2016, 9vCPU, 24GB RAM, 40 GB vdisk (master)
- Microsoft Office 2016
- Login VSI 4.1.25 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Configuration for Full Scale Testing

This test case validates twenty-eight blades mixed workloads using XenApp/XenDesktop 7.15 with 1,900 RDS sessions, 2,050 VDI Non-Persistent sessions, and 2,050 VDI Persistent sessions for a total sum of 6,000 users. Server N+1 fault tolerance is factored into this solution for each workload and infrastructure cluster.

Figure 65 Full Scale Test Configuration with Thirty Blades

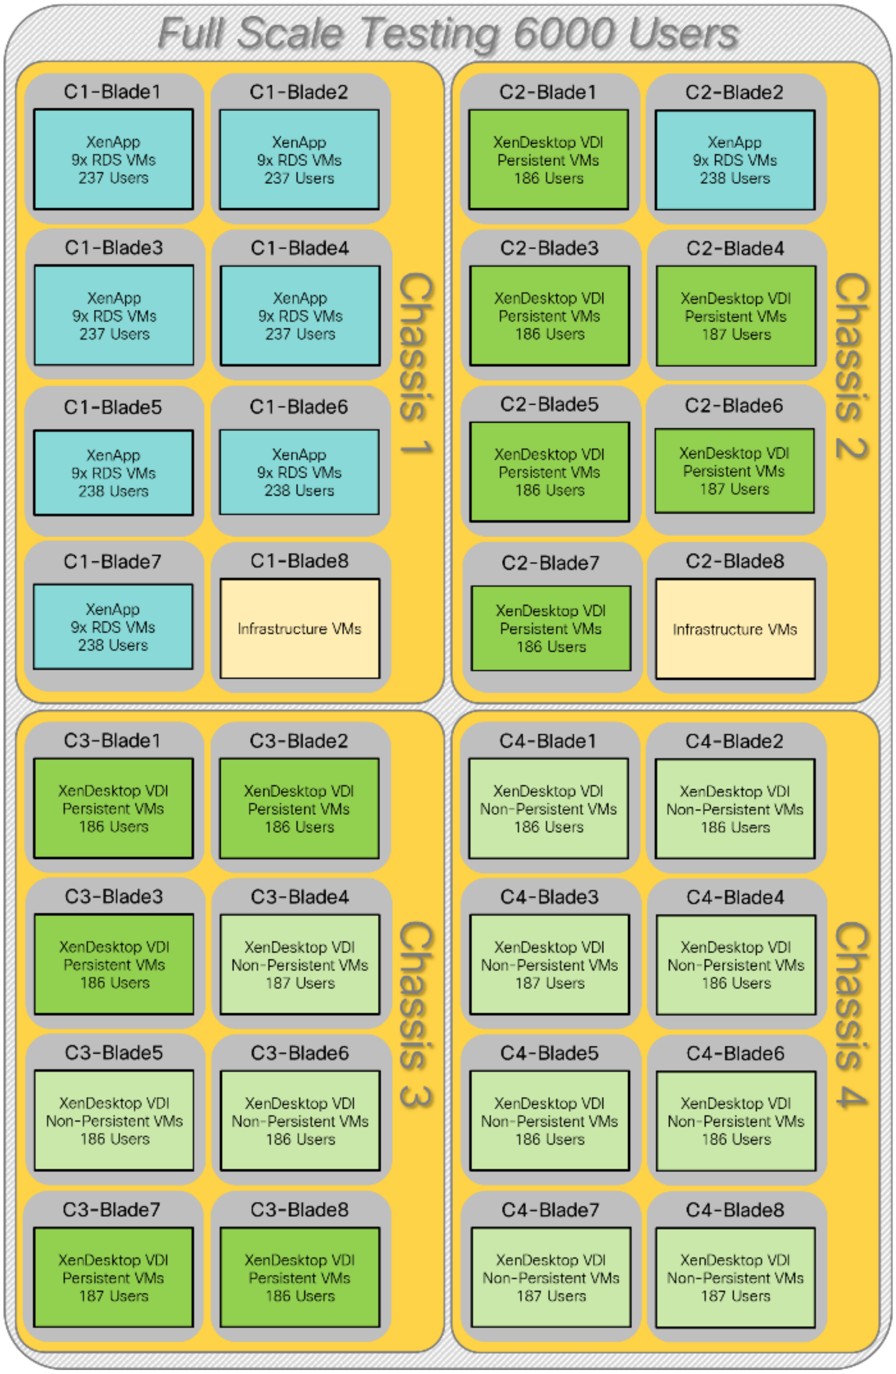

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332 - 16 UP Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.60-GHz 10-core processors, 128GB 2133MHz RAM for all host blades
- 30 (RDS/VDI Host) Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM for all host blades

- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Storage Switches
- 1 NetApp AFF A300 storage system (2x storage controllers- Active/Active High Availability pair) with 1x DS224C disk shelves, 24x 3.8TB SSD- 65TB usable / 130TB effective (2:1 efficiency)

Software components:

- Cisco UCS firmware 3.2(3d)
- VMware ESXi 6.5 Update 1 for host blades
- Citrix XenApp/XenDesktop 7.15 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.15
- Citrix User Profile Manager
- Microsoft SQL Server 2016
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB vdisk (master)
- Microsoft Windows Server 2016, 9vCPU, 24GB RAM, 40 GB vdisk (master)
- Microsoft Office 2016
- Login VSI 4.1.25 Knowledge Worker Workload (Benchmark Mode)

Testing 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 Citrix XenApp and XenDesktop Hosted Virtual Desktop and RDS Hosted Shared models under test.

Test metrics were gathered from the hypervisor, 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.

Testing 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 XenDesktop Administrator and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

All VMware ESXi VDI host blades to be tested were restarted prior to each test cycle.

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 0.5 percent 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. Complete the following steps:

1. Time 0:00:00 Start PerfMon/Esxtop/XenServer Logging on the following systems:
    a. Infrastructure and VDI Host Blades used in the test run
    b. SCVMM/vCenter used in the test run
    c. All Infrastructure VMs 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 XenDesktop Studio or View Connection server.

    The boot rate should be around 10-12 VMs per minute per server.

4. Time 0:06 First machines boot.
5. Time 0:30 Single Server or Scale target number of desktop VMs 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 VMs desktops registered on XD Studio or 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 30-40 minute rest period is sufficient.

8. Time 1:35 Start Login VSI 4.1.25 Office Worker Benchmark Mode Test, setting auto-logoff time at 900 seconds, with Single Server or Scale target number of desktop VMs utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).
9. Time 2:23 Single Server or Scale target number of desktop VMs desktops launched (48 minute benchmark launch rate).
10. Time 2:25 All launched sessions must become active.

All sessions launched must become active for a valid test run within this window.

11. Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above).
12. Time 2:55 All active sessions logged off.

13. Time 2:57 All logging terminated; Test complete.

14. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shut-down all Windows machines.

15. Time 3:30 Reboot all hypervisor hosts.

16. Time 3:45 Ready for the new test sequence.

Success Criteria

Our "pass" criteria for this testing 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 Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The Citrix Desktop Studio be monitored throughout the steady state to make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.

Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with the proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing. FlexPod Data Center with Cisco UCS and Citrix XenApp/XenDesktop 7.15 on VMware ESXi 6.5 Update 1 Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenApp Hosted Shared Desktop (RDS) and Citrix XenDesktop Hosted Virtual Desktop (VDI) randomly assigned, non-persistent with Citrix Provisioning Services 7.15 and Citrix XenDesktop Hosted Virtual Desktop (VDI) statically assigned, persistent full-clones models using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2016 sessions on Cisco UCS B200 M5 Blade Servers using a NetApp AFF300 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of Citrix products with VMware vSphere.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, linear scalability.

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 SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer, and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out 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, etc. 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 66 Sample of a VSI Max Response Time Graph, Representing a Normal Test

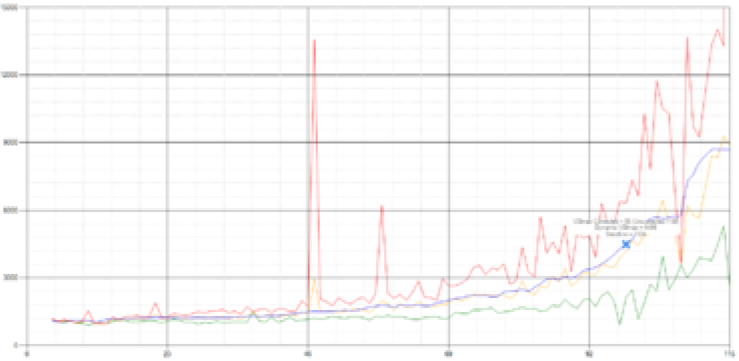

Figure 67 Sample of a VSI Test Response Time Graph with a Performance Issue

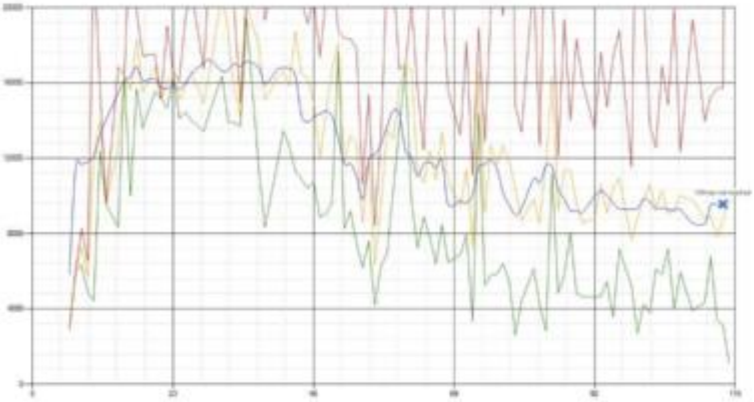

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. These response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSImax models, this weighting much better represents system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times is applied.

The following actions are part of the VSImax v4.1 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, we also created a new method to calculate the basephase of an environment. With the new workloads (Taskworker, Powerworker, etc.) enabling 'basephase' for a more reliable baseline has become obsolete. The calculation is explained below. In total the 15 lowest VSI response time samples are taken from the entire test, the lowest 2 samples are removed. and the 13 remaining samples are averaged. The result is the Baseline. To summarize:

- Take the lowest 15 samples of the complete test
- From those 15 samples remove the lowest 2
- Average the 13 results that are left is the baseline

The VSImax average response time in Login VSI 4.1.x is calculated on the amount of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40 percent of the amount of "active" sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement are 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 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 amount 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 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. This methodology gives much better insight into system performance and scales to extremely large systems.

Single-Server Recommended Maximum Workload

For both the Citrix XenDesktop 7.15 Hosted Virtual Desktop and Citrix XenApp 7.15 RDS Hosted Shared Desktop 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 insure 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 45 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 30-40 minute, <60 min)
Logon	The Login VSI phase of the test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files (typically for 15-minute duration)
Logoff	Sessions finish executing the Login VSI workload and logoff

Test Results

Single-Server Recommended Maximum Workload Testing

This section shows the key performance metrics that were captured on the Cisco UCS host blades during the single server testing to determine the Recommended Maximum Workload per host server. The single server testing comprised of three tests: 270 RDS sessions, 205 VDI Non-Persistent sessions, and 205 VDI Persistent sessions.

Single-Server Recommended Maximum Workload for RDS with 270 Users

Figure 68 Single Server Recommended Maximum Workload for RDS with 270 Users

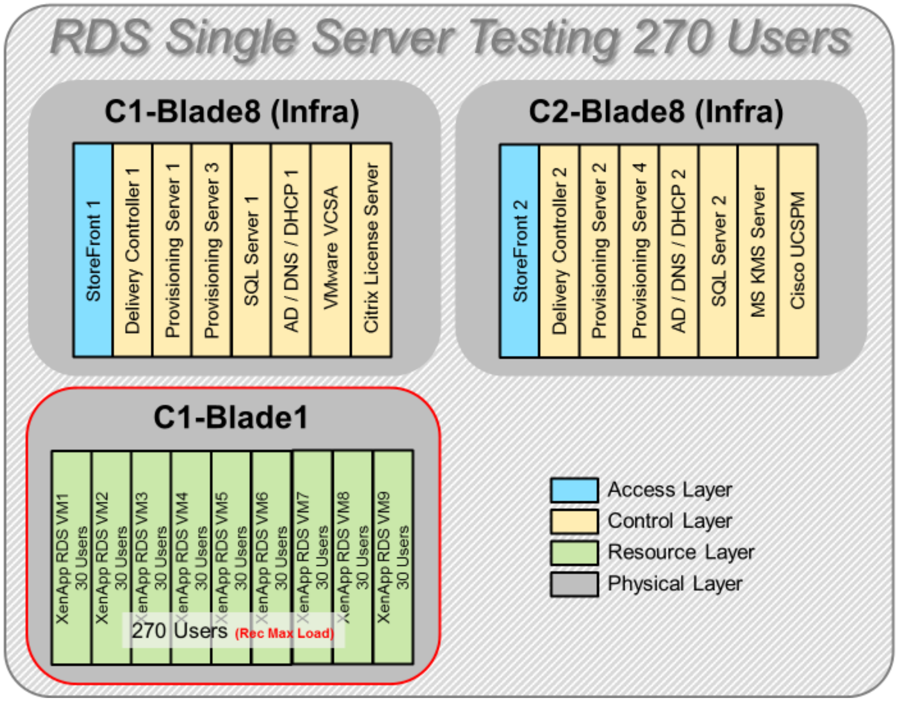

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6140 processors, 768GB 2666MHz RAM is 270 Server 2016 Hosted Shared Desktops. Each dedicated blade server ran 9 Server 2016 Virtual Machines. Each virtual server was configured with 9 vCPUs and 24GB RAM.

Figure 69 Single Server Recommended Maximum Workload | XenApp 7.15 RDS | VSI Score

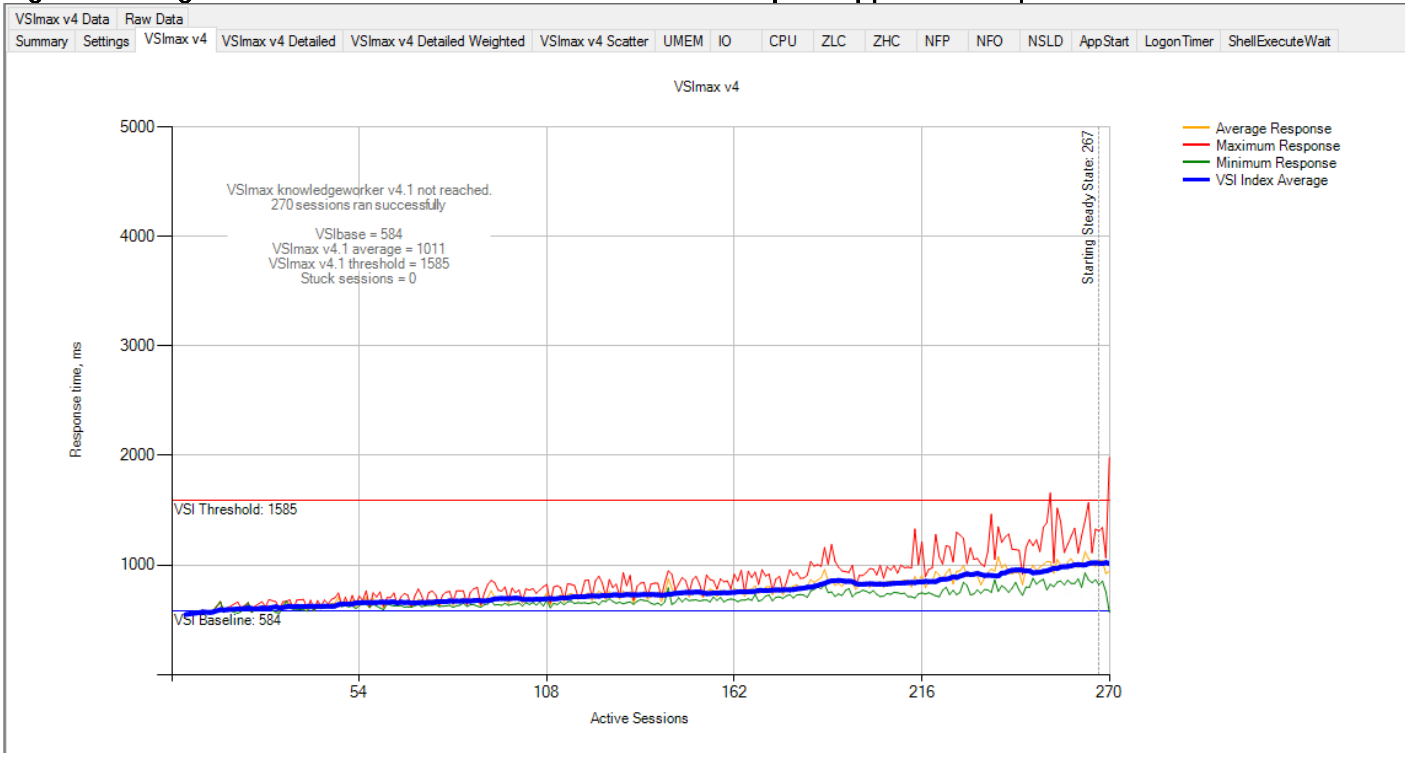

VSImax knowledgeworker v4.1 not reached.
270 sessions ran successfully

VSIbase = 584
VSImax v4.1 average = 1011
VSImax v4.1 threshold = 1585
Stuck sessions = 0

Figure 70 Single Server Recommended Maximum Workload | XenApp 7.15 RDS | VSI Repeatability

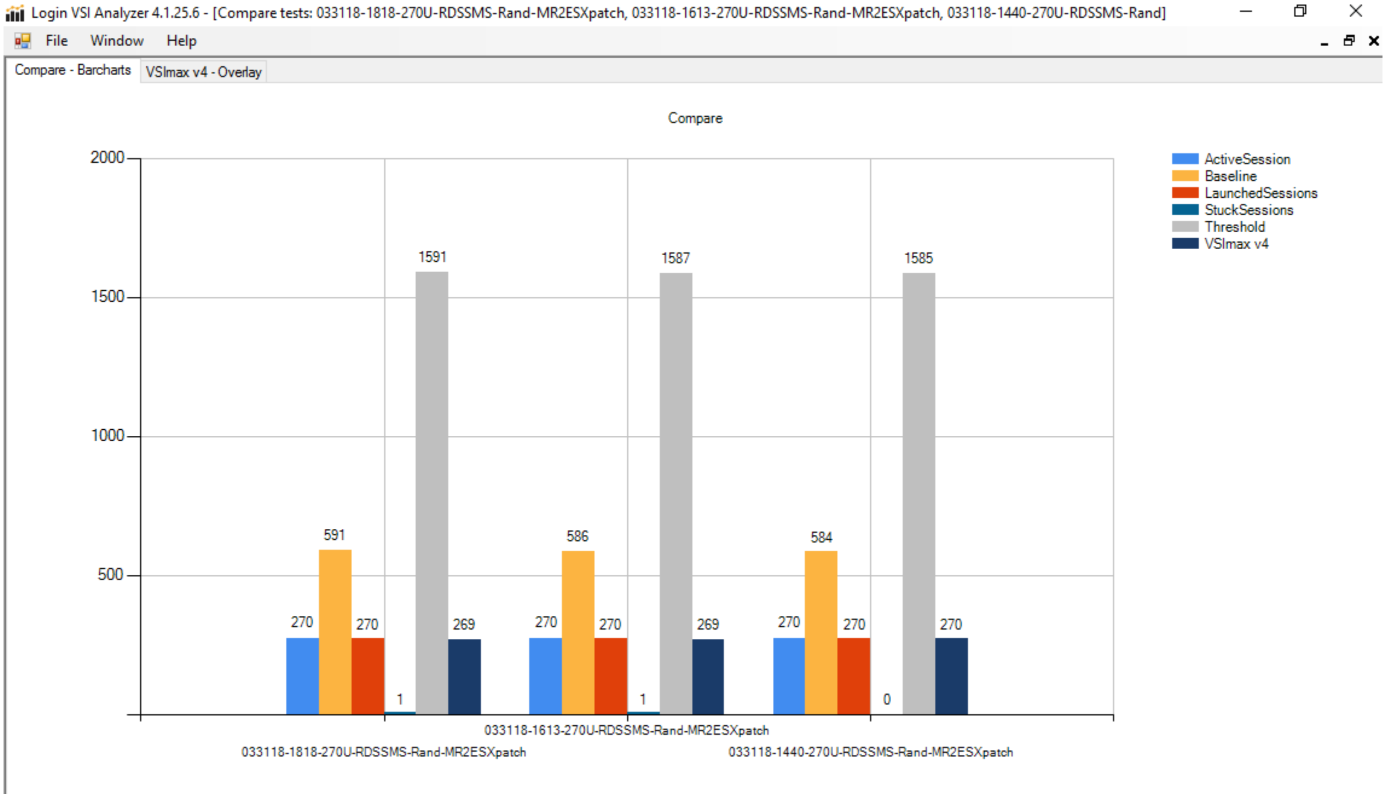

Login VSI Analyzer 4.1.25.6 - [Compare tests: 033118-1818-270U-RDSSMS-Rand-MR2ESXpatch, 033118-1613-270U-RDSSMS-Rand-MR2ESXpatch, 033118-1440-270U-RDSSMS-Rand]

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

Figure 71 Single Server Recommended Maximum Workload | XenApp 7.15RDS | Host CPU Utilization

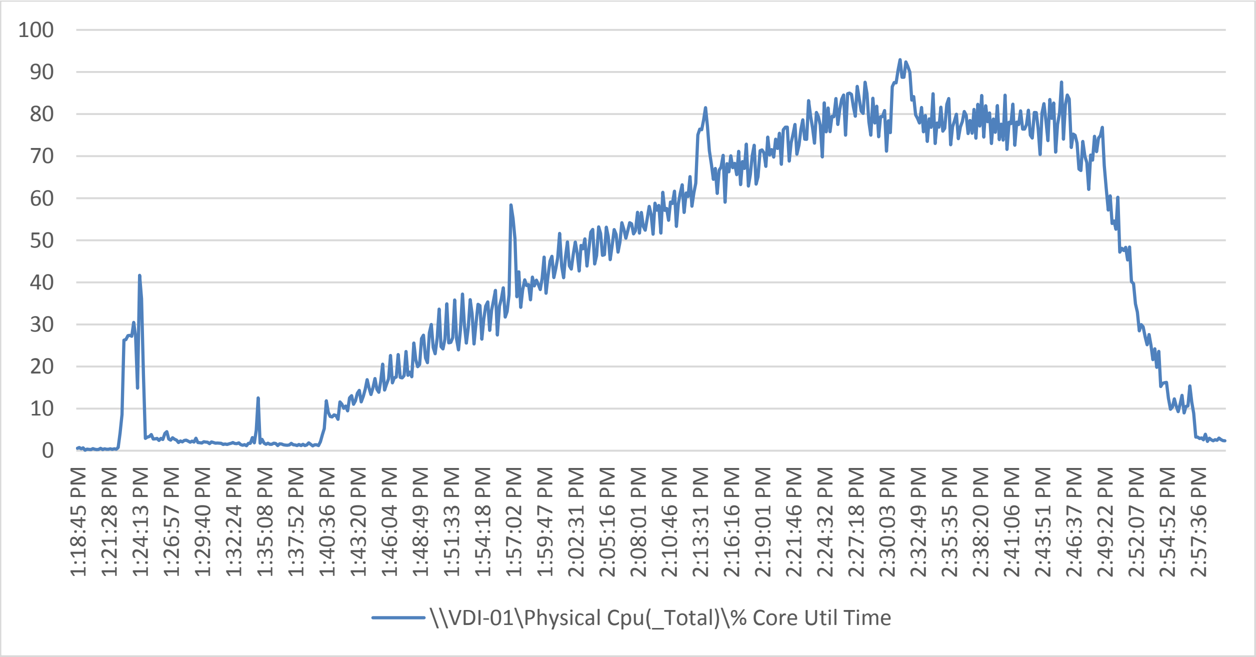

Figure 72 Single Server Recommended Maximum Workload | XenApp 7.15RDS | Host Memory Utilization

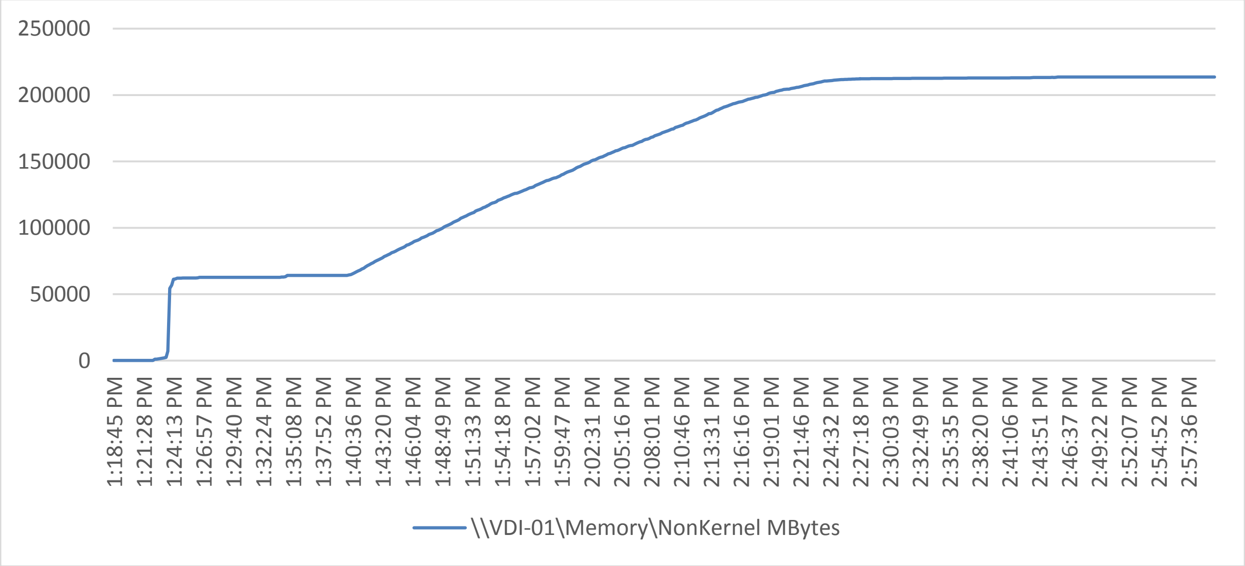

Figure 73 Single Server | XenApp 7.15RDS | Host Network Utilization

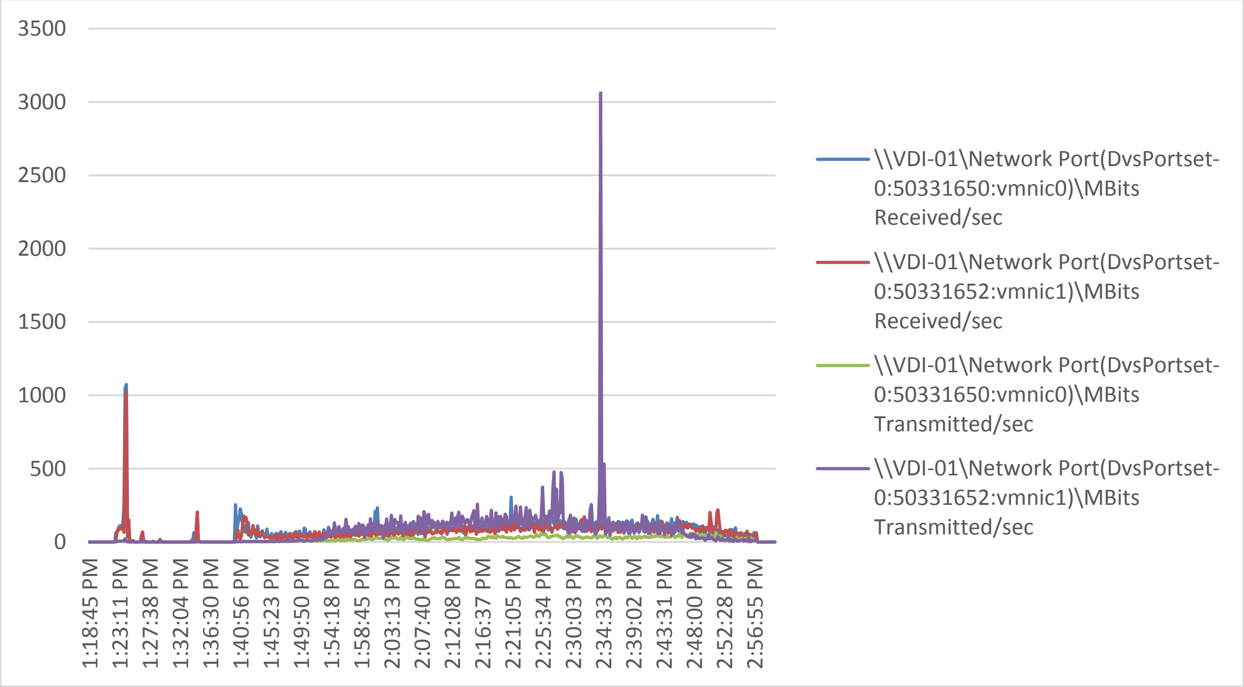

Single-Server Recommended Maximum Workload for VDI Non-Persistent with 205 Users

Figure 74 Single Server Recommended Maximum Workload for VDI Non-Persistent with 205 Users

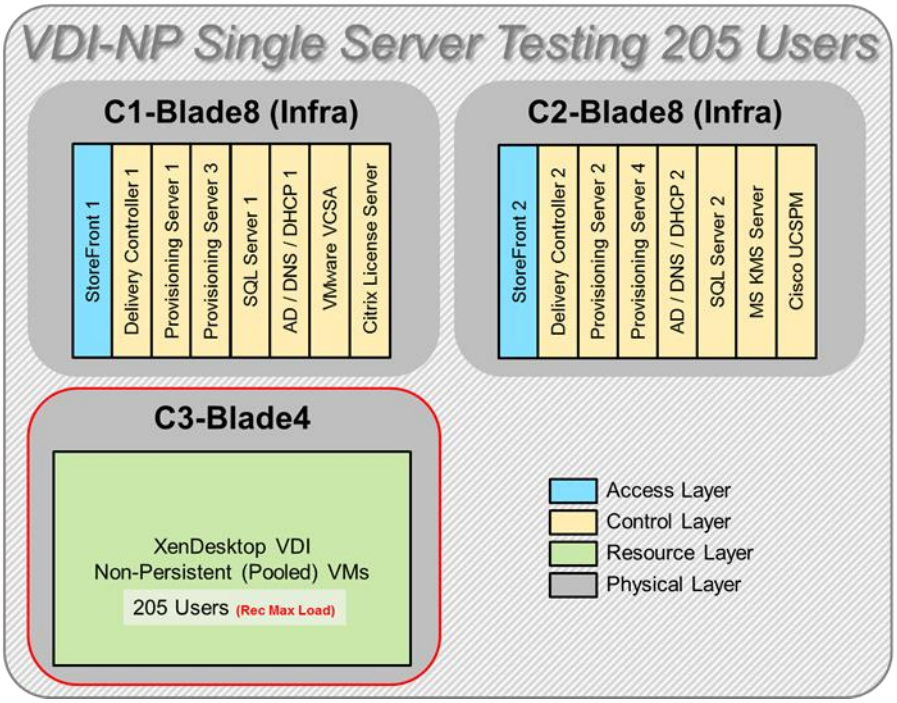

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6140 processors, 768GB 2666MHz RAM is 220 Windows 10 64-bit virtual machines with 2 vCPU and 2GB RAM. Login VSI and blade performance data is as follows.

Figure 75 Single Server | XenDesktop 7.15 VDI-NP | VSI Score

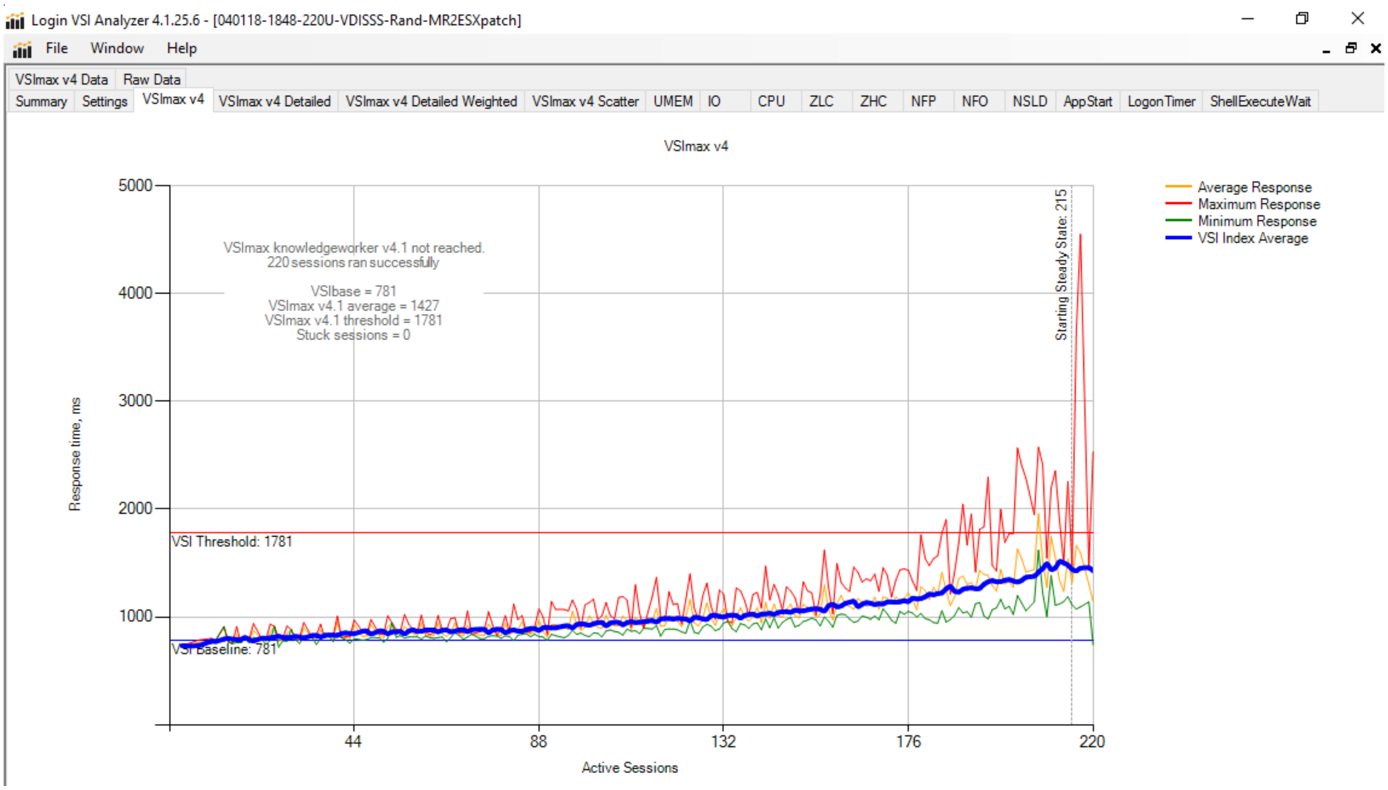

VSImax knowledgeworker v4.1 not reached.
220 sessions ran successfully

VSIbase = 781
VSImax v4.1 average = 1427
VSImax v4.1 threshold = 1781
Stuck sessions = 0

Figure 76 Single Server | XenDesktop 7.15 VDI-NP | VSI Repeatability

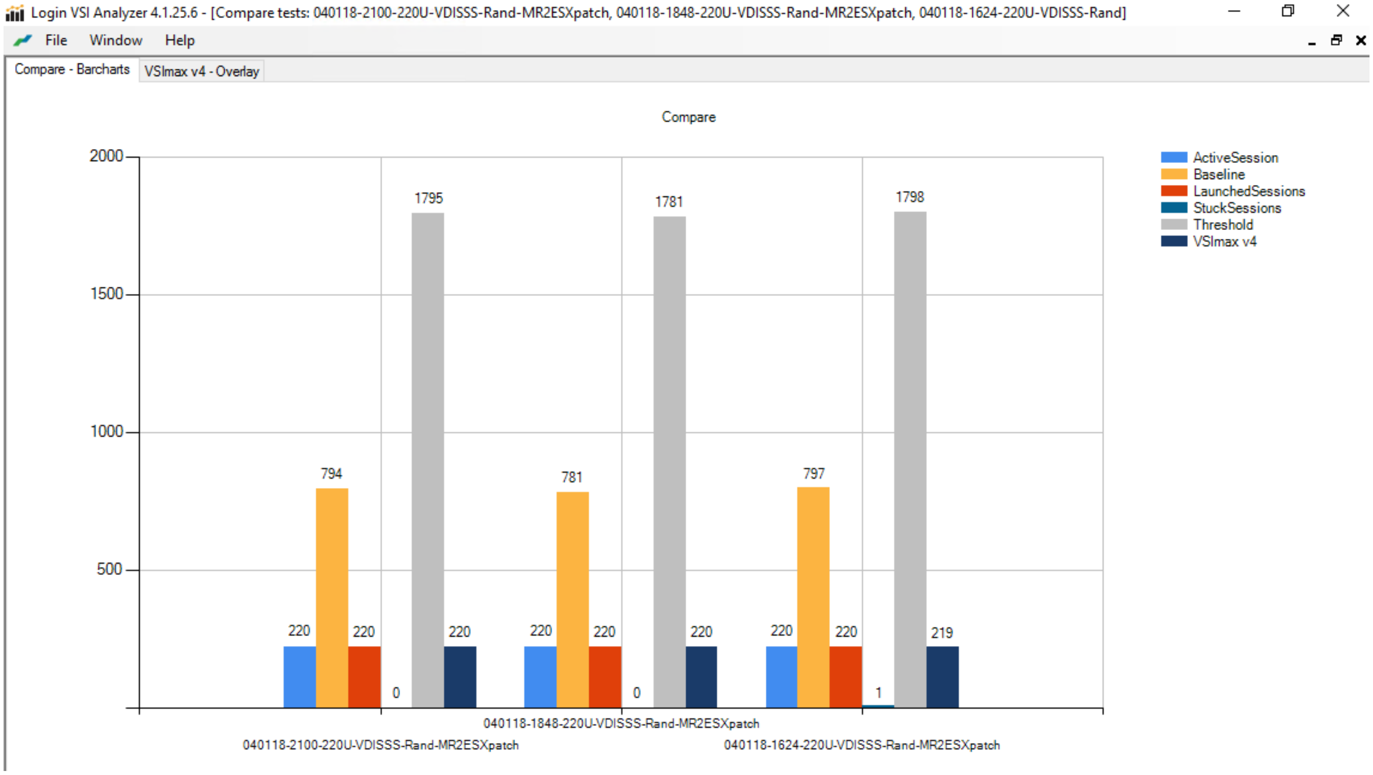
Login VSI Analyzer 4.1.25.6 - [Compare tests: 040118-2100-220U-VDISSS-Rand-MR2ESXpatch, 040118-1848-220U-VDISSS-Rand-MR2ESXpatch, 040118-1624-220U-VDISSS-Rand]
File
Window
Help
Compare - Barcharts
VSImax v4 - Overlay
Compare
2000
1500
1000
500
794
220
220
0
1795
220
781
220
220
0
1781
220
797
220
220
1
1798
219
040118-1848-220U-VDISSS-Rand-MR2ESXpatch
040118-2100-220U-VDISSS-Rand-MR2ESXpatch
040118-1624-220U-VDISSS-Rand-MR2ESXpatch
ActiveSession
Baseline
LaunchedSessions
StuckSessions
Threshold
VSImax v4

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

Figure 77 Single Server | XenDesktop 7.15 VDI-NP | Host CPU Utilization

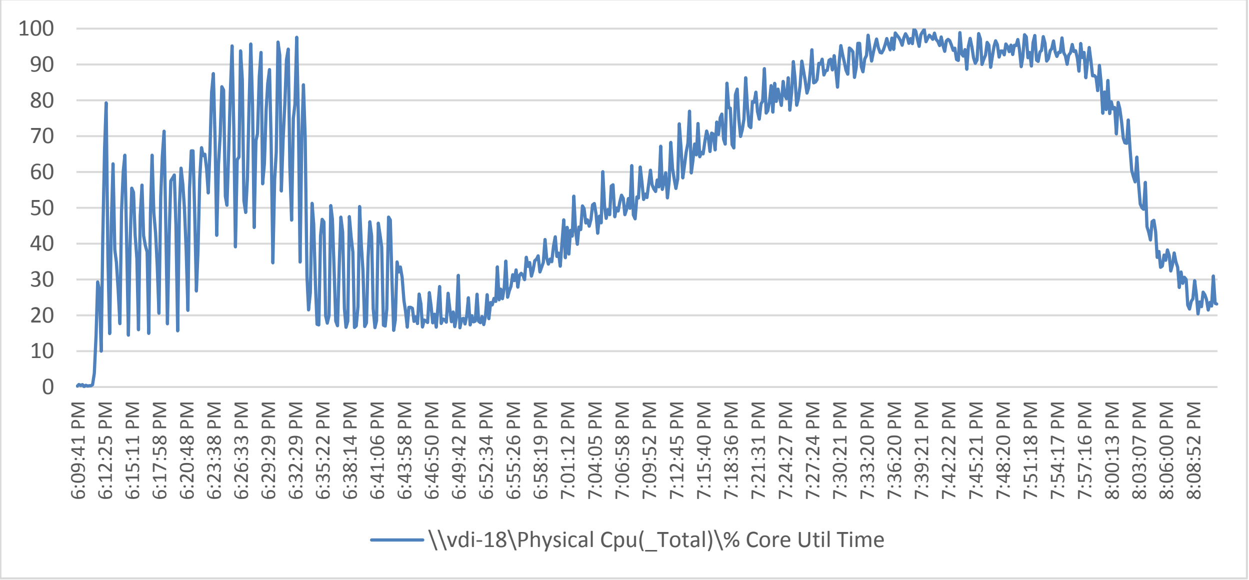

Figure 78 Single Server | XenDesktop 7.15 VDI-NP | Host Memory Utilization

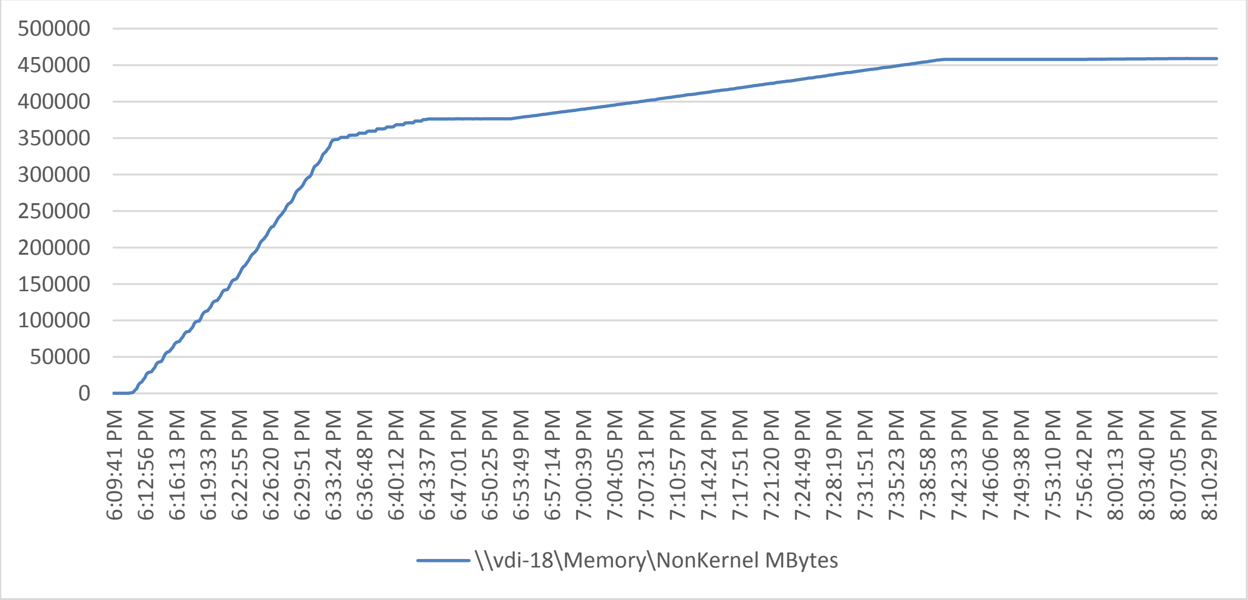

Figure 79 Single Server | XenDesktop 7.15 VDI-NP | Host Network Utilization

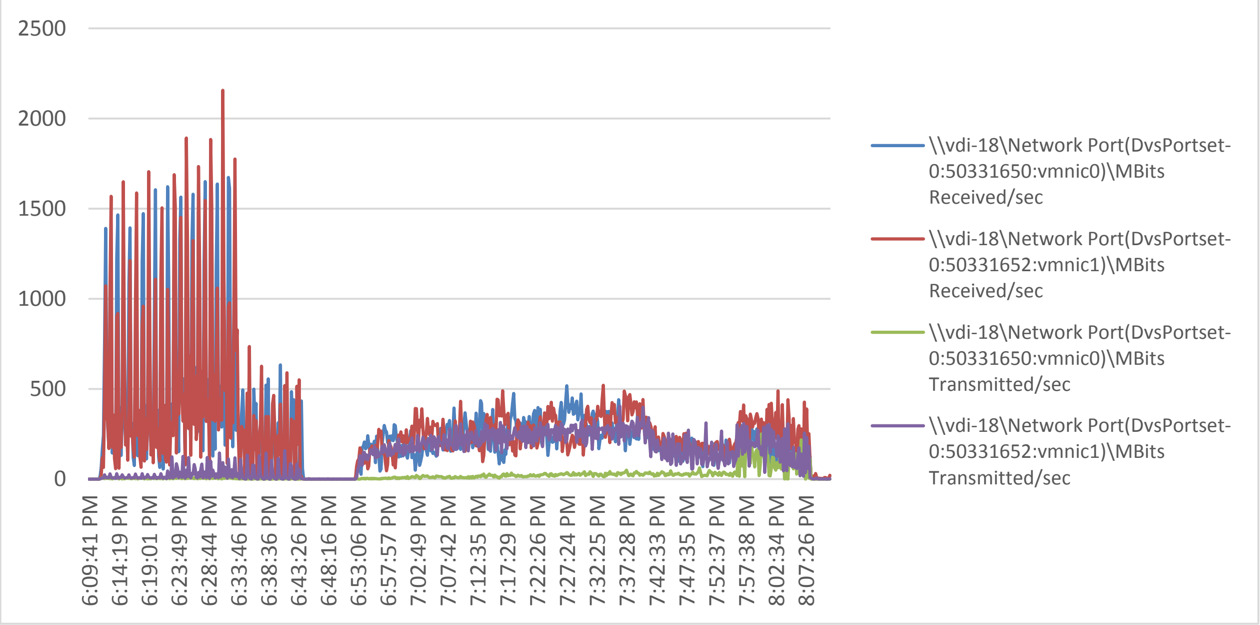

Single-Server Recommended Maximum Workload for VDI Persistent with 205 Users

Figure 80 Single Server Recommended Maximum Workload for VDI Persistent with 205 Users

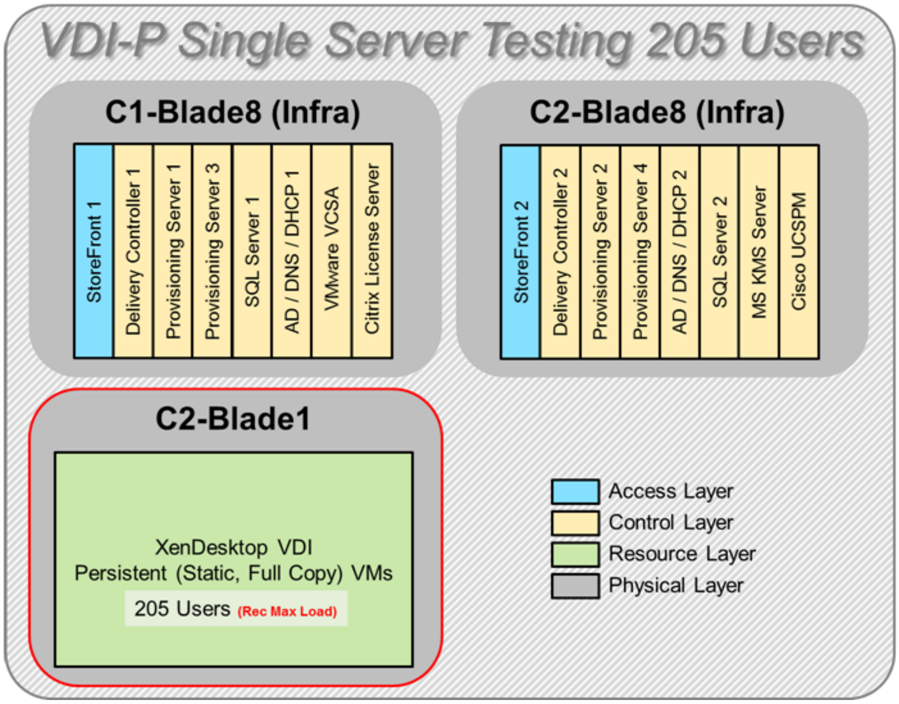

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual Intel Xeon Gold 6140 processors, 768GB 2666MHz RAM is 205 Windows 10 64-bit virtual machines with 2 vCPU and 2GB RAM. Login VSI and blade performance data is as follows.

Figure 81 Single Server | XenDesktop 7.15 VDI-P | VSI Score

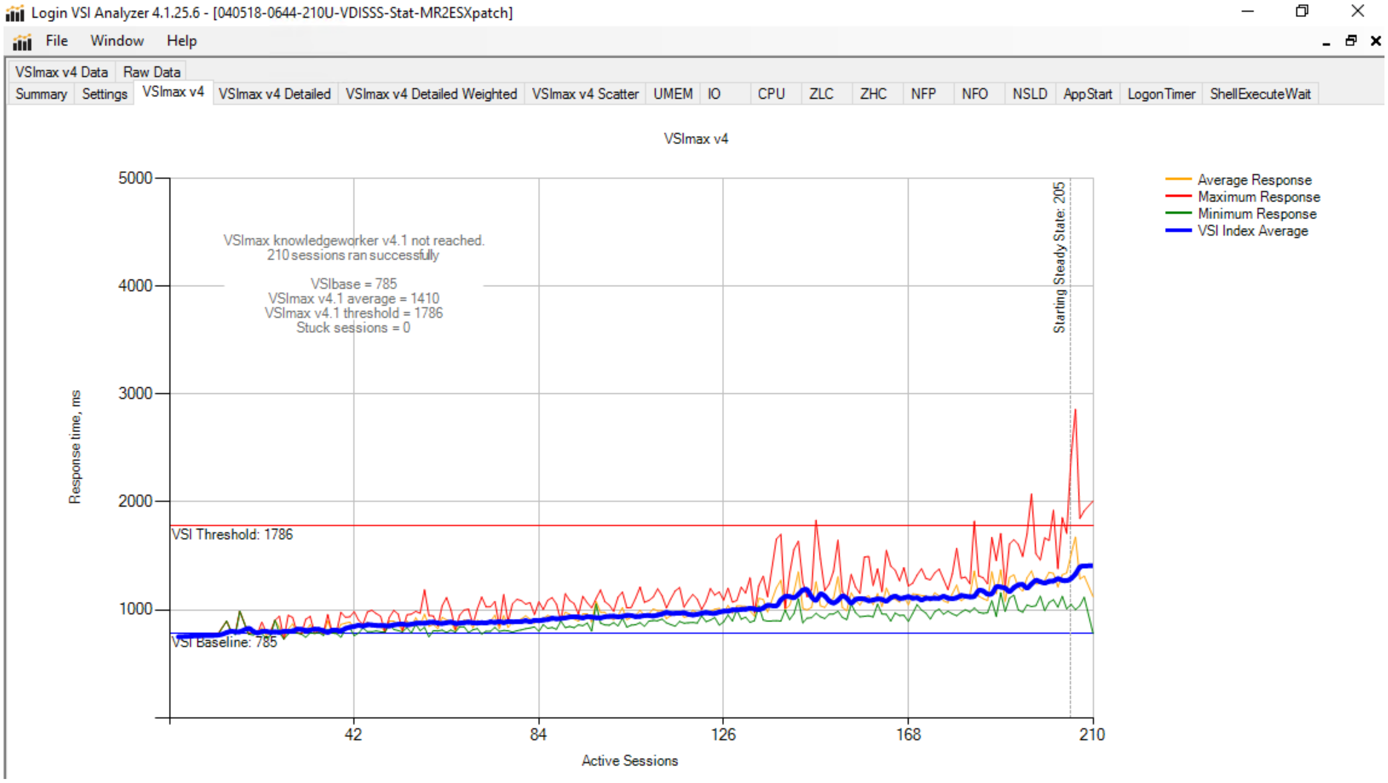

VSImax knowledgeworker v4.1 not reached.
210 sessions ran successfully

VSIbase = 785
VSImax v4.1 average = 1410
VSImax v4.1 threshold = 1786
Stuck sessions = 0

Figure 82 Single Server | XenDesktop 7.15 VDI-P | VSI Repeatability

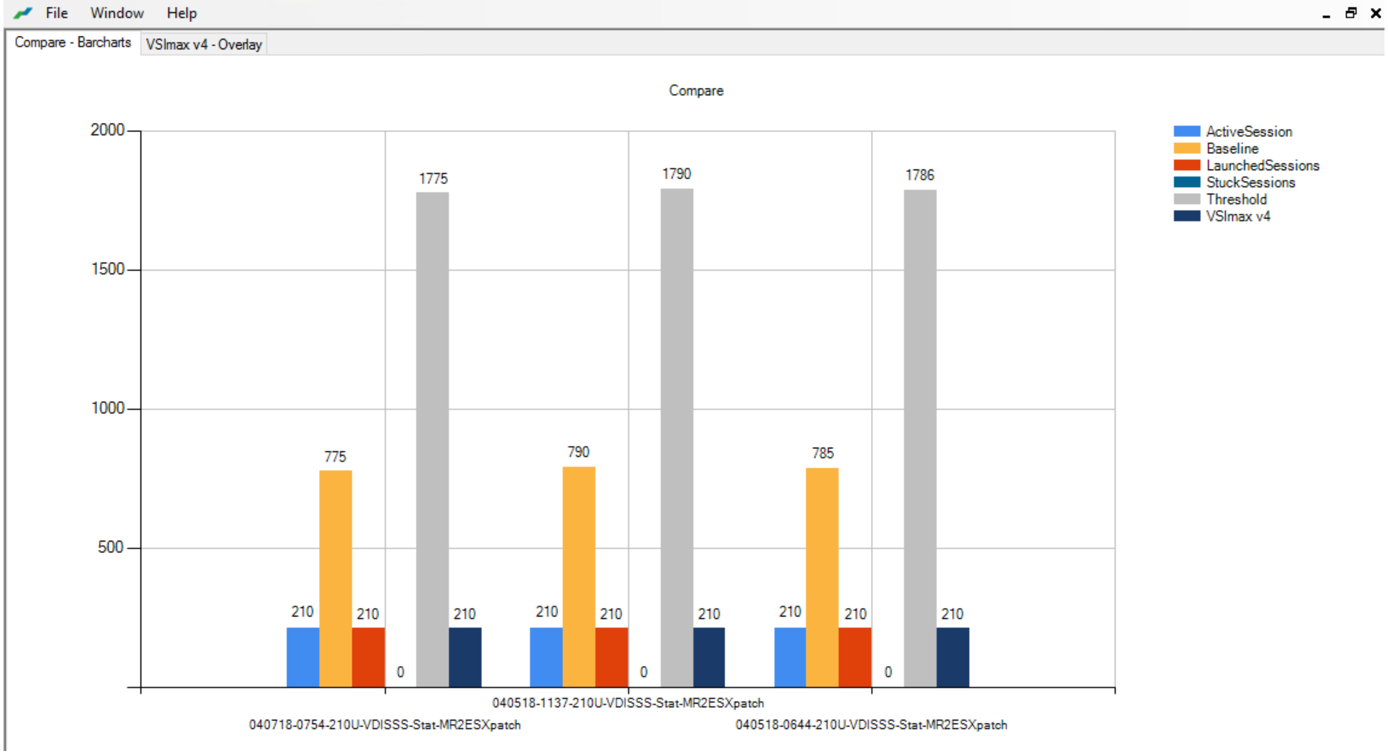

Performance data for the server running the workload is as follows:

Figure 83 Single Server | XenDesktop 7.15 VDI-P | Host CPU Utilization

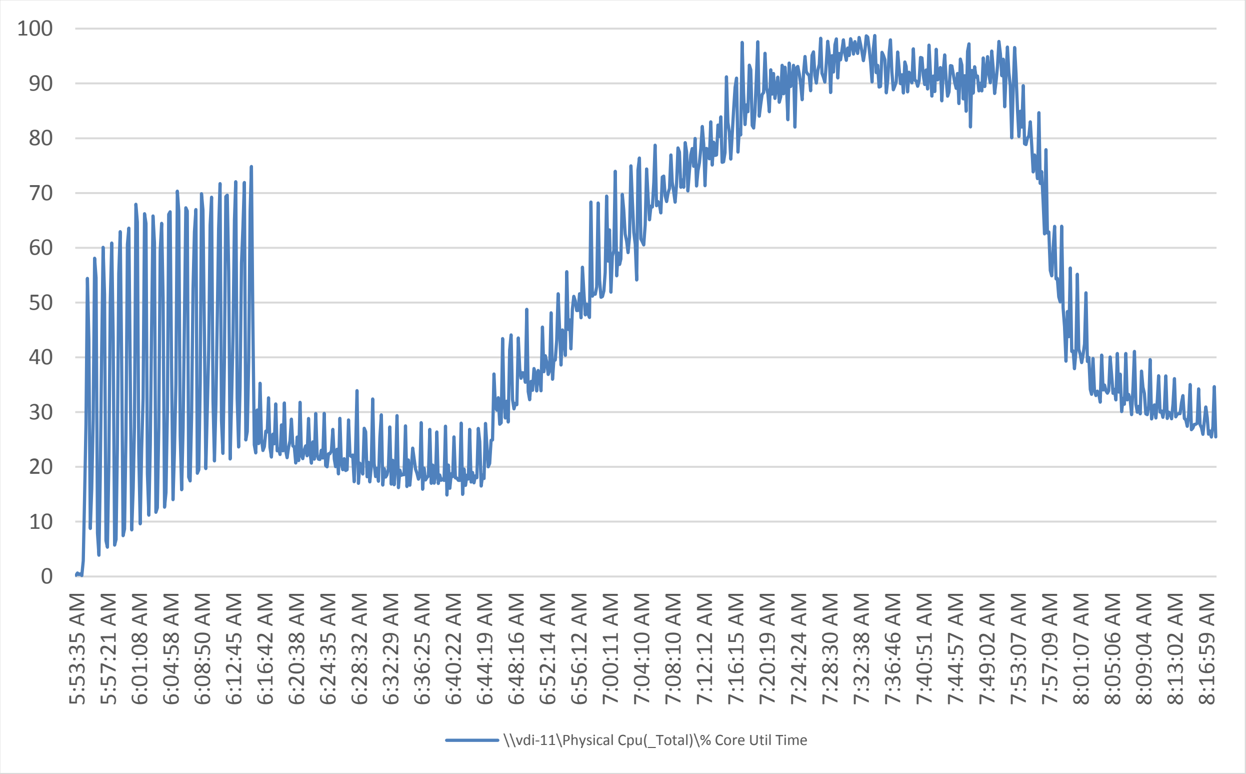

Figure 84 Single Server | XenDesktop 7.15 VDI-P | Host Memory Utilization

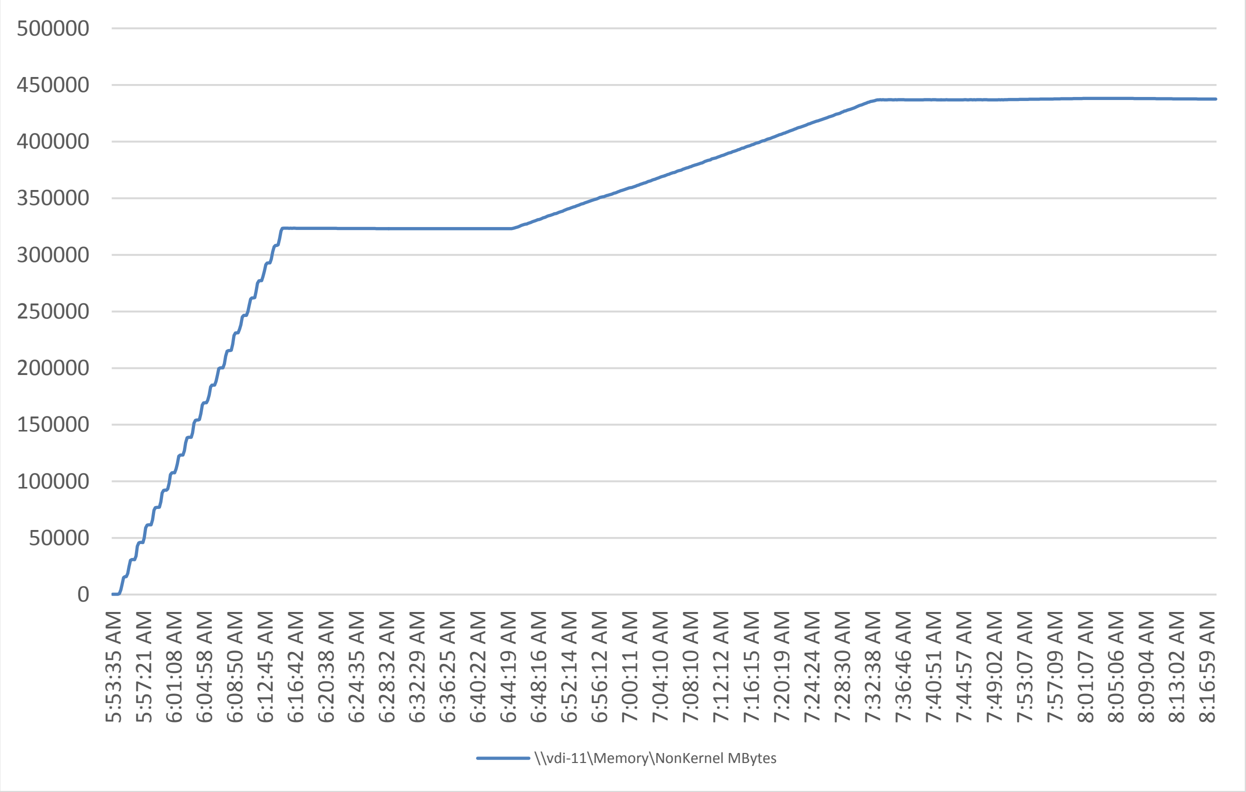

Figure 85 Single Server | XenDesktop 7.15 VDI-P | Host Network Utilization

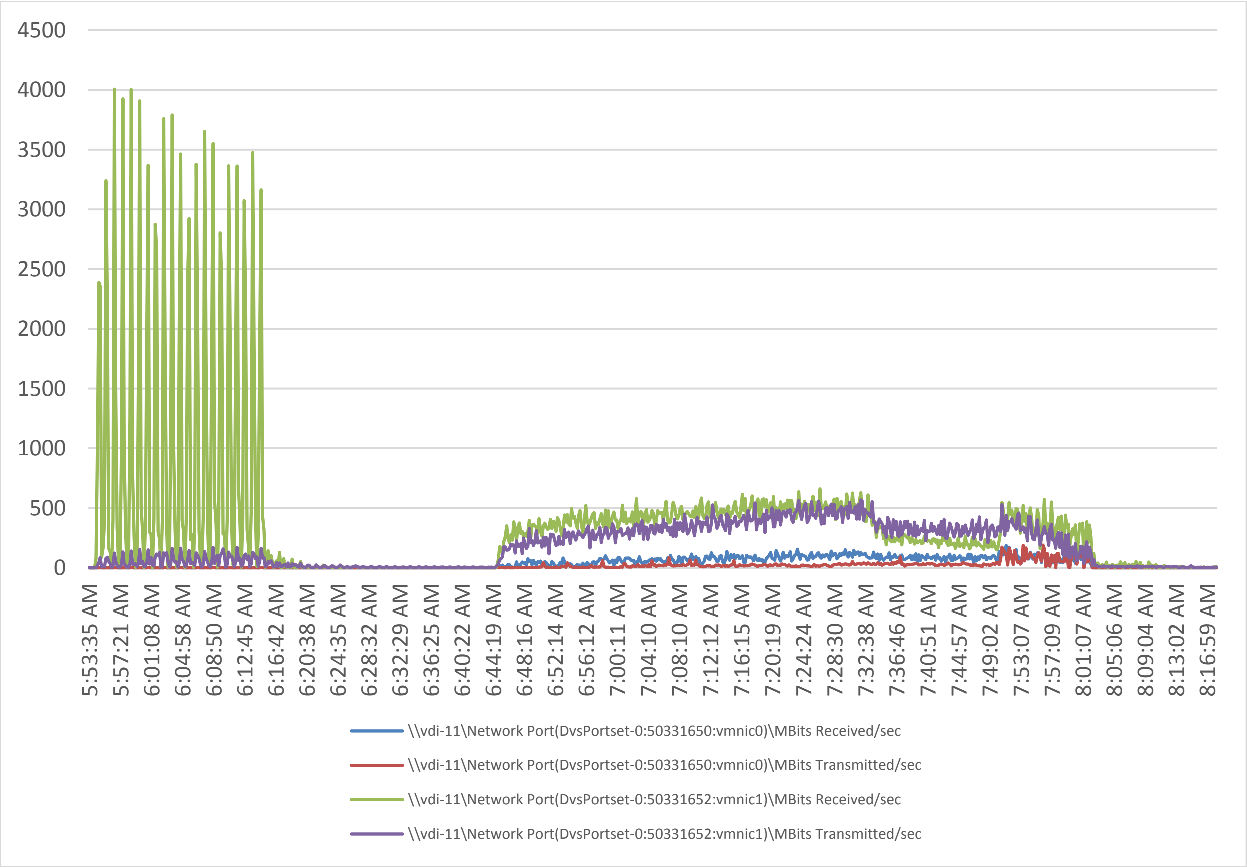

Cluster Workload Testing with 1900 RDS Users

This section describes the key performance metrics that were captured on the Cisco UCS, NetApp storage, and Infrastructure VMs during the non-persistent desktop testing. The cluster testing was comprised of 1900 RDS sessions using 8 workload blades.

Figure 86 RDS Cluster Testing with 1900 Users

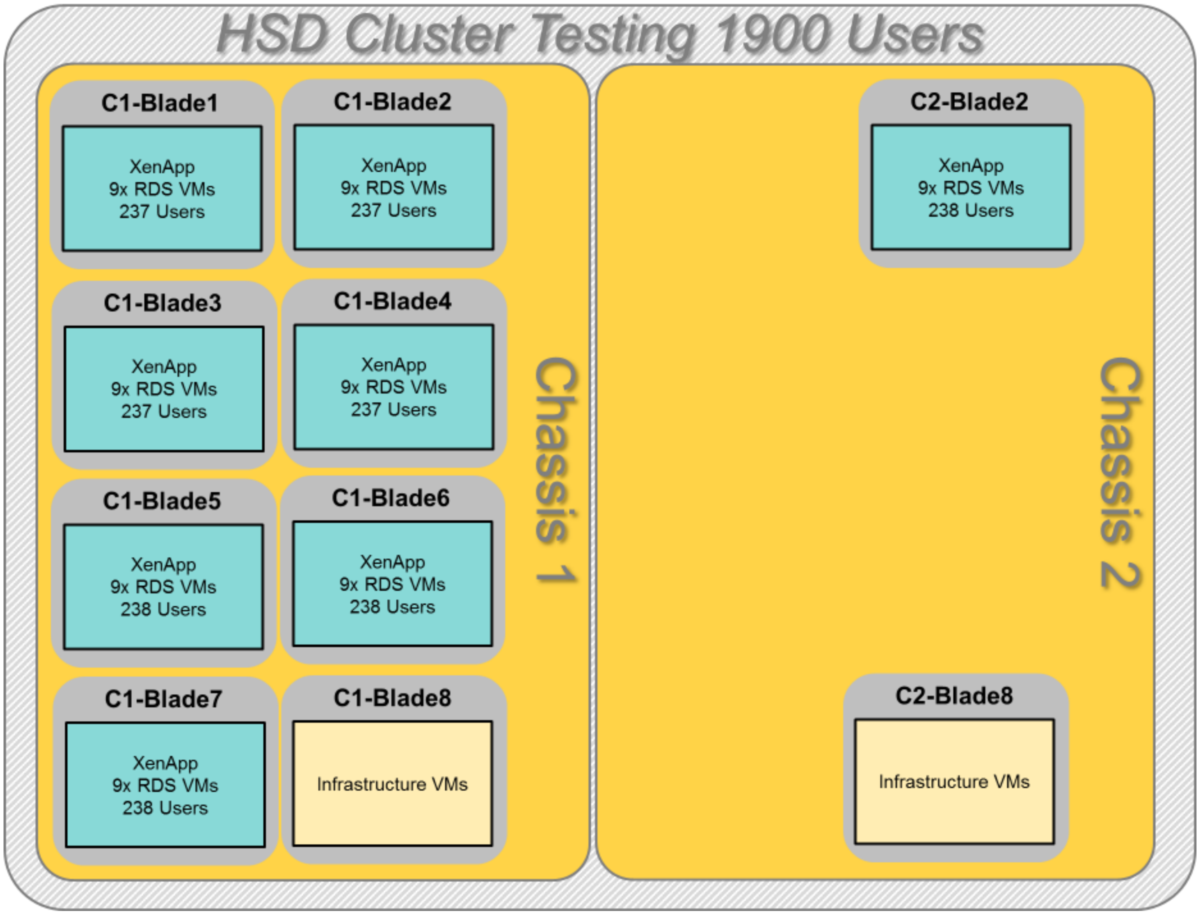

The workload for the test is 1900 RDS users. To achieve the target, sessions were launched against all workload clusters concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results.

Figure 87 Eight Node Cluster | 1900 RDS Users | VSI Score

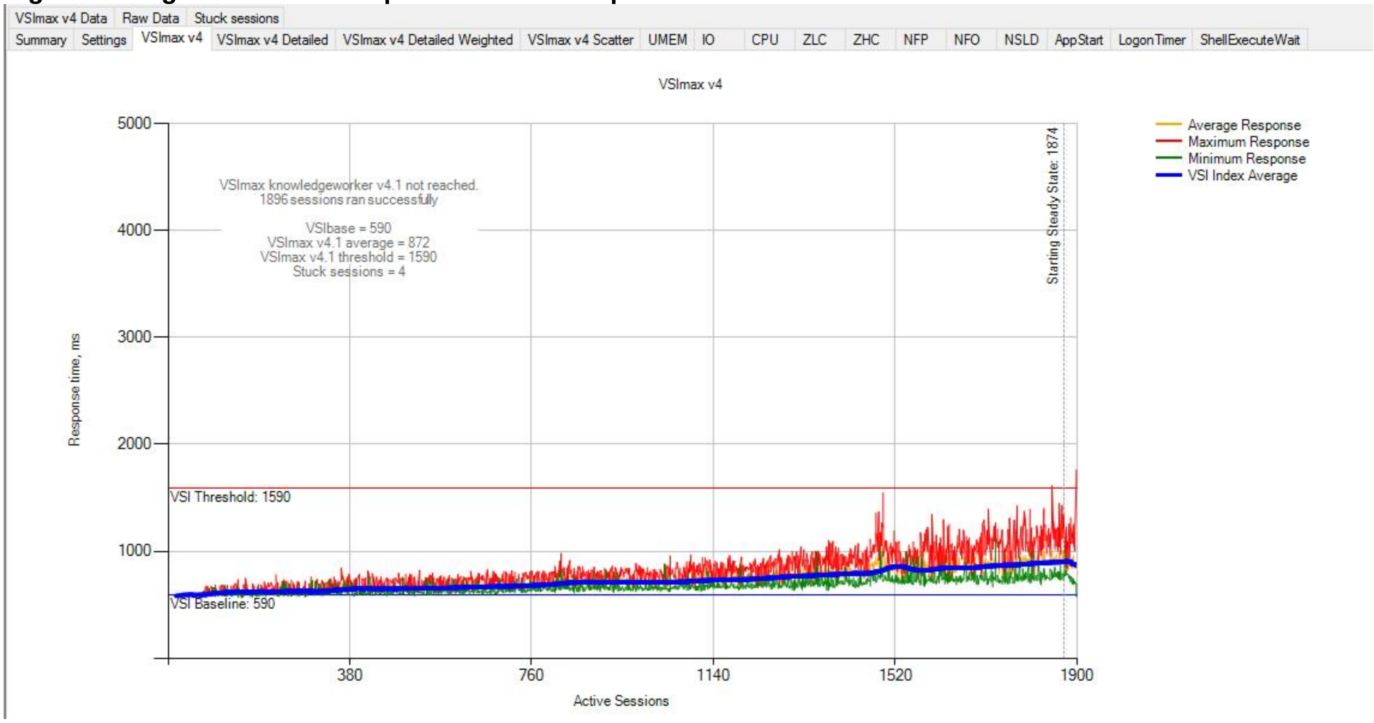

VSImax knowledgeworker v4.1 not reached.
1896 sessions ran successfully

VSIbase = 590
VSImax v4.1 average = 872
VSImax v4.1 threshold = 1590
Stuck sessions = 4

Figure 88 Eight Node Cluster | 1900 RDS Users | VSI Repeatability

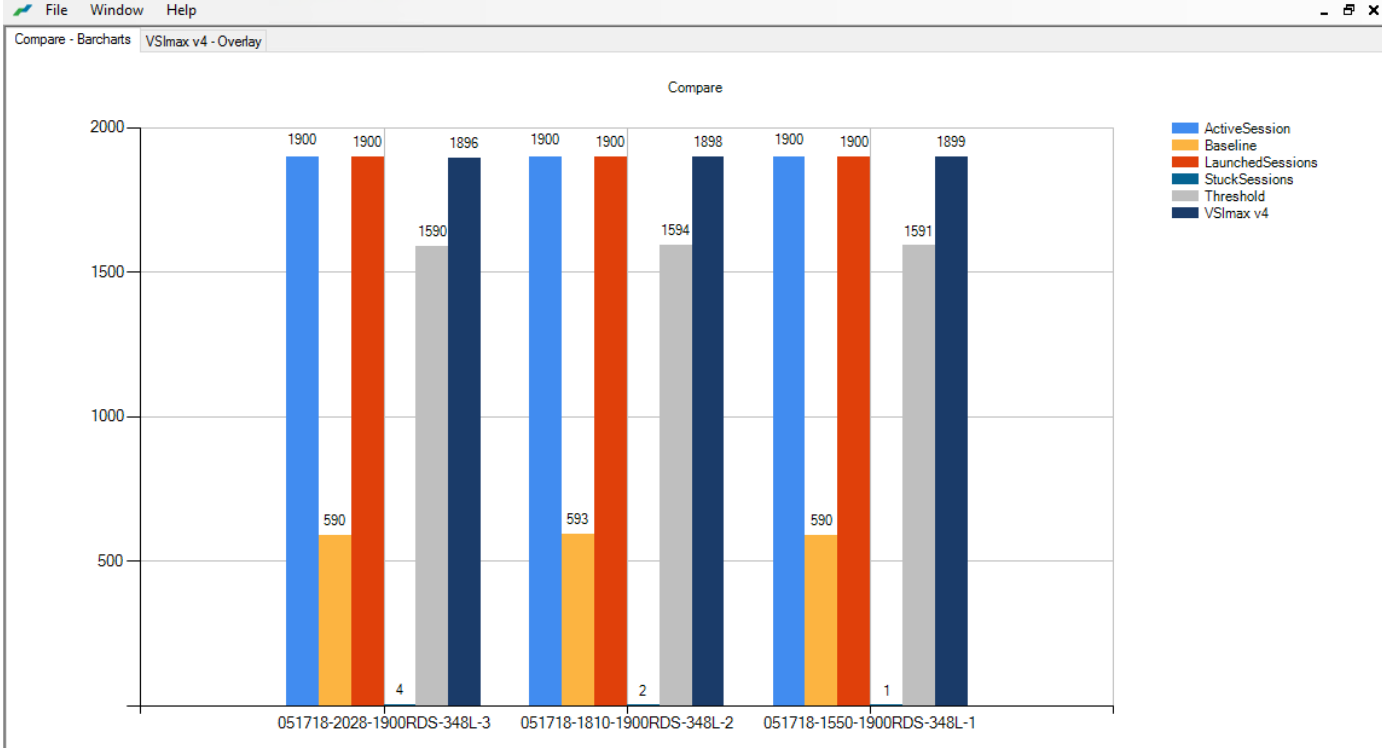

Figure 89 Cluster | 1900 RDS Users | 8 RDS Hosts | Host CPU Utilization

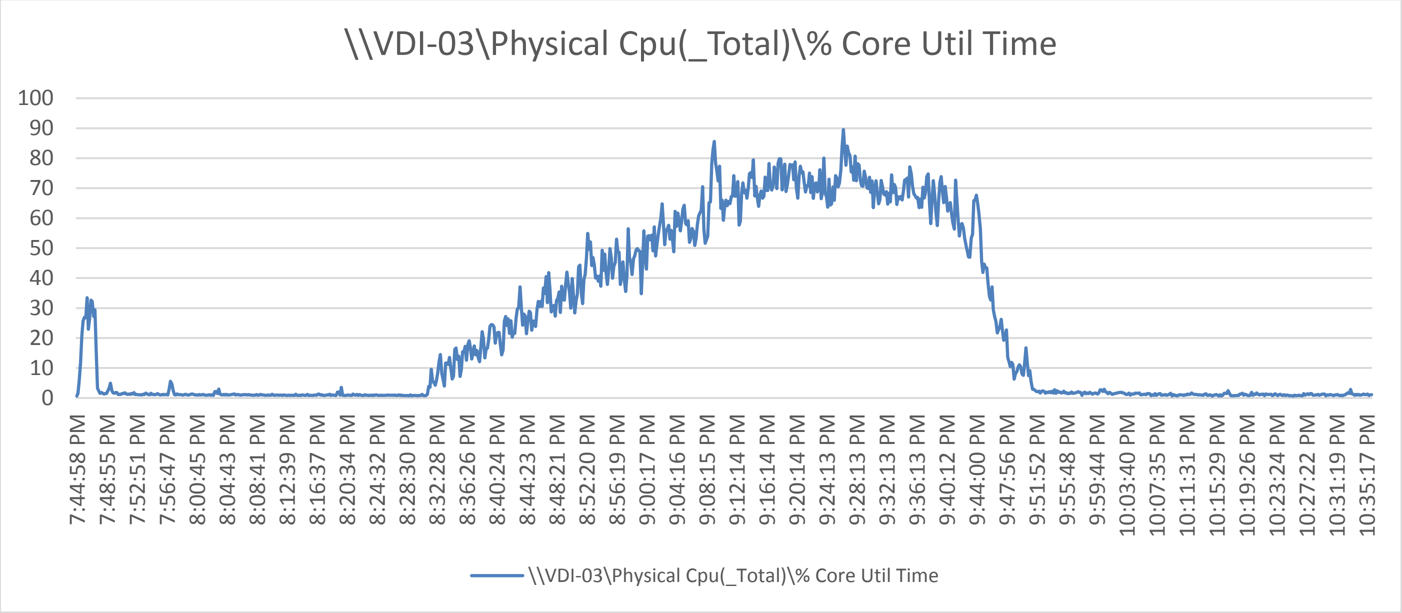

Figure 90 Cluster | 1900 RDS Users | 8 RDS Hosts | Host Memory Utilization

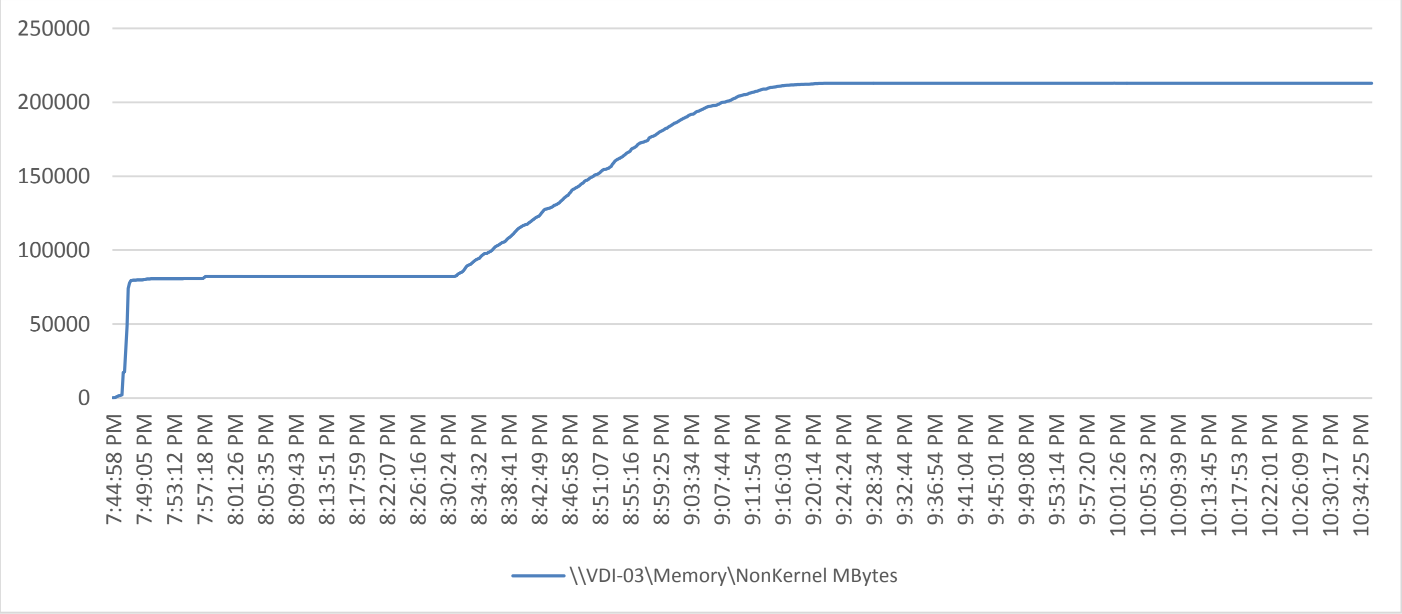

Figure 91 Cluster | 1900 RDS Users | RDS Hosts | Host System Uplink Network Utilization

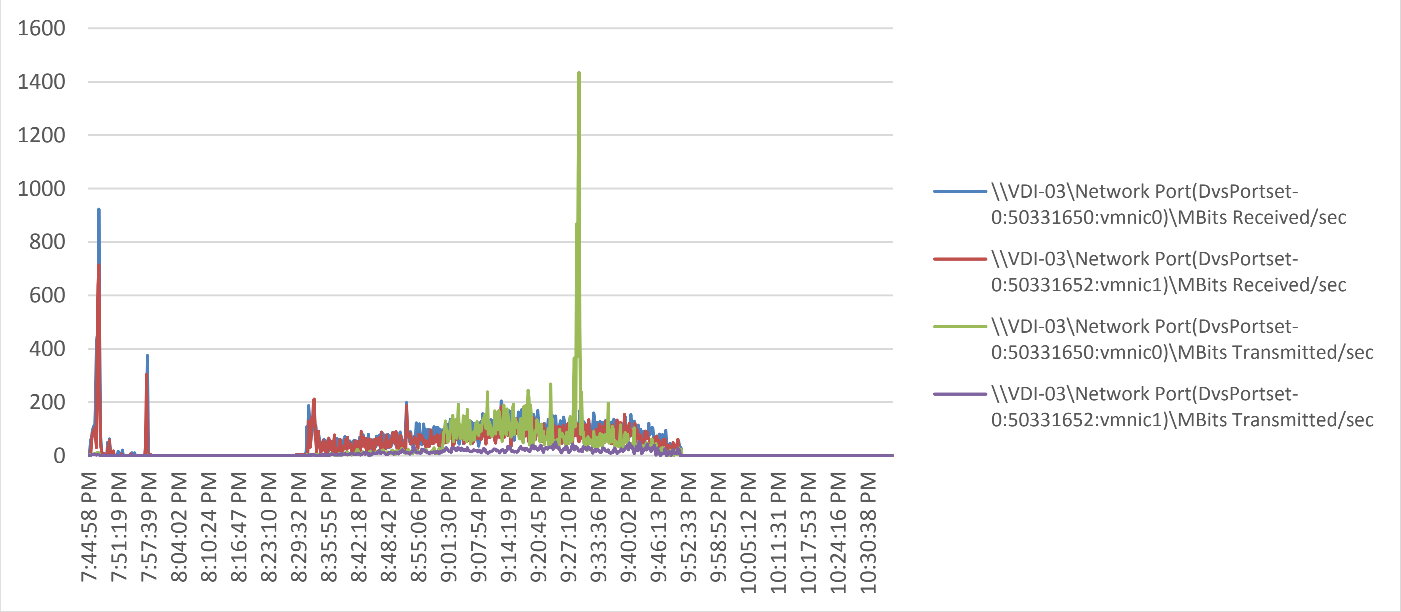

Figure 92 Cluster | 1900 RDS Users | RDS Hosts | Host Fibre Channel Utilization

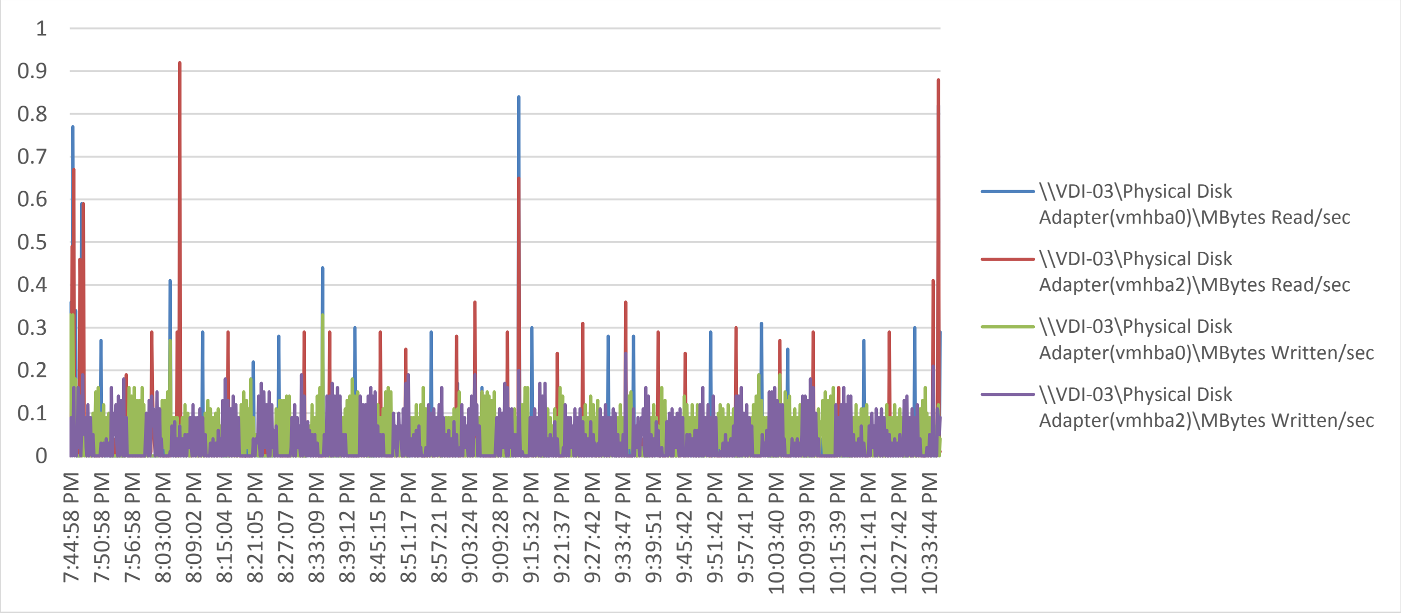

Cluster Workload Testing with 2050 Non-Persistent Desktop Users

This section describes the key performance metrics that were captured on the Cisco UCS, NetApp storage, and Infrastructure VMs during the non-persistent desktop testing. The cluster testing with comprised of 2050 HVD non-persistent desktop sessions using 11 workload blades.

