Cisco Solution for EMC VSPEX End User Computing

For 2000 VMware Horizon View 5.2 Users

Last Updated: November 21, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Ramesh Guduru, Support Engineer, Cisco Systems

Ramesh Guduru is a Virtualization System Engineer at Cisco with SSVPG. Ramesh has over 7 years of experience with VMware view thin client administration, configuration and optimization of virtual desktop environment. Ramesh is skilled on core VMware applications in the virtual environment focusing in system design and implementation of virtualization components. Ramesh holds certification in virtualization, network and Microsoft.

Hardik Patel, Support Engineer, Cisco Systems

Hardik Patel is a Virtualization System Engineer at Cisco with SSVPG. Hardik has over 9 years of experience with server virtualization and core application in the virtual environment with area of focus in design and implementation of systems and virtualization, manage and administration, Cisco Unified Computing System, storage and network configurations. Hardik holds Masters degree in Computer Science with various career oriented certification in virtualization, network and Microsoft.

Acknowledgments

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to thank:

- Mike Brennan, Manager, Technical Marketing, Cisco Systems, Inc.
- Jason Ventresco, Solutions Engineer, Strategic Solutions Engineering, EMC
- Jack McMichael, EUC Technical Enablement, VMware

About Cisco Validated Design (CVD) Program

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.

Cisco Solution for EMC VSPEX End User Computing

Overview

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information and toward shared infrastructures, to virtualized environments, and eventually to the cloud to increase agility and reduce costs.

This Cisco Solution for EMC VSPEX End User Computing reports the results of a study evaluating the scalability of VMware Horizon View 5.2 environment on Cisco UCS B-Series B200 M3 Blade Servers running VMware ESXi 5.1 hypervisor software connected to an EMC VNX 5500 Storage Array. We utilize second and third generation Cisco Unified Computing System hardware and software. We provide best practice recommendations and sizing guidelines for large scale customer deployments of VMware Horizon View 5.2 on the Cisco Unified Computing System™.

Audience

This document describes the architecture and deployment procedures of an infrastructure comprised of Cisco, EMC, and VMware hypervisor and desktop virtualization products. The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the solution described in this document.

Solution Component Benefits

Each of the components of the overall solution materially contributes to the value of functional design contained in this document.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Benefits of Cisco Unified Computing System

Cisco Unified Computing System™ is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud computing environments.

Benefits of the Unified Computing System include:

Architectural flexibility

- Third generation Cisco UCS B-Series blade servers for infrastructure and virtual workload hosting
- Third generation Cisco UCS C-Series rack-mount servers for infrastructure and virtual workload Hosting
- Cisco UCS 6200 Series second generation fabric interconnects provide unified blade, network and storage connectivity
- Cisco UCS 5108 Blade Chassis provide the perfect environment for multi-server type, multi-purpose workloads in a single containment

Infrastructure Simplicity

- Converged, simplified architecture drives increased IT productivity
- Cisco UCS management results in flexible, agile, high performance, self-integrating information technology with faster ROI
- Fabric Extender technology reduces the number of system components to purchase, configure and mainta
- Standards-based, high bandwidth, low latency virtualization-aware unified fabric delivers high density, excellent virtual desktop user-experience

Business Agility

- Model-based management means faster deployment of new capacity for rapid and accurate scalability
- Scale up to 16 Chassis and up to 128 blades in a single Cisco UCS management domain
- Leverage Cisco UCS Management Packs for System Center 2012 for integrated management

Benefits of Nexus 5548UP

The Cisco Nexus 5548UP Switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages, including:

Architectural Flexibility

- Unified ports that support traditional Ethernet, Fibre Channel (FC),and Fibre Channel over Ethernet (FCoE)
- Synchronizes system clocks with accuracy of less than one microsecond, based on IEEE 1588
- Offers converged Fabric extensibility, based on emerging standard IEEE 802.1BR, with Fabric Extender (FEX) Technology portfolio, including:

