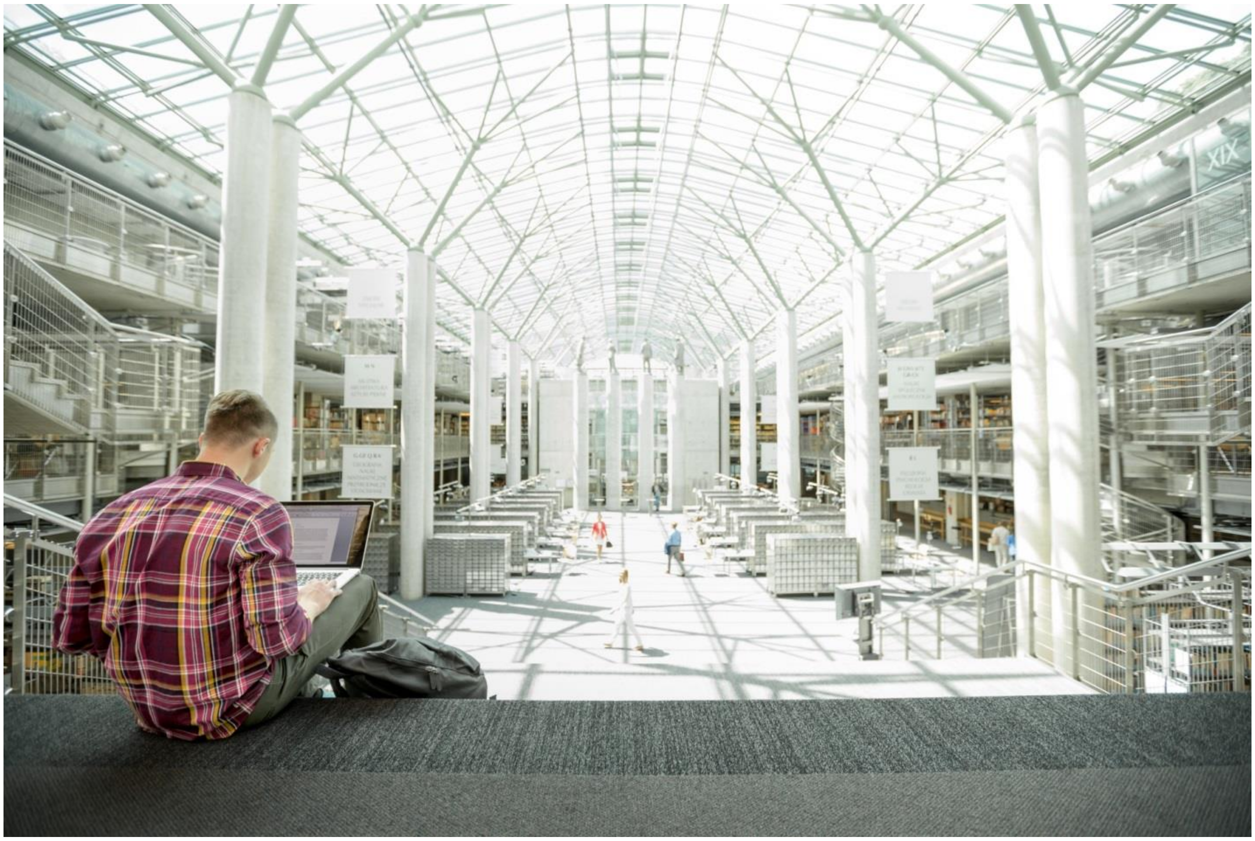

FlashStack with Cisco UCS and Pure Storage FlashArray//m for 5000 Citrix XenDesktop Users

Cisco UCS B200 M4 Blade Servers with Pure Storage FlashArray//m50 Array on Citrix XenDesktop 7.9 and VMware ESXi 6.0

Last Updated: December 22, 2016

About Cisco Validated Designs

The 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.

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

This document provides a reference architecture and design guide for up to 5000 seat mixed workload on Cisco UCS and Pure Storage FlashArray//m with Citrix XenApp server-based sessions, XenDesktop persistent Windows 10 virtual desktops and XenDesktop pooled Windows 10 virtual desktops on VMware vSphere 6. The solution is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and Pure Storage all flash array.

This solution is 100 percent virtualized on Cisco UCS B200 M4 blade server booting VMware vSphere 6.0 Update 1 via fibre channel SAN from the Pure Storage FlashArray//M50 storage array The virtual desktops are powered by Citrix Provisioning Services 7.9 with a mix of Citrix XenDesktop 7.9, which now incorporates both tradition persistent and non-persistent virtual Windows 10/8/10 desktops and hosted shared Server 2008 R2 or Server 2012 R2 server desktops, providing unparalleled scale and management simplicity. Citrix XenDesktop pooled Window 10 desktops (1200,) persistent Windows 10 desktops (1200) and Citrix XenApp Server 2012 R2 RDS server based desktop sessions (2600) were provisioned on the Pure Storage array. Where applicable the document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1 Knowledge Worker workload running in benchmark mode.

The 5000 seat solution provides a large scale building block that can be replicated to confidently scale out to tens of thousands of users.

Solution Overview

Introduction

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 and Pure Storage 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 7.9 mixed workload solution with Pure Storage FlashArray//m50, Cisco UCS Blade Servers, Cisco Nexus 9000 series ethernet switches and Cisco MDS 9000 series fibre channel switches.

What's New?

This is the second desktop virtualization Cisco Validated Design with Pure Storage. It incorporates the following features:

- Validation of Cisco Nexus 9000 with a Pure Storage all flash storage array
- Validation of Cisco MDS 9000 with a Pure Storage all flash storage array
- Support for the Cisco UCS 3.1(1) release and Cisco UCS B200-M4 servers
- Support for the latest release of Pure Storage FlashArray//m hardware and Purity Operating Environment 4.5.5 and 4.6.8.
- A Fibre Channel storage design supporting SAN LUNs
- Cisco Nexus 1000v distributed virtual switch technology
- Cisco UCS Inband KVM Access
- Cisco UCS vMedia client for vSphere Installation

- Cisco UCS Firmware Auto Sync Server policy
- VMware vSphere 6.0 Hypervisor
- Citrix XenDesktop 7.9 persistent and non-persistent Windows 10 virtual desktops and Citrix XenApp 7.9 RDS shared server sessions

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

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

The use cases include:

- Enterprise Data Center
- Service Provider Data Center

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 persistent and non-persistent virtual desktops and Citrix XenApp RDS server desktop sessions based on Microsoft Server 2012 R2. The mixed workload solution includes Pure Storage FlashArray//m storage array, 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 block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

The combination of technologies from Cisco Systems, Inc., Citrix Systems, Inc., Pure Storage 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 M4 half-width blade with dual 14-core 2.4 GHz Intel Xeon (E5-2680v4) processors and 512GB of memory for Citrix XenApp and XenDesktop hosts supports more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel Xeon E5-2680 v4 14-core 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 M4 blades for virtualized desktop and infrastructure workloads. The design provides N+1 server fault tolerance for hosted virtual desktops, hosted shared desktops and infrastructure services.

- Stress-tested to the limits during aggressive boot scenario. The 5000-user mixed hosted virtual desktop and hosted shared desktop environment booted and registered with the XenDesktop 7.9 Delivery Controllers in under 15 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.

- Stress-tested to the limits during simulated login storms. All 5000 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 less than 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 6.0. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, Provisioning Servers, SQL Servers, XenDesktop Delivery Controllers, XenDesktop VDI desktops, and XenApp RDS servers were hosted as virtual machines. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the FlashStack converged infrastructure with stateless Cisco UCS Blade Servers, and Pure Storage fibre channel storage.

- Cisco maintains industry leadership with the new Cisco UCS Manager 3.1(1) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director insure that customer environments are consistent locally, across Cisco UCS Domains and across the globe, our software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage and network.

- Our 10G unified fabric story gets additional validation on 6200 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.

- Pure Storage FlashArray//m provides industry-leading storage solutions that efficiently handle the most demanding I/O bursts (for example, login storms), profile management, and user data management, deliver simple and flexible business continuance, and help reduce storage cost per desktop.

- Pure Storage FlashArray//m provides a simple to understand storage architecture for hosting all user data components (VMs, profiles, user data) on the same storage array.

- Pure Storage Purity Operating System enables users to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.

- Pure Storage plugin for VMware vSphere hypervisor has deep integrations with vSphere, providing easy-button automation for key storage tasks such as storage repository provisioning and storage resize directly from the VCenter web client in a single pane of glass.

- Latest and greatest virtual desktop and application product. Citrix XenApp™ and XenDesktop™ 7.9 follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new XenDesktop release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of users increase. In addition, HDX enhancements help to optimize performance and improve the user experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.

- Optimized to achieve the best possible performance and scale. For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the

XenApp 7.9 RDS virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.

- Provisioning desktop machines made easy. Citrix Provisioning Services 7.9 created hosted virtual desktops as well as XenApp hosted shared desktop virtual machines for this solution using a single method for both, the "PVS XenDesktop Setup Wizard". The addition of the feature "Cache in RAM with overflow on hard disk" greatly reduced the amount of IOPS endured by the storage.

Cisco Desktop Virtualization Solutions: Data Center

The Evolving Workplace

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

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

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

Figure 1 Cisco Data Center Partner Collaboration

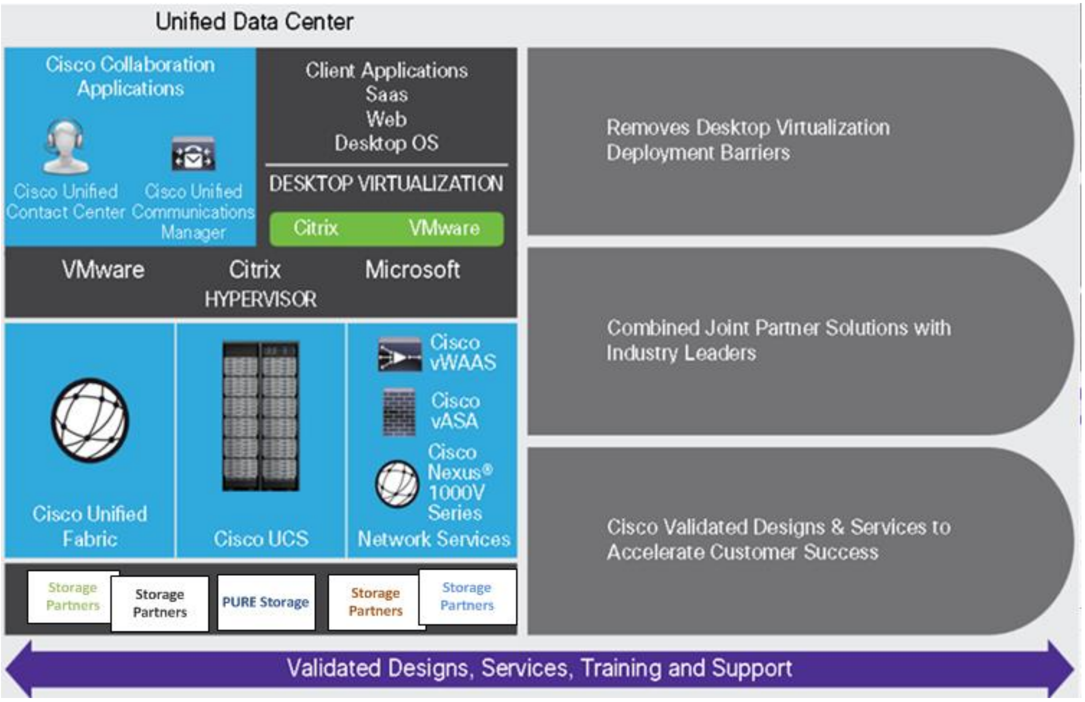
Unified Data Center
Cisco Collaboration Applications
Cisco Unified Contact Center
Cisco Unified Communications Manager
Client Applications
Saas
Web
Desktop OS
DESKTOP VIRTUALIZATION
Citrix
VMware
VMware
Citrix
Microsoft
HYPERVISOR
Cisco Unified Fabric
Cisco UCS
Cisco vWAAS
Cisco vASA
Cisco Nexus® 1000V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

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

Cisco Desktop Virtualization Focus

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

Simplified

Cisco UCS provides a radical new approach to industry-standard computing and provides the core of the data center infrastructure for desktop virtualization. Among the many features and benefits of Cisco UCS are the drastic reduction in the number of servers needed and in the number of cables used per server, and the capability to rapidly deploy or 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 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 and Pure have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlashStack. Cisco Desktop Virtualization Solutions have been tested with all the leading hypervisors, including VMware vSphere, Citrix XenServer, and Microsoft Hyper-V.

Secure

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

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

Scalable

Growth of a desktop virtualization solution is all but inevitable, so a solution must be able to scale, and scale predictably, with that growth. The Cisco Desktop Virtualization Solutions support high virtual-desktop density (desktops per server), and additional servers scale with near-linear performance. Cisco data center infrastructure 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 Pure help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs based on Citrix XenDesktop, Cisco UCS, and Pure joint solutions have demonstrated scalability and performance, with up to 5000 desktops up and running in 30 minutes.

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

Savings and Success

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

The simplified deployment of Cisco 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 Architecture

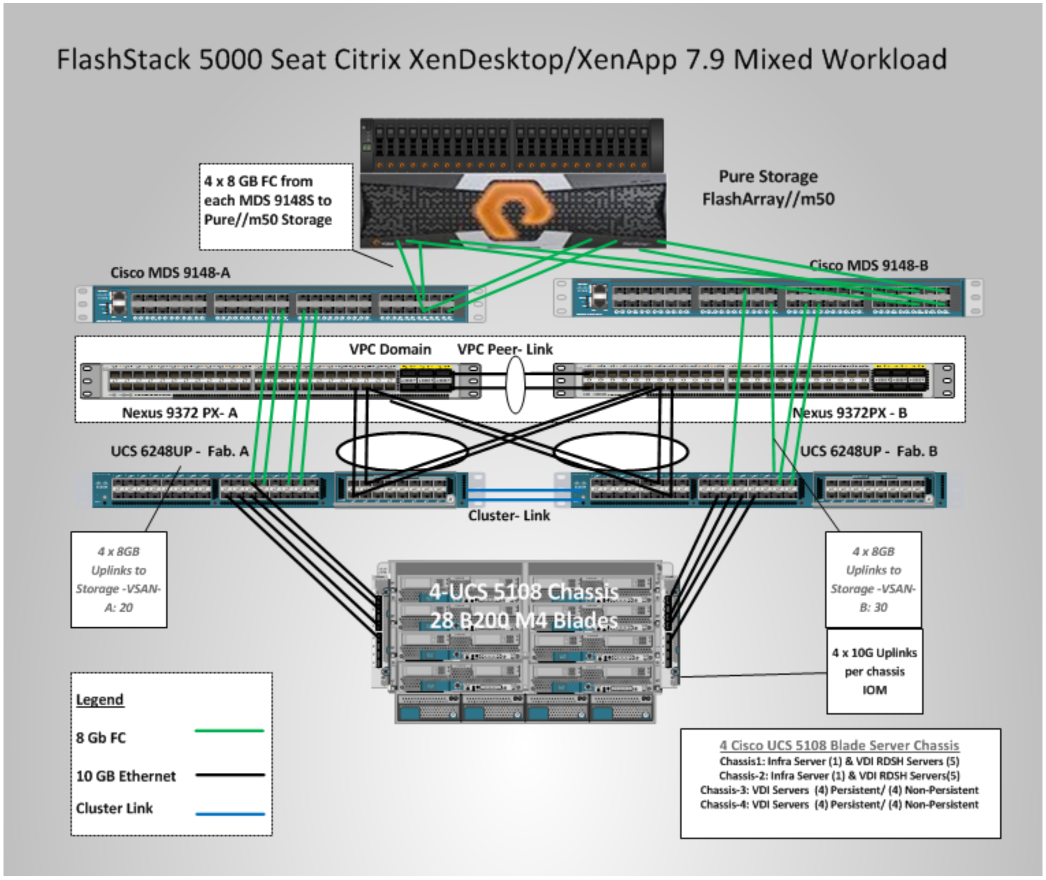
FlashStack 5000 Seat Citrix XenDesktop/XenApp 7.9 Mixed Workload
4 x 8 GB FC from each MDS 9148S to Pure//m50 Storage
Pure Storage FlashArray//m50
Cisco MDS 9148-A
Cisco MDS 9148-B
VPC Domain
VPC Peer- Link
Nexus 9372 PX- A
Nexus 9372PX - B
UCS 6248UP - Fab. A
UCS 6248UP - Fab. B
Cluster- Link
4 x 8GB Uplinks to Storage -VSAN-A: 20
4 x 8GB Uplinks to Storage -VSAN-B: 30
4-UCS 5108 Chassis
28 B200 M4 Blades
4 x 10G Uplinks per chassis IOM
Legend
8 Gb FC
10 GB Ethernet
Cluster Link
4 Cisco UCS 5108 Blade Server Chassis
Chassis1: Infra Server (1) & VDI RDSH Servers (5)
Chassis-2: Infra Server (1) & VDI RDSH Servers(5)
Chassis-3: VDI Servers (4) Persistent/ (4) Non-Persistent
Chassis-4: VDI Servers (4) Persistent/ (4) Non-Persistent

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches
- Two Cisco MDS 9148B 16GB Fibre Channel switches
- Two Cisco UCS 6248UP Fabric Interconnects
- Four Cisco UCS 5108 Blade Chassis
- Thirty Cisco UCS B200 M4 Blade Servers
- One Pure Storage FlashArray//m50 storage array
- One Pure Storage 44TB external shelf

For desktop virtualization, the deployment includes Citrix XenDesktop 7.9 running on VMware vSphere 6. The design is intended to provide a large scale building block for both XenApp RDS hosted server sessions and Windows 10 persistent and non-persistent desktops in the following ratio:

- 2600 XenApp 7.9 hosted server desktop sessions

- 1200 XenDesktop Windows 10 non-persistent virtual desktops
- 1200 XenDesktop Windows 10 persistent virtual desktops

The data provided in this document will allow our customers to adjust the mix of RDS and VDI desktops to suite their environment. For example, additional blade 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 5000 seat mixed workload virtual desktop solution with Citrix Systems and Pure Storage. Configuration guidelines are provided that refer the reader to which redundant component is being configured with each step. For example, node01 and node02 are used to identify the two Pure 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 6248UP 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.

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 10 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, 10-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 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 3 Cisco Data Center Overview

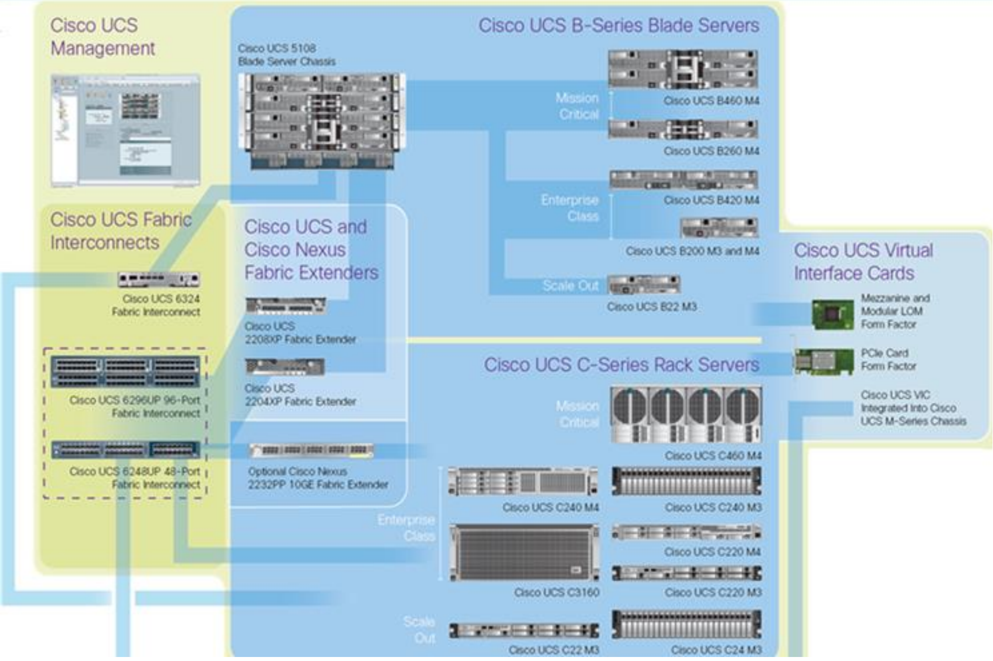
Cisco UCS Management
Cisco UCS Fabric Interconnects
Cisco UCS 6324 Fabric Interconnect
Cisco UCS 6296UP 96-Port Fabric Interconnect
Cisco UCS 6248UP 48-Port Fabric Interconnect
Cisco UCS 5108 Blade Server Chassis
Cisco UCS and Cisco Nexus Fabric Extenders
Cisco UCS 2208XP Fabric Extender
Cisco UCS 2204XP Fabric Extender
Optional Cisco Nexus 2232PP 10GE Fabric Extender
Cisco UCS B-Series Blade Servers
Mission Critical
Cisco UCS B460 M4
Cisco UCS B260 M4
Enterprise Class
Cisco UCS B420 M4
Cisco UCS B200 M3 and M4
Scale Out
Cisco UCS B22 M3
Cisco UCS Virtual Interface Cards
Mezzanine and Modular LOM Form Factor
PCIe Card Form Factor
Cisco UCS VIC Integrated Into Cisco UCS M-Series Chassis
Cisco UCS C-Series Rack Servers
Mission Critical
Cisco UCS C460 M4
Enterprise Class
Cisco UCS C240 M4
Cisco UCS C240 M3
Cisco UCS C220 M4
Cisco UCS C3160
Cisco UCS C220 M3
Scale Out
Cisco UCS C22 M3
Cisco UCS C24 M3

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 6200 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 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 6200 Series provides both LAN and SAN connectivity for all blades in the domain.

For networking, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, 1-terabit (Tb) switching capacity, and 160 Gbps of bandwidth per chassis, independent of packet size and enabled services. The product series supports Cisco low-latency, lossless, 10 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 4 Cisco UCS 6200 Series Fabric Interconnect

Cisco UCS B200 M4 Blade Server

The Cisco UCS B200 M4 Blade Server (Figures 5 and 6) is a density-optimized, half-width blade server that supports two CPU sockets for Intel Xeon processor E5-2600 v3 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 M4 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 M4 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 5 Cisco UCS B200 M4 Front View

Figure 6 Cisco UCS B200 M4 Back View

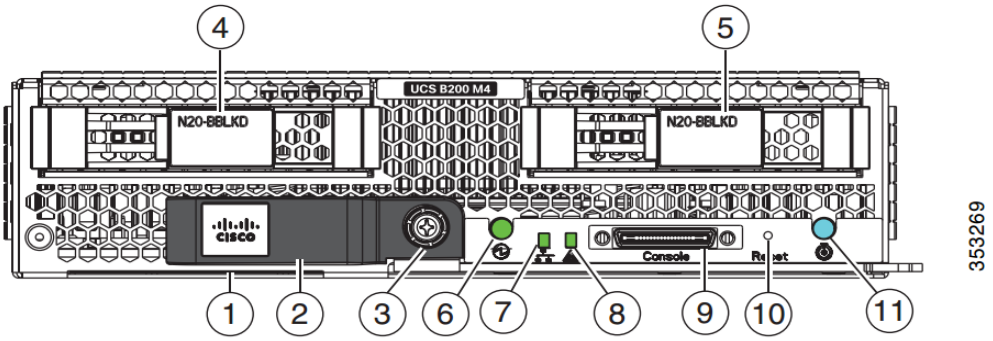

s

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 LED
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	Beaconing LED and button
6	Power button and LED	—	—

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 M4 Blade Server is one of the newest servers in the Cisco UCS portfolio.

The Cisco UCS B200 M4 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 M4 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 E5-2600 v3 and v4 product family, it offers up to 1.5TB of memory using 64GB DIMMs, up to two disk drives, and up to 80 Gbps of I/O throughput. The Cisco UCS B200 M4 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 M4 server with its leading memory-slot capacity and drive capacity.

Cisco UCS B200 M4 Features

The Cisco UCS B200 M4 provides:

- Up to two multicore Intel Xeon processor E5-2600 v3 series CPUs for up to 36 processing cores
- 24 DIMM slots for industry-standard DDR4 memory at speeds up to 2133 MHz, and up to 768 GB of total memory when using 32-GB DIMMs
- Two optional, hot-pluggable SAS and SATA HDDs or SSDs
- Cisco UCS VIC 1340, a 2-port, 40 Gigabit Ethernet and FCoE–capable modular (mLOM) mezzanine adapter
  - Provides two 40-Gbps unified I/O ports or two sets of four 10-Gbps unified I/O ports
  - Delivers 80 Gbps to the server
  - Adapts to either 10- or 40-Gbps fabric connections
    - Cisco FlexStorage local drive storage subsystem, with flexible boot and local storage capabilities that allow you to:
    - Configure the Cisco UCS B200 M4 to meet your local storage requirements without having to buy, power, and cool components that you do not need
    - Choose an enterprise-class RAID controller, or go without any controller or drive bays if you are not using local drives
    - Easily add, change, and remove Cisco FlexStorage modules

The Cisco UCS B200 M4 server is a half-width blade. Up to eight 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.

Cisco UCS B200 M4 Benefits

The Cisco UCS B200 M4 server is well suited for a broad spectrum of IT workloads, including:

- IT and web infrastructure
- Virtualized workloads
- Consolidating applications
- Virtual desktops
- Middleware
- Enterprise resource planning (ERP) and customer-relationship management (CRM) applications

Innovative Technologies

The Cisco UCS B200 M4 is one member of the Cisco UCS B-Series Blade Servers platform. As part of Cisco UCS, Cisco UCS B-Series servers incorporate many innovative Cisco technologies to help customers handle their most challenging workloads. Cisco UCS B-Series servers within a Cisco UCS management framework incorporate a standards-based unified network fabric, Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) virtualization support, Cisco UCS Manager, Cisco UCS Central Software, Cisco UCS Director software, and Cisco fabric extender architecture.

The Cisco UCS B200 M4 Blade Server delivers:

- Suitability for a wide range of applications and workload requirements
- Highest-performing CPU and memory options without constraints in configuration, power, or cooling
- Half-width form factor that offers industry-leading benefits
- Latest features of Cisco UCS VICs

For more information about the Cisco UCS B200 B4, see http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-b200-m4-blade-server/model.html

Cisco UCS VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 7) 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 M4 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 7 Cisco UCS VIC 1340

Figure 8 illustrates the Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M4 Blade Servers.

Figure 8 Cisco UCS VIC 1340 deployed in the Cisco UCS B200 M4

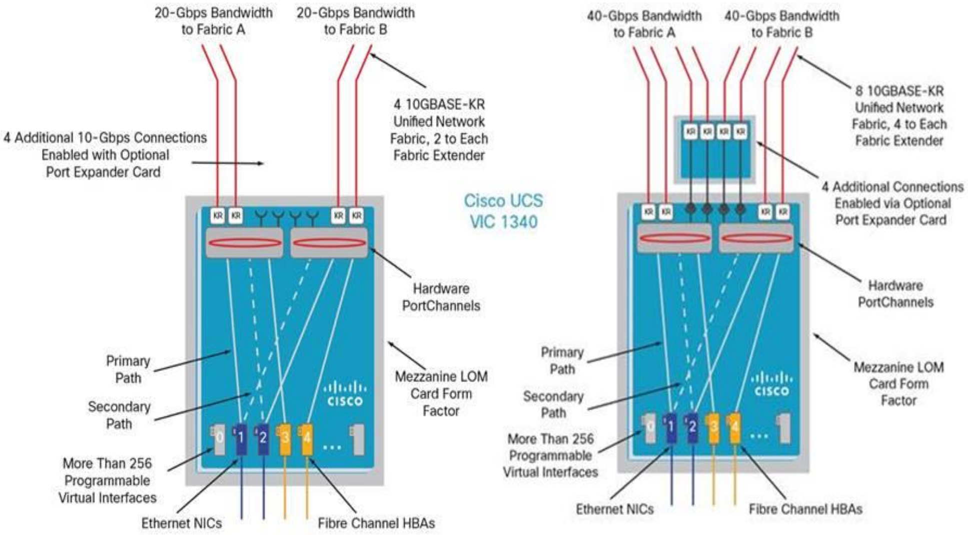
20-Gbps Bandwidth to Fabric A
20-Gbps Bandwidth to Fabric B
4 10GBASE-KR Unified Network Fabric, 2 to Each Fabric Extender
4 Additional 10-Gbps Connections Enabled with Optional Port Expander Card
KR
Cisco UCS VIC 1340
Hardware PortChannels
Mezzanine LOM Card Form Factor
Primary Path
Secondary Path
More Than 256 Programmable Virtual Interfaces
Ethernet NICs
Fibre Channel HBAs
0
1
2
3
4
cisco
40-Gbps Bandwidth to Fabric A
40-Gbps Bandwidth to Fabric B
8 10GBASE-KR Unified Network Fabric, 4 to Each 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 9372PX Switches

The Cisco Nexus 9372PX/9372PX-E Switches have 48 1/10-Gbps Small Form Pluggable Plus (SFP+) ports and 6 Quad SFP+ (QSFP+) uplink ports. All the ports are line rate, delivering 1.44 Tbps of throughput in a 1-rack-unit (1RU) form factor. Cisco Nexus 9372PX benefits are listed below.

Architectural Flexibility

- Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Leaf node support for Cisco ACI architecture is provided in the roadmap

- Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support

Feature Rich

- Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
- ACI-ready infrastructure helps users take advantage of automated policy-based systems management
- Virtual Extensible LAN (VXLAN) routing provides network services
- Cisco Nexus 9372PX-E supports IP-based endpoint group (EPG) classification in ACI mode

Highly Available and Efficient Design

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

Simplified Operations

- Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
- An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
- Python Scripting for programmatic access to the switch command-line interface (CLI)
- Hot and cold patching, and online diagnostics

Investment Protection

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

- 1.44 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 9 Cisco Nexus 9372PX Switch

Cisco Nexus 1000V Distributed Virtual Switch

Get highly secure, multitenant services by adding virtualization intelligence to your data center network with the Cisco Nexus 1000V Switch for VMware vSphere. This switch does the following:

- Extends the network edge to the hypervisor and virtual machines
- Is built to scale for cloud networks
- Forms the foundation of virtual network overlays for the Cisco Open Network Environment and Software Defined Networking (SDN)

Important Differentiators for the Cisco Nexus 1000V for VMware vSphere

The following lists the benefits of the Cisco Nexus 1000V for VMware vSphere:

- Extensive virtual network services built on Cisco advanced service insertion and routing technology
- Support for vCloud Director and vSphere hypervisor
- Feature and management consistency for easy integration with the physical infrastructure
- Exceptional policy and control features for comprehensive networking functionality
- Policy management and control by the networking team instead of the server virtualization team (separation of duties)

Virtual Networking Services

The Cisco Nexus 1000V Switch optimizes the use of Layer 4 - 7 virtual networking services in virtual machine and cloud environments through Cisco vPath architecture services.

Cisco vPath 2.0 supports service chaining so you can use multiple virtual network services as part of a single traffic flow. For example, you can simply specify the network policy, and vPath 2.0 can direct traffic through the Cisco Virtual Security Gateway for Nexus 1000V Switch for a zoning firewall.

Additionally, Cisco vPath works on VXLAN to support movement between servers in different Layer 2 domains. Together, these features promote highly secure policy, application, and service delivery in the cloud.

Cisco MDS 9348 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.

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.

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 and Desktop Broker

This Cisco Validated Design includes VMware vSphere 6 and Citrix XenDesktop 7.9.

VMware vSphere 6.0

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

VMware vSphere 6.0 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.

VMware ESXi 6.0 Hypervisor

vSphere 6.0 introduces a number of new features in the hypervisor:

- Scalability Improvements

ESXi 6.0 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 6.0 instance can support up to 480 logical CPUs, 12 terabytes (TB) of RAM, and 1024 virtual machines. By using the newest hardware advances, ESXi 6.0 enables the virtualization of applications that previously had been thought to be non-virtualizable.

Security Enhancements

- ESXi 6.0 offers these security enhancements:

- Account management: ESXi 6.0 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 6.0, 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 6.0, only one lockdown mode was available. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.0, 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 XenApp™ and XenDesktop™ 7.9

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 7.9, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

The XenDesktop 7.9 release offers these benefits:

- Comprehensive virtual desktop delivery for any use case. The XenDesktop 7.9 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 7.9 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 7.9 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 7.9 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 7.9

- Improved high-definition user experience. XenDesktop 7.9 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 are 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 7.9: 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 7.9 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 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 be repeated every five minutes until the update/restart begins.

For more information, see the Manage Machine Catalogs and Manage machines in 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 2012 Express SP2 is installed when you install the Delivery Controller. SP1 is no longer installed.

The component installers now automatically deploy newer Microsoft Visual C++ runtime versions: 32-bit and 64-bit Microsoft Visual C++ 2013, 2010 SP1, and 2008 SP1. Visual C++ 2005 is no longer deployed.

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 10 Logical Architecture of Citrix XenDesktop

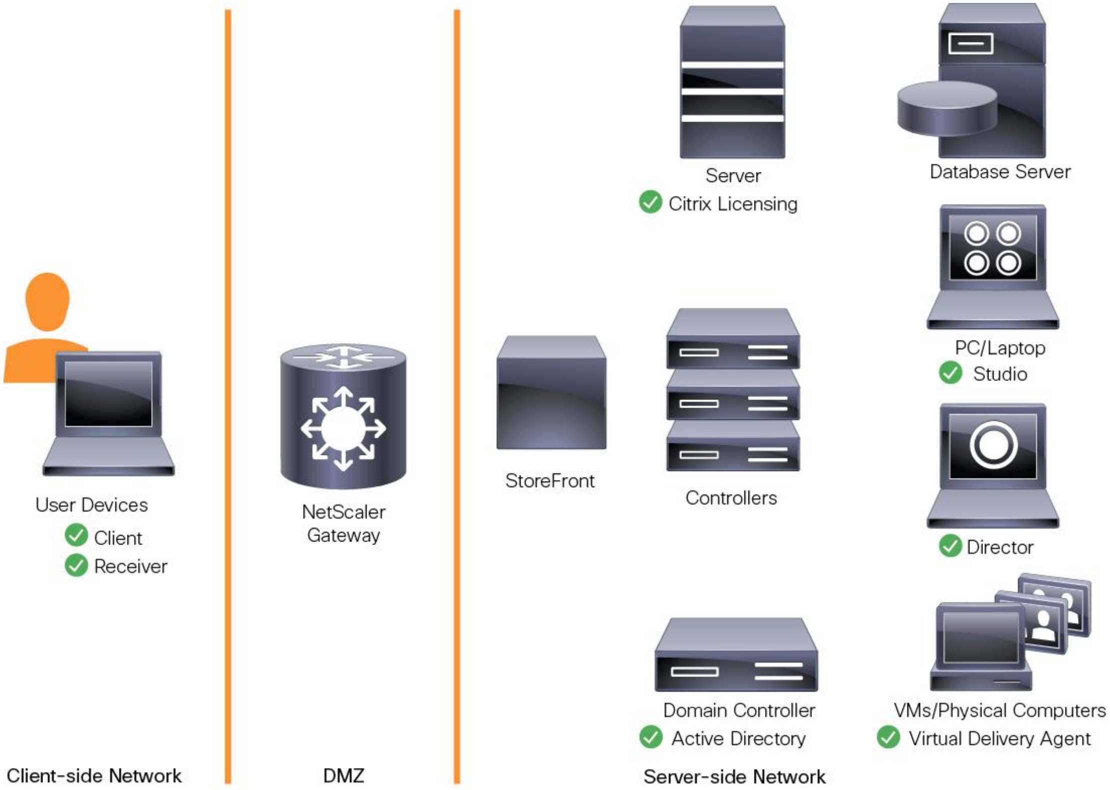

Citrix Provisioning Services 7.9

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 11 provides a high-level view of a basic PVS infrastructure and shows how PVS components might appear within that implementation.

Figure 11 Logical Architecture of Citrix Provisioning Services

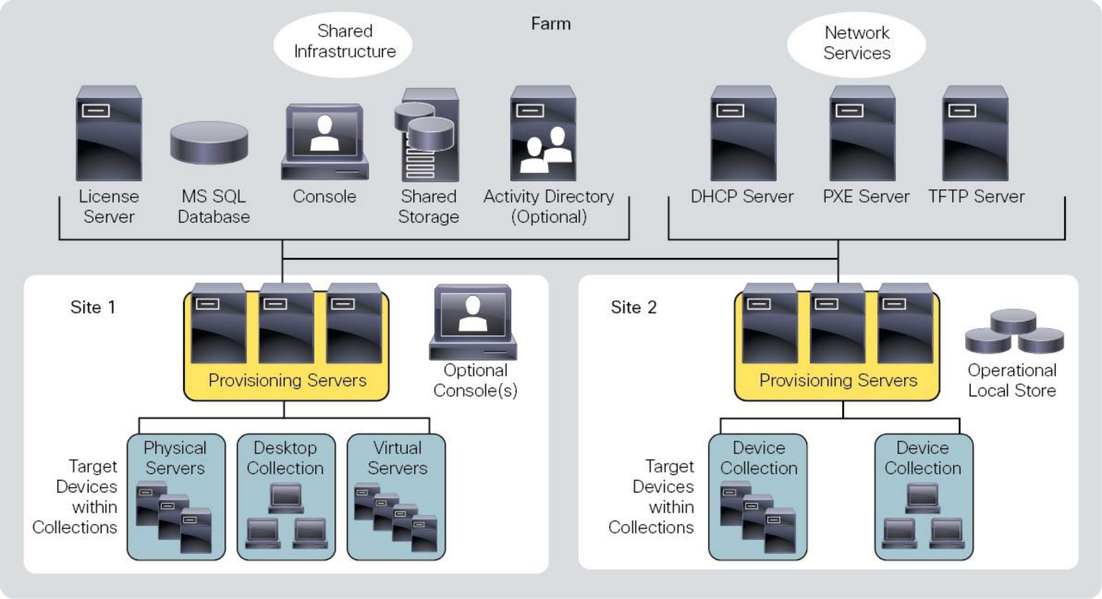

The following new features are available with Provisioning Services 7.9:

- Streaming VHDX formatted disks
- Support for Microsoft Windows 10 Enterprise and Professional editions
- Support for Unified Extensible Firmware Interface (UEFI) enhancements
- The licensing grace period for Provisioning Services has changed from 96 hours to 30 days, for consistency with XenApp and XenDesktop
- Enhancements to API
- vGPU-enabled XenDesktop machines can be provisioned using the Provisioning Services XenDesktop Setup Wizard
- Support for System Center Virtual Machine Manager Generation 2 VMs
- FIPS support
- XenApp Session Recording Enabled by Default

Pure Storage FlashArray//m50

FlashArray//m delivers breakthrough resiliency that never quits. Proven greater than 99.999 percent availability means your data is always available, always performing and always protected – with no performance loss. Now, with predictive support, you can predict vulnerability to known issues and take pre-emptive action to minimize your downtime risk.

What is FlashStack?

FlashStack CI (Converged Infrastructure) is a flexible, all-flash converged infrastructure solution that brings the flash revolution to your data center, faster. It combines the latest in compute, network, storage hardware and virtualization software, into a single, integrated architecture that speeds time to deployment, lowers overall IT costs and reduces deployment risk. Highly efficient components reduce the costs associated with power, cooling and data center space. Based on 100 percent flash storage, FlashStack CI provides the performance and reliability business-critical applications demand.

The hardware foundation of FlashStack CI includes Pure Storage FlashArrays, Cisco UCS Blade Servers, Cisco Nexus ethernet switches and Cisco MDS fibre channel switches. VMware vSphere provides the virtualization technology.

Figure 12 FlashStack Converged Infrastructure (CI)

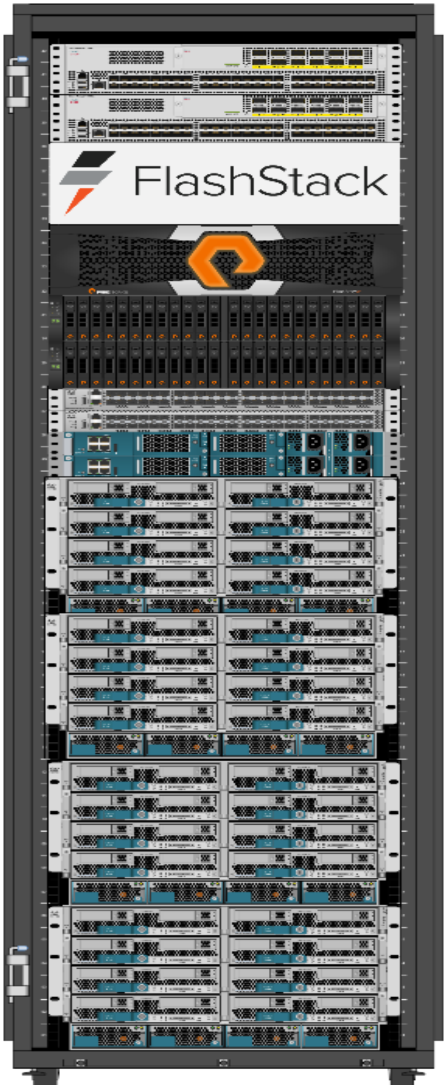

FlashStack CI is available from qualified FlashStack Partners who help to provide an excellent converged infrastructure ownership experience. FlashStack Partners have the knowledge and experience necessary to help streamline the sizing, procurement, and delivery of your entire system.

Both the hardware and software components are combined into a single integrated unit that helps in faster deployments and lowers overall IT costs.

Why FlashStack?

The following lists the benefits of FlashStack:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100% flash storage
  - Consolidate 100's of enterprise-class applications in a single rack
  - Scales easily, without disruption
  - Continuous growth through multiple FlashStack 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
  - Free FlashArray controller upgrades every three years with Forever Flash™
- Enterprise Grade Resiliency
  - Highly available architecture with no single point of failure
  - Non-disruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication
- Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases

Benefits of Pure Storage FlashArray//m Series

Who knew that moving to all-flash storage could help reduce the cost of IT? FlashArray//m makes server and workload investments more productive, while also lowering storage spend. With FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.

FlashArray//m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply were not possible with disk. All by Transforming Your Storage with FlashArray//m.

FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.

- Mini Size–Reduce power, space and complexity by 90 percent
  - 3U base chassis with 15-120+ TBs usable
  - ~1kW of power
  - 6 cables
- Mighty Performance–Transform your datacenter, cloud, or entire business
  - Up to 300,000 32K IOPS
  - Up to 9 GB/s bandwidth
  - <1ms average latency
- Modular Scale–Scale FlashArray//m inside and outside of the chassis for generations
  - Expandable to ~½ PB usable via expansion shelves
  - Upgrade controllers and drives to expand performance and/or capacity
- Meaningful Simplicity–Appliance-like deployment with worry-free operations
  - Plug-and-go deployment that takes minutes, not days

  - Non-disruptive upgrades and hot-swap everything
  - Less parts = more reliability

The FlashArray//m expands upon the FlashArray's modular, stateless architecture, designed to enable expandability and upgradability for generations. The FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your business' needs today and tomorrow.

The Pure Storage FlashArray is ideal for:

- Accelerating Databases and Applications: Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.
- Virtualizing and Consolidating Workloads: Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.
- Delivering the Ultimate Virtual Desktop Experience: Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance with simple management for under $50/desktop.
- Protecting and Recovering Vital Data Assets: Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

- Consistent Performance: FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.
- Less Cost than Disk: Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure. With VDI workloads data reduction is typically > 10:1.
- Mission-Critical Resiliency: FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.
- Disaster Recovery Built-In: FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

- Simplicity Built-In: FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, Windows PowerShell, Python, REST API, or OpenStack.

FlashArray//m Specifications

Figure 13 Pure Storage FlashArray//m Portfolio

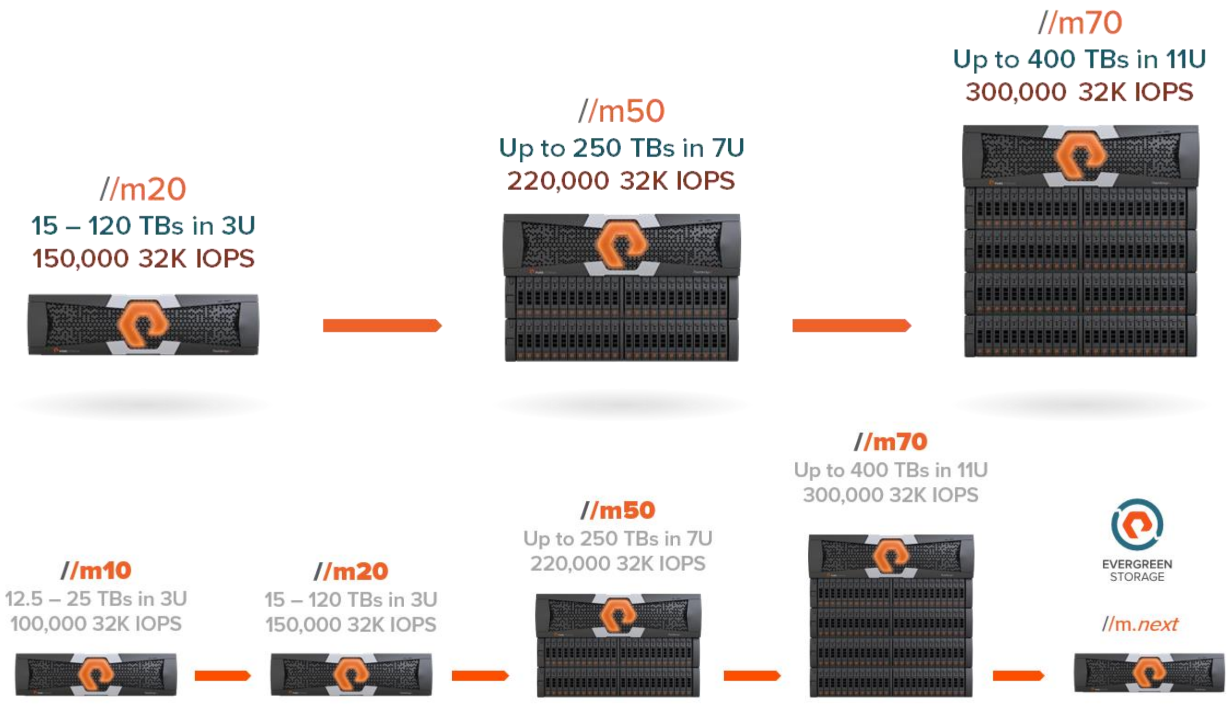

Table 1 Pure Storage FlashArray//m Series Controller Specifications

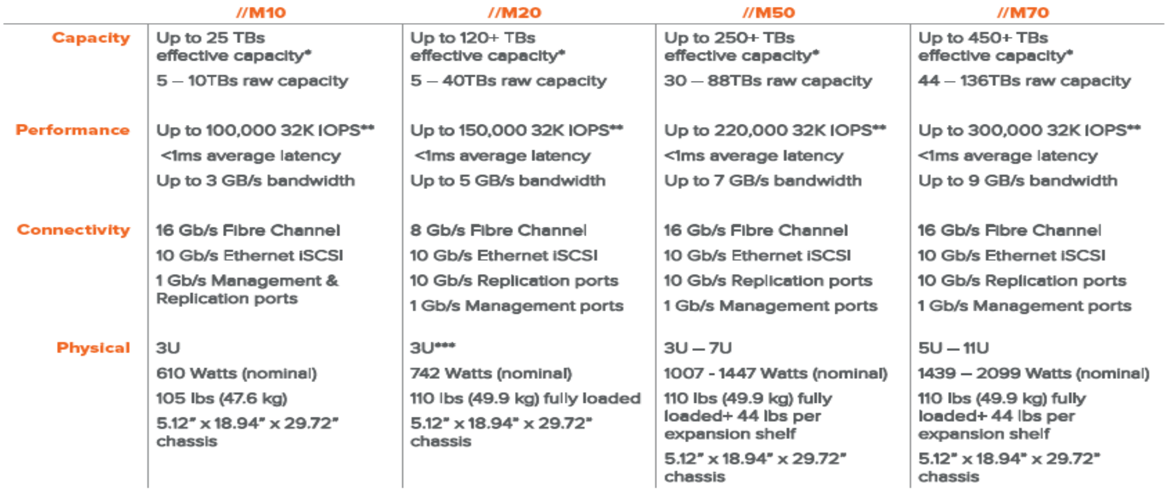

	//M10	//M20	//M50	//M70
Capacity	Up to 25 TBs effective capacity* 5 – 10TBs raw capacity	Up to 120+ TBs effective capacity* 5 – 40TBs raw capacity	Up to 250+ TBs effective capacity* 30 – 88TBs raw capacity	Up to 450+ TBs effective capacity* 44 – 136TBs raw capacity
Performance	Up to 100,000 32K IOPS** <1ms average latency Up to 3 GB/s bandwidth	Up to 150,000 32K IOPS** <1ms average latency Up to 5 GB/s bandwidth	Up to 220,000 32K IOPS** <1ms average latency Up to 7 GB/s bandwidth	Up to 300,000 32K IOPS** <1ms average latency Up to 9 GB/s bandwidth
Connectivity	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 1 Gb/s Management & Replication ports	8 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 10 Gb/s Replication ports 1 Gb/s Management ports
Physical	3U 610 Watts (nominal) 105 lbs (47.6 kg) 5.12" x 18.94" x 29.72" chassis	3U*** 742 Watts (nominal) 110 lbs (49.9 kg) fully loaded 5.12" x 18.94" x 29.72" chassis	3U – 7U 1007 - 1447 Watts (nominal) 110 lbs (49.9 kg) fully loaded+ 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis	5U – 11U 1439 – 2099 Watts (nominal) 110 lbs (49.9 kg) fully loaded+ 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis

* Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1.

** Why does Pure Storage quote 32K, not 4K IOPS? The industry commonly markets 4K IOPS benchmarks to inflate performance numbers, but real-world environments are dominated by IO sizes of 32K or larger. FlashArray adapts automatically to 512B-32KB IO for superior performance, scalability, and data reduction.

***//m20 can be expanded beyond the 3U base chassis with expansion shelves.

Purity Operating Environment

Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.

- Storage Software Built for Flash—The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer- grade MLC flash.

- Granular and Adaptive—Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.

- Best Data Reduction Available—FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

- Highly Available and Resilient—FlashProtect implements high availability, dual-parity RAID-3D, non- disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.

- Backup and Disaster Recovery Built In—FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.

Pure1

- Pure1 Manage—By combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashArray wherever you are – with just a web browser.

- Pure1 Connect–A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashArray//m to all your data center and cloud monitoring, management, and orchestration tools.
- Pure1 Support–FlashArray//m is constantly cloud- connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.
- Pure1 Collaborate–Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.

Engineered to Stay Evergreen

Never repurchase a TB of storage you already own. Forever Flash, a key business model component of Evergreen Storage, helps you Run and Upgrade your storage over time with full investment protection. With ForeverFlash we'll regularly modernize your storage for you – at no extra charge.

Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. You will get our maintenance and support for all components of your system – including flash – and you will get modern and new controller upgrades included every three years with your 3-year renewal. But your maintenance pricing will remain flat. Pure Storage's vision for Evergreen Storage is delivered by a combination of the FlashArray's stateless, modular architecture and the ForeverFlash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

Architecture and Design Considerations for Desktop Virtualization

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

- Knowledge Workers today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.

- External Contractors are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.

- Task Workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.

- Mobile Workers need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.

- Shared Workstation users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

- Traditional PC: A traditional PC is what typically constitutes a desktop environment: physical device with a locally installed operating system.

- 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 2012, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Hosted Virtual Desktop: 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 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 hosted virtual desktops and XenApp hosted shared server desktops were validated. Each of the sections provides some fundamental design decisions for this environment.

Understanding Applications and Data

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

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

Project Planning and Solution Sizing Sample Questions

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

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

- Has a VDI pilot plan been created based on the business analysis of the desktop groups, applications and data?

- Is there infrastructure and budget in place to run the pilot program?

- Are the required skill sets to execute the VDI project available? Can we hire or contract for them?
- Do we have end user experience performance metrics identified for each desktop sub-group?
- How will we measure success or failure?
- What is the future implication of success or failure?

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

- What is the desktop OS planned? Windows 10, 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 10/8/10?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?
- What is the end point graphics processing capability?
- Will XenApp RDS be used for Hosted Shared Server Desktops or exclusively XenDesktop HVD?
- Are there XenApp hosted applications planned? Are they packaged or installed?
- Will Provisioning Server, Machine Creation Services, or another method be used for virtual desktop deployment?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- 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

Citrix XenDesktop is hypervisor-agnostic, so any of the following three hypervisors can be used to host RDS- and VDI-based desktops:

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

- Hyper-V: Microsoft Windows Server with Hyper-V is available in a Standard, Server Core and free Hyper-V Server versions. More information on Hyper-V can be obtained at the Microsoft web site: http://www.microsoft.com/en-us/server-cloud/windows-server/default.aspx.

- XenServer: Citrix® XenServer® is a complete, managed server virtualization platform built on the powerful Xen® hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows and Linux virtual servers and delivers cost-effective server consolidation and business continuity. More information on XenServer can be obtained at the web site: http://www.citrix.com/products/xenserver/overview.html.

For this CVD, the hypervisor used was VMware ESXi 6.0 Update 1a.

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 7.9 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 14 shows how users access desktops and applications through machine catalogs and delivery groups.

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

Figure 14 Access Desktops and Applications through Machine Catalogs and Delivery Groups

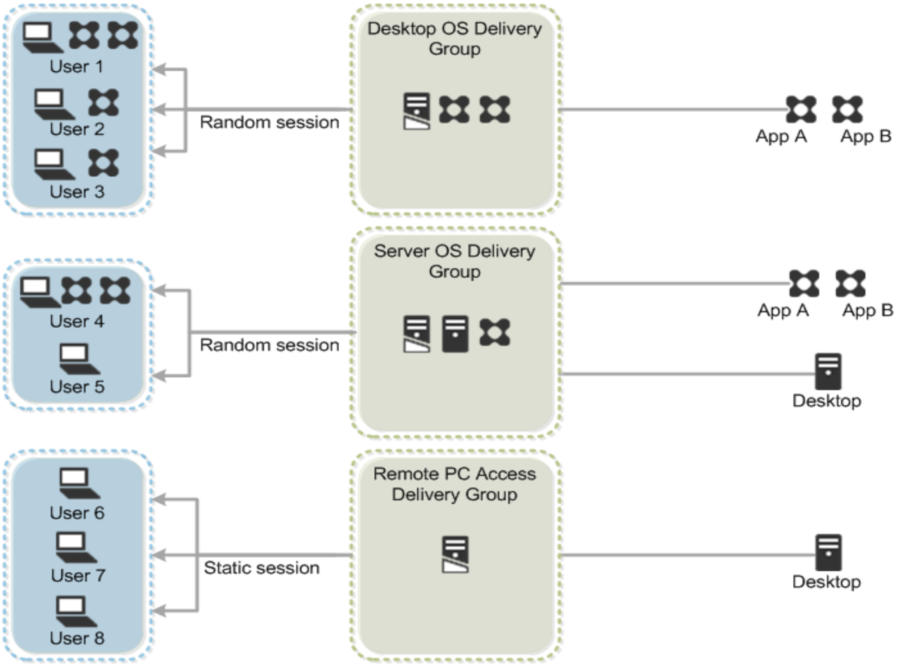

Citrix Provisioning Services

Citrix XenDesktop 7.9 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 (Step 1).

Figure 15 Citrix Provisioning Services Functionality

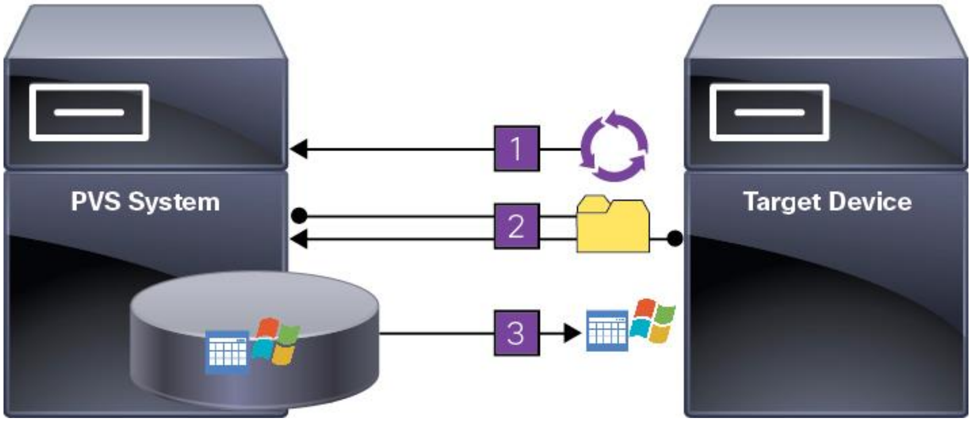

The target device downloads the boot file from a Provisioning Server (Step 2) and boots. Based on the boot configuration settings, the appropriate vDisk is mounted on the Provisioning Server (Step 3). The vDisk software is then streamed to the target device as needed, appearing as a regular hard drive 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 7.9 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 7.9 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 7.9 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 16 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.

Figure 16 Example of a Distributed Components Configuration

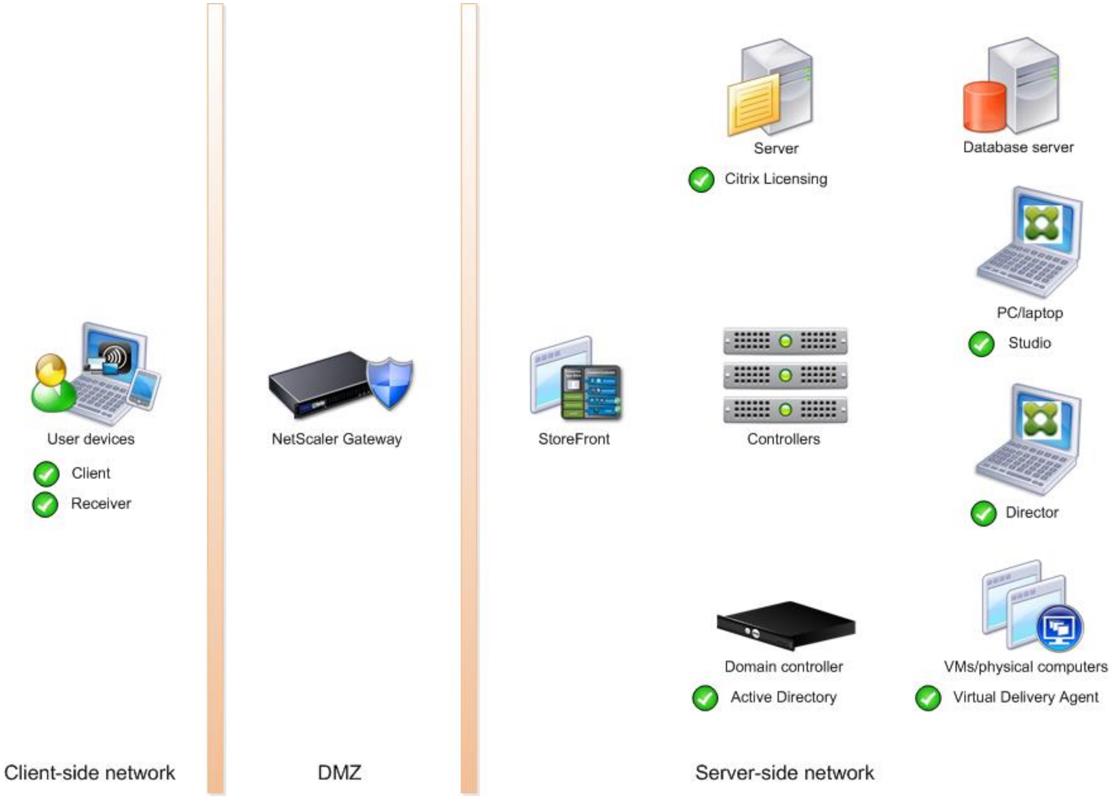

Multiple Site Configuration

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 17, 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. Two Cisco blade servers host the required infrastructure services (AD, DNS, DHCP, Profile, SQL, Citrix XenDesktop management, and web servers).

Figure 17   Multiple Sites

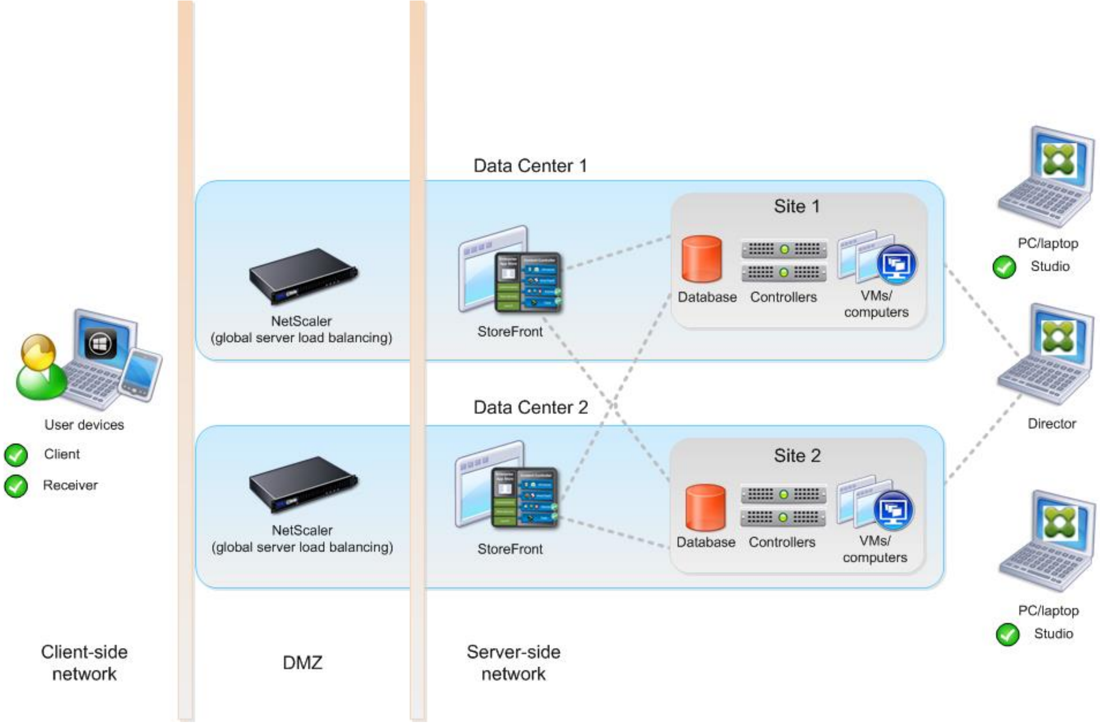

You can use StoreFront to aggregate resources from multiple sites to provide users with a single point of access with NetScaler. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

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.

Designing a XenDesktop Environment for a Mixed Workload

With Citrix XenDesktop 7.9, 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 machines	You want: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience. Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations. Application types: Any application.

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

For the Cisco Validated Design described in this document, a mix of Hosted Shared Desktops (HSDs) using RDS-based Server OS machines and Hosted Virtual Desktops (HVDs) using VDI-based Desktop OS machines were configured and tested. The mix consisted of a combination of both use cases. The following sections discuss design decisions relative to the Citrix XenDesktop deployment, including the CVD test environment.

High-Level Storage Architecture Design

This section outlines the recommended storage architecture for deploying a mix of various Citrix XenDesktop and XenApp delivery models on the same Pure Storage array. These models include persistent and non-persistent hosted VDI, hosted-shared desktops, and intelligent VDI layering, such as profile management and user data management.

For XenApp server shared desktops and the non-persistent Windows 10 virtual desktops, the following recommendations are best practices for the Provisioning Server write cache drives, user profiles, user data, and application virtualization:

- Provisioning Services (PVS) vDisk: CIFS/SMB 3 is recommended to host the PVS vDisk. CIFS/SMB 3 allows the same vDisk to be shared among multiple PVS servers while still maintaining resilience during storage node failover. This process results in significant operational savings and architecture simplicity. SMB3 is available in Windows Server 2012 or higher and provides persistent handles. SMB3 persistent handles prevents the PVS server from crashing during a storage node failover. Therefore, Windows 2012 is the required OS for a PVS server to ensure a stable PVS implementation.

- Provisioning Server write cache drives: Write cache drives should be 2 times the size of the virtual machine memory. Choose "Cache in device RAM with overflow on hard disk as the provisioning technique. Deduplication should not be enabled on this volume, because the rate of change is too high. The PVS write cache file should be set for thin provisioning at the storage layer.

- User Profiles: Use Citrix User Profile Manager to set policy and provision user profiles. Alternatively, Pure Storage has a native profile management capability, which was not used in this study.

- User Data: We recommend hosting user data on CIFS home directories to preserve data upon VM reboot or redeploy.

- Monitoring and management: We recommend using Citrix XenDesktop Director and Cisco UCS Performance Manager to provide monitoring and management of the solution.

Solution Hardware and Software

Products Deployed

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

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

This validated design document details the deployment of 5000 users for a mixed Citrix XenDesktop desktop workload featuring the following software:

This validated design document details the deployment of the multiple configurations extending to 5000 users for a mixed XenDesktop workload featuring the following software:

- Citrix XenApp 7.9 Shared Hosted Virtual Desktops (RDS) with PVS write cache on NFS storage
- Citrix XenDesktop 7.9 Non-Persistent Hosted Virtual Desktops (VDI) with PVS write cache on NFS storage
- Citrix XenDesktop 7.9 Persistent Hosted Virtual Desktops (VDI) provisioned with Citrix PVS and stored on fibre channel storage
- Citrix Provisioning Server 7.9
- Citrix User Profile Manager
- Citrix StoreFront 3
- VMware vSphere ESXi 6.0 Update 1 Hypervisor
- Microsoft Windows Server 2012 R2 and Windows 7 32-bit virtual machine Operating Systems
- Microsoft SQL Server 2012
- Cisco Nexus 1000V primary and secondary Virtual Supervisor Module

Figure 18 details the physical hardware and cabling deployed to enable the solution.

Figure 18 Virtual Desktop Workload Architecture for the 5000 Seat Citrix XenDesktop 7.9 on FlashStack

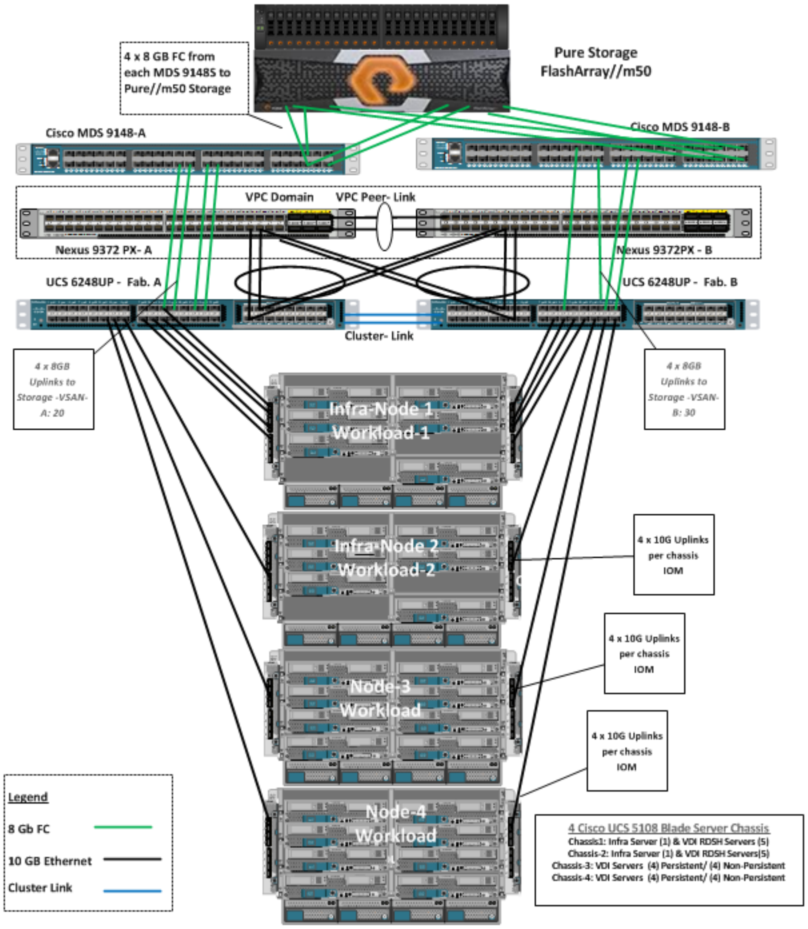
4 x 8 GB FC from each MDS 9148S to Pure//m50 Storage
Pure Storage FlashArray//m50
Cisco MDS 9148-A
Cisco MDS 9148-B
VPC Domain
VPC Peer- Link
Nexus 9372 PX- A
Nexus 9372PX - B
UCS 6248UP - Fab. A
UCS 6248UP - Fab. B
Cluster- Link
4 x 8GB Uplinks to Storage -VSAN-A: 20
4 x 8GB Uplinks to Storage -VSAN-B: 30
Infra-Node 1 Workload-1
Infra-Node 2 Workload-2
Node-3 Workload
Node-4 Workload
4 x 10G Uplinks per chassis IOM
4 x 10G Uplinks per chassis IOM
4 x 10G Uplinks per chassis IOM
Legend
8 Gb FC
10 GB Ethernet
Cluster Link
4 Cisco UCS 5108 Blade Server Chassis
Chassis1: Infra Server (1) & VDI RDSH Servers (5)
Chassis-2: Infra Server (1) & VDI RDSH Servers(5)
Chassis-3: VDI Servers (4) Persistent/ (4) Non-Persistent
Chassis-4: VDI Servers (4) Persistent/ (4) Non-Persistent

Hardware Deployed

The solution contains the following hardware as shown in Figure 18:

- Two Cisco Nexus 9372PX Layer 2 Access Switches
- Four Cisco UCS 5108 Blade Server Chassis with two UCS-IOM-2208XP IO Modules
- Two Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v4 2.0-GHz 14-core processors, 128GB 2400MHz RAM, and one Cisco VIC1340 mezzanine card for the hosted infrastructure, providing N+1 server fault tolerance
- Ten Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2680v4 2.4-GHz 14-core processors, 256GB 2400MHz RAM, and one Cisco VIC1340 mezzanine card for the XenApp Hosted Shared Desktop workload, providing N+1 server fault tolerance at the workload cluster level
- Sixteen Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2680v4 2.4-GHz 14-core processors, 512GB 2400MHz RAM, and one Cisco VIC1340 mezzanine card for the

persistent and non-persistent XenDesktop virtual desktop workload, providing N+1 server fault tolerance at the workload cluster level

- Pure Storage FlashArray//m50 dual controller storage system, one base disk shelf with 40TB raw space, one external shelf with 44TB raw space for 88TB total raw space and 8 GB ports for Fibre Channel connectivity respectively

Software Deployed

Table 2 lists the software and firmware version used in the study.

Table 2 Software and Firmware Versions

Vendor	Product	Version
Cisco	UCS Component Firmware	3.1(2b) bundle release
Cisco	UCS Manager	3.1(2b) bundle release
Cisco	UCS B200 M4 Blades	3.1(2b) bundle release
Cisco	VIC 1340	4.1(1d)
Cisco	Nexus 1000V	5.2.1
Cisco	Virtual Switch Update Manager	2.0
Cisco	UCS Performance Manager	2.0
Citrix	XenApp VDA	7.9.0.101
Citrix	XenDesktop VDA	7.9.0.101
Citrix	XenDesktop Controller	7.9.0101
Citrix	Provisioning Services	7.9.0.8201
Citrix	StoreFront Services	3.6.0.33
VMware	vCenter Server Appliance	6.0.0.10000
VMware	vSphere ESXi 6.0 Update 1a	6.0.0.4192238
Storage	Pure Storage FlashArray//m50	Purity 4.6.8

Logical Architecture

The logical architecture of this solution is designed to support up to 5000 users within four Cisco UCS 5108 Blade server chassis containing 28 blades, which provides physical redundancy for the blade servers for each workload type.

Figure 19 outlines the logical architecture of the test environment, including the Login VSI session launcher self-contained end user experience benchmarking platform

Figure 19 Logical Architecture Overview

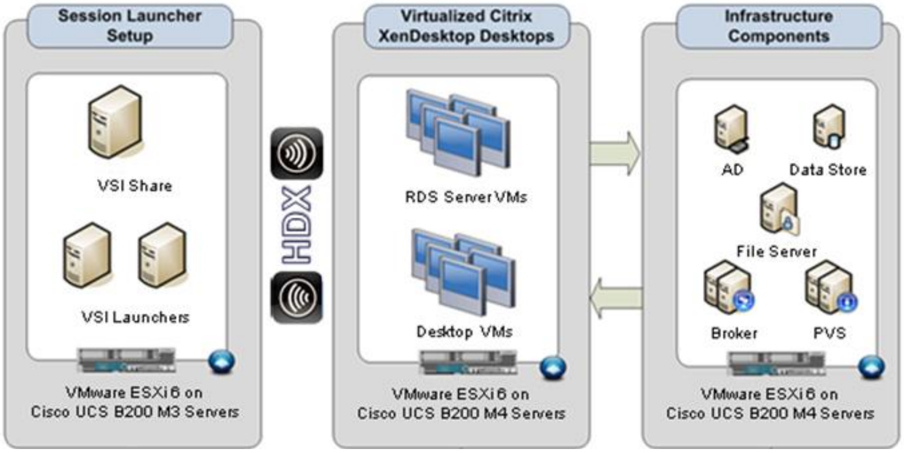
Session Launcher Setup
VSI Share
VSI Launchers
VMware ESXi 6 on Cisco UCS B200 M3 Servers
HDX
Virtualized Citrix XenDesktop Desktops
RDS Server VMs
Desktop VMs
VMware ESXi 6 on Cisco UCS B200 M4 Servers
Infrastructure Components
AD
Data Store
File Server
Broker
PVS
VMware ESXi 6 on Cisco UCS B200 M4 Servers

This document is intended to allow you to fully configure your environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. 0 through 0 lists the information you need to configure your environment.

Figure 20 identifies the server roles in the 28 server deployment to support the 5000 seat workload. We also break out the infrastructure virtual machine fault tolerant design.

Figure 20 Server, Location, and Purpose

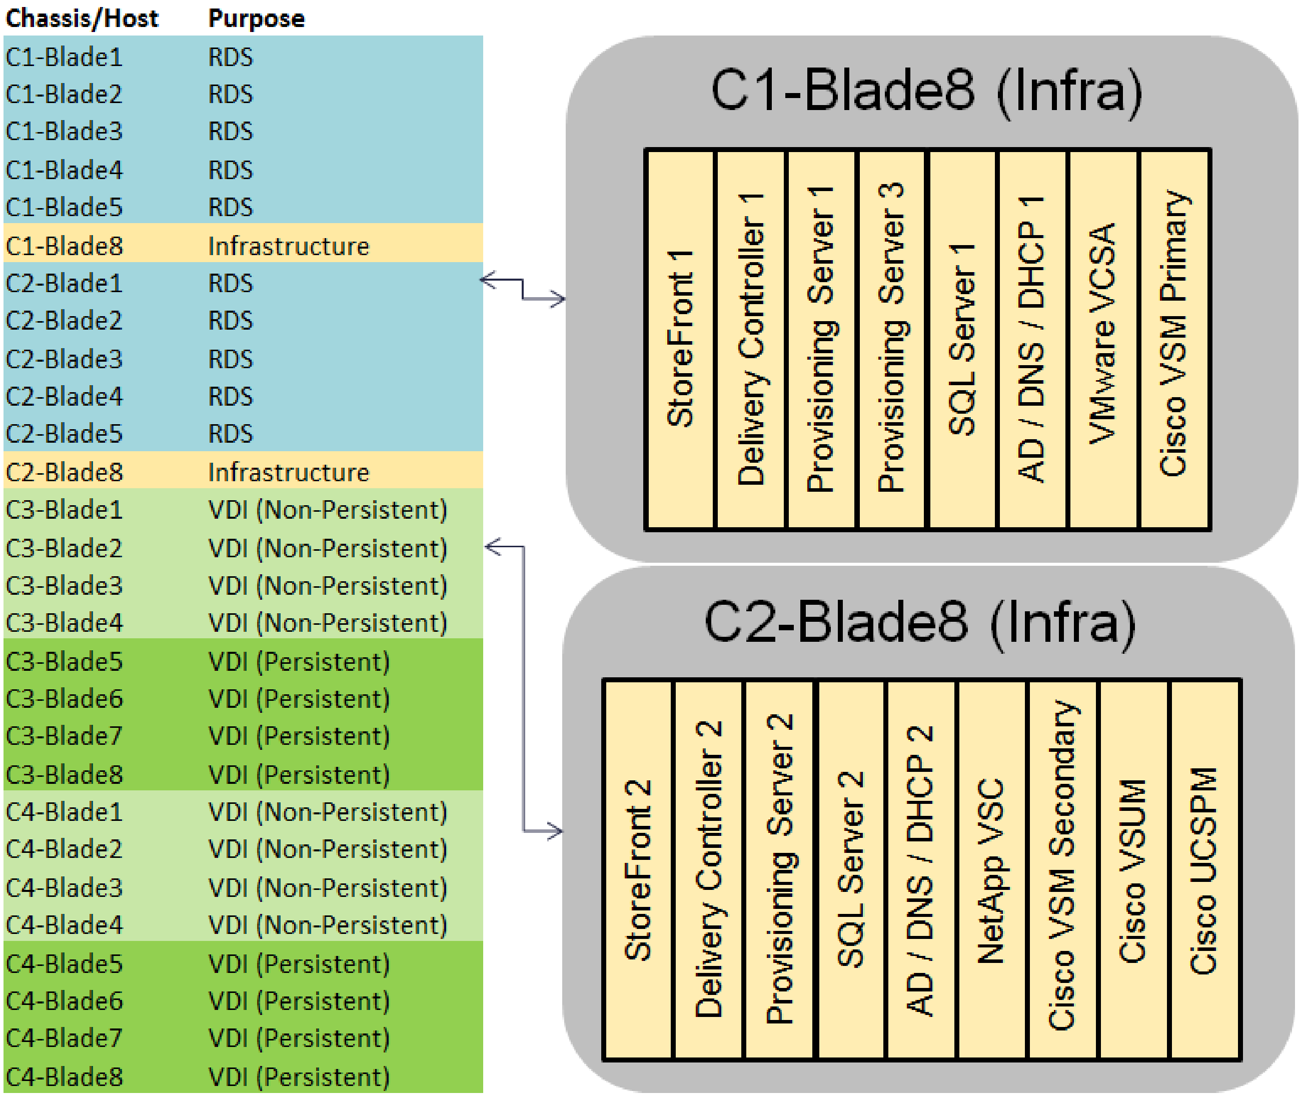

Chassis/Host	Purpose
C1-Blade1	RDS
C1-Blade2	RDS
C1-Blade3	RDS
C1-Blade4	RDS
C1-Blade5	RDS
C1-Blade8	Infrastructure
C2-Blade1	RDS
C2-Blade2	RDS
C2-Blade3	RDS
C2-Blade4	RDS
C2-Blade5	RDS
C2-Blade8	Infrastructure
C3-Blade1	VDI (Non-Persistent)
C3-Blade2	VDI (Non-Persistent)
C3-Blade3	VDI (Non-Persistent)
C3-Blade4	VDI (Non-Persistent)
C3-Blade5	VDI (Persistent)
C3-Blade6	VDI (Persistent)
C3-Blade7	VDI (Persistent)
C3-Blade8	VDI (Persistent)
C4-Blade1	VDI (Non-Persistent)
C4-Blade2	VDI (Non-Persistent)
C4-Blade3	VDI (Non-Persistent)
C4-Blade4	VDI (Non-Persistent)
C4-Blade5	VDI (Persistent)
C4-Blade6	VDI (Persistent)
C4-Blade7	VDI (Persistent)
C4-Blade8	VDI (Persistent)

The following table outlines the virtual machine deployments on the hardware platform.

Table 3 Virtual Machine Deployment Architecture

Server Name	Location	Purpose
C1-Blade8, C2-Blade8	Physical – Chassis 1, 2	ESXi 6.0 Hosts Infrastructure VMs Windows 2012-R2, VCSA, VSM, VSUM
C1-Blade1-6, C2-Blade1-6	Physical – Chassis 1, 2	ESXi 6.0 Hosts 96x XenApp RDS VMs
C3-Blade1-3, C4-Blade1-4	Physical – Chassis 3, 4	ESXi 6.0 Hosts 1200x XenDesktop VDI (Non-Persistent) VMs
C3-Blade4-6, 8 C4-Blade5, 6, 8	Physical – Chassis 3, 4	ESXi 6.0 Hosts 1200x XenDesktop VDI (Persistent) VMs

Server Name	Location	Purpose
CTX-SF1	C1-Blade8	Citrix StoreFront Server 1
CTX-XD1	C1-Blade8	XenDesktop Controller 1, Studio, Licensing
CTX-PVS1	C1-Blade8	Provisioning Services streaming server 1
CTX-PVS3	C1-Blade8	Provisioning Services streaming server 3
SQL1	C1-Blade8	SQL Server 1 (Always On)
AD-DC1	C1-Blade8	Active Directory Domain Controller 1
VCSA	C1-Blade8	VMware vCenter Server Appliance
VSM_primary	C1-Blade8	Cisco Virtual Supervisor Module
CTX-SF2	C2-Blade8	Citrix StoreFront Server 2
CTX-XD2	C2-Blade8	XenDesktop Controller 2, Director
CTX-PVS2	C2-Blade8	Provisioning Services streaming server 2
SQL2	C2-Blade8	SQL Server 2 (Always On)
AD-DC2	C2-Blade8	Active Directory Domain Controller 2
VSM_secondary	C2-Blade8	Cisco Virtual Supervisor Module
VSUM	C2-Blade8	Cisco Virtual Switch Update Manager
UCSPM	C2-Blade8	Cisco UCS Performance Manager

VLANs

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

Table 4 VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	160	VLAN for in-band management interfaces
Infra-Mgmt	161	VLAN for Virtual Infrastructure

VLAN Name	VLAN ID	VLAN Purpose
VDI	162	VLAN for VDI Traffic
vMotion	166	VLAN for VMware vMotion
OB-Mgmt	164	VLAN for out-of-band management interfaces

VSANS

We utilized two virtual SANs for communications and fault tolerance in this design:

Table 5 VASNs Configured in this Study

VSAN Name	VSAN Purpose	ID Used in Validating This Document
VSAN 1	VSAN for primary SAN communication	20
VSAN 2	VSAN for secondary SAN communication	30

VMware Clusters

The following four VMware Clusters were used in one vCenter data center to support the solution and testing environment:

- VDI Cluster Pure Storage Data Center with Cisco UCS
  - Infrastructure Cluster: Infra VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, XenDesktop Controllers, Provisioning Servers, Nexus 1000v Virtual Supervisor Modules (VSMs, etc.)
  - RDS: Citrix XenApp RDS VMs (Windows Server 2012 R2)
  - VDI Non-Persistent: Citrix XenDesktop VDI VMs (Windows 10 64-bit non-persistent virtual desktops provisioned with Citrix PVS
  - VDI Persistent: Citrix XenDesktop VDI VMs (Windows 10 64-bit persistent virtual desktops provisioned with Citrix MCS
- VSI Launchers Cluster
  - Launcher Cluster 1 and 2: Login VSI Cluster (The Login VSI launcher infrastructure was connected using the same set of switches and vCenter instance, but was hosted on separate local storage and servers.)

Figure 21 VMware vSphere Clusters on vSphere Web GUI

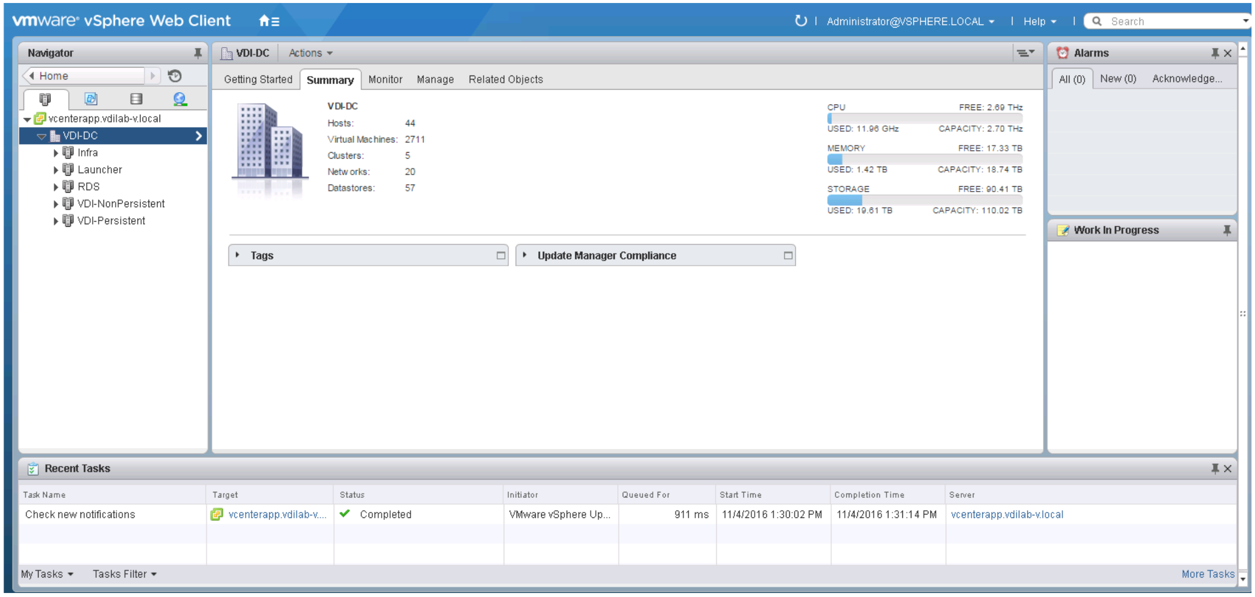

VDI-DC

Hosts:	44
Virtual Machines:	2711
Clusters:	5
Netw orks:	20
Datastores:	57

CPU FREE: 2.69 THz
USED: 11.96 GHz CAPACITY: 2.70 THz

MEMORY FREE: 17.33 TB
USED: 1.42 TB CAPACITY: 18.74 TB

STORAGE FREE: 90.41 TB
USED: 19.61 TB CAPACITY: 110.02 TB

Recent Tasks

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server
Check new notifications	vcenterapp.vdilab-v....	✓ Completed	VMware vSphere Up...	911 ms	11/4/2016 1:30:02 PM	11/4/2016 1:31:14 PM	vcenterapp.vdilab-v.local

My Tasks Tasks Filter More Tasks

Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution. Figure 22 illustrates the configuration topology for this solution.

Configuration Topology for Scalable Citrix XenDesktop Mixed Workload

Component Layers

Figure 22 Solution Component Layers

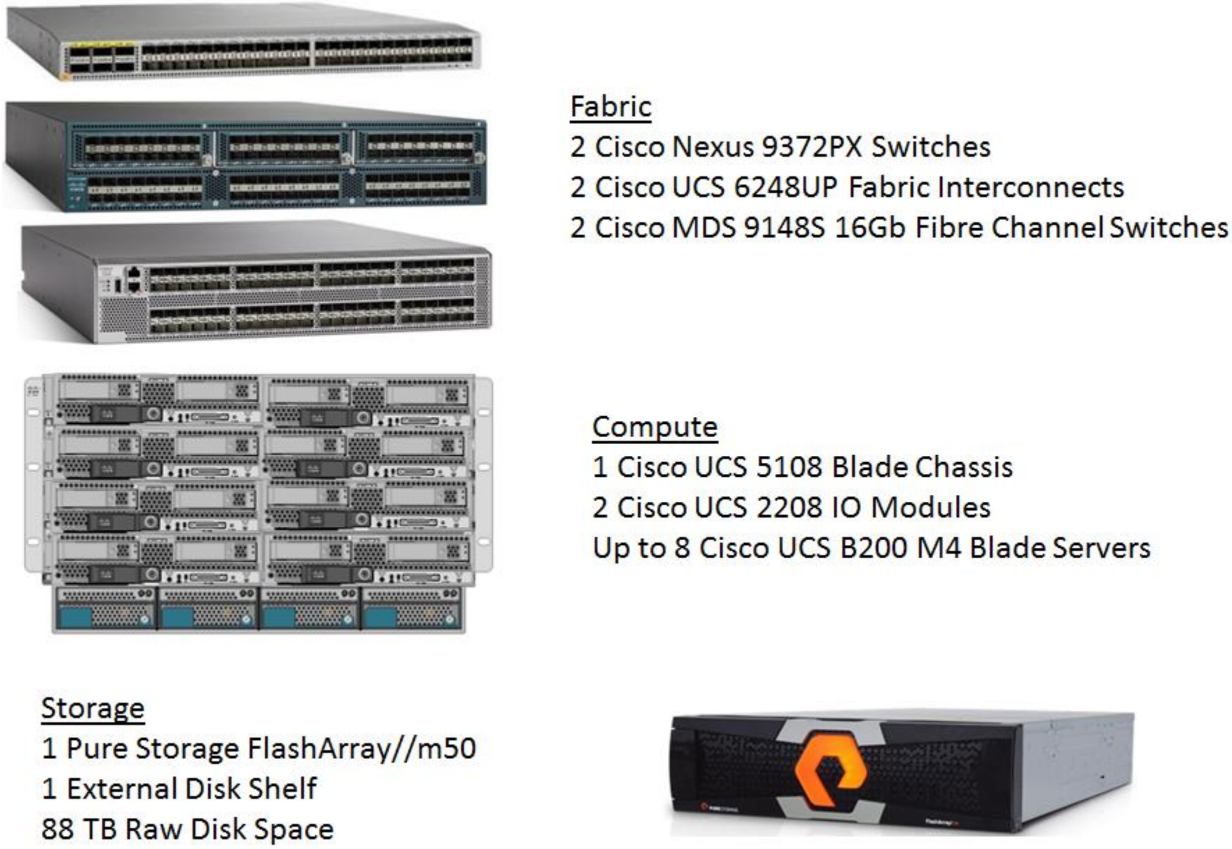
Fabric
2 Cisco Nexus 9372PX Switches
2 Cisco UCS 6248UP Fabric Interconnects
2 Cisco MDS 9148S 16Gb Fibre Channel Switches
Compute
1 Cisco UCS 5108 Blade Chassis
2 Cisco UCS 2208 IO Modules
Up to 8 Cisco UCS B200 M4 Blade Servers
Storage
1 Pure Storage FlashArray//m50
1 External Disk Shelf
88 TB Raw Disk Space

Figure 22 above 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 Pure Storage FlashArray//m50

Solution Cabling

The following subsections detail the physical connectivity configuration of the FlashStack 5000 seat Citrix XenDesktop 7.9 environment.

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

The tables in this section contain the details for the prescribed and supported configuration of the Pure Storage FlashArray//m 50 to the Cisco 6248UP Fabric Interconnects via Cisco MDS 9148 FC switches.

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

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

Figure 23 shows a cabling diagram for a Citrix XenDesktop configuration using the Cisco Nexus 9000, Cisco MDS 9400 Series and Pure Storage FlashArray//m50. The Pure Storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves.

Figure 23 FlashStack 5000 Seat Cabling Diagram

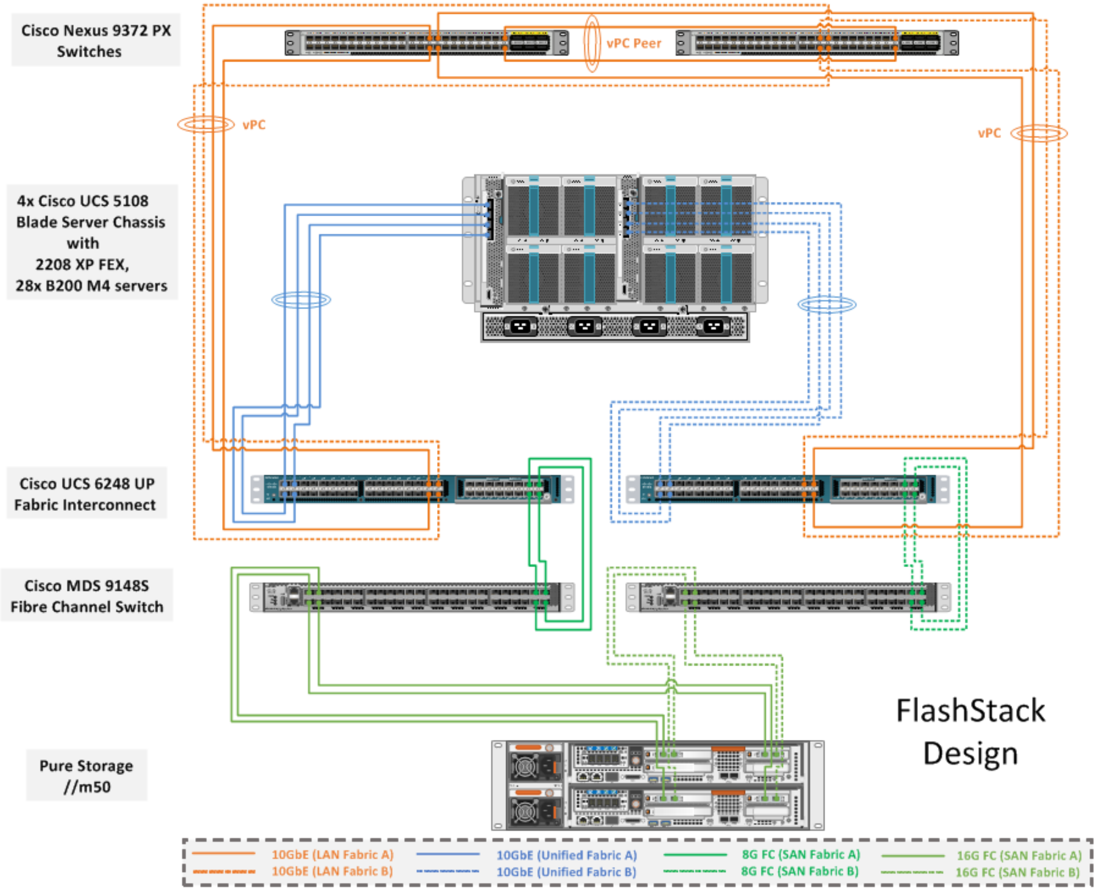

Cisco Nexus Switch Cabling Details

Table 6 Cisco Nexus 9372-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	1GbE	Pure Storage FlashArray//m50 (Console Mgmt.)	e0e
	Eth1/29	10GbE	Cisco UCS fabric interconnect A	Eth1/29
	Eth1/30	10GbE	Cisco UCS fabric interconnect A	Eth1/30
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/29
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/30

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/47	40GbE	Cisco Nexus 9372 B	Eth1/47
	Eth1/48	40GbE	Cisco Nexus 9372 B	Eth1/48
	MGMT0	GbE	GbE management switch	Any

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

Table 7 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/1	1GbE	Pure Storage FlashArray//m50 (Console Mgmt.)	e0g
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/31
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/32
	Eth1/29	10GbE	Cisco UCS fabric interconnect A	Eth1/31
	Eth1/30	10GbE	Cisco UCS fabric interconnect A	Eth1/32
	Eth1/47	40GbE	Cisco Nexus 9372 A	Eth1/47
	Eth1/48	40GbE	Cisco Nexus 9372 A	Eth1/48
	MGMT0	GbE	GbE management switch	Any

Cisco UCS 6248UP Fabric Interconnect Cabling

Table 8 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/29	10GbE	Cisco Nexus 9372 A	Eth1/29
	Eth1/30	10GbE	Cisco Nexus 9372 A	Eth1/30
	Eth1/31	10GbE	Cisco Nexus 9372 B	Eth1/29
	Eth1/32	10 GbE	Cisco Nexus 9372 B	Eth 1/30
	Eth1/1-1/8	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 1-2	IOM 1-4

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/11-1/14	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 3	IOM 1-4
	Eth 1/17-1/20	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 4	IOM 1-4
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2
	FC 2/13	8Gb FC	Cisco MDS 9148-A	FC 1/5
	FC 2/14	8Gb FC	Cisco MDS 9148-A	FC 1/6
	FC 2/15	8Gb FC	Cisco MDS 9148-A	FC 1/7
	FC 2/16	8Gb FC	Cisco MDS 9148-A	FC 1/8

Table 9 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/29	10GbE	Cisco Nexus 9372 B	Eth1/29
	Eth1/30	10GbE	Cisco Nexus 9372 B	Eth1/30
	Eth1/31	10GbE	Cisco Nexus 9372 A	Eth1/31
	Eth1/32	10GbE	Cisco Nexus 9372 A	Eth1/32
	Eth1/1-1/8	10GbE	UCS 5108 Blade Chassis IOM-B, Chassis 1-2	IOM 1-4
	Eth1/11-1/14	10GbE	UCS 5108 Blade Chassis IOM-B, Chassis 3	IOM 1-4

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth 1/17-1/20	10GbE	UCS 5108 Blade Chassis IOM-B, Chassis 4	IOM 1-4
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2
	FC 2/13	8Gb FC	Cisco MDS 9148-B	FC 1/5
	FC 2/14	8Gb FC	Cisco MDS 9148-B	FC 1/6
	FC 2/15	8Gb FC	Cisco MDS 9148-B	FC 1/7
	FC 2/16	8Gb FC	Cisco MDS 9148-B	FC 1/8

Figure 24 Cable Connectivity Between Cisco Nexus 9372PX, Cisco UCS 6248UP Fabric Interconnects and Cisco 2208 IO Modules in Cisco 5108AC Blade Chassis

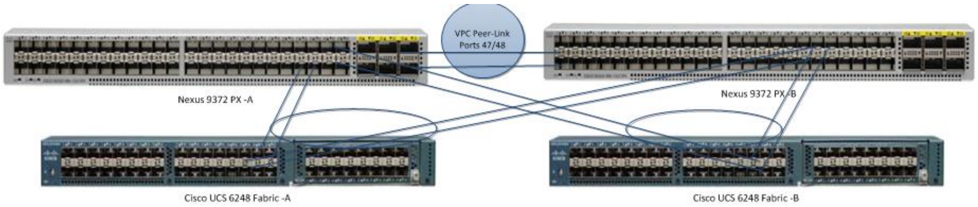

Cisco MDS 9148S Cabling

Figure 24 illustrates cable connectivity between the Cisco MDS 9148S and the Cisco 6248 Fabric Interconnects and the Pure Storage FlashArray//m50.

We used two 8Gb FC connections from each Fabric Interconnect to each MDS switch.

We utilized two 16Gb FC connections from each Pure Storage FlashArray//m50 controller to each MDS switch.

Table 10 Cisco MDS 9148S A Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148S-A	fc1/1	16Gb FC	Pure Storage FlashArray//m50	fc1/0
	fc1/2	16Gb FC	Pure Storage FlashArray//m50	fc1/1
	fc1/3	16Gb FC	Pure Storage FlashArray//m50	fc1/2
	fc1/4	16Gb FC	Pure Storage FlashArray//m50	fc1/3

Local Device	Local Port	Connection	Remote Device	Remote Port
	fc1/5	8Gb FC	Cisco 6248UP Fabric Interconnect-A	fc2/13
	fc1/6	8Gb FC	Cisco 6248UP Fabric Interconnect-A	fc2/14
	fc1/7	8Gb FC	Cisco 6248UP Fabric Interconnect-A	fc2/15
	fc1/8	8Gb FC	Cisco 6248UP Fabric Interconnect-A	fc2/16

Table 11 Cisco MDS 9148S B Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148S-B	fc1/1	16Gb FC	Pure Storage FlashArray//m50	fc2/0
	fc1/2	16Gb FC	Pure Storage FlashArray//m50	fc2/1
	fc1/3	16Gb FC	Pure Storage FlashArray//m50	fc2/2
	fc1/4	16Gb FC	Pure FlashArray//m50	fc2/3
	fc1/5	8Gb FC	Cisco 6248UP Fabric Interconnect-B	fc2/13
	fc1/6	8Gb FC	Cisco 6248UP Fabric Interconnect-B	fc2/14
	fc1/7	8Gb FC	Cisco 6248UP Fabric Interconnect-B	fc2/15
	fc1/8	8Gb FC	Cisco 6248UP Fabric Interconnect-B	fc2/16

FlashStack Fabric to FlashArray//m50 Cabling

Figure 25 FlashArray//m50 Controller-A and B Connection to MDS A and B Switches using VSAN 20 for Fabric A and VSAN 30 Configured for Fabric B Side

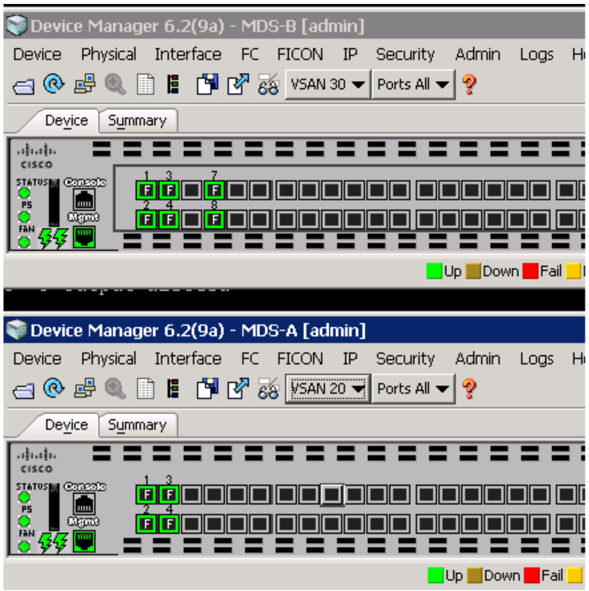
Device Manager 6.2(9a) - MDS-B [admin]

Device Manager 6.2(9a) - MDS-A [admin]

Figure 26 Fibre Channel Cable Connectivity from Pure FlashArray//m50 to Cisco MDS 9148S to Cisco 6248 Fabric Interconnects

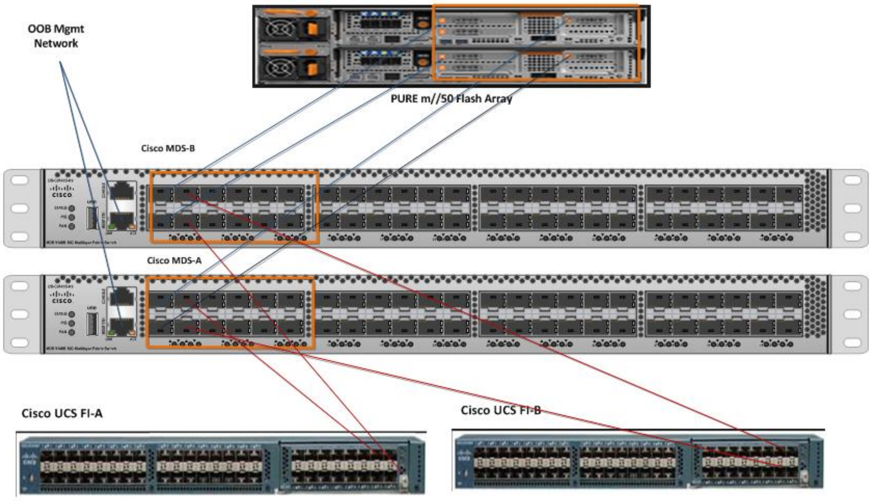

Cisco Unified Computing System Base Configuration

This section details the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the install 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 3.1(2b)

This document assumes the use of Cisco UCS Manager Software version 3.1(2b). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 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 interconnect was previously deployed and you want to erase it to redeploy, follow these steps:

   a. Login with the existing user name 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.

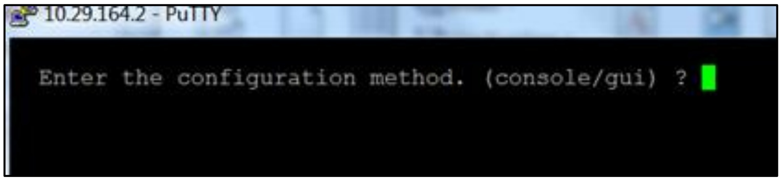

```
 Enter the configuration method. (console/gui) ?
```

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

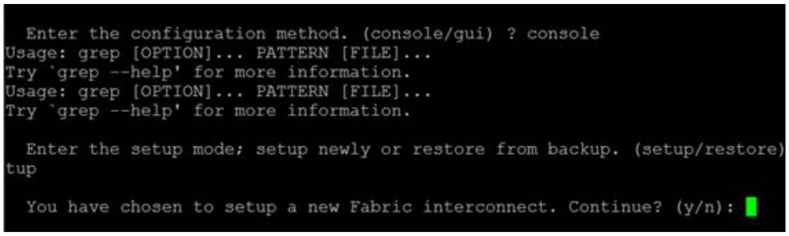

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

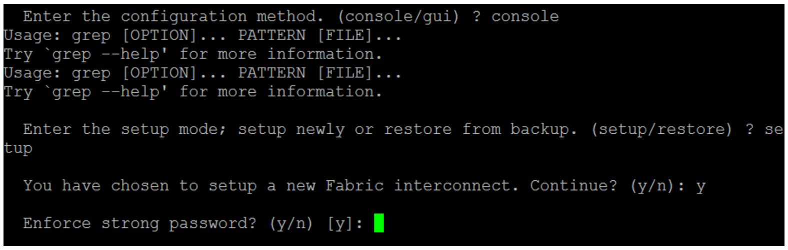

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

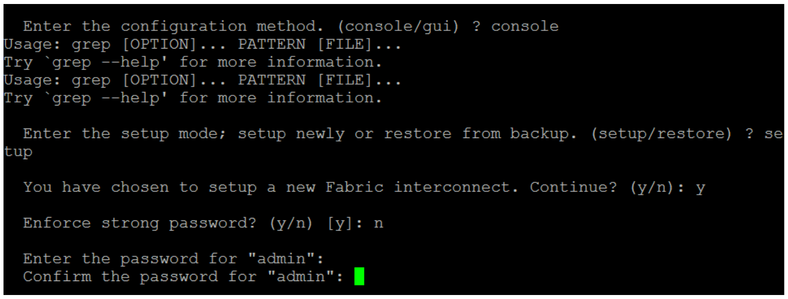

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

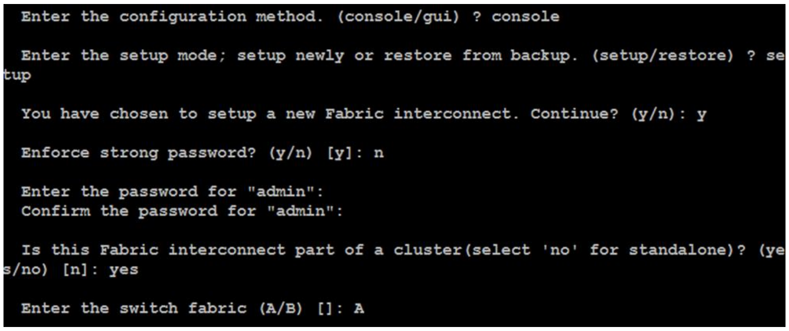

```
  Enter the configuration method. (console/gui) ? console

  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.

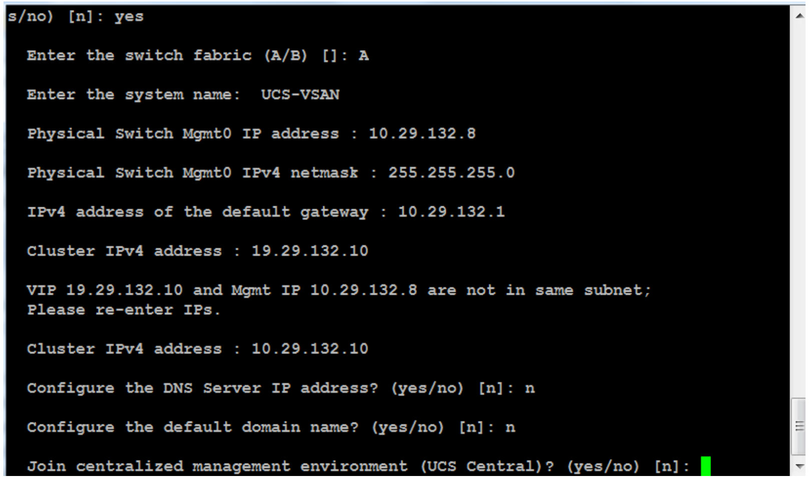

```
s/no) [n]: yes

  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”.

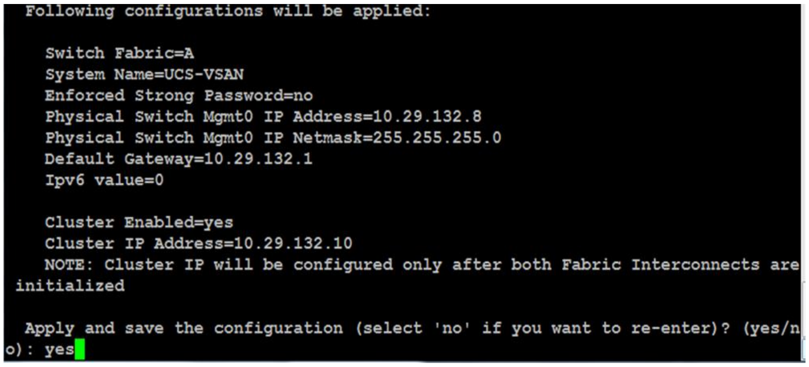

```
 Following configurations will be applied:

   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 onto second fabric interconnect, select console as the configuration method and provide the following inputs.

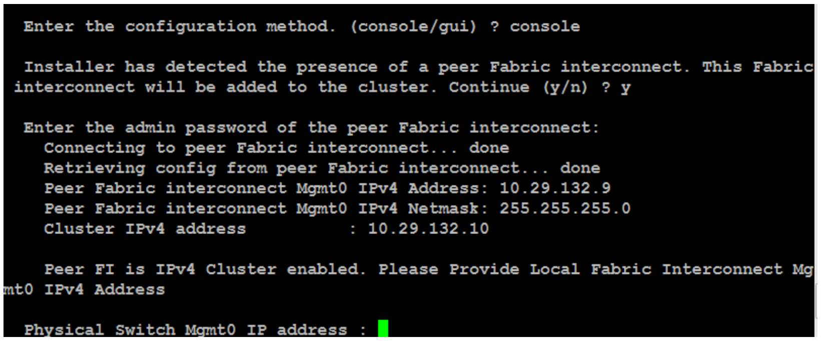

```
 Enter the configuration method. (console/gui) ? console

 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.

