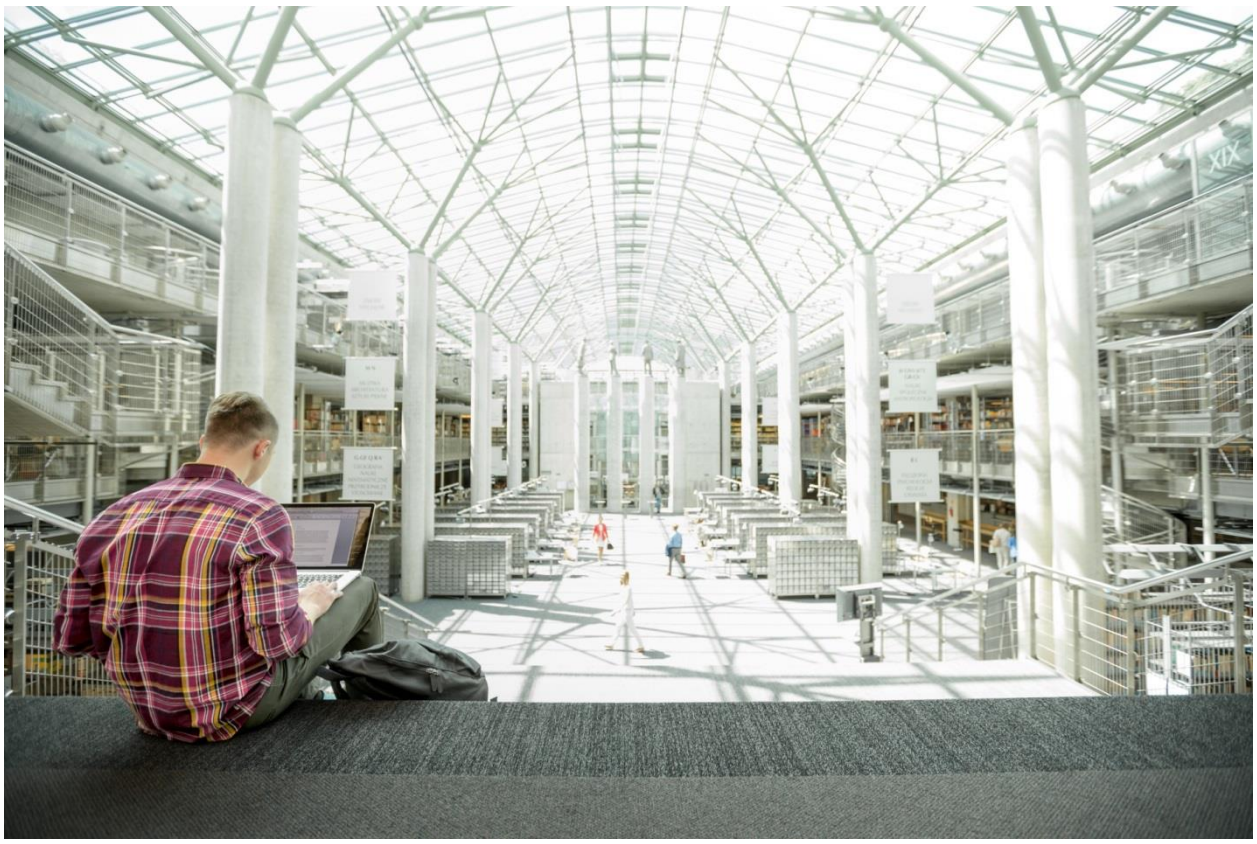

FlashStack with Cisco UCS and Pure Storage FlashArray//m for 5000 VMware Horizon View 6.2 Users

Cisco UCS B200 M4 Blade Servers with Pure Storage FlashArray//m50 Array on VMware Horizon View 6.2 and ESXi 6.0

Last Updated: January 29, 2019

CISCO VALIDATED DESIGN

About Cisco Validated Designs

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

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

Table of Contents

Executive Summary

Cisco® Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco and Pure Storage have partnered to deliver this document, which serves as a specific step by step guide for implementing this solution. This Cisco Validated Design provides an efficient architectural design that is based on customer requirements. The solution that follows is a validated approach for deploying Cisco and Pure Storage technologies as a shared, high performance, resilient, virtual desktop infrastructure.

This document provides architecture reference and design guide for up to 5000 seat mixed workload on Cisco UCS and Pure Storage FlashArray//m with VMware Horizon 6.2 RDS server-based sessions and Linked Clone Windows 7 virtual desktops on vSphere 6. The solution is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and Pure Storage all flash array.

This solution is 100 percent virtualized on Cisco UCS B200 M4 blade server booting via fibre channel SAN from Pure Storage M50 storage array running VMware vSphere 6.0 U1 hypervisor. The virtual desktops are configured with VMware Horizon with View 6.2 which now incorporated both tradition persistent and non-persistent virtual Windows 7/8/10 desktops, hosted applications and Remote desktop service (RDS) server 2012 R2 based desktops provides unparalleled scale and management simplicity. VMware Horizon with View pool linked-clone floating assignment Windows 7 desktops (3550) and RDS server based desktop sessions (1450) provisioned desktops on Pure Storage. Where applicable the document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

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

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

Solution Overview

Introduction

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

Audience

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

Purpose of this Document

This document provides a step by step design, configuration and implementation guide for the Cisco Validated Design for a large scale VMware Horizon 6.2 mixed workload solution with Pure Storage FlashArray//m50, Cisco UCS Blade Servers, Cisco Nexus 9000 series ethernet switches and Cisco MDS 9000 series fibre channel switches.

What's New?

This is the first Cisco Validated Design with Pure Storage. It incorporates the following features:

- Validation of Cisco Nexus 9000 with a Pure Storage all flash storage array
- Validation of Cisco MDS 9000 with a Pure Storage all flash storage array
- Support for the Cisco UCS 3.1(1) release and Cisco UCS B200-M4 servers
- Support for the latest release of Pure Storage FlashArray//m hardware and Purity Operating Environment 4.5.5 and 4.6.8.
- A Fibre Channel storage design supporting SAN LUNs
- Cisco Nexus 1000v distributed virtual switch technology
- Cisco UCS Inband KVM Access
- Cisco UCS vMedia client for vSphere Installation
- Cisco UCS Firmware Auto Sync Server policy
- VMware vSphere 6.0 Hypervisor
- VMware Horizon 6.2 VDI Linked Clones and RDSH shared server sessions

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

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

The use cases include:

- Enterprise data Center
- Service Provider Data Center

Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both VMware Horizon Microsoft Windows 7 virtual desktops and VMware Horizon RDSH server desktop sessions based on Microsoft Server 2012 R2. The mixed workload solution includes Pure Storage FlashArray//m storage array, Cisco Nexus® and MDS networking, the Cisco Unified Computing System (Cisco UCS®), 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.

One benefit of the Cisco Validated Design architecture is the ability to customize the environment to suit a customer's requirements. A Cisco Validated Design can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a Cisco Validated Design unit) and out (adding more Cisco Validated Design units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of a fibre channel storage solution. Figure 2 shows the VMware Horizon 6.2 on vSphere 6 built on Cisco Validated Design components and the network connections for a configuration with fibre channel-based storage. This design uses the Cisco Nexus 9000, Cisco MDS 9000, Cisco UCS B-Series blade servers and the Pure Storage FlashArray//m family of storage controllers connected in a highly available modular design. This 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.

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.

Figure 1 Cisco Data Center Collaboration

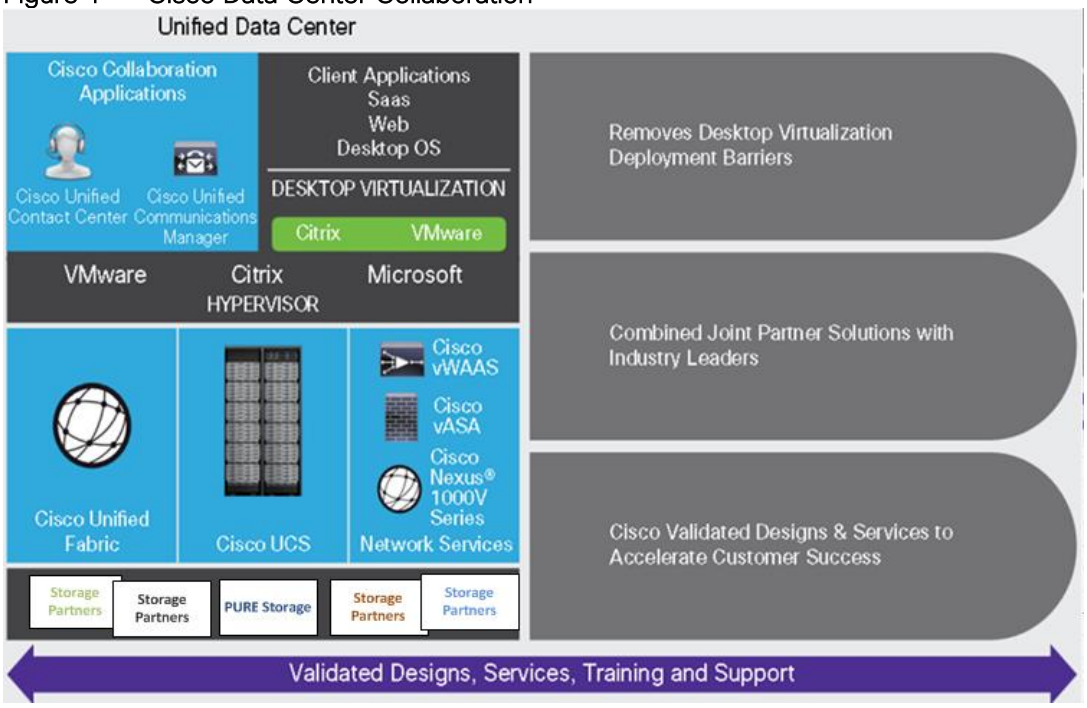

Some of the key drivers for desktop virtualization are increased data security and reduced TCO through increased control and reduced management costs.

Cisco Desktop Virtualization Focus

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

Simplified

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

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

Simplification also leads to more successful desktop virtualization implementation. Cisco and its technology partners like VMware Technologies and Pure have developed integrated, validated architectures, including predefined converged architecture infrastructure packages such as FlashStack. Cisco Desktop Virtualization Solutions have been tested with all the leading hypervisors, including VMware vSphere, Citrix XenServer, and Microsoft Hyper-V.

Secure

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

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

Scalable

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

Cisco UCS servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 1 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco storage partners Pure help maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs based on VMware Horizon View, Cisco UCS, and Pure joint solutions have demonstrated scalability and performance, with up to 5000 desktops up and running in 30 minutes.

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

Savings and Success

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

The simplified deployment of Cisco UCS for desktop virtualization accelerates the time to productivity and enhances business agility. IT staff and end users are more productive more quickly, and the business can respond to new opportunities quickly by deploying virtual desktops whenever and wherever they are needed.

The high-performance Cisco systems and network deliver a near-native end-user experience, allowing users to be productive anytime and anywhere.

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

Physical Topology

Figure 2 illustrates the physical architecture.

Figure 2 Architecture

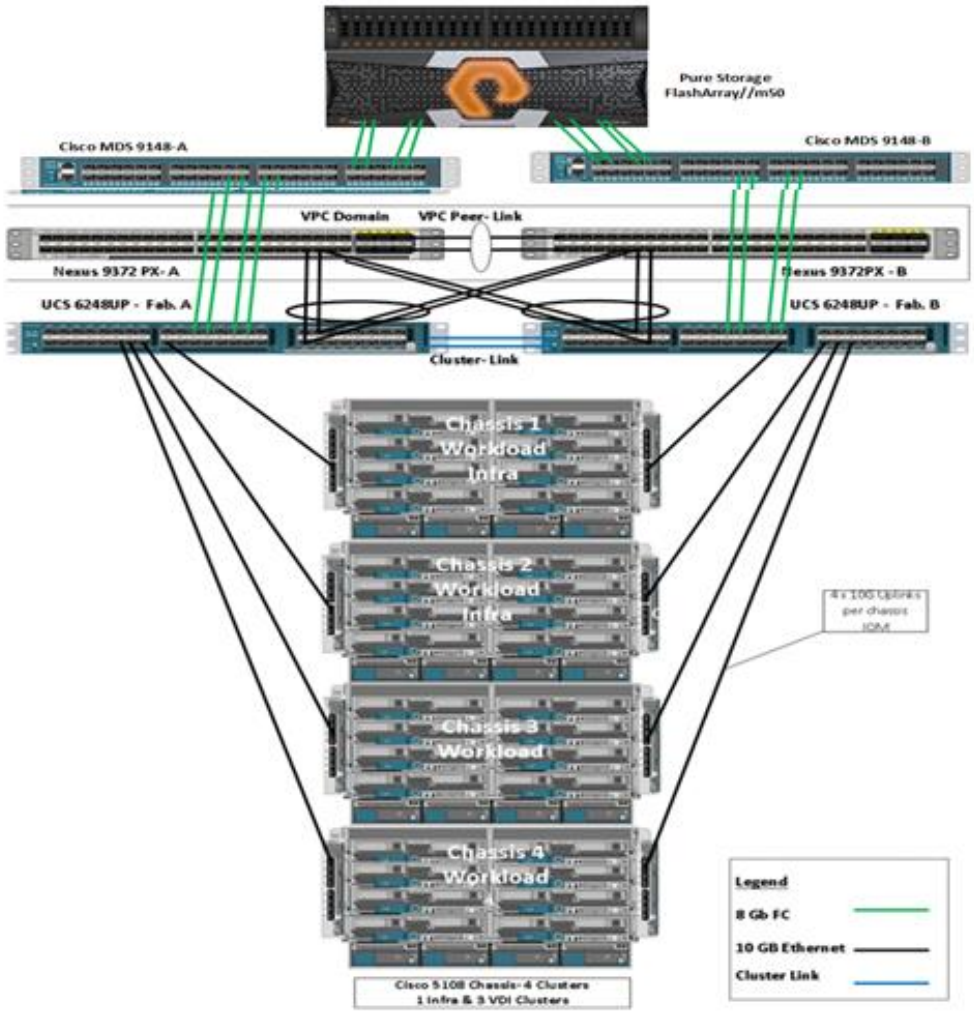

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches

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

For desktop virtualization, the deployment includes VMware Horizon 6.2 running on VMware vSphere 6. The design is intended to provide a large scale building block for both RDSH and Linked Clone desktops in the following ratio:

- 1450 Horizon 6.2 RDSH server desktop sessions
- 3550 Horizon 6.2 Windows 7 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 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 document provides details for configuring a fully redundant, highly available configuration for a Cisco Validated Design for a 5000 seat mixed workload virtual desktop solution with Pure Storage. Configuration guidelines are provided that refer to which redundant component is being configured with each step. For example, node01 and node02 are used to identify the two Pure Storage controllers that are provisioned with this document, Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured and Cisco MDS A or Cisco MDS B identifies the pair of Cisco MDS switches that are configured. The Cisco UCS 6248UP Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, VM-Host-RDSH-01, VM-Host-VDI-01 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

Solution Components

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

Cisco Unified Computing System

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

Cisco UCS is a next-generation data center platform that unites computing, networking, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency; lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Cisco Unified Computing System Components

The main components of Cisco UCS are:

- Compute: The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® processor E5-2600/4600 v3 and E7-2800 v3 family CPUs.

- Network: The system is integrated on a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables needed, and by decreasing the power and cooling requirements.

- Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.

- Storage access: The system provides consolidated access to local storage, SAN storage, and network-attached storage (NAS) over the unified fabric. With storage access unified, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Small Computer System Interface over IP (iSCSI) protocols. This capability provides customers with choice for storage access and investment protection. In addition, server administrators can preassign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management and helping increase productivity.

- Management: Cisco UCS uniquely integrates all system components, enabling the entire solution to be managed as a single entity by Cisco UCS Manager. The manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

Figure 3 Cisco Data Center Overview

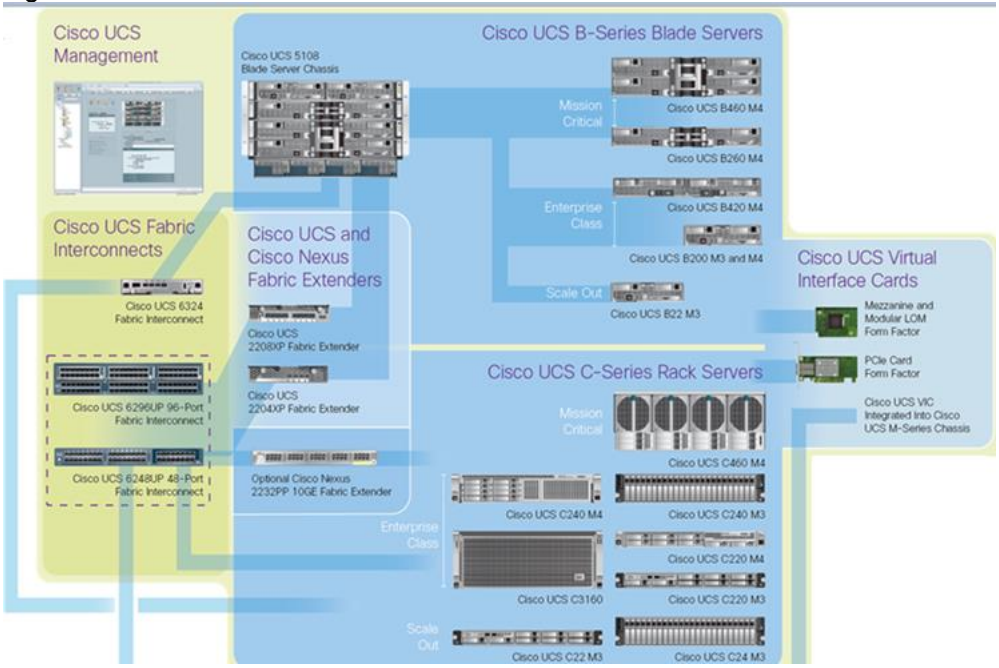

Cisco UCS is designed to deliver:

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

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

Cisco UCS Fabric Interconnect

The Cisco UCS 6200 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, FCoE, and Fibre Channel functions.

The fabric interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both LAN and SAN connectivity for all blades in the domain.

For networking, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, 1-terabit (Tb) switching capacity, and 160 Gbps of bandwidth per chassis, independent of packet size and enabled services. The product series supports Cisco low-latency, lossless, 10 Gigabit Ethernet unified network fabric capabilities, increasing the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnects support multiple traffic classes over a lossless Ethernet fabric, from the blade server through the interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Figure 4 Cisco UCS 6200 Series Fabric Interconnect

Cisco UCS B200 M4 Blade Server

The Cisco UCS B200 M4 Blade Server (Figures 5 and 6) is a density-optimized, half-width blade server that supports two CPU sockets for Intel Xeon processor E5-2600 v3 series CPUs and up to 24 DDR4 DIMMs. It supports one modular LAN-on-motherboard (LOM) dedicated slot for a Cisco virtual interface card (VIC) and one mezzanine adapter. In additions, the Cisco UCS B200 M4 supports an optional storage module that accommodates up to two SAS or SATA hard disk drives (HDDs) or solid-state disk (SSD) drives. You can install up to eight Cisco UCS B200 M4 servers in a chassis, mixing them with other models of Cisco UCS blade servers in the chassis if desired. Latest features of Cisco UCS Virtual Interface Cards (VICs)

Figure 5 Cisco UCS B200 M4 Front View

Figure 6 Cisco UCS B200 M4 Back view

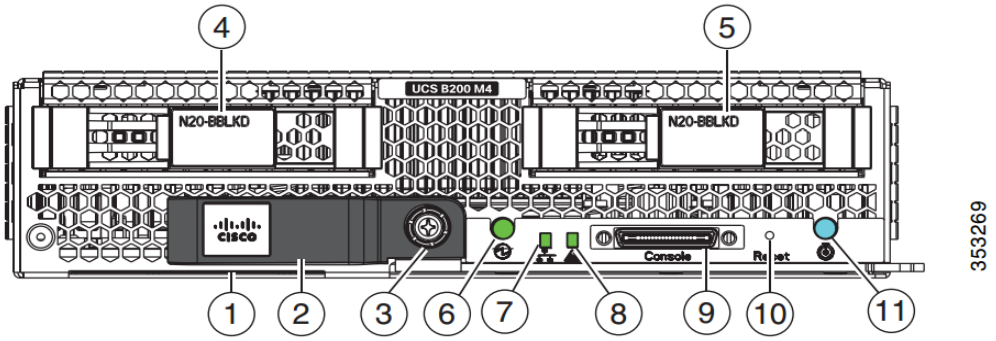

s

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

Cisco UCS combines Cisco UCS B-Series Blade Servers and C-Series Rack Servers with networking and storage access into a single converged system with simplified management, greater cost efficiency and agility, and increased visibility and control. The Cisco UCS B200 M4 Blade Server is one of the newest servers in the Cisco UCS portfolio.

The Cisco UCS B200 M4 delivers performance, flexibility, and optimization for data centers and remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads ranging from web infrastructure to distributed databases. The Cisco UCS B200 M4 can quickly deploy stateless physical and virtual workloads with the programmable ease of use of the Cisco UCS Manager software and simplified server access with Cisco® Single Connect technology. Based on the Intel Xeon processor E5-2600 v3 product family, it offers up to 768 GB of memory using 32-GB DIMMs, up to two disk drives, and up to 80 Gbps of I/O throughput. The Cisco UCS B200 M4 offers exceptional levels of performance, flexibility, and I/O throughput to run your most demanding applications.

In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches, NICs, and HBAs in each blade server chassis. With a larger power budget per blade server, it provides uncompromised expandability and capabilities, as in the new Cisco UCS B200 M4 server with its leading memory-slot capacity and drive capacity.

Cisco UCS B200 M4 Features

The Cisco UCS B200 M4 provides:

- Up to two multicore Intel Xeon processor E5-2600 v3 series CPUs for up to 36 processing cores

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

The Cisco UCS B200 M4 server is a half-width blade. Up to eight can reside in the 6-rack-unit (6RU) Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit of blade chassis in the industry.

Cisco UCS B200 M4 Benefits

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

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

Single-Instance and Distributed Databases

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

The Cisco UCS B200 M4 Blade Server delivers:

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

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

Cisco UCS VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 7) is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

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

Figure 7 Cisco UCS VIC 1340

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

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

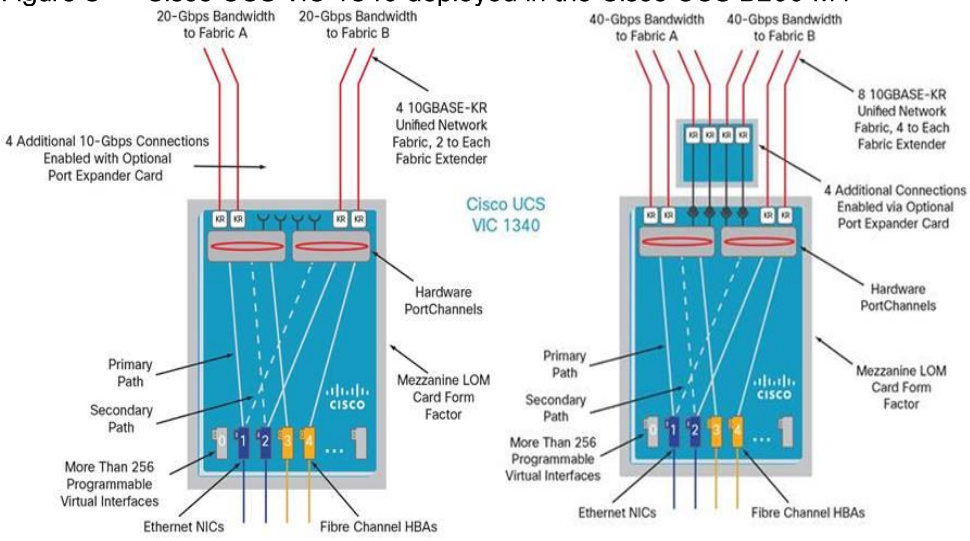

Cisco Switching

Cisco Nexus 9372PX Switches

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

Architectural Flexibility

- Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Leaf node support for Cisco ACI architecture is provided in the roadmap
- Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support

Feature Rich

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

Highly Available and Efficient Design

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

Simplified Operations

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

Investment Protection

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

- 1.44 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10-Gbps SFP+ ports
- 6 fixed 40-Gbps QSFP+ for uplink connectivity that can be turned into 10 Gb ports through a QSFP to SFP or SFP+ Adapter (QSA)
- Latency of 1 to 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 2+1 redundant fan tray

Figure 9 Cisco Nexus 9372PX Switch

Cisco Nexus 1000V Distributed Virtual Switch

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

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

Important Differentiators for the Cisco Nexus 1000V for VMware vSphere

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

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

Virtual Networking Services

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

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

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

Cisco MDS 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.

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
- Detection, mitigation, and recovery at one-minute intervals

Specifications at-a-Glance

Performance and Port Configuration

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

Advanced Functions

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

Hypervisor and Desktop Broker

This Cisco Validated Design includes VMware vSphere 6 and VMware Horizon 6.2.

VMware vSphere 6.0

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

VMware vSphere 6.0 introduces many enhancements to vSphere Hypervisor, VMware virtual machines, vCenter Server, virtual storage, and virtual networking, further extending the core capabilities of the vSphere platform.

VMware ESXi 6.0 Hypervisor

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

- Scalability Improvements

ESXi 6.0 dramatically increases the scalability of the platform. With vSphere Hypervisor 6.0, clusters can scale to as many as 64 hosts, up from 32 in previous releases. With 64 hosts in a cluster, vSphere 6.0 can support 8000 virtual machines in a single cluster. This capability enables greater consolidation ratios, more efficient

use of VMware vSphere Distributed Resource Scheduler (DRS), and fewer clusters that must be separately managed. Each vSphere Hypervisor 6.0 instance can support up to 480 logical CPUs, 12 terabytes (TB) of RAM, and 1024 virtual machines. By using the newest hardware advances, ESXi 6.0 enables the virtualization of applications that previously had been thought to be non-virtualizable.

Security Enhancements

- ESXi 6.0 offers these security enhancements:

  - Account management: ESXi 6.0 enables management of local accounts on the ESXi server using new ESXi CLI commands. The capability to add, list, remove, and modify accounts across all hosts in a cluster can be centrally managed using a vCenter Server system. Previously, the account and permission management functions for ESXi hosts were available only for direct host connections. The setup, removal, and listing of local permissions on ESXi servers can also be centrally managed.

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

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

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

  - Flexible lockdown modes: Prior to vSphere 6.0, only one lockdown mode was available. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.0, two lockdown modes are available:

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

    - In strict lockdown mode, the DCUI is stopped.

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

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

VMware Horizon View

VMware Horizon View desktop virtualization solutions built on a unified architecture so they're simple to manage and flexible enough to meet the needs of all your organization's users. 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 View 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 View and 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 View

VMware Horizon 6 version 6.2 provides the following new features and enhancements:

- 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 View 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 View 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 View 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.

- View Administrator Enhancements

  - View Administrator shows additional licensing information, including license key, named user and concurrent connection user count.

  - Pool creation is streamlined by letting View administrators clone existing pools.

- Horizon 6 Interoperability with vSphere 6 Update 1

- Horizon 6 for Linux Desktop Enhancements

  - Several new features are supported on Horizon 6 for Linux desktops, including NVIDIA GRID vGPU, vSGA, RHEL 7.1 and Ubuntu 14.04 guest operating systems, and View Agent installation of JRE 8 with no user steps required.

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 View, you can create desktop and application pools to give users remote access to virtual machine-based desktops, session-based desktops, physical computers, and applications. View 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.

Pure Storage FlashArray//m50

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

What is FlashStack?

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

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

Figure 10 FlashStack Converged Infrastructure (CI)

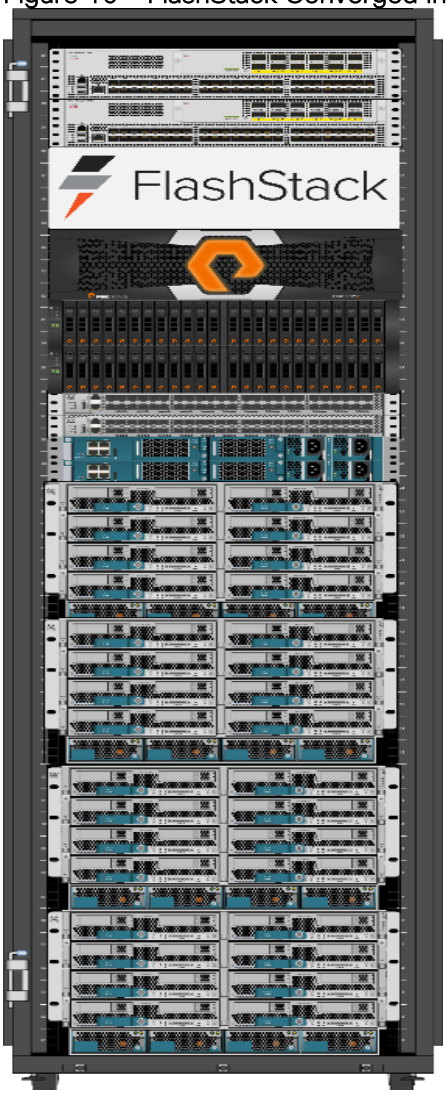

FlashStack CI is available from qualified FlashStack Partners who help to provide an excellent converged infrastructure ownership experience. FlashStack Partners have the knowledge and experience necessary to help streamline the sizing, procurement, and delivery of your entire system.

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

Why FlashStack?

The following lists the benefits of FlashStack:

- Consistent Performance and Scalability
  - Consistent sub-millisecond latency with 100% flash storage
  - Consolidate 100's of enterprise-class applications in a single rack
  - Scales easily, without disruption

  - Continuous growth through multiple FlashStack CI deployments
- Operational Simplicity
  - Fully tested, validated, and documented for rapid deployment
  - Reduced management complexity
  - Auto-aligned 512B architecture removes storage alignment issues
  - No storage tuning or tiers necessary
- Lowest TCO
  - Dramatic savings in power, cooling, and space with 100% Flash
  - Industry leading data reduction
  - Free FlashArray controller upgrades every three years with Forever Flash™
- Enterprise Grade Resiliency
  - Highly available architecture with no single point of failure
  - Non-disruptive operations with no downtime
  - Upgrade and expand without downtime or performance loss
  - Native data protection: snapshots and replication
- Suitable for even large resource-intensive workloads such as real-time analytics or heavy transactional databases

Benefits of Pure Storage FlashArray//m Series

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

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

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

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

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

The Pure Storage FlashArray is ideal for:

- Accelerating Databases and Applications Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.
- Virtualizing and Consolidating Workloads Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.

- Delivering the Ultimate Virtual Desktop Experience Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance with simple management for under $50/desktop.

- Protecting and Recovering Vital Data Assets Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

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

- Consistent Performance FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

- Less Cost than Disk Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure. With VDI workloads data reduction is typically > 10:1.

- Mission-Critical Resiliency FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

- Disaster Recovery Built-In FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

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

FlashArray//m Specifications

Figure 11 Pure Storage FlashArray//m Portfolio

Table 1 Pure Storage FlashArray//m Series Controller Specifications

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

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

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

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

Purity Operating Environment

- Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashArray//m.
- Storage Software Built for Flash—The FlashCare technology virtualizes the entire pool of flash within the FlashArray, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.
- Granular and Adaptive—Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.
- Best Data Reduction Available—FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.
- Highly Available and Resilient—FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashArray during any failure or maintenance event.
- Backup and Disaster Recovery Built In—FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss

locally and globally. All FlashProtect services are fully-integrated in the FlashArray and leverage the native data reduction capabilities.

Pure1

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

Engineered to Stay Evergreen

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

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

Architecture and Design Considerations for Desktop Virtualization

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

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

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

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

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

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

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

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

- Hosted Shared Desktop: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2012, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Hosted Virtual Desktop: A hosted virtual desktop is a virtual desktop running on a virtualization layer (ESX). The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.

- Published Applications: Published applications run entirely on the VMware RDSH Session Hosts 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 Horizon View Virtual Desktops and Remote Desktop Services sever 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, like SalesForce.com. This application and data analysis is beyond the scope of this Cisco Validated Design, but should not be omitted from the planning process. There are a variety of third party tools available to assist organizations with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

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

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

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

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

- Are the required skill sets to execute the VDI project available? Can we hire or contract for them?

- Do we have end user experience performance metrics identified for each desktop sub-group?

- How will we measure success or failure?

- What is the future implication of success or failure?

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

- What is the desktop OS planned? Windows 7, 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 7/8/10?

- How much memory per target desktop group desktop?

- Are there any rich media, Flash, or graphics-intensive workloads?
- What is the end point graphics processing capability?
- Will VMware RDSH for Remote Desktop Server Hosted Sessions used?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- Is user profile management (e.g., 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?

VMware Horizon with View Design Fundamentals

VMware Horizon View 6.2 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 View delivers a native touch-optimized experience with HDX high-definition performance, even over mobile networks.

Horizon View and RDSH Server Pools

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Desktop Pool.

In this CVD, linked clone VM provisioning relies on VMware View Composer aligning with VMware Horizon View Connection Server and vCenter Server components.

Microsoft Windows Server 2012 R2 RDSH VMs are created in vCenter using traditional cloning techniques. These VMs are created in a dedicated vCenter Cluster and are imported into a Horizon Farm. We used this technique to have more granular control over the RDSH farm VMs compared with an auto-provisioned farm.

From there Horizon Desktop and RDSH Pools are created to consume the linked clone or server desktop session resources. Figure 12 illustrates how users access desktops and applications through Unified Workspace catalogs and Single Sign-On.

Figure 12 Server OS and Desktop OS Machines Configured to Support RDS Hosted Shared Desktops and VDI Hosted Virtual Desktops

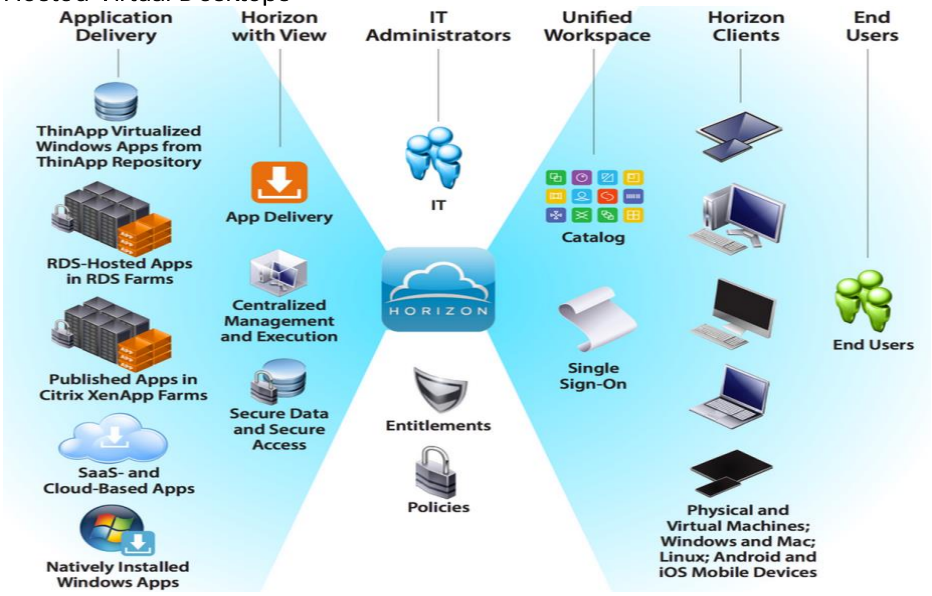

VMware View Composer

VMware Horizon View Composer is a feature in Horizon View 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 View Composer Overview

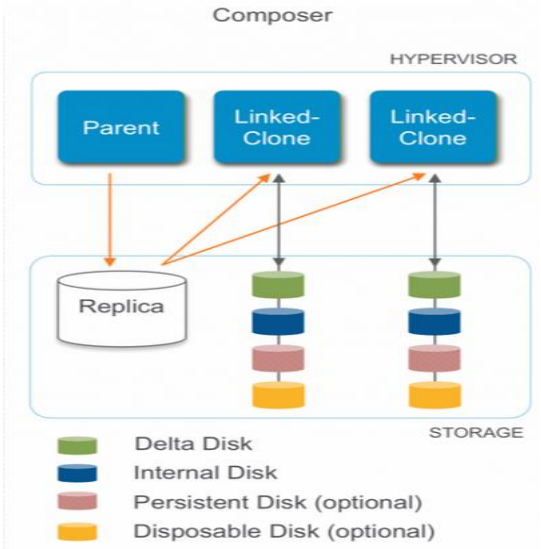

Figure 14 Example of VMware Horizon View Overview

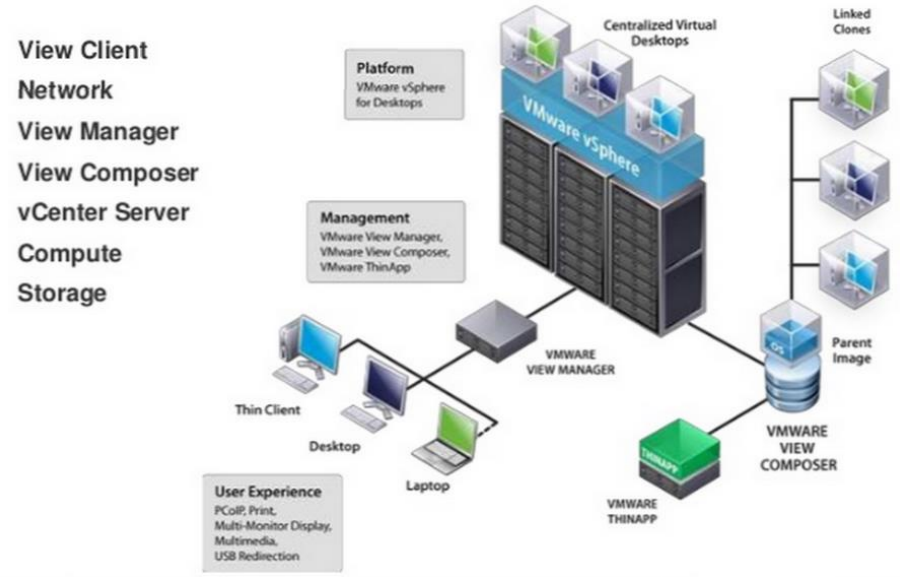

Multiple Site Configuration

If you have multiple regional sites, you can use any of the Load Balances Tools (Ex:- Big-IP Global Traffic Manager) to direct the user connections to the most appropriate site to deliver the desktops and application to users.

In Figure 15, The image depicting sites, a site was created in two data centers. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic. Two Cisco blade servers host the required infrastructure services (Domain Controllers, DNS, DHCP, Profile, SQL, VMware Horizon View Connection Servers, View Composer server and web servers).

Figure 15 Multisite Configuration Overview

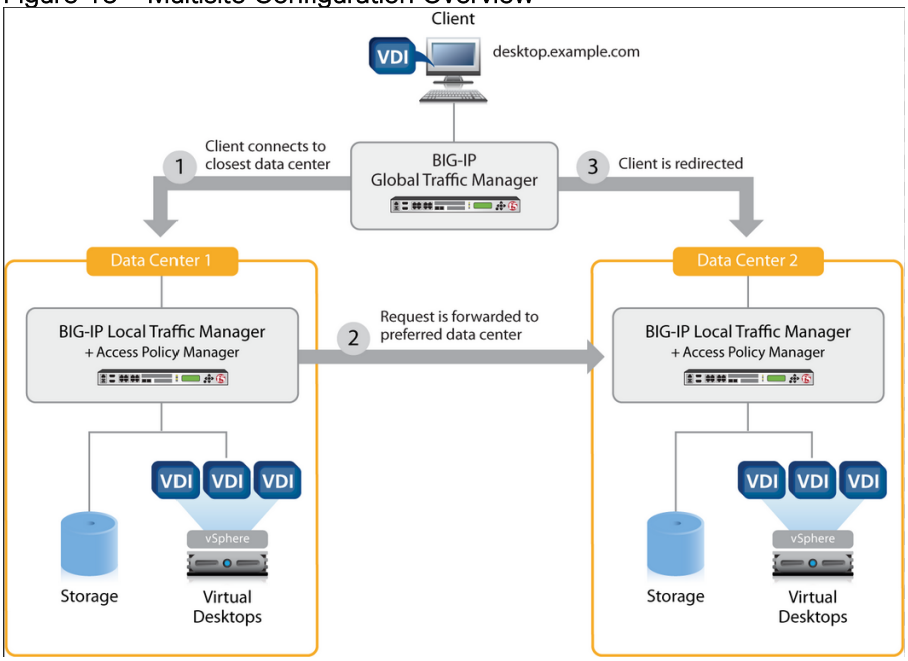

Based on the requirement and no of data centers or remote location, you can chose any of the available Load balancing software or tools accelerates the application performance, load balances servers, increases security,

and optimizes the user experience. In this example, two Big-IP Local Traffic Manager are used to provide a high availability configuration.

BIG-IP Local Traffic Manager has been shown as example for presentation purpose.

Designing a VMware Horizon View Environment for a Mixed Workload

With VMware View 6.2, the method you choose to provide applications or desktops to users depends on the types of applications and desktops you are hosting and available system resources, as well as the types of users and user experience you want to provide.

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

Desktop OS machines	You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition. Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications. Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines. Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.

Remote PC Access	You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the datacenter. Your users: Employees or contractors that have the option to work from home, but need access to specific software or data on their corporate desktops to perform their jobs remotely. Host: The same as Desktop OS machines. Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.

For this Cisco Validated Design, a mix of Remote Desktop Server Hosted sessions (RDSH) using RDS-based Server OS machines and Hosted Virtual Desktops (HVDs) using VDI-based Desktop OS machines were configured and tested. The following sections discuss design decisions relative to the VMware Horizon View deployment, including the CVD test environment.

High-Level Storage Architecture Design

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

For RDSH server shared desktops and the linked clone Windows 7 virtual desktops, the following recommendations are best practices for the delta disks, the internal disks, the disposable disks, user profiles, , user data, and application virtualization:

- User Profiles: To make sure that user profiles and settings are preserved, you can leverage the profile management server on the Pure storage Array. We did not deploy this feature in this project.
- User data: We recommend hosting user data either on CIFS home directories to preserve data upon VM reboot or redeploy.
- Monitoring and management. We recommend using VMware vCenter Operations Manager for View to provide monitoring and management of the solution.

Solution Hardware and Software

Products Deployed

The architecture deployed is highly modular. While each customer's environment might vary in its exact configuration, the reference architecture contained in this document once built, can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources within a Cisco UCS Domain) and out (adding additional Cisco UCS Domains and Pure Storage FlashArrays.) The solution includes Cisco networking, Cisco UCS and Pure Storage FlashArray//m storage, which efficiently fit into a single data center rack, including the access layer network switches.

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

- Software Deployed VMware Horizon View 6.2 Remote Desktop Server Hosted sessions(RDSH) on Pure Storage
- VMware Horizon View 6.2 non-persistent linked clone Microsoft Windows 7 Virtual Desktops (VDI) on Pure Storage
- Microsoft Windows Server 2012 for User Profile Manager
- Microsoft Windows 2012 Server for Login VSI Management Console and LVSI Share data servers to simulate real world VDI workload.
- VMware vSphere ESXi 6.0 Update 1 Hypervisor
- Microsoft Windows Server 2012 R2 Data Center Version for RDSH Servers & Microsoft Windows 7 32-bit virtual machine operating systems for VDI virtual machines
- Microsoft SQL Server 2012
- Cisco Nexus 1000V primary and secondary Virtual Supervisor Module
- VMware Horizon View 6.2 Connection Server and Replica Servers for redundancy and support 5000 seat scale
- VMware Horizon 6.2 View Composer Server

Hardware Deployed

The workload contains the following hardware as shown in Figure 2:

- Two Cisco Nexus 9372PX Layer 2 Access Switches
- Four Cisco UCS 5108 Blade Server Chassis with two UCS-IOM-2208XP IO Modules
- Two Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2660v3 2.6-GHz 10-core processors, 128GB RAM 2133-MHz, and VIC1340 mezzanine cards for the hosted infrastructure with N+1 server fault tolerance

- 30 Thirty Cisco UCS B200 M4 Blade servers with Intel Xeon E5-2680v3 2.5-GHz 12-core processors, 384GB RAM 1866-MHz, and VIC1340 mezzanine cards for the virtual desktop workloads with N+1 server fault tolerance

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

Software Deployed

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

Table 2 Software and Firmware Versions

Layer	Device	Image
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M4	2.2(6b)
Cisco	Cisco eNIC	2.3.0.7
Cisco	Cisco fNIC	1.6.0.25
Cisco	Cisco VIC 1340	
Network	Cisco Nexus 9000 NX-OS	7.0(3)I1(1a)
Network	Cisco MDS 9000	6.2.17
Nexus 1000V	Cisco Nexus 1000V	5.2(1)SV3(1.5a)
Nexus 1000V	Virtual Switch Update Manager	2.0
VMware	vSphere Client vCenter Server Appliance	6.0.0.3016447 6.0.0.3018523
VMware	vSphere ESXi6.0. Update 1a	6.0.0.3073146
View Connection Servers	VMware Horizon View Connection Servers	6.2.0.3005368
View Composer	VMware View Composer	6.2.0.3001314
Storage	Pure Storage FlashArray//m50	Purity 4.6.8

Logical Architecture

The logical architecture of this solution is designed to support up to 5000 users within four Cisco UCS 5108 Blade server chassis containing 32 blades, which provides physical redundancy for the blade servers for each workload type.

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. Table 3 through Table 5 lists the information you need to configure your environment.

Table 3 Server, Location, and Purpose

Server Name	Location	Purpose
C1-Blade 8 C2-Blade 8	Physical – Chassis 1, 2	ESXi 6.0 Hosts Infrastructure VMs WIN 2012 R2 Servers vCenter Server Appliance (VCSA), Cisco Nexus 1000v Virtual Supervisor Module (VSM) Primary & Secondary & Cisco Virtual Switch Update Manager (VSUM), View Connection Server, View Replica Servers, View Composer Server. SQL Server, Domain Controllers, (* *File Share Server, Workload monitoring servers for workload users profiles & testing* purpose)
C1-Blade 1-7,	Physical – Chassis 1	ESXi 6.0 Hosts 56x WIN 2012 R2 RDSH Servers
C2-Blade 1-7	Physical – Chassis 2	ESXi 6.0 Hosts Linked Clone WIN 7 VDI VMs (Non-Persistent, Pool 1 &2)
C3-Blade 1-7	Physical – Chassis 3,4	ESXi 6.0 Hosts Linked Clone WIN7 VDI VMs (Non-Persistent, Pool 1 & 2)
C3-Blade 8 C4-Blade 8	Physical – Chassis 3,4	ESXi 6.0 Hosts 16x WIN 2012 R2 RDSH Servers
C4-Blade 1-7	Physical –Chassis 3,4	ESXi 6.0 Hosts Linked Clone WIN7 VDI VMs (Non-Persistent, Pool 1 & 2)

VLANs

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

Table 4 VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	160	VLAN for in-band management interfaces
Infra-Mgmt	161	VLAN for Virtual Infrastructure
VDI	162	VLAN for VDI Traffic
vMotion	166	VLAN for VMware vMotion
OB-Mgmt	164	VLAN for out-of-band management interfaces

Table 5 VASNs Configured in this Study

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

VMware Clusters

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

- VDI Cluster Pure Storage Data Center with Cisco UCS
  - Infrastructure Cluster: Infra VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware Connection Servers, VMware Replica Servers, View Composer Server, and Nexus 1000v Virtual Supervisor Modules (VSMs,) etc.)
  - RDSH Cluster: VMware RDSH VMs (Windows Server 2012 R2)
  - VDI Cluster 1: VMware Linked Clones VDI VM Pools (Windows 7 SP1 32-bit)*

VMware Supports 32 Hosts in single VMware cluster with latest Horizon View version.

  - VSI Launchers Cluster
  - Launcher Cluster 1 and 2: Login VSI Cluster (The Login VSI launcher infrastructure was connected using the same set of switches and vCenter instance, but was hosted on separate local storage and servers.)

Figure 16 VMware vSphere Clusters on vSphere Web GUI

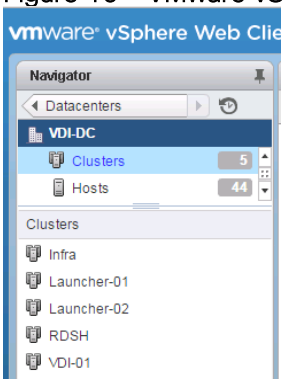

Validation

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution. Figure 17 illustrates the configuration topology for this solution.

Figure 17 Configuration Topology for Scalable VMware Horizon View 6.2 Mixed Workload

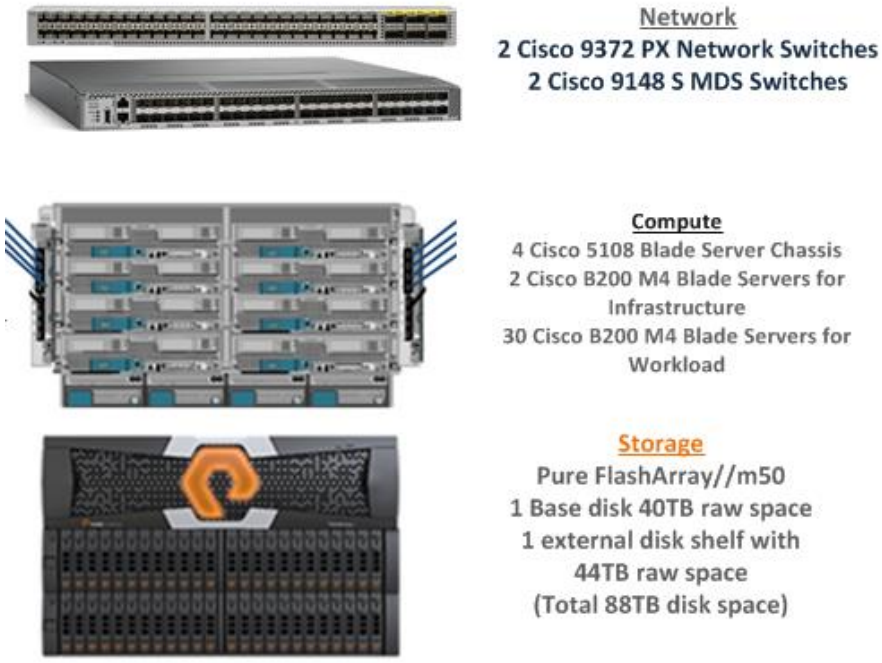

Cisco UCS Compute Platform

The following subsections detail the physical connectivity configuration of the VMware Horizon View6.2 environment.

Physical Infrastructure

Solution Cabling

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

The tables in this section contain the details for the prescribed and supported configuration of the Pure Storage FlashArray//m 50 to the Cisco 6248UP Fabric Interconnects via Cisco MDS 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 17 shows a cabling diagram for a VMware Horizon View configuration using the Cisco Nexus 9000 and Pure Storage FlashArray//m50. The Pure Storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves.

Table 6 Cisco Nexus 9372-Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	1GbE	Pure Storage FlashArray//m50 (Console Mgmt.)	e0e
	Eth1/29	10GbE	Cisco UCS fabric interconnect A	Eth1/29
	Eth1/30	10GbE	Cisco UCS fabric interconnect A	Eth1/30
	Eth1/31	10GbE	Cisco UCS fabric interconnect B	Eth1/29
	Eth1/32	10GbE	Cisco UCS fabric interconnect B	Eth1/30
	Eth1/47	40GbE	Cisco Nexus 9372 B	Eth1/47
	Eth1/48	40GbE	Cisco Nexus 9372 B	Eth1/48
	MGMT0	GbE	GbE management switch	Any

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

Table 7 Cisco Nexus 9372-B Cabling Information

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

Table 8 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/29	10GbE	Cisco Nexus 9372 A	Eth1/29
	Eth1/30	10GbE	Cisco Nexus 9372 A	Eth1/30
	Eth1/31	10GbE	Cisco Nexus 9372 B	Eth1/29
	Eth1/32	10 GbE	Cisco Nexus 9372 B	Eth 1/30
	Eth1/1-1/8	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 1-2	IOM 1-4
	Eth1/11-1/14	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 3	IOM 1-4
	Eth 1/17-1/20	10GbE	UCS 5108 Blade Chassis IOM-A, Chassis 4	IOM 1-4
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2
	FC 2/13	8Gb FC	Cisco MDS 9148S-A	FC 1/5
	FC 2/14	8Gb FC	Cisco MDS 9148S-A	FC 1/6
	FC 2/15	8Gb FC	Cisco MDS 9148S-A	FC 1/7
	FC 2/16	8Gb FC	Cisco MDS 9148S-A	FC 1/8

Table 9 Cisco UCS Fabric Interconnect B Cabling Information

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

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

Figure 18 Cable Connectivity Between Cisco Nexus 9372PX, Cisco UCS 6248UP Fabric Interconnects and Cisco 2208 IO Modules in Cisco 5108AC Blade Chassis

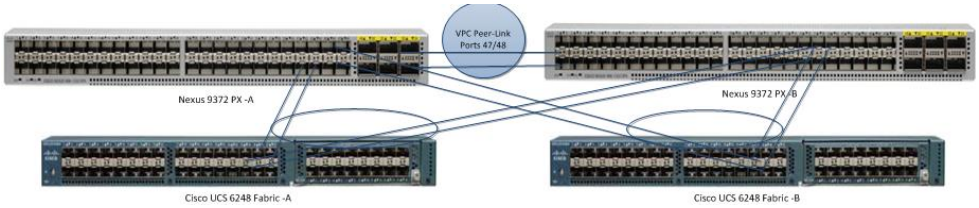

Figure 18 show cable connectivity between the Cisco MDS 9148S and the Cisco 6248 Fabric Interconnects and the Pure Storage FlashArray//m50.

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

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

Table 10 Cisco MDS 9148S A Cabling

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

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

Table 11 Cisco MDS 9148S B Cabling

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

Figure 19 FlashArray//m50 Controller-A and B Connection to MDS 9148S Switches using VSAN 20 for Fabric A and VSAN 30 Configured for Fabric B Side

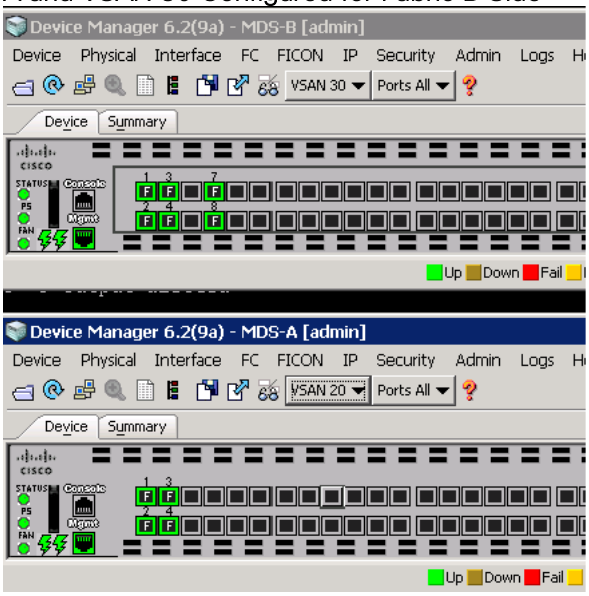

Figure 20 Fibre Channel Cable Connectivity from Pure FlashArray//m50 to Cisco MDS 9148S to Cisco 6248 Fabric Interconnects

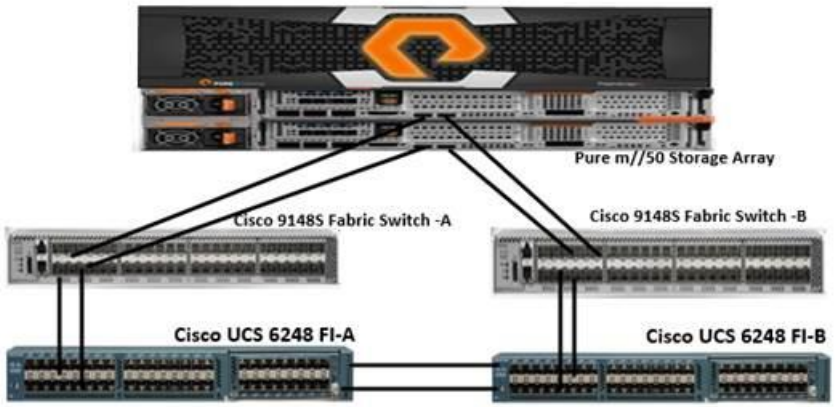

Cisco Unified Computing System Configuration

This section details the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the install guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) and it is beyond the scope of this document. For more information about each step, refer to the following documents:

Cisco UCS Manager Configuration Guides – GUI and Command Line Interface (CLI) Cisco UCS Manager - Configuration Guides - Cisco

Cisco UCS Manager Software to Version 2.2(6c)

This document assumes the use of Cisco UCS Manager Software version 2.2.6(c). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to a higher version of the firmware,) refer to Cisco UCS Manager Install and Upgrade Guides.

Configure Fabric Interconnects at Console

To configure the fabric interconnect, complete the following steps:

1. Connect a console cable to the console port on what will become the primary fabric interconnect.
2. If the fabric interconnect was previously deployed and you want to erase it to redeploy, follow these steps:
    a. Login with the existing user name and password
    b. Enter: connect local-mgmt
    c. Enter: erase config
    d. Enter: yes to confirm
3. After the fabric interconnect restarts, the out-of-box first time installation prompt appears, type “console” and press Enter.

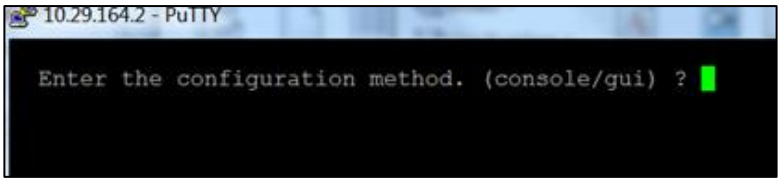

4. Type “setup” at the setup/restore prompt, then press Enter.

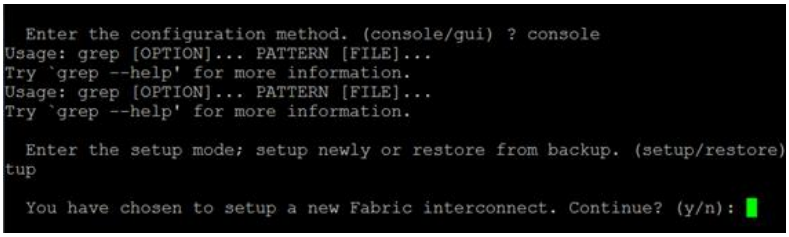

5. Type “y” then press Enter to confirm the setup.

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

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

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

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

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

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

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

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

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

  Enter the password for "admin":
  Confirm the password for "admin":
```

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

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

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

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

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

  Enter the password for "admin":
  Confirm the password for "admin":

  Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (ye
s/no) [n]: yes

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

8. Complete the setup dialog questions.

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

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

  Enter the system name:  UCS-VSAN

  Physical Switch Mgmt0 IP address : 10.29.132.8

  Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

  IPv4 address of the default gateway : 10.29.132.1

  Cluster IPv4 address : 19.29.132.10

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

  Cluster IPv4 address : 10.29.132.10

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

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

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

9. Review the selections and type “yes”.

