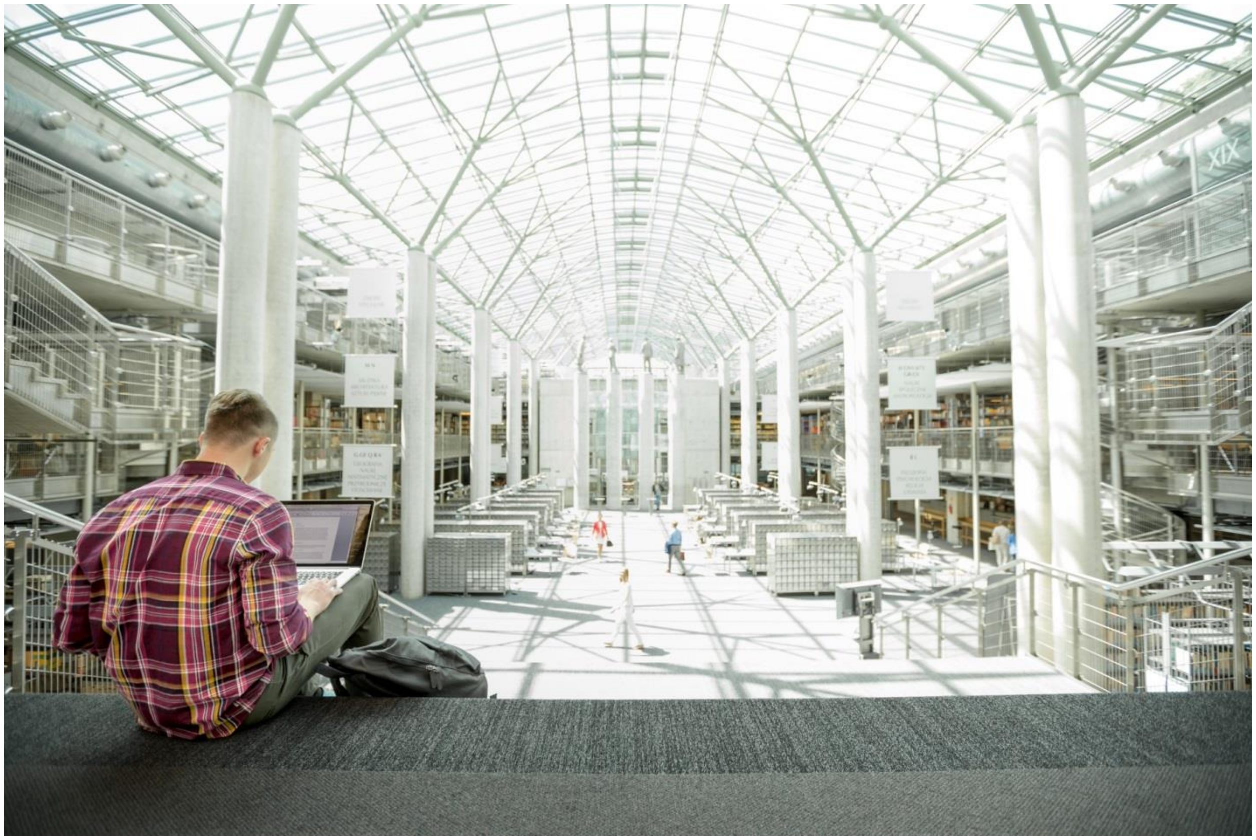

FlexPod Datacenter with VMware Horizon View 7.3 and VMware vSphere 6.5 Update 1 with Cisco UCS Manager 3.2 for 5000 Seats

Cisco Validated Design for a 5000 Seat Virtual Desktop Infrastructure Built on Cisco UCS B200 M5 and Cisco UCS Manager 3.2 with NetApp AFF A-Series on VMware Horizon View 7.3 and VMware vSphere ESXi 6.5 Update 1 Hypervisor Platform

Last Updated: January 2, 2018

About the Cisco Validated Design Program

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

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

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

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

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

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

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

This document provides a reference architecture, design guide, and deployment for up to a 5000 seat mixed workload end user computing environment on FlexPod Datacenter with Cisco UCS and NetApp® All Flash FAS (AFF) A300 storage. The solution includes VMware Horizon server-based Remote Desktop Sever Hosted sessions, VMware Horizon persistent Microsoft Windows 10 virtual desktops and VMware Horizon non-persistent Microsoft Windows 10 instant clone virtual desktops on VMware vSphere 6.5.

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

This solution is 100 percent virtualized on fifth generation Cisco UCS B200 M5 blade servers, booting VMware vSphere 6.5 Update 1 through FC SAN from the AFF A300 storage array. The virtual desktop sessions are powered by VMware Horizon 7. VMware Horizon Remote Desktop Server Hosted Sessions (1680 RDS Server sessions) and 1660 VMware Horizon Instant and 1660 Full Clones Window 10 desktops (3320 virtual desktops) were provisioned on the AFF A300 storage array. Where applicable the document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

This solution delivers the design 5000 user payload with five fewer blade servers than previous 5000 seat solutions on fourth generation Cisco UCS Blade Servers making it more efficient and cost effective in the data center.

The solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1.25.6 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, NetApp storage, and VMware have partnered to deliver this Cisco Validated Design, which uses best of breed storage, server and network components to serve as the foundation for desktop virtualization workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Audience

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

Purpose of this Document

This document provides a step-by-step design, configuration and implementation guide for the Cisco Validated Design for a large-scale VMware Horizon 7 mixed workload solution with NetApp AFF A300, Cisco UCS Blade Servers, Cisco Nexus 9000 series ethernet switches and Cisco MDS 9000 series fibre channel switches.

What's New?

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

It incorporates the following features:

- Cisco UCS B200 M5 blade servers with Intel Xeon Scalable Family processors and 2666 MHz memory
- Validation of Cisco Nexus 9000 with NetApp AFF A300 system
- Validation of Cisco MDS 9000 with NetApp AFF A300 system
- Support for the Cisco UCS 3.2(1d) release and Cisco UCS B200-M5 servers
- Support for the latest release of NetApp AFF A300 hardware and NetApp ONTAP® 9.1
- A Fibre Channel storage design supporting SAN LUNs
- Cisco UCS Inband KVM Access
- Cisco UCS vMedia client for vSphere Installation
- Cisco UCS Firmware Auto Sync Server policy

- VMware vSphere 6.5 U1 Hypervisor
- VMware Horizon 7 Server 2016 RDS Hosted server sessions
- VMware Horizon 7 non-persistent Instant Clone Windows 10 virtual machines
- VMware Horizon 7 persistent Full Clones Windows 10 virtual machines
- The data center market segment is shifting toward heavily virtualized private, hybrid and public cloud computing models running on industry-standard systems. These environments require uniform design points that can be repeated for ease of management and scalability.

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

The use cases include:

- Enterprise Datacenter
- Service Provider Datacenter
- Large Commercial Datacenter

Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both VMware Horizon RDSH server desktop sessions based on Microsoft Server 2016, VMware Horizon VDI persistent virtual machines and VMware Horizon VDI non–persistent virtual machines based on Windows 10 operating system.

The mixed workload solution includes NetApp AFF A300 storage, Cisco Nexus® and MDS networking, the Cisco Unified Computing System (Cisco UCS®), VMware Horizon and VMware vSphere software in a single package. The design is space optimized such that the network, compute, and storage required can be housed in one 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., NetApp Inc. and VMware Inc. produced a highly efficient, robust and affordable desktop virtualization solution for a hosted virtual desktop and hosted shared desktop mixed deployment supporting different use cases. Key components of the solution include the following:

- More power, same size. Cisco UCS B200 M5 half-width blade with dual 18-core 2.3 GHz Intel ® Xeon ® Gold (6140) processors and 768 GB of memory for VMware Horizon Desktop hosts supports more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel 18-core 2.3 GHz Intel ® Xeon ® Gold (6140) processors used in this study provided a balance between increased per-blade capacity and cost.

- Fewer servers. Because of the increased compute power in the Cisco UCS B200 M5 servers, we supported the 5000 seat design with 16% fewer servers compared to previous generation B200 M4s.

- Fault-tolerance with high availability built into the design. The various designs are based on using one Unified Computing System chassis with multiple Cisco UCS B200 M5 blades for virtualized desktop and infrastructure workloads. The design provides N+1 server fault tolerance for hosted virtual desktops, hosted shared desktops and infrastructure services.

- Stress-tested to the limits during aggressive boot scenario. The 5000-user mixed RDS hosted virtual sessions and VDI pooled shared desktop environment booted and registered with the VMware Horizon 7 Administrator in under 20 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.

- Stress-tested to the limits during simulated login storms. All 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 Cisco Validated Design (CVD) presents a validated design that is 100 percent virtualized on VMware ESXi 6.5. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, Provisioning Servers, SQL Servers, VMware Horizon Connection Servers, VMware Horizon Composer Server, VMware Horizon Replica Servers, VMware Horizon Remote Desktop Server Hosted sessions and VDI virtual machine desktops. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the FlexPod converged infrastructure with stateless Cisco UCS Blade servers and NetApp FC storage.

- Cisco maintains industry leadership with the new Cisco UCS Manager 3.2(1d) 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 40G unified fabric story gets additional validation on Cisco UCS 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.

- NetApp AFF A300 array provides industry-leading storage solutions that efficiently handle the most demanding I/O bursts (for example, login storms), profile management, and user data management, deliver simple and flexible business continuance, and help reduce storage cost per desktop.

- NetApp AFF A300 array provides a simple to understand storage architecture for hosting all user data components (VMs, profiles, user data) on the same storage array.

- NetApp clustered Data ONTAP software enables to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.

- NetApp Virtual Storage Console (VSC) for VMware vSphere hypervisor has deep integrations with vSphere, providing easy-button automation for key storage tasks such as storage repository provisioning, storage resize, data deduplication, directly from vCenter.

- VMware Horizon 7. Latest and greatest virtual desktop and application product. VMware Horizon 7 follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new VMware Horizon release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of users increase. In addition, Horizon enhancements help to optimize performance and improve the user experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.

- Optimized to achieve the best possible performance and scale. For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the VMware 7 RDS virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.

- Provisioning desktop machines made easy. Remote Desktop Server Hosted (RDSH) shared virtual machines and VMware Horizon 7, Microsoft Windows 10 virtual machines were created for this solution using VMware Instant and Composer pooled desktops.

Cisco Desktop Virtualization Solutions: 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

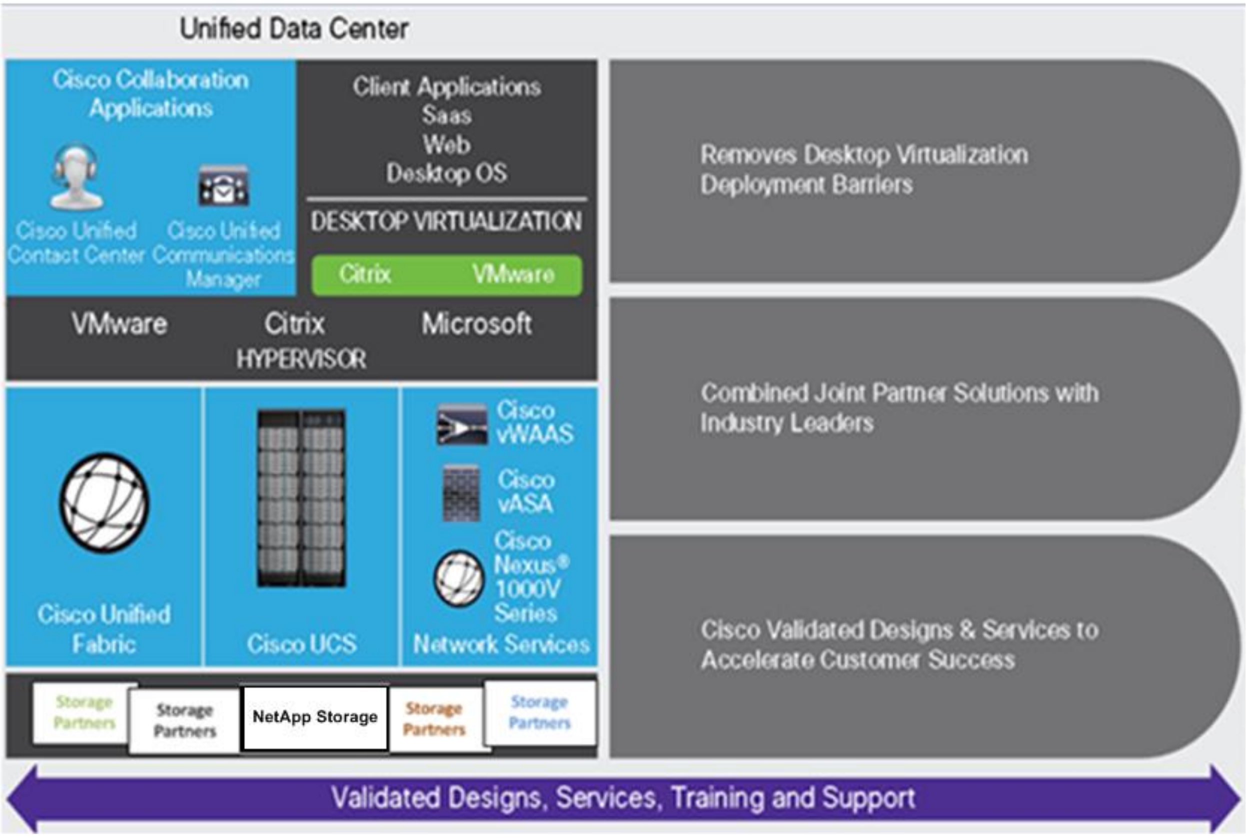
Unified Data Center
Cisco Collaboration Applications
Cisco Unified Contact Center
Cisco Unified Communications Manager
Client Applications
Saas
Web
Desktop OS
DESKTOP VIRTUALIZATION
Citrix
VMware
VMware
Citrix
Microsoft
HYPERVISOR
Cisco Unified Fabric
Cisco UCS
Cisco vWAAS
Cisco vASA
Cisco Nexus® 1000V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

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

Cisco Desktop Virtualization Focus

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

Simplified

Cisco UCS provides a radical new approach to industry-standard computing and provides the core of the data center infrastructure for desktop virtualization. Among the many features and benefits of Cisco UCS are the drastic reduction in the number of servers needed and in the number of cables used per server, and the capability to rapidly deploy or reprovision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco UCS Manager service profiles and Cisco storage partners' storage-based cloning. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco UCS Manager (UCSM) automates many mundane, error-prone data center operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS B-Series Blade Servers and C-Series Rack Servers with large memory footprints enable high desktop density that helps reduce server infrastructure requirements.

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

Secure

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

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

Scalable

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

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 1 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partners NetApp help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs for End User Computing based on FlexPod solutions have demonstrated scalability and performance, with up to 5000 desktops up and running in 20 minutes.

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

Savings and Success

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

The simplified deployment of Cisco UCS for desktop virtualization accelerates the time to productivity and enhances business agility. IT staff and end users are more productive more quickly, and the business can respond to new opportunities quickly by deploying virtual desktops whenever and wherever they are needed. The high-performance Cisco systems and network deliver a near-native end-user experience, allowing users to be productive anytime and anywhere.

The ultimate measure of desktop virtualization for any organization is its efficiency and effectiveness in both the near term and the long term. The Cisco Desktop Virtualization Solutions are very efficient, allowing rapid deployment, requiring fewer devices and cables, and reducing costs. The solutions are also very effective, providing the services that end users need on their devices of choice while improving IT operations, control, and data security. Success is bolstered through Cisco's best-in-class partnerships with leaders in virtualization and storage, and through tested and validated designs and services to help customers throughout the solution lifecycle. Long-term success is enabled through the use of Cisco's scalable, flexible, and secure architecture as the platform for desktop virtualization.

Physical Topology

Figure 2 illustrates the physical architecture.

Figure 2   Physical Architecture

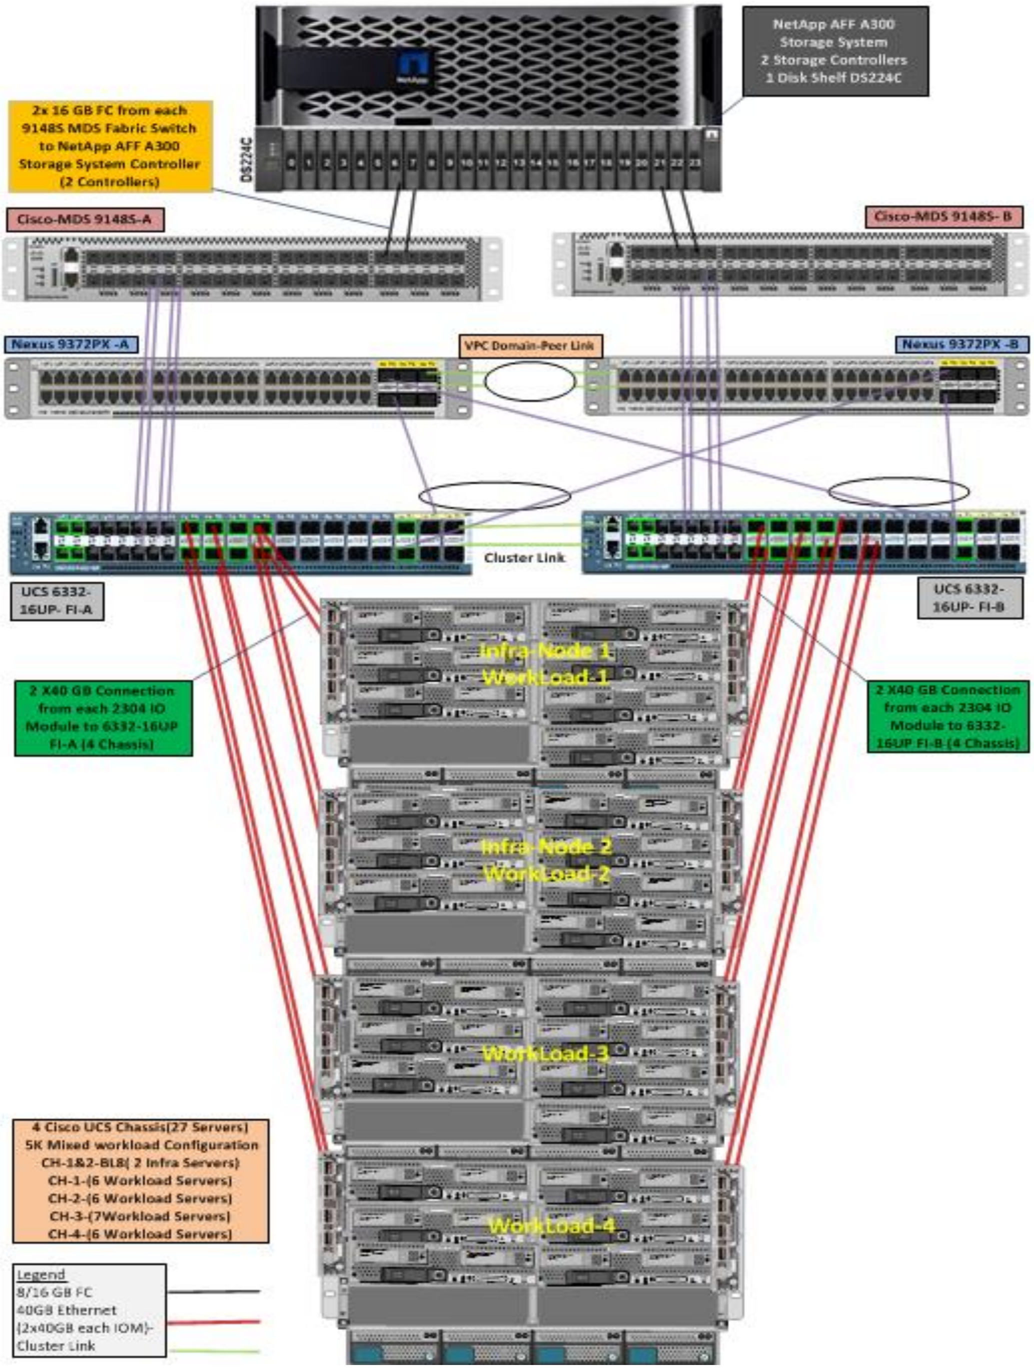

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches
- Two Cisco MDS 9148S 16GB Fibre Channel switches
- Two Cisco UCS 6332-16UP Fabric Interconnects
- Four Cisco UCS 5108 Blade Chassis

- Two Cisco UCS B200 M4 Blade Servers (2 Infra Server hosting Infrastructure VMs)
- 25 Cisco UCS B200 M5 Blade Servers (25 servers for workload)
- One NetApp AFF A300 Storage System
- One NetApp DS224C Disk Shelf

For desktop virtualization, the deployment includes VMware Horizon 7 running on VMware vSphere 6.5.

The design is intended to provide a large-scale building block for both VMware Horizon RDS Hosted server sessions and Windows 10 non-persistent and persistent VDI desktops in the following ratio:

- 1680 Remote Desktop Server Hosted (RDSH) desktop sessions
- 1660 VMware Horizon Windows 10 non-persistent virtual desktops
- 1660 VMware Horizon Windows 10 persistent virtual desktops

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

Configuration Guidelines

This Cisco Validated Design provides details for deploying a fully redundant, highly available 5000 seat mixed workload virtual desktop solution with VMware on a FlexPod Datacenter architecture. Configuration guidelines are provided that refer the reader to which redundant component is being configured with each step. For example, storage controller 01and storage controller 02 are used to identify the two AFF A300 storage controllers that are provisioned with this document, Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured and Cisco MDS A or Cisco MDS B identifies the pair of Cisco MDS switches that are configured.

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

Solution Components

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

What is FlexPod?

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

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

Figure 3 FlexPod Component Families

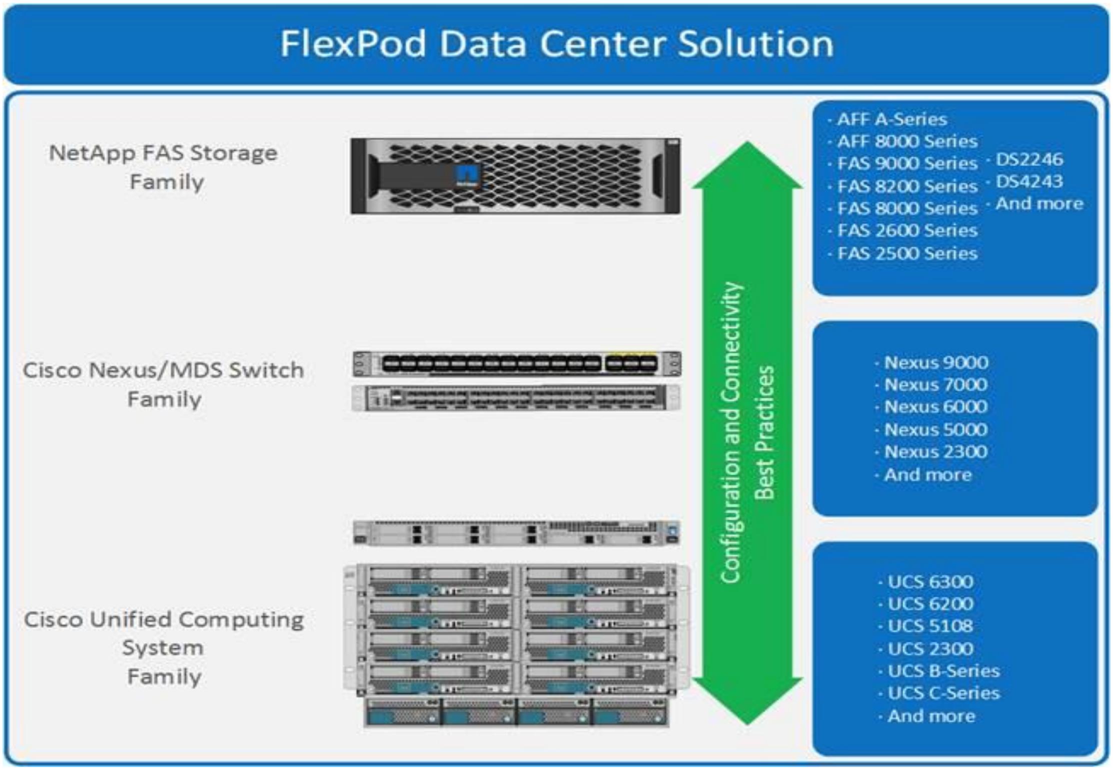

These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually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

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

Why FlexPod?

The following lists the benefits of FlexPod:

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

Cisco Unified Computing System

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

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

Cisco Unified Computing System Components

The main components of Cisco UCS are:

- Compute: The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® processor E5-2600/4600 v3 and E7-2800 v3 family CPUs.
- Network: The system is integrated on a low-latency, lossless, 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables needed, and by decreasing the power and cooling requirements.
- Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access: The system provides consolidated access to local storage, SAN storage, and network-attached storage (NAS) over the unified fabric. With storage access unified, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Small Computer System Interface over IP (iSCSI) protocols. This capability provides customers with choice for storage access and investment protection. In addition, server administrators can preassign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management and helping increase productivity.
- Management: Cisco UCS uniquely integrates all system components, enabling the entire solution to be managed as a single entity by Cisco UCS Manager. The manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

Figure 4 Cisco Data Center Overview

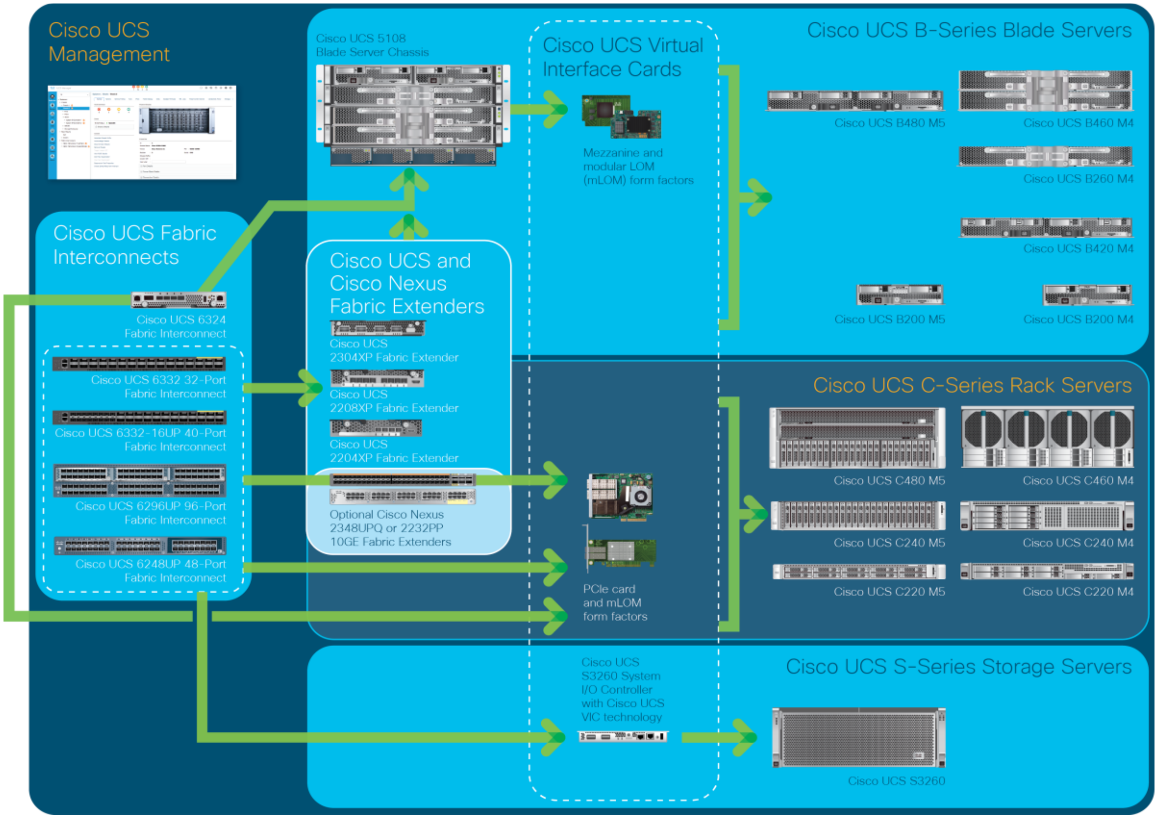
Cisco UCS Management
Cisco UCS Fabric Interconnects
Cisco UCS 6324 Fabric Interconnect
Cisco UCS 6332 32-Port Fabric Interconnect
Cisco UCS 6332-16UP 40-Port Fabric Interconnect
Cisco UCS 6296UP 96-Port Fabric Interconnect
Cisco UCS 6248UP 48-Port Fabric Interconnect
Cisco UCS 5108 Blade Server Chassis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

Cisco UCS is designed to deliver:

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

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

Cisco UCS Fabric Interconnect

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

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

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

Figure 5 Cisco UCS 6300 Series Fabric Interconnect

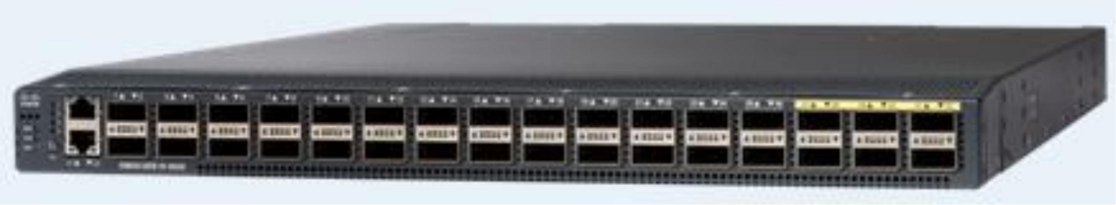

Cisco UCS B200 M5 Blade Server

The Cisco UCS B200 M5 Blade Server (Figure 6 and Figure 7) is a density-optimized, half-width blade server that supports two CPU sockets for Intel Xeon processor 6140 Gold series CPUs and up to 24 DDR4 DIMMs. It supports one modular LAN-on-motherboard (LOM) dedicated slot for a Cisco virtual interface card (VIC) and one mezzanine adapter. In additions, the Cisco UCS B200 M5 supports an optional storage module that accommodates up to two SAS or SATA hard disk drives (HDDs) or solid-state disk (SSD) drives. You can install up to eight Cisco UCS B200 M5 servers in a chassis, mixing them with other models of Cisco UCS blade servers in the chassis if desired. Latest features of Cisco UCS Virtual Interface Cards (VICs)

Figure 6 Cisco UCS B200 M5 Front View

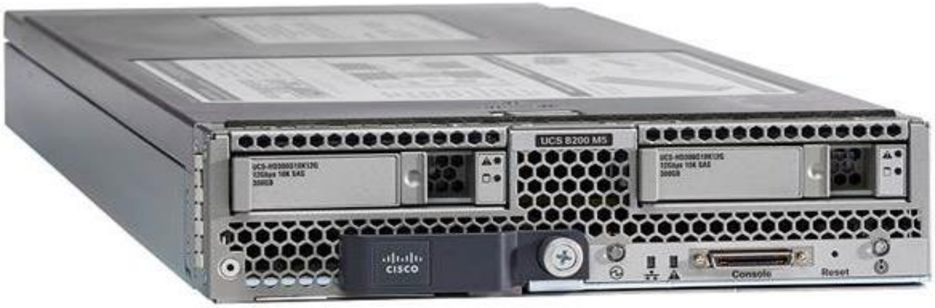

Figure 7 Cisco UCS B200 M5 Back View

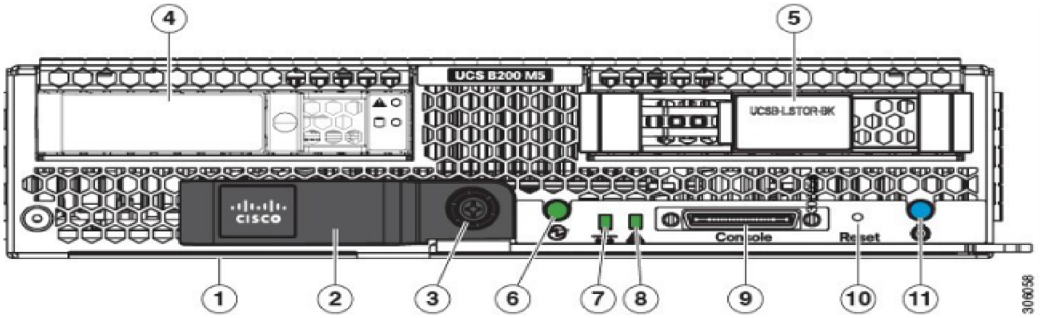

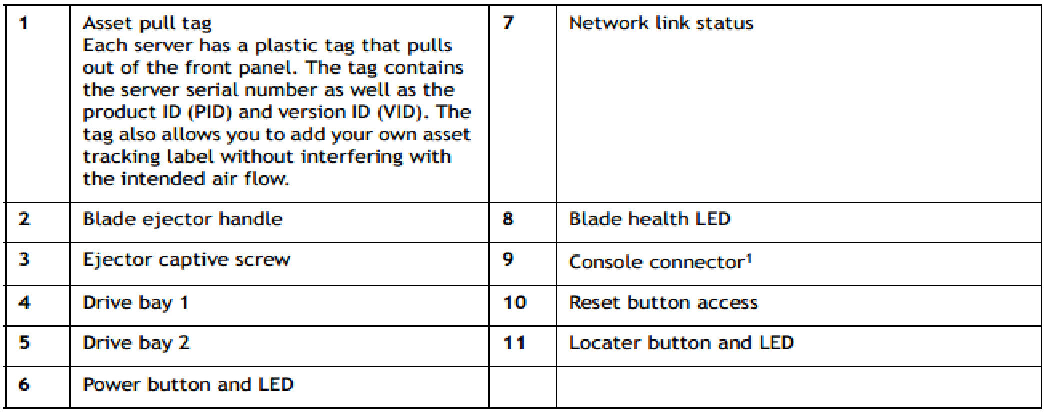

1	Asset pull tag Each server has a plastic tag that pulls out of the front panel. The tag contains the server serial number as well as the product ID (PID) and version ID (VID). The tag also allows you to add your own asset tracking label without interfering with the intended air flow.	7	Network link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

Notes:

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

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

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

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

Product Overview

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

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

Main Features

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

The Cisco UCS B200 M5 provides these main features:

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

Table 1 Product Specifications

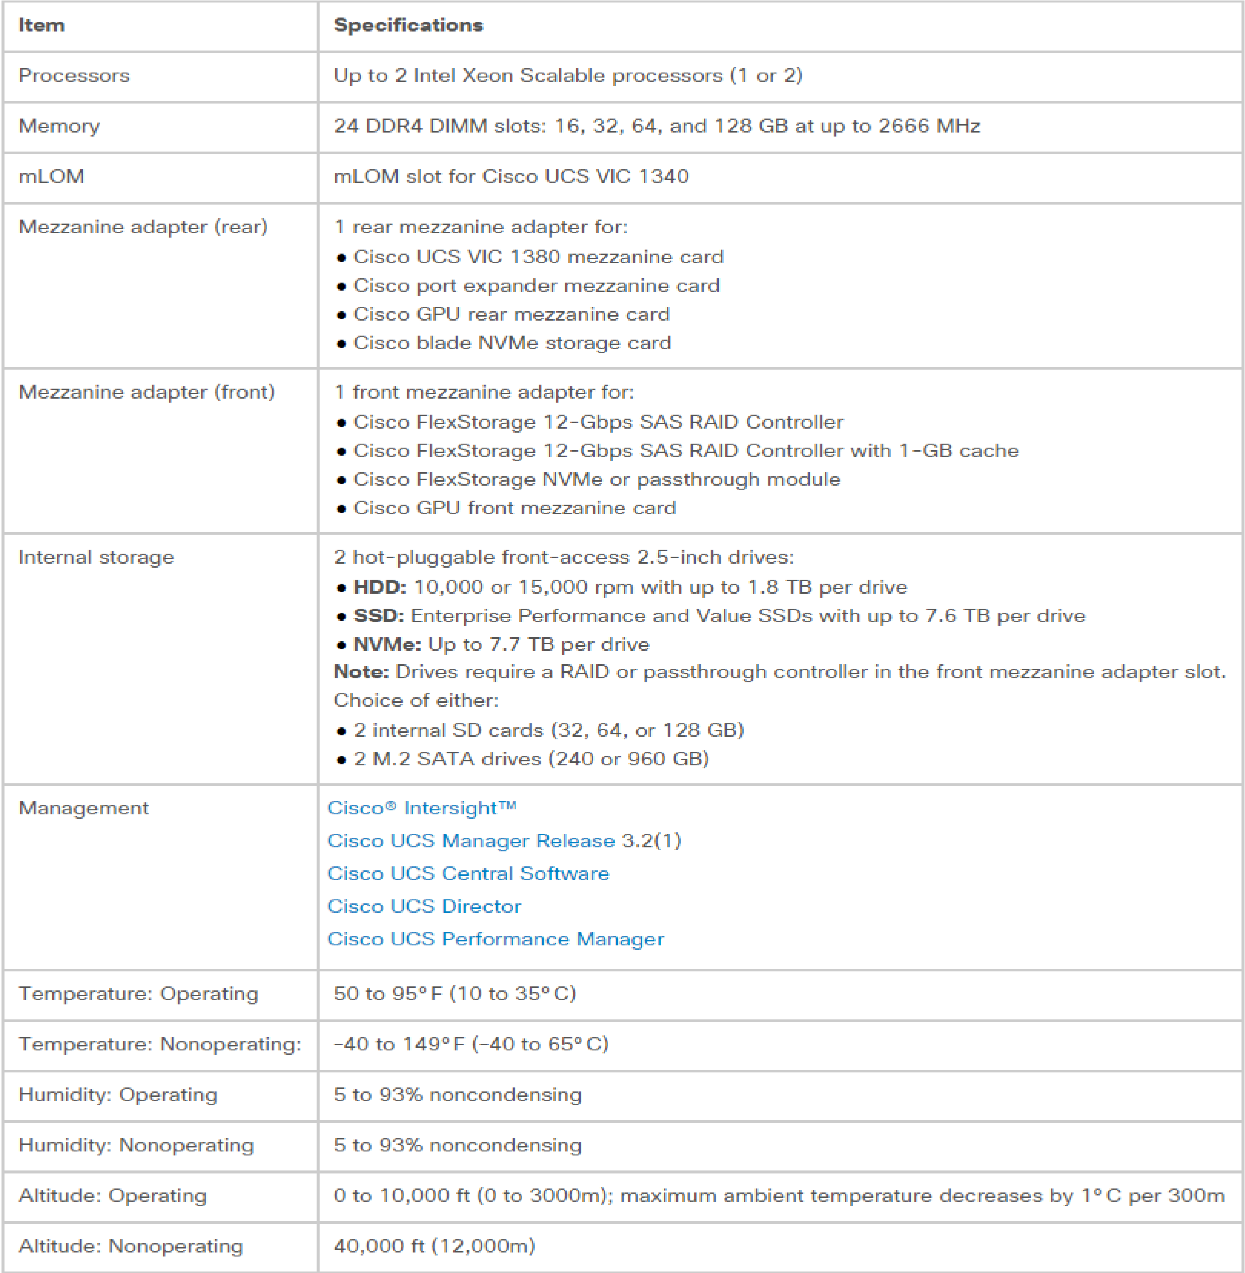

Item	Specifications
Processors	Up to 2 Intel Xeon Scalable processors (1 or 2)
Memory	24 DDR4 DIMM slots: 16, 32, 64, and 128 GB at up to 2666 MHz
mLOM	mLOM slot for Cisco UCS VIC 1340
Mezzanine adapter (rear)	1 rear mezzanine adapter for: • Cisco UCS VIC 1380 mezzanine card • Cisco port expander mezzanine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

Table 2 System Requirements

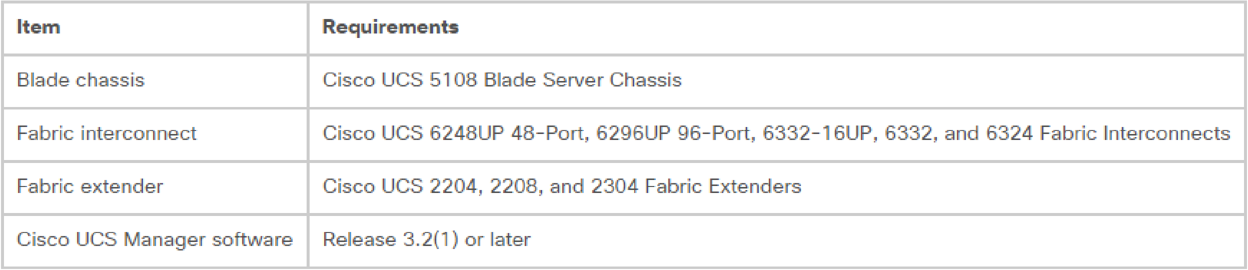

Item	Requirements
Blade chassis	Cisco UCS 5108 Blade Server Chassis
Fabric interconnect	Cisco UCS 6248UP 48-Port, 6296UP 96-Port, 6332-16UP, 6332, and 6324 Fabric Interconnects
Fabric extender	Cisco UCS 2204, 2208, and 2304 Fabric Extenders
Cisco UCS Manager software	Release 3.2(1) or later

Table 3 Ordering Information

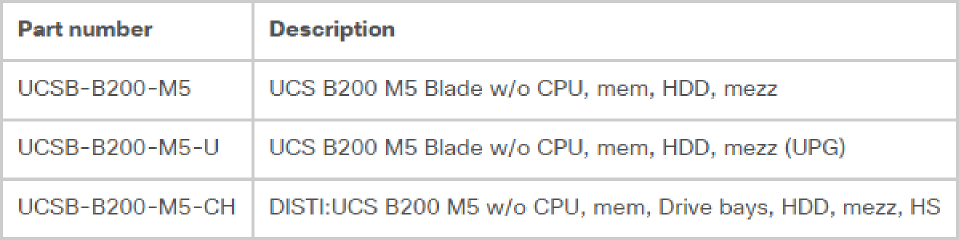

Part number	Description
UCSB-B200-M5	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz
UCSB-B200-M5-U	UCS B200 M5 Blade w/o CPU, mem, HDD, mezz (UPG)
UCSB-B200-M5-CH	DISTI:UCS B200 M5 w/o CPU, mem, Drive bays, HDD, mezz, HS

Table 4 Capabilities and Features

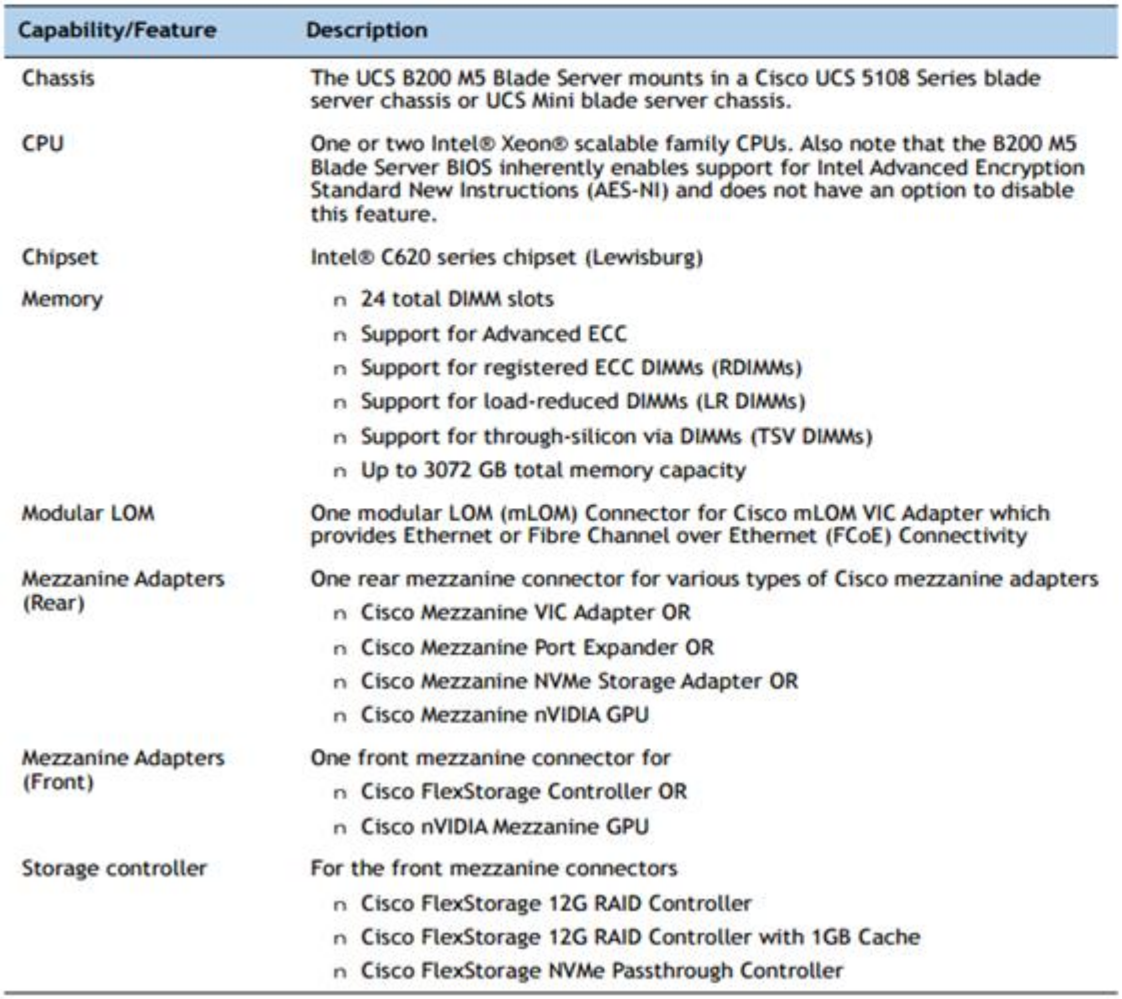

Capability/Feature	Description
Chassis	The UCS B200 M5 Blade Server mounts in a Cisco UCS 5108 Series blade server chassis or UCS Mini blade server chassis.
CPU	One or two Intel® Xeon® scalable family CPUs. Also note that the B200 M5 Blade Server BIOS inherently enables support for Intel Advanced Encryption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

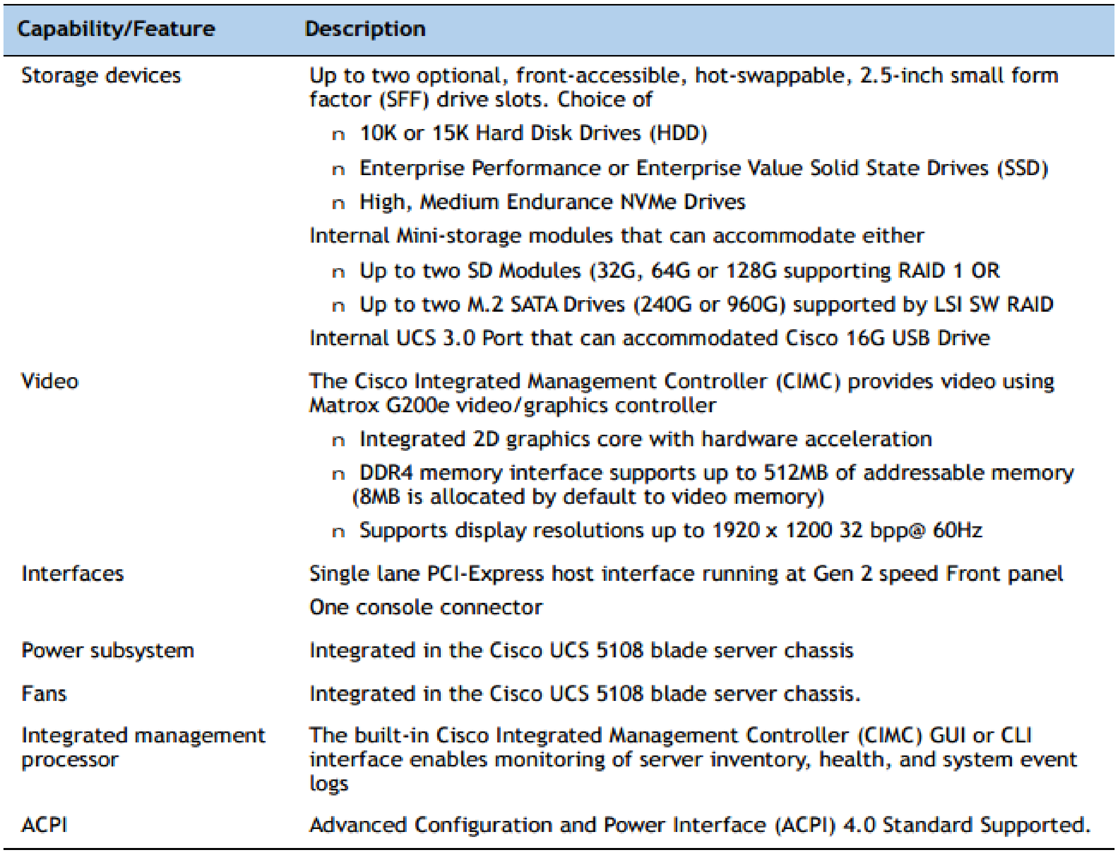

Capability/Feature	Description
Storage devices	Up to two optional, front-accessible, hot-swappable, 2.5-inch small form factor (SFF) drive slots. Choice of n 10K or 15K Hard Disk Drives (HDD) n Enterprise Performance or Enterprise Value Solid State Drives (SSD) n High, Medium Endurance NVMe Drives Internal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

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

Cisco UCS VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 8) is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

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

Figure 8 Cisco UCS VIC 1340

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

Figure 9 Cisco UCS VIC 1340 Deployed in the Cisco UCS B200 M5

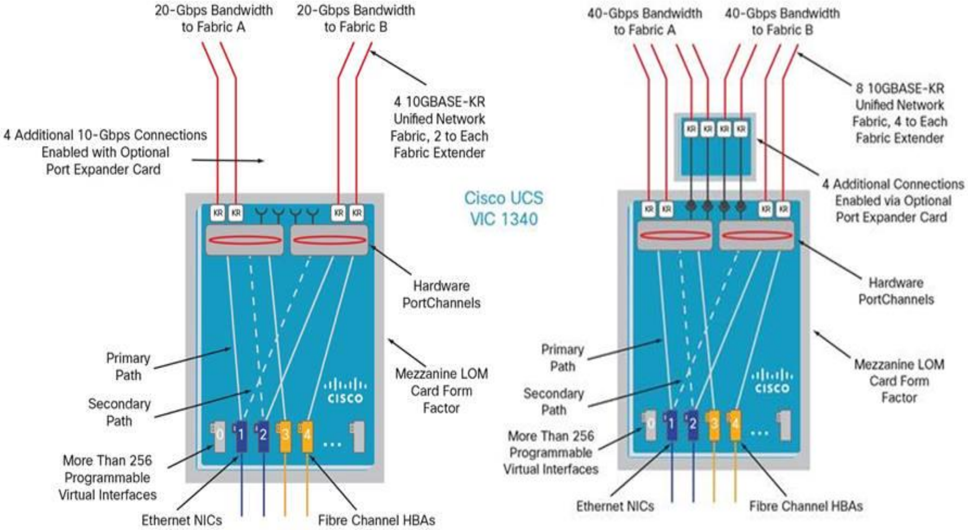
20-Gbps Bandwidth to Fabric A
20-Gbps Bandwidth to Fabric B
4 10GBASE-KR Unified Network Fabric, 2 to Each Fabric Extender
4 Additional 10-Gbps Connections Enabled with Optional Port Expander Card
Cisco UCS VIC 1340
KR KR KR KR
Hardware PortChannels
Primary Path
Secondary Path
Mezzanine LOM Card Form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

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 10 Nexus 9372PX Switch

The Nexus 9372PX switch will go end-of-sale in February 2018. For an all 40GB infrastructure, we recommend substituting the Nexus 9332PQ switch featuring 32 x 40GB ports.

Cisco MDS 9148S Fiber Channel Switch

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

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

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

Figure 11 Cisco 9148S MDS Fibre Channel Switch

Features and Capabilities

Benefits

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

Features

- PowerOn Auto Provisioning and intelligent diagnostics
- In-Service Software Upgrade and dual redundant hot-swappable power supplies for high availability
- Role-based authentication, authorization, and accounting services to support regulatory requirements

- High-performance interswitch links with multipath load balancing
- Smart zoning and virtual output queuing
- Hardware-based slow port detection and recovery

Specifications at-a-Glance

Performance and Port Configuration

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

Advanced Functions

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

Hypervisor and Desktop Broker

This Cisco Validated Design includes VMware vSphere 6.5 and VMware Horizon 7.3.

VMware vSphere 6.5

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

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

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

- Migration

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

vSphere Client

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

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

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

VMware ESXi 6.5 Hypervisor

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

- Scalability Improvements

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

- Security Enhancements
  - Account management: ESXi 6.5 enables management of local accounts on the ESXi server using new ESXi CLI commands. The capability to add, list, remove, and modify accounts across all hosts

in a cluster can be centrally managed using a vCenter Server system. Previously, the account and permission management functions for ESXi hosts were available only for direct host connections. The setup, removal, and listing of local permissions on ESXi servers can also be centrally managed.

- Account lockout: ESXi Host Advanced System Settings have two new options for the management of failed local account login attempts and account lockout duration. These parameters affect Secure Shell (SSH) and vSphere Web Services connections, but not ESXi direct console user interface (DCUI) or console shell access.
- Password complexity rules: In previous versions of ESXi, password complexity changes had to be made by manually editing the /etc/pam.d/passwd file on each ESXi host. In vSphere 6.0, an entry in Host Advanced System Settings enables changes to be centrally managed for all hosts in a cluster.
- Improved auditability of ESXi administrator actions: Prior to vSphere 6.0, actions at the vCenter Server level by a named user appeared in ESXi logs with the vpxuser username: for example, [user=vpxuser]. In vSphere 6.5, all actions at the vCenter Server level for an ESXi server appear in the ESXi logs with the vCenter Server username: for example, [user=vpxuser: DOMAIN\User]. This approach provides a better audit trail for actions run on a vCenter Server instance that conducted corresponding tasks on the ESXi hosts.
- Flexible lockdown modes: Prior to vSphere 6.5, only one lockdown mode was available. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.5, two lockdown modes are available:
  - In normal lockdown mode, DCUI access is not stopped, and users on the DCUI access list can access the DCUI.
  - In strict lockdown mode, the DCUI is stopped.
  - Exception users: vSphere 6.0 offers a new function called exception users. Exception users are local accounts or Microsoft Active Directory accounts with permissions defined locally on the host to which these users have host access. These exception users are not recommended for general user accounts, but they are recommended for use by third-party applications–for service accounts, for example–that need host access when either normal or strict lockdown mode is enabled. Permissions on these accounts should be set to the bare minimum required for the application to perform its task and with an account that needs only read-only permissions on the ESXi host.
- Smart card authentication to DCUI: This function is for U.S. federal customers only. It enables DCUI login access using a Common Access Card (CAC) and Personal Identity Verification (PIV). The ESXi host must be part of an Active Directory domain.

VMware Horizon Version 7

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

VMware Horizon

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

- VMware Horizon Virtual machines and RDSH known as server-based hosted sessions: These are applications hosted from Microsoft Windows servers to any type of device, including Windows PCs, Macs, smartphones, and tablets. Some VMware editions include technologies that further optimize the experience of using Windows applications on a mobile device by automatically translating native mobile-device display, navigation, and controls to Windows applications; enhancing performance over mobile networks; and enabling developers to optimize any custom Windows application for any mobile environment.
- VMware Horizon RDSH session users also known as server-hosted desktops: These are inexpensive, locked-down Windows virtual desktops hosted from Windows server operating systems. They are well suited for users, such as call center employees, who perform a standard set of tasks.

Advantages of Using VMware Horizon

VMware Horizon 7 provides the following new features and enhancements:

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

    - Reduced device power consumption for longer battery life
    - NVIDIA GRID acceleration for more graphical workloads per server, better performance, and a superior remote user experience
- True SSO
  - For VMware Identity Manager integration, True SSO streamlines the end-to-end login experience. After users log in to VMware Identity Manager using a smart card or an RSA SecurID or RADIUS token, users are not required to also enter Active Directory credentials in order to use a remote desktop or application.
  - Uses a short-lived Horizon virtual certificate to enable a password-free Windows login.
  - Supports using either a native Horizon Client or HTML Access.
  - System health status for True SSO appears in the Horizon Administrator dashboard.
  - Can be used in a single domain, in a single forest with multiple domains, and in a multiple-forest, multiple-domain setup.
- Smart Policies
  - Control of the clipboard cut-and-paste, client drive redirection, USB redirection, and virtual printing desktop features through defined policies.
  - PCoIP session control through PCoIP profiles.
  - Conditional policies based on user location, desktop tagging, pool name, and Horizon Client registry values.
- Configure the clipboard memory size for VMware Blast and PCoIP sessions

  Horizon administrators can configure the server clipboard memory size by setting GPOs for VMware Blast and PCoIP sessions. Horizon Client 4.1 users on Windows, Linux, and Mac OS X systems can configure the client clipboard memory size. The effective memory size is the lesser of the server and client clipboard memory size values.
- VMware Blast network recovery enhancements

  Network recovery is now supported for VMware Blast sessions initiated from iOS, Android, Mac OS X, Linux, and Chrome OS clients. Previously, network recovery was supported only for Windows client sessions. If you lose your network connection unexpectedly during a VMware Blast session, Horizon Client attempts to reconnect to the network and you can continue to use your remote desktop or application. The network recovery feature also supports IP roaming, which means you can resume your VMware Blast session after switching to a WiFi network.
- Configure Horizon Administrator to not remember the login name

  Horizon administrators can configure not to display the Remember user name check box and therefore not remember the administrator's login name.
- Allow Mac OS X users to save credentials

Horizon administrators can configure Connection Server to allow Horizon Client Mac OS X systems to remember a user's user name, password, and domain information. If users choose to have their credentials saved, the credentials are added to the login fields in Horizon Client on subsequent connections.

- Microsoft Windows 10
  - Windows 10 is supported as a desktop guest operating system
  - Horizon Client runs on Windows 10
  - Smart card is supported on Windows 10
  - The Horizon User Profile Migration tool migrates Windows 7, 8/8.1, Server 2008 R2, or Server 2012 R2 user profiles to Windows 10 user profiles.
- RDS Desktops and Hosted Apps
  - View Composer. View Composer and linked clones provide automated and efficient management of RDS server farms.
  - Graphics Support. Existing 3D vDGA and GRID vGPU graphics solutions on VDI desktops have been extended to RDS hosts, enabling graphics-intensive applications to run on RDS desktops and Hosted Apps.
  - Enhanced Load Balancing. A new capability provides load balancing of server farm applications based on memory and CPU resources.
  - One-Way AD Trusts. One-way AD trust domains are now supported. This feature enables environments with limited trust relationships between domains without requiring Horizon Connection Server to be in an external domain.
- Cloud Pod Architecture (CPA) Enhancements
  - Hosted App Support. Support for application remoting allows applications to be launched using global entitlements across a pod federation.
  - HTML Access (Blast) Support. Users can use HTML Access to connect to remote desktops and applications in a Cloud Pod Architecture deployment.
- Access Point Integration

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

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

  - AMD vDGA enables vDGA pass-through graphics for AMD graphics hardware.
  - 4K resolution monitors (3840x2160) are supported.
- Horizon Administrator Enhancements
  - Horizon Administrator shows additional licensing information, including license key, named user and concurrent connection user count.
  - Pool creation is streamlined by letting Horizon administrators to clone existing pools.
- Additional Features
  - Support for IPv6 with VMware Blast Extreme on security servers.
  - Horizon Administrator security protection layer. See VMware Knowledge Base (KB) article 2144303 for more information: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2144303
  - Protection against inadvertent pool deletion.
  - RDS per-device licensing improvements.
  - Support for Intel vDGA.
  - Support for AMD Multiuser GPU Using vDGA.
  - More resilient upgrades.
  - Display scaling for Windows Horizon Clients.
  - DPI scaling is supported if it is set at the system level and the scaling level is greater than 100.

What are VMware RDS Hosted Sessions?

The following describes the VMware RDS Hosted Sessions:

- An RDS host is a server computer that hosts applications and desktop sessions for remote access. An RDS host can be a virtual machine or a physical server.
- An RDS host has the Microsoft Remote Desktop Services role, the Microsoft Remote Desktop Session Host service, and Horizon Agent installed. Remote Desktop Services was previously known as Terminal Services. The Remote Desktop Session Host service allows a server to host applications and remote desktop sessions. With Horizon Agent installed on an RDS host, users can connect to applications and desktop sessions by using the display protocol PCoIP or Blast Extreme. Both protocols provide an optimized user experience for the delivery of remote content, including images, audio and video.
- The performance of an RDS host depends on many factors. For information on how to tune the performance of different versions of Windows Server, see

  http://msdn.microsoft.com/library/windows/hardware/gg463392.aspx.

- Horizon 7 supports at most one desktop session and one application session per user on an RDS host.
- When users submit print jobs concurrently from RDS desktops or applications that are hosted on the same RDS host, the ThinPrint server on the RDS host processes the print requests serially rather than in parallel. This can cause a delay for some users. Note that the print server does not wait for a print job to complete before processing the next one. Print jobs that are sent to different printers will print in parallel.
- If a user launches an application and also an RDS desktop, and both are hosted on the same RDS host, they share the same user profile. If the user launches an application from the desktop, conflicts may result if both applications try to access or modify the same parts of the user profile, and one of the applications may fail to run properly.
- The process of setting up applications or RDS desktops for remote access involves the following tasks:
  - Installing Applications

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

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

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

Farms, RDS Hosts, and Desktop and Application Pools

With VMware Horizon, you can create desktop and application pools to give users remote access to virtual machine-based desktops, session-based desktops, physical computers, and applications. Horizon takes advantage of Microsoft Remote Desktop Services (RDS) and VMware PC-over-IP (PCoIP) technologies to provide high-quality remote access to users.

- RDS Hosts

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

- Desktop Pools

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

- Application Pools

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

- Farms

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

Some of the latest VMware Horizon features and enhancements are:

- Flash Redirection
  - You can compile a black list to ensure that the URLs specified in the list will not be able to redirect Flash content. You must enable the GPO setting FlashMMRUrlListEnableType to use either a white list or black list.
- Horizon Agent Policy Settings
  - The VMwareAgentCIT policy setting enables remote connections to Internet Explorer to use the Client's IP address instead of the IP address of the remote desktop machine.
  - The FlashMMRUrlListEnableType and FlashMMRUrlList policy settings specify and control the white list or black list that enables or disables the list of URLs from using Flash Redirection.
- Horizon PowerCLI
  - View PowerCLI is deprecated. Horizon PowerCLI replaces View PowerCLI and includes cmdlets that you can use with VMware PowerCLI.
  - For more information about Horizon PowerCLI cmdlets, read the *VMware PowerCLI Cmdlets Reference.*
  - For information on the API specifications to create advanced functions and scripts to use with Horizon PowerCLI, see the API Reference at the VMware Developer Center
  - For more information on sample scripts that you can use to create your own Horizon PowerCLI scripts, see the Horizon PowerCLI community on GitHub.
- Horizon 7 for Linux desktops enhancements

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

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

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

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

  - Additional platforms support
    RHEL 7.4 x64 and CentOS 7.4 x64 are now supported.

- Remote Desktop Operating System

  - Windows 10 version 1607 Long-Term Servicing Branch (LTSB)

  - Windows Server 2016

Horizon Agent

- HTML5 Multimedia Redirection

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

- SHA-256 support

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

- Improved USB redirection with User Environment Manager

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

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

- Composer
  - For enhanced security, you can enable the digest access authentication method for Composer.
- Persona Management
  - Persona Management supports guest operating systems that use the "v6" version of the user profile.
  - You can use the migration tool to migrate the "v2" and "v5" user profiles versions to the "v6" user profile version. The tool is installed with the Persona binary file.

Horizon Connection Server Enhanced Features

- Horizon Help Desk Tool
  - View application and process names and resource use within a virtual or published desktop to identify which applications and process are using up machine resources.
  - View event log information about the user's activities.
  - View updated metrics such as Horizon Client version and the Blast protocol.
  - View additional session metrics such as the VM information, CPU, or memory usage.
  - You can assign predefined administrator roles to Horizon Help Desk Tool administrators to delegate the troubleshooting tasks between administrator users. You can also create custom roles and add privileges based on the predefined administrator roles.
  - You can verify the product license key for Horizon Help Desk Tool and apply a valid license.
- Monitoring
  - If the event database shuts down, Horizon Administrator maintains an audit trail of the events that occur before and after the event database shutdown.
- Instant Clones
  - You can create dedicated instant-clone desktop pools.
  - Windows Server operating systems are supported for instant clones in this release. For an updated list of supported Windows Server operating systems, see the VMware Knowledge Base (KB) article 2150295.
  - You can copy, paste, or enter the path for the AD tree in the AD container field when you create an instant-clone desktop pool.
  - If there are no internal VMs in all four internal folders created in vSphere Web Client, these folders are unprotected and you can delete these folders.
  - You can use the enhanced instant-clone maintenance utility IcUnprotect.cmd to unprotect or delete template, replica, or parent VMs or folders from vSphere hosts.

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

- Cloud Pod Architecture
  - The total session limit is increased to 140,000.
  - The site limit is increased to 7.
  - You can configure Windows Start menu shortcuts for global entitlements. When an entitled user connects to a Connection Server instance in the pod federation, Horizon Client for Windows places these shortcuts in the Start menu on the user's Windows client device.
- Published Desktops and Application Pools
  - You can restrict access to entitled desktop pools, application pools, global entitlements, and global application entitlements from certain client computers.
  - You can configure Windows start menu shortcuts for entitled desktop and application pools. When an entitled user connects to a Connection Server instance, Horizon Client for Windows places these shortcuts in the Start menu on the user's Windows client device.
- Virtual Desktops and Desktop Pools
  - Blast Extreme provides network continuity during momentary network loss on Windows clients.
  - Performance counters displayed using PerfMon on Windows agents for Blast session, imaging, audio, CDR, USB, and virtual printing provide an accurate representation of the current state of the system that also updates at a constant rate.
- Customer Experience Improvement Program
  - Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust Assurance Center.
- Security
  - With the USB over Session Enhancement SDK feature, you do not need to open TCP port 32111 for USB traffic in a DMZ-based security server deployment. This feature is supported for both virtual desktops and published desktops on RDS hosts.
- Database Support
  - The Always On Availability Groups feature for Microsoft SQL Server 2014 is supported in this release of Horizon 7. For more information, refer to the Release Notes.

Supported Windows 10 Operating Systems

Horizon 7 version 7.3.1 supports the following Windows 10 operating systems:

- Windows 10 version 1507 (RTM) Long-Term Servicing Branch (LTSB)
- Windows 10 version 1607 Long-Term Servicing Branch (LTSB)

- Windows 10 version 1607 Enterprise Current Branch (CBB)
- Windows 10 version 1703 Semi Annual Channel (broad deployment) Current Branch (CBB)

For the complete list of supported Windows 10 on Horizon including all VDI (Full Clones, Linked and Instant clones) click the following links: https://kb.vmware.com/s/article/2149393 and https://kb.vmware.com/s/article/2150295?r=2&Quarterback.validateRoute=1&KM_Utility.getArticleData=1&KM_Utility.getGUser=1&KM_Utility.getArticleLanguage=2&KM_Utility.getArticle=1

Windows 10 LTSB version 1607 being used in this study.

Figure 12 Logical Architecture of VMware Horizon

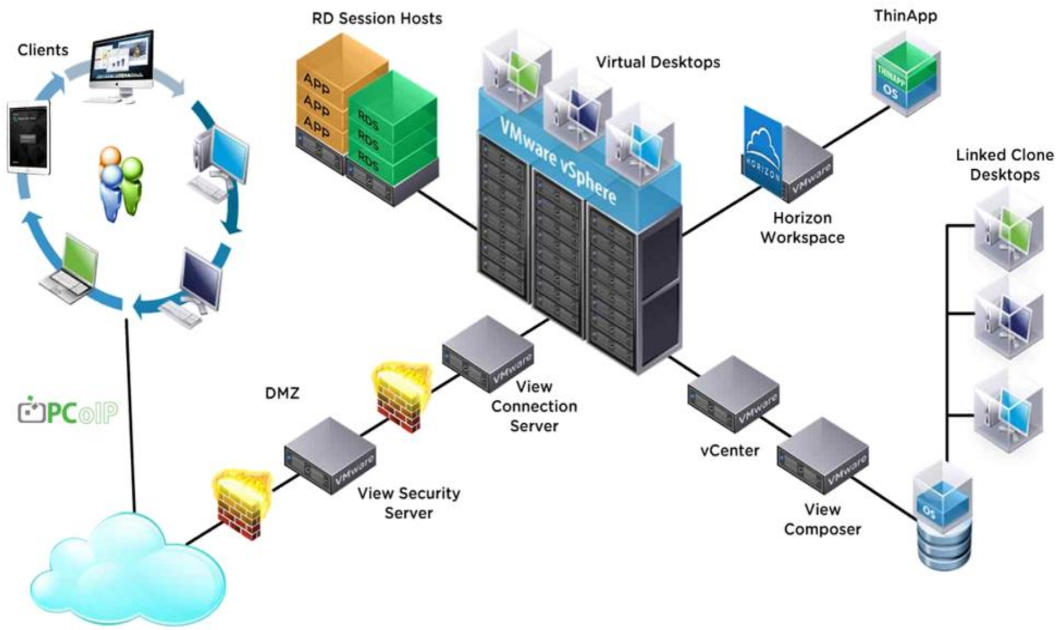

VMware Horizon Composer

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

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

Figure 13 VMware Horizon Composer

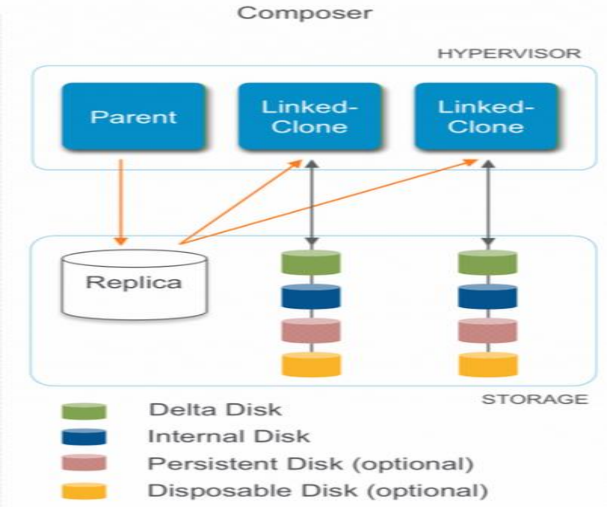

Desktop Virtualization Design Fundamentals

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

VMware Horizon Design Fundamentals

VMware Horizon 7 integrates Remote Desktop Server Hosted sessions users and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, mixed users and manageable solution for delivering Windows applications and desktops as a service.

Users can select applications from an easy-to-use “store” that is accessible from tablets, smartphones, PCs, Macs, and thin clients. VMware Horizon delivers a native touch-optimized experience via PCoIP or Blast Extreme high-definition performance, even over mobile networks.

Horizon VDI Pool and RDSH Servers Pool

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Desktop Pool. In this CVD, VM provisioning relies on VMware View Composer aligning with VMware Horizon Connection Server with vCenter Server components. In this CVD, machines in the Pools are configured to run either a Windows Server 2016 OS (for RDS Hosted shared sessions) and a Windows 10 Desktop OS (for pooled VDI desktops).

Figure 14 VMware Horizon Design Overview

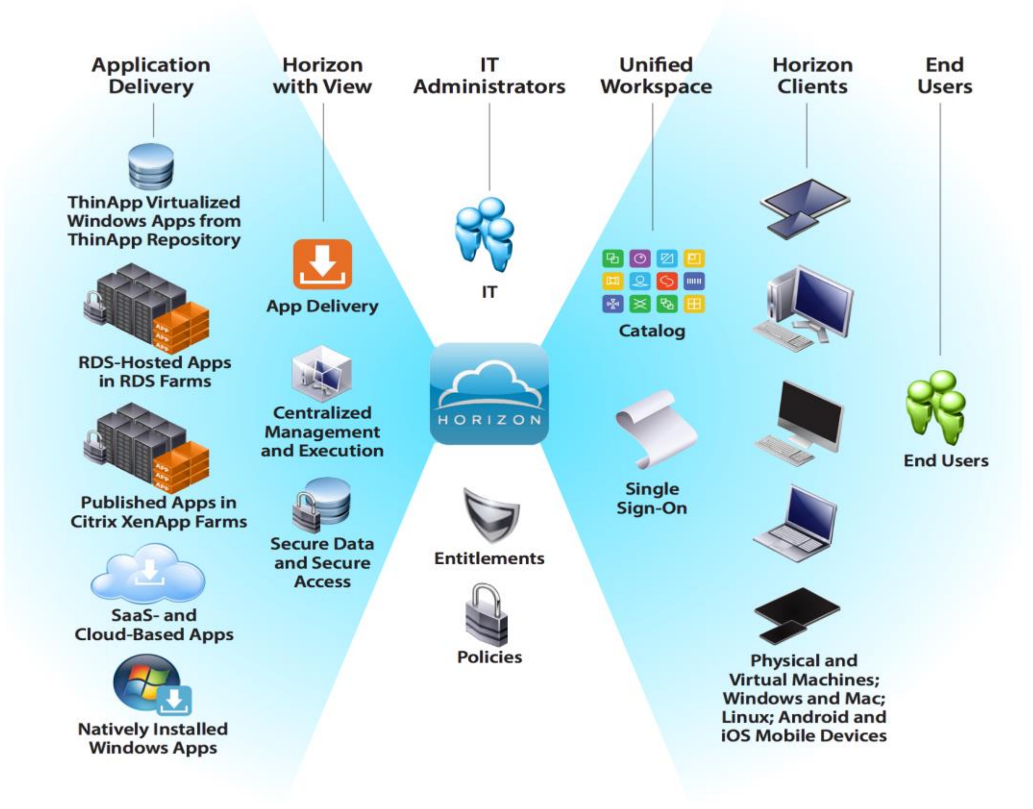
Application Delivery
Horizon with View
IT Administrators
Unified Workspace
Horizon Clients
End Users
ThinApp Virtualized Windows Apps from ThinApp Repository
RDS-Hosted Apps in RDS Farms
Published Apps in Citrix XenApp Farms
SaaS- and Cloud-Based Apps
Natively Installed Windows Apps
App Delivery
Centralized Management and Execution
Secure Data and Secure Access
IT
HORIZON
Entitlements
Policies
Catalog
Single Sign-On
Physical and Virtual Machines; Windows and Mac; Linux; Android and iOS Mobile Devices
End Users

Figure 15 Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PCoIP/Blast/RDP)

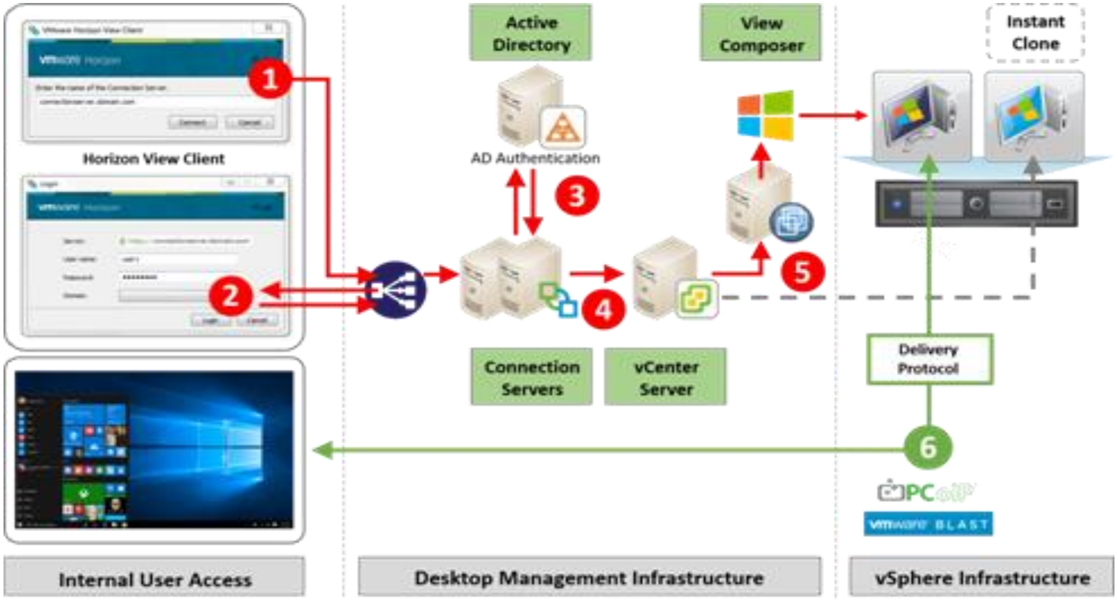

NetApp A-Series All Flash FAS

With the new NetApp A-Series All Flash FAS (AFF) controller lineup, NetApp provides industry-leading performance while continuing to provide a full suite of enterprise-grade data management and data protection features. The A-Series lineup offers double the IOPS, while decreasing the latency. The AFF A-

Series lineup includes the A200, A300, A700, and A700s. These controllers and their specifications listed in Table 5 . For more information about the A-Series AFF controllers, see:

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

Table 5 NetApp A-Series Controller Specifications

	AFF A200	AFF A300	AFF A700	AFF A700s
NAS Scale-out	2-8 nodes	2-24 nodes	2-24 nodes	2-24 nodes
SAN Scale-out	2-8 nodes	2-12 nodes	2-12 nodes	2-12 nodes
Per HA Pair Specifications (Active-Active Dual Controller)				
Maximum SSDs	144	384	480	216
Maximum Raw Capacity	2.2PB	5.9PB	7.3PB	3.3PB
Effective Capacity	8.8PB	23.8PB	29.7PB	13PB
Chassis Form Factor	2U chassis with two HA controllers and 24 SSD slots	3U chassis with two HA controllers	8u chassis with two HA controllers	4u chassis with two HA controllers and 24 SSD slots

This solution utilizes the NetApp AFF A300, seen in Figure 8 and Figure 9. This controller provides the high-performance benefits of 40GbE and all flash SSDs, offering better performance than previous models, and occupying only 3U of rack space versus 6U with the AFF8040. When combined with the 2U disk shelf of 3.8TB disks, this solution can provide ample horsepower and over 90TB of raw capacity, all while occupying only 5U of valuable rack space. This makes it an ideal controller for a shared workload converged infrastructure. The A700s would be an ideal fit for situations where more performance is needed

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

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

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

Figure 16 NetApp AFF A300 Front View

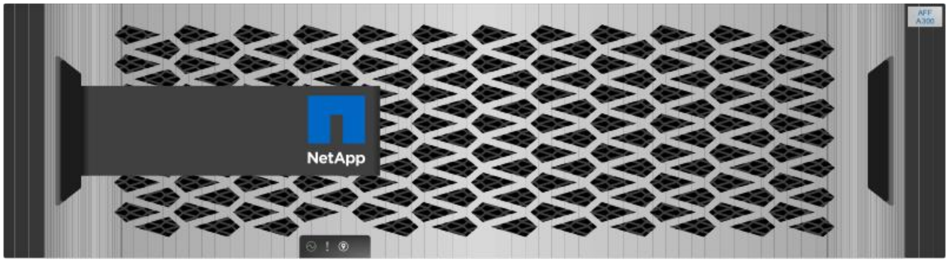

Figure 17 NetApp AFF A300 Rear View

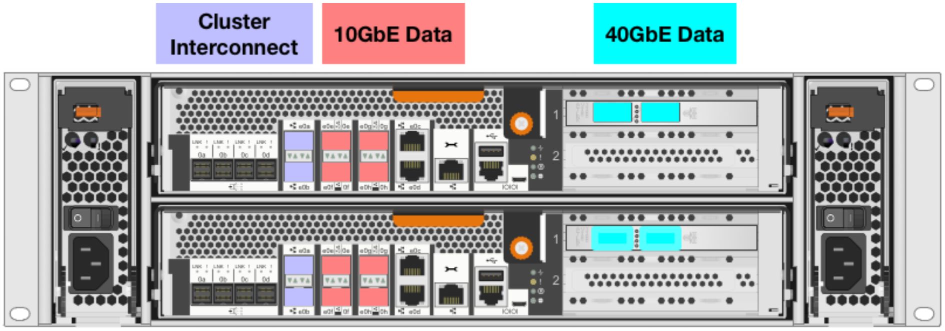

NetApp ONTAP 9.1

Storage Efficiency

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

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

Compaction, which is introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the NetApp WAFL® file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk-to-save space. This process is illustrated in Figure 18.

Figure 18 Storage Efficiency

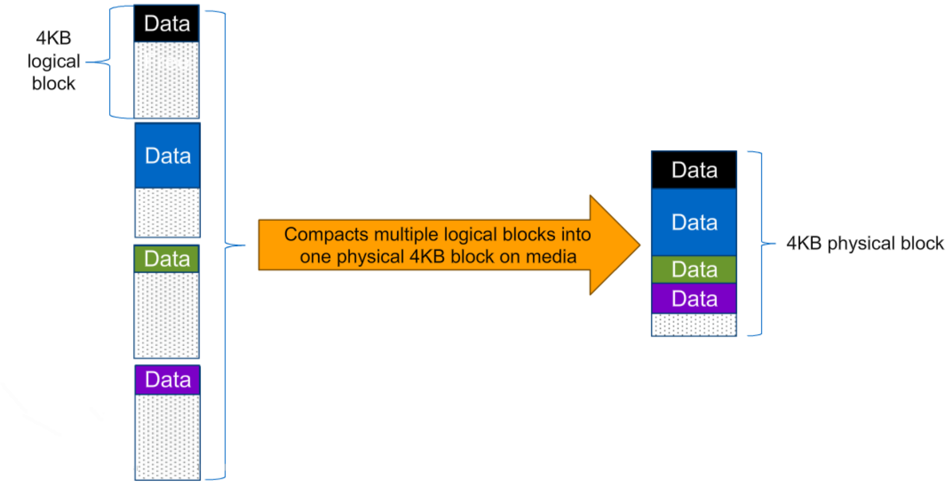
4KB logical block
Data
Data
Data
Data
Compacts multiple logical blocks into one physical 4KB block on media
Data
Data
Data
Data
4KB physical block

NetApp Storage Virtual Machine (SVM)

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

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

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

SAN Boot

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

In FC SAN boot, each Cisco UCS server boots by connecting the NetApp All Flash FAS storage to the Cisco MDS switch. The 16G FC storage ports, in this example 0g and 0h, are connected to Cisco MDS switch. The FC LIFs are created on the physical ports and each FC LIF is uniquely identified by its target WWPN. The storage system target WWPNs can be zoned with the server initiator WWPNs in the Cisco MDS switches. The FC boot LUN is exposed to the servers through the FC LIF using the MDS switch; this enables only the authorized server to have access to the boot LUN. Refer Figure 19 for the port and LIF layout

Figure 19 FC - SVM ports and LIF layout

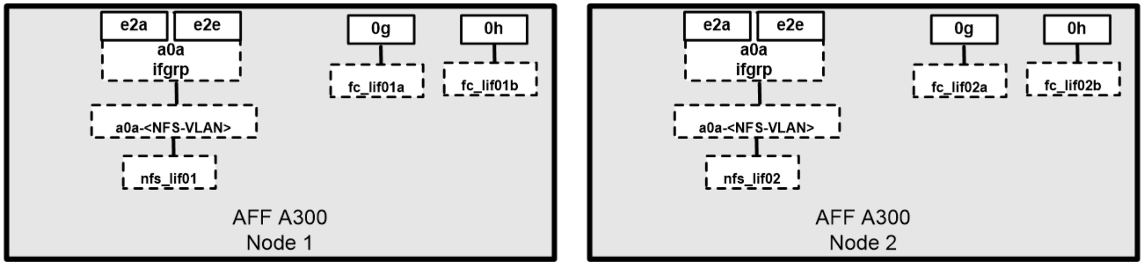

Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead if a network interface becomes unavailable, the host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp used to provide information about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

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 2016, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remoted Desktop Server Hosted Server sessions: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- Published Applications: Published applications run entirely on the VMware Horizon RDS hosted server virtual machines and the user interacts through a delivery protocol. With published applications, a

single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.

- Streamed Applications: Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly but the resources may only available while they are connected to the network.
- Local Virtual Desktop: A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a type 1 hypervisor and is synced with the data center when the device is connected to the network.
- For the purposes of the validation represented in this document, both VMware Horizon hosted virtual desktops and Remote Desktop Server Hosted sessions were validated. Each of the sections provides some fundamental design decisions for this environment.

Understanding Applications and Data

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

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

Project Planning and Solution Sizing Sample Questions

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

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

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

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

- What is the desktop OS planned? Windows 8 or Windows 10?

- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 8/10?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?
- Will VMware Horizon RDSH be used for Hosted Shared Server applications planned? Are they are any applications installed?
- What is the desktop OS planned for RDS Server Roles? Windows server 2012 or Server 2016?
- Will VMware Horizon Composer or Instant Clones or another method be used for virtual desktop deployment?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- What is the SQL server version for database? SQL server 2012 or 2016?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere has been identified the hypervisor for both RDS Hosted Sessions and VDI based desktops:

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

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

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

Designing a VMware Horizon Environment for a Mixed Workload

With VMware Horizon 7 the method you choose to provide applications or desktops to users depends on the types of applications and desktops you are hosting and available system resources, as well as the types of users and user experience you want to provide.

Table 6 Designing a VMware Horizon Environment

Server OS machines	You want: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience. Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations. Application types: Any application.
Desktop OS machines	You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition. Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications. Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines. Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.
Remote PC Access	You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop im-

	ages 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 Remote Desktop Server Hosted sessions (RDSH) using RDS based Server OS and VMware Horizon pooled Linked Clone Virtual Machine Desktops using VDI based desktop OS machines were configured and tested.

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

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 NetApp AFF A300 storage arrays).

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

This CVD document details the deployment of 5000 users for a mixed VMware Horizon desktop workload featuring the following software:

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

- VMware vSphere ESXi 6.5 Update 1 Hypervisor
- Microsoft SQL Server 2016
- VMware Horizon 7 Shared Remote Desktop Server Hosted Sessions (RDSH) on NetApp AFF A300 FC storage
- VMware Horizon 7 Non-Persistent Virtual Desktops (VDI) on NetApp AFF A300 on FC storage
- VMware Horizon 7 Persistent Virtual Desktops (VDI) on NetApp AFF A300 on FC storage
- VMware Horizon 7 Connection Server and Additional Replica Servers
- VMware Horizon 7 Composer Server
- Microsoft Windows Server 2016 for Infrastructure
- Microsoft Windows Server 2016 for RDS Server Roles Configuration
- Windows 10 64-bit virtual machine Operating Systems for Non- Persistent and Persistent virtual machine users

Figure 20 details the physical hardware and cabling deployed to enable this solution.

Figure 20 Virtual Desktop Workload Architecture for the 5000 seat on VMware Horizon 7 on FlexPod

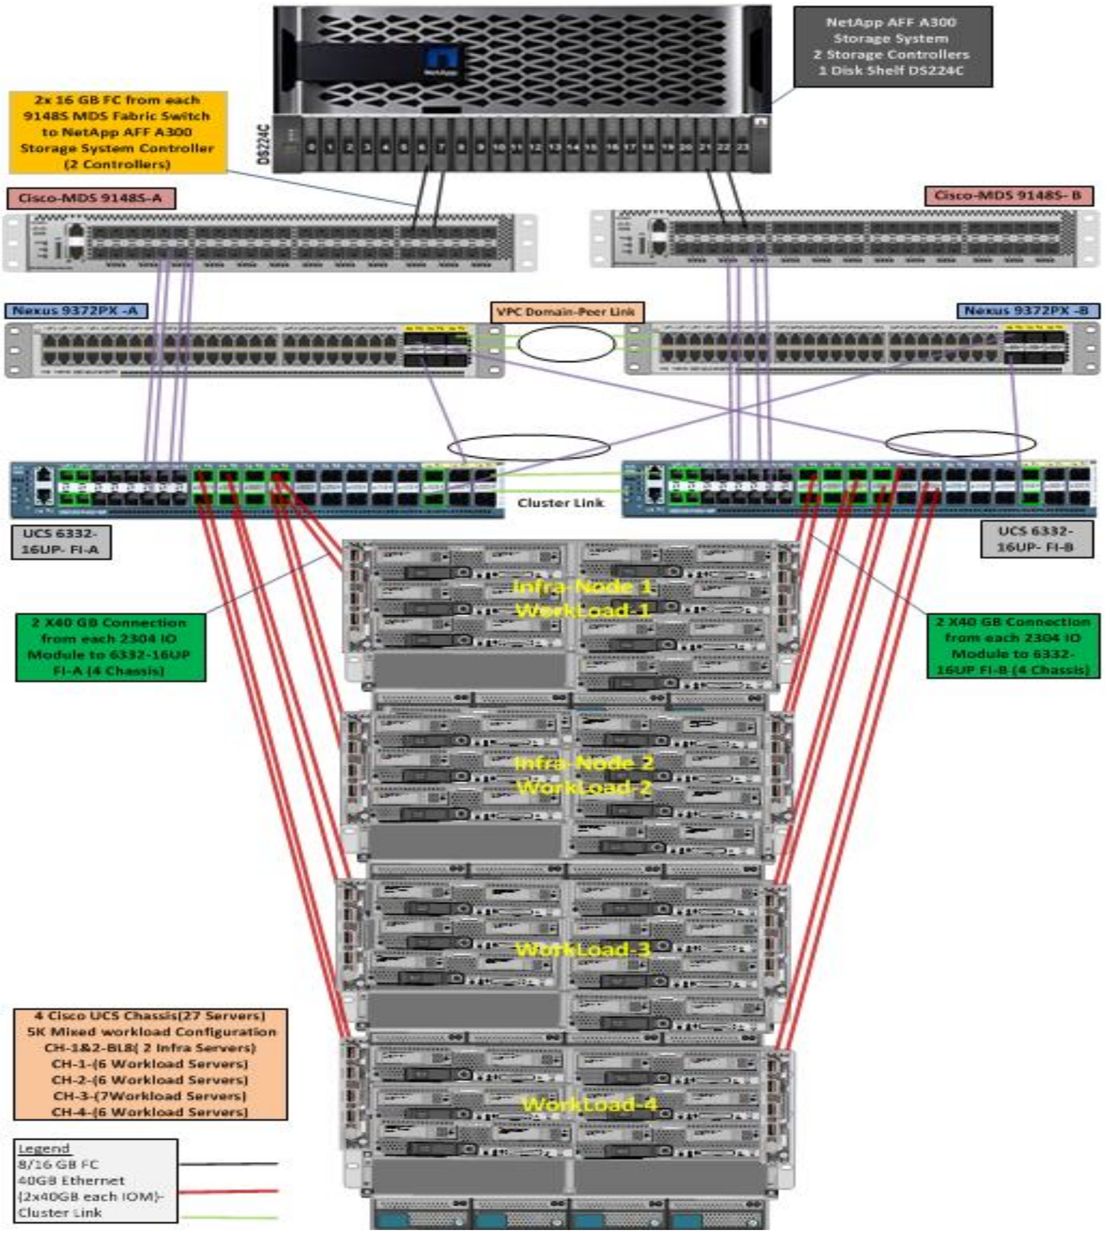

Hardware Deployed

The solution contains the following hardware as shown in Figure 20:

- Two Cisco Nexus 9372PX Layer 2 Access Switches
- Two Cisco MDS 9148S 16Gb Fibre Channel Switches
- Four Cisco UCS 5108 Blade Server Chassis with two Cisco UCS-IOM-2304 IO Modules
- Two Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.60-GHz 10-core processors, 128GB 2133MHz RAM, and one Cisco VIC1340 mezzanine card for the hosted infrastructure, providing N+1 server fault tolerance

- Seven Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the VMware Horizon Remote Desktop Server Hosted Sessions workload, providing N+1 server fault tolerance at the workload cluster level
- Nine Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the VMware Horizon Instant Clones VDI desktops workload, providing N+1 server fault tolerance at the workload cluster level
- Nine Cisco UCS B200 M5 Blade Servers with Intel Xeon Gold 6140 2.30-GHz 18-core processors, 768GB 2666MHz RAM, and one Cisco VIC1340 mezzanine card for the VMware Horizon Full Clones VDI desktops workload, providing N+1 server fault tolerance at the workload cluster level
- NetApp AFF A300 Storage System with dual redundant controllers, 1x DS224C disk shelf, and 24x 3.8 TB solid-state drives.

Table 7 lists the software and firmware version used in the study.

Table 7 Software and Firmware Versions used

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	3.2(1b) bundle release
Cisco	UCS Manager	3.2(1b) bundle release
Cisco	UCS B200 M4 Blades	3.2(1b) bundle release
Cisco	VIC 1340	4.1(1d)
VMware	VMware Horizon	7.3.1
VMware	VMware Composer Server	7.3.1
VMware	vCenter Server Appliance	6.5.0.Build 5973321
VMware	vSphere ESXi 6.5 Update 1	6.5.0. Build 5705665
NetApp	AFF A300	9.1

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 27 blades, which provides physical redundancy for the blade servers for each workload type.

Figure 21 outlines the logical architecture of the test environment, including the Login VSI session launcher self-contained end user experience benchmarking platform.

Figure 21 Logical Architecture Overview

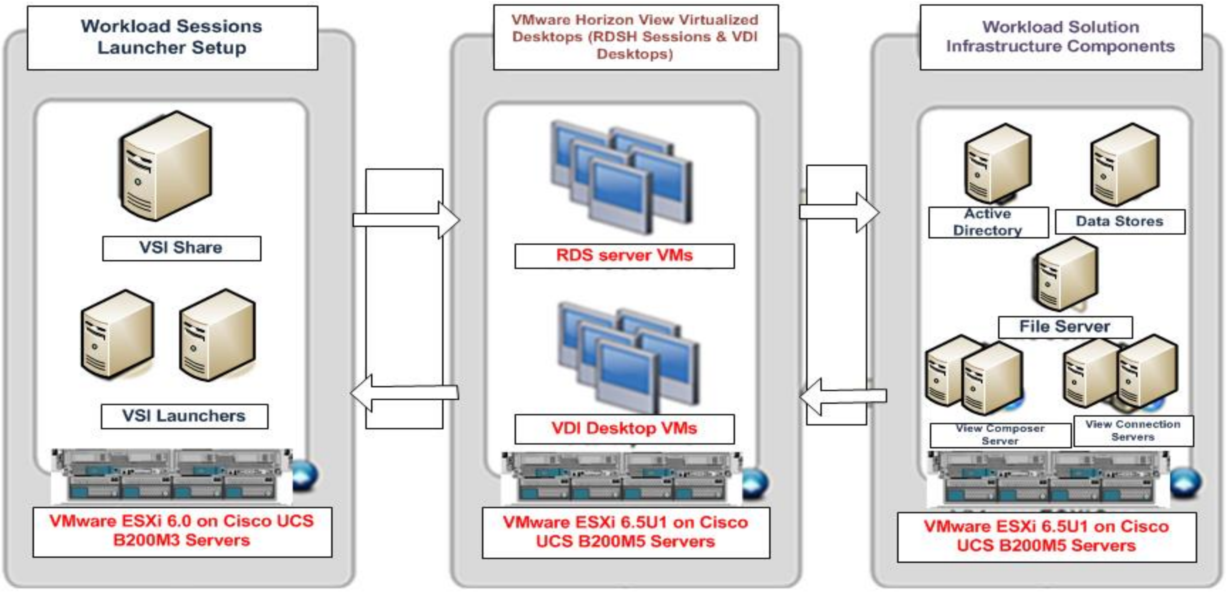
Workload Sessions Launcher Setup
VSI Share
VSI Launchers
VMware ESXi 6.0 on Cisco UCS B200M3 Servers
VMware Horizon View Virtualized Desktops (RDSH Sessions & VDI Desktops)
RDS server VMs
VDI Desktop VMs
VMware ESXi 6.5U1 on Cisco UCS B200M5 Servers
Workload Solution Infrastructure Components
Active Directory
Data Stores
File Server
View Composer Server
View Connection Servers
VMware ESXi 6.5U1 on Cisco UCS B200M5 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.

Figure 22 identifies the server roles in the 27 server deployment to support the 5000 seat workload. We also break out the infrastructure virtual machine fault tolerant design.

Figure 22 Server, Location, and Purpose

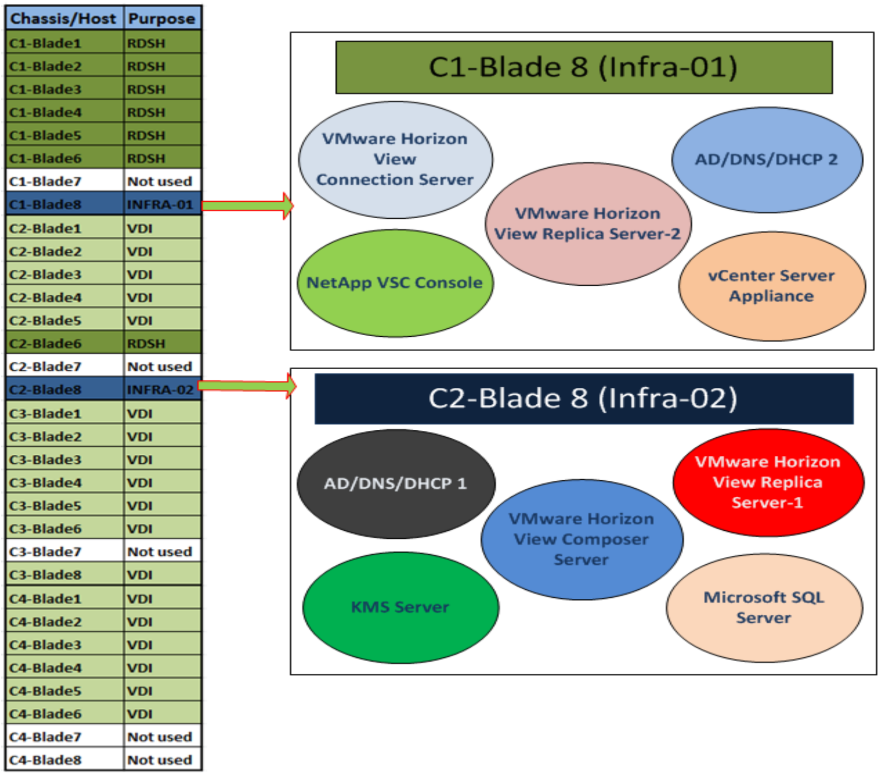

Chassis/Host	Purpose
C1-Blade1	RDSH
C1-Blade2	RDSH
C1-Blade3	RDSH
C1-Blade4	RDSH
C1-Blade5	RDSH
C1-Blade6	RDSH
C1-Blade7	Not used
C1-Blade8	INFRA-01
C2-Blade1	VDI
C2-Blade2	VDI
C2-Blade3	VDI
C2-Blade4	VDI
C2-Blade5	VDI
C2-Blade6	RDSH
C2-Blade7	Not used
C2-Blade8	INFRA-02
C3-Blade1	VDI
C3-Blade2	VDI
C3-Blade3	VDI
C3-Blade4	VDI
C3-Blade5	VDI
C3-Blade6	VDI
C3-Blade7	Not used
C3-Blade8	VDI
C4-Blade1	VDI
C4-Blade2	VDI
C4-Blade3	VDI
C4-Blade4	VDI
C4-Blade5	VDI
C4-Blade6	VDI
C4-Blade7	Not used
C4-Blade8	Not used

C1-Blade 8 (Infra-01)
VMware Horizon View Connection Server
AD/DNS/DHCP 2
VMware Horizon View Replica Server-2
NetApp VSC Console
vCenter Server Appliance
C2-Blade 8 (Infra-02)
AD/DNS/DHCP 1
VMware Horizon View Replica Server-1
VMware Horizon View Composer Server
KMS Server
Microsoft SQL Server

Table 8 lists the virtual machine deployments on the hardware platform.

Table 8 Virtual Machine Deployment Architecture

Server Name	Location	Purpose
C1-Blade 8 C2-Blade 8	Physical – Chassis 1, 2	ESXi 6.5 Hosts Infrastructure VMs Windows 2016, vCenter Server Appliance, VMware Horizon Connection Servers, Horizon Replica Servers, Horizon Composer Server, Active Directory Domain Controllers, SQL Server and Key Management Server.
C1-Blade1-6 C2-Blade 6	Physical – Chassis 1, 2	ESXi 6.5 Hosts 72x VMware Horizon Server 2016 RDSH Server VMs (1680 RDS Server Sessions)
C2-Blade1-5 C3-Blade1-7 C4-Blade1-6	Physical – Chassis 2,3, 4	ESXi 6.5 Hosts 3320x VMware Horizon VDI (2 Pools consist of Non-Persistent and Persistent virtual machines) VMs

Chassis-1/2/3, Slot 7 and Chassis-4 Slot 7 and 8 are not being used.

VLANs

The VLAN configuration recommended for the environment includes a total of seven VLANs as outlined in Table 9 .

Table 9 VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt.	60	VLAN for in-band management interfaces
Infra-Mgmt.	61	VLAN for Virtual Infrastructure
NFS	62	VLAN NFS Traffic
CIFS	63	VLAN for CIFS Share User Profile
vMotion	66	VLAN for VMware vMotion
VDI	102	VLAN for VDI Traffic
OB-Mgmt.	164	VLAN for out-of-band management interfaces

VSANs

Two virtual SANs for communications and fault tolerance were used in this design.

Table 10 VASNs Configured in this Study

VSAN Name	VSAN Purpose	ID Used in Validating this Document
VSAN 1	VSAN for primary SAN communication	400
VSAN 2	VSAN for secondary SAN communication	401

VMware Clusters

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

- VDI Cluster: NetApp AFF A300 storage with Cisco UCS
  - Infrastructure Cluster: Infra VMs (vCenter Appliance, Active Directory (2), DNS, DHCP, VMware Horizon Connection Servers, VMware Horizon Replica Servers, VMware Horizon Composer Server, Microsoft SQL Server.
  - RDSH: VMware Horizon RDSH (Remote Desktop Server Hosted) VMs (Windows Server 2016 RDS Roles) provisioned with VMware View Composer.

- VDI Non-Persistent: VMware Horizon VDI VMs (Windows 10 64-bit Non-Persistent Instant Clones virtual desktops provisioned.
- VDI Persistent: VMware Horizon VDI VMs (Windows 10 64-bit persistent virtual desktops provisioned with VMware Horizon Composer.

- VSI Launchers Cluster
  - Launcher Cluster: 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 23 VMware vSphere Clusters on vSphere Web GUI

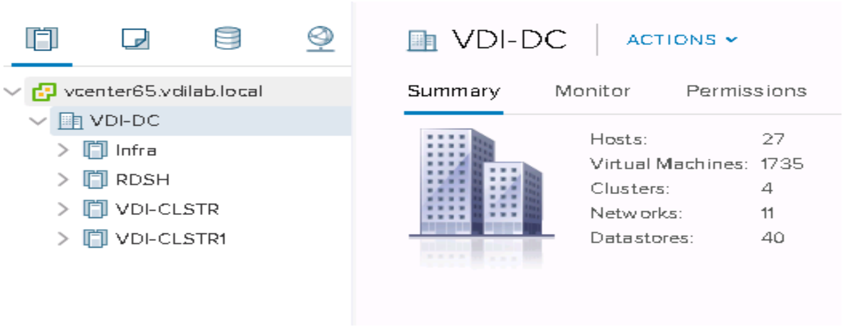

Hosts:	27
Virtual Machines:	1735
Clusters:	4
Networks:	11
Datastores:	40

Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution. Figure 24 illustrates the configuration topology for this solution.

Configuration Topology for Scalable VMware Horizon Mixed Workload

Component Layers

Figure 24 Solution Component Layers

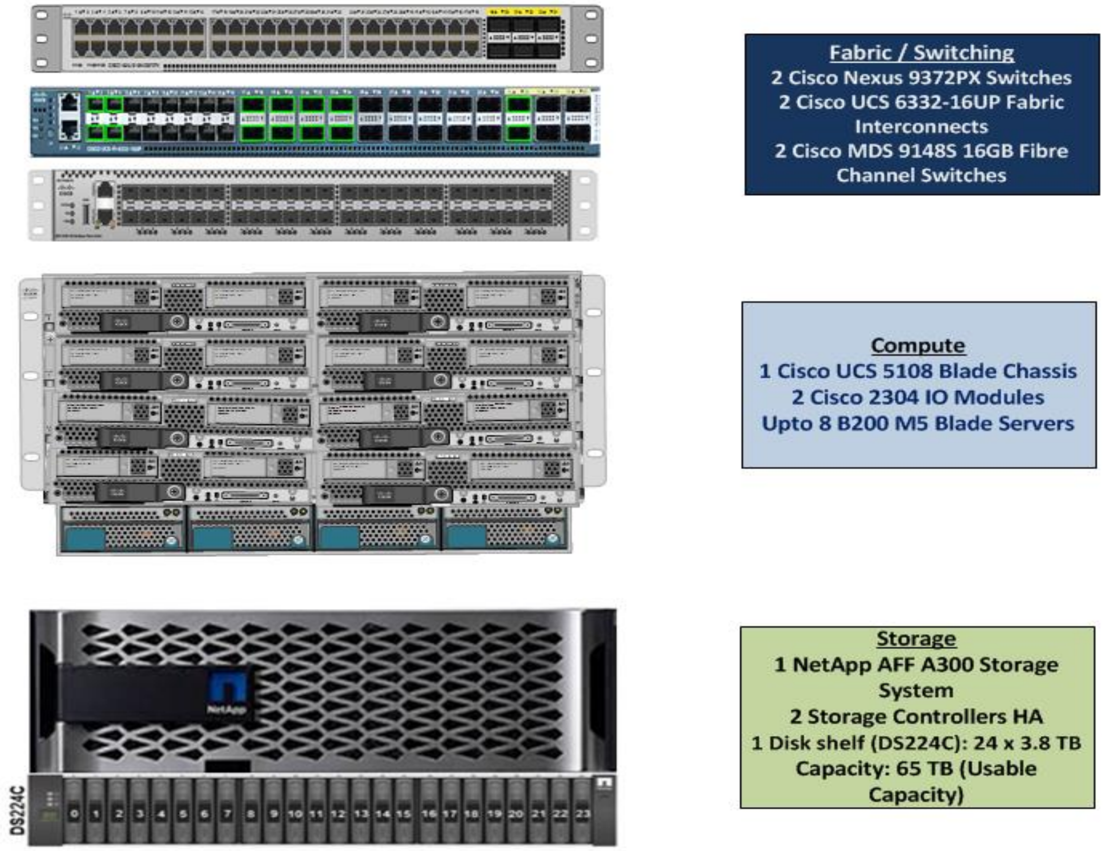
Fabric / Switching
2 Cisco Nexus 9372PX Switches
2 Cisco UCS 6332-16UP Fabric Interconnects
2 Cisco MDS 9148S 16GB Fibre Channel Switches
Compute
1 Cisco UCS 5108 Blade Chassis
2 Cisco 2304 IO Modules
Upto 8 B200 M5 Blade Servers
Storage
1 NetApp AFF A300 Storage System
2 Storage Controllers HA
1 Disk shelf (DS224C): 24 x 3.8 TB
Capacity: 65 TB (Usable Capacity)
DS224C
NetApp

Figure 24 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 AFF A300 array

Solution Cabling

The following sections detail the physical connectivity configuration of the FlexPod 5000 seat VMware Horizon 7 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 AFF A300 to the Cisco 6332-16UP Fabric Interconnects via Cisco MDS 9148S 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 25 illustrates a cabling diagram for a VMware Horizon configuration using the Cisco Nexus 9000, Cisco MDS 9100 Series, and NetApp AFF A300 array.

Figure 25 FlexPod 5000 Seat Cabling Diagram

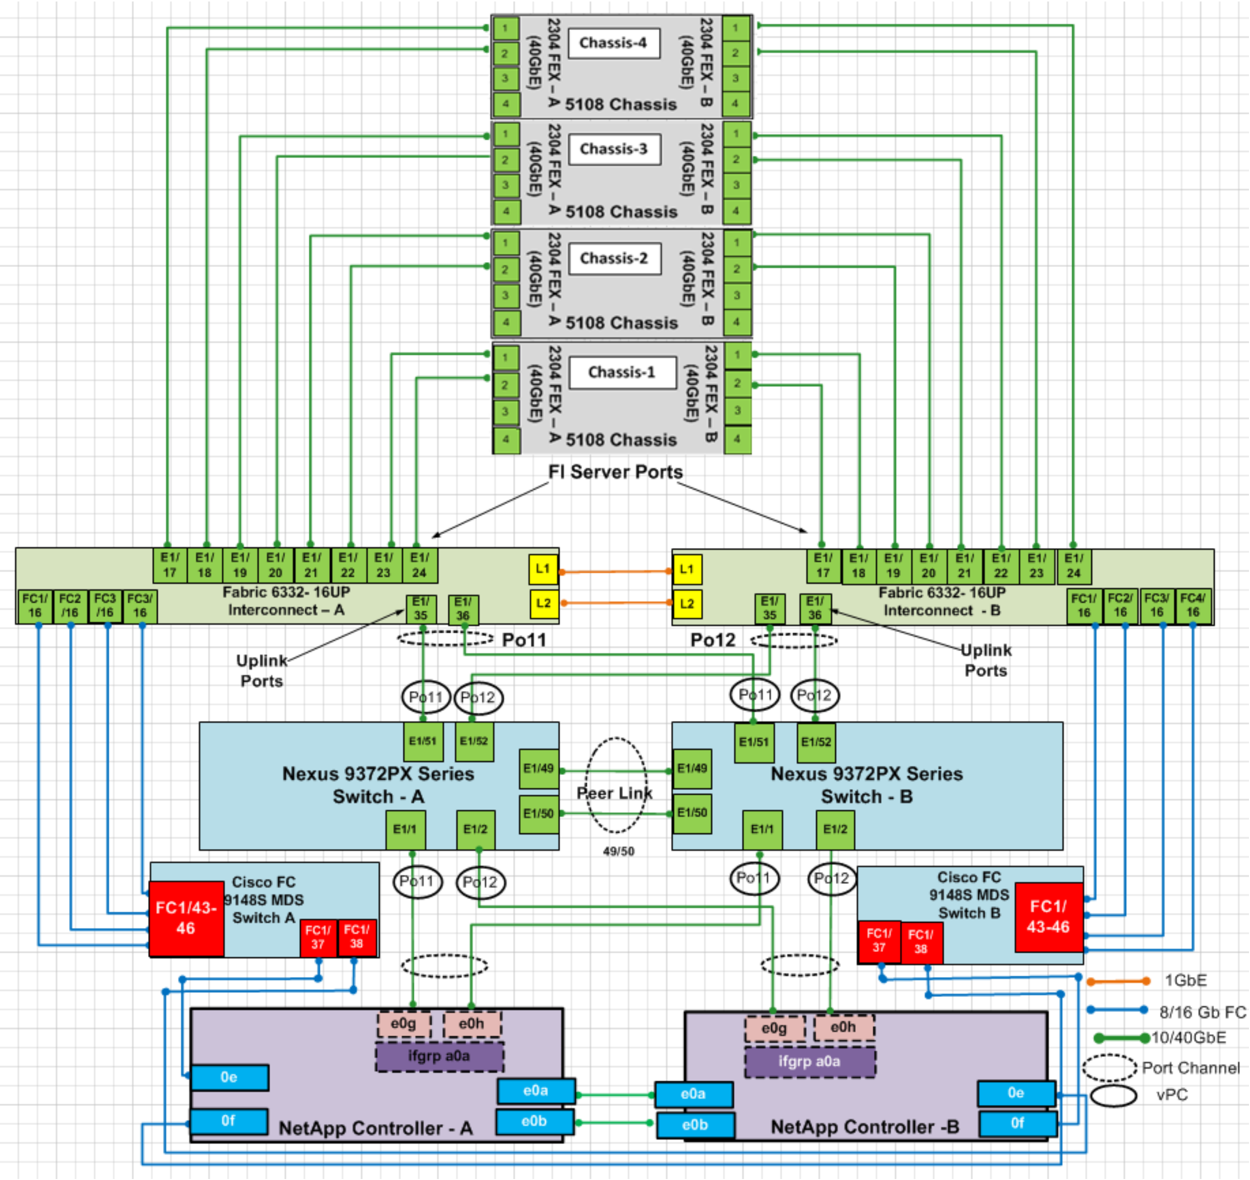

Cisco Nexus Switch Cabling Details

Table 11 Cisco Nexus 9372-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/51	40GbE	Cisco UCS fabric interconnect A	Eth1/35
	Eth1/52	40GbE	Cisco UCS fabric interconnect B	Eth1/36
	Eth1/49	40GbE	Cisco Nexus 9372 B	Eth1/49

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/50	40GbE	Cisco Nexus 9372 B	Eth1/50
	MGMT0	GbE	GbE management switch	Any

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

Table 12 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/51	40GbE	Cisco UCS fabric interconnect A	Eth1/35
	Eth1/52	40GbE	Cisco UCS fabric interconnect B	Eth1/36
	Eth1/49	40GbE	Cisco Nexus 9372 A	Eth1/49
	Eth1/50	40GbE	Cisco Nexus 9372 A	Eth1/50
	MGMT0	GbE	GbE management switch	Any

Cisco UCS 6332-16UPUP Fabric Interconnect Cabling

Table 13 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/35	40GbE	Cisco Nexus 9372 A	Eth1/51
	Eth1/36	40GbE	Cisco Nexus 9372 B	Eth1/52
	Eth1/17-1/18	40GbE	UCS 5108 Blade Chassis IOM-A, Chassis 1	IOM 1-2
	Eth1/19-1/20	40GbE	UCS 5108 Blade Chassis IOM-A, Chassis 2	IOM 1-2
	Eth1/21-1/22	40GbE	UCS 5108 Blade Chassis IOM-A, Chassis 3	IOM 1-2
	Eth 1/23-/24	40GbE	UCS 5108 Blade Chassis IOM-A, Chassis 4	IOM 1-2
	MGMT0	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2
	FC 1/1	16Gb FC	Cisco MDS 9148S-A	FC 1/43
	FC 1/2	16Gb FC	Cisco MDS 9148S-A	FC 1/44
	FC 1/3	16Gb FC	Cisco MDS 9148S-A	FC 1/45
	FC 1/4	16Gb FC	Cisco MDS 9148S-A	FC 1/46

Table 14 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/35	40GbE	Cisco Nexus 9372 A	Eth1/51
	Eth1/36	40GbE	Cisco Nexus 9372 B	Eth1/52
	Eth1/17-1/18	40GbE	UCS 5108 Blade Chassis IOM-B, Chassis 1	IOM 1-2
	Eth1/19-1/20	40GbE	UCS 5108 Blade Chassis IOM-B, Chassis 2	IOM 1-2
	Eth1/21-1/22	40GbE	UCS 5108 Blade Chassis IOM-B, Chassis 3	IOM 1-2
	Eth 1/23-1/24	40GbE	UCS 5108 Blade Chassis IOM-B, Chassis 4	IOM 1-2
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2
	FC 1/1	16Gb FC	Cisco MDS 9148S-B	FC 1/43
	FC1/2	16Gb FC	Cisco MDS 9148S-B	FC 1/44
	FC 1/3	16Gb FC	Cisco MDS 9148S-B	FC 1/45
	FC 1/4	16Gb FC	Cisco MDS 9148S-B	FC 1/46

Figure 26 Cable connectivity between Cisco UCS 6332-16UP Fabric Interconnects and Cisco Nexus 9372PX Switches

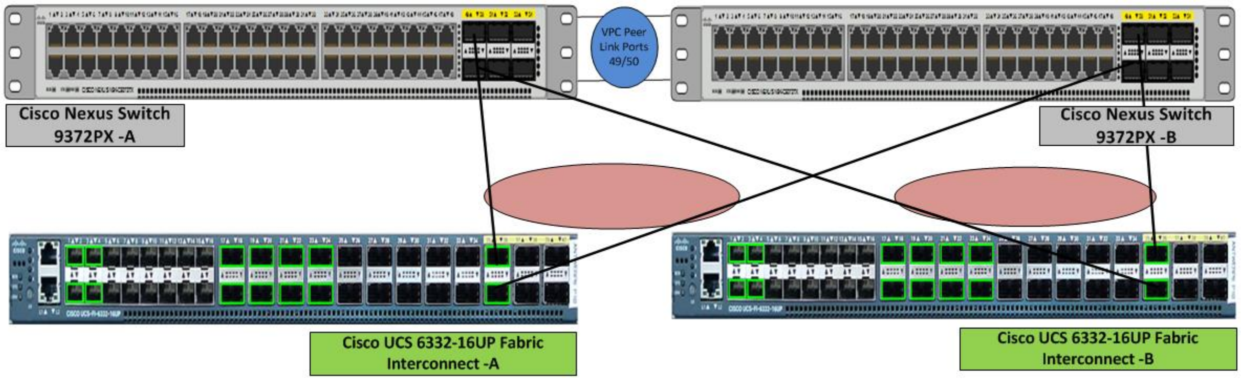
VPC Peer Link Ports 49/50
Cisco Nexus Switch 9372PX -A
Cisco Nexus Switch 9372PX -B
Cisco UCS 6332-16UP Fabric Interconnect -A
Cisco UCS 6332-16UP Fabric Interconnect -B

Cisco MDS 9148S Cabling

Figure 26 illustrates the cable connectivity between the Cisco MDS 9148S and the Cisco 6332 Fabric Interconnects and the AFF A300 storage.

We used four 16Gb FC connections from each Fabric Interconnect to each MDS switch and utilized two 16Gb FC connections from the AFF A300 storage controller to each MDS switch.

Table 15 Cisco MDS 9148S A Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148S-A	fc1/37	16Gb FC	AFF A300-01	Port 0g
	fc1/38	16Gb FC	AFF A300-02	Port 0g
	fc1/43	16Gb FC	Cisco 6332-16UP Fabric Interconnect-A	Fc 1/1
	fc1/44	16Gb FC	Cisco 6332-16UP Fabric Interconnect-A	Fc 1/2
	fc1/45	16Gb FC	Cisco 6332-16UP Fabric Interconnect-A	Fc 1/3
	fc1/46	16Gb FC	Cisco 6332-16UP Fabric Interconnect-A	Fc 1/4

Table 16 Cisco MDS 9148S B Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9148S-B	fc1/37	16Gb FC	AFF A300-01	Port 0h

Local Device	Local Port	Connection	Remote Device	Remote Port
	fc1/38	16Gb FC	AFF A300-02	Port 0h
	fc1/43	16Gb FC	Cisco 6332-16UPUP Fabric Interconnect-B	Fc 2/1
	fc1/44	16Gb FC	Cisco 6332-16UPUP Fabric Interconnect-B	Fc 2/2
	fc1/45	16Gb FC	Cisco 6332-16UPUP Fabric Interconnect-B	Fc 2/3
	fc1/46	16Gb FC	Cisco 6332-16UPUP Fabric Interconnect-B	Fc 3/4

AFF A300 to MDS SAN Fabric Connectivity

Figure 27 illustrates the NetApp AFF- MDS A and B Switches using VSAN 400 for Fabric A and VSAN 401 Configured for Fabric B.

Figure 27 AFF A300 storage connectivity to Cisco MDS FC Switches

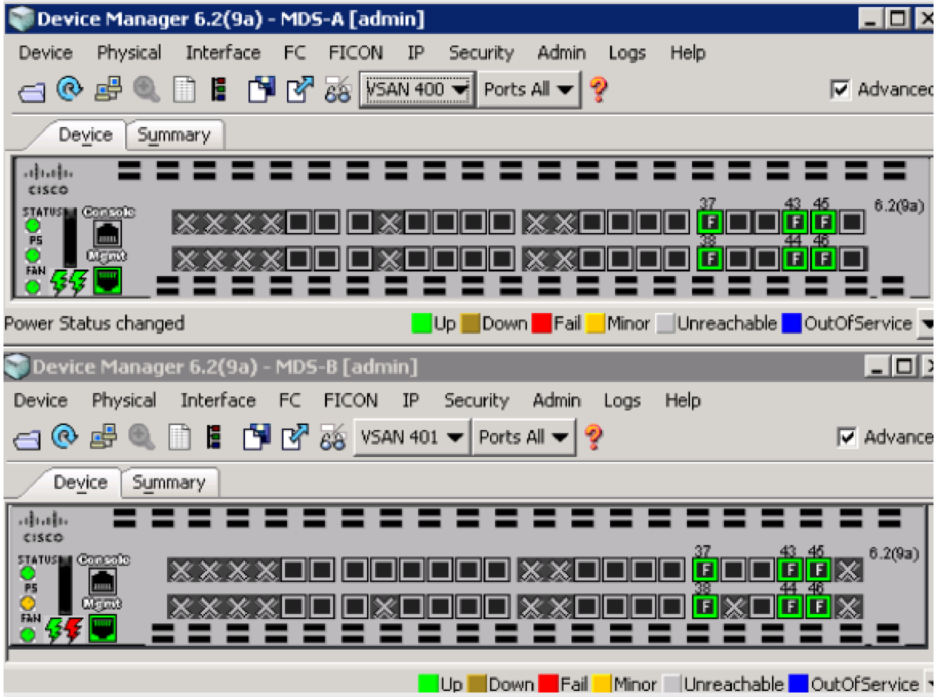
Device Manager 6.2(9a) - MDS-A [admin]
Device Physical Interface FC FICON IP Security Admin Logs Help
VSAN 400
Ports All
☑ Advanced
Device Summary
cisco
STATUS
Console
PS
Mgmt
FAN
37 43 45
38 44 46
6.2(9a)
Power Status changed
🟩Up 🟫Down 🟥Fail 🟨Minor ⬜Unreachable 🟦OutOfService
Device Manager 6.2(9a) - MDS-B [admin]
Device Physical Interface FC FICON IP Security Admin Logs Help
VSAN 401
Ports All
☑ Advanced
Device Summary
cisco
STATUS
Console
PS
Mgmt
FAN
37 43 45
38 44 46
6.2(9a)
🟩Up 🟫Down 🟥Fail 🟨Minor ⬜Unreachable 🟦OutOfService

Figure 28 Fibre Channel Cable Connectivity from AFF A300 to Cisco MDS 9148S to Cisco 6332-16UP Fabric Interconnects

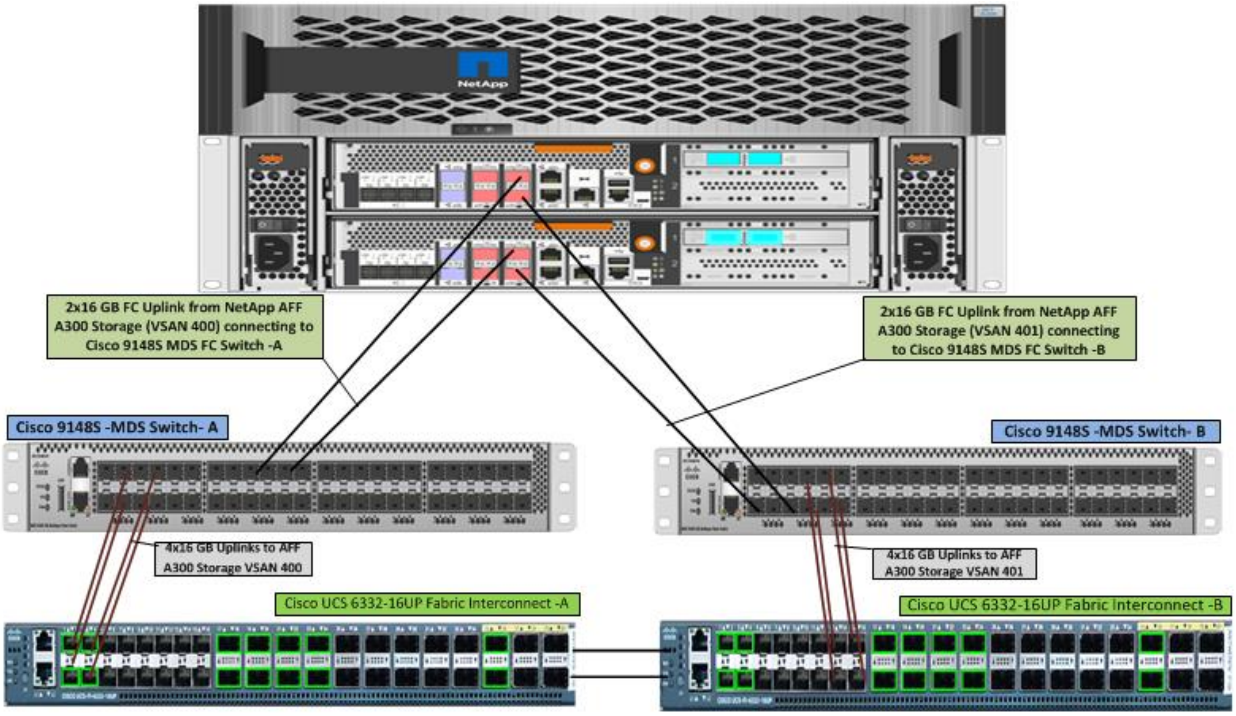
NetApp
2x16 GB FC Uplink from NetApp AFF A300 Storage (VSAN 400) connecting to Cisco 9148S MDS FC Switch -A
2x16 GB FC Uplink from NetApp AFF A300 Storage (VSAN 401) connecting to Cisco 9148S MDS FC Switch -B
Cisco 9148S -MDS Switch- A
Cisco 9148S -MDS Switch- B
4x16 GB Uplinks to AFF A300 Storage VSAN 400
4x16 GB Uplinks to AFF A300 Storage VSAN 401
Cisco UCS 6332-16UP Fabric Interconnect -A
Cisco UCS 6332-16UP Fabric Interconnect -B

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

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

Cisco UCS Manager Software Version 3.2(1d)

This document assumes the use of Cisco UCS Manager Software version 3.2(1d). To upgrade the Cisco UCS Manager software and the Cisco UCS 6332-16UP Fabric Interconnect software to a higher version of the firmware,) refer to Cisco UCS Manager Install and Upgrade Guides.

Configure Fabric Interconnects at Console

To configure the fabric interconnect, complete the following steps:

1. Connect a console cable to the console port on what will become the primary fabric interconnect.
2. If the fabric interconnects was previously deployed and you want to erase it to redeploy, follow these steps:
   a. Login with the existing 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.