Figure 93 VDI Non-Persistent Cluster Testing with 2050 Users

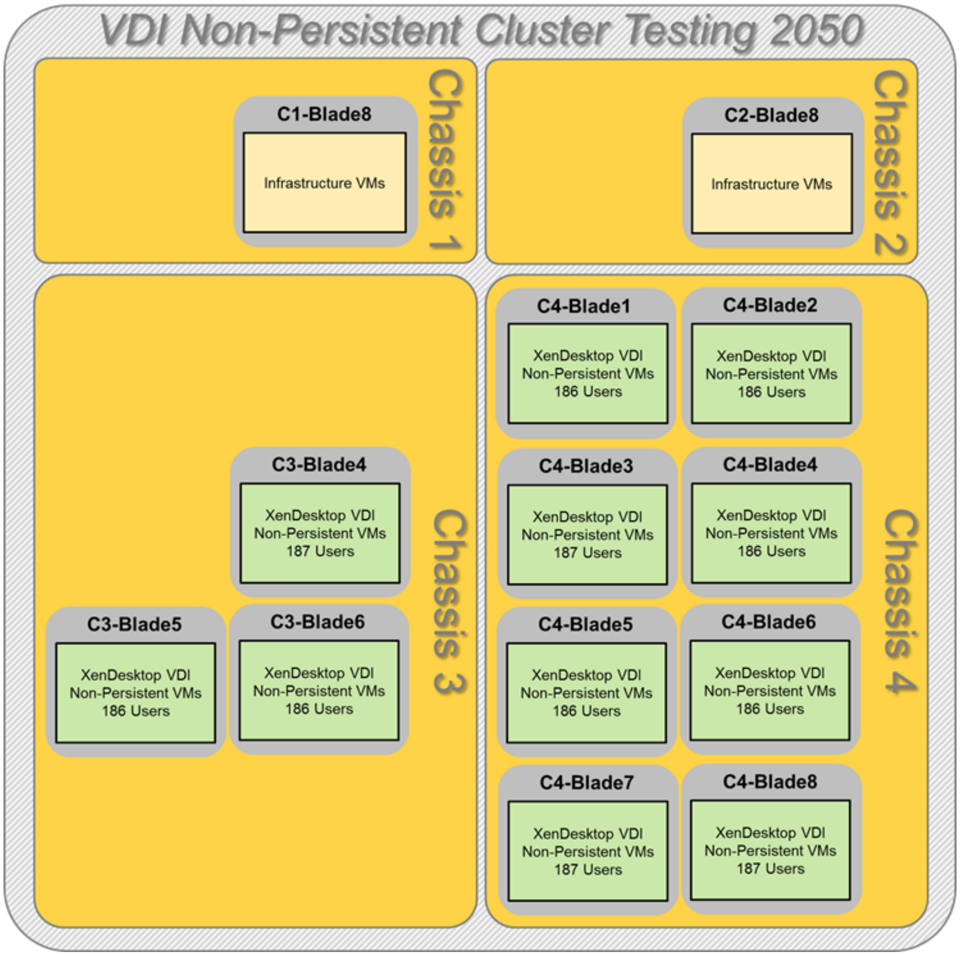
VDI Non-Persistent Cluster Testing 2050
Chassis 1
C1-Blade8
Infrastructure VMs
Chassis 2
C2-Blade8
Infrastructure VMs
Chassis 3
C3-Blade4
XenDesktop VDI Non-Persistent VMs 187 Users
C3-Blade5
XenDesktop VDI Non-Persistent VMs 186 Users
C3-Blade6
XenDesktop VDI Non-Persistent VMs 186 Users
Chassis 4
C4-Blade1
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade2
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade3
XenDesktop VDI Non-Persistent VMs 187 Users
C4-Blade4
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade5
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade6
XenDesktop VDI Non-Persistent VMs 186 Users
C4-Blade7
XenDesktop VDI Non-Persistent VMs 187 Users
C4-Blade8
XenDesktop VDI Non-Persistent VMs 187 Users

The workload for the test is 2050 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 94 Cluster | 2050 VDI-NP Users | VSI Score

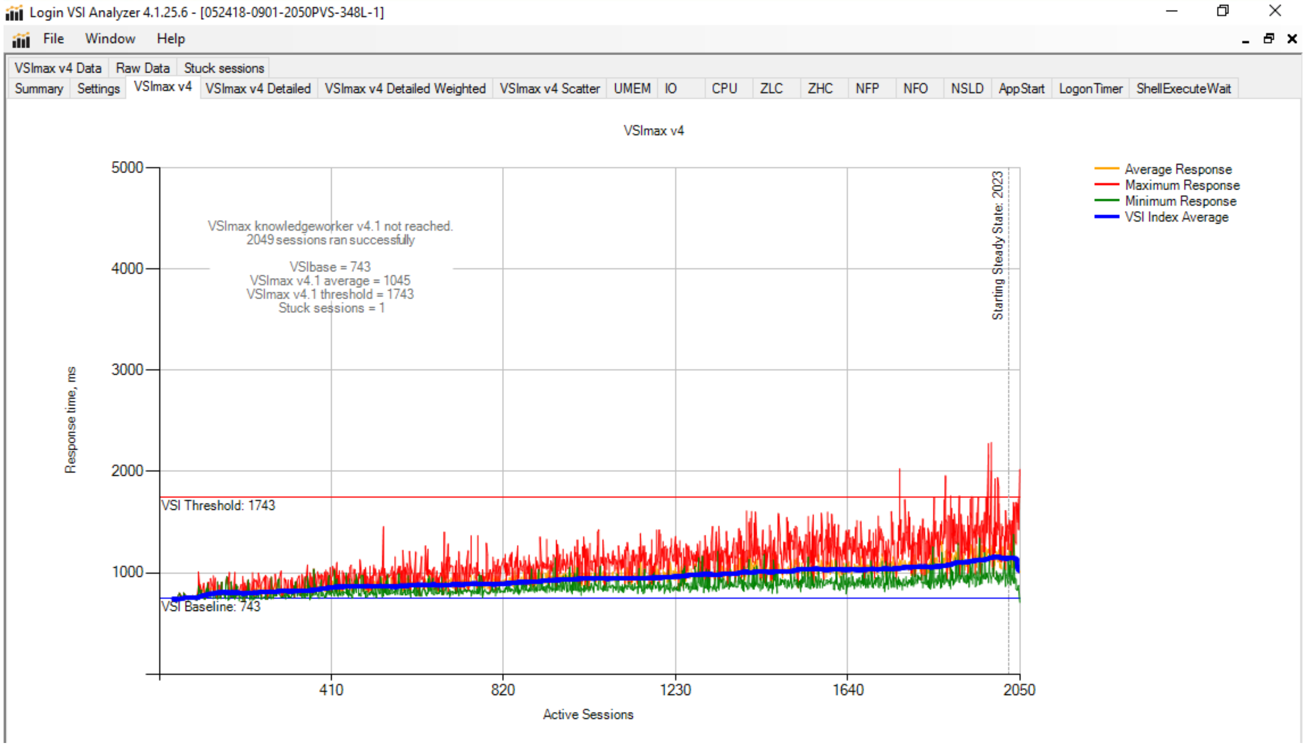

VSImax knowledgeworker v4.1 not reached.
2049 sessions ran successfully

VSIbase = 743
VSImax v4.1 average = 1045
VSImax v4.1 threshold = 1743
Stuck sessions = 1

Figure 95 Cluster | 2050 VDI-NP Users | VSI Repeatability

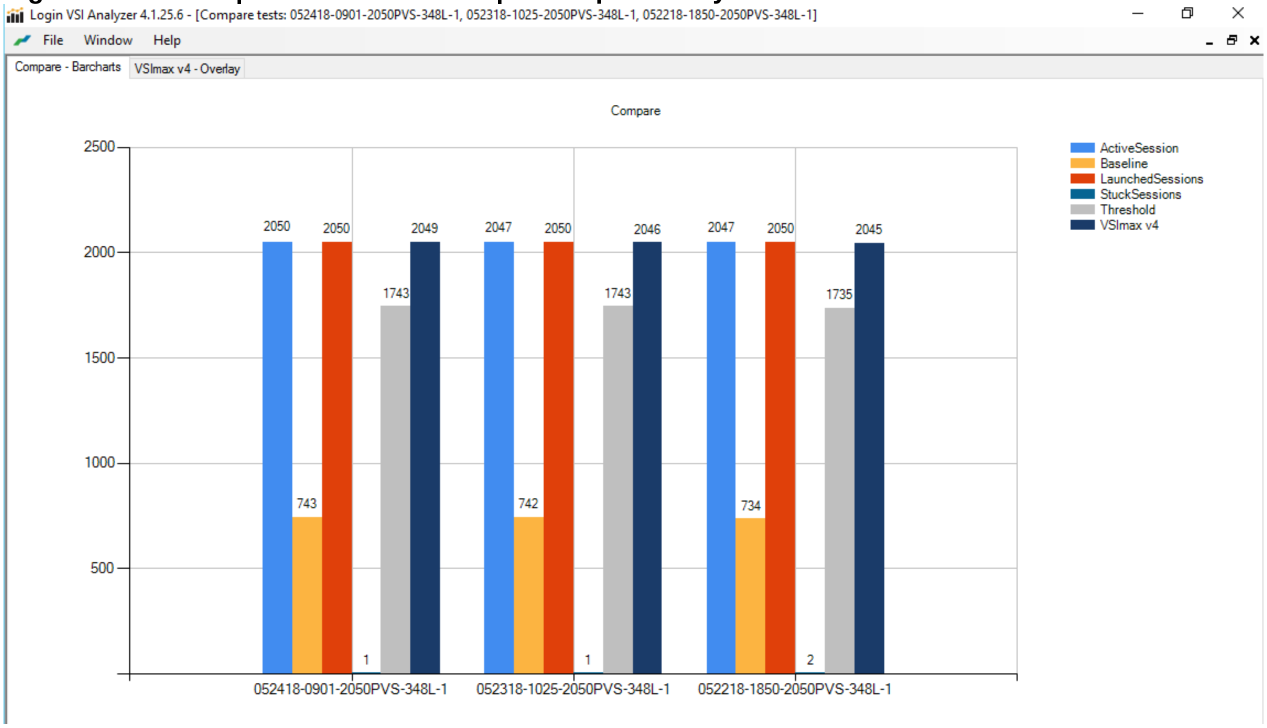

Figure 96 Cluster | 2050 VDI-NP Users | Non-Persistent Hosts | Host CPU Utilization

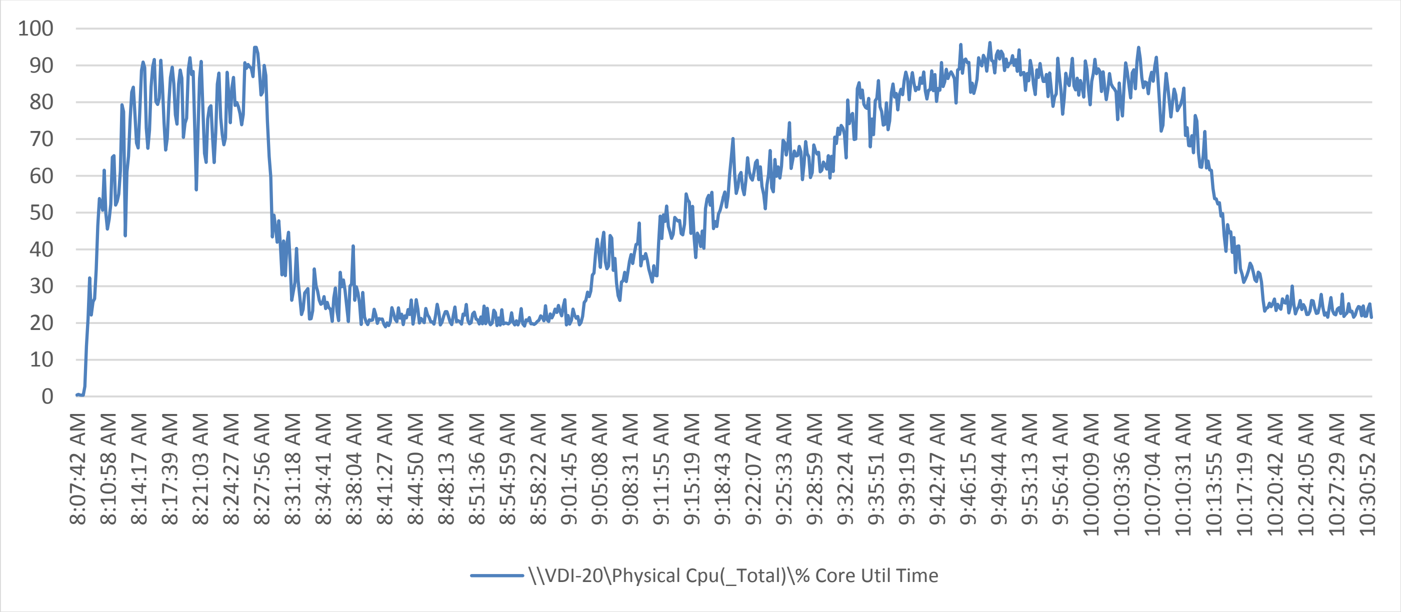

Figure 97 Cluster | 2050 VDI-NP Users | Non-Persistent Hosts | Host Memory Utilization

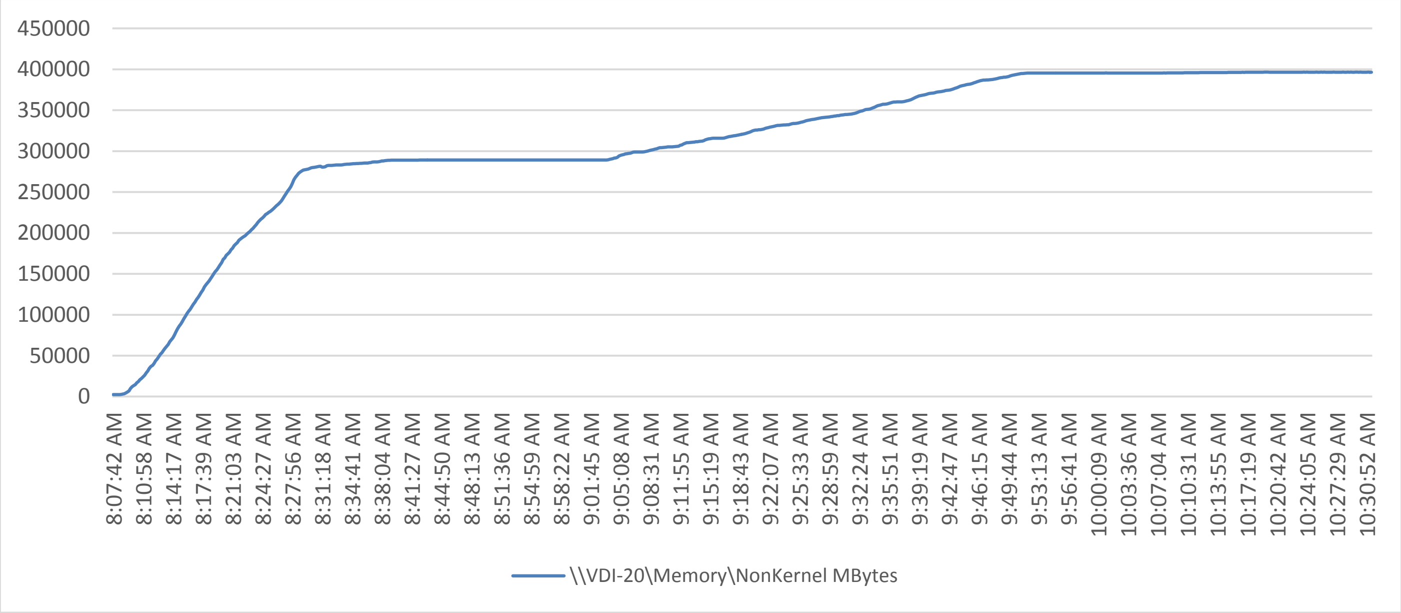

Figure 98 Cluster | 2050 VDI-NP Users | Non-Persistent Hosts | Host Network Utilization

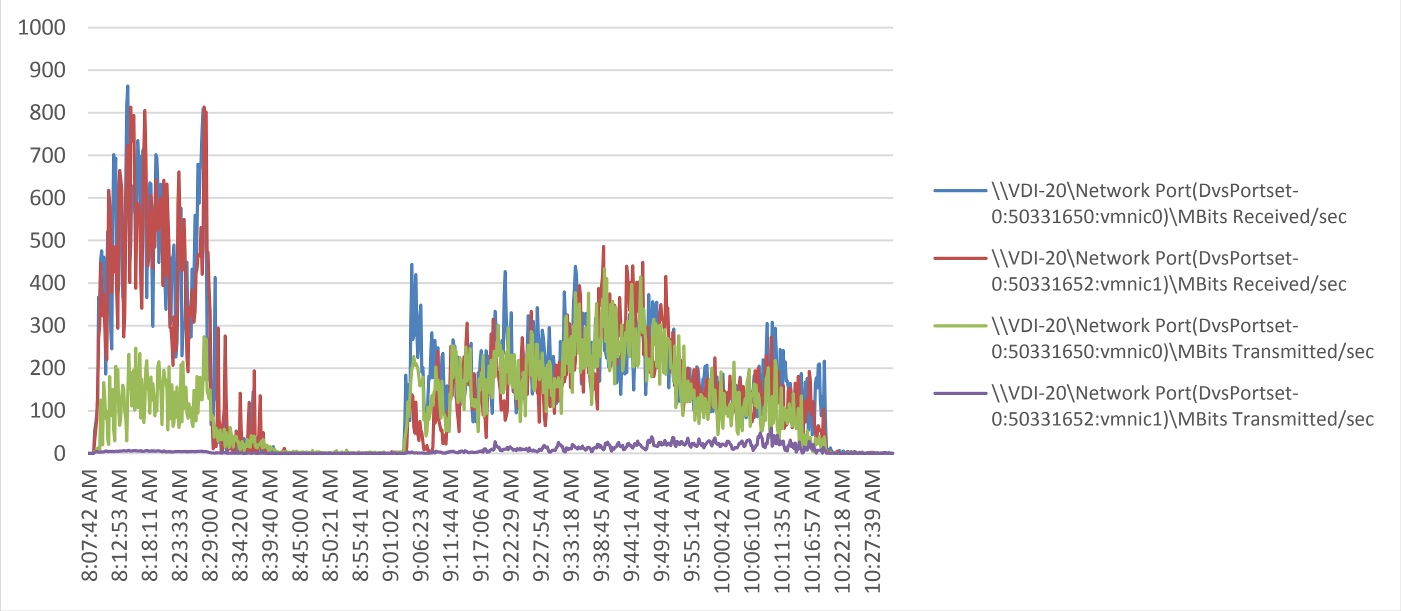

Figure 99 Cluster | 2050 VDI-NP Users | Non-Persistent Hosts | Host Fibre Channel Utilization

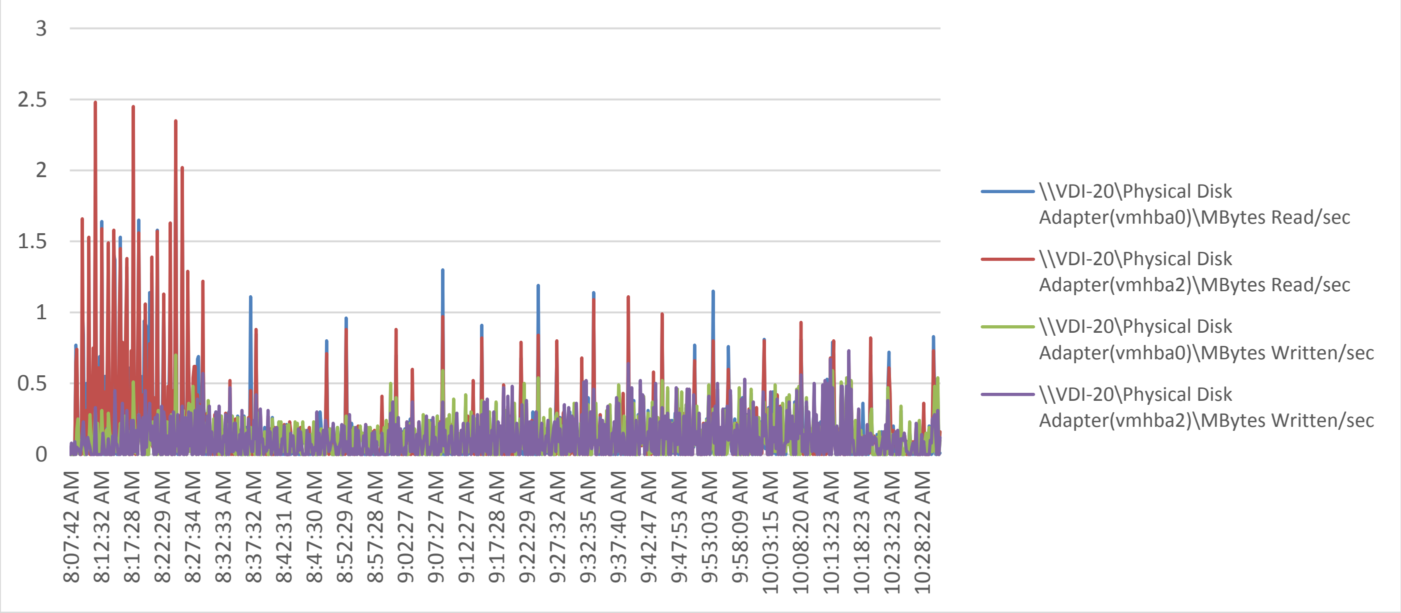

Cluster Workload Testing with 2050 Persistent Desktop Users

This section describes the key performance metrics that were captured on the Cisco UCS, NetApp storage, and Infrastructure VMs during the persistent desktop testing. The cluster testing with comprised of 2050 HVD Persistent desktop sessions using 11 workload blades.

Figure 100 VDI Persistent Cluster Testing with 2050 Users

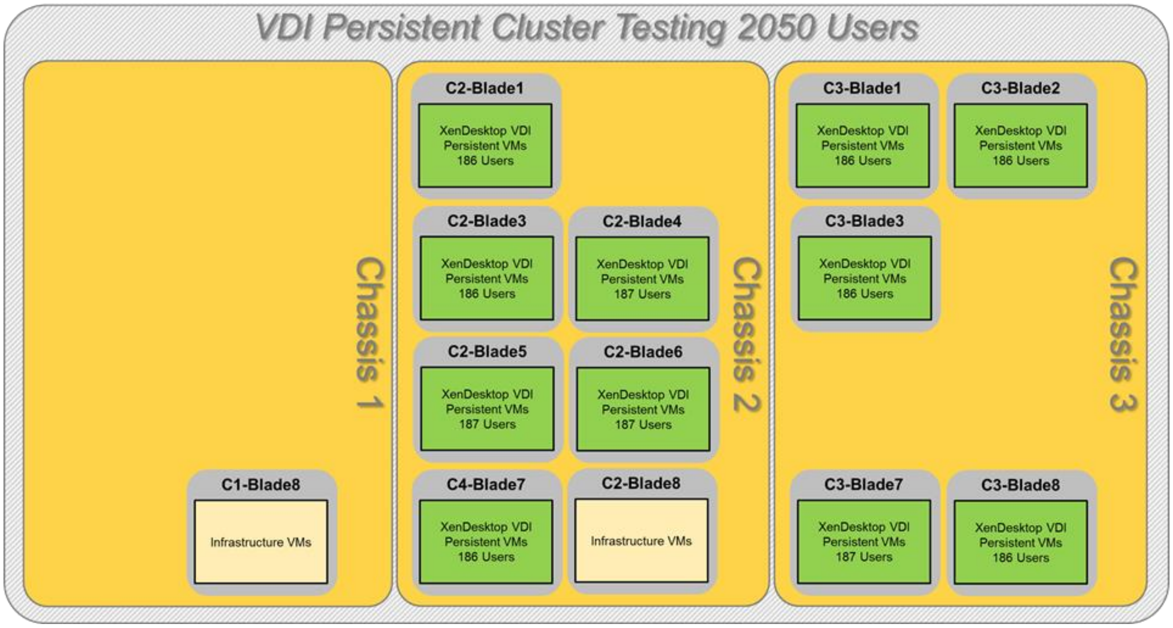
VDI Persistent Cluster Testing 2050 Users
Chassis 1
C1-Blade8
Infrastructure VMs
Chassis 2
C2-Blade1
XenDesktop VDI Persistent VMs 186 Users
C2-Blade3
XenDesktop VDI Persistent VMs 186 Users
C2-Blade4
XenDesktop VDI Persistent VMs 187 Users
C2-Blade5
XenDesktop VDI Persistent VMs 187 Users
C2-Blade6
XenDesktop VDI Persistent VMs 187 Users
C4-Blade7
XenDesktop VDI Persistent VMs 186 Users
C2-Blade8
Infrastructure VMs
Chassis 3
C3-Blade1
XenDesktop VDI Persistent VMs 186 Users
C3-Blade2
XenDesktop VDI Persistent VMs 186 Users
C3-Blade3
XenDesktop VDI Persistent VMs 186 Users
C3-Blade7
XenDesktop VDI Persistent VMs 187 Users
C3-Blade8
XenDesktop VDI Persistent VMs 186 Users

The workload for the test is 2050 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 101 Cluster | 2050 VDI-P Users | VSI Score

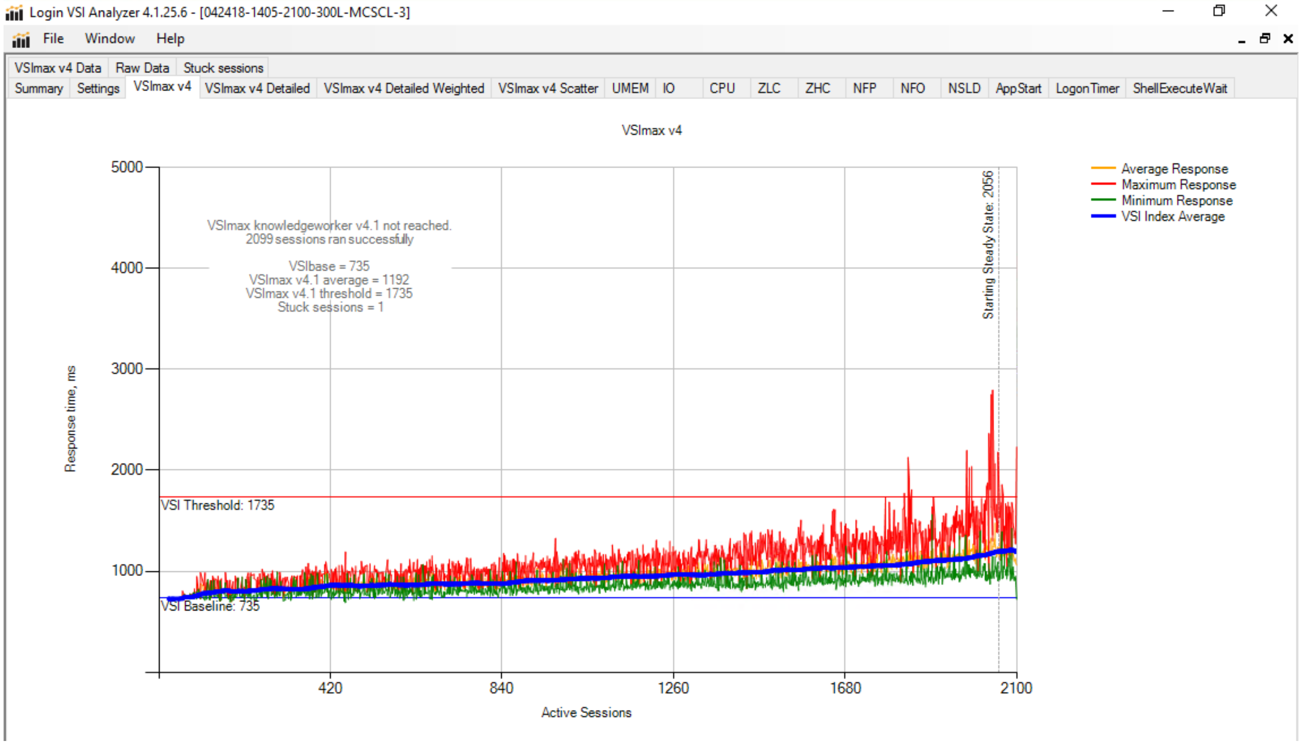
Login VSI Analyzer 4.1.25.6 - [042418-1405-2100-300L-MCSCL-3]

VSImax knowledgeworker v4.1 not reached.
2099 sessions ran successfully

VSIbase = 735
VSImax v4.1 average = 1192
VSImax v4.1 threshold = 1735
Stuck sessions = 1

Figure 102 Cluster | 2050 VDI-P Users | VSI Repeatability

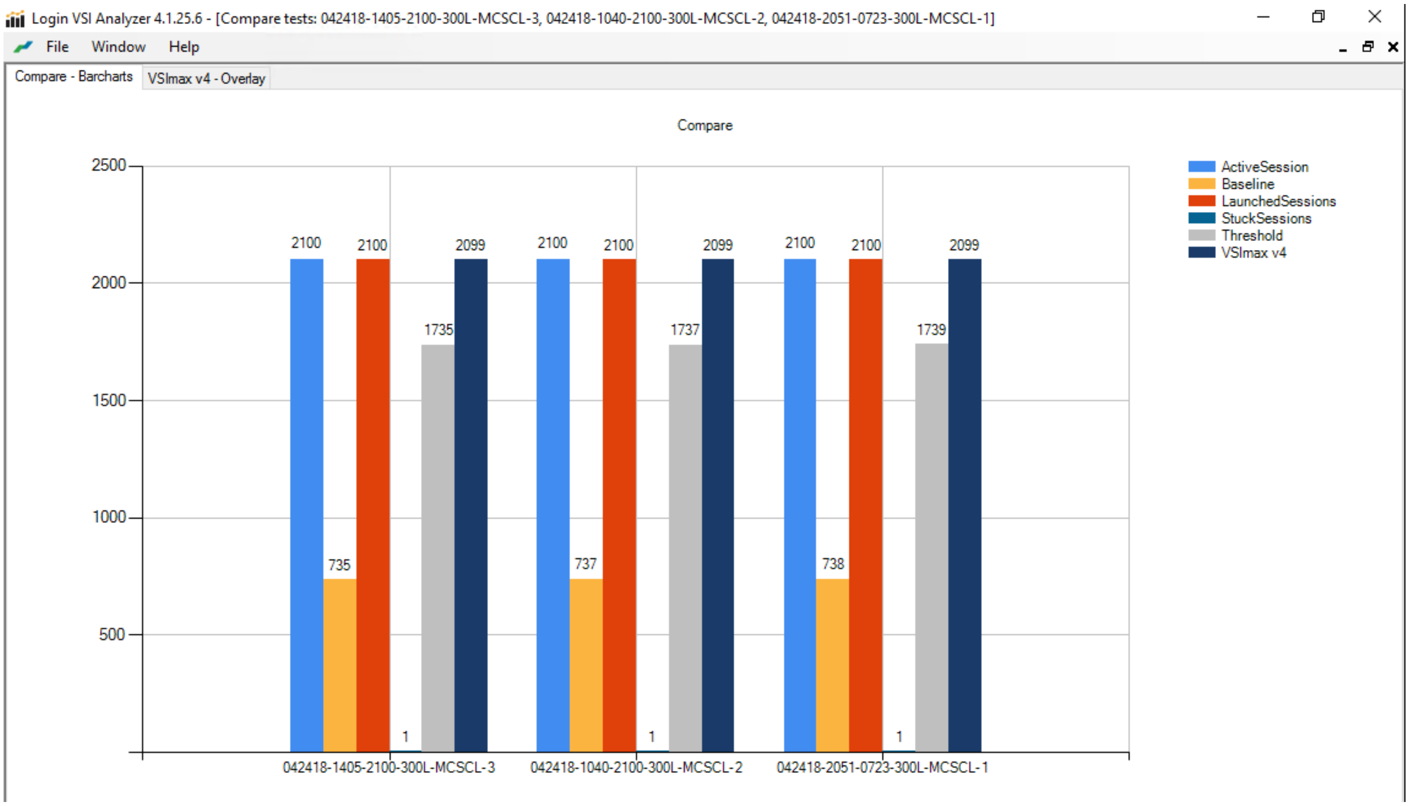
Login VSI Analyzer 4.1.25.6 - [Compare tests: 042418-1405-2100-300L-MCSCL-3, 042418-1040-2100-300L-MCSCL-2, 042418-2051-0723-300L-MCSCL-1]
File
Window
Help
Compare - Barcharts
VSImax v4 - Overlay
Compare
2500
2000
1500
1000
500
2100
735
2100
1
1735
2099
2100
737
2100
1
1737
2099
2100
738
2100
1
1739
2099
042418-1405-2100-300L-MCSCL-3
042418-1040-2100-300L-MCSCL-2
042418-2051-0723-300L-MCSCL-1
ActiveSession
Baseline
LaunchedSessions
StuckSessions
Threshold
VSImax v4

Figure 103 Cluster | 2050 VDI-P Users | Persistent Hosts | Host CPU Utilization

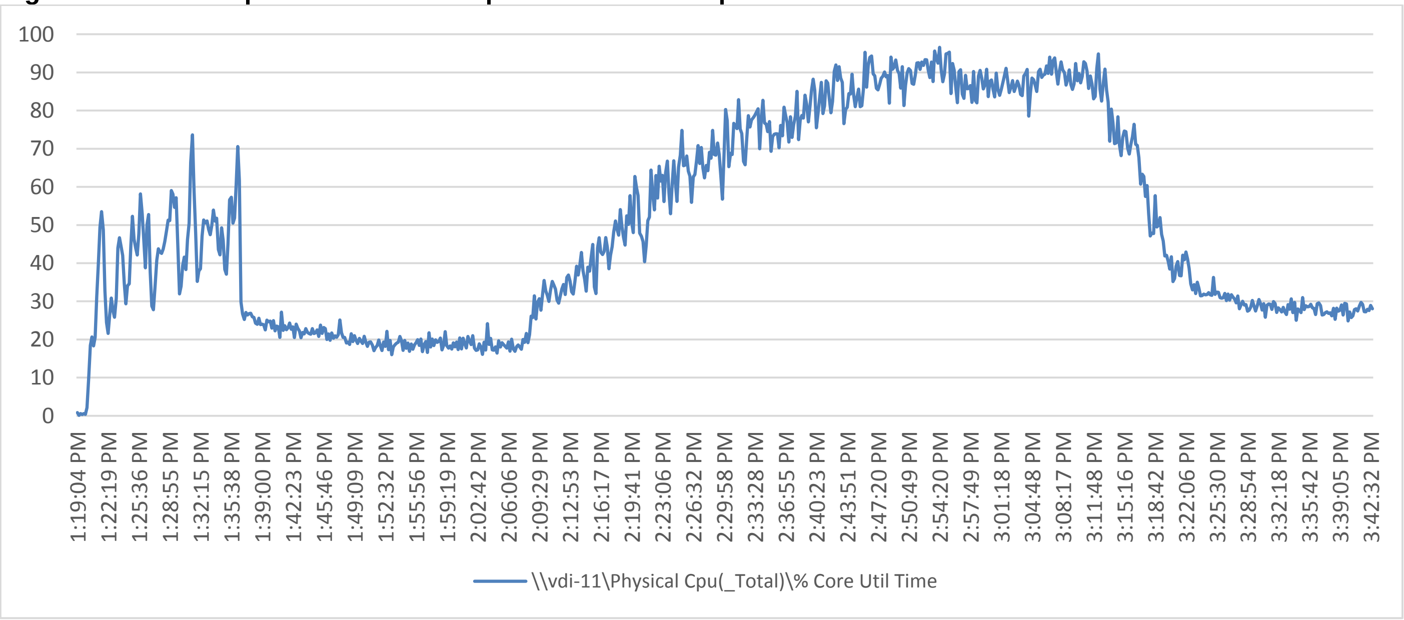

Figure 104 Cluster | 2050 VDI-P Users | Persistent Hosts | Host Memory Utilization

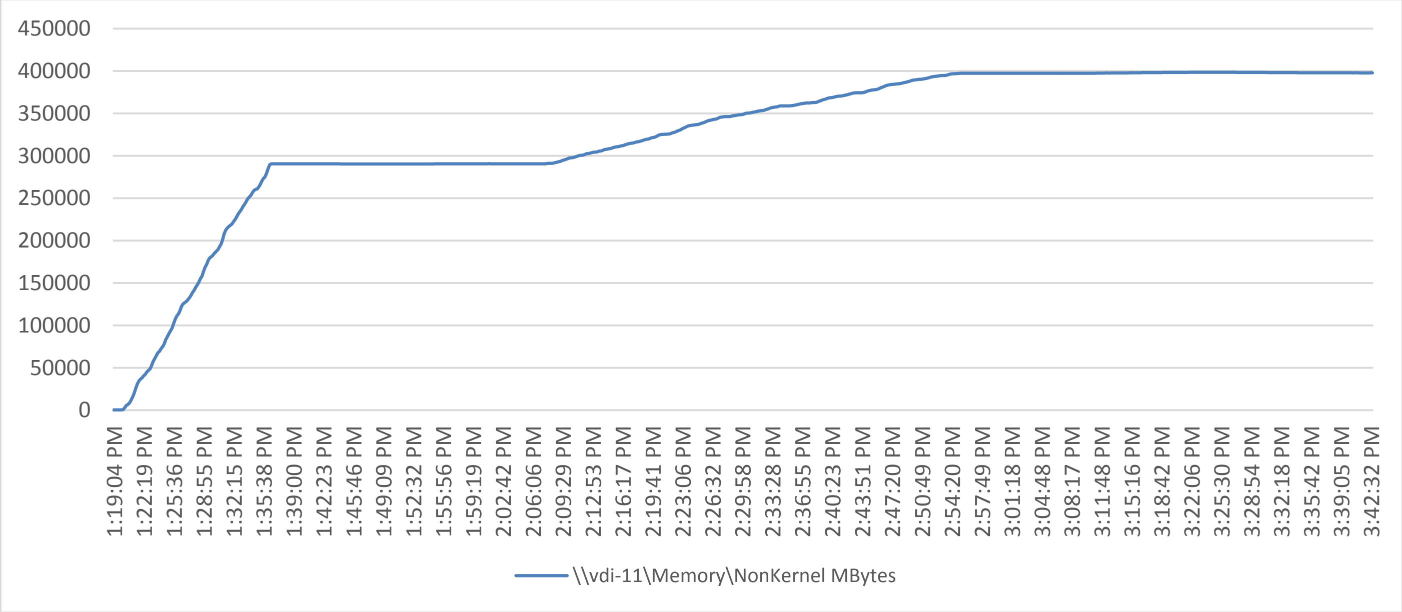

Figure 105 Cluster | 2050 VDI-P Users | Persistent Hosts | Host Network Utilization

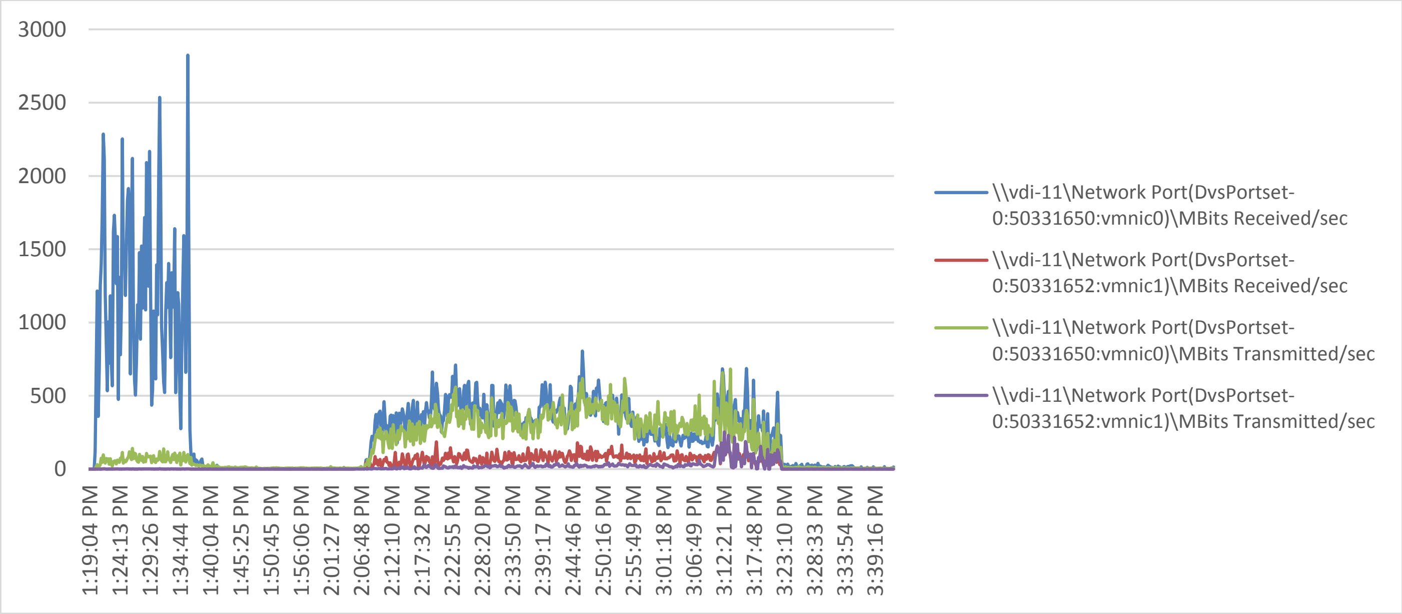

Figure 106 Cluster | 2050 VDI-P Users | Persistent Hosts | Host Fibre Channel Utilization

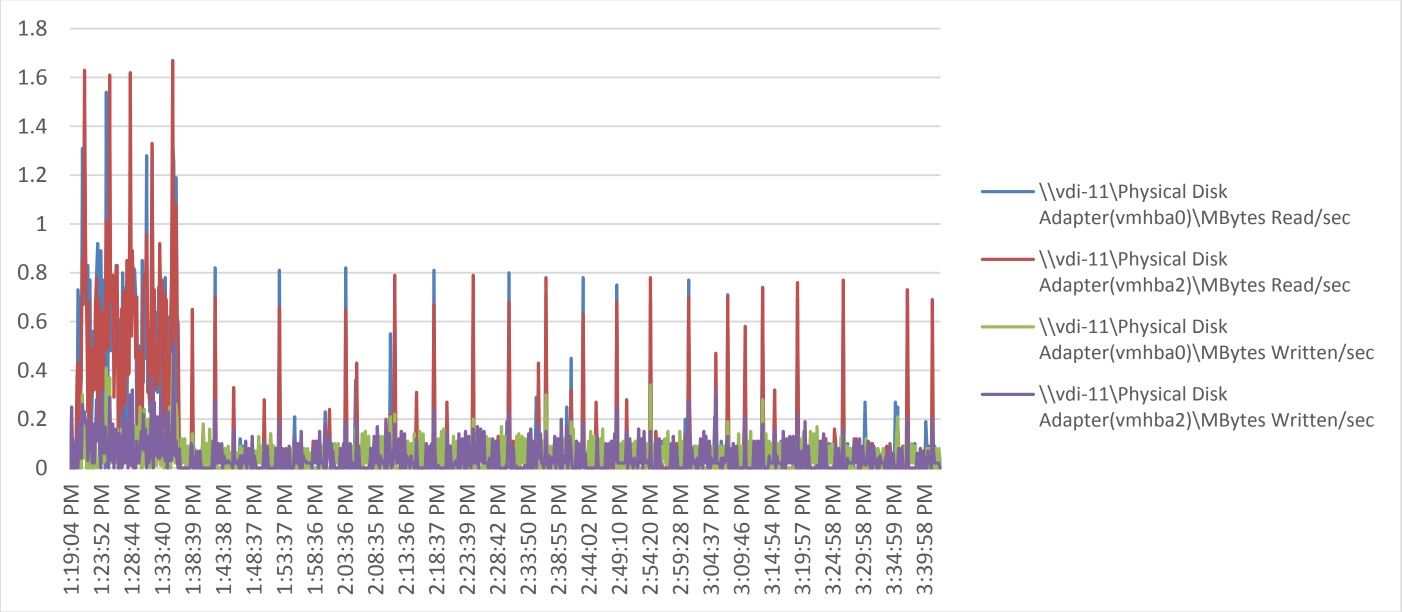

Full Scale Mixed Workload Testing with 6000 Users

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. The full-scale testing with 6000 users comprised of: 1900 Hosted Shared Desktop Sessions using 8 blades, 2050 HVD Non-Persistent sessions using 11 blades, and 2050 HVDI Persistent sessions using 11 blades.

The combined mixed workload for the solution is 6000 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 Testing 6000 Users

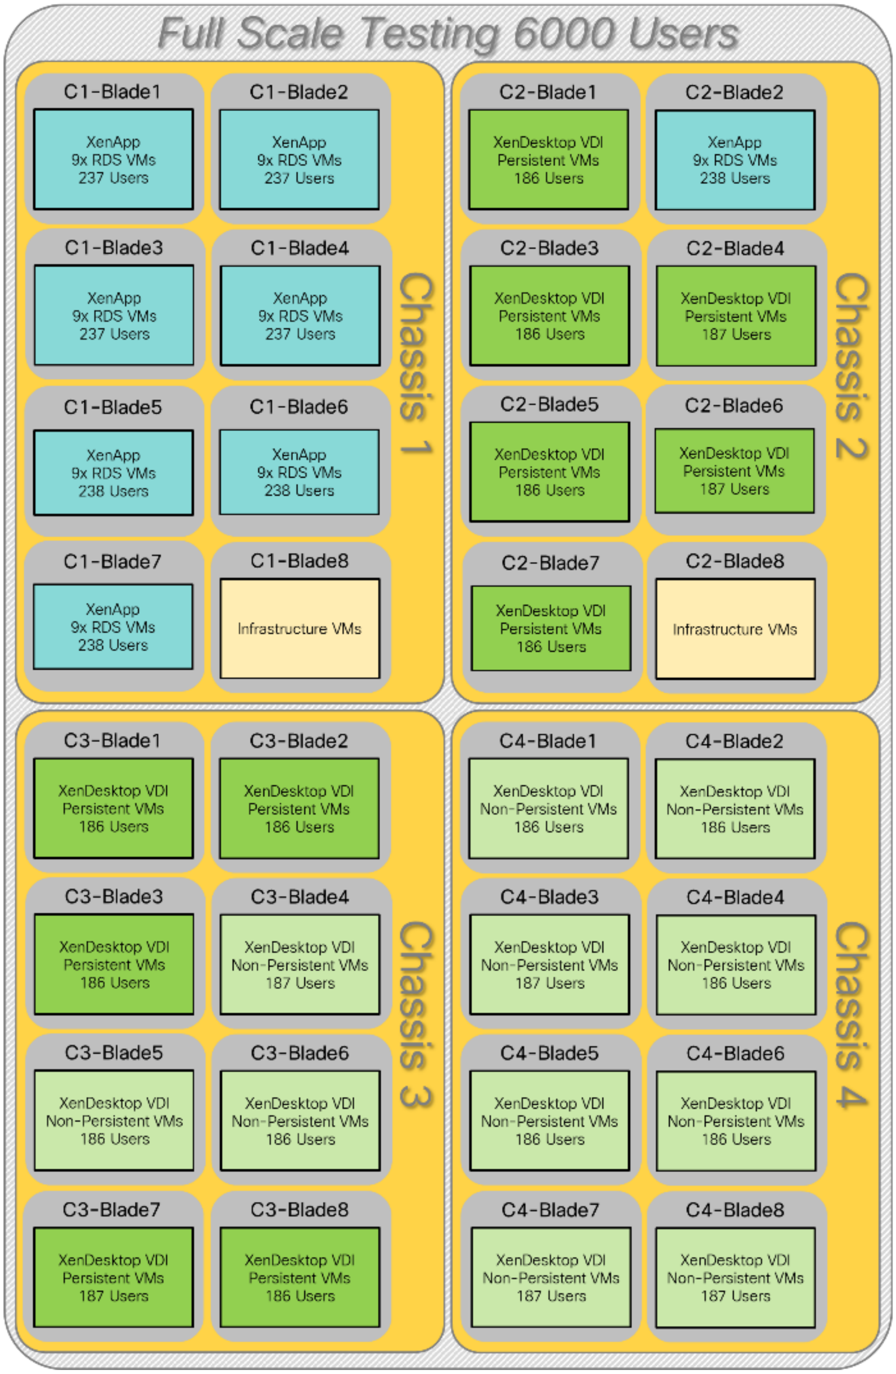

The configured system efficiently and effectively delivered the following results.

Figure 108 Full Scale | 6000 Mixed Users | VSI Score

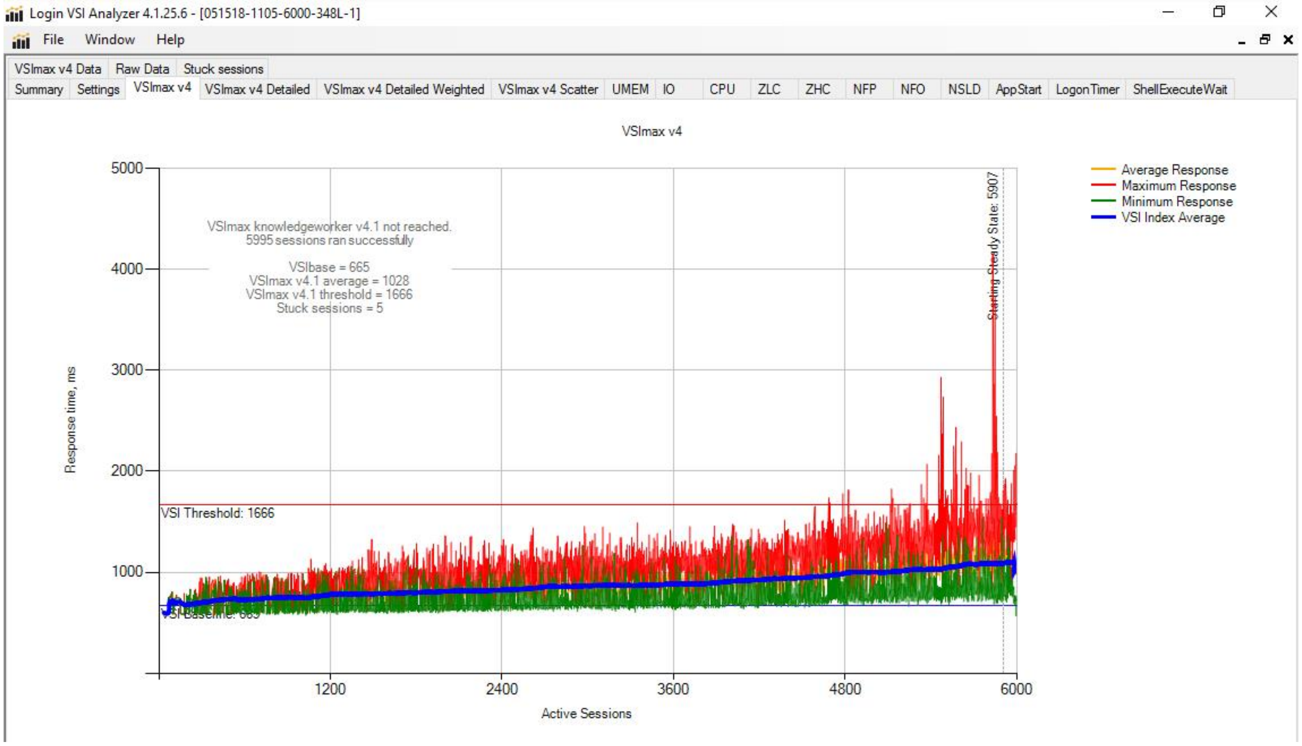

Login VSI Analyzer 4.1.25.6 - [051518-1105-6000-348L-1]

VSImax knowledgeworker v4.1 not reached.
5995 sessions ran successfully

VSIbase = 665
VSImax v4.1 average = 1028
VSImax v4.1 threshold = 1666
Stuck sessions = 5

Figure 109 Full Scale | 6000 Mixed Users | VSI Repeatability

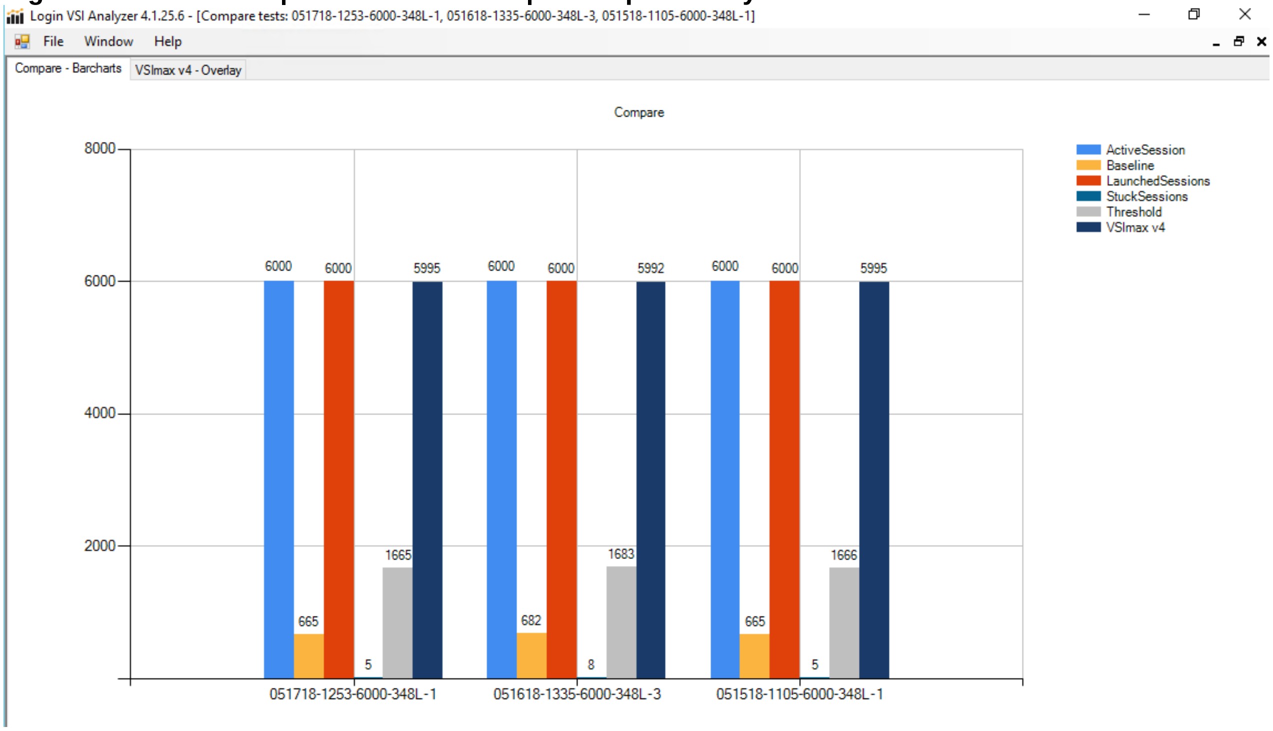

Login VSI Analyzer 4.1.25.6 - [Compare tests: 051718-1253-6000-348L-1, 051618-1335-6000-348L-3, 051518-1105-6000-348L-1]

Figure 110 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

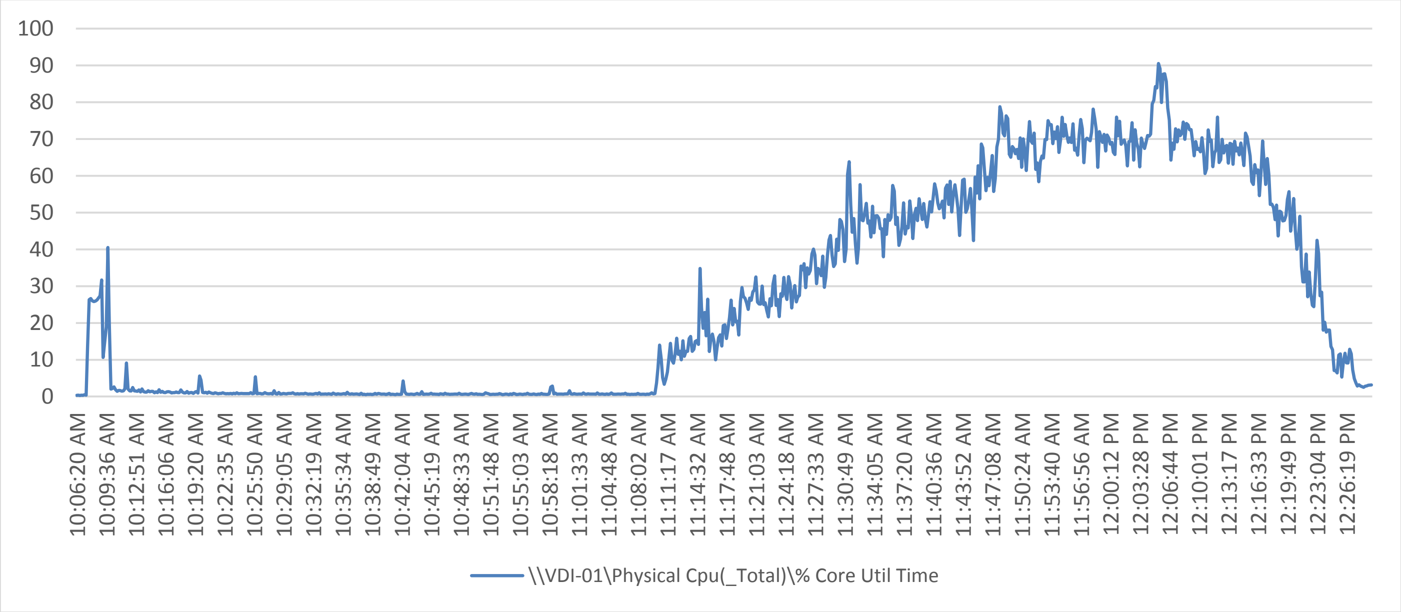

Figure 111 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

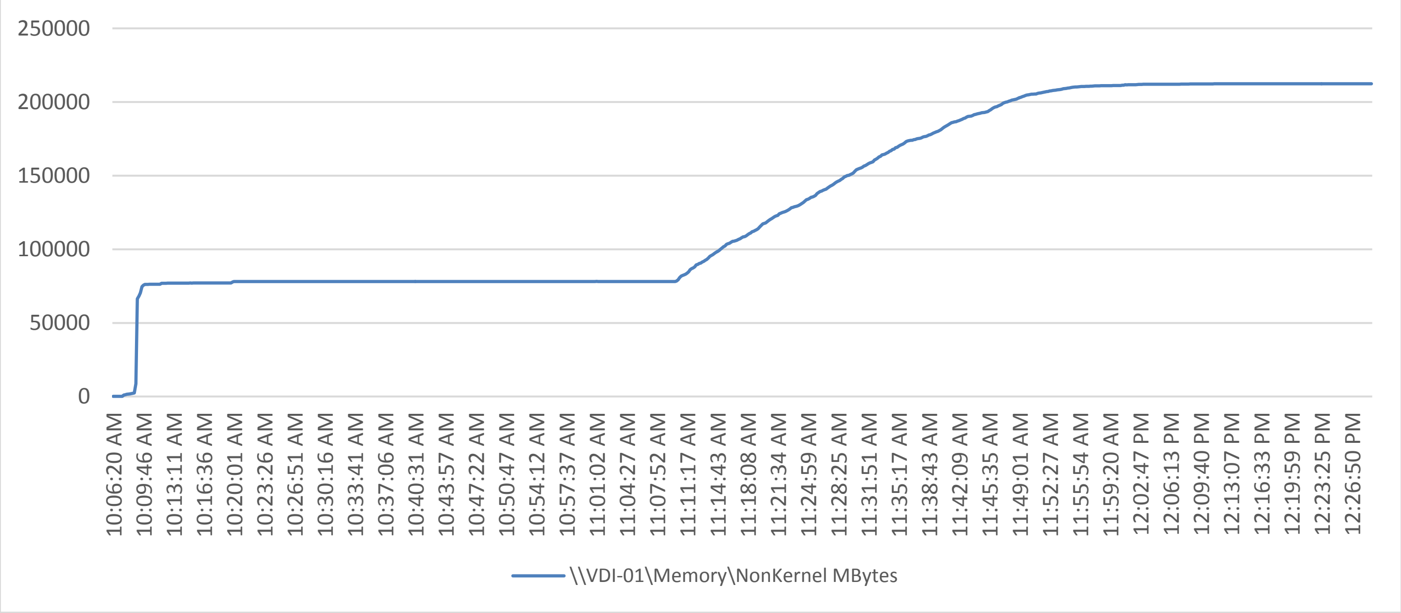

Figure 112 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

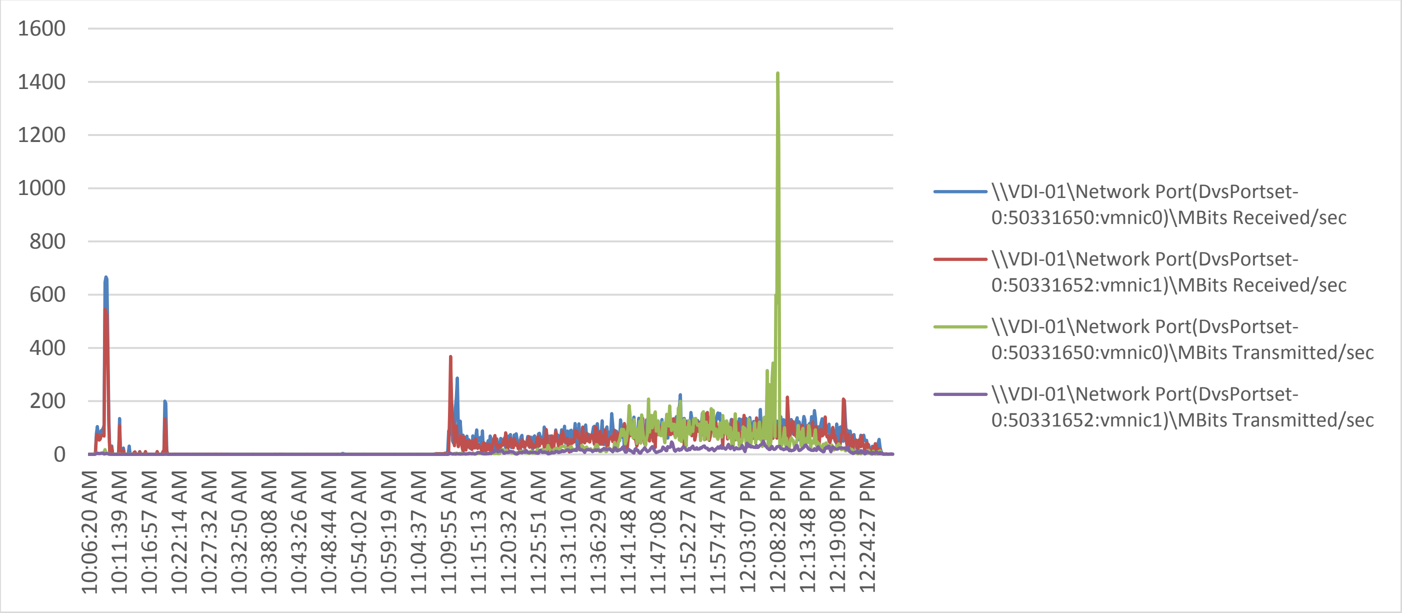

Figure 113 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

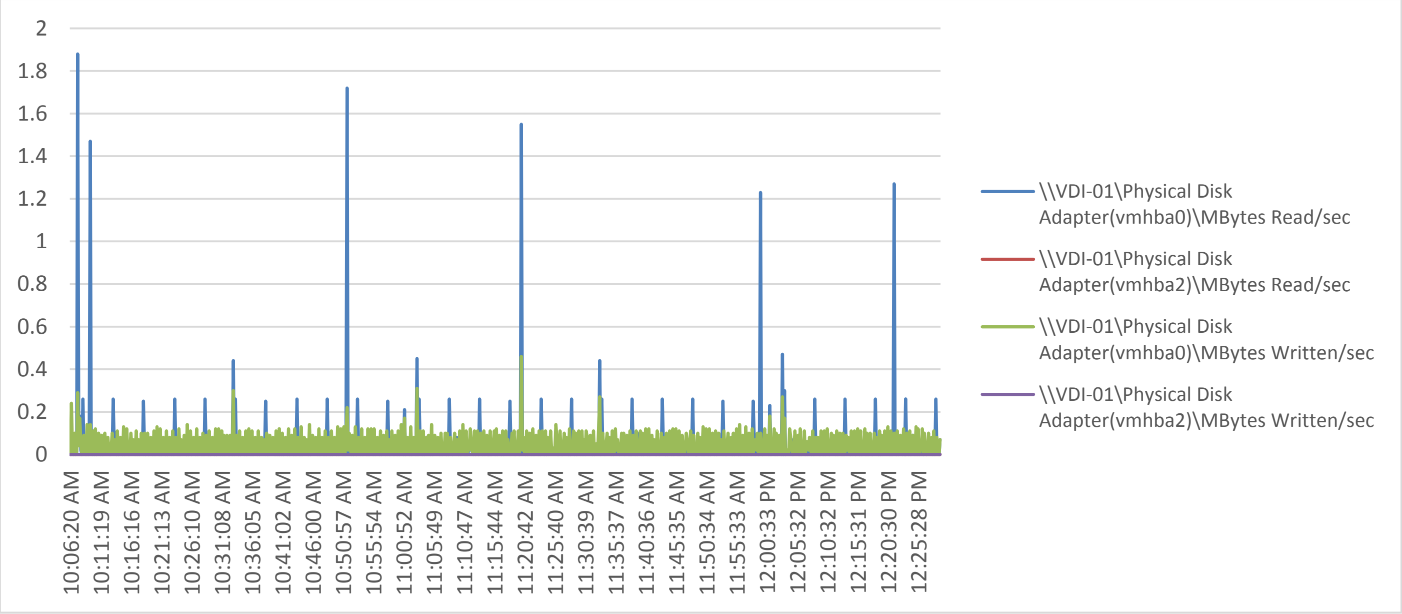

Figure 114 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

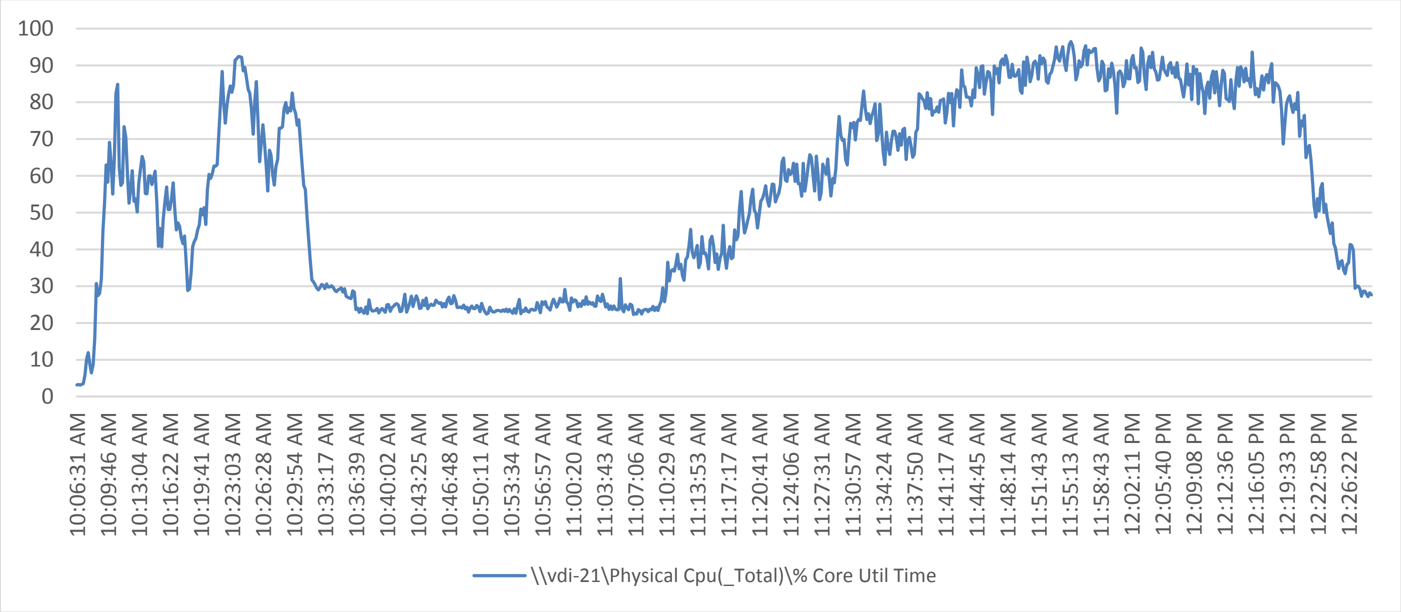

Figure 115 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

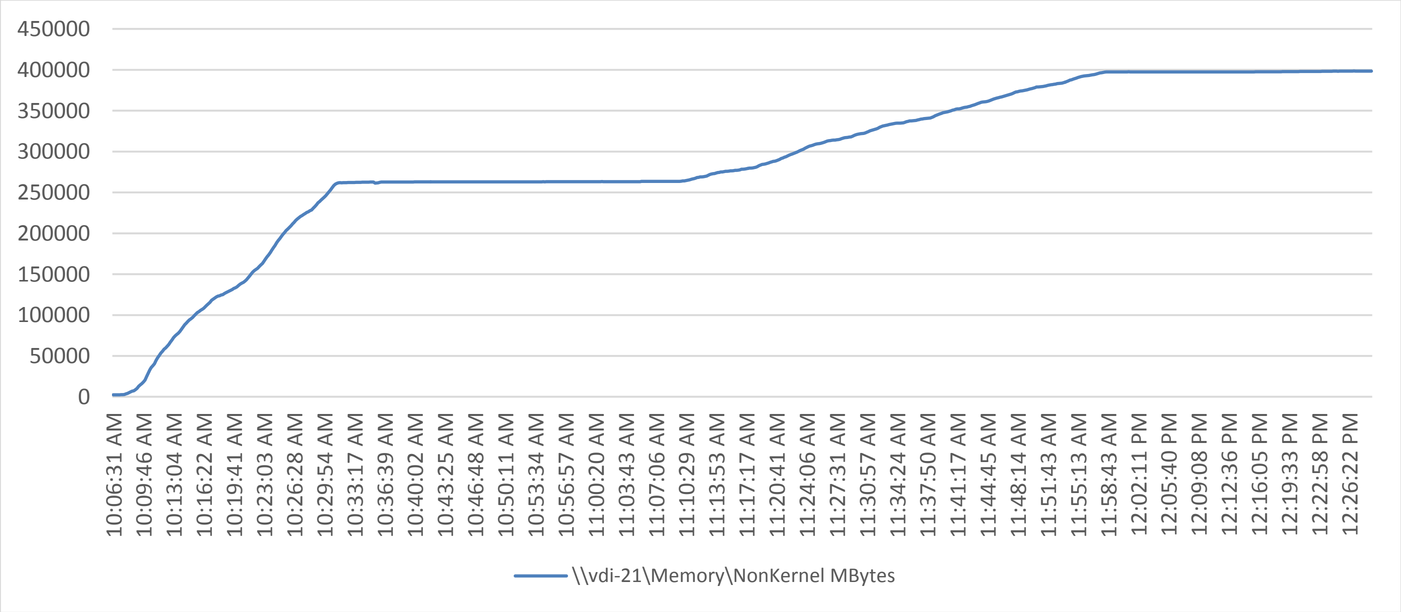

Figure 116 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

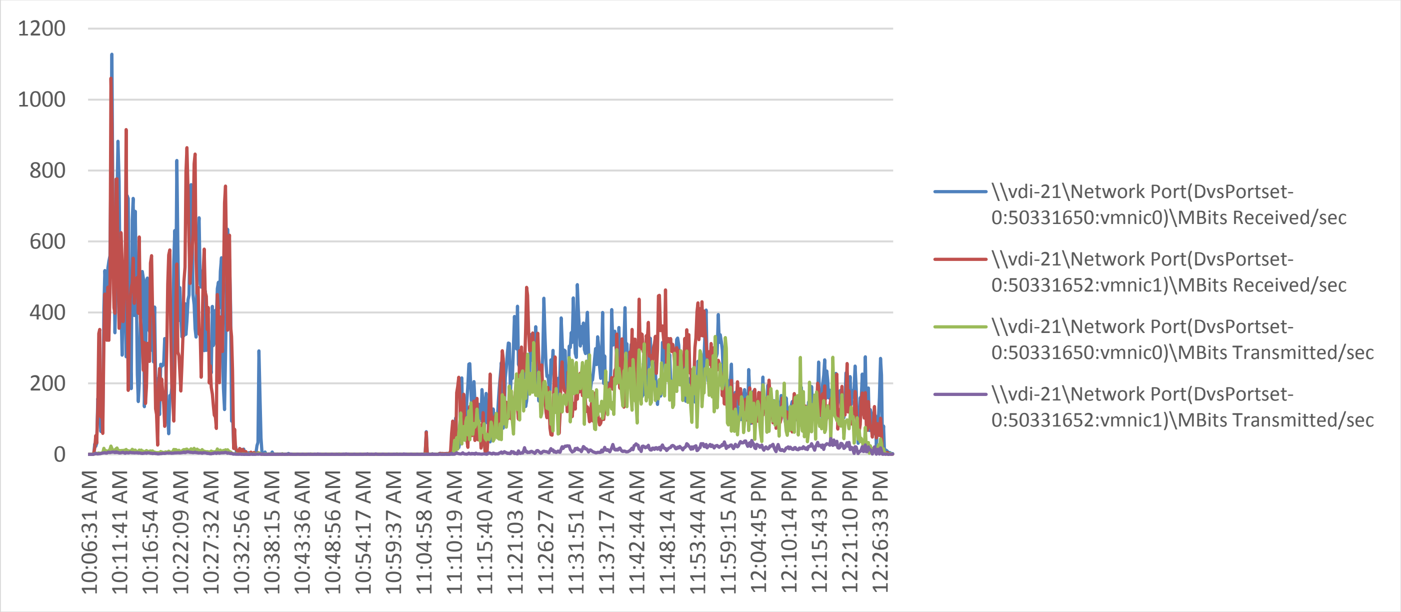

Figure 117 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

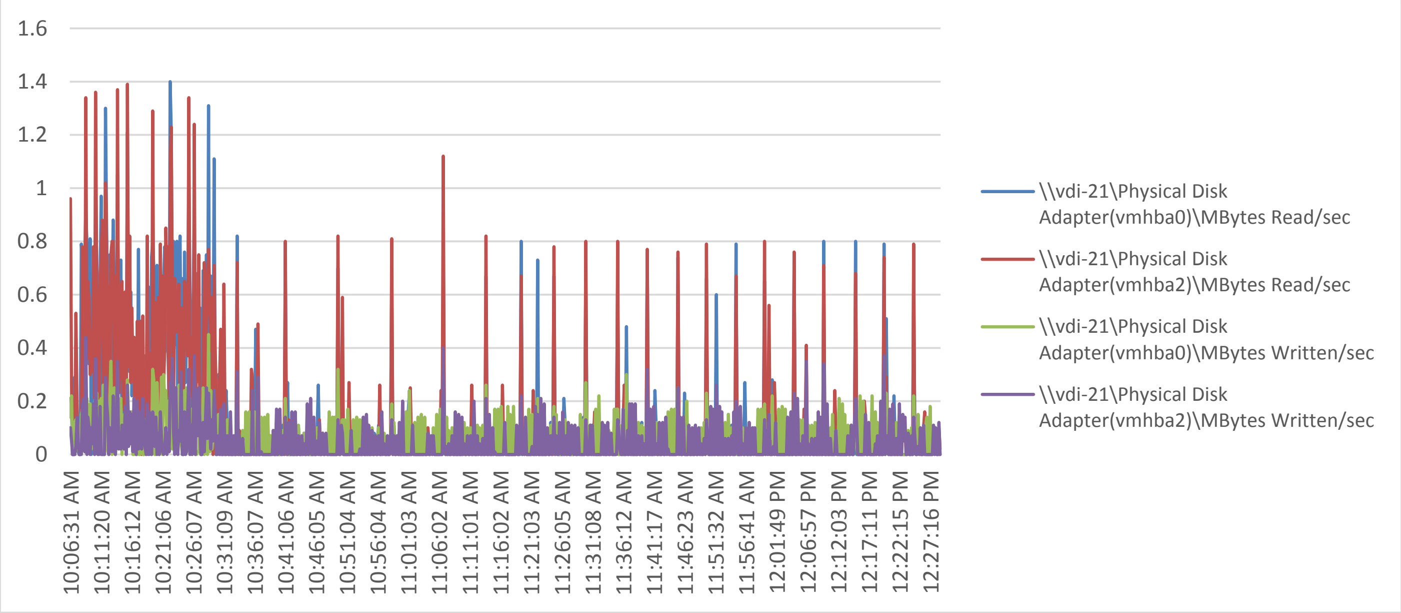

Figure 118 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

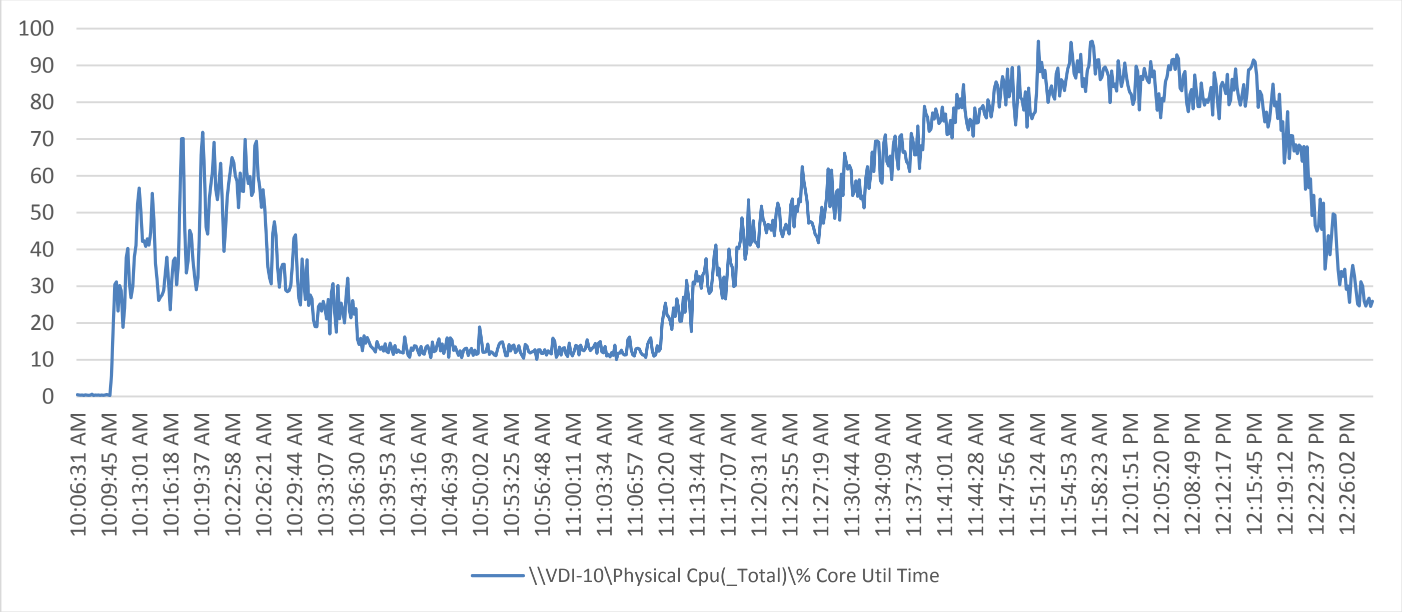

Figure 119 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

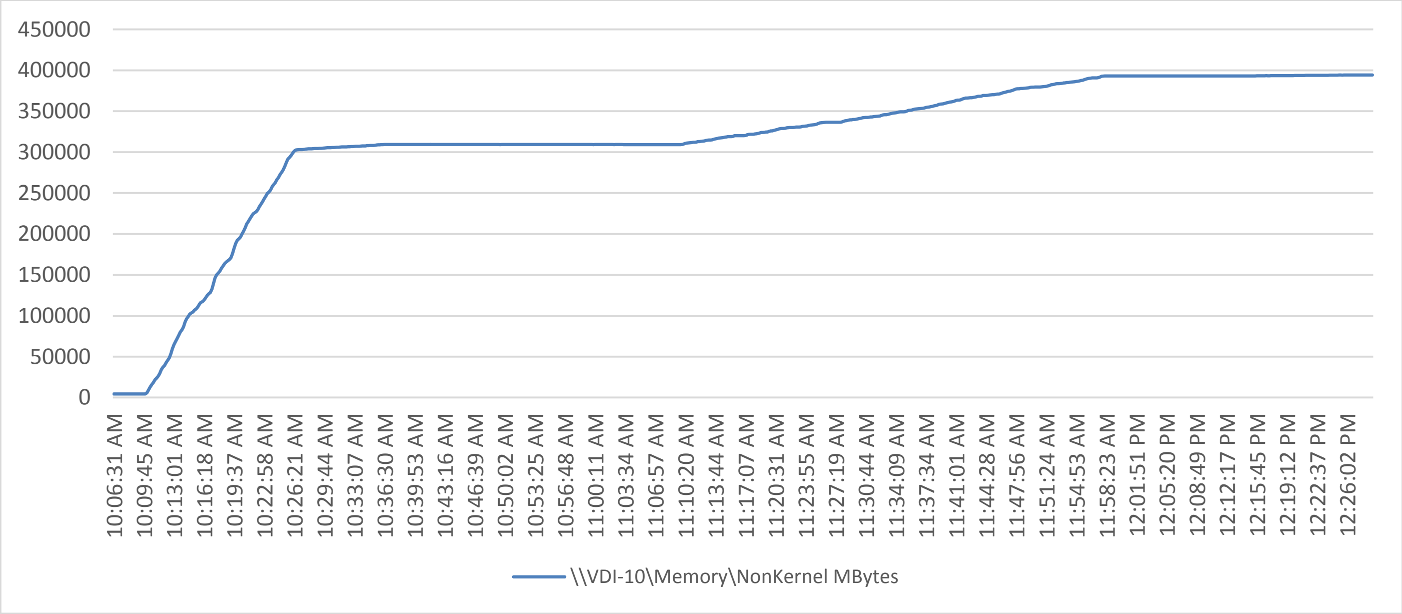

Figure 120 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

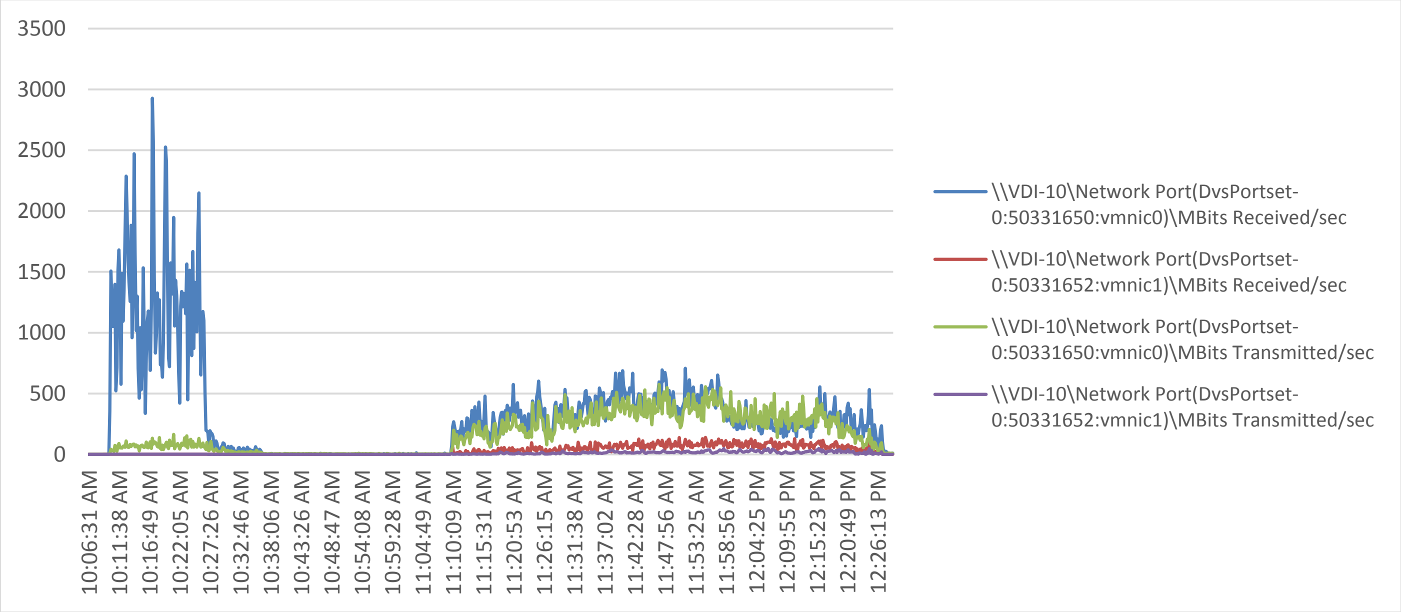

Figure 121 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

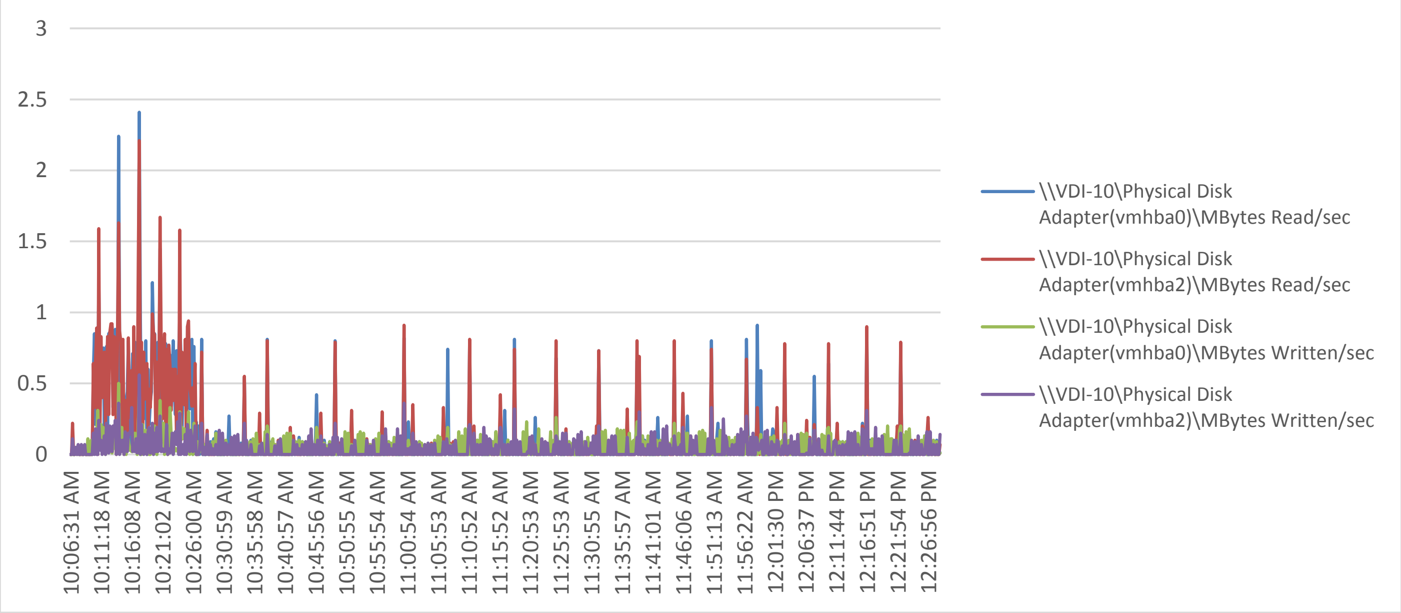

AFF A300 Storage Detailed Test Results for Cluster Scalability Test

This section highlights and provides analysis of the NetApp AFF A300 storage system performance results for each of the Citrix software module testing (HSD, PVS, Persistent), which we call cluster testing and they are identified previously in this document. Specifically, it depicts and discusses the results for the following test case scenarios:

- 1900 Windows Server 2016 Citrix Hosted Shared desktops (XenApp)
- 2050 Windows 10 x64 Citrix PVS Non-Persistent desktops
- 2050 Windows 10 x64 Citrix Persistent Full-Clone desktops

From a storage perspective, it is critical to maintain a latency of less than a millisecond for an optimal end-user experience no matter the IOPS and bandwidth being driven. The test results indicate that the AFF A300 storage delivers that essential minimum level of latency despite driving a substantial amount of IOPs and bandwidth for the thousands of desktops hosted on the AFF A300 system.