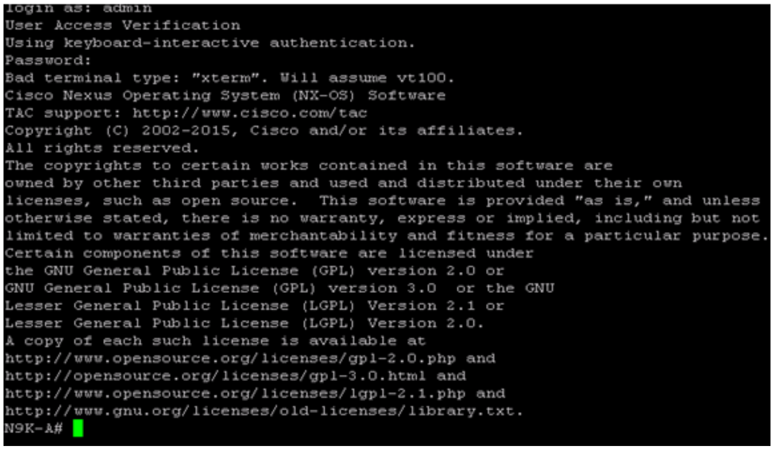

```
login as: admin
User Access Verification
Using keyboard-interactive authentication.
Password:
Bad terminal type: "xterm". Will assume vt100.
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2015, Cisco and/or its affiliates.
All rights reserved.
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 6248 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 user name 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.

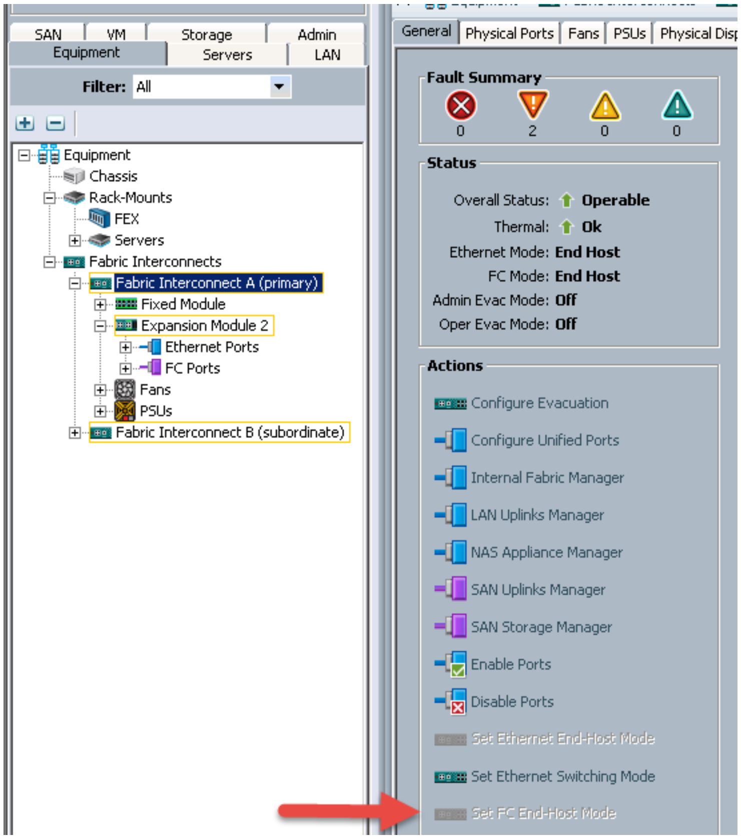

Overall Status: Operable
Thermal: Ok
Ethernet Mode: End Host
FC Mode: End Host
Admin Evac Mode: Off
Oper Evac Mode: Off

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.

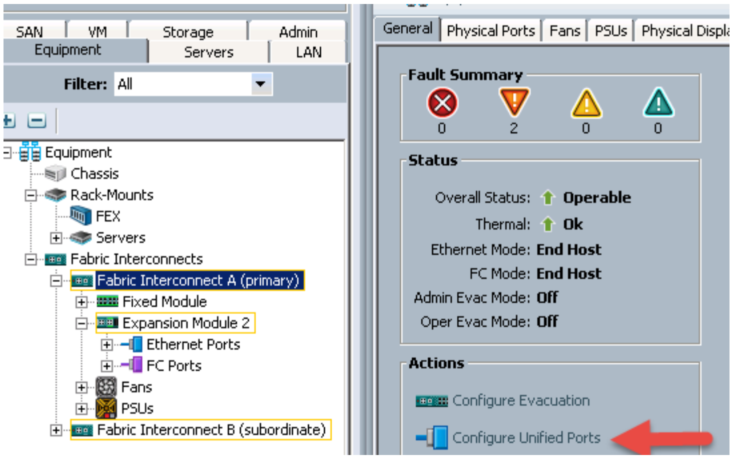

Status

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

3. Click Configure Expansion Module Ports.

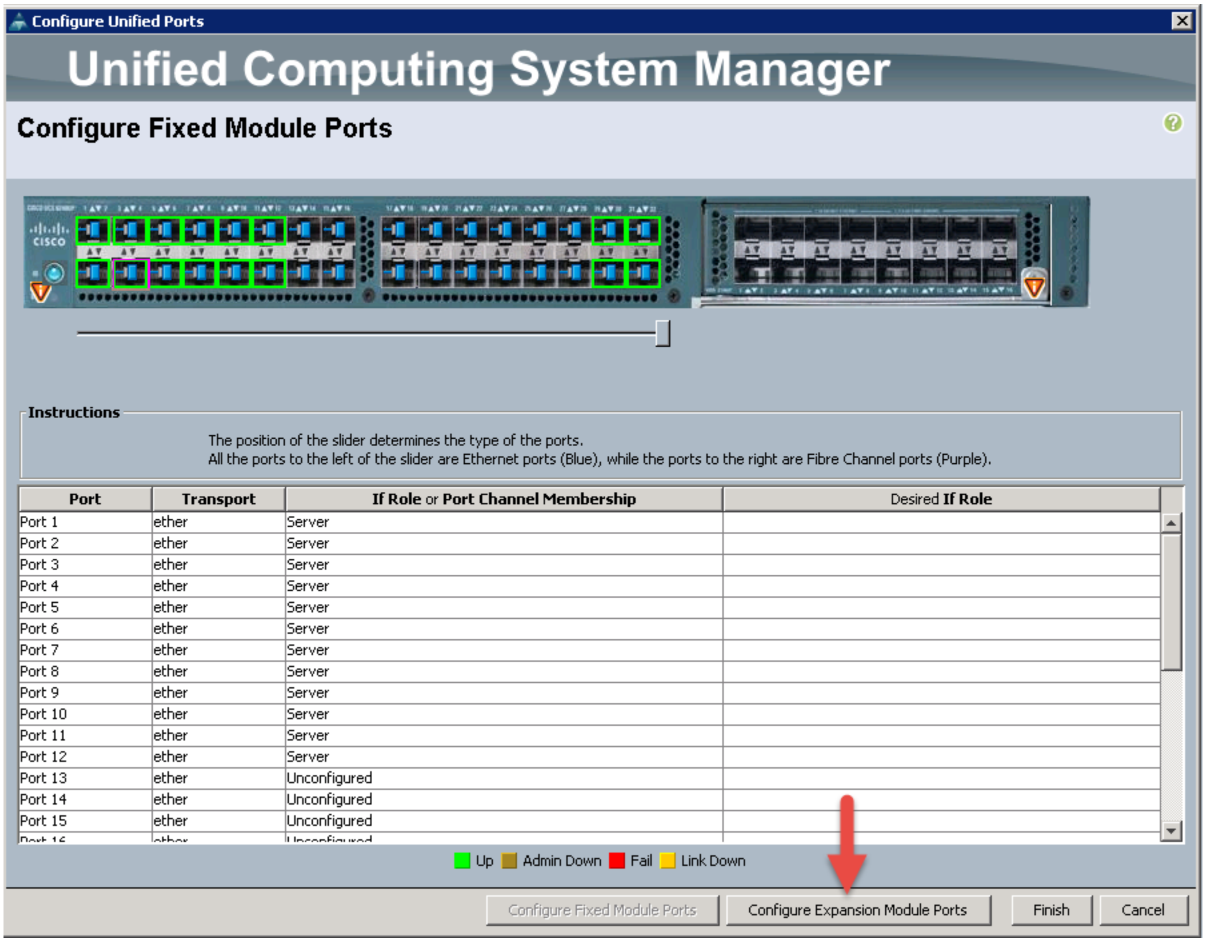
Configure Unified Ports
Unified Computing System Manager
Configure Fixed Module Ports
CISCO

Instructions

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

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

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down

Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

4. Move the slider to the left.

Ports to the right of the slider will become FC ports. For our study, we configured the last four ports on the Expansion Module as FC ports.

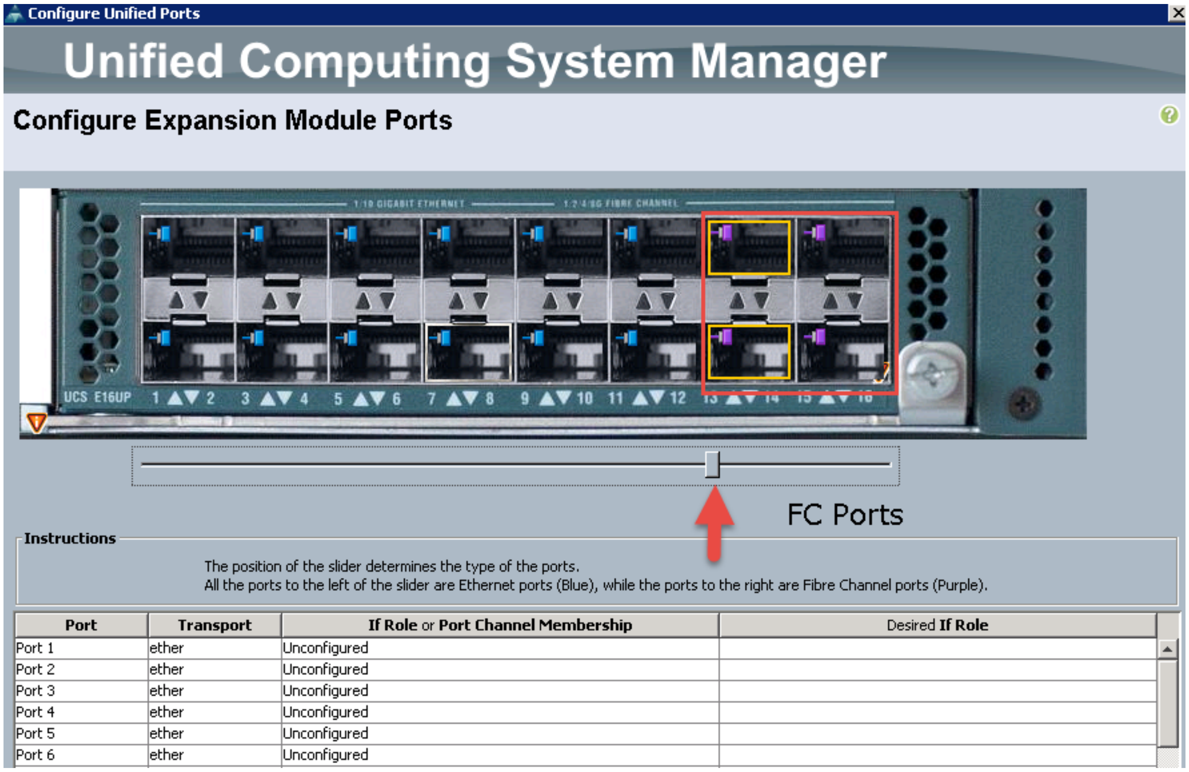

Configure Unified Ports

Unified Computing System Manager

Configure Expansion Module Ports

Instructions

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

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	
Port 2	ether	Unconfigured	
Port 3	ether	Unconfigured	
Port 4	ether	Unconfigured	
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	

5. Click Finish, then click Yes to confirm. This action will cause a reboot of the Expansion Module.

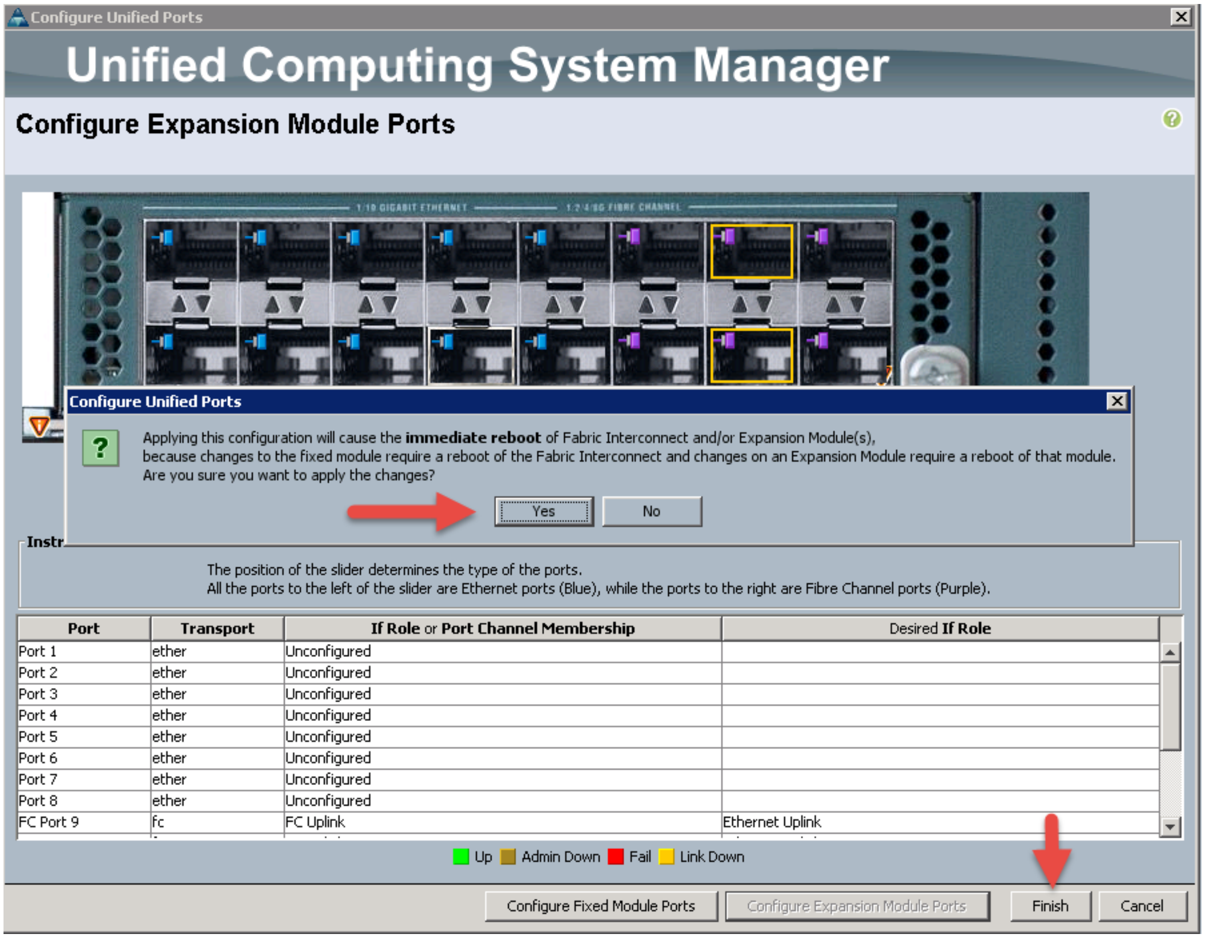

Configure Unified Ports

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?

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

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	
Port 2	ether	Unconfigured	
Port 3	ether	Unconfigured	
Port 4	ether	Unconfigured	
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
FC Port 9	fc	FC Uplink	Ethernet Uplink

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down

Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

After the expansion module reboot, your FC Ports configuration should look like the figure below:

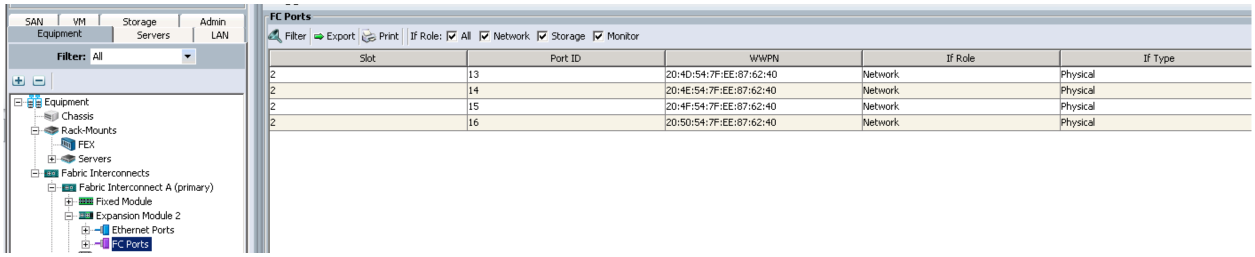

FC Ports

If Role: ☑ All ☑ Network ☑ Storage ☑ Monitor

Slot	Port ID	WWPN	If Role	If Type
2	13	20:4D:54:7F:EE:87:62:40	Network	Physical
2	14	20:4E:54:7F:EE:87:62:40	Network	Physical
2	15	20:4F:54:7F:EE:87:62:40	Network	Physical
2	16	20:50:54:7F:EE:87:62:40	Network	Physical

6. Repeat this procedure for Fabric Interconnect B.

7. Insert Cisco SFP 8 Gbps FC (DS-SFP-FC8-SW) modules into ports 13 through 16 on both Fabric Interconnects and cable as prescribed later in this document.

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 4-link.

4. Set the Link Grouping Preference to Port Channel.

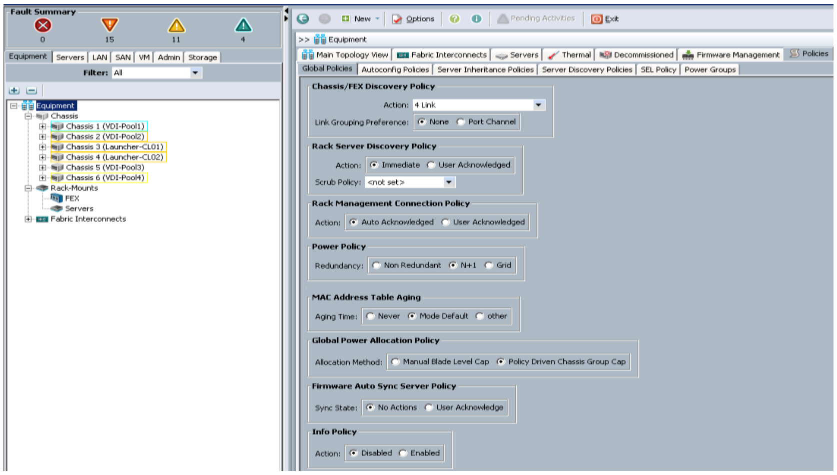

Chassis/FEX Discovery Policy

Action: 4 Link

Link Grouping Preference: ● None ○ Port Channel

Rack Server Discovery Policy

Action: ● Immediate ○ User Acknowledged

Scrub Policy: <not set>

Rack Management Connection Policy

Action: ● Auto Acknowledged ○ User Acknowledged

Power Policy

Redundancy: ○ Non Redundant ● N+1 ○ Grid

MAC Address Table Aging

Aging Time: ○ Never ● Mode Default ○ other

Global Power Allocation Policy

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

Firmware Auto Sync Server Policy

Sync State: ● No Actions ○ User Acknowledge

Info Policy

Action: ● Disabled ○ Enabled

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.

Figure 27 Acknowledge Cisco UCS Chassis

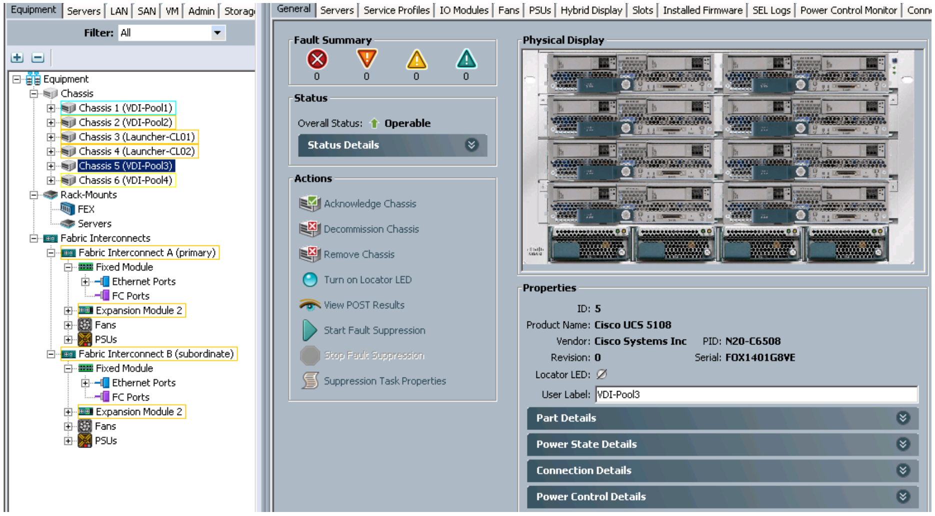

Overall Status: Operable

ID: 5

Product Name: Cisco UCS 5108

Vendor: Cisco Systems Inc PID: N20-C6508

Revision: 0 Serial: FOX1401G8VE

Locator LED:

User Label: VDI-Pool3

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

5. Repeat for each of the remaining chassis.

Add a Block of IP Addresses for Out-of-Band KVM Access

To create a block of IP addresses for server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:

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

2. Select Pools > root > IP Pools > IP Pool ext-mgmt.

3. In the Actions pane, select Create Block of IP Addresses.

4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

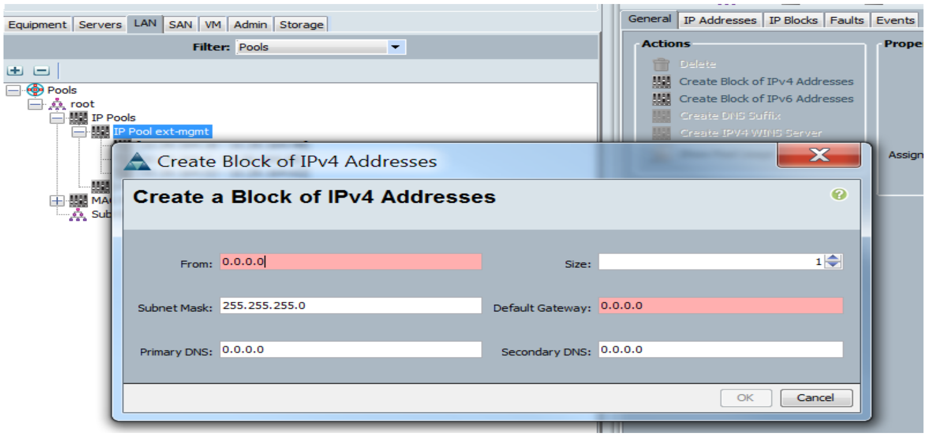

Create a Block of IPv4 Addresses

Create a Block of IPv4 Addresses

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

5. Click OK to create the IP block.

6. Click OK in the confirmation message.

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.

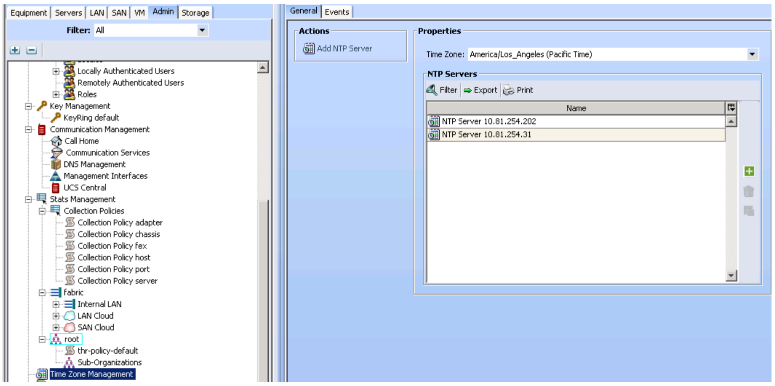

Time Zone: America/Los_Angeles (Pacific Time)

Name
NTP Server 10.81.254.202
NTP Server 10.81.254.31

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 > Fixed Module.

3. Expand Ethernet Ports.

4. Select ports 1 through 28 (Ports highlighted 9, 10, 15 & 16 not configured) 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.

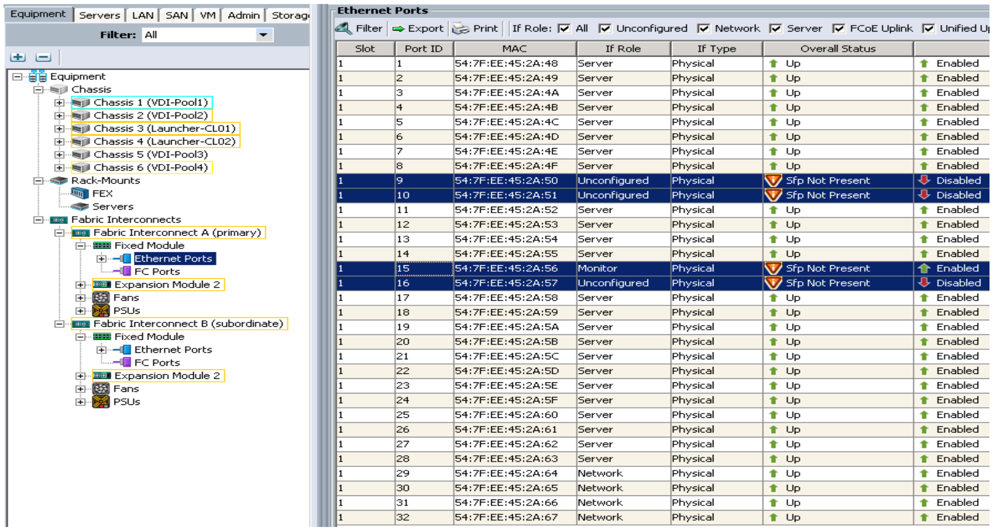

If Role: ☑ All ☑ Unconfigured ☑ Network ☑ Server ☑ FCoE Uplink ☑ Unified U

Slot	Port ID	MAC	If Role	If Type	Overall Status	
1	1	54:7F:EE:45:2A:48	Server	Physical	↑ Up	↑ Enabled
1	2	54:7F:EE:45:2A:49	Server	Physical	↑ Up	↑ Enabled
1	3	54:7F:EE:45:2A:4A	Server	Physical	↑ Up	↑ Enabled
1	4	54:7F:EE:45:2A:4B	Server	Physical	↑ Up	↑ Enabled
1	5	54:7F:EE:45:2A:4C	Server	Physical	↑ Up	↑ Enabled
1	6	54:7F:EE:45:2A:4D	Server	Physical	↑ Up	↑ Enabled
1	7	54:7F:EE:45:2A:4E	Server	Physical	↑ Up	↑ Enabled
1	8	54:7F:EE:45:2A:4F	Server	Physical	↑ Up	↑ Enabled
1	9	54:7F:EE:45:2A:50	Unconfigured	Physical	▽ Sfp Not Present	↓ Disabled
1	10	54:7F:EE:45:2A:51	Unconfigured	Physical	▽ Sfp Not Present	↓ Disabled
1	11	54:7F:EE:45:2A:52	Server	Physical	↑ Up	↑ Enabled
1	12	54:7F:EE:45:2A:53	Server	Physical	↑ Up	↑ Enabled
1	13	54:7F:EE:45:2A:54	Server	Physical	↑ Up	↑ Enabled
1	14	54:7F:EE:45:2A:55	Server	Physical	↑ Up	↑ Enabled
1	15	54:7F:EE:45:2A:56	Monitor	Physical	▽ Sfp Not Present	↑ Enabled
1	16	54:7F:EE:45:2A:57	Unconfigured	Physical	▽ Sfp Not Present	↓ Disabled
1	17	54:7F:EE:45:2A:58	Server	Physical	↑ Up	↑ Enabled
1	18	54:7F:EE:45:2A:59	Server	Physical	↑ Up	↑ Enabled
1	19	54:7F:EE:45:2A:5A	Server	Physical	↑ Up	↑ Enabled
1	20	54:7F:EE:45:2A:5B	Server	Physical	↑ Up	↑ Enabled
1	21	54:7F:EE:45:2A:5C	Server	Physical	↑ Up	↑ Enabled
1	22	54:7F:EE:45:2A:5D	Server	Physical	↑ Up	↑ Enabled
1	23	54:7F:EE:45:2A:5E	Server	Physical	↑ Up	↑ Enabled
1	24	54:7F:EE:45:2A:5F	Server	Physical	↑ Up	↑ Enabled
1	25	54:7F:EE:45:2A:60	Server	Physical	↑ Up	↑ Enabled
1	26	54:7F:EE:45:2A:61	Server	Physical	↑ Up	↑ Enabled
1	27	54:7F:EE:45:2A:62	Server	Physical	↑ Up	↑ Enabled
1	28	54:7F:EE:45:2A:63	Server	Physical	↑ Up	↑ Enabled
1	29	54:7F:EE:45:2A:64	Network	Physical	↑ Up	↑ Enabled
1	30	54:7F:EE:45:2A:65	Network	Physical	↑ Up	↑ Enabled
1	31	54:7F:EE:45:2A:66	Network	Physical	↑ Up	↑ Enabled
1	32	54:7F:EE:45:2A:67	Network	Physical	↑ Up	↑ Enabled

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

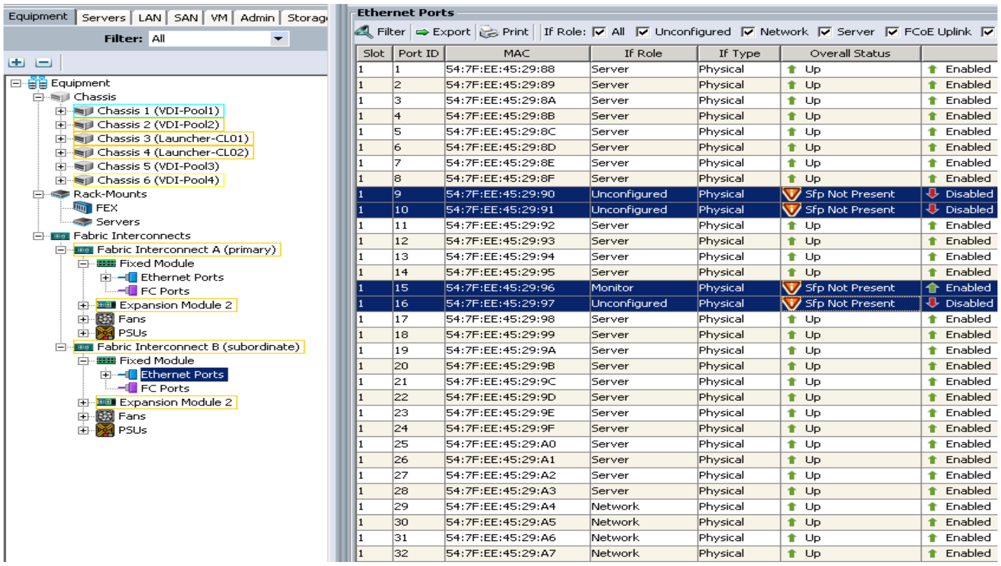

Slot	Port ID	MAC	If Role	If Type	Overall Status	
1	1	54:7F:EE:45:29:88	Server	Physical	Up	Enabled
1	2	54:7F:EE:45:29:89	Server	Physical	Up	Enabled
1	3	54:7F:EE:45:29:8A	Server	Physical	Up	Enabled
1	4	54:7F:EE:45:29:8B	Server	Physical	Up	Enabled
1	5	54:7F:EE:45:29:8C	Server	Physical	Up	Enabled
1	6	54:7F:EE:45:29:8D	Server	Physical	Up	Enabled
1	7	54:7F:EE:45:29:8E	Server	Physical	Up	Enabled
1	8	54:7F:EE:45:29:8F	Server	Physical	Up	Enabled
1	9	54:7F:EE:45:29:90	Unconfigured	Physical	Sfp Not Present	Disabled
1	10	54:7F:EE:45:29:91	Unconfigured	Physical	Sfp Not Present	Disabled
1	11	54:7F:EE:45:29:92	Server	Physical	Up	Enabled
1	12	54:7F:EE:45:29:93	Server	Physical	Up	Enabled
1	13	54:7F:EE:45:29:94	Server	Physical	Up	Enabled
1	14	54:7F:EE:45:29:95	Server	Physical	Up	Enabled
1	15	54:7F:EE:45:29:96	Monitor	Physical	Sfp Not Present	Enabled
1	16	54:7F:EE:45:29:97	Unconfigured	Physical	Sfp Not Present	Disabled
1	17	54:7F:EE:45:29:98	Server	Physical	Up	Enabled
1	18	54:7F:EE:45:29:99	Server	Physical	Up	Enabled
1	19	54:7F:EE:45:29:9A	Server	Physical	Up	Enabled
1	20	54:7F:EE:45:29:9B	Server	Physical	Up	Enabled
1	21	54:7F:EE:45:29:9C	Server	Physical	Up	Enabled
1	22	54:7F:EE:45:29:9D	Server	Physical	Up	Enabled
1	23	54:7F:EE:45:29:9E	Server	Physical	Up	Enabled
1	24	54:7F:EE:45:29:9F	Server	Physical	Up	Enabled
1	25	54:7F:EE:45:29:A0	Server	Physical	Up	Enabled
1	26	54:7F:EE:45:29:A1	Server	Physical	Up	Enabled
1	27	54:7F:EE:45:29:A2	Server	Physical	Up	Enabled
1	28	54:7F:EE:45:29:A3	Server	Physical	Up	Enabled
1	29	54:7F:EE:45:29:A4	Network	Physical	Up	Enabled
1	30	54:7F:EE:45:29:A5	Network	Physical	Up	Enabled
1	31	54:7F:EE:45:29:A6	Network	Physical	Up	Enabled
1	32	54:7F:EE:45:29:A7	Network	Physical	Up	Enabled

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

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

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

3. Expand Ethernet Ports.

4. Select ports 29 through 32 that are connected to the Nexus 9172PX switches, right-click them, and select Configure as Network Port.

5. Click Yes to confirm 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.

7. Verify the Ports connected to Cisco Nexus upstream swithces are now configured as network ports.

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

9. Successful configuration should result in ports 29-23 configured as network ports as shown in the screen shot below:

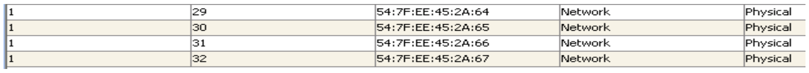

1	29	54:7F:EE:45:2A:64	Network	Physical
1	30	54:7F:EE:45:2A:65	Network	Physical
1	31	54:7F:EE:45:2A:66	Network	Physical
1	32	54:7F:EE:45:2A:67	Network	Physical

Create Uplink Port Channels to Cisco Nexus 9372PX Switches

In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 9372PX switches and one from Fabric B to both Cisco Nexus 9372PX 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 29 as the unique ID of the port channel.
6. Enter FI-A-Uplink as the name of the port channel.
7. Click Next.

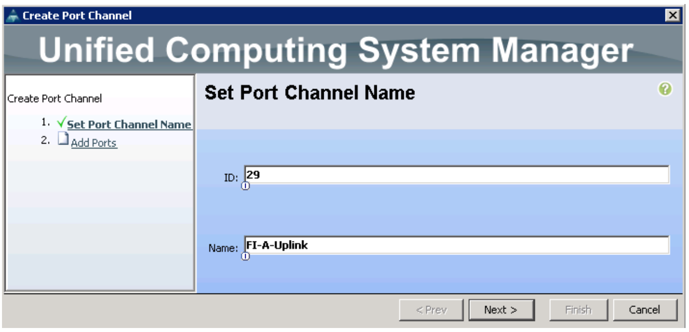

Set Port Channel Name

ID: 29

Name: FI-A-Uplink

8. Select ethernet ports 29-32 for the port channel
9. Click Finish.

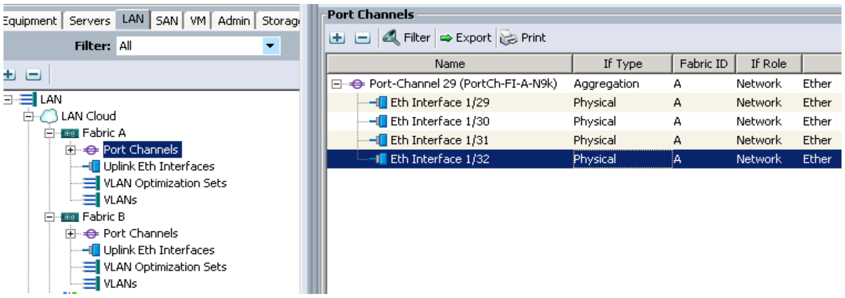

Name	If Type	Fabric ID	If Role	
Port-Channel 29 (PortCh-FI-A-N9k)	Aggregation	A	Network	Ether
Eth Interface 1/29	Physical	A	Network	Ether
Eth Interface 1/30	Physical	A	Network	Ether
Eth Interface 1/31	Physical	A	Network	Ether
Eth Interface 1/32	Physical	A	Network	Ether

10. Repeat steps 1-9 for Fabric Interconnect B, substituting 31 for the port channel number and FI-B-Uplink for the name. The resulting configuration should look like the screen shot below:

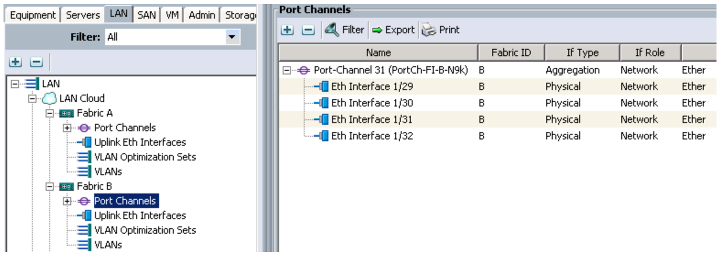

Name	Fabric ID	If Type	If Role	
Port-Channel 31 (PortCh-FI-B-N9k)	B	Aggregation	Network	Ether
Eth Interface 1/29	B	Physical	Network	Ether
Eth Interface 1/30	B	Physical	Network	Ether
Eth Interface 1/31	B	Physical	Network	Ether
Eth Interface 1/32	B	Physical	Network	Ether

Create Uplink Port Channels to Cisco MDS 9148 Switches

In this procedure, two port channels are created: One from Fabric A to Cisco MDS 9148 switch A and one from Fabric B to Cisco MDS 9148 switch B.

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Under SAN > SAN Cloud, right-click Fabric A Create Resource Pools

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.

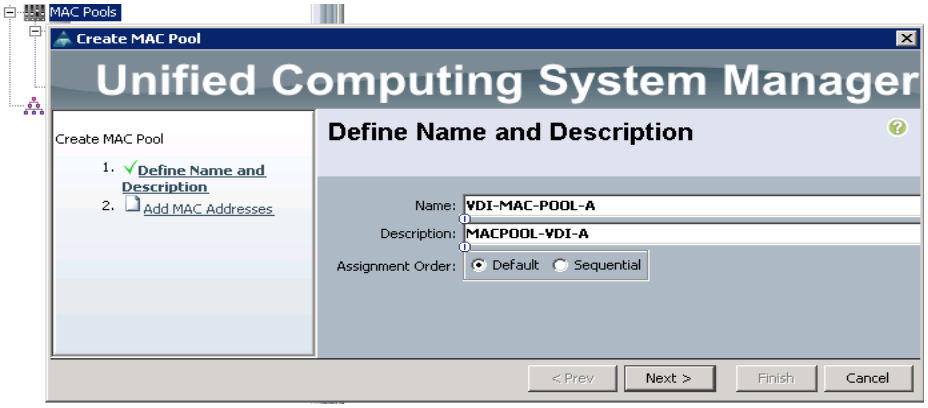

Define Name and Description

Name: VDI-MAC-POOL-A

Description: MACPOOL-VDI-A

Assignment Order: ● Default ○ Sequential

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

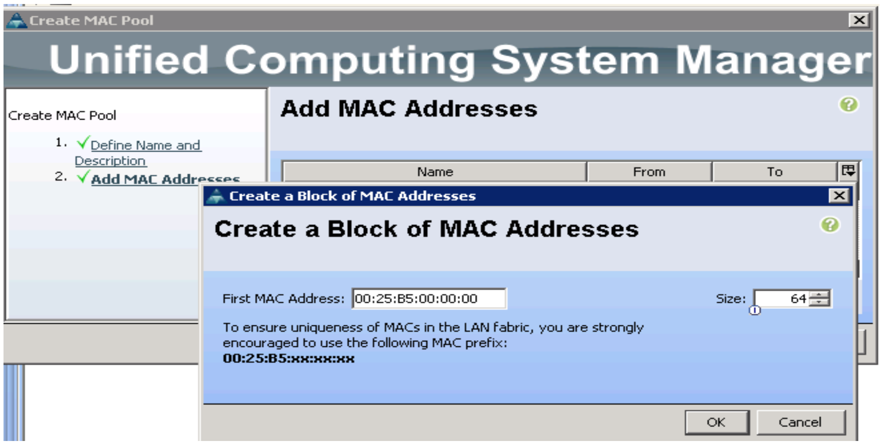
Create MAC Pool
Unified Computing System Manager
Create MAC Pool
1. Define Name and Description
2. Add MAC Addresses
Add MAC Addresses
Name
From
To
Create a Block of MAC Addresses
OK
Cancel

Create a Block of MAC Addresses

First MAC Address: 00:25:B5:00:00:00

Size: 64

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

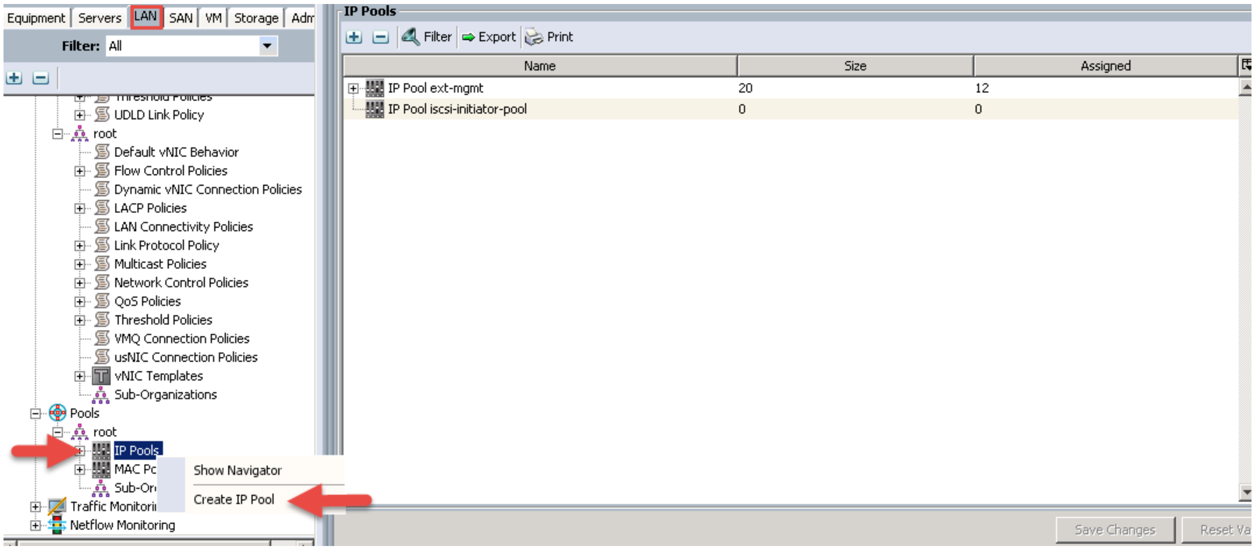

Name	Size	Assigned
IP Pool ext-mgmt	20	12
IP Pool iscsi-initiator-pool	0	0

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

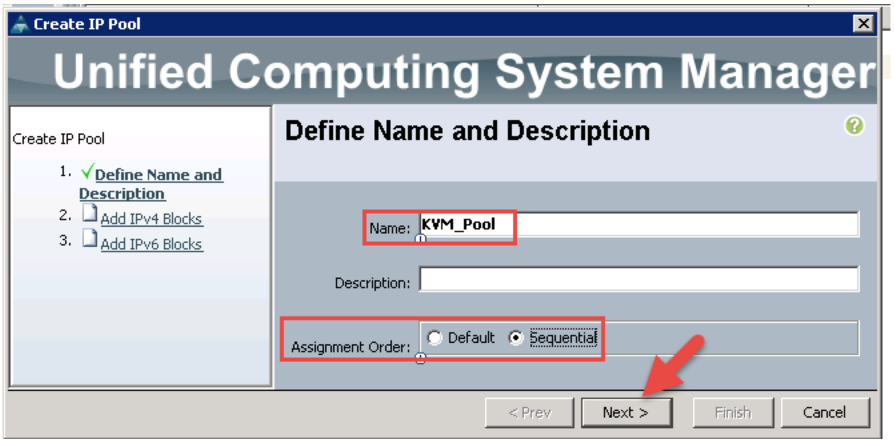
Create IP Pool
Unified Computing System Manager
Create IP Pool
1. ✓ Define Name and Description
2. Add IPv4 Blocks
3. Add IPv6 Blocks
Define Name and Description
< Prev
Next >
Finish
Cancel

Name: KVM_Pool

Description:

Assignment Order: ○ Default ● Sequential

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

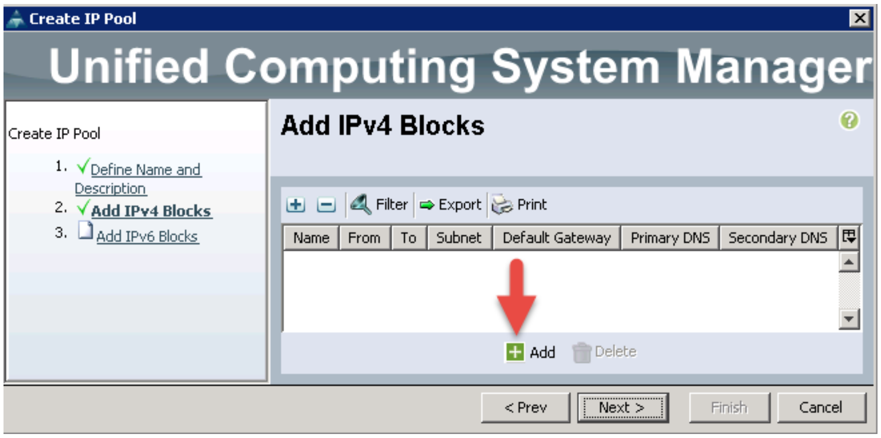

Name	From	To	Subnet	Default Gateway	Primary DNS	Secondary DNS

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

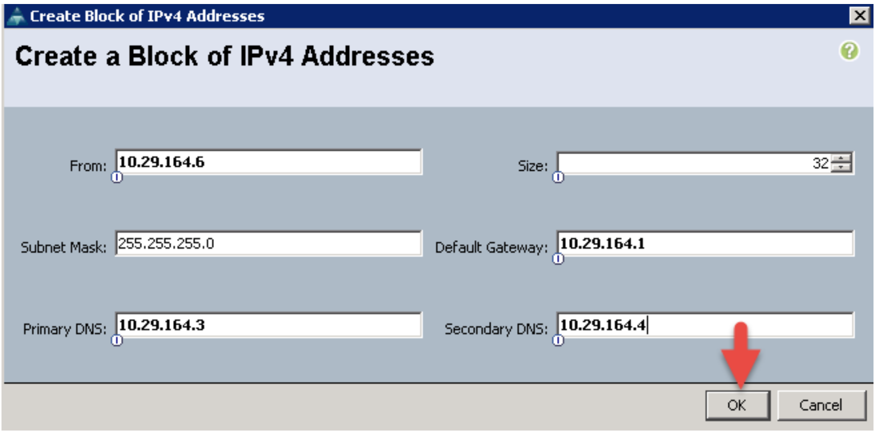
Create Block of IPv4 Addresses
Create a Block of IPv4 Addresses
OK
Cancel

From: 10.29.164.6
Size: 32
Subnet Mask: 255.255.255.0
Default Gateway: 10.29.164.1
Primary DNS: 10.29.164.3
Secondary DNS: 10.29.164.4

5. Click Next.

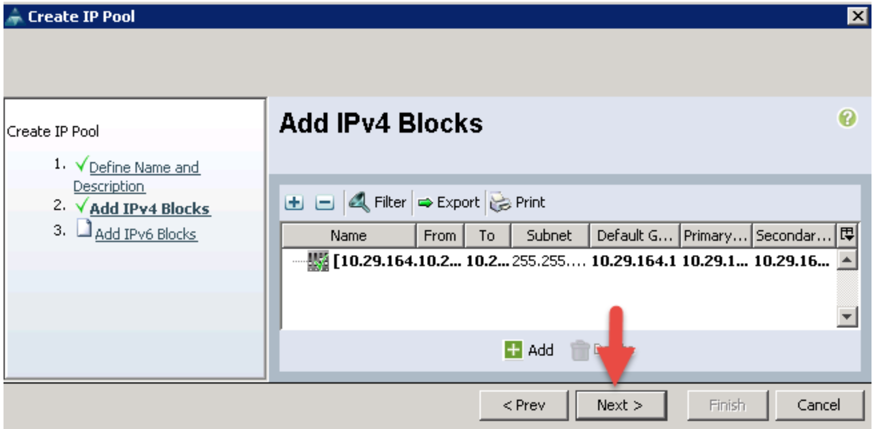

Name	From	To	Subnet	Default G...	Primary...	Secondar...
[10.29.164.10.2...		10.2...	255.255....	10.29.164.1	10.29.1...	10.29.16...

6. Click Finish.

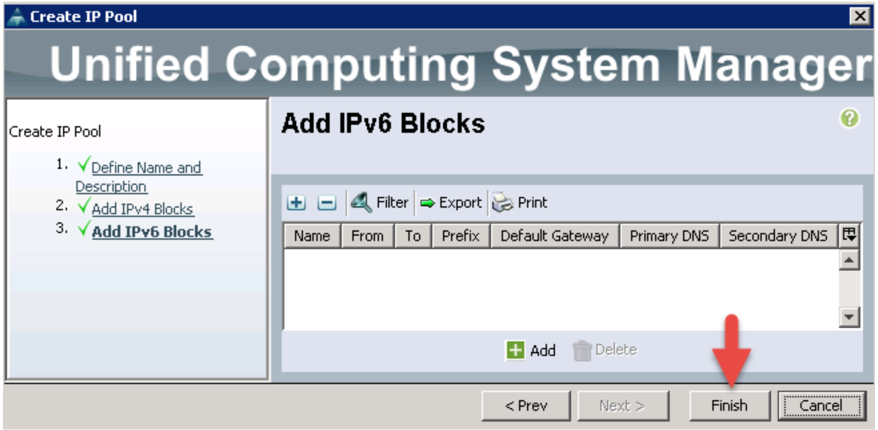

Name	From	To	Prefix	Default Gateway	Primary DNS	Secondary DNS

7. Click OK on the Success pop-up.

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.

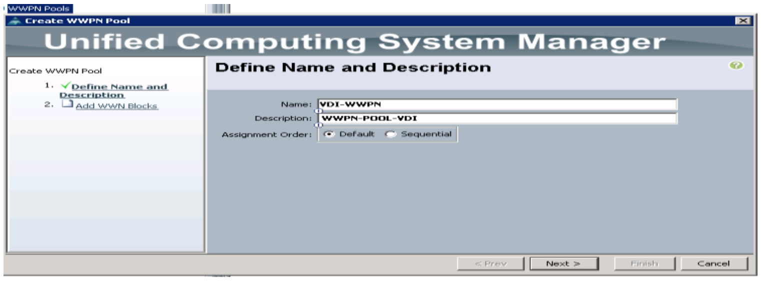
WWPN Pools
Create WWPN Pool
Unified Computing System Manager
Create WWPN Pool
1. Define Name and Description
2. Add WWN Blocks
Define Name and Description
< Prev
Next >
Finish
Cancel

Name: VDI-WWPN

Description: WWPN-POOL-VDI

Assignment Order: ● Default ○ Sequential

5. Assignment order can remain Default.

6. Click Next.

7. Click Add to add block of Ports.

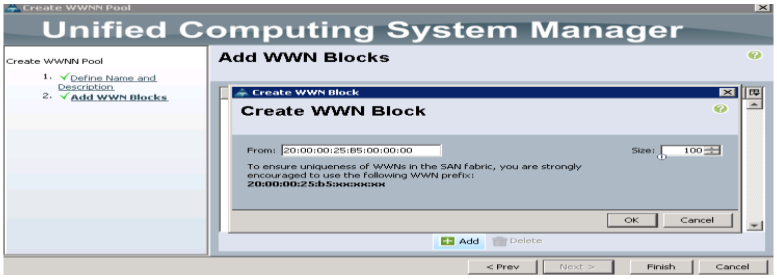

From: 20:00:00:25:B5:00:00:00

Size: 100

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

8. Enter number of WWNNs. For this study we did 100.

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.

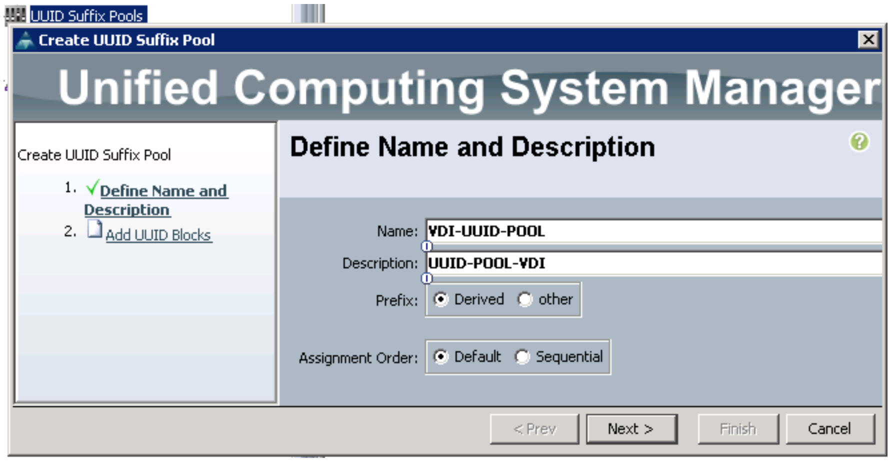

Name: VDI-UUID-POOL

Description: UUID-POOL-VDI

Prefix: ● Derived ○ other

Assignment Order: ● Default ○ Sequential

5. Enter UUID_Pool 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.

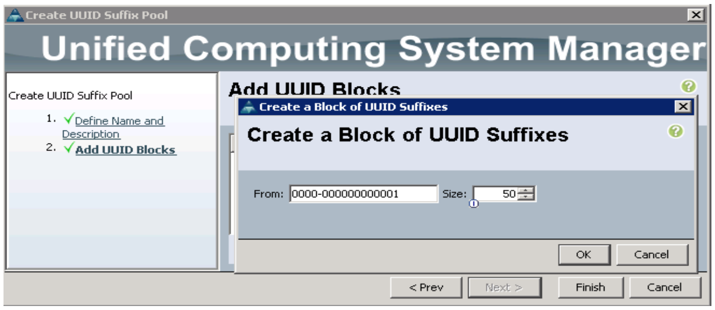
Create UUID Suffix Pool
Unified Computing System Manager
Create UUID Suffix Pool
1. Define Name and Description
2. Add UUID Blocks
Add UUID Blocks
Create a Block of UUID Suffixes
Create a Block of UUID Suffixes
OK
Cancel
< Prev
Next >
Finish
Cancel

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 XenDesktop VDI servers and XenApp RDS servers.

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, six unique VLANs are created. Refer to Table 12.

Table 12 VLANs Created

VLAN Name	VLAN ID	VLAN Purpose	vNIC Assignment
Default	1	Native VLAN	vNIC-Template-A vNIC-Template-B
In-Band-Mgmt	160	VLAN for in-band management interfaces	vNIC-Template-A vNIC-Template-B
Infra-Mgmt	161	VLAN for Virtual Infrastructure	vNIC-Template-A vNIC-Template-B
vMotion	166	VLAN for VMware vMotion	vNIC-Template-A vNIC-Template-B
VDI	162	Virtual Desktop traffic	vNIC-Template-A vNIC-Template-B
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 160 as the ID of the management VLAN.

8. Keep the Sharing Type as None.

9. Click OK, and then click OK again.

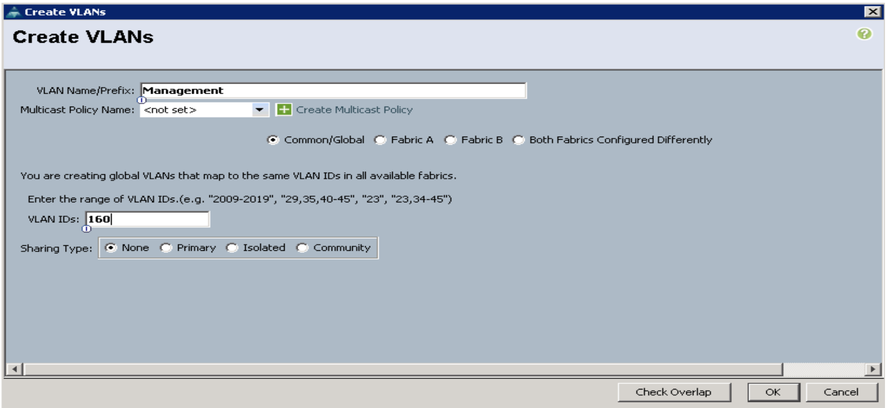

Create VLANs

VLAN Name/Prefix: Management

Multicast Policy Name: <not set>

● 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: 160

Sharing Type: ● None ○ Primary ○ Isolated ○ Community

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

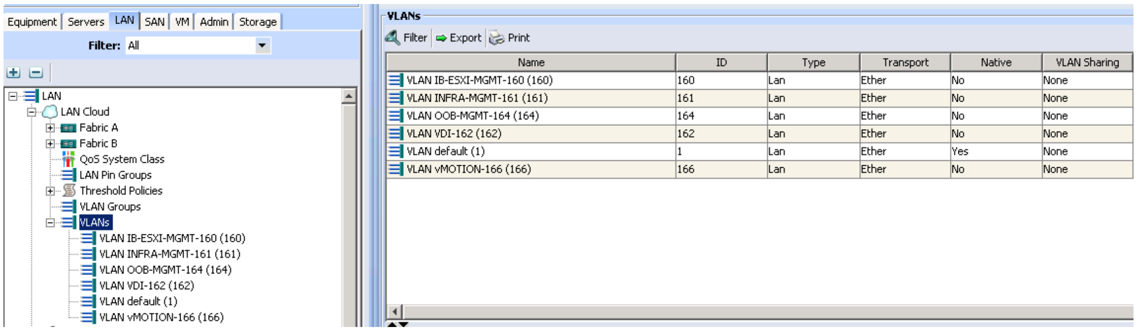

VLANs

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN IB-ESXI-MGMT-160 (160)	160	Lan	Ether	No	None
VLAN INFRA-MGMT-161 (161)	161	Lan	Ether	No	None
VLAN OOB-MGMT-164 (164)	164	Lan	Ether	No	None
VLAN VDI-162 (162)	162	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN vMOTION-166 (166)	166	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 port-channel uplink created earlier.

2. Select SAN > SAN Cloud.

3. Under Fabric A, right-click VSANs.

4. Select Create VSANs.

5. Enter VSAN20 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 20 as the ID of the VSAN.

8. Click OK, and then click OK again.

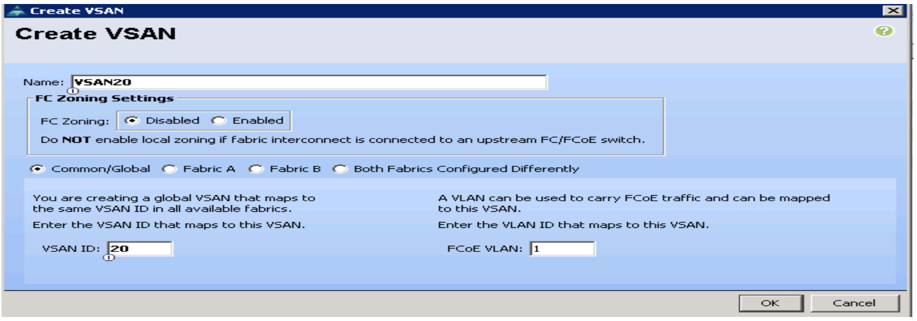

Create VSAN

Name: VSAN20

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: 20

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 VSAN30 to create the VSANs necessary for this solution.

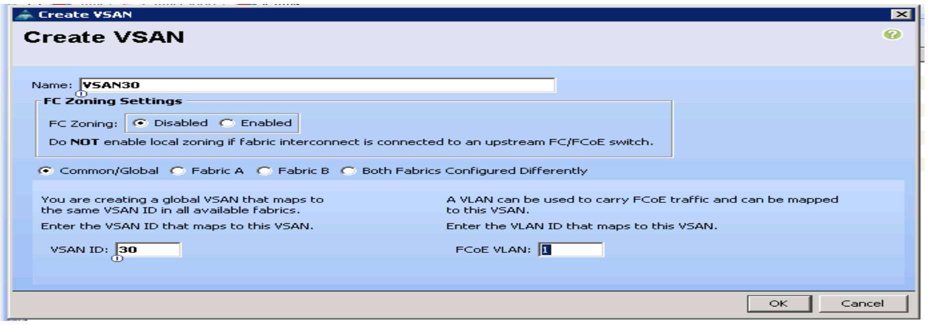

Create VSAN

Name: VSAN30

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: 30

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 20 and 30 are configured as shown below:

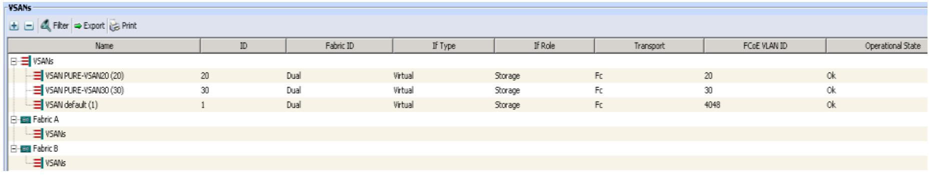

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational State
VSANs							
VSAN PURE-VSAN20 (20)	20	Dual	Virtual	Storage	Fc	20	Ok
VSAN PURE-VSAN30 (30)	30	Dual	Virtual	Storage	Fc	30	Ok
VSAN default (1)	1	Dual	Virtual	Storage	Fc	4048	Ok
Fabric A							
VSANs							
Fabric B							
VSANs							

10. After configuring VSANs both sides, go into the port-channel created earlier in the section ‘Create uplinks for MDS 9148’ and add the respective VSANs to their port channels. VSAN20 in this study is assigned to Fabric A and VSAN30 is assigned to Fabric B. (VSAN20 Should only be on Fabric A and VSAN30 on Fabric B).

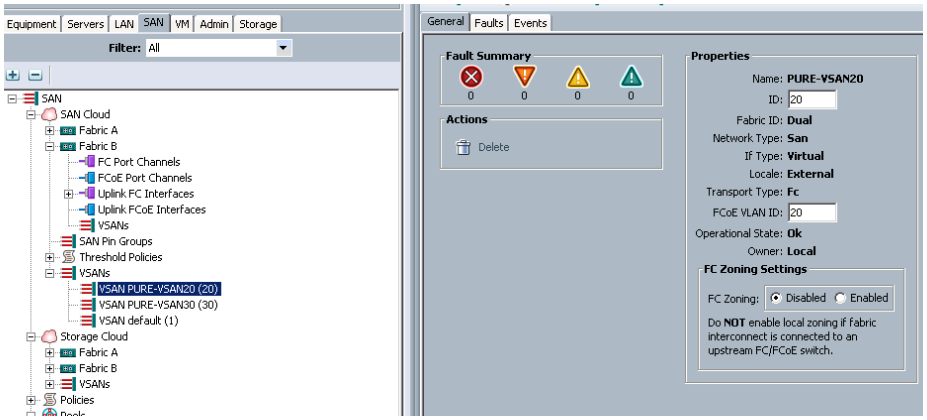

Properties

Name: PURE-VSAN20
ID: 20
Fabric ID: Dual
Network Type: San
If Type: Virtual
Locale: External
Transport Type: Fc
FCoE VLAN ID: 20
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 Port-Channel for each Fabric and assign the VSAN appropriately.

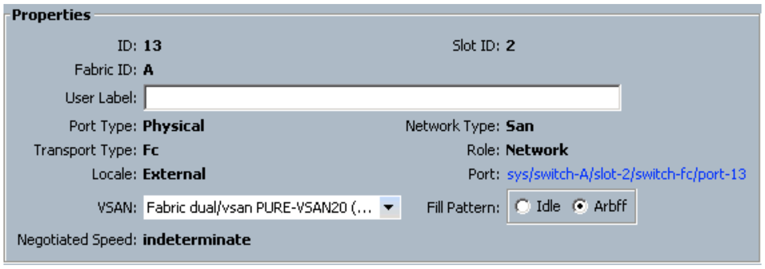

Properties

ID: 13 Slot ID: 2
Fabric ID: A
User Label:
Port Type: Physical Network Type: San
Transport Type: Fc Role: Network
Locale: External Port: sys/switch-A/slot-2/switch-fc/port-13
VSAN: Fabric dual/vsan PURE-VSAN20 (... Fill Pattern: ○ Idle ● Arbff
Negotiated Speed: indeterminate

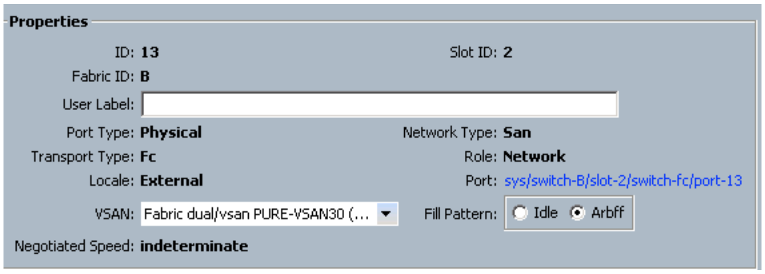

Properties

ID: 13 Slot ID: 2
Fabric ID: B
User Label:
Port Type: Physical Network Type: San
Transport Type: Fc Role: Network
Locale: External Port: sys/switch-B/slot-2/switch-fc/port-13
VSAN: Fabric dual/vsan PURE-VSAN30 (... Fill Pattern: ○ Idle ● Arbff
Negotiated Speed: indeterminate

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 2.2. (6c) for both the Blade Package

8. Click OK to create the host firmware package.

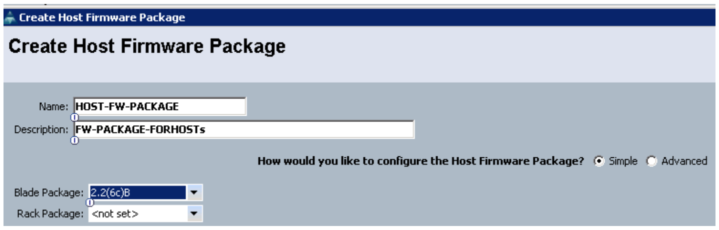

Create Host Firmware Package

Name: HOST-FW-PACKAGE

Description: FW-PACKAGE-FORHOSTs

How would you like to configure the Host Firmware Package? ● Simple ○ Advanced

Blade Package: 2.2(6c)B

Rack Package: <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.

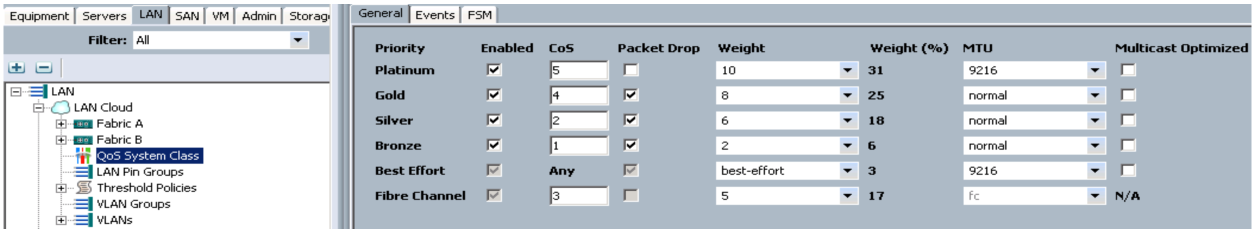

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☑	5	□	10	31	9216	□
Gold	☑	4	☑	8	25	normal	□
Silver	☑	2	☑	6	18	normal	□
Bronze	☑	1	☑	2	6	normal	□
Best Effort	☑	Any	☑	best-effort	3	9216	□
Fibre Channel	☑	3	□	5	17	fc	N/A

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.

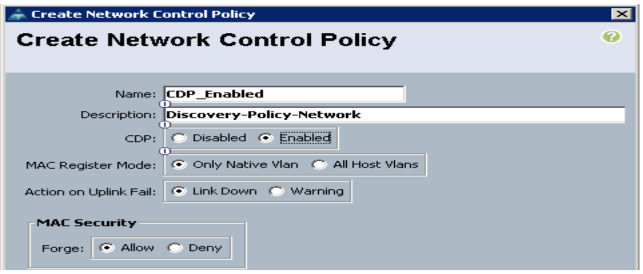

Create Network Control Policy

Name: CDP_Enabled

Description: Discovery-Policy-Network

CDP: ○ Disabled ● Enabled

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

Action on Uplink Fail: ● Link Down ○ Warning

MAC Security

Forge: ● Allow ○ Deny

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.

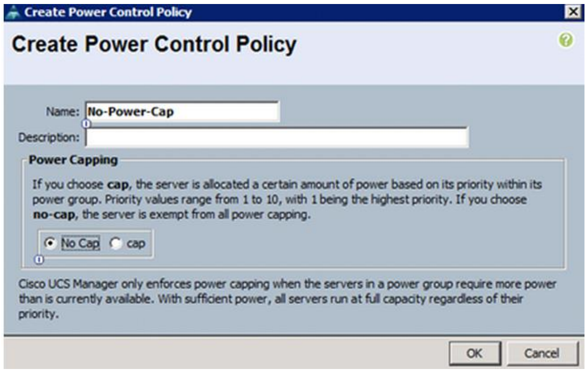

Create Power Control Policy

Name: No-Power-Cap

Description:

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-M4-BIOS as the BIOS policy name.

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

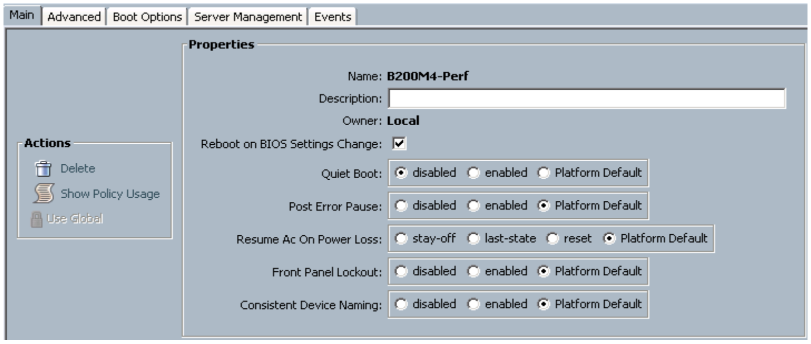

Properties

Name: B200M4-Perf

Description:

Owner: Local

Reboot on BIOS Settings Change: ☑

Quiet Boot: ● disabled ○ enabled ○ Platform Default

Post Error Pause: ○ disabled ○ enabled ● Platform Default

Resume Ac On Power Loss: ○ stay-off ○ last-state ○ reset ● Platform Default

Front Panel Lockout: ○ disabled ○ enabled ● Platform Default

Consistent Device Naming: ○ disabled ○ enabled ● Platform Default

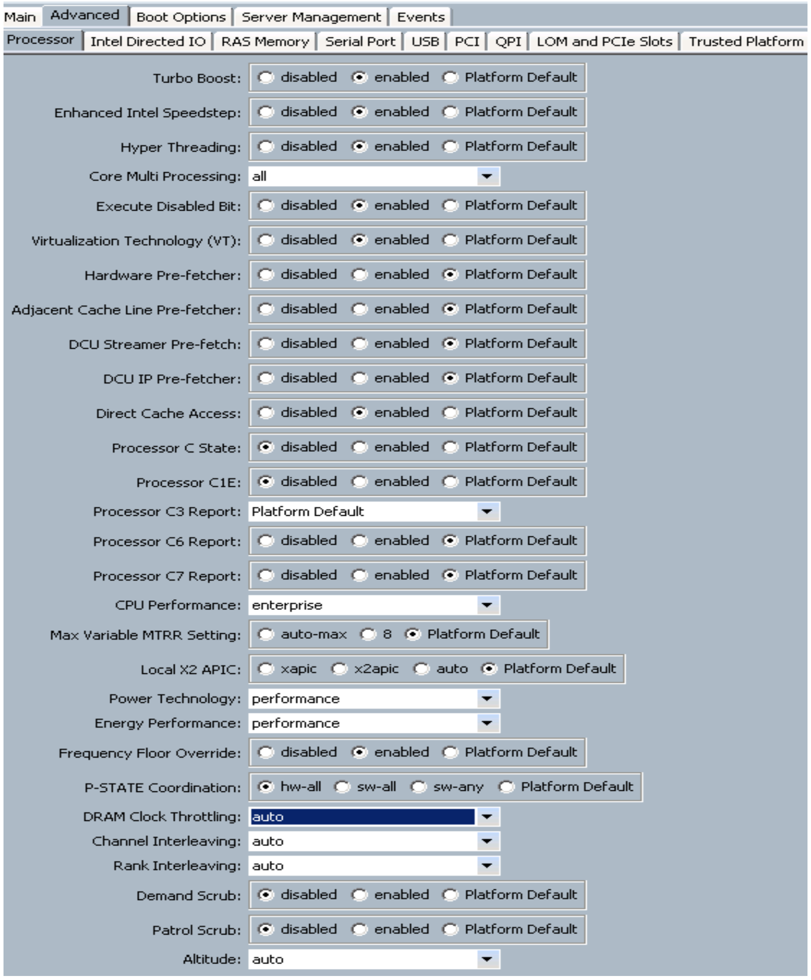

Turbo Boost: ○ disabled ● enabled ○ Platform Default

Enhanced Intel Speedstep: ○ disabled ● enabled ○ Platform Default

Hyper Threading: ○ disabled ● enabled ○ Platform Default

Core Multi Processing: all

Execute Disabled Bit: ○ disabled ● enabled ○ Platform Default

Virtualization Technology (VT): ○ disabled ● enabled ○ Platform Default

Hardware Pre-fetcher: ○ disabled ○ enabled ● Platform Default

Adjacent Cache Line Pre-fetcher: ○ disabled ○ enabled ● Platform Default

DCU Streamer Pre-fetch: ○ disabled ○ enabled ● Platform Default

DCU IP Pre-fetcher: ○ disabled ○ enabled ● Platform Default

Direct Cache Access: ○ disabled ● enabled ○ Platform Default

Processor C State: ● disabled ○ enabled ○ Platform Default

Processor C1E: ● disabled ○ enabled ○ Platform Default

Processor C3 Report: Platform Default

Processor C6 Report: ○ disabled ○ enabled ● Platform Default

Processor C7 Report: ○ disabled ○ enabled ● Platform Default

CPU Performance: enterprise

Max Variable MTRR Setting: ○ auto-max ○ 8 ● Platform Default

Local X2 APIC: ○ xapic ○ x2apic ○ auto ● Platform Default

Power Technology: performance

Energy Performance: performance

Frequency Floor Override: ○ disabled ● enabled ○ Platform Default

P-STATE Coordination: ● hw-all ○ sw-all ○ sw-any ○ Platform Default

DRAM Clock Throttling: auto

Channel Interleaving: auto

Rank Interleaving: auto

Demand Scrub: ● disabled ○ enabled ○ Platform Default

Patrol Scrub: ● disabled ○ enabled ○ Platform Default

Altitude: auto

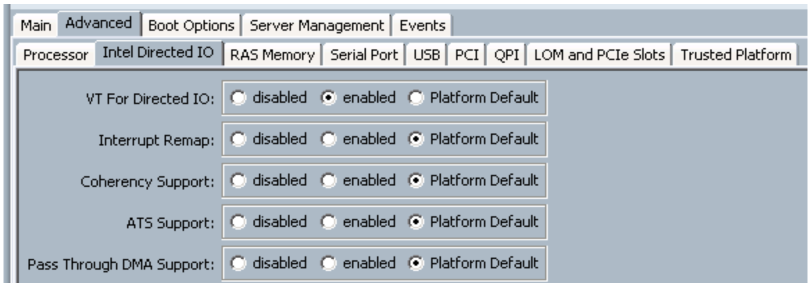

VT For Directed IO: ○ disabled ● enabled ○ Platform Default

Interrupt Remap: ○ disabled ○ enabled ● Platform Default

Coherency Support: ○ disabled ○ enabled ● Platform Default

ATS Support: ○ disabled ○ enabled ● Platform Default

Pass Through DMA Support: ○ disabled ○ enabled ● 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. Click Save Changes.

6. Click OK to accept the change.

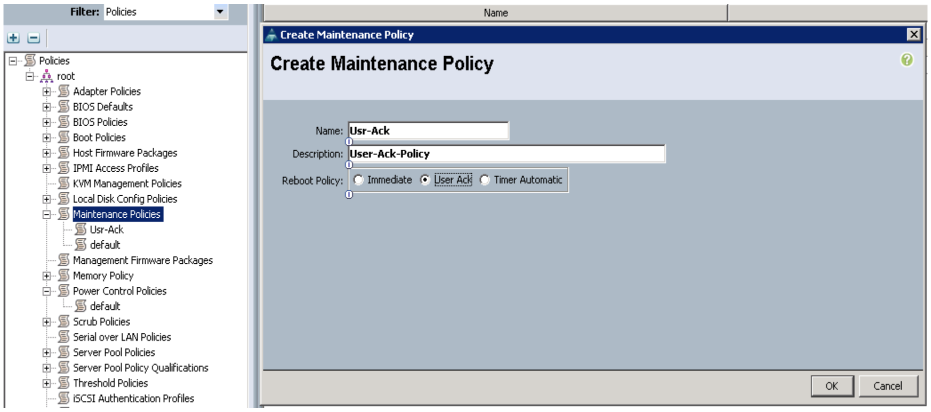

Create Maintenance Policy

Name: Usr-Ack

Description: User-Ack-Policy

Reboot Policy: ○ Immediate ● User Ack ○ Timer Automatic

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, VDI, Infra, 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.

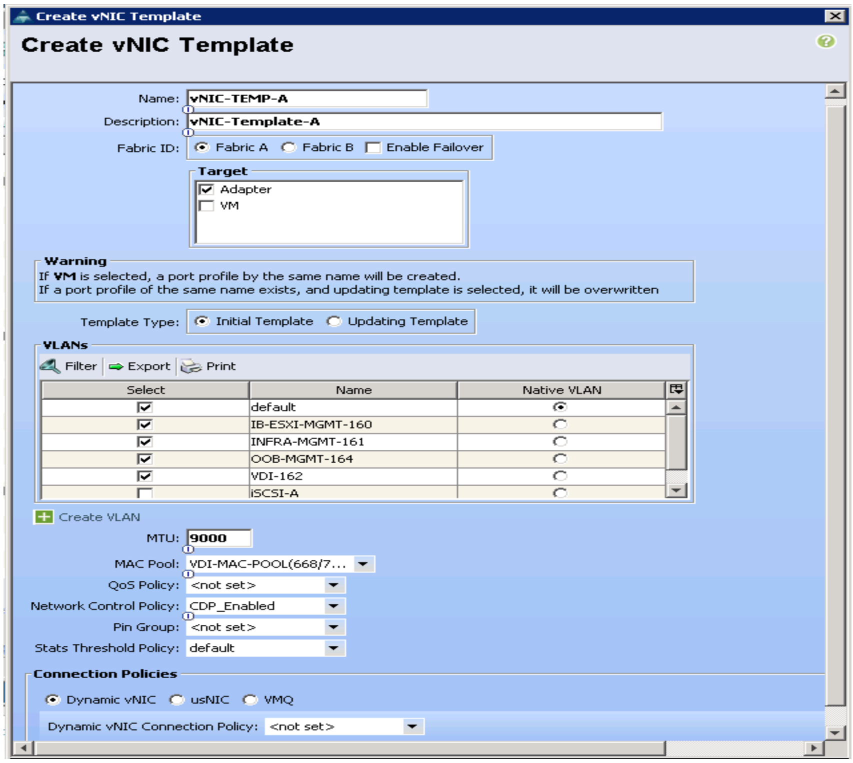

Create vNIC Template

Name: vNIC-TEMP-A

Description: vNIC-Template-A

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

Target

☑ Adapter

□ VM

Warning

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

Template Type: ● Initial Template ○ Updating Template

VLANs

Select	Name	Native VLAN
☑	default	●
☑	IB-ESXI-MGMT-160	○
☑	INFRA-MGMT-161	○
☑	OOB-MGMT-164	○
☑	VDI-162	○
□	iSCSI-A	○

MTU: 9000

MAC Pool: VDI-MAC-POOL(668/7...

QoS Policy: <not set>

Network Control Policy: CDP_Enabled

Pin Group: <not set>

Stats Threshold Policy: default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy: <not set>

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 VDI-vHBA_FI-A as the vHBA template name.

6. Keep Fabric A selected.

7. Select VSAN20 for Fabric A from the drop down.

8. Change to Updating Template.

9. For Max Data Field keep 2048.

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

11. Leave the remaining as is.

12. Click OK.

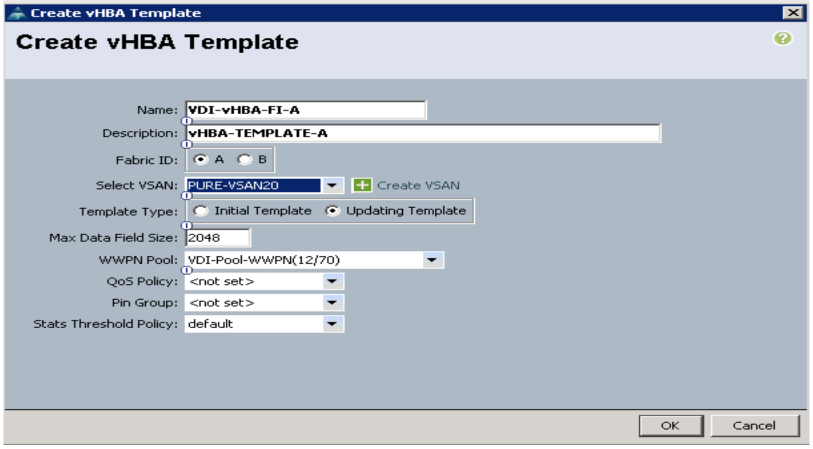

Create vHBA Template

Name: VDI-vHBA-FI-A

Description: vHBA-TEMPLATE-A

Fabric ID: ● A ○ B

Select VSAN: PURE-VSAN20

Template Type: ○ Initial Template ● Updating Template

Max Data Field Size: 2048

WWPN Pool: VDI-Pool-WWPN(12/70)

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 VDI-vHBA_FI-B as the vHBA template name.

18. Select Fabric B.

19. Select VSAN30 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 host 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.’

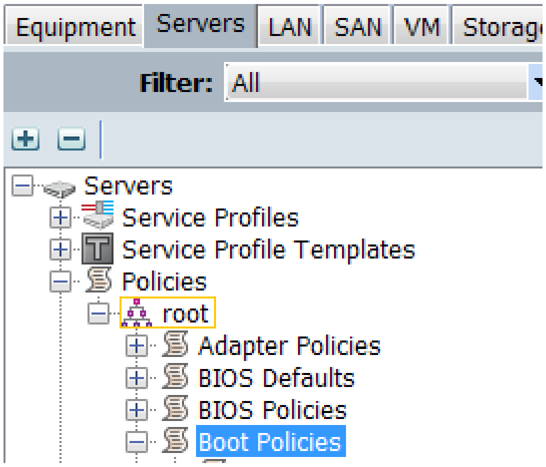

2. Name the boot policy and expand the ‘vHBAs’ menu as shown below:

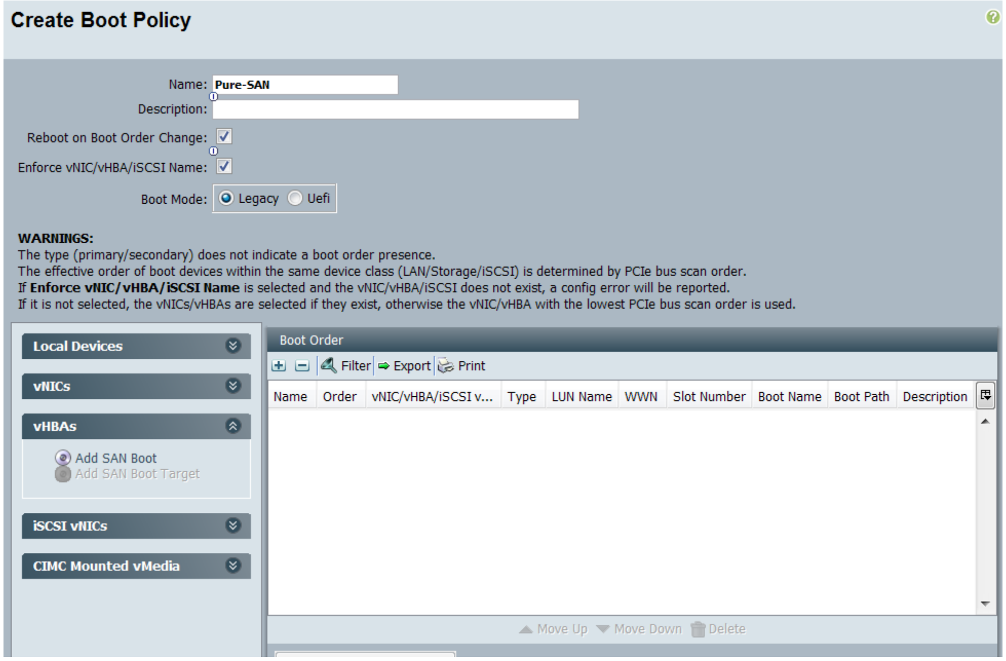

Create Boot Policy

Name: Pure-SAN

Description:

Reboot on Boot Order Change: ☑

Enforce vNIC/vHBA/iSCSI Name: ☑

Boot Mode: ● Legacy ○ Uefi

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

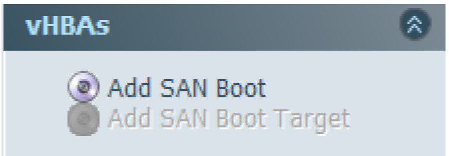

3. After selecting the ‘Add SAN Boot’ option, add the primary vHBA as shown below. Note that the vHBA name needs to match exactly. We will use the vHBA templates created in the previous step.

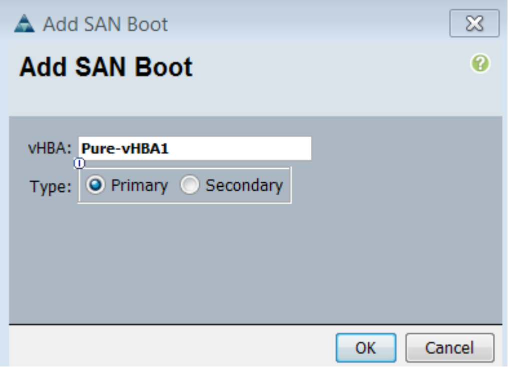

Add SAN Boot

vHBA: Pure-vHBA1

Type: ● Primary ○ Secondary

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

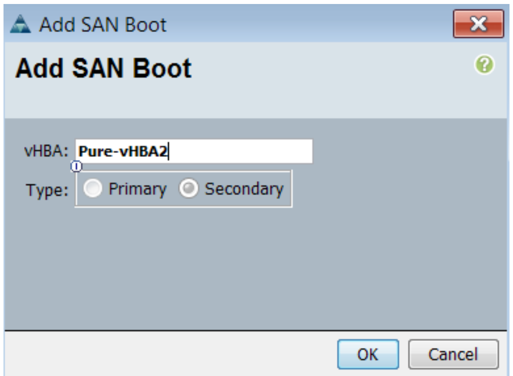

Add SAN Boot

Add SAN Boot

vHBA: Pure-vHBA2

Type: ○ Primary ● Secondary

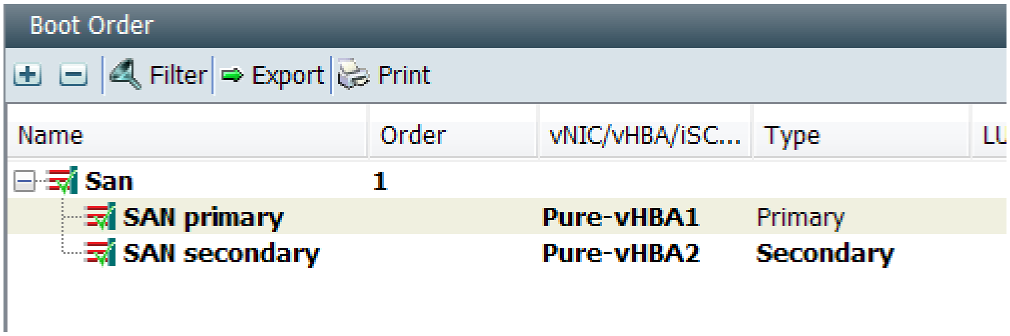

Boot Order

Name	Order	vNIC/vHBA/iSC...	Type	LU
San	1			
SAN primary		Pure-vHBA1	Primary	
SAN secondary		Pure-vHBA2	Secondary	

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.

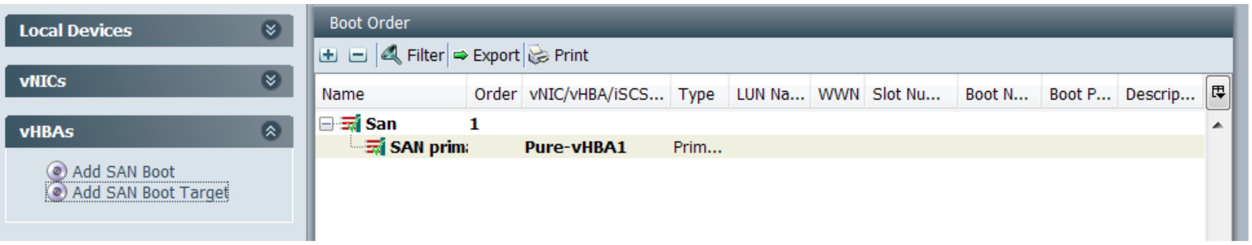

Boot Order

Name	Order	vNIC/vHBA/iSCS...	Type	LUN Na...	WWN	Slot Nu...	Boot N...	Boot P...	Descrip...
San	1								
SAN prim:		Pure-vHBA1	Prim...						

6. Highlight the SAN primary and select ‘Add SAN Boot Target to SAN Primary’

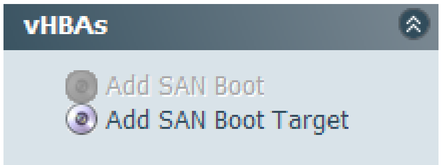

7. From the Pure Storage GUI, find and enter the Pure Storage WWN for Controller 0, Fibre Channel Port 0. This information can be found in the Pure Storage GUI under System -> Host Connections at the bottom of the screen under ‘Target Ports’:

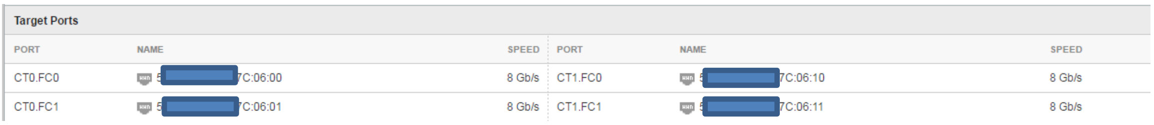

Target Ports

PORT	NAME	SPEED	PORT	NAME	SPEED
CT0.FC0	5[redacted]7C:06:00	8 Gb/s	CT1.FC0	5[redacted]7C:06:10	8 Gb/s
CT0.FC1	5[redacted]7C:06:01	8 Gb/s	CT1.FC1	5[redacted]7C:06:11	8 Gb/s

8. When the Pure WWNs have been recorded, use port CT0.FC0 for the first Boot Target WWPN:

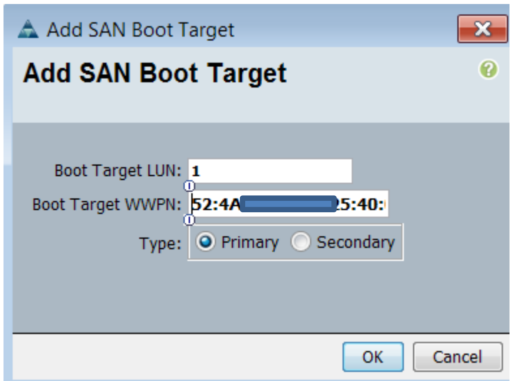

Add SAN Boot Target

Add SAN Boot Target

Boot Target LUN: 1

Boot Target WWPN: 52:4A[redacted]5:40:

Type: ● Primary ○ Secondary

9. Add a secondary SAN Boot Target by again clicking the ‘Add SAN Boot Target to SAN Primary’ while the primary SAN Boot option is highlighted. This time enter the Pure Storage WWPN for CT1.FC0.

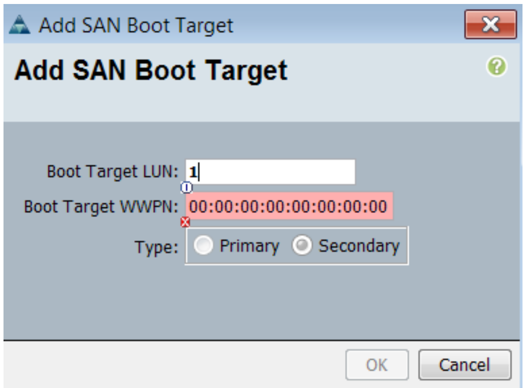

Add SAN Boot Target

Add SAN Boot Target

Boot Target LUN: 1

Boot Target WWPN: 00:00:00:00:00:00:00:00

Type: ○ Primary ● Secondary

10. Repeat this procedure for the secondary SAN boot target and use WWPN for CT0.FC1 and CT1.FC1 in the primary and secondary SAN boot options.

11. Below you can see a properly configured boot from SAN UCS policy. The next step is to create and attach the boot volumes to the hosts from within the Pure GUI:

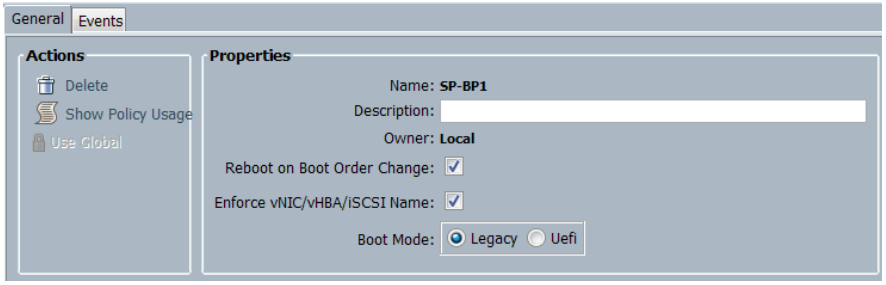

Properties

Name: SP-BP1

Description:

Owner: Local

Reboot on Boot Order Change: ☑

Enforce vNIC/vHBA/iSCSI Name: ☑

Boot Mode: ● Legacy ○ Uefi

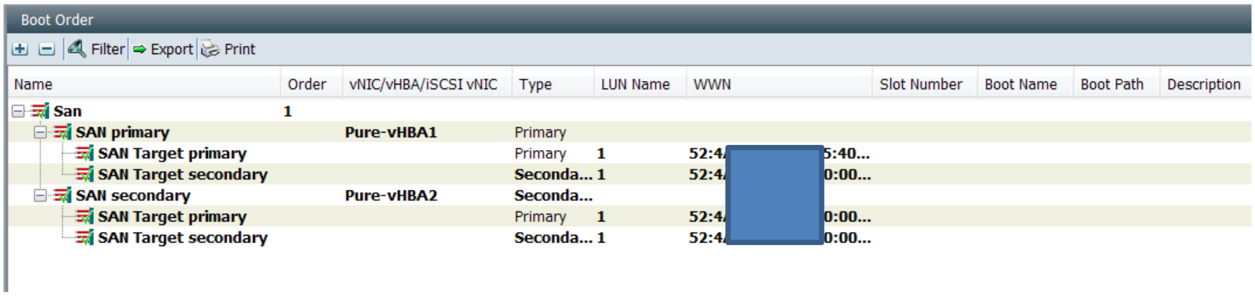

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	Slot Number	Boot Name	Boot Path	Description
San	1								
SAN primary		Pure-vHBA1	Primary						
SAN Target primary			Primary	1	52:4[redacted]5:40...				
SAN Target secondary			Seconda...	1	52:4[redacted]0:00...				
SAN secondary		Pure-vHBA2	Seconda...						
SAN Target primary			Primary	1	52:4[redacted]0:00...				
SAN Target secondary			Seconda...	1	52:4[redacted]0:00...				

12. Provisioning a volume for SAN boot for Pure Storage is simple. The only difference between a SAN volume and a vCenter datastore is that you will only connect the SAN boot volume to the single host that will be booting from the volume.

13. Within the Pure Storage GUI, select the + icon next to Volumes under the Storage tab.

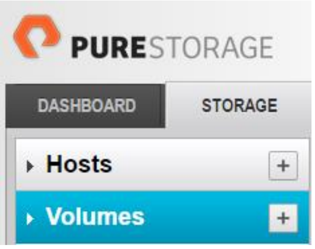

14. Create multiple volumes since there are a large number of ESXi hosts and you need to create separate boot volumes for each one.

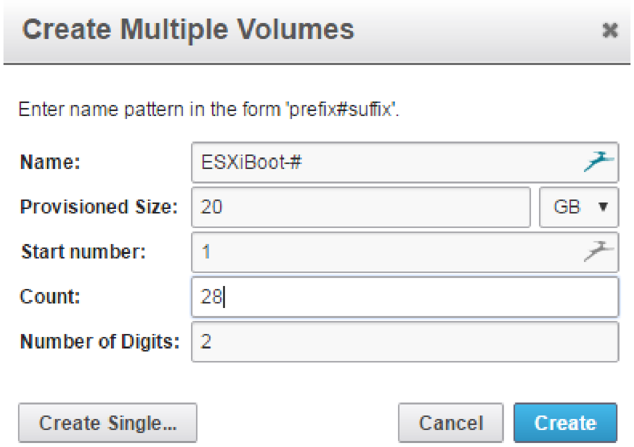

Enter name pattern in the form 'prefix#suffix'.

Name:	ESXiBoot-#
Provisioned Size:	20 GB ▾
Start number:	1
Count:	28
Number of Digits:	2

15. To connect the LUN to an ESXi host, select a newly created volume under the Volumes tab and then select ‘Connect Hosts’ from the options under the gear to the right:

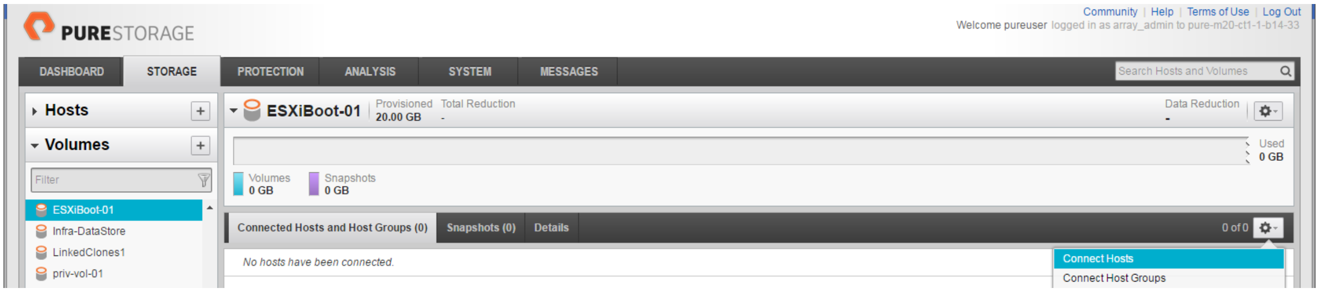

16. Click the single host you wish to connect to this volume and then click Confirm.

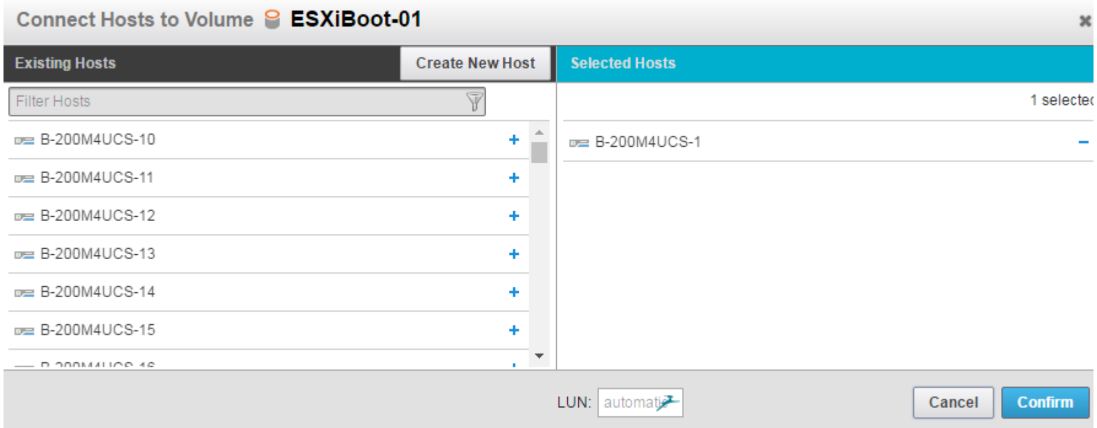

17. Repeat this procedure with all ESXi hosts. Remember it is important to only connect a single host to each boot LUN.

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

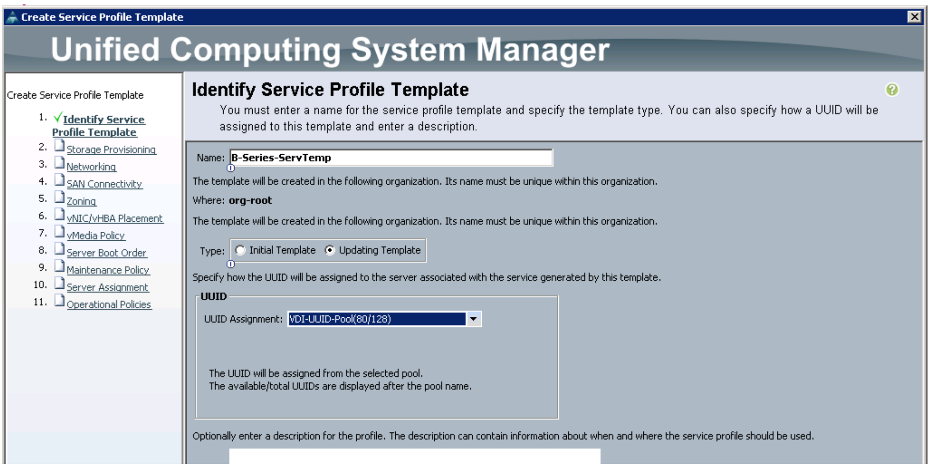

Create Service Profile Template

Unified Computing System Manager

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

Identify 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: B-Series-ServTemp

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: VDI-UUID-Pool(80/128)

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.

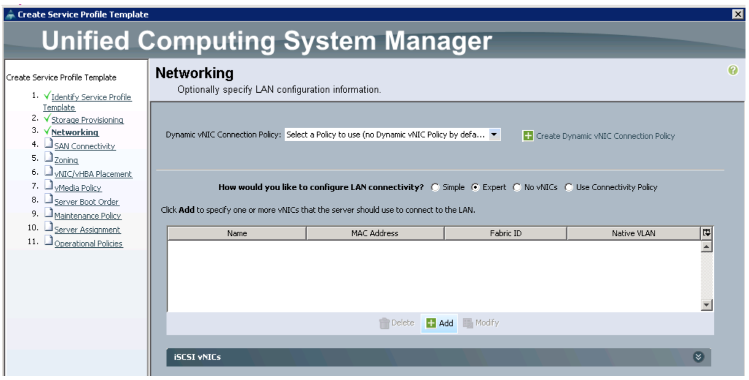

Create Service Profile Template

Unified Computing System Manager

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

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by defa...

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

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.

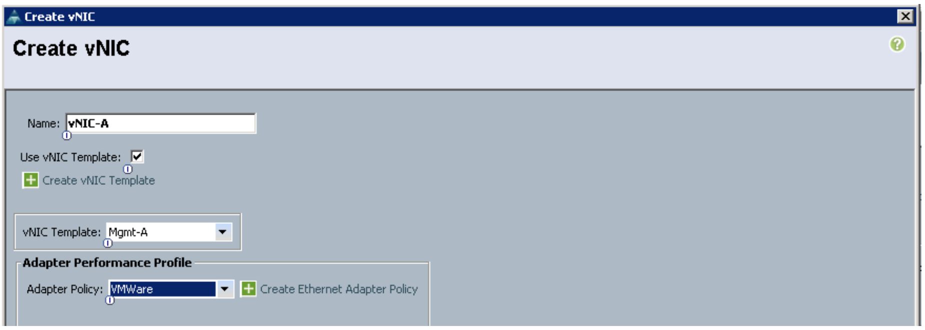
Create vNIC

Name: vNIC-A

Use vNIC Template: ☑

vNIC Template: Mgmt-A

Adapter Performance Profile

Adapter Policy: VMWare

13. Repeat networking steps for vNIC-B.

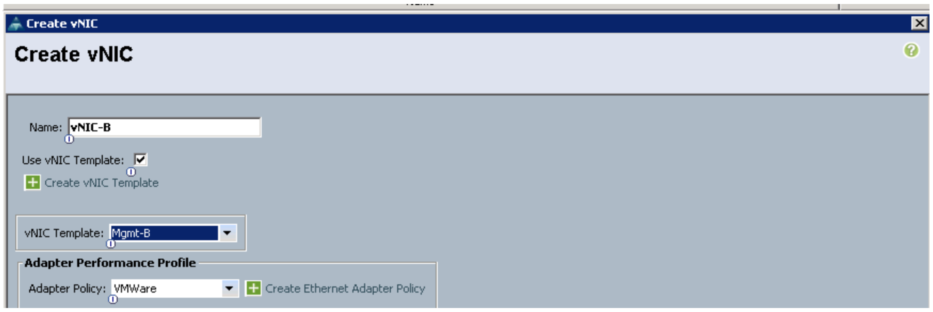
Create vNIC

Name: vNIC-B

Use vNIC Template: ☑

vNIC Template: Mgmt-B

Adapter Performance Profile

Adapter Policy: VMWare

14. Click Next.

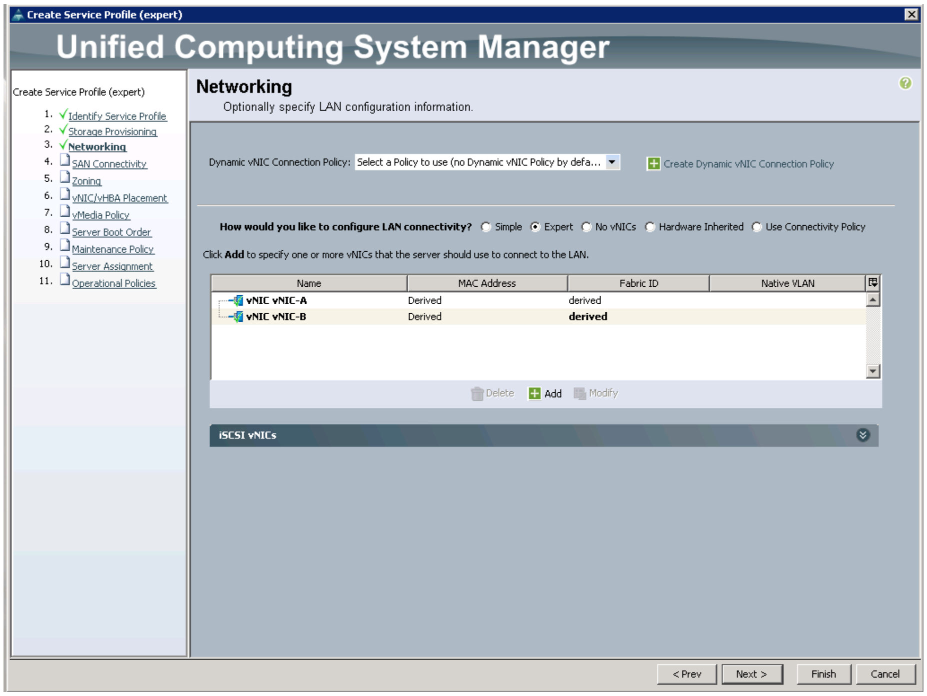
Create Service Profile (expert)

Unified Computing System Manager

Create Service Profile (expert)

1. ✓ Identify Service Profile
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

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by defa...

How would you like to configure LAN connectivity? ○ Simple ● Expert ○ No vNICs ○ Hardware Inherited ○ 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-A	Derived	derived	
vNIC vNIC-B	Derived	derived	

15. Click Next.

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

17. Select WWNN Assignment from the Pool created earlier.

18. Click Add.

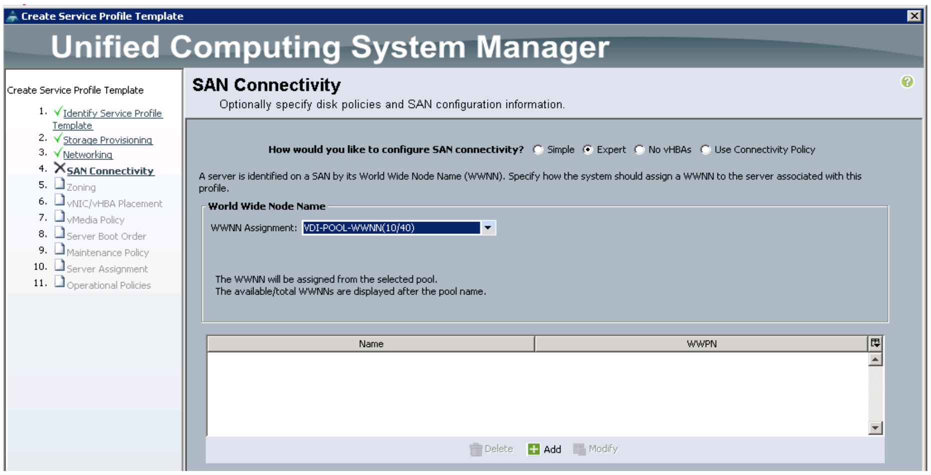

Create Service Profile Template

Unified Computing System Manager

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

SAN Connectivity

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-POOL-WWNN(10/40)

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

Name	WWPN

19. Name the adapter vHBA-A.

20. Click Use vHBA Template

21. Select vHBA Template: vHBA-A.

22. Select Adapter Policy : VMWare.

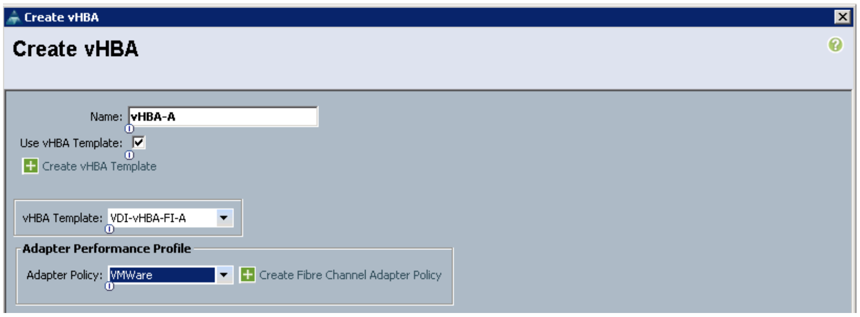

Create vHBA

Create vHBA

Name: vHBA-A

Use vHBA Template: ☑

vHBA Template: VDI-vHBA-FI-A

Adapter Performance Profile

Adapter Policy: VMWare

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

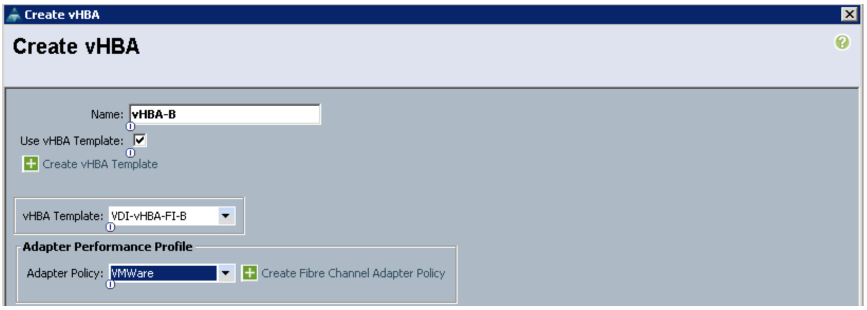

Create vHBA

Create vHBA

Name: vHBA-B

Use vHBA Template: ☑

vHBA Template: VDI-vHBA-FI-B

Adapter Performance Profile

Adapter Policy: VMWare

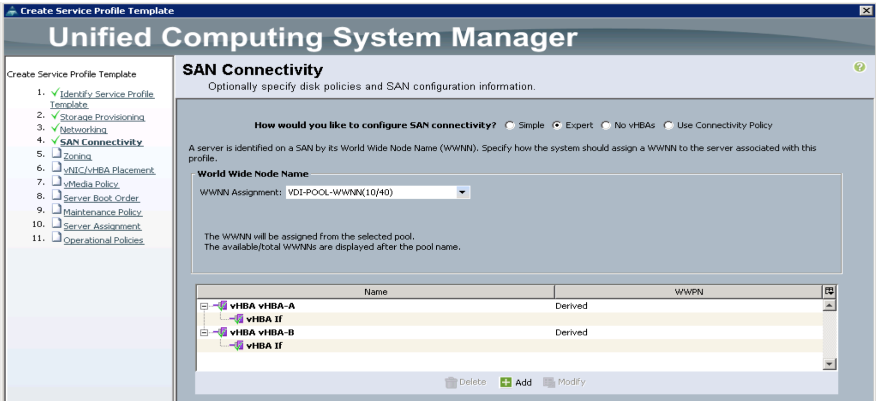

SAN Connectivity

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-POOL-WWNN(10/40)

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

Name	WWPN
vHBA vHBA-A	Derived
vHBA If	
vHBA vHBA-B	Derived
vHBA If	

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.

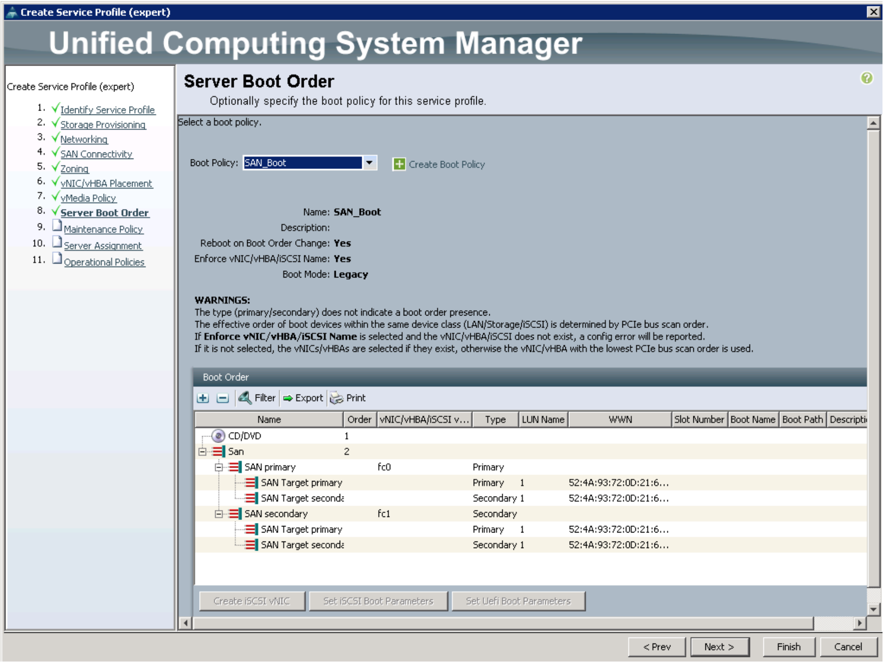

Server Boot Order

Optionally specify the boot policy for this service profile.

Select a boot policy.

Boot Policy: SAN_Boot

Name: **SAN_Boot**
Description:
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.

Name	Order	vNIC/vHBA/iSCSI v...	Type	LUN Name	WWN	Slot Number	Boot Name	Boot Path	Descripti
CD/DVD	1								
San	2								
SAN primary		fc0	Primary						
SAN Target primary			Primary	1	52:4A:93:72:0D:21:6...				
SAN Target seconda			Secondary	1	52:4A:93:72:0D:21:6...				
SAN secondary		fc1	Secondary						
SAN Target primary			Primary	1	52:4A:93:72:0D:21:6...				
SAN Target seconda			Secondary	1	52:4A:93:72:0D:21:6...				

Create Service Profiles

To create service profiles for each of the blades in the FlashStack 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 step above, then click Create Service Profiles from Template.

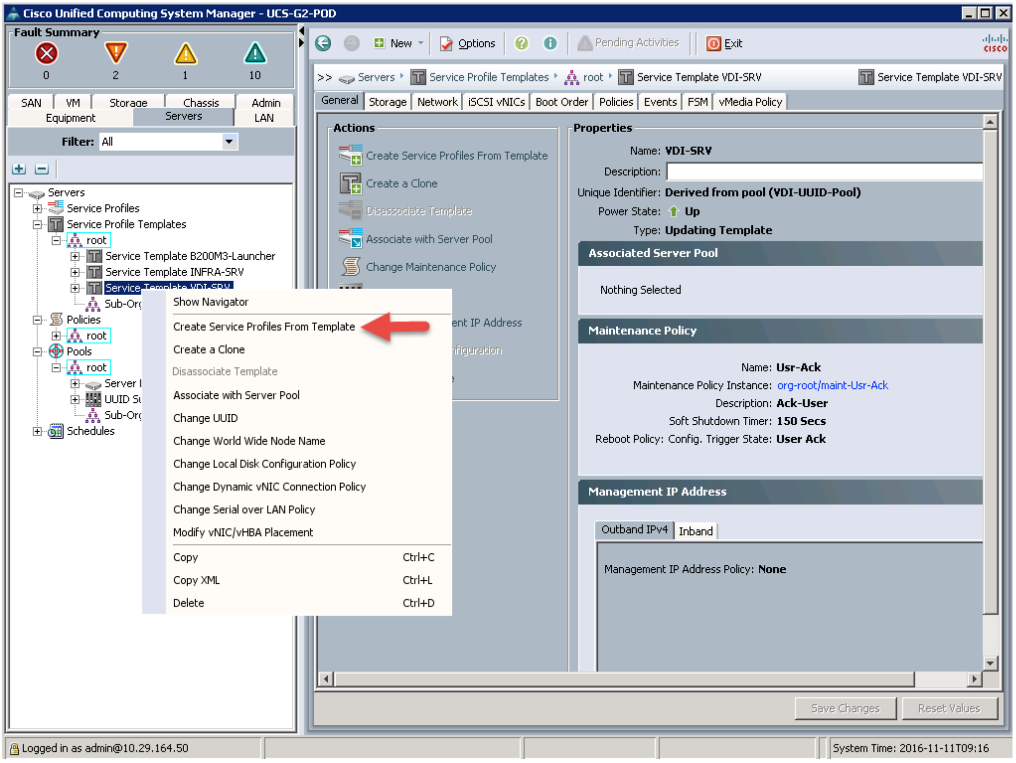

Name: VDI-SRV
Description:
Unique Identifier: Derived from pool (VDI-UUID-Pool)
Power State: Up
Type: Updating Template

Nothing Selected

Name: Usr-Ack
Maintenance Policy Instance: org-root/maint-Usr-Ack
Description: Ack-User
Soft Shutdown Timer: 150 Secs
Reboot Policy: Config. Trigger State: User Ack

Management IP Address Policy: None

Logged in as admin@10.29.164.50
System Time: 2016-11-11T09:16

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

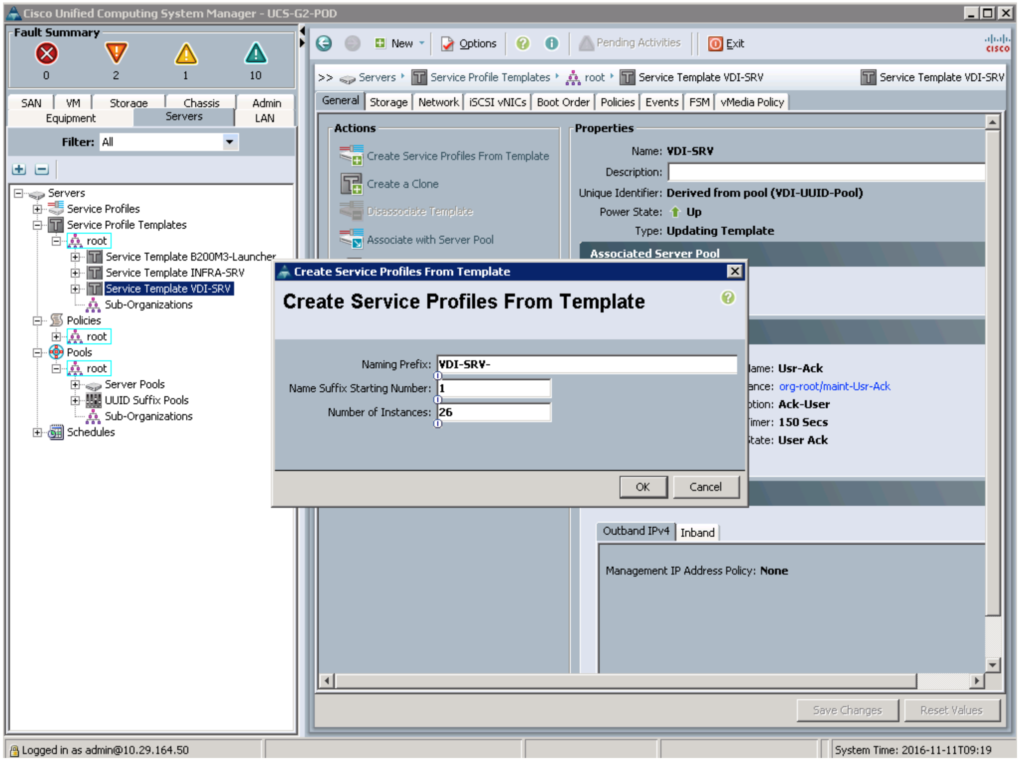

Name: VDI-SRV
Description:
Unique Identifier: Derived from pool (VDI-UUID-Pool)
Power State: Up
Type: Updating Template

Create Service Profiles From Template

Naming Prefix: VDI-SRV-
Name Suffix Starting Number: 1
Number of Instances: 26

lame: Usr-Ack
ance: org-root/maint-Usr-Ack
ption: Ack-User
imer: 150 Secs
tate: User Ack

Management IP Address Policy: None

Logged in as admin@10.29.164.50
System Time: 2016-11-11T09:19

The requested number of service profiles are created in the Service Profiles root organization.

Pure Storage FlashArray//m50 Configuration for Cisco Validated Design

The Pure Storage FlashArray//m50 contains no special configurations, tuning or value changes from the default values. For the validated design, the FlashArray//m50 contains twenty drive bays fully populated with two 1 GB SSDs each as well as a 44TB expansion shelf with two NVRAM devices for 84TB of raw space in total.

The expansion shelf is connected with four 12GB SAS cables (two per controller) in order to ensure full redundancy, availability and performance of the shelf in the event of a single controller reboot or failure.

A Pure Storage Systems Engineer or authorized partner will perform the installation and initial setup of the array. Setup (racking, power, cabling and array setup) is typically completed in less than one hour and will only require between 12-16 cables (that number is variable based upon number of SAN FC connections) for the base unit and expansion shelf. Arrays without an expansion shelf require as little as six cables in total.

The Cisco UCS hosts are redundantly connected to the array controllers (via the Cisco MDS 9148S) with four FC connections to each controller from two HBAs on each Cisco UCS host over the 16GB Fibre Channel protocol for a total of eight logical paths for 128GB/s of total throughput with no single point of failure. The FlashArray//m50 array will support up to eight FC connections in total in the baseline system and using all eight connections is recommended in production deployments in order to maximize resiliency and throughput capabilities of the array.

Pure Storage FlashArray//m Configuration

The Pure Storage FlashArray//m50 used in this validated design required a total of 5 rack units: 3 for the FlashArray//m50 base unit and 2 for the 44TB expansion shelf in a standard length rack. The front view of the array can be seen in Figure 28 below:

Figure 28 Pure Storage FlashArray//m Storage Front View

Behind the bezel on the FlashArray//m50 the 2 NVRAM modules and 20 drive bays can be accessed as shown in the screenshot below:

Figure 29 Pure Storage FlashArray //m Array 2 NV RAM Modules Overview

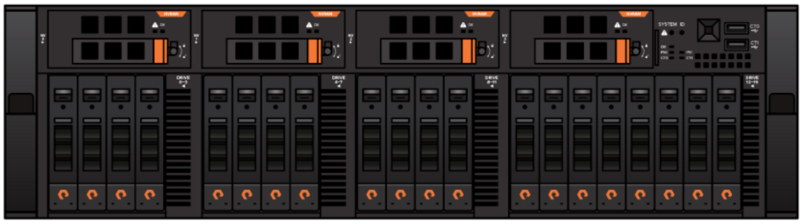

The figure below shows the back of the array which houses the two controllers. The rear of the array in this example includes Fibre-Channel connectivity (8x 16GB/s ports), though 10GB iSCSI is also supported.

Figure 30 Pure Storage FlashArray//m50 Controller 0 and 1 Overview

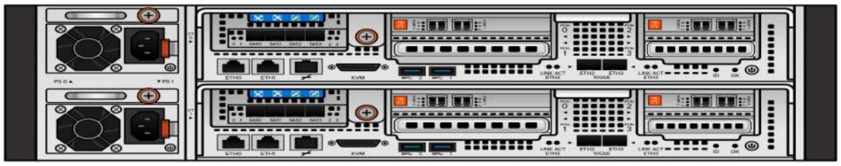

The table below lists the default port services associated with the rear view of the array. Pure Storage FlashArray//m Default Ports Overview:

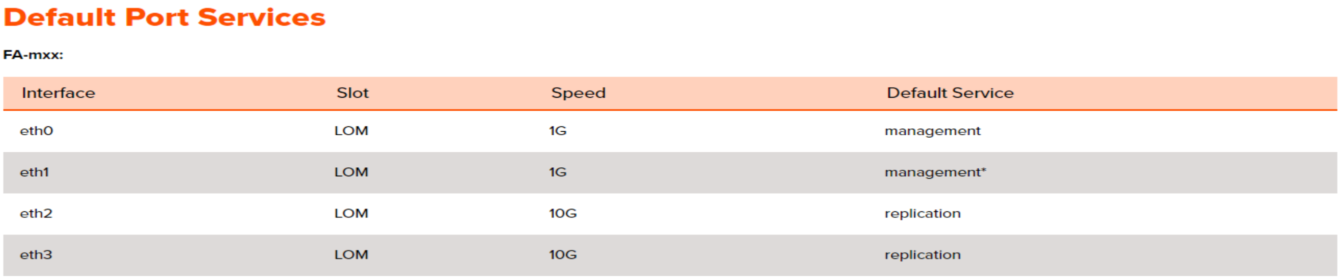

Default Port Services

FA-mxx:

Interface	Slot	Speed	Default Service
eth0	LOM	1G	management
eth1	LOM	1G	management*
eth2	LOM	10G	replication
eth3	LOM	10G	replication

The default port configurations for SAN switch connectivity can be found below, as well as further expansion options. Note that the port configurations of the new //m10 array is identical to the //m20 array.

Figure 31 Pure Storage FlashArray//m Ports Per Controller Overview

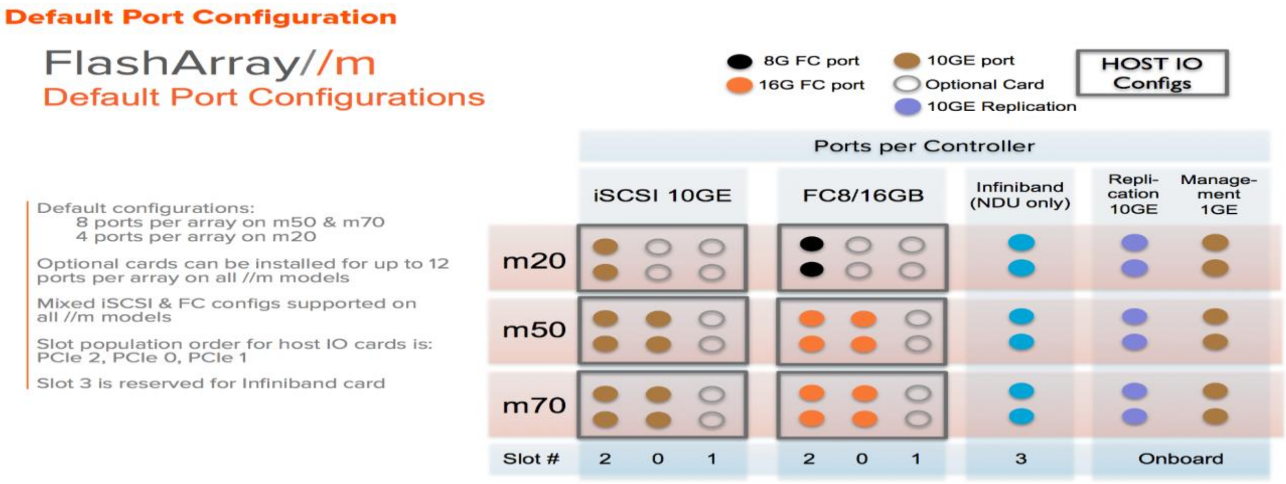

Default Port Configuration

FlashArray//m

Default Port Configurations

⚫ 8G FC port
🟠 16G FC port
🟤 10GE port
⚪ Optional Card
🟣 10GE Replication

HOST IO Configs

Default configurations:
8 ports per array on m50 & m70
4 ports per array on m20

Optional cards can be installed for up to 12 ports per array on all //m models

Mixed iSCSI & FC configs supported on all //m models

Slot population order for host IO cards is: PCIe 2, PCIe 0, PCIe 1

Slot 3 is reserved for Infiniband card

Ports per Controller

	iSCSI 10GE			FC8/16GB			Infiniband (NDU only)	Repli-cation 10GE	Manage-ment 1GE
m20	🟤 🟤	⚪ ⚪	⚪ ⚪	⚫ ⚫	⚪ ⚪	⚪ ⚪	🔵 🔵	🟣 🟣	🟤 🟤
m50	🟤 🟤	🟤 🟤	⚪ ⚪	🟠 🟠	🟠 🟠	⚪ ⚪	🔵 🔵	🟣 🟣	🟤 🟤
m70	🟤 🟤	🟤 🟤	⚪ ⚪	🟠 🟠	🟠 🟠	⚪ ⚪	🔵 🔵	🟣 🟣	🟤 🟤
Slot #	2	0	1	2	0	1	3	Onboard	

The rear of the expansion shelf shown below includes redundant SAS connectivity to each controller and the ability to daisy-chain additional shelfs for added capacity non-disruptively.

Figure 32 Pure Storage FlashArray External Shelf SAS Connectivity for each Controller Overview

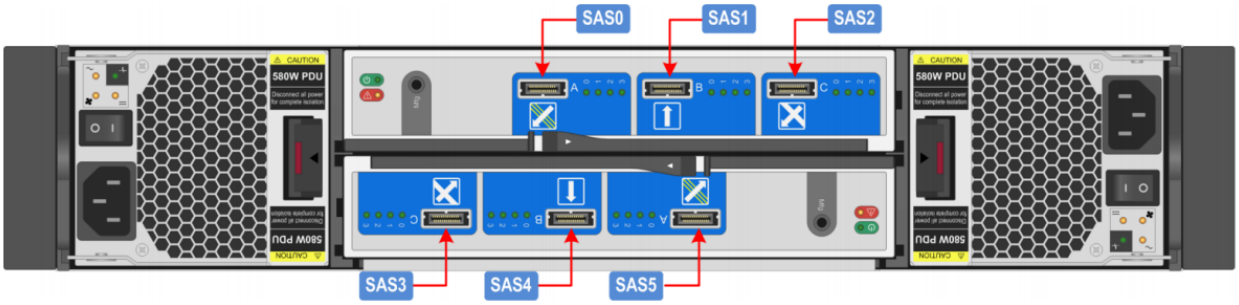

Connectivity to Cisco MDS 9148S

The connectivity diagram below illustrates the FlashArray//m50 interconnections to the 44TB expansion shelf as well as the 16GB FC connections to the redundantly paired Cisco MDS 9148S switches, yet again taking care to make certain there is no single point of path failure within the Cisco Validated Design. Not shown in this diagram are power and array/SAN management port connections.

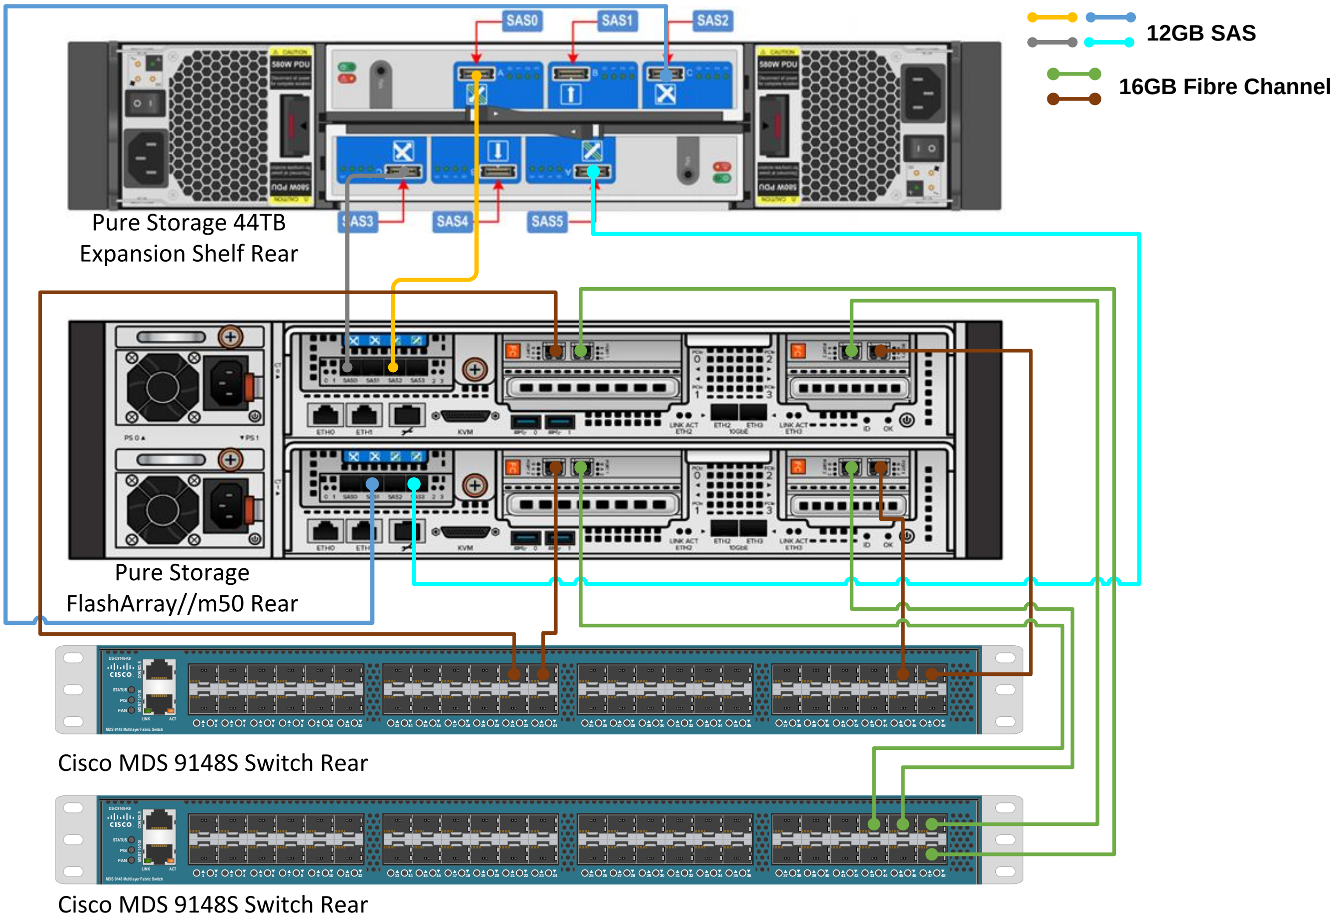
SAS0
SAS1
SAS2
SAS3
SAS4
SAS5
12GB SAS
16GB Fibre Channel
Pure Storage 44TB Expansion Shelf Rear
Pure Storage FlashArray//m50 Rear
Cisco MDS 9148S Switch Rear
Cisco MDS 9148S Switch Rear

All Pure Storage FlashArray//m systems require a minimum of three IP addresses for management. One IP address is assigned to each controller Ethernet port (eth0) and a virtual IP is assigned between the two controllers. Initial setup of the array and assignment of these IP addresses is accomplished by connecting to the console via either of the KVM ports shown above on the back of the array.

All GUI-based operations shown in this section can also be accomplished through the Purity command line interface. For the sake of consistency, we will show all configuration steps via the GUI but the array is so simple to use that the entire user manual fits on a double-sided folded business card as shown below in Figure 33. In addition, Purity features a restful API for additional extensibility and usage across a multitude of platforms and languages including PowerShell, Python and the VMware vRealize Automation Suite, among many others.

The FlashArray//m50 ships with a card showing key command line options for getting started.

Figure 33 Pure Storage FlashArray//m50 Getting Started and Troubleshooting Card

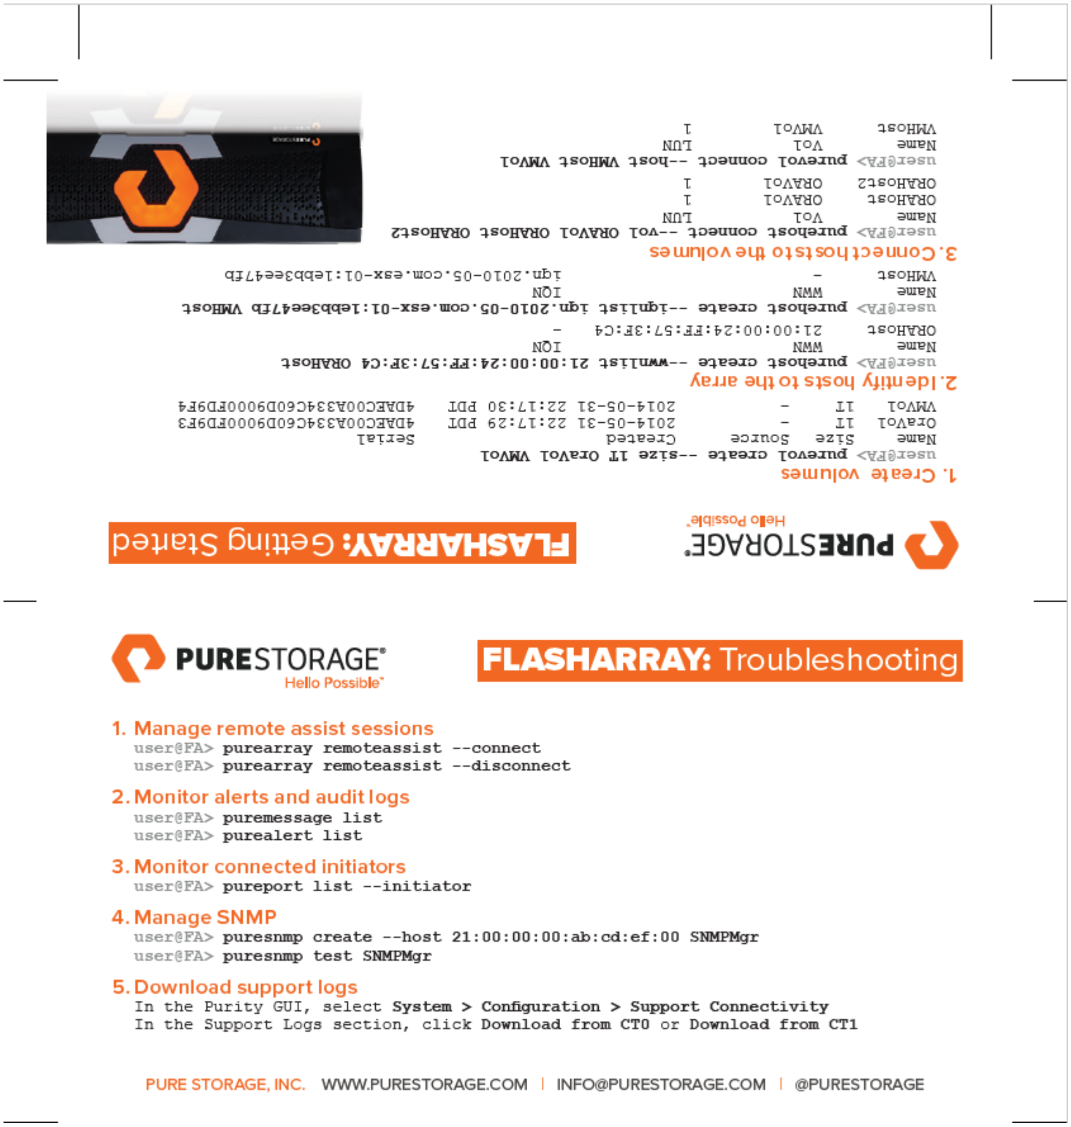

PURESTORAGE® Hello Possible™

FLASHARRAY: Getting Started

1. Create volumes

```
user@FA> purevol create --size 1T OraVol VMVol
Name     Size  Source     Created                  Serial
OraVol   1T    -          2014-05-31 22:17:29 PDT  4DAEC00A334C60D9000FD9F3
VMVol    1T    -          2014-05-31 22:17:30 PDT  4DAEC00A334C60D9000FD9F4
```

2. Identify hosts to the array

```
user@FA> purehost create --wwnlist 21:00:00:24:FF:57:3F:C4 ORAHost
Name       WWN                              IQN
ORAHost    21:00:00:24:FF:57:3F:C4          -

user@FA> purehost create --iqnlist iqn.2010-05.com.esx-01:1ebb3ee47fb VMHost
Name       WWN                              IQN
VMHost     -                                iqn.2010-05.com.esx-01:1ebb3ee47fb
```

3. Connect hosts to the volumes

```
user@FA> purehost connect --vol ORAVol ORAHost ORAHost2
Name       Vol        LUN
ORAHost    ORAVol     1
ORAHost2   ORAVol     1

user@FA> purevol connect --host VMHost VMVol
Name       Vol        LUN
VMHost     VMVol      1
```

PURESTORAGE® Hello Possible™

FLASHARRAY: Troubleshooting

1. Manage remote assist sessions

```
user@FA> purearray remoteassist --connect
user@FA> purearray remoteassist --disconnect
```

2. Monitor alerts and audit logs

```
user@FA> puremessage list
user@FA> purealert list
```

3. Monitor connected initiators

```
user@FA> pureport list --initiator
```

4. Manage SNMP

```
user@FA> puresnmp create --host 21:00:00:00:ab:cd:ef:00 SNMPMgr
user@FA> puresnmp test SNMPMgr
```

5. Download support logs

```
In the Purity GUI, select System > Configuration > Support Connectivity
In the Support Logs section, click Download from CT0 or Download from CT1
```

PURE STORAGE, INC. WWW.PURESTORAGE.COM | INFO@PURESTORAGE.COM | @PURESTORAGE

Figure 34 Pure Storage FlashArray//m50 Daily Management Reference

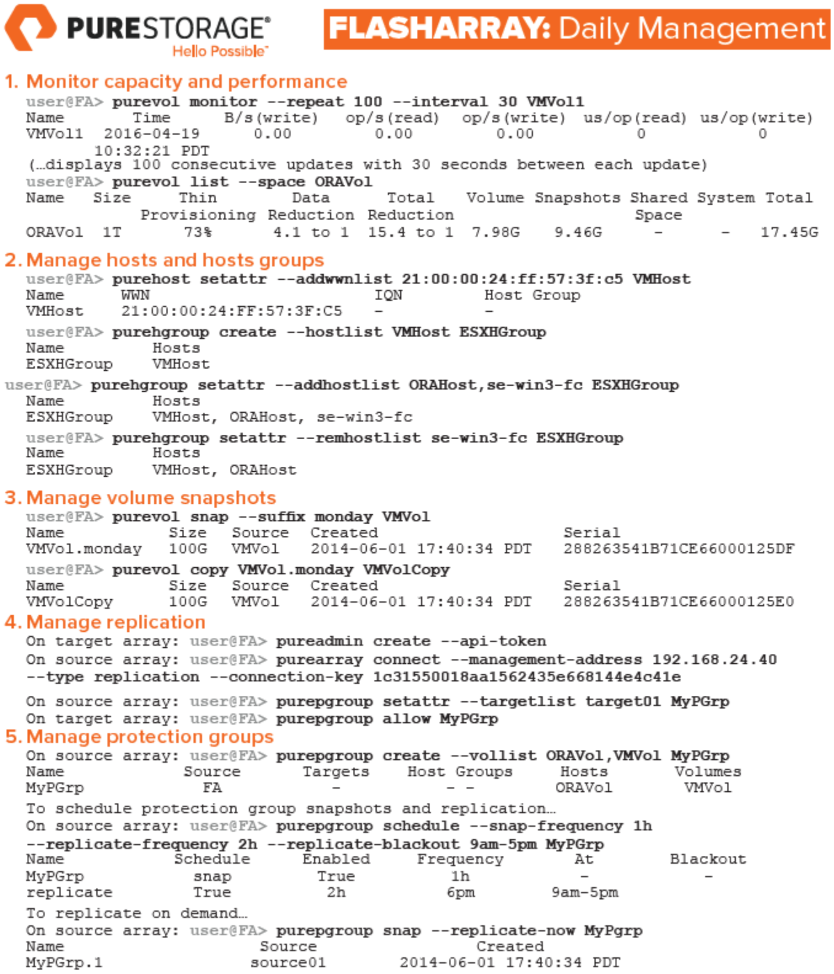

PURESTORAGE® Hello Possible™

FLASHARRAY: Daily Management

1. Monitor capacity and performance

```
user@FA> purevol monitor --repeat 100 --interval 30 VMVol1
Name       Time       B/s(write)   op/s(read)   op/s(write)  us/op(read)  us/op(write)
VMVol1  2016-04-19      0.00         0.00         0.00           0            0
       10:32:21 PDT
(…displays 100 consecutive updates with 30 seconds between each update)
user@FA> purevol list --space ORAVol
Name   Size    Thin       Data      Total    Volume Snapshots Shared System Total
            Provisioning Reduction Reduction                  Space
ORAVol  1T      73%      4.1 to 1  15.4 to 1  7.98G    9.46G     -      -    17.45G
```

2. Manage hosts and hosts groups

```
user@FA> purehost setattr --addwwnlist 21:00:00:24:ff:57:3f:c5 VMHost
Name      WWN                       IQN         Host Group
VMHost    21:00:00:24:FF:57:3F:C5   -           -

user@FA> purehgroup create --hostlist VMHost ESXHGroup
Name         Hosts
ESXHGroup    VMHost
user@FA> purehgroup setattr --addhostlist ORAHost,se-win3-fc ESXHGroup
Name         Hosts
ESXHGroup    VMHost, ORAHost, se-win3-fc

user@FA> purehgroup setattr --remhostlist se-win3-fc ESXHGroup
Name         Hosts
ESXHGroup    VMHost, ORAHost
```

3. Manage volume snapshots

```
user@FA> purevol snap --suffix monday VMVol
Name           Size  Source  Created                     Serial
VMVol.monday   100G  VMVol   2014-06-01 17:40:34 PDT     288263541B71CE66000125DF

user@FA> purevol copy VMVol.monday VMVolCopy
Name           Size  Source  Created                     Serial
VMVolCopy      100G  VMVol   2014-06-01 17:40:34 PDT     288263541B71CE66000125E0
```

4. Manage replication

```
On target array: user@FA> pureadmin create --api-token
On source array: user@FA> purearray connect --management-address 192.168.24.40
--type replication --connection-key 1c31550018aa1562435e668144e4c41e

On source array: user@FA> purepgroup setattr --targetlist target01 MyPGrp
On target array: user@FA> purepgroup allow MyPGrp
```

5. Manage protection groups

```
On source array: user@FA> purepgroup create --vollist ORAVol,VMVol MyPGrp
Name             Source      Targets    Host Groups     Hosts       Volumes
MyPGrp            FA            -           - -         ORAVol       VMVol

To schedule protection group snapshots and replication…
On source array: user@FA> purepgroup schedule --snap-frequency 1h
--replicate-frequency 2h --replicate-blackout 9am-5pm MyPGrp
Name            Schedule      Enabled     Frequency        At          Blackout
MyPGrp            snap          True         1h            -              -
replicate         True          2h           6pm        9am-5pm

To replicate on demand…
On source array: user@FA> purepgroup snap --replicate-now MyPgrp
Name                   Source             Created
MyPGrp.1               source01       2014-06-01 17:40:34 PDT
```

Pure Storage GUI Overview and ESXi Host Setup

When initial setup is complete and management IP addresses have been assigned via the console, users can then access the Pure GUI from almost any modern web browser.

To access the Pure GUI, complete the following steps:

1. Navigate to an IP address assigned for array management from the web browser. The following login screen should appear:

Figure 35 Pure Storage Web GUI or Login

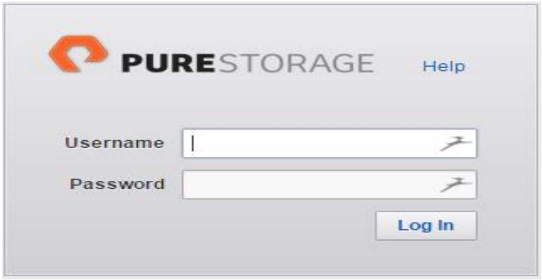

User name: pureuser

Password: pureuser

After login, the Pure GUI should appear. Note that a newly deployed array will have 0% space utilization and not be driving any IO.

Figure 36   PureStorage GUI

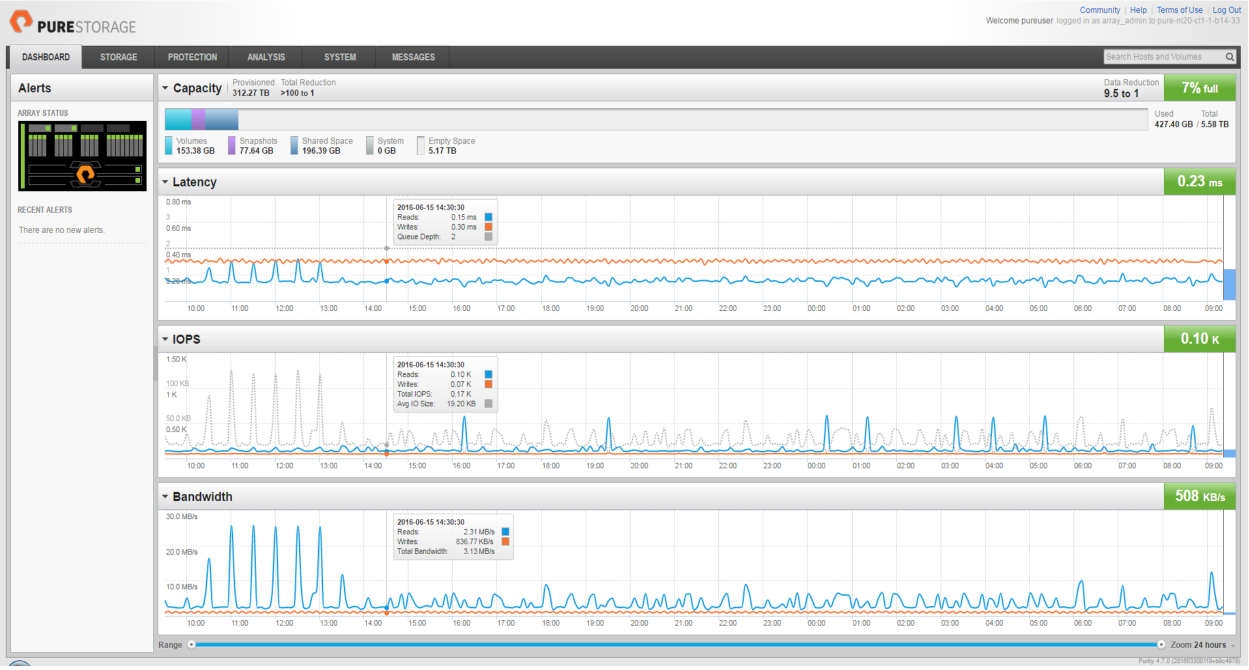

Welcome pureuser logged in as array_admin to pure-m20-ct1-1-b14-33

There are no new alerts.

Capacity: Provisioned 312.27 TB | Total Reduction >100 to 1 | Data Reduction 9.5 to 1 | 7% full

Used 427.40 GB / Total 5.58 TB

Volumes 153.38 GB | Snapshots 77.64 GB | Shared Space 196.39 GB | System 0 GB | Empty Space 5.17 TB

Latency: 0.23 ms

2016-06-15 14:30:30
Reads: 0.15 ms
Writes: 0.30 ms
Queue Depth: 2

IOPS: 0.10 K

2016-06-15 14:30:30
Reads: 0.10 K
Writes: 0.07 K
Total IOPS: 0.17 K
Avg IO Size: 19.20 KB

Bandwidth: 508 KB/s

2016-06-15 14:30:30
Reads: 2.31 MB/s
Writes: 836.77 KB/s
Total Bandwidth: 3.13 MB/s

Zoom 24 hours

Purity 4.7.0 (201603300119+b9c4978)

The Pure Storage GUI shows a variety of useful performance statistics and overall status for the array. On the left, the Array Status can be seen which shows a top-level health status of the backplane, NVRam modules, controllers and individual solid state drives. To the right of that, the amount of storage space being used, data reduction and percentage of space in use is shown across the top. Below that, key metrics including Latency, IOPs and Bandwidth are shown in real-time and the zoom can be adjusted to as fine of an interval as 15 minutes or as wide of an interval of 30 days. You have the capability to see more granular performance and capacity metrics for individual, or groups of datastores over a wider timespan under the ‘Analysis’ tab for up to a year.

When the array has been brought online, networking information can be confirmed by going to the ‘System’ tab which will show the following window. Note that this window shows a much more detailed health report for all array components.

Figure 37 FlashArray//m50 System Tab

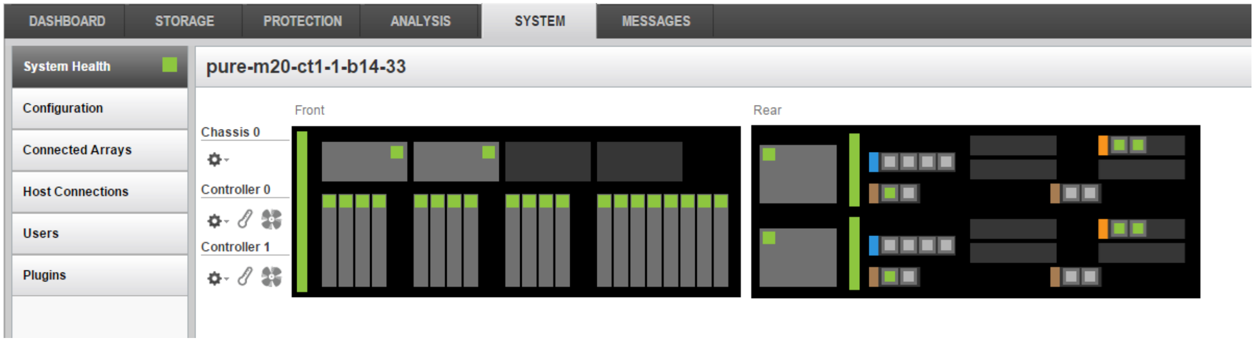

2. Clicking the 'Configuration' icon above will expand the menu and clicking the 'Networking' option will bring up the management networking configuration screen:

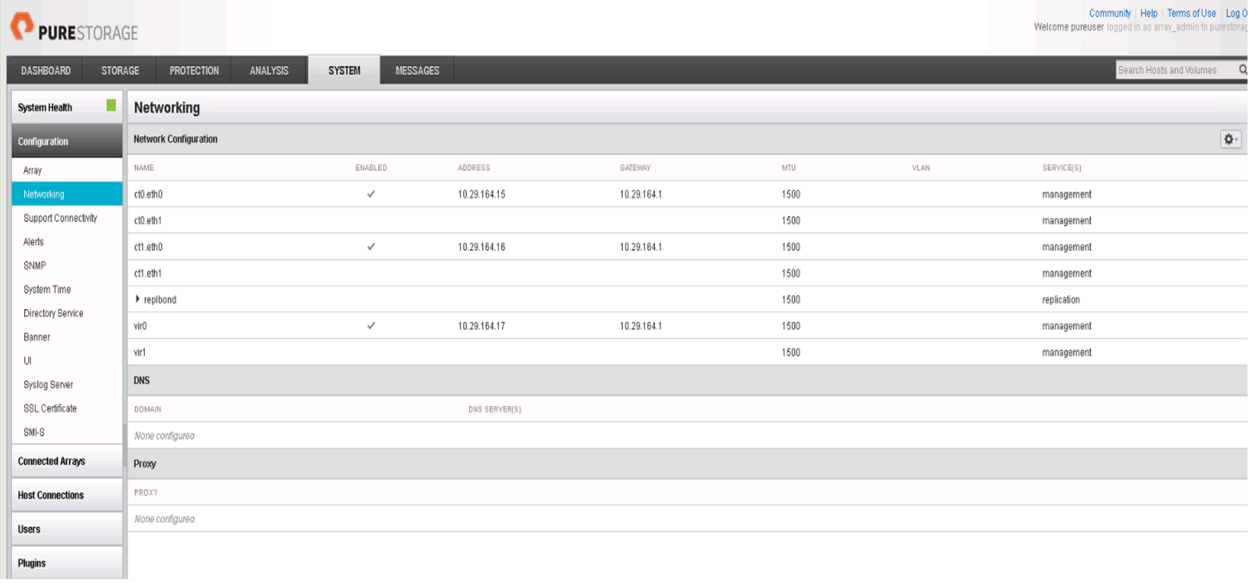

Networking

Network Configuration

NAME	ENABLED	ADDRESS	GATEWAY	MTU	VLAN	SERVICE(S)
ct0.eth0	✓	10.29.164.15	10.29.164.1	1500		management
ct0.eth1				1500		management
ct1.eth0	✓	10.29.164.16	10.29.164.1	1500		management
ct1.eth1				1500		management
▸ replbond				1500		replication
vir0	✓	10.29.164.17	10.29.164.1	1500		management
vir1				1500		management

DNS

DOMAIN	DNS SERVER(S)
None configured	

Proxy

PROXY
None configured

Networking information can be updated by hovering the mouse cursor over the Ethernet port you wish to change and clicking the 'edit' option on the gear that appears. The following screen shots an example of editing a network interface:

Figure 38 Edit Network Interface

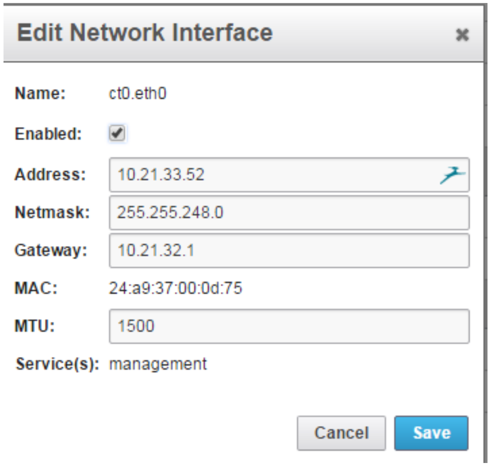

Edit Network Interface

Name: ct0.eth0

Enabled: ☑

Address: 10.21.33.52

Netmask: 255.255.248.0

Gateway: 10.21.32.1

MAC: 24:a9:37:00:0d:75

MTU: 1500

Service(s): management

When the management network has been confirmed as being properly setup, the next step is to complete the SAN zoning of the Cisco MCS 9148S switches so that the storage array and UCS servers are able to communicate.

3. Record the WWN's of the array from within the Pure GUI. This can be found under System -> Configuration -> Host Connections as shown below:

Figure 39 Host Connections

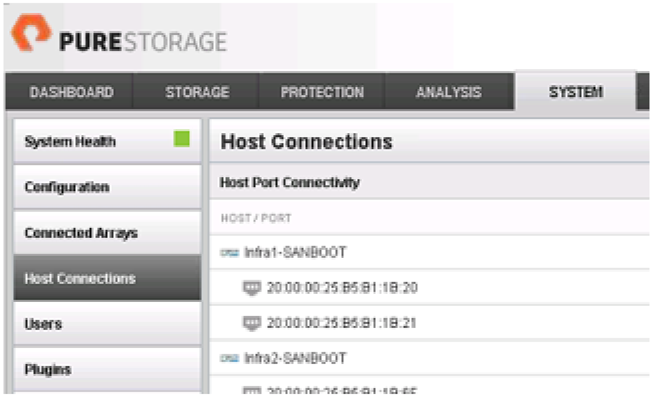

Host Connections

Host Port Connectivity

HOST / PORT
Infra1-SANBOOT
20:00:00:25:B5:B1:1B:20
20:00:00:25:B5:B1:1B:21
Infra2-SANBOOT

4. The WWN's of the Pure Storage array are displayed. Note that only 6 out of 8 are shown below. Record your unique WWN values so that they can be included in the Cisco MDS zonesets along with the WWNs that were created when the UCS Service Profiles were built.

Figure 40 WWNs on the FlashArray//m50

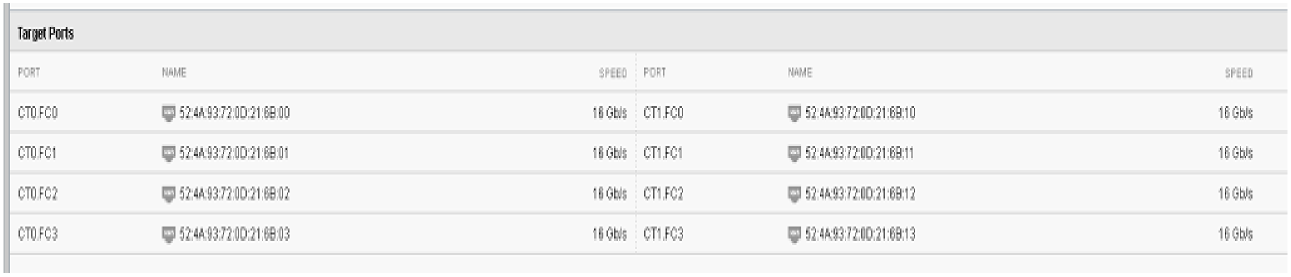

Target Ports

PORT	NAME	SPEED	PORT	NAME	SPEED
CT0.FC0	52:4A:93:72:0D:21:6B:00	16 Gb/s	CT1.FC0	52:4A:93:72:0D:21:6B:10	16 Gb/s
CT0.FC1	52:4A:93:72:0D:21:6B:01	16 Gb/s	CT1.FC1	52:4A:93:72:0D:21:6B:11	16 Gb/s
CT0.FC2	52:4A:93:72:0D:21:6B:02	16 Gb/s	CT1.FC2	52:4A:93:72:0D:21:6B:12	16 Gb/s
CT0.FC3	52:4A:93:72:0D:21:6B:03	16 Gb/s	CT1.FC3	52:4A:93:72:0D:21:6B:13	16 Gb/s

When the Cisco MDS 8148S zonesets are created and activated, the WWNs of the UCS servers should automatically become visible to the Pure Storage array (note that it can sometimes take up to 30 minutes for the UCS WWNs to become visible).

Adding a Cisco UCS Host

The next step is to create hosts within the Pure Storage GUI. Multiple hosts can be grouped together into Host Groups to provide a higher level of abstraction for management. Volumes need to be connected to hosts and/or host groups in order for to two to communicate.

To create a Host within the Pure Storage GUI, complete the following steps:

1. Navigate to the storage tab.

Figure 41   Add Hosts to the FlashArray//m50

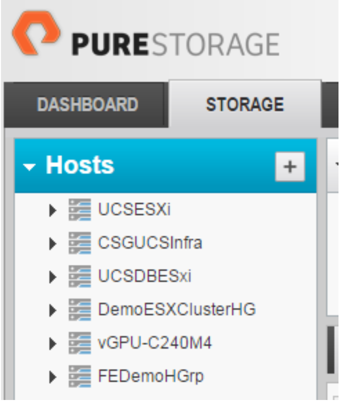

2. Click the '+' sign; you can create single, or multiple hosts. Since you are using multiple hosts in the design, you will create several hosts at once.

Figure 42   Create a Single Host

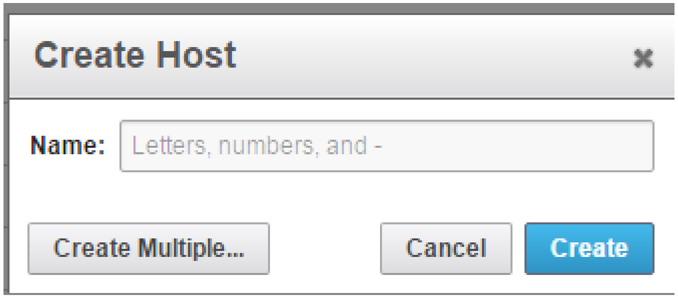

3. Create two hosts.

Figure 43 Create Multiple Hosts

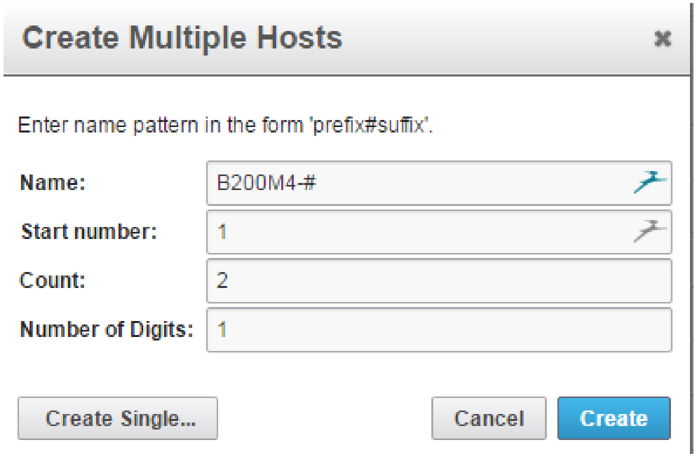

Create Multiple Hosts

Enter name pattern in the form 'prefix#suffix'.

Name:	B200M4-#
Start number:	1
Count:	2
Number of Digits:	1

4. In the figure below, you can see that two hosts have been created and the next step is to connect their WWNs. WWNs can be matched up with specific Cisco UCS servers from within Cisco UCS Manager.

Figure 44 Hosts created

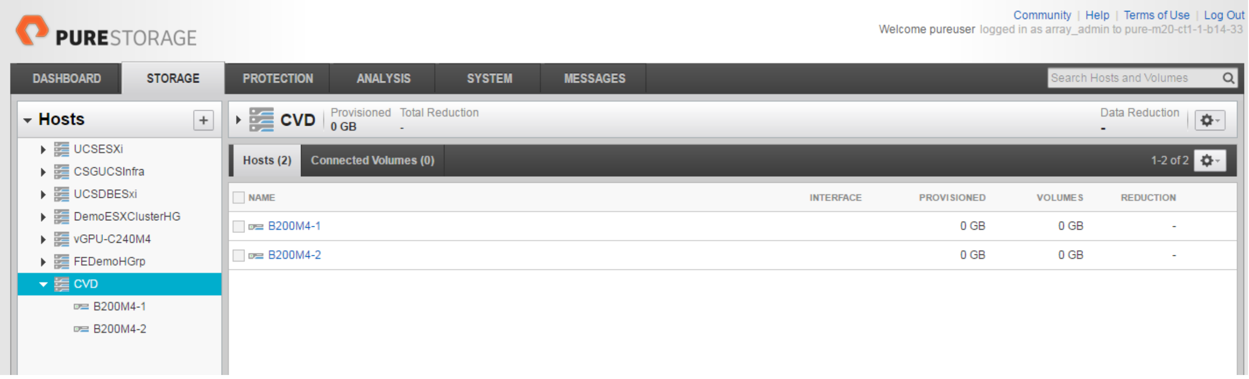

Welcome pureuser logged in as array_admin to pure-m20-ct1-1-b14-33

CVD Provisioned 0 GB Total Reduction - Data Reduction -

1-2 of 2

NAME	INTERFACE	PROVISIONED	VOLUMES	REDUCTION
□ B200M4-1		0 GB	0 GB	-
□ B200M4-2		0 GB	0 GB	-

5. To configure the FC WWNs, highlight a host, select the ‘Host Ports’ tab and then select the ‘Configure Fibre Channel WWNs’ button as shown in the figure below:

Figure 45 Configure host WWNs

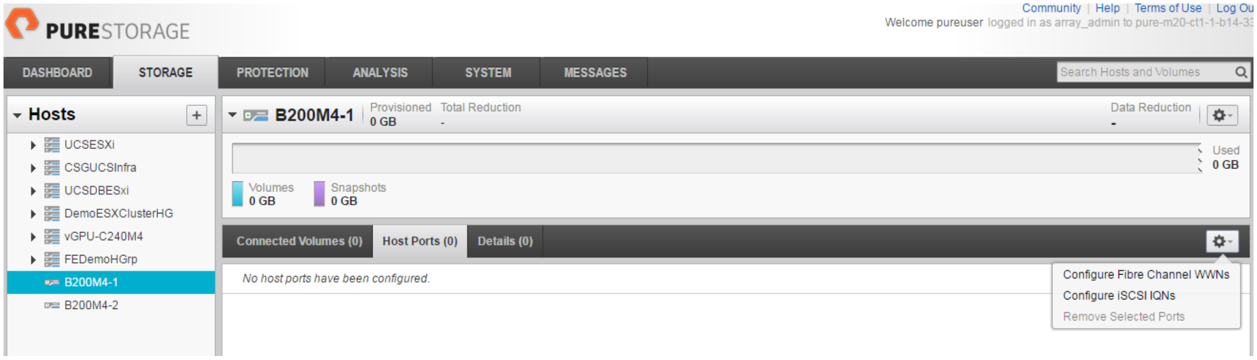

Welcome pureuser logged in as array_admin to pure-m20-ct1-1-b14-33

B200M4-1 Provisioned 0 GB Total Reduction - Data Reduction -

Used 0 GB

Volumes 0 GB

Snapshots 0 GB

No host ports have been configured.

The available WWNs will appear on the left-hand column.

Figure 46   Available WWNs

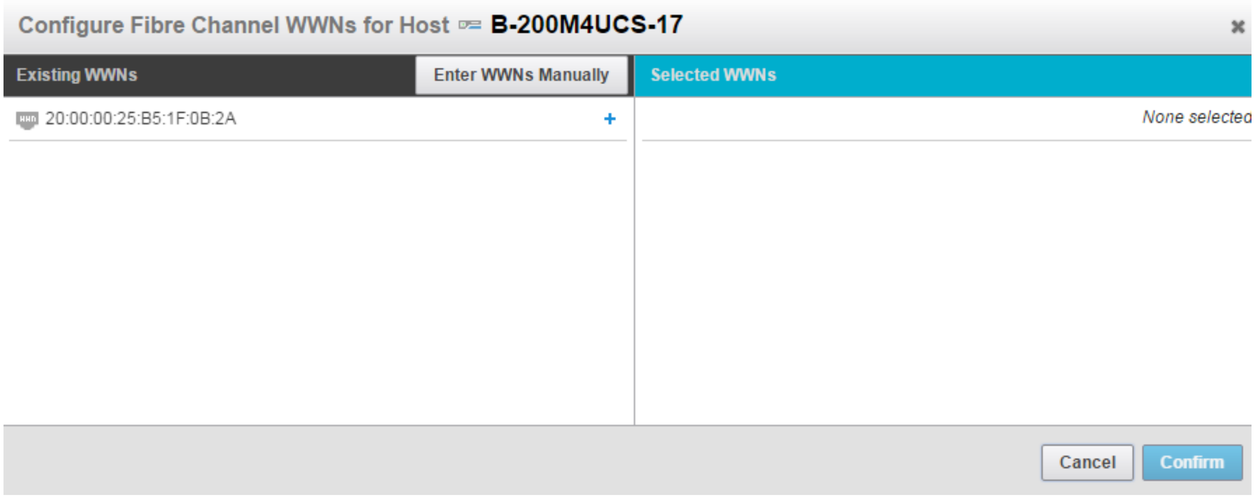
Configure Fibre Channel WWNs for Host B-200M4UCS-17

6. Click the WWNs you want to select and move them over to the right-hand column and click Confirm when completed. In this Cisco Validated Design, each UCS host has two separate vHBAs to provide both performance and resiliency.

Below, you can see a properly setup UCS host within the Pure GUI. Repeat the steps above for to create hosts and configuring the Fibre Channel WWNs for all UCS servers.

Figure 47   Host with Two WWNs

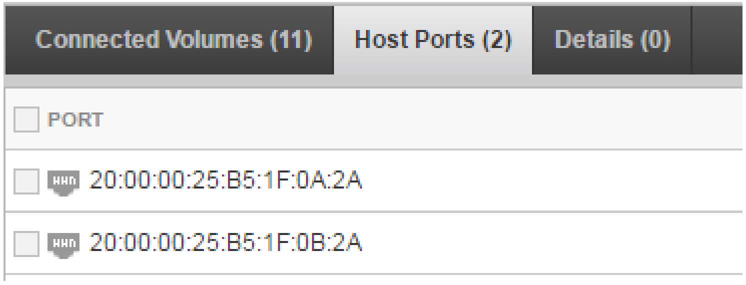

□ PORT
□ 20:00:00:25:B5:1F:0A:2A
□ 20:00:00:25:B5:1F:0B:2A

7. When the hosts are created and the WWNs have been assigned to them, click the Hosts button and select Create Host Group.

Figure 48   Create Host Group

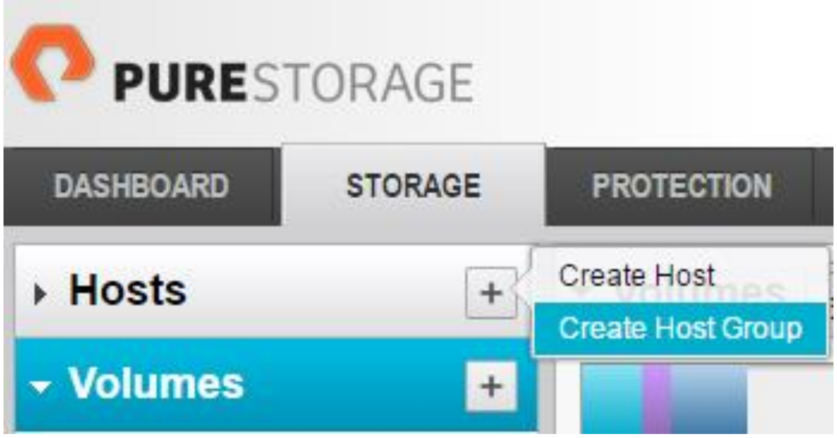

Figure 49 Create Host Group Name

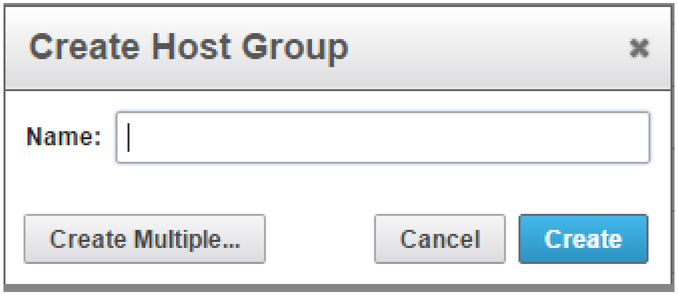

8. Add the 2 example hosts previously created to this newly created Host Group. To accomplish this, highlight the Host Group, click the gear to the right and select Add Hosts.

Figure 50 Add Hosts to Host Group

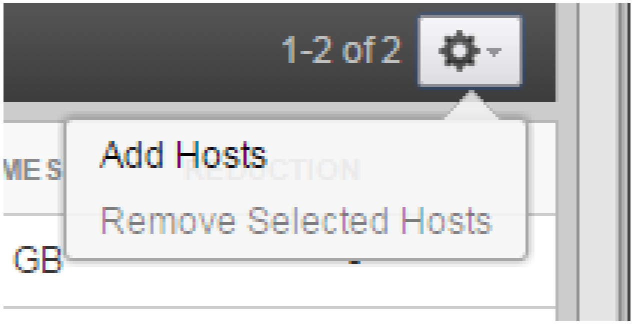

9. Click the hosts you wish to add to the Host Group on the left column and move them to the right. When all hosts have been included, click Confirm.

Figure 51 Add Hosts to Host Group by Selecting Existing Hosts

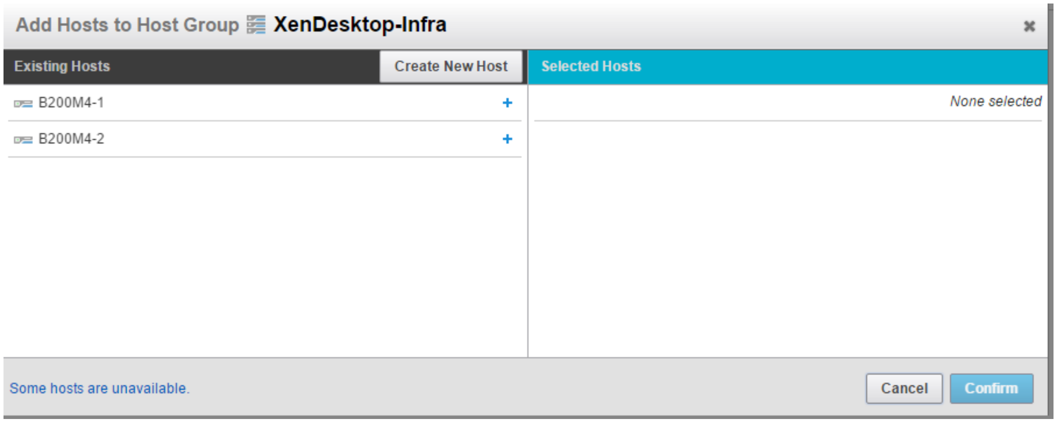

Figure 52 Hosts Added to Host Group

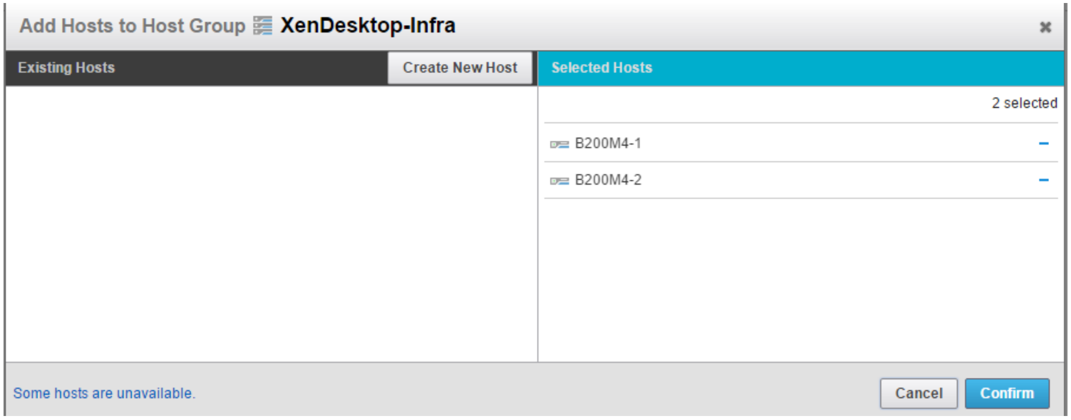

Some hosts are unavailable.

The screenshot below shows the populated Host Group with two hosts as members:

Figure 53 Populated Host Group

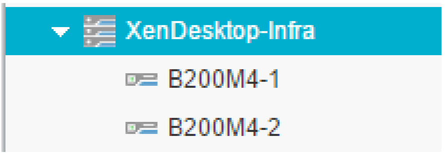

Repeat this procedure for all Host Groups that are required. For this Cisco Validated Design we created three separate Host Groups for the following components:

- Citrix XenDesktop Infrastructure (2 Hosts)
- RDS Servers (12 Hosts)
- Citrix XenDesktop PVS Desktops (16 Hosts)

With the Hosts and Host Group setup completed, the array is now ready to serve Citrix XenDesktop virtual machines and the associated infrastructure.

Pure Storage Data Storage Layout

A key benefit of the Pure Storage FlashArray//m is that it rapidly enables customers to quickly and logically design their Citrix XenDesktop datastore implementation. The Purity Operating Environment automatically handles 100% of the tuning, encryption, SSD wear-leveling and data resiliency (among other things) at the array level so Citrix administrators only need to name the datastore, input the size and then attach it to the appropriate host group(s). In addition, the Pure Storage array includes a vSphere Web-client plug-in that enables the administration of the array completely within the vSphere Web-client if they so choose.

For this Cisco Validated Design, we will break up our datastores by logical use case. This is often an overlooked benefit of an All-Flash Array – the ability to design your datastores with a logical layout rather than being forced to adhere to VM per datastore sizing limitations based upon storage performance bottlenecks. That is, several thousand VMs can be run from a single datastore without issue, though we recommend including fault tolerance into your design.

For the Citrix XenDesktop implementation used in this document we elected to use the following datastore setup (Table 13):

Table 13 Layout for the Pure Storage FlashArray//m50

VDI Component	Datastore Contents	# of Datastores	Datastore Size
Citrix XenDesktop Infrastructure Datastore	2 Active Directory/DNS/DHCP Servers 2 XenDesktop Delivery Controller Servers 2 StoreFront Servers 3 XenDesktop PVS Servers 3 Windows File Servers 1 VMware vCenter 6 Appliance 1 SQL Database Server 2 VSM Appliances (Cisco Nexus 1000V) 1 VSUM (Cisco Virtual Switch Update Manager)	1	4 TB
RDS PVS Datastore	60 Windows Server 2012R2 RDS Servers	1	10 TB
Non Persistent MCS Desktop Datastore	1200 Non-Persistent PVS desktops	1	30 TB
Persistent MCS Desktop Datastore	1200 Persistent PVS desktops	1	30 TB
PVS vDisk Datastore	1 Windows 10x64 vDisk 1 Windows 7x64 vDisk	1	1 TB
User HomeDrives Datastore	5000 User Homedrives served from Windows2012R2 Server	1	5 TB
User Profiles Datastore	5000 User Profiles served from Windows 2012R2 Server	1	5 TB
ESXi Boot Datastore	ESXi for UCS servers	30	10 GB
VSI Logs	ESXTOP log repository	1	2 TB

Figure 54 through Figure 57 provides examples of datastores created for data and for SAN boot.

Figure 54 Data stores Created for Data and for SAN Boot

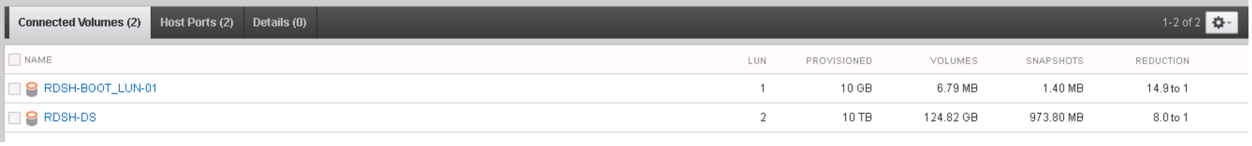

NAME	LUN	PROVISIONED	VOLUMES	SNAPSHOTS	REDUCTION
RDSH-BOOT_LUN-01	1	10 GB	6.79 MB	1.40 MB	14.9 to 1
RDSH-DS	2	10 TB	124.82 GB	973.80 MB	8.0 to 1

Figure 55 Datastores Created for Data and for SAN Boot

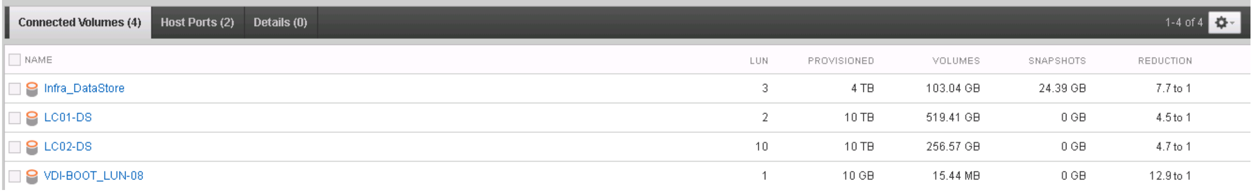

NAME	LUN	PROVISIONED	VOLUMES	SNAPSHOTS	REDUCTION
Infra_DataStore	3	4 TB	103.04 GB	24.39 GB	7.7 to 1
LC01-DS	2	10 TB	519.41 GB	0 GB	4.5 to 1
LC02-DS	10	10 TB	256.57 GB	0 GB	4.7 to 1
VDI-BOOT_LUN-08	1	10 GB	15.44 MB	0 GB	12.9 to 1

Figure 56   Datastores Created for Data and for SAN Boot

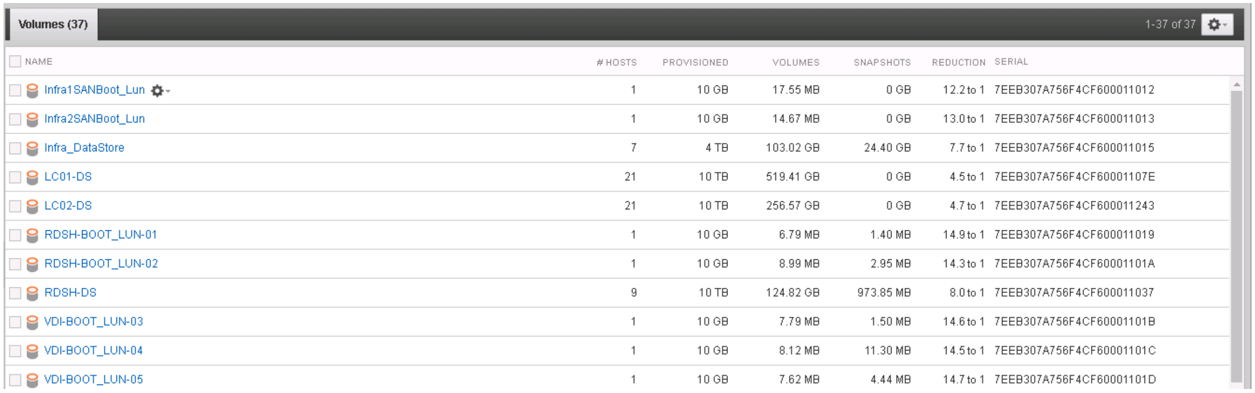

Volumes (37)   1-37 of 37

NAME	# HOSTS	PROVISIONED	VOLUMES	SNAPSHOTS	REDUCTION	SERIAL
Infra1SANBoot_Lun	1	10 GB	17.55 MB	0 GB	12.2 to 1	7EEB307A756F4CF600011012
Infra2SANBoot_Lun	1	10 GB	14.67 MB	0 GB	13.0 to 1	7EEB307A756F4CF600011013
Infra_DataStore	7	4 TB	103.02 GB	24.40 GB	7.7 to 1	7EEB307A756F4CF600011015
LC01-DS	21	10 TB	519.41 GB	0 GB	4.5 to 1	7EEB307A756F4CF60001107E
LC02-DS	21	10 TB	256.57 GB	0 GB	4.7 to 1	7EEB307A756F4CF600011243
RDSH-BOOT_LUN-01	1	10 GB	6.79 MB	1.40 MB	14.9 to 1	7EEB307A756F4CF600011019
RDSH-BOOT_LUN-02	1	10 GB	8.99 MB	2.95 MB	14.3 to 1	7EEB307A756F4CF60001101A
RDSH-DS	9	10 TB	124.82 GB	973.85 MB	8.0 to 1	7EEB307A756F4CF600011037
VDI-BOOT_LUN-03	1	10 GB	7.79 MB	1.50 MB	14.6 to 1	7EEB307A756F4CF60001101B
VDI-BOOT_LUN-04	1	10 GB	8.12 MB	11.30 MB	14.5 to 1	7EEB307A756F4CF60001101C
VDI-BOOT_LUN-05	1	10 GB	7.62 MB	4.44 MB	14.7 to 1	7EEB307A756F4CF60001101D

Figure 57   Pure storage Dashboard on login to Web GUI

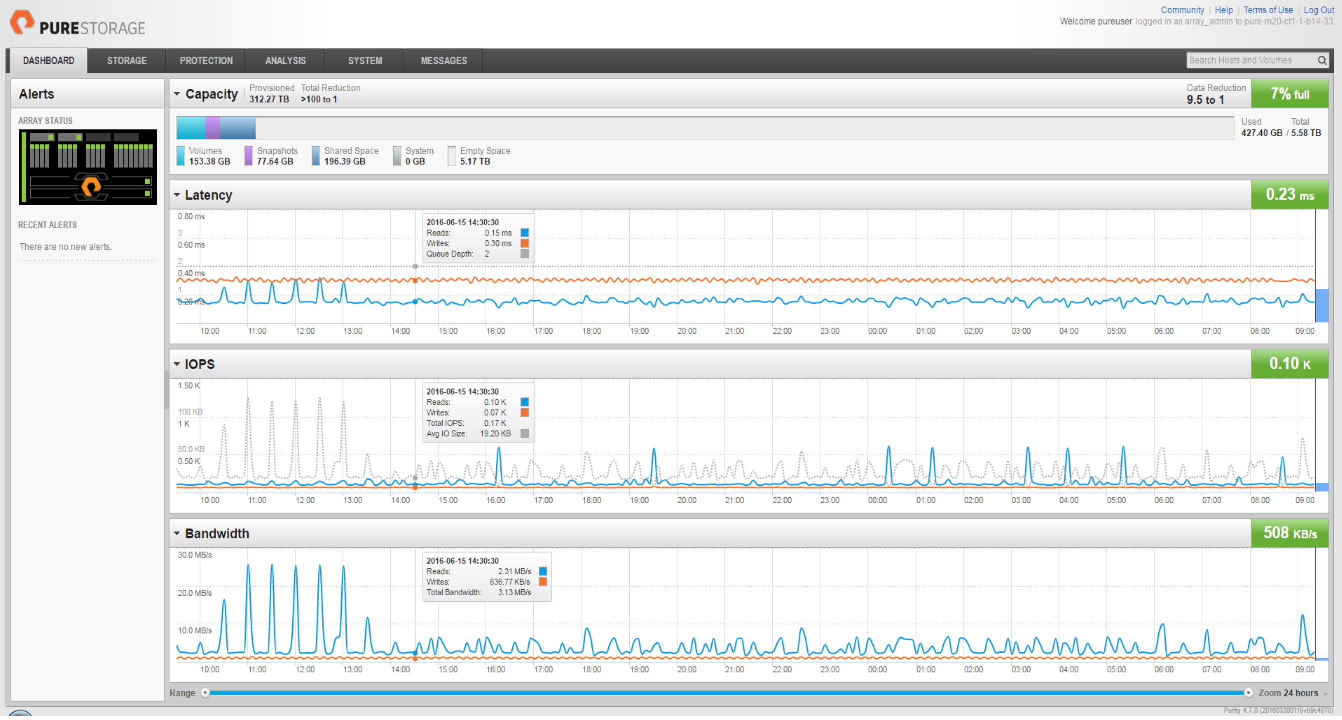

There are no new alerts.

Capacity — Provisioned 312.27 TB, Total Reduction >100 to 1; Data Reduction 9.5 to 1; 7% full; Used 427.40 GB / Total 5.58 TB

Volumes 153.38 GB; Snapshots 77.64 GB; Shared Space 196.39 GB; System 0 GB; Empty Space 5.17 TB

Latency 0.23 ms — 2016-06-15 14:30:30; Reads: 0.15 ms; Writes: 0.30 ms; Queue Depth: 2

IOPS 0.10 K — 2016-06-15 14:30:30; Reads: 0.10 K; Writes: 0.07 K; Total IOPS: 0.17 K; Avg IO Size: 19.20 KB

Bandwidth 508 KB/s — 2016-06-15 14:30:30; Reads: 2.31 MB/s; Writes: 836.77 KB/s; Total Bandwidth: 3.13 MB/s

The Pure Storage GUI shows a variety of metrics and status for the array. On the left, the Array Status can be seen which shows a top-level health status of the backplane, NVRam modules, controllers and individual solid state drive bays. To the right, the amount of storage space being used, data reduction and percentage of space in use is shown across the top. Below, key metrics including Latency, IOPs and Bandwidth are shown in real-time and the zoom can be adjusted to as fine an interval as 15 minutes or as wide of an interval of 30 days. You have the capability to see additional performance and capacity metrics for individual, or groups of datastores over a wider timespan under the Analysis tab.

When the array is brought online, networking information can be confirmed by completing the following steps:

1. Go to the System tab. The initial view shows a much more detailed health view for all array components.

Figure 58 System tab to shows the Overall Storage Health

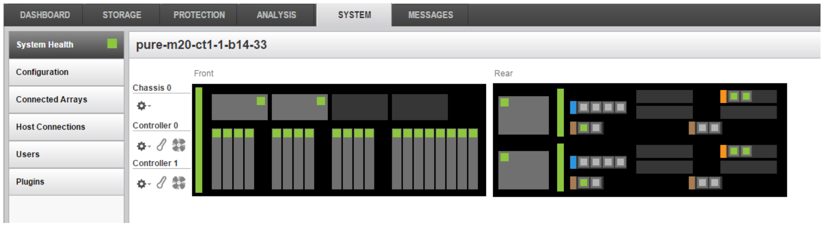

2. Click Configuration to expand the menu and click Networking to bring up the management networking configuration screen:

Figure 59 Pure Storage FlashArray//m50 Networking Information

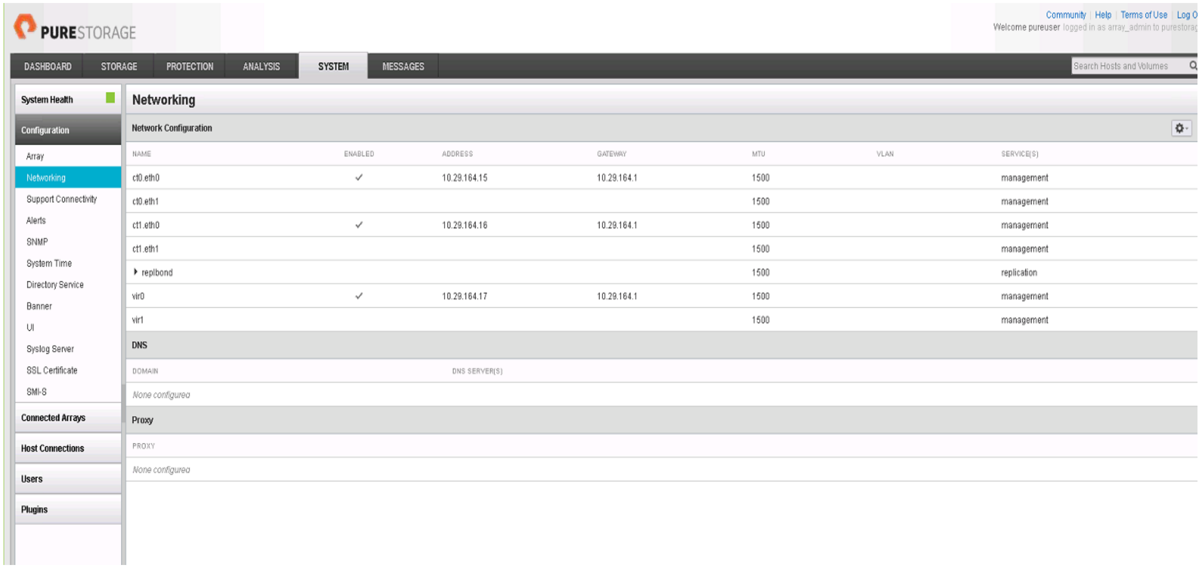

Networking

Network Configuration

NAME	ENABLED	ADDRESS	GATEWAY	MTU	VLAN	SERVICE(S)
ct0.eth0	✓	10.29.164.15	10.29.164.1	1500		management
ct0.eth1				1500		management
ct1.eth0	✓	10.29.164.16	10.29.164.1	1500		management
ct1.eth1				1500		management
▸ replbond				1500		replication
vir0	✓	10.29.164.17	10.29.164.1	1500		management
vir1				1500		management

DNS

DOMAIN	DNS SERVER(S)
None configured	

Proxy

PROXY
None configured

3. To update network information, hover the mouse cursor over the Ethernet port you wish to change and click Edit on the gear that appears. The following screenshots provide an example of editing a network interface:

Figure 60 Pure Storage FlashArray//m50 Networking Information Editing

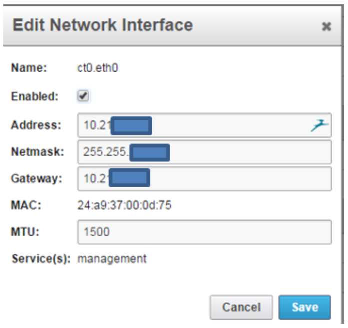

Edit Network Interface

Name: ct0.eth0

Enabled: ☑

Address: 10.2[redacted]

Netmask: 255.255.[redacted]

Gateway: 10.2[redacted]

MAC: 24:a9:37:00:0d:75

MTU: 1500

Service(s): management

When the management network is properly setup, the next step is to complete the SAN zoning of the Cisco MCS 9148S switches so that the storage array and Cisco UCS servers are able to communicate. To do the SAN zoning, complete the following steps:

1. Record the WWN's of the array from within the Pure GUI. This can be found under System > Configuration > Host Connections as shown below:

Figure 61 ESXi Host Connected View

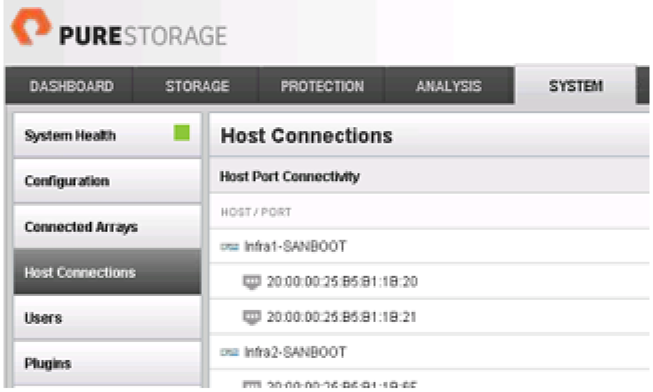

Host Connections

Host Port Connectivity

HOST / PORT

Infra1-SANBOOT

20:00:00:25:B5:B1:1B:20

20:00:00:25:B5:B1:1B:21

Infra2-SANBOOT

2. At the bottom of the above screen, the WWN's of the Pure Storage array are displayed. Note that only 6 out of 8 are being shown below. Record your unique WWN values so that they can be included in the Cisco MDS zonesets along with the WWNs that were created when the Cisco UCS Service Profiles were built.

Figure 62 Pure Storage Controllers Configured with WWN Overview

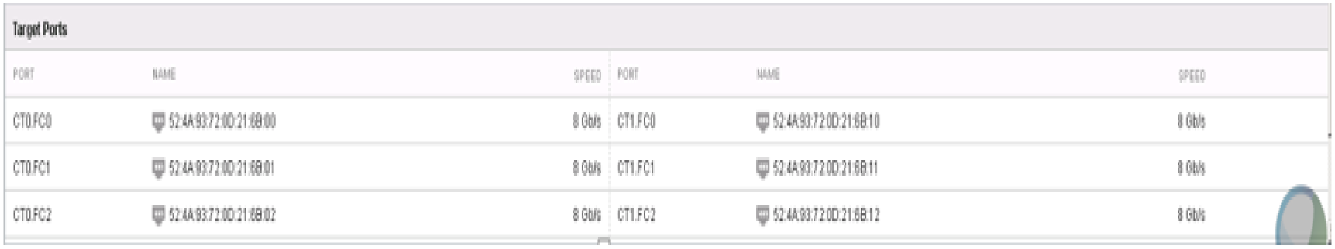

Target Ports

PORT	NAME	SPEED	PORT	NAME	SPEED
CT0.FC0	52:4A:93:72:0D:21:6B:00	8 Gb/s	CT1.FC0	52:4A:93:72:0D:21:6B:10	8 Gb/s
CT0.FC1	52:4A:93:72:0D:21:6B:01	8 Gb/s	CT1.FC1	52:4A:93:72:0D:21:6B:11	8 Gb/s
CT0.FC2	52:4A:93:72:0D:21:6B:02	8 Gb/s	CT1.FC2	52:4A:93:72:0D:21:6B:12	8 Gb/s

When the Cisco MDS 9148S zonesets have been created and activated, the WWNs of the Cisco UCS servers should automatically become visible to the Pure Storage array (note that it can sometimes take up to 30 minutes for the UCS WWNs to become visible).

Creating Boot LUNs for Cisco UCS B200 M4 Servers

Provisioning a volume for SAN boot for Pure Storage is simple. The only difference between a SAN volume and a vCenter datastore is that you will connect the SAN boot volume to the single host that will be booting from the volume.

To create Boot LUNs for Cisco UCS B200 M4 servers, complete the following steps:

1. Login in to Pure Storage GUI, select the + icon next to Volumes under the Storage tab.

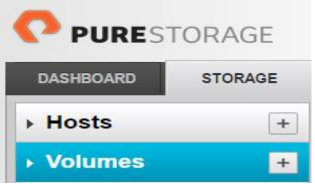

2. Create multiple volumes since there are a large number of ESXi hosts and you need to create separate boot volumes for each one.

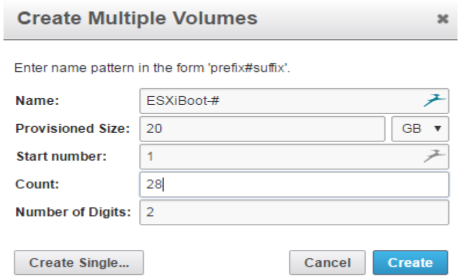

Create Multiple Volumes

Enter name pattern in the form 'prefix#suffix'.

Name: ESXiBoot-#

Provisioned Size: 20 GB

Start number: 1

Count: 28

Number of Digits: 2

3. To connect the LUN to an ESXi host, select a newly created volume under the Volumes tab and then select Connect Hosts from the options under the gear located on the right:

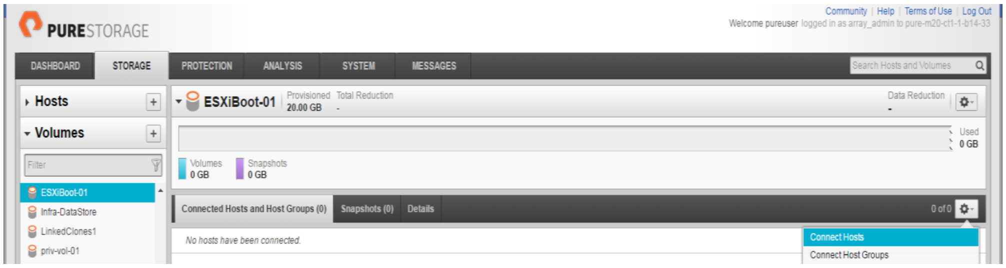

4. Click the single host you to connect to this volume and click Confirm.

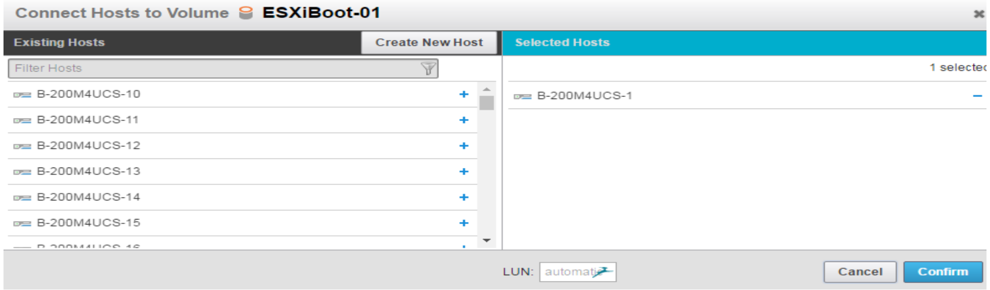

Connect Hosts to Volume ESXiBoot-01

Existing Hosts

- B-200M4UCS-10 +
- B-200M4UCS-11 +
- B-200M4UCS-12 +
- B-200M4UCS-13 +
- B-200M4UCS-14 +
- B-200M4UCS-15 +

Selected Hosts

1 selected

- B-200M4UCS-1 –

LUN: automatic

5. Repeat this procedure with all ESXi hosts. Remember, it is important to only connect a single host to each boot LUN.

Adding a Cisco UCS B200 M4 Server for ESXi Host Service

The next step is to create hosts within the Pure Storage GUI. Hosts can be grouped together into Host Groups to provide a higher level of abstraction for management. Volumes need to be connected to hosts and/or host groups in order for to two to communicate.

To add a Cisco UCS ESXi Host, complete the following steps:

1. To create a Host within the Pure Storage GUI, first navigate to the storage tab and you will see next to the Hosts menu a '+' sign:

Figure 63 Adding ESXi Host

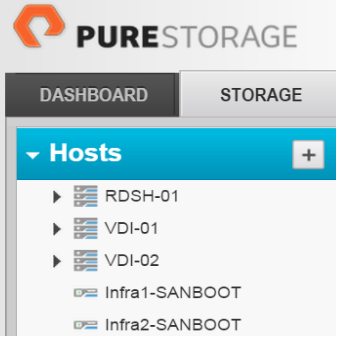

2. Click the '+' sign will spawn a window for you to create single, or multiple hosts. Since we are using multiple hosts in our design we will show the method for creating several hosts at once.

Figure 64 Create a Host Name or Multiple Hosts

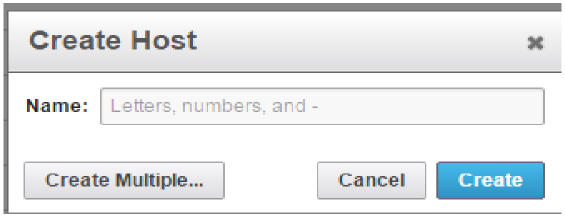

3. After clicking the Create Multiple, the following screen displays which enables batch creation of multiple hosts. In this example, you are creating two hosts.

Figure 65 Overview of one Cisco UCS B200 M4 Host being Created

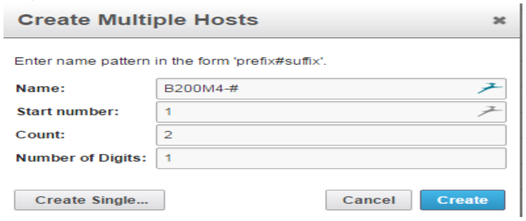
Create Multiple Hosts

Enter name pattern in the form 'prefix#suffix'.

Name: B200M4-#
Start number: 1
Count: 2
Number of Digits: 1

Figure 66 shows that two hosts have been created and the next step is to connect their WWNs. WWNs can be matched up with specific Cisco UCS servers from within Cisco UCS Manager.

Figure 66 Configuring the ESXi host WWNs from Cisco UCS for the Hosts

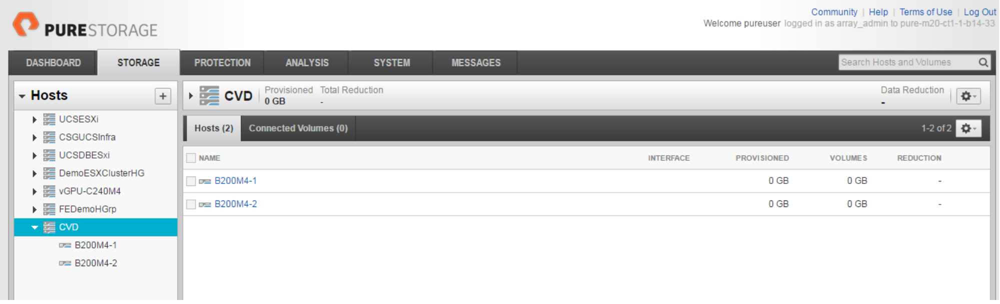

NAME	INTERFACE	PROVISIONED	VOLUMES	REDUCTION
B200M4-1		0 GB	0 GB	-
B200M4-2		0 GB	0 GB	-

4. To configure the FC WWNs, highlight a host, select the ‘Host Ports’ tab and then select the ‘Configure Fibre Channel WWNs’ button as shown in the below screenshot.

Figure 67 Configuring ESXi Host Ports with WWNs

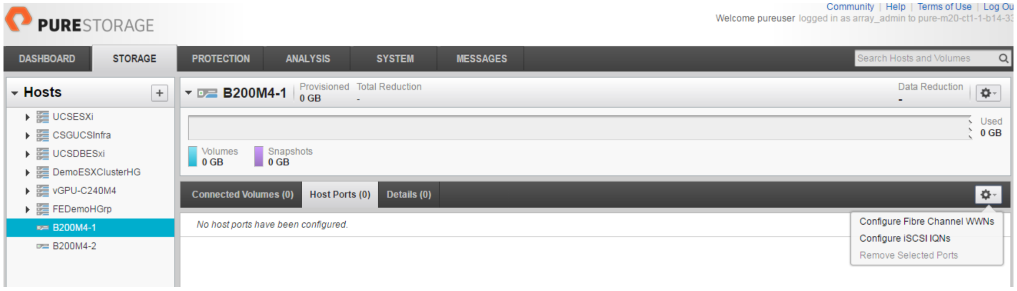

No host ports have been configured.

Available WWNs will appear in the left-hand column.

Figure 68 WWNs Configured for Hosts

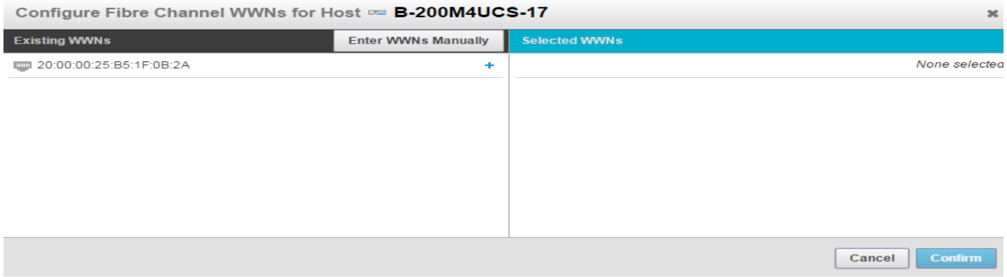

20:00:00:25:B5:1F:0B:2A

None selected

5. Click the WWNs and move them over to the right-hand column and click Confirm when completed. In this Cisco Validated Design, each Cisco UCS host will have two separate vHBAs to provide both performance and resiliency.

Figure 69 shows a properly setup UCS host within the Pure GUI. Repeat the above procedure for host creation and configuring the Fibre Channel WWNs for all Cisco UCS servers.

Figure 69 Overview of a Configured ESXi Host

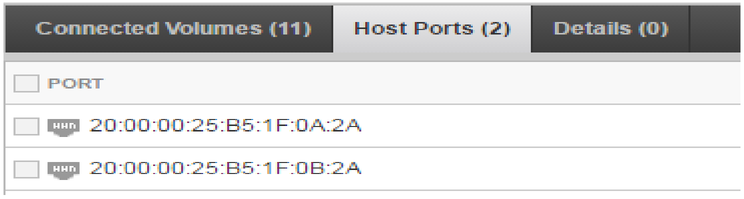

□ PORT
□ 20:00:00:25:B5:1F:0A:2A
□ 20:00:00:25:B5:1F:0B:2A

6. When the hosts have been created and the WWNs have been assigned to them, then click the Hosts button and select Create Host Group.

Figure 70 Creating a Host Group

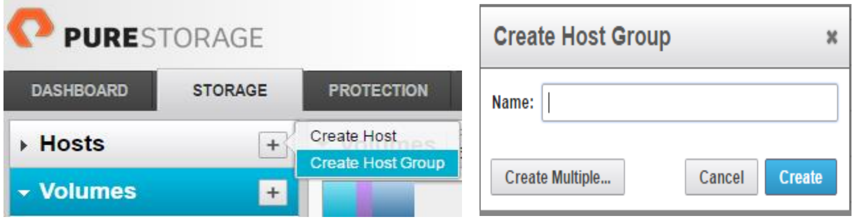

7. The next step is to add the two example hosts we previously created to this newly created Host Group. To accomplish this, highlight the Host Group, click the gear to the right and select Add Hosts.

Figure 71 Add Host to the Host Group Created

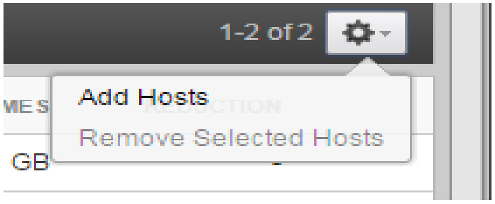

8. The following screenshots display; click the hosts you wish to add to the Host Group on the left column in order to move them to the right. When all hosts have been included, click Confirm.

Figure 72 ESXi Hosts to be added to the Host Group

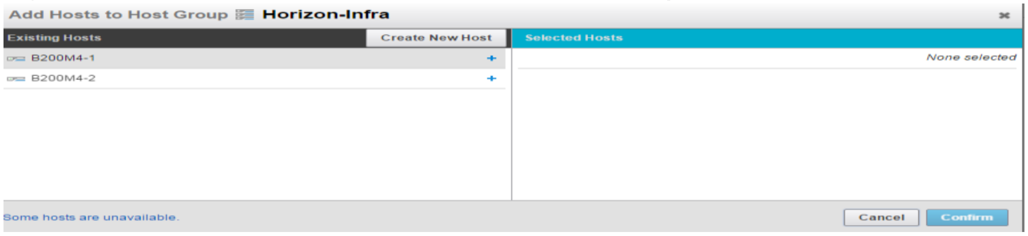
Add Hosts to Host Group Horizon-Infra

Existing Hosts	Selected Hosts
B200M4-1 +	None selected
B200M4-2 +	

Some hosts are unavailable.

Figure 73   ESXi Host added to the Host Group

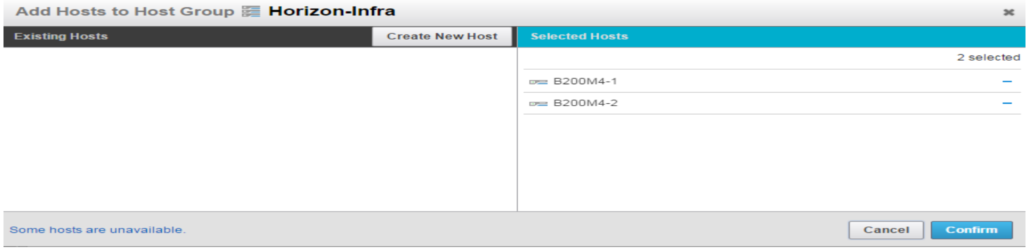
Add Hosts to Host Group Horizon-Infra

Existing Hosts	Selected Hosts
	2 selected
	B200M4-1
	B200M4-2

Some hosts are unavailable.

The following screenshot shows the populated Host Group with two hosts as members.

Figure 74   ESXi Hosts have been added to the Host Group

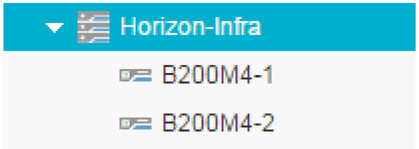
Horizon-Infra
- B200M4-1
- B200M4-2

Repeat this procedure for all Host Groups that are required. For this Cisco Validated Design we created three separate Host Groups for the following components:

- Citrix XenDesktop Infrastructure (xx Hosts)
- Citrix XenApp Hosts (xx Hosts)
- Citrix XenDesktop Persistent and Pooled Hosts (xx Hosts)

With the Hosts and Host Group setup completed, the array is now ready to serve Citrix XenDesktop and XenApp desktops and the associated infrastructure.

Volume and Data stores creation on Pure Storage FlashArray//m50

Creating and even resizing a datastore on the Pure Storage array is incredibly simple and can be accomplished in just a few clicks.

To create and resized a datastore, complete the following steps:

1. Click the Storage tab from within the Pure GUI:

Figure 75   Pure Storage GUI

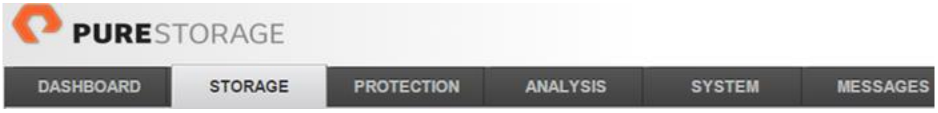

2. On the collapsible Volumes menu on the left, select the ‘+’ button:

Figure 76   Select Volume

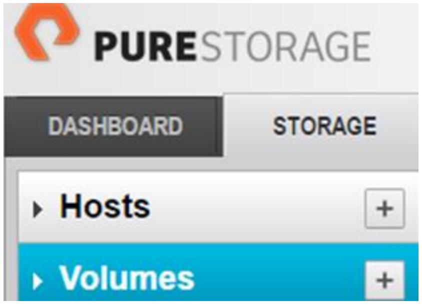

3. Name and size the Volume to complete its creation:

Figure 77   Provide a Name to the Volume and Size

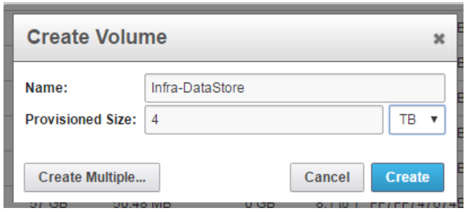

Name: Infra-DataStore

Provisioned Size: 4 TB

4. Select the newly created Volume under the Volumes collapsible menu on the left-hand side of the GUI and click on the gear icon to connect it to host(s) or host group(s):

Figure 78   Select the Volume Created to be Associated with the Host

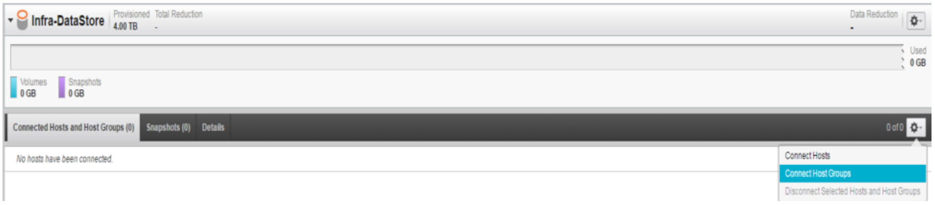

Provisioned: 4.00 TB

Total Reduction: -

Used: 0 GB

Volumes: 0 GB

Snapshots: 0 GB

No hosts have been connected.

5. Click the appropriate Host Groups that require access to the Volume so that they are under the 'Select Host Groups' list and then click Confirm when they have been added:

Figure 79   Add Host Group or Group Hosts Configured for the Volume

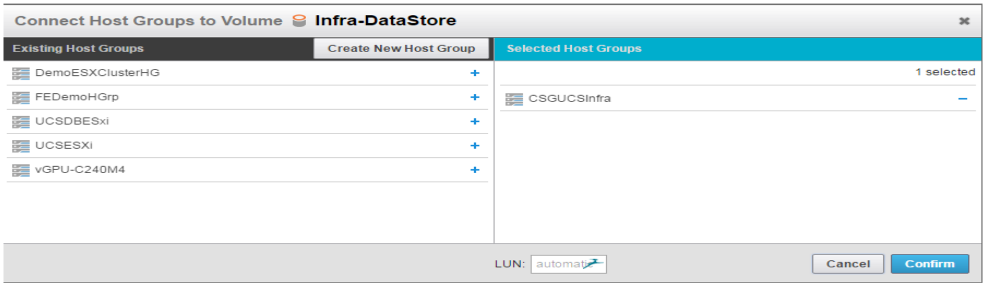

1 selected

CSGUCSInfra

LUN: automatic

6. Storage setup has been completed and the datastore now needs to be added from within vSphere. From within the vSphere client, select a host from the Host Group that was connected to the volume in the previous step, go to Related Objects and then Datastore and click the icon to add the new datastore.

Figure 80 vSphere GUI to Create the Datastore

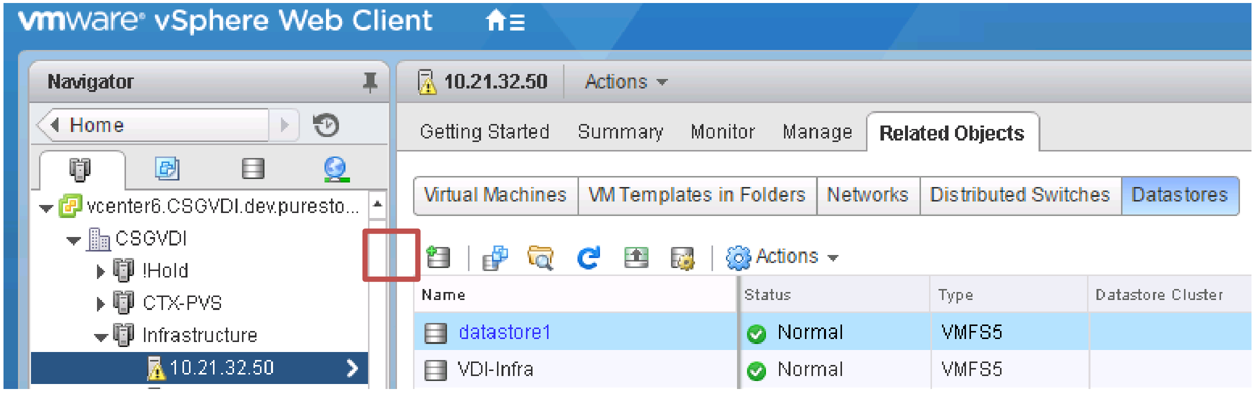

Name	Status	Type	Datastore Cluster
datastore1	Normal	VMFS5	
VDI-Infra	Normal	VMFS5	

7. Select the VMFS option and click Next.

Figure 81 Select VMFS Datastore

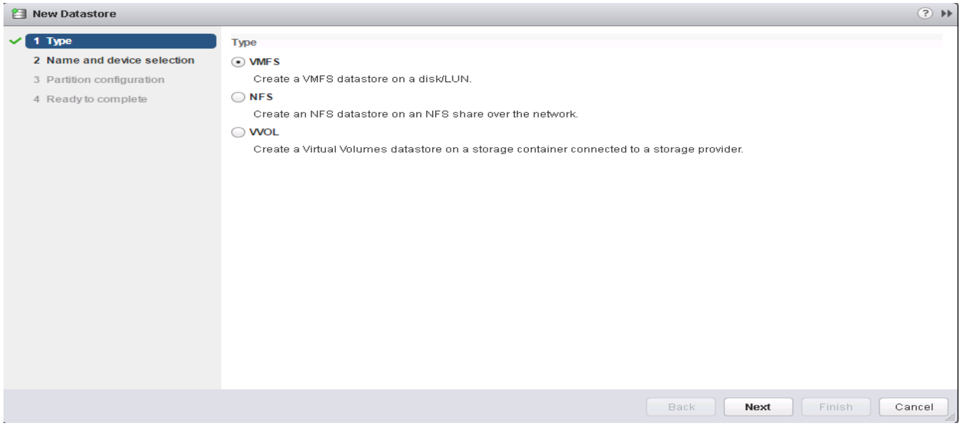

Type

● VMFS
Create a VMFS datastore on a disk/LUN.

○ NFS
Create an NFS datastore on an NFS share over the network.

○ VVOL
Create a Virtual Volumes datastore on a storage container connected to a storage provider.

8. Name the datastore and select the appropriate volume from the list shown. It might be necessary to rescan the vHBA on the host in order to see a newly created datastore.

Figure 82 Name and Device Selection for the Datastore

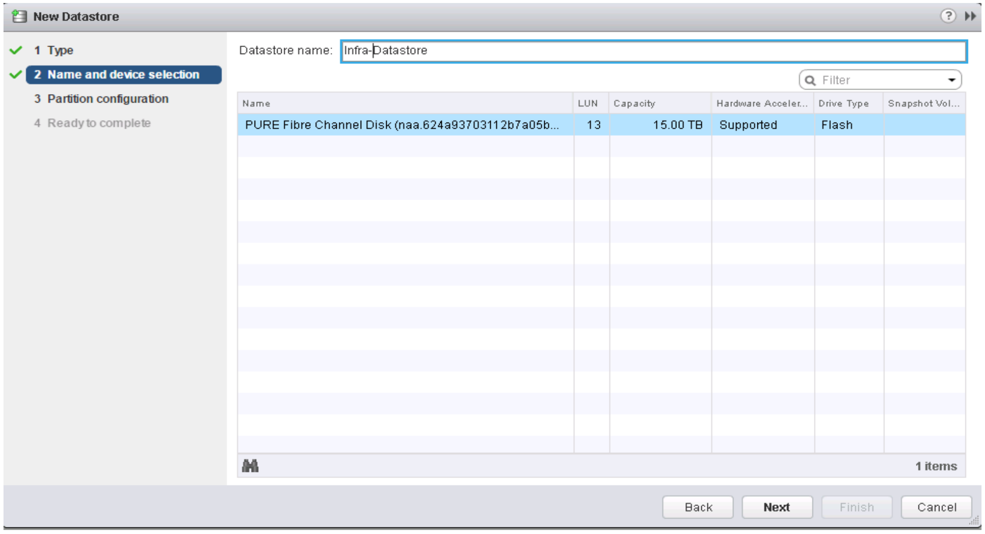

Datastore name: Infra-Datastore

Name	LUN	Capacity	Hardware Acceler...	Drive Type	Snapshot Vol...
PURE Fibre Channel Disk (naa.624a93703112b7a05b...	13	15.00 TB	Supported	Flash	

1 items

9. Leave the partition configuration options at default values and click Next.

Figure 83   Configuration Parameters for Datastore

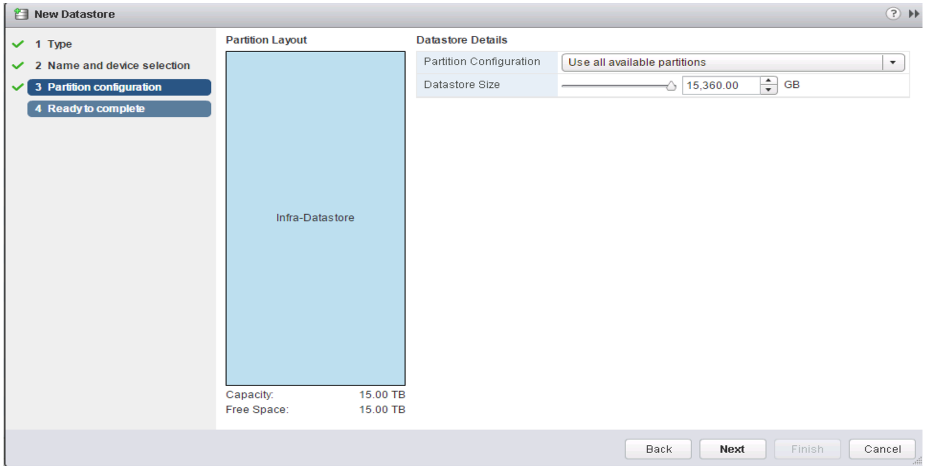

Datastore Details

Partition Configuration	Use all available partitions
Datastore Size	15,360.00 GB

Capacity: 15.00 TB
Free Space: 15.00 TB

10. Click Finish to complete the datastore setup.

Figure 84   Datastore is created

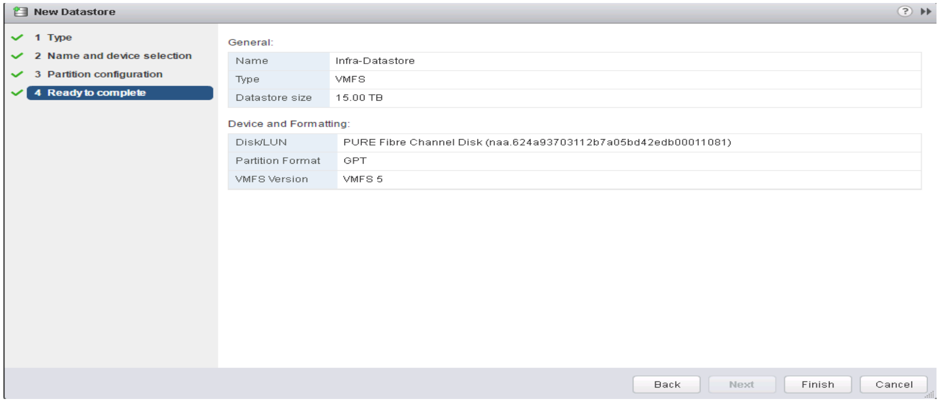

General:

Name	Infra-Datastore
Type	VMFS
Datastore size	15.00 TB

Device and Formatting:

Disk/LUN	PURE Fibre Channel Disk (naa.624a93703112b7a05bd42edb00011081)
Partition Format	GPT
VMFS Version	VMFS 5

Repeat this process for each of the datastores listed in the Datastore Table 14.

Datastore Snapshot on Pure Storage Array

A critical factor to the success of any VDI project is having a robust backup policy in place. Having the ability to revert an entire datastore to an earlier version if a mistake or catastrophic event occurs can save hundreds of hours and protect proprietary data from loss or corruption. Pure Storage includes snapshots and replication for free as a feature of the Purity Operating Environment and worth noting is that all future software features will be included at no additional charge.

We used the backup schedule shown in Table 14 for the datastores defined in the previous section. Important to note is that since the linked-clone desktops were set as being non-persistent, we elected not to assign any snapshot schedule to either datastore since no persistent data was being saved between user sessions. For environments with persistent desktops, it is a recommended practice to use a separate datastore with a defined snapshot schedule to protect against the loss of any user data.

Table 14 Example of Datastores Created on Pure Storage for Snapshot

VDI Component	# of Datastores	Datastore Size	Snapshot Schedule	Snapshot Rationale
Infrastructure Datastore	1	4 TB	Create snapshot every day, then retrain 1 snapshot per day for 5 days	Core infrastructure should be backed up regularly
User HomeDrives Datastore	1	5TB	Create snapshot every day, then retrain 1 snapshot per day for 5 days	Datastore level backup for user home drives.
User Profiles Datastore	1	5TB	Create snapshot every day, then retrain 1 snapshot per day for 5 days	Datastore level backup for user profiles.
PVS vDisk Datastore	1	1TB	Create snapshot every day, then retrain 1 snapshot per day for 3 days	Protection against accidental data loss or patching/update issue with vDisk
VDI Component	# of Data stores	Datastore Size	Snapshot Schedule	Snapshot Rationale
XenDesktop Infrastructure Datastore	1	10 TB	Create snapshot every day, then retrain 1 snapshot per day for 5 days	Core infrastructure should be backed up regularly
XenApp Datastore	1	10 TB	Create snapshot every 2 days, then retain 1 snapshot per day for 2 more days	Given the level of customization required for each RDS 2012R2 server, occasional backups needed
Linked-Clone Datastore	2-4*	5TB	No snapshots	Non-persistent VMs should be restored from parent VM on Infrastructure Datastore if LUN accidently destroyed
ESXi Boot LUNS	1 per UCS host	10GB	Create a snapshot on source every 5 days then retain 1 per day for 1 more day	Not much change on ESXi boot LUNs, hence occasional backups only

Similar to datastore creation, setting up a snapshot schedule is also accomplished easily and intuitively from within the Pure Storage GUI.

To setup a snapshot schedule, complete the following steps:

1. Click the Protection tab.

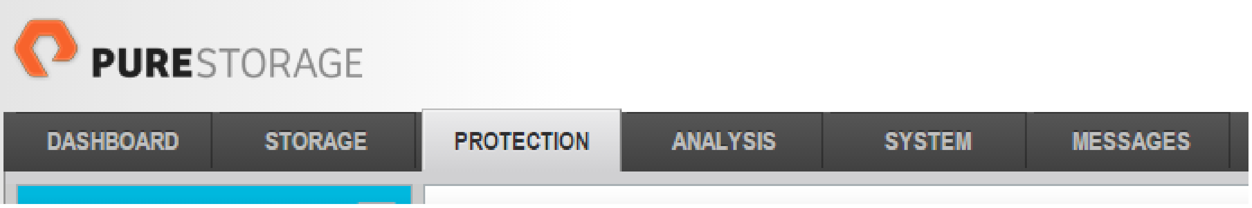

2. Click the ‘+’ sign on the Source Groups menu to the left:

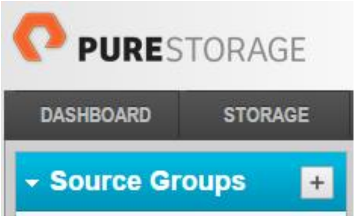

3. Give the Protection Group a recognizable name.

Figure 85 Pure GUI to Create a Snapshot Protection for Host or Host Group

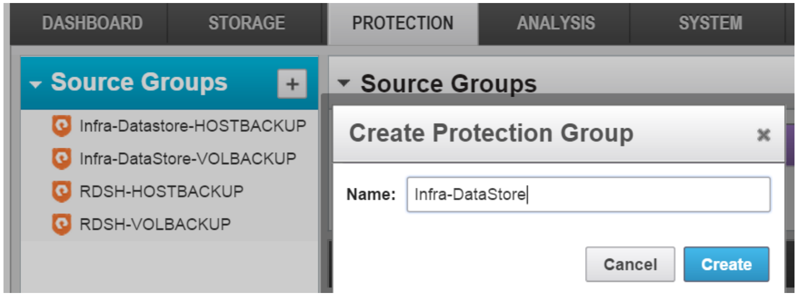

Name: Infra-DataStore

4. Using the Snapshot schedule for the Infrastructure datastore, assign the following values to create the snapshot schedule and retention policy:

Figure 86 Snapshot Schedule

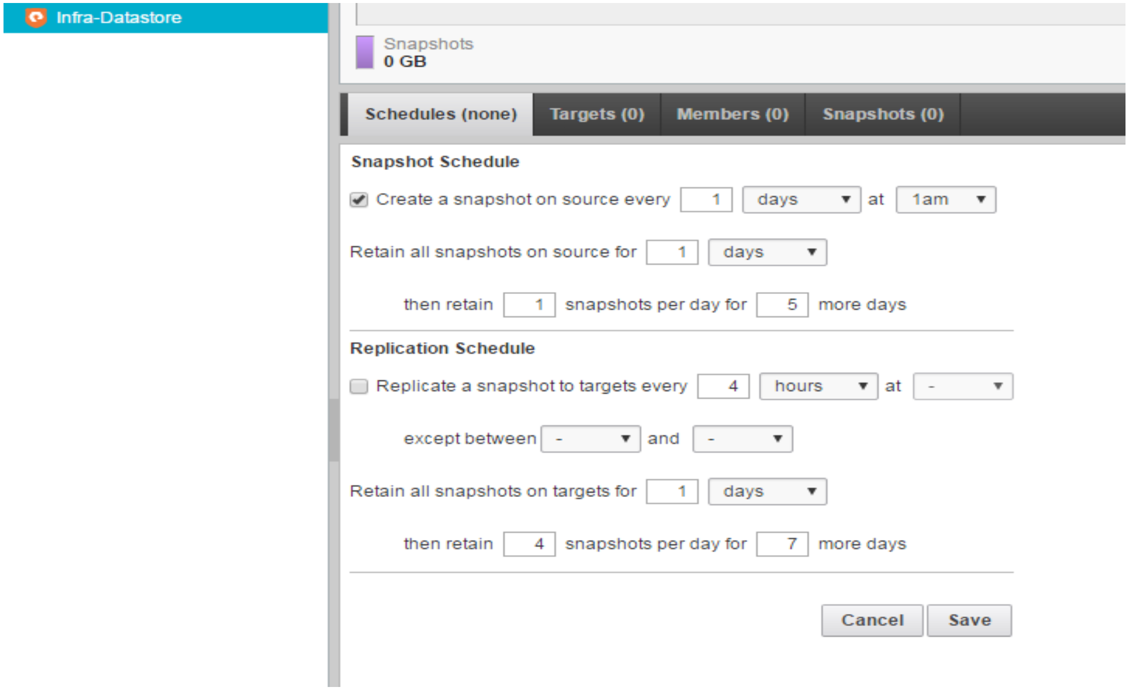

Snapshot Schedule

☑ Create a snapshot on source every 1 days at 1am

Retain all snapshots on source for 1 days

then retain 1 snapshots per day for 5 more days

Replication Schedule

☐ Replicate a snapshot to targets every 4 hours at -

except between - and -

Retain all snapshots on targets for 1 days

then retain 4 snapshots per day for 7 more days

5. When the Snapshot and retention policy has been defined, add the Datastore to the policy by clicking the ‘Members’ tab of the Source Group:

Figure 87   Add Datastore or Volume to the Snapshot Policy

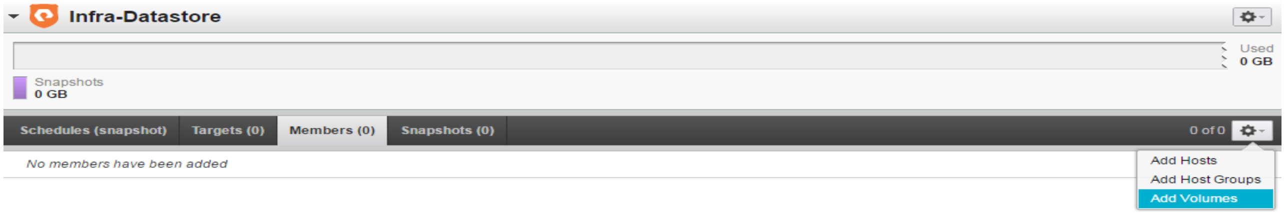

No members have been added

6. Click the Volume(s) you want to include in the Snapshot policy to move them to the right. Click Confirm when completed.

Figure 88   Adding Volumes to the Snapshot Policy Configured

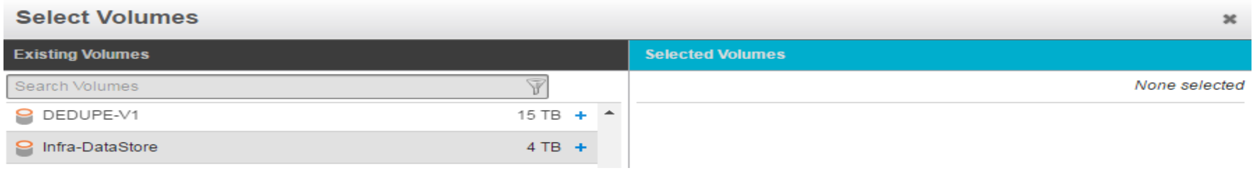

DEDUPE-V1	15 TB
Infra-DataStore	4 TB

None selected

7. After clicking Confirm, the Datastore should be listed as a Snapshot policy member:

Figure 89   Infra-Datastore Volume Applied to the Newly Created Snapshot Retention Policy for Backup

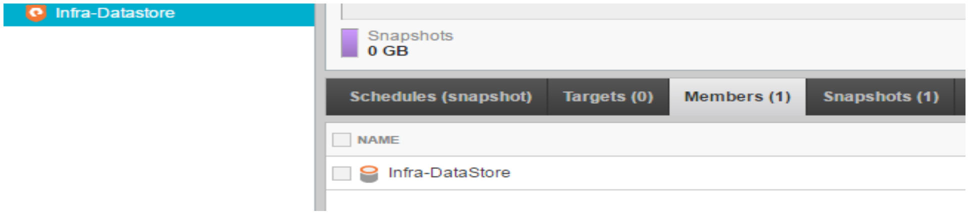

NAME
Infra-DataStore

8. Repeat these steps for the RDSH and ESXi Boot LUNs Data stores. It is recommended to include all ESXi boot LUNs in a single Snapshot policy for simplicity.

ESXi Best Practice Configuration

Due to the simplicity of both the Pure Storage FlashArray and the server, configuring VMware ESXi best-practice is also simple. ESXi uses its Native Multipathing Plugin architecture to manage I/O multipathing to underlying SAN storage volumes. Pure Storage FlashArray volumes are claimed by default by the Storage Array Type Plugin (SATP) for ALUA devices and inherit the Most Recently Used (MRU) Path Selection Policy (PSP). The FlashArray is not an ALUA array—it has an active/active frontend controller architecture. This VMware default device claiming behavior limits I/O to a single path and, if unchanged, is notably detrimental to I/O performance as only leveraging a single path/port eliminates the advantages of the active/active nature of the FlashArray.

Therefore, all ESXi servers were configured to change the default PSP for Pure Storage FlashArray devices from MRU to Round Robin (with an advanced configuration to alternate to logical paths after every I/O). The following command was run on each ESXi server prior to the presentation of FlashArray devices:

```
esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "PURE" -M
"FlashArray" - P

"VMW_PSP_RR" -O "iops=1"
```

A PowerShell script has been created to check for, and apply these best practices automatically to the entire vCenter cluster. Please see the attachment in Appendix B for further information.

Configure User Profile Manager Share on Pure Storage FlashArray//m50

A VDI user profile share was built using Windows Server 2012 R2 to closely mirror a production environment where user profiles are stored in a network location. The user profiles are created for the Login VSI users who sign in during the automated Knowledge Worker benchmark test runs. The server was hosted on the infrastructure datastore where our snapshot policies were applied for backing up the system.

The profiles were stored on a separate D: partition on the Windows 2012 server.

VDI- User Profiles

Figure 90   Profile Share for the Users Configured on the Pure Storage

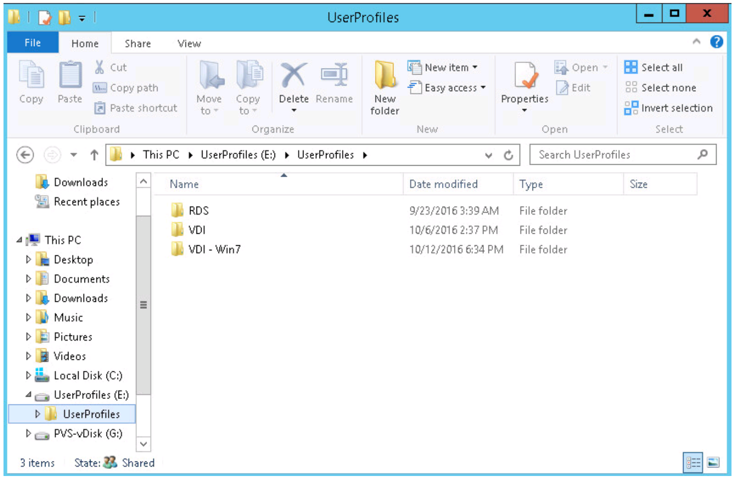

Name	Date modified	Type	Size
RDS	9/23/2016 3:39 AM	File folder	
VDI	10/6/2016 2:37 PM	File folder	
VDI - Win7	10/12/2016 6:34 PM	File folder	

Figure 91   XenApp Session User Profiles

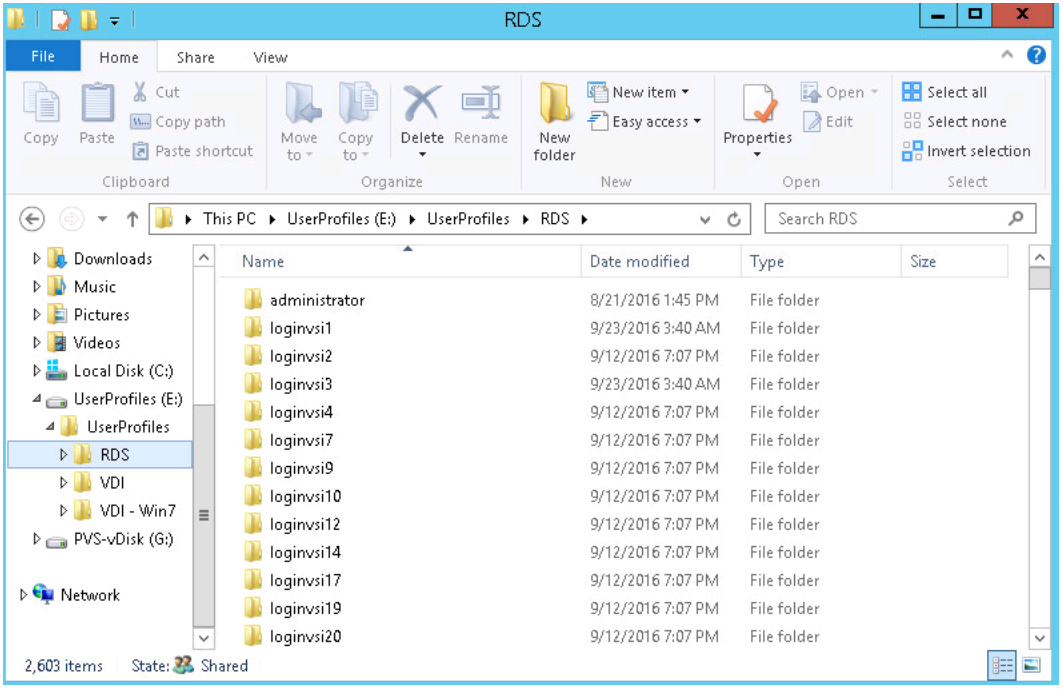

Name	Date modified	Type	Size
administrator	8/21/2016 1:45 PM	File folder	
loginvsi1	9/23/2016 3:40 AM	File folder	
loginvsi2	9/12/2016 7:07 PM	File folder	
loginvsi3	9/23/2016 3:40 AM	File folder	
loginvsi4	9/12/2016 7:07 PM	File folder	
loginvsi7	9/12/2016 7:07 PM	File folder	
loginvsi9	9/12/2016 7:07 PM	File folder	
loginvsi10	9/12/2016 7:07 PM	File folder	
loginvsi12	9/12/2016 7:07 PM	File folder	
loginvsi14	9/12/2016 7:07 PM	File folder	
loginvsi17	9/12/2016 7:07 PM	File folder	
loginvsi19	9/12/2016 7:07 PM	File folder	
loginvsi20	9/12/2016 7:07 PM	File folder	

Figure 92   XenDesktop Windows 10 User Profiles

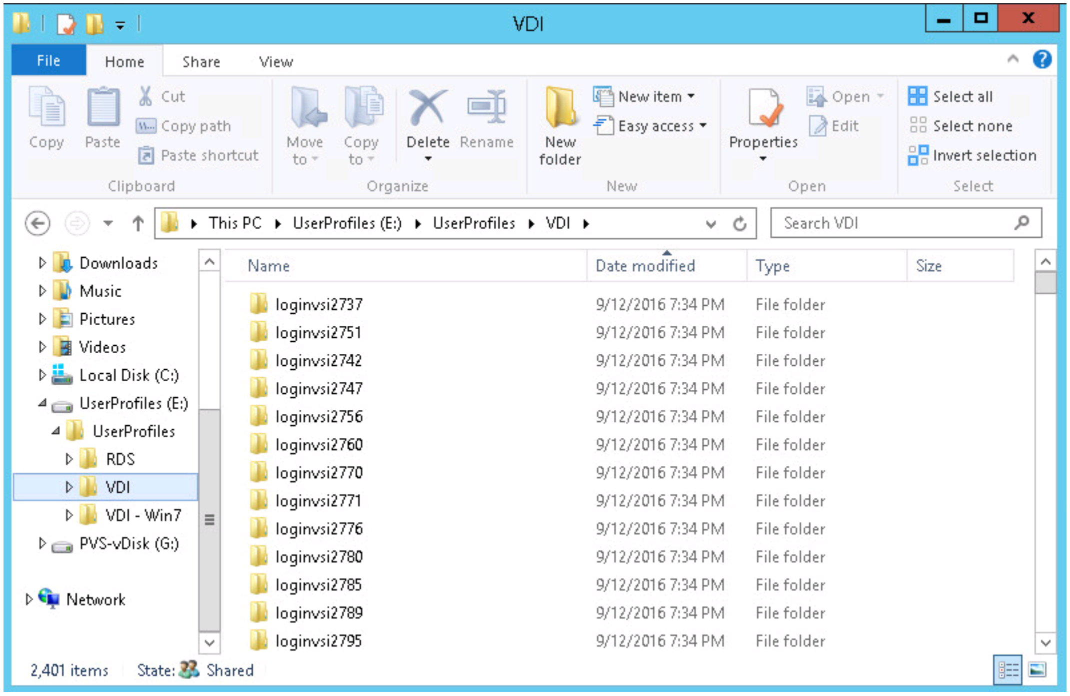

Name	Date modified	Type	Size
loginvsi2737	9/12/2016 7:34 PM	File folder	
loginvsi2751	9/12/2016 7:34 PM	File folder	
loginvsi2742	9/12/2016 7:34 PM	File folder	
loginvsi2747	9/12/2016 7:34 PM	File folder	
loginvsi2756	9/12/2016 7:34 PM	File folder	
loginvsi2760	9/12/2016 7:34 PM	File folder	
loginvsi2770	9/12/2016 7:34 PM	File folder	
loginvsi2771	9/12/2016 7:34 PM	File folder	
loginvsi2776	9/12/2016 7:34 PM	File folder	
loginvsi2780	9/12/2016 7:34 PM	File folder	
loginvsi2785	9/12/2016 7:34 PM	File folder	
loginvsi2789	9/12/2016 7:34 PM	File folder	
loginvsi2795	9/12/2016 7:34 PM	File folder	

Configure MDS 9100 Series

To configure the MDS 9100 series, complete the following steps:

1. In this solution we utilized the Cisco MDS 9148 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

2. When the MDS switch is racked and can be logged into it can now be configured to communicate with the Cisco UCS Fabric Interconnects.

3. In this study, we used two separate fabrics each with their own unique VSAN. Fabric A is configured for VSAN20 while Fabric B for VSAN30. In our initial Cisco UCS configuration you will see where we configured fiber cables on ports 13 and 14 and configured a FC port-channel. FI-A's FC port channel is configured for VSAN20 and FI-B's FC port-channel for VSAN30.

Figure 93 VSAN 20 configured for Fabric A

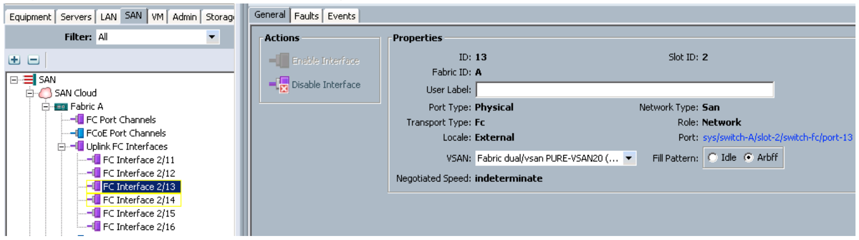

ID: 13
Slot ID: 2
Fabric ID: A
User Label:
Port Type: Physical
Network Type: San
Transport Type: Fc
Role: Network
Locale: External
Port: sys/switch-A/slot-2/switch-fc/port-13
VSAN: Fabric dual/vsan PURE-VSAN20 (...
Fill Pattern: ○ Idle ● Arbff
Negotiated Speed: indeterminate

Figure 94 VSAN 30 configured for Fabric B

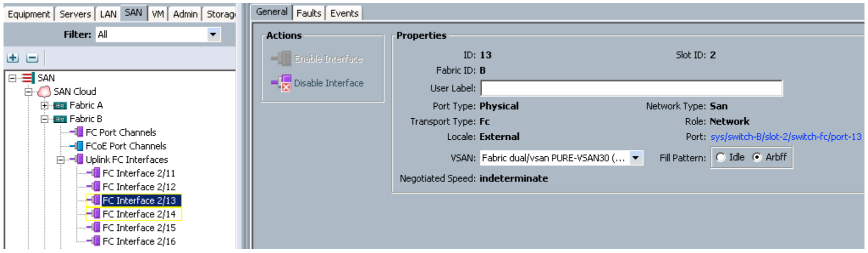

ID: 13
Slot ID: 2
Fabric ID: B
User Label:
Port Type: Physical
Network Type: San
Transport Type: Fc
Role: Network
Locale: External
Port: sys/switch-B/slot-2/switch-fc/port-13
VSAN: Fabric dual/vsan PURE-VSAN30 (...
Fill Pattern: ○ Idle ● Arbff
Negotiated Speed: indeterminate

Physically, the Fabric Interconnects extended ports 13 and 14 run to the MDS switch ports 1 and 2.

Figure 95 MDS Switch VSAN Configuration Connectivity

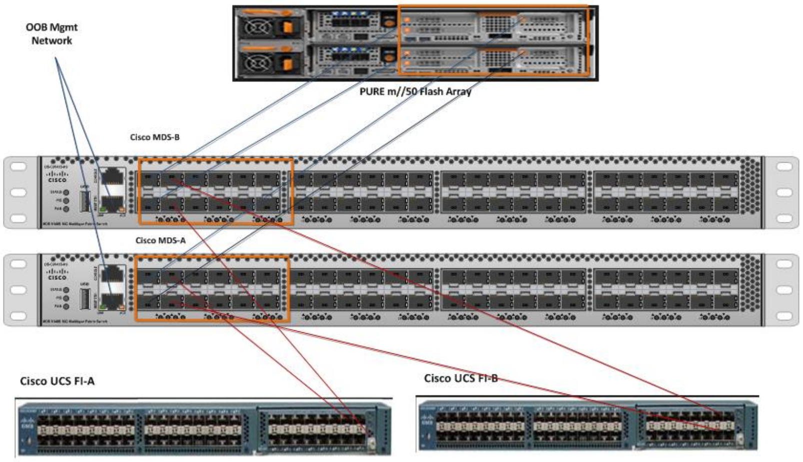

We used a total of 8 16Gb FC links, four from each FlashArray//m50 controller, to the MDS 9148 SAN Switch for high availability and maximum throughput.

After the ports and port channels are configured, the next steps are to configure the zones and Active Zoneset Database in the MDS switches. The commands listed below show how to add in a single host on both MDS A and B. You will need to configure all hosts that will access the Pure Array in these commands. Then entire MDS switch configuration is included in this document in Appendix A.

MDS-A

```
Configure Terminal
Zoneset name VDI-Infra-A vsan 20
Zone name {ESXi hostname-fc0} vsan 20
Member pwwn {ESXi Host pwwn for fc0}
Member pwwn {PURE pwwn Controller A, Port 1}
Member pwwn {PURE pwwn Controller B, Port 1}
Zone commit vsan 20
Zoneset name VDI-Infra-A vsan 20
Member {ESXi hostname-fc0}
Exit
Zoneset activate name VDI-Infra-A vsan 20
Zone commit vsan 20
Exit
Copy running-config startup-config
```

MDS-B

```
Configure Terminal
Zoneset name VDI-Infra-B vsan 30
Zone name {ESXi hostname-fc1} vsan 30
Member pwwn {ESXi Host pwwn for fc1}
Member pwwn {PURE pwwn Controller A, Port 2}
Member pwwn {PURE pwwn Controller B, Port 2}
Zone commit vsan 30
Zoneset name VDI-Infra-B vsan 30
Member {ESXi hostname-fc1}
Exit
Zoneset activate name VDI-Infra-B vsan 30
Zone commit vsan 30
Exit
Copy running-config startup-config
```

Install and Configure VMware ESXi 6.0 U2

Installing and Configuring VMware ESXi 6.0

This section provides detailed instructions for installing VMware ESXi 6 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 Update2

To download the Cisco Custom Image for ESXi 6 Update 2, 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 on the following link: https://my.vmware.com/web/vmware/details?downloadGroup=OEM-ESXI60U2-CISCO&productId=491
4. Click Download Now.
5. Save it to your destination folder.

This ESXi 6.0 Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.3.0.10; fNIC: 1.6.0.28

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 clicking 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 PURE LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

5. Select the appropriate keyboard layout and press Enter.

6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.
8. 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.

9. 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, simply click Yes to unmap the image.
10. 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 add a management network for the VMware hosts, complete the following steps on each ESXi 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.

Because 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.

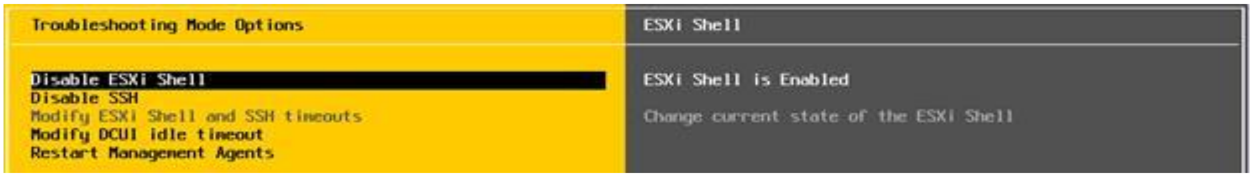

```
Troubleshooting Mode Options                ESXi Shell

Disable ESXi Shell                          ESXi Shell is Enabled
Disable SSH                                 Change current state of the ESXi Shell
Modify ESXi Shell and SSH timeouts
Modify DCUI idle timeout
Restart Management Agents
```

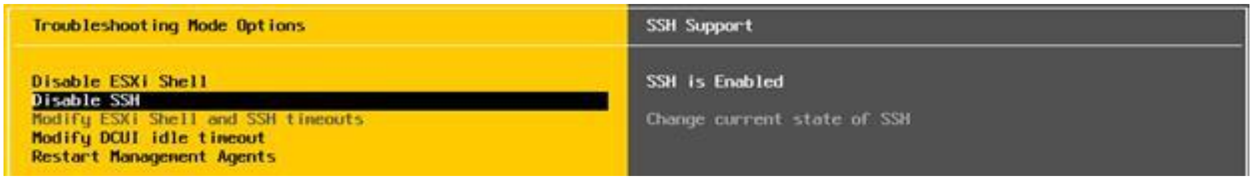

```
Troubleshooting Mode Options                SSH Support

Disable ESXi Shell                          SSH is Enabled
Disable SSH                                 Change current state of SSH
Modify ESXi Shell and SSH timeouts
Modify DCUI idle timeout
Restart Management Agents
```

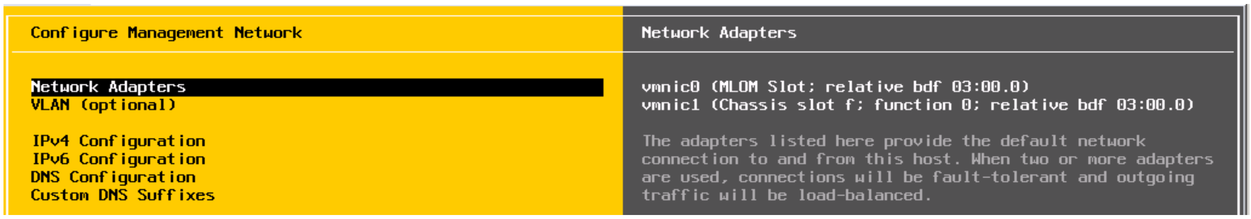

```
Configure Management Network

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 03:00.0)

The adapters listed here provide the default network
connection to and from this host. When two or more adapters
are used, connections will be fault-tolerant and outgoing
traffic will be load-balanced.
```

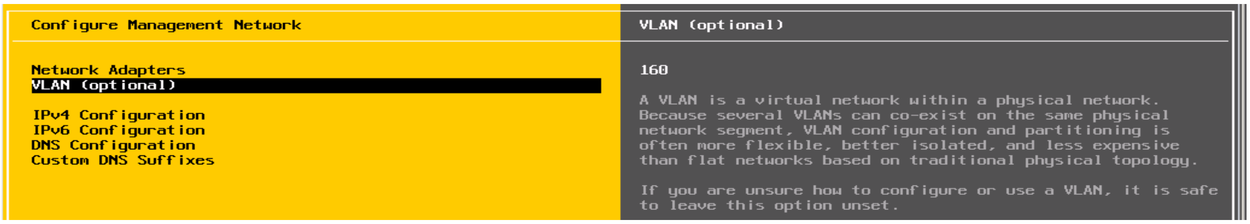

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
VLAN (optional)

160

A VLAN is a virtual network within a physical network.
Because several VLANs can co-exist on the same physical
network segment, VLAN configuration and partitioning is
often more flexible, better isolated, and less expensive
than flat networks based on traditional physical topology.

If you are unsure how to configure or use a VLAN, it is safe
to leave this option unset.
```

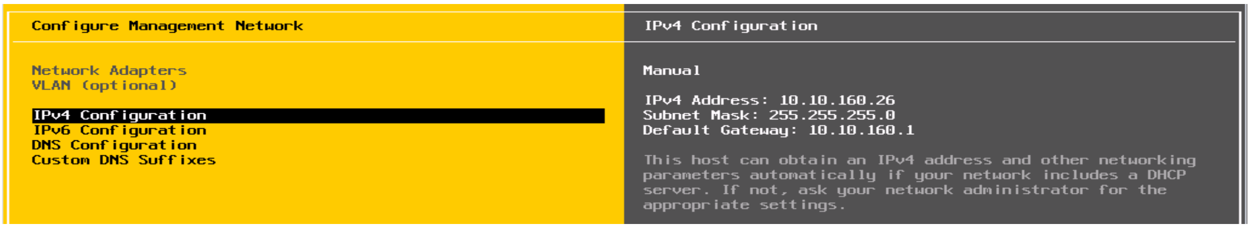

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
IPv4 Configuration

Manual

IPv4 Address: 10.10.160.26
Subnet Mask: 255.255.255.0
Default Gateway: 10.10.160.1

This host can obtain an IPv4 address and other networking
parameters automatically if your network includes a DHCP
server. If not, ask your network administrator for the
appropriate settings.
```

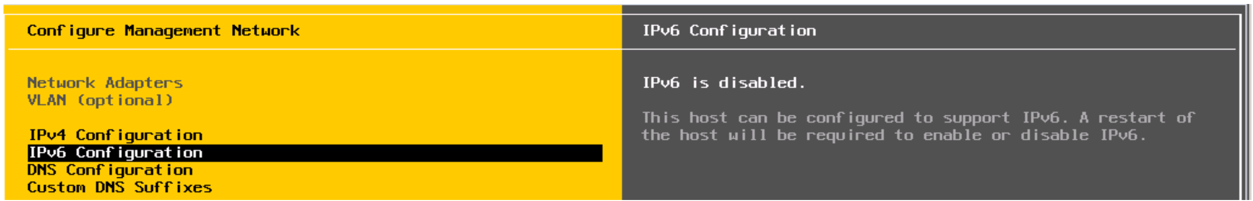

```
Configure Management Network

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 of
the host will be required to enable or disable IPv6.
```

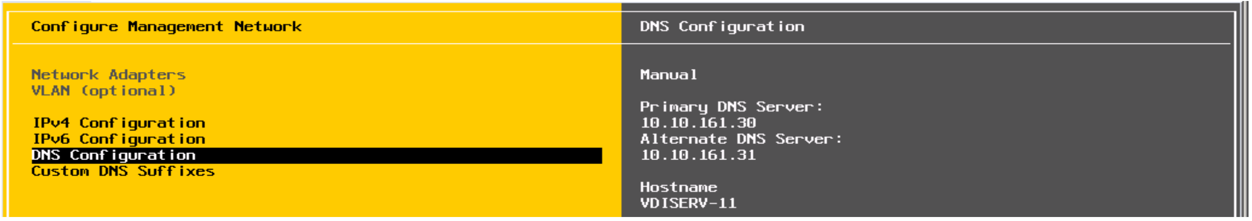

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
DNS Configuration

Manual

Primary DNS Server:
10.10.161.30
Alternate DNS Server:
10.10.161.31

Hostname
VDISERV-11
```

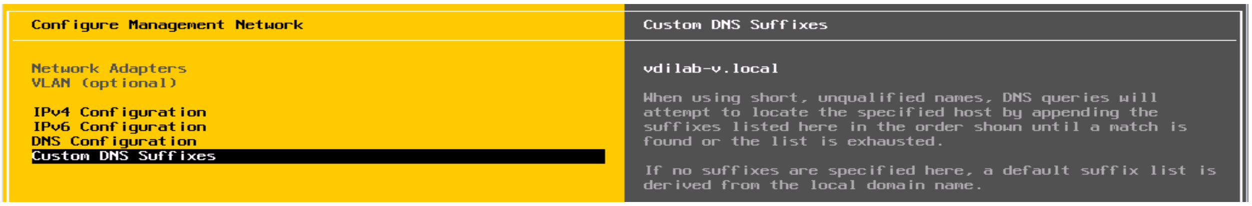

```
Configure Management Network

Network Adapters
VLAN (optional)

IPv4 Configuration
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
Custom DNS Suffixes

vdilab-v.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 is
found or the list is exhausted.

If no suffixes are specified here, a default suffix list is
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, complete the following steps:

1. Click the following link VMware vSphere CLI 6.0
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 6.0 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 Updated Cisco VIC eNIC Drivers

To download the Cisco virtual interface card (VIC) eNIC and fNIC drivers, complete the following steps:

The eNIC version is 2.3.0.10 and the fNIC version is 1.6.0.28 were used in this configuration

1. Open a Web browser on the management workstation and navigate to:
2. https://my.vmware.com/web/vmware/details?downloadGroup=ESXI60U2&productId=491&rPId=12932#drivers_tools
3. Download the Cisco eNIC and fNIC driver bundle.
4. Open the eNIC 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.

If the link above has changed, go to www.cisco.com for the latest ISO image of Cisco UCS-related drivers. This ISO will either have the drivers included or may have an HTML file with the location of the latest network drivers.

Load Updated Cisco VIC eNIC and fNIC Drivers

To install VMware VIC Drivers on the ESXi host servers, complete the following steps:

1. From each 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
   - fnic_driver_1.6.0.28-4179603.zip or
   - ESXi6.0_enic-2.3.0.10-4303638.zip
6. Click Open on each and click Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.
8. Make sure the files have been uploaded to both 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/ESXi6.0_enic-2.3.0.10-offline_bundle-
4303638.zip

esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/
fnic_driver_1.6.0.28-offline_bundle-4179603.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 enic driver version installed by entering `vmkload_mod -s enic` and `vmkload_mod -s fnic` at the command prompt.

Install and Configure VMware vCenter Appliance

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

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. Open the vSphere ISO via Windows Explorer and double-click the vcsa-setup.htm file.

Figure 96 Install VCSA Appliance from Installer

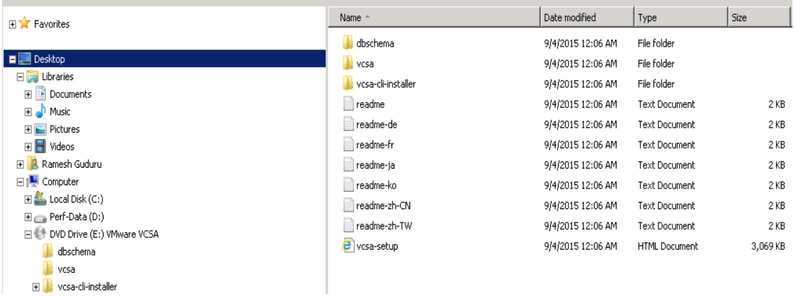

Favorites

Desktop
Libraries
Documents
Music
Pictures
Videos
Ramesh Guduru
Computer
Local Disk (C:)
Perf-Data (D:)
DVD Drive (E:) VMware VCSA
dbschema
vcsa
vcsa-cli-installer

Name	Date modified	Type	Size
dbschema	9/4/2015 12:06 AM	File folder	
vcsa	9/4/2015 12:06 AM	File folder	
vcsa-cli-installer	9/4/2015 12:06 AM	File folder	
readme	9/4/2015 12:06 AM	Text Document	2 KB
readme-de	9/4/2015 12:06 AM	Text Document	2 KB
readme-fr	9/4/2015 12:06 AM	Text Document	2 KB
readme-ja	9/4/2015 12:06 AM	Text Document	2 KB
readme-ko	9/4/2015 12:06 AM	Text Document	2 KB
readme-zh-CN	9/4/2015 12:06 AM	Text Document	2 KB
readme-zh-TW	9/4/2015 12:06 AM	Text Document	2 KB
vcsa-setup	9/4/2015 12:06 AM	HTML Document	3,069 KB

A browser will open with an option to Install.

3. Click Install.

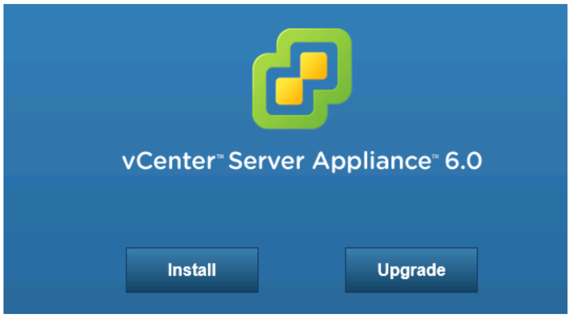

4. Follow the onscreen prompts. Accept EULA.

Figure 97 Accept EULA Agreement

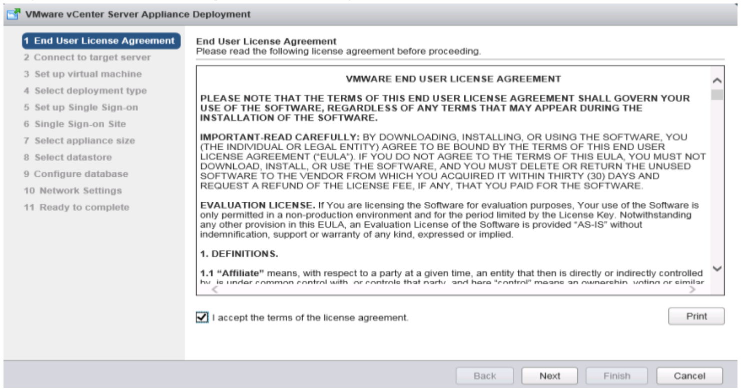

VMware vCenter Server Appliance Deployment

1 End User License Agreement
2 Connect to target server
3 Set up virtual machine
4 Select deployment type
5 Set up Single Sign-on
6 Single Sign-on Site
7 Select appliance size
8 Select datastore
9 Configure database
10 Network Settings
11 Ready to complete

End User License Agreement
Please read the following license agreement before proceeding.

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.

1.1 "Affiliate" means, with respect to a party at a given time, an entity that then is directly or indirectly controlled by, is under common control with, or controls that party, and here "control" means an ownership, voting or similar

☑ I accept the terms of the license agreement.

Print

Back Next Finish Cancel

5. Enter the IP of the ESXi host the vCenter Appliance will reside. Click Next.

Figure 98 Provide Host IP or FQDN and User Name, Password Credentials of the Host to Connect

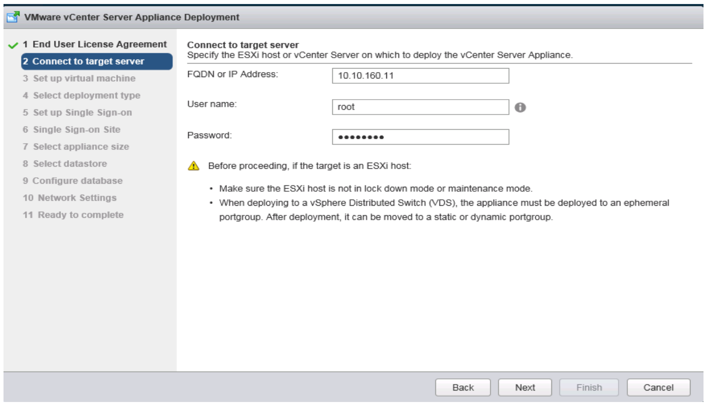

Connect to target server
Specify the ESXi host or vCenter Server on which to deploy the vCenter Server Appliance.

FQDN or IP Address: 10.10.160.11

User name: root

Password: ••••••••

Before proceeding, if the target is an ESXi host:

- Make sure the ESXi host is not in lock down mode or maintenance mode.
- When deploying to a vSphere Distributed Switch (VDS), the appliance must be deployed to an ephemeral portgroup. After deployment, it can be moved to a static or dynamic portgroup.

6. Click Yes to accept Certificate Warning.

Figure 99 Click Certificate

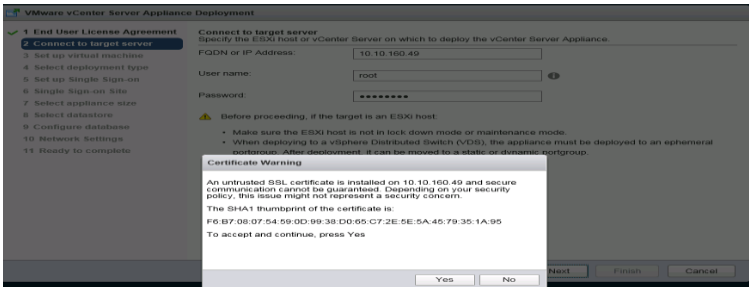

Connect to target server
Specify the ESXi host or vCenter Server on which to deploy the vCenter Server Appliance.

FQDN or IP Address: 10.10.160.49

User name: root

Password: ••••••••

Before proceeding, if the target is an ESXi host:

- Make sure the ESXi host is not in lock down mode or maintenance mode.
- When deploying to a vSphere Distributed Switch (VDS), the appliance must be deployed to an ephemeral portgroup. After deployment, it can be moved to a static or dynamic portgroup.

Certificate Warning

An untrusted SSL certificate is installed on 10.10.160.49 and 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:

F6:B7:08:07:54:59:0D:99:38:D0:65:C7:2E:5E:5A:45:79:35:1A:95

To accept and continue, press Yes

7. Provide a name for the vCenter appliance, then click Next to continue.

Figure 100 Provide a Name and Password for VC Appliance

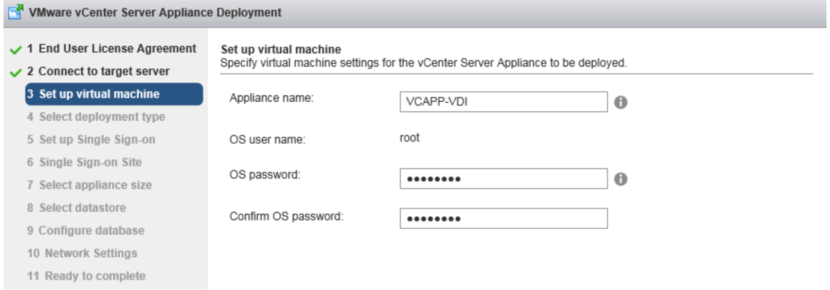

Set up virtual machine
Specify virtual machine settings for the vCenter Server Appliance to be deployed.

Appliance name: VCAPP-VDI

OS user name: root

OS password: ••••••••

Confirm OS password: ••••••••

8. Select Install vCenter Server with and Embedded Platform Services Controller (unless your environment already has a PSC).

Figure 101 Select Platform Services Controller Applicable

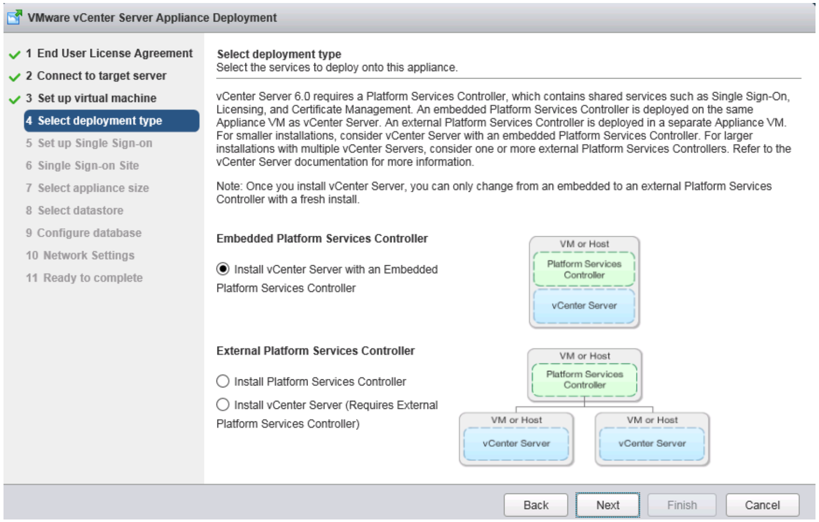

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- 4 Select deployment type
- 5 Set up Single Sign-on
- 6 Single Sign-on Site
- 7 Select appliance size
- 8 Select datastore
- 9 Configure database
- 10 Network Settings
- 11 Ready to complete

Select deployment type
Select the services to deploy onto this appliance.

vCenter Server 6.0 requires a Platform Services Controller, which contains shared services such as Single Sign-On, Licensing, and Certificate Management. An embedded Platform Services Controller is deployed on the same Appliance VM as vCenter Server. An external Platform Services Controller is deployed in a separate Appliance VM. For smaller installations, consider vCenter Server with an embedded Platform Services Controller. For larger installations with multiple vCenter Servers, consider one or more external Platform Services Controllers. Refer to the vCenter Server documentation for more information.

Note: Once you install vCenter Server, you can only change from an embedded to an external Platform Services Controller with a fresh install.

Embedded Platform Services Controller

● Install vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Install Platform Services Controller

○ Install vCenter Server (Requires External Platform Services Controller)

9. Create a new SSO domain (unless your environment already has and SSO domain. Multiple SSO domains can co-exist).

Figure 102 Provide Single Sign On Password and Site Name Credentials

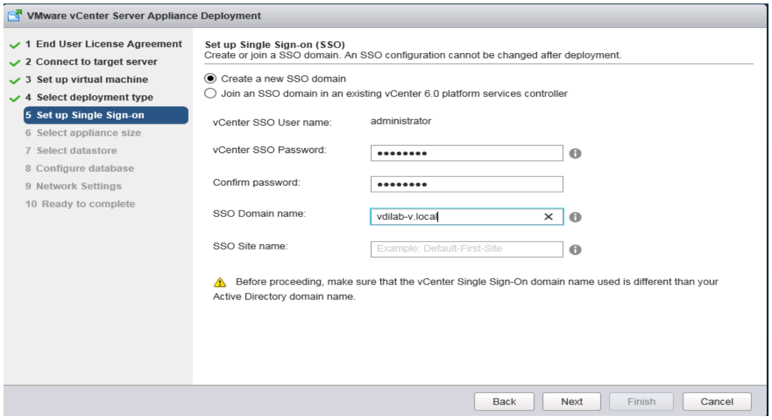

VMware vCenter Server Appliance Deployment

- ✓ 1 End User License Agreement
- ✓ 2 Connect to target server
- ✓ 3 Set up virtual machine
- ✓ 4 Select deployment type
- 5 Set up Single Sign-on
- 6 Select appliance size
- 7 Select datastore
- 8 Configure database
- 9 Network Settings
- 10 Ready to complete

Set up Single Sign-on (SSO)
Create or join a SSO domain. An SSO configuration cannot be changed after deployment.

● Create a new SSO domain
○ Join an SSO domain in an existing vCenter 6.0 platform services controller

vCenter SSO User name: administrator

vCenter SSO Password: ••••••••

Confirm password: ••••••••

SSO Domain name: vdilab-v.local

SSO Site name: Example: Default-First-Site

⚠ Before proceeding, make sure that the vCenter Single Sign-On domain name used is different than your Active Directory domain name.

10. Select the proper appliance size for your deployment. In our study, Large was selected.

Figure 103 Select the Appropriate Appliance Size

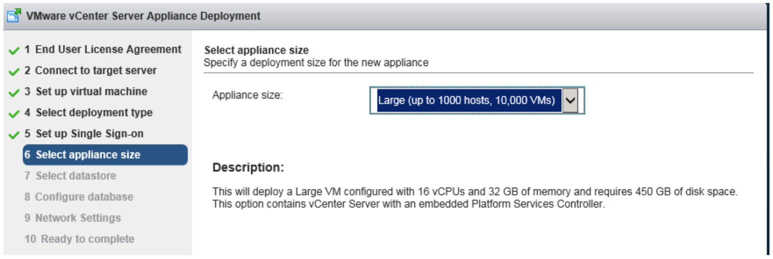

Select appliance size
Specify a deployment size for the new appliance

Appliance size: Large (up to 1000 hosts, 10,000 VMs)

Description:

This will deploy a Large VM configured with 16 vCPUs and 32 GB of memory and requires 450 GB of disk space. This option contains vCenter Server with an embedded Platform Services Controller.

11. Select the Data store

Figure 104 Select the Data store to Install Appliance

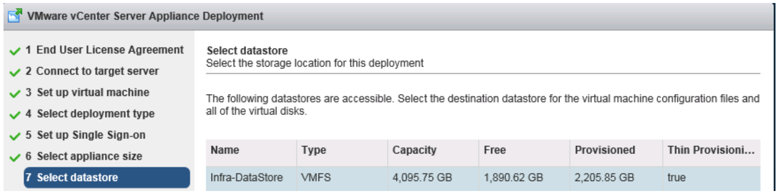

Select datastore
Select the storage location for this deployment

The following datastores are accessible. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Type	Capacity	Free	Provisioned	Thin Provisioni...
Infra-DataStore	VMFS	4,095.75 GB	1,890.62 GB	2,205.85 GB	true

12. In our study we used the embedded PostgreSQL database.

Figure 105 PostgreSQL

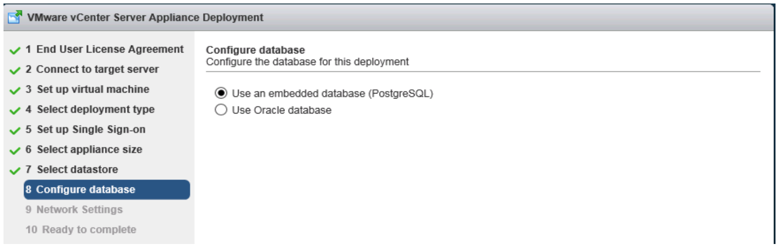

Configure database
Configure the database for this deployment

● Use an embedded database (PostgreSQL)
○ Use Oracle database

13. Enter Network Settings for 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.

Figure 106 Provide the Necessary Network Gateways and DNS Server Information

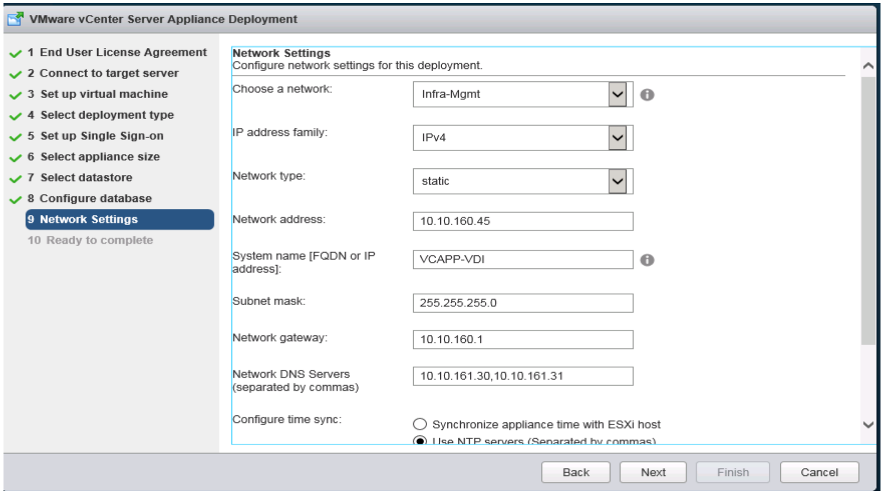

Network Settings
Configure network settings for this deployment.

Choose a network: Infra-Mgmt

IP address family: IPv4

Network type: static

Network address: 10.10.160.45

System name [FQDN or IP address]: VCAPP-VDI

Subnet mask: 255.255.255.0

Network gateway: 10.10.160.1

Network DNS Servers (separated by commas): 10.10.161.30,10.10.161.31

Configure time sync: ○ Synchronize appliance time with ESXi host
● Use NTP servers (Separated by commas)

14. Review the Install Settings and click Finish.

15. Click Next to complete installing appliance.

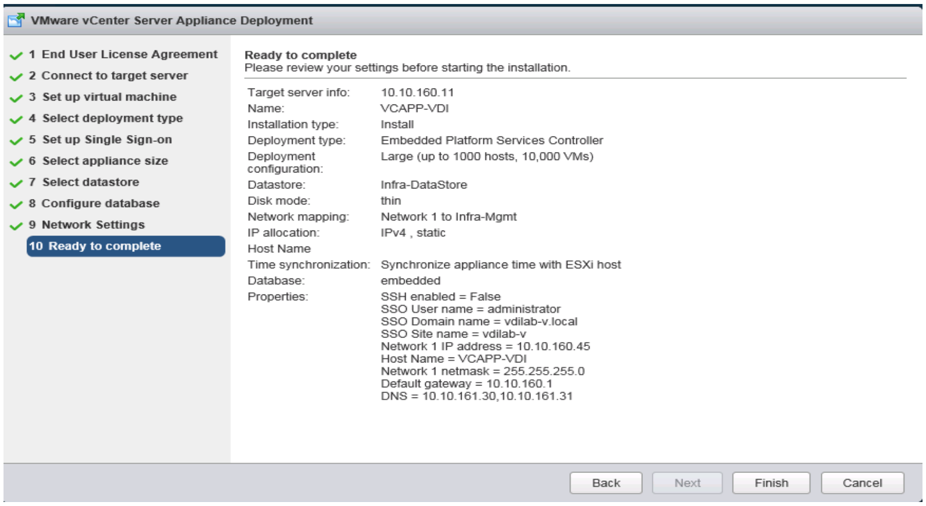

Ready to complete
Please review your settings before starting the installation.

Target server info:	10.10.160.11
Name:	VCAPP-VDI
Installation type:	Install
Deployment type:	Embedded Platform Services Controller
Deployment configuration:	Large (up to 1000 hosts, 10,000 VMs)
Datastore:	Infra-DataStore
Disk mode:	thin
Network mapping:	Network 1 to Infra-Mgmt
IP allocation:	IPv4 , static
Host Name	
Time synchronization:	Synchronize appliance time with ESXi host
Database:	embedded
Properties:	SSH enabled = False SSO User name = administrator SSO Domain name = vdilab-v.local SSO Site name = vdilab-v Network 1 IP address = 10.10.160.45 Host Name = VCAPP-VDI Network 1 netmask = 255.255.255.0 Default gateway = 10.10.160.1 DNS = 10.10.161.30,10.10.161.31

16. When your install completes successfully, you can now login to your Web Client and begin adding hosts and configuring clusters.

17. Login in to VCenter Appliances Web GUI https://10.10.160.49/vsphere-client.

Figure 107 VCenter Appliance Web GUI

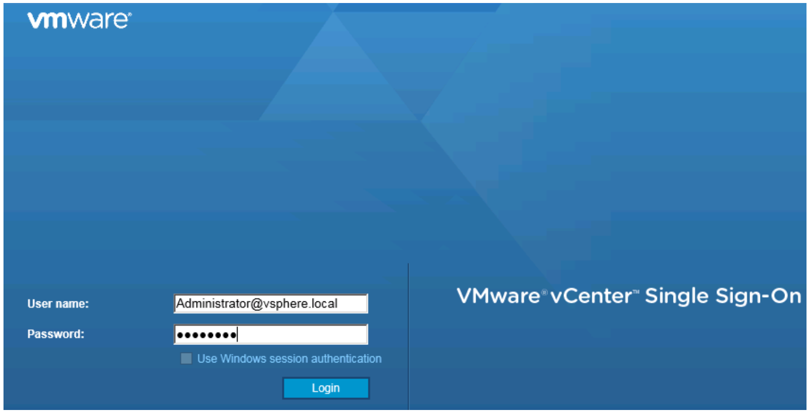

vmware

User name: Administrator@vsphere.local

Password: ••••••••

□ Use Windows session authentication

Login

VMware® vCenter™ Single Sign-On

18. Log into the vSphere Web Client.

Figure 108 Login using IP Address of the Appliance and Download vSphere

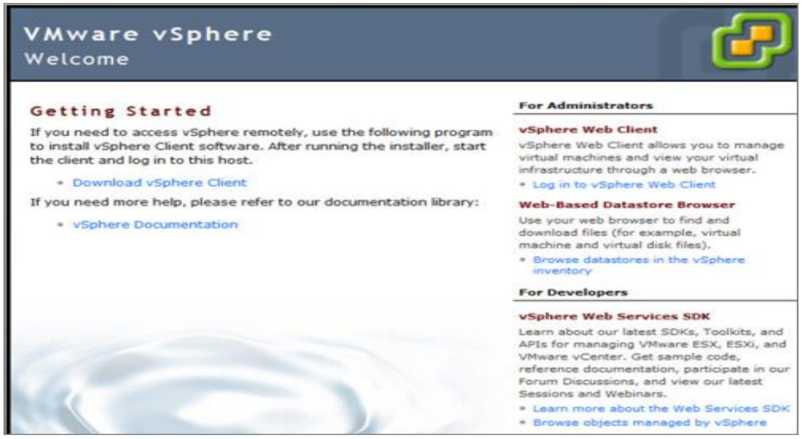

VMware vSphere
Welcome

Getting Started

If you need to access vSphere remotely, use the following program to install vSphere Client software. After running the installer, start the client and log in to this host.

- Download vSphere Client

If you need more help, please refer to our documentation library:

- vSphere Documentation

For Administrators

vSphere Web Client

vSphere Web Client allows you to manage virtual machines and view your virtual infrastructure through a web browser.

- Log in to vSphere Web Client

Web-Based Datastore Browser

Use your web browser to find and download files (for example, virtual machine and virtual disk files).

- Browse datastores in the vSphere inventory

For Developers

vSphere Web Services SDK

Learn about our latest SDKs, Toolkits, and APIs for managing VMware ESX, ESXi, and VMware vCenter. Get sample code, reference documentation, participate in our Forum Discussions, and view our latest Sessions and Webinars.

- Learn more about the Web Services SDK
- Browse objects managed by vSphere

19. Click the link labeled Log in to vSphere Web Client.

20. If prompted, run the VMWare Remote Console Plug-in.

21. Log in using the root user name and password.

22. Click the vCenter link on the left panel.

Figure 109 Login vSphere Web GUI

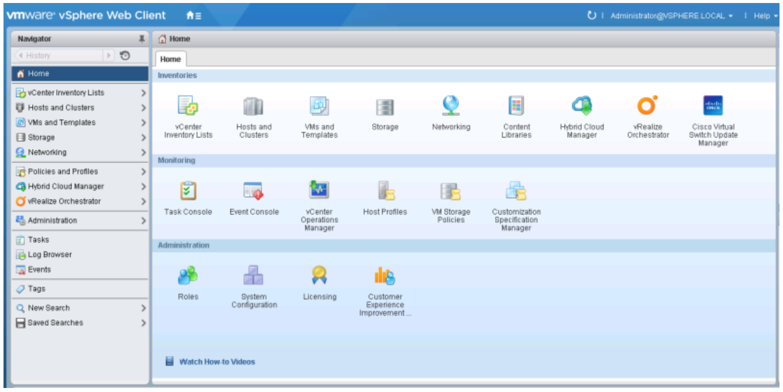

23. Click the Datacenters link on the left panel.

24. To create a Datacenter, click the icon in the center pane which has the green plus symbol above it.

Figure 110 Create Data Center, example: VDI-DC

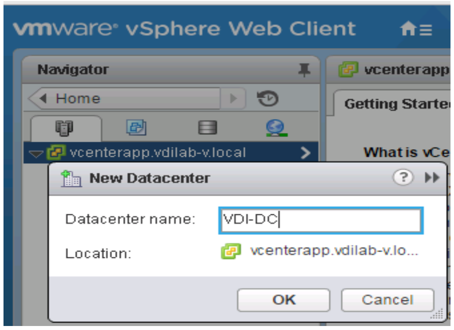

Datacenter name: VDI-DC

Location: vcenterapp.vdilab-v.lo...

25. Type VDI-DC as the Datacenter name.

26. Click the vCenter server available in the list. Click OK to continue.

Figure 111 Create a Cluster

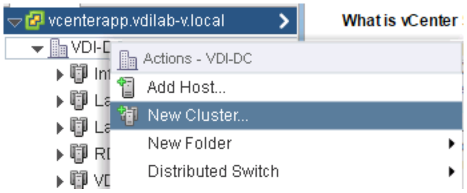

27. Right-click Datacenters > VDI-DC in the list in the center pane, then click New Cluster.

28. Name the cluster Infra.

29. Select DRS. Retain the default values.

30. Select vSphere HA. Retain the default values.

Figure 112 Configure Cluster Specific Setting

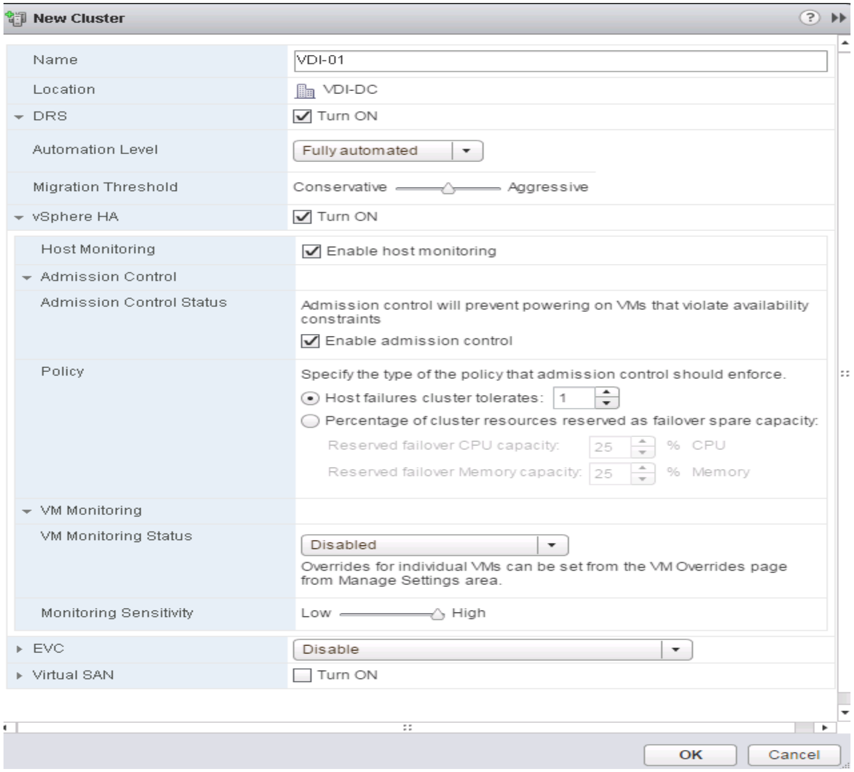

New Cluster

Name: VDI-01

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

Admission Control Status: Admission control will prevent powering on VMs that violate availability constraints

☑ Enable admission control

Policy: Specify the type of the policy that admission control should enforce.

● Host failures cluster tolerates: 1

○ Percentage of cluster resources reserved as failover spare capacity:

Reserved failover CPU capacity: 25 % CPU

Reserved failover Memory capacity: 25 % Memory

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

Virtual SAN: □ Turn ON

If mixing Cisco UCS B 200 M4 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.

31. Click OK to create the new cluster.

32. Click VDI-DC in the left pane.

Figure 113 Add a ESXi Host

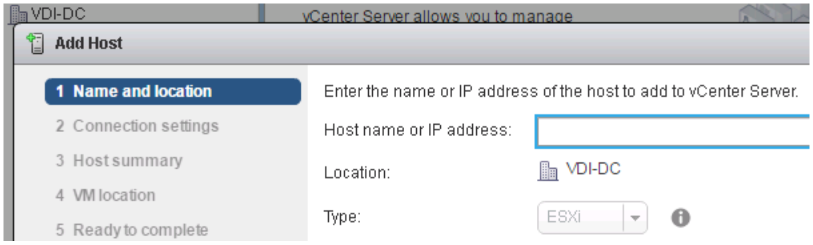

Enter the name or IP address of the host to add to vCenter Server.

Host name or IP address:

Location: VDI-DC

Type: ESXi

33. Right-click Infra in the center pane and click Add Host.

34. Type the host IP address and click Next.

35. Type root as the user name and root password as the password. Click Next to Continue.

Figure 114 ESXi Host Certificate

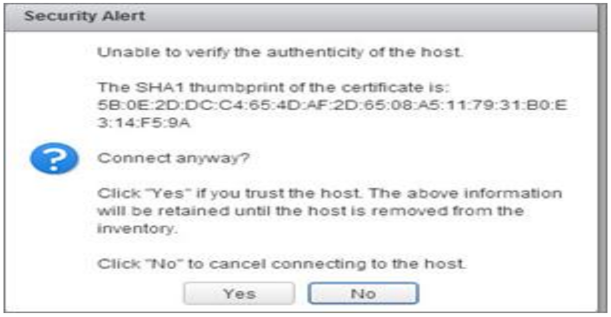

Unable to verify the authenticity of the host.

The SHA1 thumbprint of the certificate is:
5B:0E:2D:DC:C4:65:4D:AF:2D:65:08:A5:11:79:31:B0:E3:14: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.

36. Click Yes to accept the certificate.

37. Review the host details, and click Next to continue.

38. Assign a license, and click Next to continue.

39. Click Next to continue.

40. Click Next to continue.

41. Review the configuration parameters then click Finish to add the host.

42. Repeat this for the other hosts and clusters

When completed, the vCenter cluster configuration is comprised of the following clusters, including a cluster to manage the workload launcher hosts:

Figure 115 Solution vCenter Cluster Configuration

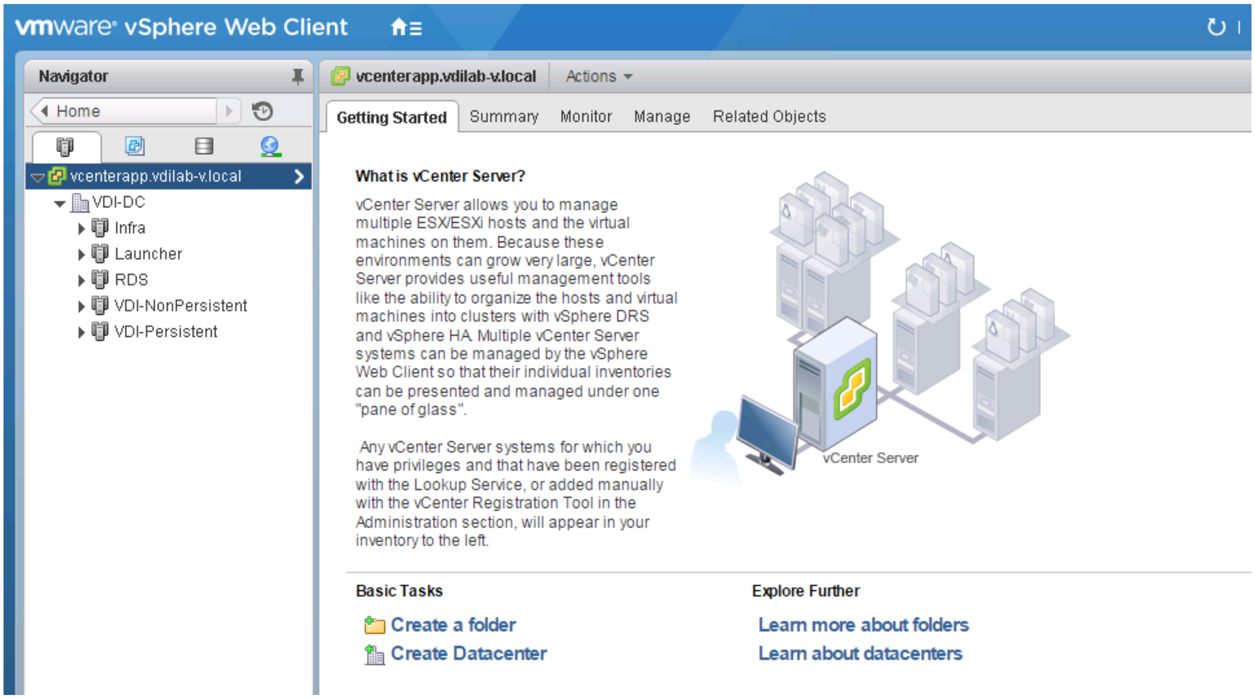

What is vCenter Server?

vCenter Server allows you to manage multiple ESX/ESXi hosts and the virtual machines on them. Because these environments can grow very large, vCenter Server provides useful management tools like the ability to organize the hosts and virtual machines into clusters with vSphere DRS and vSphere HA. Multiple vCenter Server systems can be managed by the vSphere Web Client so that their individual inventories can be presented and managed under one "pane of glass".

Any vCenter Server systems for which you have privileges and that have been registered with the Lookup Service, or added manually with the vCenter Registration Tool in the Administration section, will appear in your inventory to the left.

Install and Configure VSUM and Cisco Nexus 1000v

Install Cisco Virtual Switch Update Manager

Verifying the Authenticity of the Cisco-Signed Image (Optional)

Before you install the Cisco Nexus1000v-vsum.1.5.x-pkg.zip image, you have the option to validate its authenticity. In the zip file, there is a signature.txt file that contains an SHA-512 signature and an executable script that can be used to verify the authenticity of the Nexus1000v-vsum.1.5.x-pkg.zip image.

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:

Verifying the authenticity of an image is optional. You can still install the image without validating its authenticity.

1. Copy the following files to a directory on the Linux machine:
   - Nexus1000v-vsum.1.5.x-pkg.zip image
   - signature.txt file
   - cisco_n1k_image_validation_v_1_5_x script
2. Make sure the script is executable.
   - chmod 755 cisco_n1k_image_validation_v_1_5_x
3. Run the script.
   - ./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
4. Run the script.
   - ./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
5. Check the output. If the validation is successful, the following message displays:

   Authenticity of Cisco-signed image Nexus1000v-vsum.1.5.x-pkg.zip has been successfully verified!

Install Cisco Virtual Switch Update Manager

VMware vSphere Web Client

To install the Cisco Virtual Switch Upgrade Manager from OVA in the VMware virtual environment, complete the following steps:

1. Log into the VMware vSphere Web Client.
2. In the pane on the right, click VMs and Templates.
3. In the center pane, select Actions > Deploy OVF Template.
4. Select Browse and browse to and select the Nexus1000v-vsum.1.5.x.ova file.
5. Click Open.
6. Click Next.

Figure 116 Select the Cisco Nexus 1000v OVF File to Install

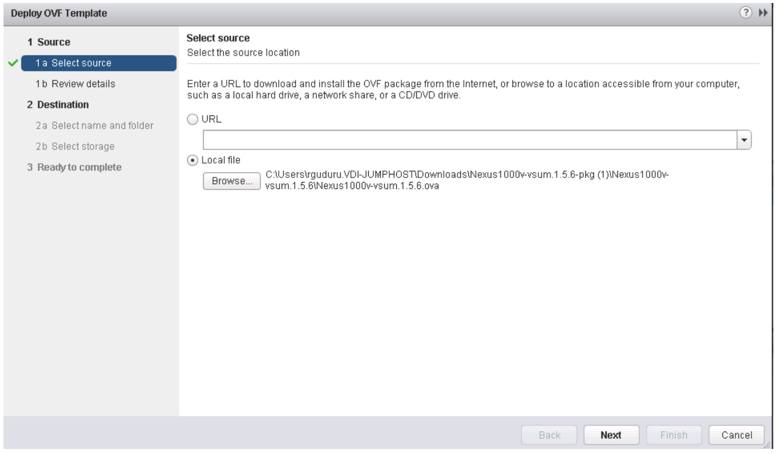

Deploy OVF Template

1 Source

✓ 1a Select source

1b Review details

2 Destination

2a Select name and folder

2b Select storage

3 Ready to complete

Select source

Select the source location

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

● Local file

Browse... C:\Users\rguduru.VDI-JUMPHOST\Downloads\Nexus1000v-vsum.1.5.6-pkg (1)\Nexus1000v-vsum.1.5.6\Nexus1000v-vsum.1.5.6.ova

7. Review the details and click Next.
8. Review and click Next.

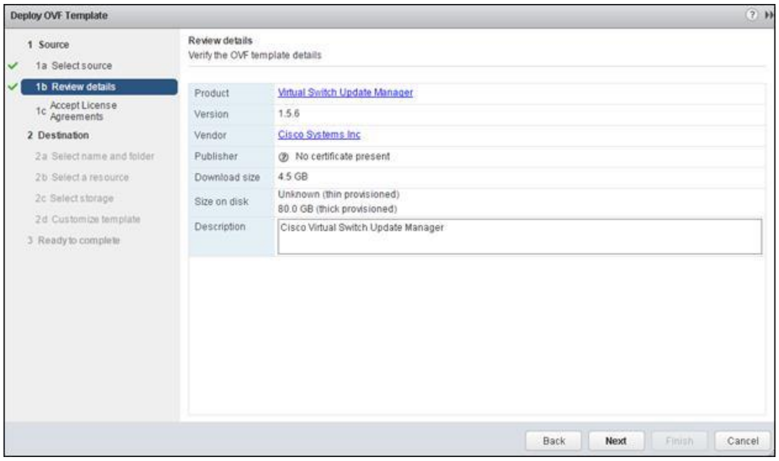

Review details
Verify the OVF template details

Product	Virtual Switch Update Manager
Version	1.5.6
Vendor	Cisco Systems Inc
Publisher	No certificate present
Download size	4.5 GB
Size on disk	Unknown (thin provisioned) 80.0 GB (thick provisioned)
Description	Cisco Virtual Switch Update Manager

9. Click Accept to accept the License Agreement and click Next.

10. Name the Virtual Machine, select the VDI-DC datacenter and click Next.

Figure 117 Provide Name for VSM to be Configured

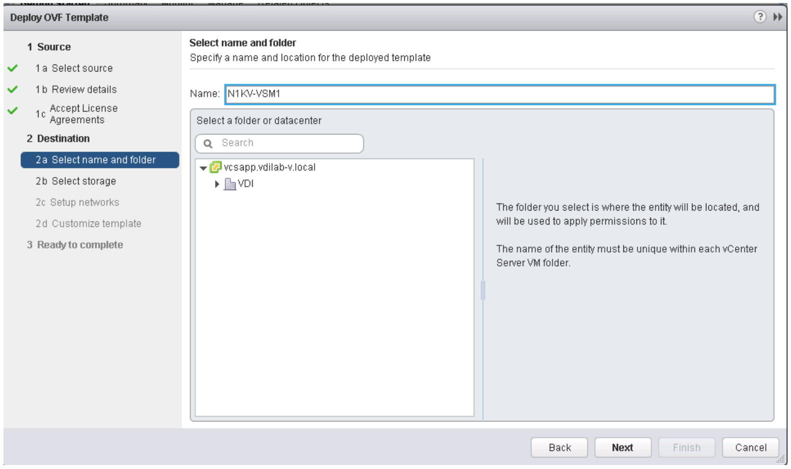

Select name and folder
Specify a name and location for the deployed template

Name: N1KV-VSM1

Select a folder or datacenter

The folder you select is where the entity will be located, and will be used to apply permissions to it.

The name of the entity must be unique within each vCenter Server VM folder.

11. Select the Infra cluster and click Next.

12. Select Infra-Datastore and the Thin Provision virtual disk format and click Next.

Figure 118 Select the Datastore

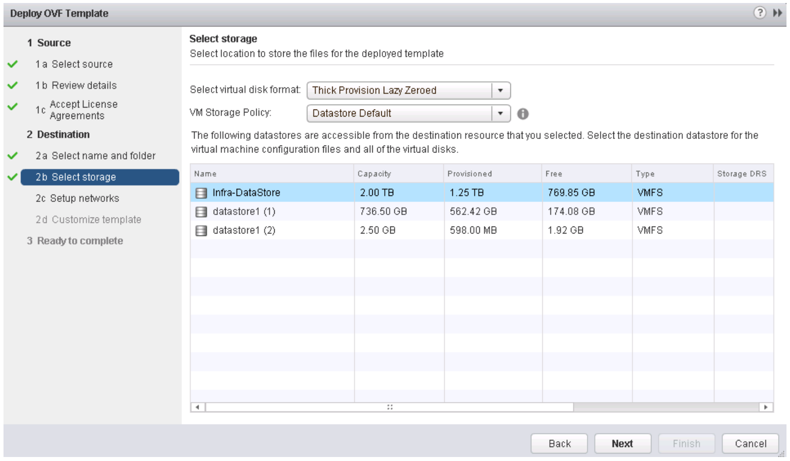

Select storage
Select location to store the files for the deployed template

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
Infra-DataStore	2.00 TB	1.25 TB	769.85 GB	VMFS	
datastore1 (1)	736.50 GB	562.42 GB	174.08 GB	VMFS	
datastore1 (2)	2.50 GB	598.00 MB	1.92 GB	VMFS	

13. Select the MGMT Network and click Next.

14. Fill in the Networking Properties.

15. Expand the vCenter Properties and fill in the fields.

16. Click Next.

17. Review all settings and click Finish.

18. Wait for the Deploy OVF template task to complete.

19. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.

20. Expand the Infrastructure cluster and select the Virtual Switch Update Manager VM.

21. In the center pane, select Launch Remote Console. If a security warning pops up, click Allow.

22. If a security certificate warning pops up, click Connect Anyway.

23. Power on the Virtual Switch Update Manager VM.

24. When the VM has completely booted up, log out and log back into the VMware vSphere Web Client.

25. Review and click Next to install the click Nexus 1000V.

Figure 119 Customize Template

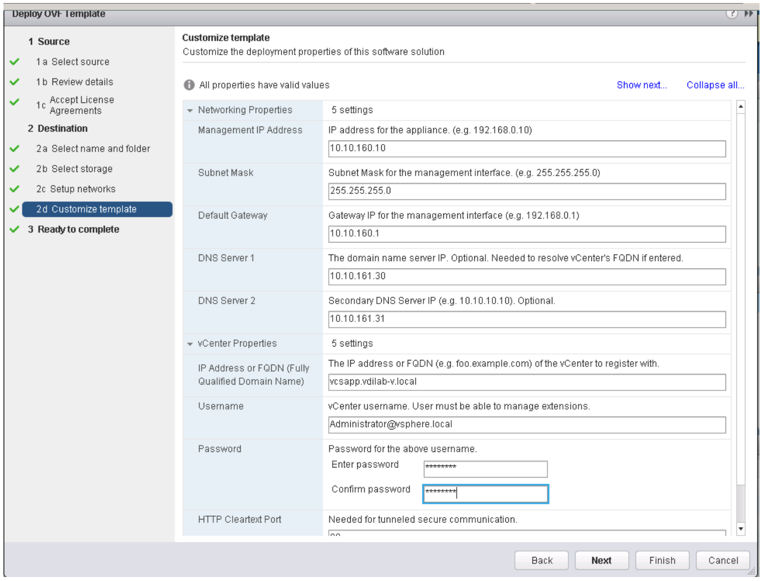

Customize template

Customize the deployment properties of this software solution

All properties have valid values

Networking Properties — 5 settings

Management IP Address — IP address for the appliance. (e.g. 192.168.0.10): 10.10.160.10

Subnet Mask — Subnet Mask for the management interface. (e.g. 255.255.255.0): 255.255.255.0

Default Gateway — Gateway IP for the management interface (e.g. 192.168.0.1): 10.10.160.1

DNS Server 1 — The domain name server IP. Optional. Needed to resolve vCenter's FQDN if entered.: 10.10.161.30

DNS Server 2 — Secondary DNS Server IP (e.g. 10.10.10.10). Optional.: 10.10.161.31

vCenter Properties — 5 settings

IP Address or FQDN (Fully Qualified Domain Name) — The IP address or FQDN (e.g. foo.example.com) of the vCenter to register with.: vcsapp.vdilab-v.local

Username — vCenter username. User must be able to manage extensions.: Administrator@vsphere.local

Password — Password for the above username.
Enter password: *********
Confirm password: *********

HTTP Cleartext Port — Needed for tunneled secure communication.

About the Cisco VSUM GUI

The following lists the details of the Cisco VSUM GUI:

- Cisco VSUM is a virtual appliance that is registered as a plug-in to the VMware vCenter Server.
- The Cisco VSUM is the GUI that you use to install, migrate, monitor, and upgrade the VSMs in high availability (HA) or standalone mode and the VEMs on ESX/ESXi hosts.

Figure 120 VMware vSphere Web Client—Home Page

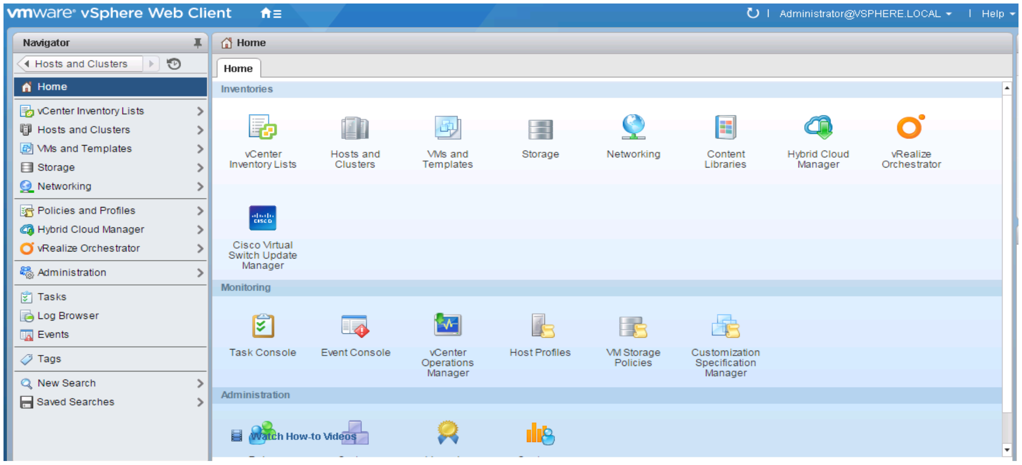

Figure 121 Cisco VSUM—Home Page

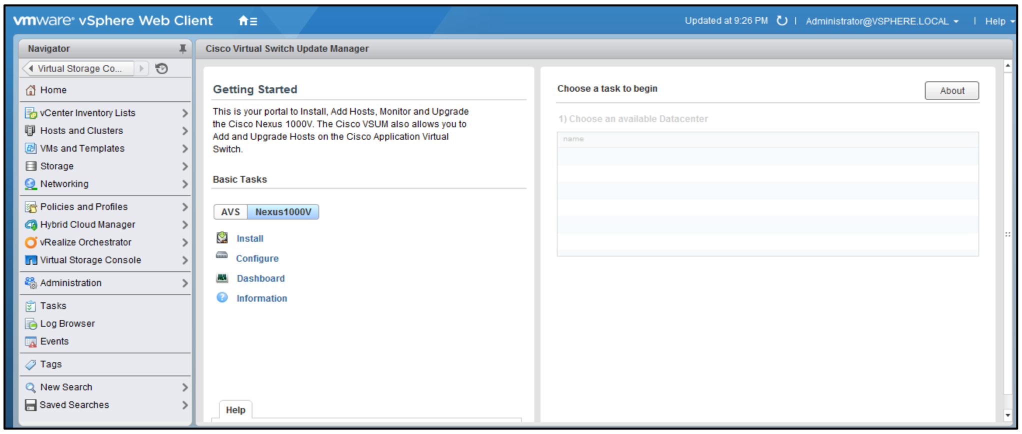

Getting Started

This is your portal to Install, Add Hosts, Monitor and Upgrade the Cisco Nexus 1000V. The Cisco VSUM also allows you to Add and Upgrade Hosts on the Cisco Application Virtual Switch.

Basic Tasks

Choose a task to begin

1) Choose an available Datacenter

Install Cisco Nexus 1000V using Cisco VSUM

VMware vSphere Web Client

To install the Cisco Nexus 1000V switch by creating a new VSM, complete the following steps:

Optionally, an existing VSM can be used that is provided by a Cisco Nexus Cloud Services Platform (CSP).

1. Log in to VMware vSphere Web Client and choose Home > Cisco Virtual Switch Update Manager > Nexus 1000V > Install, and then choose the data center. The installation screen appears.

Figure 122 **Configuring Virtual Ethernet Modules**

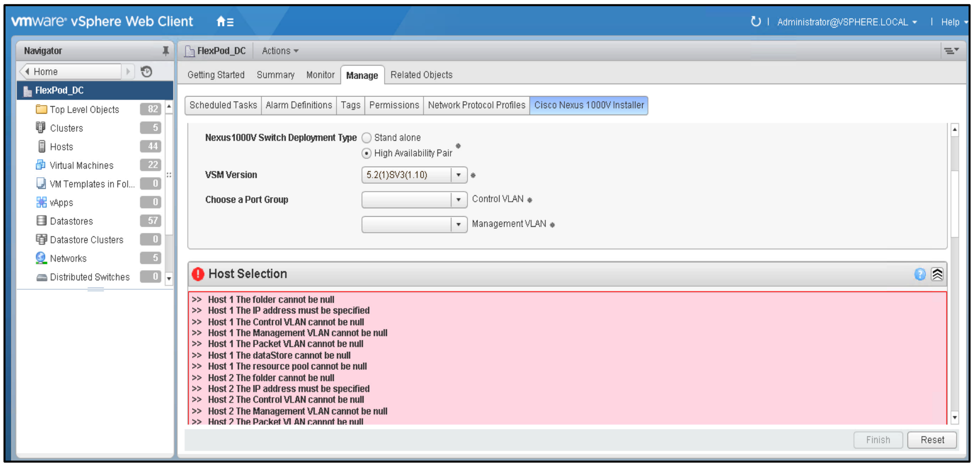

Nexus1000V Switch Deployment Type ○ Stand alone ● High Availability Pair *

VSM Version 5.2(1)SV3(1.10) *

Choose a Port Group: Control VLAN * ; Management VLAN *

```
>>  Host 1 The folder cannot be null
>>  Host 1 The IP address must be specified
>>  Host 1 The Control VLAN cannot be null
>>  Host 1 The Management VLAN cannot be null
>>  Host 1 The Packet VLAN cannot be null
>>  Host 1 The dataStore cannot be null
>>  Host 1 The resource pool cannot be null
>>  Host 2 The folder cannot be null
>>  Host 2 The IP address must be specified
>>  Host 2 The Control VLAN cannot be null
>>  Host 2 The Management VLAN cannot be null
>>  Host 2 The Packet VLAN cannot be null
```

2. In the Nexus 1000v Switch Deployment area, choose I want to deploy new control plane (VSM).

3. In the Cisco Nexus 1000V Switch Deployment Type area, install the switches as an HA pair. By default, the High Availability Pair is selected.

4. Choose the control port group for the switch.

5. Choose the management port group for the switch.

> The Cisco Nexus 1000V VSM uses the management network to communicate with vCenter Server and ESXi. The management and control port group can use the same VLAN.

6. In the Host Selection area, click Suggest to choose two hosts based on the details provided in the Cisco Nexus 1000V Switch Deployment Type area. The IP address of the hosts on which the switch will be deployed.

7. The primary switch is deployed on Infrastructure Host 1 and the secondary switch is deployed on Infrastructure Host 2. Click Pick a Host to override the system choices.

8. Choose the system-selected datastore that you want to override. Choose PURE Infra-Datastore as the datastore for each host.

9. Provide Host IP address where the Virtual Ethernet Modules to be created. (note it requires two ESXI hosts for installing VEM primary and secondary modules for redundancy purpose)

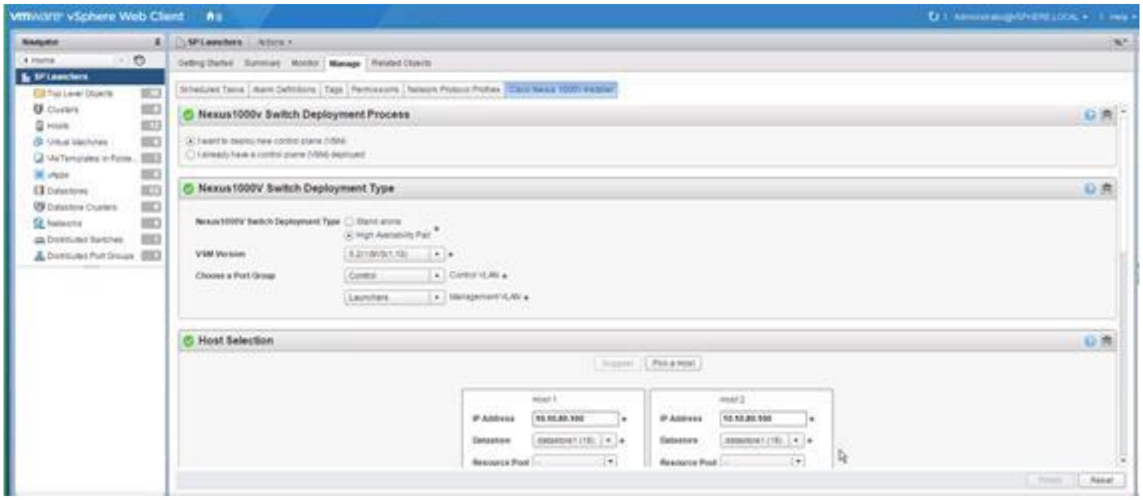

10. In the Switch Configuration area, enter 70 as the domain ID for the switch.

11. The domain ID is common for both the primary and secondary switches and it should be unique for every new switch. The range for the domain is from 1 to 1023.

12. In the Virtual Supervisor Module (VSM) configuration area, enter the Switch Name, IP Address, Subnet Mask, and Gateway Address.

13. Do not select Default Port Profiles.

14. Enter the Password and Confirm Password for Admin.

15. Provide switch name, password and IP address.

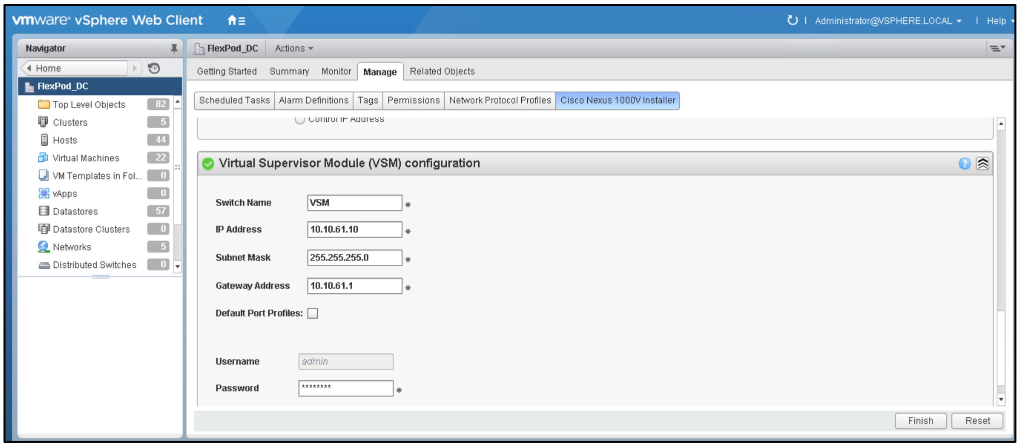

Virtual Supervisor Module (VSM) configuration

Switch Name: VSM

IP Address: 10.10.61.10

Subnet Mask: 255.255.255.0

Gateway Address: 10.10.61.1

Default Port Profiles: □

Username: admin

Password: ********

16. Click Finish to install the Cisco Nexus 1000V switch.

The Cisco Nexus 1000V installation is confirmed when the primary task Create Nexus 1000v Switch has the status Completed. A typical installation of the switch takes about 4 minutes.

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands:

Any VLAN that has a VMKernal port should be assigned as a system VLAN on both the **up-link** and the **vEthernet** ports of the virtual switch.

```
config t
ntp server <<var_switch_a_ntp_ip>> use-vrf management
ntp server <<var_switch_b_ntp_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id>> 160
name IB-MGMT-VLAN
vlan <<var_vmotion_vlan_id>> 166
name vMotion-VLAN
```

The Cisco Nexus 1000V is currently limited to 1024 Max ports per profile. This solution is comprised of 3500 plus virtual desktop machines for the user workload and requires four dedicated port-profiles(VDI,VDI-1,VDI-2,VDI-3).

```
vlan <<var_vdi_vlan_id>> 162
name VDI
vlan <<var_vdi_vlan_id>> 162
```

```
name VDI-1
vlan <<var_vdi_vlan_id>> 162
name VDI-2
vlan <<var_vdi_vlan_id>> 162
name VDI-3
vlan <<var_vm-traffic_vlan_id>> 161
name Infra
vlan <<var_vm-traffic_vlan_id>> 164
name OB-Mgmt
vlan <<var_native_vlan_id>> 1
name Native-VLAN
exit
port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>> 1
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>,
<<var_nfs_vlan_id>>, <<var_vmotion_vlan_id>>, <<var_vm-
infra_vlan_id>> 160-166
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-infra_vlan_id>> 160-166
system mtu 9000
state enabled
port-profile type vethernet IB-MGMT
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>> 160
no shutdown
system vlan <<var_ib-mgmt_vlan_id>> 160
state enabled
port-profile type vethernet vMotion
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>> 166
no shutdown
system vlan <<var_vmotion_vlan_id>> 166
state enabled
port-profile type vethernet INFRA
vmware port-group
switchport mode access
switchport access vlan <<var_vm-infra_vlan_id>> 161
```

```
no shutdown
system vlan <<var_vm-infra_vlan_id>> 161
state enabled
port-profile type vethernet n1kv-L3
capability l3control
vmware port-group
switchport mode access
port-profile type vethernet VDI
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 162
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 162
state enabled
port-profile type vethernet VDI-1
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 162
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 162
state enabled
port-profile type vethernet VDI-2
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 162
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 162
state enabled
port-profile type vethernet VDI-3
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 162
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 162
state enabled
switchport access vlan <<var_OB-MGMT_vlan_id>> 164
no shutdown
system vlan <<var_OB-MGMT_vlan_id>> 164
state enabled
```

```
exit
copy run start
```

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

To and VMware ESXi hosts, complete the following steps:

1. Back in the VMware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the VDI-DC datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the Infrastructure ESXi Cluster and select one of the Infrastructure Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile.

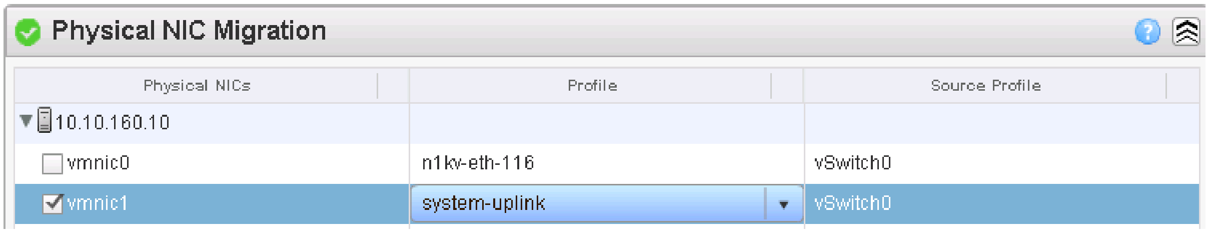

Physical NICs	Profile	Source Profile
▼ 10.10.160.10		
□ vmnic0	n1kv-eth-116	vSwitch0
☑ vmnic1	system-uplink ▾	vSwitch0

12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.
13. All VMkernel ports should already have the appropriate checkboxes selected.

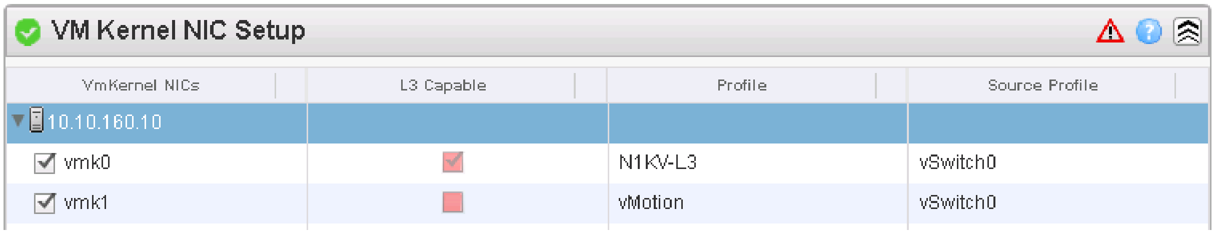

VM Kernel NIC Setup

VmKernel NICs	L3 Capable	Profile	Source Profile
▼ 10.10.160.10			
☑ vmk0	☑	N1KV-L3	vSwitch0
☑ vmk1	☐	vMotion	vSwitch0

14. Scroll down to VM Migration and expand both ESXi hosts.

15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.

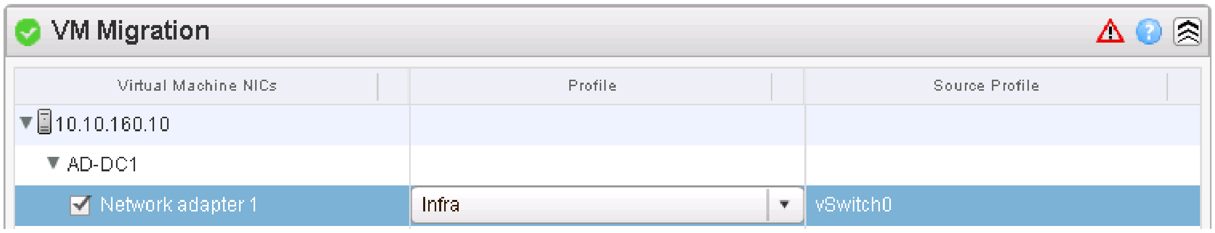

VM Migration

Virtual Machine NICs	Profile	Source Profile
▼ 10.10.160.10		
▼ AD-DC1		
☑ Network adapter 1	Infra ▾	vSwitch0

16. Click Finish.

The progress of the virtual switch installation can be monitored from the c# interface.Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

To migrate the ESXi host redundant network ports, complete the following steps:

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.

2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.

3. In the center pane, select the Manage tab, then select Networking.

4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.

5. Click Yes to remove vSwitch0. It may be necessary to refresh the Web Client to see the deletion.

6. The Nexus 1000V VSM should now be the only virtual switch. Select it and select the third icon above it under Virtual switches (Manage the physical network adapters connected to the selected switch).

7. Click the green plus sign to add an adapter.

8. For UpLink01, select the system-uplink port group and make sure vmnic0 is the Network adapter. Click OK.

9. Click OK to complete adding the Uplink. It may be necessary to refresh the Web Client to see the addition.

Figure 123 VSphere Web Client Checking the Uplinks Status

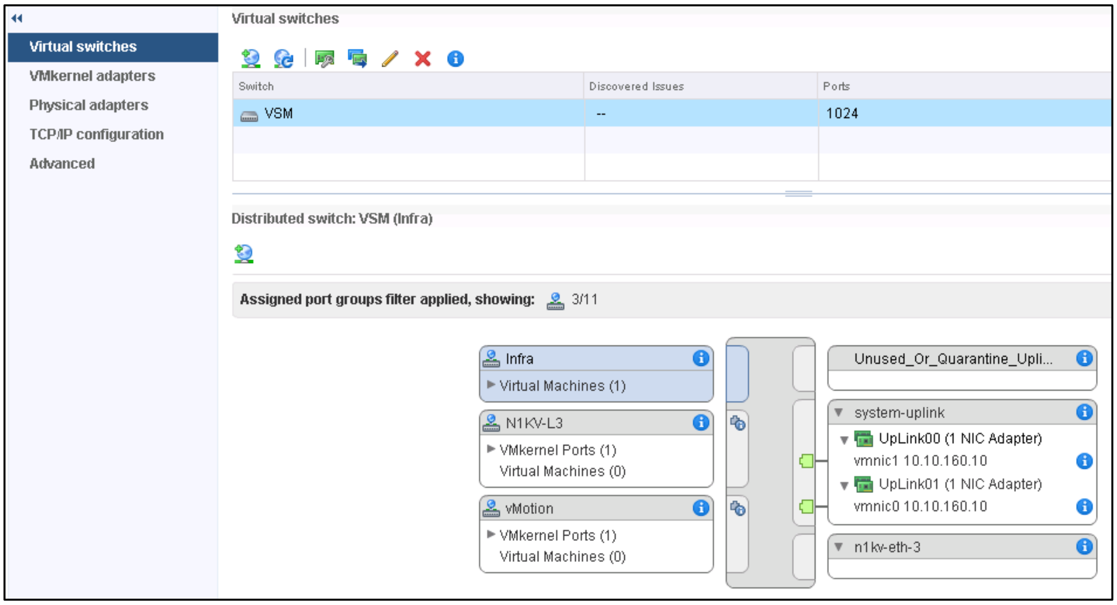

Switch	Discovered Issues	Ports
VSM	--	1024

10. Repeat this procedure for the other ESXi host.

11. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.

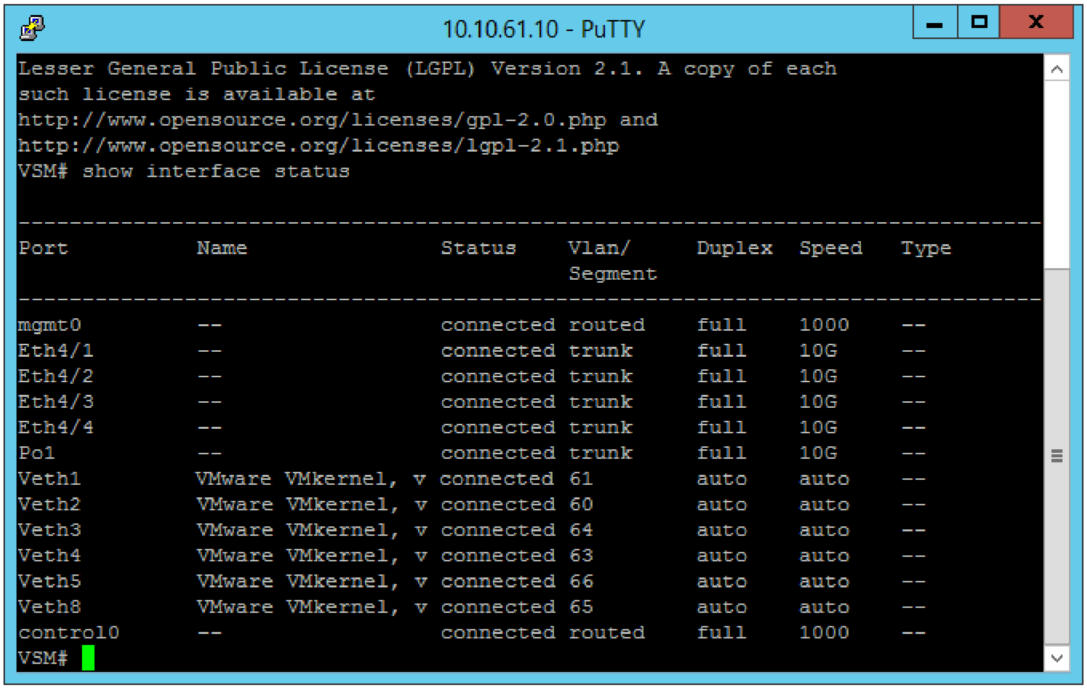

```
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
VSM# show interface status

--------------------------------------------------------------------------------
Port          Name               Status    Vlan/     Duplex  Speed   Type
                                           Segment
--------------------------------------------------------------------------------
mgmt0         --                 connected routed    full    1000    --
Eth4/1        --                 connected trunk     full    10G     --
Eth4/2        --                 connected trunk     full    10G     --
Eth4/3        --                 connected trunk     full    10G     --
Eth4/4        --                 connected trunk     full    10G     --
Po1           --                 connected trunk     full    10G     --
Veth1         VMware VMkernel, v connected 61        auto    auto    --
Veth2         VMware VMkernel, v connected 60        auto    auto    --
Veth3         VMware VMkernel, v connected 64        auto    auto    --
Veth4         VMware VMkernel, v connected 63        auto    auto    --
Veth5         VMware VMkernel, v connected 66        auto    auto    --
Veth8         VMware VMkernel, v connected 65        auto    auto    --
control0      --                 connected routed    full    1000    --
VSM#
```

12. Run show module and verify that the one ESXi host is present as a module.

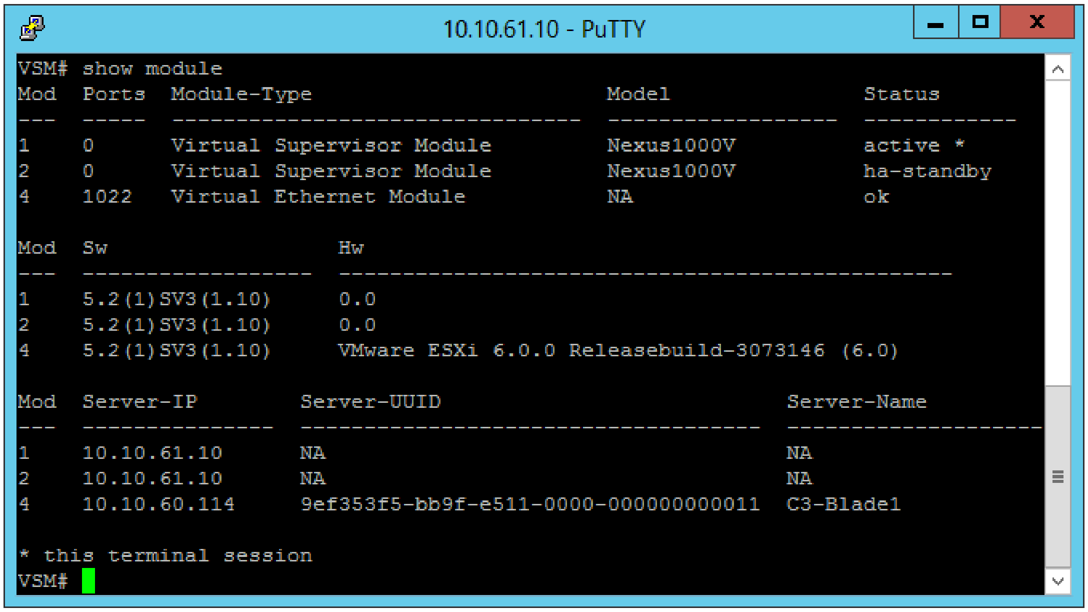

```
VSM# show module
Mod  Ports  Module-Type                         Model               Status
---  -----  ----------------------------------  ------------------  ------------
1    0      Virtual Supervisor Module           Nexus1000V          active *
2    0      Virtual Supervisor Module           Nexus1000V          ha-standby
4    1022   Virtual Ethernet Module             NA                  ok

Mod  Sw                  Hw
---  ------------------  ------------------------------------------------
1    5.2(1)SV3(1.10)     0.0
2    5.2(1)SV3(1.10)     0.0
4    5.2(1)SV3(1.10)     VMware ESXi 6.0.0 Releasebuild-3073146 (6.0)

Mod  Server-IP        Server-UUID                           Server-Name
---  ---------------  ------------------------------------  --------------------
1    10.10.61.10      NA                                    NA
2    10.10.61.10      NA                                    NA
4    10.10.60.114     9ef353f5-bb9f-e511-0000-000000000011  C3-Blade1

* this terminal session
VSM#
```

13. Repeat the above steps to migrate the remaining ESXi hosts to the Nexus 1000V.

14. Run: copy run start.

Cisco Nexus 1000V vTracker

The vTracker provides various views that are based on the data sourced from the vCenter, the Cisco Discovery Protocol (CDP), and other related systems connected with the virtual switch. You can use vTracker to troubleshoot, monitor, and maintain the systems.

Using vTracker show commands, you can access consolidated network information across the following views:

- Upstream View–Provides information on all the virtual ports connected to an upstream physical switch. The view is from top of the network to the bottom.
- VM View–Supports two sets of data:
  - VM vNIC View–Provides information about the virtual machines (VMs) that are managed by the Cisco Nexus 1000V switch. The vNIC view is from the bottom to the top of the network.
  - VM Info View–VM Info View–Provides information about all the VMs that run on each server module.
- Module pNIC View–Provides information about the physical network interface cards (pNIC) that are connected to each Virtual Ethernet Module (VEM).
- VLAN View–Provides information about all the VMs that are connected to specific VLANs.
- vMotion View–Provides information about all the ongoing and previous VM migration events.

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network. environment. To connect SSH to the primary VSM, complete the following step:

1. From an SSH interface connected to the Cisco Nexus 1000V VSM, enter the following:

```
config t
feature vtracker
copy run start
show vtracker upstream-view
show vtracker vm-view vnic
show vtracker vm-view info
show vtracker module-view pnic
show vtracker vlan-view
copy run start
```

Figure 124 vLAN 160 check

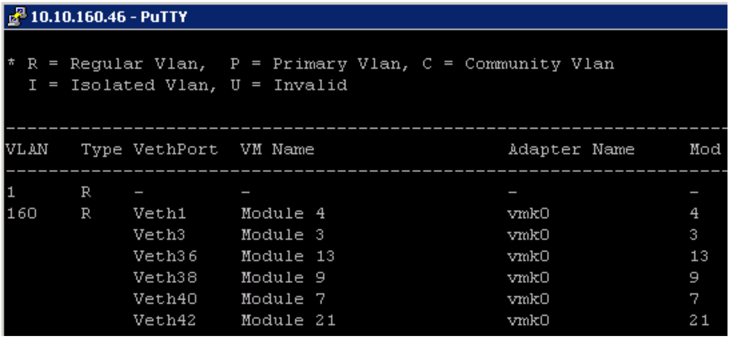

```
* R = Regular Vlan,  P = Primary Vlan, C = Community Vlan
  I = Isolated Vlan, U = Invalid

--------------------------------------------------------------------
VLAN   Type VethPort  VM Name                    Adapter Name     Mod
--------------------------------------------------------------------
1      R    -         -                          -                -
160    R    Veth1     Module 4                   vmk0             4
            Veth3     Module 3                   vmk0             3
            Veth36    Module 13                  vmk0             13
            Veth38    Module 9                   vmk0             9
            Veth40    Module 7                   vmk0             7
            Veth42    Module 21                  vmk0             21
```

Building the Virtual Machines and Environment for Workload Testing

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 the table below:

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Operating system	Microsoft Windows Server 2012 R2	Microsoft Windows Server 2012 R2
Virtual CPU amount	4	4

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Memory amount	8 GB	8 GB
Network	VMXNET3 VM-INFRA-vLAN (VSM)	VMXNET3 VM-INFRA-vLAN (VSM)
Disk-1 (OS) size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Disk-2 size and location	–	500 GB PVS-vDisk volume using CIFS
Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2012 R2	VCSA – SUSE Linux
Virtual CPU amount	4	8
Memory amount	4 GB	24 GB
Network	VMXNET3 VM-INFRA-vLAN (VSM)	VMXNET3 VM-INFRA-vLAN (VSM)
Disk size and location	40 GB Infra-DS volume	460 GB (across 11 VMDKs) Infra-DS volume
Configuration	Microsoft SQL Server Virtual Machine	
Operating system	Microsoft Windows Server 2012 R2	

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
	Microsoft SQL Server 2012 SP1	
Virtual CPU amount	4	
Memory amount	4 GB	
Network	VMXNET3 VM-INFRA-vLAN (VSM)	
Disk-1 (OS) size and location	40 GB Infra-DS volume	
Disk-2 size and location	100 GB Infra-DS volume SQL Logs	
Disk-3 size and location	150 GB Infra-DS volume SQL Databases	

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 HVD(VDI) and HSD(RDS) VMs were configured as listed in the table below:

Configuration	VDI Virtual Machines	RDS Virtual Machines
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2012 R2

Configuration	VDI Virtual Machines	RDS Virtual Machines
Virtual CPU amount	2	6
Memory amount	2.0 GB (reserved)	24 GB reserve for all guest memory
Network	VMXNET3 VDI vLAN (VSM)	VMXNET3 VDI vLAN (VSM)
Citrix PVS vDisk size and location	24 GB (thick) PVS-vDisk volume	100 GB (thick) PVS-vDisk volume
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 4.1.5 (Knowledge Worker Workload)	Microsoft Office 2016 Login VSI 4.1.5 (Knowledge Worker Workload)

Installing and Configuring XenDesktop and XenApp

This section details the installation of the core components of the XenDesktop/XenApp 7.9 system. This CVD installs two XenDesktop Delivery Controllers to support both hosted shared desktops (RDS), non-persistent virtual desktops (VDI), and persistent virtual desktops (VDI).

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.

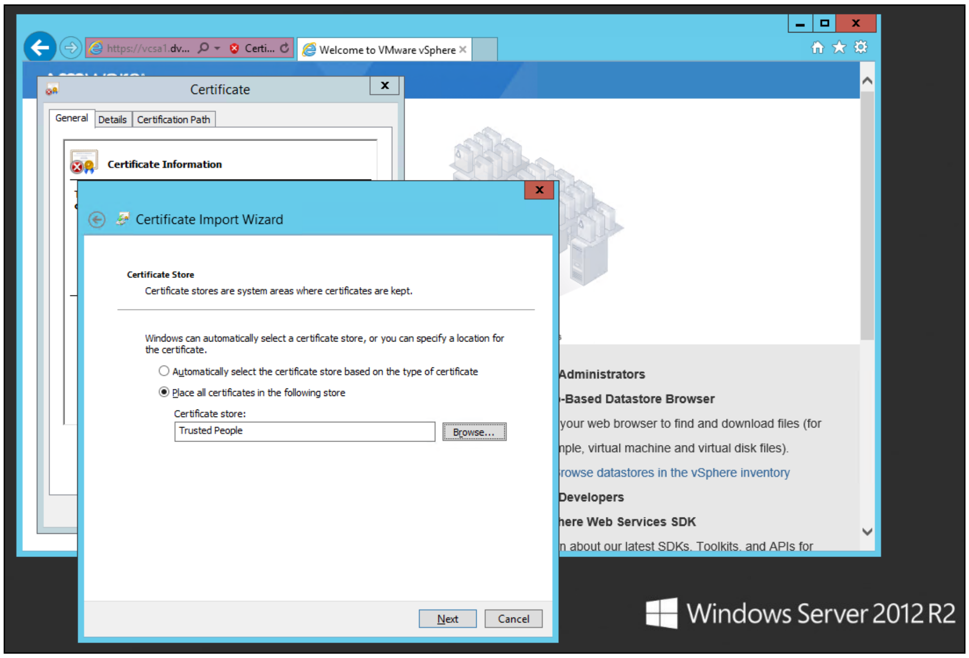

Certificate Import Wizard

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

https://vcsa1.dv...
Certi...
Welcome to VMware vSphere
Certificate
General
Details
Certification Path
Certificate Information
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.

1. To begin the installation, connect to the first XenDesktop server and launch the installer from the Citrix XenDesktop 7.9 ISO.

2. Click Start.

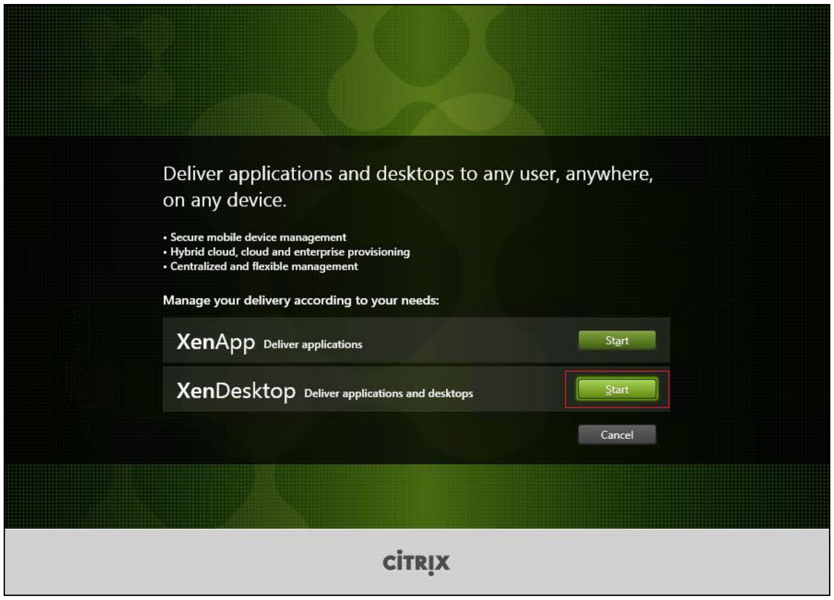

Deliver applications and desktops to any user, anywhere, on any device.

- 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

The installation wizard presents a menu with three subsections.

3. Click "Get Started - Delivery Controller."

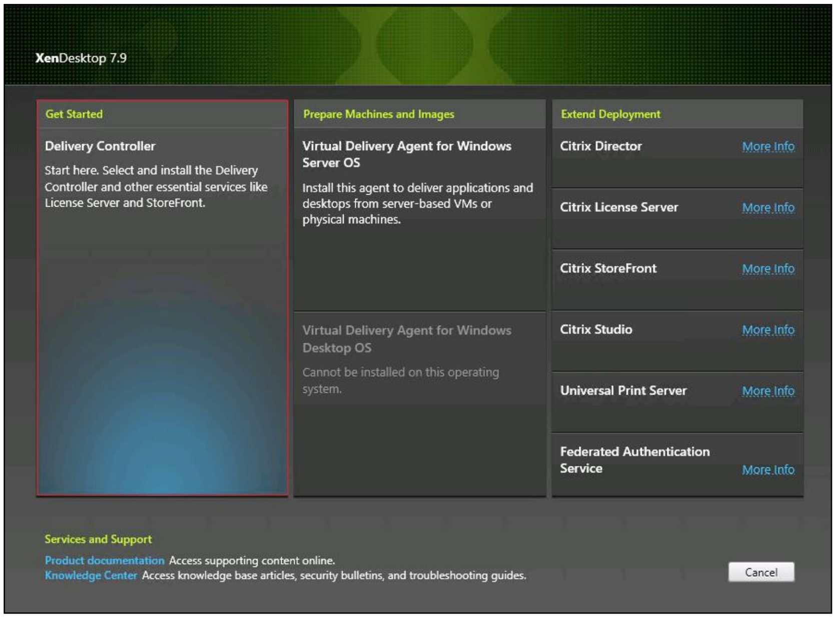

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server and StoreFront.

Virtual Delivery Agent for Windows Server OS

Install this agent to deliver applications and desktops from server-based VMs or physical machines.

Virtual Delivery Agent for Windows Desktop OS

Cannot be installed on this operating system.

Citrix Director More Info
Citrix License Server More Info
Citrix StoreFront More Info
Citrix Studio More Info
Universal Print Server More Info
Federated Authentication Service More Info
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.

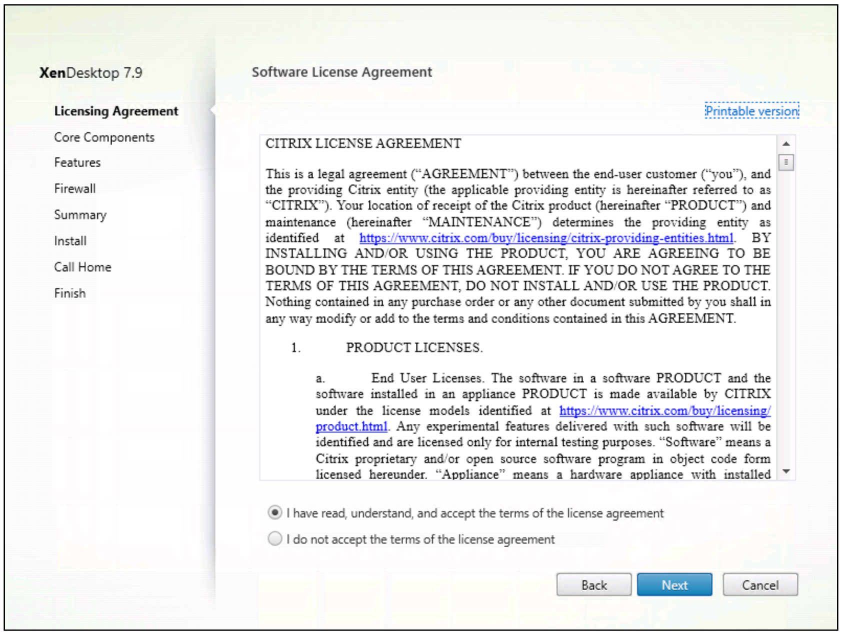

Software License Agreement

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 https://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

● 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:

   a. Delivery Controller

   a. Studio

   b. License Server

   c. StoreFront

8. Click Next.

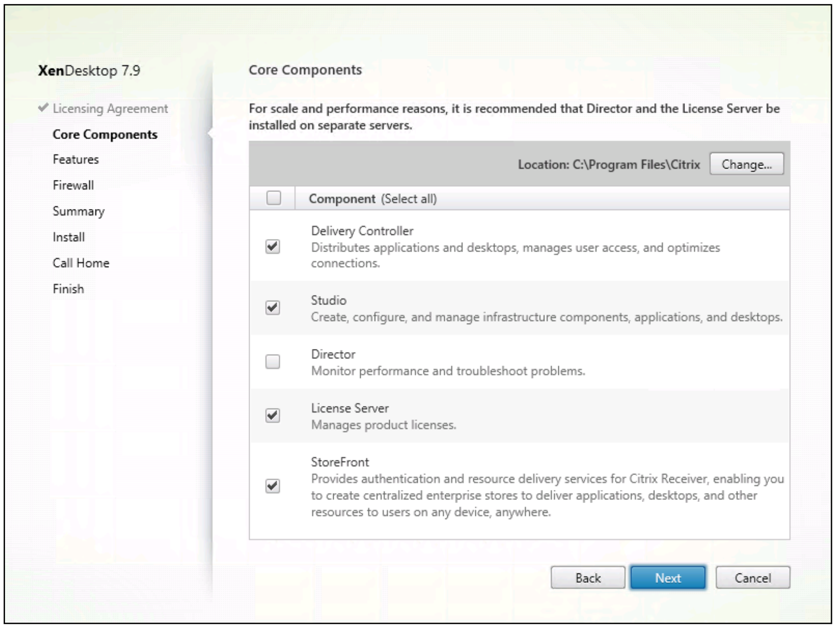

Core Components

For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.

Location: C:\Program Files\Citrix

☐	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 Manages product licenses.
☑	StoreFront 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.

Dedicated StoreFront servers should be implemented for large scale deployments.

9. Since a SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2012 SP1 Express” unchecked.

10. Click Next.

11. Select the default ports and automatically configured firewall rules.

12. Click Next.

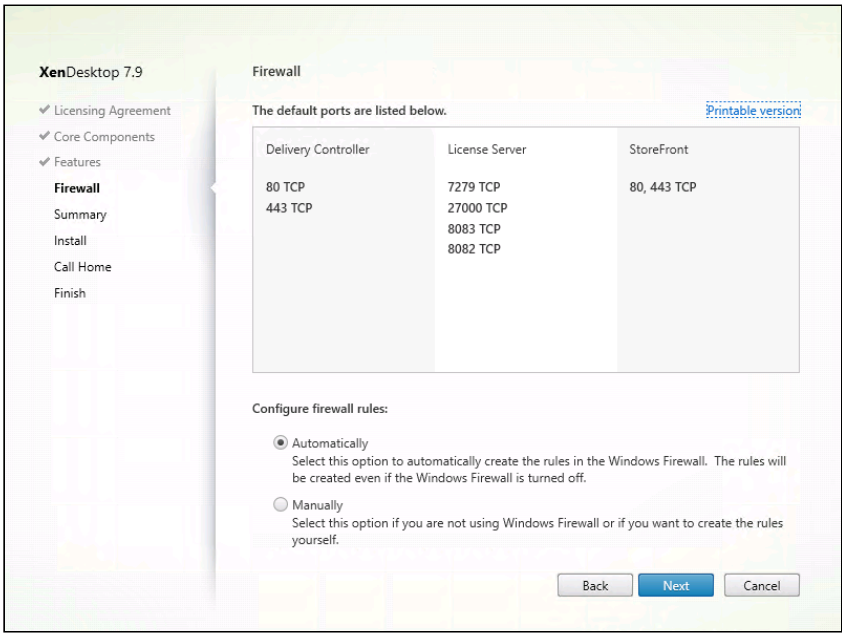

Firewall

The default ports are listed below.

Delivery Controller	License Server	StoreFront
80 TCP 443 TCP	7279 TCP 27000 TCP 8083 TCP 8082 TCP	80, 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.

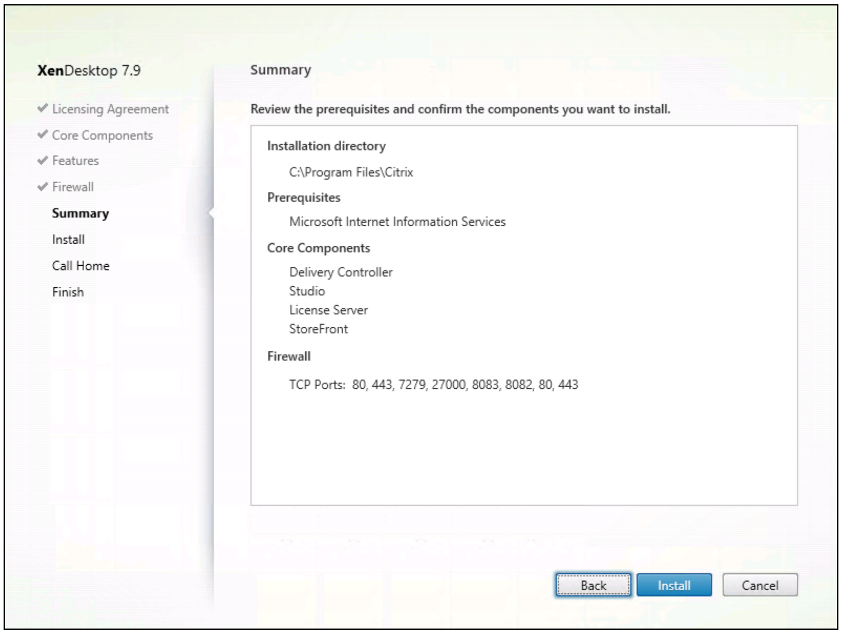

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory
C:\Program Files\Citrix

Prerequisites
Microsoft Internet Information Services

Core Components
Delivery Controller
Studio
License Server
StoreFront

Firewall
TCP Ports: 80, 443, 7279, 27000, 8083, 8082, 80, 443

14. (Optional) Click the Call Home participation.

15. Click Finish.

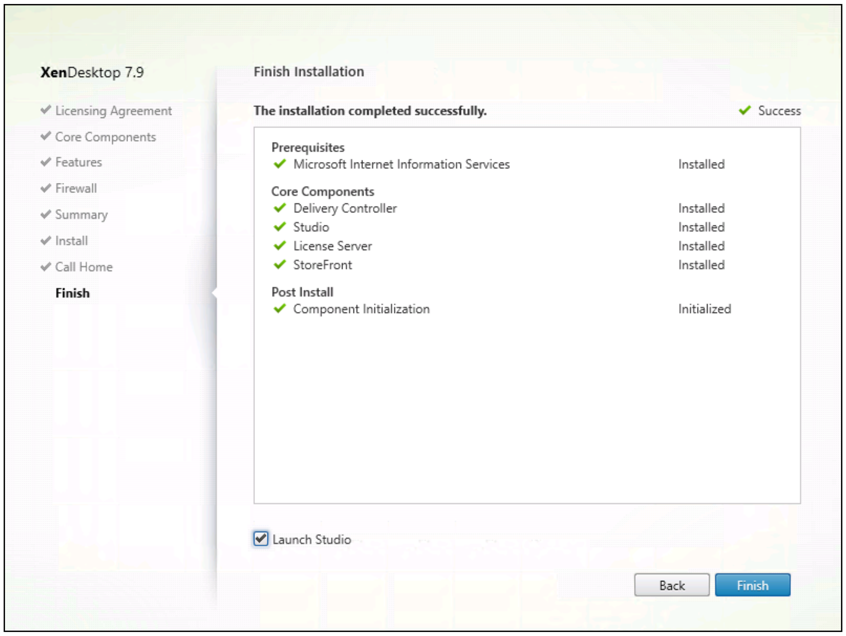

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites	
✓ Microsoft Internet Information Services	Installed
Core Components	
✓ Delivery Controller	Installed
✓ Studio	Installed
✓ License Server	Installed
✓ StoreFront	Installed
Post Install	
✓ Component Initialization	Initialized

☑ Launch Studio

16. (Optional) Check Launch Studio to launch Citrix Studio Console.

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.

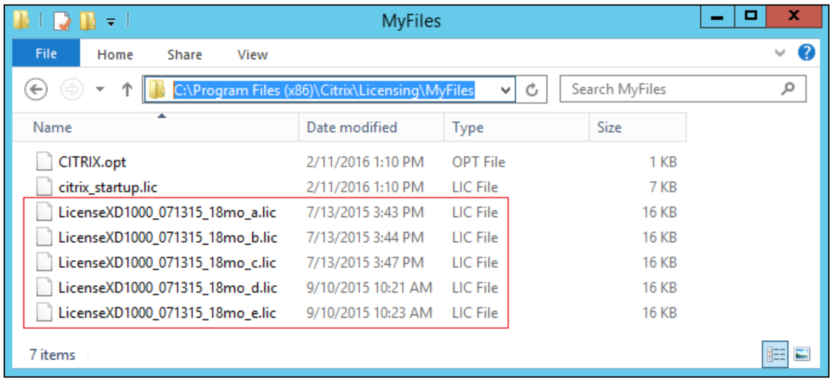

C:\Program Files (x86)\Citrix\Licensing\MyFiles

Name	Date modified	Type	Size
CITRIX.opt	2/11/2016 1:10 PM	OPT File	1 KB
citrix_startup.lic	2/11/2016 1:10 PM	LIC File	7 KB
LicenseXD1000_071315_18mo_a.lic	7/13/2015 3:43 PM	LIC File	16 KB
LicenseXD1000_071315_18mo_b.lic	7/13/2015 3:44 PM	LIC File	16 KB
LicenseXD1000_071315_18mo_c.lic	7/13/2015 3:47 PM	LIC File	16 KB
LicenseXD1000_071315_18mo_d.lic	9/10/2015 10:21 AM	LIC File	16 KB
LicenseXD1000_071315_18mo_e.lic	9/10/2015 10:23 AM	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.

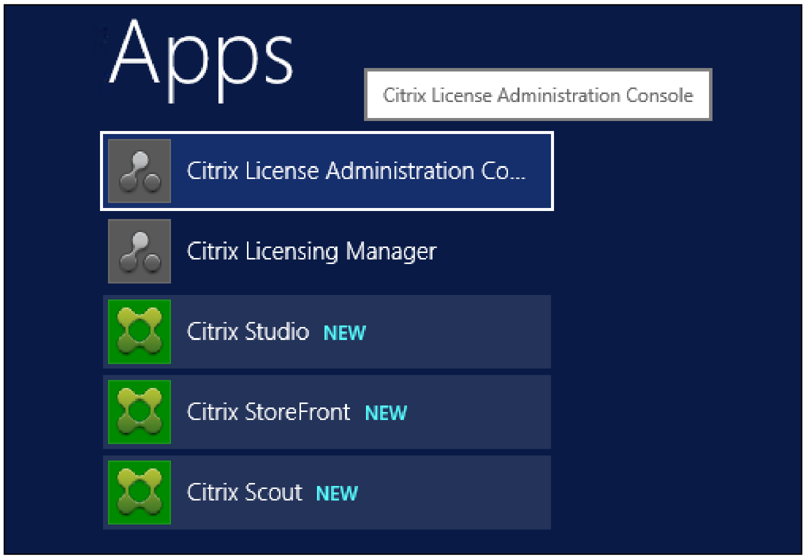

4. Confirm that the license files have been read and enabled correctly.

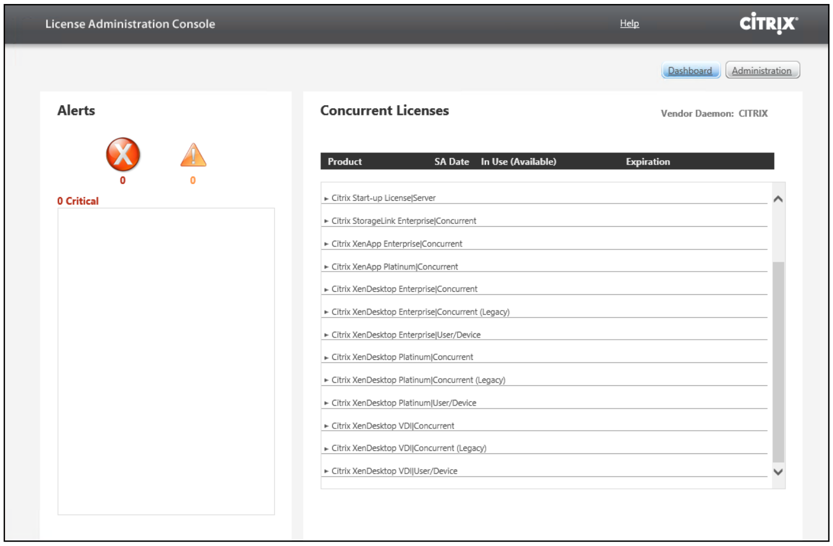

Concurrent Licenses

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			

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 7.9 environment 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.

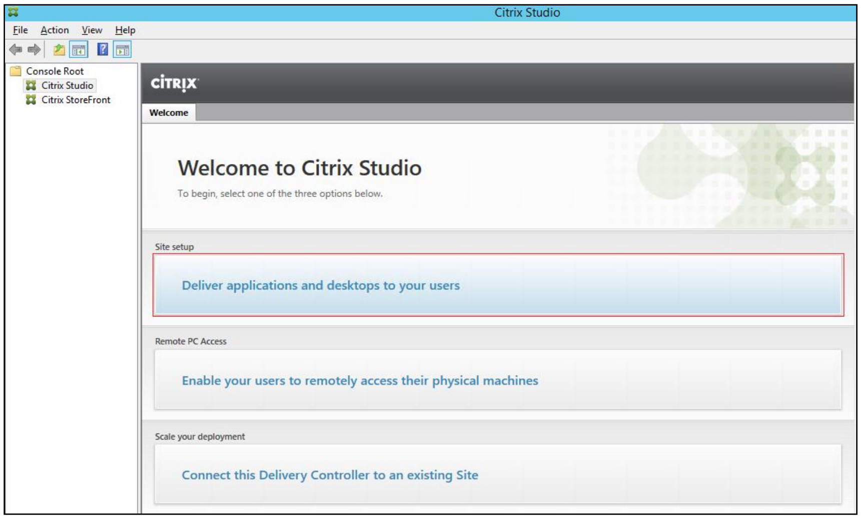

Welcome to Citrix Studio

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

2. Select the “An empty, unconfigured Site” radio button.

3. Enter a site name.

4. Click Next

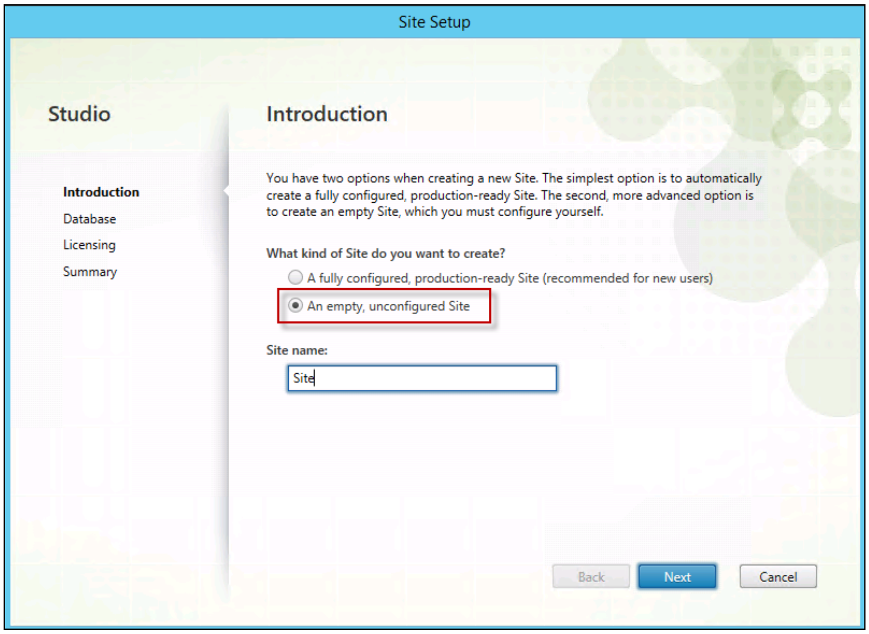

Introduction

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: Site

5. Provide the Database Server Locations for each data type and click Next.

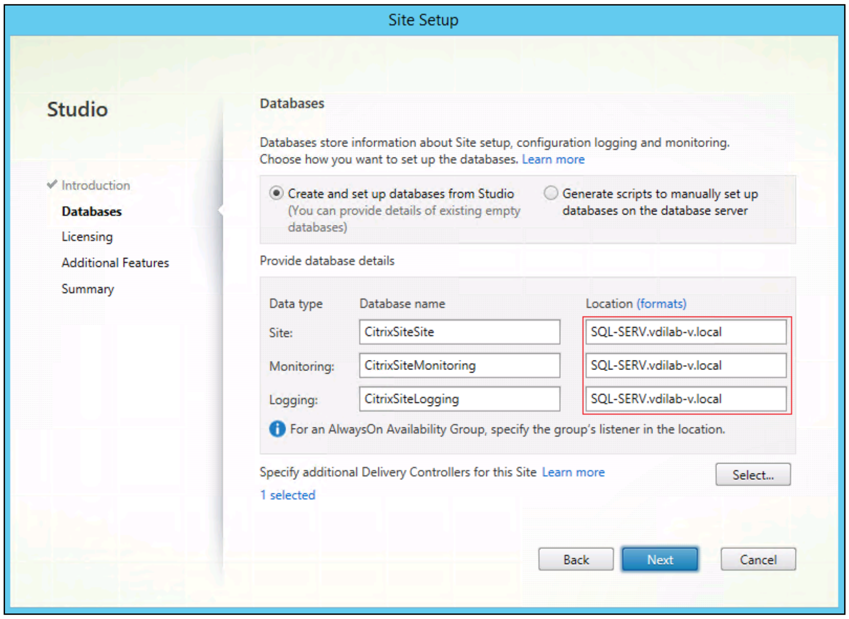

Databases

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:	CitrixSiteSite	SQL-SERV.vdilab-v.local
Monitoring:	CitrixSiteMonitoring	SQL-SERV.vdilab-v.local
Logging:	CitrixSiteLogging	SQL-SERV.vdilab-v.local

For an AlwaysOn Availability Group, specify the group's listener in the location.

Specify additional Delivery Controllers for this Site Learn more

1 selected

6. Provide the FQDN of the license server.

7. 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.

8. Select the appropriate product edition using the license radio button

9. Click Next

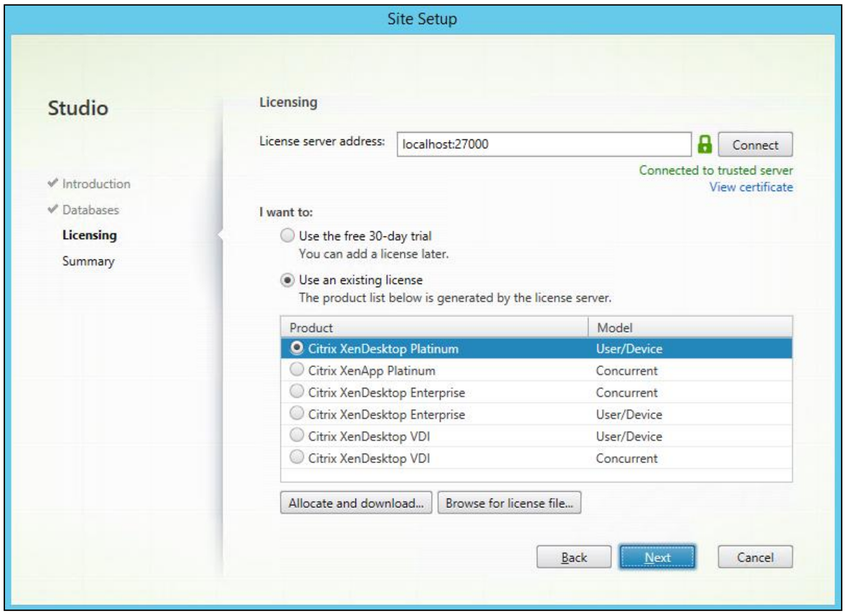

Licensing

License server address: localhost:27000

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
○ Citrix XenDesktop Enterprise	Concurrent
○ Citrix XenDesktop Enterprise	User/Device
○ Citrix XenDesktop VDI	User/Device
○ Citrix XenDesktop VDI	Concurrent

10. Click Finish to complete initial setup.

High availability will be available for the databases when added to the SQL AlwaysOn Availability Group.

11. Click Test site to determine the site creation success.

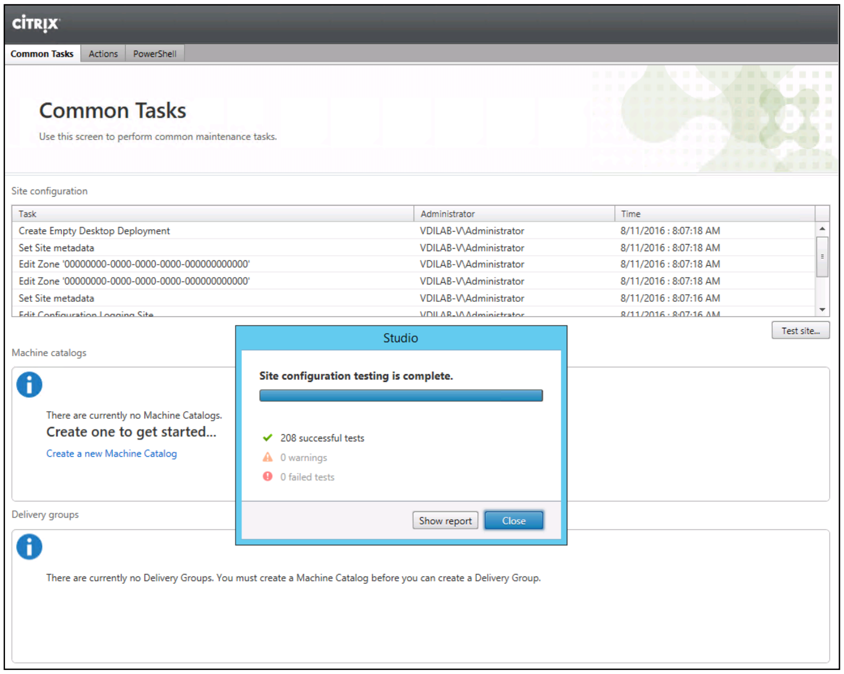

Common Tasks

Use this screen to perform common maintenance tasks.

Site configuration

Task	Administrator	Time
Create Empty Desktop Deployment	VDILAB-V\Administrator	8/11/2016 : 8:07:18 AM
Set Site metadata	VDILAB-V\Administrator	8/11/2016 : 8:07:18 AM
Edit Zone '00000000-0000-0000-0000-000000000000'	VDILAB-V\Administrator	8/11/2016 : 8:07:18 AM
Edit Zone '00000000-0000-0000-0000-000000000000'	VDILAB-V\Administrator	8/11/2016 : 8:07:18 AM
Set Site metadata	VDILAB-V\Administrator	8/11/2016 : 8:07:16 AM
Edit Configuration Logging Site	VDILAB-V\Administrator	8/11/2016 : 8:07:16 AM

Machine catalogs

There are currently no Machine Catalogs.

Create one to get started...

Create a new Machine Catalog

Delivery groups

There are currently no Delivery Groups. You must create a Machine Catalog before you can create a Delivery Group.

Site configuration testing is complete.

✓ 208 successful tests

⚠ 0 warnings

❗ 0 failed tests

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 7.9 ISO.

2. Click Start.

3. Click Delivery Controller.

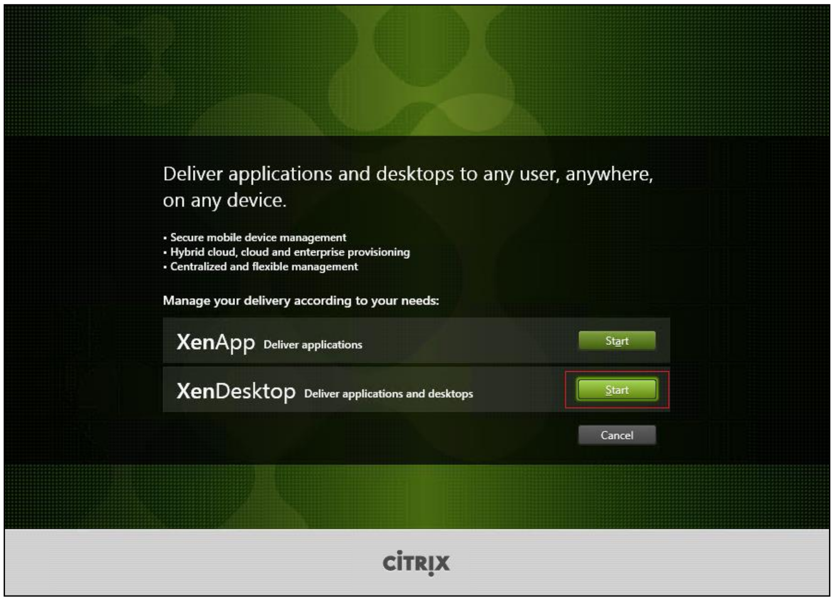

Deliver applications and desktops to any user, anywhere, on any device.

- 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

4. Select the components to be installed:

5. Delivery Controller:

   a. Studio

   b. Director

   c. StoreFront (This solution uses two dedicated StoreFront servers)

6. Click Next.

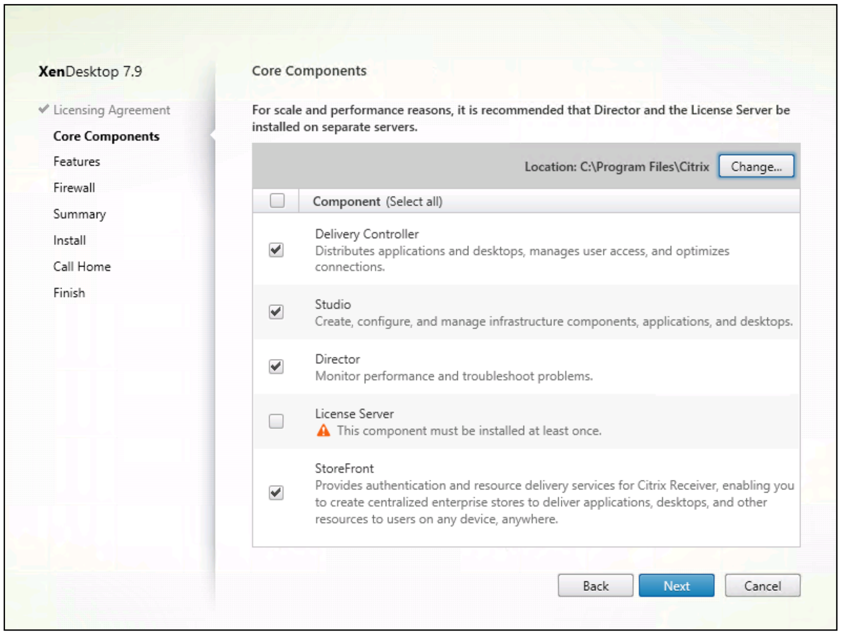

Core Components

For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.

Location: C:\Program Files\Citrix Change...

	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 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. 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.

8. Review the Summary configuration.

9. Click Install.

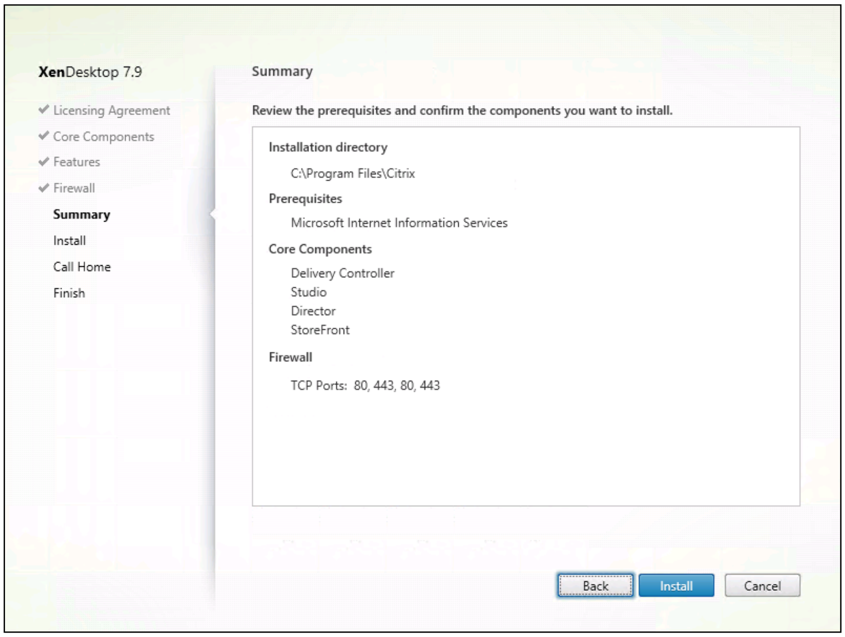

XenDesktop 7.9

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

Delivery Controller
Studio
Director
StoreFront

Firewall

TCP Ports: 80, 443, 80, 443

10. Confirm all selected components were successfully installed.

11. Verify the Launch Studio checkbox is checked.

12. Click Finish.

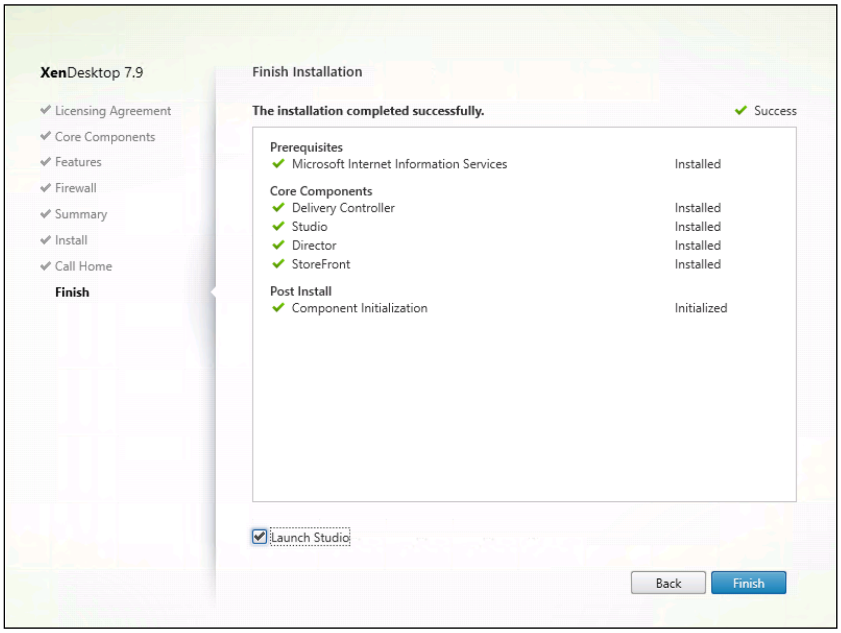

XenDesktop 7.9

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites

✓ Microsoft Internet Information Services	Installed

Core Components

✓ Delivery Controller	Installed
✓ Studio	Installed
✓ Director	Installed
✓ StoreFront	Installed

Post Install

✓ Component Initialization	Initialized

☑ Launch Studio

Add the Second Delivery Controller to the XenDesktop Site

To add the second Delivery Controller to the XenDesktop Site, complete the following steps:

1. Click the Connect this Delivery Controller to an existing Site button.

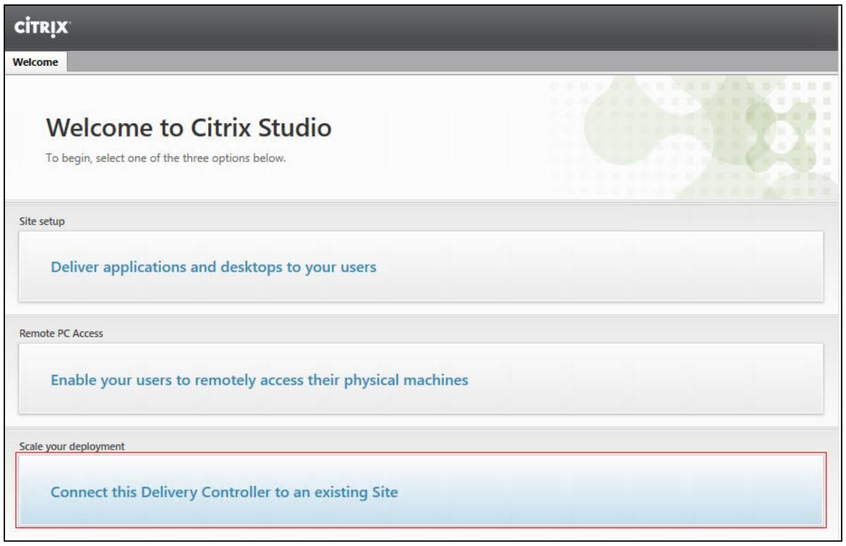

Welcome to Citrix Studio

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

2. Enter the FQDN of the first delivery controller.

3. Click OK.

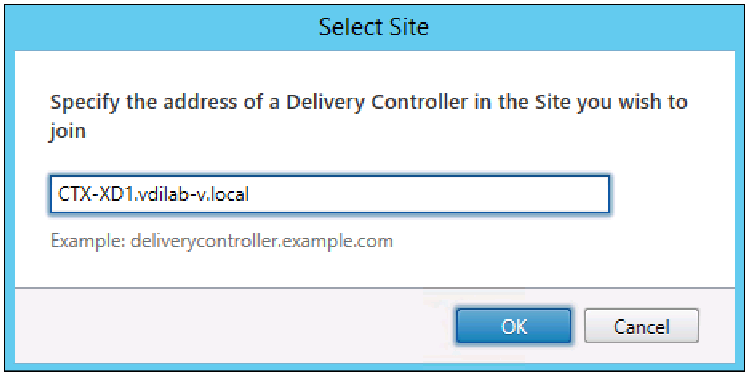

Specify the address of a Delivery Controller in the Site you wish to join

CTX-XD1.vdilab-v.local

Example: deliverycontroller.example.com

4. Click Yes to allow the database to be updated with this controller's information automatically.

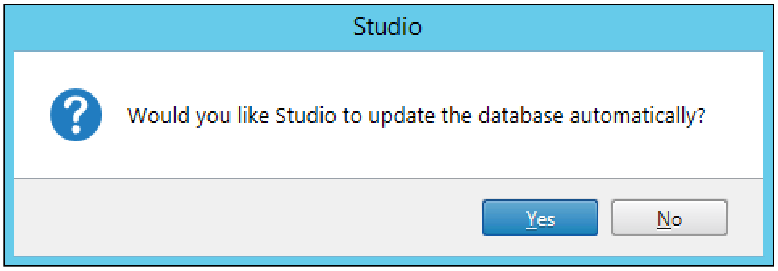

Would you like Studio to update the database automatically?

5. When complete, test the site configuration and verify the Delivery Controller has been added to the list of Controllers.

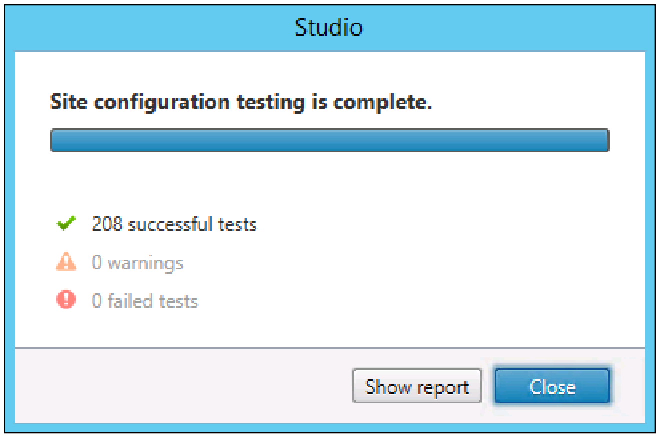

Studio

Site configuration testing is complete.

✓ 208 successful tests

⚠ 0 warnings

❗ 0 failed tests

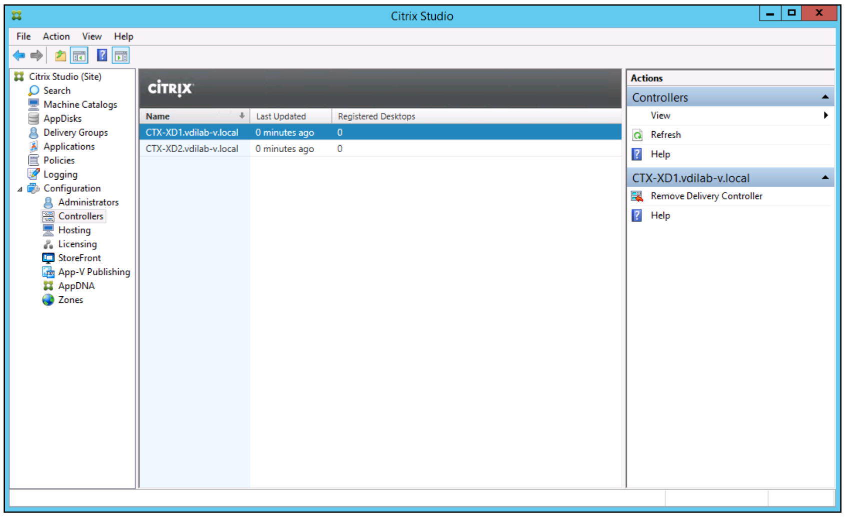

Name	Last Updated	Registered Desktops
CTX-XD1.vdilab-v.local	0 minutes ago	0
CTX-XD2.vdilab-v.local	0 minutes ago	0

Create Host Connections with Citrix Studio

Citrix Studio provides wizards to guide the process of setting up an environment and creating desktops. To set up a host connection for a cluster of VMs for the HSD and VDI desktops, complete the following steps:

> The instructions below outline the procedure to add a host connection and resources for HSD and VDI desktops.

1. Connect to the XenDesktop server and launch Citrix Studio.

2. From the Configuration menu, right-click Hosting and select Add Connection and Resources.

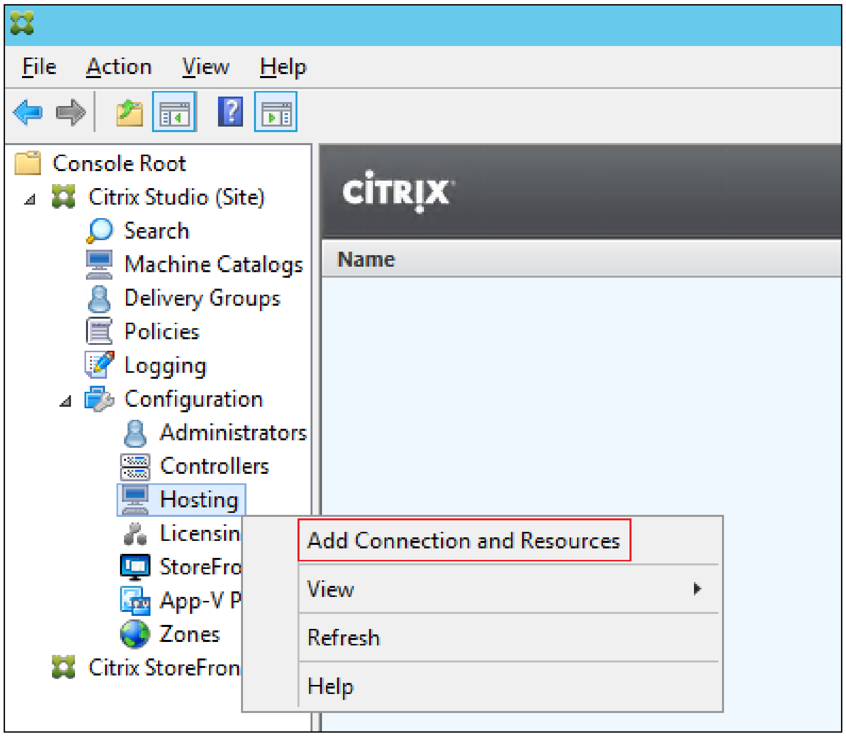

3. Select the Host Type of VMware vSphere®.
4. Enter the FQDN of the vCenter server.
5. Enter the username (in domain\username format) for the vSphere account.
6. Provide the password for the vSphere account.
7. Provide a connection name.
8. Select the Other tools radio button since Provisioning Services will be used.
9. Click Next.

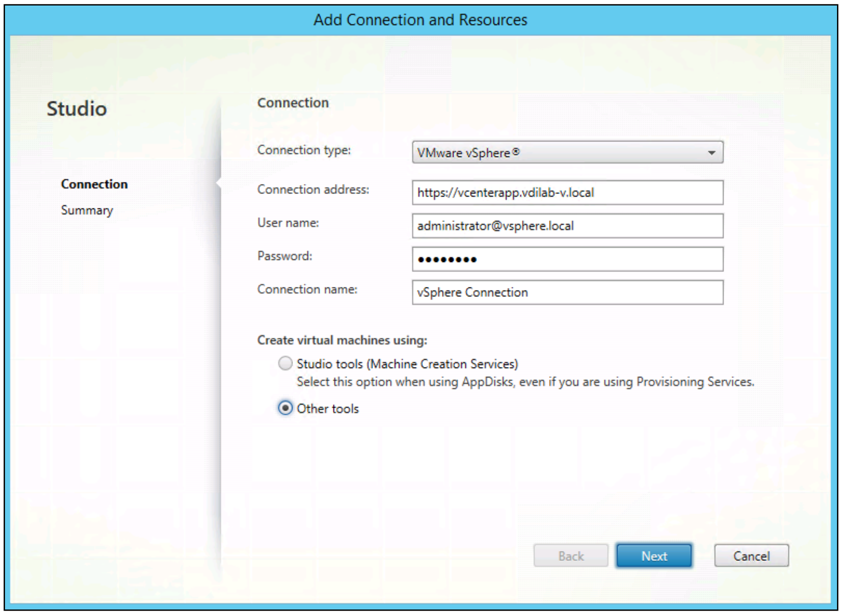

Connection

Connection type: VMware vSphere®

Connection address: https://vcenterapp.vdilab-v.local

User name: administrator@vsphere.local

Password: ••••••••

Connection name: vSphere Connection

Create virtual machines using:

○ Studio tools (Machine Creation Services)
Select this option when using AppDisks, even if you are using Provisioning Services.

● Other tools

10. Review the Summary.

11. Click Finish

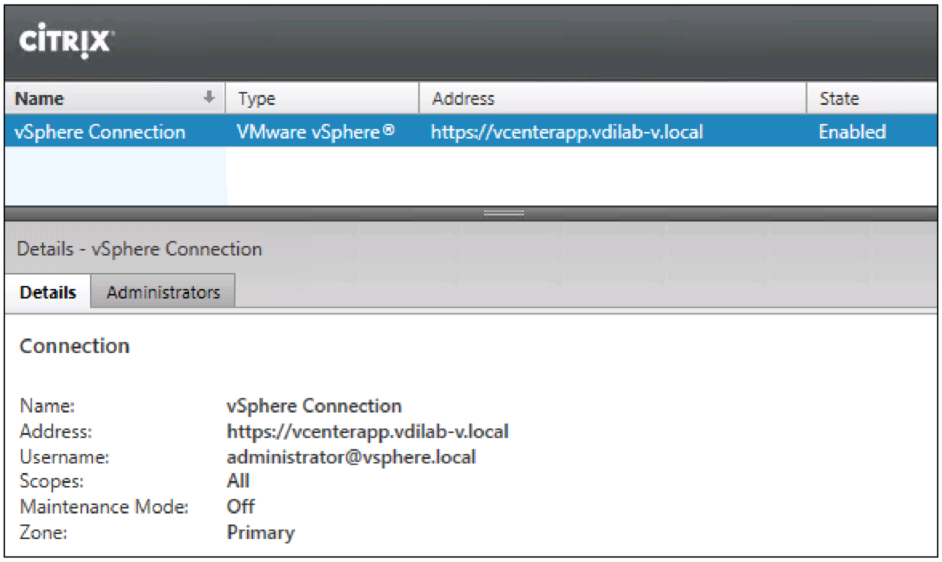

Name	Type	Address	State
vSphere Connection	VMware vSphere®	https://vcenterapp.vdilab-v.local	Enabled

Details - vSphere Connection

Connection

Name:	vSphere Connection
Address:	https://vcenterapp.vdilab-v.local
Username:	administrator@vsphere.local
Scopes:	All
Maintenance Mode:	Off
Zone:	Primary

Configuring StoreFront

Citrix StoreFront stores aggregate desktops and applications from XenDesktop sites, making resources readily available to users. In this CVD, StoreFront is installed on the Delivery Controllers virtual machine as part of the initial Delivery Controller installation. Most of the StoreFront configuration is automatically done as part of the installer. To finalize the StoreFront configuration log into the second Delivery Controller and launch the StoreFront Console.

To configure StoreFront, complete the following steps:

1. From the StoreFront Console on the second server select “Join existing server group”.

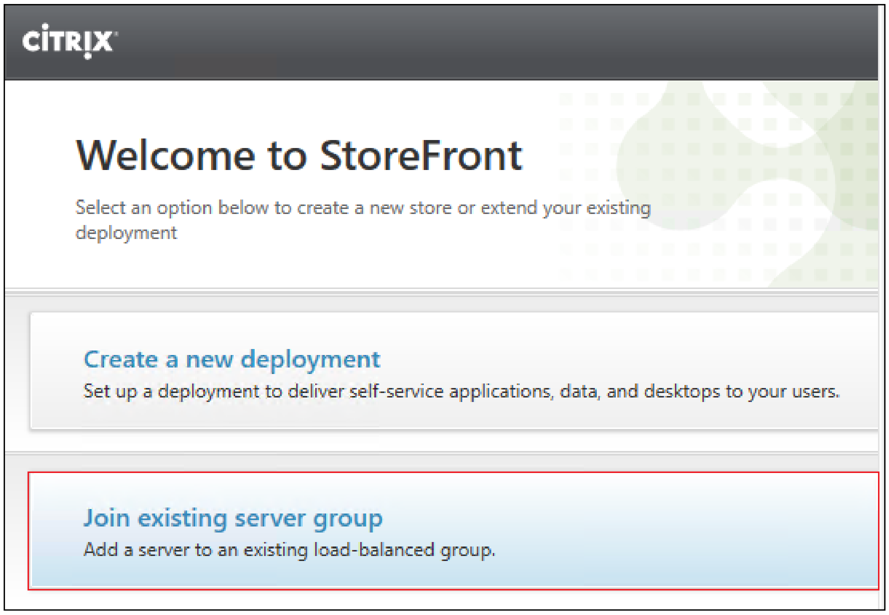

Welcome to StoreFront

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.

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.

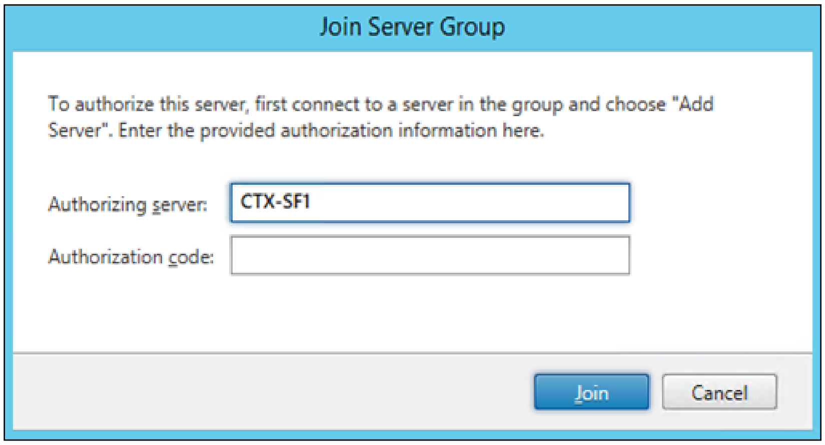

Join Server Group

To authorize this server, first connect to a server in the group and choose "Add Server". Enter the provided authorization information here.

Authorizing server: CTX-SF1

Authorization code:

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.

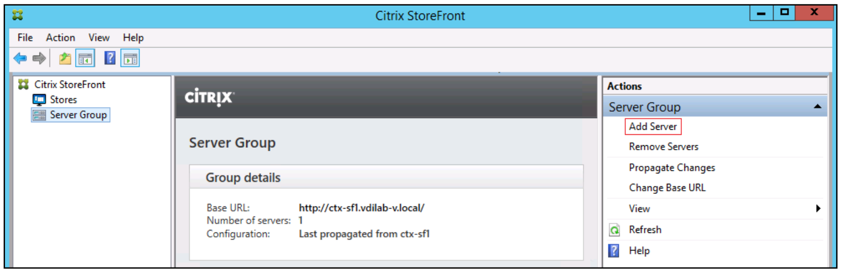

Server Group

Group details

Base URL: http://ctx-sf1.vdilab-v.local/
Number of servers: 1
Configuration: Last propagated from ctx-sf1

8. Copy the Authorization code from the Add Server dialog.

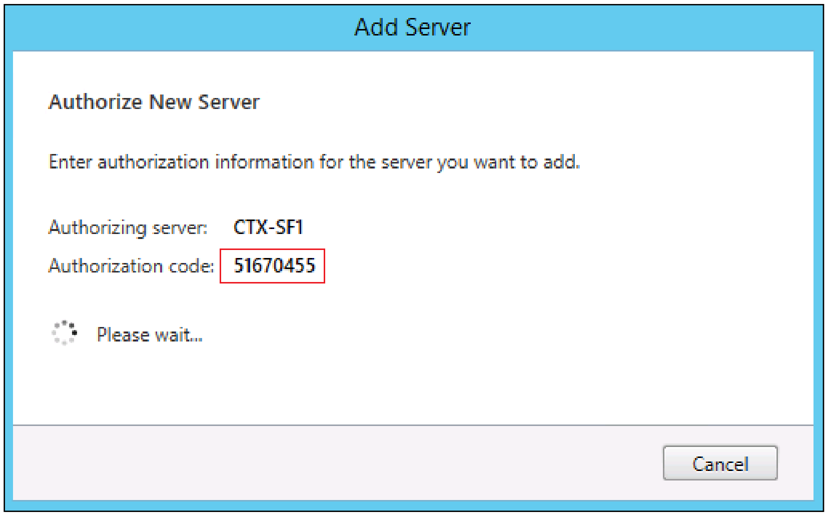

Authorize New Server

Enter authorization information for the server you want to add.

Authorizing server: CTX-SF1
Authorization code: 51670455

Please wait...

9. Connect to the second Storefront server and paste the Authorization code into the Join Server Group dialog.

10. Click Join.

11. A message appears when the second server has joined successfully.

12. Click OK.

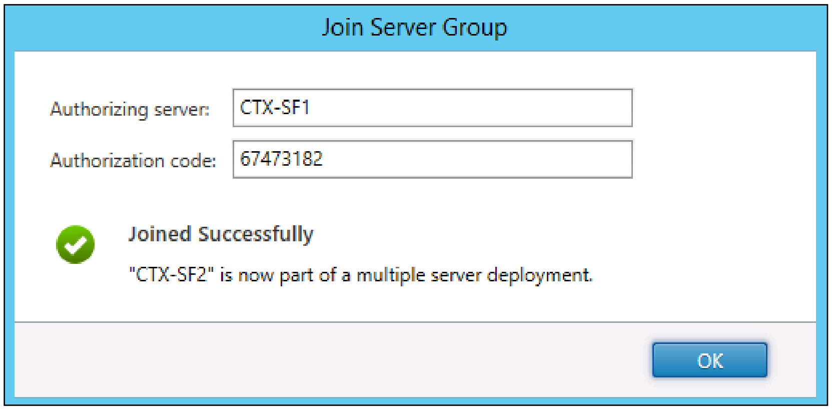

Authorizing server: CTX-SF1
Authorization code: 67473182

Joined Successfully

"CTX-SF2" is now part of a multiple server deployment.

13. The Server Group now lists both StoreFront servers in the group.

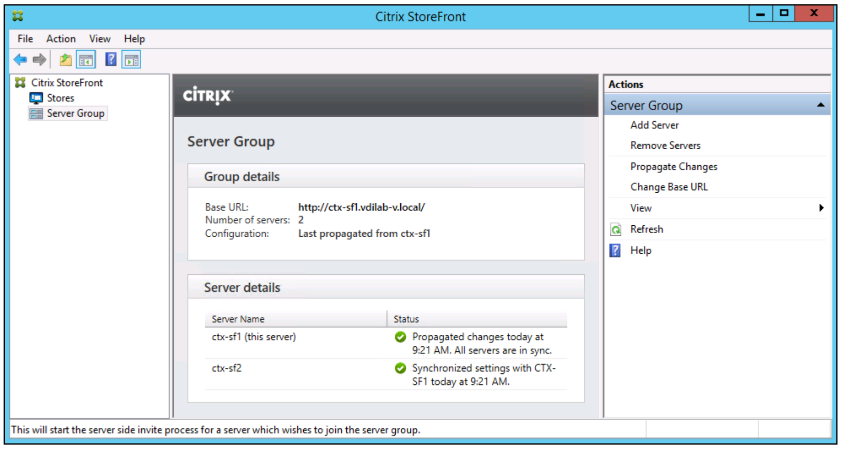

Server Group

Group details

Base URL: http://ctx-sf1.vdilab-v.local/
Number of servers: 2
Configuration: Last propagated from ctx-sf1

Server details

Server Name	Status
ctx-sf1 (this server)	Propagated changes today at 9:21 AM. All servers are in sync.
ctx-sf2	Synchronized settings with CTX-SF1 today at 9:21 AM.

This will start the server side invite process for a server which wishes to join the server group.

Installing and Configuring Citrix Provisioning Server 7.9

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 at: http://docs.citrix.com/en-us/provisioning/7-7.html

Prerequisites

Set the following Scope Options on the DHCP server hosting the PVS target machines (for example, VDI, RDS):

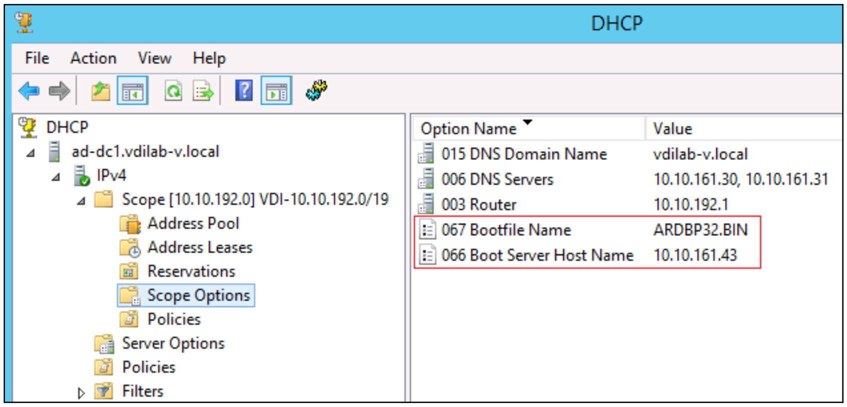

Option Name	Value
015 DNS Domain Name	vdilab-v.local
006 DNS Servers	10.10.161.30, 10.10.161.31
003 Router	10.10.192.1
067 Bootfile Name	ARDBP32.BIN
066 Boot Server Host Name	10.10.161.43

As a Citrix best-practice cited in this CTX article, apply the following registry setting on 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 MS SQL 2008, MS SQL 2008 R2, MS SQL 2012, MS SQL 2012 R2 and MS SQL 2014 Server (32 or 64-bit editions) databases can be used for the Provisioning Services database: SQL Server Express Edition, SQL Server Workgroup Edition, SQL Server Standard Edition, SQL Server Enterprise Edition. Microsoft SQL 2012 R2 was installed separately for this CVD.

To install and configure Citrix Provisioning Service 7.9, complete the following steps:

1. Insert the Citrix Provisioning Services 7.9 ISO and let AutoRun launch the installer.

2. Click the Console Installation button.

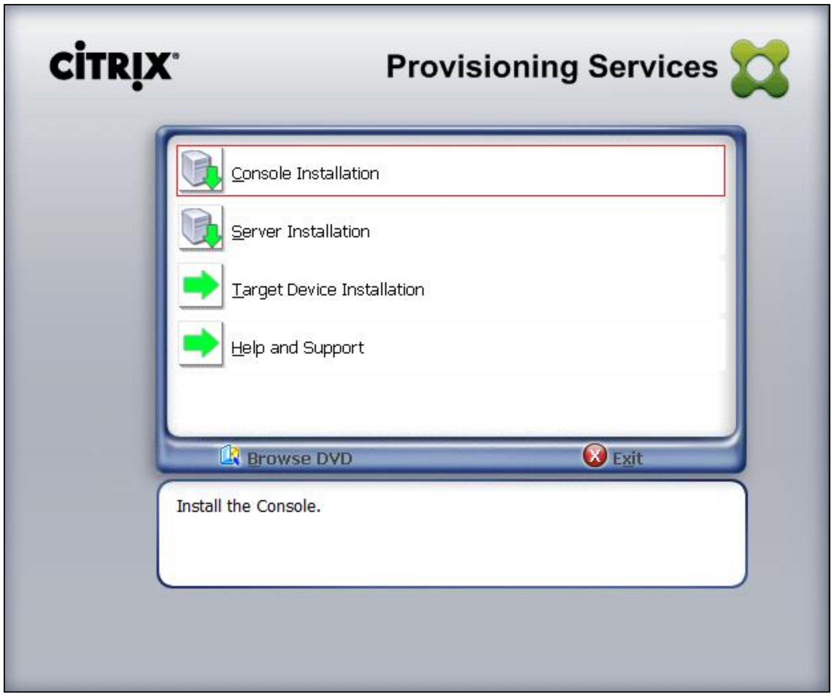

Install the Console.

3. Click Install to install the required prerequisites.

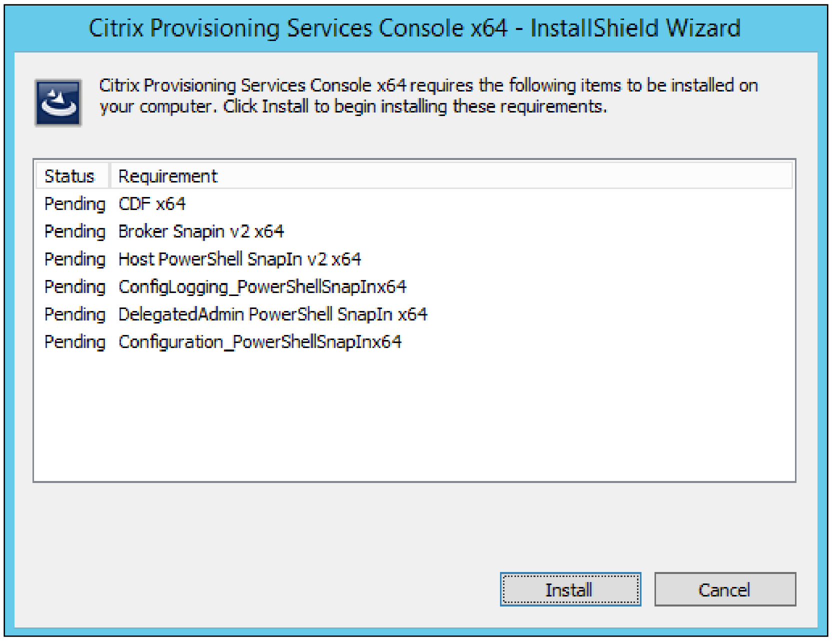
Citrix Provisioning Services Console x64 - InstallShield Wizard
Install
Cancel

Citrix 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.

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

8. Optionally provide User Name and Organization.

9. Click Next.

10. Accept the default path.

11. Click Next.

12. Click Install to start the console installation.

13. From the main installation screen, select Server Installation.

14. The installation wizard will check to resolve dependencies and then begin the PVS server installation process.

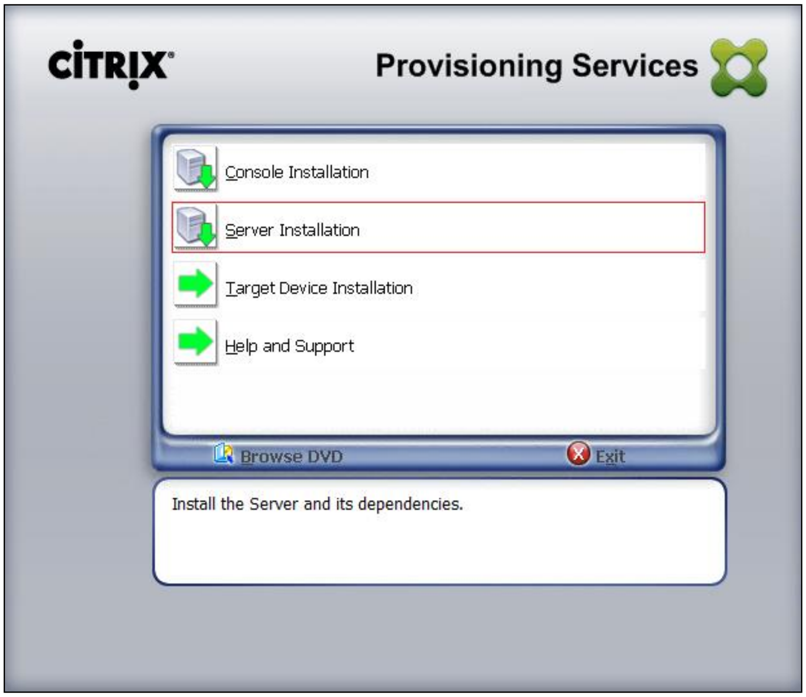

Install the Server and its dependencies.

15. Click Install on the prerequisites dialog.

16. Click Yes when prompted to install the SQL Native Client.

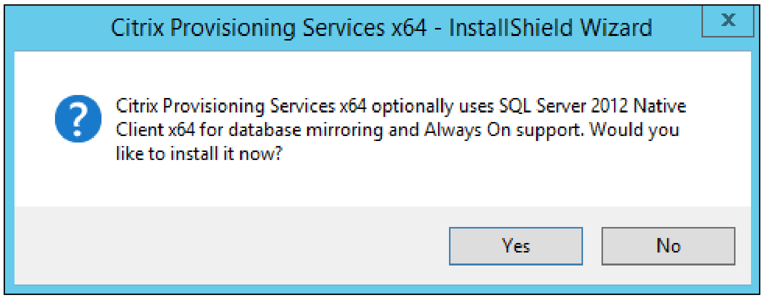

Citrix Provisioning Services x64 optionally uses SQL Server 2012 Native Client x64 for database mirroring and Always On support. Would you like to install it now?

17. Click Next when the Installation wizard starts.

18. Review the license agreement terms.

19. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”

20. Click Next.

21. Provide User Name, and Organization information. Select who will see the application.

22. Click Next.

23. Accept the default installation location.

24. Click Next.

25. Click Install to begin the installation.

26. Click Finish when the install is complete.

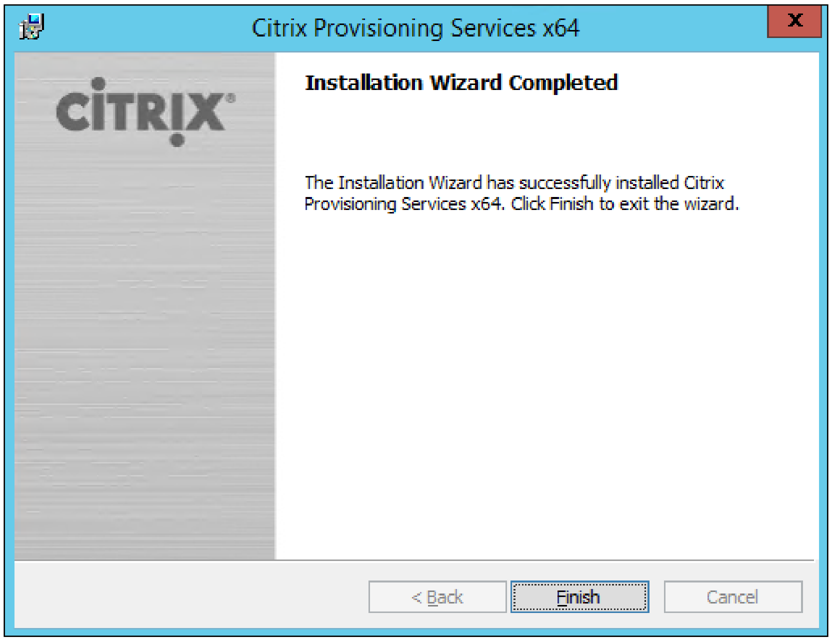

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix Provisioning Services x64. Click Finish to exit the wizard.

27. The PVS Configuration Wizard starts automatically.

28. Click Next.

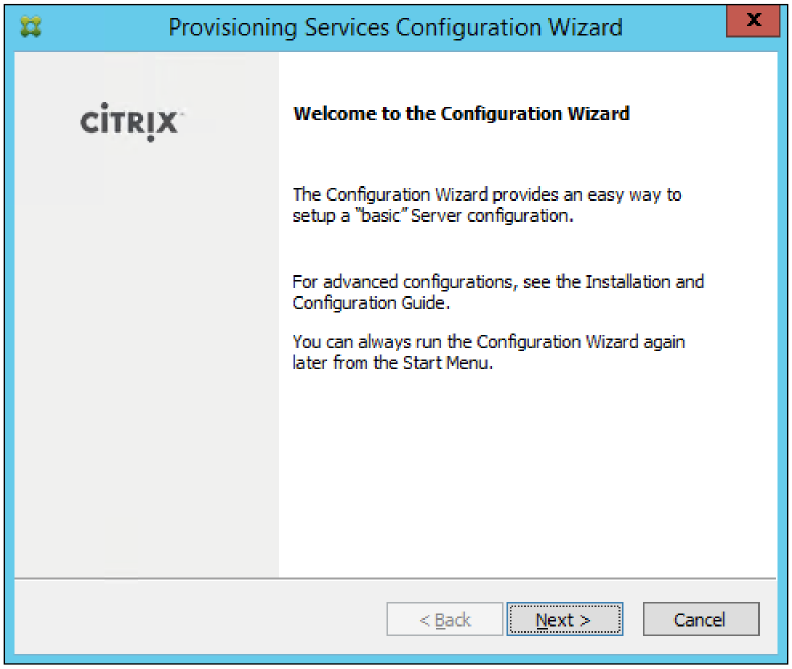

Welcome to the Configuration Wizard

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.

29. Since the PVS server is not the DHCP server for the environment, select the radio button labeled, “The service that runs on another computer.”

30. Click Next.

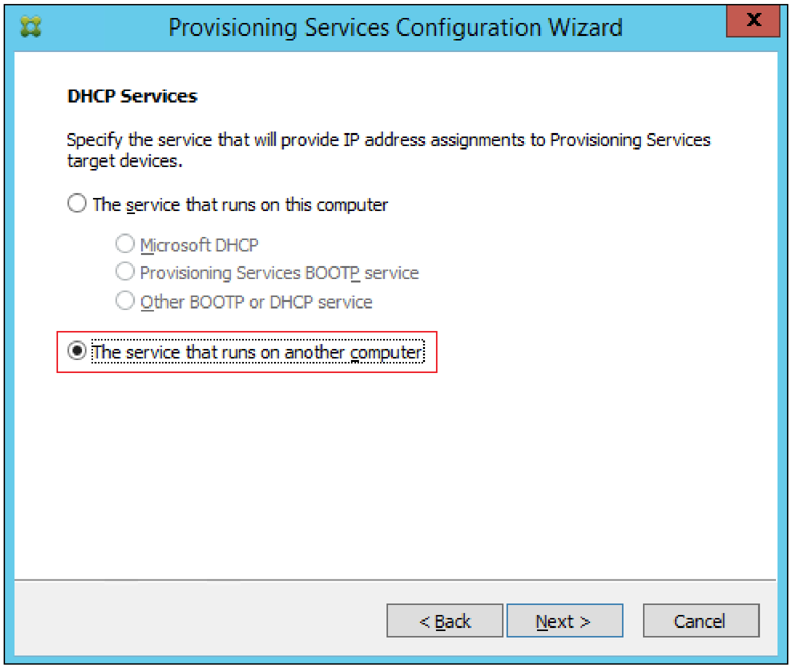

Provisioning Services Configuration Wizard

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

31. Since DHCP boot options 66 and 67 are used for TFTP services, select the radio button labeled, “The service that runs on another computer.”

32. Click Next.

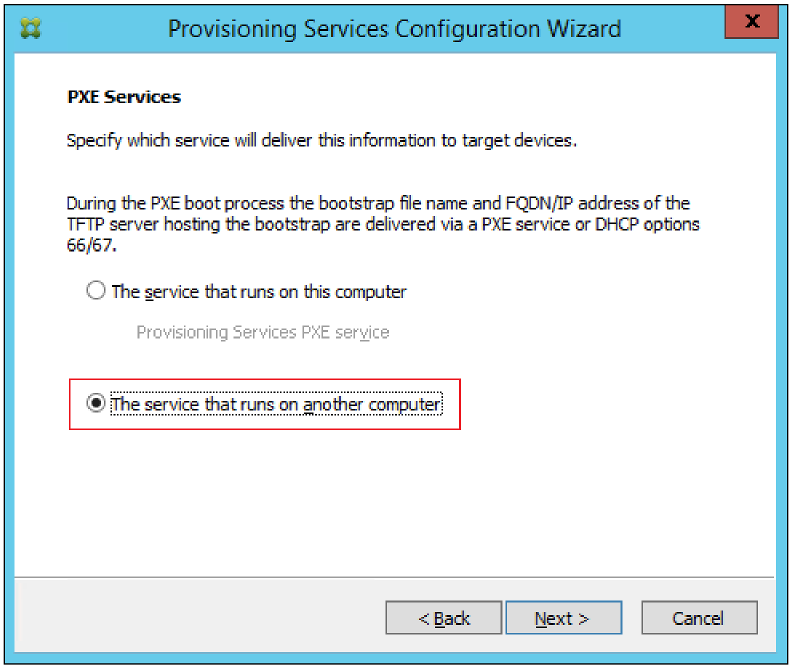

Provisioning Services Configuration Wizard

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.

○ The service that runs on this computer

Provisioning Services PXE service

● The service that runs on another computer

33. Since this is the first server in the farm, select the radio button labeled, "Create farm."

34. Click Next

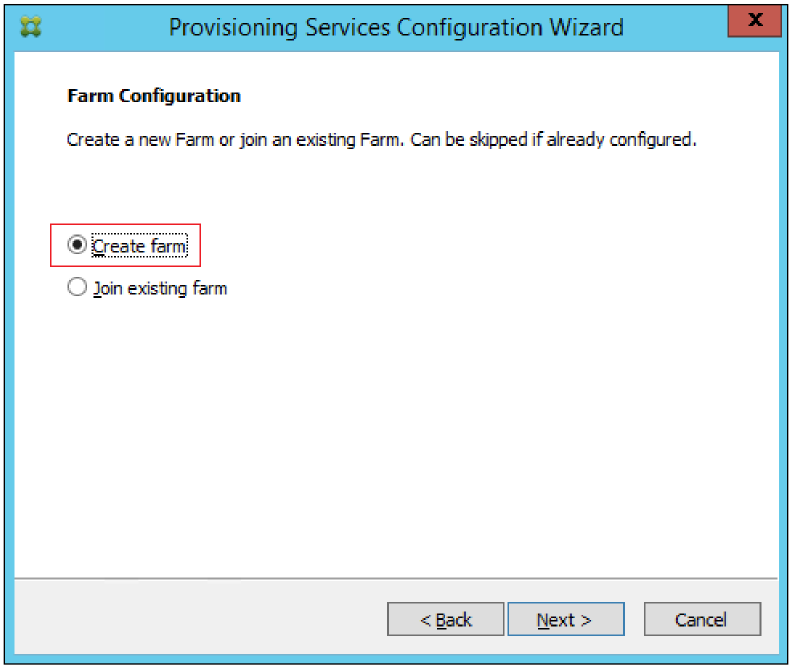

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

● Create farm

○ Join existing farm

35. Enter the FQDN of the SQL server.

36. Click Next.

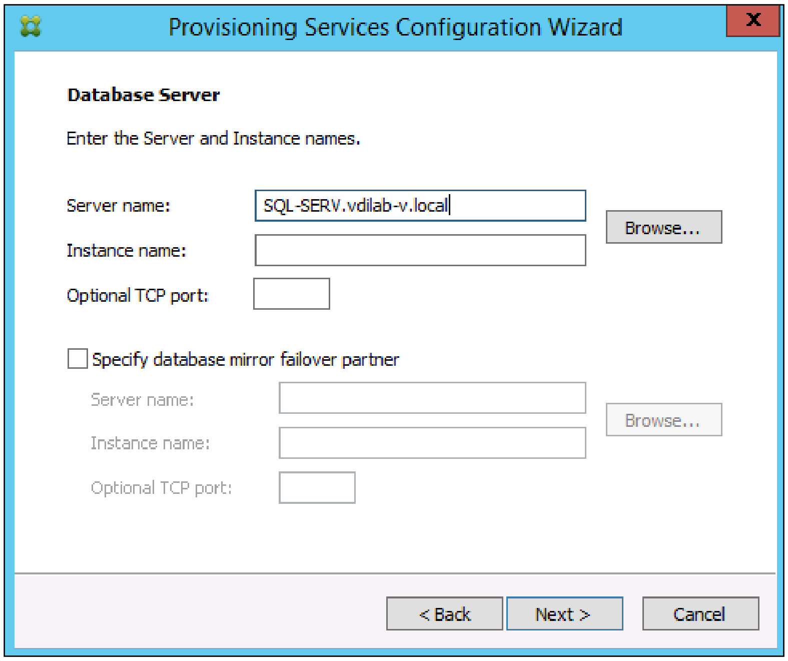

Database Server

Enter the Server and Instance names.

Server name: SQL-SERV.vdilab-v.local

Instance name:

Optional TCP port:

□ Specify database mirror failover partner

Server name:

Instance name:

Optional TCP port:

37. Provide the Database, Farm, Site, and Collection names.

38. Click Next.

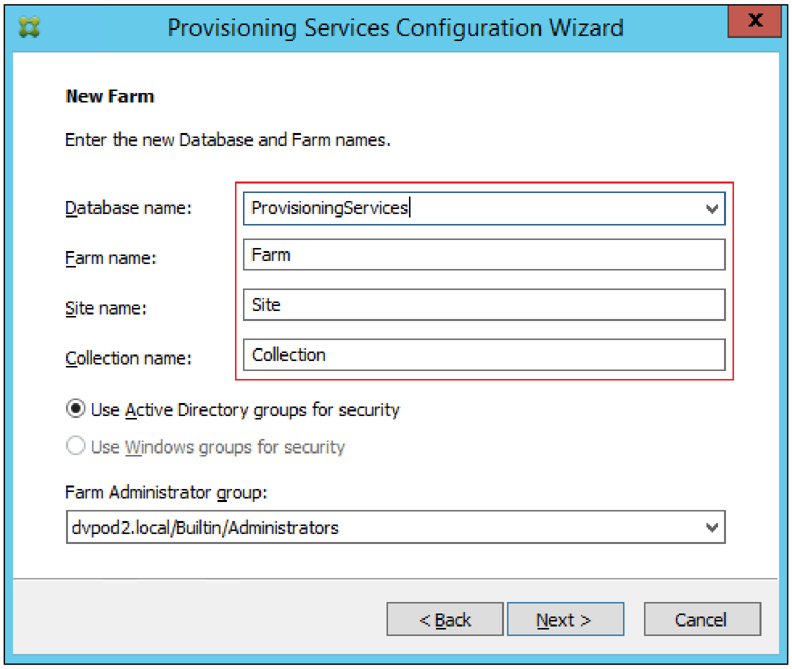

Provisioning Services Configuration Wizard

New Farm

Enter the new Database and Farm names.

Database name: ProvisioningServices

Farm name: Farm

Site name: Site

Collection name: Collection

● Use Active Directory groups for security

○ Use Windows groups for security

Farm Administrator group: dvpod2.local/Builtin/Administrators

39. Provide a vDisk Store name and the storage path to the Pure Storage vDisk share.

Create the share using support for CIFS/SMB3.

40. Click Next.

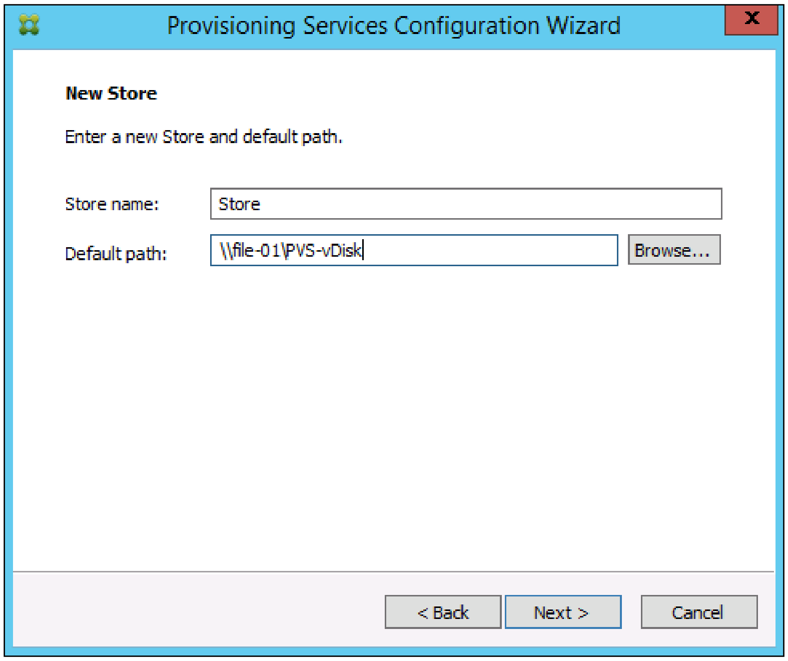

Provisioning Services Configuration Wizard

New Store

Enter a new Store and default path.

Store name: Store

Default path: \\file-01\PVS-vDisk

41. Provide the FQDN of the license server.

42. Optionally, provide a port number if changed on the license server.

43. Click Next.

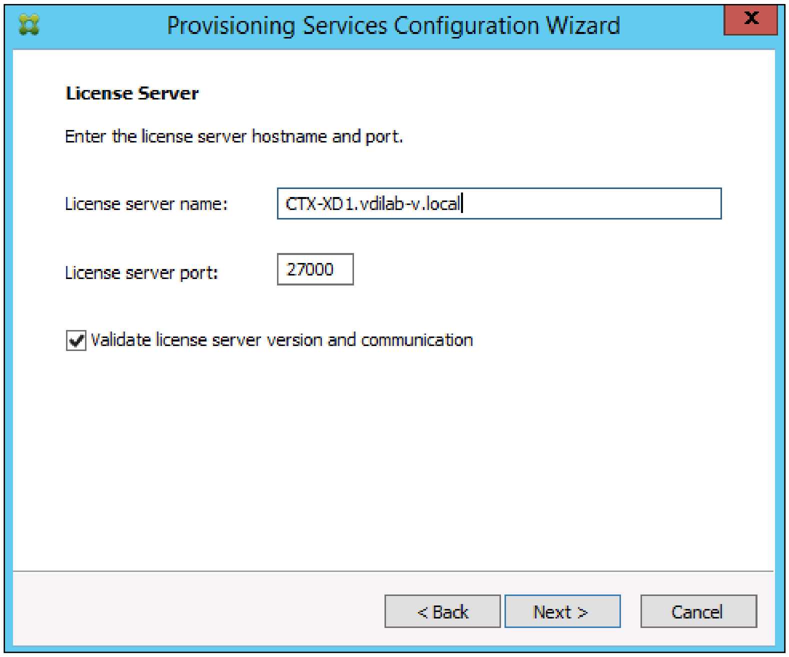

Provisioning Services Configuration Wizard

License Server

Enter the license server hostname and port.

License server name: CTX-XD1.vdilab-v.local

License server port: 27000

☑ Validate license server version and communication

44. 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.

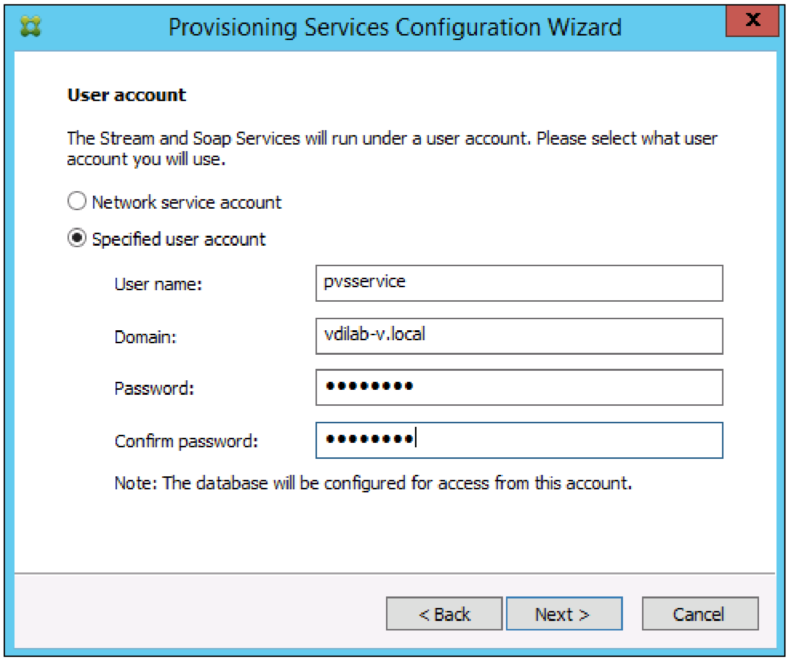

User account

The Stream and Soap Services will run under a user account. Please select what user account you will use.

○ Network service account

● Specified user account

User name: pvsservice

Domain: vdilab-v.local

Password: ••••••••

Confirm password: ••••••••

Note: The database will be configured for access from this account.

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.

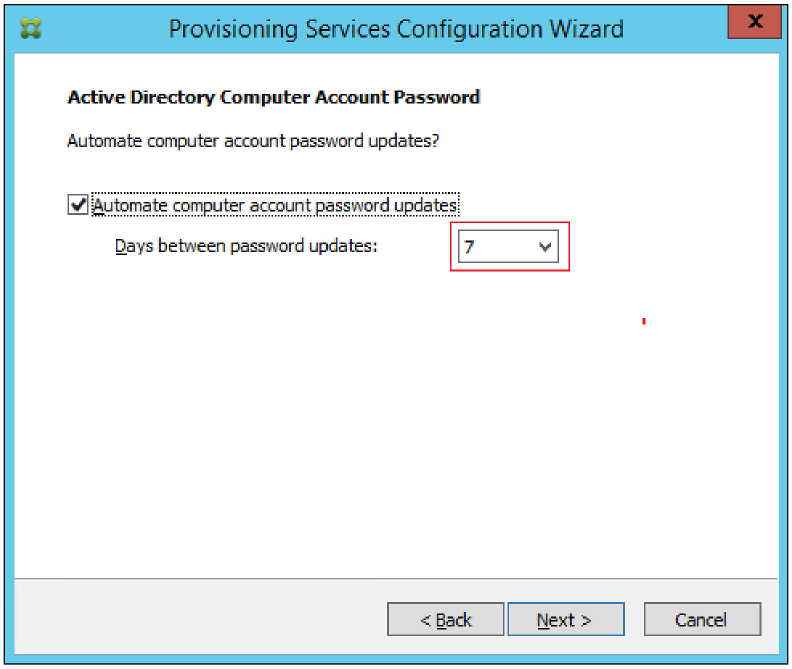

Active Directory Computer Account Password

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.

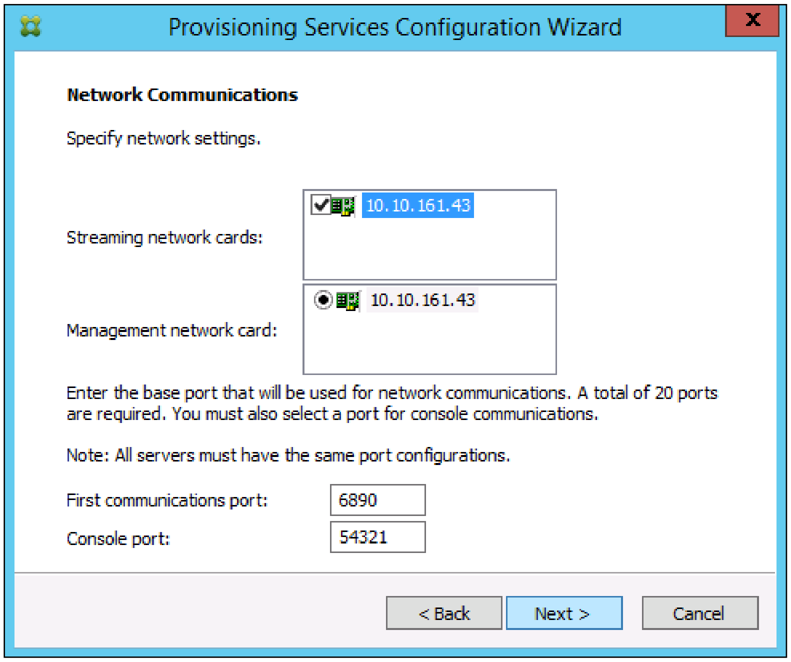

Provisioning Services Configuration Wizard

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.10.161.43

Management network card: ● 10.10.161.43

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.

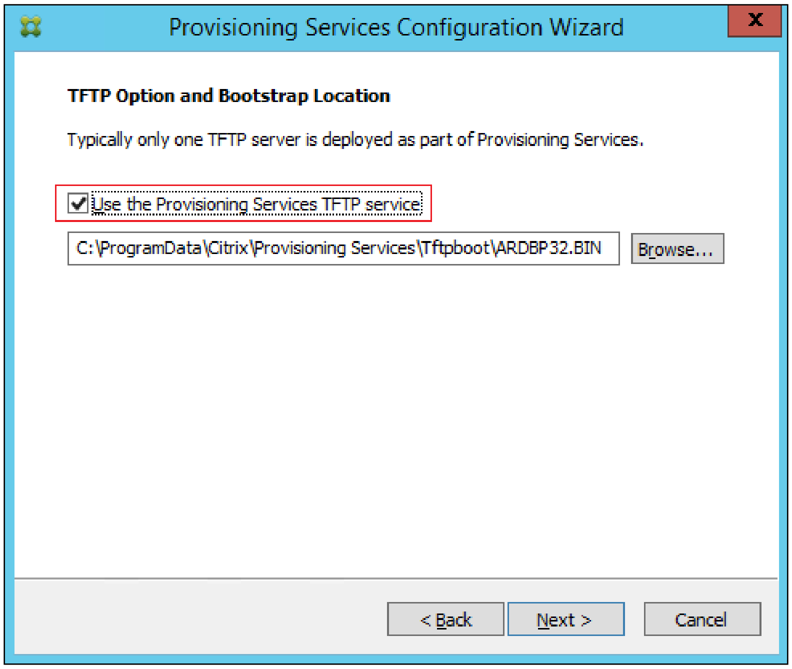

Provisioning Services Configuration Wizard

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.

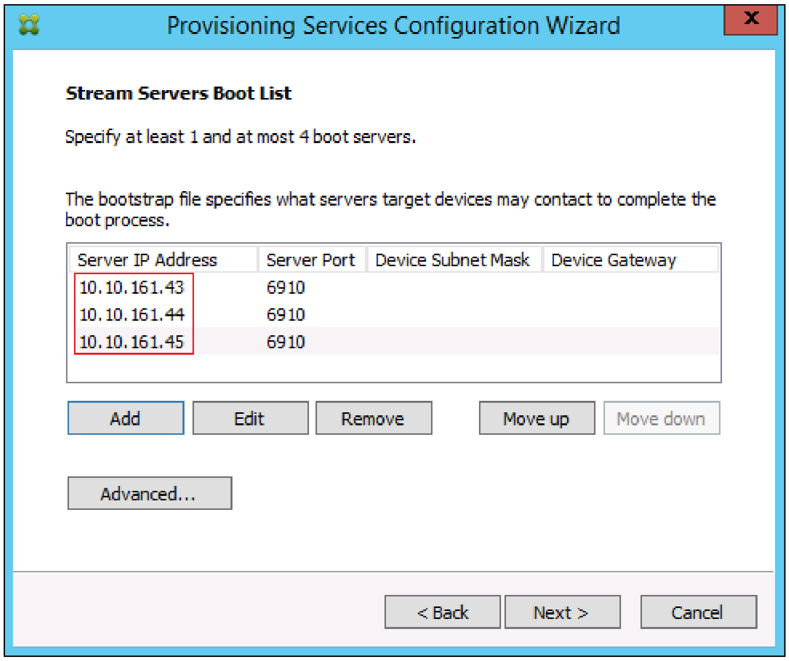

Stream Servers Boot List

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
10.10.161.43	6910		
10.10.161.44	6910		
10.10.161.45	6910		

56. Click Finish to start installation.

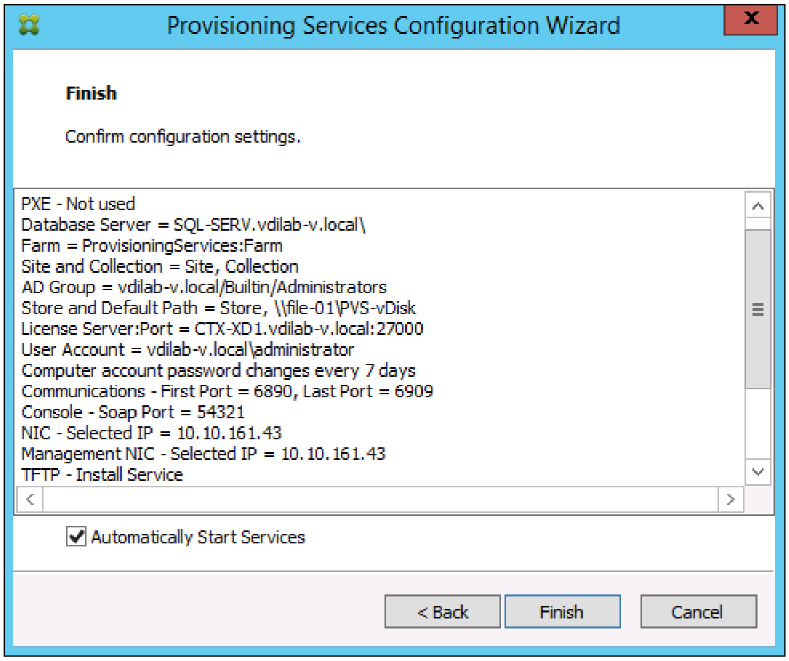

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = SQL-SERV.vdilab-v.local\
Farm = ProvisioningServices:Farm
Site and Collection = Site, Collection
AD Group = vdilab-v.local/Builtin/Administrators
Store and Default Path = Store, \\file-01\PVS-vDisk
License Server:Port = CTX-XD1.vdilab-v.local:27000
User Account = vdilab-v.local\administrator
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.10.161.43
Management NIC - Selected IP = 10.10.161.43
TFTP - Install Service
```