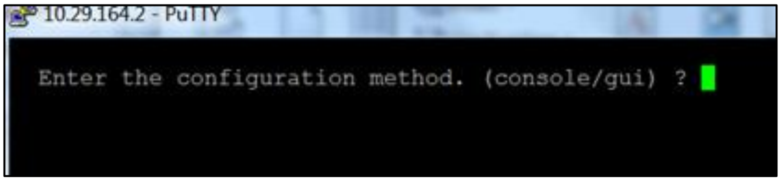

```
 Enter the configuration method. (console/gui) ?
```

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

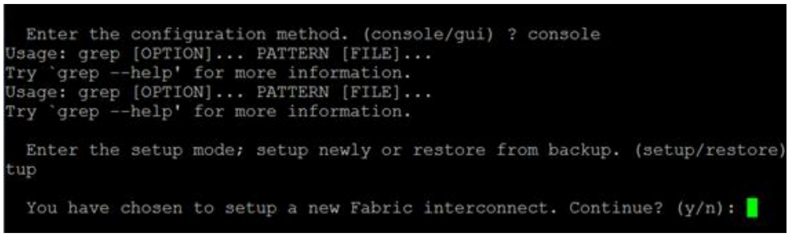

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

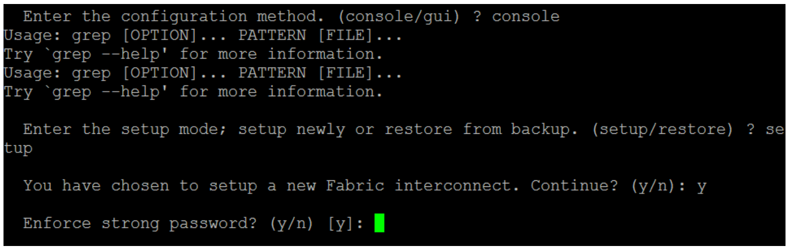

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

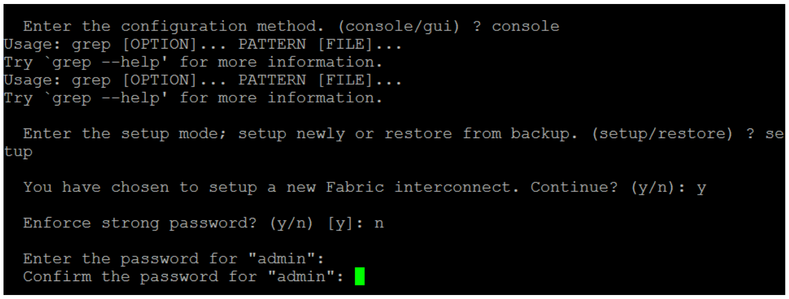

```
  Enter the configuration method. (console/gui) ? console
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.
Usage: grep [OPTION]... PATTERN [FILE]...
Try `grep --help' for more information.

  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.

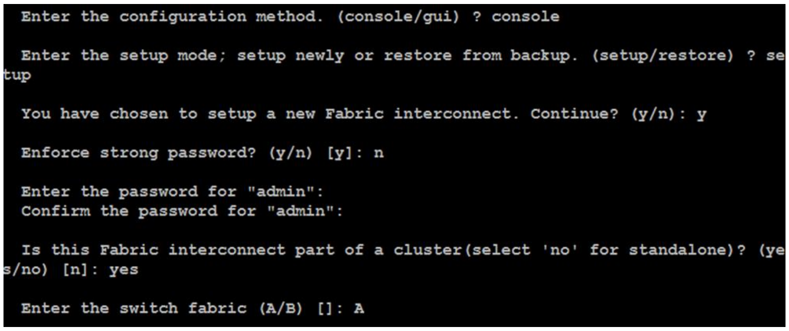

```
  Enter the configuration method. (console/gui) ? console

  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.

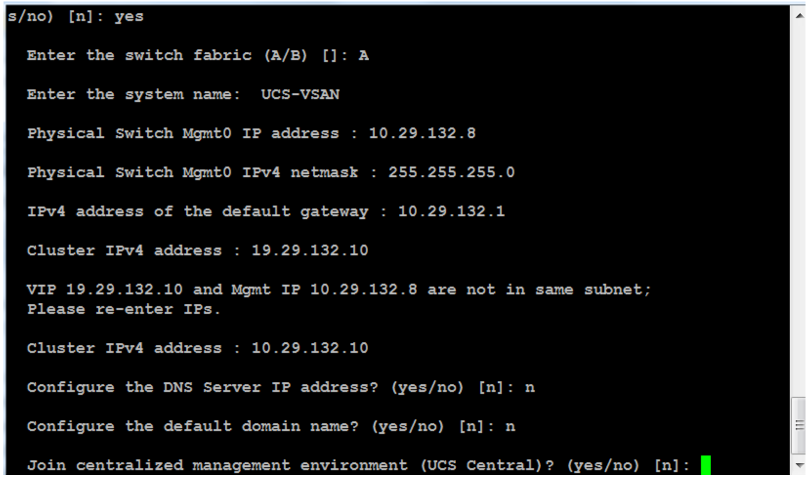

```
s/no) [n]: yes

  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”.

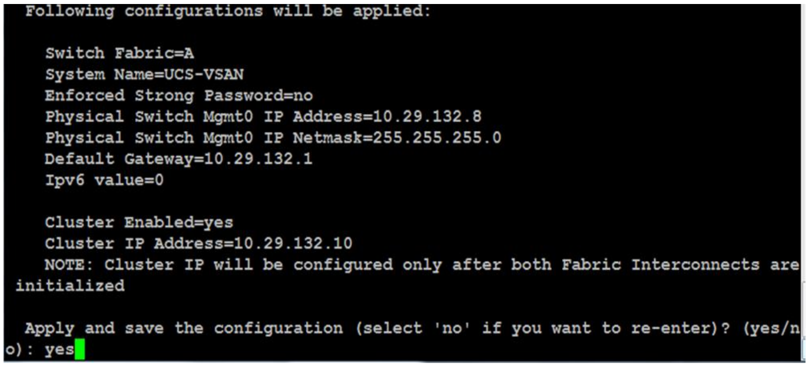

```
 Following configurations will be applied:

   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.

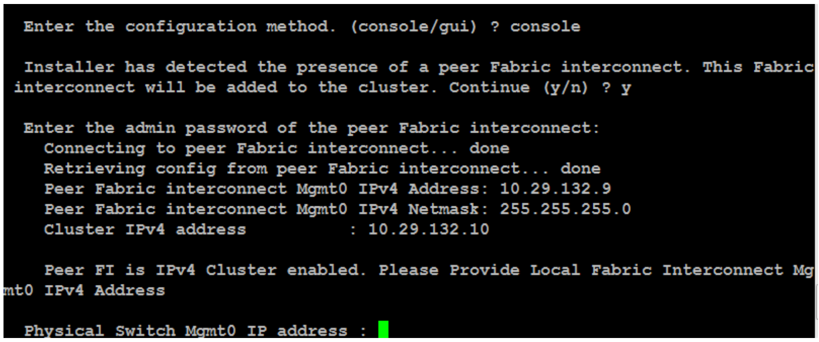

```
 Enter the configuration method. (console/gui) ? console

 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.

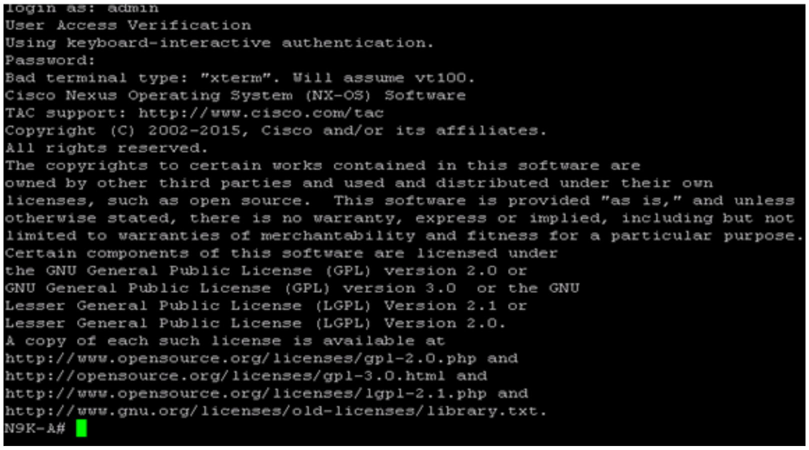

```
login as: admin
User Access Verification
Using keyboard-interactive authentication.
Password:
Bad terminal type: "xterm". Will assume vt100.
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2015, Cisco and/or its affiliates.
All rights reserved.
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 6332-16UP Fabric Interconnect cluster address.

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

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

4. When prompted, enter admin as the 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.

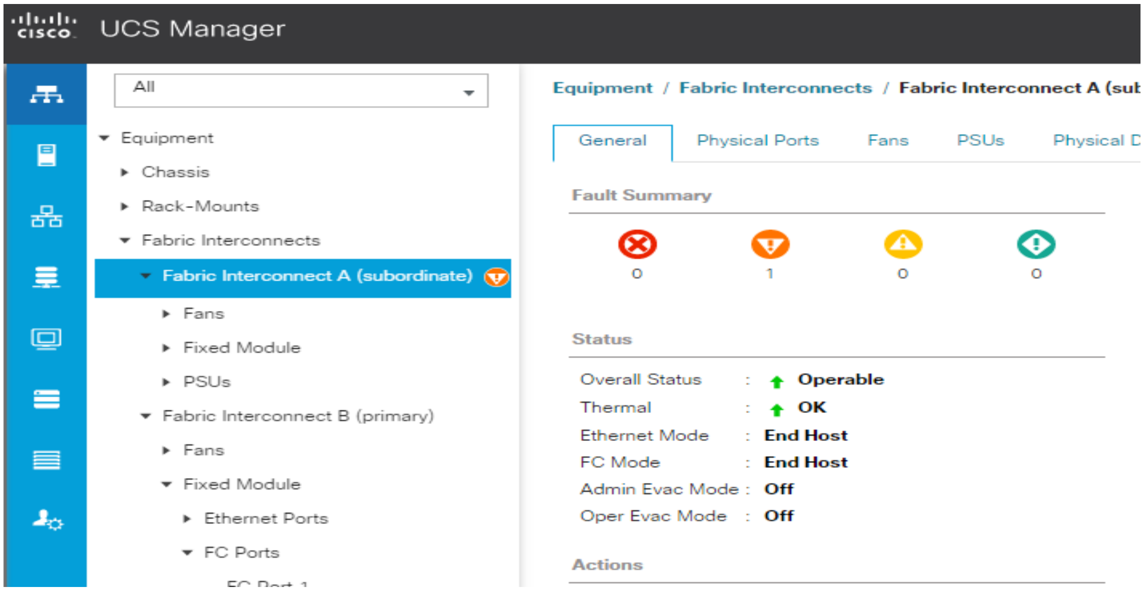

Status

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

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

Configure Fibre Channel Uplink Ports

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

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

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

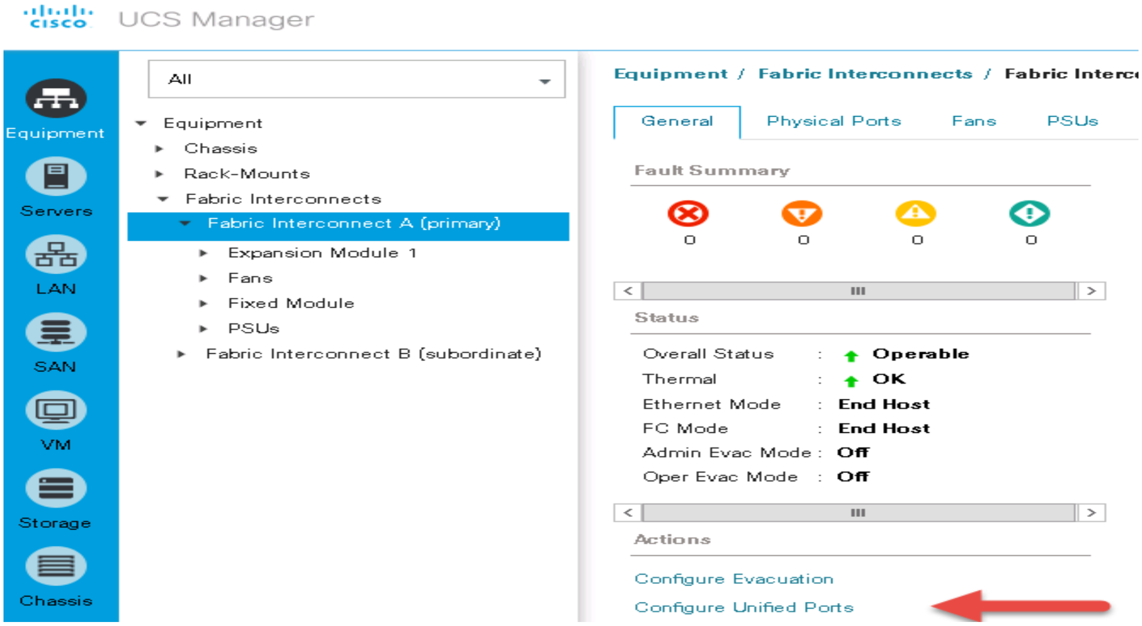

Overall Status : Operable
Thermal : OK
Ethernet Mode : End Host
FC Mode : End Host
Admin Evac Mode : Off
Oper Evac Mode : Off

3. Click Configure Expansion Module Ports.

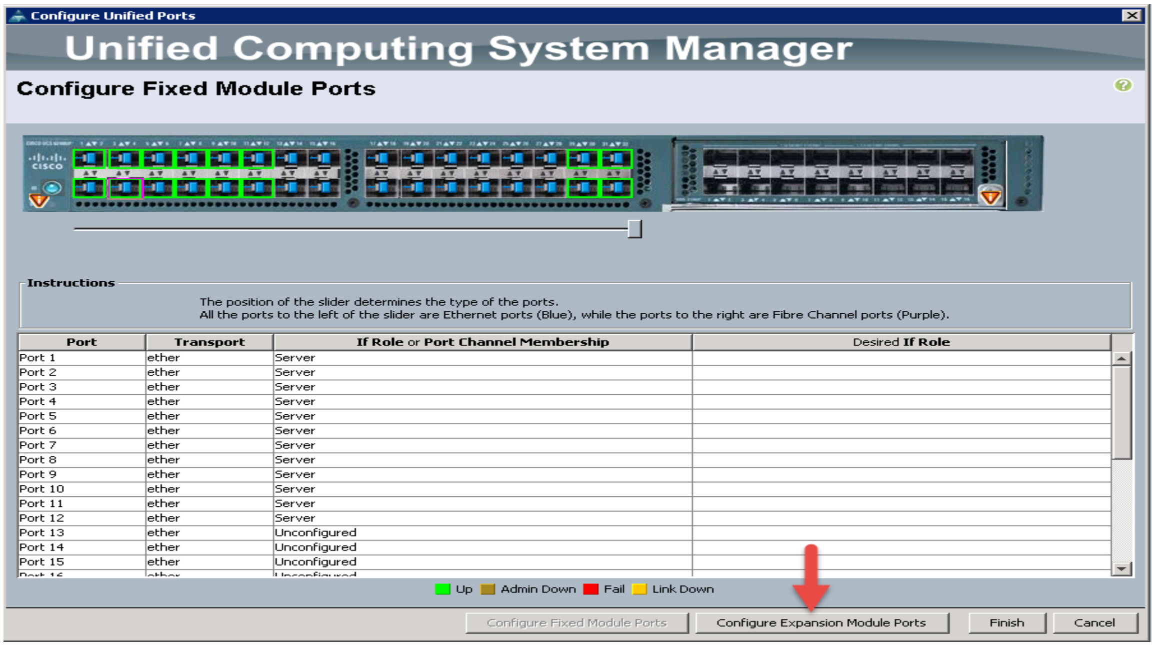
Configure Unified Ports
Unified Computing System Manager
Up
Admin Down
Fail
Link Down
Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

Configure Fixed 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	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	

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.

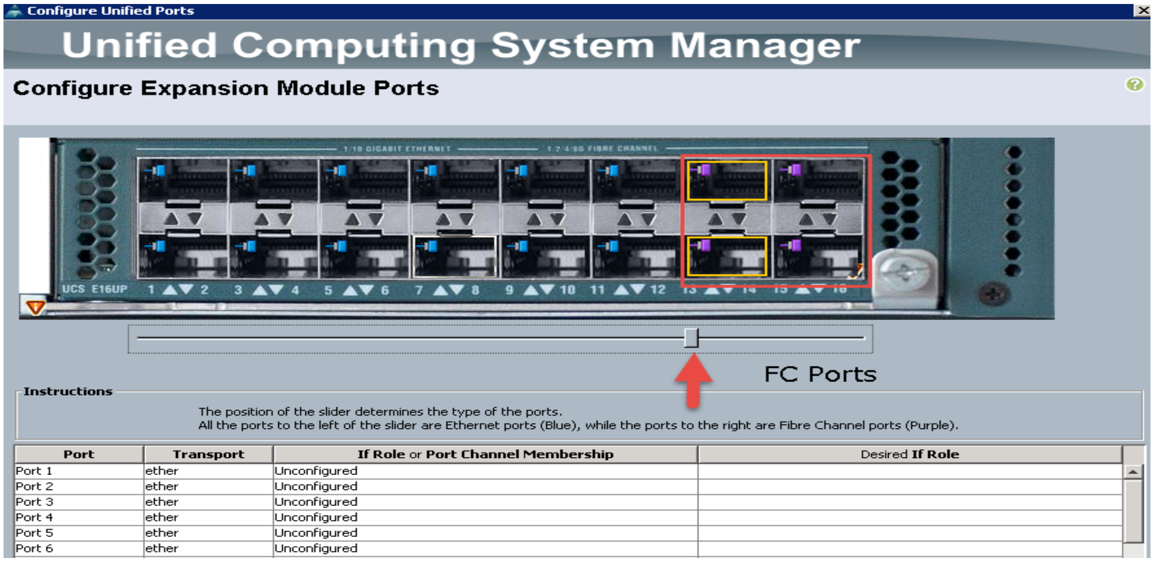

Configure Expansion Module Ports

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.

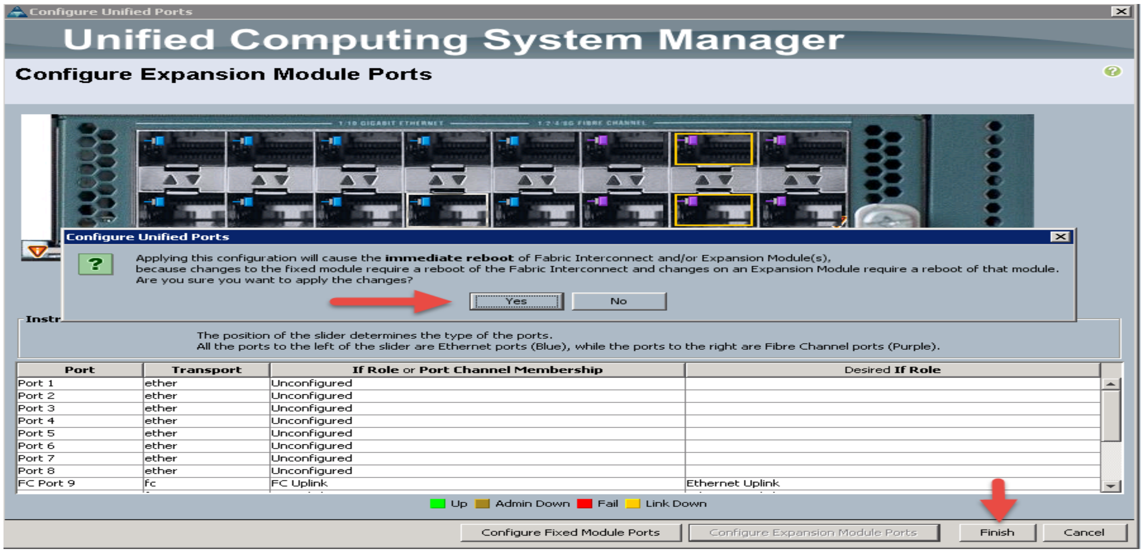
Configure Unified Ports
Unified Computing System Manager
1/10 GIGABIT ETHERNET
1/2/4/8G FIBRE CHANNEL
Instr
Yes
No
Configure Fixed Module Ports
Configure Expansion Module Ports
Finish
Cancel

Configure Expansion Module Ports

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

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

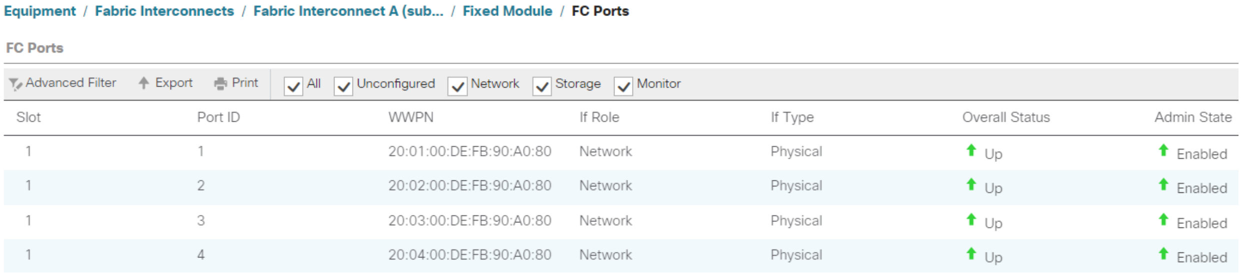

Equipment / Fabric Interconnects / Fabric Interconnect A (sub... / Fixed Module / FC Ports

FC Ports

☑ All ☑ Unconfigured ☑ Network ☑ Storage ☑ Monitor

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20:01:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
1	2	20:02:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
1	3	20:03:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled
1	4	20:04:00:DE:FB:90:A0:80	Network	Physical	↑ Up	↑ Enabled

6. Repeat this procedure for Fabric Interconnect B.

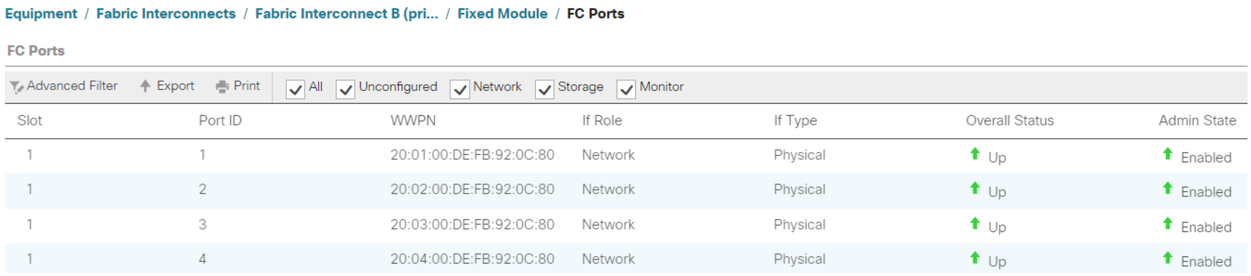
Equipment / Fabric Interconnects / Fabric Interconnect B (pri... / Fixed Module / FC Ports
Advanced Filter
Export
Print
☑ All ☑ Unconfigured ☑ Network ☑ Storage ☑ Monitor

FC Ports

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20:01:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
1	2	20:02:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
1	3	20:03:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled
1	4	20:04:00:DE:FB:92:0C:80	Network	Physical	↑ Up	↑ Enabled

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

Edit Chassis Discovery Policy

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

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

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

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

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

4. Set the Link Grouping Preference to Port Channel.

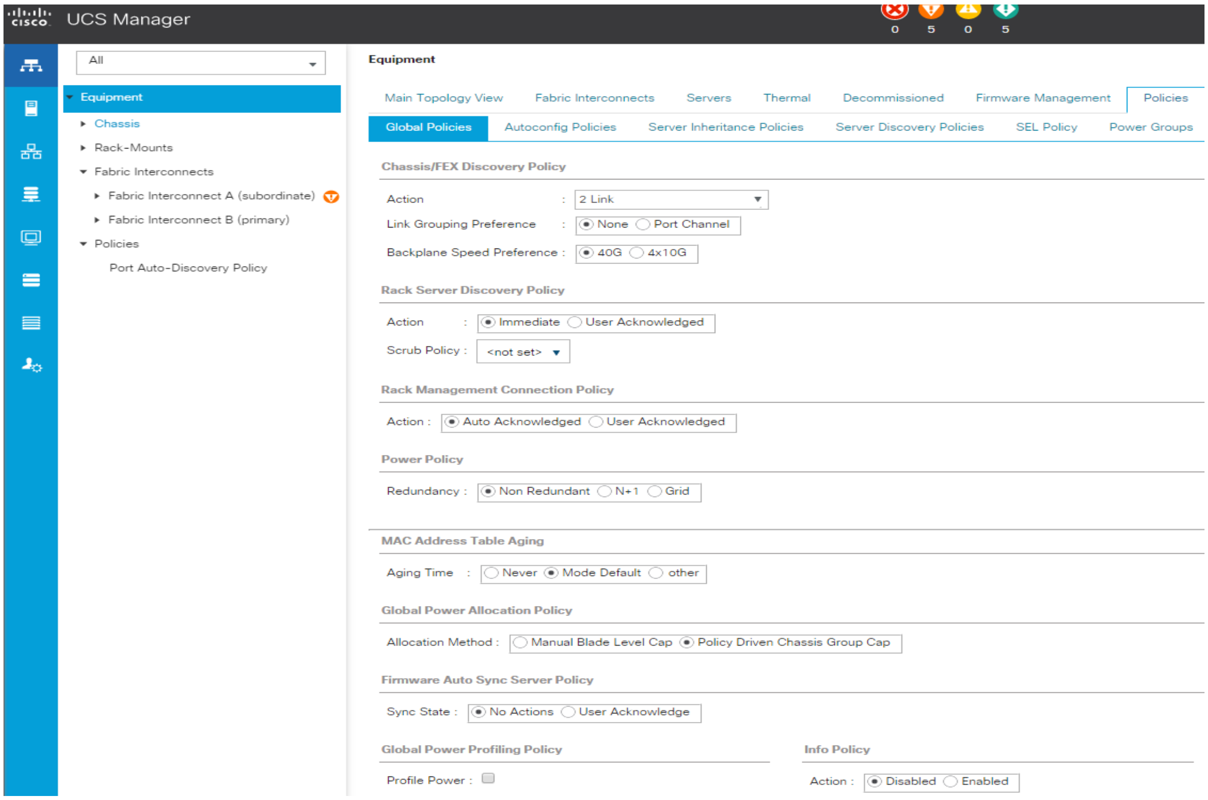

Equipment

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : ● None ○ Port Channel

Backplane Speed Preference : ● 40G ○ 4x10G

Rack Server Discovery Policy

Action : ● Immediate ○ User Acknowledged

Scrub Policy : <not set>

Rack Management Connection Policy

Action : ● Auto Acknowledged ○ User Acknowledged

Power Policy

Redundancy : ● Non Redundant ○ N+1 ○ Grid

MAC Address Table Aging

Aging Time : ○ Never ● Mode Default ○ other

Global Power Allocation Policy

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

Firmware Auto Sync Server Policy

Sync State : ● No Actions ○ User Acknowledge

Global Power Profiling Policy

Profile Power : □

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.

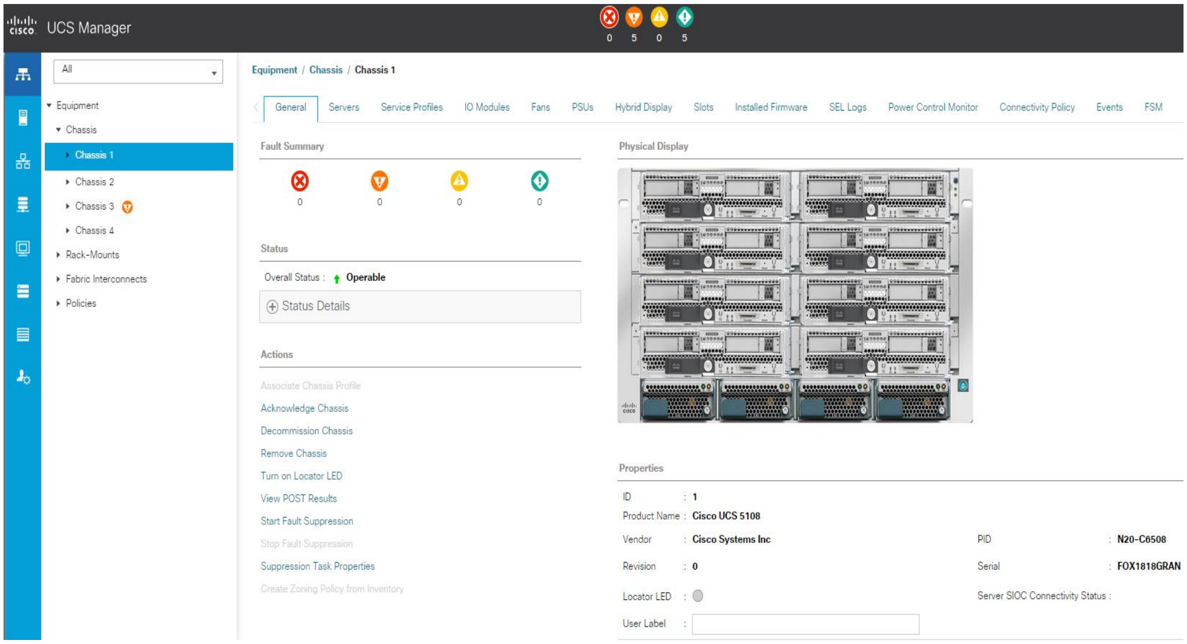

Overall Status : Operable

ID : 1
Product Name : Cisco UCS 5108
Vendor : Cisco Systems Inc
Revision : 0
Locator LED : ○
User Label :
PID : N20-C6508
Serial : FOX1818GRAN
Server SIOC Connectivity Status :

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

5. Repeat for each of the remaining chassis.

Synchronize Cisco UCS to NTP

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

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

2. Select All > Timezone Management.

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

4. Click Save Changes and then click OK.

5. Click Add NTP Server.

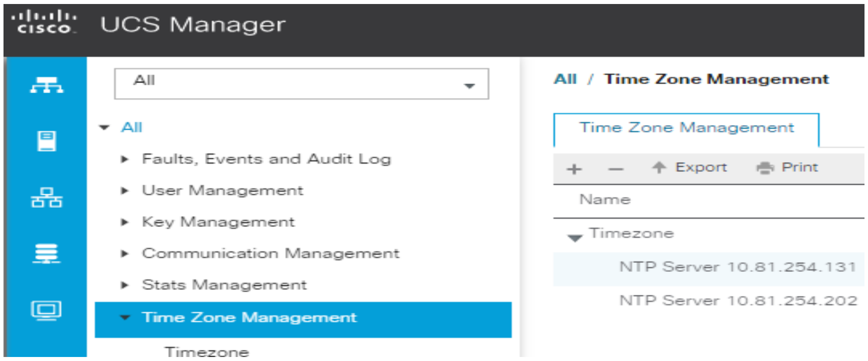

NTP Server 10.81.254.131

NTP Server 10.81.254.202

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

7. Click OK.

Enable Server and Ethernet Uplink Ports

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

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

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

3. Expand Ethernet Ports.

4. Select ports 17 through 24 that are connected to the Cisco IO Modules of the four B-Series 5108 Chassis, right-click them, and select Configure as Server Port.

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

6. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that ports have been configured correctly in the in the Role column.

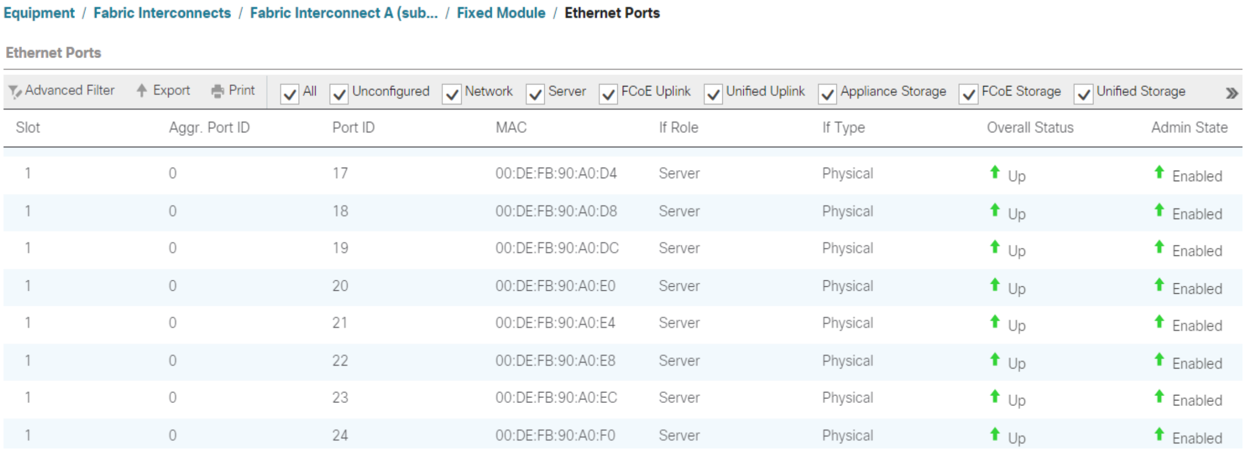

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	17	00:DE:FB:90:A0:D4	Server	Physical	Up	Enabled
1	0	18	00:DE:FB:90:A0:D8	Server	Physical	Up	Enabled
1	0	19	00:DE:FB:90:A0:DC	Server	Physical	Up	Enabled
1	0	20	00:DE:FB:90:A0:E0	Server	Physical	Up	Enabled
1	0	21	00:DE:FB:90:A0:E4	Server	Physical	Up	Enabled
1	0	22	00:DE:FB:90:A0:E8	Server	Physical	Up	Enabled
1	0	23	00:DE:FB:90:A0:EC	Server	Physical	Up	Enabled
1	0	24	00:DE:FB:90:A0:F0	Server	Physical	Up	Enabled

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

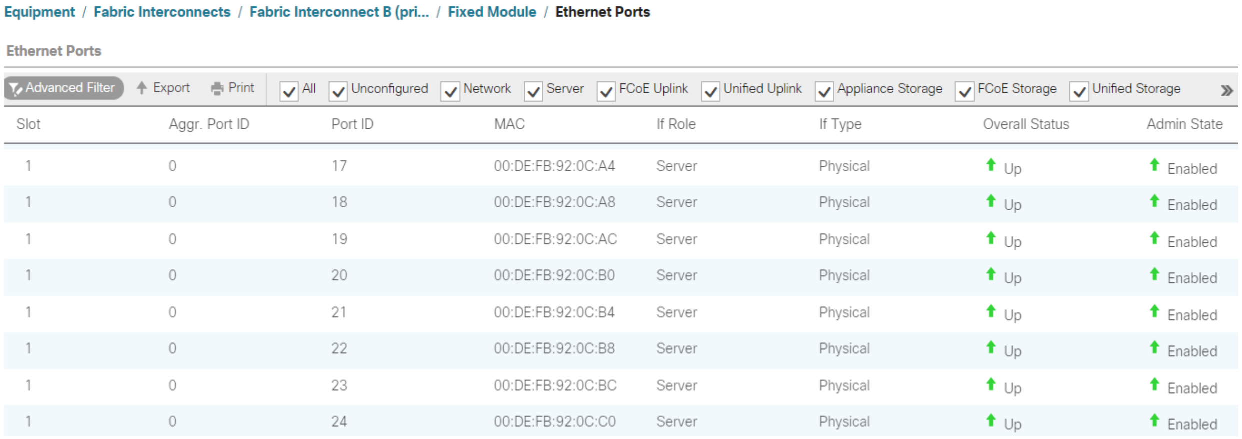

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	17	00:DE:FB:92:0C:A4	Server	Physical	⬆ Up	⬆ Enabled
1	0	18	00:DE:FB:92:0C:A8	Server	Physical	⬆ Up	⬆ Enabled
1	0	19	00:DE:FB:92:0C:AC	Server	Physical	⬆ Up	⬆ Enabled
1	0	20	00:DE:FB:92:0C:B0	Server	Physical	⬆ Up	⬆ Enabled
1	0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

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

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

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

3. Expand Ethernet Ports.

4. Select ports 35 through 36 that are connected to the Nexus 9372PX 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 switches are now configured as network ports.

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

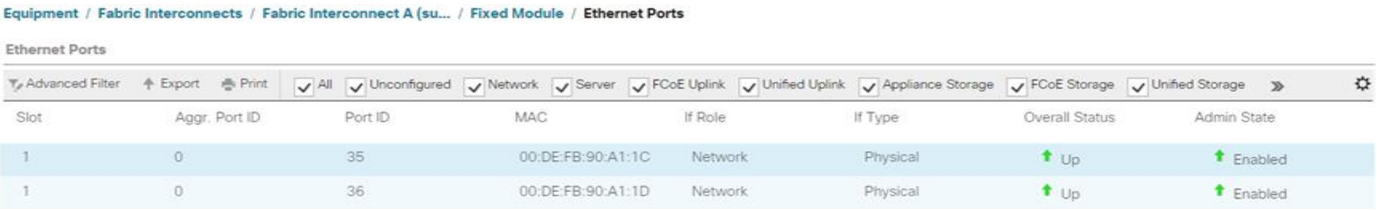

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	35	00:DE:FB:90:A1:1C	Network	Physical	⬆ Up	⬆ Enabled
1	0	36	00:DE:FB:90:A1:1D	Network	Physical	⬆ Up	⬆ Enabled

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

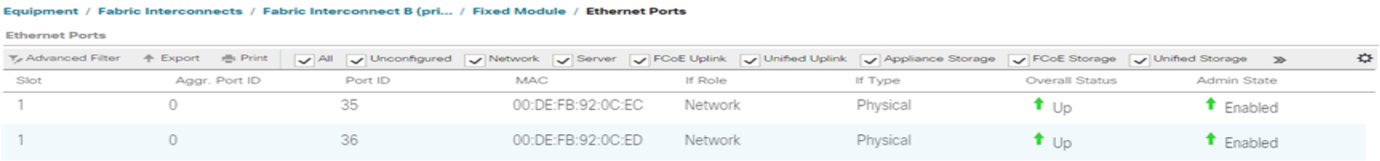

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	35	00:DE:FB:92:0C:EC	Network	Physical	⬆ Up	⬆ Enabled
1	0	36	00:DE:FB:92:0C:ED	Network	Physical	⬆ Up	⬆ Enabled

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 11 as the unique ID of the port channel.
6. Enter FI-A-Uplink as the name of the port channel.
7. Click Next.

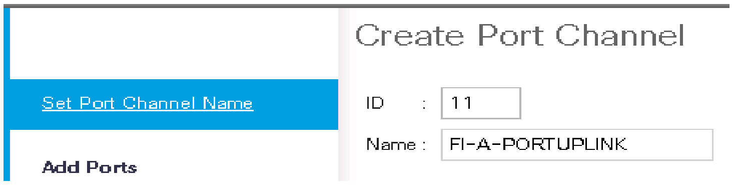

Create Port Channel

ID : 11

Name : FI-A-PORTUPLINK

8. Select ethernet ports 35-36 for the port channel.
9. Click Finish.

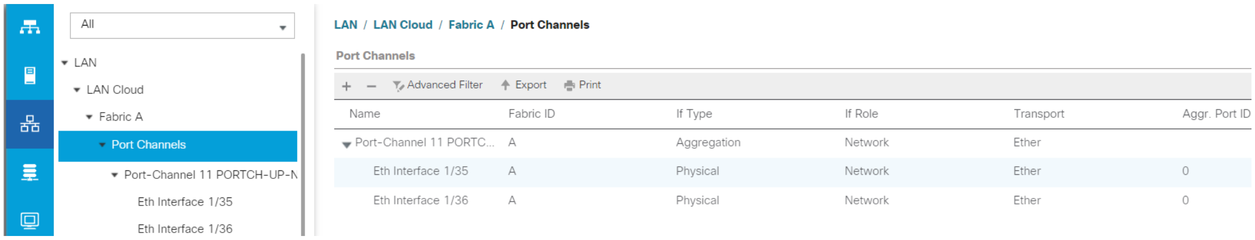

LAN / LAN Cloud / Fabric A / Port Channels

Port Channels

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

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

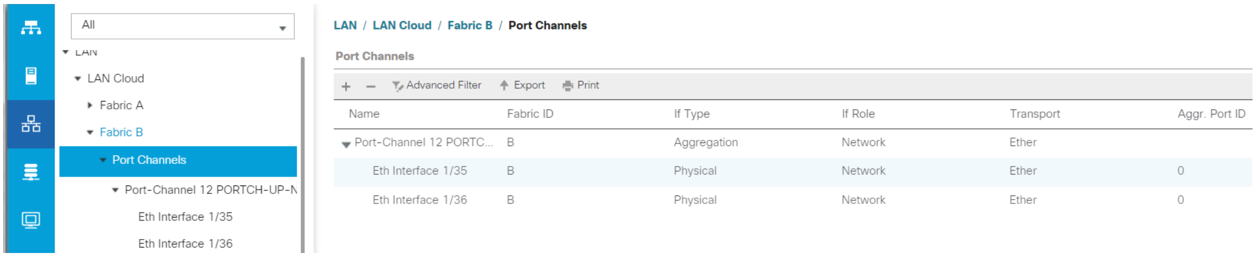

LAN / LAN Cloud / Fabric B / Port Channels

Port Channels

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

Create Uplink Port Channels to Cisco MDS 9148S Switches

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

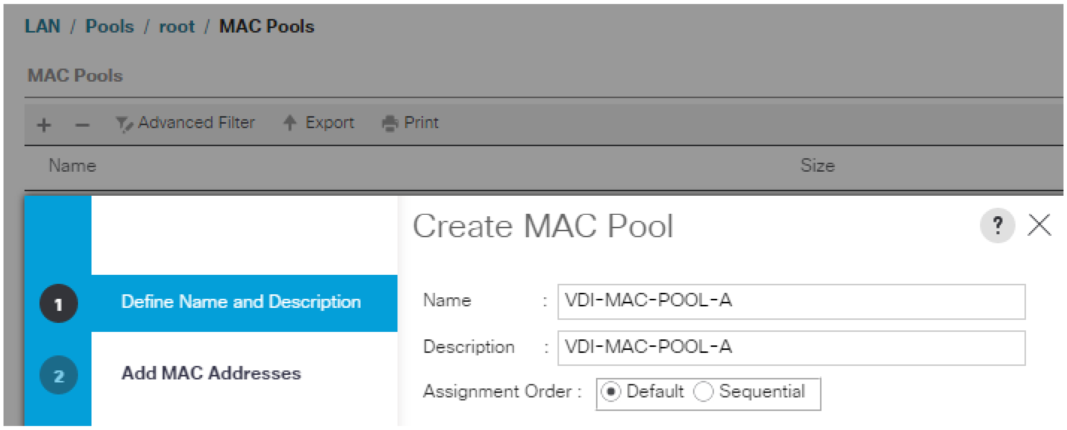
LAN / Pools / root / MAC Pools
MAC Pools
+ − Advanced Filter Export Print
Name
Size
1 Define Name and Description
2 Add MAC Addresses

Create MAC Pool

Name : VDI-MAC-POOL-A

Description : VDI-MAC-POOL-A

Assignment Order : ● Default ○ Sequential

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

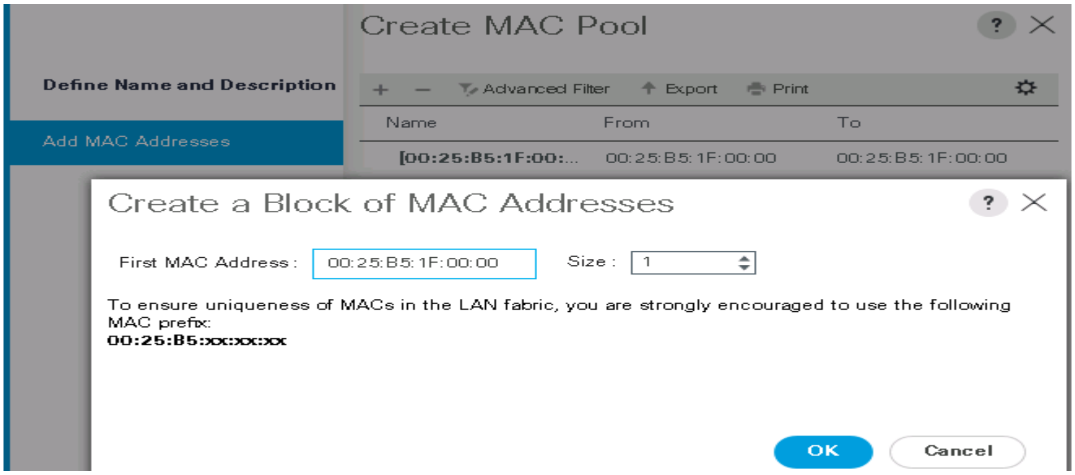

Create MAC Pool

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

Create a Block of MAC Addresses

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

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

8. Click OK, then click Finish.

9. In the confirmation message, click OK.

Create KVM IP Address Pool

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

To create the pool, complete the following steps:

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

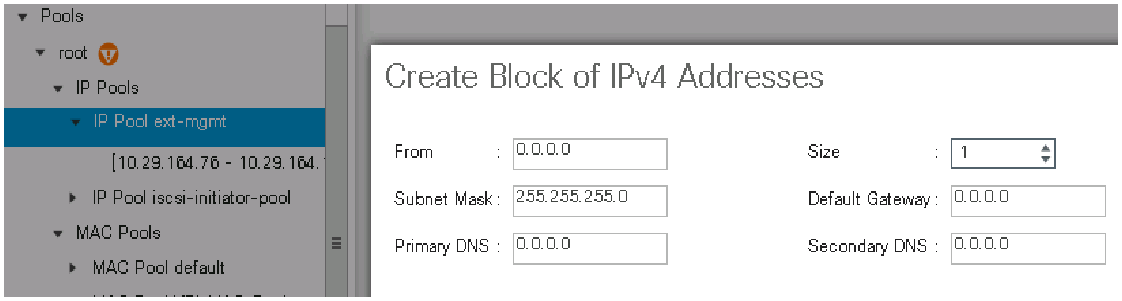

Create Block of IPv4 Addresses

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

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

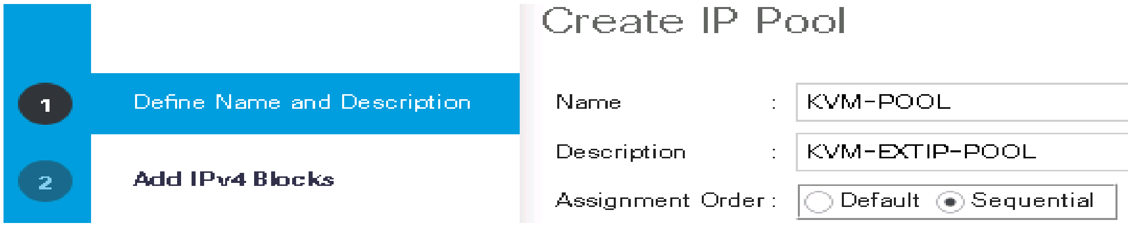

Create IP Pool

Name : KVM-POOL
Description : KVM-EXTIP-POOL
Assignment Order : ○ Default ● Sequential

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

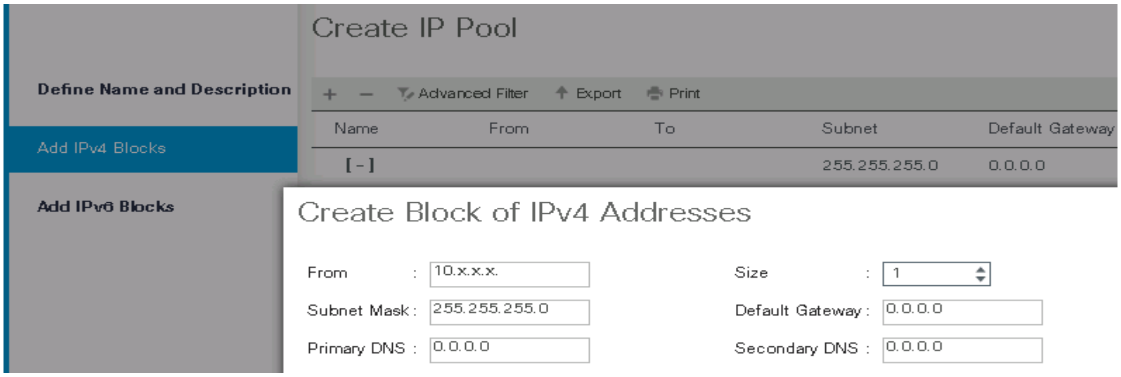

Create IP Pool

Name	From	To	Subnet	Default Gateway
[-]			255.255.255.0	0.0.0.0

Create Block of IPv4 Addresses

From :	10.x.x.x.	Size :	1
Subnet Mask :	255.255.255.0	Default Gateway :	0.0.0.0
Primary DNS :	0.0.0.0	Secondary DNS :	0.0.0.0

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

5. Click Finish.

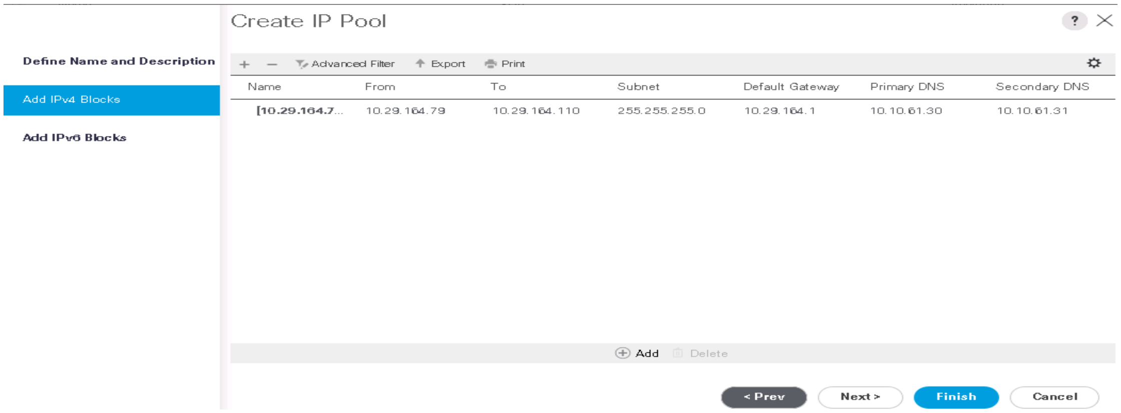

Create IP Pool

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

6. Click OK.

Create WWPN Pools

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

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

2. Select Pools > root.

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

4. Assign a name and optional description.

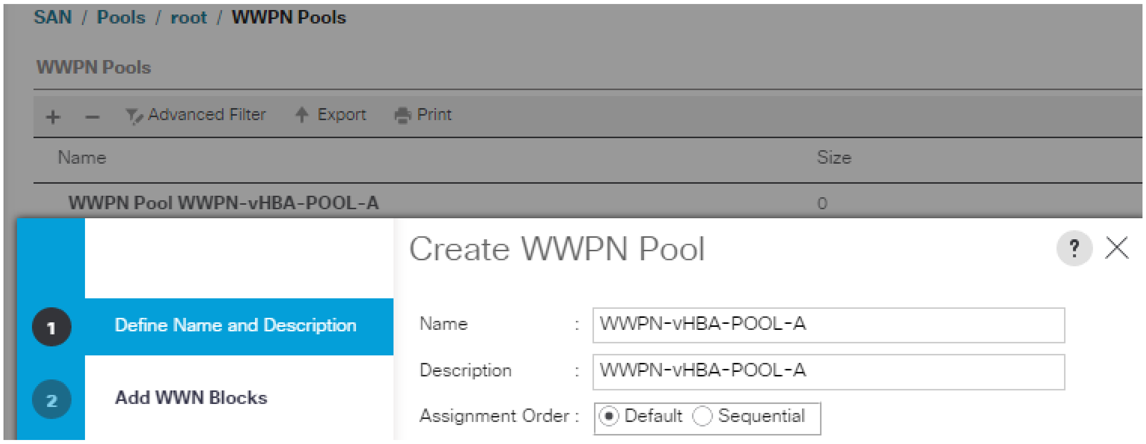
SAN / Pools / root / WWPN Pools
WWPN Pools
+ — Advanced Filter Export Print
Name
Size
WWPN Pool WWPN-vHBA-POOL-A
0
1 Define Name and Description
2 Add WWN Blocks

Create WWPN Pool

Name : WWPN-vHBA-POOL-A

Description : WWPN-vHBA-POOL-A

Assignment Order : ● Default ○ Sequential

5. Assignment order can remain Default.

6. Click Next.

7. Click Add to add block of Ports.

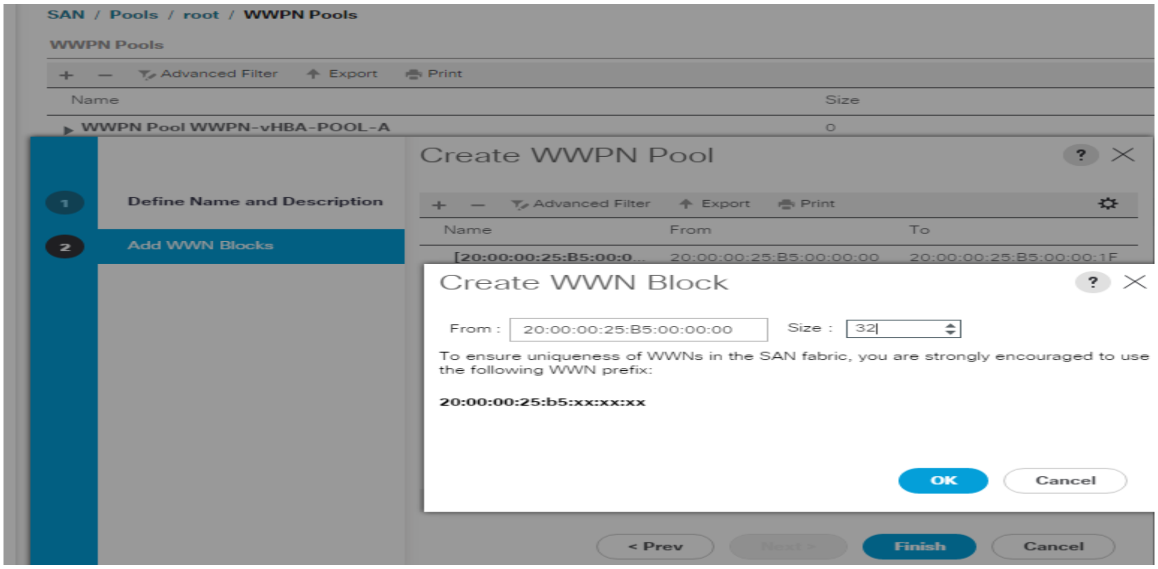
SAN / Pools / root / WWPN Pools
WWPN Pools
+ — Advanced Filter Export Print
Name
Size
WWPN Pool WWPN-vHBA-POOL-A
0
Create WWPN Pool
1 Define Name and Description
2 Add WWN Blocks
+ — Advanced Filter Export Print
OK
Cancel
< Prev
Next >
Finish
Cancel

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

Create WWN Block

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

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

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

8. Enter number of WWNNs. For this study we had 32 WWNNs.

9. Click Finish.

Create UUID Suffix Pool

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

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

2. Select Pools > root.

3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.

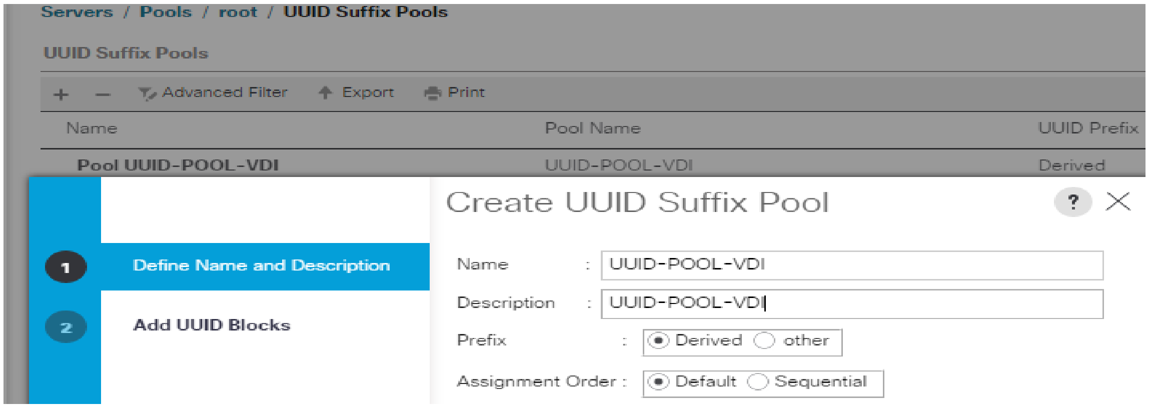
Servers / Pools / root / UUID Suffix Pools
UUID Suffix Pools
+ — Advanced Filter Export Print
1 Define Name and Description
2 Add UUID Blocks
?

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

Create UUID Suffix Pool

Name : UUID-POOL-VDI

Description : UUID-POOL-VDI

Prefix : ● Derived ○ other

Assignment Order : ● Default ○ Sequential

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

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

7. Keep the prefix at the derived option.

8. Click Next.

9. Click Add to add a block of UUIDs.

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

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

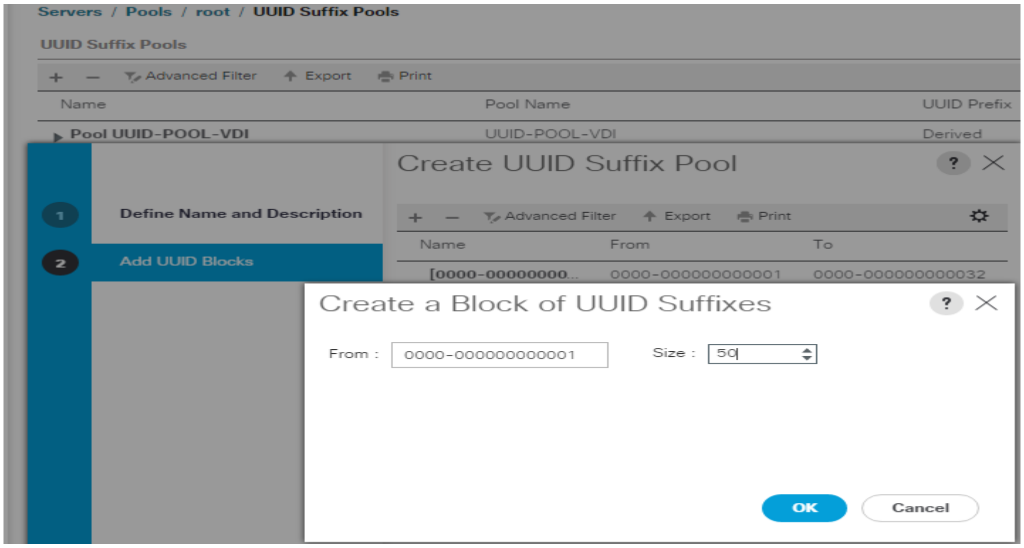
Servers / Pools / root / UUID Suffix Pools
UUID Suffix Pools
+ — Advanced Filter Export Print
1 Define Name and Description
2 Add UUID Blocks
+ — Advanced Filter Export Print
?
OK
Cancel

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

Create UUID Suffix Pool

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

Create a Block of UUID Suffixes

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 Horizon Linked Clone servers and Horizon RDSH 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, eight unique VLANs are created. Refer to Table 17 .

Table 17 VLANs Created

VLAN Name	VLAN ID	VLAN Purpose	vNIC Assignment
Default	1	Native VLAN	vNIC-Template-A vNIC-Template-B
In-Band-Mgmt	60	VLAN for in-band management interfaces	vNIC-Template-A vNIC-Template-B
Infra-Mgmt	61	VLAN for Virtual Infrastructure	vNIC-Template-A vNIC-Template-B
NFS-Vlan	62	VLAN for NFS Share	vNIC-Template-A vNIC-Template-B
CIFS-Vlan	63	VLAN-CIFS Share User Profiles	vNIC-Template-A vNIC-Template-B
vMotion	66	VLAN for VMware vMotion	vNIC-Template-A vNIC-Template-B
VDI	102	Virtual Desktop traffic	vNIC-Template-A vNIC-Template-B
OB-Mgmt	164	VLAN for out-of-band management interfaces	vNIC-Template-A vNIC-Template-B

2. Select LAN > LAN Cloud.

3. Right-click VLANs.

4. Select Create VLANs

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

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

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

8. Keep the Sharing Type as None.

9. Click OK and then click OK again.

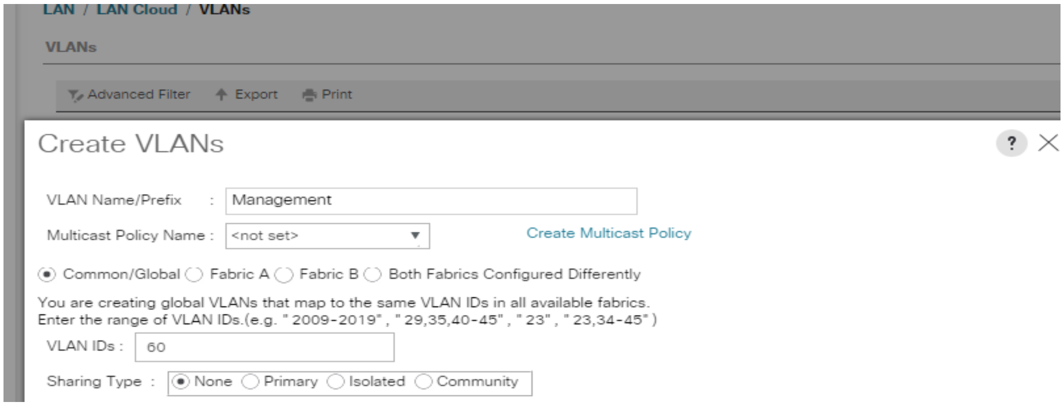

Create VLANs

VLAN Name/Prefix : Management

Multicast Policy Name : <not set>

Create Multicast Policy

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

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

VLAN IDs : 60

Sharing Type : ● None ○ Primary ○ Isolated ○ Community

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

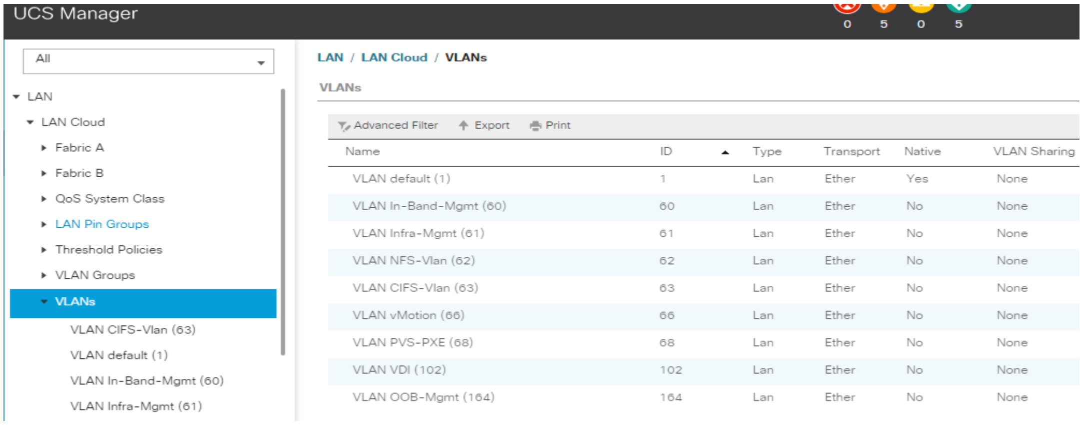

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN In-Band-Mgmt (60)	60	Lan	Ether	No	None
VLAN Infra-Mgmt (61)	61	Lan	Ether	No	None
VLAN NFS-Vlan (62)	62	Lan	Ether	No	None
VLAN CIFS-Vlan (63)	63	Lan	Ether	No	None
VLAN vMotion (66)	66	Lan	Ether	No	None
VLAN PVS-PXE (68)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

Create VSANs

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

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

In this procedure, two VSANs are created. When these VSANs are created, be sure to add them to the port-channel uplink created earlier.

2. Select SAN > SAN Cloud.

3. Under Fabric A, right-click VSANs.

4. Select Create VSANs.

5. Enter VSAN-400-A as the name of the VSAN to be used for in-band management traffic.

6. Select Fabric A for the scope of the VSAN.

7. Enter 400 as the ID of the VSAN.

8. Click OK, and then click OK again.

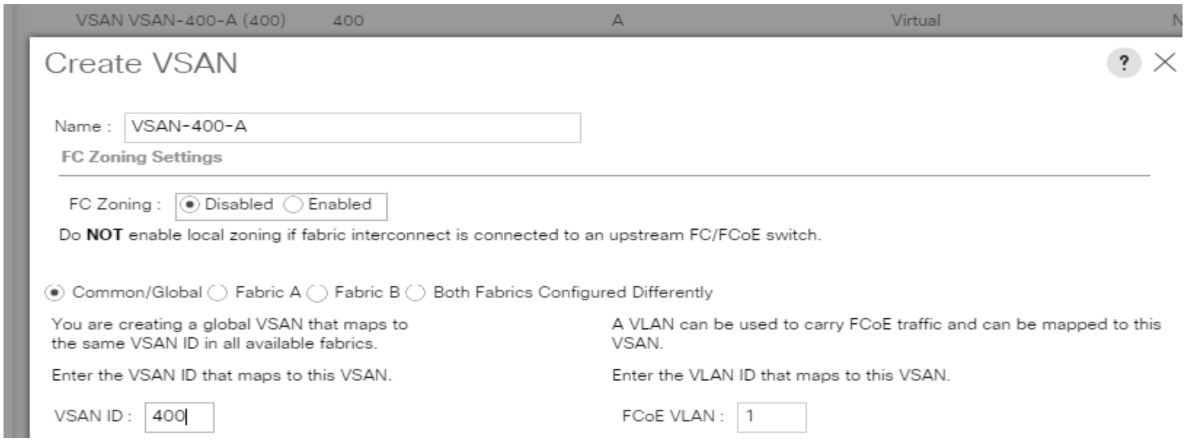

VSAN VSAN-400-A (400) 400 A Virtual N

Create VSAN

Name : VSAN-400-A

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

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

You are creating a global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 400

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 1

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

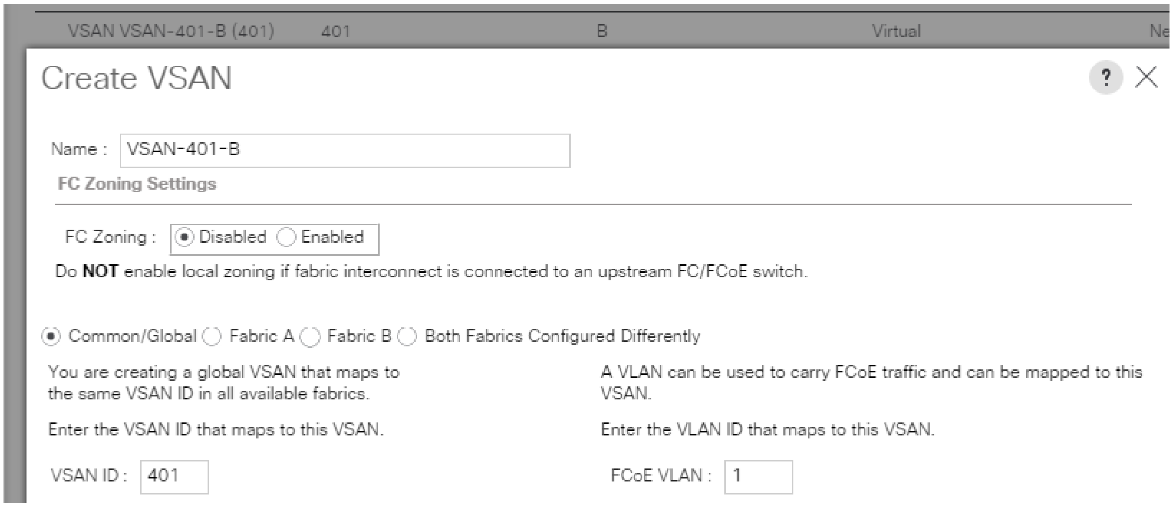

VSAN VSAN-401-B (401) 401 B Virtual Ne

Create VSAN

Name : VSAN-401-B

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

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

You are creating a global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID : 401

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN : 1

VSAN 400 and 401 are configured as shown below:

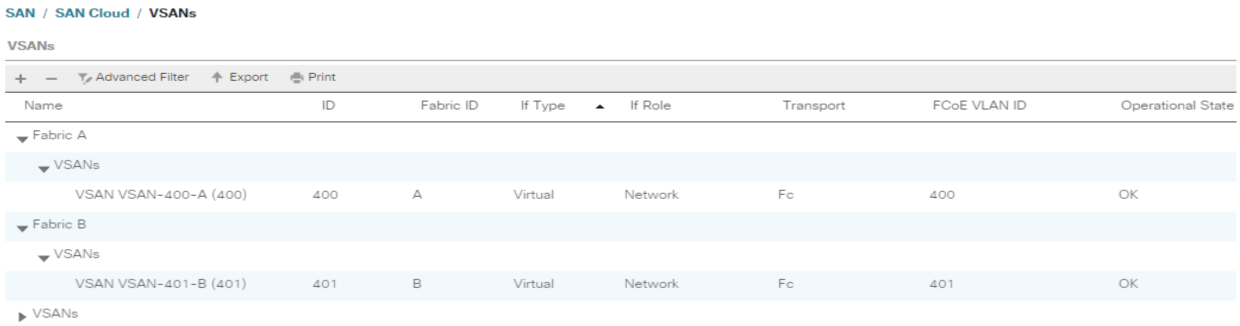

SAN / SAN Cloud / VSANs

VSANs

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

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

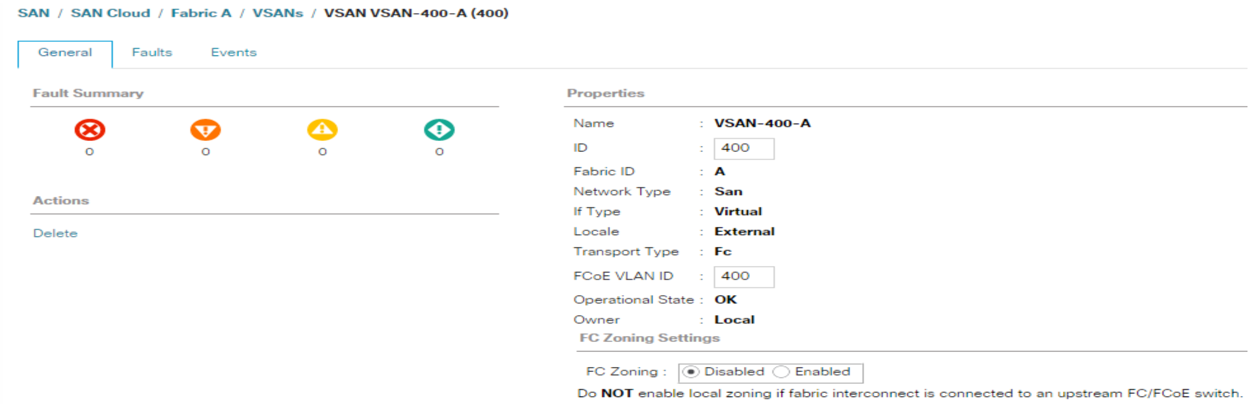

SAN / SAN Cloud / Fabric A / VSANs / VSAN VSAN-400-A (400)

Properties

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

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

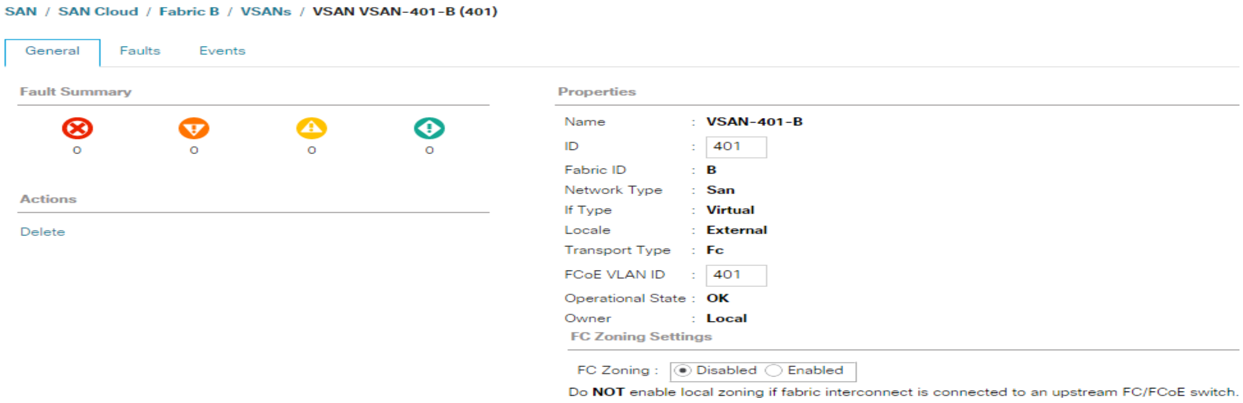

SAN / SAN Cloud / Fabric B / VSANs / VSAN VSAN-401-B (401)

Properties

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

FC Zoning Settings

FC Zoning : ● Disabled ○ Enabled

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

11. Go to the Port-Channel for each Fabric and assign the VSAN appropriately.

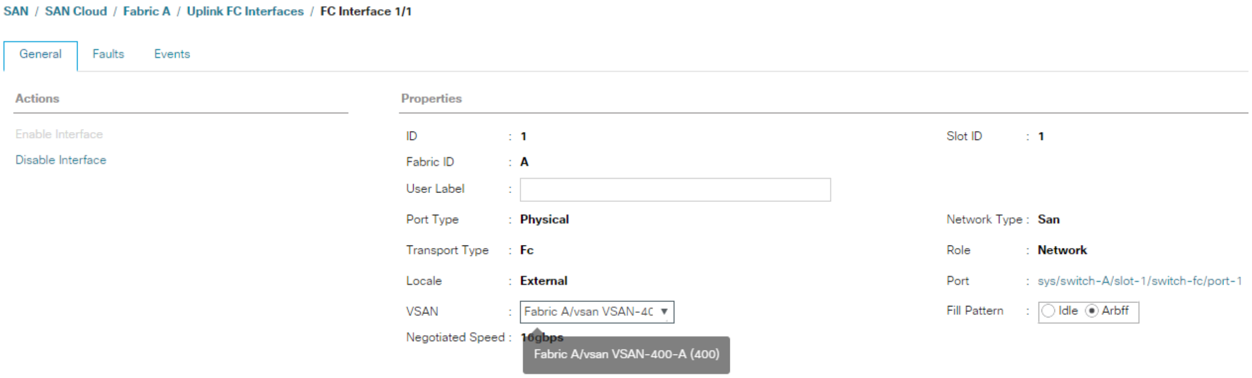

SAN / SAN Cloud / Fabric A / Uplink FC Interfaces / FC Interface 1/1

ID : 1
Fabric ID : A
User Label :
Port Type : Physical
Transport Type : Fc
Locale : External
VSAN : Fabric A/vsan VSAN-4C ▼
Negotiated Speed : 16gbps
Fabric A/vsan VSAN-400-A (400)
Slot ID : 1
Network Type : San
Role : Network
Port : sys/switch-A/slot-1/switch-fc/port-1
Fill Pattern : ○ Idle ● Arbff

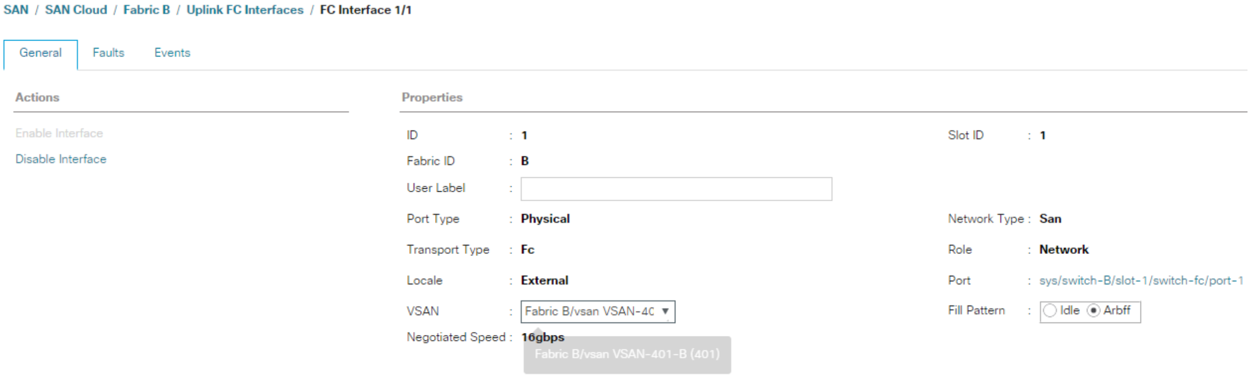

SAN / SAN Cloud / Fabric B / Uplink FC Interfaces / FC Interface 1/1

ID : 1
Fabric ID : B
User Label :
Port Type : Physical
Transport Type : Fc
Locale : External
VSAN : Fabric B/vsan VSAN-4C ▼
Negotiated Speed : 16gbps
Fabric B/vsan VSAN-401-B (401)
Slot ID : 1
Network Type : San
Role : Network
Port : sys/switch-B/slot-1/switch-fc/port-1
Fill Pattern : ○ Idle ● Arbff

Create Host Firmware Package

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

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

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

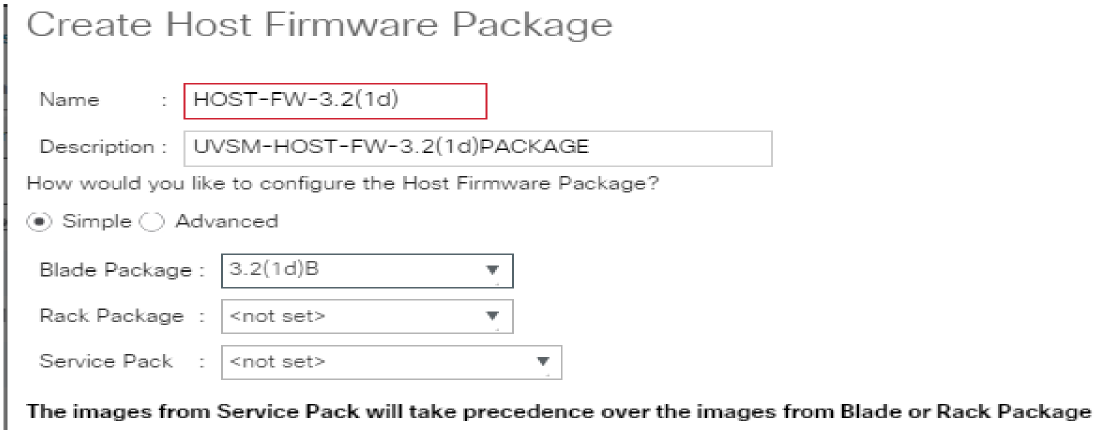

Create Host Firmware Package

Name : HOST-FW-3.2(1d)

Description : UVSM-HOST-FW-3.2(1d)PACKAGE

How would you like to configure the Host Firmware Package?

● Simple ○ Advanced

Blade Package : 3.2(1d)B

Rack Package : <not set>

Service Pack : <not set>

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

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.

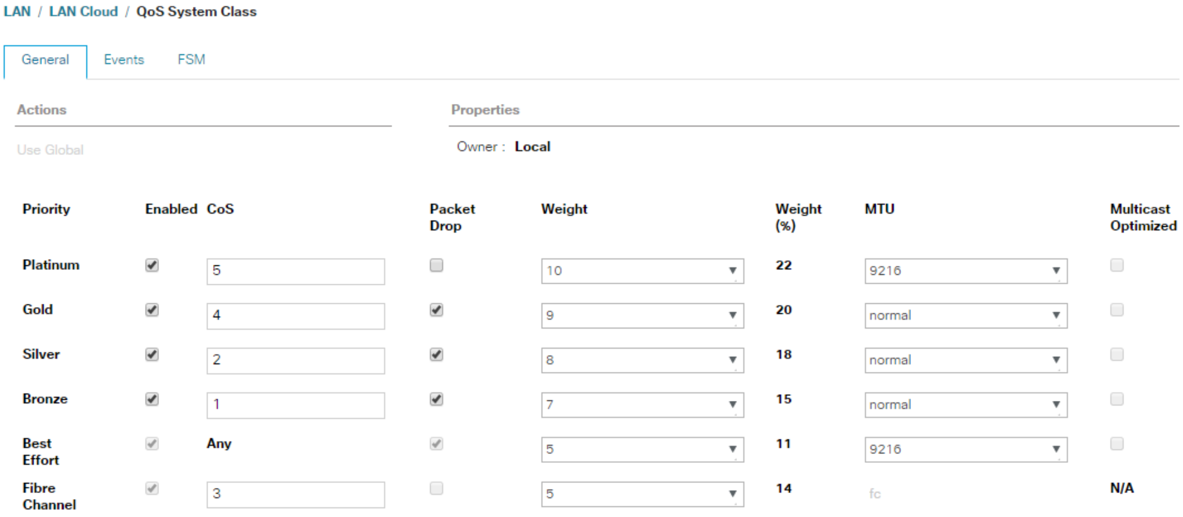

LAN / LAN Cloud / QoS System Class

Owner : Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☑	5	□	10	22	9216	□
Gold	☑	4	☑	9	20	normal	□
Silver	☑	2	☑	8	18	normal	□
Bronze	☑	1	☑	7	15	normal	□
Best Effort	☑	Any	☑	5	11	9216	□
Fibre Channel	☑	3	□	5	14	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.

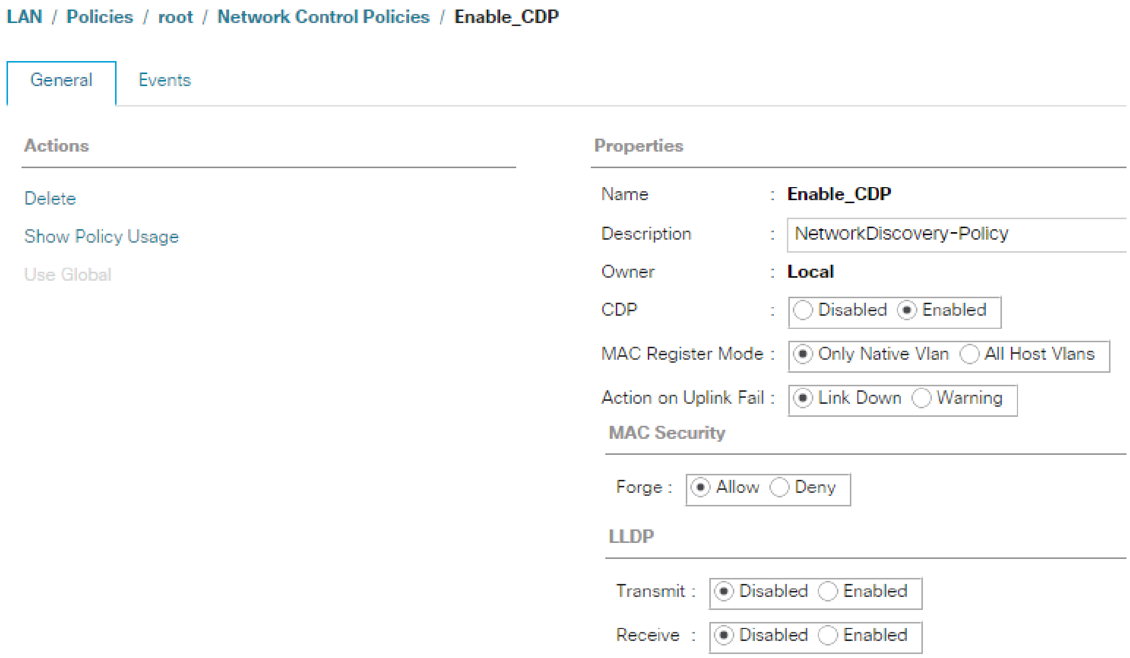

LAN / Policies / root / Network Control Policies / Enable_CDP

Name : Enable_CDP
Description : NetworkDiscovery-Policy
Owner : Local
CDP : ○ Disabled ● Enabled
MAC Register Mode : ● Only Native Vlan ○ All Host Vlans
Action on Uplink Fail : ● Link Down ○ Warning

MAC Security

Forge : ● Allow ○ Deny

LLDP

Transmit : ● Disabled ○ Enabled
Receive : ● Disabled ○ Enabled

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.

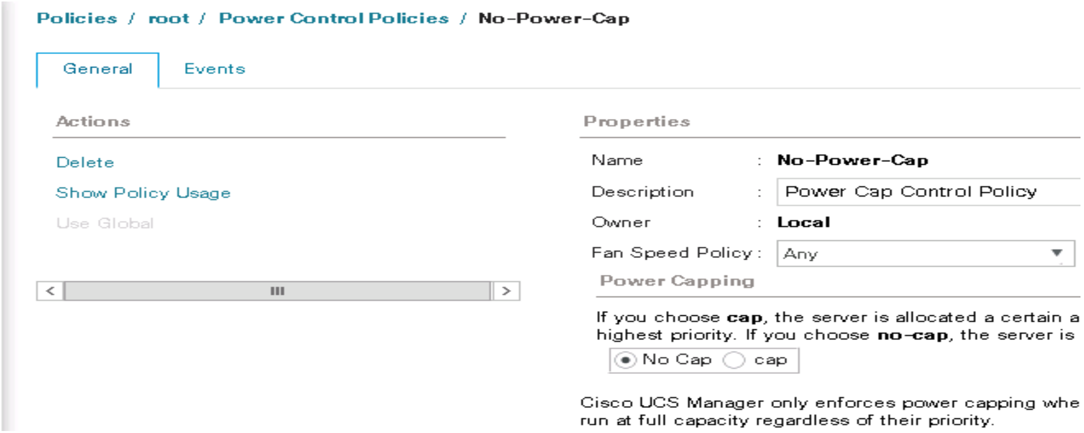

Policies / root / Power Control Policies / No-Power-Cap

Name : **No-Power-Cap**
Description : Power Cap Control Policy
Owner : **Local**
Fan Speed Policy : Any

Power Capping

If you choose **cap**, the server is allocated a certain a
highest priority. If you choose **no-cap**, the server is

● No Cap ○ cap

Cisco UCS Manager only enforces power capping whe
run at full capacity regardless of their priority.

Cisco UCS System Configuration for Cisco UCS B-Series

Create Server BIOS Policy

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

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

2. Select Policies > root.

3. Right-click BIOS Policies.

4. Select Create BIOS Policy.

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

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

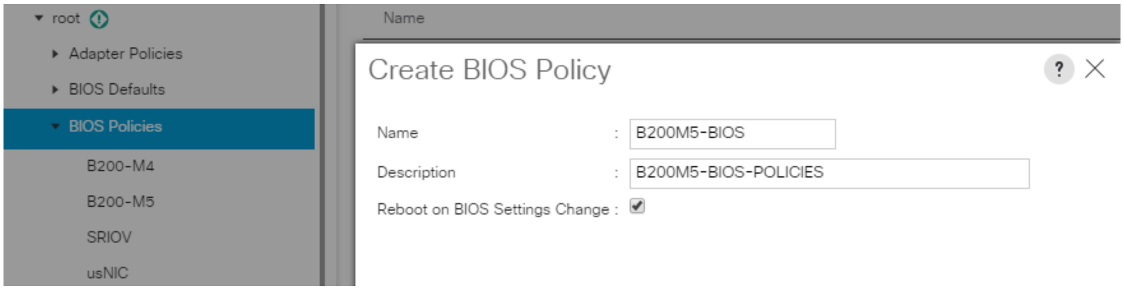

Create BIOS Policy

Name : B200M5-BIOS
Description : B200M5-BIOS-POLICIES
Reboot on BIOS Settings Change : ☑

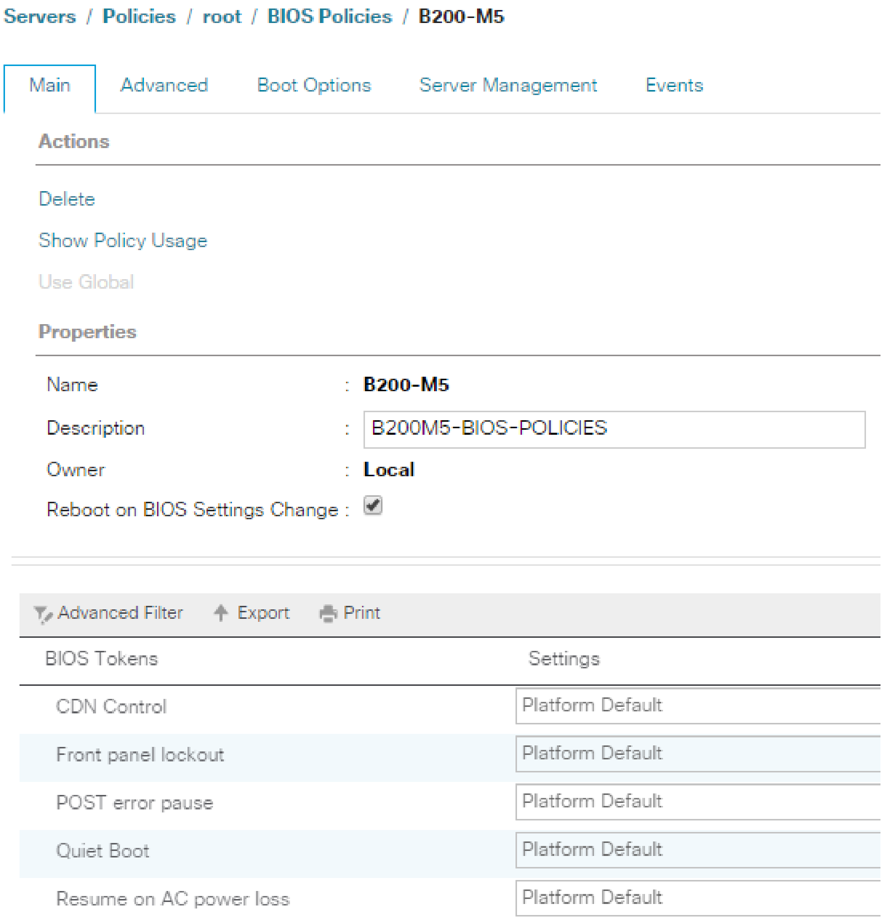

Properties

Name : **B200-M5**

Description : B200M5-BIOS-POLICIES

Owner : **Local**

Reboot on BIOS Settings Change : ☑

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

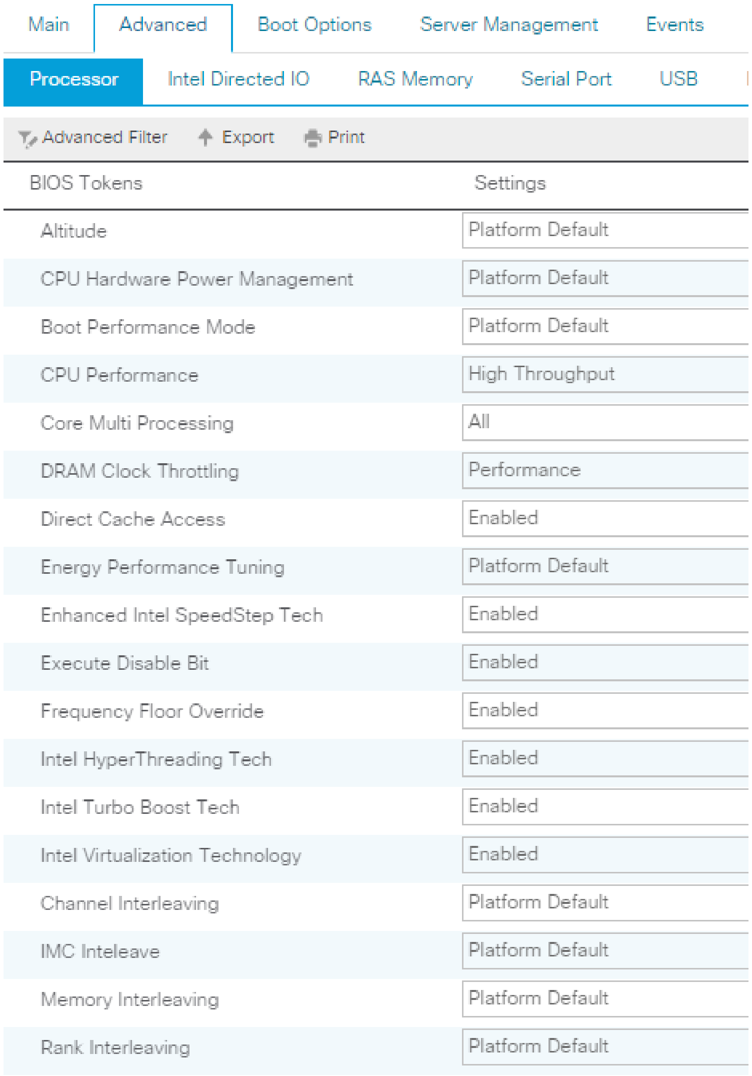

BIOS Tokens	Settings
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Platform Default
CPU Performance	High Throughput
Core Multi Processing	All
DRAM Clock Throttling	Performance
Direct Cache Access	Enabled
Energy Performance Tuning	Platform Default
Enhanced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

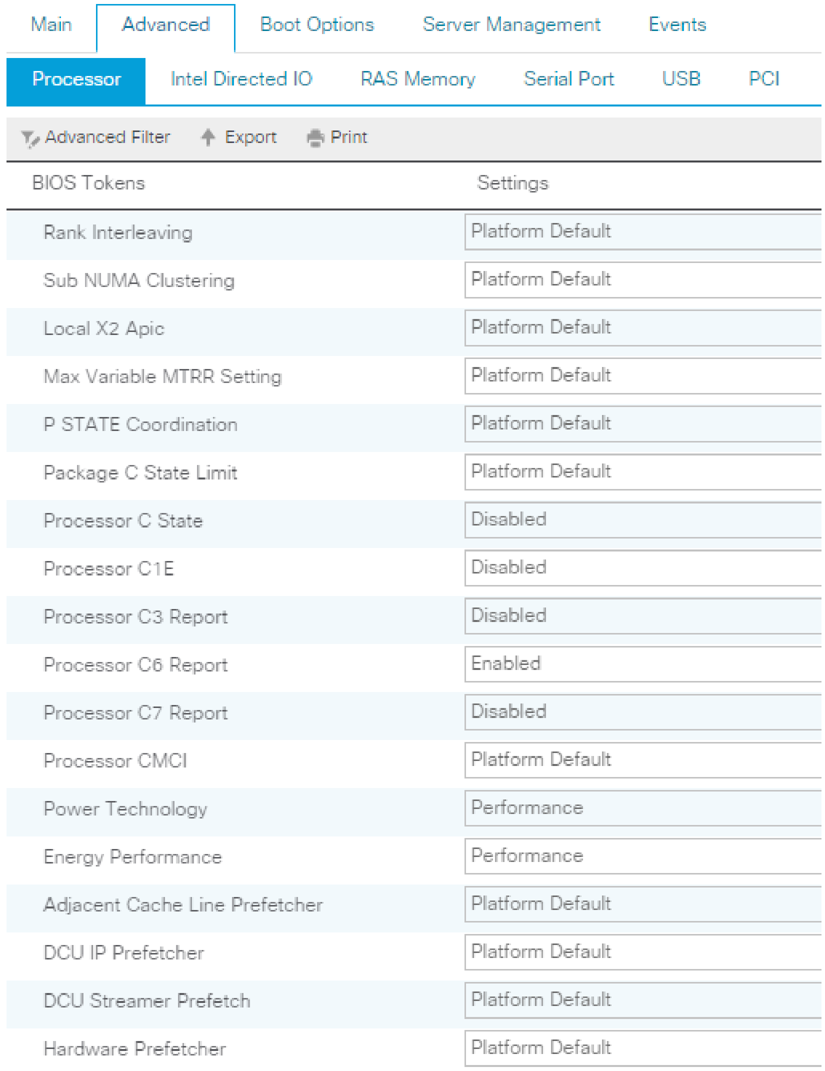

BIOS Tokens	Settings
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor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

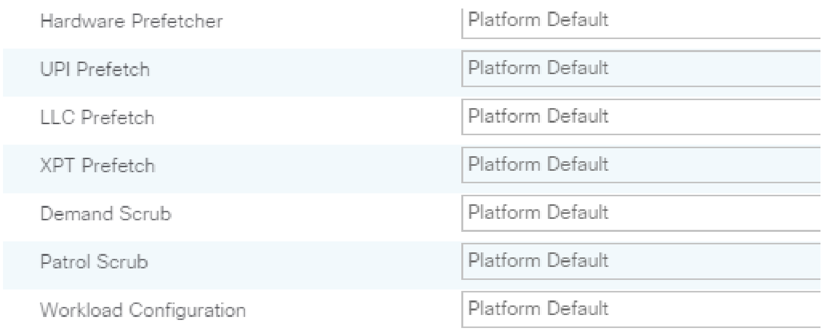

Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	Platform Default

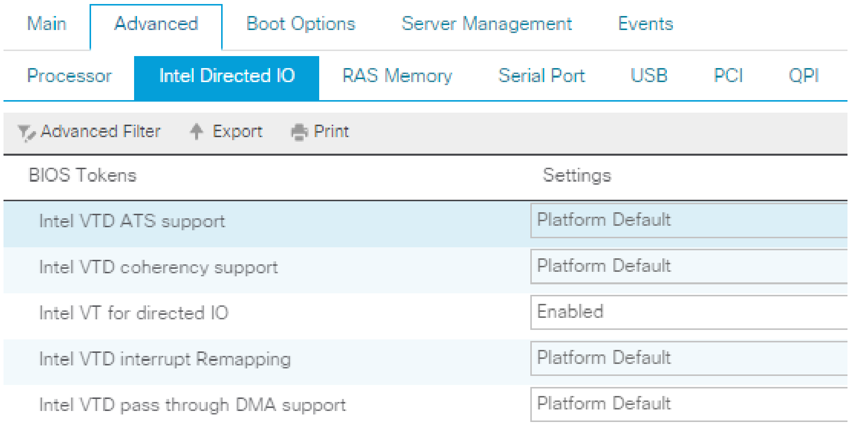

BIOS Tokens	Settings
Intel VTD ATS support	Platform Default
Intel VTD coherency support	Platform Default
Intel VT for directed IO	Enabled
Intel VTD interrupt Remapping	Platform Default
Intel VTD pass through DMA support	Platform Default

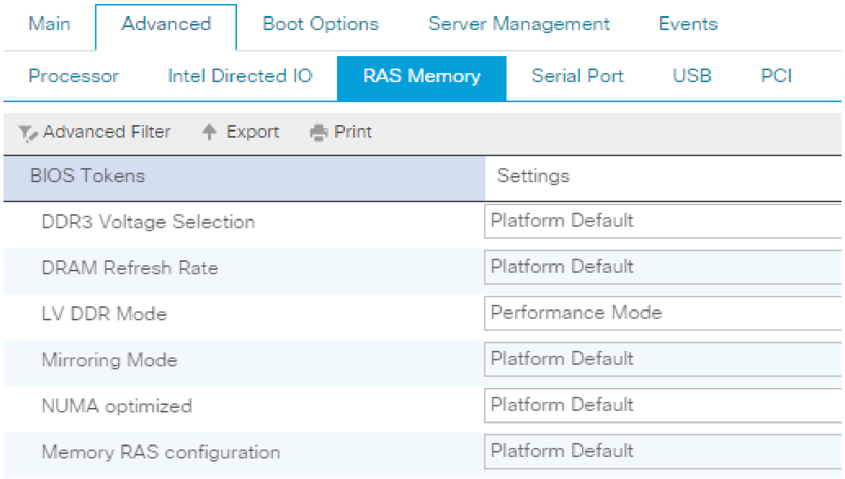

BIOS Tokens	Settings
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

7. Click Finish.

Configure Update Default Maintenance Policy

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

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

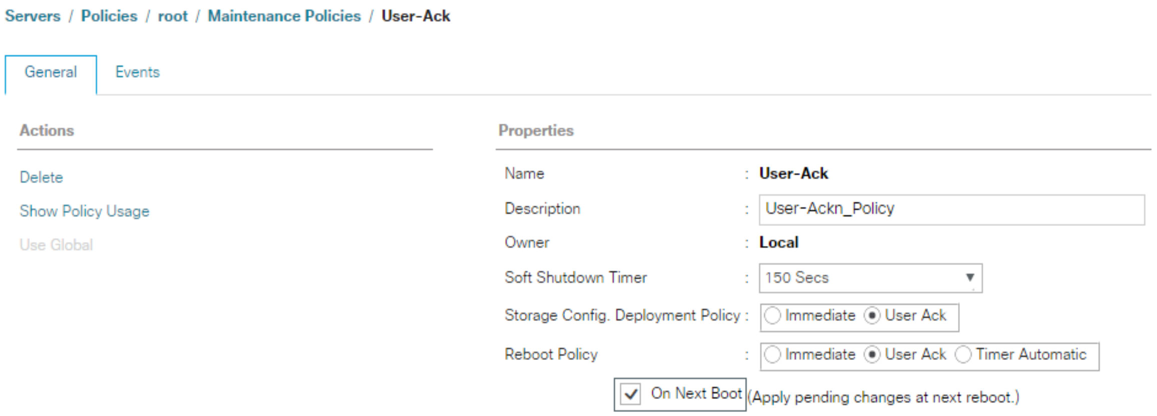

Servers / Policies / root / Maintenance Policies / User-Ack

Name : User-Ack

Description : User-Ackn_Policy

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : ○ Immediate ● User Ack

Reboot Policy : ○ Immediate ● User Ack ○ Timer Automatic

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

Create vNIC Templates for Cisco UCS B-Series

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

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

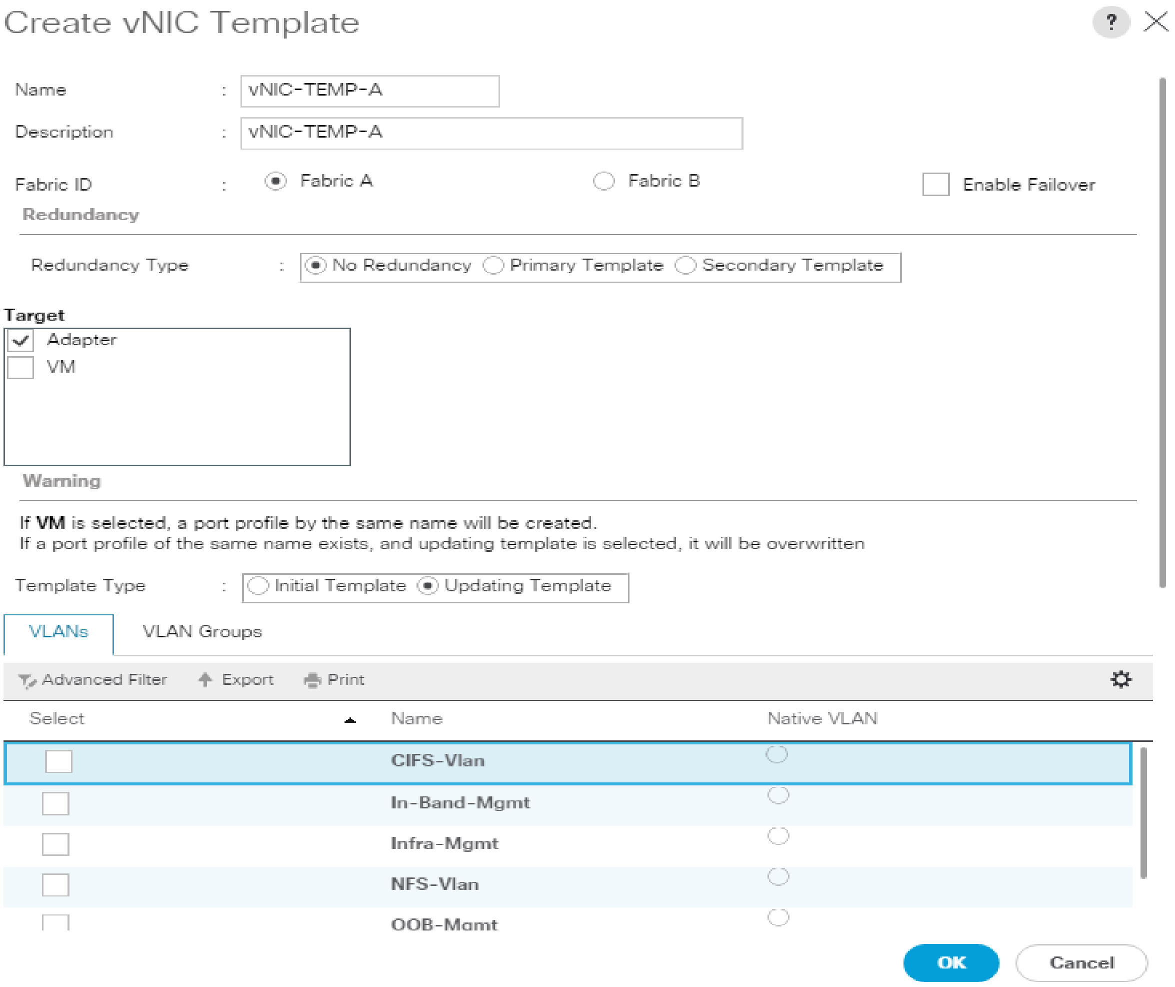

Create vNIC Template

Name : vNIC-TEMP-A

Description : vNIC-TEMP-A

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

Redundancy

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

Target

☑ Adapter
□ VM

Warning

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

Template Type : ○ Initial Template ● Updating Template

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

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

18. Select Policies > root.

19. Right-click vNIC Templates.

20. Select Create vNIC Template.

21. Enter vNIC_Template_B as the vNIC template name.

22. Select Fabric B.

23. Do not select the Enable Failover checkbox.

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

25. Select Updating Template as the template type.

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

27. Set Native-VLAN as the native VLAN.

28. For MTU, enter 9000.

29. In the MAC Pool list, select MAC_Pool_B.

30. In the Network Control Policy list, select CDP_Enabled.

31. Click OK to create the vNIC template.

32. Click OK.

Create vHBA Templates for Cisco UCS B-Series

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

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

2. Select Policies > root.

3. Right-click vHBA Templates.

4. Select Create vHBA Template.

5. Enter vHBA-FAB-A as the vHBA template name.

6. Keep Fabric A selected.

7. Select VSAN-400-A for Fabric A from the drop down.

8. Change to Updating Template.

9. For Max Data Field keep 2048.

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

11. Leave the remaining as is.

12. Click OK.

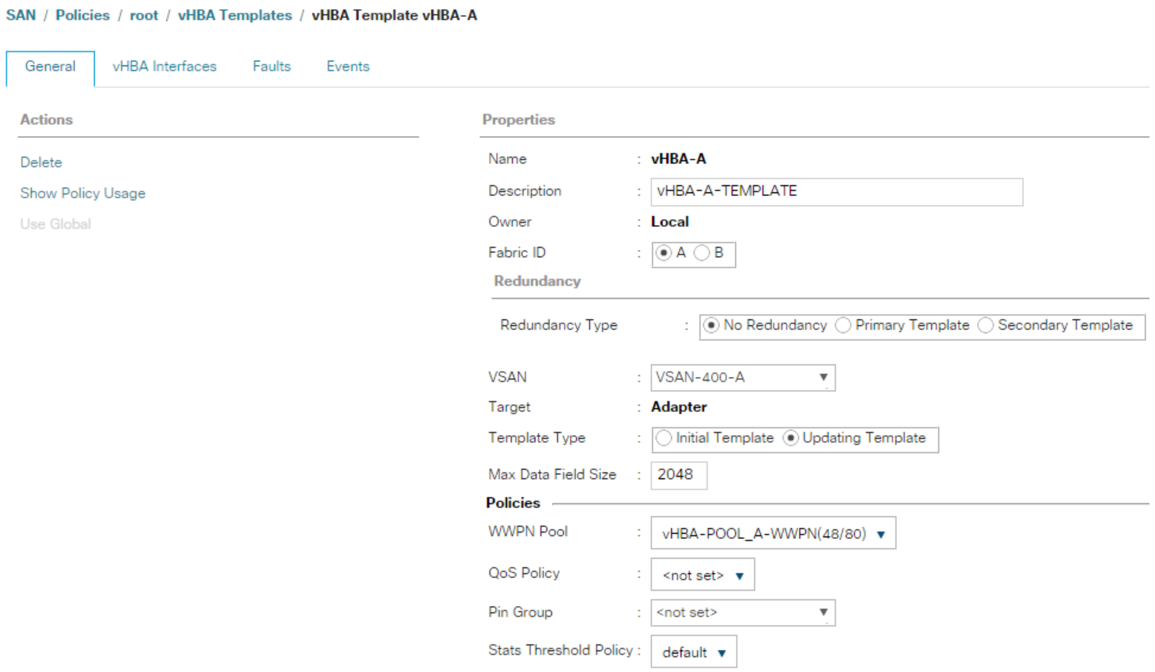

SAN / Policies / root / vHBA Templates / vHBA Template vHBA-A

General | vHBA Interfaces | Faults | Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name	: vHBA-A
Description	: vHBA-A-TEMPLATE
Owner	: Local
Fabric ID	: ● A ○ B

Redundancy

Redundancy Type	: ● No Redundancy ○ Primary Template ○ Secondary Template
VSAN	: VSAN-400-A ▾
Target	: Adapter
Template Type	: ○ Initial Template ● Updating Template
Max Data Field Size	: 2048

Policies

WWPN Pool	: vHBA-POOL_A-WWPN(48/80) ▾
QoS Policy	: <not set> ▾
Pin Group	: <not set> ▾
Stats Threshold Policy	: default ▾

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

14. Select Policies > root.

15. Right-click vHBA Templates.

16. Select Create vHBA Template.

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

18. Select Fabric B.

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

20. Change to Updating Template.

21. For Max Data Field keep 2048.

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

23. Leave the remaining as is.

24. Click OK.

Configure Boot from SAN

All ESXi 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.'

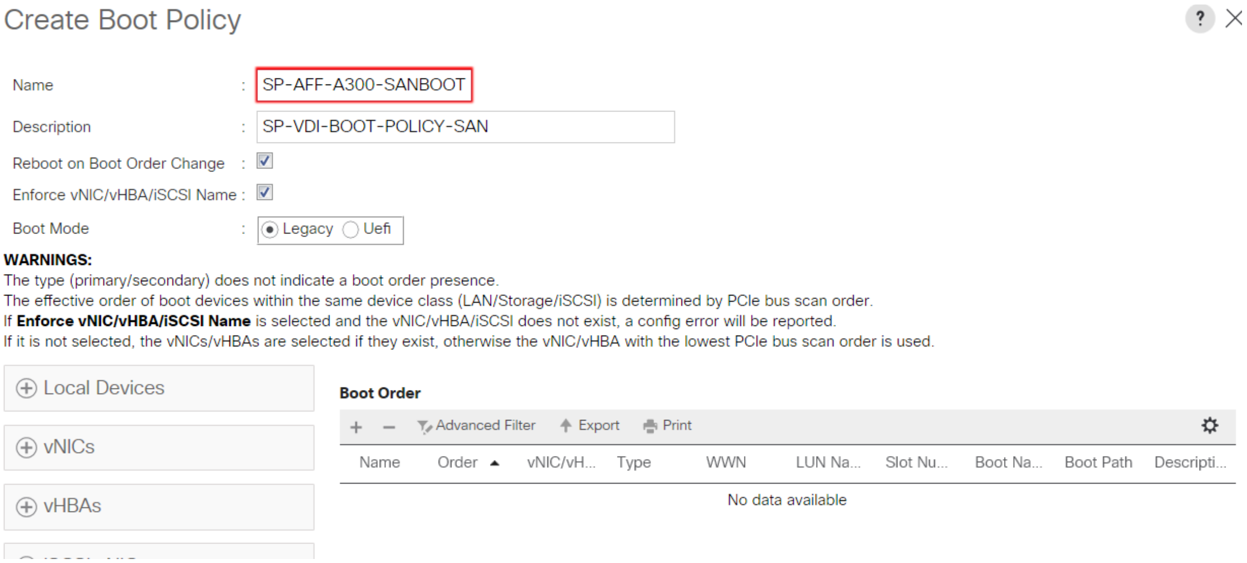

Create Boot Policy

Name : SP-AFF-A300-SANBOOT

Description : SP-VDI-BOOT-POLICY-SAN

Reboot on Boot Order Change : ☑

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

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

Boot Order

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

No data available

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

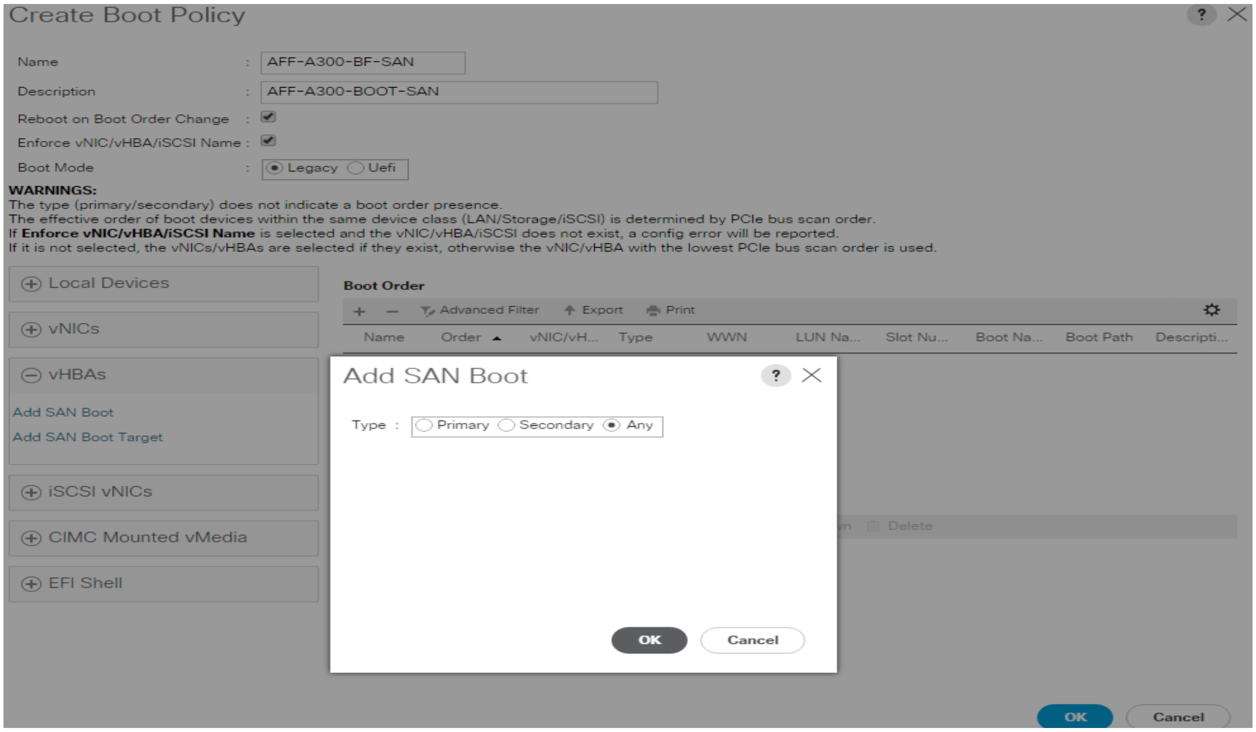

Create Boot Policy

Name : AFF-A300-BF-SAN

Description : AFF-A300-BOOT-SAN

Reboot on Boot Order Change : ☑

Enforce vNIC/vHBA/iSCSI Name : ☑

Boot Mode : ● Legacy ○ Uefi

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

Boot Order

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

Add SAN Boot

Type : ○ Primary ○ Secondary ● Any

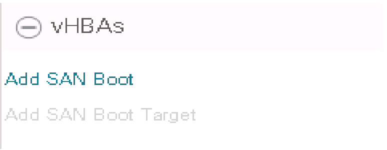

3. After selecting the ‘Add SAN Boot’ option, add the primary vHBA as shown below. Note that the vHBA name needs to match exactly. We will use the vHBA templates created in the previous step.

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

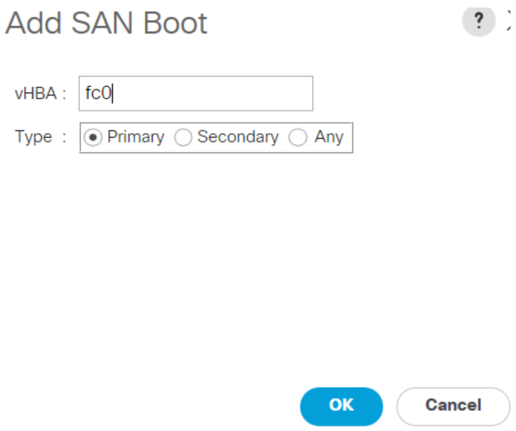

Add SAN Boot

vHBA : fc0

Type : ● Primary ○ Secondary ○ Any

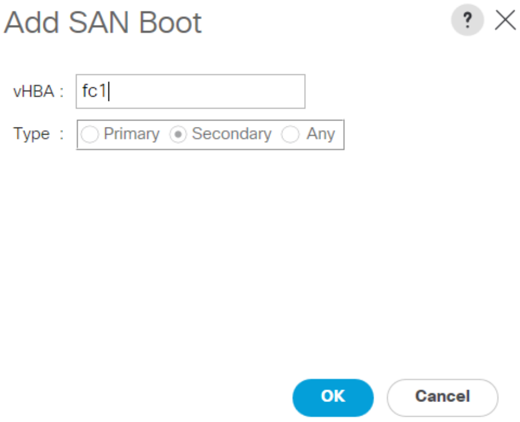

Add SAN Boot

vHBA : fc1

Type : ○ Primary ● Secondary ○ Any

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

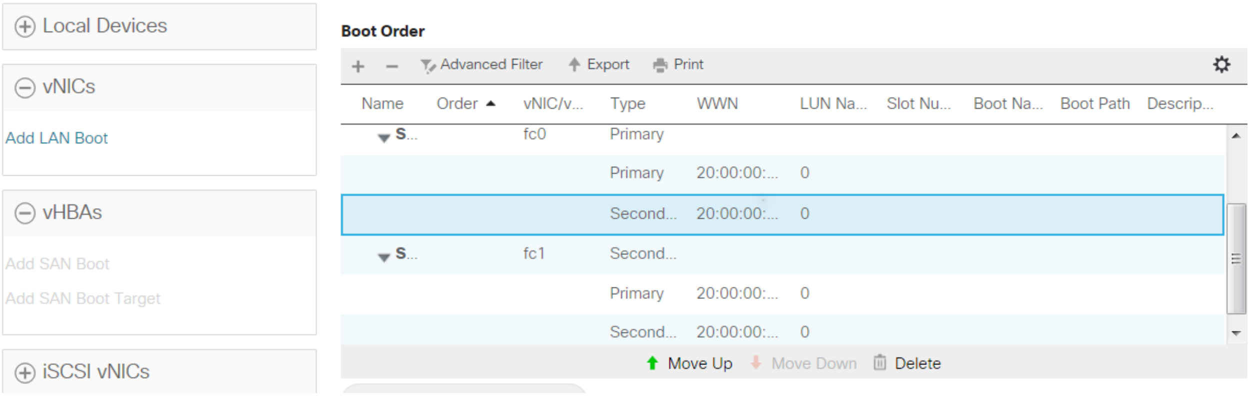

Boot Order

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

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

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

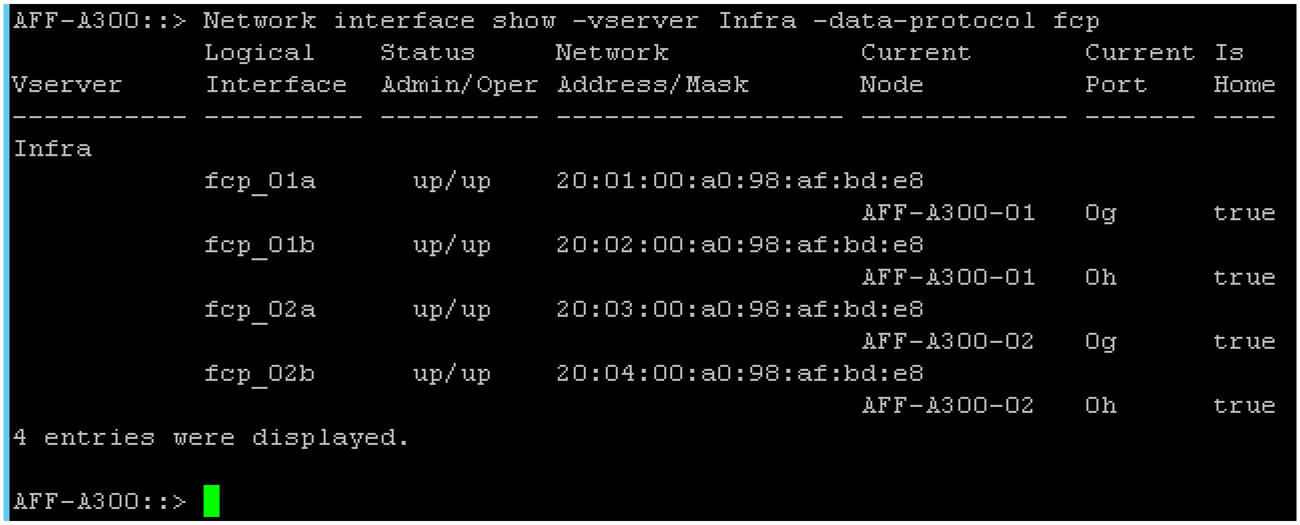

```
AFF-A300::> Network interface show -vserver Infra -data-protocol fcp
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- -------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

AFF-A300::>
```

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

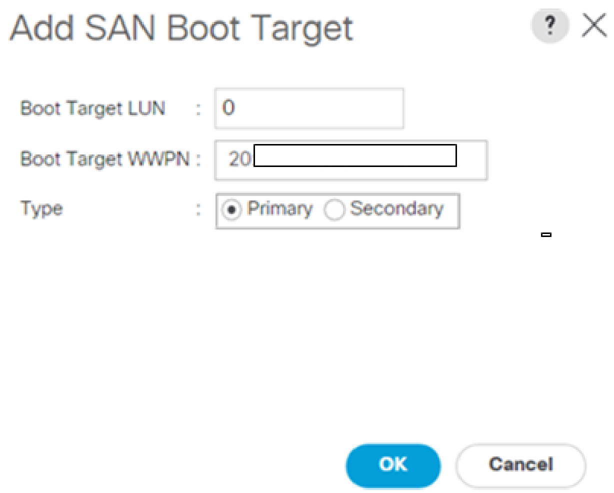

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20

Type : ● Primary ○ Secondary

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

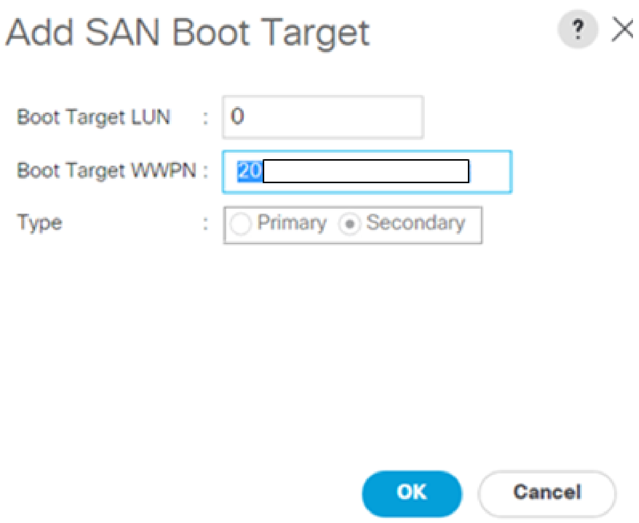

Add SAN Boot Target

Boot Target LUN : 0

Boot Target WWPN : 20[redacted]

Type : ○ Primary ● Secondary

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

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

Create Service Profile Templates for Cisco UCS B-Series

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

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

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

3. Name the template B-Series.

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

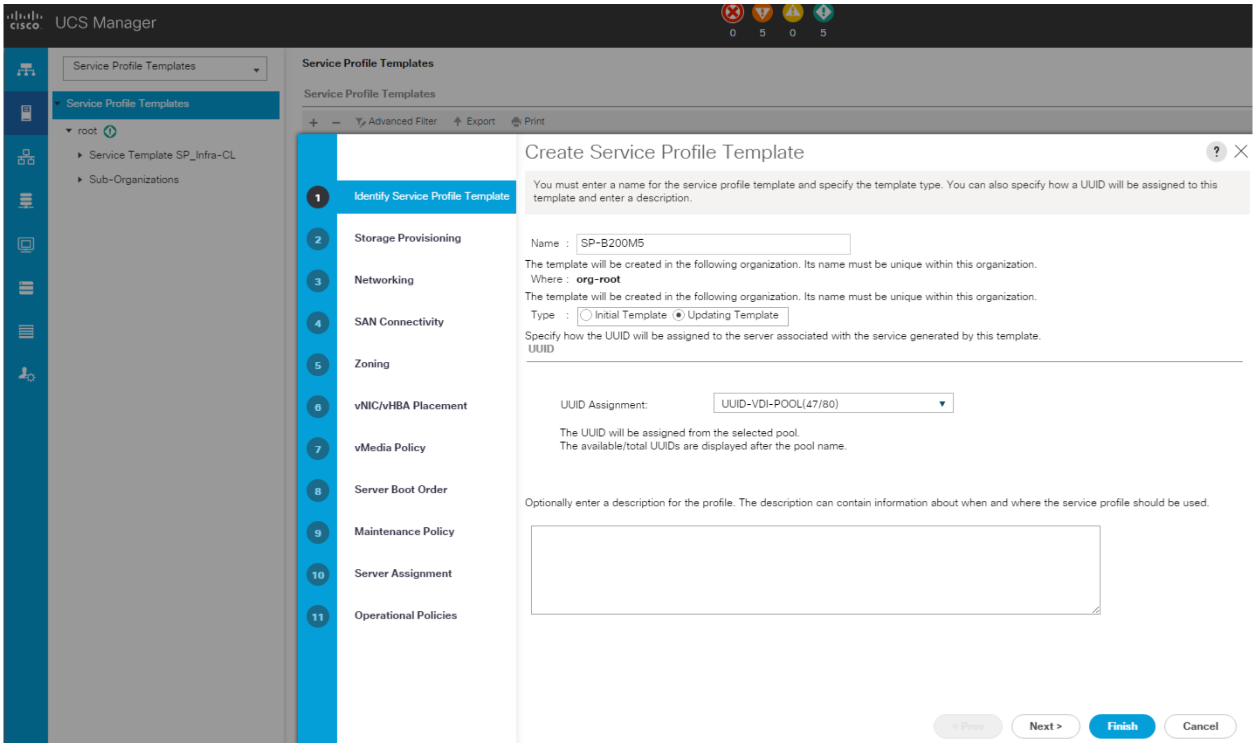

Create Service Profile Template

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

Name : SP-B200M5

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

Where : **org-root**

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

Type : ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: UUID-VDI-POOL(47/80)

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

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

5. Click Next.

6. Click Next through Storage Provisioning.

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

8. Click Add.

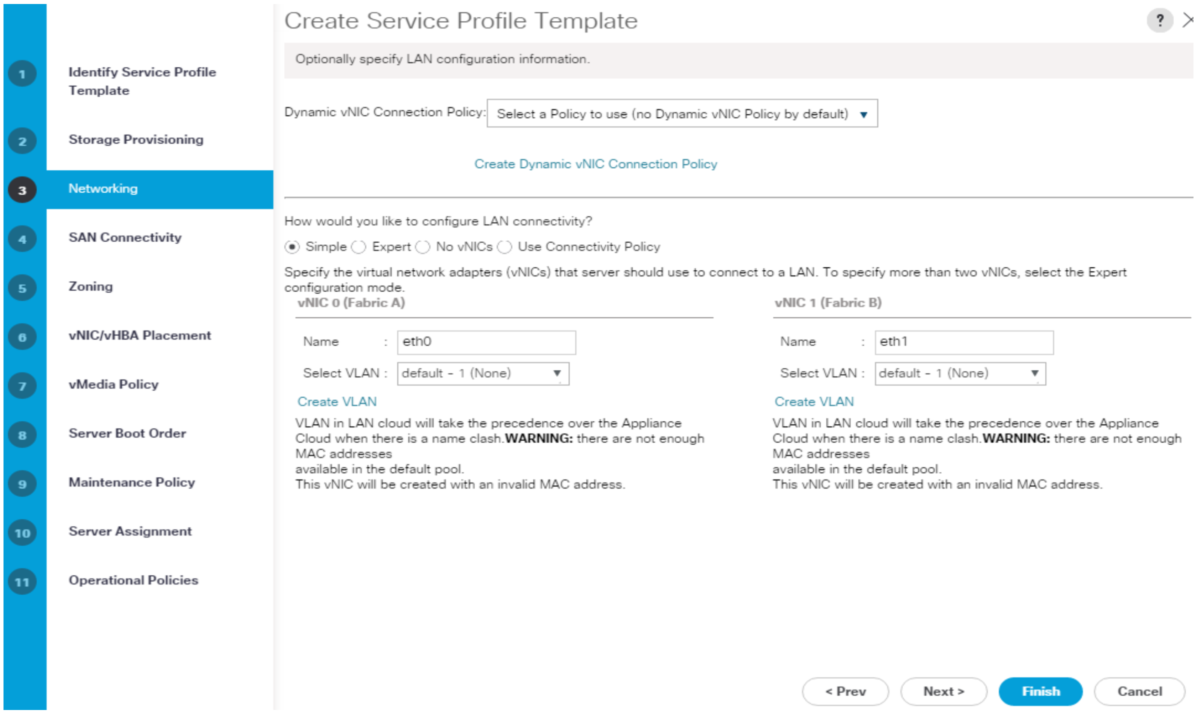

Create Service Profile Template

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

Optionally specify LAN configuration information.