The sections that follow show screenshots of the AFF A300 storage test data for all three use case scenarios with information about IOPS, bandwidth, and latency at the peak of each use case test. In all three use cases (cluster level testing with HSD PVS VDI, full clone persistent desktops and full-scale mixed workload sessions, the criteria followed prior to launching Login VSI workload test are the same.

AFF A300 Storage Test Results for 1900 Citrix HSD (XenApp) Windows 2016 Sessions

This test uses Login VSI as the workload generator in Benchmark mode with the Knowledge Worker user type and with Citrix Hosted Shared Desktops (RDSH) Sessions as the VDI delivery mechanism. This first highlighted cluster test shows that the AFF A300 can easily handle this workload with exceptional end-user experience as confirmed from Login VSI.

The first AFF A300 GUI screenshot shows the performance during the 1900 HSD (XenApp) sessions. As with all scenarios, there were three 1900 HSD simulation test runs completed in total, all with very similar results. As indicated in charts from one of these simulations, we maintained latency of less than or close to one millisecond (ms) for both read and write operations throughout this entire run. This resulted in a confirmed outstanding end-user experience for the simulated Citrix HSD users independently verified by Login VSI. It has been observed during this component of the testing with observed peak total values 3,000 IOPS during steady state and 10,500 IOPS during logoff. The latency is less than 1ms which lends itself to a great end-user experience.

1900 Citrix HSD (XenApp) Cluster Test: Storage Charts

Figure 122 AFF A300 | 1900 Users HSD Cluster Test | Read/Write IOPS

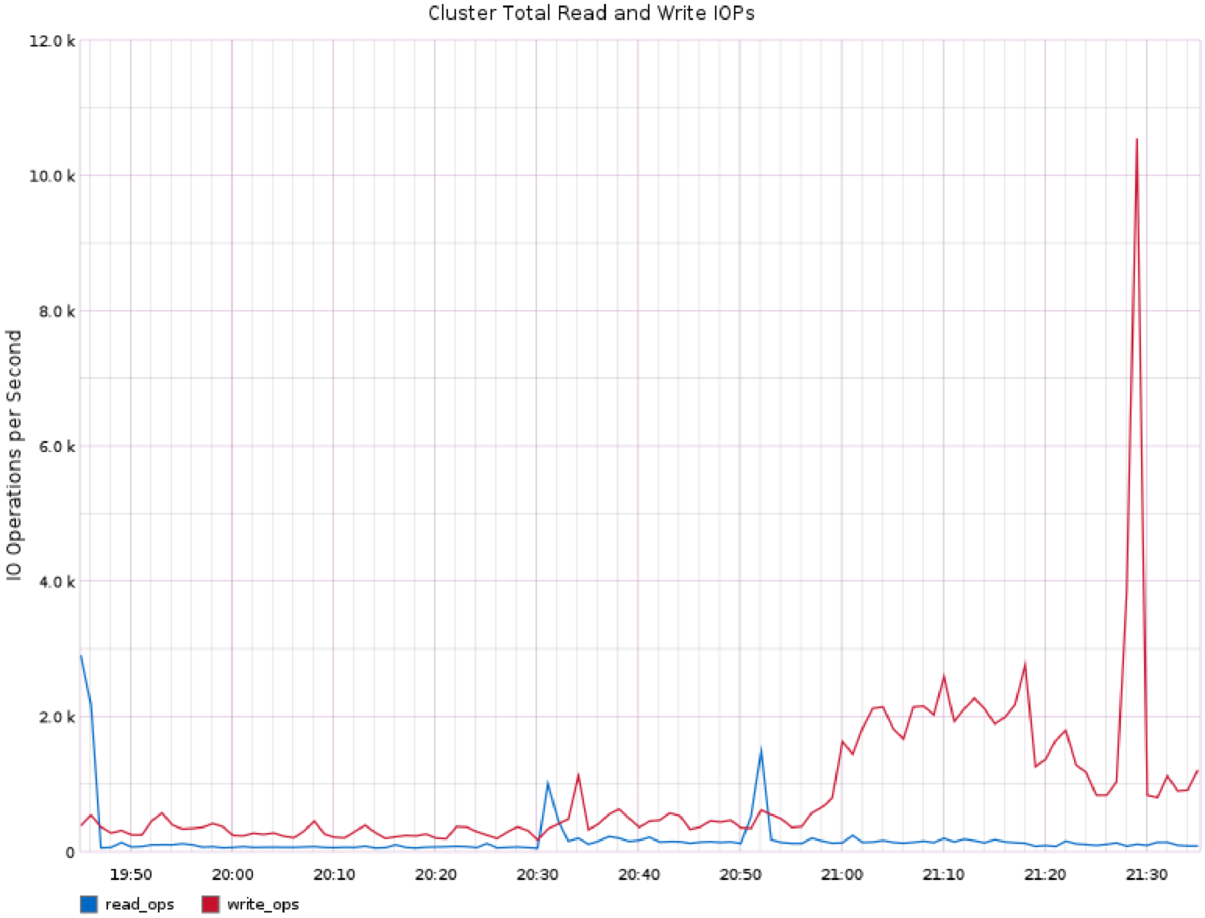

Figure 123 AFF A300 | 1900 Users Citrix HSD (XenApp) Cluster Test | Total Throughput

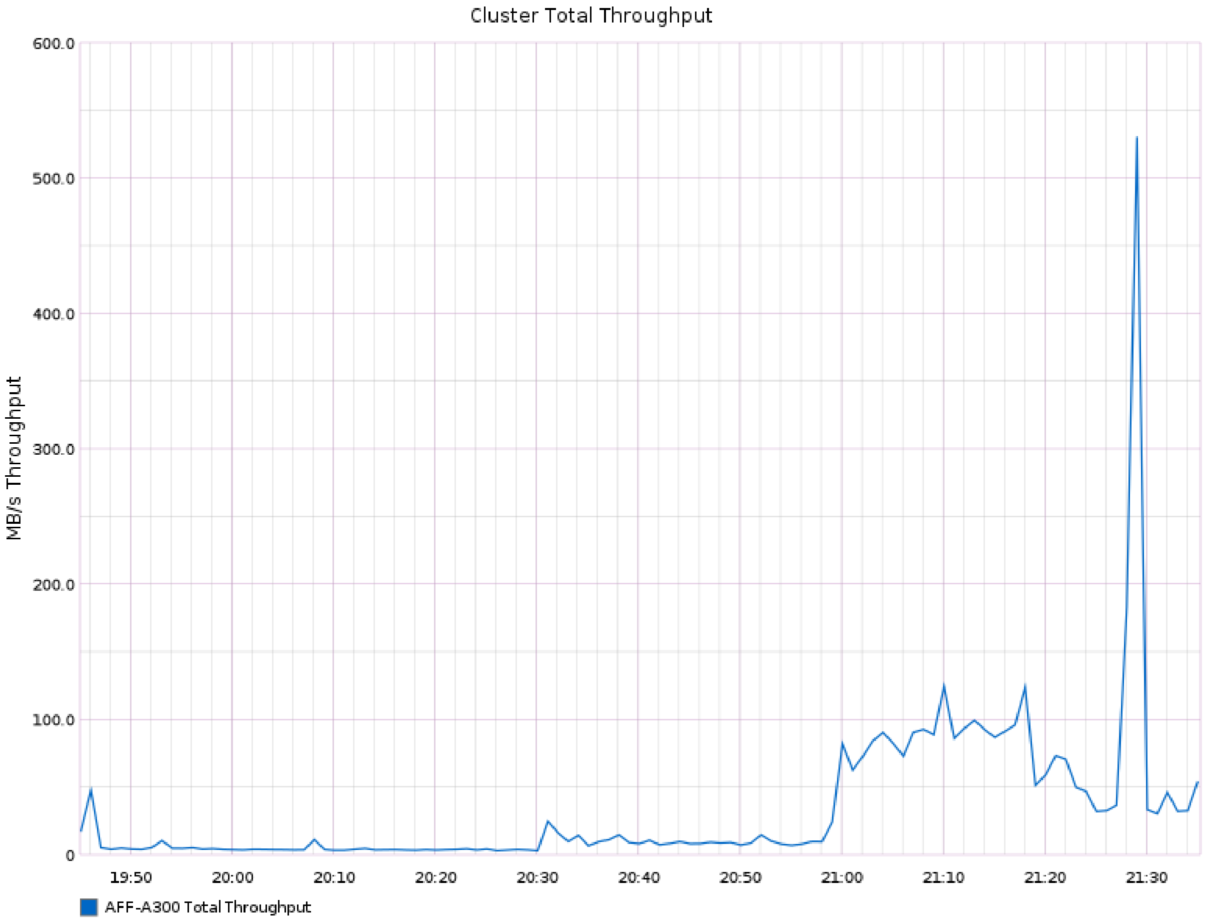

Figure 124 AFF A300 | 1900 Citrix HSD (XenApp) Cluster Test | Read/Write Latency

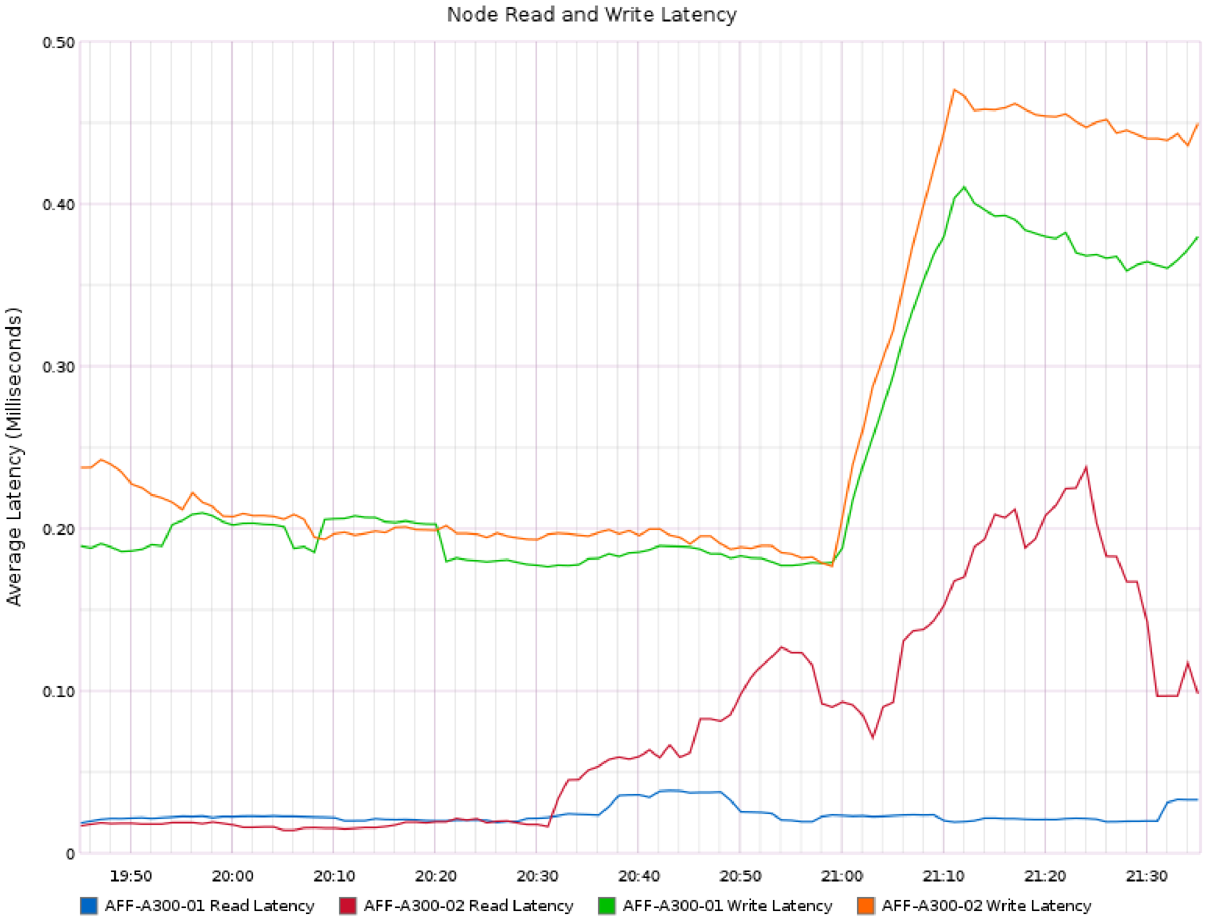

Figure 125 AFF A300 | 1900 Citrix HSD (XenApp) Cluster Test | Controller Average Processor Busy

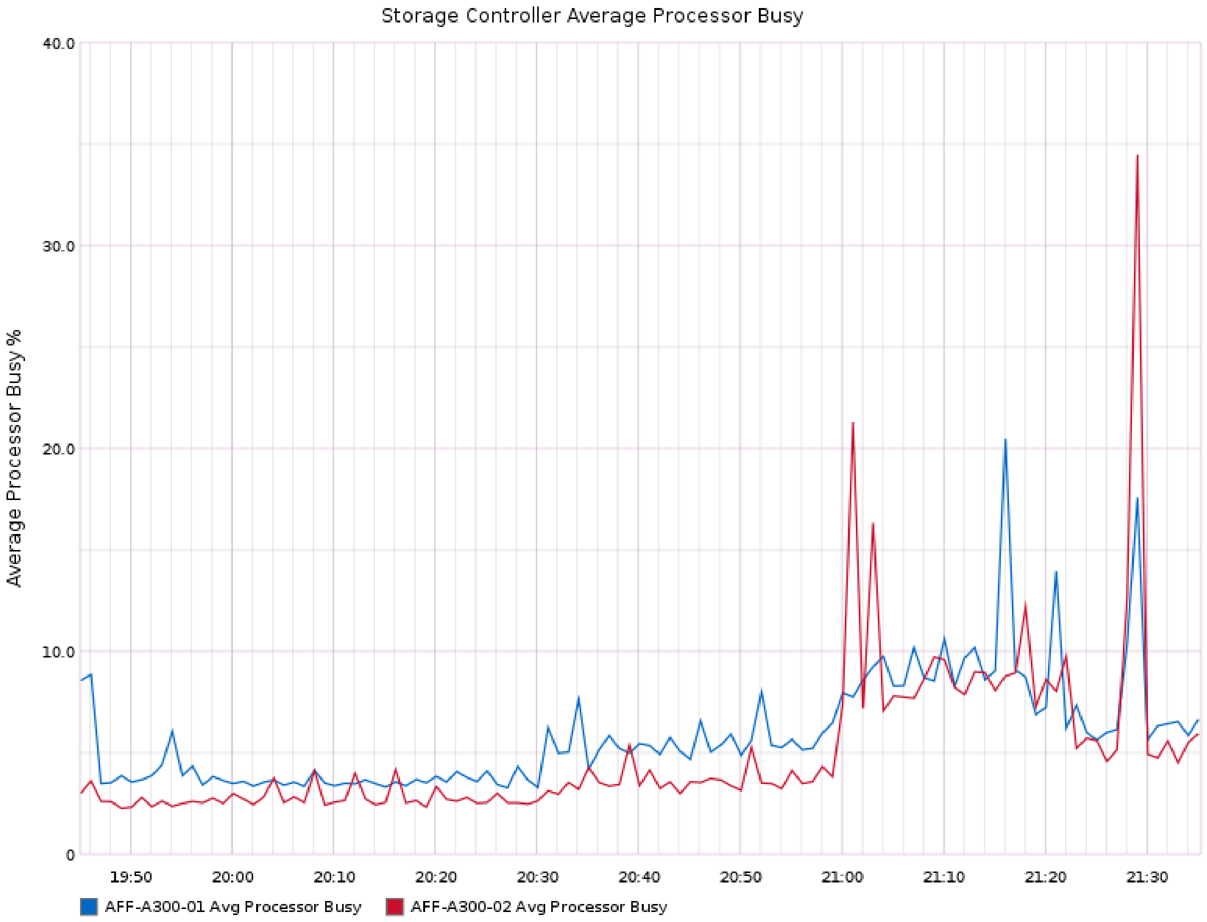

Figure 126 AFF A300 | 1900 Citrix HSD (XenApp) Cluster Test | CPU Headroom

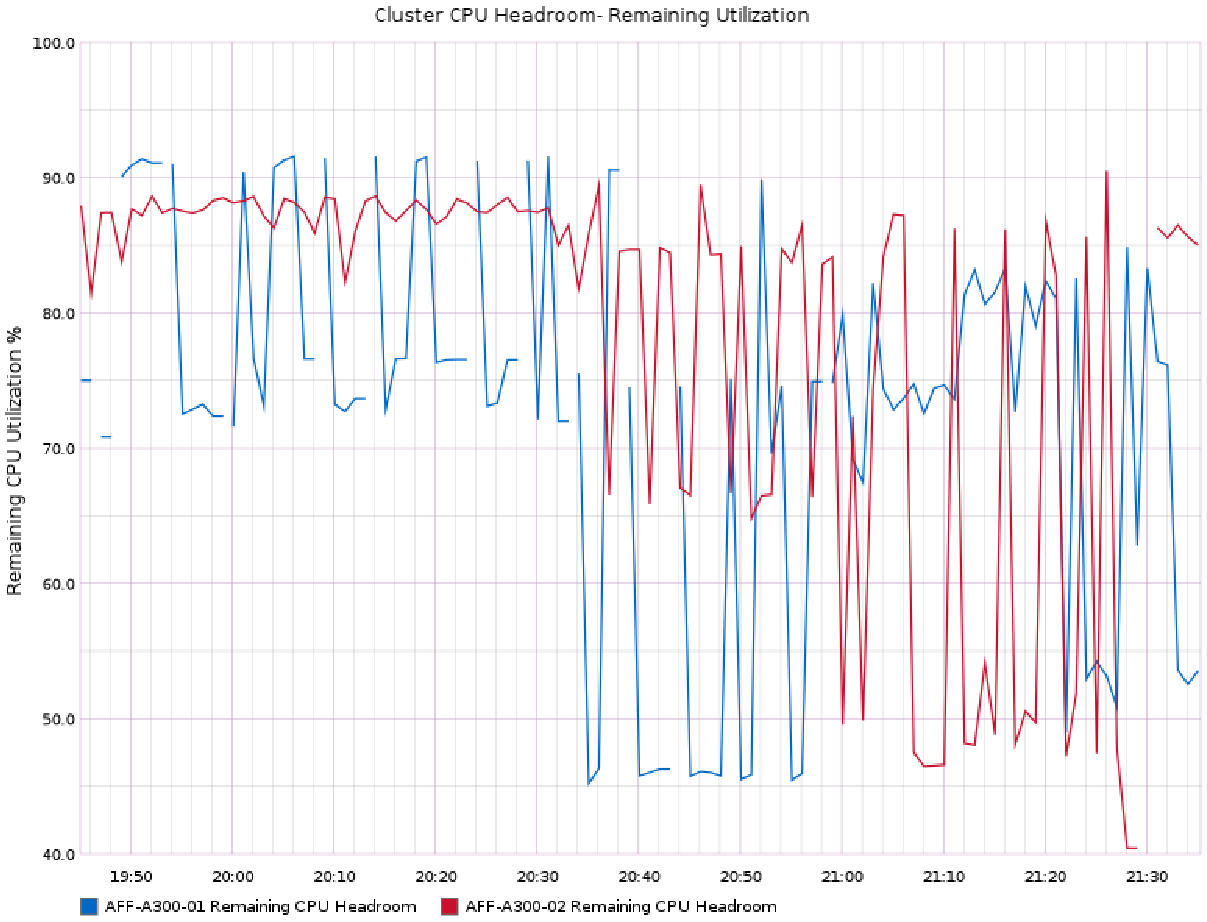

2050 Users Persistent Desktops Cluster Test

NetApp AFF A300 Test Results for 2050 Persistent Windows 10 x64 Citrix MCS Desktops

The next cluster-level simulation was to run 2050 non persistent Windows 10 x64 Citrix MCS desktops against the same AFF- A300. All Login VSI parameters were kept consistent with bringing up all 2050 desktops streamed using Citrix Machine Creation Services (MCS). As indicated by the following storage metrics, the AFF A300 system was clearly able to handle this workload and continued to provide sub milli second latency for another impressive Login VSI result. It has been observed during this component of the testing with observed peak total values 80,00 IOPS. The latency is less than 1ms which results in an excellent end user experience.

Figure 127 NetApp AFF A300 | 2050 Users VDI Persistent Cluster Test | Read/Write IOPS

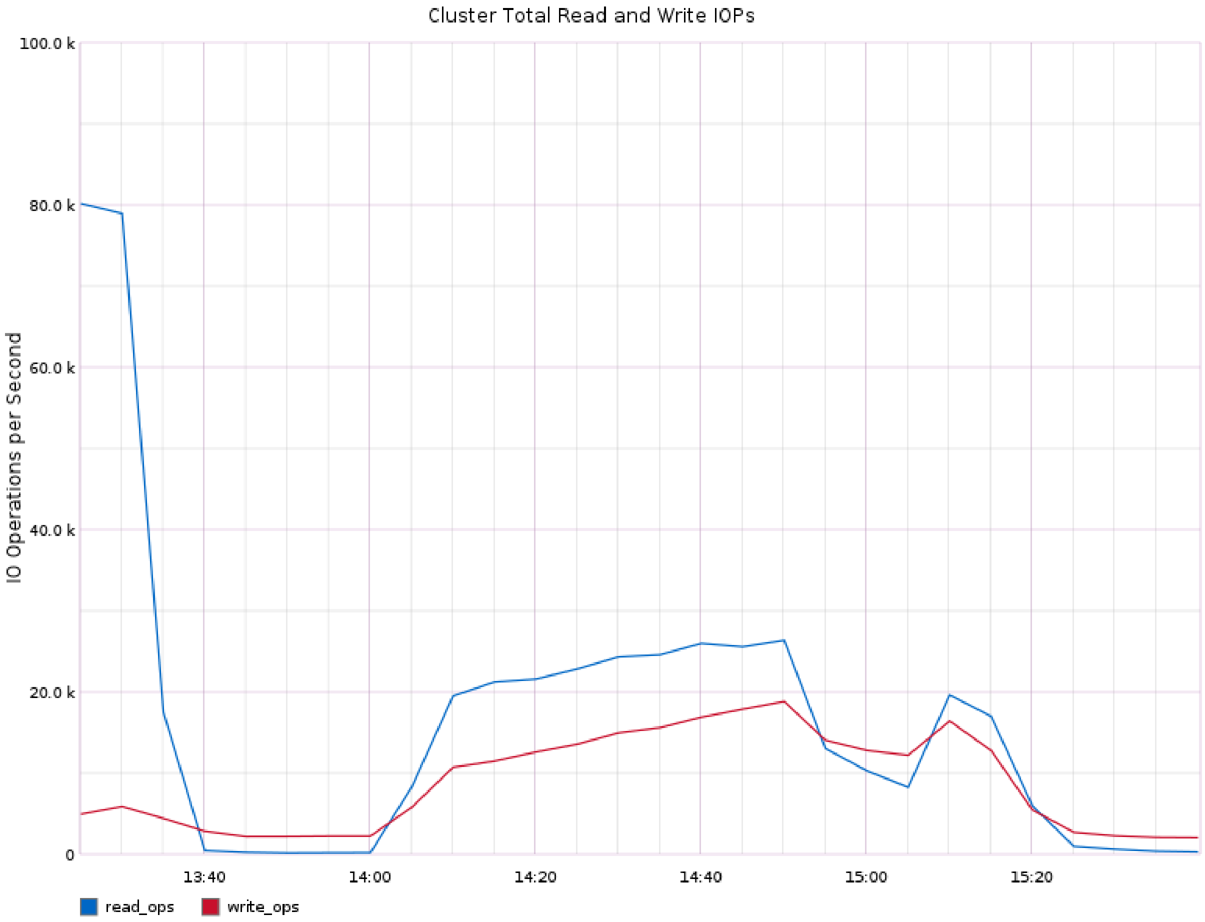

Figure 128 AFF A300 | 2050 Users VDI Persistent Cluster Test | Total Throughput

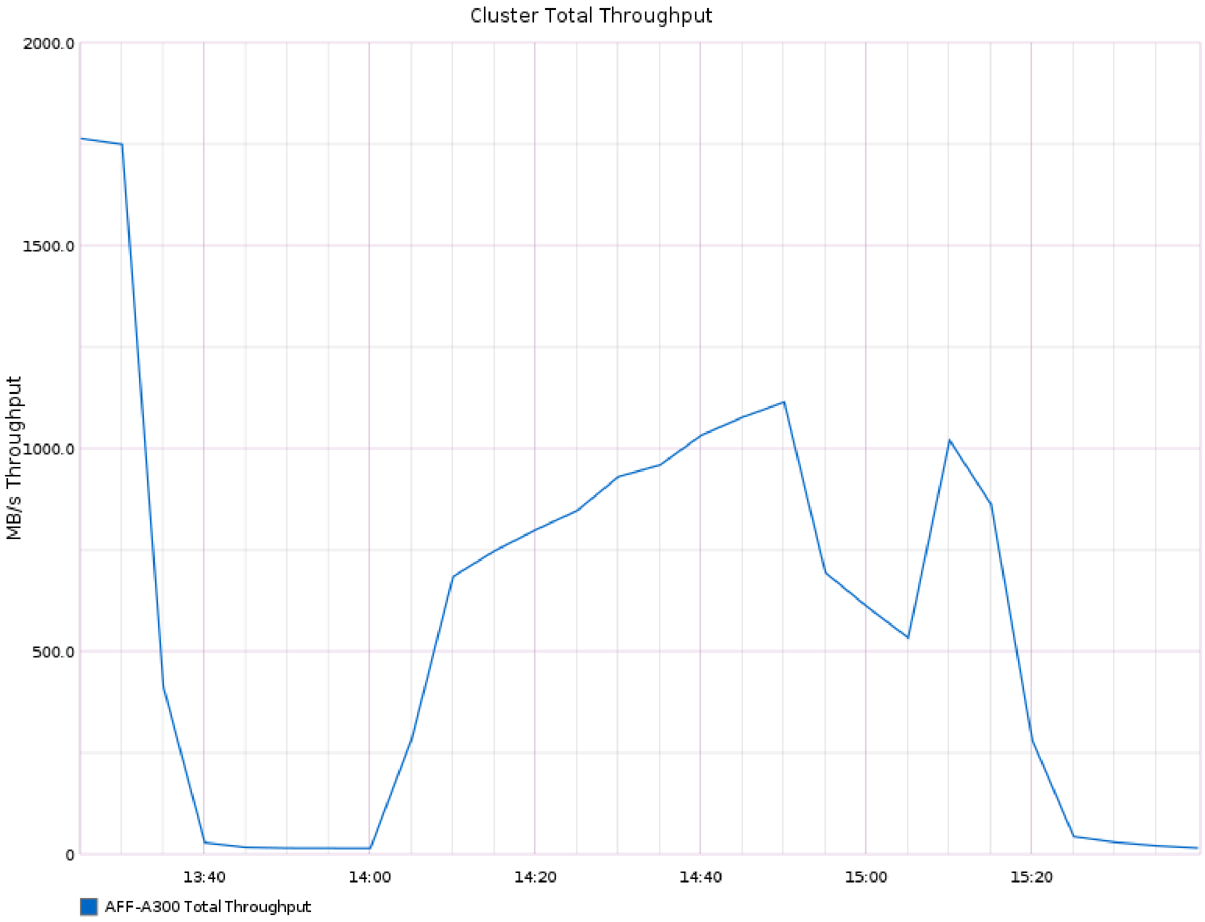

Figure 129 AFF A300 | 2050 Users VDI Persistent Cluster Test | Read/Write Latency (ms)

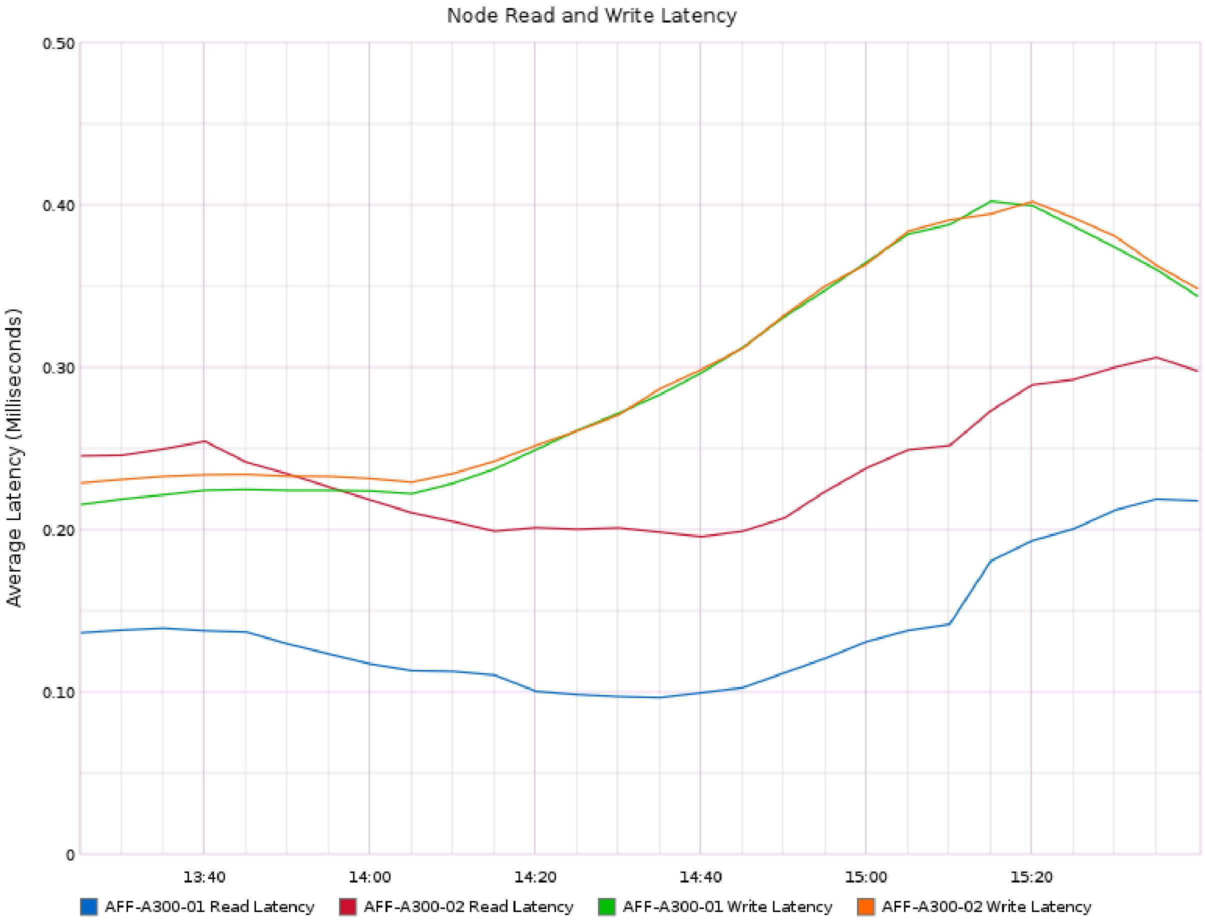

Figure 130 NetApp AFF A300| 2050 Users VDI Persistent Cluster Test | Controller Average Processor Busy

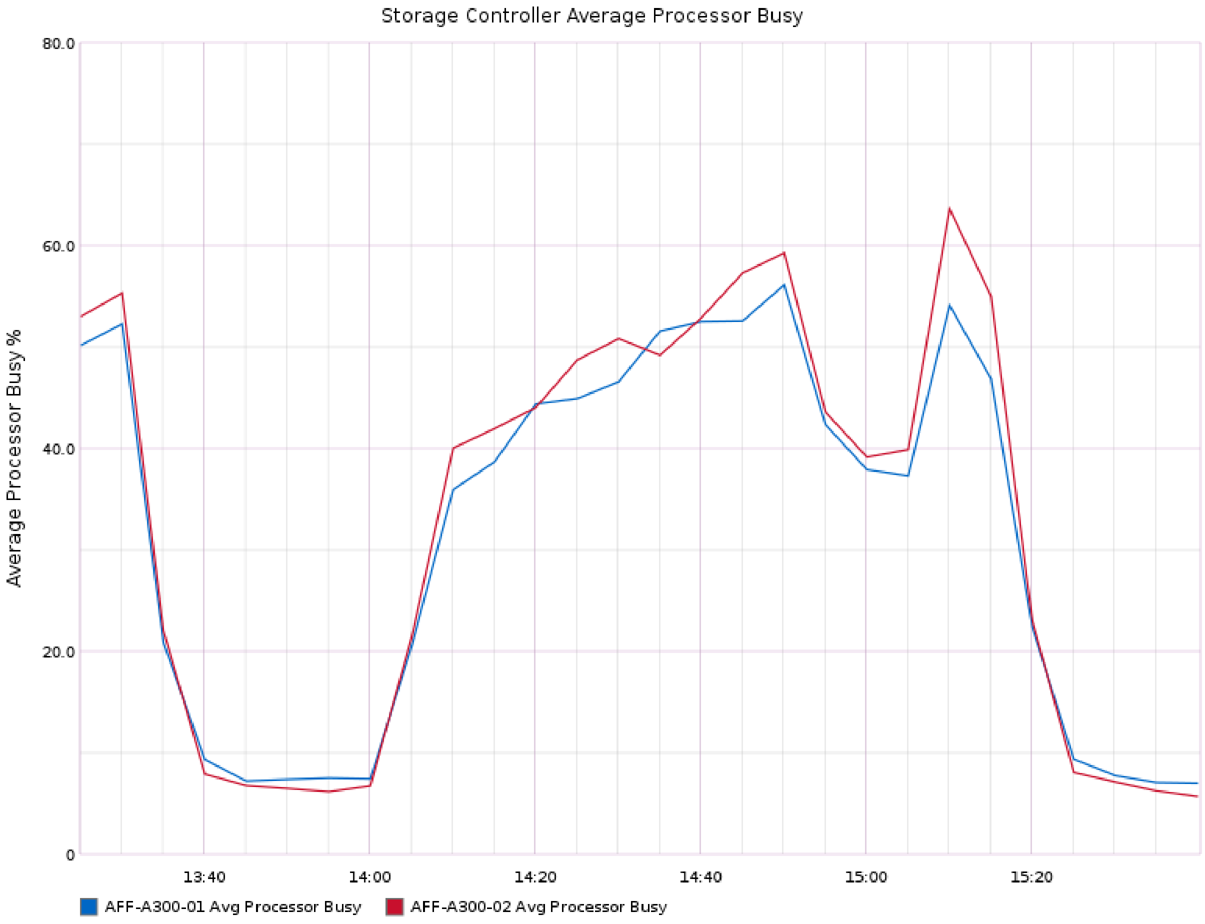

Figure 131 NetApp AFF A300 | 2050 Users VDI Persistent Cluster Test | Controller CPU Headroom

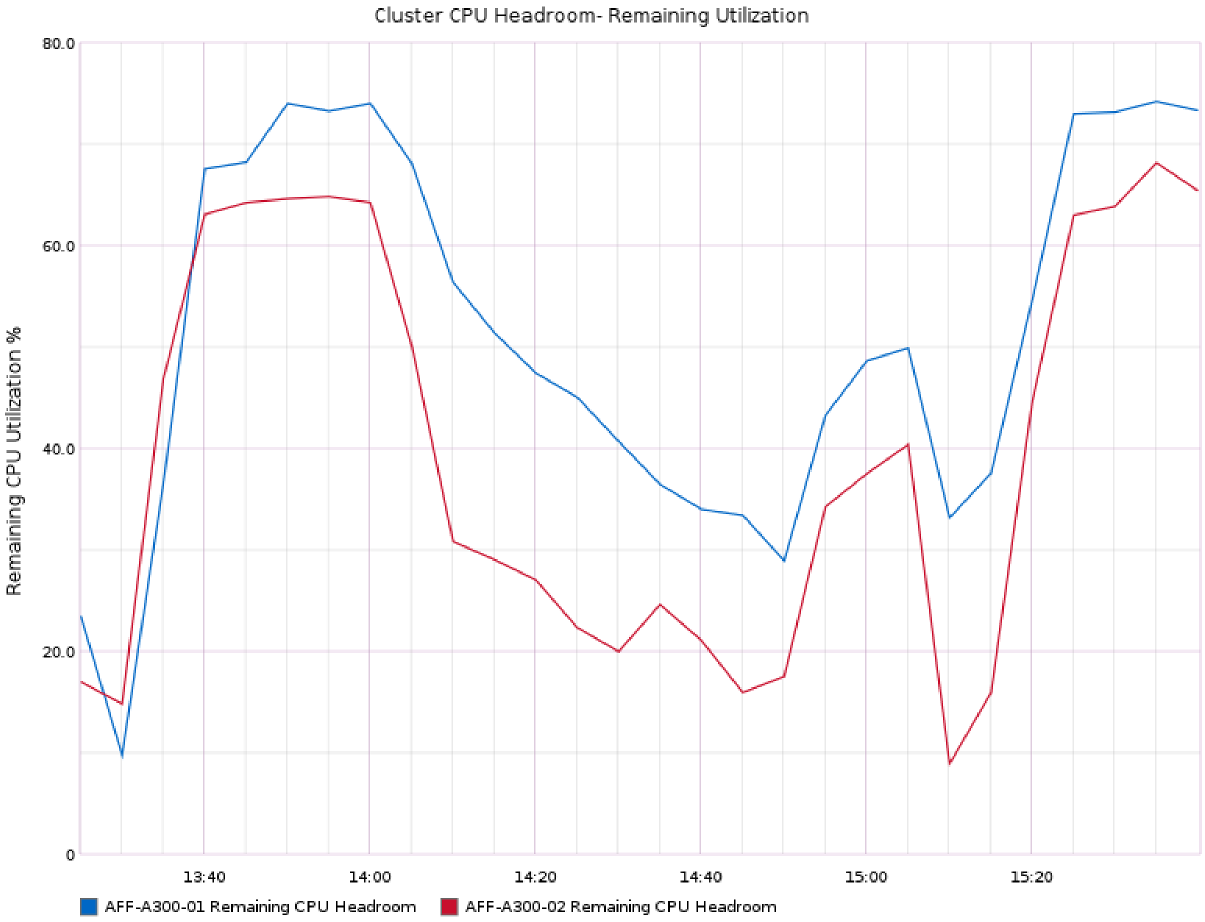

2050 Users PVS Non-Persistent desktops Cluster Test

NetApp AFF A300 Test Results for 2050 PVS Non-Persistent Windows 10 x64 Desktops

The next cluster-level simulation was to run 2050 non persistent Windows 10 x64 Citrix PVS desktops against the same AFF- A300. All Login VSI parameters were kept consistent with bringing up all 2050 desktops streamed using Citrix Provisioning Server (PVS). As indicated by the following storage metrics, the AFF A300 system was clearly able to handle this workload and continued to provide sub milli second latency for another impressive Login VSI result. It has been observed during this component of the testing with observed peak total values 80,00 IOPS. The latency is less than 1ms which results in an excellent end user experience.

Figure 132 NetApp AFF A300 | 2050 Users VDI PVS Non-Persistent Cluster Test | Read/Write IOPS

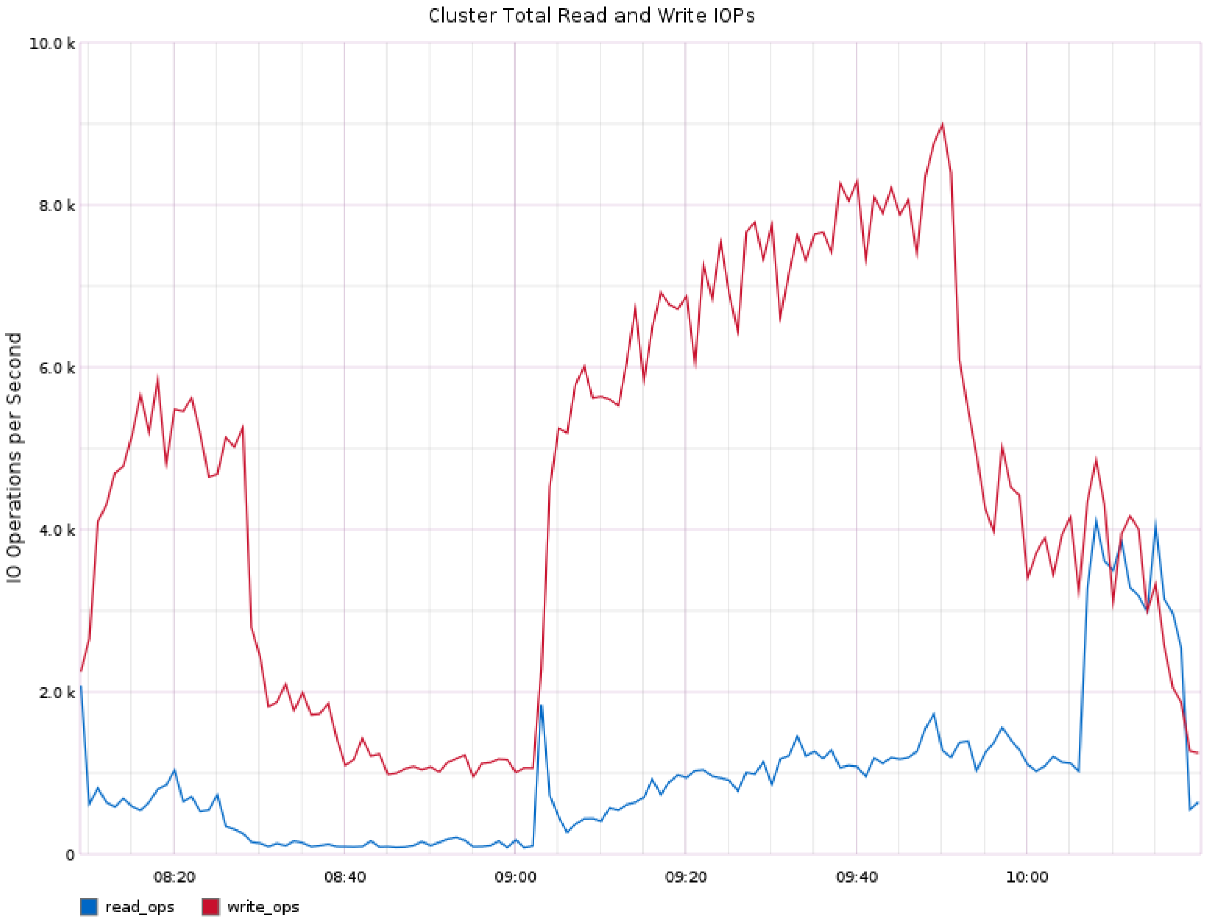

Figure 133 AFF A300 | 2050 Users VDI Persistent Cluster Test | Total Throughput

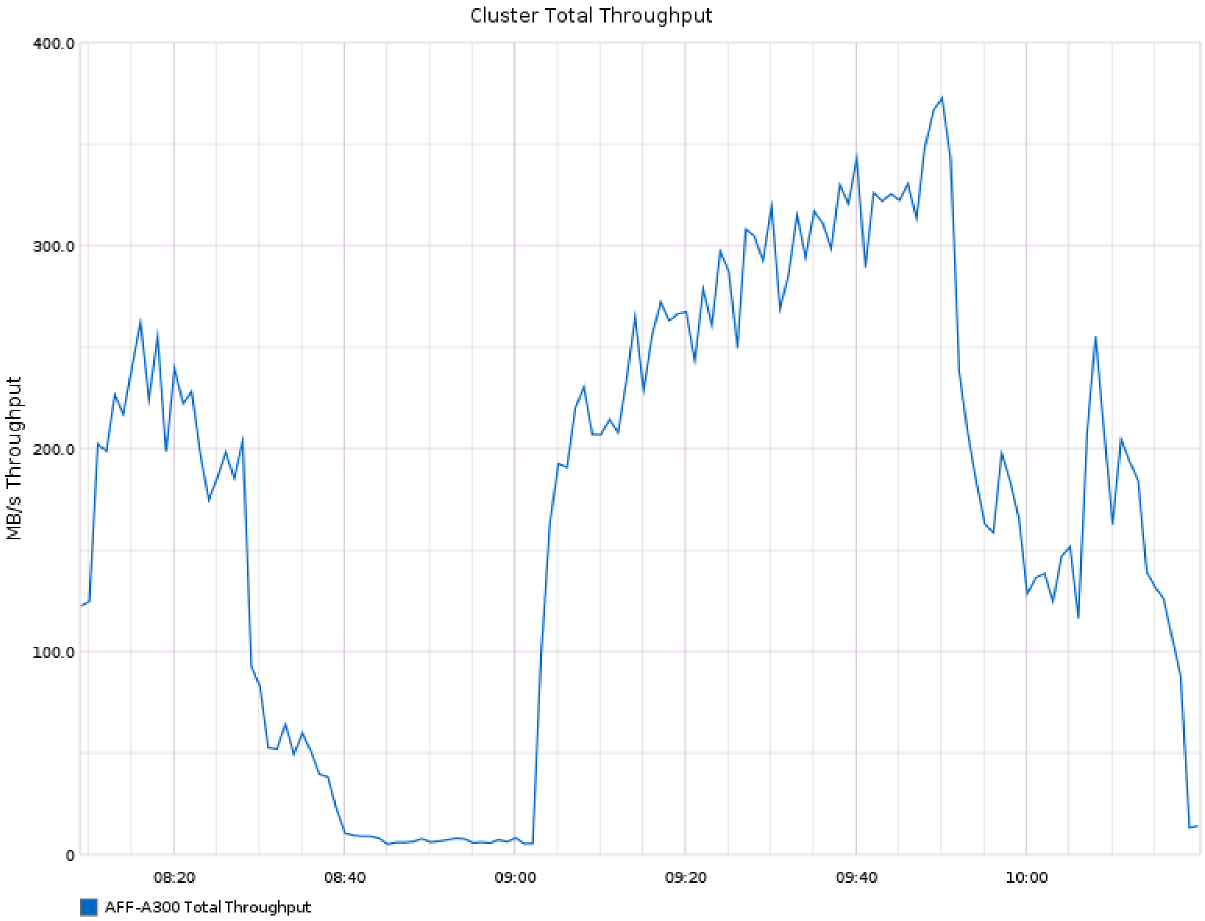

Figure 134 AFF A300 | 2050 Users VDI Persistent Cluster Test | Read/Write Latency (ms)

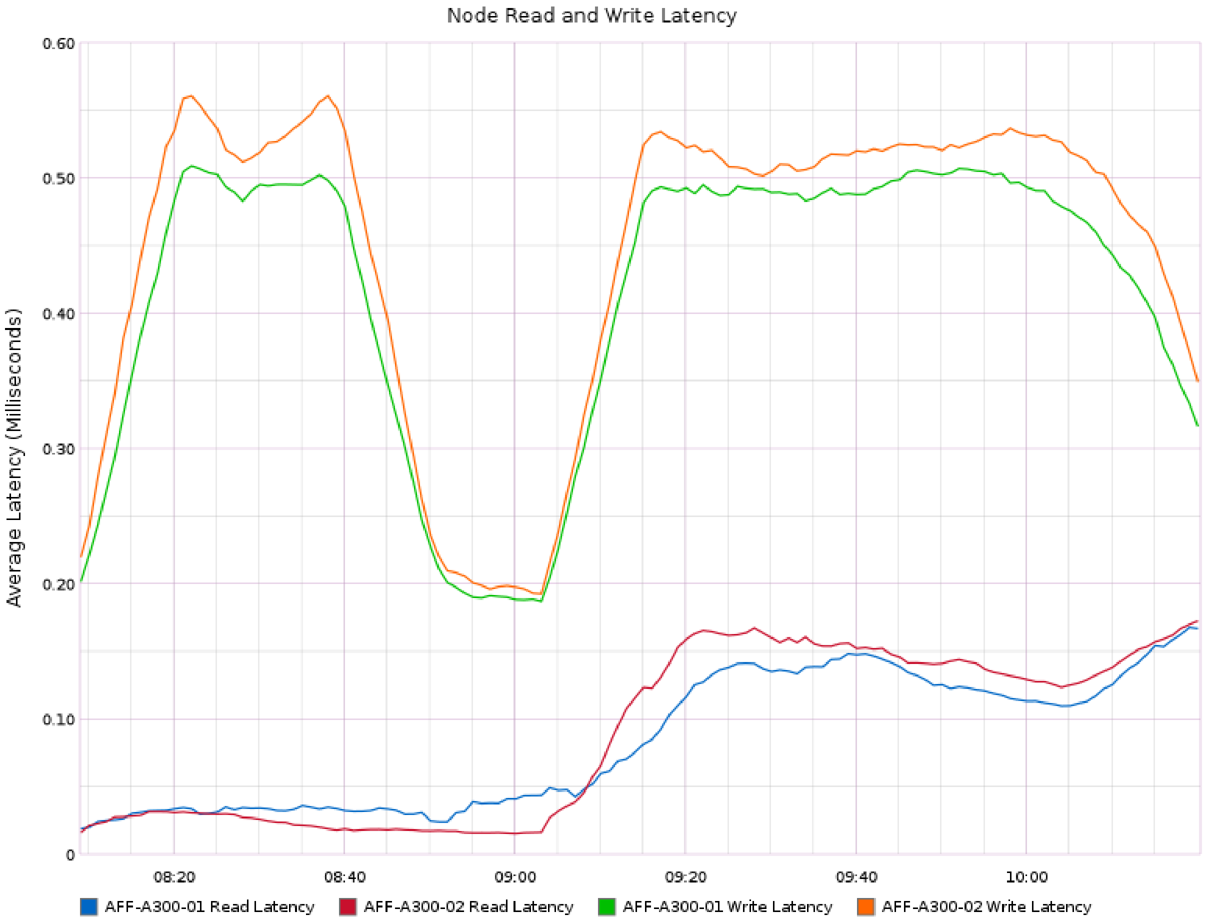

Figure 135 NetApp AFF A300| 2050 Users VDI Persistent Cluster Test | Controller Average Processor Busy

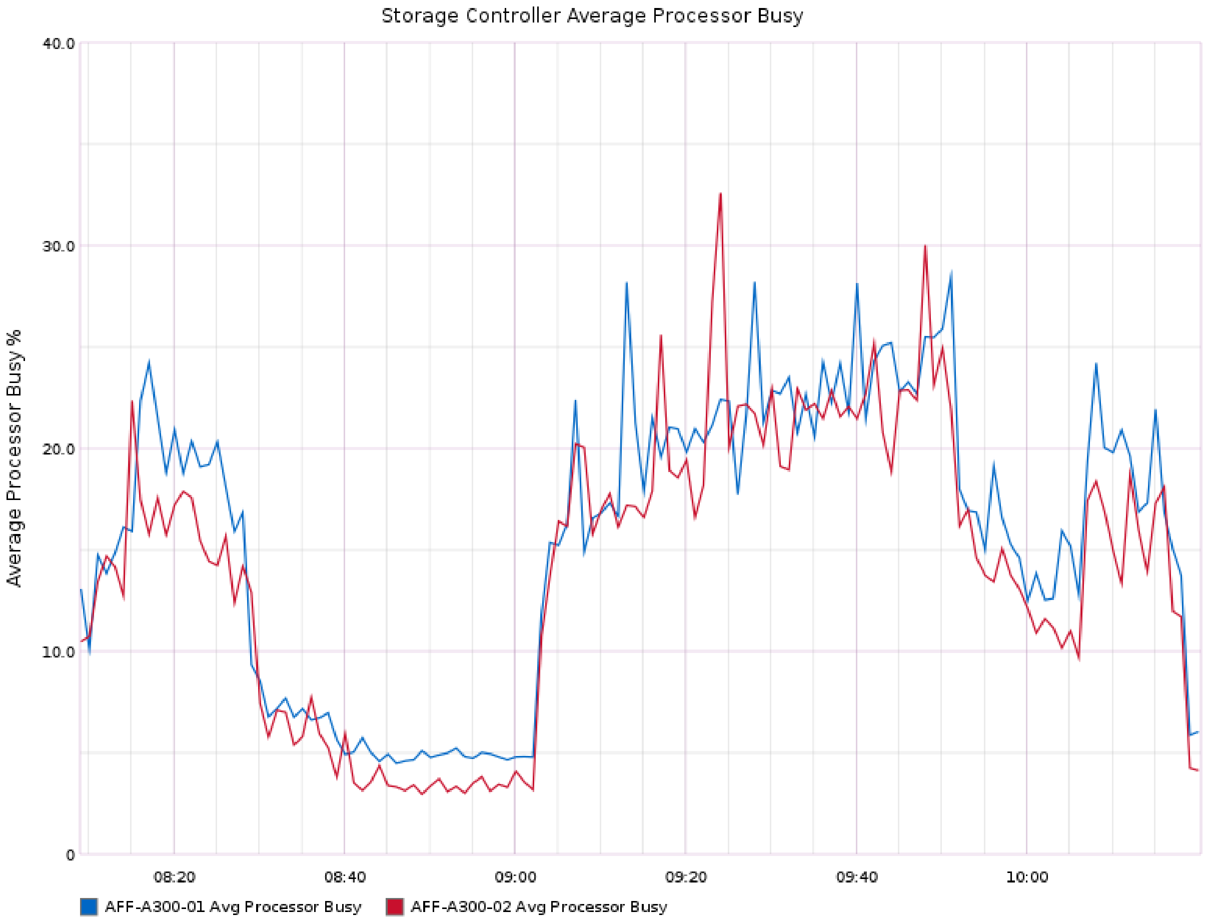

Figure 136 NetApp AFF A300 | 2050 Users VDI Persistent Cluster Test | Controller CPU Headroom

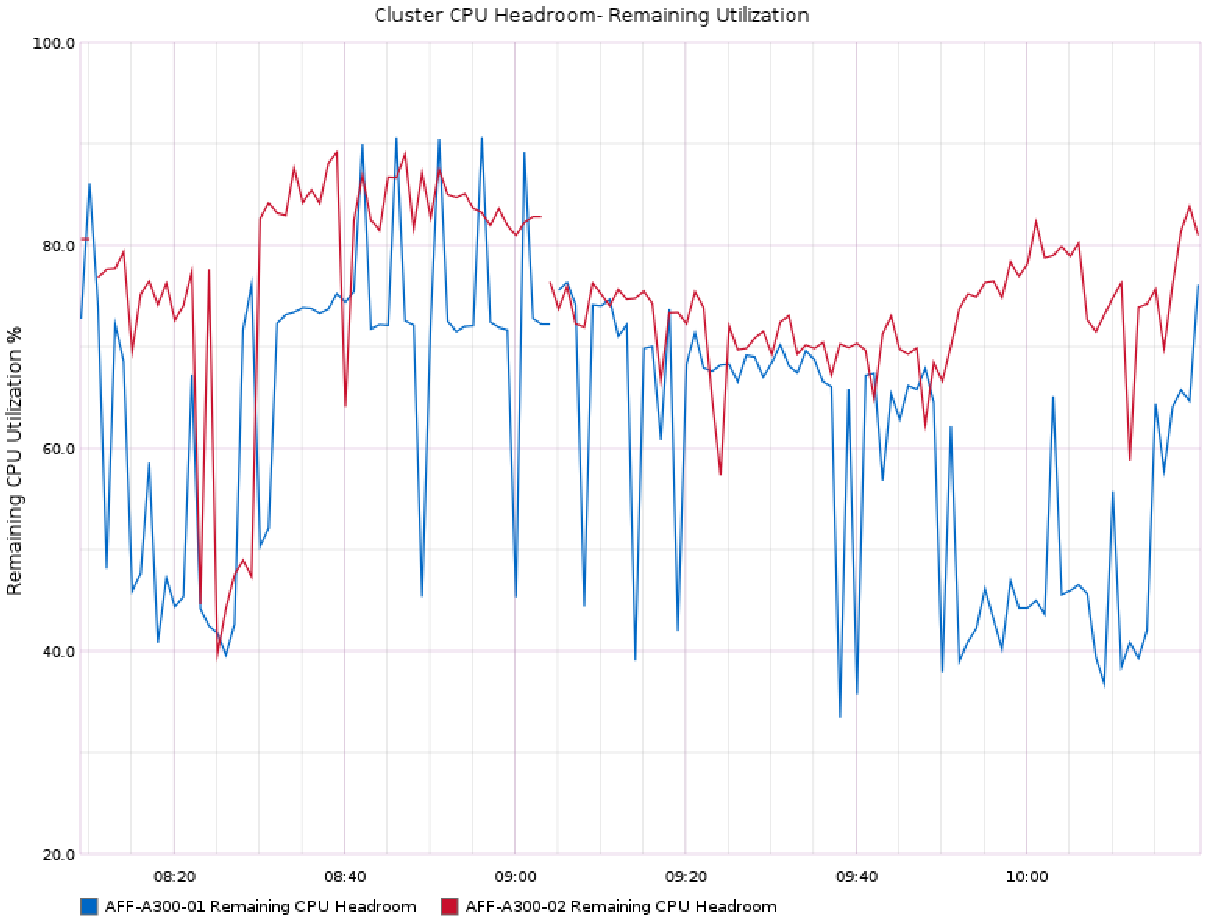

NetApp AFF A300 Storage Test Results for 6000 User Full Scale, Mixed Workload Scalability

The next simulation shows the results of combining all of our previous cluster tests of 1900 Citrix HDS (XenApp) sessions and 2050 Citrix PVS non-persistent VDI desktops and 2050 Full Clone virtual machines sessions for a full 6000 user Citrix XenDesktop VDI simulation on the same AFF A300 array.

The performance and consistent results indicate an outstanding user experience on a very large scale.

The screenshot below shows the AFF A300 GUI with the cursor providing detailed metrics at the start of the simulation during the boot storm of the desktops. Despite driving nearly 2GB/s in bandwidth during the test itself, we maintain the desktop responsiveness and performance of low latency throughout the entire test.

It has been observed during this component of the testing with observed peak total values 83,000 IOPS and a bandwidth of 1750 MB/s. The latency is again less than 1ms, which is an excellent end-user experience on mixed use case scale testing scenario.

Figure 137 AFF A300 | 6000 Users Mixed Workload Full-Scale Test | Total Read/Write IOPS

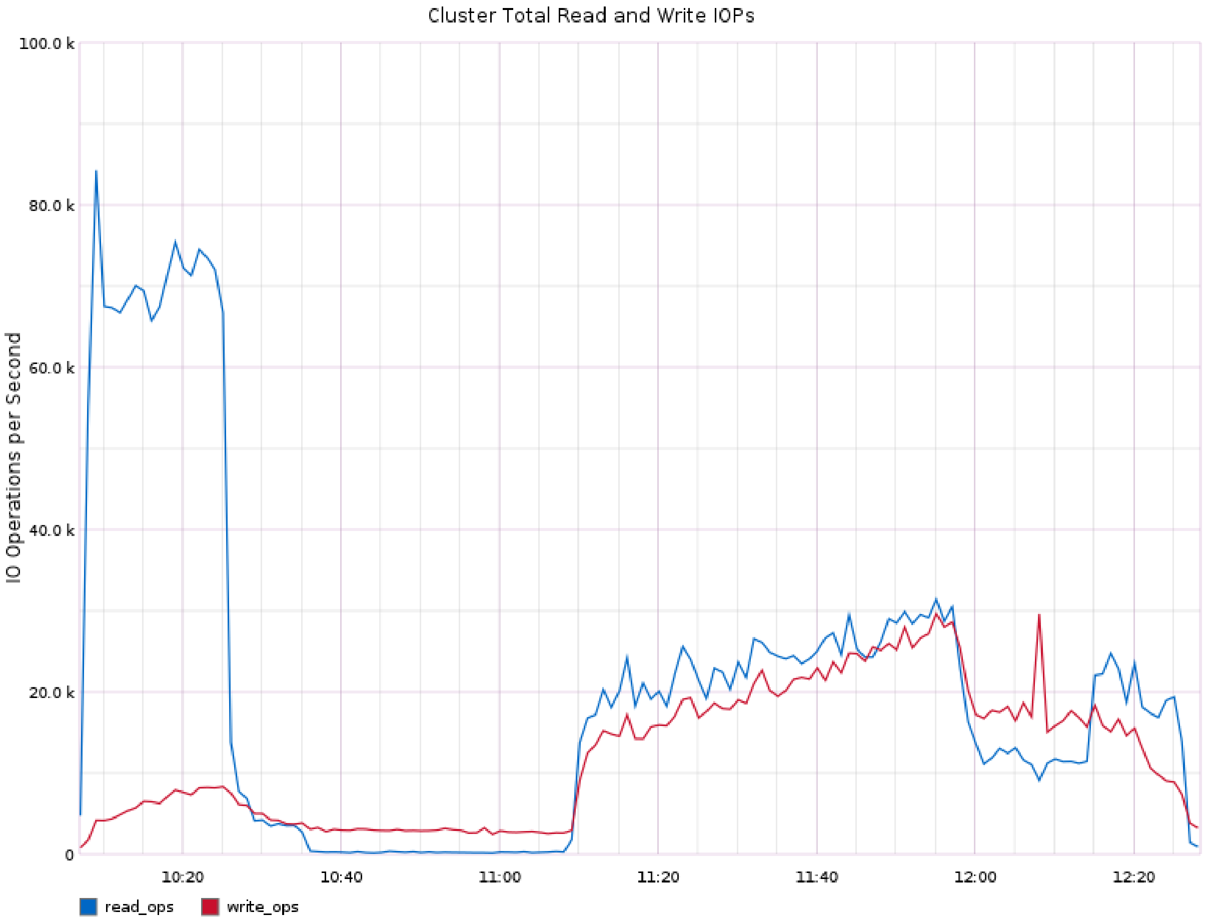

Figure 138 AFF A300 | 6000 Users Mixed Workload Full-Scale Test | Total Throughput

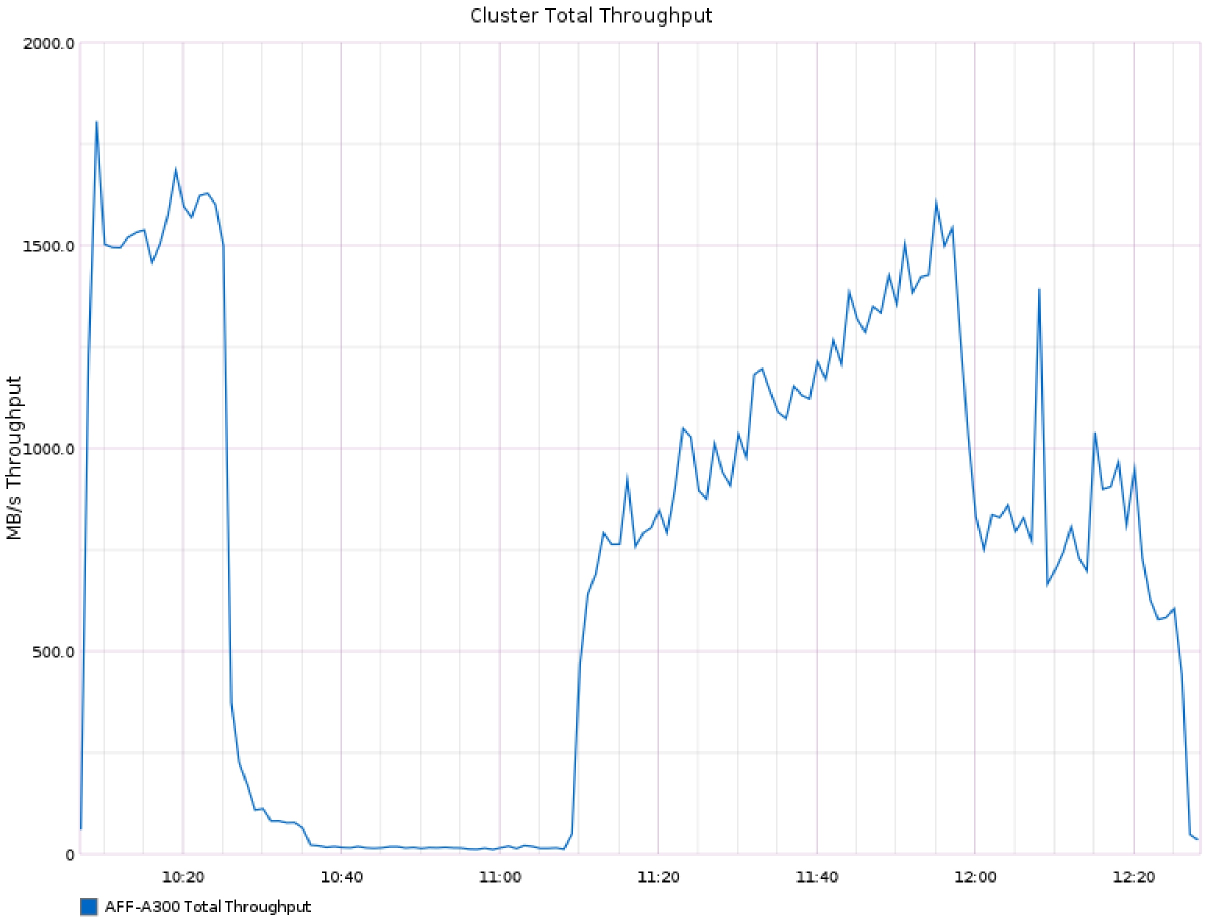

Figure 139 AFF A300 | 6000 Users Mixed Workload Test | Read/Write Latency (ms)

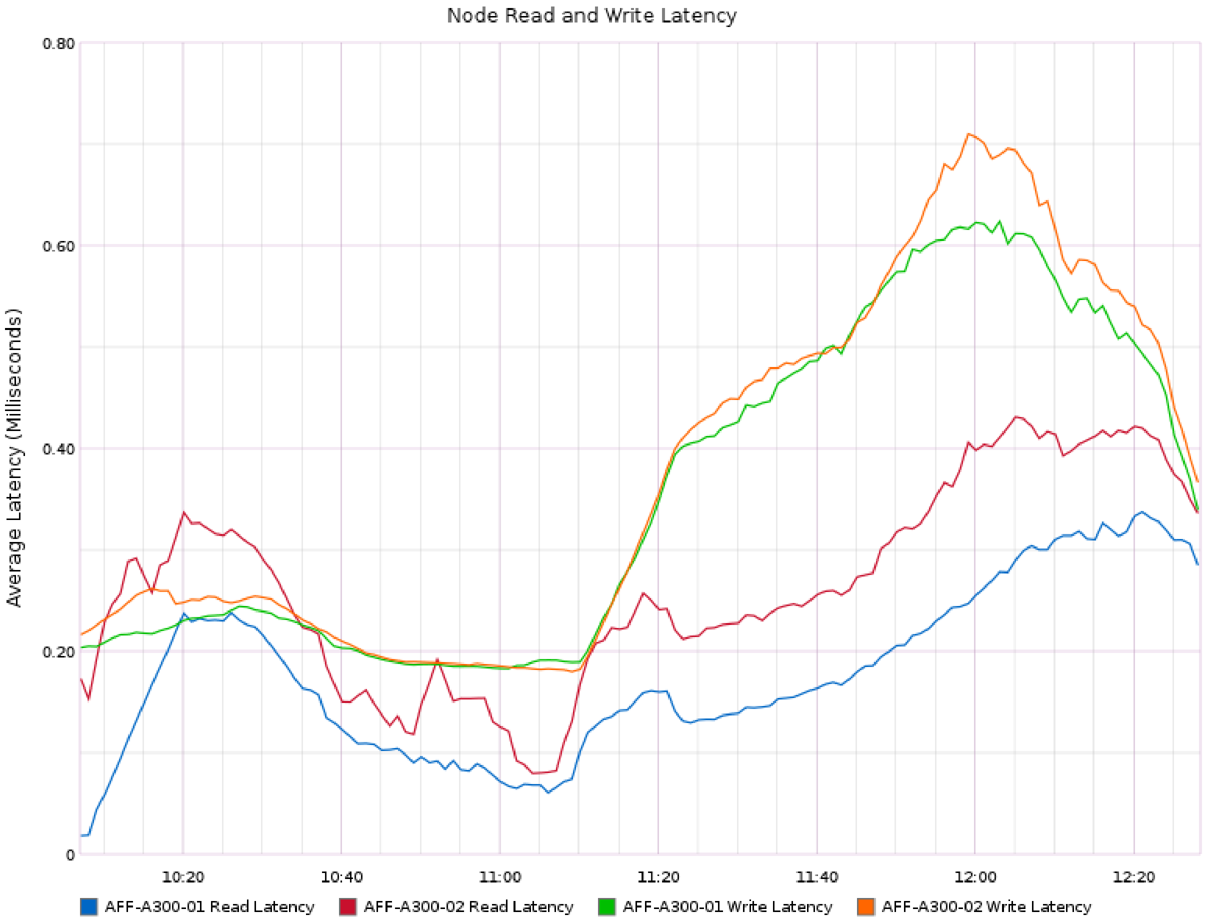

Figure 140 NetApp AFF A300| 6000 Users Mixed Workload Full-Scale Cluster Test | Controller Average Processor Busy

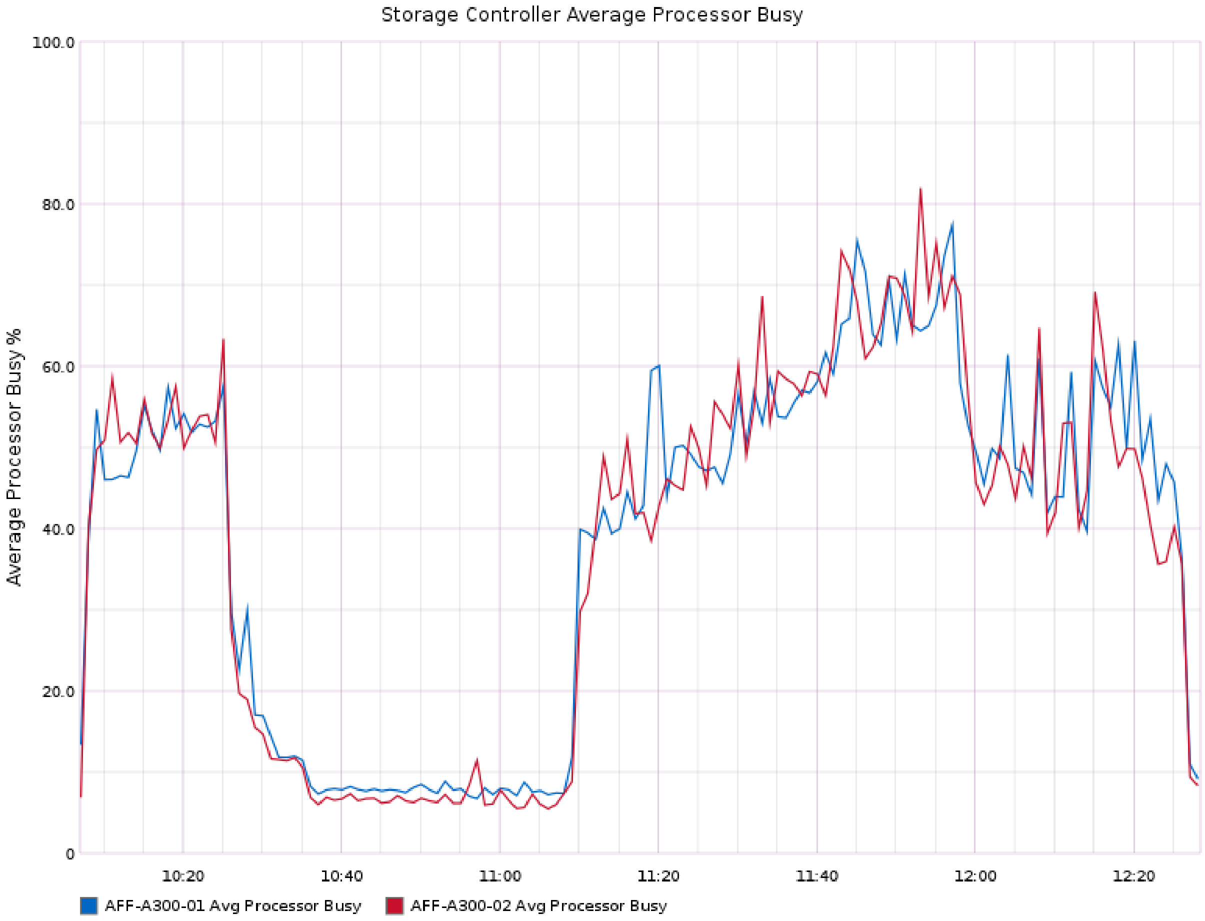

Figure 141 NetApp AFF A300 | 6000 Users Mixed Workload Full-Scale Test | Controller CPU Headroom

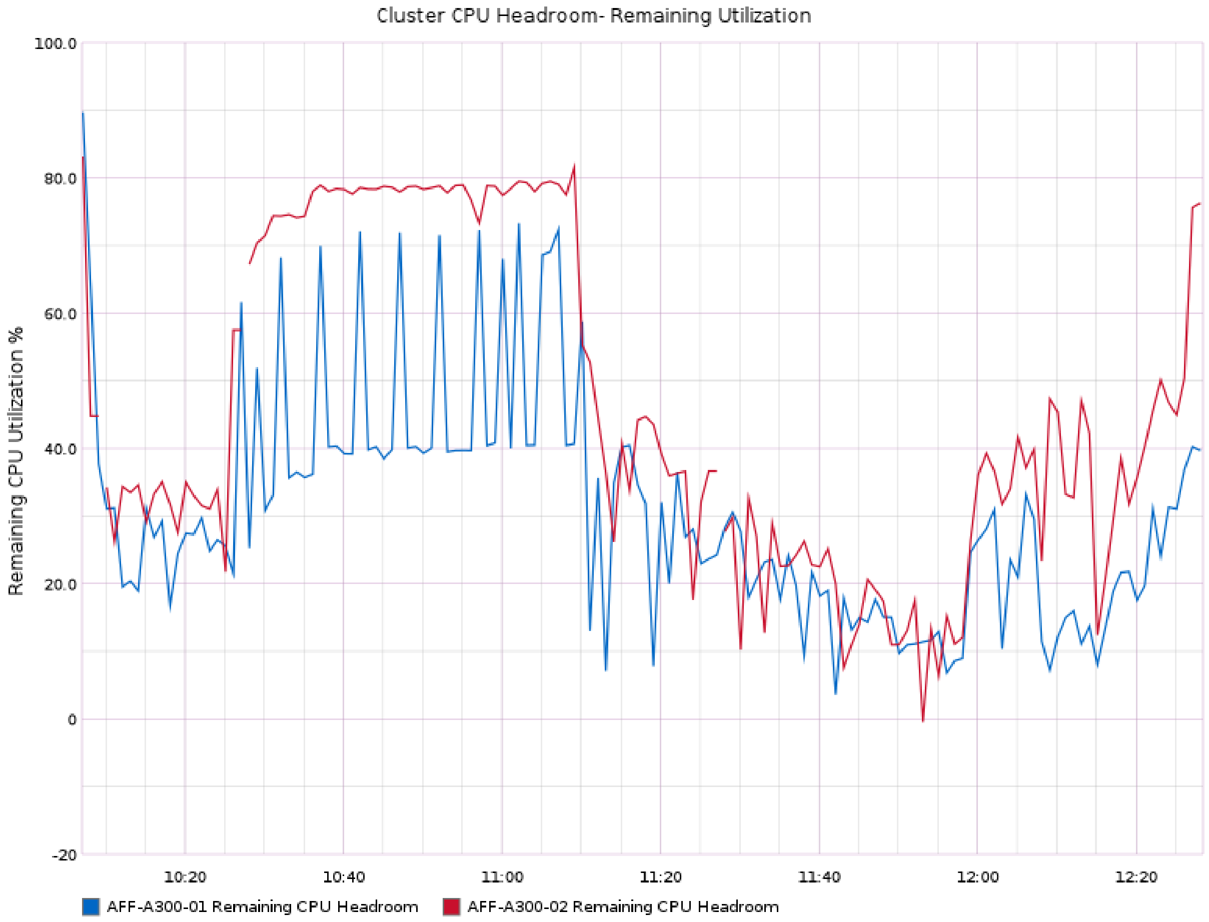

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 6000 Users, which this reference architecture has successfully tested. This 6000-seat solution provides a large-scale building block that can be replicated to confidently scale-out to tens of thousands of users.

Cisco UCS System Scalability

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested.

- Cisco UCS Manager Software supports up to 20 Cisco UCS chassis within a single Cisco UCS domain with Cisco UCS 6332UP Fabric Interconnect. A single UCS domain can grow to 160 blades for an enterprise deployment.
- Cisco UCS Central, the manager of managers, extends UCS domains and vastly increases the reach of the Cisco UCS system. Simplify daily operations by centrally managing and automating routine tasks and expediting problem resolution. Our powerful platform eliminates disparate management environments. Use it to support up to 10,000 Cisco UCS servers (blade, rack, composable, and Mini) and manage multiple Cisco UCS instances or domains across globally-distributed locations.
- As scale grows, the value of the combined UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100 percent of the time.

- To accommodate the Cisco Nexus 9000 upstream connectivity in the way we describe in the network configuration section, two Ethernet uplinks are needed to be configured on the Cisco UCS 6332UP Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp scaling section. Please refer the NetApp section that follows this one for scalability guidelines.

NetApp FAS Storage Guidelines for Mixed Desktop Virtualization Workloads

Storage sizing has three steps:

- Gathering solution requirements
- Estimating storage capacity and performance
- Obtaining recommendations for the storage configuration

Solution Assessment

Assessment is an important first step. Liquidware Labs Stratusphere FIT and Lakeside VDI Assessment are recommended to collect network, server, and storage requirements. NetApp has contracted with Liquidware Labs to provide free licenses to NetApp employees and channel partners. For information on how to obtain software and licenses, refer to this FAQ. Liquidware Labs also provides a storage template that fits the NetApp system performance modeler. For guidelines on how to use Stratusphere FIT and the NetApp custom report template, refer to TR-3902: Guidelines for Virtual Desktop Storage Profiling.

Virtual desktop sizing depends on the following:

- The number of the seats
- The VM workload (applications, VM size, and VM OS)
- The connection broker (Citrix XenDesktop)
- The hypervisor type (vSphere, XenServer, or Hyper-V)
- The provisioning method (NetApp clone, Linked clone, PVS, and MCS)
- Future storage growth
- Disaster recovery requirements
- User home directories

NetApp has developed a sizing tool called the System Performance Modeler (SPM) that simplifies the process of performance sizing for NetApp systems. It has a step-by-step wizard to support varied workload requirements and provides recommendations for meeting your performance needs.

Storage sizing has two factors: capacity and performance. NetApp recommends using the NetApp SPM tool to size the virtual desktop solution. To use this tool, contact NetApp partners and NetApp sales engineers who have the access to SPM. When using the NetApp SPM to size a solution, NetApp recommends separately sizing the VDI workload (including the write cache and personal vDisk if used), and the CIFS profile and home directory workload. When sizing CIFS, NetApp recommends sizing with a heavy user workload. Eighty percent concurrency was assumed in this solution.