- Cisco Nexus 1000V Virtual Distributed Switch
- Cisco Nexus 2000 FEX
- Adapter FEX
- VM-FEX

Infrastructure Simplicity

- Common high-density, high-performance, data-center-class, fixed-form-factor platform
- Consolidates LAN and storage
- Supports any transport over an Ethernet-based fabric, including Layer 2 and Layer 3 Traffic
- Supports storage traffic, including iSCSI, NAS, FC, RoE, and IBoE
- Reduces management points with FEX Technology

Business Agility

- Meets diverse data center deployments on one platform
- Provides rapid migration and transition for traditional and evolving technologies
- Offers performance and scalability to meet growing business needs

Specifications At-a-Glance

- A 1 rack-unit, 1/10 Gigabit Ethernet switch
- 32 fixed Unified Ports on base chassis and one expansion slot totaling 48 ports
- The slot can support any of the three modules: Unified Ports, 1/2/4/8 native Fibre Channel, and Ethernet or FCoE
- Throughput of up to 960 Gbps

Benefits of EMC VNX Family of Storage Controller

The EMC VNX Family delivers industry leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

All of this is available in a choice of systems ranging from affordable entry-level solutions to high performance, petabyte-capacity configurations servicing the most demanding application requirements. The VNX family includes the VNXe Series, purpose-built for the IT generalist in smaller environments, and the VNX Series , designed to meet the high-performance, high scalability, requirements of midsize and large enterprises.

VNX Series—Simple, Efficient, Powerful

A robust platform for consolidation of legacy block storage, file-servers, and direct-attached application storage, the VNX series enables organizations to dynamically grow, share, and cost-effectively manage multi-protocol file systems and multi-protocol block storage access. The VNX Operating environment enables Microsoft Windows and Linux/UNIX clients to share files in multi-protocol (NFS and CIFS) environments. At the same time it supports iSCSI, Fiber Channel, and FCoE access for high bandwidth and latency-sensitive block applications. The combination of EMC Atmos Virtual Edition software and VNX storage supports object-based storage and enables customers to manage web applications from

EMC Unisphere. The VNX series next generation storage platform is powered by Intel quad-core Xeon 5600 series with a 6 –Gb/s SAS drive back-end and delivers demonstrable performance improvements over the previous generation mid-tier storage:

- Run Microsoft SQL and Oracle 3x to 10x faster
- Enable 2x system performance in less than 2 minutes –non-disruptively
- Provide up to 10 GB/s bandwidth for data warehouse applications

Benefits of VMware ESXi 5.1

Virtualization is now a critical component to an overall IT strategy, it is important to choose the right vendor. VMware is the leading business virtualization infrastructure provider, offering the most trusted and reliable platform for building private clouds and federating to public clouds.

Find out how VMware delivers on the core requirements for a business virtualization infrastructure solution:

- Built on a robust, reliable foundation
- Delivers a complete virtualization platform from desktop through the datacenter out to the public cloud
- Provides the most comprehensive virtualization and cloud management
- Integrates with your overall IT infrastructure
- Proven with over 350,000 customers

And best of all, VMware provides:

- Low total-cost-of-ownership (TCO)

For detailed information about vSphere 5.1, go to:

http://www.vmware.com/files/pdf/products/vsphere/vmware-what-is-new-vsphere51.pdf

Benefits of VMware Horizon View 5.2

Deliver rich, personalized virtual desktops as a managed service from a virtualization platform built to deliver the entire desktop, including the operating system, applications and data. With Horizon View, desktop administrators virtualize the operating system, applications, and user data and deliver modern desktops to end-users. Get centralized automated management of these components for increased control and cost savings. Improve business agility while providing a flexible high performance desktop experience with VMware Horizon View, desktop administrators virtualize the operating system, applications, and user data and deliver modern desktops to end-users, across a variety of network conditions.

Deliver Business Agility

Bring the agility and availability of cloud computing to the desktop and applications with Horizon View. Built on VMware vSphere, Horizon View delivers desktops from a single integrated platform as part of your cloud services.