☑ Automatically Start Services

57. When the installation is completed, click Done.

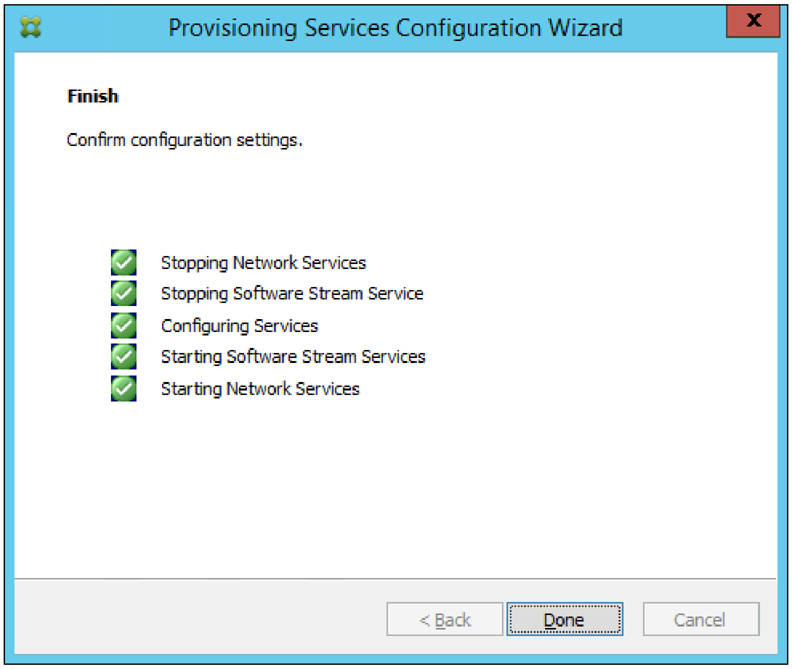

Provisioning Services Configuration Wizard

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 servers, complete the following steps:

1. On the Farm Configuration dialog, select “Join existing farm.”
2. Click Next.

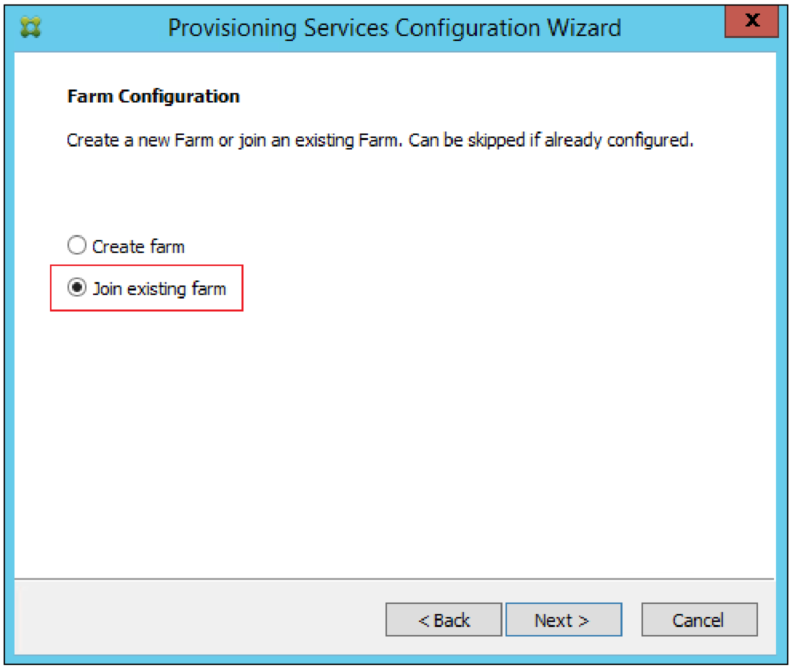

Provisioning Services Configuration Wizard

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

○ Create farm

● Join existing farm

3. Provide the FQDN of the SQL Server.

4. Click Next.

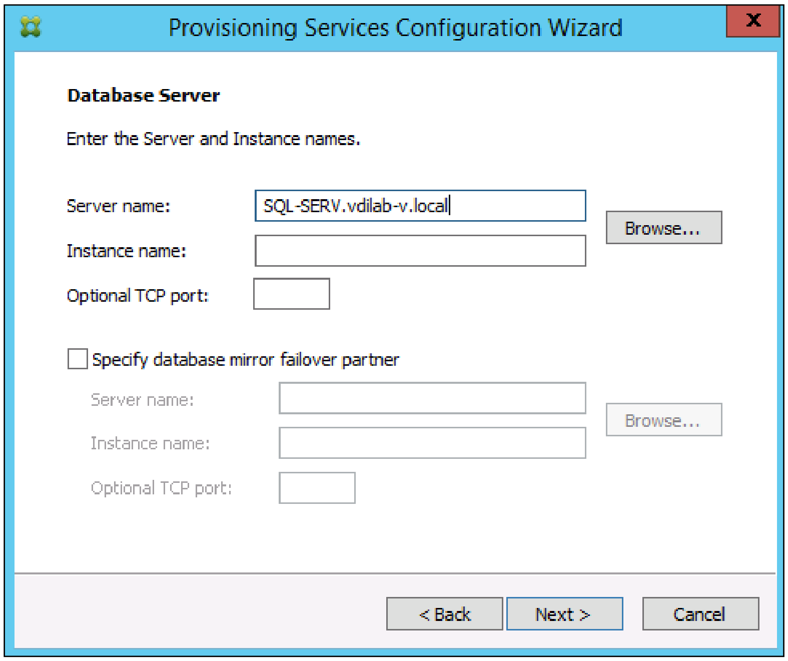

Provisioning Services Configuration Wizard

Database Server

Enter the Server and Instance names.

Server name: SQL-SERV.vdilab-v.local

Instance name:

Optional TCP port:

□ Specify database mirror failover partner

Server name:

Instance name:

Optional TCP port:

5. Accept the Farm Name.

6. Click Next.

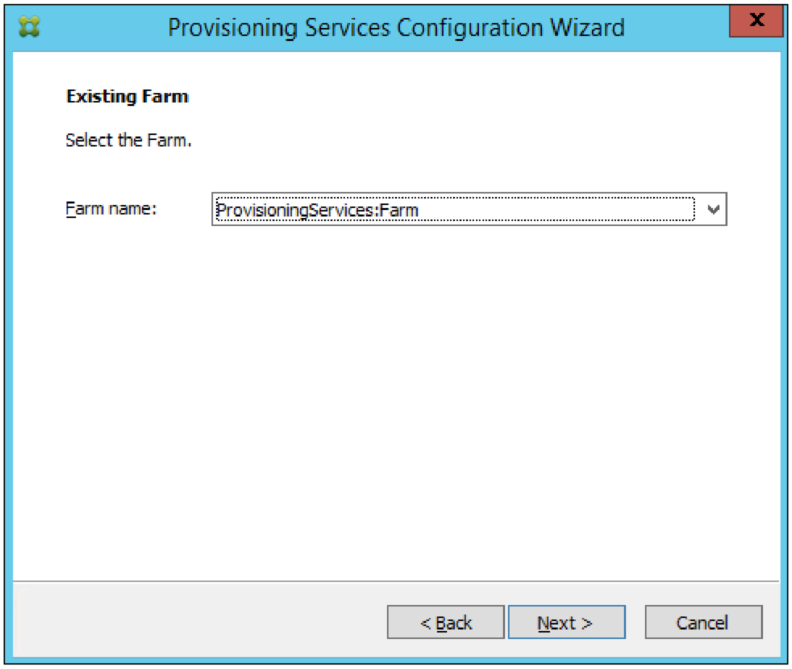

Provisioning Services Configuration Wizard

Existing Farm

Select the Farm.

Farm name: ProvisioningServices:Farm

7. Accept the Existing Site.

8. Click Next.

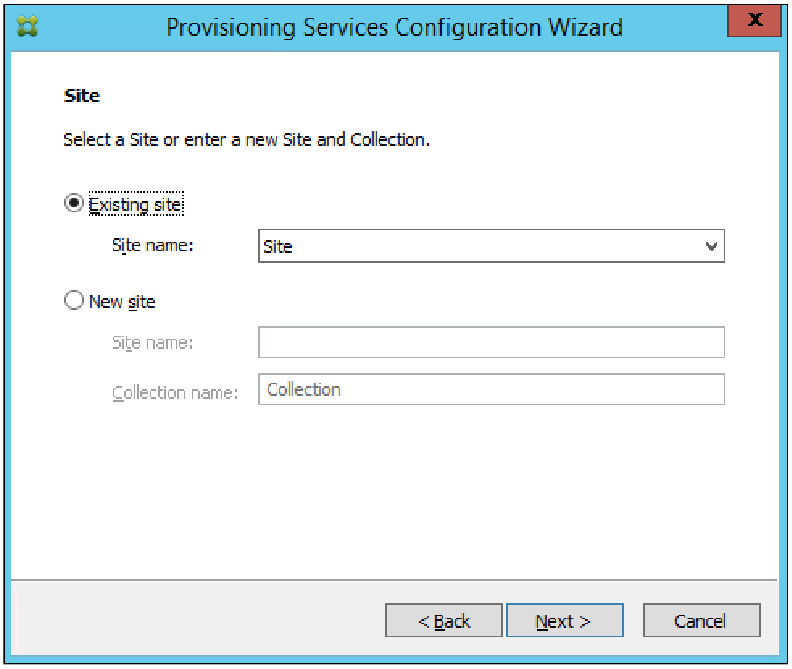

Provisioning Services Configuration Wizard

Site

Select a Site or enter a new Site and Collection.

● Existing site

Site name: Site

○ New site

Site name:

Collection name: Collection

9. Accept the existing vDisk store.

10. Click Next.

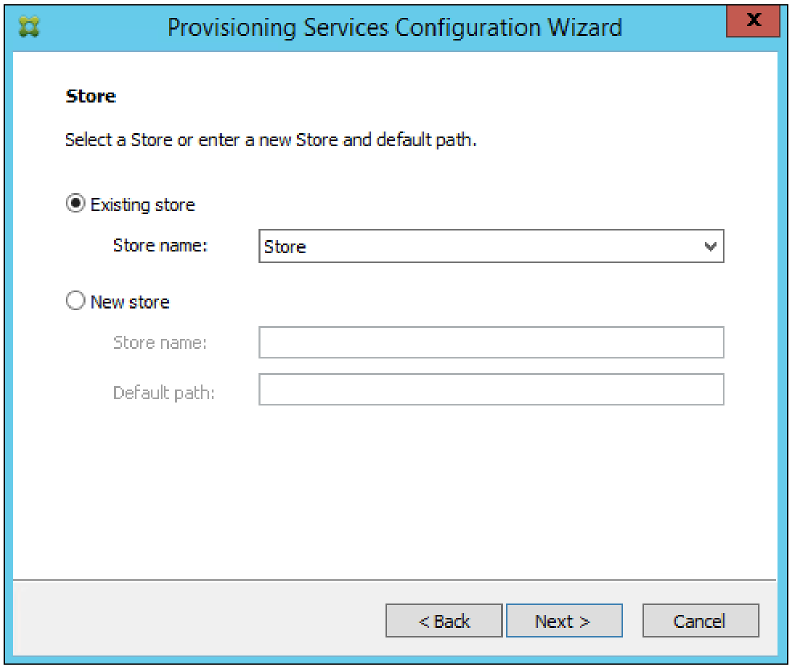

Provisioning Services Configuration Wizard

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.

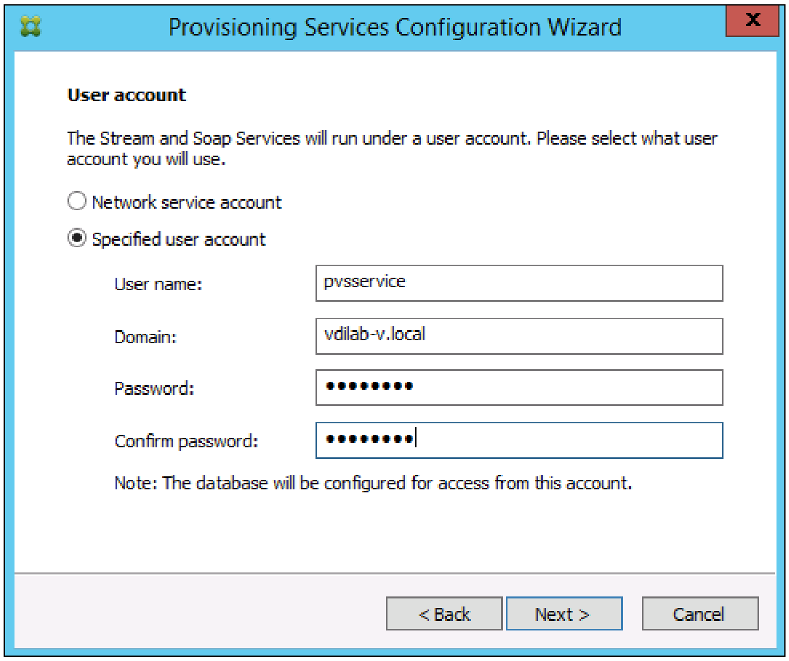

Provisioning Services Configuration Wizard

User account

The Stream and Soap Services will run under a user account. Please select what user account you will use.

○ Network service account

● Specified user account

User name: pvsservice

Domain: vdilab-v.local

Password: ●●●●●●●●

Confirm password: ●●●●●●●●

Note: The database will be configured for access from this account.

13. Set the Days between password updates to 7.

14. Click Next.

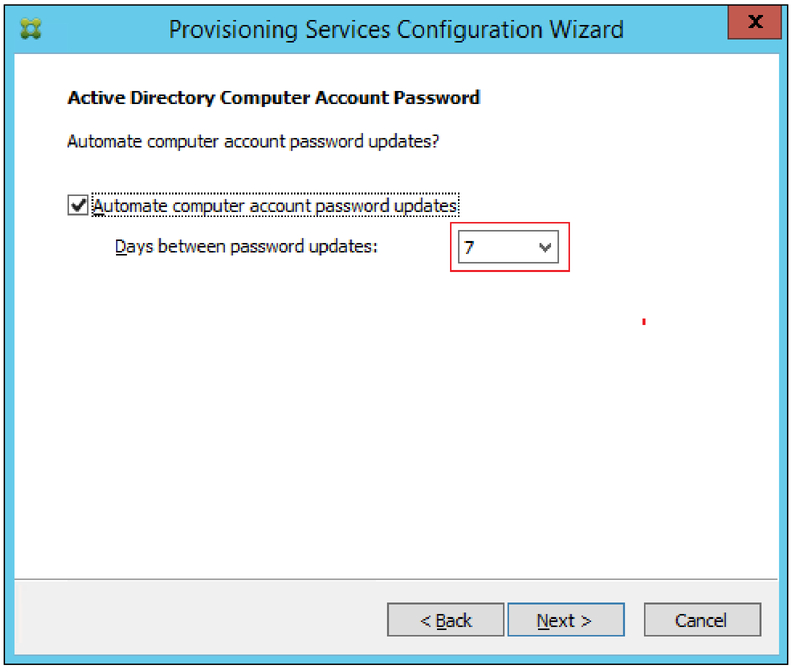

Provisioning Services Configuration Wizard

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.

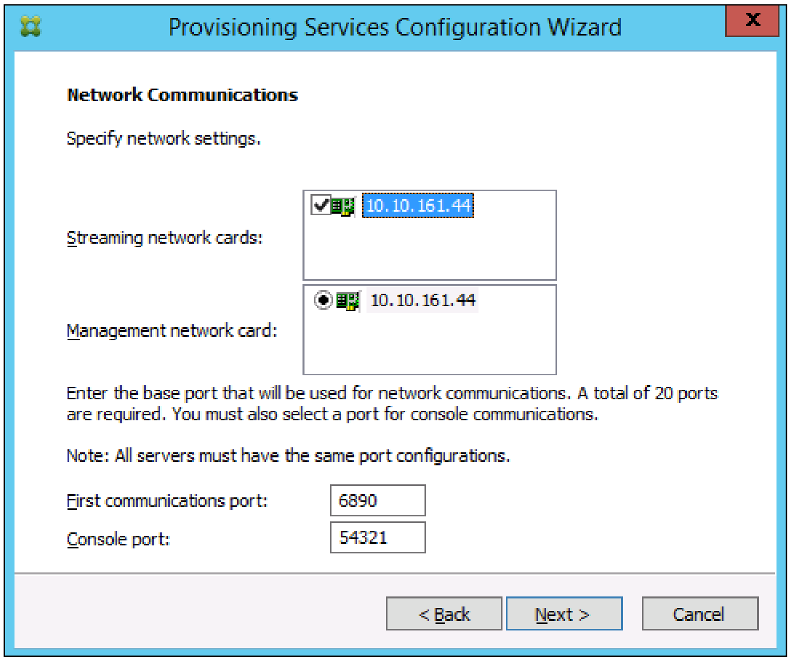

Provisioning Services Configuration Wizard

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.10.161.44

Management network card: ● 10.10.161.44

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.

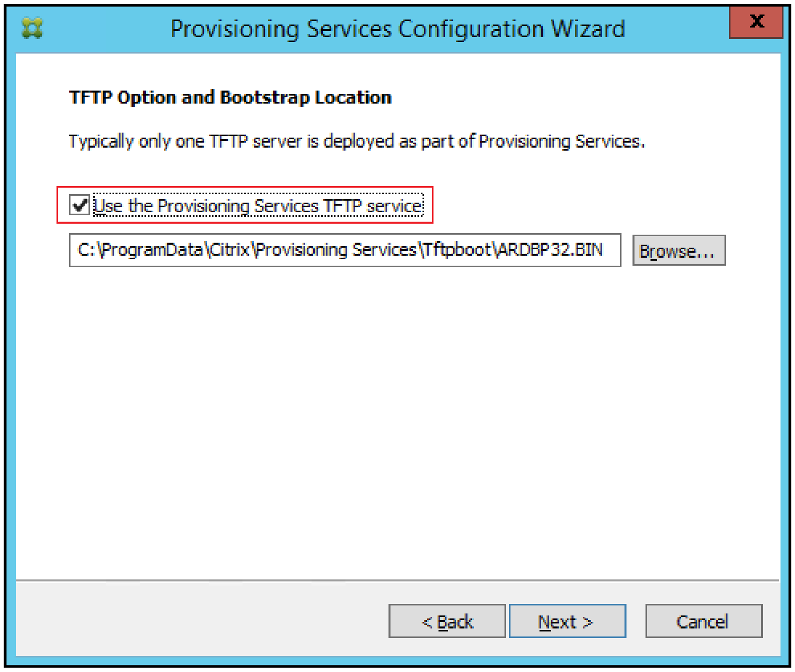

Provisioning Services Configuration Wizard

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

19. Make sure that the IP Addresses for all PVS servers are listed in the Stream Servers Boot List.

20. Click Next.

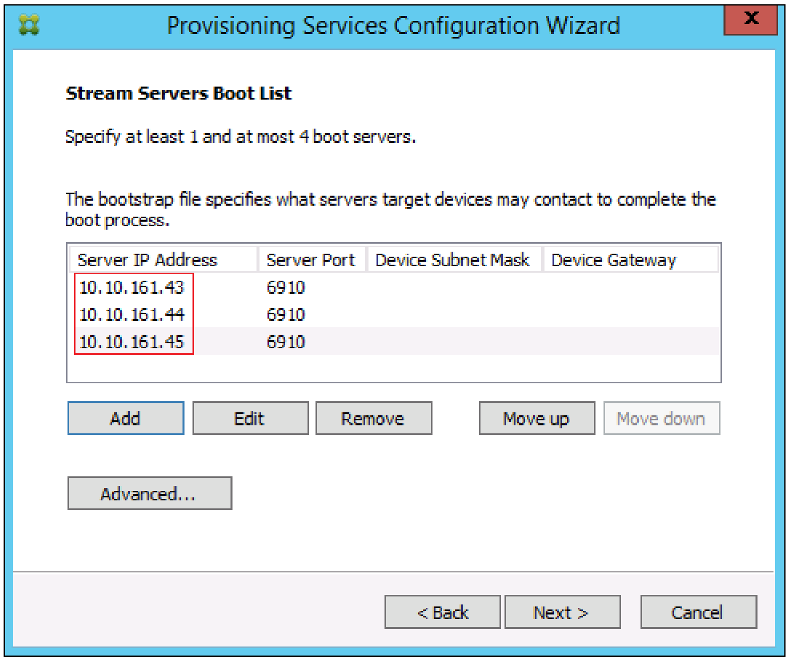

Stream Servers Boot List

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
10.10.161.43	6910		
10.10.161.44	6910		
10.10.161.45	6910		

21. Click Finish to start the installation process.

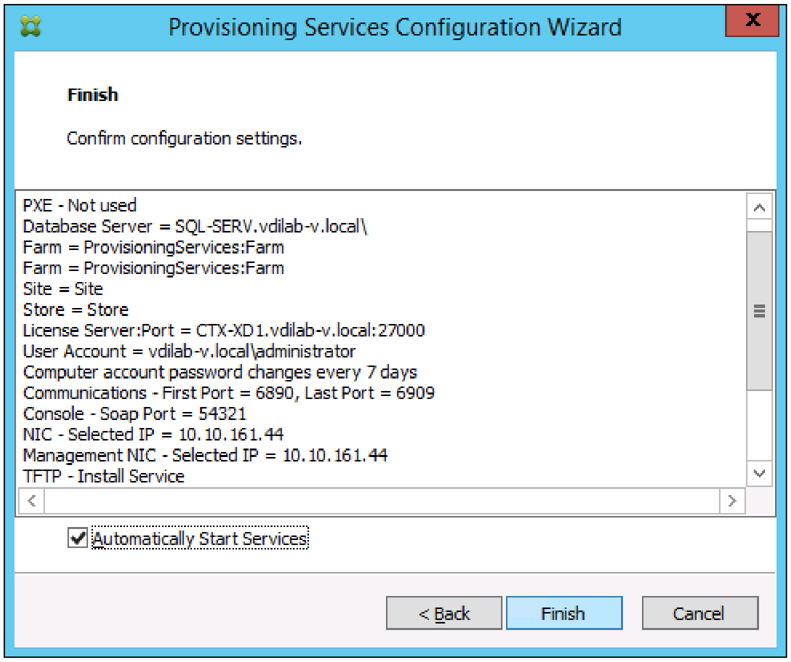

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = SQL-SERV.vdilab-v.local\
Farm = ProvisioningServices:Farm
Farm = ProvisioningServices:Farm
Site = Site
Store = Store
License Server:Port = CTX-XD1.vdilab-v.local:27000
User Account = vdilab-v.local\administrator
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.10.161.44
Management NIC - Selected IP = 10.10.161.44
TFTP - Install Service
```

☑ Automatically Start Services

22. Click Done when the installation finishes.

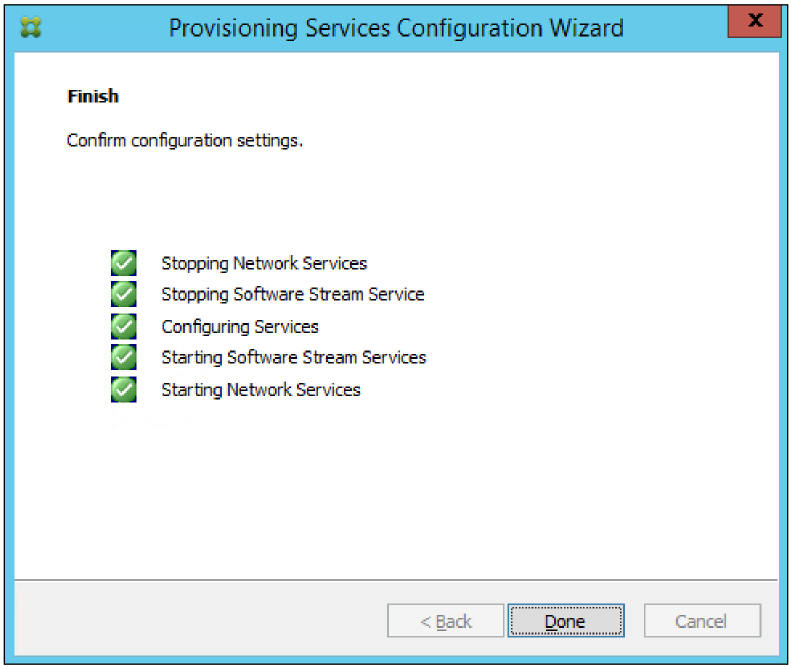

Provisioning Services Configuration Wizard

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 second PVS server, launch the Provisioning Services Console to verify that the PVS Servers and Stores are configured and that DHCP boot options are defined.

23. Launch Provisioning Services Console and select Connect to Farm.

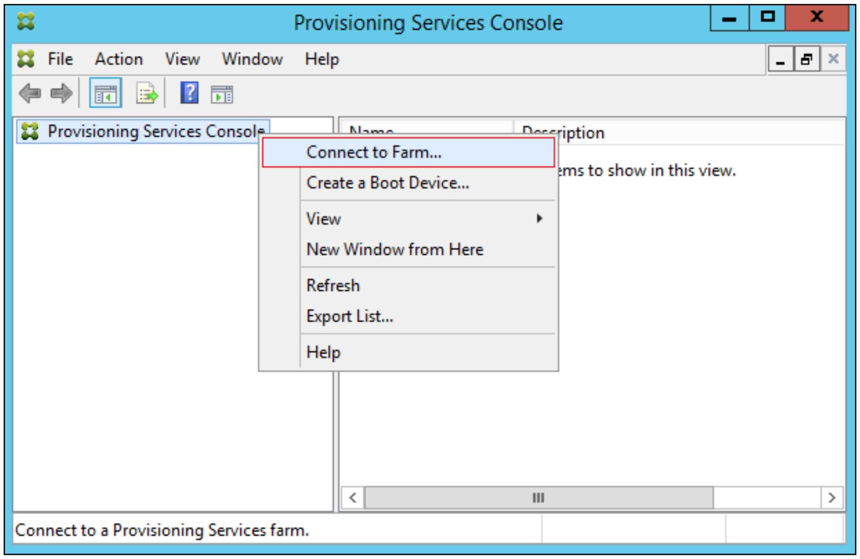

24. Enter localhost for the PVS1 server.

25. Click Connect.

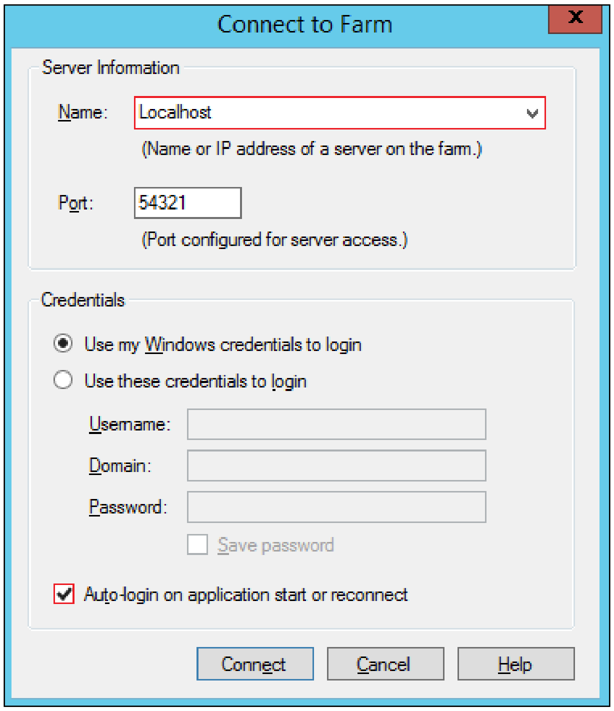

Name: Localhost

Port: 54321

● Use my Windows credentials to login

○ Use these credentials to login

□ Save password

☑ Auto-login on application start or reconnect

26. Select Store Properties from the drop-down menu.

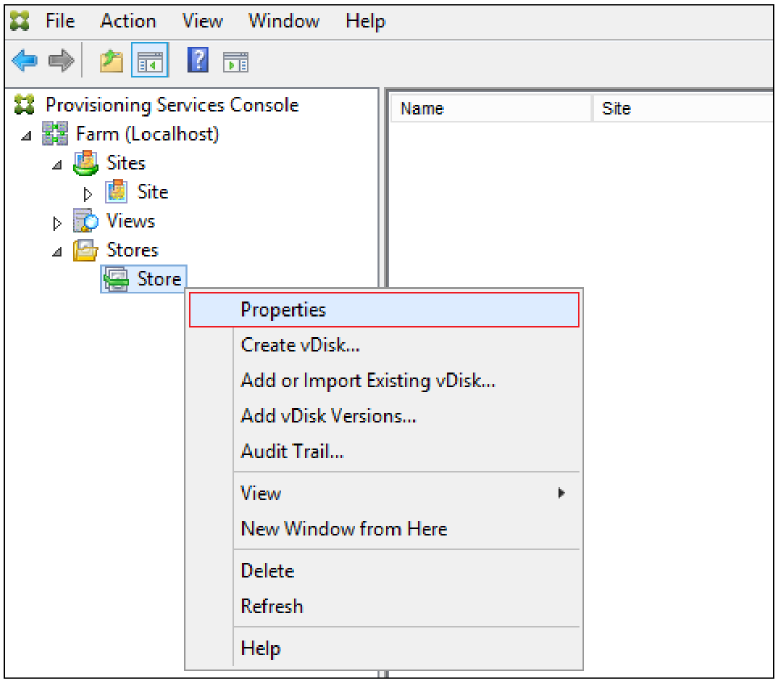

27. In the Store Properties dialog, add the Default store path to the list of Default write cache paths.

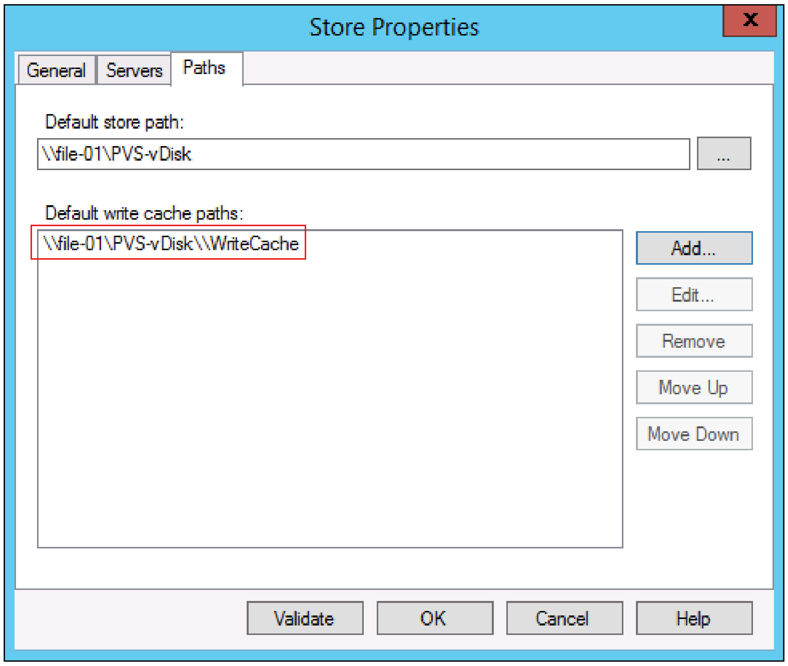

Default store path: \\file-01\PVS-vDisk

Default write cache paths: \\file-01\PVS-vDisk\\WriteCache

28. Click Validate. If the validation is successful, click OK to continue.

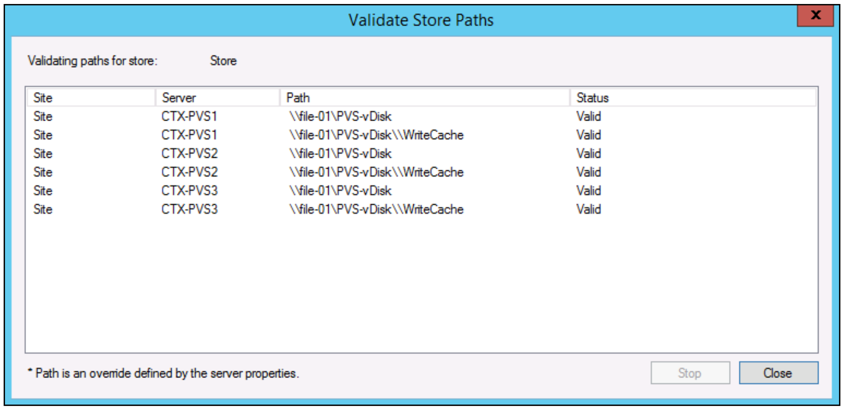

Validating paths for store: Store

Site	Server	Path	Status
Site	CTX-PVS1	\\file-01\PVS-vDisk	Valid
Site	CTX-PVS1	\\file-01\PVS-vDisk\\WriteCache	Valid
Site	CTX-PVS2	\\file-01\PVS-vDisk	Valid
Site	CTX-PVS2	\\file-01\PVS-vDisk\\WriteCache	Valid
Site	CTX-PVS3	\\file-01\PVS-vDisk	Valid
Site	CTX-PVS3	\\file-01\PVS-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 7.9 ISO.

2. Click Start on the Welcome screen.

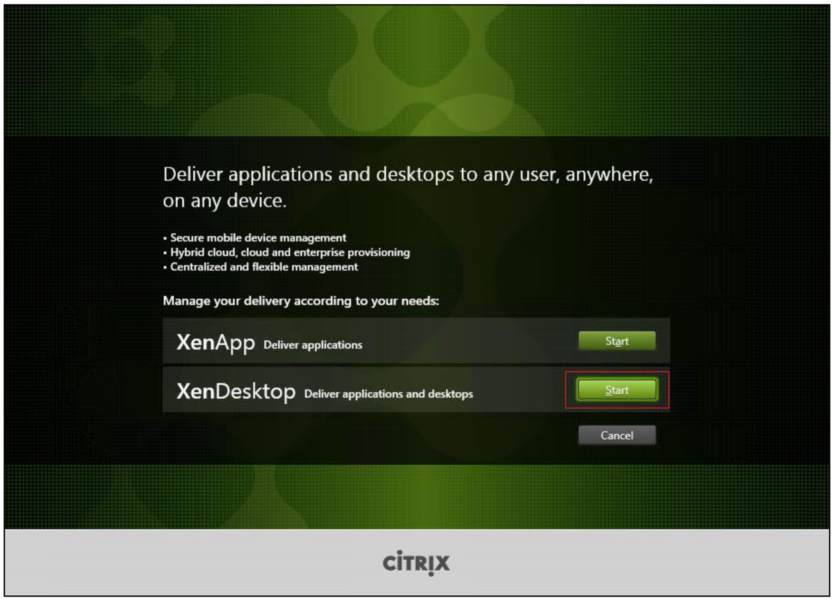

Deliver applications and desktops to any user, anywhere, on any device.

- 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

3. To install the VDA for the Hosted Virtual Desktops (VDI), select Virtual Delivery Agent for Windows Desktop OS. 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.

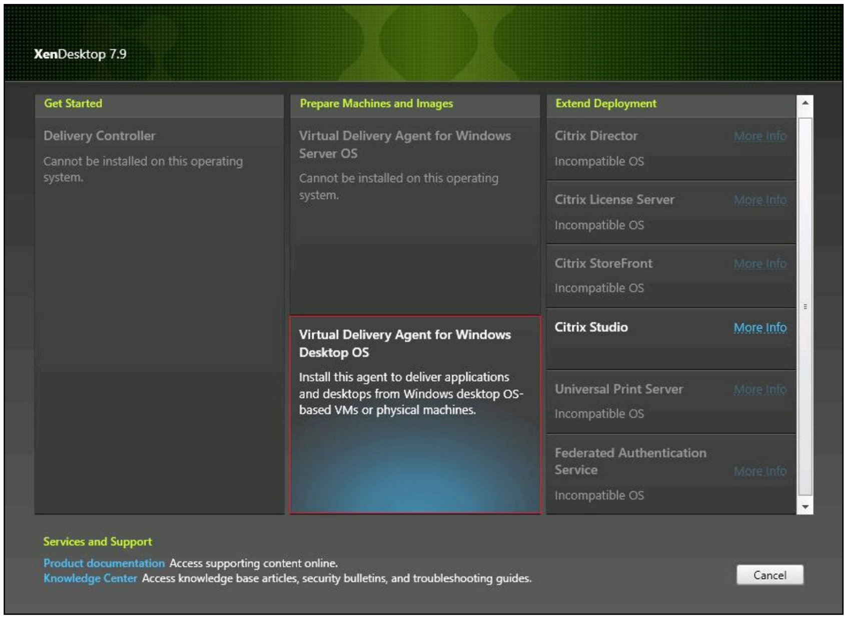

XenDesktop 7.9

Get Started

Delivery Controller
Cannot be installed on this operating system.

Prepare Machines and Images

Virtual Delivery Agent for Windows Server OS
Cannot be installed on this operating system.

Virtual Delivery Agent for Windows Desktop OS
Install this agent to deliver applications and desktops from Windows desktop OS-based VMs or physical machines.

Extend Deployment

Citrix Director — More Info
Incompatible OS

Citrix License Server — More Info
Incompatible OS

Citrix StoreFront — More Info
Incompatible OS

Citrix Studio — More Info

Universal Print Server — More Info
Incompatible OS

Federated Authentication Service — More Info
Incompatible OS

Services and Support

Product documentation Access supporting content online.
Knowledge Center Access knowledge base articles, security bulletins, and troubleshooting guides.

4. Select “Create a Master Image.”

5. Click Next.

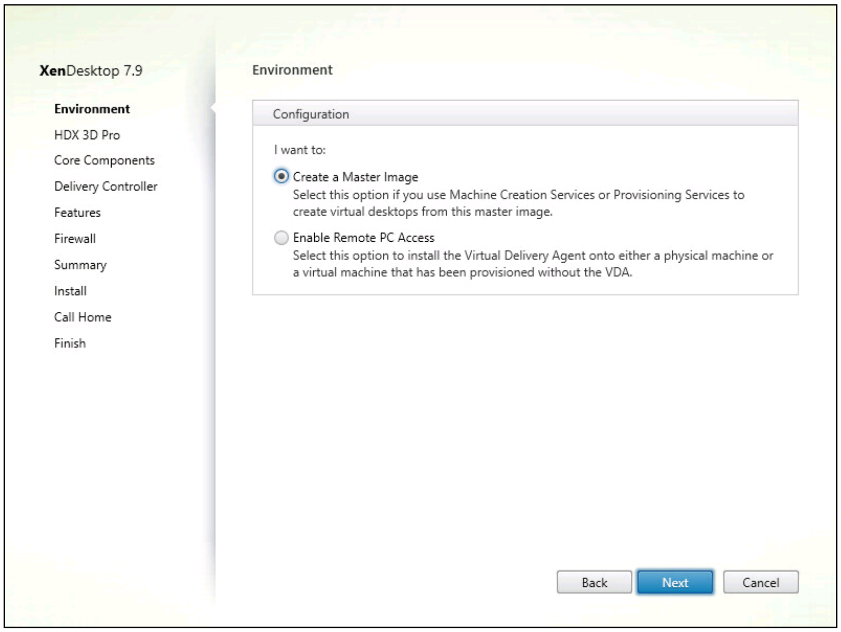

XenDesktop 7.9

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.

6. For the VDI vDisk, select “No, install the standard VDA.”

7. Click Next.

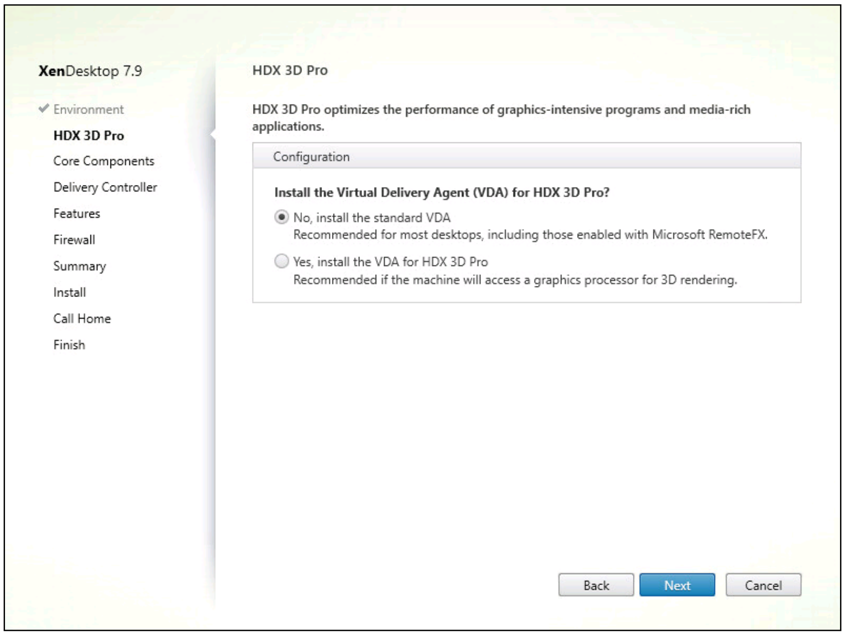

XenDesktop 7.9

HDX 3D Pro

HDX 3D Pro optimizes the performance of graphics-intensive programs and media-rich applications.

Configuration

Install the Virtual Delivery Agent (VDA) for HDX 3D Pro?

● No, install the standard VDA
Recommended for most desktops, including those enabled with Microsoft RemoteFX.

○ Yes, install the VDA for HDX 3D Pro
Recommended if the machine will access a graphics processor for 3D rendering.

8. Optional: Select Citrix Receiver.

9. Click Next.

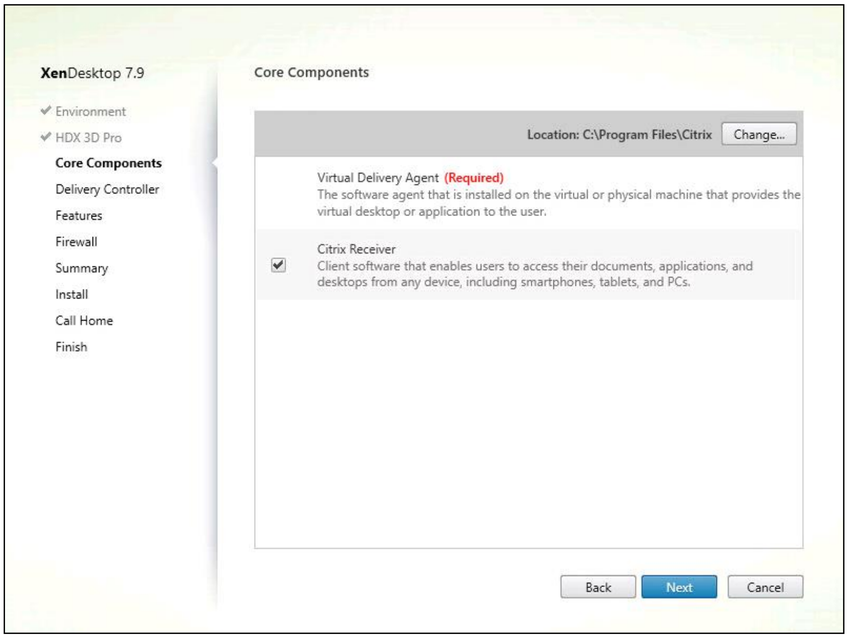

XenDesktop 7.9

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.

10. Select “Do it manually” and specify the FQDN of the Delivery Controllers.

11. Click Next.

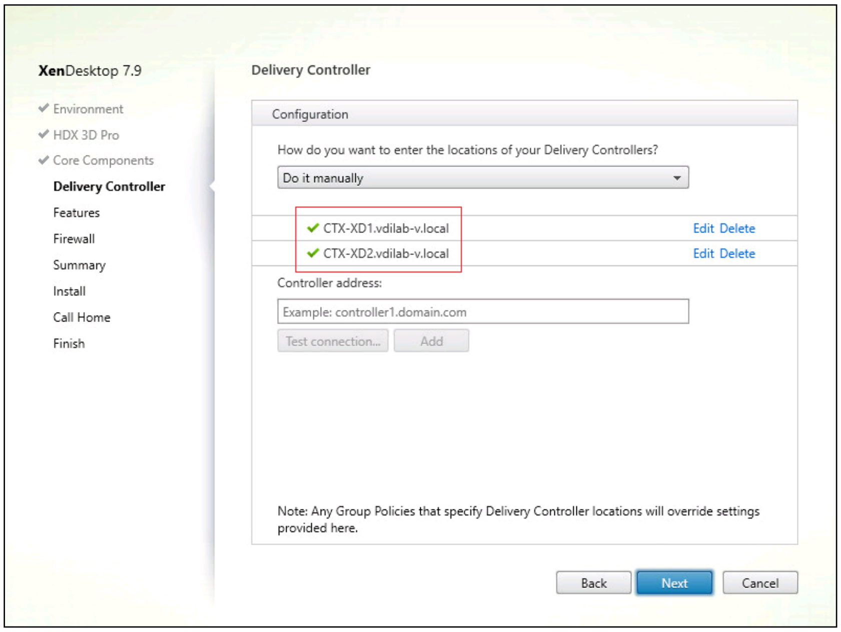

XenDesktop 7.9

Delivery Controller

Configuration

How do you want to enter the locations of your Delivery Controllers?

Do it manually

✔ CTX-XD1.vdilab-v.local

✔ CTX-XD2.vdilab-v.local

Controller address:

Example: controller1.domain.com

Note: Any Group Policies that specify Delivery Controller locations will override settings provided here.

12. Accept the default features.

13. Click Next.

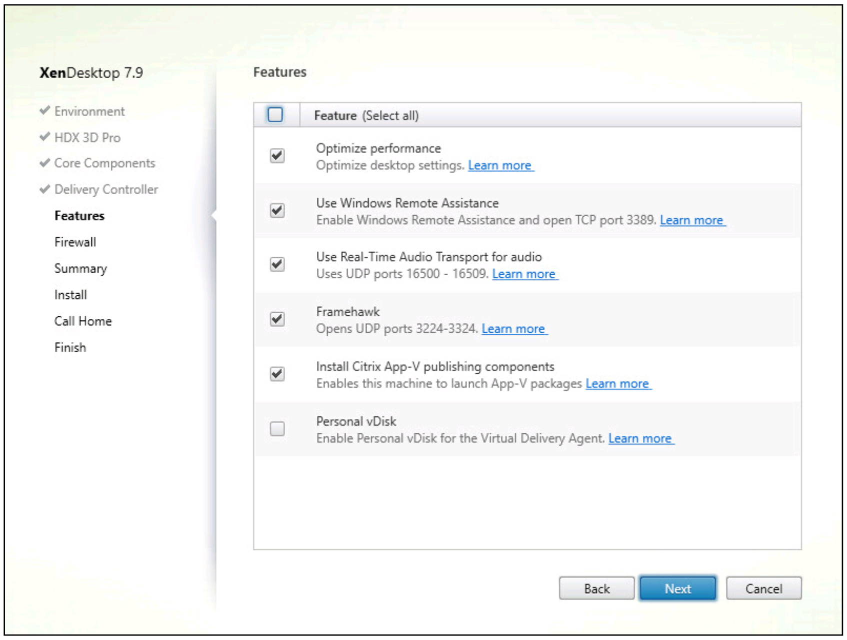

Features

□	Feature (Select all)
☑	Optimize performance Optimize desktop settings. Learn more
☑	Use Windows Remote Assistance Enable Windows Remote Assistance and open TCP port 3389. Learn more
☑	Use Real-Time Audio Transport for audio Uses UDP ports 16500 - 16509. Learn more
☑	Framehawk Opens UDP ports 3224-3324. Learn more
☑	Install Citrix App-V publishing components Enables this machine to launch App-V packages Learn more
□	Personal vDisk Enable Personal vDisk for the Virtual Delivery Agent. Learn more

14. Allow the firewall rules to be configured Automatically.

15. Click Next.

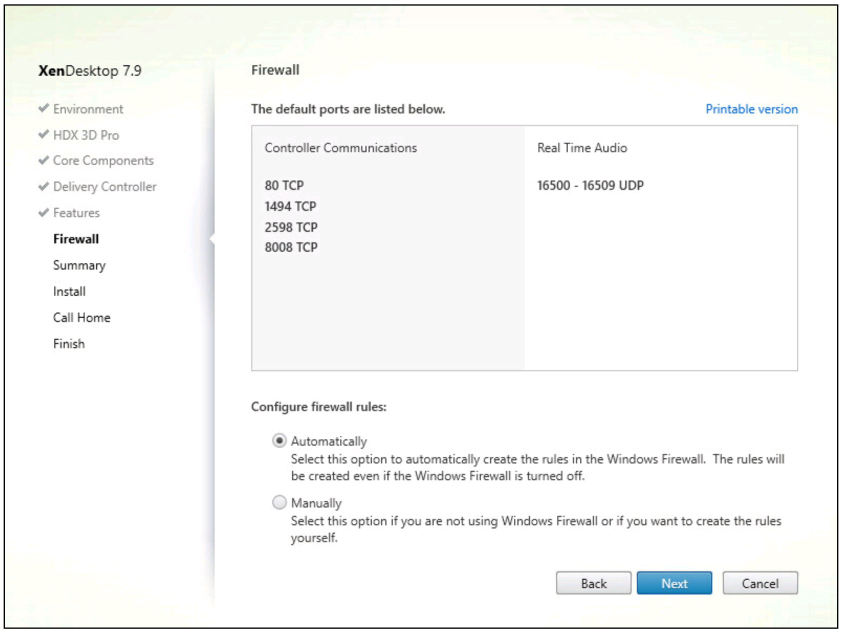

Firewall

The default ports are listed below.

Controller Communications	Real Time Audio
80 TCP 1494 TCP 2598 TCP 8008 TCP	16500 - 16509 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.

16. Verify the Summary and click Install.

17. (Optional) Select Call Home participation.

18. Check “Restart Machine.”

19. Click Finish and the machine will reboot automatically.

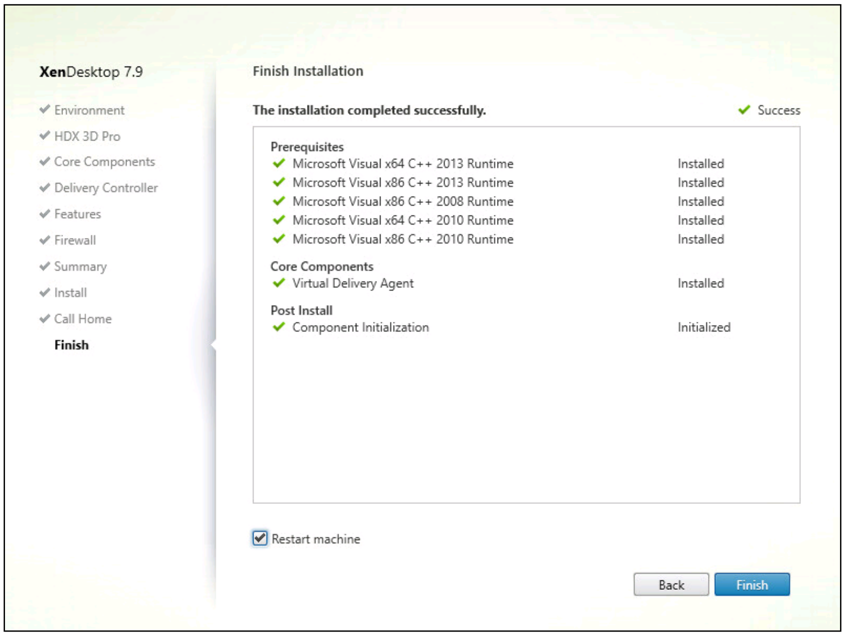

XenDesktop 7.9

✓ Environment
✓ HDX 3D Pro
✓ Core Components
✓ Delivery Controller
✓ Features
✓ Firewall
✓ Summary
✓ Install
✓ Call Home
Finish

Finish Installation

The installation completed successfully. ✓ Success

Prerequisites	
✓ Microsoft Visual x64 C++ 2013 Runtime	Installed
✓ Microsoft Visual x86 C++ 2013 Runtime	Installed
✓ Microsoft Visual x86 C++ 2008 Runtime	Installed
✓ Microsoft Visual x64 C++ 2010 Runtime	Installed
✓ Microsoft Visual x86 C++ 2010 Runtime	Installed
Core Components	
✓ Virtual Delivery Agent	Installed
Post Install	
✓ Component Initialization	Initialized

☑ Restart machine

Repeat the procedure so that VDAs are installed for both VDI (using the Windows 10 OS image) and the RDS desktops (using the Windows Server 2012 R2 image).

Install the Citrix Provisioning Services 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. On the Window 10 Master Target Device, launch the PVS installer from the Provisioning Services 7.9 ISO.

2. Click the Target Device Installation button.

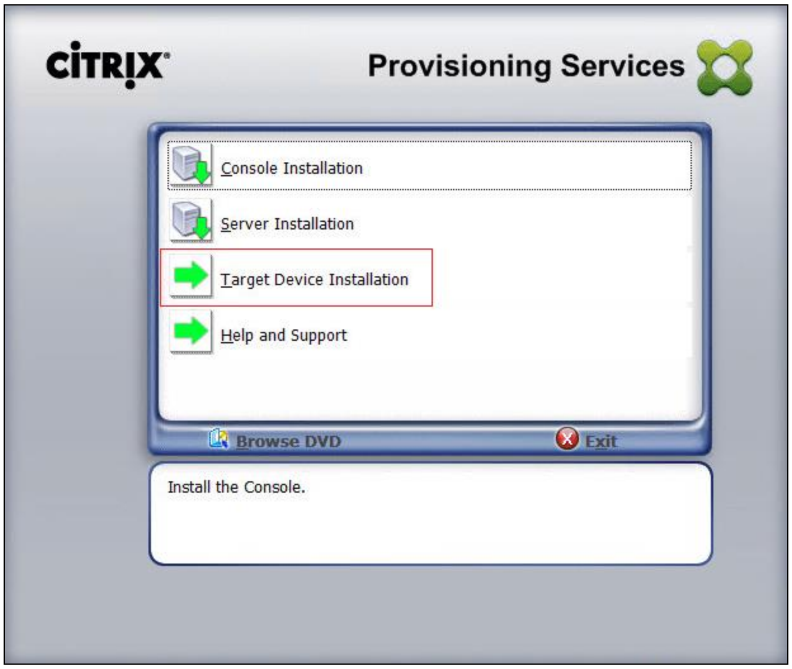

Install the Console.

The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

3. Click Next.

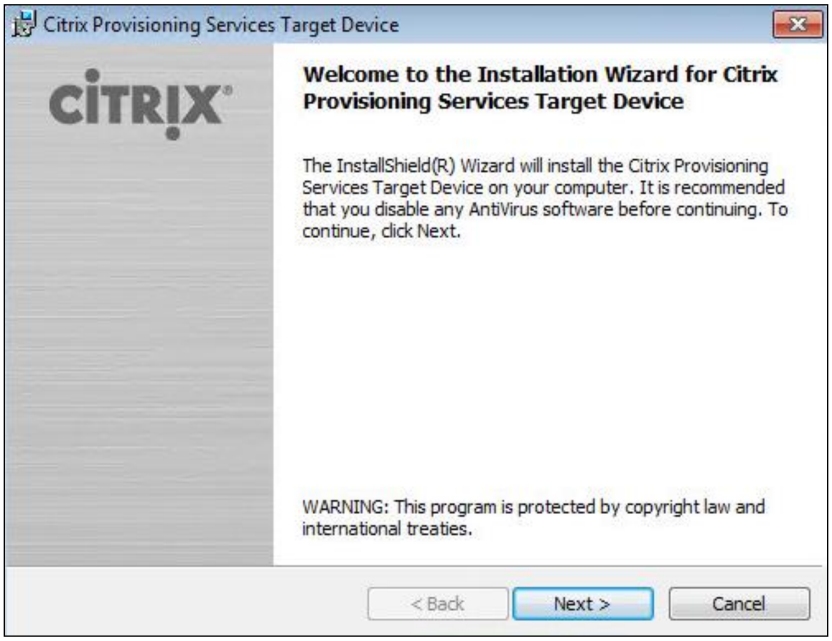

Welcome to the Installation Wizard for Citrix Provisioning Services Target Device

The InstallShield(R) Wizard will install the Citrix Provisioning Services Target Device 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. Confirm the installation settings and click Install.

5. Deselect the checkbox to launch the Imaging Wizard and click Finish.

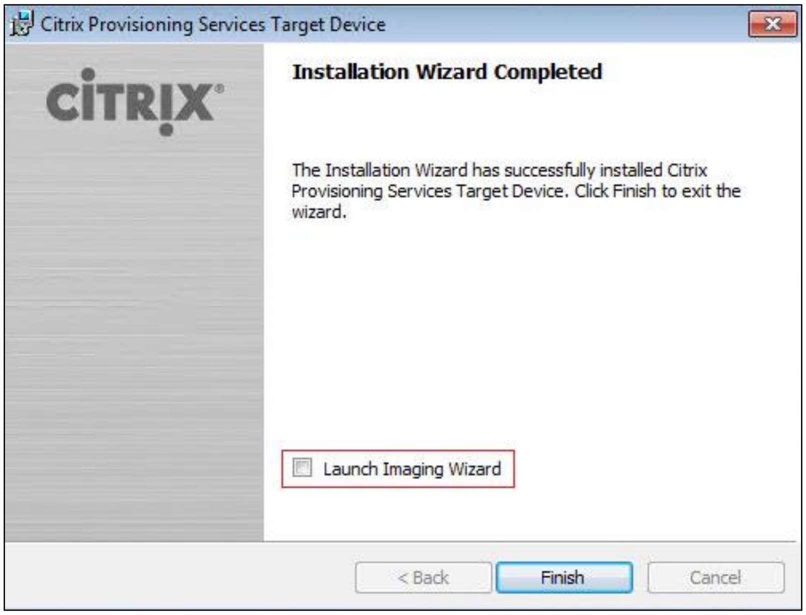

Citrix Provisioning Services Target Device

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix Provisioning Services Target Device. Click Finish to exit the wizard.

□ Launch Imaging Wizard

6. Reboot the machine.

Create Citrix Provisioning Services 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.

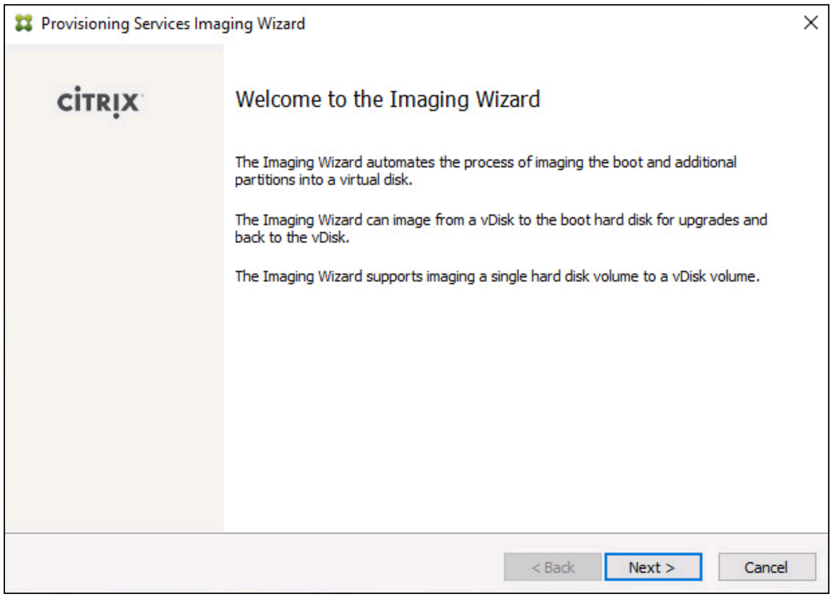

Provisioning Services Imaging Wizard

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.

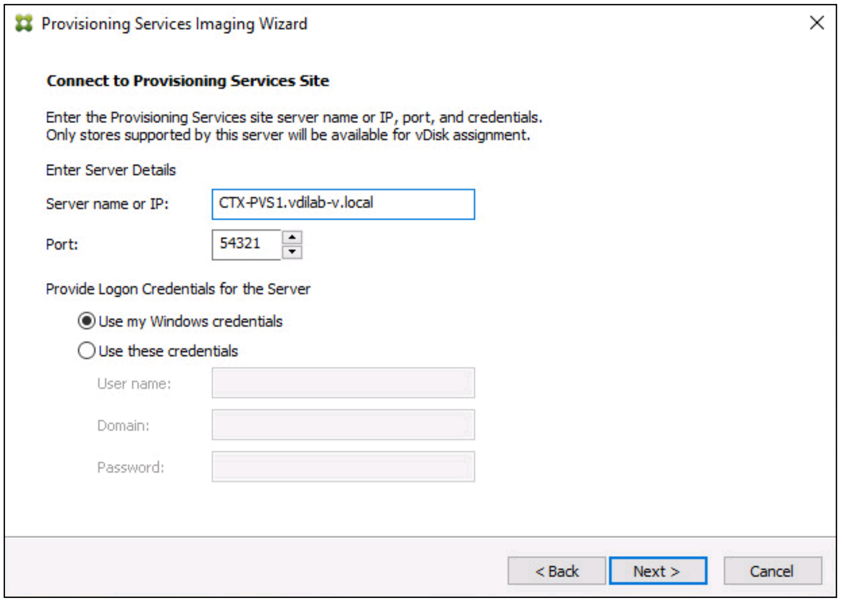

Provisioning Services Imaging Wizard

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: CTX-PVS1.vdilab-v.local

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.

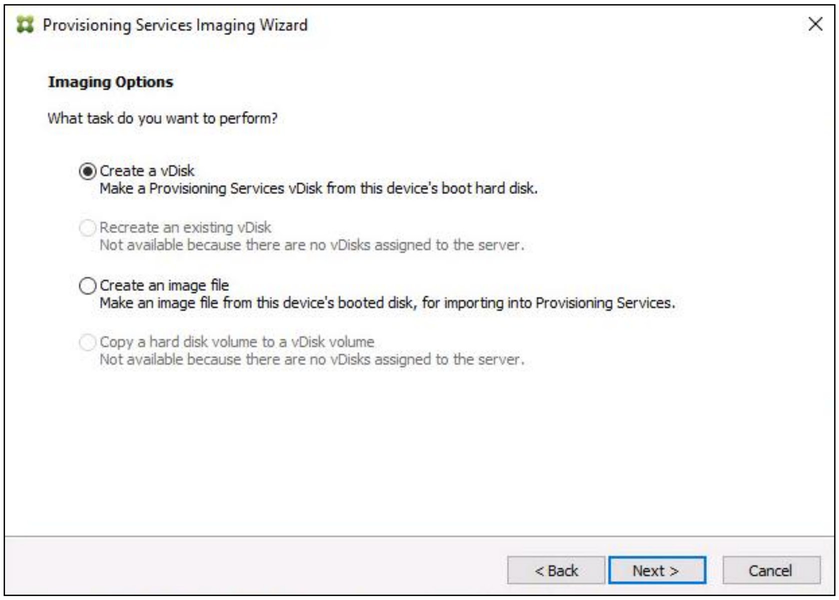

Provisioning Services Imaging Wizard

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.

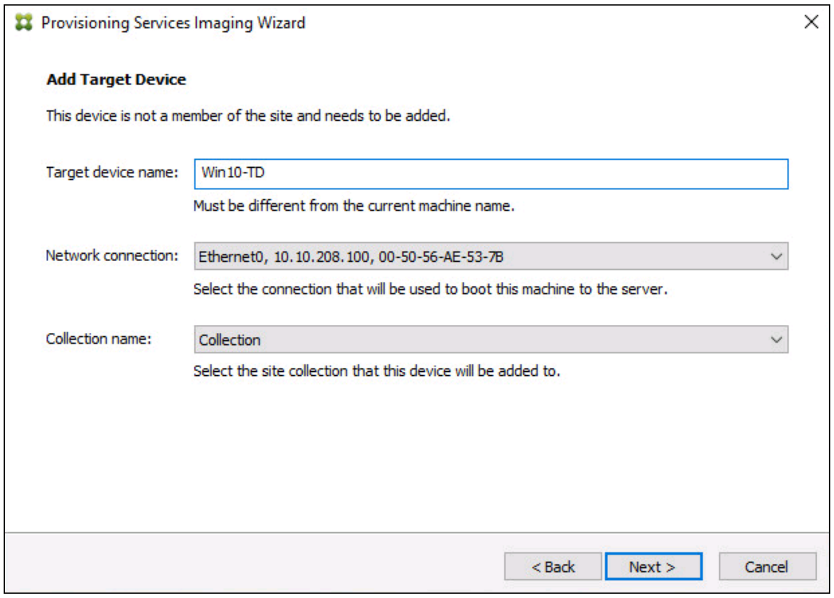

Provisioning Services Imaging Wizard

Add Target Device

This device is not a member of the site and needs to be added.

Target device name: Win10-TD
Must be different from the current machine name.

Network connection: Ethernet0, 10.10.208.100, 00-50-56-AE-53-7B
Select the connection that will be used to boot this machine to the server.

Collection name: Collection
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 menu. (This CVD used Dynamic rather than Fixed vDisks.)

13. Click Next.

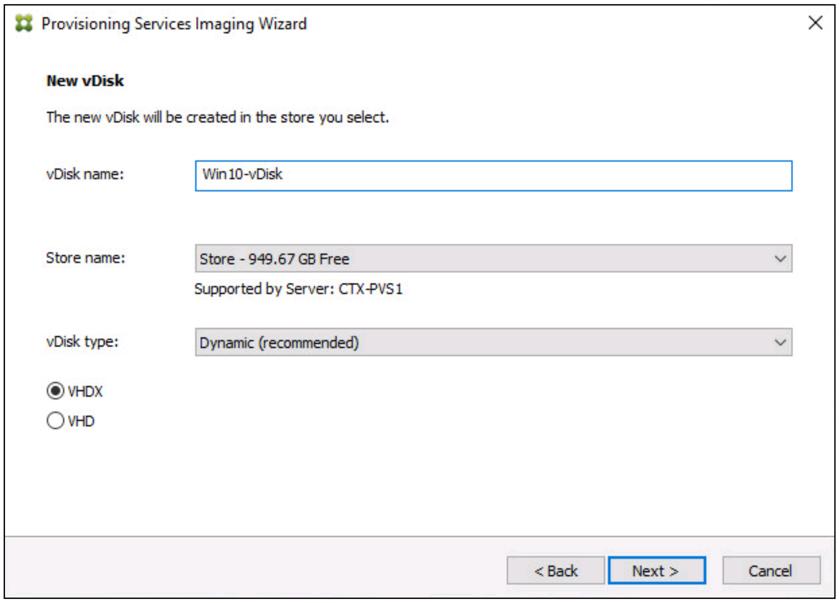

Provisioning Services Imaging Wizard

New vDisk

The new vDisk will be created in the store you select.

vDisk name: Win10-vDisk

Store name: Store - 949.67 GB Free

Supported by Server: CTX-PVS1

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 select None.

15. Click Next.

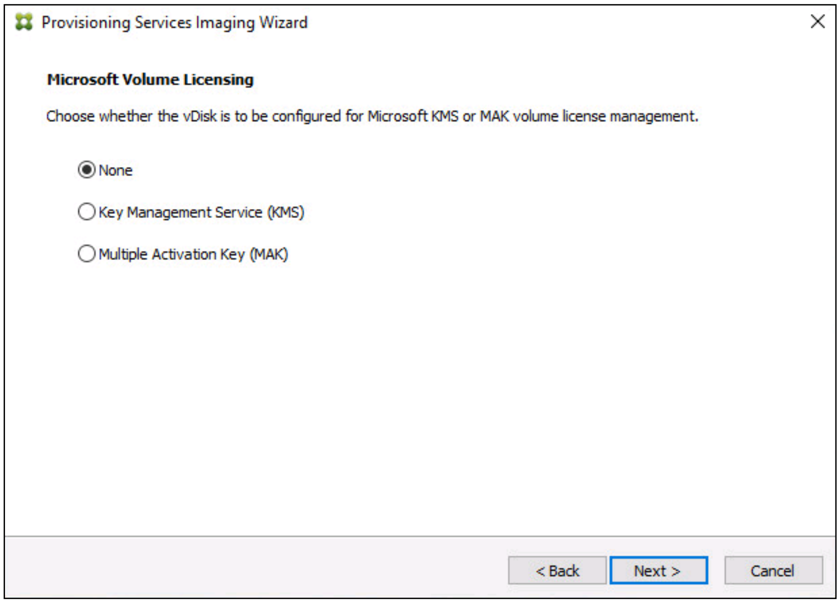

Provisioning Services Imaging Wizard

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.

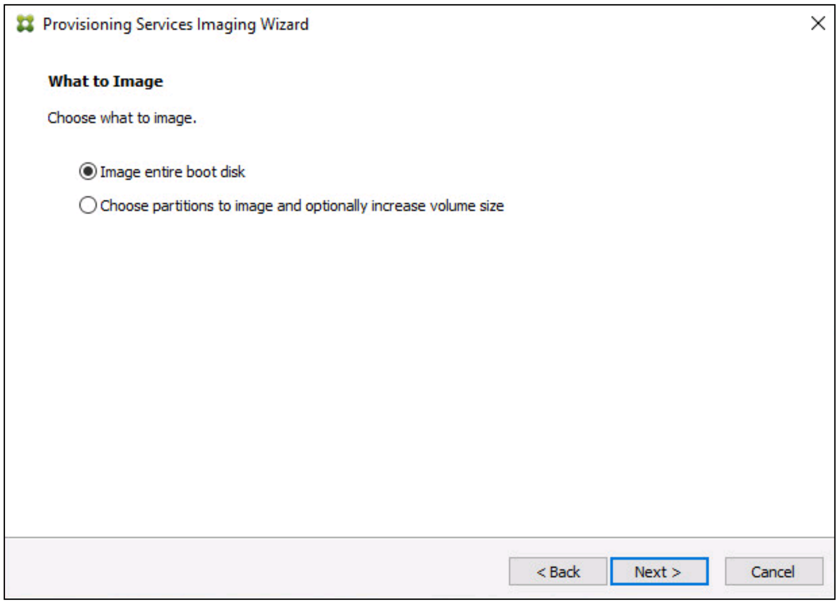

Provisioning Services Imaging Wizard

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.

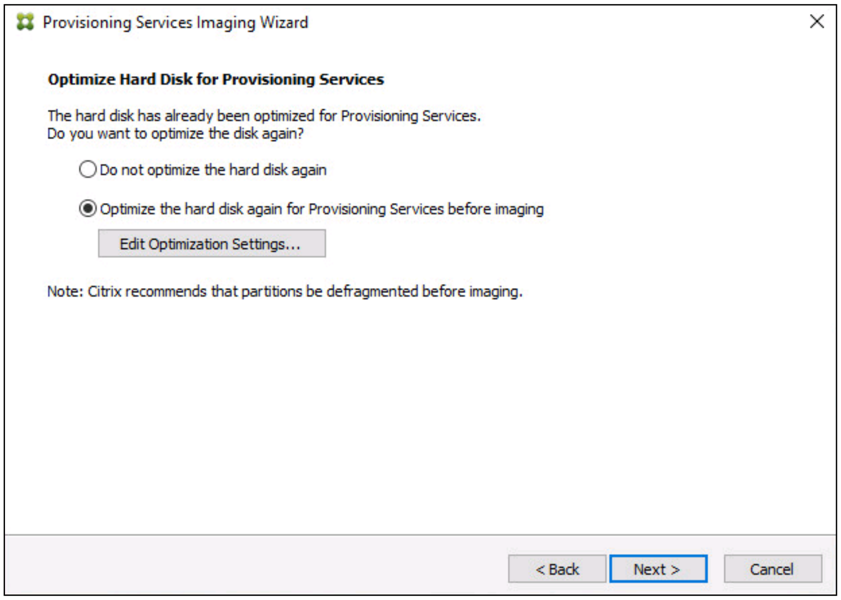

Provisioning Services Imaging Wizard

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.

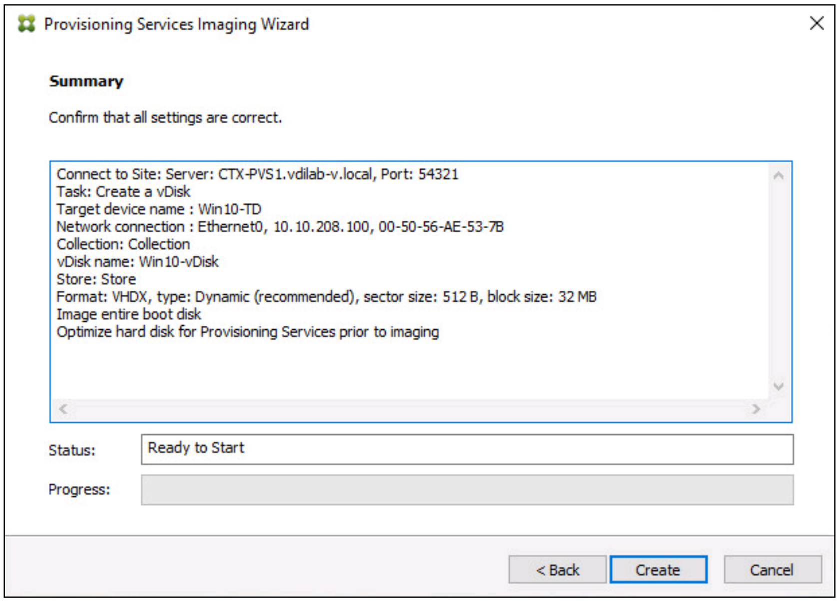

Provisioning Services Imaging Wizard

Summary

Confirm that all settings are correct.

```
Connect to Site: Server: CTX-PVS1.vdilab-v.local, Port: 54321
Task: Create a vDisk
Target device name : Win10-TD
Network connection : Ethernet0, 10.10.208.100, 00-50-56-AE-53-7B
Collection: Collection
vDisk name: Win10-vDisk
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.

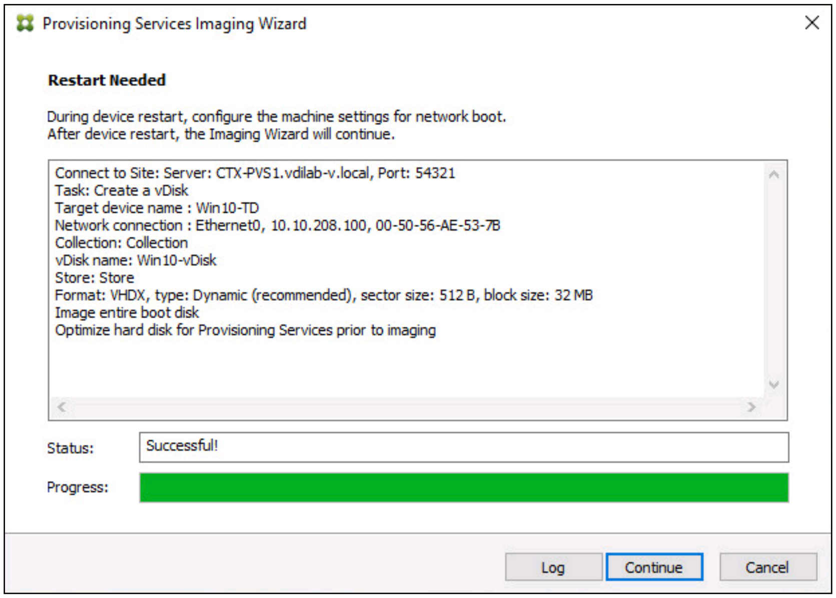

Provisioning Services Imaging Wizard

Restart Needed

During device restart, configure the machine settings for network boot.
After device restart, the Imaging Wizard will continue.

```
Connect to Site: Server: CTX-PVS1.vdilab-v.local, Port: 54321
Task: Create a vDisk
Target device name : Win10-TD
Network connection : Ethernet0, 10.10.208.100, 00-50-56-AE-53-7B
Collection: Collection
vDisk name: Win10-vDisk
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.

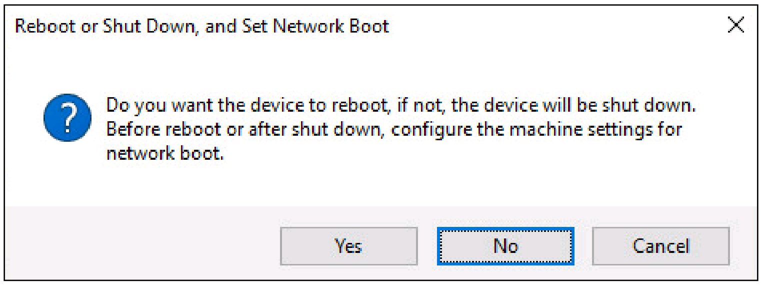

Reboot or Shut Down, and Set Network Boot

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.

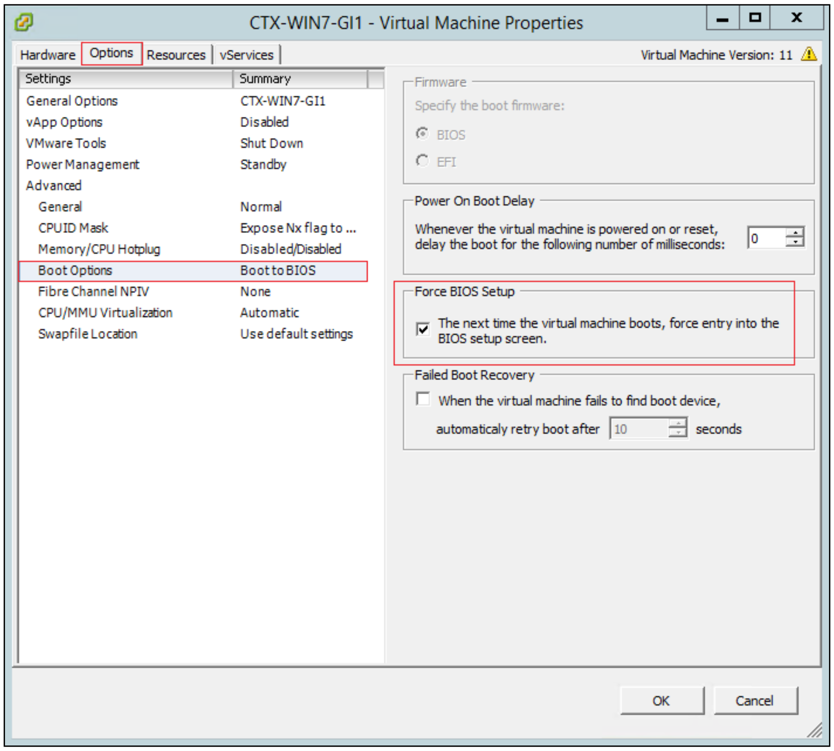

CTX-WIN7-GI1 - Virtual Machine Properties

Virtual Machine Version: 11

Settings	Summary
General Options	CTX-WIN7-GI1
vApp Options	Disabled
VMware Tools	Shut Down
Power Management	Standby
Advanced	
General	Normal
CPUID Mask	Expose Nx flag to ...
Memory/CPU Hotplug	Disabled/Disabled
Boot Options	Boot to BIOS
Fibre Channel NPIV	None
CPU/MMU Virtualization	Automatic
Swapfile Location	Use default settings

Firmware

Specify the boot firmware:

● BIOS

○ EFI

Power On Boot Delay

Whenever the virtual machine is powered on or reset, delay the boot for the following number of milliseconds: 0

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 boot device, automaticaly retry boot after 10 seconds

24. Restart Virtual Machine.

25. Configure the BIOS/VM settings for PXE/network boot, putting Network boot from VMware VMXNET3 at the top of the boot device list.

26. Select Exit Saving Changes.

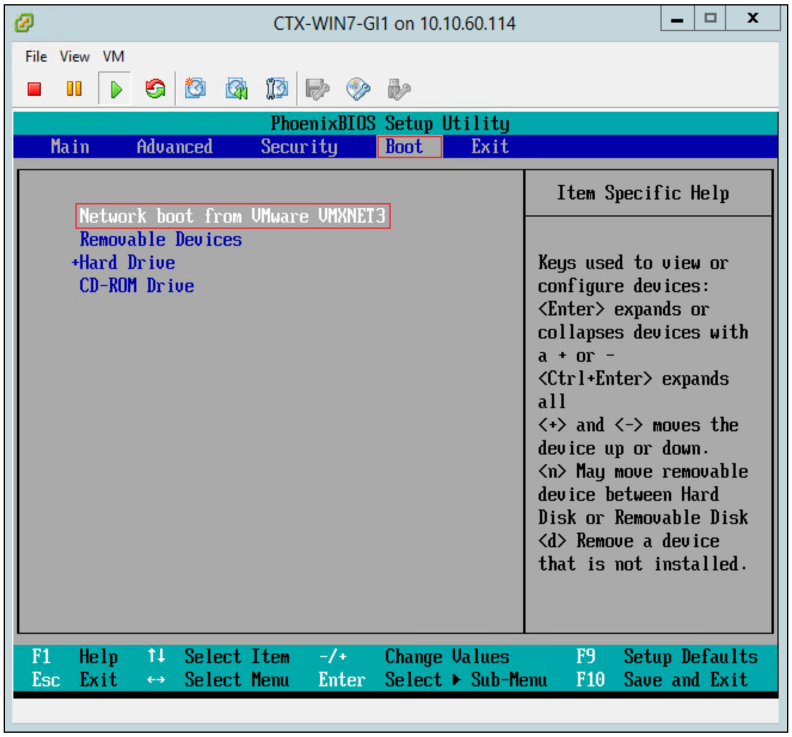

CTX-WIN7-GI1 on 10.10.60.114
File View VM
PhoenixBIOS Setup Utility
Main Advanced Security Boot Exit
Item Specific Help
F1 Help ↑↓ Select Item -/+ Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

```
Network boot from VMware VMXNET3
Removable Devices
+Hard Drive
CD-ROM Drive
```

```
Keys used to view or
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.
```

After restarting the VM, log into the VDI or RDS master target. The PVS imaging process begins, copying the contents of the C: drive to the PVS vDisk located on the server.

27. If prompted to Restart select Restart Later.

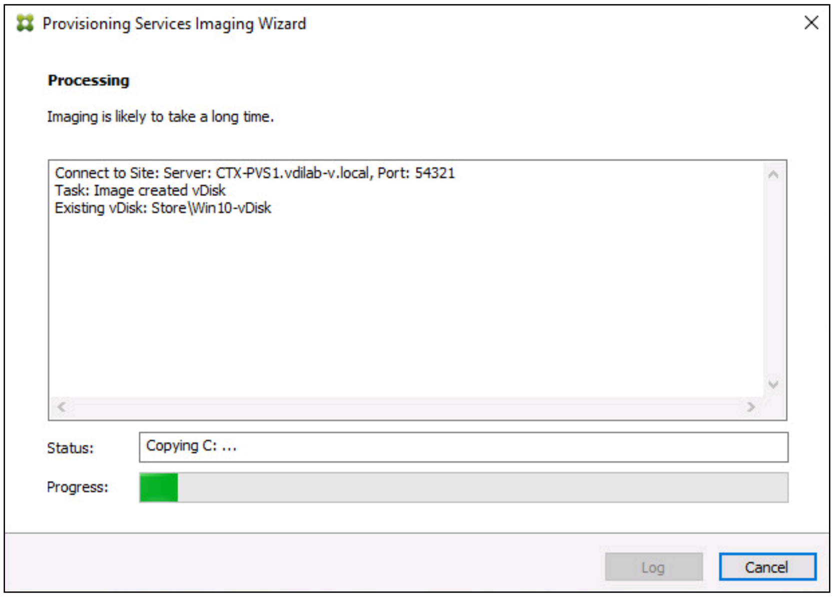

Provisioning Services Imaging Wizard
Processing
Imaging is likely to take a long time.
Log
Cancel

```
Connect to Site: Server: CTX-PVS1.vdilab-v.local, Port: 54321
Task: Image created vDisk
Existing vDisk: Store\Win10-vDisk
```

Status: Copying C: ...

Progress:

28. A message is displayed when the conversion is complete, click Done.

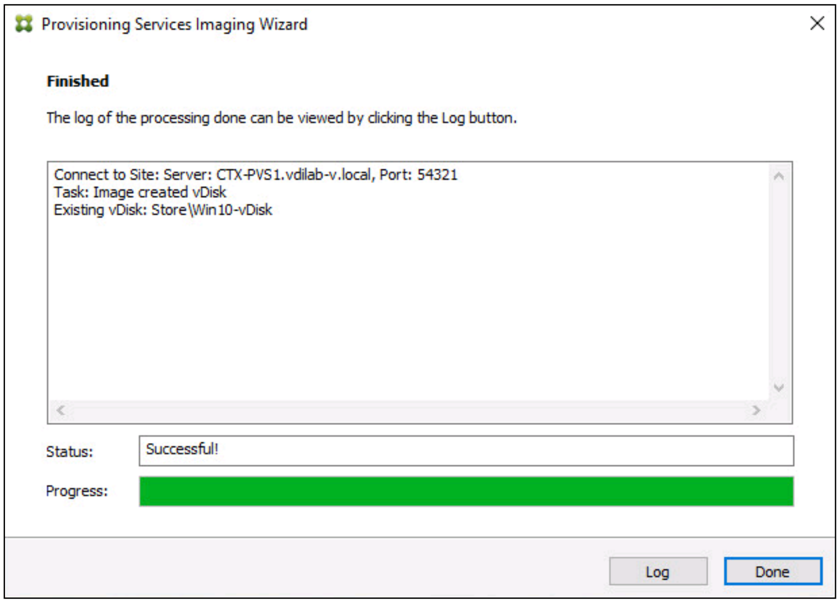

Finished

The log of the processing done can be viewed by clicking the Log button.

```
Connect to Site: Server: CTX-PVS1.vdilab-v.local, Port: 54321
Task: Image created vDisk
Existing vDisk: Store\Win10-vDisk
```

Status: Successful!

Progress:

29. Shutdown the VM used as the VDI or RDS master target.

30. Connect to the PVS server and validate that the vDisk image is available in the Store.

31. Right-click the newly created vDisk and select Properties.

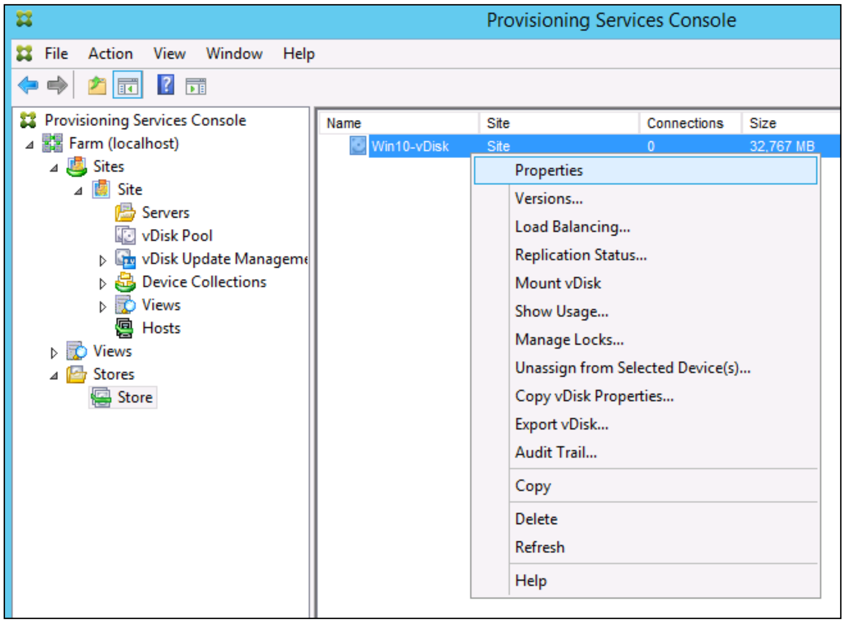

Name	Site	Connections	Size
Win10-vDisk	Site	0	32,767 MB

32. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access).”

33. Set the Cache Type to “Cache in device RAM with overflow on hard disk.”

34. Set Maximum RAM size (MBs): 64 for VDI and set 1024 MB for RDS vDisk.

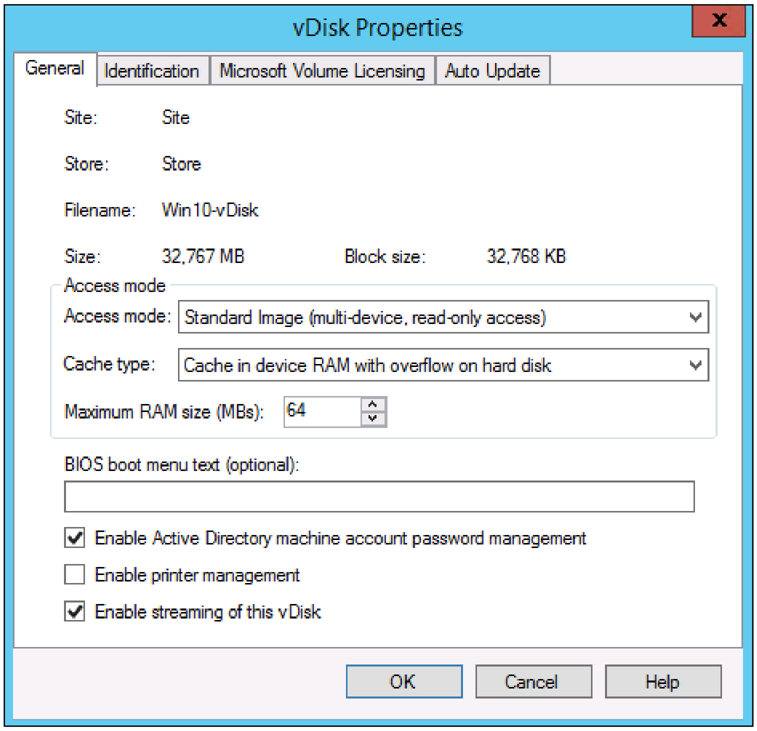

Site: Site

Store: Store

Filename: Win10-vDisk

Size: 32,767 MB Block size: 32,768 KB

Access mode

Access mode: Standard Image (multi-device, read-only 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

35. 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 2012 R2 image).

Provision Virtual Desktop Machines

To create VDI and RDS machines, complete the following steps:

1. Select the Master Target Device VM from the vSphere Client.
2. Right-click the VM and select Clone.
3. Name the cloned VM Desktop-Template.
4. Select the cluster and datastore where the first phase of provisioning will occur.

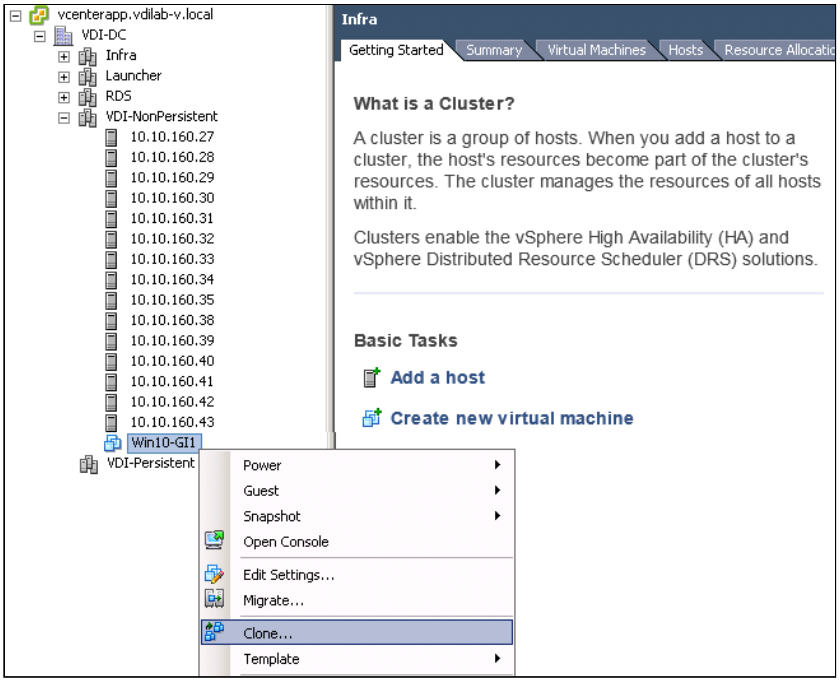

What is a Cluster?

A cluster is a group of hosts. When you add a host to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it.

Clusters enable the vSphere High Availability (HA) and vSphere Distributed Resource Scheduler (DRS) solutions.

Basic Tasks

Add a host

Create new virtual machine

5. Remove Hard disk 1 from the Template VM.

Hard disk 1 is not required to provision desktop machines as the XenDesktop Setup Wizard dynamically creates the write cache disk.

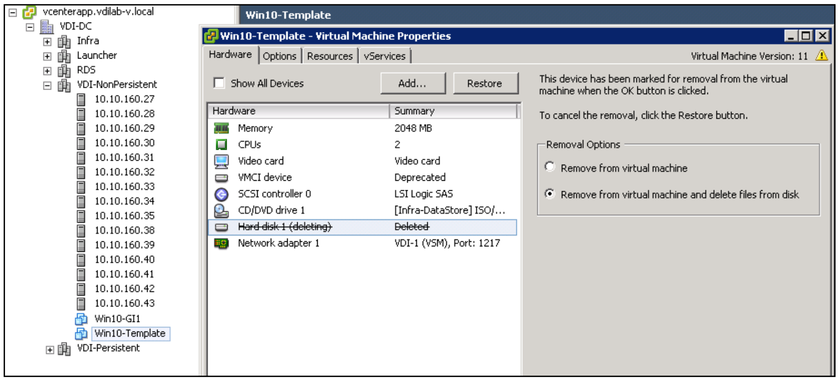

Hardware	Summary
Memory	2048 MB
CPUs	2
Video card	Video card
VMCI device	Deprecated
SCSI controller 0	LSI Logic SAS
CD/DVD drive 1	[Infra-DataStore] ISO/…
~~Hard disk 1 (deleting)~~	~~Deleted~~
Network adapter 1	VDI-1 (VSM), Port: 1217

This device has been marked for removal from the virtual machine when the OK button is clicked.

To cancel the removal, click the Restore button.

Removal Options

○ Remove from virtual machine

● Remove from virtual machine and delete files from disk

6. Convert to the Desktop-Template VM to a Template.

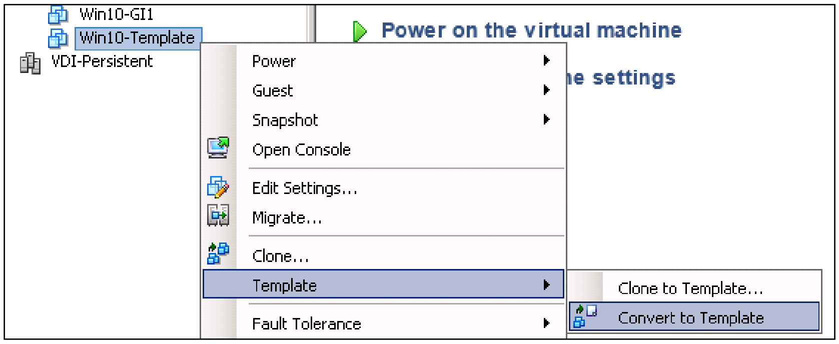

7. From Citrix Studio on the Desktop Controller, select Hosting and Add Connection and Resources.

8. Select Use an existing Connection and click Next.

9. Correspond the name of the resource with desktop machine clusters.

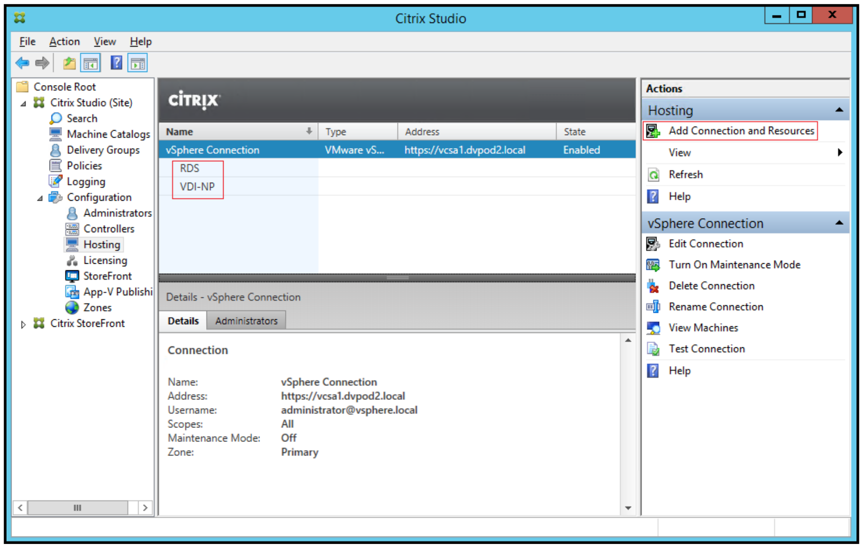

Name	Type	Address	State
vSphere Connection	VMware vS...	https://vcsa1.dvpod2.local	Enabled
RDS			
VDI-NP			

Connection

Name: vSphere Connection
Address: https://vcsa1.dvpod2.local
Username: administrator@vsphere.local
Scopes: All
Maintenance Mode: Off
Zone: Primary

10. Browse and select the vSphere clusters for desktop provisioning and use the default storage method Use storage shared by hypervisors.

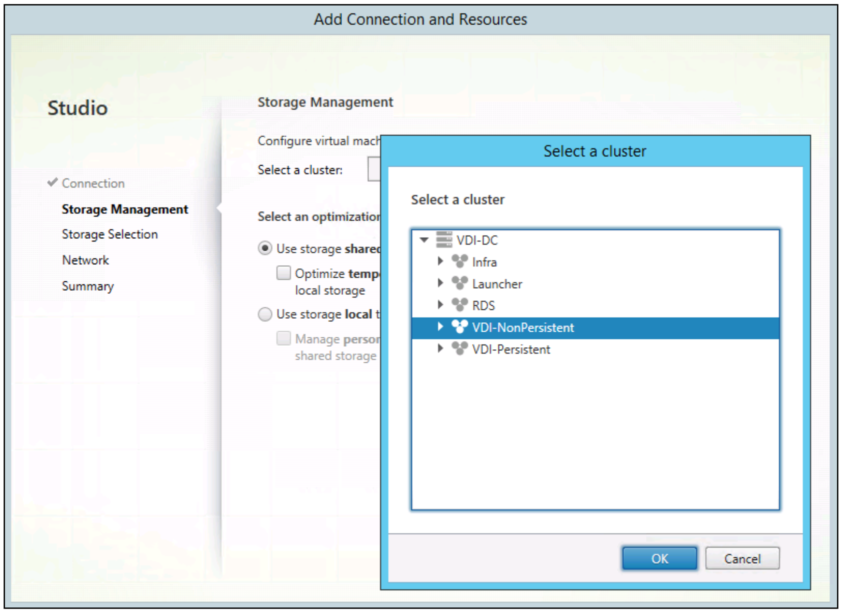

11. Select the data storage location for the corresponding resource.

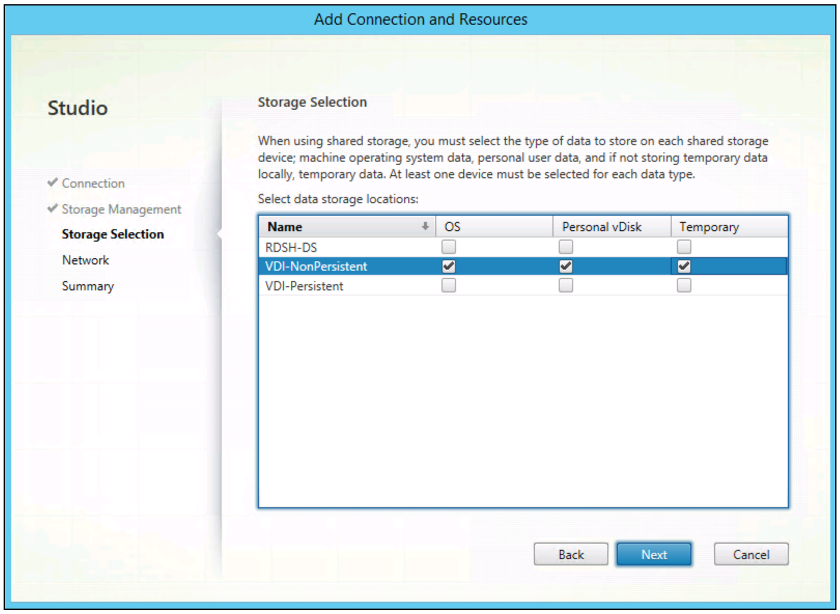

Storage Selection

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
RDSH-DS	□	□	□
VDI-NonPersistent	☑	☑	☑
VDI-Persistent	□	□	□

12. Select the VDI networks for the desktop machines and click Next.

13. Select the first datastore for desktop provisioning.

14. Click Finish.

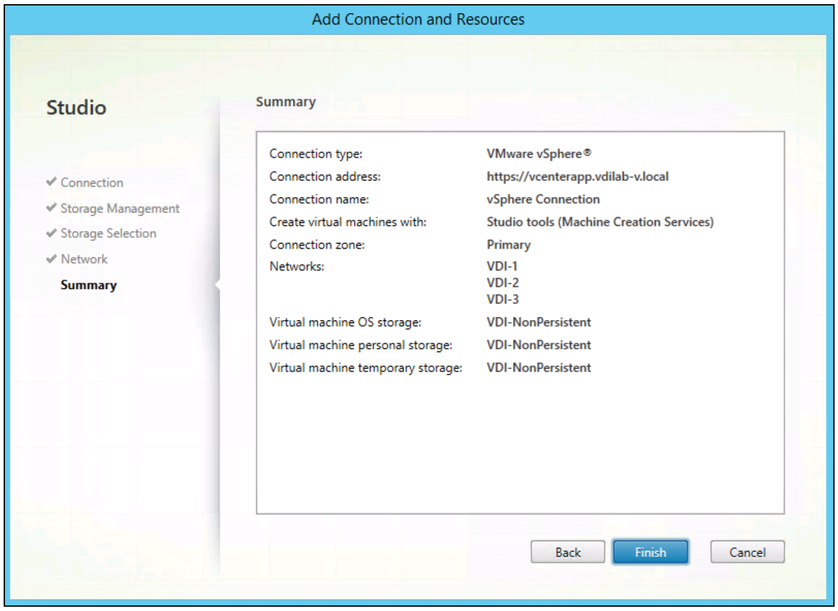

Summary

Connection type:	VMware vSphere®
Connection address:	https://vcenterapp.vdilab-v.local
Connection name:	vSphere Connection
Create virtual machines with:	Studio tools (Machine Creation Services)
Connection zone:	Primary
Networks:	VDI-1 VDI-2 VDI-3
Virtual machine OS storage:	VDI-NonPersistent
Virtual machine personal storage:	VDI-NonPersistent
Virtual machine temporary storage:	VDI-NonPersistent

Return to these settings to alter the datastore selection for each set of provisioned desktop machines.

15. Start the XenDesktop Setup Wizard from the Provisioning Services Console.

16. Right-click the Site.

17. Choose XenDesktop Setup Wizard... from the context menu.

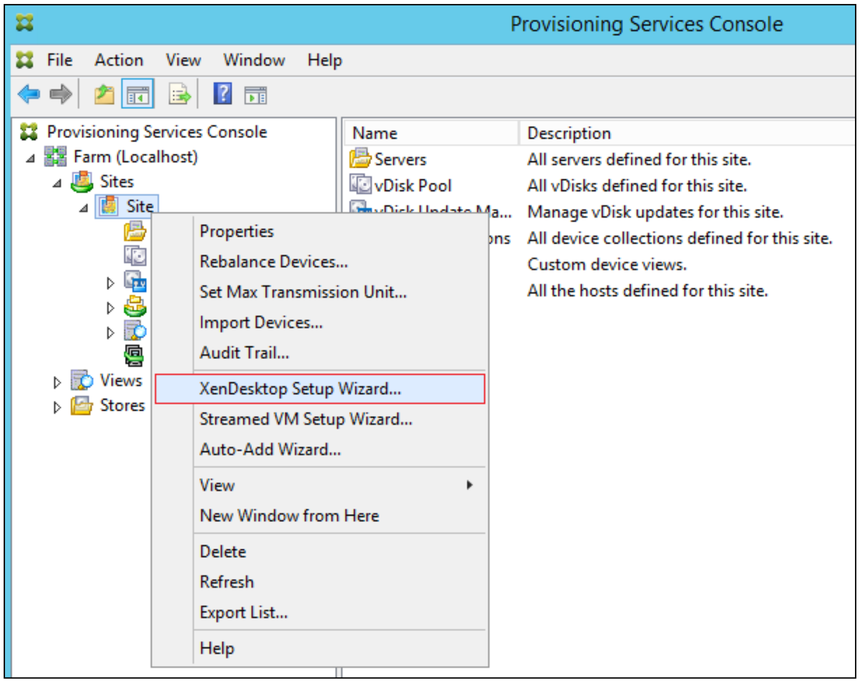

Name	Description
Servers	All servers defined for this site.
vDisk Pool	All vDisks defined for this site.
vDisk Update Ma...	Manage vDisk updates for this site.
ons	All device collections defined for this site.
	Custom device views.
	All the hosts defined for this site.

18. Click Next.

19. Enter the XenDesktop Controller address that will be used for the wizard operations.

20. Click Next.

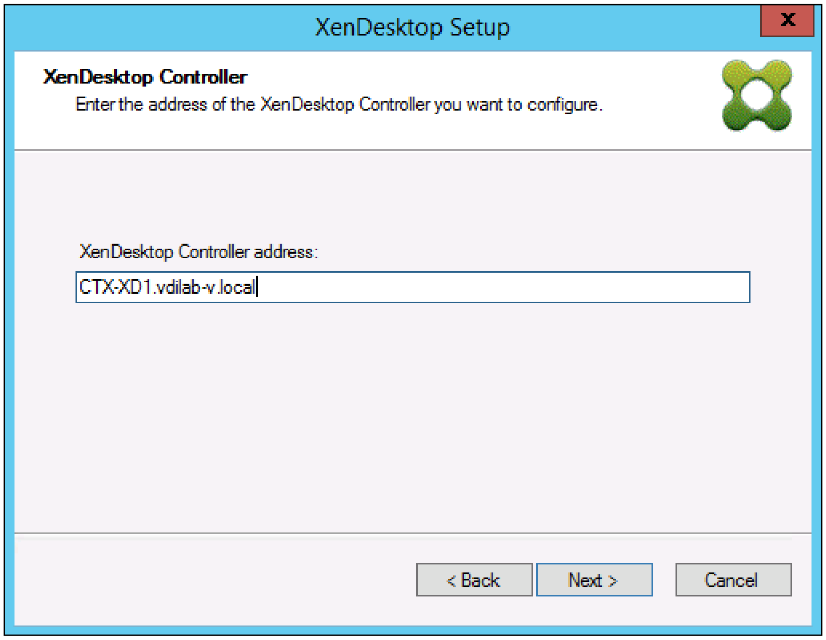

XenDesktop Controller

Enter the address of the XenDesktop Controller you want to configure.

XenDesktop Controller address: CTX-XD1.vdilab-v.local

21. Select the Host Resources on which the virtual machines will be created.

22. Click Next.

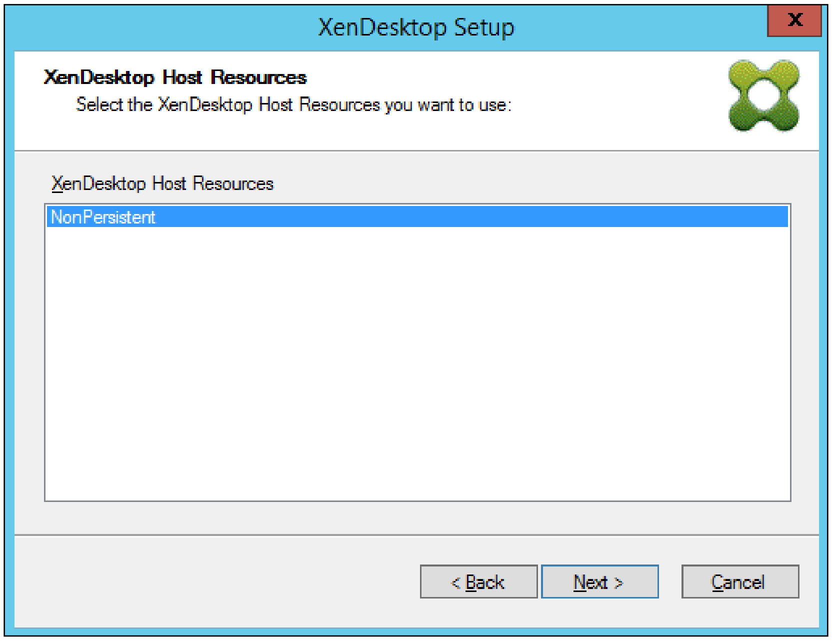

XenDesktop Setup

XenDesktop Host Resources

Select the XenDesktop Host Resources you want to use:

XenDesktop Host Resources

NonPersistent

23. Provide the Host Resources Credentials (Username and Password) to the XenDesktop controller when prompted.

24. Click OK.

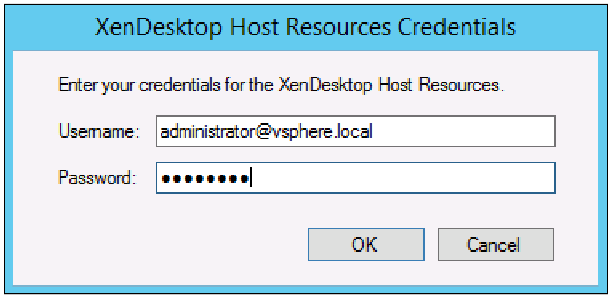

XenDesktop Host Resources Credentials

Enter your credentials for the XenDesktop Host Resources.

Username: administrator@vsphere.local

Password: ••••••••

25. Select the Template created earlier.

26. Click Next.

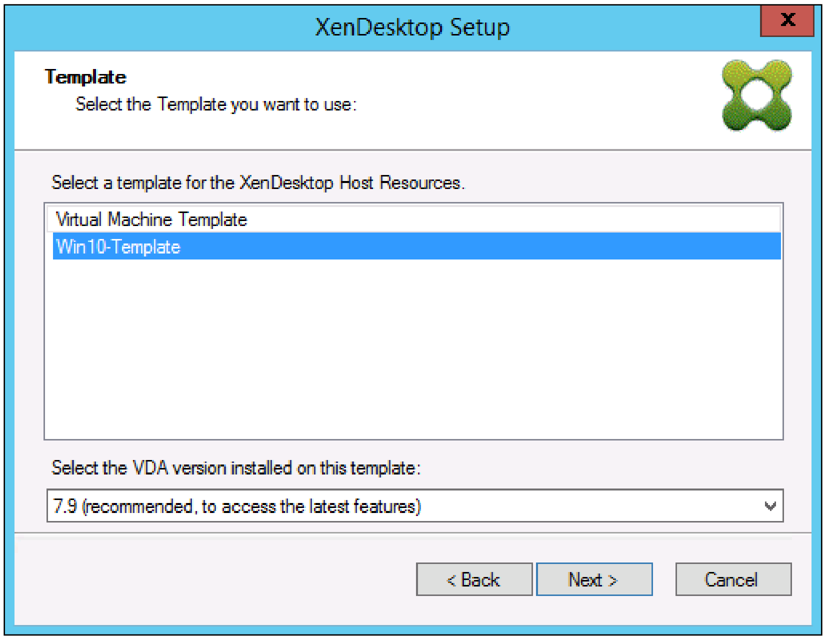

XenDesktop Setup

Template

Select the Template you want to use:

Select a template for the XenDesktop Host Resources.

- Virtual Machine Template
- Win10-Template (selected)

Select the VDA version installed on this template:

7.9 (recommended, to access the latest features)

27. Select the network that will be used for the provisioned virtual machines.

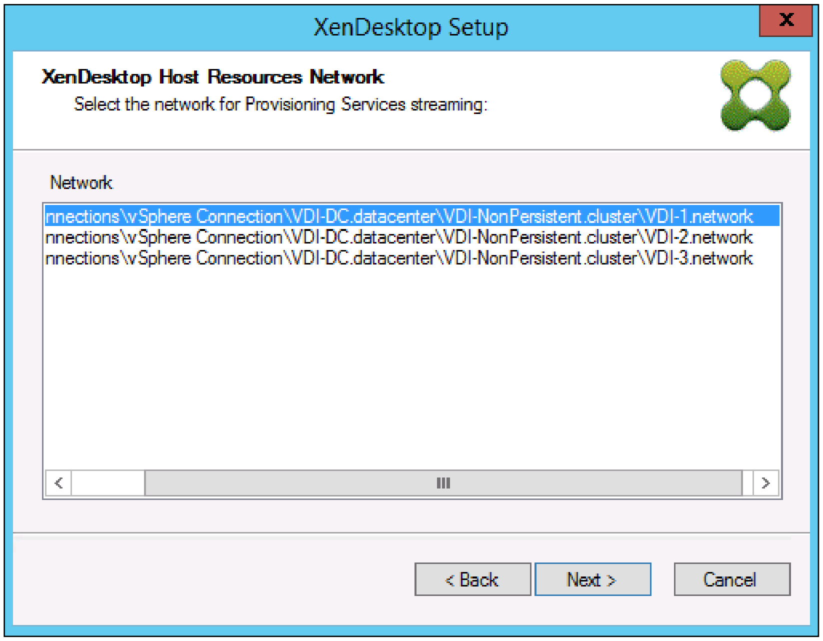

XenDesktop Setup

XenDesktop Host Resources Network

Select the network for Provisioning Services streaming:

Network