Capacity Considerations

Deploying XenDesktop with PVS imposes the following capacity considerations:

- vDisk. The size of the vDisk depends on the OS and the number of applications installed. It is a best practice to create vDisks larger than necessary in order to leave room for any additional application installations or patches. Each organization should determine the space requirements for its vDisk images.
- As an example, a 20GB vDisk with a Windows 10 image is used. NetApp deduplication can be used for space savings.
- Write cache file. NetApp recommends a size range of 4 to 18GB for each user. Write cache size is based on what type of workload and how often the VM is rebooted. In this example, 4GB is used for the write-back cache. Since NFS is thin provisioned by default, only the space currently used by the VM will be consumed on the NetApp storage. If iSCSI or FCP is used, N x 4GB would be consumed as soon as a new virtual machine is created.
- PvDisk. Normally, 5 to 10GB is allocated, depending on the application and the size of the profile. Use 20 percent of the master image as a starting point. NetApp recommends running deduplication.
- CIFS home directory. Various factors must be considered for each home directory deployment. The key considerations for architecting and sizing a CIFS home directory solution include the number of users, the number of concurrent users, the space requirement for each user, and the network load. Run deduplication to obtain space savings.
- Infrastructure. Host XenDesktop, PVS, SQL Server, DNS, and DHCP.

The space calculation formula for a 2000-seat deployment is as follows:

> Number of vDisk x 20GB + 2000 x 4GB write cache + 2000 x 10GB PvDisk + 2000 x 5GB user home directory x 70% + 2000 x 1GB vSwap + 500GB infrastructure

Performance Considerations

The collection of performance requirements is a critical step. After using Liquidware Labs Stratusphere FIT and Lakeside VDI Assessment to gather I/O requirements, contact the NetApp account team to obtain recommended software and hardware configurations.

Size, the read/write ratio, and random or sequential reads comprise the I/O considerations. We use 90 percent write and 10 percent read for PVS workload. Storage CPU utilization must also be considered. Table 46   can be used as guidance for your sizing calculations for a PVS workload when using a LoginVSI heavy workload.

Table 46   Typical IOPS Without RamCache Plus Overflow Feature

	Boot IOPS	Login IOPS	Steady IOPS
Write Cache (NFS)	8-10	9	7.5
vDisk (CIFS SMB 3)	0.5	0	0
Infrastructure (NFS)	2	1.5	0

Scalability of Citrix XenDesktop 7.15 Configuration

XenDesktop environments can scale to large numbers. When implementing Citrix XenDesktop, consider the following in scaling the number of hosted shared and hosted virtual desktops:

- Types of storage in your environment
- Types of desktops that will be deployed
- Data protection requirements
- For Citrix Provisioning Server pooled desktops, the write cache sizing and placement

"Citrix VDI Handbook and Best Practices XenApp XenDesktop 7.15 LTSR v2.10" document was used in this solution to determine various aspects of scalability and configuration of the XenDesktop components.

When designing and deploying this CVD environment Cisco and Citrix recommend using N+1 schema for virtualization host servers to accommodate resiliency. In all Reference Architectures (such as this CVD), this recommendation is applied to all host servers.

Appendix A Cisco Switch Configuration

Network Configuration

N93180YC-FX -A Configuration

```
!Command: show running-config

!Time: Sun Jun 24 13:46:03 2018


version 7.0(3)I7(2)

switchname DV-Pod-2-N9K-A

class-map type network-qos class-platinum

match qos-group 2

class-map type network-qos class-all-flood

match qos-group 2

class-map type network-qos system_nq_policy

match qos-group 2

class-map type network-qos class-ip-multicast

match qos-group 2

policy-map type network-qos jumbo

  class type network-qos class-platinum

    mtu 9216

  class type network-qos class-default

    mtu 9216

vdc DV-Pod-2-N9K-A id 1

  limit-resource vlan minimum 16 maximum 4094

  limit-resource vrf minimum 2 maximum 4096

  limit-resource port-channel minimum 0 maximum 511

  limit-resource u4route-mem minimum 248 maximum 248

  limit-resource u6route-mem minimum 96 maximum 96

  limit-resource m4route-mem minimum 58 maximum 58

  limit-resource m6route-mem minimum 8 maximum 8
```

```
feature telnet
feature nxapi
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock timezone PST -7 0

no password strength-check
username admin password 5 $1$tYYajkfc$7P7nLjWYvfTWAlvFDnwJZ.  role network-admin
ip domain-lookup
system default switchport
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-A_64
  10 permit ip 10.10.64.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-B_65
  10 permit ip 10.10.65.0 255.255.255.0 any
  20 deny ip any any
class-map type qos match-any class-platinum
  match cos 5
policy-map type qos jumbo
  class class-platinum
    set qos-group 2
  class class-default
```

```
    set qos-group 0
system qos
  service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0xa075be936e36177e1912888e7aed3223
 priv 0xa075be936e36177e1912888e7aed3223 localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 10.10.160.2 use-vrf default
ntp peer 10.10.160.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8

vlan 1-2,60-70,102,164,264
vlan 60
  name In-Band-Mgmt
vlan 61
  name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name NFS
vlan 64
  name iSCSI-A
vlan 65
  name iSCSI-B
vlan 66
  name vMotion
```

```
vlan 67
  name N1KV
vlan 68
  name LauncherPXE
vlan 69
  name Launcher81
vlan 70
  name other-3
vlan 102
  name VDI
vlan 164
  name Out-Of-Band-Mgmt
vlan 264
  name OOB-Mgmt

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.66 source 10.29.164.65
  delay restore 150
  peer-gateway
  auto-recovery
```

```
interface Vlan1
  no shutdown
  no ip redirects
  ip address 10.29.164.2/24
  no ipv6 redirects

interface Vlan2
  description Default native vlan 2
  no ip redirects
  no ipv6 redirects

interface Vlan60
  description Out of Band Management vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 60
    preempt
    priority 110
    ip 10.10.60.1

interface Vlan61
  description Infrastructure vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
```

```
    preempt
    ip 10.10.61.1

interface Vlan62
  description CIFS vlan 62
  no shutdown
  no ip redirects
  ip address 10.10.62.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 62
    preempt
    priority 110
    ip 10.10.62.1

interface Vlan63
  description NFS vlan 63
  no shutdown
  no ip redirects
  ip address 10.10.63.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 63
    preempt
    ip 10.10.63.1

interface Vlan64
  description iSCSI Fabric A path vlan 64
  no shutdown
  no ip redirects
  ip address 10.10.64.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 64
    preempt
    priority 110
    ip 10.10.64.1

interface Vlan65
  description iSCSI Fabric B path vlan 65
  no shutdown
  no ip redirects
  ip address 10.10.65.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 65
    preempt
    ip 10.10.65.1

interface Vlan66
  description vMotion network vlan 66
  no shutdown
  no ip redirects
  ip address 10.10.66.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1

interface Vlan67
  description Nexus 1000v vlan 67
  no shutdown
  no ip redirects
  ip address 10.10.67.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 67
    preempt
    ip 10.10.67.1

interface Vlan68
  description LoginVSI Launchers vlan 68
  no shutdown
  no ip redirects
  ip address 10.10.68.2/23
  no ipv6 redirects
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
  ip dhcp relay address 10.10.61.30

interface Vlan69
  description LoginVSI Launchers 10.10.81-network vlan 69
  no shutdown
  no ip redirects
  ip address 10.10.81.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 69
    preempt
    ip 10.10.81.1

interface Vlan102
  description VDI vlan 102
  no shutdown
  no ip redirects
  ip address 10.2.0.2/19
  no ipv6 redirects
  hsrp version 2
  hsrp 102
    preempt delay minimum 240
    priority 110
    timers  1  3
    ip 10.2.0.1
  ip dhcp relay address 10.10.61.30

interface port-channel10
  description VPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FI-A_9k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
  description FI-B_9k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel15
  description FI-A_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 15

interface port-channel16
  description FI-B_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 16

interface port-channel51
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 51

interface port-channel52
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  speed 40000
  no negotiate auto
  vpc 52

interface port-channel53
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 53

interface port-channel54
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 54

interface Ethernet1/1
  description NetApp_AFF8080_Node-02_port_e0e_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/2
  description NetApp_AFF8080_Node-02_port_e1a_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/3
  description NetApp_AFF8080_Node-01_port_e0e_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216
```

```
interface Ethernet1/4
  description NetApp_AFF8080_Node-01_port_e4a_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/5
  description NetApp_AFF8080_Node-02_port_e0f_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/6
  description NetApp_AFF8080_Node-02_port_e4a_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/7
  description NetApp_AFF8080_Node-01_port_e0f_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/8
  description NetApp_AFF8080_Node-01_port_e1a_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/9
```

```
  description Jumphost ToR
  switchport access vlan 60
  spanning-tree port type edge
  speed 1000

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23
interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35
  description Uplink_from_FI-A_to-N9KA
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/36
  description Uplink_from_FI-B_to-N9KB
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 15 mode active

interface Ethernet1/46
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 16 mode active

interface Ethernet1/47

interface Ethernet1/48
  description TOR
  switchport access vlan 264

interface Ethernet1/49
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  channel-group 10 mode active

interface Ethernet1/50
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  channel-group 10 mode active

interface Ethernet1/51
  description FI-A-N9K-UPlink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 51 mode active

interface Ethernet1/52
  description FI-B-N9K-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  speed 40000
  no negotiate auto
  channel-group 52 mode active

interface Ethernet1/53
  switchport mode trunk
  mtu 9216
  channel-group 53 mode active

interface Ethernet1/54
  switchport mode trunk
  mtu 9216
  channel-group 54 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.65/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
no system default switchport shutdown
```

N93180YC-FX -B Configuration

```
!Command: show running-config
!Time: Sun Jun 24 13:51:59 2018

version 7.0(3)I7(2)
switchname DV-Pod-2-N9K-B
class-map type network-qos class-platinum
match qos-group 2
class-map type network-qos class-all-flood
```

```
match qos-group 2
class-map type network-qos system_nq_policy
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
policy-map type network-qos jumbo
  class type network-qos class-platinum
    mtu 9216
  class type network-qos class-default
    mtu 9216
vdc DV-Pod-2-N9K-B id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
feature nxapi
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock timezone PST -7 0
```

```
no password strength-check
username admin password 5 $1$fp3LrGLC$PF8eML85qkPBgdH/bZAKK/  role network-admin
ip domain-lookup
system default switchport
ip access-list NFS_VLAN63
  10 permit ip 10.10.63.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-A_64
  10 permit ip 10.10.64.0 255.255.255.0 any
  20 deny ip any any
ip access-list iSCSI-B_65
  10 permit ip 10.10.65.0 255.255.255.0 any
  20 deny ip any any
class-map type qos match-any class-platinum
  match cos 5
policy-map type qos jumbo
  class class-platinum
    set qos-group 2
  class class-default
    set qos-group 0
system qos
  service-policy type network-qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0x142e177306873a75257c9a8388b47fb7
 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp peer 10.10.160.2 use-vrf default
ntp server 10.10.160.3 use-vrf default
ntp server 72.163.32.44 use-vrf management
ntp logging
ntp master 8

vlan 1-2,60-70,102,164,264
vlan 60
  name In-Band-Mgmt
vlan 61
  name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name NFS
vlan 64
  name iSCSI-A
vlan 65
  name iSCSI-B
vlan 66
  name vMotion
vlan 67
  name N1KV
vlan 68
  name LauncherPXE
vlan 69
  name Launcher81
vlan 70
  name other-3
vlan 102
  name VDI
vlan 164
  name Out-Of-Band-Mgmt
vlan 264
```

```
  name OOB-Mgmt

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.65 source 10.29.164.66
  delay restore 150
  peer-gateway
  auto-recovery


interface Vlan1
  no shutdown
  no ip redirects
  ip address 10.29.164.3/24
  no ipv6 redirects

interface Vlan2
  description Default native vlan 2
  no ip redirects
  no ipv6 redirects

interface Vlan60
```

```
  description Out of Band Management vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 60
    preempt
    priority 110
    ip 10.10.60.1

interface Vlan61
  description Infrastructure vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
    preempt
    ip 10.10.61.1

interface Vlan62
  description CIFS vlan 62
  no shutdown
  no ip redirects
  ip address 10.10.62.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 62
    preempt
    priority 110
    ip 10.10.62.1

interface Vlan63
  description NFS vlan 63
  no shutdown
  mtu 9216
  no ip redirects
  ip address 10.10.63.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 63
    preempt
    ip 10.10.63.1

interface Vlan64
  description iSCSI Fabric A path vlan 64
  no shutdown
  no ip redirects
  ip address 10.10.64.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 64
    preempt
    priority 110
    ip 10.10.64.1

interface Vlan65
  description iSCSI Fabric B path vlan 65
  no shutdown
  no ip redirects
  ip address 10.10.65.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 65
    preempt
    ip 10.10.65.1

interface Vlan66
  description vMotion network vlan 66
  no shutdown
  no ip redirects
  ip address 10.10.66.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1

interface Vlan67
  description Nexus 1000v vlan 67
  no shutdown
  no ip redirects
  ip address 10.10.67.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 67
    preempt
    ip 10.10.67.1

interface Vlan68
  description LoginVSI Launchers vlan 68
  no shutdown
  no ip redirects
  ip address 10.10.68.3/23
  no ipv6 redirects
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
  ip dhcp relay address 10.10.61.30

interface Vlan69
  description LoginVSI Launchers 10.10.81-network vlan 69
  no shutdown
  no ip redirects
  ip address 10.10.81.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 69
    preempt
    ip 10.10.81.1

interface Vlan102
  description VDI vlan 102
  no shutdown
  no ip redirects
  ip address 10.2.0.3/19
  no ipv6 redirects
  hsrp version 2
  hsrp 102
    preempt delay minimum 240
    priority 110
    timers  1  3
    ip 10.2.0.1
  ip dhcp relay address 10.10.61.30
```

```
interface port-channel10
  description VPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FI-A_9k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
  description FI-B_9k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel15
  description FI-A_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 15

interface port-channel16
```

```
  description FI-B_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 16

interface port-channel51
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  speed 40000
  no negotiate auto
  vpc 51

interface port-channel52
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  speed 40000
  no negotiate auto
  vpc 52

interface port-channel53
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 53

interface port-channel54
```

```
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 54

interface Ethernet1/1
  description NetApp_AFF8080_Node-02_port_e0g_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/2
  description NetApp_AFF8080_Node-02_port_e1b_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/3
  description NetApp_AFF8080_Node-01_port_e0g_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/4
  description NetApp_AFF8080_Node-01_port_e4b_NFS
  switchport mode trunk
  switchport trunk allowed vlan 63
  mtu 9216

interface Ethernet1/5
  description NetApp_AFF8080_Node-02_port_e0h_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/6
  description NetApp_AFF8080_Node-02_port_e4b_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/7
  description NetApp_AFF8080_Node-01_port_e0h_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/8
  description NetApp_AFF8080_Node-01_port_e1b_CIFS
  switchport mode trunk
  switchport trunk allowed vlan 62,64-65
  mtu 9216

interface Ethernet1/9
  description Jumphost ToR
  switchport access vlan 60
  spanning-tree port type edge
  speed 1000

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12
interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28
interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35
  description Uplink_from_FI-A_to_N9KA
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/36
  description Uplink_from_FI-B_to-N9KB
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39
interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45
  description Uplink_from_LoginVSI_Launchers_FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 15 mode active

interface Ethernet1/46
  description Uplink_from_LoginVSI_Launchers_FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 16 mode active

interface Ethernet1/47

interface Ethernet1/48
  description TOR
  switchport access vlan 264

interface Ethernet1/49
```

```
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  channel-group 10 mode active

interface Ethernet1/50
  description VPC Peer Link between 9ks
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164,264
  channel-group 10 mode active

interface Ethernet1/51
  description FI-B-N9KUPlink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  speed 40000
  no negotiate auto
  channel-group 51 mode active

interface Ethernet1/52
  description FI-B-N9KUPlink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  speed 40000
  no negotiate auto
  channel-group 52 mode active

interface Ethernet1/53
  switchport mode trunk
  mtu 9216
  channel-group 53 mode active

interface Ethernet1/54
  switchport mode trunk
  mtu 9216
  channel-group 54 mode active

interface mgmt0
  vrf member management
  ip address 10.29.164.66/24
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
no system default switchport shutdown Fibre Channel Configuration
```

Fibre Channel Configuration

Cisco MDS 9148S - A Configuration

```
!Command: show running-config
!Time: Sun Jun 24 21:46:22 2018

version 8.1(1)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
no password strength-check
```

```
username admin password 5 $1$DDq8vF1x$EwCSM0O3dlXZ4jlPy9ZoC.  role network-admin
ip domain-lookup
ip host MDS-A  10.29.164.238
aaa group server radius radius
snmp-server contact jnichols
snmp-server user admin network-admin auth md5 0x2efbf582e573df2038164f1422c231fe
 priv 0x2efbf582e573df2038164f1422c231fe localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1163
snmp-server host 10.29.132.18 traps version 2c public udp-port 1163
snmp-server host 10.29.164.130 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
snmp-server community public group network-operator
vsan database
  vsan 1 name "Synergy"
  vsan 3 name "SP-Launcher-A"
  vsan 100 name "FlashStack-VCC-CVD-Fabric-A"
  vsan 400 name "FlexPod-A"
device-alias database
  device-alias name C480M5-P0 pwwn 21:00:00:0e:1e:10:a2:c0
  device-alias name VDI-1-hba1 pwwn 20:00:00:25:b5:3a:00:3f
  device-alias name VDI-2-hba1 pwwn 20:00:00:25:b5:3a:00:0f
  device-alias name VDI-3-hba1 pwwn 20:00:00:25:b5:3a:00:1f
  device-alias name VDI-4-hba1 pwwn 20:00:00:25:b5:3a:00:4e
  device-alias name VDI-5-hba1 pwwn 20:00:00:25:b5:3a:00:2e
  device-alias name VDI-6-hba1 pwwn 20:00:00:25:b5:3a:00:3e
  device-alias name VDI-7-hba1 pwwn 20:00:00:25:b5:3a:00:0e
  device-alias name VDI-9-hba1 pwwn 20:00:00:25:b5:3a:00:4d
  device-alias name a300-01-0g pwwn 20:01:00:a0:98:af:bd:e8
```

```
device-alias name a300-02-0g pwwn 20:03:00:a0:98:af:bd:e8
device-alias name CS700-FC1-1 pwwn 56:c9:ce:90:0d:e8:24:02
device-alias name CS700-FC2-1 pwwn 56:c9:ce:90:0d:e8:24:06
device-alias name VDI-10-hba1 pwwn 20:00:00:25:b5:3a:00:2d
device-alias name VDI-11-hba1 pwwn 20:00:00:25:b5:3a:00:3d
device-alias name VDI-12-hba1 pwwn 20:00:00:25:b5:3a:00:0d
device-alias name VDI-13-hba1 pwwn 20:00:00:25:b5:3a:00:1d
device-alias name VDI-14-hba1 pwwn 20:00:00:25:b5:3a:00:4c
device-alias name VDI-15-hba1 pwwn 20:00:00:25:b5:3a:00:2c
device-alias name VDI-17-hba1 pwwn 20:00:00:25:b5:3a:00:0c
device-alias name VDI-18-hba1 pwwn 20:00:00:25:b5:3a:00:1c
device-alias name VDI-19-hba1 pwwn 20:00:00:25:b5:3a:00:4b
device-alias name VDI-20-hba1 pwwn 20:00:00:25:b5:3a:00:2b
device-alias name VDI-21-hba1 pwwn 20:00:00:25:b5:3a:00:3b
device-alias name VDI-22-hba1 pwwn 20:00:00:25:b5:3a:00:0b
device-alias name VDI-23-hba1 pwwn 20:00:00:25:b5:3a:00:1b
device-alias name VDI-24-hba1 pwwn 20:00:00:25:b5:3a:00:4a
device-alias name VDI-25-hba1 pwwn 20:00:00:25:b5:3a:00:2a
device-alias name VDI-26-hba1 pwwn 20:00:00:25:b5:3a:00:3a
device-alias name VDI-27-hba1 pwwn 20:00:00:25:b5:3a:00:0a
device-alias name VDI-28-hba1 pwwn 20:00:00:25:b5:3a:00:1a
device-alias name VDI-29-hba1 pwwn 20:00:00:25:b5:3a:00:49
device-alias name VDI-30-hba1 pwwn 20:00:00:25:b5:3a:00:39
device-alias name VDI-31-hba1 pwwn 20:00:00:25:b5:3a:00:1e
device-alias name VDI-32-hba1 pwwn 20:00:00:25:b5:3a:00:3c
device-alias name X70-CT0-FC0 pwwn 52:4a:93:75:dd:91:0a:00
device-alias name X70-CT0-FC2 pwwn 52:4a:93:75:dd:91:0a:02
device-alias name X70-CT1-FC1 pwwn 52:4a:93:75:dd:91:0a:11
device-alias name X70-CT1-FC3 pwwn 52:4a:93:75:dd:91:0a:13
device-alias name Infra01-8-hba1 pwwn 20:00:00:25:b5:3a:00:4f
device-alias name Infra02-16-hba1 pwwn 20:00:00:25:b5:3a:00:2f
device-alias 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 name AAD-16-CH1-BL1-FC0 pwwn 20:00:00:25:b5:9a:a0:00
```

```
device-alias name AAD-16-CH1-BL2-FC0 pwwn 20:00:00:25:b5:9a:a0:02
device-alias name AAD-16-CH1-BL3-FC0 pwwn 20:00:00:25:b5:9a:a0:04
device-alias name AAD-16-CH1-BL4-FC0 pwwn 20:00:00:25:b5:9a:a0:06
device-alias name AAD-16-CH1-BL5-FC0 pwwn 20:00:00:25:b5:9a:a0:08
device-alias name AAD-16-CH1-BL6-FC0 pwwn 20:00:00:25:b5:9a:a0:0a
device-alias name AAD-16-CH1-BL7-FC0 pwwn 20:00:00:25:b5:9a:a0:0c
device-alias name AAD-16-CH1-BL8-FC0 pwwn 20:00:00:25:b5:9a:a0:0e
device-alias name AAD-16-CH2-BL1-FC0 pwwn 20:00:00:25:b5:9a:a0:10
device-alias name AAD-16-CH2-BL2-FC0 pwwn 20:00:00:25:b5:9a:a0:12
device-alias name AAD-16-CH2-BL3-FC0 pwwn 20:00:00:25:b5:9a:a0:14
device-alias name AAD-16-CH2-BL4-FC0 pwwn 20:00:00:25:b5:9a:a0:16
device-alias name AAD-16-CH2-BL5-FC0 pwwn 20:00:00:25:b5:9a:a0:18
device-alias name AAD-16-CH2-BL6-FC0 pwwn 20:00:00:25:b5:9a:a0:1a
device-alias name AAD-16-CH2-BL7-FC0 pwwn 20:00:00:25:b5:9a:a0:1c
device-alias name AAD-16-CH2-BL8-FC0 pwwn 20:00:00:25:b5:9a:a0:1e
device-alias name AAD-16-CH3-BL1-FC0 pwwn 20:00:00:25:b5:17:aa:00
device-alias name AAD-16-CH3-BL2-FC0 pwwn 20:00:00:25:b5:17:aa:02
device-alias name AAD-16-CH3-BL3-FC0 pwwn 20:00:00:25:b5:17:aa:04
device-alias name AAD-16-CH3-BL4-FC0 pwwn 20:00:00:25:b5:17:aa:06
device-alias name AAD-16-CH3-BL5-FC0 pwwn 20:00:00:25:b5:17:aa:08
device-alias name AAD-16-CH3-BL6-FC0 pwwn 20:00:00:25:b5:17:aa:0a
device-alias name AAD-16-CH3-BL7-FC0 pwwn 20:00:00:25:b5:17:aa:0c
device-alias name AAD-16-CH3-BL8-FC0 pwwn 20:00:00:25:b5:17:aa:0e
device-alias name AAD-16-CH4-BL1-FC0 pwwn 20:00:00:25:b5:17:aa:10
device-alias name AAD-16-CH4-BL2-FC0 pwwn 20:00:00:25:b5:17:aa:12
device-alias name AAD-16-CH4-BL3-FC0 pwwn 20:00:00:25:b5:17:aa:14
device-alias name AAD-16-CH4-BL4-FC0 pwwn 20:00:00:25:b5:17:aa:16
device-alias name AAD-16-CH4-BL5-FC0 pwwn 20:00:00:25:b5:17:aa:18
device-alias name AAD-16-CH4-BL6-FC0 pwwn 20:00:00:25:b5:17:aa:1a
device-alias name AAD-16-CH4-BL7-FC0 pwwn 20:00:00:25:b5:17:aa:1c
device-alias name AAD-16-CH4-BL8-FC0 pwwn 20:00:00:25:b5:17:aa:1e
```

```
device-alias commit


fcdomain fcid database
  vsan 3 wwn 20:00:00:25:b5:17:aa:0e fcid 0x301b03 dynamic
!           [AAD-16-CH3-BL8-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:00 fcid 0x301c02 dynamic
!           [AAD-16-CH3-BL1-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:0c fcid 0x301c05 dynamic
!           [AAD-16-CH3-BL7-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:14 fcid 0x301b04 dynamic
!           [AAD-16-CH4-BL3-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:06 fcid 0x301d05 dynamic
!           [AAD-16-CH3-BL4-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:16 fcid 0x301c09 dynamic
!           [AAD-16-CH4-BL4-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:08 fcid 0x301d09 dynamic
!           [AAD-16-CH3-BL5-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:0a fcid 0x301e09 dynamic
!           [AAD-16-CH3-BL6-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:12 fcid 0x301e02 dynamic
!           [AAD-16-CH4-BL2-FC0]
  vsan 3 wwn 56:c9:ce:90:0d:e8:24:06 fcid 0x301700 dynamic
!           [CS700-FC2-1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:1d fcid 0x680509 dynamic
!             [VDI-13-hba1]
  vsan 3 wwn 20:00:00:25:b5:17:aa:02 fcid 0x301b06 dynamic
!           [AAD-16-CH3-BL2-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:1a fcid 0x301b01 dynamic
!           [AAD-16-CH4-BL6-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:04 fcid 0x301e01 dynamic
!           [AAD-16-CH3-BL3-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:1c fcid 0x301c08 dynamic
```

```
!           [AAD-16-CH4-BL7-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:18 fcid 0x301d01 dynamic
!           [AAD-16-CH4-BL5-FC0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:1e fcid 0x301d04 dynamic
!           [AAD-16-CH4-BL8-FC0]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:06 fcid 0x810000 dynamic
  vsan 400 wwn 20:00:00:25:b5:3a:00:4c fcid 0x68040a dynamic
!             [VDI-14-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:2c fcid 0x680209 dynamic
!             [VDI-15-hba1]
  vsan 1 wwn 52:4a:93:75:dd:91:0a:06 fcid 0x291000 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:07 fcid 0x810100 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:16 fcid 0x810200 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:17 fcid 0x810300 dynamic
  vsan 3 wwn 20:4d:54:7f:ee:83:42:00 fcid 0x301b00 dynamic
  vsan 3 wwn 20:4e:54:7f:ee:83:42:00 fcid 0x301c00 dynamic
  vsan 3 wwn 20:4f:54:7f:ee:83:42:00 fcid 0x301d00 dynamic
  vsan 3 wwn 20:50:54:7f:ee:83:42:00 fcid 0x301e00 dynamic
  vsan 1 wwn 20:04:00:de:fb:92:8d:00 fcid 0x291100 dynamic
  vsan 1 wwn 20:02:00:de:fb:92:8d:00 fcid 0x291200 dynamic
  vsan 1 wwn 20:03:00:de:fb:92:8d:00 fcid 0x291300 dynamic
  vsan 1 wwn 20:01:00:de:fb:92:8d:00 fcid 0x291400 dynamic
  vsan 400 wwn 20:00:00:25:b5:3a:00:49 fcid 0x680205 dynamic
!             [VDI-29-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:4b fcid 0x680201 dynamic
!             [VDI-19-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:0c fcid 0x680309 dynamic
!             [VDI-17-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:39 fcid 0x680307 dynamic
!             [VDI-30-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:3f fcid 0x680409 dynamic
!             [VDI-1-hba1]
```

```
  vsan 400 wwn 20:00:00:25:b5:3a:00:2e fcid 0x680508 dynamic
!             [VDI-5-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:0a fcid 0x680206 dynamic
!             [VDI-27-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:1e fcid 0x680207 dynamic
!             [VDI-31-hba1]
  vsan 100 wwn 20:02:00:de:fb:92:8d:00 fcid 0x810400 dynamic
  vsan 100 wwn 20:03:00:de:fb:92:8d:00 fcid 0x810500 dynamic
  vsan 400 wwn 20:00:00:25:b5:3a:00:3a fcid 0x680305 dynamic
!             [VDI-26-hba1]
  vsan 100 wwn 20:04:00:de:fb:92:8d:00 fcid 0x810600 dynamic
  vsan 400 wwn 20:00:00:25:b5:3a:00:0f fcid 0x680202 dynamic
!             [VDI-2-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:1c fcid 0x680301 dynamic
!             [VDI-18-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:0e fcid 0x680308 dynamic
!             [VDI-7-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:2a fcid 0x680402 dynamic
!             [VDI-25-hba1]
  vsan 100 wwn 20:01:00:de:fb:92:8d:00 fcid 0x810700 dynamic
  vsan 400 wwn 20:00:00:25:b5:3a:00:1f fcid 0x680506 dynamic
!             [VDI-3-hba1]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1e fcid 0x810601 dynamic
!             [VCC-Infra01-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:02 fcid 0x810706 dynamic
!             [VCC-WLHost02-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0a fcid 0x810711 dynamic
!             [VCC-WLHost06-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0e fcid 0x810604 dynamic
!             [VCC-WLHost08-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:26 fcid 0x810509 dynamic
!             [VCC-WLHost19-HBA0]
```

```
  vsan 100 wwn 20:00:00:25:b5:aa:17:2e fcid 0x81070d dynamic
!          [VCC-WLHost23-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:22 fcid 0x81060b dynamic
!          [VCC-WLHost17-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:28 fcid 0x810703 dynamic
!          [VCC-WLHost20-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:24 fcid 0x810512 dynamic
!          [VCC-WLHost18-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:06 fcid 0x810606 dynamic
!          [VCC-WLHost04-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0c fcid 0x81050e dynamic
!          [VCC-WLHost07-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:08 fcid 0x810402 dynamic
!          [VCC-WLHost05-HBA0]
  vsan 400 wwn 20:00:00:25:b5:3a:00:4e fcid 0x680507 dynamic
!          [VDI-4-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:1a fcid 0x680406 dynamic
!          [VDI-28-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:3e fcid 0x680504 dynamic
!          [VDI-6-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:0d fcid 0x680302 dynamic
!          [VDI-12-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:3d fcid 0x680401 dynamic
!          [VDI-11-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:1b fcid 0x680407 dynamic
!          [VDI-23-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:2f fcid 0x680304 dynamic
!          [Infra02-16-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:0b fcid 0x680403 dynamic
!          [VDI-22-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:4f fcid 0x680404 dynamic
!          [Infra01-8-hba1]
```

```
  vsan 1 wwn 52:4a:93:75:dd:91:0a:07 fcid 0x291500 dynamic
  vsan 1 wwn 52:4a:93:75:dd:91:0a:16 fcid 0x291600 dynamic
  vsan 1 wwn 52:4a:93:75:dd:91:0a:17 fcid 0x291700 dynamic
  vsan 100 wwn 20:00:00:25:b5:aa:17:00 fcid 0x81070b dynamic
!              [VCC-WLHost01-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:04 fcid 0x810701 dynamic
!              [VCC-WLHost03-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:16 fcid 0x81050d dynamic
!              [VCC-WLHost12-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:10 fcid 0x810607 dynamic
!              [VCC-WLHost09-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:12 fcid 0x810705 dynamic
!              [VCC-WLHost10-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:18 fcid 0x810702 dynamic
!              [VCC-WLHost13-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:20 fcid 0x810708 dynamic
!              [VCC-WLHost16-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:2c fcid 0x810508 dynamic
!              [VCC-WLHost22-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:2a fcid 0x810413 dynamic
!              [VCC-WLHost21-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1a fcid 0x81070c dynamic
!              [VCC-WLHost14-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1c fcid 0x810414 dynamic
!              [VCC-WLHost15-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:14 fcid 0x810411 dynamic
!              [VCC-WLHost11-HBA0]
  vsan 3 wwn 20:00:00:25:b5:17:aa:10 fcid 0x301e03 dynamic
!            [AAD-16-CH4-BL1-FC0]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:02 fcid 0x810800 dynamic
!              [X70-CT0-FC2]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:03 fcid 0x810900 dynamic
  vsan 100 wwn 52:4a:93:75:dd:91:0a:13 fcid 0x810a00 dynamic
!           [X70-CT1-FC3]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:12 fcid 0x810b00 dynamic
  vsan 100 wwn 20:00:00:25:b5:aa:17:3e fcid 0x810406 dynamic
!           [VCC-Infra02-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:3f fcid 0x810713 dynamic
!           [VCC-Infra02-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1f fcid 0x810511 dynamic
!           [VCC-Infra01-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:03 fcid 0x810504 dynamic
!           [VCC-WLHost02-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:01 fcid 0x810608 dynamic
!           [VCC-WLHost01-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:05 fcid 0x810704 dynamic
!           [VCC-WLHost03-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:07 fcid 0x810610 dynamic
!           [VCC-WLHost04-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0b fcid 0x810611 dynamic
!           [VCC-WLHost06-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:09 fcid 0x810709 dynamic
!           [VCC-WLHost05-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0d fcid 0x810605 dynamic
!           [VCC-WLHost07-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:0f fcid 0x810407 dynamic
!           [VCC-WLHost08-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:21 fcid 0x810401 dynamic
!           [VCC-WLHost16-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:23 fcid 0x81050a dynamic
!           [VCC-WLHost17-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:25 fcid 0x810403 dynamic
!           [VCC-WLHost18-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:27 fcid 0x81060a dynamic
```

```
!          [VCC-WLHost19-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:29 fcid 0x810603 dynamic
!          [VCC-WLHost20-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:2b fcid 0x81060d dynamic
!          [VCC-WLHost21-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:2d fcid 0x81040d dynamic
!          [VCC-WLHost22-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:2f fcid 0x810505 dynamic
!          [VCC-WLHost23-HBA2]
  vsan 400 wwn 20:00:00:25:b5:3a:00:4d fcid 0x680203 dynamic
!          [VDI-9-hba1]
  vsan 400 wwn 20:00:00:25:b5:3a:00:3c fcid 0x680204 dynamic
!          [VDI-32-hba1]
  vsan 100 wwn 20:00:00:25:b5:aa:17:11 fcid 0x810506 dynamic
!          [VCC-WLHost09-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:13 fcid 0x81050b dynamic
!          [VCC-WLHost10-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:15 fcid 0x810409 dynamic
!          [VCC-WLHost11-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:17 fcid 0x81050f dynamic
!          [VCC-WLHost12-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:19 fcid 0x810710 dynamic
!          [VCC-WLHost13-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1b fcid 0x81040b dynamic
!          [VCC-WLHost14-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:1d fcid 0x810404 dynamic
!          [VCC-WLHost15-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:34 fcid 0x810408 dynamic
!          [VCC-WLHost26-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:32 fcid 0x81040a dynamic
!          [VCC-WLHost25-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:33 fcid 0x810405 dynamic
```

```
!              [VCC-WLHost25-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:35 fcid 0x81040c dynamic
!              [VCC-WLHost26-HBA2]
  vsan 400 wwn 20:00:00:25:b5:3a:00:2d fcid 0x680501 dynamic
!              [VDI-10-hba1]
  vsan 100 wwn 20:00:00:25:b5:aa:17:38 fcid 0x810507 dynamic
!              [VCC-WLHost28-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:39 fcid 0x81070f dynamic
!              [VCC-WLHost28-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:30 fcid 0x810502 dynamic
!              [VCC-WLHost24-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:3a fcid 0x81060e dynamic
!              [VCC-WLHost29-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:36 fcid 0x810503 dynamic
!              [VCC-WLHost27-HBA0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:3c fcid 0x810609 dynamic
!              [VCC-WLHost30-HBA0]
  vsan 400 wwn 20:00:00:25:b5:3a:00:3b fcid 0x680502 dynamic
!              [VDI-21-hba1]
  vsan 100 wwn 20:00:00:25:b5:aa:17:3d fcid 0x81070a dynamic
!              [VCC-WLHost30-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:3b fcid 0x810501 dynamic
!              [VCC-WLHost29-HBA2]
  vsan 100 wwn 20:00:00:25:b5:aa:17:37 fcid 0x810602 dynamic
!              [VCC-WLHost27-HBA2]
  vsan 400 wwn 20:00:00:25:b5:3a:00:4a fcid 0x680306 dynamic
!              [VDI-24-hba1]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:00 fcid 0x301b02 dynamic
!            [AAD-16-CH1-BL1-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:02 fcid 0x301c03 dynamic
!            [AAD-16-CH1-BL2-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:06 fcid 0x301e05 dynamic
```

```
!          [AAD-16-CH1-BL4-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:08 fcid 0x301b09 dynamic
!          [AAD-16-CH1-BL5-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:0a fcid 0x301c06 dynamic
!          [AAD-16-CH1-BL6-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:0e fcid 0x301e06 dynamic
!          [AAD-16-CH1-BL8-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:12 fcid 0x301c04 dynamic
!          [AAD-16-CH2-BL2-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:0c fcid 0x301d02 dynamic
!          [AAD-16-CH1-BL7-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:10 fcid 0x301b07 dynamic
!          [AAD-16-CH2-BL1-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:04 fcid 0x301c07 dynamic
!          [AAD-16-CH1-BL3-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:1c fcid 0x301d07 dynamic
!          [AAD-16-CH2-BL7-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:14 fcid 0x301e07 dynamic
!          [AAD-16-CH2-BL3-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:16 fcid 0x301b08 dynamic
!          [AAD-16-CH2-BL4-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:18 fcid 0x301d03 dynamic
!          [AAD-16-CH2-BL5-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:1a fcid 0x301d08 dynamic
!          [AAD-16-CH2-BL6-FC0]
  vsan 3 wwn 20:00:00:25:b5:9a:a0:1e fcid 0x301e08 dynamic
!          [AAD-16-CH2-BL8-FC0]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:00 fcid 0x810c00 dynamic
!            [X70-CT0-FC0]
  vsan 100 wwn 52:4a:93:75:dd:91:0a:11 fcid 0x810d00 dynamic
!            [X70-CT1-FC1]
  vsan 1 wwn 21:00:00:0e:1e:10:a2:c0 fcid 0x290000 dynamic
```

```
!           [C480M5-P0]
  vsan 3 wwn 21:00:00:0e:1e:10:a2:c0 fcid 0x300000 dynamic
!           [C480M5-P0]
  vsan 100 wwn 20:00:00:25:b5:aa:17:31 fcid 0x81060c dynamic
!             [VCC-WLHost24-HBA2]
  vsan 400 wwn 20:00:00:25:b5:3a:00:2b fcid 0x680503 dynamic
!             [VDI-20-hba1]
  vsan 1 wwn 20:01:00:de:fb:8f:60:00 fcid 0x290100 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:05 fcid 0x290101 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:01 fcid 0x290102 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:03 fcid 0x290103 dynamic
  vsan 3 wwn 20:01:00:de:fb:8f:60:00 fcid 0x300100 dynamic
  vsan 3 wwn 20:00:00:25:b5:0a:00:05 fcid 0x300101 dynamic
  vsan 3 wwn 20:00:00:25:b5:0a:00:03 fcid 0x300102 dynamic
  vsan 3 wwn 20:00:00:25:b5:0a:00:01 fcid 0x300103 dynamic
  vsan 400 wwn 50:0a:09:84:80:d3:67:d3 fcid 0x680000 dynamic
  vsan 400 wwn 20:03:00:a0:98:af:bd:e8 fcid 0x680001 dynamic
!             [a300-02-0g]
  vsan 400 wwn 50:0a:09:84:80:13:41:27 fcid 0x680100 dynamic
  vsan 400 wwn 20:01:00:a0:98:af:bd:e8 fcid 0x680101 dynamic
!             [a300-01-0g]
  vsan 400 wwn 20:02:00:de:fb:90:a0:80 fcid 0x680200 dynamic
  vsan 400 wwn 20:03:00:de:fb:90:a0:80 fcid 0x680300 dynamic
  vsan 400 wwn 20:04:00:de:fb:90:a0:80 fcid 0x680400 dynamic
  vsan 400 wwn 20:01:00:de:fb:90:a0:80 fcid 0x680500 dynamic
  vsan 3 wwn 56:c9:ce:90:0d:e8:24:02 fcid 0x301600 dynamic
!           [CS700-FC1-1]
!Active Zone Database Section for vsan 1
zone name xs1-fc0 vsan 1
    member pwwn 20:00:00:25:b5:0a:00:03
    member pwwn 56:c9:ce:90:0d:e8:24:02
!              [CS700-FC1-1]
```

```
  member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]

zone name xs2-fc0 vsan 1
  member pwwn 20:00:00:25:b5:0a:00:05
  member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]
  member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]

zoneset name Synergy-A vsan 1
  member xs1-fc0
  member xs2-fc0

zoneset activate name Synergy-A vsan 1
do clear zone database vsan 1
!Full Zone Database Section for vsan 1
zone name xs1-fc0 vsan 1
  member pwwn 20:00:00:25:b5:0a:00:03
  member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]
  member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]

zone name xs2-fc0 vsan 1
  member pwwn 20:00:00:25:b5:0a:00:05
  member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]
  member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]

zoneset name Synergy-A vsan 1
```

```
    member xs1-fc0
    member xs2-fc0

!Active Zone Database Section for vsan 3
zone name SP-Launcher-01-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:00
!               [AAD-16-CH3-BL1-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-02-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:02
!               [AAD-16-CH3-BL2-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-03-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:04
!               [AAD-16-CH3-BL3-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-04-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:06
!               [AAD-16-CH3-BL4-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!            [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!            [CS700-FC2-1]

zone name SP-Launcher-05-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:08
!            [AAD-16-CH3-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!            [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!            [CS700-FC2-1]

zone name SP-Launcher-06-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0a
!            [AAD-16-CH3-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!            [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!            [CS700-FC2-1]

zone name SP-Launcher-07-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0c
!            [AAD-16-CH3-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!            [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!            [CS700-FC2-1]

zone name SP-Launcher-08-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0e
!            [AAD-16-CH3-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-09-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:10
!               [AAD-16-CH4-BL1-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-10-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:12
!               [AAD-16-CH4-BL2-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-11-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:14
!               [AAD-16-CH4-BL3-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-12-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:16
!               [AAD-16-CH4-BL4-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-13-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:18
!               [AAD-16-CH4-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-14-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1a
!               [AAD-16-CH4-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-15-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1c
!               [AAD-16-CH4-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-16-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1e
!               [AAD-16-CH4-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-01-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:00
!               [AAD-16-CH1-BL1-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-02-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:02
!               [AAD-16-CH1-BL2-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-03-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:04
!               [AAD-16-CH1-BL3-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-04-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:06
!               [AAD-16-CH1-BL4-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!          [CS700-FC1-1]
  member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name FP-Launcher-05-FC0 vsan 3
  member pwwn 20:00:00:25:b5:9a:a0:08
!          [AAD-16-CH1-BL5-FC0]
  member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
  member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name FP-Launcher-06-FC0 vsan 3
  member pwwn 20:00:00:25:b5:9a:a0:0a
!          [AAD-16-CH1-BL6-FC0]
  member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
  member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name FP-Launcher-07-FC0 vsan 3
  member pwwn 20:00:00:25:b5:9a:a0:0c
!          [AAD-16-CH1-BL7-FC0]
  member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
  member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name FP-Launcher-08-FC0 vsan 3
  member pwwn 20:00:00:25:b5:9a:a0:0e
!          [AAD-16-CH1-BL8-FC0]
  member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-09-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:10
!               [AAD-16-CH2-BL1-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-10-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:12
!               [AAD-16-CH2-BL2-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-11-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:14
!               [AAD-16-CH2-BL3-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-12-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:16
!               [AAD-16-CH2-BL4-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-13-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:18
!               [AAD-16-CH2-BL5-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-14-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:1a
!               [AAD-16-CH2-BL6-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-15-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:1c
!               [AAD-16-CH2-BL7-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-16-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:1e
!               [AAD-16-CH2-BL8-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name C480M5-P0 vsan 3
    member pwwn 21:00:00:0e:1e:10:a2:c0
!               [C480M5-P0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name xs1-fc0 vsan 3
    member pwwn 20:00:00:25:b5:0a:00:03
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name xs2-fc0 vsan 3
    member pwwn 20:00:00:25:b5:0a:00:05
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]

zoneset name Launcher_FabricA vsan 3
    member SP-Launcher-01-FC0
    member SP-Launcher-02-FC0
    member SP-Launcher-03-FC0
    member SP-Launcher-04-FC0
    member SP-Launcher-05-FC0
    member SP-Launcher-06-FC0
    member SP-Launcher-07-FC0
    member SP-Launcher-08-FC0
    member SP-Launcher-09-FC0
    member SP-Launcher-10-FC0
    member SP-Launcher-11-FC0
    member SP-Launcher-12-FC0
    member SP-Launcher-13-FC0
    member SP-Launcher-14-FC0
    member SP-Launcher-15-FC0
    member SP-Launcher-16-FC0
    member FP-Launcher-01-FC0
    member FP-Launcher-02-FC0
    member FP-Launcher-03-FC0
    member FP-Launcher-04-FC0
    member FP-Launcher-05-FC0
    member FP-Launcher-06-FC0
    member FP-Launcher-07-FC0
    member FP-Launcher-08-FC0
    member FP-Launcher-09-FC0
    member FP-Launcher-10-FC0
    member FP-Launcher-11-FC0
    member FP-Launcher-12-FC0
    member FP-Launcher-13-FC0
    member FP-Launcher-14-FC0
    member FP-Launcher-15-FC0
    member FP-Launcher-16-FC0
    member C480M5-P0
    member xs1-fc0
    member xs2-fc0

  zoneset activate name Launcher_FabricA vsan 3
```

```
do clear zone database vsan 3
!Full Zone Database Section for vsan 3
zone name SP-Launcher-01-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:00
!          [AAD-16-CH3-BL1-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name SP-Launcher-02-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:02
!          [AAD-16-CH3-BL2-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name SP-Launcher-03-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:04
!          [AAD-16-CH3-BL3-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!          [CS700-FC2-1]

zone name SP-Launcher-04-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:06
!          [AAD-16-CH3-BL4-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!          [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!               [CS700-FC2-1]

zone name SP-Launcher-05-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:08
!               [AAD-16-CH3-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-06-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0a
!               [AAD-16-CH3-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-07-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0c
!               [AAD-16-CH3-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-08-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:0e
!               [AAD-16-CH3-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!               [CS700-FC2-1]

zone name SP-Launcher-09-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:10
!               [AAD-16-CH4-BL1-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-10-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:12
!               [AAD-16-CH4-BL2-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-11-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:14
!               [AAD-16-CH4-BL3-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name SP-Launcher-12-FC0 vsan 3
    member pwwn 20:00:00:25:b5:17:aa:16
!               [AAD-16-CH4-BL4-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!           [CS700-FC2-1]

zone name SP-Launcher-13-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:18
!           [AAD-16-CH4-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]

zone name SP-Launcher-14-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1a
!           [AAD-16-CH4-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]

zone name SP-Launcher-15-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1c
!           [AAD-16-CH4-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!           [CS700-FC2-1]

zone name SP-Launcher-16-FC0 vsan 3
   member pwwn 20:00:00:25:b5:17:aa:1e
!           [AAD-16-CH4-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!           [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!               [CS700-FC2-1]

zone name FP-Launcher-01-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:00
!               [AAD-16-CH1-BL1-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-02-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:02
!               [AAD-16-CH1-BL2-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-03-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:04
!               [AAD-16-CH1-BL3-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-04-FC0 vsan 3
    member pwwn 20:00:00:25:b5:9a:a0:06
!               [AAD-16-CH1-BL4-FC0]
    member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!               [CS700-FC2-1]

zone name FP-Launcher-05-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:08
!               [AAD-16-CH1-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-06-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:0a
!               [AAD-16-CH1-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-07-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:0c
!               [AAD-16-CH1-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name FP-Launcher-08-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:0e
!               [AAD-16-CH1-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!             [CS700-FC2-1]

zone name FP-Launcher-09-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:10
!             [AAD-16-CH2-BL1-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-10-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:12
!             [AAD-16-CH2-BL2-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-11-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:14
!             [AAD-16-CH2-BL3-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-12-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:16
!             [AAD-16-CH2-BL4-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!             [CS700-FC2-1]

zone name FP-Launcher-13-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:18
!             [AAD-16-CH2-BL5-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-14-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:1a
!             [AAD-16-CH2-BL6-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-15-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:1c
!             [AAD-16-CH2-BL7-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!             [CS700-FC2-1]

zone name FP-Launcher-16-FC0 vsan 3
   member pwwn 20:00:00:25:b5:9a:a0:1e
!             [AAD-16-CH2-BL8-FC0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!             [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
!               [CS700-FC2-1]

zone name C480M5-P0 vsan 3
   member pwwn 21:00:00:0e:1e:10:a2:c0
!               [C480M5-P0]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name xs1-fc0 vsan 3
   member pwwn 20:00:00:25:b5:0a:00:03
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]

zone name xs2-fc0 vsan 3
   member pwwn 20:00:00:25:b5:0a:00:05
   member pwwn 56:c9:ce:90:0d:e8:24:06
!               [CS700-FC2-1]
   member pwwn 56:c9:ce:90:0d:e8:24:02
!               [CS700-FC1-1]

zoneset name Launcher_FabricA vsan 3
   member SP-Launcher-01-FC0
   member SP-Launcher-02-FC0
   member SP-Launcher-03-FC0
   member SP-Launcher-04-FC0
   member SP-Launcher-05-FC0
   member SP-Launcher-06-FC0
   member SP-Launcher-07-FC0
```

```
    member SP-Launcher-08-FC0
    member SP-Launcher-09-FC0
    member SP-Launcher-10-FC0
    member SP-Launcher-11-FC0
    member SP-Launcher-12-FC0
    member SP-Launcher-13-FC0
    member SP-Launcher-14-FC0
    member SP-Launcher-15-FC0
    member SP-Launcher-16-FC0
    member FP-Launcher-01-FC0
    member FP-Launcher-02-FC0
    member FP-Launcher-03-FC0
    member FP-Launcher-04-FC0
    member FP-Launcher-05-FC0
    member FP-Launcher-06-FC0
    member FP-Launcher-07-FC0
    member FP-Launcher-08-FC0
    member FP-Launcher-09-FC0
    member FP-Launcher-10-FC0
    member FP-Launcher-11-FC0
    member FP-Launcher-12-FC0
    member FP-Launcher-13-FC0
    member FP-Launcher-14-FC0
    member FP-Launcher-15-FC0
    member FP-Launcher-16-FC0
    member C480M5-P0
    member xs1-fc0
    member xs2-fc0

  zoneset name Laucner_FabricA vsan 3
    member xs2-fc0
```

```
!Active Zone Database Section for vsan 100
zone name FlaskStack-VCC-CVD-WLHost01 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
```

```
    member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
```

```
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100
```

```
  member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
```

```
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
```

```
  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]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]
```

```
zone name FlaskStack-VCC-CVD-WLHost16 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
```

```
    member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
```

```
  member pwwn 52:4a:93:75:dd:91:0a:13
!         [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!         [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!         [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!         [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!         [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!         [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!         [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!         [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!         [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 100
```

```
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
```

```
  member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
```

```
    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]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!             [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!             [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!             [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!             [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!             [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!             [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!             [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!             [X70-CT1-FC3]
    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]
```

```
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

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 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!            [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!            [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!            [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!            [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
```

```
!               [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
```

```
!              [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!              [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!              [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!              [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!              [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!              [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!              [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!              [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!              [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!              [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
    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]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100
    member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
    member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
    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]
```

```
zone name FlaskStack-VCC-CVD-WLHost12 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
```

```
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!               [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!               [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!               [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!               [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
```

```
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
```

```
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
```

```
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!             [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!             [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!             [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!             [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!             [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!             [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!             [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!             [X70-CT1-FC3]
   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]
```

```
zone name FlaskStack-VCC-CVD-WLHost27 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
  member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
  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]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 100
  member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:02
```

```
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
!           [X70-CT1-FC3]
   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]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100
   member pwwn 52:4a:93:75:dd:91:0a:00
!           [X70-CT0-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:02
!           [X70-CT0-FC2]
   member pwwn 52:4a:93:75:dd:91:0a:11
!           [X70-CT1-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:13
```

```
!             [X70-CT1-FC3]
   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]

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


!Active Zone Database Section for vsan 400

zone name a300_VDI-1-hba1 vsan 400

    member pwwn 20:00:00:25:b5:3a:00:3f

!               [VDI-1-hba1]

    member pwwn 20:01:00:a0:98:af:bd:e8

!               [a300-01-0g]

    member pwwn 20:03:00:a0:98:af:bd:e8

!               [a300-02-0g]


zone name a300_VDI-2-hba1 vsan 400

    member pwwn 20:01:00:a0:98:af:bd:e8

!               [a300-01-0g]

    member pwwn 20:03:00:a0:98:af:bd:e8

!               [a300-02-0g]

    member pwwn 20:00:00:25:b5:3a:00:0f

!               [VDI-2-hba1]


zone name a300_VDI-3-hba1 vsan 400

    member pwwn 20:01:00:a0:98:af:bd:e8

!               [a300-01-0g]

    member pwwn 20:03:00:a0:98:af:bd:e8

!               [a300-02-0g]

    member pwwn 20:00:00:25:b5:3a:00:1f

!               [VDI-3-hba1]
```

```
zone name a300_VDI-4-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:4e
!               [VDI-4-hba1]

zone name a300_VDI-5-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:2e
!               [VDI-5-hba1]

zone name a300_VDI-6-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:3e
!               [VDI-6-hba1]

zone name a300_VDI-7-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:0e
!               [VDI-7-hba1]
```

```
zone name a300_Infra01-8-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4f
!           [Infra01-8-hba1]

zone name a300_VDI-9-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4d
!           [VDI-9-hba1]

zone name a300_VDI-10-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:2d
!           [VDI-10-hba1]

zone name a300_VDI-11-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:3d
!           [VDI-11-hba1]
```

```
zone name a300_VDI-12-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:0d
!              [VDI-12-hba1]

zone name a300_VDI-13-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:1d
!              [VDI-13-hba1]

zone name a300_VDI-14-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:4c
!              [VDI-14-hba1]

zone name a300_VDI-15-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:2c
!              [VDI-15-hba1]
```

```
zone name a300_Infra02-16-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:2f
!           [Infra02-16-hba1]

zone name a300_VDI-17-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:0c
!           [VDI-17-hba1]

zone name a300_VDI-18-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:1c
!           [VDI-18-hba1]

zone name a300_VDI-19-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4b
!           [VDI-19-hba1]
```

```
zone name a300_VDI-20-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:2b
!               [VDI-20-hba1]

zone name a300_VDI-21-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:3b
!               [VDI-21-hba1]

zone name a300_VDI-22-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:0b
!               [VDI-22-hba1]

zone name a300_VDI-23-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1b
!               [VDI-23-hba1]
```

```
zone name a300_VDI-24-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:4a
!           [VDI-24-hba1]

zone name a300_VDI-25-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:2a
!           [VDI-25-hba1]

zone name a300_VDI-26-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:3a
!           [VDI-26-hba1]

zone name a300_VDI-27-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:0a
!           [VDI-27-hba1]
```

```
zone name a300_VDI-28-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1a
!               [VDI-28-hba1]

zone name a300_VDI-29-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:49
!               [VDI-29-hba1]

zone name a300_VDI-30-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:39
!               [VDI-30-hba1]

zone name a300_VDI-31-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1e
!               [VDI-31-hba1]
```

```
zone name a300_VDI-32-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:3c
!               [VDI-32-hba1]

zoneset name FlexPod_FabricA vsan 400
    member a300_VDI-1-hba1
    member a300_VDI-2-hba1
    member a300_VDI-3-hba1
    member a300_VDI-4-hba1
    member a300_VDI-5-hba1
    member a300_VDI-6-hba1
    member a300_VDI-7-hba1
    member a300_Infra01-8-hba1
    member a300_VDI-9-hba1
    member a300_VDI-10-hba1
    member a300_VDI-11-hba1
    member a300_VDI-12-hba1
    member a300_VDI-13-hba1
    member a300_VDI-14-hba1
    member a300_VDI-15-hba1
    member a300_Infra02-16-hba1
    member a300_VDI-17-hba1
    member a300_VDI-18-hba1
    member a300_VDI-19-hba1
    member a300_VDI-20-hba1
    member a300_VDI-21-hba1
    member a300_VDI-22-hba1
    member a300_VDI-23-hba1
    member a300_VDI-24-hba1
    member a300_VDI-25-hba1
    member a300_VDI-26-hba1
    member a300_VDI-27-hba1
    member a300_VDI-28-hba1
    member a300_VDI-29-hba1
    member a300_VDI-30-hba1
    member a300_VDI-31-hba1
    member a300_VDI-32-hba1

zoneset activate name FlexPod_FabricA vsan 400
do clear zone database vsan 400
!Full Zone Database Section for vsan 400
zone name a300_VDI-1-hba1 vsan 400
    member pwwn 20:00:00:25:b5:3a:00:3f
!               [VDI-1-hba1]
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]

zone name a300_VDI-2-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:0f
!               [VDI-2-hba1]

zone name a300_VDI-3-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:1f
!           [VDI-3-hba1]

zone name a300_VDI-4-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:4e
!           [VDI-4-hba1]

zone name a300_VDI-5-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:2e
!           [VDI-5-hba1]

zone name a300_VDI-6-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!           [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!           [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:3e
!           [VDI-6-hba1]

zone name a300_VDI-7-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!              [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:0e
!              [VDI-7-hba1]

zone name a300_Infra01-8-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1e
!              [VDI-31-hba1]

zone name a300_VDI-9-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:4d
!              [VDI-9-hba1]

zone name a300_VDI-10-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!              [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!              [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:2d
!              [VDI-10-hba1]

zone name a300_VDI-11-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!        [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!        [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:3d
!        [VDI-11-hba1]

zone name a300_VDI-12-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!        [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!        [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:0d
!        [VDI-12-hba1]

zone name a300_VDI-13-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!        [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!        [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:1d
!        [VDI-13-hba1]

zone name a300_VDI-14-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!        [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!        [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4c
!        [VDI-14-hba1]

zone name a300_VDI-15-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:2c
!               [VDI-15-hba1]

zone name a300_Infra02-16-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:2f
!               [Infra02-16-hba1]

zone name a300_VDI-17-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:0c
!               [VDI-17-hba1]

zone name a300_VDI-18-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1c
!               [VDI-18-hba1]

zone name a300_VDI-19-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!          [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!          [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:4b
!          [VDI-19-hba1]

zone name a300_VDI-20-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!          [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!          [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:2b
!          [VDI-20-hba1]

zone name a300_VDI-21-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!          [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!          [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:3b
!          [VDI-21-hba1]

zone name a300_VDI-22-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
!          [a300-01-0g]
  member pwwn 20:03:00:a0:98:af:bd:e8
!          [a300-02-0g]
  member pwwn 20:00:00:25:b5:3a:00:0b
!          [VDI-22-hba1]

zone name a300_VDI-23-hba1 vsan 400
  member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!               [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:1b
!               [VDI-23-hba1]

zone name a300_VDI-24-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:4a
!               [VDI-24-hba1]

zone name a300_VDI-25-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:2a
!               [VDI-25-hba1]

zone name a300_VDI-26-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:3a
!               [VDI-26-hba1]

zone name a300_VDI-27-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!       [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!       [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:0a
!       [VDI-27-hba1]

zone name a300_VDI-28-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!       [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!       [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:1a
!       [VDI-28-hba1]

zone name a300_VDI-29-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!       [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!       [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:49
!       [VDI-29-hba1]

zone name a300_VDI-30-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
!       [a300-01-0g]
   member pwwn 20:03:00:a0:98:af:bd:e8
!       [a300-02-0g]
   member pwwn 20:00:00:25:b5:3a:00:39
!       [VDI-30-hba1]

zone name a300_VDI-31-hba1 vsan 400
   member pwwn 20:01:00:a0:98:af:bd:e8
```

```
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:1e
!               [VDI-31-hba1]

zone name a300_VDI-32-hba1 vsan 400
    member pwwn 20:01:00:a0:98:af:bd:e8
!               [a300-01-0g]
    member pwwn 20:03:00:a0:98:af:bd:e8
!               [a300-02-0g]
    member pwwn 20:00:00:25:b5:3a:00:3c
!               [VDI-32-hba1]

zoneset name FlexPod_FabricA vsan 400
    member a300_VDI-1-hba1
    member a300_VDI-2-hba1
    member a300_VDI-3-hba1
    member a300_VDI-4-hba1
    member a300_VDI-5-hba1
    member a300_VDI-6-hba1
    member a300_VDI-7-hba1
    member a300_Infra01-8-hba1
    member a300_VDI-9-hba1
    member a300_VDI-10-hba1
    member a300_VDI-11-hba1
    member a300_VDI-12-hba1
    member a300_VDI-13-hba1
    member a300_VDI-14-hba1
    member a300_VDI-15-hba1
    member a300_Infra02-16-hba1
    member a300_VDI-17-hba1
    member a300_VDI-18-hba1
    member a300_VDI-19-hba1
    member a300_VDI-20-hba1
    member a300_VDI-21-hba1
    member a300_VDI-22-hba1
    member a300_VDI-23-hba1
    member a300_VDI-24-hba1
    member a300_VDI-25-hba1
    member a300_VDI-26-hba1
    member a300_VDI-27-hba1
    member a300_VDI-28-hba1
    member a300_VDI-29-hba1
    member a300_VDI-30-hba1
    member a300_VDI-31-hba1
    member a300_VDI-32-hba1


interface mgmt0
  ip address 10.29.164.238 255.255.255.0
vsan database
  vsan 3 interface fc1/13
  vsan 3 interface fc1/14
  vsan 3 interface fc1/19
  vsan 3 interface fc1/20
  vsan 3 interface fc1/21
  vsan 3 interface fc1/22
  vsan 3 interface fc1/23
  vsan 3 interface fc1/24
  vsan 100 interface fc1/25
  vsan 100 interface fc1/26
  vsan 100 interface fc1/27
  vsan 100 interface fc1/28
  vsan 100 interface fc1/29
  vsan 100 interface fc1/30
  vsan 100 interface fc1/31
  vsan 100 interface fc1/32
  vsan 400 interface fc1/37
  vsan 400 interface fc1/38
  vsan 400 interface fc1/43
  vsan 400 interface fc1/44
  vsan 400 interface fc1/45
  vsan 400 interface fc1/46
switchname MDS-A
no terminal log-all
line console
  terminal width  80
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.8.1.1.bin
boot system bootflash:/m9100-s5ek9-mz.8.1.1.bin
interface fc1/14
  switchport speed 8000
interface fc1/15
  switchport speed 8000
interface fc1/16
  switchport speed 8000
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/13
interface fc1/17
interface fc1/18
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/47
interface fc1/48
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46

interface fc1/1
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/2
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/3
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/4
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/5
  port-license acquire
  no shutdown

interface fc1/6
  port-license acquire
  no shutdown

interface fc1/7
  port-license acquire
  no shutdown

interface fc1/8
  port-license acquire
  no shutdown

interface fc1/9
  port-license acquire

interface fc1/10
  port-license acquire
```

```
interface fc1/11
  port-license acquire

interface fc1/12
  port-license acquire

interface fc1/13
  port-license acquire
  no shutdown

interface fc1/14
  port-license acquire
  no shutdown

interface fc1/15
  port-license acquire
  no shutdown

interface fc1/16
  port-license acquire
  no shutdown

interface fc1/17
  port-license acquire
  no shutdown

interface fc1/18
  port-license acquire
  no shutdown

interface fc1/19
```

```
  switchport trunk allowed vsan 3
  switchport description CS700 CTRL-A:02
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/20
  switchport trunk allowed vsan 3
  switchport description CS700 CTRL-A:06
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/21
  switchport trunk allowed vsan 3
  switchport description Launcher-FIA
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/22
  switchport trunk allowed vsan 3
  switchport description Launcher-FIA
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/23
  switchport trunk allowed vsan 3
  switchport description Launcher-FIA
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/24
  switchport trunk allowed vsan 3
  switchport description Launcher-FIA
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/25
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/26
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/27
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/28
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/29
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/30
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/31
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/32
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/33
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/34
```

```
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/35
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/36
  switchport trunk allowed vsan 100
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/37
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/38
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/39
  port-license acquire
  no shutdown

interface fc1/40
```

```
  port-license acquire

  no shutdown


interface fc1/41

  port-license acquire

  no shutdown


interface fc1/42

  port-license acquire

  no shutdown


interface fc1/43

  port-license acquire

  no shutdown


interface fc1/44

  port-license acquire

  no shutdown


interface fc1/45

  port-license acquire

  no shutdown


interface fc1/46

  port-license acquire

  no shutdown


interface fc1/47

  port-license acquire

  no shutdown


interface fc1/48
```

```
  port-license acquire
  no shutdown
ip default-gateway 10.29.164.1
```