Dynamically allocate resources to enable highly responsive environment to end users. Scale up and down desktop services on demand to quickly meet changing business needs and proactively protect against planned and unplanned downtime. Run your desktops as business critical services for your workforce.

Easily Control and Manage Desktops

Increase control of desktops, applications and data by delivering and managing them as centralized services.

A single, powerful administrative console provides oversight of desktop services while enabling IT to simply execute previously cumbersome tasks like provisioning, updates and patches. Easily apply policies, quickly enable and disable users all from a centralized console for optimal business response. Free up time from maintenance for technology innovation.

Deliver a Better Desktop Experience

Unlike traditional PCs, View desktops are not tied to the physical computer. Instead, they reside in your cloud and end-users can access their View desktop when needed.

Horizon View with PCoIP delivers the richest, most flexible and adaptive experience for end-users around the world in a variety of network conditions. Business happens everywhere and whether online or offline, desktops or mobile devices, LAN or WAN, Horizon View delivers maximum workplace productivity.

Automate Desktop Operations Management

VMware vCenter Operations Manager for View allows administrators to gain insight into desktop and infrastructure performance, quickly pinpoint and troubleshoot issues. Administrators can optimize resource utilization, and proactively manage the desktop environment through the management dashboards. vCenter Operations Manager for View is an optional add-on for Horizon View customers. You can also leverage PCoIP Extension Services to collect Horizon View statistics into your existing WMI tool.

Built-in Security

Maintain control over data and intellectual property by keeping it secure in the datacenter. Encrypted protocol traffic provides secure end-users access virtual desktops inside or outside of the corporate network. Integration with vShield Endpoint enables offloaded and centralized anti-virus and anti-malware (AV) solutions. This integration helps to eliminate agent sprawl and AV storm issues while minimizing the risk of malware infection and simplifying AV administration. VMware View also supports integration with RSA SecureID for 2-factor authentication requirements.

Summary of Main Findings

The combination of technologies from Cisco Systems, Inc, VMware and EMC produced a highly efficient, robust and scalable Desktop Virtualization (DV) infrastructure for a hosted virtual desktop deployment. Key components of the solution included:

- The combined power of the Cisco Unified Computing System, Nexus switching and EMC storage hardware with VMware ESXi 5.1, and VMware Horizon View 5.2 software produces a high density per blade and per chassis Virtual Desktop delivery system.
- Cisco UCS B200 M3 half-width blade with dual 8-core processors and 256GB of memory running at 1600 MHz supports 30% more virtual desktop workloads than the previously studied full width blade using a new medium workload with flash. In addition, density achieved with Horizon View 5.2 is equivalent to a prior study on the same platform with View 5.1 Update 2.

- The study design based on two Unified Computing System chassis, each with seven Cisco UCS B200 M3 blades, each with dual 8-core processors and 256GB of memory running at 1600 MHz and a Cisco VIC 1240 converged network adapter supports 2000 virtual desktop workloads running the new medium workload with flash providing outstanding End User Experience with average response times under 1.75 seconds at full load.
- Able to boot the full complement of 2000 virtual desktops (ready to login) in under 20 minutes.
- Able to ramp up (log in and start workloads) to steady state with all 2000 users running a knowledge worker workload with flash in 30 minutes without pegging the processor, exhausting memory or storage subsystems.
- Our design provides N+1 server fault tolerance for the 2000 virtual desktop system, making the design fully fault tolerant from end to end.
- Compared to previous studies with full width blades, the rack space required to support 2000 users was reduced from 30 Rack Units to 12 Rack Units.
- Pure Virtualization: We continue to present a validated design that is 100% virtualized on ESXi 5.1. All of the Windows 7 SP1 virtual desktops and supporting infrastructure components, including vCenter, Active Directory, Profile Servers, SQL Servers, and Horizon View 5.2 components were hosted as virtual servers.
- Maintained our industry leadership with our new Cisco UCS Manager 2.1(1b) software that makes scaling simple, consistency guaranteed and maintenance simple.
- Our 10G unified fabric story gets additional validation on second generation 6200 Series Fabric Interconnects and second generation Nexus 5500 Series access switches as we run more challenging workload testing, maintaining unsurpassed user response times.
- EMC's VNX5500 system provides storage consolidation and outstanding efficiency. Both block and file bases storage resources are available on a single system, utilizing EMC Fast Cache technology.
- VMware Horizon View 5.2 with the Sparse Virtual Disk feature used for floating assignment linked clones provided better disk performance and space efficiency.