```
nnections\vSphere Connection\VDI-DC.datacenter\VDI-NonPersistent.cluster\VDI-1.network
nnections\vSphere Connection\VDI-DC.datacenter\VDI-NonPersistent.cluster\VDI-2.network
nnections\vSphere Connection\VDI-DC.datacenter\VDI-NonPersistent.cluster\VDI-3.network
```

A single VLAN was created for the VDI and RDS VMs, however, the Nexus 1000V is limited to 1024 ports per interface. Three port-profiles where created to accommodate this CVD.

28. Select the vDisk that will be used to stream virtual machines.

29. Click Next.

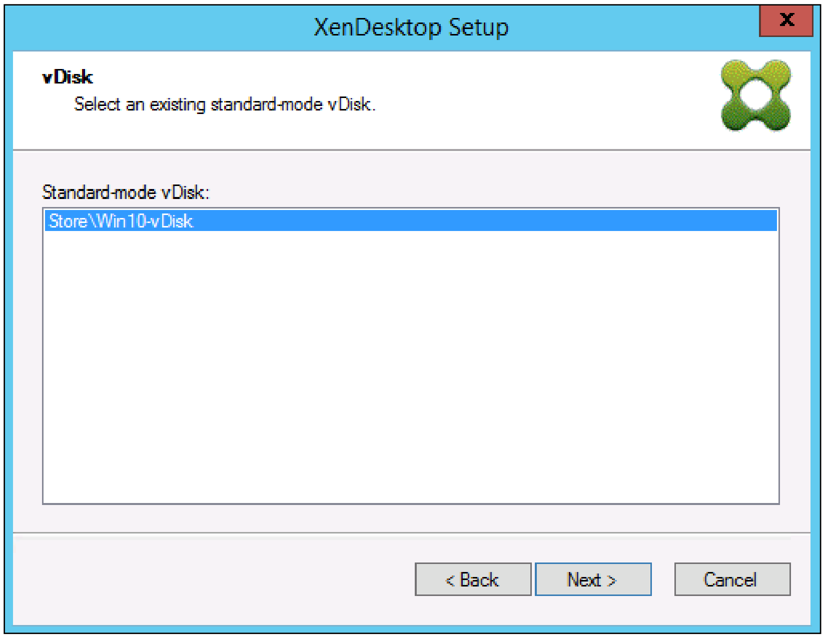

vDisk
Select an existing standard-mode vDisk.

Standard-mode vDisk:
Store\Win10-vDisk

30. Select “Create a new catalog.”

The catalog name is also used as the collection name in the PVS site.

31. Click Next.

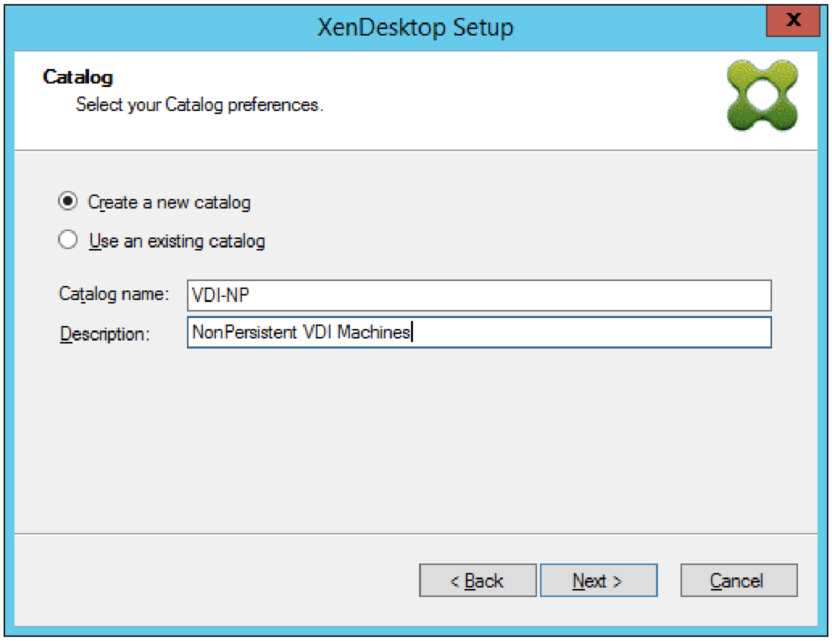

Catalog
Select your Catalog preferences.

● Create a new catalog
○ Use an existing catalog

Catalog name: VDI-NP
Description: NonPersistent VDI Machines

32. 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.

33. Click Next.

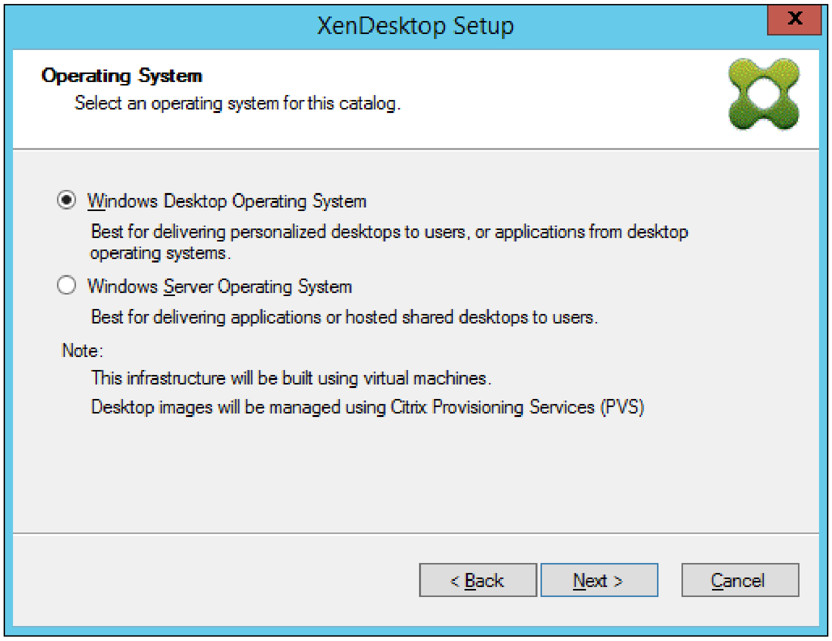

XenDesktop Setup

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)

34. 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.”

35. Click Next.

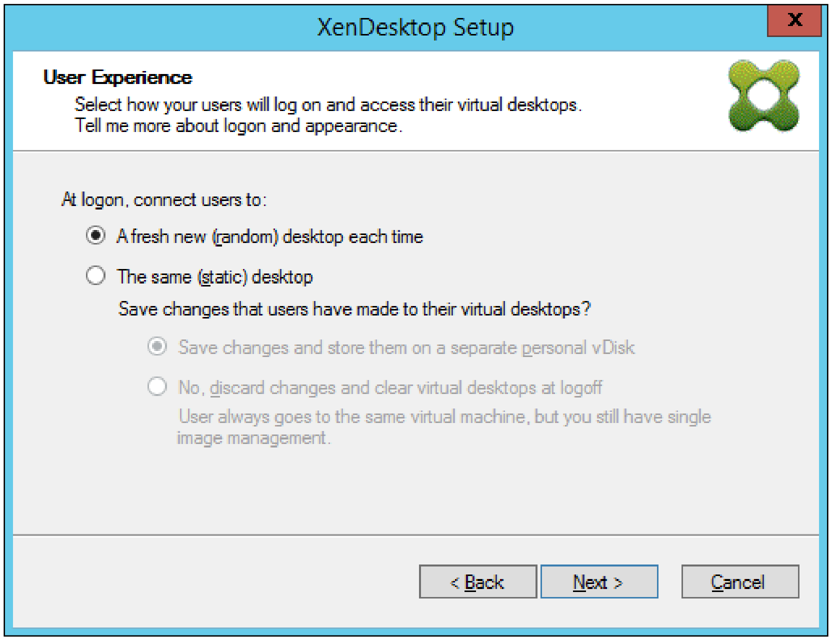

XenDesktop Setup

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.

36. 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 VDI, 6 for RDS)
- The amount of memory for the VM (1.7GB for VDI, 24GB for RDS)
- The write-cache disk size (10GB for VDI, 30GB for RDS)
- PXE boot as the Boot Mode

37. Click Next.

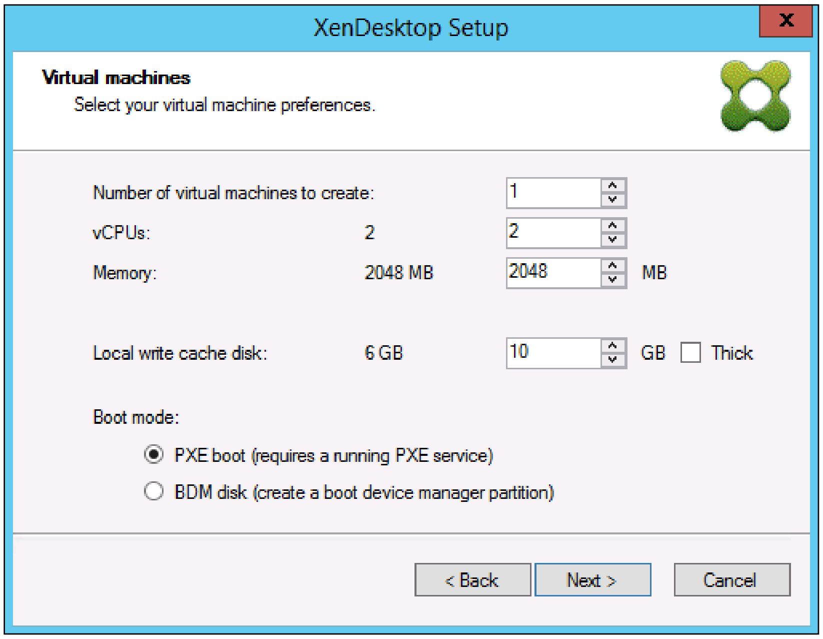

XenDesktop Setup

Virtual machines

Select your virtual machine preferences.

Number of virtual machines to create:		1
vCPUs:	2	2
Memory:	2048 MB	2048 MB
Local write cache disk:	6 GB	10 GB □ Thick

Boot mode:

● PXE boot (requires a running PXE service)

○ BDM disk (create a boot device manager partition)

38. Select the Create new accounts radio button.

39. Click Next.

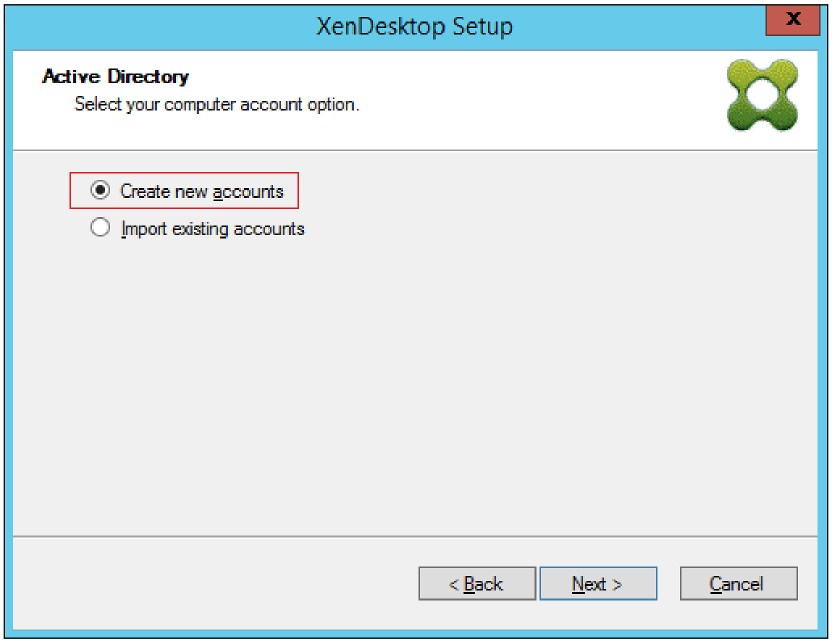

XenDesktop Setup

Active Directory

Select your computer account option.

● Create new accounts

○ Import existing accounts

40. Specify the Active Directory Accounts and Location. This is where the wizard should create the computer accounts.

41. Provide the Account naming scheme. An example name is shown in the text box below the name scheme selection location.

42. Click Next.

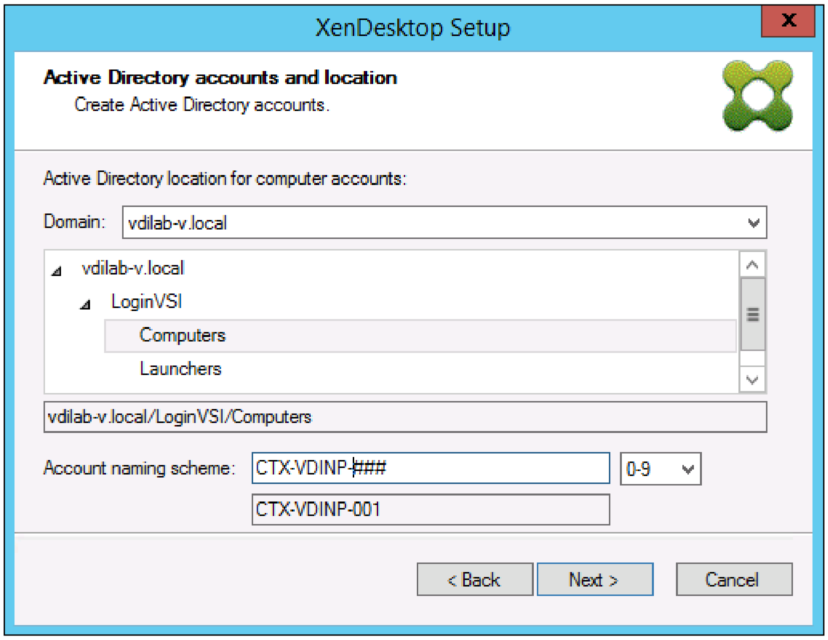

XenDesktop Setup

Active Directory accounts and location

Create Active Directory accounts.

Active Directory location for computer accounts:

Domain: vdilab-v.local

- vdilab-v.local
  - LoginVSI
    - Computers
    - Launchers

vdilab-v.local/LoginVSI/Computers

Account naming scheme: CTX-VDINP-### 0-9

CTX-VDINP-001

43. Click Finish to begin the virtual machine creation.

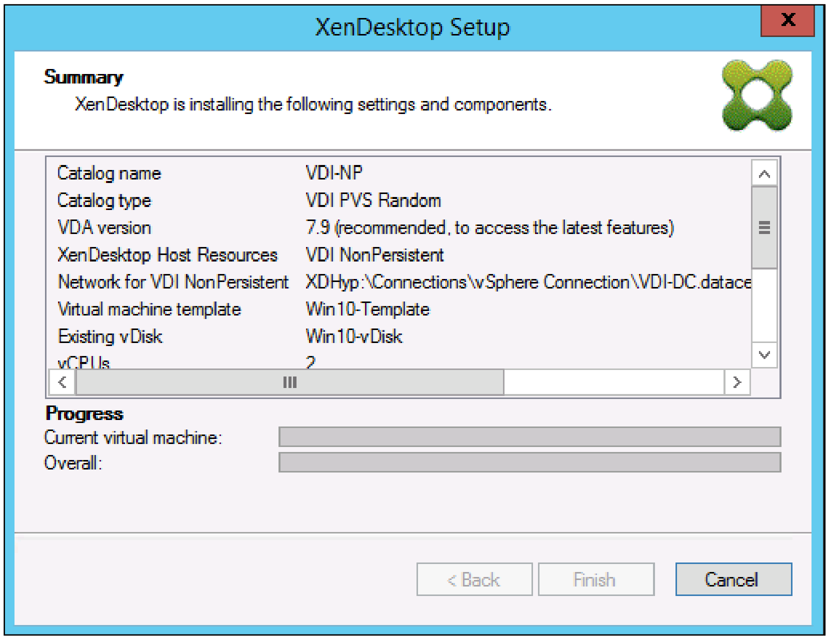

Summary

XenDesktop is installing the following settings and components.

Catalog name	VDI-NP
Catalog type	VDI PVS Random
VDA version	7.9 (recommended, to access the latest features)
XenDesktop Host Resources	VDI NonPersistent
Network for VDI NonPersistent	XDHyp:\Connections\vSphere Connection\VDI-DC.datace
Virtual machine template	Win10-Template
Existing vDisk	Win10-vDisk
vCPUs	2

Progress

Current virtual machine:

Overall:

44. When the wizard is done provisioning the virtual machines, click Done.

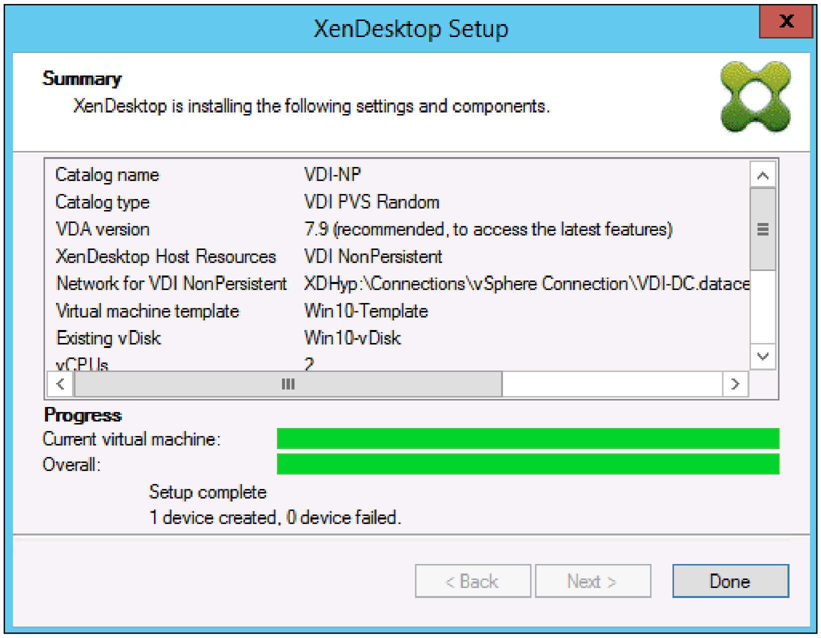

Summary

XenDesktop is installing the following settings and components.

Catalog name	VDI-NP
Catalog type	VDI PVS Random
VDA version	7.9 (recommended, to access the latest features)
XenDesktop Host Resources	VDI NonPersistent
Network for VDI NonPersistent	XDHyp:\Connections\vSphere Connection\VDI-DC.datace
Virtual machine template	Win10-Template
Existing vDisk	Win10-vDisk
vCPUs	2

Progress

Current virtual machine:

Overall:

Setup complete

1 device created, 0 device failed.

Provisioning process takes ~10 seconds per machine.

45. Verify the desktop machines were successfully created in the following locations:

- PVS1 > Provisioning Services Console > Farm > Site > Device Collections > VDI-NP > CTX-VDI-001

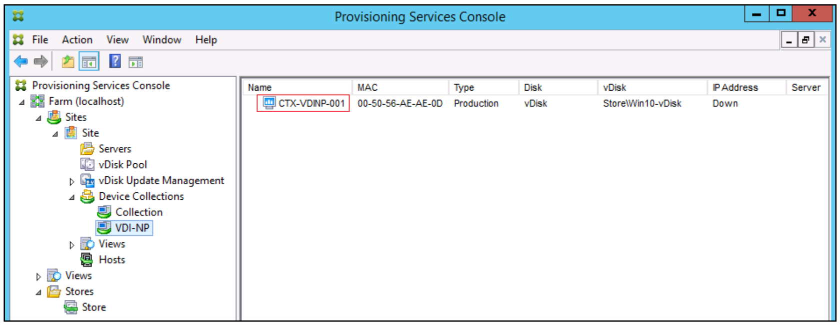

Name	MAC	Type	Disk	vDisk	IP Address	Server
CTX-VDINP-001	00-50-56-AE-AE-0D	Production	vDisk	Store\Win10-vDisk	Down	

— CTX-XD1 > Citrix Studio > Machine Catalogs > VDI-NP

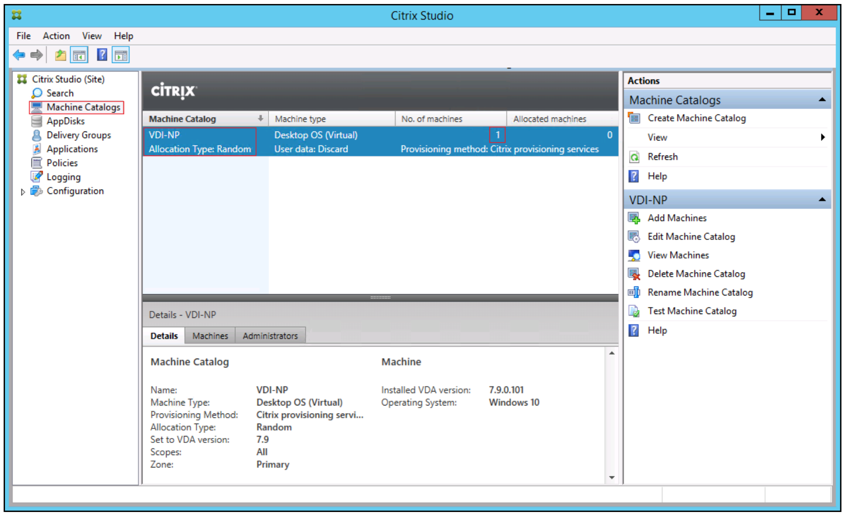

Machine Catalog	Machine type	No. of machines	Allocated machines
VDI-NP Allocation Type: Random	Desktop OS (Virtual) User data: Discard	1 Provisioning method: Citrix provisioning services	0

Machine Catalog

Name: VDI-NP
Machine Type: Desktop OS (Virtual)
Provisioning Method: Citrix provisioning servi...
Allocation Type: Random
Set to VDA version: 7.9
Scopes: All
Zone: Primary

Machine

Installed VDA version: 7.9.0.101
Operating System: Windows 10

— AD-DC1 > Active Directory Users and Computers > dvpod2.local > Computers > CTX-VDI-001

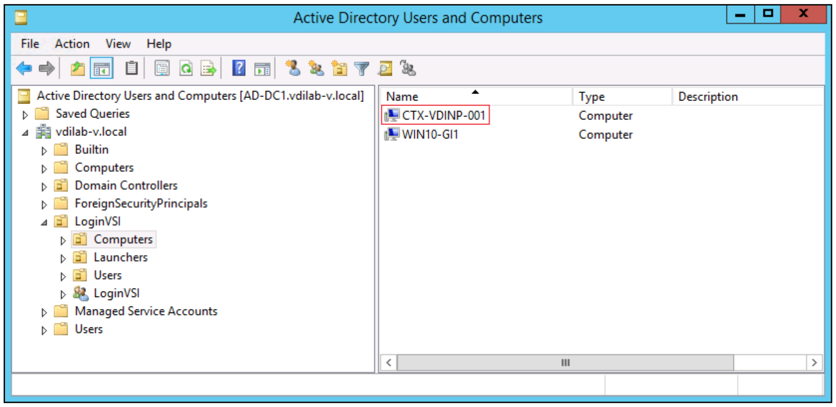

Name	Type	Description
CTX-VDINP-001	Computer	
WIN10-GI1	Computer	

46. Logon to newly provisioned desktop machine, using the Virtual Disk Status verify the image mode is set to Ready Only and the cache type as Device Ram with overflow on local hard drive.

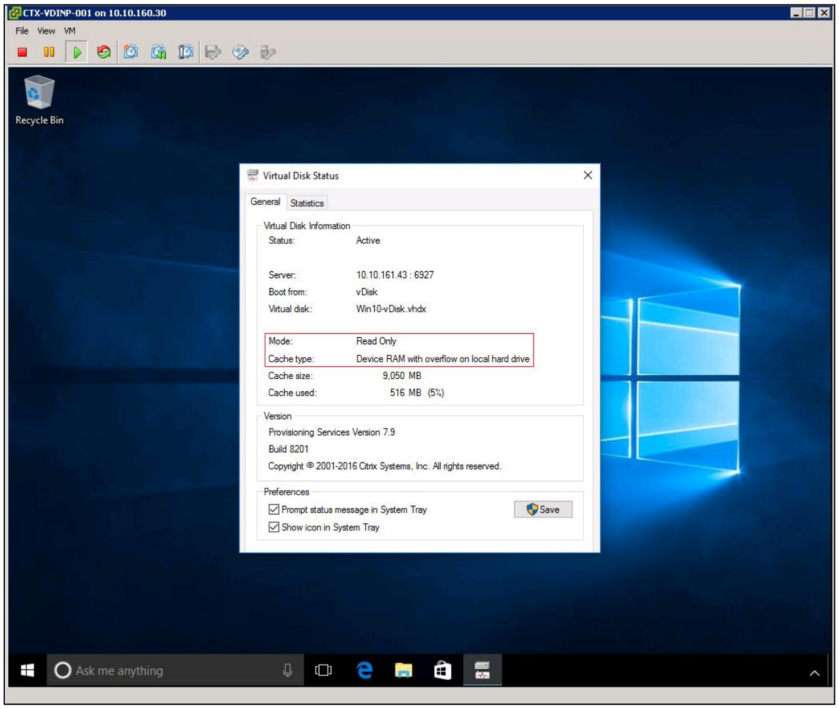

Virtual Disk Status

General | Statistics

Virtual Disk Information

Status: Active

Server: 10.10.161.43 : 6927
Boot from: vDisk
Virtual disk: Win10-vDisk.vhdx

Mode: Read Only
Cache type: Device RAM with overflow on local hard drive
Cache size: 9,050 MB
Cache used: 516 MB (5%)

Version

Provisioning Services Version 7.9
Build 8201
Copyright © 2001-2016 Citrix Systems, Inc. All rights reserved.

Preferences

☑ Prompt status message in System Tray
☑ Show icon in System Tray
Save

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 menu.

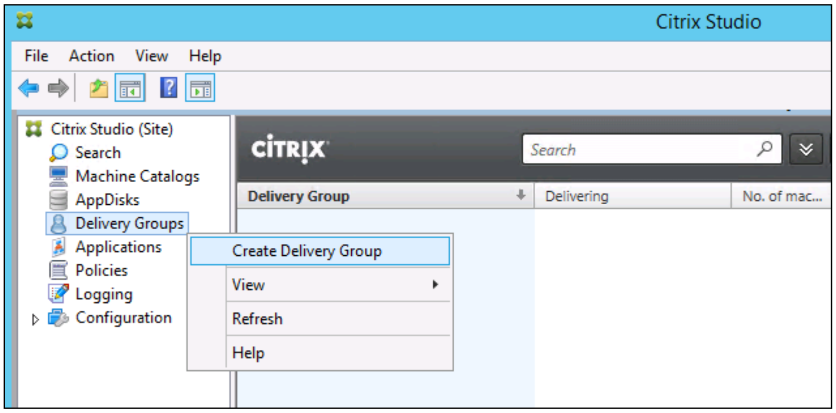

3. Specify the Machine Catalog and increment the number of machines to add.

4. Click Next.

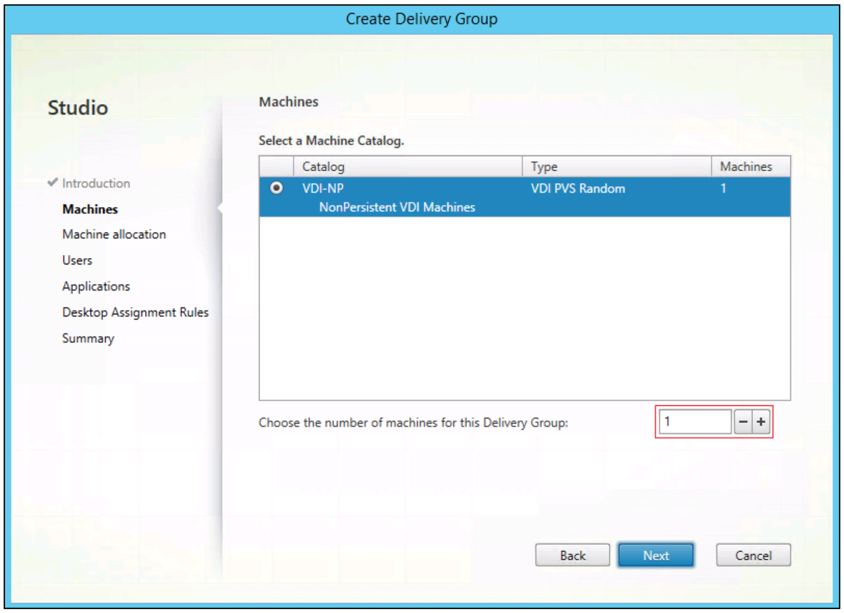

Create Delivery Group

Studio

Machines

Select a Machine Catalog.

	Catalog	Type	Machines
●	VDI-NP NonPersistent VDI Machines	VDI PVS Random	1

Choose the number of machines for this Delivery Group: 1

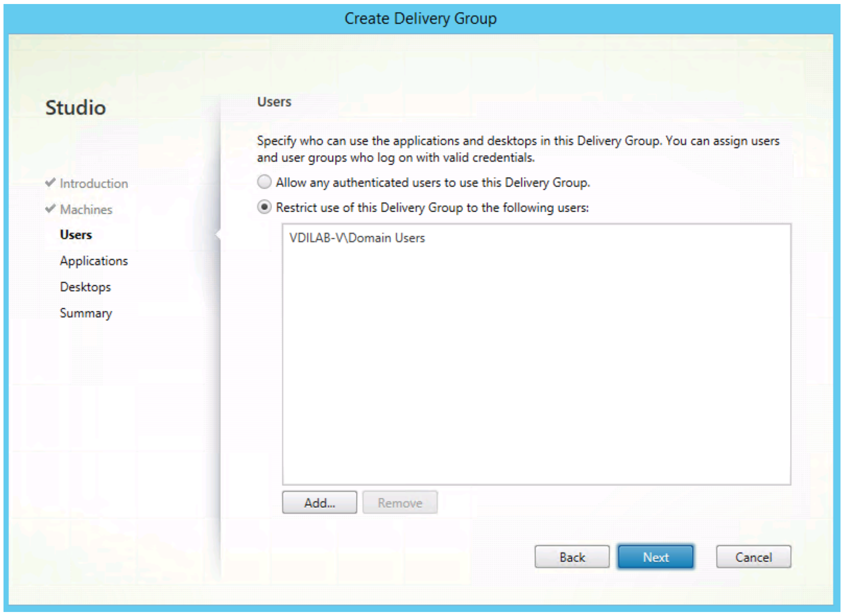

Create Delivery Group

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:

VDILAB-V\Domain Users

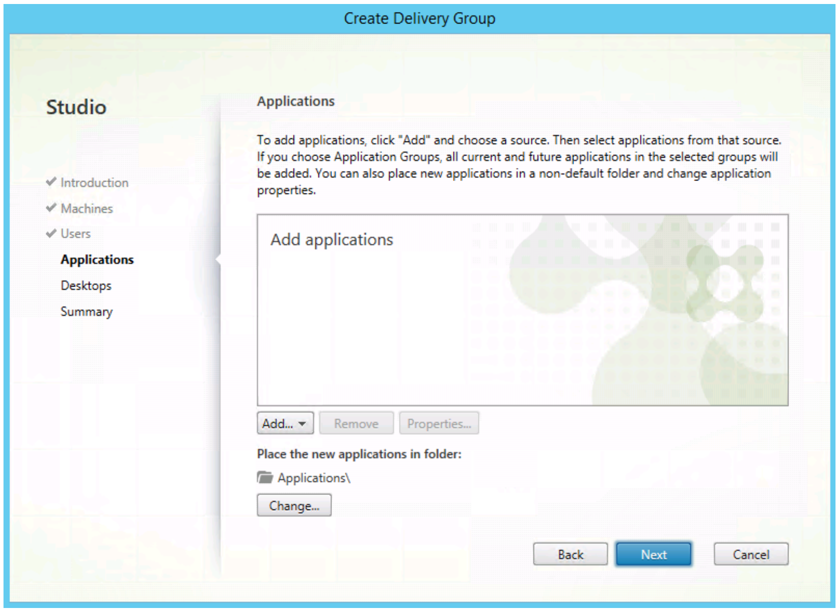

Create Delivery Group

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\

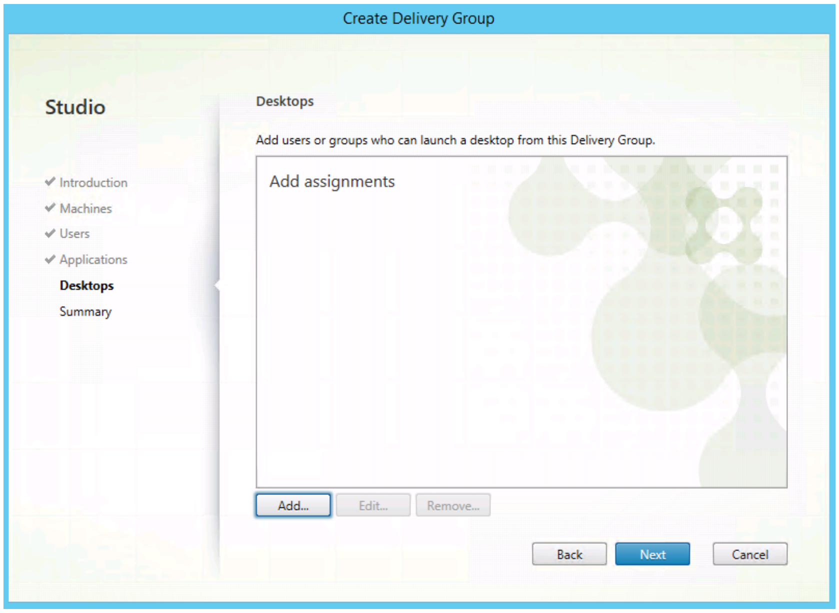

Create Delivery Group

Studio

Desktops

Add users or groups who can launch a desktop from this Delivery Group.

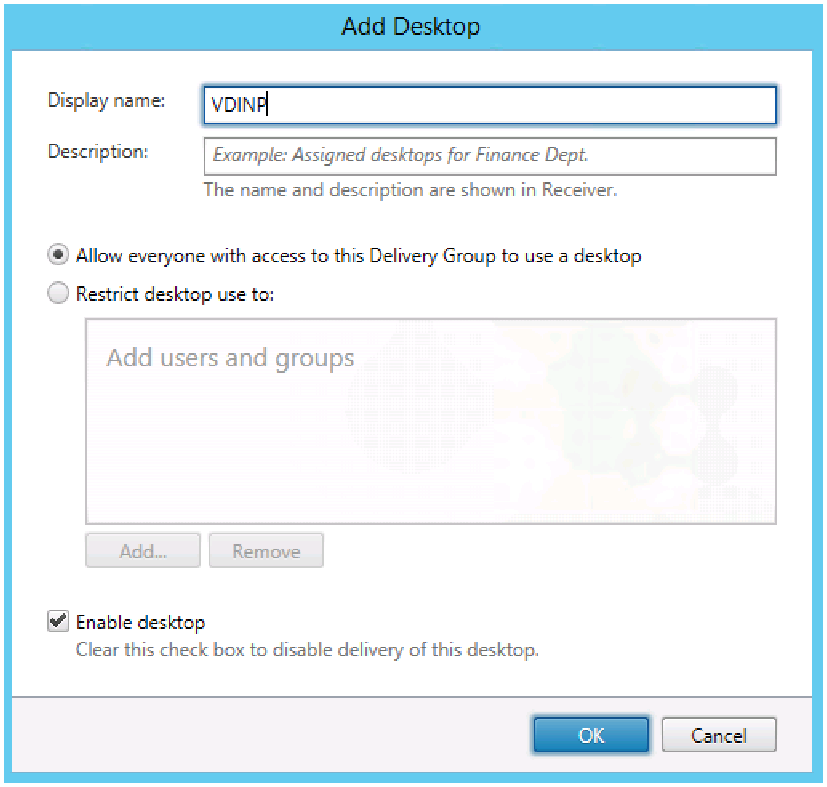

Add Desktop

Display name: VDINP

Description: *Example: Assigned desktops for Finance Dept.*

The name and description are shown in Receiver.

● Allow everyone with access to this Delivery Group to use a desktop

○ Restrict desktop use to:

Add users and groups

☑ Enable desktop

Clear this check box to disable delivery of this desktop.

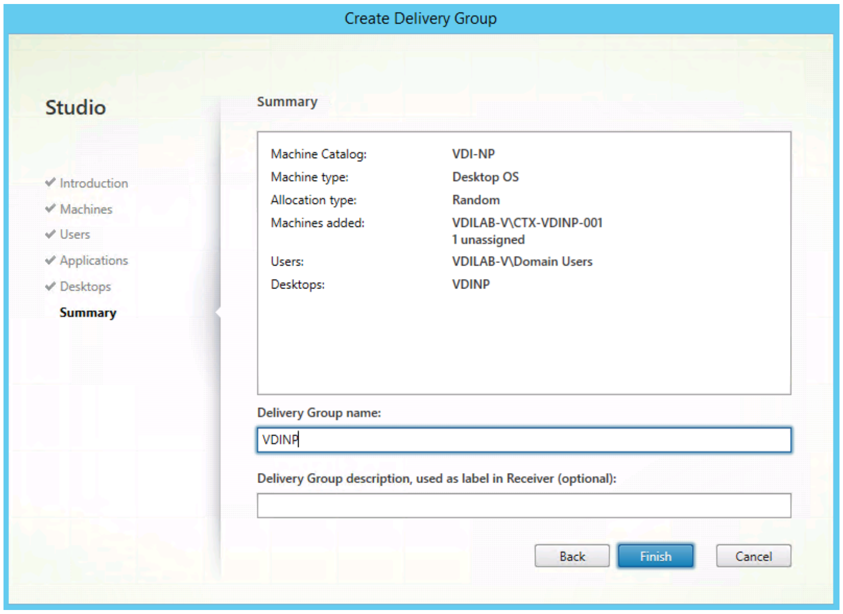

Create Delivery Group

Studio

- ✓ Introduction
- ✓ Machines
- ✓ Users
- ✓ Applications
- ✓ Desktops
- **Summary**

Summary

Machine Catalog:	VDI-NP
Machine type:	Desktop OS
Allocation type:	Random
Machines added:	VDILAB-V\CTX-VDINP-001 1 unassigned
Users:	VDILAB-V\Domain Users
Desktops:	VDINP

Delivery Group name: VDINP

Delivery Group description, used as label in Receiver (optional):

5. Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.

6. Select Desktops.

7. Click Next.

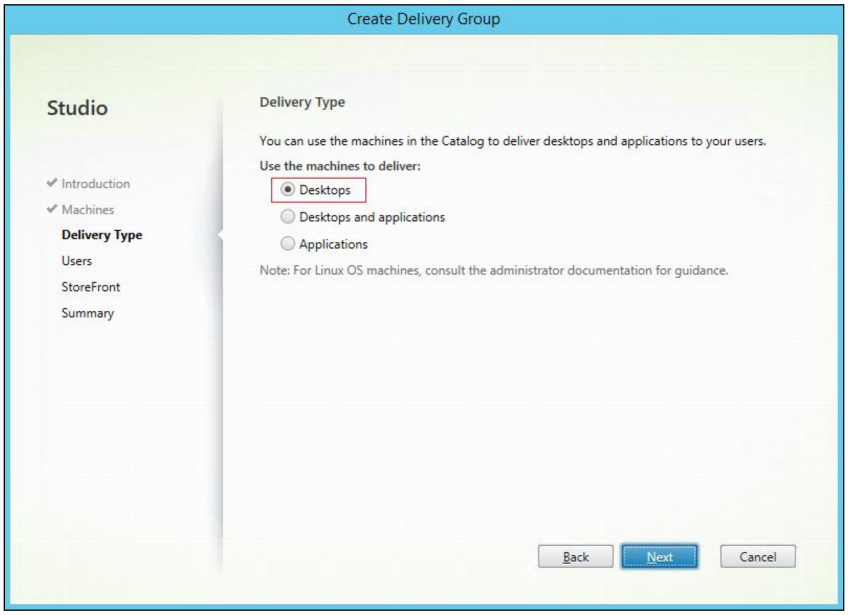

Delivery Type

You can use the machines in the Catalog to deliver desktops and applications to your users.

Use the machines to deliver:

● Desktops
○ Desktops and applications
○ Applications

Note: For Linux OS machines, consult the administrator documentation for guidance.

8. To make the Delivery Group accessible, you must add users, click Add…

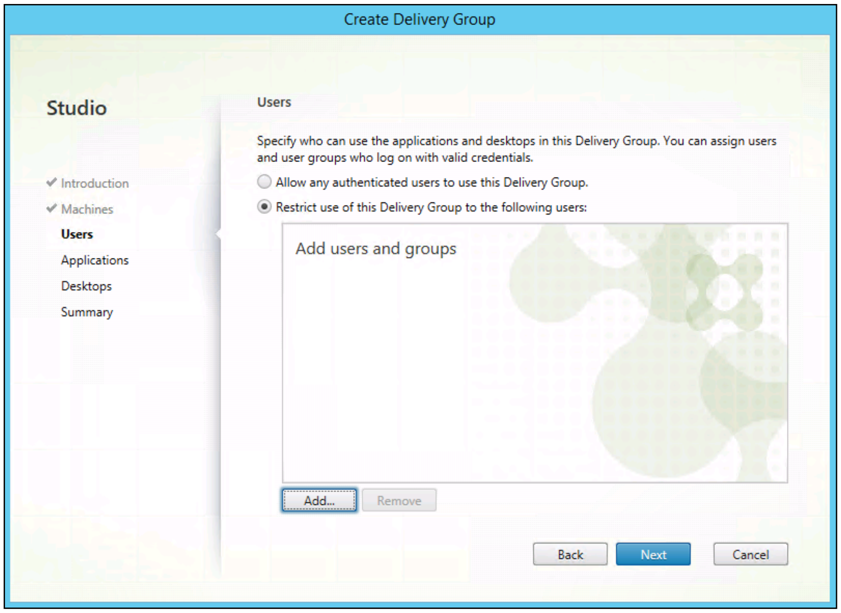

Users

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:

9. In the Select Users or Groups dialog, add users or groups.

10. Click OK. When users have been added, click Next on the Assign dialog (shown above).

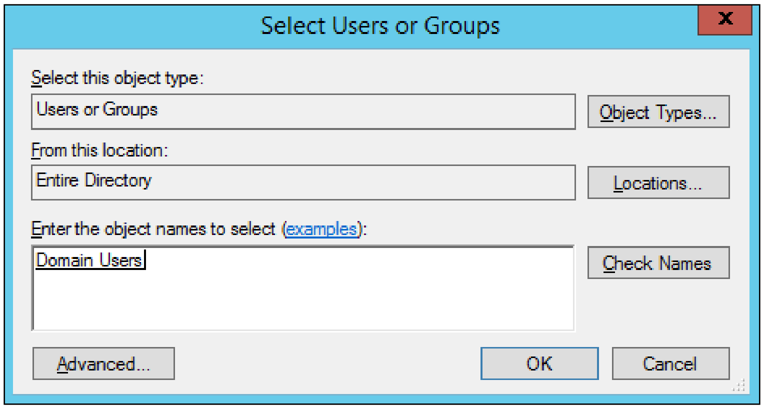

Select Users or Groups

Select this object type:

Users or Groups

From this location:

Entire Directory

Enter the object names to select (examples):

Domain Users

11. Enter the StoreFront configuration for how Receiver will be installed on the machines in this Delivery Group. Click “Manually, using a StoreFront server address that I will provide later.”

12. Click Next.

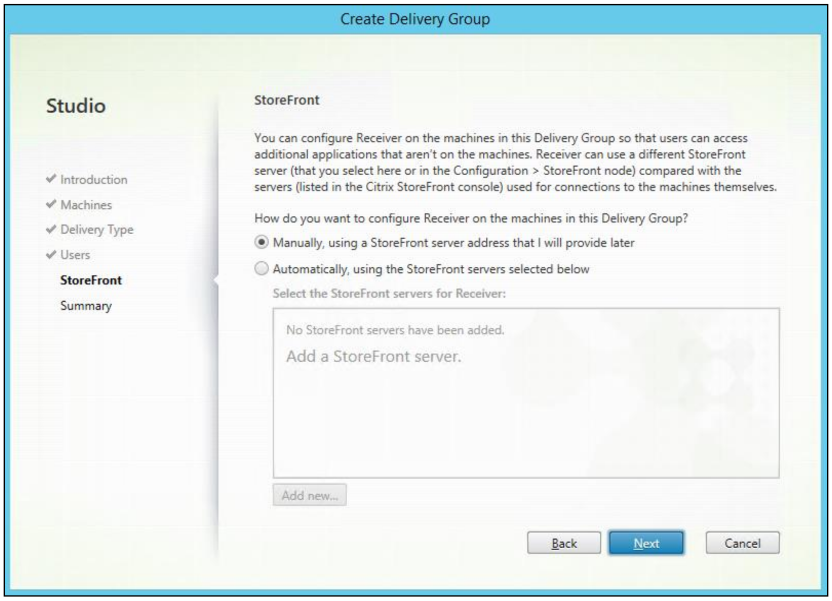

Create Delivery Group

Studio

- ✓ Introduction
- ✓ Machines
- ✓ Delivery Type
- ✓ Users
- **StoreFront**
- Summary

StoreFront

You can configure Receiver on the machines in this Delivery Group so that users can access additional applications that aren't on the machines. Receiver can use a different StoreFront server (that you select here or in the Configuration > StoreFront node) compared with the servers (listed in the Citrix StoreFront console) used for connections to the machines themselves.

How do you want to configure Receiver on the machines in this Delivery Group?

● Manually, using a StoreFront server address that I will provide later

○ Automatically, using the StoreFront servers selected below

Select the StoreFront servers for Receiver:

No StoreFront servers have been added.

Add a StoreFront server.

13. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Display name (for example, VDI or RDS).

14. Click Finish.

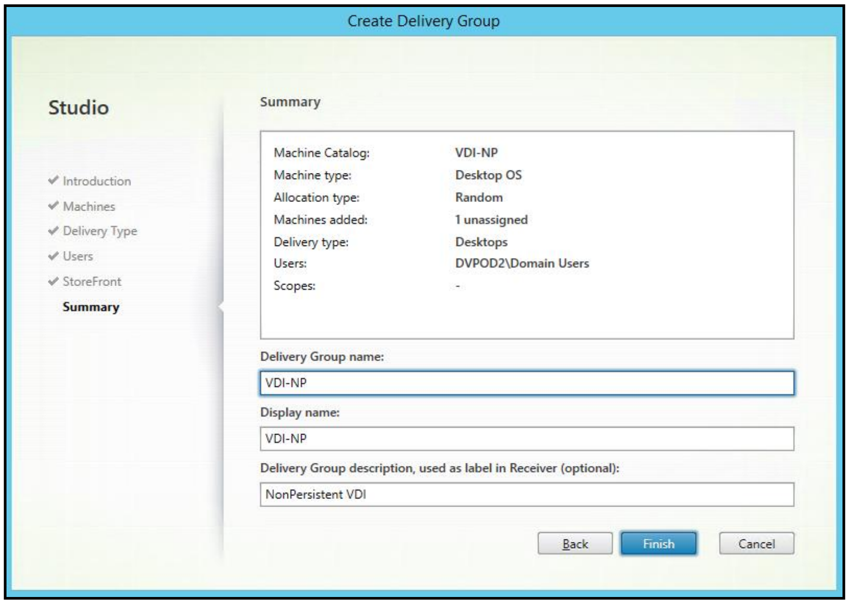

Summary

Machine Catalog:	VDI-NP
Machine type:	Desktop OS
Allocation type:	Random
Machines added:	1 unassigned
Delivery type:	Desktops
Users:	DVPOD2\Domain Users
Scopes:	-

Delivery Group name: VDI-NP

Display name: VDI-NP

Delivery Group description, used as label in Receiver (optional): NonPersistent VDI

15. Citrix Studio lists the created Delivery Groups and the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab.

16. On the pull-down menu, select “Turn on Maintenance Mode.”

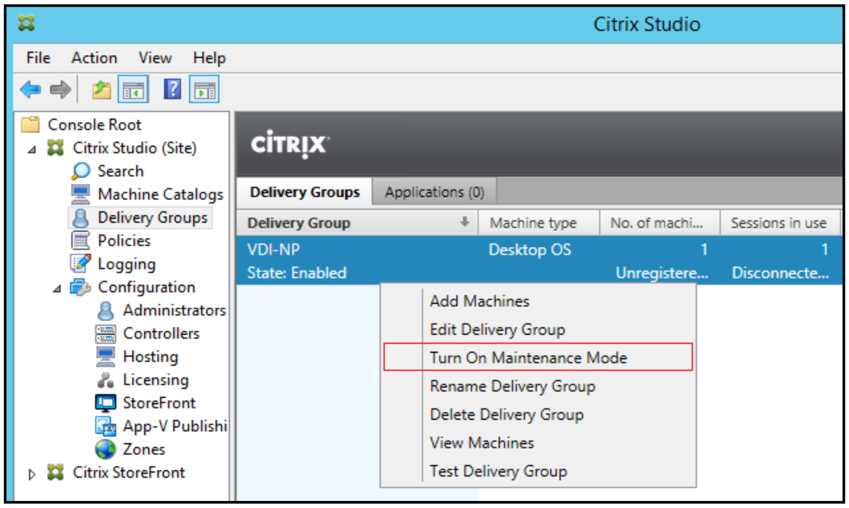

Delivery Group	Machine type	No. of machi...	Sessions in use
VDI-NP State: Enabled	Desktop OS	1 Unregistere...	1 Disconnecte...

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 125 XenDesktop Policy

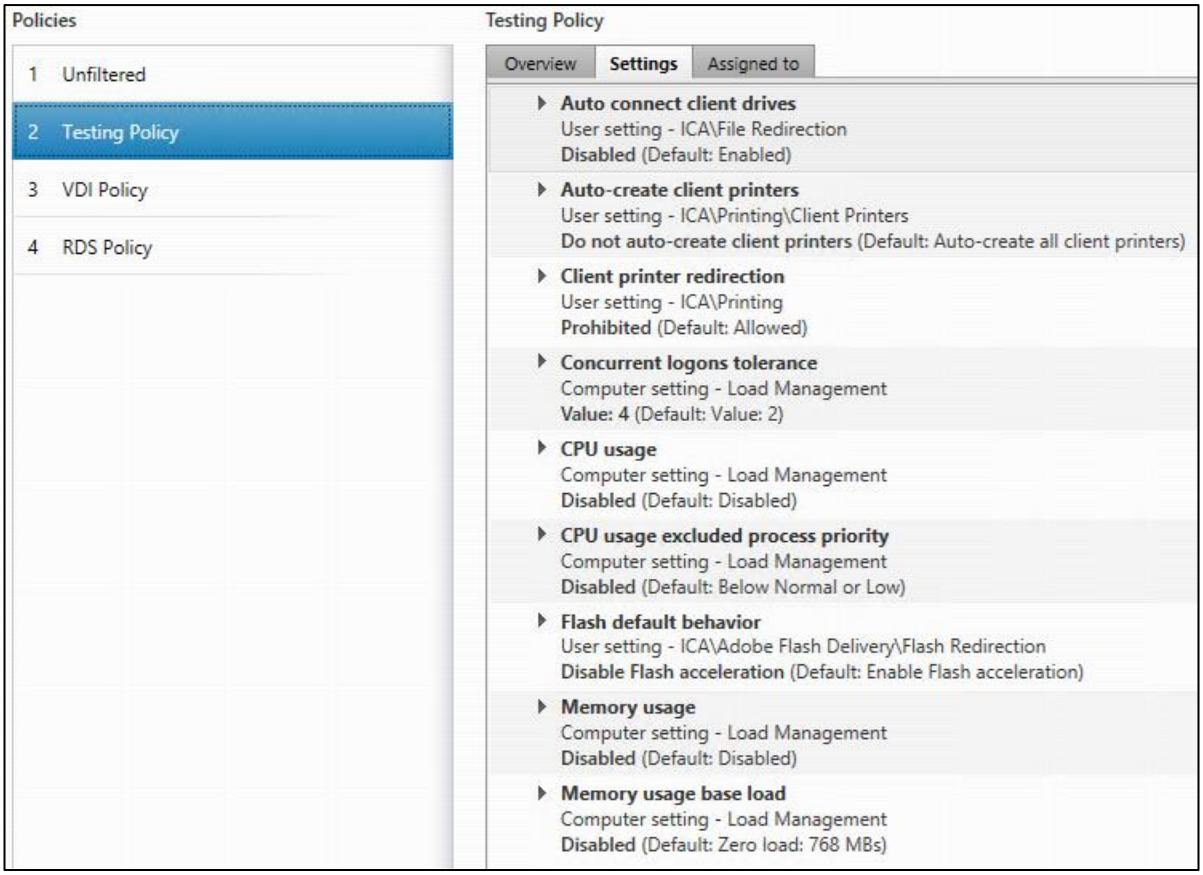

Auto connect client drives
User setting - ICA\File Redirection
Disabled (Default: Enabled)

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)

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:

http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-7.html

Figure 126 VDI User Profile Manager Policy

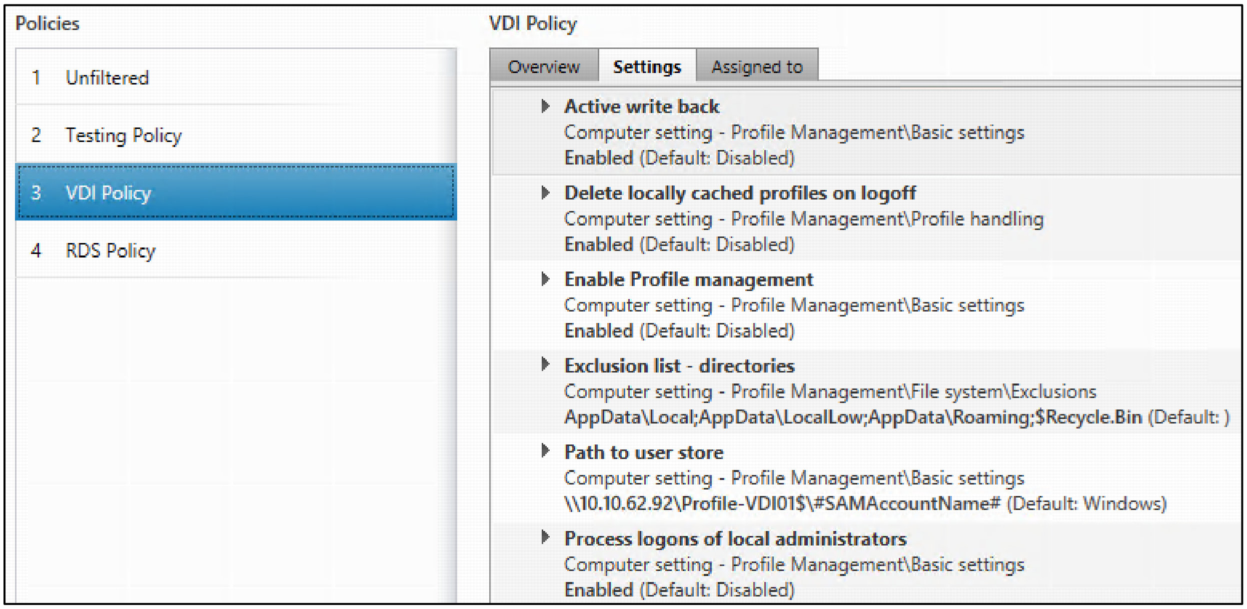

Active write back
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

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 127 RDS User Profile Manager Policy

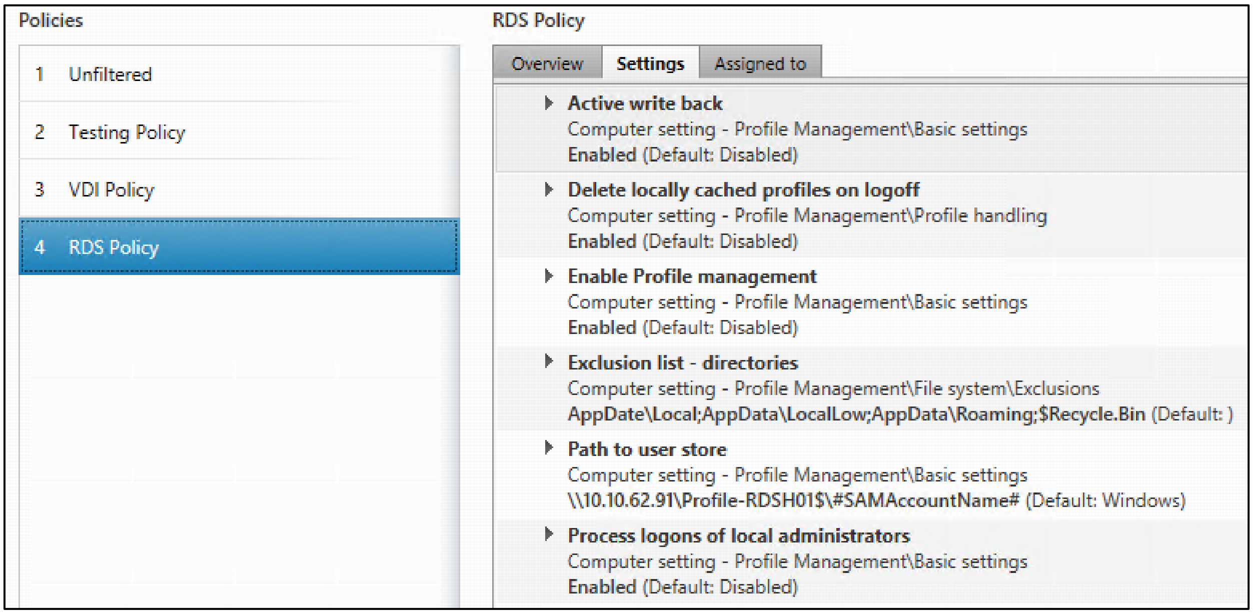

1 Unfiltered

2 Testing Policy

3 VDI Policy

4 RDS 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 M6 Card

This section focuses on installing and configuring the NVIDIA M6 cards with the Cisco UCS B200 M4 servers to deploy vGPU enabled virtual desktops.

Physical Install of M6 Card into B200 M4 Server

The NVIDIA M6 graphics processing unit (GPU) provides graphics and computing capabilities to the server. The GPU package consists of the three elements shown in the following figure.

Figure 128 NVIDIA M6 GPU Package

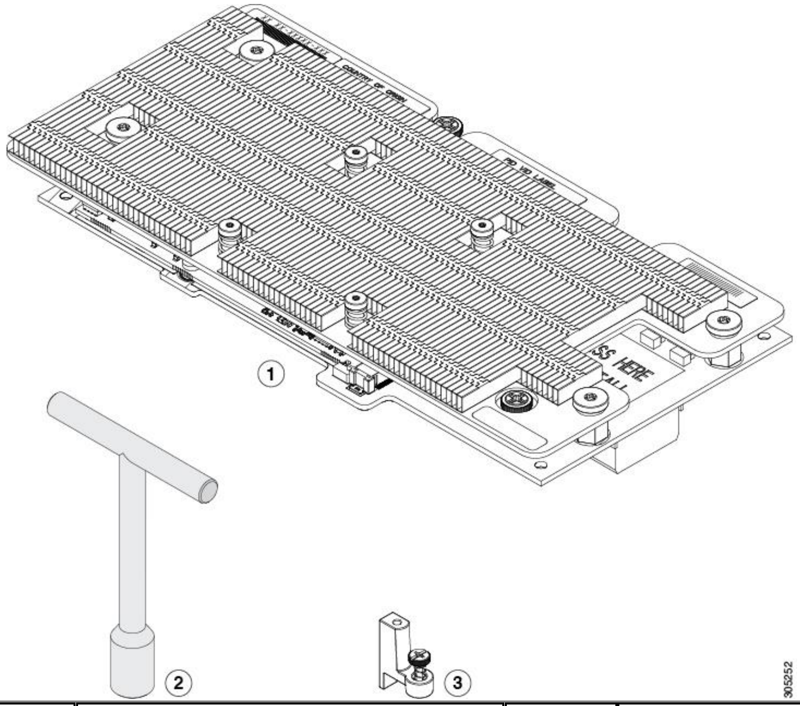

1	NVIDIA M6 GPU (CPU and heat sink)	2	T-shaped wrench
3	Custom standoff		

Before You Begin

Before installing the NVIDIA M6 GPU:

- Remove any adapter card, such as a VIC 1380, VIC 1280, or PT extender card from slot 2. You cannot use any other card in slot 2 when the NVIDIA M6 GPU is installed.
- Upgrade your Cisco UCS system 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.

Procedure

Step 1 Use the T-shaped wrench that comes with the GPU to remove the existing standoff at the back end of the motherboard.

Step 2 Install the custom standoff in the same location at the back end of the motherboard.

Step 3 Position the GPU over the connector on the motherboard and align all captive screws to the standoff posts (callout 1).

Step 4 Tighten the captive screws (callout 2).

Figure 129 Installing the NVIDIA M6 GPU

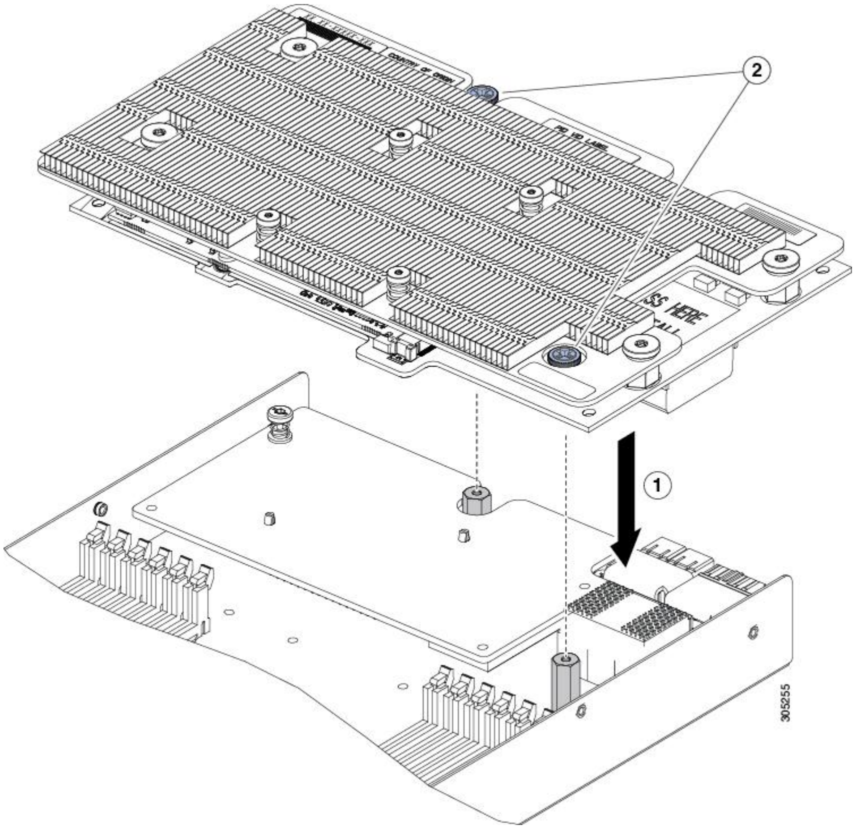

The following figure shows a GPU installed in a Cisco UCS B200 M4 blade server.

Figure 130 Installed NVIDIA M6 GPU

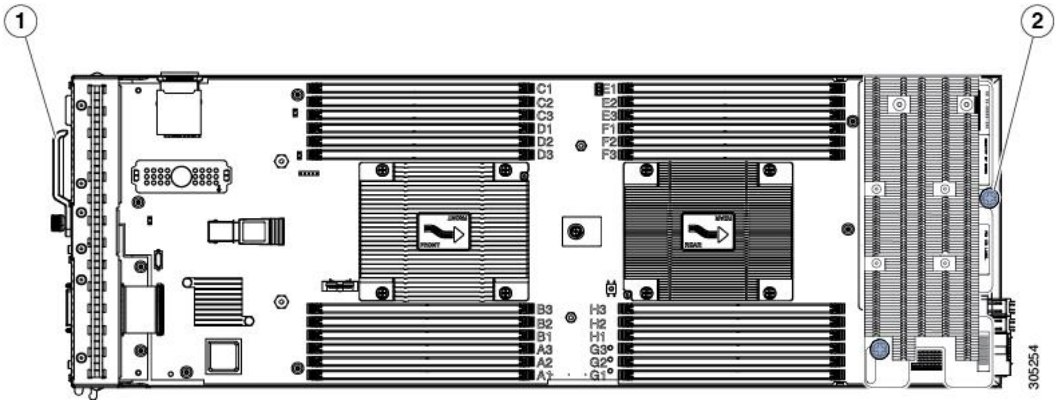

1	Front of server	2	Custom standoff screw

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.

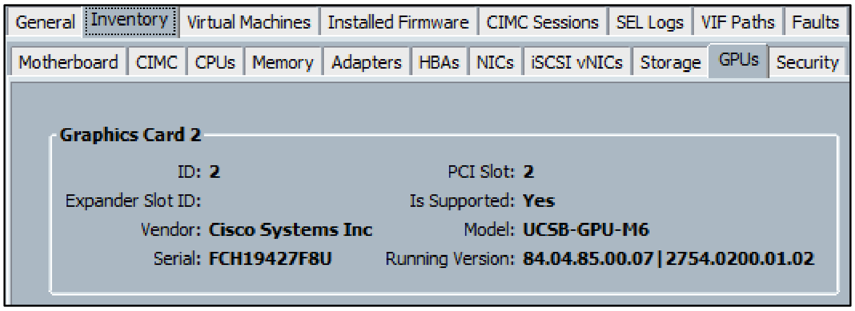

Graphics Card 2

ID: **2** PCI Slot: **2**
Expander Slot ID: Is Supported: **Yes**
Vendor: **Cisco Systems Inc** Model: **UCSB-GPU-M6**
Serial: **FCH19427F8U** Running Version: **84.04.85.00.07 | 2754.0200.01.02**

2. Download the latest drivers and software packages from NVidia's Web Site.

3. Upload the VIB file to the /tmp directory of the ESXi host.

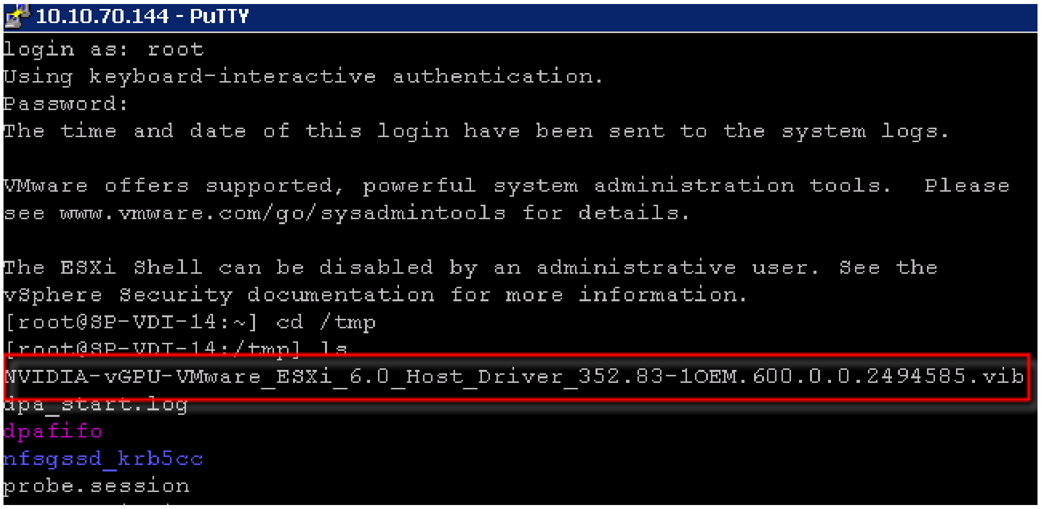

```
login as: root
Using keyboard-interactive authentication.
Password:
The time and date of this login have been sent to the system logs.

VMware offers supported, powerful system administration tools.  Please
see www.vmware.com/go/sysadmintools for details.

The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@SP-VDI-14:~] cd /tmp
[root@SP-VDI-14:/tmp] ls
NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
dpa_start.log
dpafifo
nfsgssd_krb5cc
probe.session
```

4. Install the latest driver: esxcli software vib install –v /tmp/{Latest Driver Package Name}

Host must be in Maintenance Mode to install.

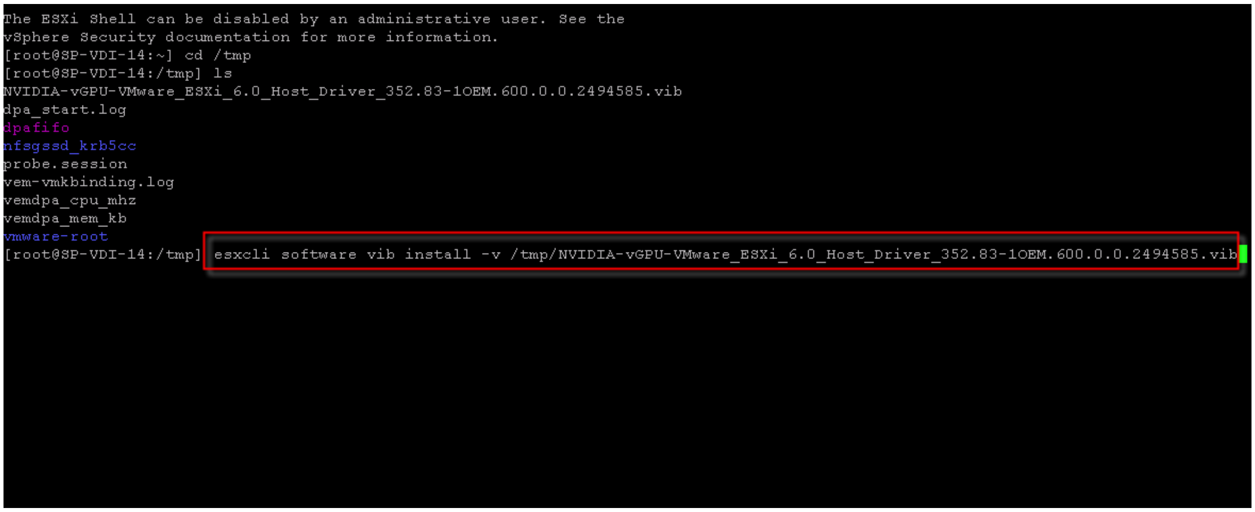

```
The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@SP-VDI-14:~] cd /tmp
[root@SP-VDI-14:/tmp] ls
NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
dpa_start.log
dpafifo
nfsgssd_krb5cc
probe.session
vem-vmkbinding.log
vemdpa_cpu_mhz
vemdpa_mem_kb
vmware-root
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
```

5. A message should validate that the vib installed correctly.

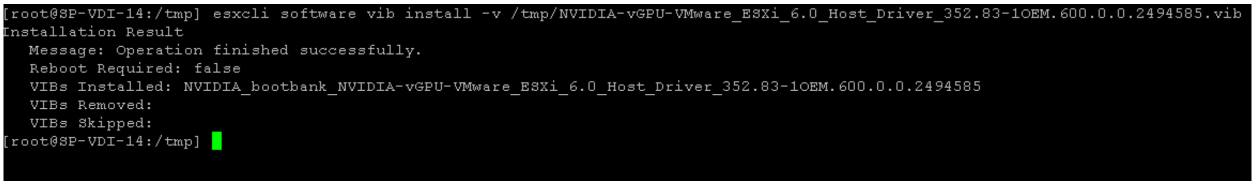

```
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
Installation Result
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: NVIDIA_bootbank_NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585
   VIBs Removed:
   VIBs Skipped:
[root@SP-VDI-14:/tmp]
```

6. Validate the driver was installed by running the command ‘nvidia-smi’ command.

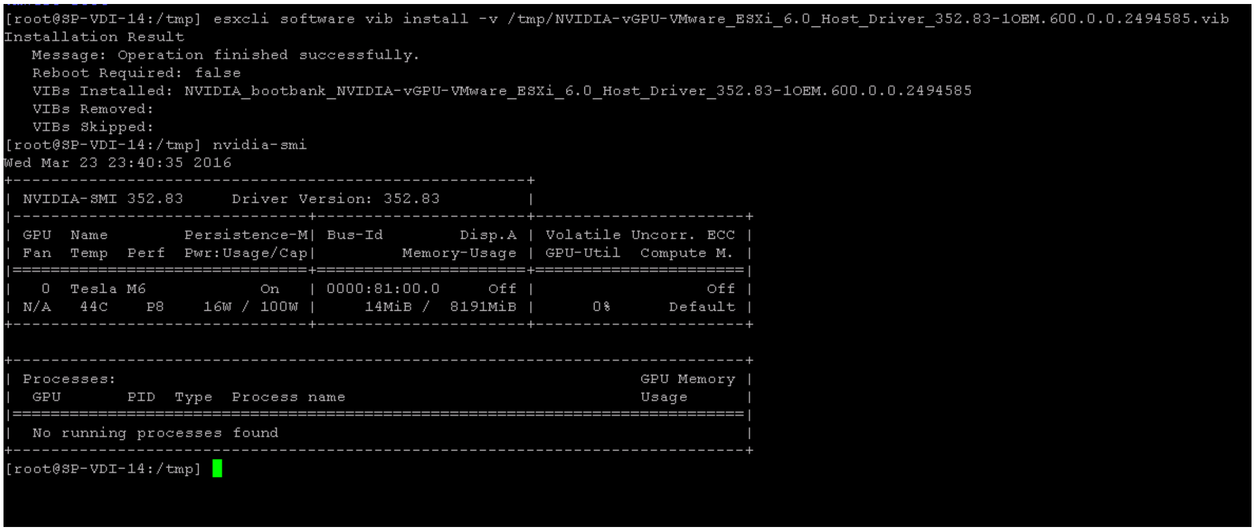

```
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
Installation Result
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: NVIDIA_bootbank_NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585
   VIBs Removed:
   VIBs Skipped:
[root@SP-VDI-14:/tmp] nvidia-smi
Wed Mar 23 23:40:35 2016
+------------------------------------------------------+
| NVIDIA-SMI 352.83     Driver Version: 352.83         |
|-------------------------------+----------------------+----------------------+
| 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 M6            On   | 0000:81:00.0     Off |                  Off |
| N/A   44C    P8    16W / 100W |     14MiB /  8191MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID  Type  Process name                               Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
[root@SP-VDI-14:/tmp]
```

7. By Default the M6 cards come in Compute mode. We will utilize them in Graphics mode in this study. You will need to download the gpumodeswitch utility from NVidia’s web site.   In this exercise, we used the boot ISO which loads a Linux environment with the gpumodeswitch utility already loaded.

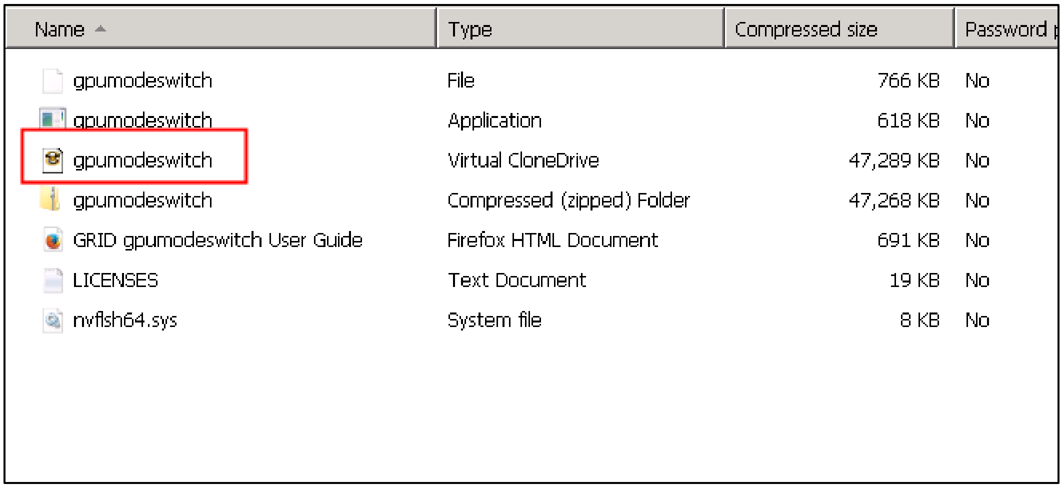

Name	Type	Compressed size	Password p
gpumodeswitch	File	766 KB	No
gpumodeswitch	Application	618 KB	No
gpumodeswitch	Virtual CloneDrive	47,289 KB	No
gpumodeswitch	Compressed (zipped) Folder	47,268 KB	No
GRID gpumodeswitch User Guide	Firefox HTML Document	691 KB	No
LICENSES	Text Document	19 KB	No
nvflsh64.sys	System file	8 KB	No

8. Mount the ISO file through the UCSM KVM and reboot the host.
9. When the Linux shell loads, enter the command: gpumodeswitch –gpumode graphics
10. Type ‘Y’ when prompted to switch all adapters to Graphics. When it completes, reboot back into ESXi.

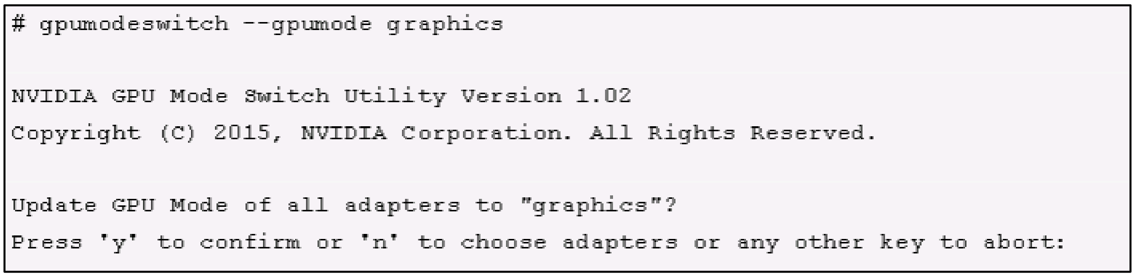

```
# gpumodeswitch --gpumode graphics

NVIDIA GPU Mode Switch Utility Version 1.02
Copyright (C) 2015, NVIDIA Corporation. All Rights Reserved.

Update GPU Mode of all adapters to "graphics"?
Press 'y' to confirm or 'n' to choose adapters or any other key to abort:
```

Configure a VM with a vGPU

To configure a vGPU for a VM, complete the following steps:

1. Select ‘Edit Settings’ in the VSphere Web client for the VM you want to add the vGPU.
2. Select the ‘Virtual Hardware’ tab.
3. In the ‘New device’ section, select ‘Shared PCI Device’ to add the NVIDIA GRID Card.

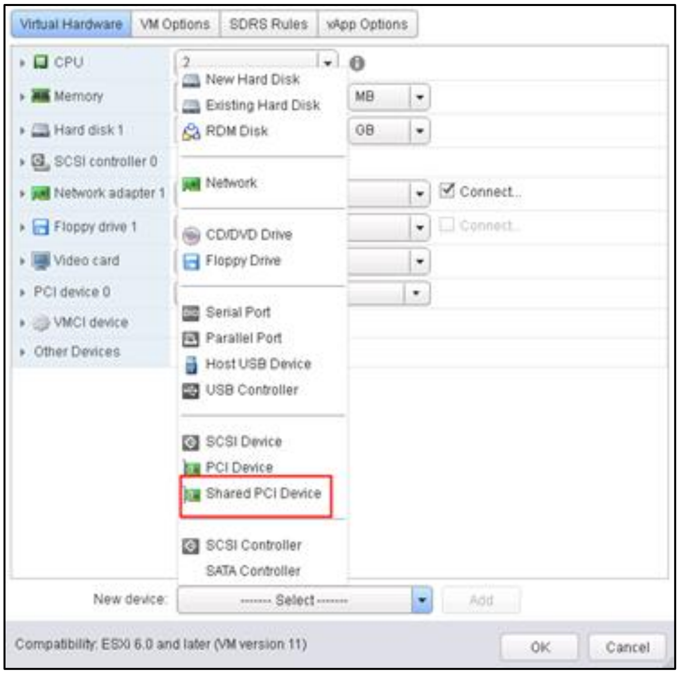

Compatibility: ESXi 6.0 and later (VM version 11)

4. Select the GPU Profile you want to run. In this study, we wanted to achieve a density of 16 vGPU machines on this host so we chose Profile 'grid_m6-0b' which allocates 512Mb per VM for a total of 16 per blade with the M6 Card.

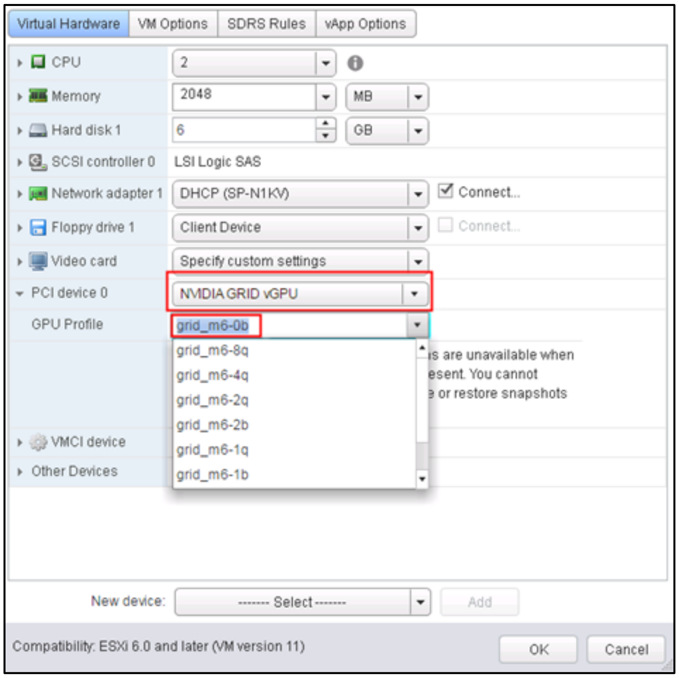

CPU: 2

Memory: 2048 MB

Hard disk 1: 6 GB

SCSI controller 0: LSI Logic SAS

Network adapter 1: DHCP (SP-N1KV)

Floppy drive 1: Client Device

Video card: Specify custom settings

PCI device 0: NVIDIA GRID vGPU

GPU Profile: grid_m6-0b

s are unavailable when
esent. You cannot
e or restore snapshots

Compatibility: ESXi 6.0 and later (VM version 11)

GPU Profiles for the M6 are as follows:

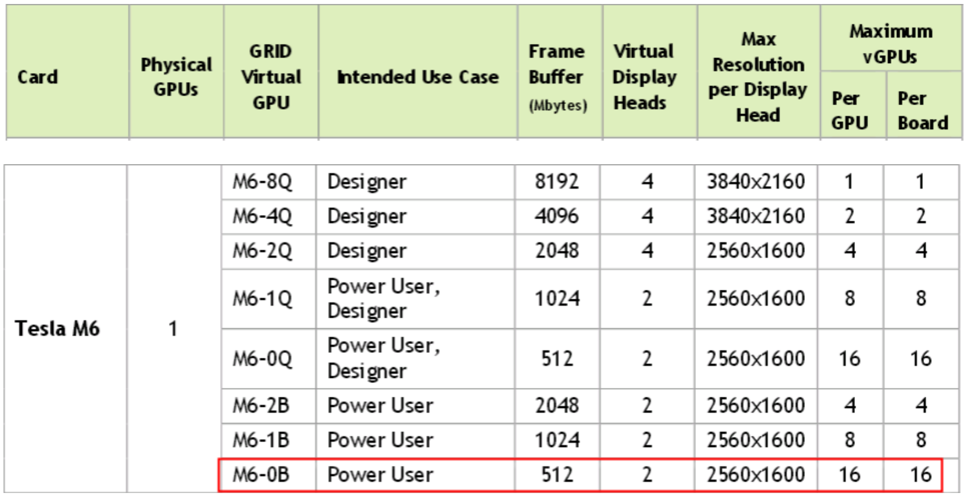

Card	Physical GPUs	GRID Virtual GPU	Intended Use Case	Frame Buffer (Mbytes)	Virtual Display Heads	Max Resolution per Display Head	Maximum vGPUs	
							Per GPU	Per Board
Tesla M6	1	M6-8Q	Designer	8192	4	3840x2160	1	1
		M6-4Q	Designer	4096	4	3840x2160	2	2
		M6-2Q	Designer	2048	4	2560x1600	4	4
		M6-1Q	Power User, Designer	1024	2	2560x1600	8	8
		M6-0Q	Power User, Designer	512	2	2560x1600	16	16
		M6-2B	Power User	2048	2	2560x1600	4	4
		M6-1B	Power User	1024	2	2560x1600	8	8
		M6-0B	Power User	512	2	2560x1600	16	16

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. Since our image is deployed through Citrix PVS, first place the image in Private Mode.

2. Double-click file ‘354.80_grid_win8_win7_international’

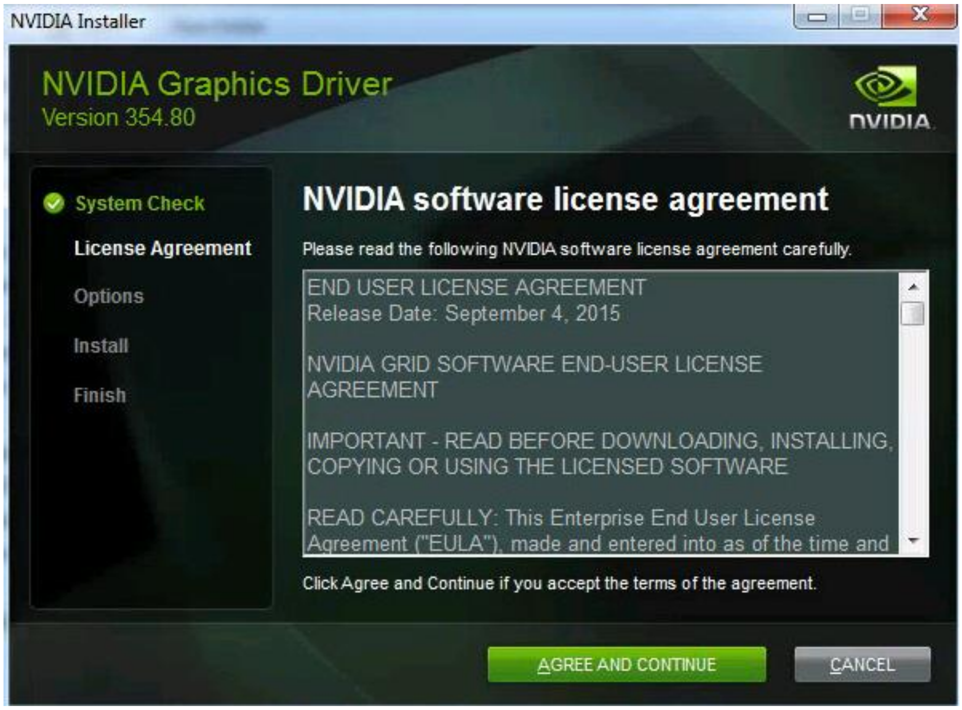

NVIDIA software license agreement

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

Click Agree and Continue if you accept the terms of the agreement.

3. Select Agree and Continue.

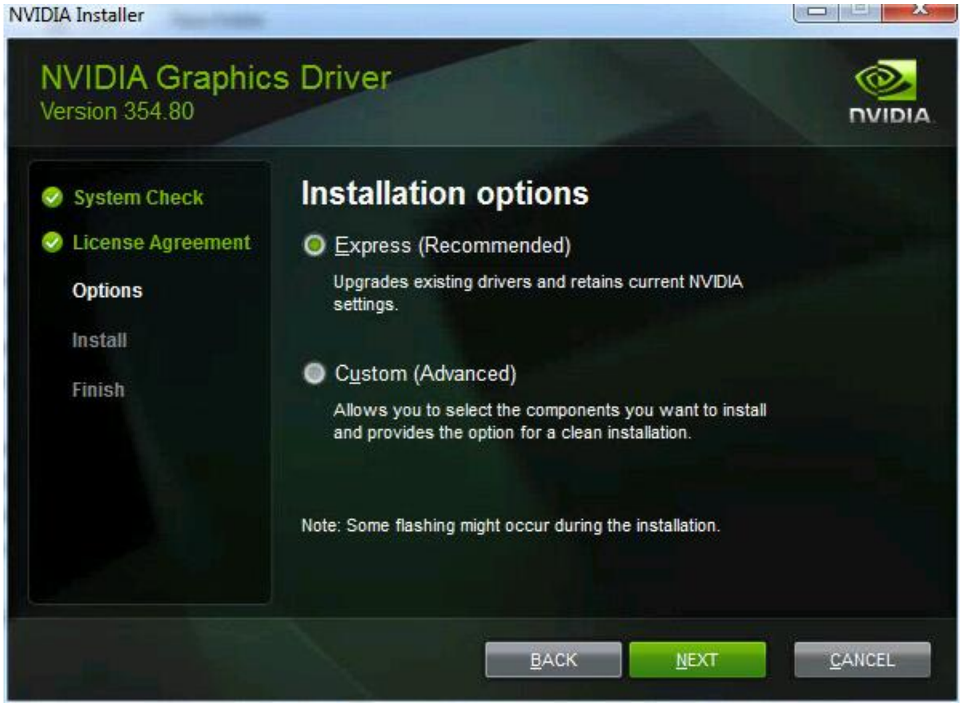

Installation options

● 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.

4. Click Next to use Express Installation.

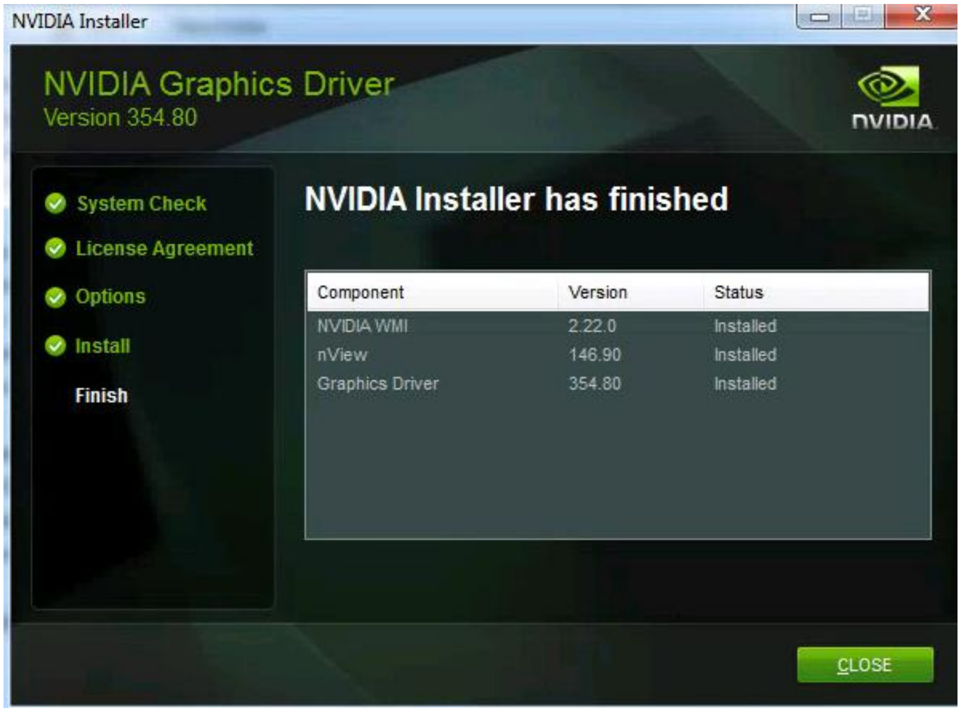

NVIDIA Installer has finished

Component	Version	Status
NVIDIA WMI	2.22.0	Installed
nView	146.90	Installed
Graphics Driver	354.80	Installed

5. The driver and software will be installed and click ‘Finish’ to complete install.

Install and Configure NVIDIA Grid License Server

To use NVIDIA’s vGPU features we must setup a Grid Licensing server. The detailed instructions for setting up a Grid License server can be found in the Grid Quick Start guide. (http://images.nvidia.com/content/grid/pdf/grid-2.0-quick-start-guide.pdf)

The license server requires a fixed IP address. The IP address may be assigned through DHCP or can be statically configured. The server’s Ethernet MAC address is used as a unique identifier when registering the server and generating licenses in NVIDIA’s licensing portal. The server runs on either Windows or Linux.

To create a server interface, complete the following steps:

1. Select Create License Server from under GRID Licensing in the left pane of the NVIDIA Software Licensing Center page to display the Create Server page.

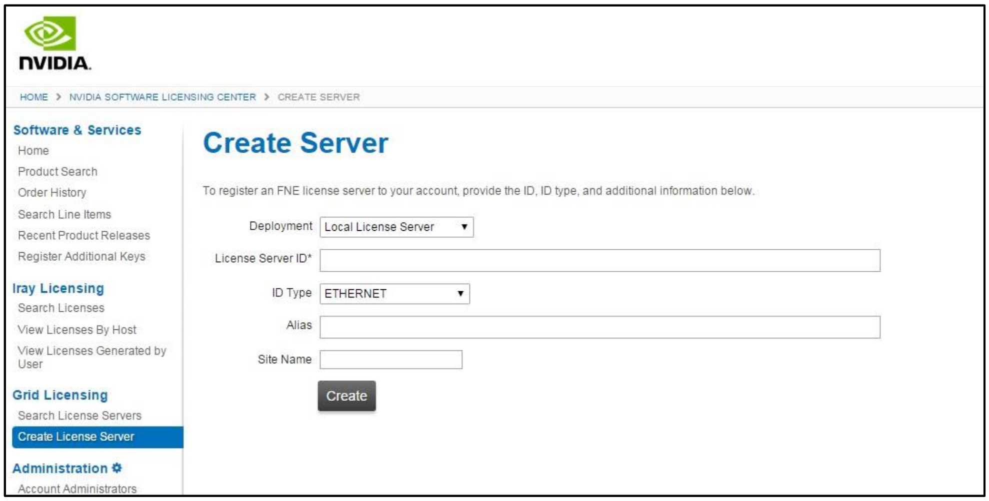

Create Server

To register an FNE license server to your account, provide the ID, ID type, and additional information below.

Deployment: Local License Server

License Server ID*:

ID Type: ETHERNET

Alias:

Site Name:

2. Fill in your server details on the Create Server page.

The License Server ID field is the MAC address of the VM of the License server.

3. Save the .bin file onto your license server for installation.

Java is required to install the NVIDIA GRID License Server. The package comes in a**.zip**file.

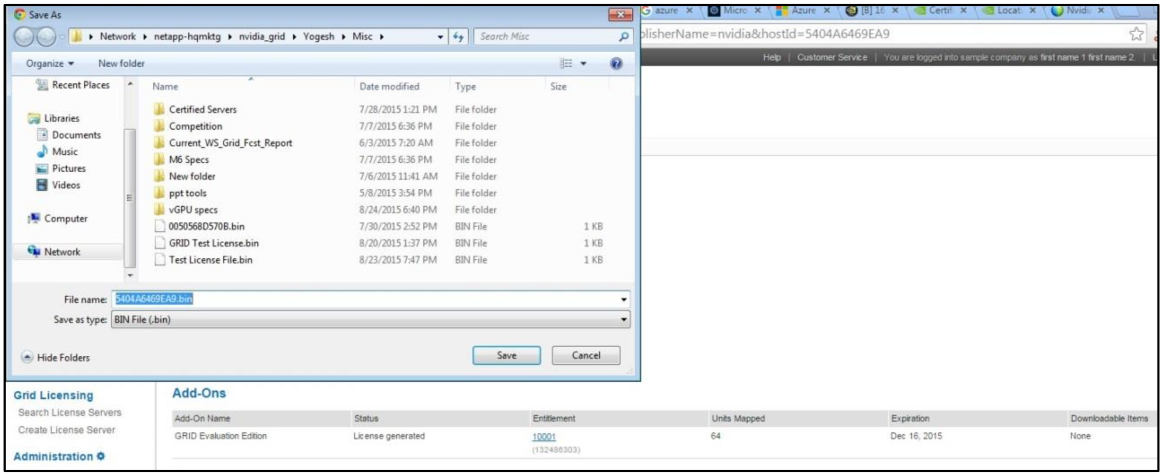

plisherName=nvidia&hostId=5404A6469EA9

Name	Date modified	Type	Size
Certified Servers	7/28/2015 1:21 PM	File folder	
Competition	7/7/2015 6:36 PM	File folder	
Current_WS_Grid_Fcst_Report	6/3/2015 7:20 AM	File folder	
M6 Specs	7/7/2015 6:36 PM	File folder	
New folder	7/6/2015 11:41 AM	File folder	
ppt tools	5/8/2015 3:54 PM	File folder	
vGPU specs	8/24/2015 6:40 PM	File folder	
0050568D570B.bin	7/30/2015 2:52 PM	BIN File	1 KB
GRID Test License.bin	8/20/2015 1:37 PM	BIN File	1 KB
Test License File.bin	8/23/2015 7:47 PM	BIN File	1 KB

File name: 5404A6469EA9.bin

Save as type: BIN File (.bin)

Add-Ons

Add-On Name	Status	Entitlement	Units Mapped	Expiration	Downloadable Items
GRID Evaluation Edition	License generated	10001 (132486303)	64	Dec 16, 2015	None

4. Unzip the license server installer.

5. Run setup.exe and follow the installation wizard.

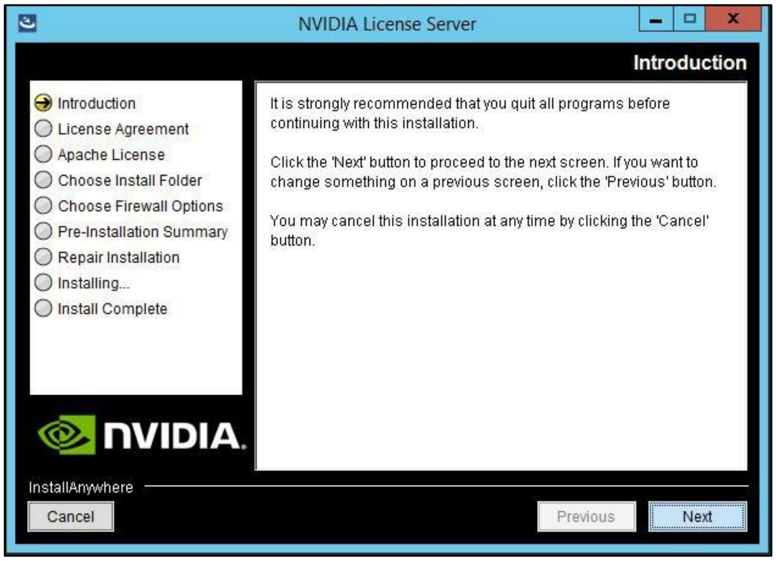

It is strongly recommended that you quit all programs before continuing with this installation.

Click the 'Next' button to proceed to the next screen. If you want to change something on a previous screen, click the 'Previous' button.

You may cancel this installation at any time by clicking the 'Cancel' button.

6. Go to http://<FQDN of the license Server>:8080/licserver to display the License Server Configuration page. You will need the License Server's MAC Address to generate a license .bin file on the portal.

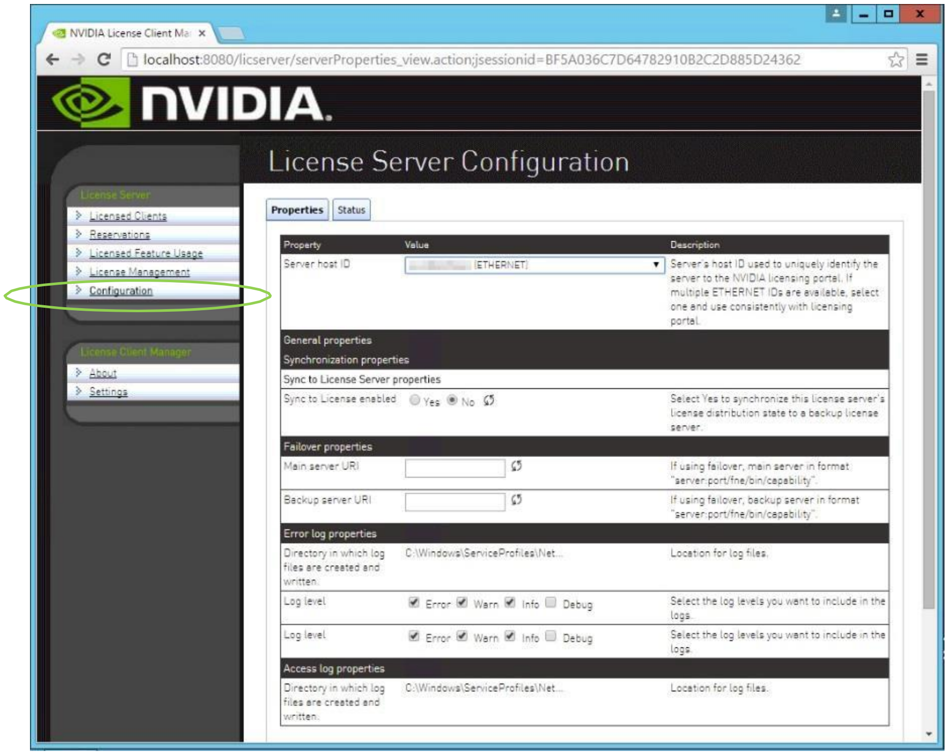

License Server Configuration

Property	Value	Description
Server host ID	[redacted] (ETHERNET) ▾	Server's host ID used to uniquely identify the server to the NVIDIA licensing portal. If multiple ETHERNET IDs are available, select one and use consistently with licensing portal.
General properties		
Synchronization properties		
Sync to License Server properties		
Sync to License enabled	○ Yes ● No	Select Yes to synchronize this license server's license distribution state to a backup license server.
Failover properties		
Main server URI		If using failover, main server in format "server:port/fne/bin/capability".
Backup server URI		If using failover, backup server in format "server:port/fne/bin/capability".
Error log properties		
Directory in which log files are created and written.	C:\Windows\ServiceProfiles\Net...	Location for log files.
Log level	☑ Error ☑ Warn ☑ Info ☐ Debug	Select the log levels you want to include in the logs.
Log level	☑ Error ☑ Warn ☑ Info ☐ Debug	Select the log levels you want to include in the logs.
Access log properties		
Directory in which log files are created and written.	C:\Windows\ServiceProfiles\Net...	Location for log files.

7. Select Configuration from the menu in the left pane.

8. Use the License Server Configuration menu to install the .bin file:

   a. Select Choose File.

   b. Use the file browser to locate the .bin file downloaded from the licensing portal web site.

9. 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.

   a. In Windows – Control Panel, double click the NVidia Control Panel.

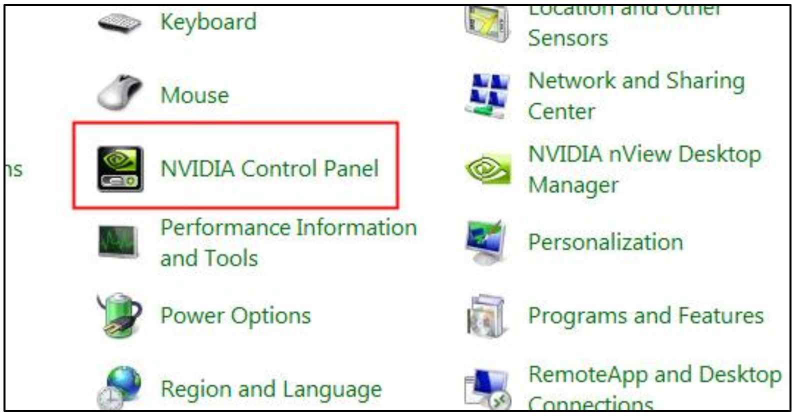
Keyboard
Location and Other Sensors
Mouse
Network and Sharing Center
ns
NVIDIA Control Panel
NVIDIA nView Desktop Manager
Performance Information and Tools
Personalization
Power Options
Programs and Features
Region and Language
RemoteApp and Desktop Connections

b. In the Control Panel, enter the IP or FQDN of the Grid License Server. You should receive a result similar to the below image.

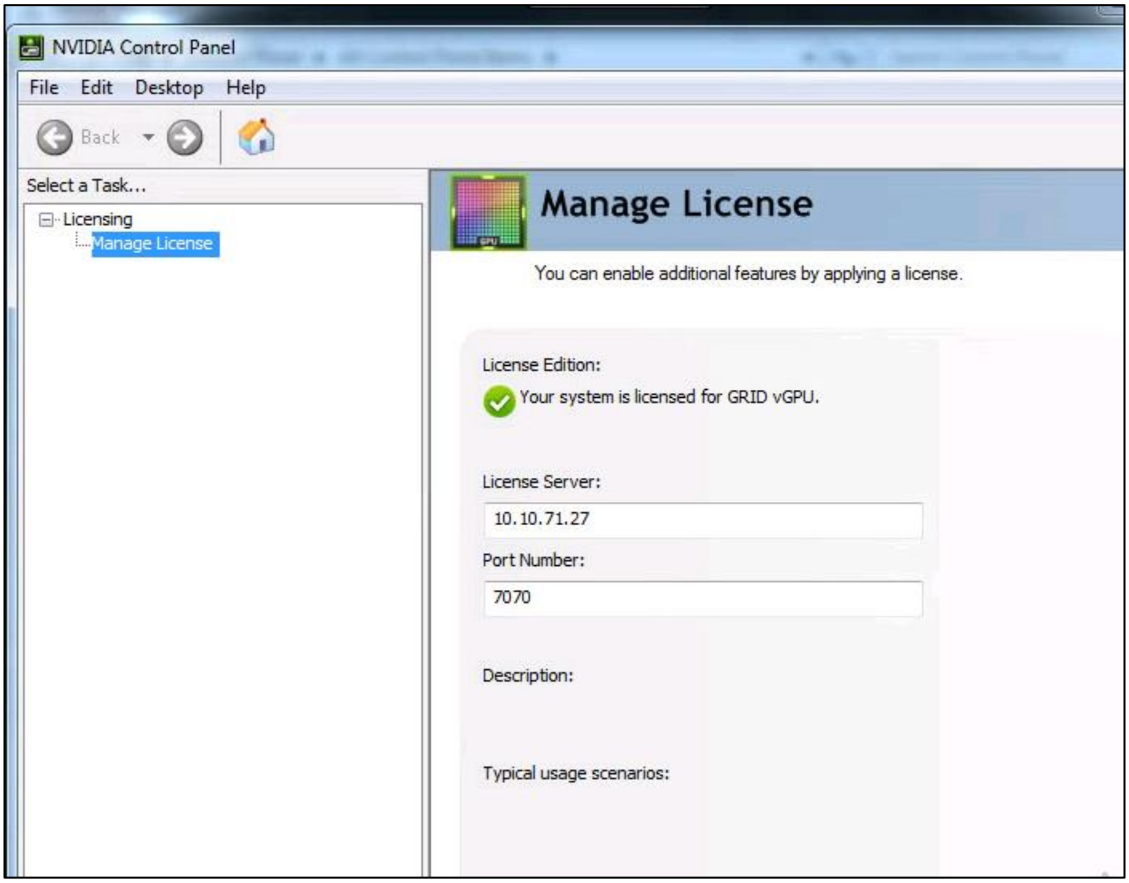

Manage License

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.

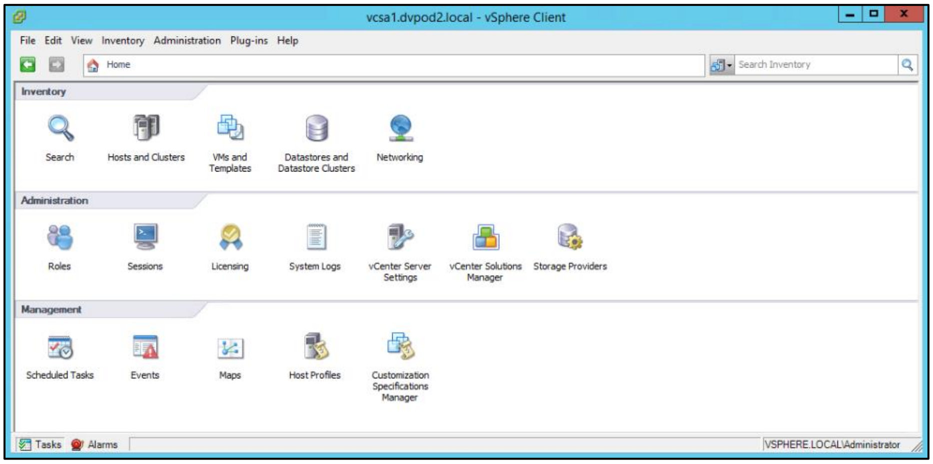

3. From the File menu, select Deploy OVF Template....

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.

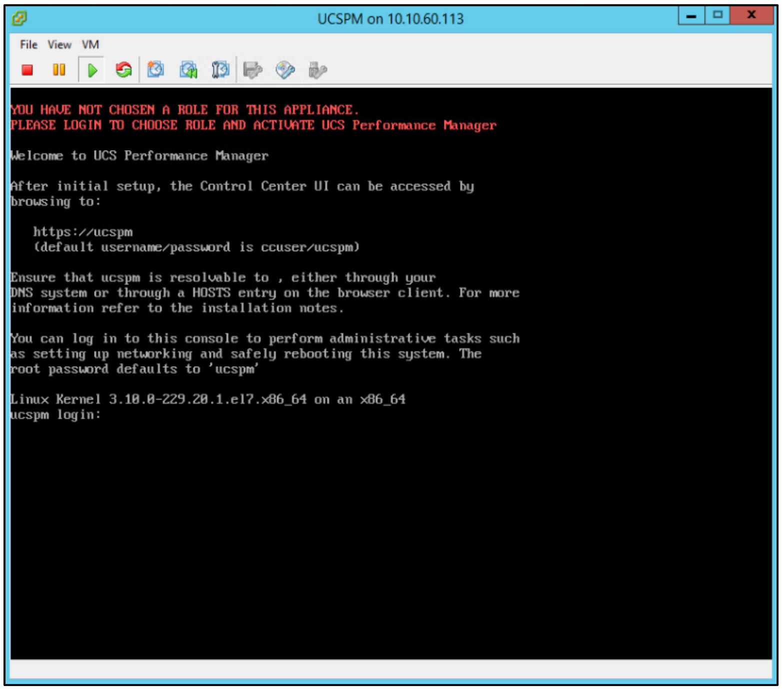

```
YOU HAVE NOT CHOSEN A ROLE FOR THIS APPLIANCE.
PLEASE LOGIN TO CHOOSE ROLE AND ACTIVATE UCS Performance Manager

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 acount is the default account for gaining access to the Control Center browser interface.

6. Select the master role for the host.

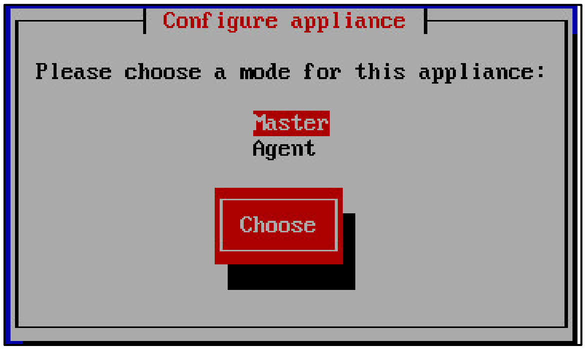

Configure appliance

Please choose a mode for this appliance:

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.

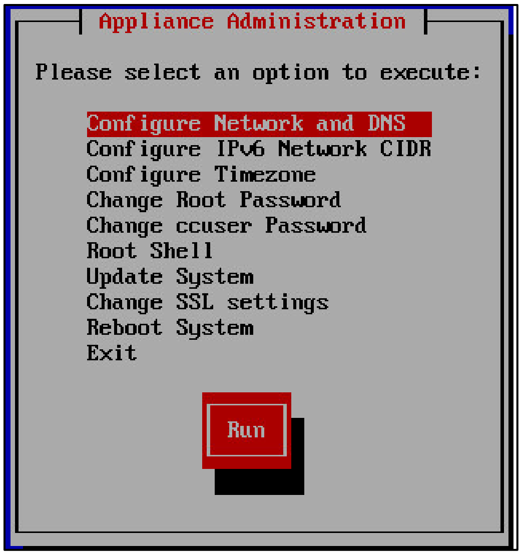

2. Select the NetworkManager TUI menu.

   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.

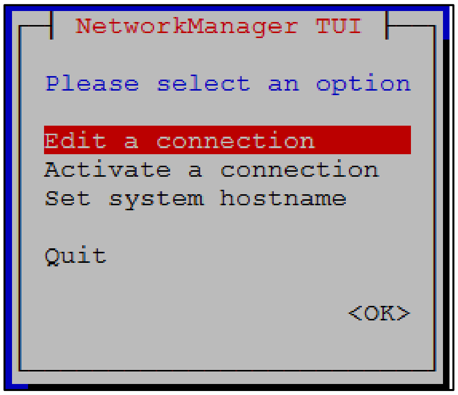

3. On the NetworkManager TUI menu, select Edit a connection, and then press the Return key.

The TUI displays the connections that are available on this host.