```
  Following configurations will be applied:

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

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

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

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

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

  Installer has detected the presence of a peer Fabric interconnect. This Fabric
interconnect will be added to the cluster. Continue (y/n) ? y

  Enter the admin password of the peer Fabric interconnect:
    Connecting to peer Fabric interconnect... done
    Retrieving config from peer Fabric interconnect... done
    Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.132.9
    Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
    Cluster IPv4 address          : 10.29.132.10

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

  Physical Switch Mgmt0 IP address :
```

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

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

Base Cisco UCS System Configuration

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

1. Open a web browser and navigate to the Cisco UCS 6248 Fabric Interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.

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

Set Fabric Interconnects to Fibre Channel End Host Mode

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

1. On the Equipment tab, expand the Fabric Interconnects node and click Fabric Interconnect A.
2. On the General tab in the Actions pane, click Set FC End Host mode.
3. Follow the dialogs to complete the change.

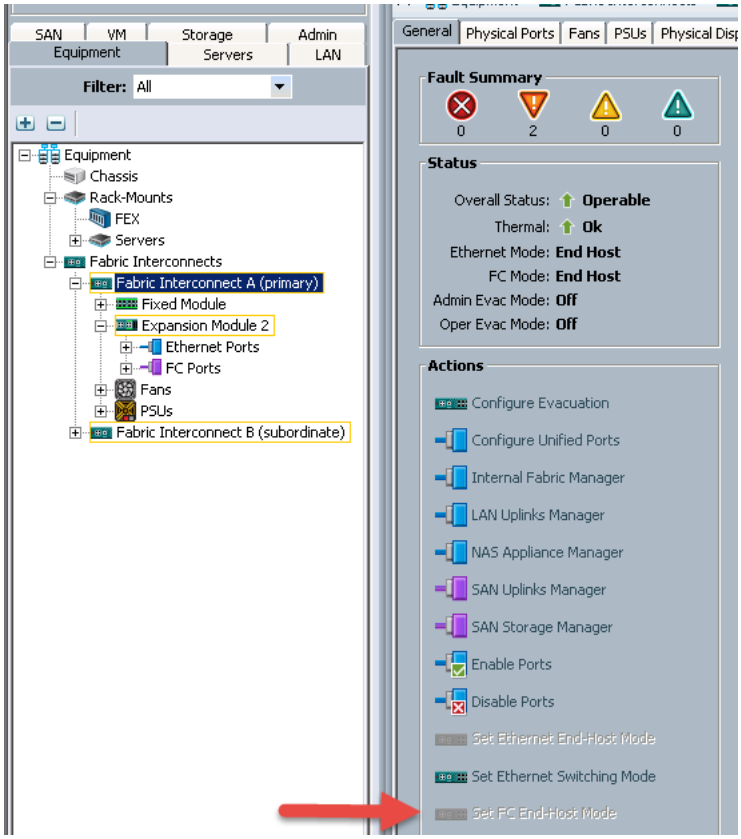

Both Fabric Interconnects automatically reboot sequentially when you confirm you want to operate in this mode.

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.

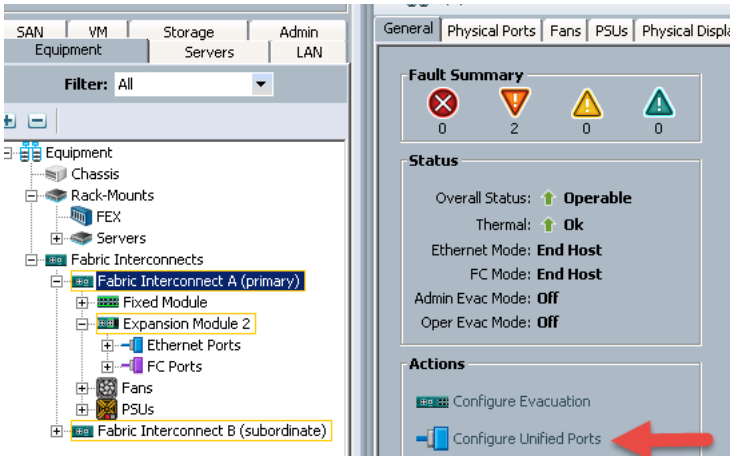

3. Click Configure Expansion Module Ports.

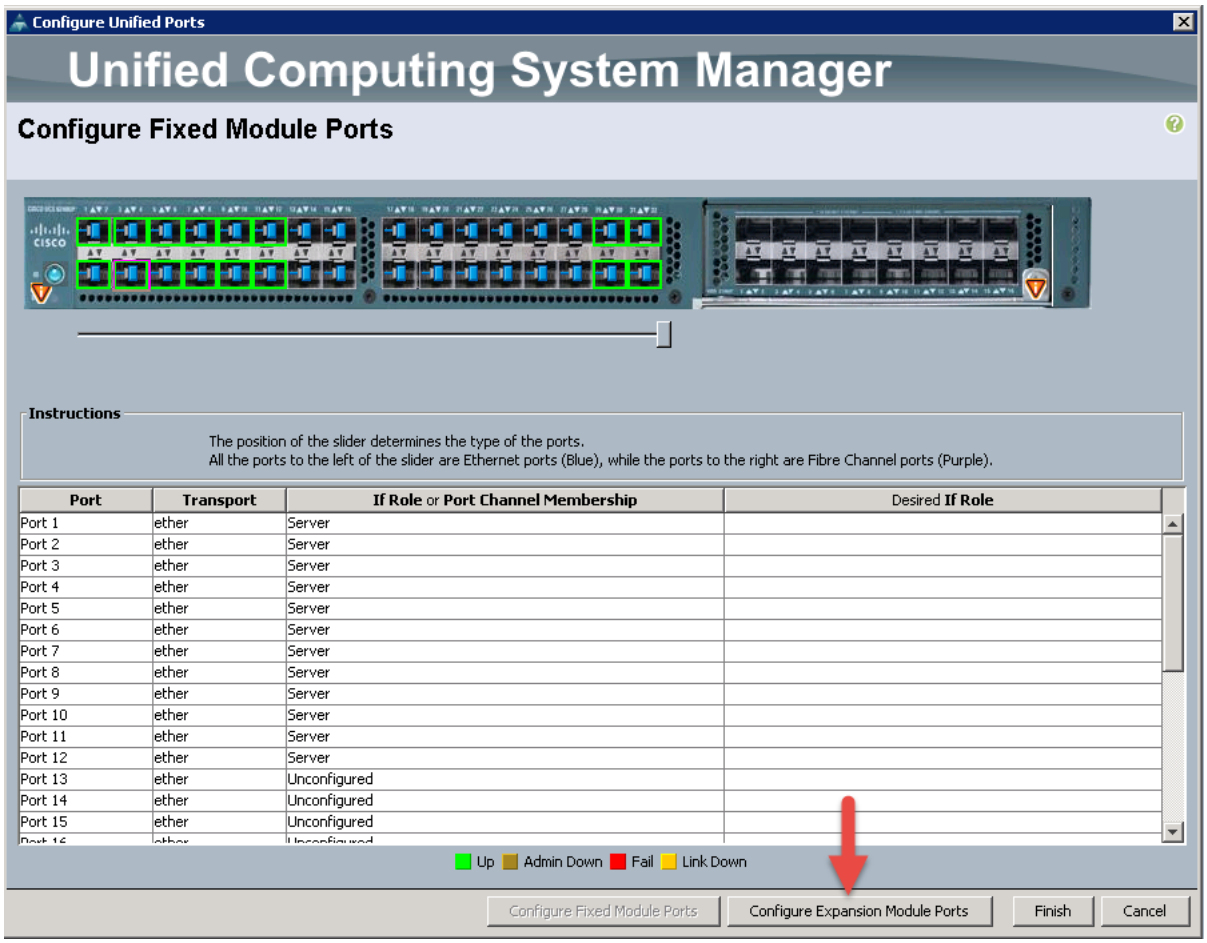

4. Move the slider to the left.

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.

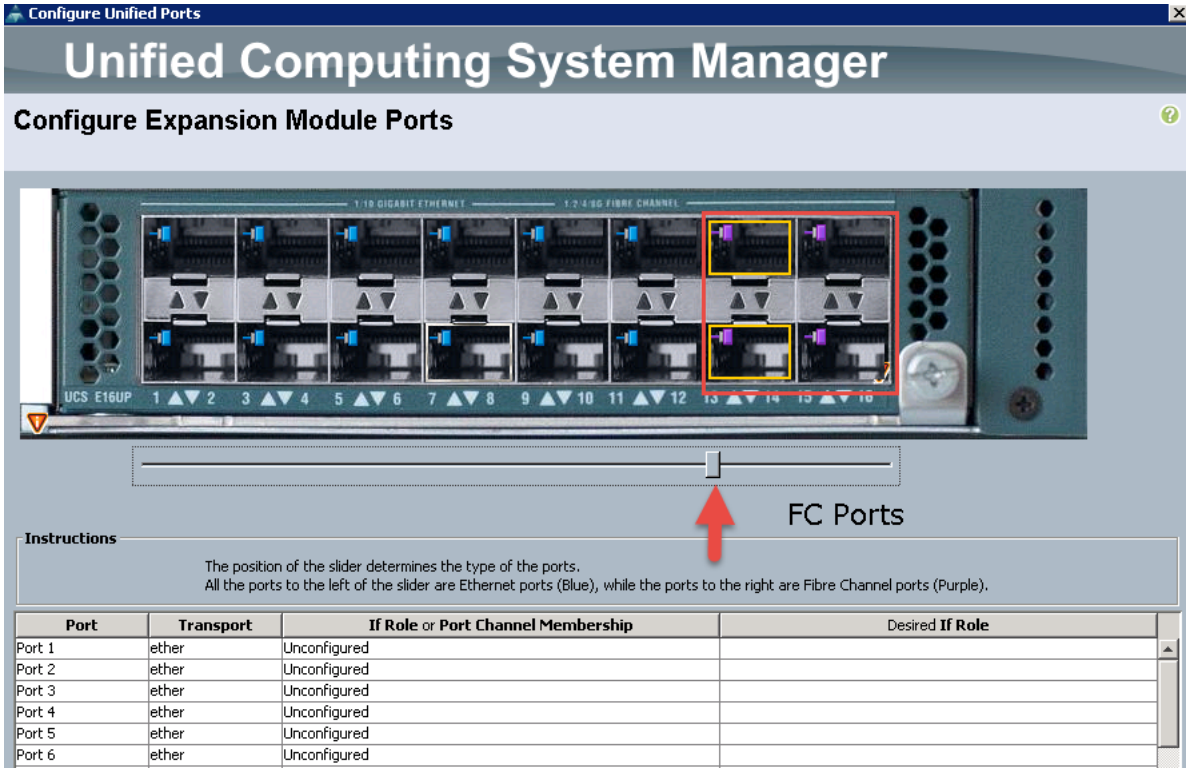

5. Click Finish, then click Yes to confirm. This action will cause a reboot of the Expansion Module.

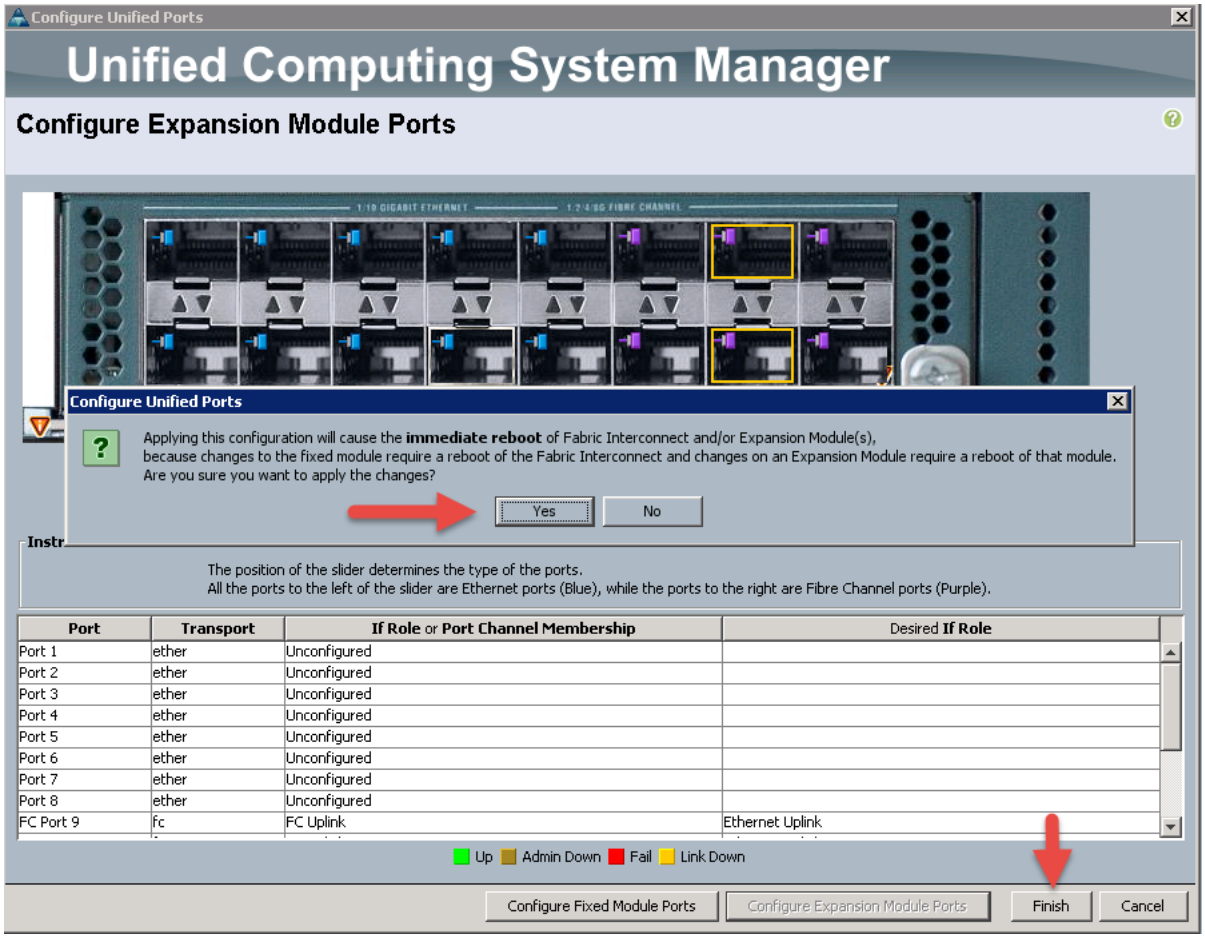

After the expansion module reboot, your FC Ports configuration should look like the figure below:

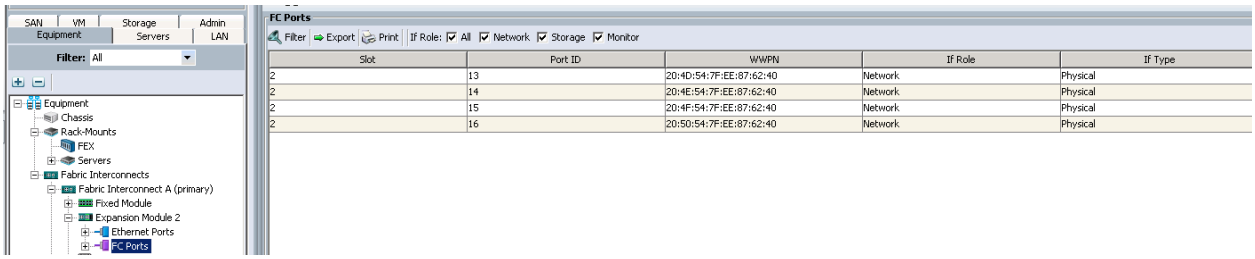

6. Repeat this procedure for Fabric Interconnect B.

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

Edit Chassis Discovery Policy

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

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

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

2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to 4-link.
4. Set the Link Grouping Preference to Port Channel.

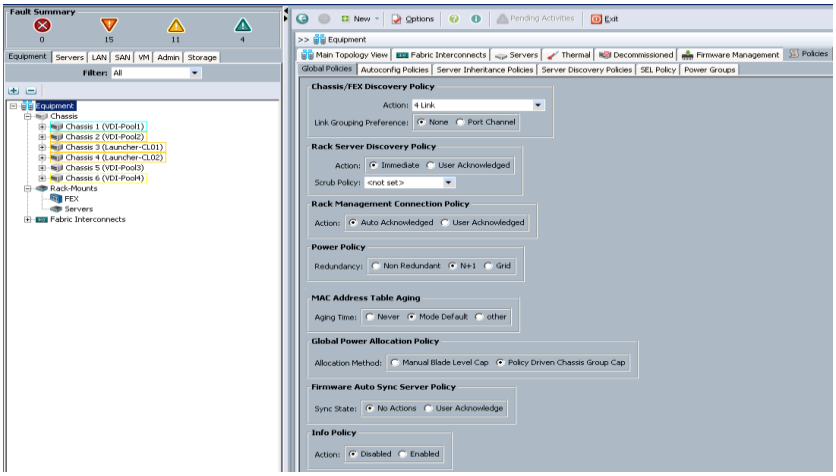

5. Click Save Changes.
6. Click OK.

Acknowledge Cisco UCS Chassis

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

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.

Figure 21 Acknowledge Cisco UCS Chassis

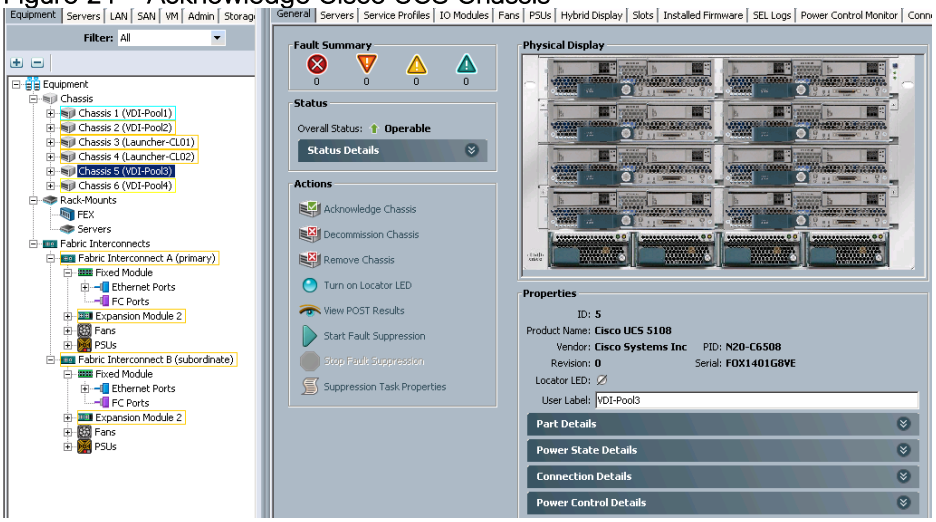

4. Click Yes and then click OK to complete acknowledging the chassis.

5. Repeat for each of the remaining chassis.

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

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

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Pools > root > IP Pools > IP Pool ext-mgmt.
3. In the Actions pane, select Create Block of IP Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

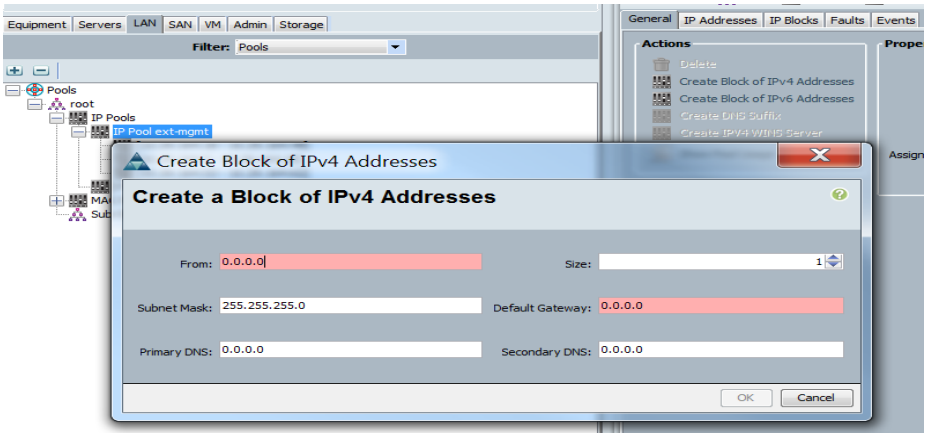

5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, in the navigation pane, click the Admin tab.
2. Select All > Timezone Management.
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.

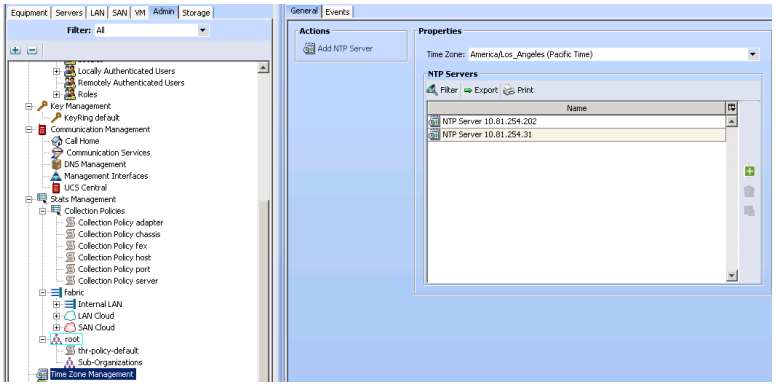

6. Enter the NTP server IP address and click OK.

7. Click OK.

Enable Server and Ethernet Uplink Ports

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

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

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

3. Expand Ethernet Ports.

4. Select ports 1 through 28 (Ports highlighted 9, 10, 15 & 16 not configured) that are connected to the Cisco IO Modules of the four B-Series 5108 Chassis, right-click them, and select Configure as Server Port.

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

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

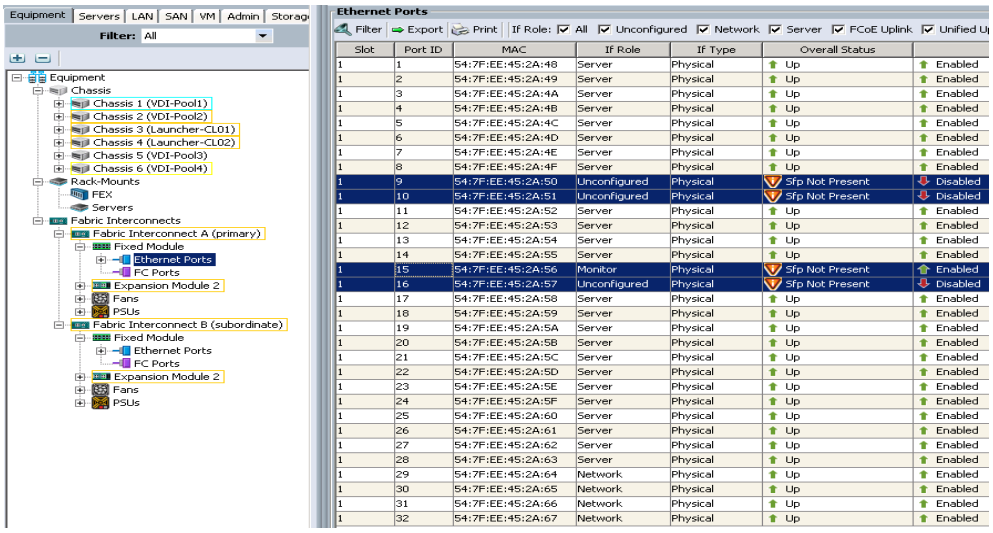

7. Repeat the above steps for Fabric Interconnect B. The screenshot below shows the server ports for Fabric B.

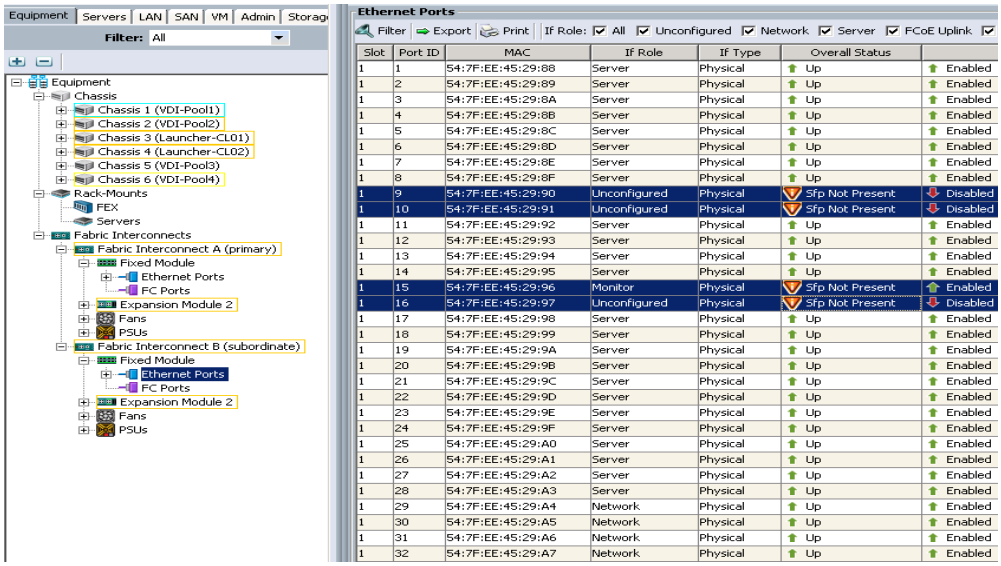

To configure network ports used to uplink the Fabric Interconnects to the Cisco Nexus 9172PX switches, follow these steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports 29 through 32 that are connected to the Nexus 9172PX switches, right-click them, and select Configure as Network Port.
5. Click Yes to confirm ports and click OK.
6. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that ports have been configured correctly in the in the Role column.
7. Verify the Ports connected to Cisco Nexus upstream swithces are now configured as network ports.
8. Repeat the above steps for Fabric Interconnect B. The screenshot shows the network uplink ports for Fabric B.
9. Successful configuration should result in ports 29-23 configured as network ports as shown in the screen shot below:

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

Create Uplink Port Channels to Cisco Nexus 9372PX Switches

In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 9372PX switches and one from Fabric B to both Cisco Nexus 9372PX switches.

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, expand node Fabric A tree:
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 29 as the unique ID of the port channel.
6. Enter FI-A-Uplink as the name of the port channel.
7. Click Next.

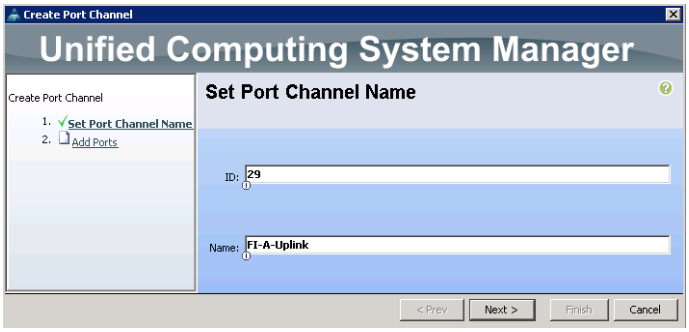

8. Select ethernet ports 29-32 for the port channel
9. Click Finish.

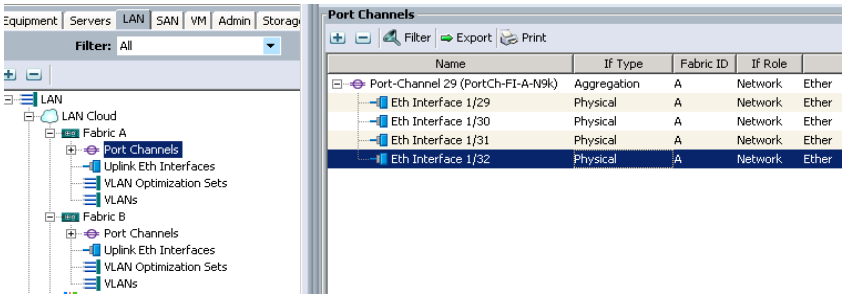

10. Repeat steps 1-9 for Fabric Interconnect B, substituting 31 for the port channel number and FI-B-Uplink for the name. The resulting configuration should look like the screen shot below:

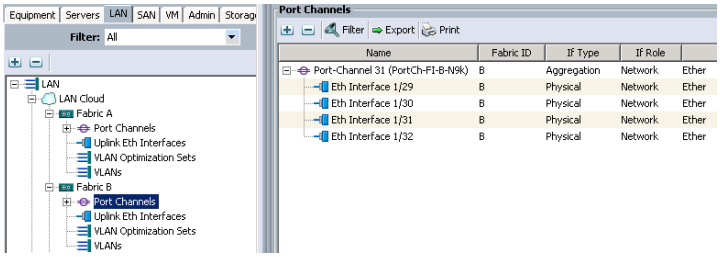

Create Uplink Port Channels to Cisco MDS 9148S Switches

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.

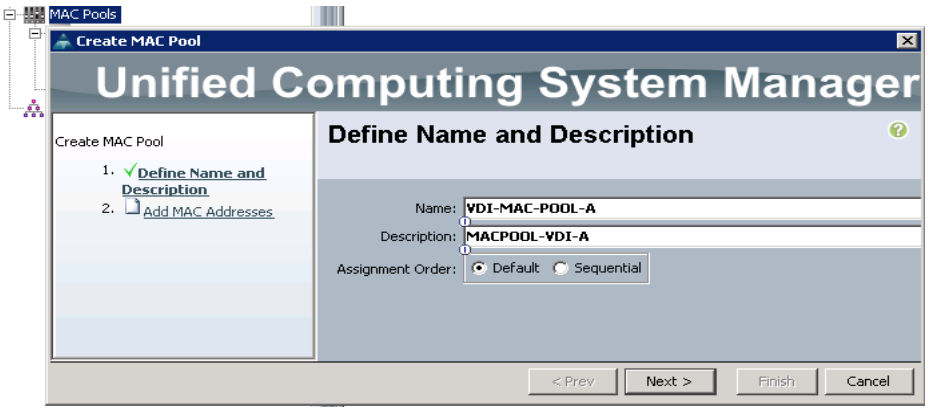

7. Enter the seed MAC address and provide the number of MAC addresses to be provisioned.

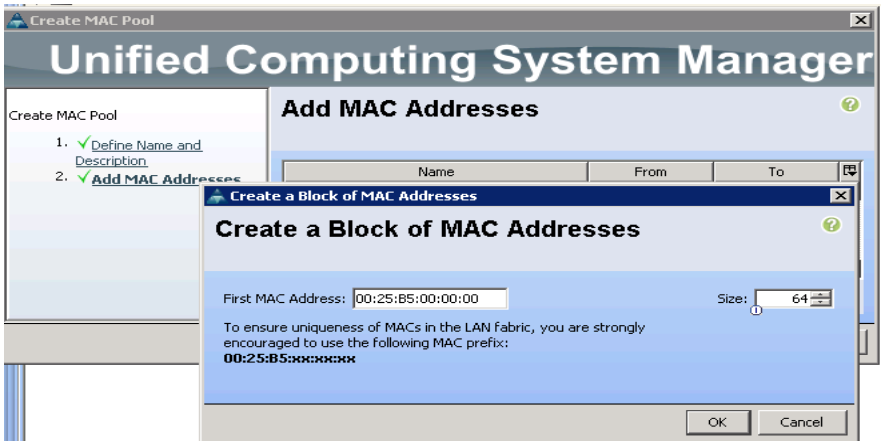

8. Click OK, then click Finish.

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, follow these 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

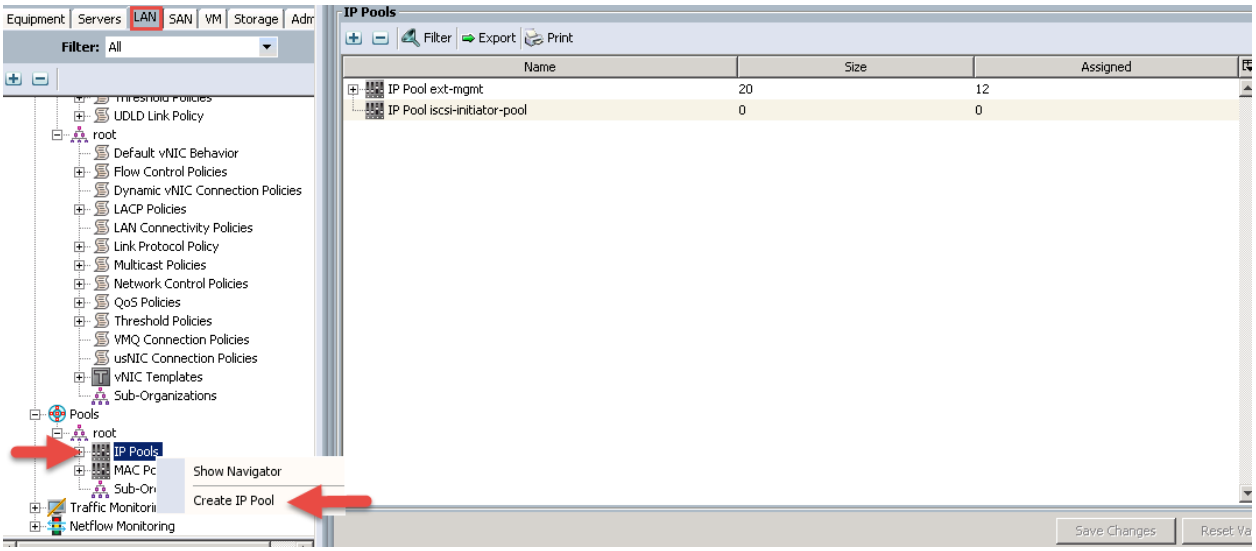

2. Provide a Name, choose Default or Sequential, and then click Next.

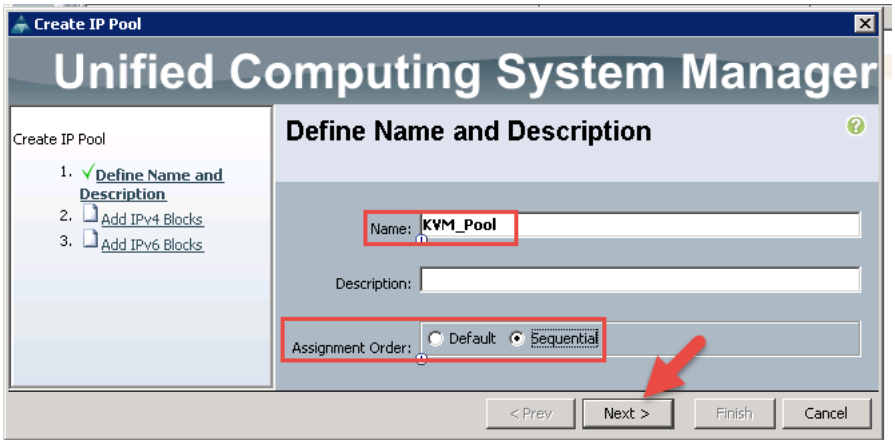

3. Click the green + sign to add an IPv4 address block.

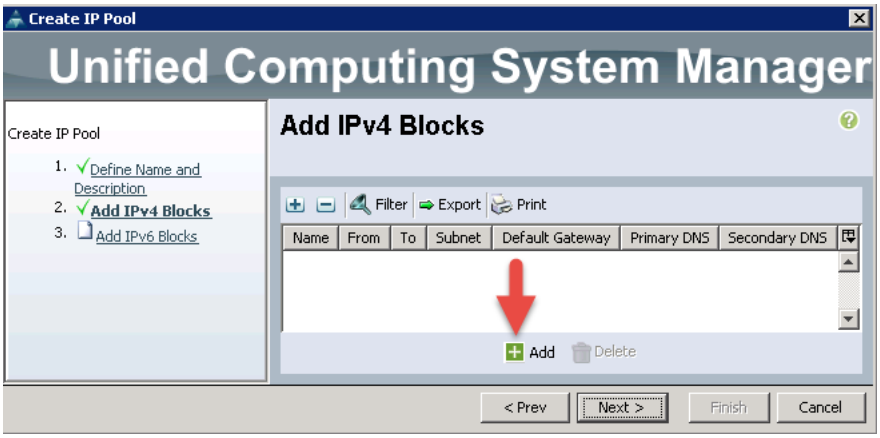

4. Complete the starting IP address, size, subnet mask, default gateway, primary and secondary DNS values for your network, then click OK.

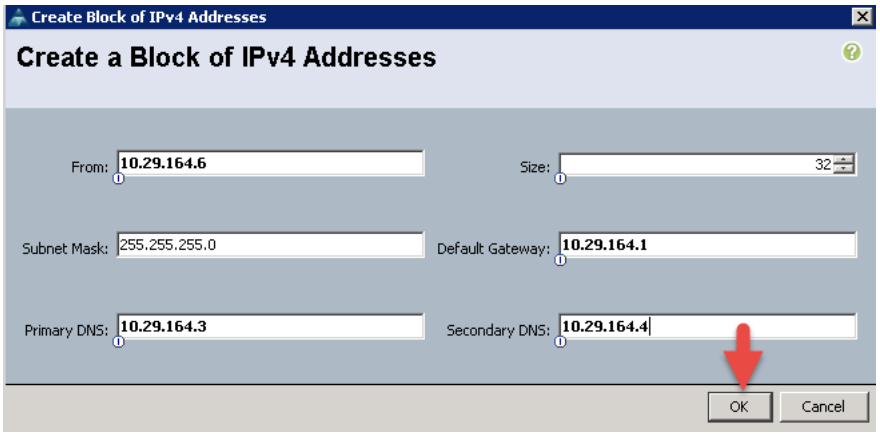

5. Click Next.

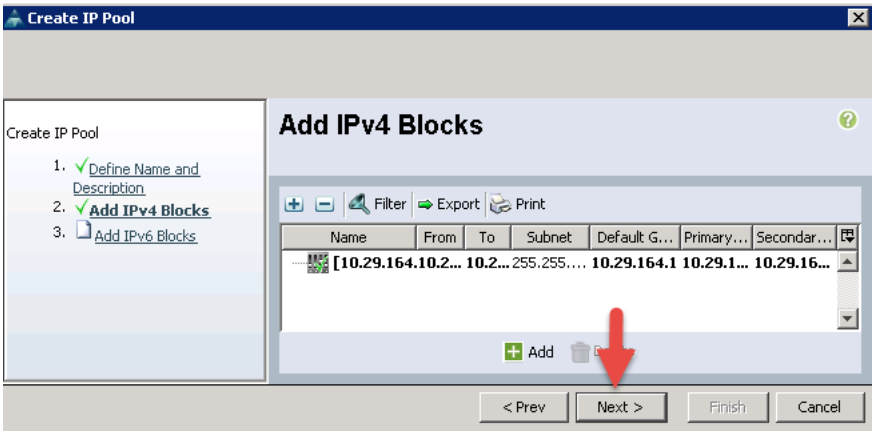

6. Click Finish.

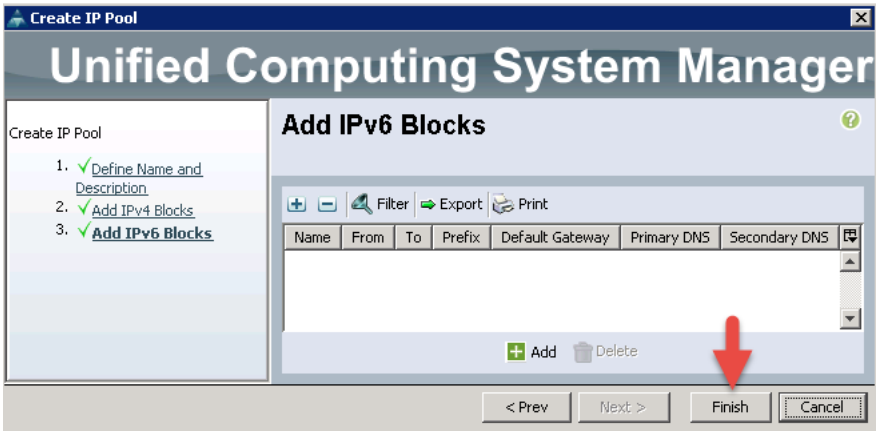

7. Click OK on the success pop-up.

Create WWPN Pools

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Under WWPN Pools, right click WWPN Pools and select Create WWPN Pool.
4. Assign a name and optional description.

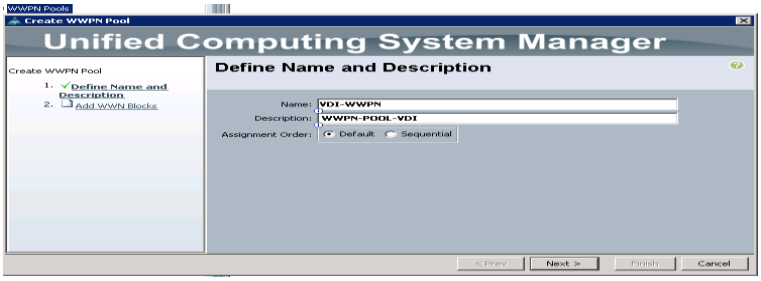

5. Assignment order can remain Default.

6. Click Next.

7. Click Add to add block of Ports.

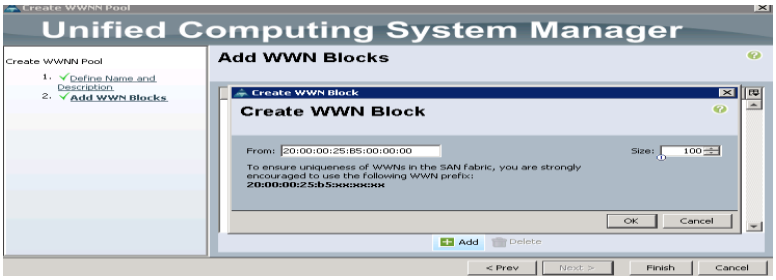

8. Enter number of WWNNs. For this study we did 100.

9. Click Finish.

Create UUID Suffix Pool

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

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

2. Select Pools > root.

3. Right-click UUID Suffix Pools.

4. Select Create UUID Suffix Pool.

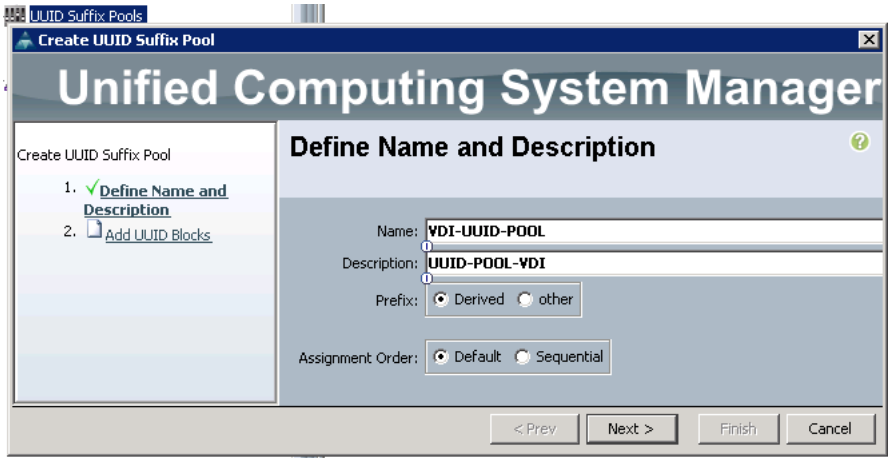

5. Enter UUID_Pool as the name of the UUID suffix pool.

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

7. Keep the prefix at the derived option.

8. Click Next.

9. Click Add to add a block of UUIDs.

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

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

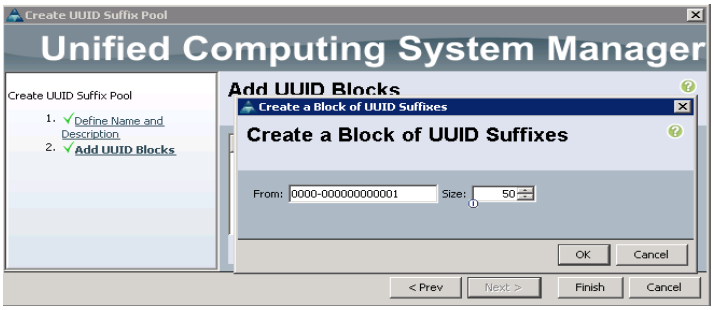

Create Server Pool

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

Table 12 VLANs Created

VLAN Name	VLAN ID	VLAN Purpose	vNIC Assignment
Default	1	Native VLAN	vNIC-Template-A vNIC-Template-B
In-Band-Mgmt	160	VLAN for in-band management interfaces	vNIC-Template-A vNIC-Template-B
Infra-Mgmt	161	VLAN for Virtual Infrastructure	vNIC-Template-A vNIC-Template-B
vMotion	166	VLAN for VMware vMotion	vNIC-Template-A vNIC-Template-B
VDI	162	Virtual Desktop traffic	vNIC-Template-A vNIC-Template-B
OB-Mgmt	164	VLAN for out-of-band management interfaces	vNIC-Template-A vNIC-Template-B

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs
5. Enter MGMT as the name of the VLAN to be used for in-band management traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 160 as the ID of the management VLAN.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

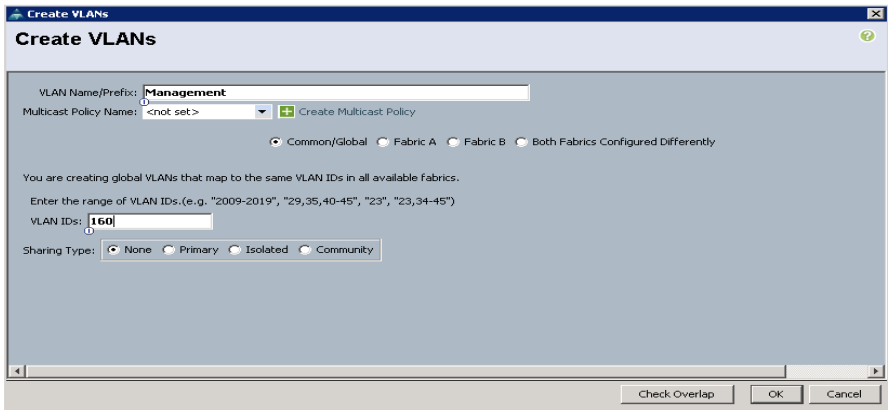

10. Repeat the above steps to create all VLANs and configure the Default VLAN as native.

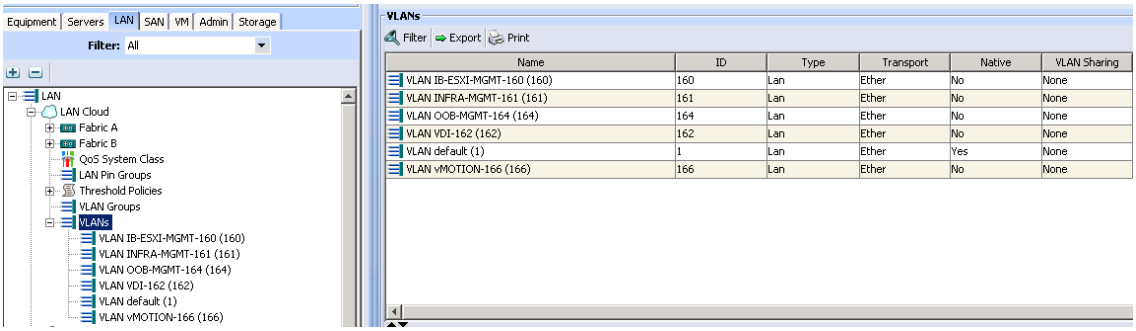

Create VSANs

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

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

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

2. Select SAN > SAN Cloud.
3. Under Fabric A, right-click VSANs.
4. Select Create VSANs.
5. Enter VSAN20 as the name of the VSAN to be used for in-band management traffic.
6. Select Fabric A for the scope of the VSAN.
7. Enter 20 as the ID of the VSAN.
8. Click OK, and then click OK again.

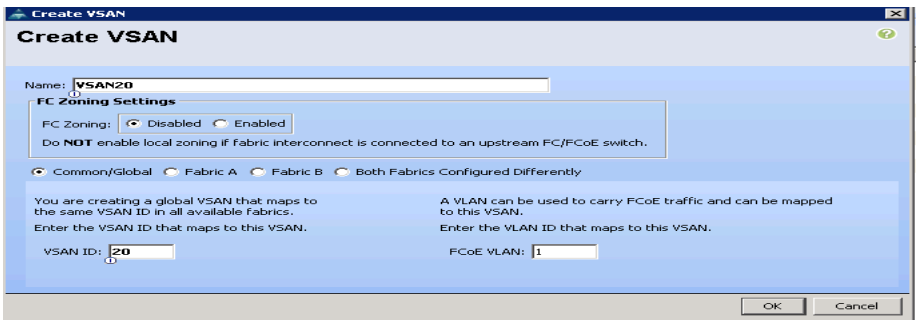

9. Repeat the above steps on Fabric B with VSAN30 to create the VSANs necessary for this solution.

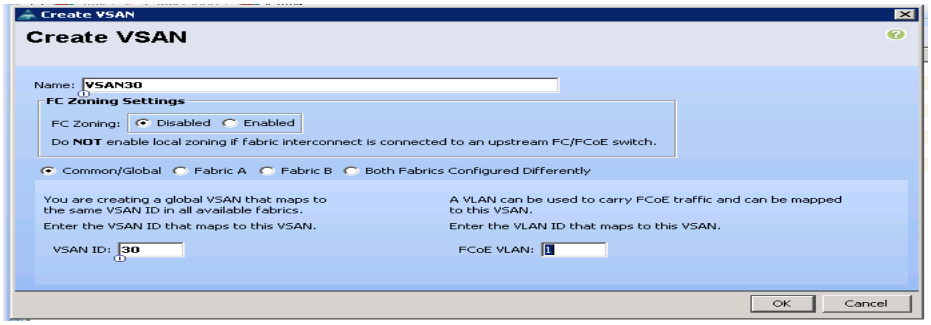

VSAN 20 and 30 are configured as shown below:

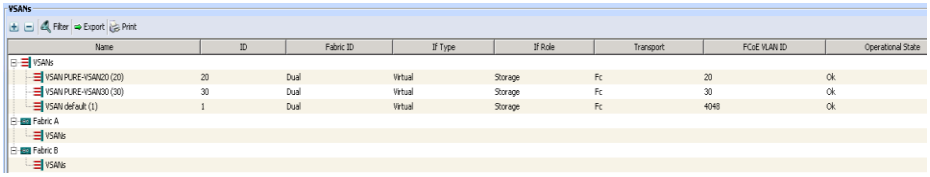

10. After configuring VSANs both sides, go into the port-channel created earlier in the section ‘Create up-links for MDS 9148S and add the respective VSANs to their port channels. VSAN20 in this study is assigned to Fabric A and VSAN30 is assigned to Fabric B. (VSAN20 Should only be on Fabric A and VSAN30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

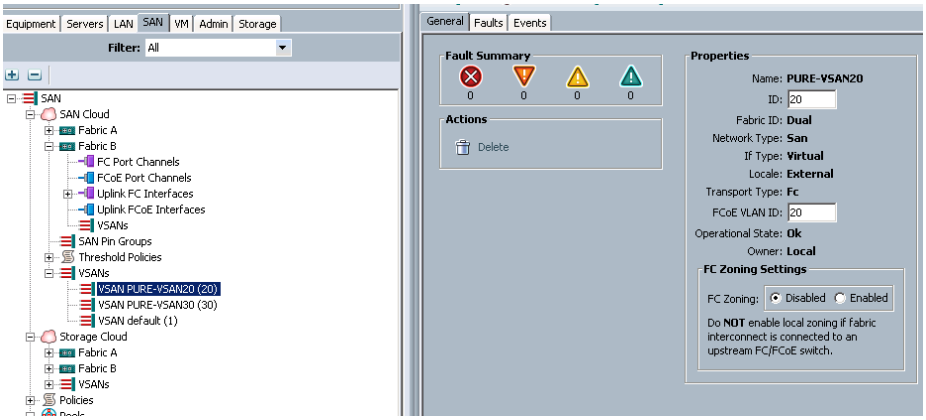

11. Go to the Port-Channel for each Fabric and assign the VSAN appropriately.

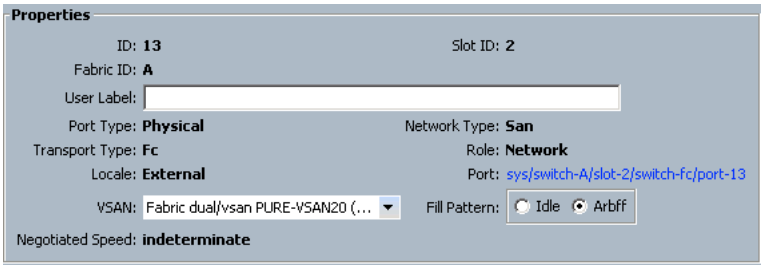

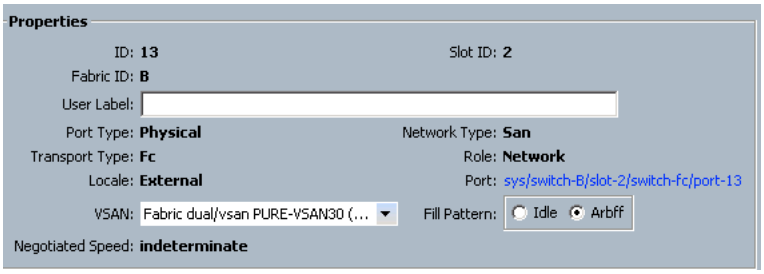

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter VM-Host as the name of the host firmware package.
6. Leave Simple selected.
7. Select the version 2.2. (6c) for both the Blade Package
8. Click OK to create the host firmware package.

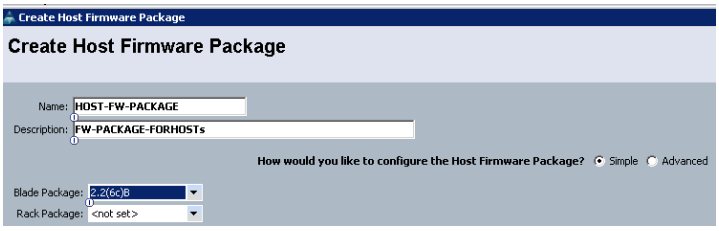

Set Jumbo Frames in Cisco UCS Fabric

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.

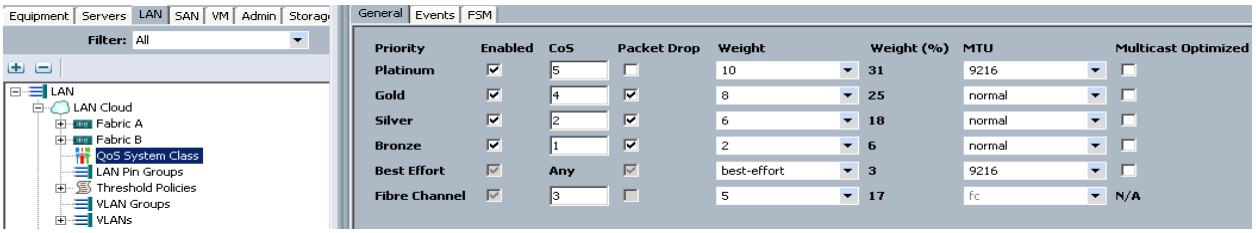

Create Network Control Policy for Cisco Discovery Protocol

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

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

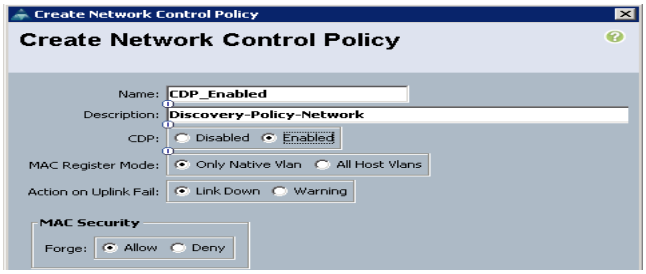

Create Power Control Policy

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.

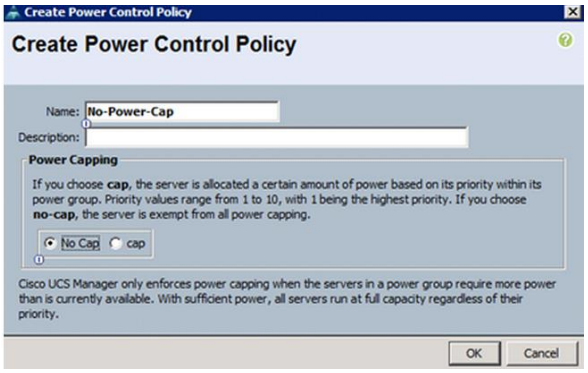

Cisco UCS System Configuration for Cisco UCS B-Series

Create Server BIOS Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter B200-M4-BIOS as the BIOS policy name.
6. Configure the remaining BIOS policies as follows and click Finish.

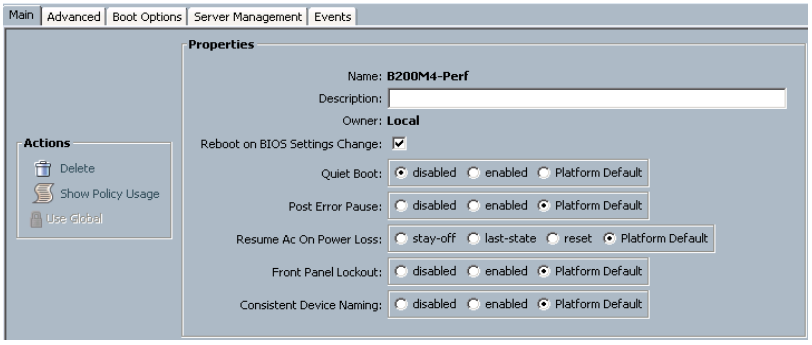

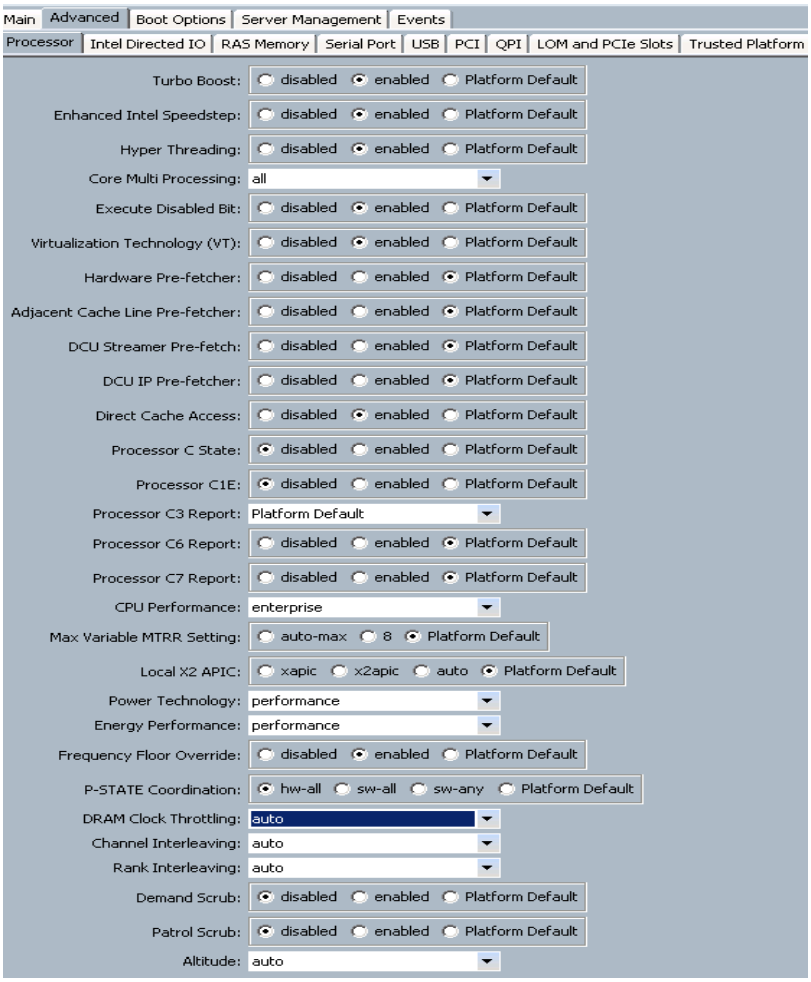

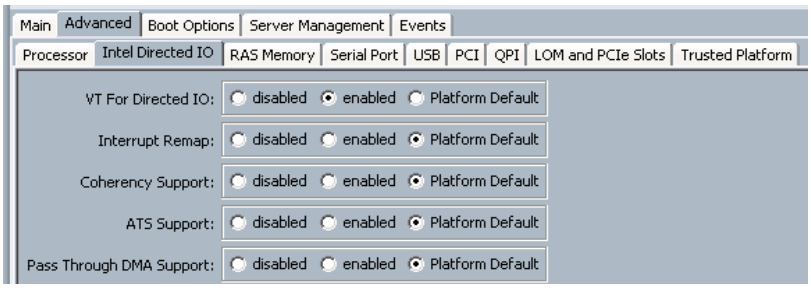

7. Click Finish.

Configure Update Default Maintenance Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.

5. Click Save Changes.
6. Click OK to accept the change.

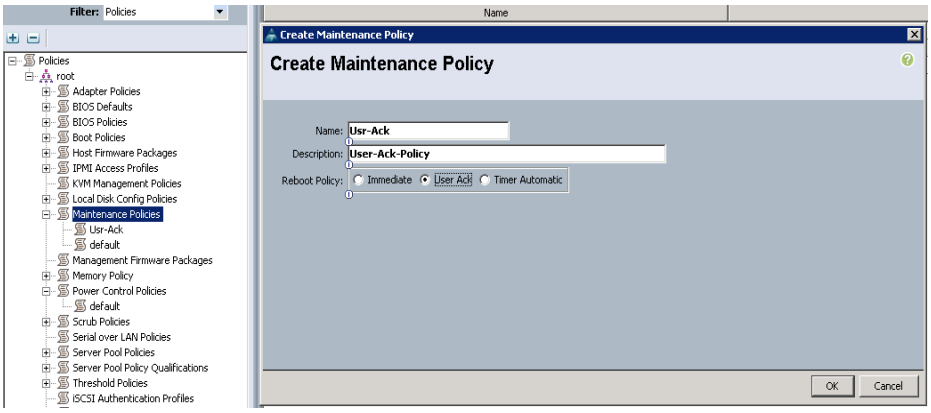

Create vNIC Templates for Cisco UCS B-Series

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

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

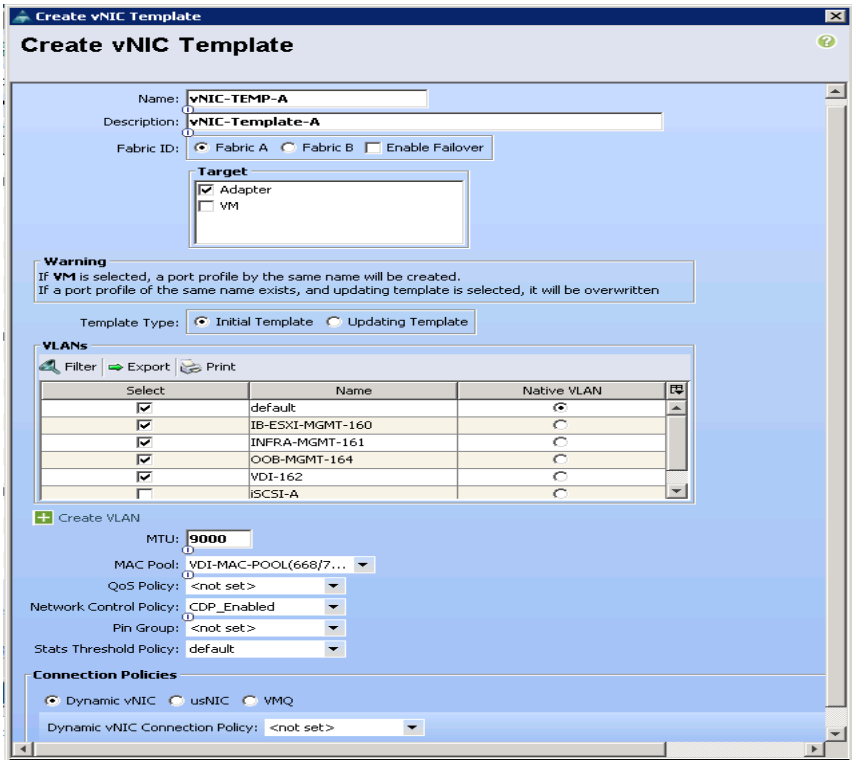

17. In the navigation pane, select the LAN tab.
18. Select Policies > root.
19. Right-click vNIC Templates.
20. Select Create vNIC Template.
21. Enter vNIC_Template_B as the vNIC template name.
22. Select Fabric B.
23. Do not select the Enable Failover checkbox.
24. Under Target, make sure the VM checkbox is not selected.
25. Select Updating Template as the template type.
26. Under VLANs, select the checkboxes for MGMT, Default, VDI, Infra, and vMotion.
27. Set Native-VLAN as the native VLAN.
28. For MTU, enter 9000.
29. In the MAC Pool list, select MAC_Pool_B.
30. In the Network Control Policy list, select CDP_Enabled.
31. Click OK to create the vNIC template.

32. Click OK.

Create vHBA Templates for Cisco UCS B-Series

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter VDI-vHBA_FI-A as the vHBA template name.
6. Keep Fabric A selected.
7. Select VSAN20 for Fabric A from the drop down.
8. Change to Updating Template.
9. For Max Data Field keep 2048.
10. Select VDI-Pool-WWPN (created earlier) for our WWPN Pool.
11. Leave the remaining as is.
12. Click OK.

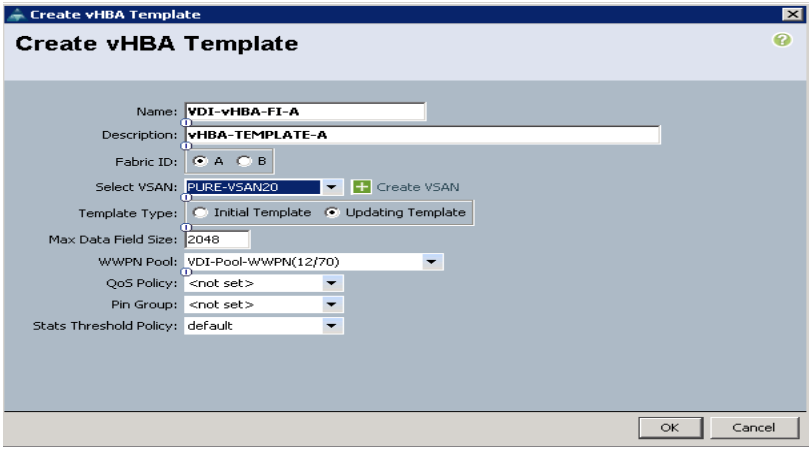

13. In the navigation pane, select the LAN tab.
14. Select Policies > root.
15. Right-click vHBA Templates.
16. Select Create vHBA Template.

17. Enter VDI-vHBA_FI-B as the vHBA template name.
18. Select Fabric B.
19. Select VSAN30 for Fabric B from the drop down.
20. Change to Updating Template.
21. For Max Data Field keep 2048.
22. Select VDI-Pool-WWPN (created earlier) for our WWPN Pool.
23. Leave the remaining as is.
24. Click OK.

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. Change to Updating Template.
5. Select UUID pool created earlier.

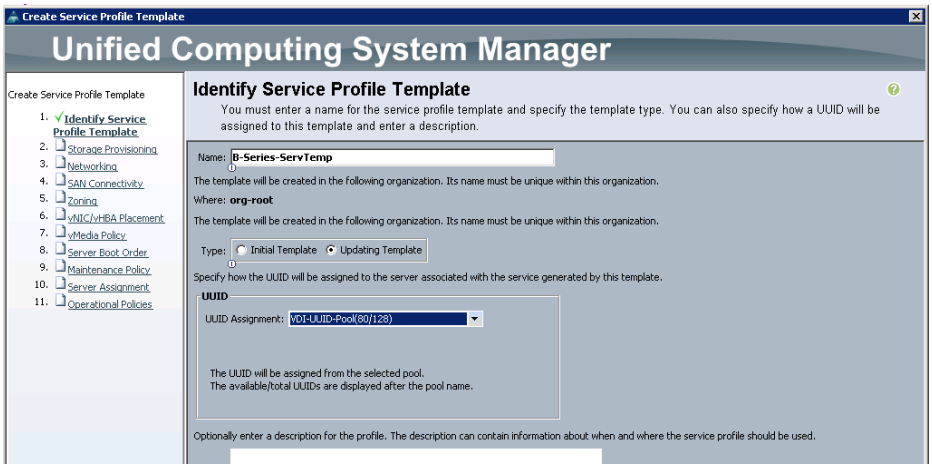

6. Click Next.
7. Click Next through Storage Provisioning.
8. Under Networking, Select Expert.
9. Click Add.

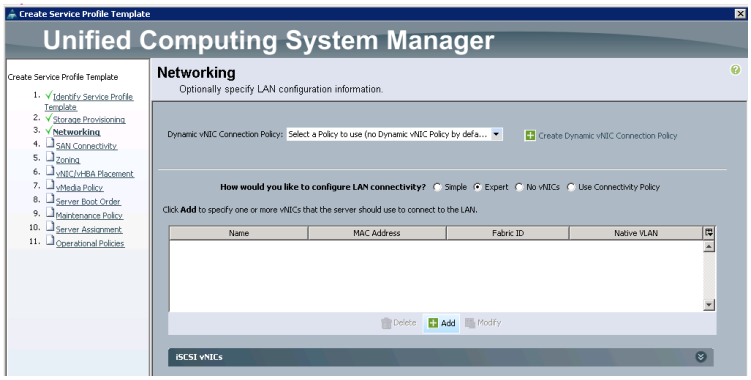

10. Name it vNIC-A.

11. Select check box for Use vNIC Template.

12. Under vNIC template select the vNIC-A.

13. For Adapter Policy select VMware.

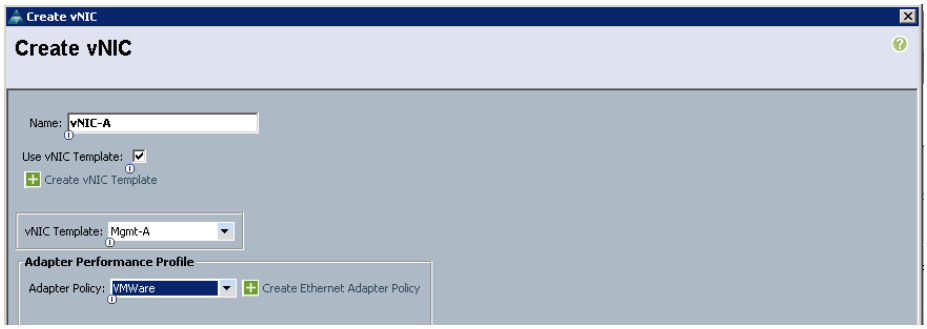

14. Repeat networking steps for vNIC-B.

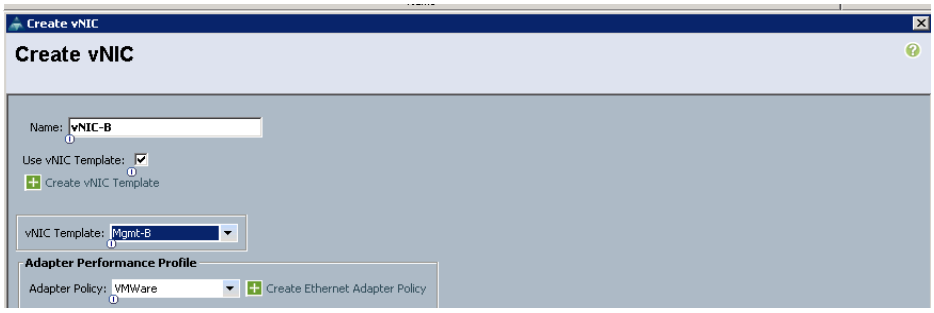

15. Click Next.

16. Under SAN Connectivity, select Expert.

17. Select WWNN Assignment from the Pool created earlier.

18. Click Add.

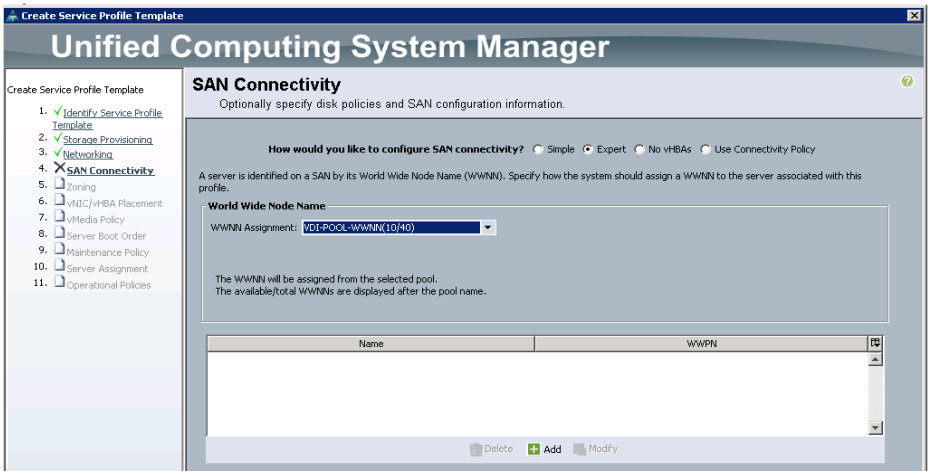

19. Name the adapter vHBA-A.

20. Select vHBA Template: vHBA-A.

21. Select Adapter Policy : VMWare.

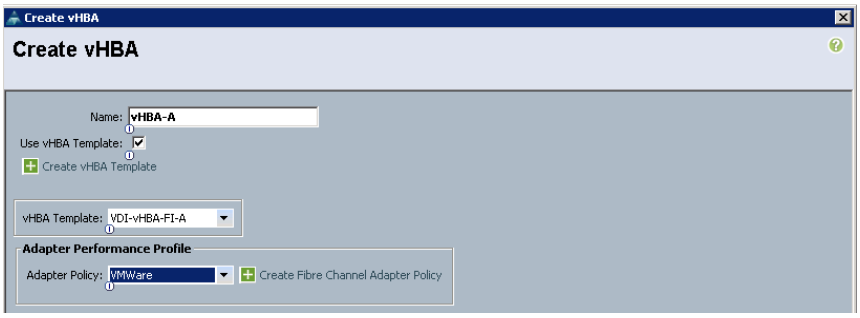

22. Repeat steps for vHBA-B on Fabric B.

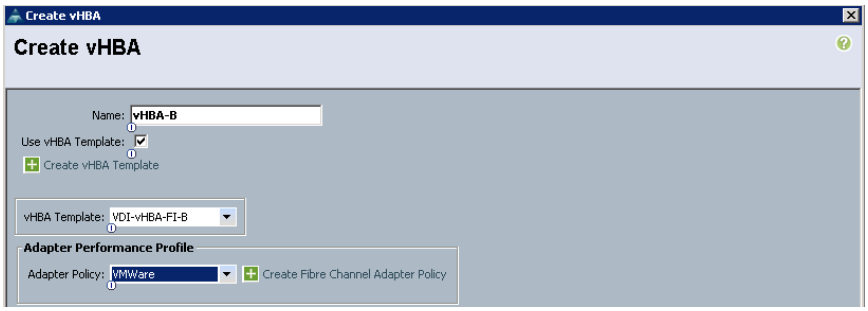

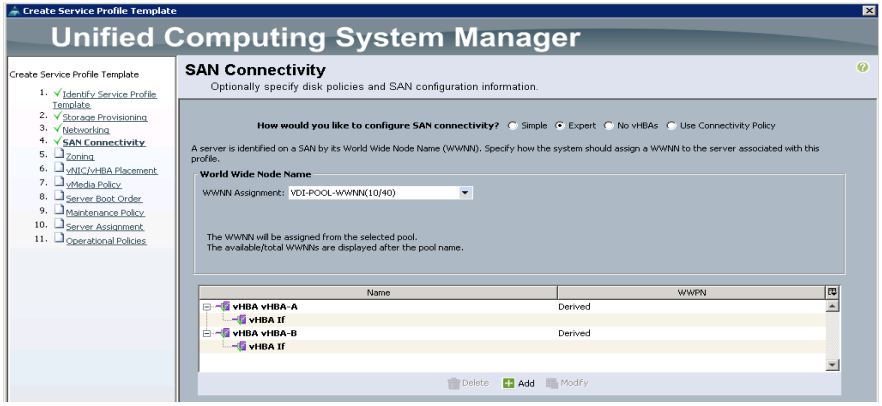

23. No Zoning will be used.

24. Click Next through vNIC/vHBA Placement policy.

25. Click Next through vMedia Policy.

26. Click Create Boot Policy to create a Boot From SAN policy.

27. Add SAN primary target and second target WWN.

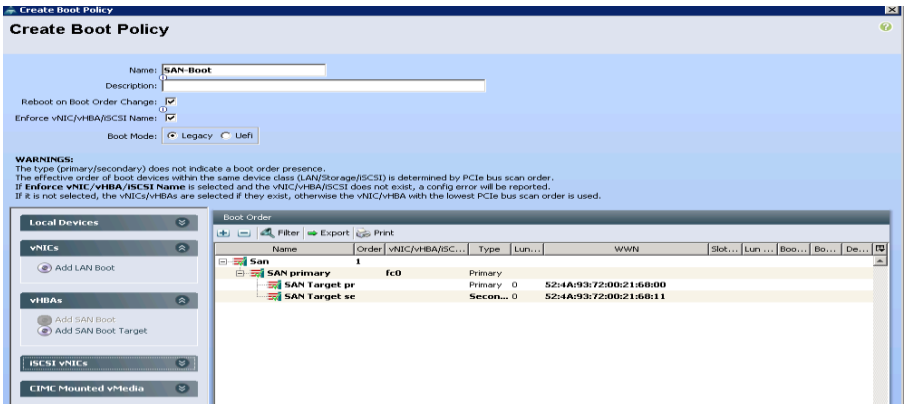

28. Add the SAN Secondary target.

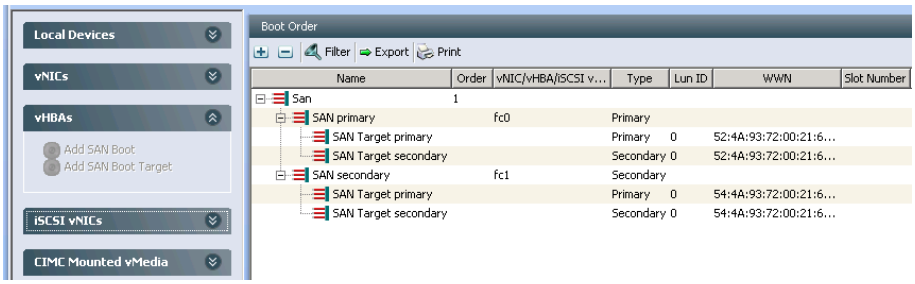

SAN Boot policy and advantages are detailed in a later section.

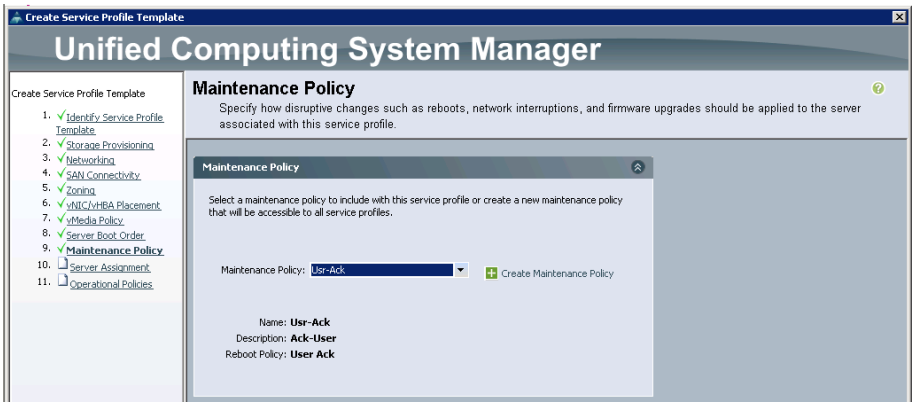

Configuring ESXI host to Boot from Pure Storage FlashArray//m50 SAN

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, go to Cisco UCS Manager, complete the following steps:

1. Right-click the Boot Policies option shown below and select Create Boot Policy.

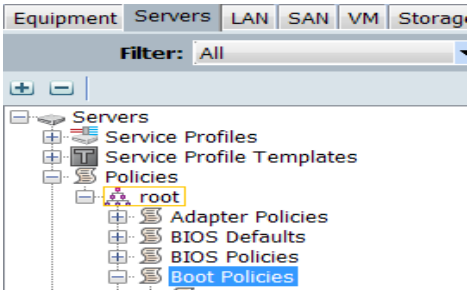

2. Enter a name for the boot policy and expand the vHBA menu shown below:

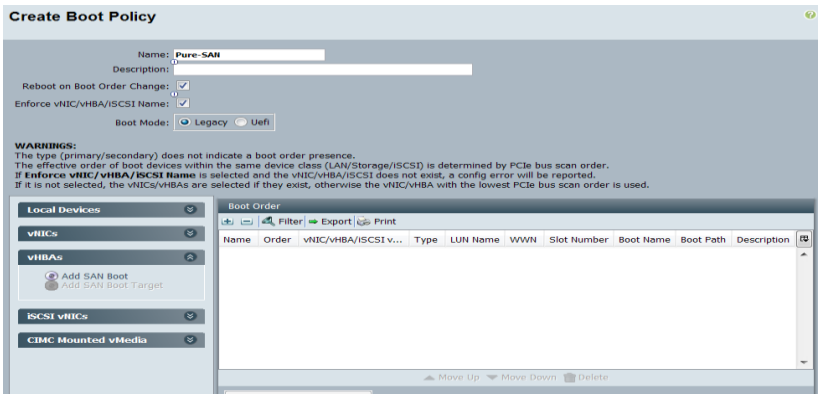

3. Add SAN boot.

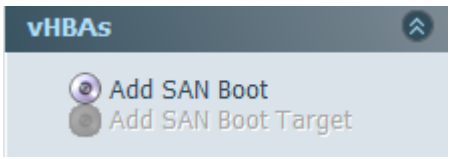

4. After selecting Add SAN Boot, add the primary vHBA as shown below. Note that the vHBA name needs to match exactly.

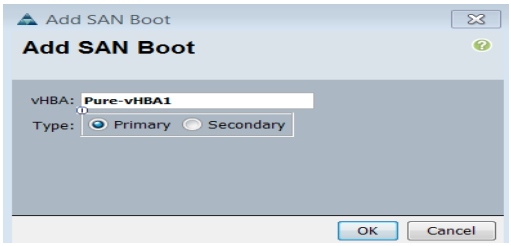

5. Repeat the procedure to add a secondary SAN Boot option.

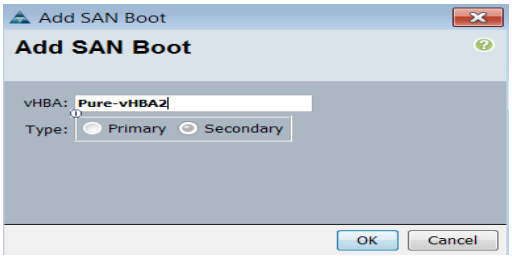

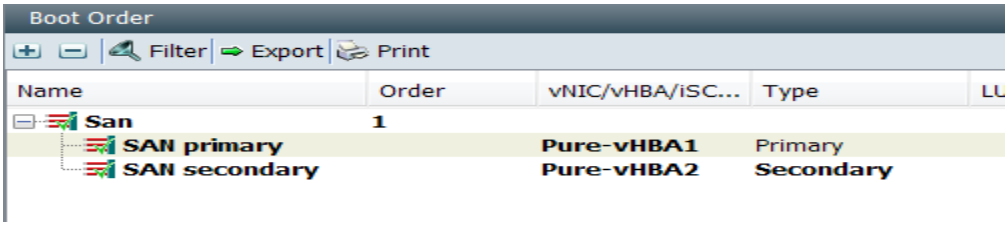

6. Add SAN Boot Targets to the primary and secondary. The SAN boot targets also include primary and secondary options in order to maximize resiliency and number of paths.

7. Highlight the SAN primary and select Add SAN Boot Target to SAN Primary.

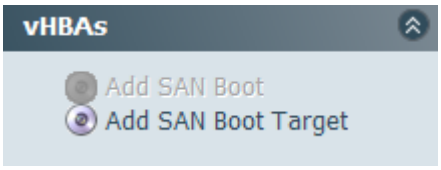

8. From the Pure Storage GUI, find and enter the Pure Storage WWN for Controller 0, Fibre Channel Port 0. This information can be found in the Pure Storage GUI under System > Host Connections at the bottom of the screen under Target Ports:

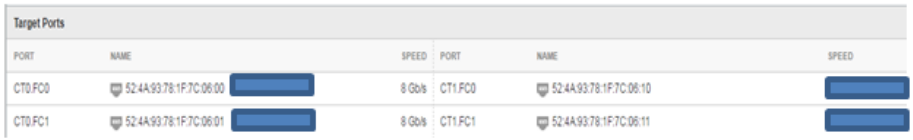

9. When the Pure WWNs are recorded, use port CT0.FC0 for the first Boot Target WWPN:

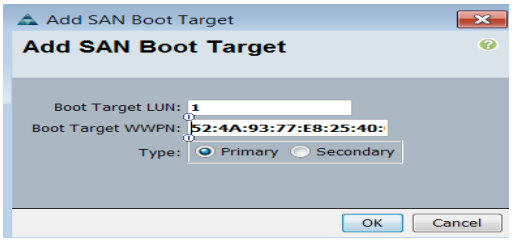

10. Add a secondary SAN Boot Target by clicking the Add SAN Boot Target to SAN Primary while the primary SAN Boot option is highlighted. Enter the Pure Storage WWPN for CT1.FC0.

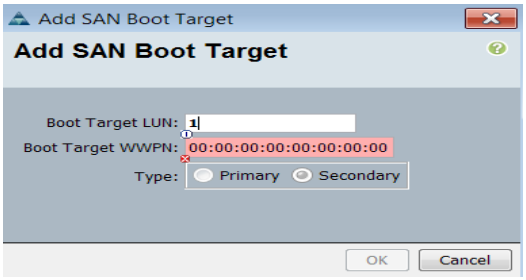

11. Repeat this procedure for the secondary SAN boot target and use WWPN for CT0.FC1 and CT1.FC1 in the primary and secondary SAN boot options.

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

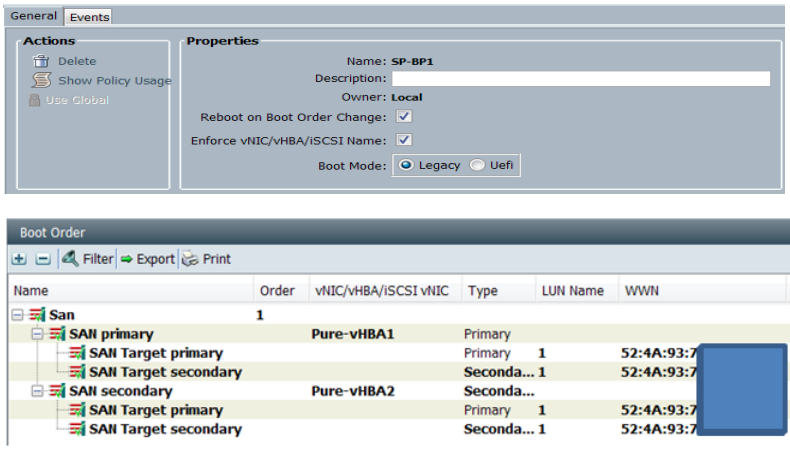

Provisioning a volume for SAN boot for Pure Storage is extremely simple. The only difference between a SAN volume and a vCenter datastore is that you will connect the SAN boot volume to the single host that will be booting from the volume.

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

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

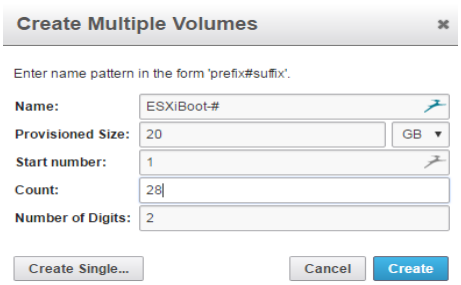

15. To connect the LUN to an ESXi host, select a newly created volume under the Volumes tab and then select Connect Hosts from the options under the gear located on the right:

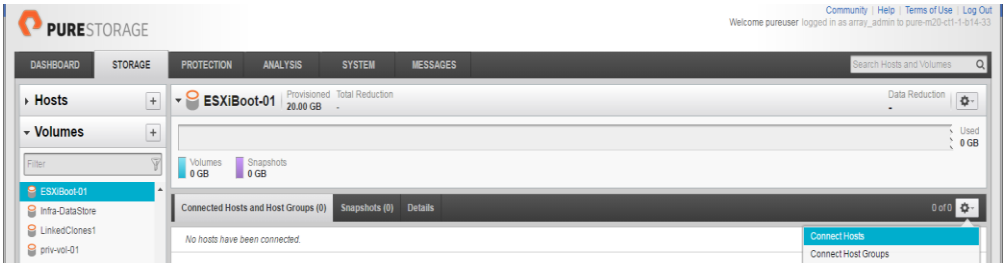

16. Click the single host you to connect to this volume and click Confirm.

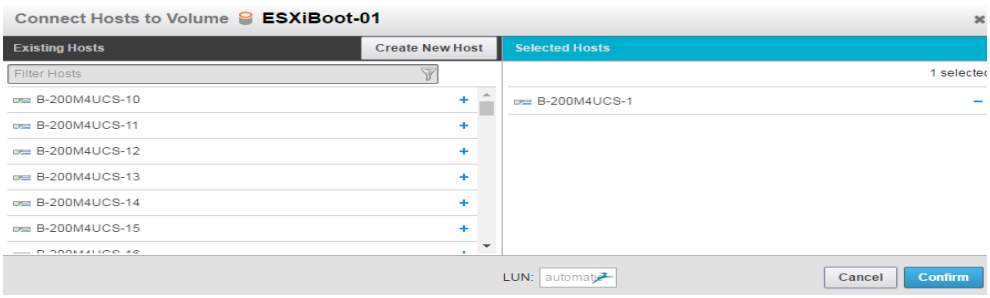

17. Repeat this procedure with all ESXi hosts. Remember, it is important to only connect a single host to each boot LUN.

Pure Storage FlashArray//m50 Configuration for Cisco Validated Design

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

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

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

The Cisco UCS hosts are redundantly connected to the array controllers (via the Cisco MDS 9148S) with two FC connections to each controller from two HBAs on each Cisco UCS host over the 8GB Fibre Channel protocol for a total of four logical paths for 32GB/s of total throughput with no single point of failure. 16GB is also supported by default on the FlashArray//m50 but was not used within this Cisco Validated Design. The FlashArray//m50 array will support up to eight FC connections in total in the baseline system and using all eight connections is recommended in production deployments in order to maximize resiliency and throughput capabilities of the array. However, we will show that a large scale VDI deployment can be accomplished with only half of the available FC ports in use to further prove the resiliency and performance of the FlashArray//m50.

Pure Storage FlashArray//m Configuration

The Pure Storage FlashArray//m50 used in this validated design required a total of 5 rack units: 3 for the FlashArray//m50 base unit and 2 for the 44TB expansion shelf in a standard length rack. The front view of the array can be seen in Figure 22 below:

Figure 22 Pure Storage FlashArray//m Storage Front View

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

Figure 23 Pure Storage FlashArray //m Array 2 NV RAM Modules Overview

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

Figure 24 Pure Storage FlashArray//m50 Controller 0 and 1 Overview

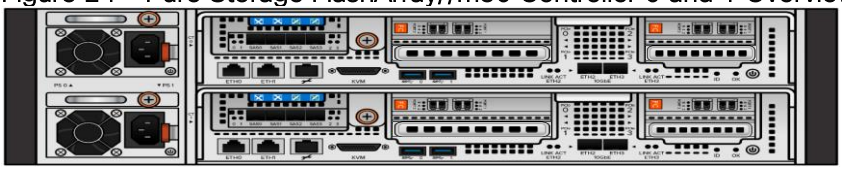

The table below lists the default port services associated with the rear view of the array. Pure Storage FlashArray//m Default Ports Overview

Default Port Services

FA-mxx:

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

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

Figure 25 Pure Storage FlashArray//m Ports Per Controller Overview

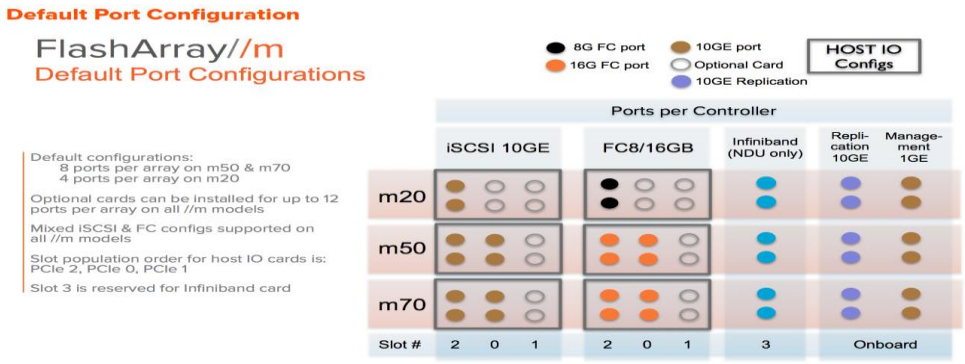

Lastly, the rear of the expansion shelf shown below includes redundant SAS connectivity to each controller and the ability to daisy-chain additional shelfs for added capacity non-disruptively.

Figure 26 Pure Storage FlashAray External Shelf SAS Connectivity for each Controller Overview

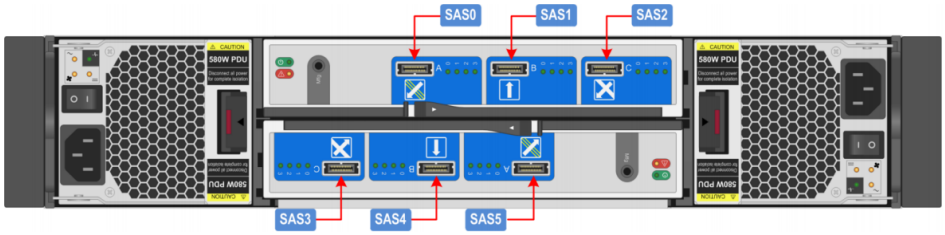

Connectivity to Cisco MDS 9148S

Figure 27 shows the FlashArray//m50 interconnections to the 44TB expansion shelf as well as the 8GB FC connections to the redundantly paired Cisco MDS 9148S switches, yet again taking care to make certain there is no single point of path failure within the Cisco Validated Design. Not shown in the diagram below are power and array/SAN management port connections.

Figure 27 Pure Storage FlashArray//m50 Storage Connectivity to Cisco MDS 9148S Switches Overview

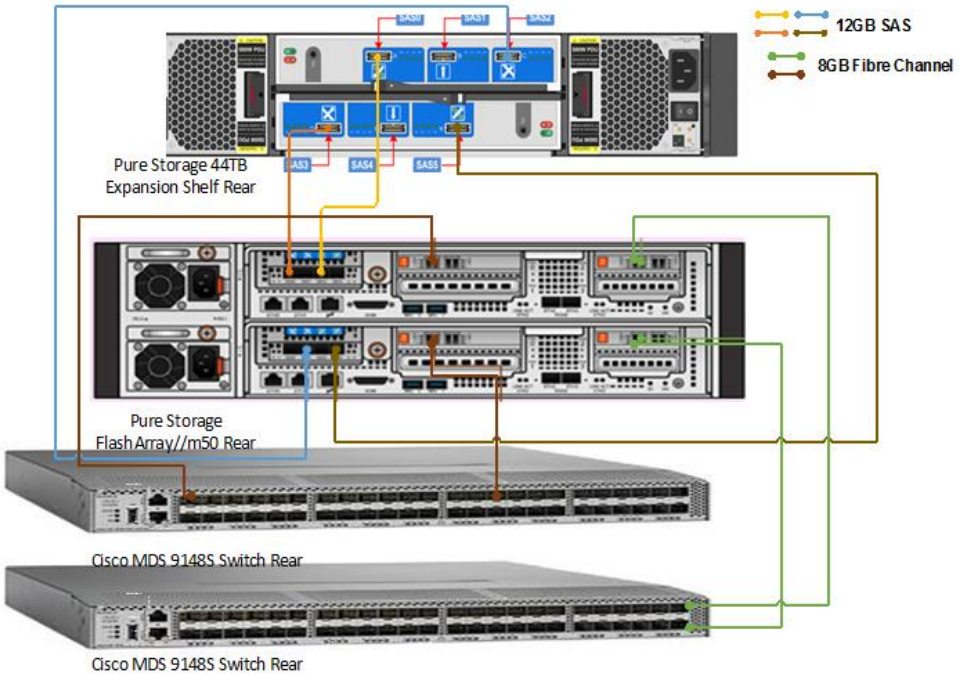

All Pure Storage FlashArray//m systems require a minimum of three static IP addresses for management. One IP address is assigned to each controller Ethernet port (eth0) and a virtual IP is assigned between the two controllers. Initial setup of the array and assignment of these IP addresses is accomplished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

All GUI-based operations shown in this section can also be accomplished via the Purity command line. For the sake of consistency, we will show all configuration steps via the GUI but the array is so simple to use that the entire user manual fits on a double-sided folded business card as shown below in Figure 28. In addition, Purity features a restful API for additional extensibility and usage across a multitude of platforms and languages including PowerShell and the VMware vRealize Automation Suite, among many others.

Pure Storage GUI and SAN Zoning

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

To access the Pure GUI, complete the following steps:

1. To do so, navigate to an IP address assigned for array management from the web browser. The following login screen should appear:

Figure 28 Pure Storage Web GUI or Login

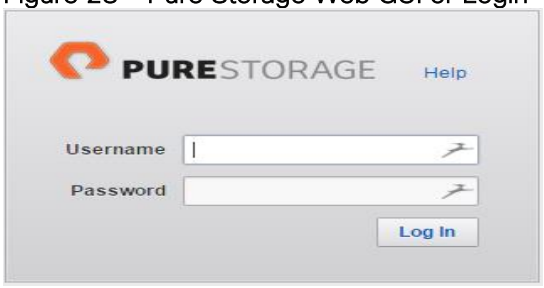

After login, the Pure GUI should appear. Note that a newly deployed array will have 0% space utilization and not be driving any IO.

Data Storage Layout

For this solution, the layout for the Pure Storage FlashArray//m50 Flash Array are listed in the following table. This is the recommended sizes and Performance Policies for best performance in this VDI solution.

Table 13 Layout for the Pure Storage FlashArray//m50

Volumes	Name	Volume Size	VMware / OS Connect / Boot	Performance Policy
Infrastructure Boot 1	Infra1-SANBoot_Lun	10 GB	VMware	VMware ESXi
Infrastructure Boot 2	Infra2-SANBoot_Lun	10 GB	VMware	VMware ESXi
Infrastructure ESXi Servers: (VCSA, Horizon View ConnectionServer, View Replica Server & View Composer Servers, File share, User Profiles, Nexus 1000V VSUM and Login VSI Workload Servers etc.)	Infra_DataStore	4 TB	VMware	VMware ESXi
RDSH-Hosts from 1 to 9 (Total of 9 Cisco B2004 M4 Servers booting from SAN for RDSH Server Roles configuration)	RDSH-BOOT_LUN-01 to 09	10 GB (each server boot lun)	VMware	VMware ESXi
RDSH-DS (Total of 72 Win 2012 R2 Servers for RDSH Session Users)	RDSH-DS	10 TB	VMware	VMware ESXi
VDI-Servers From 1 to 21 (Total of 21 Cisco B2004 M4 Servers for SAN boot)	VDI-BOOT_LUN-01 to 21	10 GB(each server boot lun)	VMware	VMware ESXi
Win7 Linked Clones Non-persistent virtual machines Data store (Virtual Machines from Pool 1 & 2 approx. 1775 VMs equally distributed)	LC01-DS	10 TB	VMware	VMware ESXi
Win7 Linked Clones Non-persistent virtual machines Data store (Virtual Machines from Pool 1 & 2 approx. 1775 VMs equally distributed)	LC02-DS	10 TB	VMware	VMware ESXi

Figure 29 through Figure 32 shows examples of datastores created for data and for SAN boot.

Figure 29 Data stores Created for Data and for SAN Boot

Connected Volumes (2) | Host Ports (2) | Details (0) — 1-2 of 2

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

Figure 30 Datastores Created for Data and for SAN Boot

Connected Volumes (4) | Host Ports (2) | Details (0) — 1-4 of 4

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

Figure 31 Datastores Created for Data and for SAN Boot

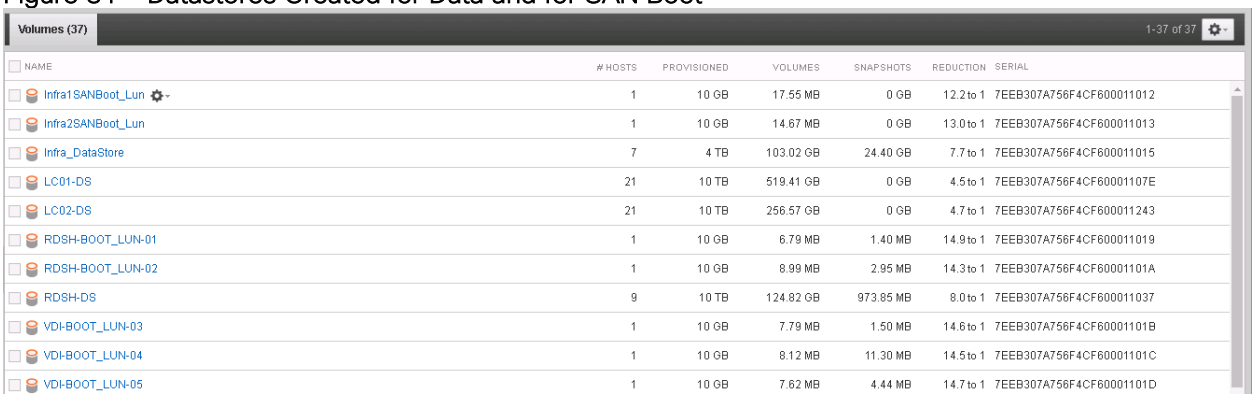

Volumes (37) — 1-37 of 37

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

Figure 32 Pure storage Dashboard on login to Web GUI

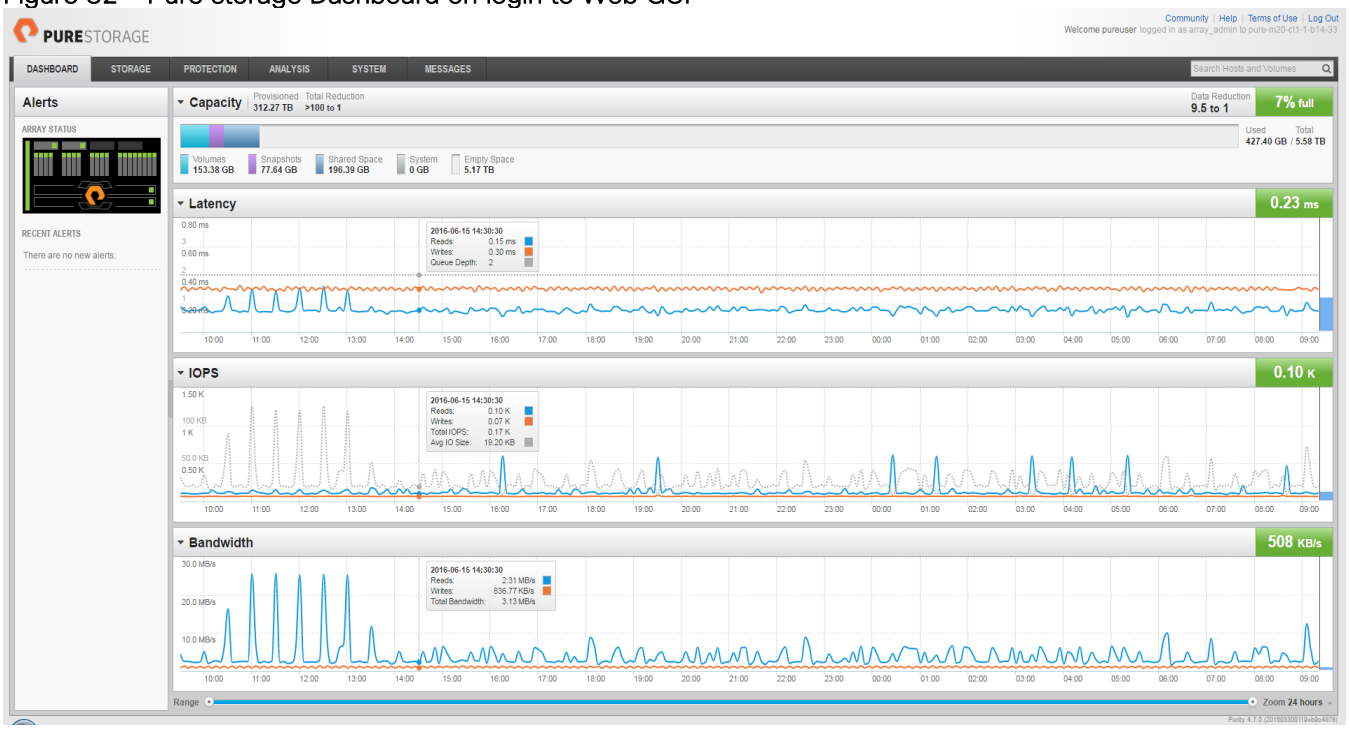

The Pure Storage GUI shows a variety of metrics and status for the array. On the left, the Array Status can be seen which shows a top-level health status of the backplane, NVRam modules, controllers and individual solid state drive bays. To the right of that, the amount of storage space being used, data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

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

1. Go to the System tab which will show the following window. Note that the initial view shows a much more detailed health view for all array components.

Figure 33 System tab to shows the Overall Storage Health

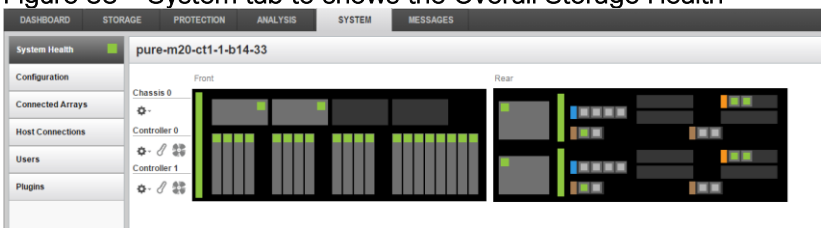

2. Click Configuration to expand the menu and click Networking to bring up the management networking configuration screen:

Figure 34 Pure Storage FlashArray//m50 Networking Information

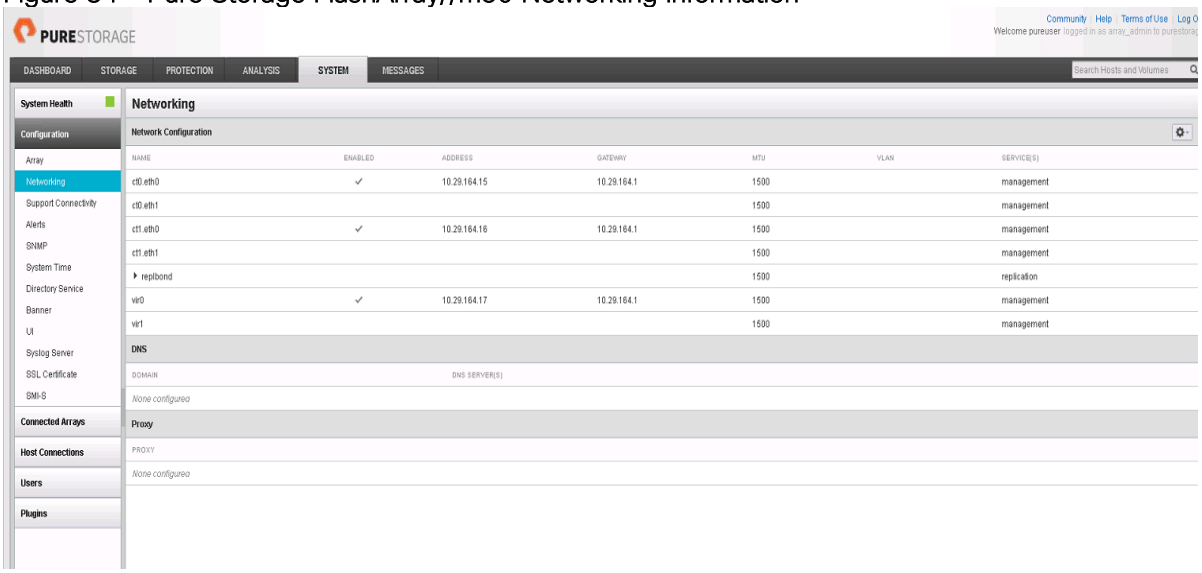

3. To update network information, hover the mouse cursor over the Ethernet port you wish to change and click Edit on the gear that appears. The following screen shots an example of editing a network interface:

Figure 35 Pure Storage FlashArray//m50 Networking Information Editing

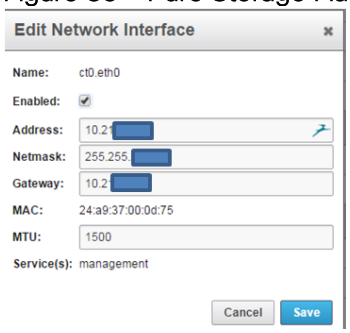

When the management network has been confirmed as being properly setup, the next step is to complete the SAN zoning of the Cisco MCS 9148S switches so that the storage array and Cisco UCS servers are able to communicate. To do the SAN zoning, complete the following steps:

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

Figure 36 ESXi Host Connected View

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

Figure 37 Pure Storage Controllers Configured with WWN Overview

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

Adding a Cisco UCS ESXi Host

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

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

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

Figure 38 Adding ESXi Host

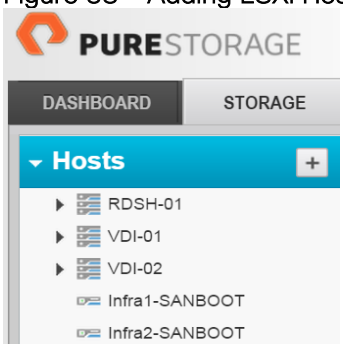

2. Click the '+' sign will spawn a window for you to create single, or multiple hosts. Since we are using multiple hosts in our design we will show the method for creating several hosts at once.

Figure 39 Create a Host Name or Multiple Hosts

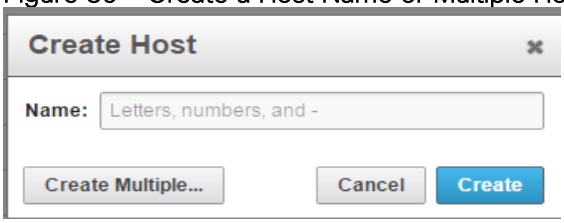

3. After clicking the 'Create Multiple..' option above, the following screen displays which enables batch creation of multiple hosts. In this example, we are creating two hosts.

Figure 40 Overview of one Cisco UCS B200 M4 Host being Created

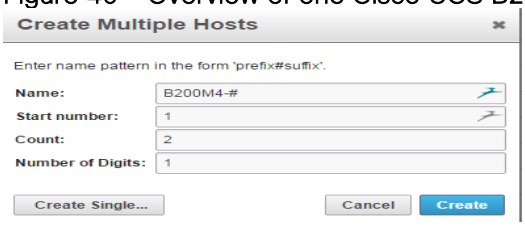

Below you can see that our two hosts have been created and the next step is to connect their WWNs. WWNs can be matched up with specific Cisco UCS servers from within Cisco UCS Manager.

Figure 41 Configuring the ESXi host WWNs from Cisco UCS for the Hosts

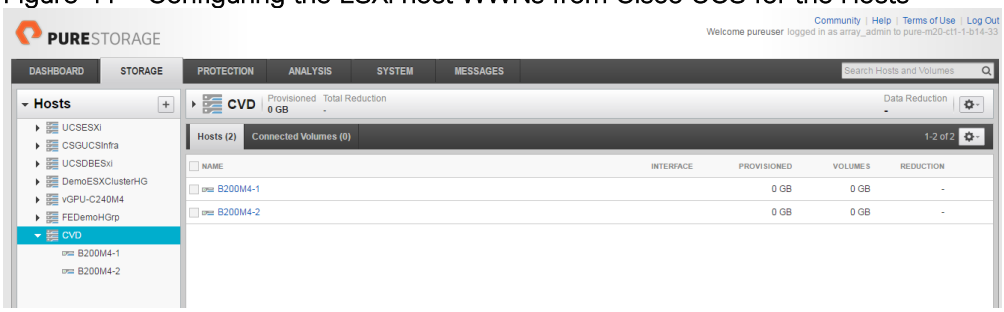

4. To configure the FC WWNs, highlight a host, select the ‘Host Ports’ tab and then select the ‘Configure Fibre Channel WWNs’ button as shown in the below screenshot.

Figure 42 Configuring ESXi Host Ports with WWNs

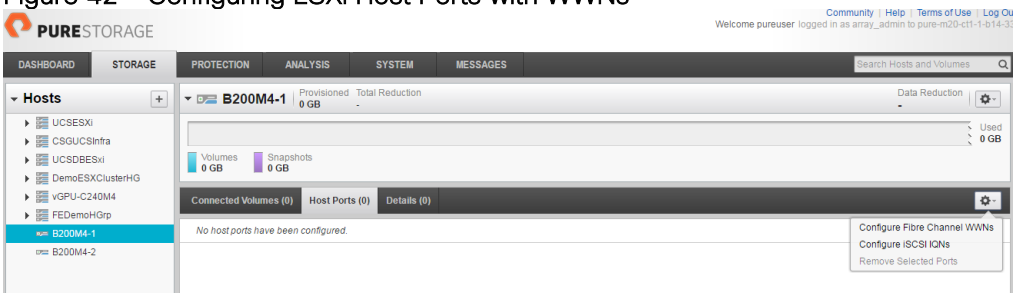

Available WWNs will appear in the left-hand column.

Figure 43 WWNs Configured for Hosts

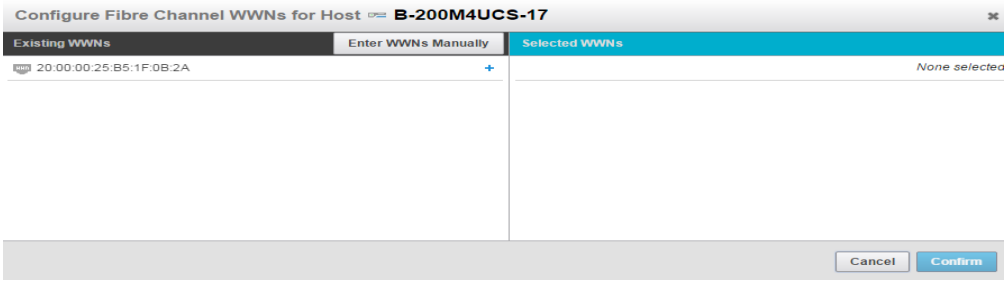

5. Click them to select and move them over to the right-hand column and click ‘Confirm’ when completed. In this Cisco Validated Design, each Cisco UCS host will have two separate vHBAs to provide both performance and resiliency.

Below, you see a properly setup UCS host within the Pure GUI. Repeat the above procedure for host creation and configuring the Fibre Channel WWNs for all Cisco UCS servers.

Figure 44 Overview of a Configured ESXi Host

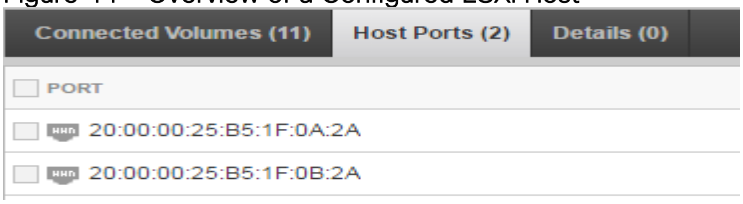

6. When the hosts have been created and the WWNs have been assigned to them, then click the ‘Hosts’ button again and this time select ‘Create Host Group’.

Figure 45 Creating a Host Group

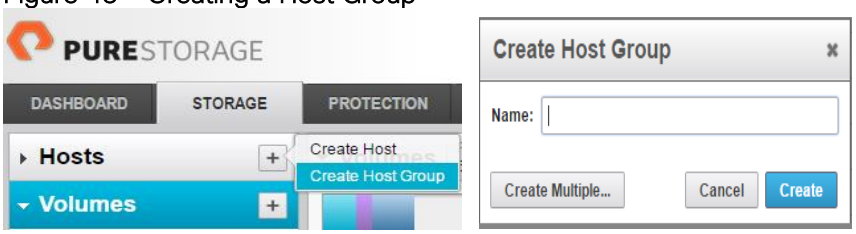

7. The next step is to add the two example hosts we previously created to this newly created Host Group. To accomplish this, highlight the Host Group, click the gear to the right and select 'Add Hosts':

Figure 46 Add Host to the Host Group Created

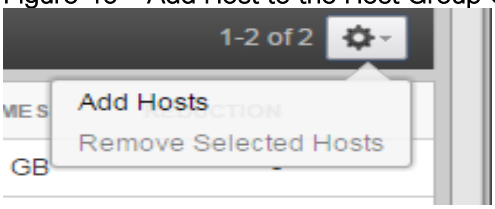

8. The following screenshots display; click the hosts you wish to add to the Host Group on the left column in order to move them to the right. When all hosts have been included, click 'Confirm.'

Figure 47 ESXi Hosts to be added to the Host Group

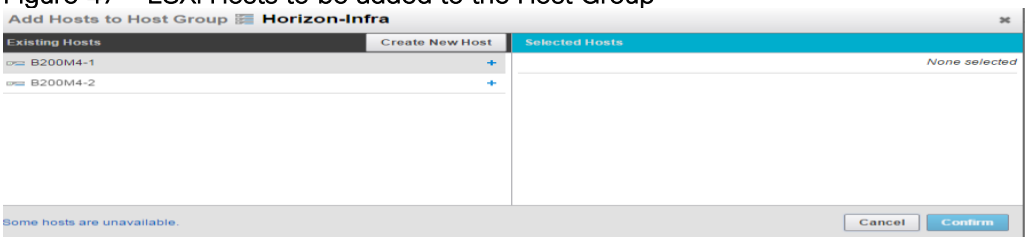

Figure 48 ESXi Host added to the Host Group

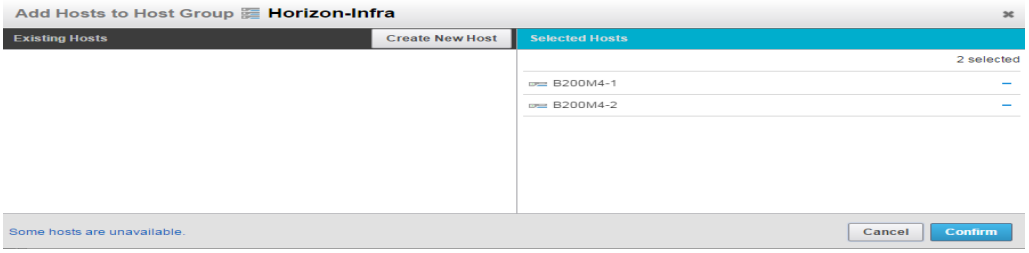

The following screenshot shows the populated Host Group with two hosts as members.

Figure 49 ESXi Hosts have been added to the Host Group

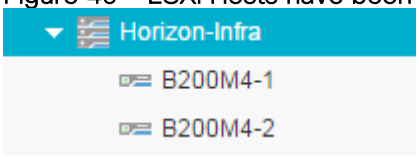

Repeat this procedure for all Host Groups that are required. For this Cisco Validated Design we created three separate Host Groups for the following components:

- VMware Horizon Infrastructure (xx Hosts)
- RSDH Hosts (xx Hosts)
- VMware Horizon Linked-Clones (xx Hosts)

With the Hosts and Host Group setup completed, the array is now ready to serve VMware Horizon desktops and the associated infrastructure.

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

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

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

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

Figure 50 Pure Storage GUI

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

Figure 51 Select Volume

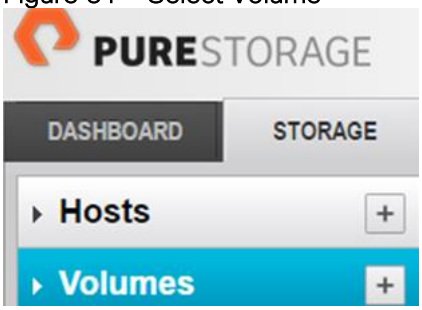

3. Name and size the Volume to complete its creation:

Figure 52 Provide a Name to the Volume and Size

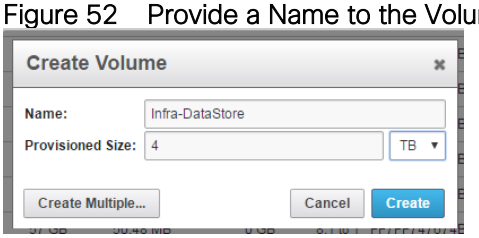

4. Select the newly created Volume under the Volumes collapsible menu on the left-hand side of the GUI and click on the gear icon to connect it to host(s) or host group(s):

Figure 53 Select the Volume Created to be Associated with the Host

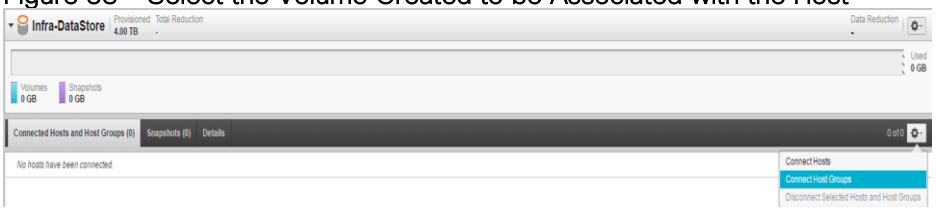

5. Click the appropriate Host Groups that require access to the Volume so that they are under the 'Select Host Groups' list and then click Confirm when they have been added:

Figure 54 Add Host Group or Group Hosts Configured for the Volume

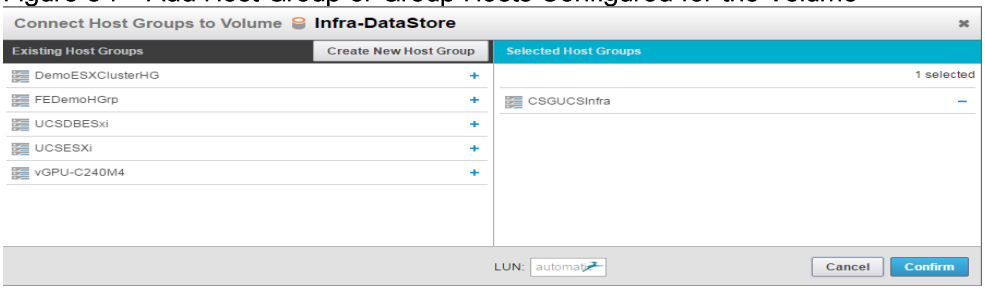

6. Storage setup has been completed and the datastore now needs to be added from within vSphere. From within the vSphere client, select a host from the Host Group that was connected to the volume in the previous step, go to Related Objects and then Datastore and click the icon to add the new datastore:

Figure 55 vSphere GUI to Create the Datastore

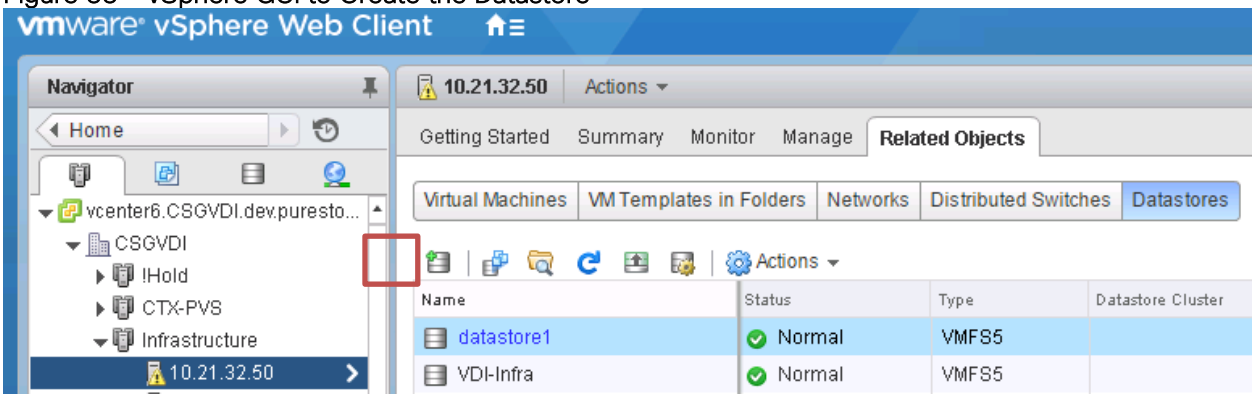

7. Select the VMFS option and click Next.

Figure 56 Select VMFS Datastore

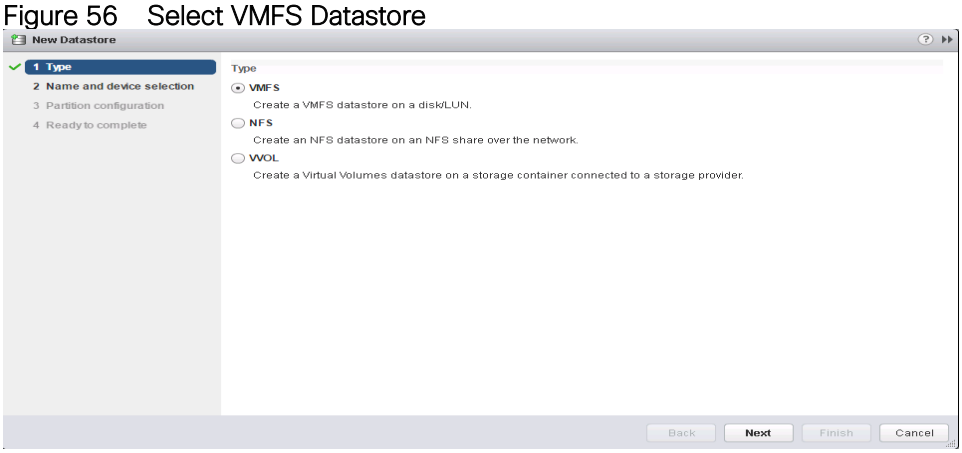

8. Name the datastore and select the appropriate volume from the list shown. It might be necessary to rescan the vHBA on the host in order to see a newly created datastore.

Figure 57 Name and Device Selection for the Datastore

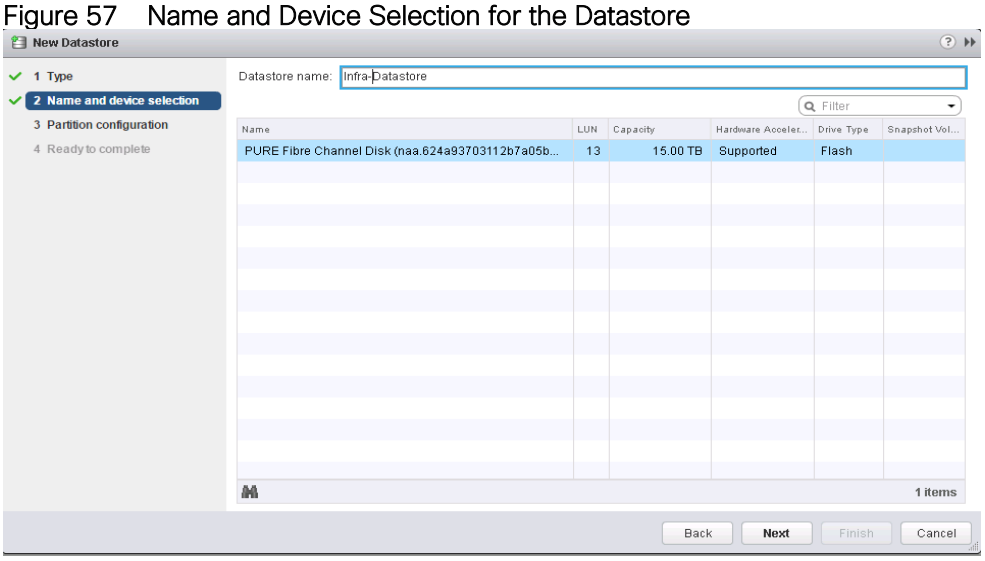

9. Leave the partition configuration options at default values and click Next.

Figure 58 Configuration Parameters for Datastore

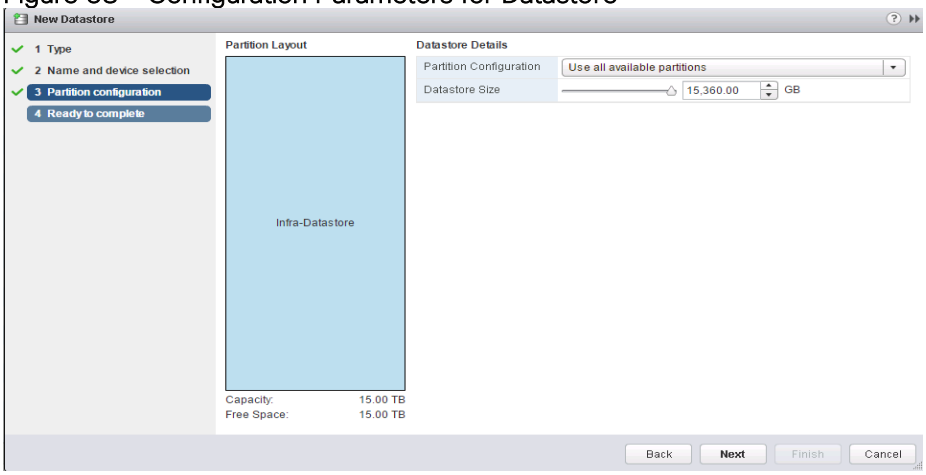

10. Click Finish to complete the datastore setup.

Figure 59 Datastore is created

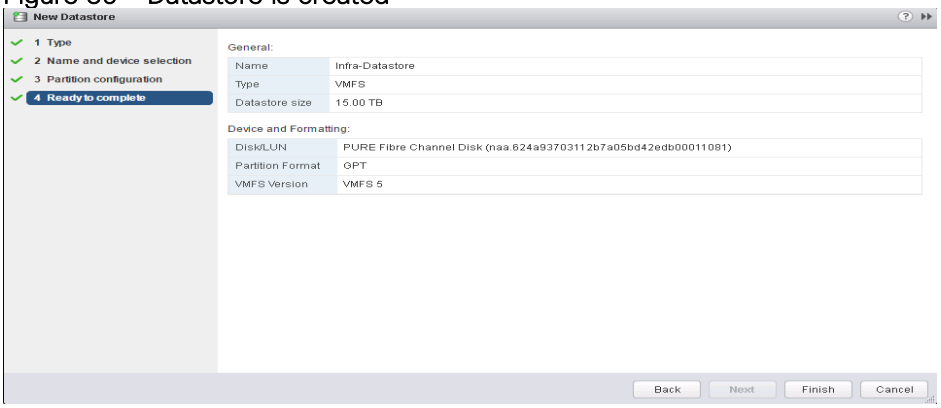

Repeat this process for each of the datastores listed in the Datastore Table 14.

Datastore Snapshot on Pure Storage Array

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

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

Table 14 Example of Data stores Created on Pure Storage for Snapshot

VDI Component	# of Data stores	Datastore Size	Snapshot Schedule	Snapshot Rationale
Horizon View Infra-structure Datastore	1	10 TB	Create snapshot every day, then retrain 1 snap-shot per day for 5 days	Core infrastructure should be backed up regularly
RDSH Datastore	1	10 TB	Create snapshot every 2 days, then retain 1 snap-shot per day for 2 more days	Given the level of customiza-tion required for each RDS 2012R2 server, occasional backups needed
Linked-Clone Datastore	2-4*	5TB	No snapshots	Non-persistent VMs should be restored from parent VM on Infrastructure Datastore if LUN accidently destroyed
ESXi Boot LUNS	1 per UCS host	10GB	Create a snapshot on source every 5 days then retain 1 per day for 1 more day	Not much change on ESXi boot LUNs, hence occasional backups only

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

To setup a snapshot schedule, complete the following steps:

1. Click the 'Protection' tab.

Figure 60 Pure GUI to Create a Snapshot Protection for Host or Host Group

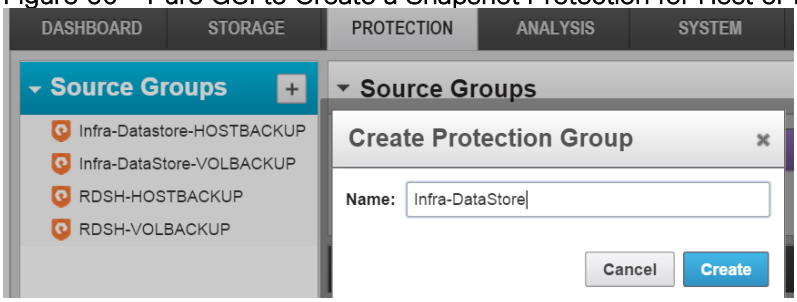

2. Using the Snapshot schedule for the Infrastructure datastore, assign the following values to create the snapshot schedule and retention policy:

Figure 61 Snapshot Schedule

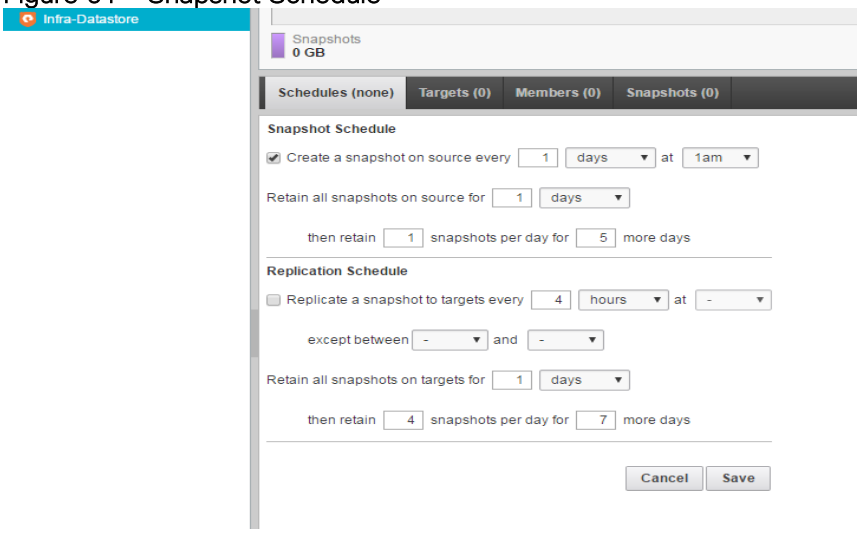

3. When the Snapshot and retention policy has been defined, add the Datastore to the policy by clicking the 'Members' tab of the Source Group:

Figure 62 Add Datastore or Volume to the Snapshot Policy Configured

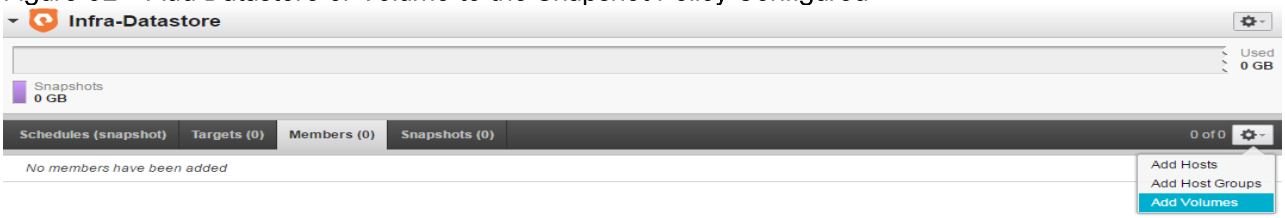

4. Click the Volume(s) you want to include in the Snapshot policy to move them to the right. Click Confirm when completed.

Figure 63 Adding Volumes to the Snapshot Policy Configured

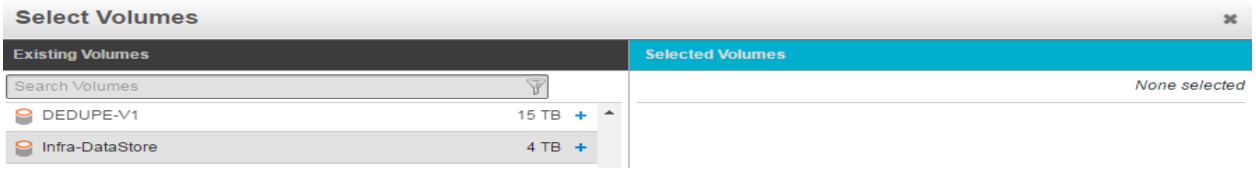

5. After clicking Confirm, the Datastore should be listed as a Snapshot policy member:

Figure 64 Infra-Datastore Volume Applied to the Newly Created Snapshot Retention Policy for Backup

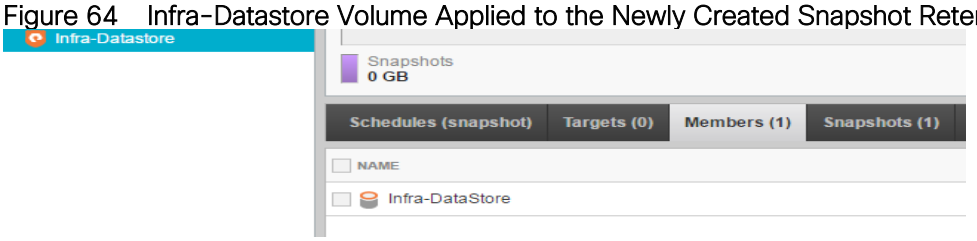

Repeat these steps for the RDSH and ESXi Boot LUNs Data stores. It's recommended to include all ESXi boot LUNs in a single Snapshot policy for simplicity.

ESXi Best Practice Configuration

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

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

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

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

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

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

A VDI user profile share was built using Windows Server 2008 R2 to closely mirror a production environment where user profiles are stored in a network location. The user profiles are created for the Login VSI users who sign in during the automated Knowledge Worker benchmark test runs. The server was hosted on the infrastructure datastore where our snapshot policies were applied for backing up the system.

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

VDI- User Profiles

Figure 65 Profile Share for the Users Configured on the Pure Storage

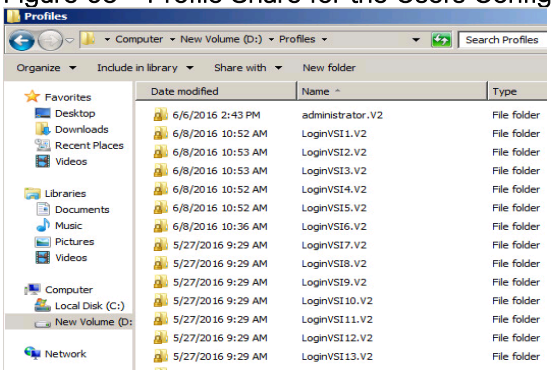

Figure 66 RDSH Session User Profiles

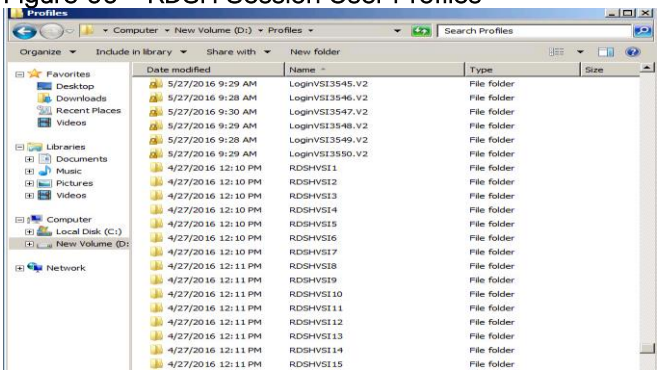

Configure MDS 9100 Series

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

1. 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

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

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

Figure 67 VSAN 20 configured for Fabric A

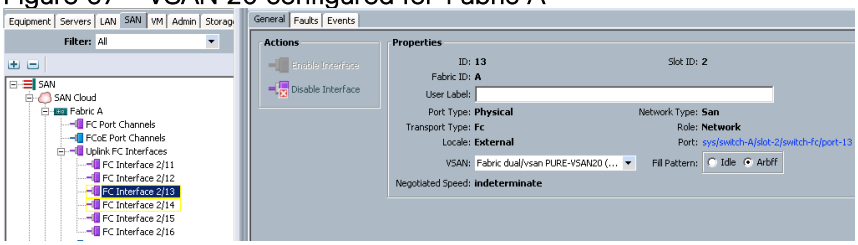

Figure 68 VSAN 30 configured for Fabric B

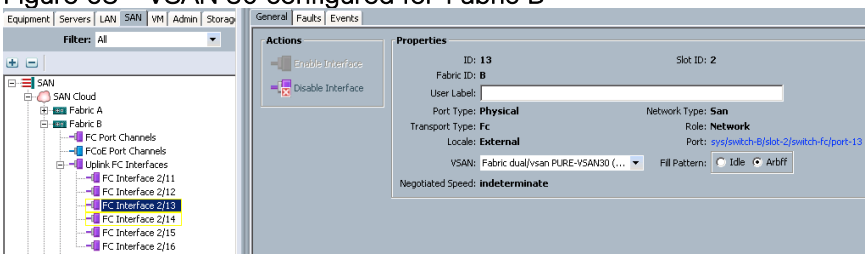

Physically, the Fabric Interconnects extended ports 13 and 14 run to the MDS switch ports 1 and 2.

Figure 69 MDS Switch VSAN Configuration Connectivity

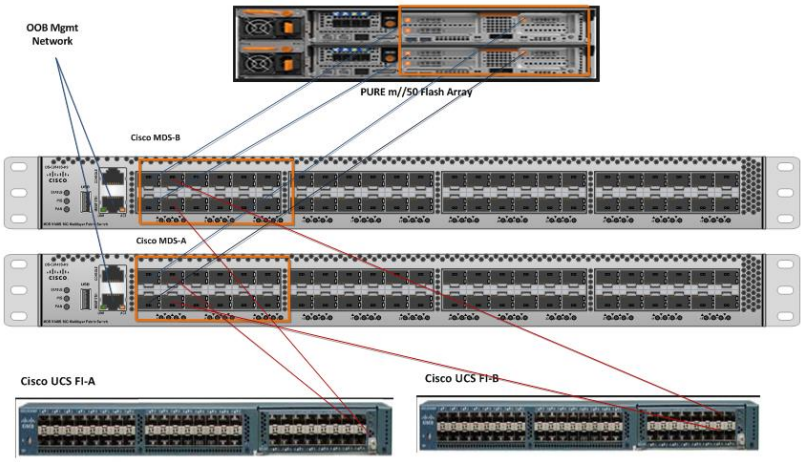

After the ports and port channels are configured, the next steps are to configure the zones and Active Zoneset Database in the MDS switches. The below commands show how to add in a single host on both MDS A and B. You will need to configure all hosts that will access the Pure Array in these commands.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

9148S-A

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

9148S-B

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