Cisco MDS 9148S - B Configuration

```
!Command: show running-config
!Time: Sun Jun 24 22:00:37 2018

version 8.1(1)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
no password strength-check
username admin password 5 $1$OPnyy3RN$s8SLqLN3W3JPvf4rEb2CD0  role network-admin
ip domain-lookup
ip host MDS-B  10.29.164.239
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xc9e1af5dbb0bbac72253a1bef037bbbe
 priv 0xc9e1af5dbb0bbac72253a1bef037bbbe localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1164
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
```

```
snmp-server community public group network-operator
vsan database
  vsan 4 name "SP-Launcher-B"
  vsan 101 name "FlashStack-VCC-CVD-Fabric-B"
  vsan 401 name "FlexPod-B"
fcdroplatency network 2000 vsan 1
device-alias database
  device-alias name C480M5-P1 pwwn 21:00:00:0e:1e:10:a2:c1
  device-alias name VDI-1-hba2 pwwn 20:00:00:25:d5:06:00:3f
  device-alias name VDI-2-hba2 pwwn 20:00:00:25:d5:06:00:0f
  device-alias name VDI-3-hba2 pwwn 20:00:00:25:d5:06:00:1f
  device-alias name VDI-4-hba2 pwwn 20:00:00:25:d5:06:00:4e
  device-alias name VDI-5-hba2 pwwn 20:00:00:25:d5:06:00:2e
  device-alias name VDI-6-hba2 pwwn 20:00:00:25:d5:06:00:3e
  device-alias name VDI-7-hba2 pwwn 20:00:00:25:d5:06:00:0e
  device-alias name VDI-9-hba2 pwwn 20:00:00:25:d5:06:00:4d
  device-alias name a300-01-0h pwwn 20:02:00:a0:98:af:bd:e8
  device-alias name a300-02-0h pwwn 20:04:00:a0:98:af:bd:e8
  device-alias name CS700-FC1-2 pwwn 56:c9:ce:90:0d:e8:24:01
  device-alias name CS700-FC2-2 pwwn 56:c9:ce:90:0d:e8:24:05
  device-alias name VDI-10-hba2 pwwn 20:00:00:25:d5:06:00:2d
  device-alias name VDI-11-hba2 pwwn 20:00:00:25:d5:06:00:3d
  device-alias name VDI-12-hba2 pwwn 20:00:00:25:d5:06:00:0d
  device-alias name VDI-13-hba2 pwwn 20:00:00:25:d5:06:00:1d
  device-alias name VDI-14-hba2 pwwn 20:00:00:25:d5:06:00:4c
  device-alias name VDI-15-hba2 pwwn 20:00:00:25:d5:06:00:2c
  device-alias name VDI-17-hba2 pwwn 20:00:00:25:d5:06:00:0c
  device-alias name VDI-18-hba2 pwwn 20:00:00:25:d5:06:00:1c
  device-alias name VDI-19-hba2 pwwn 20:00:00:25:d5:06:00:4b
  device-alias name VDI-20-hba2 pwwn 20:00:00:25:d5:06:00:2b
  device-alias name VDI-21-hba2 pwwn 20:00:00:25:d5:06:00:3b
  device-alias name VDI-22-hba2 pwwn 20:00:00:25:d5:06:00:6b
```

```
device-alias name VDI-23-hba2 pwwn 20:00:00:25:d5:06:00:1b

device-alias name VDI-24-hba2 pwwn 20:00:00:25:d5:06:00:4a

device-alias name VDI-25-hba2 pwwn 20:00:00:25:d5:06:00:2a

device-alias name VDI-26-hba2 pwwn 20:00:00:25:d5:06:00:3a

device-alias name VDI-27-hba2 pwwn 20:00:00:25:d5:06:00:0a

device-alias name VDI-28-hba2 pwwn 20:00:00:25:d5:06:00:1a

device-alias name VDI-29-hba2 pwwn 20:00:00:25:d5:06:00:49

device-alias name VDI-30-hba2 pwwn 20:00:00:25:d5:06:00:39

device-alias name VDI-31-hba2 pwwn 20:00:00:25:d5:06:00:1e

device-alias name VDI-32-hba2 pwwn 20:00:00:25:d5:06:00:3c

device-alias name X70-CT0-FC1 pwwn 52:4a:93:75:dd:91:0a:01

device-alias name X70-CT0-FC3 pwwn 52:4a:93:75:dd:91:0a:03

device-alias name X70-CT1-FC0 pwwn 52:4a:93:75:dd:91:0a:10

device-alias name X70-CT1-FC2 pwwn 52:4a:93:75:dd:91:0a:12

device-alias name Infra01-8-hba2 pwwn 20:00:00:25:d5:06:00:4f

device-alias name Infra02-16-hba2 pwwn 20:00:00:25:d5:06:00:2f

device-alias 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 name AAD-16-CH1-BL1-FC1 pwwn 20:00:00:25:b5:9a:a0:01
device-alias name AAD-16-CH1-BL2-FC1 pwwn 20:00:00:25:b5:9a:a0:03
device-alias name AAD-16-CH1-BL3-FC1 pwwn 20:00:00:25:b5:9a:a0:05
device-alias name AAD-16-CH1-BL4-FC1 pwwn 20:00:00:25:b5:9a:a0:07
device-alias name AAD-16-CH1-BL5-FC1 pwwn 20:00:00:25:b5:9a:a0:09
device-alias name AAD-16-CH1-BL6-FC1 pwwn 20:00:00:25:b5:9a:a0:0b
device-alias name AAD-16-CH1-BL7-FC1 pwwn 20:00:00:25:b5:9a:a0:0d
device-alias name AAD-16-CH1-BL8-FC1 pwwn 20:00:00:25:b5:9a:a0:0f
device-alias name AAD-16-CH2-BL1-FC1 pwwn 20:00:00:25:b5:9a:a0:11
device-alias name AAD-16-CH2-BL2-FC1 pwwn 20:00:00:25:b5:9a:a0:13
device-alias name AAD-16-CH2-BL3-FC1 pwwn 20:00:00:25:b5:9a:a0:15
device-alias name AAD-16-CH2-BL4-FC1 pwwn 20:00:00:25:b5:9a:a0:17
device-alias name AAD-16-CH2-BL5-FC1 pwwn 20:00:00:25:b5:9a:a0:19
device-alias name AAD-16-CH2-BL6-FC1 pwwn 20:00:00:25:b5:9a:a0:1b
device-alias name AAD-16-CH2-BL7-FC1 pwwn 20:00:00:25:b5:9a:a0:1d
device-alias name AAD-16-CH2-BL8-FC1 pwwn 20:00:00:25:b5:9a:a0:1f
```

```
  device-alias name AAD-16-CH3-BL1-FC1 pwwn 20:00:00:25:b5:17:aa:01
  device-alias name AAD-16-CH3-BL2-FC1 pwwn 20:00:00:25:b5:17:aa:03
  device-alias name AAD-16-CH3-BL3-FC1 pwwn 20:00:00:25:b5:17:aa:05
  device-alias name AAD-16-CH3-BL4-FC1 pwwn 20:00:00:25:b5:17:aa:07
  device-alias name AAD-16-CH3-BL5-FC1 pwwn 20:00:00:25:b5:17:aa:09
  device-alias name AAD-16-CH3-BL6-FC1 pwwn 20:00:00:25:b5:17:aa:0b
  device-alias name AAD-16-CH3-BL7-FC1 pwwn 20:00:00:25:b5:17:aa:0d
  device-alias name AAD-16-CH3-BL8-FC1 pwwn 20:00:00:25:b5:17:aa:0f
  device-alias name AAD-16-CH4-BL1-FC1 pwwn 20:00:00:25:b5:17:aa:11
  device-alias name AAD-16-CH4-BL2-FC1 pwwn 20:00:00:25:b5:17:aa:13
  device-alias name AAD-16-CH4-BL3-FC1 pwwn 20:00:00:25:b5:17:aa:15
  device-alias name AAD-16-CH4-BL4-FC1 pwwn 20:00:00:25:b5:17:aa:17
  device-alias name AAD-16-CH4-BL5-FC1 pwwn 20:00:00:25:b5:17:aa:19
  device-alias name AAD-16-CH4-BL6-FC1 pwwn 20:00:00:25:b5:17:aa:1b
  device-alias name AAD-16-CH4-BL7-FC1 pwwn 20:00:00:25:b5:17:aa:1d
  device-alias name AAD-16-CH4-BL8-FC1 pwwn 20:00:00:25:b5:17:aa:1f

device-alias commit

fcdomain fcid database
  vsan 4 wwn 20:00:00:25:b5:17:aa:01 fcid 0x5b1909 dynamic
!           [AAD-16-CH3-BL1-FC1]
  vsan 4 wwn 56:c9:ce:90:0d:e8:24:05 fcid 0x5b1400 dynamic
!           [CS700-FC2-2]
  vsan 401 wwn 20:00:00:25:d5:06:00:2e fcid 0x870503 dynamic
!             [VDI-5-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0e fcid 0x870303 dynamic
!             [VDI-7-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1d fcid 0x870204 dynamic
!             [VDI-13-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:2c fcid 0x870302 dynamic
!             [VDI-15-hba2]
```

```
  vsan 4 wwn 20:00:00:25:b5:17:aa:03 fcid 0x5b1906 dynamic
!           [AAD-16-CH3-BL2-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:05 fcid 0x5b1a01 dynamic
!           [AAD-16-CH3-BL3-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:0f fcid 0x5b190a dynamic
!           [AAD-16-CH3-BL8-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:07 fcid 0x5b190b dynamic
!           [AAD-16-CH3-BL4-FC1]
  vsan 1 wwn 52:4a:93:75:dd:91:0a:02 fcid 0xb61700 dynamic
  vsan 1 wwn 52:4a:93:75:dd:91:0a:03 fcid 0xb61800 dynamic
!           [X70-CT0-FC3]
  vsan 4 wwn 20:00:00:25:b5:17:aa:09 fcid 0x5b1904 dynamic
!           [AAD-16-CH3-BL5-FC1]
  vsan 1 wwn 52:4a:93:75:dd:91:0a:12 fcid 0xb61900 dynamic
!           [X70-CT1-FC2]
  vsan 4 wwn 20:00:00:25:b5:17:aa:0d fcid 0x5b1c01 dynamic
!           [AAD-16-CH3-BL7-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:0b fcid 0x5b1b09 dynamic
!           [AAD-16-CH3-BL6-FC1]
  vsan 4 wwn 20:4d:00:de:fb:18:3c:00 fcid 0x5b1900 dynamic
  vsan 4 wwn 20:00:00:25:b5:17:aa:17 fcid 0x5b1c0a dynamic
!           [AAD-16-CH4-BL4-FC1]
  vsan 4 wwn 20:4f:00:de:fb:18:3c:00 fcid 0x5b1a00 dynamic
  vsan 4 wwn 20:4e:00:de:fb:18:3c:00 fcid 0x5b1b00 dynamic
  vsan 4 wwn 20:50:00:de:fb:18:3c:00 fcid 0x5b1c00 dynamic
  vsan 1 wwn 20:01:00:de:fb:90:a4:40 fcid 0xb61300 dynamic
  vsan 1 wwn 20:03:00:de:fb:90:a4:40 fcid 0xb61400 dynamic
  vsan 1 wwn 20:04:00:de:fb:90:a4:40 fcid 0xb61500 dynamic
  vsan 1 wwn 20:02:00:de:fb:90:a4:40 fcid 0xb61600 dynamic
  vsan 401 wwn 20:00:00:25:d5:06:00:4c fcid 0x870409 dynamic
!             [VDI-14-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:2f fcid 0x870403 dynamic
```

```
!          [Infra02-16-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0c fcid 0x870202 dynamic
!          [VDI-17-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1b fcid 0x870507 dynamic
!          [VDI-23-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:39 fcid 0x870407 dynamic
!          [VDI-30-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:3f fcid 0x870305 dynamic
!          [VDI-1-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1e fcid 0x870203 dynamic
!          [VDI-31-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1c fcid 0x870405 dynamic
!          [VDI-18-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:2a fcid 0x870308 dynamic
!          [VDI-25-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1f fcid 0x870309 dynamic
!          [VDI-3-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:4e fcid 0x870307 dynamic
!          [VDI-4-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:1a fcid 0x870502 dynamic
!          [VDI-28-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:3e fcid 0x870208 dynamic
!          [VDI-6-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0d fcid 0x870505 dynamic
!          [VDI-12-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:3d fcid 0x870506 dynamic
!          [VDI-11-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:3a fcid 0x870206 dynamic
!          [VDI-26-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0a fcid 0x870501 dynamic
!          [VDI-27-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:4b fcid 0x870304 dynamic
```

```
!             [VDI-19-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0f fcid 0x87020a dynamic
!             [VDI-2-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:49 fcid 0x870408 dynamic
!             [VDI-29-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:3b fcid 0x870306 dynamic
!             [VDI-21-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:0b fcid 0x870504 dynamic
  vsan 401 wwn 20:00:00:25:d5:06:00:4f fcid 0x870508 dynamic
!             [Infra01-8-hba2]
  vsan 1 wwn 52:4a:93:75:dd:91:0a:13 fcid 0xb61a00 dynamic
  vsan 4 wwn 20:00:00:25:b5:17:aa:15 fcid 0x5b1902 dynamic
!           [AAD-16-CH4-BL3-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:13 fcid 0x5b1c03 dynamic
!           [AAD-16-CH4-BL2-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:1b fcid 0x5b1b02 dynamic
!           [AAD-16-CH4-BL6-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:19 fcid 0x5b1903 dynamic
!           [AAD-16-CH4-BL5-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:1d fcid 0x5b1b05 dynamic
!           [AAD-16-CH4-BL7-FC1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:1f fcid 0x5b1905 dynamic
!           [AAD-16-CH4-BL8-FC1]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:02 fcid 0x2e0000 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:03 fcid 0x2e0100 dynamic
!             [X70-CT0-FC3]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:12 fcid 0x2e0200 dynamic
!             [X70-CT1-FC2]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:13 fcid 0x2e0300 dynamic
  vsan 101 wwn 20:04:00:de:fb:90:a4:40 fcid 0x2e0400 dynamic
  vsan 101 wwn 20:02:00:de:fb:90:a4:40 fcid 0x2e0500 dynamic
  vsan 101 wwn 20:03:00:de:fb:90:a4:40 fcid 0x2e0600 dynamic
```

```
   vsan 101 wwn 20:01:00:de:fb:90:a4:40 fcid 0x2e0700 dynamic
   vsan 101 wwn 20:00:00:25:b5:bb:17:1e fcid 0x2e0406 dynamic
 !          [VCC-Infra01-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:02 fcid 0x2e0709 dynamic
 !          [VCC-WLHost02-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:0a fcid 0x2e0602 dynamic
 !          [VCC-WLHost06-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:0e fcid 0x2e0604 dynamic
 !          [VCC-WLHost08-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:26 fcid 0x2e060a dynamic
 !          [VCC-WLHost19-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:2e fcid 0x2e040d dynamic
 !          [VCC-WLHost23-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:22 fcid 0x2e0707 dynamic
 !          [VCC-WLHost17-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:28 fcid 0x2e0705 dynamic
 !          [VCC-WLHost20-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:24 fcid 0x2e060f dynamic
 !          [VCC-WLHost18-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:06 fcid 0x2e040e dynamic
 !          [VCC-WLHost04-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:0c fcid 0x2e060b dynamic
 !          [VCC-WLHost07-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:08 fcid 0x2e0505 dynamic
 !          [VCC-WLHost05-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:00 fcid 0x2e0710 dynamic
 !          [VCC-WLHost01-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:04 fcid 0x2e0706 dynamic
 !          [VCC-WLHost03-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:16 fcid 0x2e0608 dynamic
 !          [VCC-WLHost12-HBA1]
   vsan 101 wwn 20:00:00:25:b5:bb:17:10 fcid 0x2e0410 dynamic
```

```
!          [VCC-WLHost09-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:12 fcid 0x2e0603 dynamic
!          [VCC-WLHost10-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:18 fcid 0x2e0704 dynamic
!          [VCC-WLHost13-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:20 fcid 0x2e0407 dynamic
!          [VCC-WLHost16-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:2c fcid 0x2e050b dynamic
!          [VCC-WLHost22-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:2a fcid 0x2e0511 dynamic
!          [VCC-WLHost21-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:1a fcid 0x2e0405 dynamic
!          [VCC-WLHost14-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:1c fcid 0x2e0601 dynamic
!          [VCC-WLHost15-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:14 fcid 0x2e0408 dynamic
!          [VCC-WLHost11-HBA1]
  vsan 4 wwn 20:00:00:25:b5:17:aa:11 fcid 0x5b1a02 dynamic
!         [AAD-16-CH4-BL1-FC1]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:07 fcid 0x2e0800 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:06 fcid 0x2e0900 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:16 fcid 0x2e0a00 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:17 fcid 0x2e0b00 dynamic
  vsan 101 wwn 20:00:00:25:b5:bb:17:3e fcid 0x2e0402 dynamic
!          [VCC-Infra02-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3f fcid 0x2e0609 dynamic
!          [VCC-Infra02-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:1f fcid 0x2e0611 dynamic
!          [VCC-Infra01-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:03 fcid 0x2e050f dynamic
!          [VCC-WLHost02-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:01 fcid 0x2e040f dynamic
```

```
!         [VCC-WLHost01-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:05 fcid 0x2e060d dynamic
!         [VCC-WLHost03-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:07 fcid 0x2e0513 dynamic
!         [VCC-WLHost04-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:0b fcid 0x2e060c dynamic
!         [VCC-WLHost06-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:09 fcid 0x2e0708 dynamic
!         [VCC-WLHost05-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:0d fcid 0x2e0607 dynamic
!         [VCC-WLHost07-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:0f fcid 0x2e070d dynamic
!         [VCC-WLHost08-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:21 fcid 0x2e0612 dynamic
!         [VCC-WLHost16-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:23 fcid 0x2e0404 dynamic
!         [VCC-WLHost17-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:25 fcid 0x2e070a dynamic
!         [VCC-WLHost18-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:27 fcid 0x2e0409 dynamic
!         [VCC-WLHost19-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:29 fcid 0x2e040b dynamic
!         [VCC-WLHost20-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:2b fcid 0x2e0510 dynamic
!         [VCC-WLHost21-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:2d fcid 0x2e0509 dynamic
!         [VCC-WLHost22-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:2f fcid 0x2e060e dynamic
!         [VCC-WLHost23-HBA3]
  vsan 401 wwn 20:00:00:25:d5:06:00:4d fcid 0x870406 dynamic
!         [VDI-9-hba2]
  vsan 401 wwn 20:00:00:25:d5:06:00:2d fcid 0x870201 dynamic
```

```
!          [VDI-10-hba2]
 vsan 401 wwn 20:00:00:25:d5:06:00:3c fcid 0x870402 dynamic
!          [VDI-32-hba2]
 vsan 101 wwn 20:00:00:25:b5:bb:17:11 fcid 0x2e0503 dynamic
!          [VCC-WLHost09-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:13 fcid 0x2e0701 dynamic
!          [VCC-WLHost10-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:15 fcid 0x2e050e dynamic
!          [VCC-WLHost11-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:17 fcid 0x2e050a dynamic
!          [VCC-WLHost12-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:19 fcid 0x2e070b dynamic
!          [VCC-WLHost13-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:1b fcid 0x2e0702 dynamic
!          [VCC-WLHost14-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:1d fcid 0x2e0403 dynamic
!          [VCC-WLHost15-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:34 fcid 0x2e0507 dynamic
!          [VCC-WLHost26-HBA1]
 vsan 101 wwn 20:00:00:25:b5:bb:17:32 fcid 0x2e0401 dynamic
!          [VCC-WLHost25-HBA1]
 vsan 101 wwn 20:00:00:25:b5:bb:17:33 fcid 0x2e070c dynamic
!          [VCC-WLHost25-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:35 fcid 0x2e040c dynamic
!          [VCC-WLHost26-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:38 fcid 0x2e0506 dynamic
!          [VCC-WLHost28-HBA1]
 vsan 101 wwn 20:00:00:25:b5:bb:17:39 fcid 0x2e070e dynamic
!          [VCC-WLHost28-HBA3]
 vsan 101 wwn 20:00:00:25:b5:bb:17:30 fcid 0x2e050c dynamic
!          [VCC-WLHost24-HBA1]
 vsan 101 wwn 20:00:00:25:b5:bb:17:3a fcid 0x2e040a dynamic
```

```
!          [VCC-WLHost29-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:36 fcid 0x2e0703 dynamic
!          [VCC-WLHost27-HBA1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3c fcid 0x2e0712 dynamic
!          [VCC-WLHost30-HBA1]
  vsan 401 wwn 20:00:00:25:d5:06:00:2b fcid 0x870205 dynamic
!          [VDI-20-hba2]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3d fcid 0x2e0501 dynamic
!          [VCC-WLHost30-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:3b fcid 0x2e0610 dynamic
!          [VCC-WLHost29-HBA3]
  vsan 101 wwn 20:00:00:25:b5:bb:17:37 fcid 0x2e0502 dynamic
!          [VCC-WLHost27-HBA3]
  vsan 401 wwn 20:00:00:25:d5:06:00:4a fcid 0x870401 dynamic
!          [VDI-24-hba2]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:01 fcid 0x5b1b0b dynamic
!        [AAD-16-CH1-BL1-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:03 fcid 0x5b1b03 dynamic
!        [AAD-16-CH1-BL2-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:07 fcid 0x5b1c02 dynamic
!        [AAD-16-CH1-BL4-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:09 fcid 0x5b1a0b dynamic
!        [AAD-16-CH1-BL5-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:0b fcid 0x5b1b06 dynamic
!        [AAD-16-CH1-BL6-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:0f fcid 0x5b1c06 dynamic
!        [AAD-16-CH1-BL8-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:13 fcid 0x5b1a06 dynamic
!        [AAD-16-CH2-BL2-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:0d fcid 0x5b1a04 dynamic
!        [AAD-16-CH1-BL7-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:11 fcid 0x5b1a09 dynamic
```

```
!        [AAD-16-CH2-BL1-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:05 fcid 0x5b1b07 dynamic
!        [AAD-16-CH1-BL3-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:1d fcid 0x5b1a07 dynamic
!        [AAD-16-CH2-BL7-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:15 fcid 0x5b1c07 dynamic
!        [AAD-16-CH2-BL3-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:17 fcid 0x5b1b01 dynamic
!        [AAD-16-CH2-BL4-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:19 fcid 0x5b1c05 dynamic
!        [AAD-16-CH2-BL5-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:1b fcid 0x5b1a08 dynamic
!        [AAD-16-CH2-BL6-FC1]
  vsan 4 wwn 20:00:00:25:b5:9a:a0:1f fcid 0x5b1c08 dynamic
!        [AAD-16-CH2-BL8-FC1]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:01 fcid 0x2e0c00 dynamic
!          [X70-CT0-FC1]
  vsan 101 wwn 52:4a:93:75:dd:91:0a:11 fcid 0x2e0d00 dynamic
  vsan 101 wwn 52:4a:93:75:dd:91:0a:10 fcid 0x2e0e00 dynamic
!          [X70-CT1-FC0]
  vsan 4 wwn 21:00:00:0e:1e:10:a2:c1 fcid 0x5b0000 dynamic
!        [C480M5-P1]
  vsan 101 wwn 20:00:00:25:b5:bb:17:31 fcid 0x2e0504 dynamic
!          [VCC-WLHost24-HBA3]
  vsan 1 wwn 20:01:00:de:fb:92:93:80 fcid 0xb60000 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:04 fcid 0xb60001 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:00 fcid 0xb60002 dynamic
  vsan 1 wwn 20:00:00:25:b5:0a:00:02 fcid 0xb60003 dynamic
  vsan 401 wwn 50:0a:09:83:80:d3:67:d3 fcid 0x870000 dynamic
  vsan 401 wwn 20:04:00:a0:98:af:bd:e8 fcid 0x870001 dynamic
!          [a300-02-0h]
  vsan 401 wwn 50:0a:09:83:80:13:41:27 fcid 0x870100 dynamic
  vsan 401 wwn 20:02:00:a0:98:af:bd:e8 fcid 0x870101 dynamic
!            [a300-01-0h]
  vsan 401 wwn 20:01:00:de:fb:92:0c:80 fcid 0x870200 dynamic
  vsan 401 wwn 20:02:00:de:fb:92:0c:80 fcid 0x870300 dynamic
  vsan 401 wwn 20:03:00:de:fb:92:0c:80 fcid 0x870400 dynamic
  vsan 401 wwn 20:04:00:de:fb:92:0c:80 fcid 0x870500 dynamic
  vsan 4 wwn 56:c9:ce:90:0d:e8:24:01 fcid 0x5b1300 dynamic
!          [CS700-FC1-2]
!Active Zone Database Section for vsan 4
zone name SP-Launcher-01-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:01
!             [AAD-16-CH3-BL1-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!             [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!             [CS700-FC2-2]

zone name SP-Launcher-02-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:03
!             [AAD-16-CH3-BL2-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!             [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!             [CS700-FC2-2]

zone name SP-Launcher-03-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:05
!             [AAD-16-CH3-BL3-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!             [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!             [CS700-FC2-2]
```

```
zone name SP-Launcher-04-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:07
!              [AAD-16-CH3-BL4-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name SP-Launcher-05-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:09
!              [AAD-16-CH3-BL5-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name SP-Launcher-06-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:0b
!              [AAD-16-CH3-BL6-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name SP-Launcher-07-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:0d
!              [AAD-16-CH3-BL7-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]
```

```
zone name SP-Launcher-08-FC1 vsan 4
  member pwwn 20:00:00:25:b5:17:aa:0f
!           [AAD-16-CH3-BL8-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name SP-Launcher-09-FC1 vsan 4
  member pwwn 20:00:00:25:b5:17:aa:11
!           [AAD-16-CH4-BL1-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name SP-Launcher-10-FC1 vsan 4
  member pwwn 20:00:00:25:b5:17:aa:13
!           [AAD-16-CH4-BL2-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name SP-Launcher-11-FC1 vsan 4
  member pwwn 20:00:00:25:b5:17:aa:15
!           [AAD-16-CH4-BL3-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]
```

```
zone name SP-Launcher-12-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:17
!            [AAD-16-CH4-BL4-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name SP-Launcher-13-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:19
!            [AAD-16-CH4-BL5-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name SP-Launcher-14-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:1b
!            [AAD-16-CH4-BL6-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name SP-Launcher-15-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:1d
!            [AAD-16-CH4-BL7-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]
```

```
zone name SP-Launcher-16-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:1f
!                [AAD-16-CH4-BL8-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!                [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!                [CS700-FC2-2]

zone name FP-Launcher-01-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:01
!                [AAD-16-CH1-BL1-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!                [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!                [CS700-FC2-2]

zone name FP-Launcher-02-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:03
!                [AAD-16-CH1-BL2-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!                [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!                [CS700-FC2-2]

zone name FP-Launcher-03-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:05
!                [AAD-16-CH1-BL3-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!                [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!                [CS700-FC2-2]
```

```
zone name FP-Launcher-04-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:07
!            [AAD-16-CH1-BL4-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name FP-Launcher-05-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:09
!            [AAD-16-CH1-BL5-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name FP-Launcher-06-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:0b
!            [AAD-16-CH1-BL6-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]

zone name FP-Launcher-07-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:0d
!            [AAD-16-CH1-BL7-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!            [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!            [CS700-FC2-2]
```

```
zone name FP-Launcher-08-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:0f
!              [AAD-16-CH1-BL8-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name FP-Launcher-09-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:11
!              [AAD-16-CH2-BL1-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name FP-Launcher-10-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:13
!              [AAD-16-CH2-BL2-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]

zone name FP-Launcher-11-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:15
!              [AAD-16-CH2-BL3-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!              [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!              [CS700-FC2-2]
```

```
zone name FP-Launcher-12-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:17
!          [AAD-16-CH2-BL4-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!          [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!          [CS700-FC2-2]

zone name FP-Launcher-13-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:19
!          [AAD-16-CH2-BL5-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!          [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!          [CS700-FC2-2]

zone name FP-Launcher-14-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:1b
!          [AAD-16-CH2-BL6-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!          [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!          [CS700-FC2-2]

zone name FP-Launcher-15-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:1d
!          [AAD-16-CH2-BL7-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!          [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!          [CS700-FC2-2]
```

```
zone name FP-Launcher-16-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:1f
!               [AAD-16-CH2-BL8-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name C480M5-P1 vsan 4
    member pwwn 21:00:00:0e:1e:10:a2:c1
!               [C480M5-P1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zoneset name Launcher_FabricB vsan 4
    member SP-Launcher-01-FC1
    member SP-Launcher-02-FC1
    member SP-Launcher-03-FC1
    member SP-Launcher-04-FC1
    member SP-Launcher-05-FC1
    member SP-Launcher-06-FC1
    member SP-Launcher-07-FC1
    member SP-Launcher-08-FC1
    member SP-Launcher-09-FC1
    member SP-Launcher-10-FC1
    member SP-Launcher-11-FC1
    member SP-Launcher-12-FC1
    member SP-Launcher-13-FC1
    member SP-Launcher-14-FC1
    member SP-Launcher-15-FC1
    member SP-Launcher-16-FC1
    member FP-Launcher-01-FC1
    member FP-Launcher-02-FC1
    member FP-Launcher-03-FC1
    member FP-Launcher-04-FC1
    member FP-Launcher-05-FC1
    member FP-Launcher-06-FC1
    member FP-Launcher-07-FC1
    member FP-Launcher-08-FC1
    member FP-Launcher-09-FC1
    member FP-Launcher-10-FC1
    member FP-Launcher-11-FC1
    member FP-Launcher-12-FC1
    member FP-Launcher-13-FC1
    member FP-Launcher-14-FC1
    member FP-Launcher-15-FC1
    member FP-Launcher-16-FC1
    member C480M5-P1

zoneset activate name Launcher_FabricB vsan 4
do clear zone database vsan 4
!Full Zone Database Section for vsan 4
zone name SP-Launcher-01-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:01
!               [AAD-16-CH3-BL1-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-02-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:17:aa:03
!               [AAD-16-CH3-BL2-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-03-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:05
!               [AAD-16-CH3-BL3-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-04-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:07
!               [AAD-16-CH3-BL4-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-05-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:09
!               [AAD-16-CH3-BL5-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-06-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:17:aa:0b
!               [AAD-16-CH3-BL6-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-07-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:0d
!               [AAD-16-CH3-BL7-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-08-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:0f
!               [AAD-16-CH3-BL8-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-09-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:11
!               [AAD-16-CH4-BL1-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-10-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:17:aa:13
!               [AAD-16-CH4-BL2-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-11-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:15
!               [AAD-16-CH4-BL3-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-12-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:17
!               [AAD-16-CH4-BL4-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-13-FC1 vsan 4
    member pwwn 20:00:00:25:b5:17:aa:19
!               [AAD-16-CH4-BL5-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name SP-Launcher-14-FC1 vsan 4
```

```
   member pwwn 20:00:00:25:b5:17:aa:1b
!           [AAD-16-CH4-BL6-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name SP-Launcher-15-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:1d
!           [AAD-16-CH4-BL7-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name SP-Launcher-16-FC1 vsan 4
   member pwwn 20:00:00:25:b5:17:aa:1f
!           [AAD-16-CH4-BL8-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-01-FC1 vsan 4
   member pwwn 20:00:00:25:b5:9a:a0:01
!           [AAD-16-CH1-BL1-FC1]
   member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
   member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-02-FC1 vsan 4
```

```
  member pwwn 20:00:00:25:b5:9a:a0:03
!           [AAD-16-CH1-BL2-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-03-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:05
!           [AAD-16-CH1-BL3-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-04-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:07
!           [AAD-16-CH1-BL4-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-05-FC1 vsan 4
  member pwwn 20:00:00:25:b5:9a:a0:09
!           [AAD-16-CH1-BL5-FC1]
  member pwwn 56:c9:ce:90:0d:e8:24:01
!           [CS700-FC1-2]
  member pwwn 56:c9:ce:90:0d:e8:24:05
!           [CS700-FC2-2]

zone name FP-Launcher-06-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:9a:a0:0b
!               [AAD-16-CH1-BL6-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-07-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:0d
!               [AAD-16-CH1-BL7-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-08-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:0f
!               [AAD-16-CH1-BL8-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-09-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:11
!               [AAD-16-CH2-BL1-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-10-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:9a:a0:13
!               [AAD-16-CH2-BL2-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-11-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:15
!               [AAD-16-CH2-BL3-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-12-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:17
!               [AAD-16-CH2-BL4-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-13-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:19
!               [AAD-16-CH2-BL5-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-14-FC1 vsan 4
```

```
    member pwwn 20:00:00:25:b5:9a:a0:1b
!               [AAD-16-CH2-BL6-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-15-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:1d
!               [AAD-16-CH2-BL7-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name FP-Launcher-16-FC1 vsan 4
    member pwwn 20:00:00:25:b5:9a:a0:1f
!               [AAD-16-CH2-BL8-FC1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zone name C480M5-P1 vsan 4
    member pwwn 21:00:00:0e:1e:10:a2:c1
!               [C480M5-P1]
    member pwwn 56:c9:ce:90:0d:e8:24:01
!               [CS700-FC1-2]
    member pwwn 56:c9:ce:90:0d:e8:24:05
!               [CS700-FC2-2]

zoneset name Launcher_FabricB vsan 4
```

```
    member SP-Launcher-01-FC1
    member SP-Launcher-02-FC1
    member SP-Launcher-03-FC1
    member SP-Launcher-04-FC1
    member SP-Launcher-05-FC1
    member SP-Launcher-06-FC1
    member SP-Launcher-07-FC1
    member SP-Launcher-08-FC1
    member SP-Launcher-09-FC1
    member SP-Launcher-10-FC1
    member SP-Launcher-11-FC1
    member SP-Launcher-12-FC1
    member SP-Launcher-13-FC1
    member SP-Launcher-14-FC1
    member SP-Launcher-15-FC1
    member SP-Launcher-16-FC1
    member FP-Launcher-01-FC1
    member FP-Launcher-02-FC1
    member FP-Launcher-03-FC1
    member FP-Launcher-04-FC1
    member FP-Launcher-05-FC1
    member FP-Launcher-06-FC1
    member FP-Launcher-07-FC1
    member FP-Launcher-08-FC1
    member FP-Launcher-09-FC1
    member FP-Launcher-10-FC1
    member FP-Launcher-11-FC1
    member FP-Launcher-12-FC1
    member FP-Launcher-13-FC1
    member FP-Launcher-14-FC1
    member FP-Launcher-15-FC1
    member FP-Launcher-16-FC1
    member C480M5-P1

!Active Zone Database Section for vsan 101
zone name FlaskStack-VCC-CVD-WLHost01 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101
```

```
  member pwwn 52:4a:93:75:dd:91:0a:01
!        [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!        [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!        [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!        [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!        [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!        [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!        [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!        [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!        [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!        [X70-CT0-FC3]
```

```
  member pwwn 52:4a:93:75:dd:91:0a:10
!         [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!         [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!         [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!         [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!         [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!         [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!         [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!         [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!         [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!         [X70-CT1-FC2]
```

```
    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]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]
```

```
zone name FlaskStack-VCC-CVD-WLHost10 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
```

```
    member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
```

```
  member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101
```

```
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
```

```
  member pwwn 52:4a:93:75:dd:91:0a:10
!          [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!          [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!          [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!          [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!          [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!          [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!          [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!          [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!          [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!          [X70-CT1-FC2]
```

```
   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]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
   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]
```

```
zone name FlaskStack-VCC-CVD-WLHost25 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
```

```
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
```

```
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!            [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!            [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!            [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!            [X70-CT1-FC2]
   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]

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

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 52:4a:93:75:dd:91:0a:01
!        [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!        [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!        [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!        [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!        [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!        [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!        [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!        [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]
```

```
zone name FlaskStack-VCC-CVD-WLHost05 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
```

```
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
```

```
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!          [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!          [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!          [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!          [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!          [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!          [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!          [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!          [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
```

```
!               [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
```

```
!             [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
!             [X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
!             [X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
!             [X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
!             [X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]
```

```
zone name FlaskStack-VCC-CVD-WLHost20 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
!           [X70-CT0-FC3]
  member pwwn 52:4a:93:75:dd:91:0a:10
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 101
  member pwwn 52:4a:93:75:dd:91:0a:01
!           [X70-CT0-FC1]
  member pwwn 52:4a:93:75:dd:91:0a:03
```

```
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
```

```
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
!               [X70-CT0-FC1]
    member pwwn 52:4a:93:75:dd:91:0a:03
!               [X70-CT0-FC3]
    member pwwn 52:4a:93:75:dd:91:0a:10
!               [X70-CT1-FC0]
    member pwwn 52:4a:93:75:dd:91:0a:12
!               [X70-CT1-FC2]
    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]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101
    member pwwn 52:4a:93:75:dd:91:0a:01
```

```
! 		[X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
! 		[X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
! 		[X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
! 		[X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
! 		[X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
! 		[X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
! 		[X70-CT1-FC0]
   member pwwn 52:4a:93:75:dd:91:0a:12
! 		[X70-CT1-FC2]
   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]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101
   member pwwn 52:4a:93:75:dd:91:0a:01
! 		[X70-CT0-FC1]
   member pwwn 52:4a:93:75:dd:91:0a:03
! 		[X70-CT0-FC3]
   member pwwn 52:4a:93:75:dd:91:0a:10
```

```
!           [X70-CT1-FC0]
  member pwwn 52:4a:93:75:dd:91:0a:12
!           [X70-CT1-FC2]
  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]

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

!Active Zone Database Section for vsan 401
zone name a300_VDI-1-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3f
!               [VDI-1-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-2-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0f
!               [VDI-2-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-3-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1f
!               [VDI-3-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-4-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4e
!               [VDI-4-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-5-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2e
!               [VDI-5-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-6-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3e
!               [VDI-6-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-7-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0e
!               [VDI-7-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_Infra01-8-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4f
!               [Infra01-8-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-9-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4d
!               [VDI-9-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-10-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2d
!               [VDI-10-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-11-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3d
!               [VDI-11-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-12-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0d
!               [VDI-12-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-13-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1d
!               [VDI-13-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-14-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4c
!               [VDI-14-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-15-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2c
!               [VDI-15-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_Infra02-16-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2f
!               [Infra02-16-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-17-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0c
!               [VDI-17-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-18-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1c
!               [VDI-18-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-19-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4b
!               [VDI-19-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-20-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2b
!               [VDI-20-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-21-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3b
!               [VDI-21-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-22-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:6b
!               [VDI-22-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-23-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1b
!               [VDI-23-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-24-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4a
!               [VDI-24-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-25-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2a
!               [VDI-25-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-26-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3a
!               [VDI-26-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-27-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0a
!               [VDI-27-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
```

```
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-28-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:1a
!           [VDI-28-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-29-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:49
!           [VDI-29-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-30-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:39
!           [VDI-30-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-31-hba2 vsan 401
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
  member pwwn 20:00:00:25:d5:06:00:1e
!        [VDI-31-hba2]

zone name a300_VDI-32-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:3c
!        [VDI-32-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!        [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!        [a300-02-0h]

zoneset name FlexPod_FabricB vsan 401
  member a300_VDI-1-hba2
  member a300_VDI-2-hba2
  member a300_VDI-3-hba2
  member a300_VDI-4-hba2
  member a300_VDI-5-hba2
  member a300_VDI-6-hba2
  member a300_VDI-7-hba2
  member a300_Infra01-8-hba2
  member a300_VDI-9-hba2
  member a300_VDI-10-hba2
  member a300_VDI-11-hba2
  member a300_VDI-12-hba2
  member a300_VDI-13-hba2
  member a300_VDI-14-hba2
  member a300_VDI-15-hba2
  member a300_Infra02-16-hba2
  member a300_VDI-17-hba2
  member a300_VDI-18-hba2
  member a300_VDI-19-hba2
  member a300_VDI-20-hba2
  member a300_VDI-21-hba2
  member a300_VDI-22-hba2
  member a300_VDI-23-hba2
  member a300_VDI-24-hba2
  member a300_VDI-25-hba2
  member a300_VDI-26-hba2
  member a300_VDI-27-hba2
  member a300_VDI-28-hba2
  member a300_VDI-29-hba2
  member a300_VDI-30-hba2
  member a300_VDI-31-hba2
  member a300_VDI-32-hba2

zoneset activate name FlexPod_FabricB vsan 401
do clear zone database vsan 401
!Full Zone Database Section for vsan 401
zone name a300_VDI-1-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:3f
!         [VDI-1-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!         [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!         [a300-02-0h]

zone name a300_VDI-2-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0f
!         [VDI-2-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!         [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!         [a300-02-0h]
```

```
zone name a300_VDI-3-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1f
!               [VDI-3-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-4-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4e
!               [VDI-4-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-5-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2e
!               [VDI-5-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-6-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3e
!               [VDI-6-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]
```

```
zone name a300_VDI-7-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:0e
!               [VDI-7-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_Infra01-8-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:1e
!               [VDI-31-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-9-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4d
!               [VDI-9-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-10-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2d
!               [VDI-10-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]
```

```
zone name a300_VDI-11-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:3d
!          [VDI-11-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-12-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0d
!          [VDI-12-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-13-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:1d
!          [VDI-13-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-14-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:4c
!          [VDI-14-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-15-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:2c
!           [VDI-15-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_Infra02-16-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:2f
!           [Infra02-16-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-17-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:0c
!           [VDI-17-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zone name a300_VDI-18-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:1c
!           [VDI-18-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
```

```
zone name a300_VDI-19-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:4b
!               [VDI-19-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-20-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:2b
!               [VDI-20-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-21-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:3b
!               [VDI-21-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]

zone name a300_VDI-22-hba2 vsan 401
    member pwwn 20:00:00:25:d5:06:00:6b
!               [VDI-22-hba2]
    member pwwn 20:02:00:a0:98:af:bd:e8
!               [a300-01-0h]
    member pwwn 20:04:00:a0:98:af:bd:e8
!               [a300-02-0h]
```

```
zone name a300_VDI-23-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:1b
!              [VDI-23-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!              [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!              [a300-02-0h]

zone name a300_VDI-24-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:4a
!              [VDI-24-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!              [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!              [a300-02-0h]

zone name a300_VDI-25-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:2a
!              [VDI-25-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!              [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!              [a300-02-0h]

zone name a300_VDI-26-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:3a
!              [VDI-26-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!              [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!              [a300-02-0h]
```

```
zone name a300_VDI-27-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:0a
!          [VDI-27-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-28-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:1a
!          [VDI-28-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-29-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:49
!          [VDI-29-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]

zone name a300_VDI-30-hba2 vsan 401
  member pwwn 20:00:00:25:d5:06:00:39
!          [VDI-30-hba2]
  member pwwn 20:02:00:a0:98:af:bd:e8
!          [a300-01-0h]
  member pwwn 20:04:00:a0:98:af:bd:e8
!          [a300-02-0h]
```

```
zone name a300_VDI-31-hba2 vsan 401
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]
   member pwwn 20:00:00:25:d5:06:00:1e
!           [VDI-31-hba2]

zone name a300_VDI-32-hba2 vsan 401
   member pwwn 20:00:00:25:d5:06:00:3c
!           [VDI-32-hba2]
   member pwwn 20:02:00:a0:98:af:bd:e8
!           [a300-01-0h]
   member pwwn 20:04:00:a0:98:af:bd:e8
!           [a300-02-0h]

zoneset name FlexPod_FabricB vsan 401
   member a300_VDI-1-hba2
   member a300_VDI-2-hba2
   member a300_VDI-3-hba2
   member a300_VDI-4-hba2
   member a300_VDI-5-hba2
   member a300_VDI-6-hba2
   member a300_VDI-7-hba2
   member a300_Infra01-8-hba2
   member a300_VDI-9-hba2
   member a300_VDI-10-hba2
   member a300_VDI-11-hba2
   member a300_VDI-12-hba2
   member a300_VDI-13-hba2
   member a300_VDI-14-hba2
   member a300_VDI-15-hba2
```

```
    member a300_Infra02-16-hba2
    member a300_VDI-17-hba2
    member a300_VDI-18-hba2
    member a300_VDI-19-hba2
    member a300_VDI-20-hba2
    member a300_VDI-21-hba2
    member a300_VDI-22-hba2
    member a300_VDI-23-hba2
    member a300_VDI-24-hba2
    member a300_VDI-25-hba2
    member a300_VDI-26-hba2
    member a300_VDI-27-hba2
    member a300_VDI-28-hba2
    member a300_VDI-29-hba2
    member a300_VDI-30-hba2
    member a300_VDI-31-hba2
    member a300_VDI-32-hba2



interface mgmt0
  ip address 10.29.164.239 255.255.255.0
vsan database
  vsan 4 interface fc1/13
  vsan 4 interface fc1/19
  vsan 4 interface fc1/20
  vsan 4 interface fc1/21
  vsan 4 interface fc1/22
  vsan 4 interface fc1/23
  vsan 4 interface fc1/24
  vsan 101 interface fc1/25
  vsan 101 interface fc1/26
  vsan 101 interface fc1/27
  vsan 101 interface fc1/28
  vsan 101 interface fc1/29
  vsan 101 interface fc1/30
  vsan 101 interface fc1/31
  vsan 101 interface fc1/32
  vsan 101 interface fc1/33
  vsan 101 interface fc1/34
  vsan 101 interface fc1/35
  vsan 101 interface fc1/36
  vsan 401 interface fc1/37
  vsan 401 interface fc1/38
  vsan 401 interface fc1/43
  vsan 401 interface fc1/44
  vsan 401 interface fc1/45
  vsan 401 interface fc1/46
switchname MDS-B
no terminal log-all
line console
  terminal width  80
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.8.1.1.bin
boot system bootflash:/m9100-s5ek9-mz.8.1.1.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46

interface fc1/1
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/2
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/3
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/4
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/5
  port-license acquire
  no shutdown

interface fc1/6
  port-license acquire
  no shutdown

interface fc1/7
  port-license acquire
  no shutdown

interface fc1/8
  port-license acquire
  no shutdown

interface fc1/9
  port-license acquire
  no shutdown

interface fc1/10
  port-license acquire
  no shutdown

interface fc1/11
  port-license acquire

interface fc1/12
  port-license acquire

interface fc1/13
  port-license acquire
  no shutdown

interface fc1/14
  port-license acquire
  no shutdown

interface fc1/15
  port-license acquire
  no shutdown

interface fc1/16
  port-license acquire
  no shutdown

interface fc1/17
  port-license acquire
  no shutdown

interface fc1/18
  port-license acquire
  no shutdown
```

```
interface fc1/19
  switchport trunk allowed vsan 4
  switchport description CS700 CTRL-A:01
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/20
  switchport trunk allowed vsan 4
  switchport description CS700 CTRL-A:05
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/21
  switchport trunk allowed vsan 4
  switchport description Launcher-FIB
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/22
  switchport trunk allowed vsan 4
  switchport description Launcher-FIB
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/23
  switchport trunk allowed vsan 4
  switchport description Launcher-FIB
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/24
  switchport trunk allowed vsan 4
  switchport description Launcher-FIB
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/25
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/26
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/27
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/28
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/29
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/30
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/31
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/32
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/33
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/34
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/35
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/36
  switchport trunk allowed vsan 101
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/37
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/38
  switchport trunk mode off
  port-license acquire
  no shutdown

interface fc1/39
  port-license acquire
  no shutdown
```

```
interface fc1/40
  port-license acquire
  no shutdown

interface fc1/41
  port-license acquire
  no shutdown

interface fc1/42
  port-license acquire
  no shutdown

interface fc1/43
  port-license acquire
  no shutdown

interface fc1/44
  port-license acquire
  no shutdown

interface fc1/45
  port-license acquire
  no shutdown

interface fc1/46
  port-license acquire
  no shutdown

interface fc1/47
  port-license acquire
  no shutdown
```

```
interface fc1/48
  port-license acquire
  no shutdown
ip default-gateway 10.29.164.1
```

Appendix B NetApp AFF300 Monitoring with PowerShell Scripts

NetApp AFF A300 Monitoring with Powershell Scripts

NetApp offers a wide range of methods for connecting and operating its storage controllers. One of such method is the Data ONTAP PowerShell Toolkit, which is available free of charge and can be downloaded from the NetApp Communities.

This reference architecture uses the PowerShell Toolkit for collecting performance information. Specifically, we use Invoke-NcSysstat, a cmdlet designed to report live performance data for the cluster. Like Invoke-NaSysstat, Invoke-NcSysstat monitors several performance counters to report on components of the system. The following performance statistics can be retrieved: System, FCP, NFSv3, NFSv4, CIFS, iSCSI, Volume, Ifnet, LUN, and Disk.

To use Invoke-NcSysstat, provide a switch indicating which type of performance statistics to retrieve. Use the Name parameter to monitor only specific performance objects. For example, to monitor volumes named `powershell` in the cluster, run the following command:

```
PS C:\Toolkit\3.0> Invoke-NcSysstat -Volume -Name powershell -Count 5

Name                                 RdOps  WrOps TotOps   RdLat   WrLat  TotLat    Read Written
----                                 -----  ----- ------   -----   -----  ------    ---- -------
powershell                               0    430    434     0.0     0.5     0.6    1 KB   27 MB
powershell                               0    515    519    13.3     0.7     0.8    1 KB   32 MB
powershell                               0    622    631     0.1     0.8     0.9    2 KB   39 MB
powershell                               1    600    604     4.6     1.0     1.0    2 KB   37 MB
powershell                               0    590    599     0.1     0.9     1.0    2 KB   37 MB
```

The Name parameter also accepts wildcards.

```
PS C:\Toolkit\3.0> Invoke-NcSysstat -Lun –Name /vol/powershell/*

    Read  Written  RdOps   WrOps TotOps  TotLat LunPath
    ----  -------  -----   ----- ------  ------ -------
       0        0      0       0      0     0.0 /vol/powershell/disk1_gpt
  135 B         0      0       0      1     2.0 /vol/powershell/luns/cluster
       0   128 B       0     128    128   217.0 /vol/powershell/luns/disk0
```

Invoke-NcSysstat works in both the cluster and SVM context for Data ONTAP 8.2 and later. For Data ONTAP versions earlier than 8.2, Invoke-NcSysstat must be run in the cluster context. The following performance statistics can be retrieved in the SVM context: FCP, NFSv3, NFSv4, CIFS, iSCSI, Volume, and LUN.

When run in the cluster context, select monitored performance objects can be filtered by node or SVM. Invoke-NcSysstat then only monitors the performance objects associated with the given node or SVM. For example, to monitor all of the volumes on a specific SVM, run the following command:

```
PS C:\Toolkit\3.0> Invoke-NcSysstat -Volume -Vserver beam01

Name                                 RdOps  WrOps TotOps   RdLat   WrLat  TotLat    Read Written
----                                 -----  ----- ------   -----   -----  ------    ---- -------
beam01_root_vol                          0      0      0     0.0     0.0     0.0       0       0
clusterdisks                             0      0      0     0.0     0.0     0.0       0       0
davidCModeIscsiMoun1                     0      0      0     0.0     0.0     0.0       0       0
ndmp_destination                         0      0      0     0.0     0.0     0.0       0       0
powershell                               0    370    370     0.0     0.2     0.2       0   23 MB
testvol                                  0      0      0     0.0     0.0     0.0       0       0
v1NfsSrCMode                             0      0      0     0.0     0.0     0.0       0       0
vmstorage                                0      0      0     0.0     0.0     0.0       0       0
```

The following performance statistics can be filtered by SVM: FCP, Volume, CIFS, and iSCSI. The following performance statistics can be filtered by node: FCP, Ifnet, Disk, Volume, CIFS, iSCSI, and System.

Additionally, Invoke-NcSysstat can aggregate the performance statistics for select objects by SVM or node. Instead of displaying the performance results for each individual object, the performance statistics are aggregated for all the objects on the given SVM or node. The following example aggregates the volume performance for all of the volumes in SVM `beam01`:

```
PS C:\Toolkit\3.0> Invoke-NcSysstat -Volume -Vserver beam01 -Aggregated -Count 5

Name                                 RdOps  WrOps TotOps   RdLat   WrLat  TotLat    Read Written
----                                 -----  ----- ------   -----   -----  ------    ---- -------
beam01                                   0    257    266     0.0     0.2     0.2       0   16 MB
beam01                                   0    606    614     0.0     0.2     0.3       0   38 MB
beam01                                   0    357    363     0.0     0.3     0.3       0   22 MB
beam01                                   1     22     24     0.0     0.2     2.6   341 B    1 MB
beam01                                   0      1      8     0.0     0.1     0.1       0    2 KB
```

The following example aggregates the CIFS performance for all of the CIFS servers on the node `MFIT-01`:

```
PS C:\Toolkit\3.0> Invoke-NcSysstat -Cifs -Node MFIT-01 -Aggregated -Count 5

Name                                 RdOps  WrOps TotOps   RdLat   WrLat  TotLat
----                                 -----  ----- ------   -----   -----  ------
MFIT-01                                  0    152    155     0.0     0.3     0.3
MFIT-01                                  0    153    156     0.0     0.3     0.3
MFIT-01                                  0    121    123     0.0     0.3     0.3
MFIT-01                                  0    154    158     0.0     0.3     0.3
MFIT-01                                  0    155    157     0.0     0.3     0.3
```

The CIFS, iSCSI, FCP, and Volume performance statistics can be aggregated by node or by SVM.

Provisioning Persistent Desktops Powershell Script that Utilizes Storage Copy Offload with VAAI

The following is a script that can be used to provision persistent desktops verses using Citrix Machine Creation Services (MCS). MCS uses VMware Hypervisor snapshots and does not use VAAI integration with the storage array. The following Powershell script will utilize VAAI and reducing a 5-minute copy time down to seconds. If you are creating many persistent desktops at one time, this script will be useful.

```
# Name: Clone Desktops Powershell script
# Date: 1/02/2018
# Description: Clones persistent desktops from a template and utilizes VAAI
#
# Author: Dave Arnette, NetApp, Inc.
#
# Revisions:
#
# Starting script
[CmdletBinding()]
Param(

    [Parameter()]
    [switch]$deploy = $false,

    [Parameter()]
    [switch]$changeNetwork = $false,

    [Parameter()]
    [switch]$test = $false

    )

$TargetCluster = Get-Cluster -Name "TMEonly"
$targetFolder = "Launchers"
$vmHost = "192.168.201.101"
$SourceVMTemplate = Get-Template -Name "W2016_template_0510"
$SourceCustomSpec = Get-OSCustomizationSpec -Name "desktop-v1"
```

```
$datastore = get-datastore -name "TMEinfraONLY"
$domaincred = (New-Object System.Management.Automation.PSCredential "vdi\administrator",(get-content
c:\vdi_domaincred.txt | convertto-securestring) )


$servers = (1,3,5,7,8)
$basename = "launcher"


$vmList = @()

$servers | foreach-object {

    $vm = New-Object System.Object

    if ( $_ -le 9) {
        $name = "${basename}-00${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.10${_}"

        }
    elseif ( $_ -ge 10 -and $_ -le 99 ) {
        $name = "${basename}-0${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.1${_}"

        }
    else {
        $name = "${basename}-${_}"
        $vm | add-member -MemberType NoteProperty -name "Name" -value "${name}"
        $vm | add-member -MemberType NoteProperty -name "IP" -value "172.18.0.${_}"

        }
        $vmList = $vmList + $vm
}



function change_dns {
    Write-host "Waiting for VMware Tools to start on $vmname"
    wait-tools -vm $vmname -TimeoutSeconds 300

    write-host " Changing DNS server on $vmname"

    $dnsChangeCmd = { get-DNSClientServerAddress -interfaceAlias "Ethernet0" | set-DNSClientServerAddress
-serveraddresses "192.168.201.41" }
    invoke-command -computername $vmname -credential $domaincred -scriptblock $dnsChangeCmd

    }

function change_ip {

    write-host " Changing IP address on $vmname"

    $ipChangeCmd = { new-netIPaddress -interfaceAlias "Ethernet0" -ipaddress $Using:ip -PrefixLength "16"
-DefaultGateway "172.18.255.254" }

    invoke-command -computername $vmname -credential $domaincred -scriptblock $ipChangeCmd -AsJob

    }

if ($deploy) {
    if ($test) {
        $vmList| foreach-object {
            $vmname = $_.name

            write-host " Creating VM $vmname from template $sourceVMtemplate " -ForegroundColor Green
            }
        }
    else {

        #create the VMs asynchronously
        $createJobList = @()
        $vmList| foreach-object {
            #$vmname = ""
            $vmname = $_.Name

            write-host " Creating VM $vmname from template $sourceVMtemplate " -ForegroundColor Green
            new-vm -name $vmname -resourcepool $targetCluster -vmhost $vmhost -location $targetFolder -
template $sourceVMtemplate -OSCustomizationSpec $SourceCustomSpec -datastore $datastore -runasync -
OutVariable createJob

            $createJob | add-member -MemberType NoteProperty -name "VM" -value "$vmname"

            $createJobList = $createJobList + $createJob


        }


    # Wait for each clone job to finish, then start the VM


        $createJobList| foreach-object {
            $vmname = $_.VM

            write-host " Waiting for VM $vmname clone process to finish " -ForegroundColor Green

            do { start-sleep -seconds 5 }
            until ( (get-task -id $_.ID).state -eq "Success" )

            write-host " Starting VM $vmname " -ForegroundColor Green
            start-vm -vm $vmname -runAsync


        }

    }
}
if ($changeNetwork) {

    if ($test) {
        $vmList| foreach-object {
            $vmname = $_.Name
            $ip = $_.IP

            write-host " Changing DNS on $vmname" -ForegroundColor Green
            Write-host " Changing IP on $vmname to $ip"
            }
        }
    else {

        $vmList | foreach-object {
            $vmname = $_.name
            $ip = $_.IP

            change_dns
            change_ip

            }
        }
}
```

Unfortunately, VMware’s ESXtop tool will only show you block protocol copy offload bytes and it does not capture NFS copy offload byte metrics. Therefore, too validate that the cloning of the desktop (provisioning) is actually being offloaded to the storage with the above Powershell script, you can use the following NetApp storage commands:

```
set -priv diag"
```

```
statistics start -object copy_manager

stats stop

statistics show -object copy_manager -counter sce_kb_copied
```

If you want to see all of the VAAI stats, type:

```
statistics show -object copy_manager
```

Creating User Home Directory Folders with a Powershell Script

There are many tools to create user home directory folders. For example, you can use a Microsoft Powershell script to automatically create the folder at first login. You can also use a Powershell script. For this reference architecture, we used the Powershell script method to create the user home directory folders. The following console text shows the power shell script commands used to create the home directories.

```
# Name: CreateHomeDirFolders Powershell script
# Date: 10/02/2014
# Description: Creates the initial home directory folders with names supplied in CreateHomeDirFolders.txt
text file.
#              If Folder exists, it will leave the folder alone and move to the next home directory name
in the text file.
#
# Author: C-Rod, NetApp, Inc.
#
# Revisions:
#
# Starting script
Set-ExecutionPolicy -ExecutionPolicy Unrestricted -Scope Process -Force

clear
Write-Host "Create Home Directory Folders Powershell Script"
        Write-Host ""

$ExecutionPath = split-path $SCRIPT:MyInvocation.MyCommand.Path -parent
$HomeDirTextFile = $ExecutionPath + "`\CreateHomeDirFolders.txt"
$HomeDirFolderNames = (Get-Content $HomeDirTextFile)
$inputanswer = "n"

function create_folders {
    do {
        Write-Host ""
        Write-Host "Ensure the Top Level Home directory folder is mapped in Windows to a drive letter
`(e.g. \\servername\Common\Home H:`)"
        $HD_Location = (Read-Host "Enter Drive letter and folder to create the home directories `(e.g.
H:\Common\Home`)").Trim()
        Write-Host ""
        Write-Host ""
        If (!(Test-Path $HD_Location)) {
            Write-Host "Home Directory Folder Share $HD_Location does not exist"
            Write-Host ""
            $inputanswer = "n"
        } else {
            Write-Host "Summary of Input Parameters:"
            Write-Host "HomeDir Folder Name: $HD_Location"
            Write-Host ""
            $inputanswer = (Read-Host "Is this value correct? [y] or [n]").Trim()
        }
    } while ($inputanswer -ne "y")

    foreach($HD_Name in $HomeDirFolderNames)
    {
        $NewFolderName = $HD_Location + "`\" + $HD_Name
        If (!(Test-Path $NewFolderName)) {
            md $NewFolderName
        }
    }
}

create_folders
```