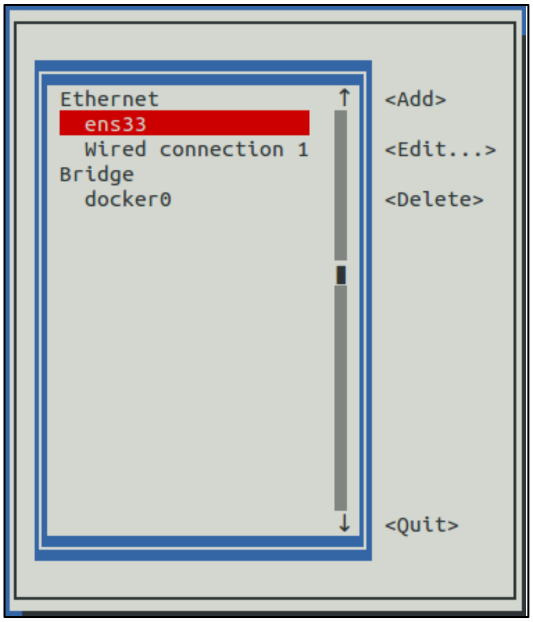

4. Use the down-arrow key to select Wired Connection 1, and then press the Return key.

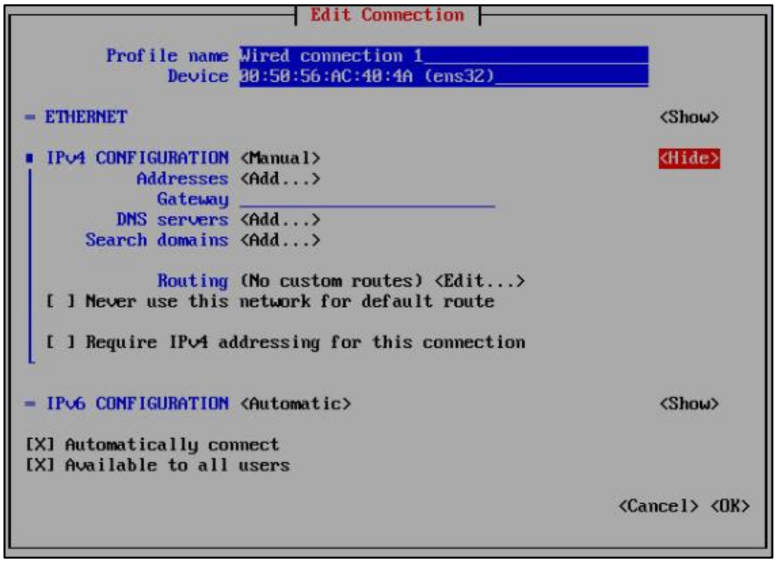

```
                    Edit Connection

      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

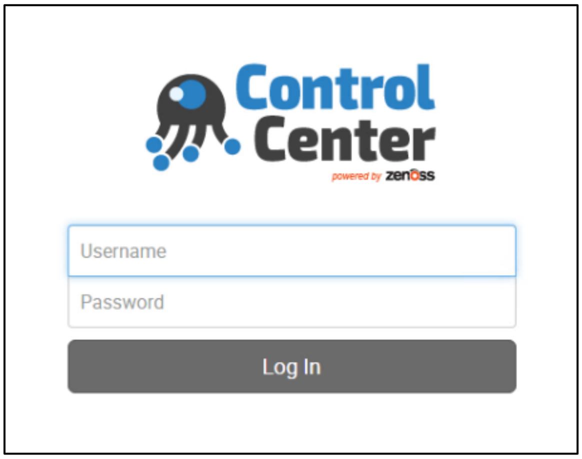

3. At the login page, enter ccuser and its password.

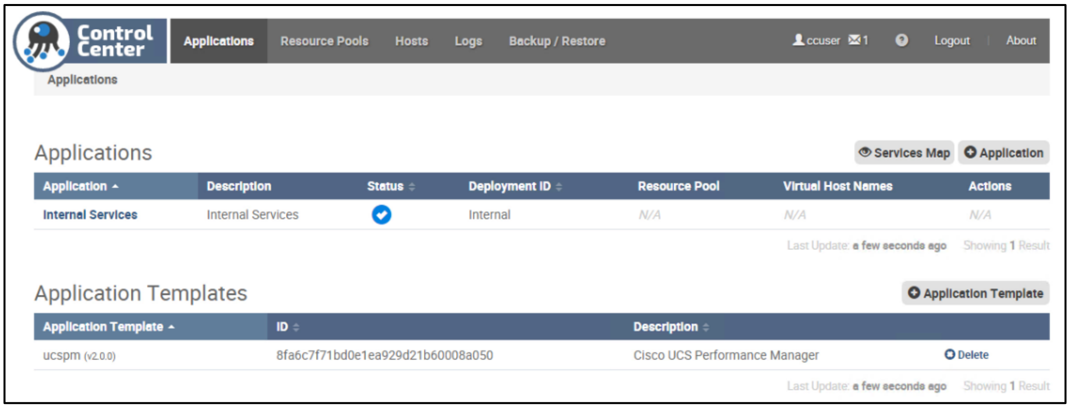

Applications

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.

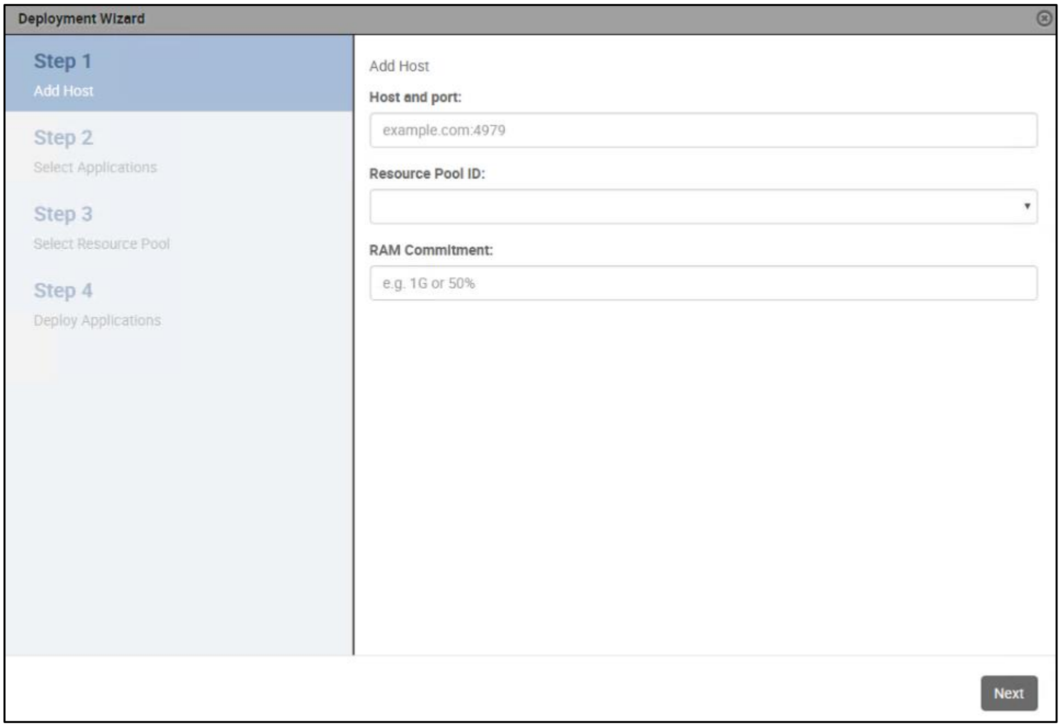

Add Host

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.

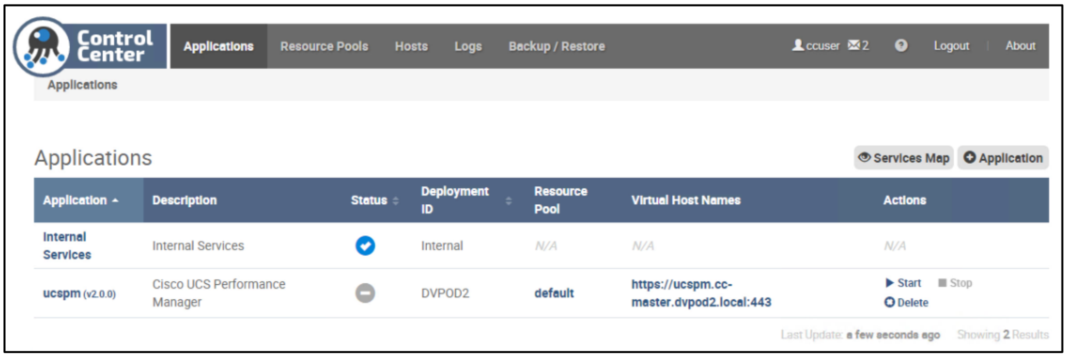

Applications

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.

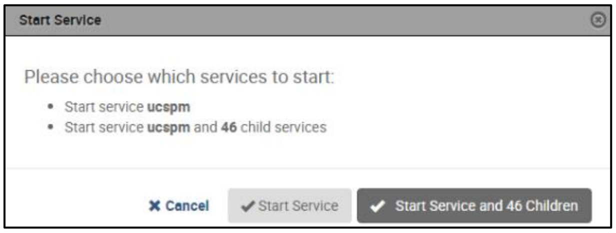

Please choose which services to start:

- 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.

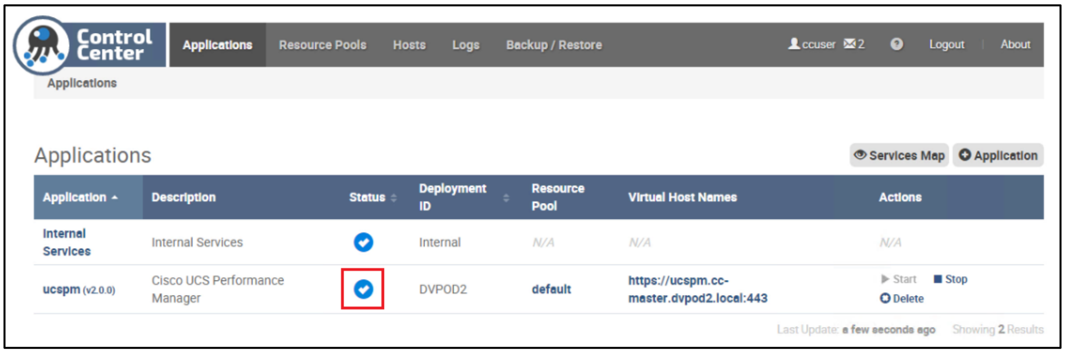

Applications

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

Configuring 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 check box on the left side, and then click the Accept License button on the right side.

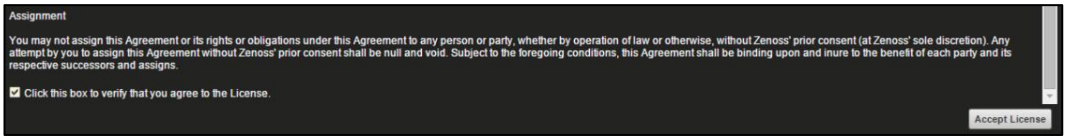

Assignment

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!

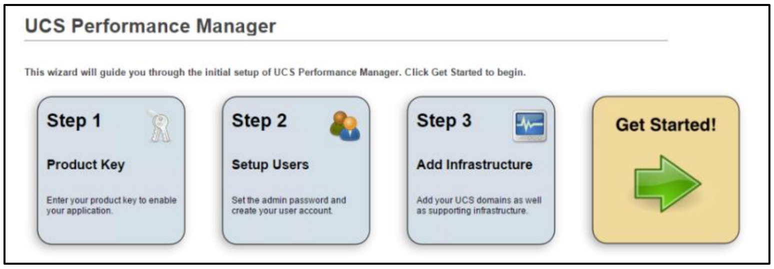

UCS Performance Manager

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.

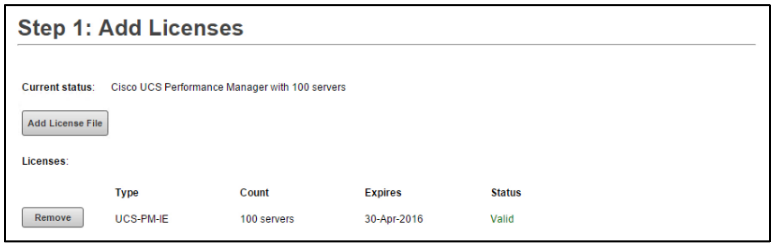

Step 1: Add Licenses

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.

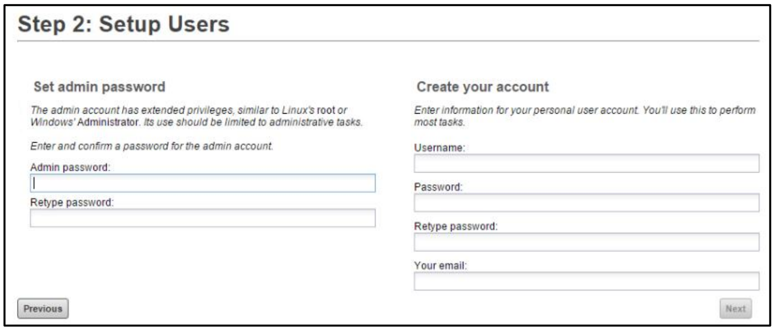

Step 2: Setup Users

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.

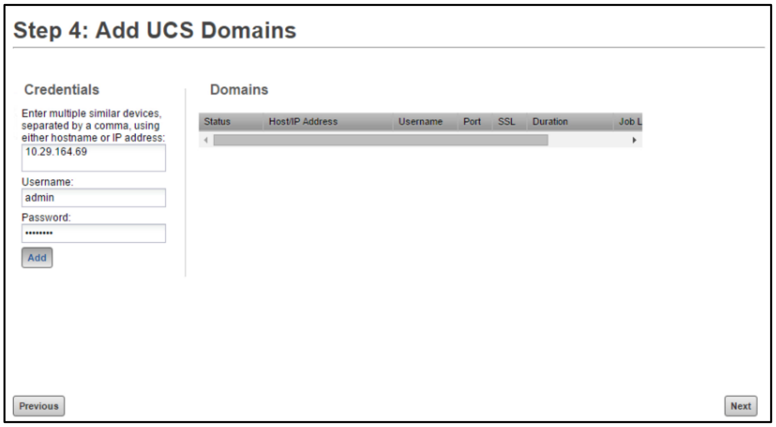

Step 4: Add UCS Domains

Credentials

Enter multiple similar devices, separated by a comma, using either hostname or IP address: 10.29.164.69

Username: admin

Password: ••••••••

Domains

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.

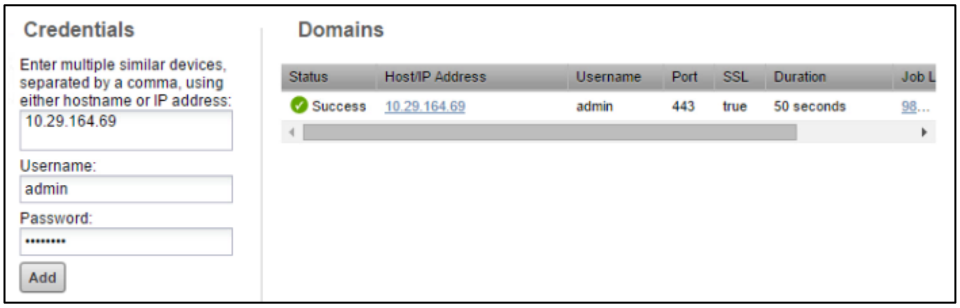

Credentials

Enter multiple similar devices, separated by a comma, using either hostname or IP address: 10.29.164.69

Username: admin

Password: ••••••••

Domains

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:

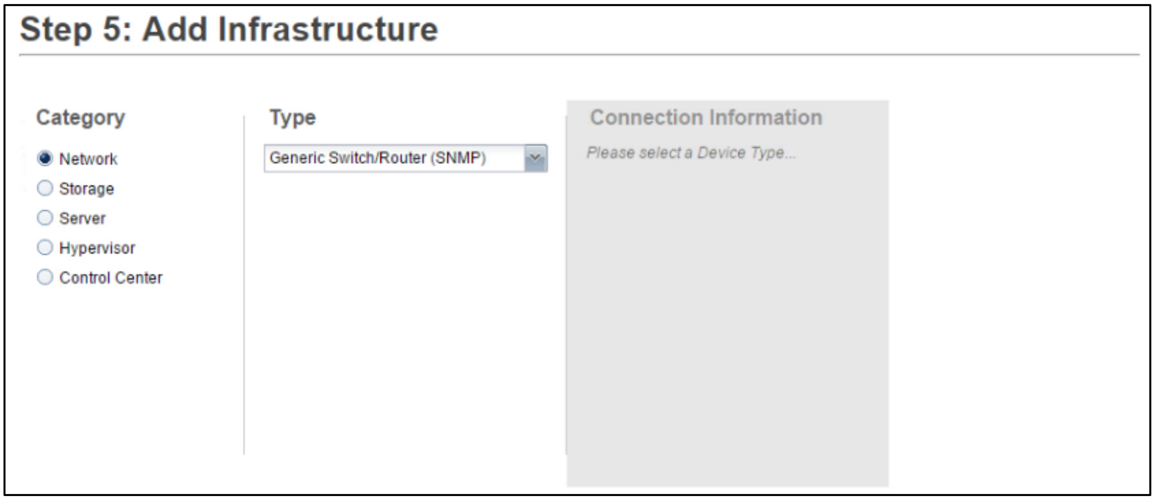

Step 5: Add Infrastructure

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, ESXi hosts using SOAP, and Windows Servers using SNMP or WinRM.

Add Nexus 9000 Series Switches

To add the Infrastructure Devices to Cisco UCS Performance Manager after completing the initial setup configuration, complete the following steps:

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 9372 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.

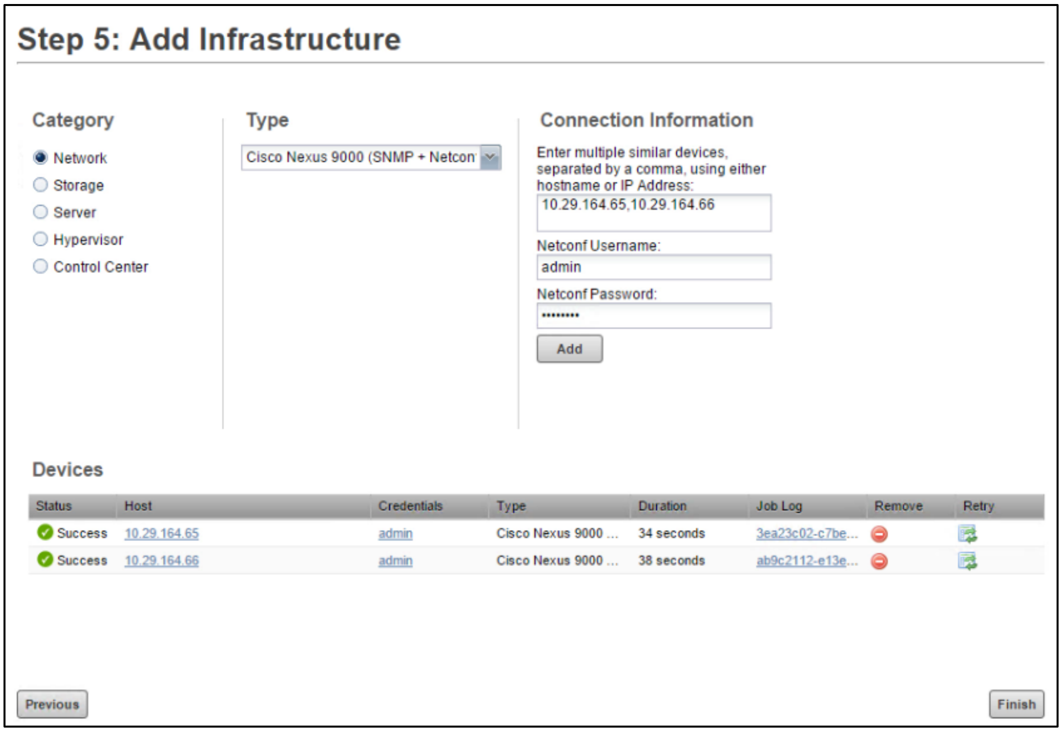

Step 5: Add Infrastructure

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 Sample Test Data

The following samples represent just some of the useful data that can be obtained using Cisco UCS Performance Manager.

The chart shows the network usage from a Fabric Interconnect uplink vPC to a Nexus 9372 switch (Fabric A) during the virtual machine boot storm. All RDS and VDI VMs pertaining to the 5,000 user environment where started in a 15 minute period.

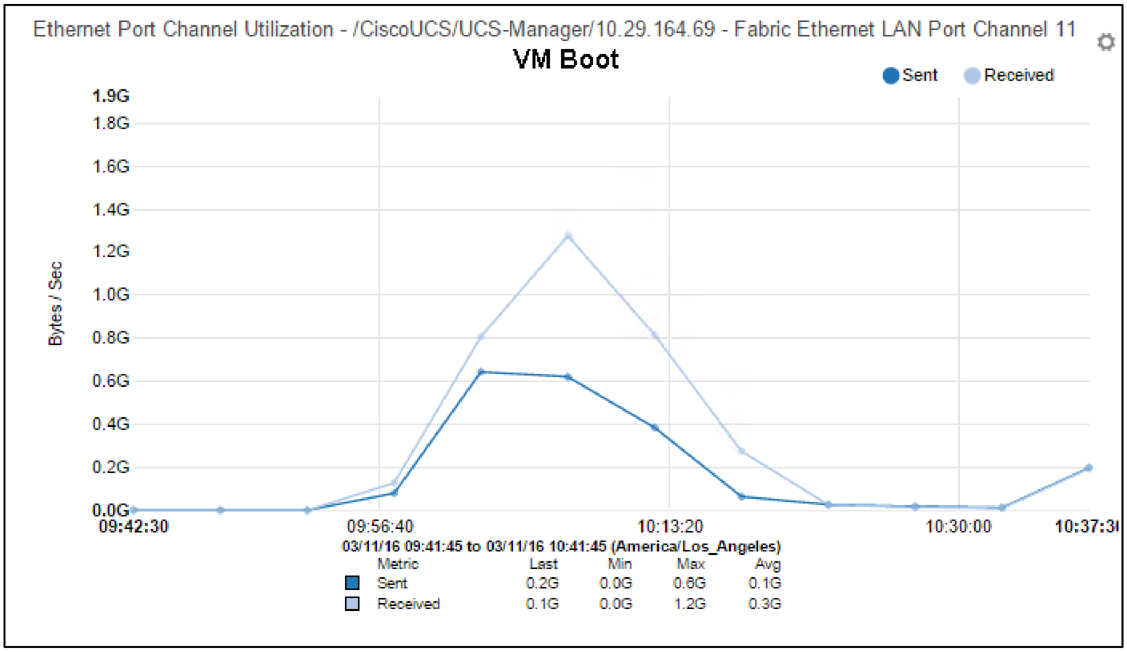

	Metric	Last	Min	Max	Avg
🟦	Sent	0.2G	0.0G	0.6G	0.1G
▢	Received	0.1G	0.0G	1.2G	0.3G

Test Setup and Configurations

In this solution, we tested a single UCS B200 M4 blade to validate against the performance of one blade and twenty-eight B200 M4 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 7.9 with 280 RDS sessions, 190 VDI Non-Persistent sessions, and 190 VDI Persistent sessions.

Figure 131 Cisco UCS B200 M4 Blade Server for Single Server Scalability XenApp 7.9 RDS with PVS 7.9

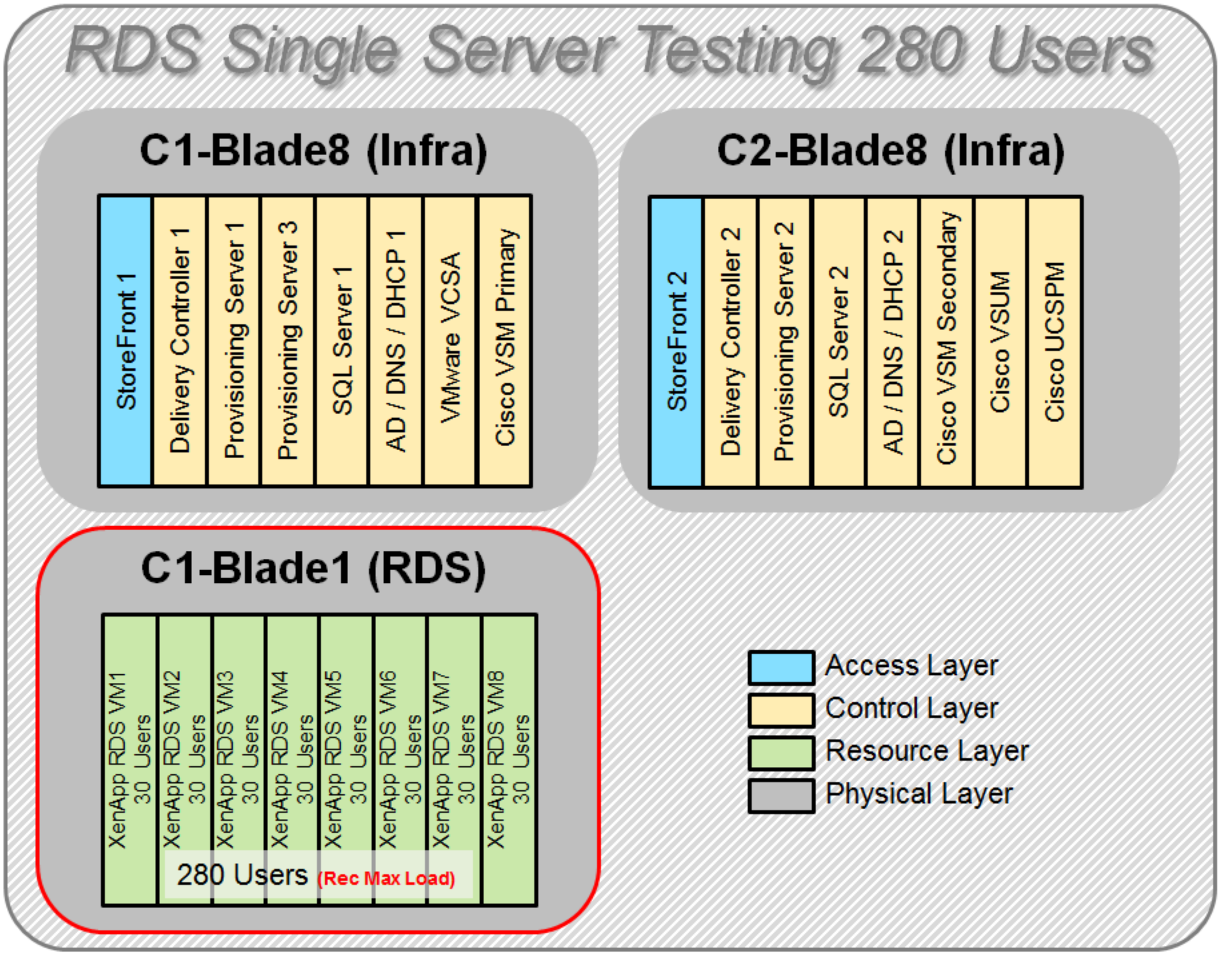

Figure 132 Cisco UCS B200 M4 Blade Server for Single Server Scalability XenDesktop 7.9 VDI (Non-Persistent) with PVS 7.9

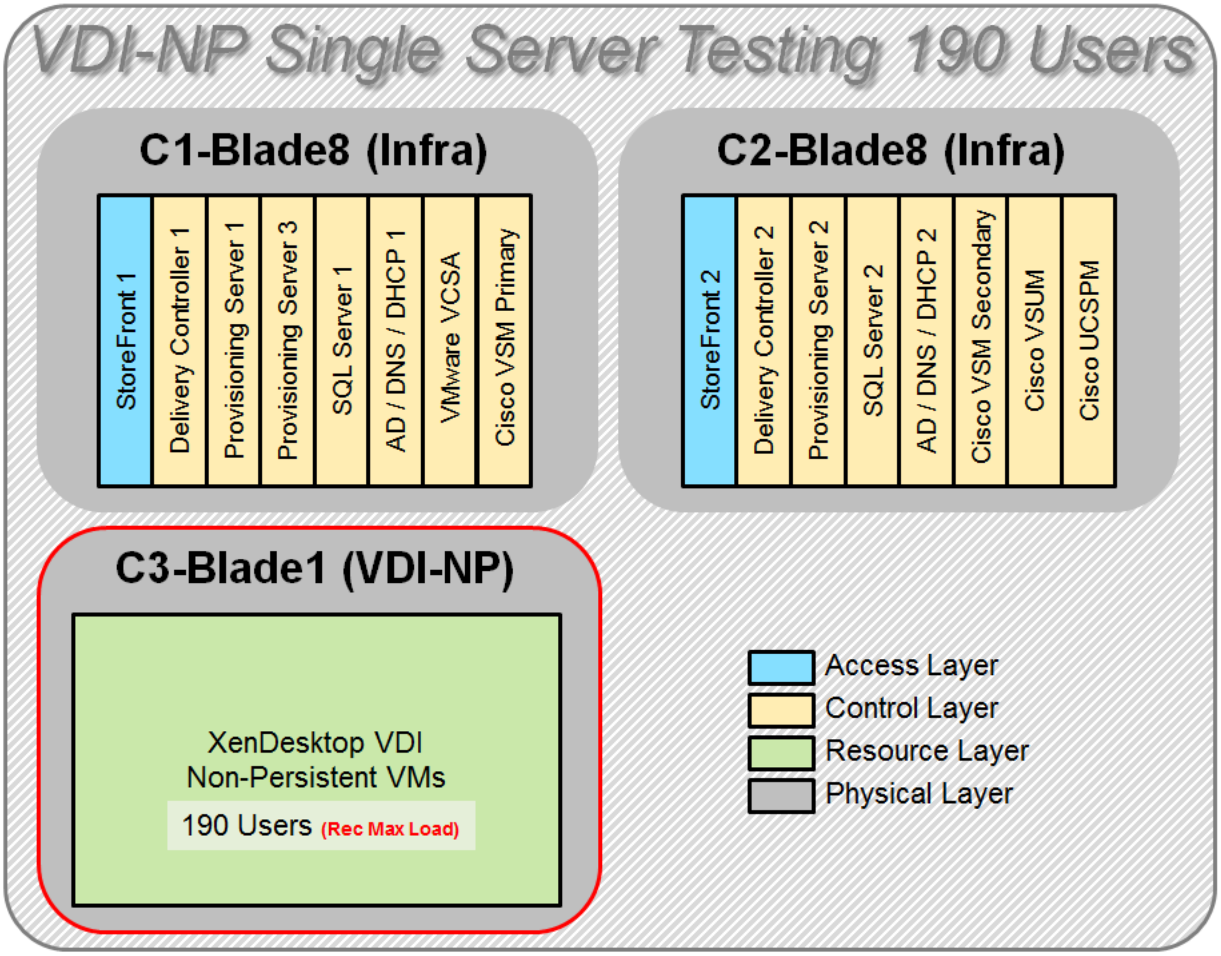

Figure 133 Cisco UCS B200 M4 Blade Server for Single Server Scalability XenDesktop 7.9 VDI (Persistent) with Citrix Machine Creation Services

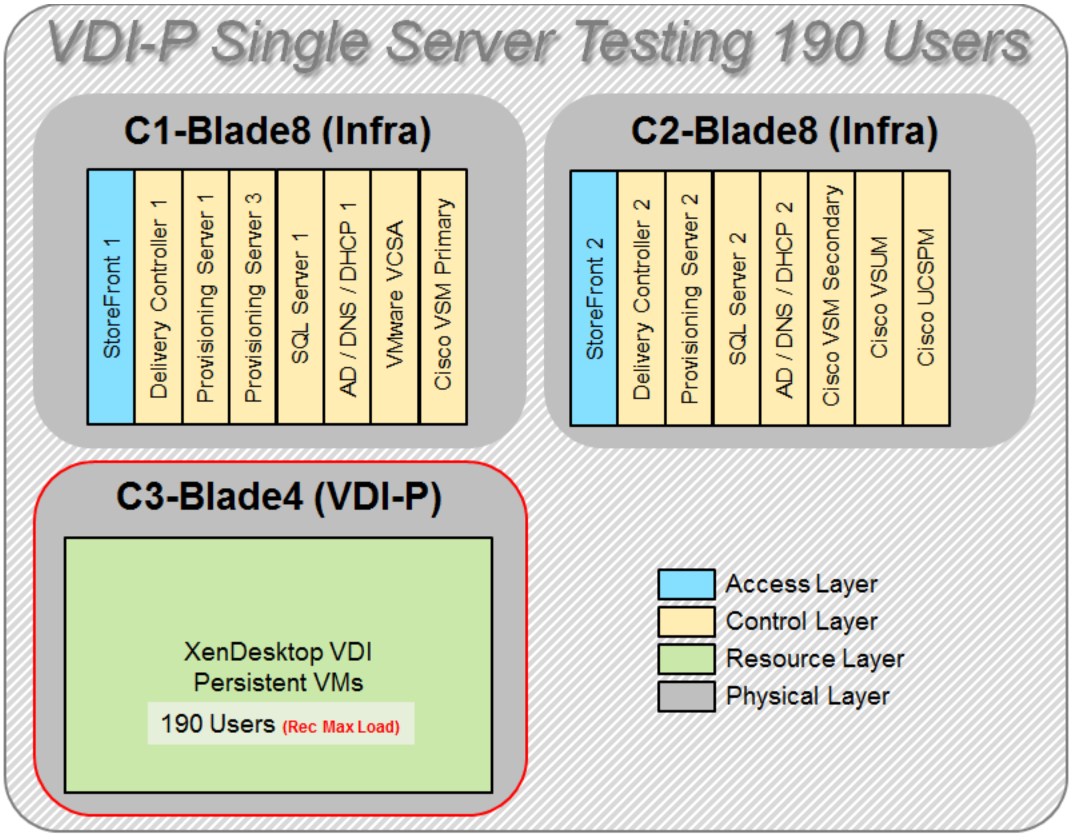

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6248 Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v4 CPUs at 2.0 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz]) for infrastructure host blades
- 1 RDS Cisco UCS B200 M4 Blade Server (2 Intel Xeon processor E5-2680 v4 CPUs at 2.4 GHz, with 256GB of memory per blade server [32 GB x 8 DIMMs at 2400 MHz]) for workload host blades

or

- 1 Persistent/non-persistent VDI Cisco UCS B200 M4 Blade Server (2 Intel Xeon processor E5-2680 v4 CPUs at 2.4 GHz, with 512GB of memory per blade server [32 GB x 16DIMMs at 2400 MHz]) for workload host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148 Fibre Channel Switches

- 1 Pure Storage FlashArray//m50 storage system with 1x 44TB Raw disk shelf

Software components:

- Cisco UCS firmware 3.1(2b)
- VMware ESXi 6.0 Update 1a for host blades
- Citrix XenApp/XenDesktop 7.9 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.9
- Citrix User Profile Manager
- Microsoft SQL Server 2012
- Microsoft Windows 10 64 bit, 2vCPU, 1.7 GB RAM, 24 GB vdisk
- Microsoft Windows Server 2012 R2, 6vCPU, 24GB RAM, 40 GB vdisk
- Microsoft Office 2016
- Login VSI 4.1.5 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Configuration for Cluster Testing

This test case validates three workload clusters using XenApp/XenDesktop 7.9 with 2,600 RDS sessions, 1,200 VDI Non-Persistent sessions, and 1,200 VDI Persistent sessions. Server N+1 fault tolerance is factored into this test scenario for each workload and infrastructure cluster.

Figure 134 RDS Cluster Test Configuration with Ten Blades

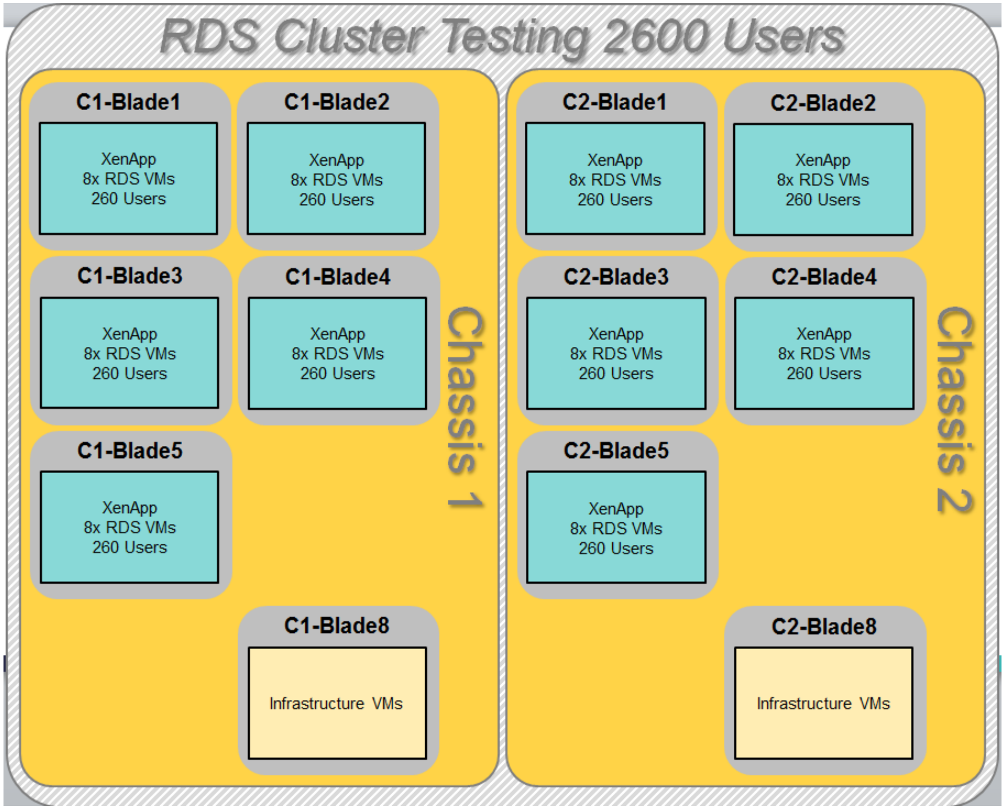

Figure 135 VDI Non-Persistent Cluster Test Configuration with Eight Blades

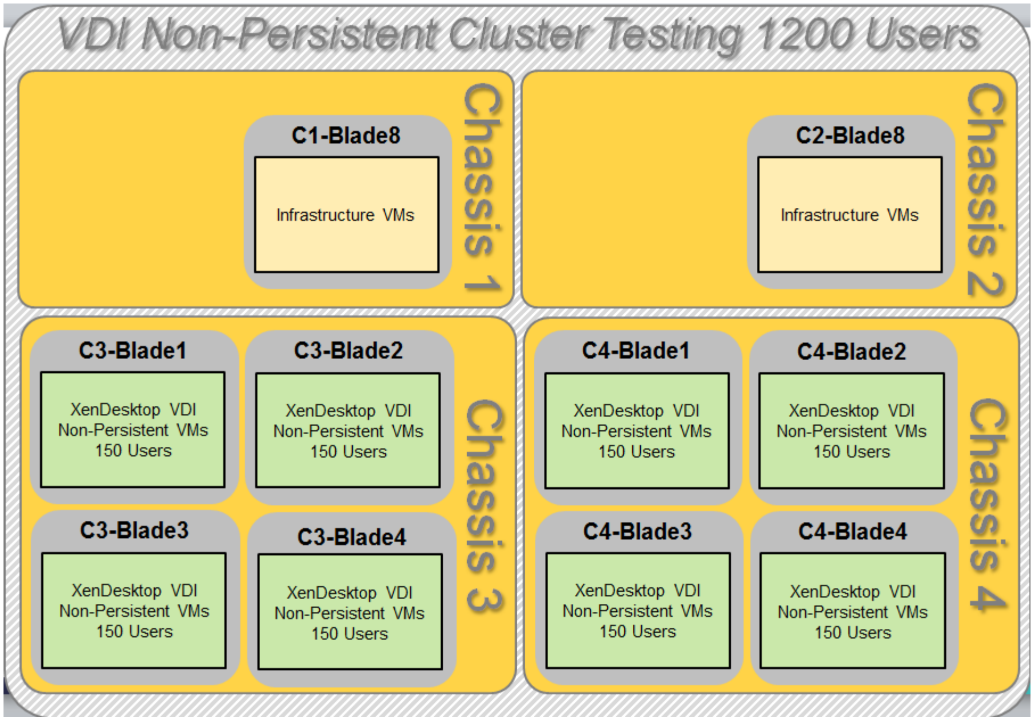

Figure 136 VDI Persistent Cluster Test Configuration with Eight Blades

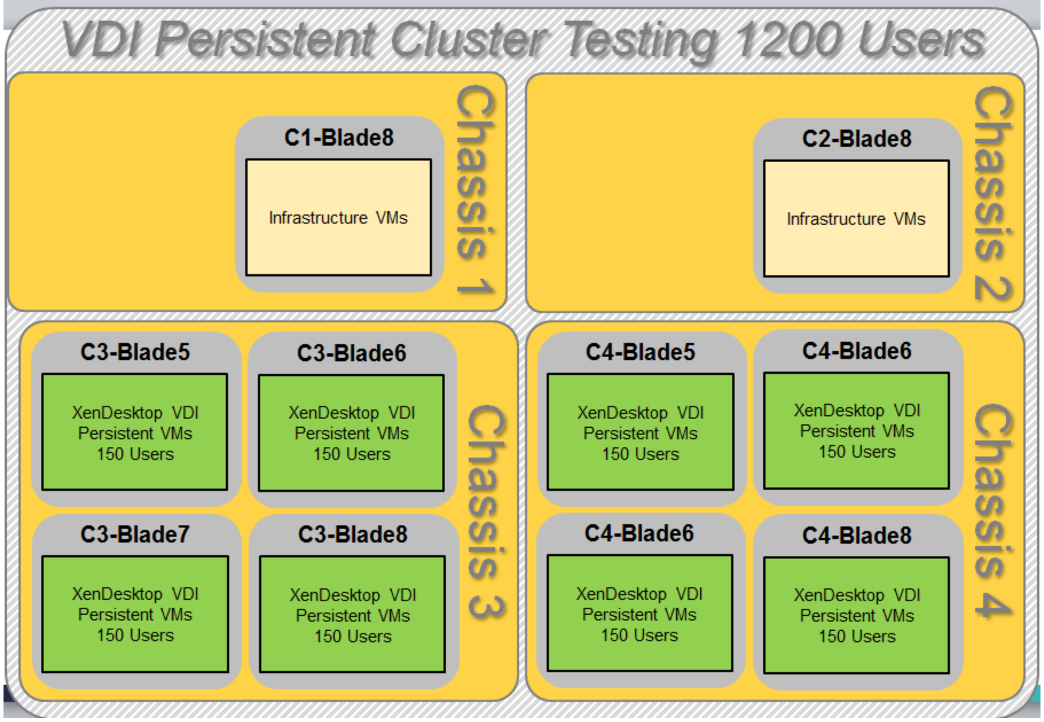
VDI Persistent Cluster Testing 1200 Users
Chassis 1
C1-Blade8
Infrastructure VMs
Chassis 2
C2-Blade8
Infrastructure VMs
Chassis 3
C3-Blade5
XenDesktop VDI Persistent VMs 150 Users
C3-Blade6
XenDesktop VDI Persistent VMs 150 Users
C3-Blade7
XenDesktop VDI Persistent VMs 150 Users
C3-Blade8
XenDesktop VDI Persistent VMs 150 Users
Chassis 4
C4-Blade5
XenDesktop VDI Persistent VMs 150 Users
C4-Blade6
XenDesktop VDI Persistent VMs 150 Users
C4-Blade6
XenDesktop VDI Persistent VMs 150 Users
C4-Blade8
XenDesktop VDI Persistent VMs 150 Users

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6248 Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v4 CPUs at 2.0 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz]) for infrastructure host blades
- 26 Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2680 v4 CPUs at 2.4 GHz, with 512GB of memory per blade server [32 GB x 16DIMMs at 2400 MHz]) for workload host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 1 Pure Storage FlashArray//m50 storage system with 1x 44TB Raw disk shelf

Software components:

- Cisco UCS firmware 3.1(1e)
- VMware ESXi 6.0 Update 1a for host blades

- Citrix XenApp/XenDesktop 7.9 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.9
- Citrix User Profile Manager
- Microsoft SQL Server 2012
- Microsoft Windows 10 64 bit, 2vCPU, 1.7 GB RAM, 24 GB vdisk
- Microsoft Windows Server 2012 R2, 6vCPU, 24GB RAM, 40 GB vdisk
- Microsoft Office 2016
- Login VSI 4.1.5 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Configuration for Full Scale Testing

This test case validates twenty-eight blades mixed workloads using XenApp/XenDesktop 7.9 with 2,600 RDS sessions, 1,200 VDI Non-Persistent sessions, and 1,200 VDI Persistent sessions for a total sum of 5,000 users. Server N+1 fault tolerance is factored into this solution for each workload and infrastructure cluster.

Figure 137 Full Scale Test Configuration with Twenty-Eight Blades

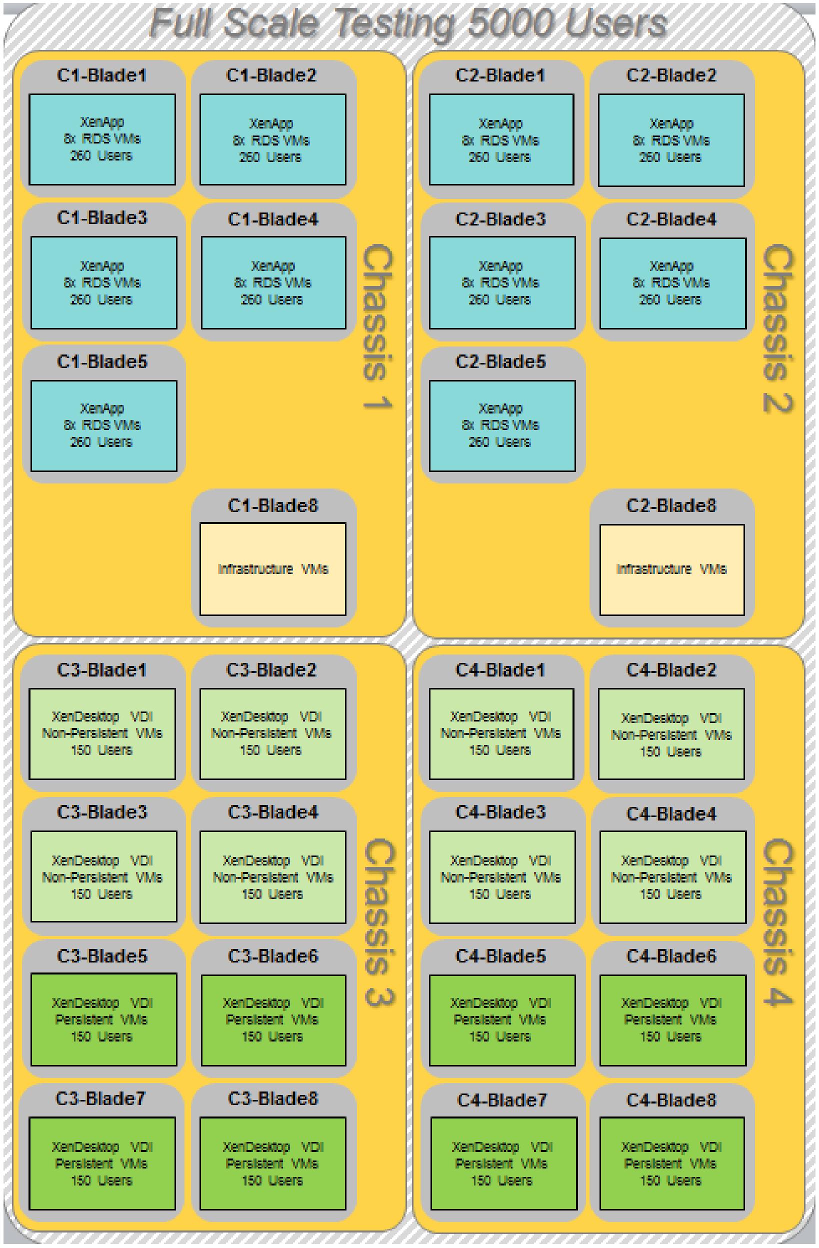
Full Scale Testing 5000 Users
Chassis 1
C1-Blade1
XenApp 8x RDS VMs 260 Users
C1-Blade2
XenApp 8x RDS VMs 260 Users
C1-Blade3
XenApp 8x RDS VMs 260 Users
C1-Blade4
XenApp 8x RDS VMs 260 Users
C1-Blade5
XenApp 8x RDS VMs 260 Users
C1-Blade8
Infrastructure VMs
Chassis 2
C2-Blade1
XenApp 8x RDS VMs 260 Users
C2-Blade2
XenApp 8x RDS VMs 260 Users
C2-Blade3
XenApp 8x RDS VMs 260 Users
C2-Blade4
XenApp 8x RDS VMs 260 Users
C2-Blade5
XenApp 8x RDS VMs 260 Users
C2-Blade8
Infrastructure VMs
Chassis 3
C3-Blade1
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade2
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade3
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade4
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade5
XenDesktop VDI Persistent VMs 150 Users
C3-Blade6
XenDesktop VDI Persistent VMs 150 Users
C3-Blade7
XenDesktop VDI Persistent VMs 150 Users
C3-Blade8
XenDesktop VDI Persistent VMs 150 Users
Chassis 4
C4-Blade1
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade2
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade3
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade4
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade5
XenDesktop VDI Persistent VMs 150 Users
C4-Blade6
XenDesktop VDI Persistent VMs 150 Users
C4-Blade7
XenDesktop VDI Persistent VMs 150 Users
C4-Blade8
XenDesktop VDI Persistent VMs 150 Users

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6248 Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v4 CPUs at 2.0 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz]) for infrastructure host blades

- 26 Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2680 v4 CPUs at 2.4 GHz, with 512GB of memory per blade server [32 GB x 16DIMMs at 2400 MHz]) for workload host blades
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 1 Pure Storage FlashArray//m50 storage system with 1x 44TB Raw disk shelf

Software components:

- Cisco UCS firmware 3.1(2b)
- VMware ESXi 6.0 Update 2 for host blades
- Citrix XenApp/XenDesktop 7.9 VDI Hosted Virtual Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.9
- Citrix User Profile Manager
- Microsoft SQL Server 2012
- Microsoft Windows 10 64-bit, 2vCPU, 1.7 GB RAM, 24 GB vdisk
- Microsoft Windows Server 2012 R2, 6vCPU, 24GB RAM, 40 GB vdisk
- Microsoft Office 2016
- Login VSI 4.1.5 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 RDSH Servers Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from http://www.loginvsi.com.

Testing Procedure

The following protocol was used for each test cycle in this study to insure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

All machines were shut down utilizing the Citrix XenDesktop Studio console.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 900 full scale test users, to become active within two minutes after the last session is launched.

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 Logging on the following systems:

   - Infrastructure and VDI Host Blades used in test run

    - All Infrastructure VMs used in test run (AD, SQL, View Connection brokers, image mgmt., etc.)

2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.

3. Time 0:05: Boot RDS Machines using Citrix XenDesktop Studio.

4. Time 0:06 First machines boot.

5. Time 0:35 Single Server or Scale target number of RDS Servers registered on XD.

No more than 60 Minutes of rest time is allowed after the last desktop is registered and available on Citrix XenDesktop Studio dashboard. Typically a 20-30 minute rest period for Windows 10 desktops and 10 minutes for RDS VMs is sufficient.

6. Time 1:35 Start Login VSI 4.1.5 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).

7. Time 2:23 Single Server or Scale target number of desktop VMs desktops launched (48 minute benchmark launch rate).

8. Time 2:25 All launched sessions must become active.

All sessions launched must become active for a valid test run within this window.

9. Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above).

10. Time 2:55 All active sessions logged off.

All sessions launched and active must be logged off for a valid test run. The Citrix XenDesktop Studio must show that all desktops have been returned to the registered/available state as evidence of this condition being met.

11. Time 2:57 All logging terminated; Test complete.

12. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows 10 machines.

13. Time 3:30 Reboot all hypervisors.

14. Time 3:45 Ready for 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 4.1.5 Knowledge 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 must 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.

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 proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing FlashStack with Cisco UCS B200 M4 and Citrix XenApp/XenDesktop 7.9 on VMware ESXi 6.0 Update 2 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) models with Citrix Provisioning Services 7.9 using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2012 R2 sessions on Cisco UCS B200 M4 Blade Servers using a Pure Storage FlashArray//m 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.

Three 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 to 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 138 Sample of a VSI max response time graph, representing a normal test

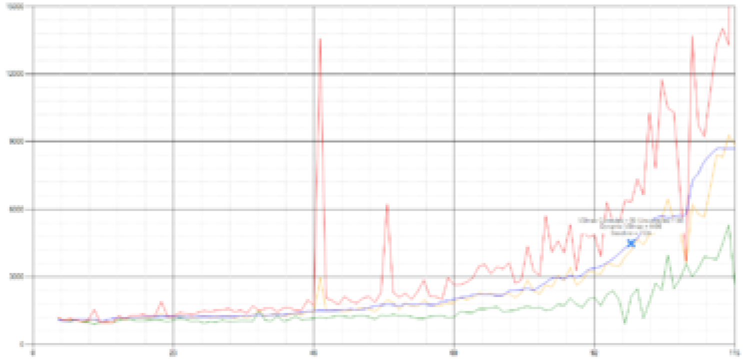

Figure 139 Sample of a VSI test response time graph where there was a clear performance issue

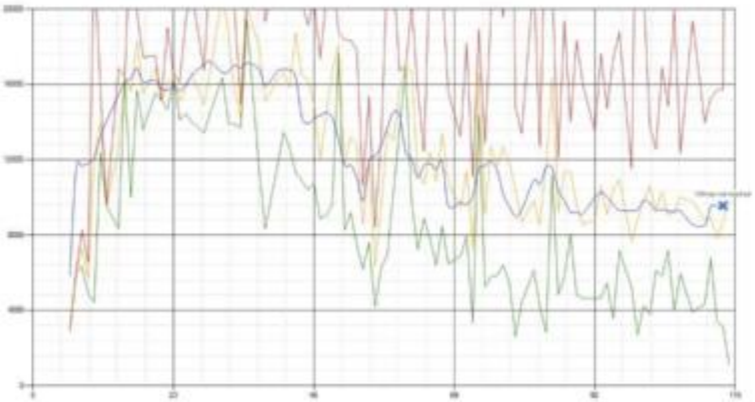

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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 represent system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times are 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 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. In short:

- 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 in system performance and scales to extremely large systems.

Single-Server Recommended Maximum Workload

For both the Citrix XenDesktop 7.9 Hosted Virtual Desktop and Citrix XenApp 7.9 RDS Hosted Shared Desktop use cases, a recommended maximum workload was determined that was based on both Login VSI Medium workload with flash end user experience measures 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%. (Memory should never be oversubscribed for Desktop Virtualization workloads.)

Callouts have been added throughout the data charts to indicate each phase of testing.

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time
Logon	The Login VSI phase of 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: 280 RDS sessions, 190 VDI Non-Persistent sessions, and 190 VDI Persistent sessions.

Single-Server Recommended Maximum Workload for RDS with 280 Users

Figure 140 Single Server Recommended Maximum Workload for RDS with 280 Users

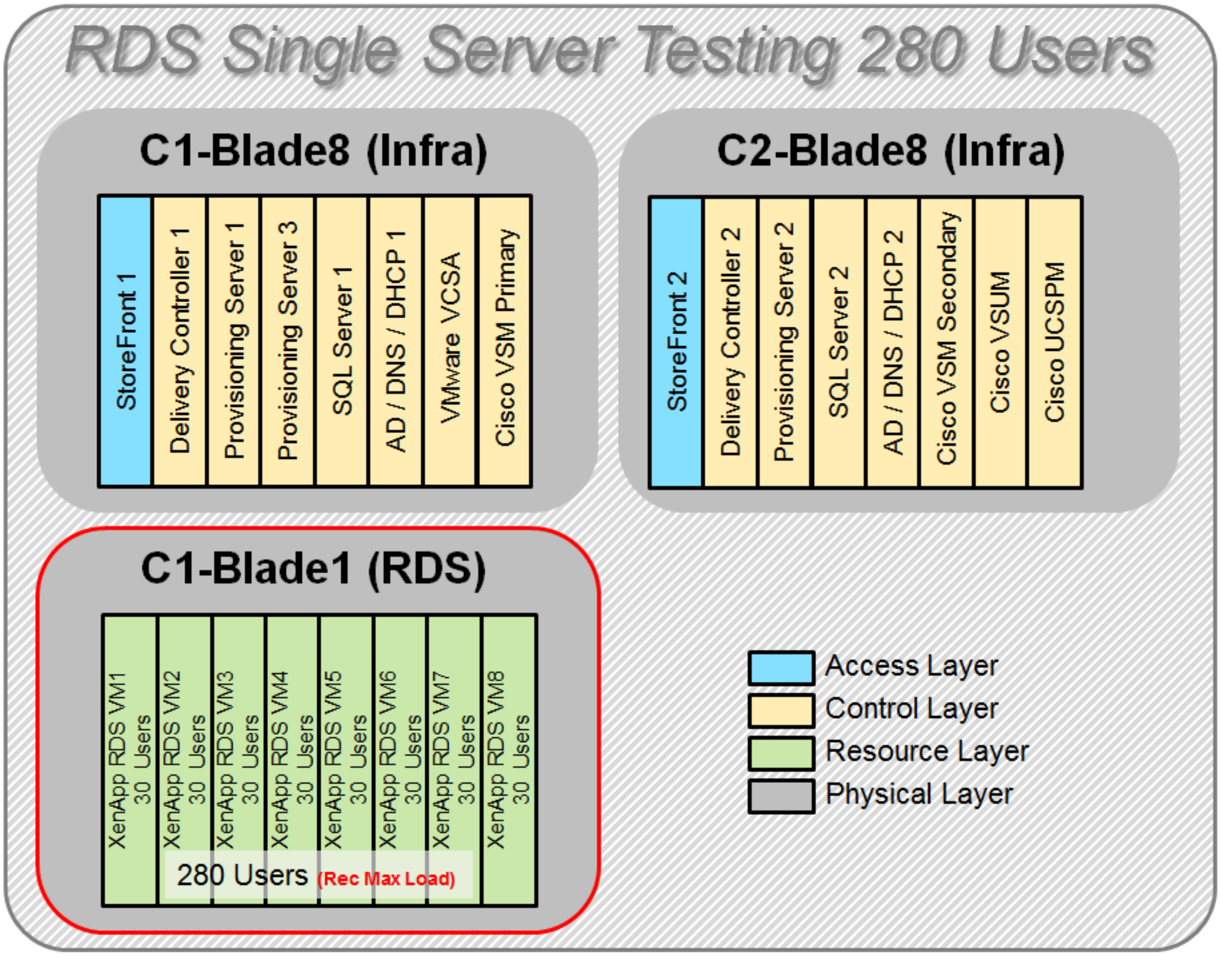

The recommended maximum workload for a B200 M4 blade server with dual E5-2680 v4 processors and 256GB of RAM is 280 Server 2012 R2 Hosted Shared Desktops. Each dedicated blade server ran 8 Server 2012 R2 Virtual Machines. Each virtual server was configured with 6 vCPUs and 24GB RAM.

Figure 141 Single Server | XenApp 7.9 RDS | VSI Score

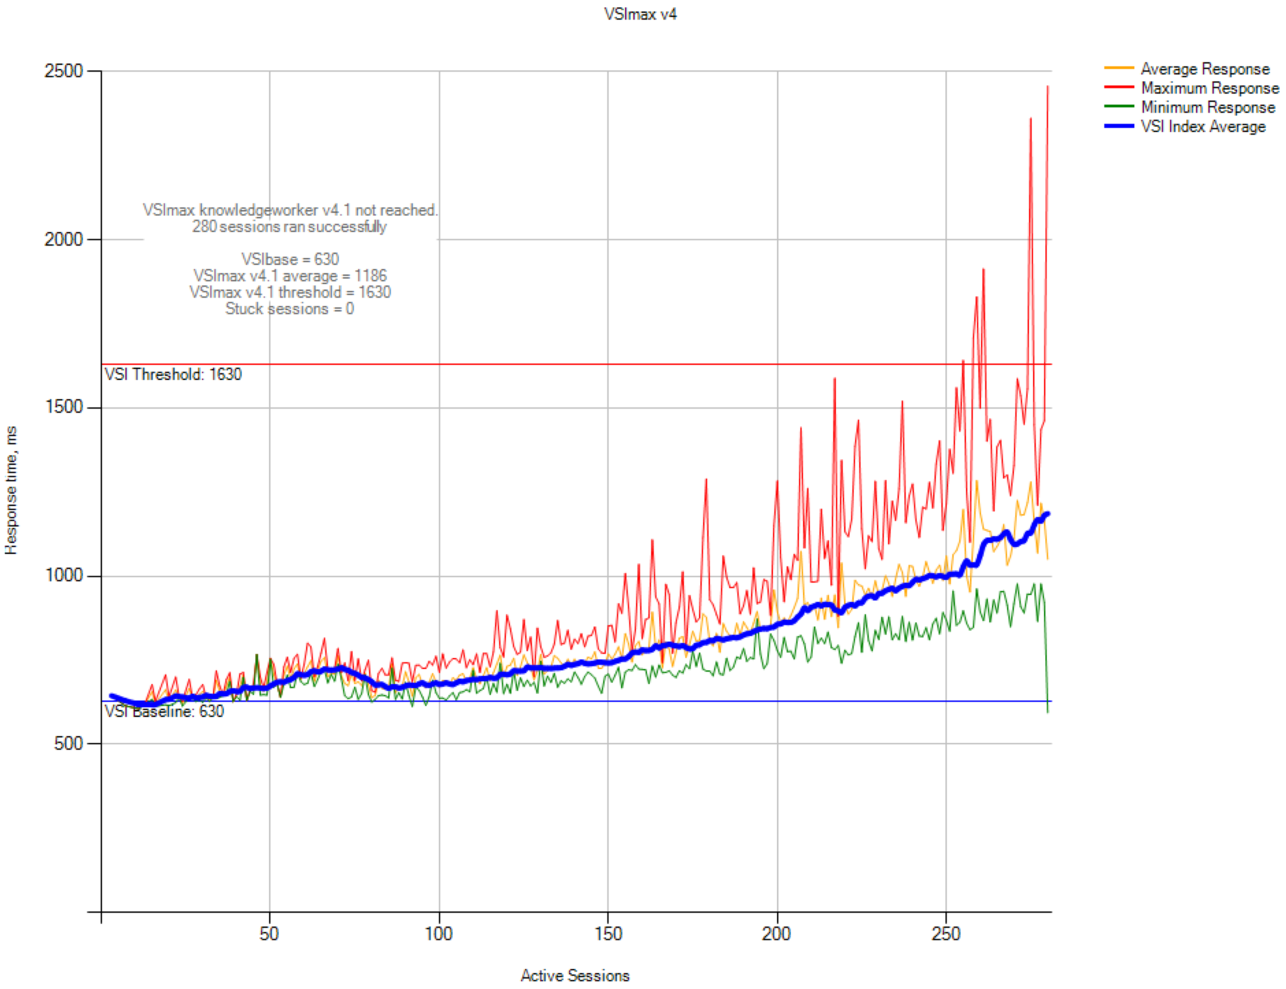

Performance data for the server running the workload follows:

Figure 142 Single Server | XenApp 7.9 RDS | Host CPU Utilization

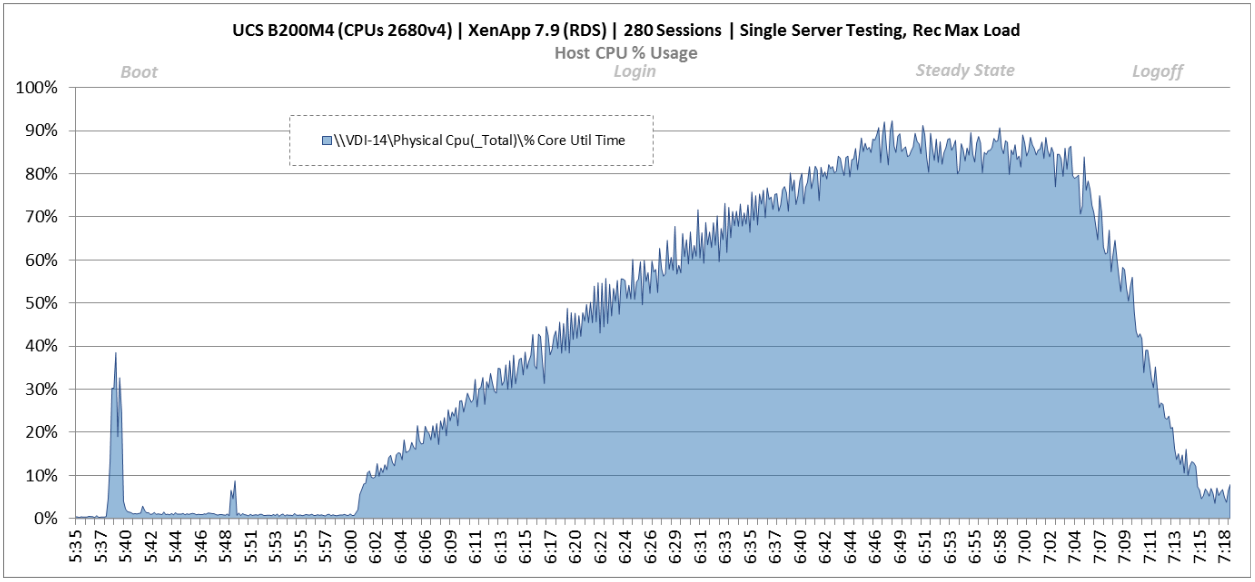

Figure 143 Single Server | XenApp 7.9 RDS | Host Memory Utilization

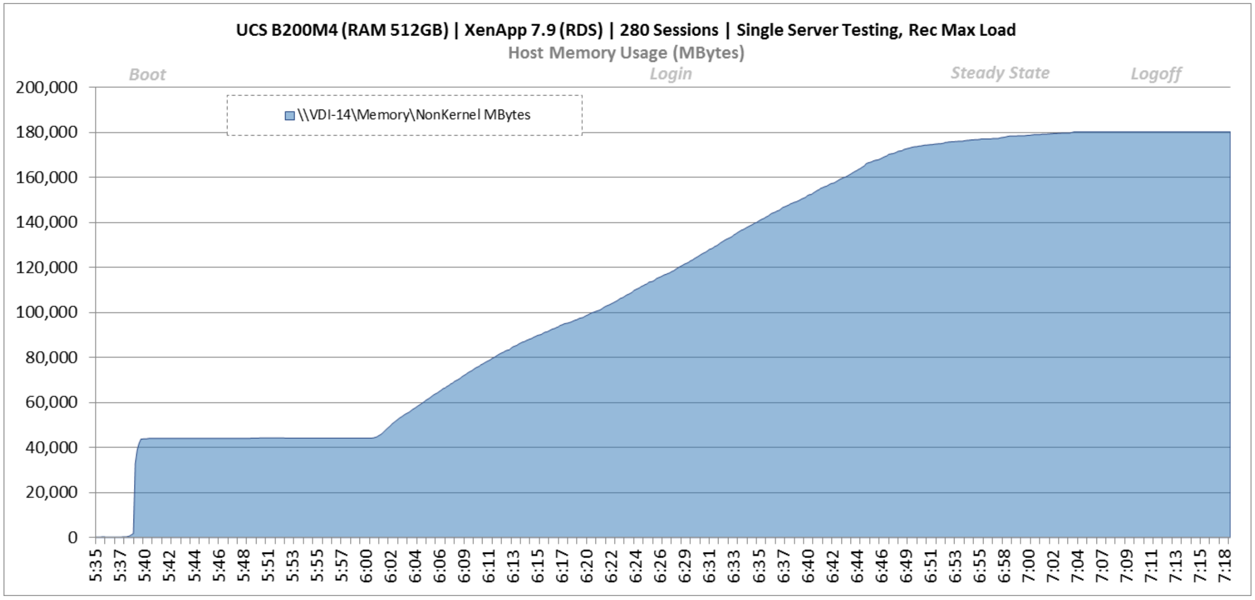

Figure 144 Single Server | XenApp 7.9 RDS | Host Network Utilization

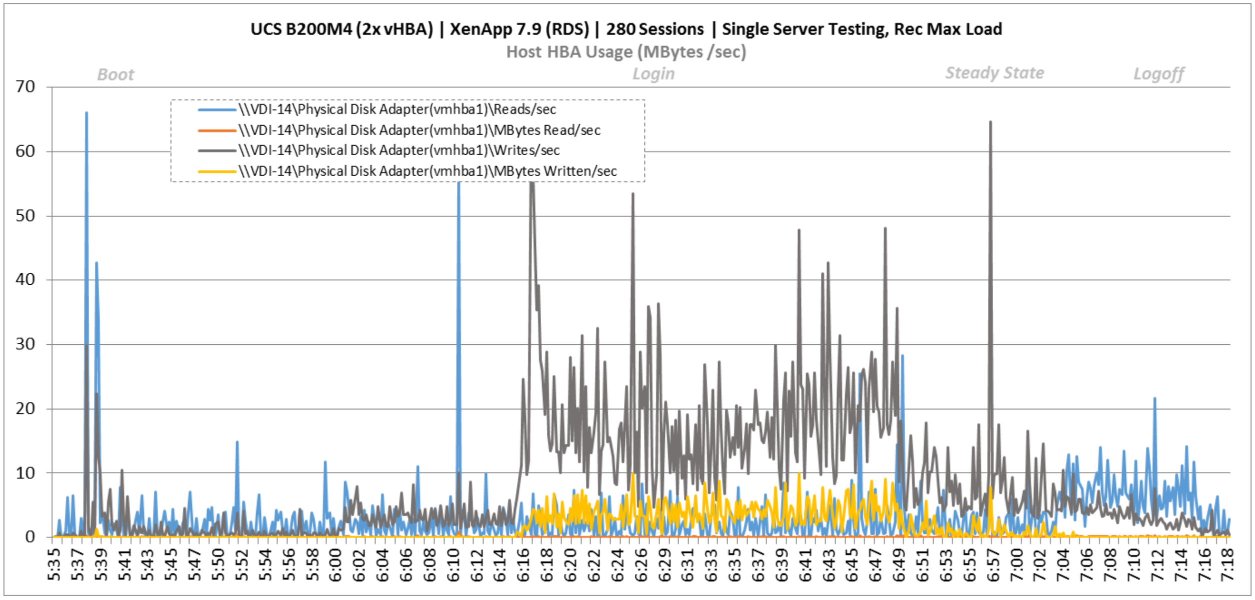

Single-Server Recommended Maximum Workload for VDI Non-Persistent with 190 Users

Figure 145 Single Server Recommended Maximum Workload for VDI Non-Persistent with 190 Users

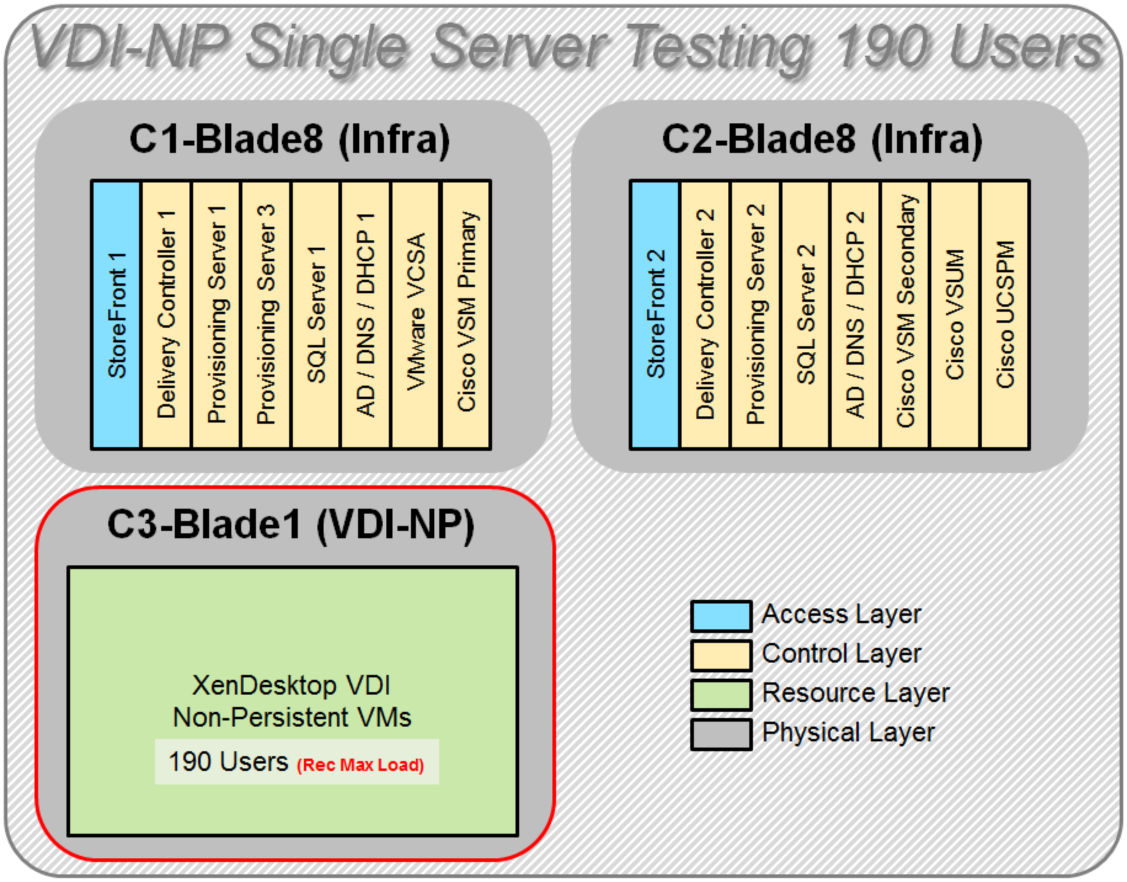

The recommended maximum workload for a B200 M4 blade server with dual E5-2680 v4 processors and 512GB of RAM is 190 Windows 10 64-bit virtual machines with 2 vCPU and 2GB RAM. Login VSI and blade performance data follows.

Figure 146 Single Server | XenDesktop 7.9 VDI-NP | VSI Score

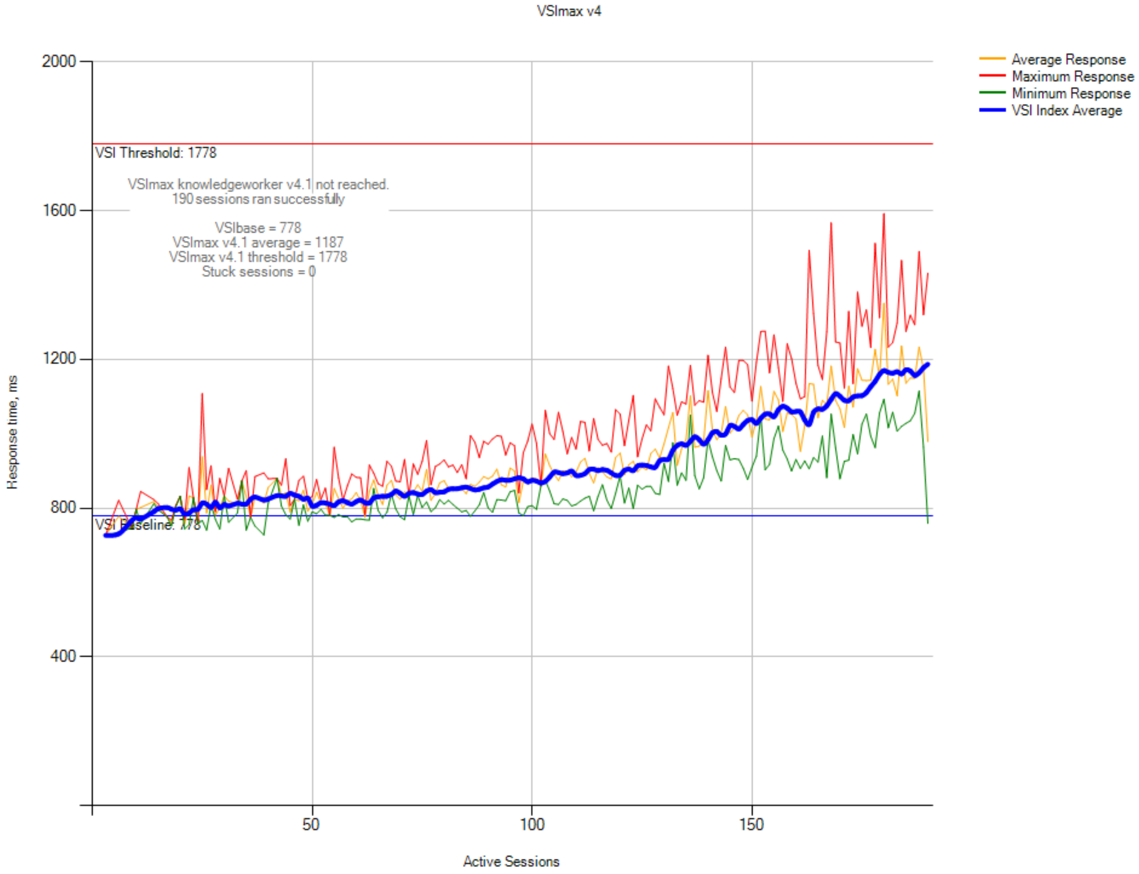

Performance data for the server running the workload follows:

Figure 147 Single Server | XenDesktop 7.9 VDI-NP | Host CPU Utilization

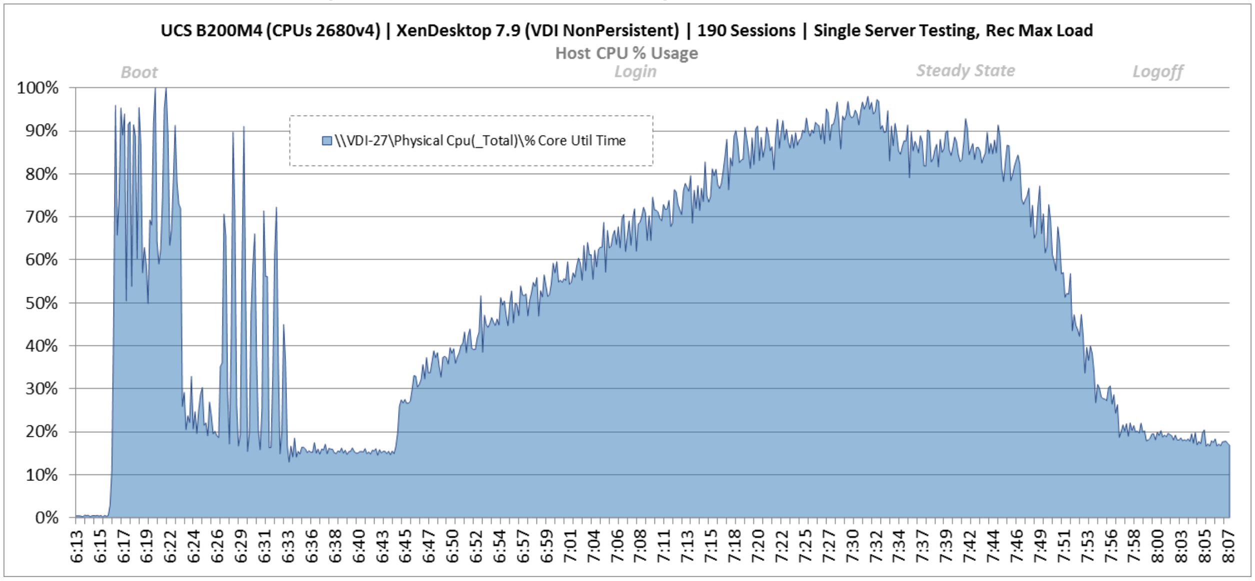

Figure 148 Single Server | XenDesktop 7.9 VDI-NP | Host Memory Utilization

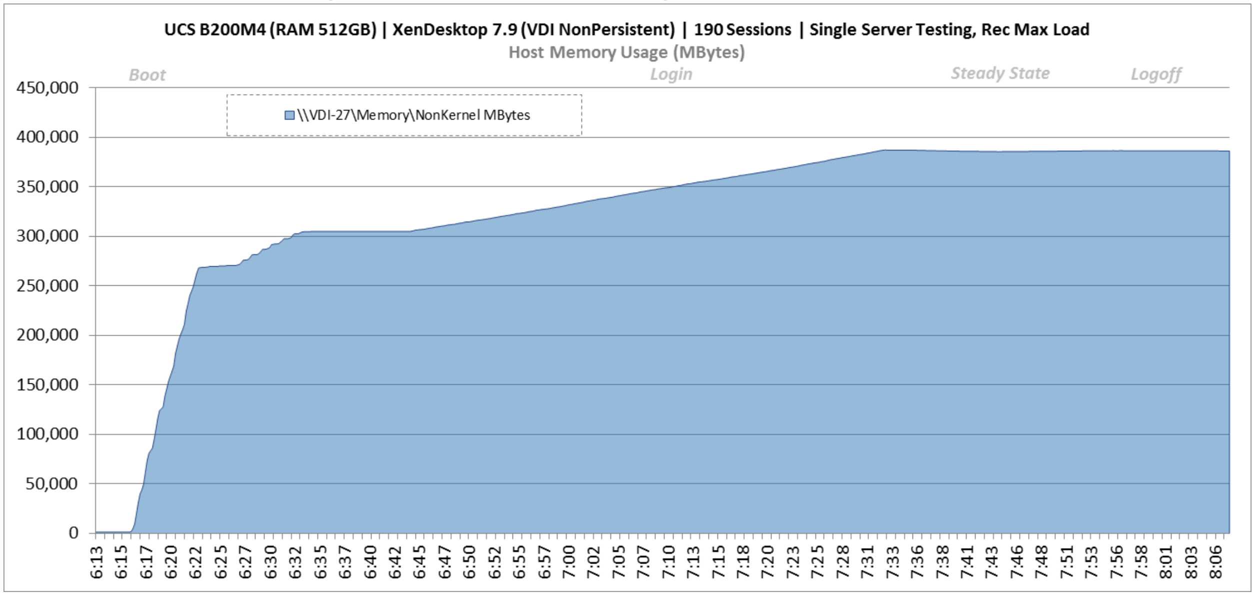

Figure 149 Single Server | XenDesktop 7.9 VDI-NP | Host Network Utilization

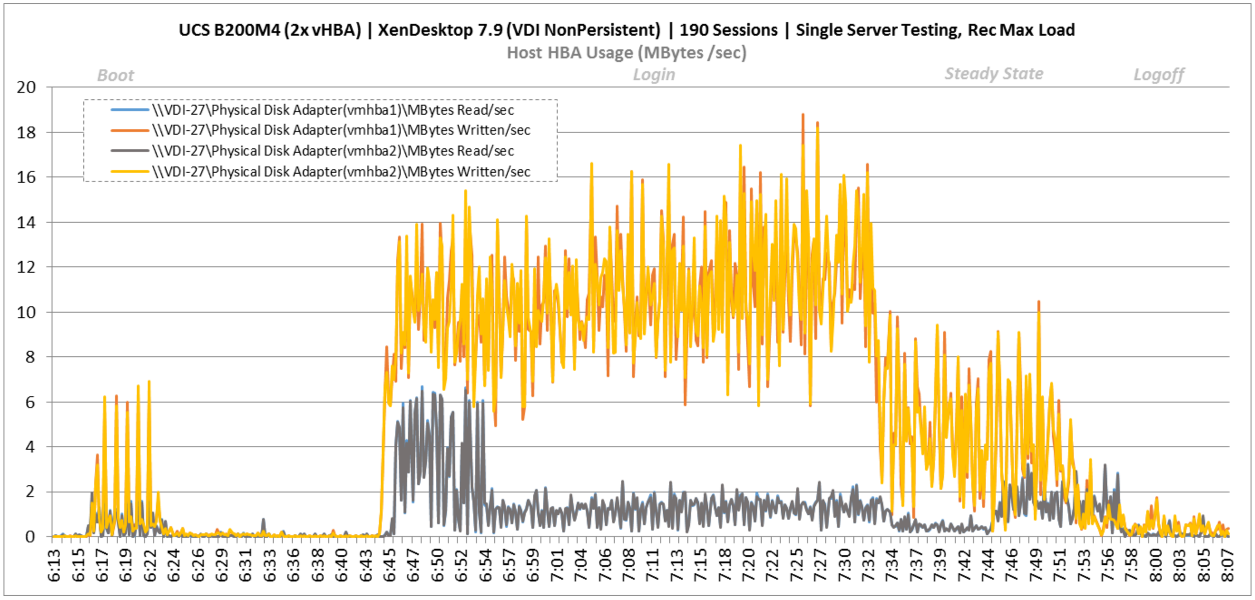

Single-Server Recommended Maximum Workload for VDI Persistent with 190 Users

Figure 150 Single Server Recommended Maximum Workload for VDI Persistent with 190 Users

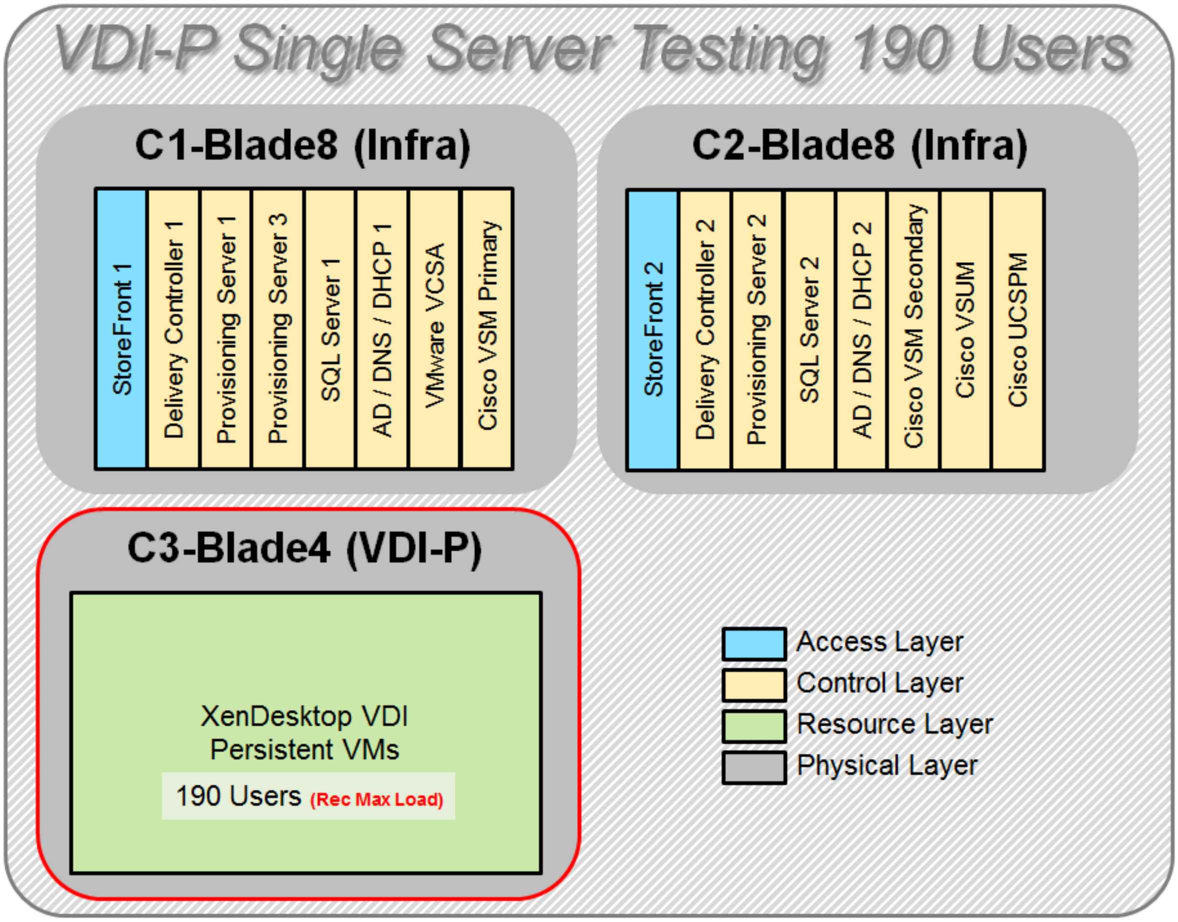

The recommended maximum workload for a B200 M4 blade server with dual E5-2680 v4 processors and 512GB of RAM is 190 Windows 10 64-bit virtual machines with 2 vCPU and 2GB RAM. Login VSI and blade performance data follows.

Figure 151 Single Server | XenDesktop 7.9 VDI-P | VSI Score

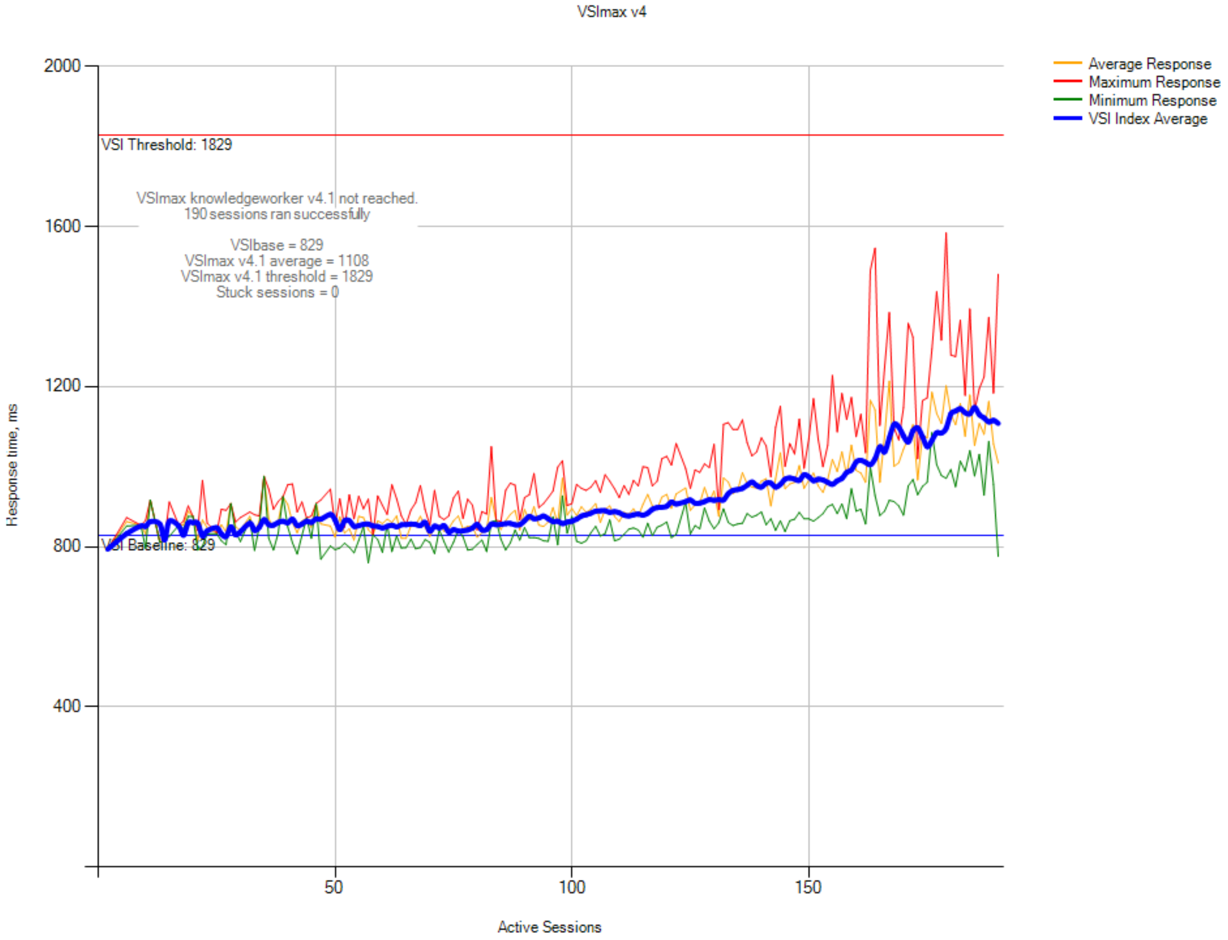

Performance data for the server running the workload follows:

Figure 152 Single Server | XenDesktop 7.9 VDI-P | Host CPU Utilization

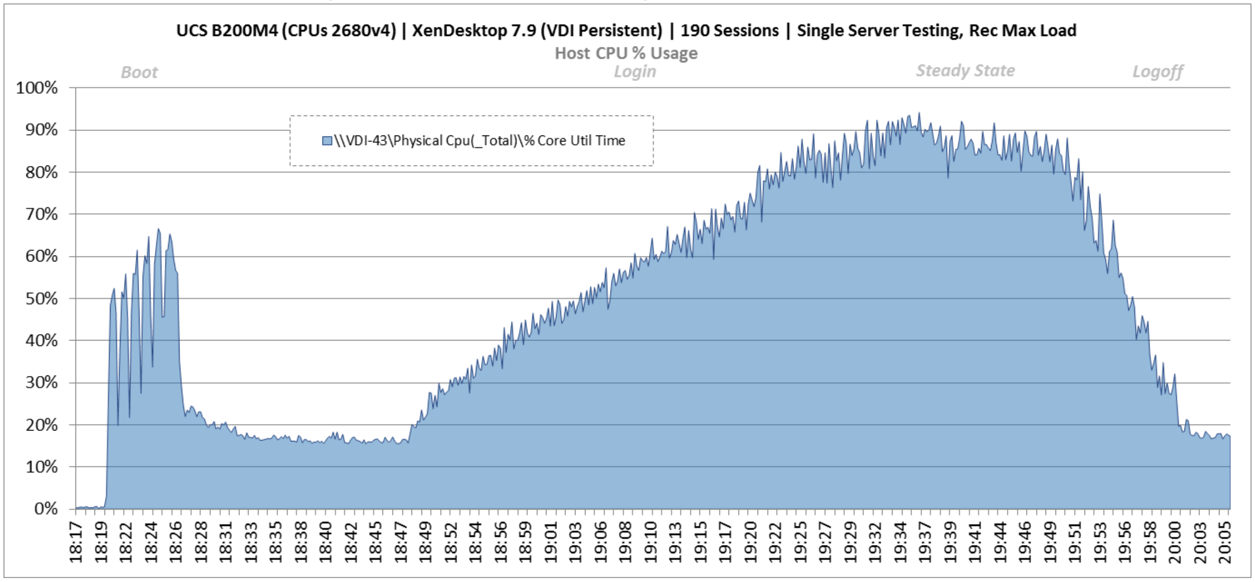

Figure 153 Single Server | XenDesktop 7.9 VDI-P | Host Memory Utilization

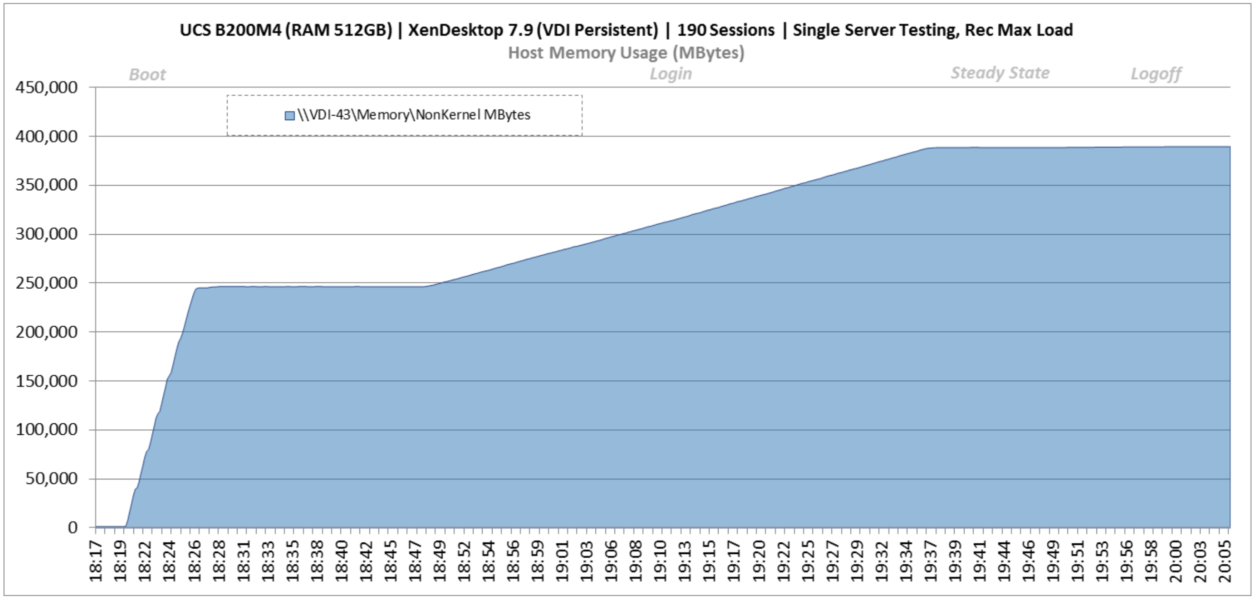

Figure 154 Single Server | XenDesktop 7.9 VDI-P | Host Network Utilization

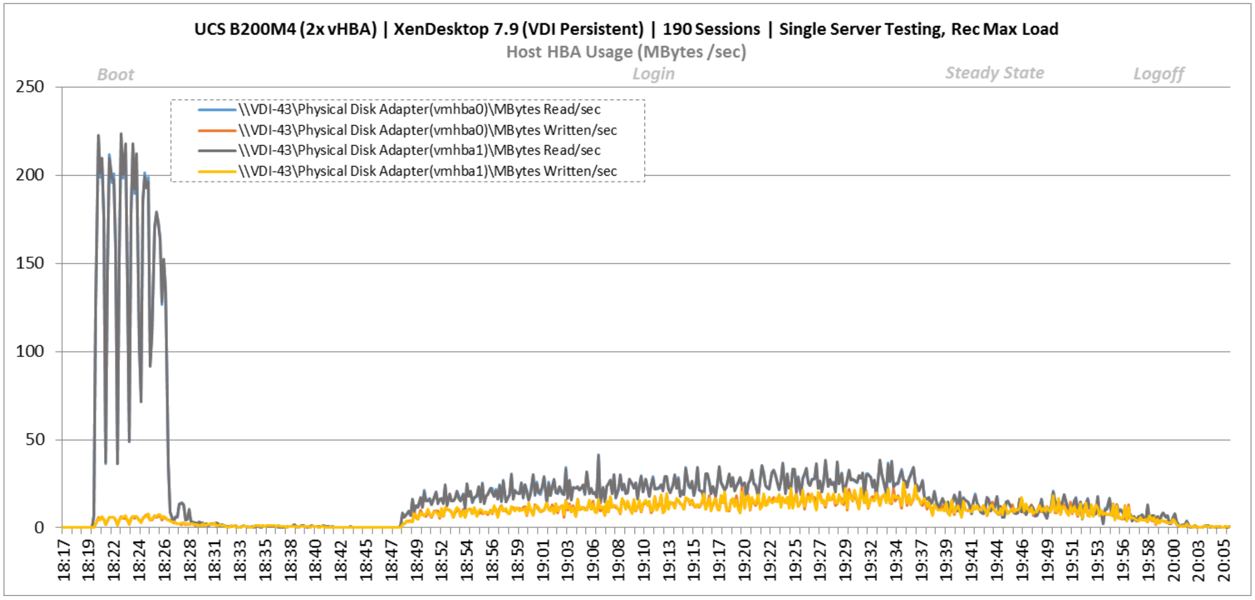

Cluster Workload Testing with 2600 RDS Users

This section shows the key performance metrics that were captured on the Cisco UCS, Pure Storage FlashArray//m, and Infrastructure VMs during the non-persistent desktop testing. The cluster testing with comprised of 2600 RDS sessions using 10 workload blades.

Figure 155 RDS Cluster Testing with 2600 Users

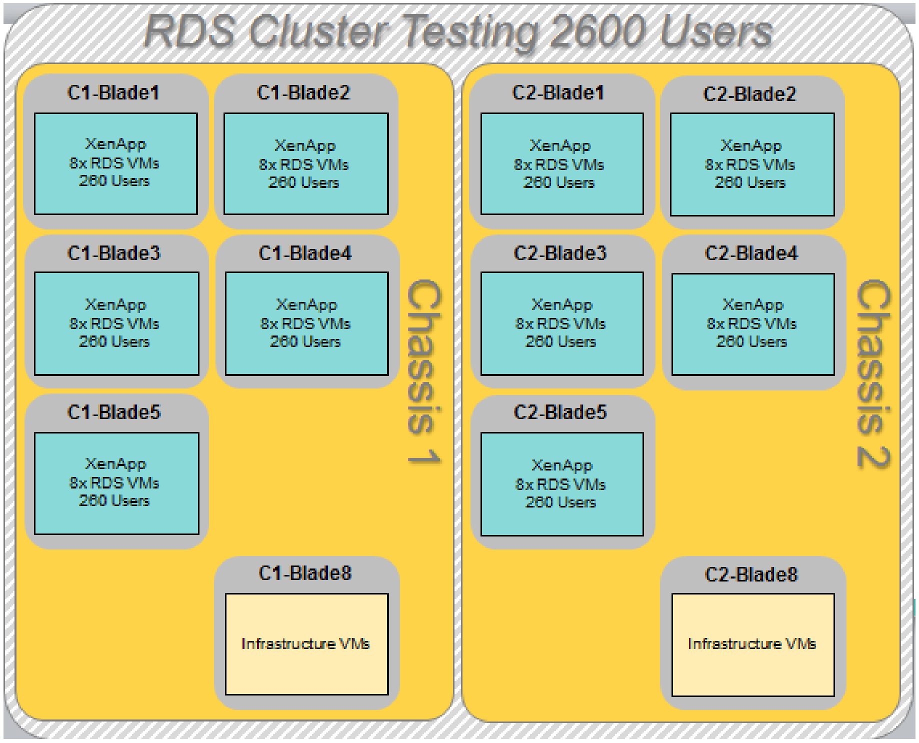

The workload for the test is 2600 RDS users. To achieve the target, sessions were launched against the single RDS cluster only. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

The configured system efficiently and effectively delivered the following results:

Figure 156 Cluster | 2600 RDS Users | VSI Score

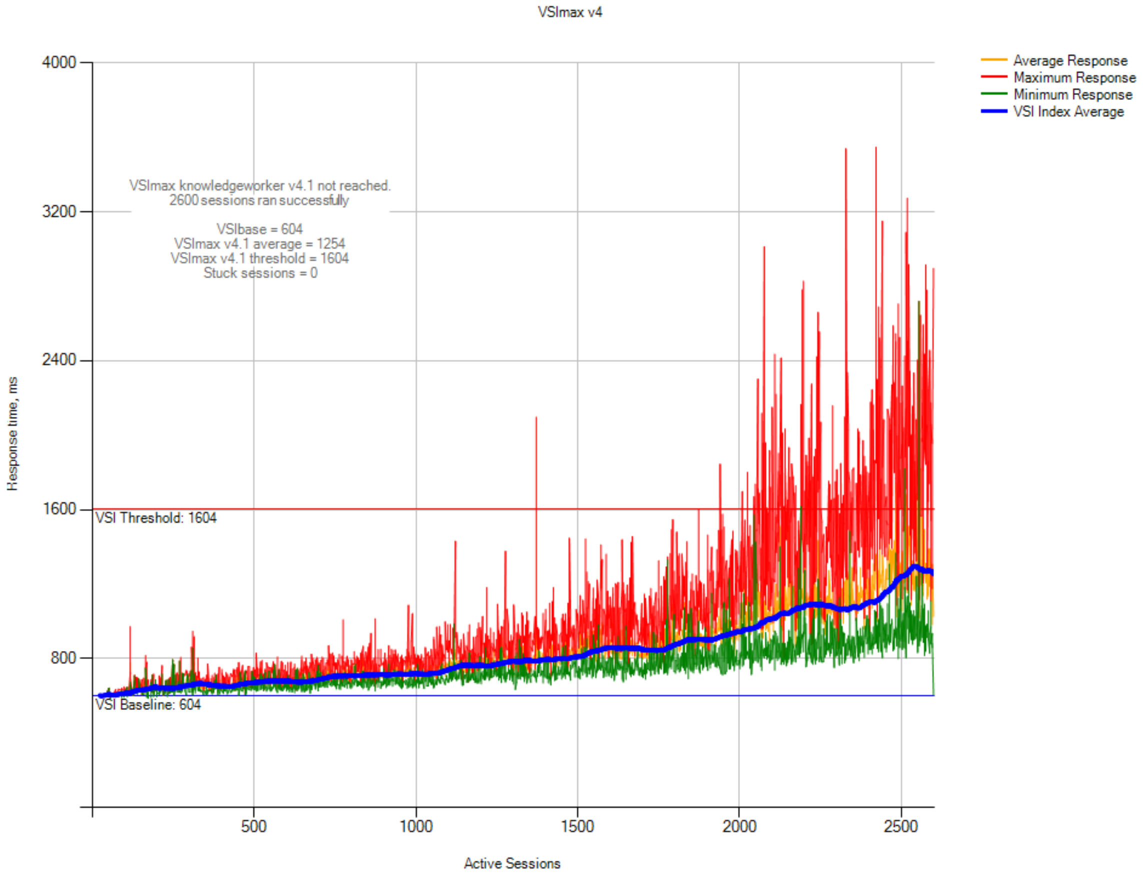

Figure 157 Cluster | 2600 RDS Users | Infrastructure Hosts | Host CPU Utilization

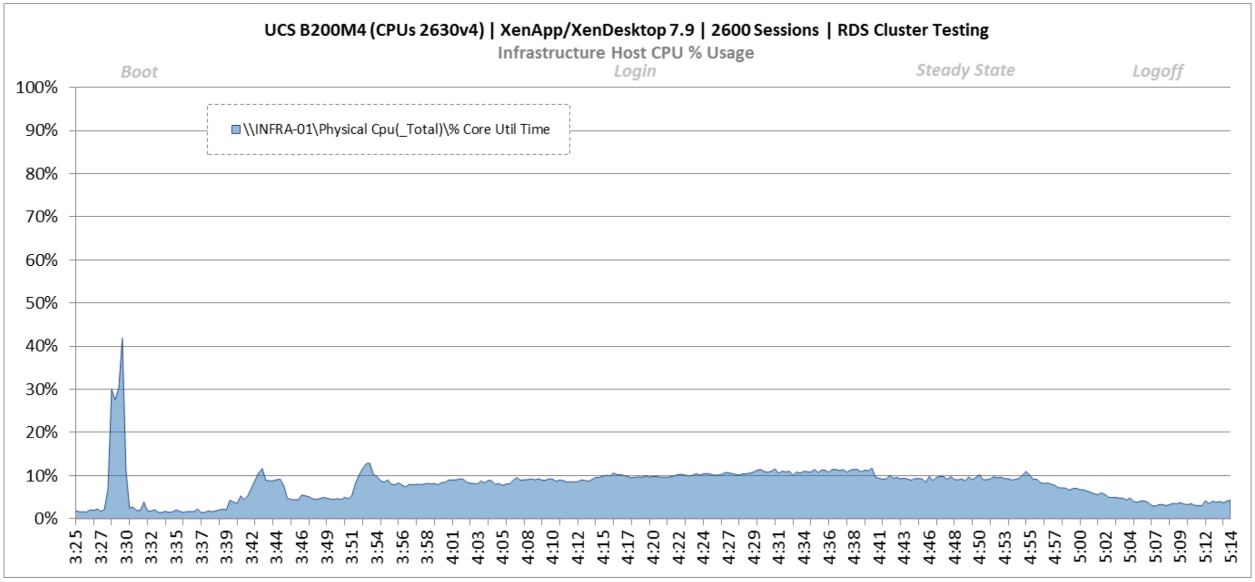

Figure 158 Cluster | 2600 RDS Users | 2 Infrastructure Hosts | Host Memory Utilization

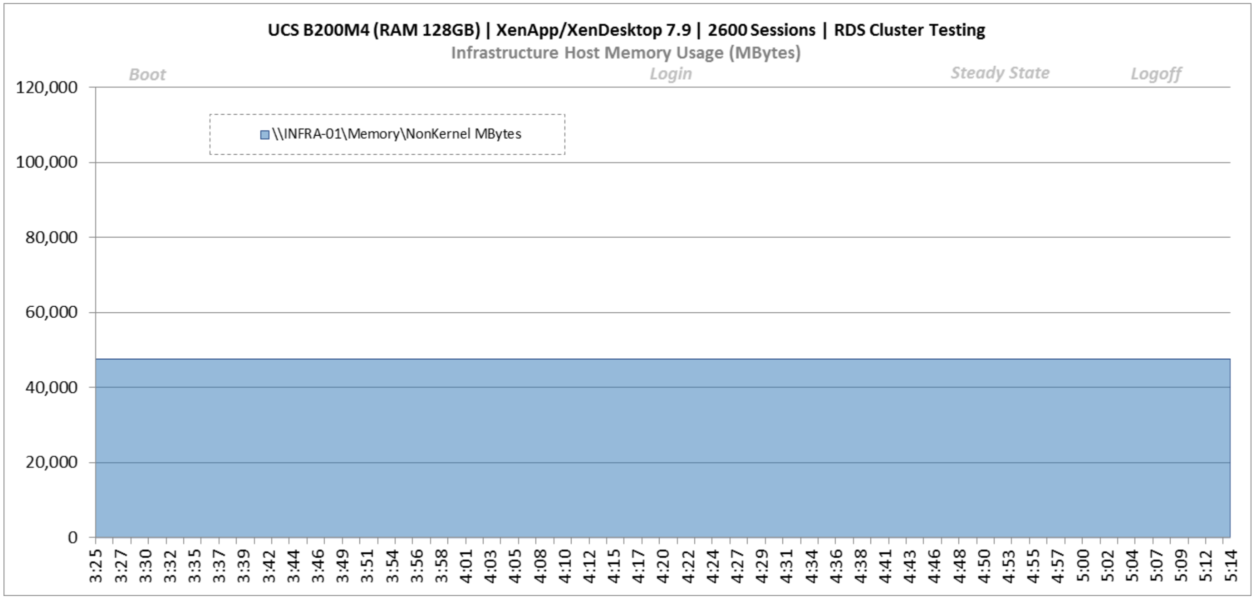

Figure 159 Cluster | 2600 RDS Users | 2 Infrastructure Hosts | Host System Network Utilization

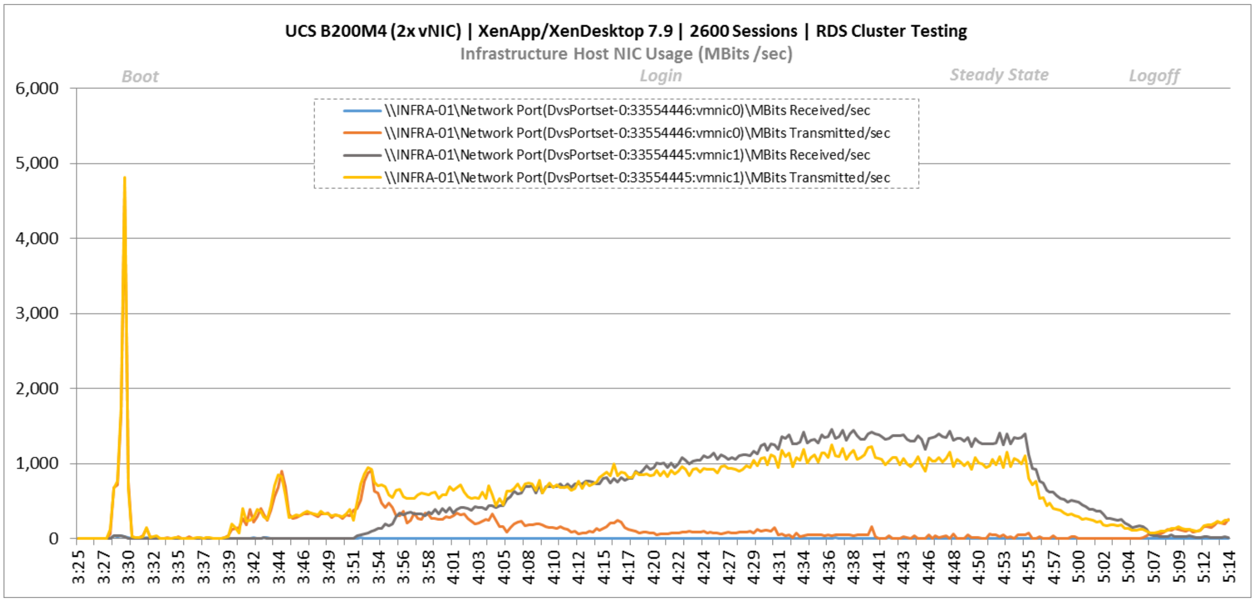

Figure 160 Cluster | 2600 RDS Users | 2 Infrastructure Hosts | Host vHBA Utilization

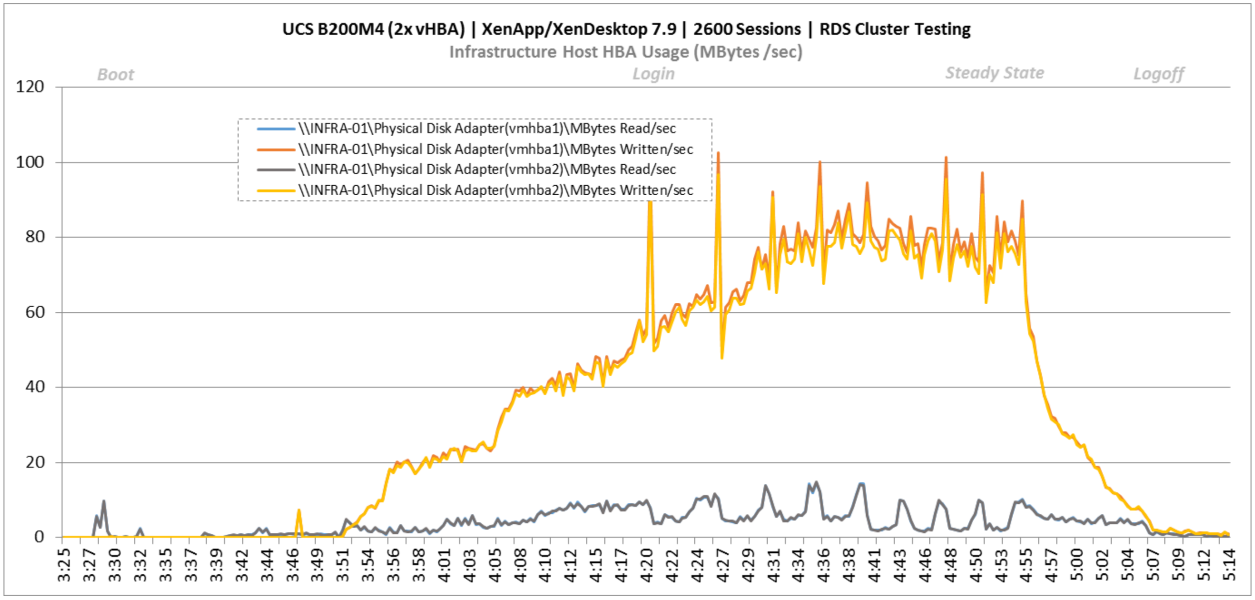

Figure 161 Cluster | 2600 RDS Users | 10 RDS Hosts | Sample Host CPU Utilization

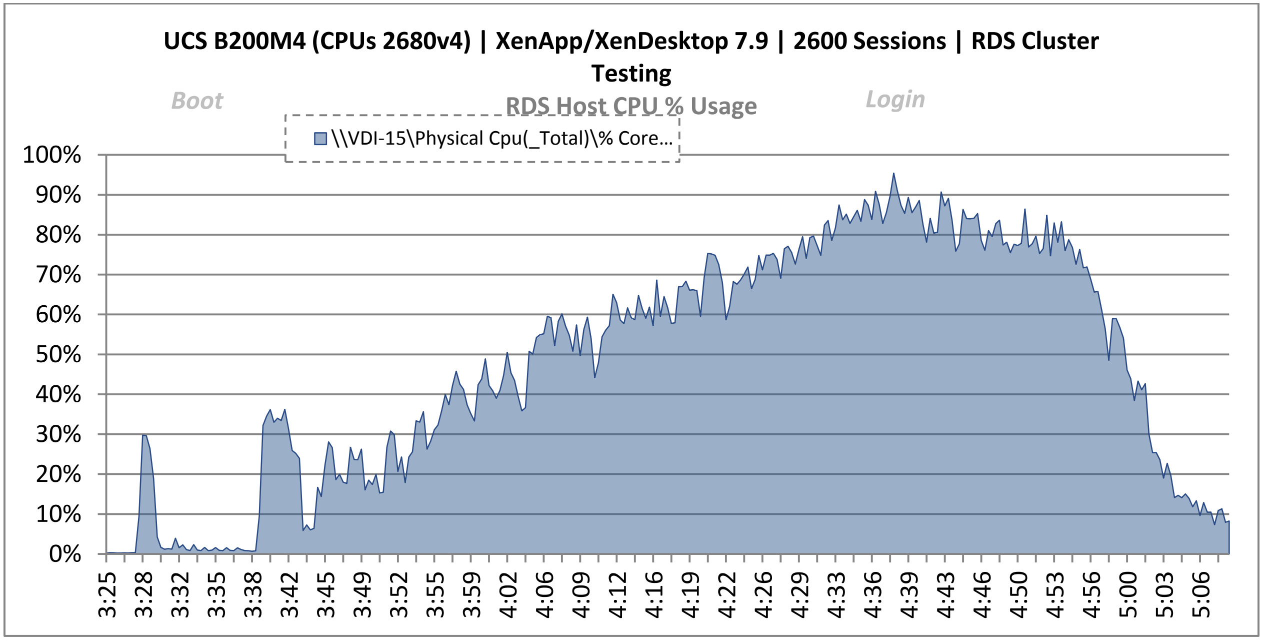

Figure 162 Cluster | 2600 RDS Users | 10 RDS Hosts | Sample Host Memory Utilization

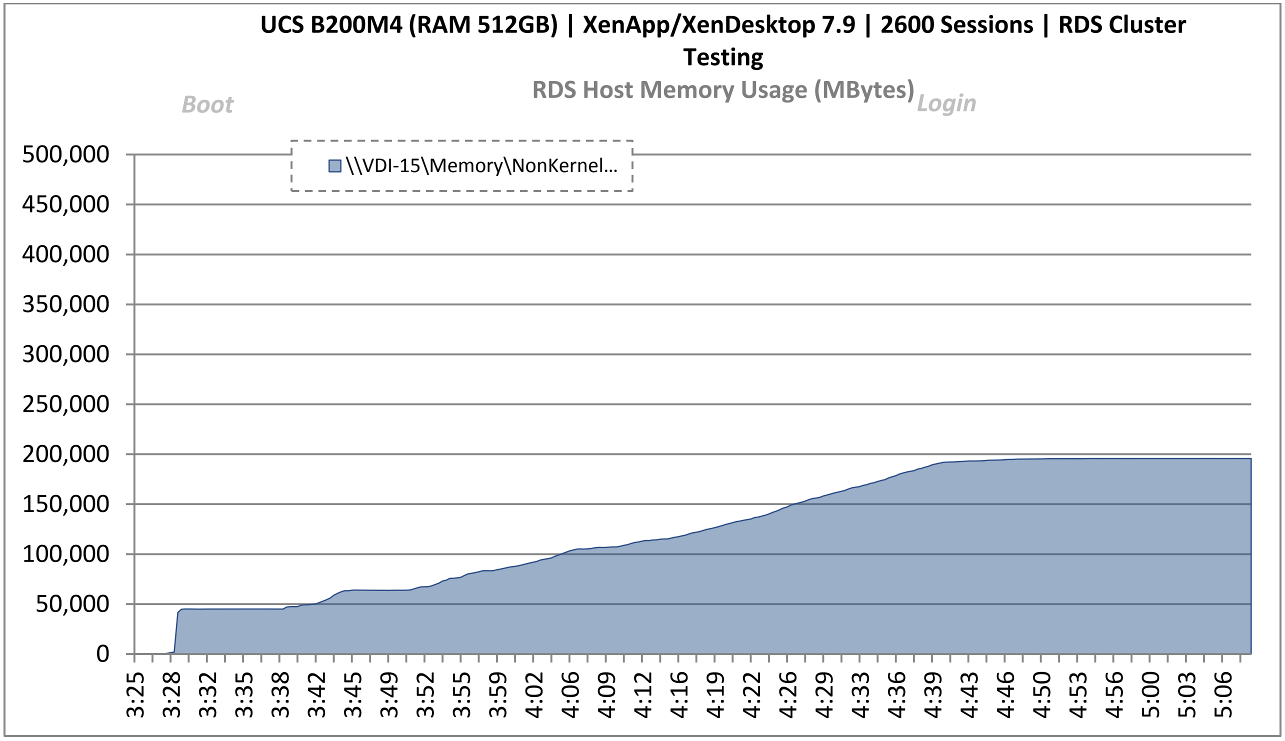

Figure 163 Cluster | 2600 RDS Users | 10 RDS Hosts | Sample Host Network Utilization

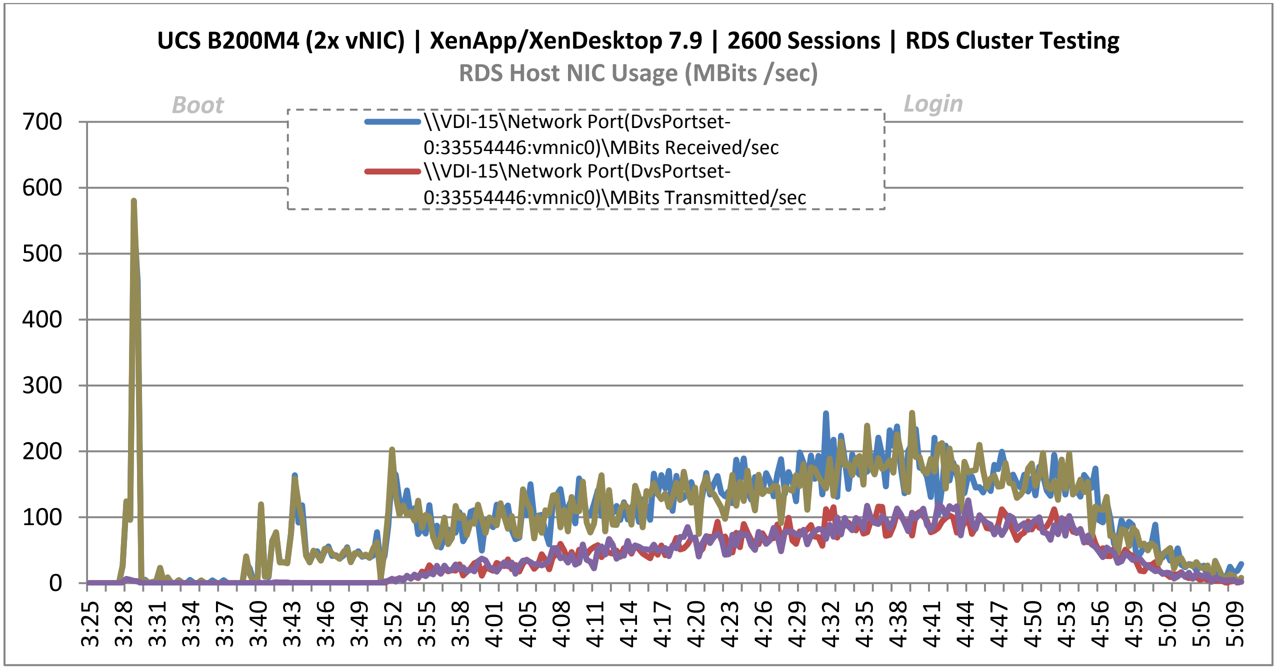

Figure 164 Cluster | 2600 RDS Users | 10 RDS Hosts | Sample Host Fibre Channel Network Utilization

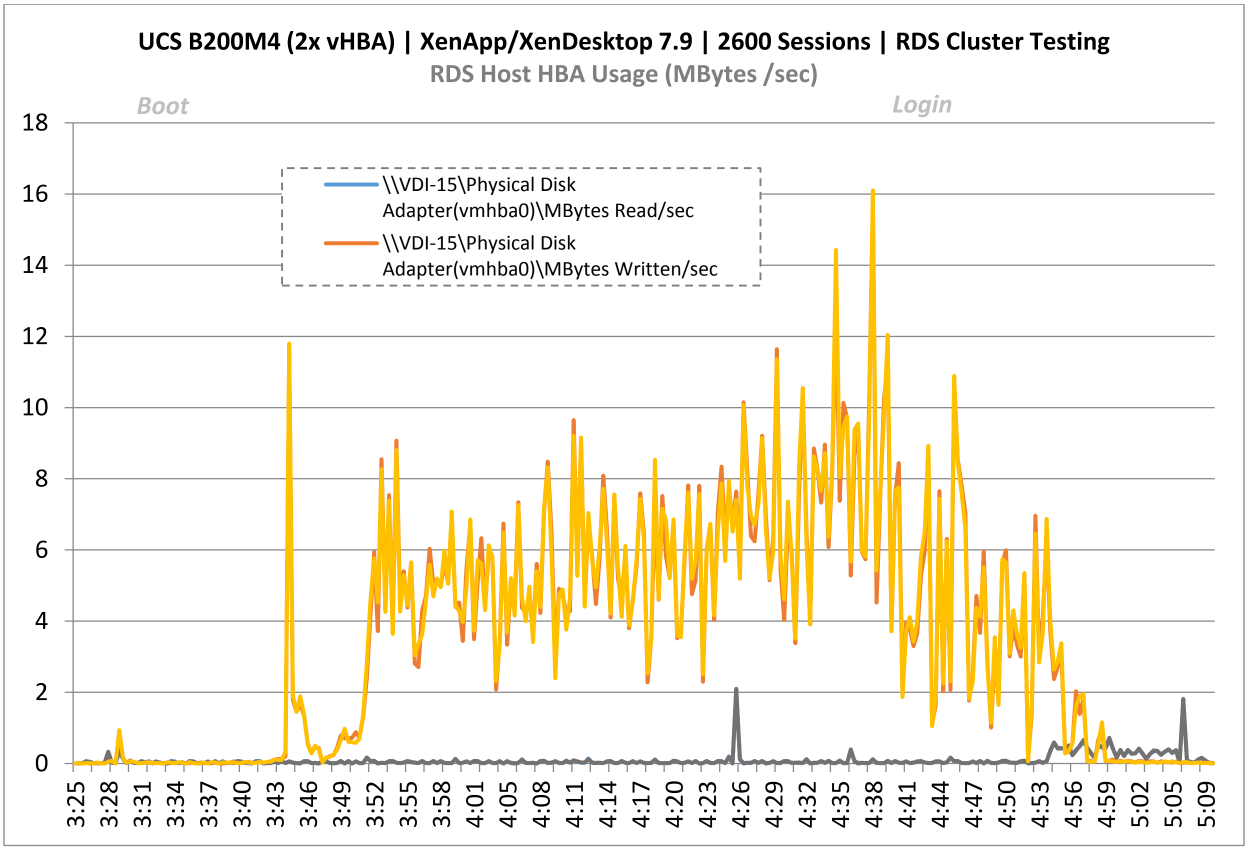

Cluster Workload Testing with 1200 Persistent Desktop Users

This section shows the key performance metrics that were captured on the Cisco UCS, Pure Storage FlashArray//m, and Infrastructure VMs during the persistent desktop testing. The cluster testing with comprised of 1200 VDI Persistent desktop sessions using 8 workload blades.

Figure 165 VDI Persistent Cluster Testing with 1200 Users

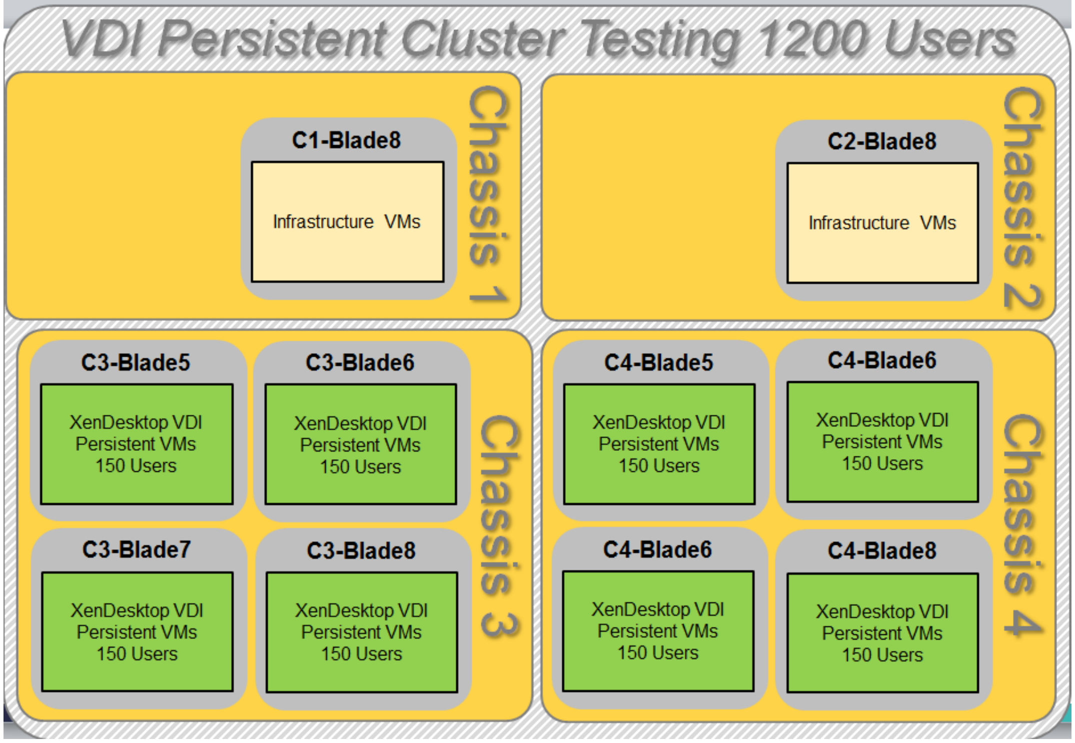
VDI Persistent Cluster Testing 1200 Users
C1-Blade8
Infrastructure VMs
Chassis 1
C2-Blade8
Infrastructure VMs
Chassis 2
C3-Blade5
XenDesktop VDI Persistent VMs 150 Users
C3-Blade6
XenDesktop VDI Persistent VMs 150 Users
C3-Blade7
XenDesktop VDI Persistent VMs 150 Users
C3-Blade8
XenDesktop VDI Persistent VMs 150 Users
Chassis 3
C4-Blade5
XenDesktop VDI Persistent VMs 150 Users
C4-Blade6
XenDesktop VDI Persistent VMs 150 Users
C4-Blade6
XenDesktop VDI Persistent VMs 150 Users
C4-Blade8
XenDesktop VDI Persistent VMs 150 Users
Chassis 4

The workload for the test is 1200 persistent desktop users. To achieve the target, sessions were launched against the single persistent cluster only. 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 166 Cluster | 1200 VDI-P Users | VSI Score

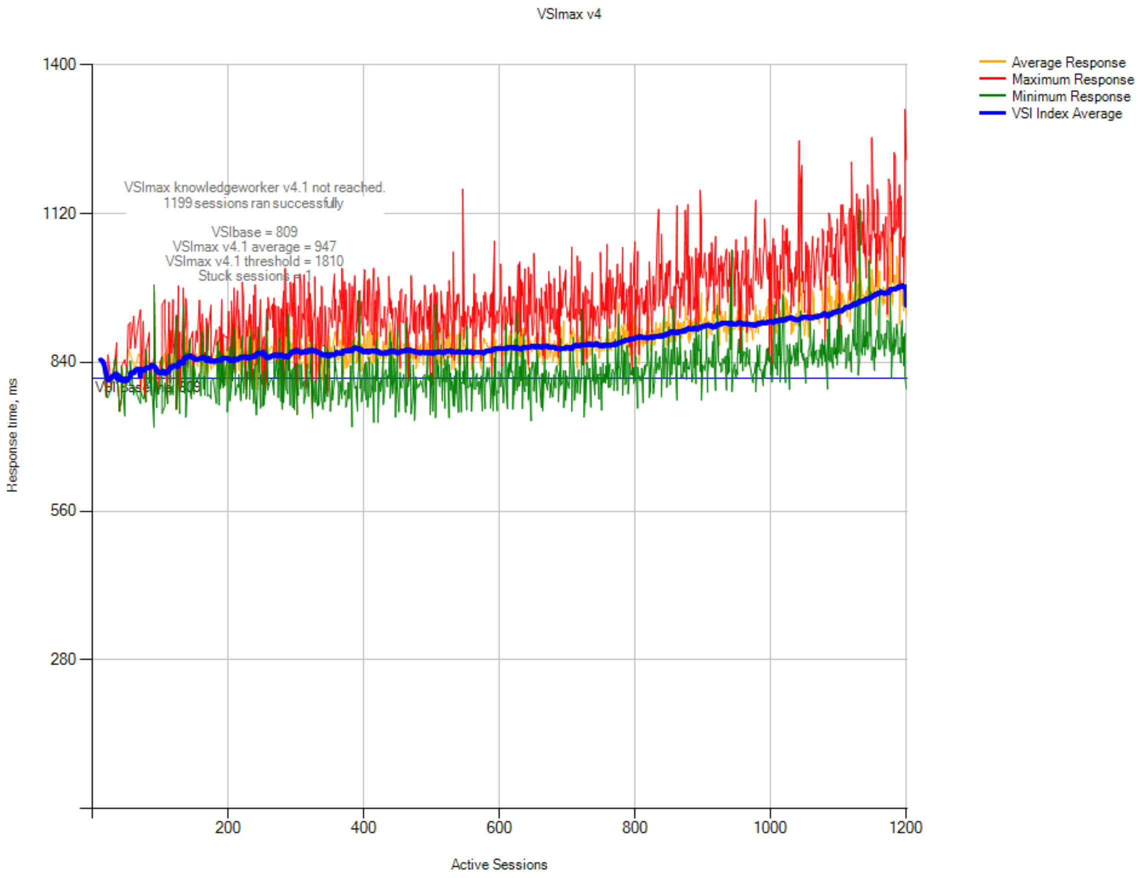

Figure 167 Cluster | 1200 VDI-P Users | Infrastructure Hosts | Host CPU Utilization

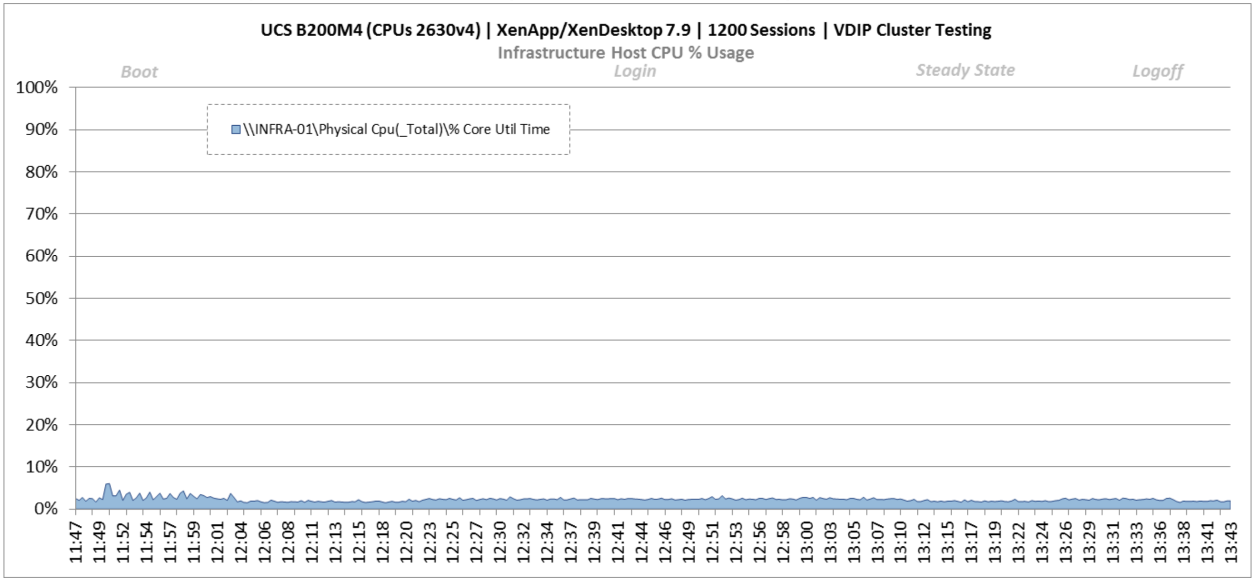

Figure 168 Cluster | 1200 VDI-P Users | Infrastructure Hosts | Host Memory Utilization

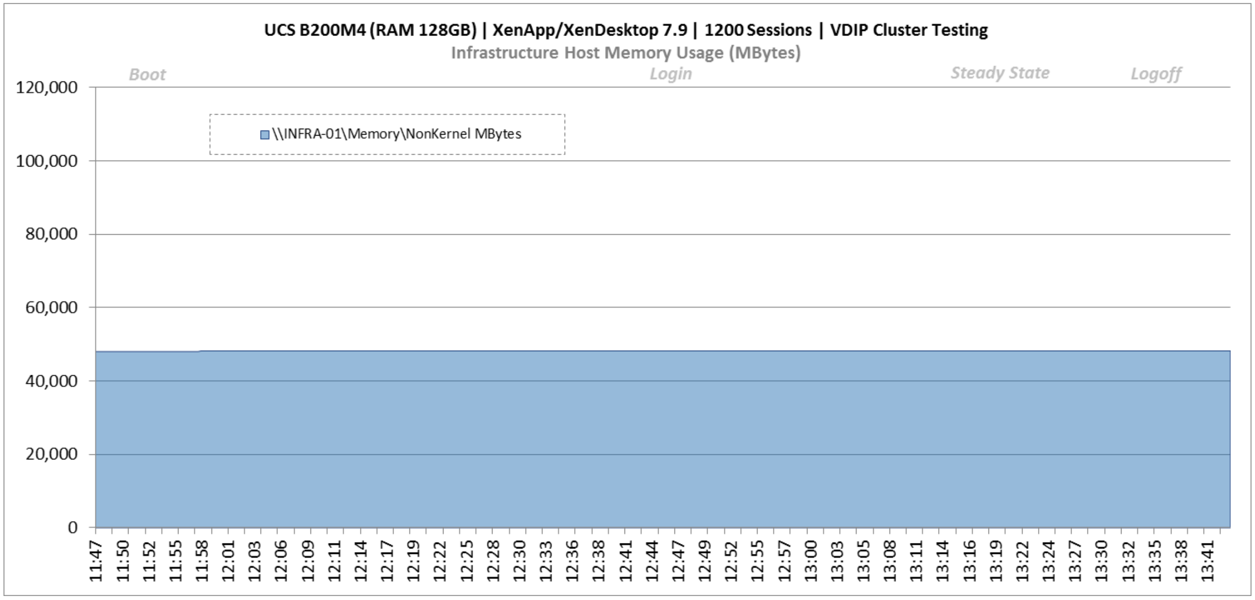

Figure 169 Cluster | 1200 VDI-P Users | Infrastructure Hosts | Host Network Utilization

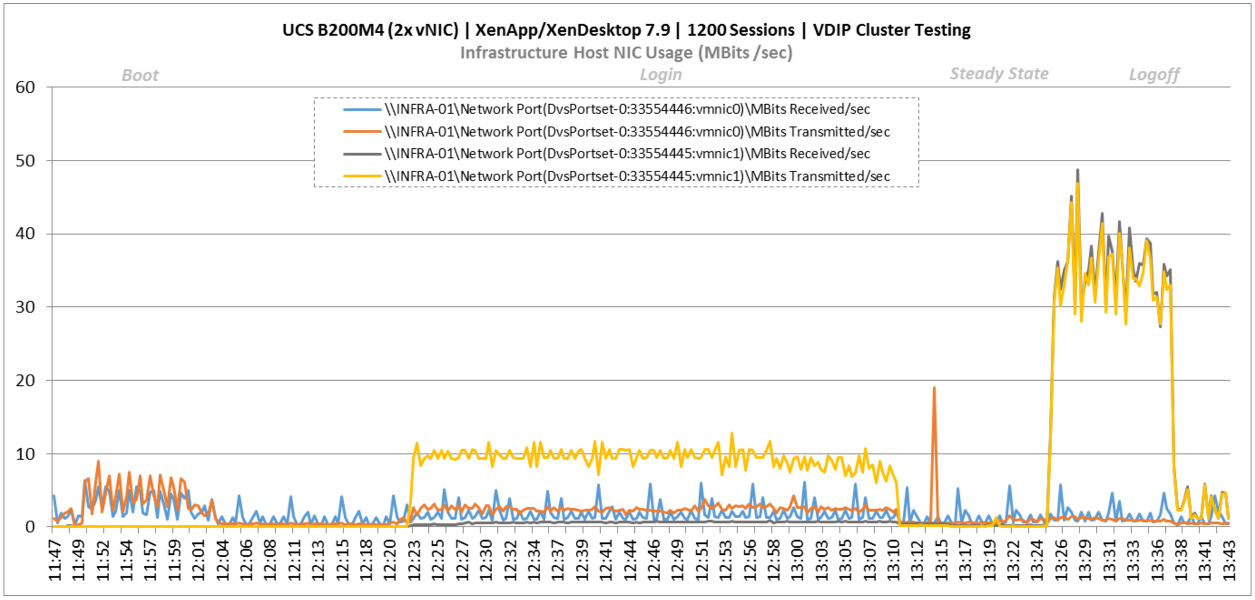

Figure 170 Cluster | 1200 VDI-P Users | Infrastructure Hosts | Host Fibre Channel Network Utilization

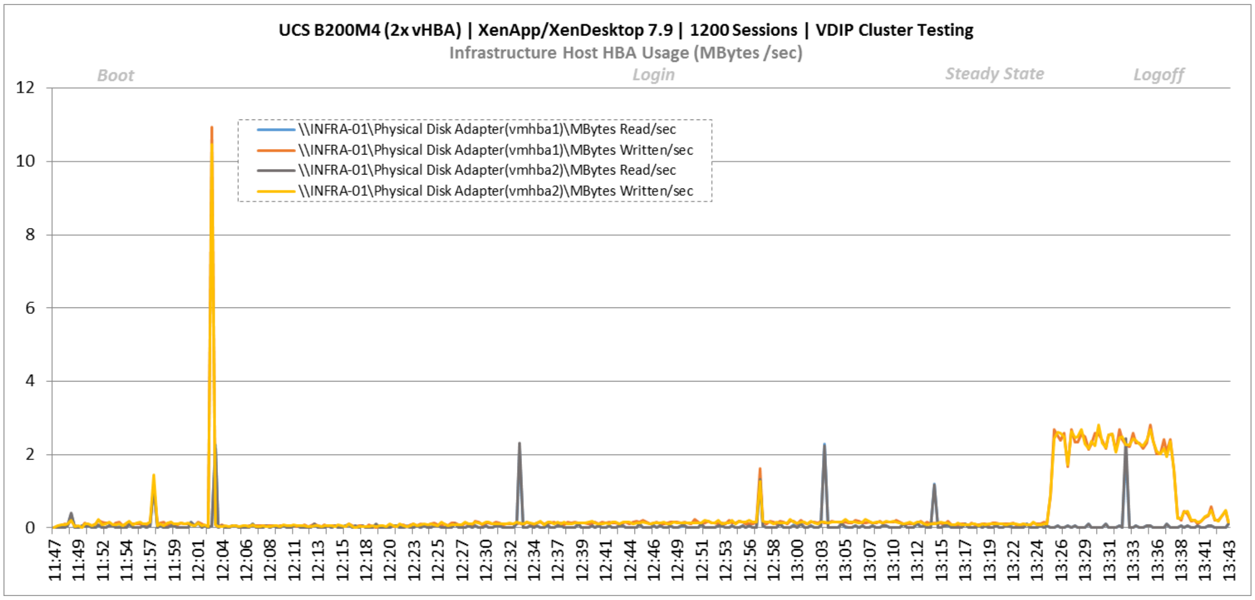

Figure 171 Cluster | 1200 VDI-P Users | Persistent Hosts | Sample Host CPU Utilization

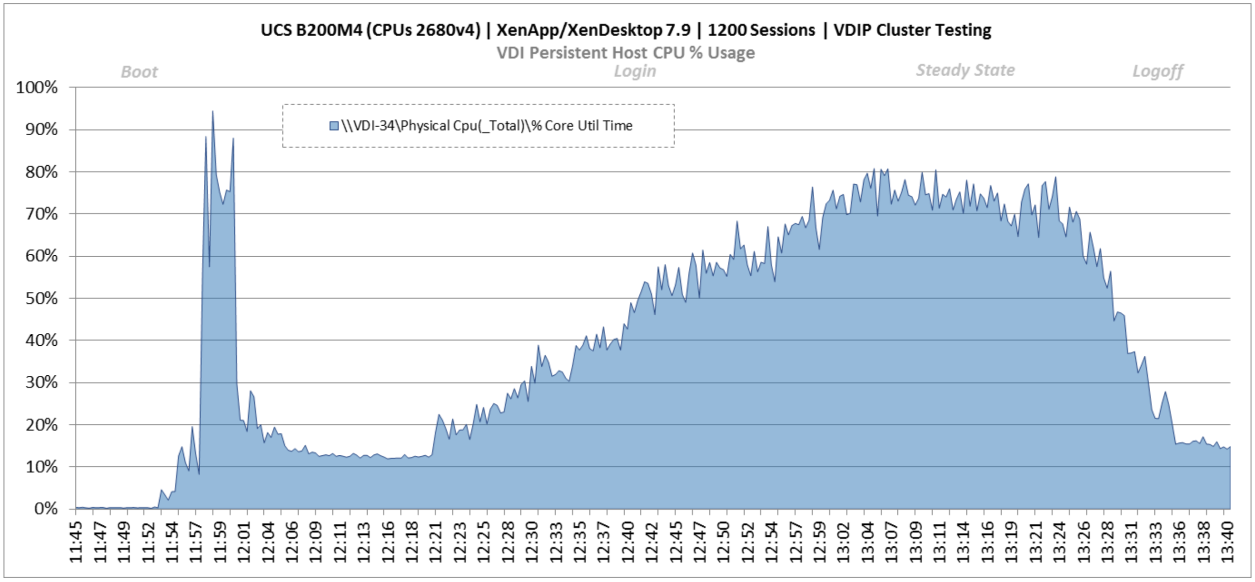

Figure 172 Cluster | 1200 VDI-P Users | Persistent Hosts | Sample Host Memory Utilization

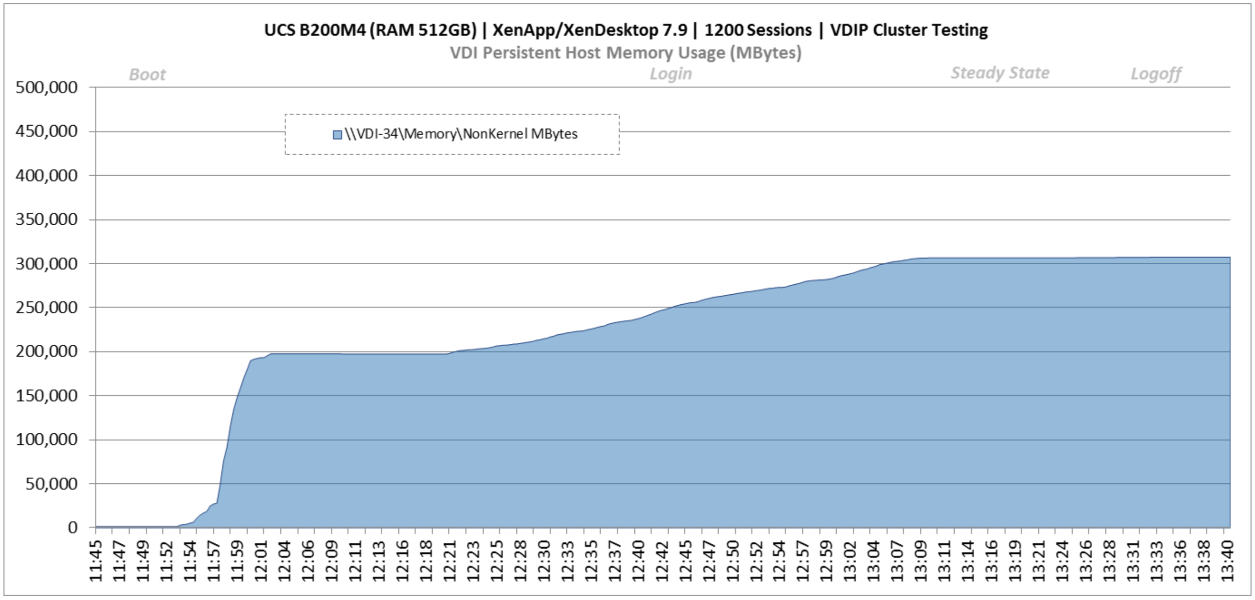

Figure 173 Cluster | 1200 VDI-P Users | Persistent Hosts | Sample Host Network Utilization

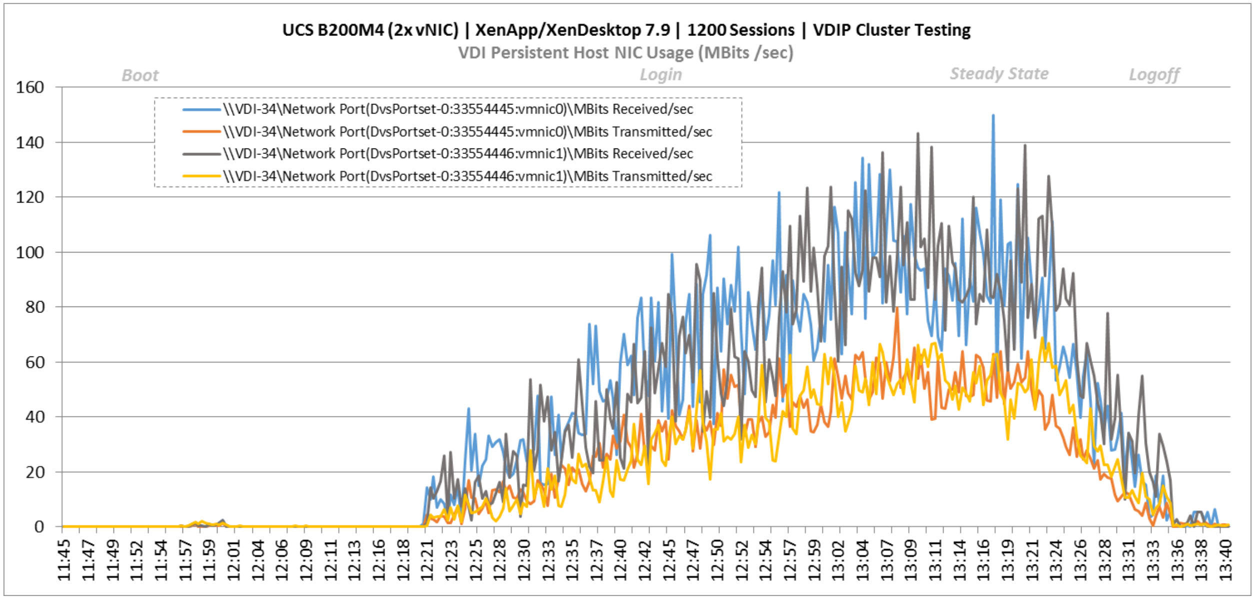

Figure 174 Cluster | 1200 VDI-P Users | Persistent Hosts | Sample Host Fibre Channel Network Utilization

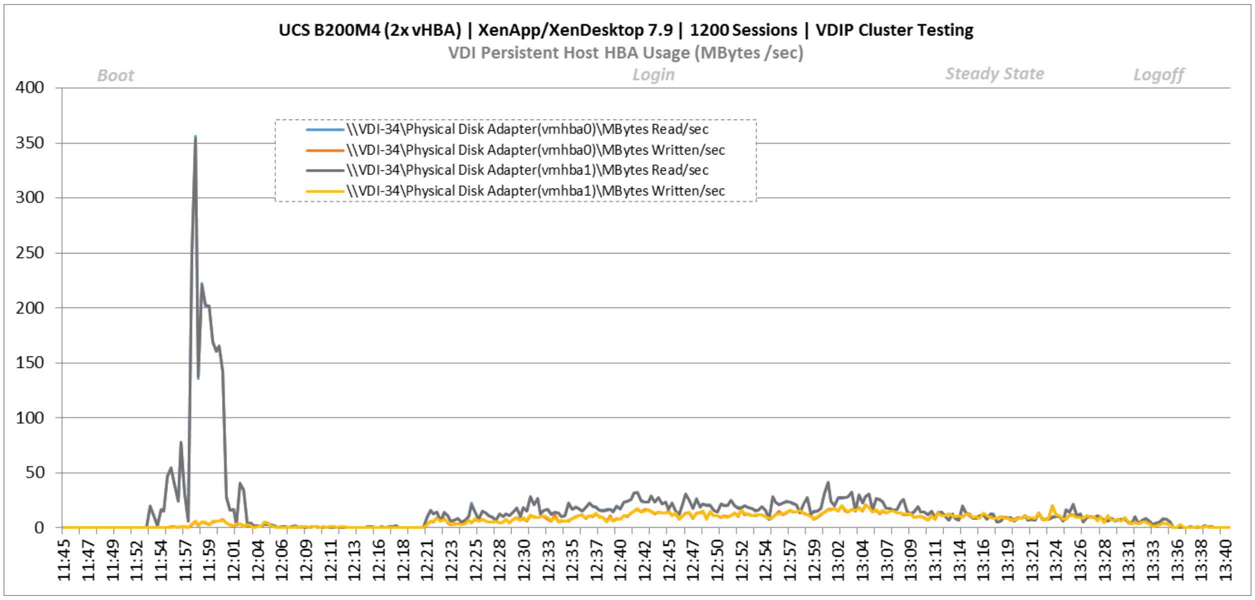

Cluster Workload Testing with 1200 Non-Persistent Desktop Users

This section shows the key performance metrics that were captured on the Cisco UCS, Pure Storage FlashArray//m, and Infrastructure VMs during the non-persistent desktop testing. The cluster testing with comprised of 1200 VDI Non-Persistent desktop sessions using 8 workload blades.

Figure 175 VDI Non-Persistent Cluster Testing with 1200 Users

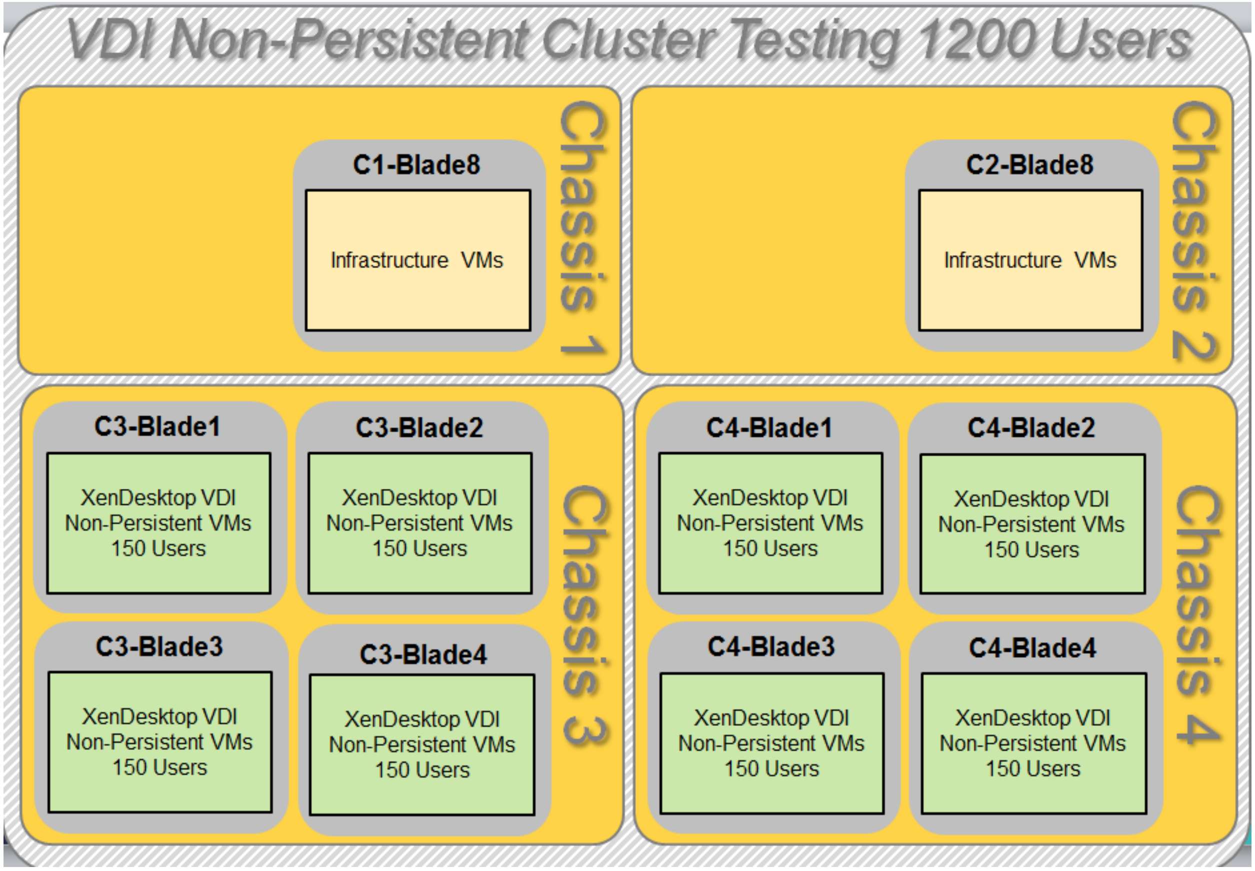
VDI Non-Persistent Cluster Testing 1200 Users
C1-Blade8
Infrastructure VMs
Chassis 1
C2-Blade8
Infrastructure VMs
Chassis 2
C3-Blade1
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade2
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade3
XenDesktop VDI Non-Persistent VMs 150 Users
C3-Blade4
XenDesktop VDI Non-Persistent VMs 150 Users
Chassis 3
C4-Blade1
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade2
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade3
XenDesktop VDI Non-Persistent VMs 150 Users
C4-Blade4
XenDesktop VDI Non-Persistent VMs 150 Users
Chassis 4

The workload for the test is 1200 non-persistent desktop users. To achieve the target, sessions were launched against the single non-persistent cluster only. 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 176 Cluster | 1200 VDI-NP Users | VSI Score

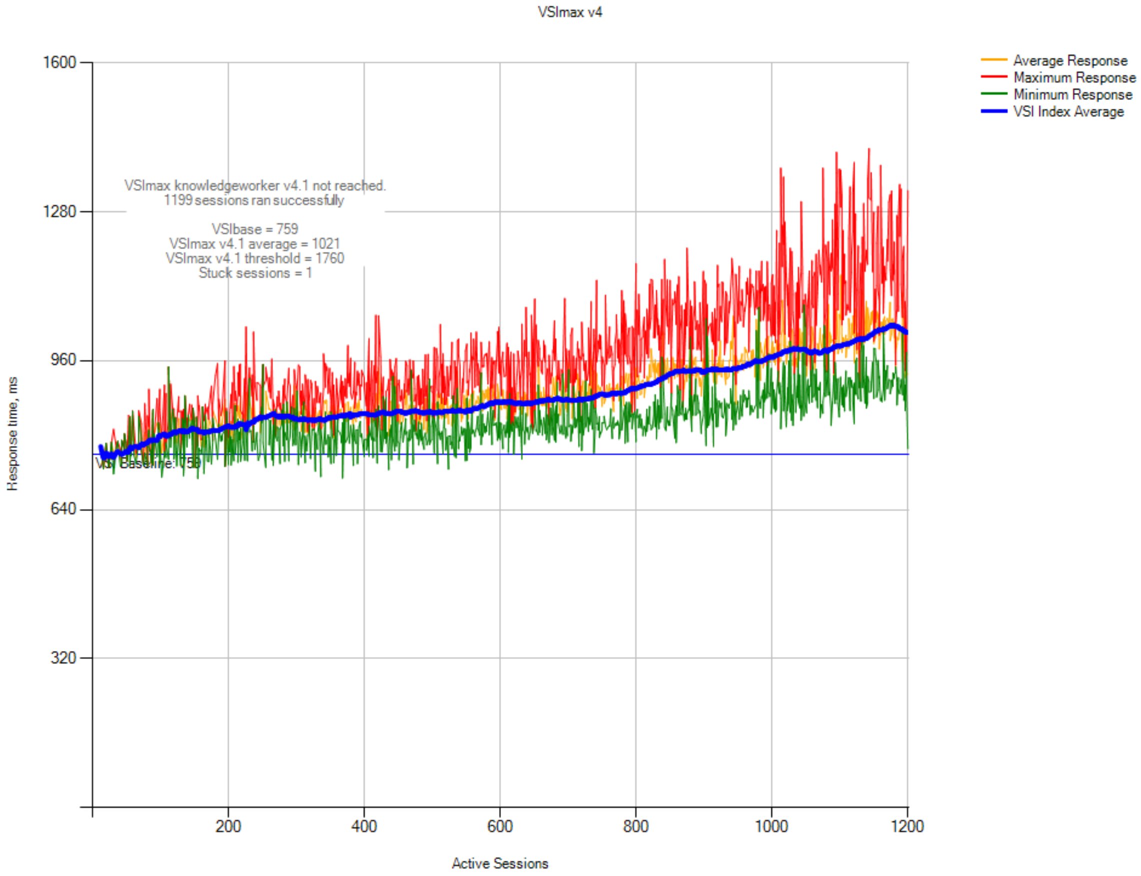

Figure 177 Cluster | 1200 VDI-NP Users | Infrastructure Hosts | Host CPU Utilization

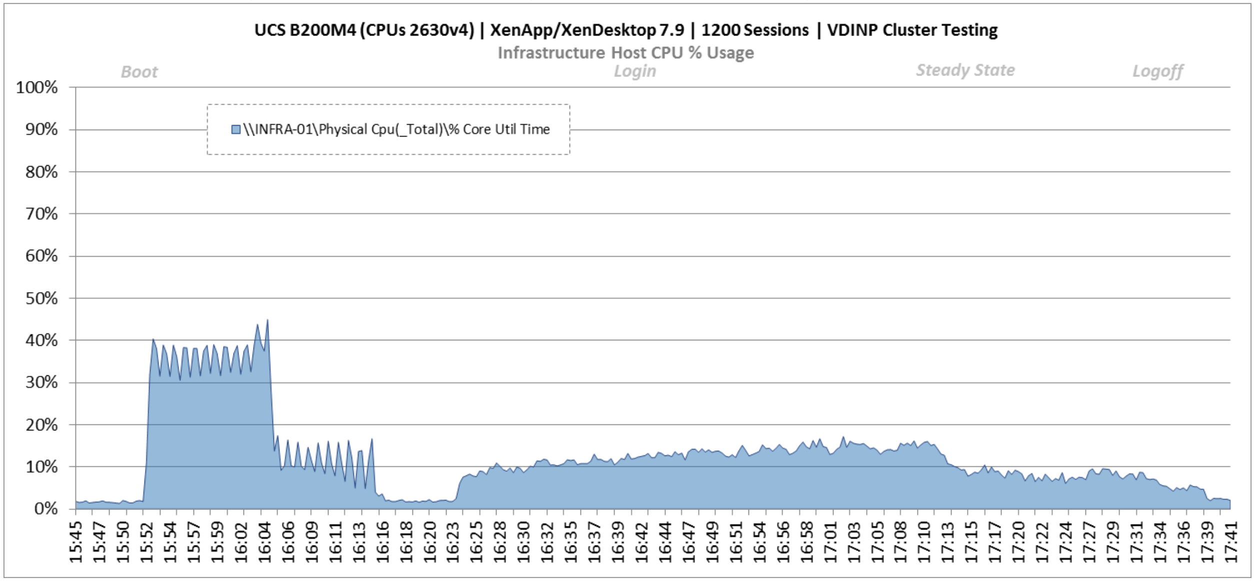

Figure 178 Cluster | 1200 VDI-NP Users | Infrastructure Hosts | Host Memory Utilization

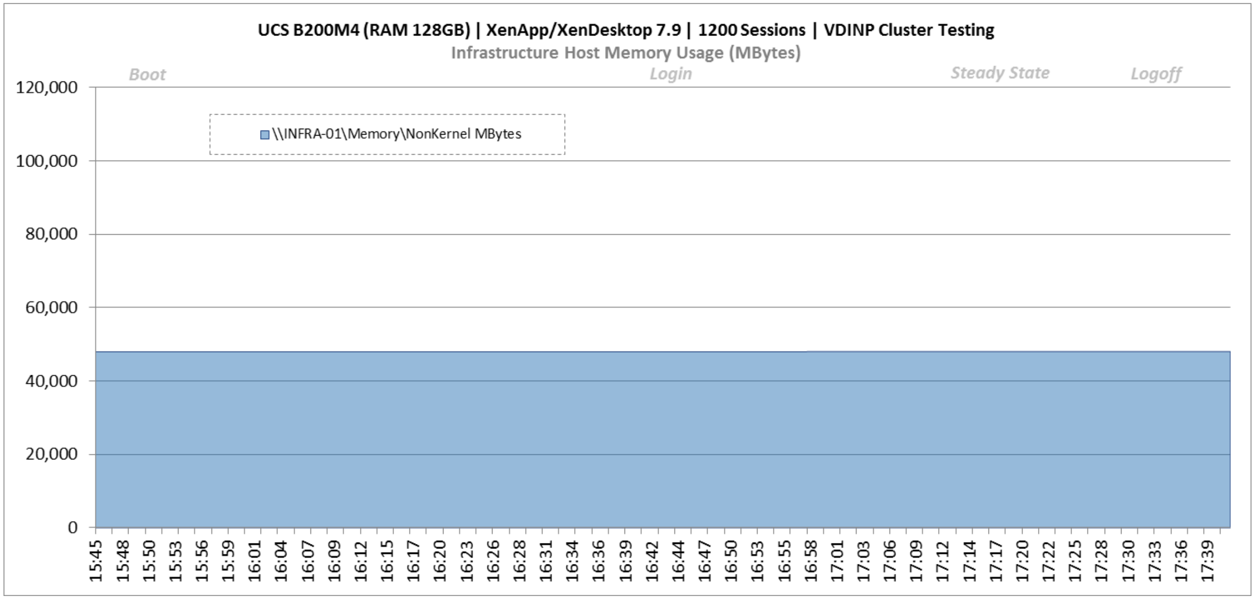

Figure 179 Cluster | 1200 VDI-NP Users | Infrastructure Hosts | Host Network Utilization

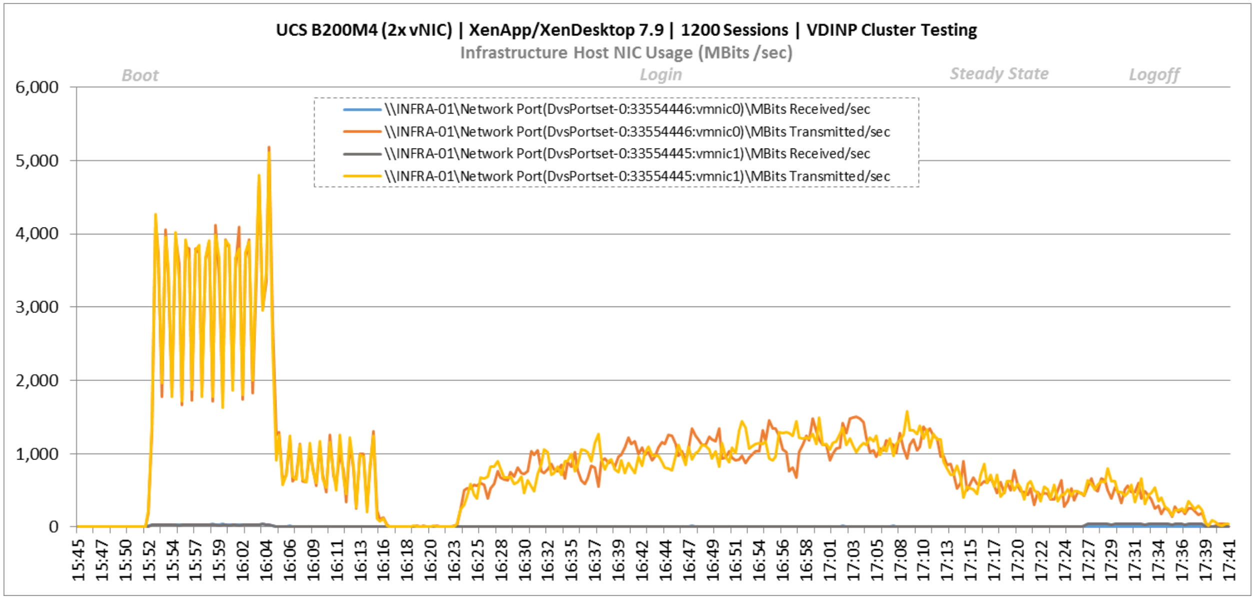

Figure 180 Cluster | 1200 VDI-NP Users | Infrastructure Hosts | Host Fibre Channel Network Utilization

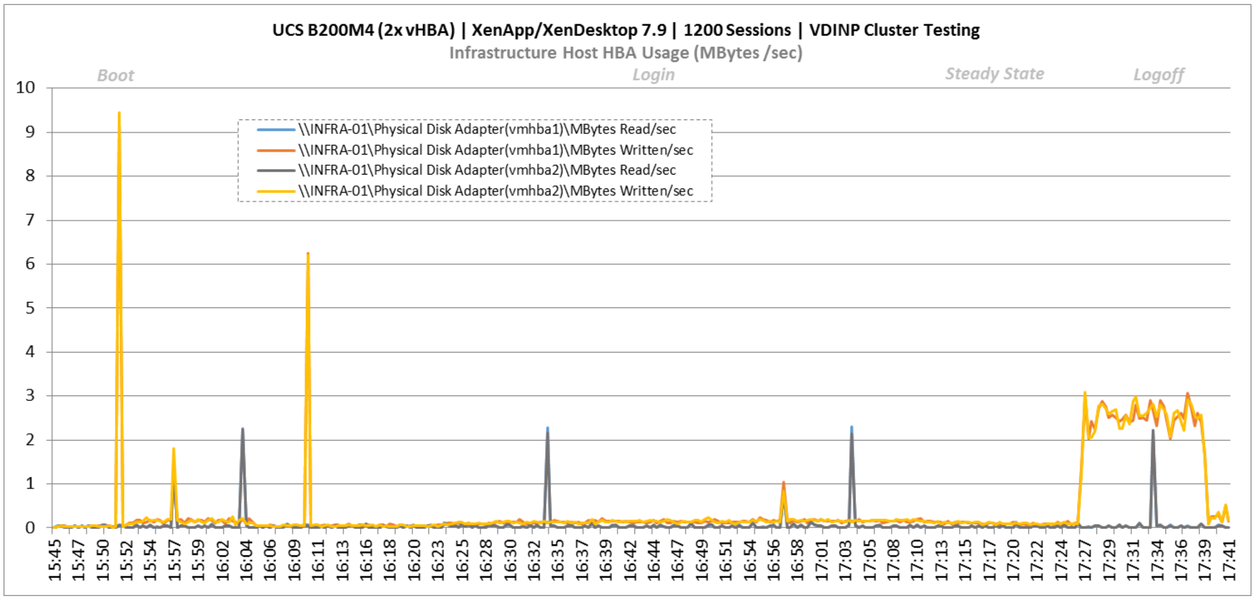

Figure 181 Cluster | 1200 VDI-NP Users | Persistent Hosts | Sample Host CPU Utilization

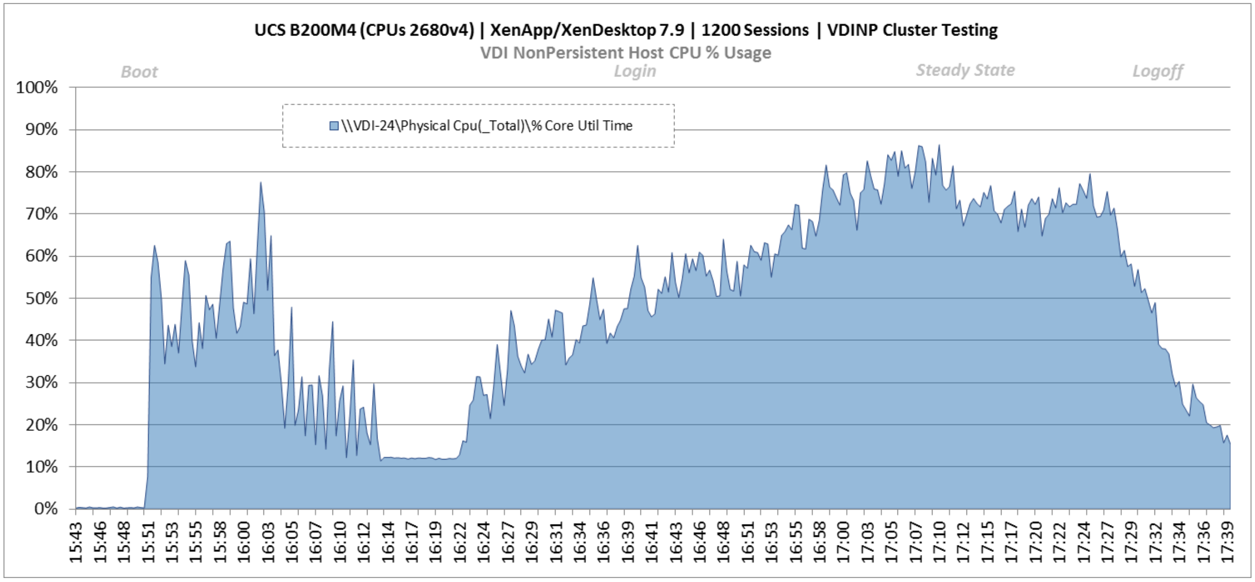

Figure 182 Cluster | 1200 VDI-NP Users | Persistent Hosts | Sample Host Memory Utilization

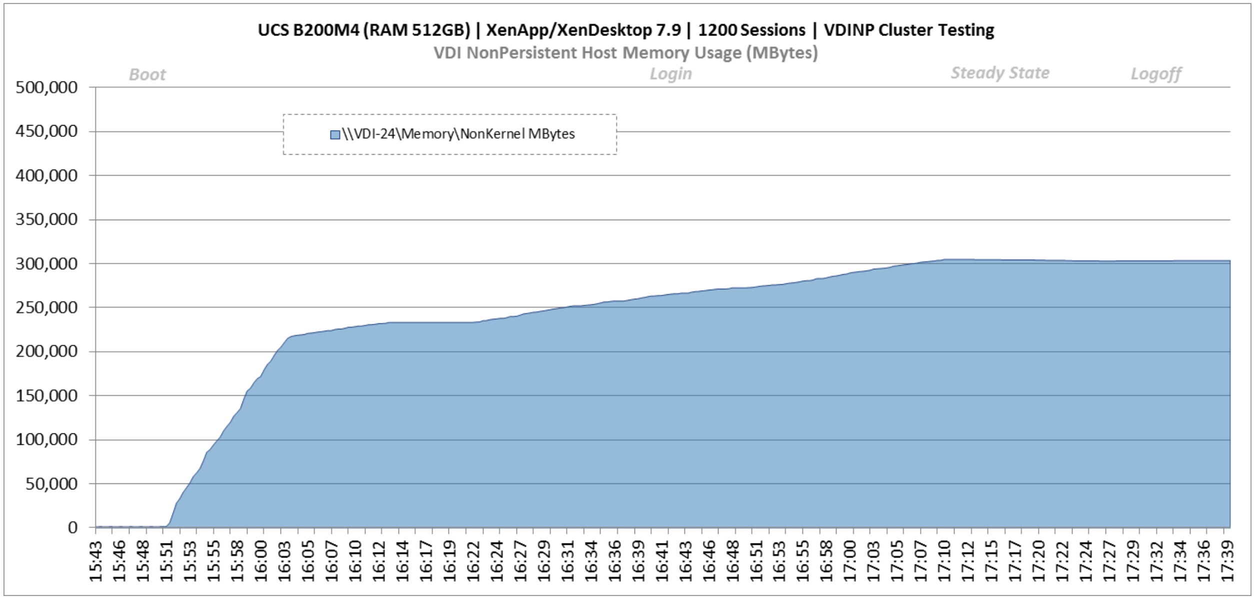

Figure 183 Cluster | 1200 VDI-NP Users | Persistent Hosts | Sample Host Network Utilization

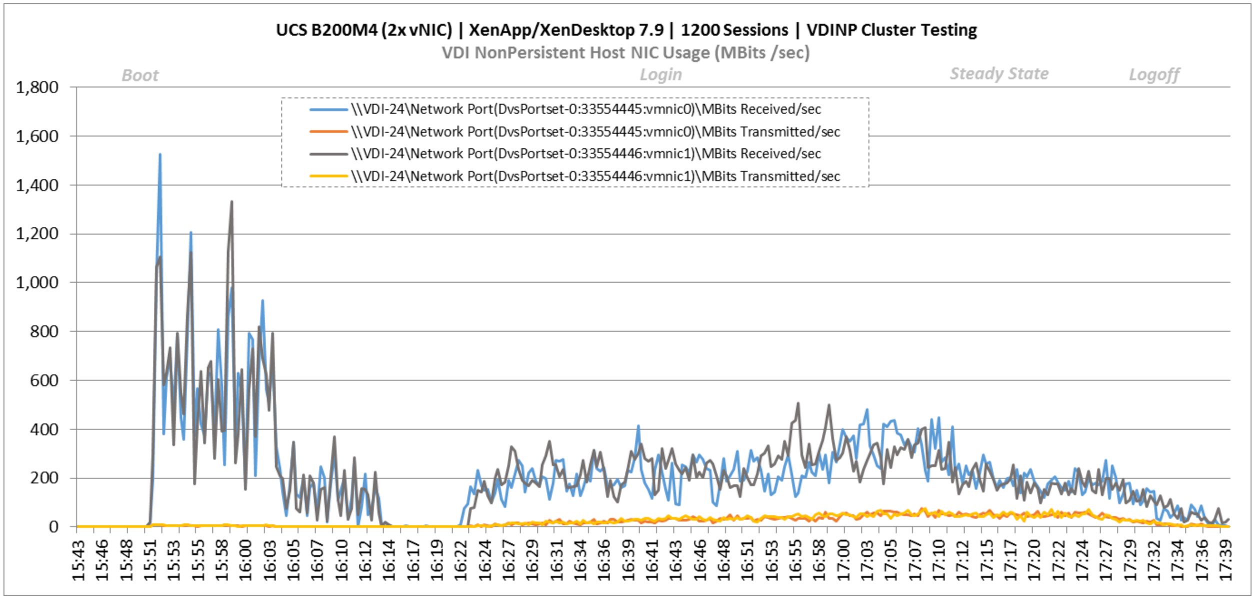

Figure 184 Cluster | 1200 VDI-NP Users | Persistent Hosts | Sample Host Fibre Channel Network Utilization

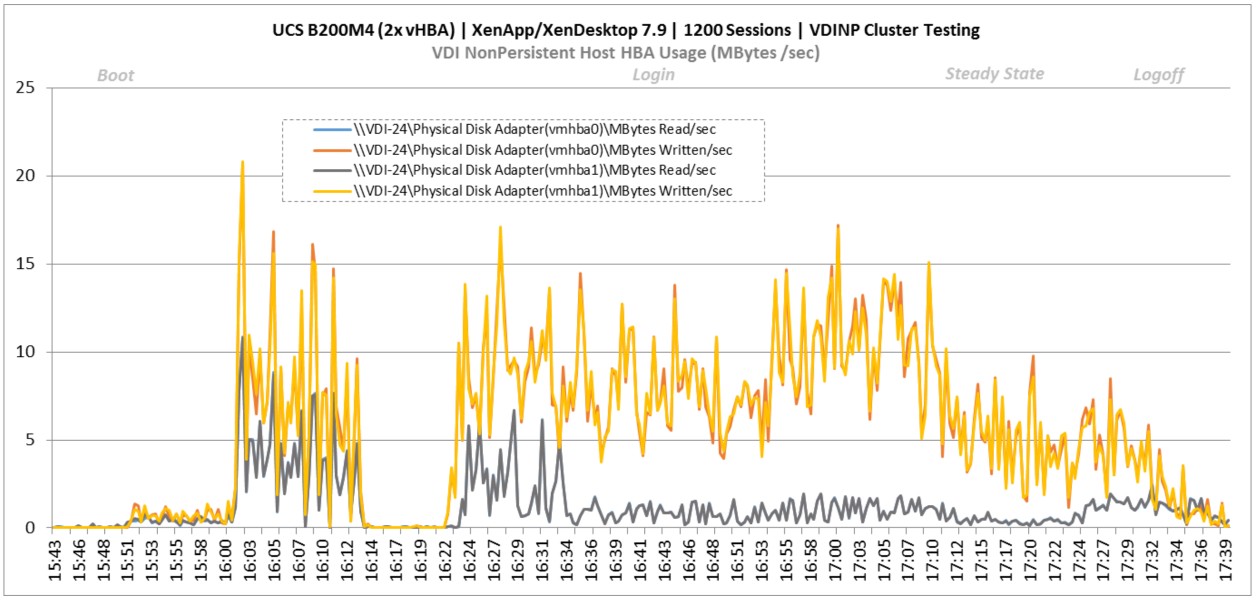

Full Scale Mixed Workload Testing with 5000 Users

This section shows the key performance metrics that were captured on the Cisco UCS, Pure Storage FlashArray//m, and Infrastructure VMs during the full-scale testing. The full-scale testing with 5000 users comprised of: 2600 RDS sessions using 10 blades, 1200 VDI Non-Persistent sessions using 8 blades, and 1200 VDI Persistent sessions using 8 blades.

Figure 185 Full Scale Mixed Test with 5000 Users

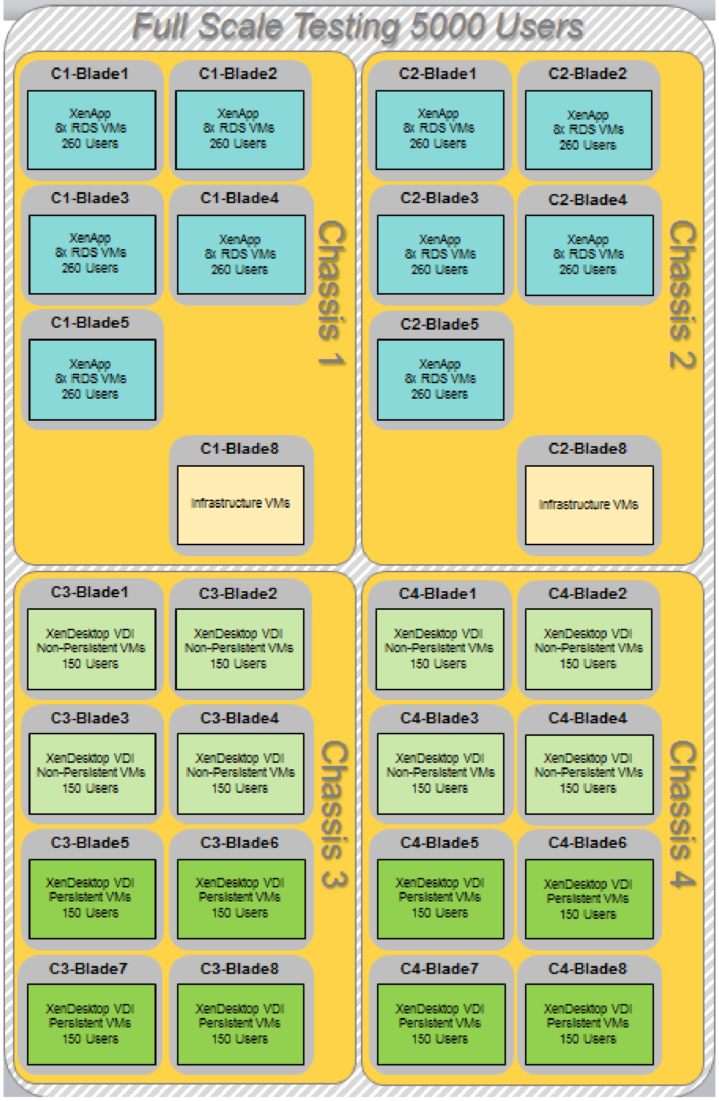

The combined mixed workload for the solution is 5000 users. To achieve the target, sessions were launched against all workload clusters concurrently. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark 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 186 Full Scale | 5000 Mixed Users | VSI Score

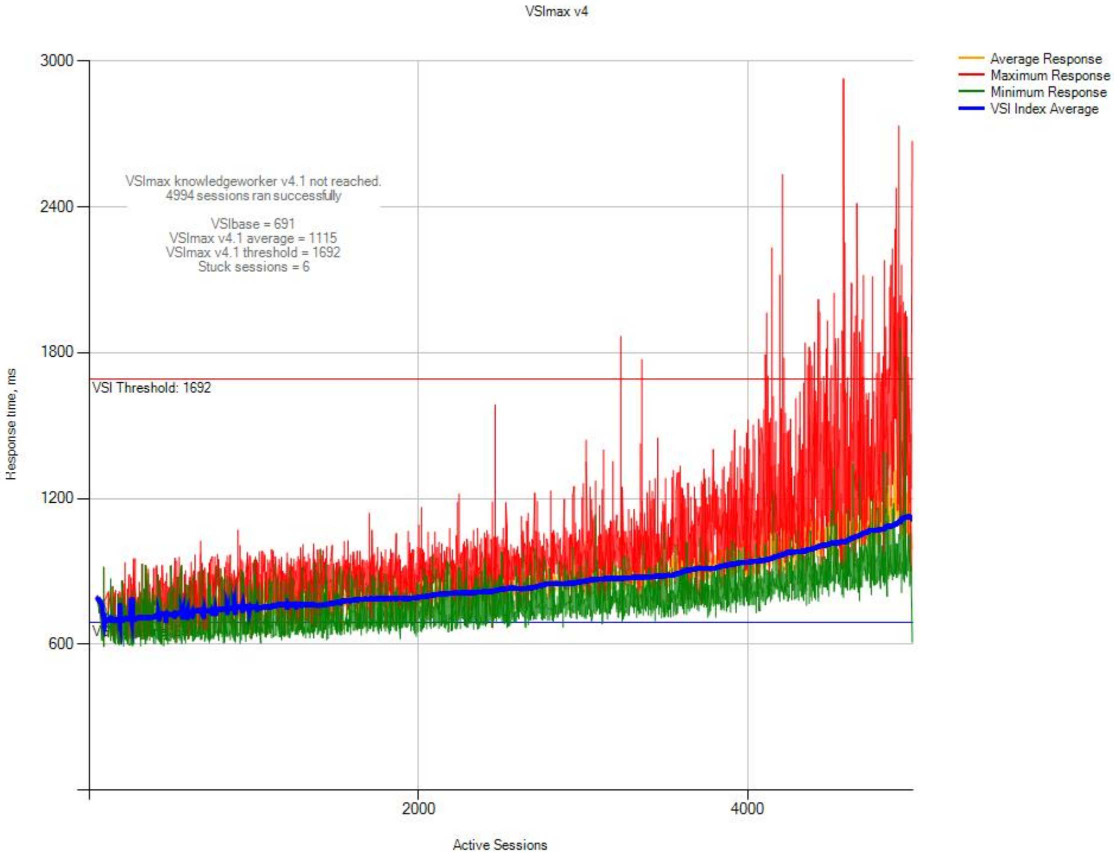

Figure 187 Full Scale | 5000 Mixed Users | Citrix Director | Logon Stats

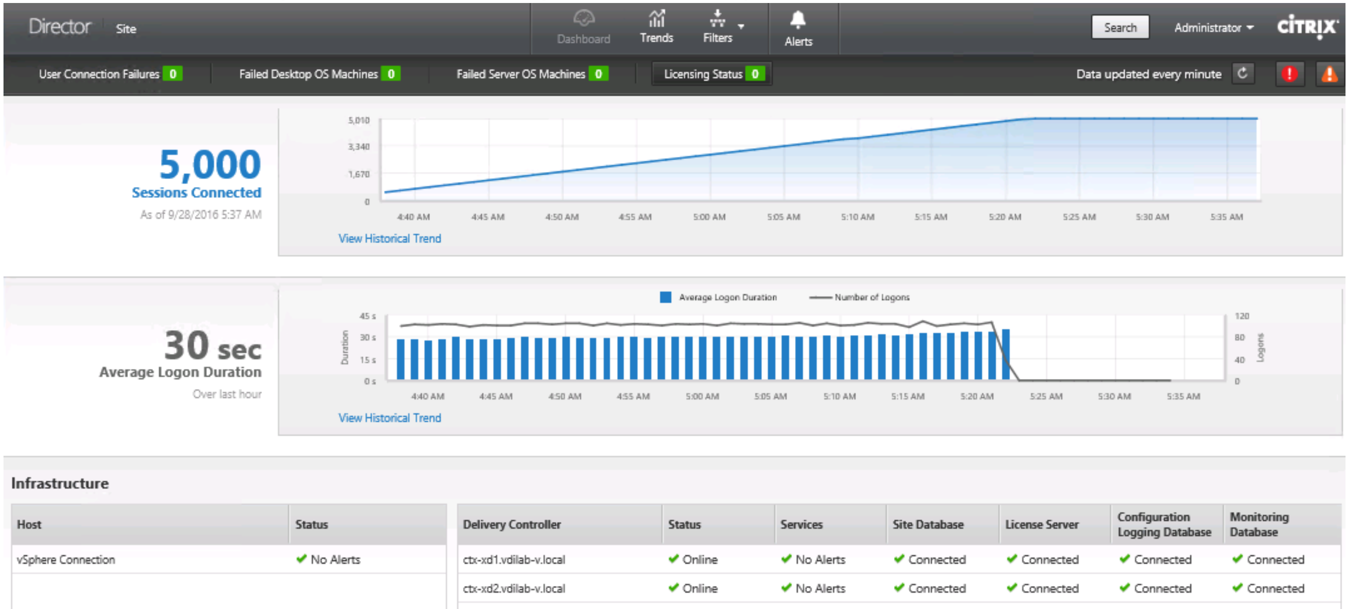

5,000
Sessions Connected
As of 9/28/2016 5:37 AM

30 sec
Average Logon Duration
Over last hour

Infrastructure

Host	Status
vSphere Connection	✓ No Alerts

Delivery Controller	Status	Services	Site Database	License Server	Configuration Logging Database	Monitoring Database
ctx-xd1.vdilab-v.local	✓ Online	✓ No Alerts	✓ Connected	✓ Connected	✓ Connected	✓ Connected
ctx-xd2.vdilab-v.local	✓ Online	✓ No Alerts	✓ Connected	✓ Connected	✓ Connected	✓ Connected

Figure 188 Full Scale | 5000 Mixed Users | Citrix PVS | Balanced Load

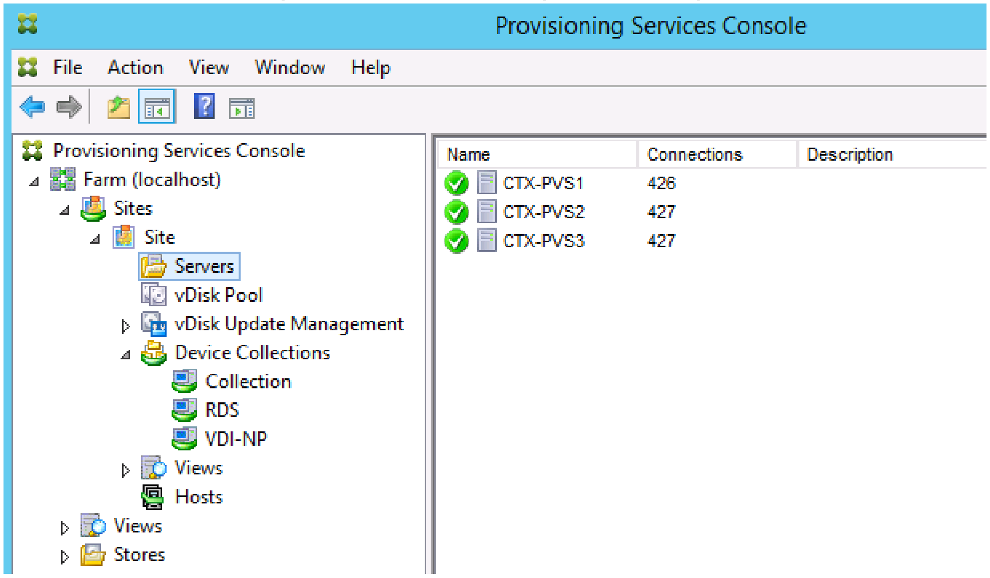

Name	Connections	Description
CTX-PVS1	426	
CTX-PVS2	427	
CTX-PVS3	427	

Figure 189 Full Scale | 5000 Mixed Users | Citrix Studio | 5000 Mixed Session Active

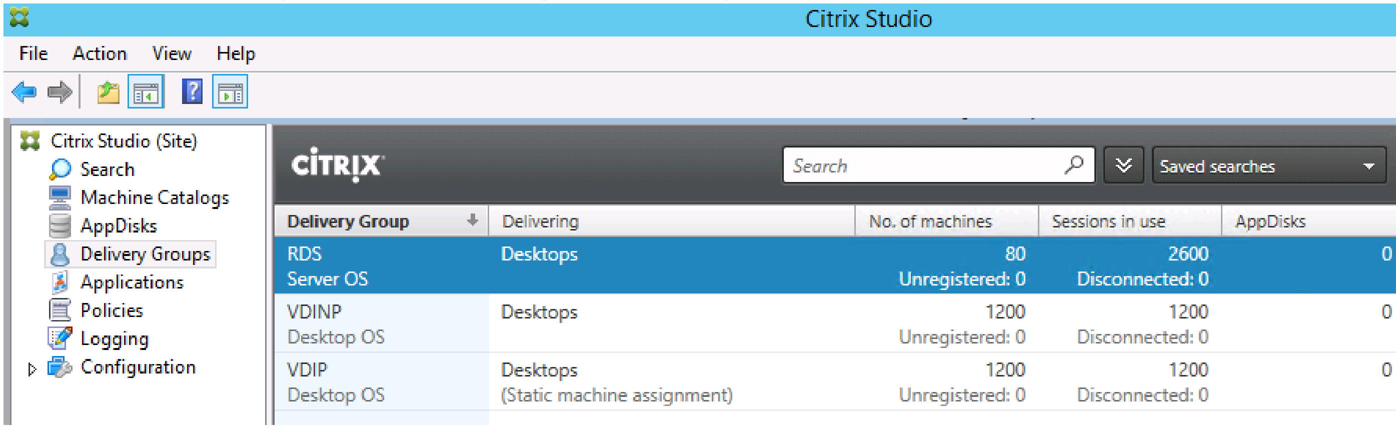

Delivery Group	Delivering	No. of machines	Sessions in use	AppDisks
RDS Server OS	Desktops	80 Unregistered: 0	2600 Disconnected: 0	0
VDINP Desktop OS	Desktops	1200 Unregistered: 0	1200 Disconnected: 0	0
VDIP Desktop OS	Desktops (Static machine assignment)	1200 Unregistered: 0	1200 Disconnected: 0	0

Figure 190 Full Scale | 5000 Mixed Users | Infrastructure Hosts | Host CPU Utilization

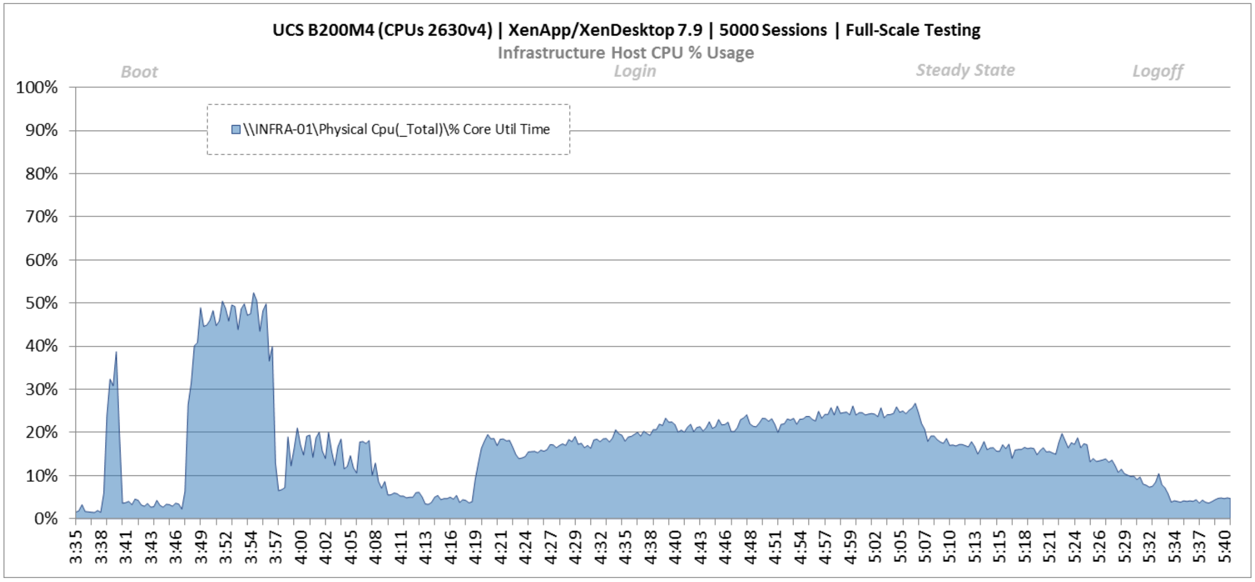

Figure 191 Full Scale | 5000 Mixed Users | Infrastructure Hosts | Host Memory Utilization

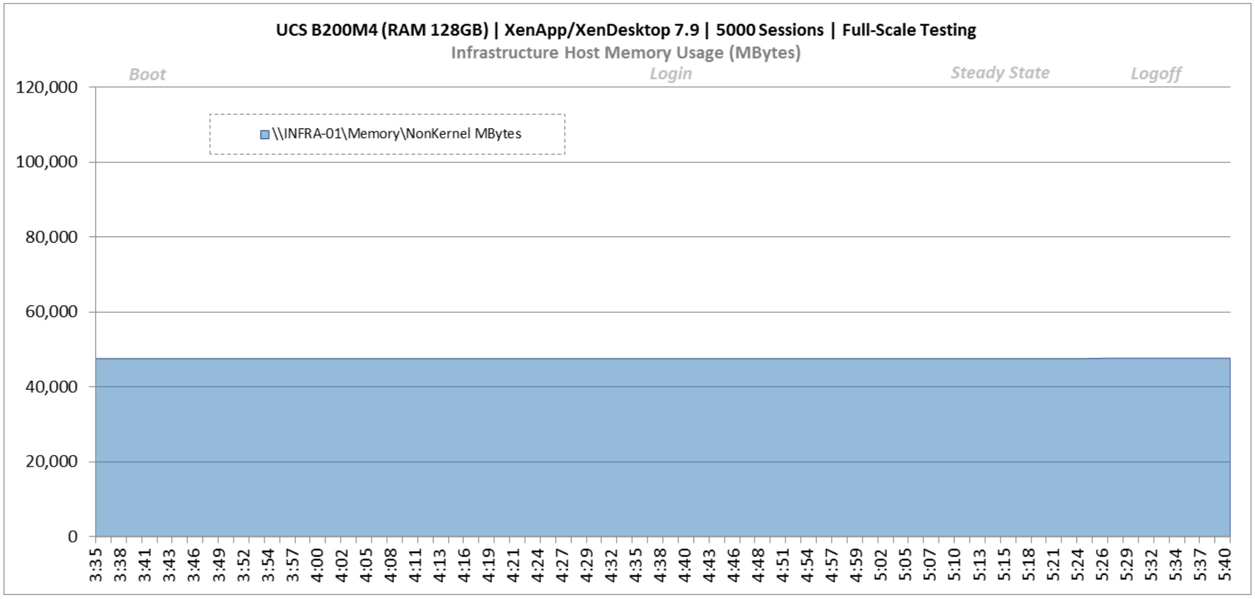

Figure 192 Full Scale | 5000 Mixed Users | Infrastructure Hosts | Host Network Utilization

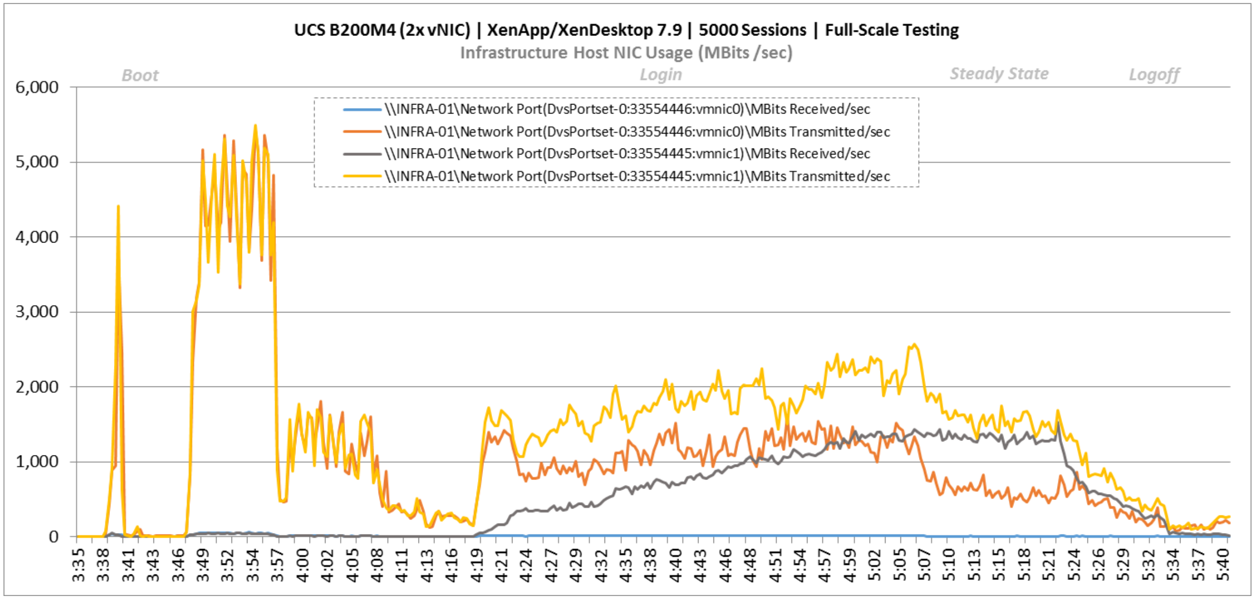

Figure 193 Full Scale | 5000 Mixed Users | Infrastructure Hosts | Host Fibre Channel Network Utilization

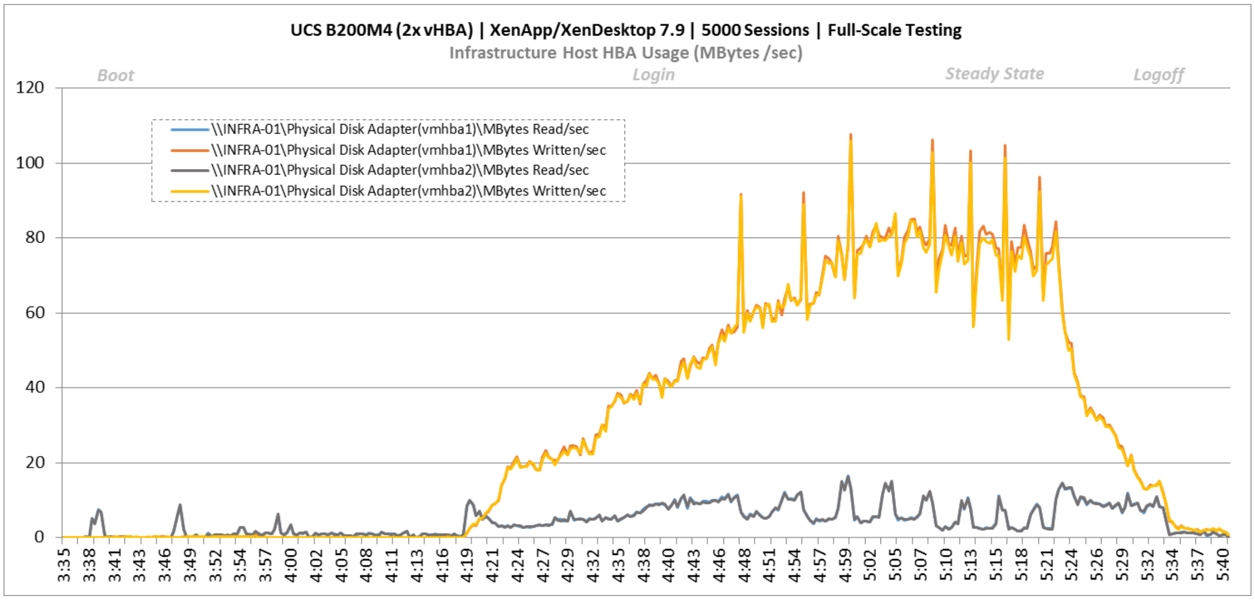

Figure 194 Full Scale | 5000 Mixed Users | RDS Hosts | Sample Host CPU Utilization

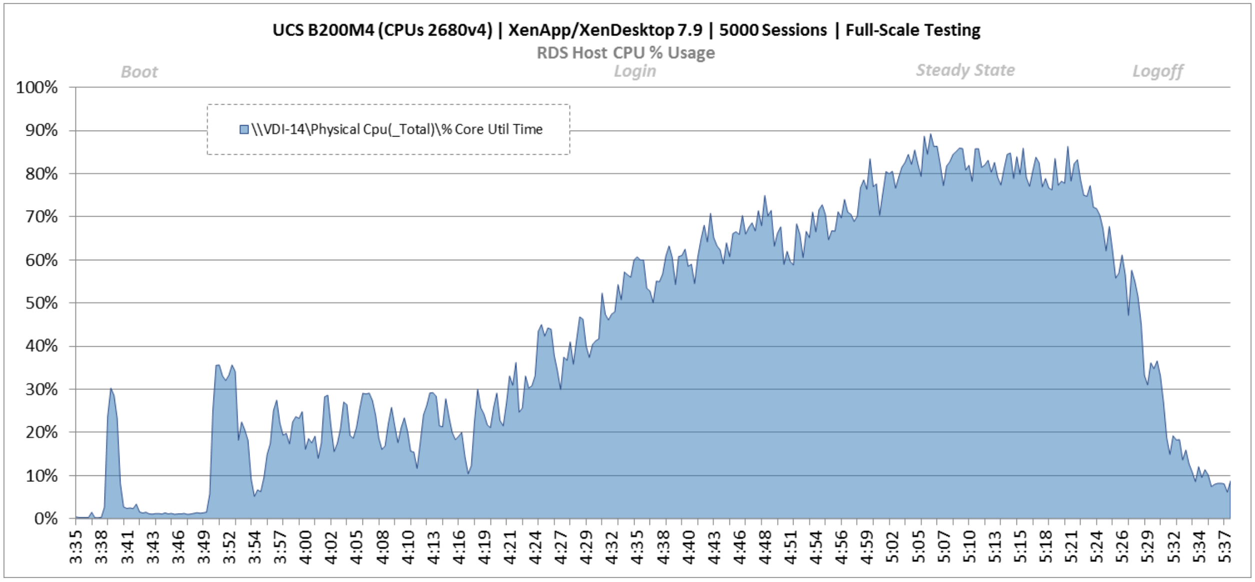

Figure 195 Full Scale | 5000 Mixed Users | RDS Hosts | Sample Host Memory Utilization

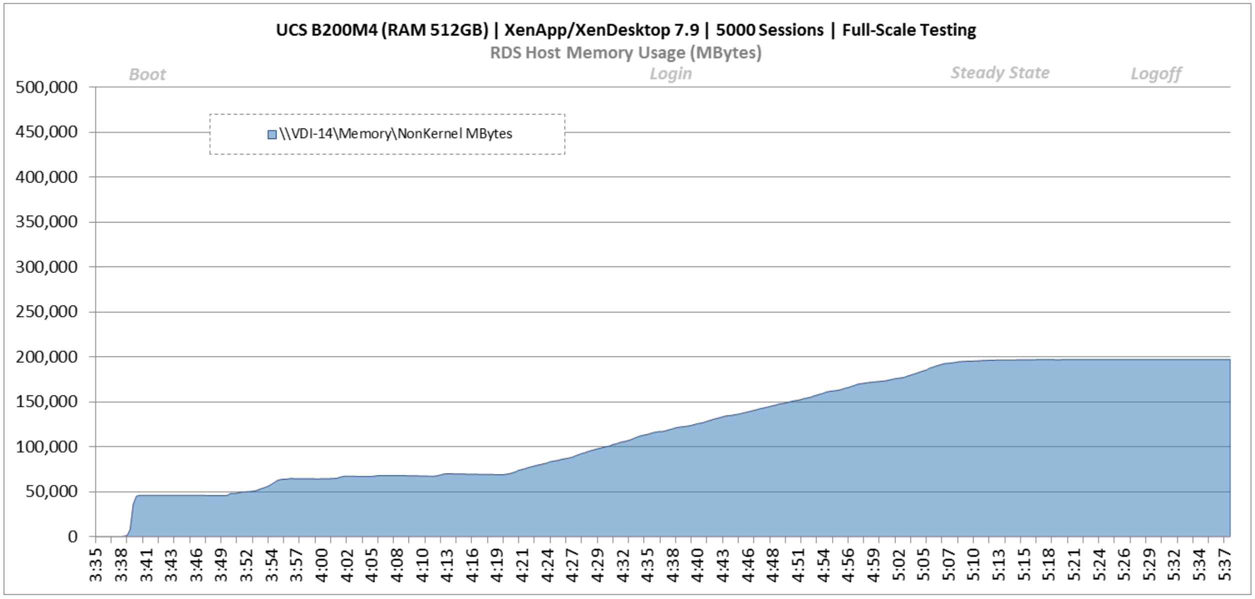

Figure 196 Full Scale | 5000 Mixed Users | RDS Hosts | Sample Host System Network Utilization

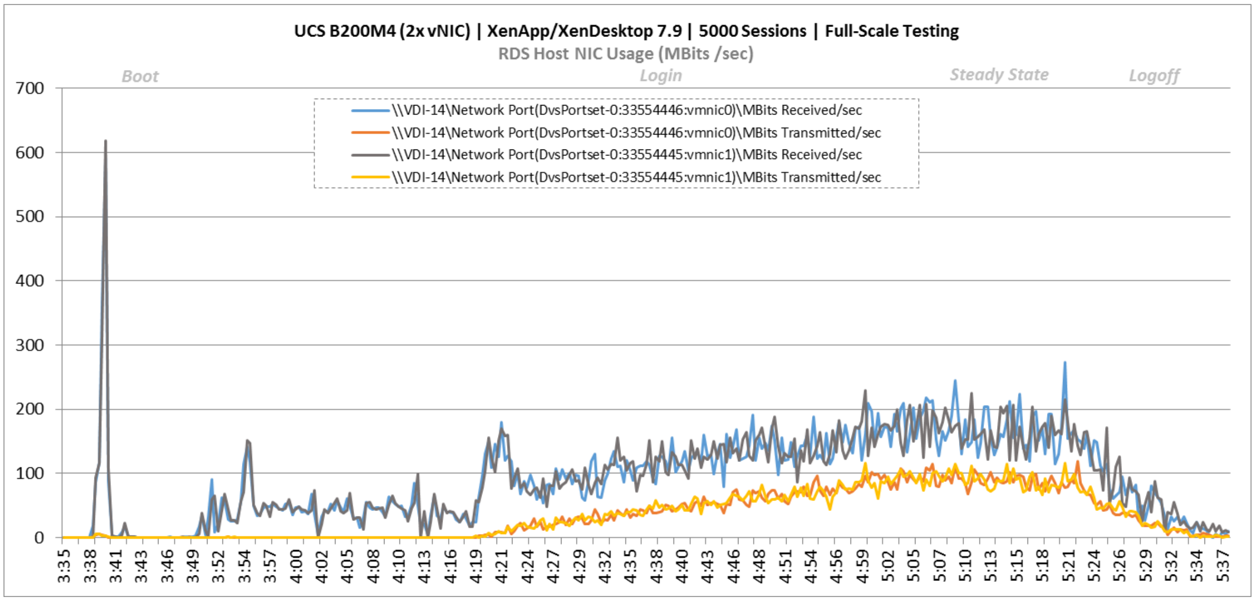

Figure 197 Full Scale | 5000 Mixed Users | RDS Hosts | Sample Host Fibre Channel Network Utilization

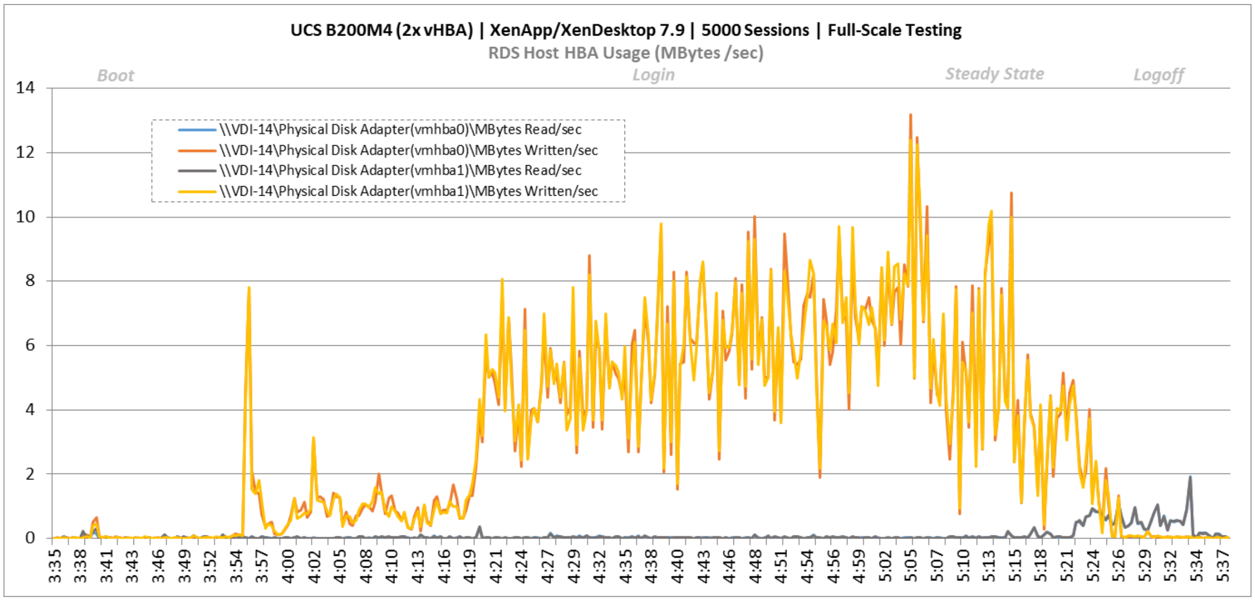

Figure 198 Full Scale | 5000 Mixed Users | Non-Persistent Hosts | Sample Host CPU Utilization

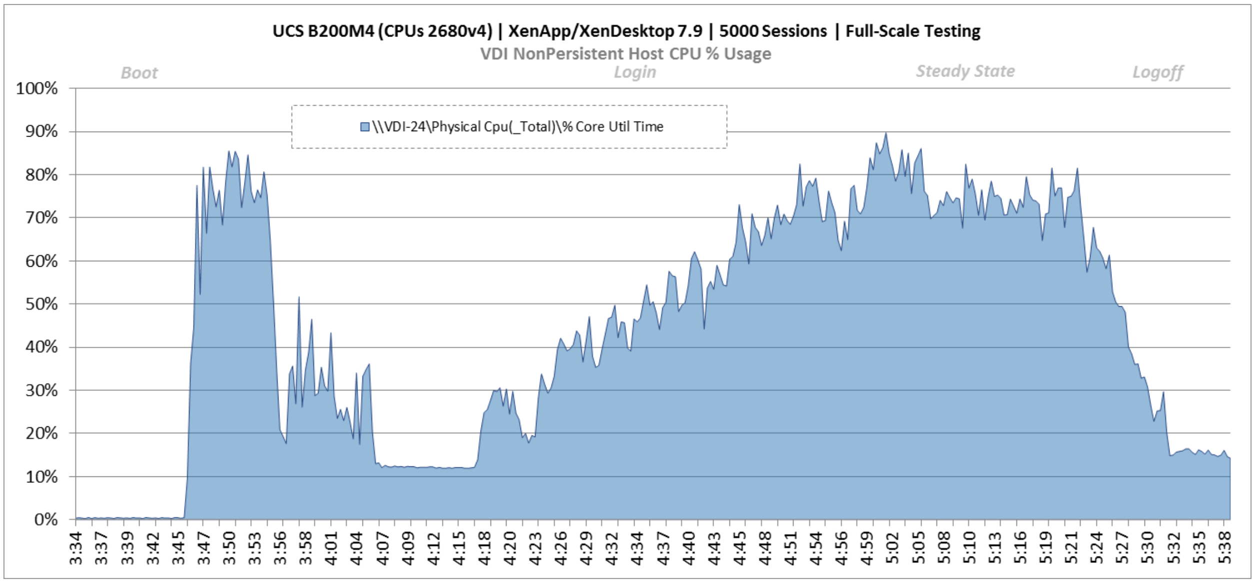

Figure 199 Full Scale | 5000 Mixed Users | Non-Persistent Hosts | Sample Host Memory Utilization

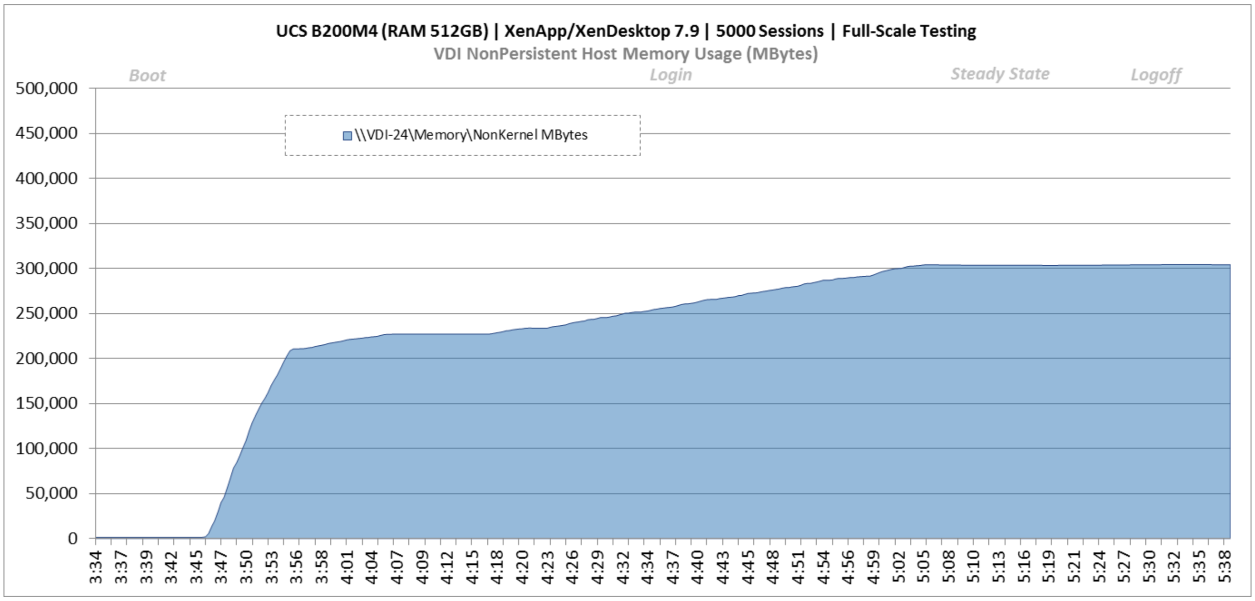

Figure 200 Full Scale | 5000 Mixed Users | Non-Persistent Hosts | Sample Host Network Utilization

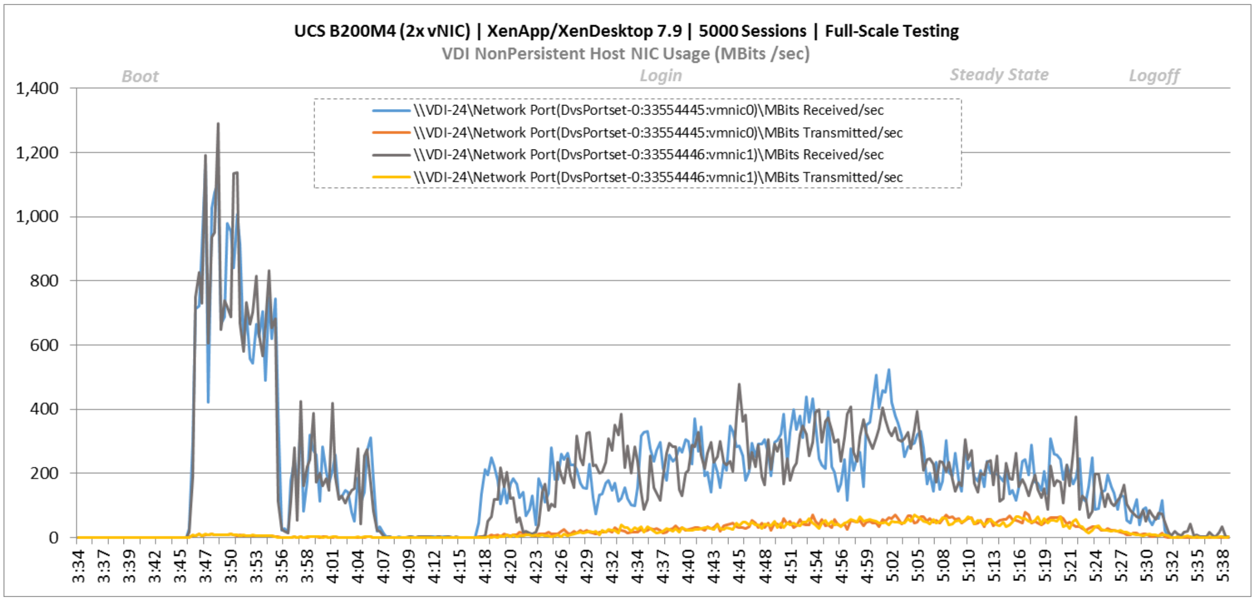

Figure 201 Full Scale | 5000 Mixed Users | Non-Persistent Hosts | Sample Host Fibre Channel Network Utilization

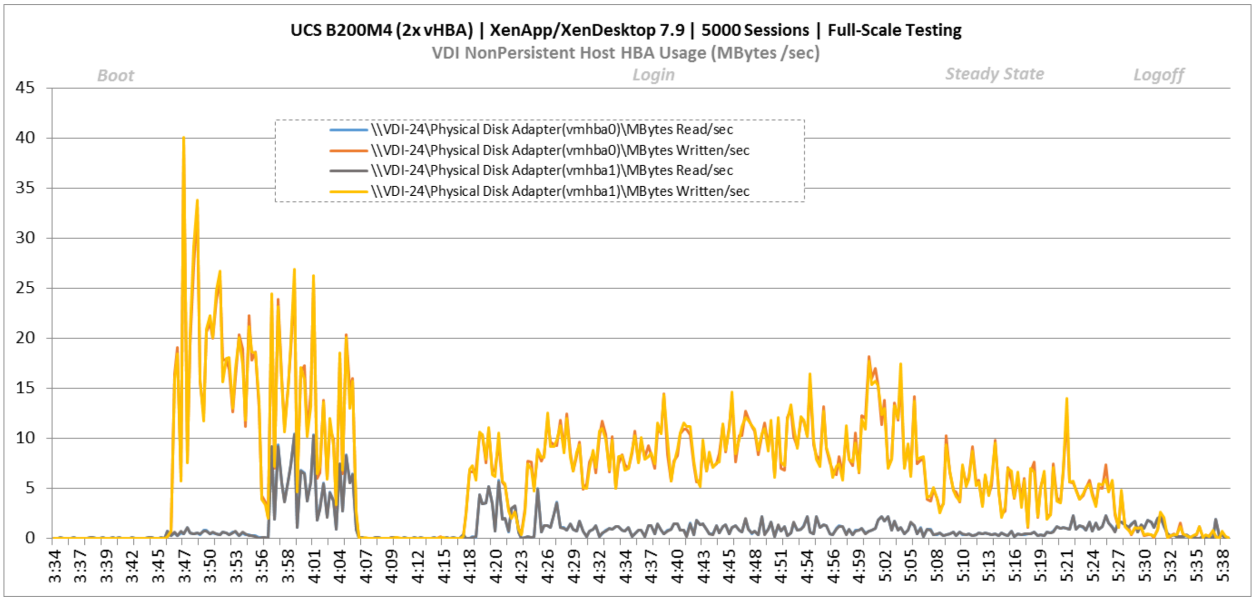

Figure 202 Full Scale | 5000 Mixed Users | Persistent Hosts | Sample Host CPU Utilization

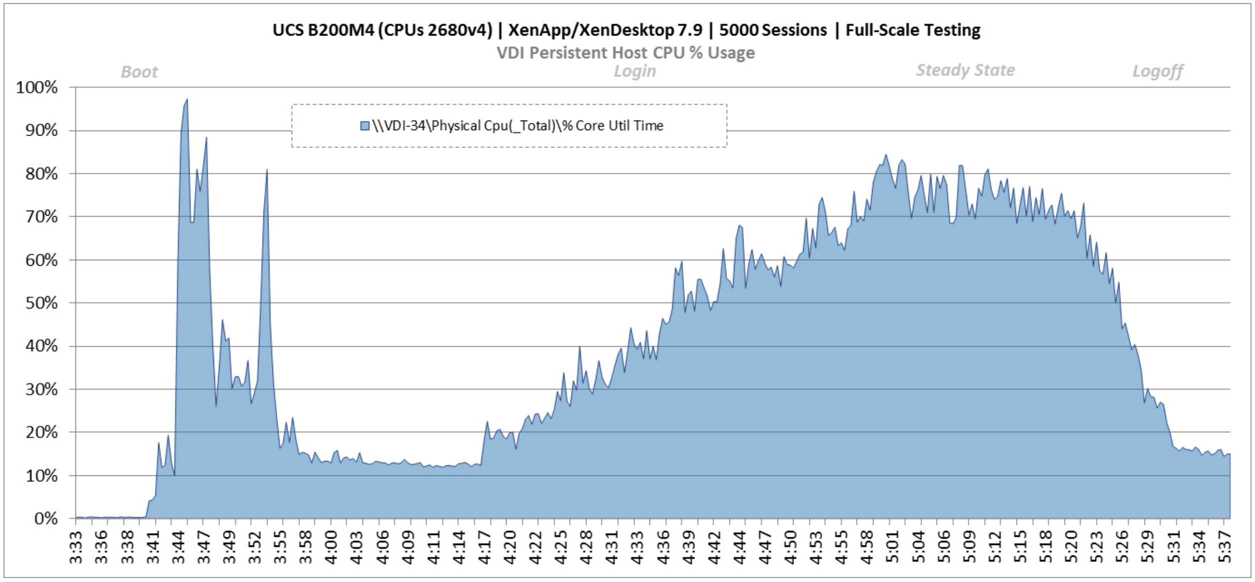

Figure 203 Full Scale | 5000 Mixed Users | Persistent Hosts | Sample Host Memory Utilization

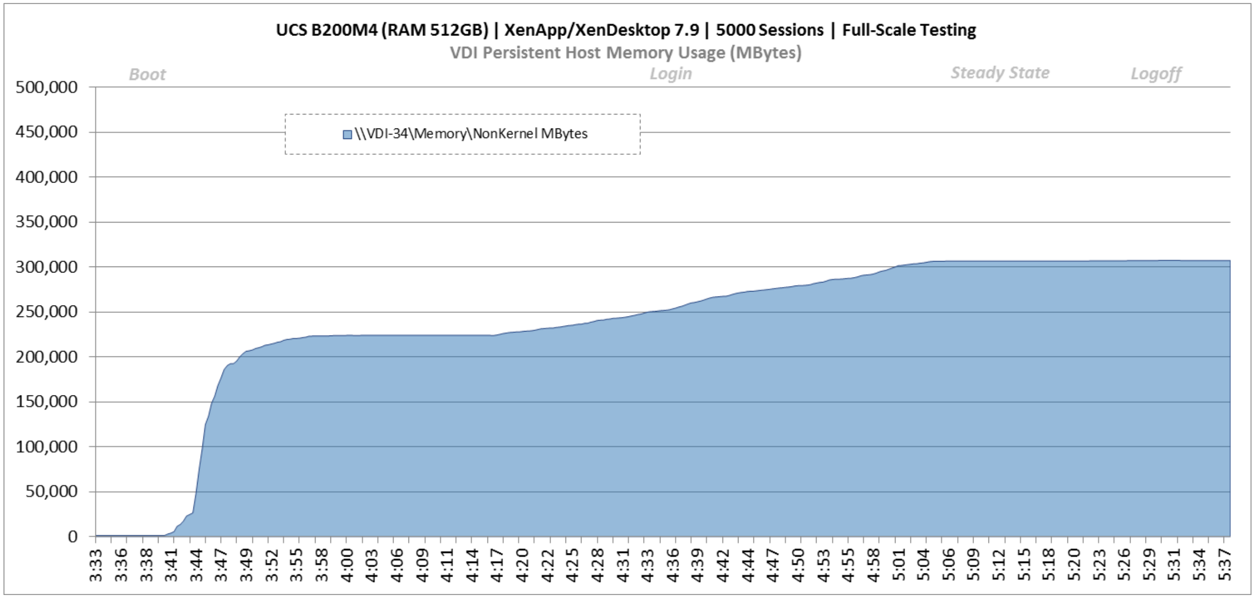

Figure 204 Full Scale | 5000 Mixed Users | Persistent Hosts | Sample Host Network Utilization

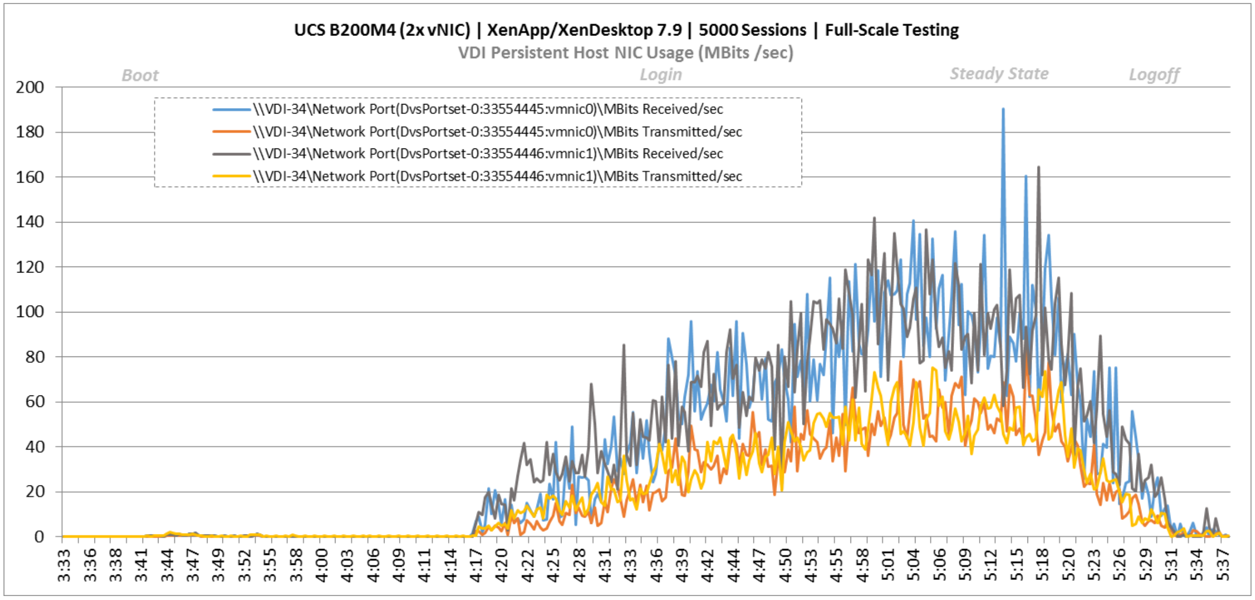

Figure 205 Full Scale | 5000 Mixed Users | Persistent Hosts | Sample Host Fibre Channel Network Utilization

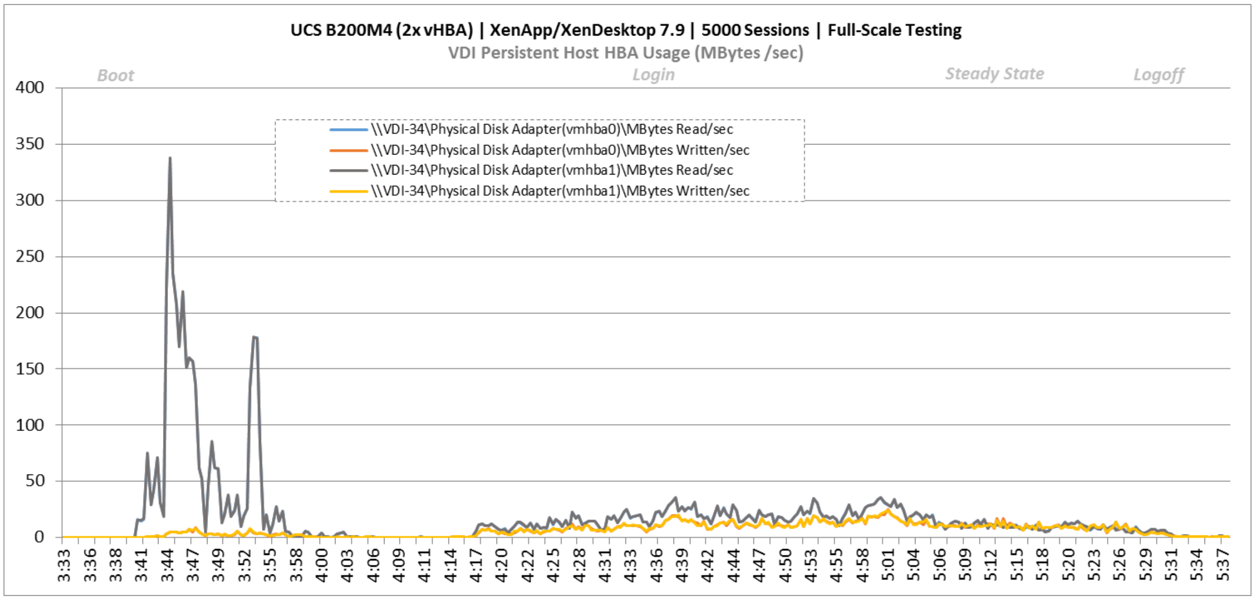

Note: Performance data for all 26 workload hosts are provided in Appendix D.

Figure 206 Pure Storage FlashArray//m50 Read/Write Latency

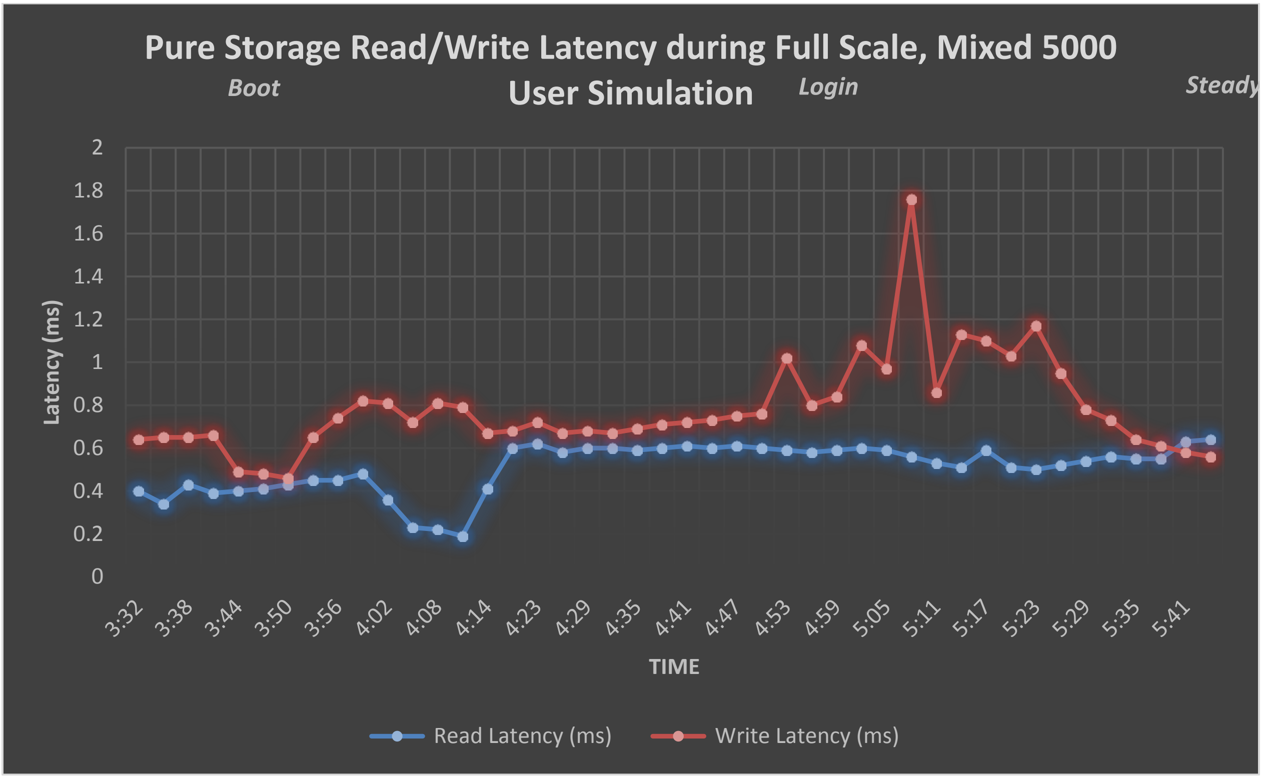

Figure 207 Pure Storage FlashArray//m50 Read/Write Bandwidth

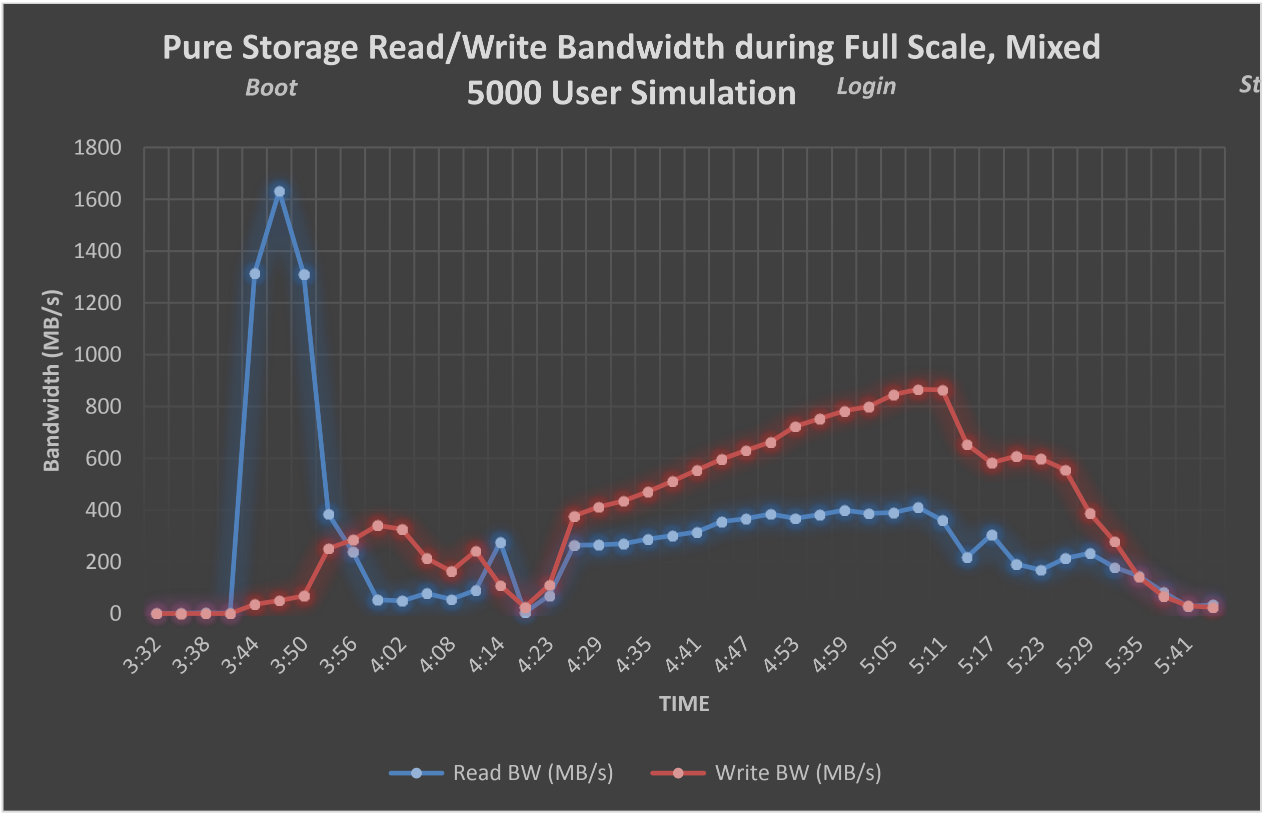

Figure 208 Pure Storage FlashArray//m50 Read/Write IOPS

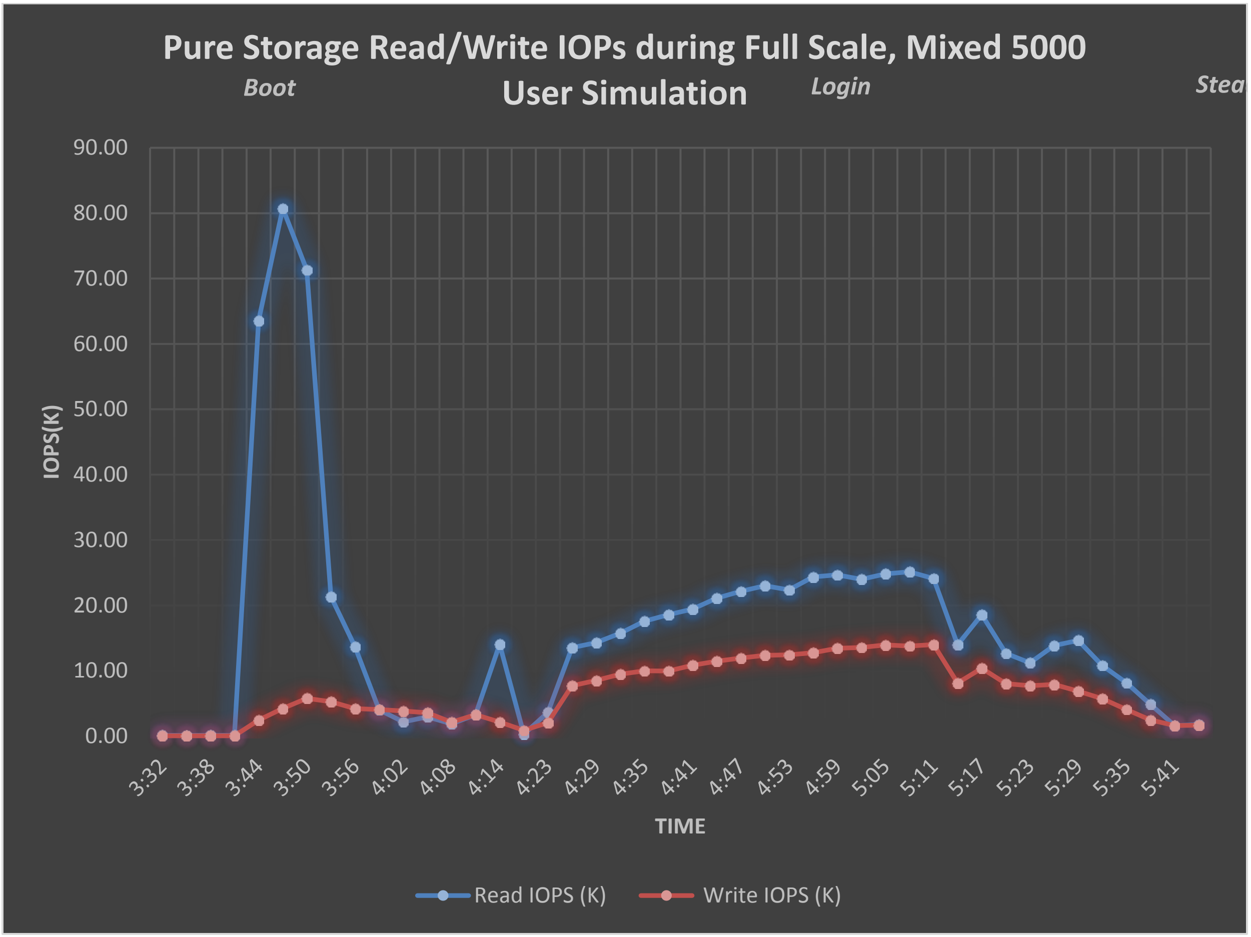

Key Infrastructure VM Server Performance Metrics during Full Scale Testing

It is important to verify that the key infrastructure servers are performing optimally during the scale test run. The following performance parameters were collected and charted; they validate that the designed infrastructure supports the mixed workload.

Figure 209 Full Scale | 5000 Mixed Users | Active Directory Domain Controllers | CPU Utilization

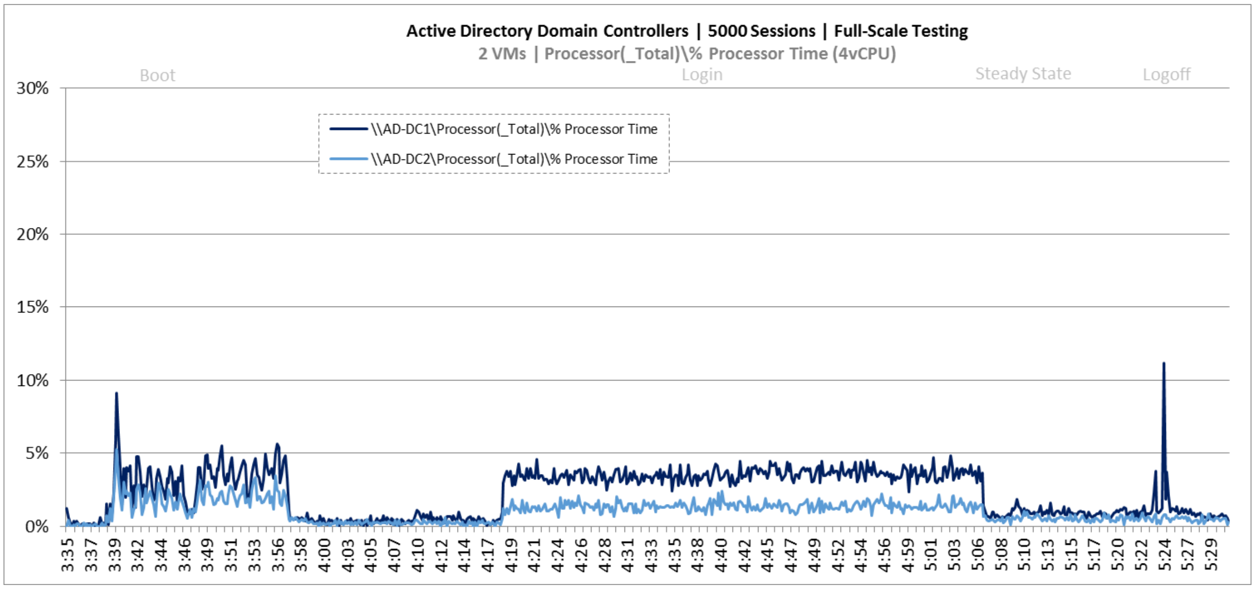

Figure 210 Full Scale | 5000 Mixed Users | Active Directory Domain Controllers | Memory Utilization

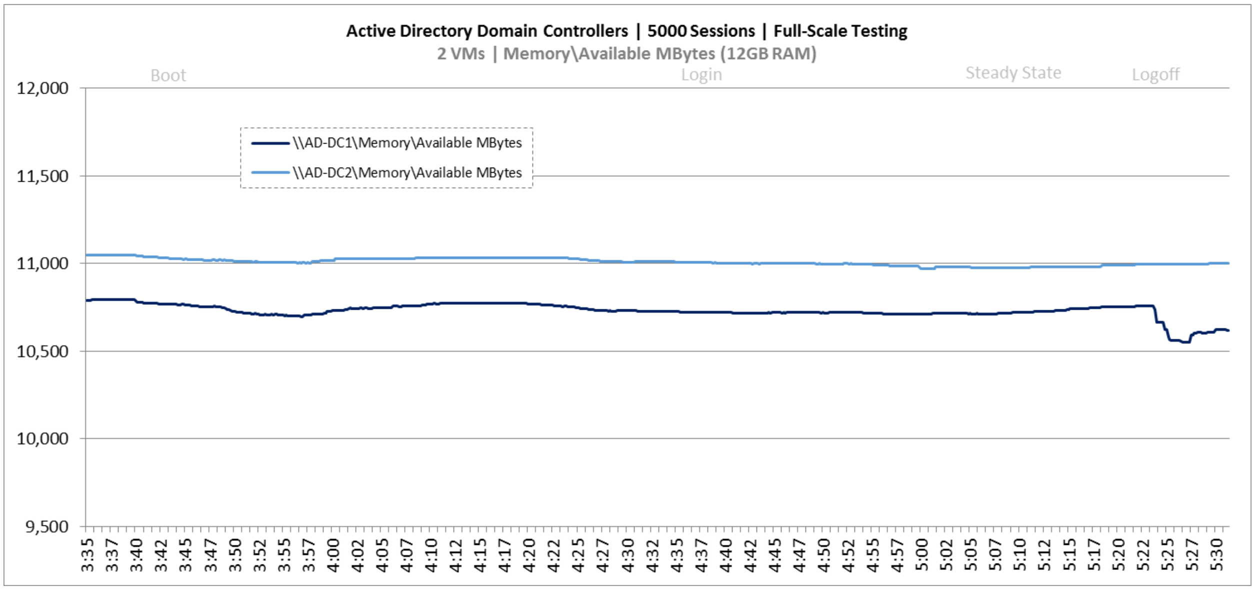

Figure 211 Full Scale | 5000 Mixed Users | Active Directory Domain Controllers | Network Utilization

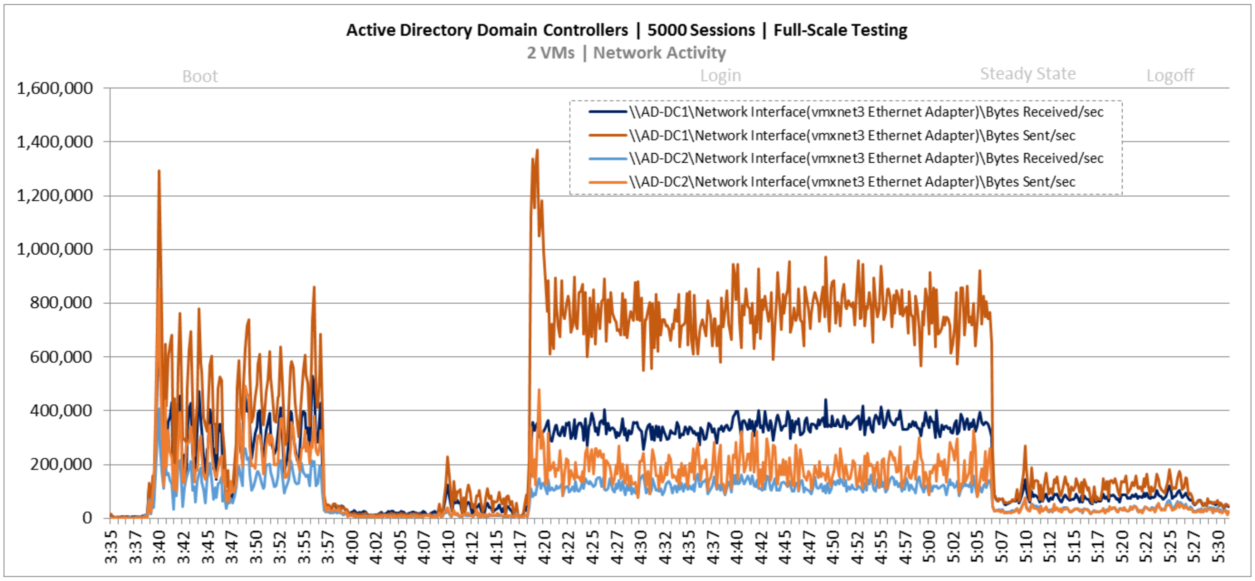

Figure 212 Full Scale | 5000 Mixed Users | Active Directory Domain Controllers | Disk Queue Lengths

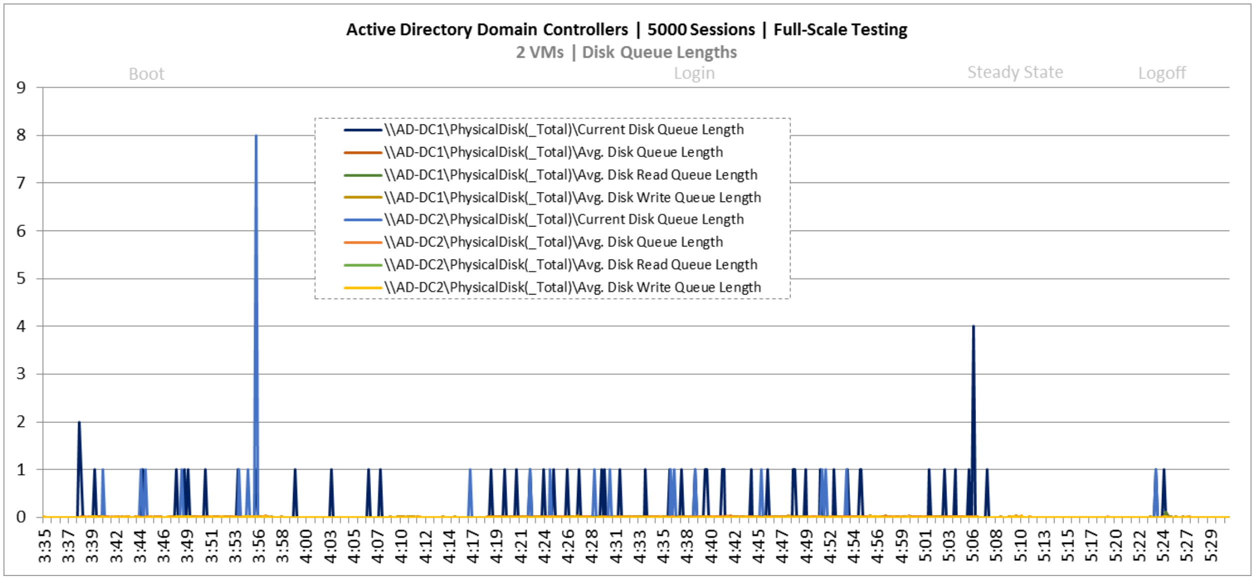

Figure 213 Full Scale | 5000 Mixed Users | Active Directory Domain Controllers | Disk IO Operations

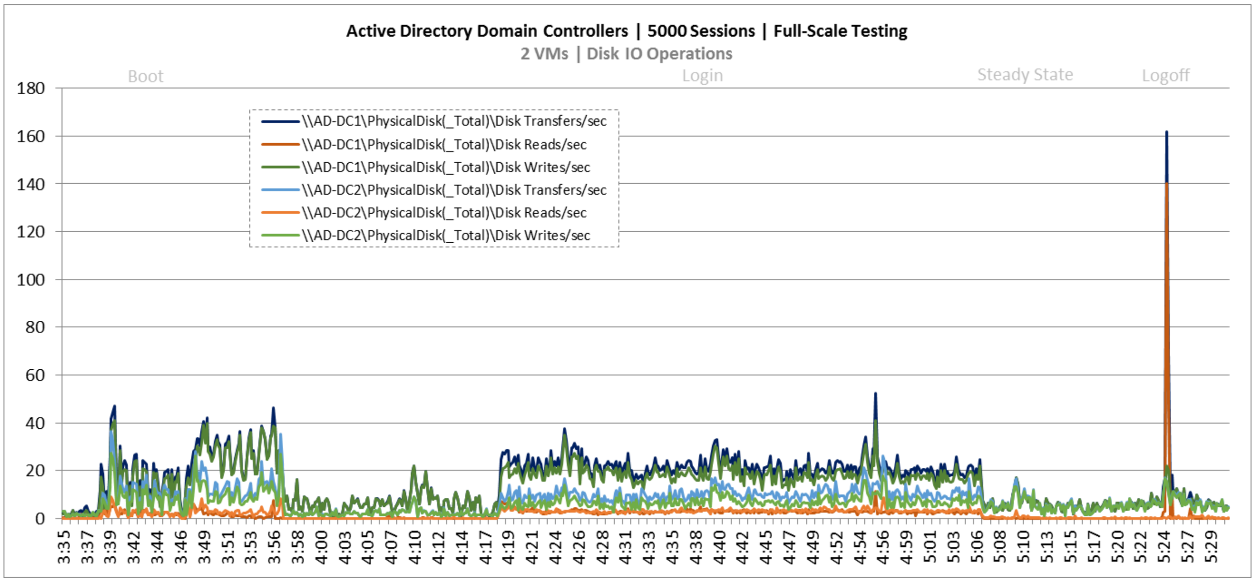

Figure 214 Full Scale | 5000 Mixed Users | SQL Server | CPU Utilization

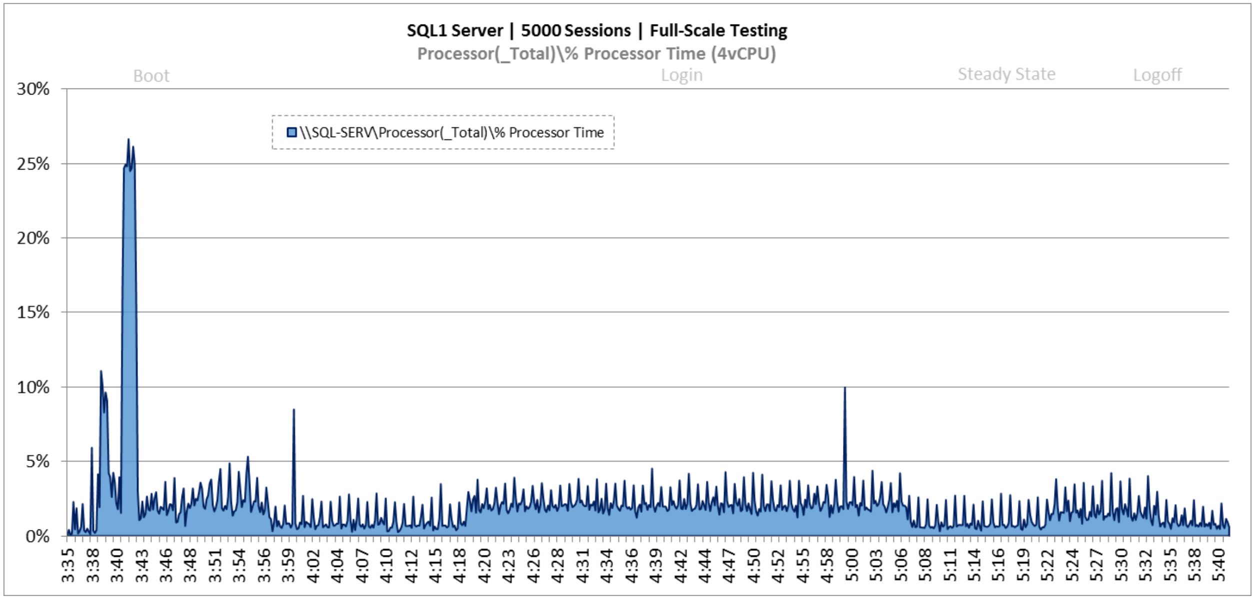

Figure 215 Full Scale | 5000 Mixed Users | SQL Server | Memory Utilization

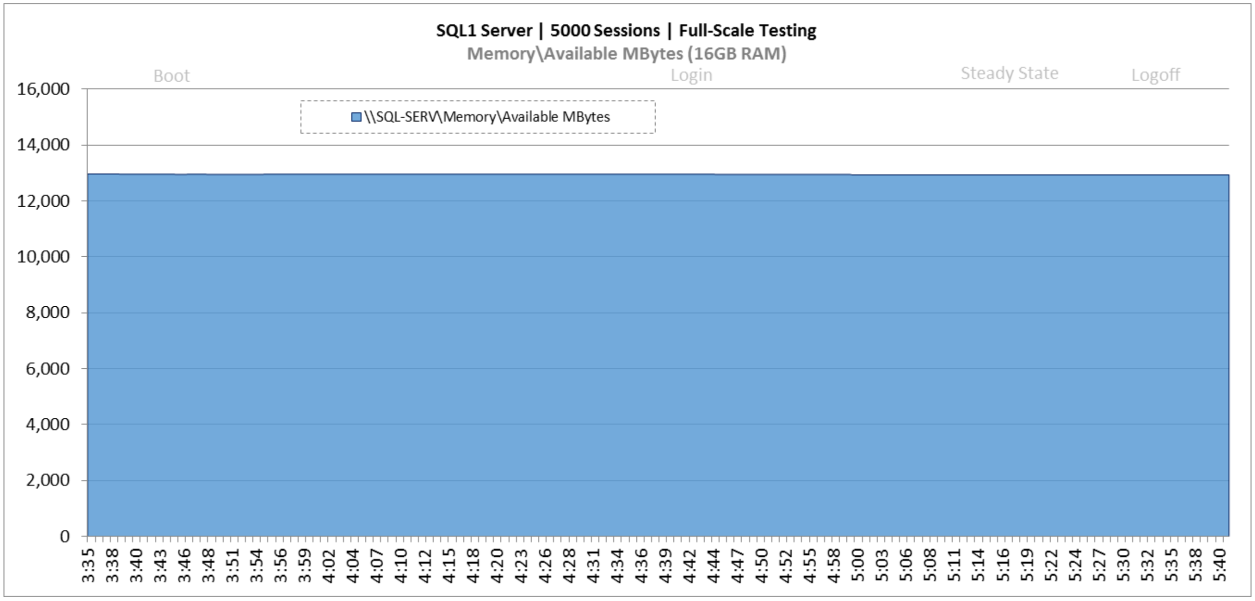

Figure 216 Full Scale | 5000 Mixed Users | SQL Server | Network Utilization

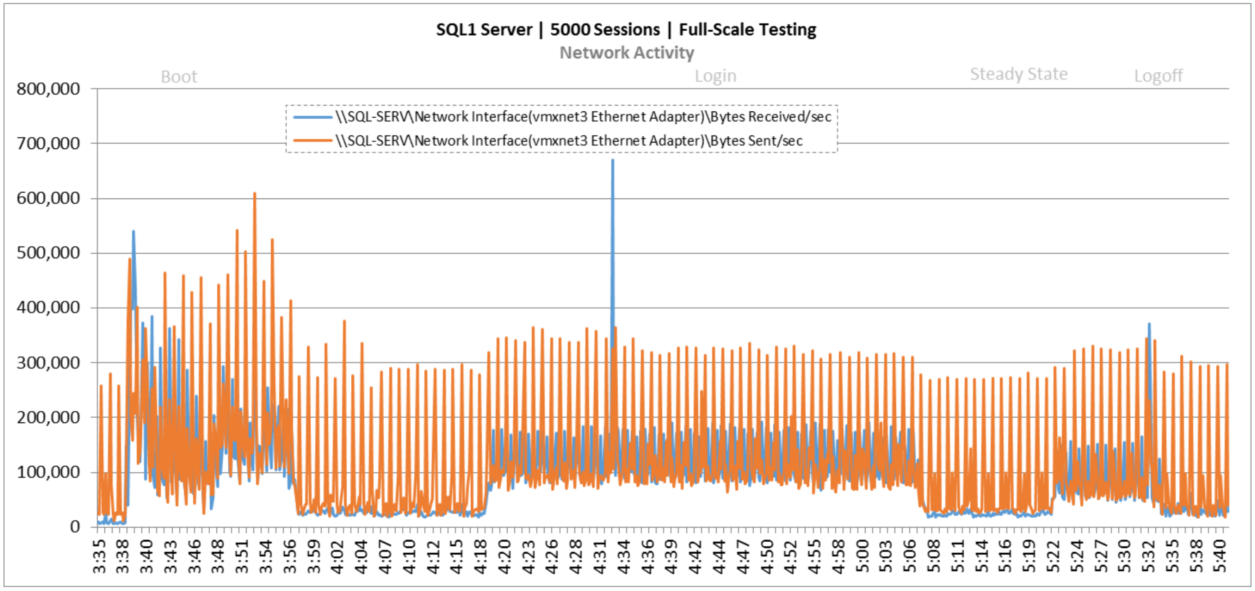

Figure 217 Full Scale | 5000 Mixed Users | SQL Server | Disk Queue Lengths

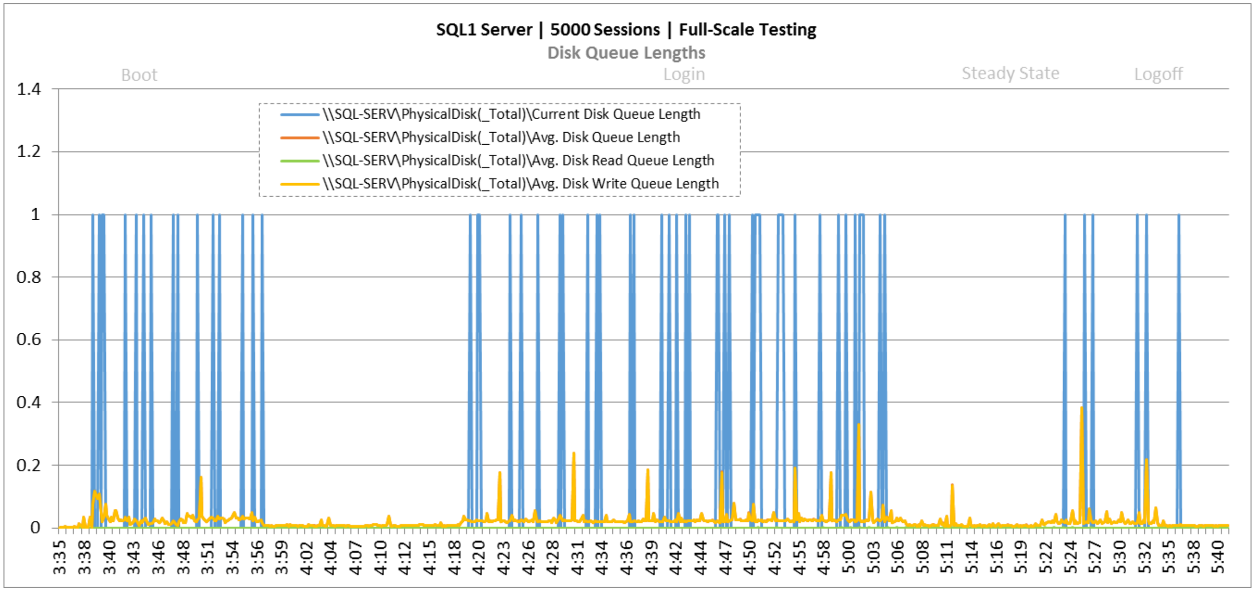

Figure 218 Full Scale | 5000 Mixed Users | SQL Server | Disk IO

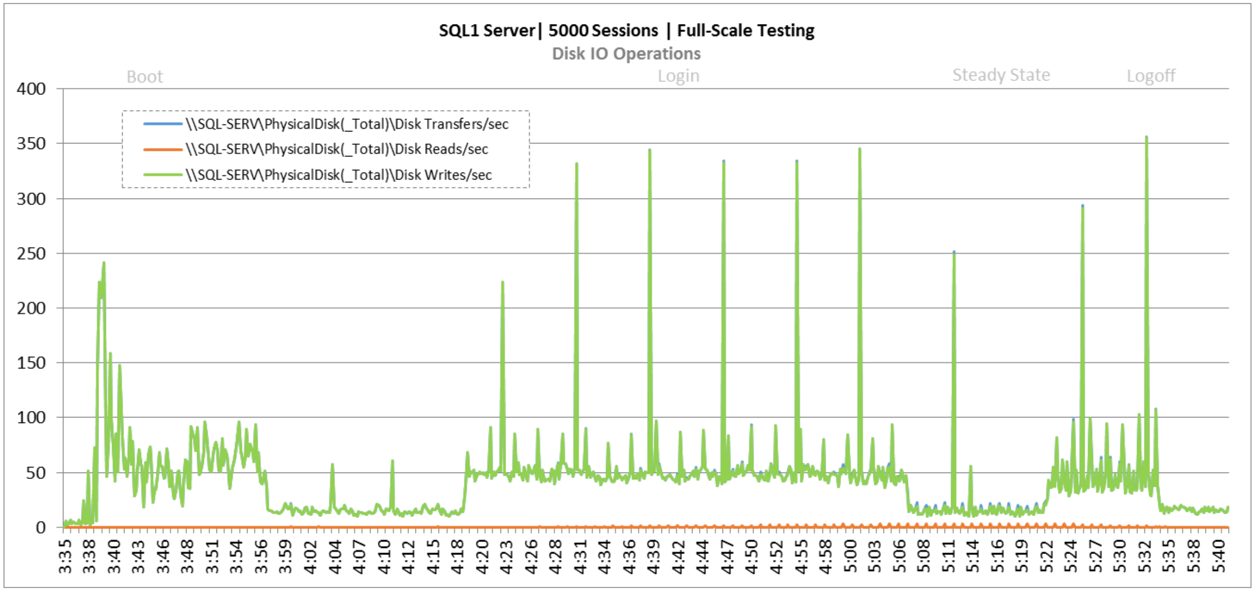

Pure Storage FlashArray//m50 Detailed Test Results for Cluster and Use Case Scalability

The following section will highlight and provide analysis of the Pure Storage FlashArray//m50 performance results for each of the cluster test cases identified earlier in this document. Specifically, we will show and discuss results for the following six test scenarios:

1. 2600 Windows 10 x64 RDS XenApp Sessions

2. 1200 Windows 10 x64 Non-Persistent XenDesktop PVS Sessions

3. 1200 Windows 10 x64 Persistent XenDesktop PVS Sessions

4. 5000 Mixed Workload XenApp and XenDesktop Sessions (Combination of Scenarios 1-3)

5. 3800 Mixed Workload XenDesktop Sessions + Non-disruptive Pure Storage Array Upgrade

6. 1200 Windows 7 x64 Non-Persistent PVS Sessions

From a storage perspective, it is critical to maintain a latency of less than a millisecond in order to guarantee a good end-user experience no matter what amount of IOPS and bandwidth are being driven. As we will see, Pure 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 FlashArray//m50.

The upcoming screenshots of the Pure Storage GUI during each simulation run showcase the important statistics during the various phases of each test run. The different phases of each Login VSI test are denoted across the top of each graph and are individually identified. The first phase (green arrows and text box) are the desktops booting up and the Citrix VDAs registering against the Delivery Controller. Followed by this stage is a twenty minute window for all VMs to settle and then we begin the 2880 second Login VSI simulation phase when all sessions are ramping up and logging in. Once all sessions are connected we hold the simulation in steady-state which is denoted by the grey arrows and text box. Finally, the black arrow to the right shows the end of the simulation where simulated users are logged out of the environment. For brevity, we generally did not show the entire logout operation as array activity is relatively minimal during that time and the Login VSI simulation had completed.

In addition, more granular statistics were obtained by extracting front-end telemetry data from the array storage logs and show equivalent performance that matches with the Pure GUI screenshots. These results were plotted with our metrics of interest and are described in Appendix C.

Pure Storage FlashArray//m50 Test Results for 2600 RDS Windows 10 x64 Sessions

Using Login VSI as the workload generator in Benchmark mode with the Knowledge Worker user type and with XenApp as the VDI delivery mechanism, our first highlighted cluster test shows that the FlashArray//m50 can easily handle this workload with exceptional end-user experience confirmed from Login VSI.

The first Pure Storage GUI screenshot show the FlashArray//m50 performance during the 2600 XenApp RDS sessions running on top of 80 Windows 2012R2 servers. As with all scenarios, there were three separate 2600 RDS simulation runs completed in total, all with very similar results. As we can see in the below chart from one of these simulations - we maintained latency of less than or close to one millisecond for both read and write operations throughout this entire run. This resulted in a confirmed outstanding end-user experience for the simulated XenApp RDS users independently verified by Login VSI. We noticed a very high write-to-read ratio (an average of approximately 85:15) during this component of the testing with observed peak total values 3.42K IOPS and 370.26 MB/s.

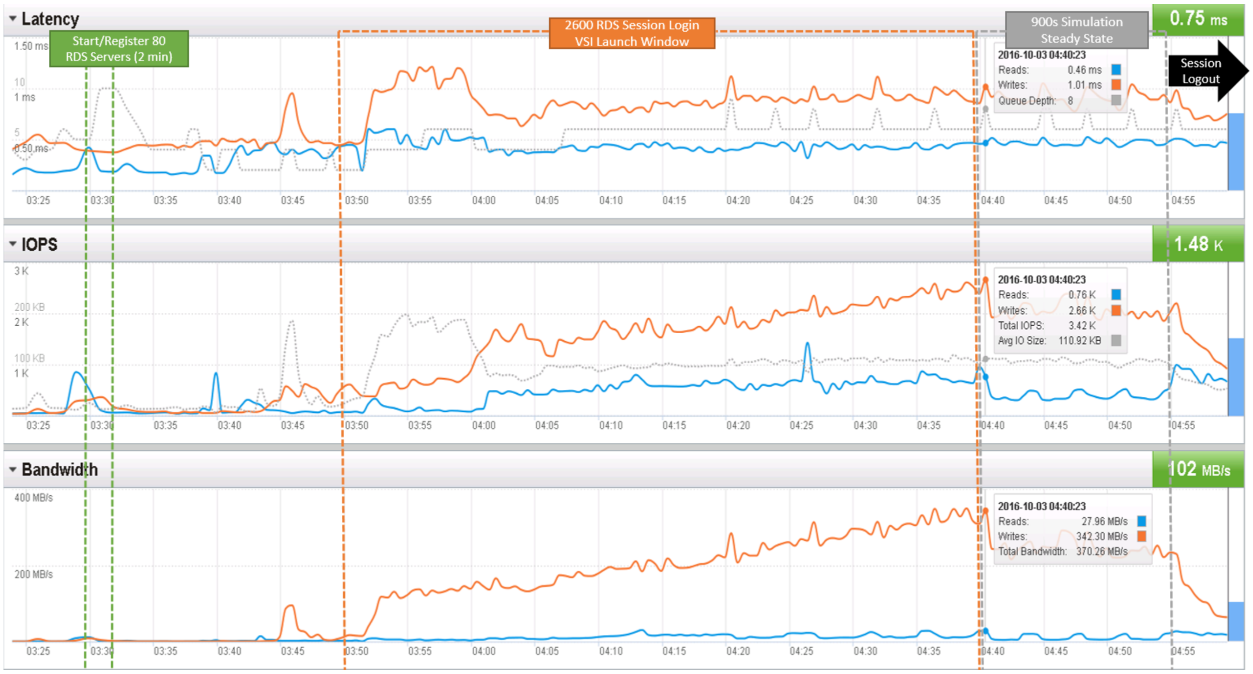

Pure Storage FlashArray//m50 Test Results for 1200 Non-Persistent Windows 10 x64 PVS Desktops

The next cluster-level simulation was to run 1200 non-persistent Windows 10 x64 desktops against the same FlashArray//m50. All Login VSI parameters were kept consistent with the previous RDS test with the only change being to use 1200 desktops created via XenDesktop Provisioning Services rather than XenApp. As can be seen in the below storage metrics, the Pure Storage FlashArray//m50 was clearly able to handle this workload and continued to provide sub-millisecond latency for another impressive Login VSI result.

Another item worth noting from the GUI screen capture shows that latency was consistently sub-millisecond throughout all phases of the Login VSI simulation despite driving hundreds of megabytes of sustained write bandwidth. The brief dip in all performance metrics from approximately 16:15 to 16:25 occurs after the 1200 desktops have registered the VDA and are allowed to settle prior to the Login VSI ramp up simulation beginning.

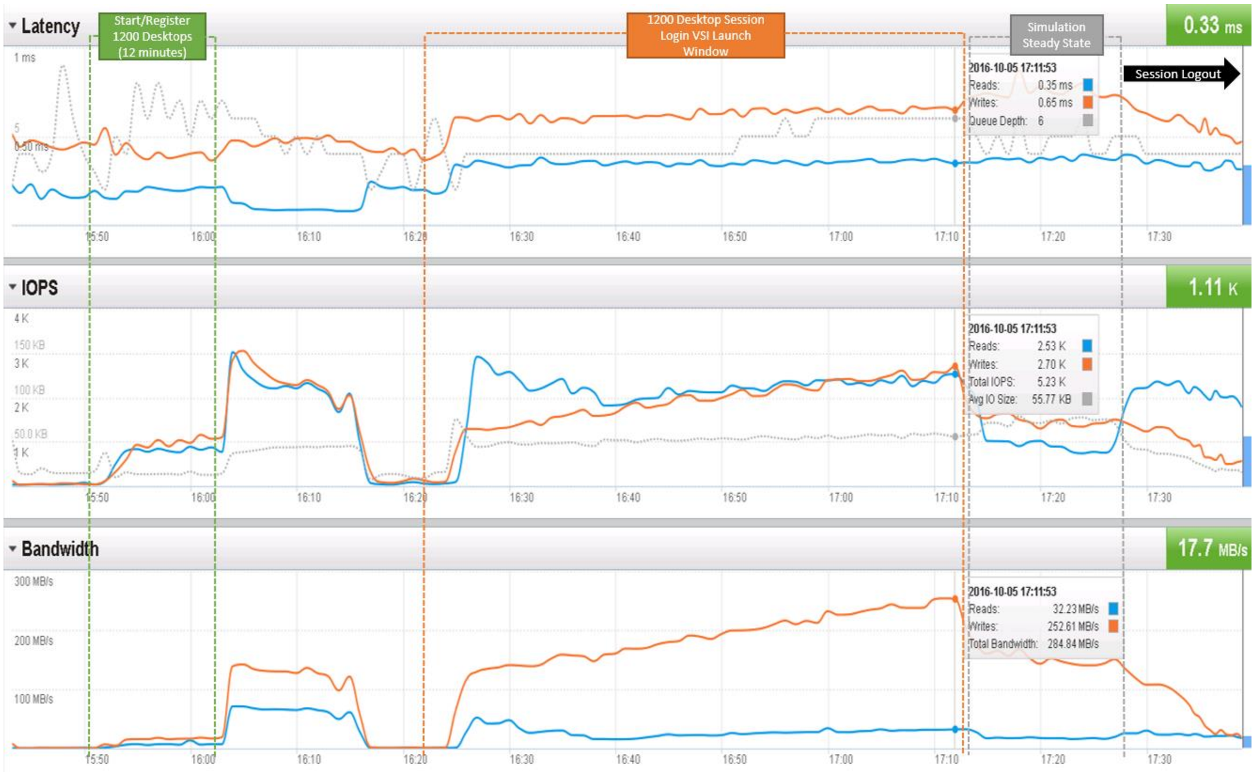

Pure Storage FlashArray//m50 Test Results for 1200 Persistent Windows 10 x64 PVS Desktops

The subsequent cluster-level simulation was to run 1200 persistent Windows 10 x64 desktops. All Login VSI parameters were kept consistent with the previous 1200 non-persistent desktop test with the only change being to use desktops that are persistent rather than non-persistent. These persistent desktops pointed to a Windows CIFS share hosted on Pure Storage for the user profiles managed by Citrix Profile Manager. As can be seen in the below storage metrics, the Pure Storage FlashArray//m50 was clearly able to handle this workload as well and continued to provide sub-millisecond latency for an impressive Login VSI result.

On the left of the GUI screenshot below, we can see that the 1200 desktop persistent bootstorm resulted in a very large amount of read bandwidth and IOPS and was able to successfully register all desktop VDAs in approximately 12 minutes. After a 20 minutes interval in which the VMs were allowed to settle the simulation was started. Yet again we see these desktops driving a large amount of bandwidth (observed peak value of 720.06 MB/s) and IOPS (observed peak value of 34.27K) during the test.

An interesting observation for this simulation relative to the last few is that while the write-IOPs was similar to our non-persistent desktop and RDS experiments, we noticed a much larger amount of read-IOPS/bandwidth from the persistent desktops due to fetching persistent user data from the profile server.

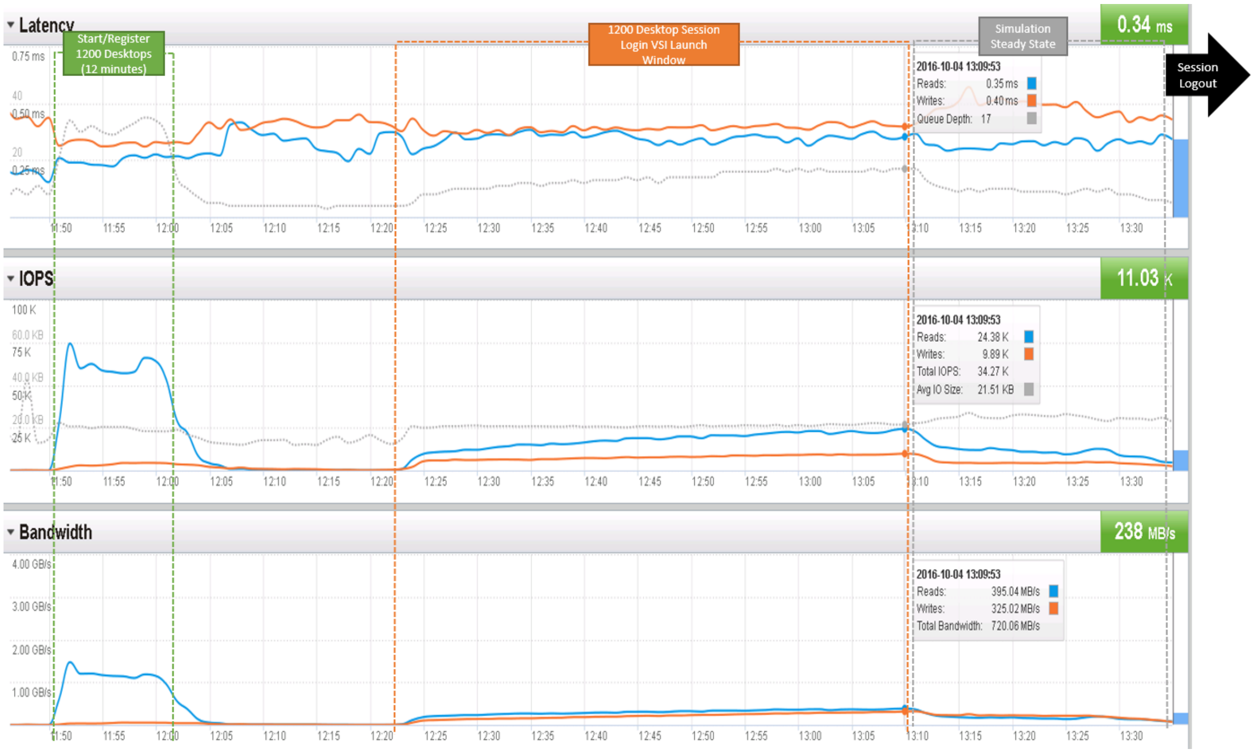

2016-10-04 13:09:53
Reads: 0.35 ms
Writes: 0.40 ms
Queue Depth: 17

2016-10-04 13:09:53
Reads: 24.38 K
Writes: 9.89 K
Total IOPS: 34.27 K
Avg IO Size: 21.51 KB

2016-10-04 13:09:53
Reads: 395.04 MB/s
Writes: 325.02 MB/s
Total Bandwidth: 720.06 MB/s

Pure Storage FlashArray//m50 Test Results for 5000 User Full Scale, Mixed Workload Scalability

The next simulation shows the results of combining all of our previous cluster tests of 2600 XenApp RDS sessions, the 1200 Non-Persistent and the 1200 persistent XenDesktop PVS-based sessions for a full 5000 user Citrix XenApp and XenDesktop simulation on the same Pure Storage FlashArray//m50 array. Yet again we see performant and consistent results that showcase an outstanding user experience at a very large scale.

The GUI screenshot below shows the Pure Storage GUI with the cursor providing more detailed metrics at the start of the simulation during the bootstorm of the desktops. Despite driving nearly 2GB/s in bandwidth during the bootstorm and later close to 1.5GB/s total bandwidth during the test itself, we maintain the desktop responsiveness and performance of low latency throughout the entire test.

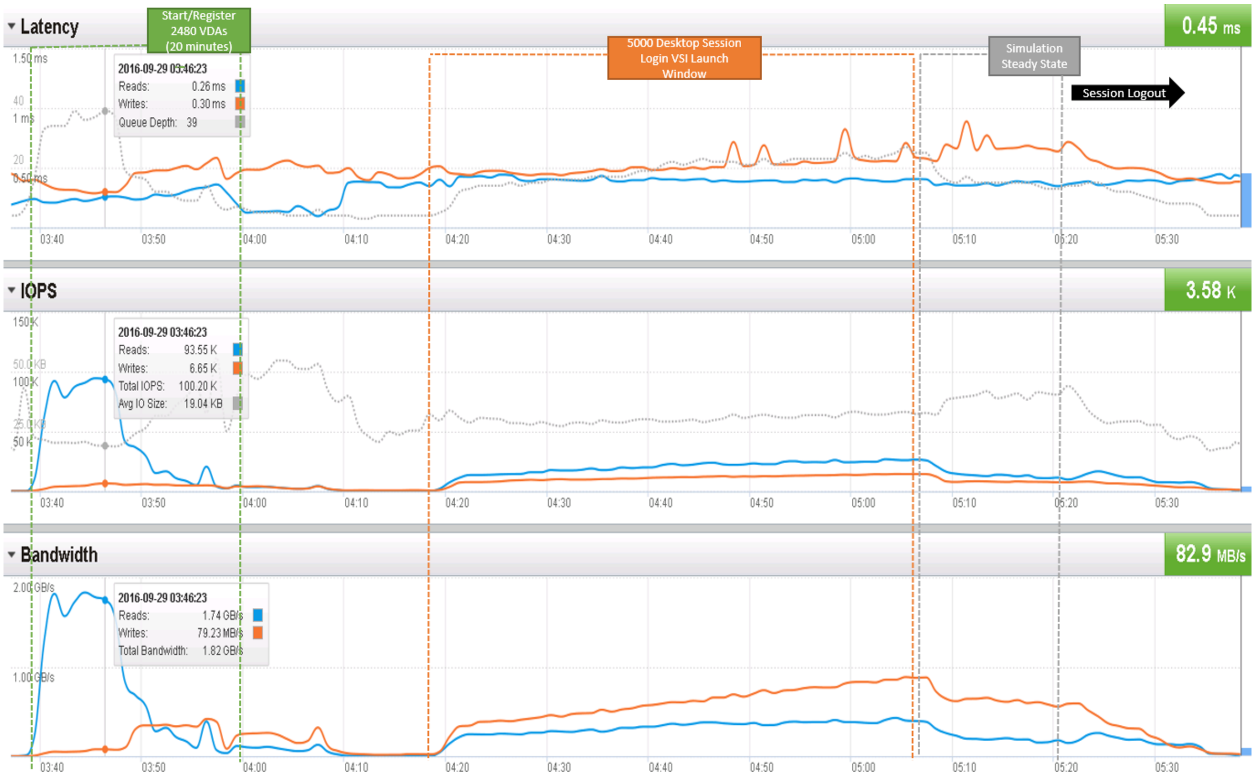

Keeping in mind that this was a pristine lab environment, we can see that our data reduction and overall array utilization during this full test scenario was extremely impressive with only 5% of the overall array space being utilized, allowing substantial room for additional user applications, data and even additional workloads to be hosted on this array without any capacity concerns.

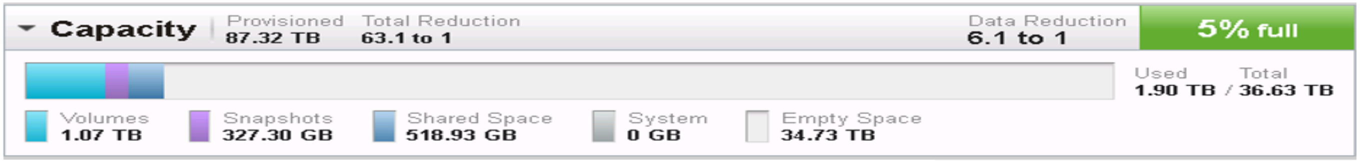

Capacity

Provisioned	Total Reduction	Data Reduction	5% full
87.32 TB	63.1 to 1	6.1 to 1	

Used 1.90 TB / Total 36.63 TB

Volumes	Snapshots	Shared Space	System	Empty Space
1.07 TB	327.30 GB	518.93 GB	0 GB	34.73 TB

Pure Storage FlashArray//m50 Test Results for Large Scale, Mixed Workload Resiliency Testing

We elected to highlight the tremendous resiliency of the Pure Storage FlashArray//m50 by performing an upgrade of the Purity Operating Environment in parallel with a running desktop mixed-workload simulation at scale during steady-state operations. This is to demonstrate that a traditionally disruptive operation like a storage firmware upgrade can be done transparently on Pure Storage with active VDI users being none the wiser. In this test, we upgraded from Purity v4.6.8 to Purity v4.6.11 while 3800 desktops were actively running common knowledge worker tasks against the array.

Worth noting is that this exact procedure would also be followed during a generational controller upgrade as part of the Evergreen Storage business model where customers receive the latest Pure Storage controllers for free every three years so long as a valid maintenance agreement is in place. Another supported scenario following this process would be non-disruptively upgrading to the more performant FlashArray//m70. Pure Storage controllers are stateless and do not require any additional setup other than inserting them into the chassis and connecting cables. Relatedly, capacity expansions (adding an external shelf or denser SSD drives) are also accomplished without any downtime in the middle of production operations.

Upgrading the Purity Operating Environment is typically handled by Pure Storage support for arrays that are managed via Cloud Assist. Customers have the ability to open an SSH tunnel into the array that Pure support can access remotely and carry out the upgrade process. This procedure can also be accomplished locally by KVM or console connection to the array as well by a Pure Storage or other authorized support customer or partner.

To carry out the upgrade, complete the following steps:

1. Copy the upgraded Purity Operating Environment code to both array controllers through FTP or SCP.

2. Install the new Purity code on secondary controller through the CLI.

3. Reboot the secondary controller and wait for it to come back online.

4. Install the new Purity code on primary controller through the CLI.

5. Reboot the primary controller; the array non-disruptively fails over to the upgraded secondary controller which becomes the primary and the new version of Purity OE is in use;

6. The rebooted primary controller comes up as secondary running the upgraded Purity code.

The Pure Storage GUI shown below during the Login VSI simulation run indicates that there was no performance impact at all during the Purity upgrade, which was completed from approximately 12:00PM – 12:20PM while all 3600 desktops were exercising common Knowledge worker applications and operations.

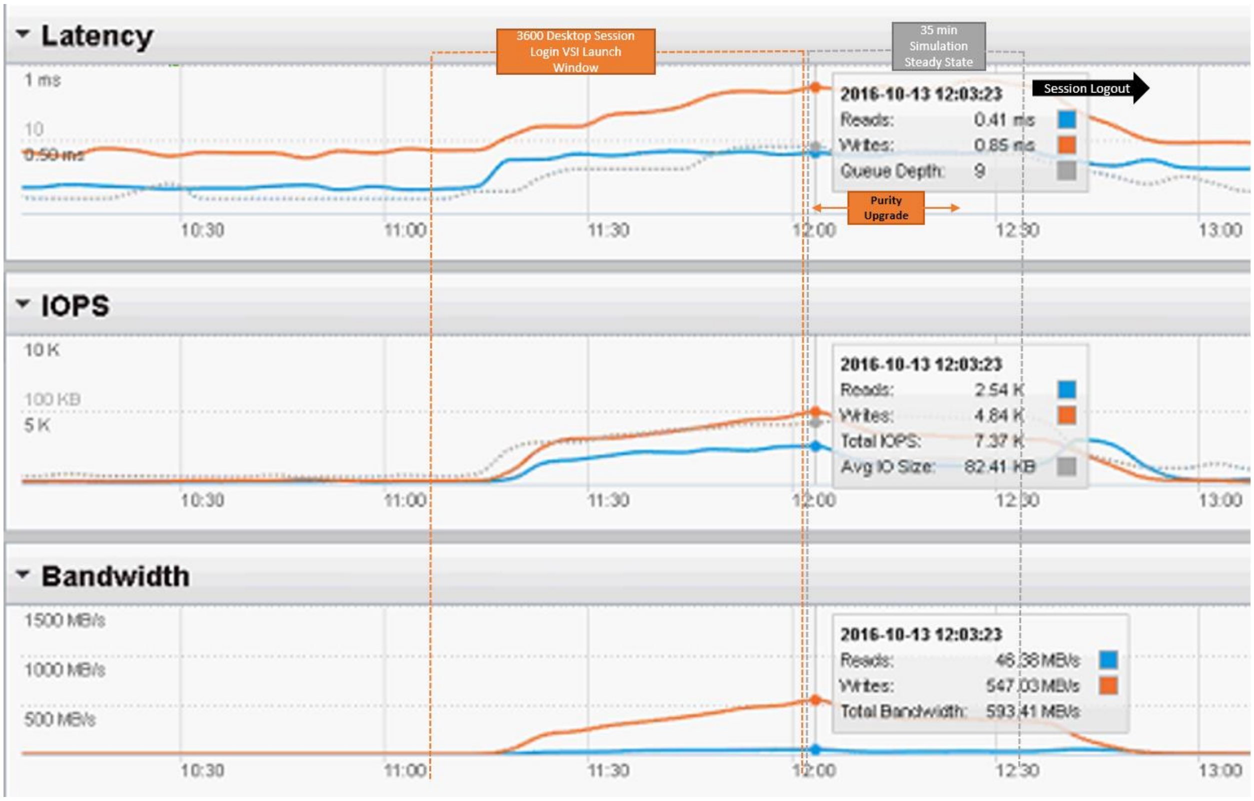

Since this is a more involved operation we wanted to show a deeper level of detail from the Pure Storage array for this particular test. The following charts were built from front-end telemetry data pulled off the array and they show a more thorough picture of the array's behavior during the Purity code upgrade. Immediately worth noting is that the drop in reported metrics to 0 during the upgrade process was a bug in the Purity code that has since been correct in version 4.7.0 and above (for this test we upgraded from Purity 4.6.8 to Purity 4.6.11). Despite some metrics dropping to zero briefly below, the desktops continued to operate without issue as was verified by the Login VSI test and the above Pure Storage GUI screenshot.

The below chart shows the latency of the array during the Purity upgrade process window shown by the orange box. Overall the upgrade process took approximately 20 minutes to complete including controller reboots.

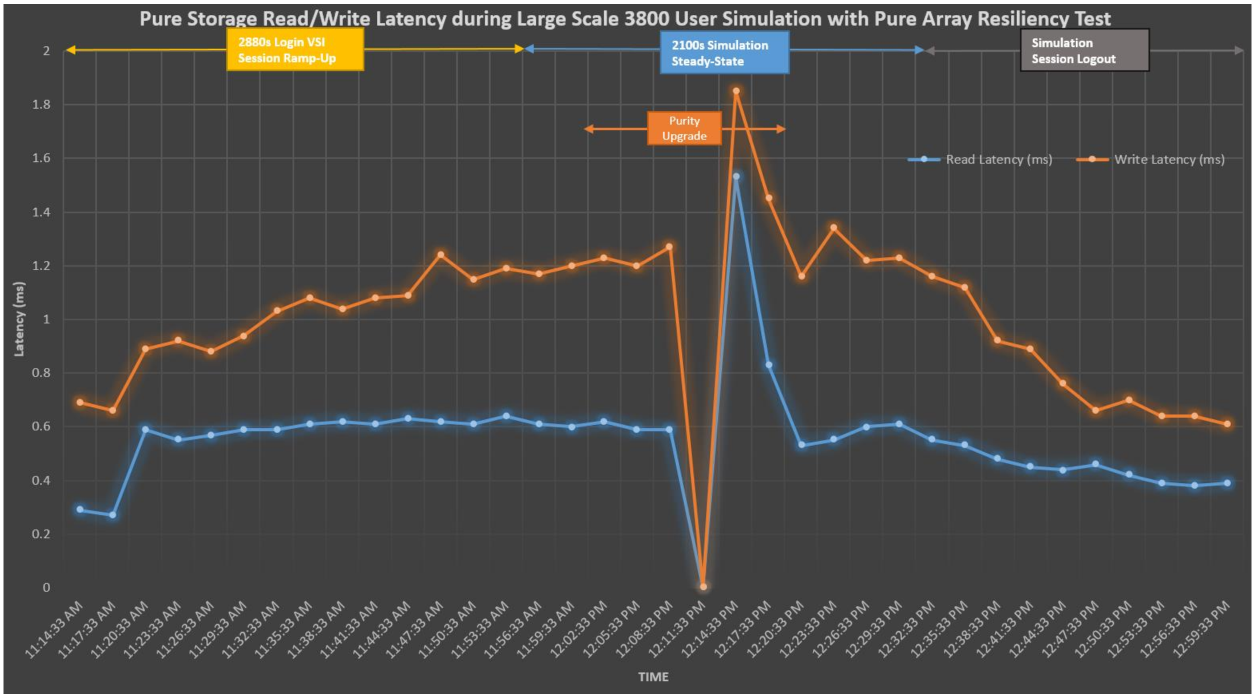

Next, we show IOPs being served during the 3600 mixed workload simulation with the Purity upgrade.

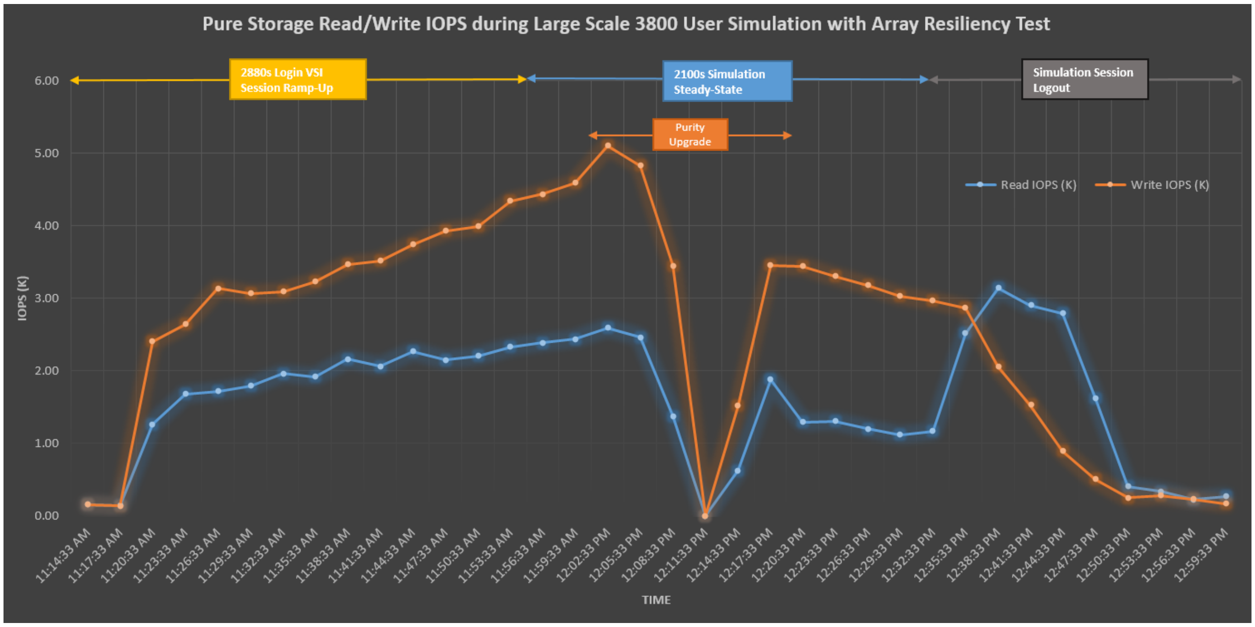

Lastly, this chart shows our bandwidth during the mixed workload + Pure Storage FlashArray//m50 upgrade process.

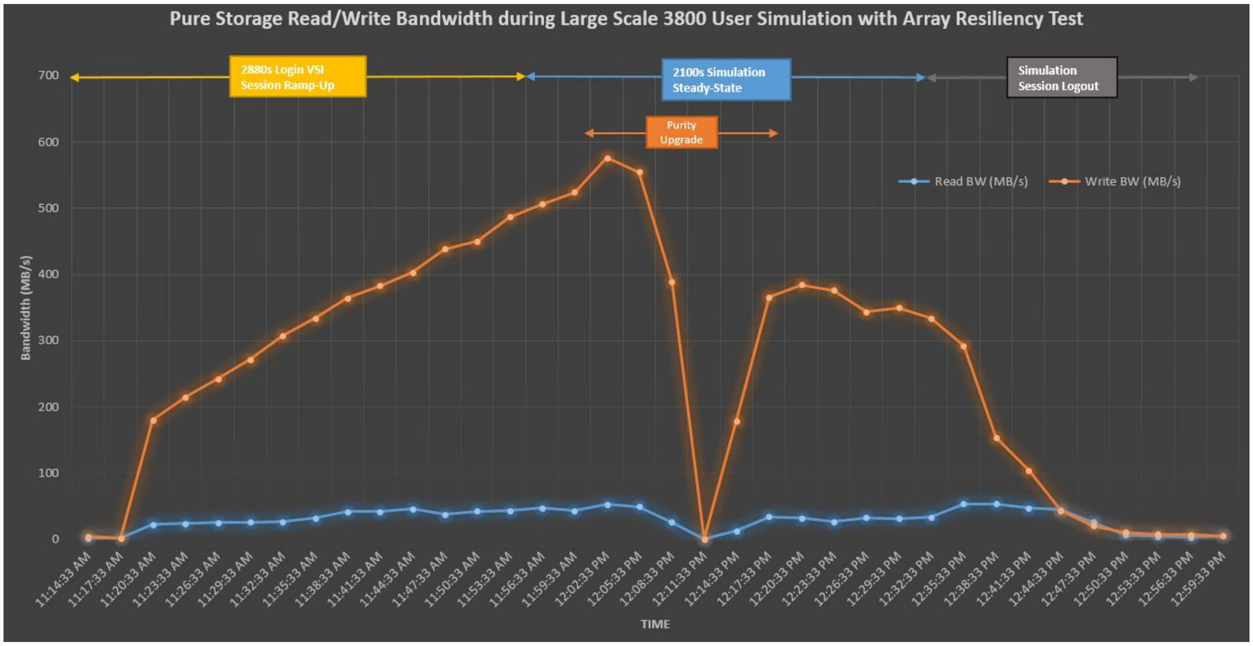

Login VSI provides impartial verification that the Purity upgrade performed in this section was non-disruptive.

The ability of Pure Storage to provide non-disruptive operations, even across generations of hardware, such as the above example, is a key differentiator of this storage array. No more forklift upgrades. No longer do storage administrators need to schedule downtime with the Citrix XenDesktop and/or other team(s) for these kinds of tasks – maintenance operations can now be accomplished in the middle of the workday without causing any interruption to a large number of active Citrix XenDesktop users.

Pure Storage FlashArray//m50 Test Results for 1200 Non-Persistent Windows7x64 PVS Users

While Microsoft has continued to aggressively roll out the Windows 10 OS across their install base, it remains readily apparent that a large percentage of customers remain on the Windows 7 desktop OS until such a time as they are able to qualify and migrate to the Windows 10 OS for use within their unique corporate environment.

To that end, we elected to perform cluster-level testing using 1200 non-persistent PVS-based Windows 7 x64 desktops in order to provide a measure of insight of what to expect from an operational performance perspective when making this transition. Care was taken to make certain that the only variable changed for this test was the base OS. That is, the same size write-cache, Office version, amount of RAM and number of vCPUs were used in our Windows 7 x64 experiment as was used in the previous equivalent Windows 10 x64 test. This section will show storage performance results as well as a comparison between the two operating systems from that perspective.

The screenshot below of the Pure Storage GUI shows the characteristics of the 1200 desktops during the Windows7 Login VSI Knowledge Worker test.

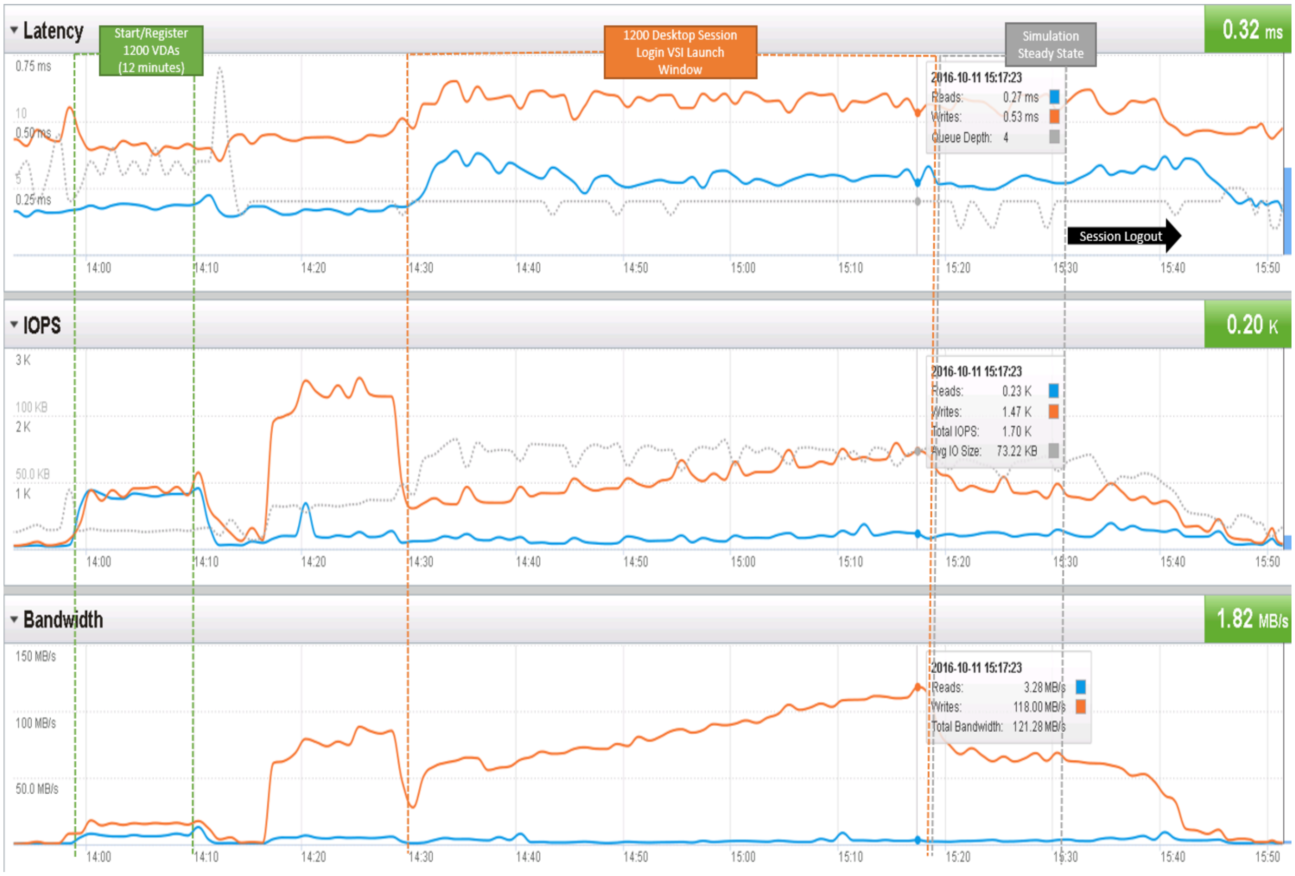

The table below shows the observed maximum latency, IOPS and bandwidth values from the Pure Storage array during both the Windows 7 as well as the Windows 10 non-persistent PVS simulation runs. It is very clear that Windows 10 pulls substantially more storage bandwidth and IOPS relative to Windows 7. Despite these increased storage performance requirements of Windows 10, Pure Storage was able to provide the necessary sub-millisecond latency needed to give a highly performing end-user desktop in both these as well as all other test scenarios covered within this guide.

Guest OS	# of VMs	Read Latency (ms)	Write Latency (ms)	Read IOPS (K)	Write IOPS (K)	Read Bandwidth (MB/s)	Write Bandwidth (MB/s)	Total Bandwidth (MB/s)
Windows 7 x64	1200	0.27	0.53	0.23	1.47	3.28	118	121.28
Windows 10 x64	1200	0.35	0.65	2.53	2.7	32.23	252.61	284.84

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 1000 Users, two chassis 8 mixed workload VDI/HSD host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 1000 user system.

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 6248UP 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 6248UP Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the Pure Storage scaling section. Please refer the Pure Storage web site for scalability guidelines.

Scalability of Citrix XenDesktop 7.9 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

These and other various aspects of scalability considerations are described in greater detail in "XenDesktop - Modular Reference Architecture" document and should be a part of any XenDesktop design.

When designing and deploying this CVD environment, best practices were followed including the following:

- Citrix recommends using N+1 schema for virtualization host servers to accommodate resiliency. In all Reference Architectures (such as this CVD), this recommendation is applied to all host servers.

- All Provisioning Server Network Adapters are configured to have a static IP and management.

We used the XenDesktop Setup Wizard in PVS. Wizard does an excellent job of creating the desktops automatically and it's possible to run multiple instances of the wizard, provided the deployed desktops are placed in different catalogs and have different naming conventions. To use the PVS XenDesktop Setup Wizard, at a minimum you need to install the Provisioning Server, the XenDesktop Controller, and configure hosts, as well as create VM templates on all datastores where desktops will be deployed.

Summary

FlashStack delivers a platform for Enterprise VDI deployments and cloud datacenters using Cisco UCS Blade and Rack Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches and fibre channel-attached Pure Storage FlashArray//m. FlashStack is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlashStack provides to customers wishing to deploy enterprise-class VDI for 5000 users at a time.

Get More Business Value with Services

Whether you are planning your next-generation environment, need specialized know-how for a major deployment, or want to get the most from your current storage, Cisco Advanced Services, Pure Storage and our certified partners can help. We collaborate with you to enhance your IT capabilities through a full portfolio of services that covers your IT lifecycle with:

- Strategy services to align IT with your business goals:
- Design services to architect your best storage environment
- Deploy and transition services to implement validated architectures and prepare your storage environment
- Operations services to deliver continuous operations while driving operational excellence and efficiency.

In addition, Cisco Advanced Services and Pure Storage provides in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Authors

Mike Brennan, Product and Technical Marketing Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Mike is the Desktop Virtualization Solutions Product Manager and TME Manager in Cisco's Computer Systems Product Group.

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:

- Kyle Grossmiller, Solutions Architect, Customer Solutions Group, Pure Storage, Inc. Kyle provided key technical guidance with the FlashArray//m50 and significant contributions to this document.

- Frank Anderson, former Cisco Desktop Virtualization Architect. Frank currently works at VMware's End User Computing group in a similar role. Frank executed all of the setup, testing and produced the graphics in this Cisco Validated Design.

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-b200-m4-blade-server/index.html
- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html

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
- http://www.cisco.com/c/en/us/products/switches/nexus-9372px-switch/index.html
- http://www.cisco.com/c/en/us/products/switches/nexus-1000v-switch-vmware-vsphere/index.html

Citrix References

- http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-9.html
- https://docs.citrix.com/en-us/provisioning/7-9.html
- https://www.citrix.com/go/jmp/upm.html

- https://www.citrix.com/virtualization/hdx/
- https://www.citrix.com/products/xenapp-xendesktop/hdx-3d-pro.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

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

Pure Storage Reference Documents

- Pure Storage FlashArray//m Datasheet
  - http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/PureStorage_FlashArraym-Brochure.pdf
- Pure Storage FlashStack Converged Infrastructure Solutions
  - http://www.purestorage.com/solutions/infrastructure/flashstack.html
- Pure Storage Best Practices Guide for VMware vSphere
  - http://www.purestorage.com/resources/type-a/WP-PureStorageandVMwarevSphereBestPracticesGuide_Request.html
- Pure Storage and VMware Storage APIs for Array Integration
  - http://www.purestorage.com/resources/type-b/WP-PureStorageandVAAI_Request.html
- FlashStack Converged Infrastructure for Citrix XenDesktop 7.7 Design Guide

– http://www.purestorage.com/resources/type-a/pure-storage-flashstack-citrix-xd-7-7.html

- Consolidating Workloads with VMware and Pure Storage
  - http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/ESG_Lab_Validation_Summary_Pure_Storage_Sep_2015.pdf

Appendix A – Cisco Nexus Ethernet and MDS Fibre Channel Switch Configurations

Ethernet Network Configuration

The following section provides a detailed procedure for configuring the Cisco Nexus 9000 and 1000V Switches used in this study.

Cisco Nexus 9172PX-A Configuration