```
Exit
Zoneset activate name VDI-Infra-B vsan 30
Zone commit vsan 30
Exit
Copy running-config startup-config
```

Install and Configure VMware ESXi 6.0 U1a

VMware ESXi 6.0

This section provides detailed instructions for installing VMware ESXi 6 Update1 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6 Update1

To download the Cisco Custom Image for ESXi 6 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 CiscoCustomImage6.0.
4. Click Download Now.
5. Save it to your destination folder.

This ESXi 6.0 Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.3.0.7; fNIC: 1.6.0.25

Install ESXi

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the PURE LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.

7. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.

8. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.

The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, simply click Yes to unmap the image.

10. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host.

To configure the ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.

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

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

4. Select the VLAN (Optional) option and press Enter.

5. Enter the VLAN in-band management ID and press Enter.

6. From the Configure Management Network menu, select IP Configuration and press Enter.

7. Select the Set Static IP Address and Network Configuration option by using the space bar.

8. Enter the IP address for managing the first ESXi host.

9. Enter the subnet mask for the first ESXi host.

10. Enter the default gateway for the first ESXi host.

11. Press Enter to accept the changes to the IP configuration.

12. Select the IPv6 Configuration option and press Enter.

13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.

14. Select the DNS Configuration option and press Enter.

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

15. Enter the IP address of the primary DNS server.

16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

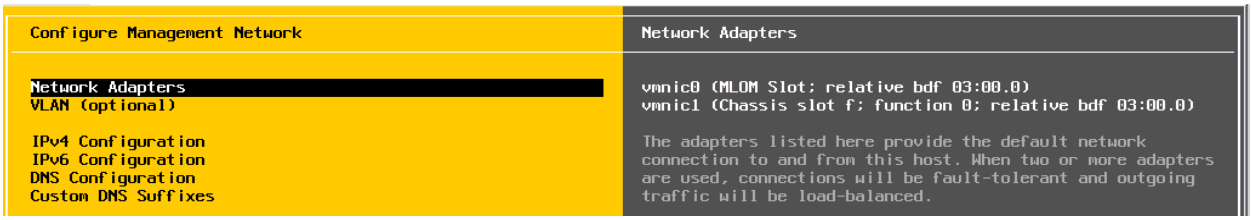

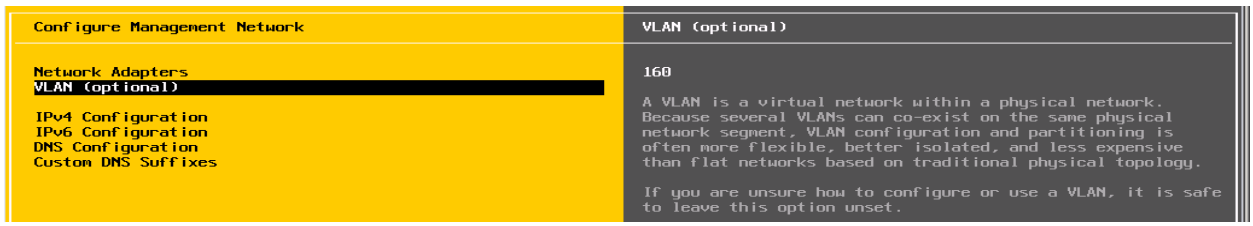

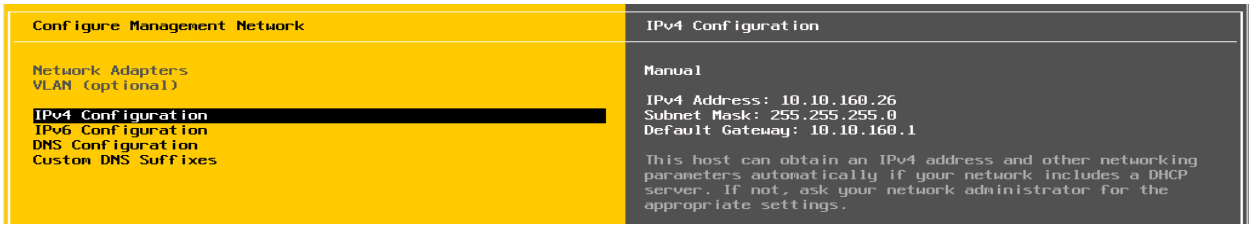

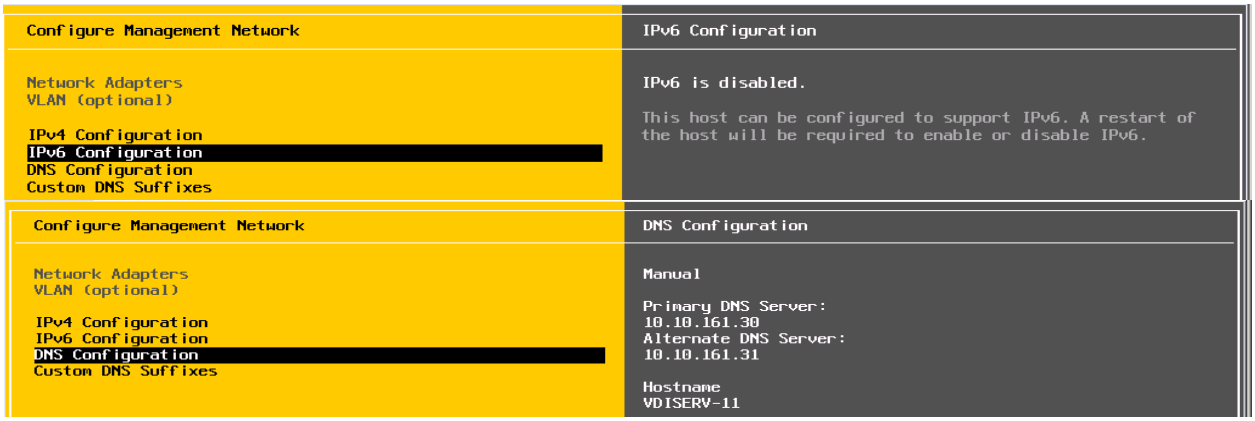

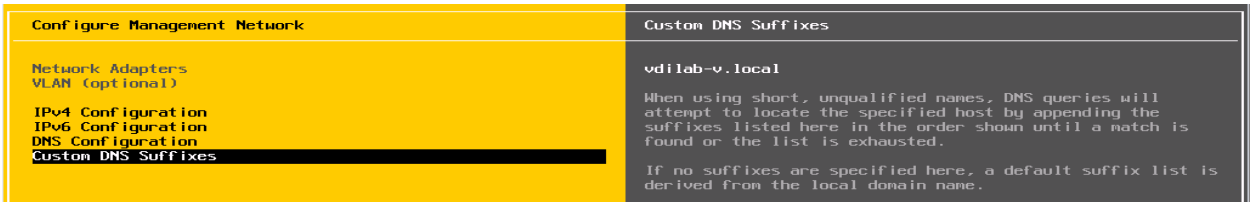

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-01 management IP address.
2. Download and install the vSphere Client.

This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 6

To download VMware vSphere CLI 6, complete the following steps:

1. Click the following link VMware vSphere CLI 6.0
2. Select your OS and Click Download.
3. Save it to your destination folder.
4. Run the VMware-vSphere-CLI.exe
5. Click Next.

6. Accept the terms for the license and click Next.
7. Click Next on the Destination Folder screen.
8. Click Install.
9. Click Finish.

Install VMware vSphere CLI 6.0 on the management workstation.

10. Log in to VMware ESXi Hosts by Using VMware vSphere Client.

ESXi Host VM-Host-01

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
2. Enter root for the user name.
3. Enter the root password.
4. Click Login to connect.

Install and Configure VMware vCenter Appliance

To build the VMWare vCenter VM, complete the following steps:

1. From the vSphere 6 download page on the VMware Web site, download the vCenter ISO file for the vCenter Server appliance onto your system.
2. Open the vSphere ISO via Windows Explorer and double-click the vcsa-setup.htm file.

Figure 70 Install VCSA Appliance from Installer

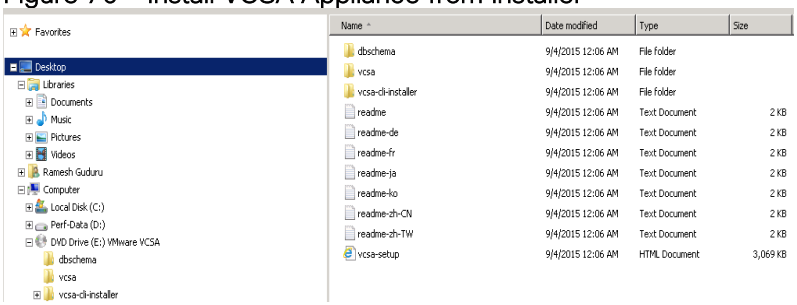

A browser will open with an option to Install.

3. Click Install.

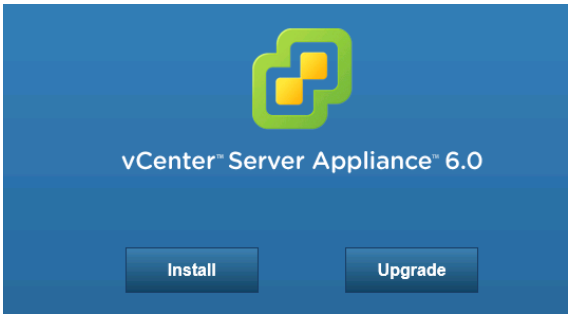

4. Follow the onscreen prompts. Accept EULA.

Figure 71 Accept EULA Agreement

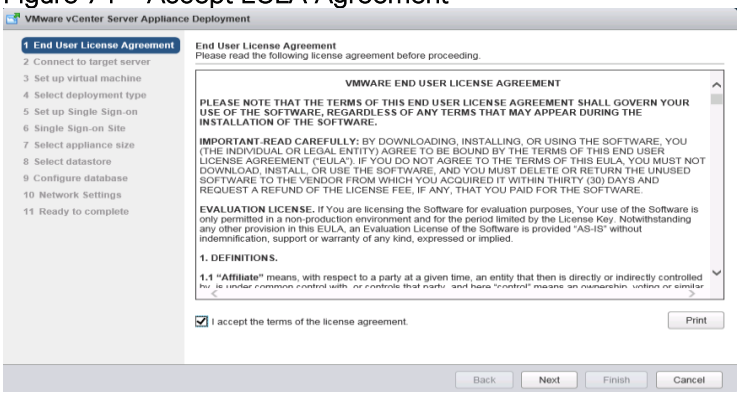

5. Enter the IP of the ESXi host the vCenter Appliance will reside. Click Next.

Figure 72 Provide Host IP or FQDN and User Name, Password Credentials of the Host to Connect

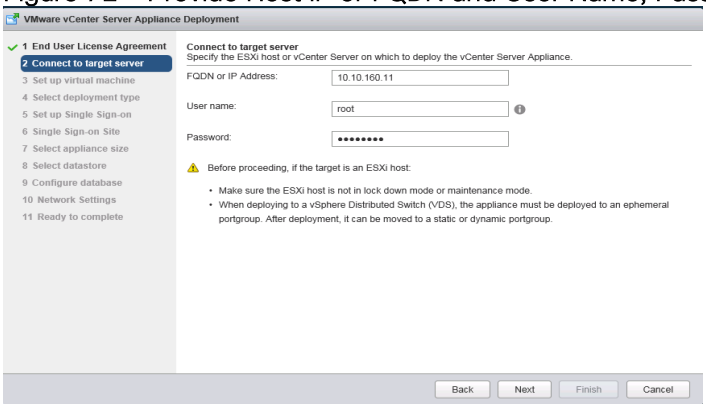

6. Click Yes to accept Certificate Warning.

Figure 73 Click Certificate

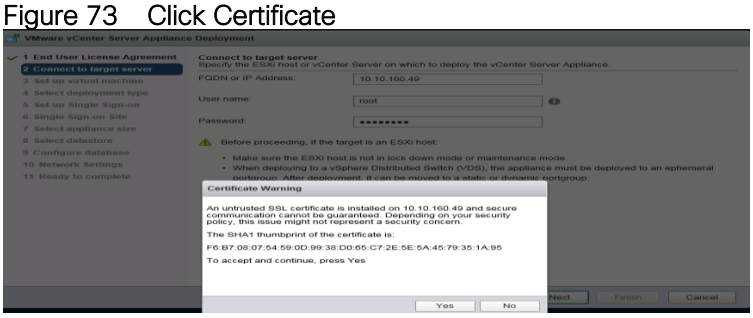

7. Provide a name for the vCenter appliance, then click Next to continue.

Figure 74 Provide a Name and Password for VC Appliance to Configure

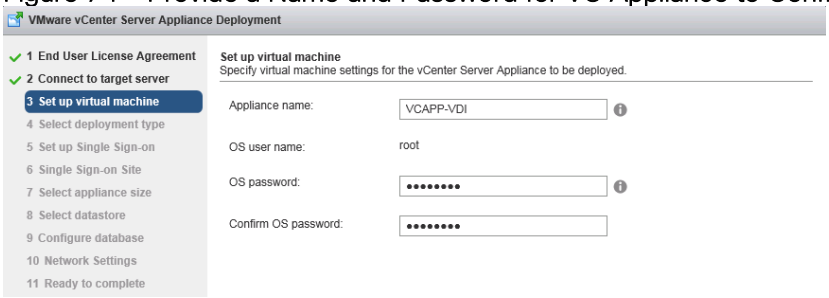

8. Select Install vCenter Server with and Embedded Platform Services Controller (unless your environment already has a PSC).

Figure 75 Select Platform Services Controller Applicable

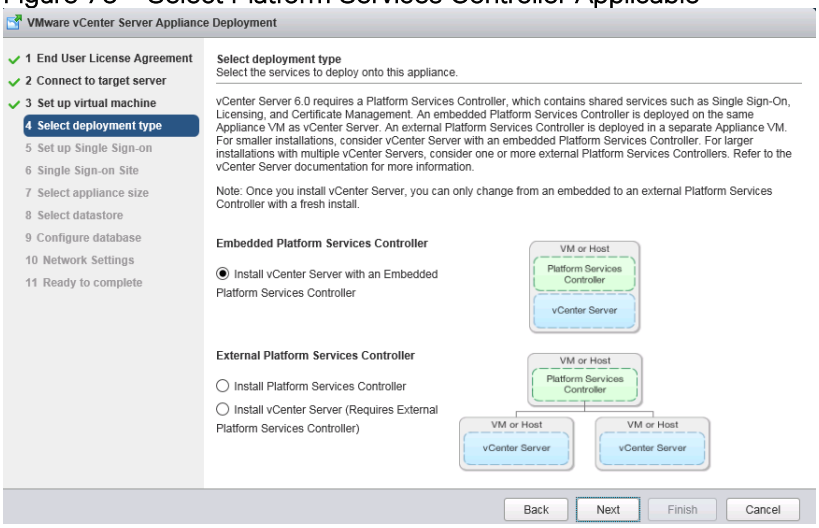

9. Create a new SSO domain (unless your environment already has and SSO domain. Multiple SSO domains can co-exist).

Figure 76 Provide Single Sign On Password and Site Name Credentials

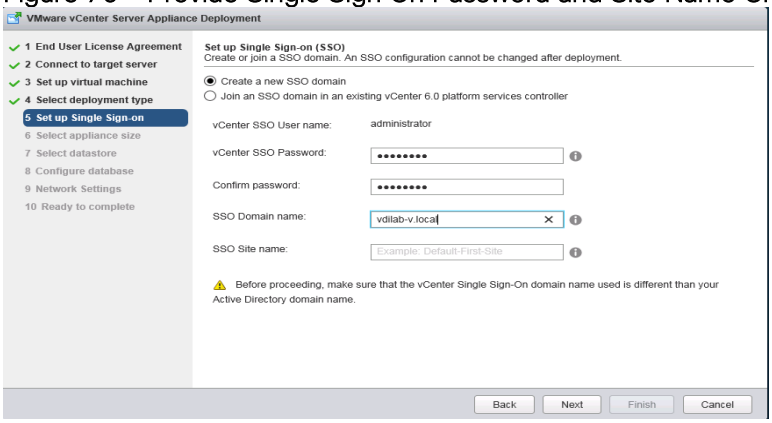

10. Select the proper appliance size for your deployment. In our study, Large was selected.

Figure 77 Select the Appropriate Appliance Size

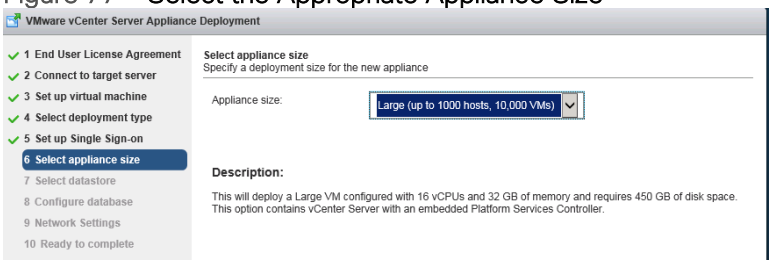

11. Select the Data store

Figure 78 Select the Data store to Install Appliance

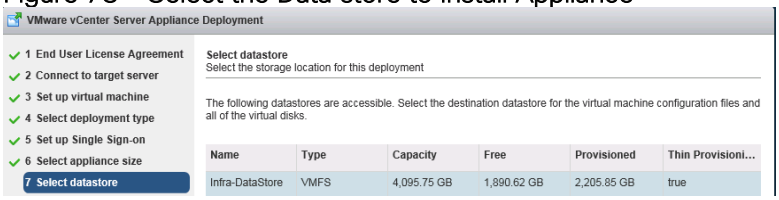

12. In our study we used the embedded PostgreSQL database.

Figure 79 PostgreSQL

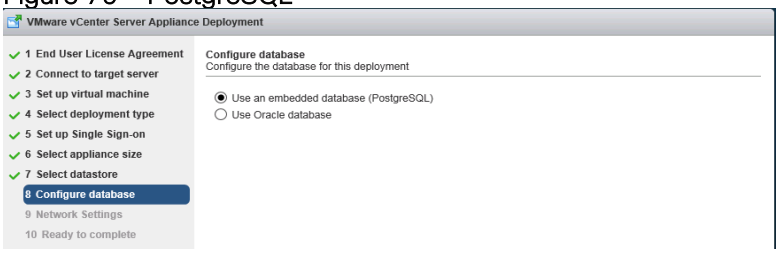

13. Enter Network Settings for appliance.

It is important to note at this step that you should create a DNS A record for your appliance prior to running the install. The services will fail to startup and your install will fail if it cannot resolve properly.

Figure 80 Provide the Necessary Network Gateways and DNS Server Information

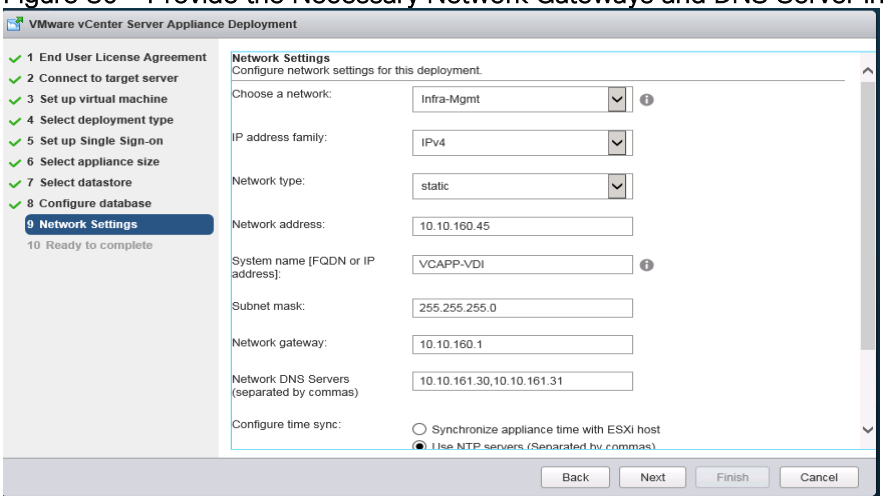

14. Review the Install Settings and click Finish.

15. Click Next to complete installing appliance.

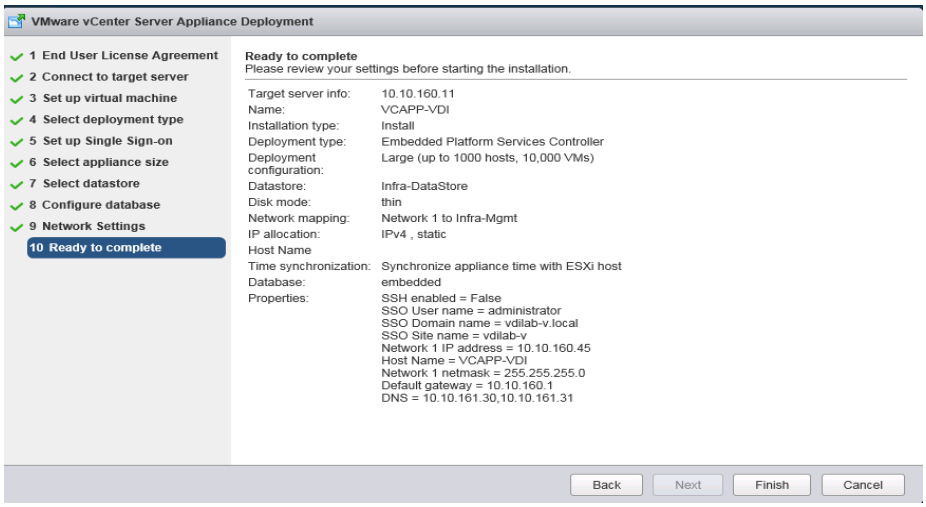

16. When your install completes successfully, you can now login to your Web Client and begin adding hosts and configuring clusters.

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

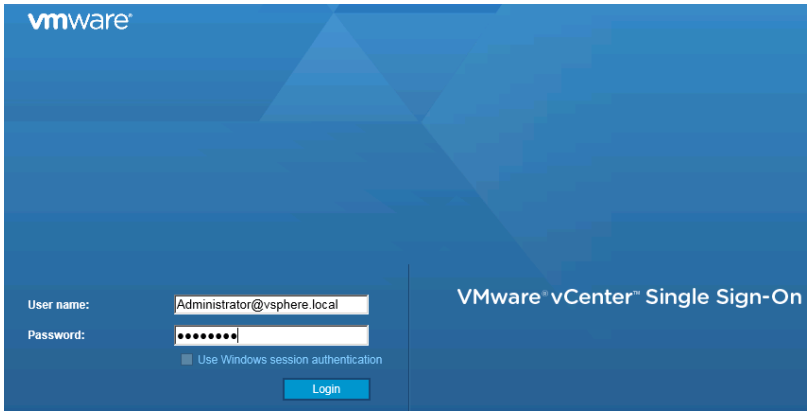

18. Log into the vSphere Web Client.

Figure 81 Login using IP Address of the Appliance and Download vSphere

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

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

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

22. Click the vCenter link on the left panel.

Figure 82 Login vSphere Web GUI

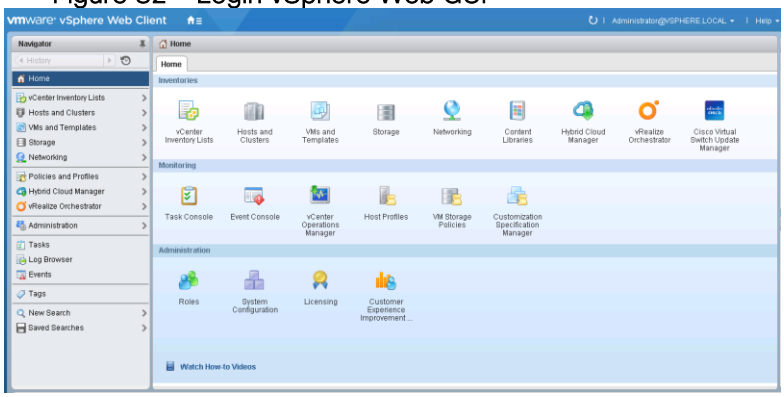

23. Click the Datacenters link on the left panel.

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

Figure 83 Create Data Center, example: VDI-DC

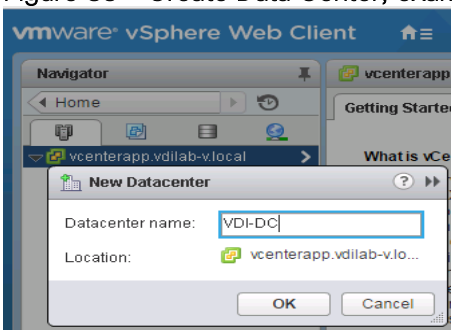

25. Type VDI-DC as the Datacenter name.

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

Figure 84 Create a Cluster

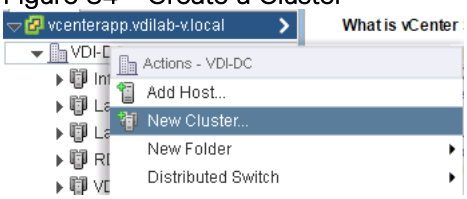

27. Right-click Datacenters > VDI-DC in the list in the center pane, then click New Cluster.

28. Name the cluster Infra.

29. Select DRS. Retain the default values.

30. Select vSphere HA. Retain the default values.

Figure 85 Configure Cluster Specific Setting

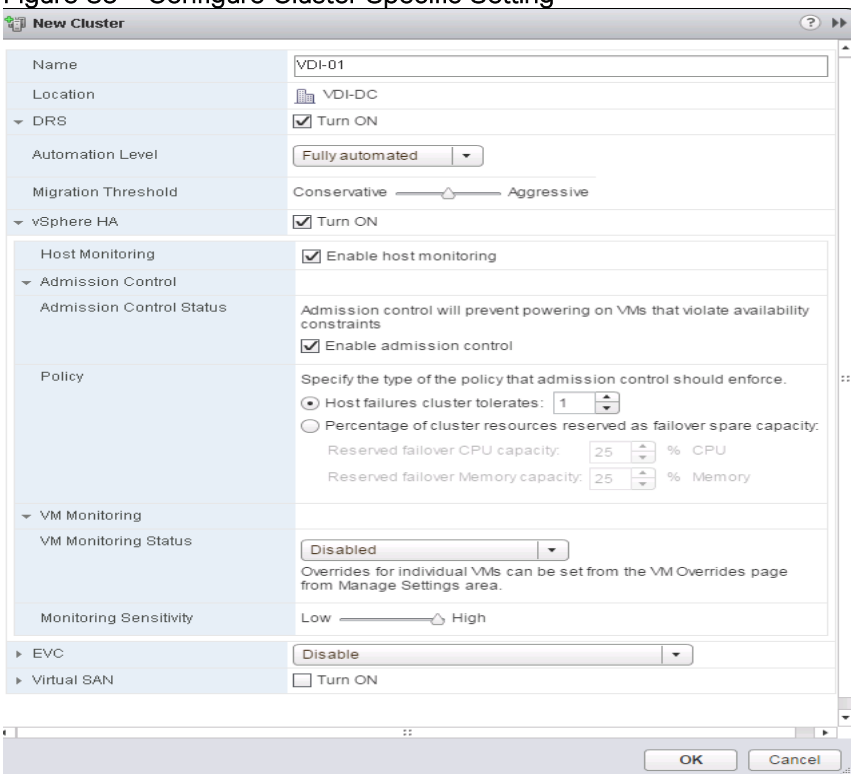

If mixing Cisco UCS B 200 M4 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

31. Click OK to create the new cluster.

32. Click VDI-DC in the left pane.

Figure 86 Add a ESXi Host

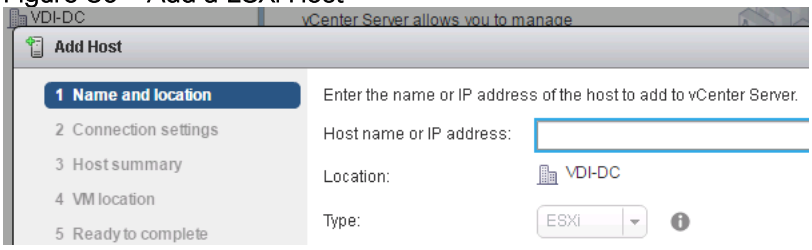

33. Right-click Infra in the center pane and click Add Host.

34. Type the host IP address and click Next.

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

Figure 87 ESXi Host Certificate

36. Click Yes to accept the certificate.
37. Review the host details, and click Next to continue.
38. Assign a license, and click Next to continue.
39. Click Next to continue.
40. Click Next to continue.
41. Review the configuration parameters then click Finish to add the host.
42. Repeat this for the other hosts.

Install and Configure VSUM and Cisco Nexus 1000v

Install Cisco Virtual Switch Update Manager

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

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

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

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

1. Copy the following files to a directory on the Linux machine:
   - Nexus1000v-vsum.1.5.x-pkg.zip image
   - signature.txt file
   - cisco_n1k_image_validation_v_1_5_x script
2. Make sure the script is executable.
   - chmod 755 cisco_n1k_image_validation_v_1_5_x

3. Run the script.
   - ./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
4. Run the script.
   - ./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
5. Check the output. If the validation is successful, the following message displays:

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

Install Cisco Virtual Switch Update Manager

VMware vSphere Web Client

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

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

Figure 88 Select the Cisco Nexus 1000v OVF File to Install

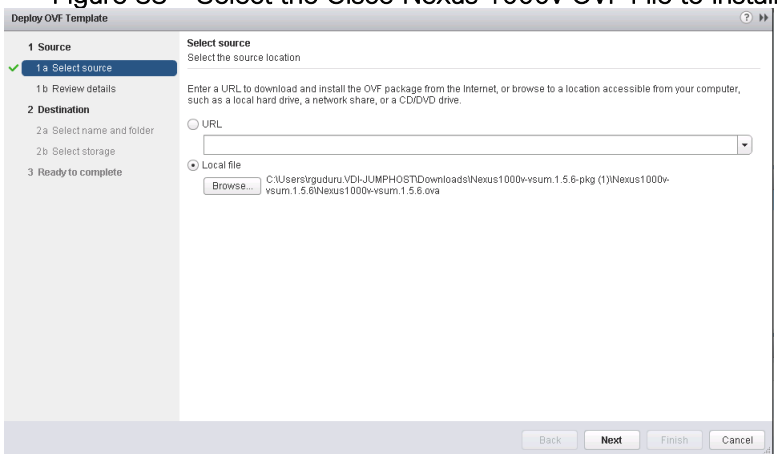

7. Review the details and click Next.
8. Review and click Next.

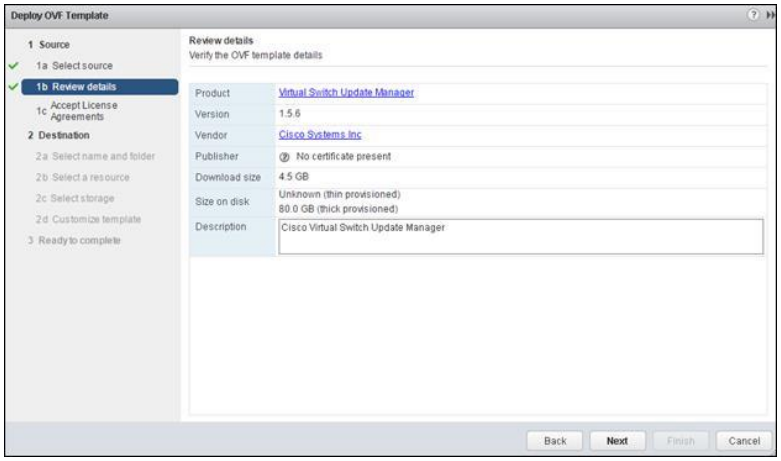

9. Click Accept to accept the License Agreement and click Next.

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

Figure 89 Provide Name for VSM to be Configured

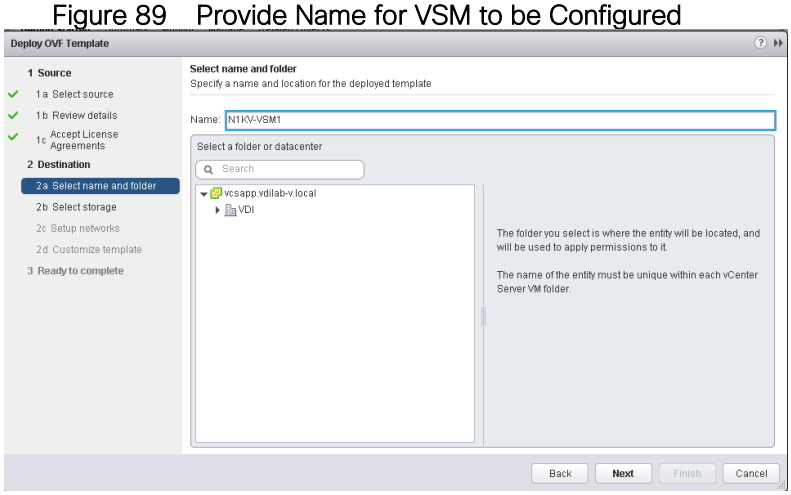

11. Select the Infra cluster and click Next.

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

Figure 90 Select the Datastore

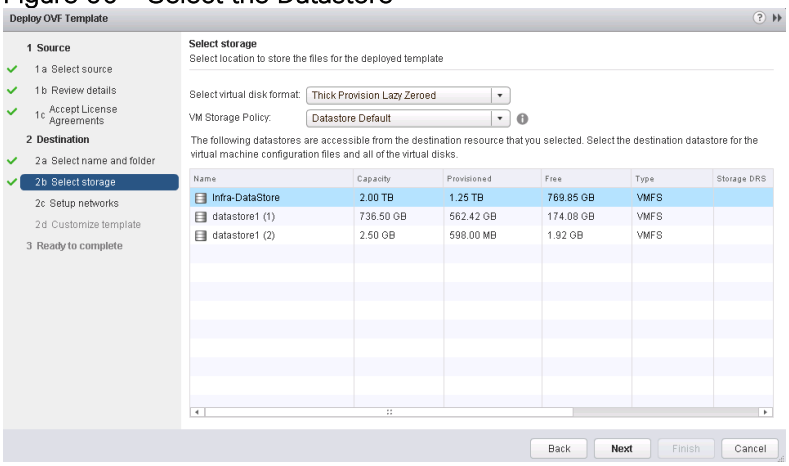

13. Select the MGMT Network and click Next.
14. Fill in the Networking Properties.
15. Expand the vCenter Properties and fil those in.
16. Click Next.
17. Review all settings and click Finish.
18. Wait for the Deploy OVF template task to complete.
19. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.
20. Expand the Infrastructure cluster and select the Virtual Switch Update Manager VM.
21. In the center pane, select Launch Remote Console. If a security warning pops up, click Allow.
22. If a security certificate warning pops up, click Connect Anyway.
23. Power on the Virtual Switch Update Manager VM.
24. When the VM has completely booted up, log out and log back into the VMware vSphere Web Client.
25. Review and click Next to install the click Nexus 1000V.

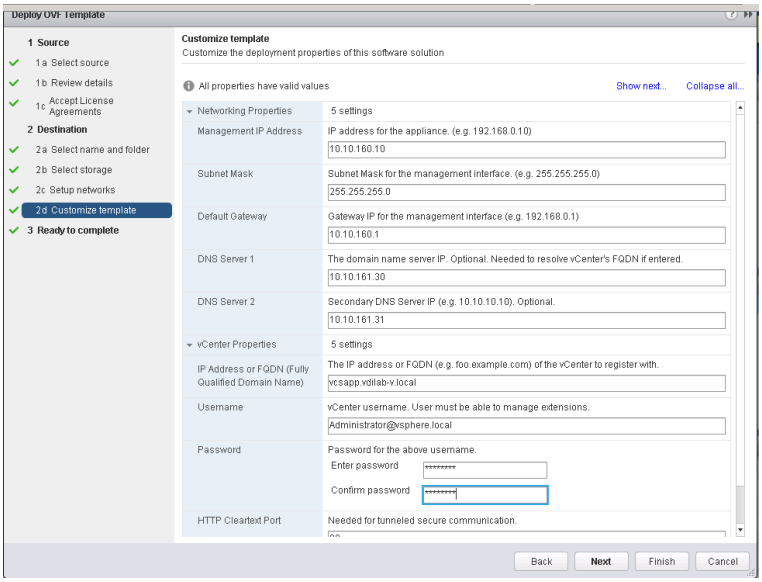

About the Cisco VSUM GUI

The following lists the details of the Cisco VSUM GUI:

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

Figure 91 VMware vSphere Web Client–Home Page

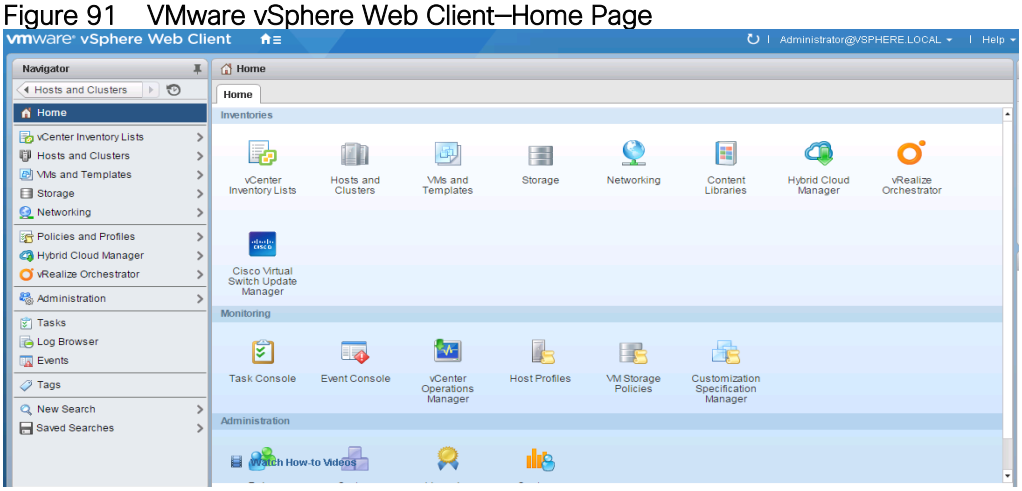

Figure 92 Cisco VSUM–Home Page

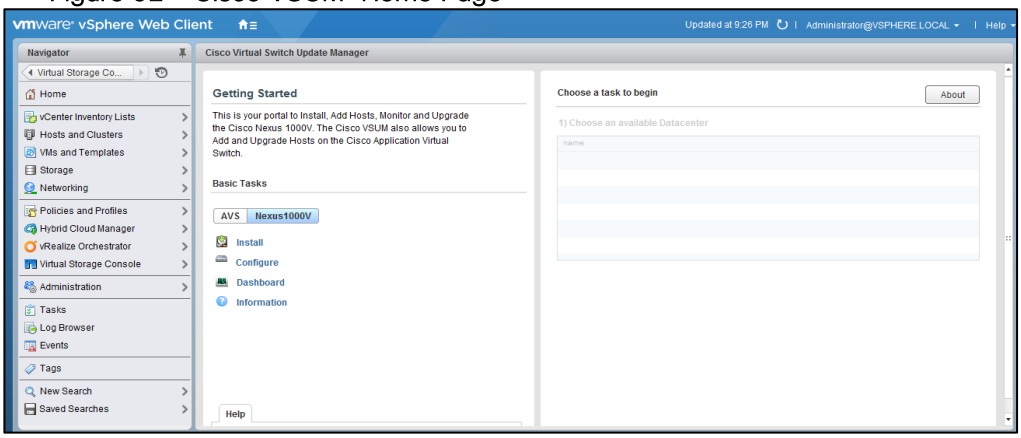

Install Cisco Nexus 1000V using Cisco VSUM

VMware vSphere Web Client

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

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

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

Figure 93 Configuring Virtual Ethernet Modules

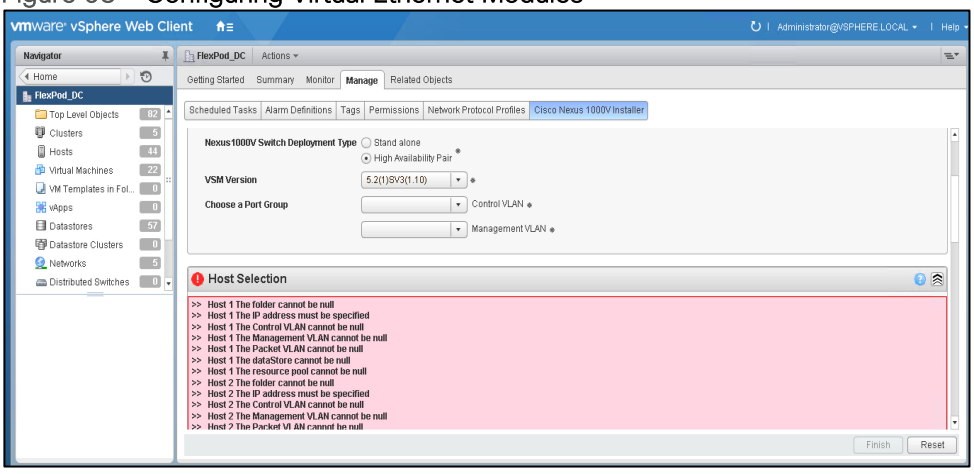

2. In the Nexus 1000v Switch Deployment area, choose I want to deploy new control plane (VSM).

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

4. Choose the control port group for the switch.

5. Choose the management port group for the switch.

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

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

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

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

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

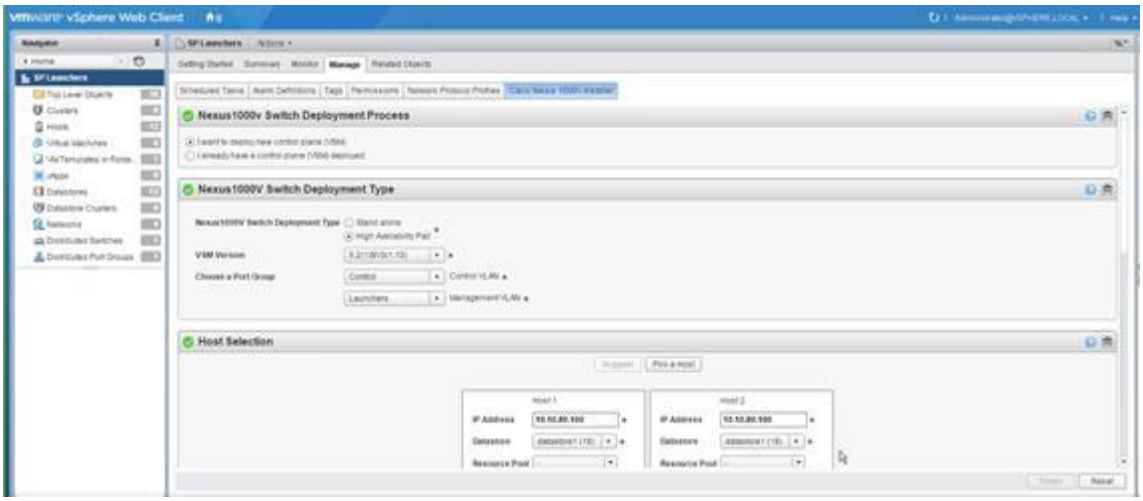

10. In the Switch Configuration area, enter 70 as the domain ID for the switch.

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

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

13. Do not select Default Port Profiles.

14. Enter the Password and Confirm Password for Admin.

15. Provide switch name, password and IP address.

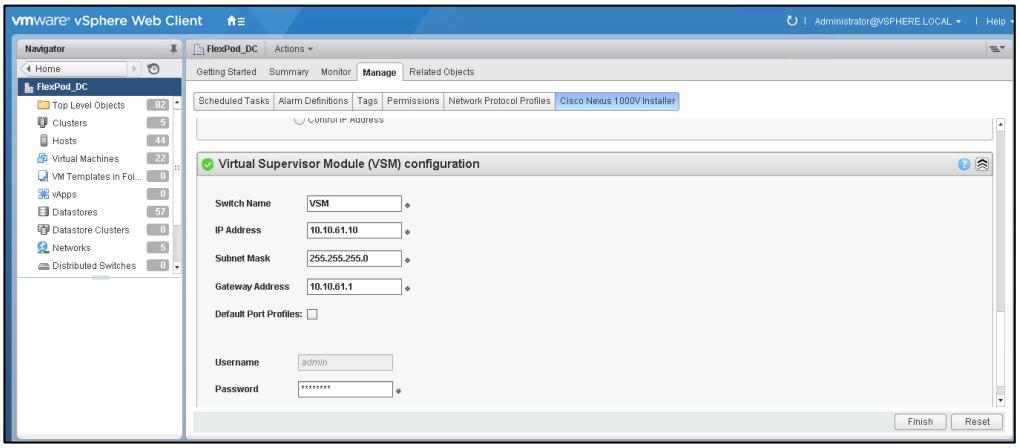

16. Click Finish to install the Cisco Nexus 1000V switch.

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

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

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

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

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

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

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