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

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

● Simple ○ Expert ○ No vNICs ○ Use Connectivity Policy

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

vNIC 0 (Fabric A)

Name : eth0

Select VLAN : default - 1 (None) ▼

Create VLAN

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

vNIC 1 (Fabric B)

Name : eth1

Select VLAN : default - 1 (None) ▼

Create VLAN

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

9. Name it vNIC-A.

10. Select check box for Use vNIC Template.

11. Under vNIC template select the vNIC-A.

12. For Adapter Policy select VMware.

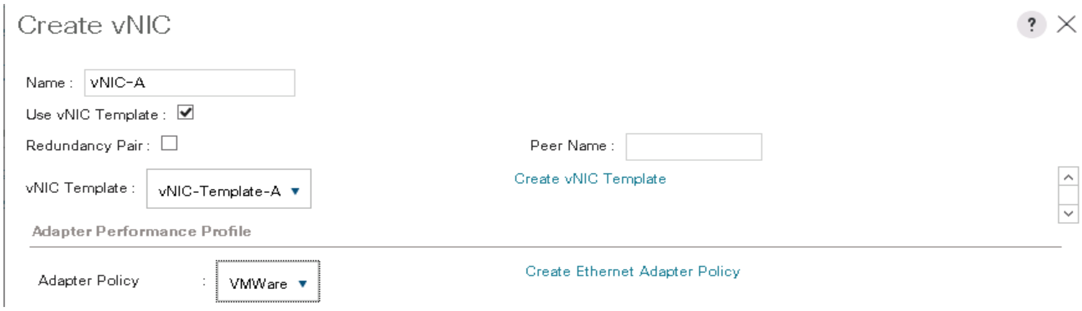

Create vNIC

Name : vNIC-A

Use vNIC Template : ☑

Redundancy Pair : ☐

Peer Name :

vNIC Template : vNIC-Template-A ▼

Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Ethernet Adapter Policy

13. Repeat networking steps for vNIC-B.

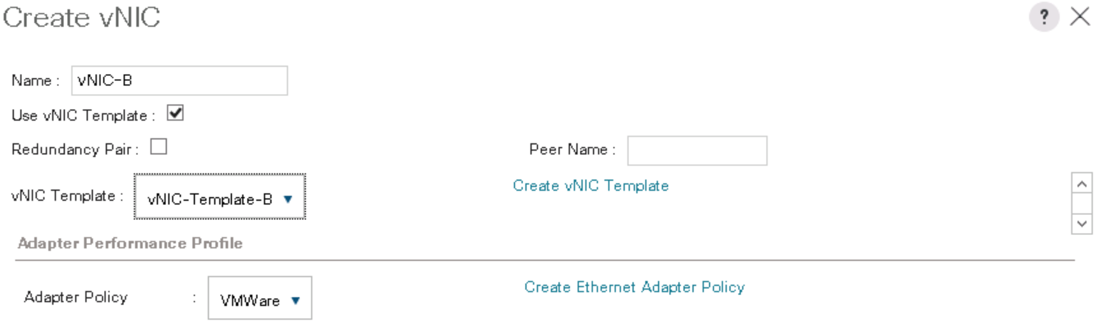

Create vNIC

Name : vNIC-B

Use vNIC Template : ☑

Redundancy Pair : ☐

Peer Name :

vNIC Template : vNIC-Template-B ▼

Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Ethernet Adapter Policy

14. Click Next.

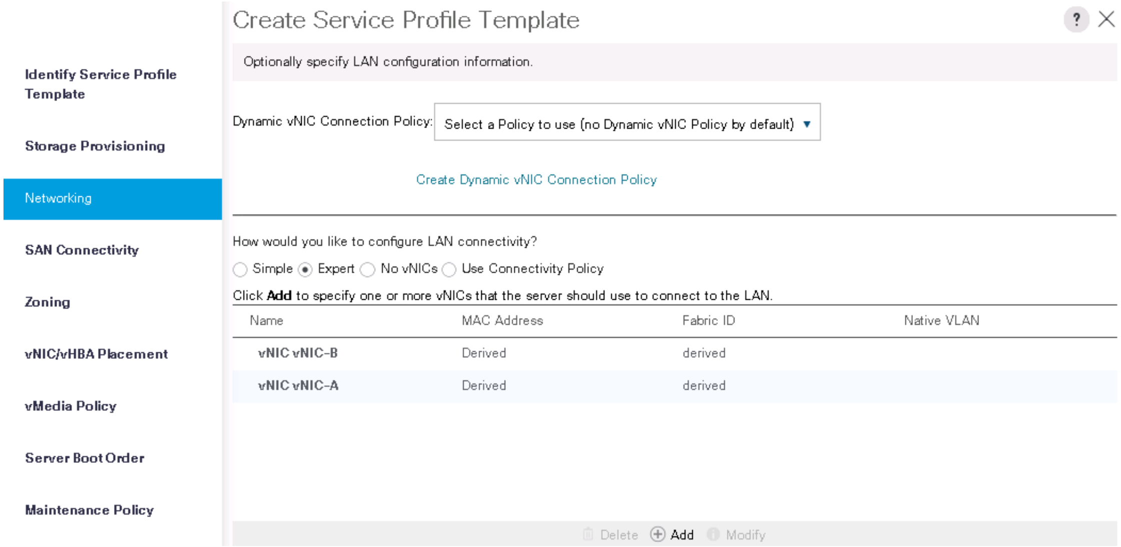

Identify Service Profile Template

Storage Provisioning

Networking

SAN Connectivity

Zoning

vNIC/vHBA Placement

vMedia Policy

Server Boot Order

Maintenance Policy

Create Service Profile Template

Optionally specify LAN configuration information.

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

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

○ Simple ● Expert ○ No vNICs ○ Use Connectivity Policy

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

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

15. Click Next.

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

17. Select WWNN Assignment from the Pool created earlier.

18. Click Add.

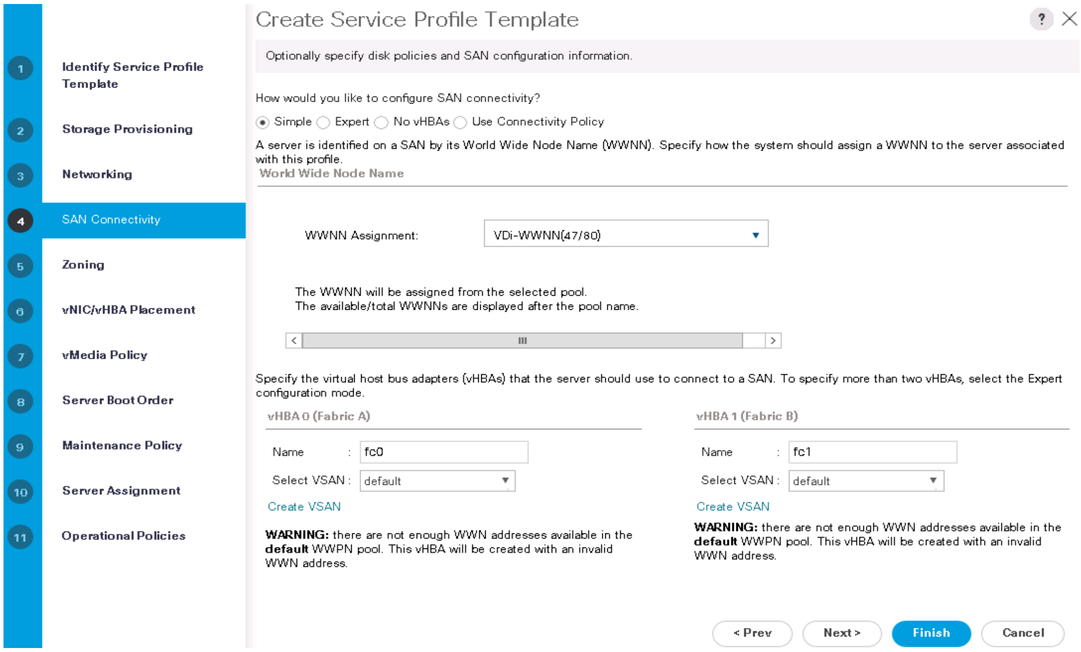

1 Identify Service Profile Template

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

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

● Simple ○ Expert ○ No vHBAs ○ Use Connectivity Policy

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

World Wide Node Name

WWNN Assignment: VDi-WWNN(47/80) ▾

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

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

vHBA 0 (Fabric A)

Name : fc0

Select VSAN : default ▾

Create VSAN

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

vHBA 1 (Fabric B)

Name : fc1

Select VSAN : default ▾

Create VSAN

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

19. Name the adapter vHBA-A.

20. Click Use vHBA Template.

21. Select vHBA Template: vHBA-A.

22. Select Adapter Policy: VMWare.

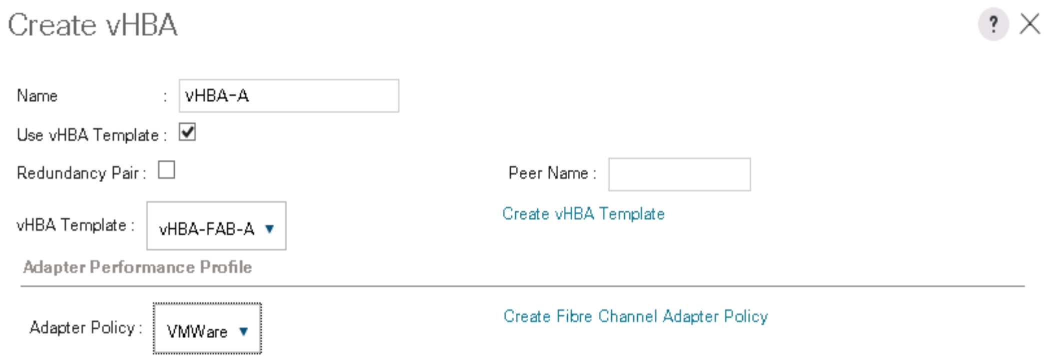

Create vHBA

Name : vHBA-A

Use vHBA Template : ☑

Redundancy Pair : □ Peer Name :

vHBA Template : vHBA-FAB-A ▾

Adapter Performance Profile

Adapter Policy : VMWare ▾

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

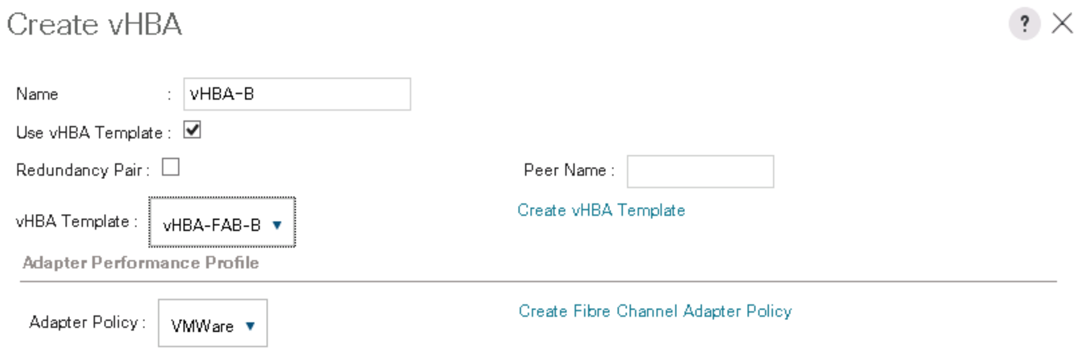

Create vHBA

Name : vHBA-B

Use vHBA Template : ☑

Redundancy Pair : □ Peer Name :

vHBA Template : vHBA-FAB-B ▾

Adapter Performance Profile

Adapter Policy : VMWare ▾

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 list to select the Boot Policy created earlier, then click Finish.

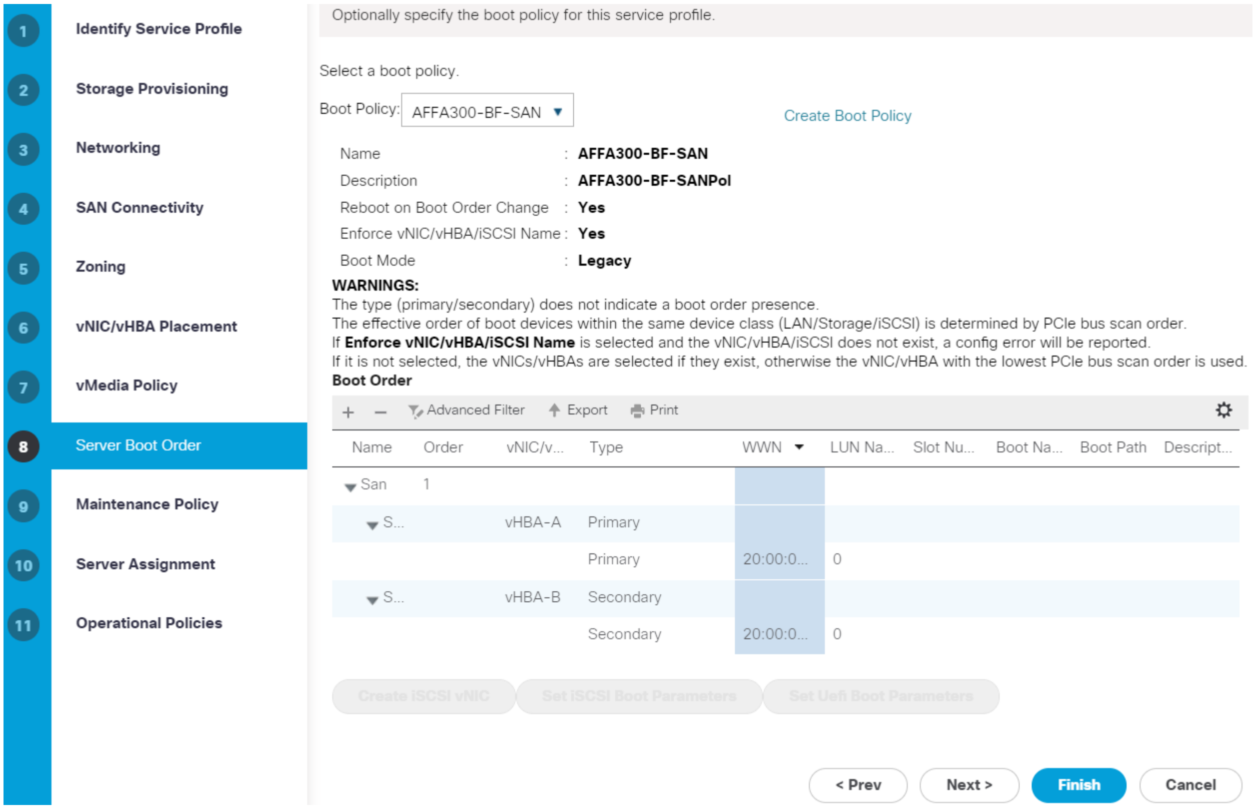

Optionally specify the boot policy for this service profile.

Select a boot policy.

Boot Policy: AFFA300-BF-SAN ▼

Name	:	**AFFA300-BF-SAN**
Description	:	**AFFA300-BF-SANPol**
Reboot on Boot Order Change	:	**Yes**
Enforce vNIC/vHBA/iSCSI Name	:	**Yes**
Boot Mode	:	**Legacy**

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

Boot Order

Name	Order	vNIC/v...	Type	WWN	LUN Na...	Slot Nu...	Boot Na...	Boot Path	Descript...
San	1								
S...		vHBA-A	Primary						
			Primary	20:00:0...	0				
S...		vHBA-B	Secondary						
			Secondary	20:00:0...	0				

28. Select maintenance Policy and Server Assignment.

29. Click Finish and complete the Service Profile creation.

Create Service Profiles

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

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

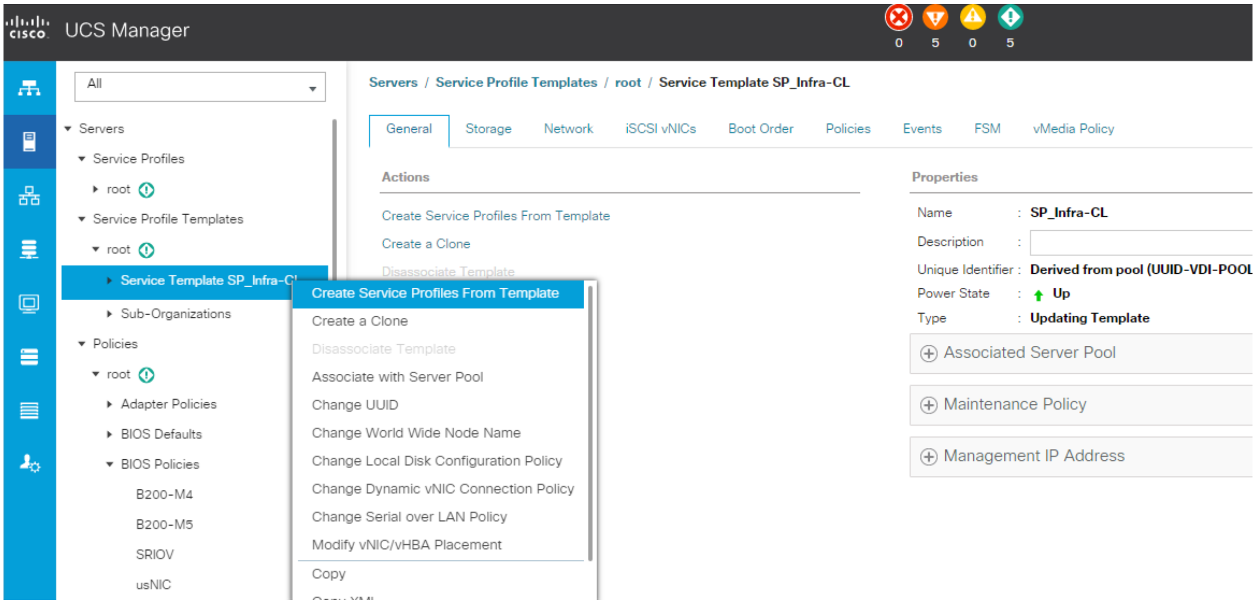
UCS Manager
0 5 0 5
All
Servers
Service Profiles
root
Service Profile Templates
root
Service Template SP_Infra-CL
Sub-Organizations
Policies
root
Adapter Policies
BIOS Defaults
BIOS Policies
B200-M4
B200-M5
SRIOV
usNIC
Servers / Service Profile Templates / root / Service Template SP_Infra-CL
General
Storage
Network
iSCSI vNICs
Boot Order
Policies
Events
FSM
vMedia Policy
Actions
Create Service Profiles From Template
Create a Clone
Disassociate Template
Create Service Profiles From Template
Create a Clone
Disassociate Template
Associate with Server Pool
Change UUID
Change World Wide Node Name
Change Local Disk Configuration Policy
Change Dynamic vNIC Connection Policy
Change Serial over LAN Policy
Modify vNIC/vHBA Placement
Copy
Properties
Associated Server Pool
Maintenance Policy
Management IP Address

Name : SP_Infra-CL

Description :

Unique Identifier : Derived from pool (UUID-VDI-POOL

Power State : Up

Type : Updating Template

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

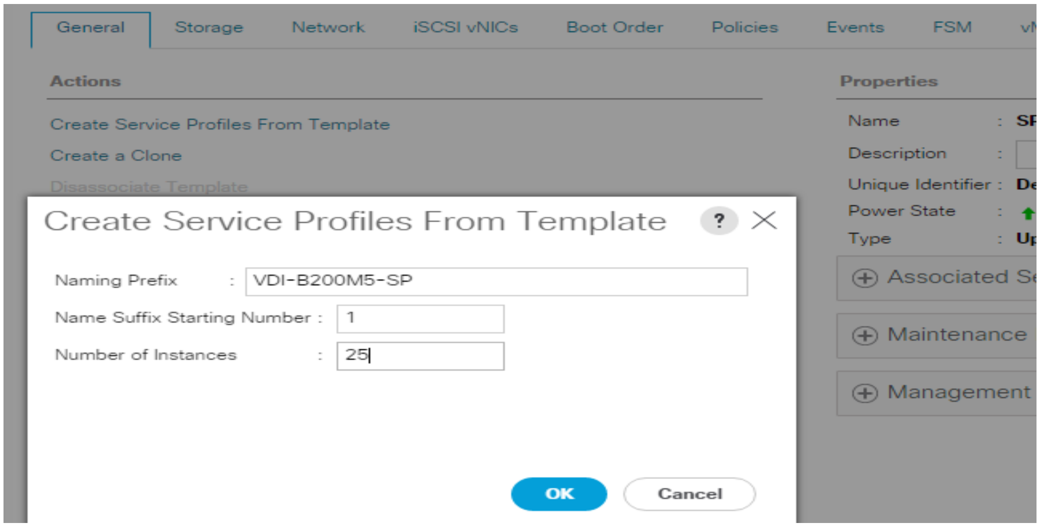

Naming Prefix : VDI-B200M5-SP

Name Suffix Starting Number : 1

Number of Instances : 25

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

NetApp A300 Storage System Configuration

The following section includes instructions on the steps necessary to perform initial setup and configuration of the NetApp A300 storage system. Specific details of the configuration as tested can be found in the NetApp A300 Storage Architecture Design section below.

NetApp All Flash FAS A300 Controllers

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

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

NetApp Hardware Universe

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

1. Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the HWU application at the NetApp Support site.

2. Access the HWU application to view the System Configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

3. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

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

Disk Shelves

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

When using SAS disk shelves with NetApp storage controllers, refer to the SAS Disk Shelves Universal SAS and ACP Cabling Guide for cabling guidelines.

NetApp ONTAP 9.1

Complete Configuration Worksheet

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

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the ONTAP 9.1 Software Setup Guide to learn about configuring ONTAP. Table 18 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 18 ONTAP Software Installation Prerequisites

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

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

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

4. To install new software, select option 7.

5. Enter `y` to perform an upgrade.

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

7. Enter `y` to reboot now.

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

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

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

This web server must be pingable.

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

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

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

12. Enter `y` to reboot the node.

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

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

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

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

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

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

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

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

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

4. To install new software, select option 7.

5. Enter y to perform an upgrade.

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

7. Enter y to reboot now.

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

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

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

This web server must be pingable.

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

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

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

12. Enter y to reboot the node.

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

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

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

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

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

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

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

Set Up Node

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

1. Follow the prompts to set up node 01:

```
Welcome to node setup.

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

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

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

Type yes to confirm and continue {yes}: yes

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

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

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

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

Table 19 Cluster Create in ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	`<clustername>`
ONTAP base license	`<cluster-base-license-key>`
Cluster management IP address	`<clustermgmt-ip>`
Cluster management netmask	`<clustermgmt-mask>`
Cluster management gateway	`<clustermgmt-gateway>`
Cluster node 01 IP address	`<node01-mgmt-ip>`
Cluster node 01 netmask	`<node01-mgmt-mask>`
Cluster node 01 gateway	`<node01-mgmt-gateway>`
Cluster node 02 IP address	`<node02-mgmt-ip>`
Cluster node 02 netmask	`<node02-mgmt-mask>`
Cluster node 02 gateway	`<node02-mgmt-gateway>`
Node 01 service processor IP address	`<node01-SP-ip>`
Node 02 service processor IP address	`<node02-SP-ip>`

Cluster Detail	Cluster Detail Value
DNS domain name	`<dns-domain-name>`
DNS server IP address	`<dns-ip>`
NTP server IP address	`<ntp-ip>`

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

3. Click Guided Setup on the Welcome screen.

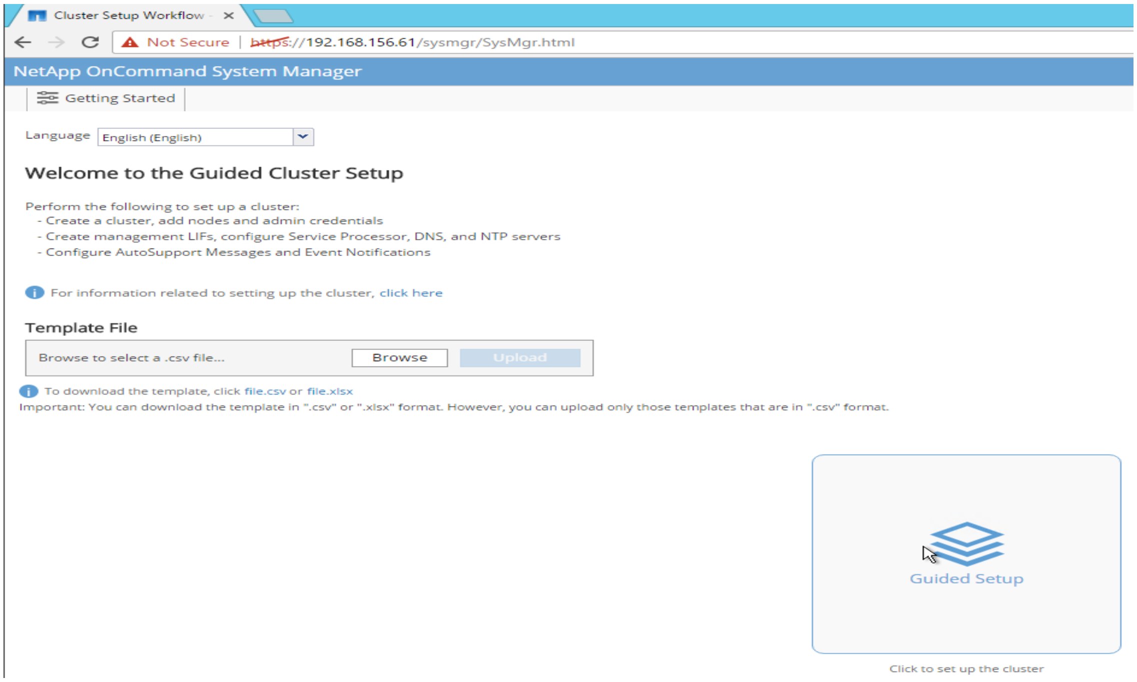

Welcome to the Guided Cluster Setup

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

For information related to setting up the cluster, click here

Template File

Browse to select a .csv file...

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

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

Click to set up the cluster

4. In the Cluster screen, do the following:

   a. Enter the cluster and node names.

   b. Select the cluster configuration.

   c. Enter and confirm the password.

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

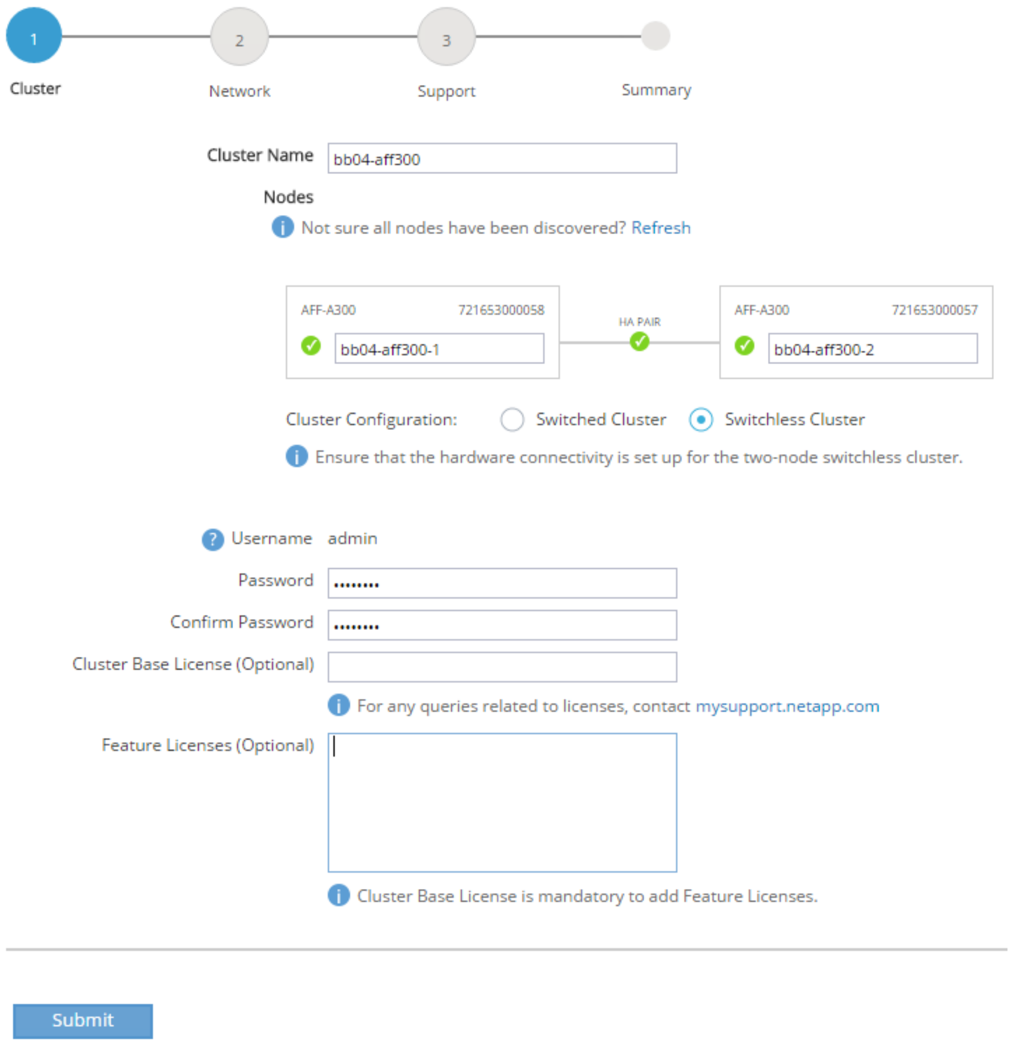

Cluster Name: bb04-aff300

Nodes

Not sure all nodes have been discovered? Refresh

AFF-A300	721653000058	bb04-aff300-1
AFF-A300	721653000057	bb04-aff300-2

Cluster Configuration: ○ Switched Cluster ● Switchless Cluster

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

Username: admin

Password: ••••••••

Confirm Password: ••••••••

Cluster Base License (Optional):

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

Feature Licenses (Optional):

Cluster Base License is mandatory to add Feature Licenses.

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

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

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

5. Click Submit.

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

- Cluster Management

  - Enter the IP address, netmask, gateway and port details.

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

7. Click Submit.

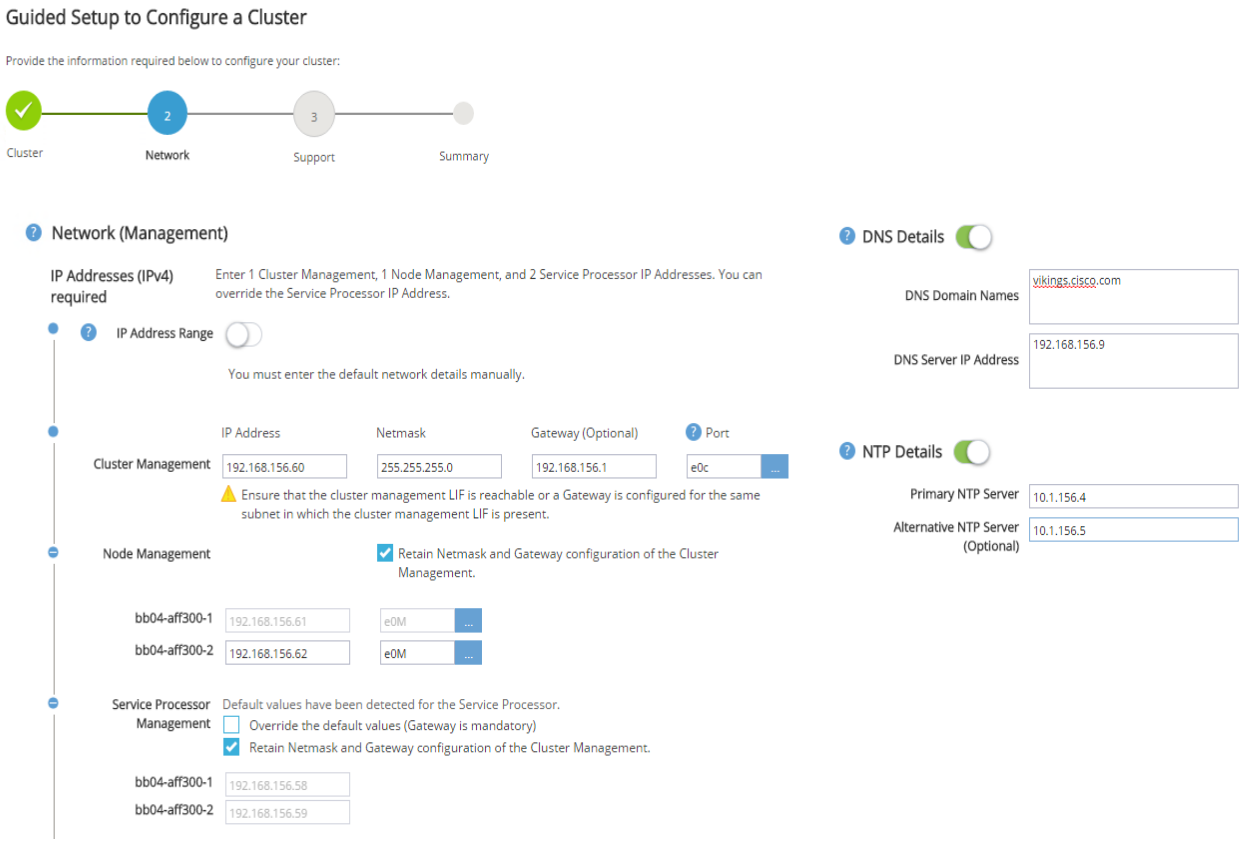

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Network (Management)

IP Addresses (IPv4) required: Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range

You must enter the default network details manually.

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

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

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

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

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

□ Override the default values (Gateway is mandatory)

☑ Retain Netmask and Gateway configuration of the Cluster Management.

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

DNS Details

DNS Domain Names: vikings.cisco.com

DNS Server IP Address: 192.168.156.9

NTP Details

Primary NTP Server: 10.1.156.4

Alternative NTP Server (Optional): 10.1.156.5

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

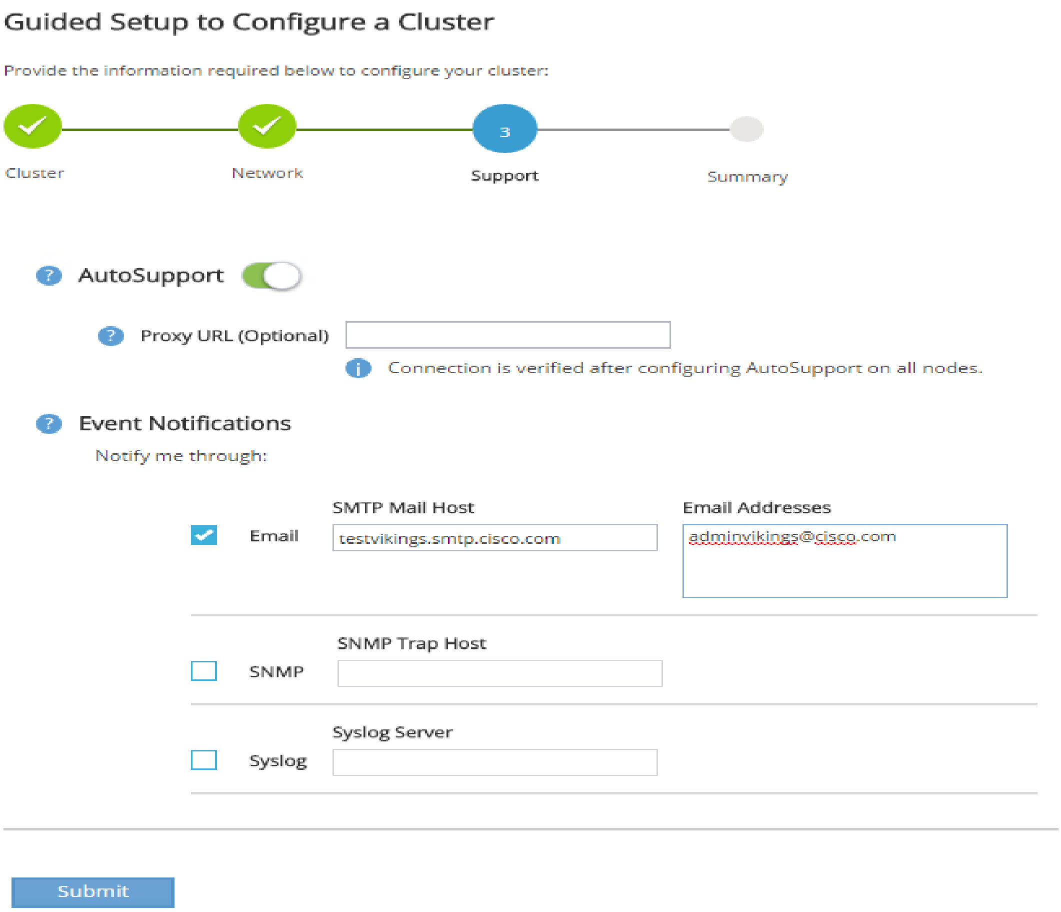

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

AutoSupport: on

Proxy URL (Optional):

Connection is verified after configuring AutoSupport on all nodes.

Event Notifications

Notify me through:

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

□ SNMP — SNMP Trap Host:

□ Syslog — Syslog Server:

9. Click Submit.

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

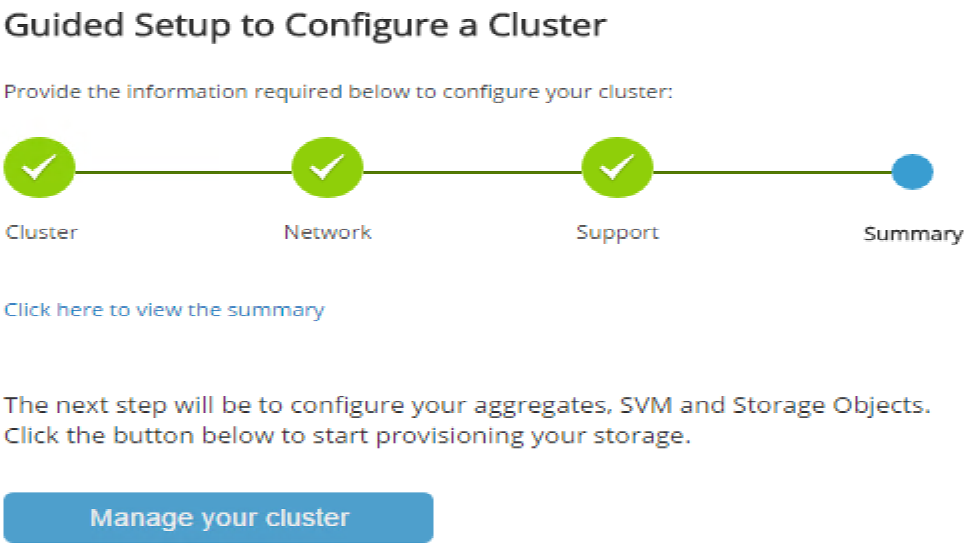

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Click here to view the summary

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

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

Log into the Cluster

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

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

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

Zero All Spare Disks

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

```
disk zerospares
```

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

Set Onboard Unified Target Adapter 2 Port Personality

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

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

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

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

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

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

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

Set Auto-Revert on Cluster Management

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

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

Run the following command:

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

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, e`0d`, `e2a`, and `e2e`) should be removed from the default broadcast domain, leaving just the management network ports (`e0c` and `e0M`). To perform this task, run the following commands:

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

Set Up Service Processor Network Interface

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

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

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

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

Create Aggregates

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

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

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

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

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

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

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

Verify Storage Failover

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

1. Verify the status of the storage failover.

```
storage failover show
```

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

2. Enable failover on one of the two nodes.

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

Enabling failover on one node enables it for both nodes.

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

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

```
cluster ha show
```

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

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

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

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

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

Disable Flow Control on 10GbE and 40GbE Ports

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

1. Run the following commands to configure node 01:

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

2. Run the following commands to configure node 02:

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

Disable Unused FCoE Capability on CNA Ports

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

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

Configure Network Time Protocol

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

1. Set the time zone for the cluster.

```
timezone <timezone>
```

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

2. Set the date for the cluster.

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

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

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

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

Configure Simple Network Management Protocol

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

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

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

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

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

Configure SNMPv1 Access

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

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

Configure AutoSupport

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

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

Enable Cisco Discovery Protocol

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

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

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

Create Jumbo Frame MTU Broadcast Domains in ONTAP

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

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

Create Interface Groups

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

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

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

ifgrp show
```

Create VLANs

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

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

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

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

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

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

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

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

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

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

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

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

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

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

Create Load-Sharing Mirrors of SVM Root Volume

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

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

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

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

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

3. Create the mirroring relationships.

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

4. Initialize the mirroring relationship.

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

Create Block Protocol (FC) Service

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

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

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

Configure HTTPS Access

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

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

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

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

```
security certificate show
```

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

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

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

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

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

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

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

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

6. Disable HTTP cluster management access.

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

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

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

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

Configure NFSv3

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

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

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

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

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

Create FlexVol Volumes

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

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

FlexVol volumes are created to house boot LUNs for ESXi servers, datastore LUNs for virtual desktops and RDS hosts, and for CIFS shares hosting user profile data. For specific details about the volumes created during this validation, see the Storage Configuration section below.

To create FlexVol volumes, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_ds01 -aggregate aggr1_AFF300_01 -size 6TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume rdsh_1 -aggregate aggr1_AFF300_02 -size 6TB -state online -
policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_AFF300_01 -size 500GB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_01 -aggregate aggr1_AFF300_01 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_02 -aggregate aggr1_AFF300_02 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_03 -aggregate aggr1_AFF300_01 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_04 -aggregate aggr1_AFF300_02 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_05 -aggregate aggr1_AFF300_01 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume vdi_fc_06 -aggregate aggr1_AFF300_02 -size 11TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path Infra-SVM:rootvol
```

Create Boot LUNs

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

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 10GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-02 -size 10GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VDI-1 -size 10GB -ostype vmware -space-reserve
disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VDI-2 -size 10GB -ostype vmware -space-reserve
disabled

…
lun create -vserver Infra-SVM -volume vdi_fc_01 -lun vdi_fc_01 -size 10TB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume vdi_fc_02 -lun vdi_fc_02 -size 10TB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume vdi_fc_03 -lun vdi_fc_03 -size 10TB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume vdi_fc_04 -lun vdi_fc_04 -size 10TB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume vdi_fc_05 -lun vdi_fc_05 -size 10TB -ostype vmware -space-
reserve disabled
lun create -vserver Infra-SVM -volume vdi_fc_06 -lun vdi_fc_06 -size 10TB -ostype vmware -space-
reserve disabled
```

Create igroups

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

To create igroups, run the following commands:

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol fcp –ostype
vmware –initiator <vm-host-infra-01-iqn-a>,<vm-host-infra-01-iqn-b>
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-02 –protocol fcp –ostype
vmware –initiator <vm-host-infra-02-iqn-a>,<vm-host-infra-02-iqn-b>
igroup create –vserver Infra-SVM –igroup VDI-1 –protocol fcp –ostype vmware –
initiator <vm-host-VDI-1-iqn-a>,<vm-host-VDI-1-iqn-b>
igroup create –vserver Infra-SVM –igroup VDI-2 –protocol fcp –ostype vmware –
initiator <vm-host-VDI-2-iqn-a>,<vm-host-VDI-2-iqn-b>
igroup create –vserver Infra-SVM –igroup VDI-3 –protocol fcp –ostype vmware –
initiator <vm-host-VDI-3-iqn-a>,<vm-host-VDI-3-iqn-b>
…
igroup create –vserver Infra-SVM –igroup VDI-Cluster –protocol fcp –ostype
vmware –initiator <vm-host-VDI-1-iqn-a>,<vm-host-VDI-1-iqn-b>,<vm-host-VDI-2-
iqn-a>,<vm-host-VDI-2-iqn-b>,<vm-host-VDI-3-iqn-a>,<vm-host-VDI-3-iqn-b>,<…>
```

To view igroups, type igroup show.

Map Boot LUNs to igroups

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

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-01 –igroup VM-Host-Infra-01 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-02 –igroup VM-Host-Infra-02 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-1 -igroup VDI-1 -lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-2 -igroup VDI-2 -lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VDI-3 -igroup VDI-3 -lun-id 0
…
lun map –vserver Infra-SVM –volume infra_ds01 –lun infra_ds01 –igroup VDI-Cluster –lun-id 1
lun map -vserver Infra-SVM -volume rdsh_1 -lun rdsh_1 -igroup VDI-Cluster -lun-id 2
lun map -vserver Infra-SVM -volume vdi_fc_01 -lun vdi_fc_01 -igroup VDI-Cluster -lun-id 3
lun map -vserver Infra-SVM -volume vdi_fc_02 -lun vdi_fc_02 -igroup VDI-Cluster -lun-id 4
lun map -vserver Infra-SVM -volume vdi_fc_03 -lun vdi_fc_03 -igroup VDI-Cluster -lun-id 5
lun map -vserver Infra-SVM -volume vdi_fc_04 -lun vdi_fc_04 -igroup VDI-Cluster -lun-id 6
lun map -vserver Infra-SVM -volume vdi_fc_05 -lun vdi_fc_05 -igroup VDI-Cluster -lun-id 7
lun map -vserver Infra-SVM -volume vdi_fc_06 -lun vdi_fc_06 -igroup VDI-Cluster -lun-id 8
```

Schedule Deduplication

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

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

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

Create FC LIFs

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

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

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

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

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

network interface show
```

Create NFS LIF

To create an NFS LIF, run the following commands:

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

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

network interface show
```

Add Infrastructure SVM Administrator

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

1. Run the following commands:

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

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

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

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

network route show
```

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

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

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

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

NetApp A300 Storage Configuration

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

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

Cluster Details

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

Table 20 lists the cluster details.

Table 20 Cluster Details

Cluster Name	ONTAP Version	Node Count	Data SVM Count	Cluster Raw Capacity
AFF A300	9.1P5	2	1	83.84TB

Storage Details

Table 21 lists the storage details for each HA pair.

Table 21 Storage Details

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

Raw capacity is not the same as usable capacity.

Drive Allocation Details

Table 22 lists the drive allocation details for each node.

Table 22 Drive Allocation Details

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

Raw capacity is not the same as usable capacity.

Adapter Card Details

Table 23 lists the adapter cards present in each node.

Table 23 Adapter Card Details

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

Firmware Details

Table 24 lists the relevant firmware details for each node.

Table 24 Firmware Details

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

Network Port Settings

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

Error! Reference source not found.lists the network port settings.

Table 25 Network Port Settings for ONTAP

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

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

Network Port Interface Group Settings

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

Error! Reference source not found.lists the network port ifgrp settings.

Table 26 Network Port Ifgrp Settings

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

Network Routes

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

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

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

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

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

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

Table 27 Network Routes

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

Network Port Broadcast Domains

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

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

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

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

Table 28 Network Port Broadcast Domains

Cluster Name	Broadcast Domain	IPspace Name	MTU Size	Subnet Names	Port List	Failover Group Names
AFF A300	Cifs	Default	1500		AFF-A300-01:a0a-62 AFF-A300-02:a0a-62	cifs
AFF A300	Cluster	Cluster	9000		AFF-A300-01:e0a AFF-A300-01:e0b AFF-A300-02:e0a AFF-A300-02:e0b	Cluster
AFF A300	Default	Default	1500		AFF-A300-01:e0c AFF-A300-01:e0M AFF-A300-02:e0c AFF-A300-02:e0M	Default
AFF A300	IB	Default	1500		AFF-A300-01:a0a-61 AFF-A300-02:a0a-61	IB

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

Aggregate Configuration

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

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

Table 29 contains all aggregate configuration information.

Table 29 Aggregate Configuration

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

Storage Virtual Machines

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

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

SVM Configuration

Table 30 lists the SVM configuration.

Table 30 SVM Configuration

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

SVM Storage Configuration

Table 31 lists the SVM storage configuration.

Table 31 SVM Storage Configuration

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

Volume Configuration

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

Error! Reference source not found.the FlexVol configuration.

Table 32 FlexVol Configuration

Cluster Name	SVM Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
AFF A300	Infra	esxi_boot	aggr1_AFF300_01	RW	default	default	unix	500.00 GB
AFF A300	Infra	home	aggr1_AFF300_02	RW	default	default	ntfs	1.00TB
AFF A300	Infra	infra_ds01	aggr1_AFF300_02	RW	default	default	unix	6.00TB
AFF A300	Infra	rdsh_1	aggr1_AFF300_01	RW	default	default	unix	6.00TB

Cluster Name	SVM Name	Volume Name	Containing Aggregate	Type	Snapshot Policy	Export Policy	Security Style	Size Nominal
AFF A300	Infra	vdi_ds01	aggr1_AFF300_02	RW	default	default	unix	6.00TB
AFF A300	Infra	vdi_fc_ds01	aggr1_AFF300_01	RW	default	default	unix	11.00TB
AFF A300	Infra	vdi_fc_ds02	aggr1_AFF300_01	RW	default	default	unix	11.00TB
AFF A300	Infra	vdi_fc_ds03	aggr1_AFF300_01	RW	default	default	unix	11.00TB
AFF A300	Infra	vdi_fc_ds04	aggr1_AFF300_02	RW	default	default	unix	11.00TB
AFF A300	Infra	vdi_fc_ds05	aggr1_AFF300_02	RW	default	default	unix	11.00TB
AFF A300	Infra	vdi_fc_ds06	aggr1_AFF300_02	RW	default	default	unix	11.00TB

Protocol Configuration

NAS

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

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

Logical Interfaces

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

LIFs can be hosted on the following ports:

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

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

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

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

NAS Logical Interface Settings

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

Table 33 NAS LIF Settings

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

Windows File Services

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

CIFS Servers

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

Table 34 lists CIFS server configuration information.

Table 34 CIFS Servers

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

CIFS Options

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

Table 35 lists CIFS options.

Table 35 CIFS Options

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

CIFS Local Users and Groups

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

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

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

Privileges do not provide ONTAP general administrative capabilities.

CIFS Shares

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

Table 36 lists the CIFS shares.

Table 36 CIFS Shares

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

Cluster Name	SVM Name	Share Name	Path	Share Properties	Symlink Properties	Share ACL
AFF A300	Infra	Profile$	/home/Profile	oplocks browsable changenotify show_previous_version s	symlinks	Everyone:Full Control

CIFS Home Directory Search Paths

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

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

Table 37 lists the CIFS home directory search paths.

Table 37   CIFS Home Directory Search Paths

Cluster Name	SVM Name	Position	Path
AFF A300	Infra	1	/home/LoginVSI
AFF A300	Infra	2	/home/RDSHVSI

SAN

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

LUNs

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

Table 38 lists the LUN details.

Table 38   LUN Configuration

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF A300	Infra	/vol/esxi_boot/infra_host_01	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/infra_host_02	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-1	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-2	True	True	vmware	False	10.00GB

Cluster Name	SVM Name	Path	Mapped	Online	Protocol Type	Read Only	Size
AFF A300	Infra	/vol/esxi_boot/VDI-3	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-4	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-5	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-6	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-7	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-9	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-10	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-11	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-12	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-13	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-14	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-15	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-17	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-18	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-19	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-20	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-21	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-22	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-23	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-24	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/esxi_boot/VDI-25	True	True	vmware	False	10.00GB
AFF A300	Infra	/vol/infra_ds01/infra_ds01	True	True	vmware	False	5.00TB
AFF A300	Infra	/vol/rdsh_1/RDSH_ds01	True	True	vmware	False	5.00TB
AFF A300	Infra	/vol/vdi_ds01/vdi_ds01	True	True	vmware	False	5.00TB
AFF A300	Infra	/vol/vdi_fc_ds01/vdi_fc_ds01	True	True	vmware	False	10.00TB
AFF A300	Infra	/vol/vdi_fc_ds02/vdi_fc_ds02	True	True	vmware	False	10.00TB
AFF A300	Infra	/vol/vdi_fc_ds03/vdi_fc_ds03	True	True	vmware	False	10.00TB
AFF A300	Infra	/vol/vdi_fc_ds04/vdi_fc_ds04	True	True	vmware	False	10.00TB
AFF A300	Infra	/vol/vdi_fc_ds05/vdi_fc_ds05	True	True	vmware	False	10.00TB

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

Initiator Groups

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

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

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

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

Table 39 lists the igroups that have been created.

Table 39 Initiator Groups

Cluster Name	SVM Name	Initiator Group Name	Protocol	OS Type	ALUA	Initiators Logged In
AFF A300	Infra	Infra_cluster	fcp	vmware	True	full
AFF A300	Infra	SP_Infra1	fcp	vmware	True	full
AFF A300	Infra	SP_Infra2	fcp	vmware	True	full
AFF A300	Infra	VDI_cluster	fcp	vmware	True	full
AFF A300	Infra	VDI-1	fcp	vmware	True	full
AFF A300	Infra	VDI-2	fcp	vmware	True	full
AFF A300	Infra	VDI-3	fcp	vmware	True	full
AFF A300	Infra	VDI-4	fcp	vmware	True	full
AFF A300	Infra	VDI-5	fcp	vmware	True	full
AFF A300	Infra	VDI-6	fcp	vmware	True	full
AFF A300	Infra	VDI-7	fcp	vmware	True	full
AFF A300	Infra	VDI-9	fcp	vmware	True	full
AFF A300	Infra	VDI-10	fcp	vmware	True	full
AFF A300	Infra	VDI-11	fcp	vmware	True	full
AFF A300	Infra	VDI-12	fcp	vmware	True	full
AFF A300	Infra	VDI-13	fcp	vmware	True	full
AFF A300	Infra	VDI-14	fcp	vmware	True	full
AFF A300	Infra	VDI-15	fcp	vmware	True	full
AFF A300	Infra	VDI-17	fcp	vmware	True	full
AFF A300	Infra	VDI-18	fcp	vmware	True	full
AFF A300	Infra	VDI-19	fcp	vmware	True	full
AFF A300	Infra	VDI-20	fcp	vmware	True	full
AFF A300	Infra	VDI-21	fcp	vmware	True	full
AFF A300	Infra	VDI-22	fcp	vmware	True	full
AFF A300	Infra	VDI-23	fcp	vmware	True	full
AFF A300	Infra	VDI-24	fcp	vmware	True	full
AFF A300	Infra	VDI-25	fcp	vmware	True	full

FCP Logical Interface Settings

Table 40 lists the FCP LIF settings.

Table 40 FCP LIF Settings

Cluster Name	SVM Name	Interface Name	Status (Admin/Oper)	Port Name	Current Node	Current Port	Is Home
AFF A300	Infra	fcp_01a	up/up	20:01:00:a0:98:af:bd:e8	AFF-A300-01	0g	True
AFF A300	Infra	fcp_01b	up/up	20:02:00:a0:98:af:bd:e8	AFF-A300-01	0h	True
AFF A300	Infra	fcp_02a	up/up	20:03:00:a0:98:af:bd:e8	AFF-A300-02	0g	True
AFF A300	Infra	fcp_02b	up/up	20:04:00:a0:98:af:bd:e8	AFF-A300-02	0h	True

FCP / FCoE

FCP Service Configuration

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

Table 41 lists the FCP service configuration details.

Table 41 FCP Service Configuration

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

FCP Adapter Configuration

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

Table 42 lists the details of the storage controller target ports and the WWPN address of each.

Table 42 FCP Adapter Configuration

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

Node Name	Adapter Name	State	Data Link Rate	Media Type	Speed	Port Name
AFF-A300-02	0f	offlined by user/system	0	ptp	auto	50:0a:09:81:80:d3:67:d3
AFF-A300-02	0g	online	8	ptp	auto	50:0a:09:84:80:d3:67:d3
AFF-A300-02	0h	online	8	ptp	auto	50:0a:09:83:80:d3:67:d3

Storage Efficiency and Space Management

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

Volume Efficiency

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

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

Table 43 lists the volume efficiency settings.

Table 43 Volume Efficiency Settings

Cluster Name	SVM Name	Volume Name	Space Guaran-tee	Dedupe	Schedule Or Policy Name	Compres-sion	Inline Compres-sion
AFF A300	Infra	esxi_boot	none	True	sun-sat@1	True	True
AFF A300	Infra	home	none	True	inline-only	True	True
AFF A300	Infra	infra_ds01	none	True	inline-only	True	True
AFF A300	Infra	rdsh_1	none	True	inline-only	True	True
AFF A300	Infra	svm_root	volume		-		
AFF A300	Infra	vdi_ds01	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds01	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds02	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds03	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds04	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds05	none	True	inline-only	True	True
AFF A300	Infra	vdi_fc_ds06	none	True	inline-only	True	True

LUN Efficiency

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

With thin provisioning, you can present more storage space to the hosts connecting to the SVM than what is actually available on the SVM. Storage provisioning with thinly provisioned LUNs enables storage administrators to provide actual storage that the LUN needs. As ONTAP writes blocks to the LUN, the LUN increases in size automatically.

Table 44 lists the LUN efficiency settings.

Table 44   LUN Efficiency Settings

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF A300	Infra	/vol/esxi_boot/infra_host_01	False	False
AFF A300	Infra	/vol/esxi_boot/infra_host_02	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-1	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-2	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-3	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-4	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-5	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-6	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-7	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-9	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-10	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-11	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-12	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-13	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-14	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-15	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-17	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-18	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-19	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-20	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-21	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-22	False	False

Cluster Name	SVM Name	Path	Space Reservation Enabled	Space Allocation Enabled
AFF A300	Infra	/vol/esxi_boot/VDI-23	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-24	False	False
AFF A300	Infra	/vol/esxi_boot/VDI-25	False	False
AFF A300	Infra	/vol/infra_ds01/infra_ds01	True	False
AFF A300	Infra	/vol/rdsh_1/rdsh_1	False	False
AFF A300	Infra	/vol/vdi_ds01/vdi_ds01	False	False
AFF A300	Infra	/vol/vdi_fc_ds01/vdi_fc_ds01	False	False
AFF A300	Infra	/vol/vdi_fc_ds02/vdi_fc_ds02	False	False
AFF A300	Infra	/vol/vdi_fc_ds03/vdi_fc_ds03	False	False
AFF A300	Infra	/vol/vdi_fc_ds04/vdi_fc_ds04	False	False
AFF A300	Infra	/vol/vdi_fc_ds05/vdi_fc_ds05	False	False
AFF A300	Infra	/vol/vdi_fc_ds06/vdi_fc_ds06	False	False

Configure MDS 9100 Series

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

In this solution, we utilized the Cisco MDS 9148S Switches for Fiber Channel Switching. For racking, cable and initial setup of the MDS switches, please refer to the Quick Start Guide: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/hw/9148/quick/guide/MDS_9148_QSG.pdf

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

2. In this study, we used two separate fabrics each with their own unique VSAN. Fabric A is configured for VSAN400 while Fabric B for VSAN401. In our initial Cisco UCS configuration, you will see where we configured fiber cables on ports 1 and 4 and configured a FC port-channel. FI-A’s FC port channel is configured for VSAN400 and FI-B’s FC port-channel for VSAN401.

Figure 29 VSAN 400 Configured for Fabric A

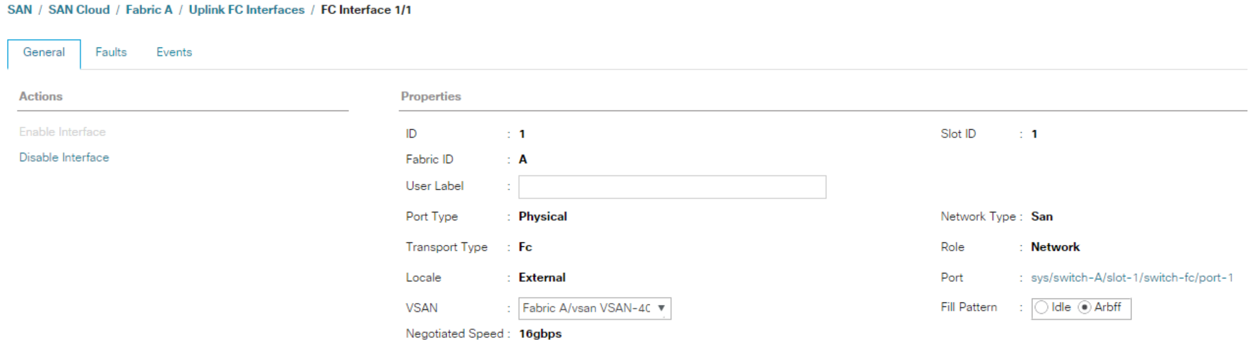

SAN / SAN Cloud / Fabric A / Uplink FC Interfaces / **FC Interface 1/1**

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID	: **1**	Slot ID	: **1**
Fabric ID	: **A**		
User Label	:		
Port Type	: **Physical**	Network Type	: **San**
Transport Type	: **Fc**	Role	: **Network**
Locale	: **External**	Port	: sys/switch-A/slot-1/switch-fc/port-1
VSAN	: Fabric A/vsan VSAN-4C ▼	Fill Pattern	: ○ Idle ● Arbff
Negotiated Speed	: **16gbps**		

Figure 30 VSAN 401 Configured for Fabric B

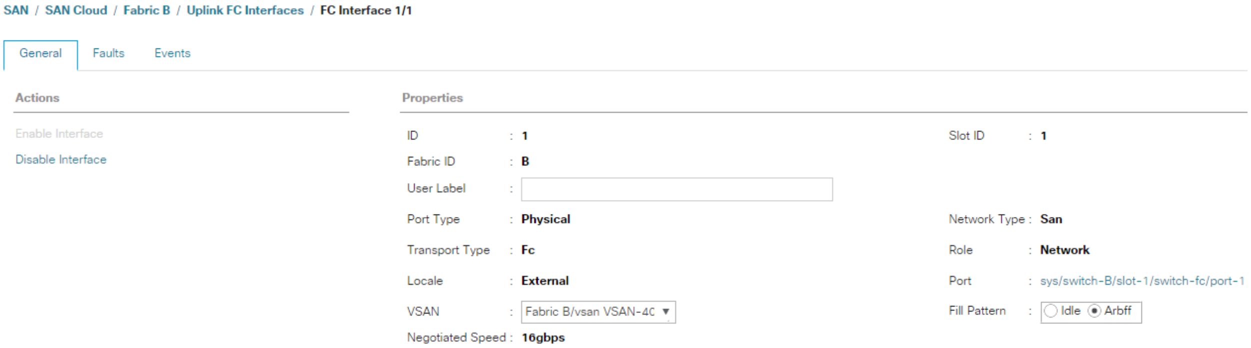

SAN / SAN Cloud / Fabric B / Uplink FC Interfaces / **FC Interface 1/1**

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID	: **1**	Slot ID	: **1**
Fabric ID	: **B**		
User Label	:		
Port Type	: **Physical**	Network Type	: **San**
Transport Type	: **Fc**	Role	: **Network**
Locale	: **External**	Port	: sys/switch-B/slot-1/switch-fc/port-1
VSAN	: Fabric B/vsan VSAN-4C ▼	Fill Pattern	: ○ Idle ● Arbff
Negotiated Speed	: **16gbps**		

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

Figure 31 MDS Switch VSAN Configuration Connectivity

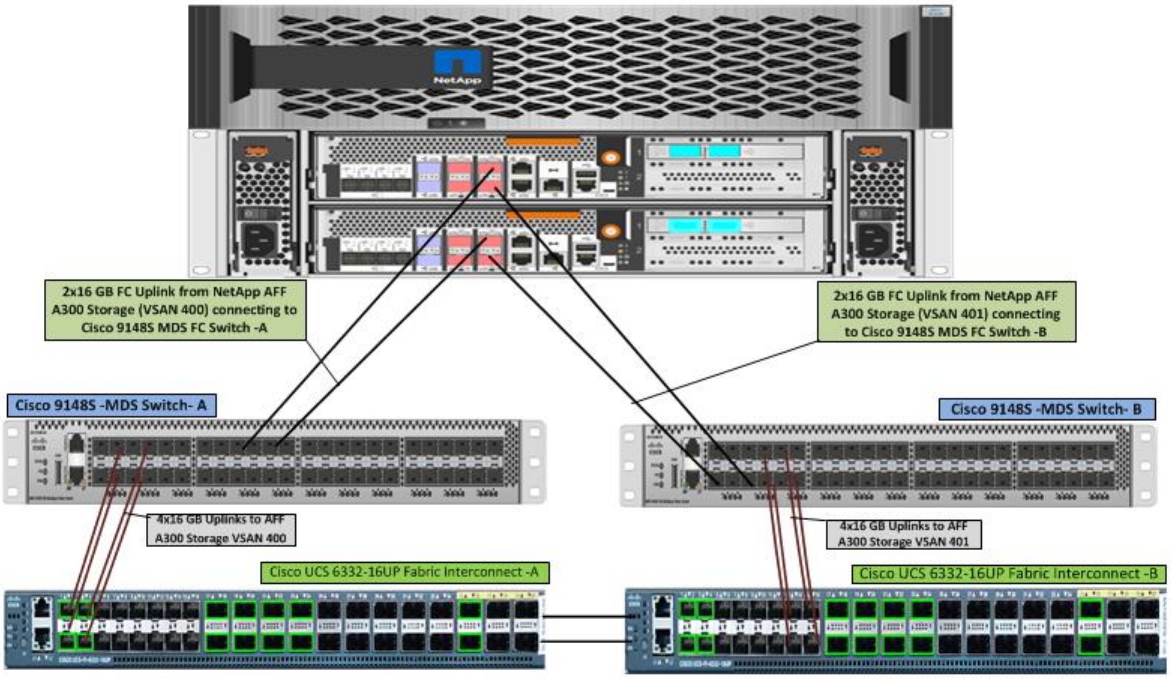
NetApp
2x16 GB FC Uplink from NetApp AFF A300 Storage (VSAN 400) connecting to Cisco 9148S MDS FC Switch -A
2x16 GB FC Uplink from NetApp AFF A300 Storage (VSAN 401) connecting to Cisco 9148S MDS FC Switch -B
Cisco 9148S -MDS Switch- A
Cisco 9148S -MDS Switch- B
4x16 GB Uplinks to AFF A300 Storage VSAN 400
4x16 GB Uplinks to AFF A300 Storage VSAN 401
Cisco UCS 6332-16UP Fabric Interconnect -A
Cisco UCS 6332-16UP Fabric Interconnect -B

We used a total of 8 16Gb FC links, four connections from each storage controller to each the MDS 9148S 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 detail how to add a single host on both 9148S A and B. You will need to configure all hosts that will access the NetApp array in these commands. The entire MDS 9148S FC switch configuration is included in Appendix A .

MDS-A

```
Zoneset name AFF-A300_VDI vsan 400
Member {ESXi hostname-fc0}
Exit
Zoneset activate name AFF-A300_VDI vsan 400
Zone commit vsan 400
Exit
Copy running-config startup-config
```

MDS-B

```
Zoneset name AFF-A300_VDI vsan 401
Member {ESXi hostname-fc1}
Exit
Zoneset activate name AFF-A300_VDI vsan 401
Zone commit vsan 401
Exit
Copy running-config startup-config
```

Installing and Configuring VMware ESXi 6.5

This section provides detailed instructions for installing VMware ESXi 6.5 Update1 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6.5 Update 1

To download the Cisco Custom Image for ESXi 6.5 Update 1, complete the following steps:

1. Click the following link vmware login page.

2. Type your email or customer number and the password and then click Log in.

3. Click on the following link:

https://my.vmware.com/web/vmware/details?productId=614&downloadGroup=ESXI65U1

4. Click Download Now.

5. Save it to your destination folder.

This ESXi 6.5 Cisco custom image includes updates for the fNIC and neNIC drivers. The versions that are part of this image are: neNIC: 1.0.6.0-1; fNIC: 1.6.0.34

KVM Access to Hosts

To log in to the Cisco UCS environment, complete the following steps:

1. Log in to Cisco UCS Manager.
2. The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.
3. Open a Web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
4. Log in to Cisco UCS Manager by using the admin user name and password.
5. From the main menu, click the Servers tab.
6. Select Servers > Service Profiles > root > VM-Host-01.
7. Right-click VM-Host-01 and select KVM Console.
8. Repeat steps for 4-6 for all host servers.

Set Up VMware ESXi Installation

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click the Virtual Media tab.
2. Click Add Image.
3. Browse to the ESXi installer ISO image file and click Open.
4. Select the Mapped checkbox to map the newly added image.
5. Click the KVM tab to monitor the server boot.
6. Boot the server by selecting Boot Server and click OK, then click OK again.

Install ESXi

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.

3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

4. Select the AFF A300 boot LUN.

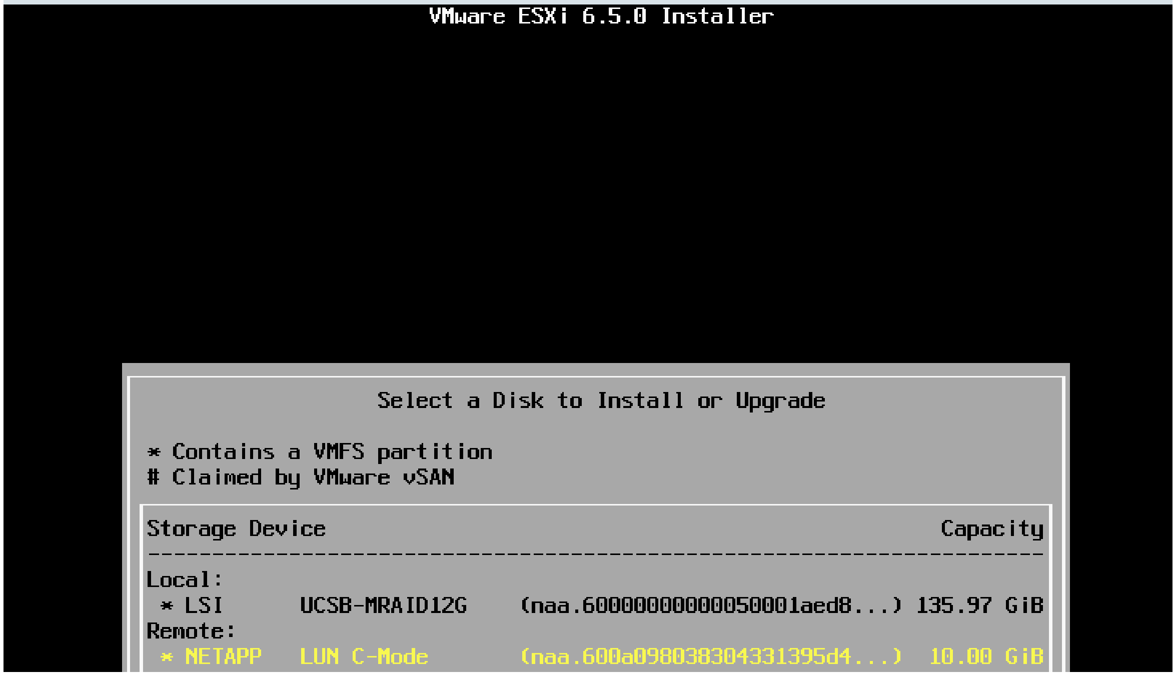

```
VMware ESXi 6.5.0 Installer

Select a Disk to Install or Upgrade

* Contains a VMFS partition
# Claimed by VMware vSAN

Storage Device                                                  Capacity
-------------------------------------------------------------------------
Local:
 * LSI      UCSB-MRAID12G    (naa.60000000000050001aed8...) 135.97 GiB
Remote:
 * NETAPP   LUN C-Mode       (naa.600a098038304331395d4...)  10.00 GiB
```

5. (NetApp LUN C-Mode(naa.600a098038304331395d4) 10 GB that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

6. Select the appropriate keyboard layout and press Enter.

7. Enter and confirm the root password and press Enter.

8. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.

9. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

10. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, click Yes to unmap the image.

11. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host.

To configure the ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the VLAN in-band management ID and press Enter.
6. From the Configure Management Network menu, select IP Configuration and press Enter.
7. Select the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the first ESXi host.
9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Select the IPv6 Configuration option and press Enter.
13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
14. Select the DNS Configuration option and press Enter.

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

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.

24. Press Enter to exit the window.

25. Press Esc to log out of the VMware console.

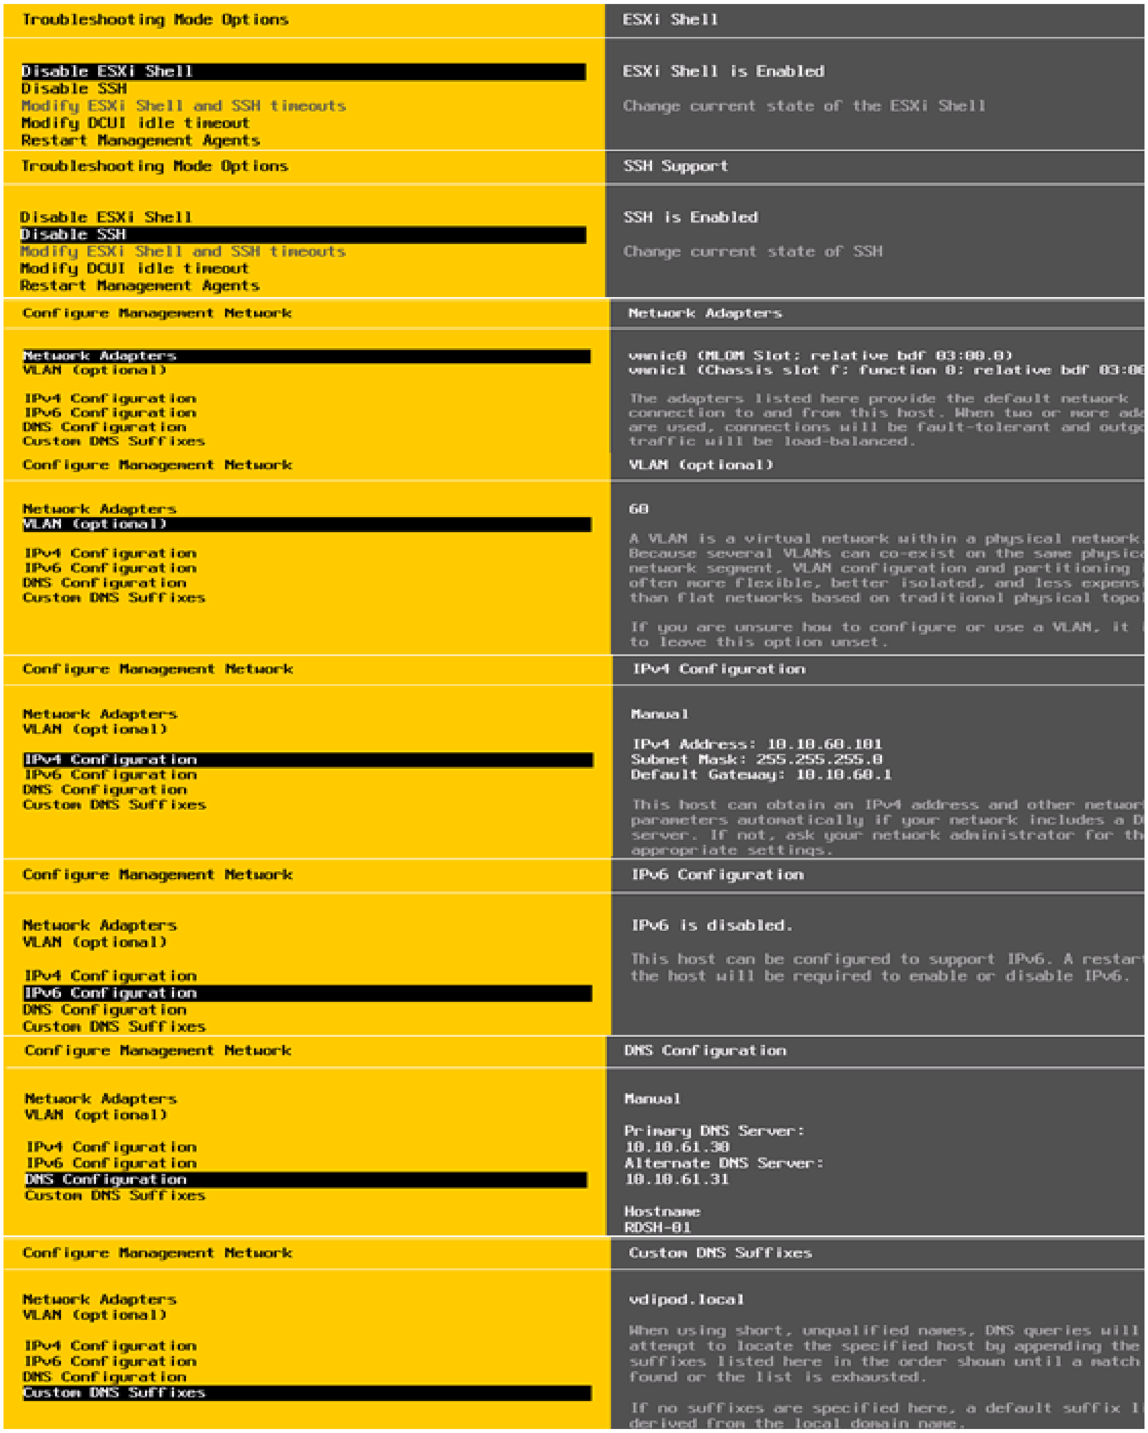

```
Troubleshooting Mode Options              ESXi Shell

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

```
Troubleshooting Mode Options              SSH Support

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

```
Configure Management Network              Network Adapters

Network Adapters                          vmnic0 (MLOM Slot: relative bdf 03:00.0)
VLAN (optional)                           vmnic1 (Chassis slot f: function 0: relative bdf 03:00

IPv4 Configuration                        The adapters listed here provide the default network
IPv6 Configuration                        connection to and from this host. When two or more ada
DNS Configuration                         are used, connections will be fault-tolerant and outgo
Custom DNS Suffixes                       traffic will be load-balanced.
```

```
Configure Management Network              VLAN (optional)

Network Adapters                          60
VLAN (optional)
                                          A VLAN is a virtual network within a physical network.
IPv4 Configuration                        Because several VLANs can co-exist on the same physica
IPv6 Configuration                        network segment, VLAN configuration and partitioning i
DNS Configuration                         often more flexible, better isolated, and less expensi
Custom DNS Suffixes                       than flat networks based on traditional physical topol

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

```
Configure Management Network              IPv4 Configuration

Network Adapters                          Manual
VLAN (optional)
                                          IPv4 Address: 10.10.60.101
IPv4 Configuration                        Subnet Mask: 255.255.255.0
IPv6 Configuration                        Default Gateway: 10.10.60.1
DNS Configuration
Custom DNS Suffixes                       This host can obtain an IPv4 address and other networ
                                          parameters automatically if your network includes a D
                                          server. If not, ask your network administrator for th
                                          appropriate settings.
```

```
Configure Management Network              IPv6 Configuration

Network Adapters                          IPv6 is disabled.
VLAN (optional)
                                          This host can be configured to support IPv6. A restart
IPv4 Configuration                        the host will be required to enable or disable IPv6.
IPv6 Configuration
DNS Configuration
Custom DNS Suffixes
```

```
Configure Management Network              DNS Configuration

Network Adapters                          Manual
VLAN (optional)
                                          Primary DNS Server:
IPv4 Configuration                        10.10.61.30
IPv6 Configuration                        Alternate DNS Server:
DNS Configuration                         10.10.61.31
Custom DNS Suffixes
                                          Hostname
                                          RDSH-01
```

```
Configure Management Network              Custom DNS Suffixes

Network Adapters                          vdipod.local
VLAN (optional)
                                          When using short, unqualified names, DNS queries will
IPv4 Configuration                        attempt to locate the specified host by appending the
IPv6 Configuration                        suffixes listed here in the order shown until a match
DNS Configuration                         found or the list is exhausted.
Custom DNS Suffixes
                                          If no suffixes are specified here, a default suffix li
                                          derived from the local domain name.
```

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-01 management IP address.

2. Download and install the vSphere Client.

> This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 6.5

To download VMware vSphere CLI 6.5, complete the following steps:

1. Click the following link: https://my.vmware.com/en/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5

2. Select your OS and click Download.

3. Save it to your destination folder.

4. Run the VMware-vSphere-CLI.exe

5. Click Next.

6. Accept the terms for the license and click Next.

7. Click Next on the Destination Folder screen.

8. Click Install.

9. Click Finish.

> Install VMware vSphere CLI 6.5 on the management workstation.

10. Log in to VMware ESXi Hosts by Using VMware vSphere Client.

Log in to VMware ESXi Hosts by using VMware vSphere Client

To log in to the `VM-Host-01` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-01` as the host you are trying to connect to: `<<var_vm_host_01_ip>>`.

2. Enter `root` for the user name.

3. Enter the root password.

4. Click Login to connect.

Download Updated Cisco VIC eNIC Drivers

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

The neNIC version is 1.0.6.0-1 and the fNIC version is 1.6.0.34 were used in this configuration.

1. Open a Web browser on the management workstation and navigate to:

https://my.vmware.com/web/vmware/details?downloadGroup=ESXI65U1&productId=614&rPId=18152

https://my.vmware.com/web/vmware/details?productId=614&downloadGroup=DT-ESXI65-CISCO-NENIC-1060

https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI65-CISCO-NENIC-1020&productId=614

https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI60-CISCO-FNIC-16034&productId=491

2. Download the Cisco eNIC and fNIC driver bundle.
3. Open the neNIC driver bundle. This bundle includes the VMware driver bundle which will be uploaded to ESXi hosts.
4. Open the fNIC driver bundle. This bundle includes the VMware driver bundle which will be uploaded to ESXi hosts.
5. 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 neNIC 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:
   a. VMware ESXi 6.5 NIC nenic 1.0.6.0 Driver for Cisco nenic
   b. VMware ESXi 6.0 fnic 1.6.0.34 FC Driver for Cisco
6. Click Open on each and click Yes to upload the file to datastore1.

7. Click the fourth button and select Upload File.

8. 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/ESXi6.5_nenic-
1.0.6.0-offline_bundle-5894048.zip
```

```
esxcli -s <<var_vm_host_ip>> -u root -p <<var_password>> --thumbprint
<host_thumbprint> software vib update -d /vmfs/volumes/datastore1/
fnic_driver_1.6.0.34-offline_bundle-5367434.zip
```

11. Back in the vSphere Client for each host, right click the host and select Reboot.

12. Click Yes and OK to reboot the host.

13. Log back into each host with vSphere Client.

Verify the neNIC driver version installed by entering `vmkload_mod –s nenic` and `vmkload_mod –s fNIC` at the command prompt.

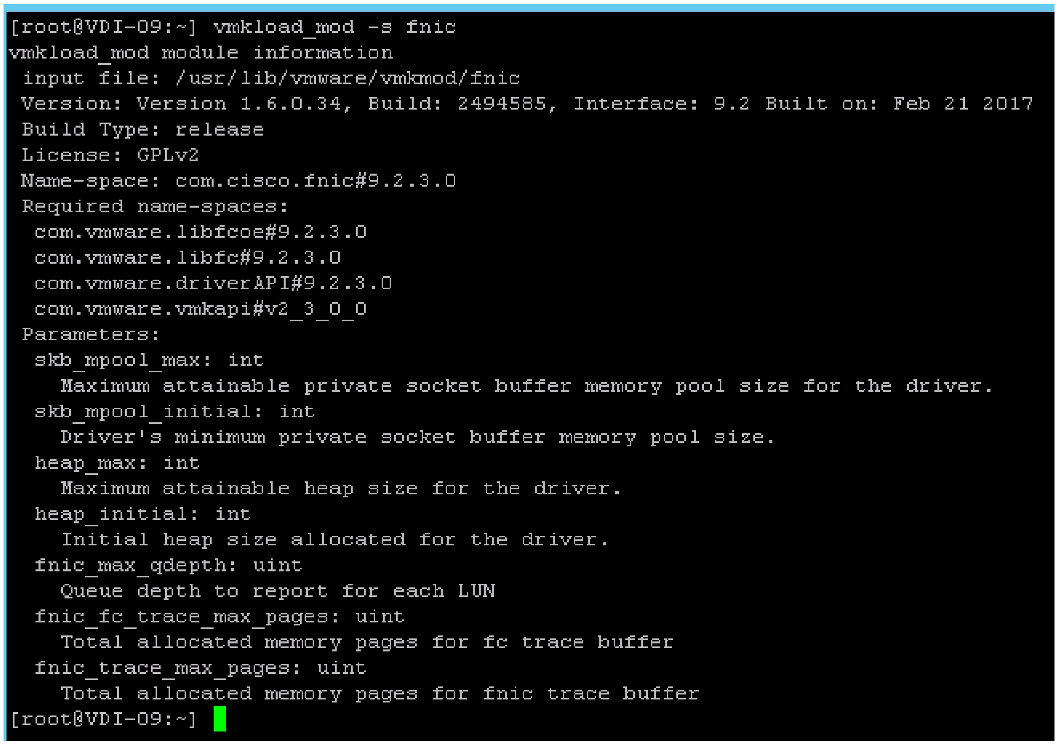

```
[root@VDI-09:~] vmkload_mod -s fnic
vmkload_mod module information
 input file: /usr/lib/vmware/vmkmod/fnic
 Version: Version 1.6.0.34, Build: 2494585, Interface: 9.2 Built on: Feb 21 2017
 Build Type: release
 License: GPLv2
 Name-space: com.cisco.fnic#9.2.3.0
 Required name-spaces:
  com.vmware.libfcoe#9.2.3.0
  com.vmware.libfc#9.2.3.0
  com.vmware.driverAPI#9.2.3.0
  com.vmware.vmkapi#v2_3_0_0
 Parameters:
  skb_mpool_max: int
    Maximum attainable private socket buffer memory pool size for the driver.
  skb_mpool_initial: int
    Driver's minimum private socket buffer memory pool size.
  heap_max: int
    Maximum attainable heap size for the driver.
  heap_initial: int
    Initial heap size allocated for the driver.
  fnic_max_qdepth: uint
    Queue depth to report for each LUN
  fnic_fc_trace_max_pages: uint
    Total allocated memory pages for fc trace buffer
  fnic_trace_max_pages: uint
    Total allocated memory pages for fnic trace buffer
[root@VDI-09:~]
```

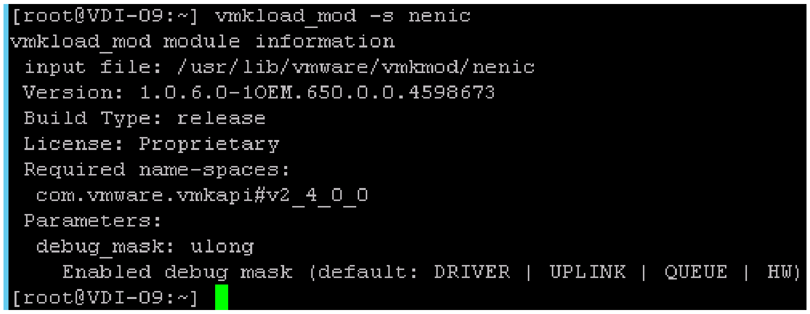