```
!Command: show running-config
!Time: Fri Nov  4 08:01:27 2016

version 7.0(3)I2(2e)
switchname N9K-A
vdc N9K-A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp

feature vpc
feature lldp
clock protocol none vdc 1

no password strength-check
username admin password 5 $1$7LGx6.qz$3xYFFA2B9CgCGt0n3EOm60  role
network-admin
ssh key rsa 2048
ip domain-lookup
no service unsupported-transceiver
copp profile strict
snmp-server user admin network-admin auth md5
0x4112400e904ff685e2625e9e302ec9ad
 priv 0x4112400e904ff685e2625e9e302ec9ad localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner
PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner
PMON@INFO
ntp server 10.10.160.2
ntp peer 10.10.160.3
ntp server 171.68.38.66 use-vrf management

vlan 1,160-163,166-167
vlan 160
  name In-Band-Mgmt
vlan 161
  name Infra-Mgmt
vlan 162
  name VDI
vlan 163
  name Storage
vlan 166
  name vMotion
vlan 167
  name Launcher

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
hardware qos ns-buffer-profile mesh
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.12 source 10.29.164.11
  delay restore 150
  peer-gateway
  auto-recovery

interface Vlan1
  description default native vlan
  no shutdown
  no ip redirects
  ip address 10.29.164.253/24
  no ipv6 redirects

interface Vlan160
  description In Band Mmgmt vlan 160
  no shutdown
  no ip redirects
  ip address 10.10.160.2/24
  hsrp version 2
  hsrp 11
    preempt
    ip 10.10.160.1

interface Vlan161
  description Infrastructure Mgmt vlan 161
  no shutdown
  no ip redirects
  ip address 10.10.161.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 12
    preempt
    ip 10.10.161.1

interface Vlan162
  description VDI vlan 162
  no shutdown
  ip address 10.10.192.2/19
  no ipv6 redirects
  hsrp version 2
  hsrp 13
    preempt
    priority 110
    ip 10.10.192.1
  ip dhcp relay address 10.10.161.30
  ip dhcp relay address 10.10.161.31

interface Vlan167
  description VSI Launchers vlan 167
  no shutdown
  ip address 10.10.167.2/24
  hsrp version 2
  hsrp 15
    preempt
    ip 10.10.167.1
  ip dhcp relay address 10.10.167.20

interface port-channel10
  description VPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type network
  vpc peer-link

interface port-channel29
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type edge trunk
  mtu 9216
  vpc 29

interface port-channel31
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type edge trunk
  mtu 9216
  vpc 31

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19
  description OOB-Mgmt_access

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
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 29 mode active

interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 29 mode active

interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 31 mode active

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 31 mode active

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39
interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47
  description VPC Peer N9K-B:1/47
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  channel-group 10 mode active

interface Ethernet1/48
  description VPC Peer N9K-B:1/48
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  channel-group 10 mode active

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 10.29.164.11/24
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
  session-limit 16
boot nxos bootflash:/nxos.7.0.3.I2.2e.bin
```

Cisco Nexus 9172PX-B Configuration

```
!Command: show running-config
```

```
!Time: Fri Nov  4 08:05:10 2016

version 7.0(3)I2(2e)
switchname N9K-B
vdc N9K-B id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol none vdc 1

no password strength-check
username admin password 5
$5$CLLPGB$DsP.fsZZVX1H6IgxcPGZCyE4Y0./mCKqGdqfj30xsj2
 role network-admin
ip domain-lookup
system default switchport shutdown
no service unsupported-transceiver
copp profile strict
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
ntp server 10.10.160.3
ntp server 171.68.38.66 use-vrf management
ntp master 8

vlan 1,160-163,166-167
vlan 160
  name In-Band-Mgmt
vlan 161
  name Infra-Mgmt
vlan 162
  name VDI
```

```
vlan 163
  name Storage
vlan 166
  name vMotion
vlan 167
  name Launcher

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
hardware qos ns-buffer-profile mesh
vpc domain 10
  role priority 20
  peer-keepalive destination 10.29.164.11 source 10.29.164.12
  delay restore 150
  peer-gateway
  auto-recovery


interface Vlan1
  description default native vlan
  no shutdown
  no ip redirects
  ip address 10.29.164.254/24
  no ipv6 redirects

interface Vlan160
  description In Band Mmgmt vlan 160
  no shutdown
  no ip redirects
  ip address 10.10.160.3/24
  hsrp version 2
  hsrp 11
    preempt
    ip 10.10.160.1

interface Vlan161
  description Infrastructure Mgmt vlan 161
  no shutdown
  no ip redirects
  ip address 10.10.161.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 12
    preempt
    ip 10.10.161.1
```

```
interface Vlan162
  description VDI vlan 162
  no shutdown
  no ip redirects
  ip address 10.10.192.3/19
  no ipv6 redirects
  hsrp version 2
  hsrp 13
    preempt
    priority 110
    ip 10.10.192.1
  ip dhcp relay address 10.10.161.30
  ip dhcp relay address 10.10.161.31

interface Vlan167
  description VSI Launchers vlan 167
  no shutdown
  ip address 10.10.167.3/24
  hsrp version 2
  hsrp 15
    preempt
    ip 10.10.167.1
  ip dhcp relay address 10.10.167.20

interface port-channel10
  description VPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type network
  vpc peer-link

interface port-channel29
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type edge trunk
  mtu 9216
  vpc 29

interface port-channel31
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  spanning-tree port type edge trunk
  mtu 9216
  vpc 31

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3
interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19
  description Jumphost
  switchport access vlan 160
  spanning-tree port type edge trunk
  no shutdown

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
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 29 mode active
  no shutdown

interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 29 mode active
  no shutdown

interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 31 mode active
  no shutdown

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  mtu 9216
  channel-group 31 mode active
  no shutdown

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43
interface Ethernet1/44

interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47
  description VPC Peer N9K-A:1/47
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  channel-group 10 mode active
  no shutdown

interface Ethernet1/48
  description VPC Peer N9K-A:1/48
  switchport mode trunk
  switchport trunk allowed vlan 1,160-164,166-167
  channel-group 10 mode active
  no shutdown

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 10.29.164.12/24
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I2.2e.bin
```

Fibre Channel Network Configuration

Cisco MDS 9148-A Configuration

```
!Time: Fri Nov  4 15:15:49 2016

version 6.2(9a)

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

username admin password 5 $1$loX7vizP$00IbhSFcpx6WufBmOMKB.1  role
network-admin

ip domain-lookup

ip host MDS-A  10.29.164.64

aaa group server radius radius

snmp-server user admin network-admin auth md5
0x6c81eb7167a2e69497a60698ca3957da

 priv 0x6c81eb7167a2e69497a60698ca3957da localizedkey

snmp-server host 10.155.160.192 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

vsan database

  vsan 20

device-alias database

 vsan 20 wwn 20:00:00:25:b5:1f:1a:24 fcid 0xc20511 dynamic

  vsan 20 wwn 20:00:00:25:b5:1f:1a:0c fcid 0xc2041c dynamic

  vsan 20 wwn 52:4a:93:72:0d:21:6b:11 fcid 0xc20000 dynamic

  vsan 20 wwn 52:4a:93:72:0d:21:6b:00 fcid 0xc20100 dynamic

  vsan 20 wwn 20:4f:54:7f:ee:45:29:80 fcid 0xc20200 dynamic

  vsan 20 wwn 20:50:54:7f:ee:45:29:80 fcid 0xc20300 dynamic
  vsan 20 wwn 20:4f:54:7f:ee:45:2a:40 fcid 0xc20400 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:1d fcid 0xc20401 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:17 fcid 0xc20406 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:15 fcid 0xc20404 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:09 fcid 0xc20408 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:07 fcid 0xc20502 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:05 fcid 0xc20504 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:03 fcid 0xc20402 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:0b fcid 0xc20503 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:0d fcid 0xc20509 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:13 fcid 0xc20405 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:11 fcid 0xc20508 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:7f fcid 0xc20416 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:21 fcid 0xc20510 dynamic
  vsan 20 wwn 20:50:54:7f:ee:45:2a:40 fcid 0xc20500 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:1e fcid 0xc20407 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:1b fcid 0xc20403 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:19 fcid 0xc20506 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:00 fcid 0xc20501 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:00 fcid 0xc2040d dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:02 fcid 0xc2040e dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:04 fcid 0xc20410 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:06 fcid 0xc20507 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:08 fcid 0xc20513 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:1e fcid 0xc2050c dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:0e fcid 0xc20411 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:0a fcid 0xc2041b dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:1c fcid 0xc2040c dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:10 fcid 0xc20517 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:20 fcid 0xc20505 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:22 fcid 0xc20412 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:1a fcid 0xc20409 dynamic
```

```
  vsan 20 wwn 20:00:00:25:b5:1f:1a:16 fcid 0xc2041a dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:14 fcid 0xc2050b dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:26 fcid 0xc2050a dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:12 fcid 0xc20413 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:2a fcid 0xc2040f dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:2e fcid 0xc2050d dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:30 fcid 0xc2040b dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:2c fcid 0xc20418 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:32 fcid 0xc2050f dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:34 fcid 0xc20419 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:36 fcid 0xc2050e dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:38 fcid 0xc20515 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:18 fcid 0xc20415 dynamic
vsan 20 wwn 52:4a:93:72:0d:21:6b:13 fcid 0xc20600 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:28 fcid 0xc20414 dynamic
interface port-channel1
channel mode active
switchport rate-mode dedicated
vsan database
vsan 20 interface fc1/1
vsan 20 interface fc1/2
vsan 20 interface fc1/3
vsan 20 interface fc1/4
switchname MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
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 fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
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 fc1/43

interface fc1/44

interface fc1/45

interface fc1/46

interface fc1/47

interface fc1/48

interface fc1/1

interface fc1/2

interface fc1/11

interface fc1/12

!Active Zone Database Section for vsan 20

zone name Infra-serv1-fc0 vsan 20

member pwwn 20:00:00:25:b5:b1:1b:21

member pwwn 52:4a:93:72:0d:21:6b:11

member pwwn 52:4a:93:72:0d:21:6b:00

zone name Infra-serv2-fc0 vsan 20

member pwwn 20:00:00:25:b5:b1:1b:7f

member pwwn 52:4a:93:72:0d:21:6b:11

member pwwn 52:4a:93:72:0d:21:6b:00

zone name pure-vdi-serv1-fc0 vsan 20

   member pwwn 20:00:00:25:b5:b1:1b:00

    member pwwn 52:4a:93:72:0d:21:6b:11

    member pwwn 52:4a:93:72:0d:21:6b:00

zone name pure-vdi-serv2-fc0 vsan 20

    member pwwn 20:00:00:25:b5:1f:1a:00

    member pwwn 52:4a:93:72:0d:21:6b:11
```

```
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv3-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:02
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv5-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:06
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv6-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:08
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv8-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:0c
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv7-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:0a
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv9-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:0e
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv10-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:10
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv11-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:12
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv12-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:14
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv13-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:16
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv16-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:1c
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv17-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:1e
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv18-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:20
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv19-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:22
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv20-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:24
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv21-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:26
        member pwwn 52:4a:93:72:0d:21:6b:11
```

```
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv24-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:2c
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv25-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:2e
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv26-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:30
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv27-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:32
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv28-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:34
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv29-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:36
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv30-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:38
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv15-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:1a
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv14-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:18
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv23-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:2a
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zone name pure-vdi-serv22-fc0 vsan 20
        member pwwn 20:00:00:25:b5:1f:1a:28
        member pwwn 52:4a:93:72:0d:21:6b:11
        member pwwn 52:4a:93:72:0d:21:6b:00
    zoneset name pure-vdi-fab-A vsan 20
        member Infra-serv1-fc0
        member Infra-serv2-fc0
        member pure-vdi-serv4-fc0
        member pure-vdi-serv1-fc0
        member pure-vdi-serv2-fc0
        member pure-vdi-serv3-fc0
        member pure-vdi-serv5-fc0
        member pure-vdi-serv6-fc0
        member pure-vdi-serv8-fc0
        member pure-vdi-serv7-fc0
        member pure-vdi-serv9-fc0
        member pure-vdi-serv10-fc0
        member pure-vdi-serv11-fc0
        member pure-vdi-serv12-fc0
        member pure-vdi-serv13-fc0
        member pure-vdi-serv16-fc0
        member pure-vdi-serv17-fc0
        member pure-vdi-serv18-fc0
        member pure-vdi-serv19-fc0
        member pure-vdi-serv20-fc0
        member pure-vdi-serv21-fc0
        member pure-vdi-serv24-fc0
        member pure-vdi-serv25-fc0
        member pure-vdi-serv26-fc0
        member pure-vdi-serv27-fc0
        member pure-vdi-serv28-fc0
        member pure-vdi-serv29-fc0
        member pure-vdi-serv30-fc0
        member pure-vdi-serv15-fc0
        member pure-vdi-serv14-fc0
        member pure-vdi-serv23-fc0
        member pure-vdi-serv22-fc0
    zoneset activate name pure-vdi-fab-A vsan 20
    do clear zone database vsan 20
    !Full Zone Database Section for vsan 20
    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
interface fc1/6
port-license acquire
interface fc1/7
no port-license
no shutdown
interface fc1/8
no port-license
no shutdown
interface fc1/9
port-license acquire
no shutdown
interface fc1/10
port-license acquire
no shutdown
interface fc1/11
port-license acquire
channel-group 1 force
no shutdown
interface fc1/12
port-license acquire
channel-group 1 force
no shutdown
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
interface fc1/18
port-license acquire
interface fc1/19
port-license acquire
interface fc1/20
  port-license acquire
interface fc1/21
port-license acquire
interface fc1/22
port-license acquire
interface fc1/23
port-license acquire
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
interface fc1/26
port-license acquire
 interface fc1/27
port-license acquire
interface fc1/28
port-license acquire
interface fc1/29
port-license acquire
interface fc1/30
```

```
port-license acquire
interface fc1/31
port-license acquire
interface fc1/32
interface fc1/33
port-license acquire
interface fc1/34
port-license acquire
interface fc1/35
port-license acquire
interface fc1/36
port-license acquire
interface fc1/37
port-license acquire
interface fc1/38
port-license acquire
interface fc1/39
port-license acquire
interface fc1/40
port-license acquire
interface fc1/41
port-license acquire
interface fc1/42
port-license acquire
interface fc1/43
port-license acquire
interface fc1/44
port-license acquire
interface fc1/45
port-license acquire
interface fc1/46
port-license acquire
```

```
interface fc1/47
port-license acquire
interface fc1/48
 port-license acquire
interface mgmt0
  ip address 10.29.164.64 255.255.255.0
ip default-gateway 10.29.164.1
MDS-A#
```

Cisco MDS 9148-B Configuration

```
!Time: Fri Nov  4 15:15:49 2016
version 6.2(9a)
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
username admin password 5 $1$dcmFCg/p$0ZC5U6uhI65oOePpHfAzn0  role
network-admin
no password strength-check
ip domain-lookup
ip host MDS-B  10.29.164.128
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
snmp-server host 10.29.164.130 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

vsan database

vsan 30

device-alias database

device-alias commit

fcdomain fcid database


  vsan 30 wwn 52:4a:93:72:0d:21:6b:01 fcid 0x9d0000 dynamic

  vsan 30 wwn 52:4a:93:72:0d:21:6b:10 fcid 0x9d0100 dynamic

  vsan 30 wwn 20:4f:54:7f:ee:45:2a:40 fcid 0x9d0200 dynamic

  vsan 30 wwn 20:50:54:7f:ee:45:2a:40 fcid 0x9d0300 dynamic

  vsan 30 wwn 20:4f:54:7f:ee:45:29:80 fcid 0x9d0400 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:1c fcid 0x9d0411 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:16 fcid 0x9d0402 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:14 fcid 0x9d0506 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:12 fcid 0x9d0503 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:08 fcid 0x9d0404 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:06 fcid 0x9d0505 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:04 fcid 0x9d0405 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:02 fcid 0x9d0508 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:10 fcid 0x9d0406 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:0a fcid 0x9d0507 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:0c fcid 0x9d0504 dynamic

  vsan 30 wwn 20:50:54:7f:ee:45:29:80 fcid 0x9d0500 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:6f fcid 0x9d0413 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:20 fcid 0x9d050e dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:1f fcid 0x9d0401 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:1a fcid 0x9d0403 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:18 fcid 0x9d0408 dynamic

  vsan 30 wwn 20:00:00:25:b5:b1:1b:01 fcid 0x9d050c dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:01 fcid 0x9d041a dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:03 fcid 0x9d040f dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:05 fcid 0x9d0515 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:07 fcid 0x9d050b dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:0d fcid 0x9d050f dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:0f fcid 0x9d041b dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:0b fcid 0x9d0509 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:1d fcid 0x9d0512 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:1f fcid 0x9d0518 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:21 fcid 0x9d0414 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:23 fcid 0x9d0417 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:25 fcid 0x9d040b dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:17 fcid 0x9d040c dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:15 fcid 0x9d050d dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:27 fcid 0x9d050a dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:13 fcid 0x9d0407 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:11 fcid 0x9d0412 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:31 fcid 0x9d040d dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:2d fcid 0x9d0410 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:33 fcid 0x9d0502 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:35 fcid 0x9d040a dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:37 fcid 0x9d0501 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:39 fcid 0x9d0409 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:2b fcid 0x9d0510 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:1b fcid 0x9d0415 dynamic

  vsan 30 wwn 20:00:00:25:b5:1f:1a:19 fcid 0x9d040e dynamic

  vsan 30 wwn 52:4a:93:72:0d:21:6b:02 fcid 0x9d0600 dynamic

   vsan 30 wwn 52:4a:93:72:0d:21:6b:12 fcid 0x9d0700 dynamic

  vsan 30 wwn 52:4a:93:72:0d:21:6b:13 fcid 0x9d0800 dynamic
```

```
  vsan 30 wwn 20:00:00:25:b5:1f:1a:29 fcid 0x9d0513 dynamic
interface port-channel1
  channel mode active
  switchport rate-mode dedicated
vsan database
  vsan 30 interface fc1/1
  vsan 30 interface fc1/2
  vsan 30 interface fc1/3
  vsan 30 interface fc1/4
  vsan 30 interface fc1/7
  vsan 30 interface fc1/8
switchname MDS-B
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
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 fc1/19

interface fc1/20

interface fc1/21

interface fc1/22

interface fc1/23

interface fc1/24

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 fc1/43

interface fc1/44

interface fc1/45

interface fc1/46

interface fc1/47

interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
!Active Zone Database Section for vsan 30
zone name Infra-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:20
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
    member pwwn 20:00:00:25:b5:b1:1b:6f
    member pwwn 52:4a:93:72:0d:21:6b:01
    member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:1f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv3-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:03
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv4-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:05
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
```

```
zone name pure-vdi-serv5-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:07
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv6-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:09
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv7-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv8-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv9-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv10-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:11
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv11-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:13
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv12-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:15
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
```

```
zone name pure-vdi-serv13-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:17
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv16-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv17-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv18-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:21
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv19-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:23
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv20-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:25
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv21-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:27
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv24-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
```

```
zone name pure-vdi-serv25-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv26-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:31
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv27-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:33
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv28-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:35
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv29-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:37
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv30-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:39
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv15-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv23-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
```

```
zone name pure-vdi-serv22-fc1 vsan 30

member pwwn 20:00:00:25:b5:1f:1a:29

member pwwn 52:4a:93:72:0d:21:6b:01

member pwwn 52:4a:93:72:0d:21:6b:10

zoneset name pure-vdi-fab-B vsan 30

member Infra-serv1-fc1

member Infra-serv2-fc1

member Infra-serv3-fc1

member pure-vdi-serv1-fc1

member pure-vdi-serv2-fc1

member pure-vdi-serv3-fc1

member pure-vdi-serv4-fc1

member pure-vdi-serv5-fc1

member pure-vdi-serv6-fc1

member pure-vdi-serv7-fc1

member pure-vdi-serv8-fc1

member pure-vdi-serv9-fc1

member pure-vdi-serv10-fc1

member pure-vdi-serv11-fc1

member pure-vdi-serv12-fc1

member pure-vdi-serv13-fc1

member pure-vdi-serv16-fc1

member pure-vdi-serv17-fc1

member pure-vdi-serv18-fc1

member pure-vdi-serv19-fc1

member pure-vdi-serv20-fc1

member pure-vdi-serv21-fc1

member pure-vdi-serv24-fc1

member pure-vdi-serv25-fc1

member pure-vdi-serv26-fc1

member pure-vdi-serv27-fc1

member pure-vdi-serv28-fc1
member pure-vdi-serv29-fc1
member pure-vdi-serv30-fc1
member pure-vdi-serv15-fc1
member pure-vdi-serv23-fc1
member pure-vdi-serv22-fc1
zoneset activate name pure-vdi-fab-B vsan 30
do clear zone database vsan 30
!Full Zone Database Section for vsan 30
zone name Infra-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:20
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:6f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name vdi-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-fab-B vsan 30
zone name pure-vdi-serv7-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
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
interface fc1/6
port-license acquire
interface fc1/7
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/8
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/9
port-license acquir no shutdown
interface fc1/10
port-license acquire
no shutdown
interface fc1/11
port-license acquire
channel-group 1 force
no shutdown
interface fc1/12
port-license acquire
```

```
channel-group 1 force
no shutdown
interface fc1/13
port-license acquire
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
interface fc1/18
port-license acquire
interface fc1/19
port-license acquire
interface fc1/20
port-license acquire
interface fc1/21
port-license acquire
interface fc1/22
port-license acquire
interface fc1/23
port-license acquire
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
interface fc1/26
```

```
port-license acquire
interface fc1/27
port-license acquire
interface fc1/28
port-license acquire
interface fc1/29
port-license acquire
interface fc1/30
port-license acquire
interface fc1/31
port-license acquire
interface fc1/32
port-license acquire
interface fc1/33
port-license acquire
interface fc1/34
port-license acquire
interface fc1/35
port-license acquire
interface fc1/36
port-license acquire
interface fc1/37
port-license acquire
interface fc1/38
port-license acquire
interface fc1/39
port-license acquire
interface fc1/40
port-license acquire
interface fc1/41
port-license acquire
interface fc1/42
```

```
port-license acquire
interface fc1/43
port-license acquire
interface fc1/44
port-license acquire
interface fc1/45
port-license acquire
interface fc1/46
port-license acquire
interface fc1/47
port-license acquire
interface fc1/48
port-license acquire
interface mgmt0
ip address 10.29.164.128 255.255.255.0
ip default-gateway 10.29.164.1
MDS-B#
```

Appendix B - Pure Storage Configuration and Scripts

Two separate PowerCLI scripts were used during the setup and testing phases of this Cisco Validated Design. The first script's function was to apply all Pure Storage best practices and was used only once after all ESXi hosts had been built and placed into vCenter. The second script was exercised after the end of all Login VSI cluster tests so as to run the unmap command against all datastores in order to make certain we had accurate array space utilization and data reduction numbers recorded during each simulation.

In both instances, the only input parameters needed for the PowerCLI scripts were the vCenter instance name (IP address or FQDN) and the username and password of an administrative user over that vCenter instance.

Pure Storage ESXi Best Practice Script

Configuring the ESXi hosts in the environment for optimal use with the FlashArray//m50 was accomplished by running a single PowerShell script once after all ESXi hosts were built and datastores were provisioned. Additional datastores created later will have these best practices applied to them automatically. It is recommended to reboot each ESXi host after applying this script.

Future updates to this script as well as other useful PowerShell scripts built for use with Pure Storage can be found here:

https://github.com/codyhosterman/powercli

Below is the ESXi Best Practice PowerShell script used in this Cisco Validated Design:

```
#Enter the following required parameters. Log folder directory is just an example, change as needed.

#Put all entries inside the quotes:

#**********************************

$vcenter = ""

$vcuser = ""

$vcpass = ""

$logfolder = "C:\folder\folder\etc\"

#**********************************



<#

Optional parameters. Keep these values at default unless necessary and understood

For a different IO Operations limit beside the Pure Storage recommended value of 1, change $iopsvalue to another integer value 1-1000.
```

```
To skip changing host-wide settings for XCOPY Transfer Size and In-Guest UNMAP
change $hostwidesettings to $false

#>

$iopsvalue = 1

$hostwidesettings = $true



<#

*******Disclaimer:******************************************************

This scripts are offered "as is" with no warranty.  While this

scripts is tested and working in my environment, it is recommended that you
test

this script in a test lab before using in a production environment. Everyone
can

use the scripts/commands provided here without any written permission but I

will not be liable for any damage or loss to the system.

************************************************************************

This script will:

-Check for a SATP rule for Pure Storage FlashArrays

-Create a SATP rule for Round Robin and IO Operations Limit of 1 only for
FlashArrays

-Remove any incorrectly configured Pure Storage FlashArray rules

-Configure any existing devices properly (Pure Storage FlashArray devices only)

-Set VAAI XCOPY transfer size to 16MB

-Enable EnableBlockDelete on ESXi 6 hosts only

All change operations are logged to a file.

This can be run directly from PowerCLI or from a standard PowerShell prompt.
PowerCLI must be installed on the local host regardless.

Supports:

-FlashArray 400 Series and //m

-vCenter 5.0 and later

-PowerCLI 6.3 R1 or later required

For info, refer to www.codyhosterman.com

#>



#Create log folder if non-existent
```

```
If (!(Test-Path -Path $logfolder)) { New-Item -ItemType Directory -Path $log-
folder }

$logfile = $logfolder + (Get-Date -Format o |ForEach-Object {$_ -Replace ":",
"."}) + "setbestpractices.txt"

write-host "Checking and setting Pure Storage FlashArray Best Practices for
VMware on the ESXi hosts in this vCenter. No further information is printed to
the screen."

write-host "Script log information can be found at $logfile"



add-content $logfile '             __________________________'

add-content $logfile '            /+++++++++++++++++++++++++++\'

add-content $logfile '           /+++++++++++++++++++++++++++++\'

add-content $logfile '          /+++++++++++++++++++++++++++++++\'

add-content $logfile '         /+++++++++++++++++++++++++++++++++\'

add-content $logfile '        /+++++++++++++++++++++++++++++++++++\'

add-content $logfile '       /++++++++++++/-----------\++++++++++++\'

add-content $logfile '      /++++++++++++/             \++++++++++++\'

add-content $logfile '     /++++++++++++/               \++++++++++++\'

add-content $logfile '    /++++++++++++/                 \++++++++++++\'

add-content $logfile '   /++++++++++++/                   \++++++++++++\'

add-content $logfile '   \++++++++++++\                   /++++++++++++/'

add-content $logfile '    \++++++++++++\                 /++++++++++++/'

add-content $logfile '     \++++++++++++\               /++++++++++++/'

add-content $logfile '      \++++++++++++\             /++++++++++++/'

add-content $logfile '       \++++++++++++\           /++++++++++++/'

add-content $logfile '        \++++++++++++\'

add-content $logfile '         \++++++++++++\'

add-content $logfile '          \++++++++++++\'

add-content $logfile '           \++++++++++++\'

add-content $logfile '            \------------\'

add-content $logfile 'Pure Storage  FlashArray VMware ESXi Best Practices
Script v3.1'

add-content $logfile '----------------------------------------------------------
-------------------------------------------'
```

```
if ( !(Get-Module -Name VMware.VimAutomation.Core -ErrorAction SilentlyContin-
ue) ) {

. “C:\Program Files (x86)\VMware\Infrastructure\vSphere Power-
CLI\Scripts\Initialize-PowerCLIEnvironment.ps1” |out-null

}

set-powercliconfiguration -invalidcertificateaction "ignore" -confirm:$false
|out-null



if ((Get-PowerCLIVersion).build -lt 3737840)

{

       write-host "This version of PowerCLI is too old, version 6.3 Release 1 or
       later is required (Build 3737840)" -BackgroundColor Red

       write-host "Found the following build number:"

       write-host (Get-PowerCLIVersion).build

       write-host "Terminating Script" -BackgroundColor Red

       add-content $logfile "This version of PowerCLI is too old, version 6.3
       Release 1 or later is required (Build 3737840)"

       add-content $logfile "Found the following build number:"

       add-content $logfile (Get-PowerCLIVersion).build

       add-content $logfile "Terminating Script"

       return

}



try

{

       connect-viserver -Server $vcenter -username $vcuser -password $vcpass -
       ErrorAction Stop |out-null

}

catch

{

       write-host "Failed to connect to vCenter" -BackgroundColor Red

       write-host $Error

       write-host "Terminating Script" -BackgroundColor Red

       add-content $logfile "Failed to connect to vCenter"

       add-content $logfile $Error
```

```
        add-content $logfile "Terminating Script"

        return

}

add-content $logfile ('Connected to vCenter at ' + $vcenter)

add-content $logfile '----------------------------------------------------------------------------------------------------'



$hosts= get-vmhost



add-content $logfile "Iterating through all ESXi hosts..."



#Iterating through each host in the vCenter

foreach ($esx in $hosts)

{

        $esxcli=get-esxcli -VMHost $esx -v2

        add-content $logfile "-----------------------------------------------------------------------------------------------"

        add-content $logfile "-----------------------------------------------------------------------------------------------"

        add-content $logfile "Working on the following ESXi host:"

        add-content $logfile $esx.NetworkInfo.hostname

        add-content $logfile "-----------------------------------------------"

        if ($hostwidesettings -eq $true)

        {

        add-content $logfile "Checking host-wide setting for XCOPY and In-Guest UNMAP"

        $xfersize = $esx | Get-AdvancedSetting -Name DataMover.MaxHWTransferSize

        if ($xfersize.value -ne 16384)

        {

                add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this host is incorrect:"

                add-content $logfile $xfersize.value

                add-content $logfile "This should be set to 16386 (16 MB). Changing to 16384..."

                $xfersize |Set-AdvancedSetting -Value 16384 -Confirm:$false |out-null
                add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this
                host is now 16 MB"

        }

        else

        {

                add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this
                host is correct at 16 MB and will not be altered."

        }

        if ($esx.Version -like "6.0.*")

        {

                $enableblockdelete = ($esx | Get-AdvancedSetting -Name
                VMFS3.EnableBlockDelete).Value

                if ($enableblockdelete.Value -eq 0)

                {

                add-content $logfile "EnableBlockDelete is currently disabled. En-
                abling..."

                $enableblockdelete |Set-AdvancedSetting -Value 1 -Confirm:$false
                |out-null

                add-content $logfile "EnableBlockDelete has been set to enabled."

                }

                else

                {

                add-content $logfile "EnableBlockDelete for this host is correctly
                enabled and will not be altered."

                }

        }

        else

        {

                add-content $logfile "The current host is not version 6.0. Skip-
                ping EnableBlockDelete check."

        }

        }

        else

        {

        add-content $logfile "Not checking host wide settings for XCOPY and In-
        Guest UNMAP due to in-script override"
```

```
}

add-content $logfile "-----------------------------------------------"

$rules = $esxcli.storage.nmp.satp.rule.list.invoke() |where-object
{$_.Vendor -eq "PURE"}

$correctrule = 0

$iopsoption = "iops=" + $iopsvalue

if ($rules.Count -ge 1)

{

add-content $logfile "Found the following existing Pure Storage SATP
rules"

$rules | out-string | add-content $logfile

add-content $logfile "-----------------------------------------------"

foreach ($rule in $rules)

{

       add-content $logfile "--------------------------------------------
       ---"

       add-content $logfile "Checking the following existing rule:"

       $rule | out-string | add-content $logfile

       $issuecount = 0

       if ($rule.DefaultPSP -eq "VMW_PSP_RR")

       {

       add-content $logfile "The existing Pure Storage FlashArray rule is
       configured with the correct Path Selection Policy:"

       add-content $logfile $rule.DefaultPSP

       }

       else

       {

       add-content $logfile "The existing Pure Storage FlashArray rule is
       NOT configured with the correct Path Selection Policy:"

       add-content $logfile $rule.DefaultPSP

       add-content $logfile "The rule should be configured to Round Robin
       (VMW_PSP_RR)"

       $issuecount = 1

       }

       if ($rule.PSPOptions -eq $iopsoption)
```

```
{
add-content $logfile "The existing Pure Storage FlashArray rule is configured with the correct IO Operations Limit:"
add-content $logfile $rule.PSPOptions
}
else
{
add-content $logfile "The existing Pure Storage FlashArray rule is NOT configured with the correct IO Operations Limit:"
add-content $logfile $rule.PSPOptions
add-content $logfile "The rule should be configured to an IO Operations Limit of $iopsvalue"
$issuecount = $issuecount + 1
}
if ($rule.Model -eq "FlashArray")
{
add-content $logfile "The existing Pure Storage FlashArray rule is configured with the correct model:"
add-content $logfile $rule.Model
}
else
{
add-content $logfile "The existing Pure Storage FlashArray rule is NOT configured with the correct model:"
add-content $logfile $rule.Model
add-content $logfile "The rule should be configured with the model of FlashArray"
$issuecount = $issuecount + 1
}
if ($issuecount -ge 1)
{
$satpArgs = $esxcli.storage.nmp.satp.rule.remove.createArgs()
$satpArgs.model = $rule.Model
$satpArgs.vendor = "PURE"
$satpArgs.satp = $rule.Name
```

```
       $satpArgs.psp = $rule.DefaultPSP
       $satpArgs.pspoption = $rule.PSPOptions
       add-content $logfile "This rule is incorrect, deleting..."
       $esxcli.storage.nmp.satp.rule.remove.invoke($satpArgs)
       add-content $logfile "*****NOTE: Deleted the rule.*****"
       add-content $logfile "-----------------------------------------------"
       }
       else
       {
       add-content $logfile "This rule is correct"
       add-content $logfile "-----------------------------------------------"
       $correctrule = 1
       }
}
}
if ($correctrule -eq 0)
{
add-content $logfile "No correct SATP rule for the Pure Storage FlashArray is found. Creating a new rule to set Round Robin and an IO Operations Limit of $iopsvalue"
$satpArgs = $esxcli.storage.nmp.satp.rule.remove.createArgs()
$satpArgs.description = "Pure Storage FlashArray SATP Rule"
$satpArgs.model = "FlashArray"
$satpArgs.vendor = "PURE"
$satpArgs.satp = "VMW_SATP_ALUA"
$satpArgs.psp = "VMW_PSP_RR"
$satpArgs.pspoption = $iopsoption
$result = $esxcli.storage.nmp.satp.rule.add.invoke($satpArgs)
if ($result -eq $true)
{
       add-content $logfile "New rule created:"
       $newrule = $esxcli.storage.nmp.satp.rule.list.invoke() |where-object {$_.Vendor -eq "PURE"}
```

```
    $newrule | out-string | add-content $logfile

}

else

{

    add-content $logfile "ERROR: The rule failed to create. Manual intervention might be required."

}

}

else

{

add-content $logfile "A correct SATP rule for the FlashArray exists. No need to create a new one on this host."

}

$devices = $esx |Get-ScsiLun -CanonicalName "naa.624a9370*"

add-content $logfile "------------------------------------------------"

if ($devices.count -ge 1)

{

add-content $logfile "Looking for existing Pure Storage volumes on this host"

add-content $logfile "Found the following number of existing Pure Storage volumes on this host."

add-content $logfile $devices.count

add-content $logfile "Checking and fixing their multipathing configuration now."

add-content $logfile "------------------------------------------------"

foreach ($device in $devices)

{

    add-content $logfile "Found and examining the following FlashArray device:"

    add-content $logfile $device.CanonicalName

    if ($device.MultipathPolicy -ne "RoundRobin")

    {

    add-content $logfile "This device does not have the correct Path Selection Policy, it is set to:"

    add-content $logfile $device.MultipathPolicy

    add-content $logfile "Changing to Round Robin."
```

```
      Get-VMhost $esx |Get-ScsiLun $device |Set-ScsiLun -MultipathPolicy
      RoundRobin |out-null

      }

      else

      {

      add-content $logfile "This device's Path Selection Policy is cor-
      rectly set to Round Robin. No need to change."

      }

      $deviceargs = $esx-
      cli.storage.nmp.psp.roundrobin.deviceconfig.get.createargs()

      $deviceargs.device = $device.CanonicalName

      $deviceconfig = $esx-
      cli.storage.nmp.psp.roundrobin.deviceconfig.get.invoke($deviceargs
      )

      $nmpargs =  $esx-
      cli.storage.nmp.psp.roundrobin.deviceconfig.set.createargs()

      $nmpargs.iops = $iopsvalue

      $nmpargs.type = "iops"

      if ($deviceconfig.IOOperationLimit -ne $iopsvalue)

      {

      add-content $logfile "The current IO Operation limit for this de-
      vice is:"

      add-content $logfile $deviceconfig.IOOperationLimit

      add-content $logfile "This device's IO Operation Limit is unset or
      is not set to the value of $iopsvalue. Changing..."

      $nmpargs.device = $device.CanonicalName

      $esx-
      cli.storage.nmp.psp.roundrobin.deviceconfig.set.invoke($nmpargs)
      |out-null

      }

      else

      {

      add-content $logfile "This device's IO Operation Limit matches the
      value of $iopsvalue. No need to change."

      }

      add-content $logfile "--------------------"

}

}
```

```
                else

                {

                add-content $logfile "No existing Pure Storage volumes found on this
                host."

                }

        }

        disconnect-viserver -Server $vcenter -confirm:$false

        add-content $logfile "Disconnected vCenter connection"
```

Pure Storage UNMAP PowerCLI Script

Dead space reclamation via a SCSI unmap command is a necessary occasional operation on an all Flash Array so that the array is able to convert that dead space to usable space and be able to write to it as well as provide accurate space utilization and data reduction metrics.

In most VDI environments, we recommend running this script on either a weekly or monthly basis depending on how often and if virtual desktops are being created and destroyed upon user logoff. If non-persistent desktops are being created and deleted often, running this script on a weekly schedule is recommended during off-hours.

This script can be found at the below location:

https://github.com/codyhosterman/powercli/blob/master/unmapsdk.ps1

More information on unmap can be found here:

http://www.codyhosterman.com/category/vmware/vaai/unmap/

The following script is what was used at the conclusion of every simulation in the Cisco Validated Design.

```
#************************************************************************************************
***

#VMWARE POWERCLI AND PURE STORAGE POWERSHELL SDK MUST BE INSTALLED ON THE MACHINE THIS IS RUNNING
ON

#************************************************************************************************
***

#

#For info, refer to www.codyhosterman.com

#

#*****************************************************************

#Enter the following parameters. Put all entries inside the quotes.

#One or more FlashArrays are supported. Remove/add additional ,''s for more/less arrays.
```

```
#Remove '<array IP or FQDN>' and replace that entire string with a FlashArray IP or FQDN like
'192.168.0.10'. Separate each array by a comma.

#*****************************************************************

$vcenter = ""

$vcuser = ""

$vcpass = ""

$flasharrays = @('<array IP or FQDN>','<array IP or FQDN>')

$pureuser = ""

$pureuserpwd = ""

$logfolder = 'C:\folder\folder\etc\'



#Optional settings. Leave equal to $null if not needed. Otherwise add a IP or FQDN inside of dou-
ble quotes and a UUID for the agentID.

$loginsightserver = $null

$loginsightagentID = ""



<#

*******Disclaimer:******************************************************

This scripts are offered "as is" with no warranty.  While this

scripts is tested and working in my environment, it is recommended that you test

this script in a test lab before using in a production environment. Everyone can

use the scripts/commands provided here without any written permission but I

will not be liable for any damage or loss to the system.

************************************************************************

This script will identify Pure Storage FlashArray volumes and issue UNMAP against them. The
script uses the best practice

recommendation block count of 1% of the free capacity of the datastore. All operations are logged
to a file and also

output to the screen. REST API calls to the array before and after UNMAP will report on how much
(if any) space has been reclaimed.

This can be run directly from PowerCLI or from a standard PowerShell prompt. PowerCLI must be
installed on the local host regardless.

Supports:

-PowerShell 3.0 or later

-Pure Storage PowerShell SDK 1.5 or later
```

```
-PowerCLI 6.3 Release 1

-REST API 1.4 and later

-Purity 4.1 and later

-FlashArray 400 Series and //m

-vCenter 5.5 and later

-Each FlashArray datastore must be present to at least one ESXi version 5.5 or later host or it
will not be reclaimed

#>

#Create log folder if non-existent

If (!(Test-Path -Path $logfolder)) { New-Item -ItemType Directory -Path $logfolder }

$logfile = $logfolder + (Get-Date -Format o |ForEach-Object {$_ -Replace ':', '.'}) + "unmap.txt"


add-content $logfile '             __________________________'

add-content $logfile '            /++++++++++++++++++++++++++\'

add-content $logfile '           /++++++++++++++++++++++++++++\'

add-content $logfile '          /++++++++++++++++++++++++++++++\'

add-content $logfile '         /++++++++++++++++++++++++++++++++\'

add-content $logfile '        /++++++++++++++++++++++++++++++++++\'

add-content $logfile '       /++++++++++++/----------\++++++++++++\'

add-content $logfile '      /++++++++++++/            \++++++++++++\'

add-content $logfile '     /++++++++++++/              \++++++++++++\'

add-content $logfile '    /++++++++++++/                \++++++++++++\'

add-content $logfile '   /++++++++++++/                  \++++++++++++\'

add-content $logfile '   \++++++++++++\                  /++++++++++++/'

add-content $logfile '    \++++++++++++\                /++++++++++++/'

add-content $logfile '     \++++++++++++\              /++++++++++++/'

add-content $logfile '      \++++++++++++\            /++++++++++++/'

add-content $logfile '       \++++++++++++\          /++++++++++++/'

add-content $logfile '        \++++++++++++\'

add-content $logfile '         \++++++++++++\'

add-content $logfile '          \++++++++++++\'

add-content $logfile '           \++++++++++++\'
```

```
add-content $logfile '            \------------\'
add-content $logfile 'Pure Storage VMware ESXi UNMAP Script v4.0'
add-content $logfile '---------------------------------------------------------------------------
-------------------------'

#A function to make REST Calls to Log Insight
function logInsightRestCall
{
    $restvmfs = [ordered]@{
                    name = "Datastore"
                    content = $datastore.Name
                    }
    $restarray = [ordered]@{
                    name = "FlashArray"
                    content = $endpoint[$arraychoice].endpoint
                    }
    $restvol = [ordered]@{
                    name = "FlashArrayvol"
                    content = $purevol.name
                    }
    $restunmap = [ordered]@{
                    name = "ReclaimedSpace"
                    content = $unmapsavings
                    }
    $esxhost = [ordered]@{
                    name = "ESXihost"
                    content = $esx
                    }
    $devicenaa = [ordered]@{
                    name = "SCSINaa"
                    content = $lun
                    }
    $fields = @($restvmfs,$restarray,$restvol,$restunmap,$esxhost,$devicenaa)

    $restcall = @{

                 messages =    ([Object[]]($messages = [ordered]@{

                        text = ("Completed an UNMAP operation on the VMFS volume named " + $datastore.Name + " that is on the FlashArray named " + $endpoint[$arraychoice].endpoint + ".")

                        fields = ([Object[]]$fields)

                        }))

                } |convertto-json -Depth 4

    $resturl = ("http://" + $loginsightserver + ":9000/api/v1/messages/ingest/" + $loginsightagentID)

    add-content $logfile ""

    add-content $logfile ("Posting results to Log Insight server: " + $loginsightserver)

    try

    {

        $response = Invoke-RestMethod $resturl -Method Post -Body $restcall -ContentType 'application/json' -ErrorAction stop

        add-content $logfile "REST Call to Log Insight server successful"

        $response| out-string |add-content $logfile

    }

    catch

    {

        add-content $logfile "REST Call failed to Log Insight server"

        add-content $logfile $error[0]

        add-content $logfile $resturl

    }

}


#Connect to FlashArray via REST

$facount=0

$purevolumes=@()

$purevol=$null

$EndPoint= @()

$Pwd = ConvertTo-SecureString $pureuserpwd -AsPlainText -Force
```

```
$Creds = New-Object System.Management.Automation.PSCredential ($pureuser, $pwd)

write-host "Script information can be found at $logfile" -ForegroundColor Green



<#

Connect to FlashArray via REST with the SDK

Creates an array of connections for as many FlashArrays as you have entered into the $flasharrays variable.

Assumes the same credentials are in use for every FlashArray

#>



foreach ($flasharray in $flasharrays)

{

    if ($facount -eq 0)

    {

        try

        {

            $EndPoint += (New-PfaArray -EndPoint $flasharray -Credentials $Creds -IgnoreCertificateError -ErrorAction stop)

        }

        catch

        {

            write-host ("Connection to FlashArray " + $flasharray + " failed. Please check credentials or IP/FQDN") -BackgroundColor Red

            write-host $Error[0]

            write-host "Terminating Script" -BackgroundColor Red

            add-content $logfile ("Connection to FlashArray " + $flasharray + " failed. Please check credentials or IP/FQDN")

            add-content $logfile $Error[0]

            add-content $logfile "Terminating Script"

            return

        }

        $purevolumes += Get-PfaVolumes -Array $EndPoint[$facount]

        $tempvols = @(Get-PfaVolumes -Array $EndPoint[$facount])

        $arraysnlist = @($tempvols.serial[0].substring(0,16))
    }

    else

    {

        try

        {

            $EndPoint += New-PfaArray -EndPoint $flasharray -Credentials $Creds -
IgnoreCertificateError

            $purevolumes += Get-PfaVolumes -Array $EndPoint[$facount]

            $tempvols = Get-PfaVolumes -Array $EndPoint[$facount]

            $arraysnlist += $tempvols.serial[0].substring(0,16)

        }

        catch

        {

            write-host ("Connection to FlashArray " + $flasharray + " failed. Please check cre-
dentials or IP/FQDN") -BackgroundColor Red

            write-host $Error[0]

            add-content $logfile ("Connection to FlashArray " + $flasharray + " failed. Please
check credentials or IP/FQDN")

            add-content $logfile $Error[0]

            return

        }

    }

    $facount = $facount + 1

}


add-content $logfile 'Connected to the following FlashArray(s):'

add-content $logfile $flasharrays

add-content $logfile '-----------------------------------------------------------------------------
-------------------------'


#Important PowerCLI if not done and connect to vCenter.


if ( !(Get-Module -Name VMware.VimAutomation.Core -ErrorAction SilentlyContinue) ) {
```

```
. "C:\Program Files (x86)\VMware\Infrastructure\vSphere PowerCLI\Scripts\Initialize-
PowerCLIEnvironment.ps1"

}

Set-PowerCLIConfiguration -invalidcertificateaction 'ignore' -confirm:$false |out-null

Set-PowerCLIConfiguration -Scope Session -WebOperationTimeoutSeconds -1 -confirm:$false |out-null

if ((Get-PowerCLIVersion).build -lt 3737840)

{

    write-host "This version of PowerCLI is too old, version 6.3 Release 1 or later is required
(Build 3737840)" -BackgroundColor Red

    write-host "Found the following build number:"

    write-host (Get-PowerCLIVersion).build

    write-host "Terminating Script" -BackgroundColor Red

    write-host "Get it here: https://my.vmware.com/group/vmware/get-
download?downloadGroup=PCLI630R1"

    add-content $logfile "This version of PowerCLI is too old, version 6.3 Release 1 or later is
required (Build 3737840)"

    add-content $logfile "Found the following build number:"

    add-content $logfile (Get-PowerCLIVersion).build

    add-content $logfile "Terminating Script"

    add-content $logfile "Get it here: https://my.vmware.com/group/vmware/get-
download?downloadGroup=PCLI630R1"

    return

}



try

{

    connect-viserver -Server $vcenter -username $vcuser -password $vcpass -ErrorAction Stop |out-
null

}

catch

{

    write-host "Failed to connect to vCenter" -BackgroundColor Red

    write-host $vcenter

    write-host $Error[0]

    write-host "Terminating Script" -BackgroundColor Red
    add-content $logfile "Failed to connect to vCenter"
    add-content $logfile $vcenter
    add-content $logfile $Error[0]
    add-content $logfile "Terminating Script"
    return
}

write-host "No further information is printed to the screen."
add-content $logfile ('Connected to vCenter at ' + $vcenter)
add-content $logfile '---------------------------------------------------------------------------
-------------------------'

#Gather VMFS Datastores and identify how many are Pure Storage volumes
$datastores = get-datastore
add-content $logfile 'Found the following datastores:'
add-content $logfile $datastores
add-content $logfile '---------------------------------------------------------------------------
-------------------------'

#Starting UNMAP Process on datastores
$volcount=0
$purevol = $null
$totalspacereclaimed = 0
foreach ($datastore in $datastores)
{
    add-content $logfile (get-date)
    add-content $logfile ('The datastore named ' + $datastore + ' is being examined')
    $esx = $datastore | get-vmhost | where-object {($_.version -like '5.5.*') -or ($_.version -
like '6.0.*')}| where-object {($_.ConnectionState -eq 'Connected')} |Select-Object -last 1
    if ($datastore.Type -ne 'VMFS')
    {
        add-content $logfile ('This volume is not a VMFS volume, it is of type ' + $datas-
tore.Type + ' and cannot be reclaimed. Skipping...')
        add-content $logfile ''
        add-content $logfile '-------------------------------------------------------------------
---------------------------------'

    }

    elseif ($esx.count -eq 0)

    {

        add-content $logfile ('This datastore has no 5.5 or later hosts to run UNMAP from. Skip-
ping...')

        add-content $logfile ''

        add-content $logfile '-------------------------------------------------------------------
---------------------------------'



    }

    else

    {

        $lun = $datastore.ExtensionData.Info.Vmfs.Extent.DiskName |select-object -unique

        if ($lun.count -eq 1)

        {

            add-content $logfile "The UUID for this volume is:"

            add-content $logfile ($datastore.ExtensionData.Info.Vmfs.Extent.DiskName)

            $esxcli=get-esxcli -VMHost $esx -v2

            if ($lun -like 'naa.624a9370*')

            {

                $volserial = ($lun.ToUpper()).substring(12)

                $purevol = $purevolumes | where-object { $_.serial -eq $volserial }

                if ($purevol.name -eq $null)

                {

                   add-content $logfile 'ERROR: This volume has not been found. Please make sure
that all of the FlashArrays presented to this vCenter are entered into this script.'

                   add-content $logfile ''

                   add-content $logfile '--------------------------------------------------------
--------------------------------------------'



                }

                else
                {
                    for($i=0; $i -lt $arraysnlist.count; $i++)
                    {
                        if ($arraysnlist[$i] -eq ($volserial.substring(0,16)))
                        {
                            $arraychoice = $i
                        }
                    }
                    add-content $logfile ('The volume is on the FlashArray ' + $end-
point[$arraychoice].endpoint)
                    add-content $logfile ('This datastore is a Pure Storage volume named ' +
$purevol.name)
                    add-content $logfile ''
                    add-content $logfile ('The ESXi named ' + $esx + ' will run the UNMAP/reclaim
operation')
                    add-content $logfile ''
                    $volinfo = Get-PfaVolumeSpaceMetrics -Array $EndPoint[$arraychoice] -
VolumeName $purevol.name
                    $volreduction = '{0:N3}' -f ($volinfo.data_reduction)
                    $volphysicalcapacity = '{0:N3}' -f ($volinfo.volumes/1024/1024)
                    add-content $logfile ''
                    add-content $logfile ('The current data reduction for this volume before UN-
MAP is ' + $volreduction + " to 1")
                    add-content $logfile ('The  physical space consumption in MB of this device
before UNMAP is ' + $volphysicalcapacity)
                    add-content $logfile ''
                    #Calculating optimal block count. If VMFS is 75% full or more the count must
be 200 MB only. Ideal block count is 1% of free space of the VMFS in MB
                    if ((1 - $datastore.FreeSpaceMB/$datastore.CapacityMB) -ge .75)
                    {
                        $blockcount = 200
                        add-content $logfile 'The volume is 75% or more full so the block count
is overridden to 200 MB. This will slow down the reclaim dramatically'
                        add-content $logfile 'It is recommended to either free up space on the
volume or increase the capacity so it is less than 75% full'
                        add-content $logfile ("The block count in MB will be " + $blockcount)
                    }
                    else
                    {
                        $blockcount = [math]::floor($datastore.FreeSpaceMB * .01)
                        add-content $logfile ("The maximum allowed block count for this datastore is " + $blockcount)
                    }
                    $unmapargs = $esxcli.storage.vmfs.unmap.createargs()
                    $unmapargs.volumelabel = $datastore.Name
                    $unmapargs.reclaimunit = $blockcount
                    try
                    {
                        $esxcli.storage.vmfs.unmap.invoke($unmapargs) |out-null
                    }
                    catch
                    {
                        add-content $logfile "Failed to run UNMAP to this volume. Most common cause is the device is locked by another process."
                        add-content $logfile $Error[0]
                        add-content $logfile "Skipping volume..."
                    }
                    Start-Sleep -s 10
                    $volinfo = Get-PfaVolumeSpaceMetrics -Array $EndPoint[$arraychoice] -VolumeName $purevol.name
                    $volreduction = '{0:N3}' -f ($volinfo.data_reduction)
                    $volphysicalcapacitynew = '{0:N3}' -f ($volinfo.volumes/1024/1024)
                    $unmapsavings = [math]::Round(($volphysicalcapacity - $volphysicalcapacitynew),2)
                    $volcount=$volcount+1
                    add-content $logfile ''
                    add-content $logfile ('The new data reduction for this volume after UNMAP is ' + $volreduction + " to 1")
                    add-content $logfile ('The new physical space consumption in MB of this device after UNMAP is ' + $volphysicalcapacitynew)
                    add-content $logfile ("$unmapsavings" + ' in MB has been reclaimed from the
FlashArray from this volume')
                    $totalspacereclaimed = $totalspacereclaimed + $unmapsavings
                    if ($loginsightserver -ne $null){logInsightRestCall}
                    add-content $logfile ''
                    add-content $logfile '--------------------------------------------------------
---------------------------------------------'
                    Start-Sleep -s 5
                }
            }
            else
            {
                add-content $logfile ('The volume is not a FlashArray device, skipping the UNMAP
operation')
                add-content $logfile ''
                add-content $logfile '------------------------------------------------------------
-----------------------------------------'
            }
        }
        elseif ($lun.count -gt 1)
            {
                add-content $logfile ('The volume spans more than one SCSI device, skipping UNMAP
operation')
                add-content $logfile ''
                add-content $logfile '------------------------------------------------------------
-----------------------------------------'
            }
    }
}
add-content $logfile ("Space reclaim operation for all volumes is complete. Total immediate space
reclaimed is " + $totalspacereclaimed + " MB")
add-content $logfile "Note that more space will likely reclaim over time."
add-content $logfile ""
#disconnecting sessions
add-content $logfile ("Disconnecting vCenter and FlashArray sessions")
```

```
disconnect-viserver -Server $vcenter -confirm:$false

foreach ($flasharray in $endpoint)