```
vlan <<var_vdi_vlan_id>> 162
name VDI
vlan <<var_vdi_vlan_id>> 162
name VDI-1
```

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

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

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

To and VMware ESXi hosts, complete the following steps:

1. Back in the VMware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the VDI-DC datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the Infrastructure ESXi Cluster and select one of the Infrastructure Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile.
12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.
13. All VMkernel ports should already have the appropriate checkboxes selected.
14. Scroll down to VM Migration and expand both ESXi hosts.
15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.
16. Click Finish.

The progress of the virtual switch installation can be monitored from the c# interface.

Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

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

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.
2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.
3. In the center pane, select the Manage tab, then select Networking.
4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.

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

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

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

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

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

Figure 94 VSphere Web Client Checking the Uplinks Status

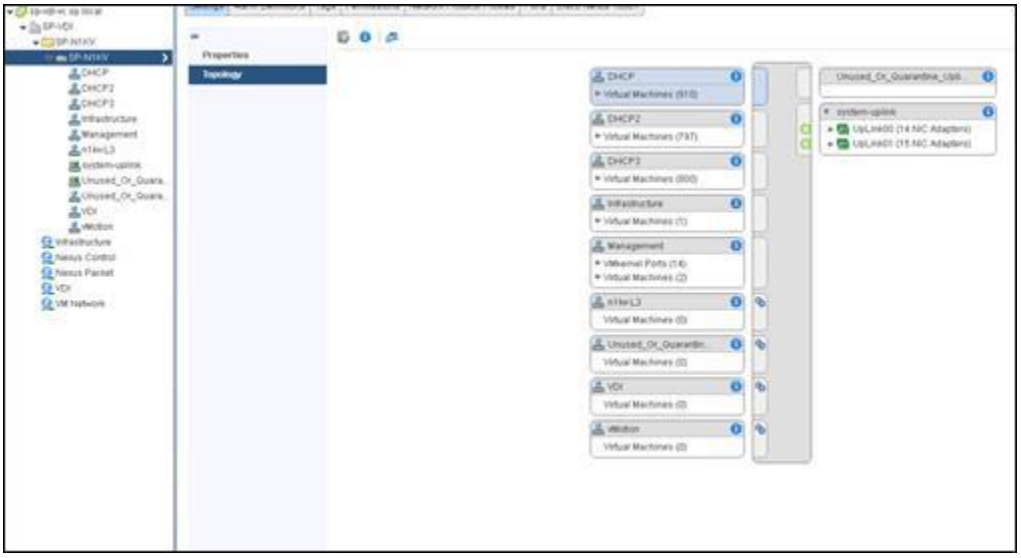

10. Repeat this procedure for the other ESXi host.

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

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

13. Run show to check the connectivity status.

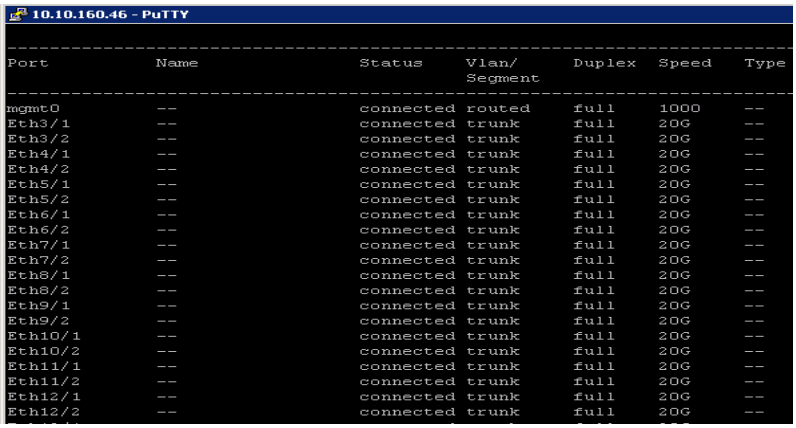

14. Repeat the above steps to migrate the remaining ESXi hosts to the Nexus 1000V.

15. Run: copy run start.

Cisco Nexus 1000V vTracker

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

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

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

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

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

```
config t
feature vtracker
copy run start
show vtracker upstream-view
show vtracker vm-view vnic
```