```
[root@VDI-09:~] vmkload_mod -s nenic
vmkload_mod module information
 input file: /usr/lib/vmware/vmkmod/nenic
 Version: 1.0.6.0-1OEM.650.0.0.4598673
 Build Type: release
 License: Proprietary
 Required name-spaces:
  com.vmware.vmkapi#v2_4_0_0
 Parameters:
  debug_mask: ulong
    Enabled debug mask (default: DRIVER | UPLINK | QUEUE | HW)
[root@VDI-09:~]
```

Install and Configure VMware vCenter Appliance

Log in to the VM-Host-01 ESXi host by using the VMware vSphere Client and complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of VM-Host-01 as the host you are trying to connect to.

2. Enter root for the user name.

3. Enter the root password.

4. Click Login to connect.

To build the VMWare vCenter VM, complete the following steps:

1. From the vSphere 6 download page on the VMware Web site, download the vCenter ISO file for the vCenter Server appliance onto your system.

2. Mount the vSphere ISO file via Windows Explorer and navigate to the folder vcsa-ui-installer/win32 and click installer file to start the VCSA Appliance from Installer.

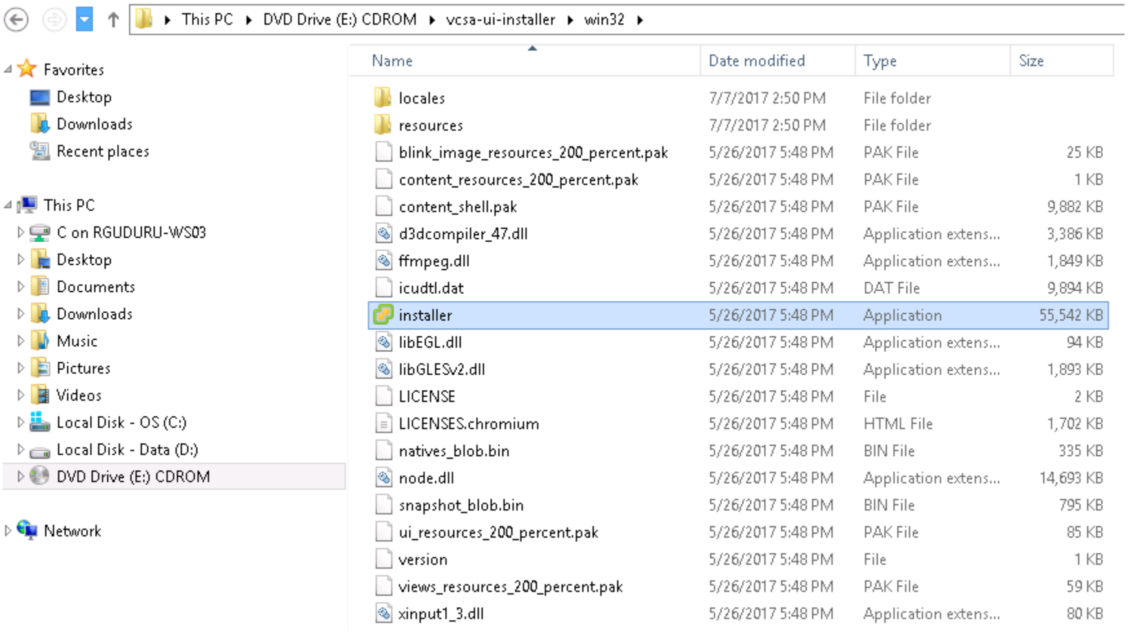

Name	Date modified	Type	Size
locales	7/7/2017 2:50 PM	File folder	
resources	7/7/2017 2:50 PM	File folder	
blink_image_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	25 KB
content_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	1 KB
content_shell.pak	5/26/2017 5:48 PM	PAK File	9,882 KB
d3dcompiler_47.dll	5/26/2017 5:48 PM	Application extens...	3,386 KB
ffmpeg.dll	5/26/2017 5:48 PM	Application extens...	1,849 KB
icudtl.dat	5/26/2017 5:48 PM	DAT File	9,894 KB
installer	5/26/2017 5:48 PM	Application	55,542 KB
libEGL.dll	5/26/2017 5:48 PM	Application extens...	94 KB
libGLESv2.dll	5/26/2017 5:48 PM	Application extens...	1,893 KB
LICENSE	5/26/2017 5:48 PM	File	2 KB
LICENSES.chromium	5/26/2017 5:48 PM	HTML File	1,702 KB
natives_blob.bin	5/26/2017 5:48 PM	BIN File	335 KB
node.dll	5/26/2017 5:48 PM	Application extens...	14,693 KB
snapshot_blob.bin	5/26/2017 5:48 PM	BIN File	795 KB
ui_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	85 KB
version	5/26/2017 5:48 PM	File	1 KB
views_resources_200_percent.pak	5/26/2017 5:48 PM	PAK File	59 KB
xinput1_3.dll	5/26/2017 5:48 PM	Application extens...	80 KB

A browser will open with an option to Install.

3. Click Install.

Install – Stage 1: Deploy vCenter Server with Embedded Platform Services Controller opens.

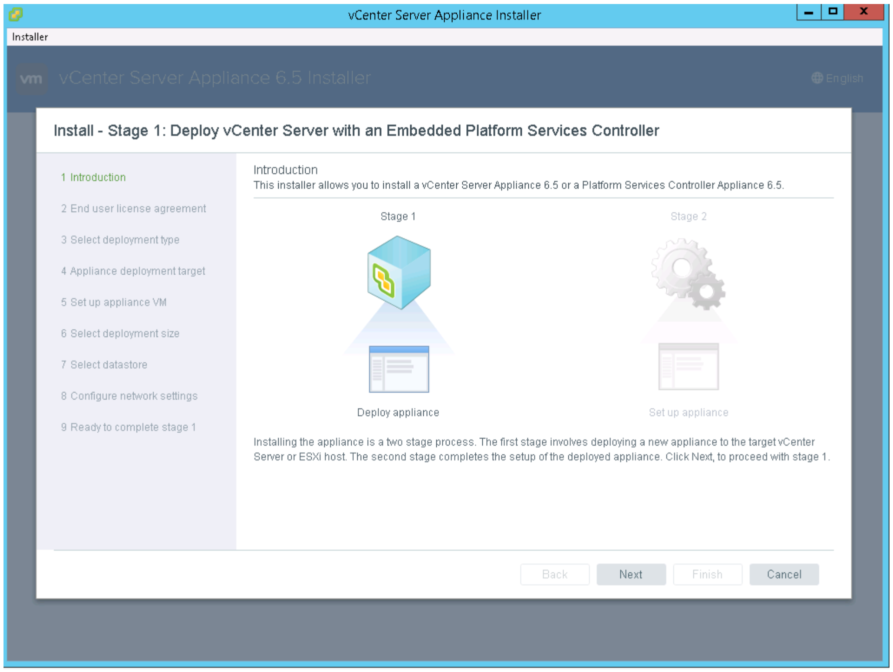
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Stage 1
Stage 2
Deploy appliance
Set up appliance
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Introduction

This installer allows you to install a vCenter Server Appliance 6.5 or a Platform Services Controller Appliance 6.5.

Installing the appliance is a two stage process. The first stage involves deploying a new appliance to the target vCenter Server or ESXi host. The second stage completes the setup of the deployed appliance. Click Next, to proceed with stage 1.

4. Follow the onscreen prompts. Accept EULA.

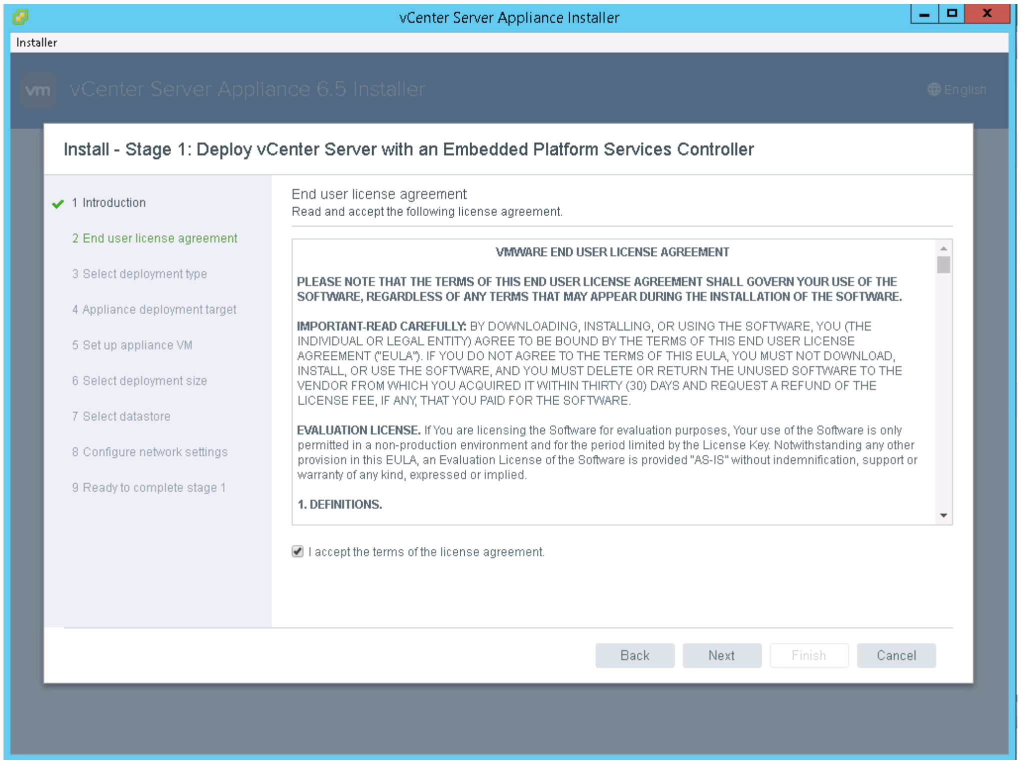
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
✓ 1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

End user license agreement

Read and accept the following license agreement.

VMWARE END USER LICENSE AGREEMENT

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

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE VENDOR FROM WHICH YOU ACQUIRED IT WITHIN THIRTY (30) DAYS AND REQUEST A REFUND OF THE LICENSE FEE, IF ANY, THAT YOU PAID FOR THE SOFTWARE.

EVALUATION LICENSE. If You are licensing the Software for evaluation purposes, Your use of the Software is only permitted in a non-production environment and for the period limited by the License Key. Notwithstanding any other provision in this EULA, an Evaluation License of the Software is provided "AS-IS" without indemnification, support or warranty of any kind, expressed or implied.

1. DEFINITIONS.

☑ I accept the terms of the license agreement.

5. Deploy the appliances. Select Embedded Platform Services Controller.

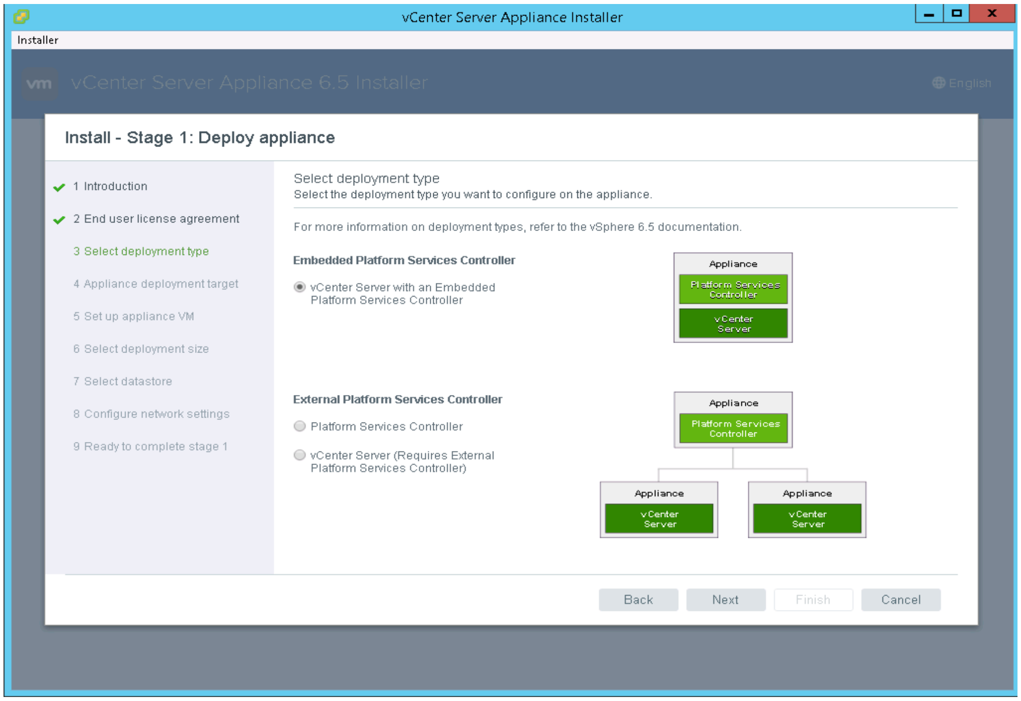
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Appliance
Platform Services Controller
vCenter Server
Back
Next
Finish
Cancel

Install - Stage 1: Deploy appliance

Select deployment type

Select the deployment type you want to configure on the appliance.

For more information on deployment types, refer to the vSphere 6.5 documentation.

Embedded Platform Services Controller

● vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

○ Platform Services Controller

○ vCenter Server (Requires External Platform Services Controller)

6. Enter the IP of the ESXi host the vCenter Appliance will reside. Click Next. Provide Host IP or FQDN and User Name, Password Credentials of the Host to Connect.

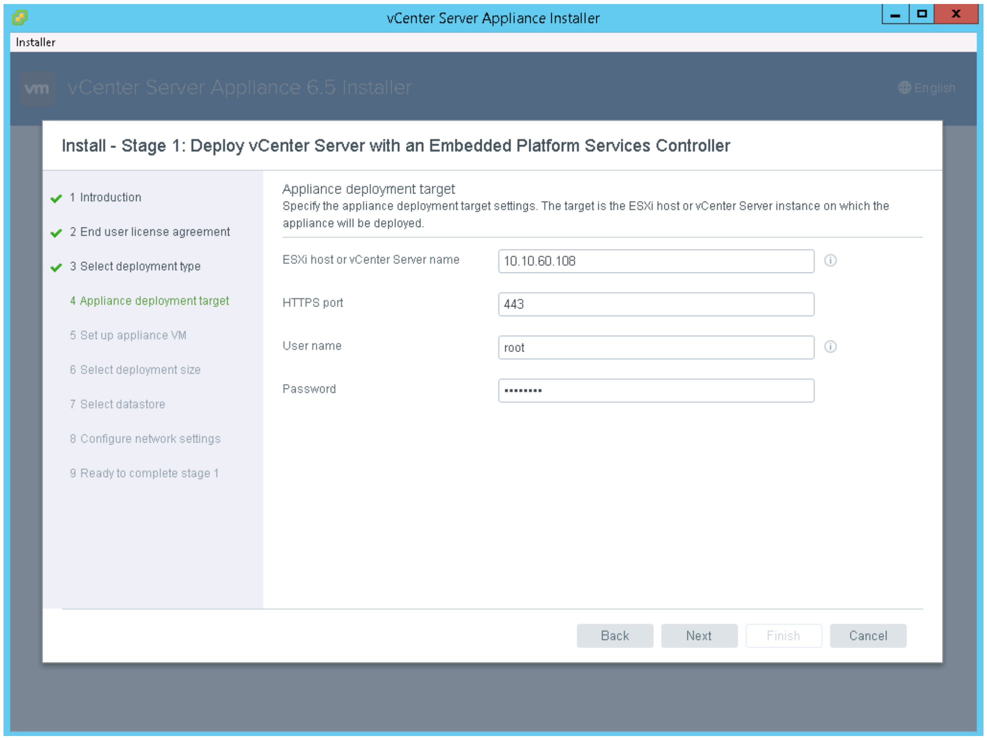
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

ESXi host or vCenter Server name	10.10.60.108
HTTPS port	443
User name	root
Password	••••••••

7. Click Yes to accept Certificate Warning.

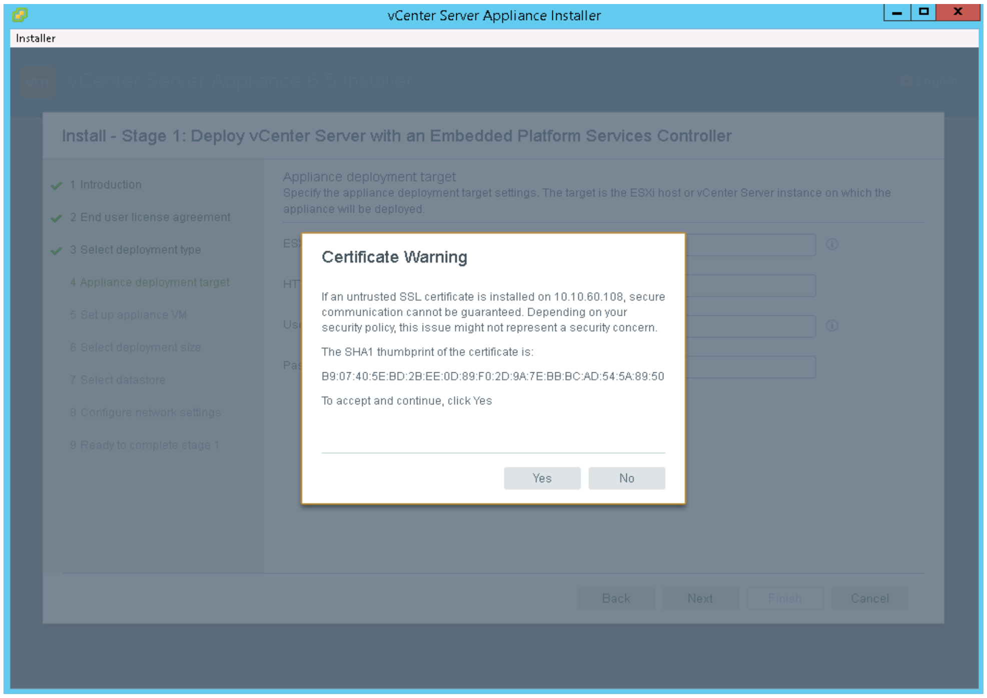
vCenter Server Appliance Installer
Installer
Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

Certificate Warning

If an untrusted SSL certificate is installed on 10.10.60.108, secure communication cannot be guaranteed. Depending on your security policy, this issue might not represent a security concern.

The SHA1 thumbprint of the certificate is:

B9:07:40:5E:BD:2B:EE:0D:89:F0:2D:9A:7E:BB:BC:AD:54:5A:89:50

To accept and continue, click Yes

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

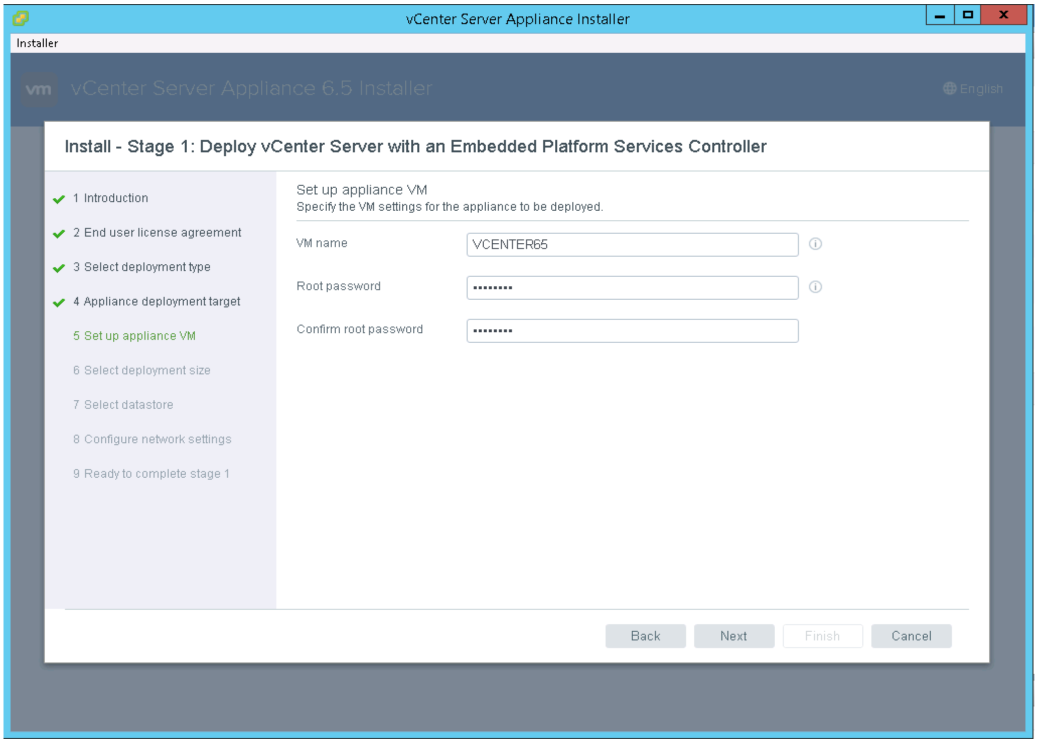
vCenter Server Appliance Installer
Installer
vm
vCenter Server Appliance 6.5 Installer
English
Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller
1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1
Back
Next
Finish
Cancel

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name: VCENTER65

Root password: ••••••••

Confirm root password: ••••••••

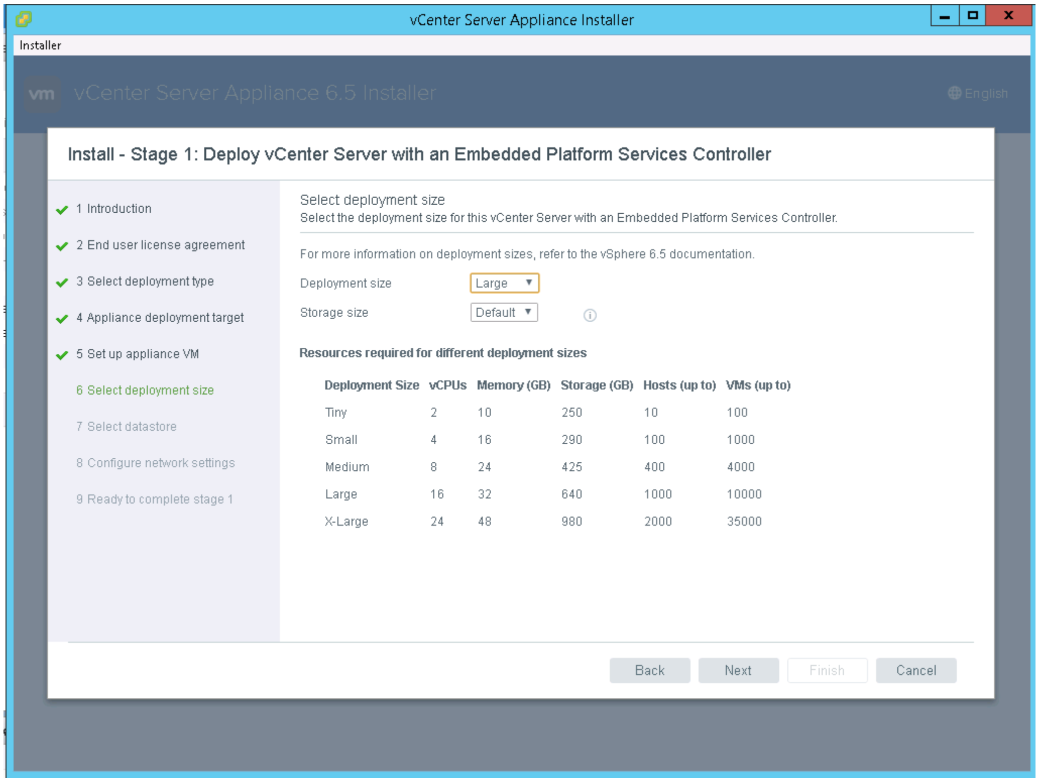
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

- ✓ 1 Introduction
- ✓ 2 End user license agreement
- ✓ 3 Select deployment type
- ✓ 4 Appliance deployment target
- ✓ 5 Set up appliance VM
- 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings
- 9 Ready to complete stage 1

Select deployment size

Select the deployment size for this vCenter Server with an Embedded Platform Services Controller.

For more information on deployment sizes, refer to the vSphere 6.5 documentation.

Deployment size: Large

Storage size: Default

Resources required for different deployment sizes

Deployment Size	vCPUs	Memory (GB)	Storage (GB)	Hosts (up to)	VMs (up to)
Tiny	2	10	250	10	100
Small	4	16	290	100	1000
Medium	8	24	425	400	4000
Large	16	32	640	1000	10000
X-Large	24	48	980	2000	35000

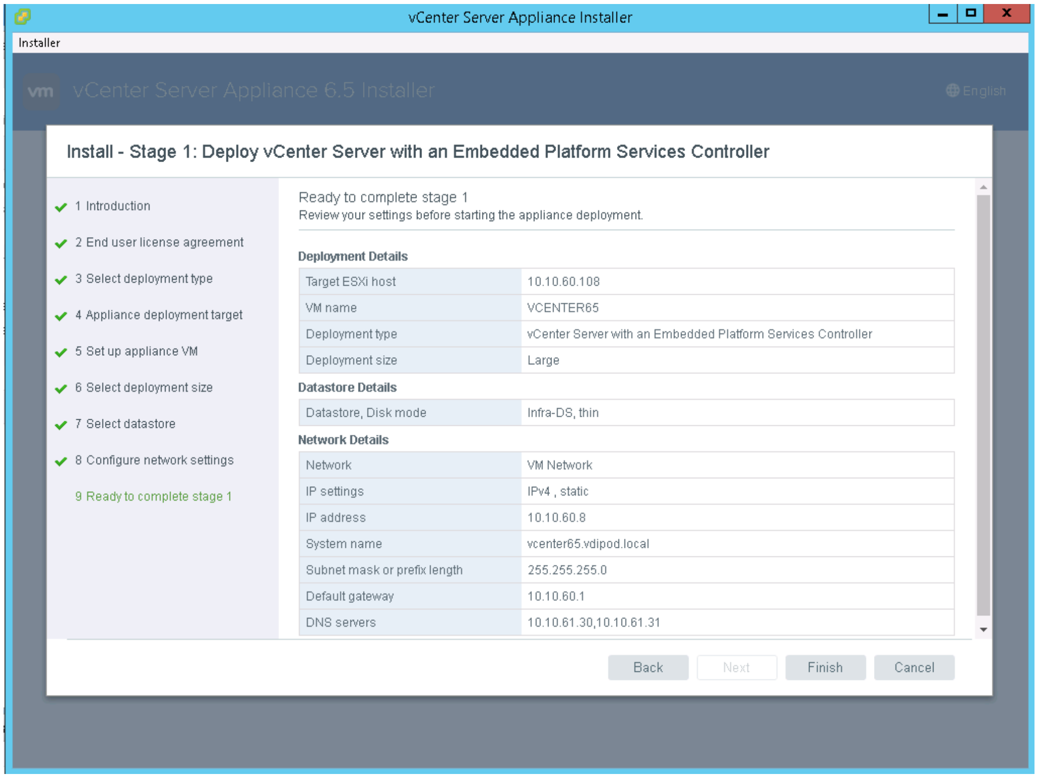
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
English
Back
Next
Finish
Cancel

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

- ✓ 1 Introduction
- ✓ 2 End user license agreement
- ✓ 3 Select deployment type
- ✓ 4 Appliance deployment target
- ✓ 5 Set up appliance VM
- ✓ 6 Select deployment size
- ✓ 7 Select datastore
- ✓ 8 Configure network settings
- 9 Ready to complete stage 1

Ready to complete stage 1

Review your settings before starting the appliance deployment.

Deployment Details

Target ESXi host	10.10.60.108
VM name	VCENTER65
Deployment type	vCenter Server with an Embedded Platform Services Controller
Deployment size	Large

Datastore Details

Datastore, Disk mode	Infra-DS, thin

Network Details

Network	VM Network
IP settings	IPv4 , static
IP address	10.10.60.8
System name	vcenter65.vdipod.local
Subnet mask or prefix length	255.255.255.0
Default gateway	10.10.60.1
DNS servers	10.10.61.30,10.10.61.31

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

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

11. Provide Single Sign On Password and Site Name Credentials.

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

13. Select the Data store.

In our study we used the embedded PostgreSQL database.

14. Select Use an embedded database (PostgreSQL).

15. Enter Network Settings for appliance.

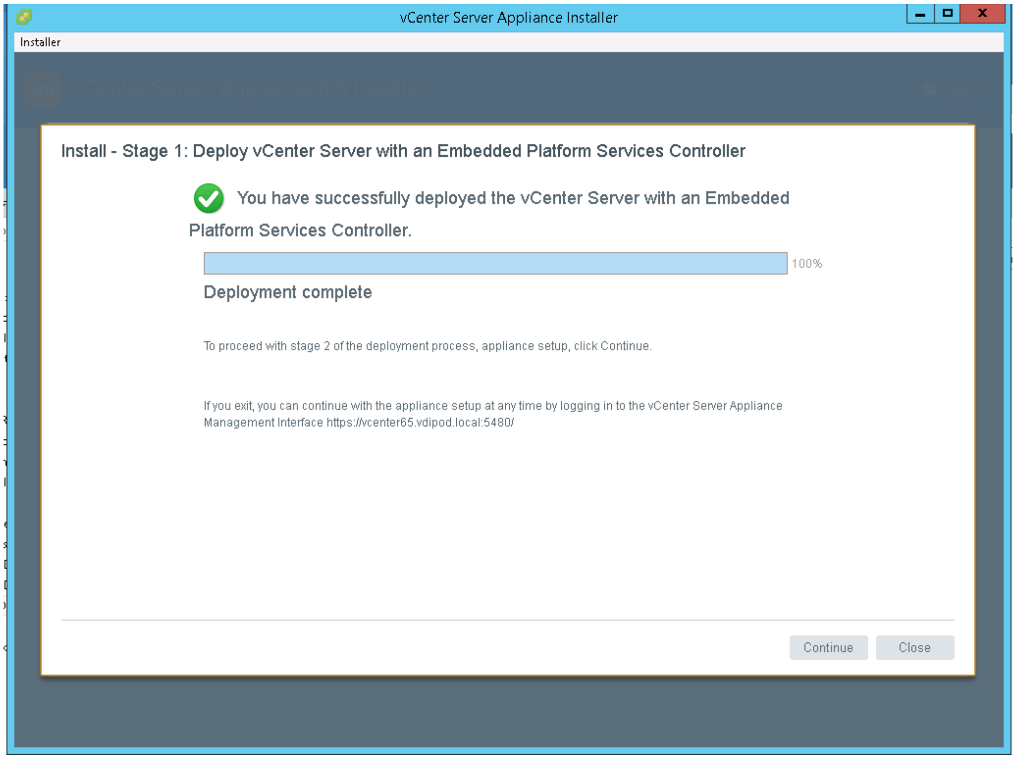

vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer
English

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

You have successfully deployed the vCenter Server with an Embedded Platform Services Controller.

100%

Deployment complete

To proceed with stage 2 of the deployment process, appliance setup, click Continue.

If you exit, you can continue with the appliance setup at any time by logging in to the vCenter Server Appliance Management Interface https://vcenter65.vdipod.local:5480/

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.

16. Provide the Necessary Network Gateways and DNS Server Information. Review the install settings and click Finish.

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

18. Configure the Embedded Server Appliance.

Install – Stage 2: Set Up vCenter Server Appliance with an Embedded PSC opens.

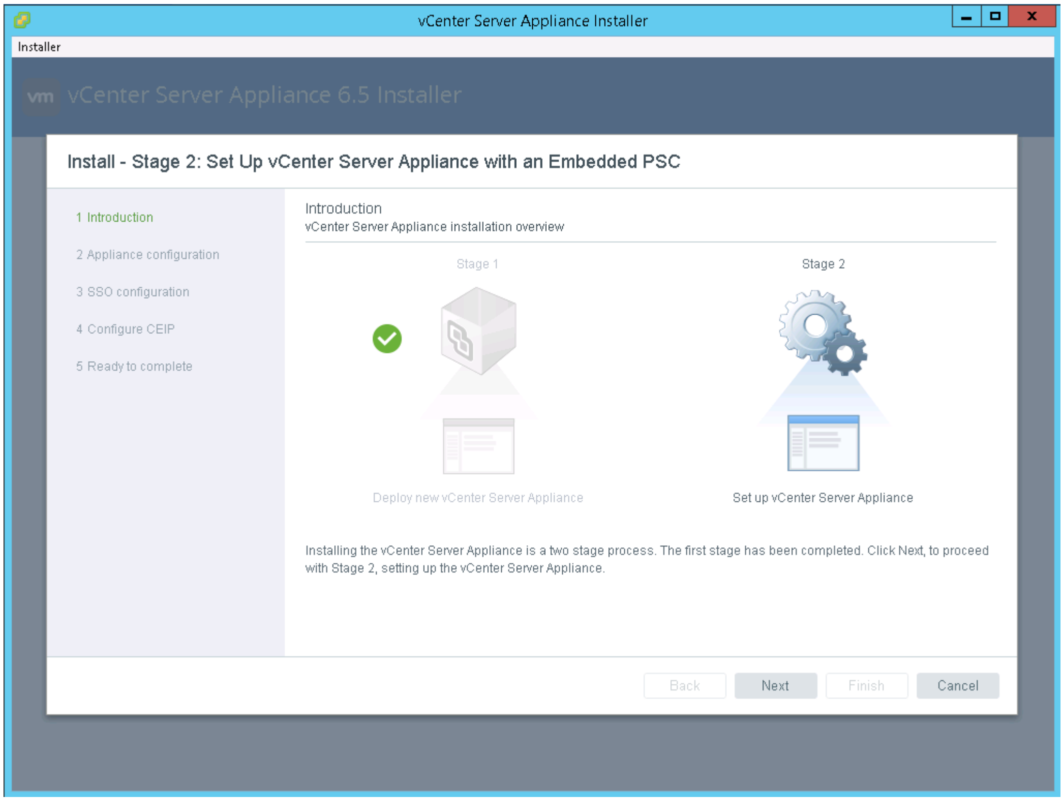
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
Stage 1
Stage 2
Deploy new vCenter Server Appliance
Set up vCenter Server Appliance
Back
Next
Finish
Cancel

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Introduction

vCenter Server Appliance installation overview

Installing the vCenter Server Appliance is a two stage process. The first stage has been completed. Click Next, to proceed with Stage 2, setting up the vCenter Server Appliance.

19. Sync the time for appliance.

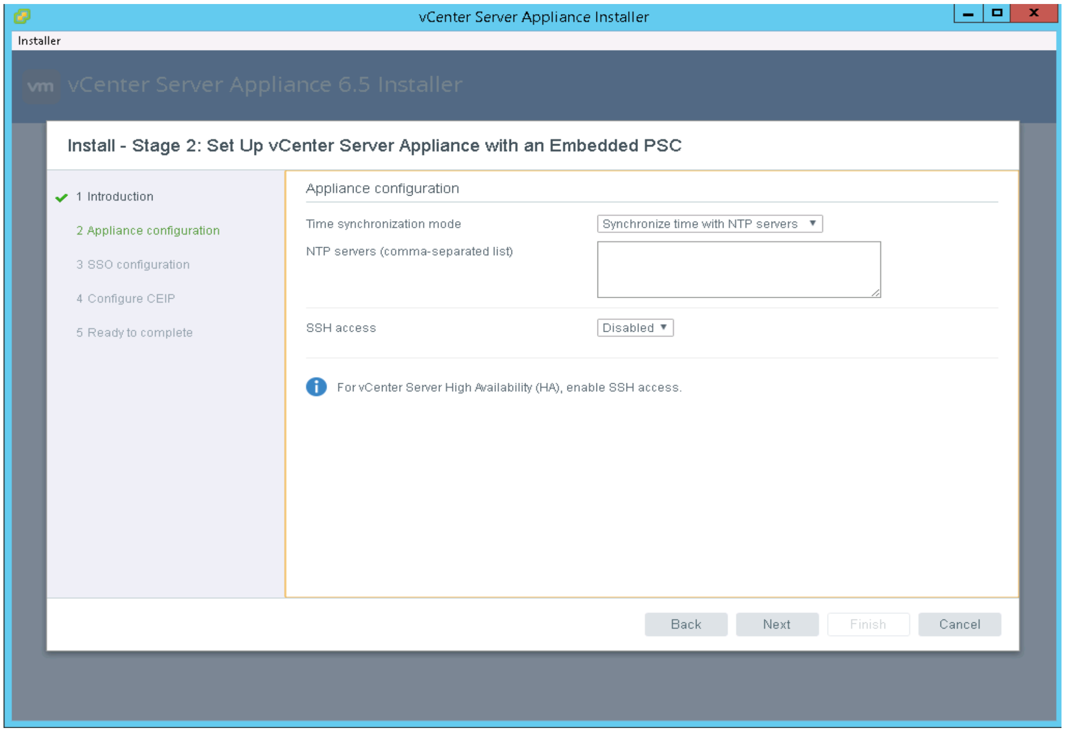
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
✓ 1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
Back
Next
Finish
Cancel

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Appliance configuration

Time synchronization mode: Synchronize time with NTP servers ▼

NTP servers (comma-separated list):

SSH access: Disabled ▼

For vCenter Server High Availability (HA), enable SSH access.

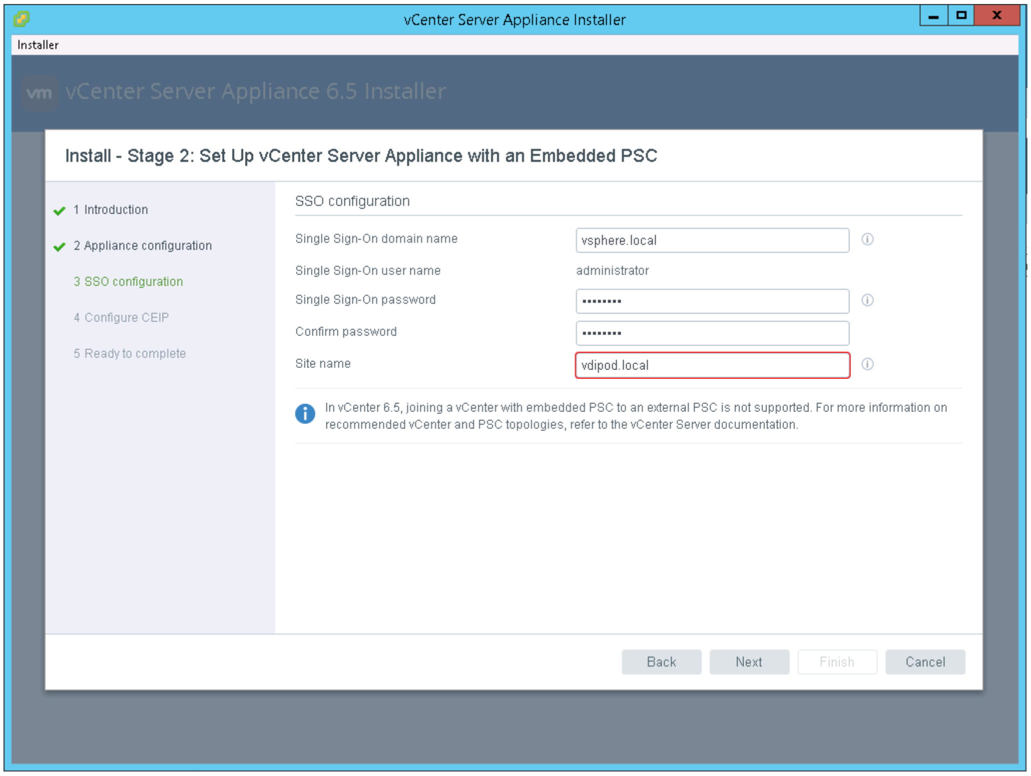
vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
Back
Next
Finish
Cancel

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

SSO configuration

Single Sign-On domain name	vsphere.local
Single Sign-On user name	administrator
Single Sign-On password	••••••••
Confirm password	••••••••
Site name	vdipod.local

In vCenter 6.5, joining a vCenter with embedded PSC to an external PSC is not supported. For more information on recommended vCenter and PSC topologies, refer to the vCenter Server documentation.

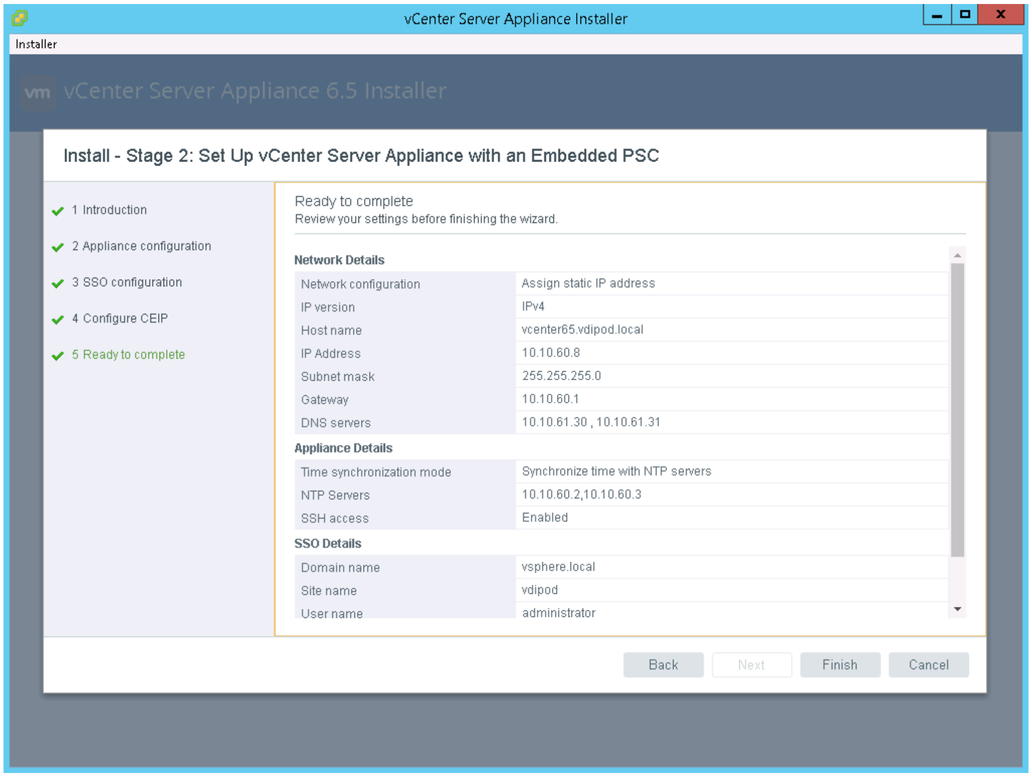
vCenter Server Appliance Installer
Installer
vCenter Server Appliance 6.5 Installer
1 Introduction
2 Appliance configuration
3 SSO configuration
4 Configure CEIP
5 Ready to complete
Back
Next
Finish
Cancel

Install - Stage 2: Set Up vCenter Server Appliance with an Embedded PSC

Ready to complete

Review your settings before finishing the wizard.

Network Details	
Network configuration	Assign static IP address
IP version	IPv4
Host name	vcenter65.vdipod.local
IP Address	10.10.60.8
Subnet mask	255.255.255.0
Gateway	10.10.60.1
DNS servers	10.10.61.30 , 10.10.61.31
Appliance Details	
Time synchronization mode	Synchronize time with NTP servers
NTP Servers	10.10.60.2,10.10.60.3
SSH access	Enabled
SSO Details	
Domain name	vsphere.local
Site name	vdipod
User name	administrator

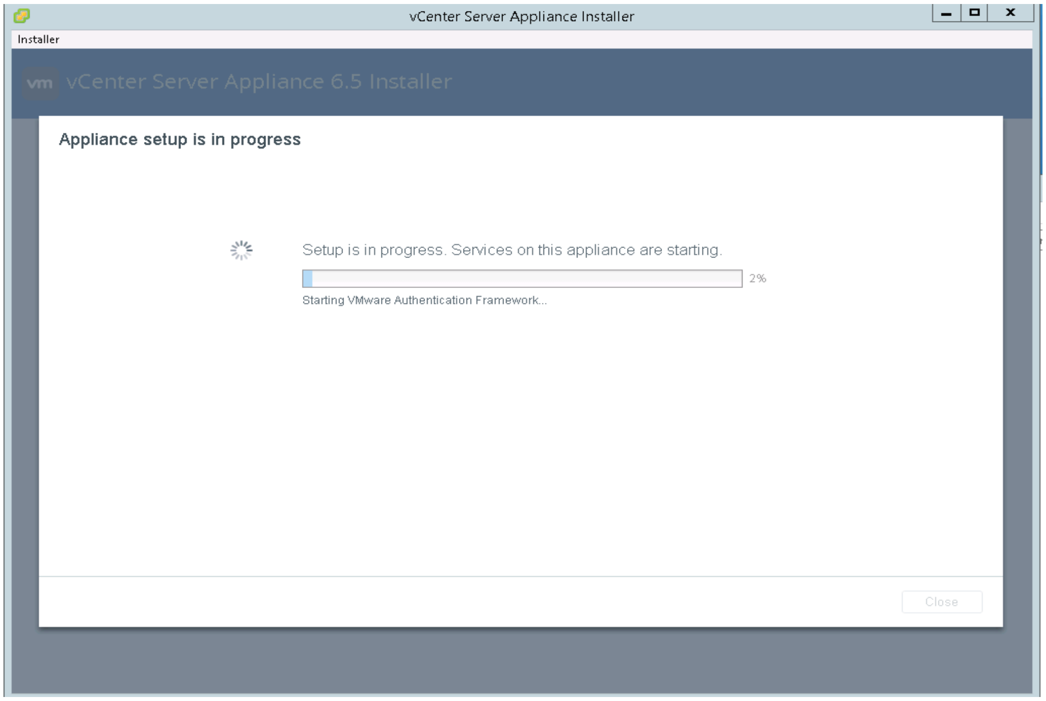
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
Close

Appliance setup is in progress

Setup is in progress. Services on this appliance are starting.

2%

Starting VMware Authentication Framework...

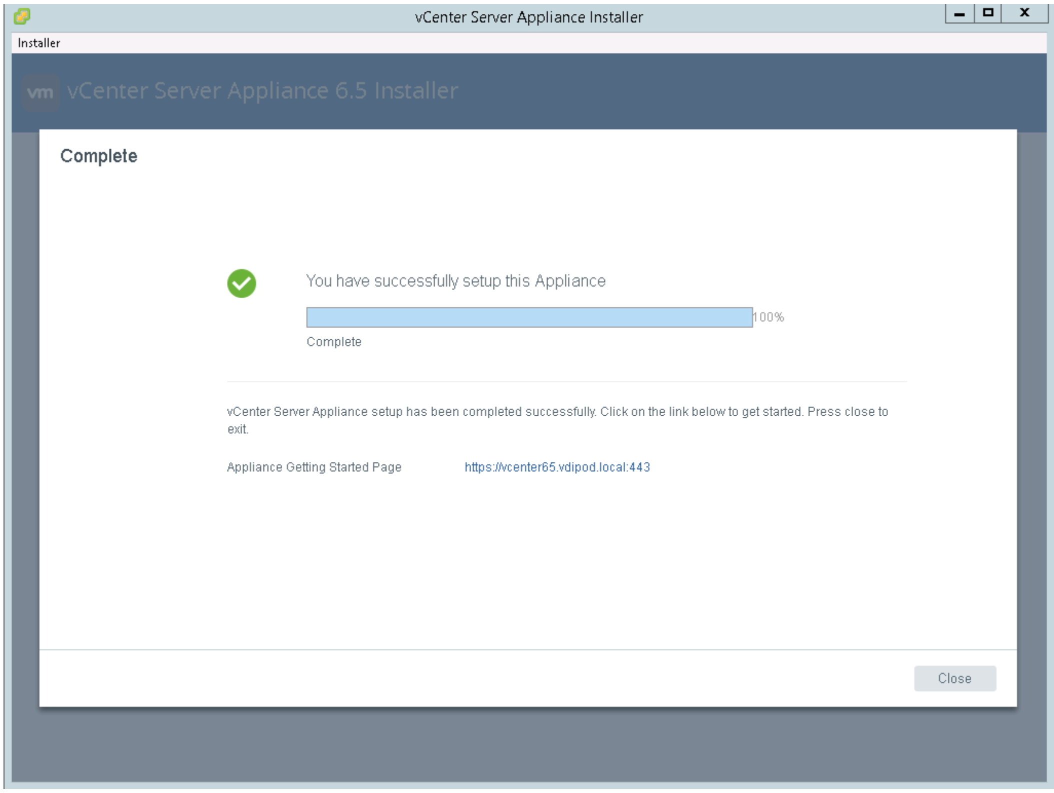
vCenter Server Appliance Installer
Installer
vm vCenter Server Appliance 6.5 Installer
Close

Complete

You have successfully setup this Appliance

100%

Complete

vCenter Server Appliance setup has been completed successfully. Click on the link below to get started. Press close to exit.

Appliance Getting Started Page https://vcenter65.vdipod.local:443

20. Login in to vCenter Appliances Web GUI.

21. Login using IP Address of the Appliance and Download vSphere

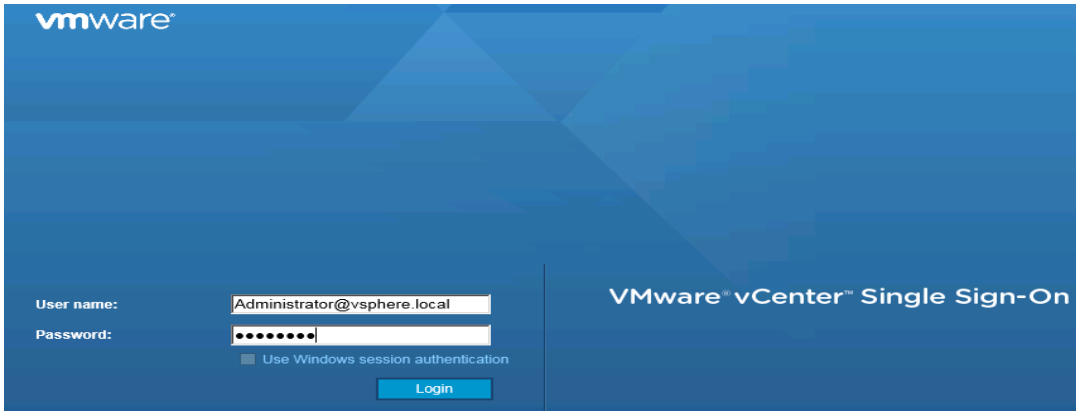
vmware
VMware® vCenter™ Single Sign-On
Use Windows session authentication
Login

User name: Administrator@vsphere.local

Password: ••••••••

22. Log into the vSphere Web Client.

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

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

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

26. Click the vCenter link on the left panel.

27. Login vSphere Web GUI.

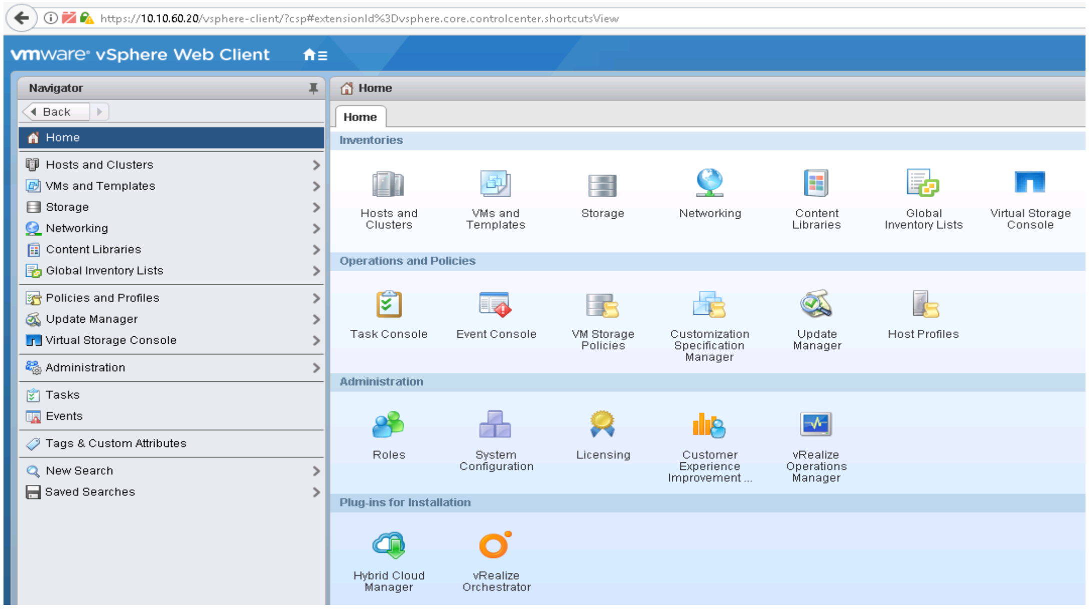

https://10.10.60.20/vsphere-client/?csp#extensionId%3Dvsphere.core.controlcenter.shortcutsView

28. Click the Datacenters link on the left panel.

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

The screenshot below shows an example of a VDI-DC Data Center.

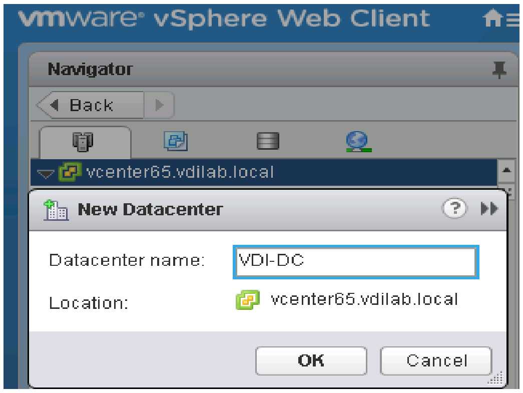

New Datacenter

Datacenter name: VDI-DC

Location: vcenter65.vdilab.local

30. Type VDI-DC as the Datacenter name.

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

32. Create a Cluster.

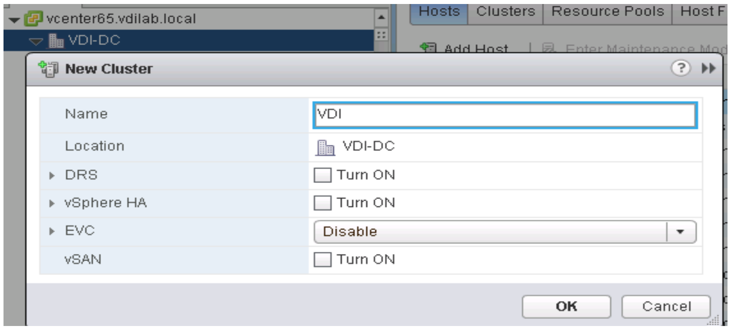

New Cluster

Name	VDI
Location	VDI-DC
DRS	□ Turn ON
vSphere HA	□ Turn ON
EVC	Disable
vSAN	□ Turn ON

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

34. Name the cluster VDI-CL.

35. Select DRS. Retain the default values.

36. Select vSphere HA. Retain the default values. Configure Cluster Specific Setting.

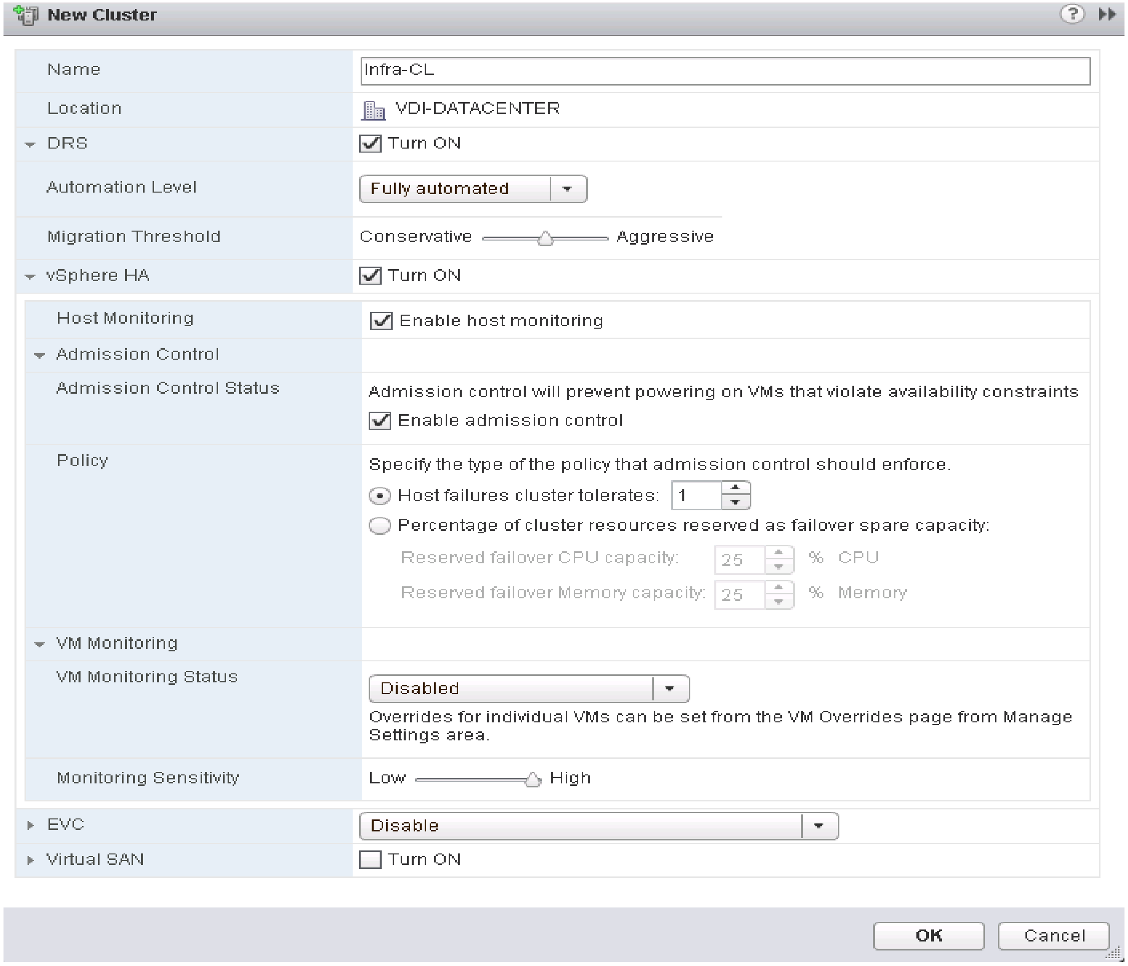

New Cluster

Name	Infra-CL
Location	VDI-DATACENTER
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 M5 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

37. Click OK to create the new cluster.

38. Click VDI-DC in the left pane.

39. Add a ESXi Host.

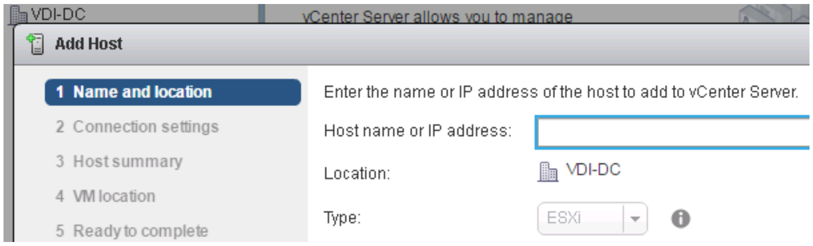

Enter the name or IP address of the host to add to vCenter Server.

Host name or IP address:

Location: VDI-DC

Type: ESXi

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

41. Type the host IP address and click Next.

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

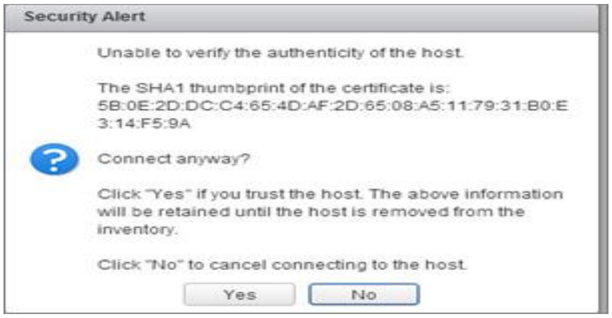

Security Alert

Unable to verify the authenticity of the host.

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

43. Click Yes to accept the certificate.

44. Review the host details and click Next to continue.

45. Assign a license and click Next to continue.

46. Click Next to continue.

47. Click Next to continue.

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

49. Repeat this for the other hosts and clusters.

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

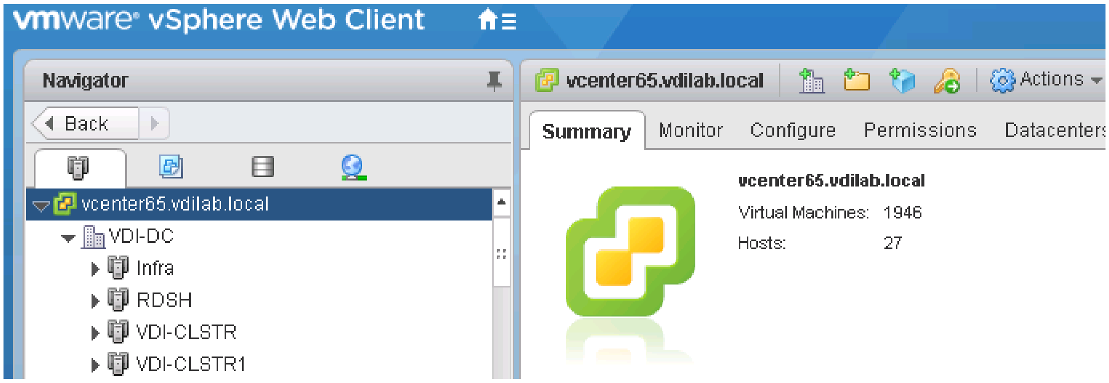

vcenter65.vdilab.local

Virtual Machines: 1946

Hosts: 27

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:

Table 45 Software Infrastructure Configuration

Configuration	Operating System	Virtual CPU	Memory	Disk Size	Network
vCenter Server Appliance	VCSA- SUSE Linux	16	32	640	MGMT-VLAN
Active Directory Domain Controllers/DHCP/DNS (2)	Microsoft Windows 2016	4	12	60	Infra-VLAN
VMware Horizon Connection Server	Microsoft Windows 2016	4	12	60	Infra-VLAN
VMware Horizon Composer Server-1	Microsoft Windows 2016	4	12	40	Infra-VLAN
VMware Horizon Replica Server (2)	Microsoft Windows 2016	4	8	40	Infra-VLAN
Microsoft SQL Server 2016	Microsoft Windows 2016	4	12	60	Infra-VLAN
KMS License Server	Microsoft Windows 2016	4	8	40	Infra-VLAN

Preparing the Master Image

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.

The master target for RDSH server roles was configured server 2016 and target master image configured for VDI VMs with Windows 10 64 bit OS as listed in the table below:

Table 46 Master Image Configuration

Configuration	Operating System	Virtual CPU	Memory	Disk Size	Network	Additional Software
RDSH Virtual Machine	Microsoft Windows Server 2016	6	24 GB	40	VMXNET3 VDI-VLAN	Microsoft Office 2016. Login VSI 4.1.25.6 (Knowledge Worker workload)
VDI Virtual Machine (Instant Clone Pool)	Microsoft Windows 10 64-Bit LTSB Version 1607	2	2 GB (Reserved)	32	VMXNET3 VDI-DVS (vDS)	Microsoft Office 2016. Login VSI 4.1.25.6 (Knowledge Worker workload)
VDI Virtual Machine(Full Clone Pool)	Microsoft Windows 10 64-Bit LTSB Version 1607	2	2 GB (Reserved)	32	VMXNET3 VDI-DVS (vDS)	Microsoft Office 2016. Login VSI 4.1.25.6 (Knowledge Worker

Configuration	Operating System	Virtual CPU	Memory	Disk Size	Network	Additional Software
						workload)

Installing and Configuring VMware Horizon Environment

This section details the installation of the VMware core components of the Horizon Connection Server and Replica Servers. This CVD installs 1 VMware Horizon Connection server and 2 VMware Horizon Replica Servers to support both remote desktop server hosted sessions (RDSH), non-persistent virtual desktops (VDI) based on the best practices from VMware. For information about sizing limits, see VMware Horizon View 6 sizing limits and recommendations.

The prerequisites for installing the Horizon Connection server or Composer server is to have Windows 2008/12 to 2016 servers ready. In this study, we have used Windows 2016 server for Horizon Connection Server, Replica Servers, and Composer Server.

VMware Horizon Connection Server Configuration

To configure the VMware Horizon Connection Server, complete the following steps:

1. Download the Horizon Connection server installer from VMware and click Install on the Connection Server Windows Server Image. In this study, we used version Connection Server 7.3.1 build.6760913. For the download, see Download VMware Horizon 7.3.1 Standard.

2. Click the Connection Server installer.

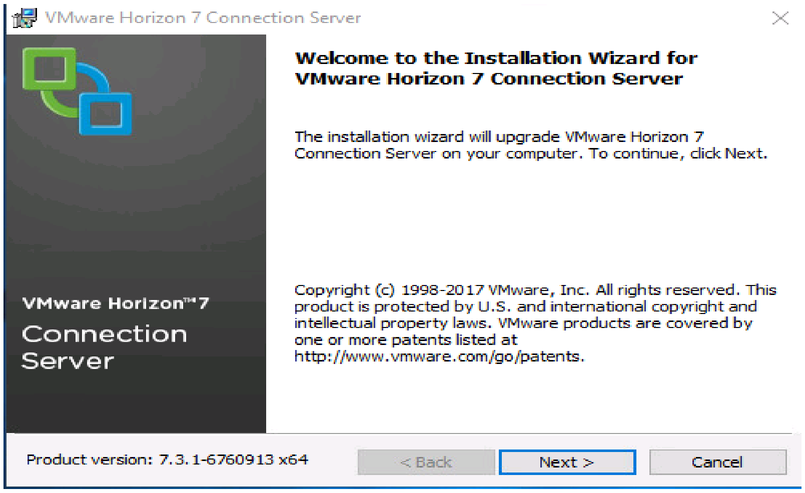

VMware Horizon 7 Connection Server

Welcome to the Installation Wizard for VMware Horizon 7 Connection Server

The installation wizard will upgrade VMware Horizon 7 Connection Server on your computer. To continue, click Next.

Copyright (c) 1998-2017 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 7.3.1-6760913 x64

3. Click Next.

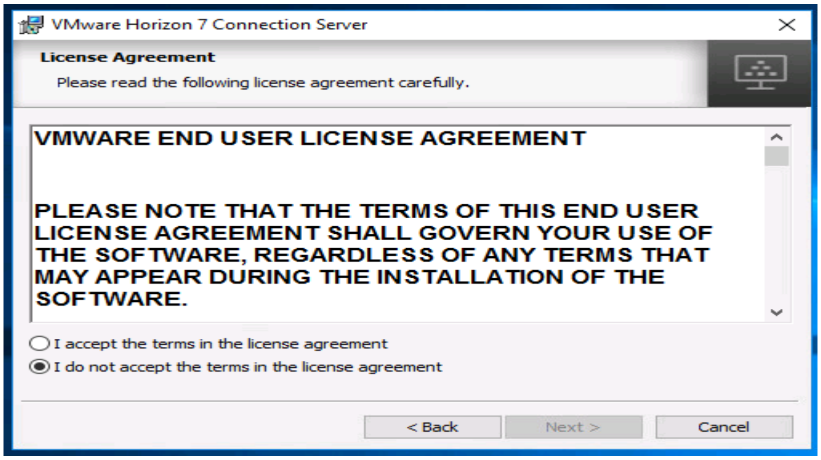

VMware Horizon 7 Connection Server

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

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

○ I accept the terms in the license agreement

● I do not accept the terms in the license agreement

4. Accept the terms in the License Agreement.

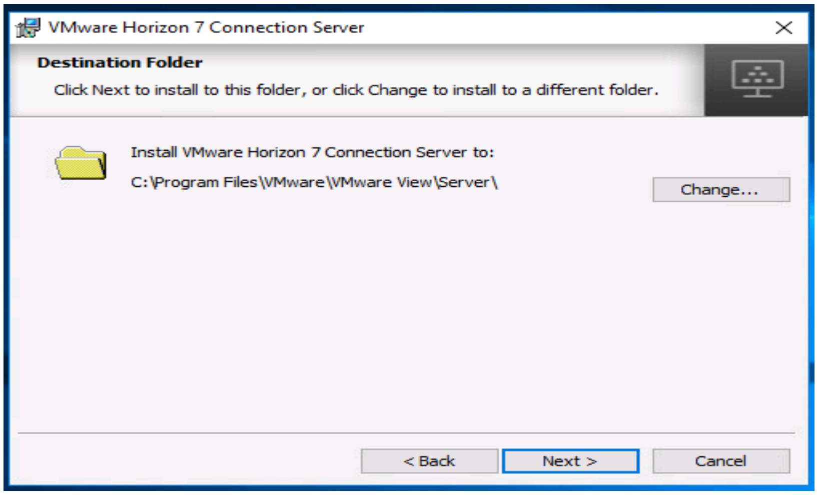

VMware Horizon 7 Connection Server

Destination Folder

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

Install VMware Horizon 7 Connection Server to:

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

5. Click Next.

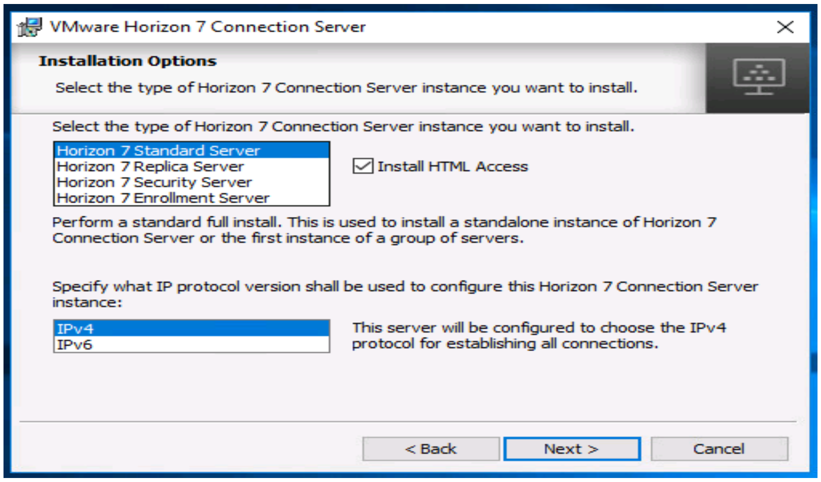

VMware Horizon 7 Connection Server

Installation Options

Select the type of Horizon 7 Connection Server instance you want to install.

Select the type of Horizon 7 Connection Server instance you want to install.

- Horizon 7 Standard Server
- Horizon 7 Replica Server
- Horizon 7 Security Server
- Horizon 7 Enrollment Server

☑ Install HTML Access

Perform a standard full install. This is used to install a standalone instance of Horizon 7 Connection Server or the first instance of a group of servers.

Specify what IP protocol version shall be used to configure this Horizon 7 Connection Server instance:

- IPv4
- IPv6

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

6. Select the Standard Server.

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

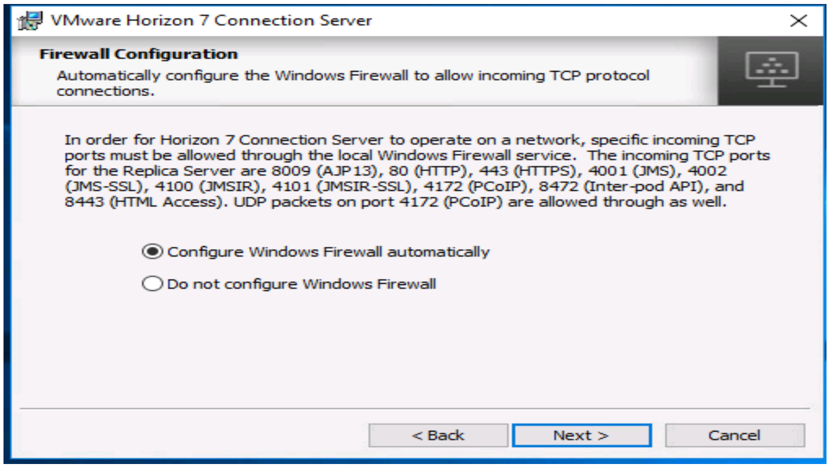

VMware Horizon 7 Connection Server

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for Horizon 7 Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Replica Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), 4002 (JMS-SSL), 4100 (JMSIR), 4101 (JMSIR-SSL), 4172 (PCoIP), 8472 (Inter-pod API), and 8443 (HTML Access). UDP packets on port 4172 (PCoIP) are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

7. Configure the Windows firewall automatically.

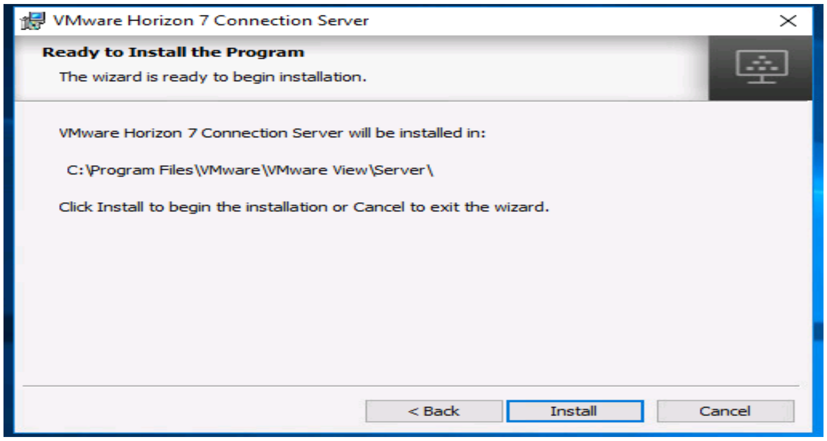

VMware Horizon 7 Connection Server

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon 7 Connection Server will be installed in:

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

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

8. Click Install.

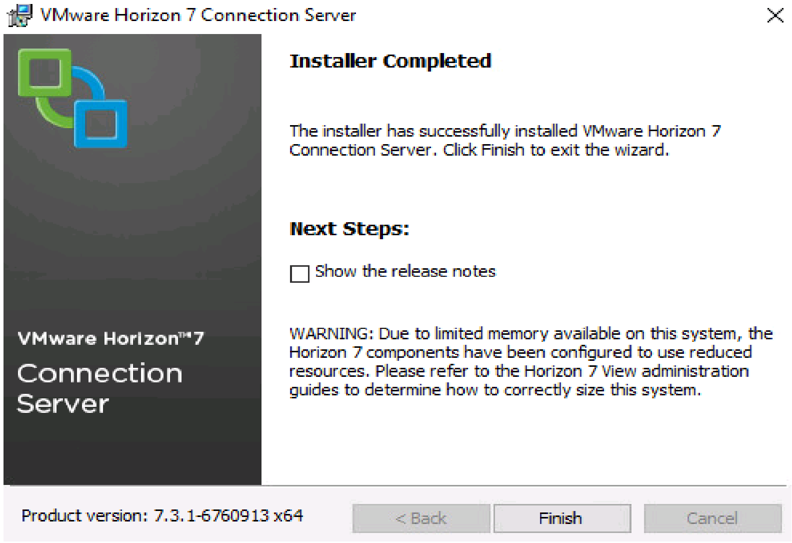

VMware Horizon 7 Connection Server

Installer Completed

The installer has successfully installed VMware Horizon 7 Connection Server. Click Finish to exit the wizard.

Next Steps:

□ Show the release notes

WARNING: Due to limited memory available on this system, the Horizon 7 components have been configured to use reduced resources. Please refer to the Horizon 7 View administration guides to determine how to correctly size this system.

Product version: 7.3.1-6760913 x64

9. Click Finish to complete the Horizon Connection Server installation.

Horizon VMware Replica Server Creation

To install Horizon Replica Server and additional Replica servers, complete the following steps:

1. Follow the steps shown in section VMware Horizon Connection Server Configuration. During the Horizon Replica Server installation, select the Replica Server option in order to configure this server as a Replica Server (shown below) and complete all other steps shown above.

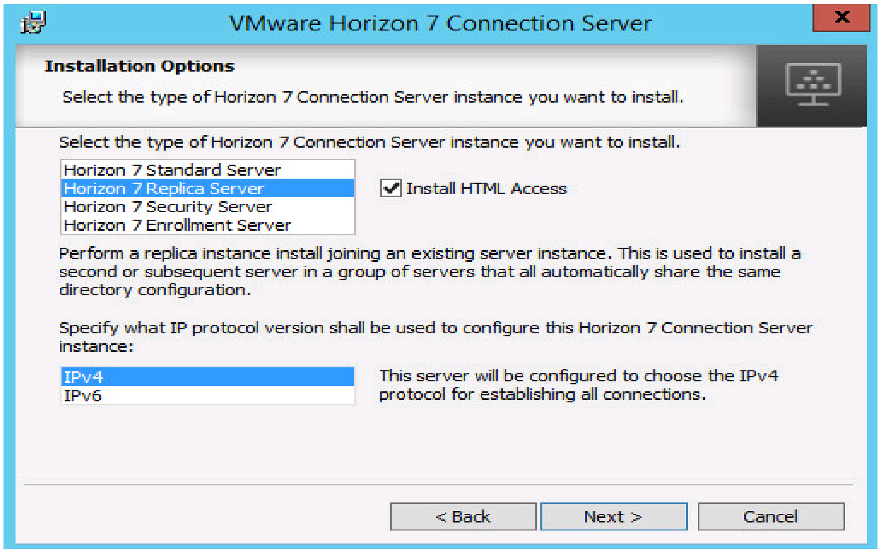

VMware Horizon 7 Connection Server

Installation Options

Select the type of Horizon 7 Connection Server instance you want to install.

Select the type of Horizon 7 Connection Server instance you want to install.

Horizon 7 Standard Server
Horizon 7 Replica Server (selected)
Horizon 7 Security Server
Horizon 7 Enrollment Server

☑ Install HTML Access

Perform a replica instance install joining an existing server instance. This is used to install a second or subsequent server in a group of servers that all automatically share the same directory configuration.

Specify what IP protocol version shall be used to configure this Horizon 7 Connection Server instance:

IPv4 (selected)
IPv6

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

2. Click Next and follow the steps (shown for Horizon Standard server) to complete installing additional Horizon Replica Servers.

Install VMware Horizon Composer Server

To install the VMware Horizon Composer Server, complete the following steps:

1. Download the Composer installer from VMware and click Install on the Composer Windows server Image. In this study, we used Composer 7.3.1 build.6744335.

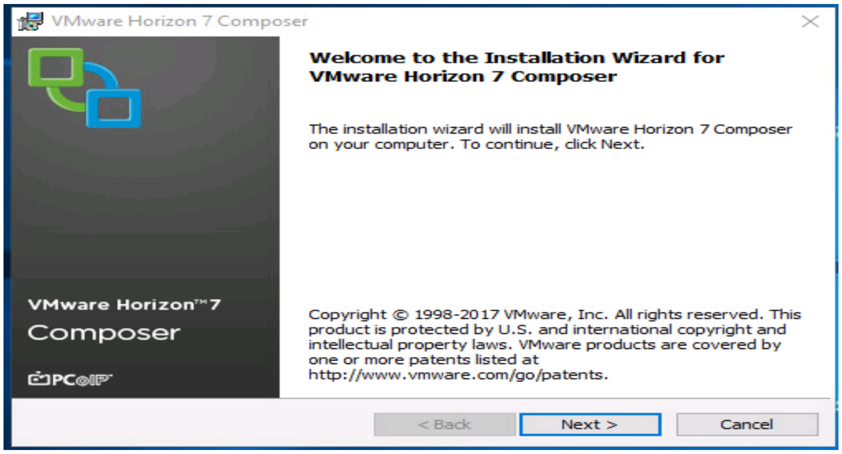

Welcome to the Installation Wizard for VMware Horizon 7 Composer

The installation wizard will install VMware Horizon 7 Composer on your computer. To continue, click Next.

Copyright © 1998-2017 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

2. Click the Install Horizon Composer installer and click Next.

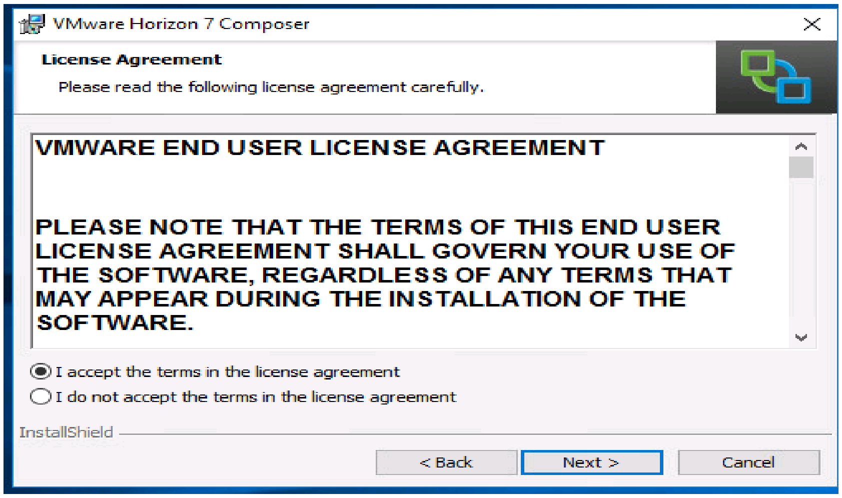

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

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

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

3. Accept the License Agreement and click Next.

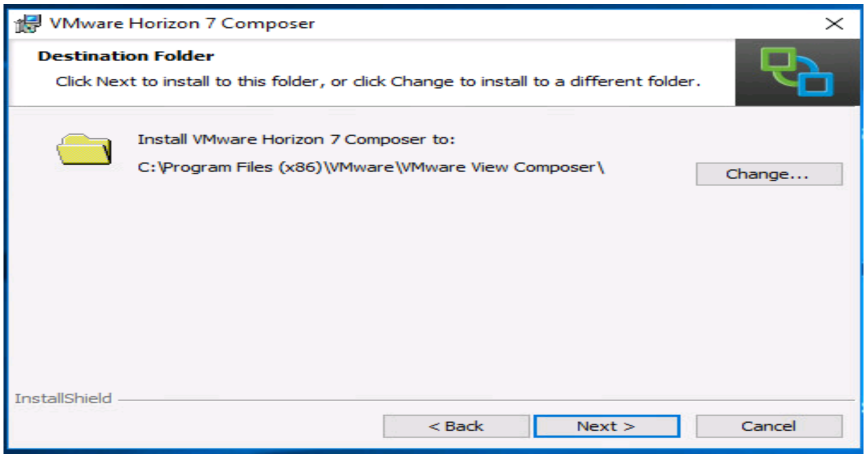

VMware Horizon 7 Composer

Destination Folder

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

Install VMware Horizon 7 Composer to:

C:\Program Files (x86)\VMware\VMware View Composer\

4. Click Install.

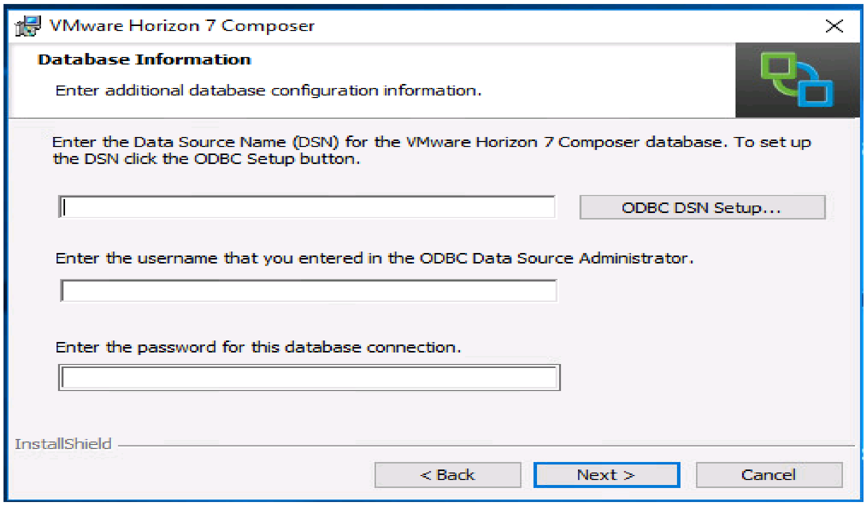

VMware Horizon 7 Composer

Database Information

Enter additional database configuration information.

Enter the Data Source Name (DSN) for the VMware Horizon 7 Composer database. To set up the DSN click the ODBC Setup button.

Enter the username that you entered in the ODBC Data Source Administrator.

Enter the password for this database connection.

5. Provide ODBC database connection details and click Next.

The VMware Horizon 7 Composer is being installed.

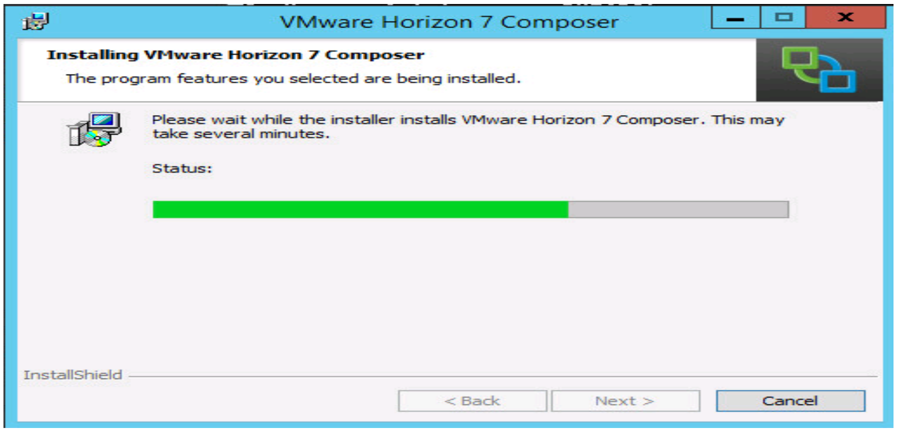

VMware Horizon 7 Composer

Installing VMware Horizon 7 Composer

The program features you selected are being installed.

Please wait while the installer installs VMware Horizon 7 Composer. This may take several minutes.

Status:

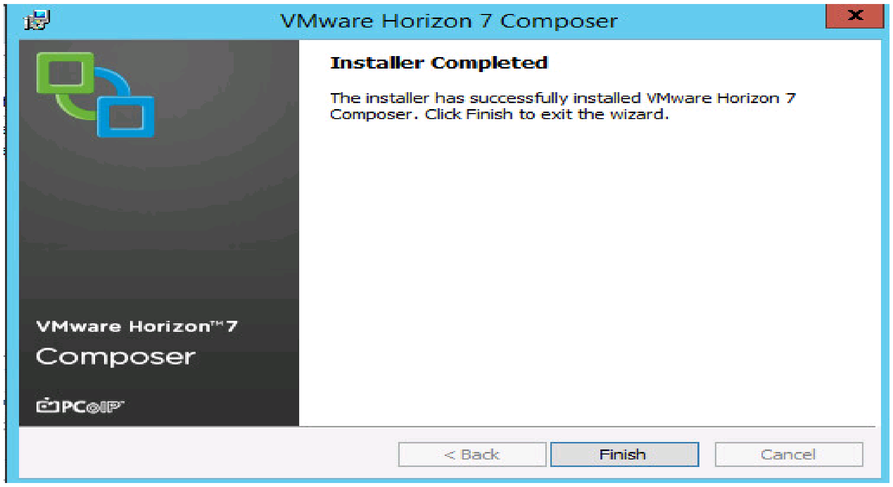

VMware Horizon 7 Composer

Installer Completed

The installer has successfully installed VMware Horizon 7 Composer. Click Finish to exit the wizard.

6. Click Finish.

Create the Golden Image for VMware Horizon RDS Deployment

You need to create the Golden Image as a prerequisite to install and configure server 2016. To create the Golden Image, complete the following steps:

We used Microsoft 2016 Standard Edition to configure the RDS Server Roles for RDSH VMs to be deployed from the Master image.

---

1. Login to Windows 2016 RDS Server base VM (Master Image) and click Server Properties and then click Add Roles and Features.

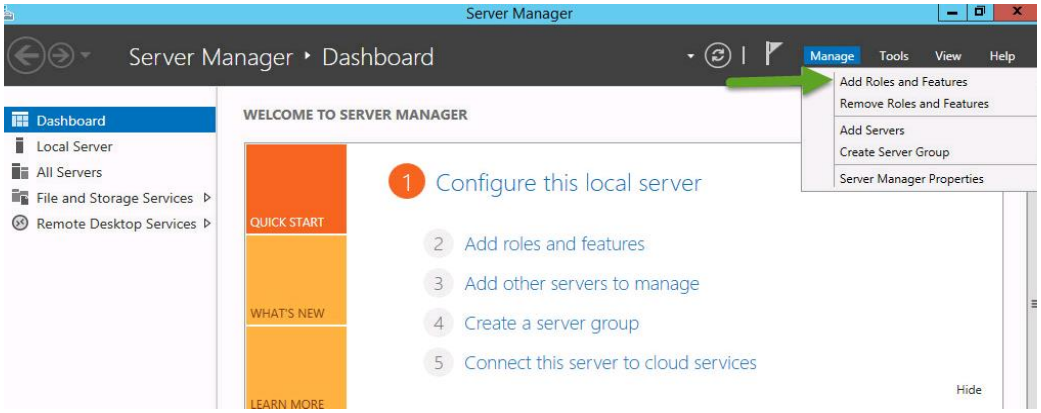

Server Manager
Server Manager ‣ Dashboard
Manage
Tools
View
Help
Add Roles and Features
Remove Roles and Features
Add Servers
Create Server Group
Server Manager Properties
Dashboard
Local Server
All Servers
File and Storage Services
Remote Desktop Services
WELCOME TO SERVER MANAGER
QUICK START
WHAT'S NEW
LEARN MORE
1 Configure this local server
2 Add roles and features
3 Add other servers to manage
4 Create a server group
5 Connect this server to cloud services
Hide

2. Select Role based or feature-based installation.

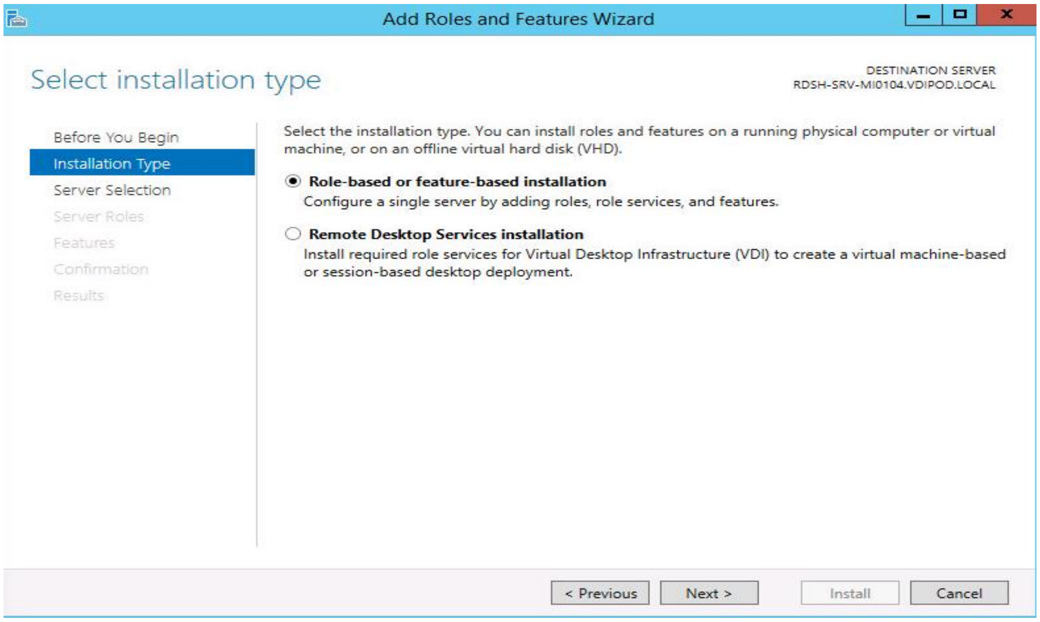

Select installation type

DESTINATION SERVER
RDSH-SRV-MI0104.VDIPOD.LOCAL

Select the installation type. You can install roles and features on a running physical computer or virtual machine, or on an offline virtual hard disk (VHD).

● **Role-based or feature-based installation**
Configure a single server by adding roles, role services, and features.

○ **Remote Desktop Services installation**
Install required role services for Virtual Desktop Infrastructure (VDI) to create a virtual machine-based or session-based desktop deployment.

3. Select Server Roles.

4. Remote Desktop Services > select Remote Desktop Session Host.

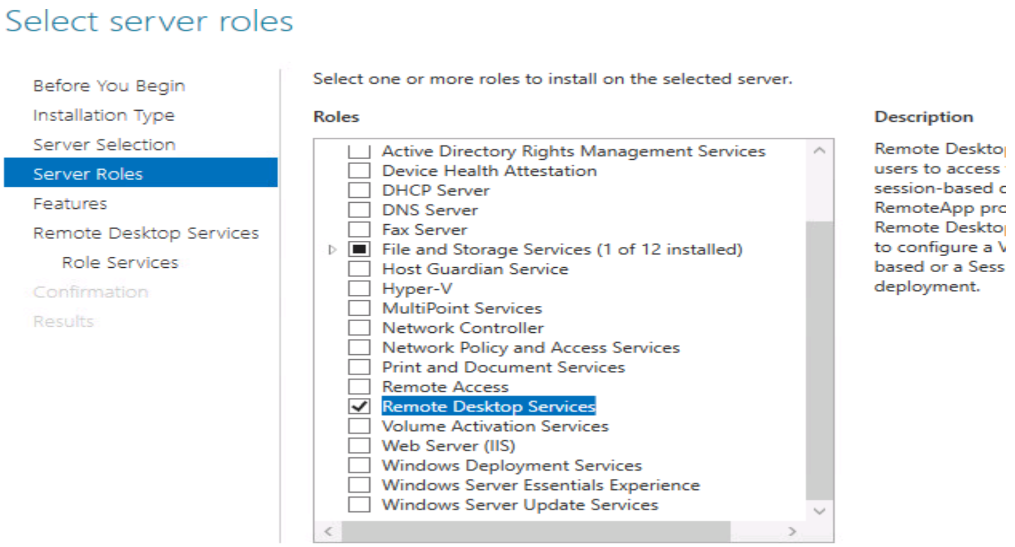
Select server roles

Select one or more roles to install on the selected server.

Roles

□ Active Directory Rights Management Services
□ Device Health Attestation
□ DHCP Server
□ DNS Server
□ Fax Server
▪ File and Storage Services (1 of 12 installed)
□ Host Guardian Service
□ Hyper-V
□ MultiPoint Services
□ Network Controller
□ Network Policy and Access Services
□ Print and Document Services
□ Remote Access
☑ Remote Desktop Services
□ Volume Activation Services
□ Web Server (IIS)
□ Windows Deployment Services
□ Windows Server Essentials Experience
□ Windows Server Update Services

Description

Remote Deskto
users to access
session-based c
RemoteApp prc
Remote Deskto
to configure a V
based or a Sess
deployment.

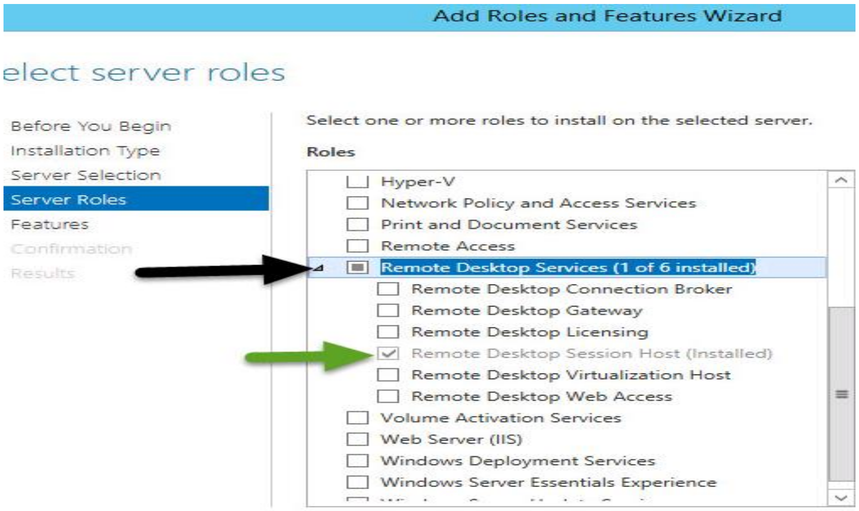

Add Roles and Features Wizard

elect server roles

Before You Begin
Installation Type
Server Selection
Server Roles
Features
Confirmation
Results

Select one or more roles to install on the selected server.

Roles

- □ Hyper-V
- □ Network Policy and Access Services
- □ Print and Document Services
- □ Remote Access
- ■ Remote Desktop Services (1 of 6 installed)
  - □ Remote Desktop Connection Broker
  - □ Remote Desktop Gateway
  - □ Remote Desktop Licensing
  - ☑ Remote Desktop Session Host (Installed)
  - □ Remote Desktop Virtualization Host
  - □ Remote Desktop Web Access
- □ Volume Activation Services
- □ Web Server (IIS)
- □ Windows Deployment Services
- □ Windows Server Essentials Experience

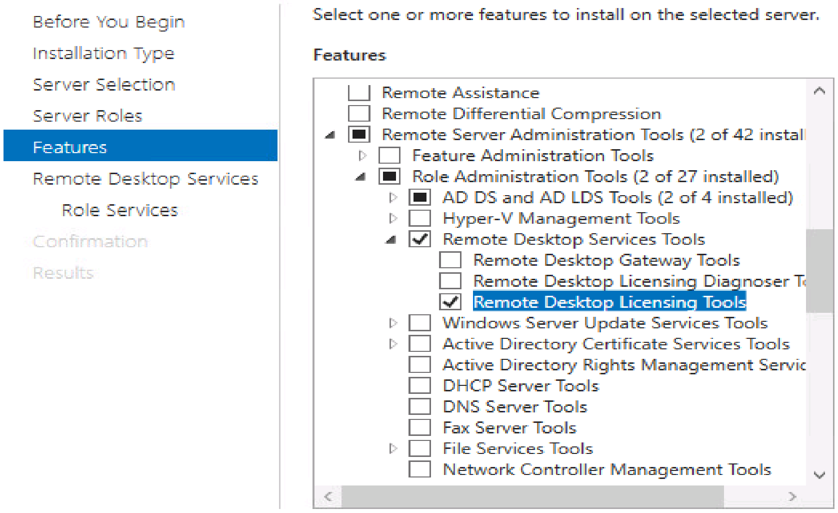

Before You Begin
Installation Type
Server Selection
Server Roles
Features
Remote Desktop Services
Role Services
Confirmation
Results

Select one or more features to install on the selected server.

Features

- □ Remote Assistance
- □ Remote Differential Compression
- ■ Remote Server Administration Tools (2 of 42 instal
  - □ Feature Administration Tools
  - ■ Role Administration Tools (2 of 27 installed)
    - ■ AD DS and AD LDS Tools (2 of 4 installed)
    - □ Hyper-V Management Tools
    - ☑ Remote Desktop Services Tools
      - □ Remote Desktop Gateway Tools
      - □ Remote Desktop Licensing Diagnoser To
      - ☑ Remote Desktop Licensing Tools
    - □ Windows Server Update Services Tools
    - □ Active Directory Certificate Services Tools
    - □ Active Directory Rights Management Servic
    - □ DHCP Server Tools
    - □ DNS Server Tools
    - □ Fax Server Tools
    - □ File Services Tools
    - □ Network Controller Management Tools

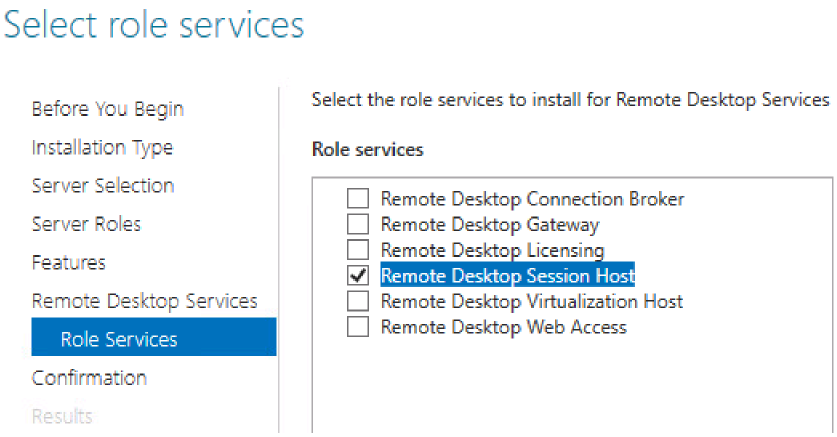

Select role services

Before You Begin
Installation Type
Server Selection
Server Roles
Features
Remote Desktop Services
Role Services
Confirmation
Results

Select the role services to install for Remote Desktop Services

Role services

- □ Remote Desktop Connection Broker
- □ Remote Desktop Gateway
- □ Remote Desktop Licensing
- ☑ Remote Desktop Session Host
- □ Remote Desktop Virtualization Host
- □ Remote Desktop Web Access

5. Click Next and complete the RDS server roles for RDSH Server Sessions enablement.

Create the Golden Image for Horizon Linked Clone Desktops

To create the Golden Image, complete the following steps:

1. Select ESXi host in Infrastructure cluster and create a virtual machine to use as the Golden Image with windows 10 OS. We used windows 10, 64 bit OS for our testing.

For the virtual machine, the following parameters were used:

Memory : 2048MB

Processor : 2vCPU

Hard Disk : 32 GB

Network Adapter : 1 VMXNET3 type attached to VDI-DVS port-group on VMware Distributed Virtual

2. Attach the already downloaded Windows 10 LTSB Version Build 1607 to the virtual machine to complete the Windows 10 Master image installation.

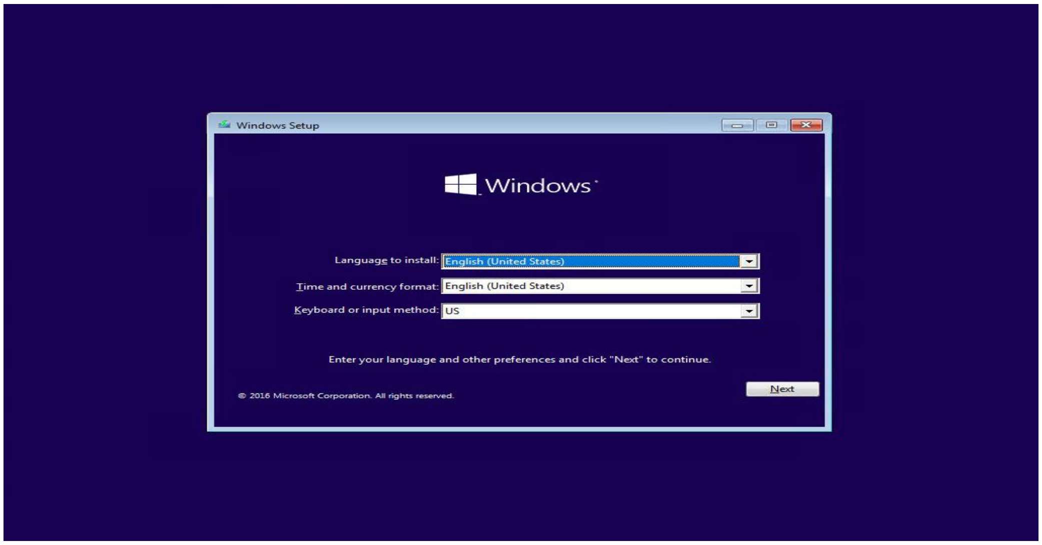

Language to install: English (United States)

Time and currency format: English (United States)

Keyboard or input method: US

Enter your language and other preferences and click "Next" to continue.

© 2016 Microsoft Corporation. All rights reserved.

3. Click Next.

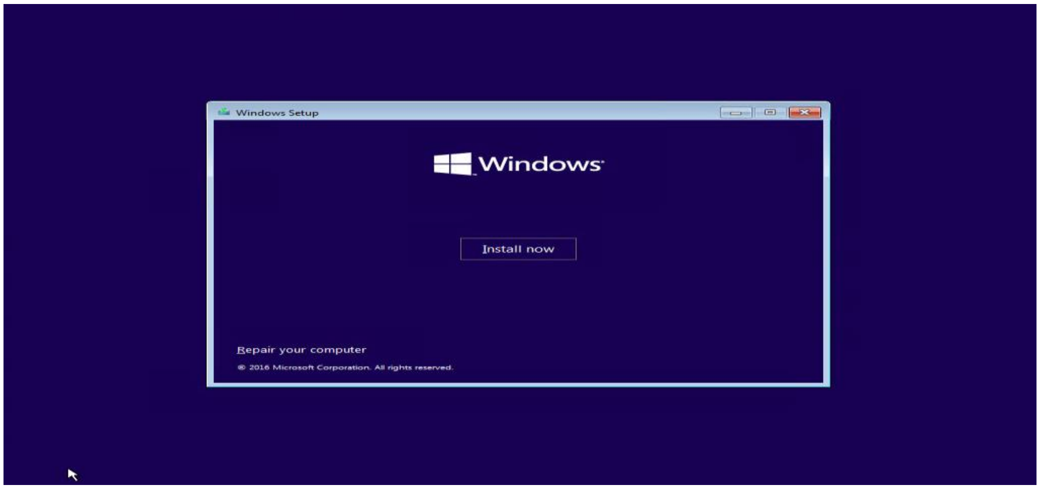

© 2016 Microsoft Corporation. All rights reserved.

4. Click Install now.

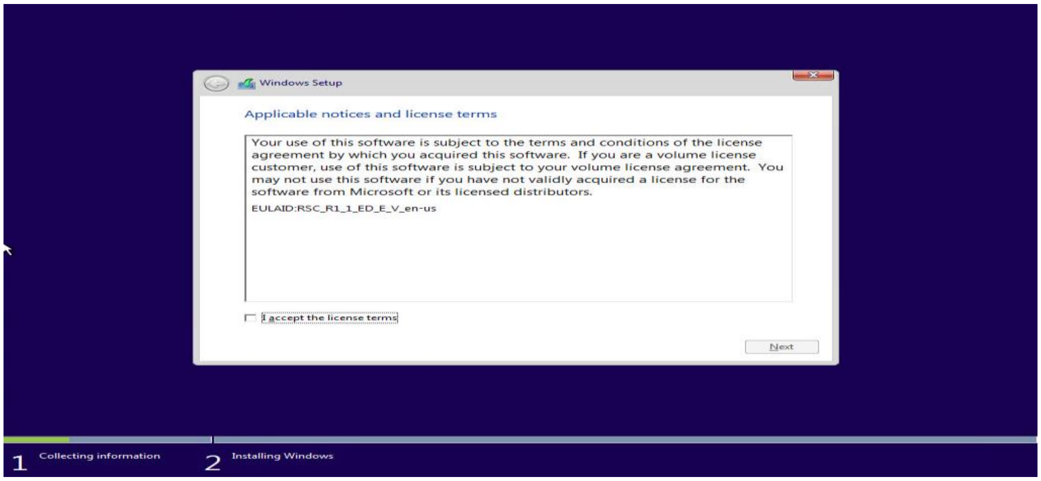

Applicable notices and license terms

Your use of this software is subject to the terms and conditions of the license agreement by which you acquired this software. If you are a volume license customer, use of this software is subject to your volume license agreement. You may not use this software if you have not validly acquired a license for the software from Microsoft or its licensed distributors.

EULAID:RSC_R1_1_ED_E_V_en-us

□ I accept the license terms

Next

5. Accept the license terms.

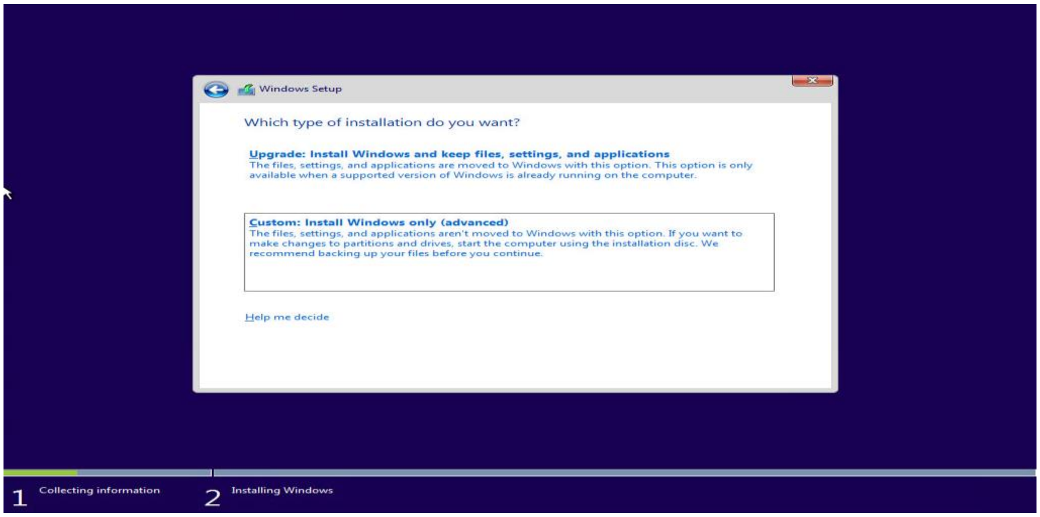

Which type of installation do you want?

Upgrade: Install Windows and keep files, settings, and applications
The files, settings, and applications are moved to Windows with this option. This option is only available when a supported version of Windows is already running on the computer.

Custom: Install Windows only (advanced)
The files, settings, and applications aren't moved to Windows with this option. If you want to make changes to partitions and drives, start the computer using the installation disc. We recommend backing up your files before you continue.

Help me decide

6. Click Install windows.

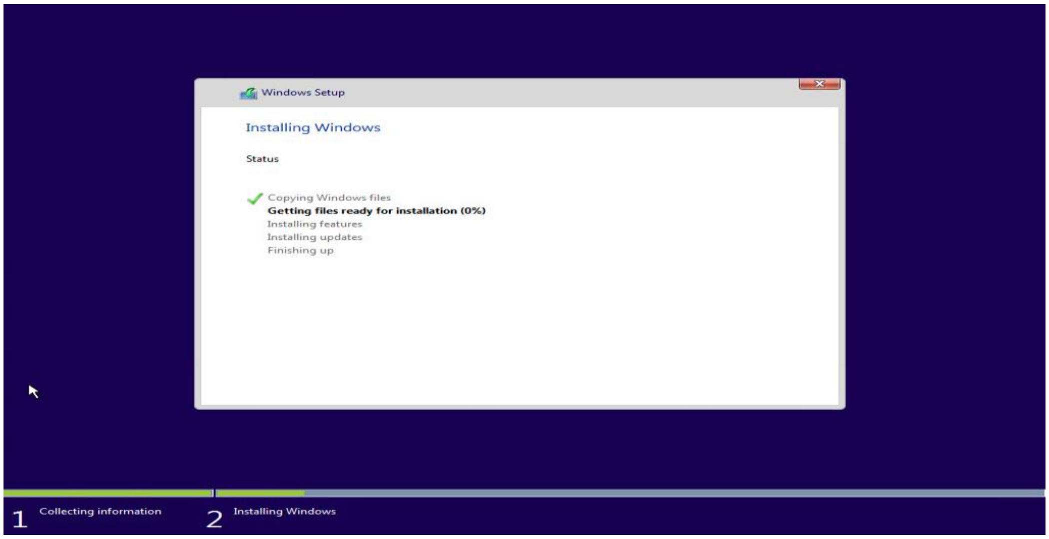

Installing Windows

Status

✓ Copying Windows files
Getting files ready for installation (0%)
Installing features
Installing updates
Finishing up

Windows is installing.

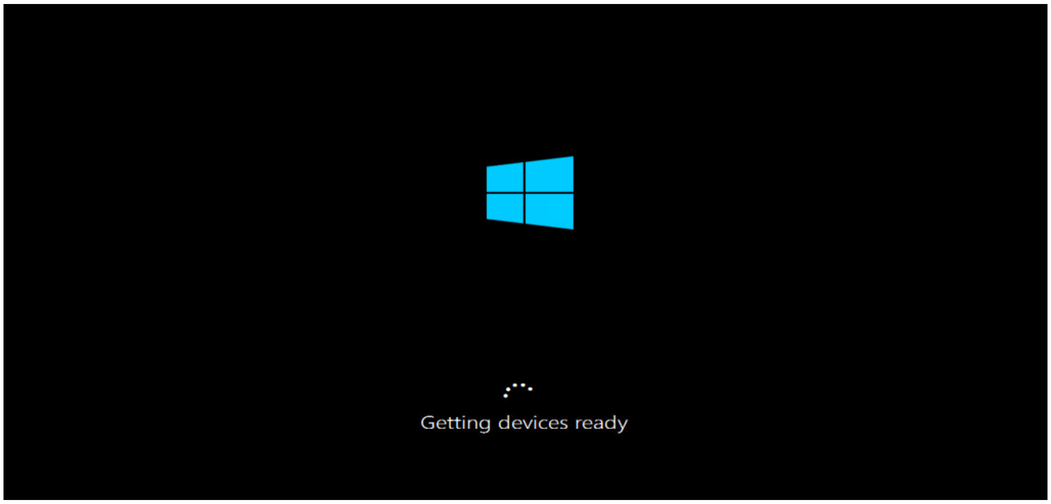

Getting devices ready

The Windows installation is complete.

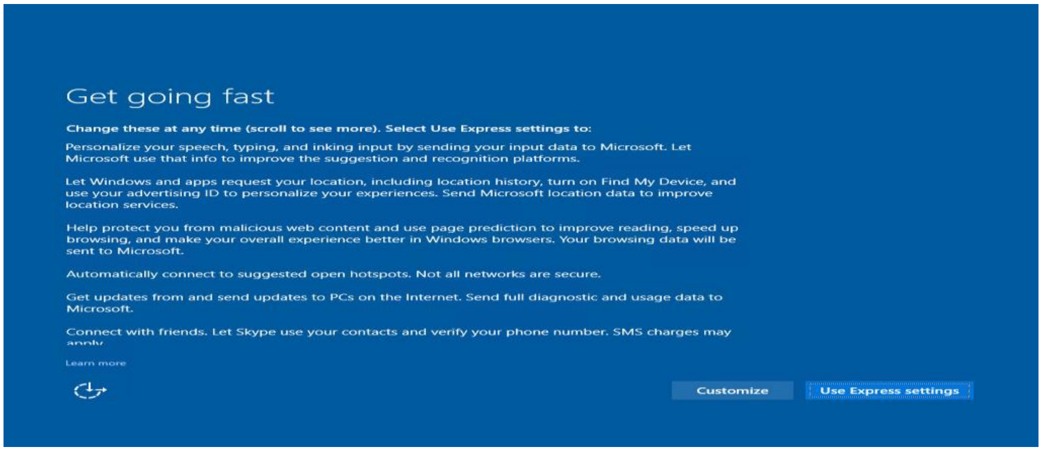

Get going fast

Change these at any time (scroll to see more). Select Use Express settings to:

Personalize your speech, typing, and inking input by sending your input data to Microsoft. Let Microsoft use that info to improve the suggestion and recognition platforms.

Let Windows and apps request your location, including location history, turn on Find My Device, and use your advertising ID to personalize your experiences. Send Microsoft location data to improve location services.

Help protect you from malicious web content and use page prediction to improve reading, speed up browsing, and make your overall experience better in Windows browsers. Your browsing data will be sent to Microsoft.

Automatically connect to suggested open hotspots. Not all networks are secure.

Get updates from and send updates to PCs on the Internet. Send full diagnostic and usage data to Microsoft.

Connect with friends. Let Skype use your contacts and verify your phone number. SMS charges may apply

7. Customize the windows setup or use Express Settings.

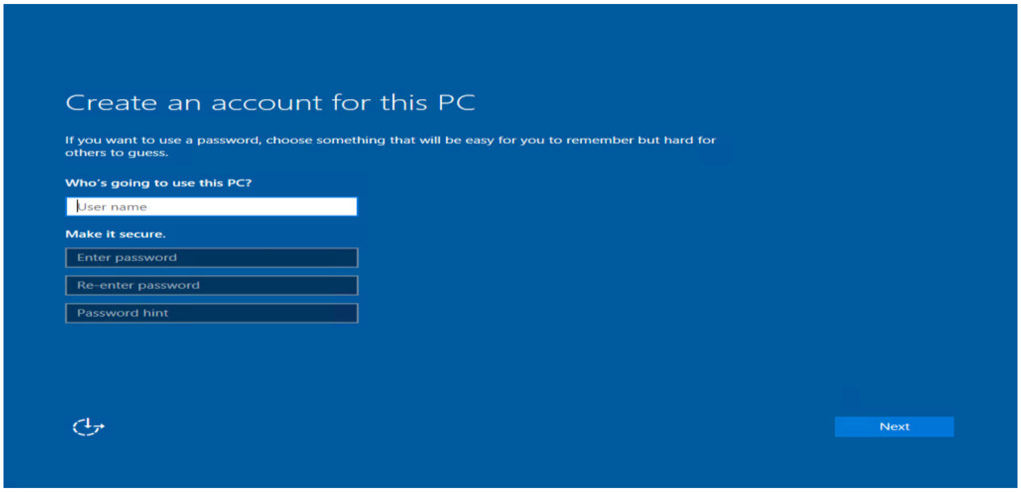

Create an account for this PC

If you want to use a password, choose something that will be easy for you to remember but hard for others to guess.

Who's going to use this PC?

Make it secure.

8. Provide the User credentials for the windows VM created.

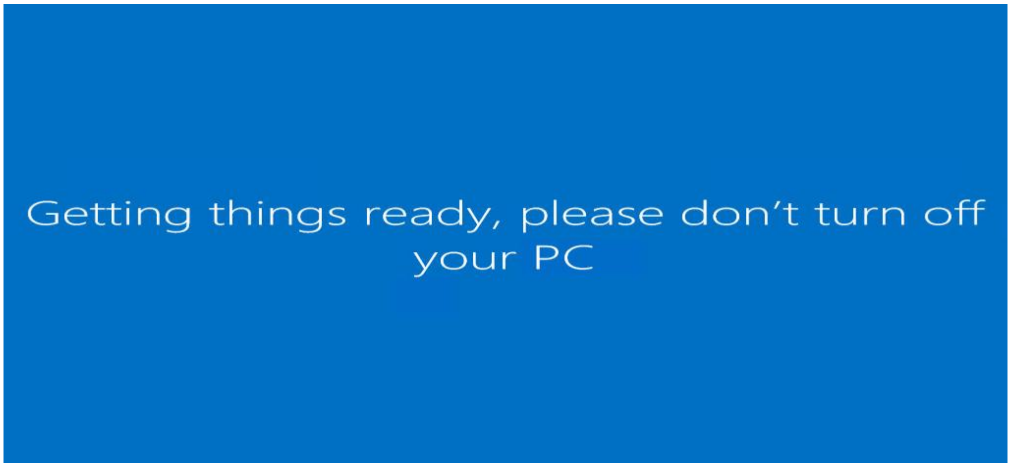

Getting things ready, please don't turn off your PC

9. Reboot the Windows VM and install additional software as applicable.

VMware Horizon Agent Installation

To install the VMware Horizon Agent, complete the following steps:

1. Download VMware-viewagent-x86_64-7.3.1-6761332 version.

https://my.vmware.com/web/vmware/downloads

https://my.vmware.com/group/vmware/details?downloadGroup=VIEW-732-STD&productId=681&download=true&fileId=977d331ce20874ba55052514e85bbab4&secureParam=dbc8854351a3ea1f0d6214f4fd9c0d59&uuId=2560f9b4-f3d3-4a9c-ae79-4ef35994ed2b&downloadType=

2. Click the VMware Horizon Agent installer.

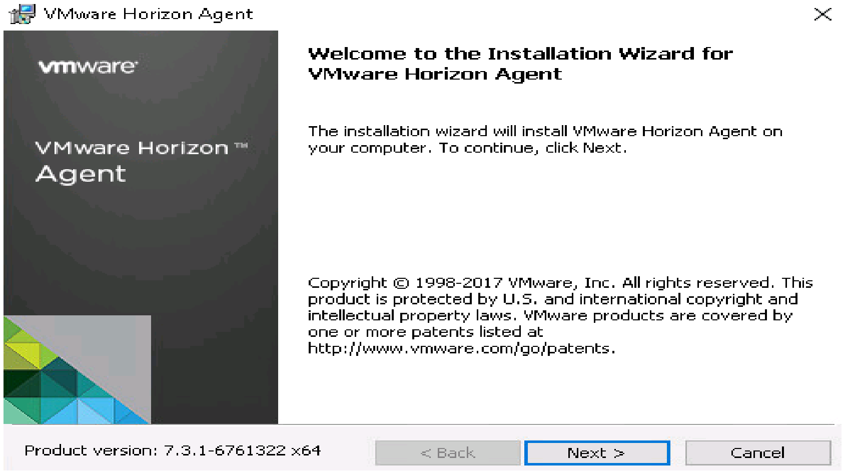

Welcome to the Installation Wizard for VMware Horizon Agent

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

Copyright © 1998-2017 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 7.3.1-6761322 x64

3. Click Next.

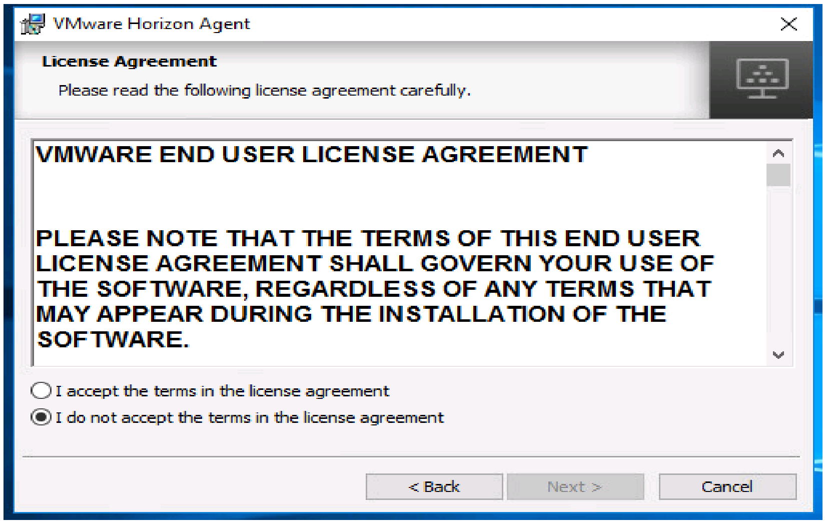

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

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

○ I accept the terms in the license agreement

● I do not accept the terms in the license agreement