{

    Disconnect-PfaArray -Array $flasharray
```

Appendix C - Pure Storage FlashArray//m50 Expanded Test Results

This section highlights and provide analysis of the Pure Storage FlashArray//m50 performance results for each of the cluster test cases identified in the Cisco Validated Design.

From a storage perspective, it is critical to maintain a latency of near to or less than a millisecond in order to guarantee a good end-user experience. As we will see, Pure Storage delivers that level of performance despite driving a substantial amount of IOPs and bandwidth for the thousands of desktops hosted on the single FlashArray//m50.

The following charts were compiled from extracting front-end array telemetry data from the storage logs and are equivalent to values shown in the Pure GUI. Results were plotted in this format to highlight individual storage performance metrics of interest during each simulation as well as clearly show the various phases of each simulation. Please note that across the top of each graph we have identified and broken up the Login VSI simulation into the three separate phases of the simulation run. The first phase (green arrows and text box) is the 2880 second Login VSI simulation phase when all sessions are ramping up and logging in. Next, the all sessions in the simulation steady-states for 600 seconds which is denoted by the yellow arrows and text box, and finally the black arrow to the right shows the end of the simulation when users begin logging out of the environment. For brevity, we generally did not show the entire logout operation as array activity is minimal during that time and the Login VSI simulation had completed.

Front-end statistics were pulled off of the Pure Storage array and plotted in the following appendix to show a more detailed summary of how the key array metrics performed during each cluster-level simulation. As in the results section, individual stages of each simulation are clearly denoted in each chart to provide more understanding of the array’s behavior during each test.

Simulation 1: 2600 Windows 10 x64 RDS XenApp Sessions

Figure 219 Latency

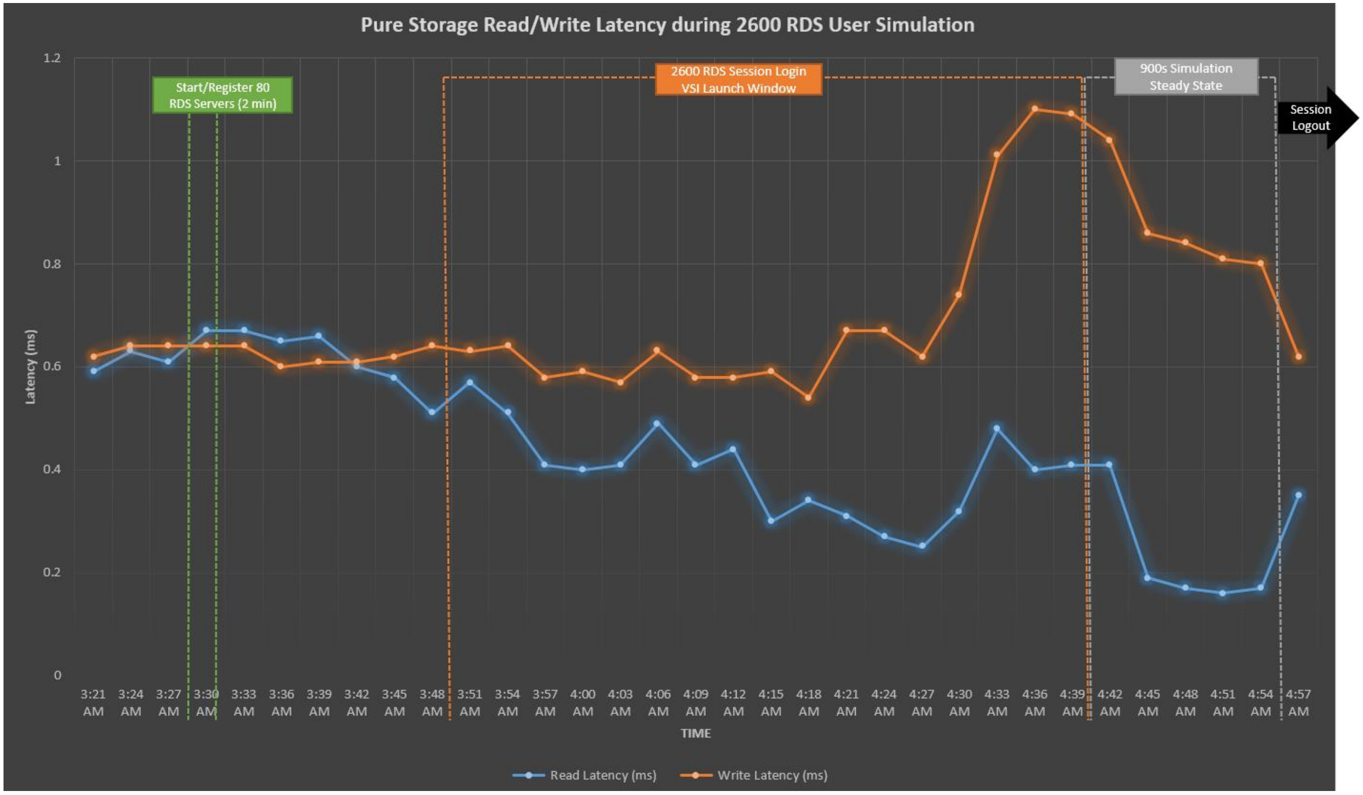

Figure 220 IOPS

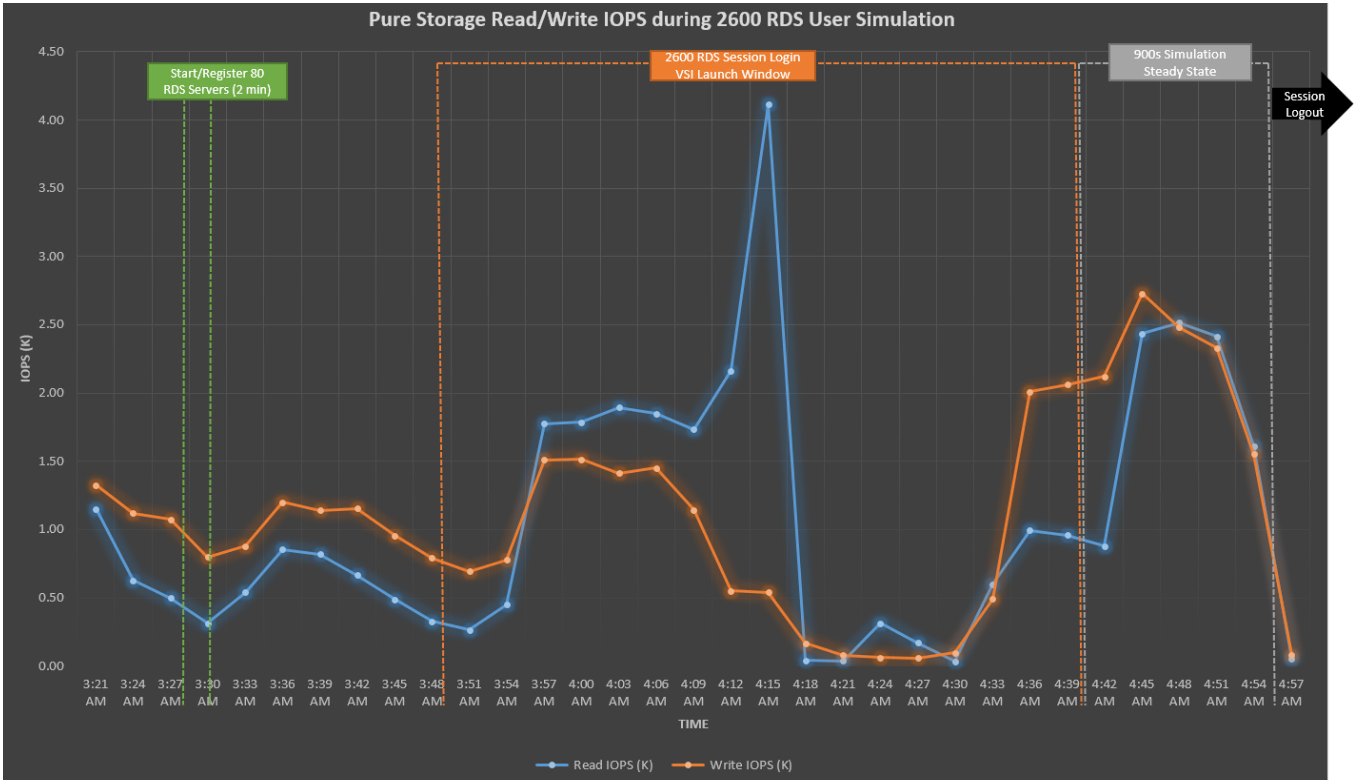

Figure 221 Bandwidth

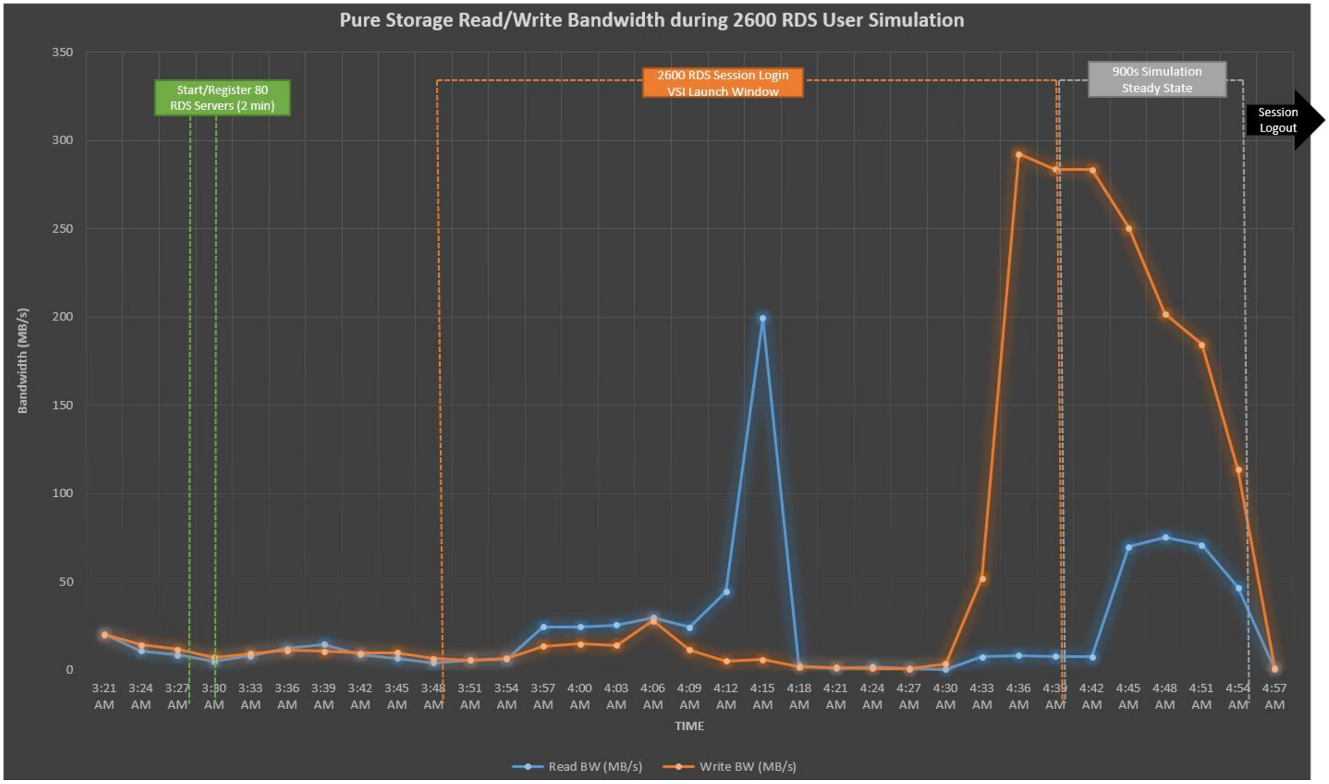

Simulation 2: 1200 Windows 10 x64 Non-Persistent XenDesktop PVS Sessions

Figure 222 Latency

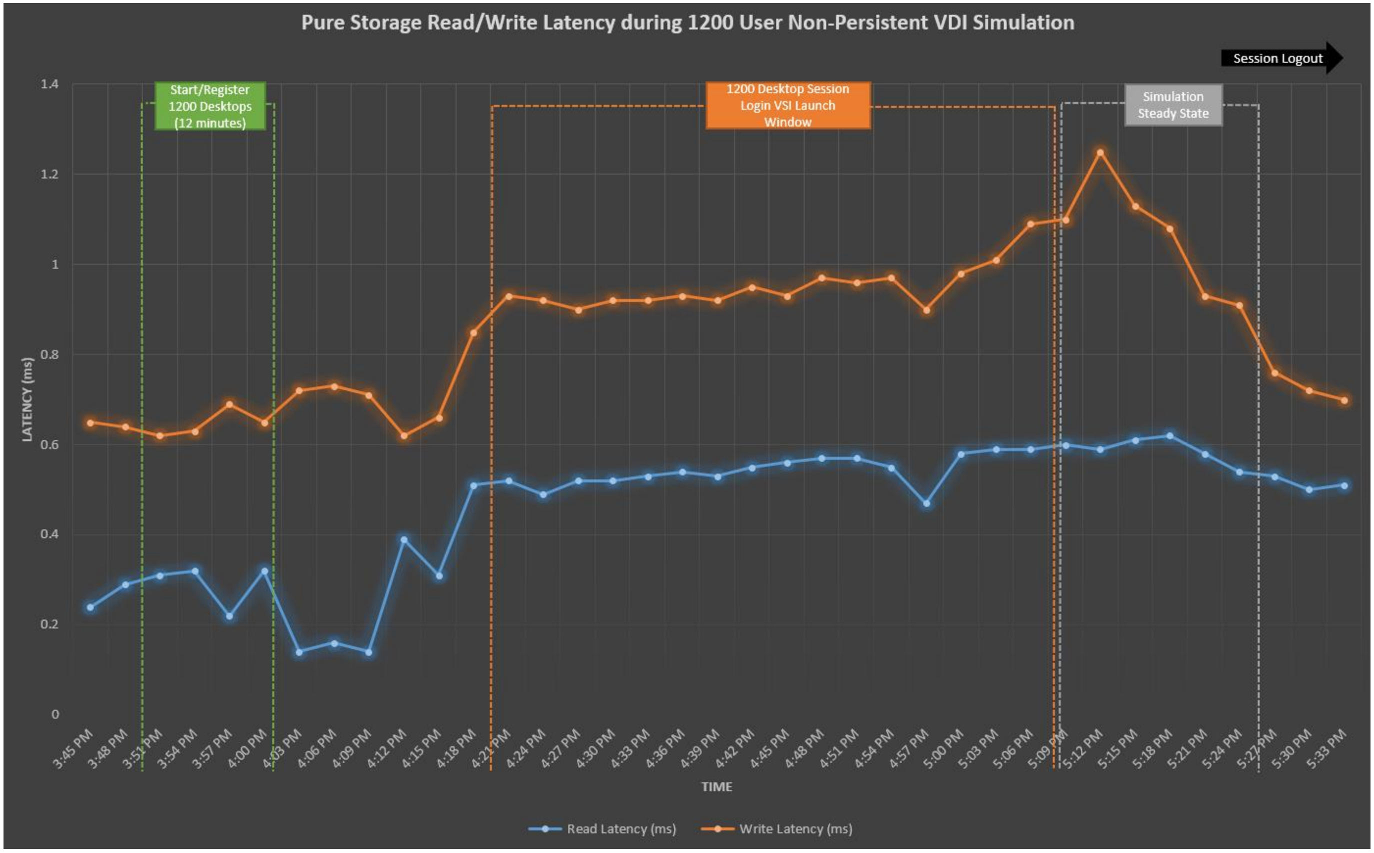

Figure 223 IOPS

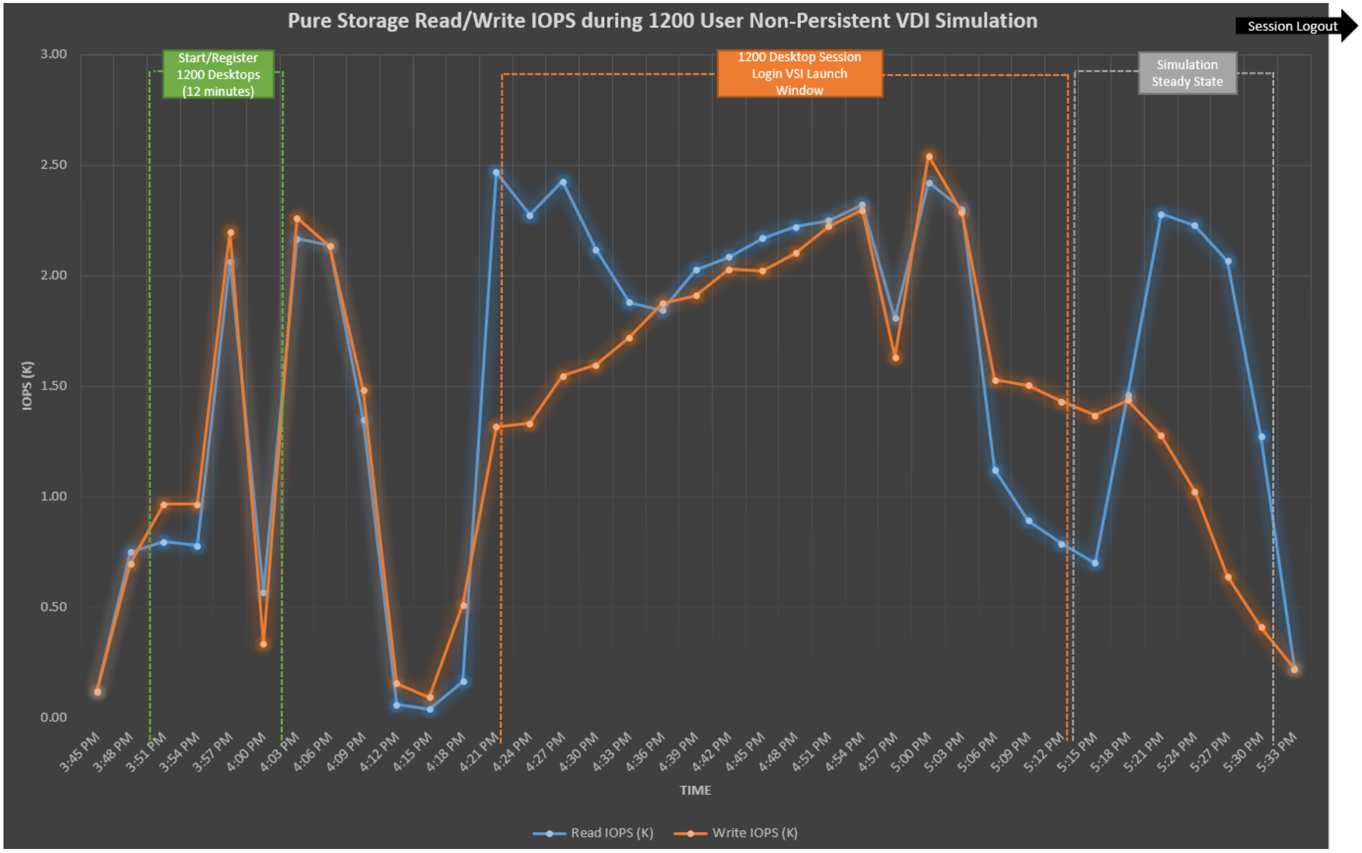

Figure 224 Bandwidth

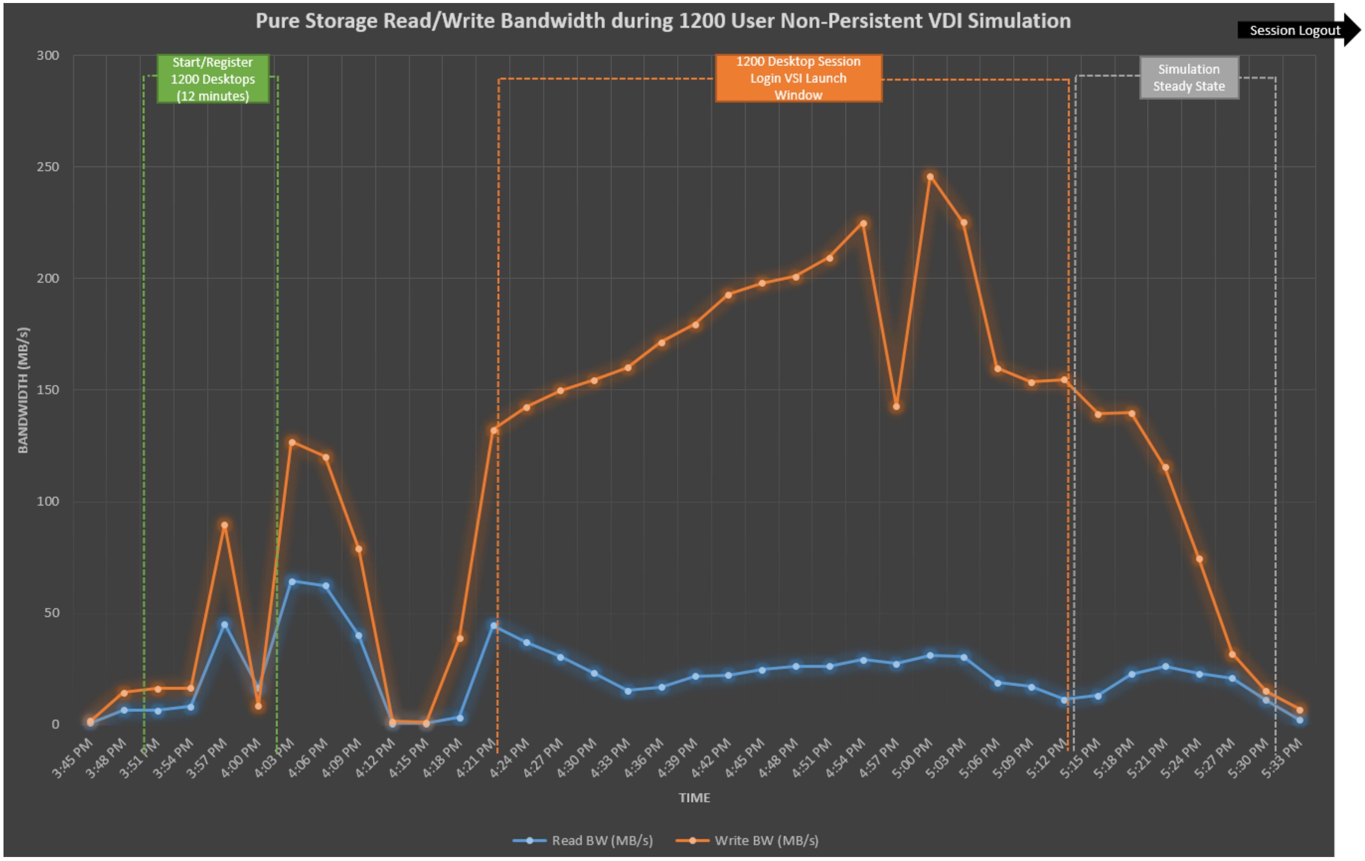

1200 Windows 10 x64 Persistent XenDesktop PVS Sessions

Figure 225 Latency

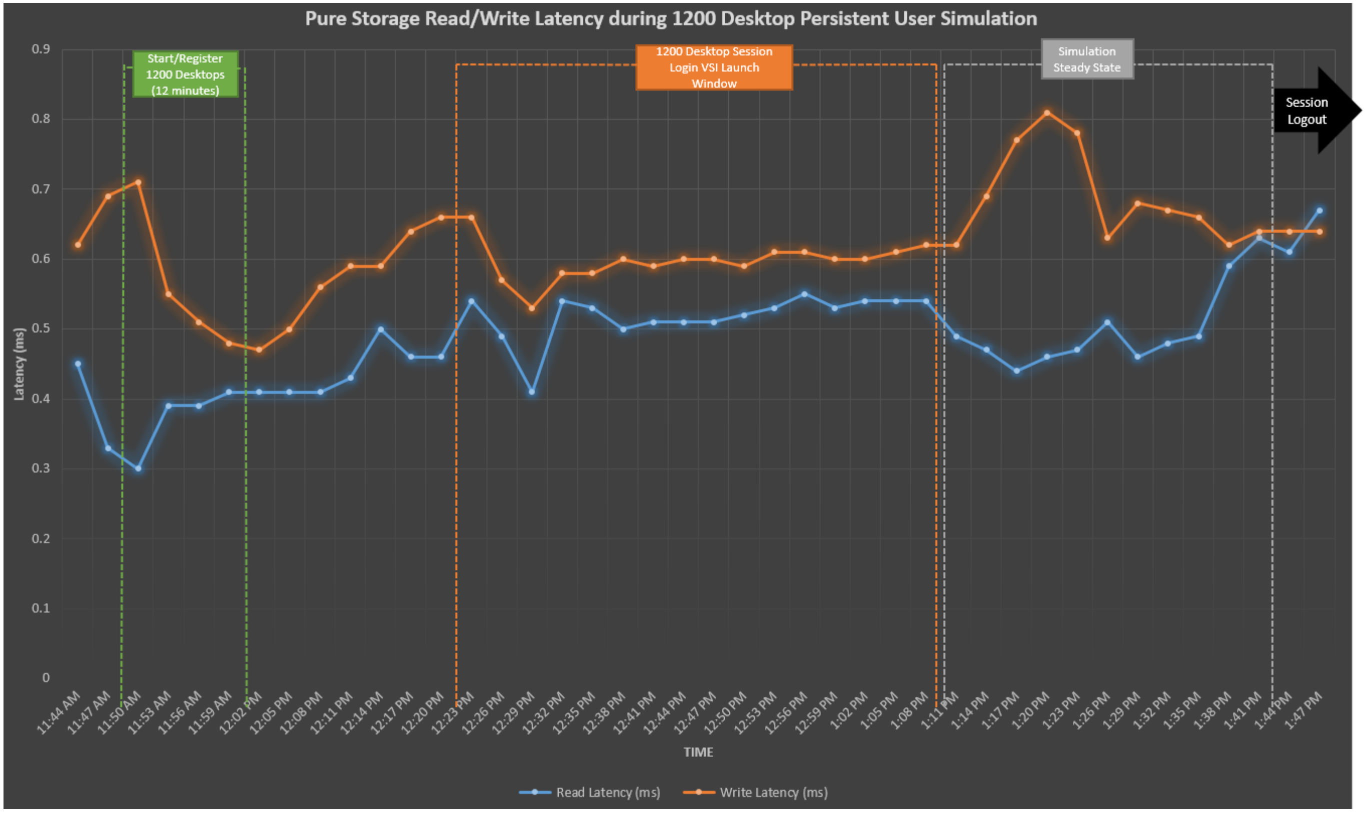

Figure 226 IOPS

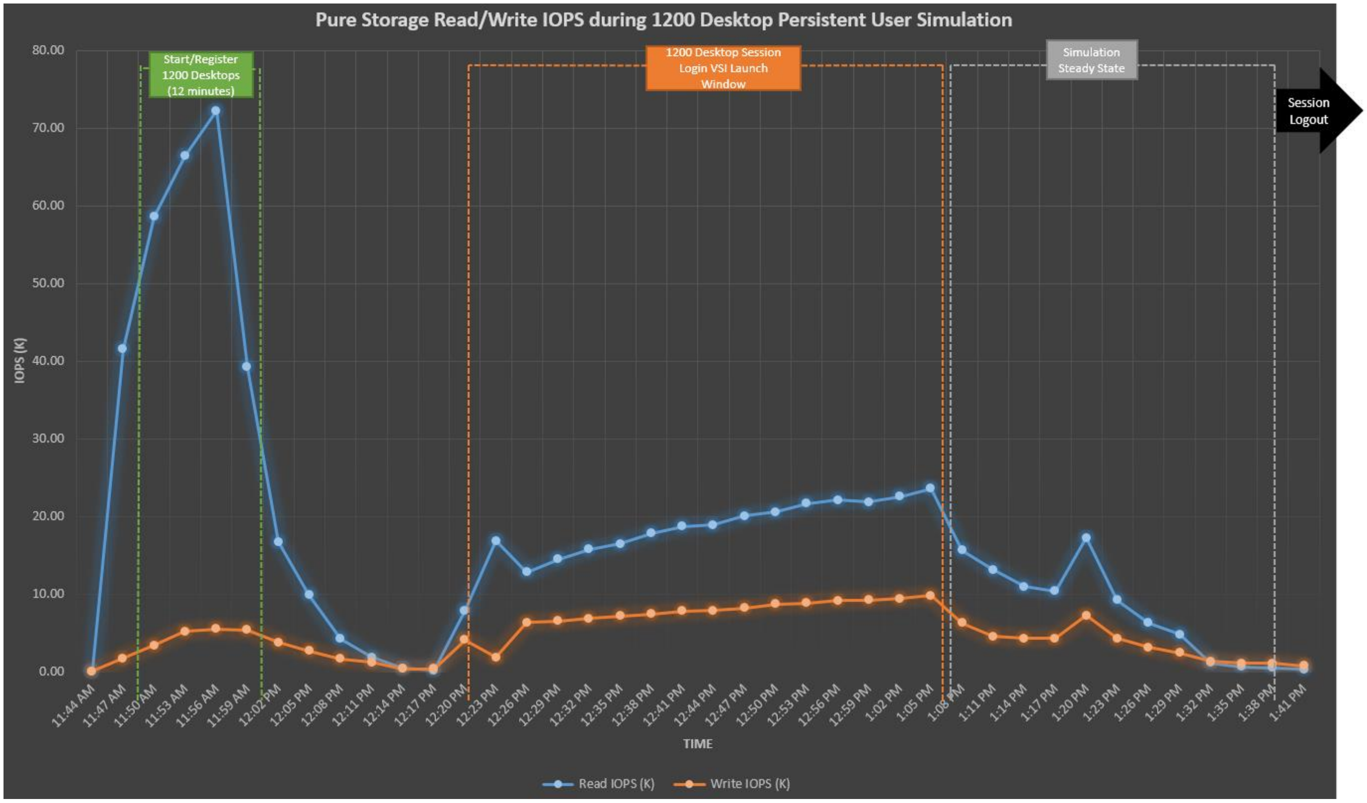

Figure 227 Bandwidth

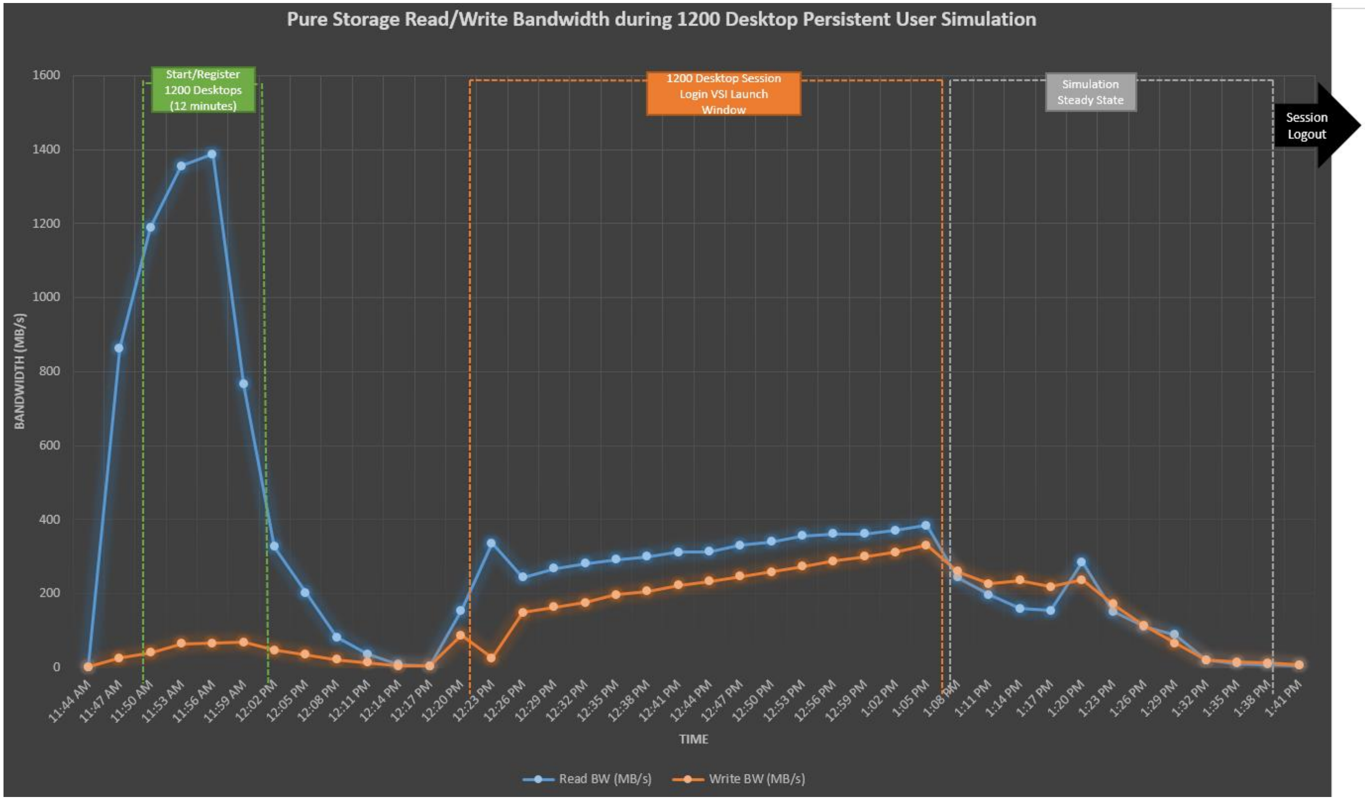

5000 Mixed Workload XenApp and XenDesktop Sessions (Combination of Scenarios 1-3)

Figure 228 Latency

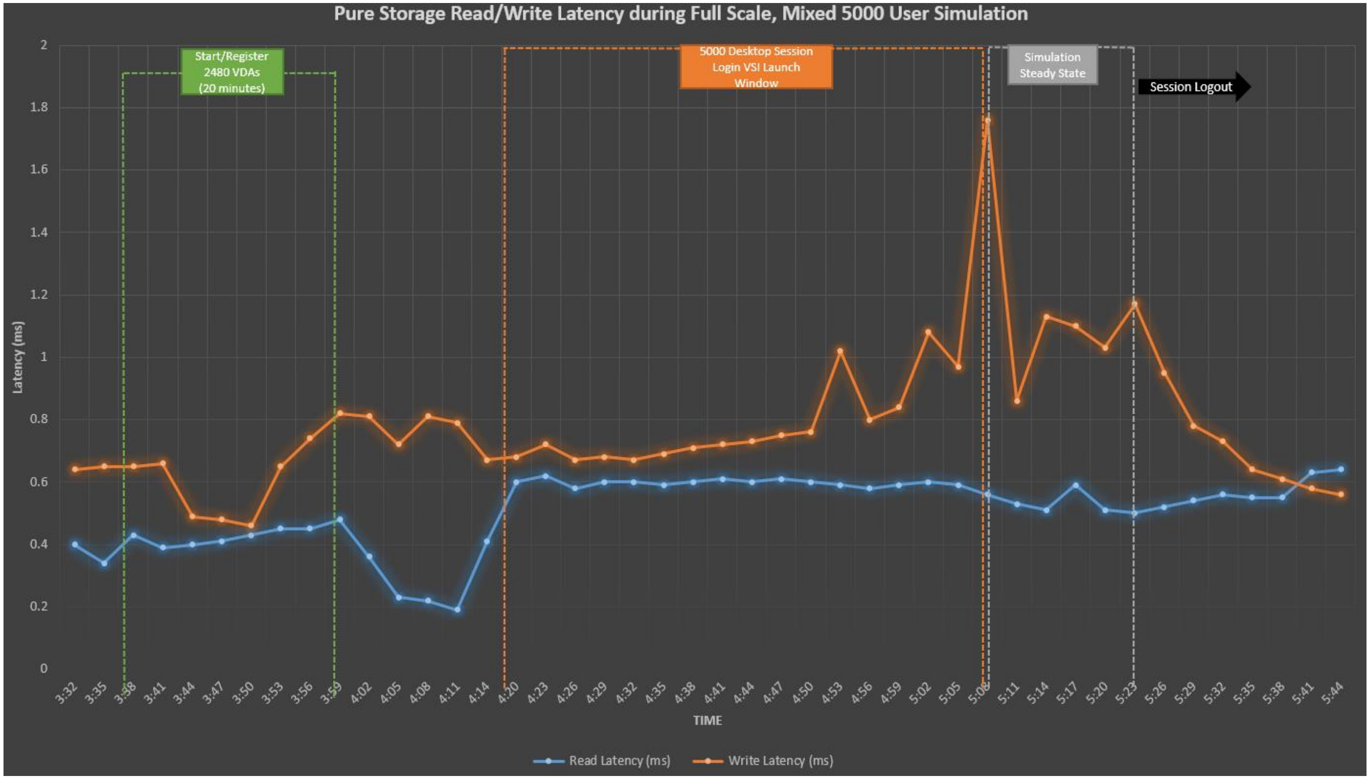

Figure 229 IOPS

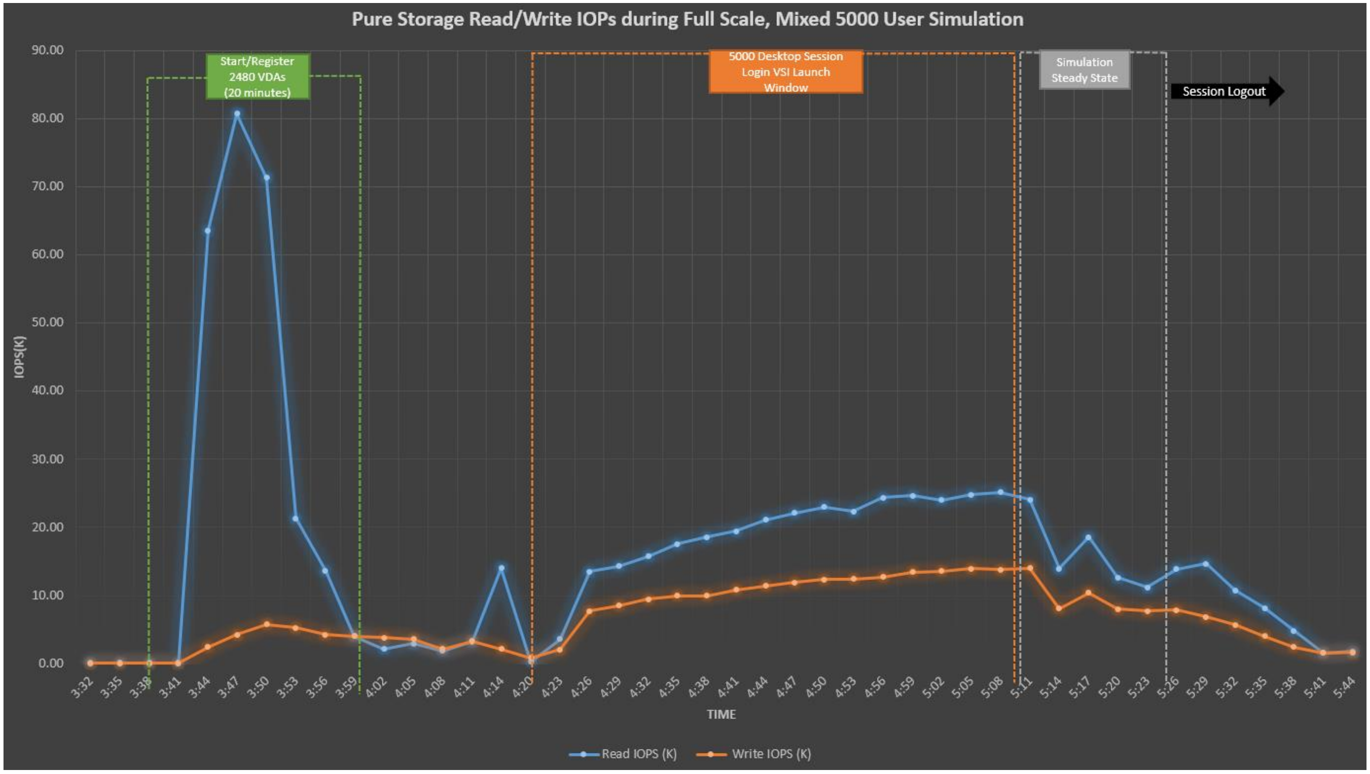

Figure 230 Bandwidth

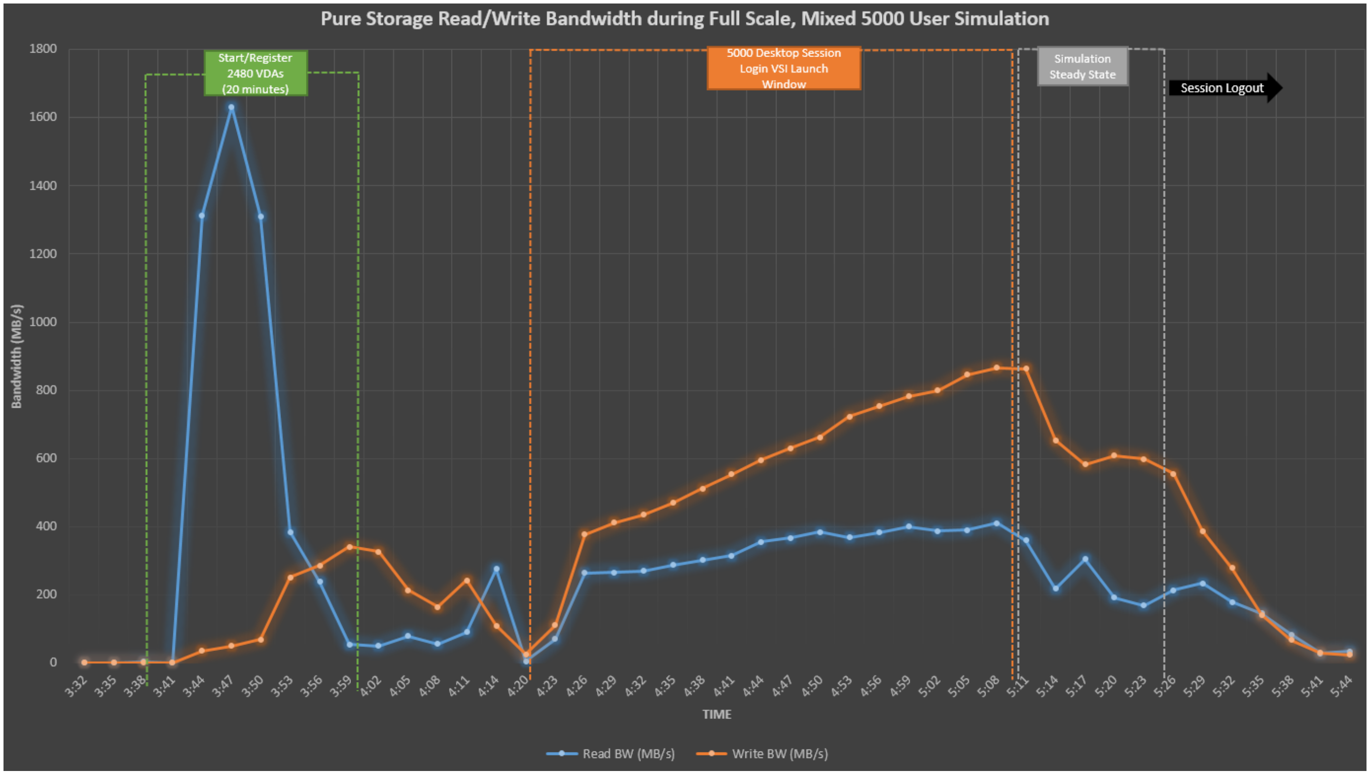

1200 Windows 7 x64 Non-Persistent PVS Sessions

Figure 231 Latency

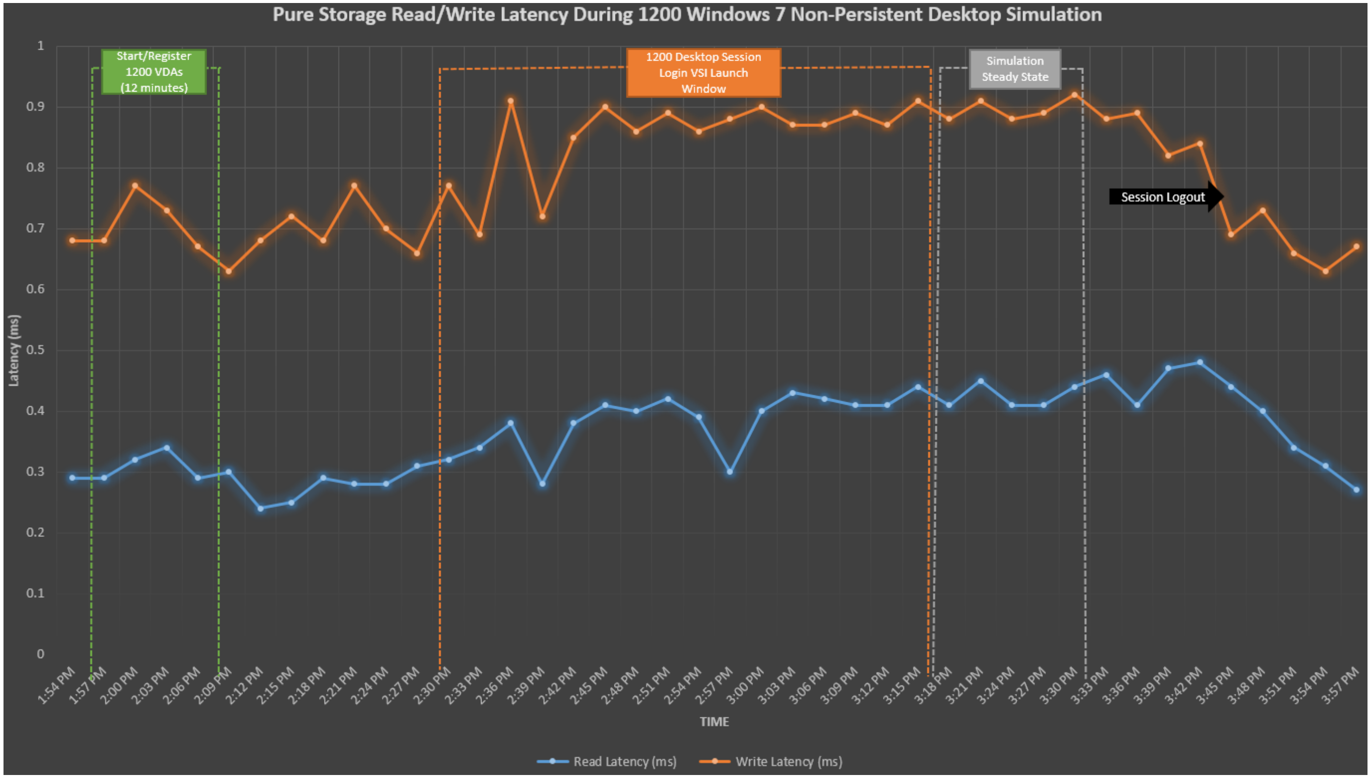

Figure 232 IOPS

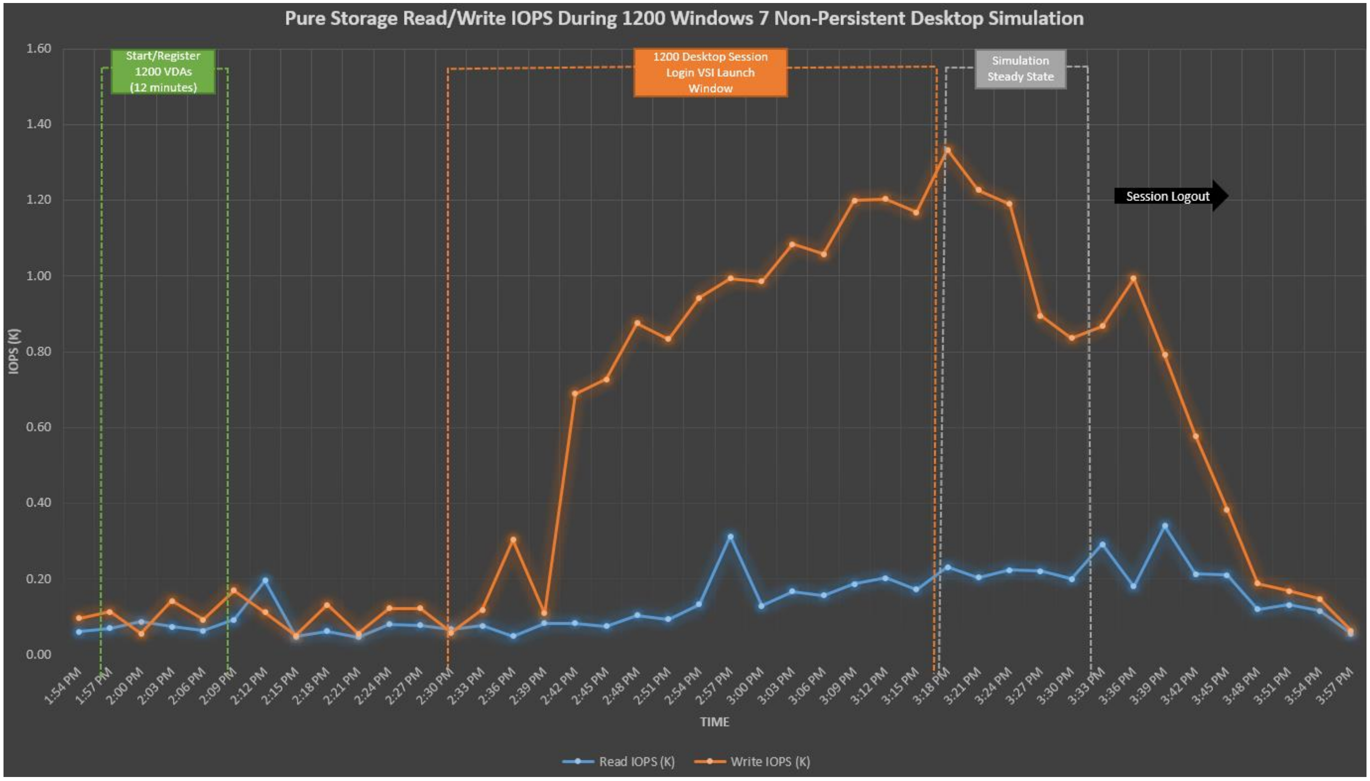

Figure 233 Bandwidth

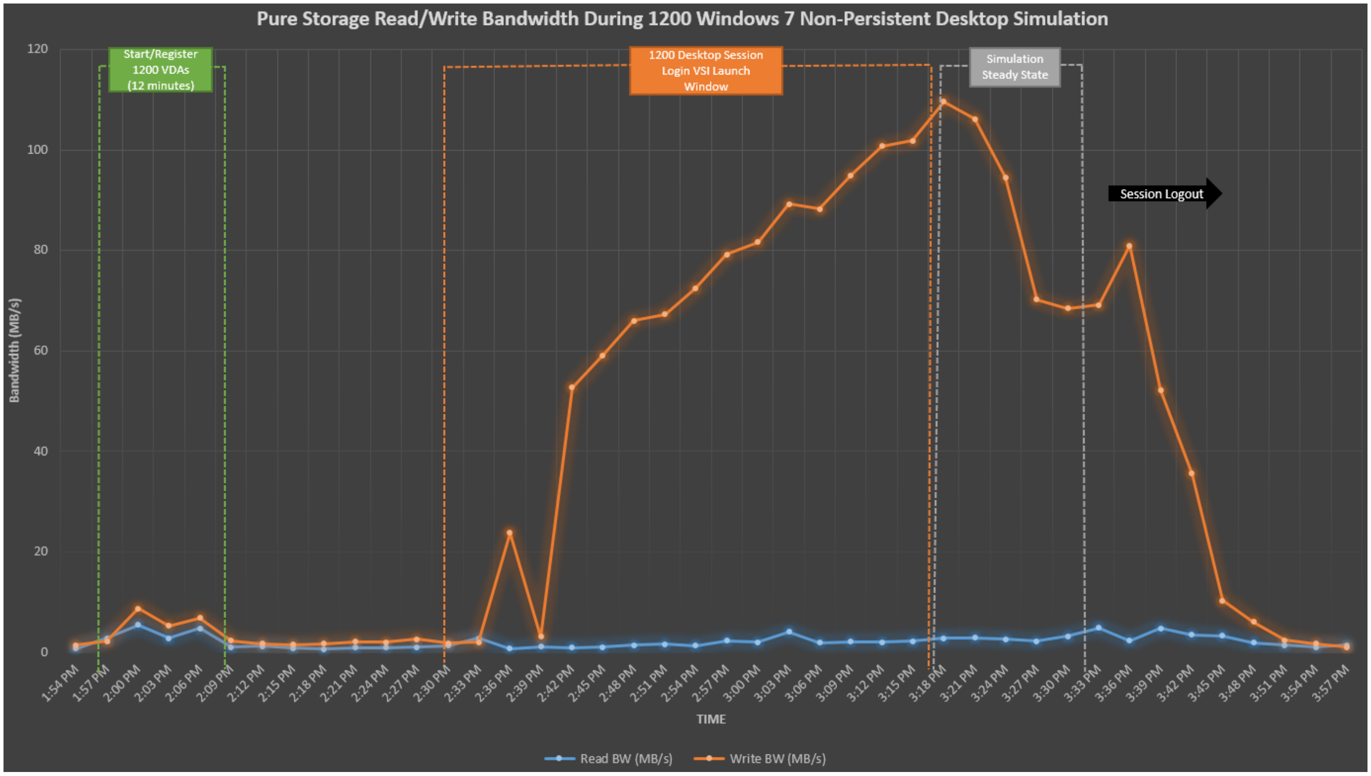

Appendix D – Workload Host Full Scale Performance Data

XenDesktop RDS 10 Host Performance Data

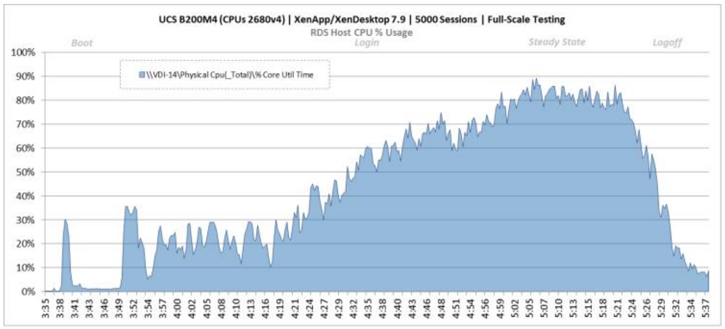

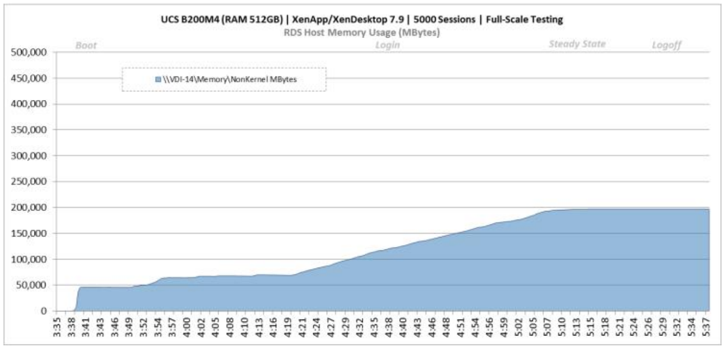

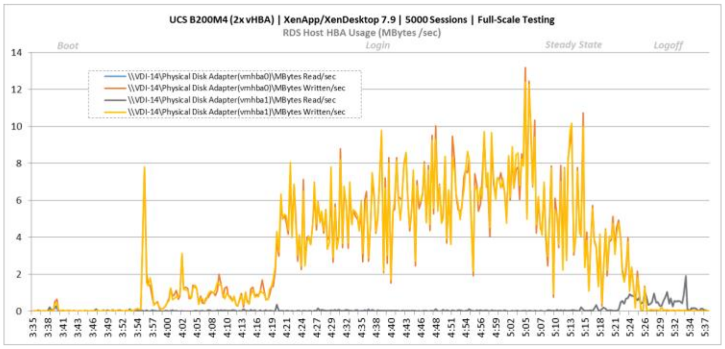

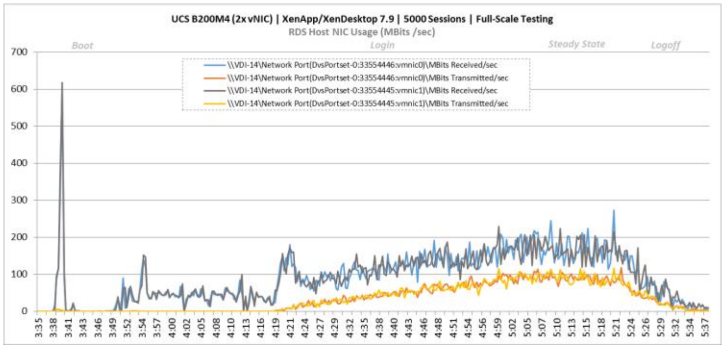

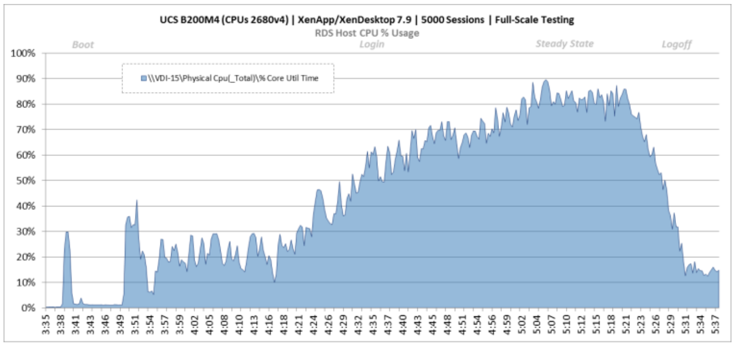
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-15\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:08 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:54 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34 5:37

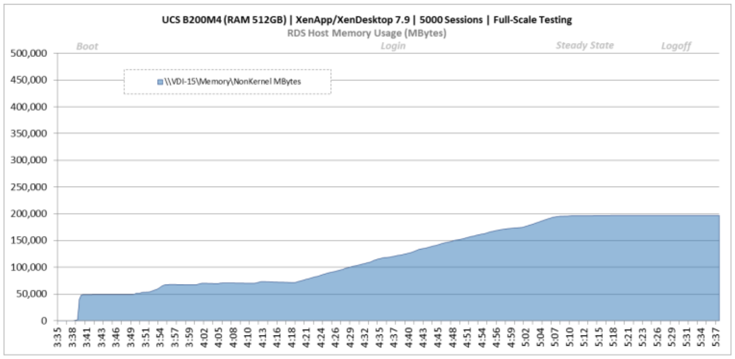
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-15\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:08 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:54 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34 5:37

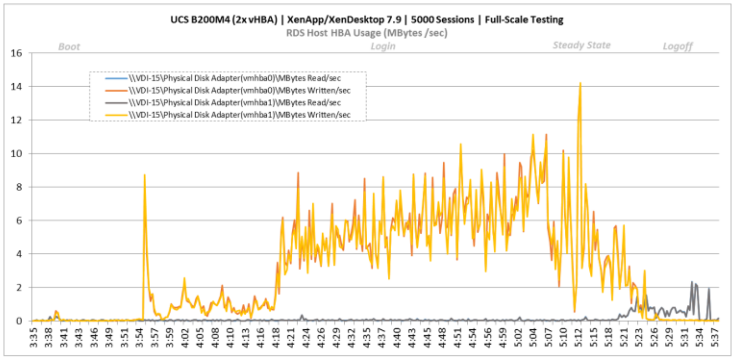
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-15\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-15\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-15\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-15\Physical Disk Adapter(vmhba1)\MBytes Written/sec
16
14
12
10
8
6
4
2
0
3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:08 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:54 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34 5:37

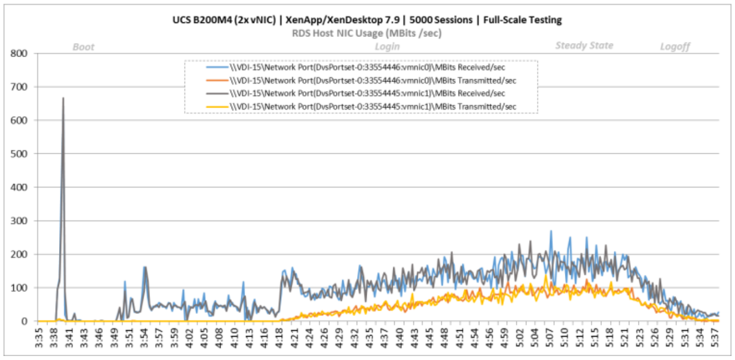
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-15\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-15\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-15\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-15\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:08 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:54 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34 5:37

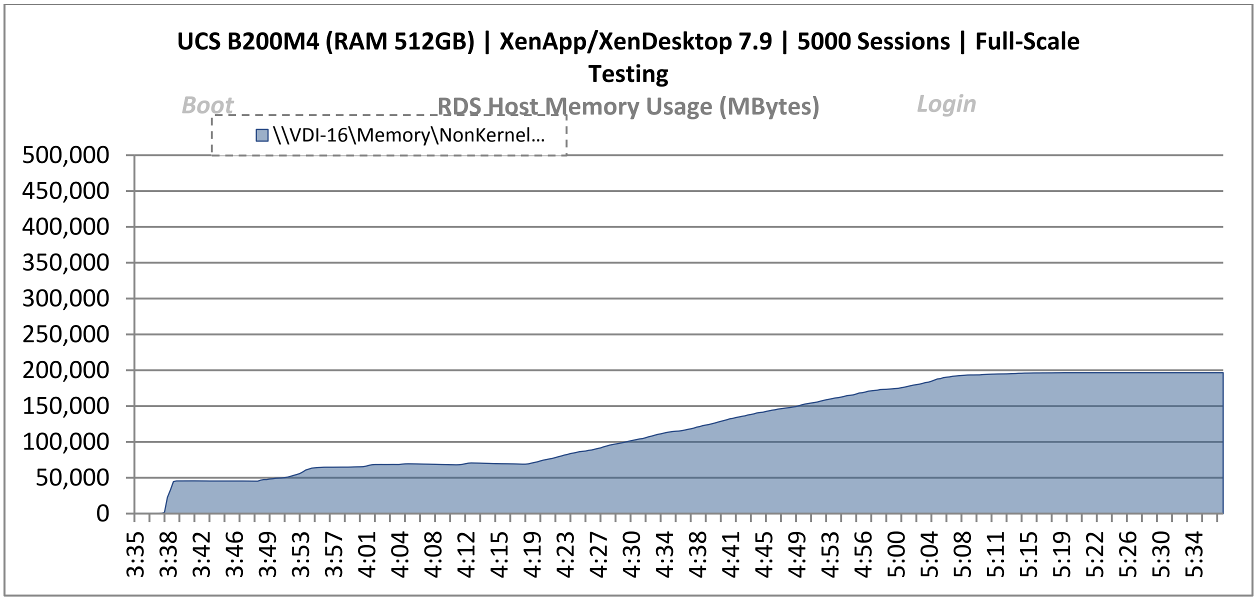
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
\\VDI-16\Memory\NonKernel...
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:35
3:38
3:42
3:46
3:49
3:53
3:57
4:01
4:04
4:08
4:12
4:15
4:19
4:23
4:27
4:30
4:34
4:38
4:41
4:45
4:49
4:53
4:56
5:00
5:04
5:08
5:11
5:15
5:19
5:22
5:26
5:30
5:34

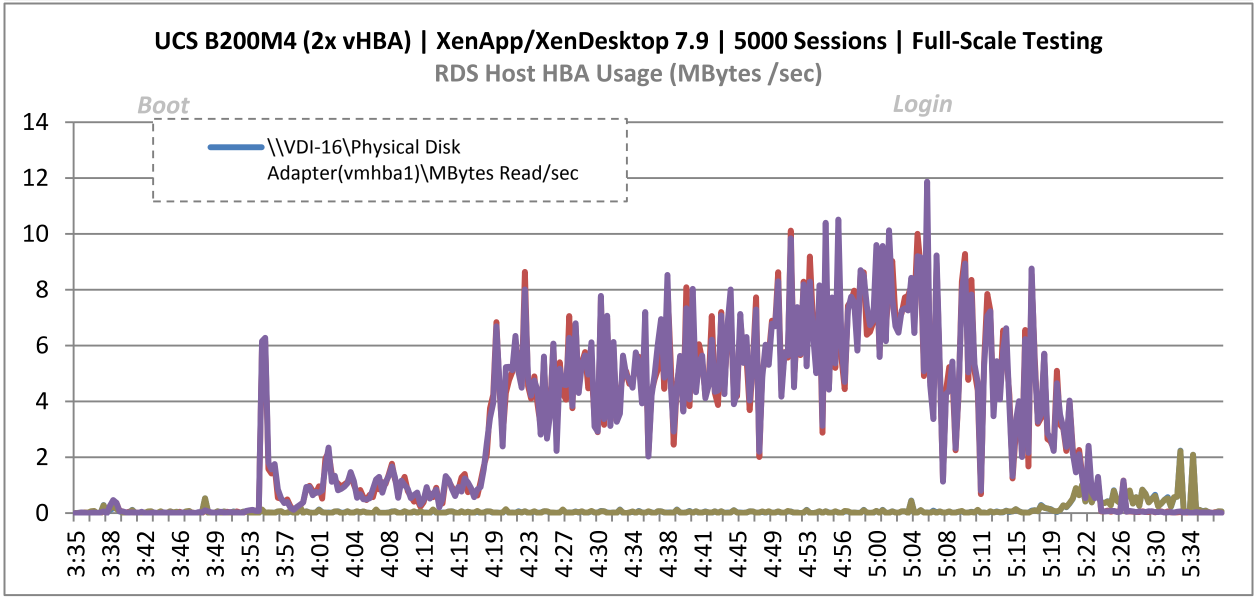
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
\\VDI-16\Physical Disk Adapter(vmhba1)\MBytes Read/sec
14
12
10
8
6
4
2
0
3:35
3:38
3:42
3:46
3:49
3:53
3:57
4:01
4:04
4:08
4:12
4:15
4:19
4:23
4:27
4:30
4:34
4:38
4:41
4:45
4:49
4:53
4:56
5:00
5:04
5:08
5:11
5:15
5:19
5:22
5:26
5:30
5:34

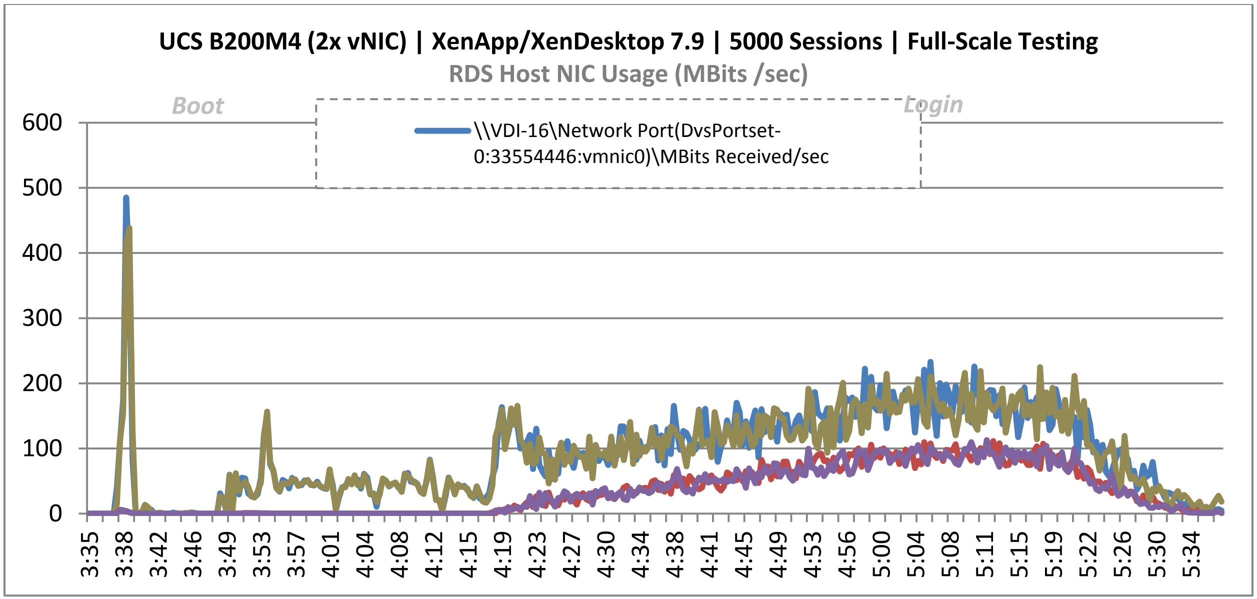
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
\\VDI-16\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
600
500
400
300
200
100
0
3:35
3:38
3:42
3:46
3:49
3:53
3:57
4:01
4:04
4:08
4:12
4:15
4:19
4:23
4:27
4:30
4:34
4:38
4:41
4:45
4:49
4:53
4:56
5:00
5:04
5:08
5:11
5:15
5:19
5:22
5:26
5:30
5:34

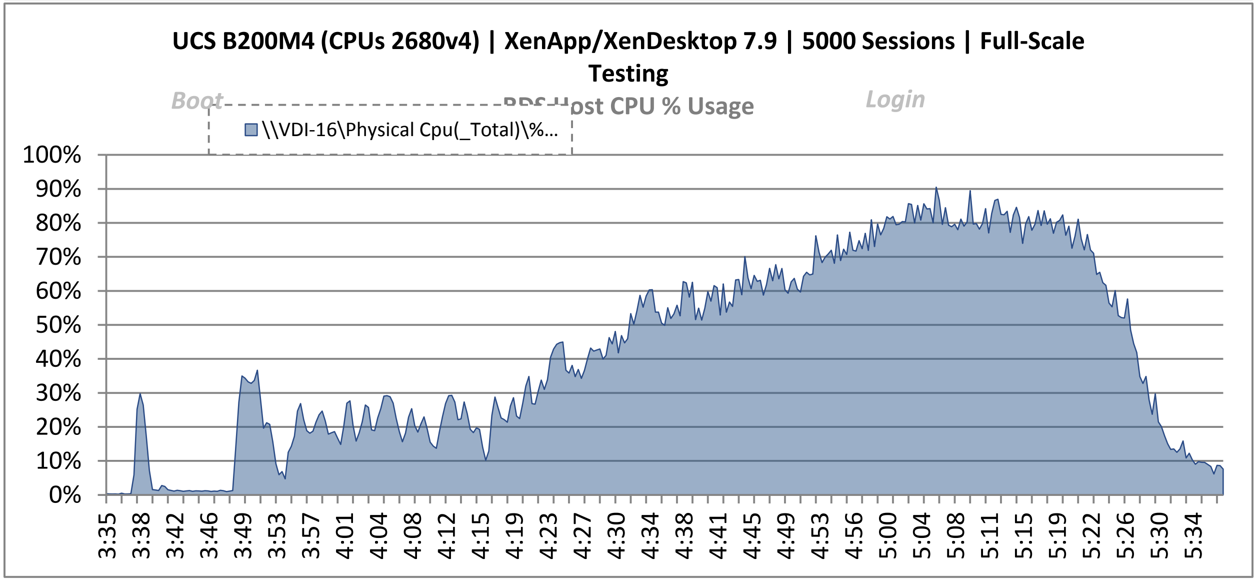
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
\\VDI-16\Physical Cpu(_Total)\%...
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:35
3:38
3:42
3:46
3:49
3:53
3:57
4:01
4:04
4:08
4:12
4:15
4:19
4:23
4:27
4:30
4:34
4:38
4:41
4:45
4:49
4:53
4:56
5:00
5:04
5:08
5:11
5:15
5:19
5:22
5:26
5:30
5:34

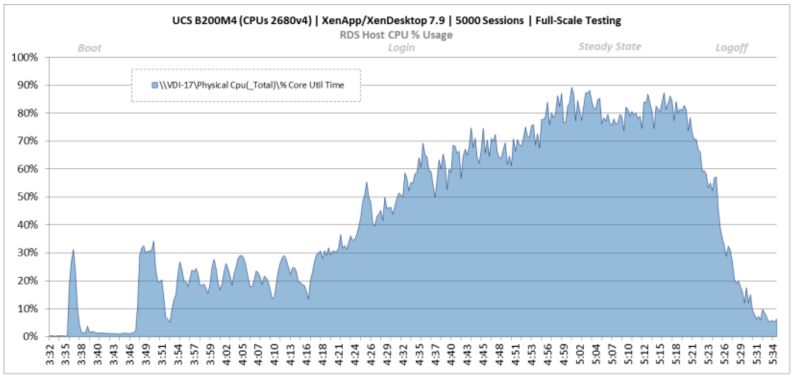
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-17\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

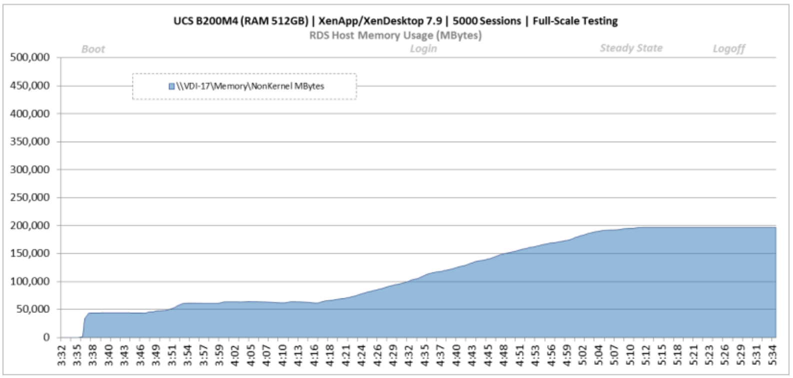
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-17\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

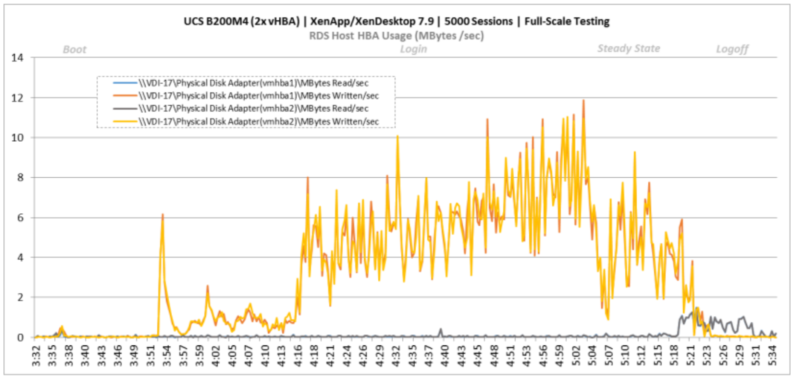
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-17\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-17\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-17\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-17\Physical Disk Adapter(vmhba2)\MBytes Written/sec
14
12
10
8
6
4
2
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

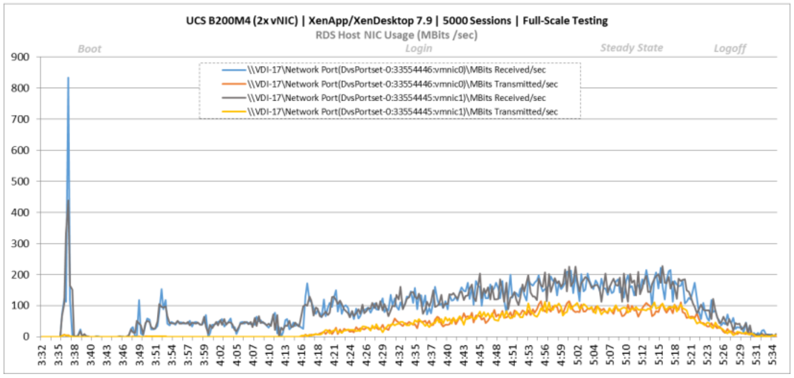
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-17\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-17\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-17\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-17\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
900
800
700
600
500
400
300
200
100
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:35 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

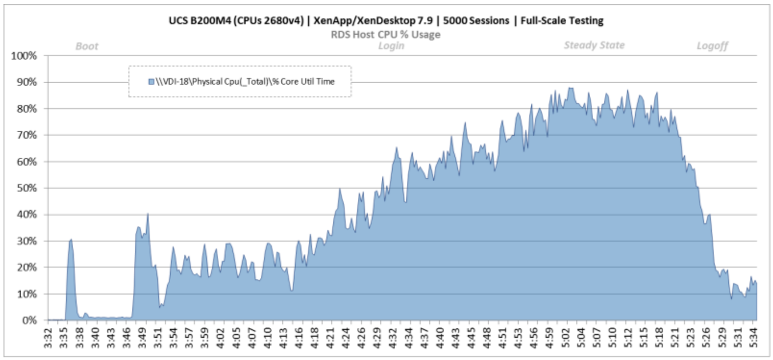
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
■ \\VDI-18\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:34 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

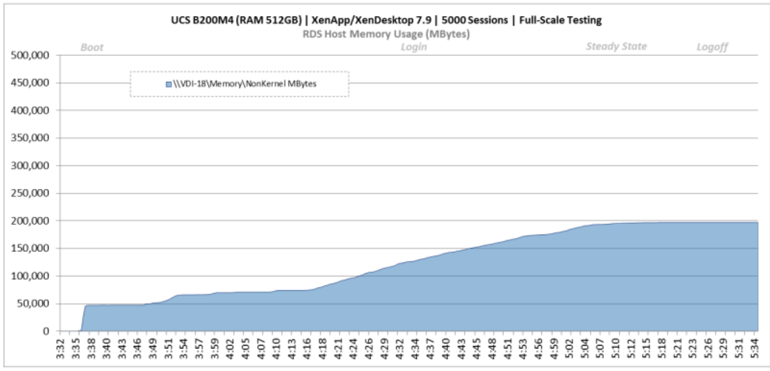
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
■ \\VDI-18\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:34 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

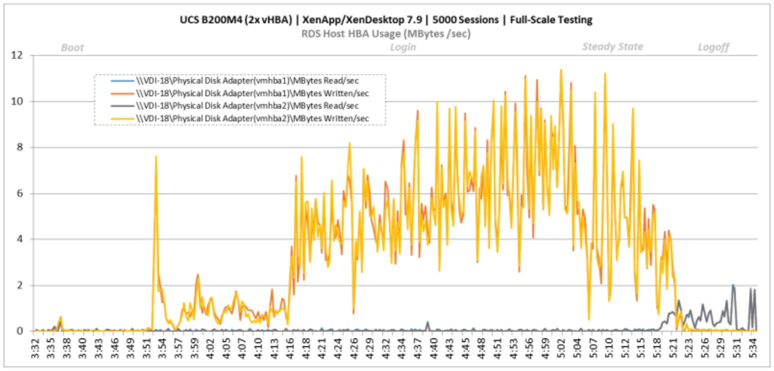
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-18\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-18\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-18\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-18\Physical Disk Adapter(vmhba2)\MBytes Written/sec
12
10
8
6
4
2
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:34 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

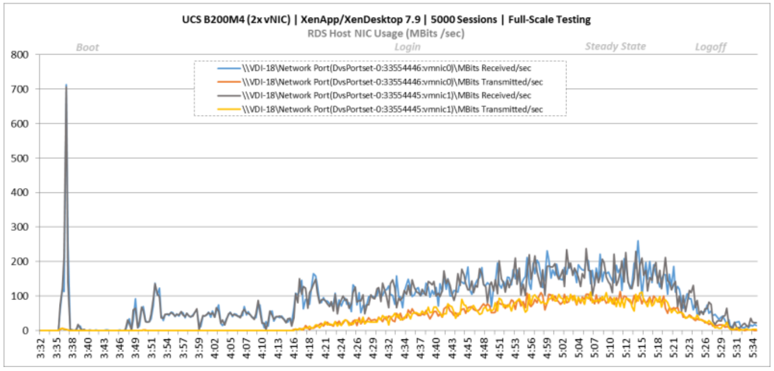
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-18\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-18\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-18\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-18\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:32 3:35 3:38 3:40 3:43 3:46 3:49 3:51 3:54 3:57 3:59 4:02 4:05 4:07 4:10 4:13 4:16 4:18 4:21 4:24 4:26 4:29 4:32 4:34 4:37 4:40 4:43 4:45 4:48 4:51 4:53 4:56 4:59 5:02 5:04 5:07 5:10 5:12 5:15 5:18 5:21 5:23 5:26 5:29 5:31 5:34

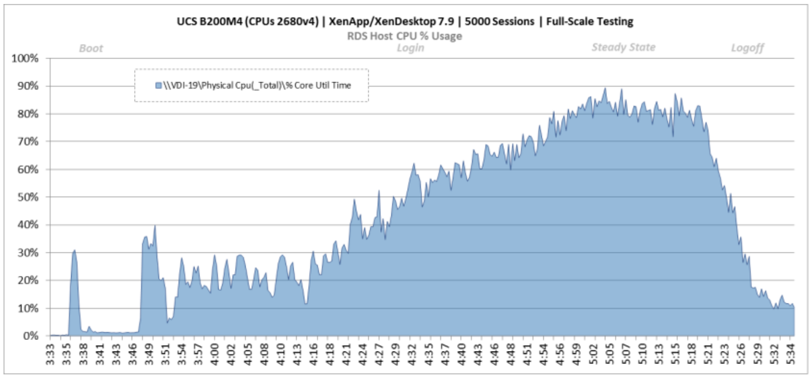
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-19\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33
3:35
3:38
3:41
3:43
3:46
3:49
3:51
3:54
3:57
4:00
4:02
4:05
4:08
4:10
4:13
4:16
4:18
4:21
4:24
4:27
4:29
4:32
4:35
4:37
4:40
4:43
4:46
4:48
4:51
4:54
4:56
4:59
5:02
5:05
5:07
5:10
5:13
5:15
5:18
5:21
5:23
5:26
5:29
5:32
5:34

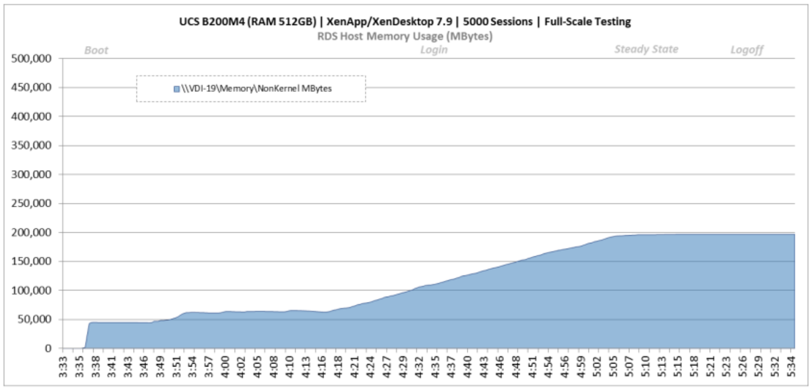
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-19\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33
3:35
3:38
3:41
3:43
3:46
3:49
3:51
3:54
3:57
4:00
4:02
4:05
4:08
4:10
4:13
4:16
4:18
4:21
4:24
4:27
4:29
4:32
4:35
4:37
4:40
4:43
4:46
4:48
4:51
4:54
4:56
4:59
5:02
5:05
5:07
5:10
5:13
5:15
5:18
5:21
5:23
5:26
5:29
5:32
5:34

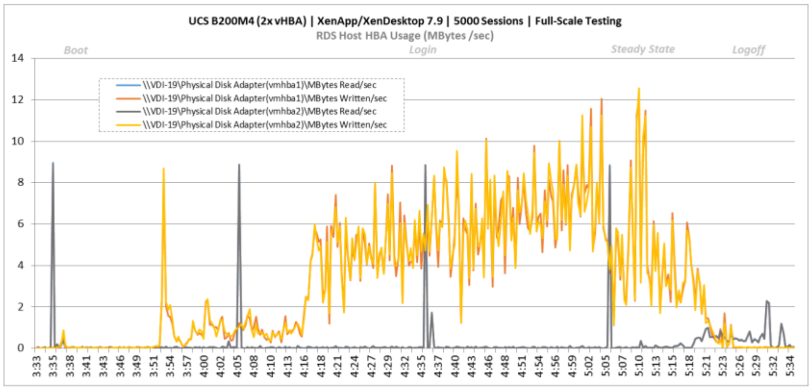
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-19\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-19\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-19\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-19\Physical Disk Adapter(vmhba2)\MBytes Written/sec
14
12
10
8
6
4
2
0
3:33
3:35
3:38
3:41
3:43
3:46
3:49
3:51
3:54
3:57
4:00
4:02
4:05
4:08
4:10
4:13
4:16
4:18
4:21
4:24
4:27
4:29
4:32
4:35
4:37
4:40
4:43
4:46
4:48
4:51
4:54
4:56
4:59
5:02
5:05
5:07
5:10
5:13
5:15
5:18
5:21
5:23
5:26
5:29
5:32
5:34

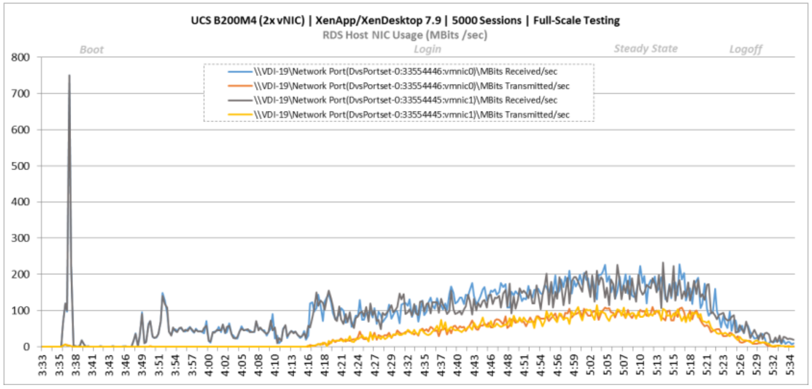
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-19\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-19\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-19\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-19\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:33
3:35
3:38
3:41
3:43
3:46
3:49
3:51
3:54
3:57
4:00
4:02
4:05
4:08
4:10
4:13
4:16
4:18
4:21
4:24
4:27
4:29
4:32
4:35
4:37
4:40
4:43
4:46
4:48
4:51
4:54
4:56
4:59
5:02
5:05
5:07
5:10
5:13
5:15
5:18
5:21
5:23
5:26
5:29
5:32
5:34

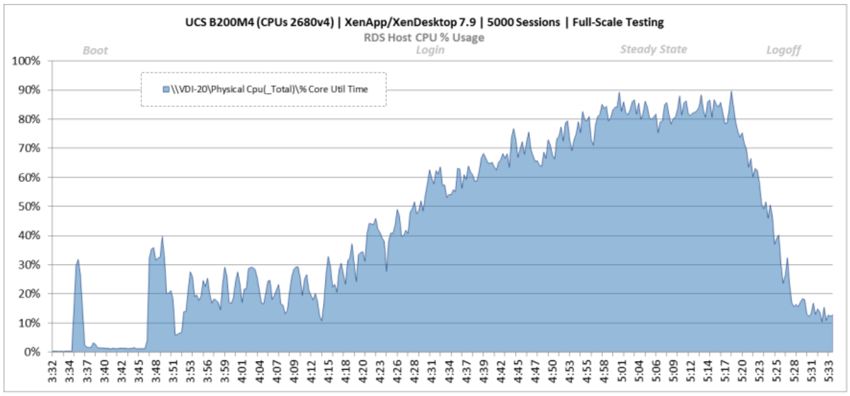
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-20\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:32 3:34 3:37 3:40 3:42 3:45 3:48 3:51 3:53 3:56 3:59 4:01 4:04 4:07 4:09 4:12 4:15 4:18 4:20 4:23 4:26 4:28 4:31 4:34 4:37 4:39 4:42 4:45 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:33

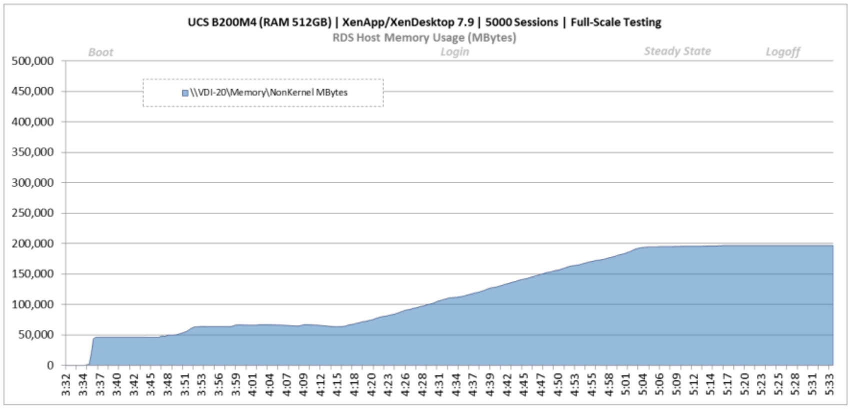
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-20\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:32 3:34 3:37 3:40 3:42 3:45 3:48 3:51 3:53 3:56 3:59 4:01 4:04 4:07 4:09 4:12 4:15 4:18 4:20 4:23 4:26 4:28 4:31 4:34 4:37 4:39 4:42 4:45 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:33

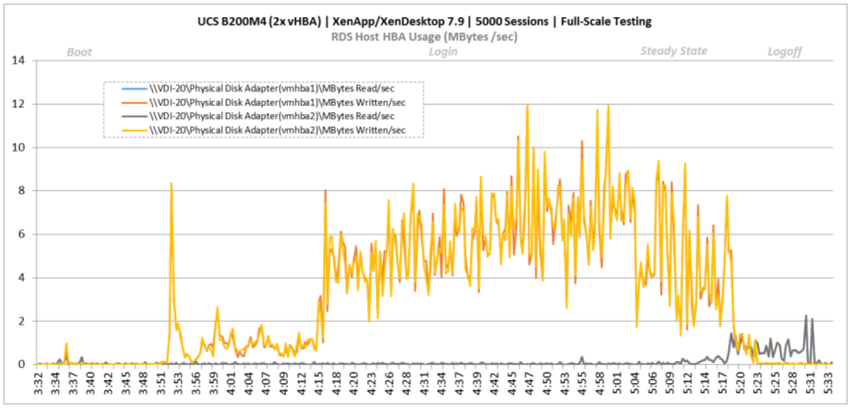
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-20\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-20\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-20\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-20\Physical Disk Adapter(vmhba2)\MBytes Written/sec
14
12
10
8
6
4
2
0
3:32 3:34 3:37 3:40 3:42 3:45 3:48 3:51 3:53 3:56 3:59 4:01 4:04 4:07 4:09 4:12 4:15 4:18 4:20 4:23 4:26 4:28 4:31 4:34 4:37 4:39 4:42 4:45 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:33

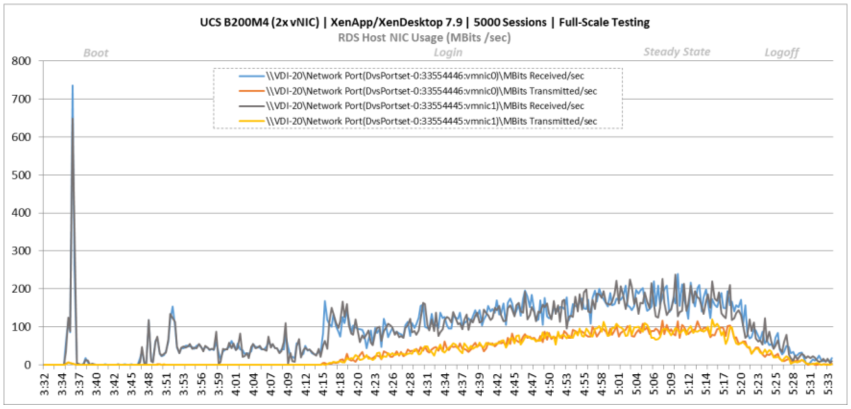
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-20\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-20\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-20\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-20\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:32 3:34 3:37 3:40 3:42 3:45 3:48 3:51 3:53 3:56 3:59 4:01 4:04 4:07 4:09 4:12 4:15 4:18 4:20 4:23 4:26 4:28 4:31 4:34 4:37 4:39 4:42 4:45 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:33

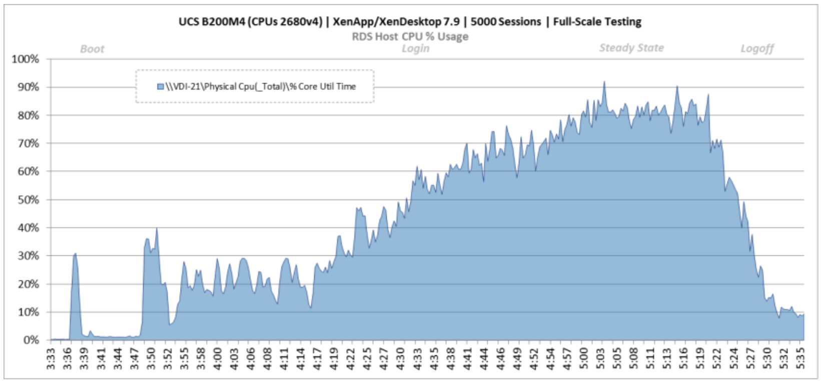
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-21\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:00 4:03 4:06 4:08 4:11 4:14 4:17 4:19 4:22 4:25 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:54 4:57 5:00 5:03 5:05 5:08 5:11 5:13 5:16 5:19 5:22 5:24 5:27 5:30 5:32 5:35

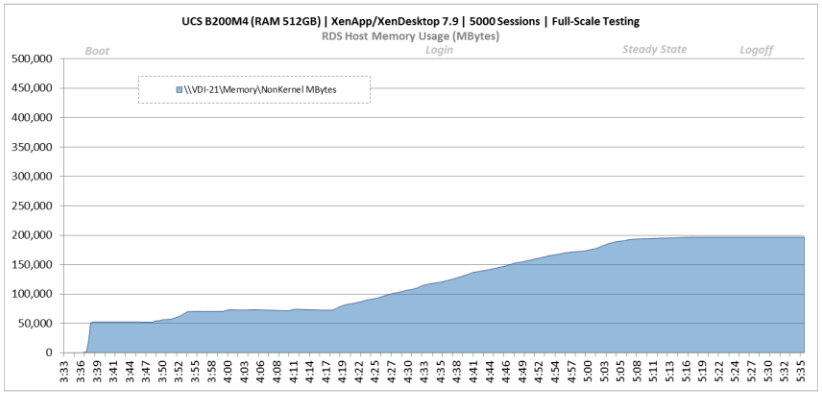
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-21\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:00 4:03 4:06 4:08 4:11 4:14 4:17 4:19 4:22 4:25 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:54 4:57 5:00 5:03 5:05 5:08 5:11 5:13 5:16 5:19 5:22 5:24 5:27 5:30 5:32 5:35

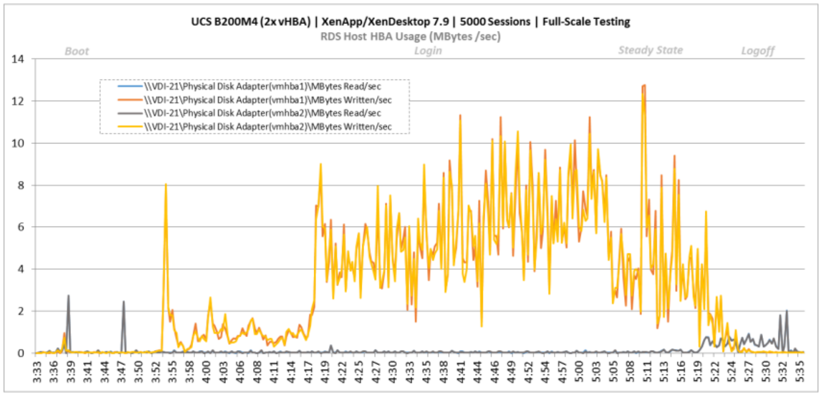
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-21\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-21\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-21\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-21\Physical Disk Adapter(vmhba2)\MBytes Written/sec
14
12
10
8
6
4
2
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:00 4:03 4:06 4:08 4:11 4:14 4:17 4:19 4:22 4:25 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:54 4:57 5:00 5:03 5:05 5:08 5:11 5:13 5:16 5:19 5:22 5:24 5:27 5:30 5:32 5:35

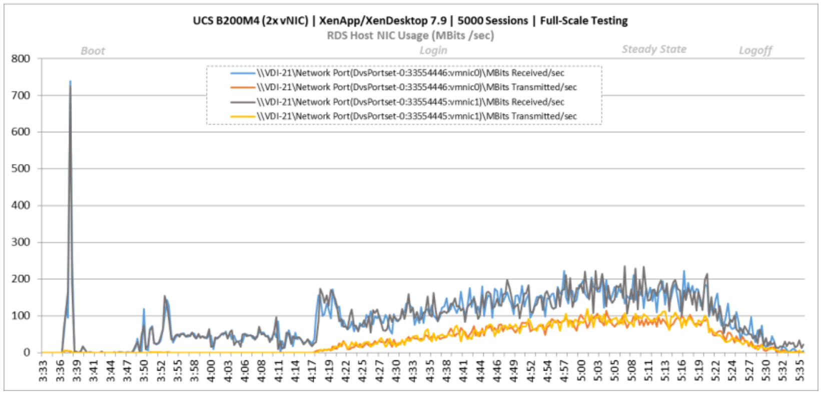
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-21\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-21\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-21\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-21\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:00 4:03 4:06 4:08 4:11 4:14 4:17 4:19 4:22 4:25 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:54 4:57 5:00 5:03 5:05 5:08 5:11 5:13 5:16 5:19 5:22 5:24 5:27 5:30 5:32 5:35

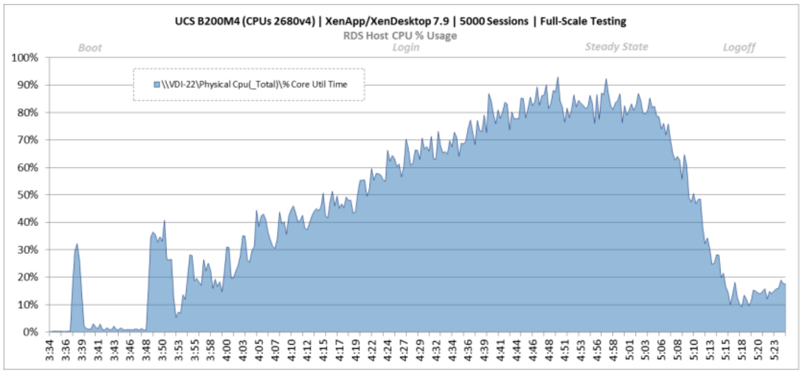
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-22\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:34
3:36
3:39
3:41
3:43
3:46
3:48
3:50
3:53
3:55
3:58
4:00
4:03
4:05
4:07
4:10
4:12
4:15
4:17
4:19
4:22
4:24
4:27
4:29
4:32
4:34
4:36
4:39
4:41
4:44
4:46
4:48
4:51
4:53
4:56
4:58
5:01
5:03
5:06
5:08
5:10
5:13
5:15
5:18
5:20
5:23

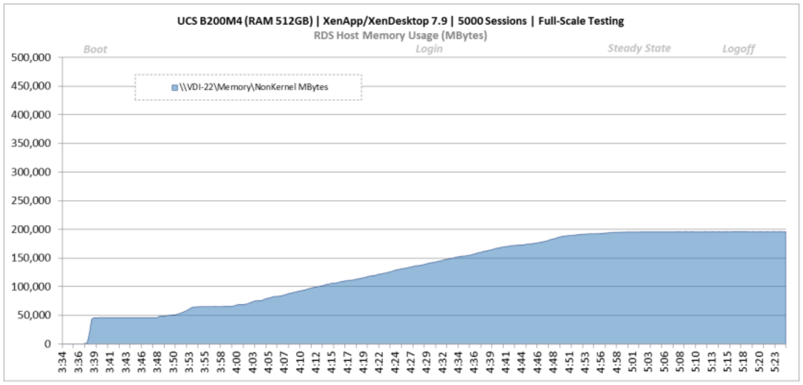
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-22\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:34
3:36
3:39
3:41
3:43
3:46
3:48
3:50
3:53
3:55
3:58
4:00
4:03
4:05
4:07
4:10
4:12
4:15
4:17
4:19
4:22
4:24
4:27
4:29
4:32
4:34
4:36
4:39
4:41
4:44
4:46
4:48
4:51
4:53
4:56
4:58
5:01
5:03
5:06
5:08
5:10
5:13
5:15
5:18
5:20
5:23

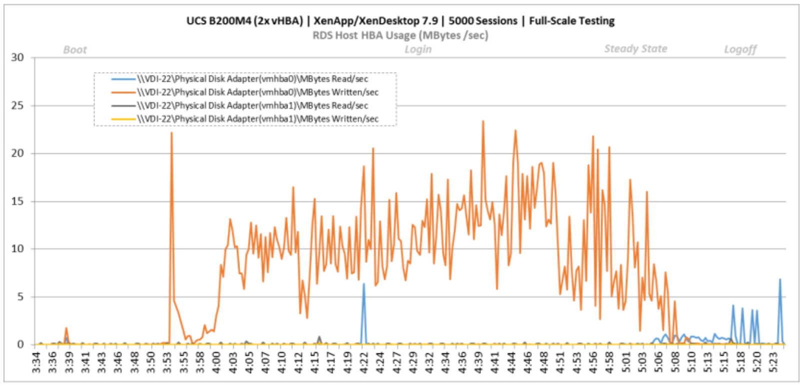
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-22\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-22\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-22\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-22\Physical Disk Adapter(vmhba1)\MBytes Written/sec
30
25
20
15
10
5
0
3:34
3:36
3:39
3:41
3:43
3:46
3:48
3:50
3:53
3:55
3:58
4:00
4:03
4:05
4:07
4:10
4:12
4:15
4:17
4:19
4:22
4:24
4:27
4:29
4:32
4:34
4:36
4:39
4:41
4:44
4:46
4:48
4:51
4:53
4:56
4:58
5:01
5:03
5:06
5:08
5:10
5:13
5:15
5:18
5:20
5:23

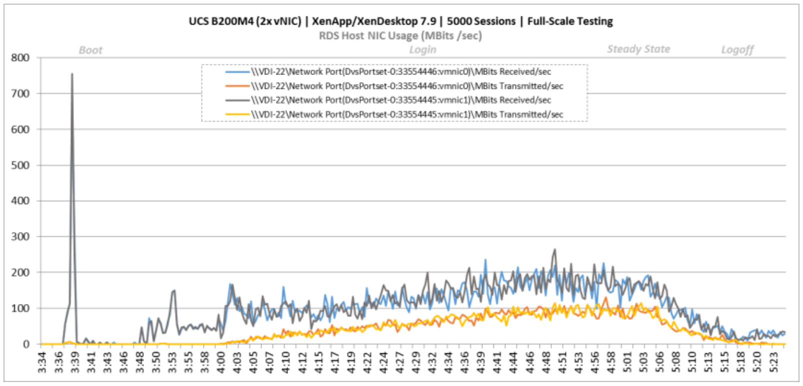
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
RDS Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-22\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-22\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-22\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-22\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
800
700
600
500
400
300
200
100
0
3:34
3:36
3:39
3:41
3:43
3:46
3:48
3:50
3:53
3:55
3:58
4:00
4:03
4:05
4:07
4:10
4:12
4:15
4:17
4:19
4:22
4:24
4:27
4:29
4:32
4:34
4:36
4:39
4:41
4:44
4:46
4:48
4:51
4:53
4:56
4:58
5:01
5:03
5:06
5:08
5:10
5:13
5:15
5:18
5:20
5:23

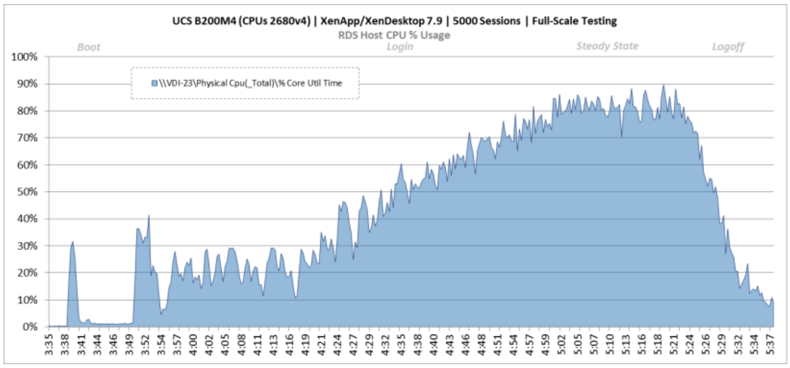

UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

RDS Host CPU % Usage

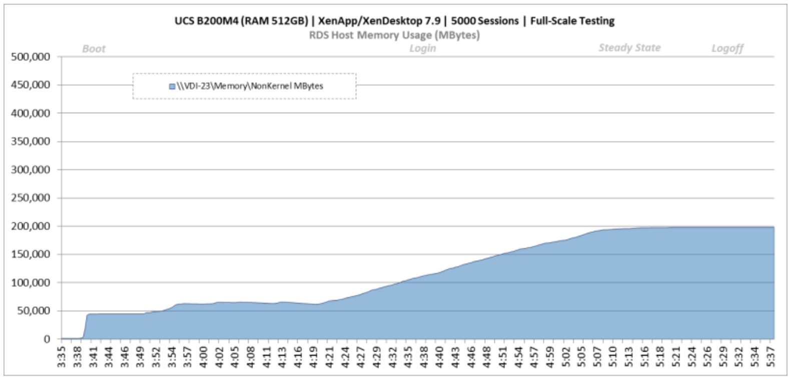

UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

RDS Host Memory Usage (MBytes)

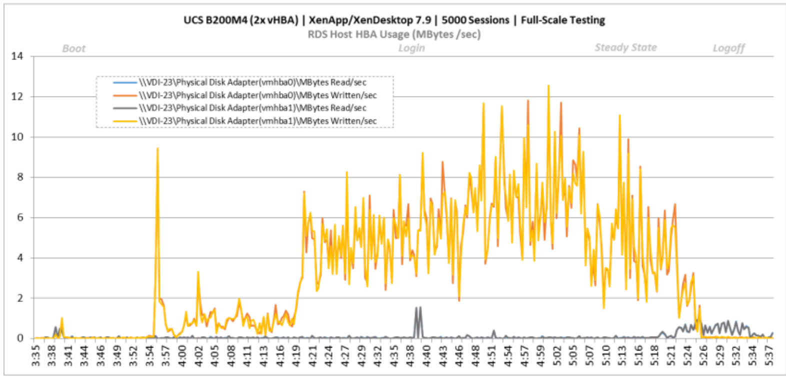

UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

RDS Host HBA Usage (MBytes /sec)

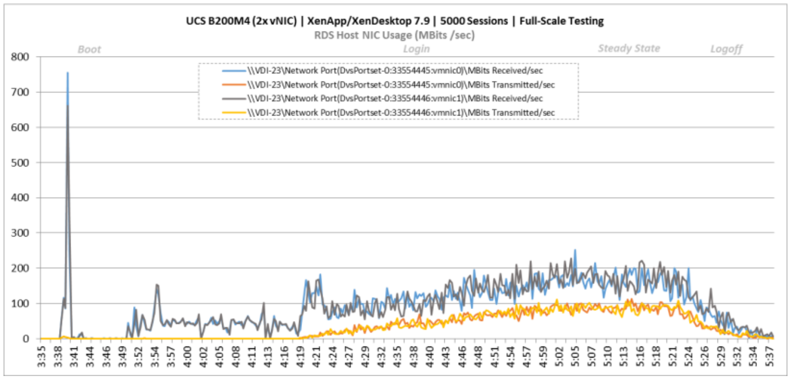

UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

RDS Host NIC Usage (MBits /sec)

XenDesktop Non-Persistent 8 Host Performance Detail

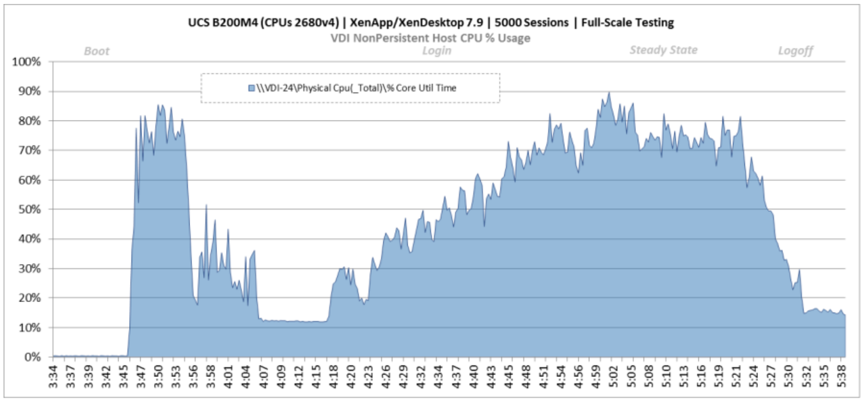

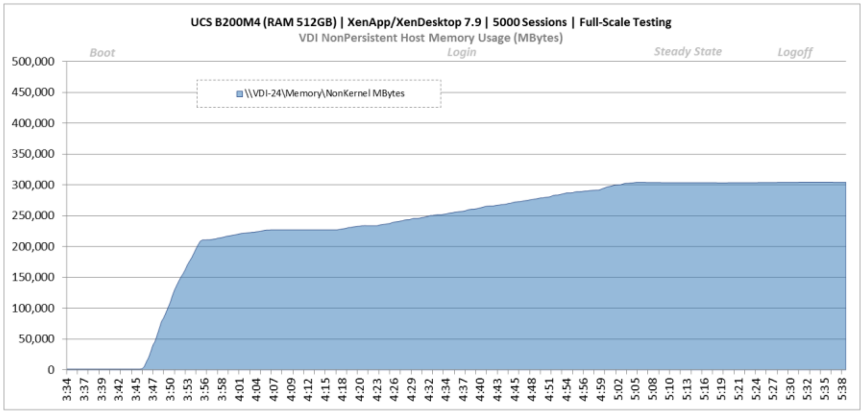

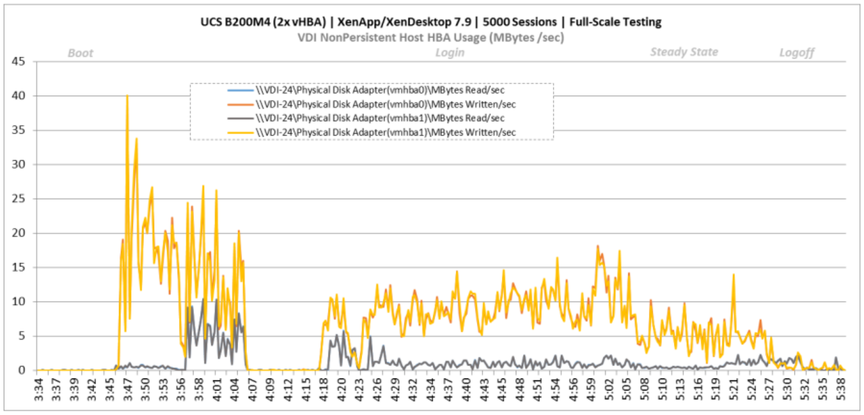

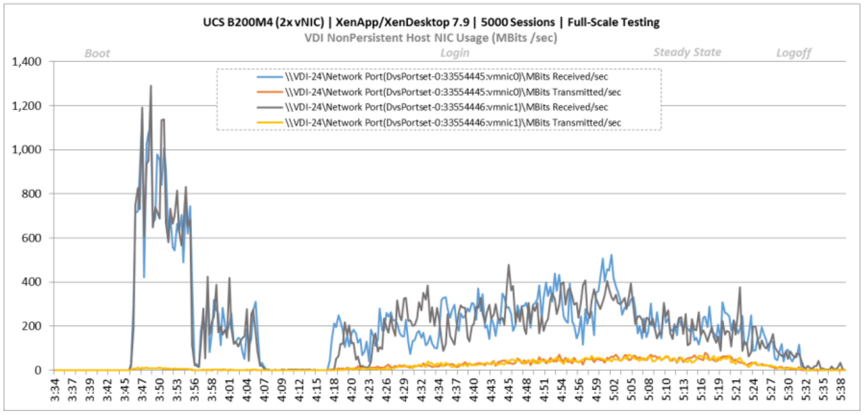

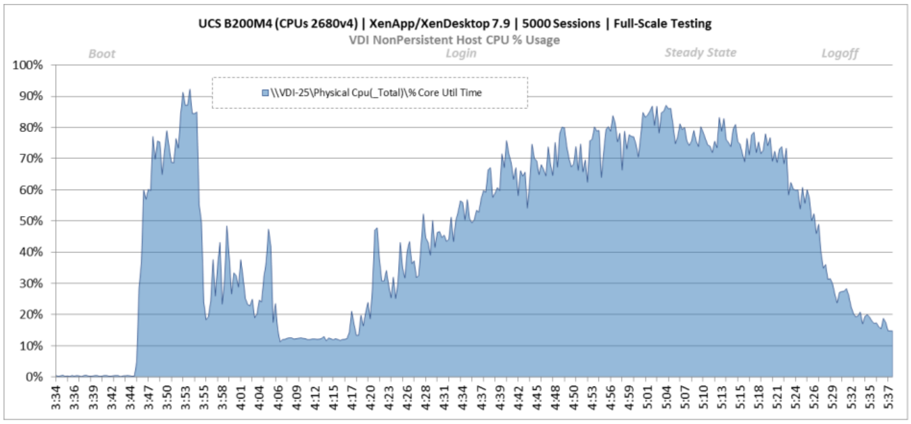
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-25\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:34
5:37

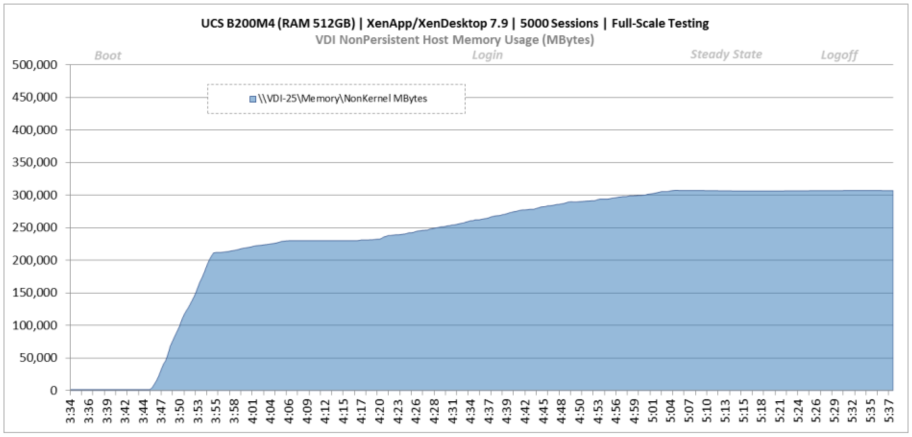
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-25\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:34
5:37

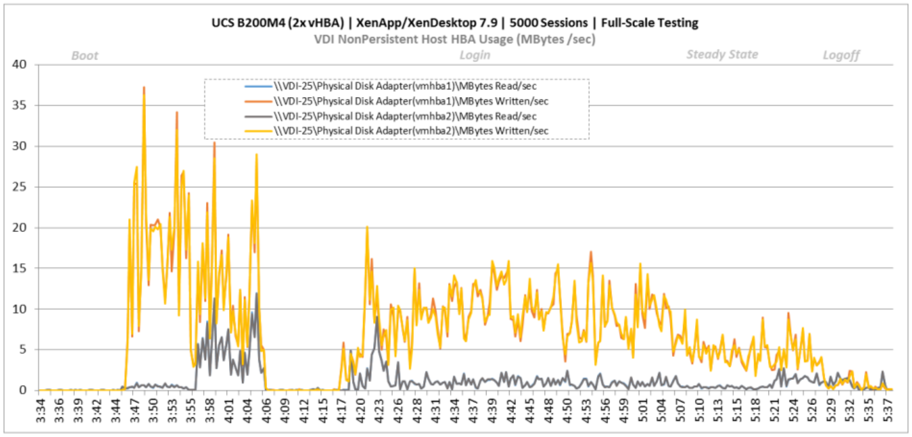
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-25\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-25\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-25\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-25\Physical Disk Adapter(vmhba2)\MBytes Written/sec
40
35
30
25
20
15
10
5
0
3:34
5:37

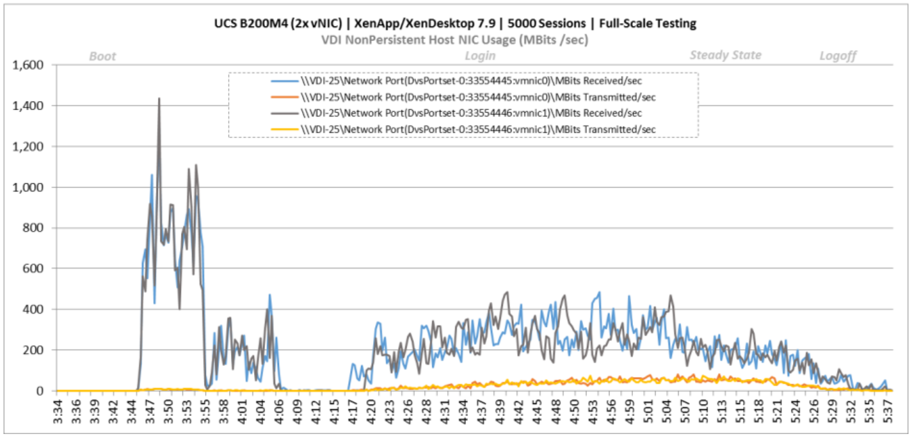
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-25\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-25\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-25\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-25\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
1,600
1,400
1,200
1,000
800
600
400
200
0
3:34
5:37

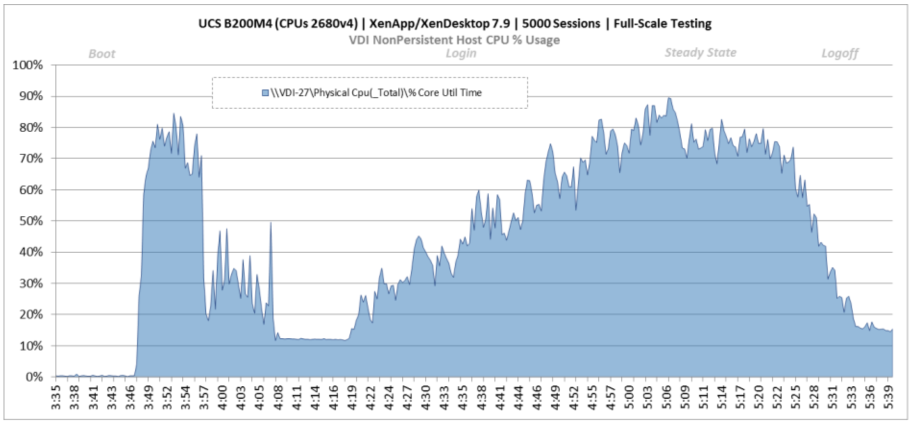
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-27\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:35 3:38 3:41 3:43 3:46 3:49 3:52 3:54 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:09 5:11 5:14 5:17 5:20 5:22 5:25 5:28 5:31 5:33 5:36 5:39

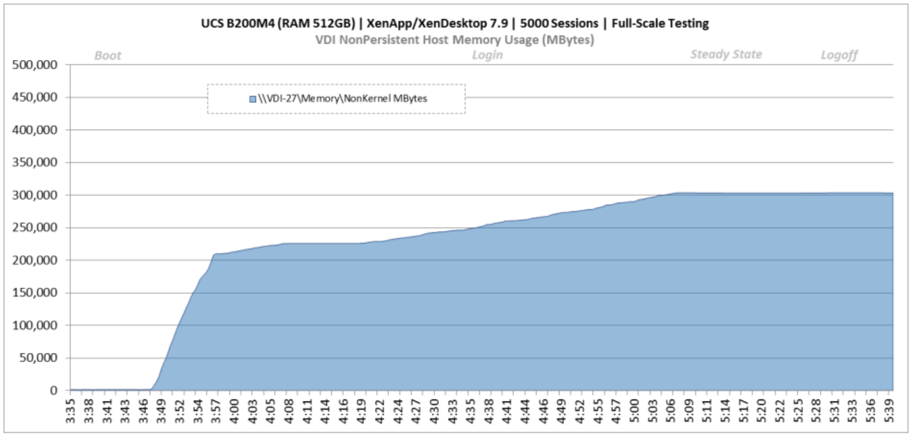
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-27\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:35 3:38 3:41 3:43 3:46 3:49 3:52 3:54 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:09 5:11 5:14 5:17 5:20 5:22 5:25 5:28 5:31 5:33 5:36 5:39

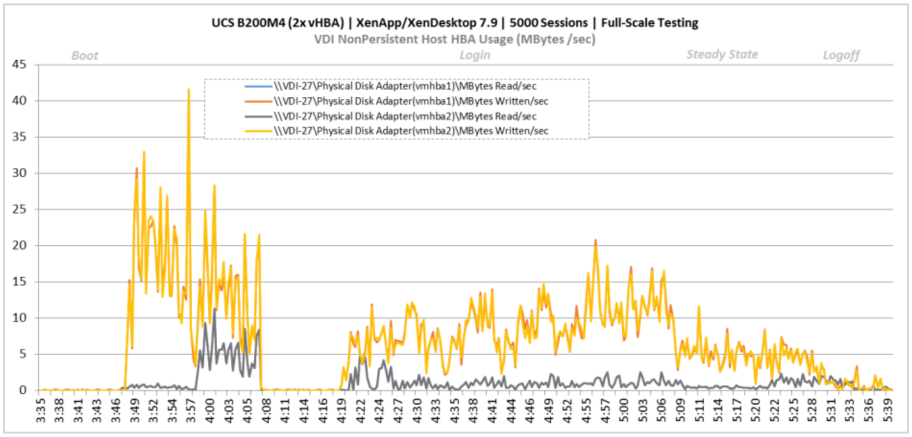
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-27\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-27\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-27\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-27\Physical Disk Adapter(vmhba2)\MBytes Written/sec
45
40
35
30
25
20
15
10
5
0
3:35 3:38 3:41 3:43 3:46 3:49 3:52 3:54 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:09 5:11 5:14 5:17 5:20 5:22 5:25 5:28 5:31 5:33 5:36 5:39

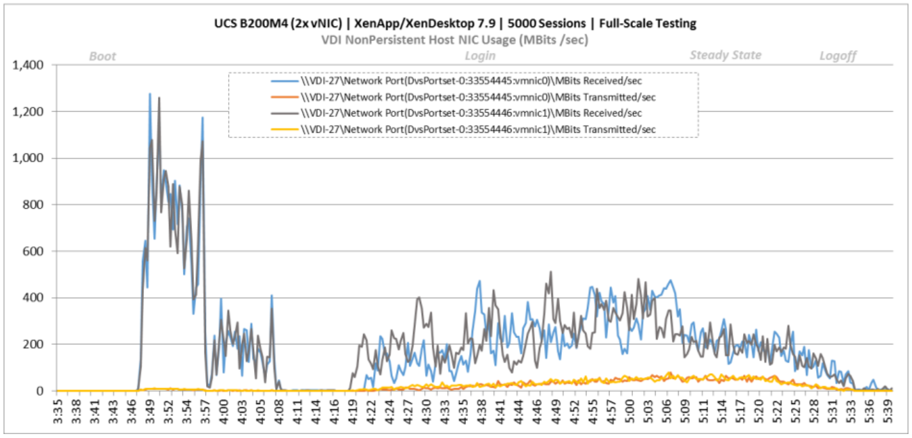
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-27\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-27\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-27\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-27\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
1,400
1,200
1,000
800
600
400
200
0
3:35 3:38 3:41 3:43 3:46 3:49 3:52 3:54 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:09 5:11 5:14 5:17 5:20 5:22 5:25 5:28 5:31 5:33 5:36 5:39

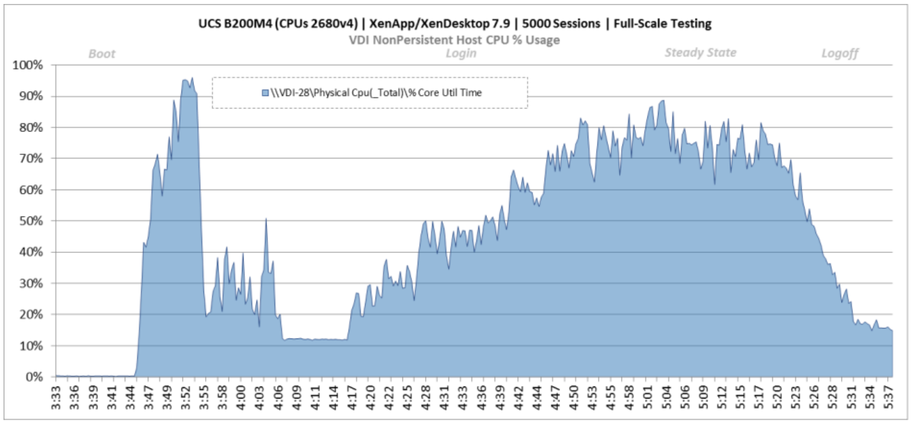
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-28\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:17 5:20 5:23 5:26 5:28 5:31 5:34 5:37

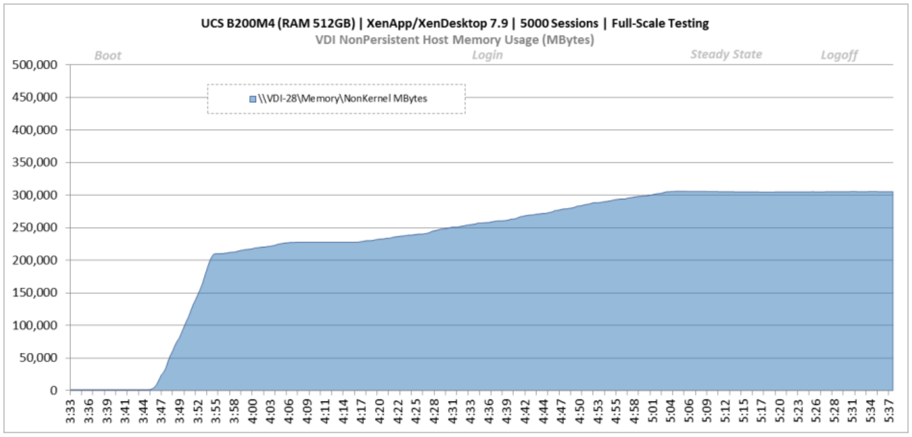
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-28\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:17 5:20 5:23 5:26 5:28 5:31 5:34 5:37

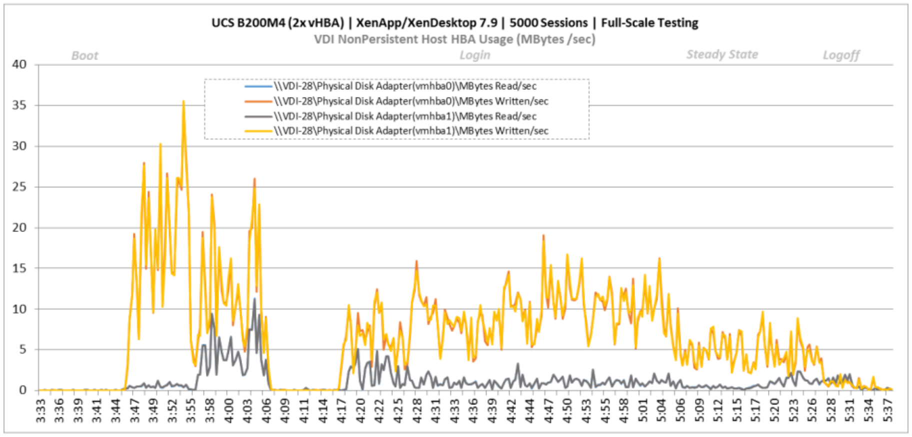
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-28\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-28\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-28\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-28\Physical Disk Adapter(vmhba1)\MBytes Written/sec
40
35
30
25
20
15
10
5
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:17 5:20 5:23 5:26 5:28 5:31 5:34 5:37

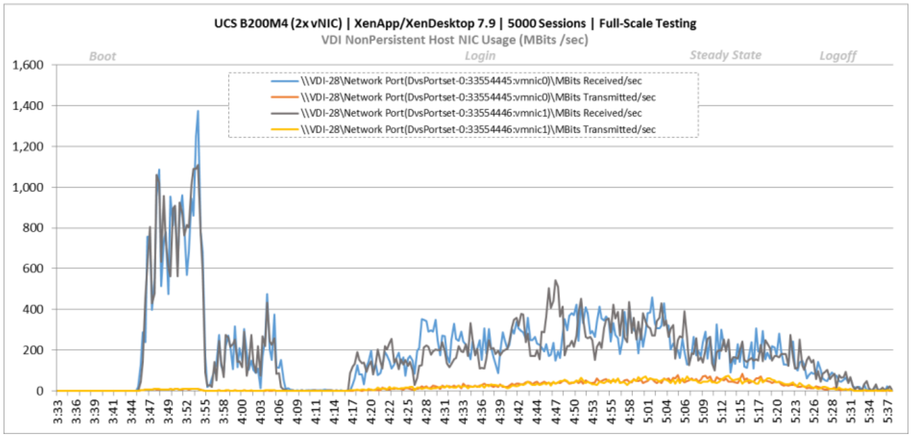
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-28\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-28\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-28\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-28\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
1,600
1,400
1,200
1,000
800
600
400
200
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:17 5:20 5:23 5:26 5:28 5:31 5:34 5:37

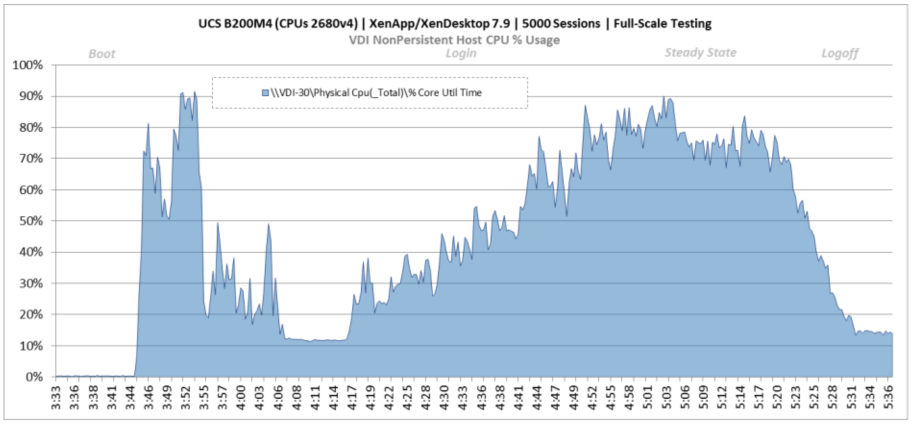

UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host CPU % Usage

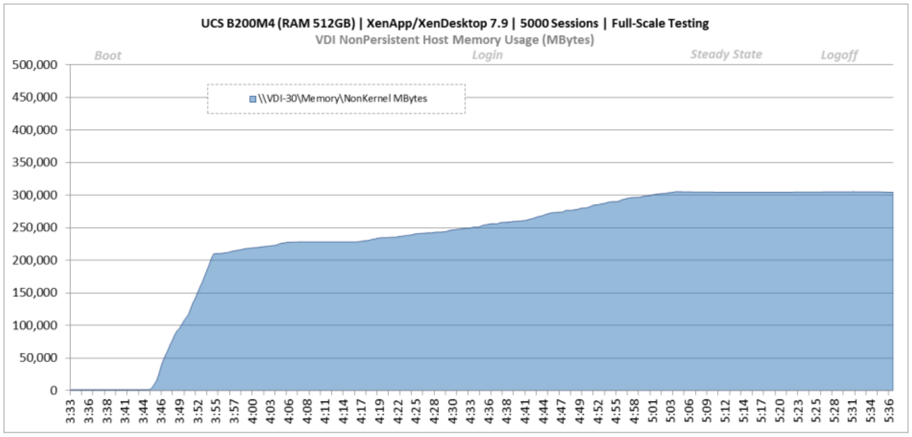

UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host Memory Usage (MBytes)

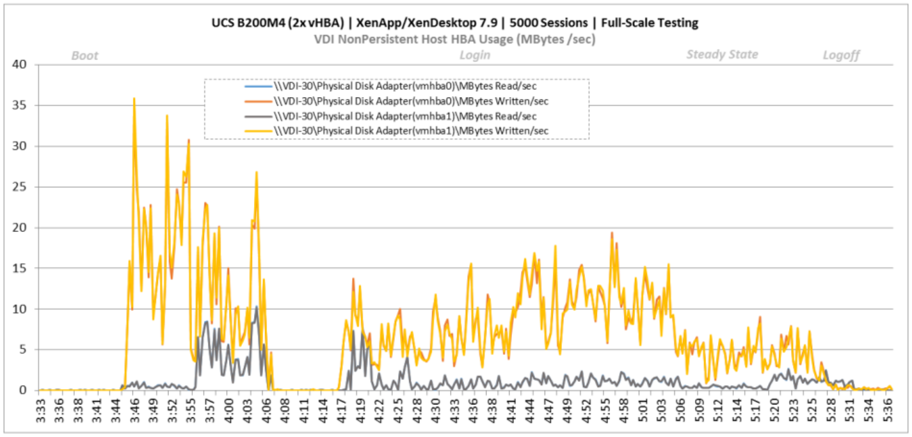

UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host HBA Usage (MBytes /sec)

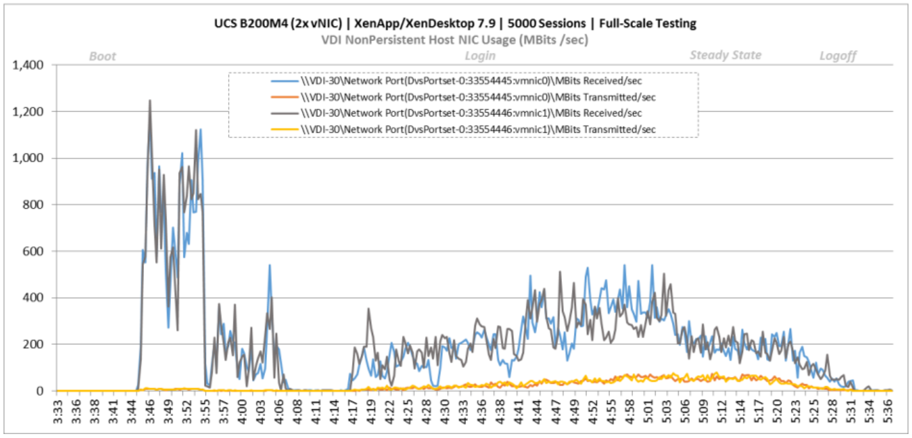

UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host NIC Usage (MBits /sec)

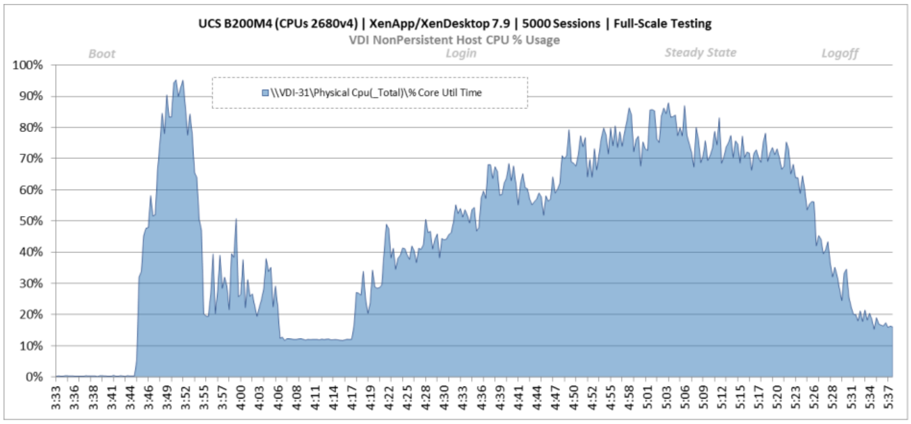

UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host CPU % Usage

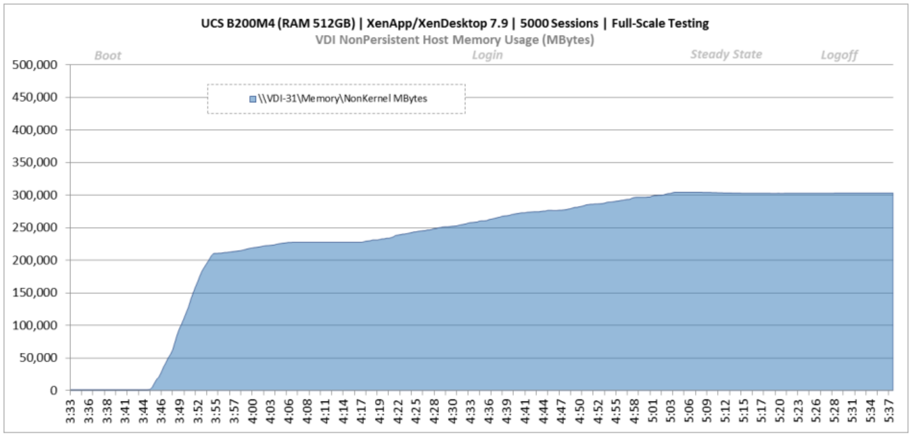

UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host Memory Usage (MBytes)

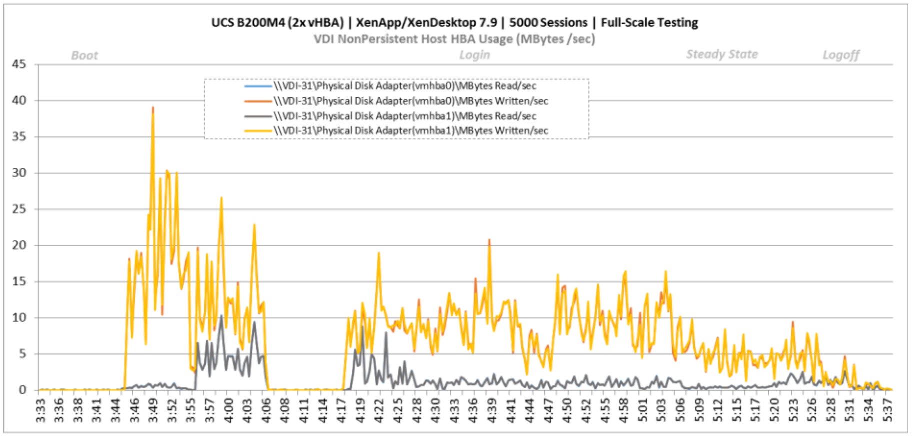

UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host HBA Usage (MBytes /sec)

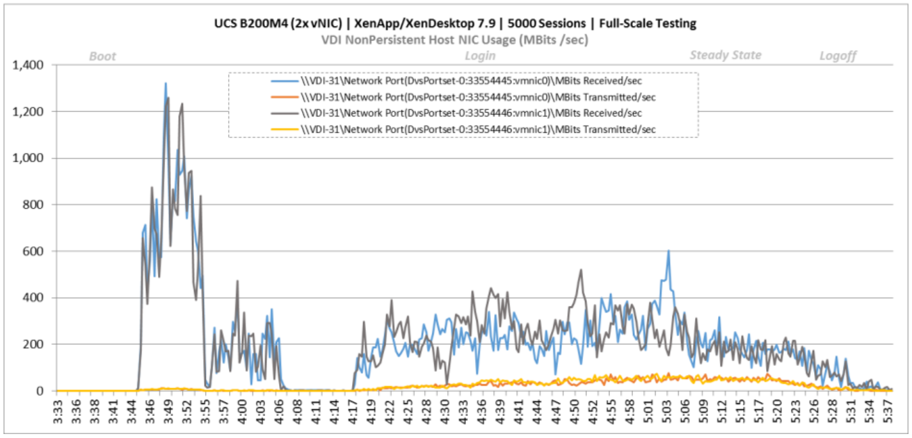

UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI NonPersistent Host NIC Usage (MBits /sec)

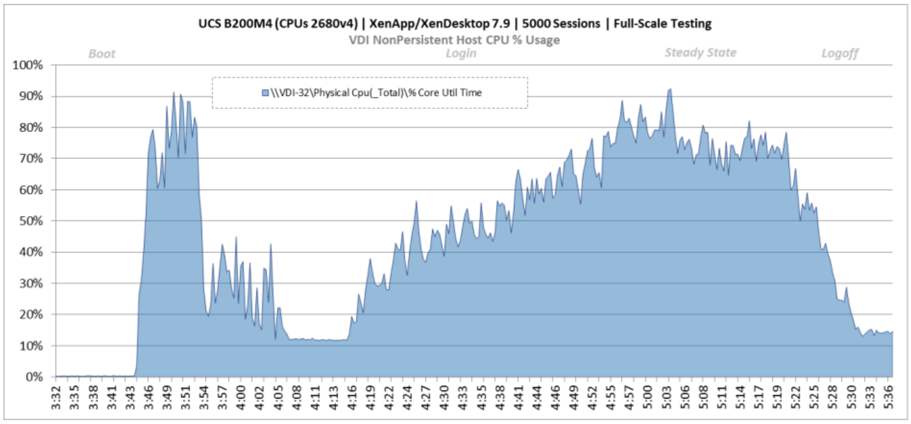
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host CPU % Usage
Boot
Login
Steady State
Logoff
■ \\VDI-32\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:32 3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 4:00 4:02 4:05 4:08 4:11 4:13 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:08 5:11 5:14 5:17 5:19 5:22 5:25 5:28 5:30 5:33 5:36

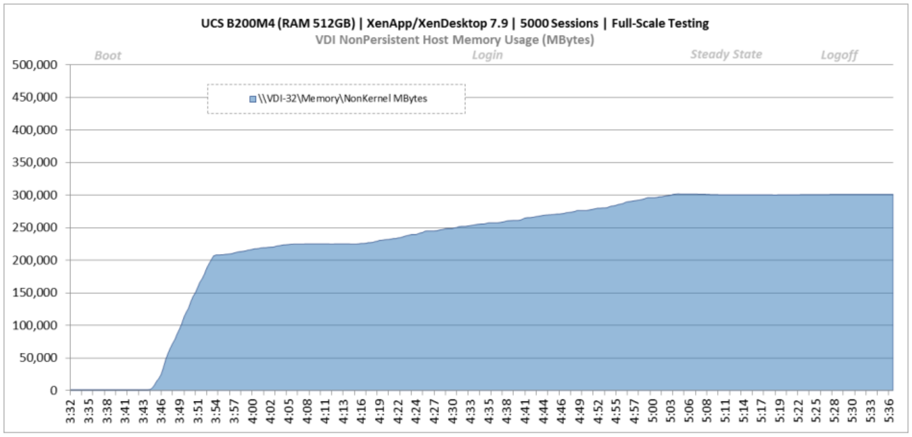
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
■ \\VDI-32\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:32 3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 4:00 4:02 4:05 4:08 4:11 4:13 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:08 5:11 5:14 5:17 5:19 5:22 5:25 5:28 5:30 5:33 5:36

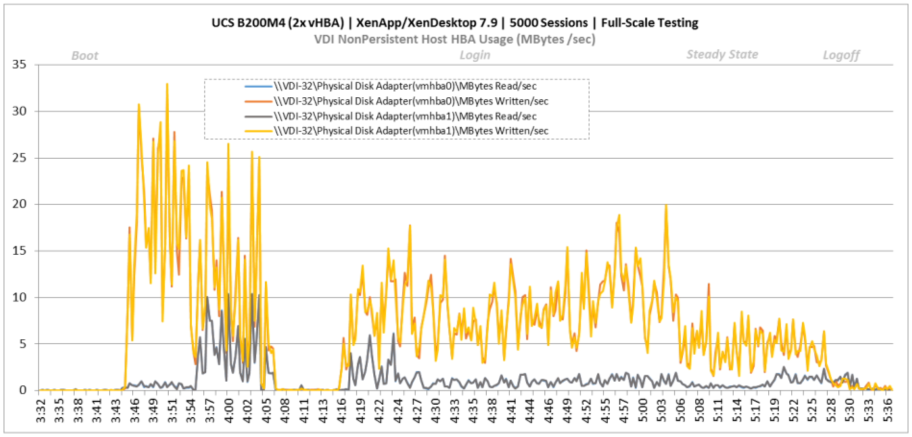
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-32\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-32\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-32\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-32\Physical Disk Adapter(vmhba1)\MBytes Written/sec
35
30
25
20
15
10
5
0
3:32 3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 4:00 4:02 4:05 4:08 4:11 4:13 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:08 5:11 5:14 5:17 5:19 5:22 5:25 5:28 5:30 5:33 5:36

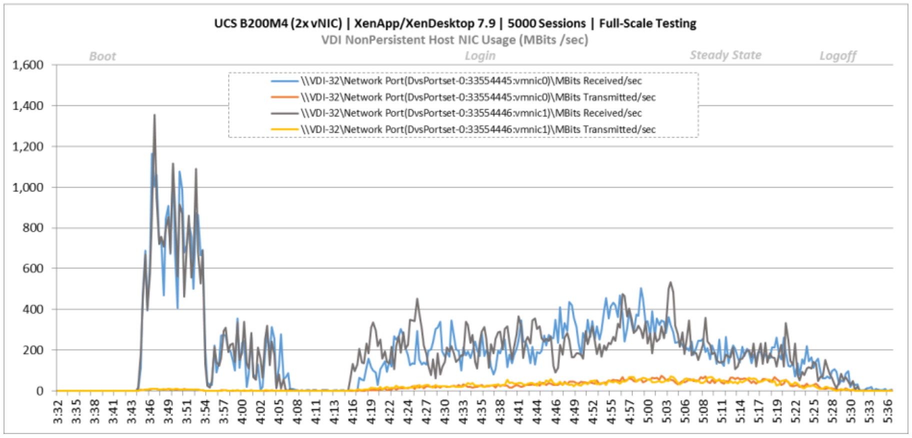
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-32\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-32\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-32\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-32\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
1,600
1,400
1,200
1,000
800
600
400
200
0
3:32 3:35 3:38 3:41 3:43 3:46 3:49 3:51 3:54 3:57 4:00 4:02 4:05 4:08 4:11 4:13 4:16 4:19 4:22 4:24 4:27 4:30 4:33 4:35 4:38 4:41 4:44 4:46 4:49 4:52 4:55 4:57 5:00 5:03 5:06 5:08 5:11 5:14 5:17 5:19 5:22 5:25 5:28 5:30 5:33 5:36

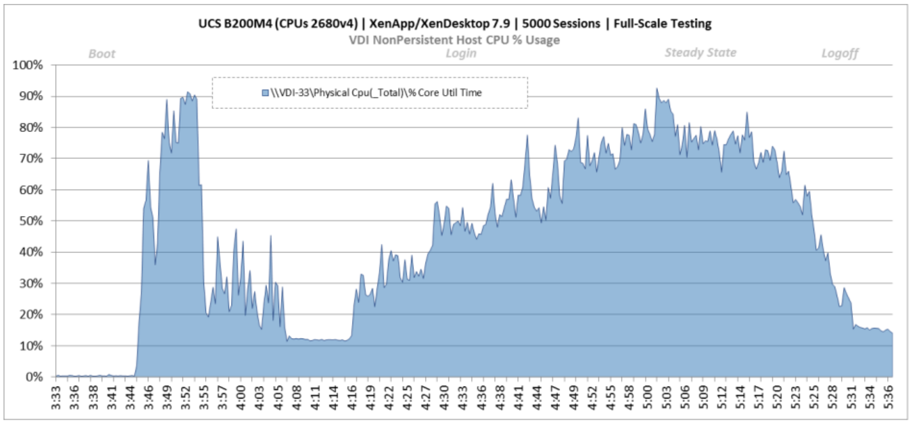
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-33\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:38 3:41 3:44 3:46 3:49 3:52 3:55 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:25 4:27 4:30 4:33 4:36 4:38 4:41 4:44 4:47 4:49 4:52 4:55 4:58 5:00 5:03 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:34 5:36

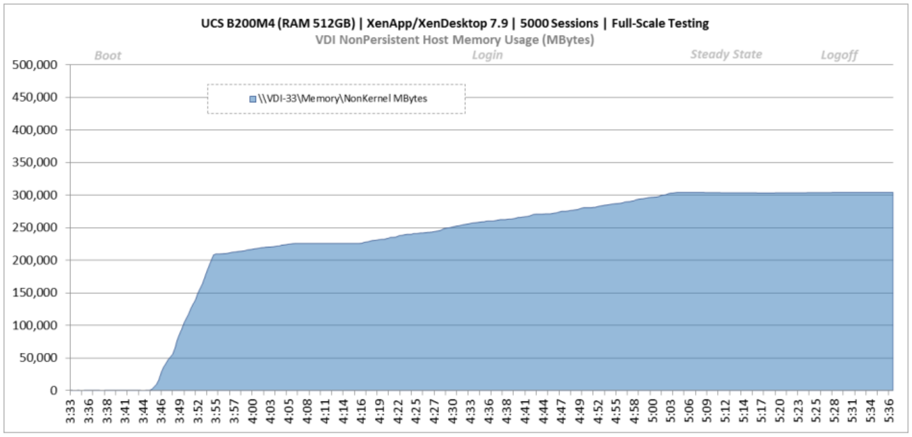
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-33\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:38 3:41 3:44 3:46 3:49 3:52 3:55 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:25 4:27 4:30 4:33 4:36 4:38 4:41 4:44 4:47 4:49 4:52 4:55 4:58 5:00 5:03 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:34 5:36

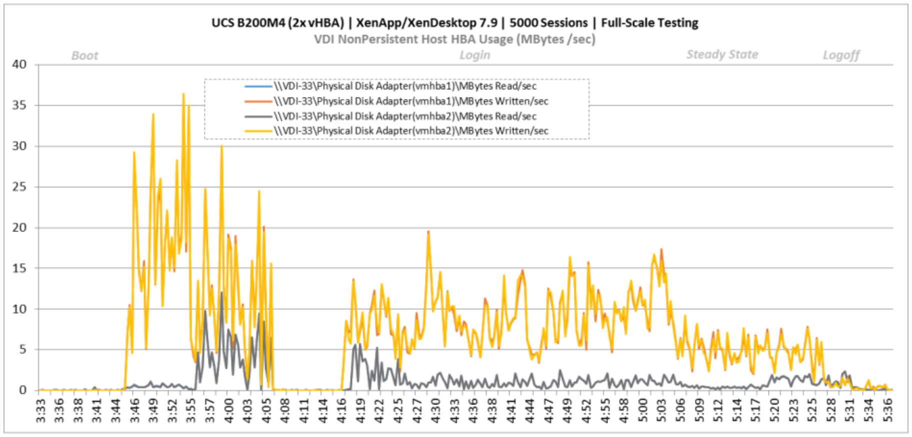
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-33\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-33\Physical Disk Adapter(vmhba1)\MBytes Written/sec
\\VDI-33\Physical Disk Adapter(vmhba2)\MBytes Read/sec
\\VDI-33\Physical Disk Adapter(vmhba2)\MBytes Written/sec
40
35
30
25
20
15
10
5
0
3:33 3:36 3:38 3:41 3:44 3:46 3:49 3:52 3:55 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:25 4:27 4:30 4:33 4:36 4:38 4:41 4:44 4:47 4:49 4:52 4:55 4:58 5:00 5:03 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:34 5:36

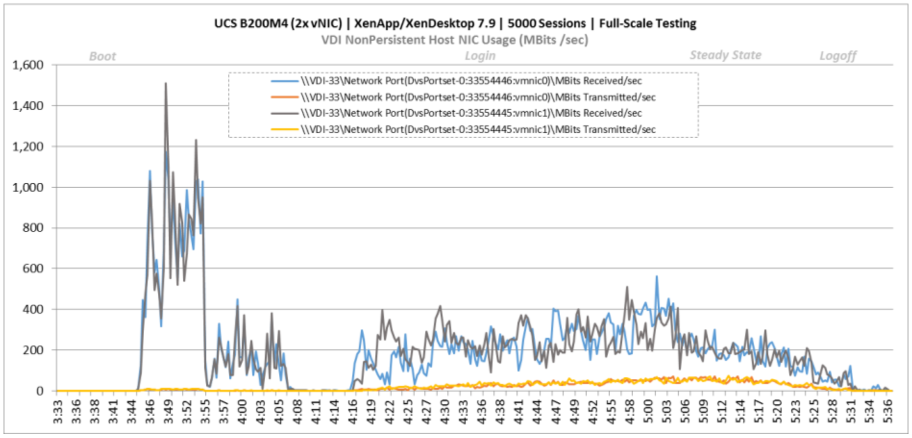
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI NonPersistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-33\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Received/sec
\\VDI-33\Network Port(DvsPortset-0:33554446:vmnic0)\MBits Transmitted/sec
\\VDI-33\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Received/sec
\\VDI-33\Network Port(DvsPortset-0:33554445:vmnic1)\MBits Transmitted/sec
1,600
1,400
1,200
1,000
800
600
400
200
0
3:33 3:36 3:38 3:41 3:44 3:46 3:49 3:52 3:55 3:57 4:00 4:03 4:05 4:08 4:11 4:14 4:16 4:19 4:22 4:25 4:27 4:30 4:33 4:36 4:38 4:41 4:44 4:47 4:49 4:52 4:55 4:58 5:00 5:03 5:06 5:09 5:12 5:14 5:17 5:20 5:23 5:25 5:28 5:31 5:34 5:36

XenDesktop Persistent Desktop Host Performance

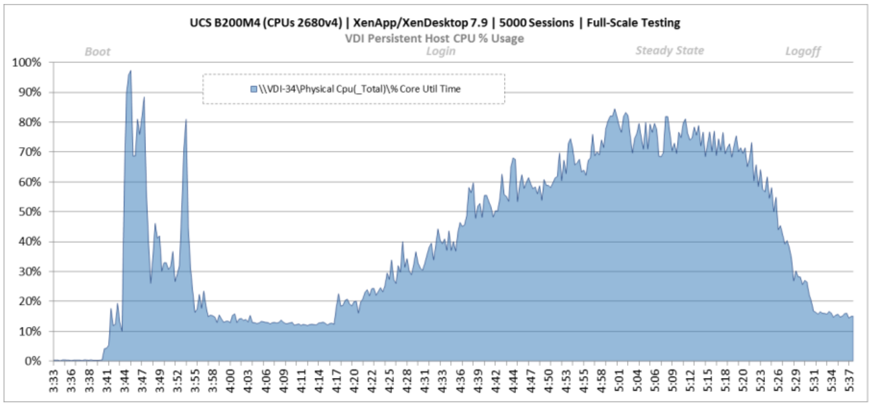

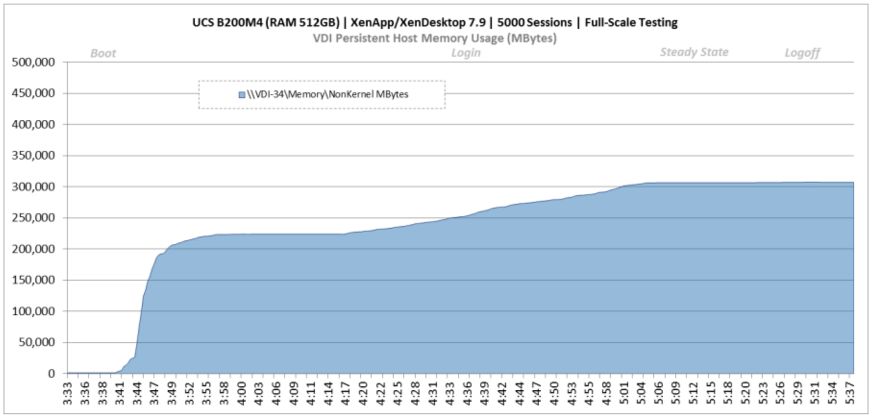

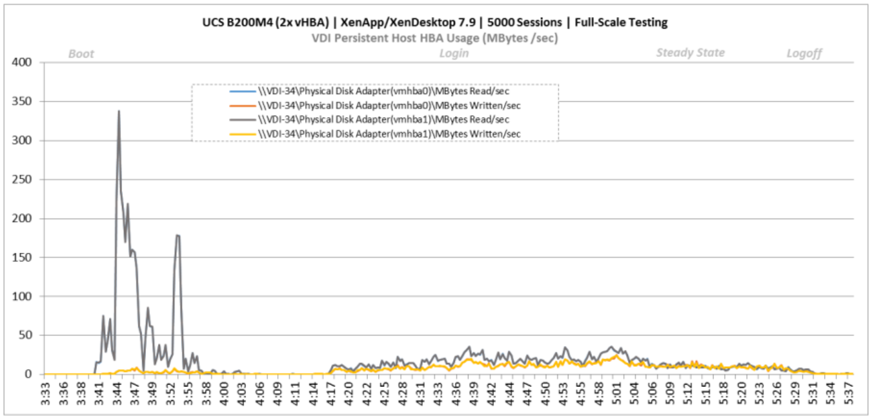

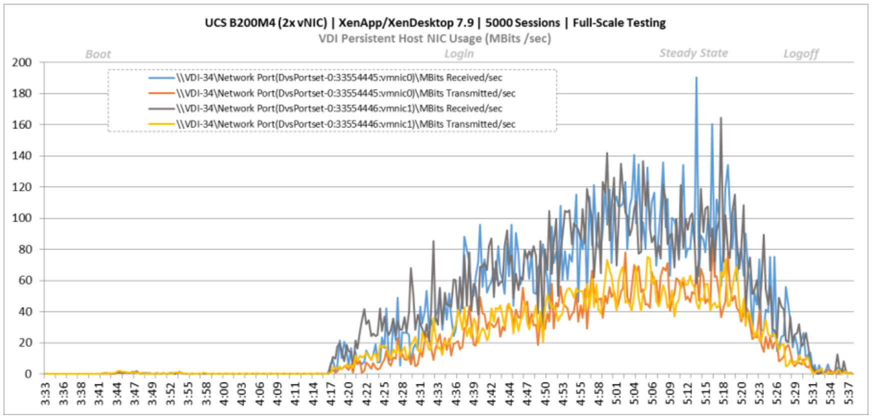

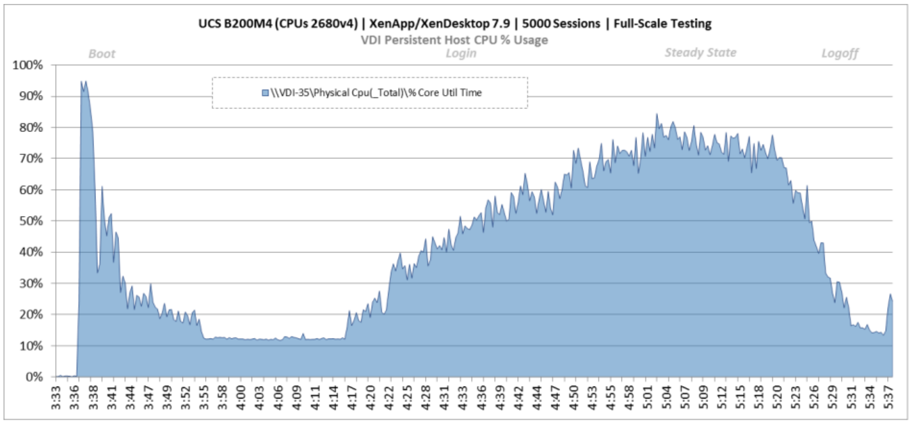
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-35\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33
3:36
3:38
3:41
3:44
3:47
3:49
3:52
3:55
3:58
4:00
4:03
4:06
4:09
4:11
4:14
4:17
4:20
4:22
4:25
4:28
4:31
4:33
4:36
4:39
4:42
4:44
4:47
4:50
4:53
4:55
4:58
5:01
5:04
5:07
5:09
5:12
5:15
5:18
5:20
5:23
5:26
5:29
5:31
5:34
5:37

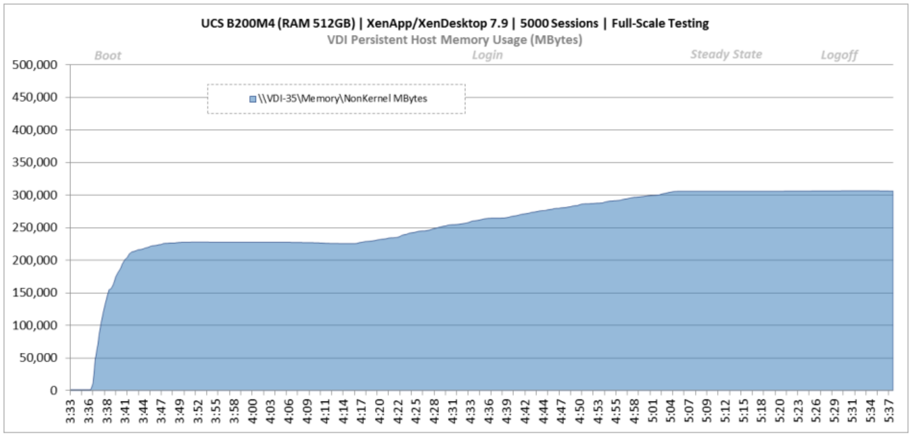
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-35\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33
3:36
3:38
3:41
3:44
3:47
3:49
3:52
3:55
3:58
4:00
4:03
4:06
4:09
4:11
4:14
4:17
4:20
4:22
4:25
4:28
4:31
4:33
4:36
4:39
4:42
4:44
4:47
4:50
4:53
4:55
4:58
5:01
5:04
5:07
5:09
5:12
5:15
5:18
5:20
5:23
5:26
5:29
5:31
5:34
5:37

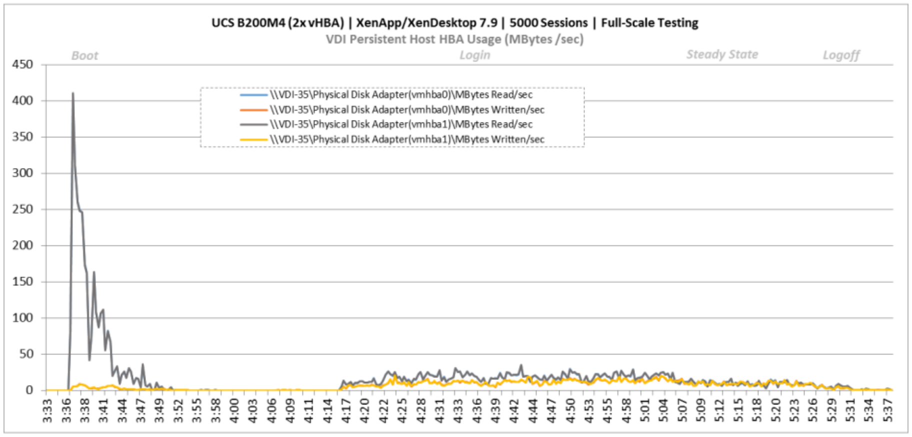
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-35\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-35\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-35\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-35\Physical Disk Adapter(vmhba1)\MBytes Written/sec
450
400
350
300
250
200
150
100
50
0
3:33
3:36
3:38
3:41
3:44
3:47
3:49
3:52
3:55
3:58
4:00
4:03
4:06
4:09
4:11
4:14
4:17
4:20
4:22
4:25
4:28
4:31
4:33
4:36
4:39
4:42
4:44
4:47
4:50
4:53
4:55
4:58
5:01
5:04
5:07
5:09
5:12
5:15
5:18
5:20
5:23
5:26
5:29
5:31
5:34
5:37

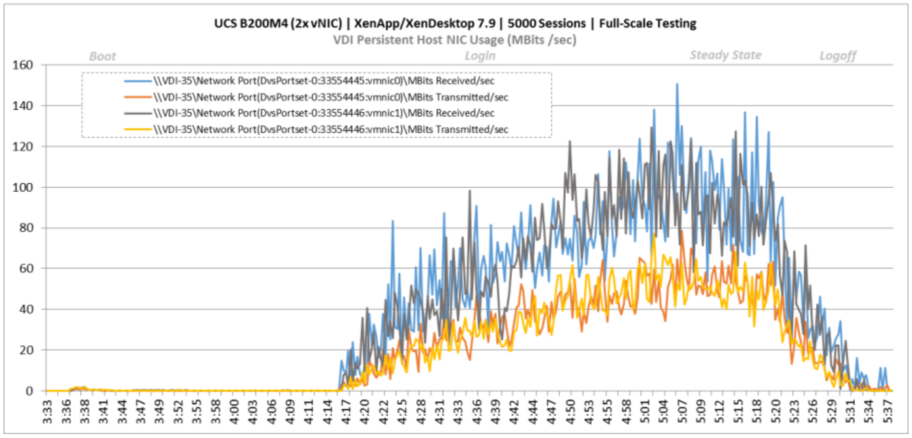
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-35\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-35\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-35\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-35\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
160
140
120
100
80
60
40
20
0
3:33
3:36
3:38
3:41
3:44
3:47
3:49
3:52
3:55
3:58
4:00
4:03
4:06
4:09
4:11
4:14
4:17
4:20
4:22
4:25
4:28
4:31
4:33
4:36
4:39
4:42
4:44
4:47
4:50
4:53
4:55
4:58
5:01
5:04
5:07
5:09
5:12
5:15
5:18
5:20
5:23
5:26
5:29
5:31
5:34
5:37

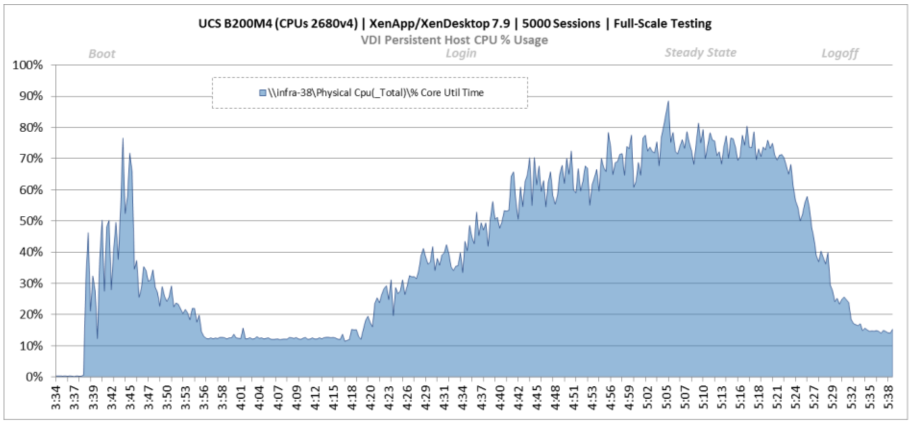
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host CPU % Usage

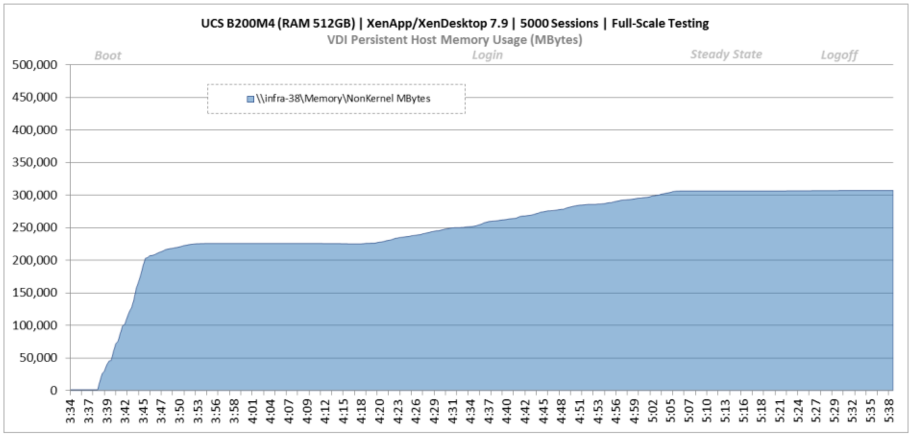
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host Memory Usage (MBytes)

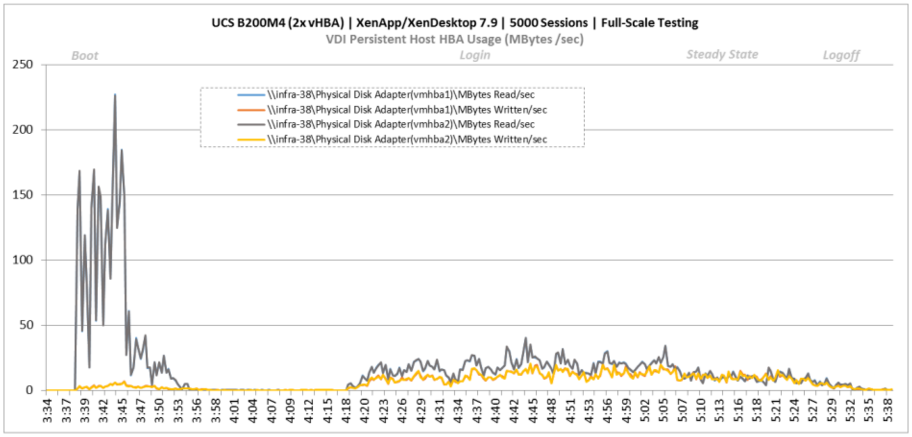
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host HBA Usage (MBytes /sec)

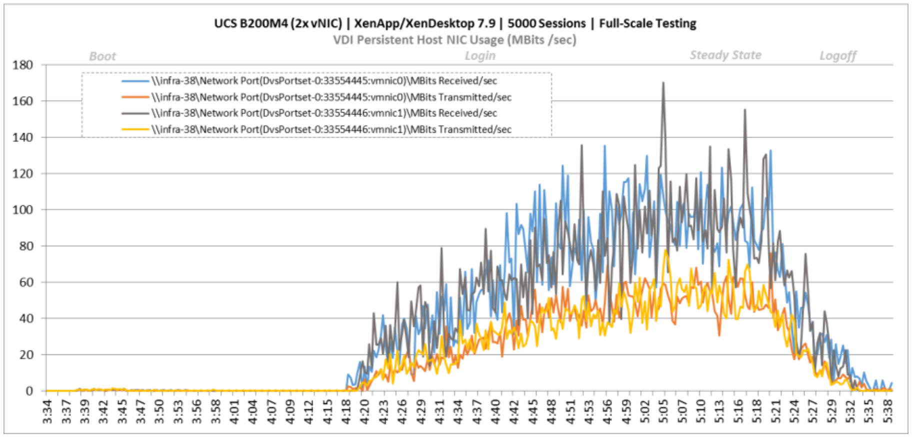
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host NIC Usage (MBits /sec)

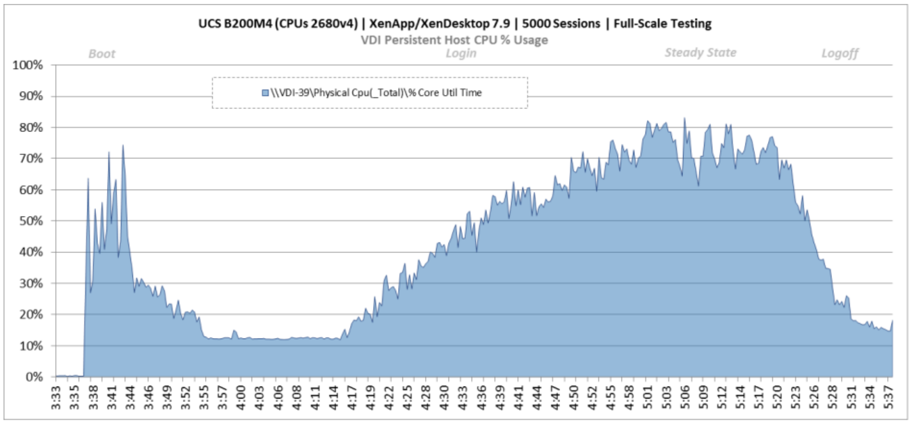

UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host CPU % Usage

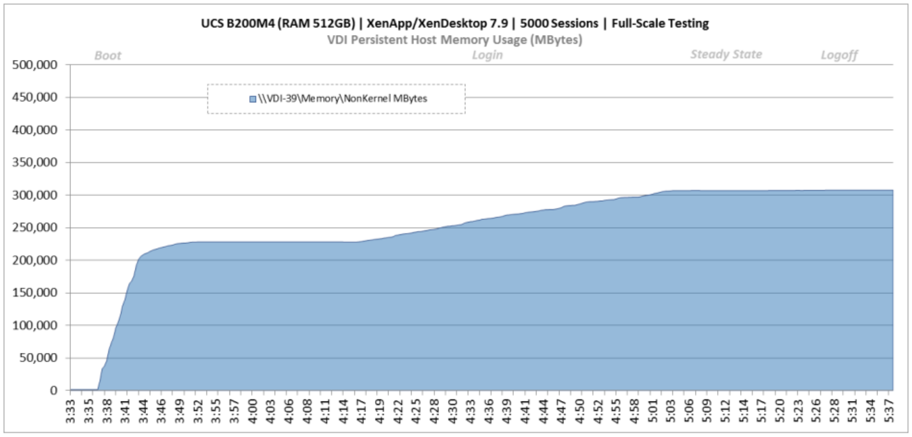

UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host Memory Usage (MBytes)

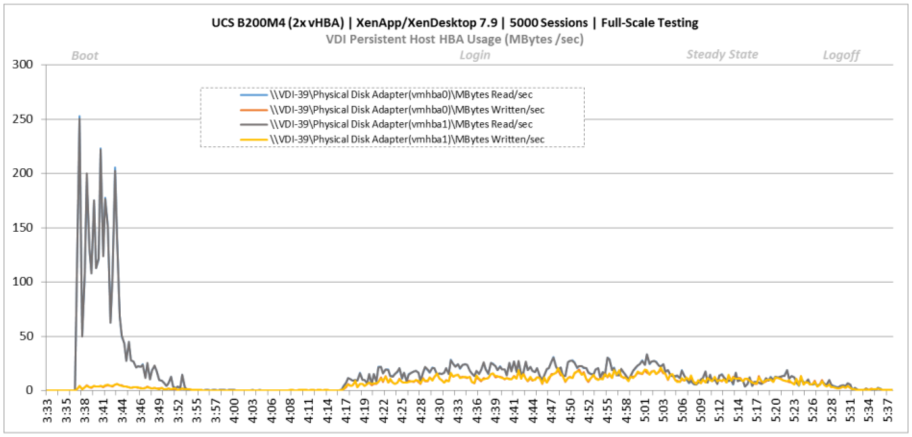

UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host HBA Usage (MBytes /sec)

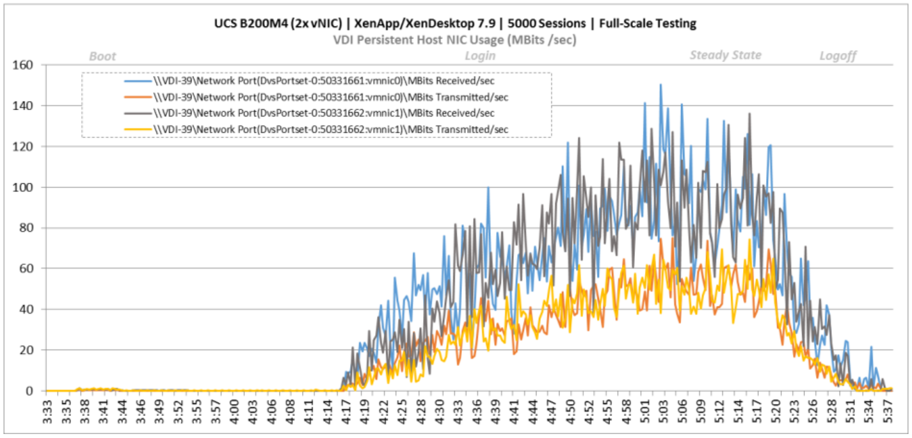

UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host NIC Usage (MBits /sec)

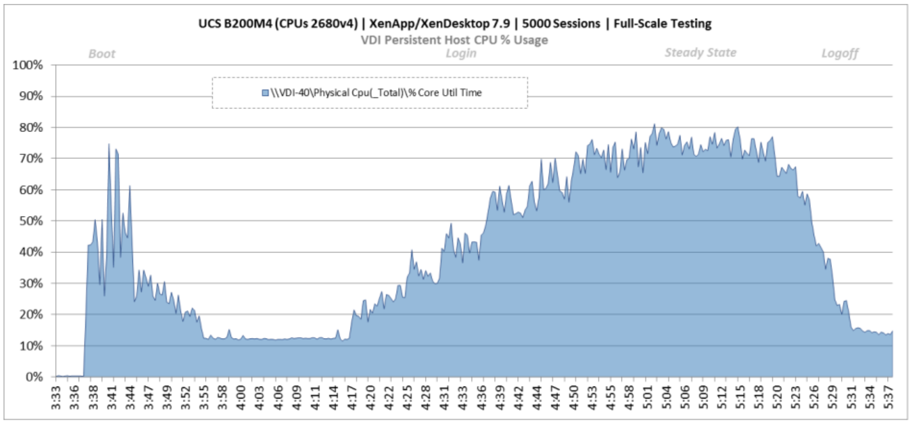
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-40\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:38 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

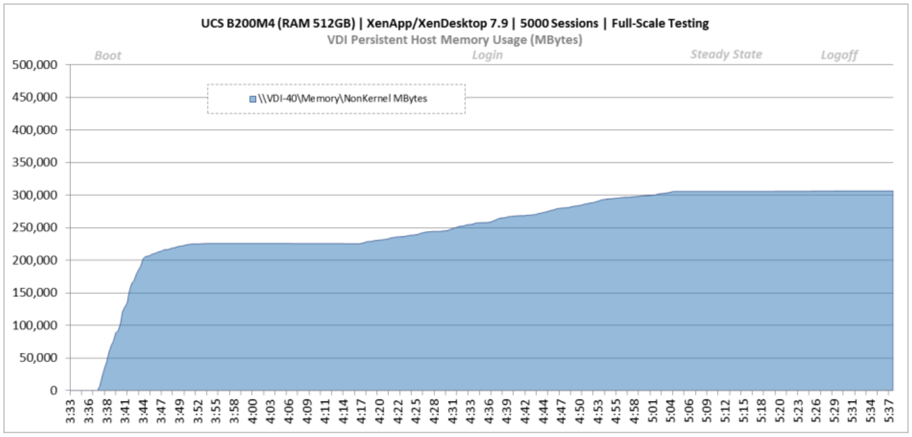
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-40\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:38 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

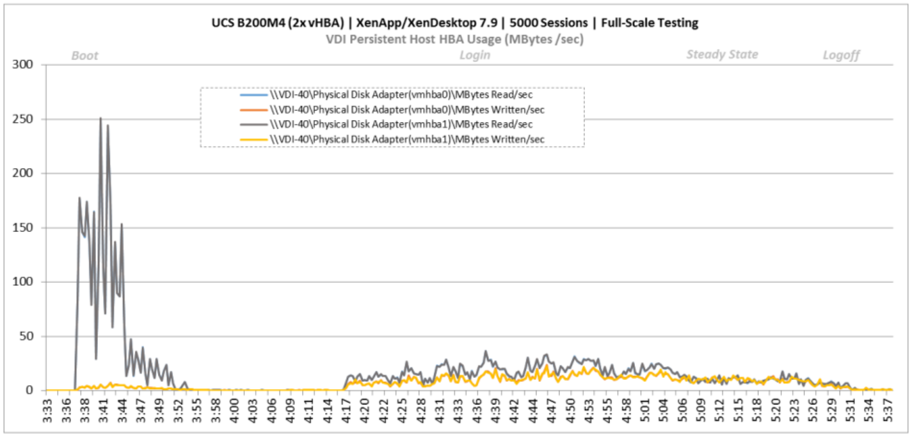
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-40\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-40\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-40\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-40\Physical Disk Adapter(vmhba1)\MBytes Written/sec
300
250
200
150
100
50
0
3:33 3:36 3:38 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

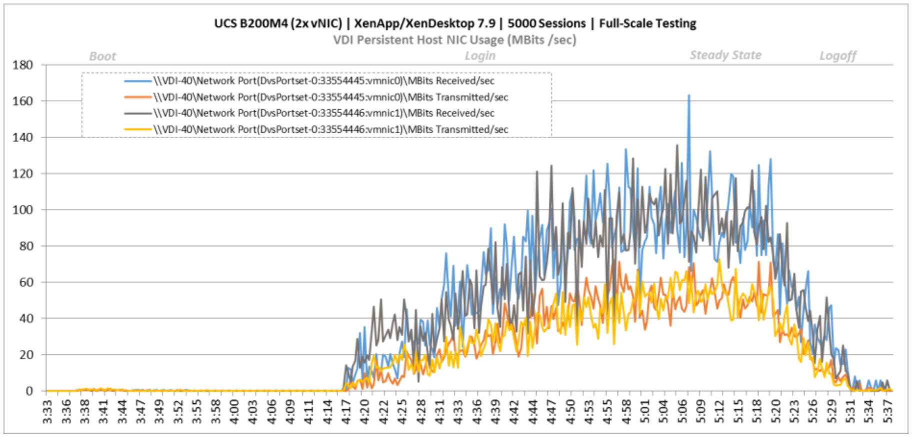
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-40\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-40\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-40\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-40\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
180
160
140
120
100
80
60
40
20
0
3:33 3:36 3:38 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:06 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

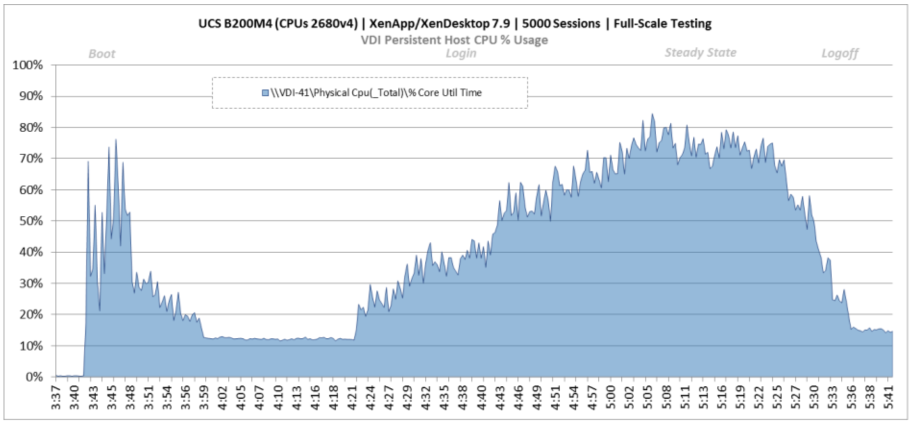
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host CPU % Usage

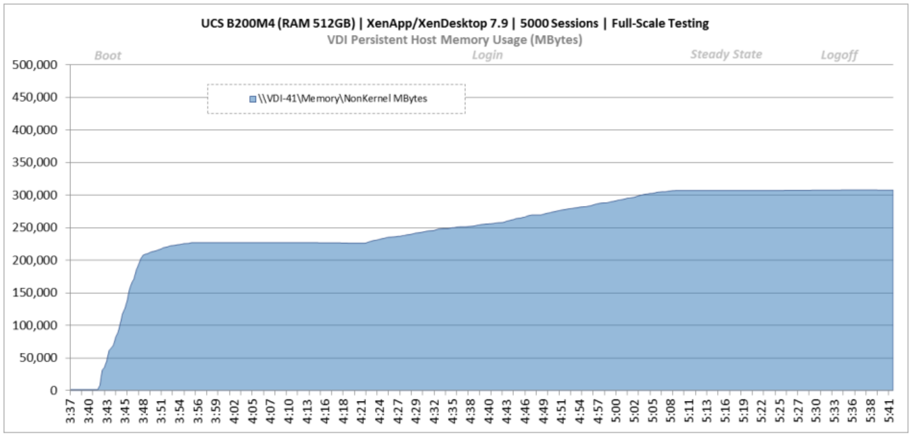
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host Memory Usage (MBytes)

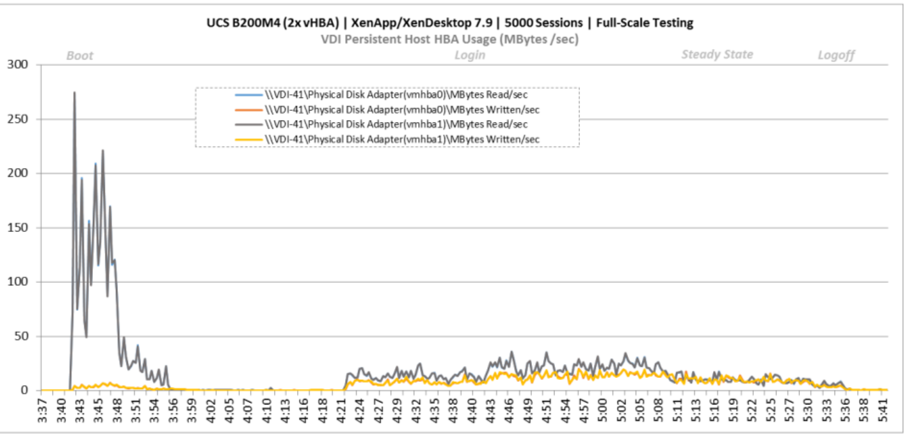
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host HBA Usage (MBytes /sec)

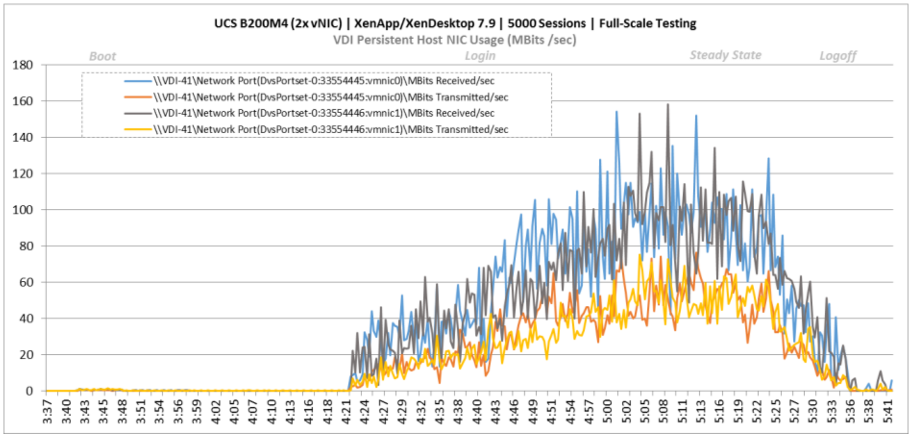
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing

VDI Persistent Host NIC Usage (MBits /sec)

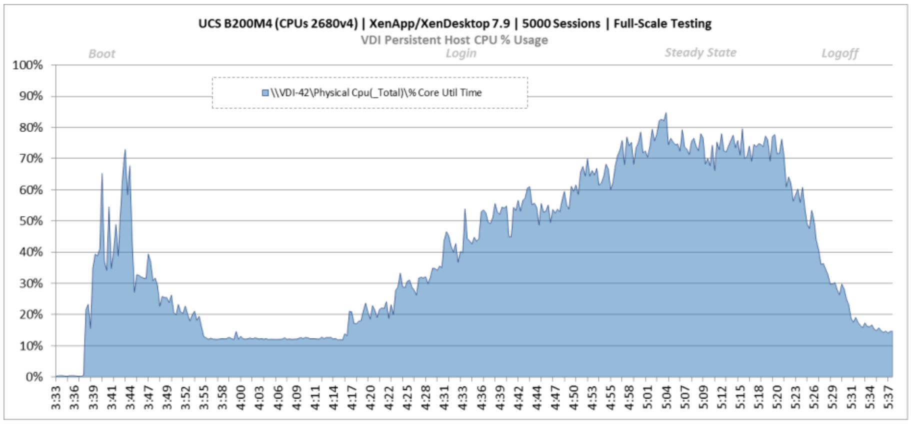
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-42\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

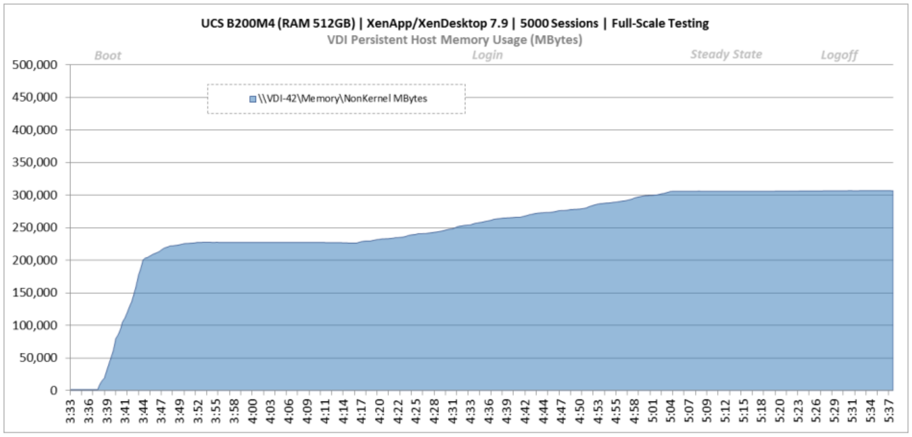
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-42\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

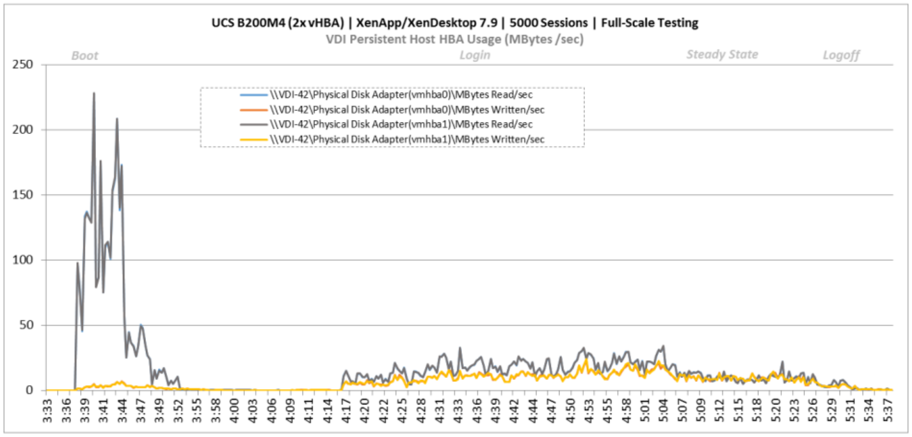
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-42\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-42\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-42\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-42\Physical Disk Adapter(vmhba1)\MBytes Written/sec
250
200
150
100
50
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

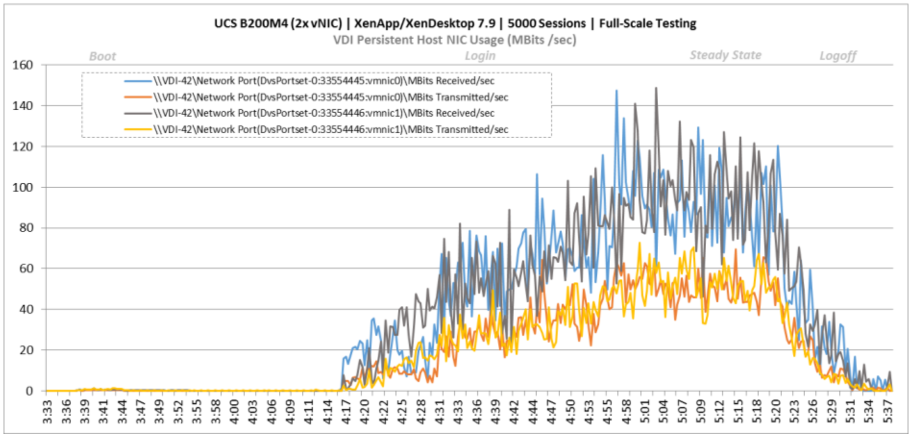
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-42\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-42\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-42\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-42\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
160
140
120
100
80
60
40
20
0
3:33 3:36 3:39 3:41 3:44 3:47 3:49 3:52 3:55 3:58 4:00 4:03 4:06 4:09 4:11 4:14 4:17 4:20 4:22 4:25 4:28 4:31 4:33 4:36 4:39 4:42 4:44 4:47 4:50 4:53 4:55 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

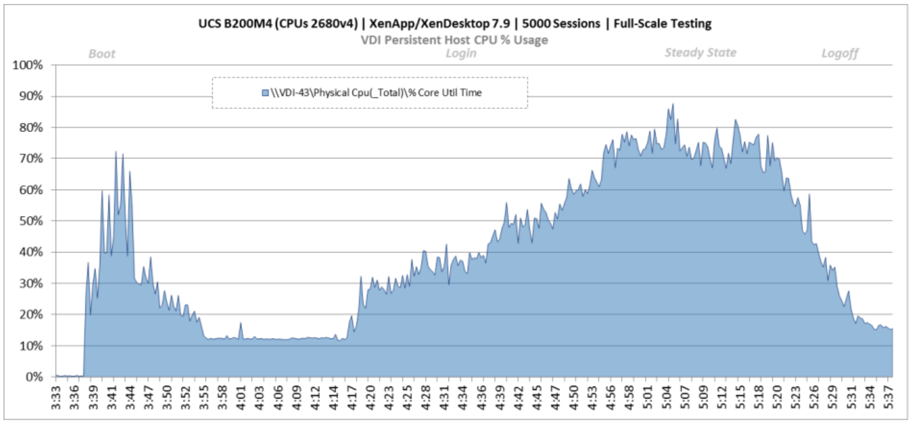
UCS B200M4 (CPUs 2680v4) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host CPU % Usage
Boot
Login
Steady State
Logoff
\\VDI-43\Physical Cpu(_Total)\% Core Util Time
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:01 4:03 4:06 4:09 4:12 4:14 4:17 4:20 4:23 4:25 4:28 4:31 4:34 4:36 4:39 4:42 4:45 4:47 4:50 4:53 4:56 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

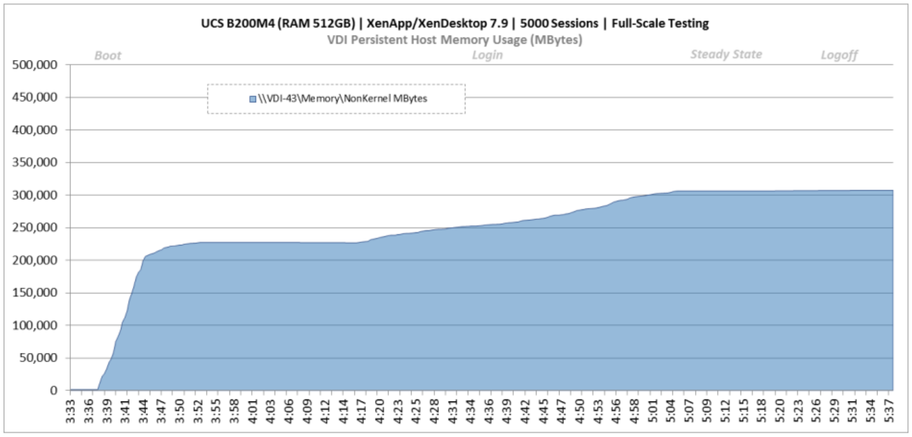
UCS B200M4 (RAM 512GB) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host Memory Usage (MBytes)
Boot
Login
Steady State
Logoff
\\VDI-43\Memory\NonKernel MBytes
500,000
450,000
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:01 4:03 4:06 4:09 4:12 4:14 4:17 4:20 4:23 4:25 4:28 4:31 4:34 4:36 4:39 4:42 4:45 4:47 4:50 4:53 4:56 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

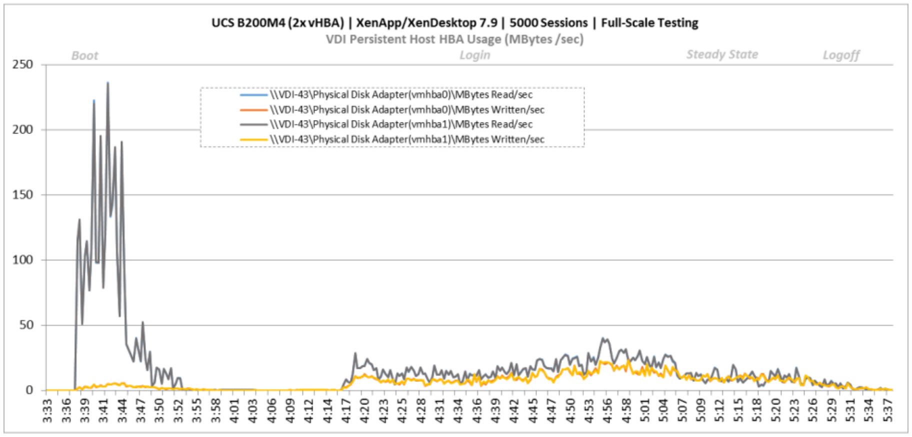
UCS B200M4 (2x vHBA) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host HBA Usage (MBytes /sec)
Boot
Login
Steady State
Logoff
\\VDI-43\Physical Disk Adapter(vmhba0)\MBytes Read/sec
\\VDI-43\Physical Disk Adapter(vmhba0)\MBytes Written/sec
\\VDI-43\Physical Disk Adapter(vmhba1)\MBytes Read/sec
\\VDI-43\Physical Disk Adapter(vmhba1)\MBytes Written/sec
250
200
150
100
50
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:01 4:03 4:06 4:09 4:12 4:14 4:17 4:20 4:23 4:25 4:28 4:31 4:34 4:36 4:39 4:42 4:45 4:47 4:50 4:53 4:56 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37

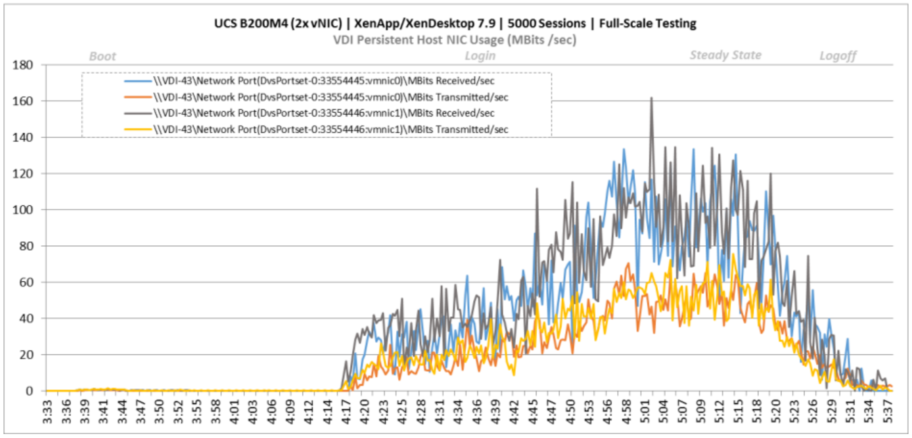
UCS B200M4 (2x vNIC) | XenApp/XenDesktop 7.9 | 5000 Sessions | Full-Scale Testing
VDI Persistent Host NIC Usage (MBits /sec)
Boot
Login
Steady State
Logoff
\\VDI-43\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Received/sec
\\VDI-43\Network Port(DvsPortset-0:33554445:vmnic0)\MBits Transmitted/sec
\\VDI-43\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Received/sec
\\VDI-43\Network Port(DvsPortset-0:33554446:vmnic1)\MBits Transmitted/sec
180
160
140
120
100
80
60
40
20
0
3:33 3:36 3:39 3:41 3:44 3:47 3:50 3:52 3:55 3:58 4:01 4:03 4:06 4:09 4:12 4:14 4:17 4:20 4:23 4:25 4:28 4:31 4:34 4:36 4:39 4:42 4:45 4:47 4:50 4:53 4:56 4:58 5:01 5:04 5:07 5:09 5:12 5:15 5:18 5:20 5:23 5:26 5:29 5:31 5:34 5:37