```
show vtracker vm-view info
show vtracker module-view pnic
show vtracker vlan-view
copy run start
```

Figure 95 vLAN 160 check

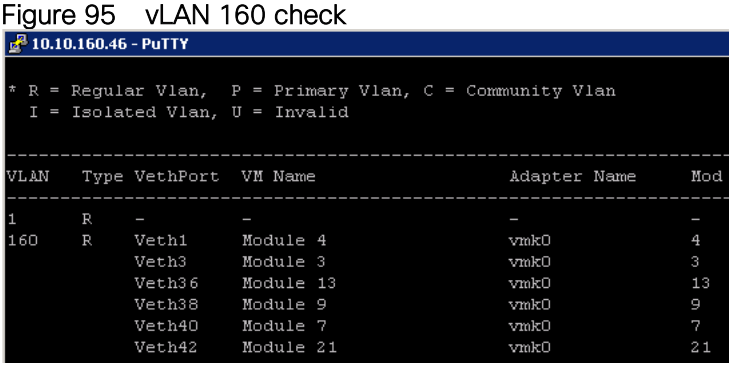

Building the Virtual Machines and Environment for Workload Testing

Software Infrastructure Configuration

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

Installing View Connection Servers and View Composer Server

The prerequisites for installing the view connection server or composer server is to have Windows 2008 or 2012 servers ready. In this study, we have used Windows 2012 server for both View connection server and view composer server.

1. Download the view composer installer from VMware and click install on the View composer Windows server Image. In this study, View Composer 6.2.0.3001314 version used

2. Click to install view composer installer.

3. Click Next to install view composer installer.

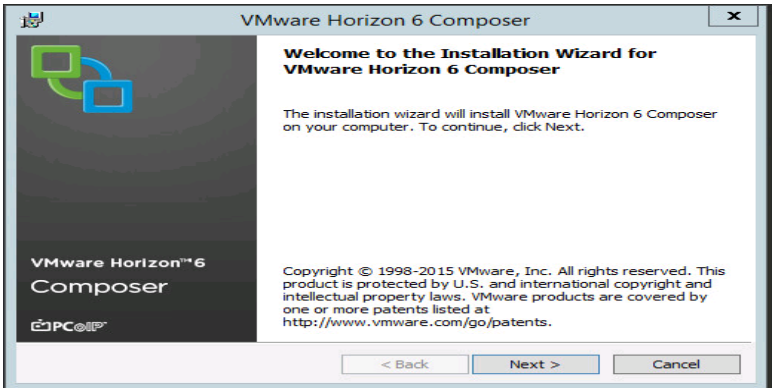

4. Click to accept the EULA.

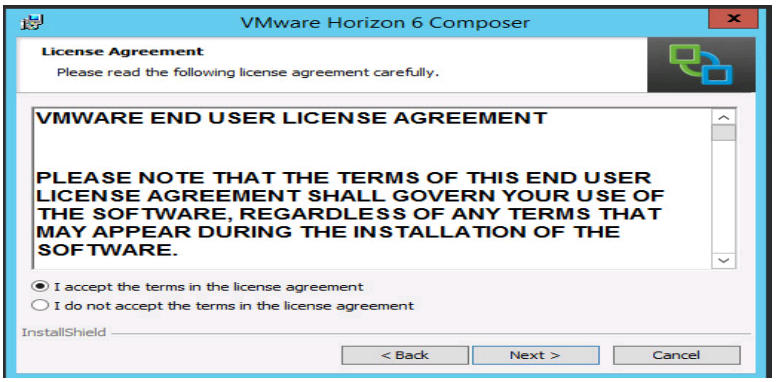

5. Click Next to install view composer.

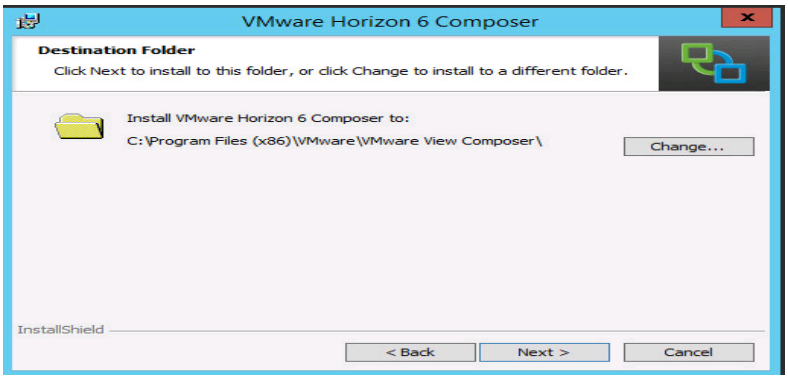

6. Provide ODBC Data Source Name.

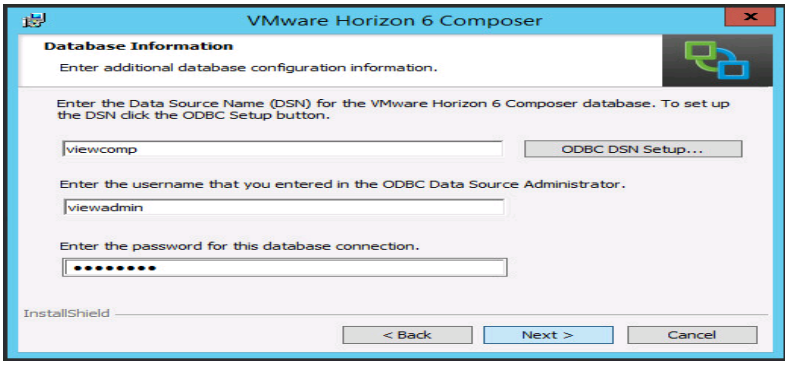

7. Configure port automatically and create default certificate.

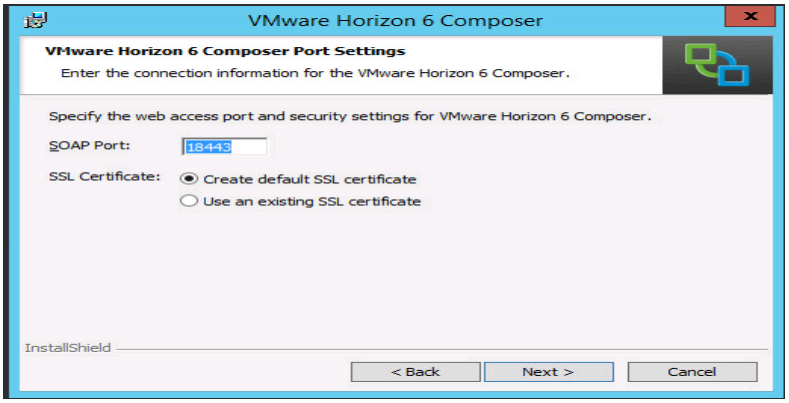

8. Click Install.

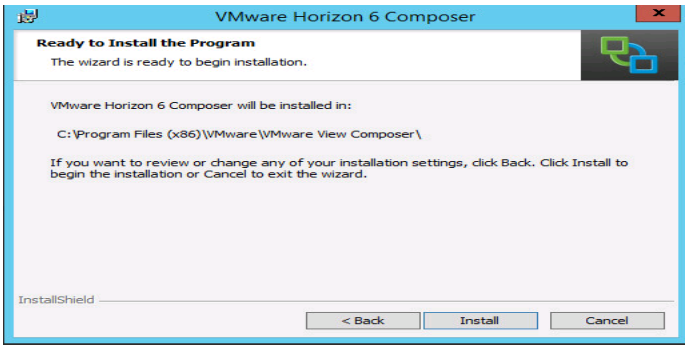

9. Installing vmware view composer.

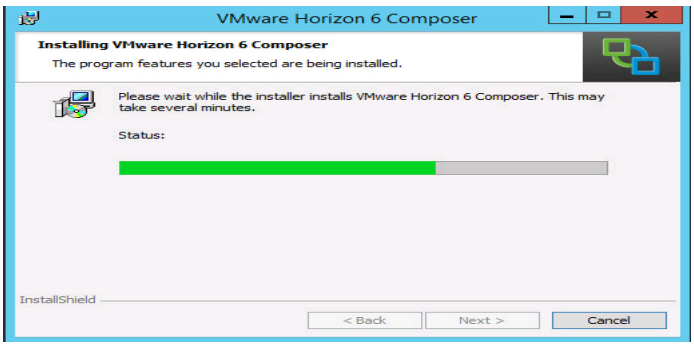

10. Completed the installation and view composer is ready to serve the composer operations or desktops provisioning

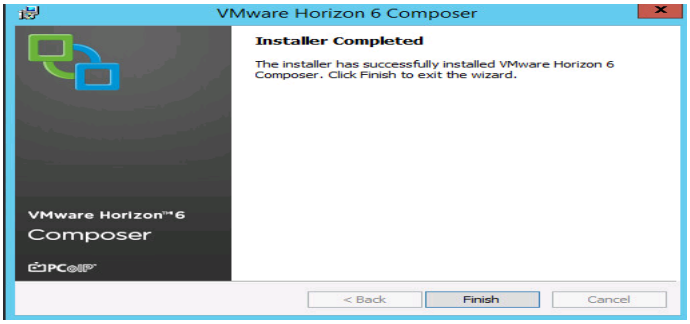

View Connection and Replica Server Installation

The prerequisites are to have Windows 2008 or 2012 Server configured and download the VMware view Connection server bits. In this study we have used VMware View Connection Server version 6.2.0-3005368.

To view connection and replica server installation, complete the following step:

1. Download the View Composer server Installer. In this study, View Connection Server 6.2.0-3005368 version.

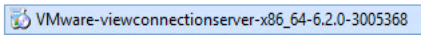

2. Click Next to install the view connection server installer.

3. Accept the EULA agrement.

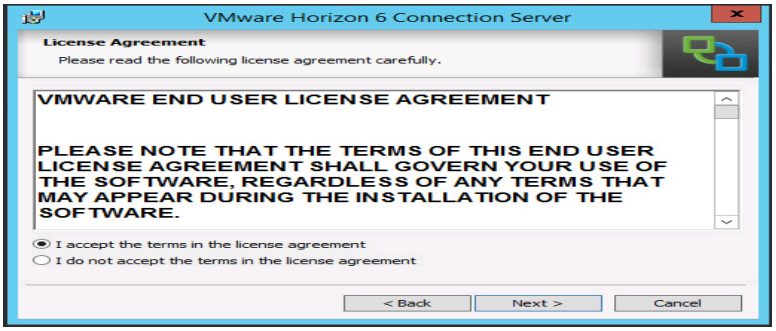

4. Select Horizon 6 standard server.

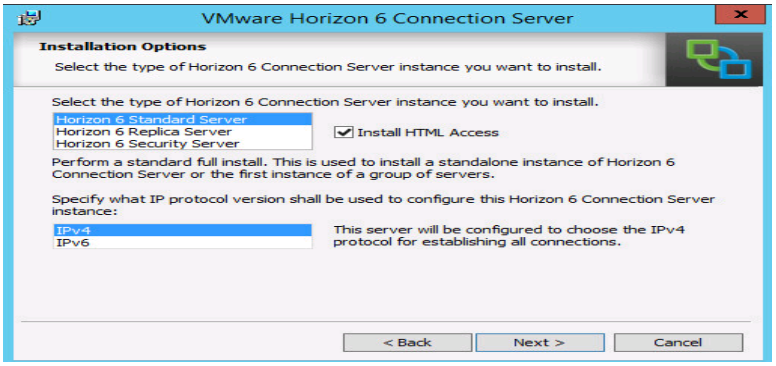

5. Provide the password information.

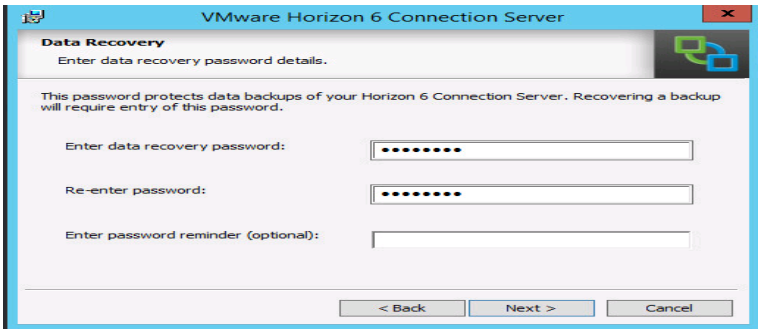

6. Configure Windows Firewall automatically.

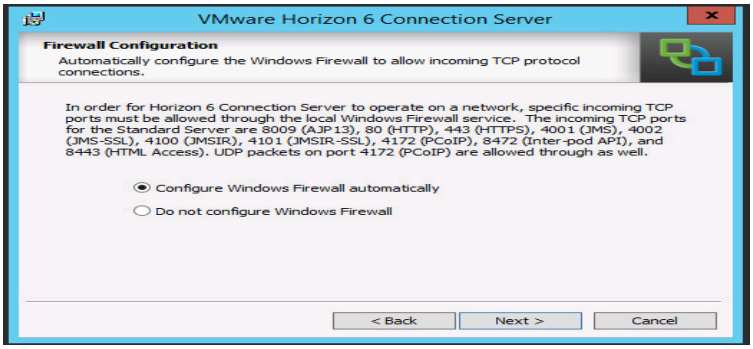

7. Provide domain admin credentials.

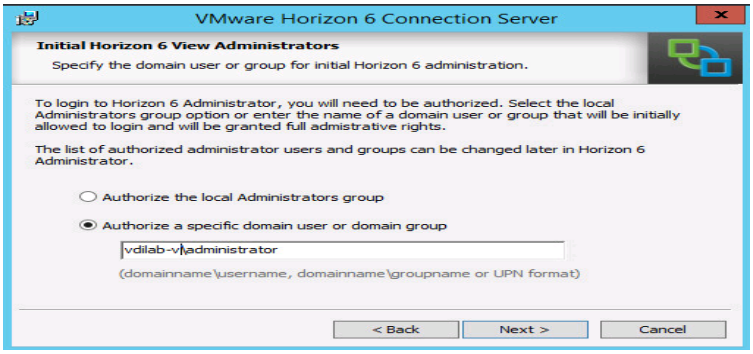

8. Click Next to install.

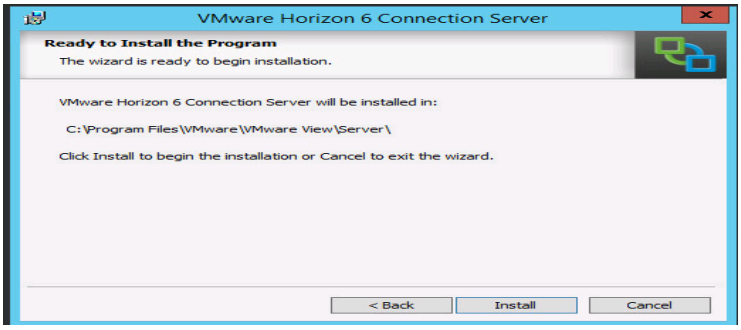

9. Click Next to install.

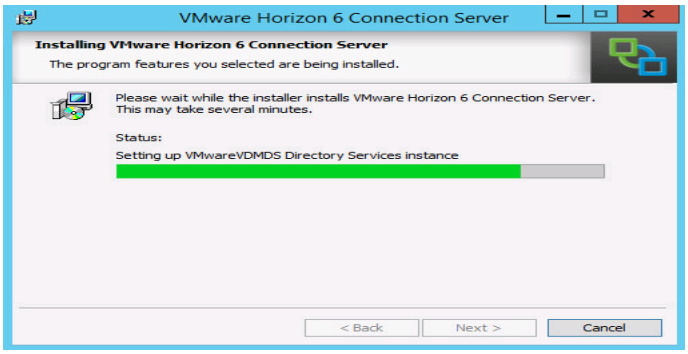

10. View connection installation is complete.

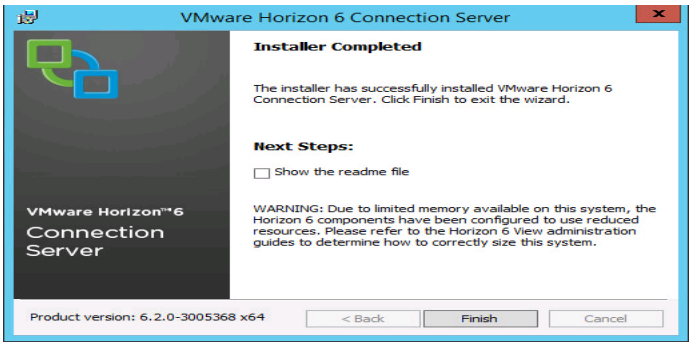

11. Click Login in to view adminsitration cosnole.

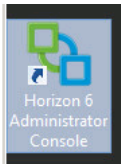

VMware Horizon View Replica Server Installation

To install Horizon View Replica Server and additional replica servers, follow the installations steps shown above installing VMware View Connection and Replica Server Installation and complete the remaining steps to configure Horizon View Replica Server.

To install additional Horizon View Replica servers, complete the following steps:

1. Select the Horizon 6 replica server.

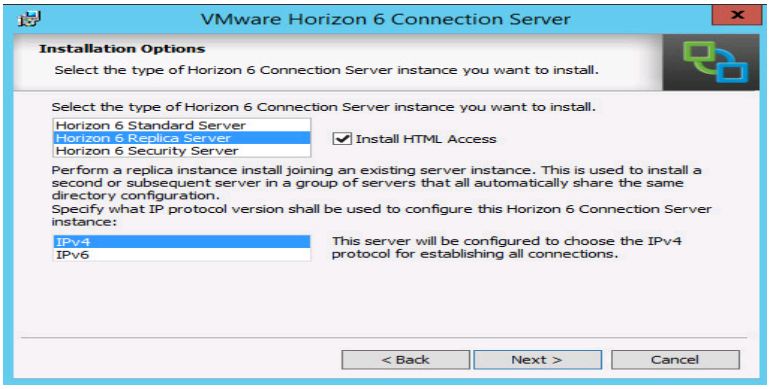

2. Provide the existing View Connection Server IP address or FQDN.

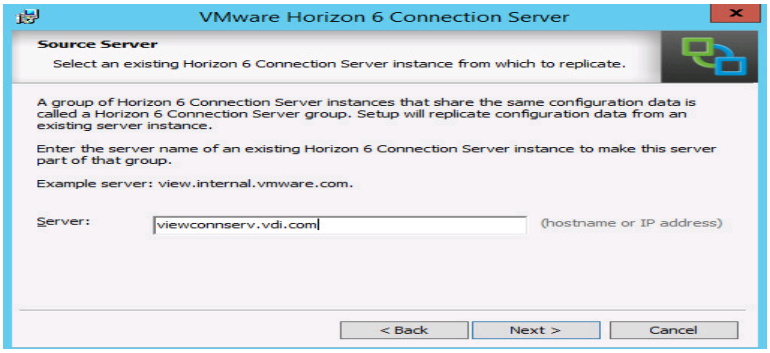

3. Proceed with the rest of the procedure to complete Replica Servers installation.

Creating 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 7 OS. We used windows 7 32 bit OS for our testing.

For the virtual machine following parameters were used:

Memory : 2048MB

Processor : 2vCPU

Hard Disk : 20 GB

Network Adapter : 1 VMXNET3 type attached to VDI port-group on Nexus 1000v

2. Create New Virtual Machine.

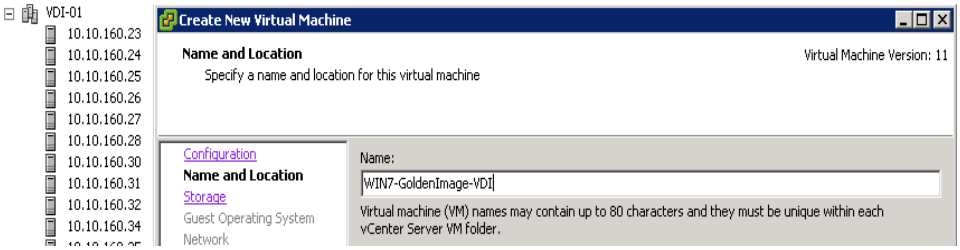

3. Name and location of the virtual machine to be created.

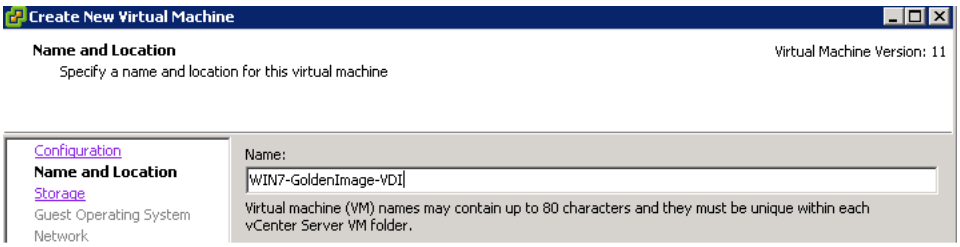

4. Select datastore.

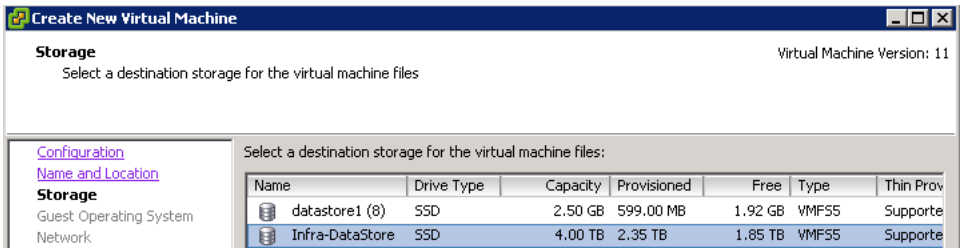

5. Select the Operatign system applicable.

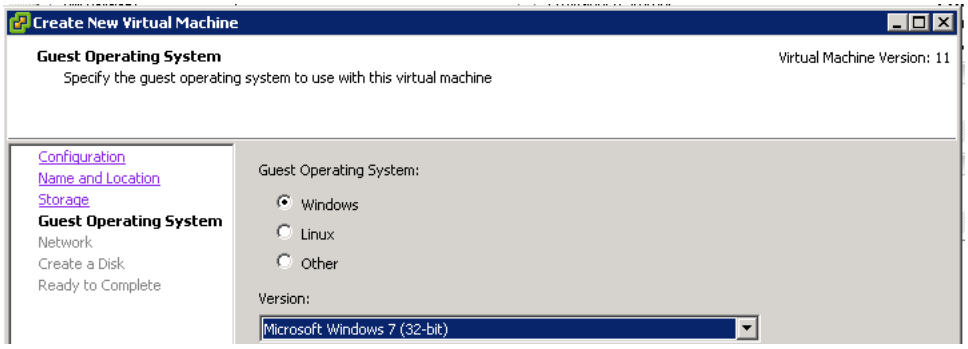

6. Select the netowrk and adapter.

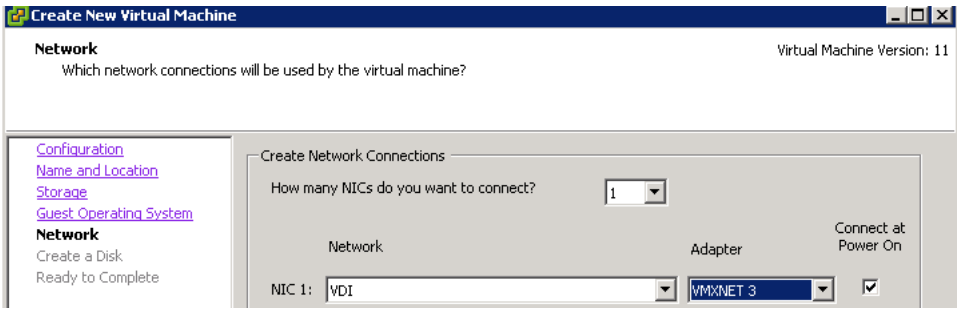

7. Select disk size.

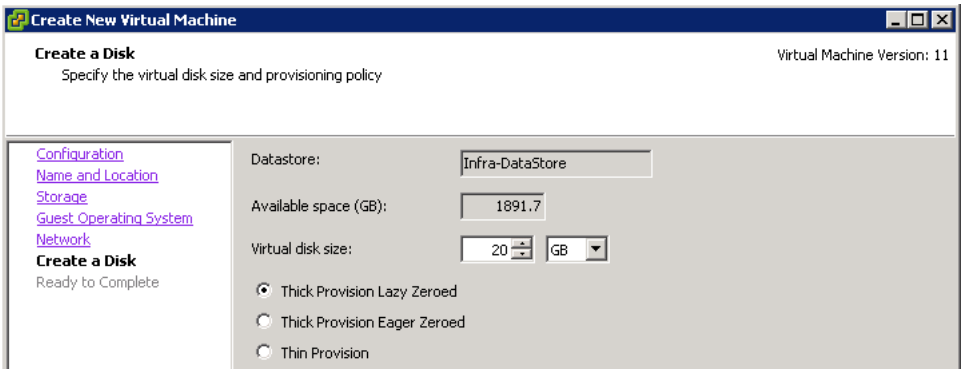

8. Review and click Next to attach the OS Image.

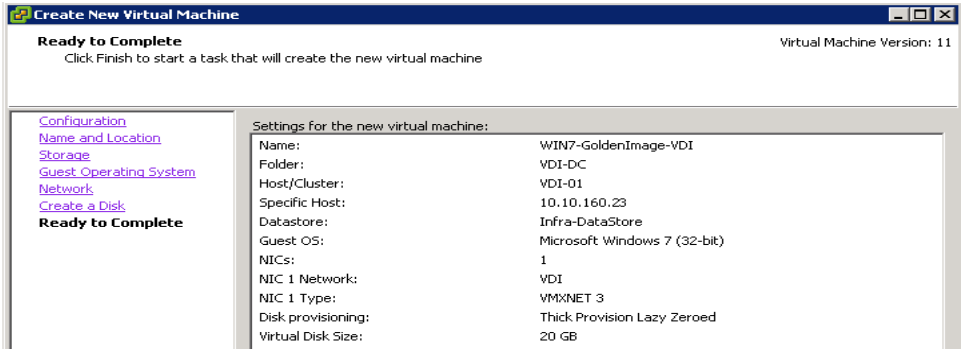

9. Attach the windows ISO image.

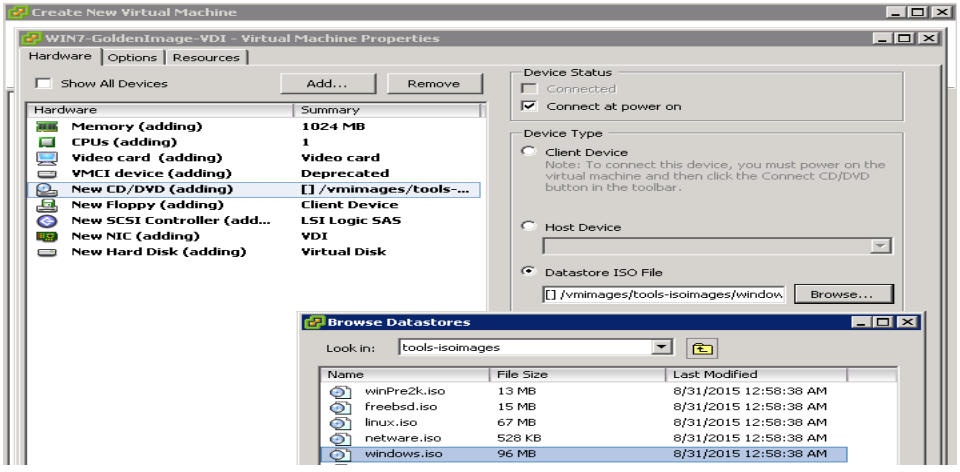

10. Complete the remaining steps to install the Win7 Operating System.

Optimize Base Windows 7 SP1 Virtual Machine

To optimize the base Windows 7 SP1 machine, complete the following steps:

1. Follow the link to optimize windows 7 SP1 32 bit virtual machine:
   https://labs.vmware.com/flings/vmware-os-optimization-tool

   www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf

2. Install View 6.2,0 Virtual Desktop Agent software:

   https://my.vmware.com/web/vmware/details?productId=529&downloadGroup=VIEW-620-ENT-GA

3. Click Next to install.

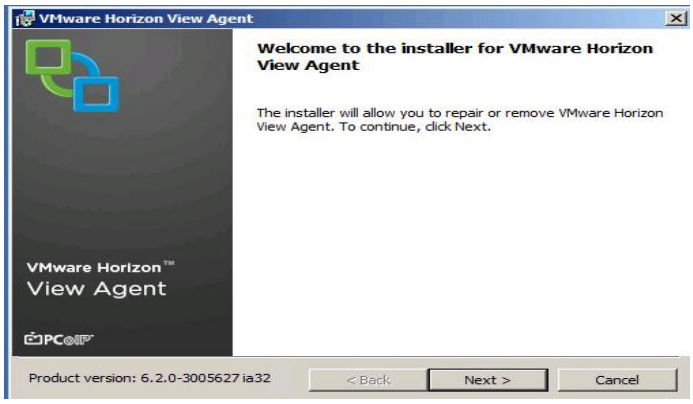

4. Accept the EULA Agreement.

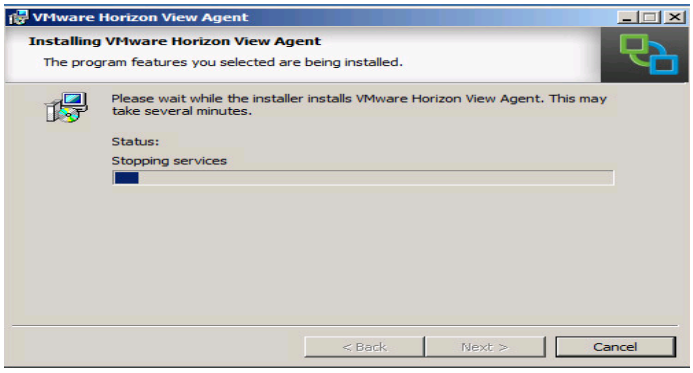

5. View agent has been installed.

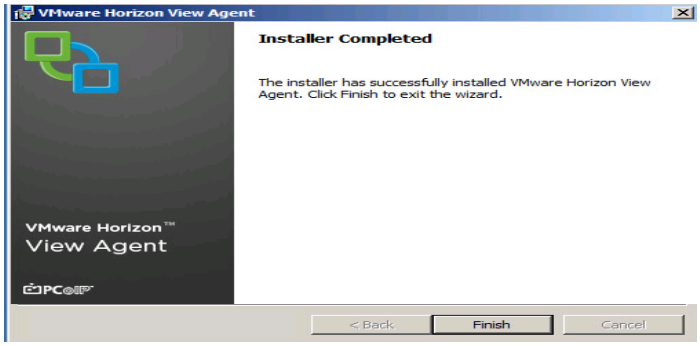

Install Additional Software

To install additional software, complete the following steps:

1. Install additional software required in your base windows image.

For the testing, we installed Office 2010.

2. Login VSI Target software package to facilitate workload testing.
3. Reboot the VM.
4. Install service packs and hot fixes required for the additional software components that were added.
5. Shut down the VM.

Create a Snapshot for Virtual Machine

To create a Snapshot for a virtual machine, complete the following steps:

1. Shut down the Windows 7 Golden Image virtual machine to take a snapshot.
2. Right-click Windows 7 Golden Image Virtual Machine Properties to take a snapshot which is required for the virtual desktop deployment.
3. Provide the name and description for the Snapshot and click OK.
4. Click OK.

5. Take a Snapshot of the configured OS.

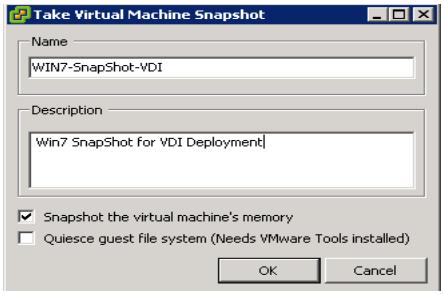

Create Customization Specification for Virtual Desktops

To create a customization specification, complete the following steps:

1. Right-click the powered off virtual machine after taking a snapshot and select Template and click Convert to template.

2. Provide a name to the template and provide the host /cluster, data store details.

3. Select Guest Customization and check the radio button for Customize using the Customization Wizard.

4. Click Next.

5. Windows customization specification.

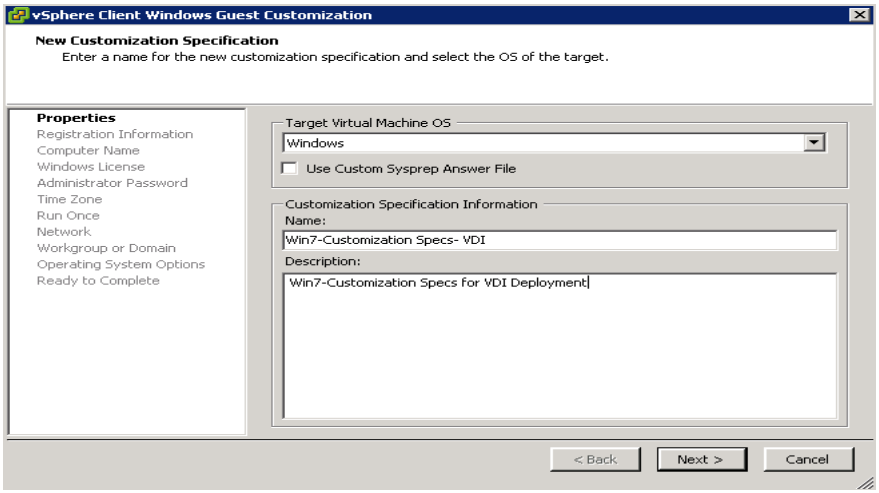

6. Provide a name and organization details.

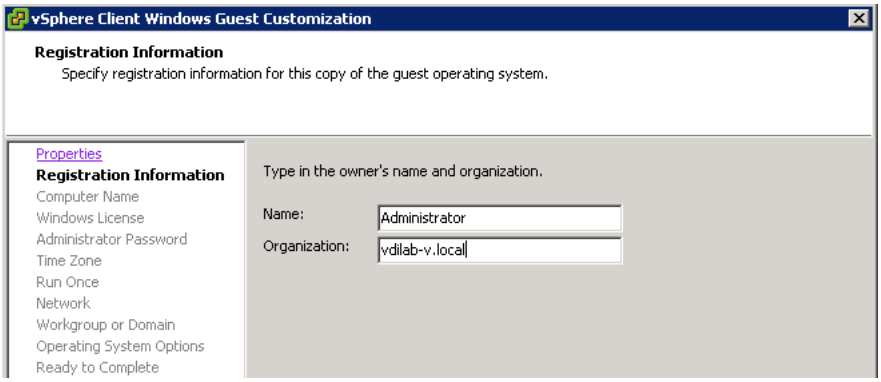

7. Provide computer name.

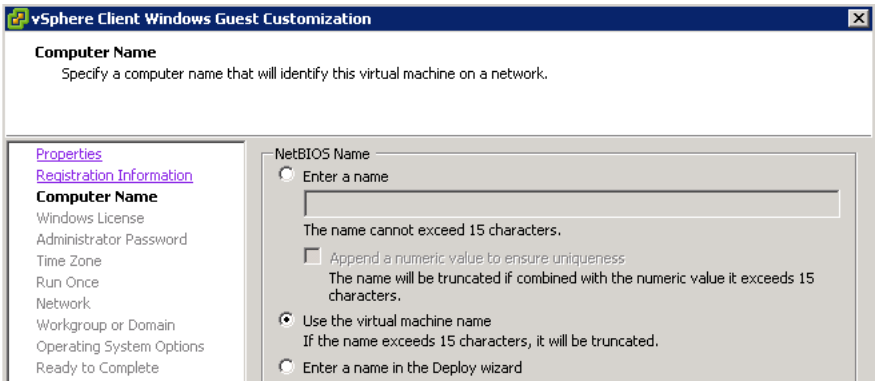

8. Provide product key.

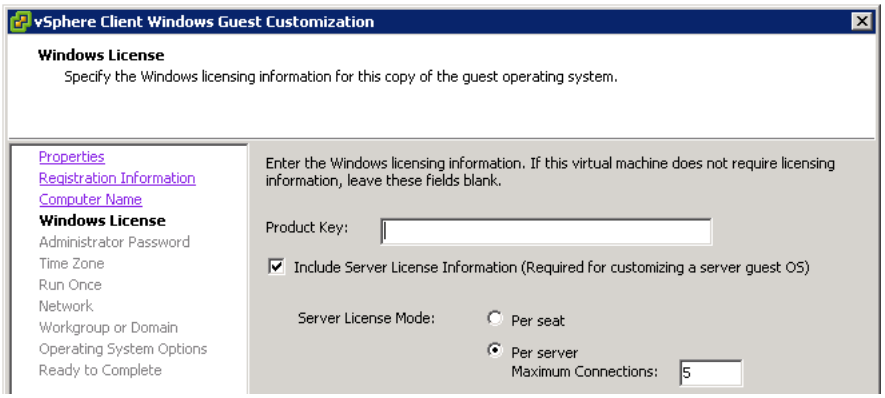

9. Provide Password credentials.

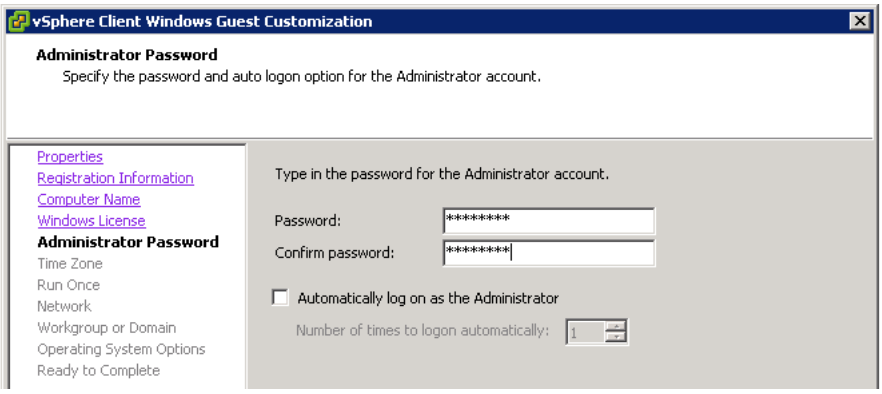

10. Provide network information.

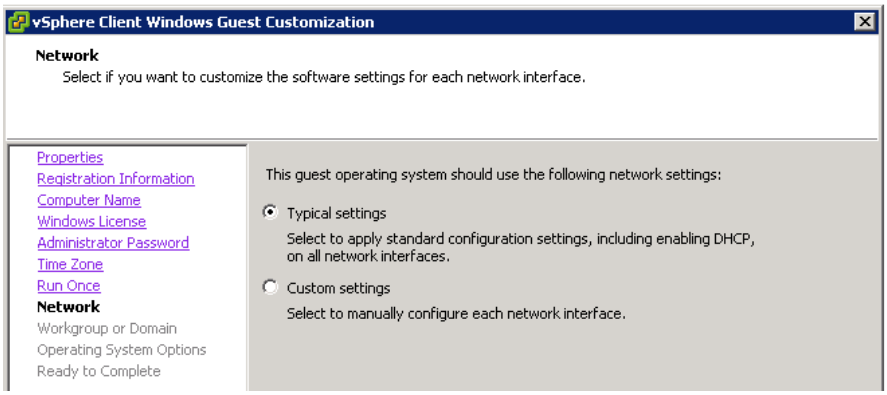

11. Provide domain name and user credentials.

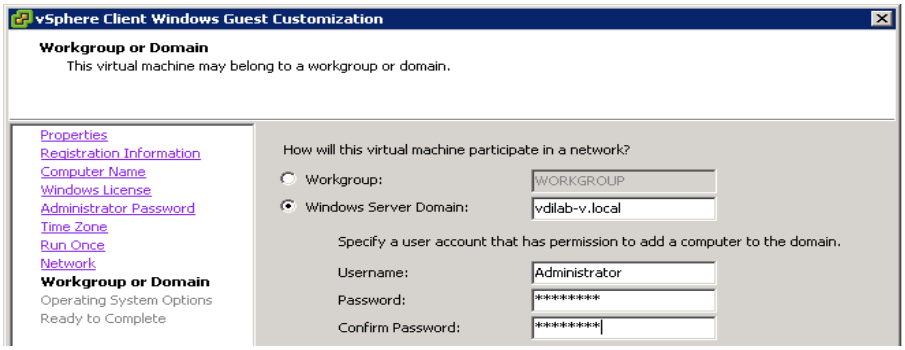

12. Generate new security ID.

13. Review and click Next to complete creating customization specs.

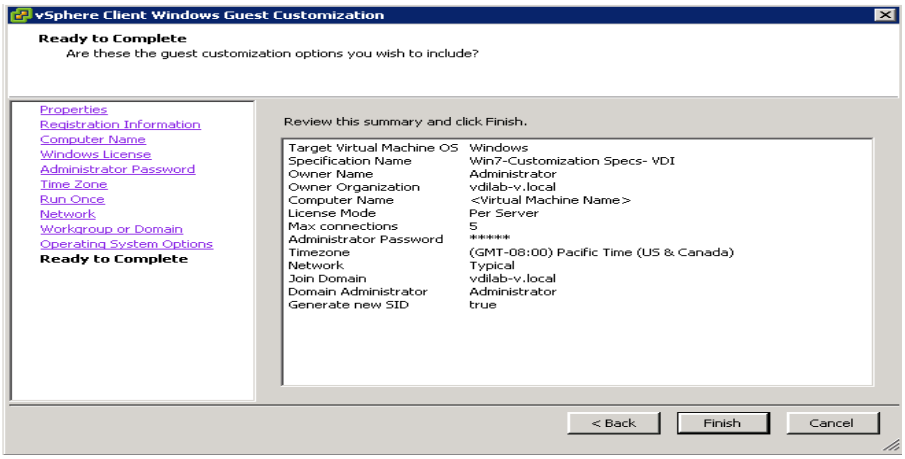

Now the Golden image is ready for VDI virtual machines deployment.

Install RDSH Server Role for RDS Hosted Session Hosts Users

This section describes the installation of RDS Role for RDSH server configuration.

Prerequisite: Configure a Windows 2012 or R2 server for configuring the RDSH server roles. In this study we have configured 2012 R2 servers for RDS server roles.

To create a new Virtual machine as a Windows 2012 R2 server, complete the following steps:

1. Install new virtual machine.

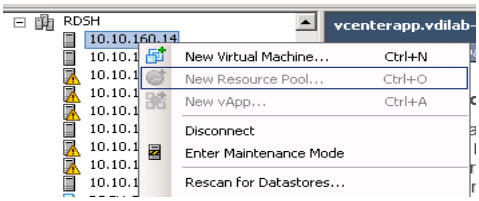

2. Select configuration.

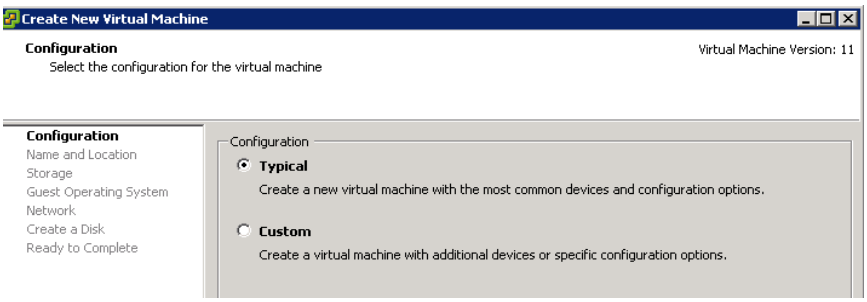

3. Provide name for the RDSH server.

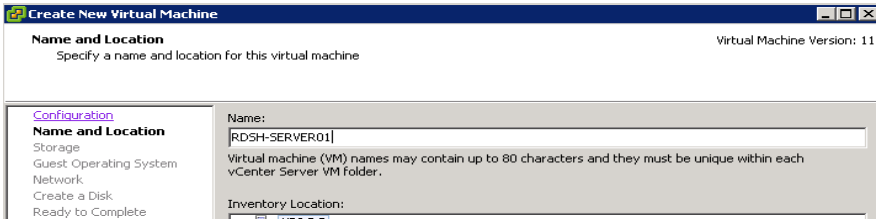

4. Select datastore.

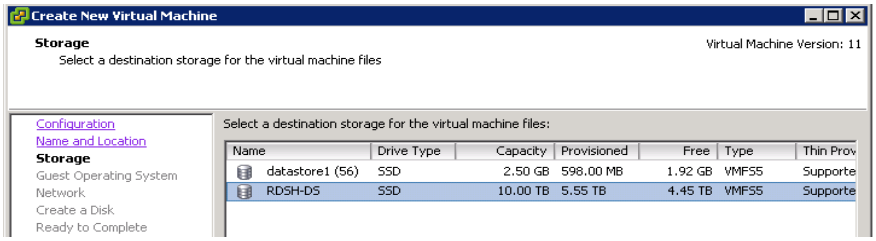

5. Select operating system.

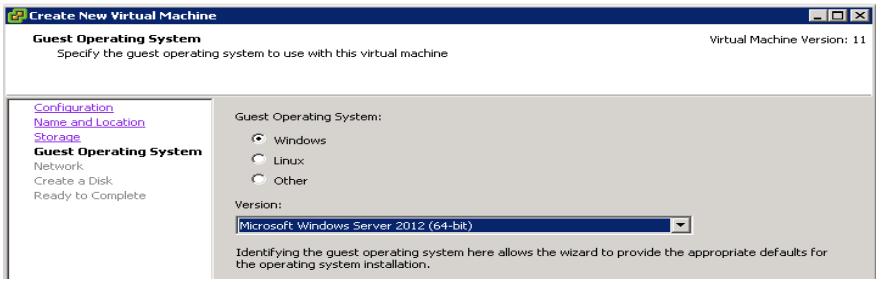

6. Disk size.

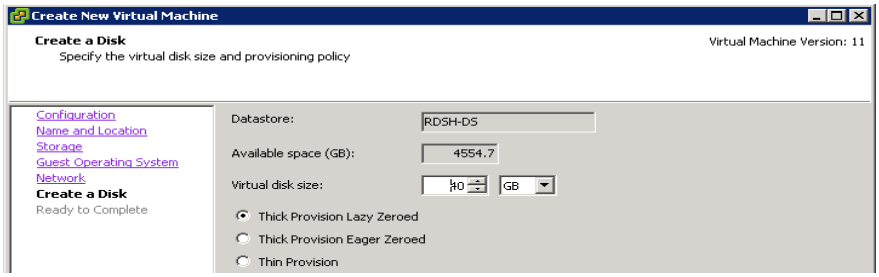

7. Select the datastore size.

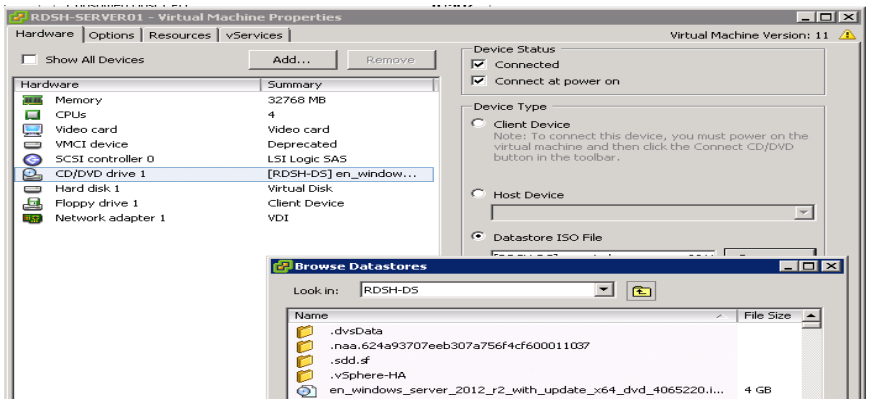

8. Windows 2012 is ready to install.

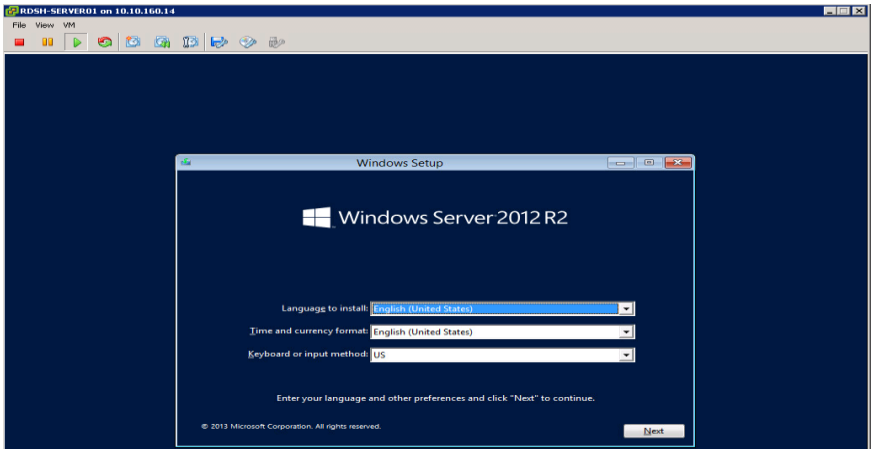

9. RDSH Server has been created.

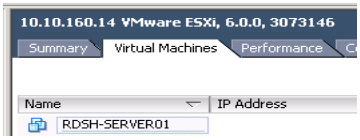

10. Follow the link provided below to optimize windows server 2012.

Complete the required Windows server 2012 creation and provide the necessary IP address or DHCP and reboot the server.

Now you need to customize the windows servers using VMware Customization Specs.

https://labs.vmware.com/flings/vmware-os-optimization-tool

Configure the RDSH Server Roles for RDS Hosted Session User Hosts

To configure the RDSH server roles for RDS hosted session user hosts, complete the following steps:

1. Login in to RDSH-SERVER01 and configure RDSH server roles.

2. Add roles and features from server manager.

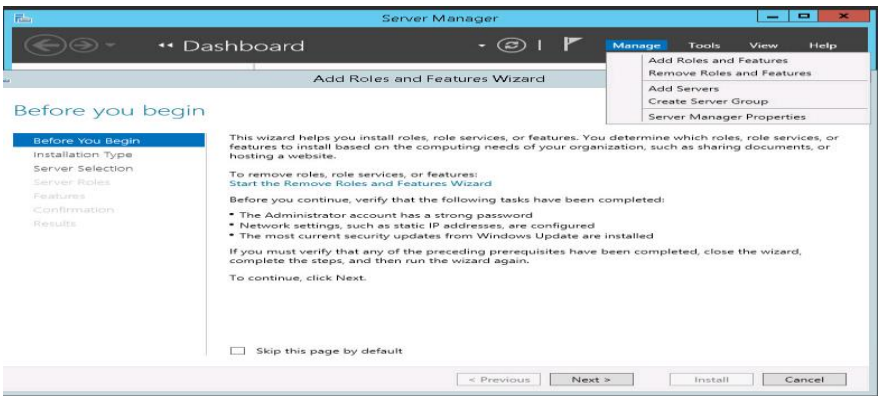

3. Select role based or feature based as applicable.

4. Select the server from server pool.

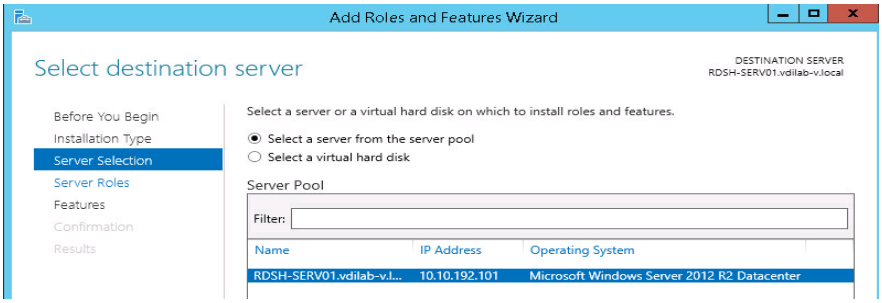

5. Select remote desktop services and select remote desktop session host.

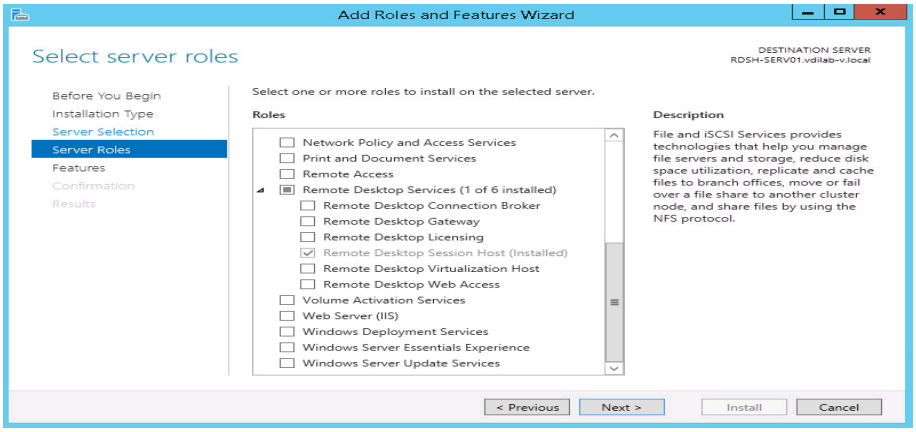

6. Select remote desktops services tools and select licensing diagnostic tool.

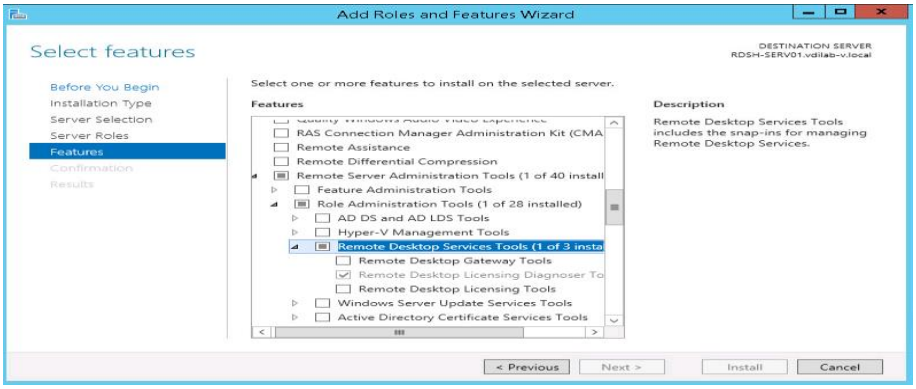

7. RDSH Server has been created and ready for RDSH server role for users to connect.

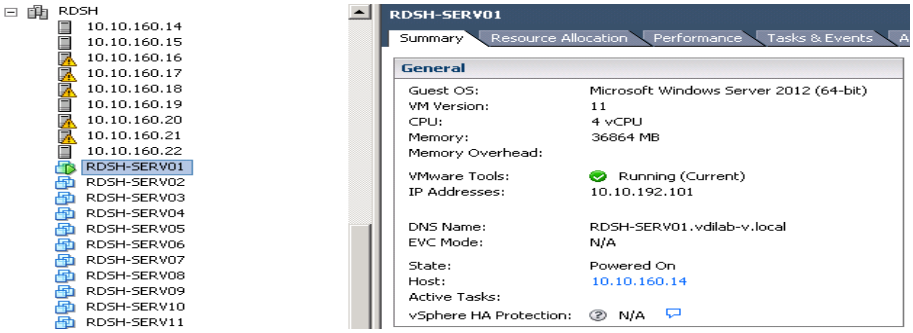

8. Reboot the Server and install other updates and necessary software.

Install additional software

To install additional software required in your base windows server image, complete the following steps:

For the testing, we installed Office 2010.

1. Login VSI Target software package to facilitate workload testing.
2. Reboot the VM.
3. Install service packs and hot fixes required for the additional software components that were added.
4. Reboot the server.

Test Configurations and Sizing Guidelines

In this project, we tested a single Cisco UCS B200 M4 blade in a single chassis and 30 Cisco UCS B200 M4 blades in four chassis to illustrate linear scalability for each workload studied.

We also tested a mixed workload scenario with 5000 VDI users running both RDS Hosted server sessions and VDI virtual machines on Cisco B200 M4 B-Series servers. We separately tested single server RDS Hosted sessions, RDSH cluster testing, single server VDI virtual machines along with VDI cluster testing all running on Cisco UCS B200 M4 servers.

The tested system was comprised of the following hardware and software components:

Hardware components:

- Cisco UCS B200 M4 Server Chassis
- 2 Cisco UCS 6248UP Fabric Interconnects
- 2 Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.6 GHz, with 128 GB of memory per blade server [16 GB x 8 DIMMs at 2133 MHz]) for two Infrastructure blades.
- 30 Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2680 v3 CPUs at 2.5 GHz, with 384 GB of memory per blade server [16 GB x 24 DIMMs at 2133 MHz]) for RDS work load blade.
- Cisco VIC CNA (1 per blade)

- 2 Cisco MDS 9148S
- 2 Cisco Nexus 9300 access layer switch
- Pure Storage FlashArray//m50 + 44TB external shelf (Purity Operating Environment 4.5.5)

Software components:

- Cisco UCS firmware 2.2.6(c)
- VMware View RDS Hosted Servers Sessions
- Windows 2012 R2 Server for RDS Hosted Sessions. Total of 9 Servers Configured.
- v-File Windows 2008 User Profile Server
- Microsoft Windows Server 2012 R2, 4CPU, 32GB RAM, 40 GB vdisk.
- Microsoft Windows 7
- Microsoft Office 2010
- Login VSI 4.1.4

Cisco UCS B200 M4 Single Server Configuration and Sizing for Horizon RDSH User Sessions

Nine Windows 2012 R2 virtual servers on one Cisco UCS B200 M4 blade server configured for Hosted Shared Desktops to host 180 User Sessions.

Figure 96 RDS Single Server Testing – Recommended workload: 180 User Sessions per Server

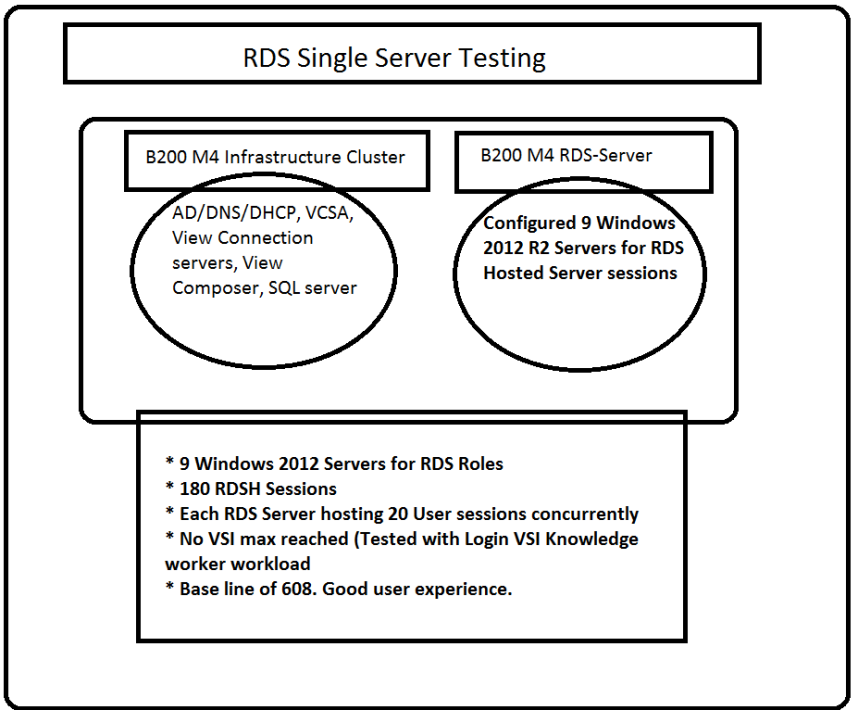

Cisco UCS B200 M4 Configuration and Sizing for Horizon RDSH User Session Cluster

Nine Cisco UCS B200 M4 blade servers running Windows 2012 R2 Servers configured for Hosted Shared Desktops to host 1450 User Sessions. N+1 server fault tolerance is provided by this cluster.

Figure 97 RDS Cluster Server Testing

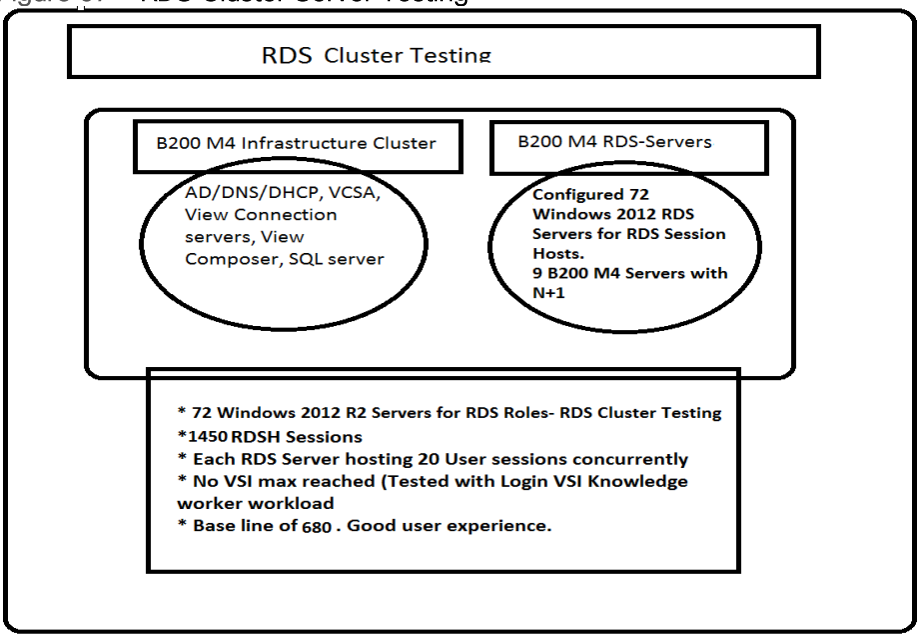

Cisco UCS B200 M4 Single Server Configuration and Sizing for Horizon Linked Clones

181 Horizon virtual Windows 7 linked clone desktops on one Cisco UCS B200 M4 blade server represents our maximum recommended workload for a single server.

Figure 98 Cisco UCS B200 M4 Server for Single Server Scalability Horizon View 6.2

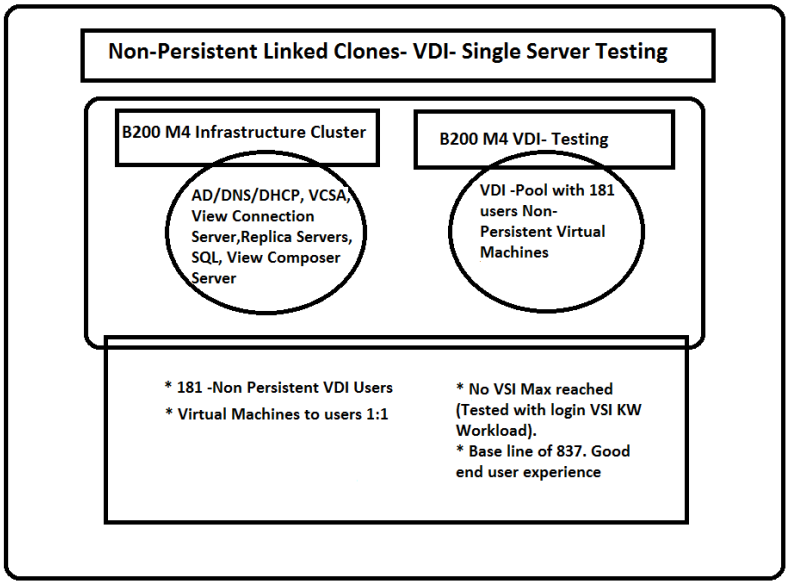

Cisco UCS B200 M4 Configuration and Sizing for Horizon Linked Clone Cluster

3550 Horizon virtual Windows 7 linked clone desktops on 21 Cisco UCS B200 M4 blade server represents our cluster scale point, providing N+1 server fault tolerance at the cluster level.

Figure 99 21x Cisco UCS B200 M4 Server for Blade Server VDI Scalability VMware Horizon View VDI Desktops

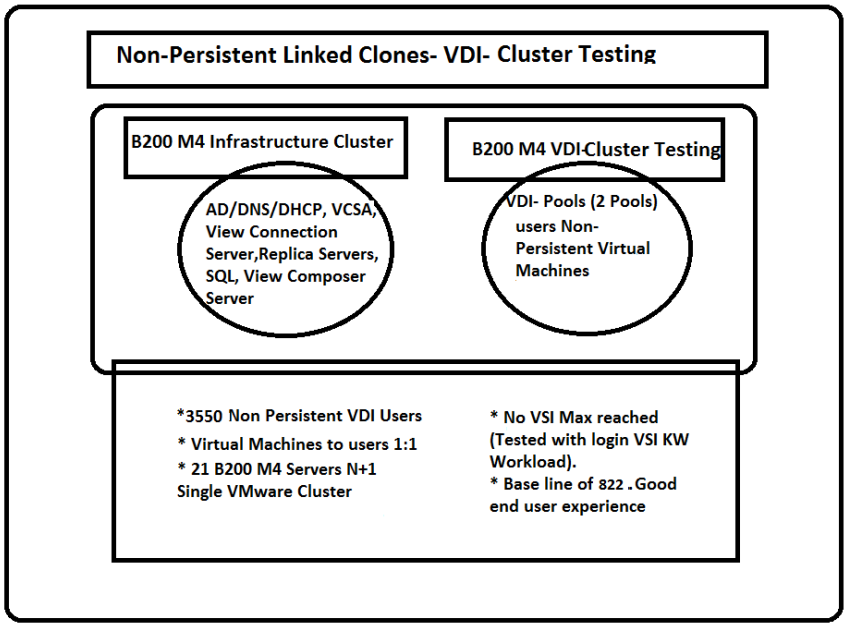

Cisco UCS B200 M4 Configuration and Sizing for 5000 User Mixed Workload Horizon RDSH and Linked Clone Scale Test

1450 Horizon RDSH sessions on 9 Cisco UCS B200 M4 blade servers running Windows 2012 R2 virtual servers configured for Hosted Shared Desktops.

3550 Horizon virtual Windows 7 linked clone desktops in two Horizon Linked Clone Desktop pools on 21 Cisco UCS B200 M4 blade server represents our cluster scale point, providing N+1 server fault tolerance at the cluster level.

Figure 100 30x Cisco UCS B200 M4 Servers for both RDSH and VDI mixed work load with server N+1 fault tolerance at the cluster level.

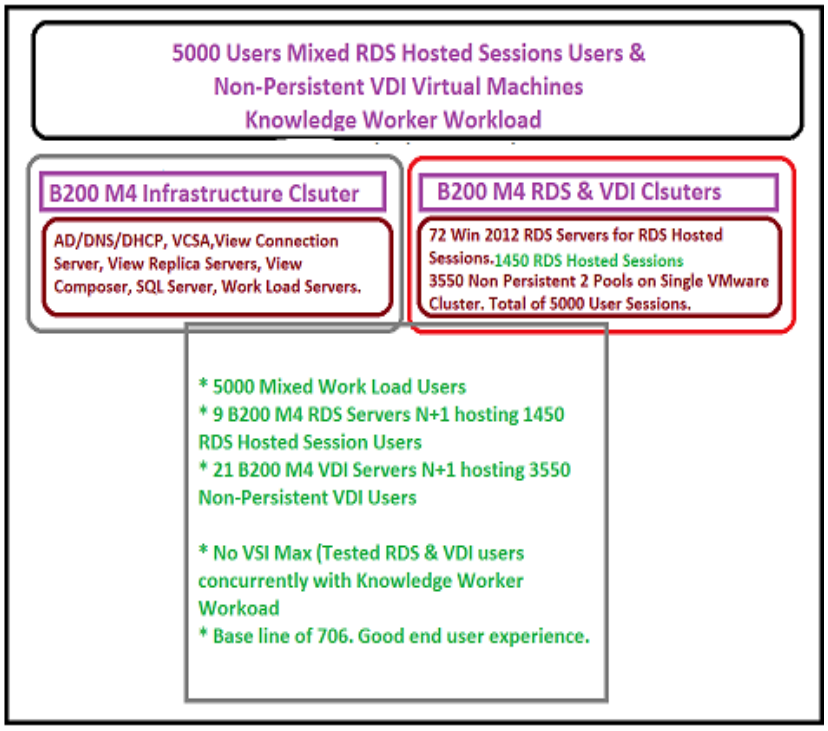

Testing Methodology and Success Criteria

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 View 6.2 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 Blades used in test run
   - All Infrastructure VMs used in test run (AD, SQL, View Connection brokers, image mgmt., etc.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot RDS Machines using VMware Horizon View 6.2 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 View 6.2 Administrator Console dashboard. Typically a 20-30 minute rest period for Windows 7 desktops and 10 minutes for RDS VMs is sufficient.

6. Time 1:35 Start Login VSI 4.1.4 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 View 6.2 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 7 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 is 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 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 with View Connection Server Dashboard will be monitored throughout the steady state to make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.

Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with 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.

The purpose of this testing is to provide the data needed to validate VMware Horizon View 6.2 Hosted Shared Desktop with VMware View 6.2 Composer provisioning using Microsoft Windows Server 2012 R2 sessions on Cisco UCS B200 M4

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, linear scalability.

VSImax 4.1.x Description

The philosophy behind Login VSI is different to conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system, and are initiated at logon within the simulated user’s desktop session context.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solutions, 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 101 Sample of a VSI Max Response Time Graph, Representing a Normal Test

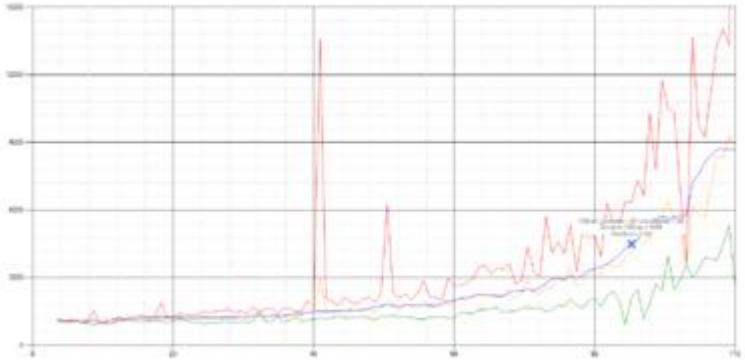

Figure 102 Sample of a VSI Test Response Time Graph with a Clear Performance Issue

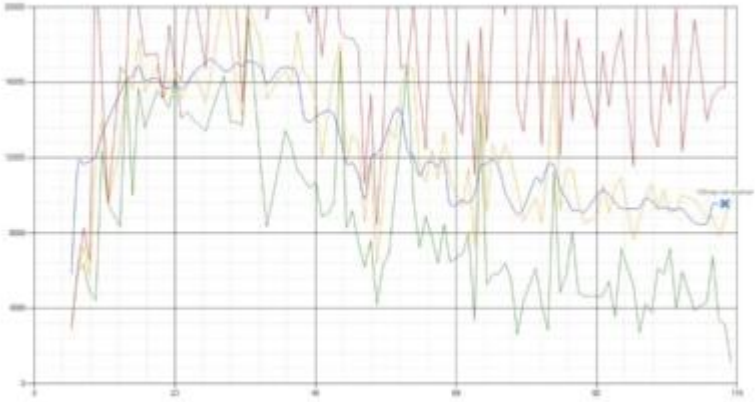

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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. This 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 'base phase' 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. The calculation is as follows:

- 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% of the amount of “active” sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40% of 60) = 31 response time measurement are used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5% and bottom 5% 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% 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.

Test Results

Single-Server Recommended Maximum Workload for Cisco UCS B200 M4 Blade Servers

For both VMware Horizon View 6.2 RDS Hosted Shared Desktop and Linked Clone use cases, the recommended maximum workload was determined based on both Login VSI Knowledge Worker workload end user experience measures and blade server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to insure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95% when running the maximum recommended load.