4. Accept the license agreement and click Next.

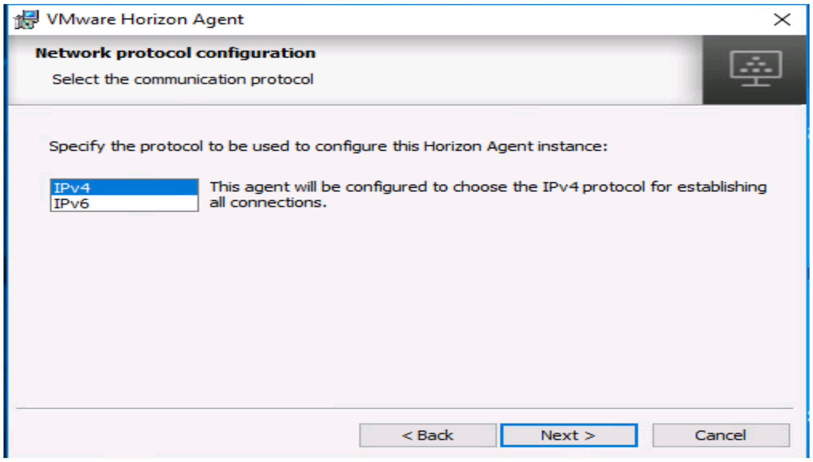

Network protocol configuration

Select the communication protocol

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

IPv4 (selected)
IPv6

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

5. Select the default IPV4 and click Next.

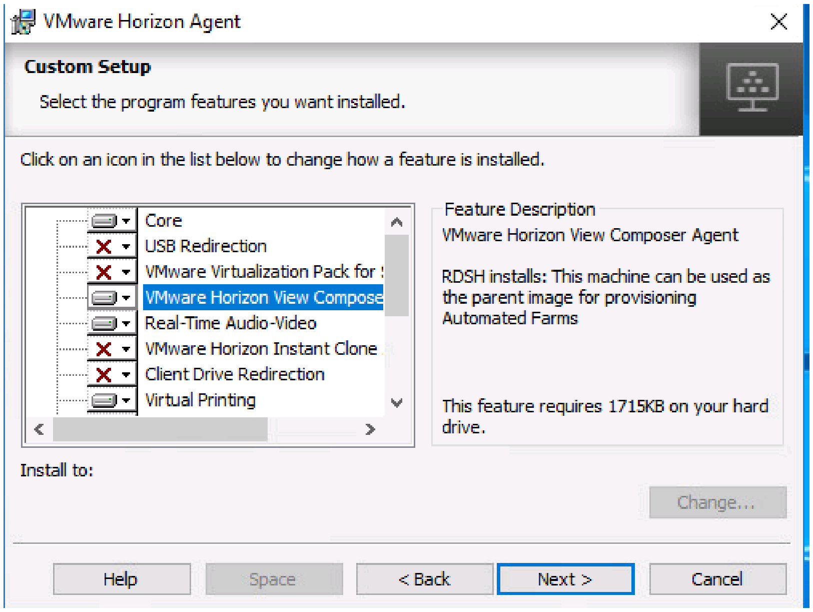

Custom Setup

Select the program features you want installed.

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

- Core
- ☒ USB Redirection
- ☒ VMware Virtualization Pack for !
- VMware Horizon View Compose
- Real-Time Audio-Video
- ☒ VMware Horizon Instant Clone
- ☒ Client Drive Redirection
- Virtual Printing

Feature Description

VMware Horizon View Composer Agent

RDSH installs: This machine can be used as the parent image for provisioning Automated Farms

This feature requires 1715KB on your hard drive.

Install to:

6. Select the features to install.

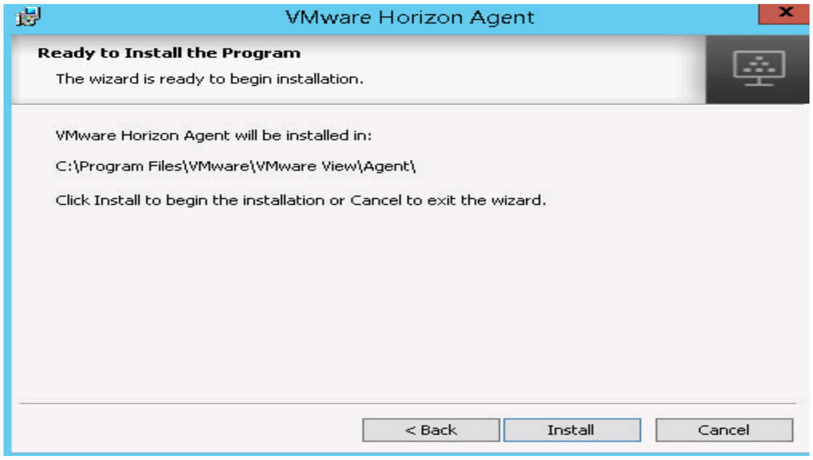

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Agent will be installed in:

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

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

7. Click Install.

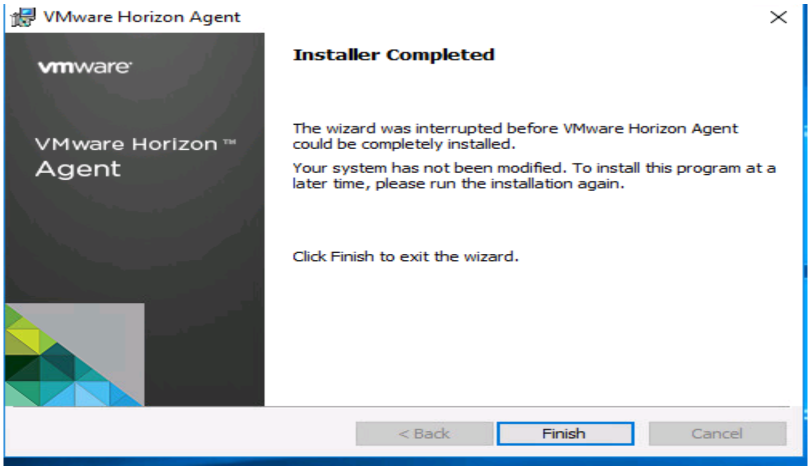

Installer Completed

The wizard was interrupted before VMware Horizon Agent could be completely installed.

Your system has not been modified. To install this program at a later time, please run the installation again.

Click Finish to exit the wizard.

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

The same agent installation steps are applicable for the RDSH Server 2016 base Image you intend to use for RDS VMs.

Prerequisites

To set the Scope Options on the DHCP server hosting the Horizon target machines (for example, RDSH and VDI virtual machines), complete the following steps:

1. From the DHCP server, navigate to the Scope folder.

2. Set your options.

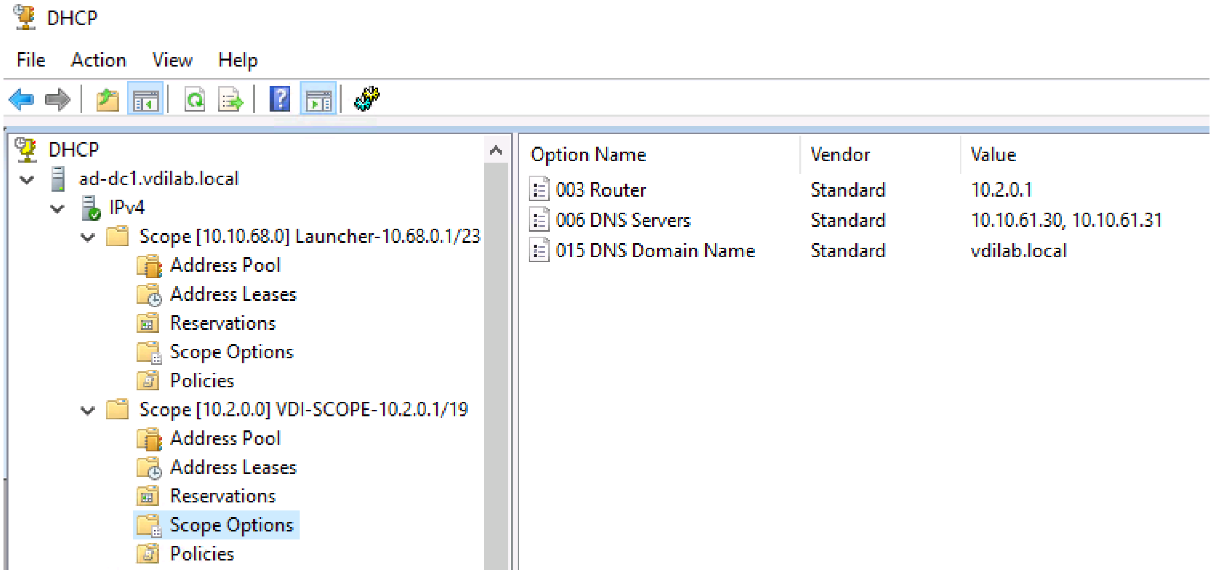

Option Name	Vendor	Value
003 Router	Standard	10.2.0.1
006 DNS Servers	Standard	10.10.61.30, 10.10.61.31
015 DNS Domain Name	Standard	vdilab.local

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.

5. In case of RDSH VM, clone the RDS Master Image.

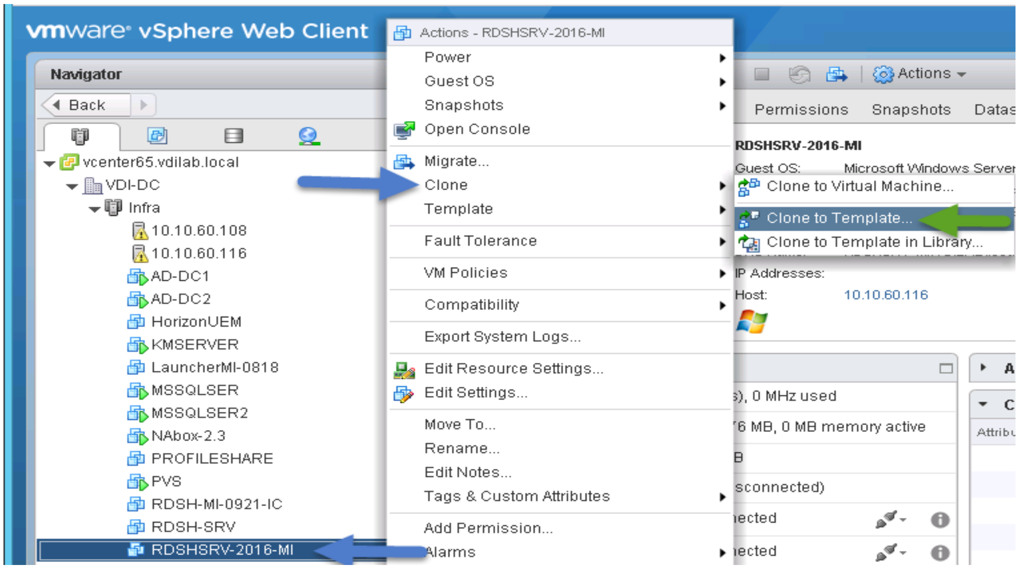

6. For Windows 10, follow the steps for cloning the Master Image for further deployment.

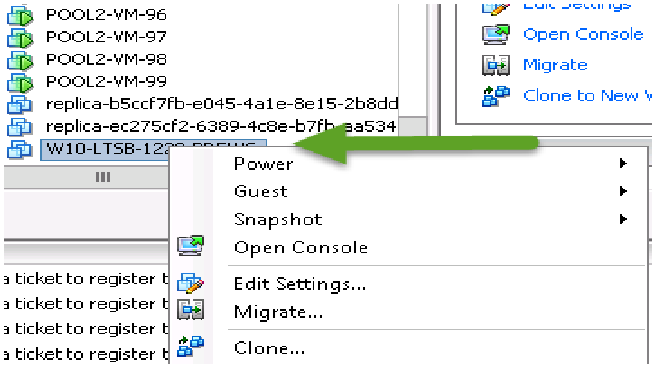

7. Change the memory to “Reserve all guest memory” (All locked).

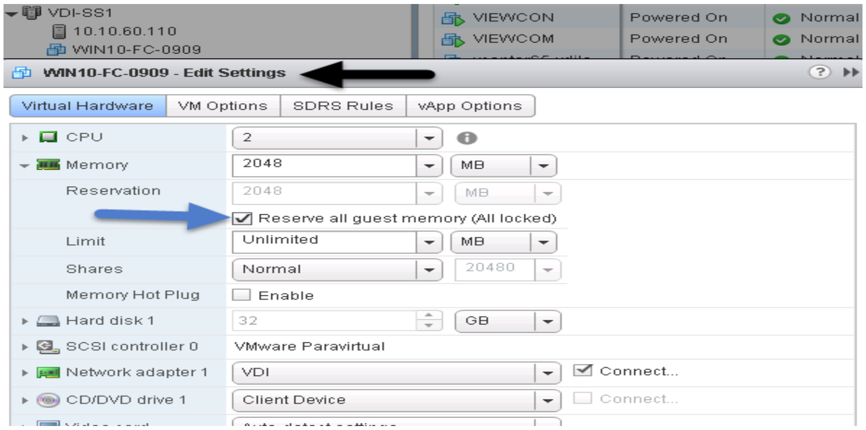

CPU: 2
Memory: 2048 MB
Reservation: 2048 MB
☑ Reserve all guest memory (All locked)
Limit: Unlimited MB
Shares: Normal 20480
Memory Hot Plug: □ Enable
Hard disk 1: 32 GB
SCSI controller 0: VMware Paravirtual
Network adapter 1: VDI ☑ Connect...
CD/DVD drive 1: Client Device □ Connect...

8. Convert the VM to a Template for additional steps.

9. When the template is ready, convert to VM and take a snapshot of the VM to deploy the VDI virtual machines from Horizon Administrator Console.

10. Right-click the Master Image and take a snapshot.

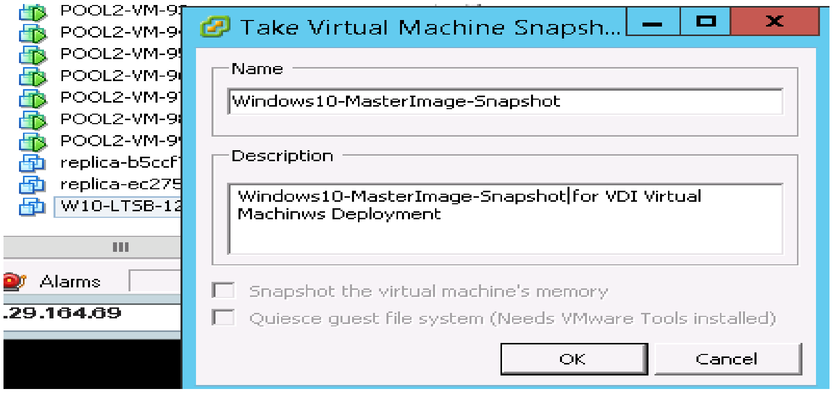

Take Virtual Machine Snapsh...

Name: Windows10-MasterImage-Snapshot

Description: Windows10-MasterImage-Snapshot for VDI Virtual Machinws Deployment

□ Snapshot the virtual machine's memory

□ Quiesce guest file system (Needs VMware Tools installed)

11. Provide a Name for the Master Image snapshot and click OK.

VMware Horizon Desktop Pool Creation

To create the VMware Horizon Desktop Pool, complete the following steps;

1. Login to Horizon 7 Administrator via a web browser. Address or FQDN>/admin.

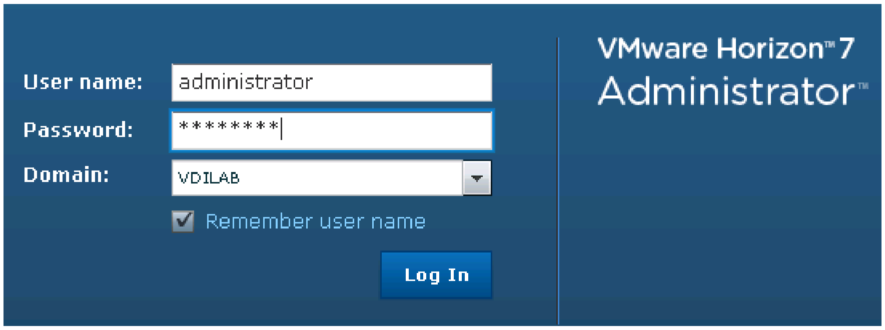

User name: administrator

Password: ********

Domain: VDILAB

☑ Remember user name

2. Login to VMware Horizon Connection Server Administrator to create the RDSH farm and RDSH pool or VDI Desktops Pools.

Create VDI Desktop Pool

1. Select Type of Desktop pool to be created.

We created an Automated Desktop Pool.

2. Click Next.

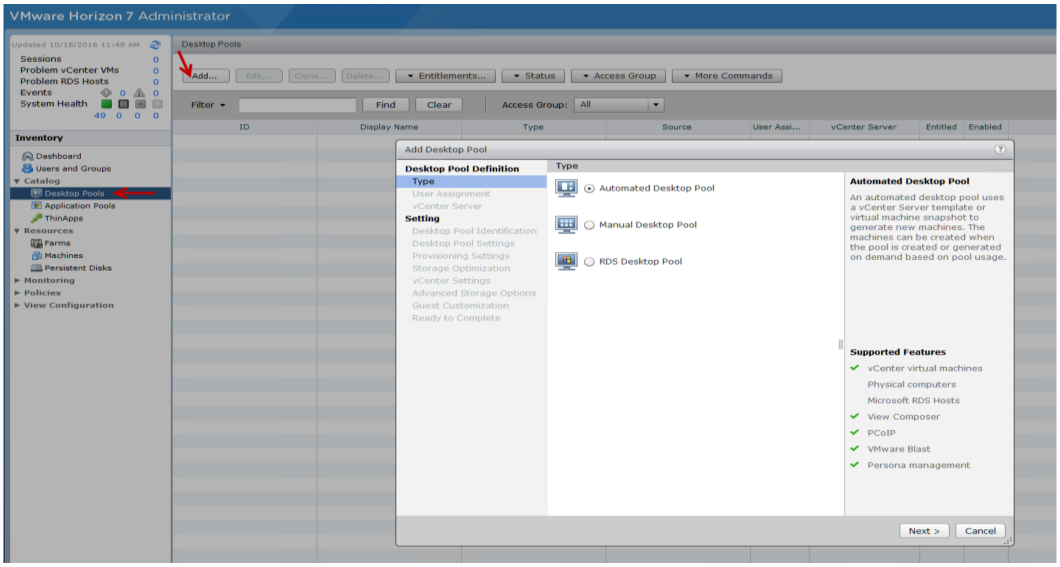

Type

● Automated Desktop Pool

○ Manual Desktop Pool

○ RDS Desktop Pool

Automated Desktop Pool

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

Supported Features

- ✓ vCenter virtual machines
- Physical computers
- Microsoft RDS Hosts
- ✓ View Composer
- ✓ PCoIP
- ✓ VMware Blast
- ✓ Persona management

3. Select the pool type by clicking the Add button (for example, automated pool or RDS pool).

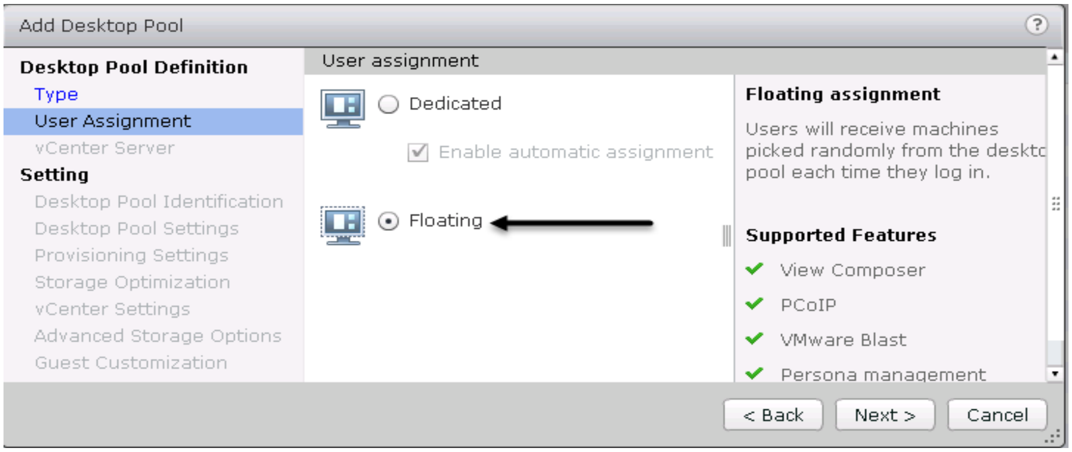

User assignment

○ Dedicated

☑ Enable automatic assignment

● Floating

Floating assignment

Users will receive machines picked randomly from the deskto pool each time they log in.

Supported Features

- ✓ View Composer
- ✓ PCoIP
- ✓ VMware Blast
- ✓ Persona management

4. Select User Assignment for desktop pool. (We created Floating assigned.) Click Next.

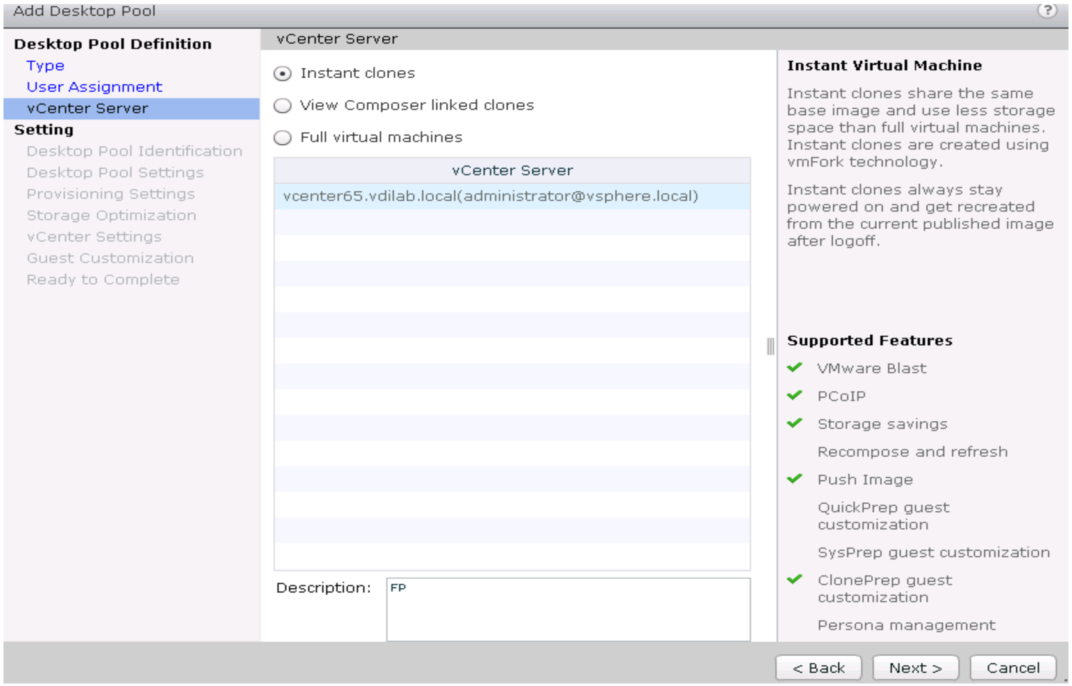

vCenter Server

● Instant clones

○ View Composer linked clones

○ Full virtual machines

vCenter Server
vcenter65.vdilab.local(administrator@vsphere.local)

Description: FP

Instant Virtual Machine

Instant clones share the same base image and use less storage space than full virtual machines. Instant clones are created using vmFork technology.

Instant clones always stay powered on and get recreated from the current published image after logoff.

Supported Features

- ✓ VMware Blast
- ✓ PCoIP
- ✓ Storage savings
- Recompose and refresh
- ✓ Push Image
- QuickPrep guest customization
- SysPrep guest customization
- ✓ ClonePrep guest customization
- Persona management

5. Select vCenter Server and the type of Desktop deployment. (We created Horizon Instant Clones).

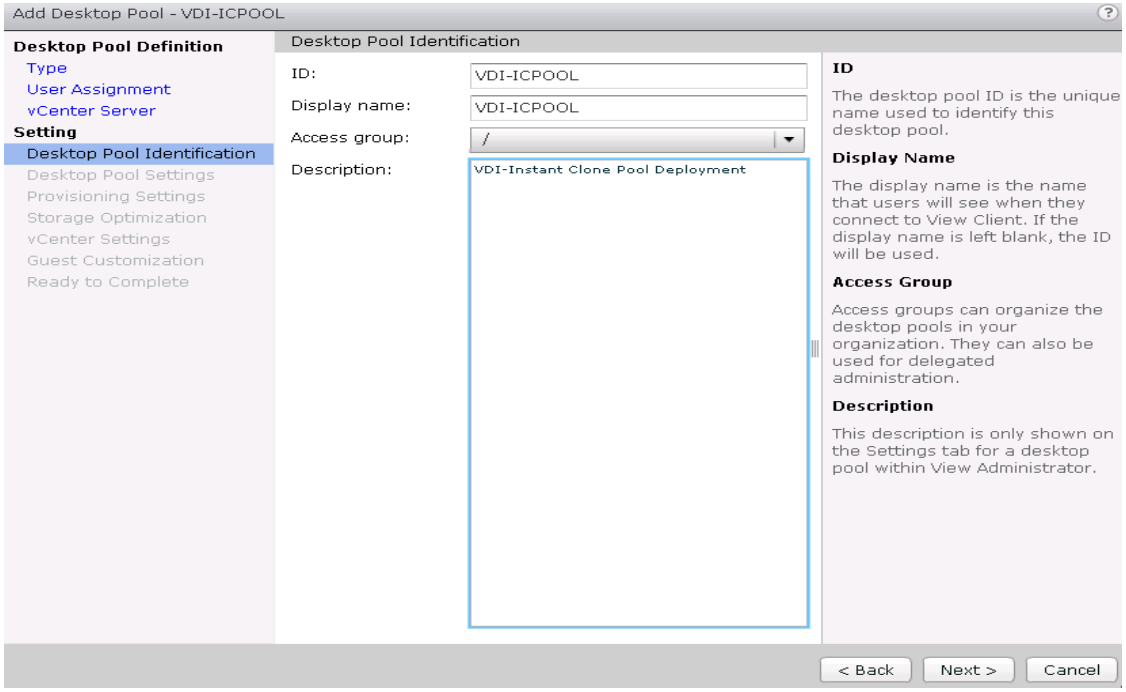

Desktop Pool Identification

ID: VDI-ICPOOL

Display name: VDI-ICPOOL

Access group: /

Description: VDI-Instant Clone Pool Deployment

ID

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

Display Name

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

Access Group

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

Description

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

6. Provide the Pool ID and virtual display name.

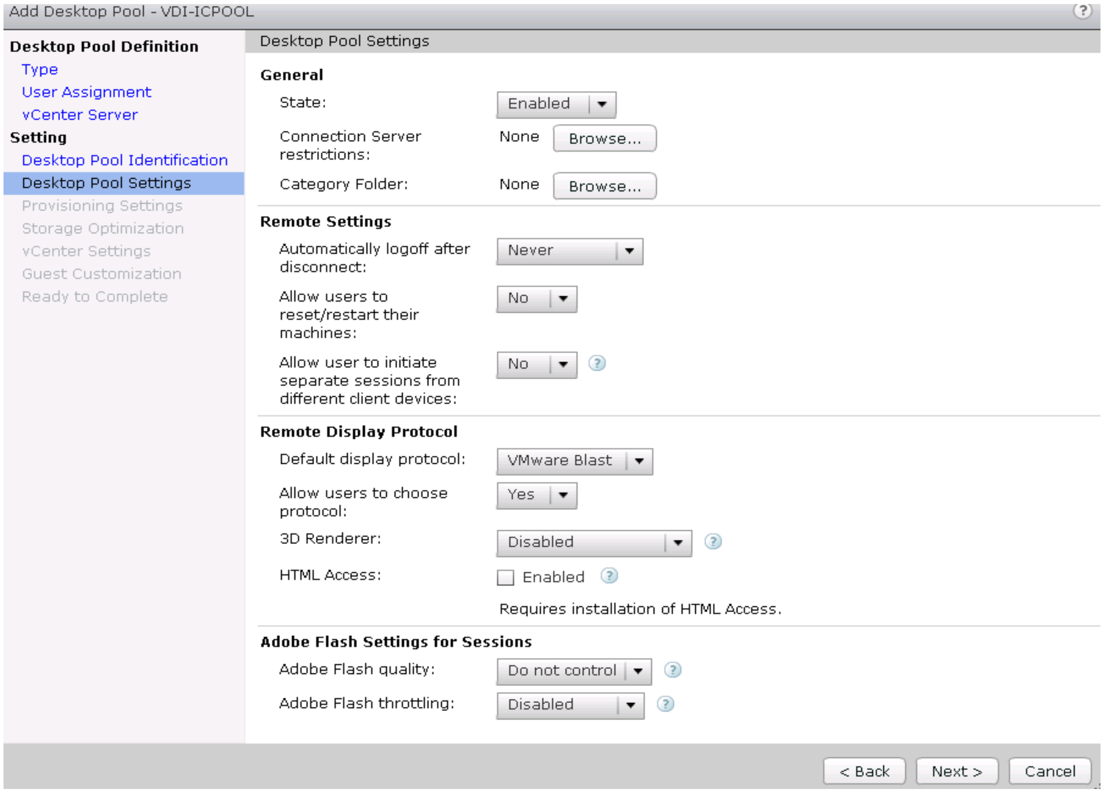

Add Desktop Pool - VDI-ICPOOL

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

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

Desktop Pool Settings

General

State: Enabled

Connection Server restrictions: None Browse...

Category Folder: None Browse...

Remote Settings

Automatically logoff after disconnect: Never

Allow users to reset/restart their machines: No

Allow user to initiate separate sessions from different client devices: No

Remote Display Protocol

Default display protocol: VMware Blast

Allow users to choose protocol: Yes

3D Renderer: Disabled

HTML Access: □ Enabled

Requires installation of HTML Access.

Adobe Flash Settings for Sessions

Adobe Flash quality: Do not control

Adobe Flash throttling: Disabled

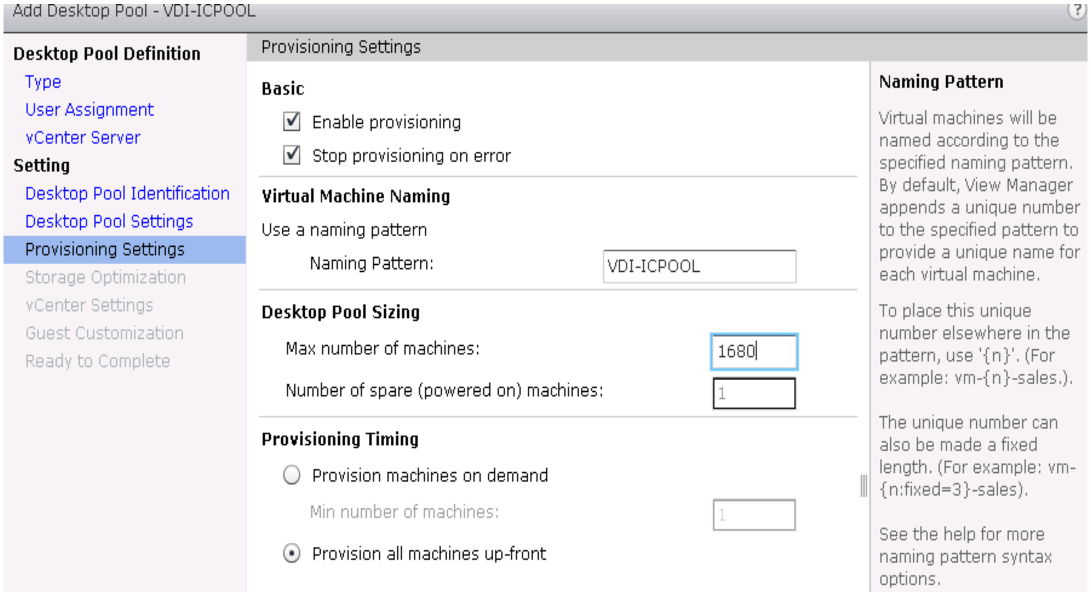

Add Desktop Pool - VDI-ICPOOL

Desktop Pool Definition

- Type
- User Assignment
- vCenter Server

Setting

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

Provisioning Settings

Basic

☑ Enable provisioning

☑ Stop provisioning on error

Virtual Machine Naming

Use a naming pattern

Naming Pattern: VDI-ICPOOL

Desktop Pool Sizing

Max number of machines: 1680

Number of spare (powered on) machines: 1

Provisioning Timing

○ Provision machines on demand

Min number of machines: 1

● Provision all machines up-front

Naming Pattern

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

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

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

See the help for more naming pattern syntax options.

7. Provide the number of desktops to be provisioned (shown above POOL1-VDI, total of 1680 desktops).

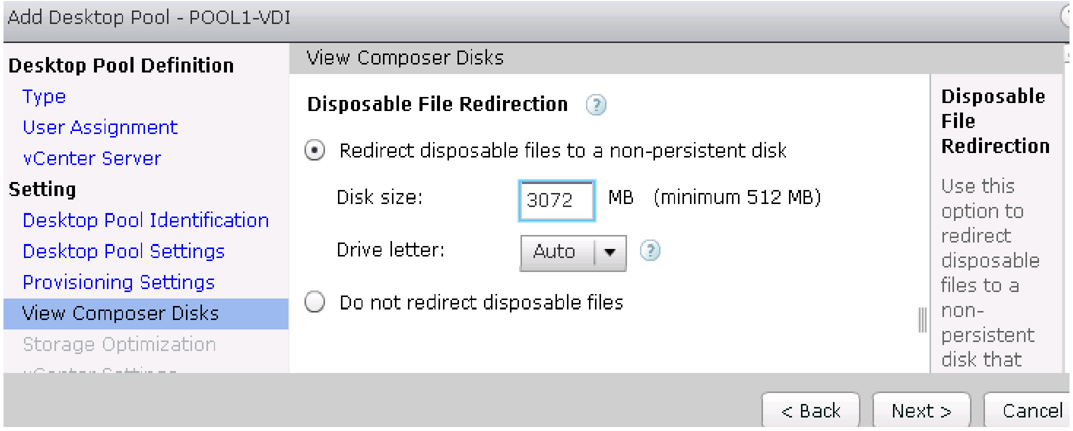

View Composer Disks

Disposable File Redirection

● Redirect disposable files to a non-persistent disk

Disk size: 3072 MB (minimum 512 MB)

Drive letter: Auto

○ Do not redirect disposable files

Disposable File Redirection

Use this option to redirect disposable files to a non-persistent disk that

8. Select the option Redirect disposable files to non-persistent disk.

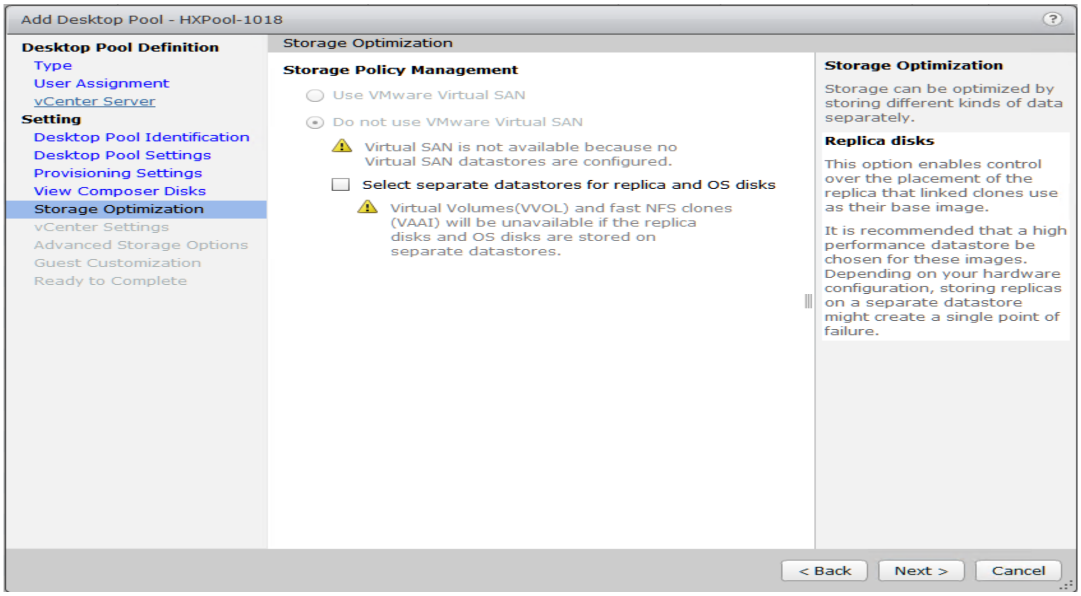

Storage Optimization

Storage Policy Management

○ Use VMware Virtual SAN

● Do not use VMware Virtual SAN

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

□ Select separate datastores for replica and OS disks

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

Storage Optimization

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

Replica disks

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

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

9. Select the required option for Storage Policy Management.

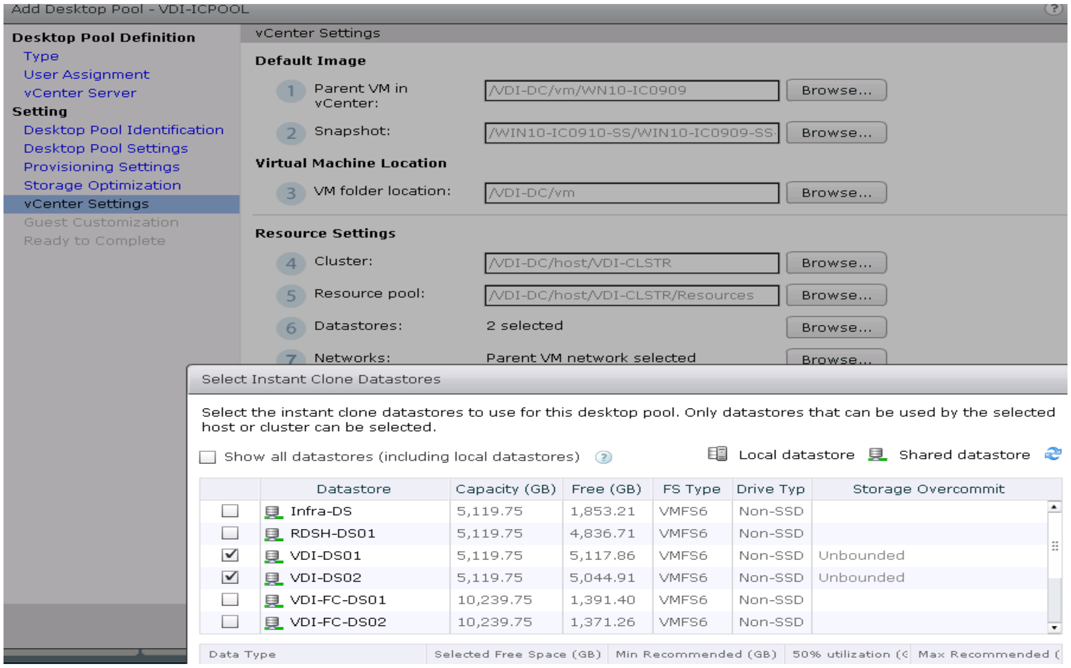

vCenter Settings

Default Image

1 Parent VM in vCenter: /VDI-DC/vm/WN10-IC0909

2 Snapshot: /WIN10-IC0910-SS/WIN10-IC0909-SS

Virtual Machine Location

3 VM folder location: /VDI-DC/vm

Resource Settings

4 Cluster: /VDI-DC/host/VDI-CLSTR

5 Resource pool: /VDI-DC/host/VDI-CLSTR/Resources

6 Datastores: 2 selected

7 Networks: Parent VM network selected

Select Instant Clone Datastores

Select the instant clone datastores to use for this desktop pool. Only datastores that can be used by the selected host or cluster can be selected.

□ Show all datastores (including local datastores)

Local datastore   Shared datastore

	Datastore	Capacity (GB)	Free (GB)	FS Type	Drive Typ	Storage Overcommit
□	Infra-DS	5,119.75	1,853.21	VMFS6	Non-SSD	
□	RDSH-DS01	5,119.75	4,836.71	VMFS6	Non-SSD	
☑	VDI-DS01	5,119.75	5,117.86	VMFS6	Non-SSD	Unbounded
☑	VDI-DS02	5,119.75	5,044.91	VMFS6	Non-SSD	Unbounded
□	VDI-FC-DS01	10,239.75	1,391.40	VMFS6	Non-SSD	
□	VDI-FC-DS02	10,239.75	1,371.26	VMFS6	Non-SSD	

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% utilization ((	Max Recommended (

10. Provide the parent VM, snapshot and host/cluster info, data store information for the virtual machines to create.

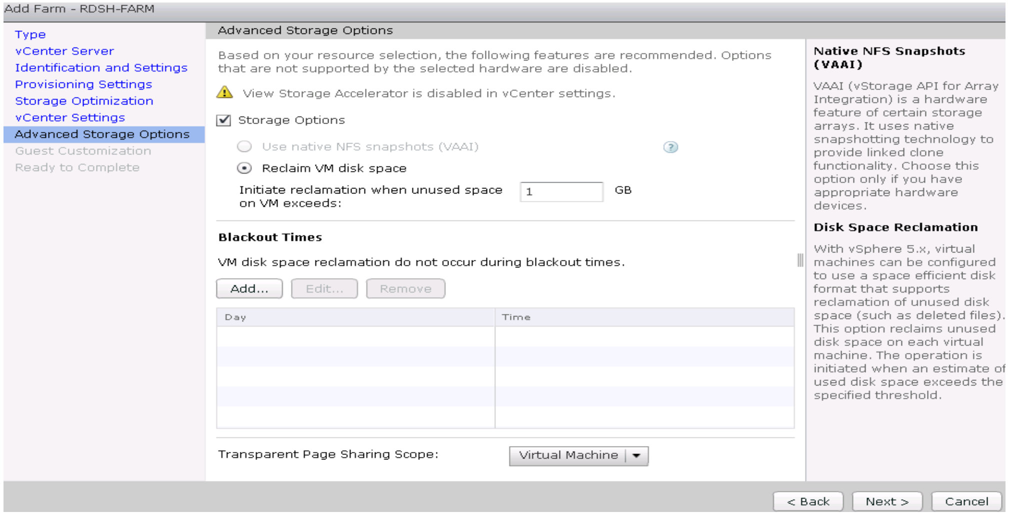

Advanced Storage Options

Based on your resource selection, the following features are recommended. Options that are not supported by the selected hardware are disabled.

View Storage Accelerator is disabled in vCenter settings.

☑ Storage Options

○ Use native NFS snapshots (VAAI)

● Reclaim VM disk space

Initiate reclamation when unused space on VM exceeds: 1 GB

Blackout Times

VM disk space reclamation do not occur during blackout times.

Day	Time

Transparent Page Sharing Scope: Virtual Machine

Native NFS Snapshots (VAAI)

VAAI (vStorage API for Array Integration) is a hardware feature of certain storage arrays. It uses native snapshotting technology to provide linked clone functionality. Choose this option only if you have appropriate hardware devices.

Disk Space Reclamation

With vSphere 5.x, virtual machines can be configured to use a space efficient disk format that supports reclamation of unused disk space (such as deleted files). This option reclaims unused disk space on each virtual machine. The operation is initiated when an estimate of used disk space exceeds the specified threshold.

11. In Advanced Storage Options, select the View Storage Accelerator and reclaim disk space parameters as required.

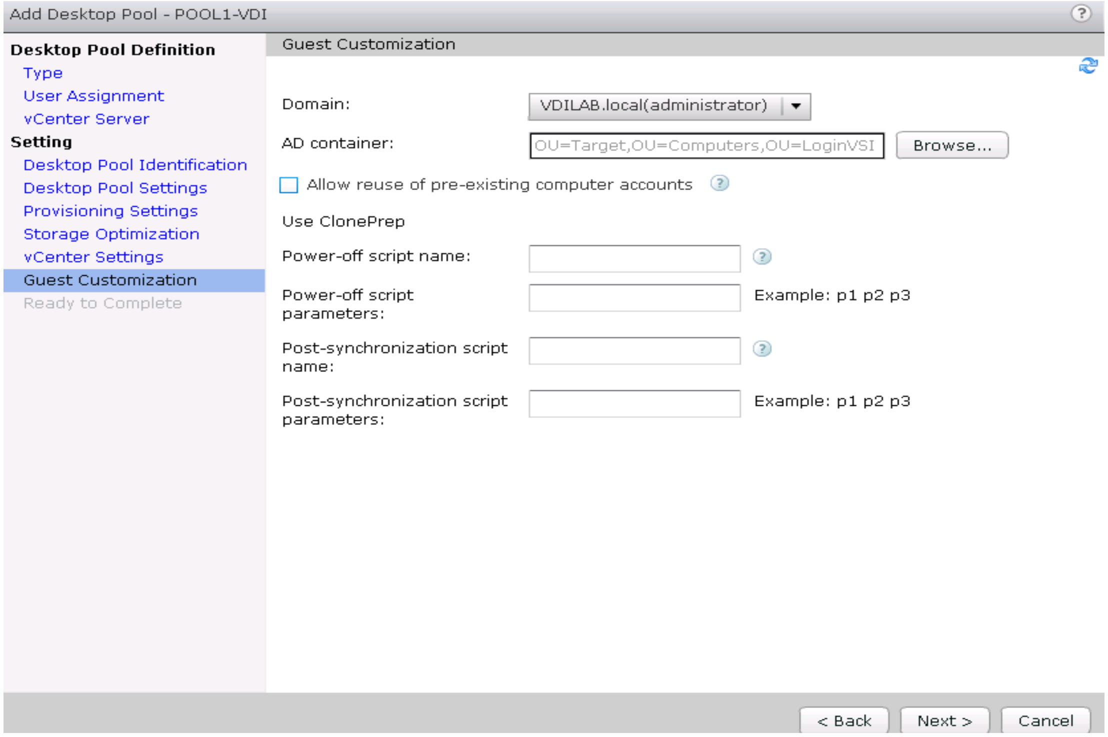

Guest Customization

Domain: VDILAB.local(administrator)

AD container: OU=Target,OU=Computers,OU=LoginVSI Browse...

□ Allow reuse of pre-existing computer accounts

Use ClonePrep

Power-off script name:

Power-off script parameters: Example: p1 p2 p3

Post-synchronization script name:

Post-synchronization script parameters: Example: p1 p2 p3

12. Select AD container for desktops to place in a Domain Controller computer location.

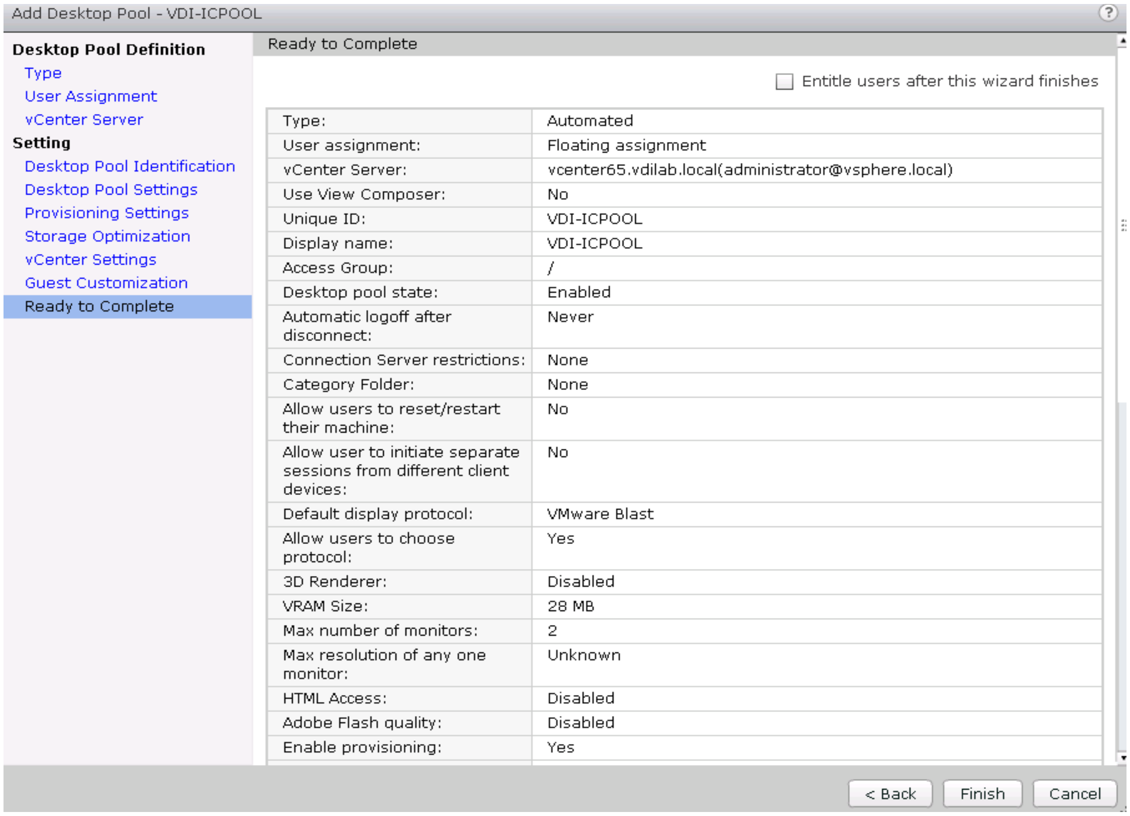

Ready to Complete

□ Entitle users after this wizard finishes

Type:	Automated
User assignment:	Floating assignment
vCenter Server:	vcenter65.vdilab.local(administrator@vsphere.local)
Use View Composer:	No
Unique ID:	VDI-ICPOOL
Display name:	VDI-ICPOOL
Access Group:	/
Desktop pool state:	Enabled
Automatic logoff after disconnect:	Never
Connection Server restrictions:	None
Category Folder:	None
Allow users to reset/restart their machine:	No
Allow user to initiate separate sessions from different client devices:	No
Default display protocol:	VMware Blast
Allow users to choose protocol:	Yes
3D Renderer:	Disabled
VRAM Size:	28 MB
Max number of monitors:	2
Max resolution of any one monitor:	Unknown
HTML Access:	Disabled
Adobe Flash quality:	Disabled
Enable provisioning:	Yes

13. Review all the deployment specifications and click Finish to complete the deployment.

14. Select Entitle users after this wizard finishes, to enable desktop user group/ users to access this pool.

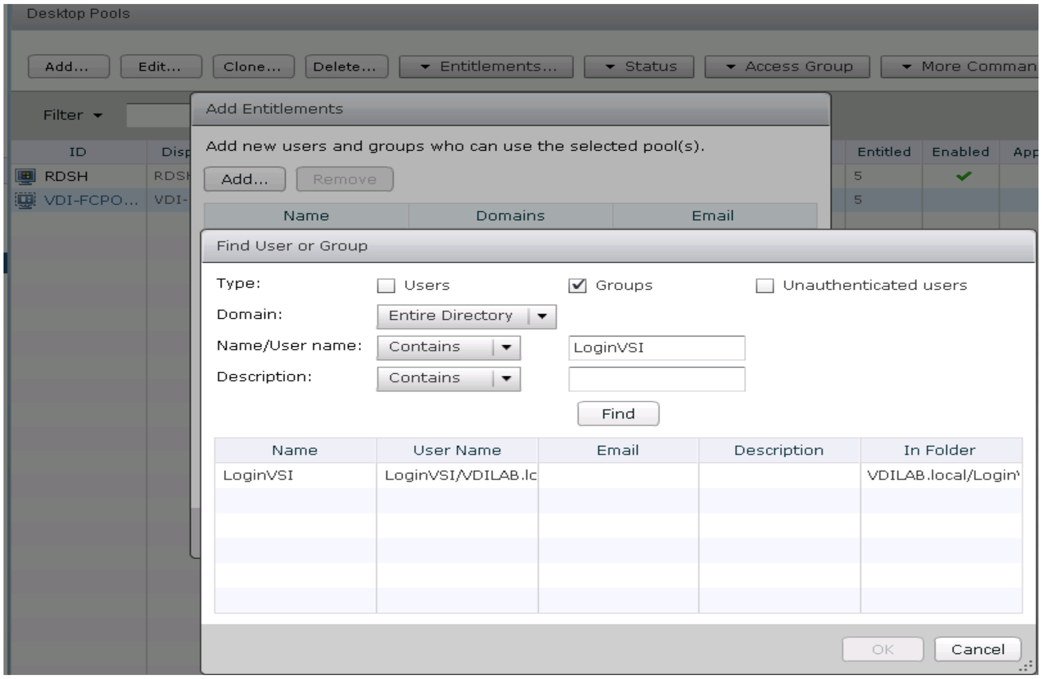
Desktop Pools
Add... Edit... Clone... Delete... Entitlements... Status Access Group More Comman
Filter
Add Entitlements
Add new users and groups who can use the selected pool(s).
Add... Remove
Find User or Group
Find
OK Cancel

ID	Disp	Entitled	Enabled	App
RDSH	RDSH	5	✓	
VDI-FCPO...	VDI-	5		

Name	Domains	Email

Type: □ Users ☑ Groups □ Unauthenticated users

Domain: Entire Directory

Name/User name: Contains LoginVSI

Description: Contains

Name	User Name	Email	Description	In Folder
LoginVSI	LoginVSI/VDILAB.lc			VDILAB.local/Login'

15. Add users group/ users to the pool:

– The computer container "LoginVSI" created on the domain controller.

– <Domain Controller> Active Directory Users and Computers > vdilab.local > LoginVSI > Computers > VDI-ICPOOL1,VDI-ICPOO L2...VDI-ICPOOL1680

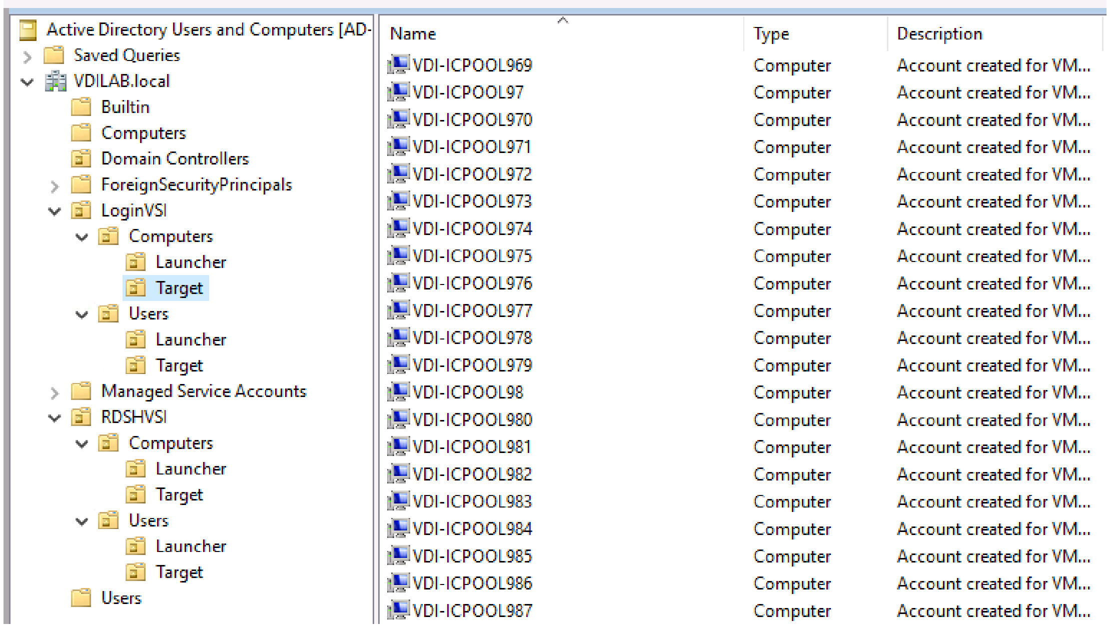

Name	Type	Description
VDI-ICPOOL969	Computer	Account created for VM...
VDI-ICPOOL97	Computer	Account created for VM...
VDI-ICPOOL970	Computer	Account created for VM...
VDI-ICPOOL971	Computer	Account created for VM...
VDI-ICPOOL972	Computer	Account created for VM...
VDI-ICPOOL973	Computer	Account created for VM...
VDI-ICPOOL974	Computer	Account created for VM...
VDI-ICPOOL975	Computer	Account created for VM...
VDI-ICPOOL976	Computer	Account created for VM...
VDI-ICPOOL977	Computer	Account created for VM...
VDI-ICPOOL978	Computer	Account created for VM...
VDI-ICPOOL979	Computer	Account created for VM...
VDI-ICPOOL98	Computer	Account created for VM...
VDI-ICPOOL980	Computer	Account created for VM...
VDI-ICPOOL981	Computer	Account created for VM...
VDI-ICPOOL982	Computer	Account created for VM...
VDI-ICPOOL983	Computer	Account created for VM...
VDI-ICPOOL984	Computer	Account created for VM...
VDI-ICPOOL985	Computer	Account created for VM...
VDI-ICPOOL986	Computer	Account created for VM...
VDI-ICPOOL987	Computer	Account created for VM...

RDSH Servers created in the RDSHVSI Computer target

16. Login to the newly provisioned VDI-ICPOOL1 desktop machine.

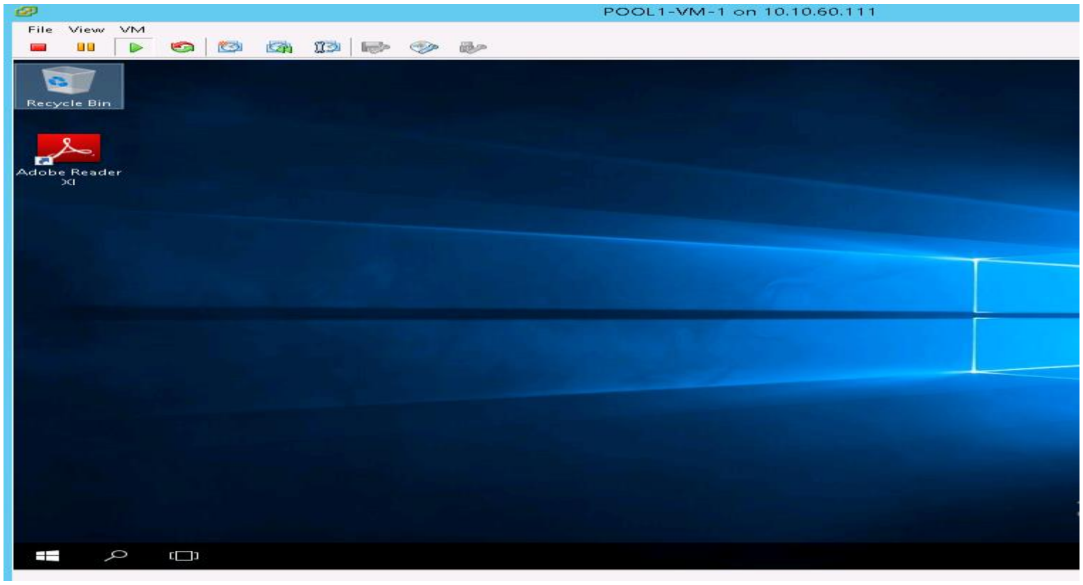

Create RDSH Farm and Pool

It is recommended to create a RDSH Farm first with specifications set for RDS Server VMs and deploying a number of RDS servers required for users. To create the RDSH Farm and Pool, complete the following steps:

1. Select the FARM when creating the RDS Pool.

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

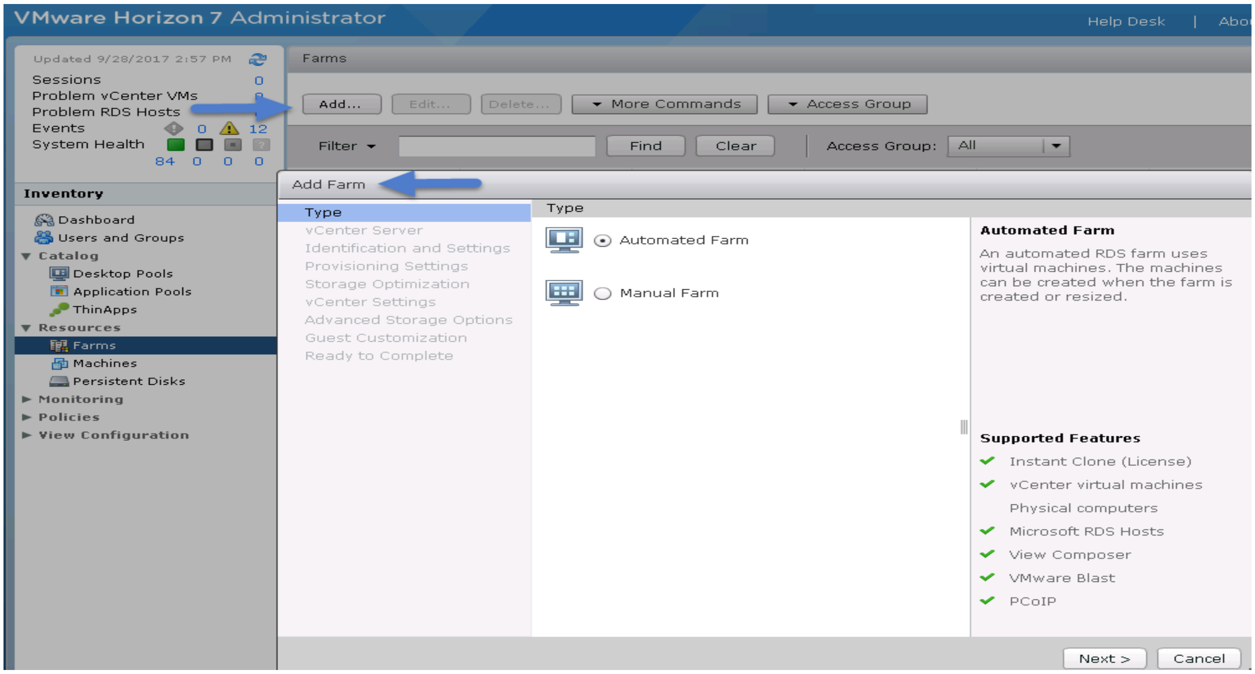

Type

- ● Automated Farm
- ○ Manual Farm

Automated Farm

An automated RDS farm uses virtual machines. The machines can be created when the farm is created or resized.

Supported Features

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

2. Click Add type of pool. We used automated pool in the study.

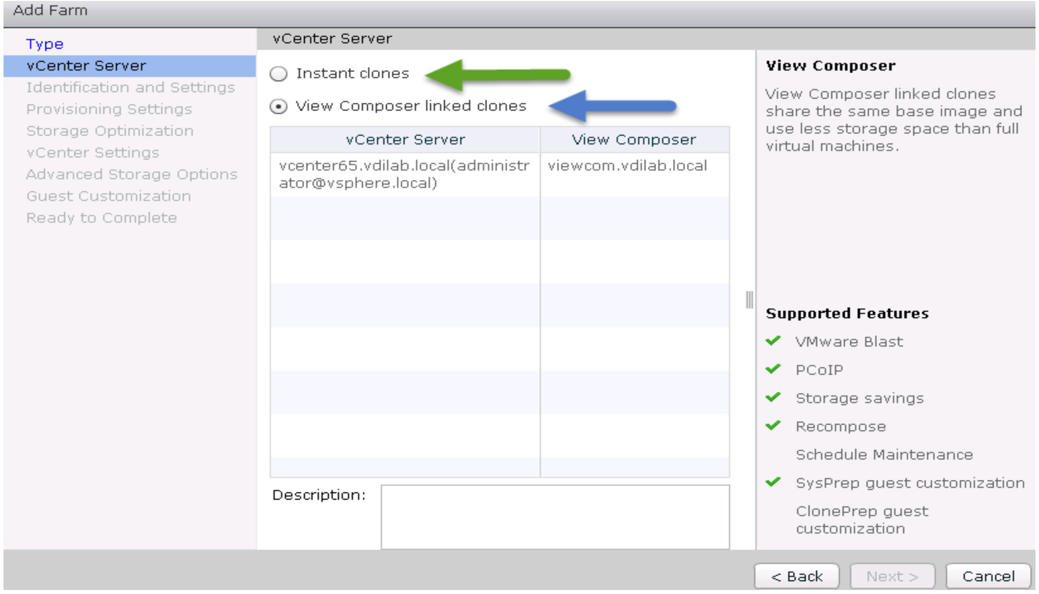

vCenter Server

○ Instant clones

● View Composer linked clones

vCenter Server	View Composer
vcenter65.vdilab.local(administrator@vsphere.local)	viewcom.vdilab.local

Description:

View Composer

View Composer linked clones share the same base image and use less storage space than full virtual machines.

Supported Features

- ✓ VMware Blast
- ✓ PCoIP
- ✓ Storage savings
- ✓ Recompose
- Schedule Maintenance
- ✓ SysPrep guest customization
- ClonePrep guest customization

3. Click Next.

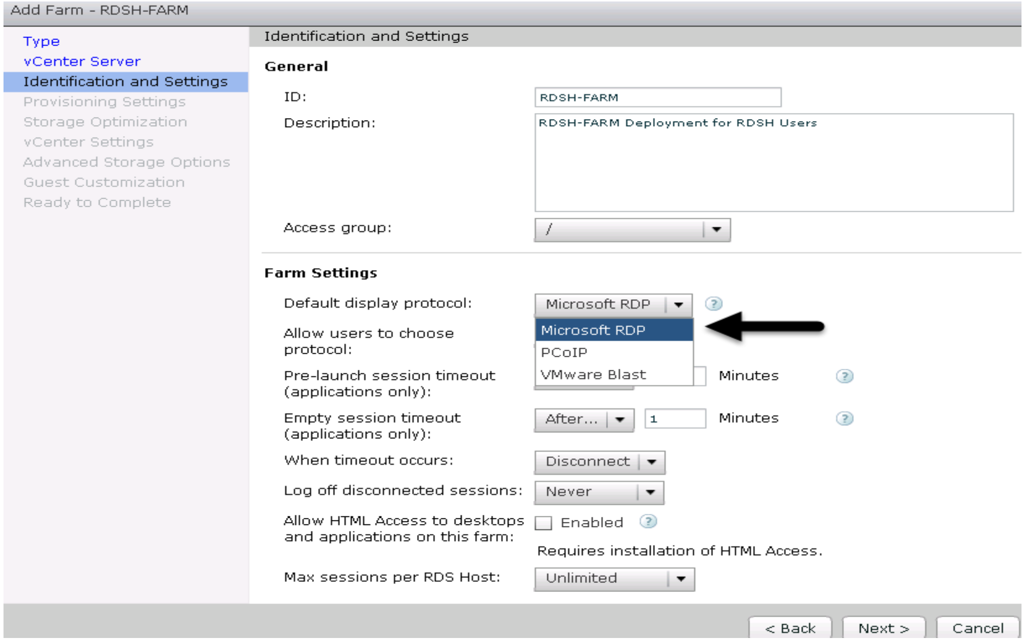

Identification and Settings

General

ID: RDSH-FARM

Description: RDSH-FARM Deployment for RDSH Users

Access group: /

Farm Settings

Default display protocol: Microsoft RDP

Allow users to choose protocol:

Pre-launch session timeout (applications only): Minutes

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

When timeout occurs: Disconnect

Log off disconnected sessions: Never

Allow HTML Access to desktops and applications on this farm: □ Enabled

Requires installation of HTML Access.

Max sessions per RDS Host: Unlimited

4. Provide ID and Description for RDS FARM. Select the Display Protocol which is required for users to connect to the RDS Sessions.

For RDS Pool, we used Microsoft RDP protocol.

5. Click Next.

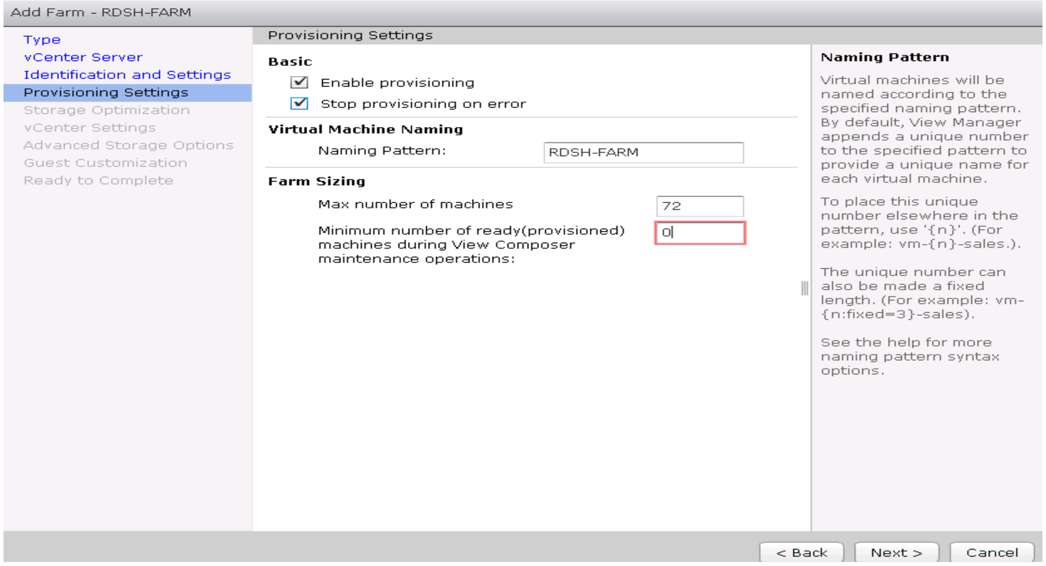

Provisioning Settings

Basic

☑ Enable provisioning

☑ Stop provisioning on error

Virtual Machine Naming

Naming Pattern: RDSH-FARM

Farm Sizing

Max number of machines: 72

Minimum number of ready(provisioned) machines during View Composer maintenance operations: 0

Naming Pattern

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

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

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

See the help for more naming pattern syntax options.

6. Provide the naming pattern for the RDS Desktops /VMs you want to create and the number of RDS VMs you want to create on the RDS host or RDS cluster. For example, 72 RDS server virtual machines created in this study.

7. Click Next.

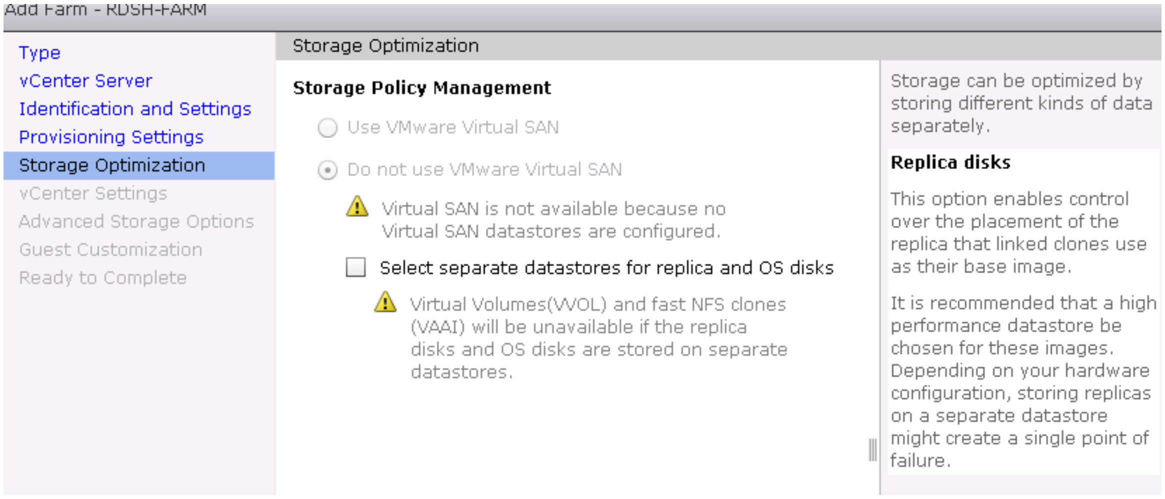

Storage Optimization

Storage Policy Management

○ Use VMware Virtual SAN

● Do not use VMware Virtual SAN

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

□ Select separate datastores for replica and OS disks

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

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

Replica disks

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

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

8. Complete the storage Optimization settings as required.

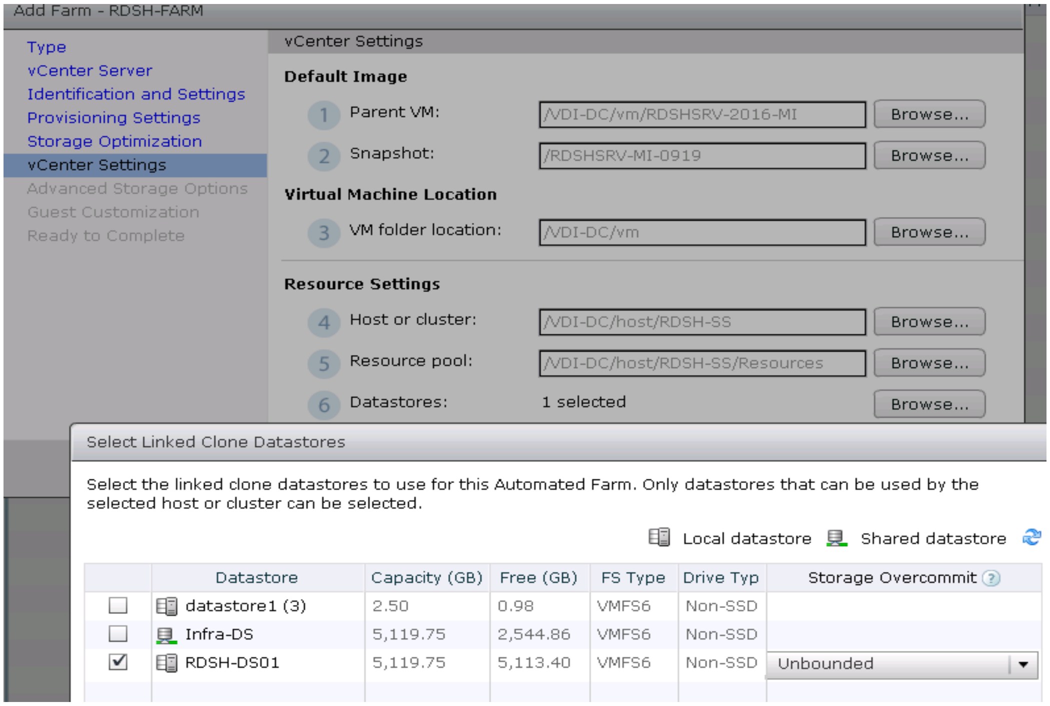

Add Farm - RDSH-FARM

vCenter Settings

Default Image

1 Parent VM: /VDI-DC/vm/RDSHSRV-2016-MI Browse...

2 Snapshot: /RDSHSRV-MI-0919 Browse...

Virtual Machine Location

3 VM folder location: /VDI-DC/vm Browse...

Resource Settings

4 Host or cluster: /VDI-DC/host/RDSH-SS Browse...

5 Resource pool: /VDI-DC/host/RDSH-SS/Resources Browse...

6 Datastores: 1 selected Browse...

Select Linked Clone Datastores

Select the linked clone datastores to use for this Automated Farm. Only datastores that can be used by the selected host or cluster can be selected.

Local datastore   Shared datastore

	Datastore	Capacity (GB)	Free (GB)	FS Type	Drive Typ	Storage Overcommit
□	datastore1 (3)	2.50	0.98	VMFS6	Non-SSD	
□	Infra-DS	5,119.75	2,544.86	VMFS6	Non-SSD	
☑	RDSH-DS01	5,119.75	5,113.40	VMFS6	Non-SSD	Unbounded

9. Provide all the information about vCenter settings with parent RDS master image, snapshot, Host /Cluster information, and Datastores for VMs to be stored.

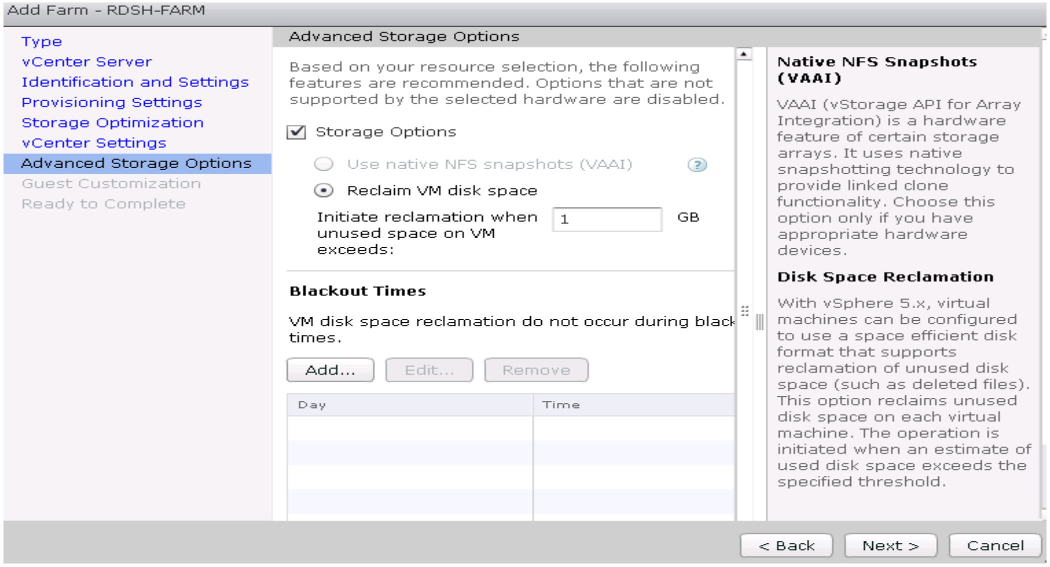

Add Farm - RDSH-FARM

Advanced Storage Options

Based on your resource selection, the following features are recommended. Options that are not supported by the selected hardware are disabled.

☑ Storage Options

○ Use native NFS snapshots (VAAI)

● Reclaim VM disk space

Initiate reclamation when unused space on VM exceeds: 1 GB

Blackout Times

VM disk space reclamation do not occur during black times.

Native NFS Snapshots (VAAI)

VAAI (vStorage API for Array Integration) is a hardware feature of certain storage arrays. It uses native snapshotting technology to provide linked clone functionality. Choose this option only if you have appropriate hardware devices.

Disk Space Reclamation

With vSphere 5.x, virtual machines can be configured to use a space efficient disk format that supports reclamation of unused disk space (such as deleted files). This option reclaims unused disk space on each virtual machine. The operation is initiated when an estimate of used disk space exceeds the specified threshold.

10. Select the required advanced storage options.

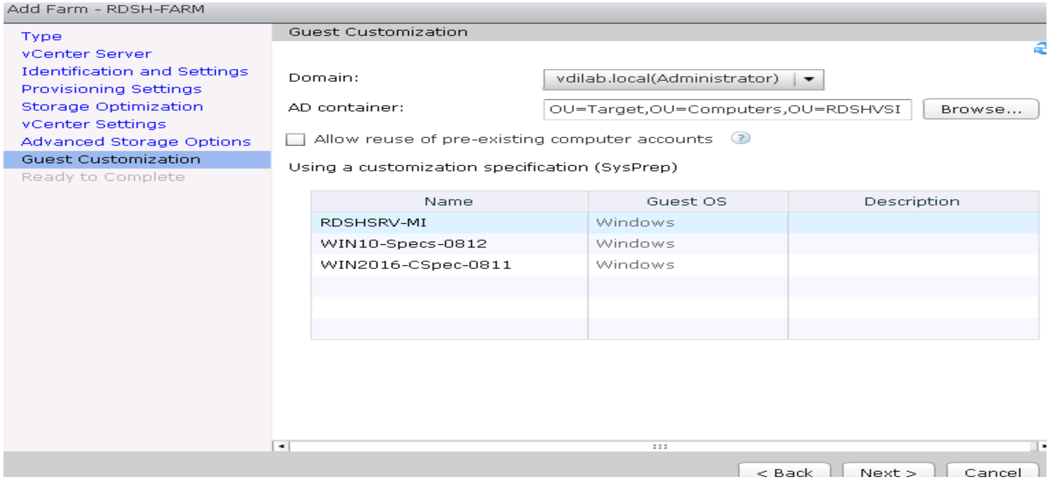

Guest Customization

Domain: vdilab.local(Administrator)

AD container: OU=Target,OU=Computers,OU=RDSHVSI

□ Allow reuse of pre-existing computer accounts

Using a customization specification (SysPrep)

Name	Guest OS	Description
RDSHSRV-MI	Windows	
WIN10-Specs-0812	Windows	
WIN2016-CSpec-0811	Windows	

11. Select Active Directory Domain controller container (VMs to be stored on the separate Computer VM (RDSHVSI) container in the Domain Controller) intended for storing RDS VMs and select the sysprep customization specs for creating VMware Composer provisioned RDS VMs.

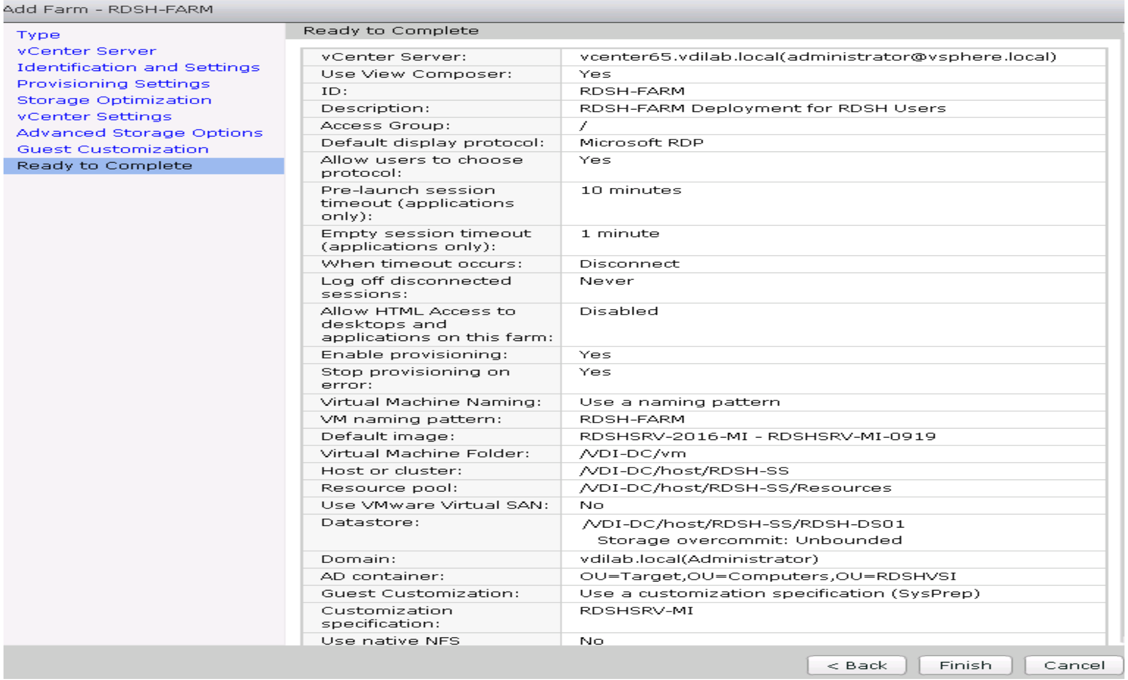

Ready to Complete

vCenter Server:	vcenter65.vdilab.local(administrator@vsphere.local)
Use View Composer:	Yes
ID:	RDSH-FARM
Description:	RDSH-FARM Deployment for RDSH Users
Access Group:	/
Default display protocol:	Microsoft RDP
Allow users to choose protocol:	Yes
Pre-launch session timeout (applications only):	10 minutes
Empty session timeout (applications only):	1 minute
When timeout occurs:	Disconnect
Log off disconnected sessions:	Never
Allow HTML Access to desktops and applications on this farm:	Disabled
Enable provisioning:	Yes
Stop provisioning on error:	Yes
Virtual Machine Naming:	Use a naming pattern
VM naming pattern:	RDSH-FARM
Default image:	RDSHSRV-2016-MI - RDSHSRV-MI-0919
Virtual Machine Folder:	/VDI-DC/vm
Host or cluster:	/VDI-DC/host/RDSH-SS
Resource pool:	/VDI-DC/host/RDSH-SS/Resources
Use VMware Virtual SAN:	No
Datastore:	/VDI-DC/host/RDSH-SS/RDSH-DS01 Storage overcommit: Unbounded
Domain:	vdilab.local(Administrator)
AD container:	OU=Target,OU=Computers,OU=RDSHVSI
Guest Customization:	Use a customization specification (SysPrep)
Customization specification:	RDSHSRV-MI
Use native NFS	No

12. Review the RDS Farm automatic deployment specifications and click Finish to complete the RDS pool.

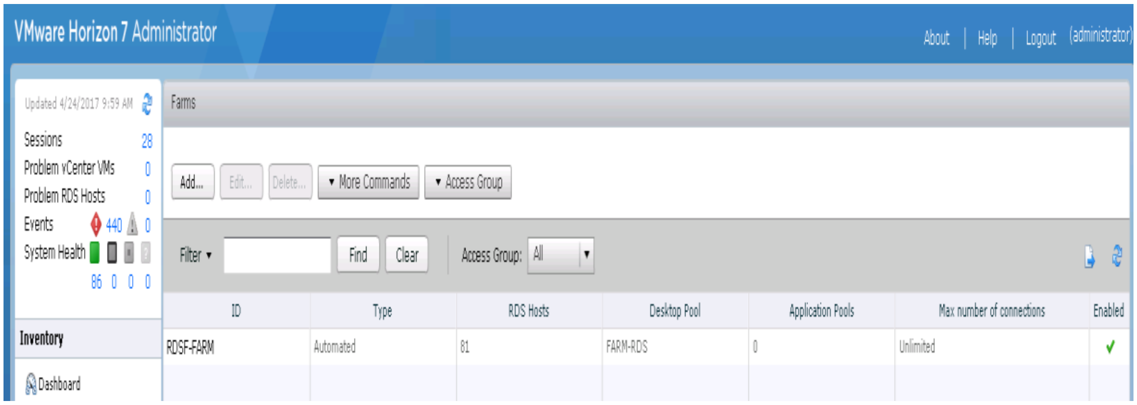

ID	Type	RDS Hosts	Desktop Pool	Application Pools	Max number of connections	Enabled
RDSF-FARM	Automated	81	FARM-RDS	0	Unlimited	✓

The RDS Farm is created in the Horizon admin console.

When the RDS FARM is created, you need to create a RDS pool to absorb the RDS VMS FARM into the Pool for further managing the RDS pool.

A Snapshot of the RDS VMs created in the AD container are selected in the deployment process.

- <Domain Controller> Active Directory Users and Computers > vdilab.local > RDSHVSI > Computers > RD1, RD2...RD-72

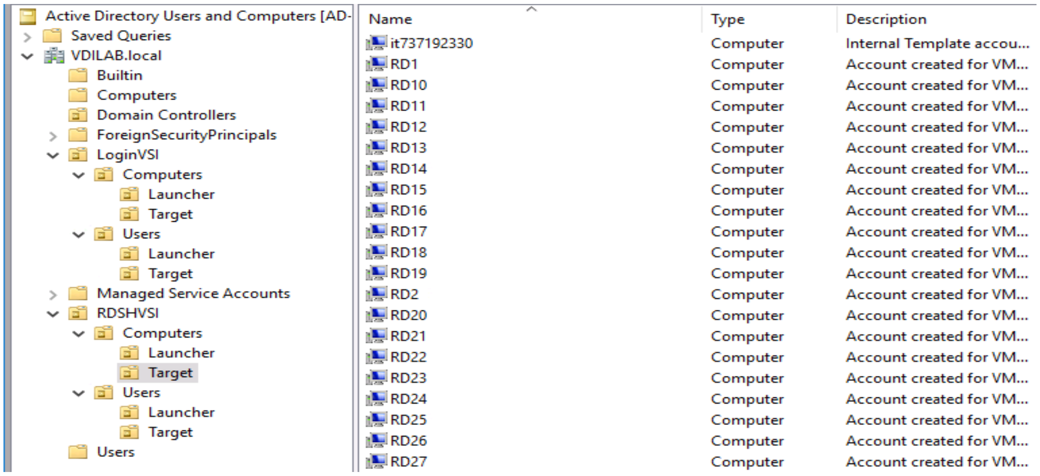

Name	Type	Description
it737192330	Computer	Internal Template accou...
RD1	Computer	Account created for VM...
RD10	Computer	Account created for VM...
RD11	Computer	Account created for VM...
RD12	Computer	Account created for VM...
RD13	Computer	Account created for VM...
RD14	Computer	Account created for VM...
RD15	Computer	Account created for VM...
RD16	Computer	Account created for VM...
RD17	Computer	Account created for VM...
RD18	Computer	Account created for VM...
RD19	Computer	Account created for VM...
RD2	Computer	Account created for VM...
RD20	Computer	Account created for VM...
RD21	Computer	Account created for VM...
RD22	Computer	Account created for VM...
RD23	Computer	Account created for VM...
RD24	Computer	Account created for VM...
RD25	Computer	Account created for VM...
RD26	Computer	Account created for VM...
RD27	Computer	Account created for VM...

Create RDS Pool

1. Click Add type of pool section on the Horizon Administrator Console and the default choice is RDS Desktop Pool.

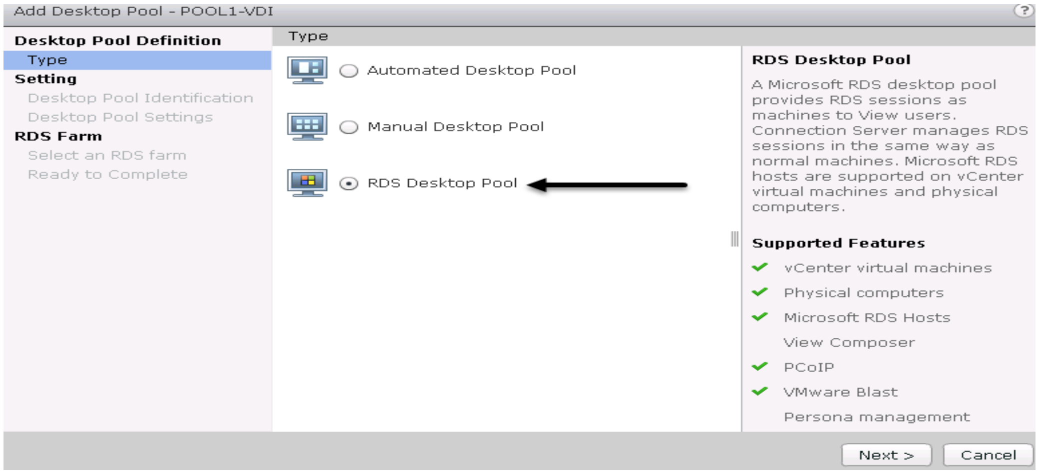

Add Desktop Pool - POOL1-VDI

RDS Desktop Pool

A Microsoft RDS desktop pool provides RDS sessions as machines to View users. Connection Server manages RDS sessions in the same way as normal machines. Microsoft RDS hosts are supported on vCenter virtual machines and physical computers.

Supported Features

- ✓ vCenter virtual machines
- ✓ Physical computers
- ✓ Microsoft RDS Hosts
- View Composer
- ✓ PCoIP
- ✓ VMware Blast
- Persona management

2. Provide ID and Display Name for the Pool. Click Next.

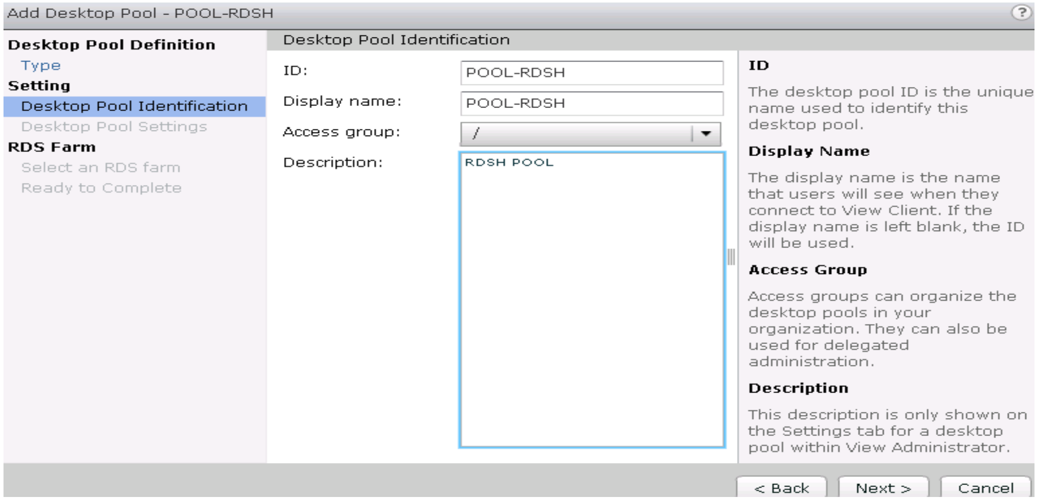

Add Desktop Pool - POOL-RDSH

ID: POOL-RDSH

Display name: POOL-RDSH

Access group: /

Description: RDSH POOL

ID

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

Display Name

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

Access Group

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

Description

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

3. Leave default settings for the Desktop Pool Settings. Click Next.

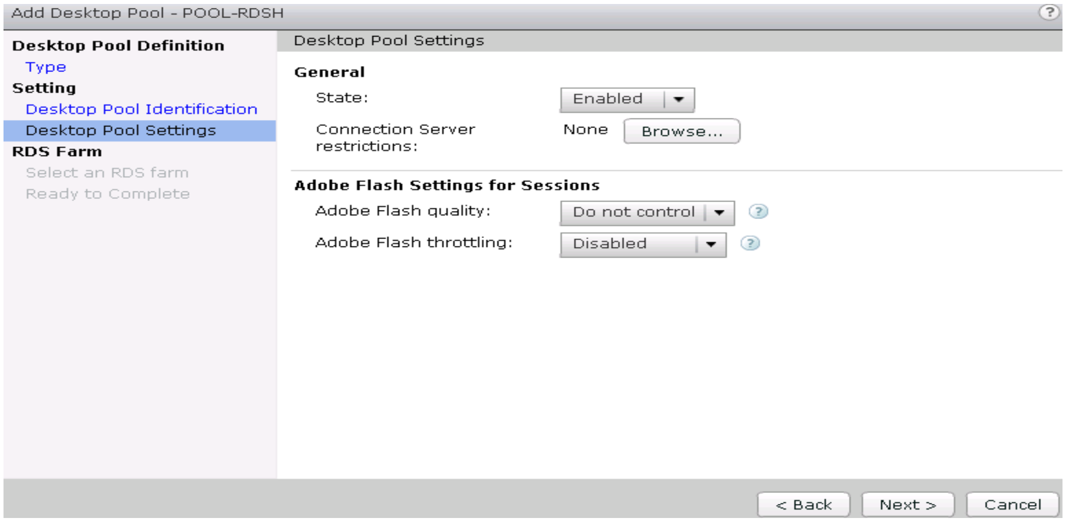

Add Desktop Pool - POOL-RDSH

General

State: Enabled

Connection Server restrictions: None

Adobe Flash Settings for Sessions

Adobe Flash quality: Do not control

Adobe Flash throttling: Disabled

4. Select the RDS Farm. You have the option to create a farm or select the farm which is already created. We chose the option of selecting an RDS farm (FA-RDS) for this desktop pool. Click Next.

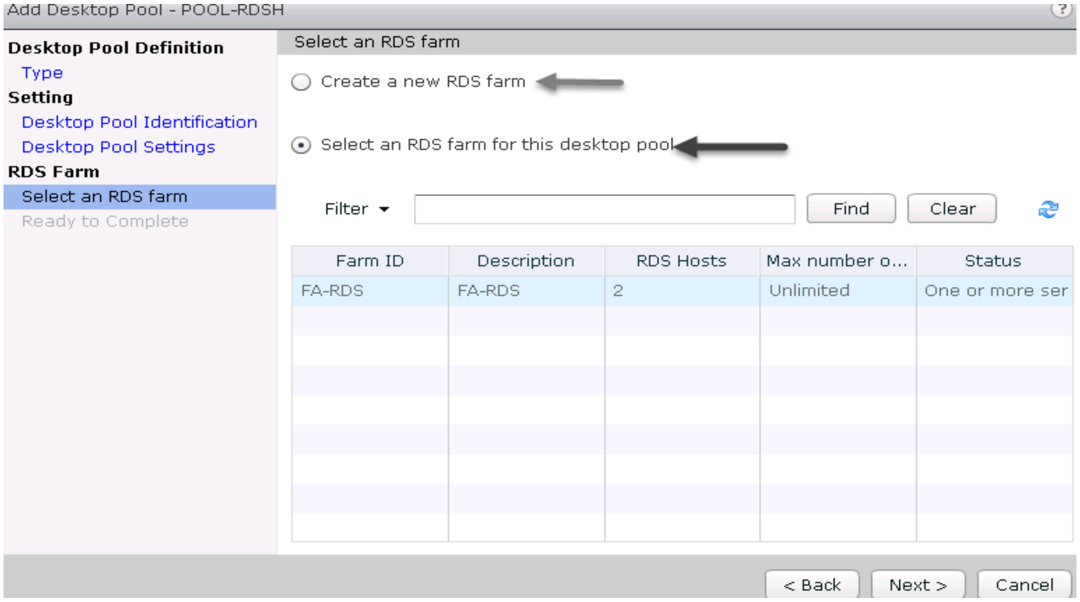

○ Create a new RDS farm

● Select an RDS farm for this desktop pool

Farm ID	Description	RDS Hosts	Max number o...	Status
FA-RDS	FA-RDS	2	Unlimited	One or more ser

5. Click Ready to Complete and select Entitle users after this wizard finishes to provide users/user group permission to access this RDS pool. Click Next.

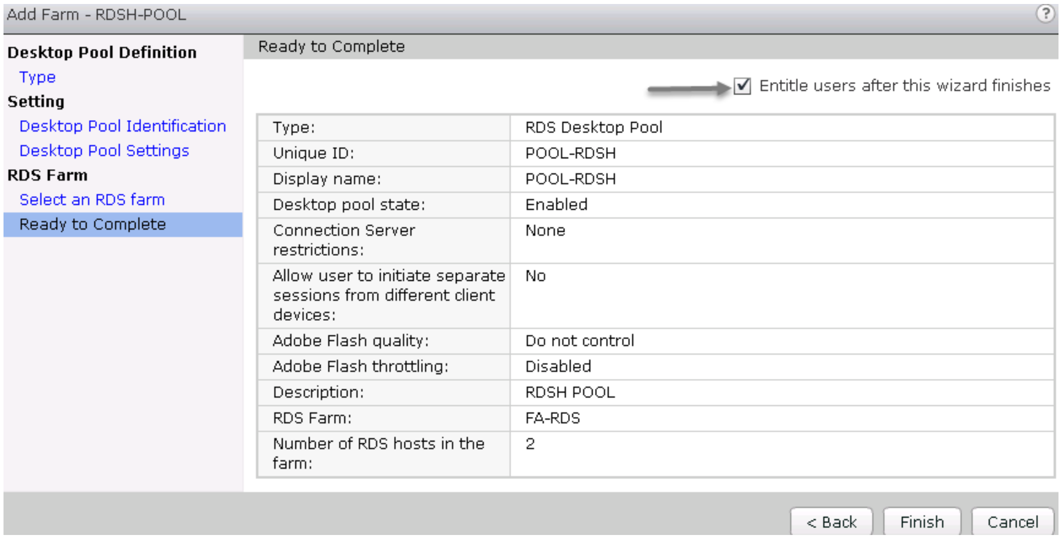

☑ Entitle users after this wizard finishes

Type:	RDS Desktop Pool
Unique ID:	POOL-RDSH
Display name:	POOL-RDSH
Desktop pool state:	Enabled
Connection Server restrictions:	None
Allow user to initiate separate sessions from different client devices:	No
Adobe Flash quality:	Do not control
Adobe Flash throttling:	Disabled
Description:	RDSH POOL
RDS Farm:	FA-RDS
Number of RDS hosts in the farm:	2

6. Select the users who can access this pool.

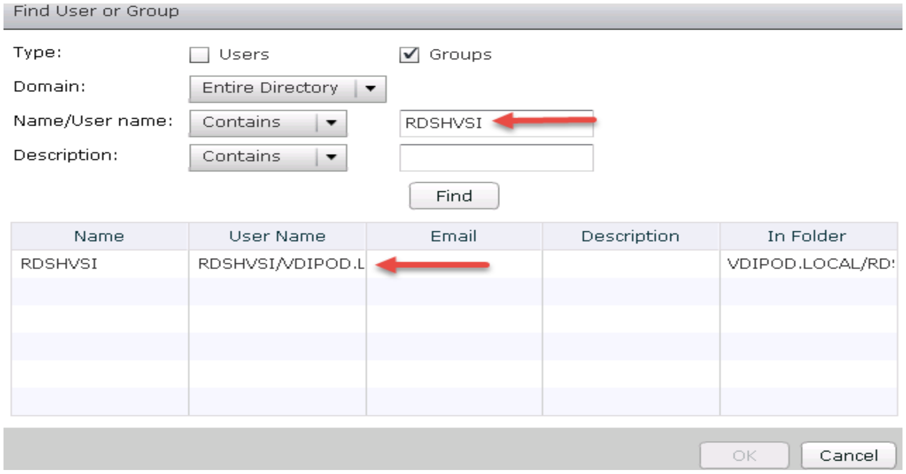

Find User or Group

Type:	□ Users	☑ Groups
Domain:	Entire Directory	
Name/User name:	Contains	RDSHVSI
Description:	Contains	

Name	User Name	Email	Description	In Folder
RDSHVSI	RDSHVSI/VDIPOD.L			VDIPOD.LOCAL/RD:

The RDS Pool is created on the Horizon Administrator console.

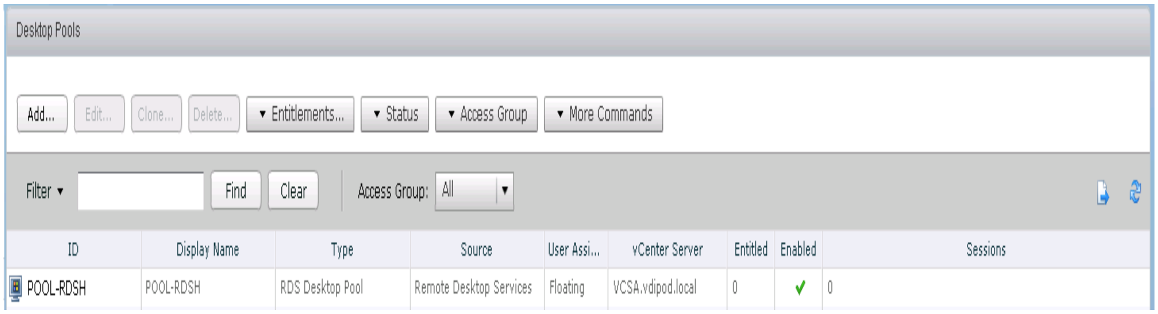

Desktop Pools

Access Group: All

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	Sessions
POOL-RDSH	POOL-RDSH	RDS Desktop Pool	Remote Desktop Services	Floating	VCSA.vdipod.local	0	✓	0

Configuring User Profile Management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting
- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for VMware Horizon desktop deployments. Screenshots of the User

Profile Management interfaces that establish policies for this CVD's RDS and VDI users (for testing purposes) are shown below. Basic profile management policy settings are documented here:

Single Server Testing

Figure 32 Cisco UCS B200 M5 Blade Server for Single Server Scalability VMware Horizon 7 Remote Desktop Server Hosted Sessions (RDSH) with Windows Server 2016

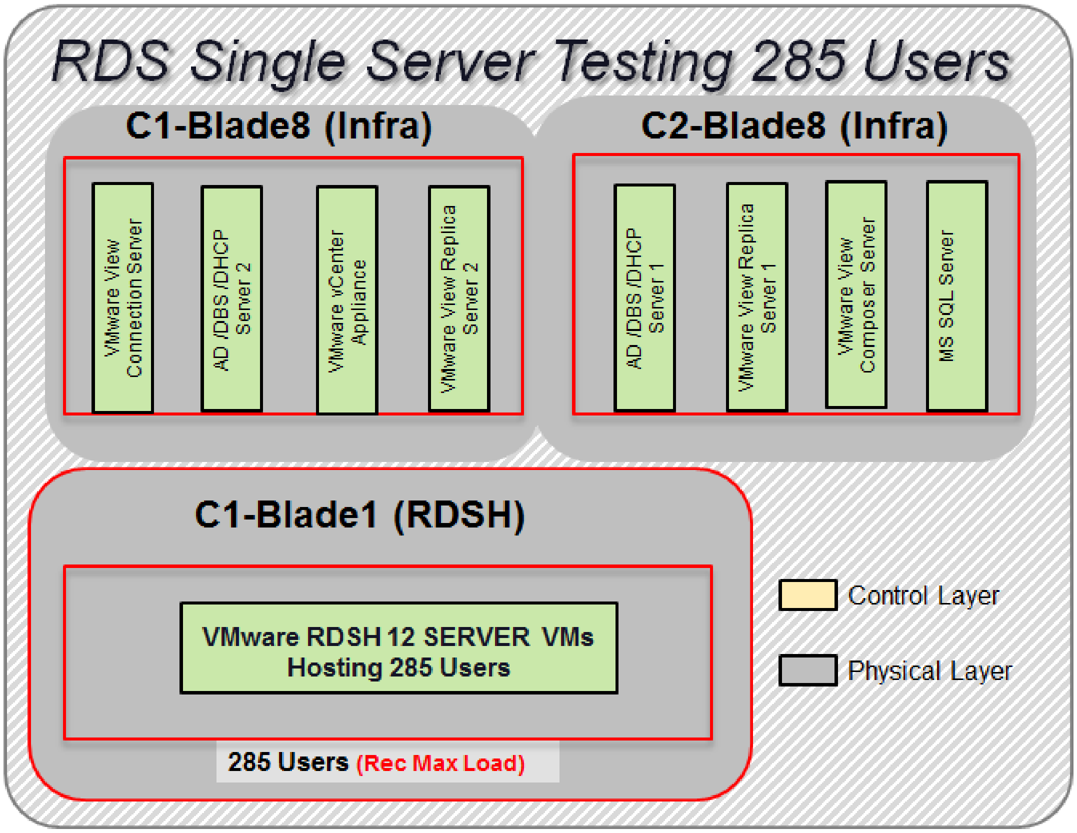
RDS Single Server Testing 285 Users
C1-Blade8 (Infra)
VMware View Connection Server
AD /DBS /DHCP Server 2
VMware vCenter Appliance
VMware View Replica Server 2
C2-Blade8 (Infra)
AD /DBS /DHCP Server 1
VMware View Replica Server 1
VMware View Composer Server
MS SQL Server
C1-Blade1 (RDSH)
VMware RDSH 12 SERVER VMs Hosting 285 Users
285 Users (Rec Max Load)
Control Layer
Physical Layer

Figure 33 Cisco UCS B200 M5 Blade Server for Single Server Scalability VMware Horizon 7 VDI (Non-Persistent) Instant Clones with Windows 10 64bit OS

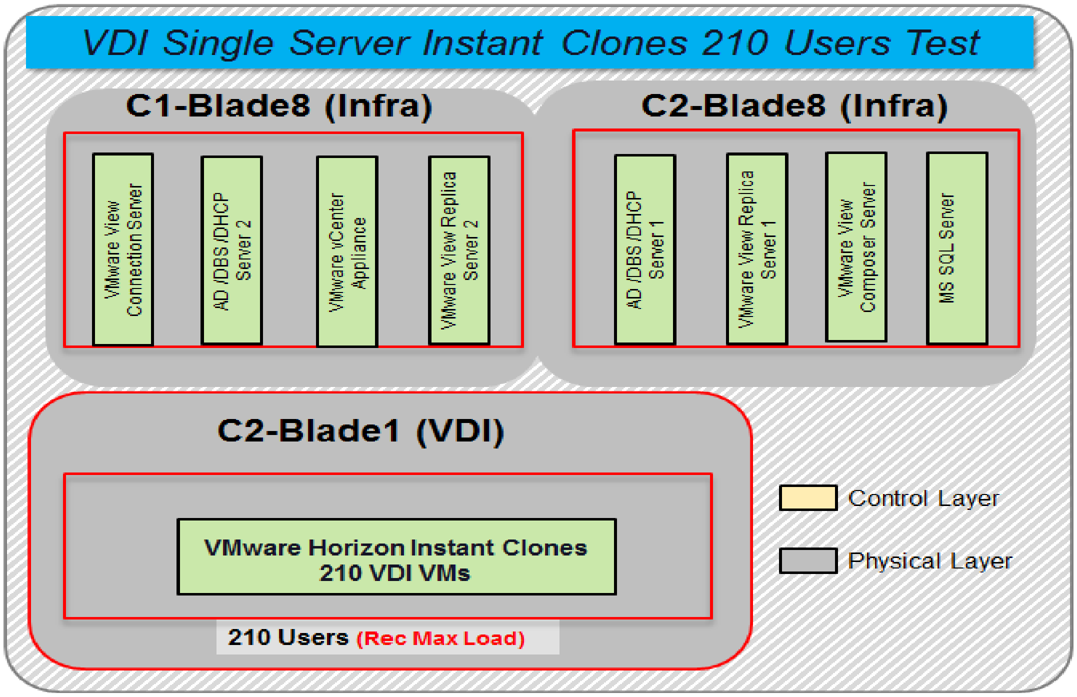
VDI Single Server Instant Clones 210 Users Test
C1-Blade8 (Infra)
VMware View Connection Server
AD /DBS /DHCP Server 2
VMware vCenter Appliance
VMware View Replica Server 2
C2-Blade8 (Infra)
AD /DBS /DHCP Server 1
VMware View Replica Server 1
VMware View Composer Server
MS SQL Server
C2-Blade1 (VDI)
VMware Horizon Instant Clones 210 VDI VMs
210 Users (Rec Max Load)
Control Layer
Physical Layer

Figure 34 Cisco UCS B200 M5 Blade Server for Single Server Scalability VMware Horizon 7 VDI (Persistent) Full Clones with Windows 10 64bit OS

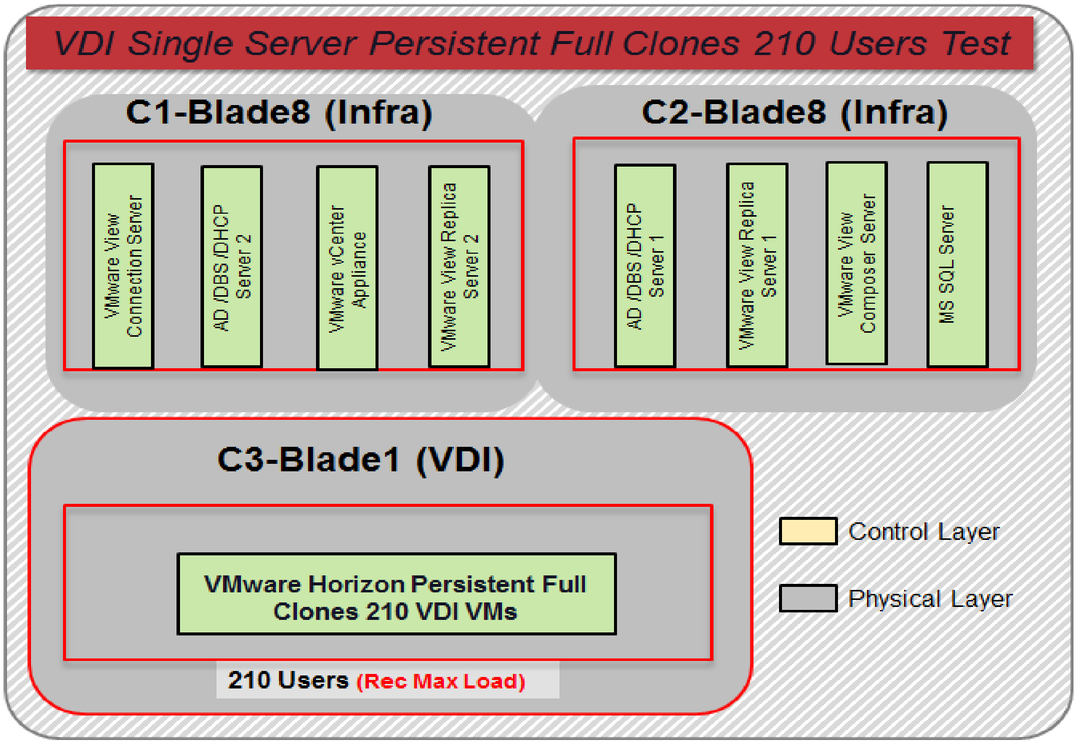
VDI Single Server Persistent Full Clones 210 Users Test
C1-Blade8 (Infra)
VMware View Connection Server
AD /DBS /DHCP Server 2
VMware vCenter Appliance
VMware View Replica Server 2
C2-Blade8 (Infra)
AD /DBS /DHCP Server 1
VMware View Replica Server 1
VMware View Composer Server
MS SQL Server
C3-Blade1 (VDI)
VMware Horizon Persistent Full Clones 210 VDI VMs
210 Users (Rec Max Load)
Control Layer
Physical Layer

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332-16UP Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.60 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz]) for infrastructure host blades
- Total number of VMware Horizon RDS server configured 12 and session 285
- 1 (one) Cisco UCS B200 M5 Blade Server (2 Intel Xeon processor 6140 Gold CPUs at 2.3 GHz, with 768GB of memory per blade server [12 x 64 GB DIMMs at 2666 MHz] for workload host blade

And

- Total number of VMware Horizon VDI Virtual Machines 210 (Non Persistent Instant Clones and Persistent Full Clones Single Server Testing)
- 1 (one) Cisco UCS B200 M5 Blade Server (2 Intel Xeon processor 6140 Gold CPUs at 2.3 GHz, with 768GB of memory per blade server [12 x 64GB DIMMs at 2666 MHz] for workload host blade
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Switches
- 1 NetApp AFF A300Storage with (24x 3.8TB) providing 65 TB Usable disk capacity

Software components:

- Cisco UCS firmware 3.2(1d)
- NetApp ONTAP 9.1
- VMware ESXi 6.5 Update 1 for host blades
- VMware Horizon RDS Hosted Shared Desktops or VMware Horizon 7 VDI Hosted Virtual Desktops
- VMware Horizon Instant Clones 7
- VMware Horizon Composer Persistent Clones 7
- Microsoft SQL Server 2016
- Microsoft Windows Server 2016, 6vCPU, 24GB RAM, 40 GB base disk for each RDS Server VM
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB disk Instant Clones Testing
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB disk for Persistent Clones Testing
- Microsoft Office 2016
- Login VSI 4.1.25.6 Knowledge Worker Workload

Cisco UCS Configuration for Cluster Testing

This test case validates two workload clusters using VMware Horizon 7 with 1680 RDS Hosted Server Sessions and 3320 VDI Instant Clone non- persistent and full clone persistent virtual machines. Server N+1 fault tolerance is factored into this test scenario for each workload and infrastructure cluster.

Figure 35 RDS Cluster Test Configuration with Seven Blades

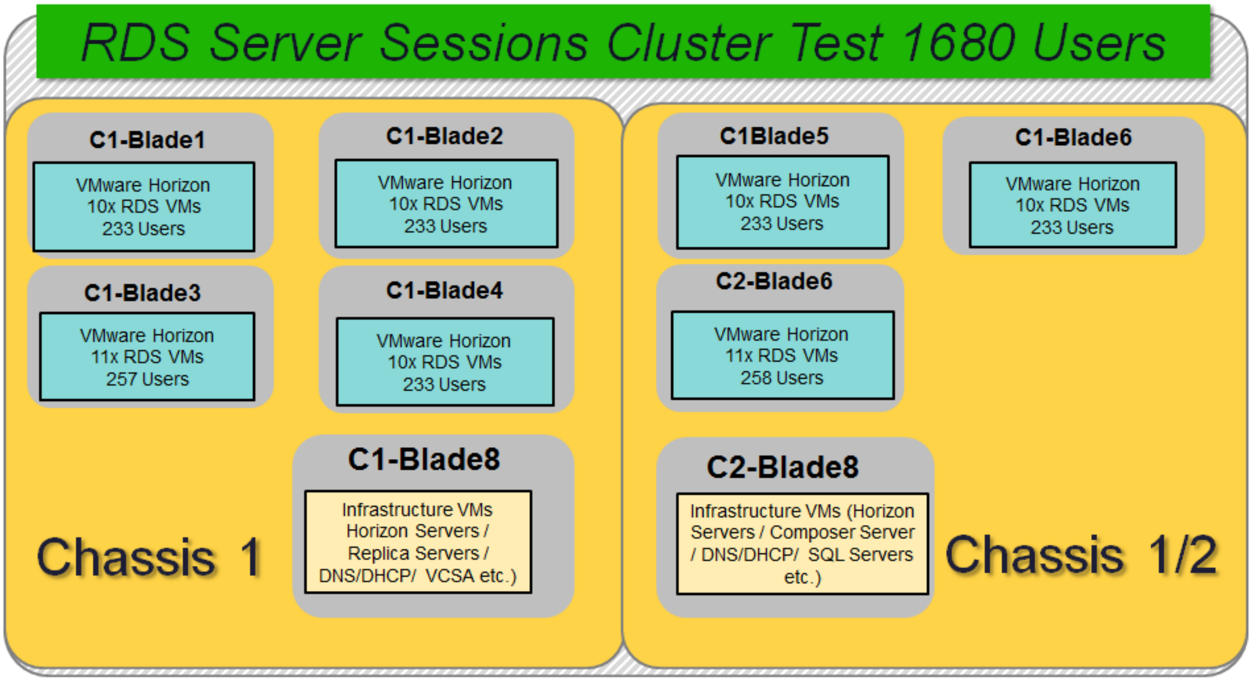
RDS Server Sessions Cluster Test 1680 Users
C1-Blade1
VMware Horizon 10x RDS VMs 233 Users
C1-Blade2
VMware Horizon 10x RDS VMs 233 Users
C1Blade5
VMware Horizon 10x RDS VMs 233 Users
C1-Blade6
VMware Horizon 10x RDS VMs 233 Users
C1-Blade3
VMware Horizon 11x RDS VMs 257 Users
C1-Blade4
VMware Horizon 10x RDS VMs 233 Users
C2-Blade6
VMware Horizon 11x RDS VMs 258 Users
C1-Blade8
Infrastructure VMs Horizon Servers / Replica Servers / DNS/DHCP/ VCSA etc.)
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Server / DNS/DHCP/ SQL Servers etc.)
Chassis 1
Chassis 1/2

Figure 36 VDI Cluster Test with VDI Instant Clone Non-Persistent Cluster Test Configuration with Nine Blades

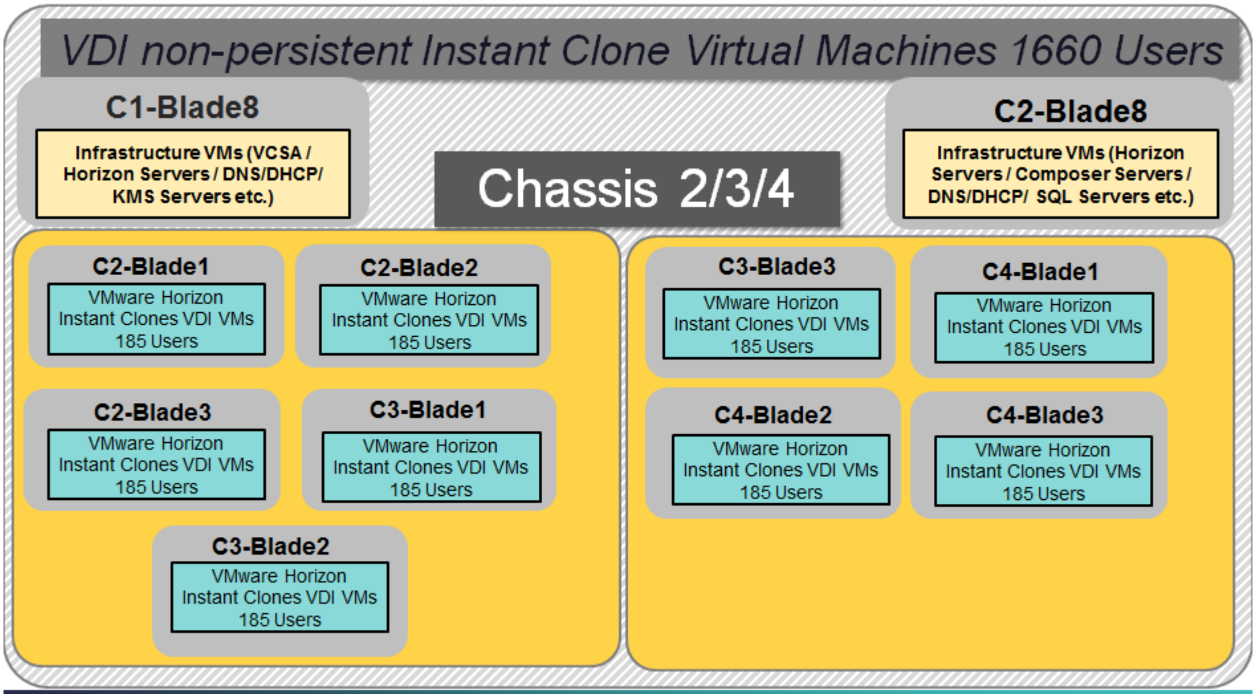
VDI non-persistent Instant Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users

Figure 37 VDI Cluster Test with VDI Full Clone Persistent Cluster Test Configuration with Nine Blades

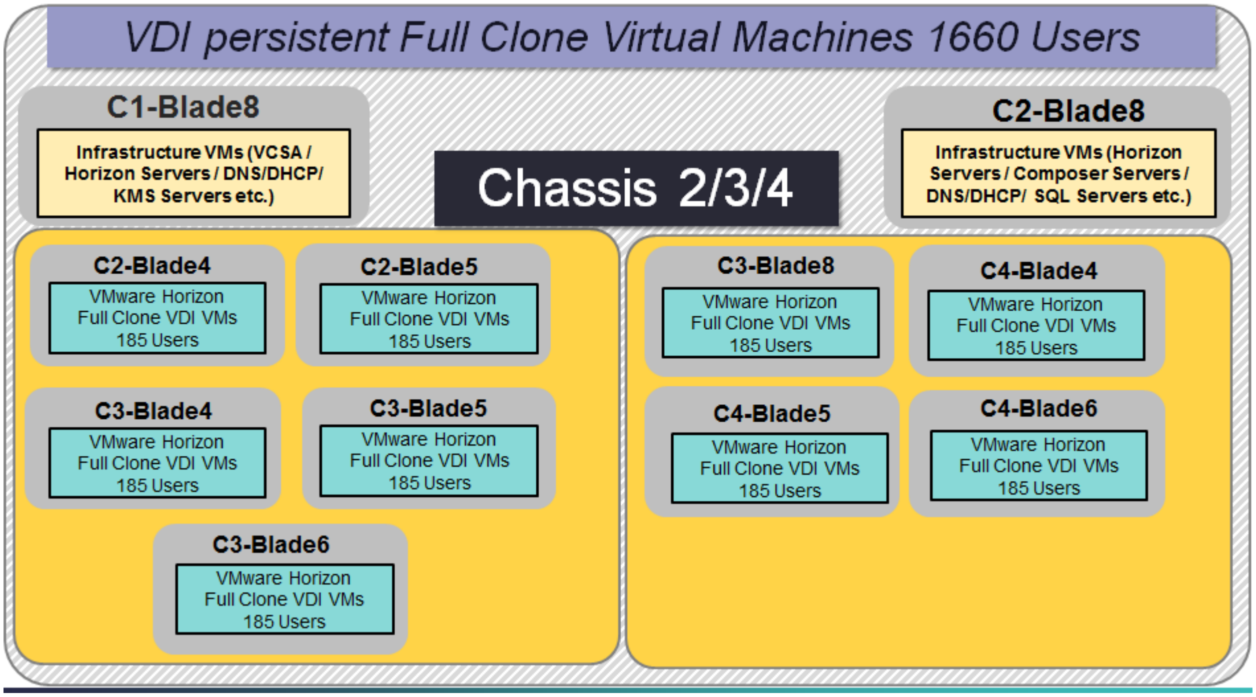
VDI persistent Full Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C2-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade8
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade6
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade6
VMware Horizon Full Clone VDI VMs 185 Users

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332-16UP Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.60 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz]) for infrastructure host blades
- Seven (7) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [12 x 64GB DIMMs at 2666 MHz]) for workload host blades
- Total of 72 VMware Horizon RDS Server VMs configured on 7 Hosts RDSH Cluster
- Total number of RDS Hosted Server Sessions 1680

or

- Nine (9) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [12 x 64GB DIMMs at 2666 MHz]) for workload host blades
- Total no of VDI Virtual Machines Configured 1660 on 1 VMware Horizon Instant Clone pool
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Switches
- 1 NetApp AFF A300Storage with (24x 3.8TB) 65 TB Usable disk capacity

And

- Nine (9) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [12 X 64 GB DIMMs at 2666 MHz] for workload host blades
- Total no of VDI Virtual Machines Configured 1660 on 1 VMware Horizon Persistent Clone pool
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Switches
- 1 NetApp AFF A300Storage with (24x3.8TB) 65 TB Usable disk capacity

Software components:

- Cisco UCS firmware 3.2(1d)
- NetApp ONTAP 9.1
- VMware ESXi 6.5 Update 1 for host blades
- VMware Horizon RDS Hosted Shared Desktops or VMware Horizon 7 VDI Hosted Virtual Desktops
- VMware Horizon Composer Server 7
- Microsoft SQL Server 2016
- Microsoft Windows Server 2012 R2, 6vCPU, 24GB RAM, 40 GB disk
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32 GB disk (32GB Instant cones/ 32GB full clones)
- Microsoft Office 2016
- Login VSI 4.1.25.6 Knowledge Worker Workload

Cisco UCS Configuration for Full Scale Testing

This test case validates thirty blades mixed workloads using VMware Horizon 7 with 1680 RDS Hosted sessions and 3320 VDI non-persistent virtual desktops 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 38 Full Scale Test Configuration with Thirty Blades

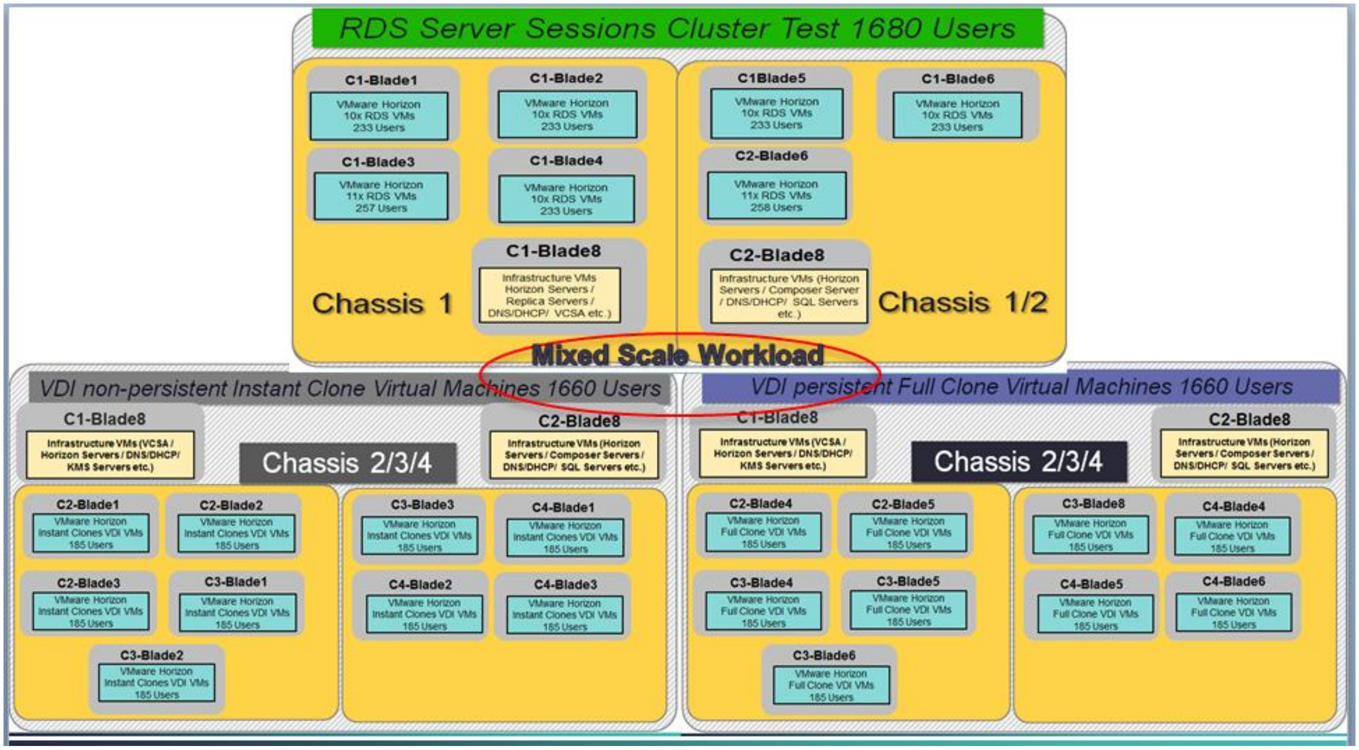
RDS Server Sessions Cluster Test 1680 Users
C1-Blade1
VMware Horizon 10x RDS VMs 233 Users
C1-Blade2