Appendix C Full Scale Mixed Workload Test Results

Figure 142 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

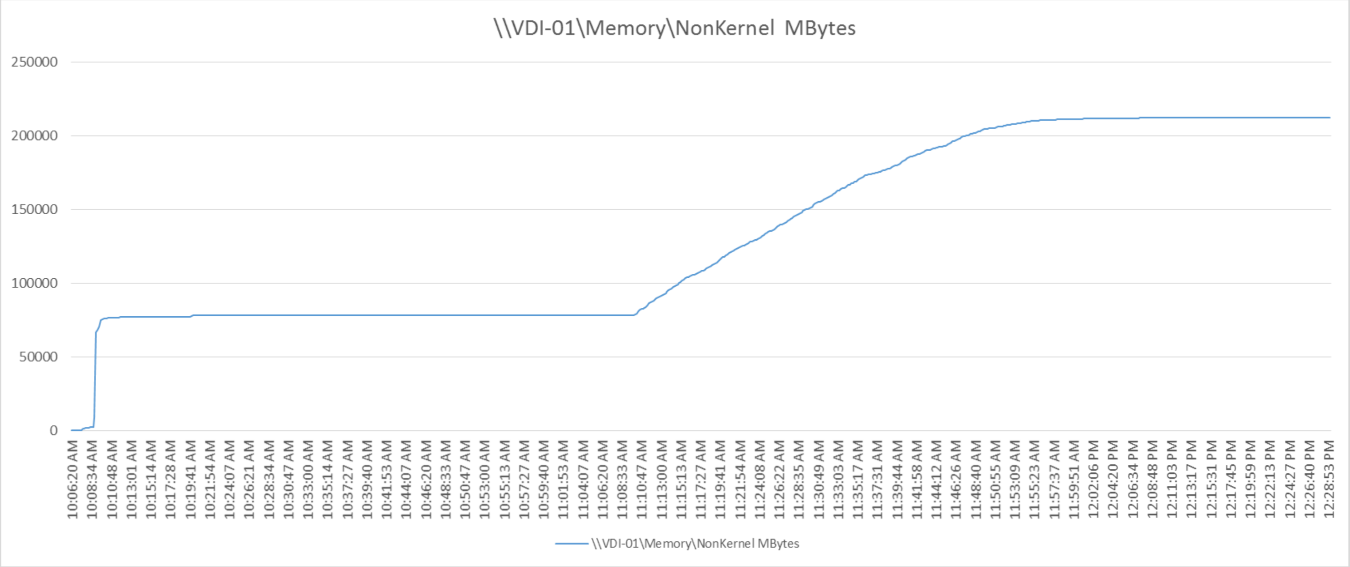

Figure 143 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

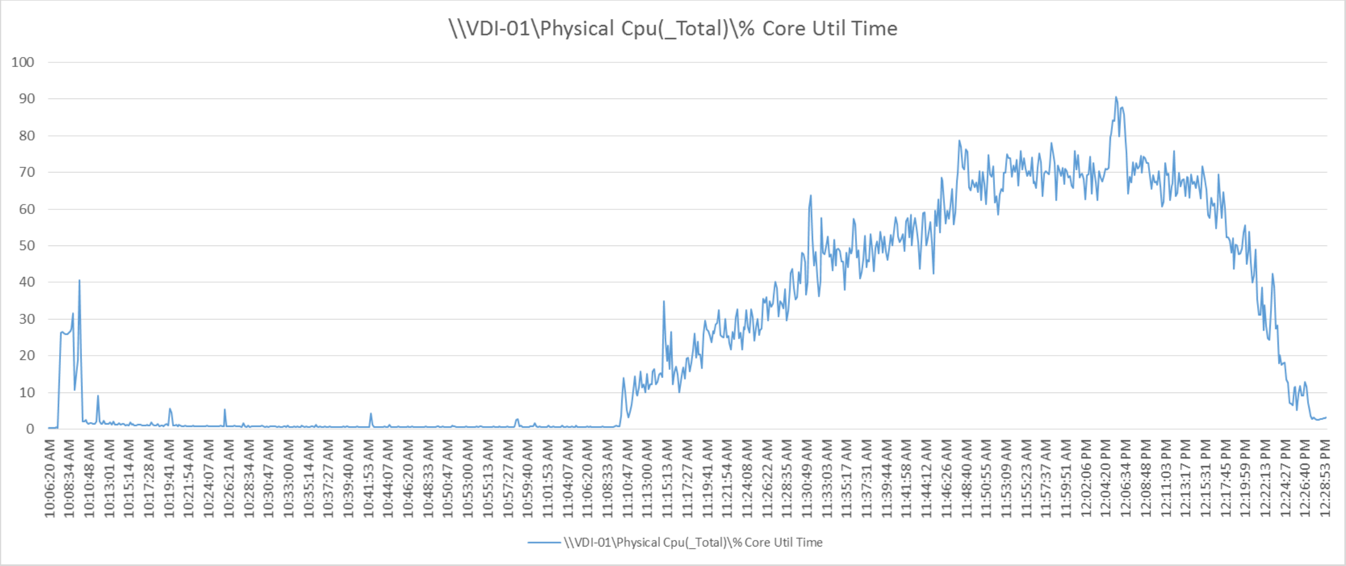

Figure 144 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

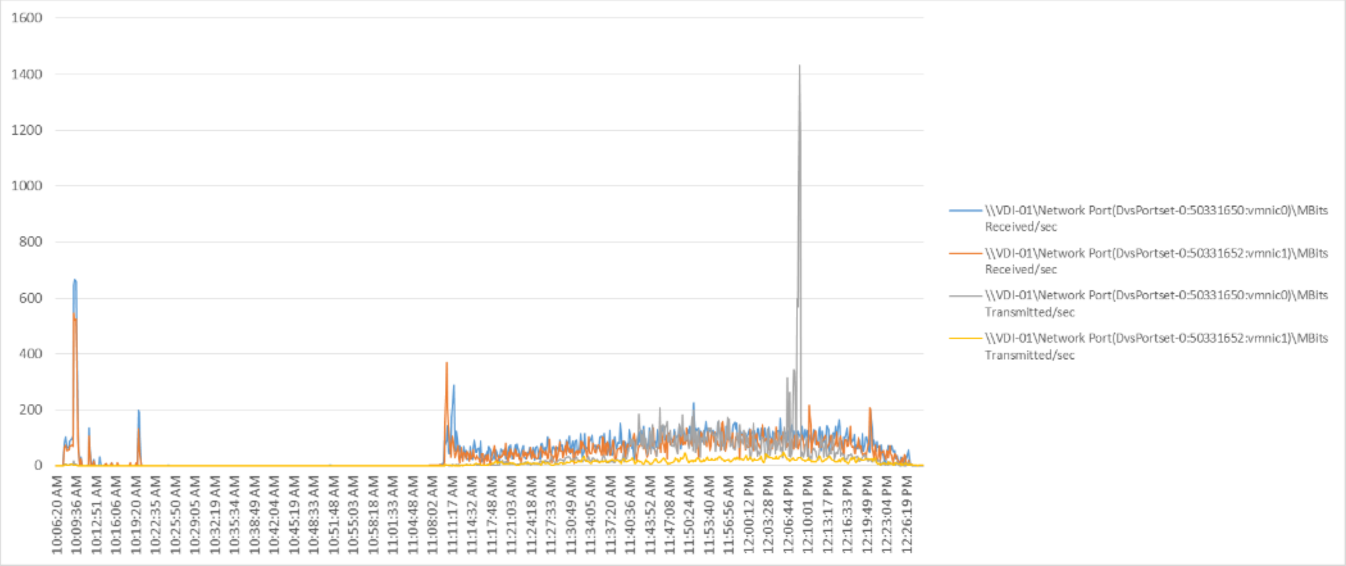

Figure 145 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

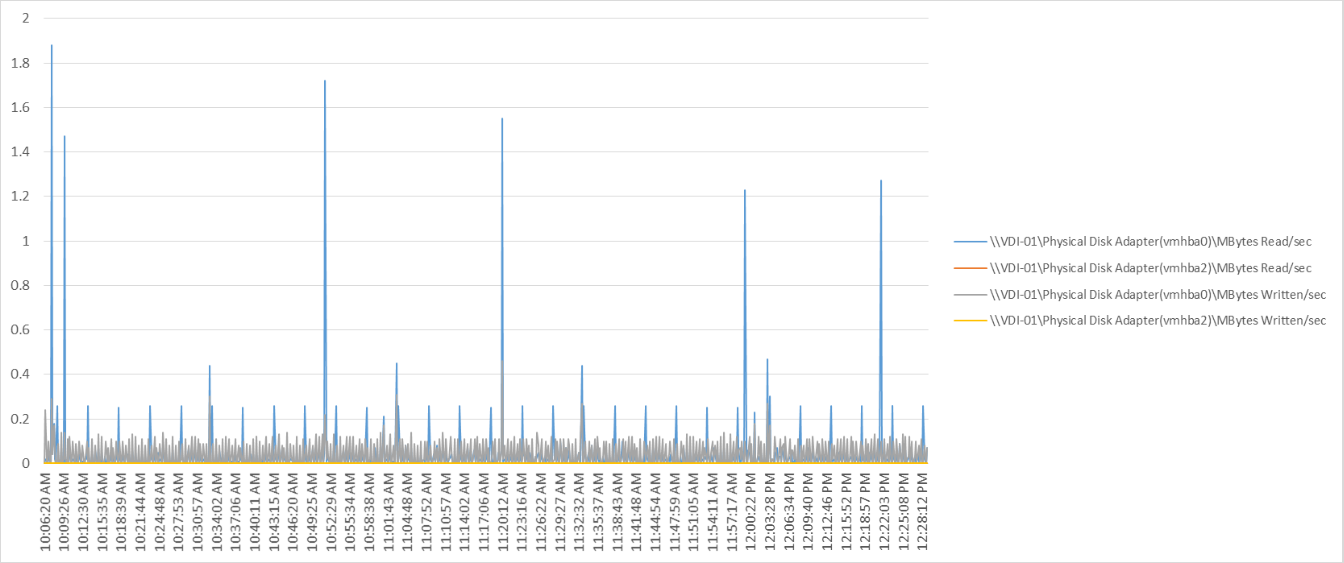

Figure 146 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

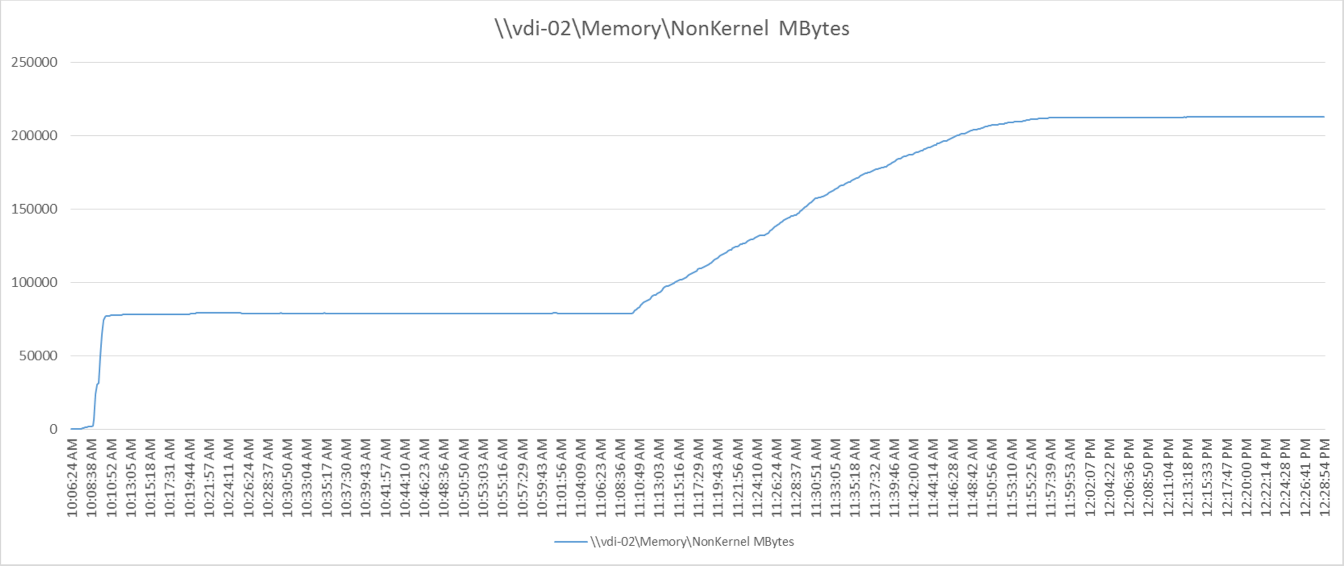

Figure 147 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

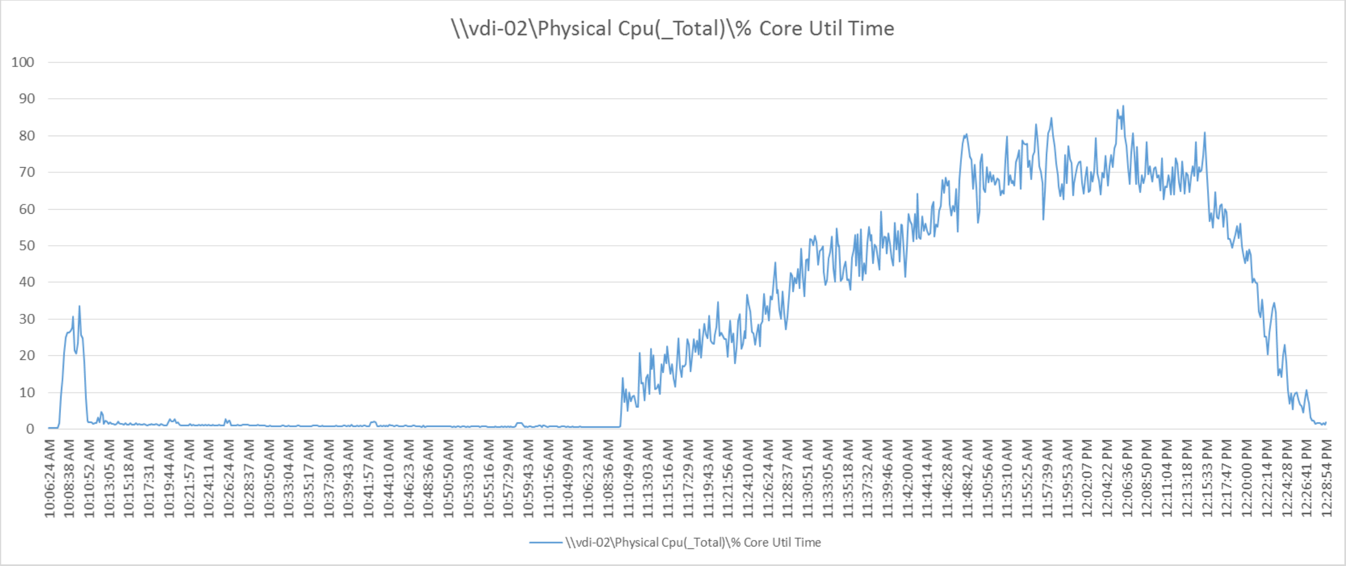

Figure 148 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

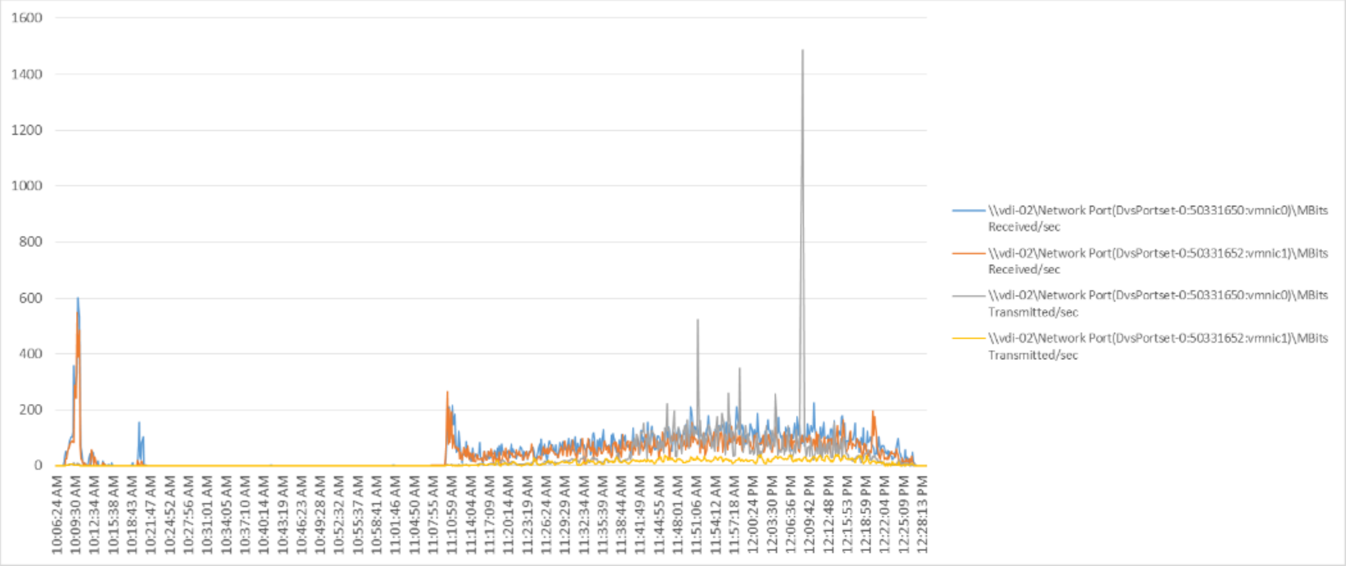

Figure 149 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

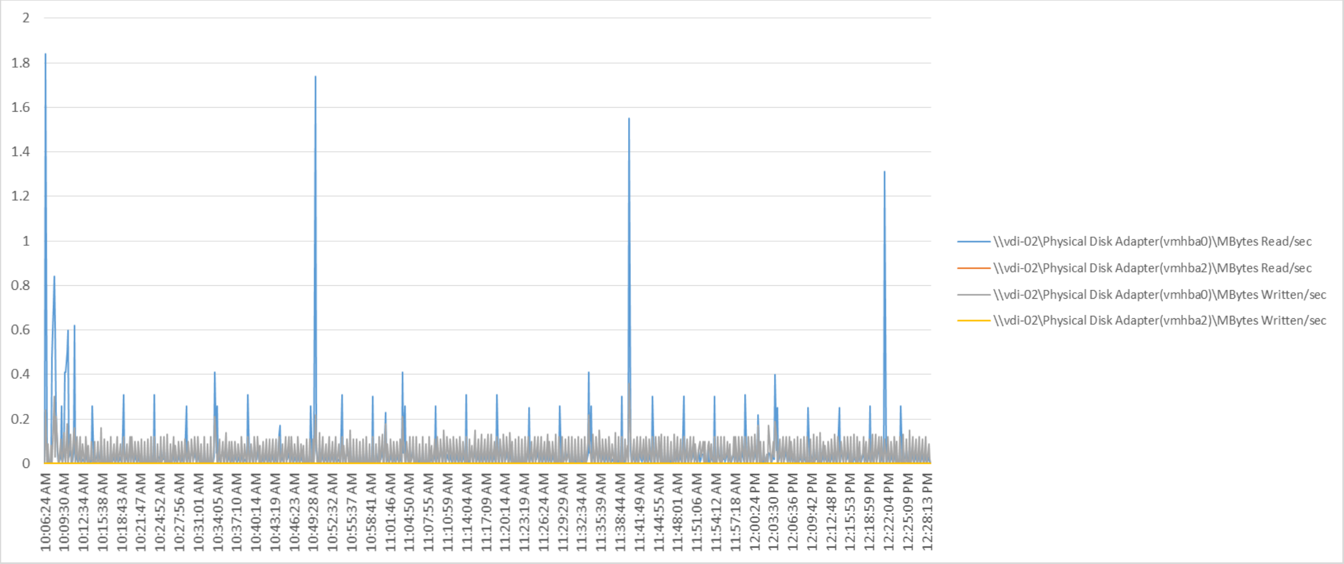

Figure 150 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

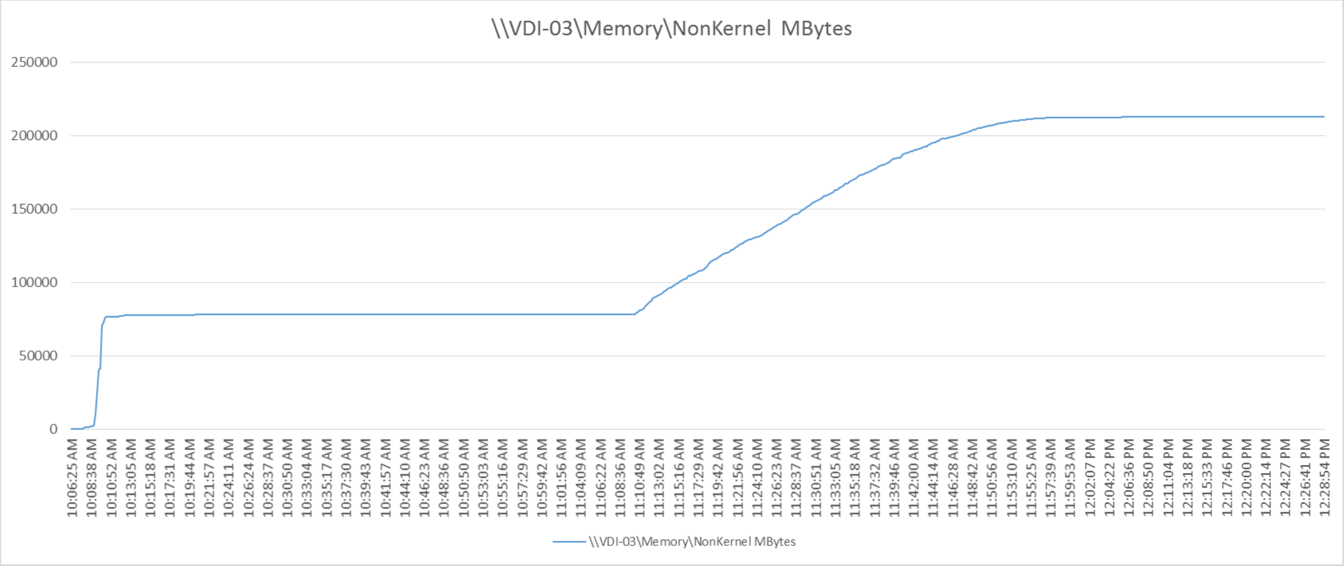

Figure 151 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

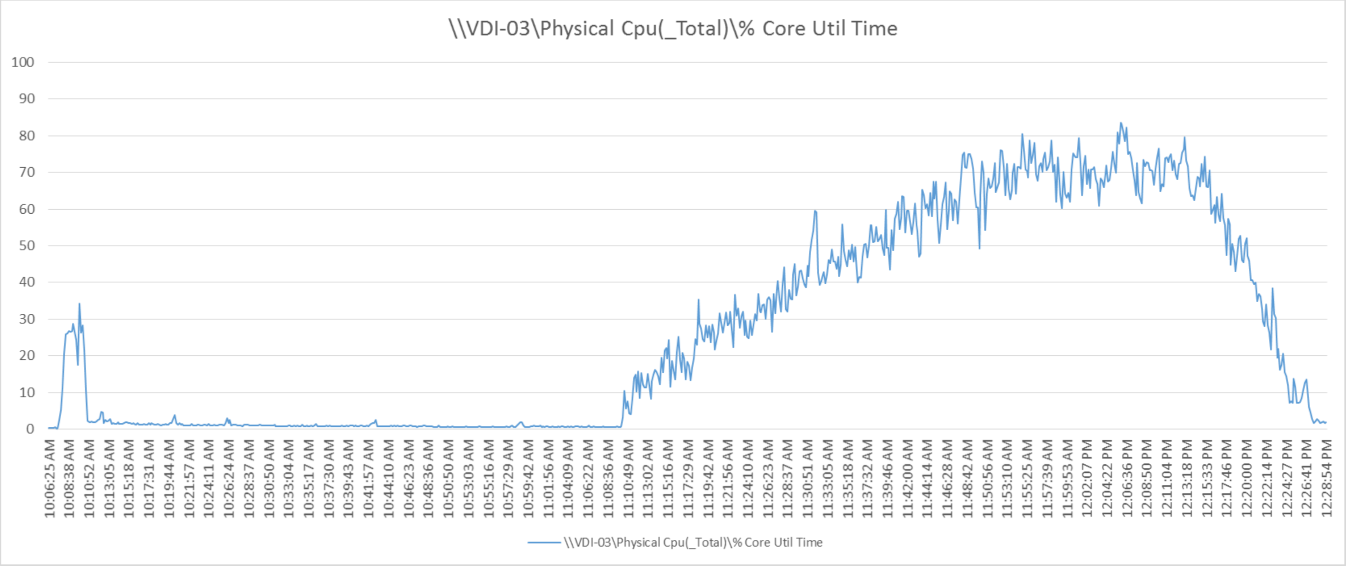

Figure 152 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

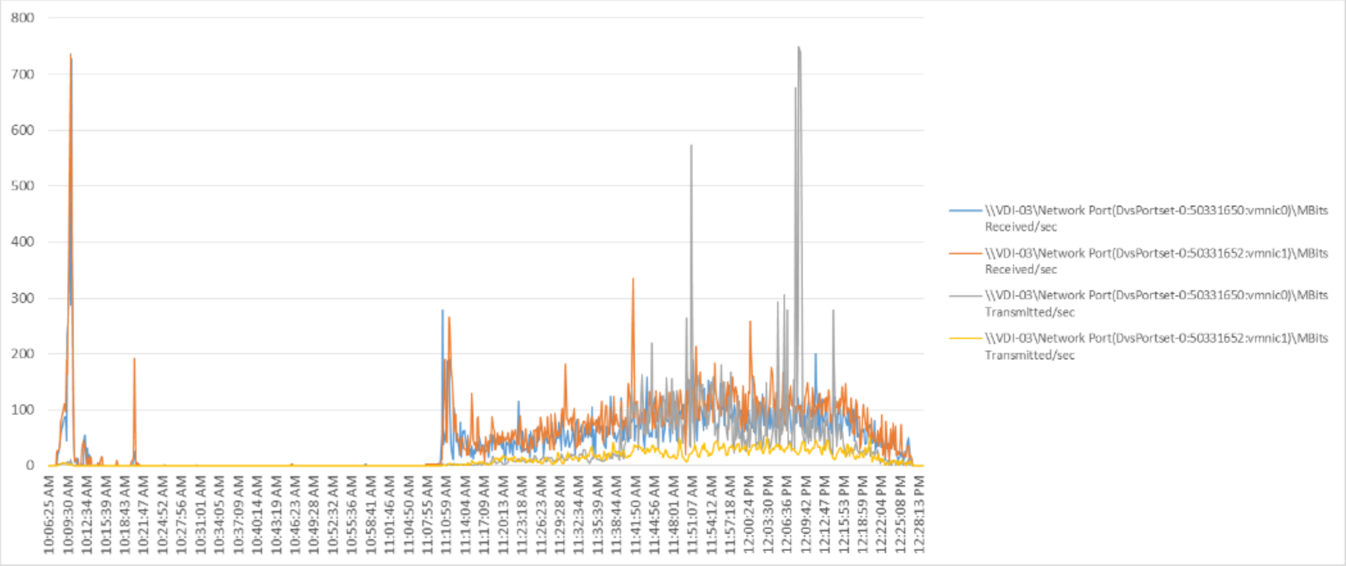

Figure 153 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

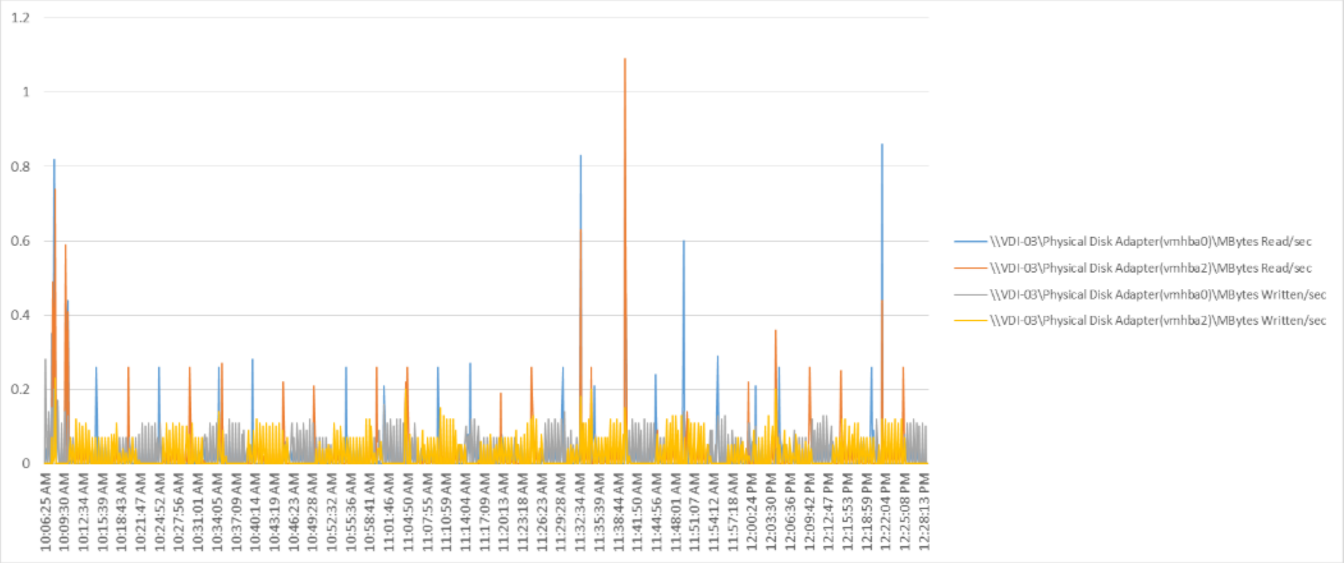

Figure 154 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

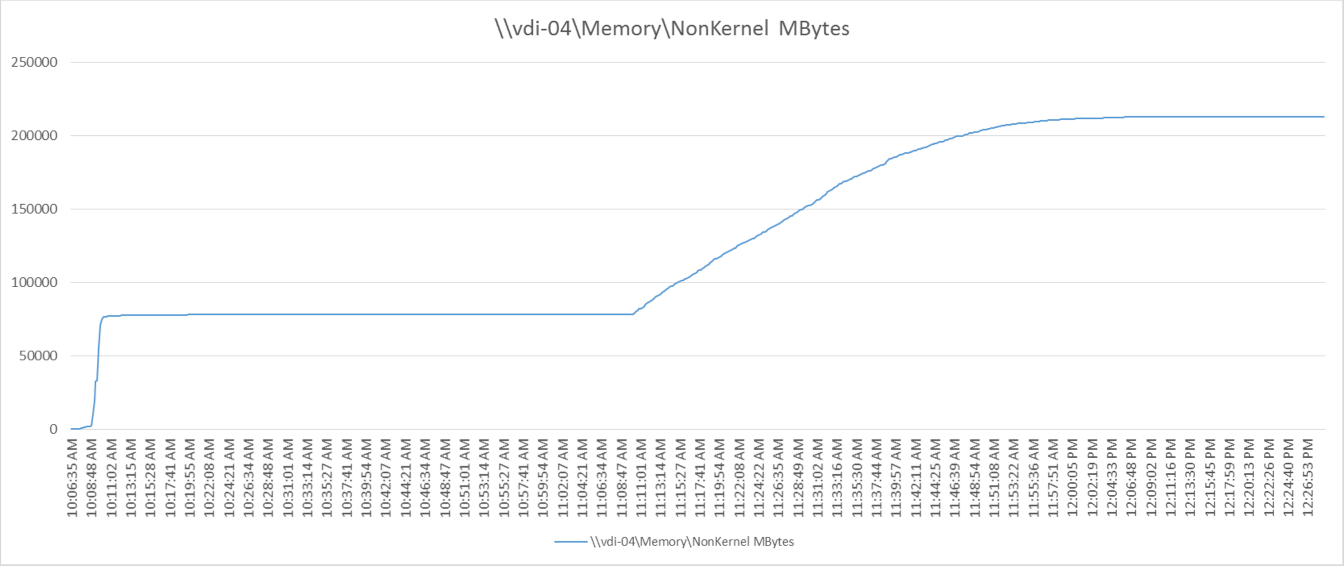

Figure 155 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

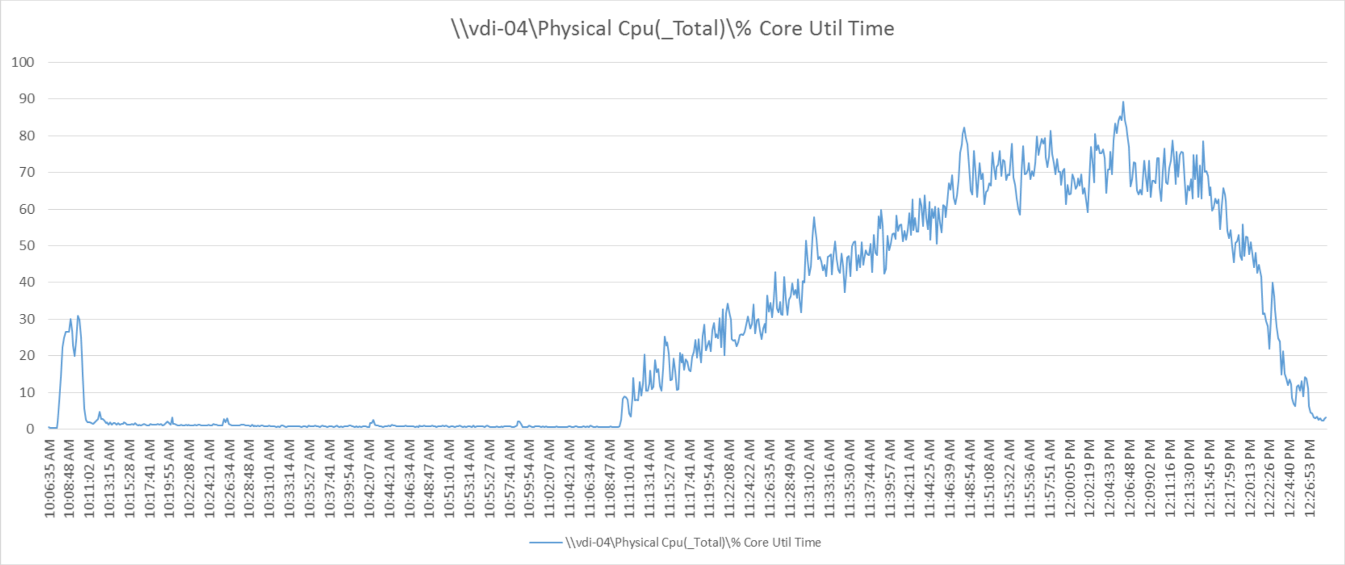

Figure 156 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

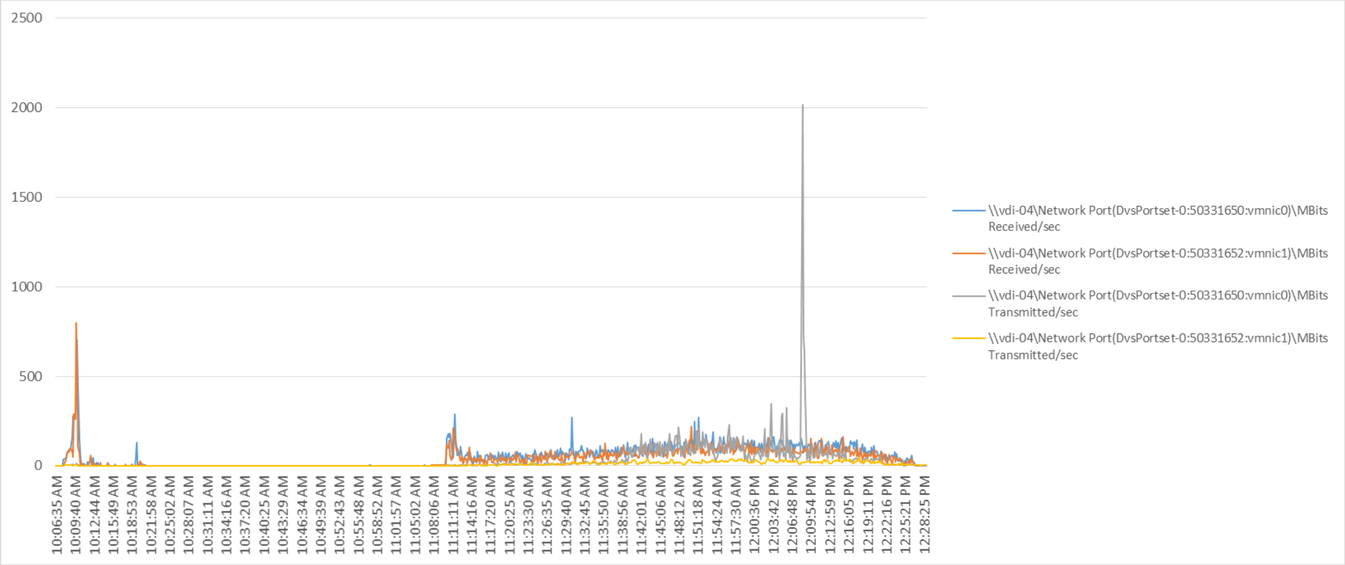

Figure 157 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

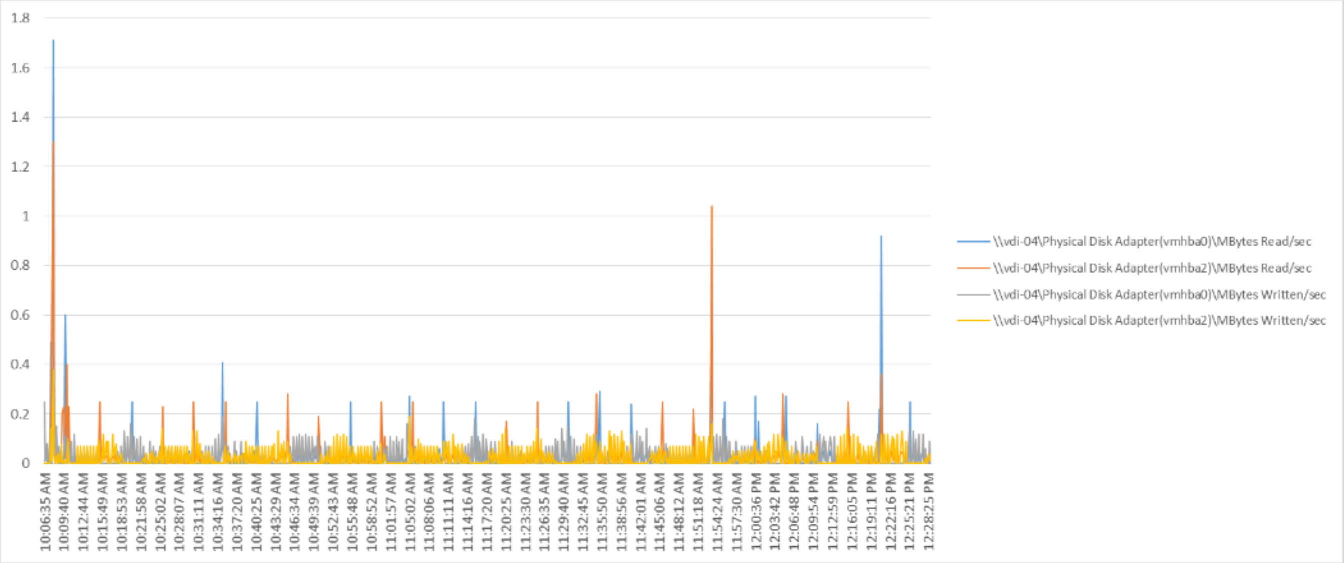

Figure 158 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

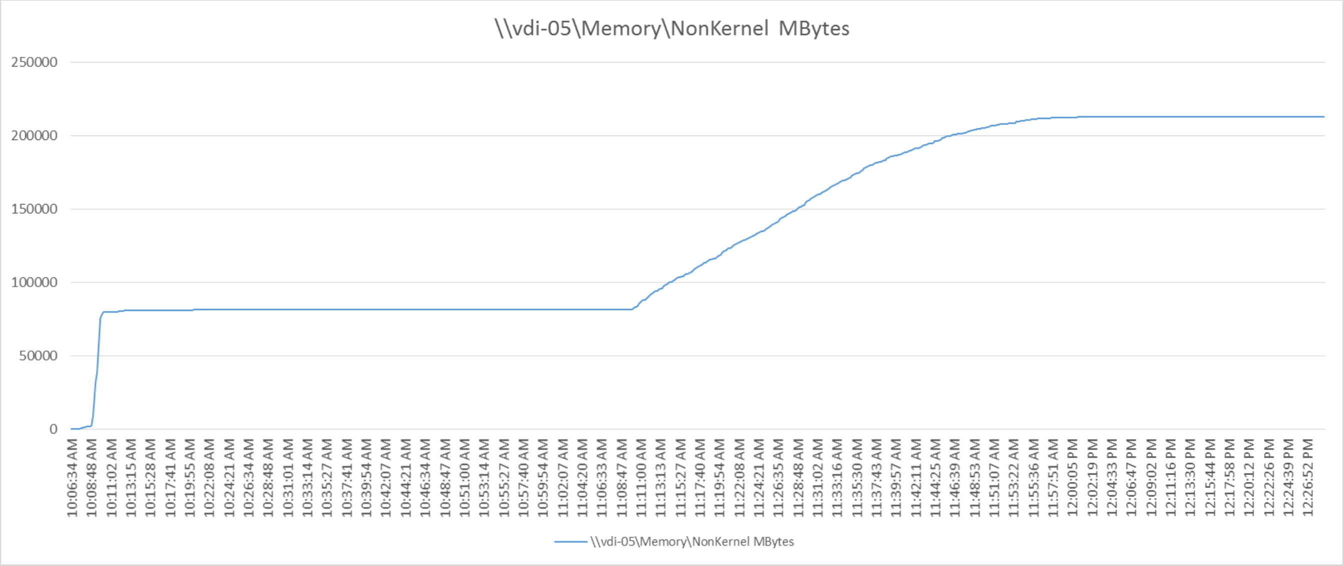

Figure 159 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

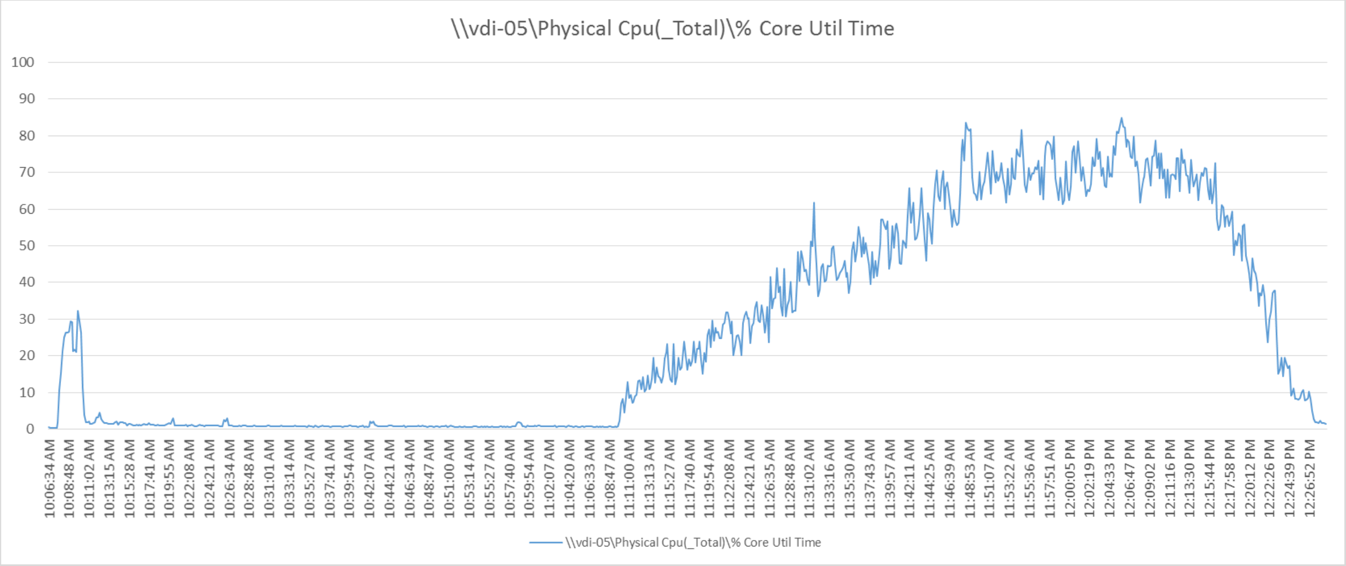

Figure 160 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

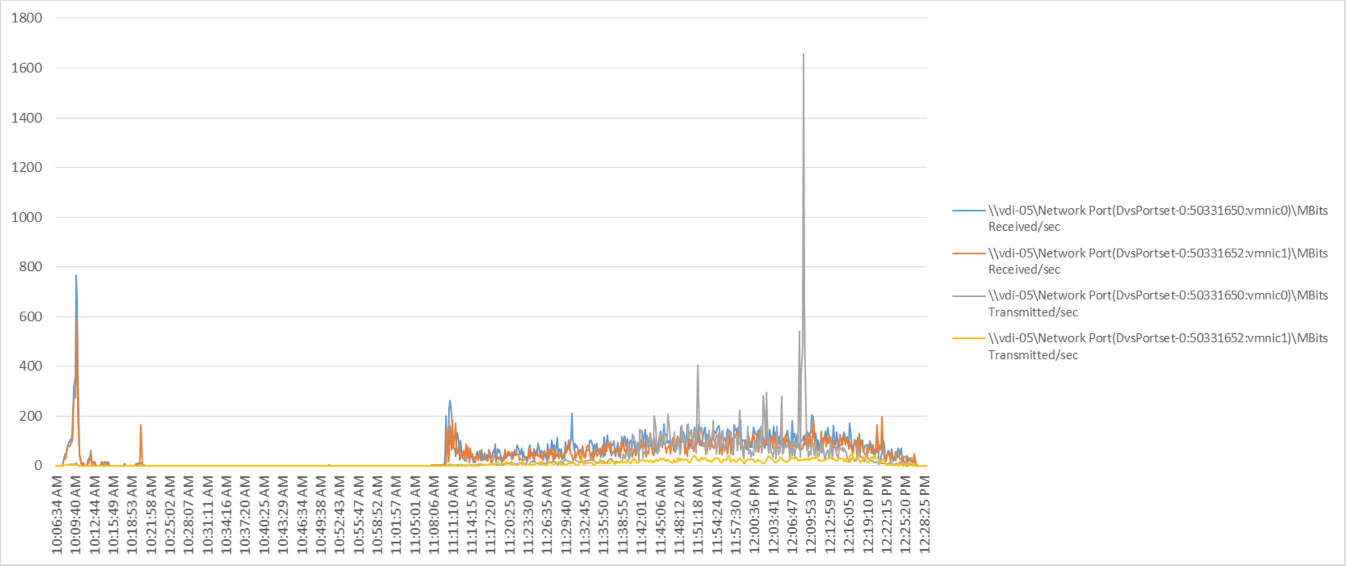

Figure 161 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

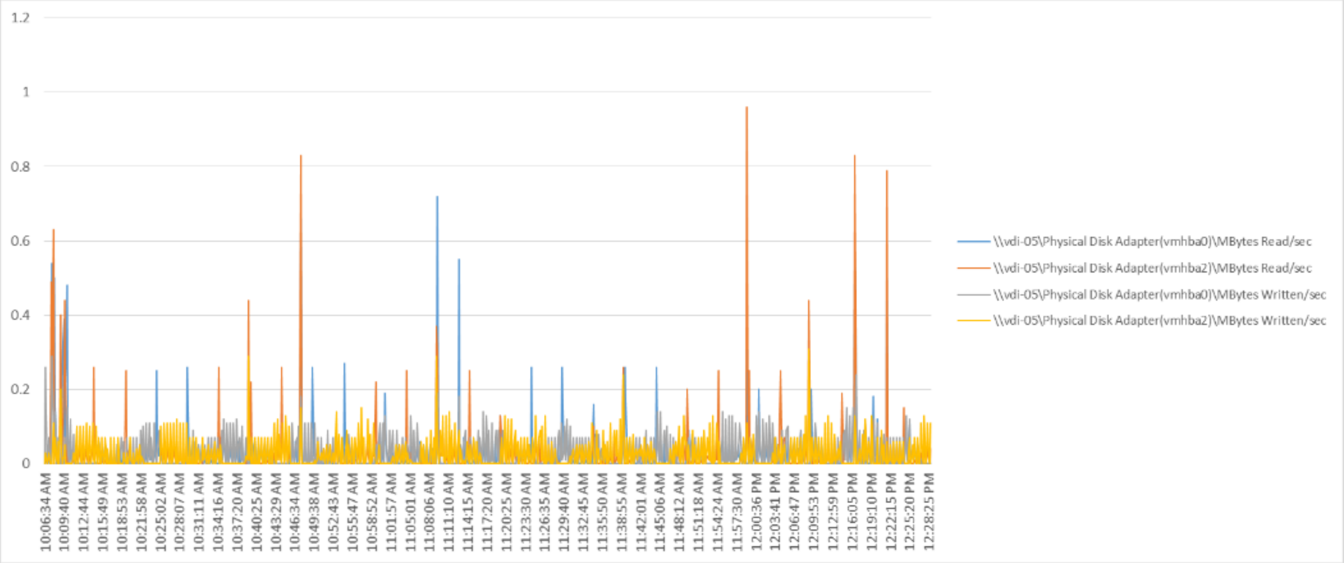

Figure 162 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

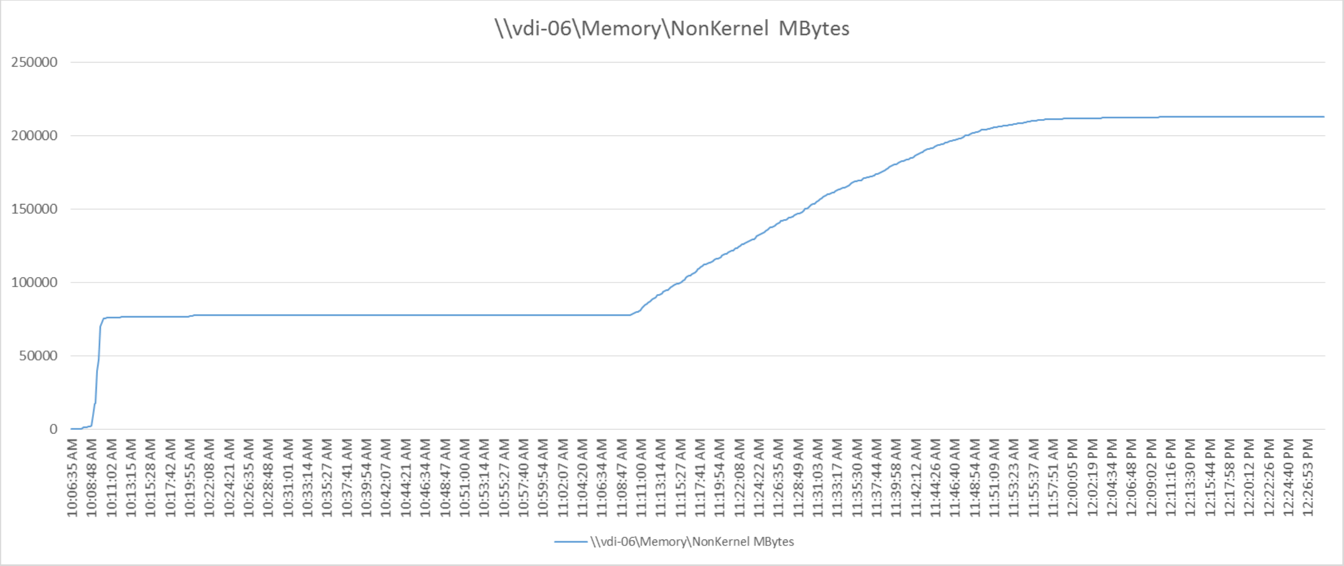

Figure 163 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

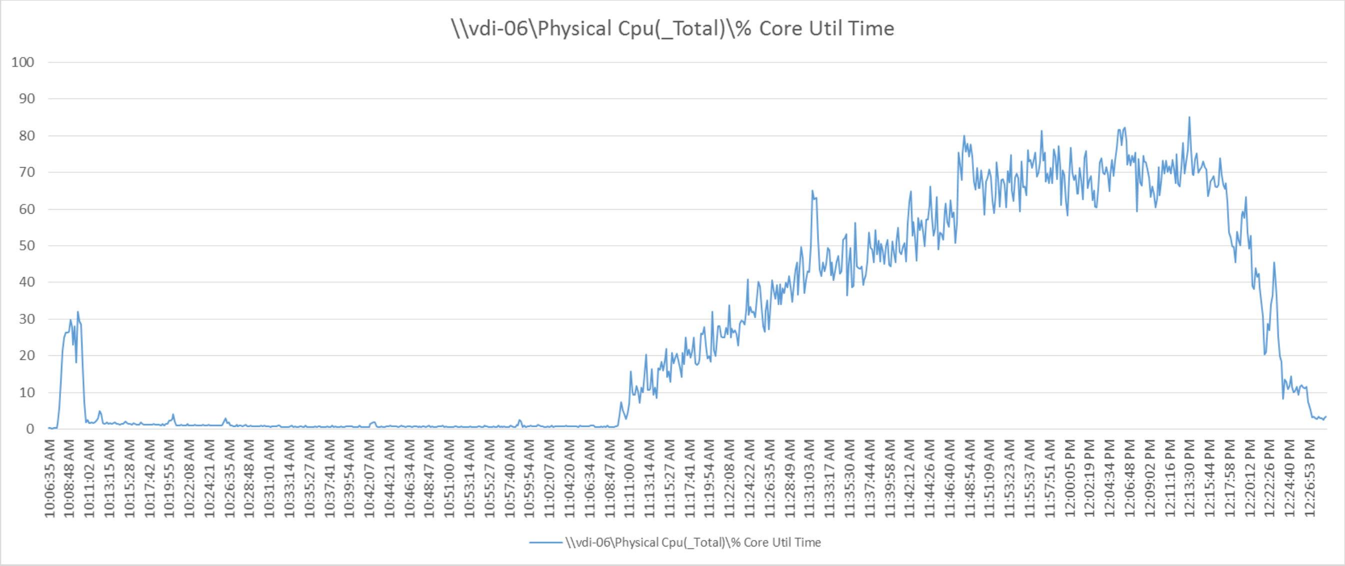

Figure 164 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

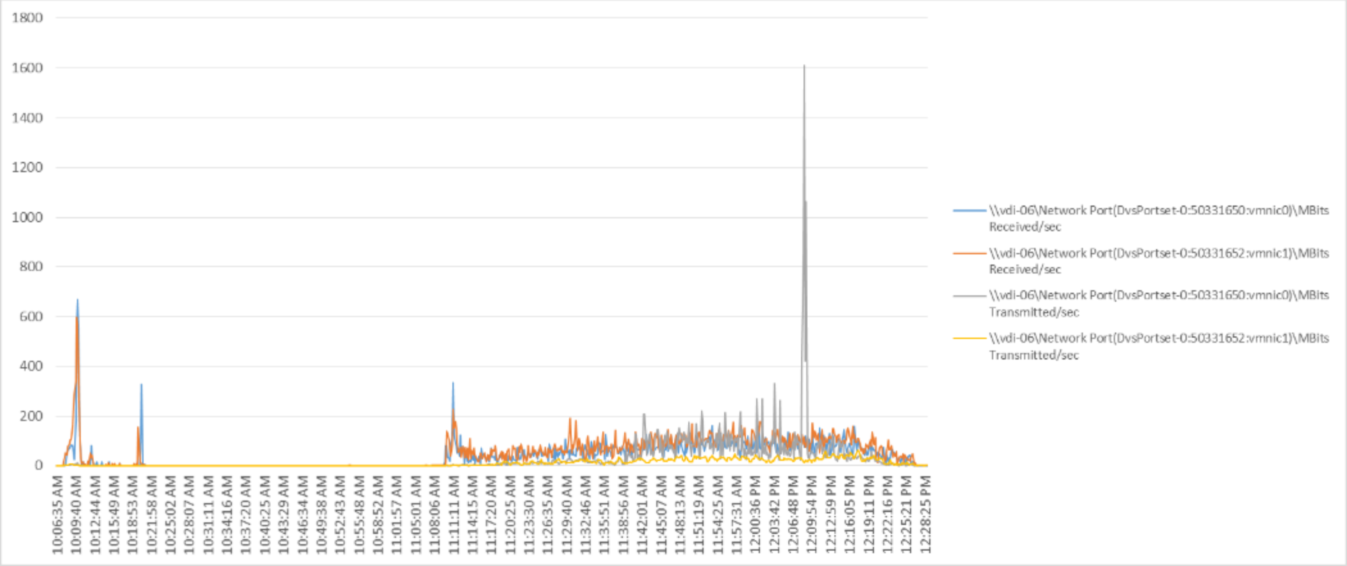

Figure 165 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

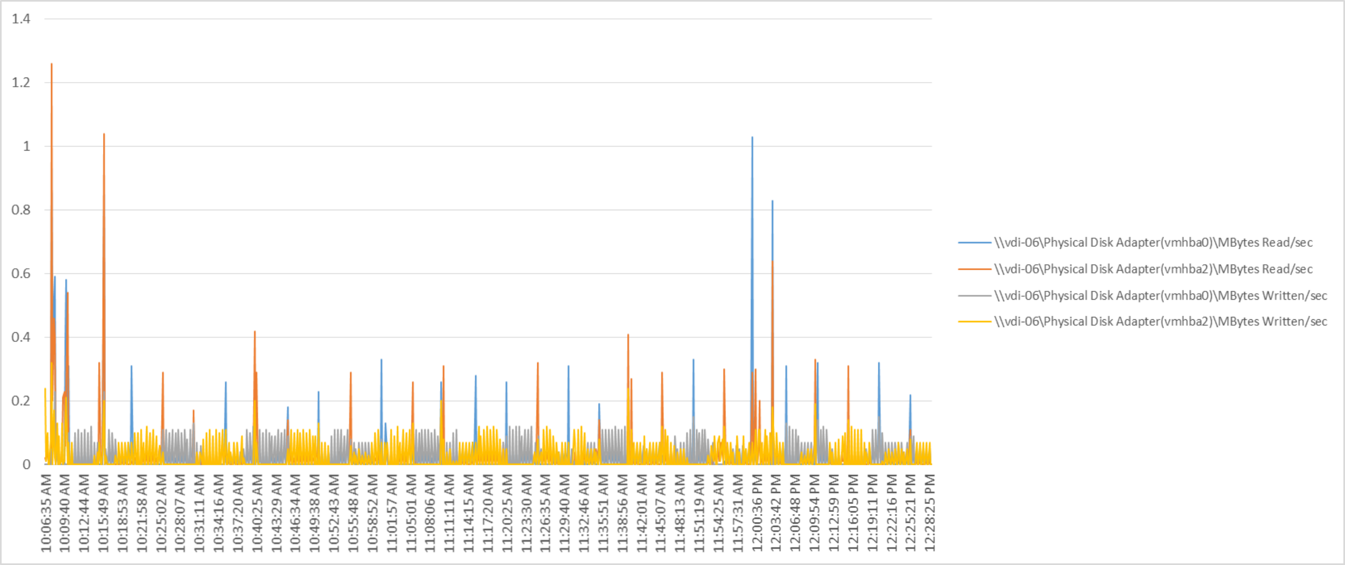

Figure 166 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

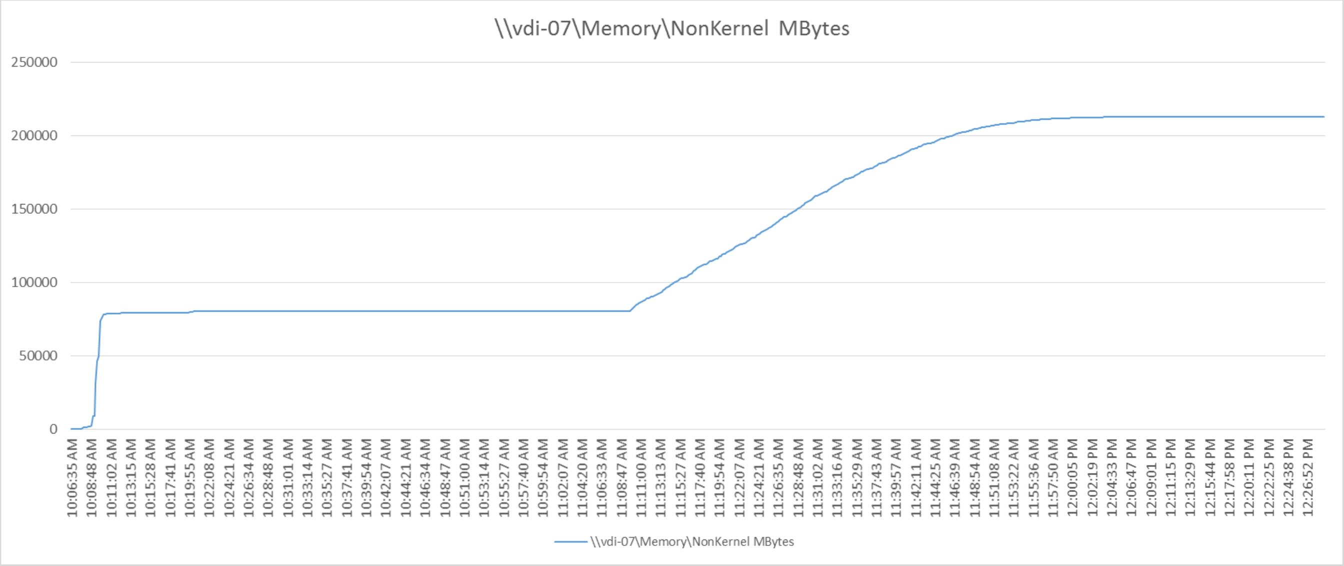

Figure 167 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

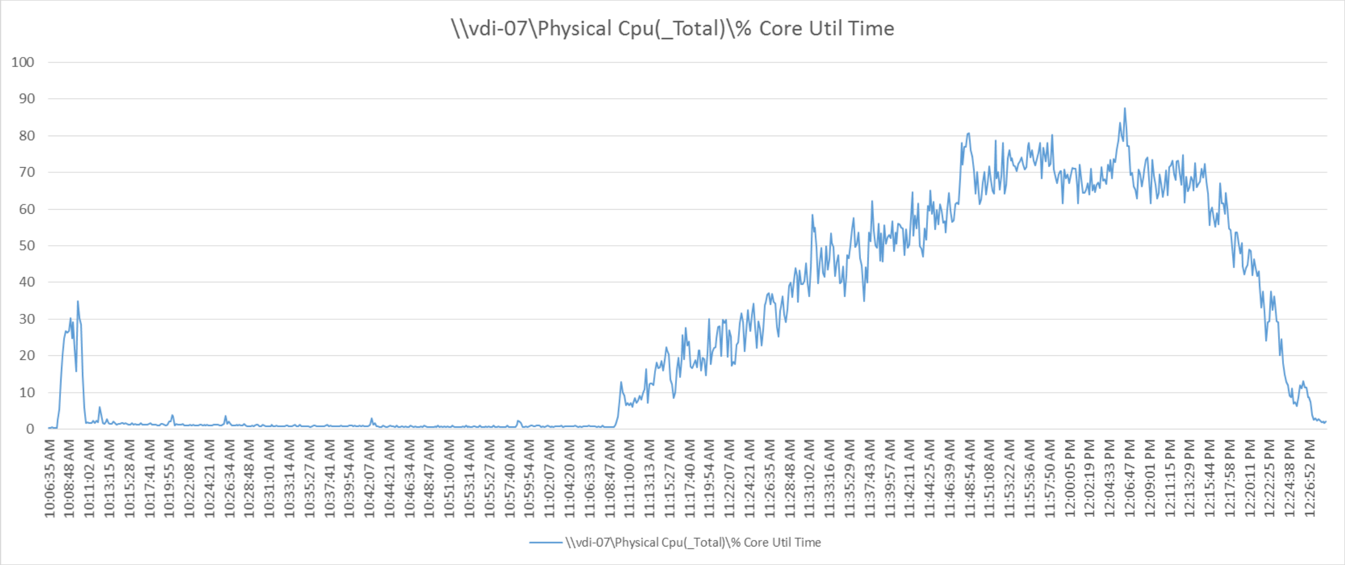

Figure 168 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

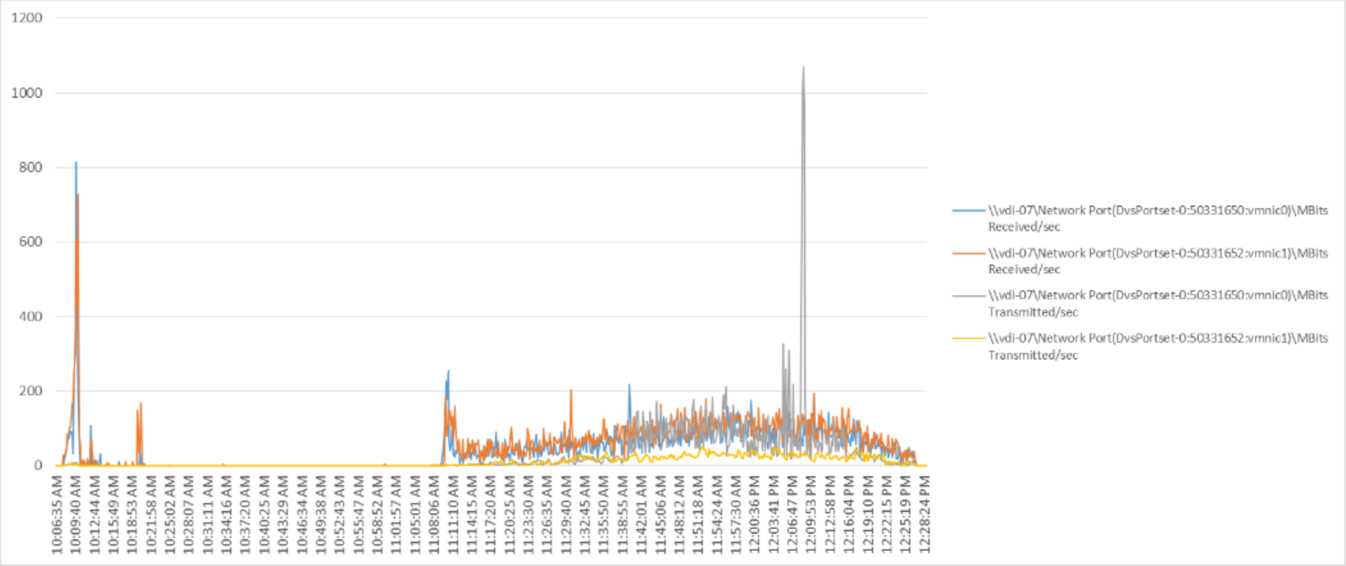

Figure 169 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

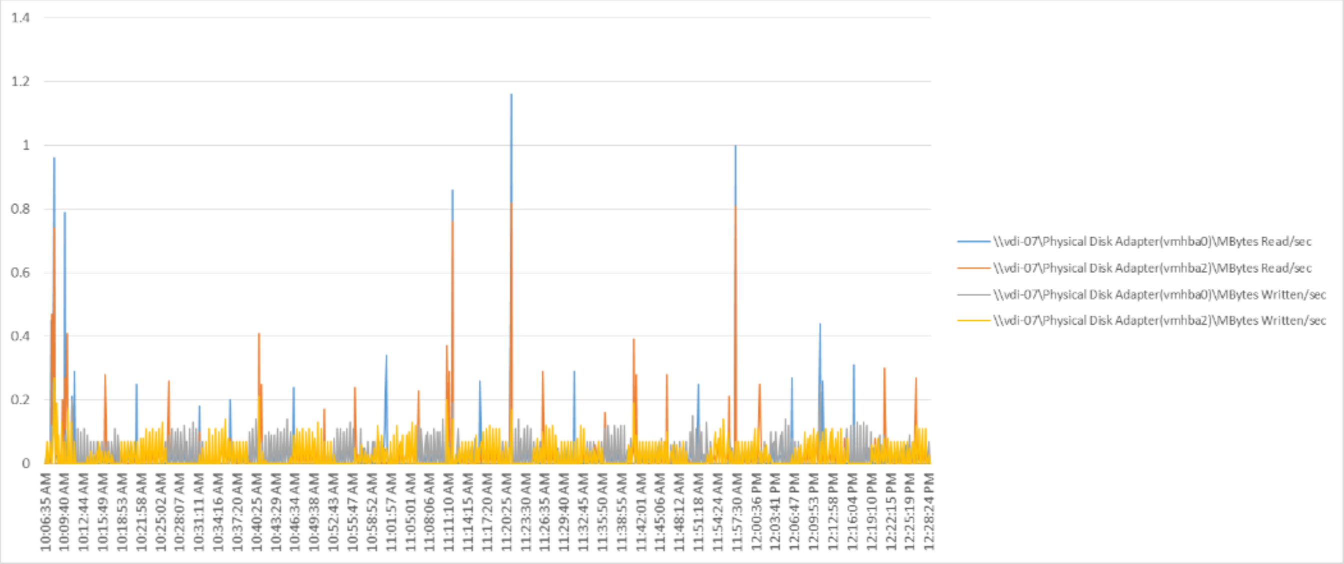

Figure 170 Full Scale | 6000 Mixed Users | RDS Hosts | Host Memory Utilization

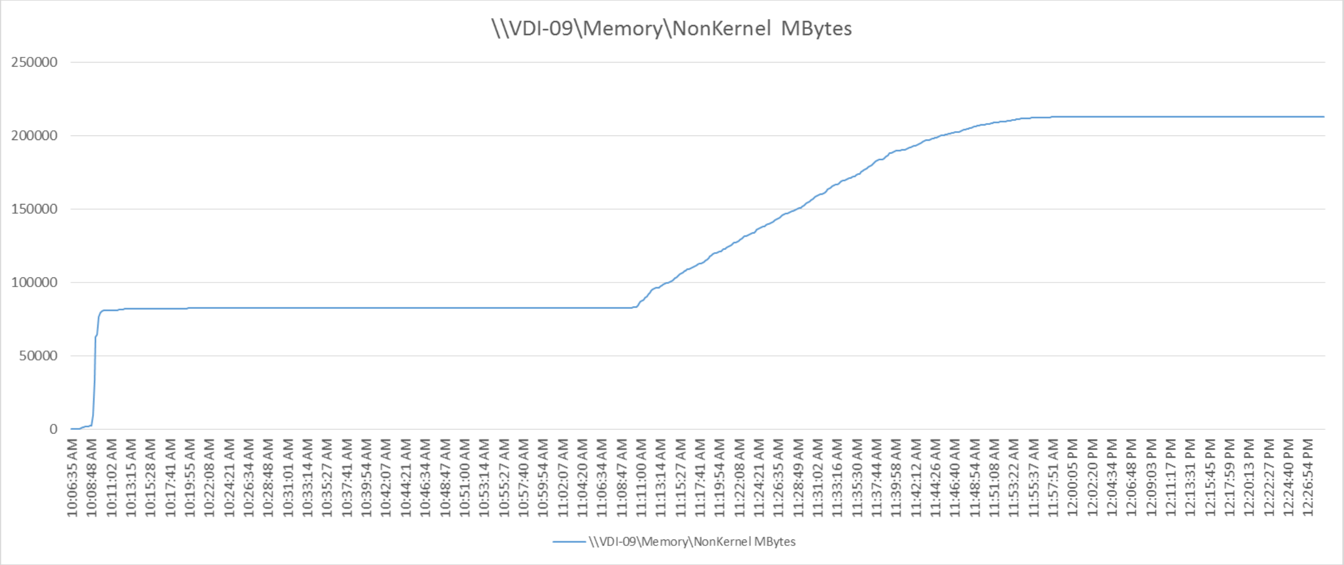

Figure 171 Full Scale | 6000 Mixed Users | RDS Hosts | Host CPU Utilization

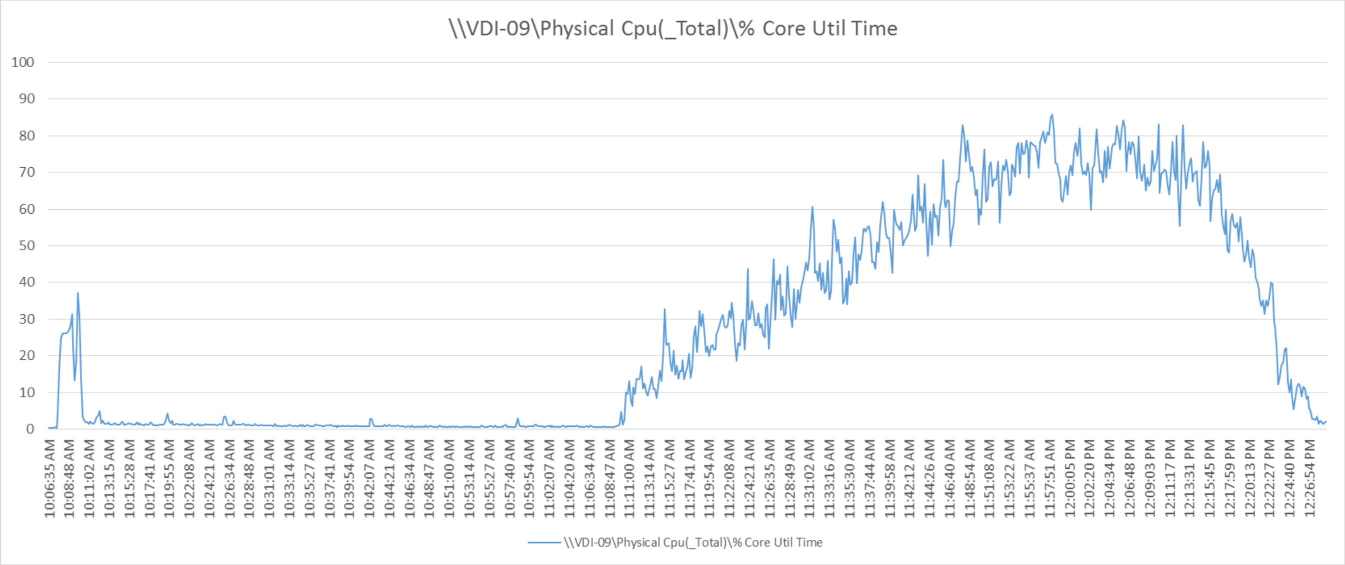

Figure 172 Full Scale | 6000 Mixed Users | RDS Hosts | Host Network Utilization

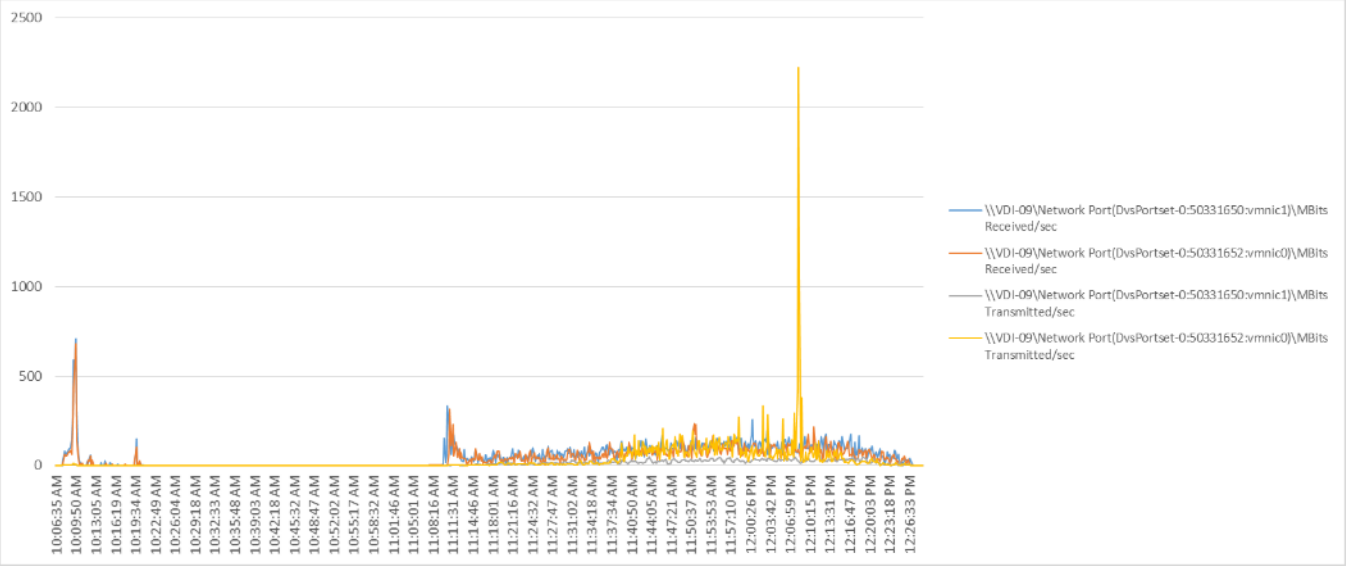

Figure 173 Full Scale | 6000 Mixed Users | RDS Hosts | Host Fibre Channel Utilization