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
Login	The Login VSI phase of test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files
Logoff	Sessions finish executing the Login VSI workload and logoff

Single Server Testing with RDSH Remote Desktop Server Hosted Session Users

The maximum recommended load for Horizon RDSH user sessions on a single Cisco UCS B200 M4 blade server is 180 sessions running on 9 virtual Windows Server 2012 R2 machines.

Figure 103 Login VSI Chart for RDSH Single Server Testing

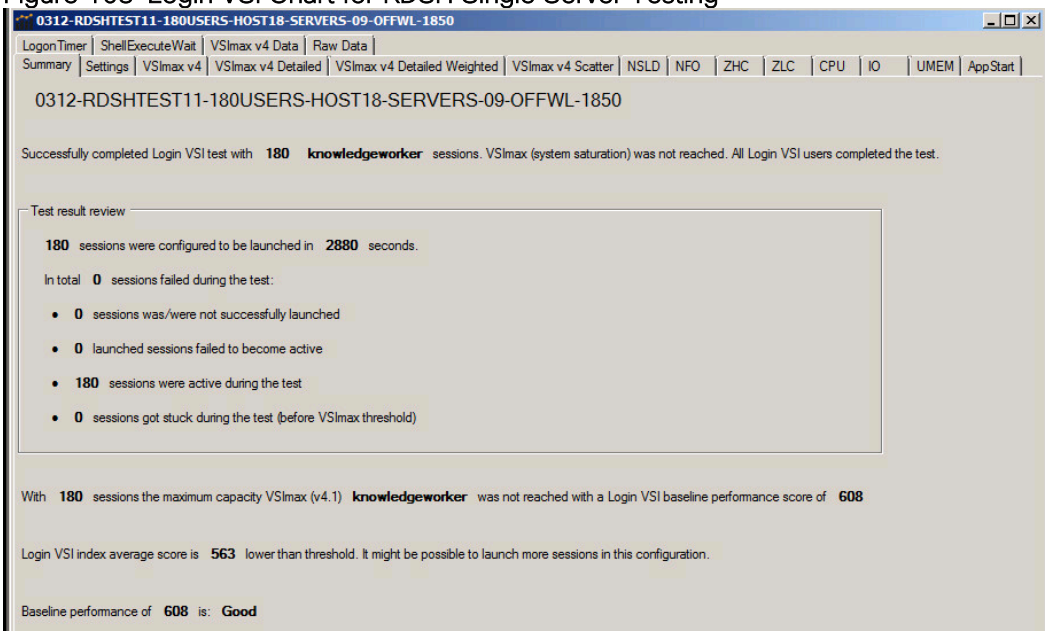

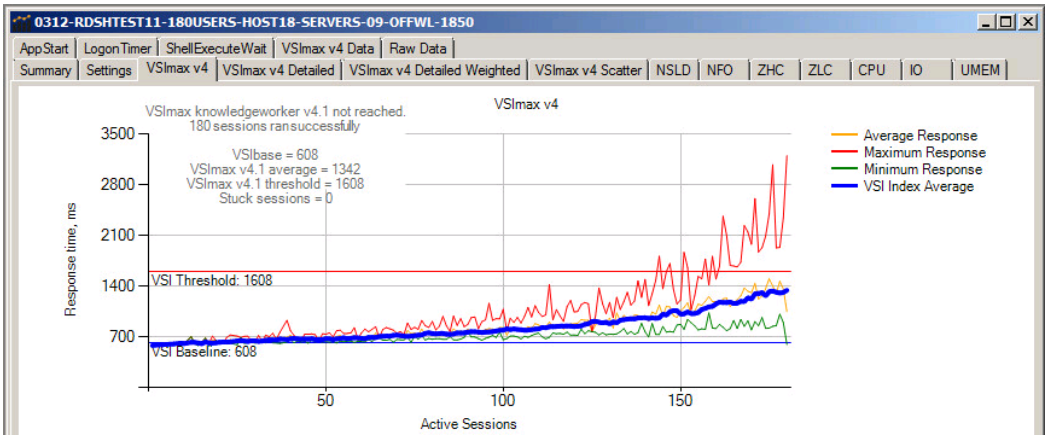

Figure 104 ESXTOP Core Utilization for RDSH Single Server testing

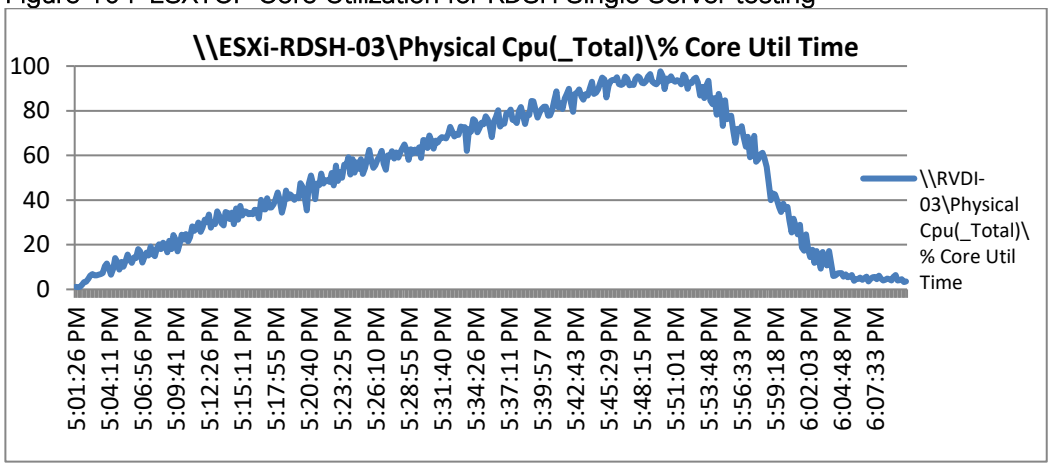

Figure 105 ESXTOP Utilization Time for RDSH Single Server testing

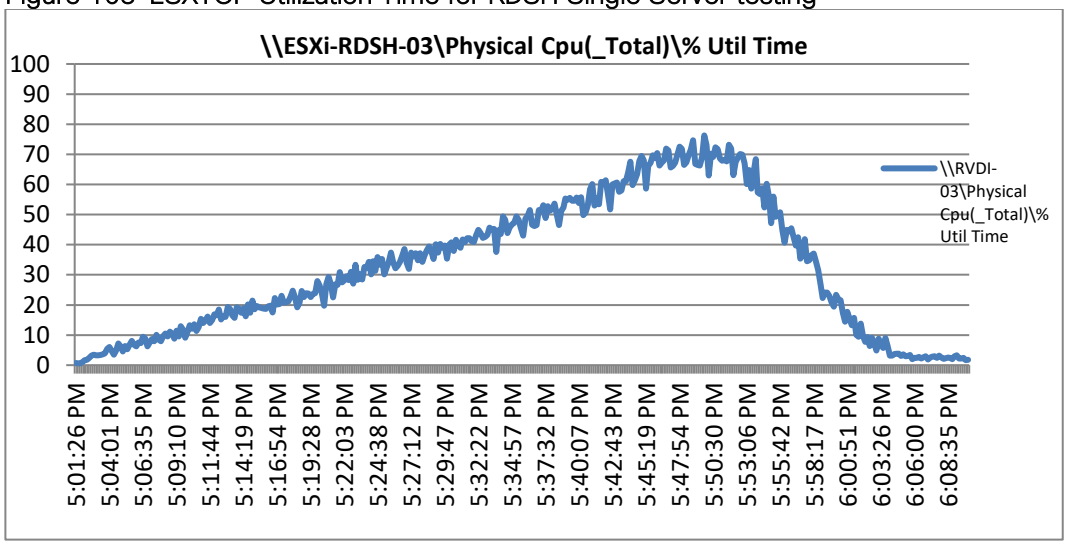

Figure 106 Network Adapter (VMNICs) MB received /Transmitted for Sec for the RDSH Single Server testing

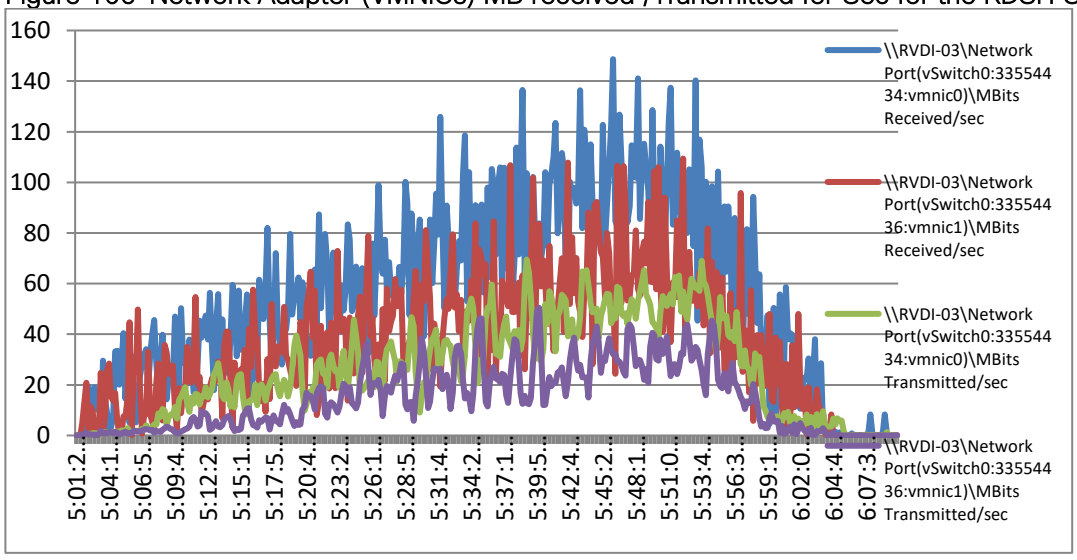

Figure 107 Disk Adapter MBytes Read / Write for Sec for the RDSH Single Server testing

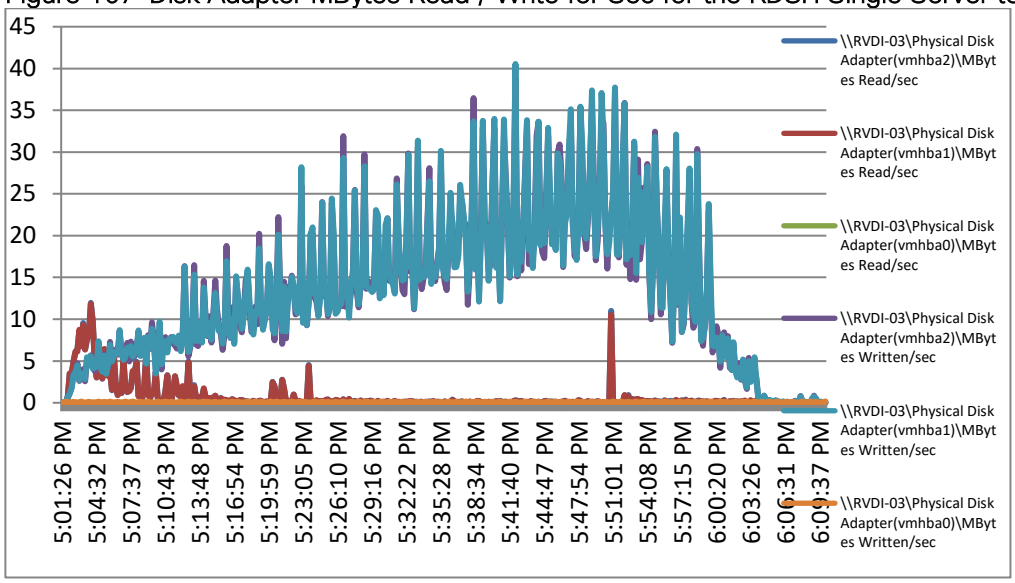

Figure 108 Non Kernel Mbytes for RDSH Single Server testing

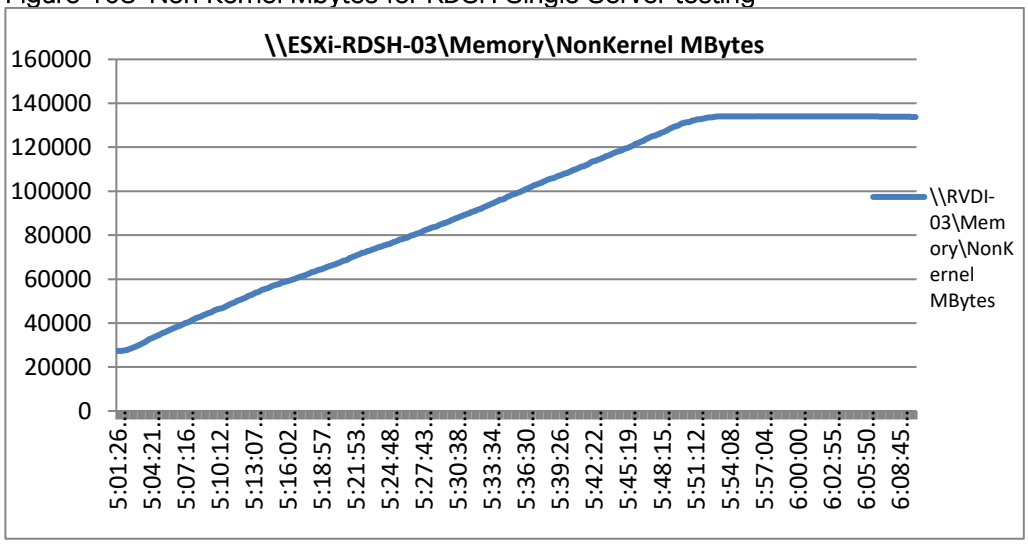

Single Server Testing with Horizon Linked Clone Users

Based on the testing performed, 181 Horizon Linked Clone virtual desktops represent the maximum recommended load for the Cisco UCS B200 M4 blade as configured.

Figure 109 Login VSI graph for Non-persistent VDI users Single Server

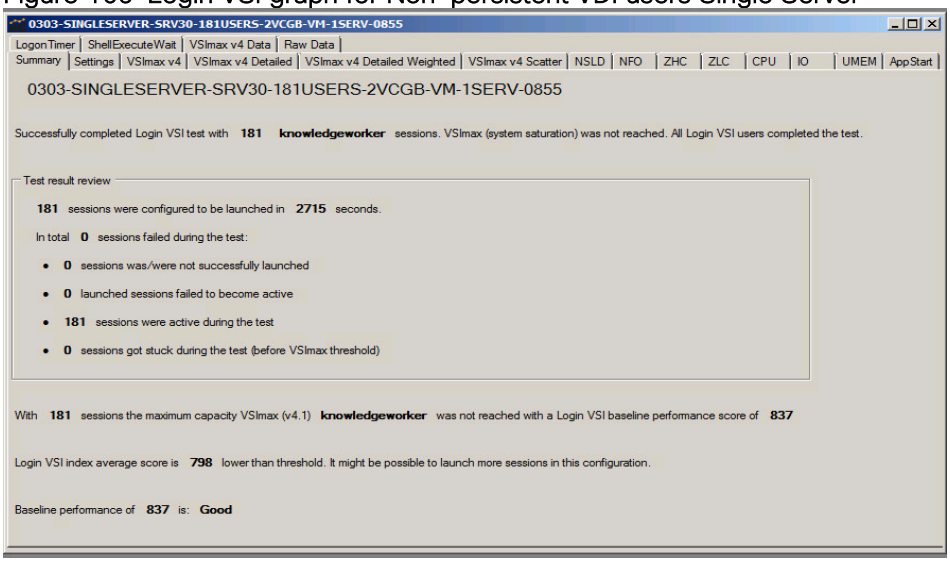

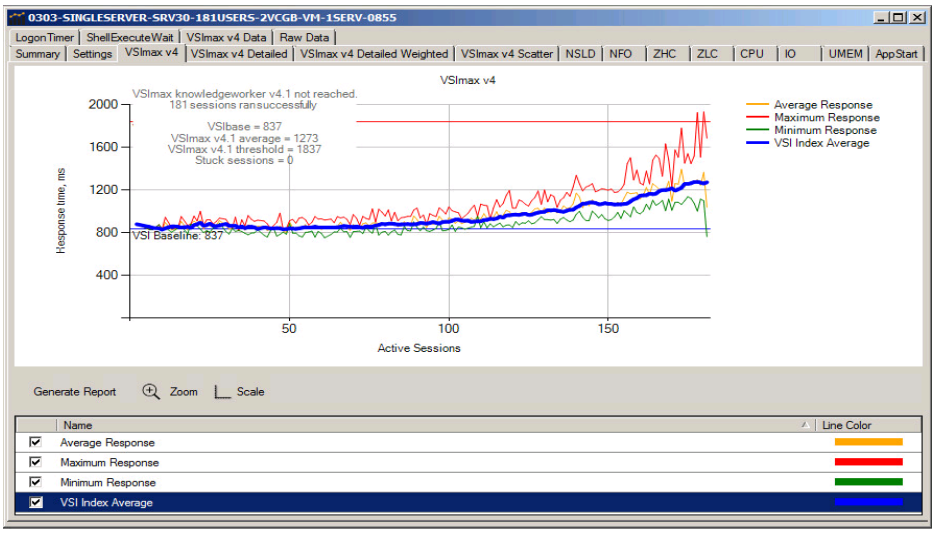

Figure 110 ESXTOP CPU Core utilization for Non-persistent VDI users Single Server testing.

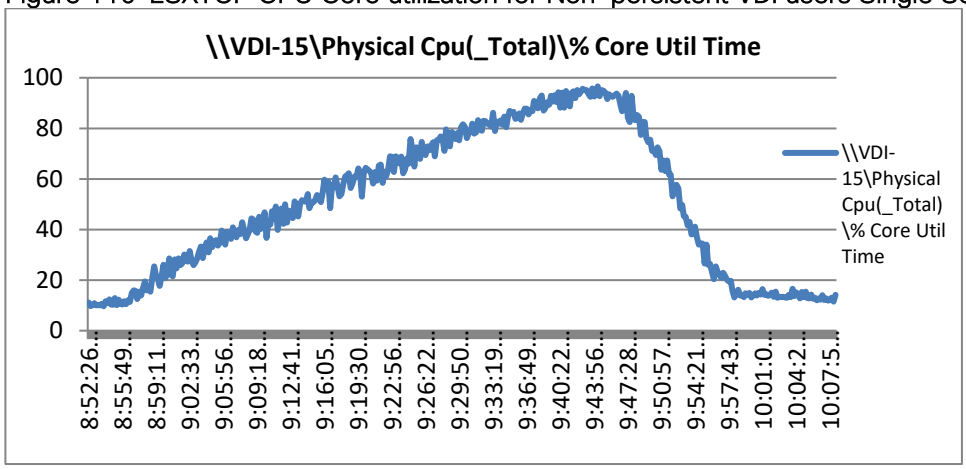

Figure 111 ESXTOP CPU Core Total utilization for Non-persistent VDI users Single Server testing

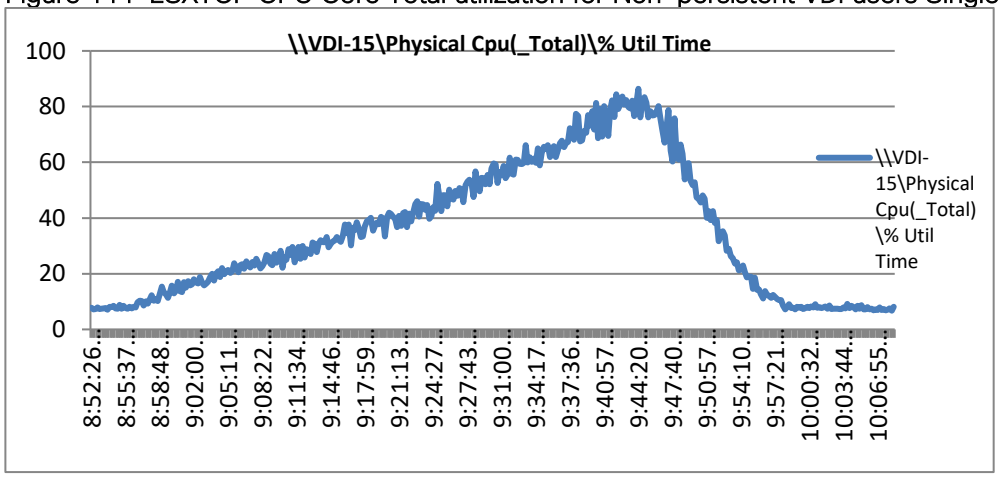

Figure 112 Esxtop for Non-persistent VDI users Single Server testing Network Adapter MBits Received /Sent /Sec

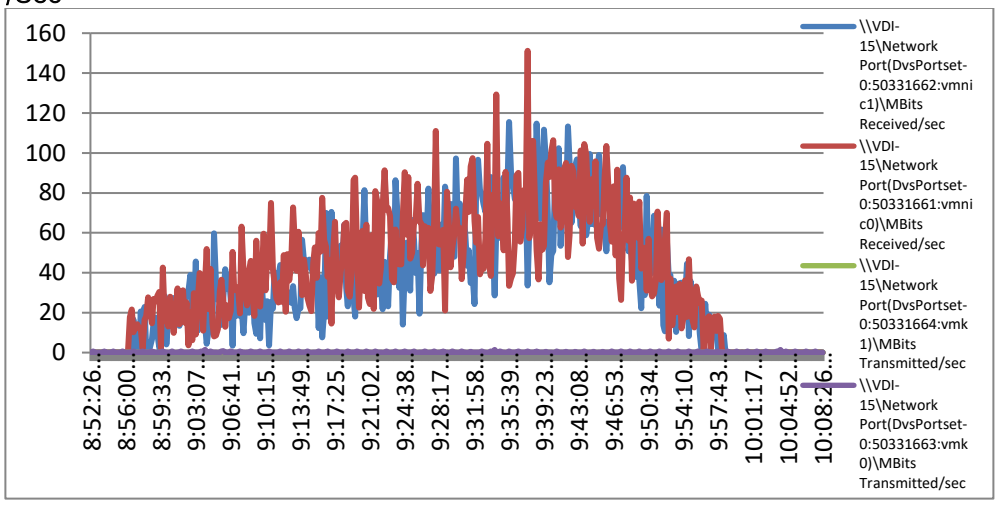

Figure 113 Esxtop for Non-persistent VDI users Single Server testing Disk Adapter Mbytes Received /Sent /Sec

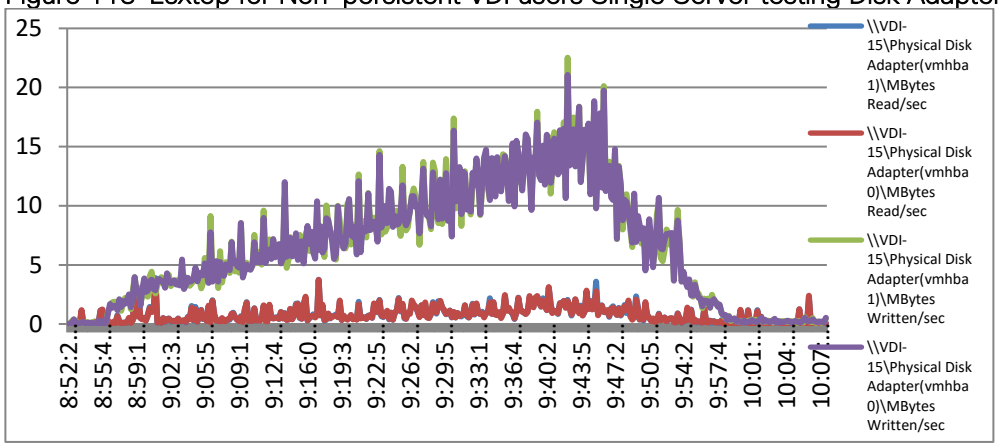

Figure 114 ESXTOP for Non-persistent VDI users Single Server testing Non-Kernel Mbytes

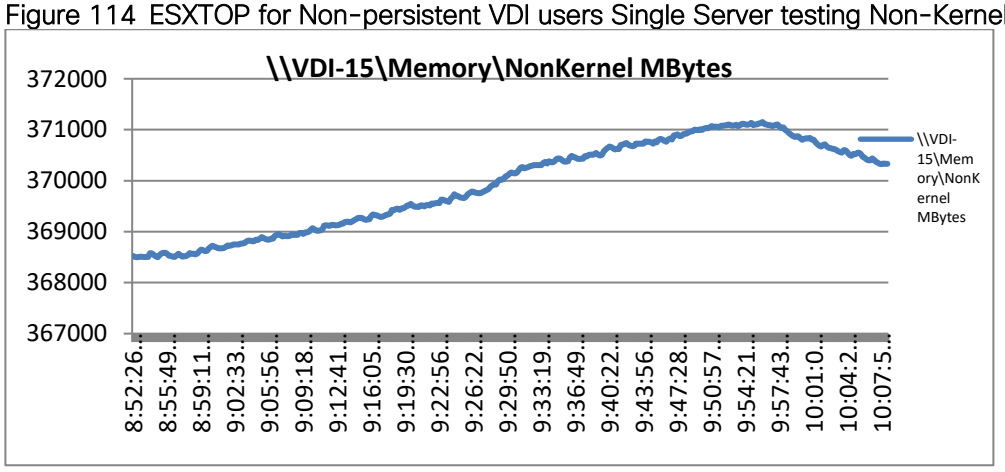

Cluster Testing for 1450 Horizon RDSH Sessions

Testing of the Horizon RDSH cluster with server N+1 fault tolerance at the cluster level demonstrated outstanding end user experience while effectively utilizing the compute, memory, network and storage resources.

Figure 115 RDS cluster configured with 72 Windows 2012 R2 RDSH virtual servers hosting 1450 User Sessions on 9 B200 M4 Servers (N+1)

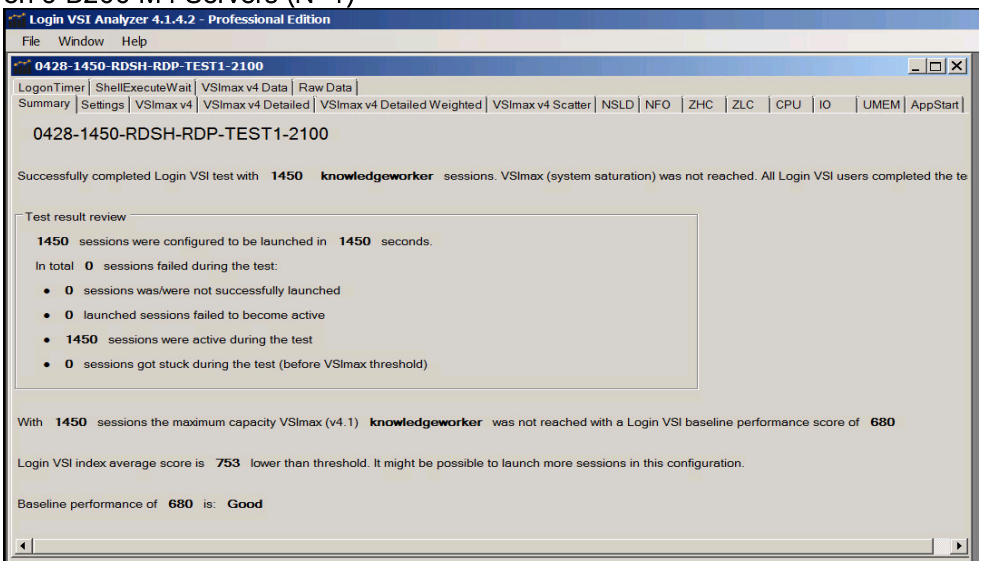

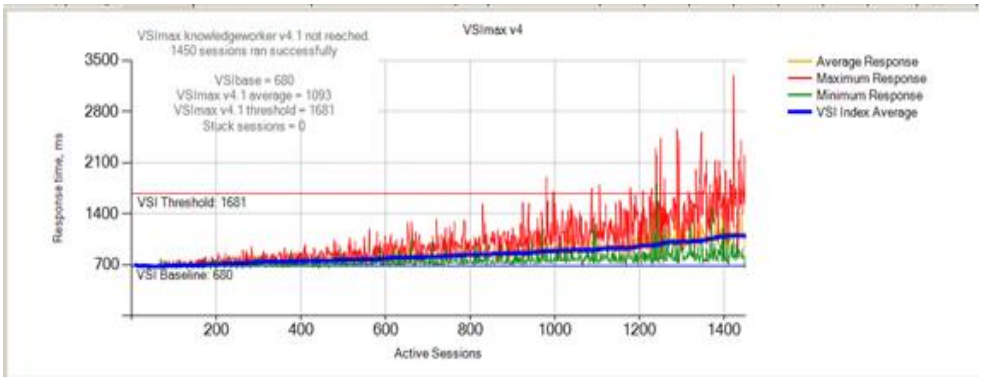

Figure 116 RDSH-Cluster testing 1450 User Sessions
ESXTOP CPU Core utilization for one of RDSH-Server on RDSH Cluster testing

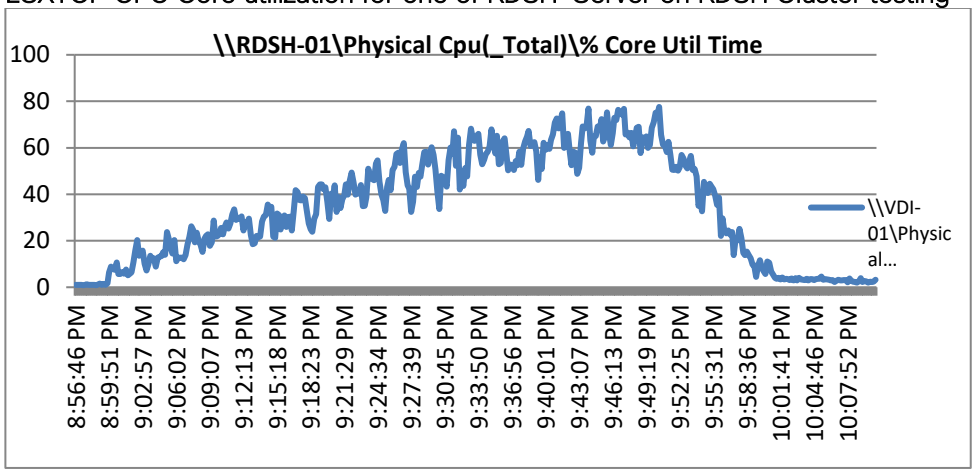

Figure 117 ESXTOP CPU Total Utilization for one of RDSH-Server on RDSH Cluster testing

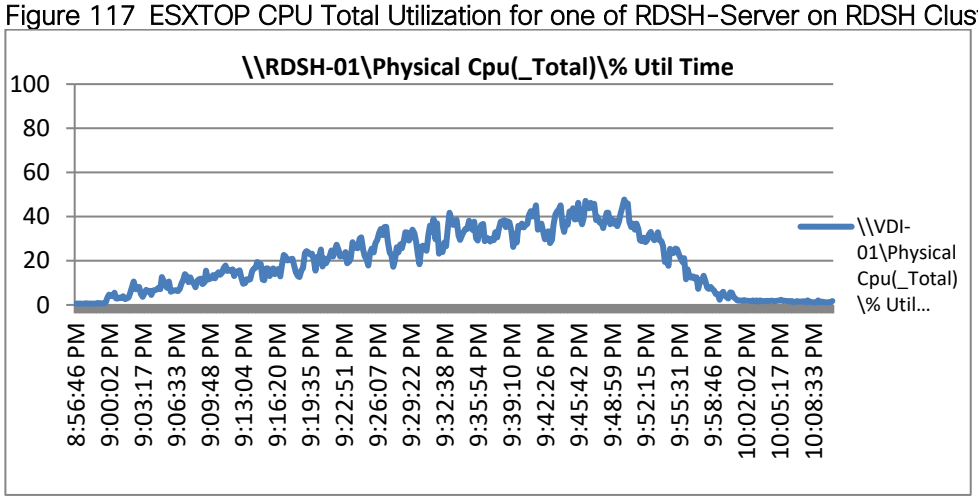

Figure 118 ESXTOP Network Adapter Mbytes Received /Sent /Sec for one of RDSH Server on Cluster testing

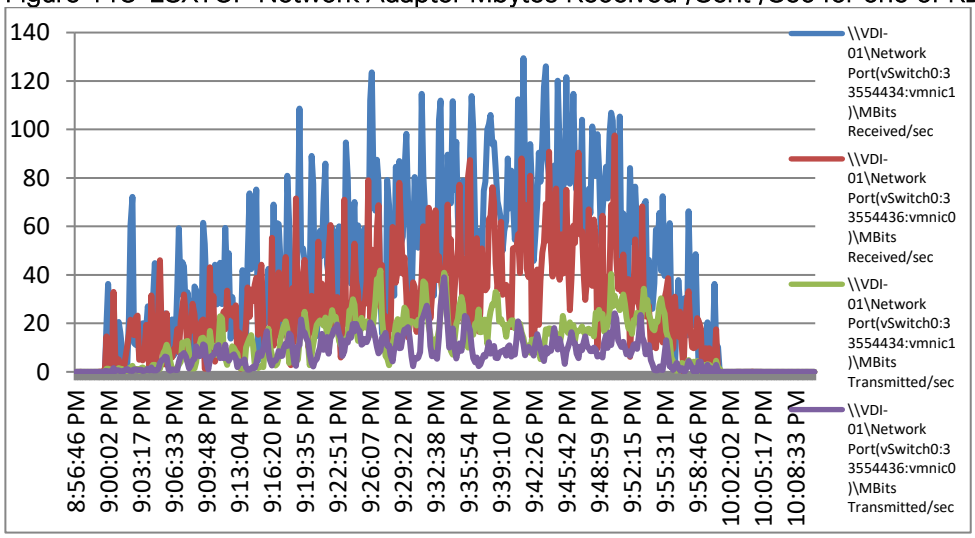

Figure 119 ESXTOP Physical Disc Adapter Mbytes Read/Written /Sec for one of RDSH-Server on RDSH Cluster testing

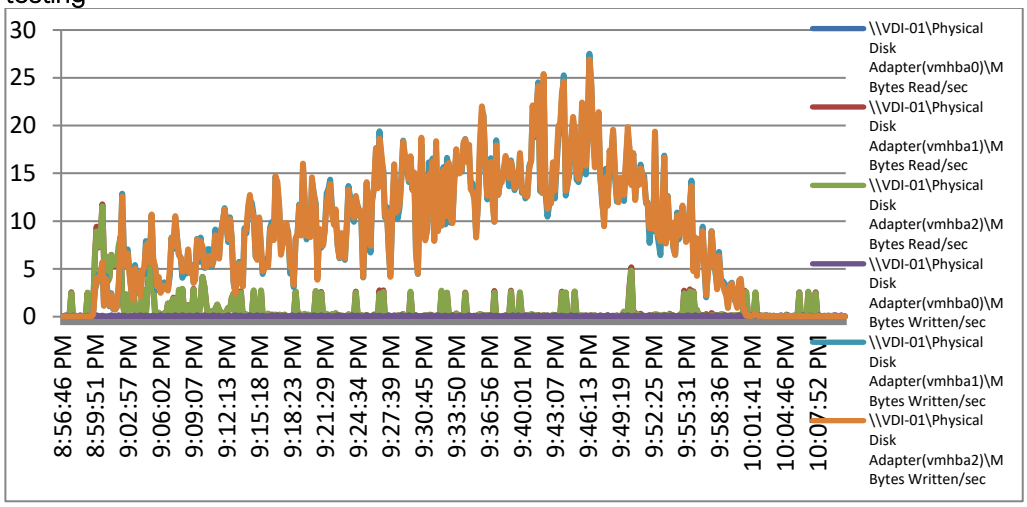

Figure 120 ESXTOP Non-Kernel Mbytes for one of RDSH-Server on RDSH Cluster testing

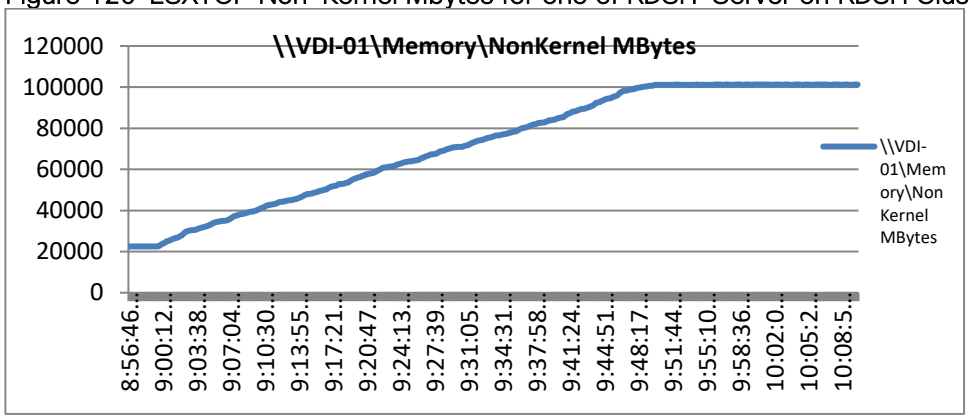

Storage Charts for 1450 User Cluster Testing

Figure 121 Read /Write Latency for 1450 Users RDSH testing

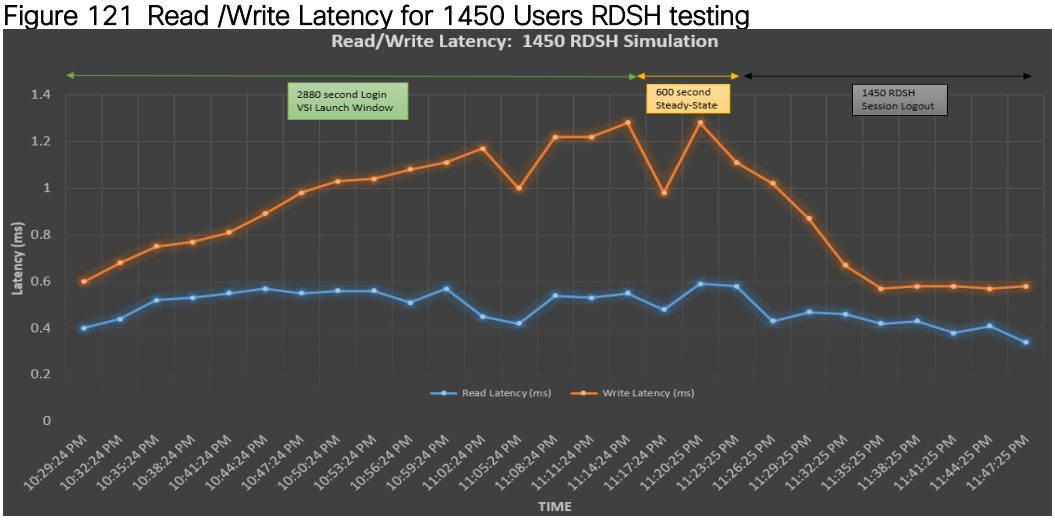

Figure 122 IOPs for 1450 Users RDSH testing

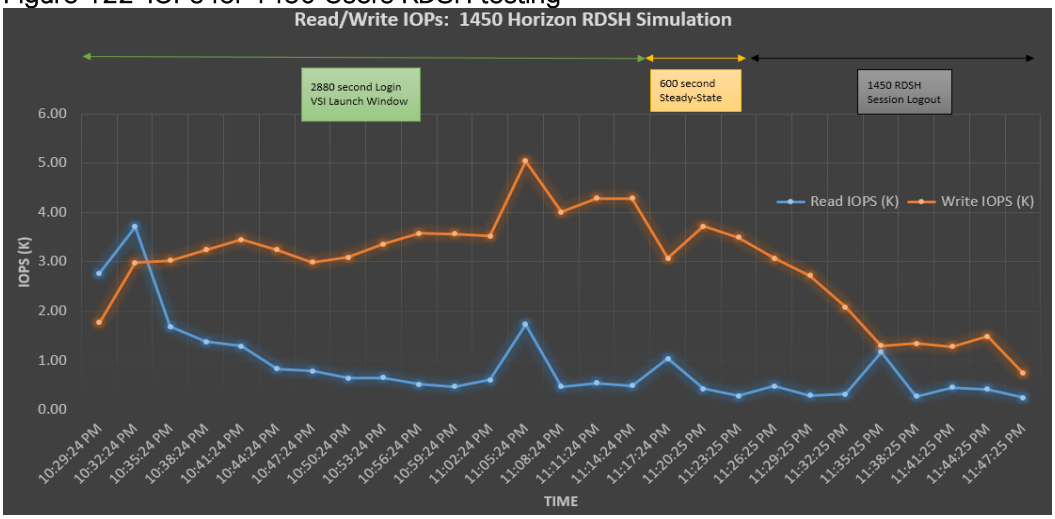

Figure 123 Read /Write Bandwidth for 1450 Users RDSH testing

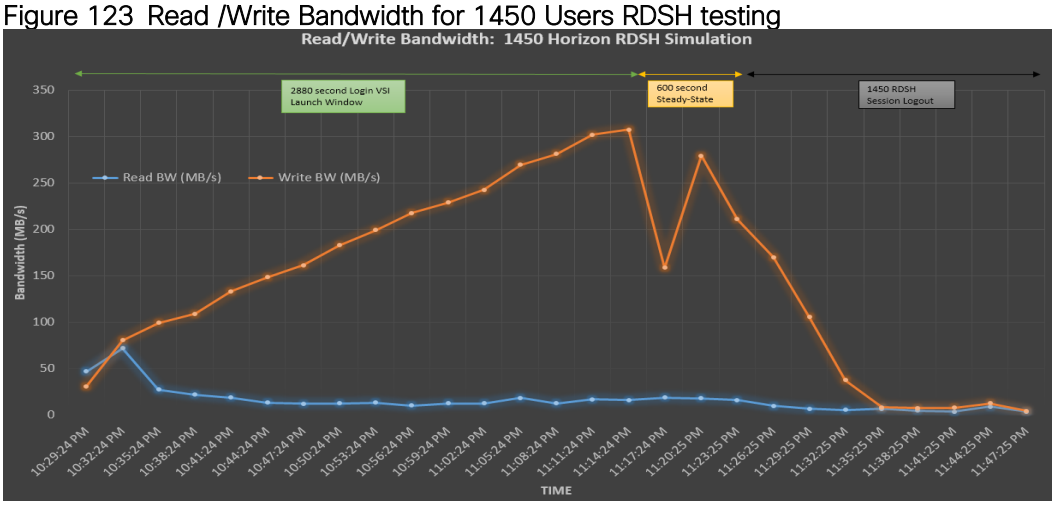

Horizon Linked Clone Cluster Testing with 3550 Users

For testing purposes we have created 2 identical VDI pools with 1775 Users on each pool. A total of 21 Cisco UCS B200 M4 servers were used to host 3550 VDI virtual machines in a single vCenter cluster. This cluster provides N+1 server fault tolerance for the Horizon Linked Clone workload.

Figure 124 Login VSI graph for 3550 Users Linked Clones Cluster testing

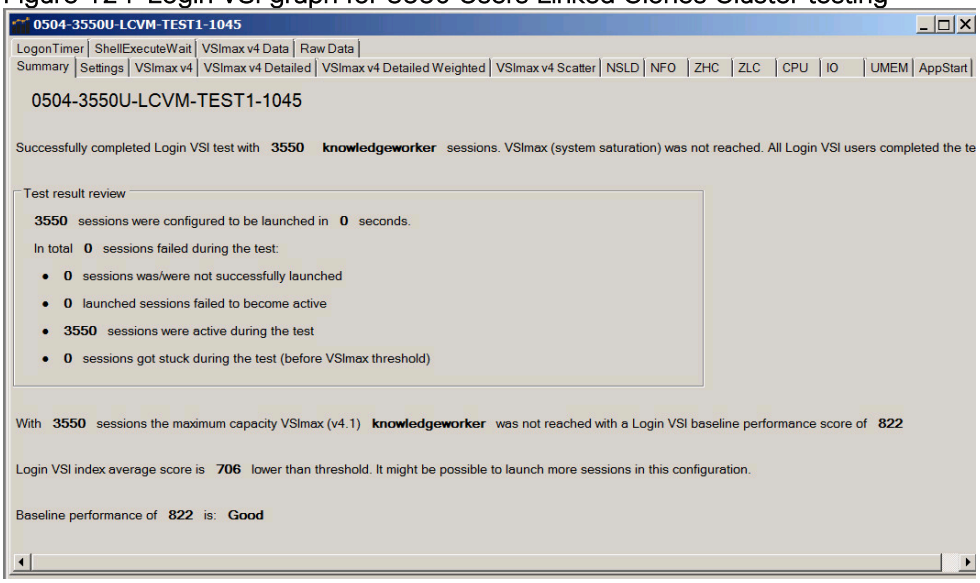

Figure 125 3550 Users Linked Clones Cluster testing Login VSI Chart

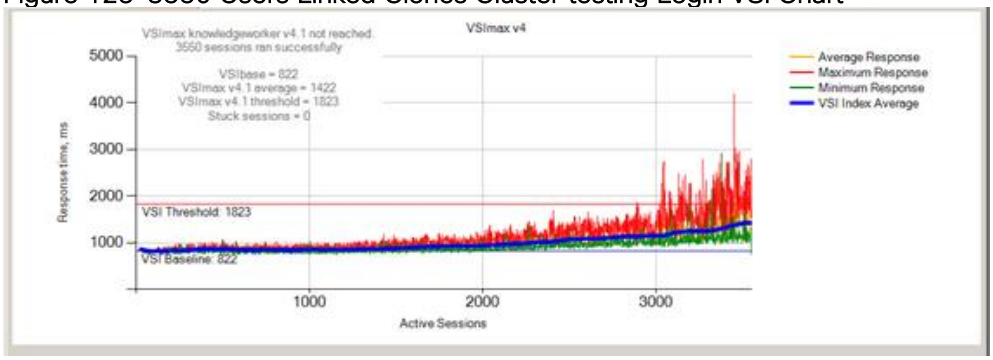

Figure 126 3550 Users CPU Core Utilization for one of VDI server

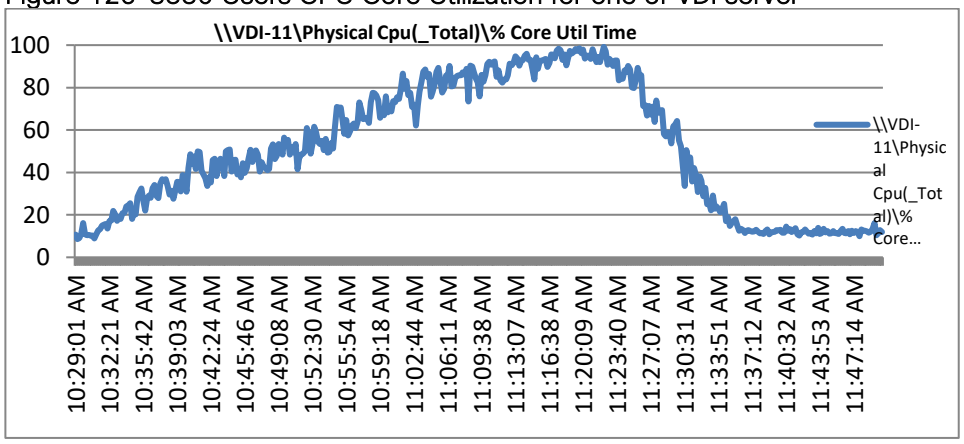

Figure 127 3550 Users CPU Total Utilization for one of VDI server

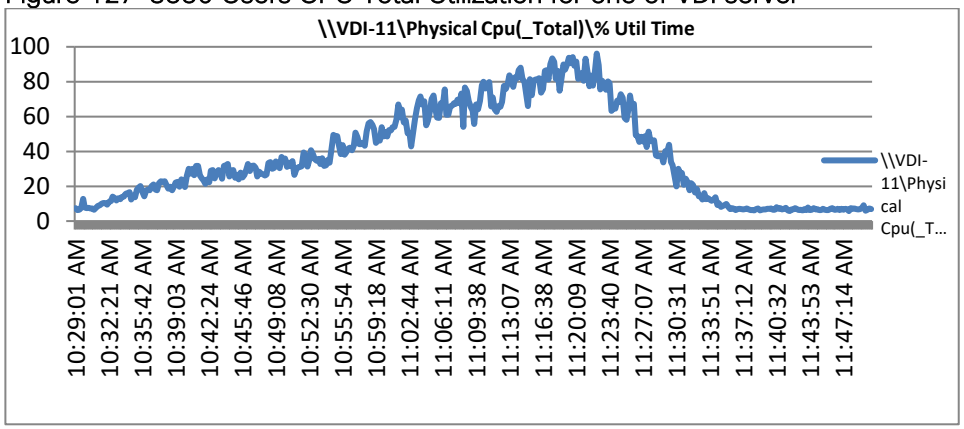

Figure 128 3550 Users network Mbytes received /transmitted per sec for one of VDI server

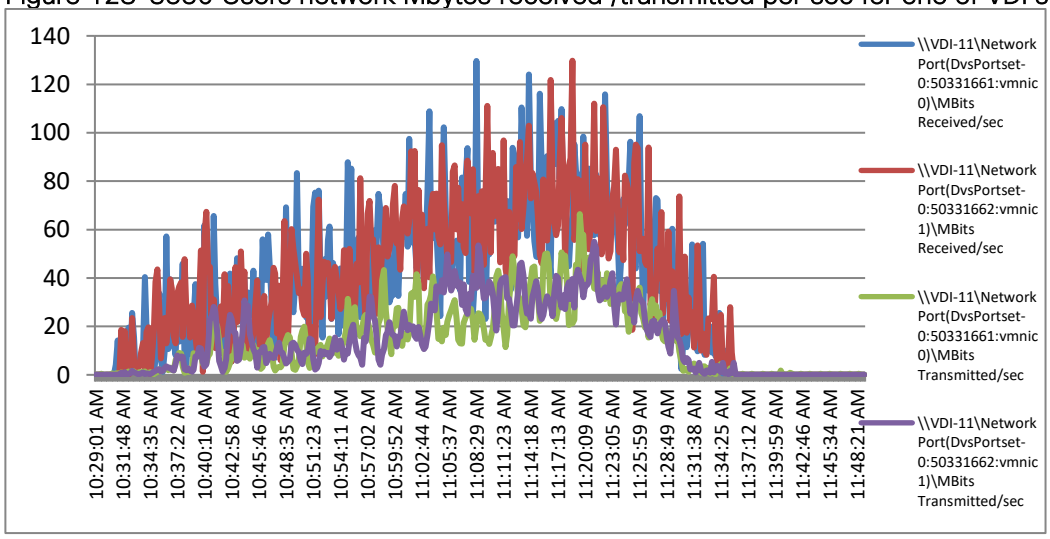

Figure 129 3550 Users vHBA disk adapter Mbytes read /written per sec for one of VDI server

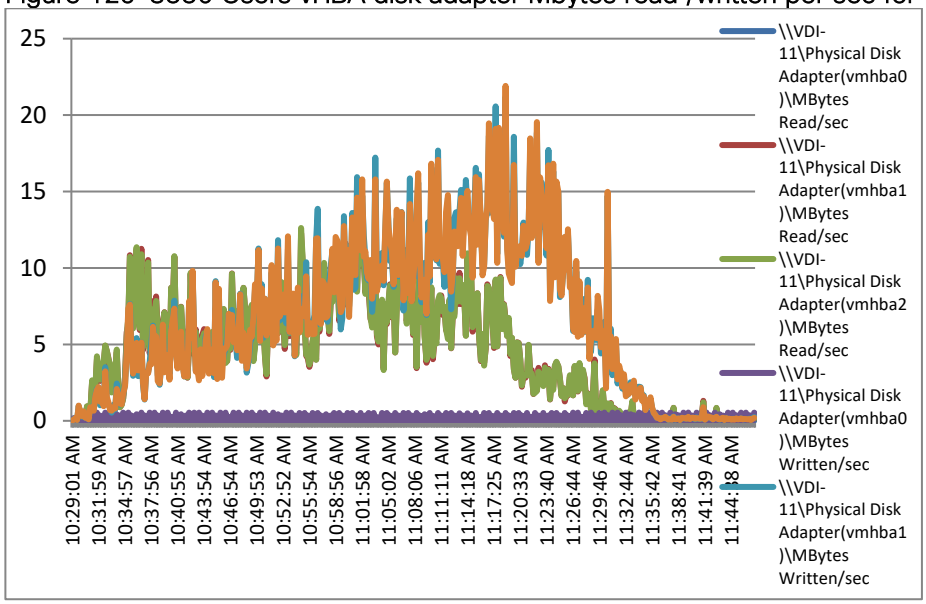

Figure 130 3550 Users Non-Kernel Mbytes for one VDI servers

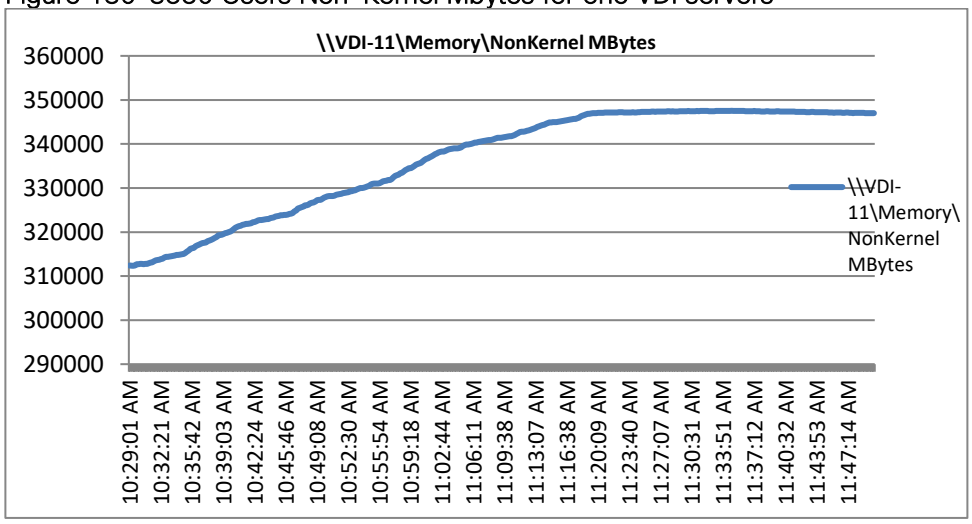

Read /Write Latency for 3550 Users Linked VDI Testing

Figure 131 Storage Charts for 3550 users Linked Clone VDI testing

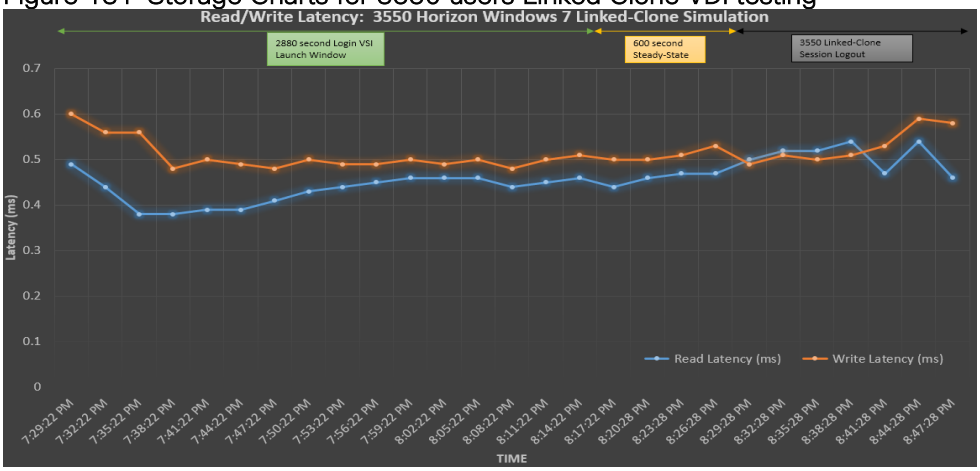

Figure 132 Read / Write IOPs for 3550 Users Linked Clone VDI testing

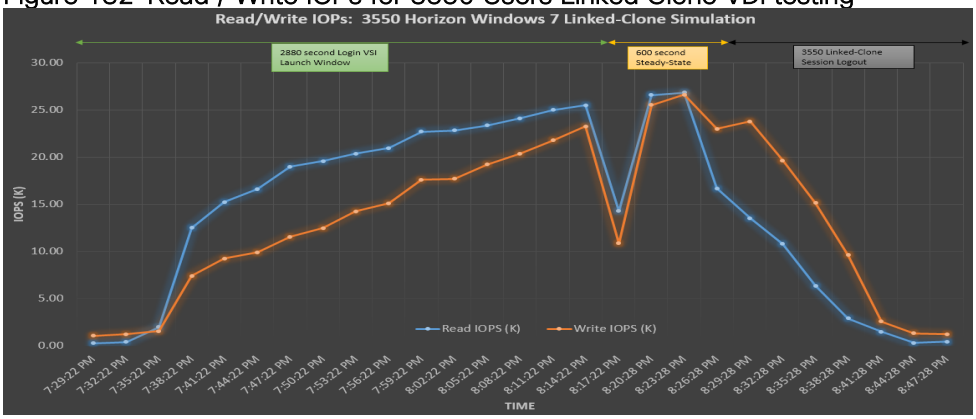

Figure 133 Read /Write bandwidth for 3550 Users Liked clones Testing

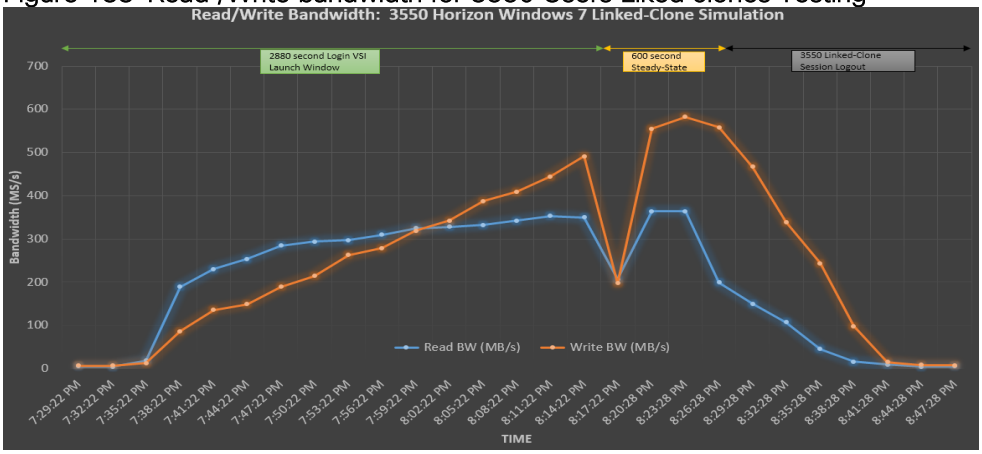

5000 Users Mixed Horizon RDSH and Linked Clone Workload Testing

1450 RDS Hosted Server Sessions along with 3550 Non-Persistent VDI-Linked Clone Virtual Machine Tested Concurrently as 5000 Users Combined Mixed Workload

Figure 134 5000 Users Mixed Work Load Login VSI Chart

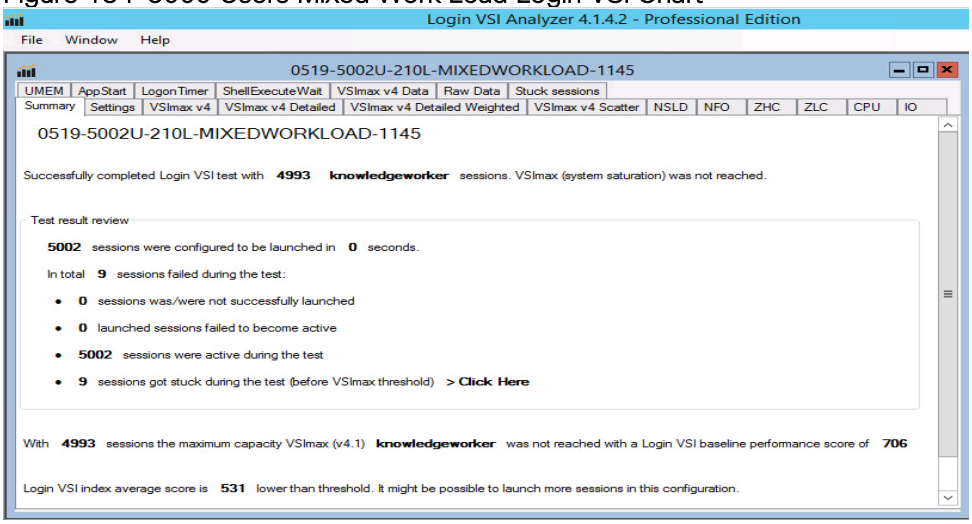

Figure 135 5000 Users Mixed workload LoginVSI Chart

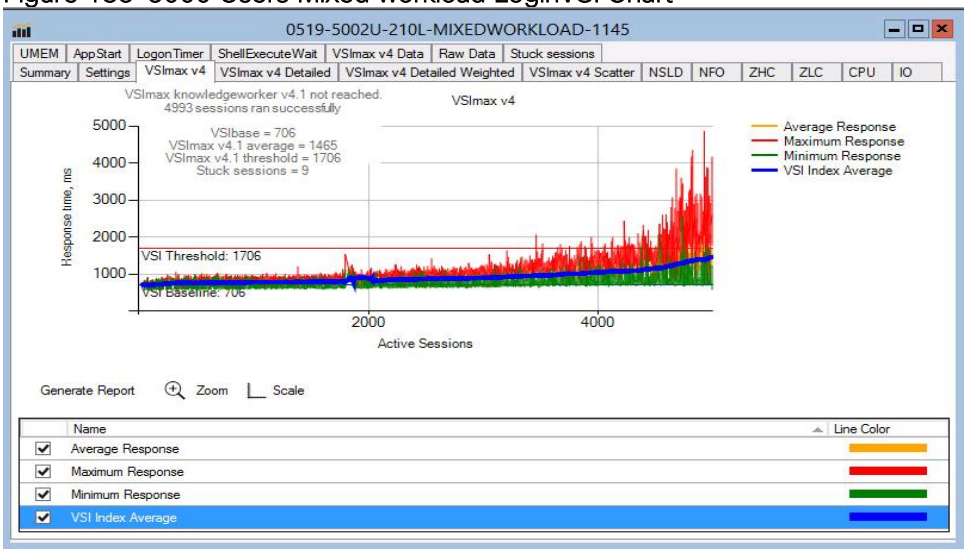

Figure 136 5000 Users Mixed workload CPU Core Util Time for VDI-Server

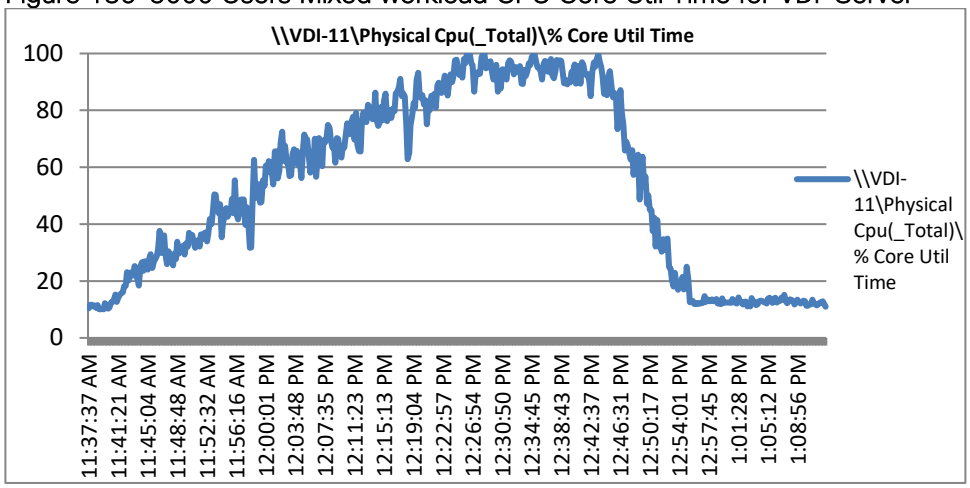

Figure 137 5000 Users Mixed workload CPU Total Util Time for VDI-Server

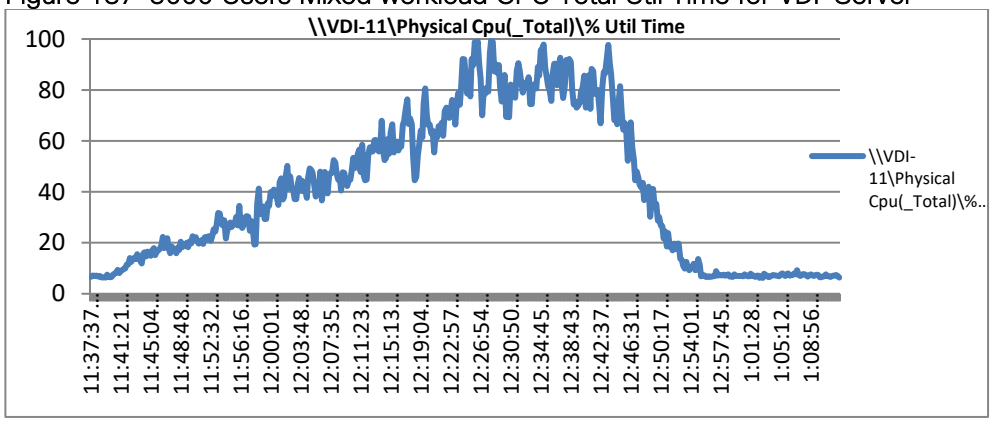

Figure 138 5000 Users Mixed workload Network Received /Transmitted sec for VDI-Server

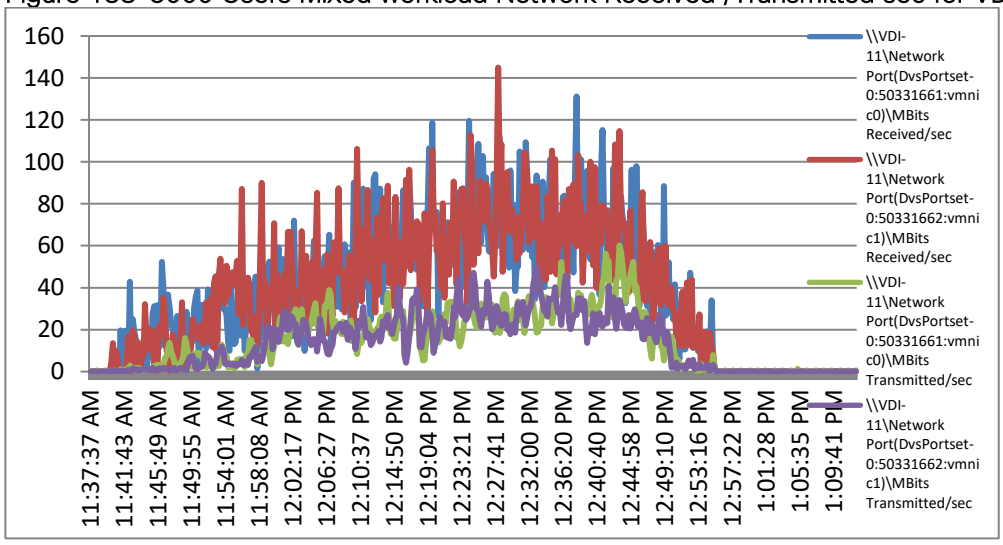

Figure 139 5000 Users Mixed workload Disk Adapter Mbytes for VDI-Server Disk Adapter

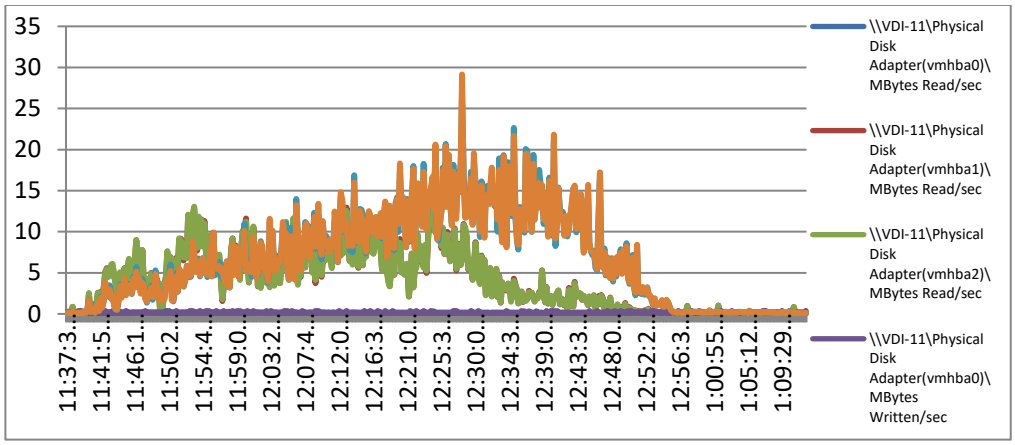

Figure 140 5000 Users Mixed workload Non-Kernel Memory Mbytes for VDI Server

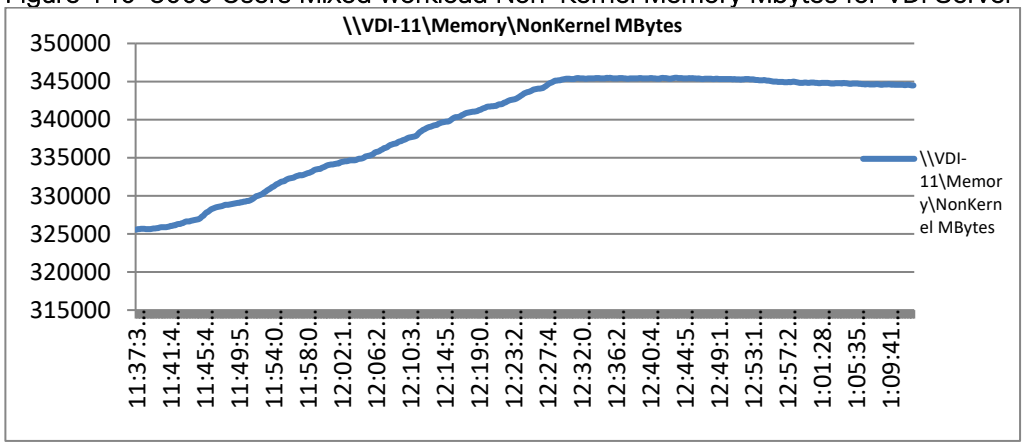

Figure 141 Storage Graphs for 5000 User Mixed Workload

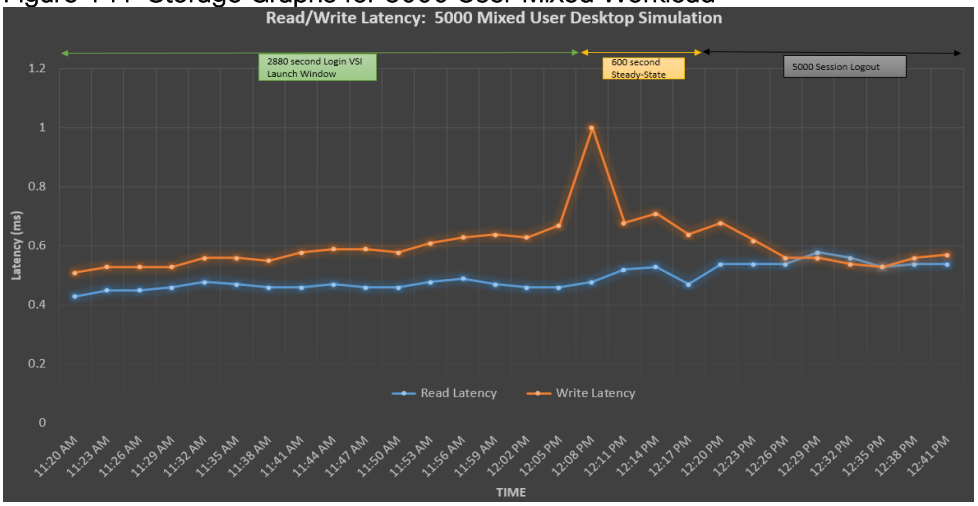

Figure 142 Read / Write IOPs for 5000 Users Mixed Workload

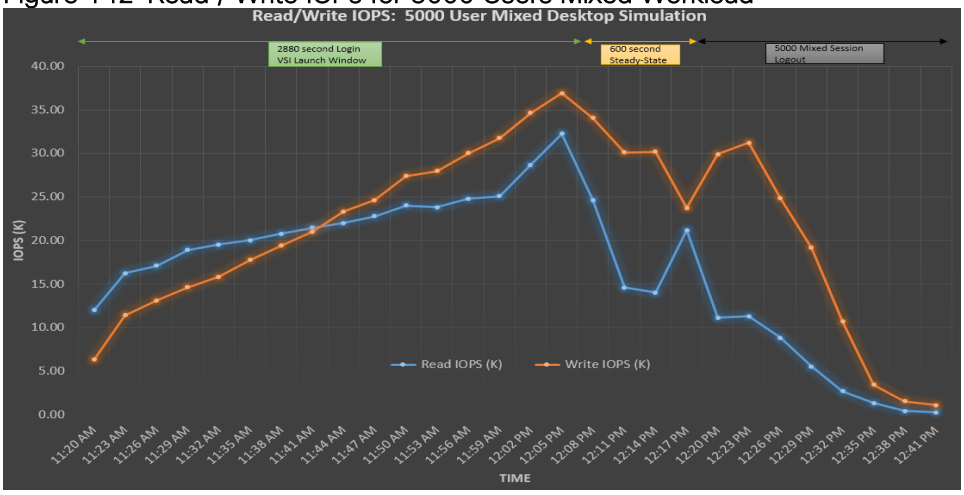

Figure 143 Read / write Bandwidth for 5000 Users Mixed Workload

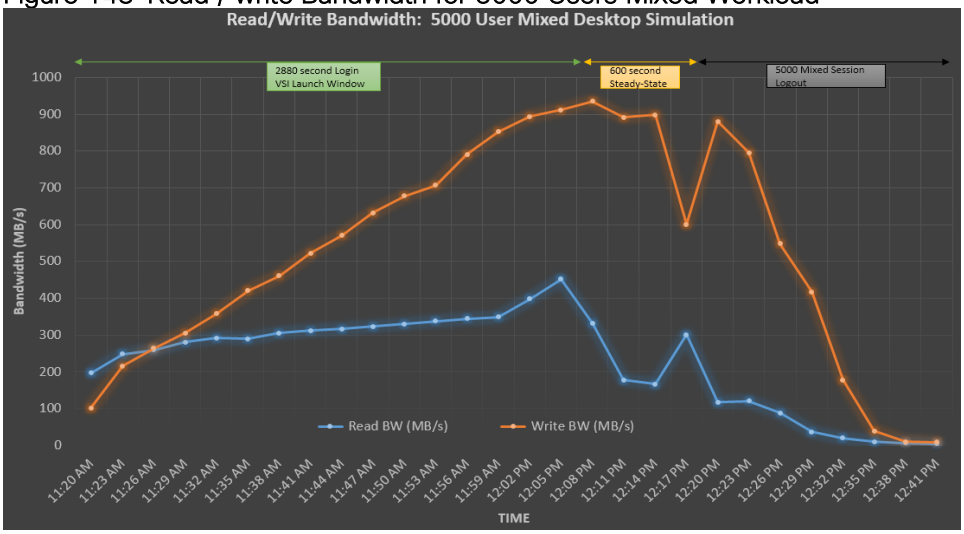

Perfmon Charts for 5000 Users Mixed Workload Test for VMware View Connection Server

View Connection Server Perfmon for 50000 Users Mixed Workload

Figure 144 5000 Users Processor User time View Connection Server Perfmon

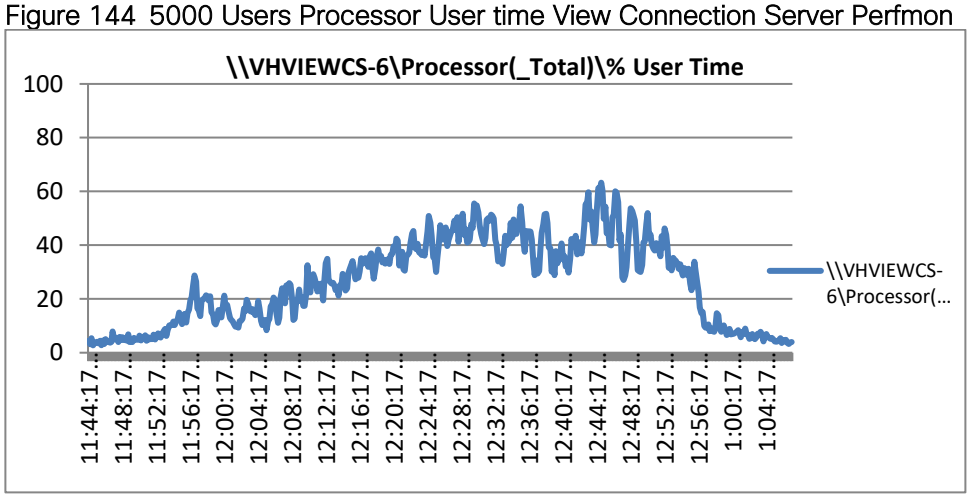

Figure 145 5000 Users Processor time View Connection Server Perfmon

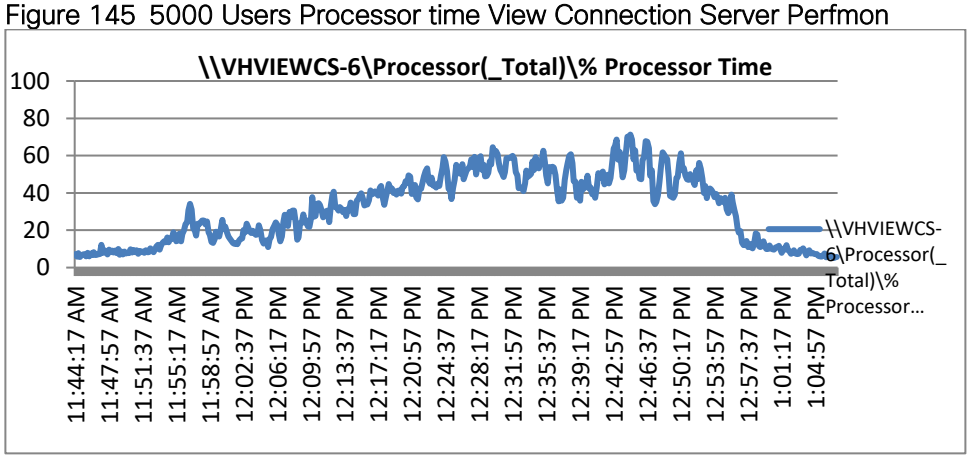

Figure 146 5000 Users Available Mbytes View Connection Server Perfmon

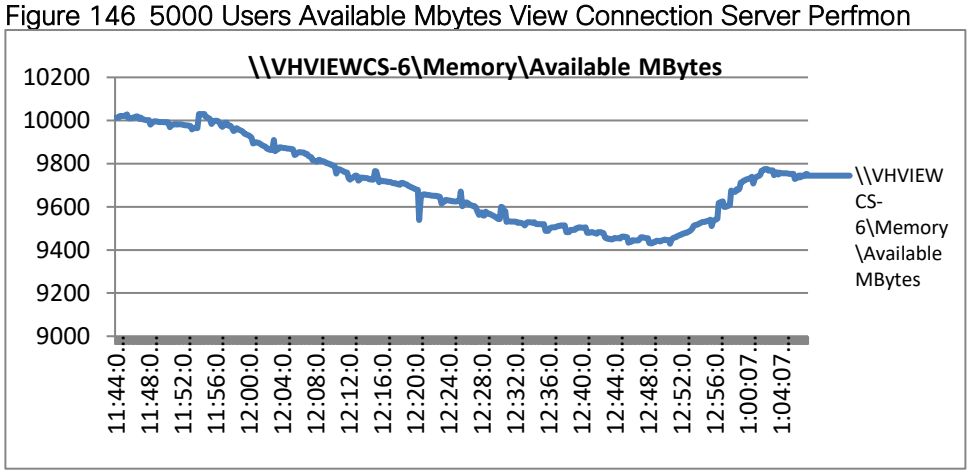

Figure 147 5000 Users View Connection Server Network Bytes Received /sec Perfmon

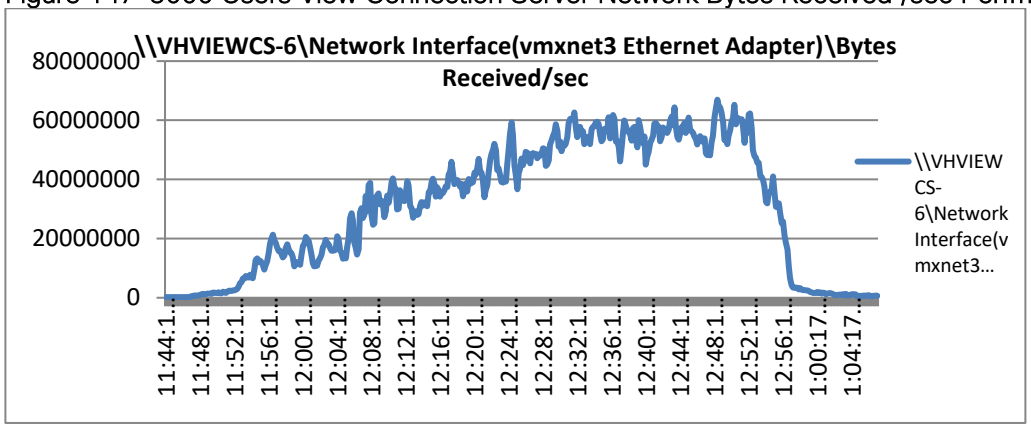

Figure 148 5000 Users View Connection Server Network Bytes Sent /sec

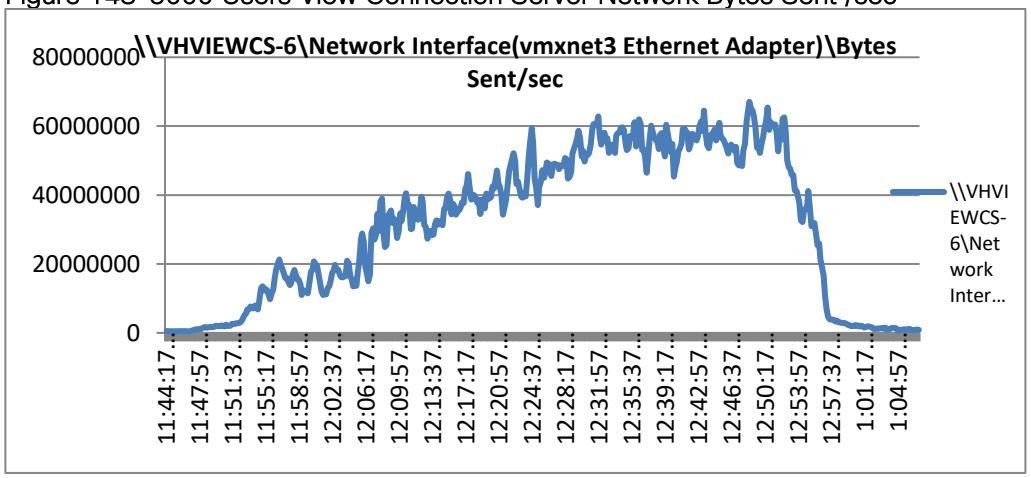

Pure Storage FlashArray//m50 Test Results for Full Scale, Mixed Workload Resiliency Testing

For the final simulation, we elected to highlight the tremendous resiliency of the Pure Storage FlashArray//m50 by performing an upgrade of the Purity Operating Environment in parallel with repeating the previous 5000 desktop mixed-workload simulation. In this test, we upgraded from Purity v4.5.5 to Purity v4.6.8 in the middle of the Login VSI simulation while desktops were still logging in to the environment and running the Login VSI Knowledge Worker workload.

Worth noting is that this exact procedure would also be followed during a controller upgrade as part of the Evergreen Storage business model where customers receive the latest Pure Storage controllers for free every three years so long as a valid maintenance agreement is in place. Another supported scenario would be non-disruptively upgrading to the more performant FlashArray//m70 following this same procedure. Pure Storage controllers are stateless and do not require any additional setup other than inserting them into the chassis and connecting cables. Relatedly, capacity expansions (adding a shelf or denser SSD drives) are also accomplished without any downtime in the middle of production operations.

Upgrading the Purity Operating Environment is typically handled by Pure Storage support for arrays that are managed via Cloud Assist. Upgrades can also be accomplished locally by KVM or console connection to the array as well by a Pure Storage or other authorized support customer. To accomplish the upgrade, the process is extremely simple, usually takes less than 20 minutes and achieved with the following steps:

1. Copy upgraded Purity Operating Environment code to both array controllers
2. Install new Purity code on secondary controller
3. Reboot secondary controller, wait for it to come back online
4. Install new Purity code on primary controller
5. Reboot primary controller, array non-disruptively fails over to upgraded secondary controller which becomes primary and new version of Purity OE is in use
6. Rebooted primary controller comes up as secondary running upgraded Purity code

The following charts show the array behavior during the Purity code upgrade. Immediately worth noting is that the drop in reported metrics to 0 during the upgrade process was a bug in the Purity code that has since been correct in version 4.7.0 (for this test we upgraded from Purity 4.5.5 to Purity 4.6.8). The Login VSI results that will be shown in the section provide independent verification that the upgrade was non-disruptive and non-impactful to the end-users running in the environment.

The chart below shows the latency of the array during the Purity upgrade process window shown by the orange box. Overall the upgrade process took approximately 14 minutes to complete including controller reboots.

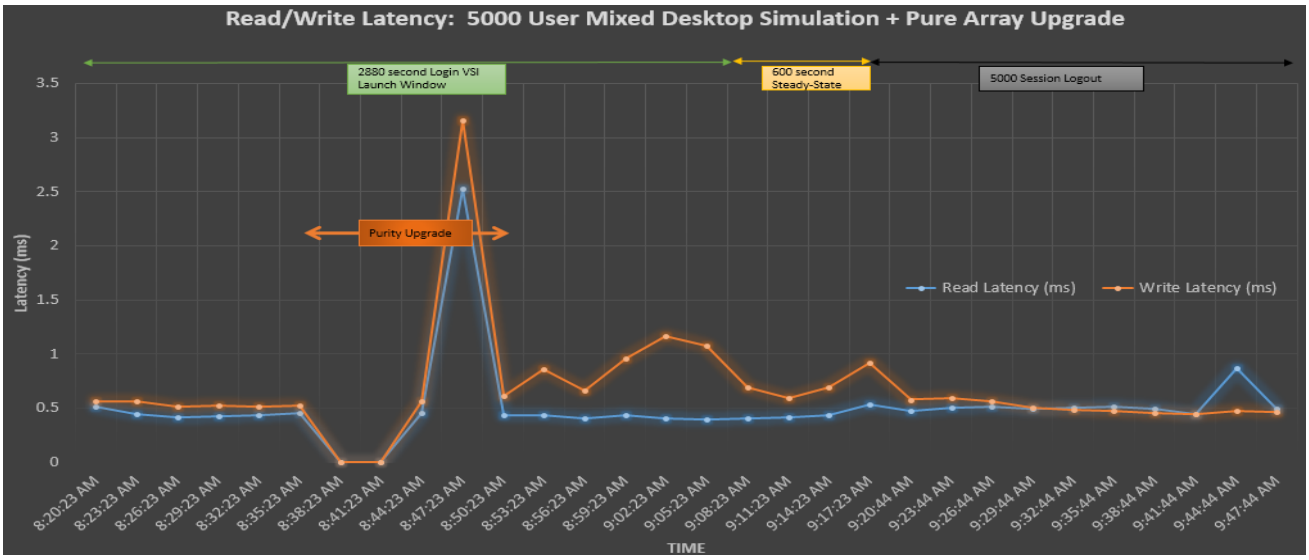

Next we show IOPs being served during the 5000 mixed workload simulation with the Purity upgrade.

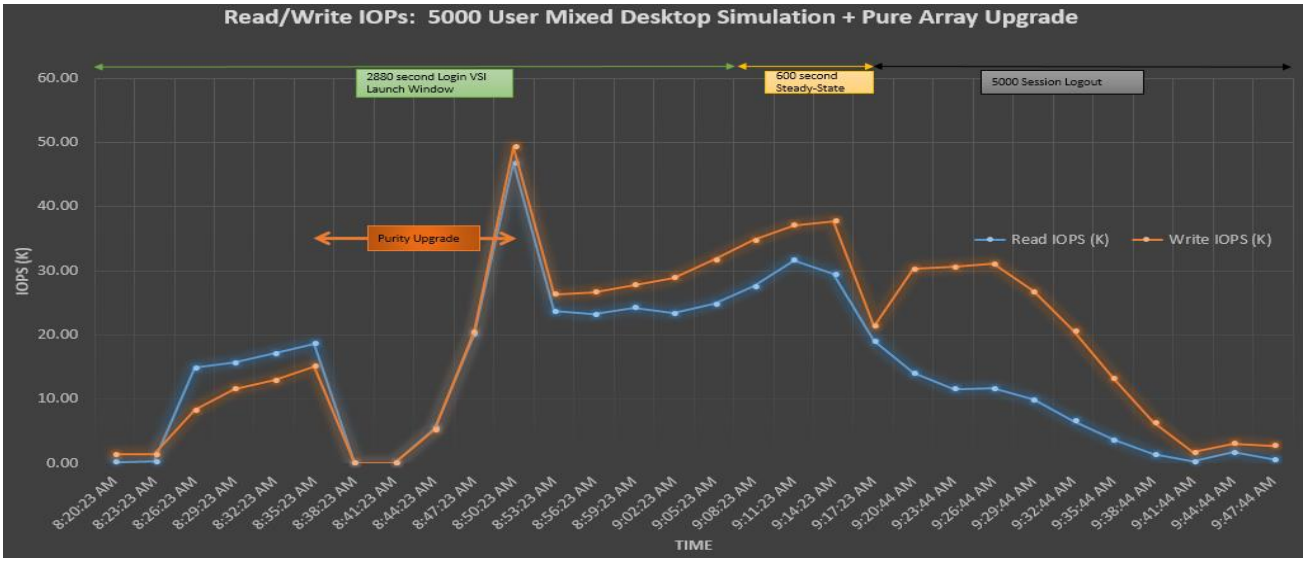

Lastly, this chart shows our bandwidth during the mixed workload + Pure Storage FlashArray//m50 upgrade process.

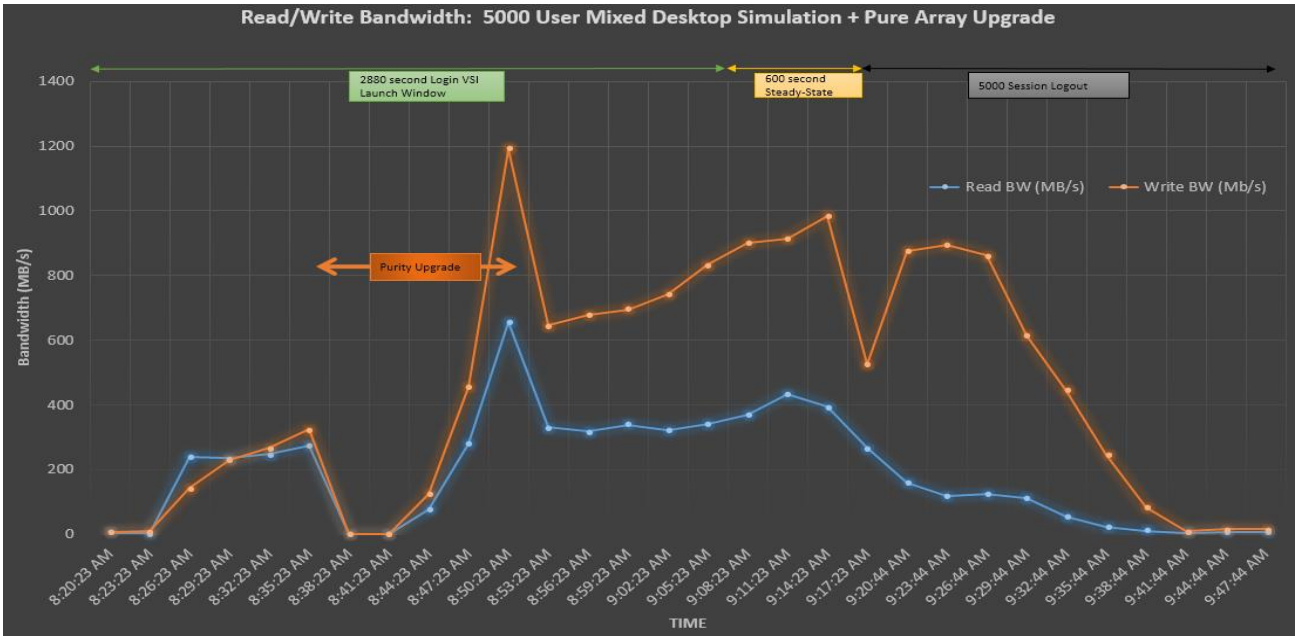

Login VSI provides impartial verification that the Purity upgrade was non-disruptive. You can see from the below performance summary chart from the tool that no end-user performance degradation or large number

of disconnected sessions occurred while the Pure Storage array was upgraded.

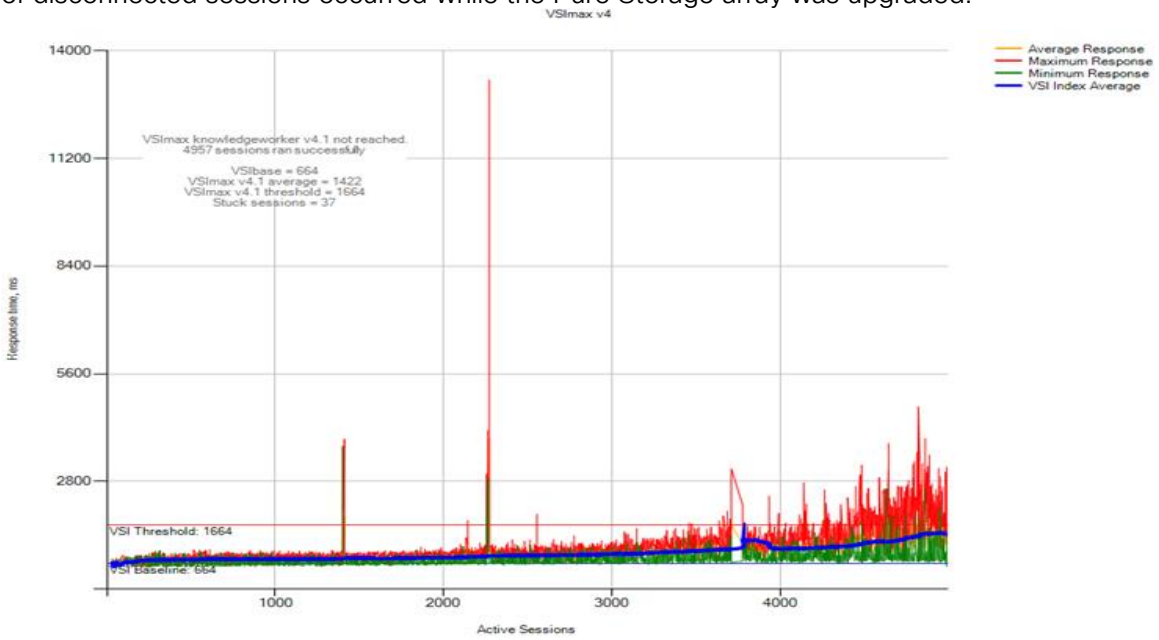

Lastly, you can see from the Login VSI performance charts that relative performance from the 5000 Mixed Workload simulation and the 5000 Mixed Workload simulation that included the Purity upgrade are essentially identical with only brief spikes in response time during the upgrade process that likely would not be noticeable by the end-user:

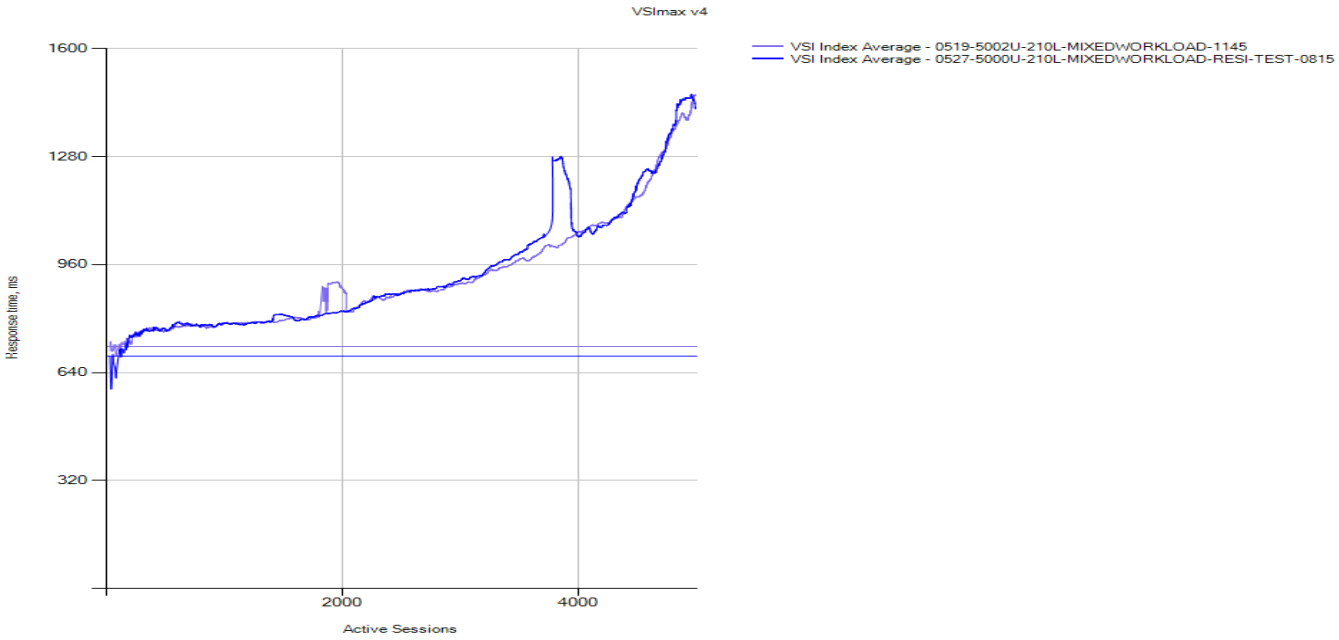

The ability of Pure Storage to provide non-disruptive operations such as the above example cannot be understated. No longer do storage administrators need to schedule downtime with the VMware Horizon team for maintenance operations and as we have clearly demonstrated above – maintenance operations can now be accomplished in the middle of the workday without causing any interruption to a huge amount of active VMware Horizon users.

Summary

FlashStack delivers a platform for Enterprise VDI deployments and cloud datacenters using Cisco UCS Blade and Rack Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches and fibre channel-attached Pure Storage FlashArray//m. FlashStack is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlashStack provides to customers wishing to deploy enterprise-class VDI for 5000 users at a time.

Get More Business Value with Services

Whether you are planning your next-generation environment, need specialized know-how for a major deployment, or want to get the most from your current storage, Pure 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, Pure provides in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Authors

Ramesh Guduru, Cisco Systems, Inc.

Ramesh Guduru is a Technical Marketing Engineer at Cisco Systems with over 12 years of experience in the field of information technology. Ramesh's focus primarily on design and implementation of Virtual Desktop Infrastructure projects for data center, VMware View thin client administration, configuration and optimization of virtual desktop environment, Data center reference architectures, Cisco Unified Computing System and Storage design and implementation. Ramesh's skill set include core VMware applications in the virtual environment focusing in system design and implementation of virtualization components, solutions validation, technical content creation and testing/benchmarking.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the significant contribution and expertise that resulted in developing this document:

- Kyle Grossmiller, Solutions Architect, Customer Solutions Group, Pure Storage, Inc.
- Mike Brennan, Manager, Technical Marketing, Desktop Virtualization Solutions Team, Computing Systems Product Group, 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-b200-m4-blade-server/index.html
- http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html

Cisco UCS Manager Configuration Guides

- http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html
- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/release/notes/ucs_2_2_rn.html
- 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
- http://www.cisco.com/c/en/us/support/storage-networking/mds-9148s-16g-multilayer-fabric-switch/model.html
- http://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9148-multilayer-fabric-switch/data_sheet_c78-571411.html
- http://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9148s-16g-multilayer-fabric-switch/datasheet-c78-731523.html

VMware References

- https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
- http://pubs.vmware.com/Release_Notes/en/horizon-6-view/horizon-62-view-release-notes.html
- https://pubs.vmware.com/horizon-62-view/index.jsp

- https://pubs.vmware.com/horizon-62-view/index.jsp?topic=%2Fcom.vmware.horizon-view.desktops.doc%2FGUID-DFAD071A-7F60-4720-86AB-8F1597BFC95C.html
- http://pubs.vmware.com/vsphere-60/index.jsp
- http://www.vmware.com/files/pdf/view/vmware-horizon-view-best-practices-performance-study.pdf
- https://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf
- https://labs.vmware.com/flings/vmware-os-optimization-tool

Microsoft References

- https://technet.microsoft.com/en-us/library/hh831620(v=ws.11).aspx
- https://technet.microsoft.com/en-us/library/dn281793(v=ws.11).aspx
- https://support.microsoft.com/en-us/kb/2833839
- https://technet.microsoft.com/en-us/library/hh831447(v=ws.11).aspx

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

Pure Storage Reference Documents

- Pure Storage FlashArray//m Datasheet
  - http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/PureStorage_FlashArraym-Brochure.pdf
- Pure Storage FlashStack Converged Infrastructure Solutions
  - http://www.purestorage.com/solutions/infrastructure/flashstack.html
- Pure Storage Best Practices Guide for VMware vSphere
  - http://www.purestorage.com/resources/type-a/WP-PureStorageandVMwarevSphereBestPracticesGuide_Request.html
- Pure Storage and VMware Storage APIs for Array Integration
  - http://www.purestorage.com/resources/type-b/WP-PureStorageandVAAI_Request.html
- FlashStack Converged Infrastructure for VMware vSphere Design Guide
  - http://www.purestorage.com/resources/type-a/WP-FlashStackRefArch-VSI_Request.html

- FlashStack Converged Infrastructure for VMware Horizon 6.2 Design Guide
  - http://info.purestorage.com/WP-FlashStackRefArch-VDI_Request.html
- Consolidating Workloads with VMware and Pure Storage
  - http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/ESG_Lab_Validation_Summary_Pure_Storage_Sep_2015.pdf

Appendix A – Cisco Nexus Ethernet and MDS Fibre Channel Switch Configurations

Ethernet Network Configuration

The following section provides a detailed procedure for configuring the Cisco Nexus 9000 and 1000V Switches used in this study.

Cisco Nexus 9172PX-A Configuration