VMware Horizon 10x RDS VMs 233 Users
C1Blade5
VMware Horizon 10x RDS VMs 233 Users
C1-Blade6
VMware Horizon 10x RDS VMs 233 Users
C1-Blade3
VMware Horizon 11x RDS VMs 257 Users
C1-Blade4
VMware Horizon 10x RDS VMs 233 Users
C2-Blade6
VMware Horizon 11x RDS VMs 258 Users
C1-Blade8
Infrastructure VMs Horizon Servers / Replica Servers / DNS/DHCP/ VCSA etc.)
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Server / DNS/DHCP/ SQL Servers etc.)
Chassis 1
Chassis 1/2
Mixed Scale Workload
VDI non-persistent Instant Clone Virtual Machines 1660 Users
VDI persistent Full Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C2-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade8
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade6
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade6
VMware Horizon Full Clone VDI VMs 185 Users

Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6332-16UP Fabric Interconnects
- 2 (Infrastructure Hosts) Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.60 GHz, with 128GB of memory per blade server [16 GB x 8 DIMMs at 2400 MHz] for infrastructure host blades
- Seven (7) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [12 X 64 GB DIMMs at 2666 MHz] for workload host blades
- Total of 72 VMware Horizon RDS Server VMs configured on 7 ESXi Hosts on the RDSH Cluster
- Total number of RDS Hosted Server Sessions 1680

And

- Nine (9) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [12 x 64 GB DIMMs at 2666 MHz] for workload host blades
- Total number of VDI Virtual Machines Configured 1660 VMware Horizon Instant Clone pool
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Switches

- 1 NetApp AFF A300Storage with (24x 3.8) 65 TB Usable disk capacity

And

- Nine (9) Cisco UCS B200 M5 Blade Server (2 Intel Xeon 6140 Gold processor CPUs at 2.3 GHz, with 768GB of memory per blade server [64 GB x 12 DIMMs at 2666 MHz]) for workload host blades
- Total number of VDI Virtual Machines Configured 1660 VMWare Horizon Full Clone pool
- Cisco VIC 1340 CNA (1 per blade)
- 2 Cisco Nexus 9372PX Access Switches
- 2 Cisco MDS 9148S Fibre Channel Switches
- 1 NetApp AFF A300Storage with (24x 3.8TB) 65 Usable disk capacity

Software components:

- Cisco UCS firmware 3.2(1d)
- NetApp ONTAP 9.1
- VMware ESXi 6.5 Update 1 for host blades
- VMware Horizon RDS Hosted Shared Desktops or VMware Horizon 7 VDI Hosted Virtual Desktops
- VMware Horizon Composer Server 7
- Microsoft SQL Server 2016
- Microsoft Windows Server 2012 R2, 6vCPU, 24GB RAM, 40 GB disk
- Microsoft Windows 10 64 bit, 2vCPU, 2 GB RAM, 32/32 GB disk
- Microsoft Office 2016
- Login VSI 4.1.25.6 Knowledge Worker Workload

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 VMware Horizon 7 Administrator 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 Blade servers used in test run
   - All Infrastructure VMs used in test run (AD, SQL, Horizon Connection brokers, Horizon Composer, etc.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot RDS Machines using VMware Horizon 7 Administrator Console.
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 VMware Horizon 7 Administrator Console 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 4.1.25.6 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 VMware Horizon 7 Administrator Dashboard 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 VMware Horizon Administrator Console or Horizon Connection Server Console 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 FlexPod with Cisco UCS B200 M4 and VMware Horizon 7 on VMware ESXi 6.5 Update 1 Test Results.

The purpose of this testing is to provide the data needed to validate VMware Horizon Remote Desktop Session Hosted (RDSH) server sessions and VMware Horizon Virtual Desktop (VDI) models with VMware Horizon Composer 7 using ESXi, vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2016 sessions on Cisco UCS B200 M5 Blade Servers using a NetApp AFF A300 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware Horizon 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. When 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 39 Sample of a VSI Max Response Time Graph, Representing a Normal Test

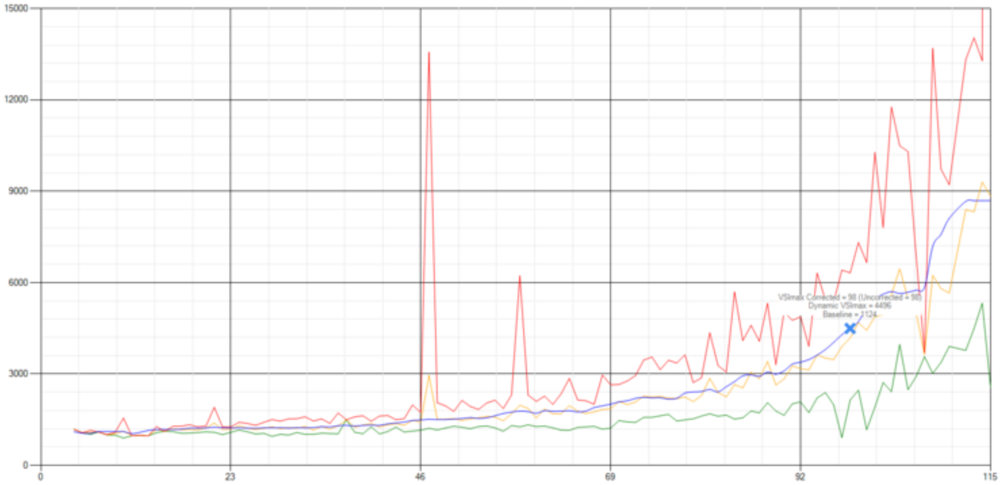

Figure 40 Sample of a VSI Test Response Time Graph (where there was a clear performance issue)

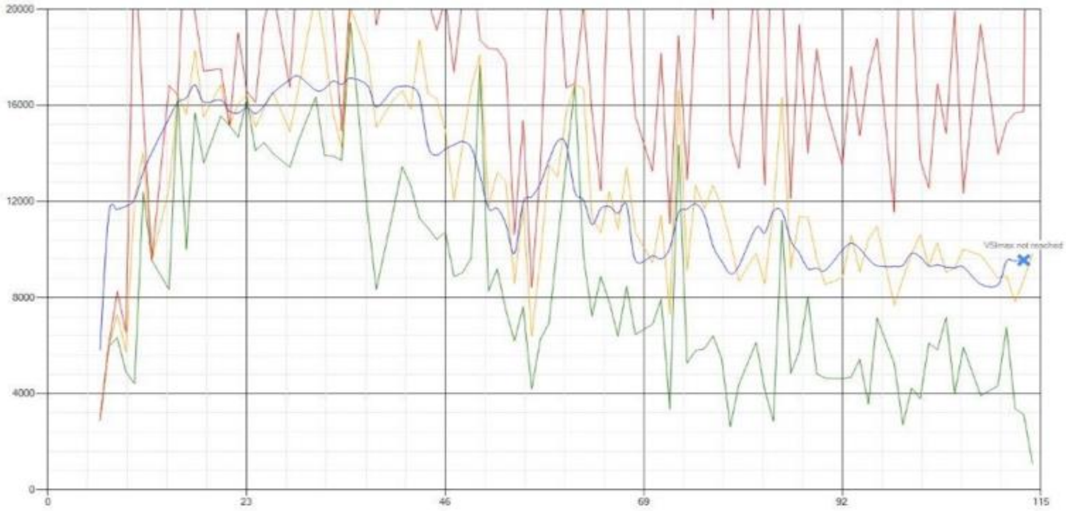

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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 base phase 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 Persistent Testing

For both the VMware Horizon 7 RDS Hosted Virtual Desktops and VDI virtual machines 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 1680 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

Test Phase	Description
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 provides 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: 285 RDS Hosted server sessions and 210 VDI non persistent and 210 VDI persistent desktops.

Single-Server Recommended Maximum Workload for RDS Persistent Hosted Server Sessions: 285 Users

Figure 41 Single Server Recommended Maximum Workload for RDS with 285 Users

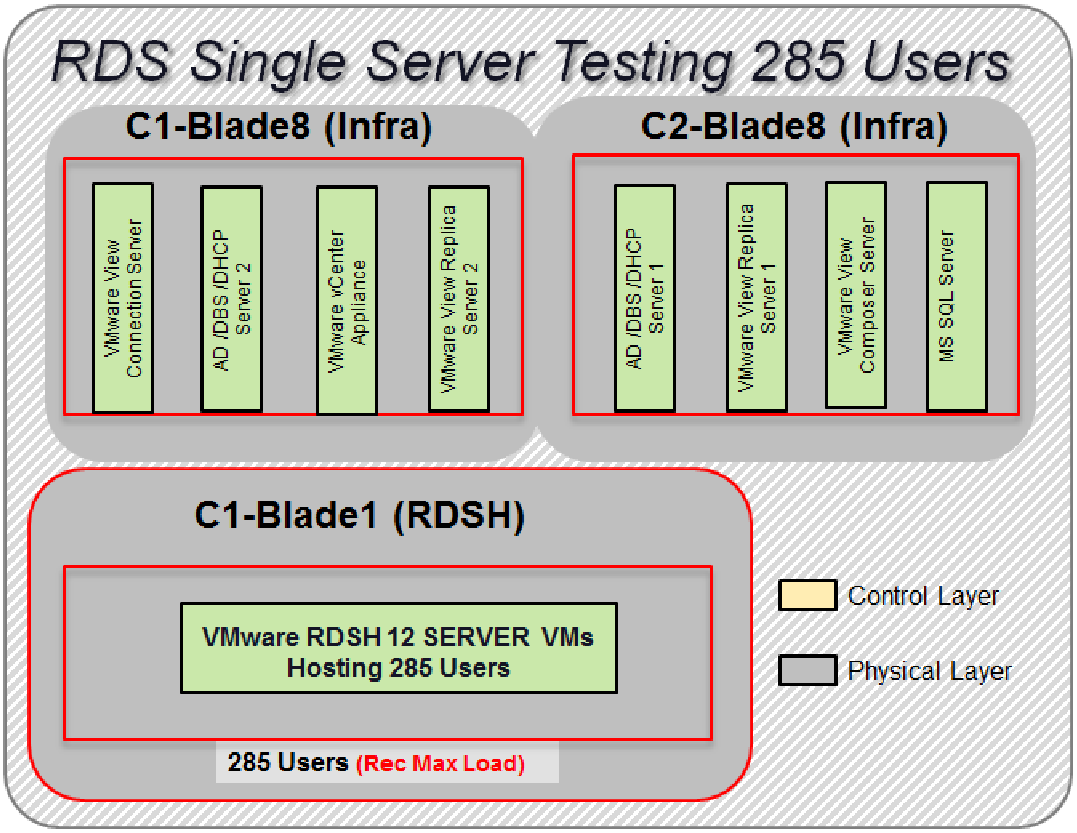
RDS Single Server Testing 285 Users
C1-Blade8 (Infra)
VMware View Connection Server
AD /DBS /DHCP Server 2
VMware vCenter Appliance
VMware View Replica Server 2
C2-Blade8 (Infra)
AD /DBS /DHCP Server 1
VMware View Replica Server 1
VMware View Composer Server
MS SQL Server
C1-Blade1 (RDSH)
VMware RDSH 12 SERVER VMs Hosting 285 Users
285 Users (Rec Max Load)
Control Layer
Physical Layer

The recommended maximum workload for a B200 M5 blade server with dual 6140 Gold processors and 768GB of RAM is 285 Server 2016 Remote Desktop Server Hosted Session Desktops. Each dedicated blade server ran 12 Server 2016 Virtual Machines. Each virtual server was configured with 6 vCPUs and 24GB RAM.

Figure 42 Single Server | VMware Horizon 7 RDS Hosted Sessions | VSI Score

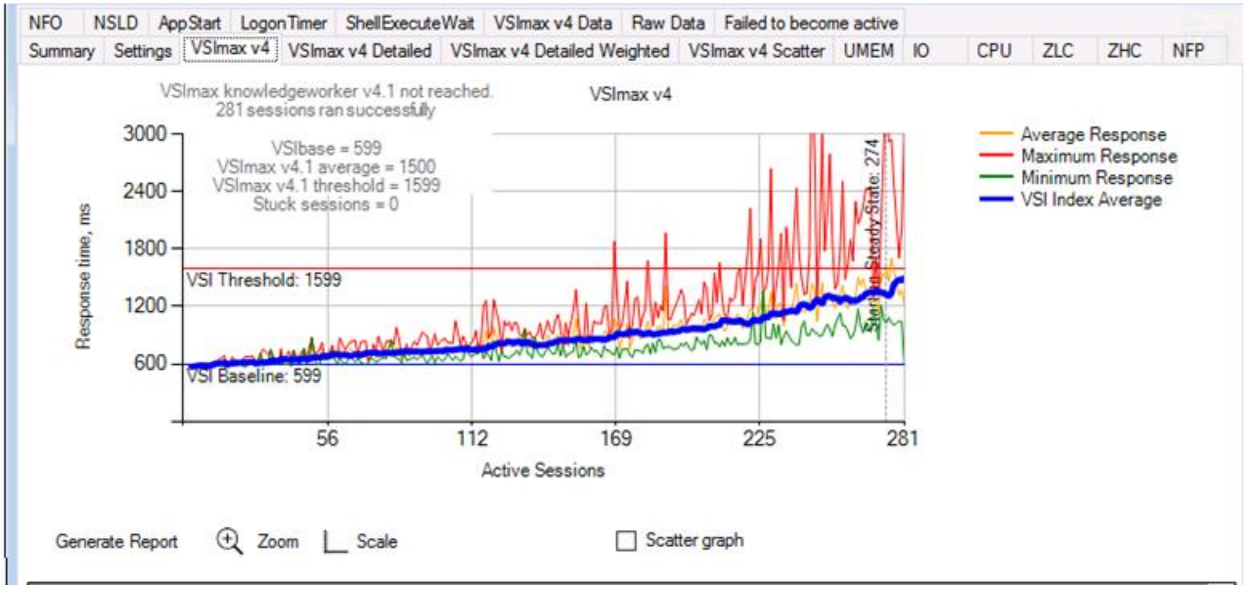

Performance data for the server running the workload is shown below:

Figure 43 Single Server | VMware Horizon 7 RDSH processor | Host CPU Utilization

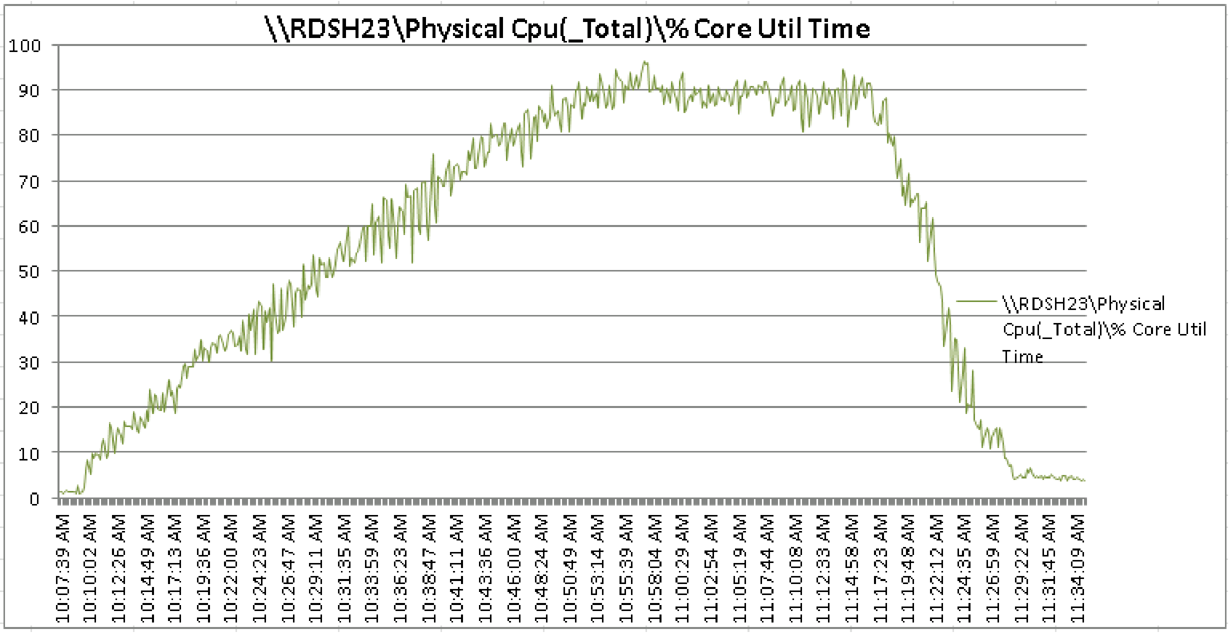

Figure 44 Single Server | VMware Horizon 7 RDSH | Host Memory Utilization

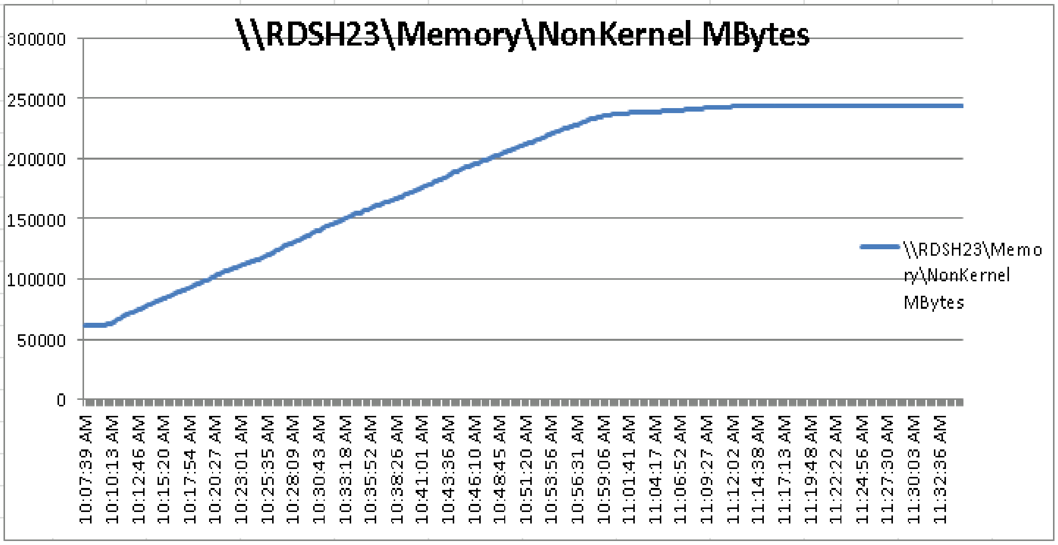

Figure 45 Single Server | VMware Horizon 7 RDSH | Host Network Utilization

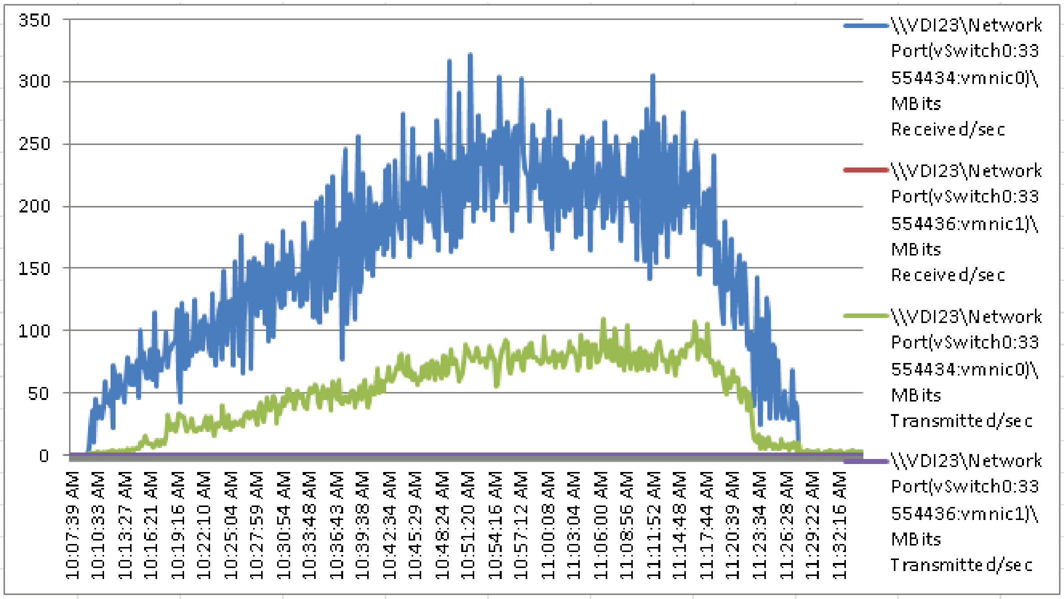

Figure 46 Single Server | VMware Horizon 7 RDSH processor (Total%) Processor Time

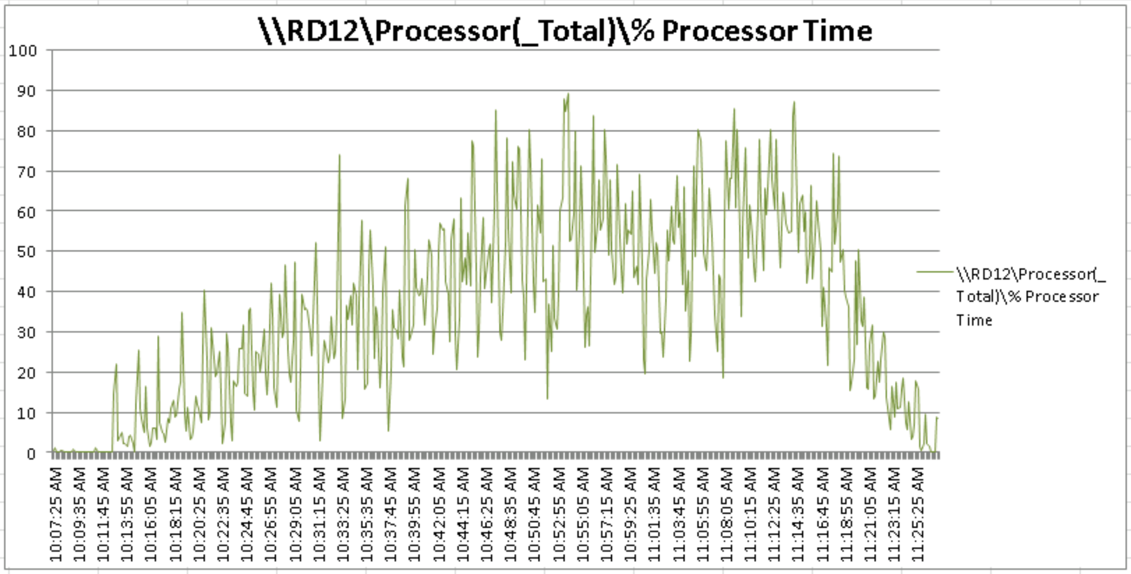

Figure 47 Single Server | VMware Horizon 7 RDSH processor (Total%) User Time

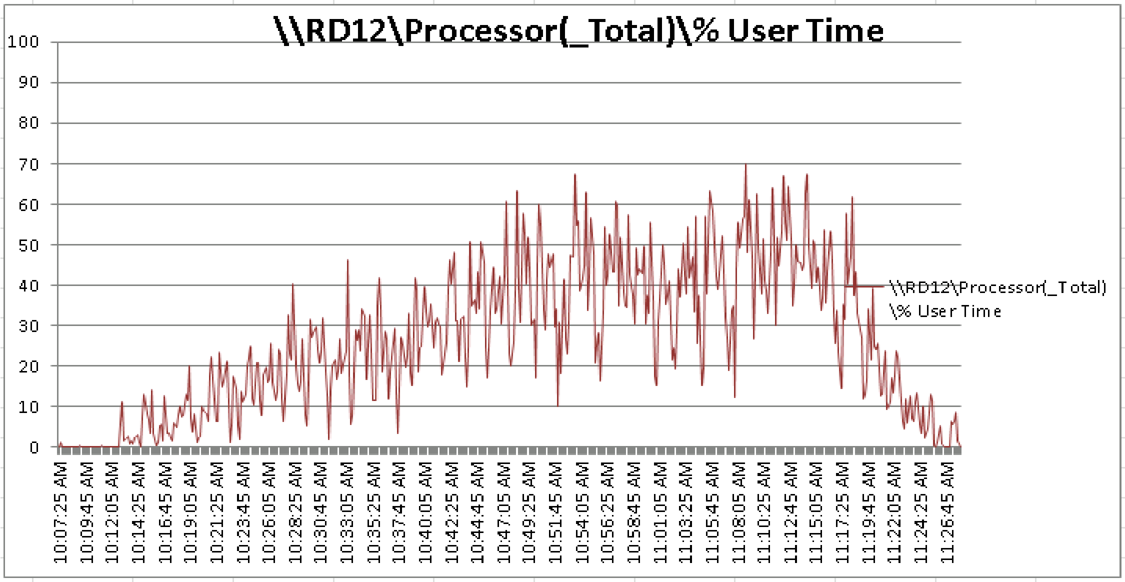

Single-Server Recommended Maximum Workload for VDI Persistent and Non-Persistent with 210 Users

Figure 48 Single Server Recommended Maximum Workload for VDI Non-Persistent with 210 Users

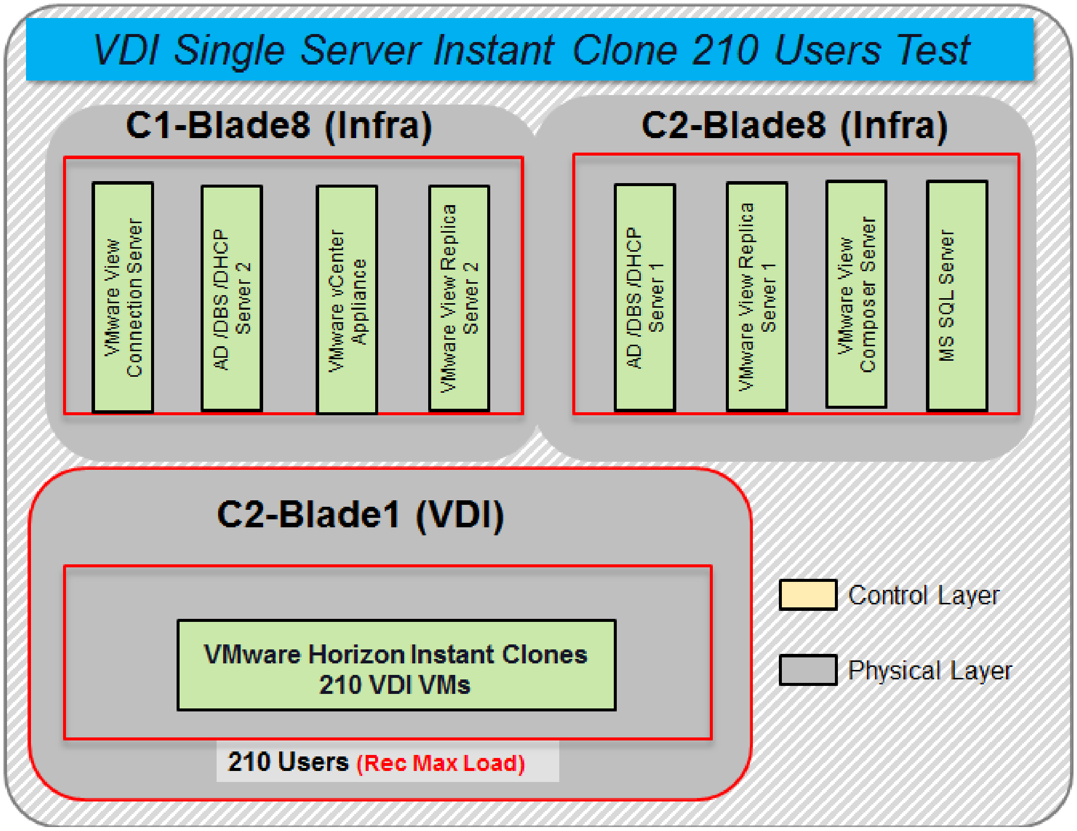

The recommended maximum workload for a Cisco UCS B200 M5 blade server with dual 6140 processors and 768GB of RAM is 210 Windows 10 64-bit virtual machines with 2 vCPU and 2GB RAM. The Login VSI and blade performance data is shown below.

Figure 49 Single Server | VMware Horizon 7 VDI-Non Persistent | VSI Score

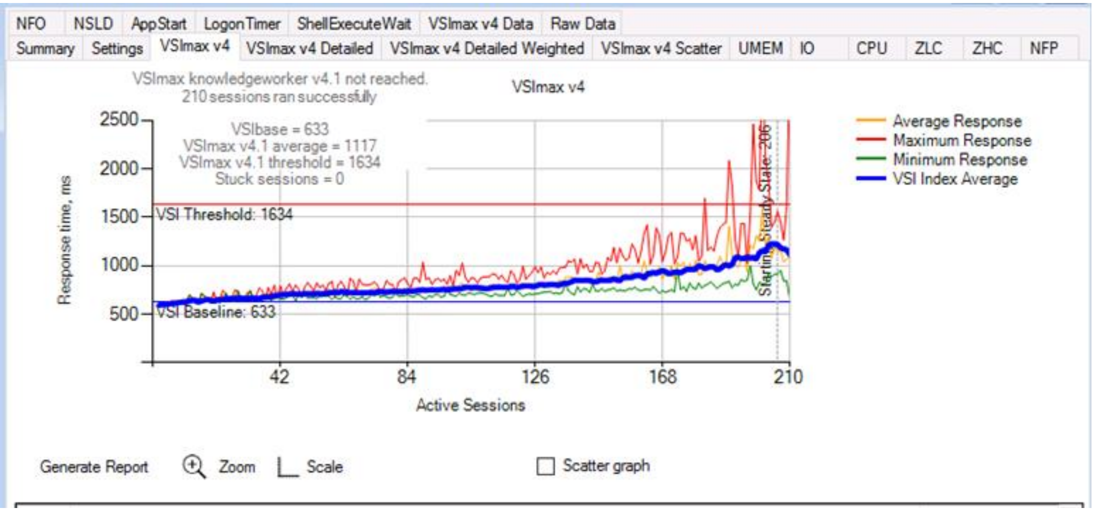

VSImax knowledgeworker v4.1 not reached.
210 sessions ran successfully

VSIbase = 633
VSImax v4.1 average = 1117
VSImax v4.1 threshold = 1634
Stuck sessions = 0

Performance data for the server running the workload is shown below:

Figure 50 Single Server | VMware Horizon VDI Non-Persistent Desktops | Host CPU Utilization

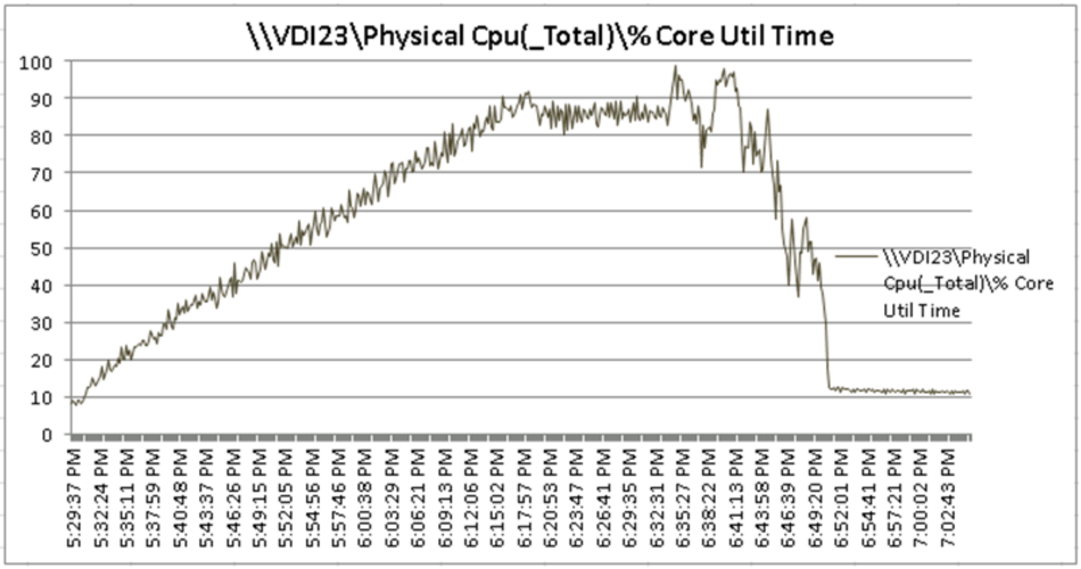

Figure 51 Single Server | VMware Horizon VDI Non -Persistent Desktops | Host Memory Utilization

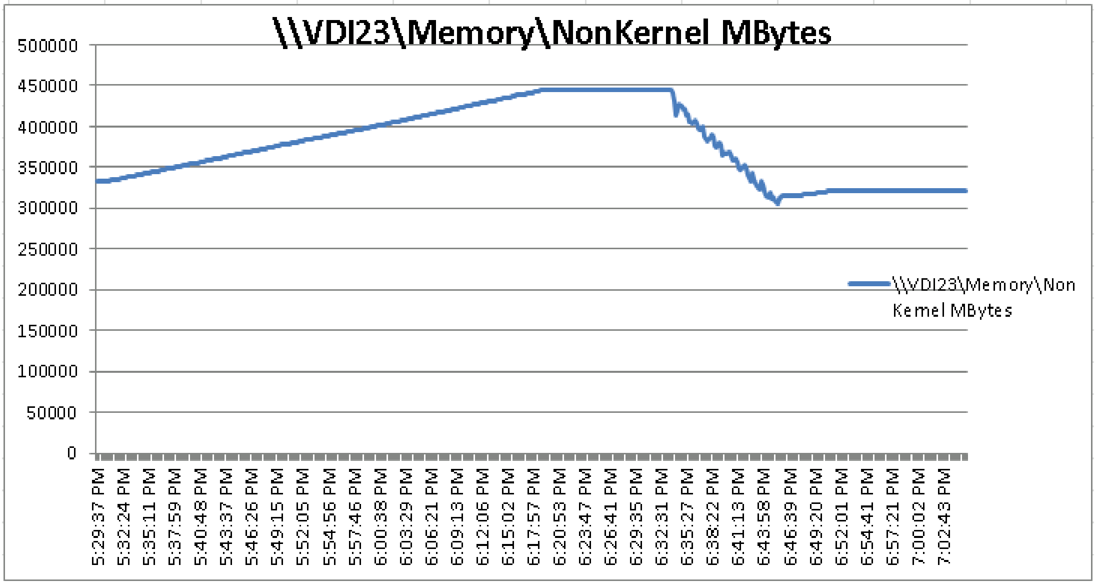

Figure 52 Single Server | VMware Horizon VDI non -persistent desktops | Host Network Utilization

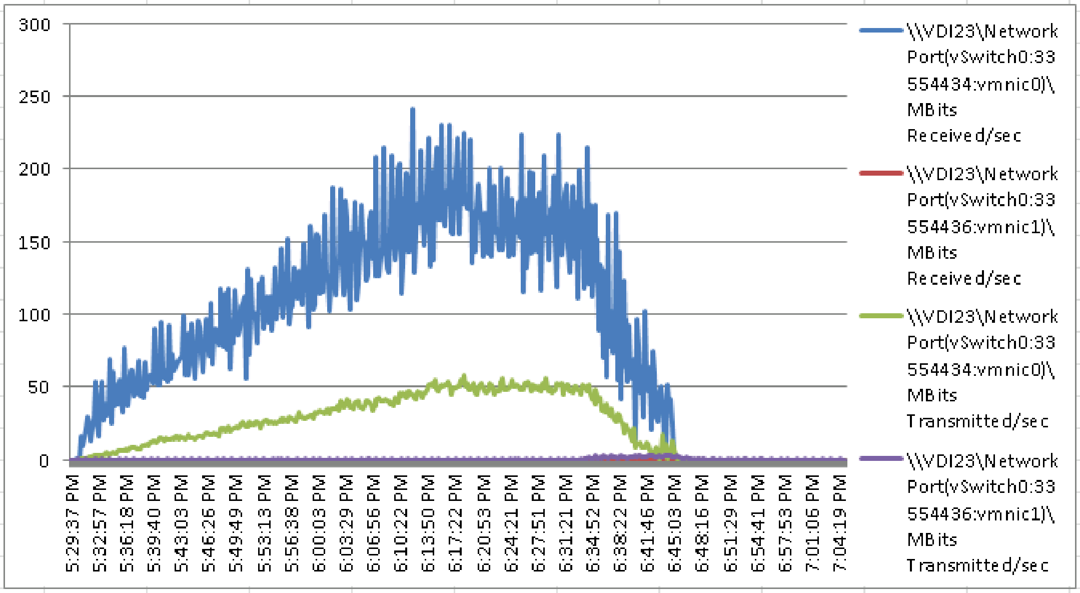

Figure 53 VDI- Single Server Recommended Maximum Workload for VDI Non-Persistent with 210 Users

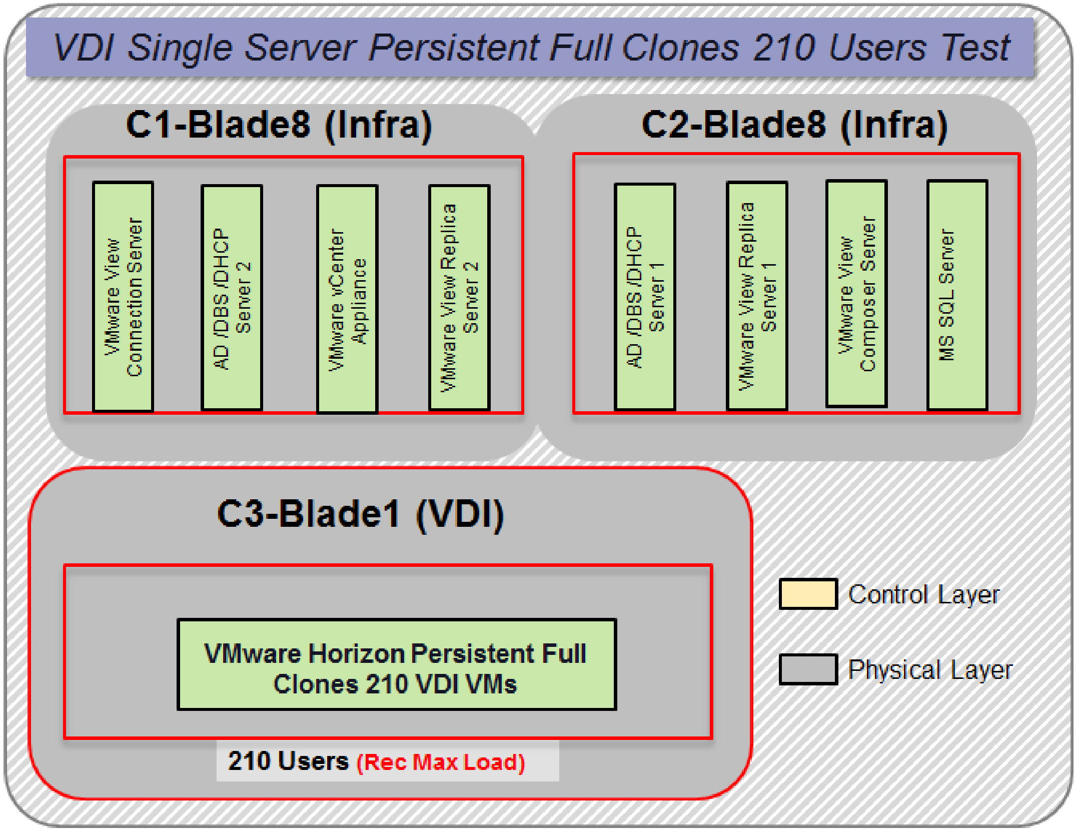
VDI Single Server Persistent Full Clones 210 Users Test
C1-Blade8 (Infra)
VMware View Connection Server
AD /DBS /DHCP Server 2
VMware vCenter Appliance
VMware View Replica Server 2
C2-Blade8 (Infra)
AD /DBS /DHCP Server 1
VMware View Replica Server 1
VMware View Composer Server
MS SQL Server
C3-Blade1 (VDI)
VMware Horizon Persistent Full Clones 210 VDI VMs
210 Users (Rec Max Load)
Control Layer
Physical Layer

Figure 54 Single Server | VMware Horizon 7 VDI-Persistent | VSI Score

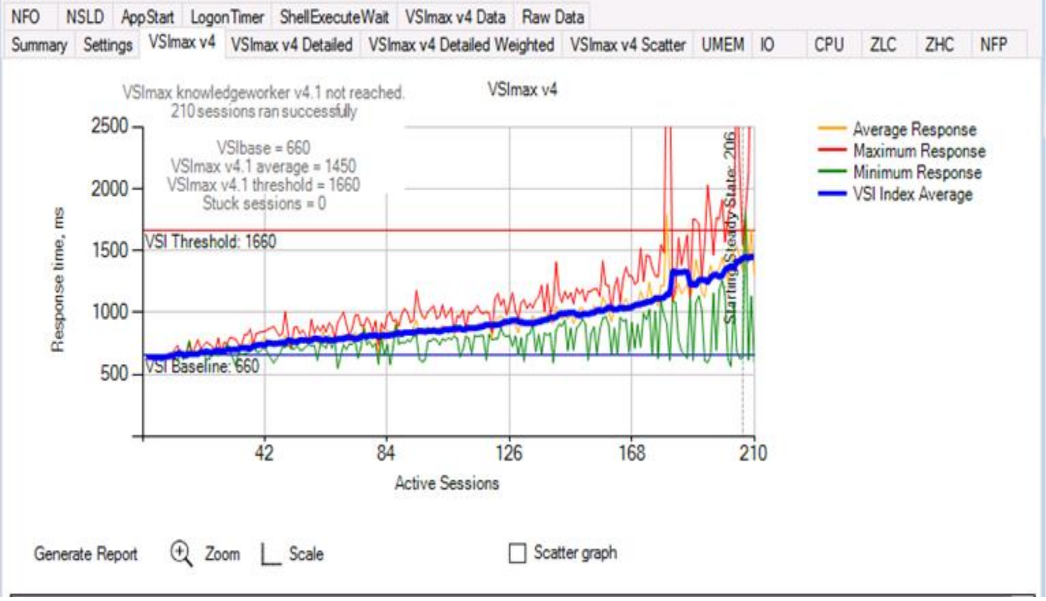

VSImax knowledgeworker v4.1 not reached.
210 sessions ran successfully

VSIbase = 660
VSImax v4.1 average = 1450
VSImax v4.1 threshold = 1660
Stuck sessions = 0

Figure 55 Single Server | VMware Horizon VDI Non-Persistent Desktops | Host CPU Utilization

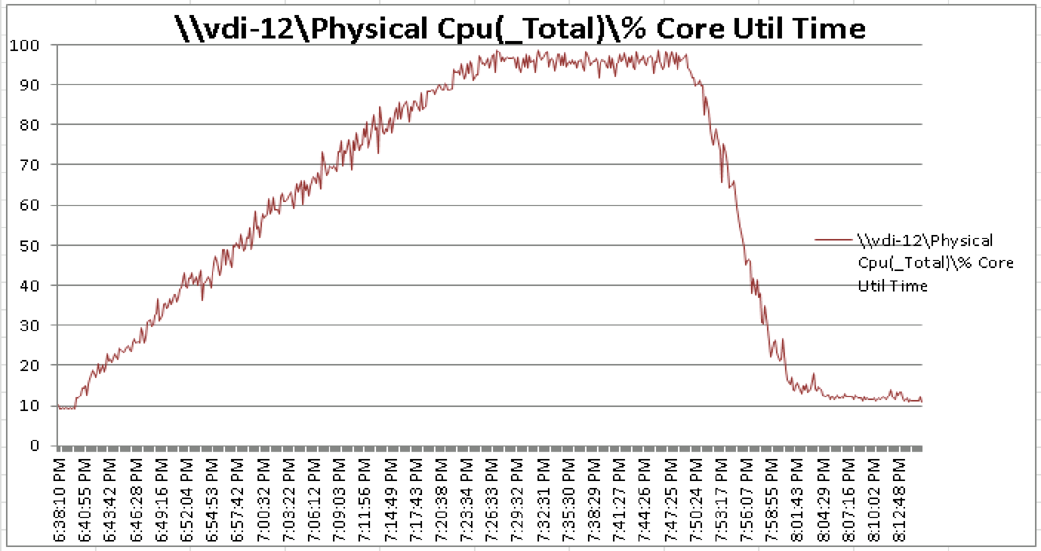

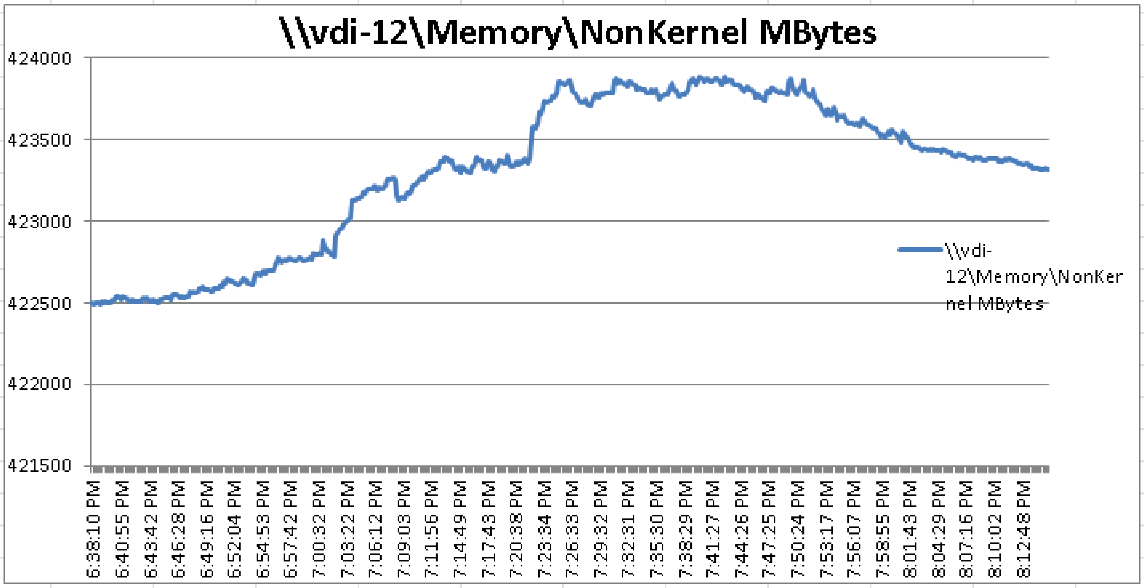

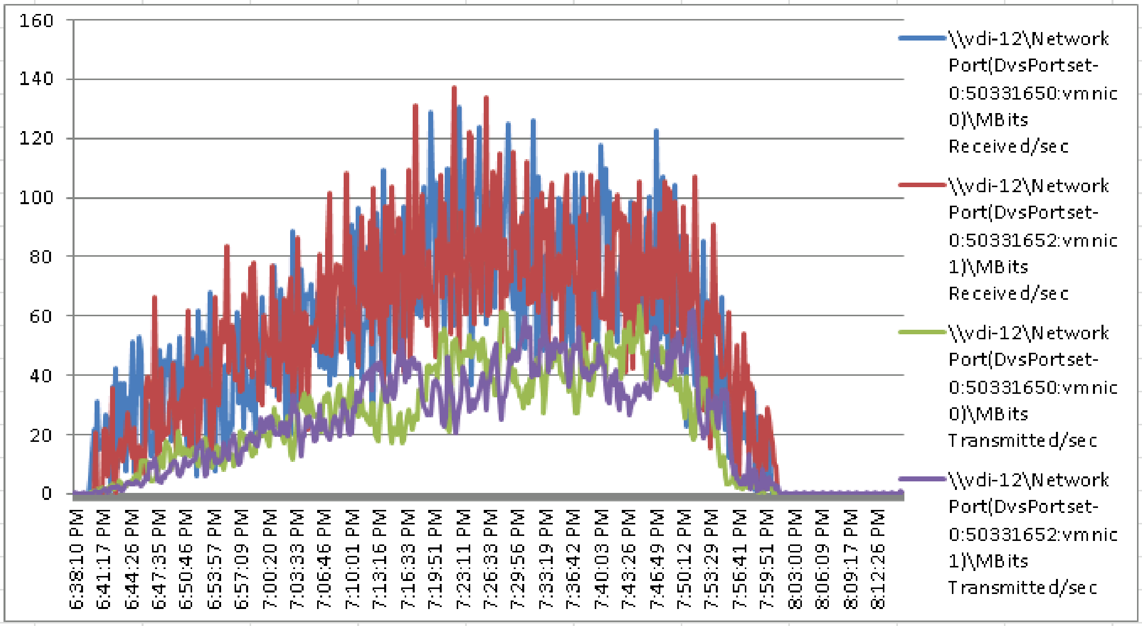

Cluster Workload Testing with 1680 Persistent RDS Users

This section provides the key performance metrics that were captured on the Cisco UCS, NetApp AFFA300 and RDS workload VMs during the RDSH Sessions testing. The cluster testing with comprised of 1680 RDS Hosted sessions using 7 Cisco UCS B200 M5 workload blades.

Figure 56 RDS Cluster Testing with 1680 Users

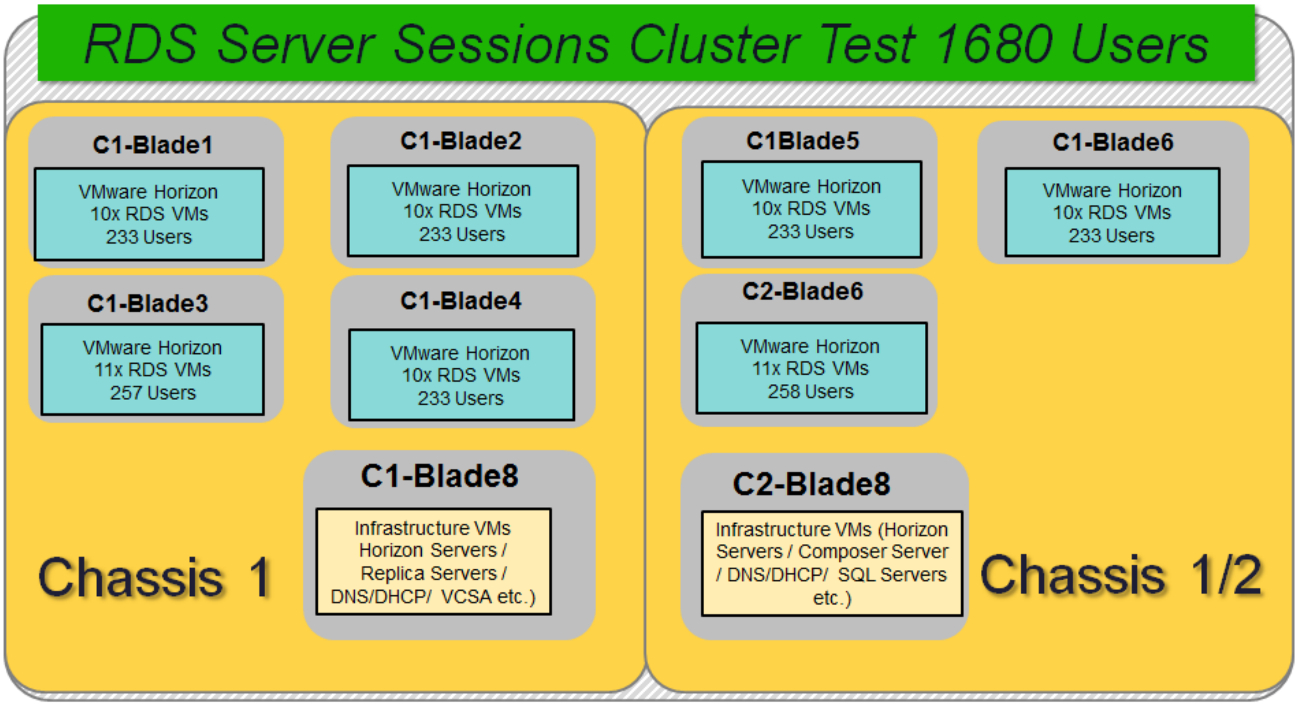
RDS Server Sessions Cluster Test 1680 Users
C1-Blade1
VMware Horizon 10x RDS VMs 233 Users
C1-Blade2
VMware Horizon 10x RDS VMs 233 Users
C1Blade5
VMware Horizon 10x RDS VMs 233 Users
C1-Blade6
VMware Horizon 10x RDS VMs 233 Users
C1-Blade3
VMware Horizon 11x RDS VMs 257 Users
C1-Blade4
VMware Horizon 10x RDS VMs 233 Users
C2-Blade6
VMware Horizon 11x RDS VMs 258 Users
C1-Blade8
Infrastructure VMs Horizon Servers / Replica Servers / DNS/DHCP/ VCSA etc.)
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Server / DNS/DHCP/ SQL Servers etc.)
Chassis 1
Chassis 1/2

The workload for the test is 1680 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 57 RDSH Cluster | 1680 RDS Users | VMware Horizon RDSH VSI Score

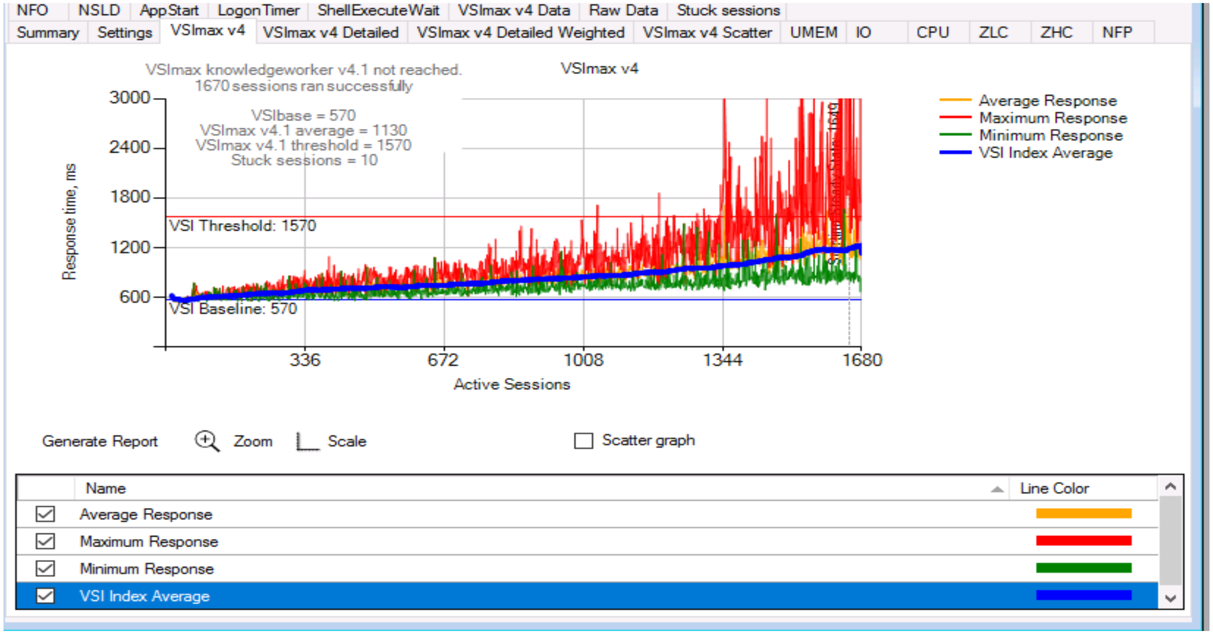

VSImax knowledgeworker v4.1 not reached.
1670 sessions ran successfully

VSIbase = 570
VSImax v4.1 average = 1130
VSImax v4.1 threshold = 1570
Stuck sessions = 10

☐ Scatter graph

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 58 RDSH Cluster | 1680 RDS Users | Workload Host | Host CPU Utilization

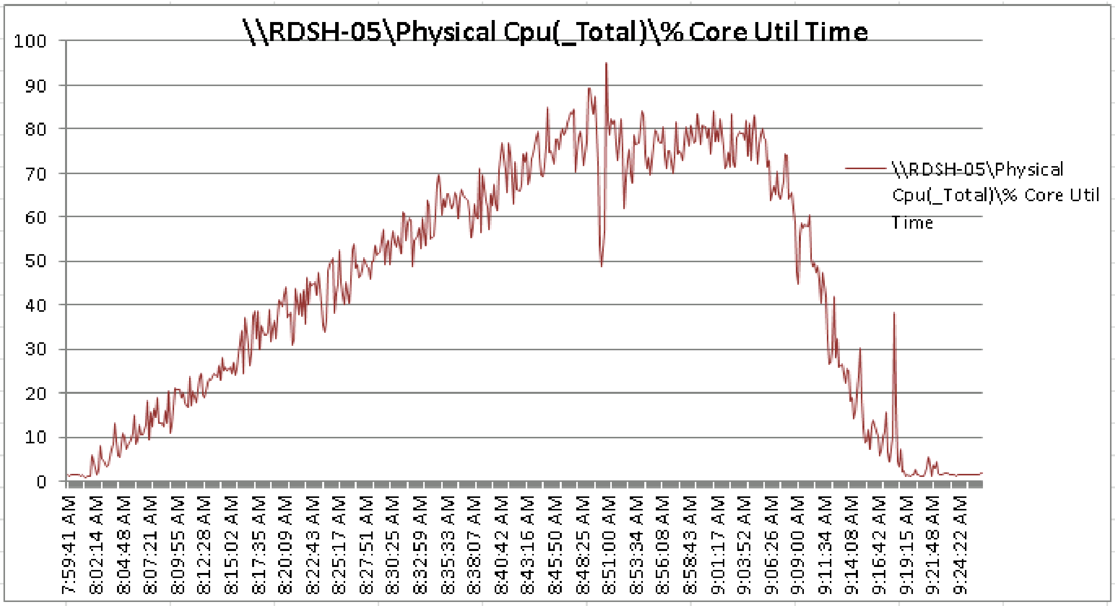

Figure 59 RDSH Cluster | 1680 RDS Users | Workload Host | Host Memory Utilization

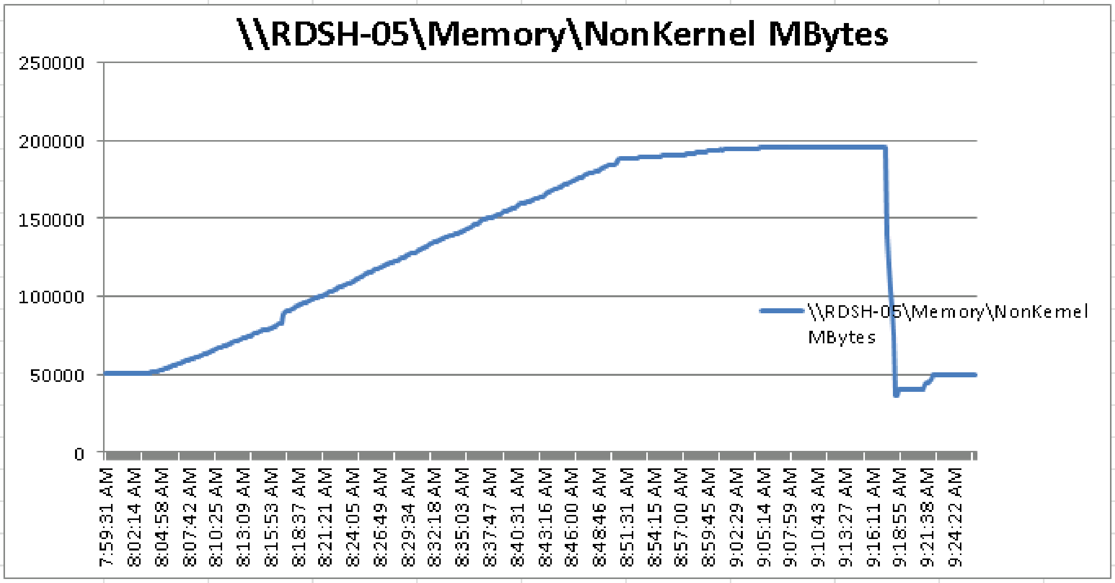

Figure 60 RDSH Cluster | 1680 RDS Users | RDS Host | Host Network Utilization

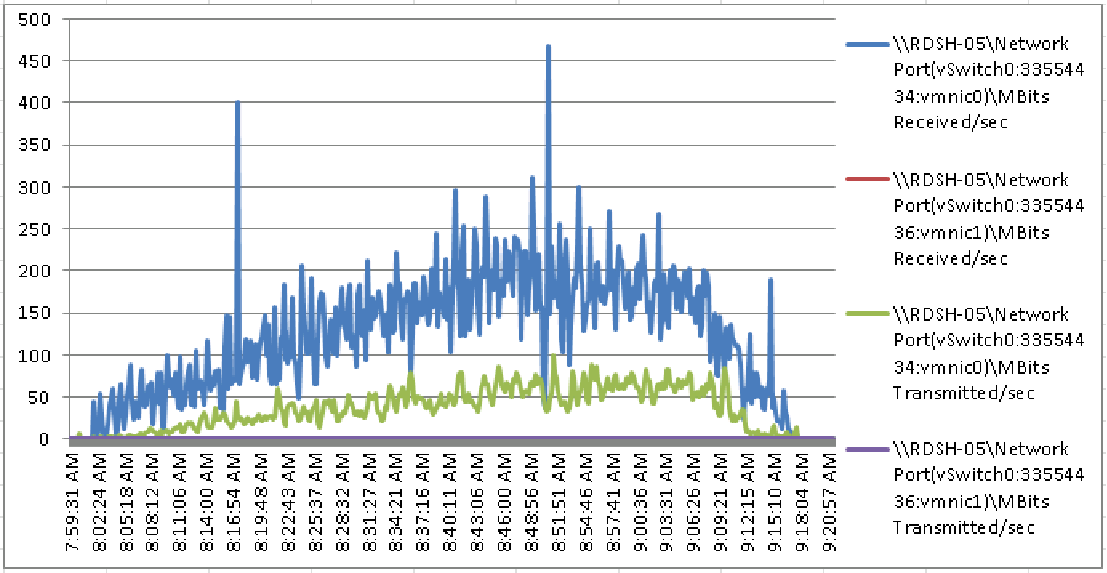

Figure 61 Cluster | 1680 RDS Users | RDS Host | Host Fibre Channel Network Utilization

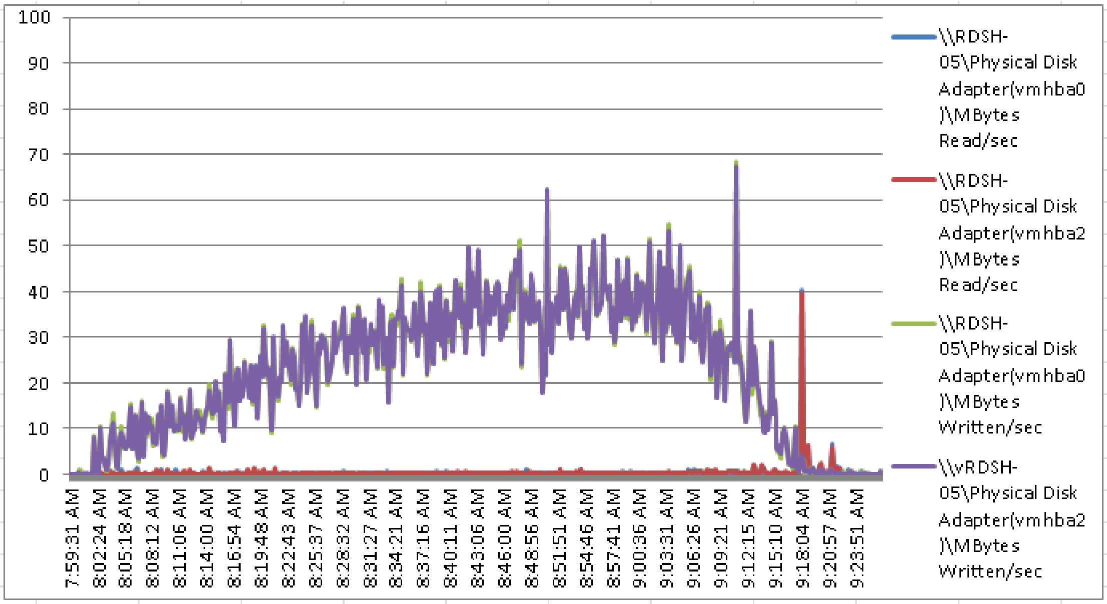

Performance Data from One RDSH Server: 1680 Users RDSH Sessions Cluster Testing

Figure 62 RDSH Server Processor (Total%) Time

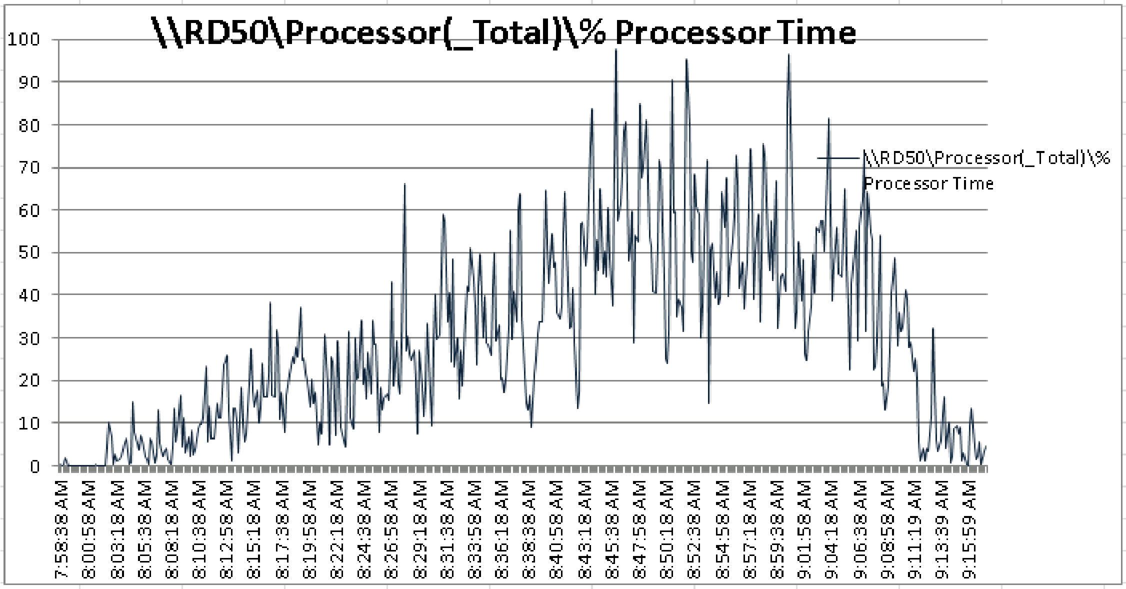

Figure 63 RDSH Server User (Total%) Time

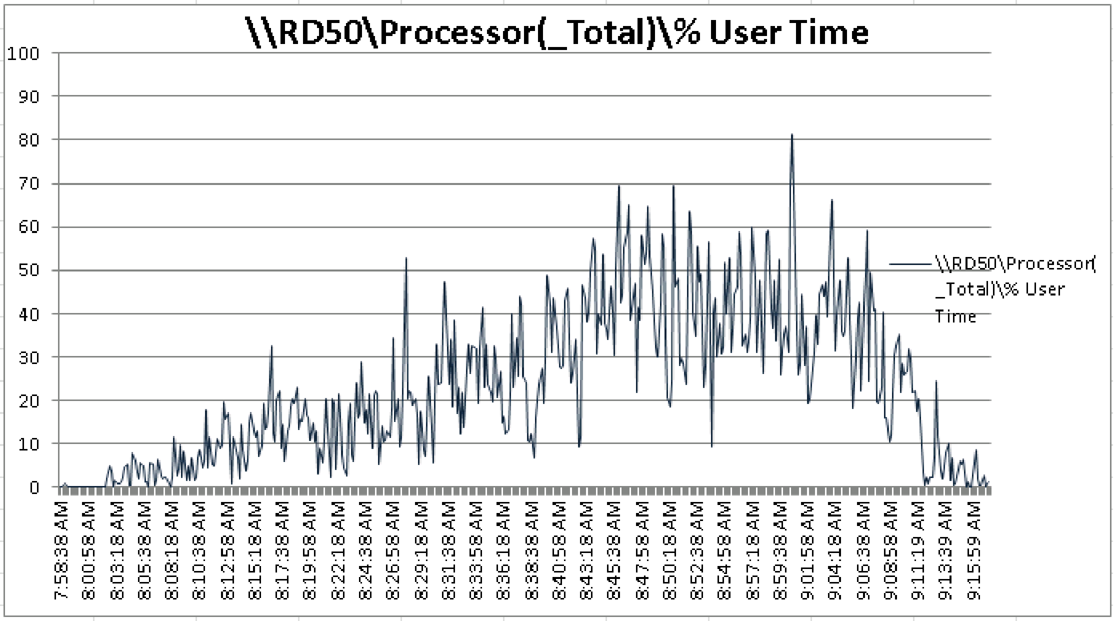

Cluster Workload Testing with 1660 Persistent and Non-Persistent VDI Instant Desktop Users

This section provides the key performance metrics that were captured on the Cisco UCS, NetApp AFF A300storage, and Infrastructure VMs during the persistent desktop testing. The cluster testing with comprised of 1660 VDI Non- Persistent desktop sessions using 9 workload blades.

Figure 64 VMware Horizon VDI Non-Persistent Cluster Testing with 1660 Users

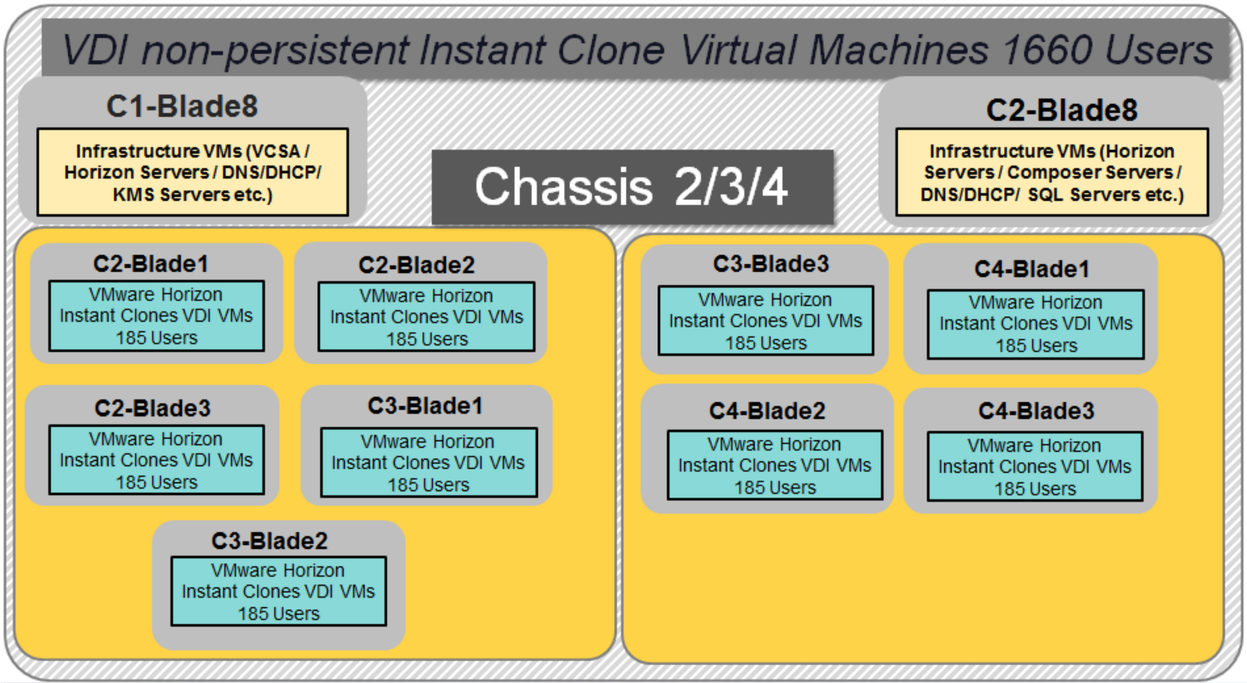
VDI non-persistent Instant Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users

The workload for the test is 1660 non-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 65 VDI Cluster | VMware Horizon 1660 VDI Non-Persistent Users | VSI Score

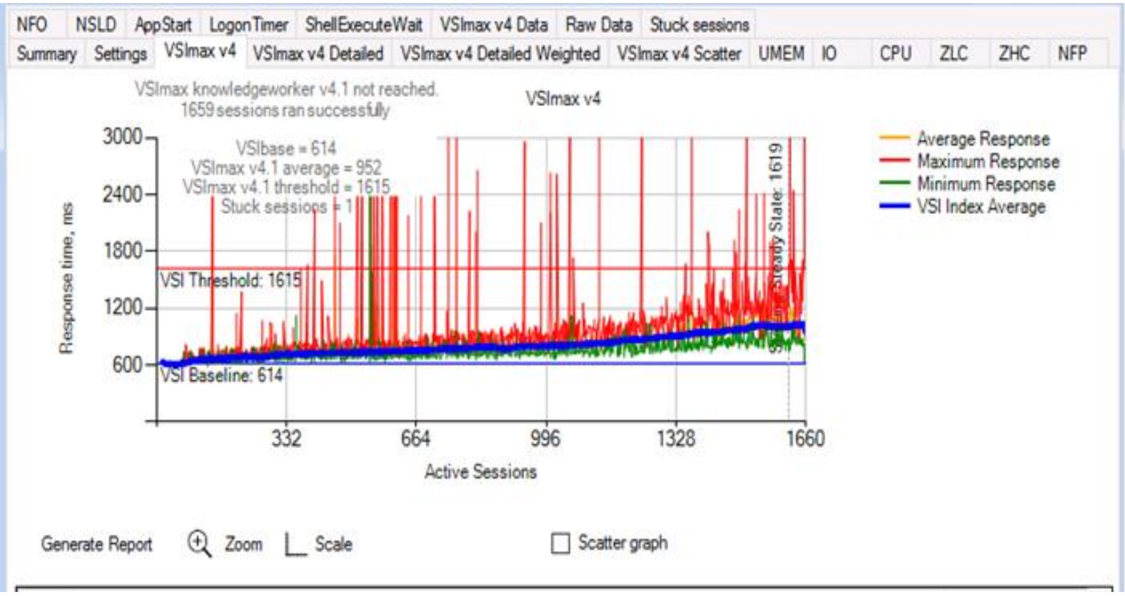

VSImax knowledgeworker v4.1 not reached.
1659 sessions ran successfully

VSIbase = 614
VSImax v4.1 average = 952
VSImax v4.1 threshold = 1615
Stuck sessions = 1

Figure 66 VDI Cluster | VDI Cluster | 1660 VDI Non-Persistent Users | VDI Host | Host CPU Utilization

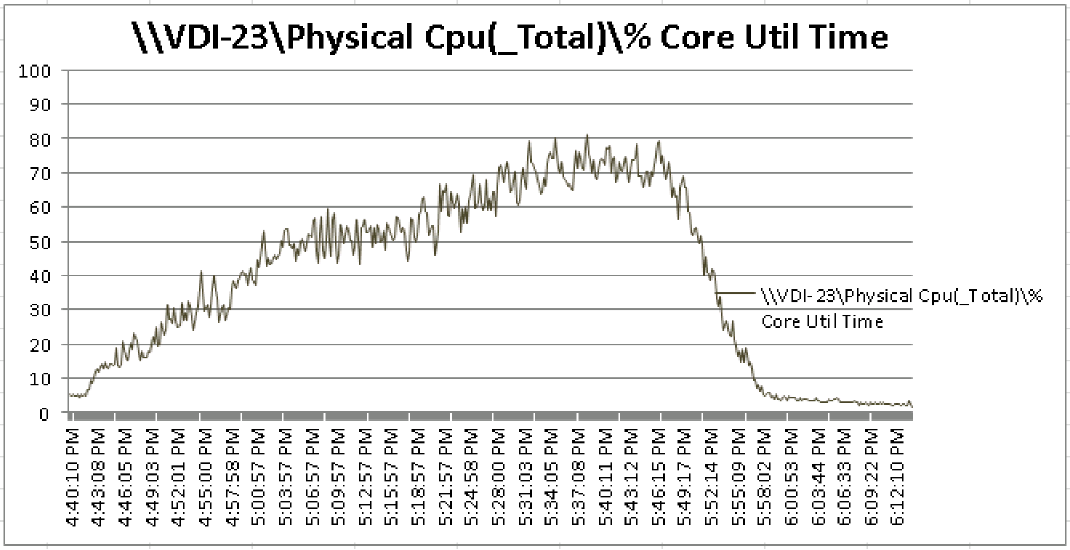

Figure 67 VDI Cluster | VDI Cluster | 1660 VDI Non-Persistent Users | VDI Host | Host Memory Utilization

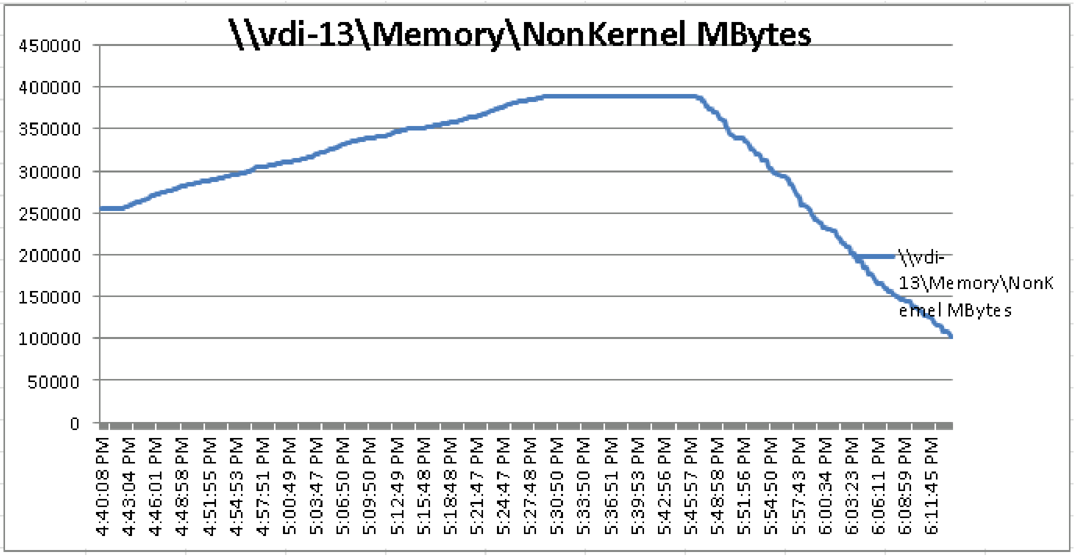

Figure 68 VDI Cluster | 1660 VDI Non-Persistent Users | VDI Host | Host Network Utilization

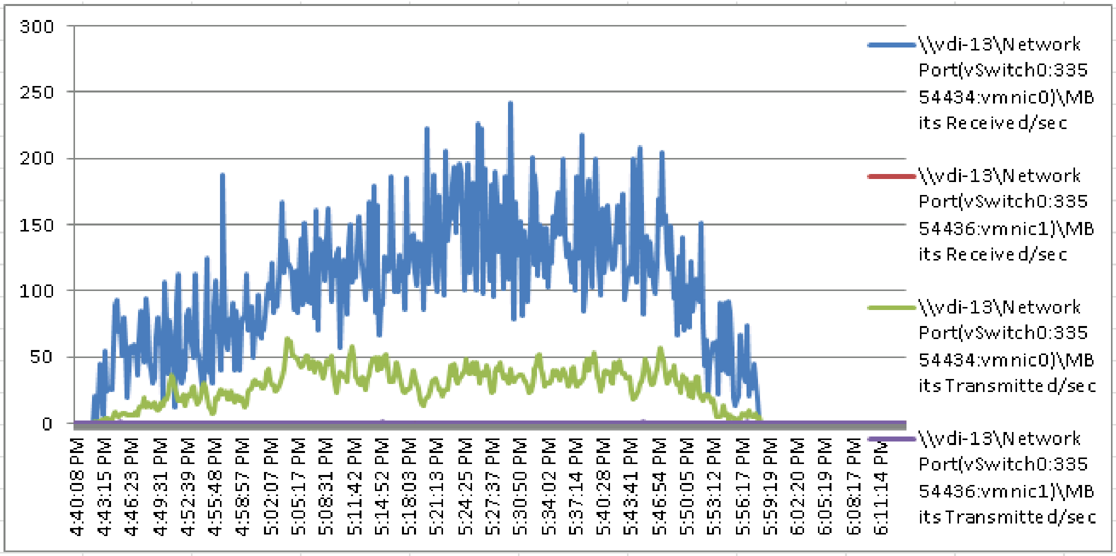

Figure 69 VDI Cluster | 1660 VDI-Non Persistent Users | VDI Host | Host Fibre Channel Network Utilization

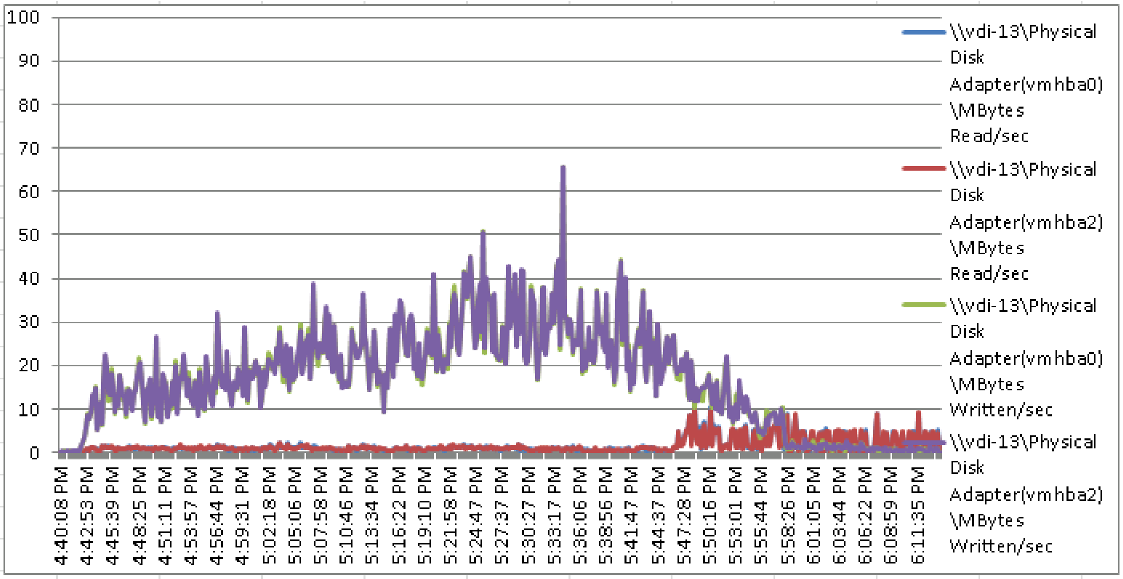

The workload for the test is 1660 persistent desktop users. To achieve the target, sessions were launched against the single persistent cluster only. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark 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 70 VMware Horizon VDI Persistent Cluster Testing with 1660 Users

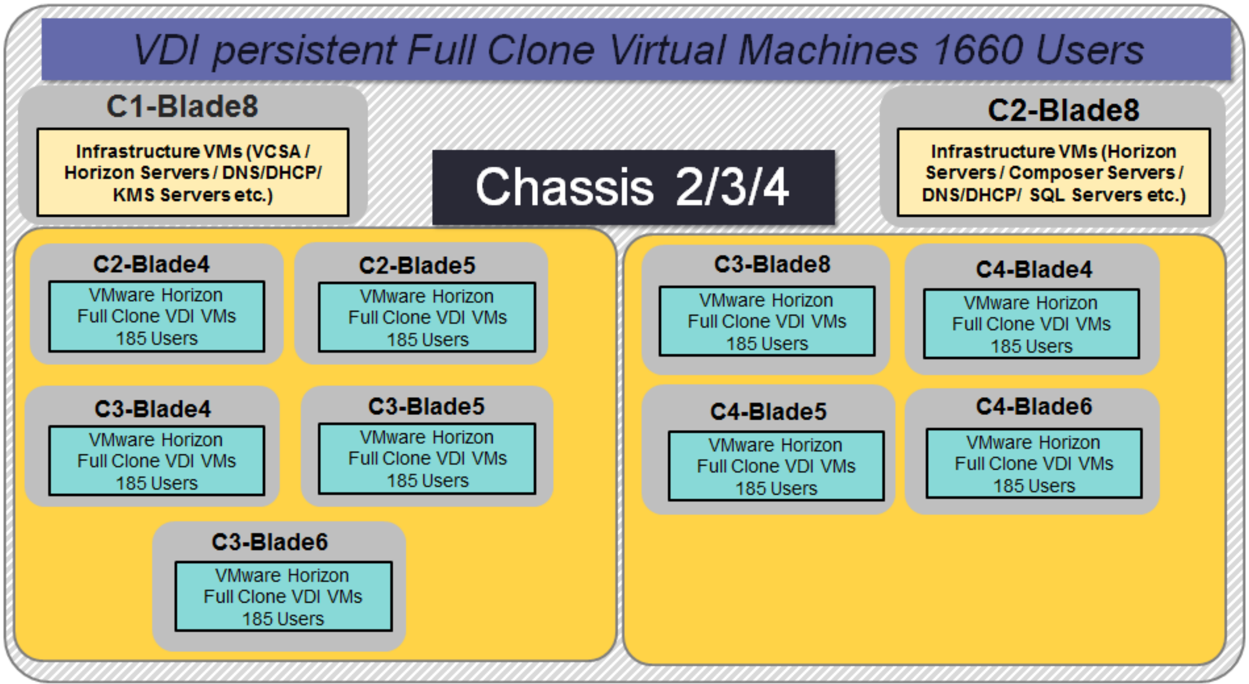
VDI persistent Full Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C2-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade8
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade6
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade6
VMware Horizon Full Clone VDI VMs 185 Users

Figure 71 VDI Cluster | VMware Horizon 1660 VDI Persistent Users | VSI Score

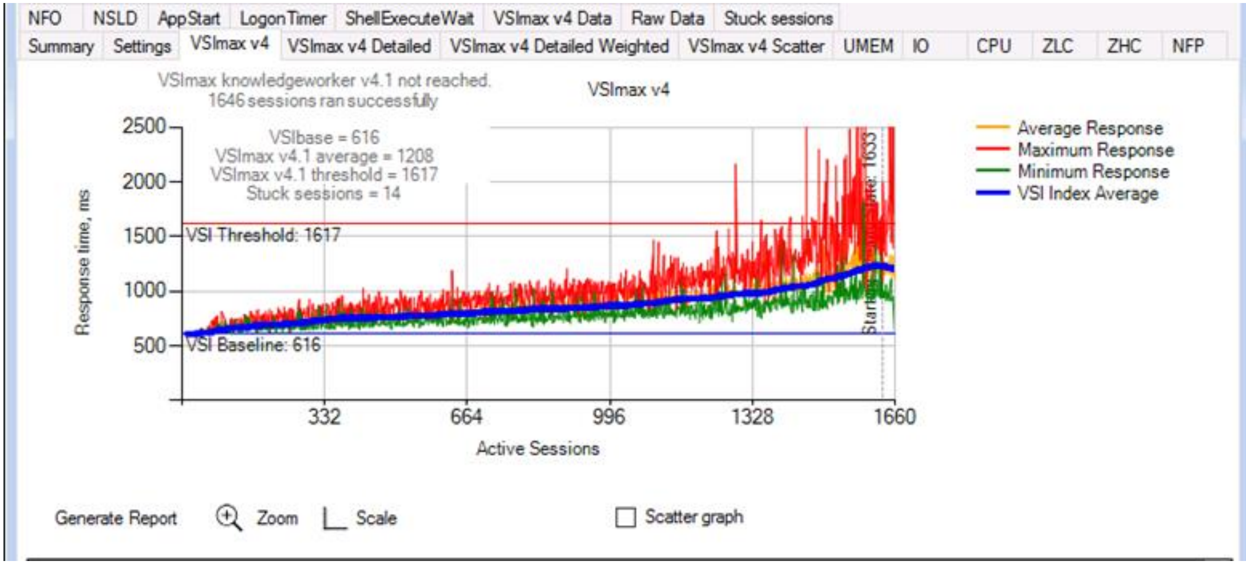
VSImax knowledgeworker v4.1 not reached.
1646 sessions ran successfully

VSIbase = 616
VSImax v4.1 average = 1208
VSImax v4.1 threshold = 1617
Stuck sessions = 14

Figure 72 VDI Cluster | VDI Cluster | 1660 VDI Persistent Users | VDI Host | Host CPU Utilization

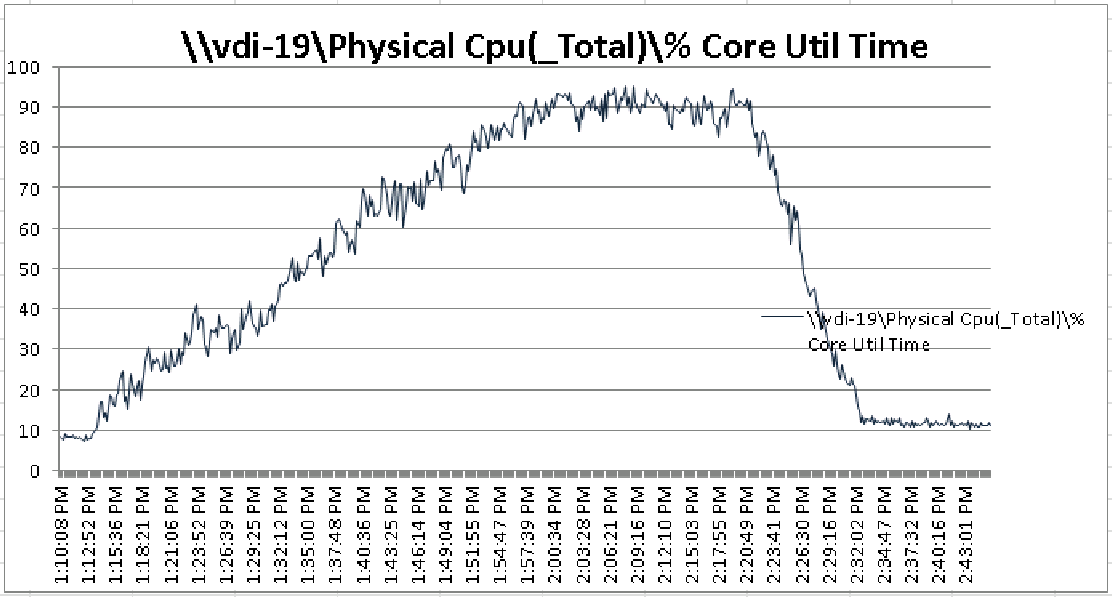

Figure 73 VDI Cluster | VDI Cluster | 1660 VDI Persistent Users | VDI Host | Host Memory Utilization

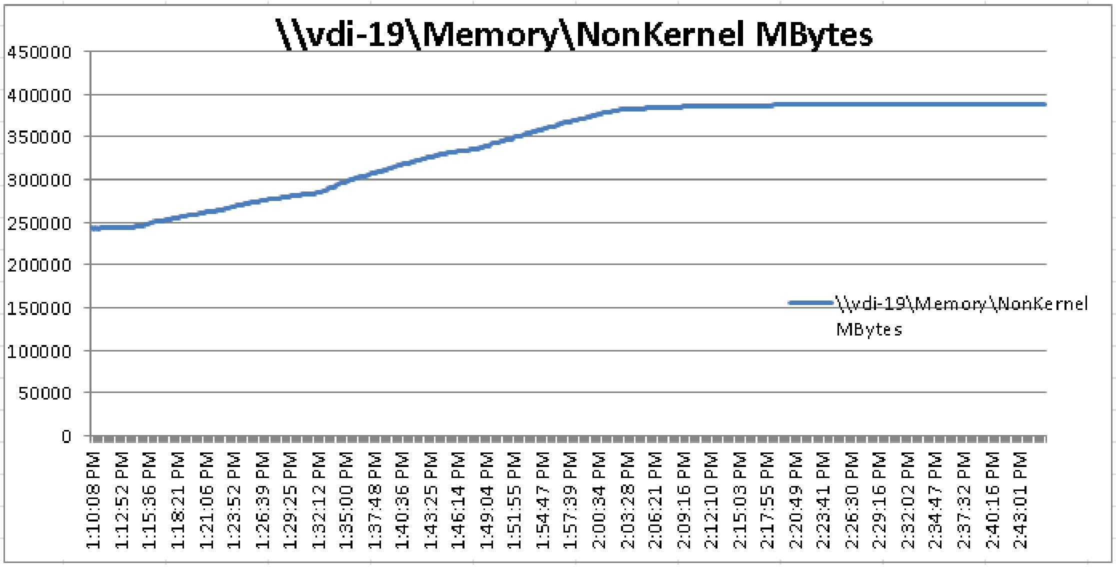

Figure 74 VDI Cluster | 1660 VDI Persistent Users | VDI Host | Host Network Utilization

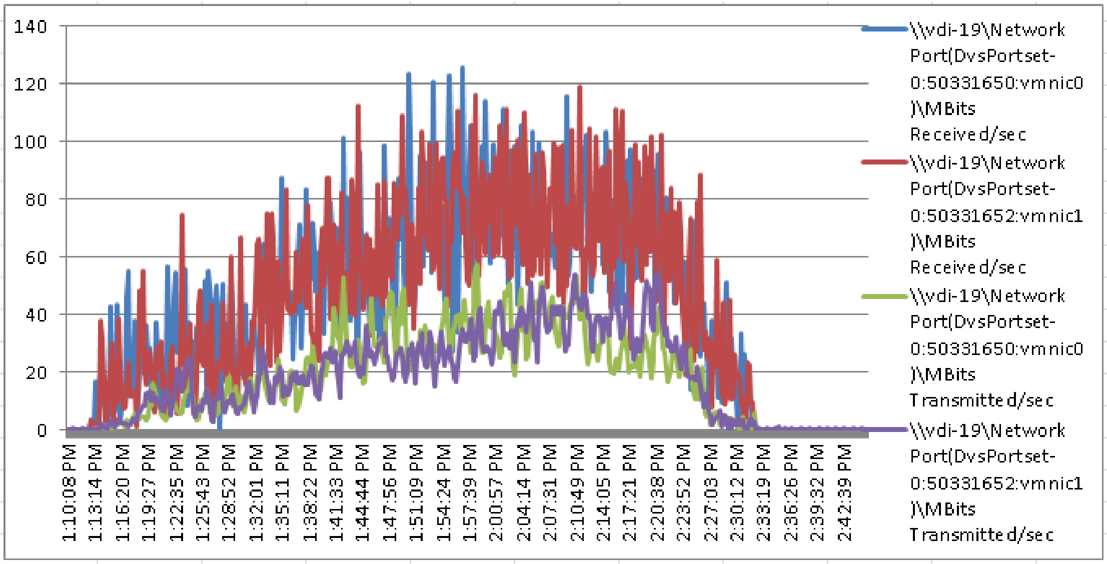

Figure 75 VDI Cluster | 1660 VDI Persistent Users | VDI Host | Host Fibre Channel Network Utilization

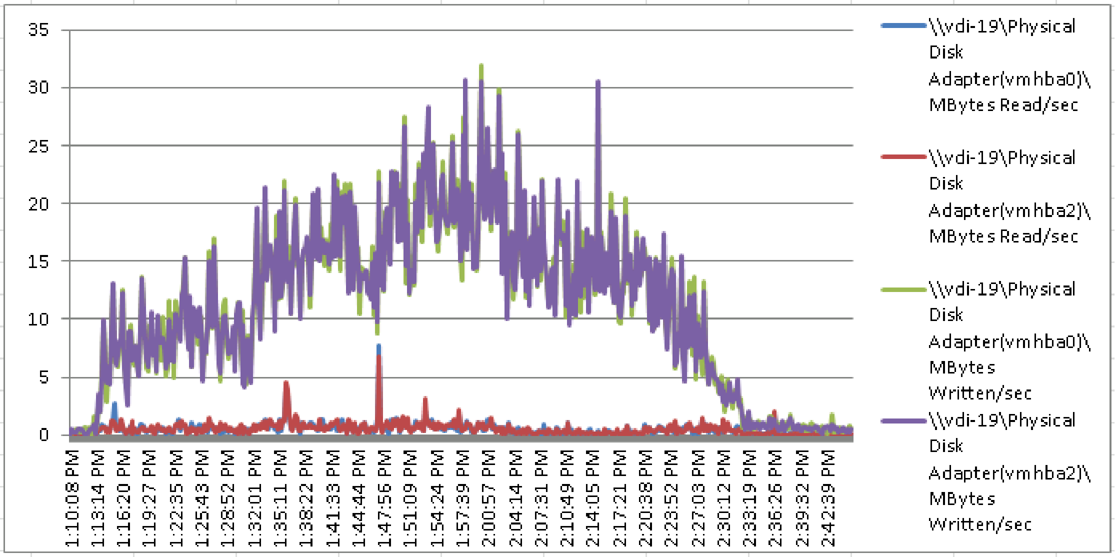

Full Scale Mixed Workload Testing with 5000 Users

This section provides the key performance metrics that were captured on the Cisco UCS, NetApp AFF A300storage, RDSH VMs and VDI non –persistent and Persistent VDI virtual machines performance monitoring during the full-scale testing. The full-scale testing with 5000 users comprised of: 1680 RDS Hosted Server Sessions using 7 Cisco UCS B200 M5 blades, 1660 VDI Non-Persistent Instant clones using 9 Cisco UCS B200 M5 blades and remaining 1660 VDI Persistent Full clones using 9 Cisco UCS B200 M5 blades.

Figure 76 Full Scale Mixed Test with 5000 Users

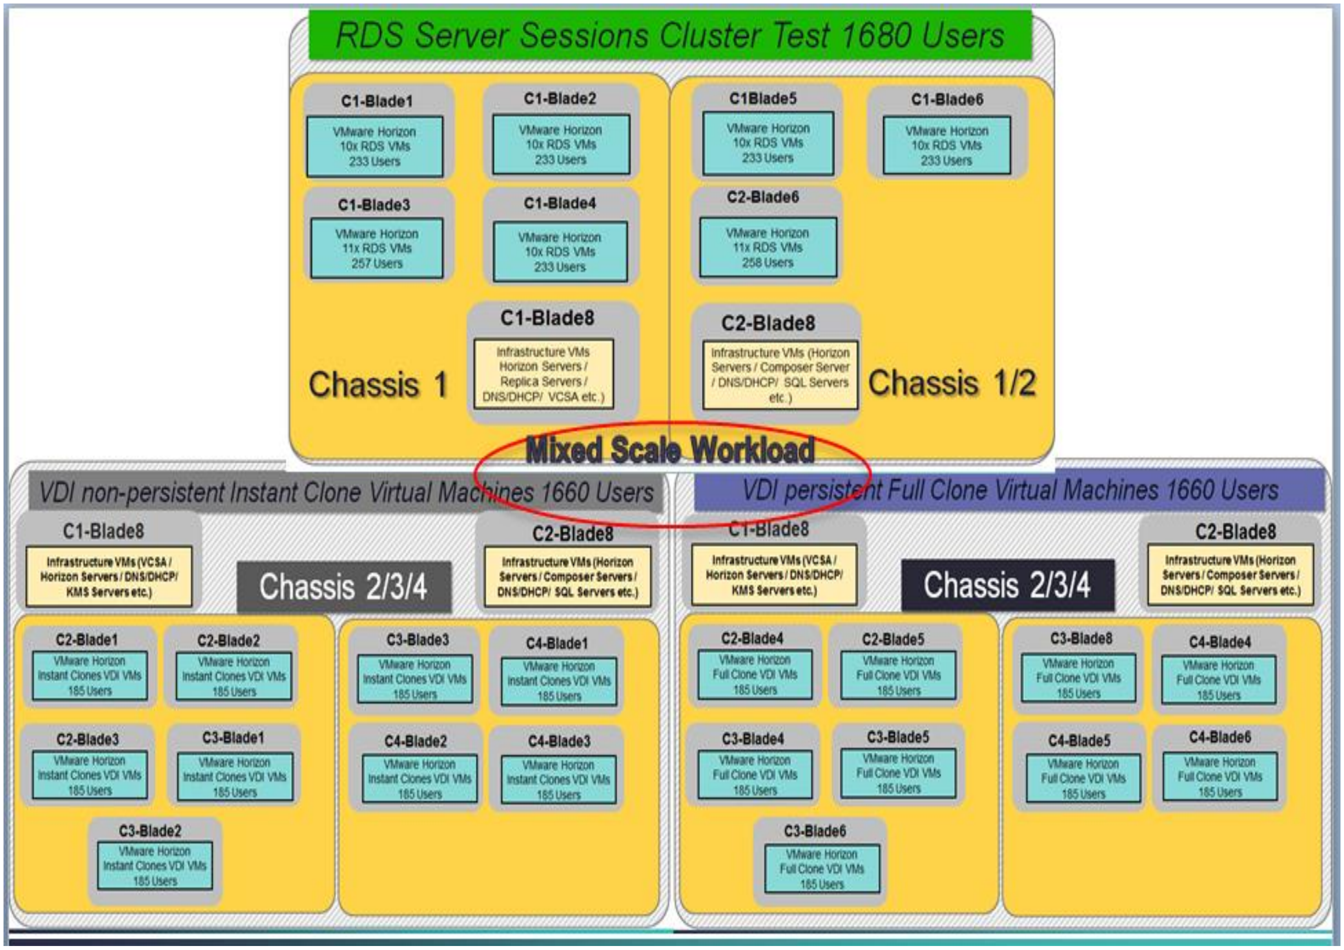
RDS Server Sessions Cluster Test 1680 Users
C1-Blade1
VMware Horizon 10x RDS VMs 233 Users
C1-Blade2
VMware Horizon 10x RDS VMs 233 Users
C1Blade5
VMware Horizon 10x RDS VMs 233 Users
C1-Blade6
VMware Horizon 10x RDS VMs 233 Users
C1-Blade3
VMware Horizon 11x RDS VMs 257 Users
C1-Blade4
VMware Horizon 10x RDS VMs 233 Users
C2-Blade6
VMware Horizon 11x RDS VMs 258 Users
C1-Blade8
Infrastructure VMs Horizon Servers / Replica Servers / DNS/DHCP/ VCSA etc.)
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Server / DNS/DHCP/ SQL Servers etc.)
Chassis 1
Chassis 1/2
Mixed Scale Workload
VDI non-persistent Instant Clone Virtual Machines 1660 Users
VDI persistent Full Clone Virtual Machines 1660 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C2-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade1
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C4-Blade3
VMware Horizon Instant Clones VDI VMs 185 Users
C3-Blade2
VMware Horizon Instant Clones VDI VMs 185 Users
C1-Blade8
Infrastructure VMs (VCSA / Horizon Servers / DNS/DHCP/ KMS Servers etc.)
Chassis 2/3/4
C2-Blade8
Infrastructure VMs (Horizon Servers / Composer Servers / DNS/DHCP/ SQL Servers etc.)
C2-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C2-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade8
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade4
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade5
VMware Horizon Full Clone VDI VMs 185 Users
C4-Blade6
VMware Horizon Full Clone VDI VMs 185 Users
C3-Blade6
VMware Horizon Full Clone VDI VMs 185 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 77 Full Scale | 5000 Mixed Users | VSI Score

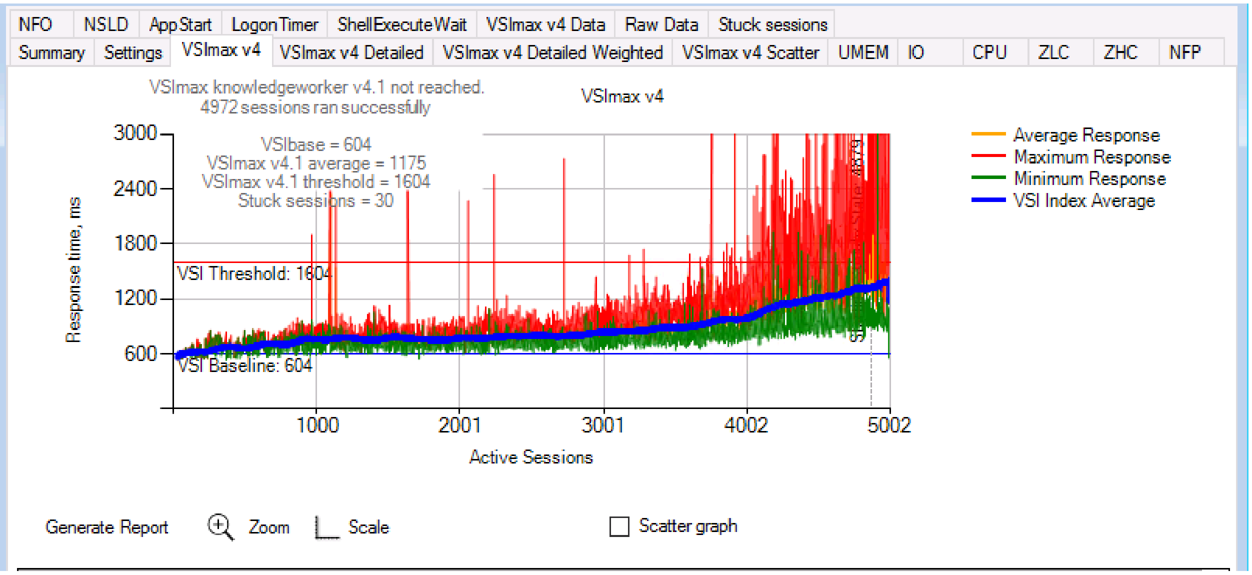

VSImax knowledgeworker v4.1 not reached.
4972 sessions ran successfully

VSIbase = 604
VSImax v4.1 average = 1175
VSImax v4.1 threshold = 1604
Stuck sessions = 30

Figure 78 Full Scale | 5000 Mixed Users RDSH Server Host | Host CPU Utilization

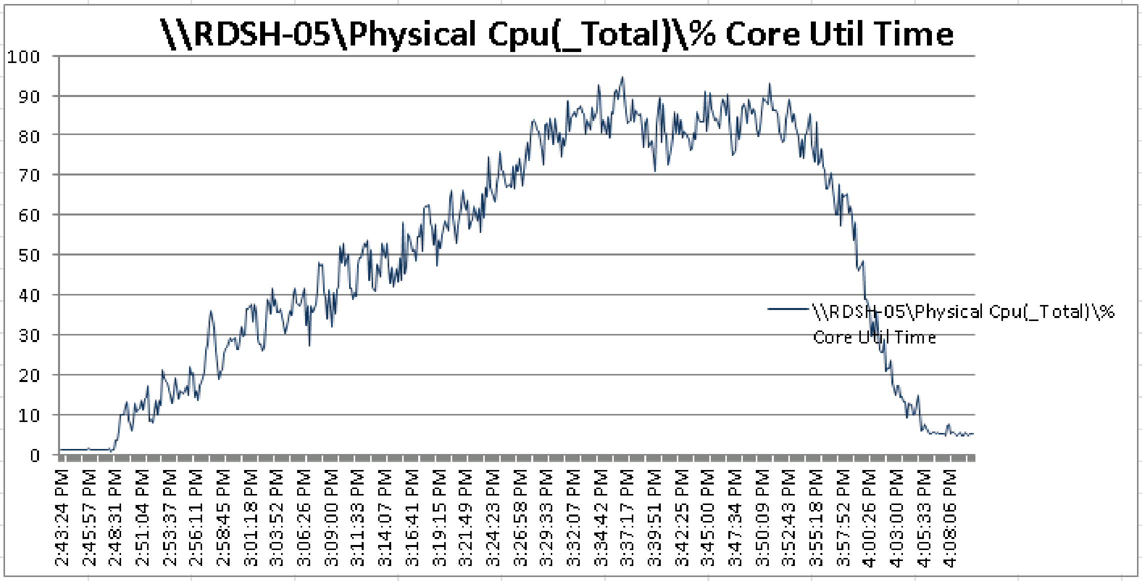

Figure 79 Full Scale | 5000 Mixed Users RDSH Server Host | Host Memory Utilization

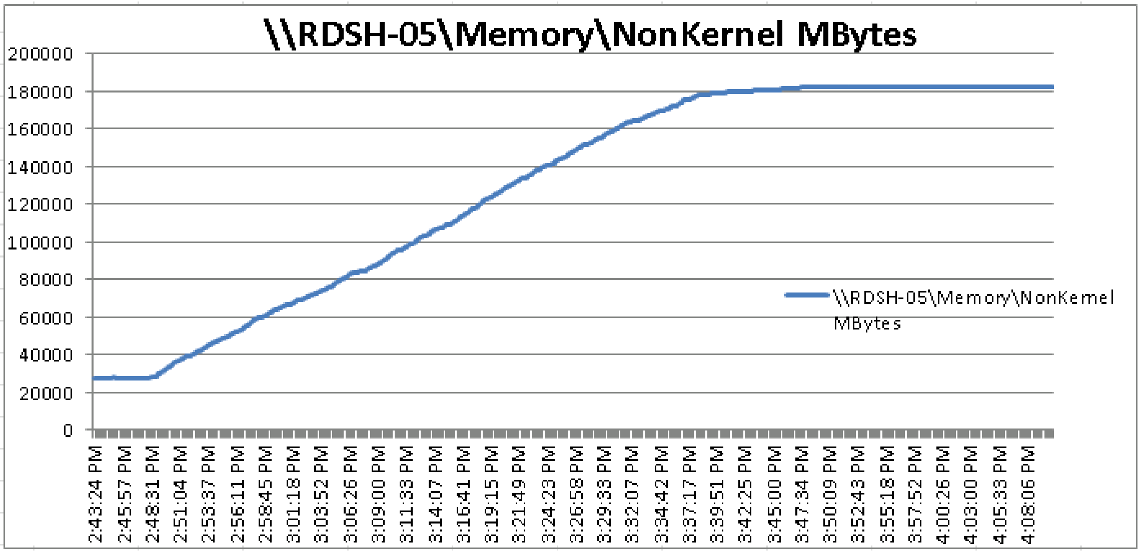

Figure 80 Full Scale | 5000 Mixed Users RDSH Server Host | Host CPU Network Utilization

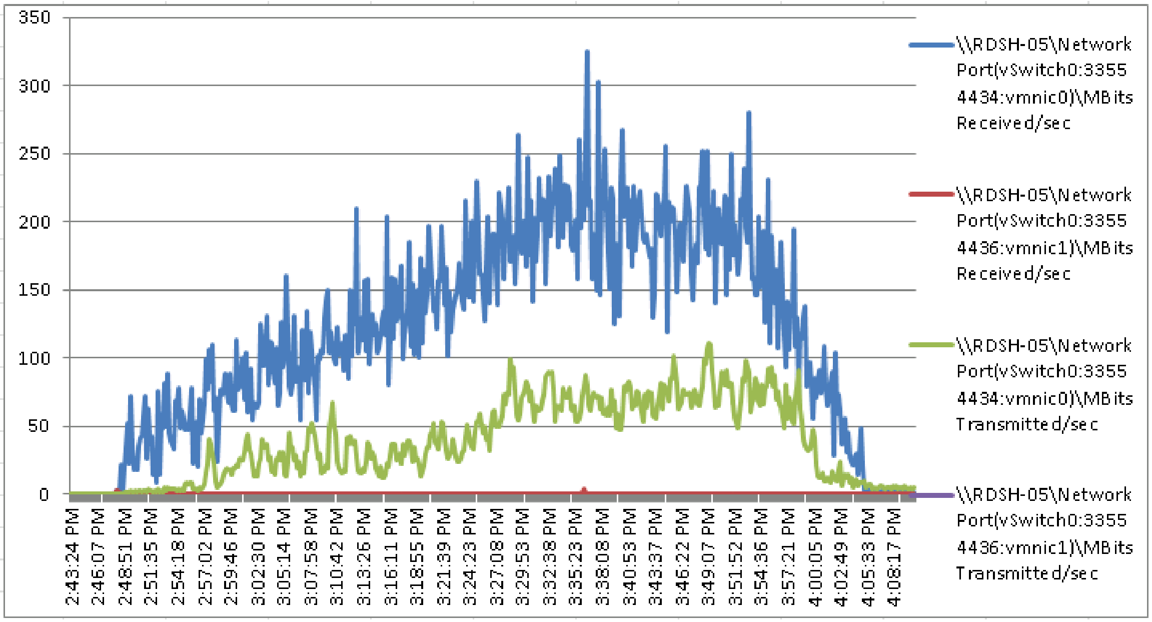

Figure 81 Full Scale | 5000 Mixed Users RDSH Server Host | Host vHBA Adapter Utilization

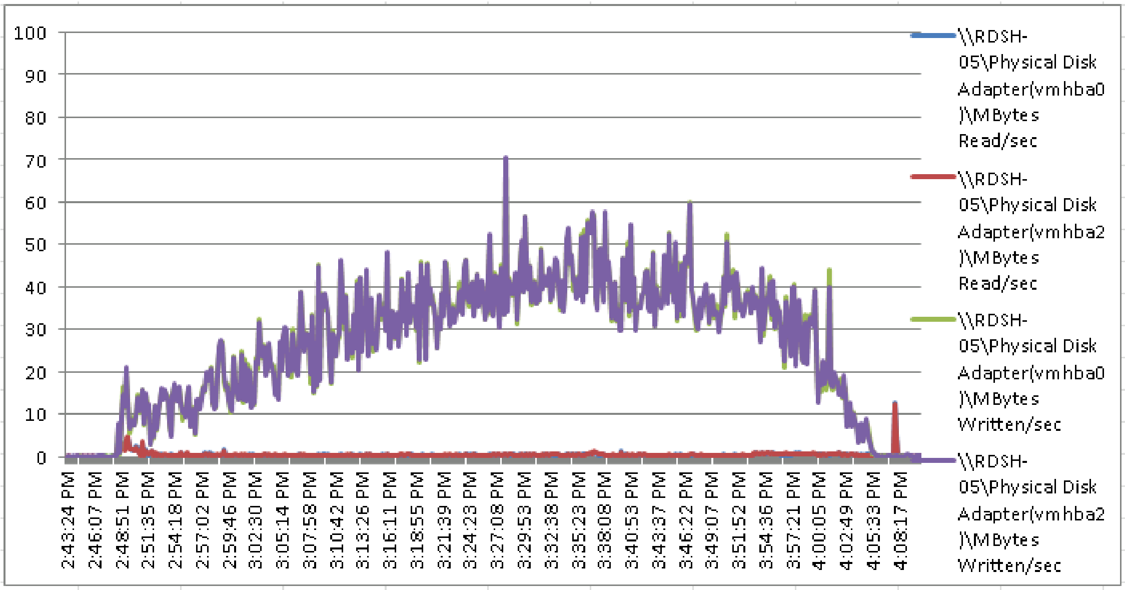

RDSH Server Performance Monitor Data for One Sample RDSH Server: 5000 Users Mixed Scale Testing

Figure 82 RDSH Server Processor (Total%) Time

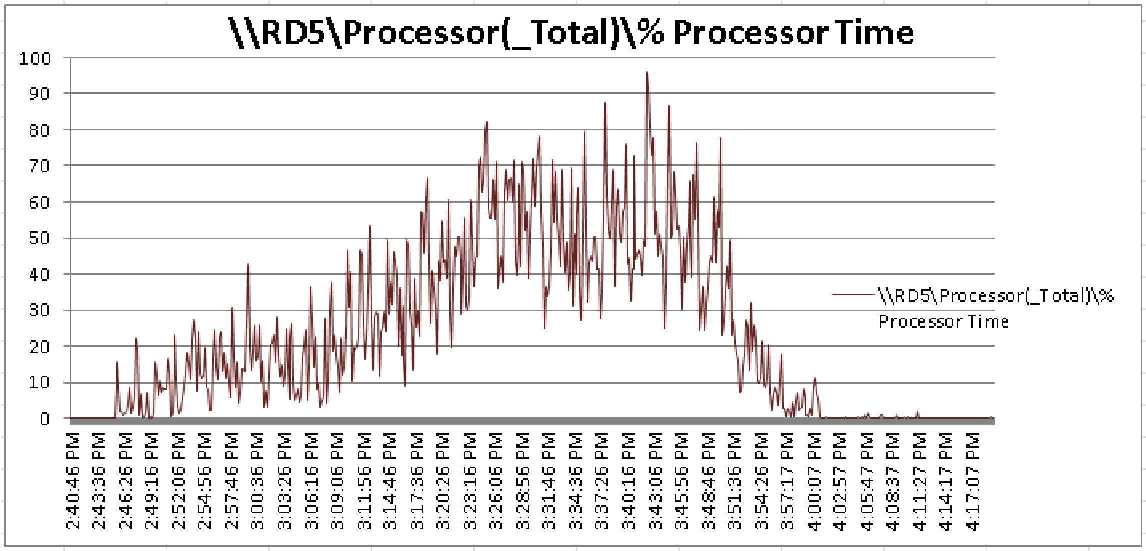

Figure 83 RDSH Server User (Total%) Time

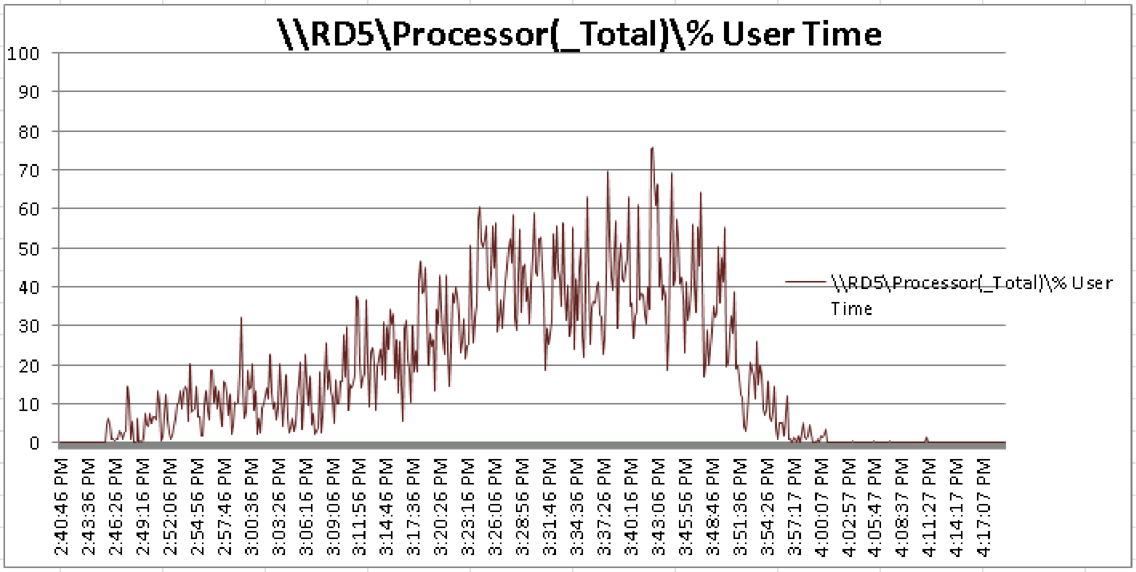

Figure 84 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

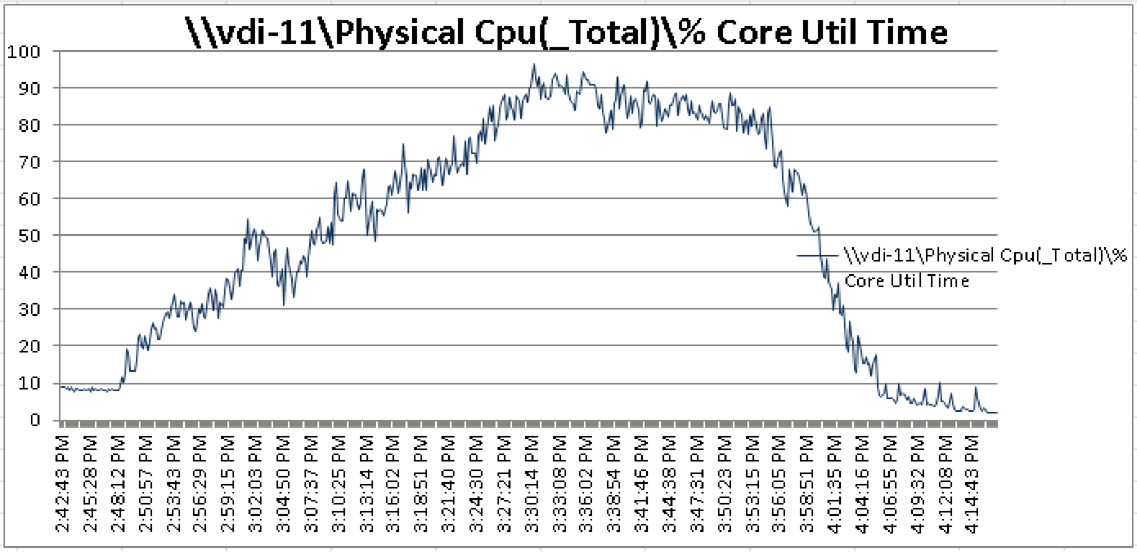

Figure 85 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

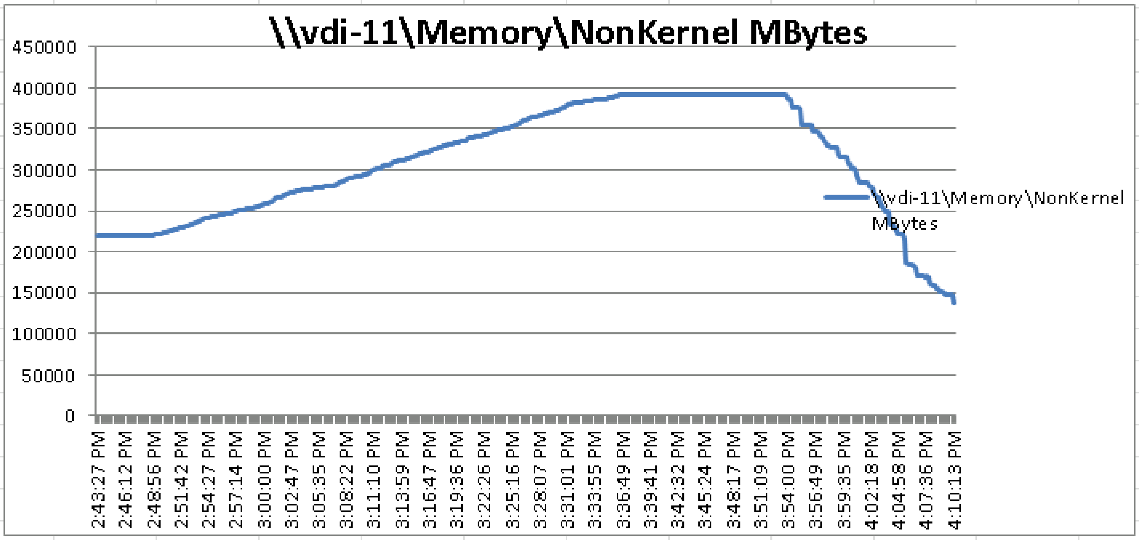

Figure 86 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

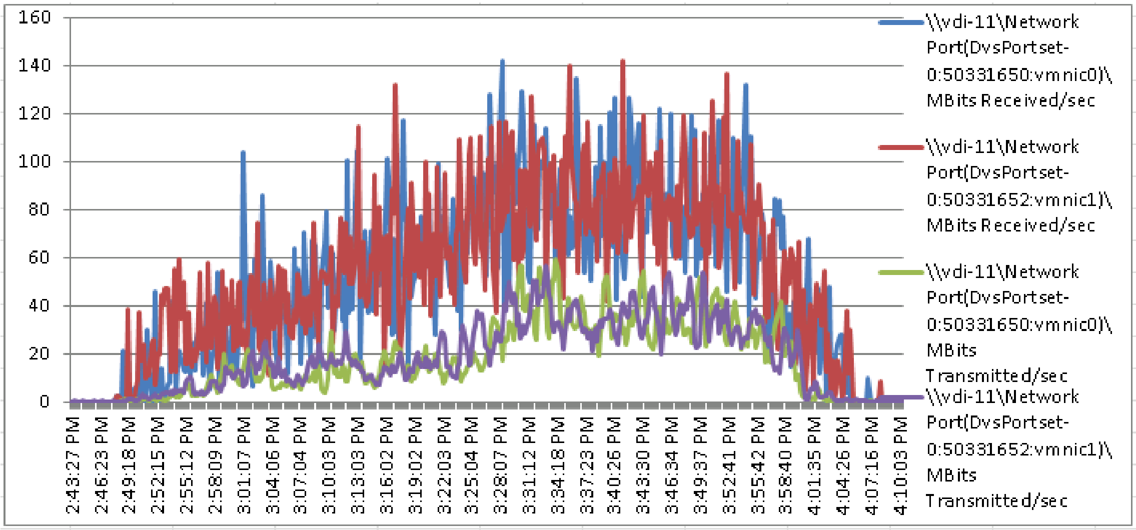

Figure 87 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

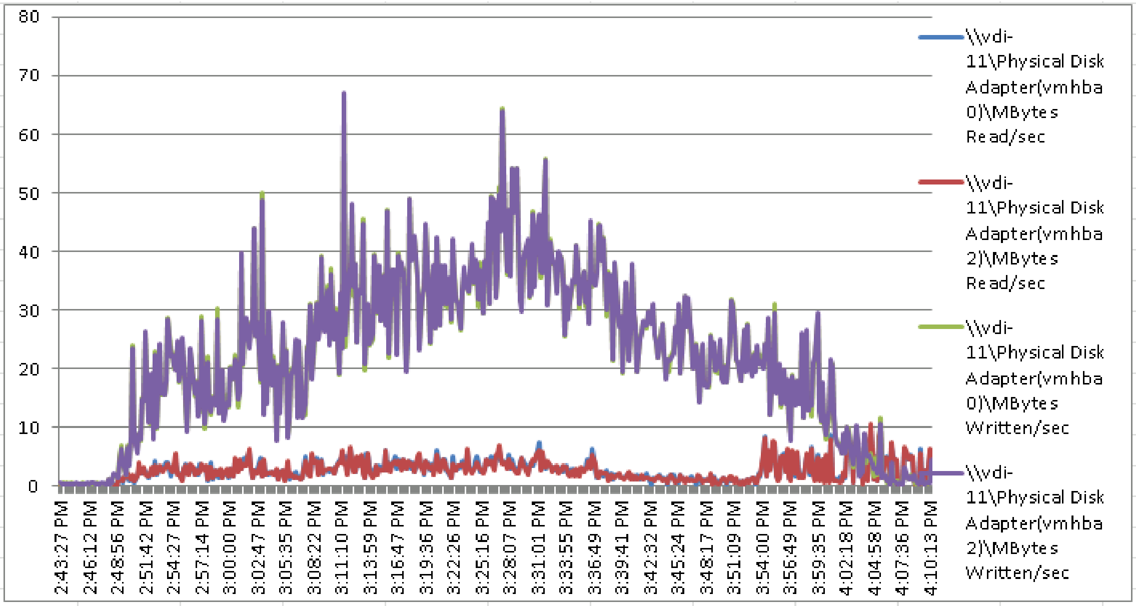

Full Clone VDI Host ESXTOP Charts

Figure 88 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

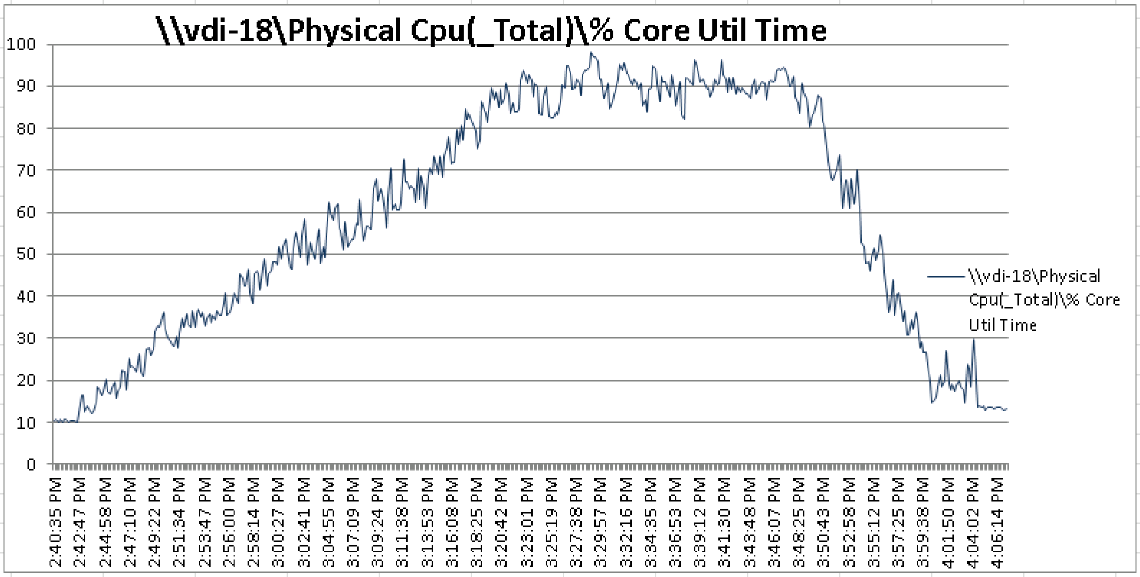

Figure 89 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

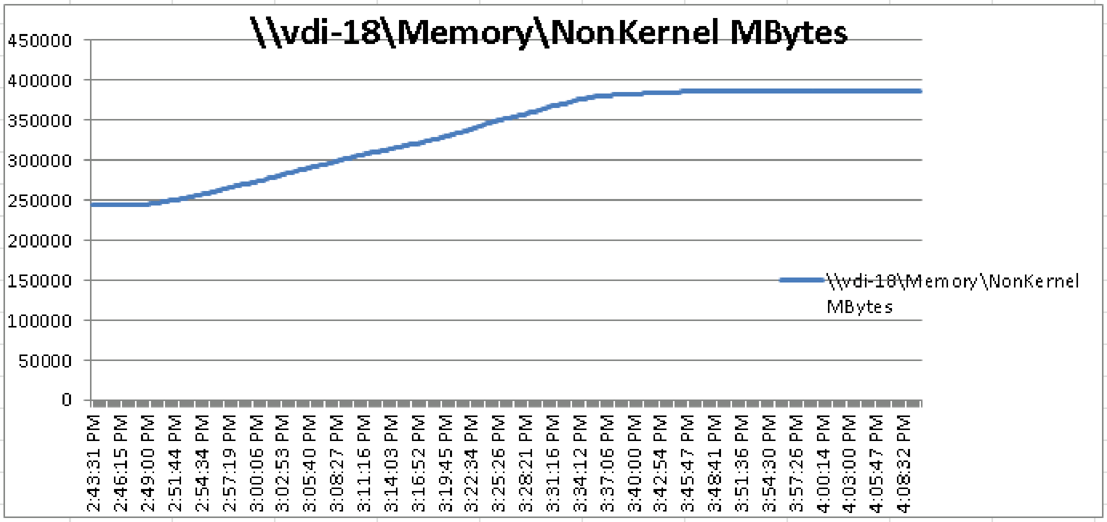

Figure 90 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

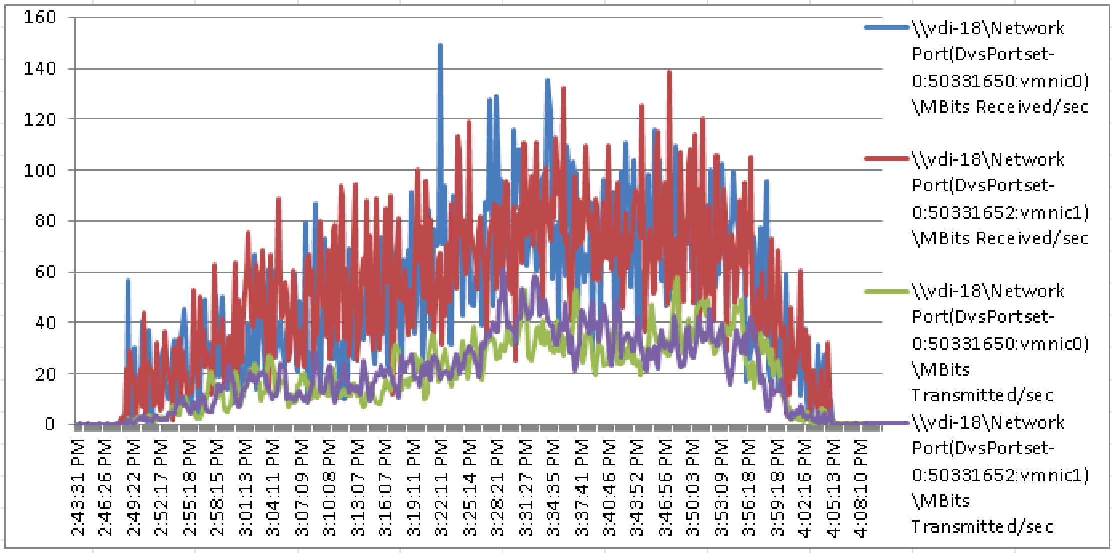

Figure 91 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

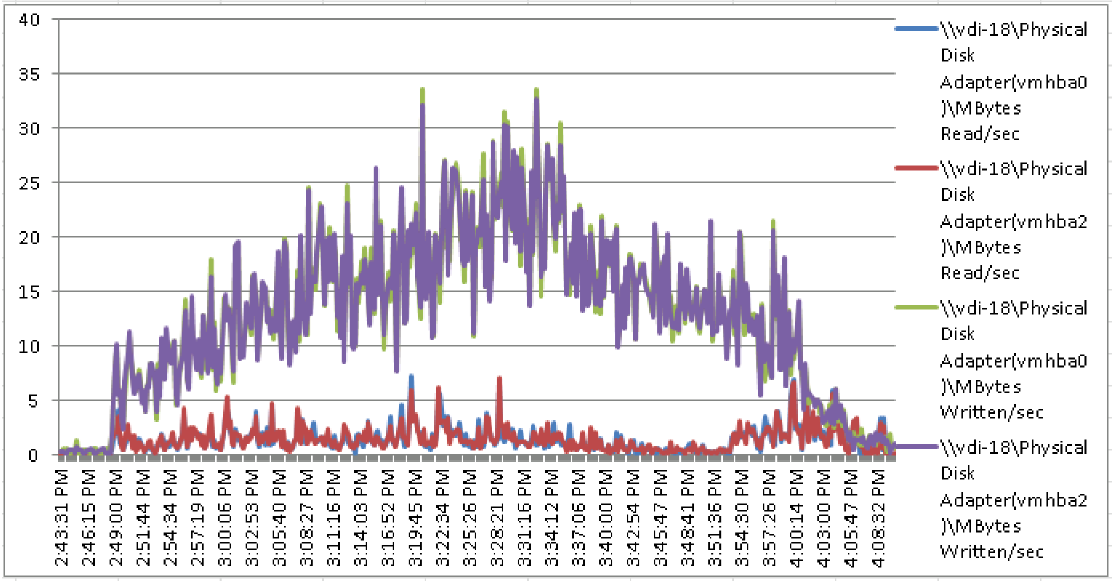

Storage Graphs for 5000 Users Mixed Workload Test

Figure 92 AFF A300 5000 Users Mixed Workload Scale Test | Read/Write Latency

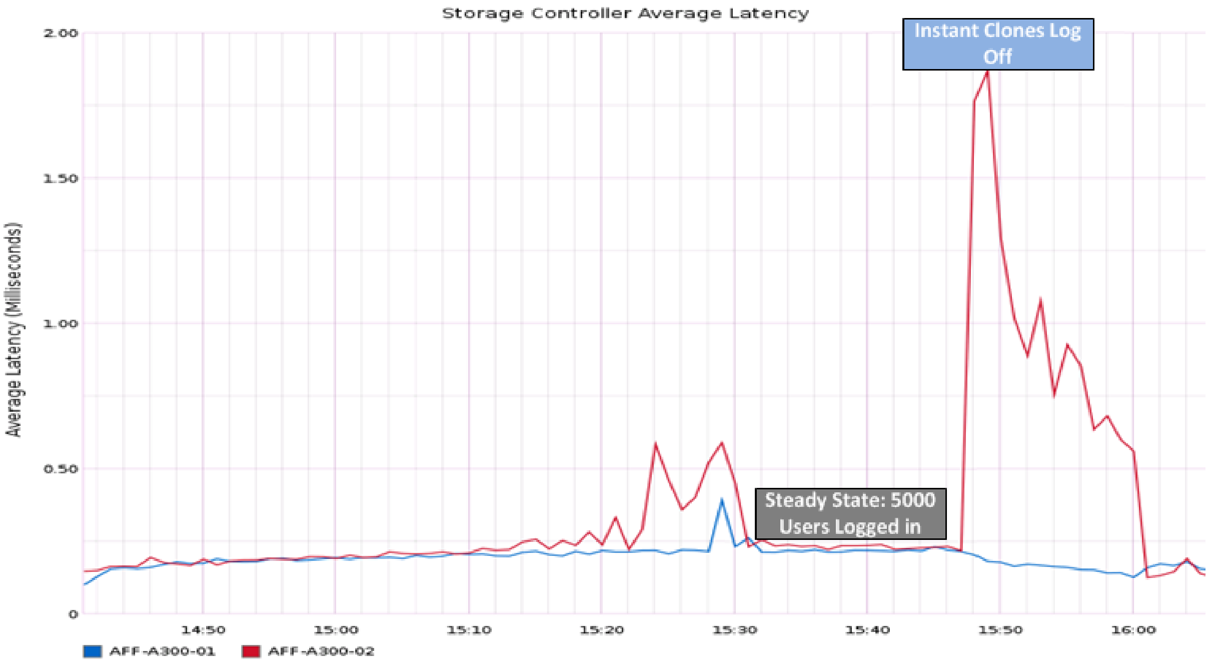

Figure 92 represents the read and write latency on both controllers in the storage system. As noted in the graph, read latency never exceeded 0.6ms and write latency never exceeded 0.4ms during the full-scale testing.

Figure 93 AFF A300 5000 Users Mixed Workload Scale Test | Storage controller Read Latency

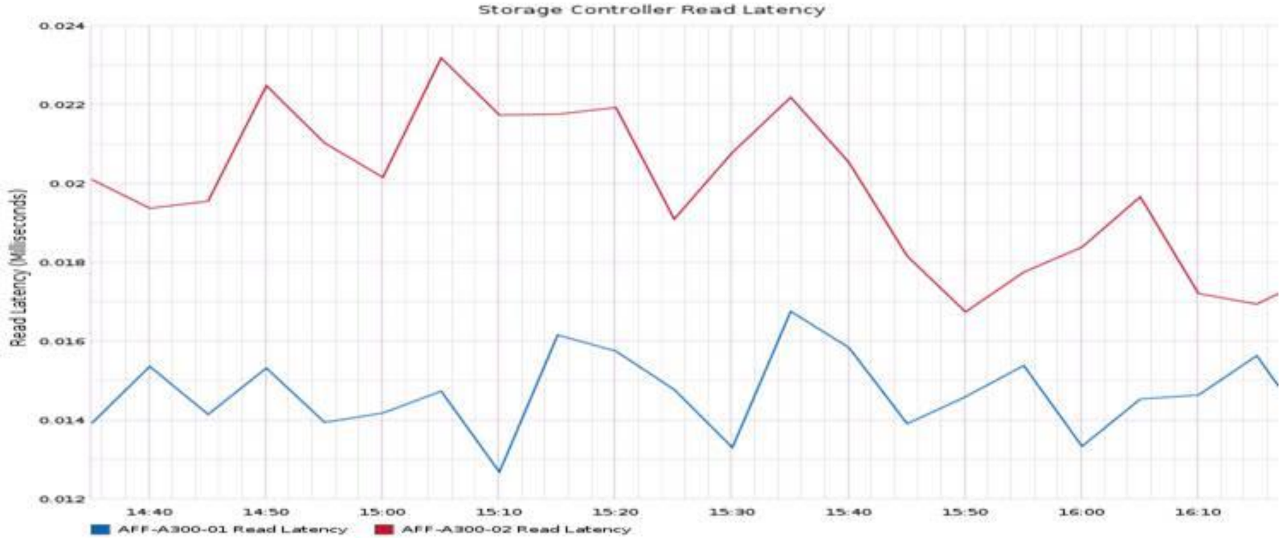

Figure 94 AFF A300 5000 Users Mixed Workload Scale Test | Storage controller Write Latency

Figure 95 AFF A300 5000 Users Mixed Workload Scale Test | Cluster Throughput

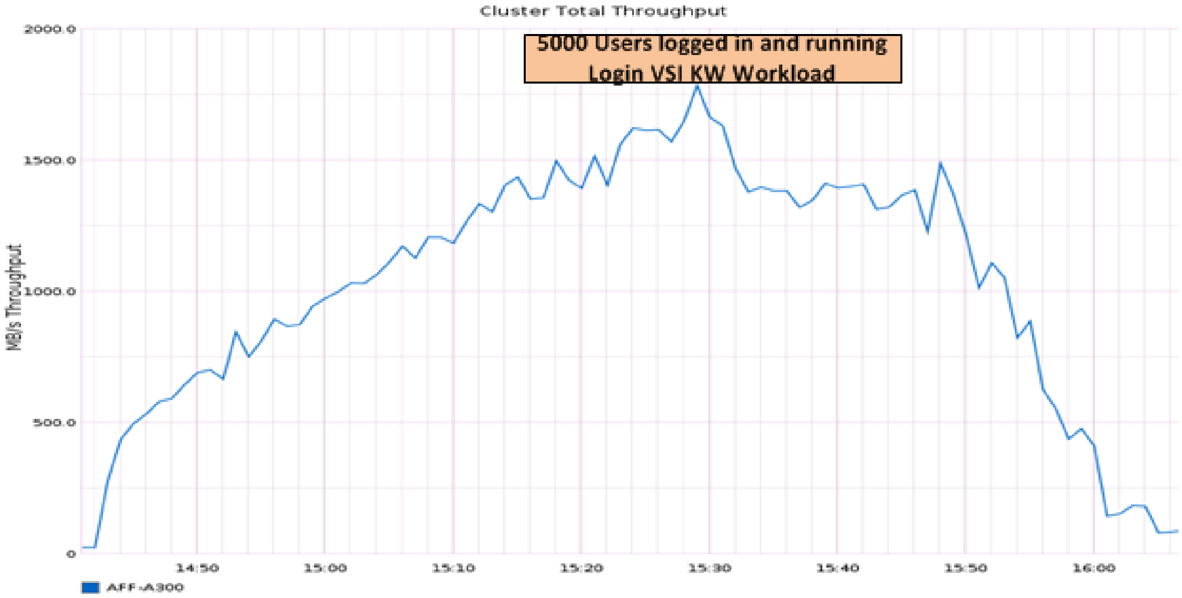

Figure 95 represents the total throughput of both storage controllers during the full-scale testing. As noted, total throughput peaked at approximately 1.75GB/sec during the full-scale testing.

Figure 96 AFF A300 5000 Users Mixed Workload Scale Test | Total IOPS

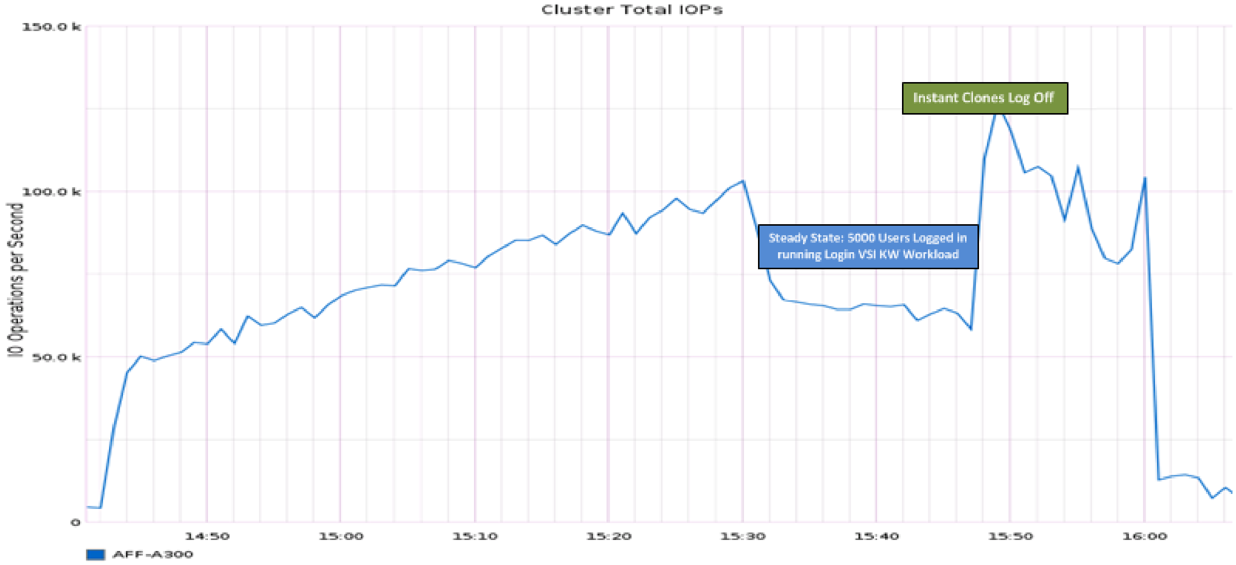

Figure 96 represents the total number of IO operations on both storage controllers during the full-scale testing. As noted in the graph, the A300 delivered approximately 145,000 IOPs at the peak of the testing.

Figure 97 Controller Headroom

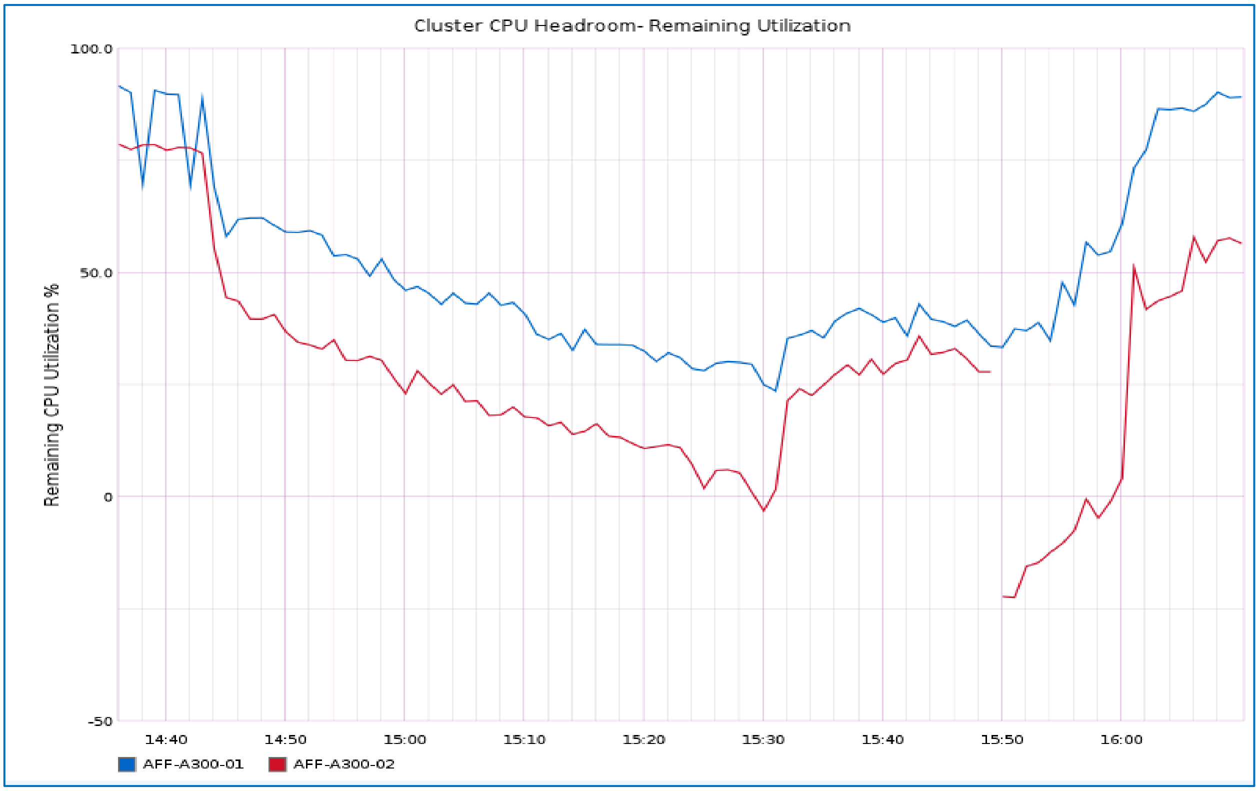

Figure 97 shows the remaining CPU performance capacity of the storage system. This metric can help administrators determine the growth capacity of an existing storage system. This metric considers HA failover requirements when determining headroom to allow for consistent performance even in the event of a controller failure. As seen in this graph, there is little remaining headroom on this storage system. If additional desktops were to be deployed, and additional A300 storage system would be required to maintain expected performance.

AFF A300 Storage Detailed Test Results for Cluster Scalability Test

This section highlights and provides analysis of the AFF A300 system performance results for each of the cluster test cases identified previously in this document. Specifically, it depicts and discusses the results for the following test scenarios:

- 1680 Windows Server 2016 RDS hosted VMware Horizon sessions
- 1660 Windows 10 x64 non- persistent VMware Horizon Users
- 1660 Windows 10 x64 persistent VMware Horizon Users

From a storage perspective, it is critical to maintain a latency of less than a millisecond for an optimal end-user experience no matter the IOPS and bandwidth being driven. The test results indicate that the AFF A300 storage delivers that essential minimum level of latency despite driving a substantial amount of IOPs and bandwidth for the thousands of desktops hosted on the AFF A300 system.

The sections that follow show screenshots of the AFF A300 storage test data for all three use case scenarios with information about IOPS, bandwidth, and latency at the peak of each use case test. In all three use cases (cluster level testing with RDSH, VDI Instant clones and full clones and mixed workload), the criteria followed prior to launching Login VSI workload test are the same.

Followed by this stage is a forty-minute window for all RDSH and VDI VMs to settle and then we begin the 2880-second Login VSI simulation phase when all sessions are ramping up and logging in.

AFF A300 Storage Test Results for 1680 RDS Windows 2016 Sessions

This test uses Login VSI as the workload generator in Benchmark mode with the Knowledge Worker user type and with VMware Remote Desktop Session Hosts (RDSH) Server Sessions as the VDI delivery mechanism. This first highlighted cluster test shows that the AFF A300 can easily handle this workload with exceptional end-user experience as confirmed from Login VSI.

The first AFF A300 GUI screenshot shows the performance during the 1680 RDSH sessions running on top of 72 Windows 2016 servers. As with all scenarios, there were three separate 1680 RDSH simulation runs completed in total, all with very similar results. As indicated in charts from one of these simulations, we maintained latency of less than or close to one millisecond (ms) for both read and write operations throughout this entire run. This resulted in a confirmed outstanding end-user experience for the simulated VMware RDSH users independently verified by Login VSI. It has been observed during this component of the testing with observed peak total values 35K IOPS and 600 MB/s. The latency is less than 0.20ms which lends itself to a very good end-user experience.

1680 RDSH Cluster Test: Storage Charts

Figure 98 AFF A300 1680 Users RDSH Cluster Test | Read/Write IOPS

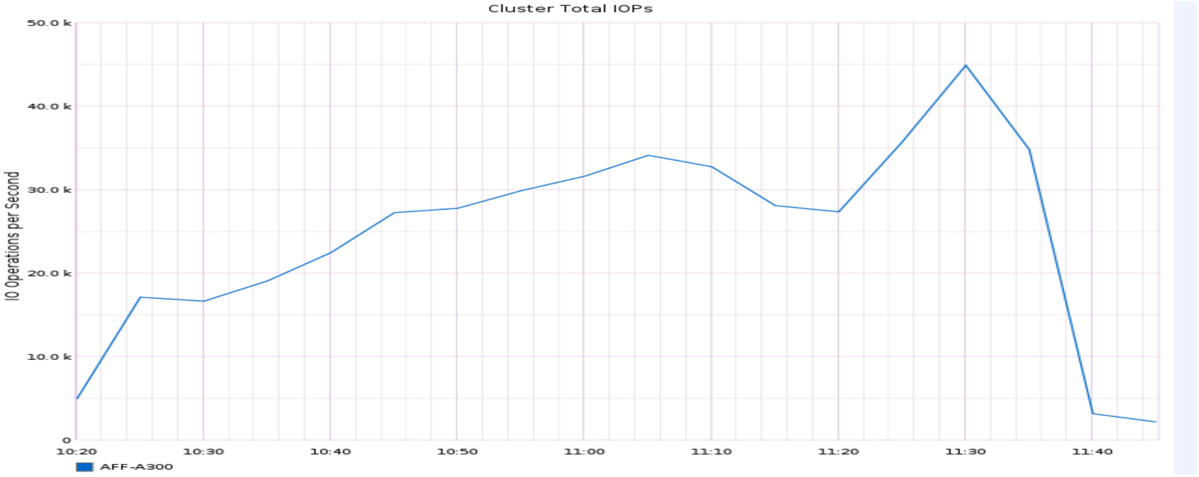

Figure 99 AFF A300 1680 Users RDSH Cluster Test | Throughput

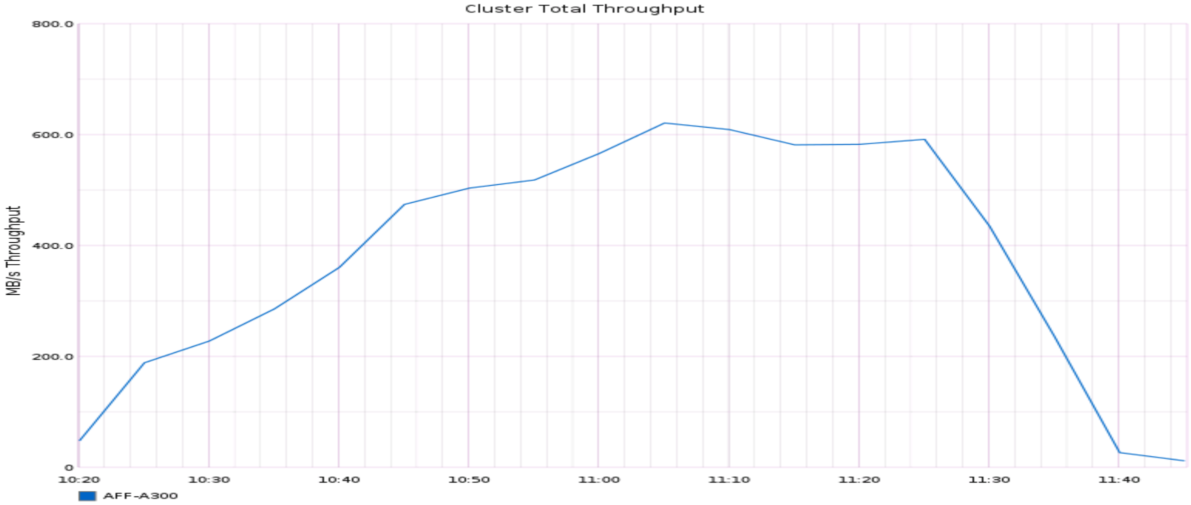

Figure 100 AFF A300 1680 RDSH Cluster Test | Read/Write Latency

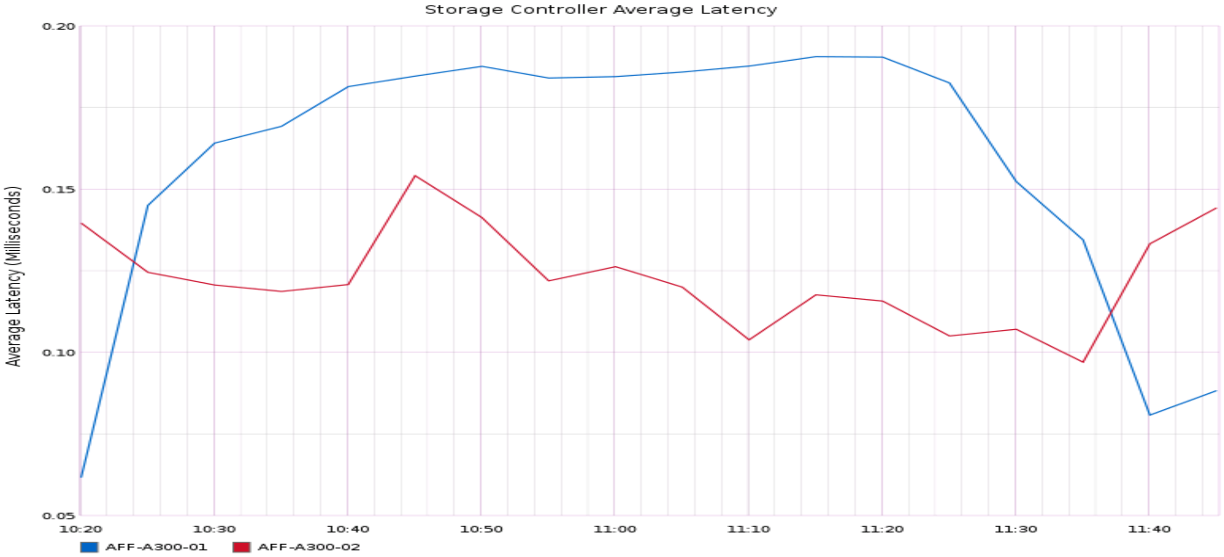

Figure 101 AFF A300 1680 RDSH Cluster Test Controller Processor Average

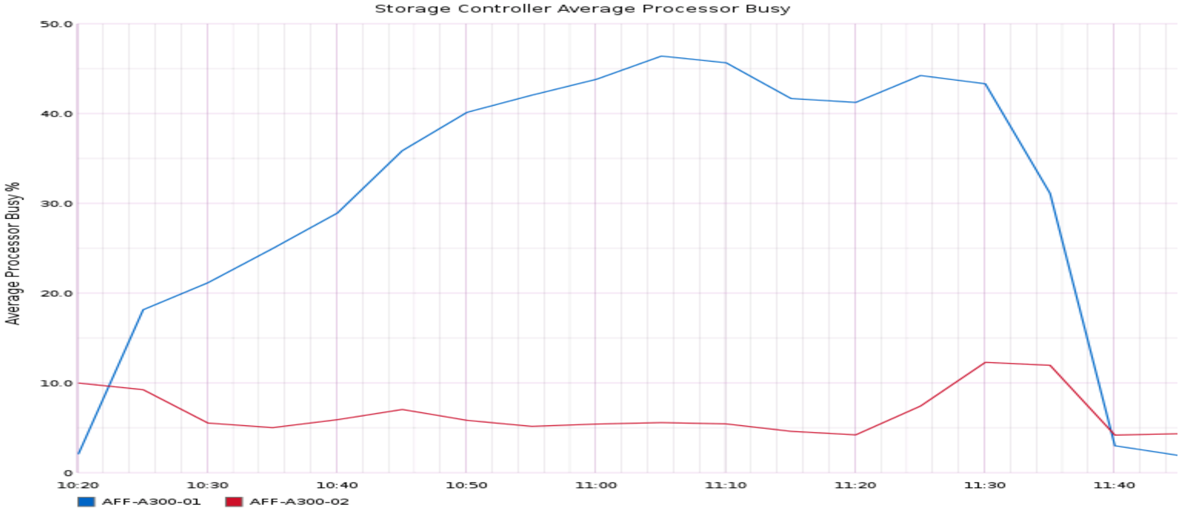

Figure 102 AFF A300 1680 RDSH Cluster Test CPU Headroom

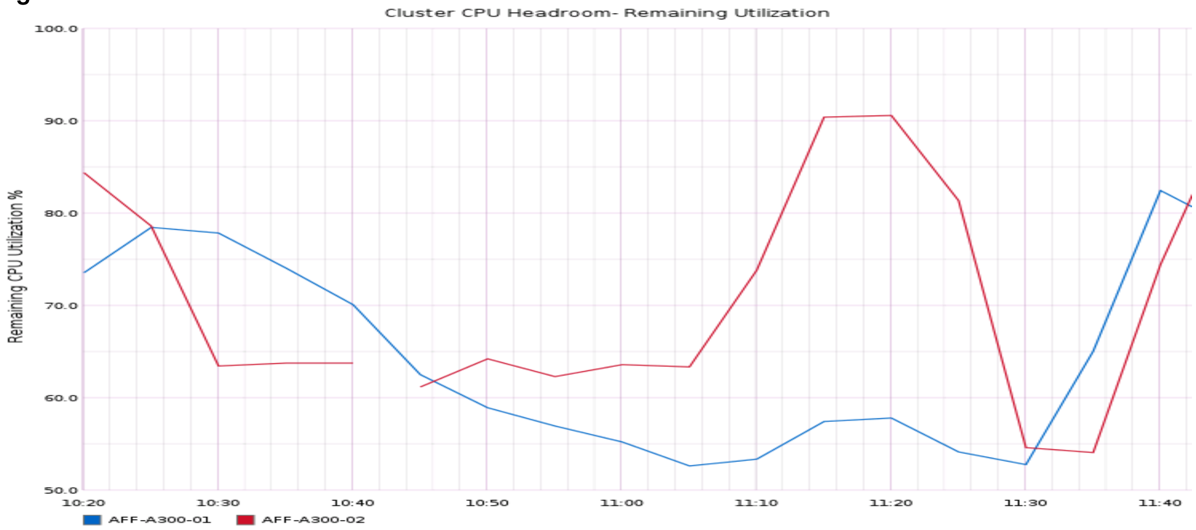

1660 Users VDI Instant Clones Cluster Test

NetApp AFF A300 Test Results for 1660 Non-Persistent Windows 10 x64 VMware Horizon Desktops

The next cluster-level simulation was to run 1660 non persistent Windows 10 x64 desktops against the same AFF- A300. All Login VSI parameters were kept consistent with bringing up all 1660 desktops created using VMware Horizon Composer Provisioning virtual machines. As indicated by the following storage metrics, the AFF A300 system was clearly able to handle this workload and continued to provide sub milli second latency for another impressive Login VSI result. It has been observed during this component of the testing with observed peak total values 42K IOPS and bandwidth of 600 MB/s. The latency is less than 0.20ms which results in a very good end user experience.

Figure 103 NetApp AFF A300| 1660 Users VDI Non-Persistent Cluster Test | Read/Write IOPS

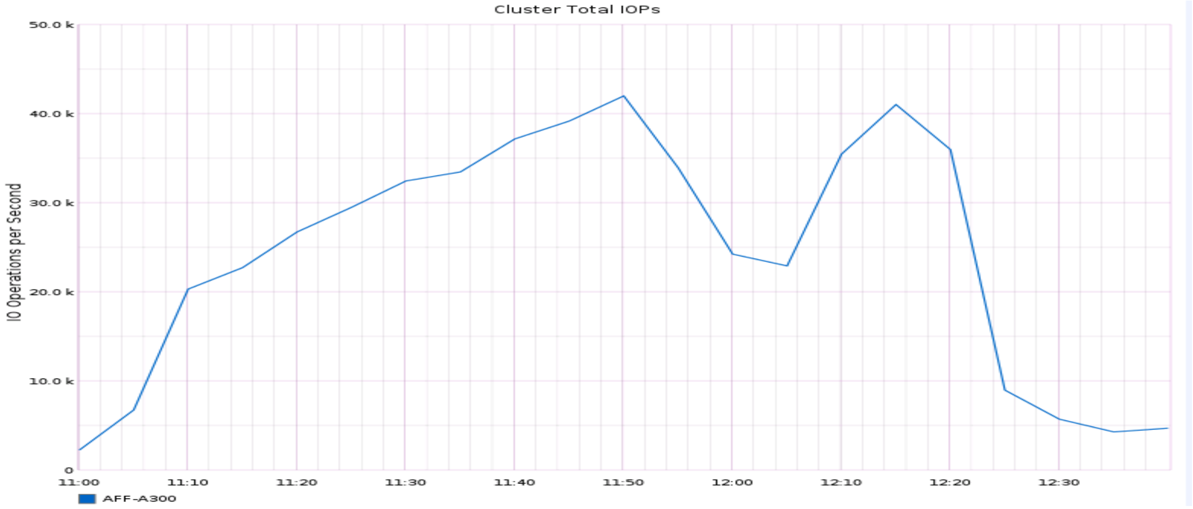

Figure 104 AFF A300| 1680 Users VDI Non-Persistent Cluster Test Bandwidth Transmitted /Received

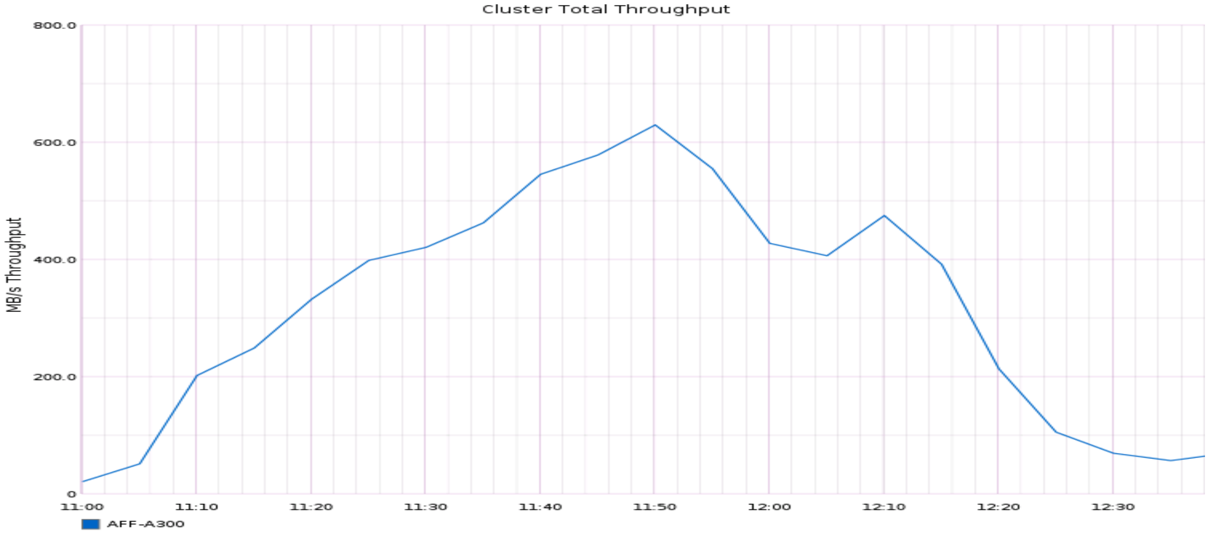

Figure 105 AFF A300| 1680 Users VDI Non-Persistent Cluster Test Latency (ms)

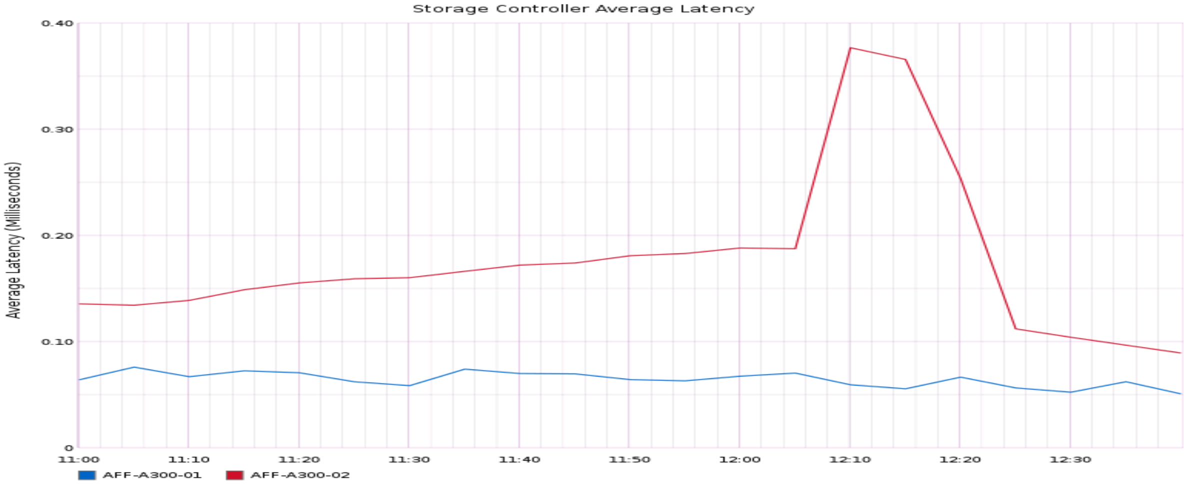

Figure 106 NetApp AFF A300| 1660 Users VDI Non-persistent Cluster Test Controller Processor Average

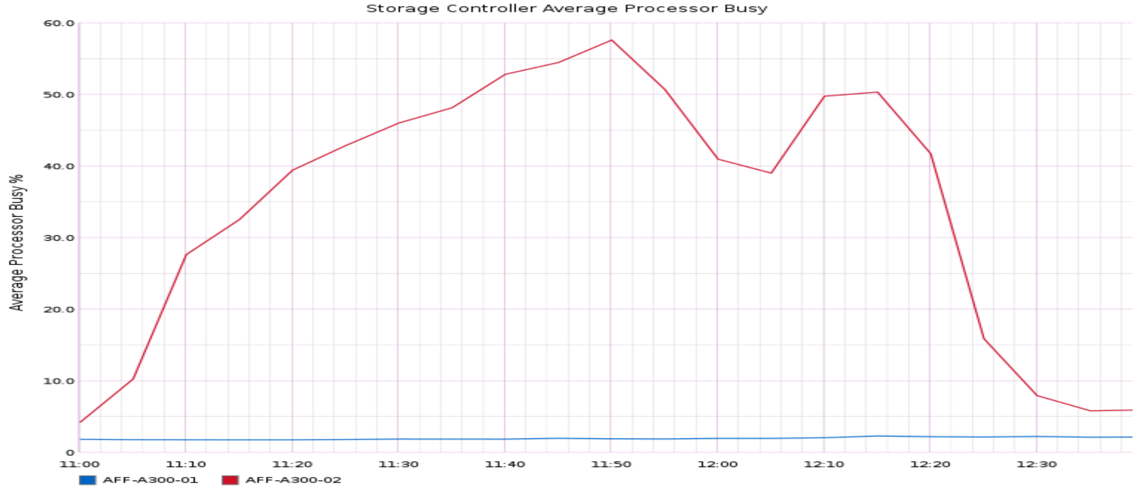

Figure 107 NetApp AFF A300| 1660 Users VDI Non-Persistent Cluster Test Controller CPU Headroom

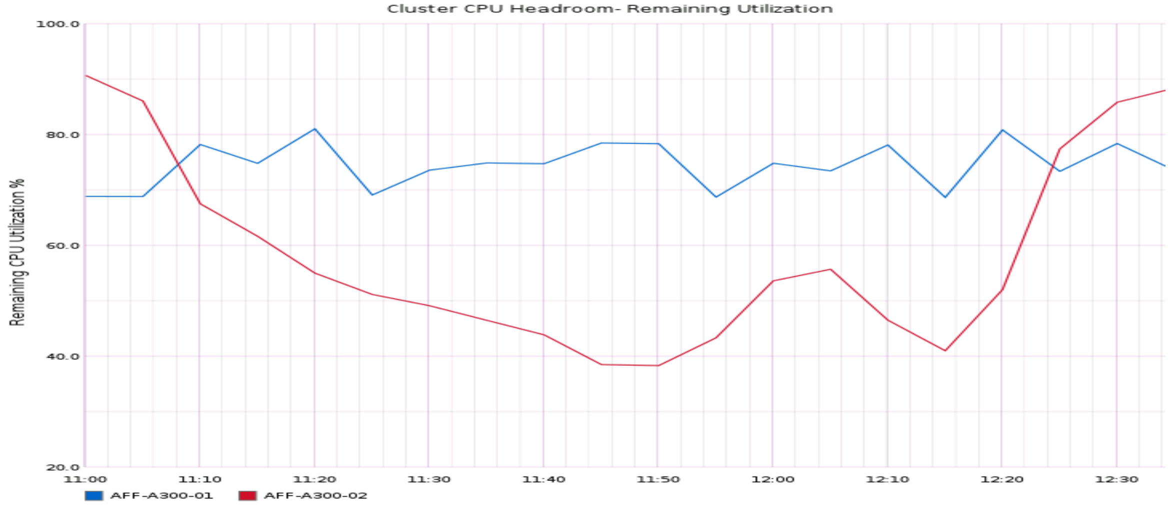

1660 Users VDI Persistent Clones Cluster Test

NetApp AFF A300 Test Results for 1660 Persistent Windows 10x64 VMware Horizon Desktops

The next cluster-level simulation was to run 1660 non persistent Windows 10x64 desktops against the same AFF- A300. All Login VSI parameters were kept consistent with bringing up all 1660 desktops created using VMware Horizon Composer Provisioning virtual machines. As indicated by the following storage metrics, the AFF A300 system was clearly able to handle this workload and continued to provide sub milli second latency for another impressive Login VSI result. It has been observed during this component of the testing with observed peak total values 21K IOPS and bandwidth of 430 MB/s. The latency is less than 0.27ms which results in a very good end user experience

Figure 108 NetApp AFF A300| 1660 Users VDI Persistent Cluster Test | Read/Write IOPS

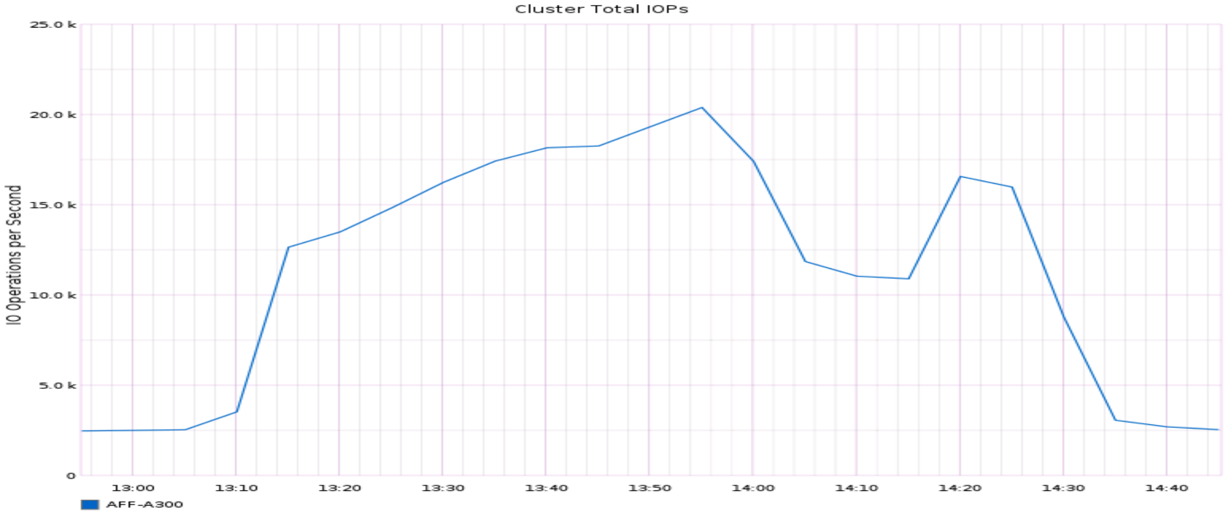

Figure 109 NetApp AFF A300| 1660 Users VDI Persistent Cluster Test | Throughput

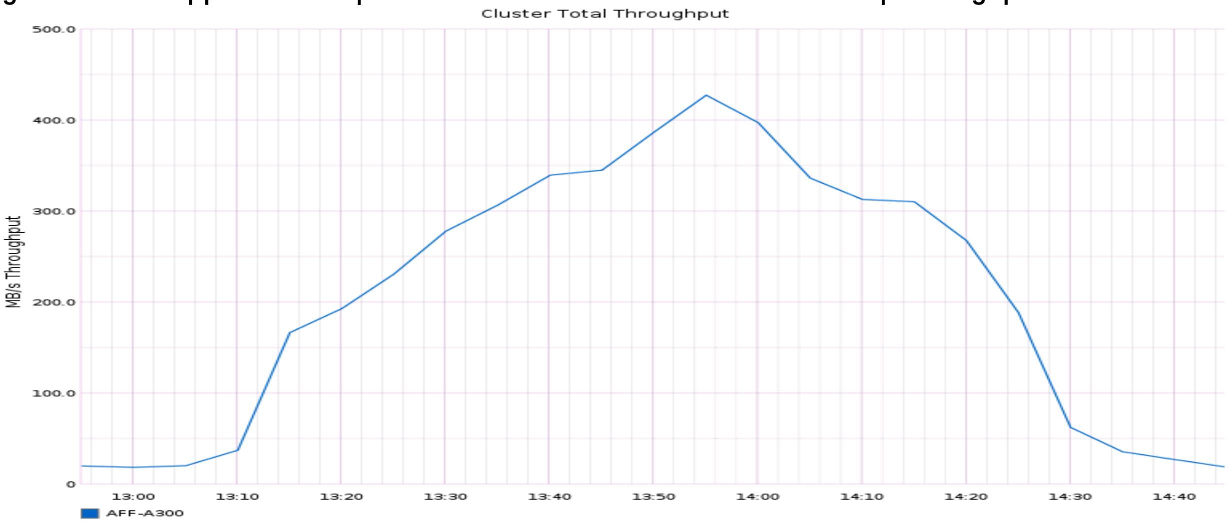

Figure 110 NetApp AFF A300| 1660 Users VDI Persistent Cluster Test | Latency (ms)

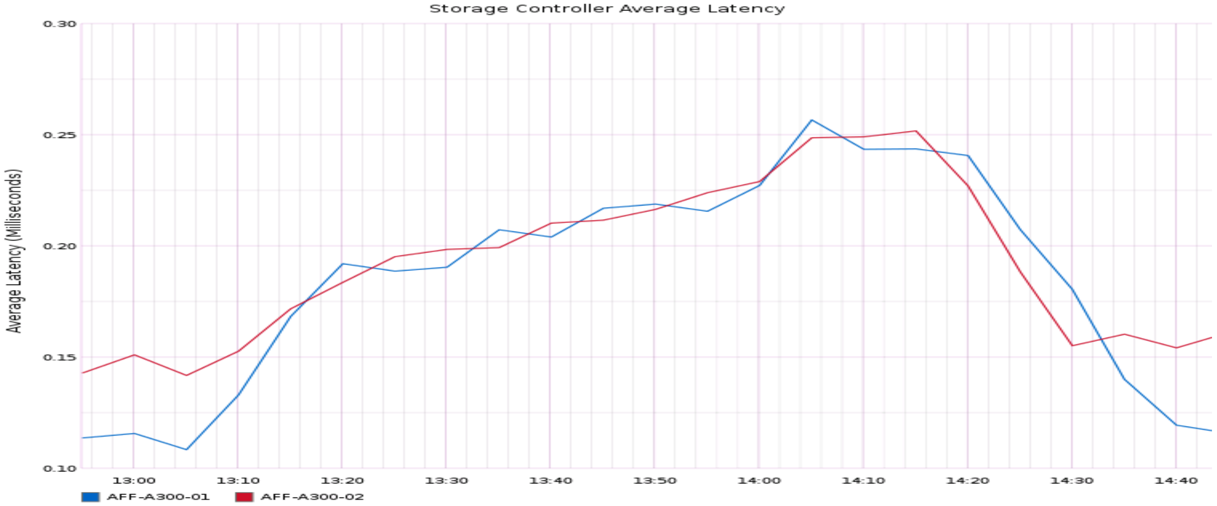

Figure 111 NetApp AFF A300| 1660 Users VDI Persistent Cluster Test | Controller Processor Average

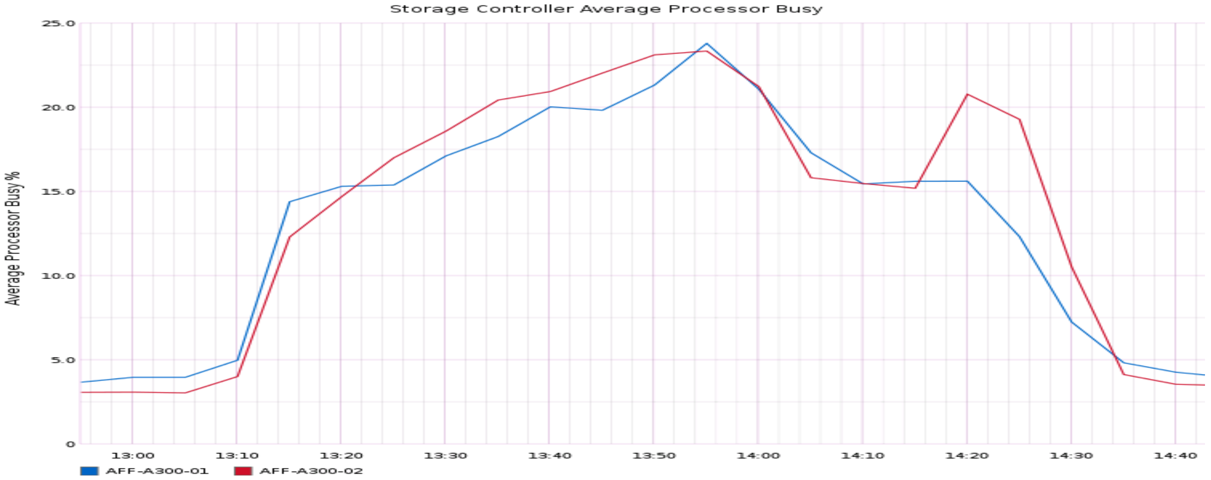

Figure 112 NetApp AFF A300| 1660 Users VDI Persistent Cluster Test | Controller CPU headroom

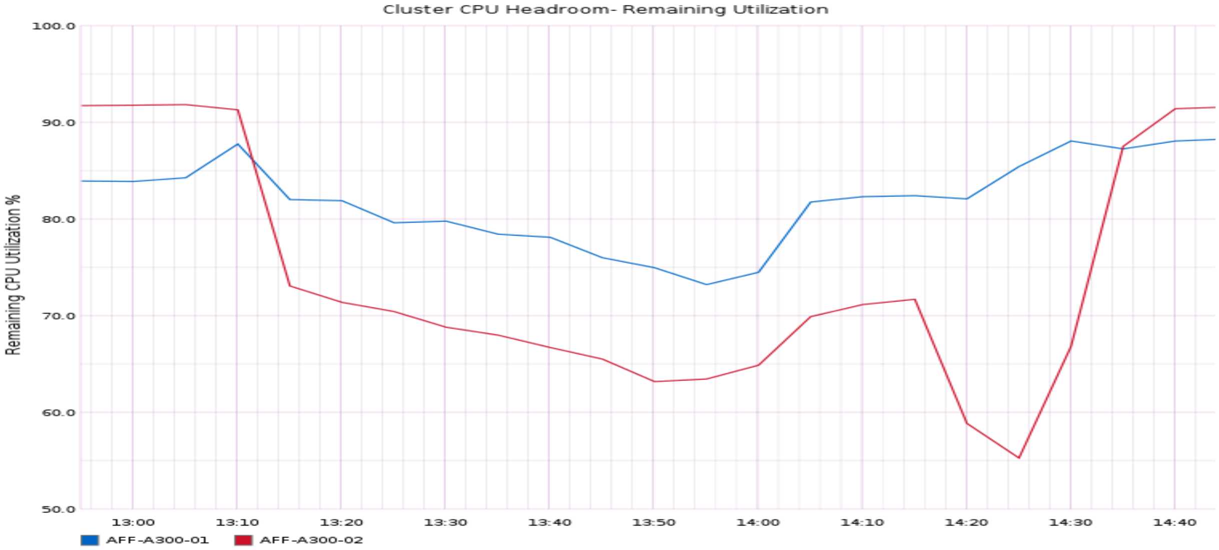

NetApp AFF A300 Storage Test Results for 5000 User Full Scale, Mixed Workload Scalability

The next simulation shows the results of combining all of our previous cluster tests of 1680 VMware Horizon RDS sessions and 1660 non persistent VDI Instant Clone and 1660 Full Clone virtual machines sessions for a full 5000 user VMware Horizon RDSH and VMware Horizon VDI simulation on the same AFF A300 array.

The performance and consistent results indicate an outstanding user experience on a very large scale.

The screenshot below shows the AFF A300 GUI with the cursor providing detailed metrics at the start of the simulation during the boot storm of the desktops. Despite driving nearly 2GB/s in bandwidth during the test itself, we maintain the desktop responsiveness and performance of low latency throughout the entire test.

It has been observed during this component of the testing with observed peak total values 110K IOPS and a bandwidth of 2000 MB/s. The latency is again less than 0.7ms, which is a very good end-user experience on mixed use case scale testing scenario.

Figure 113 AFF A300 5000 Users Mixed Workload Read/Write IOPS

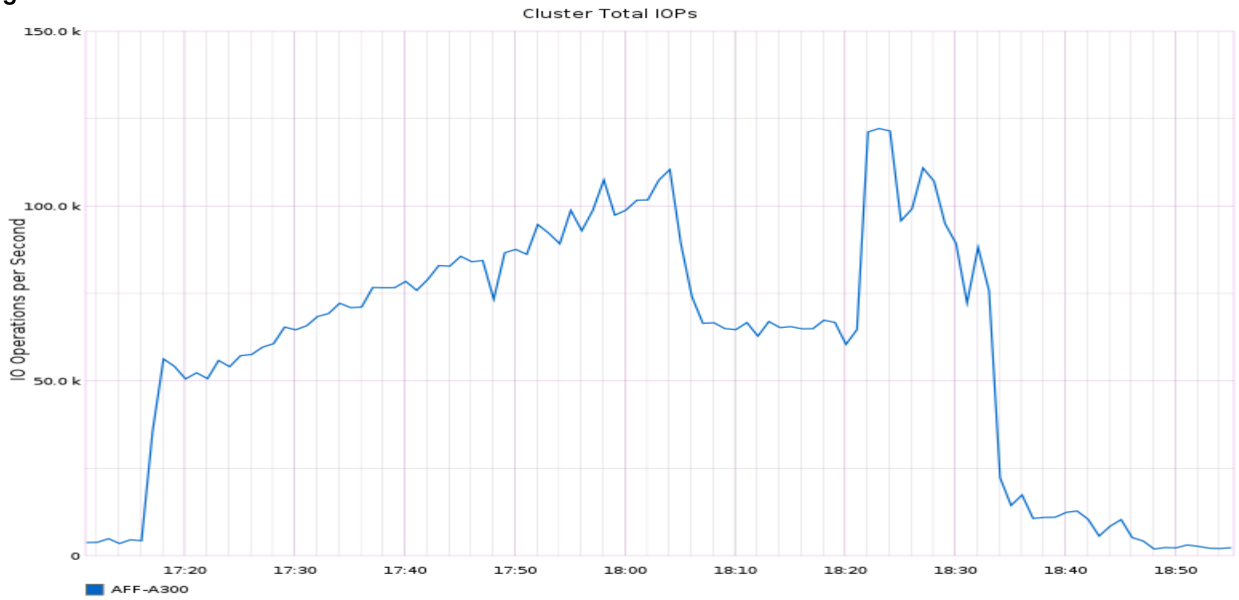

Figure 114 AFF A300 | 5000 Users Mixed Workload Scale Test | Throughput MB/S

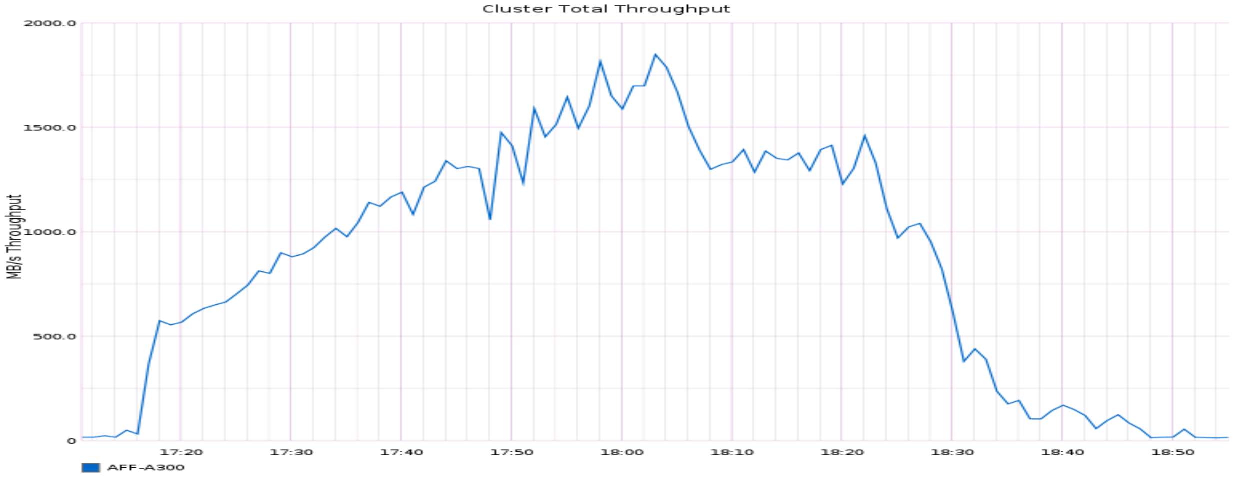

Figure 115 AFF A300 5000 Users Mixed Workload Latency (ms)

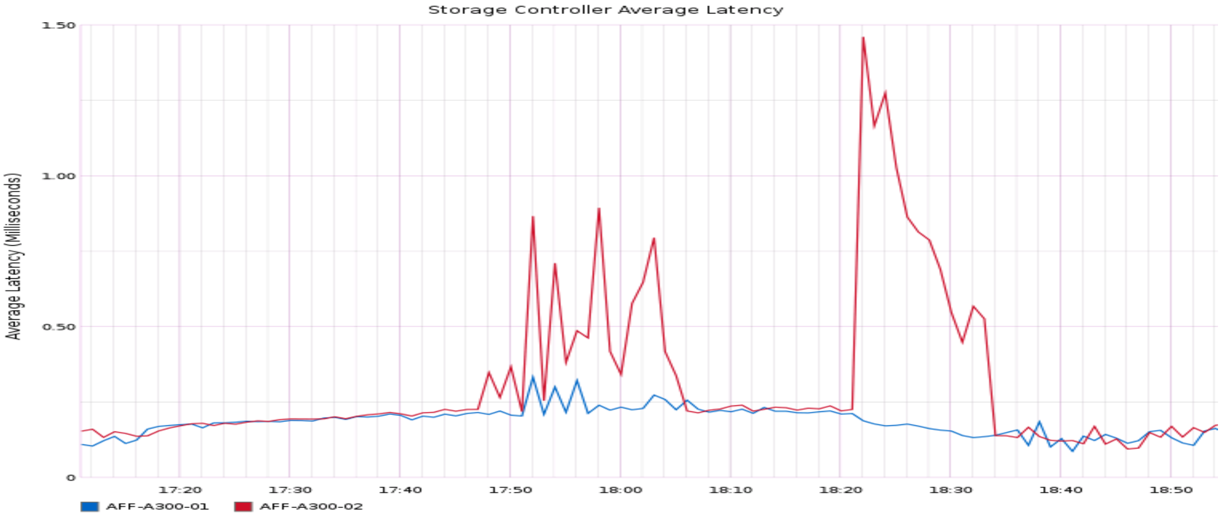

The VMware Horizon Administrator Console reports all 5000 sessions /desktops (1680 RDSH sessions and 3320 VDI virtual machines, 2 pools each with 1660 non-persistent and persistent desktops provisioned) have been logged in during the Login VSI testing.

Figure 116 VMware Horizon Administrator Console

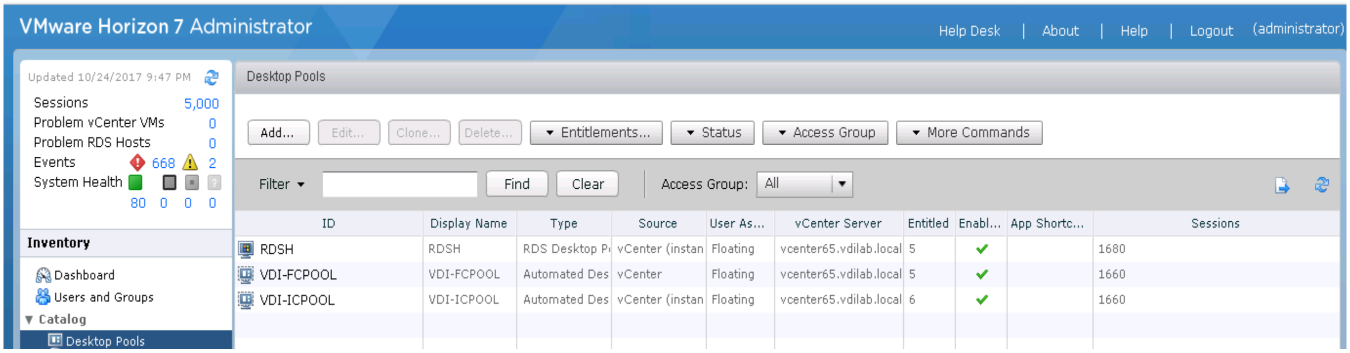

VMware Horizon 7 Administrator

Help Desk | About | Help | Logout (administrator)

Updated 10/24/2017 9:47 PM

Sessions	5,000
Problem vCenter VMs	0
Problem RDS Hosts	0
Events	668 ⚠ 2
System Health	80 0 0 0

Inventory

Dashboard
Users and Groups
▼ Catalog
Desktop Pools

Desktop Pools

Add... Edit... Clone... Delete... ▾ Entitlements... ▾ Status ▾ Access Group ▾ More Commands
Filter ▾ Find Clear

Access Group: All

ID	Display Name	Type	Source	User As...	vCenter Server	Entitled	Enabl...	App Shortc...	Sessions
RDSH	RDSH	RDS Desktop P	vCenter (instan	Floating	vcenter65.vdilab.local	5	✓		1680
VDI-FCPOOL	VDI-FCPOOL	Automated Des	vCenter	Floating	vcenter65.vdilab.local	5	✓		1660
VDI-ICPOOL	VDI-ICPOOL	Automated Des	vCenter (instan	Floating	vcenter65.vdilab.local	6	✓		1660

This was a pristine lab environment; you can see that our data reduction and overall array utilization during this full test scenario was extremely impressive with only less than 30 percent 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.

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 1000 Users, two chassis 8 mixed workload RDSH / VDI 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 6332UP Fabric Interconnect. A single Cisco UCS domain can grow to 160 blades for an enterprise deployment.
- Cisco UCS Central, the manager of managers, extends Cisco 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 Cisco UCS fabric, Cisco Nexus physical switches, and Cisco 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 or more Ethernet uplinks are needed to be configured on the Cisco UCS 6332UP Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the NetApp AFF A300 storage scaling section. Please refer the AFF A300 website for scalability guidelines.

Scalability of VMware Horizon 7 Configuration

VMware Horizon environments can scale to large numbers. When implementing VMware Horizon, 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 and Data protection requirements
- For VMware Horizon pooled desktops, the disk size and memory requirements.

These and other various aspects of scalability considerations are described in greater detail in “VMware Horizon Reference Architecture” document and should be a part of any VMware design.

When designing and deploying this CVD environment, best practices were followed including the following:

- VMware 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.

Summary

FlexPod delivers a platform for Enterprise End User Computing deployments and cloud datacenters using Cisco UCS Blade and Rack Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches and FC-attached NetApp AFF A300. FlexPod is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlexPod provides to customers wishing to deploy enterprise-class VDI 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, NetApp AFF A300 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 NetApp Advanced Consulting Services provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Authors

Ramesh Guduru, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Ramesh is a VMware Horizon and Cisco UCS subject matter expert in the Cisco Computer Systems Product Group's Solution Team based in San Jose, CA. He has authored several Cisco Validated Designs for VMware Horizon.

David Arnette, Solutions Architect, NetApp Inc.

David is a Solution Architect with NetApp's Converged Infrastructure Group based in RTP, North Carolina. He has authored CVD and NVA reference architectures for numerous enterprise workloads including VMware Horizon VDI, Microsoft SQL and SAP HANA.

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:

- Mike Brennan, Product and Technical Marketing Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

References

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

Cisco UCS B-Series Servers

- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html
- https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b200m5-specsheet.pdf
- https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/datasheet-listing.html
- https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-b200-m5-blade-server/model.html
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/blade-servers/B200M5.pdf

Cisco UCS Manager Configuration Guides

- http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
- 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

Cisco MDS 9000 Service Switch References

- http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html
- http://www.cisco.com/c/en/us/products/storage-networking/product-listing.html
- http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/datasheet-listing.html

FlexPod

- https://www.flexpod.com
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1design.html?referring_site=RE&pos=1&page=https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

VMware References

- https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
- http://pubs.vmware.com/Release_Notes/en/horizon-7-view/horizon-703-view-release-notes.html
- https://labs.vmware.com/flings/vmware-os-optimization-tool
- https://pubs.vmware.com/view-51/index.jsp?topic=%2Fcom.vmware.view.planning.doc%2FGUID-6CAFE558-A0AB-4894-A0F4-97CF556784A9.html
- https://pubs.vmware.com/horizon-7-view/index.jsp?topic=%2Fcom.vmware.horizon-view.desktops.doc%2FGUID-DFAD071A-7F60-4720-86AB-8F1597BFC95C.html

Microsoft References

- https://technet.microsoft.com/en-us/library/hh831620(v=ws.11).aspx
- https://technet.microsoft.com/en-us/library/dn281793(v=ws.11).aspx
- https://support.microsoft.com/en-us/kb/2833839
- https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

NetApp Reference Documents

- http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx
- http://www.netapp.com/us/products/data-management-software/ontap.aspx
- https://mysupport.netapp.com/documentation/docweb/index.html?productID=62379&language=en-US
- http://www.netapp.com/us/products/management-software/
- http://www.netapp.com/us/products/management-software/vsc/

Appendix A – Cisco Nexus Ethernet and MDS Fibre Channel Switch Configurations

Ethernet Network Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000 and 1000V Switches used in this study.

Cisco Nexus 9372PX-A Configuration