Architecture

Deployed Hardware

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

The 2000 User Horizon View 5.2 solution includes Cisco networking, Cisco Unified Computing System and EMC storage, all of which fits in two data center racks (one for the EMC VNX and one for the Cisco networking and Cisco UCS gear.) In fact, there is adequate rack space in the Cisco rack to add blades and chassis to support an additional 4000 users

This document details the deployment of VMware Horizon View 5.2 floating assignment linked clones on VMware ESXi 5.1. Cisco Nexus 1000V distributed switch manages the two VMware Clusters hosting the virtual desktops, insuring end-to-end Quality of Service and ease of management by the network team.

Figure 1 VMware Horizon View 5.2 2000 User Hardware Components

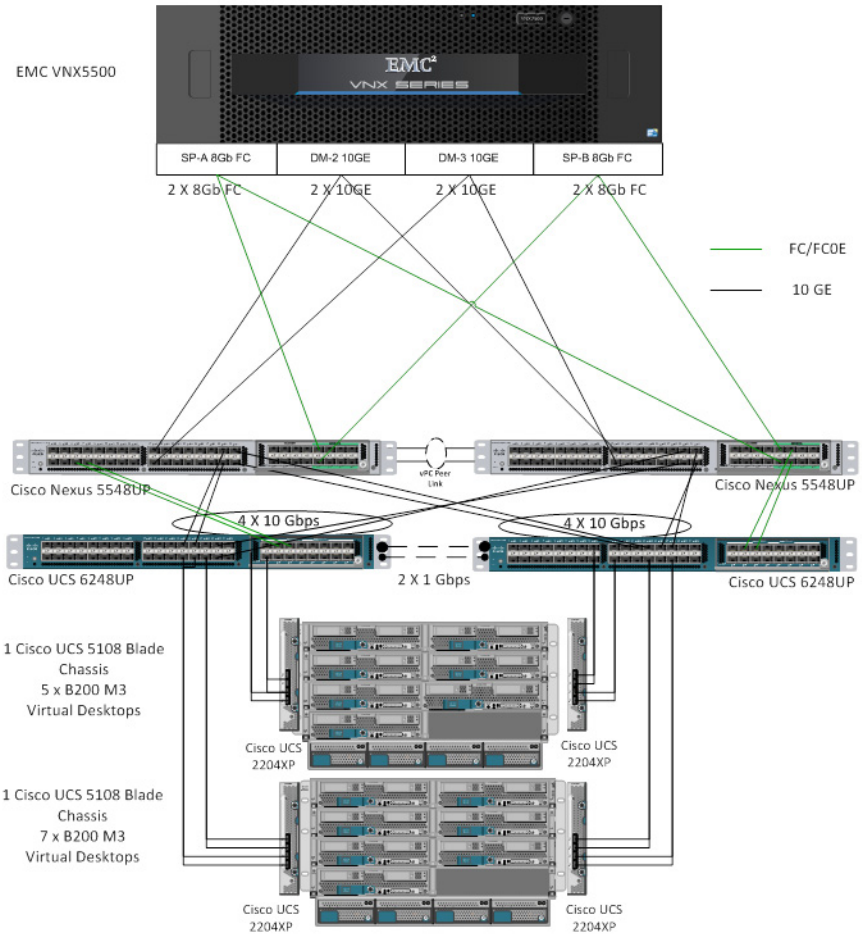

The reference configuration includes:

- Two Cisco Nexus 5548UP switches with 16-universal port Expansion Modules (Optional)
- Two Cisco UCS 6248 Series Fabric Interconnects with Cisco UCS 6200 16-universal port Expansion Modules (Optional)
- Two Cisco UCS 5108 Blade Server Chassis with two 2204XP IO Modules per chassis
- Fourteen Cisco UCS B200 M3 Blade Servers with Intel E5-2690 processors, 256 GB RAM, and VIC 1240 mezzanine cards for Horizon View 5.2 virtual desktops (providing N+1 Server fault tolerance for the system)
- One EMC VNX5500 dual controller storage system for HA
- Two Cisco UCS B200 M3 Blade Servers with Intel E5-2650 processors, 96 GB RAM, and VIC 1240 mezzanine card for infrastructure (not shown in the drawing above)

The EMC VNX5500 disk shelf, disk and Fast Cache configurations are detailed in section “Storage Architecture Design” later in this document.

Software Revisions

Table 1 Software Used in this Deployment

Layer	Compute	Version or Release	Details
Compute	Cisco UCS Fabric Interconnect Cisco UCS B200 M3	2.1 (1a) 2.1 (1b)	Embedded Management Hardware BIOS
Network	Nexus 5500 Switch	5.2(1)N1(1)	Operating System Version
Storage	EMC VNX5500	05.32.000.6.203 Block 7.1.65-8 File	Operating System Version
Software	Cisco UCS Blade Hosts Cisco Nexus 1000V	VMware ESXi 5.1 4.2(1)SV1(5.2)	Operating System Version Virtual Switch appliance version

Configuration Guidelines

The 2000 User Horizon View 5.2 solution described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided that refer to which redundant component is being configured with each step, whether that be A or B. For example, SP A and SP B are used to identify the two EMC VNX storage controllers that are provisioned with this document while Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured similarly.

This document is intended to allow the reader to configure the VMware Horizon View 5.2 customer environment as stand-alone solution.

VLANs

For the 2000 User Horizon View 5.2 solution, we utilized VLANs to isolate and apply access strategies to various types of network traffic. Table 2 details the VLANs used in this study.

Table 2 VLANs

VLAN Name	VLAN ID	Purpose	Native
VDA	122	Virtual Desktops	No
MGMT	164	ESXi, N1KV Management	Yes
INFRA	165	Infrastructure VMs	No
N1K-Control	167	N1KV Control	No
N1K-Packet	168	N1KV Packet	No
vMOTION	169	vMotion	No

VMware Clusters

Four VMware Clusters were used to support the solution and testing environment:

- Infrastructure Cluster (vCenter, Active Directory, DNS, DHCP, SQL Clusters, VMware View Connection Servers, View Composer, and Nexus 1000V Virtual Switch Manager appliances, etc.)
- VDA Clusters (2) (Windows 7 SP1 32-bit pooled virtual desktops; 1000 per cluster per VMware best practices recommended Horizon View 5.2 desktop cluster density.)
- Launcher Cluster (The Login Consultants Login VSI launcher infrastructure was hosted on the same Cisco UCS Domain sharing switching, but running on local storage).

Infrastructure Components

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

Cisco Unified Computing System

Cisco Unified Computing System (UCS) is a set of pre-integrated data center components that comprises blade servers, adapters, fabric interconnects, and extenders that are integrated under a common embedded management system. This approach results in far fewer system components and much better manageability, operational efficiencies, and flexibility than comparable data center platforms.

Cisco Unified Computing System Components

Cisco UCS components are shown in Figure 2.