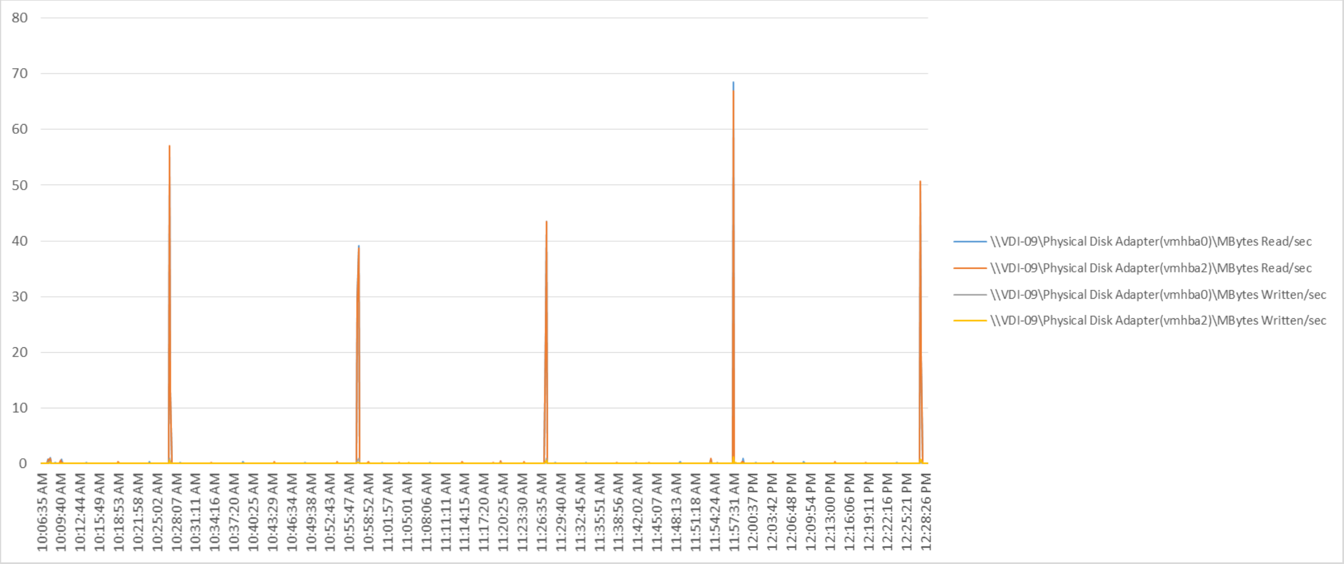

Figure 174 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

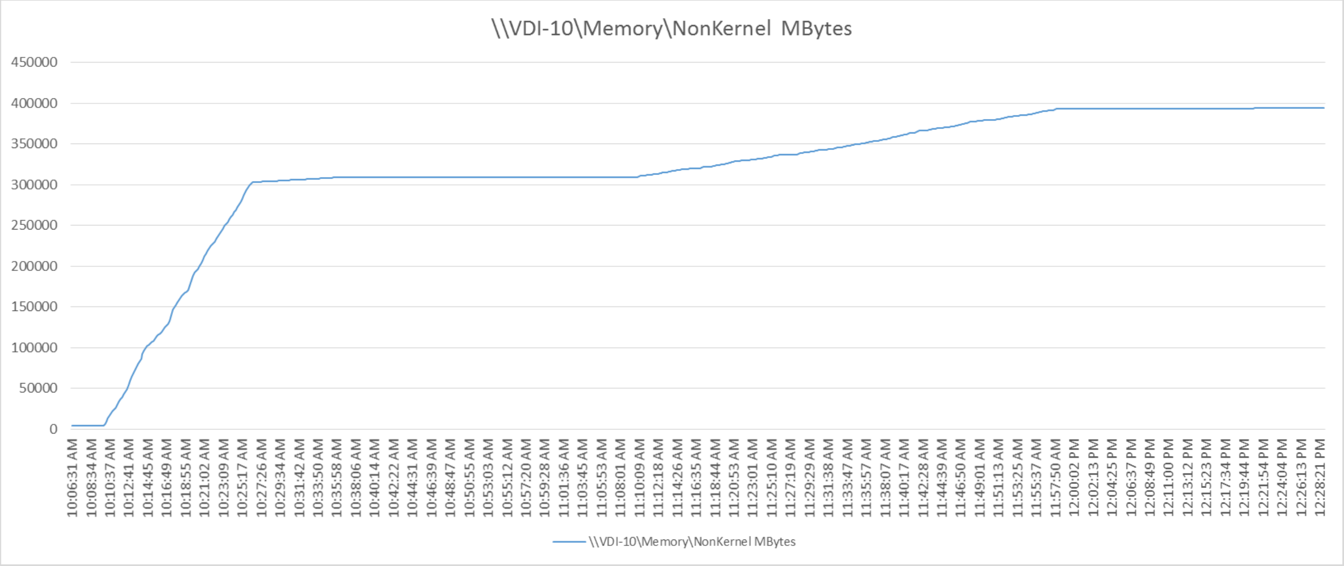

Figure 175 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

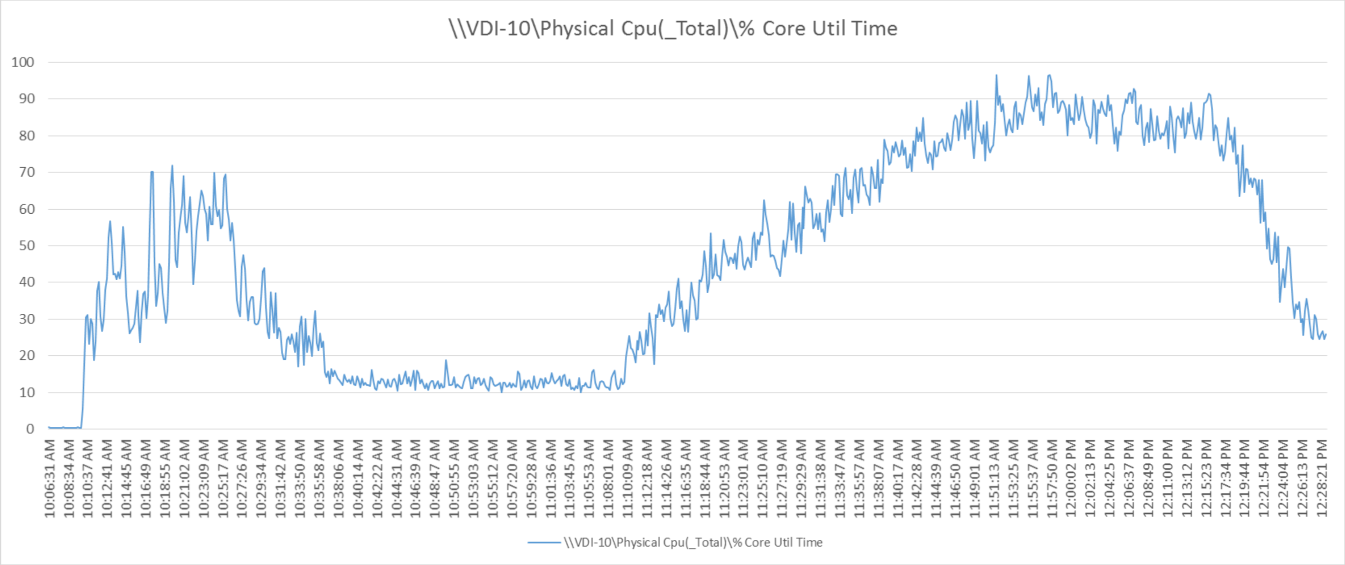

Figure 176 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

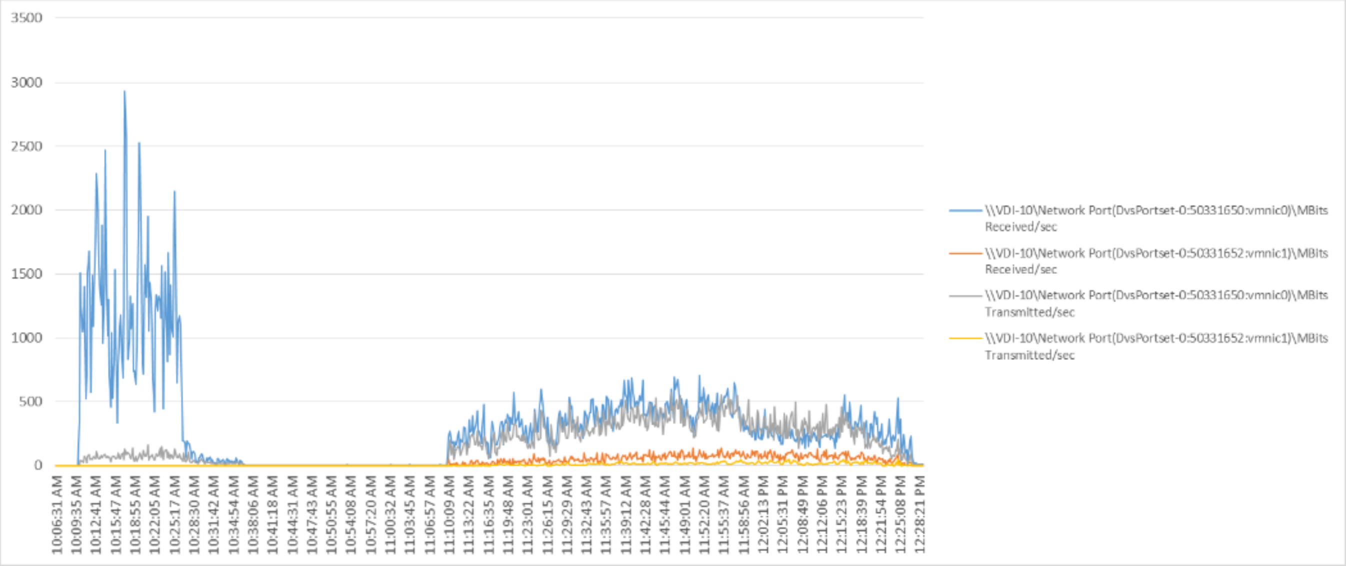

Figure 177 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

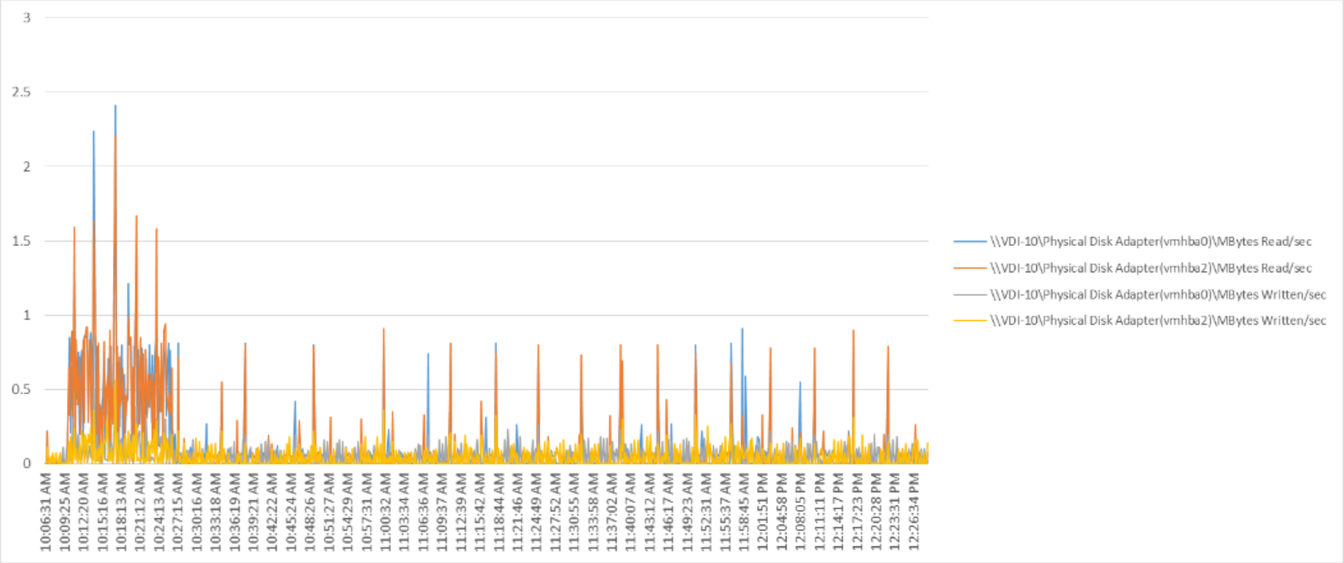

Figure 178 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

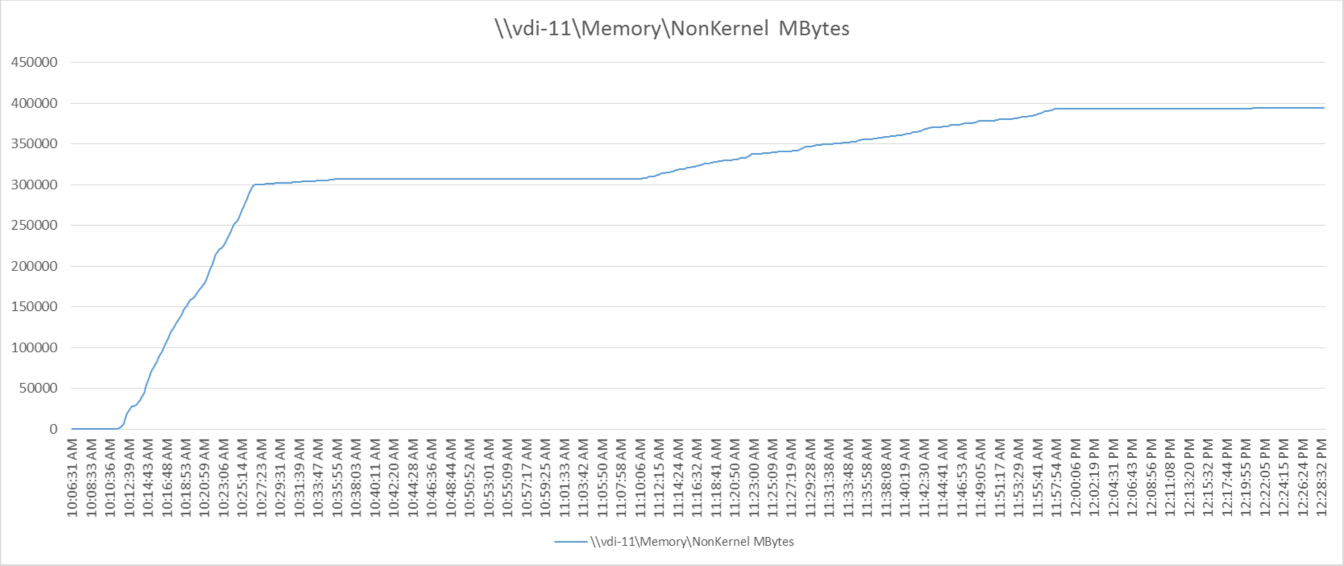

Figure 179 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

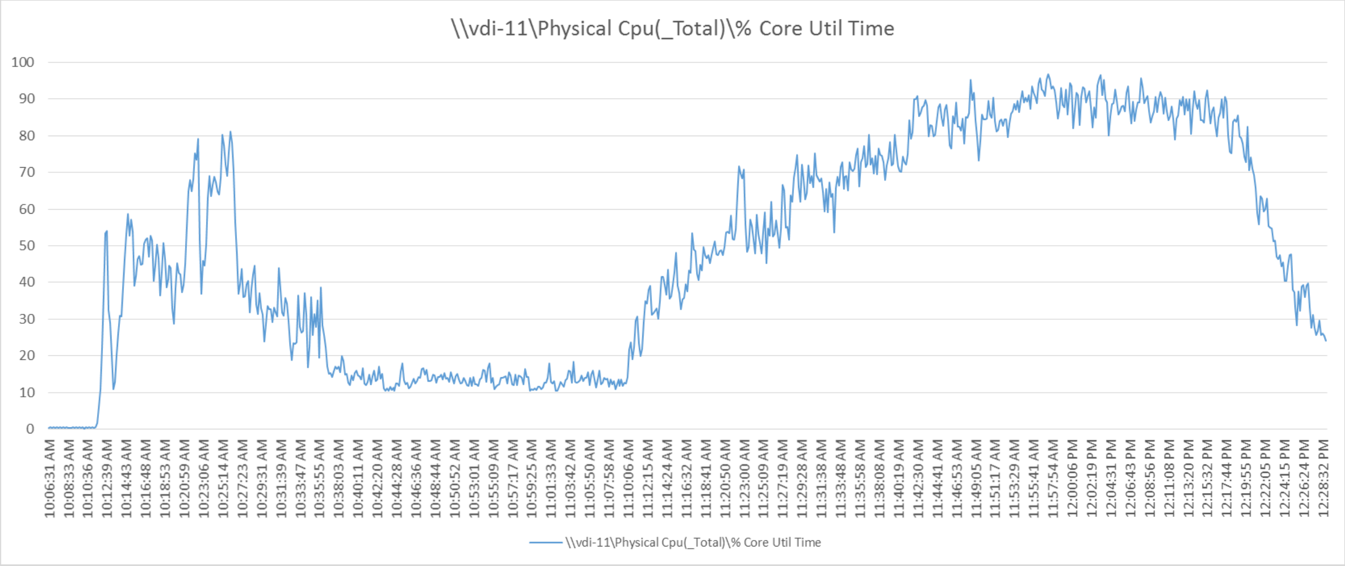

Figure 180 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

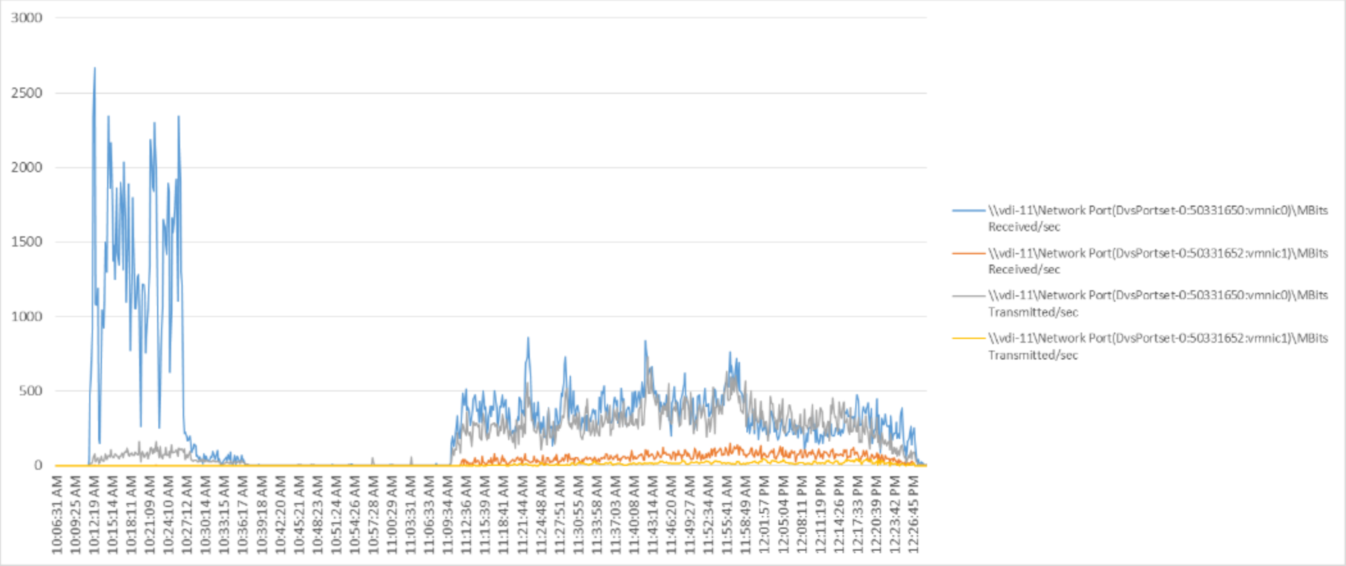

Figure 181 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

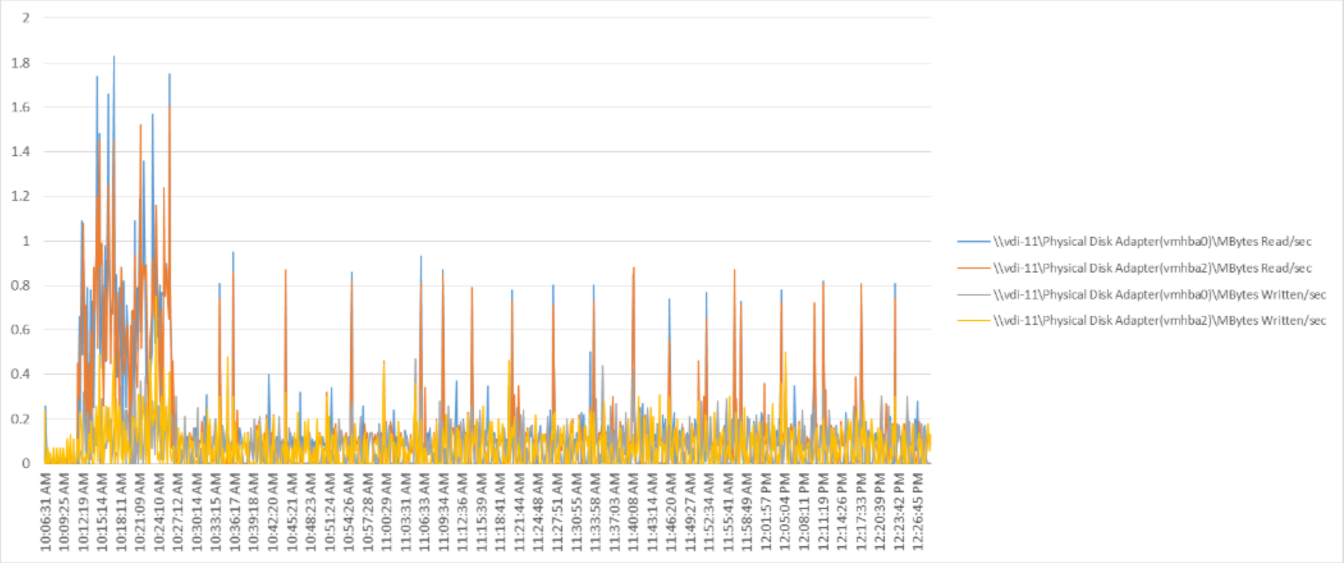

Figure 182 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

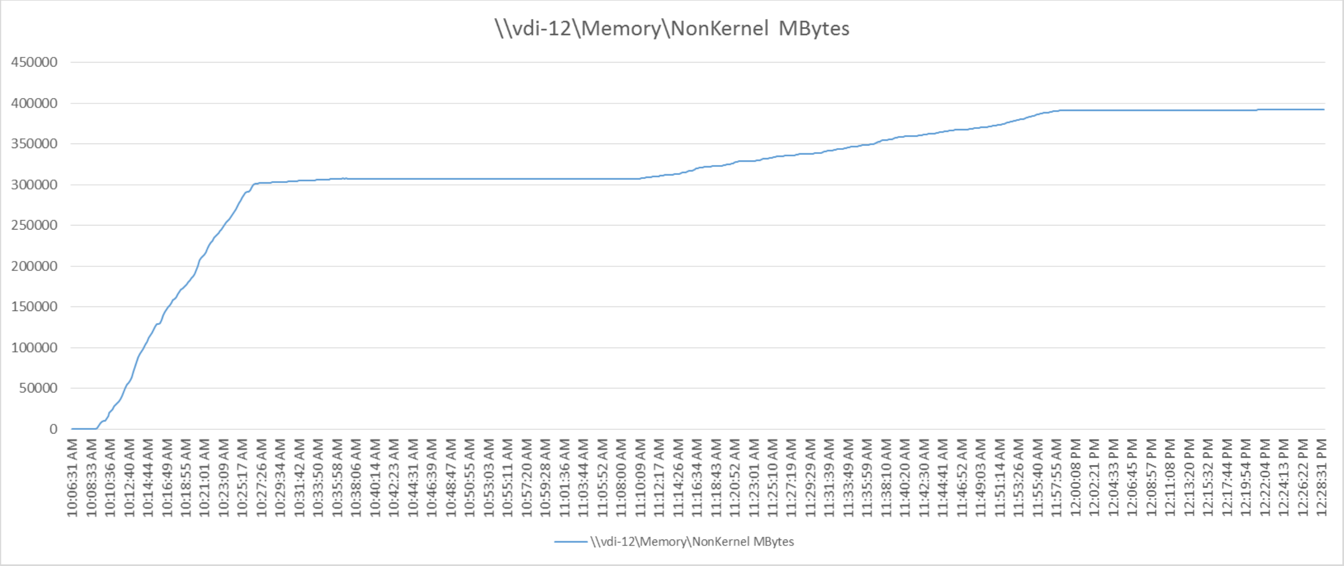

Figure 183 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

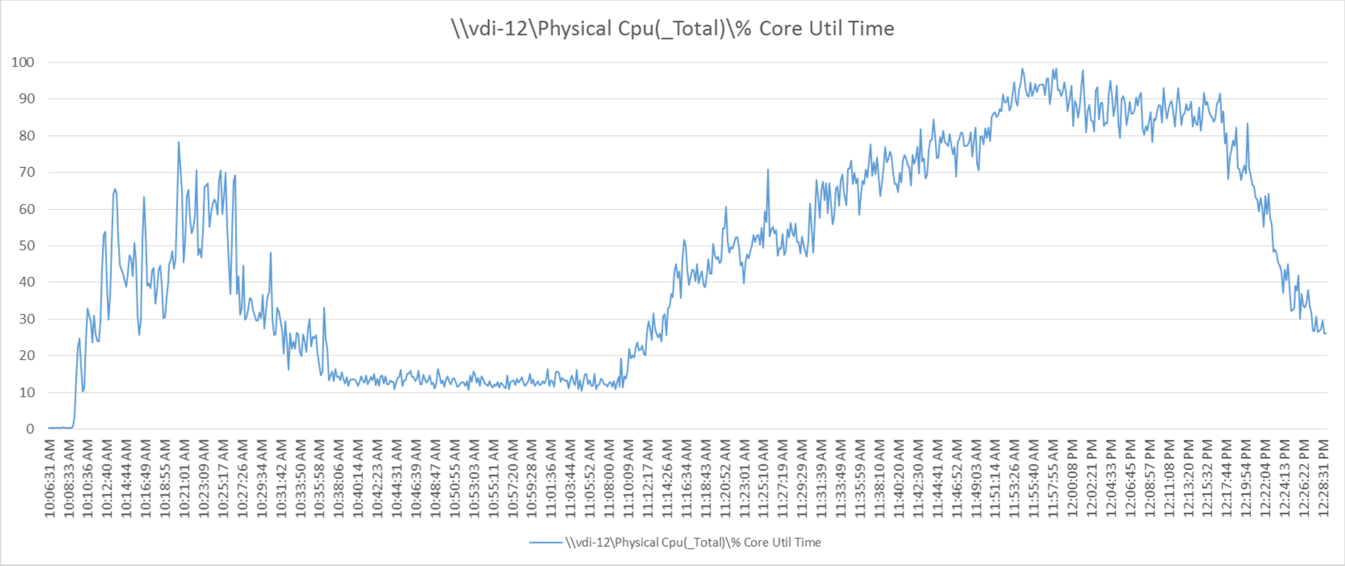

Figure 184 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

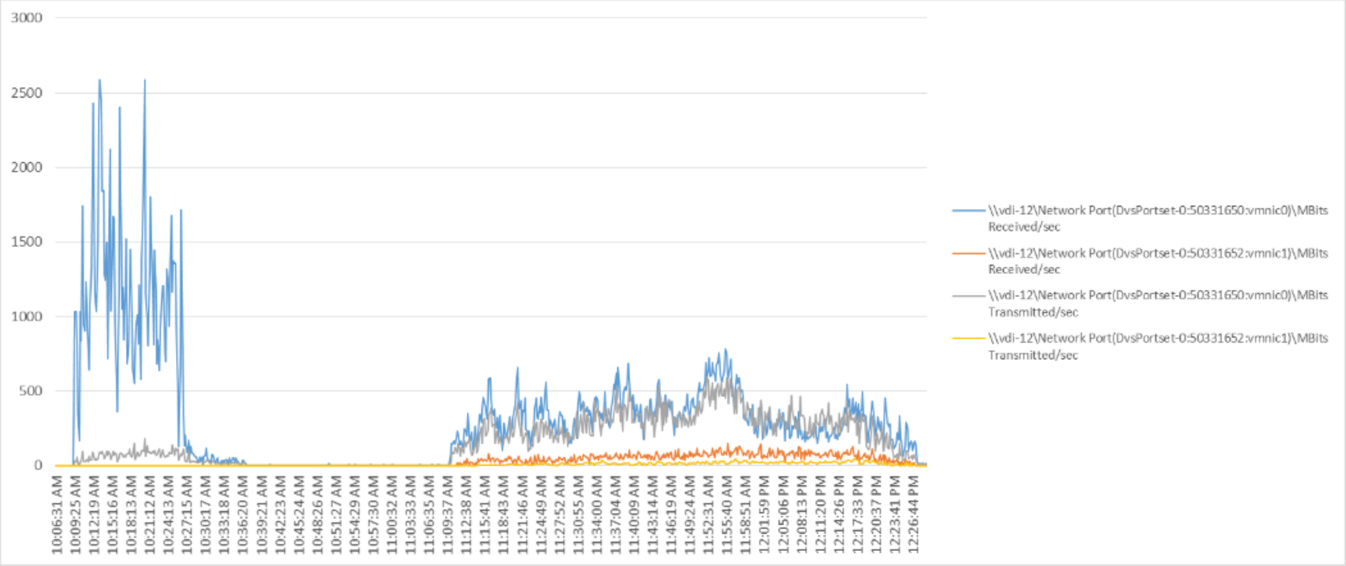

Figure 185 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

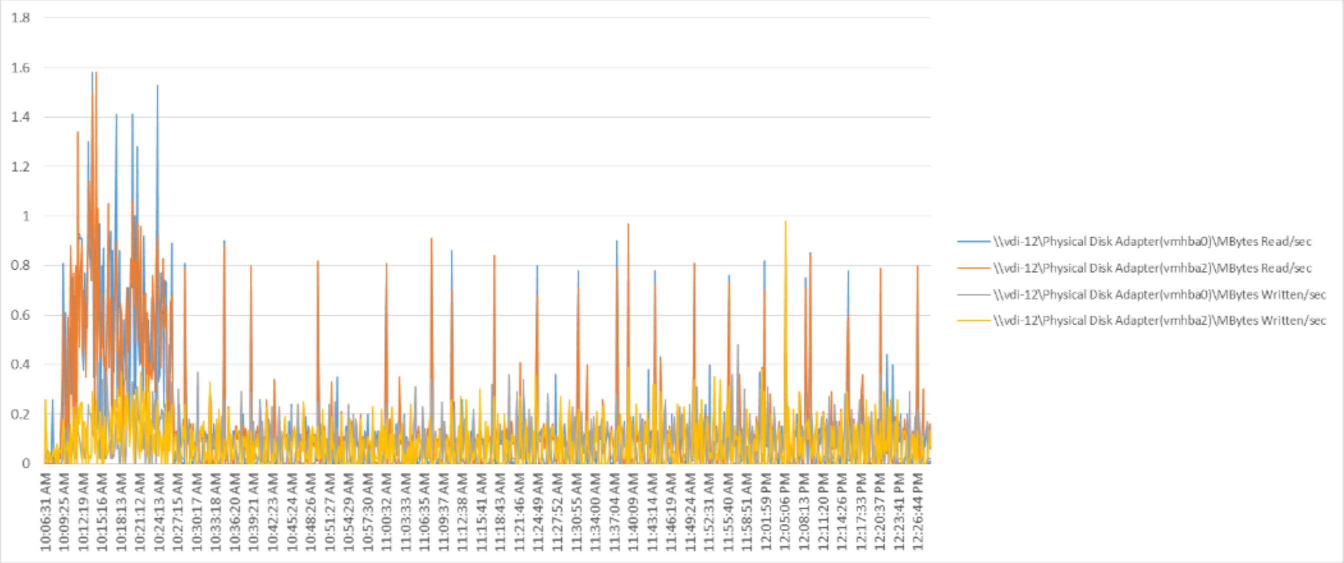

Figure 186 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

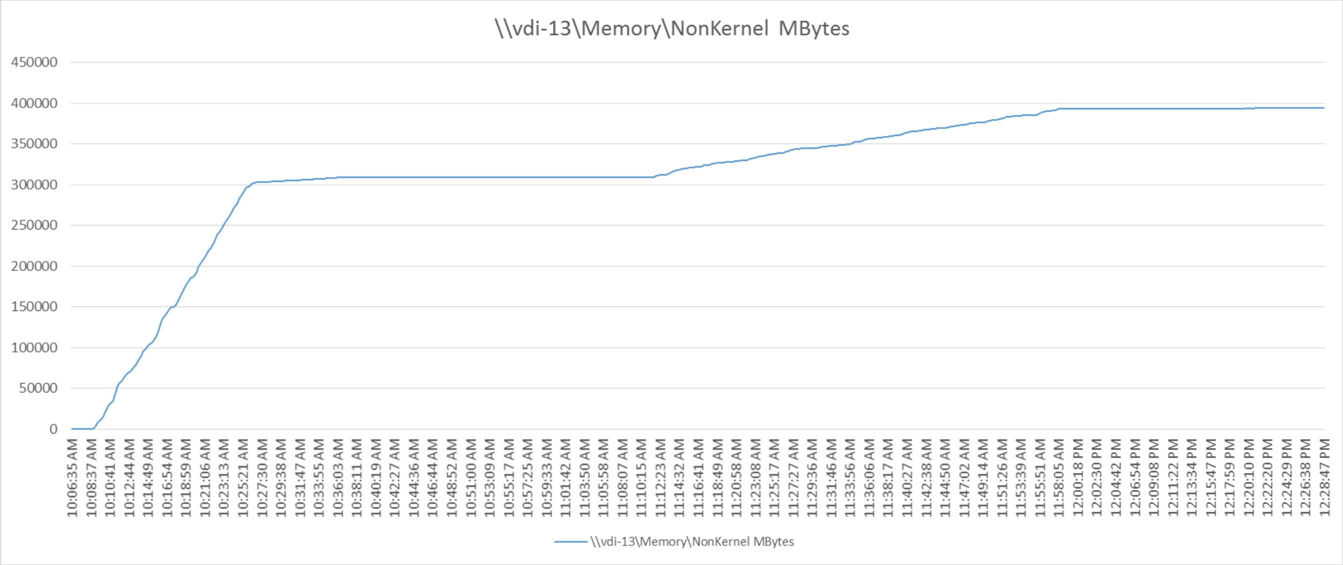

Figure 187 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

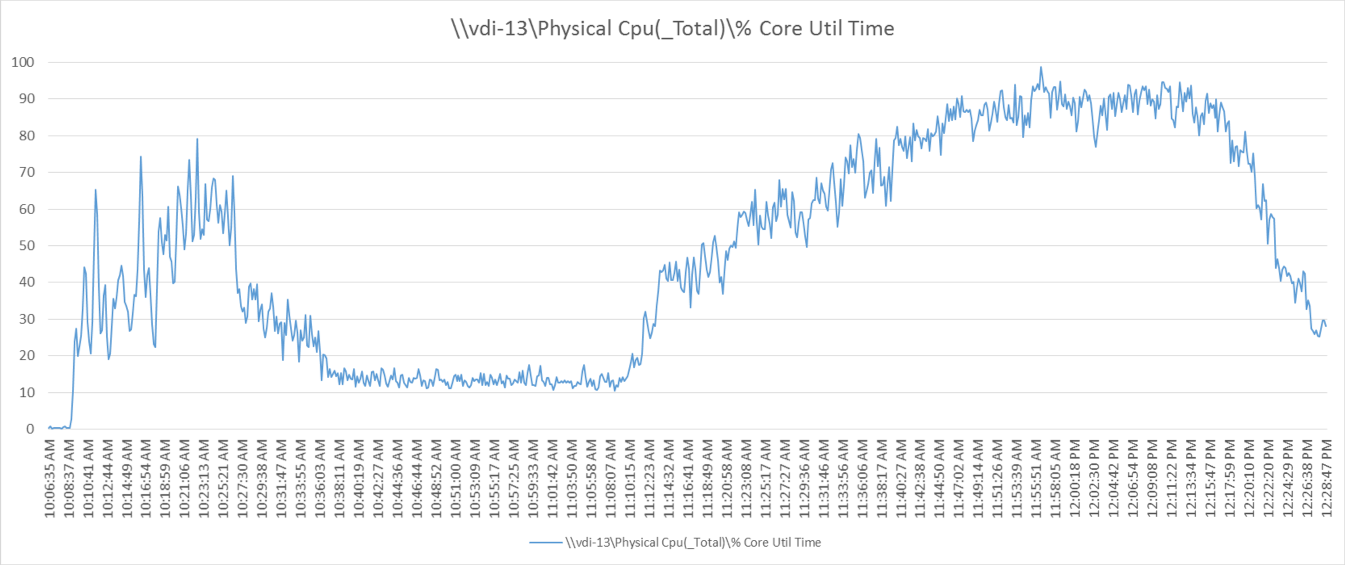

Figure 188 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

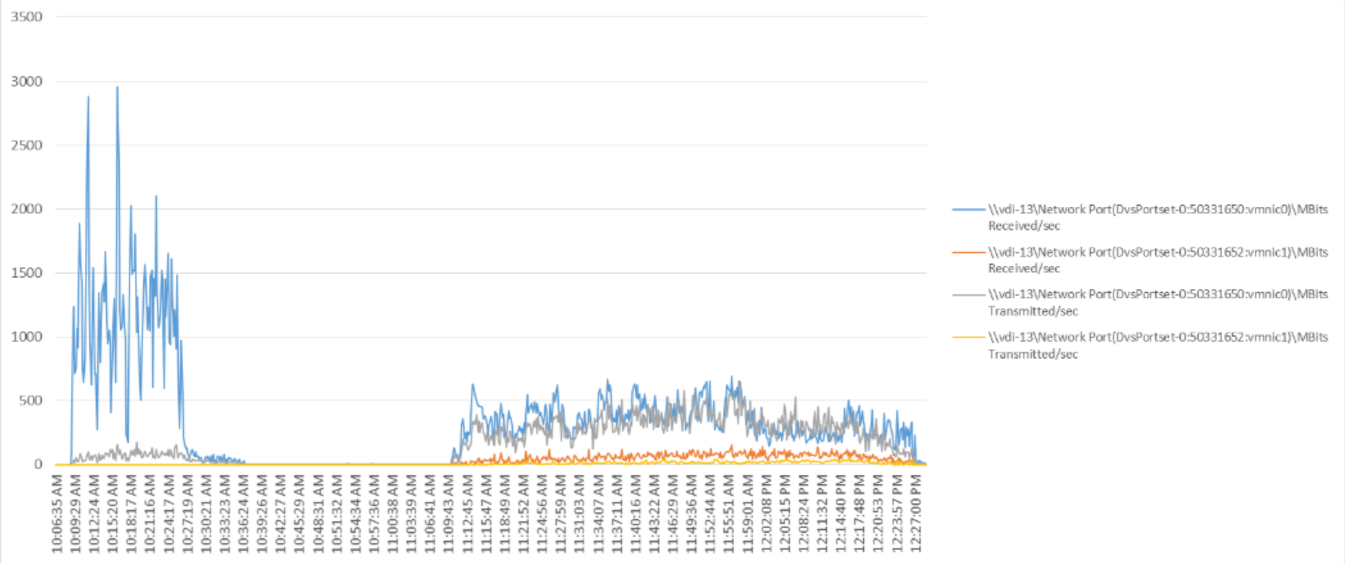

Figure 189 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

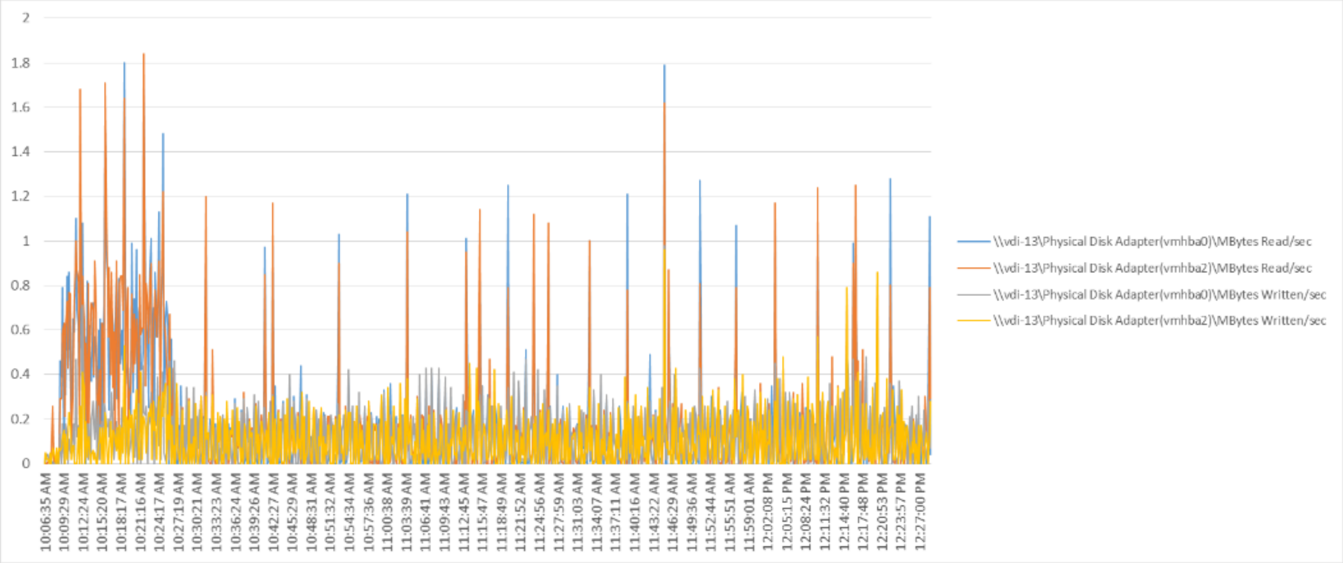

Figure 190 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

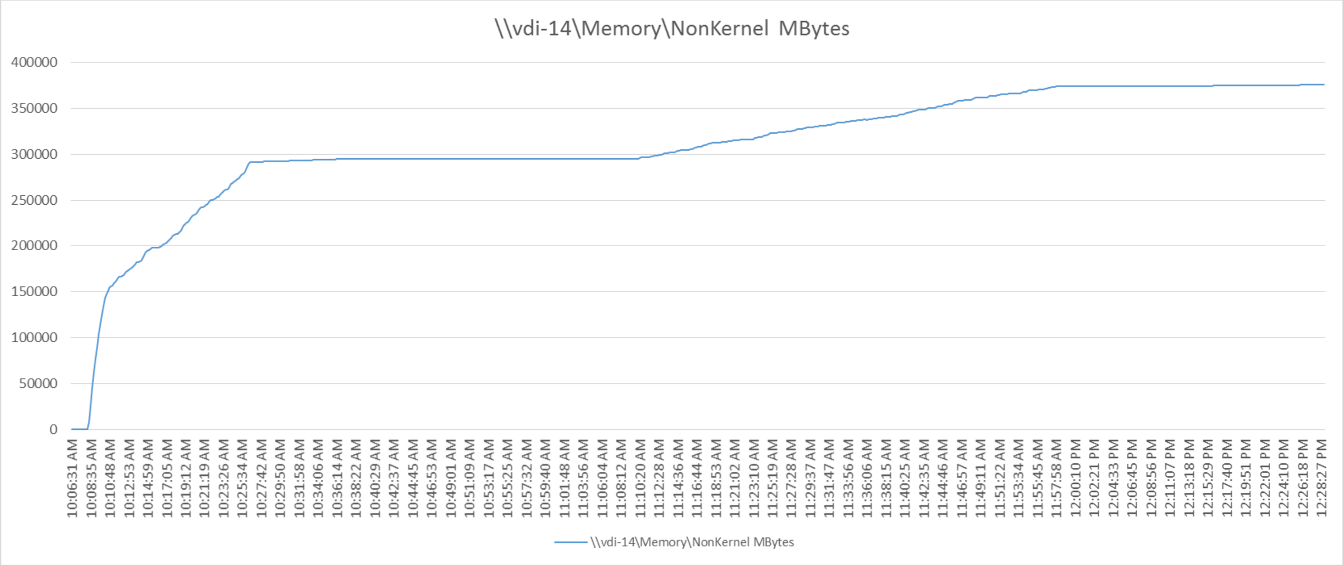

Figure 191 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

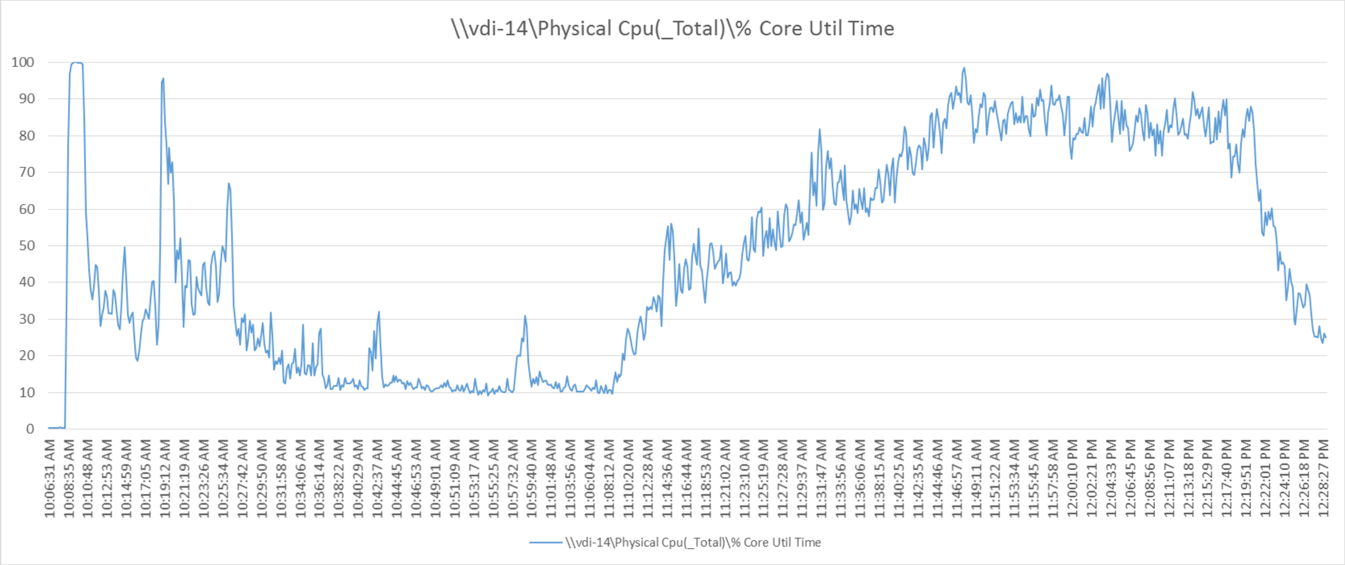

Figure 192 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

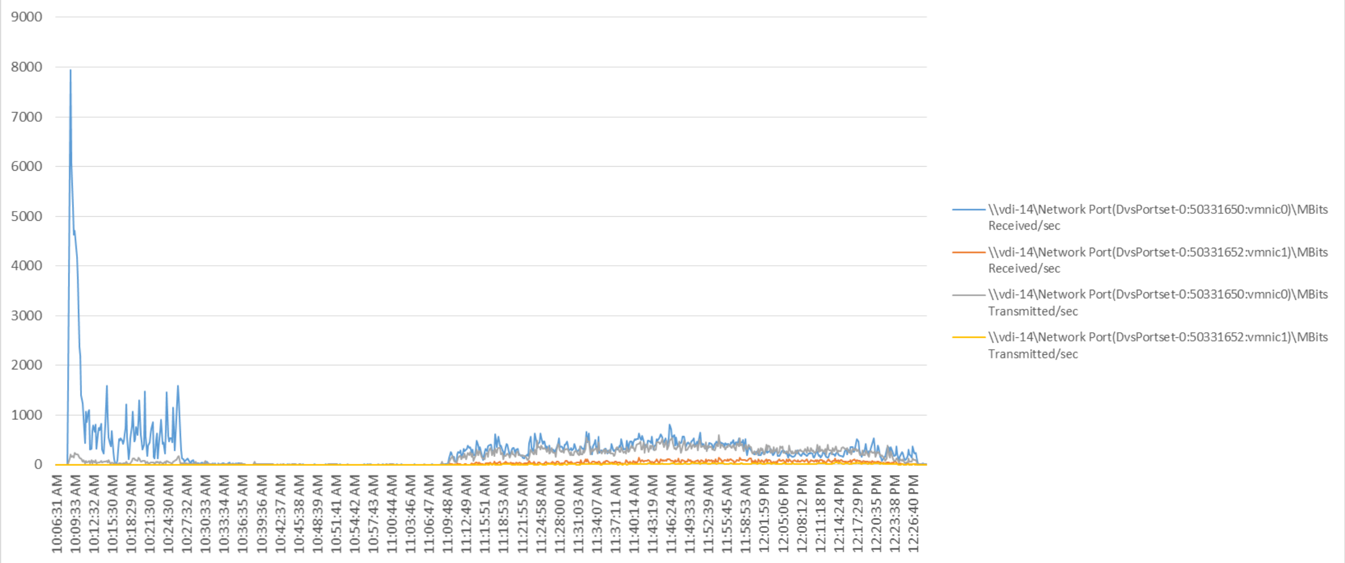

Figure 193 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

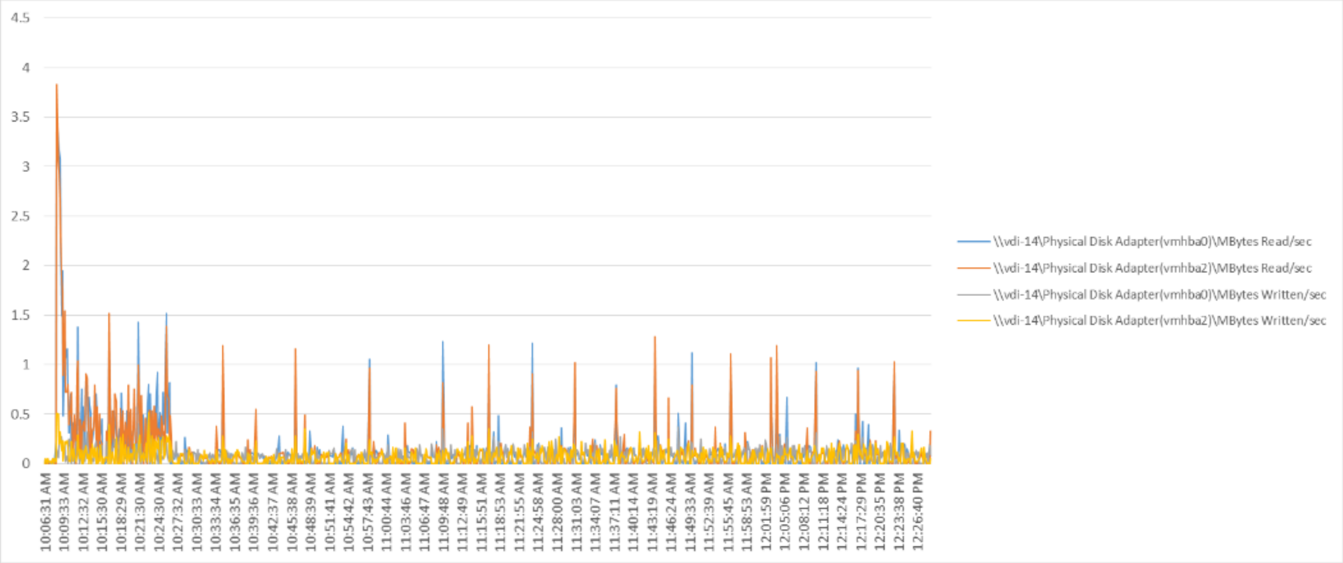

Figure 194 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

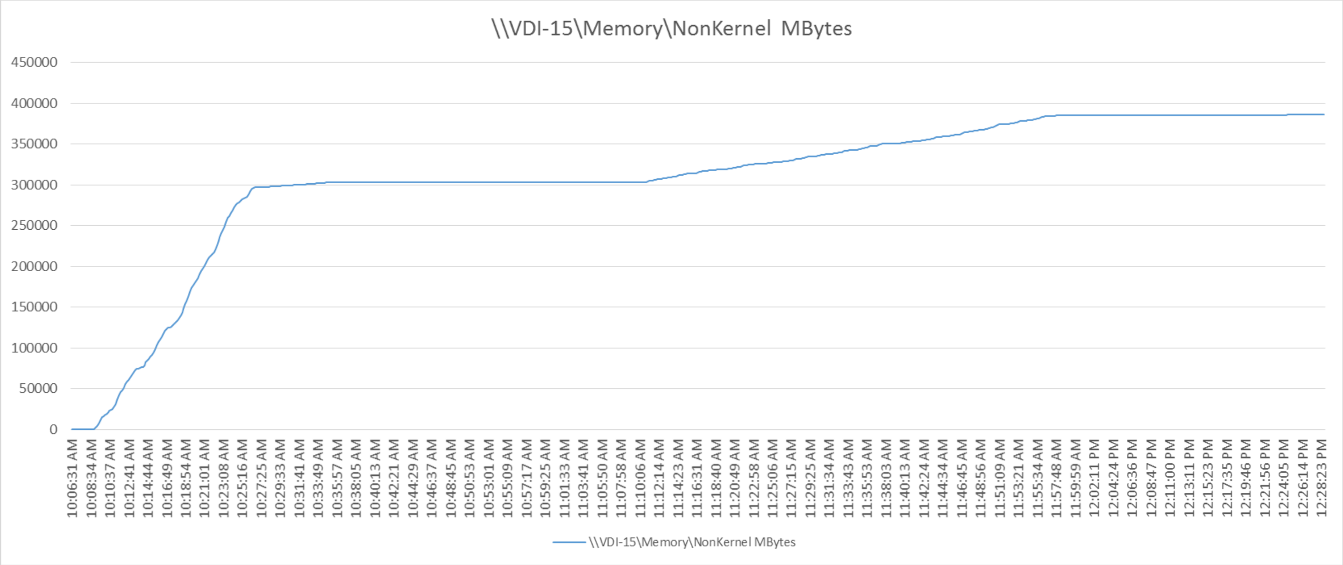

Figure 195 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

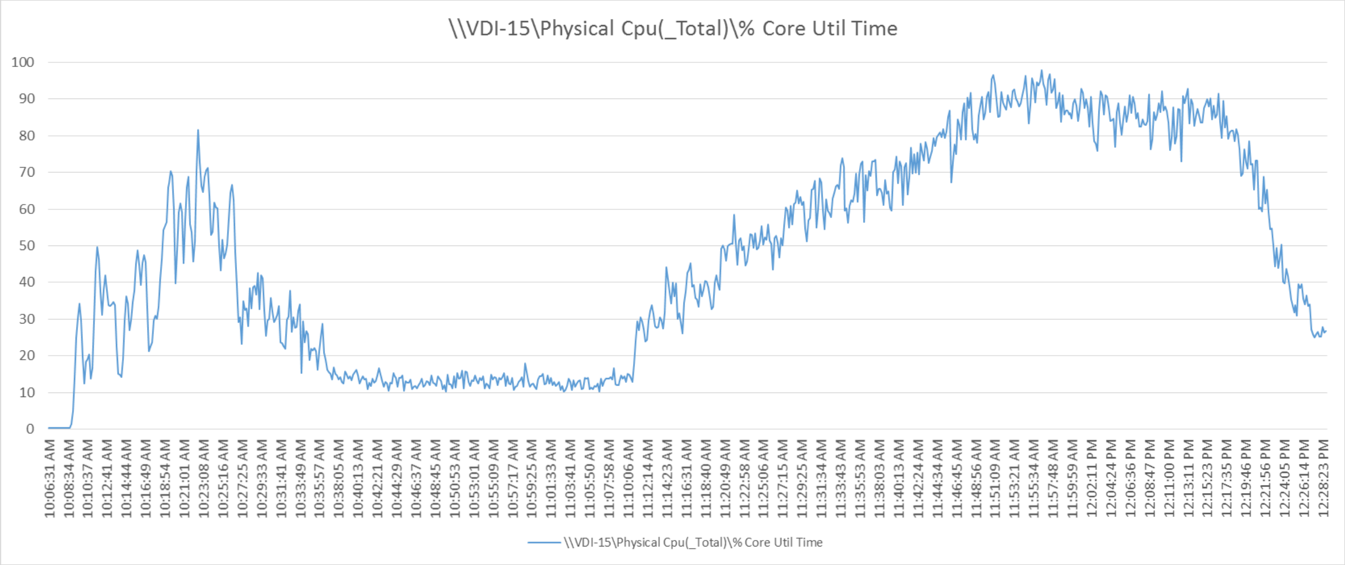

Figure 196 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

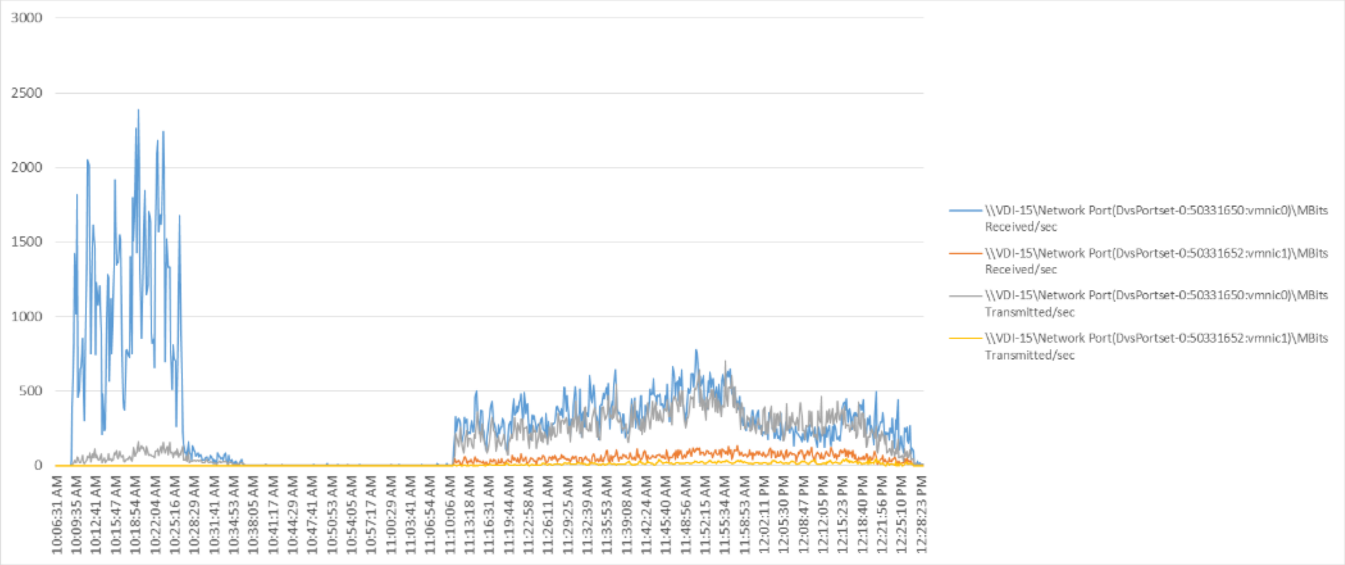

Figure 197 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

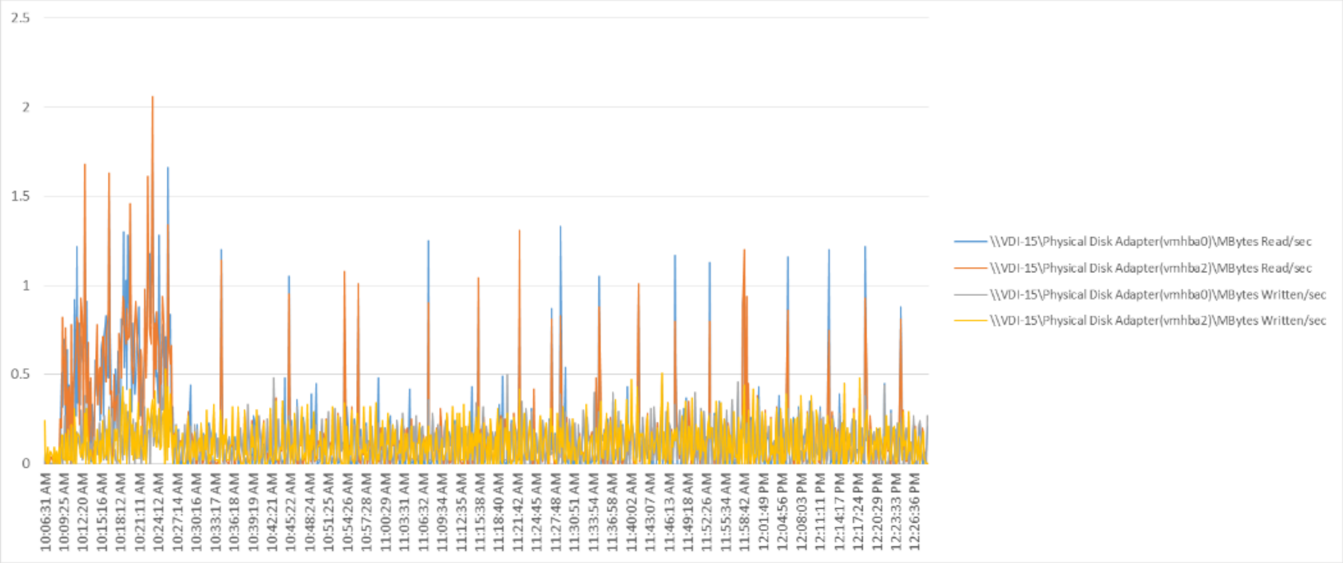

Figure 198 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

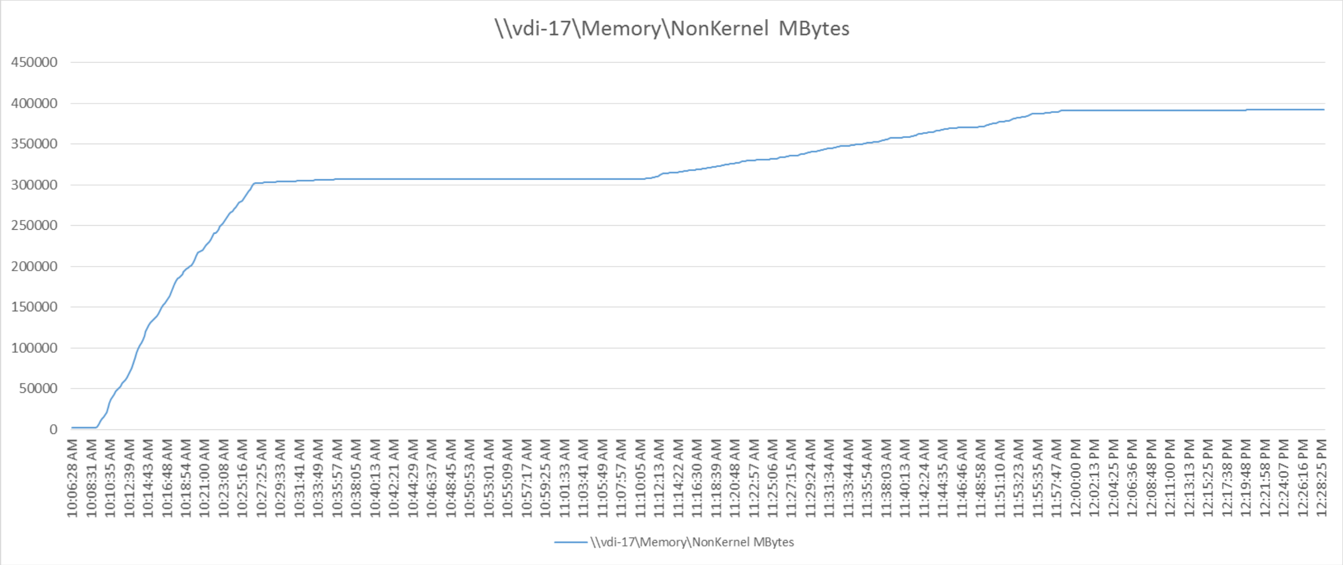

Figure 199 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

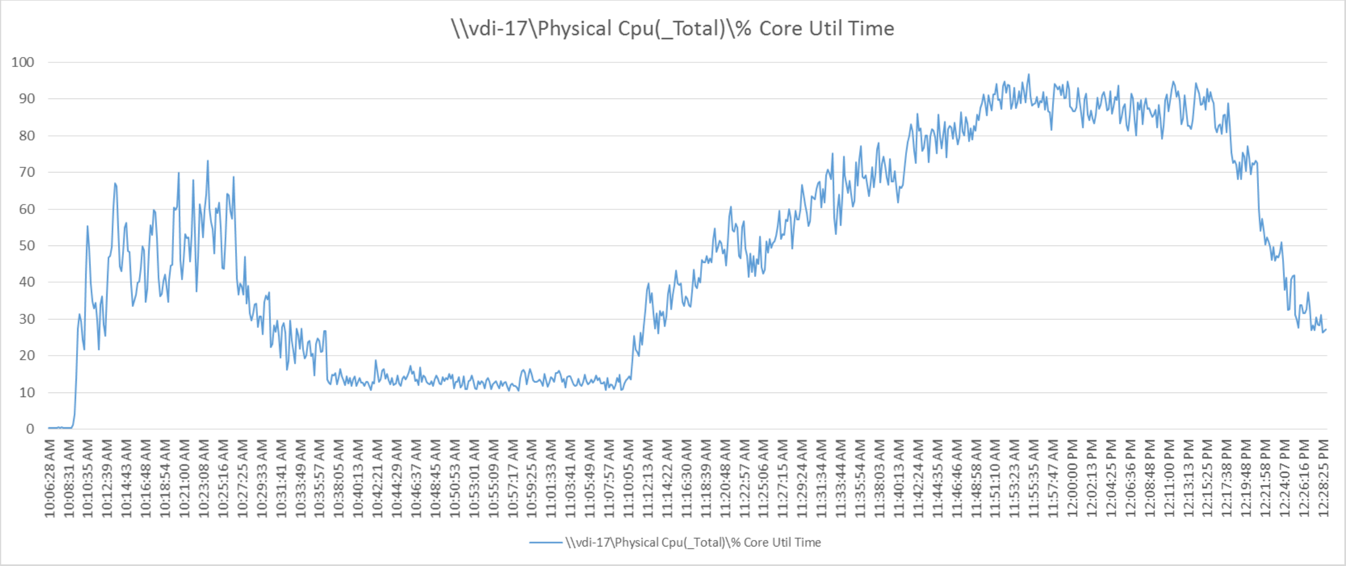

Figure 200 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

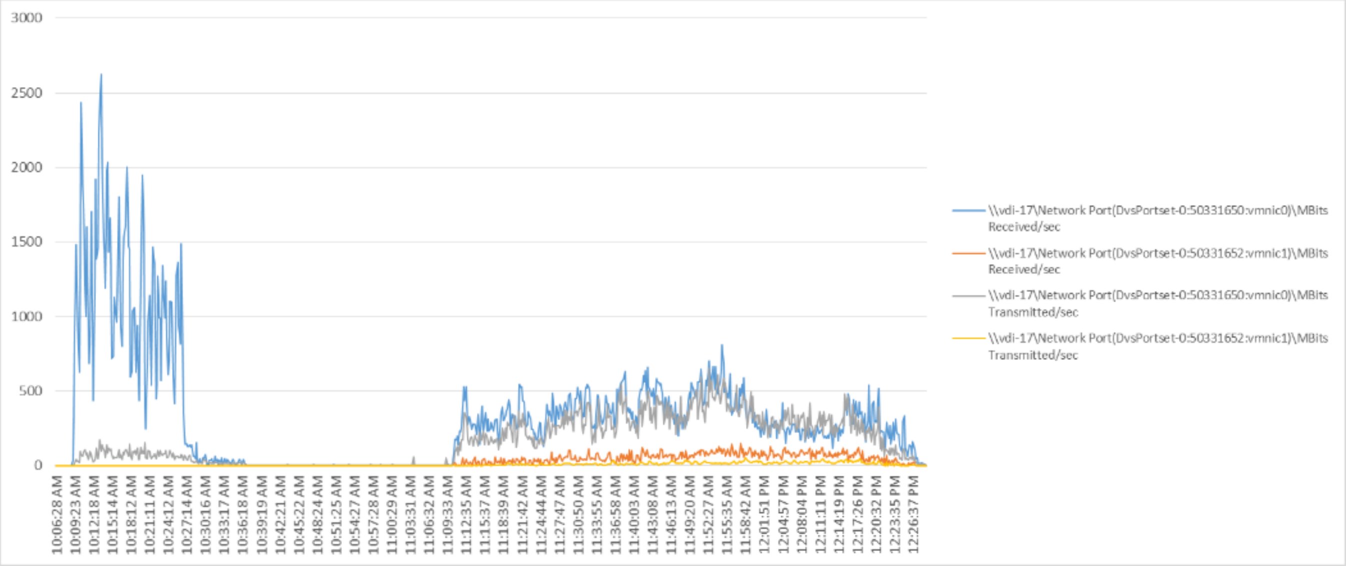

Figure 201 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

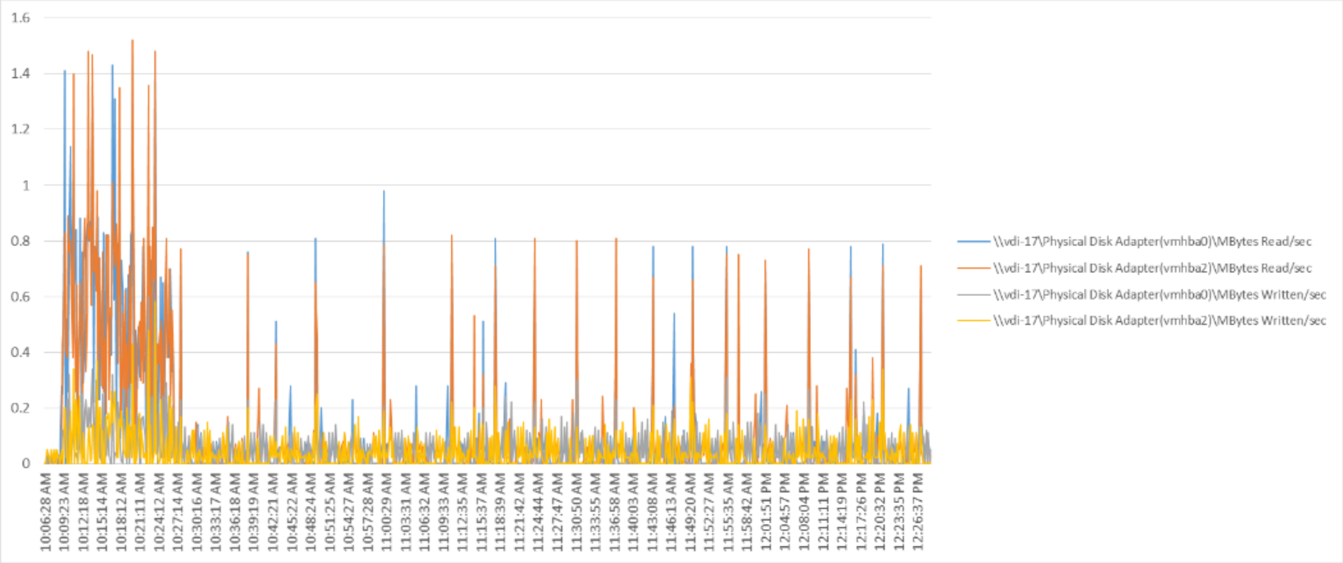

Figure 202 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

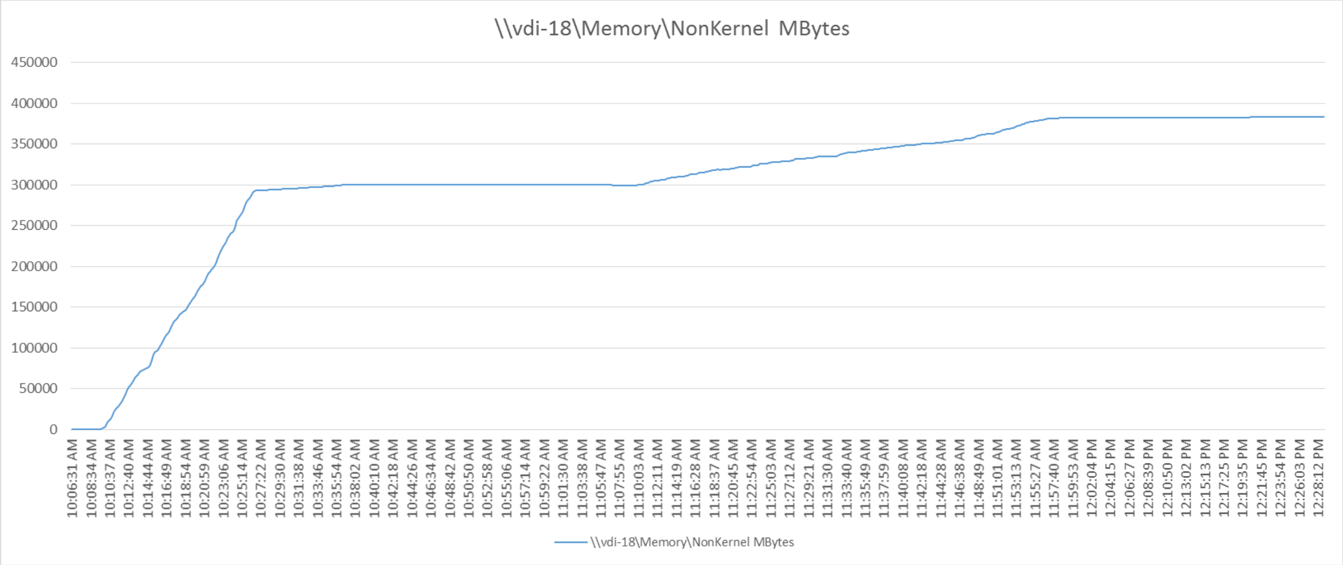

Figure 203 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

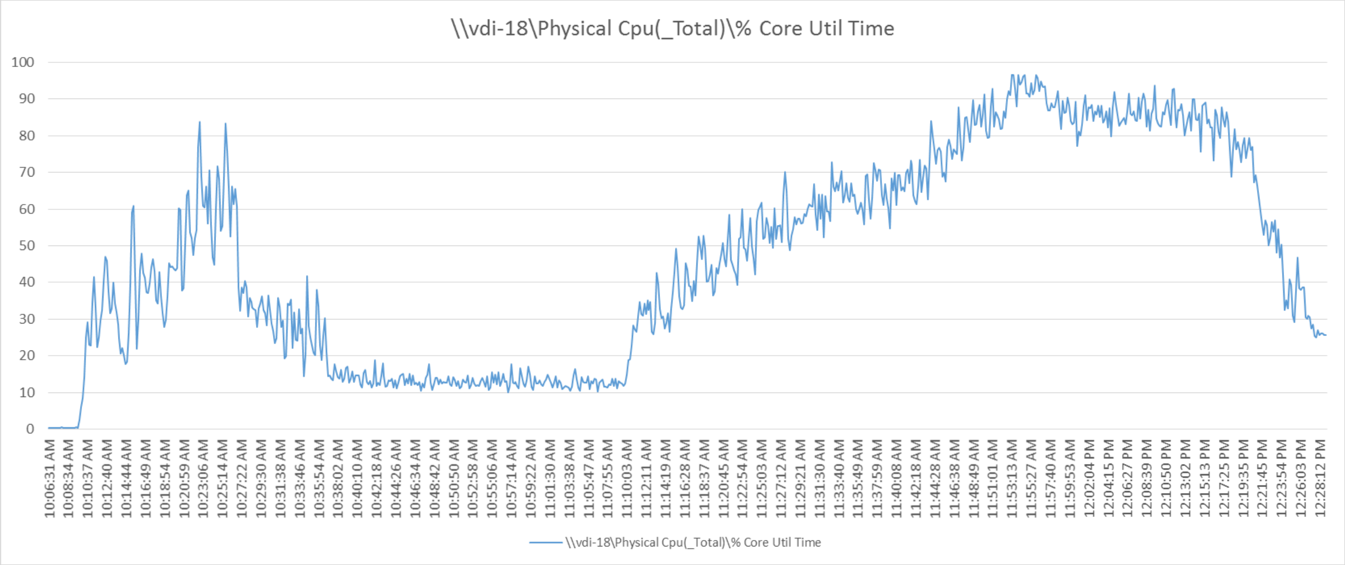

Figure 204 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

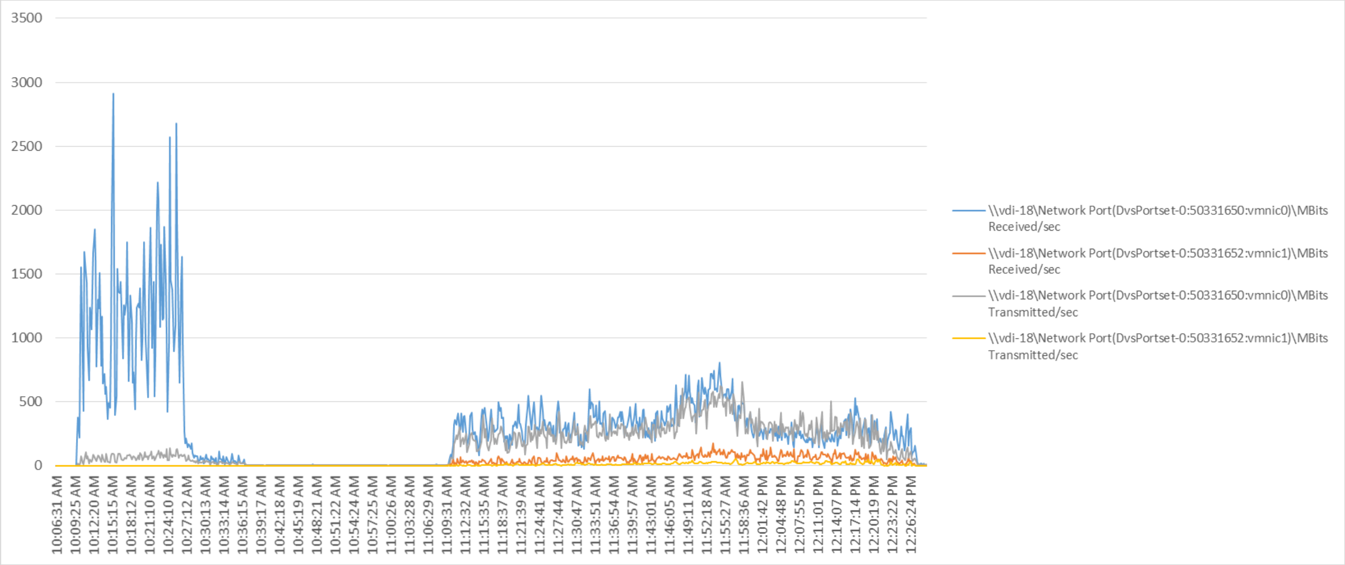

Figure 205 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

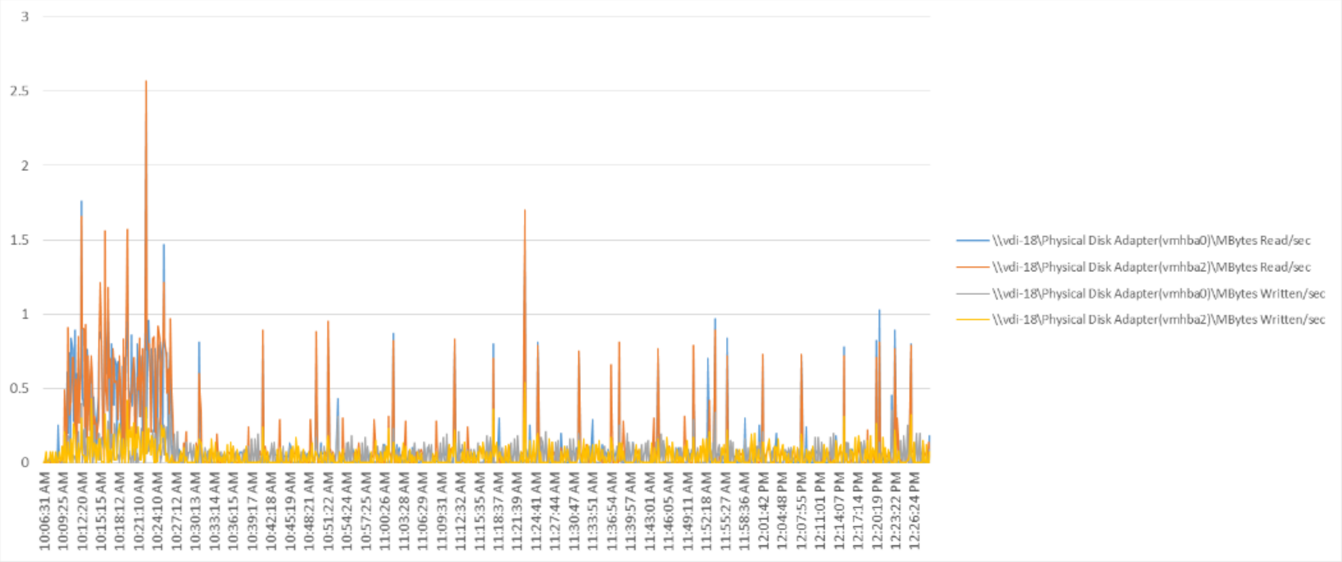

Figure 206 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

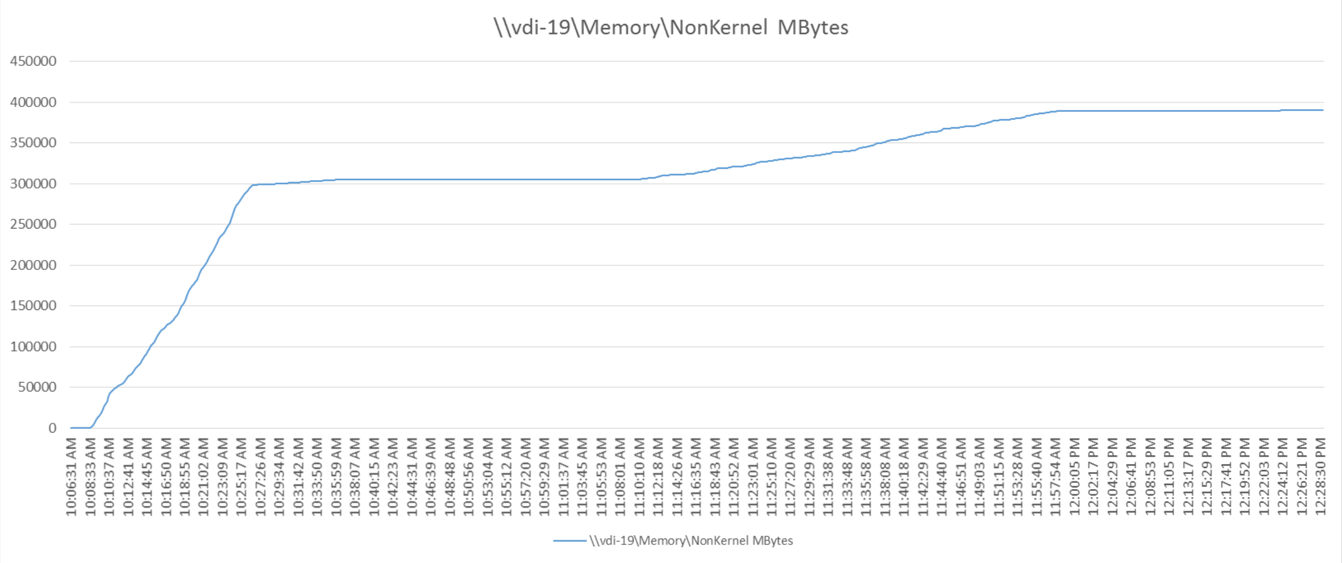

Figure 207 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

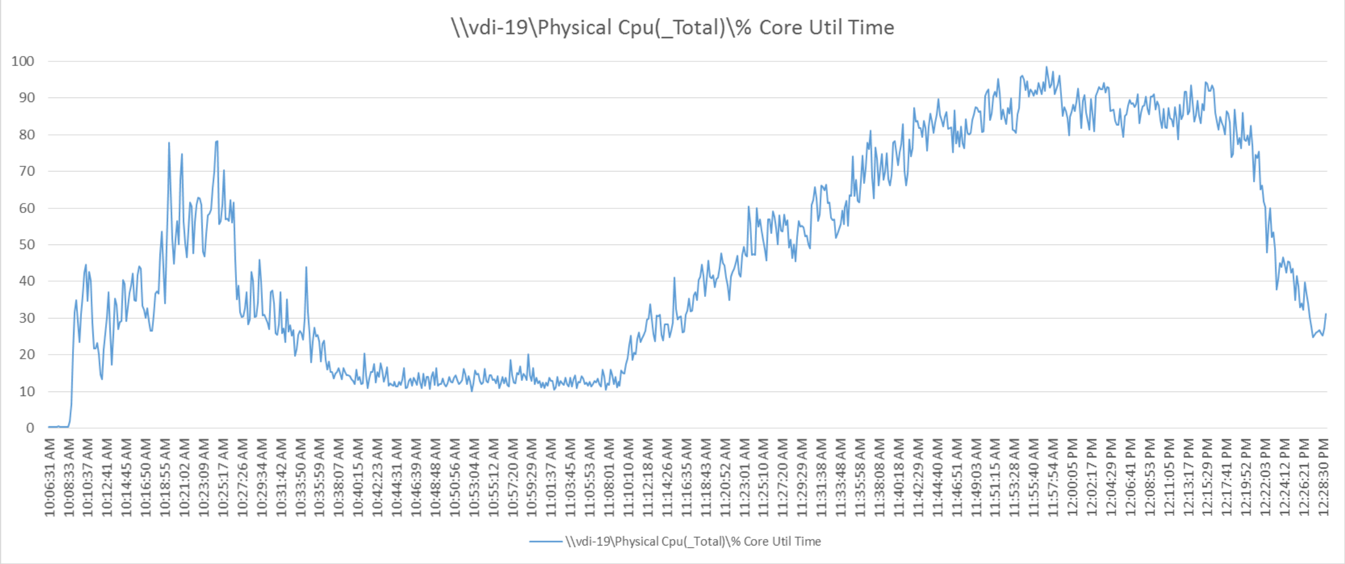

Figure 208 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

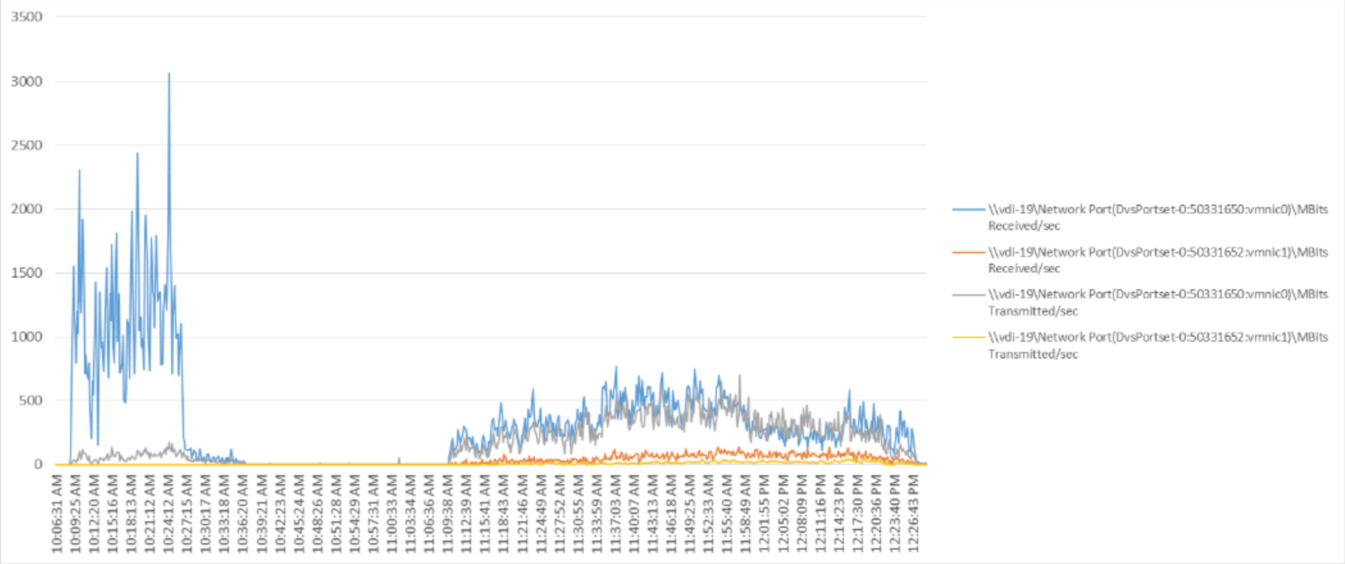

Figure 209 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

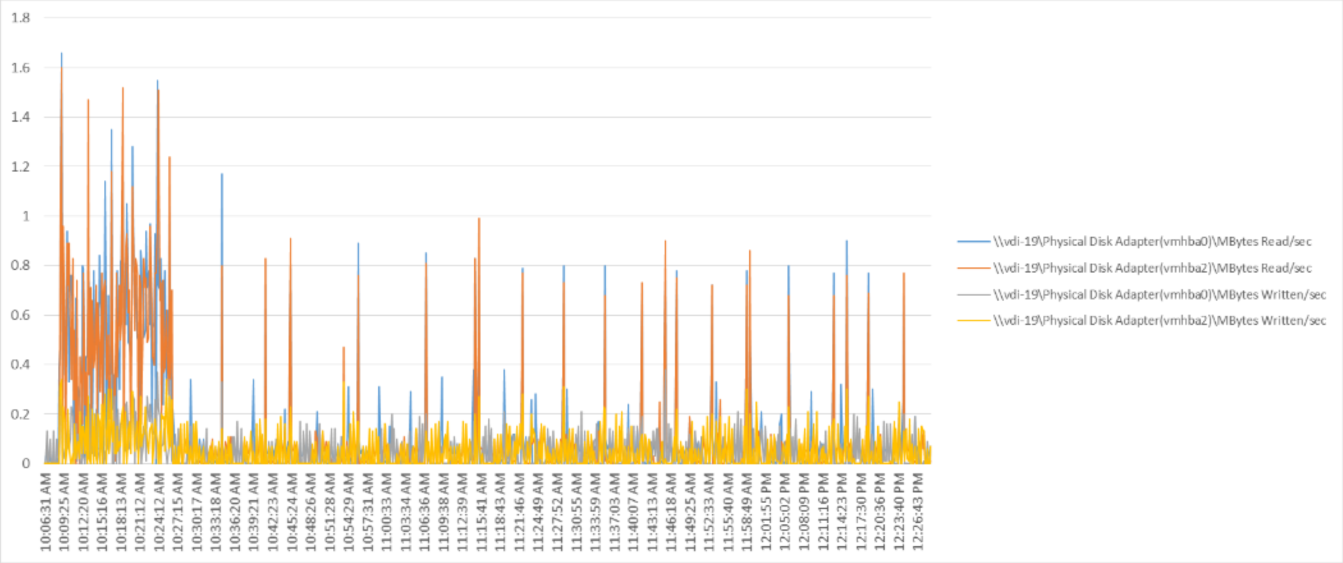

Figure 210 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

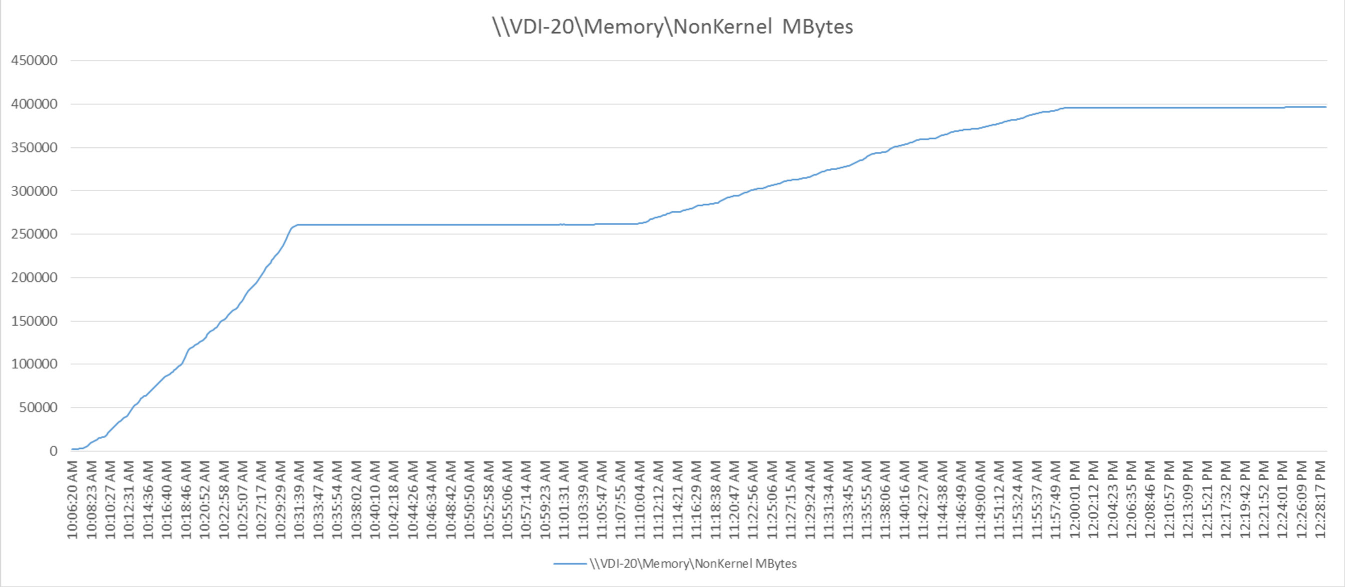

Figure 211 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

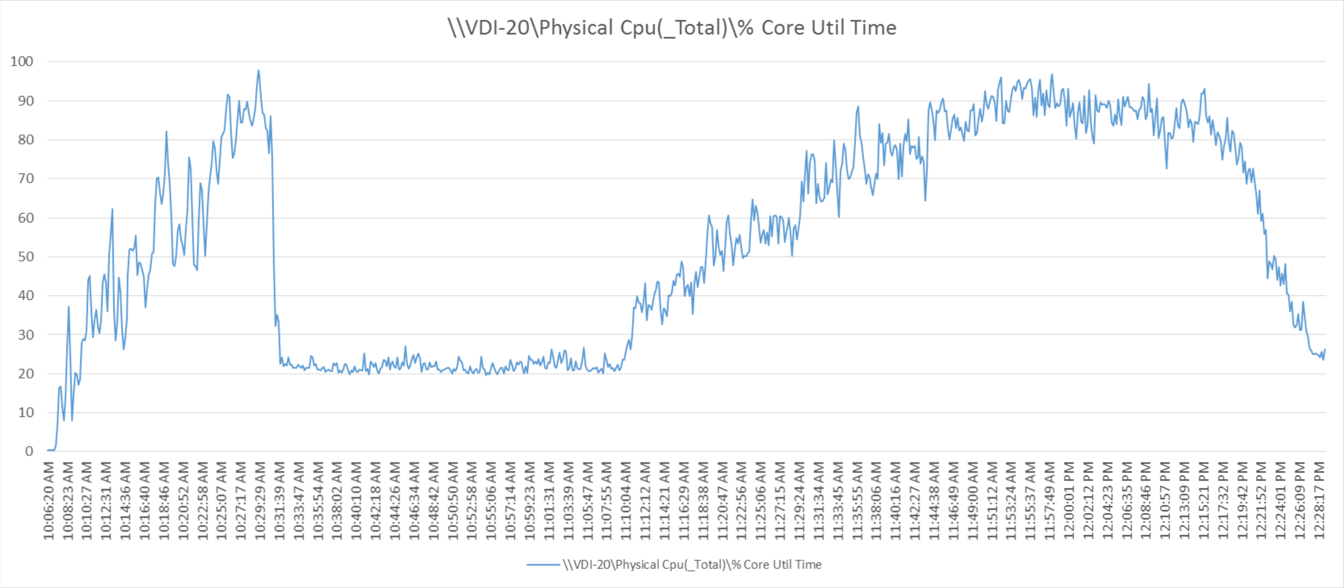

Figure 212 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

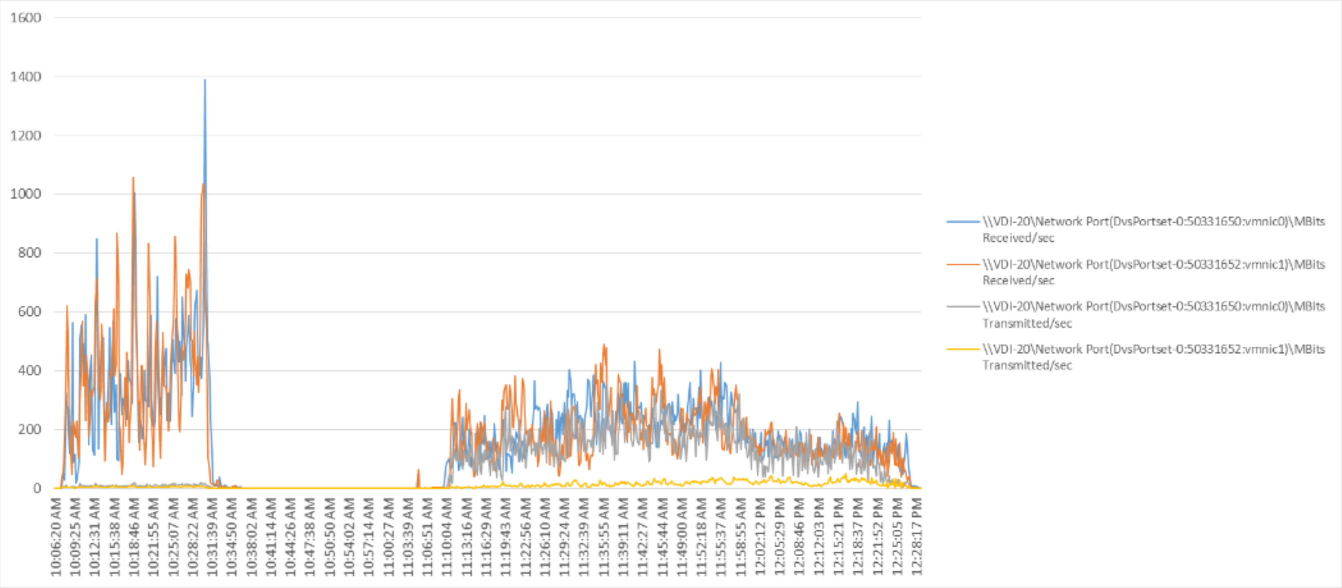

Figure 213 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

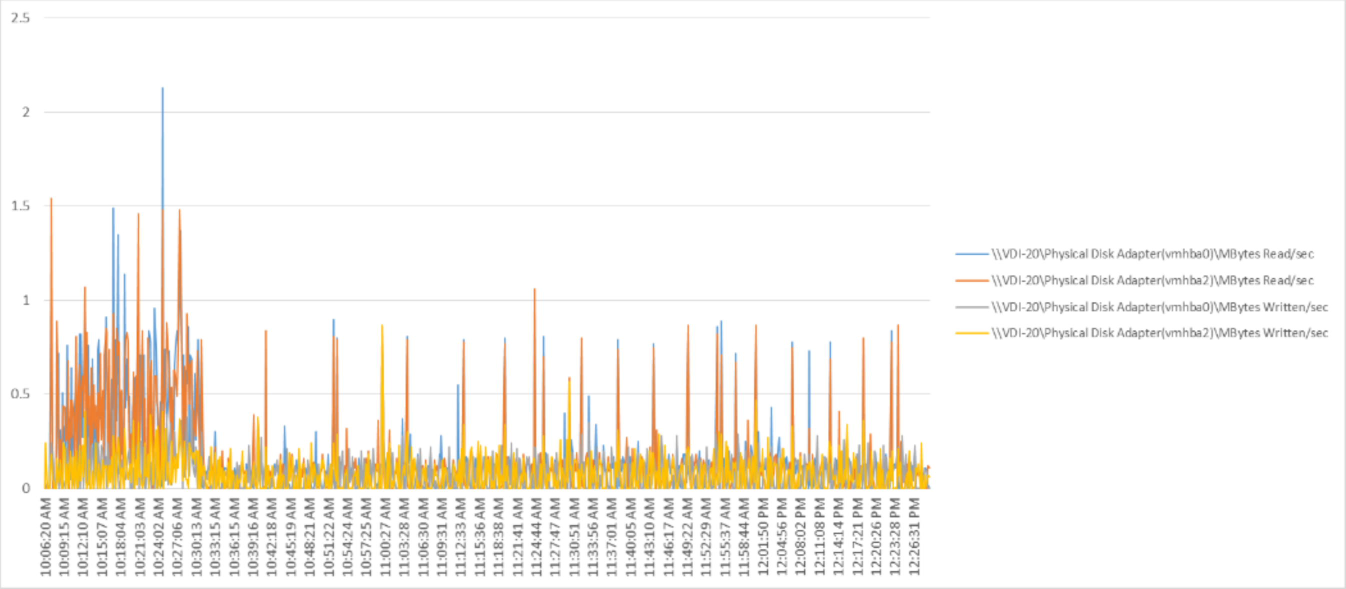

Figure 214 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

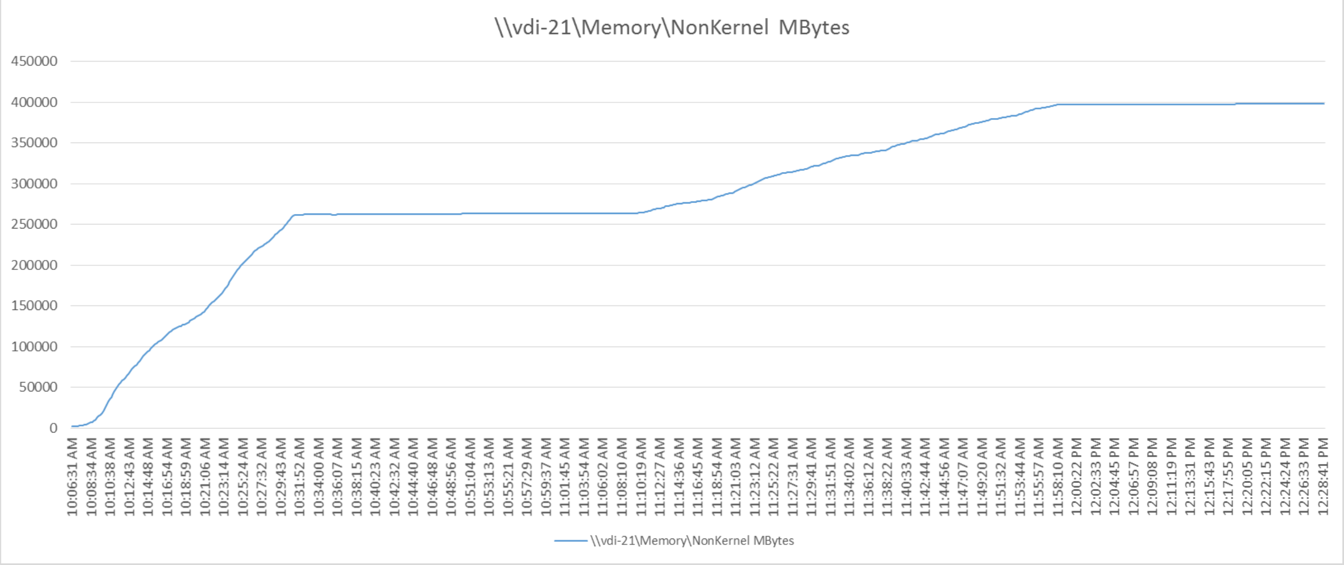

Figure 215 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

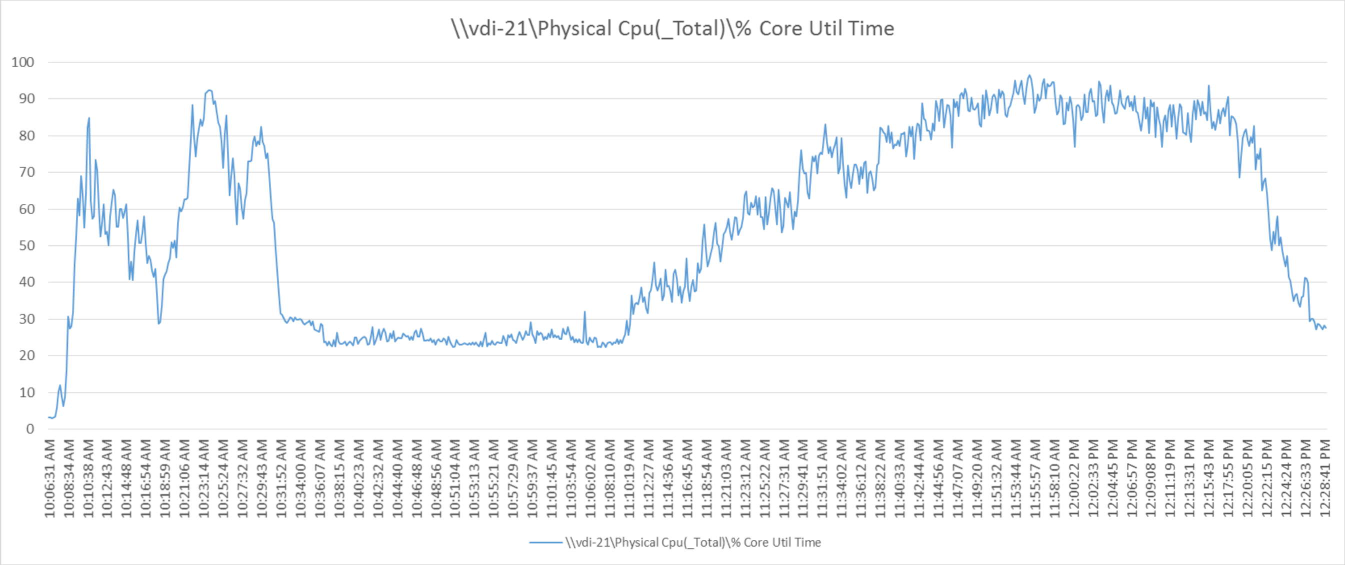

Figure 216 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

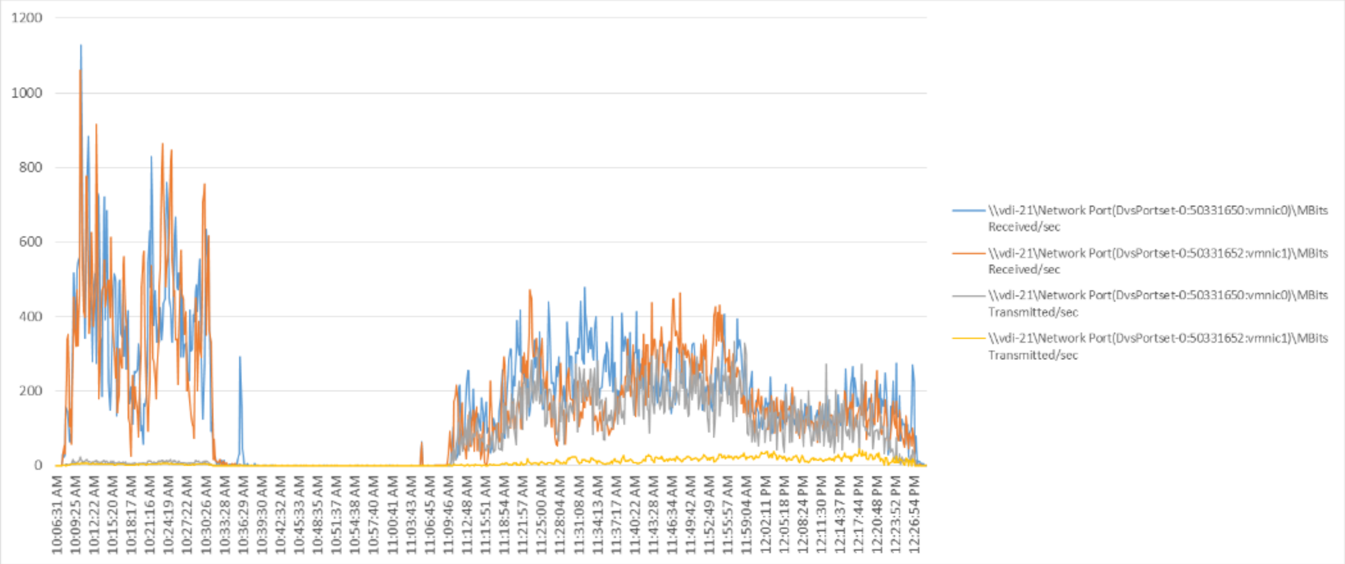

Figure 217 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

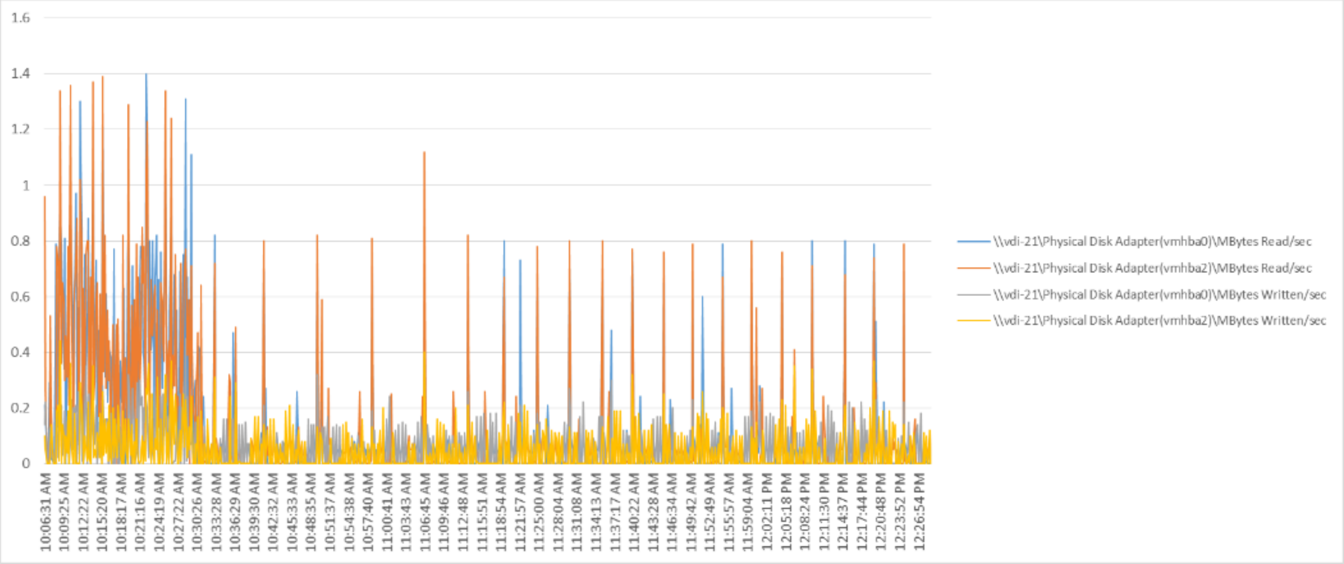

Figure 218 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

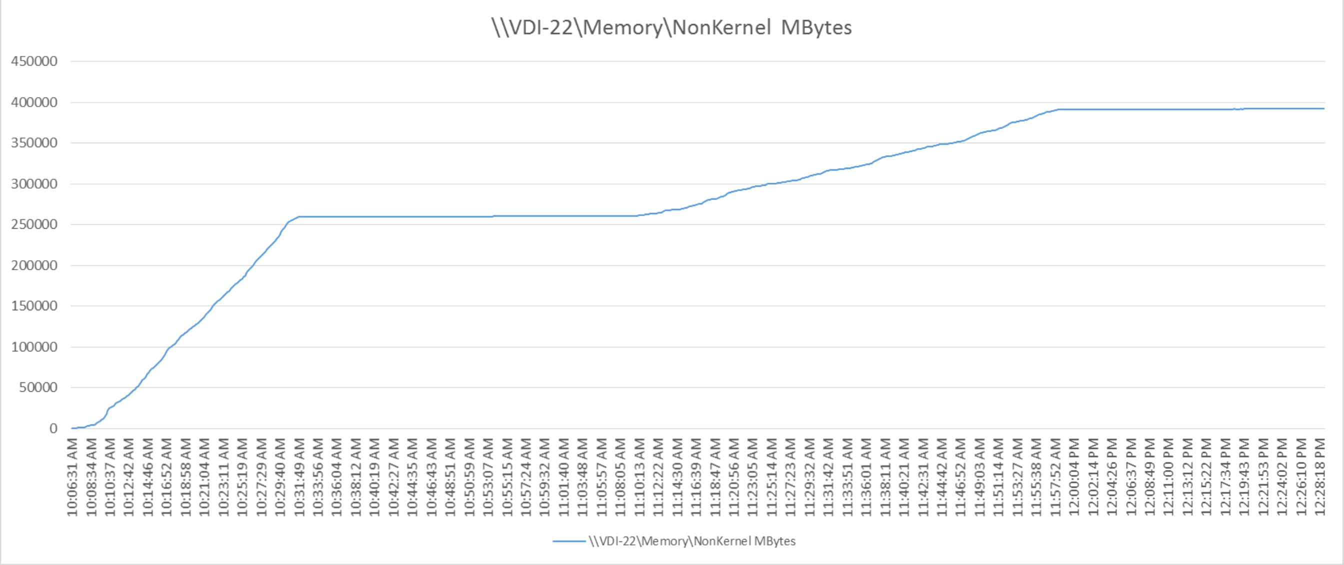

Figure 219 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

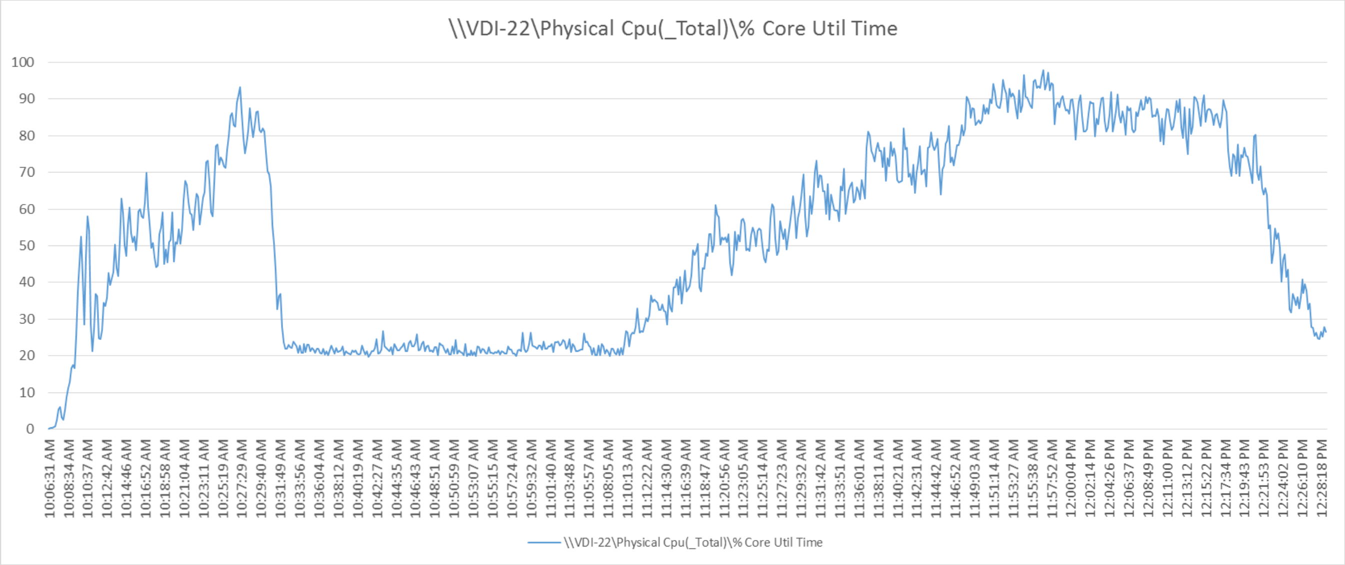

Figure 220 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

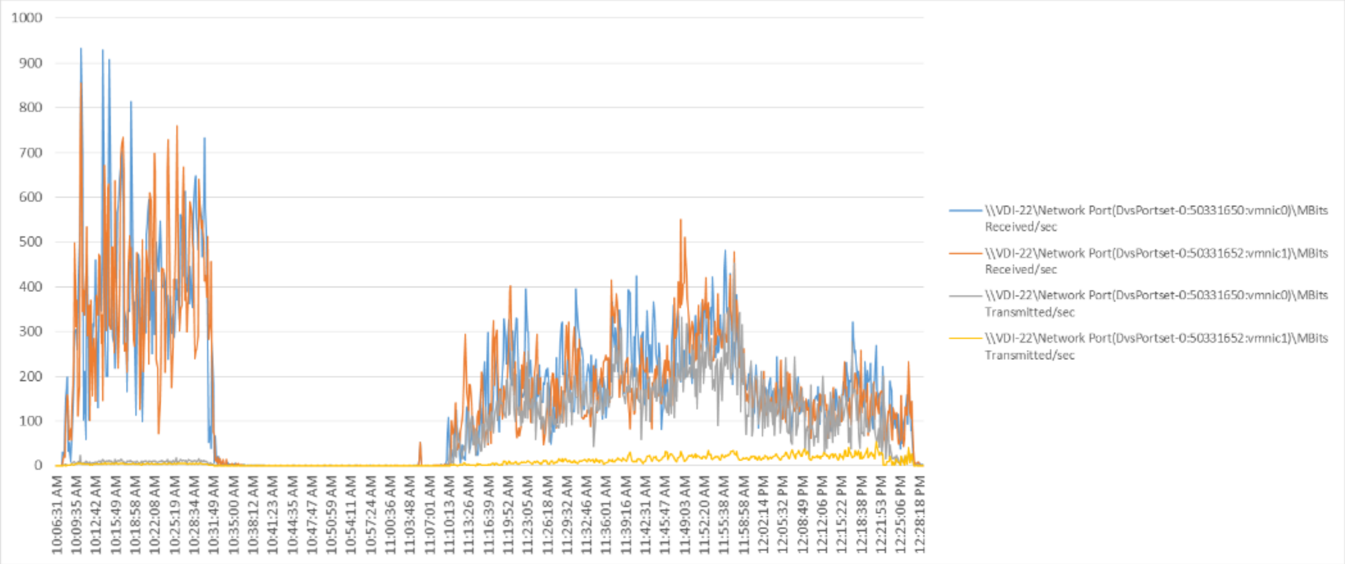

Figure 221 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

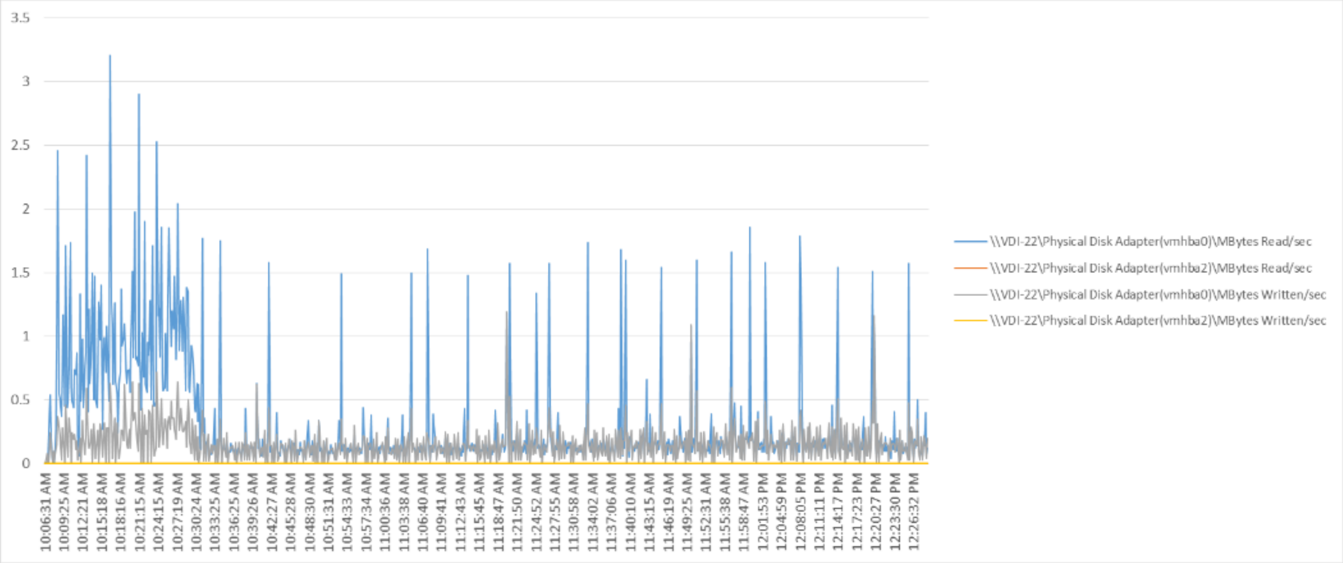

Figure 222 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

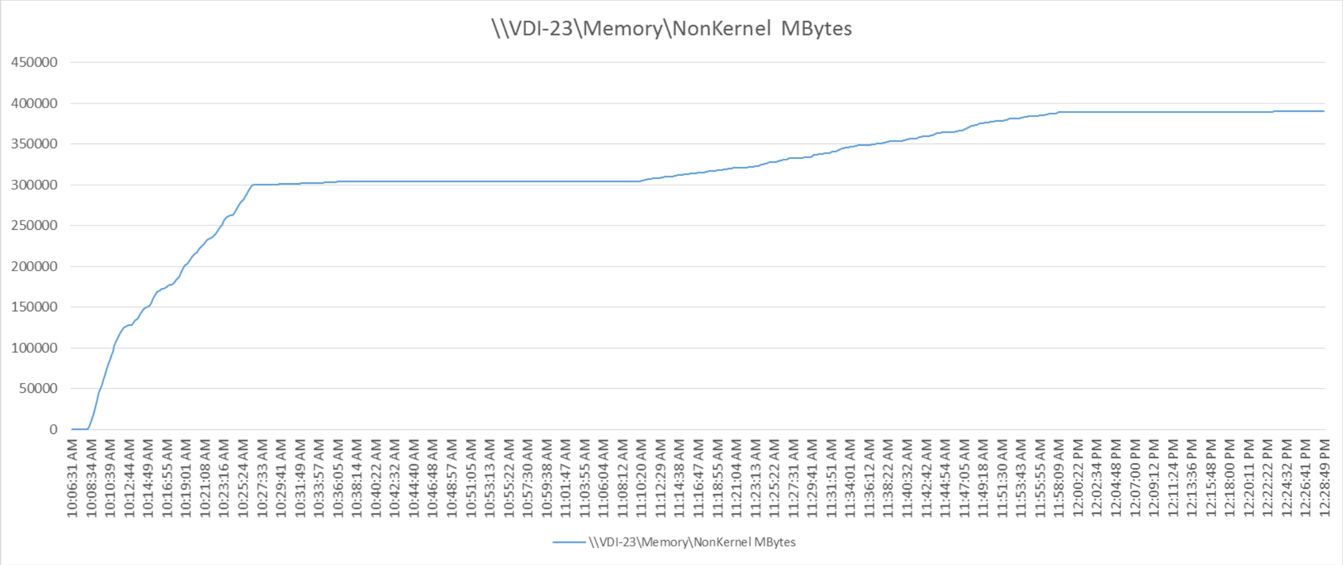

Figure 223 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

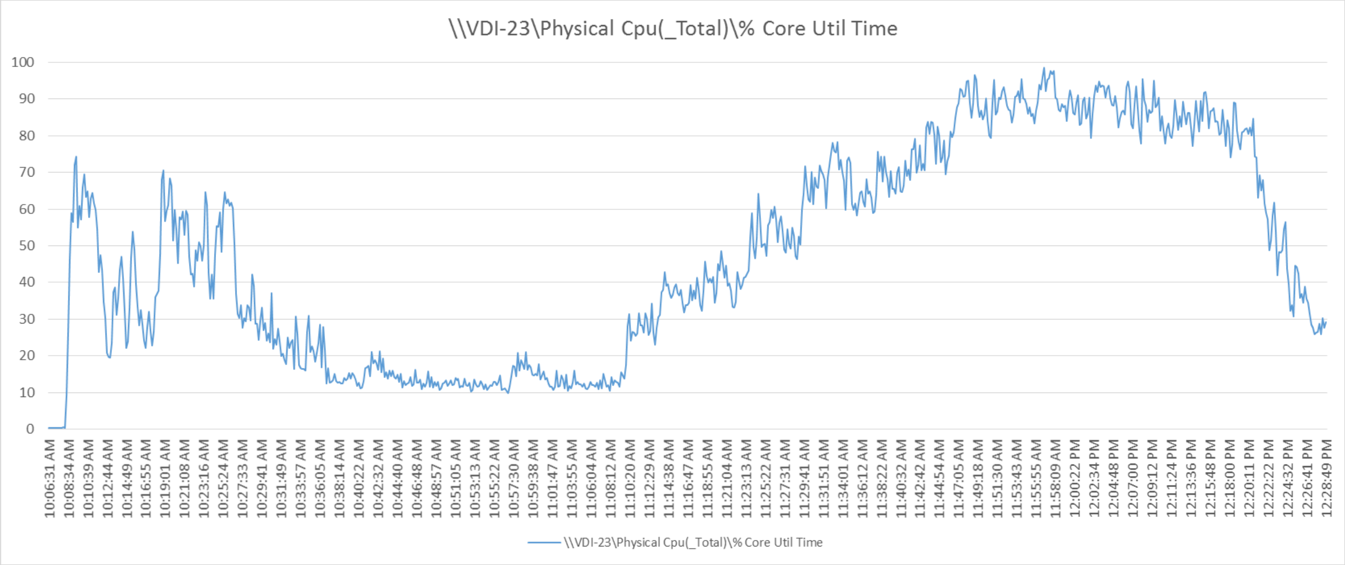

Figure 224 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

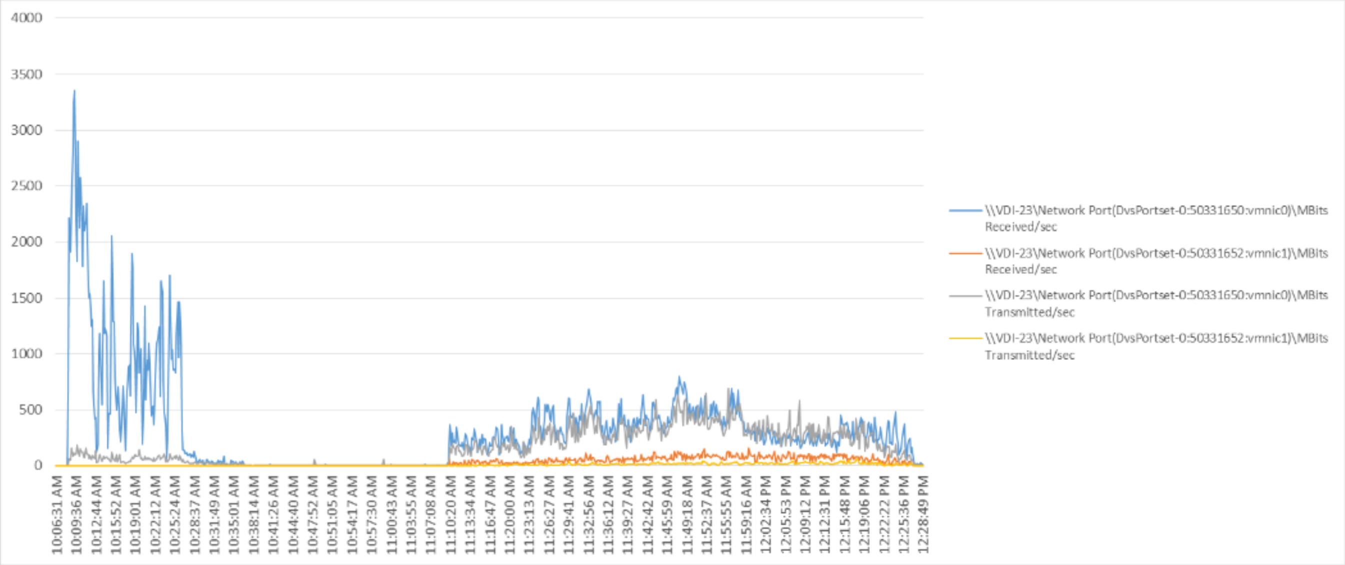

Figure 225 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

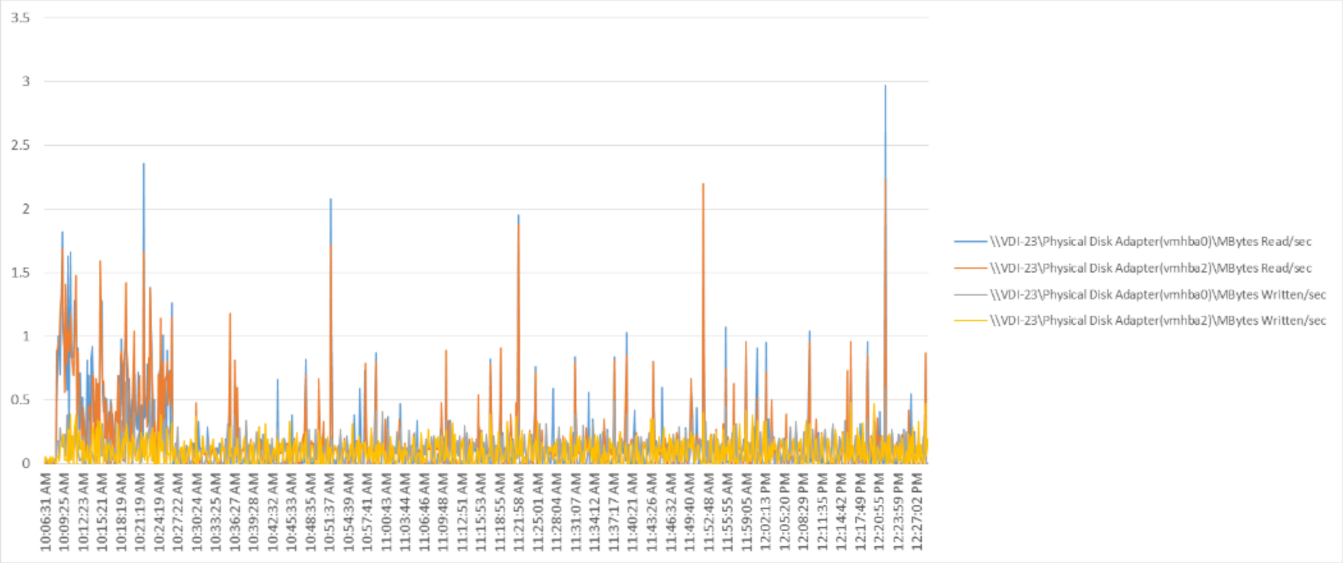

Figure 226 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Memory Utilization

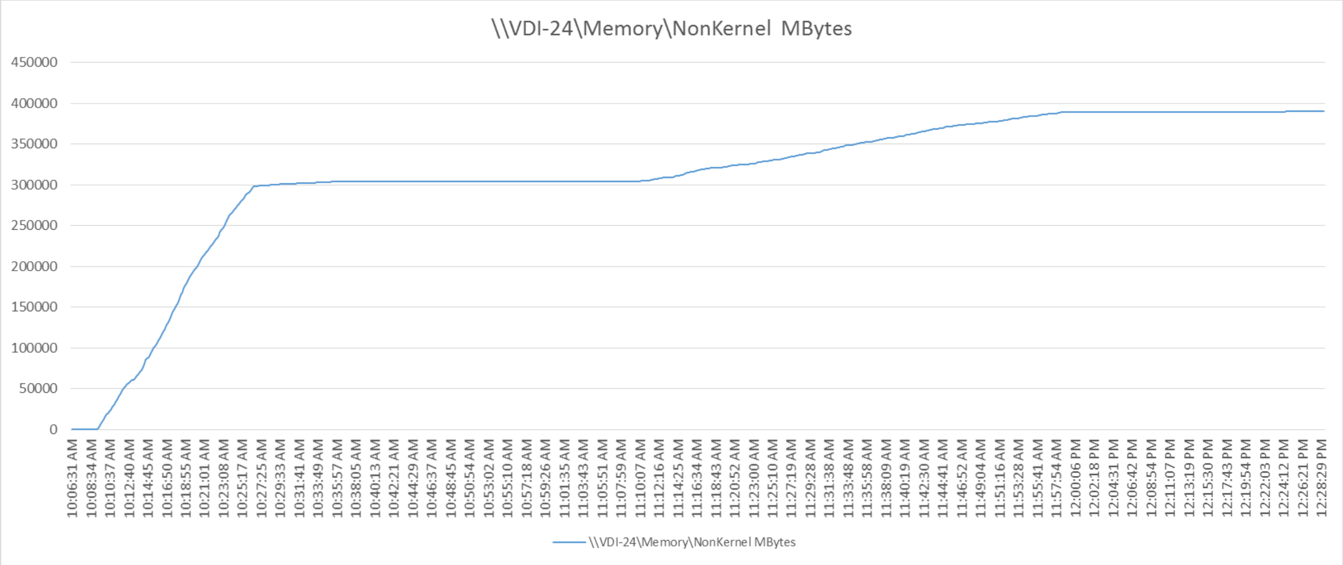

Figure 227 Full Scale | 6000 Mixed Users | Persistent Hosts | Host CPU Utilization

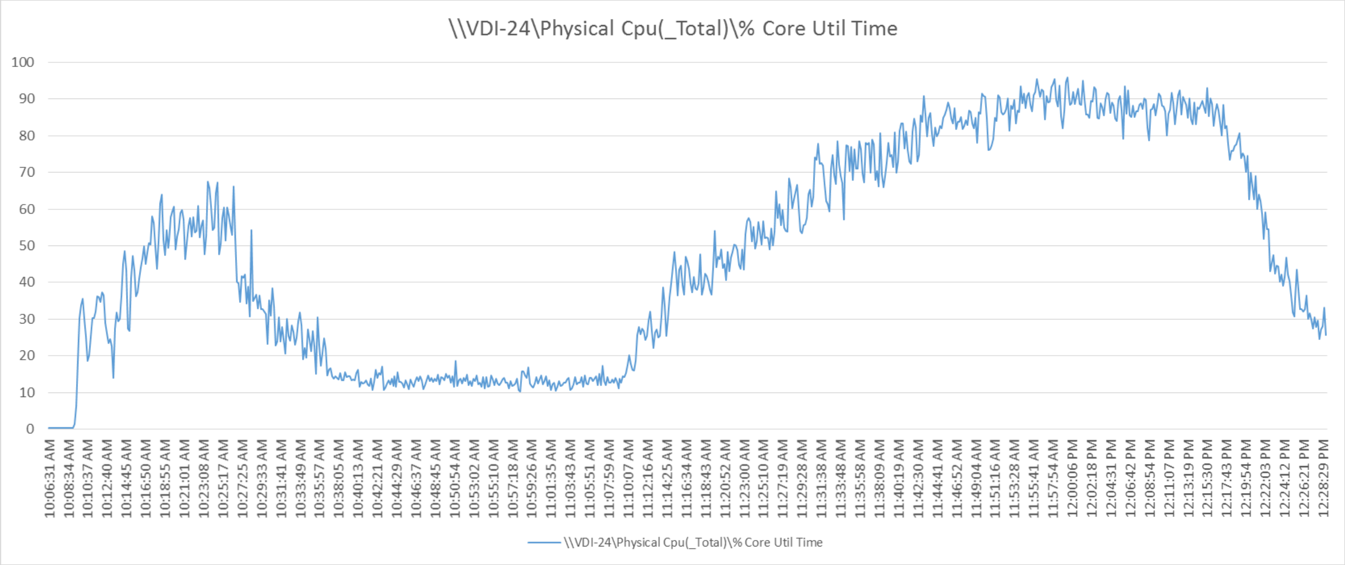

Figure 228 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Network Utilization

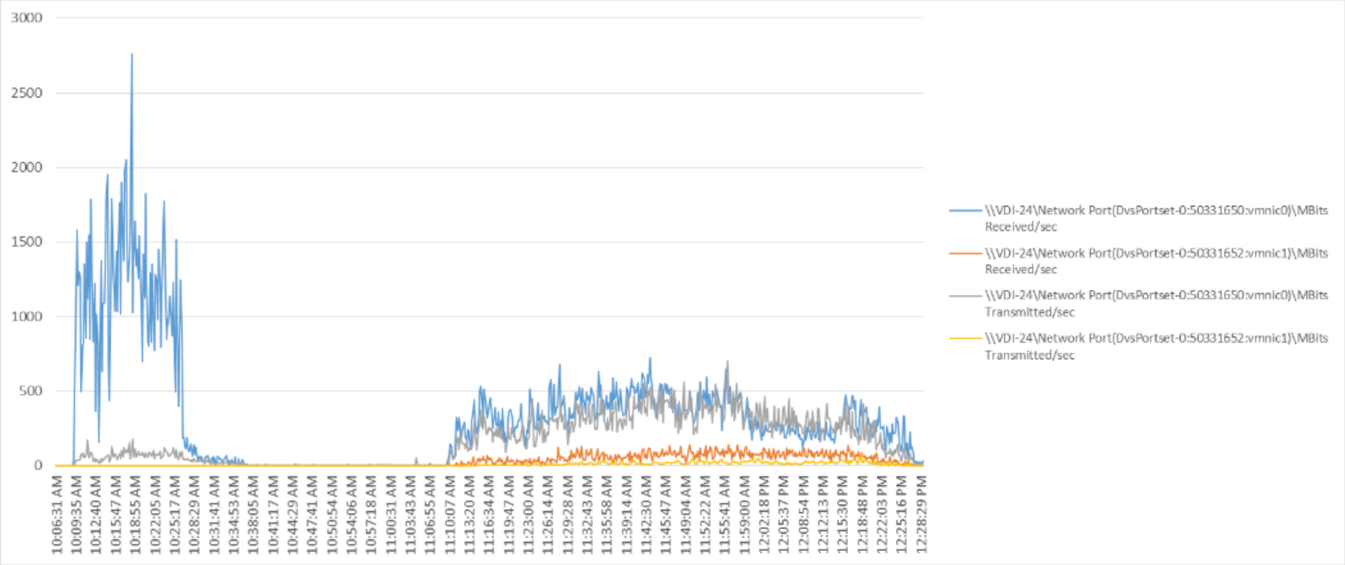

Figure 229 Full Scale | 6000 Mixed Users | Persistent Hosts | Host Fibre Channel Utilization