```
!Command: show running-config

!Time: Wed Nov  30 13:14:54 2017

version 7.0(3)I1(3b)
switchname DV-POD-2-N9K-A
vdc DV-Pod-2-N9K-A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature 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 $1$tYYajkfc$7P7nLjWYvfTWAlvFDnwJZ.  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 0xf747567d6cfecf362a9641ac6f3cefc9 priv
0xf747567d6cfecf362a9641ac6f3cefc9 localizedkey

rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.10.160.2
ntp peer 10.10.160.3
ntp server 10.81.254.202 use-vrf management
vlan 1-2,60-70,102,164
vlan 60
  name In-Band-Mgmt
```

```
vlan 61
  name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name NFS
vlan 66
  name vMotion
vlan 68
  name LauncherPXE
vlan 69
name Launcher81
vlan 102
  name VDI
vlan 164
  name Out-Of-Band-Mgmt

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
hardware qos ns-buffer-profile mesh
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.66 source 10.29.164.65
  delay restore 150
  peer-gateway
  auto-recovery

interface Vlan1
  description default native vlan
  no shutdown
  no ip redirects
  ip address 10.29.164.253/24
  no ipv6 redirects

interface Vlan60
  description In Band Mgmt vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.2/24
  hsrp version 2
  hsrp 60
    preempt
    ip 10.10.60.1

interface Vlan61
  description Infrastructure vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
    preempt
    ip 10.10.61.1
```

```
interface Vlan62
  description CIFS vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.62.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 62
    preempt
    priority 110
    ip 10.10.62.1

interface Vlan63
  description NFS vlan 63
  no shutdown
  no ip redirects
  ip address 10.10.63.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 63
    preempt
    ip 10.10.62.1
interface Vlan66
  description vMotion network vlan 66
  no shutdown
  ip address 10.10.66.2/24
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1
interface Vlan68
  description VSI Launchers vlan 68
  no shutdown
  ip address 10.10.68.2/23
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
  ip dhcp relay address 10.10.61.30
interface Vlan69
  description LoginVSI Launchers 10.10.81.network vlan 69
  no shutdown
  ip address 10.10.81.2/24
  hsrp version 2
  hsrp 69
    preempt
    ip 10.10.69.1
  ip dhcp relay address 10.10.61.30
interface Vlan102
  description VDI vlan 102
  no shutdown
  ip address 10.2.2/19
  no ipv6 redirects
  hsrp version 2
  hsrp 102
    preempt delay minimum 240
    priority 110
    ip 10.2.0.1
  ip dhcp relay address 10.10.61.30
  ip dhcp relay address 10.10.61.31

interface port-channel10
  description VPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type network
  vpc peer-link

interface port-channel1
description FI-A_6k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel2
description FI-B_6k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102-164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel15

  description FI-A_6k_Launchers-Uplink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  mtu 9216

  vpc 15

interface port-channel16

  description FI-B_6k_Launchers-Uplink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  mtu 9216

  vpc 16

interface port-channel51

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  speed 40000

  no negotiate auto

  mtu 9216

  vpc 51
```

```
interface port-channel52

 switchport mode trunk

 switchport trunk allowed vlan 1-2,60-70,102,164

 spanning-tree port type edge trunk

 speed 40000

 no negotiate auto

 mtu 9216

 vpc 52

interface port-channel53

 switchport mode trunk

 spanning-tree port type edge trunk

 mtu 9216

 vpc 53

interface port-channel54

 switchport mode trunk

 spanning-tree port type edge trunk

 mtu 9216

 vpc 54

interface Ethernet1/1
interface Ethernet1/2
interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
 description Uplink_from_FI-A_to-N9KA

 switchport mode trunk

 switchport trunk allowed vlan 1-2,60-70,102,164

 mtu 9216

 channel-group 12 mode active

interface Ethernet1/36
 description Uplink_from_FI-B_to-N9KB

 switchport mode trunk

 switchport trunk allowed vlan 1-2,60-70,102,164

 mtu 9216

 channel-group 12 mode active

interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45

  description Uplink_from_LoginVSI_Launchers_FI-A

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  mtu 9216

  channel-group 15 mode active

interface Ethernet1/46

  description Uplink_from_LoginVSI_Launchers_FI-B

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  mtu 9216


interface Ethernet1/47
interface Ethernet1/48
 description TOR
 switchport access vlan 164
interface Ethernet1/49

  description VPC Peer Link between 9ks

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  channel-group 10 mode active

interface Ethernet1/50

  description VPC Peer Link between 9ks

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  channel-group 10 mode active

interface Ethernet1/51
  description FI-A-N9K-Uplink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  speed 40000

  no negotiate mtu 9216

  channel-group 51 mode active

interface Ethernet1/52
  description FI-B-N9K-Uplink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  speed 40000

  no negotiate auto

  mtu 9216

  channel-group 52 mode active

interface Ethernet1/53
switchport mode trunk
mtu 9216
channel-group 53 mode active

interface Ethernet1/54
switchport mode trunk
```

```
mtu 9216
channel-group 54 mode active

interface mgmt0
vrf member management
ip address 10.29.164.65/24
line console
line vty

boot nxos bootflash://sup-1/n9000-dk9.7.0.3.I1.3b.bin
```

Cisco Nexus 9172PX-B Configuration

```
DV-Pod-2-N9K-B# show running-config

!Command: show running-config

!Time: Wed Nov  30 13:55:44 2017

version 7.0(3)I1(3b)

switchname DV-Pod-2-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 $1$fp3LrGLC$PF8eML85qkPBgdH/bZAKK/  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 1,60-70,102,164
vlan 60
```

```
  name In-Band-Mgmt
vlan 61
  name Infra-Mgmt
vlan 62
  name CIFS
vlan 63
  name CIFS
vlan 66
  name vMotion
vlan 68
  name LauncherPXE
vlan 69
  name Launcher81
vlan 102
  name VDI
vlan 164
  name Out-Of-Band-Mgmt

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
hardware qos ns-buffer-profile mesh
vpc domain 10
  role priority 20
  peer-keepalive destination 10.29.164.65 source 10.29.164.66
  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 Vlan60
  description In Band Mmgmt vlan 60
  no shutdown
  no ip redirects
  ip address 10.10.60.3/24
  hsrp version 2
  hsrp 60
    preempt
    ip 10.10.60.1

interface Vlan61
  description Infrastructure Mgmt vlan 61
  no shutdown
  no ip redirects
  ip address 10.10.61.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 61
    preempt
    ip 10.10.61.1
interface Vlan62
```

```
  description CIFS vlan 62
  no shutdown
  ip address 10.10.62.3/24
  no ipv redirects
  hsrp version 2
  hsrp 62
  priority 110
    preempt
    ip 10.10.62.1
interface Vlan63
  description NFS vlan 63
  no shutdown
  ip address 10.10.63.3/24
  no ipv redirects
  hsrp version 2
  hsrp 63
    preempt
    ip 10.10.63.1
interface Vlan66
  description vMotion network vlan 66
  no shutdown
  ip address 10.10.66.3/24
  hsrp version 2
  hsrp 66
    preempt
    ip 10.10.66.1
interface Vlan68
  description LoginVSI Launchers vlan 68
  no shutdown
  no ip redirects
  ip address 10.10.68.3/23
  no ipv6 redirects
  hsrp version 2
  hsrp 68
    preempt
    ip 10.10.68.1
  ip dhcp relay address 10.10.61.30

interface Vlan69
  description LoginVSI Launchers 10.10.81-network vlan 69
  no shutdown
  no ip redirects
  ip address 10.10.81.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 69
    preempt
    ip 10.10.81.1
  ip dhcp relay address 10.10.61.30

interface Vlan102
  description VDI vlan 102
  no shutdown
  no ip redirects
  ip address 10.2.0.3/19
  no ipv6 redirects
  hsrp version 2
  hsrp 102
    preempt delay minimum 240
    priority 110
    ip 10.2.0.1
  ip dhcp relay address 10.10.61.30
  ip dhcp relay address 10.10.61.31
```

```
interface port-channel10
  description VPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type network
  vpc peer-link

interface port-channel11
description FI-A_6k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
description FI-B_6k_UCS-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel15
description FI-A_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 15
interface port-channel16
description FI-B_6k_Launchers-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 16
interface port-channel51
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 51
interface port-channel52
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  spanning-tree port type edge trunk
  mtu 9216
  vpc 52

interface port-channel53
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 53
interface port-channel54
  switchport mode trunk
  spanning-tree port type edge trunk
  mtu 9216
  vpc 54

interface Ethernet1/1
interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34
interface Ethernet1/35
  description Uplink_fromFI-A-to_N9KA
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/36
  description Uplink_fromFI-B-to_N9KB
  switchport mode trunk
  switchport trunk allowed vlan 1-2,60-70,102,164
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45
description Uplink_from_LoginVSI_Launchers_FI-A

  switchport mode trunk

  switchport trunk allowed vlan 1,60-70,102,164

  mtu 9216

  channel-group 15 mode active

interface Ethernet1/46
description Uplink_from_LoginVSI_Launchers_FI-B

  switchport mode trunk

  switchport trunk allowed vlan 1,60-70,102,164

  mtu 9216

  channel-group 16 mode active

interface Ethernet1/47

interface Ethernet1/48
description TOR
switchport access vlan 164
interface Ethernet1/49

  description VPC Peer Link between 9ks

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164
  channel-group 10 mode active

interface Ethernet1/50
  description VPC Peer Link between 9ks

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  channel-group 10 mode active

interface Ethernet1/51
  description FI-A-N9KUPlink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  speed 40000

   no negotiate auto

 mtu 9216

 channel-group 51 mode active

interface Ethernet1/52
  description FI-A-N9KUPlink

  switchport mode trunk

  switchport trunk allowed vlan 1-2,60-70,102,164

  spanning-tree port type edge trunk

  speed 40000

  no negotiate auto

  mtu 9216

 channel-group 52 mode active

interface Ethernet1/53
 switchport mode trunk
 mtu 9216
 channel-group 53 mode active

interface Ethernet1/54
 switchport mode trunk
 mtu 9216
 channel-group 54 mode active

ip address 10.29.164.66/24

line console

line vty

boot nxos bootflash://sup-1/n9000-dk9.7.0.3.I1.3b.bin
```

Fibre Channel Network Configuration

Cisco MDS 9148S-A Configuration

```
=~=~=~=~=~=~=~=~=~=~=~= PuTTY log 2017.11.30 10:03:13 =~=~=~=~=~=~=~=~=~=~=~=
login as: admin
User Access Verification
Using keyboard-interactive authentication.
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
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

MDS-A# show r[13D[J
MDS-A# show rradius       rdl           role          radius-cfs    redundancy    rscn          radius-server
rmon          running-config

MDS-A# show running-config
!Command: show running-config
!Time: Thu Nov 30 17:59:15 2017
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
snmp-server community public group network-operator
vsan database
 vsan 400 name "FlexPod-A"
device-alias database
device-alias name VDI-1-hba1 pwwn 20:00:00:25:b5:3a:00:3f
device-alias name VDI-2-hba1 pwwn 20:00:00:25:b5:3a:00:0f
device-alias name VDI-3-hba1 pwwn 20:00:00:25:b5:3a:00:1f
device-alias name VDI-4-hba1 pwwn 20:00:00:25:b5:3a:00:4e
device-alias name VDI-5-hba1 pwwn 20:00:00:25:b5:3a:00:2e
device-alias name VDI-6-hba1 pwwn 20:00:00:25:b5:3a:00:3e
device-alias name VDI-7-hba1 pwwn 20:00:00:25:b5:3a:00:0e
device-alias name VDI-9-hba1 pwwn 20:00:00:25:b5:3a:00:4d
device-alias name VDI-10-hba1 pwwn 20:00:00:25:b5:3a:00:2d
```

```
device-alias name VDI-11-hba1 pwwn 20:00:00:25:b5:3a:00:3d
device-alias name VDI-12-hba1 pwwn 20:00:00:25:b5:3a:00:0d
device-alias name VDI-13-hba1 pwwn 20:00:00:25:b5:3a:00:1d
device-alias name VDI-14-hba1 pwwn 20:00:00:25:b5:3a:00:4c
device-alias name VDI-15-hba1 pwwn 20:00:00:25:b5:3a:00:2c
device-alias name VDI-17-hba1 pwwn 20:00:00:25:b5:3a:00:0c
device-alias name VDI-18-hba1 pwwn 20:00:00:25:b5:3a:00:1c
device-alias name VDI-19-hba1 pwwn 20:00:00:25:b5:3a:00:4b
device-alias name VDI-20-hba1 pwwn 20:00:00:25:b5:3a:00:2b
device-alias name VDI-21-hba1 pwwn 20:00:00:25:b5:3a:00:3b
device-alias name VDI-22-hba1 pwwn 20:00:00:25:b5:3a:00:0b
device-alias name VDI-23-hba1 pwwn 20:00:00:25:b5:3a:00:1b
device-alias name VDI-24-hba1 pwwn 20:00:00:25:b5:3a:00:4a
device-alias name VDI-25-hba1 pwwn 20:00:00:25:b5:3a:00:2a
device-alias name Infra01-8-hba1 pwwn 20:00:00:25:b5:3a:00:4f
device-alias name Infra02-16-hba1 pwwn 20:00:00:25:b5:3a:00:2f
device-alias commit
fcdomain fcid database
vsan 400 wwn 50:0a:09:84:80:d3:67:d3 fcid 0x680000 dynamic
vsan 400 wwn 20:03:00:a0:98:af:bd:e8 fcid 0x680001 dynamic [a300-02-0g]
vsan 400 wwn 50:0a:09:84:80:13:41:27 fcid 0x680100 dynamic
vsan 400 wwn 20:01:00:a0:98:af:bd:e8 fcid 0x680101 dynamic [a300-01-0g]
vsan 400 wwn 20:02:00:de:fb:90:a0:80 fcid 0x680200 dynamic
vsan 400 wwn 20:03:00:de:fb:90:a0:80 fcid 0x680300 dynamic
vsan 400 wwn 20:04:00:de:fb:90:a0:80 fcid 0x680400 dynamic
vsan 400 wwn 20:01:00:de:fb:90:a0:80 fcid 0x680500 dynamic
vsan 400 wwn 20:00:00:25:b5:3a:00:3f fcid 0x680501 dynamic[VDI-1-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:0f fcid 0x680503 dynamic[VDI-2-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:1f fcid 0x680206 dynamic[VDI-3-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:4e fcid 0x680408 dynamic[VDI-4-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:2e fcid 0x680304 dynamic[VDI-5-hba1]
vsan 400 wwn 20:00:00:25:b5:3a:00:3e fcid 0x680407 dynamic[VDI-6-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:0e fcid 0x680306 dynamic[VDI-7-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:4f fcid 0x680406 dynamic[Infra01-8-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:4d fcid 0x680302 dynamic[VDI-9-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:2d fcid 0x680305 dynamic[VDI-10-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:3d fcid 0x680207 dynamic[VDI-11-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:0d fcid 0x68020a dynamic[VDI-12-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:1d fcid 0x680507 dynamic[VDI-13-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:4c fcid 0x680208 dynamic[VDI-14-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:2c fcid 0x680402 dynamic[VDI-15-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:2f fcid 0x680307 dynamic[Infra02-16-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:0c fcid 0x680403 dynamic[VDI-17-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:1c fcid 0x680205 dynamic[VDI-18-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:49 fcid 0x680504 dynamic [VDI-19-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:2b fcid 0x680202 dynamic[VDI-20-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:3b fcid 0x680303 dynamic[VDI-21-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:0b fcid 0x68020b dynamic[VDI-22-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:1b fcid 0x680508 dynamic[VDI-23-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:4a fcid 0x680505 dynamic[VDI-24-hba1]

vsan 400 wwn 20:00:00:25:b5:3a:00:2a fcid 0x680301 dynamic [VDI-25-hba1]

interface port-channel1

channel mode active

switchport rate-mode dedicated

interface port-channel2

channel mode active

switchport rate-mode dedicated

interface port-channel30

switchport rate-mode dedicated

vsan database

vsan 400 interface fc1/37

vsan 400 interface fc1/38
vsan 400 interface fc1/43

vsan 400 interface fc1/44

vsan 400 interface fc1/45

vsan 400 interface fc1/46

switchname MDS-A

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 fc1/14

switchport speed 8000

interface fc1/15

switchport speed 8000

interface fc1/16

switchport speed 8000

interface fc1/11

interface fc1/12

interface fc1/43

interface fc1/44

interface fc1/45

interface fc1/46

interface fc1/3

interface fc1/4

interface fc1/5

interface fc1/6

interface fc1/7

interface fc1/8

interface fc1/9

interface fc1/10

interface fc1/13

interface fc1/17
interface fc1/18

interface fc1/19

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 fc1/47

interface fc1/48

interface fc1/14

interface fc1/15

interface fc1/16

interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
Active Zone Database Section for vsan 400
zone name a300_VDI-1-hba1 vsan 400
member pwwn 20:00:00:25:b5:3a:00:3f [VDI-1-hba1]
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

zone name a300_VDI-2-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:0f[VDI-2-hba1]
zone name a300_VDI-3-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:1f[VDI-3-hba1]
zone name a300_VDI-4-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4e[VDI-4-hba1]
zone name a300_VDI-5-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2e[VDI-5-hba1]
zone name a300_VDI-6-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:3e[VDI-6-hba1]

zone name a300_VDI-7-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0e[VDI-7-hba1]

zone name a300_Infra01-8-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:4f[Infra01-8-hba1]

zone name a300_VDI-9-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:4d[VDI-9-hba1]

zone name a300_VDI-10-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:2d[VDI-10-hba1]

zone name a300_VDI-11-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:3d[VDI-11-hba1]

zone name a300_VDI-12-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0d[VDI-12-hba1]

zone name a300_VDI-13-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:1d[VDI-13-hba1]

zone name a300_VDI-14-hba1 vsan 400
```

```
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4c[VDI-14-hba1]
zone name a300_VDI-15-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2c[VDI-15-hba1]
zone name a300_Infra02-16-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2f[Infra02-16-hba1]
zone name a300_VDI-17-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:0c[VDI-17-hba1]
zone name a300_VDI-18-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:1c[VDI-18-hba1]
zone name a300_VDI-19-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4b[VDI-19-hba1]
zone name a300_VDI-20-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2b[VDI-20-hba1]
zone name a300_VDI-21-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:3b[VDI-21-hba1]
```

```
zone name a300_VDI-22-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:0b[VDI-22-hba1]
zone name a300_VDI-23-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:1b[VDI-23-hba1]
zone name a300_VDI-24-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4a[VDI-24-hba1]
zone name a300_VDI-25-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2a[VDI-25-hba1]
zoneset name FlexPod_FabricA vsan 400
member a300_VDI-1-hba1
member a300_VDI-2-hba1
member a300_VDI-3-hba1
member a300_VDI-4-hba1
member a300_VDI-5-hba1
member a300_VDI-6-hba1
member a300_VDI-7-hba1
member a300_Infra01-8-hba1
member a300_VDI-9-hba1
member a300_VDI-10-hba1
member a300_VDI-11-hba1
member a300_VDI-12-hba1
member a300_VDI-13-hba1
member a300_VDI-14-hba1
member a300_VDI-15-hba1

member a300_Infra02-16-hba1

member a300_VDI-17-hba1

member a300_VDI-18-hba1

member a300_VDI-19-hba1

member a300_VDI-20-hba1

member a300_VDI-21-hba1

member a300_VDI-22-hba1

member a300_VDI-23-hba1

member a300_VDI-24-hba1

member a300_VDI-25-hba1

zoneset activate name FlexPod_FabricA vsan 400

do clear zone database vsan 400

Full Zone Database Section for vsan 400

zone name a300_VDI-1-hba1 vsan 400

member pwwn 20:00:00:25:b5:3a:00:3f[VDI-1-hba1]

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

zone name a300_VDI-2-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

 member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0f[VDI-2-hba1]

zone name a300_VDI-3-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:1f[VDI-3-hba1]

zone name a300_VDI-4-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:4e[VDI-4-hba1]

zone name a300_VDI-5-hba1 vsan 400
```

```
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:2e[VDI-5-hba1]

zone name a300_VDI-6-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:3e[VDI-6-hba1]

zone name a300_VDI-7-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0e[VDI-7-hba1]

zone name a300_Infra01-8-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:1e[VDI-31-hba1]

zone name a300_VDI-9-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:4d[VDI-9-hba1]

zone name a300_VDI-10-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:2d[VDI-10-hba1]

zone name a300_VDI-11-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:3d[VDI-11-hba1]

zone name a300_VDI-12-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0d[VDI-12-hba1]
```

```
zone name a300_VDI-13-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:1d[VDI-13-hba1]
zone name a300_VDI-14-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4c[VDI-14-hba1]
zone name a300_VDI-15-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2c[VDI-15-hba1]
zone name a300_Infra02-16-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:2f[Infra02-16-hba1]
zone name a300_VDI-17-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:0c[VDI-17-hba1]
zone name a300_VDI-18-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:1c[VDI-18-hba1]
zone name a300_VDI-19-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]
member pwwn 20:00:00:25:b5:3a:00:4b[VDI-19-hba1]
zone name a300_VDI-20-hba1 vsan 400
member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]
member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:2b[VDI-20-hba1]

zone name a300_VDI-21-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:3b[VDI-21-hba1]

zone name a300_VDI-22-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:0b[VDI-22-hba1]

zone name a300_VDI-23-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:1b[VDI-23-hba1]

zone name a300_VDI-24-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:4a[VDI-24-hba1]

zone name a300_VDI-25-hba1 vsan 400

member pwwn 20:01:00:a0:98:af:bd:e8[a300-01-0g]

member pwwn 20:03:00:a0:98:af:bd:e8[a300-02-0g]

member pwwn 20:00:00:25:b5:3a:00:2a[VDI-25-hba1]

zoneset name FlexPod_FabricA vsan 400

member a300_VDI-1-hba1

member a300_VDI-2-hba1

member a300_VDI-3-hba1

member a300_VDI-4-hba1

member a300_VDI-5-hba1

member a300_VDI-6-hba1

member a300_VDI-7-hba1

member a300_Infra01-8-hba1
member a300_VDI-9-hba1

member a300_VDI-10-hba1

member a300_VDI-11-hba1

member a300_VDI-12-hba1

member a300_VDI-13-hba1

member a300_VDI-14-hba1

member a300_VDI-15-hba1

member a300_Infra02-16-hba1

member a300_VDI-17-hba1

member a300_VDI-18-hba1

member a300_VDI-19-hba1

member a300_VDI-20-hba1

member a300_VDI-21-hba1

member a300_VDI-22-hba1

member a300_VDI-23-hba1

member a300_VDI-24-hba1

member a300_VDI-25-hba1

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
channel-group 2 force
no shutdown
interface fc1/6
port-license acquire
channel-group 2 force
no shutdown
interface fc1/7
port-license acquire
channel-group 2 force
no shutdown
interface fc1/8
port-license acquire
channel-group 2 force
no shutdown
interface fc1/9
port-license acquire
interface fc1/10
port-license acquire
interface fc1/11
port-license acquire
interface fc1/12
port-license acquire
interface fc1/13
port-license acquire
no shutdown
interface fc1/14
port-license acquire
```

```
no shutdown
interface fc1/15
port-license acquire
no shutdown
interface fc1/16
port-license acquire
no shutdown
interface fc1/17
port-license acquire
channel-group 1 force
no shutdown
interface fc1/18
port-license acquire
channel-group 1 force
no shutdown
interface fc1/19
port-license acquire
channel-group 1 force
no shutdown
interface fc1/20
port-license acquire
channel-group 1 force
no shutdown
interface fc1/21
port-license acquire
no shutdown
interface fc1/22
port-license acquire
no shutdown
interface fc1/23
port-license acquire
```

```
no shutdown
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
no shutdown
interface fc1/26
port-license acquire
no shutdown
interface fc1/27
port-license acquire
no shutdown
interface fc1/28
port-license acquire
no shutdown
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

no shutdown

interface fc1/38

port-license acquire

no shutdown

interface fc1/39

port-license acquire

no shutdown

interface fc1/40

port-license acquire

no shutdown

interface fc1/41

port-license acquire

no shutdown

interface fc1/42

port-license acquire

no shutdown

interface fc1/43

port-license acquire

no shutdown

interface fc1/44

port-license acquire

no shutdown

interface fc1/45

port-license acquire

no shutdown

interface fc1/46

port-license acquire

no shutdown

interface fc1/47
```

```
port-license acquire

no shutdown

interface fc1/48

port-license acquire

no shutdown

interface mgmt0

ip address 10.29.164.64 255.255.255.0

ip default-gateway 10.29.164.1

MDS-A#
```

Cisco MDS 9148S-B Configuration

```
=~=~=~=~=~=~=~=~=~=~=~= PuTTY log 2017.11.30 10:06:05 =~=~=~=~=~=~=~=~=~=~=~=

login as: admin

User Access Verification

Using keyboard-interactive authentication.

Password:

Cisco Nexus Operating System (NX-OS) Software

TAC support: http://www.cisco.com/tac

Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.

The copyrights to certain works contained in this software are

owned by other third parties and used and distributed under

license. Certain components of this software are licensed under

the GNU General Public License (GPL) version 2.0 or the GNU

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

MDS-B# show running config

!Command: show running-config

!Time: Thu Nov 30 18:04:42 2017
```

```
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

snmp-server community public group network-operator

vsan database

vsan 401 name "FlexPod-B"

device-alias database

device-alias name VDI-1-hba2 pwwn 20:00:00:25:d5:06:00:3f

device-alias name VDI-2-hba2 pwwn 20:00:00:25:d5:06:00:0f
```

```
device-alias name VDI-3-hba2 pwwn 20:00:00:25:d5:06:00:1f
device-alias name VDI-4-hba2 pwwn 20:00:00:25:d5:06:00:4e
device-alias name VDI-5-hba2 pwwn 20:00:00:25:d5:06:00:2e
device-alias name VDI-6-hba2 pwwn 20:00:00:25:d5:06:00:3e
device-alias name VDI-7-hba2 pwwn 20:00:00:25:d5:06:00:0e
device-alias name VDI-9-hba2 pwwn 20:00:00:25:d5:06:00:4d
device-alias name a300-01-0h pwwn 20:02:00:a0:98:af:bd:e8
device-alias name a300-02-0h pwwn 20:04:00:a0:98:af:bd:e8
device-alias name VDI-10-hba2 pwwn 20:00:00:25:d5:06:00:2d
device-alias name VDI-11-hba2 pwwn 20:00:00:25:d5:06:00:3d
device-alias name VDI-12-hba2 pwwn 20:00:00:25:d5:06:00:0d
device-alias name VDI-13-hba2 pwwn 20:00:00:25:d5:06:00:1d
device-alias name VDI-14-hba2 pwwn 20:00:00:25:d5:06:00:4c
device-alias name VDI-15-hba2 pwwn 20:00:00:25:d5:06:00:2c
device-alias name VDI-17-hba2 pwwn 20:00:00:25:d5:06:00:0c
device-alias name VDI-18-hba2 pwwn 20:00:00:25:d5:06:00:1c
device-alias name VDI-19-hba2 pwwn 20:00:00:25:d5:06:00:4b
device-alias name VDI-20-hba2 pwwn 20:00:00:25:d5:06:00:2b
device-alias name VDI-21-hba2 pwwn 20:00:00:25:d5:06:00:3b
device-alias name VDI-22-hba2 pwwn 20:00:00:25:d5:06:00:6b
device-alias name VDI-23-hba2 pwwn 20:00:00:25:d5:06:00:1b
device-alias name VDI-24-hba2 pwwn 20:00:00:25:d5:06:00:4a
device-alias name VDI-25-hba2 pwwn 20:00:00:25:d5:06:00:2a
device-alias name Infra01-8-hba2 pwwn 20:00:00:25:d5:06:00:4f
device-alias name Infra02-16-hba2 pwwn 20:00:00:25:d5:06:00:2f
device-alias commit
vsan 401 wwn 20:04:00:a0:98:af:bd:e8 fcid 0x870001 dynamic[a300-02-0h]
vsan 401 wwn 50:0a:09:83:80:13:41:27 fcid 0x870100 dynamic
vsan 401 wwn 20:02:00:a0:98:af:bd:e8 fcid 0x870101 dynamic[a300-01-0h]
vsan 401 wwn 20:01:00:de:fb:92:0c:80 fcid 0x870200 dynamic
vsan 401 wwn 20:02:00:de:fb:92:0c:80 fcid 0x870300 dynamic
vsan 401 wwn 20:03:00:de:fb:92:0c:80 fcid 0x870400 dynamic

vsan 401 wwn 20:04:00:de:fb:92:0c:80 fcid 0x870500 dynamic

vsan 401 wwn 20:00:00:25:d5:06:00:3f fcid 0x870205 dynamic[VDI-1-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:0f fcid 0x870406 dynamic[VDI-2-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:1f fcid 0x870506 dynamic[VDI-3-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4e fcid 0x870304 dynamic[VDI-4-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2e fcid 0x870305 dynamic[VDI-5-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:3e fcid 0x870405 dynamic[VDI-6-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:0e fcid 0x870503 dynamic[VDI-7-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4f fcid 0x870502 dynamic[Infra01-8-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4d fcid 0x870301 dynamic[VDI-9-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2d fcid 0x870302 dynamic[VDI-10-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:3d fcid 0x870209 dynamic[VDI-11-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:0d fcid 0x870201 dynamic[VDI-12-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:1d fcid 0x870308 dynamic[VDI-13-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4c fcid 0x870507 dynamic[VDI-14-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2c fcid 0x870407 dynamic[VDI-15-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2f fcid 0x870203 dynamic[Infra02-16-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:0c fcid 0x870204 dynamic[VDI-17-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:1c fcid 0x870206 dynamic[VDI-18-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4b fcid 0x870207 dynamic [VDI-19-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2b fcid 0x870501 dynamic[VDI-20-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:3b fcid 0x870505 dynamic[VDI-21-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:3c fcid 0x870520 dynamic[VDI-22-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:1b fcid 0x870307 dynamic[VDI-23-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:4a fcid 0x870306 dynamic[VDI-24-hba2]

vsan 401 wwn 20:00:00:25:d5:06:00:2a fcid 0x870504 dynamic[VDI-25-hba2]

interface port-channel1

channel mode active

switchport rate-mode dedicated

interface port-channel2
```

```
channel mode active

switchport rate-mode dedicated

interface port-channel11

channel mode active

switchport rate-mode dedicated

vsan database

vsan 401 interface port-channel11

vsan 401 interface fc1/37

vsan 401 interface fc1/38

vsan 401 interface fc1/43

vsan 401 interface fc1/44

vsan 401 interface fc1/45

vsan 401 interface fc1/46

switchname MDS-B

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 fc1/13

switchport speed 8000

interface fc1/14

switchport speed 8000

interface fc1/15

switchport speed 8000

interface fc1/16

switchport speed 8000

interface fc1/3

interface fc1/4

interface fc1/43

interface fc1/44

interface fc1/45
interface fc1/46

interface fc1/3

interface fc1/4

interface fc1/5

interface fc1/6

interface fc1/7

interface fc1/8

interface fc1/9

interface fc1/10

interface fc1/17

interface fc1/18

interface fc1/19

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 fc1/47

interface fc1/48

interface fc1/13

interface fc1/14

interface fc1/15

interface fc1/16

interface fc1/1

interface fc1/2

interface fc1/11

interface fc1/12

interface fc1/43

interface fc1/44

interface fc1/45

interface fc1/46

Active Zone Database Section for vsan 401

zone name a300_VDI-1-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:3f[VDI-1-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-2-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:0f[VDI-2-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-3-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1f[VDI-3-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
```

```
zone name a300_VDI-4-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4e
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-5-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2e[VDI-5-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-6-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:3e[VDI-6-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-7-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:0e[VDI-7-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_Infra01-8-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4f[Infra01-8-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-9-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4d[VDI-9-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-10-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2d[VDI-10-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-11-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:3d[VDI-11-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-12-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:0d[VDI-12-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-13-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1d[VDI-13-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-14-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:4c[VDI-14-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-15-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2c[VDI-15-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_Infra02-16-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2f[Infra02-16-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-17-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:0c[VDI-17-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-18-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1c[VDI-18-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-19-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:4b[VDI-19-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-20-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2b[VDI-20-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-21-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:3b[VDI-21-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-22-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:6b[VDI-22-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-23-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1b[VDI-23-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-24-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:4a[VDI-24-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-25-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2a[VDI-25-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zoneset name FlexPod_FabricB vsan 401

member a300_VDI-1-hba2

member a300_VDI-2-hba2

member a300_VDI-3-hba2

member a300_VDI-4-hba2
member a300_VDI-5-hba2

member a300_VDI-6-hba2

member a300_VDI-7-hba2

member a300_Infra01-8-hba2

member a300_VDI-9-hba2

member a300_VDI-10-hba2

member a300_VDI-11-hba2

member a300_VDI-12-hba2

member a300_VDI-13-hba2

member a300_VDI-14-hba2

member a300_VDI-15-hba2

member a300_Infra02-16-hba2

member a300_VDI-17-hba2

member a300_VDI-18-hba2

member a300_VDI-19-hba2

member a300_VDI-20-hba2

member a300_VDI-21-hba2

member a300_VDI-22-hba2

member a300_VDI-23-hba2

member a300_VDI-24-hba2

member a300_VDI-25-hba2

zoneset activate name FlexPod_FabricB vsan 401

do clear zone database vsan 401

Full Zone Database Section for vsan 401

zone name a300_VDI-1-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:3f [VDI-1-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-2-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:0f[VDI-2-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-3-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1f[VDI-3-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-4-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:4e[VDI-4-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-5-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2e[VDI-5-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-6-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:3e[VDI-6-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-7-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:0e[VDI-7-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_Infra01-8-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:1e[VDI-31-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-9-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:4d[VDI-9-hba2]

member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]

member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]

zone name a300_VDI-10-hba2 vsan 401

member pwwn 20:00:00:25:d5:06:00:2d[VDI-10-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-11-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:3d[VDI-11-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-12-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:0d[VDI-12-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-13-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:1d[VDI-13-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-14-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4c[VDI-14-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-15-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2c[VDI-15-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_Infra02-16-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2f[Infra02-16-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-17-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:0c[VDI-17-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-18-hba2 vsan 401
```

```
member pwwn 20:00:00:25:d5:06:00:1c[VDI-18-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-19-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4b[VDI-19-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-20-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2b[VDI-20-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-21-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:3b[VDI-21-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-22-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:6b[VDI-22-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-23-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:1b[VDI-23-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-24-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:4a[VDI-24-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
zone name a300_VDI-25-hba2 vsan 401
member pwwn 20:00:00:25:d5:06:00:2a[VDI-25-hba2]
member pwwn 20:02:00:a0:98:af:bd:e8[a300-01-0h]
member pwwn 20:04:00:a0:98:af:bd:e8[a300-02-0h]
```

```
zoneset name FlexPod_FabricB vsan 401

member a300_VDI-1-hba2

member a300_VDI-2-hba2

member a300_VDI-3-hba2

member a300_VDI-4-hba2

member a300_VDI-5-hba2

member a300_VDI-6-hba2

member a300_VDI-7-hba2

member a300_Infra01-8-hba2

member a300_VDI-9-hba2

member a300_VDI-10-hba2

member a300_VDI-11-hba2

member a300_VDI-12-hba2

member a300_VDI-13-hba2

member a300_VDI-14-hba2

member a300_VDI-15-hba2

member a300_Infra02-16-hba2

member a300_VDI-17-hba2

member a300_VDI-18-hba2

member a300_VDI-19-hba2

member a300_VDI-20-hba2

member a300_VDI-21-hba2

member a300_VDI-22-hba2

member a300_VDI-23-hba2

member a300_VDI-24-hba2

member a300_VDI-25-hba2

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
channel-group 2 force
no shutdown
interface fc1/6
port-license acquire
channel-group 2 force
no shutdown
interface fc1/7
port-license acquire
channel-group 2 force
no shutdown
interface fc1/8
port-license acquire
channel-group 2 force
no shutdown
interface fc1/9
port-license acquire
no shutdown
interface fc1/10
```

```
port-license acquire

no shutdown

interface fc1/11

port-license acquire

interface fc1/12

port-license acquire

interface fc1/13

port-license acquire

no shutdown

interface fc1/14

port-license acquire

no shutdown

interface fc1/15

port-license acquire

no shutdown

interface fc1/16

port-license acquire

no shutdown

interface fc1/17

port-license acquire

channel-group 1 force

no shutdown

interface fc1/18

port-license acquire

channel-group 1 force

no shutdown

interface fc1/19

port-license acquire

channel-group 1 force

no shutdown

interface fc1/20
```

```
port-license acquire
channel-group 1 force
no shutdown
interface fc1/21
port-license acquire
no shutdown
interface fc1/22
port-license acquire
no shutdown
interface fc1/23
port-license acquire
no shutdown
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
no shutdown
interface fc1/26
port-license acquire
no shutdown
interface fc1/27
port-license acquire
no shutdown
interface fc1/28
port-license acquire
no shutdown
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
no shutdown
interface fc1/38
port-license acquire
no shutdown
interface fc1/39
port-license acquire
no shutdown
interface fc1/40
port-license acquire
no shutdown
interface fc1/41
port-license acquire
no shutdown
interface fc1/42
port-license acquire
no shutdown
```