```
!Time: Wed Jun 29 20:44:02 2016

version 6.1(2)I3(3a)

switchname N9K-A

vdc N9K-A id 1

  allocate interface Ethernet1/1-54

  limit-resource vlan minimum 16 maximum 4094

  limit-resource vrf minimum 2 maximum 4096

  limit-resource port-channel minimum 0 maximum 512

  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

username admin password 5 $1$7LGx6.qz$3xYFFA2B9CgCGt0n3EOm60  role network-admin

ssh key rsa 2048
```

```
ip domain-lookup
policy-map type qos jumbo
copp profile strict
snmp-server user admin network-admin auth md5 0x4112400e904ff685e2625e9e302ec9ad
priv 0x4112400e904ff685e2625e9e302ec9ad localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.81.254.202
vlan 1,160-164,166
vlan 160
  name IB_Mgmt
vlan 161
  name infra
vlan 162
  name VDI
vlan 164
  name OOB-Mgmt
vlan 166
  name vMotion
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
hardware qos ns-buffer-profile mesh
```

```
vpc domain 10
  peer-switch
  role priority 10
  peer-keepalive destination 10.29.164.12 source 10.29.164.11
  delay restore 150
  peer-gateway
  auto-recovery
interface Vlan1
  no ip redirects
  no ipv6 redirects
interface Vlan 160
description IB-Mgmt vlan 160

  no ip redirects
  ip address 10.10.160.2/24
  hsrp version 2
  hsrp 11
   preempt
   priority 90
   ip 10.10.160.1
  no shutdown
interface Vlan161
description Infra vlan 161
 no ip redirects
  ip address 10.10.161.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 12
  preempt
  priority 80
  ip 10.10.161.1
  no shutdown
```

```
interface Vlan162
description VDI vlan 162
no ip redirects
 ip address 10.10.192.2/19
  hsrp version 2
  hsrp 13
    preempt
    priority 82
    ip 10.10.192.1
  ip dhcp relay address 10.10.161.30
  ip dhcp relay address 10.10.161.31
  no shutdown
interface Vlan164
  no ip redirects
  ip address 10.29.164.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 14
    preempt
    priority 80
    ip 10.29.164.1
  no shutdown
interface Vlan166
Description vMotion Vlan 166
  no ip redirects
  ip address 10.29.166.3/24
  no ipv6 redirects
  hsrp version 2
  hsrp 14
    preempt
    priority 80
    ip 10.29.166.1
```

```
  no shutdown
switchport trunk allowed vlan 160-166
switchport trunk allowed vlan 160-164,166
spanning-tree port type edge trunk
vpc 29
interface port-channel31
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  spanning-tree port type edge trunk
vpc 31
interface Ethernet1/1
interface Ethernet1/2
interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
 switchport access vlan 160
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
speed 1000
interface Ethernet1/20
```

```
interface Ethernet1/21
  switchport access vlan 161
  spanning-tree port type edge
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 29 mode active
interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 29 mode active
interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 31 mode active
interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 31 mode active
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
```

```
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
  description VPC Peer N9K-B:1/47
  switchport mode trunk
  switchport trunk allowed vlan 160-166
  channel-group 10 mode active
interface Ethernet1/48
  description VPC Peer N9K-B:1/48
  switchport mode trunk
  switchport trunk allowed vlan 160-166
  channel-group 10 mode active
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
  vrf member management
  ip address 10.29.164.11/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.3a.bin
 N9K-A#
```

Cisco Nexus 9172PX-B Configuration

```
N9K-B# show run
!Command: show running-config
!Time: Wed Jun 29 20:36:20 2016
version 6.1(2)I3(3a)
switchname N9K-B
vdc N9K-B id 1
  allocate interface Ethernet1/1-54
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 512
  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
username admin password 5 $1$yJzJso32$HOIWoF6/1m/wuLBQuVXRw/  role network-admin
ssh key rsa 2048
ip domain-lookup
system default switchport shutdown
copp profile strict
snmp-server user admin network-admin auth md5 0x5e43c86fb0bd3b4040c115c79b62ef3a
priv 0x5e43c86fb0bd3b4040c115c79b62ef3a 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
vlan 1,160-164,166
vlan 160
  name IB-Mgmt vLAN 160
vlan 161
 name infra vLAN 161
vlan 162
  name VDI vLAN 162
vlan 164
  name OOB-Mgmt vLAN 164
vlan 166
  name vMotion vLAN 166
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
 ip dhcp relay
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.164.1
port-channel load-balance src-dst l4port
hardware qos ns-buffer-profile mesh
vpc domain 10
  role priority 20
  peer-keepalive destination 10.29.164.11 source 10.29.164.12
  delay restore 150
  peer-gateway
  auto-recovery
```

```
 interface Vlan1
  no ip redirects
  no ipv6 redirects
interface Vlan160
description OOB-Mgmt vlan 160
  no ip redirects
  ip address 10.10.160.3/24
  hsrp version 2
  hsrp 11
    preempt
    priority 80
    ip 10.10.160.1
  no shutdown
interface Vlan161
description Infra Vlan 161
  no ip redirects
  ip address 10.10.161.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 12
    preempt
    priority 90
    ip 10.10.161.1
  no shutdown
interface Vlan162
description VDI Vlan 162
  no ip redirects
  ip address 10.10.192.3/19
  no ipv6 redirects
  hsrp version 2
  hsrp 13
    preempt
```

```
    priority 92
    ip 10.10.192.1
  hsrp 192
    preempt
    priority 121
  ip dhcp relay address 10.10.161.30
  ip dhcp relay address 10.10.161.31
  no shutdown
 interface Vlan164
description OOB-Vlan 164
  no ip redirects
  ip address 10.29.164.2/24
  no ipv6 redirects
  hsrp version 2
  hsrp 14
    preempt
    priority 90
    ip 10.29.164.1
  no shutdown
interface port-channel5
  description vPC peer-link
interface port-channel10
  description VPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 160-166
  spanning-tree port type network
  vpc peer-link
interface port-channel29
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  spanning-tree port type edge trunk
  vpc 29
```

```
interface port-channel31
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166-167
  spanning-tree port type edge trunk
  vpc 31
interface Ethernet1/1
  no shutdown
interface Ethernet1/2
  no shutdown
interface Ethernet1/3
  no shutdown
interface Ethernet1/4
  no shutdown
interface Ethernet1/5
  no shutdown
interface Ethernet1/6
  no shutdown
interface Ethernet1/7
  no shutdown
interface Ethernet1/8
  no shutdown
interface Ethernet1/9
  no shutdown
interface Ethernet1/10
  no shutdown
interface Ethernet1/11
  no shutdown
interface Ethernet1/12
  no shutdown
interface Ethernet1/13
  no shutdown
interface Ethernet1/14
```

```
  no shutdown
interface Ethernet1/15
  no shutdown
interface Ethernet1/16
  no shutdown
interface Ethernet1/17
  switchport access vlan 160
  no shutdown
interface Ethernet1/18
  no shutdown
interface Ethernet1/19
  switchport access vlan 160
  speed 1000
  no shutdown
interface Ethernet1/20
  no shutdown
interface Ethernet1/21
  switchport access vlan 161
  spanning-tree port type edge
  no shutdown
interface Ethernet1/22
  no shutdown
interface Ethernet1/23
  no shutdown
interface Ethernet1/24
  no shutdown
interface Ethernet1/25
  no shutdown
interface Ethernet1/26
  no shutdown
interface Ethernet1/27
  no shutdown
```

```
interface Ethernet1/28
  no shutdown
interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 29 mode active
  no shutdown
interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 29 mode active
  no shutdown
interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 31 mode active
  no shutdown
interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 160-164,166
  channel-group 31 mode active
  no shutdown
interface Ethernet1/33
  no shutdown
interface Ethernet1/34
  no shutdown
interface Ethernet1/35
  no shutdown
interface Ethernet1/36
  no shutdown
interface Ethernet1/37
  no shutdown
```

```
interface Ethernet1/38
  no shutdown
interface Ethernet1/39
  no shutdown
interface Ethernet1/40
  no shutdown
interface Ethernet1/41
  no shutdown
interface Ethernet1/42
  no shutdown
interface Ethernet1/43
  no shutdown
interface Ethernet1/44
  no shutdown
interface Ethernet1/45
  no shutdown
interface Ethernet1/46
  no shutdown
interface Ethernet1/47
  description VPC Peer N9K-A:1/47
  switchport mode trunk
  switchport trunk allowed vlan 160-167
  channel-group 10 mode active
  no shutdown
interface Ethernet1/48
  description VPC Peer N9K-A:1/48
  switchport mode trunk
  switchport trunk allowed vlan 160-167
  channel-group 10 mode active
  no shutdown
interface Ethernet1/49
  no shutdown
```