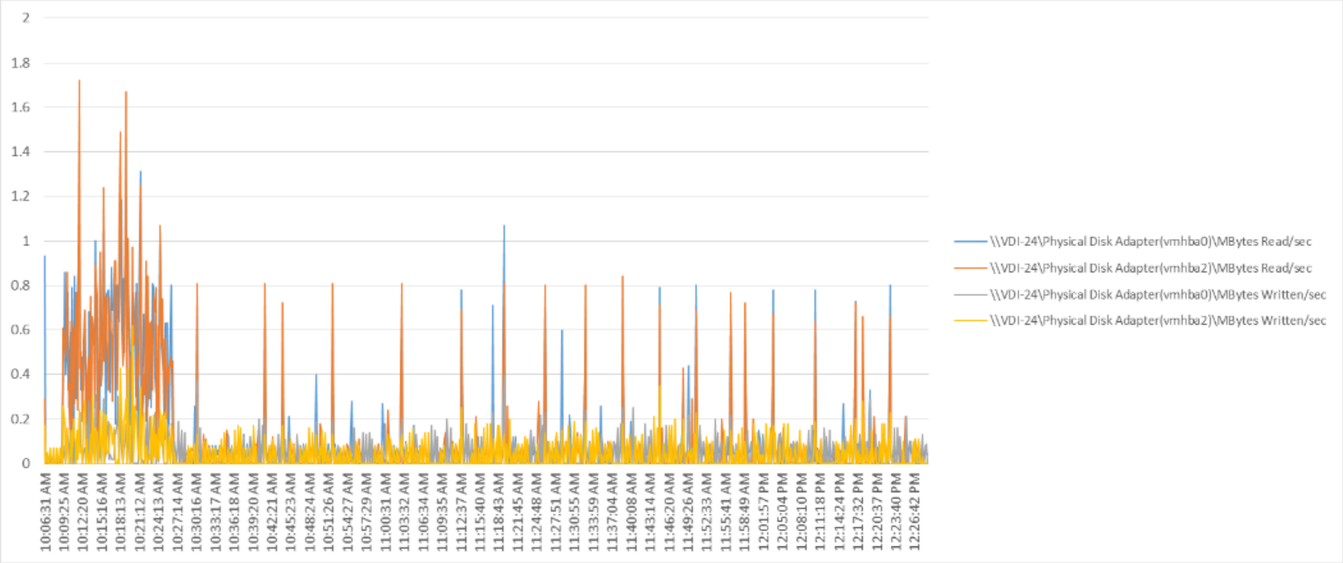

Figure 230 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

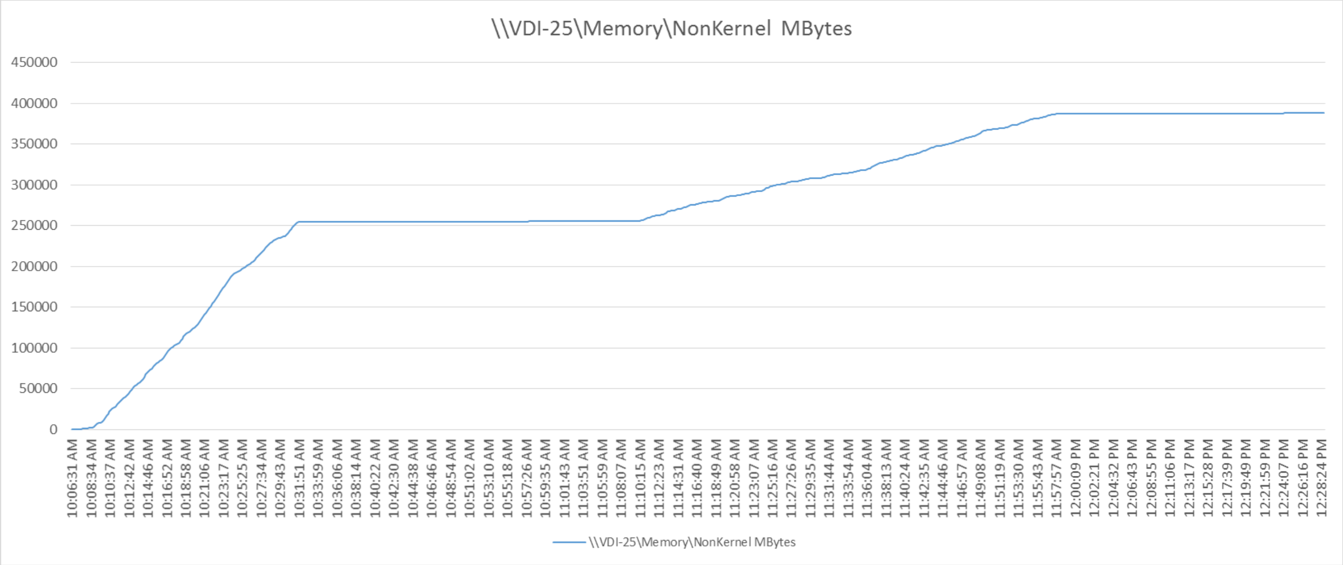

Figure 231 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

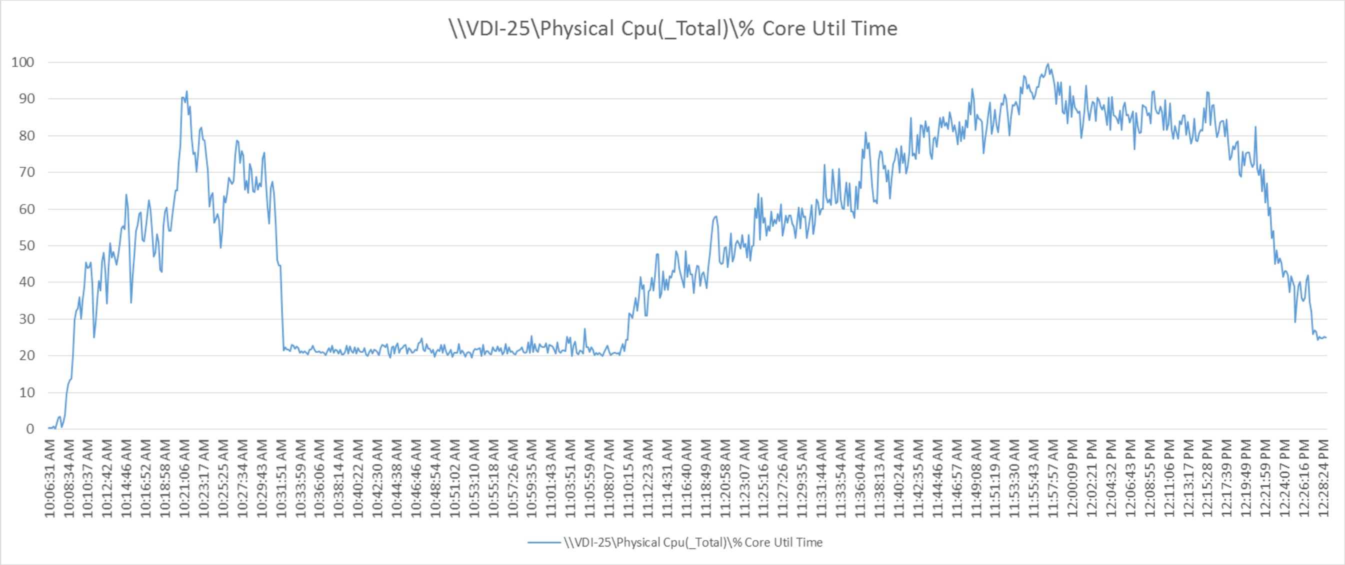

Figure 232 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

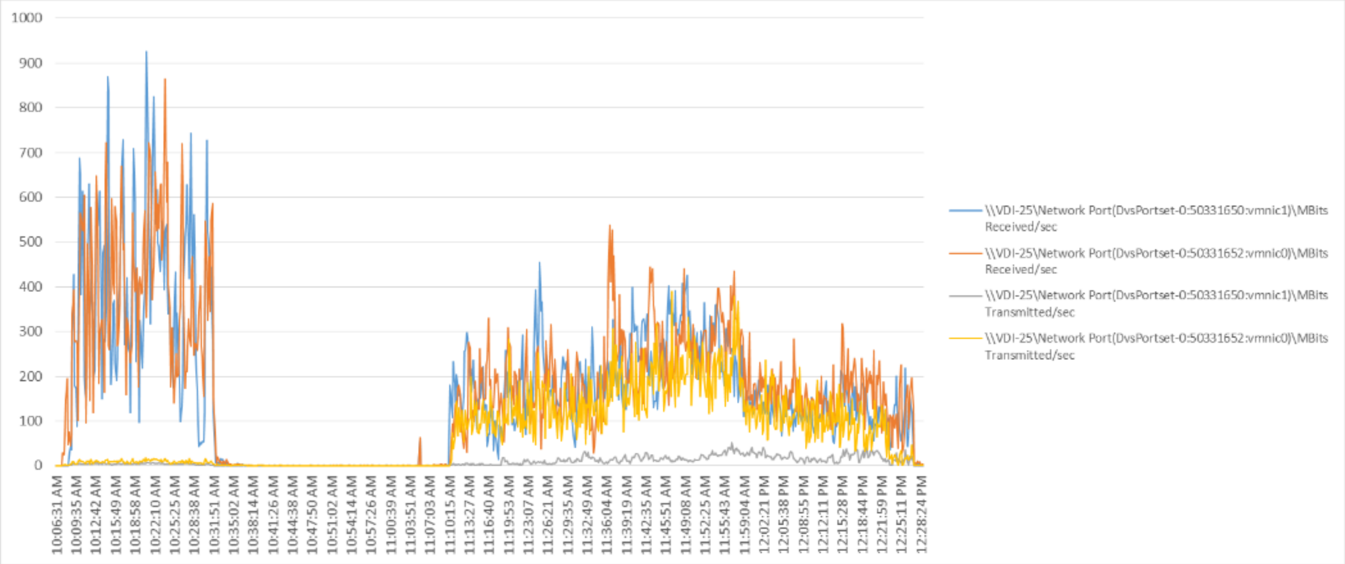

Figure 233 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

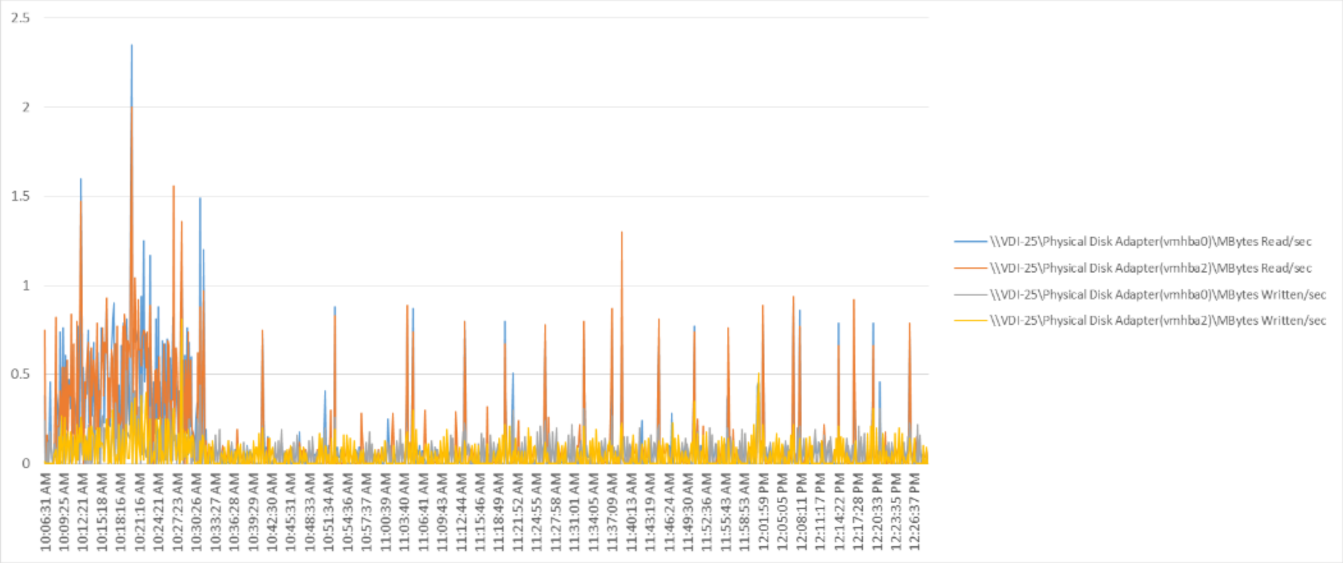

Figure 234 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

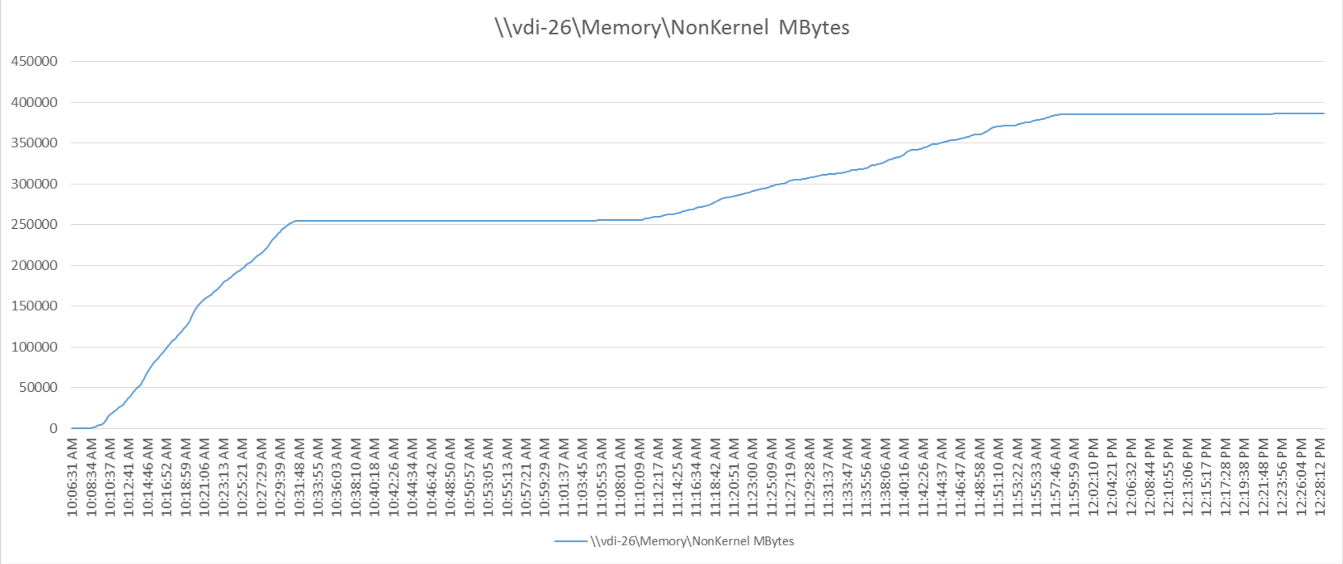

Figure 235 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

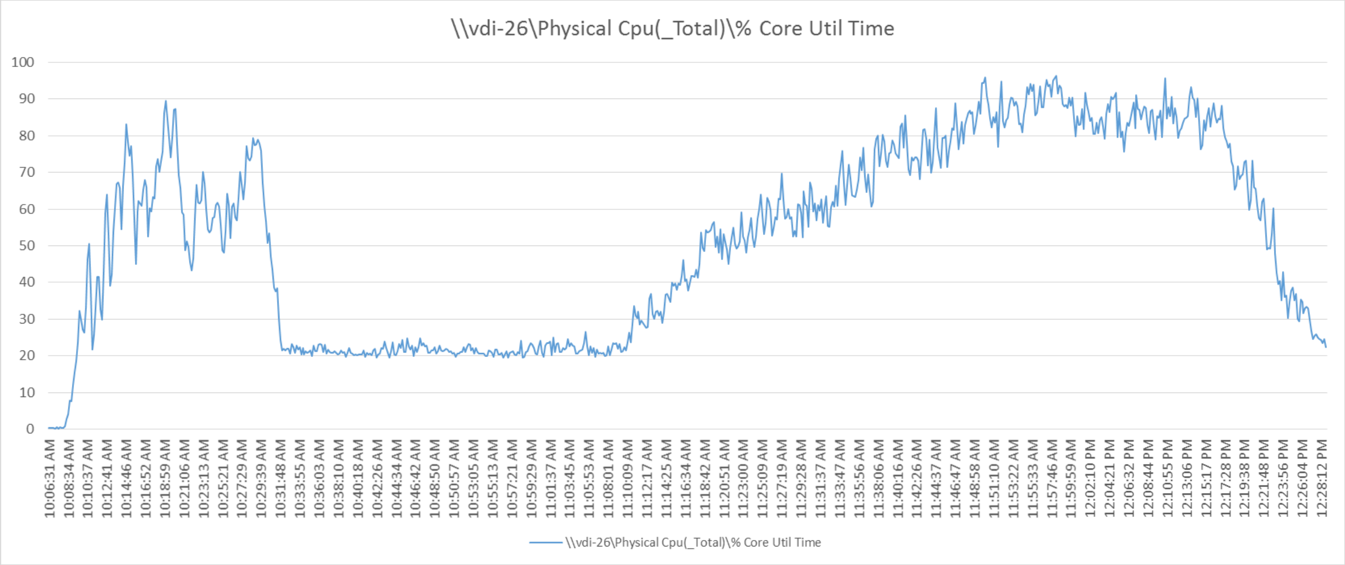

Figure 236 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

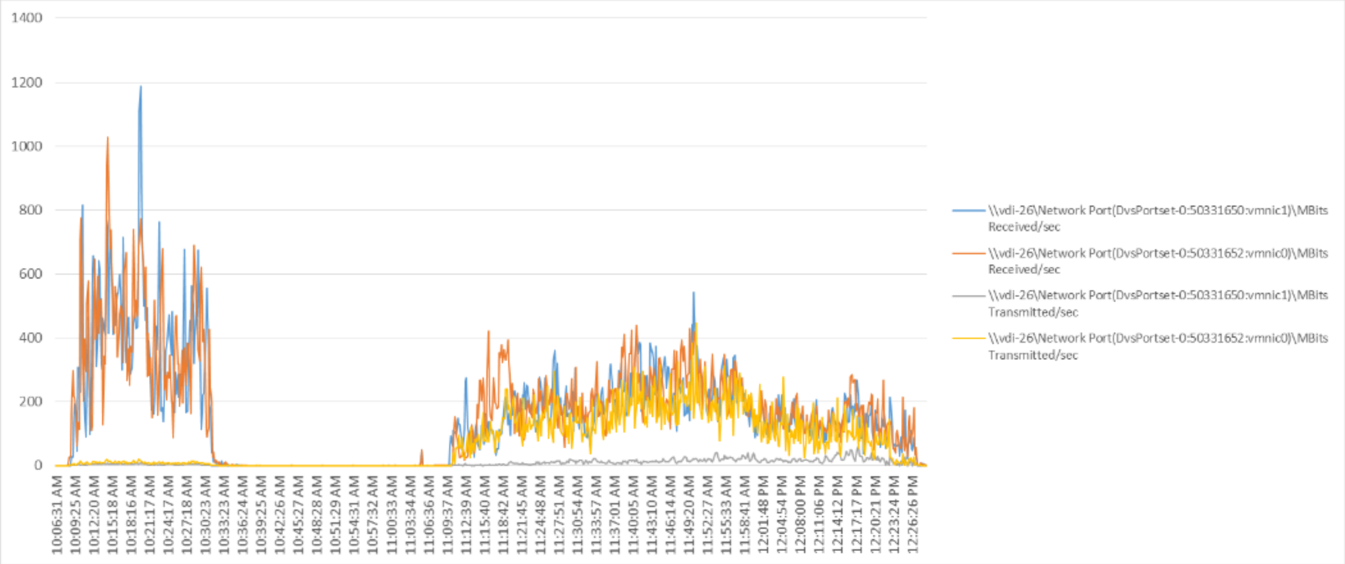

Figure 237 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

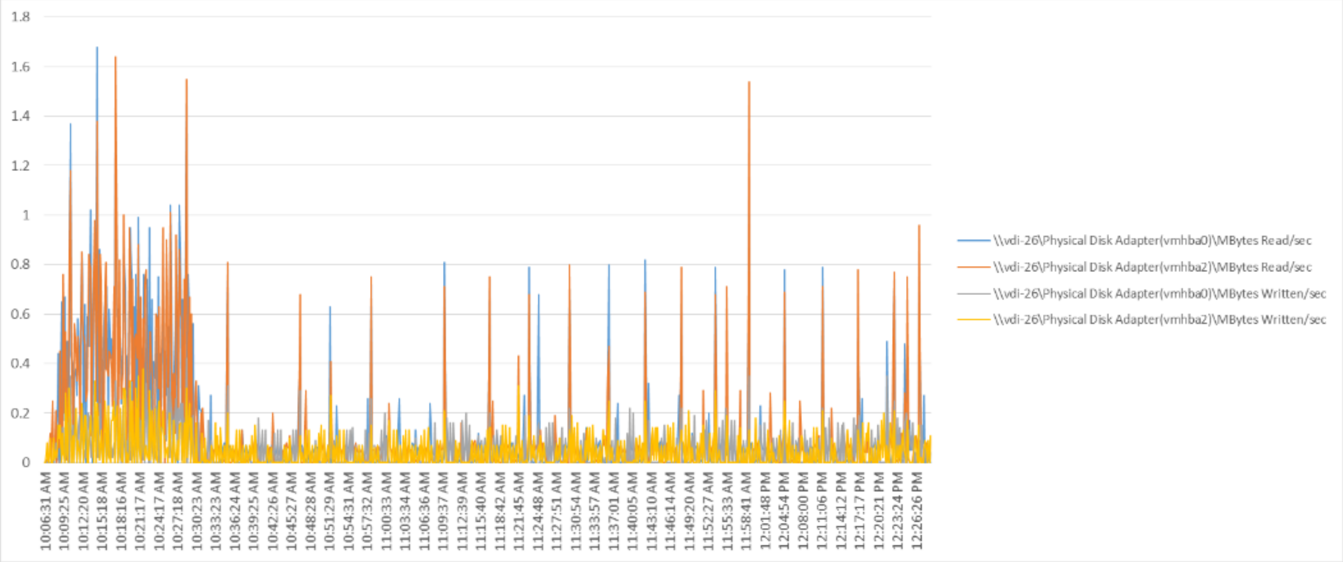

Figure 238 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

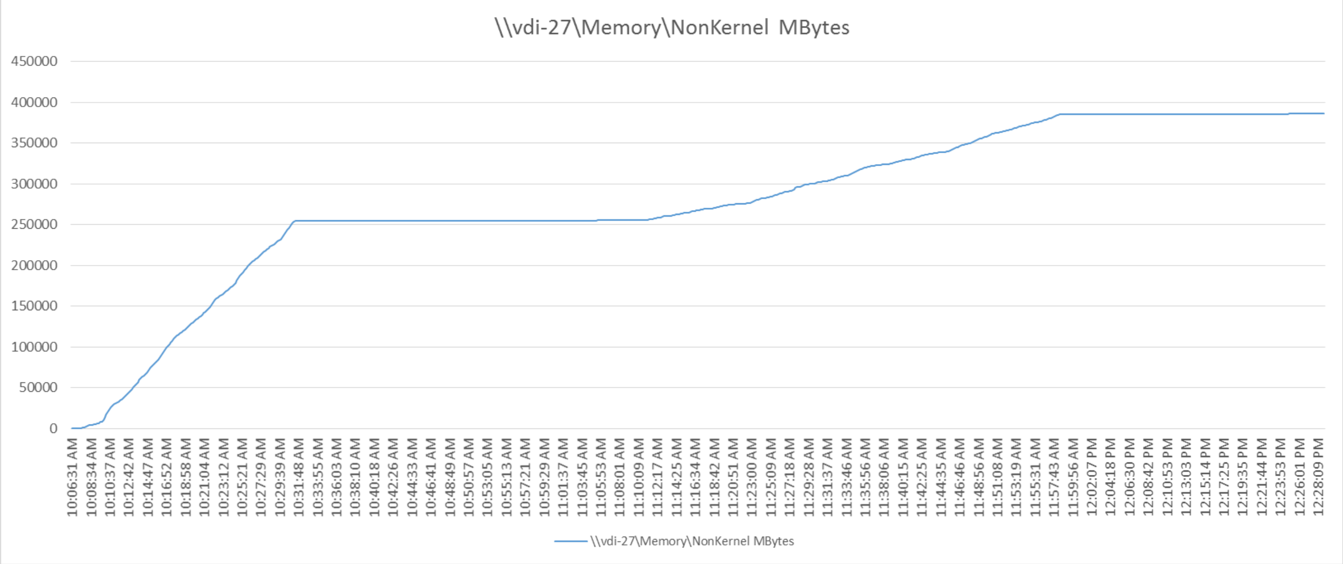

Figure 239 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

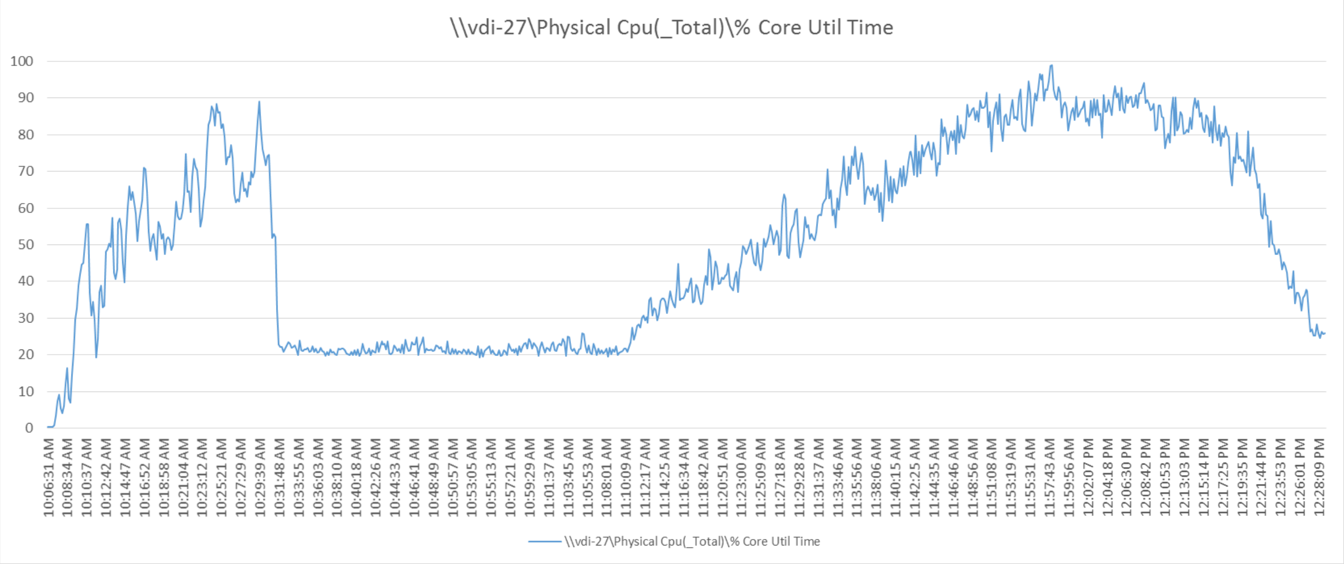

Figure 240 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

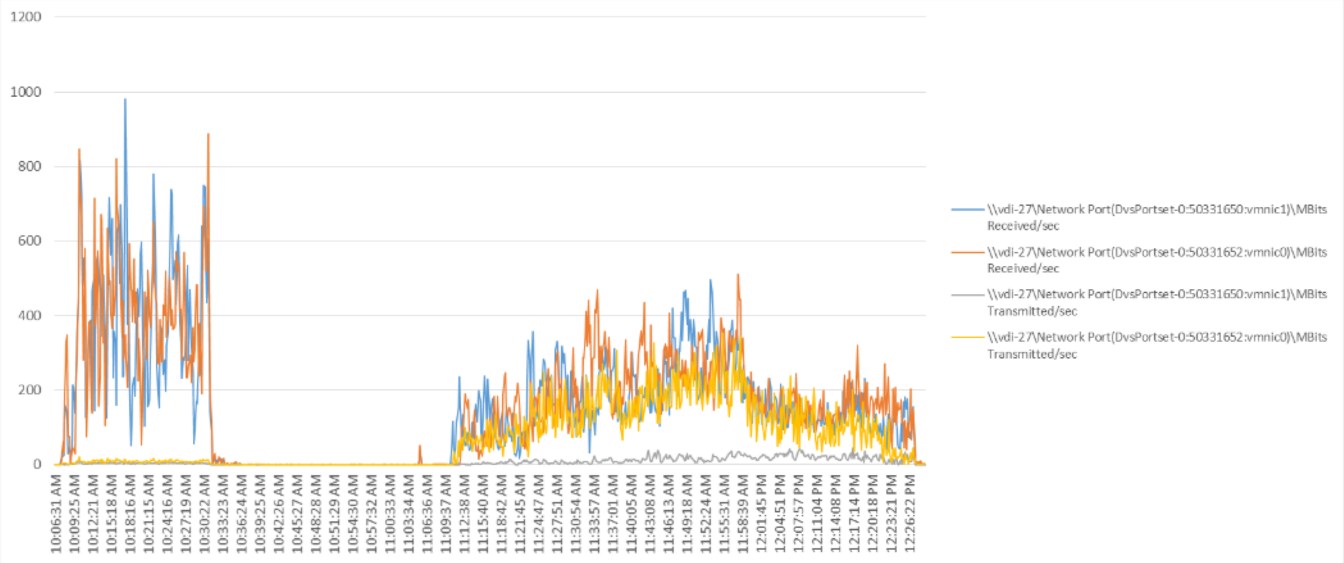

Figure 241 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

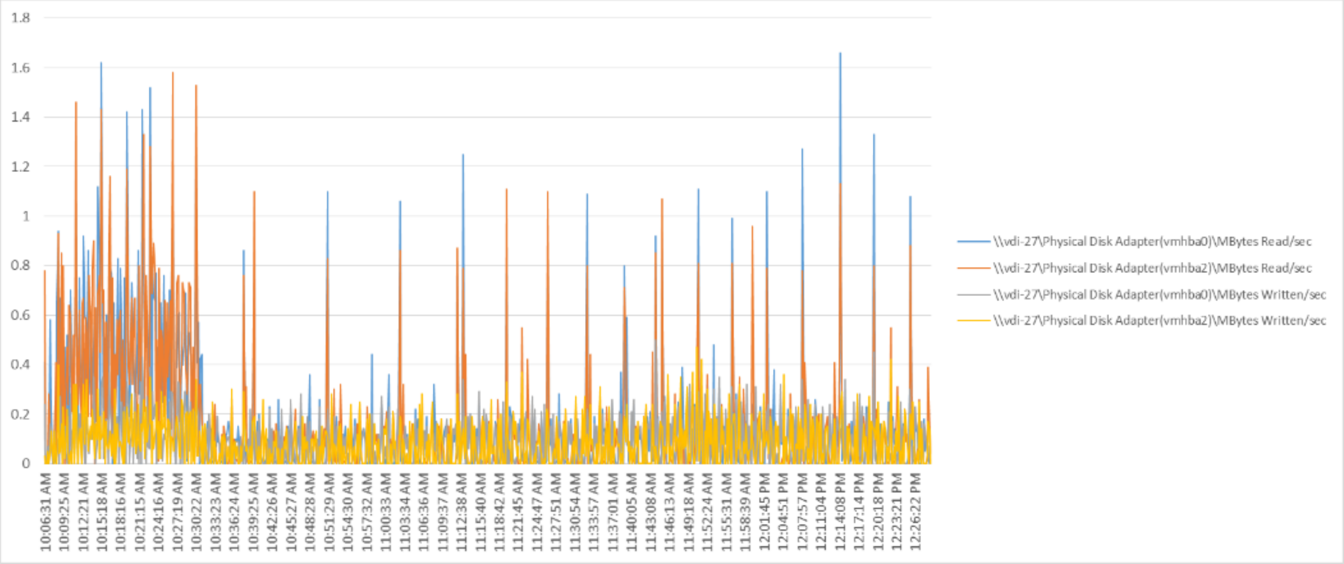

Figure 242 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

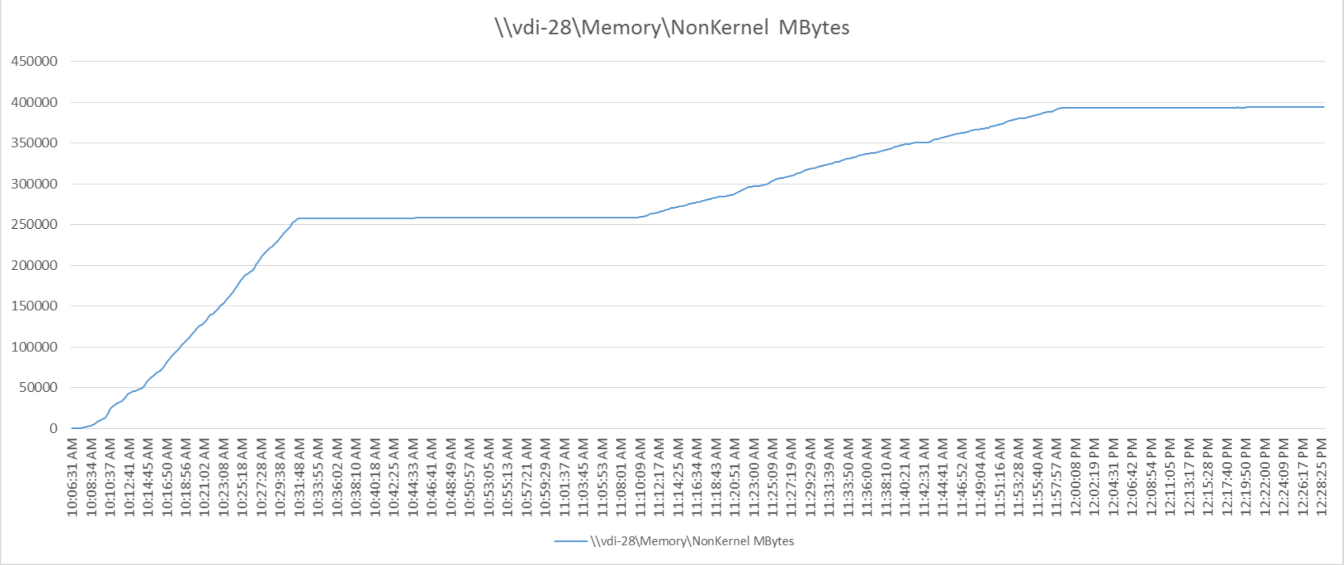

Figure 243 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

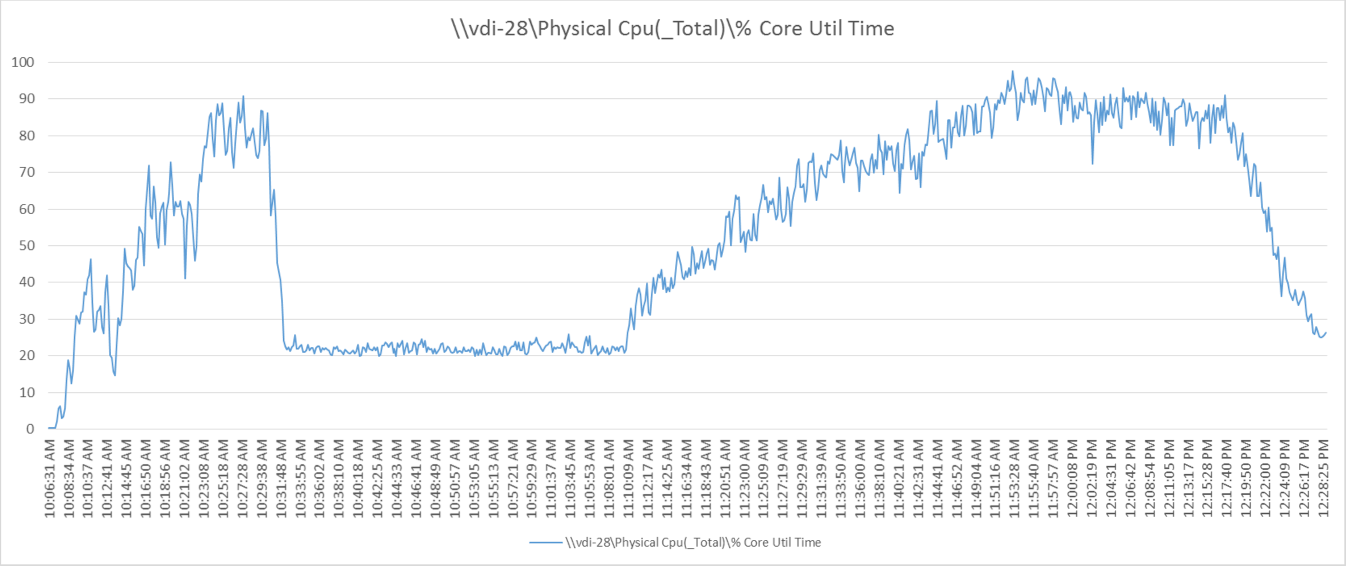

Figure 244 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

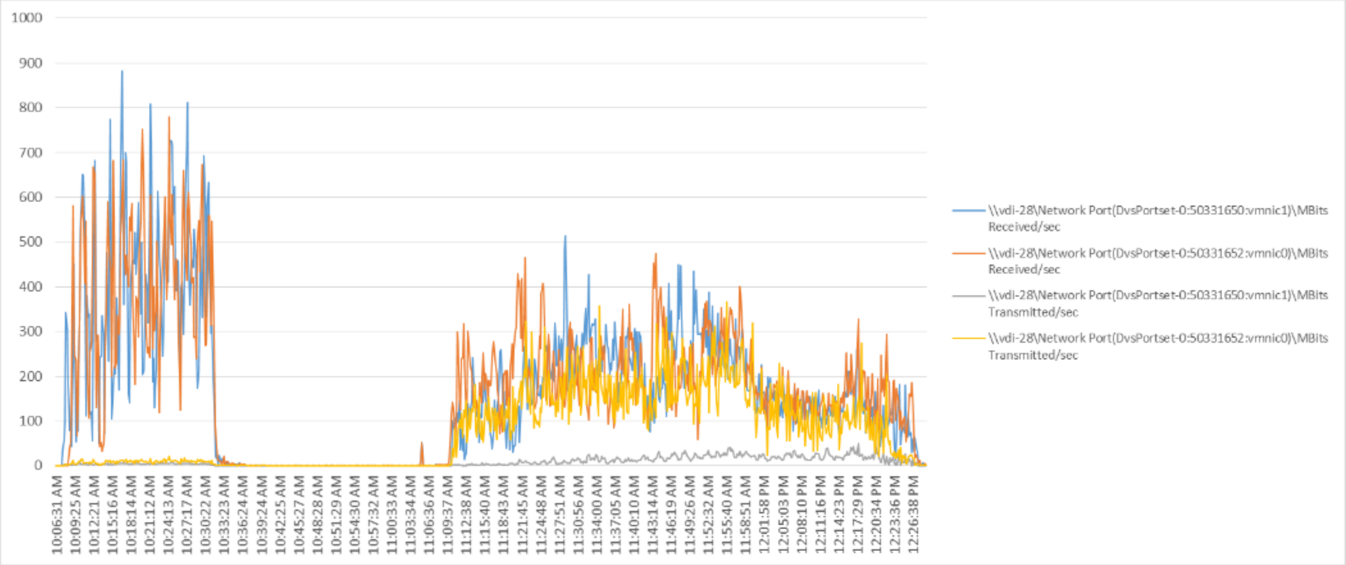

Figure 245 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

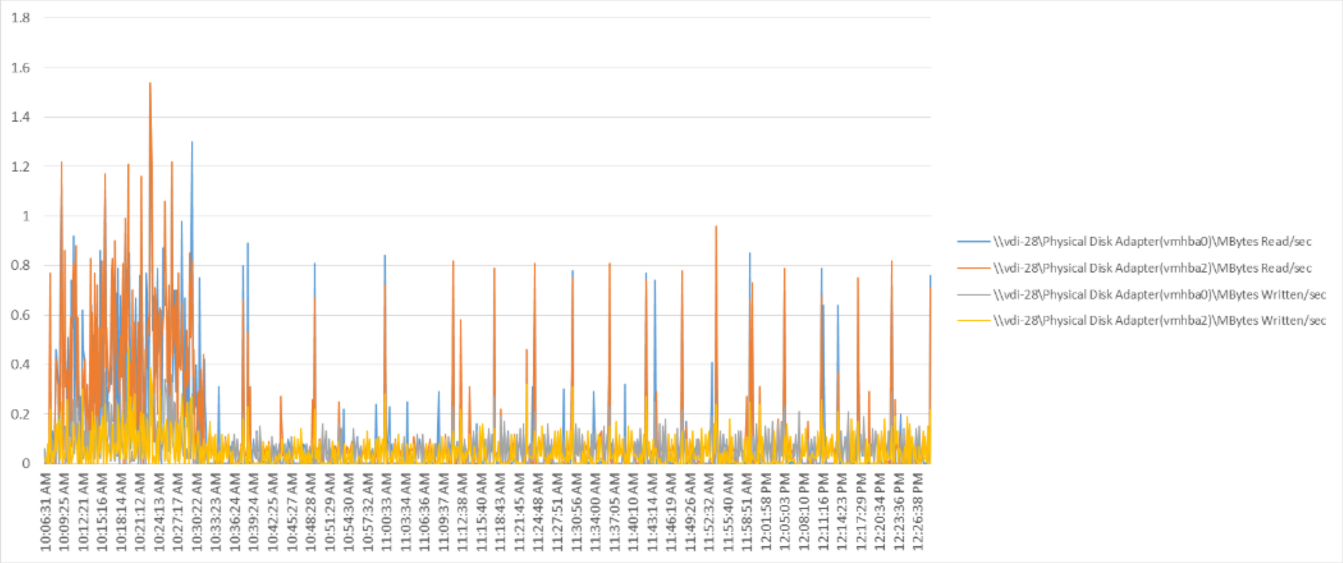

Figure 246 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

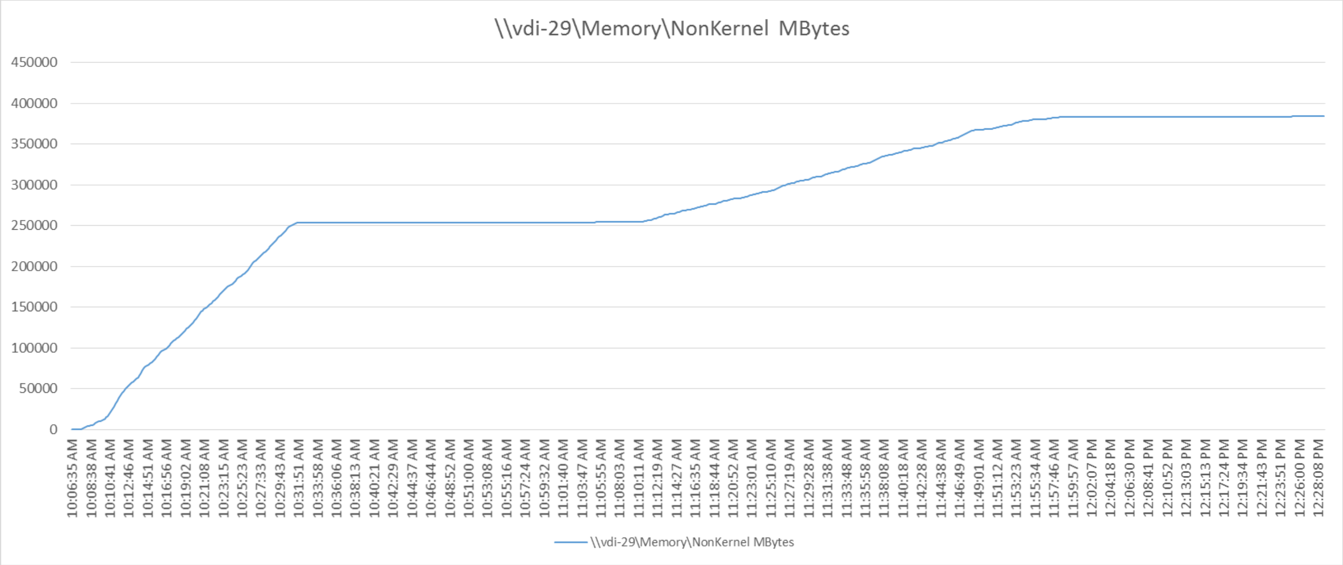

Figure 247 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

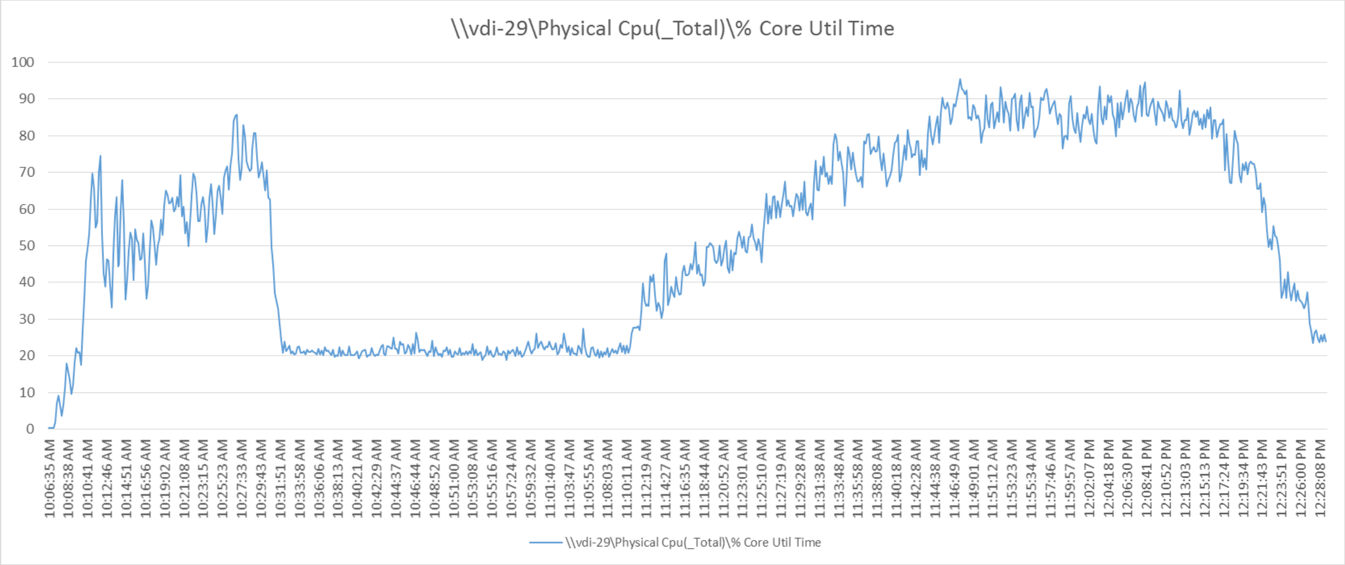

Figure 248 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

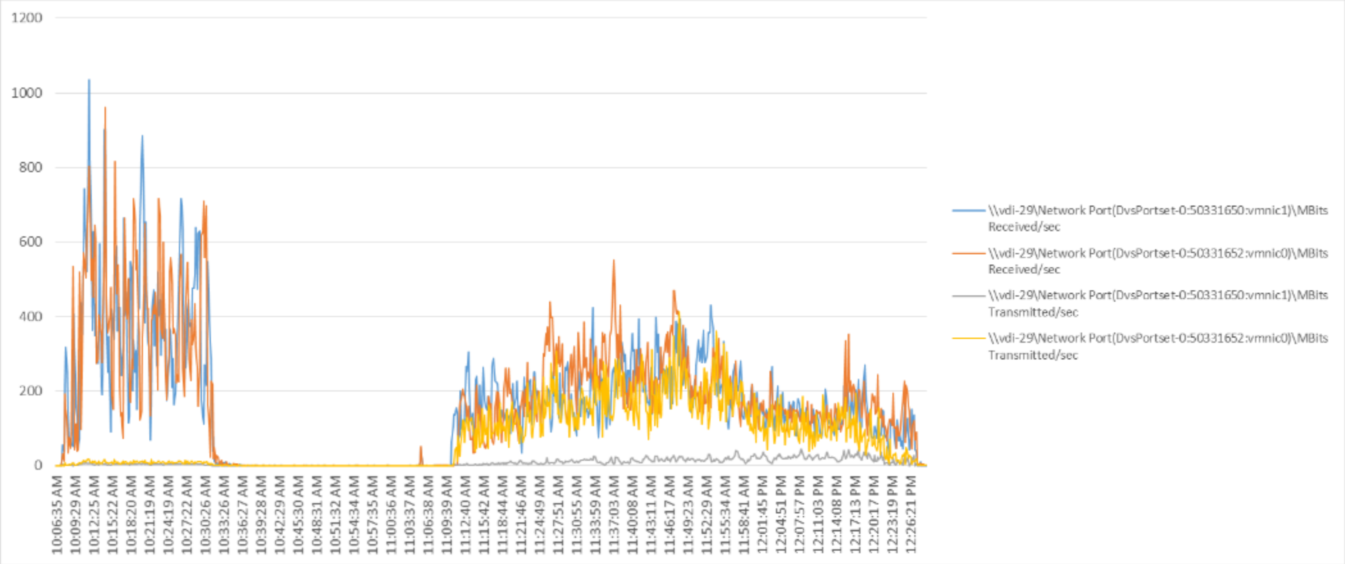

Figure 249 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

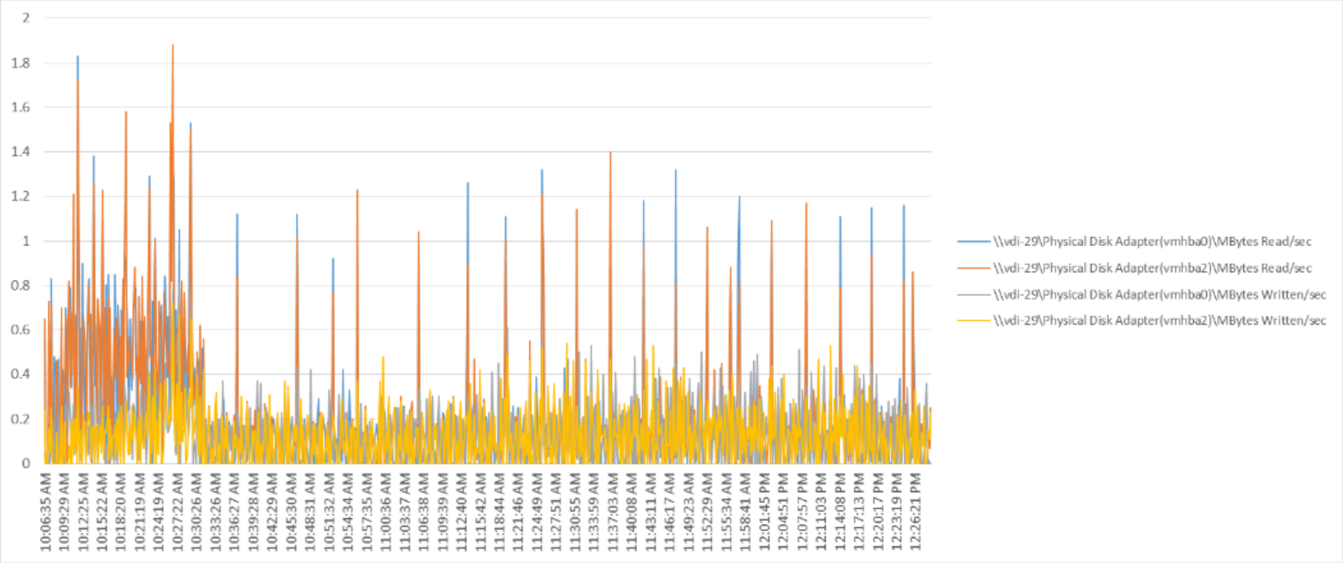

Figure 250 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

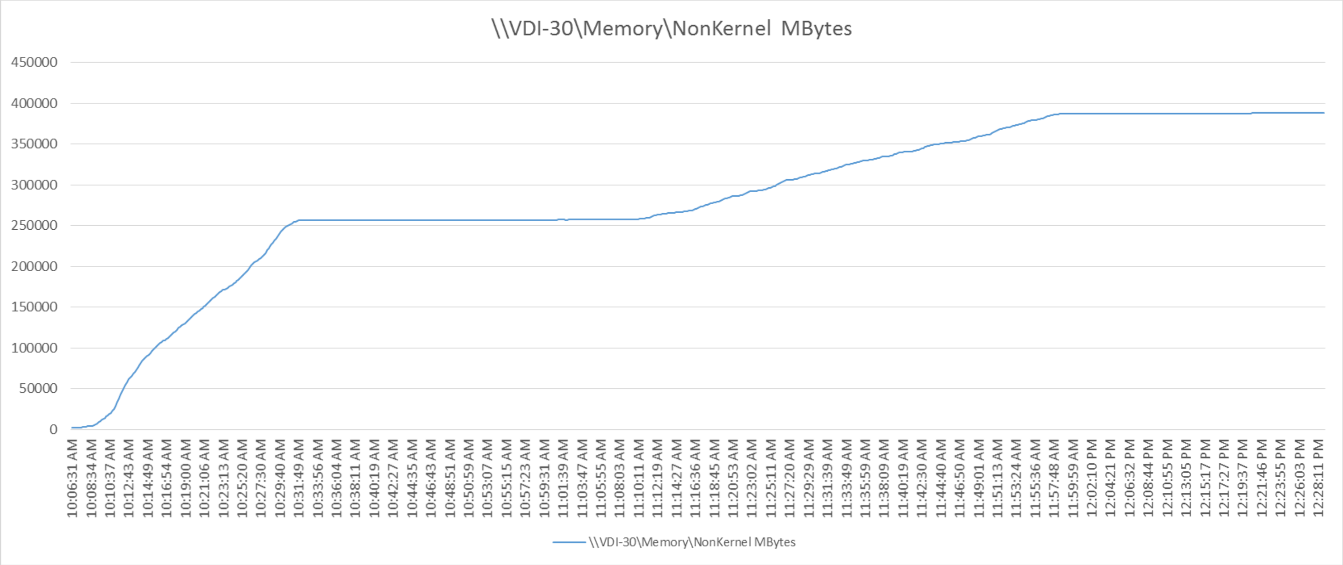

Figure 251 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

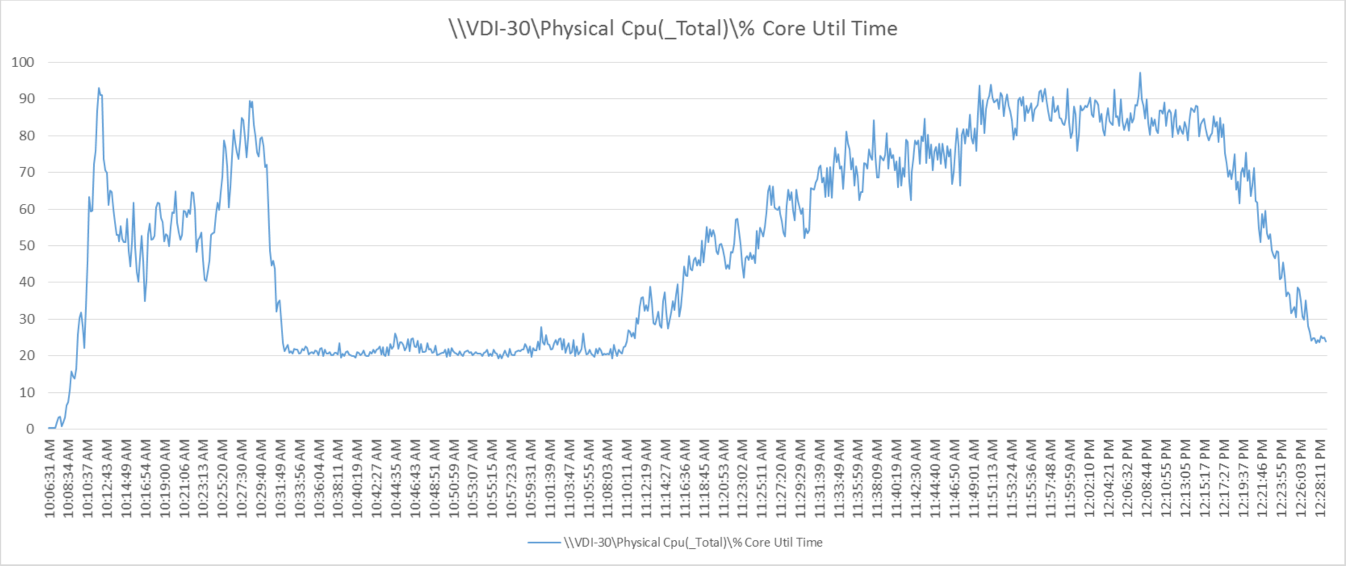

Figure 252 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

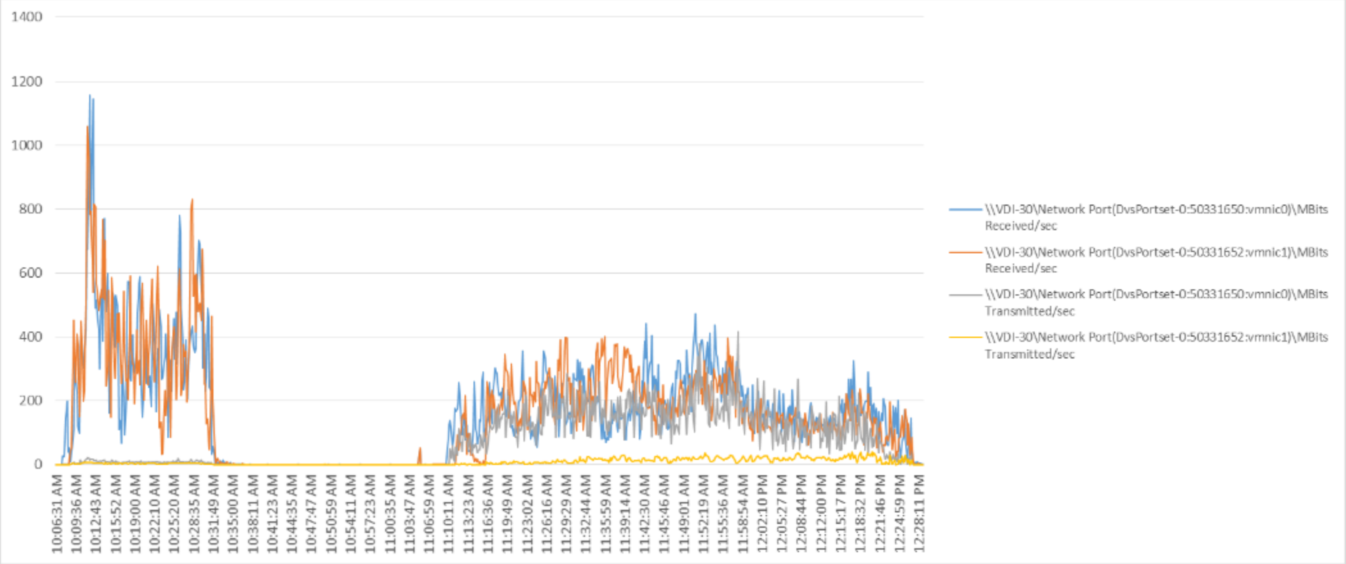

Figure 253 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

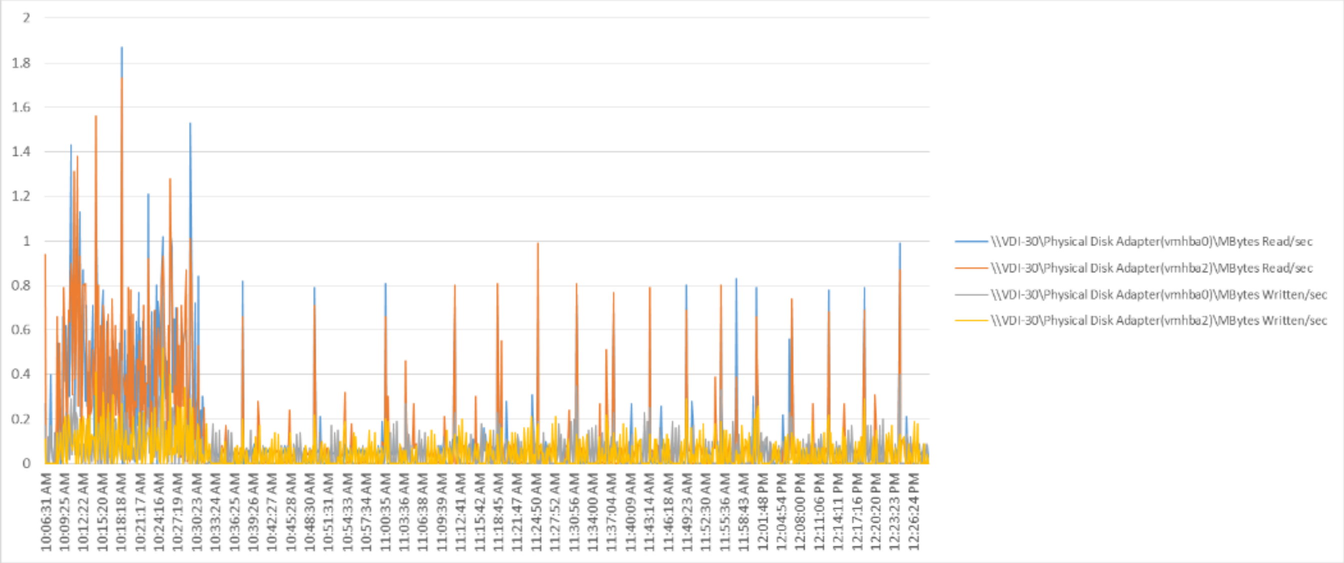

Figure 254 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

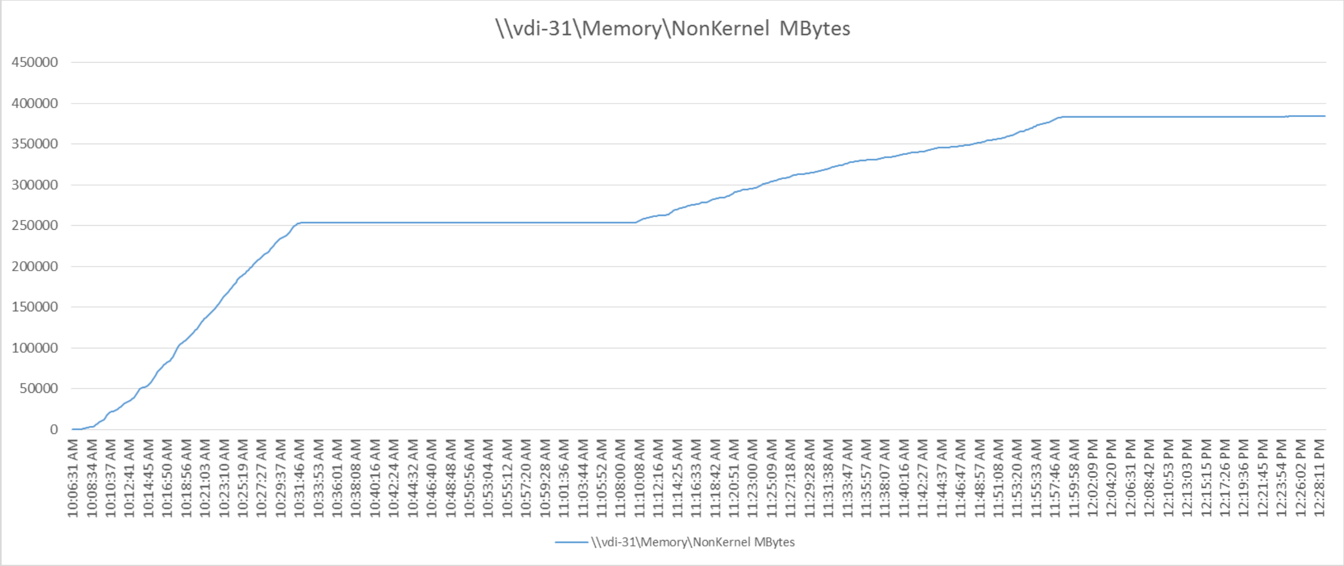

Figure 255 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

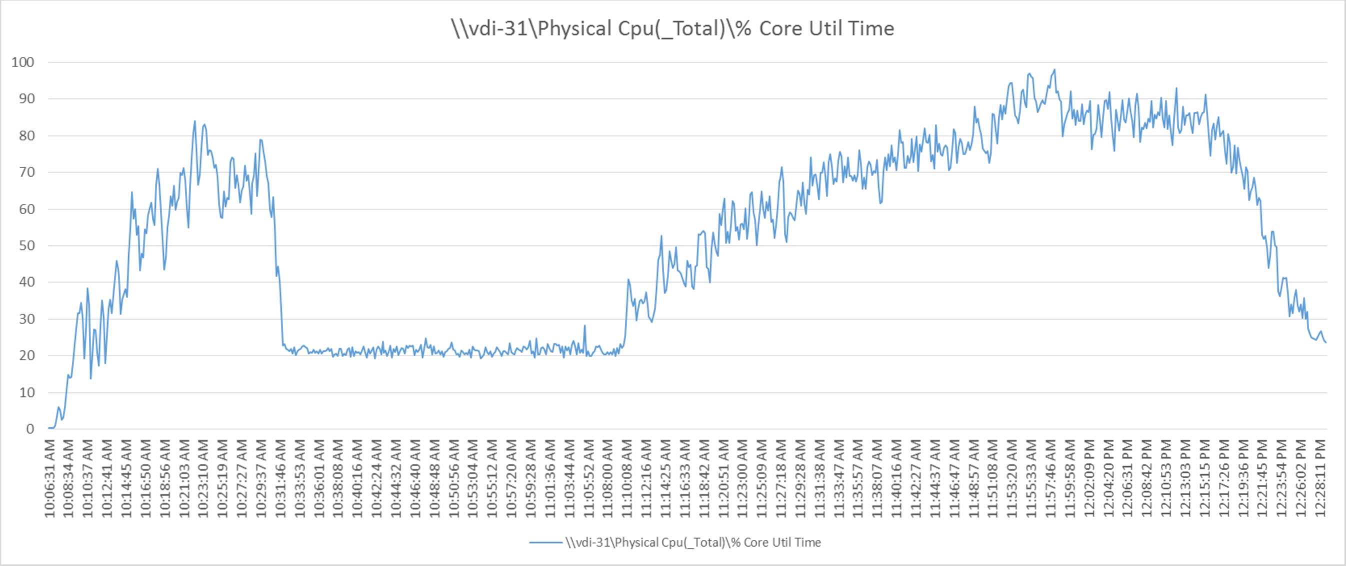

Figure 256 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

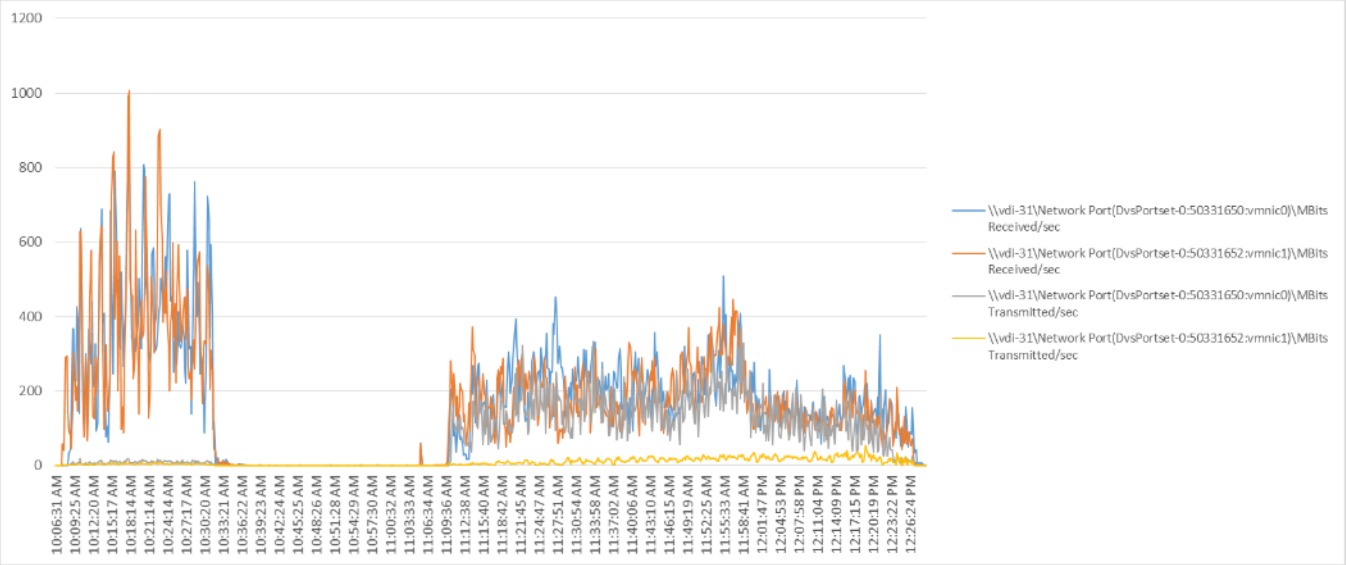

Figure 257 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

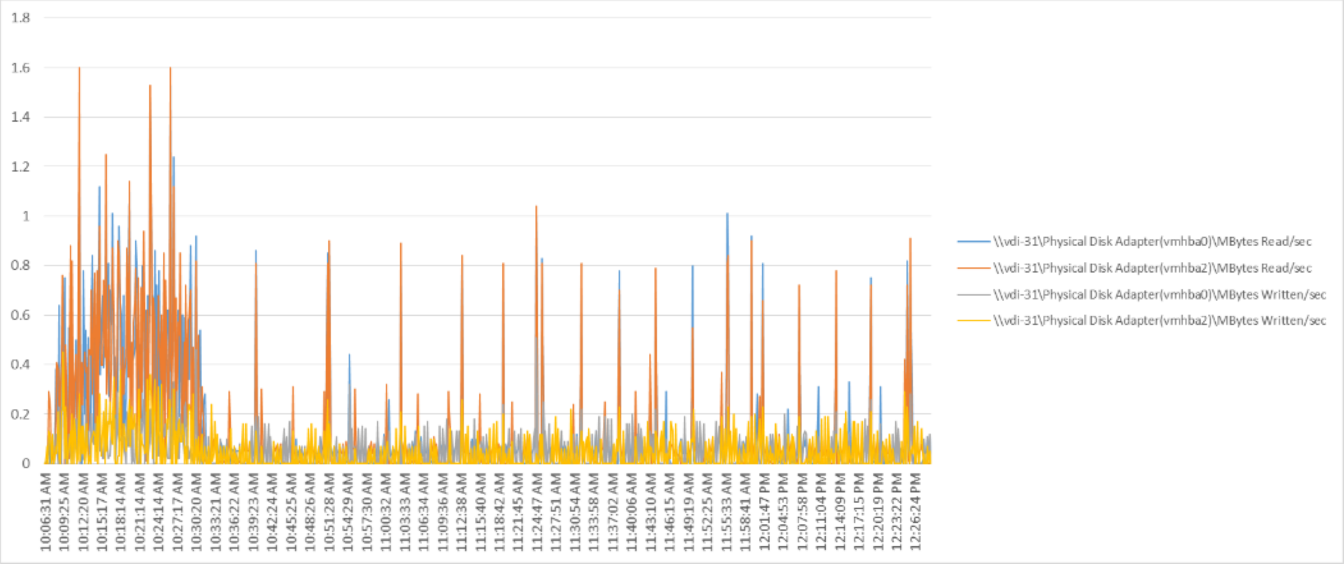

Figure 258 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Memory Utilization

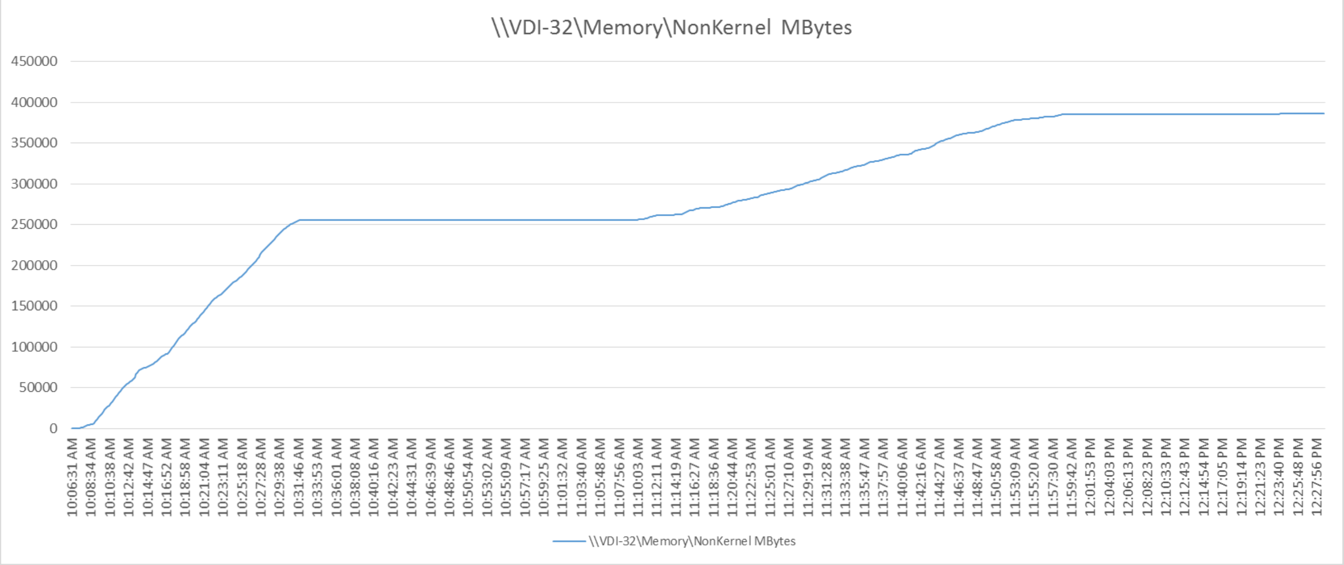

Figure 259 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host CPU Utilization

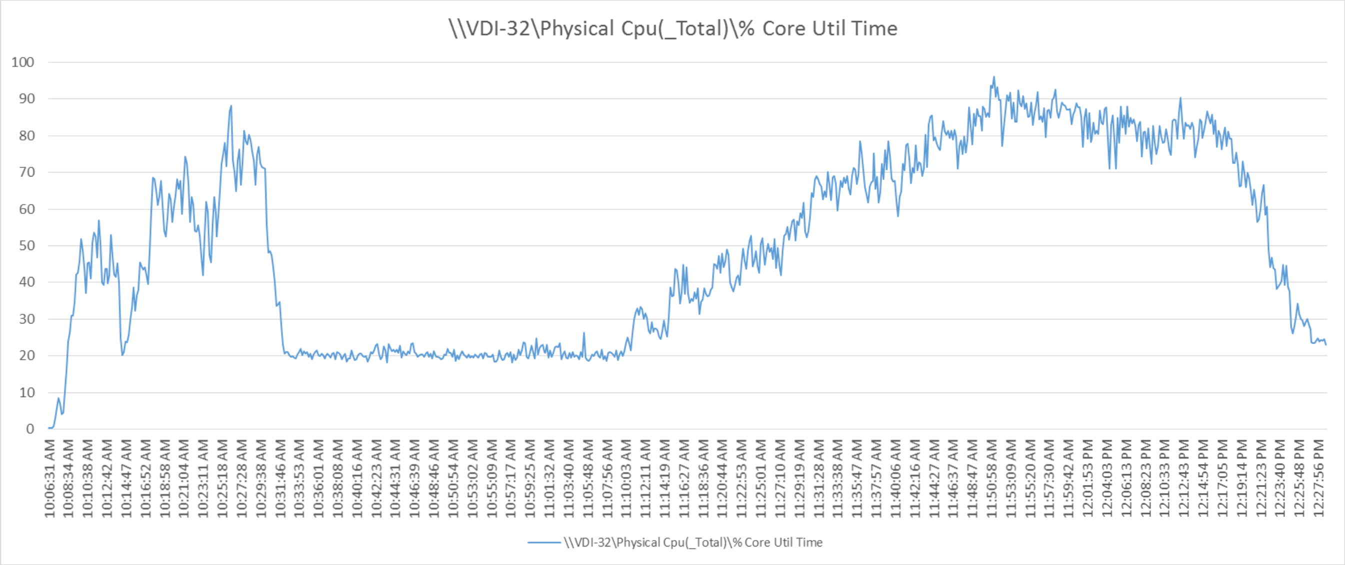

Figure 260 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Network Utilization

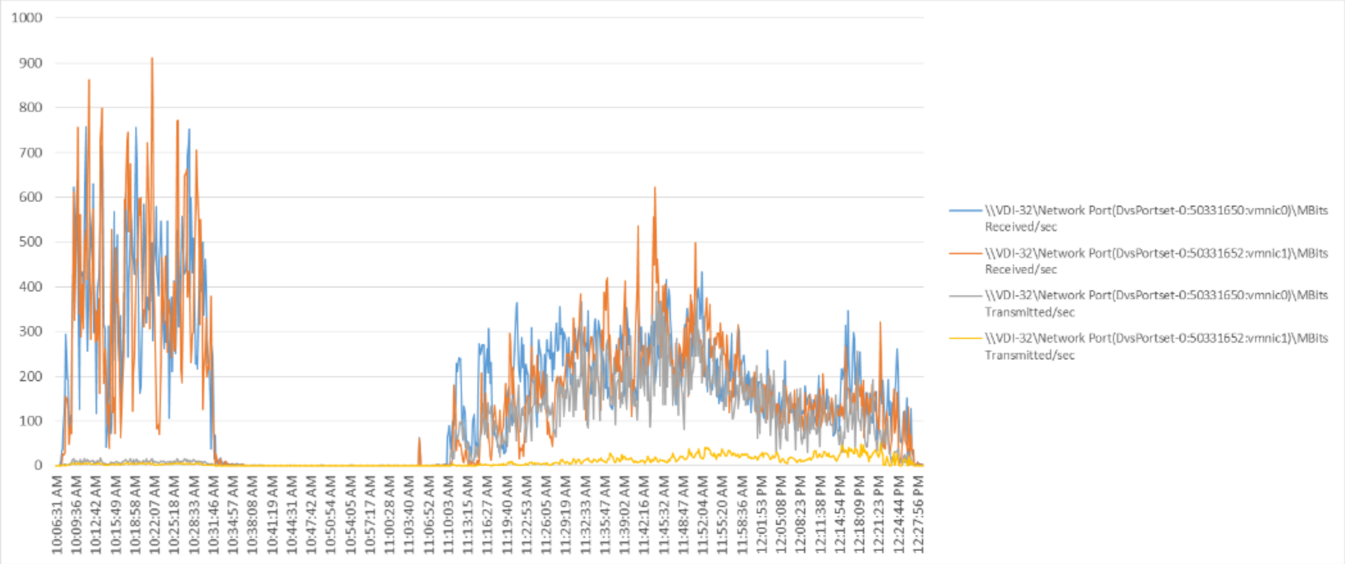

Figure 261 Full Scale | 6000 Mixed Users | Non-Persistent Hosts | Host Fibre Channel Utilization

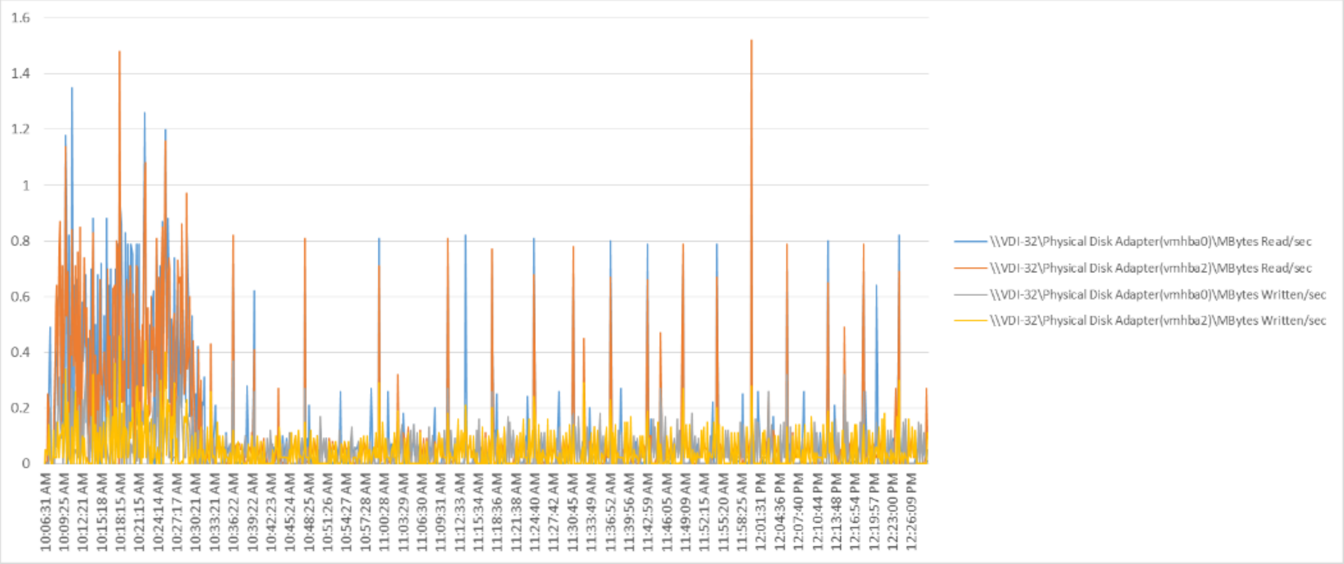

References

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
- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/release/notes/CiscoUCSManager-RN-3-1.html

Cisco UCS Virtual Interface Cards

- http://www.cisco.com/c/en/us/products/collateral/interfaces-modules/ucs-virtual-interface-card-1340/datasheet-c78-732517.html
- http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.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-ex-switch/index.html

Cisco MDS 9000 Service Switch References

- http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html
- http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html
- http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/datasheet-listing.html

Citrix References

- https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-15-ltsr.html
- https://docs.citrix.com/en-us/provisioning/7-15.html

- https://support.citrix.com/article/CTX216252?recommended
- https://support.citrix.com/article/CTX117374
- https://support.citrix.com/article/CTX202400
- https://support.citrix.com/article/CTX205488

FlexPod

- https://www.flexpod.com
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1design.html?referring_site=RE&pos=1&page=https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

VMware References

- https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
- https://labs.vmware.com/flings/vmware-os-optimization-tool
- https://pubs.vmware.com/view-51/index.jsp?topic=%2Fcom.vmware.view.planning.doc%2FGUID-6CAFE558-A0AB-4894-A0F4-97CF556784A9.html

Microsoft References

- https://technet.microsoft.com/en-us/library/hh831620(v=ws.11).aspx
- https://technet.microsoft.com/en-us/library/dn281793(v=ws.11).aspx
- https://support.microsoft.com/en-us/kb/2833839
- https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

NetApp Reference Documents

- http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx
- http://www.netapp.com/us/products/data-management-software/ontap.aspx
- https://mysupport.netapp.com/documentation/docweb/index.html?productID=62379&language=en-US
- http://www.netapp.com/us/products/management-software/
- http://www.netapp.com/us/products/management-software/vsc/

About the Authors

Vadim Lebedev, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Vadim Lebedev is a member of the Cisco's Computing Systems Product Group team focusing on design, testing, and solutions validation, technical content creation, and performance testing/benchmarking. He has years of experience in server and desktop virtualization. Vadim is a subject matter expert on Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and Citrix Certified Expert - Virtualization.

Chris Rodriguez, Senior Technical Marketing Engineer, NetApp

Chris Rodriguez (C-Rod) is a Senior Technical Marketing Engineer at NetApp, who has been involved with VDI since early 1990's. Chris has professional services experience with implementing VDI products on NetApp storage with many customers. In addition, Chris has 15 years of Enterprise Storage experience. Currently, Chris works on Cloud and VDI reference architectures for the Converged Infrastructure Engineering team at NetApp and he has been conducting reference architectures with Microsoft Azure, Microsoft Cloud Platform, Citrix Cloud Platform, Citrix XenDesktop and VMware Cloud (VMC) on NetApp storage.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the significant contribution and expertise that resulted in developing this document:

- David Arnette
- Chris Gebhardt
- Jenn Schrie
- Joe Drake