```
interface fc1/43
port-license acquire
no shutdown
interface fc1/44
port-license acquire
no shutdown
interface fc1/45
port-license acquire
no shutdown
interface fc1/46
port-license acquire
no shutdown
interface fc1/47
port-license acquire
no shutdown
interface fc1/48
port-license acquire
no shutdown
interface mgmt0
ip address 10.29.164.128 255.255.255.0
ip default-gateway 10.29.164.1
MDS-B#
```

Appendix B - 5000 Users Scale Test Host Metrics for All Workload Servers and FlexPod AFF A300 Storage Graphs

This section highlights and provides analysis of the FlexPod AFF A300storage performance results for each of the cluster test cases identified in this 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 you will see, FlexPod AFF A300storage delivers that level of performance despite driving a substantial amount of IOPs and bandwidth for the thousands of desktops hosted on the single FlexPod AFFA300Storage.

The following charts were compiled from extracting front-end array telemetry data from the storage logs and are equivalent to values shown in the FlexPod AFF A300storage. 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 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 900 seconds and users begin logging out of the environment. For brevity, we 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 AFF A300 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.

5000 Mixed Workload VMware Horizon RDSH and VDI Non–Persistent Instant Clones and Persistent Full Clones Virtual Machines Testing

Mixed Workload All 3 Clusters as Scale Test

ESXTOP Util% Charts for All RDSH Hosts

RDSH Host-01

Figure 117 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

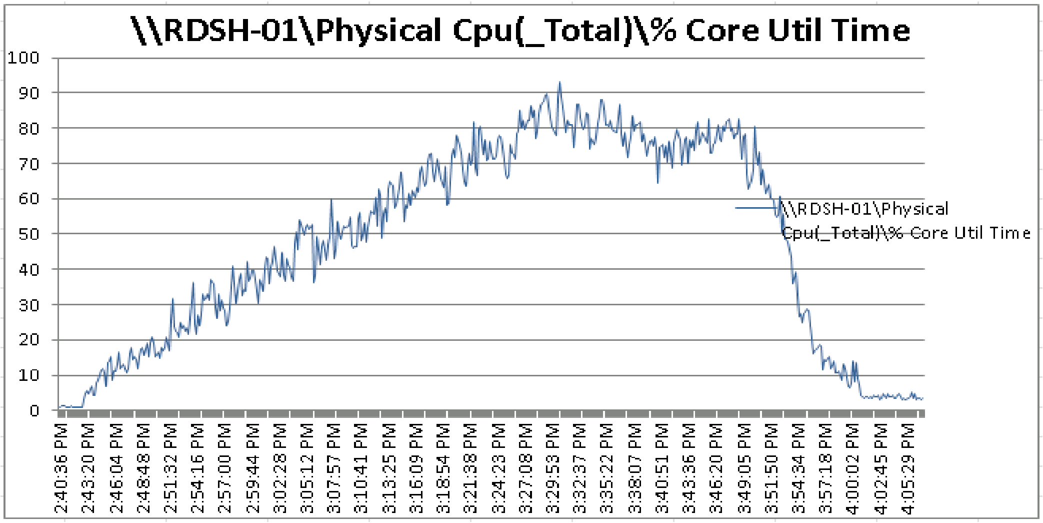

Figure 118 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

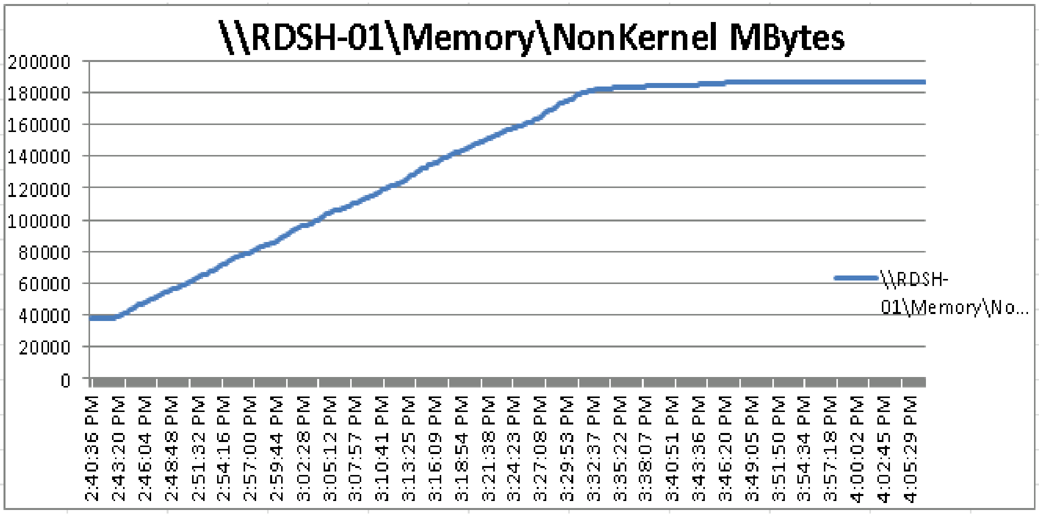

Figure 119 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

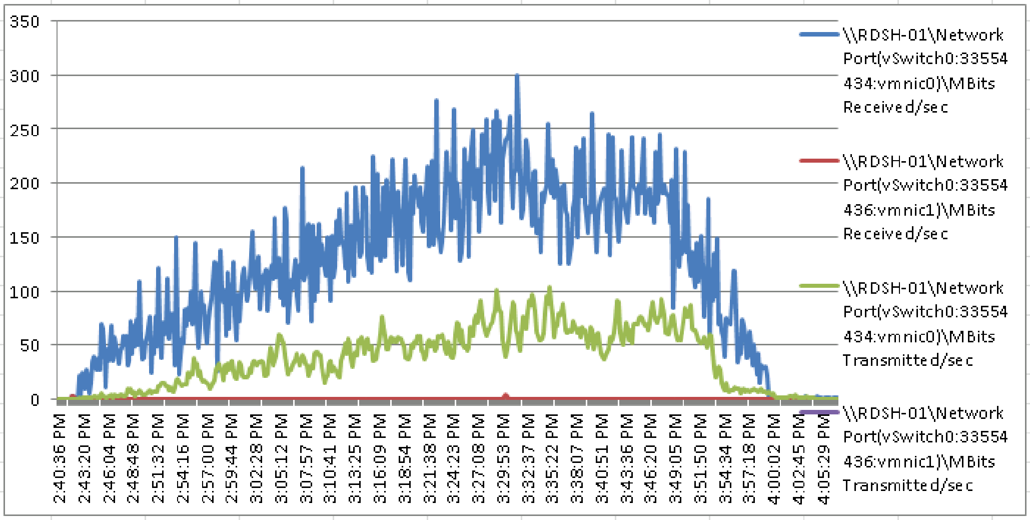

Figure 120 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

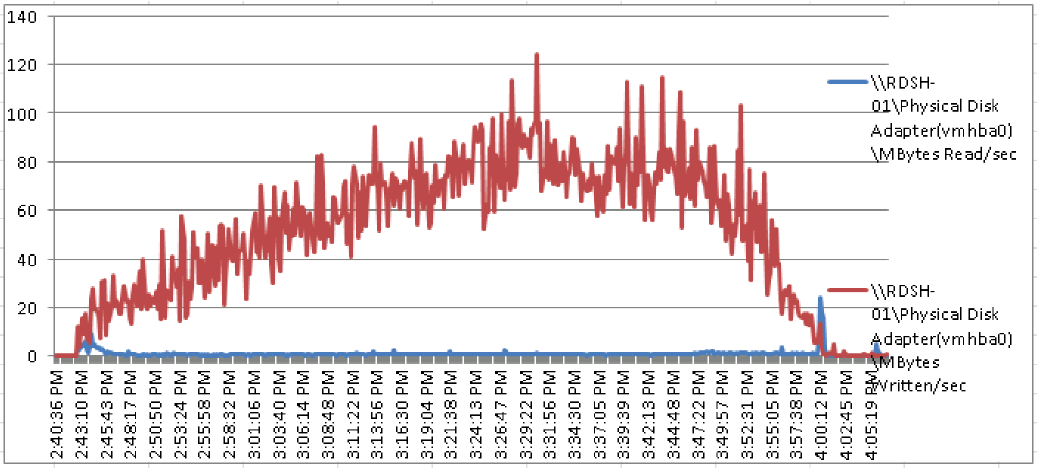

RDSH-Host 02

Figure 121 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

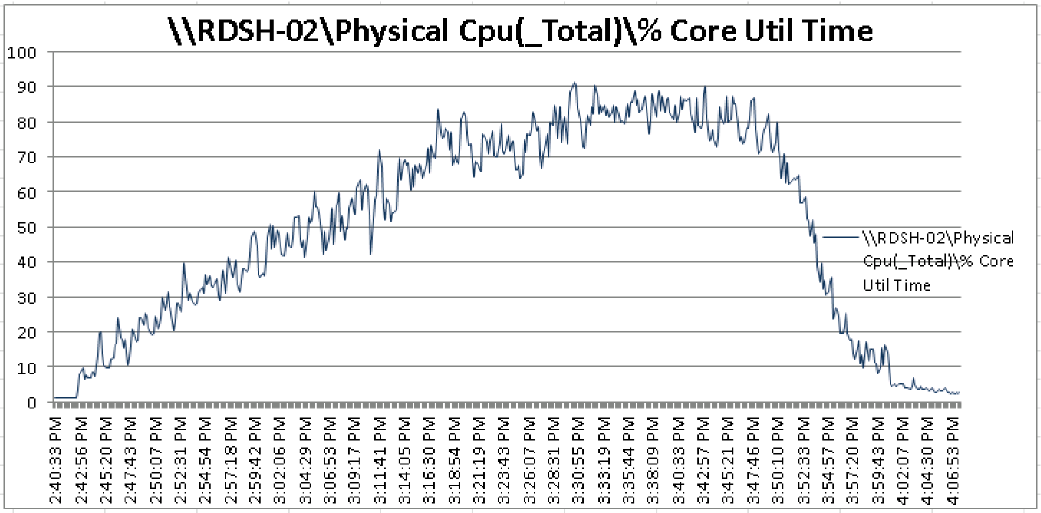

Figure 122 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

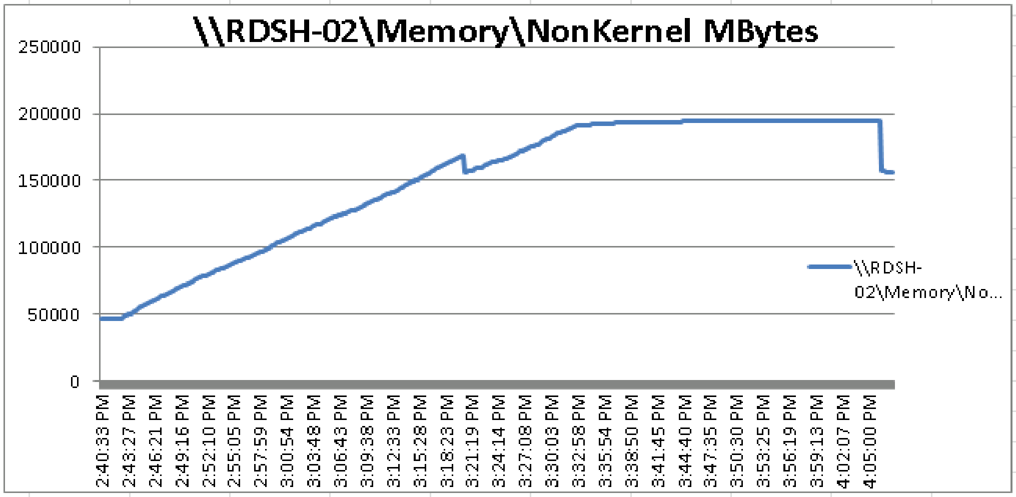

Figure 123 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

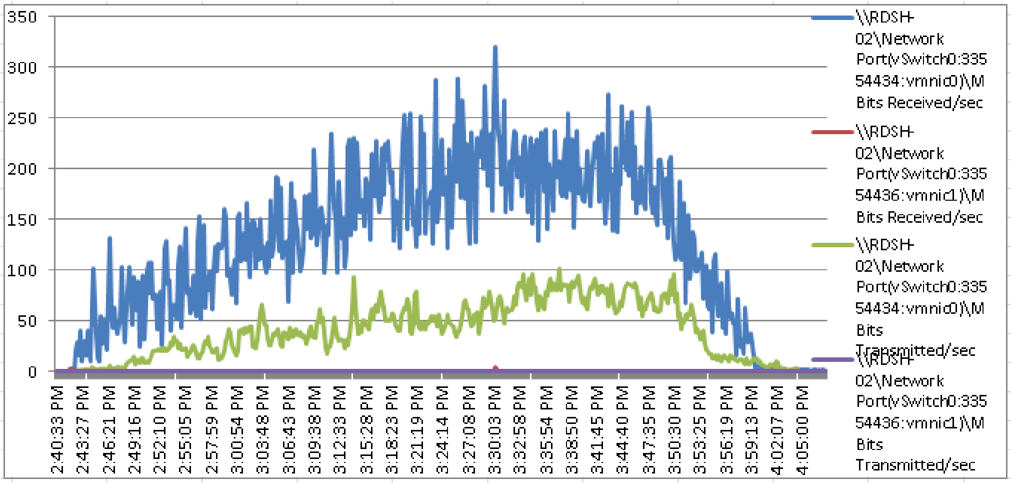

Figure 124 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

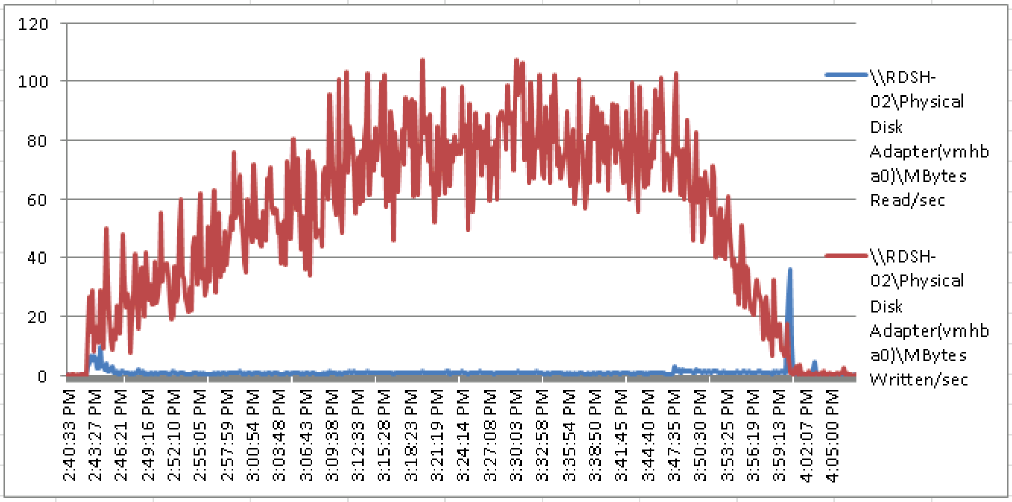

RDSH-03

Figure 125 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

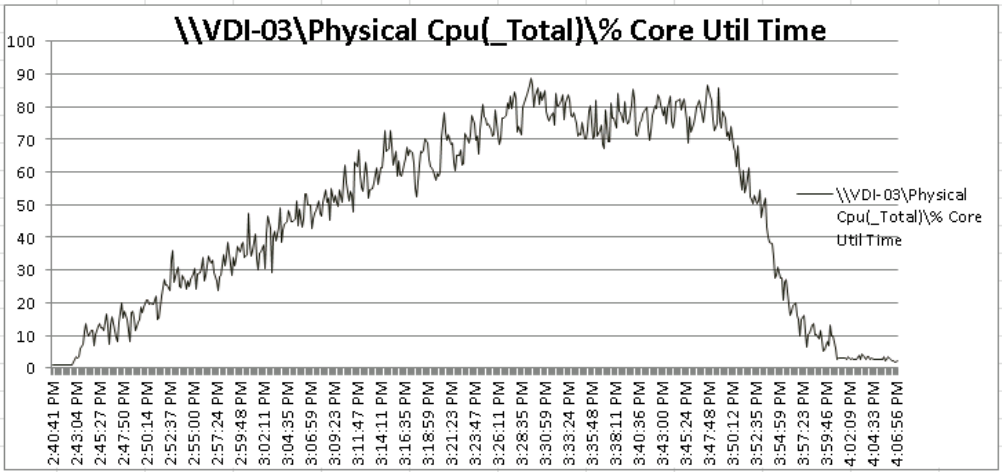

Figure 126 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

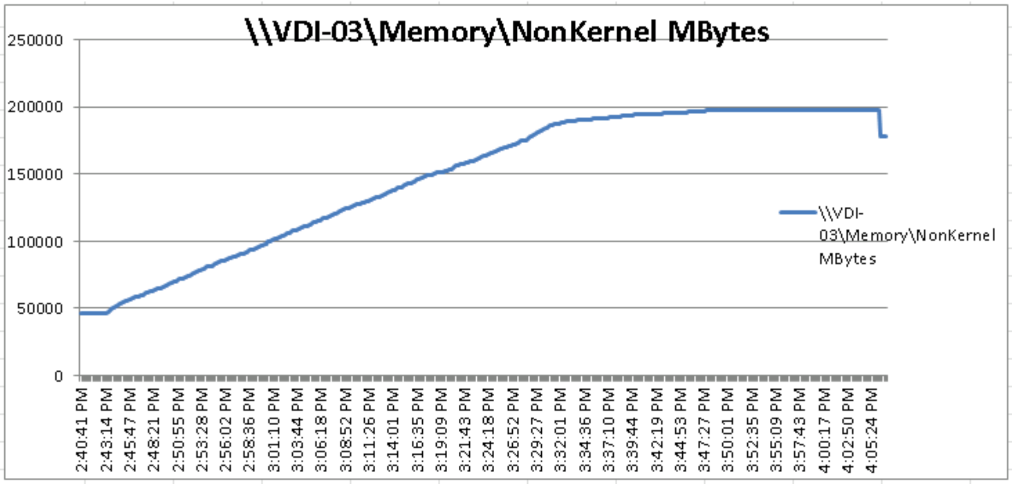

Figure 127 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

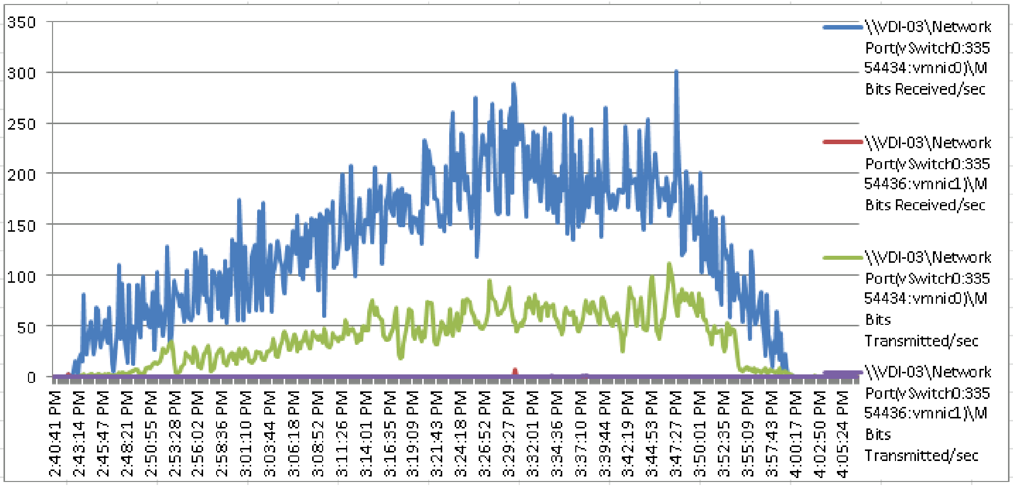

Figure 128 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

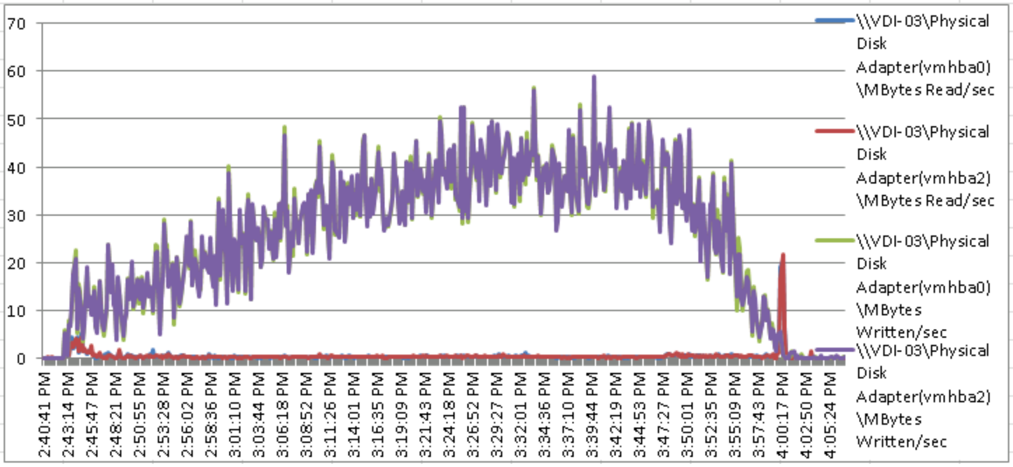

RDSH-04

Figure 129 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

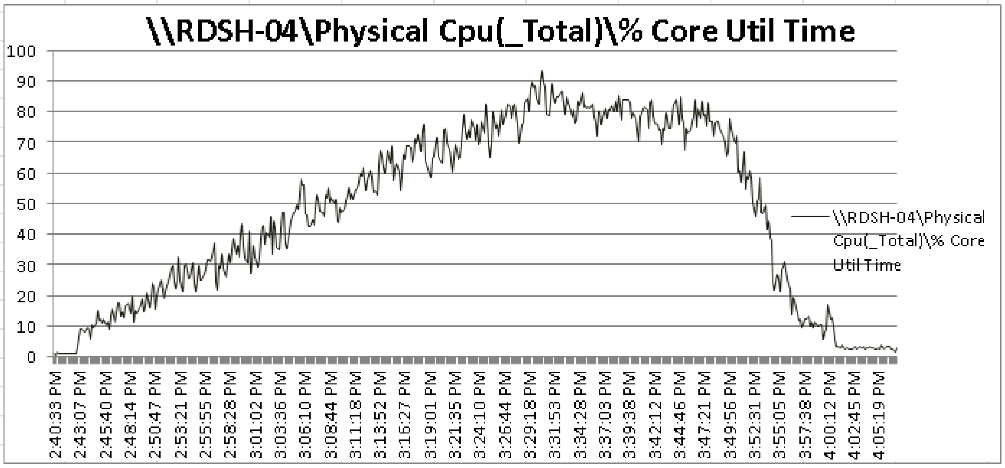

Figure 130 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

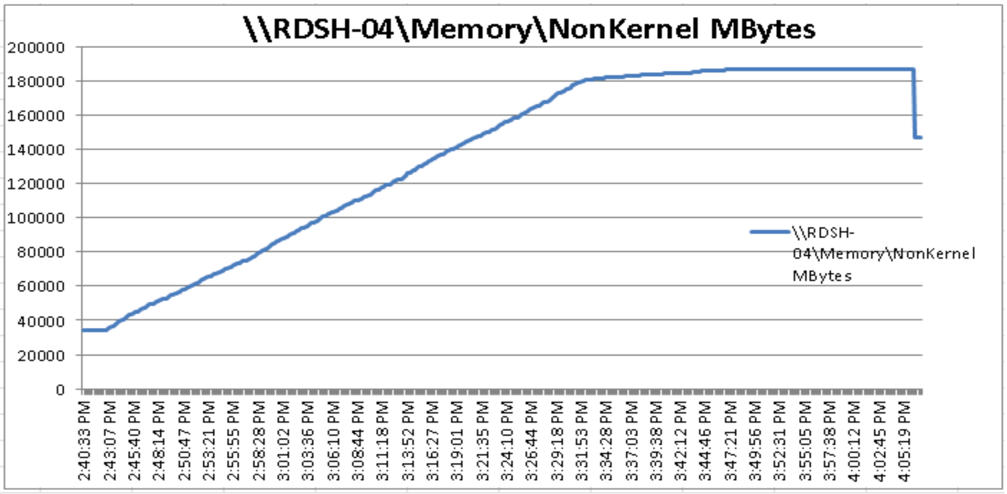

Figure 131 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

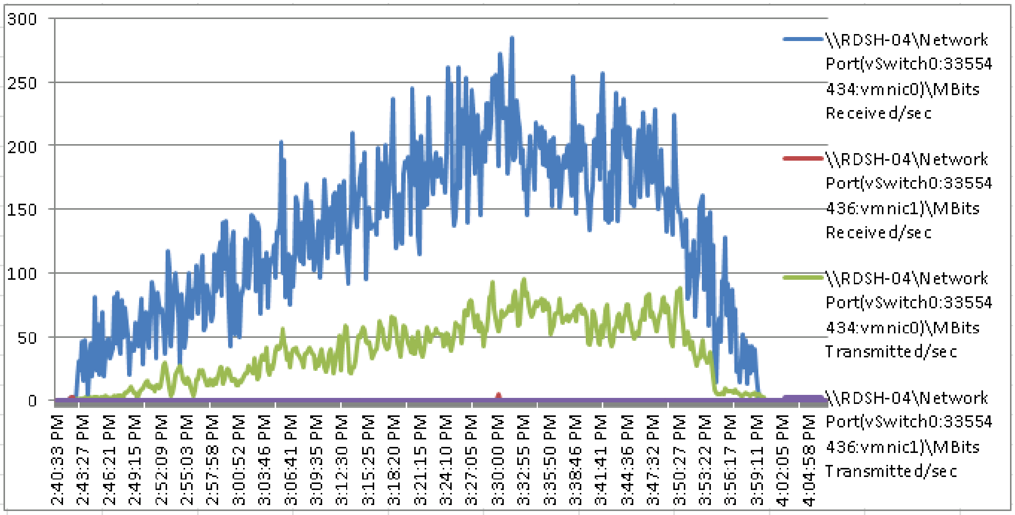

Figure 132 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

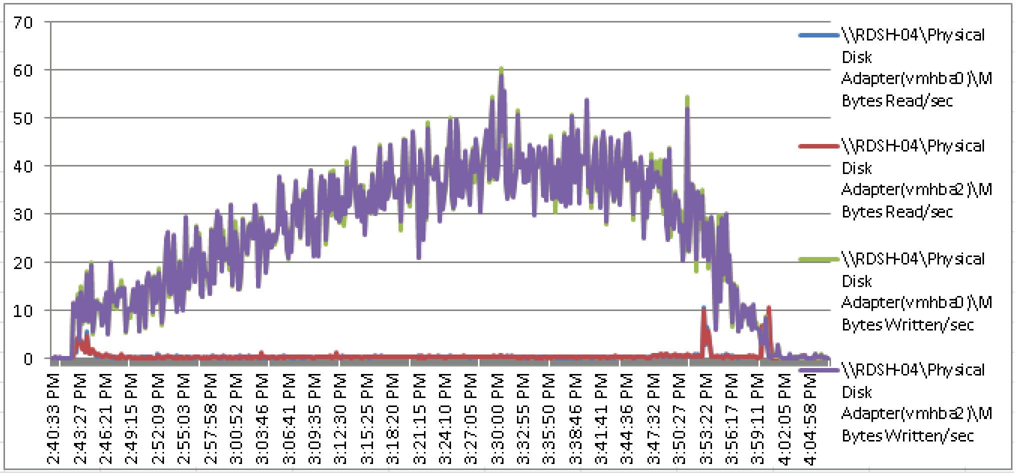

RDSH-05

Figure 133 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

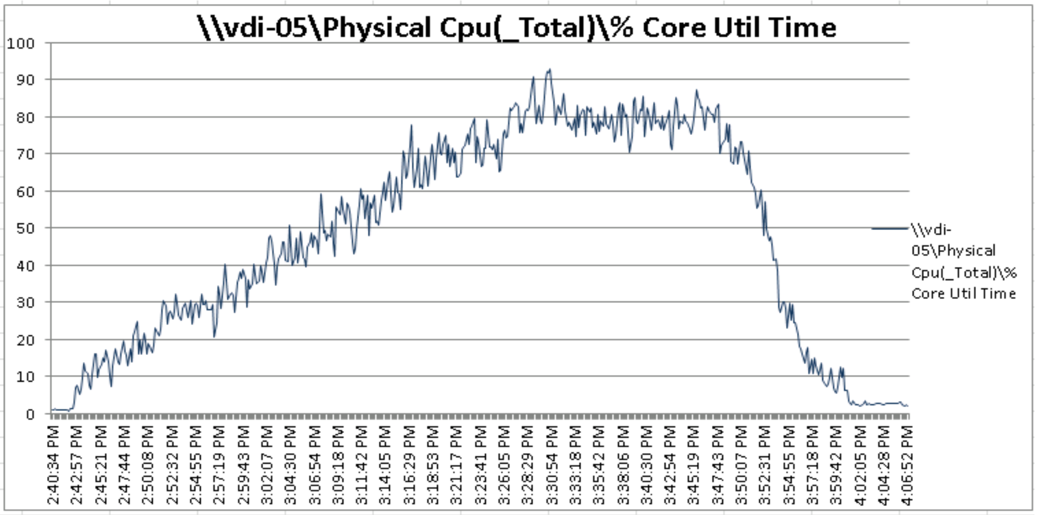

Figure 134 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

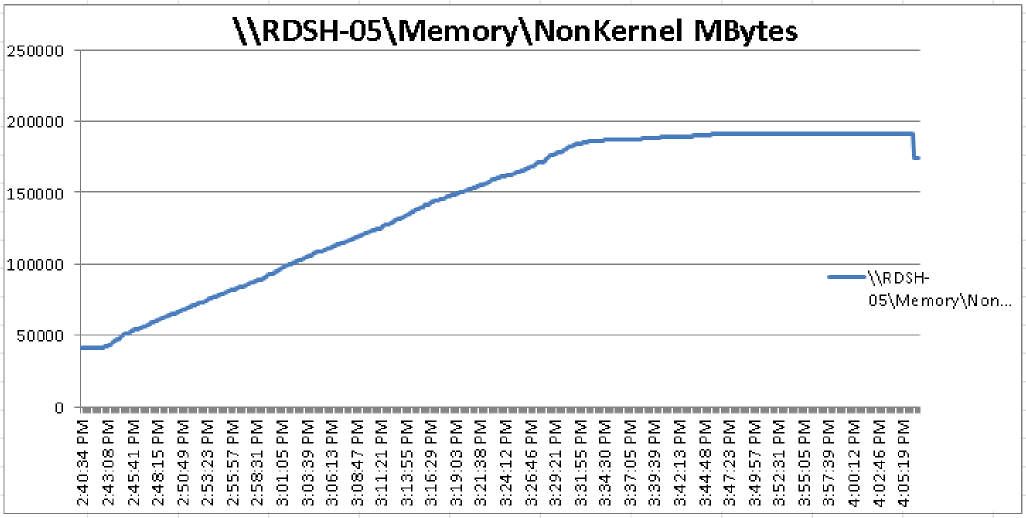

Figure 135 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

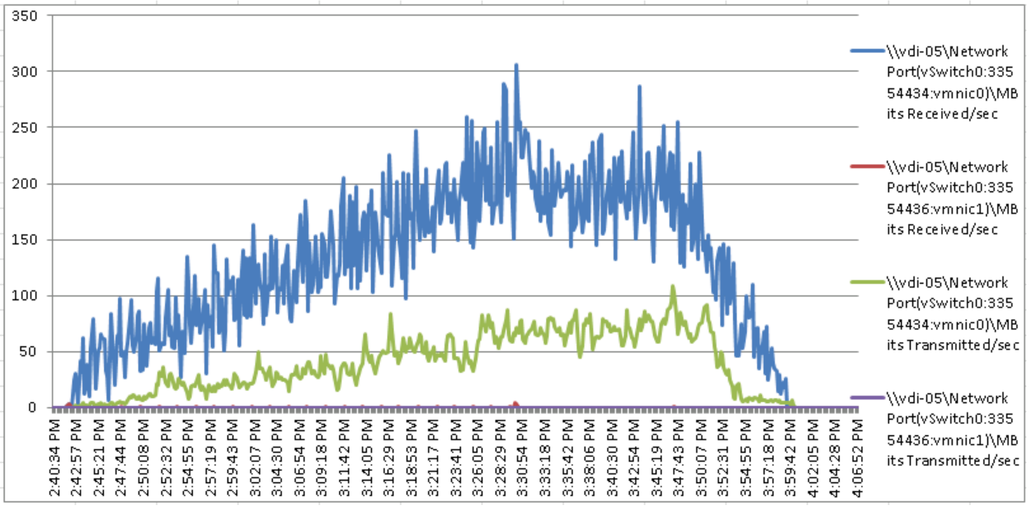

Figure 136 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

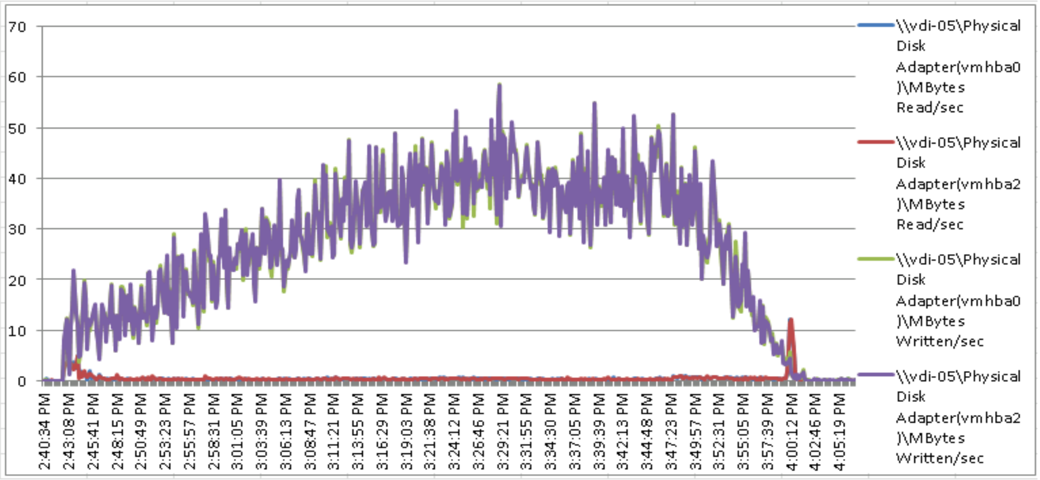

RDSH-06

Figure 137 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

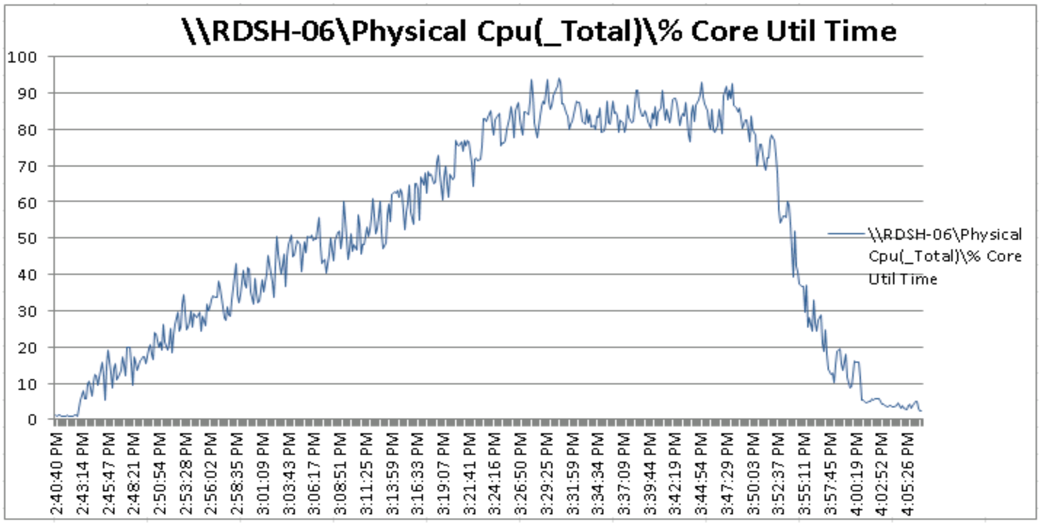

Figure 138 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

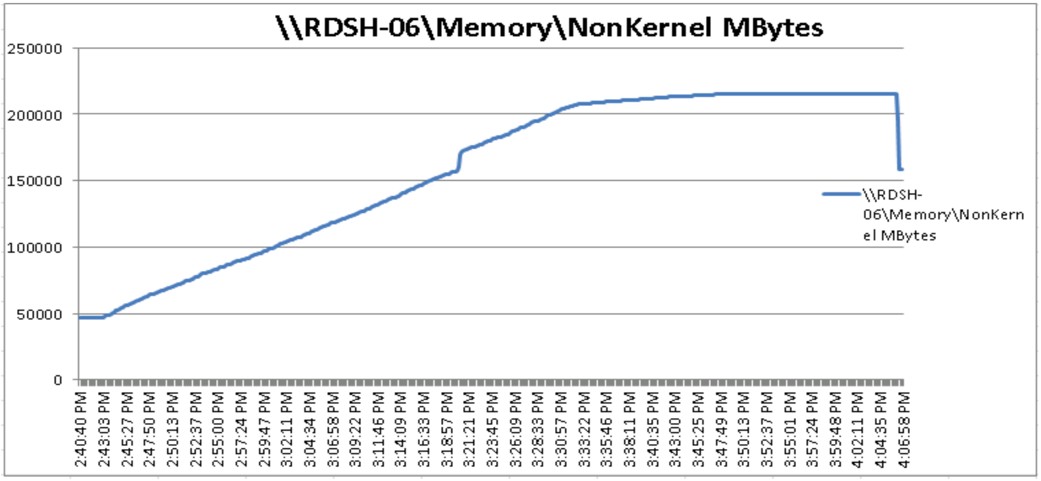

Figure 139 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

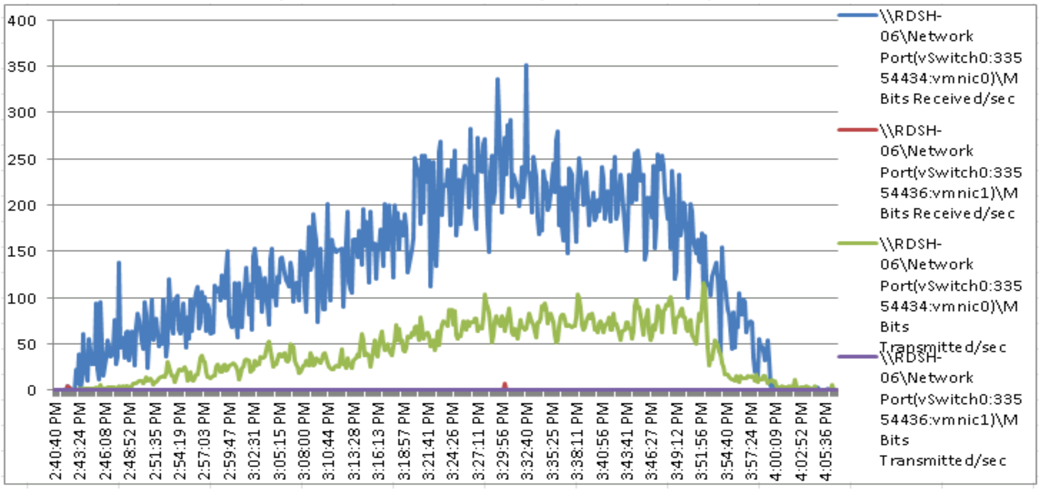

Figure 140 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

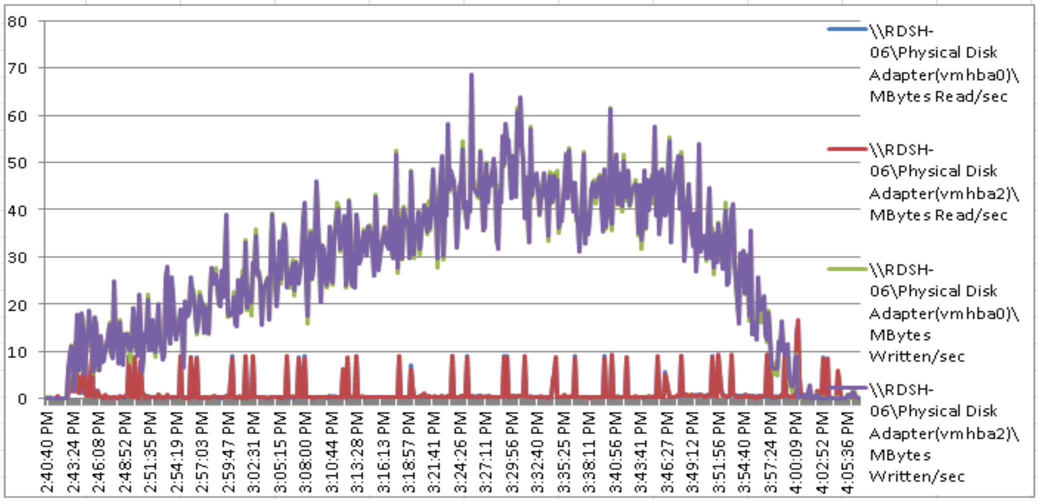

RDSH-07

Figure 141 Full Scale | 5000 Mixed Users | RDSH Host | Host CPU Utilization

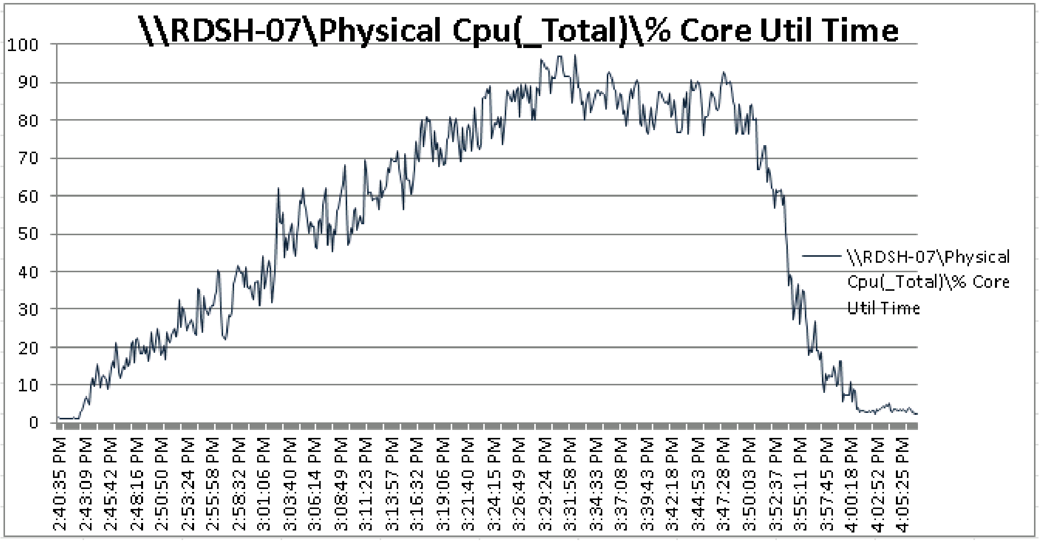

Figure 142 Full Scale | 5000 Mixed Users | RDSH Host | Host Memory Utilization

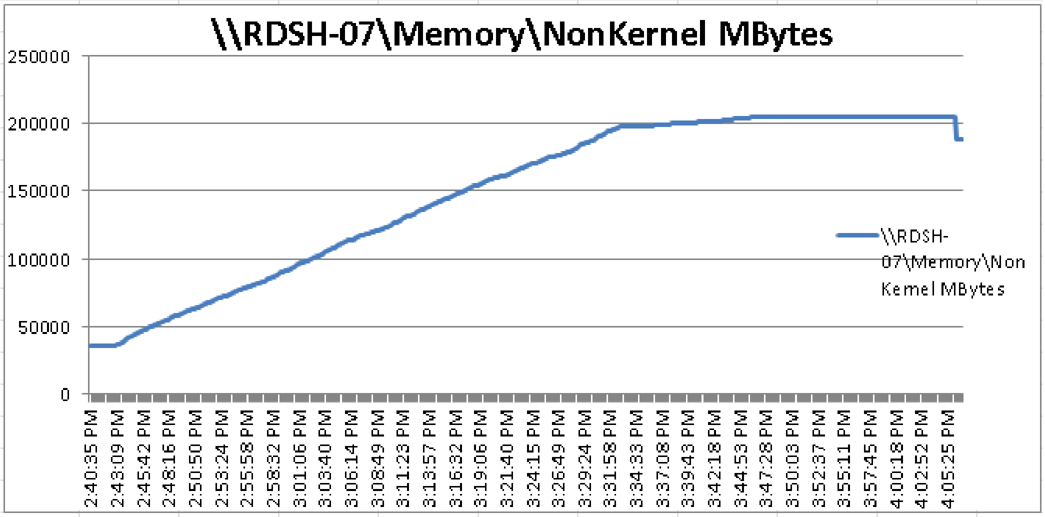

Figure 143 Full Scale | 5000 Mixed Users | RDSH Host | Host Network Utilization

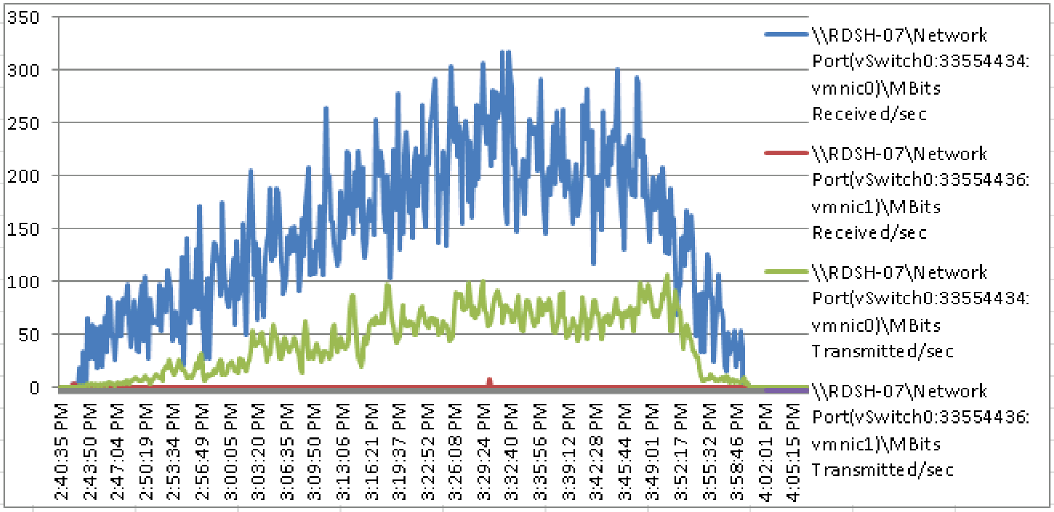

Figure 144 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

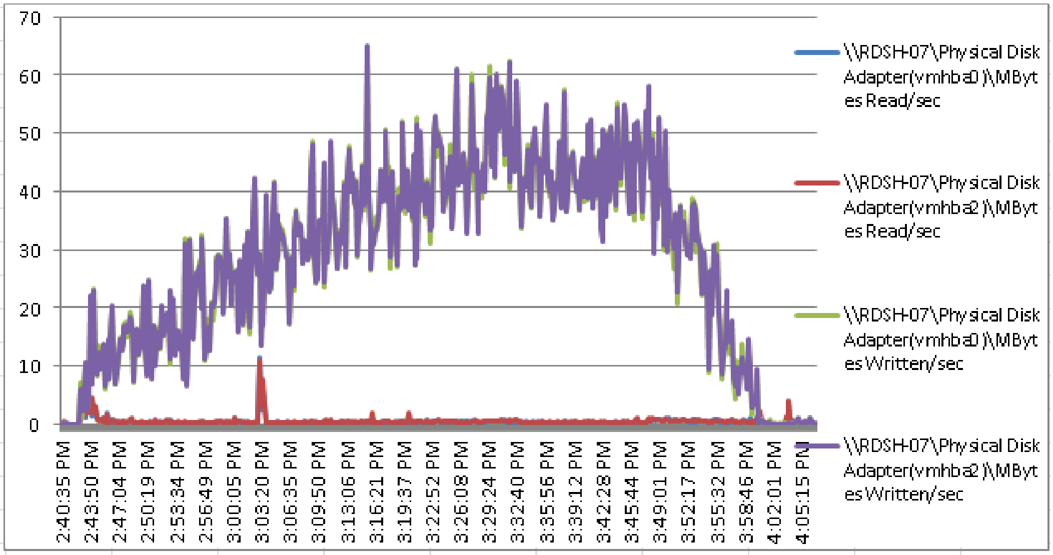

VDI Host Metrics for 5000 Users Scale Test

ESXTOP Util% Charts for All VDI (Non–Persistent and Persistent Workload Hosts)

VDI-01

Figure 145 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

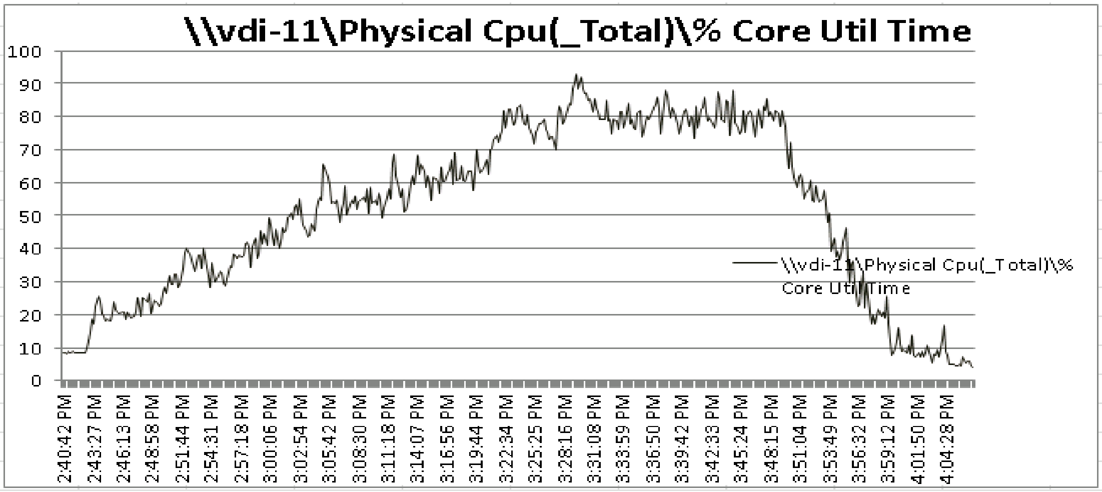

Figure 146 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

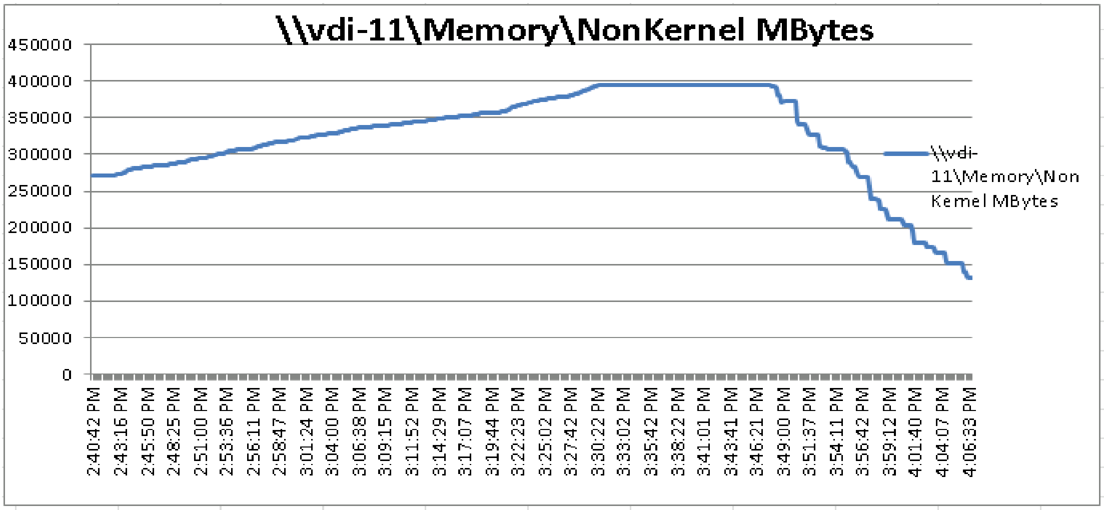

Figure 147 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

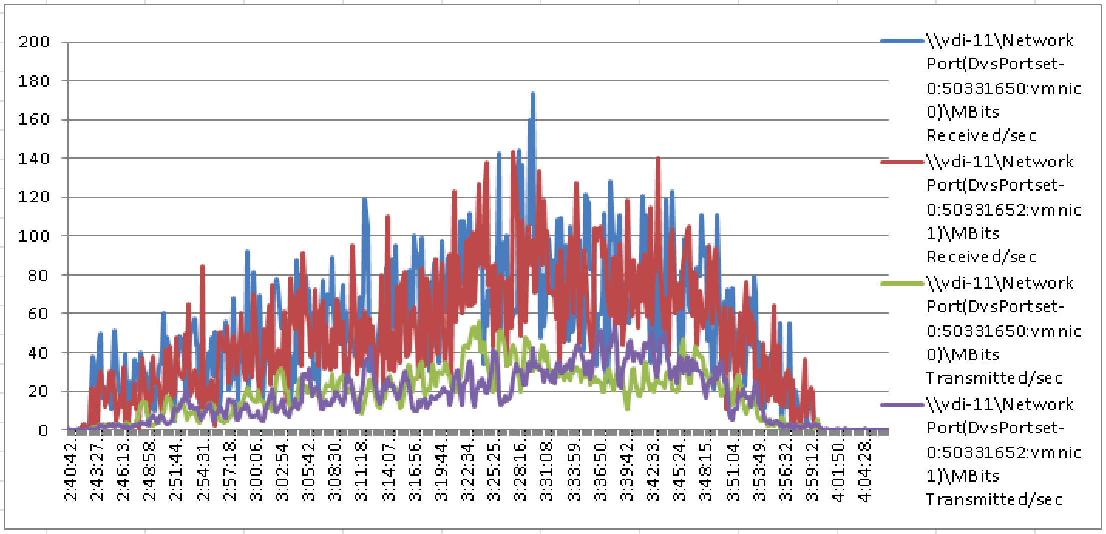

Figure 148 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

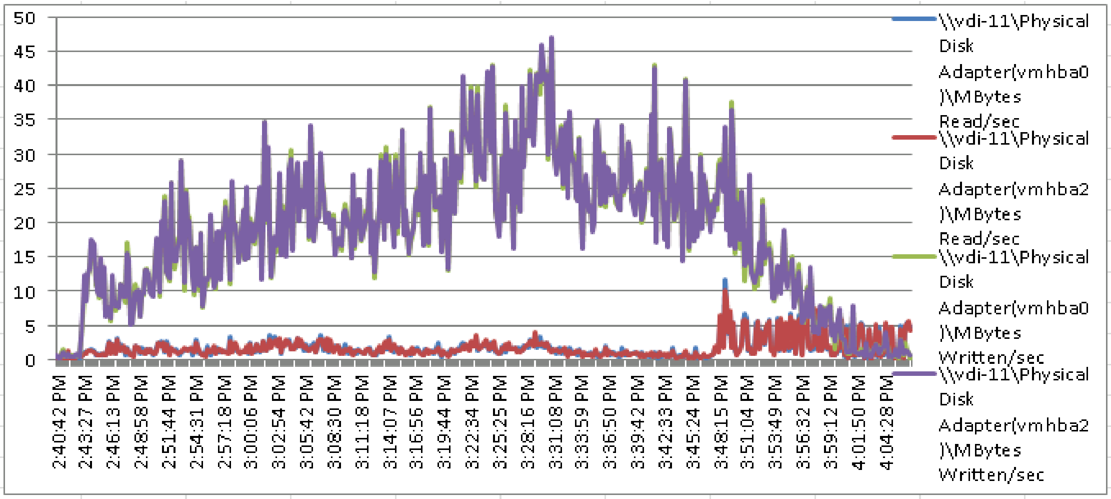

VDI Host-02

Figure 149 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

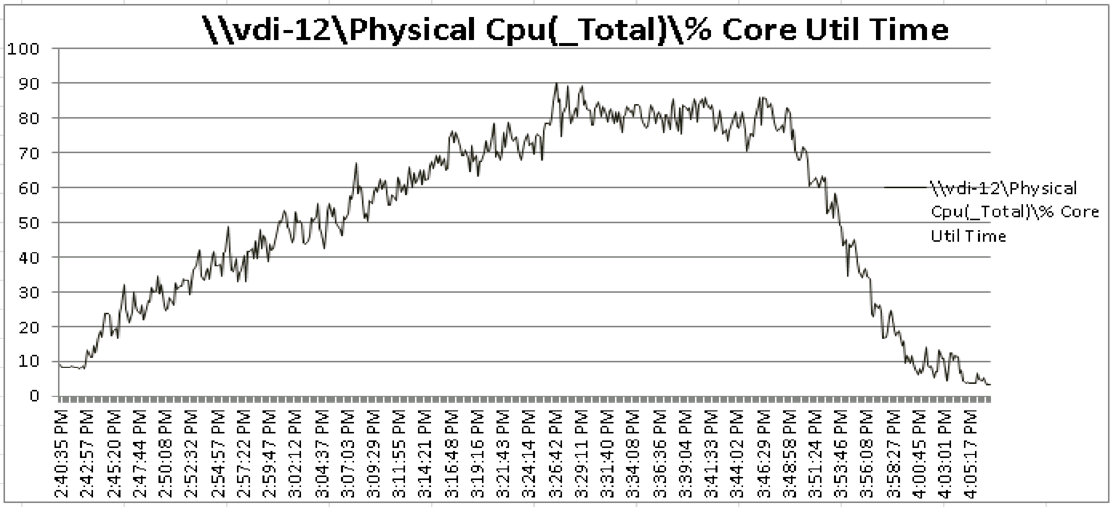

Figure 150 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

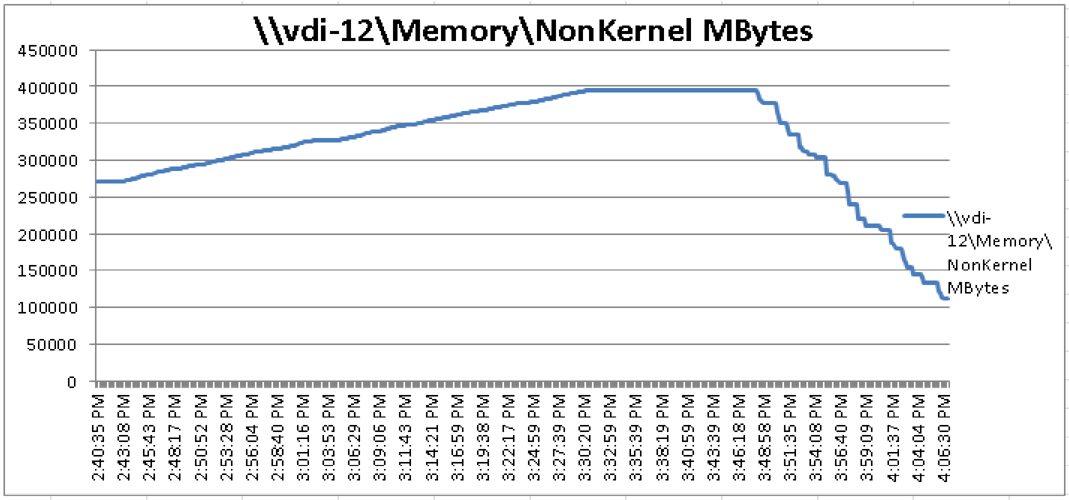

Figure 151 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

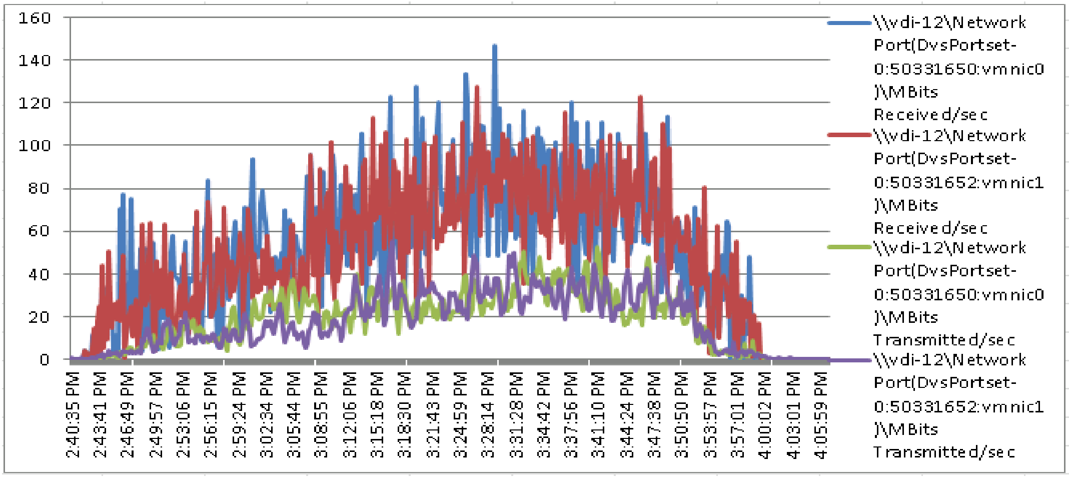

Figure 152 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

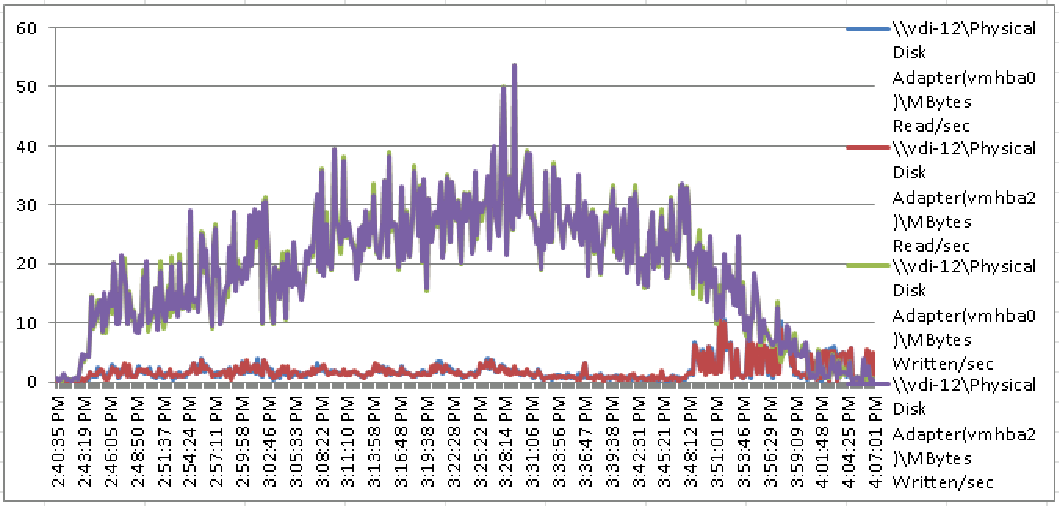

VDI Host-03

Figure 153 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

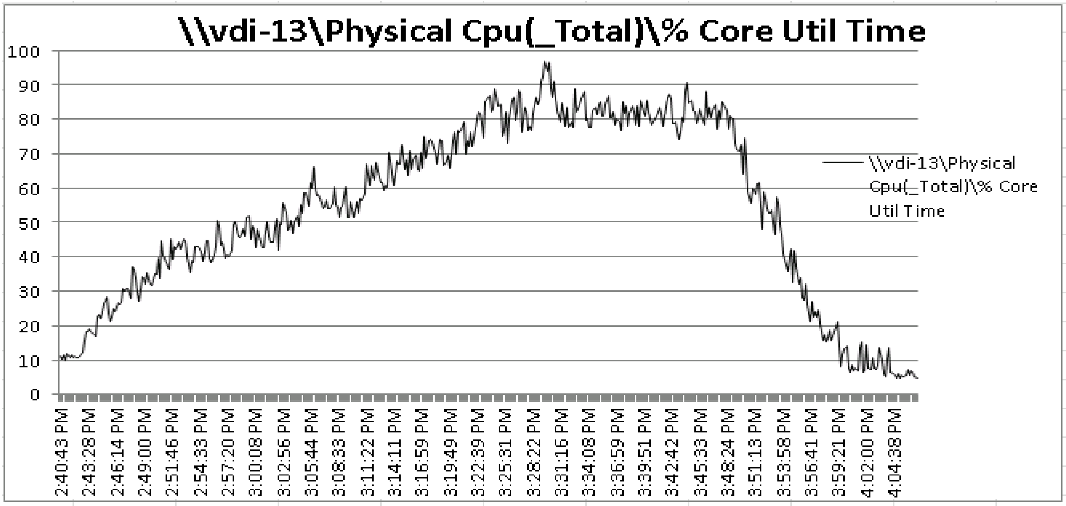

Figure 154 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

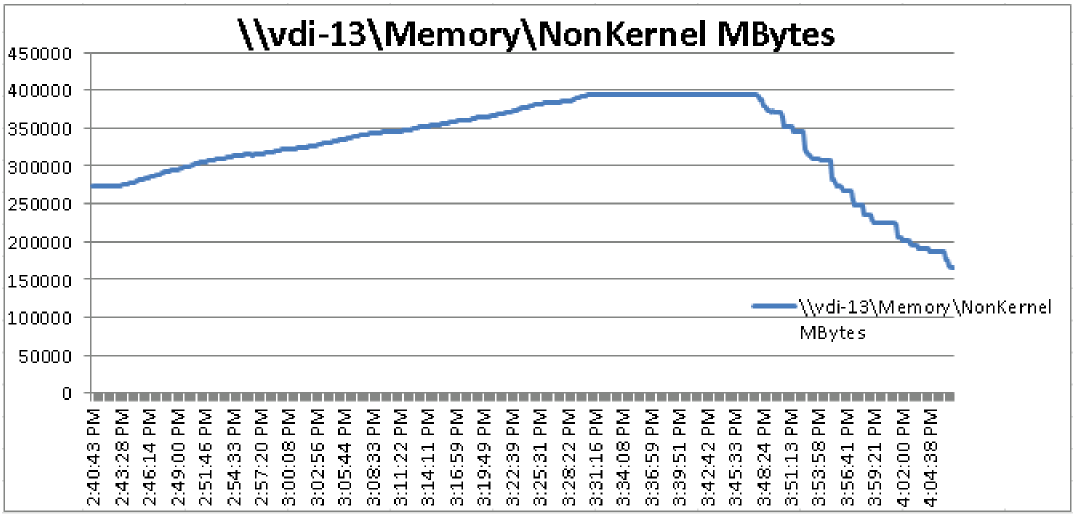

Figure 155 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

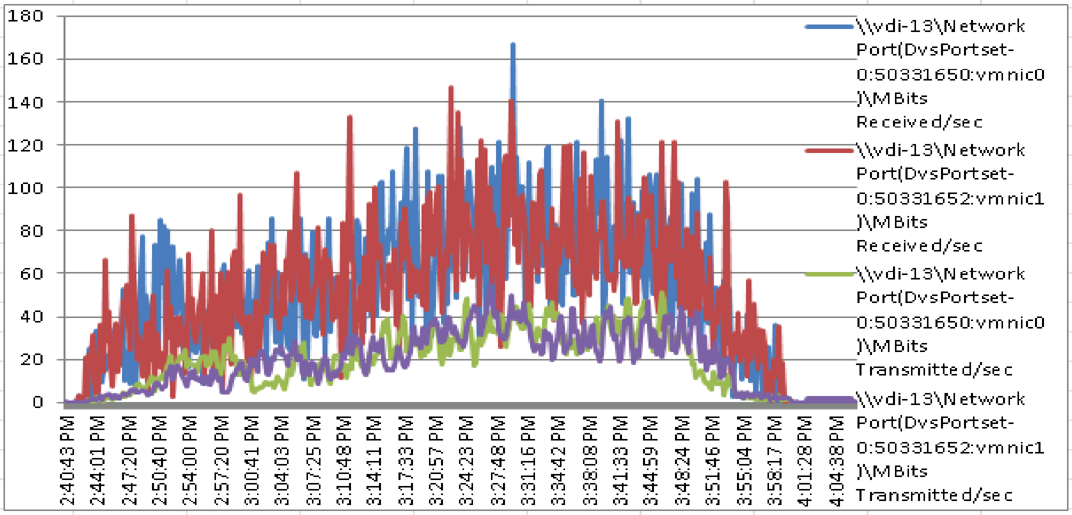

Figure 156 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

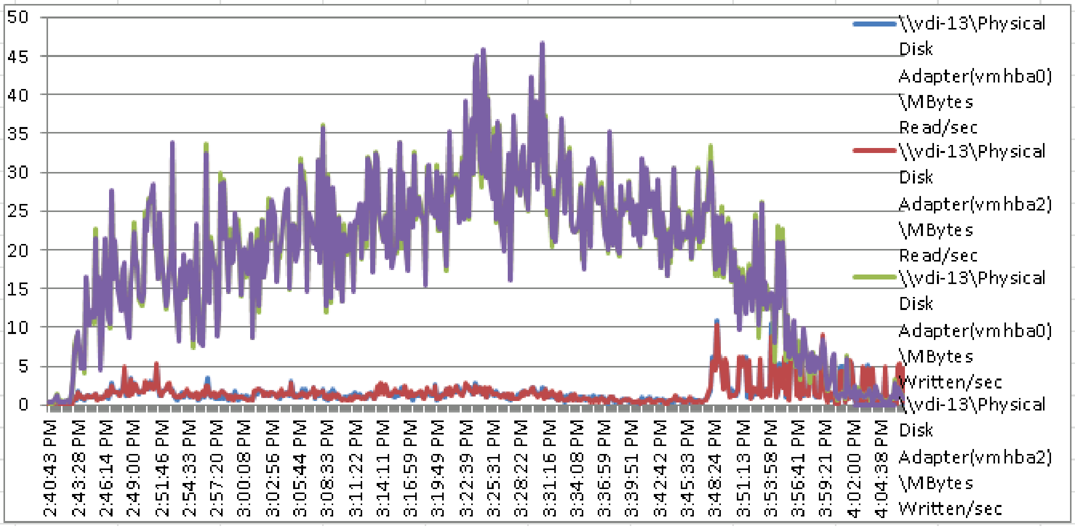

VDI Host–04

Figure 157 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

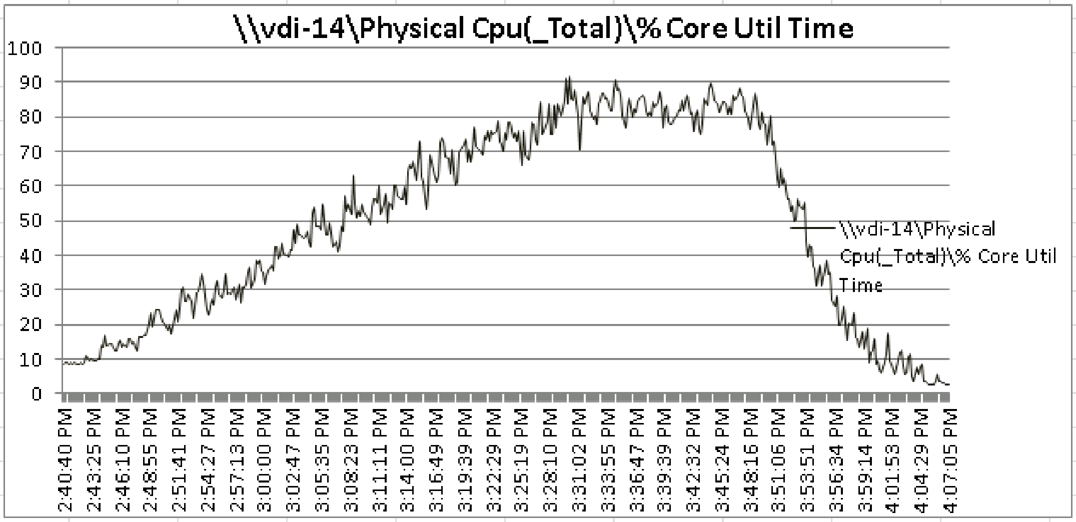

Figure 158 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

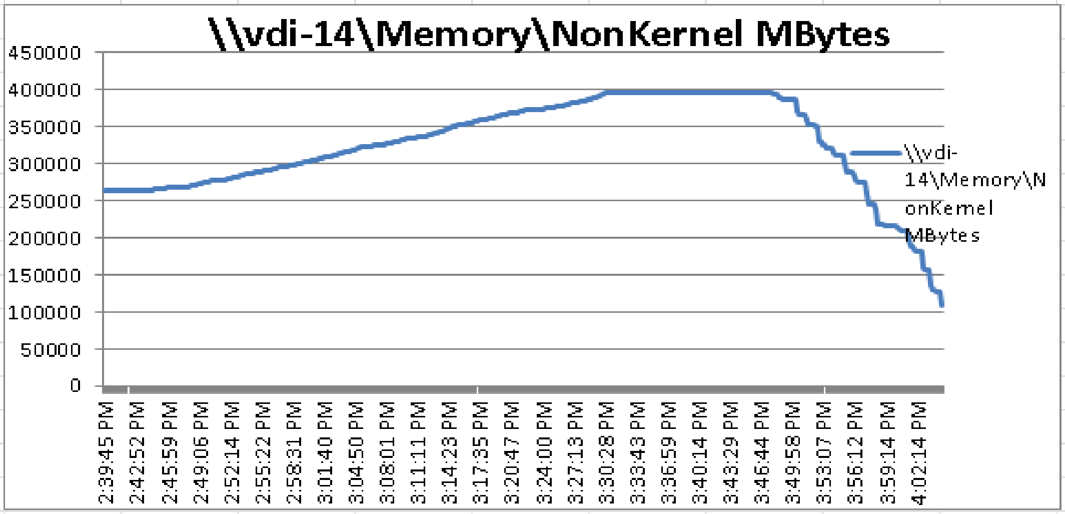

Figure 159 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

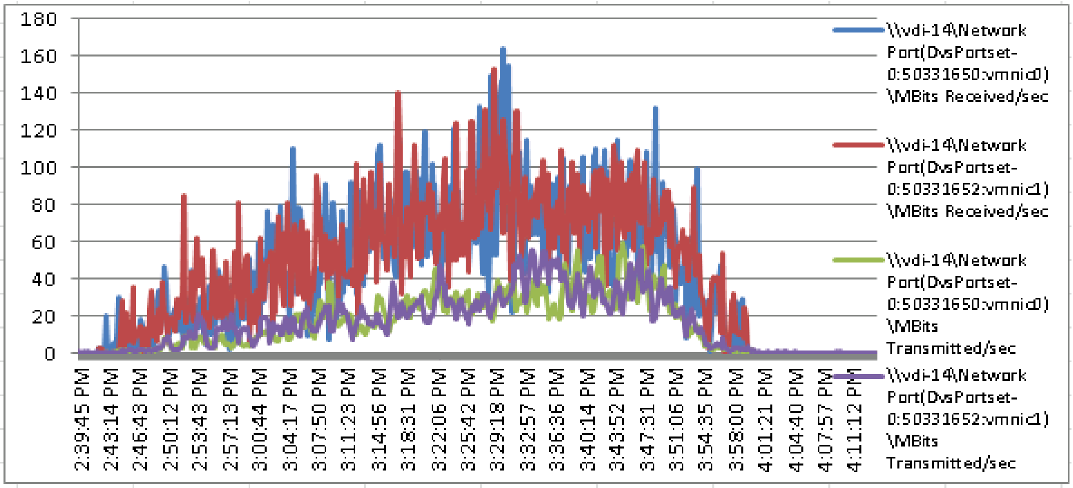

Figure 160 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

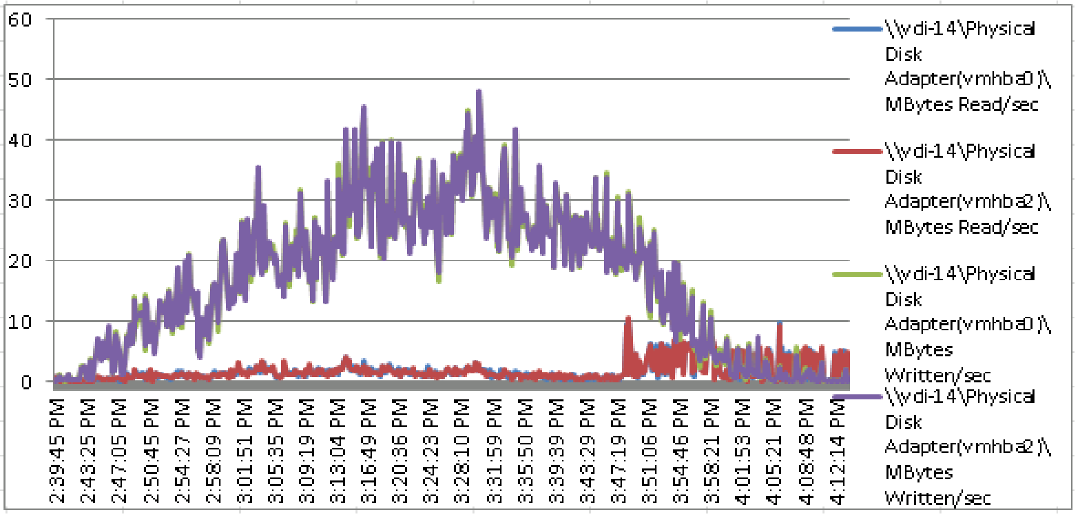

VDI Host-05

Figure 161 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

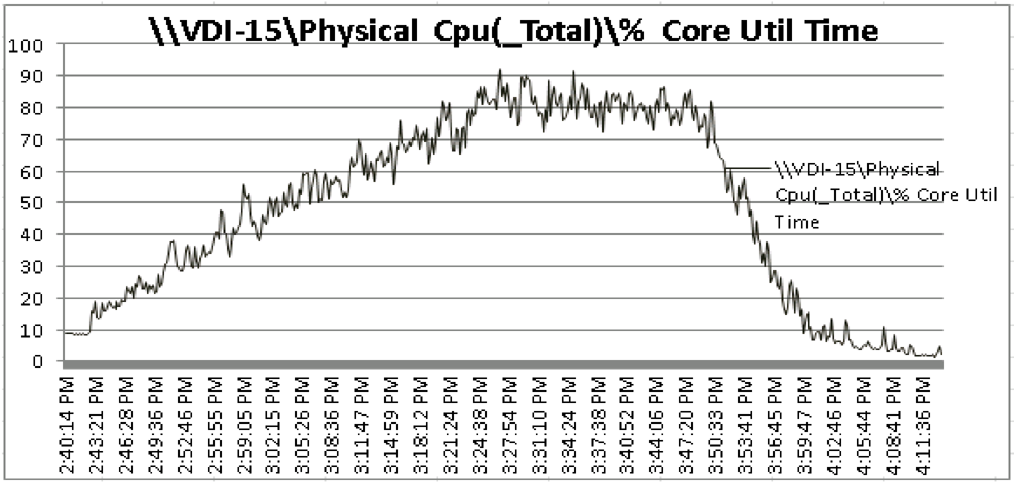

Figure 162 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

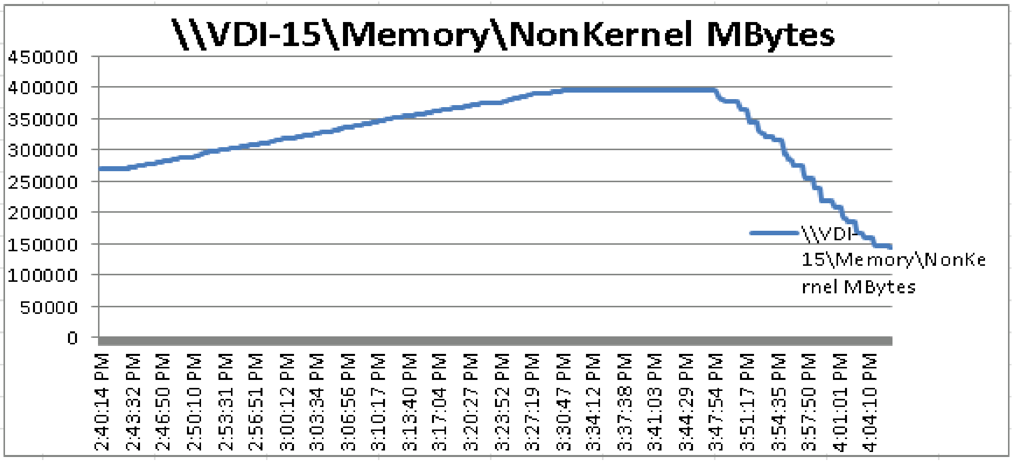

Figure 163 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

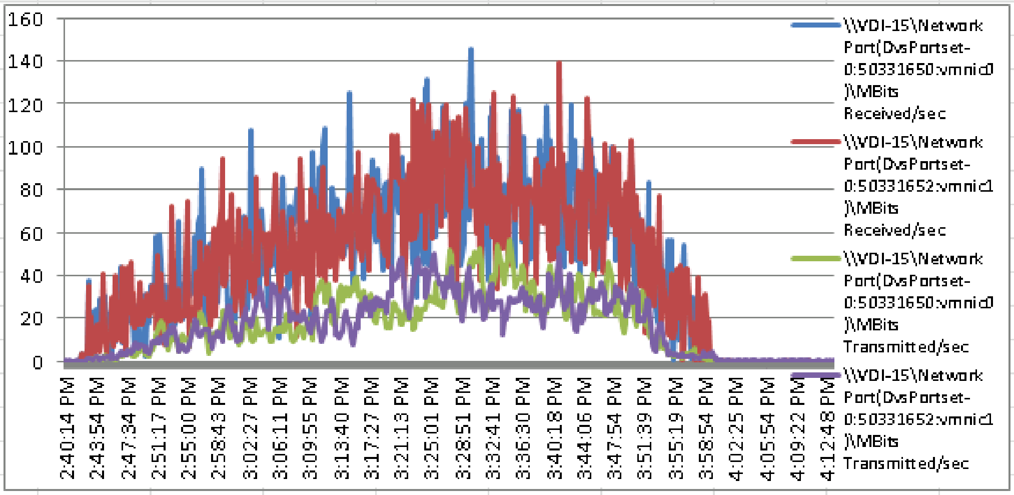

Figure 164 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

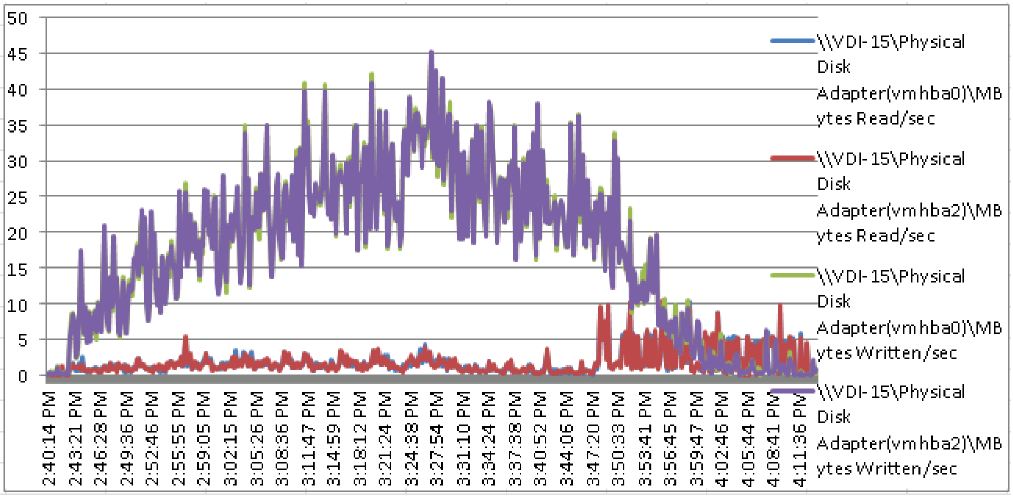

VDI Host-06

Figure 165 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

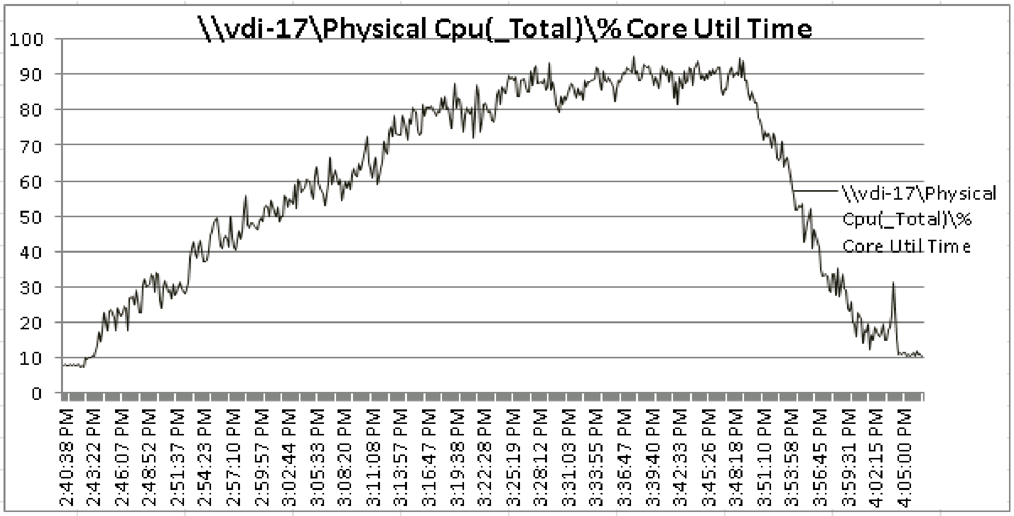

Figure 166 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

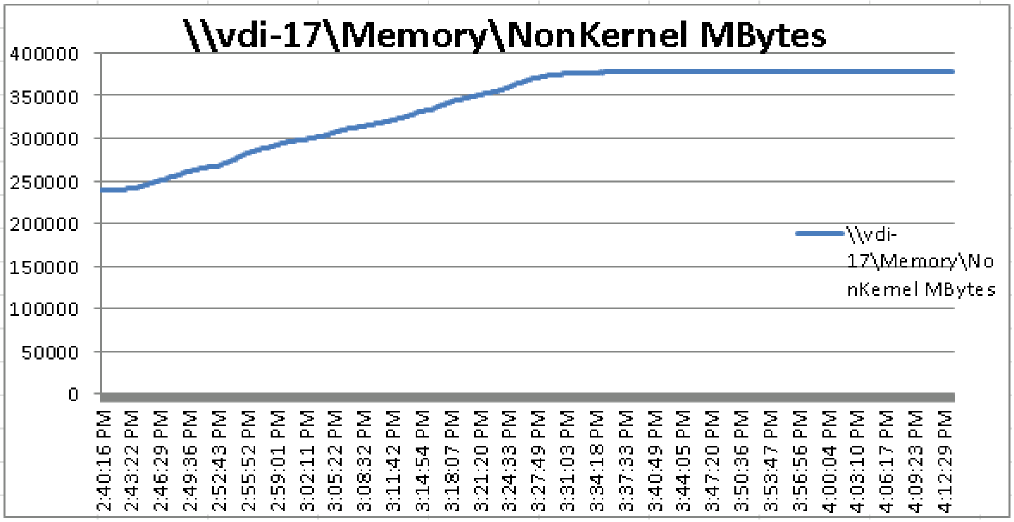

Figure 167 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

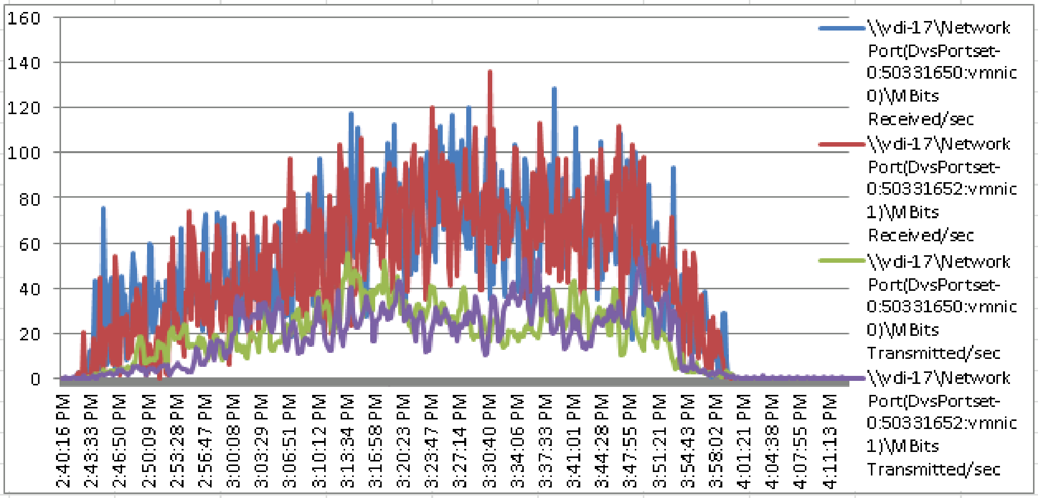

Figure 168 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

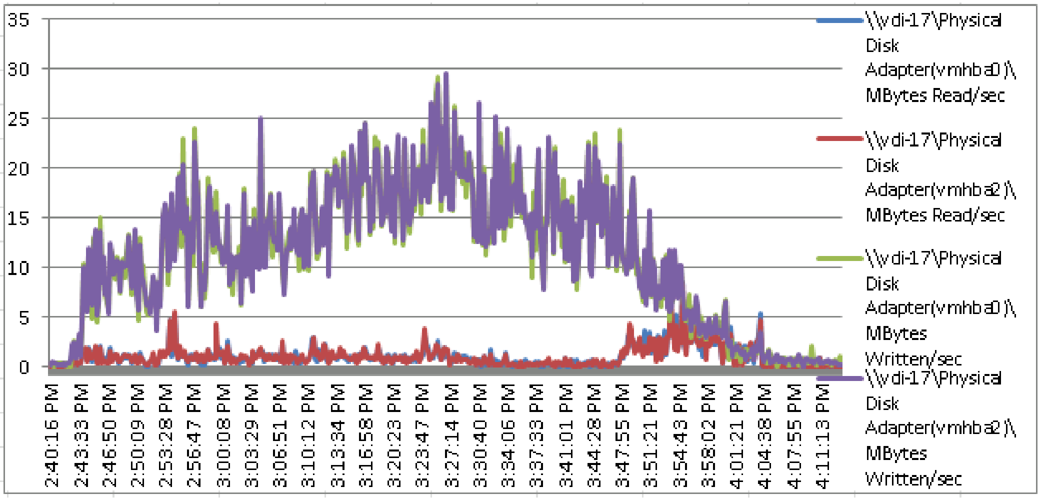

VDI-07

Figure 169 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

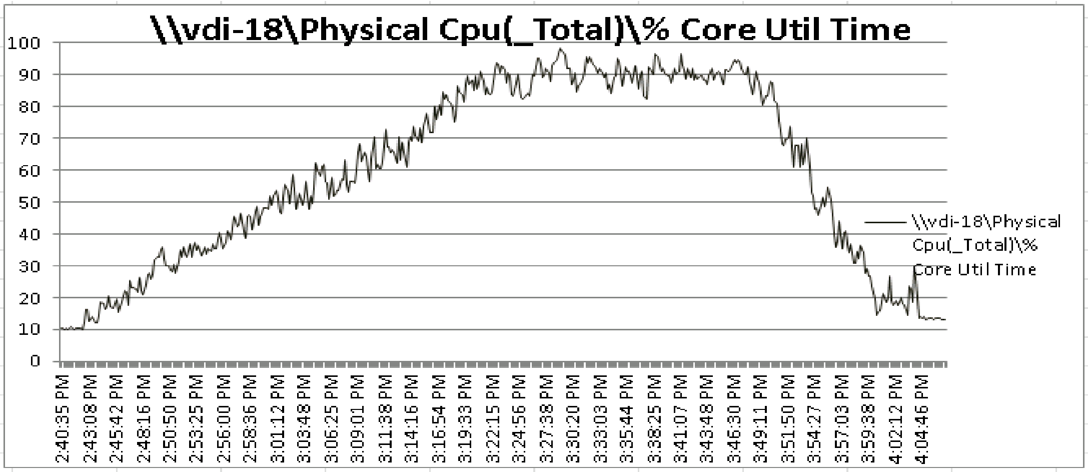

Figure 170 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

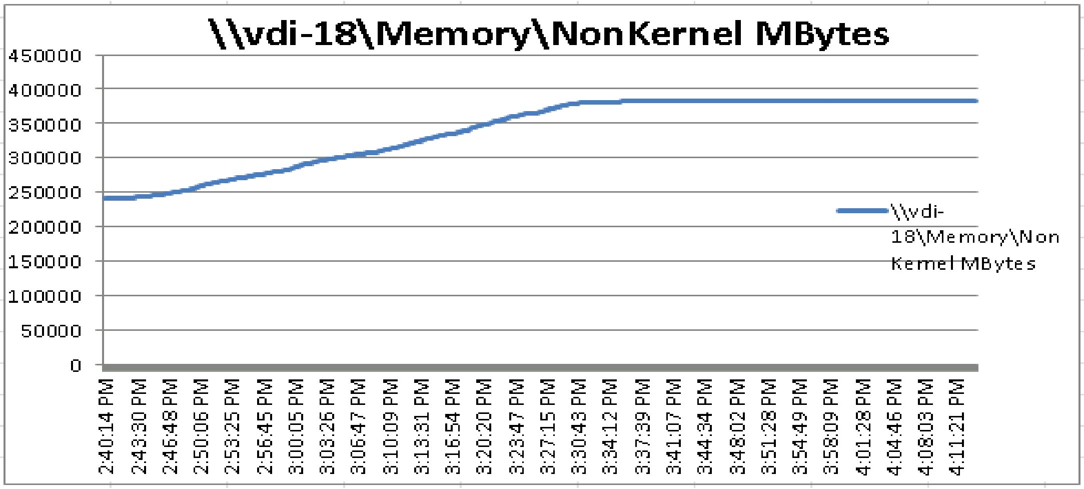

Figure 171 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

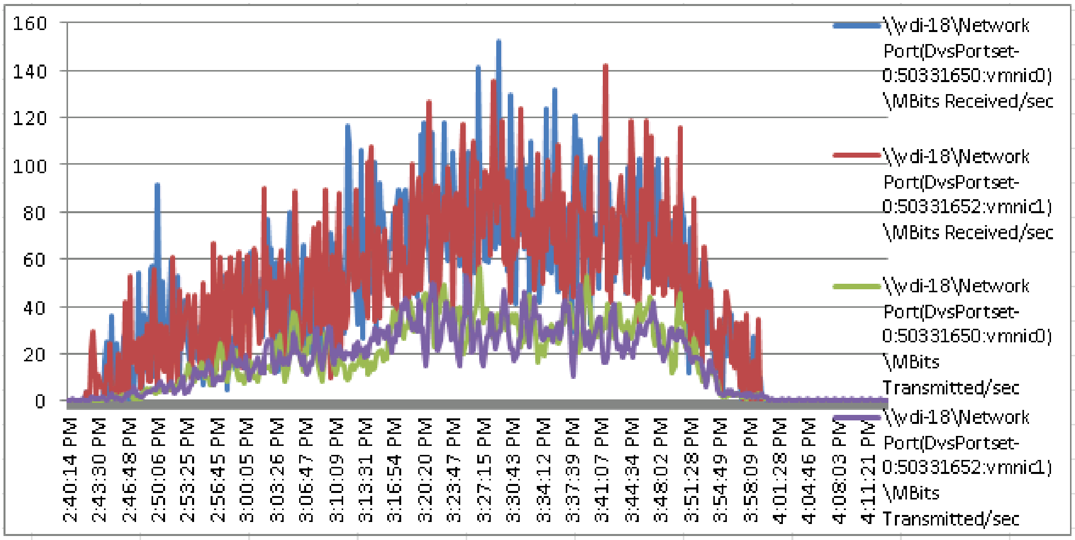

Figure 172 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

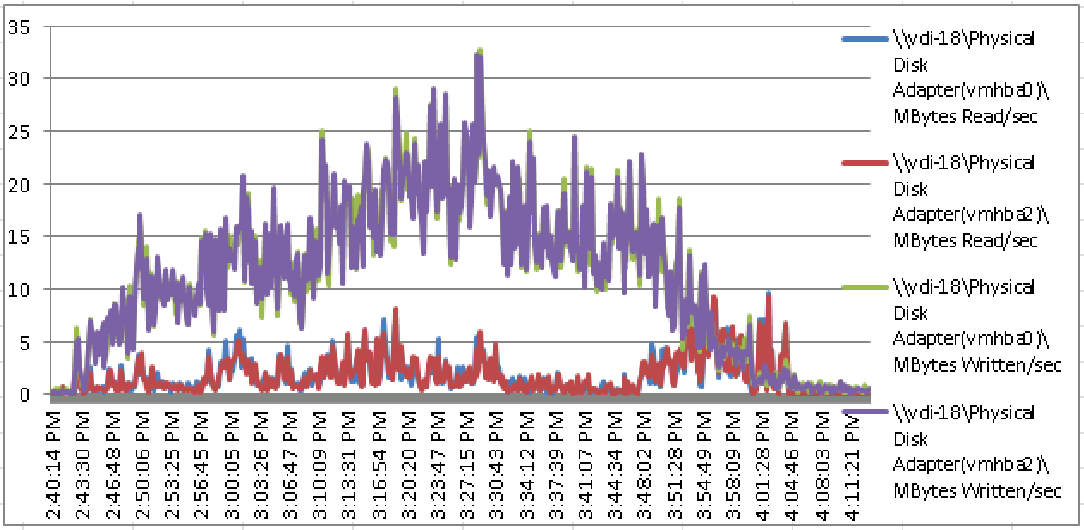

VDI Host-08

Figure 173 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

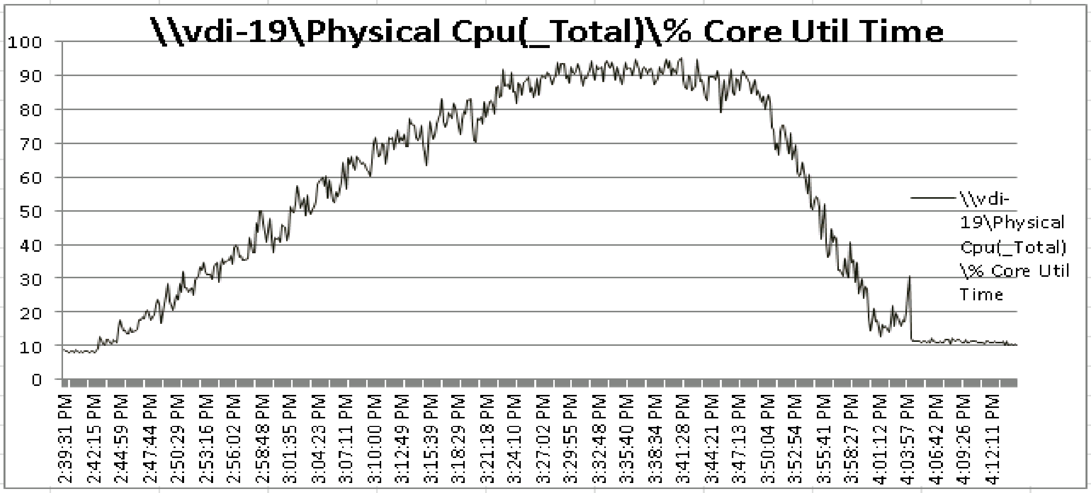

Figure 174 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

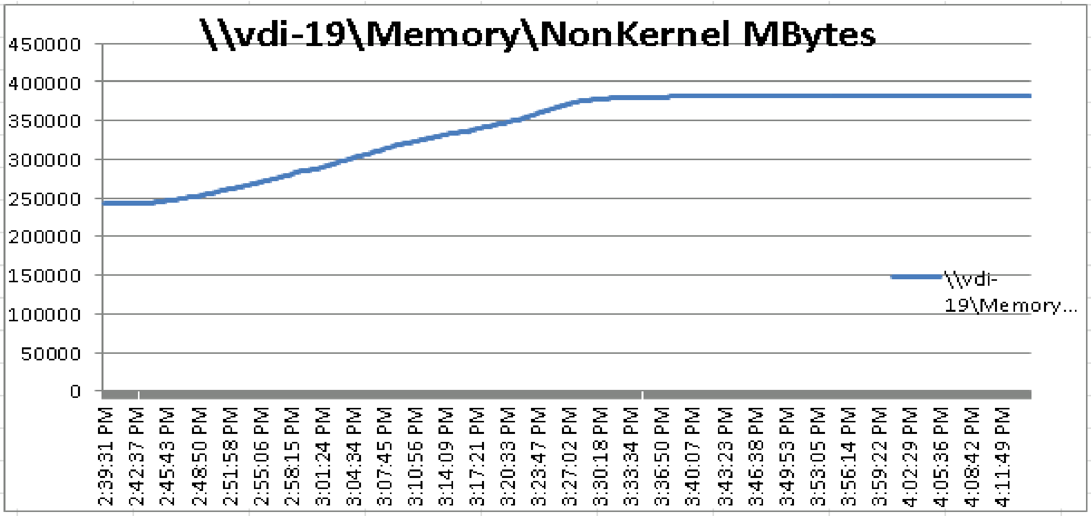

Figure 175 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

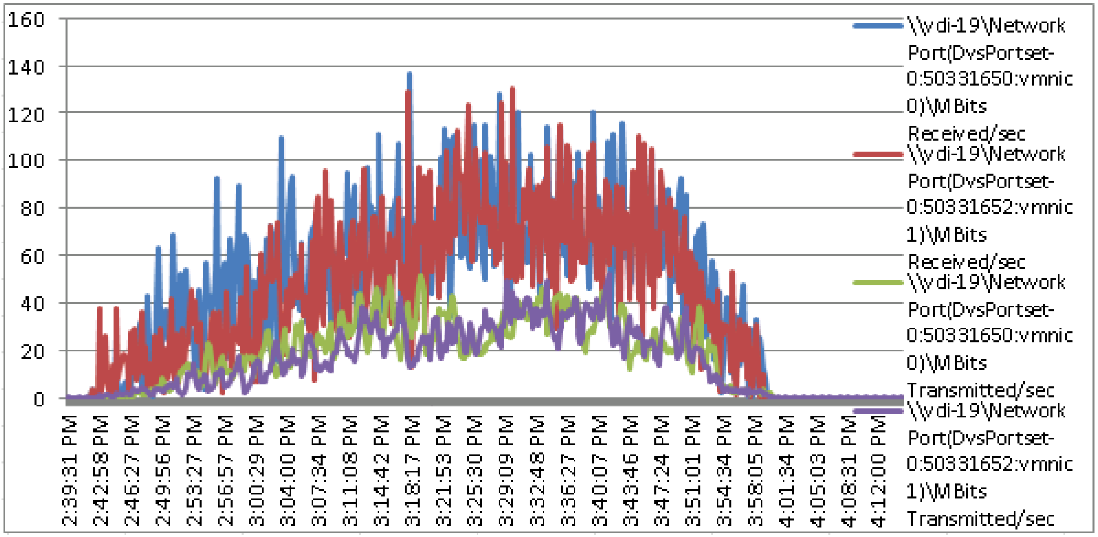

Figure 176 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

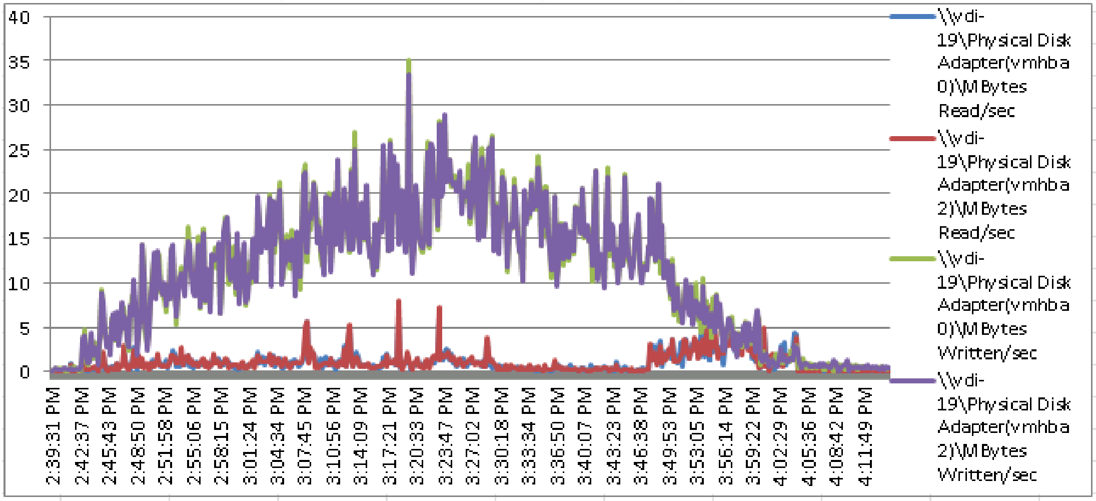

VDI-09

Figure 177 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

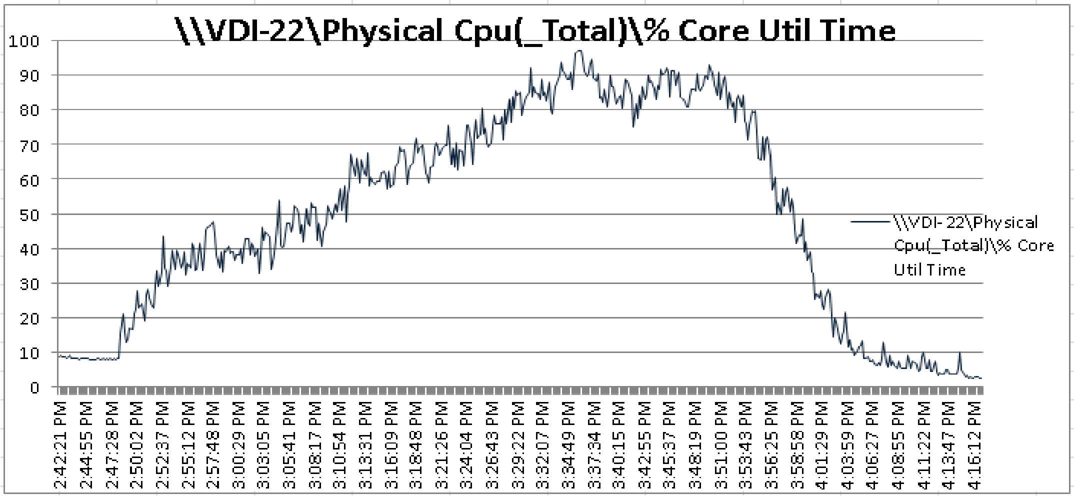

Figure 178 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

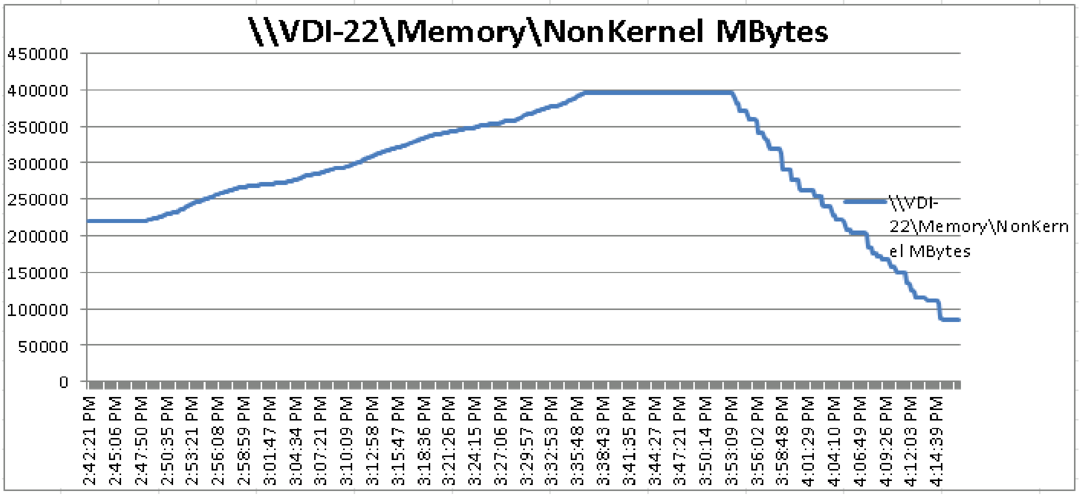

Figure 179 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

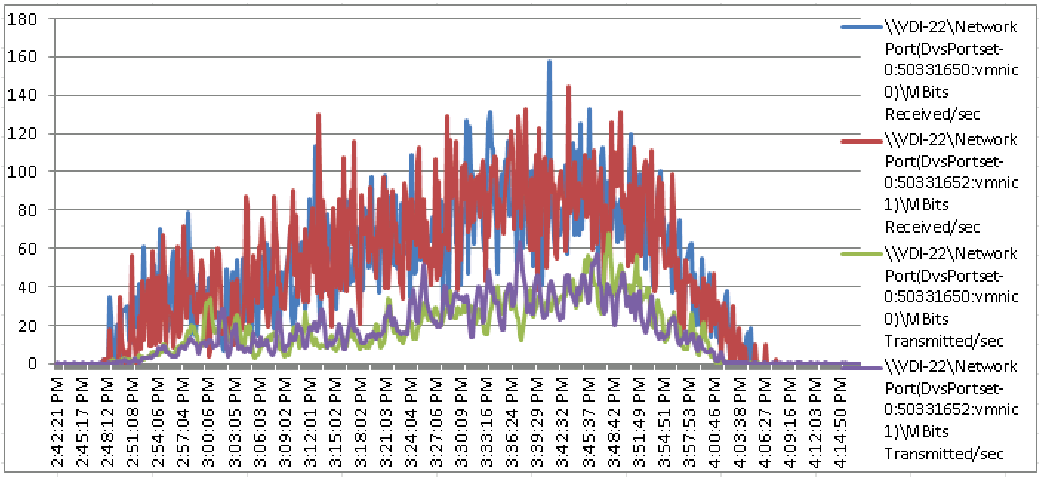

Figure 180 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

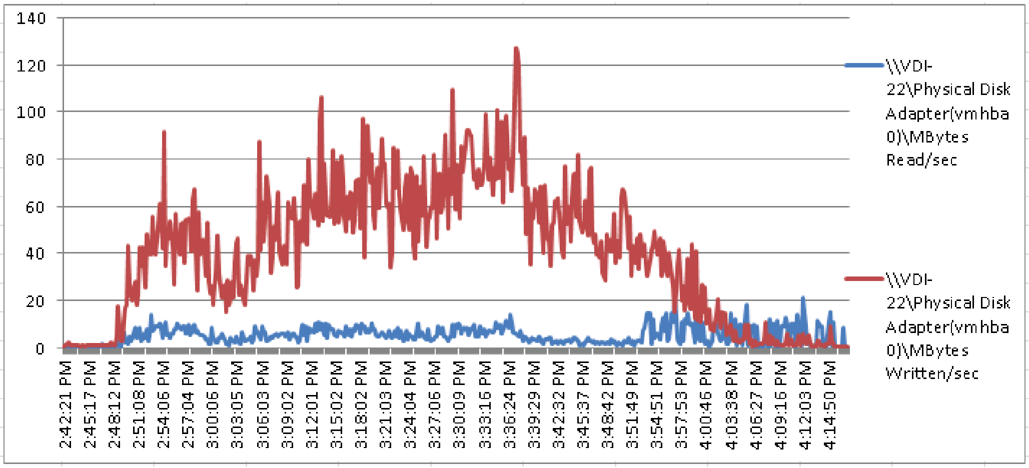

VDI Host-10

Figure 181 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

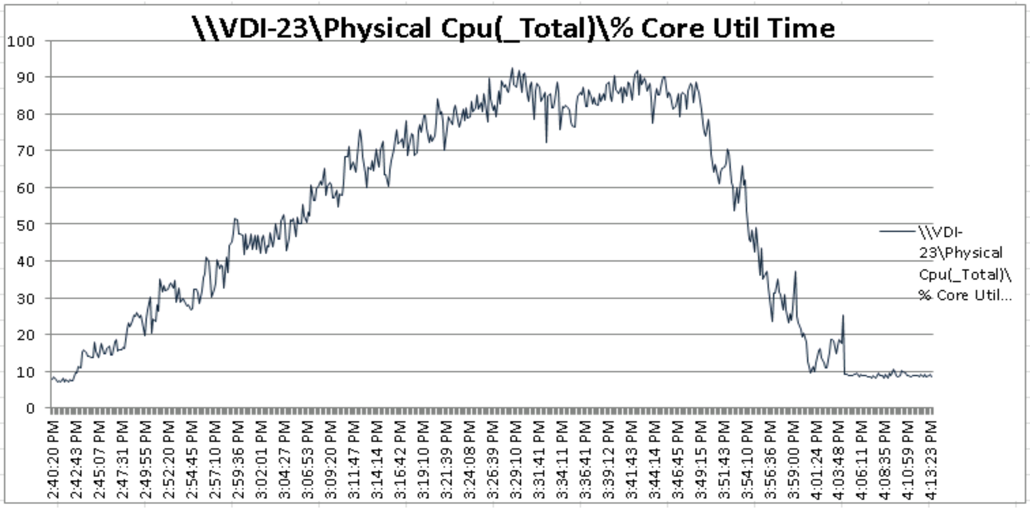

Figure 182 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

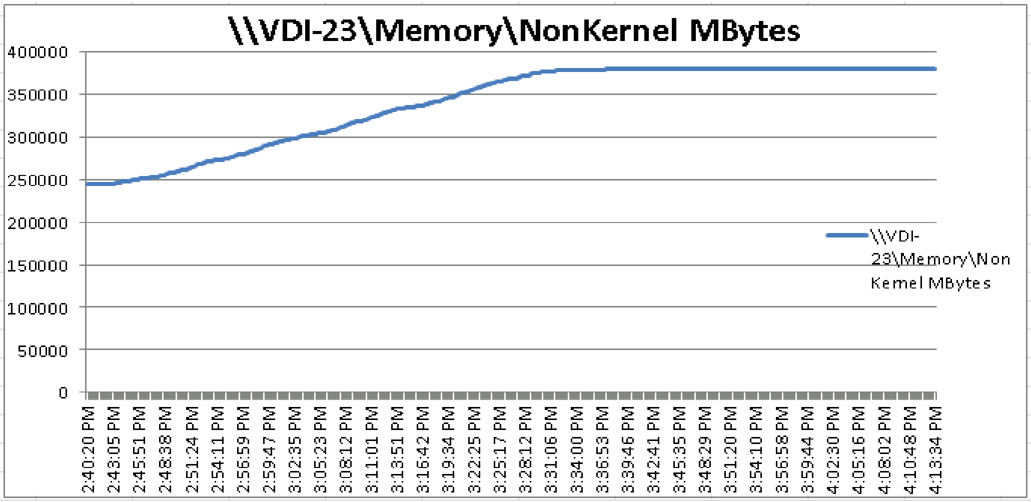

Figure 183 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

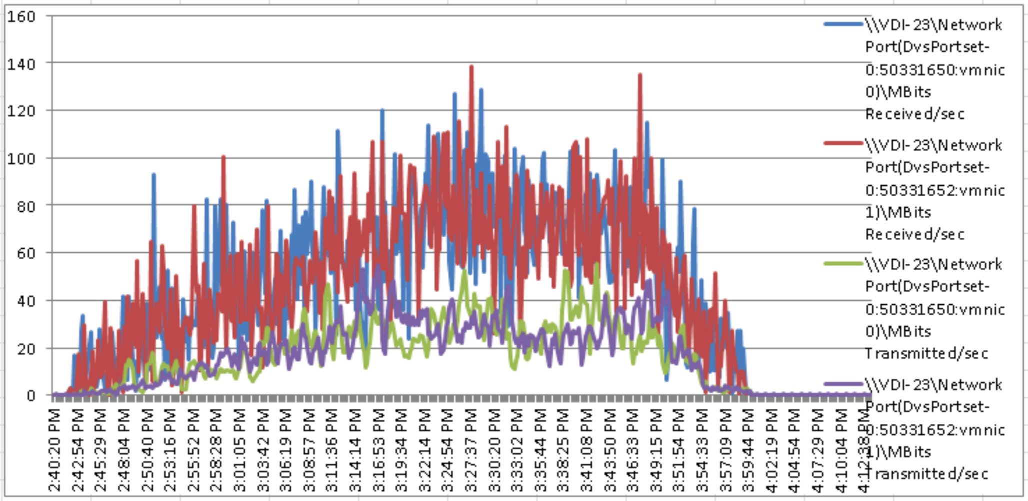

Figure 184 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

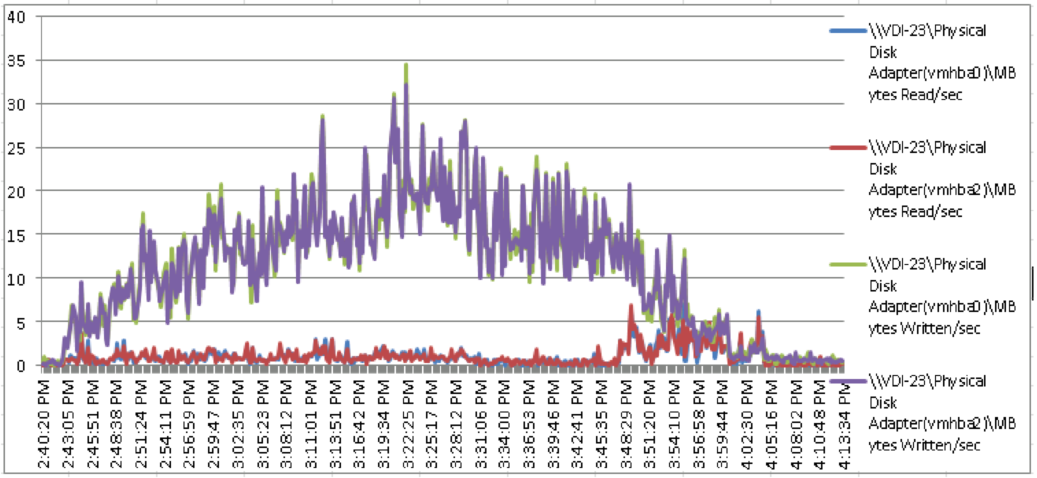

VDI Host-11

Figure 185 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

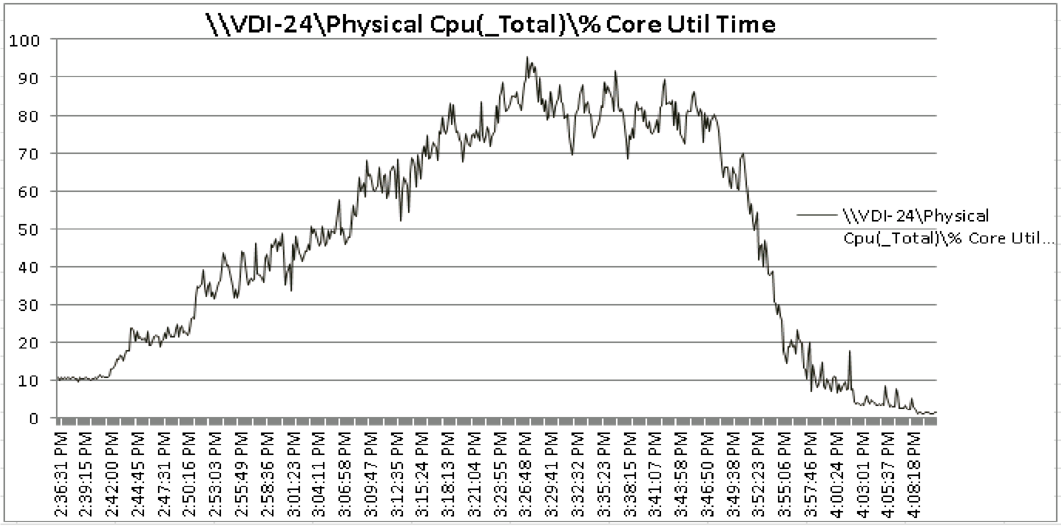

Figure 186 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

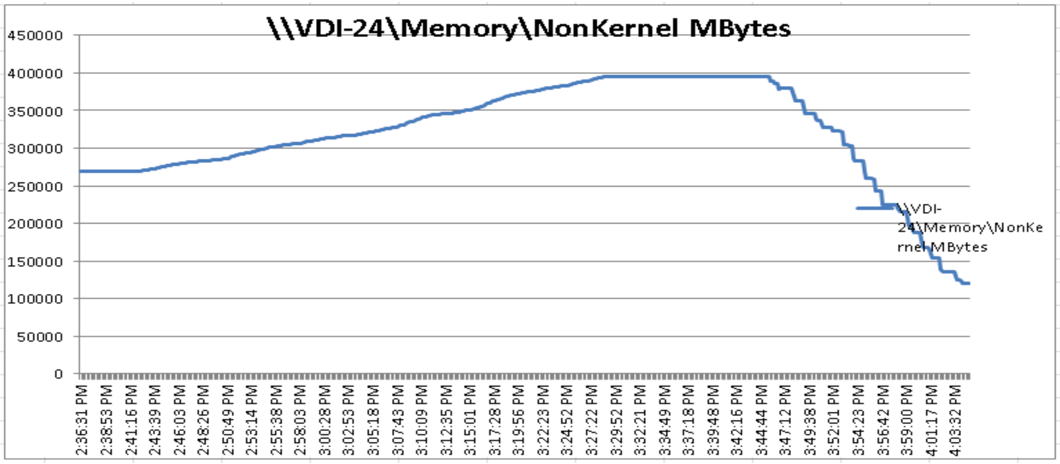

Figure 187 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

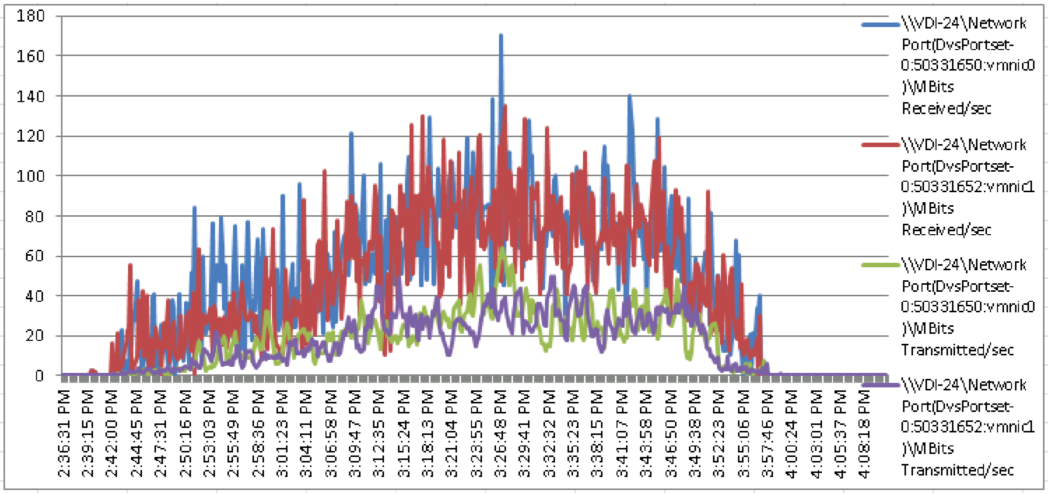

Figure 188 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

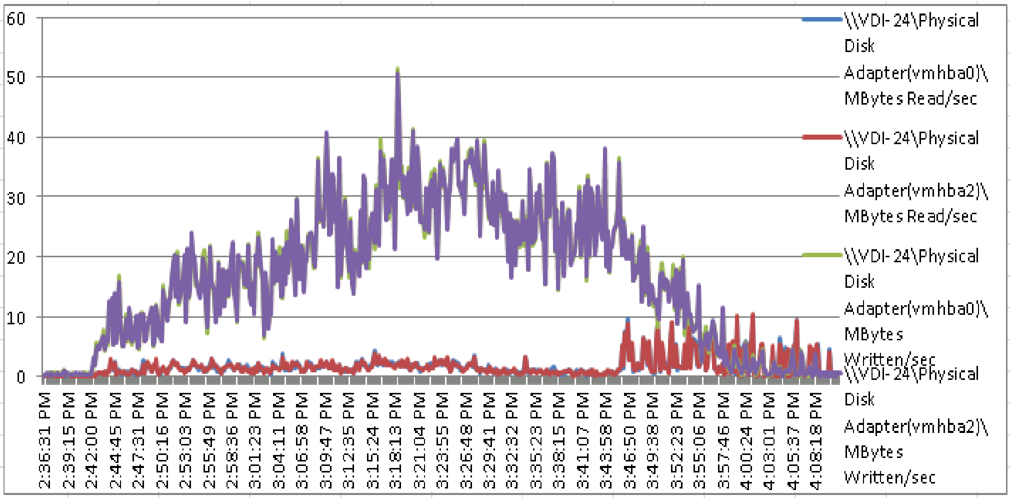

VDI Host-12

Figure 189 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

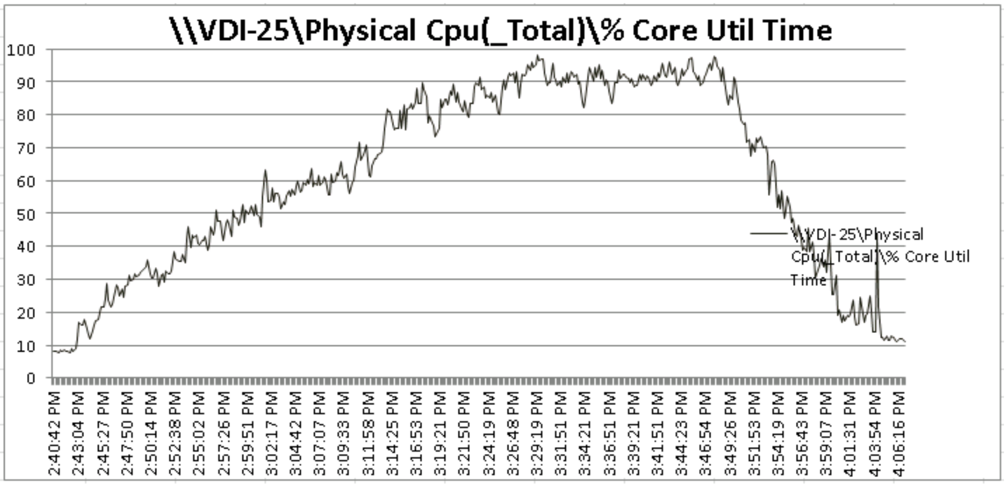

Figure 190 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

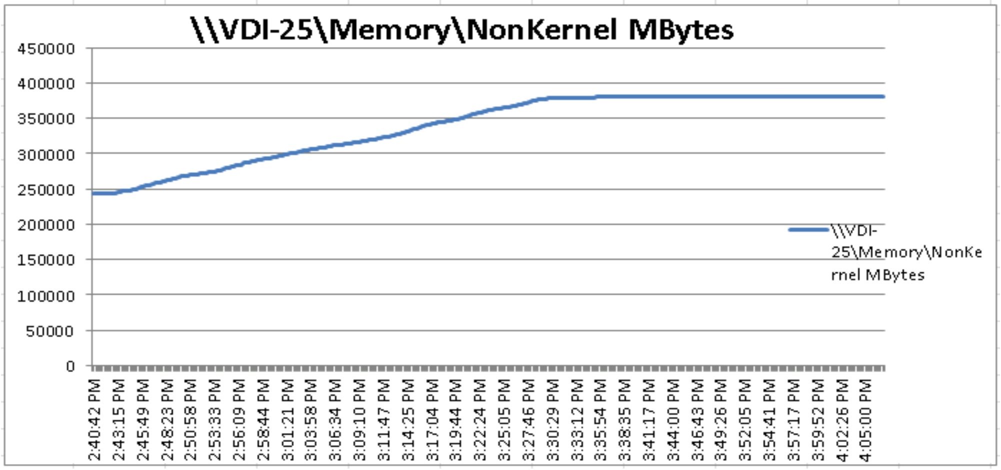

Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

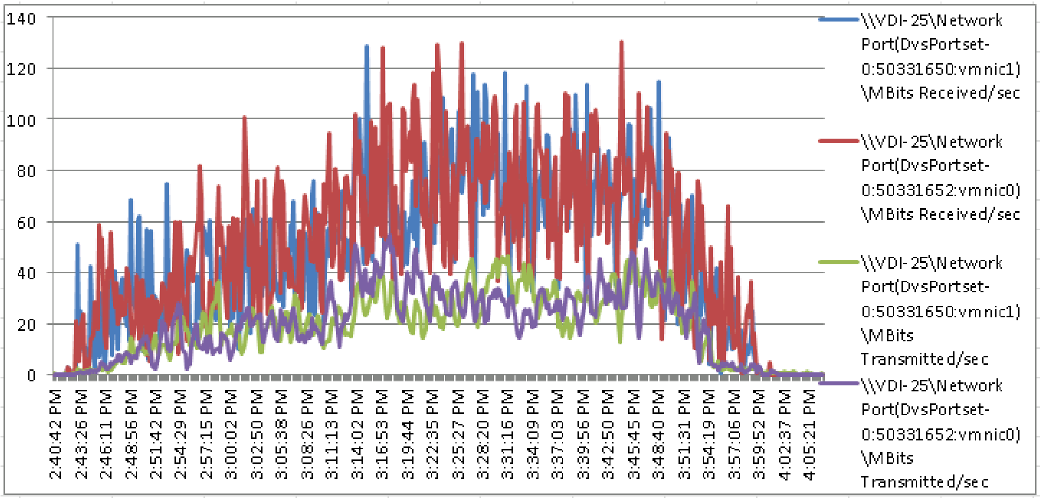

Figure 191 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

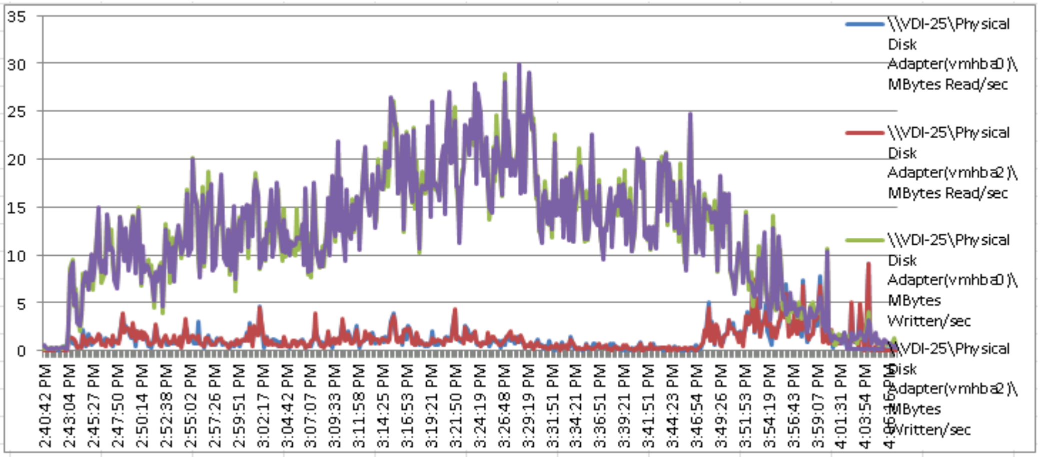

VDI Host-13

Figure 192 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

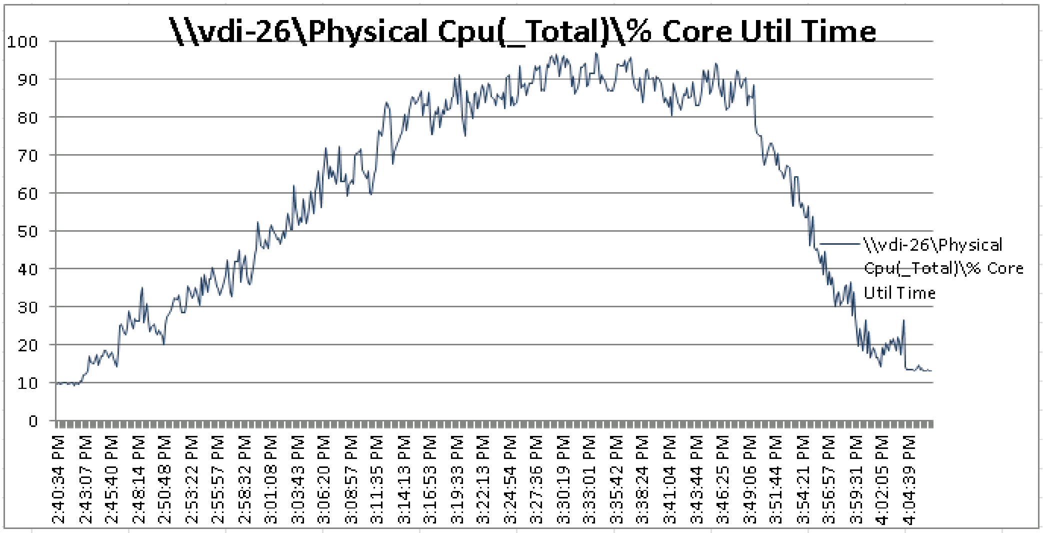

Figure 193 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

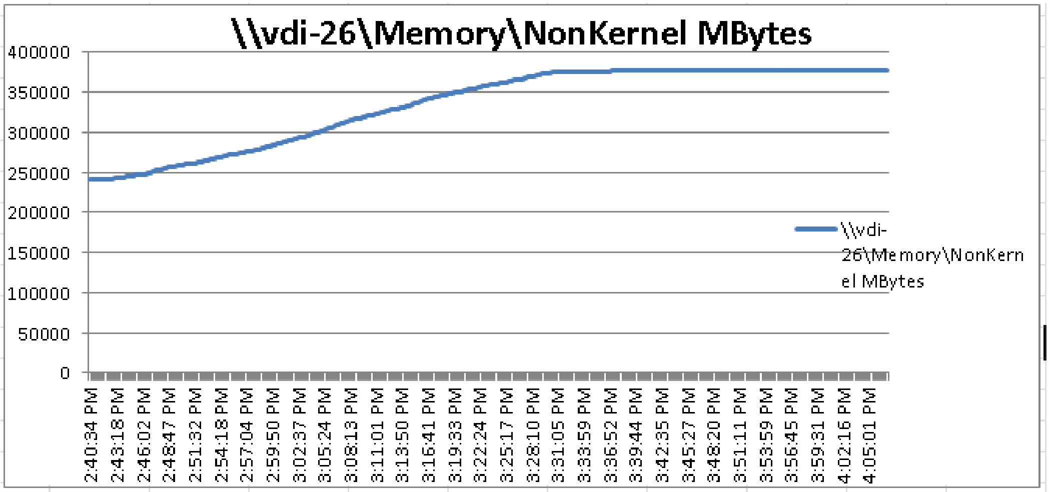

Figure 194 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

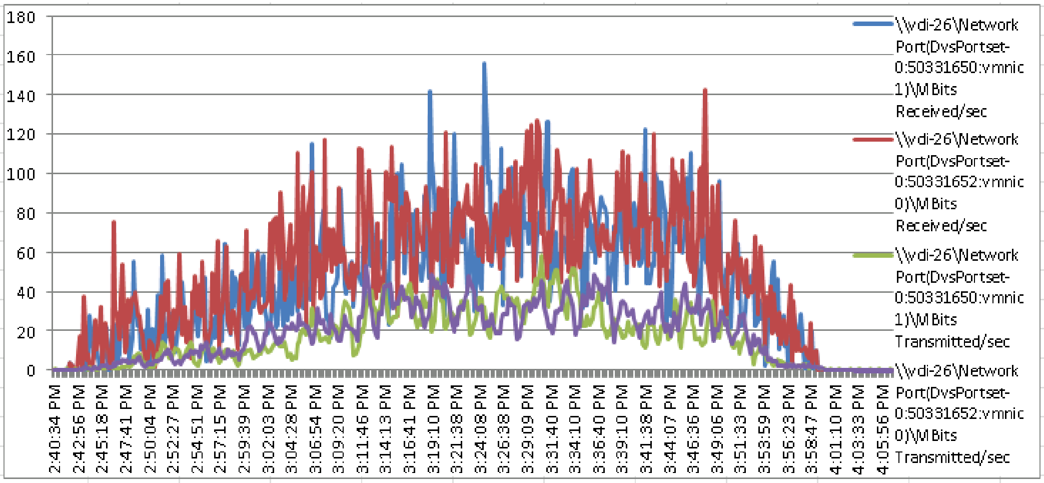

Figure 195 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

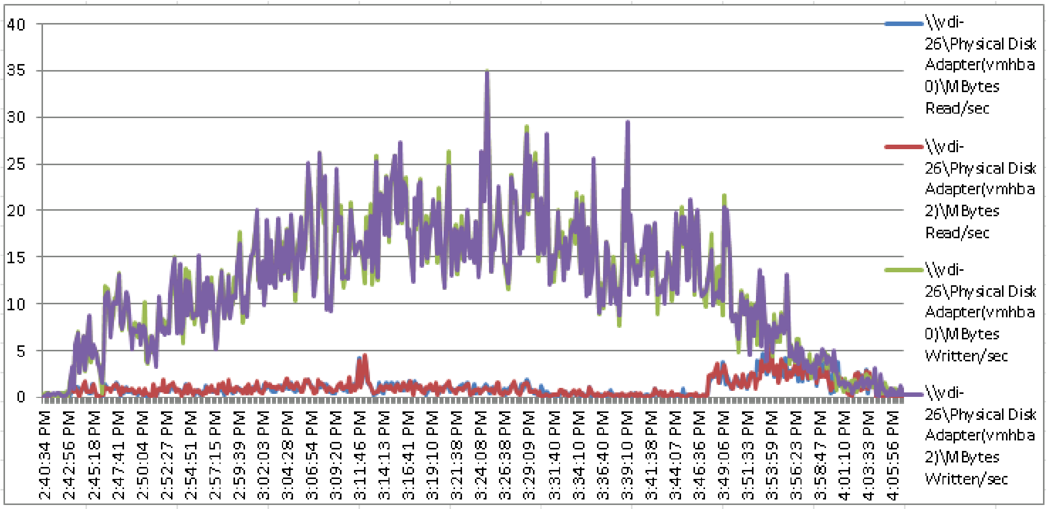

VDI Host-14

Figure 196 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

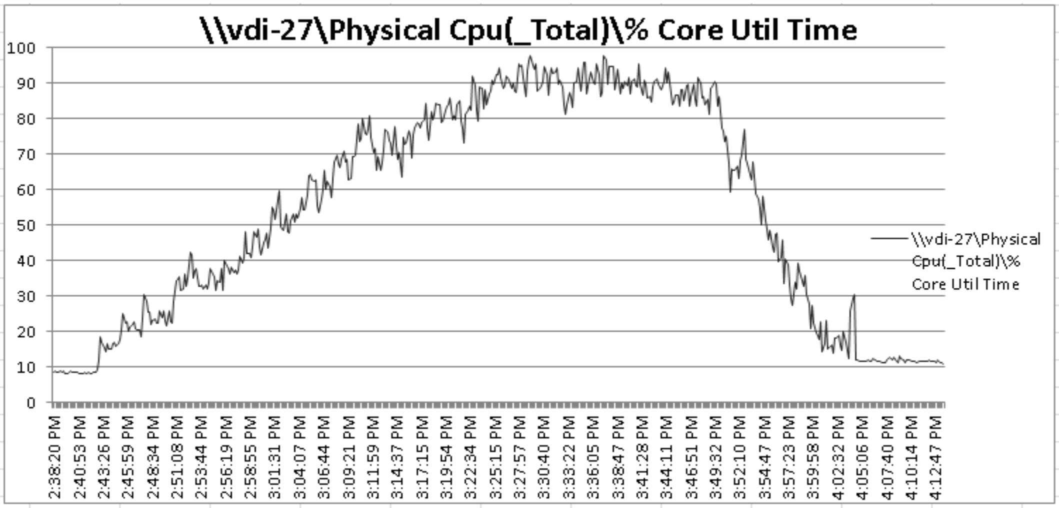

Figure 197 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

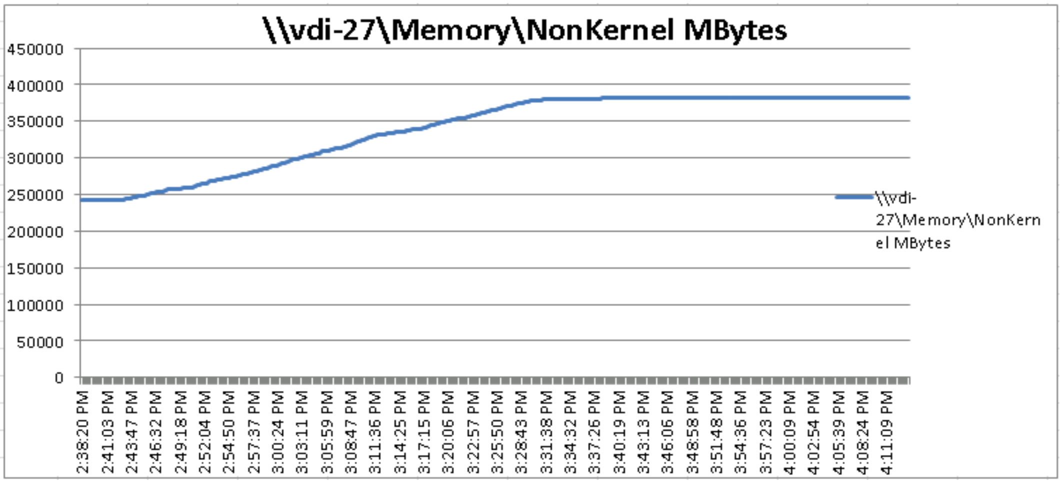

Figure 198 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

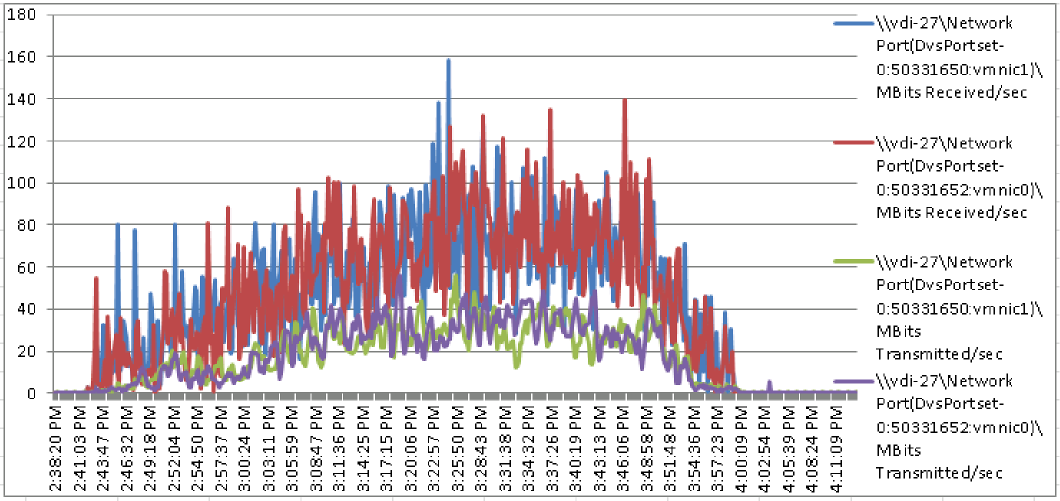

Figure 199 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

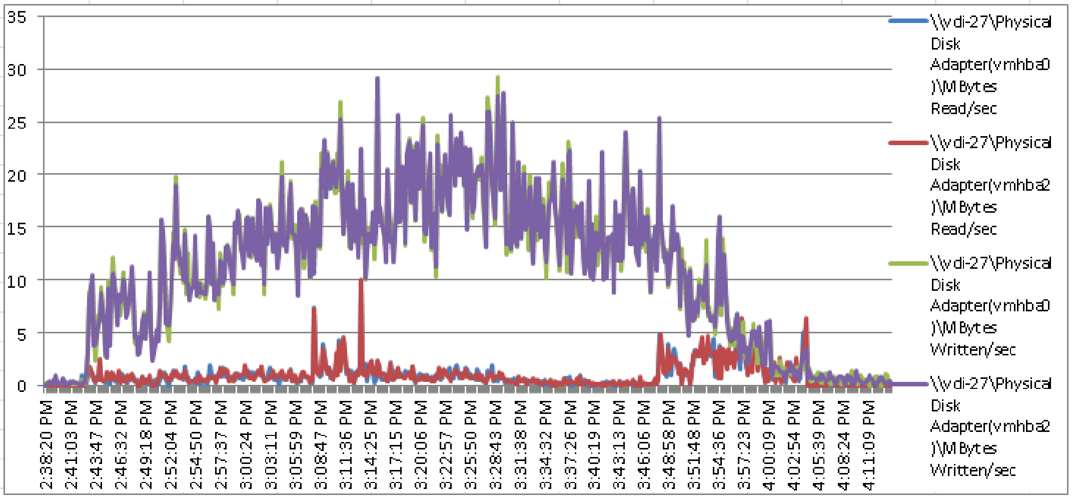

Figure 200 VDI Host-15

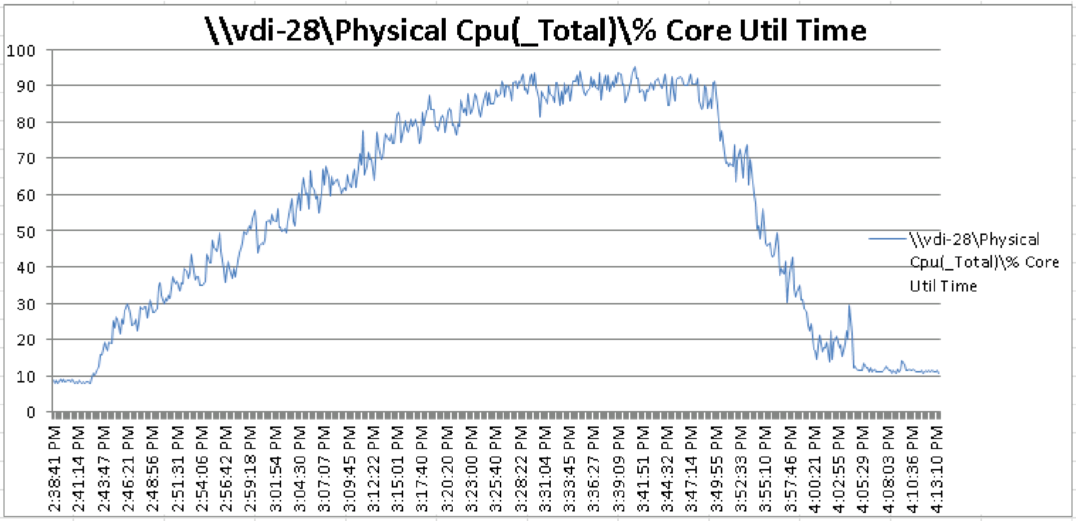

Figure 201 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

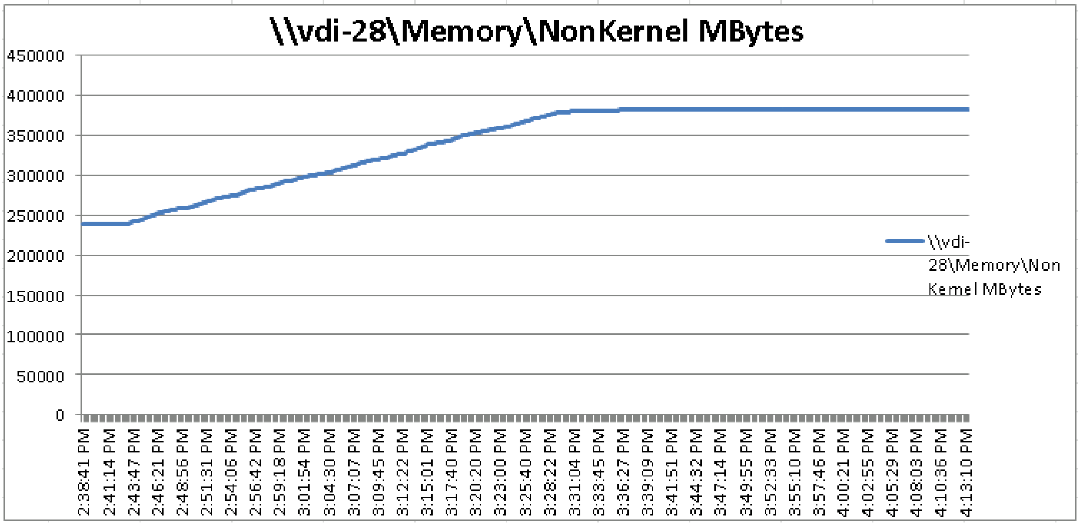

Figure 202 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

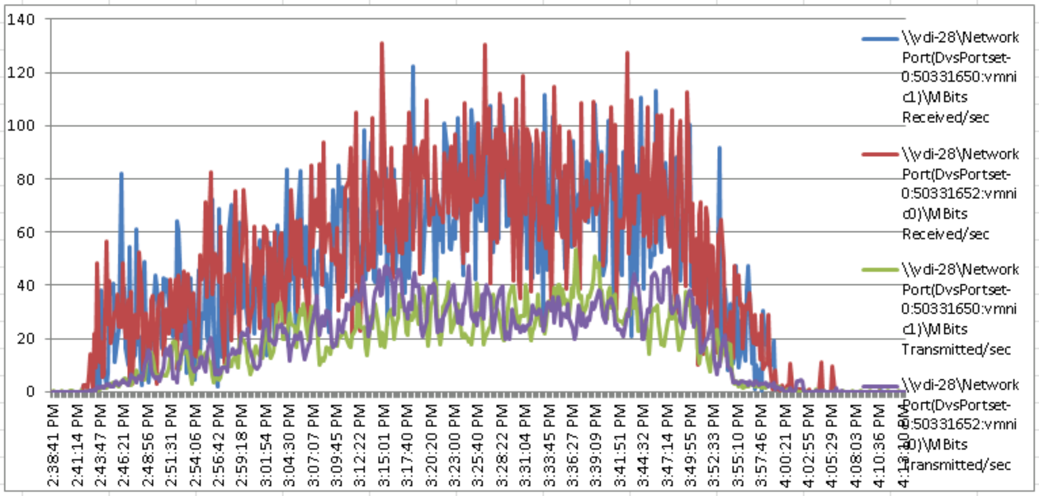

Figure 203 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

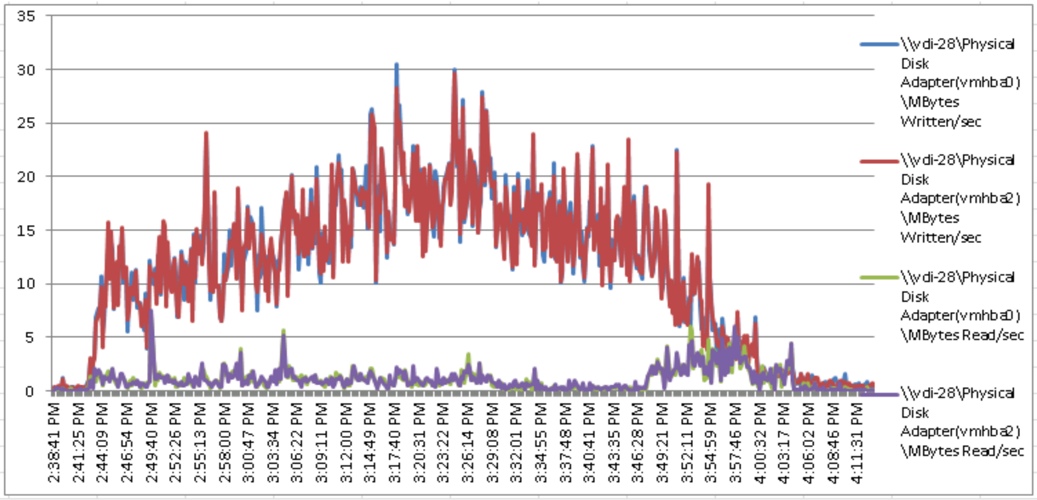

VDI Host-16

Figure 204 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

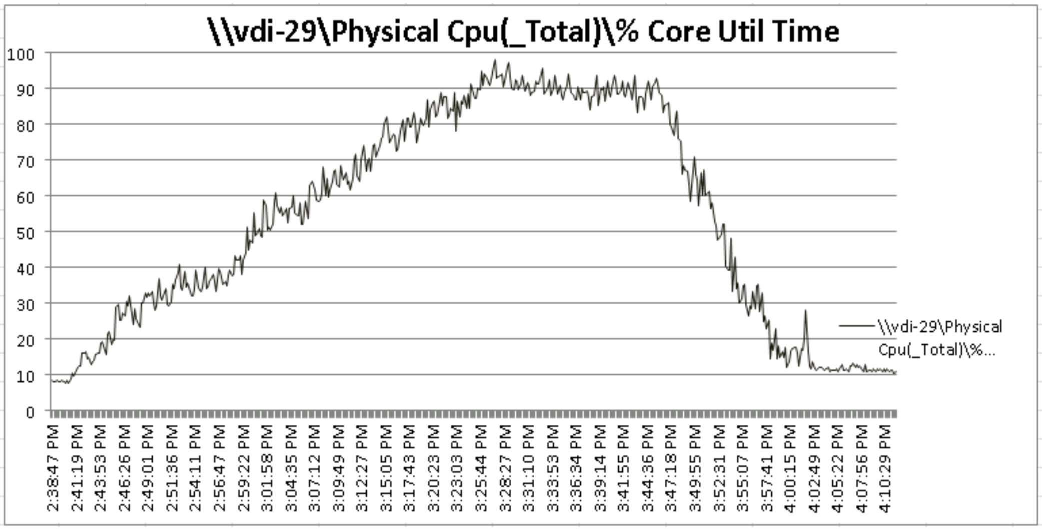

Figure 205 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

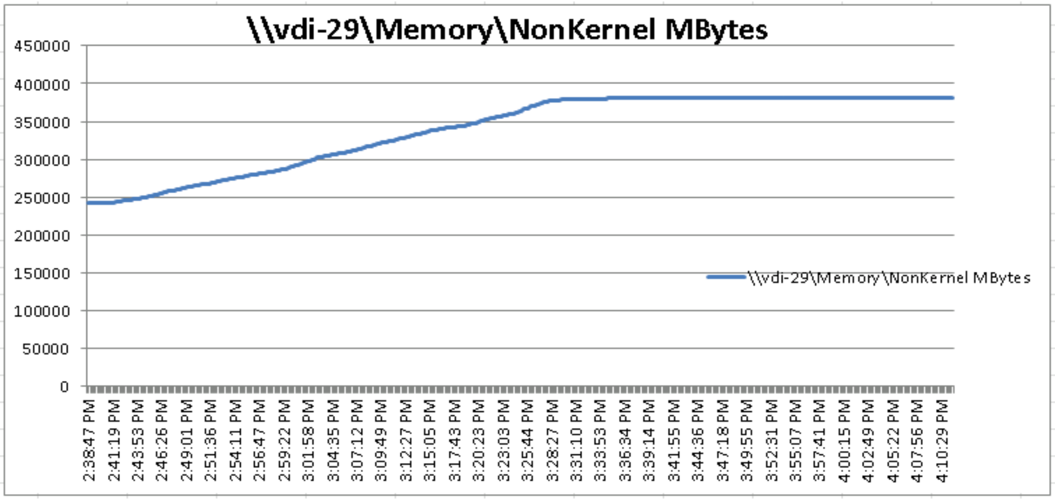

Figure 206 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

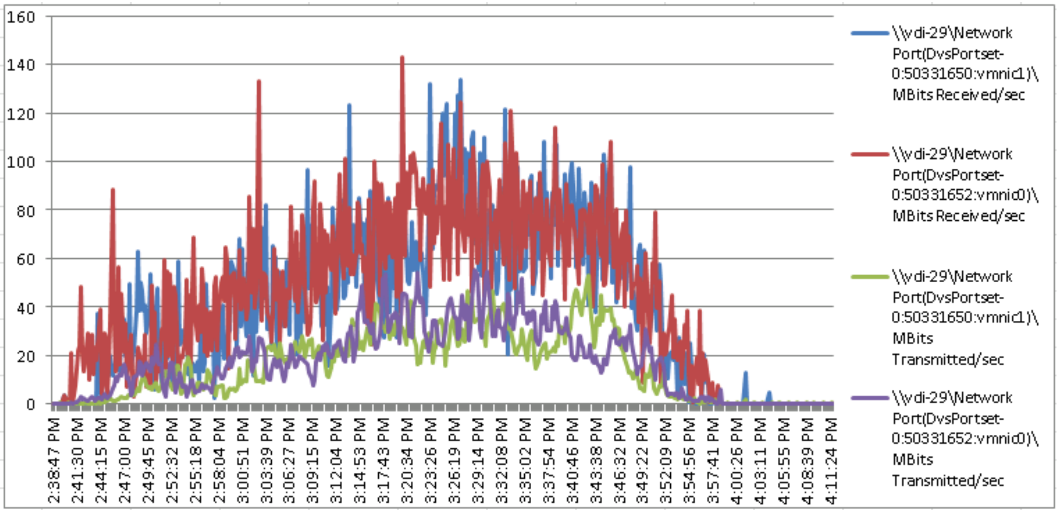

Figure 207 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

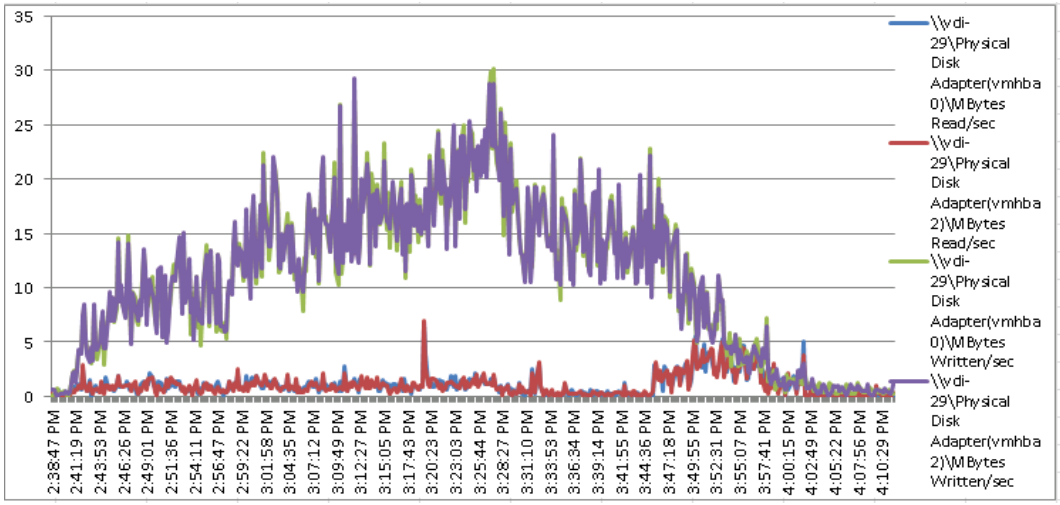

VDI Host-17

Figure 208 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

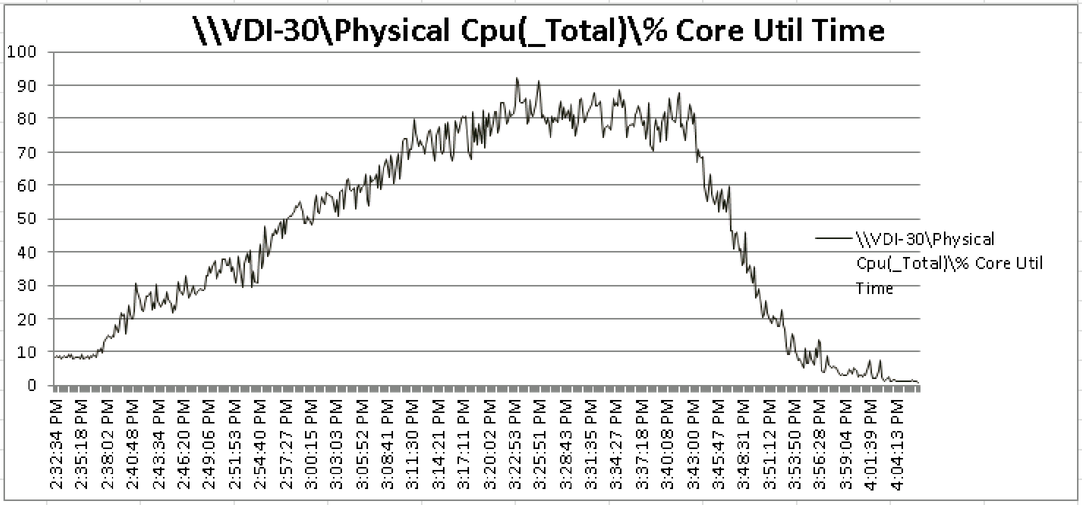

Figure 209 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

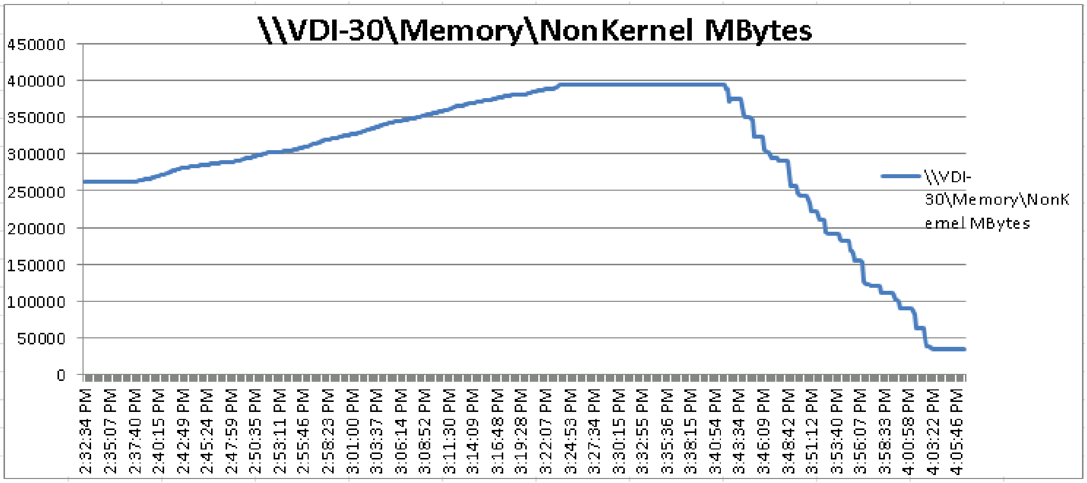

Figure 210 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

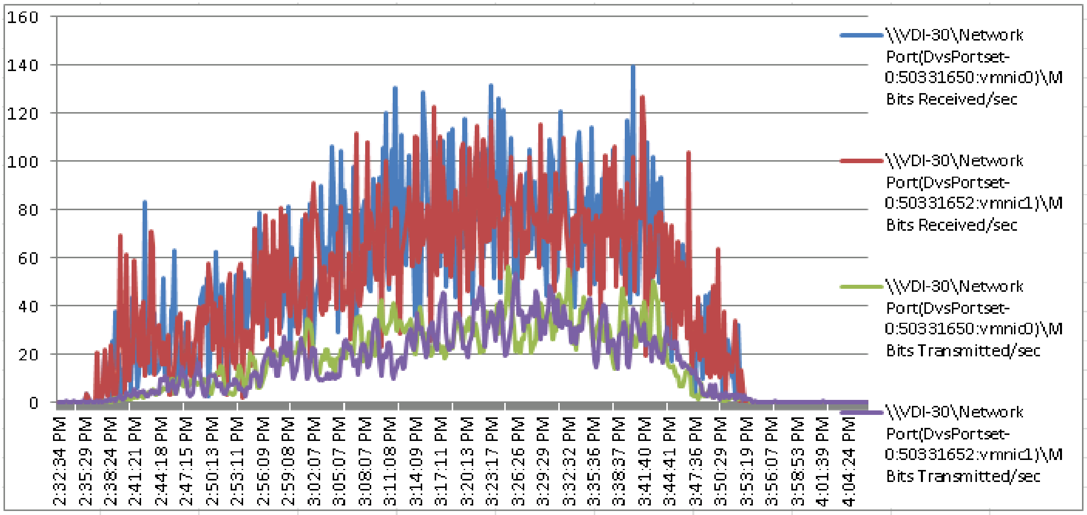

Figure 211 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

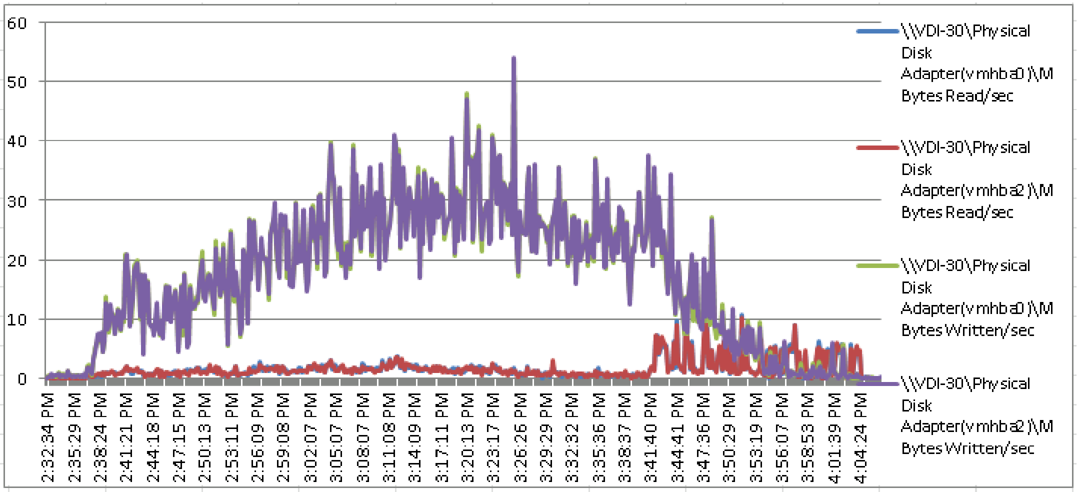

VDI Host-18

Figure 212 Full Scale | 5000 Mixed Users | VDI Host | Host CPU Utilization

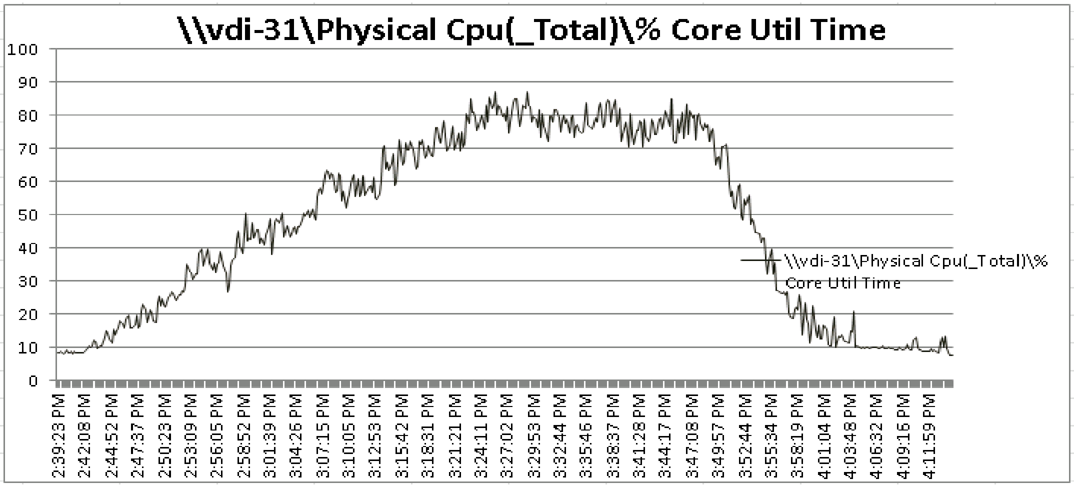

Figure 213 Full Scale | 5000 Mixed Users | VDI Host | Host Memory Utilization

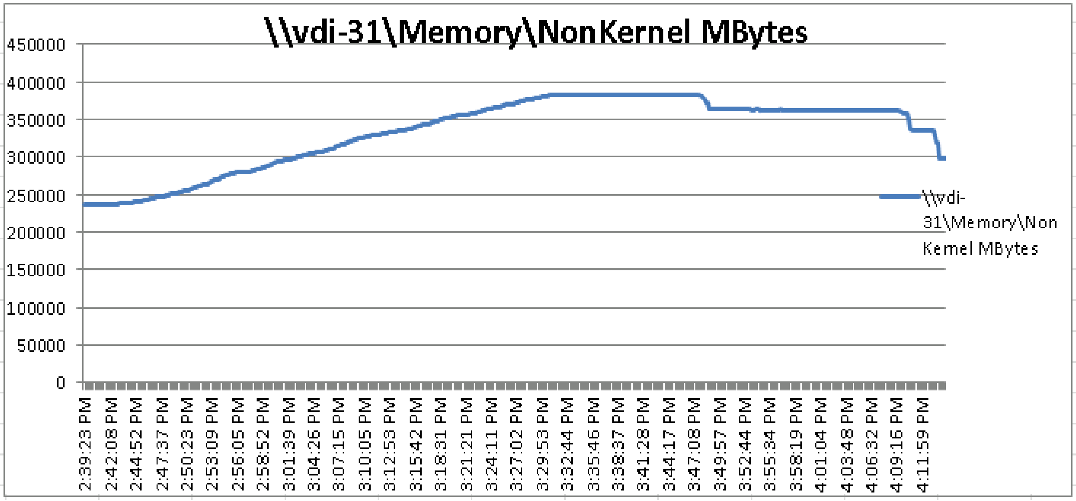

Figure 214 Full Scale | 5000 Mixed Users | VDI Host | Host Network Utilization

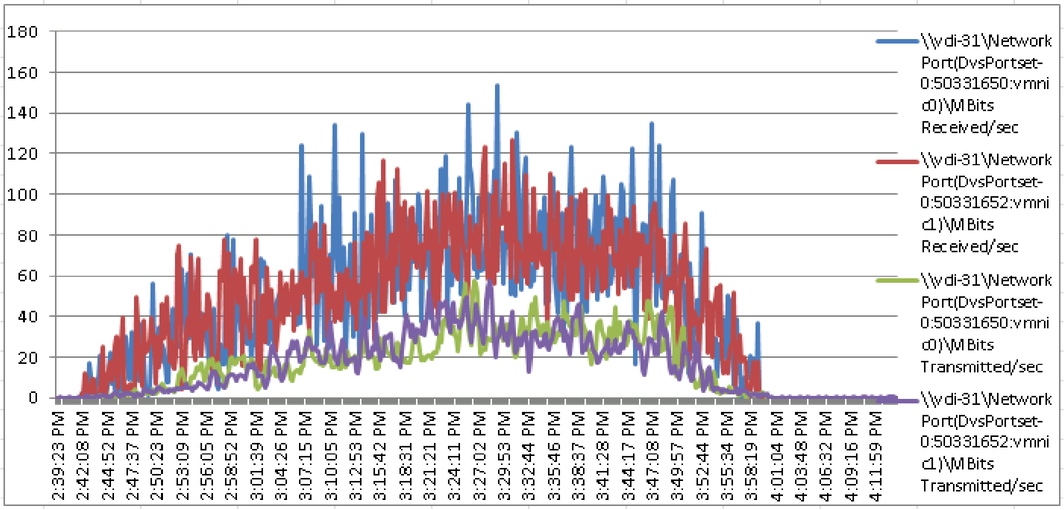

Figure 215 Full Scale | 5000 Mixed Users | VDI Host | Host Fibre Channel Network Utilization

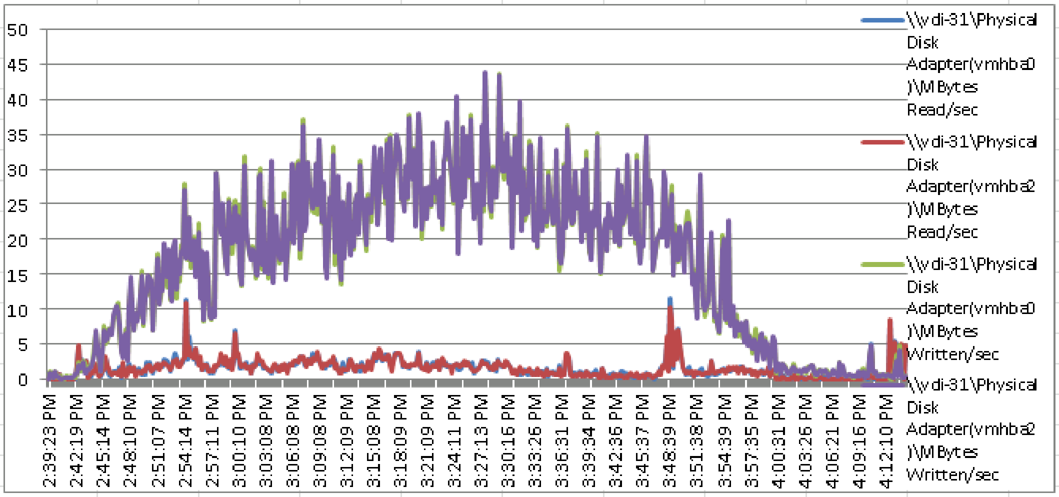

NETAPP AFF A300 Storage Charts for 5000 Users Scale Test

Figure 216 AFF A300 5000 Users Mixed Workload Scale Test | Total IOPS

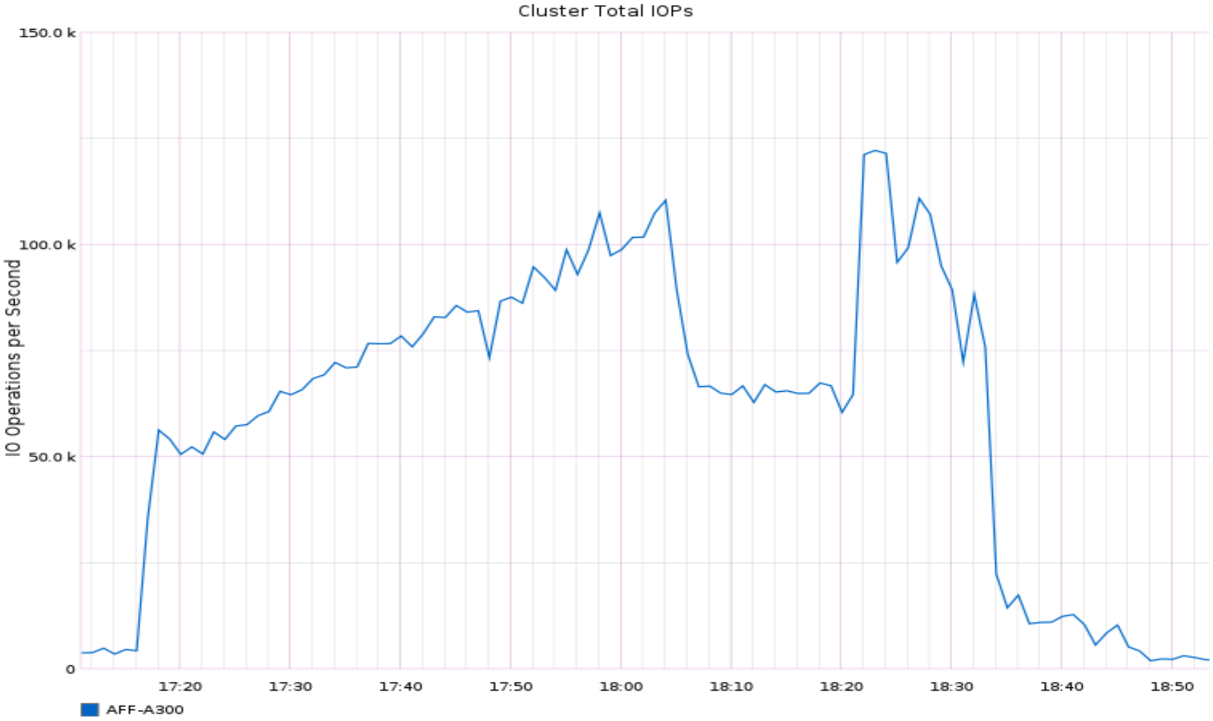

Figure 217 AFF A300 5000 Users Mixed Workload Scale Test | Total Throughput

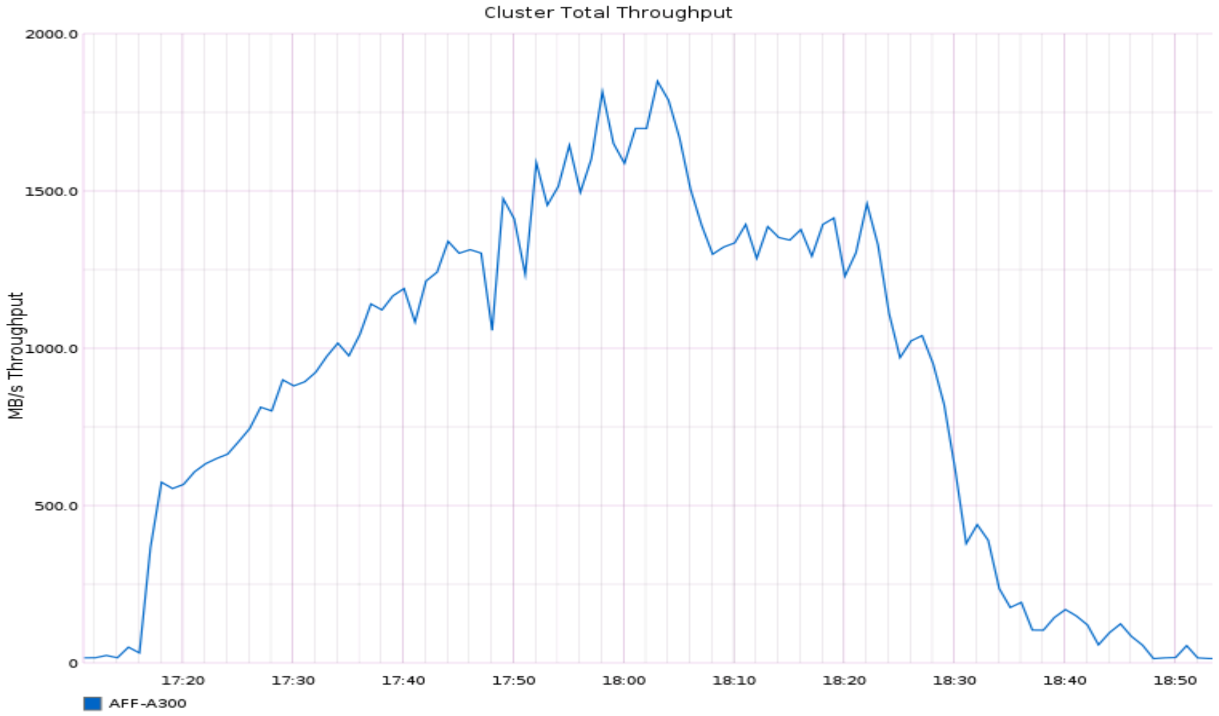

Figure 218 AFF A300 5000 Users Mixed Workload Scale Test | Controller Average Latency

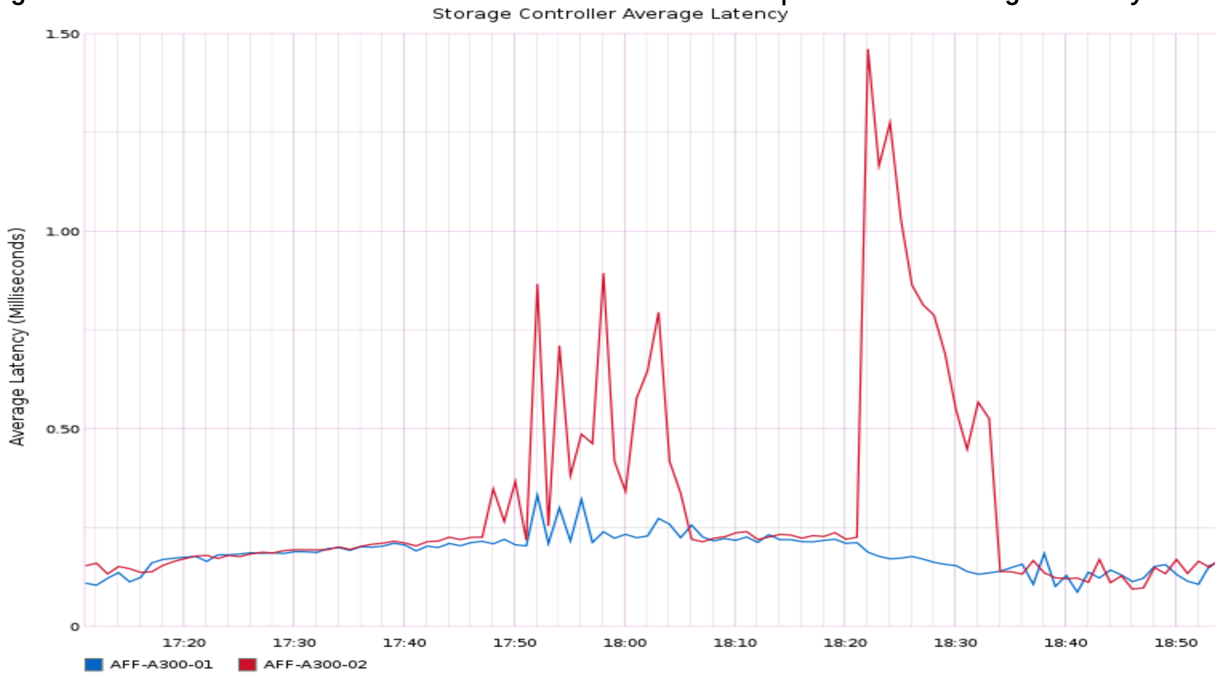

Figure 219 AFF A300 5000 Users Mixed Workload Scale Test | Controller Read / Write Latency

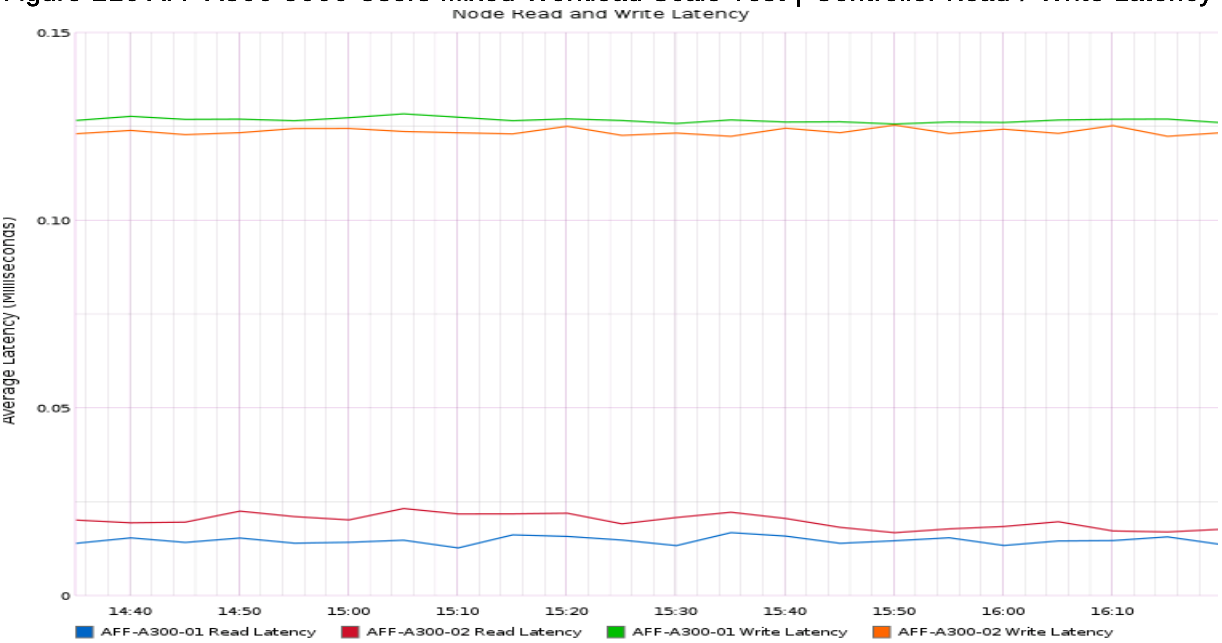

Figure 220 AFF A300 5000 Users Mixed Workload Scale Test | Controller Read Latency

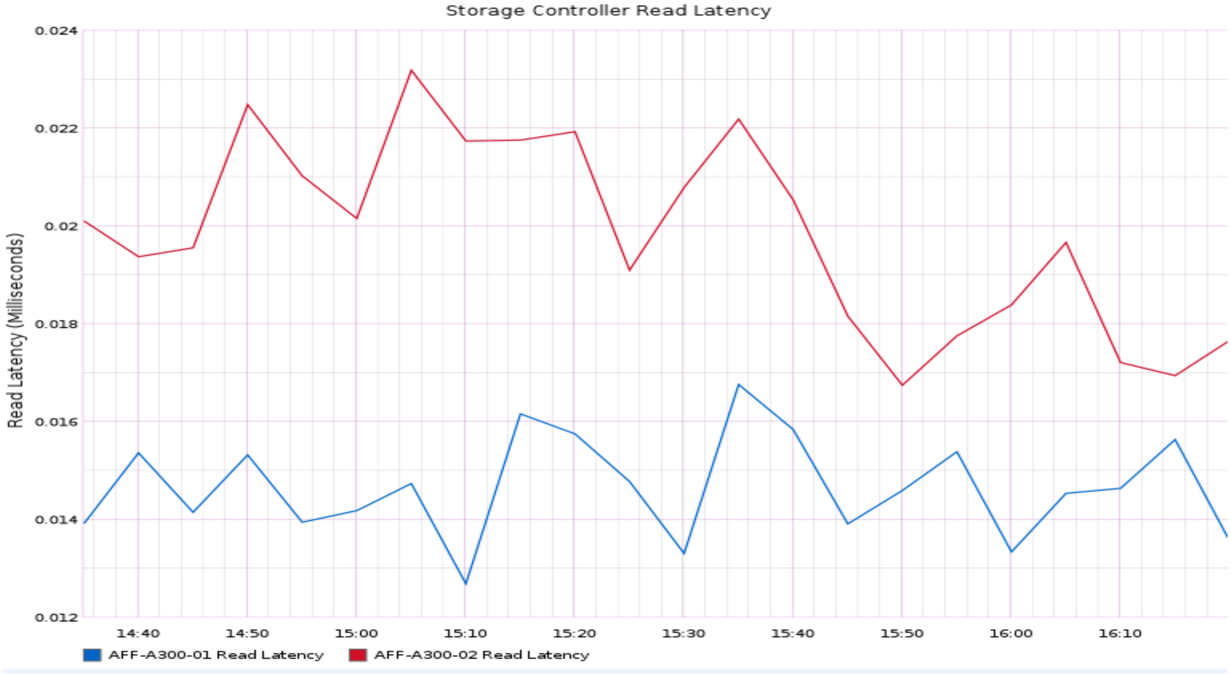

Figure 221 AFF A300 5000 Users Mixed Workload Scale Test | Controller Write Latency

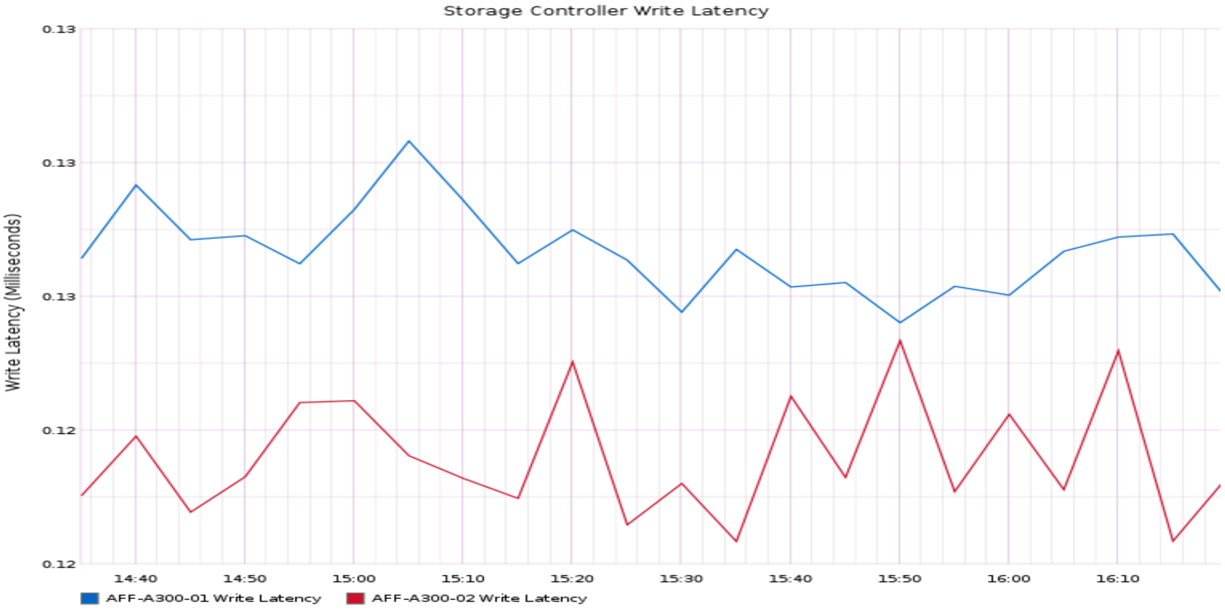

Figure 222 AFF A300 5000 Users Mixed Workload Scale Test | Controller CPU Headroom Util. %

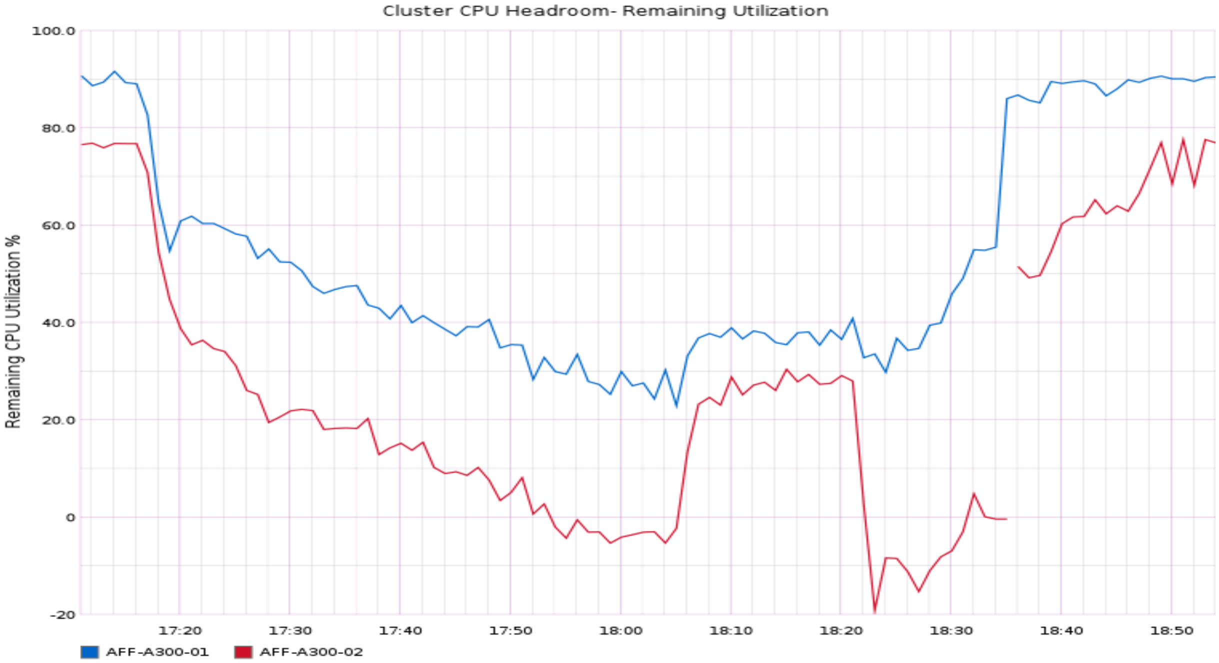

Figure 223 AFF A300 5000 Users Mixed Workload Scale Test | Controller Average Processor Busy

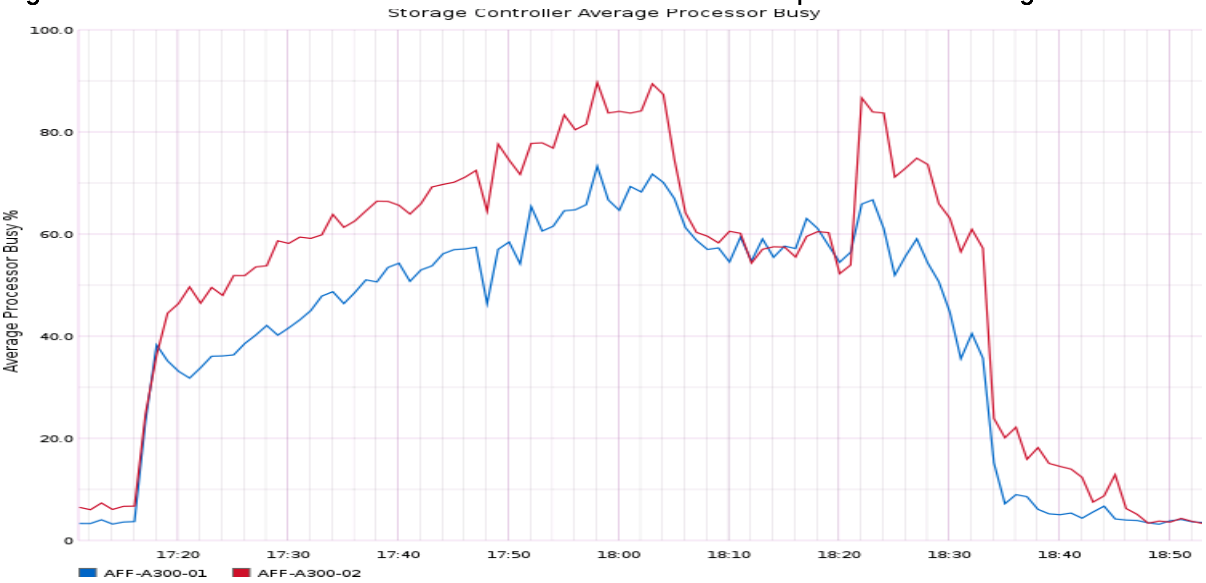
Storage Controller Average Processor Busy
Average Processor Busy %
17:20
17:30
17:40
17:50
18:00
18:10
18:20
18:30
18:40
18:50
AFF-A300-01
AFF-A300-02