```
interface Ethernet1/50
  no shutdown
interface Ethernet1/51
  no shutdown
interface Ethernet1/52
  no shutdown
interface Ethernet1/53
  no shutdown
interface Ethernet1/54
  no shutdown
interface mgmt0
  vrf member management
  ip address 10.29.164.12/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.3a.bin
 N9K-B#
```

Fibre Channel Network Configuration

Cisco MDS 9148S-A Configuration

```
 MDS-A# show run
MDS-A# show run[Jning-config  !Command: show running-config
!Time: Tue Jun 28 22:50:59 2016
version 6.2(9a)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
```

```
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $1$loX7vizP$00IbhSFcpx6WufBmOMKB.1  role network-admin
ip domain-lookup
ip host MDS-A  10.29.164.64
aaa group server radius radius
snmp-server user admin network-admin auth md5 0x6c81eb7167a2e69497a60698ca3957da
 priv 0x6c81eb7167a2e69497a60698ca3957da localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1163
snmp-server host 10.29.164.130 traps version 2c public udp-port 1163
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vsan database
  vsan 20
device-alias database
 vsan 20 wwn 20:00:00:25:b5:1f:1a:24 fcid 0xc20511 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:0c fcid 0xc2041c dynamic
  vsan 20 wwn 52:4a:93:72:0d:21:6b:11 fcid 0xc20000 dynamic
  vsan 20 wwn 52:4a:93:72:0d:21:6b:00 fcid 0xc20100 dynamic
  vsan 20 wwn 20:4f:54:7f:ee:45:29:80 fcid 0xc20200 dynamic
  vsan 20 wwn 20:50:54:7f:ee:45:29:80 fcid 0xc20300 dynamic
  vsan 20 wwn 20:4f:54:7f:ee:45:2a:40 fcid 0xc20400 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:1d fcid 0xc20401 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:17 fcid 0xc20406 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:15 fcid 0xc20404 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:09 fcid 0xc20408 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:07 fcid 0xc20502 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:05 fcid 0xc20504 dynamic
  vsan 20 wwn 20:00:00:25:b5:b1:1b:03 fcid 0xc20402 dynamic
```

```
vsan 20 wwn 20:00:00:25:b5:b1:1b:0b fcid 0xc20503 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:0d fcid 0xc20509 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:13 fcid 0xc20405 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:11 fcid 0xc20508 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:7f fcid 0xc20416 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:21 fcid 0xc20510 dynamic
vsan 20 wwn 20:50:54:7f:ee:45:2a:40 fcid 0xc20500 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:1e fcid 0xc20407 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:1b fcid 0xc20403 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:19 fcid 0xc20506 dynamic
vsan 20 wwn 20:00:00:25:b5:b1:1b:00 fcid 0xc20501 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:00 fcid 0xc2040d dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:02 fcid 0xc2040e dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:04 fcid 0xc20410 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:06 fcid 0xc20507 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:08 fcid 0xc20513 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:1e fcid 0xc2050c dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:0e fcid 0xc20411 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:0a fcid 0xc2041b dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:1c fcid 0xc2040c dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:10 fcid 0xc20517 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:20 fcid 0xc20505 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:22 fcid 0xc20412 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:1a fcid 0xc20409 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:16 fcid 0xc2041a dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:14 fcid 0xc2050b dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:26 fcid 0xc2050a dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:12 fcid 0xc20413 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:2a fcid 0xc2040f dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:2e fcid 0xc2050d dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:30 fcid 0xc2040b dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:2c fcid 0xc20418 dynamic
```

```
  vsan 20 wwn 20:00:00:25:b5:1f:1a:32 fcid 0xc2050f dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:34 fcid 0xc20419 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:36 fcid 0xc2050e dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:38 fcid 0xc20515 dynamic
  vsan 20 wwn 20:00:00:25:b5:1f:1a:18 fcid 0xc20415 dynamic
vsan 20 wwn 52:4a:93:72:0d:21:6b:13 fcid 0xc20600 dynamic
vsan 20 wwn 20:00:00:25:b5:1f:1a:28 fcid 0xc20414 dynamic
interface port-channel1
channel mode active
switchport rate-mode dedicated
vsan database
vsan 20 interface fc1/1
vsan 20 interface fc1/2
vsan 20 interface fc1/3
vsan 20 interface fc1/4
switchname MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
```

```
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
```

```
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
!Active Zone Database Section for vsan 20
zone name Infra-serv1-fc0 vsan 20
member pwwn 20:00:00:25:b5:b1:1b:21
member pwwn 52:4a:93:72:0d:21:6b:11
member pwwn 52:4a:93:72:0d:21:6b:00
zone name Infra-serv2-fc0 vsan 20
member pwwn 20:00:00:25:b5:b1:1b:7f
member pwwn 52:4a:93:72:0d:21:6b:11
member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv1-fc0 vsan 20
   member pwwn 20:00:00:25:b5:b1:1b:00
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv2-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:00
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv3-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:02
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv5-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:06
    member pwwn 52:4a:93:72:0d:21:6b:11
```

```
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv6-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:08
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv8-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:0c
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv7-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:0a
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv9-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:0e
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv10-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:10
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv11-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:12
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv12-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:14
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv13-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:16
    member pwwn 52:4a:93:72:0d:21:6b:11
```

```
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv16-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:1c
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv17-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:1e
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv18-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:20
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv19-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:22
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv20-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:24
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv21-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:26
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv24-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:2c
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv25-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:2e
    member pwwn 52:4a:93:72:0d:21:6b:11
```

```
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv26-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:30
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv27-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:32
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv28-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:34
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv29-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:36
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv30-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:38
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv15-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:1a
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv14-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:18
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv23-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:2a
    member pwwn 52:4a:93:72:0d:21:6b:11
```

```
    member pwwn 52:4a:93:72:0d:21:6b:00
zone name pure-vdi-serv22-fc0 vsan 20
    member pwwn 20:00:00:25:b5:1f:1a:28
    member pwwn 52:4a:93:72:0d:21:6b:11
    member pwwn 52:4a:93:72:0d:21:6b:00
zoneset name pure-vdi-fab-A vsan 20
    member Infra-serv1-fc0
    member Infra-serv2-fc0
    member pure-vdi-serv4-fc0
    member pure-vdi-serv1-fc0
    member pure-vdi-serv2-fc0
    member pure-vdi-serv3-fc0
    member pure-vdi-serv5-fc0
    member pure-vdi-serv6-fc0
    member pure-vdi-serv8-fc0
    member pure-vdi-serv7-fc0
    member pure-vdi-serv9-fc0
    member pure-vdi-serv10-fc0
    member pure-vdi-serv11-fc0
    member pure-vdi-serv12-fc0
    member pure-vdi-serv13-fc0
    member pure-vdi-serv16-fc0
    member pure-vdi-serv17-fc0
    member pure-vdi-serv18-fc0
    member pure-vdi-serv19-fc0
    member pure-vdi-serv20-fc0
    member pure-vdi-serv21-fc0
    member pure-vdi-serv24-fc0
    member pure-vdi-serv25-fc0
    member pure-vdi-serv26-fc0
    member pure-vdi-serv27-fc0
    member pure-vdi-serv28-fc0
```

```
    member pure-vdi-serv29-fc0
    member pure-vdi-serv30-fc0
    member pure-vdi-serv15-fc0
    member pure-vdi-serv14-fc0
    member pure-vdi-serv23-fc0
    member pure-vdi-serv22-fc0
zoneset activate name pure-vdi-fab-A vsan 20
do clear zone database vsan 20
!Full Zone Database Section for vsan 20
interface fc1/1
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/2
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/3
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/4
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/5
port-license acquire
interface fc1/6
port-license acquire
interface fc1/7
no port-license
no shutdown
```

```
interface fc1/8
no port-license
no shutdown
interface fc1/9
port-license acquire
no shutdown
interface fc1/10
port-license acquire
no shutdown
interface fc1/11
port-license acquire
channel-group 1 force
no shutdown
interface fc1/12
port-license acquire
channel-group 1 force
no shutdown
interface fc1/13
port-license acquire
no shutdown
interface fc1/14
port-license acquire
no shutdown
interface fc1/15
port-license acquire
no shutdown
interface fc1/16
port-license acquire
no shutdown
interface fc1/17
port-license acquire
interface fc1/18
```

```
port-license acquire
interface fc1/19
port-license acquire
interface fc1/20
  port-license acquire
interface fc1/21
port-license acquire
interface fc1/22
port-license acquire
interface fc1/23
port-license acquire
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
interface fc1/26
port-license acquire
 interface fc1/27
port-license acquire
interface fc1/28
port-license acquire
interface fc1/29
port-license acquire
interface fc1/30
port-license acquire
interface fc1/31
port-license acquire
interface fc1/32
interface fc1/33
port-license acquire
interface fc1/34
port-license acquire
```

```
interface fc1/35
port-license acquire
interface fc1/36
port-license acquire
interface fc1/37
port-license acquire
interface fc1/38
port-license acquire
interface fc1/39
port-license acquire
interface fc1/40
port-license acquire
interface fc1/41
port-license acquire
interface fc1/42
port-license acquire
interface fc1/43
port-license acquire
interface fc1/44
port-license acquire
interface fc1/45
port-license acquire
interface fc1/46
port-license acquire
interface fc1/47
port-license acquire
interface fc1/48
 port-license acquire
interface mgmt0
  ip address 10.29.164.64 255.255.255.0
ip default-gateway 10.29.164.1
MDS-A#
```

Cisco MDS 9148S-B Configuration

```
MDS-B# show run[15D[J MDS-B# show run[Jning-config  !Command: show running-config
!Time: Tue Jun 28 22:50:59 2016
version 6.2(9a)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
 description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $1$dcmFCg/p$0ZC5U6uhI65oOePpHfAzn0  role network-admin
no password strength-check
ip domain-lookup
ip host MDS-B  10.29.164.128
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xc9e1af5dbb0bbac72253a1bef037bbbe
 priv 0xc9e1af5dbb0bbac72253a1bef037bbbe localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1164
snmp-server host 10.29.164.130 traps version 2c public udp-port 1164
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vsan database
vsan 30
device-alias database
device-alias commit
```

```
fcdomain fcid database

  vsan 30 wwn 52:4a:93:72:0d:21:6b:01 fcid 0x9d0000 dynamic
  vsan 30 wwn 52:4a:93:72:0d:21:6b:10 fcid 0x9d0100 dynamic
  vsan 30 wwn 20:4f:54:7f:ee:45:2a:40 fcid 0x9d0200 dynamic
  vsan 30 wwn 20:50:54:7f:ee:45:2a:40 fcid 0x9d0300 dynamic
  vsan 30 wwn 20:4f:54:7f:ee:45:29:80 fcid 0x9d0400 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:1c fcid 0x9d0411 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:16 fcid 0x9d0402 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:14 fcid 0x9d0506 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:12 fcid 0x9d0503 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:08 fcid 0x9d0404 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:06 fcid 0x9d0505 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:04 fcid 0x9d0405 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:02 fcid 0x9d0508 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:10 fcid 0x9d0406 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:0a fcid 0x9d0507 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:0c fcid 0x9d0504 dynamic
  vsan 30 wwn 20:50:54:7f:ee:45:29:80 fcid 0x9d0500 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:6f fcid 0x9d0413 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:20 fcid 0x9d050e dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:1f fcid 0x9d0401 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:1a fcid 0x9d0403 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:18 fcid 0x9d0408 dynamic
  vsan 30 wwn 20:00:00:25:b5:b1:1b:01 fcid 0x9d050c dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:01 fcid 0x9d041a dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:03 fcid 0x9d040f dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:05 fcid 0x9d0515 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:07 fcid 0x9d050b dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:0d fcid 0x9d050f dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:0f fcid 0x9d041b dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:0b fcid 0x9d0509 dynamic
```

```
  vsan 30 wwn 20:00:00:25:b5:1f:1a:1d fcid 0x9d0512 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:1f fcid 0x9d0518 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:21 fcid 0x9d0414 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:23 fcid 0x9d0417 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:25 fcid 0x9d040b dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:17 fcid 0x9d040c dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:15 fcid 0x9d050d dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:27 fcid 0x9d050a dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:13 fcid 0x9d0407 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:11 fcid 0x9d0412 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:31 fcid 0x9d040d dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:2d fcid 0x9d0410 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:33 fcid 0x9d0502 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:35 fcid 0x9d040a dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:37 fcid 0x9d0501 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:39 fcid 0x9d0409 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:2b fcid 0x9d0510 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:1b fcid 0x9d0415 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:19 fcid 0x9d040e dynamic
  vsan 30 wwn 52:4a:93:72:0d:21:6b:02 fcid 0x9d0600 dynamic
   vsan 30 wwn 52:4a:93:72:0d:21:6b:12 fcid 0x9d0700 dynamic
  vsan 30 wwn 52:4a:93:72:0d:21:6b:13 fcid 0x9d0800 dynamic
  vsan 30 wwn 20:00:00:25:b5:1f:1a:29 fcid 0x9d0513 dynamic
interface port-channel1
  channel mode active
  switchport rate-mode dedicated
vsan database
  vsan 30 interface fc1/1
  vsan 30 interface fc1/2
  vsan 30 interface fc1/3
  vsan 30 interface fc1/4
  vsan 30 interface fc1/7
```

```
  vsan 30 interface fc1/8
switchname MDS-B
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
```

```
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
!Active Zone Database Section for vsan 30
zone name Infra-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:20
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
    member pwwn 20:00:00:25:b5:b1:1b:6f
```

```
    member pwwn 52:4a:93:72:0d:21:6b:01
    member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:1f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv3-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:03
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv4-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:05
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv5-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:07
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv6-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:09
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv7-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0b
```

```
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv8-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv9-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv10-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:11
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv11-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:13
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv12-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:15
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv13-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:17
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv16-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv17-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1f
```

```
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv18-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:21
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv19-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:23
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv20-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:25
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv21-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:27
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv24-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2d
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv25-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv26-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:31
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv27-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:33
```

```
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv28-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:35
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv29-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:37
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv30-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:39
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv15-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:1b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv23-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:2b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-serv22-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:29
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zoneset name pure-vdi-fab-B vsan 30
member Infra-serv1-fc1
member Infra-serv2-fc1
member Infra-serv3-fc1
member pure-vdi-serv1-fc1
member pure-vdi-serv2-fc1
```

```
member pure-vdi-serv3-fc1
member pure-vdi-serv4-fc1
member pure-vdi-serv5-fc1
member pure-vdi-serv6-fc1
member pure-vdi-serv7-fc1
member pure-vdi-serv8-fc1
member pure-vdi-serv9-fc1
member pure-vdi-serv10-fc1
member pure-vdi-serv11-fc1
member pure-vdi-serv12-fc1
member pure-vdi-serv13-fc1
member pure-vdi-serv16-fc1
member pure-vdi-serv17-fc1
member pure-vdi-serv18-fc1
member pure-vdi-serv19-fc1
member pure-vdi-serv20-fc1
member pure-vdi-serv21-fc1
member pure-vdi-serv24-fc1
member pure-vdi-serv25-fc1
member pure-vdi-serv26-fc1
member pure-vdi-serv27-fc1
member pure-vdi-serv28-fc1
member pure-vdi-serv29-fc1
member pure-vdi-serv30-fc1
member pure-vdi-serv15-fc1
member pure-vdi-serv23-fc1
member pure-vdi-serv22-fc1
zoneset activate name pure-vdi-fab-B vsan 30
do clear zone database vsan 30
!Full Zone Database Section for vsan 30
zone name Infra-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:20
```

```
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name Infra-serv2-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:6f
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name vdi-serv1-fc1 vsan 30
member pwwn 20:00:00:25:b5:b1:1b:01
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
zone name pure-vdi-fab-B vsan 30
zone name pure-vdi-serv7-fc1 vsan 30
member pwwn 20:00:00:25:b5:1f:1a:0b
member pwwn 52:4a:93:72:0d:21:6b:01
member pwwn 52:4a:93:72:0d:21:6b:10
interface fc1/1
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/2
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/3
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/4
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/5
```

```
port-license acquire
interface fc1/6
port-license acquire
interface fc1/7
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/8
switchport trunk mode off
port-license acquire
no shutdown
interface fc1/9
port-license acquir no shutdown
interface fc1/10
port-license acquire
no shutdown
interface fc1/11
port-license acquire
channel-group 1 force
no shutdown
interface fc1/12
port-license acquire
channel-group 1 force
no shutdown
interface fc1/13
port-license acquire
interface fc1/14
port-license acquire
no shutdown
interface fc1/15
port-license acquire
no shutdown
```

```
interface fc1/16
port-license acquire
no shutdown
interface fc1/17
port-license acquire
interface fc1/18
port-license acquire
interface fc1/19
port-license acquire
interface fc1/20
port-license acquire
interface fc1/21
port-license acquire
interface fc1/22
port-license acquire
interface fc1/23
port-license acquire
interface fc1/24
port-license acquire
interface fc1/25
port-license acquire
interface fc1/26
port-license acquire
interface fc1/27
port-license acquire
interface fc1/28
port-license acquire
interface fc1/29
port-license acquire
interface fc1/30
port-license acquire
interface fc1/31
```

```
port-license acquire
interface fc1/32
port-license acquire
interface fc1/33
port-license acquire
interface fc1/34
port-license acquire
interface fc1/35
port-license acquire
interface fc1/36
port-license acquire
interface fc1/37
port-license acquire
interface fc1/38
port-license acquire
interface fc1/39
port-license acquire
interface fc1/40
port-license acquire
interface fc1/41
port-license acquire
interface fc1/42
port-license acquire
interface fc1/43
port-license acquire
interface fc1/44
port-license acquire
interface fc1/45
port-license acquire
interface fc1/46
port-license acquire
interface fc1/47
```

```
port-license acquire
interface fc1/48
port-license acquire
interface mgmt0
ip address 10.29.164.128 255.255.255.0
ip default-gateway 10.29.164.1
MDS-B#
```

Appendix B – Pure Storage Configuration and Scripts

Configuring the ESXi hosts in the environment for optimal use with the FlashArray//m50 was accomplished by running a single PowerShell script once after all ESXi hosts were built and datastores were provisioned. Additional datastores created later will have these best practices applied to them automatically. It is recommended to reboot each ESXi host after applying this script.

Future updates to this script as well as other useful PowerShell scripts built for use with Pure Storage can be found here:

https://github.com/codyhosterman/powercli

Below is the ESXi Best Practice PowerShell script used in this Cisco Validated Design:

```
#Enter the following required parameters. Log folder directory is just an example, change as needed.

#Put all entries inside the quotes:

#**********************************

$vcenter = ""

$vcuser = ""

$vcpass = ""

$logfolder = "C:\folder\folder\etc\"

#**********************************



<#

Optional parameters. Keep these values at default unless necessary and understood

For a different IO Operations limit beside the Pure Storage recommended value of 1, change $iopsvalue
to another integer value 1-1000.

To skip changing host-wide settings for XCOPY Transfer Size and In-Guest UNMAP change $hostwideset-
tings to $false

#>

$iopsvalue = 1

$hostwidesettings = $true



<#

*******Disclaimer:******************************************************

This scripts are offered "as is" with no warranty.  While this

scripts is tested and working in my environment, it is recommended that you test

this script in a test lab before using in a production environment. Everyone can

use the scripts/commands provided here without any written permission but I

will not be liable for any damage or loss to the system.

************************************************************************
```

```
This script will:

-Check for a SATP rule for Pure Storage FlashArrays

-Create a SATP rule for Round Robin and IO Operations Limit of 1 only for FlashArrays

-Remove any incorrectly configured Pure Storage FlashArray rules

-Configure any existing devices properly (Pure Storage FlashArray devices only)

-Set VAAI XCOPY transfer size to 16MB

-Enable EnableBlockDelete on ESXi 6 hosts only

All change operations are logged to a file.

This can be run directly from PowerCLI or from a standard PowerShell prompt. PowerCLI must be in-
stalled on the local host regardless.

Supports:

-FlashArray 400 Series and //m

-vCenter 5.0 and later

-PowerCLI 6.3 R1 or later required

For info, refer to www.codyhosterman.com

#>


#Create log folder if non-existent

If (!(Test-Path -Path $logfolder)) { New-Item -ItemType Directory -Path $logfolder }

$logfile = $logfolder + (Get-Date -Format o |ForEach-Object {$_ -Replace ":", "."}) + "setbestprac-
tices.txt"

write-host "Checking and setting Pure Storage FlashArray Best Practices for VMware on the ESXi hosts
in this vCenter. No further information is printed to the screen."

write-host "Script log information can be found at $logfile"


add-content $logfile '             _____________________________'

add-content $logfile '            /++++++++++++++++++++++++++++\'

add-content $logfile '           /++++++++++++++++++++++++++++++\'

add-content $logfile '          /++++++++++++++++++++++++++++++++\'

add-content $logfile '         /++++++++++++++++++++++++++++++++++\'

add-content $logfile '        /++++++++++++++++++++++++++++++++++++\'

add-content $logfile '       /++++++++++++/----------\++++++++++++\'

add-content $logfile '      /++++++++++++/            \++++++++++++\'

add-content $logfile '     /++++++++++++/              \++++++++++++\'

add-content $logfile '    /++++++++++++/                \++++++++++++\'

add-content $logfile '   /++++++++++++/                  \++++++++++++\'
```

```
add-content $logfile '   \++++++++++++\                  /++++++++++++/'

add-content $logfile '    \++++++++++++\                /++++++++++++/'

add-content $logfile '     \++++++++++++\              /++++++++++++/'

add-content $logfile '      \++++++++++++\            /++++++++++++/'

add-content $logfile '       \++++++++++++\          /++++++++++++/'

add-content $logfile '        \++++++++++++\'

add-content $logfile '         \++++++++++++\'

add-content $logfile '          \++++++++++++\'

add-content $logfile '           \++++++++++++\'

add-content $logfile '            \------------\'

add-content $logfile 'Pure Storage  FlashArray VMware ESXi Best Practices Script v3.1'

add-content $logfile '-------------------------------------------------------------------------------
----------------------'



if ( !(Get-Module -Name VMware.VimAutomation.Core -ErrorAction SilentlyContinue) ) {

. "C:\Program Files (x86)\VMware\Infrastructure\vSphere PowerCLI\Scripts\Initialize-PowerCLIEnviron-
ment.ps1" |out-null

}

set-powercliconfiguration -invalidcertificateaction "ignore" -confirm:$false |out-null



if ((Get-PowerCLIVersion).build -lt 3737840)

{

        write-host "This version of PowerCLI is too old, version 6.3 Release 1 or later is required
        (Build 3737840)" -BackgroundColor Red

        write-host "Found the following build number:"

        write-host (Get-PowerCLIVersion).build

        write-host "Terminating Script" -BackgroundColor Red

        add-content $logfile "This version of PowerCLI is too old, version 6.3 Release 1 or later is
        required (Build 3737840)"

        add-content $logfile "Found the following build number:"

        add-content $logfile (Get-PowerCLIVersion).build

        add-content $logfile "Terminating Script"

        return

}



try

{
```

```
        connect-viserver -Server $vcenter -username $vcuser -password $vcpass -ErrorAction Stop |out-
        null

}

catch

{

        write-host "Failed to connect to vCenter" -BackgroundColor Red

        write-host $Error

        write-host "Terminating Script" -BackgroundColor Red

        add-content $logfile "Failed to connect to vCenter"

        add-content $logfile $Error

        add-content $logfile "Terminating Script"

        return

}

add-content $logfile ('Connected to vCenter at ' + $vcenter)

add-content $logfile '-------------------------------------------------------------------------------
----------------------'



$hosts= get-vmhost



add-content $logfile "Iterating through all ESXi hosts..."



#Iterating through each host in the vCenter

foreach ($esx in $hosts)

{

        $esxcli=get-esxcli -VMHost $esx -v2

        add-content $logfile "-------------------------------------------------------------------------
        ------------------------"

        add-content $logfile "-------------------------------------------------------------------------
        ------------------------"

        add-content $logfile "Working on the following ESXi host:"

        add-content $logfile $esx.NetworkInfo.hostname

        add-content $logfile "------------------------------------------------"

        if ($hostwidesettings -eq $true)

        {

        add-content $logfile "Checking host-wide setting for XCOPY and In-Guest UNMAP"

        $xfersize = $esx | Get-AdvancedSetting -Name DataMover.MaxHWTransferSize

        if ($xfersize.value -ne 16384)
```

```
{
        add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this host is incorrect:"
        add-content $logfile $xfersize.value
        add-content $logfile "This should be set to 16386 (16 MB). Changing to 16384..."
        $xfersize |Set-AdvancedSetting -Value 16384 -Confirm:$false |out-null
        add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this host is now 16 MB"
}
else
{
        add-content $logfile "The VAAI XCOPY MaxHWTransferSize for this host is correct at 16
        MB and will not be altered."
}
if ($esx.Version -like "6.0.*")
{
        $enableblockdelete = ($esx | Get-AdvancedSetting -Name VMFS3.EnableBlockDelete).Value
        if ($enableblockdelete.Value -eq 0)
        {
        add-content $logfile "EnableBlockDelete is currently disabled. Enabling..."
        $enableblockdelete |Set-AdvancedSetting -Value 1 -Confirm:$false |out-null
        add-content $logfile "EnableBlockDelete has been set to enabled."
        }
        else
        {
        add-content $logfile "EnableBlockDelete for this host is correctly enabled and will
        not be altered."
        }
}
else
{
        add-content $logfile "The current host is not version 6.0. Skipping EnableBlockDelete
        check."
}
}
else
{
add-content $logfile "Not checking host wide settings for XCOPY and In-Guest UNMAP due to in-
script override"
```

```
}

add-content $logfile "----------------------------------------------"

$rules = $esxcli.storage.nmp.satp.rule.list.invoke() |where-object {$_.Vendor -eq "PURE"}

$correctrule = 0

$iopsoption = "iops=" + $iopsvalue

if ($rules.Count -ge 1)

{

add-content $logfile "Found the following existing Pure Storage SATP rules"

$rules | out-string | add-content $logfile

add-content $logfile "----------------------------------------------"

foreach ($rule in $rules)

{

        add-content $logfile "----------------------------------------------"

        add-content $logfile "Checking the following existing rule:"

        $rule | out-string | add-content $logfile

        $issuecount = 0

        if ($rule.DefaultPSP -eq "VMW_PSP_RR")

        {

        add-content $logfile "The existing Pure Storage FlashArray rule is configured with the
        correct Path Selection Policy:"

        add-content $logfile $rule.DefaultPSP

        }

        else

        {

        add-content $logfile "The existing Pure Storage FlashArray rule is NOT configured with
        the correct Path Selection Policy:"

        add-content $logfile $rule.DefaultPSP

        add-content $logfile "The rule should be configured to Round Robin (VMW_PSP_RR)"

        $issuecount = 1

        }

        if ($rule.PSPOptions -eq $iopsoption)

        {

        add-content $logfile "The existing Pure Storage FlashArray rule is configured with the
        correct IO Operations Limit:"

        add-content $logfile $rule.PSPOptions

        }
```

```
else

{

add-content $logfile "The existing Pure Storage FlashArray rule is NOT configured with
the correct IO Operations Limit:"

add-content $logfile $rule.PSPOptions

add-content $logfile "The rule should be configured to an IO Operations Limit of $iop-
svalue"

$issuecount = $issuecount + 1

}

if ($rule.Model -eq "FlashArray")

{

add-content $logfile "The existing Pure Storage FlashArray rule is configured with the
correct model:"

add-content $logfile $rule.Model

}

else

{

add-content $logfile "The existing Pure Storage FlashArray rule is NOT configured with
the correct model:"

add-content $logfile $rule.Model

add-content $logfile "The rule should be configured with the model of FlashArray"

$issuecount = $issuecount + 1

}

if ($issuecount -ge 1)

{

$satpArgs = $esxcli.storage.nmp.satp.rule.remove.createArgs()

$satpArgs.model = $rule.Model

$satpArgs.vendor = "PURE"

$satpArgs.satp = $rule.Name

$satpArgs.psp = $rule.DefaultPSP

$satpArgs.pspoption = $rule.PSPOptions

add-content $logfile "This rule is incorrect, deleting..."

$esxcli.storage.nmp.satp.rule.remove.invoke($satpArgs)

add-content $logfile "*****NOTE: Deleted the rule.*****"

add-content $logfile "------------------------------------------------"

}

else
```

```
        {
        add-content $logfile "This rule is correct"
        add-content $logfile "----------------------------------------------"
        $correctrule = 1
        }
}
}
if ($correctrule -eq 0)
{
add-content $logfile "No correct SATP rule for the Pure Storage FlashArray is found. Creating
a new rule to set Round Robin and an IO Operations Limit of $iopsvalue"
$satpArgs = $esxcli.storage.nmp.satp.rule.remove.createArgs()
$satpArgs.description = "Pure Storage FlashArray SATP Rule"
$satpArgs.model = "FlashArray"
$satpArgs.vendor = "PURE"
$satpArgs.satp = "VMW_SATP_ALUA"
$satpArgs.psp = "VMW_PSP_RR"
$satpArgs.pspoption = $iopsoption
$result = $esxcli.storage.nmp.satp.rule.add.invoke($satpArgs)
if ($result -eq $true)
{
        add-content $logfile "New rule created:"
        $newrule = $esxcli.storage.nmp.satp.rule.list.invoke() |where-object {$_.Vendor -eq
        "PURE"}
        $newrule | out-string | add-content $logfile
}
else
{
        add-content $logfile "ERROR: The rule failed to create. Manual intervention might be
        required."
}
}
else
{
add-content $logfile "A correct SATP rule for the FlashArray exists. No need to create a new
one on this host."
}
```

```
$devices = $esx |Get-ScsiLun -CanonicalName "naa.624a9370*"

add-content $logfile "----------------------------------------------"

if ($devices.count -ge 1)

{

add-content $logfile "Looking for existing Pure Storage volumes on this host"

add-content $logfile "Found the following number of existing Pure Storage volumes on this
host."

add-content $logfile $devices.count

add-content $logfile "Checking and fixing their multipathing configuration now."

add-content $logfile "----------------------------------------------"

foreach ($device in $devices)

{

        add-content $logfile "Found and examining the following FlashArray device:"

        add-content $logfile $device.CanonicalName

        if ($device.MultipathPolicy -ne "RoundRobin")

        {

        add-content $logfile "This device does not have the correct Path Selection Policy, it
        is set to:"

        add-content $logfile $device.MultipathPolicy

        add-content $logfile "Changing to Round Robin."

        Get-VMhost $esx |Get-ScsiLun $device |Set-ScsiLun -MultipathPolicy RoundRobin |out-
        null

        }

        else

        {

        add-content $logfile "This device's Path Selection Policy is correctly set to Round
        Robin. No need to change."

        }

        $deviceargs = $esxcli.storage.nmp.psp.roundrobin.deviceconfig.get.createargs()

        $deviceargs.device = $device.CanonicalName

        $deviceconfig = $esxcli.storage.nmp.psp.roundrobin.deviceconfig.get.invoke($de-
        viceargs)

        $nmpargs =  $esxcli.storage.nmp.psp.roundrobin.deviceconfig.set.createargs()

        $nmpargs.iops = $iopsvalue

        $nmpargs.type = "iops"

        if ($deviceconfig.IOOperationLimit -ne $iopsvalue)

        {
```

```
                add-content $logfile "The current IO Operation limit for this device is:"

                add-content $logfile $deviceconfig.IOOperationLimit

                add-content $logfile "This device's IO Operation Limit is unset or is not set to the
                value of $iopsvalue. Changing..."

                $nmpargs.device = $device.CanonicalName

                $esxcli.storage.nmp.psp.roundrobin.deviceconfig.set.invoke($nmpargs) |out-null

                }

                else

                {

                add-content $logfile "This device's IO Operation Limit matches the value of $iop-
                svalue. No need to change."

                }

                add-content $logfile "--------------------"

        }

        }

        else

        {

        add-content $logfile "No existing Pure Storage volumes found on this host."

        }

}

disconnect-viserver -Server $vcenter -confirm:$false

add-content $logfile "Disconnected vCenter connection"
```

Appendix C - Pure Storage FlashArray//m50 Expanded Test Results

This section highlights and provide analysis of the Pure Storage FlashArray//m50 performance results for each of the cluster test cases identified in the Cisco Validated Design.

From a storage perspective, it is critical to maintain a latency of near to or less than a millisecond in order to guarantee a good end-user experience. As we will see, Pure Storage delivers that level of performance despite driving a substantial amount of IOPs and bandwidth for the thousands of desktops hosted on the single FlashArray//m50.

The following charts were compiled from extracting front-end array telemetry data from the storage logs and are equivalent to values shown in the Pure GUI. Results were plotted in this format to highlight individual storage performance metrics of interest during each simulation as well as clearly show the various phases of each simulation. Please note that across the top of each graph we have identified and broken up the Login VSI simulation into the three separate phases of the simulation run. The first phase (green arrows and text box) is the 2880 second Login VSI simulation phase when all sessions are ramping up and logging in. Next, the all sessions in the simulation steady-states for 600 seconds which is denoted by the yellow arrows and text box, and finally the black arrow to the right shows the end of the simulation when users begin logging out of the environment. For brevity, we generally did not show the entire logout operation as array activity is minimal during that time and the Login VSI simulation had completed.

Pure Storage FlashArray//m50 Test Results for 1450 RDSH Sessions

Using Login VSI as the workload generator in Benchmark mode with the Knowledge Worker workload, our first highlighted cluster test shows that the FlashArray//m50 can easily handle this workload with exceptional end-user experience confirmed from Login VSI.

The first chart below shows the latency of the Pure Storage FlashArray//m50 during the 1450 Horizon RDSH sessions running on top of 72 Windows 2012 R2 servers. There were three separate 1450 Horizon RDSH simulation runs completed in total, all with very similar results. As we can see in the below chart, we maintained latency of less than or close to one millisecond for both read and write operations throughout this entire run, which resulted in a confirmed outstanding end-user experience for the simulated RDSH users.

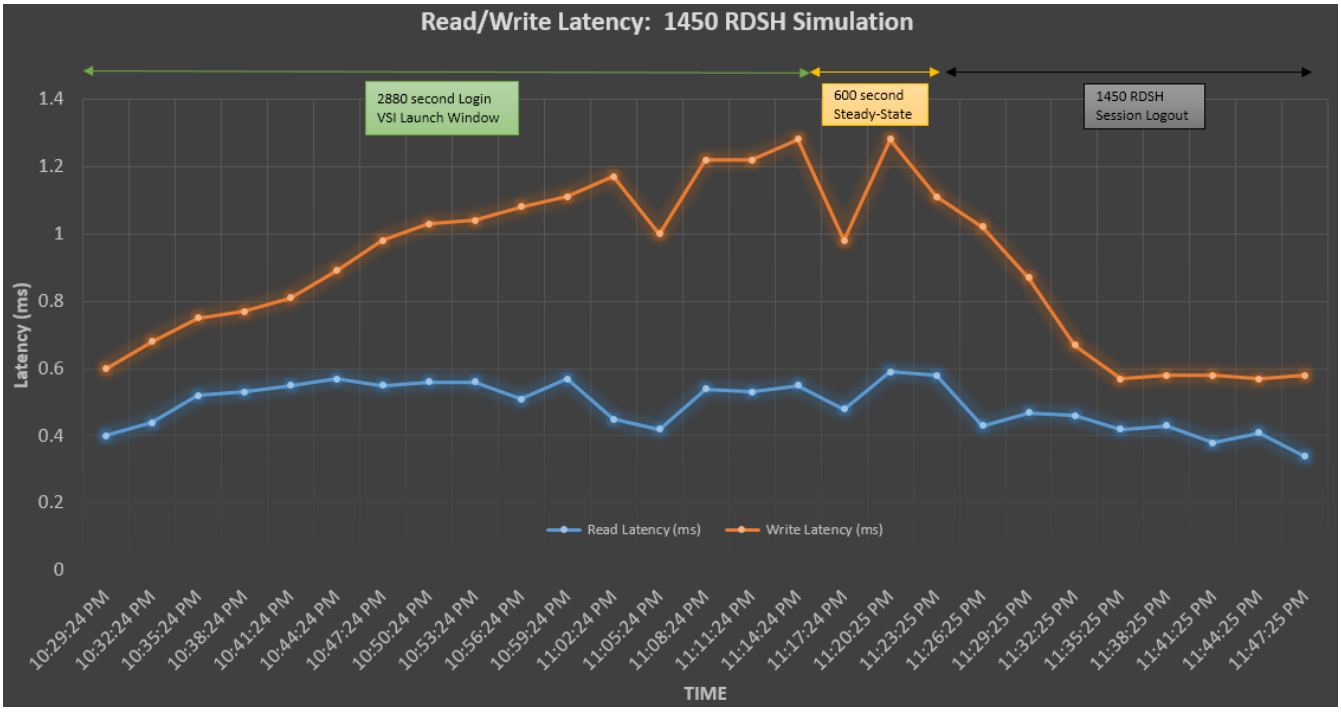

The next chart shows our read and write IOPs during the selected 1450 Horizon RDSH simulation.

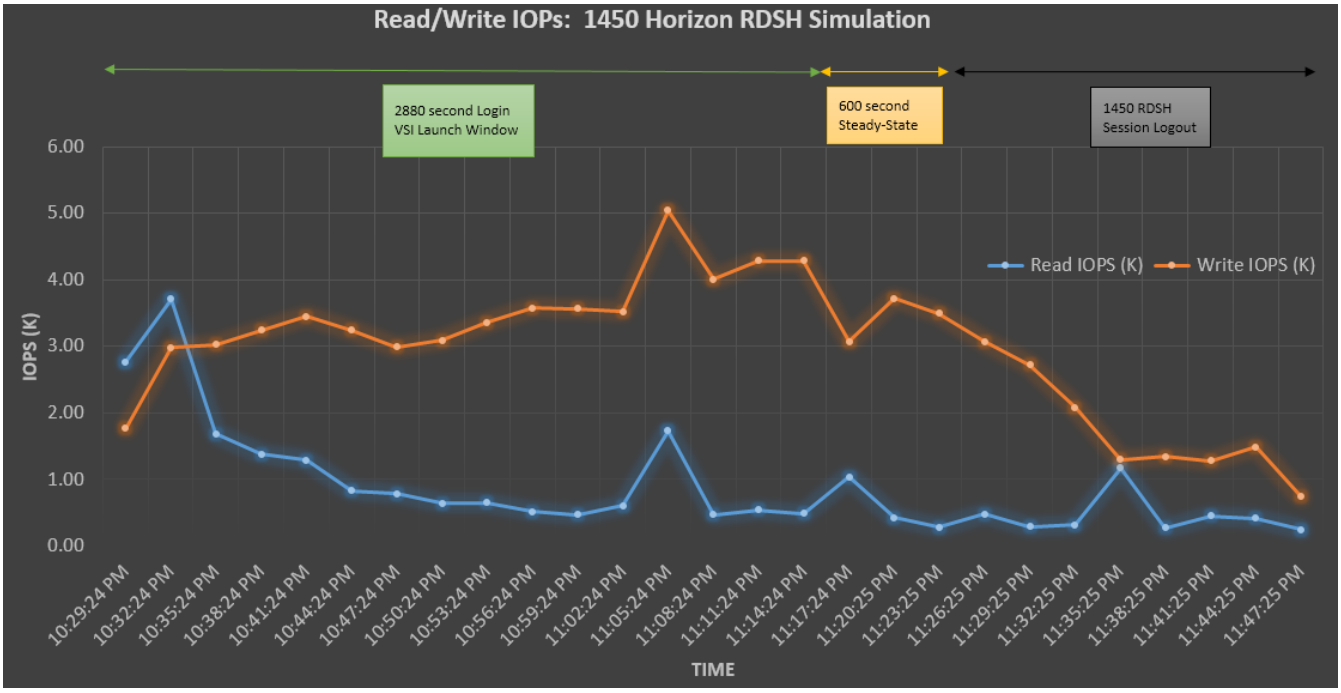

Finally, we can see bandwidth against the array ramping up, hitting steady-state and then dropping as sessions logout in parallel with the 1450 Horizon RDSH sessions in the last chart from this simulation below.

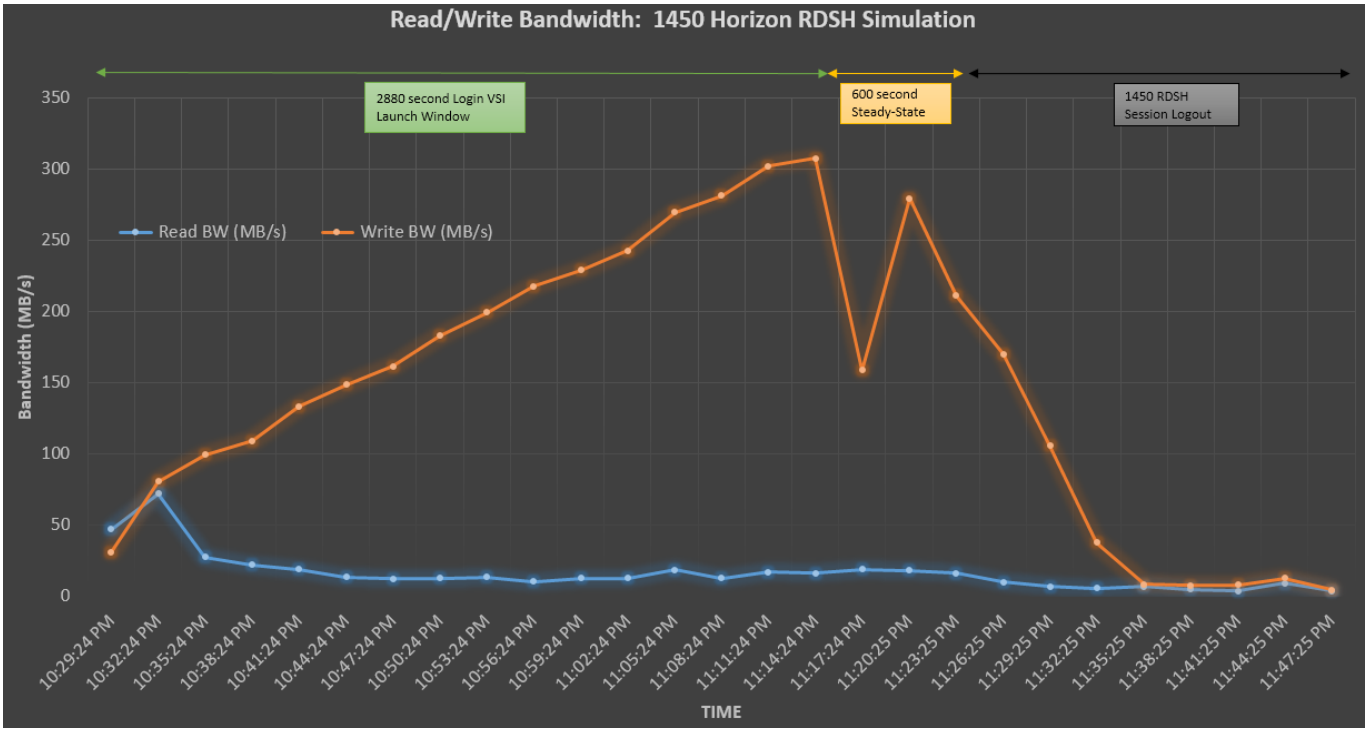

Below we see the utilization of the Horizon RDSH datastore after the Login VSI simulation. Worth noting is that the data reduction numbers are artificially low since Login VSI writes primarily non-reducible data to the desktop.

NAME	# HOSTS	PROVISIONED	VOLUMES	SNAPSHOTS	REDUCTION
RDSH-DS	9	10 TB	117.55 GB	24.46 GB	8.3 to 1

Pure Storage FlashArray//m50 Test Results for 3550 Linked-Clone Windows 7 Sessions

The next cluster-level simulation was to run 3550 Horizon linked-clone Windows 7 desktops against the same FlashArray//m50. All Login VSI parameters were kept consistent with the previous RDSH test with the only change being to use 3550 linked-clone desktops. As can be seen in the below storage metrics, the Pure Storage FlashArray//m50 was clearly able to handle this workload and continued to provide sub-millisecond latency for an impressive Login VSI result.

Firstly, latency was consistently sub-millisecond throughout all phases of the Login VSI simulation despite driving tens of thousands of IOPs and hundreds of megabytes of read and write bandwidth.

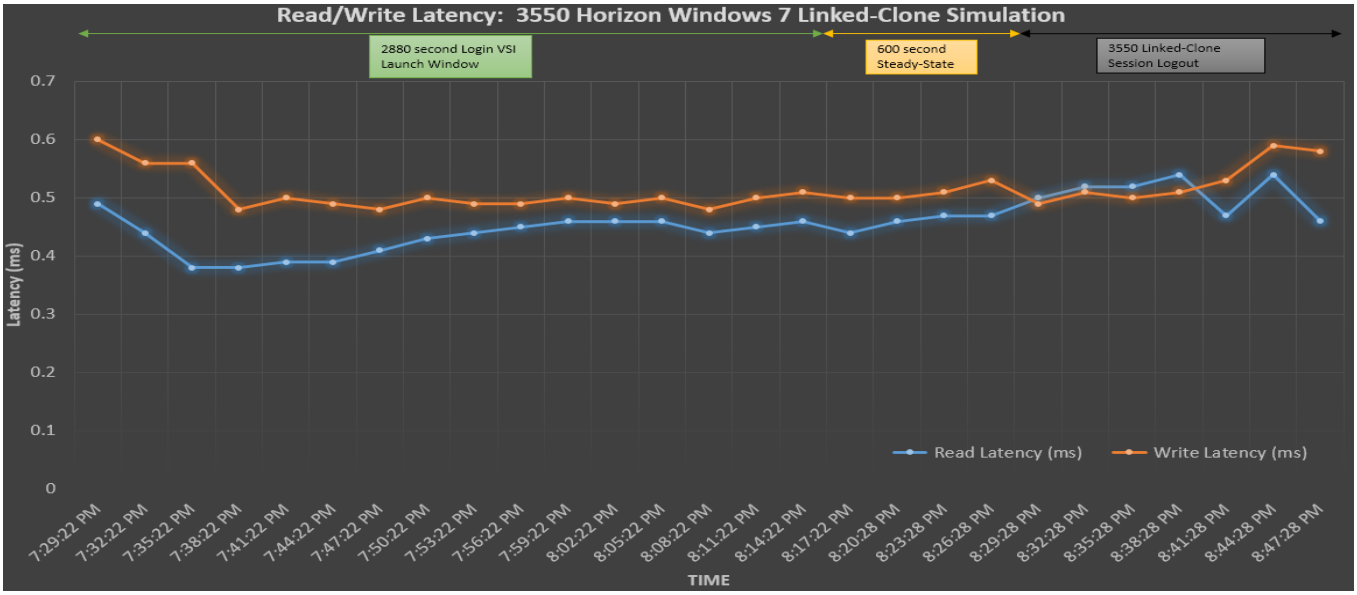

This next chart shows IOPS ramping up as linked-clone sessions are added to the simulation. The brief dip occurs when all sessions login and begin to steady-state and then increases again as linked-clone desktops log out and are refreshed back to the template image.

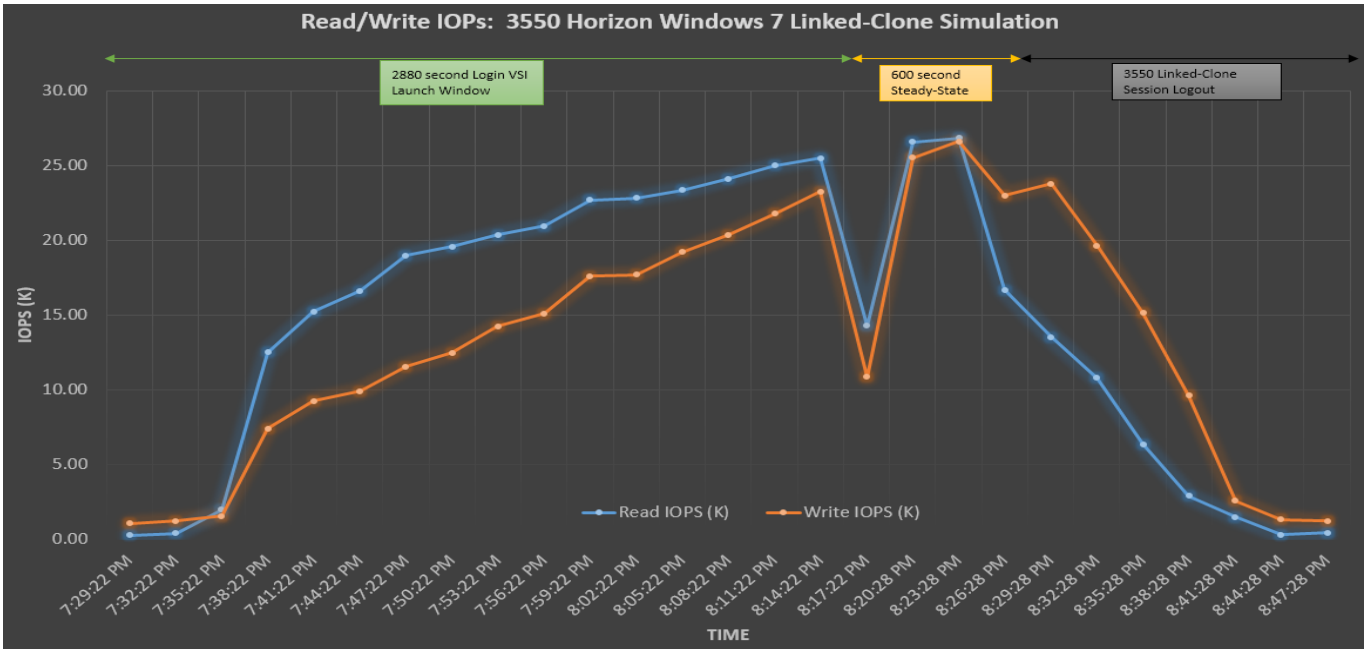

The bandwidth against the array followed a similar performance curve throughout the simulation with it ramping up initially, temporarily dropping due to simulation steady-state and then write bandwidth increasing dramatically as desktops are refreshed at logout.

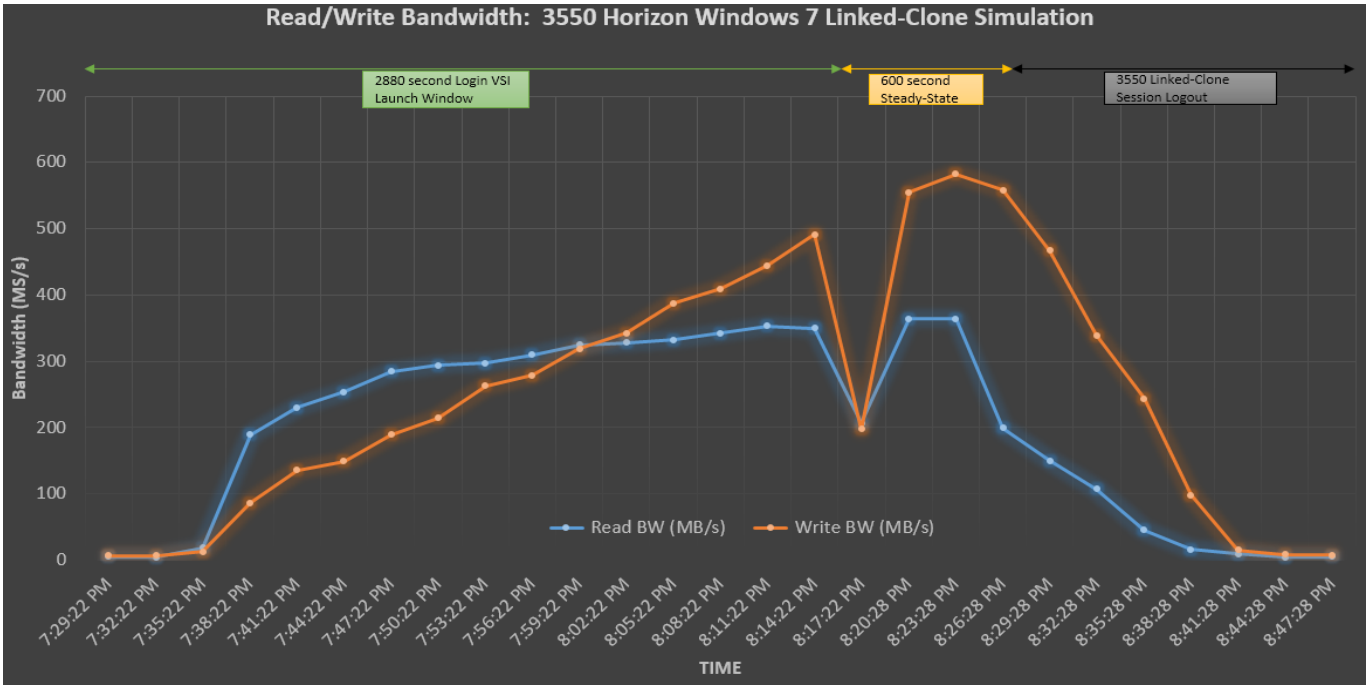

Pure Storage FlashArray//m50 Test Results for Full Scale, Mixed Workload Scalability

The next simulation shows the results of combining the earlier 1450 Horizon RDSH sessions with the 3550 Horizon linked-clone sessions for a 5000 user VMware Horizon simulation all on the same Pure Storage FlashArray//m50 array. Yet again we see performant and consistent results that prove an outstanding user experience – even at this large of a user scale.

The below two screenshots shows the Pure Storage GUI with the cursor providing more detailed metrics at both the start of the simulation and at the end. Despite driving nearly 1GB/s in bandwidth we maintain the responsiveness of low latency throughout the entire test.

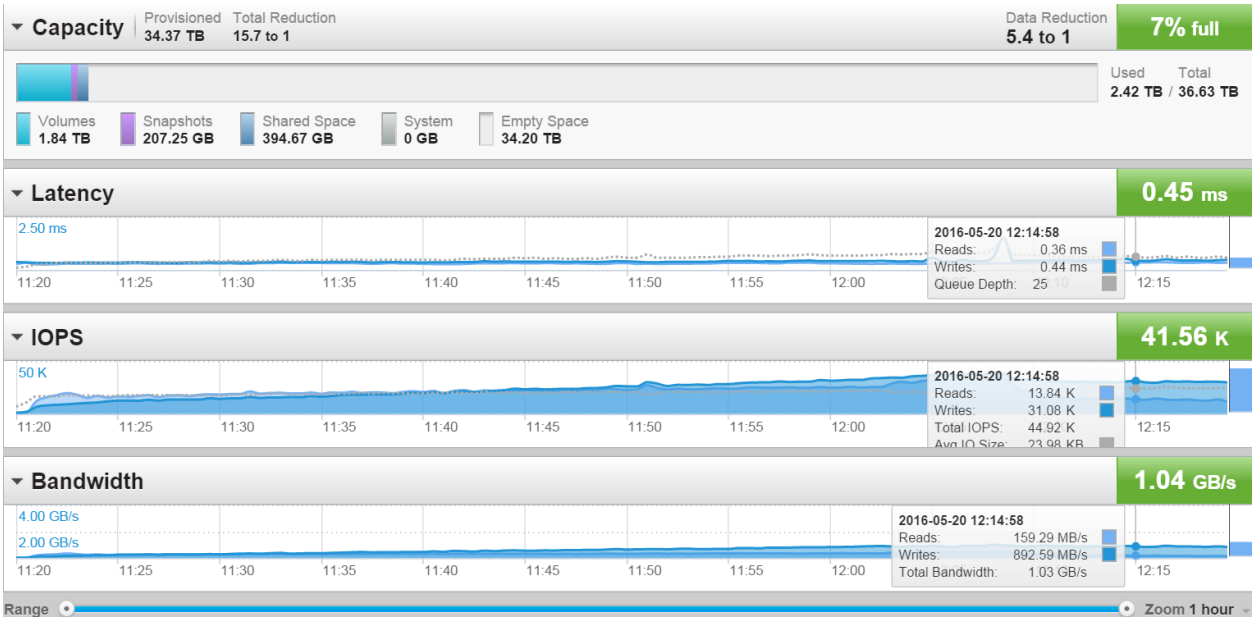

Latency results shown in the above GUI are exactly mirrored from logs pulled from the //m50 array and plotted in similar fashion to earlier tests. We can see yet again consistent low latency throughout the test run that show clear evidence of great VDI experience for the end-user.

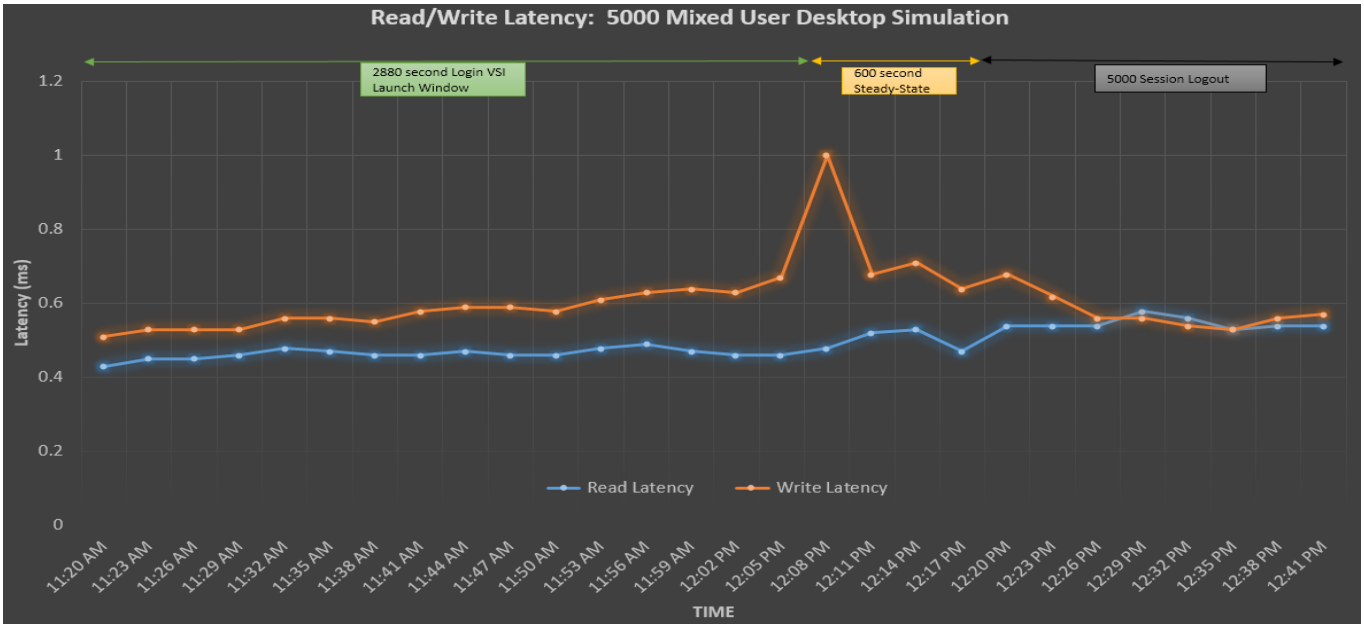

Read and Write IOPS can be seen dramatically increasing as linked-clone and RSDH sessions are added during the 2880 second simulation.

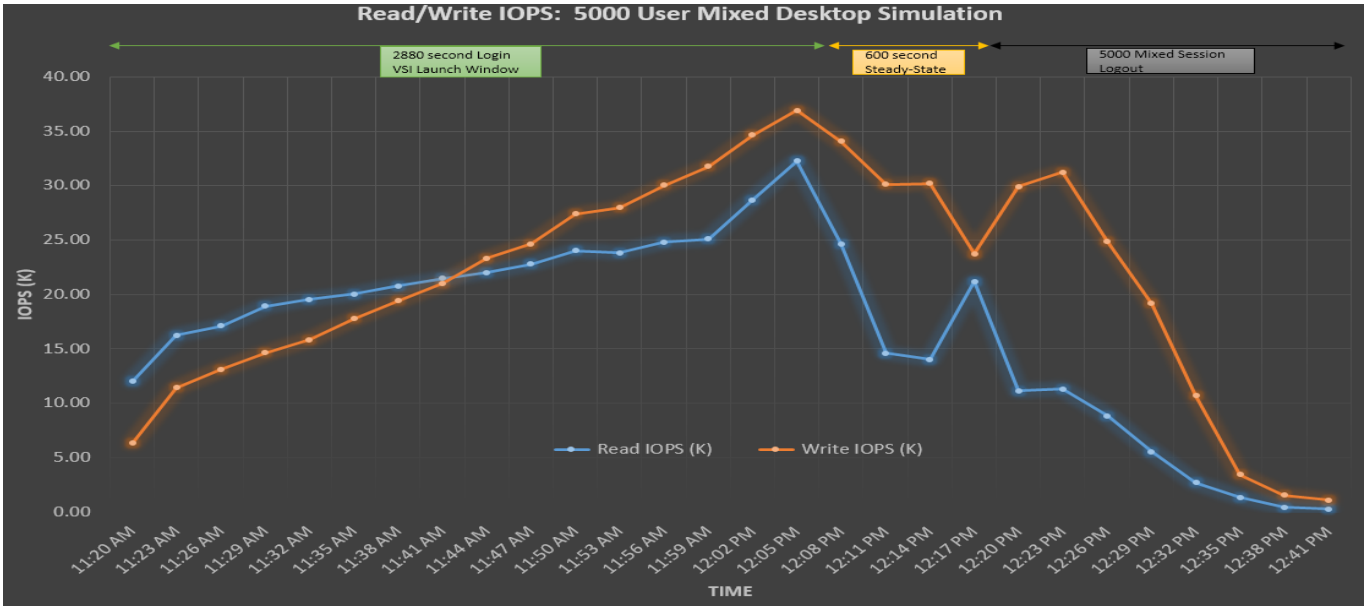

Similarly, bandwidth increases during the 2880 ramp-up timeframe of the simulation, then falls briefly as the test hits steady-state and then begins to drop as sessions are logged out.