Figure 2 Cisco Unified Computing System Components

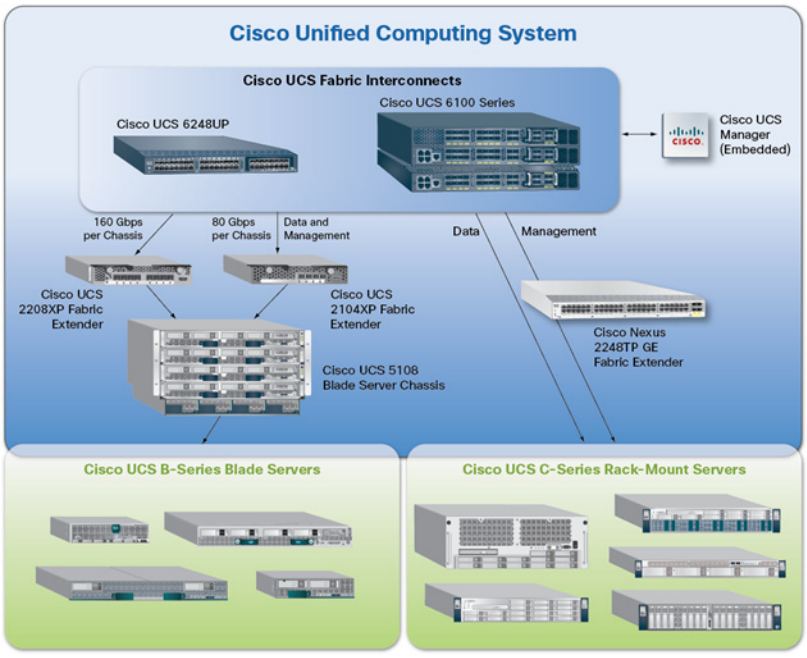

The Cisco Unified Computing System is designed from the ground up to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards, even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration, associate a model's service profile with hardware resources and the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender technology reduces the number of system components to purchase, configure, manage, and maintain by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly as physical networks are, but with massive scalability. This represents a radical simplification over traditional systems, reducing capital and operating costs while increasing business agility, simplifying and speeding deployment, and improving performance.

Fabric Interconnect

Cisco UCS Fabric Interconnects create a unified network fabric throughout the Cisco Unified Computing System. They provide uniform access to both networks and storage, eliminating the barriers to deploying a fully virtualized environment based on a flexible, programmable pool of resources.

Cisco Fabric Interconnects comprise a family of line-rate, low-latency, lossless 10-GE, Cisco Data Center Ethernet, and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus 5000 Series, Cisco UCS 6000 Series Fabric Interconnects provide the additional features and management capabilities that make them the central nervous system of Cisco Unified Computing System.

The Cisco UCS Manager software runs inside the Cisco UCS Fabric Interconnects. The Cisco UCS 6000 Series Fabric Interconnects expand the Cisco UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blades and Cisco UCS Blade Server Chassis.

All chassis and all blades that are attached to the Fabric Interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6200 Series provides the flexibility to support LAN and SAN connectivity for all blades within its domain right at configuration time. Typically deployed in redundant pairs, the Cisco UCS Fabric Interconnect provides uniform access to both networks and storage, facilitating a fully virtualized environment.

The Cisco UCS Fabric Interconnect family is currently comprised of the Cisco 6100 Series and Cisco 6200 Series of Fabric Interconnects.

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a 1 RU, 10-GE, Cisco Data Center Ethernet, FCoE interconnect providing more than 1Tbps throughput with low latency. It has 32 fixed ports of Fibre Channel, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

One expansion module slot can be up to sixteen additional ports of Fibre Channel, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+.

Cisco UCS 6248UP 48-Port Fabric Interconnects were used in this study.

Cisco UCS 2200 Series IO Module

The Cisco UCS 2100/2200 Series FEX multiplexes and forwards all traffic from blade servers in a chassis to a parent Cisco UCS Fabric Interconnect over from 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis, or VMs on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the Fabric Interconnect. At the core of the Cisco UCS Fabric Extender are ASIC processors developed by Cisco that multiplex all traffic.

Note Up to two fabric extenders can be placed in a blade chassis.

Cisco UCS 2104 has eight 10GBASE-KR connections to the blade chassis mid-plane, with one connection per fabric extender for each of the chassis' eight half slots. This gives each half-slot blade server access to each of two 10-Gbps unified fabric-based networks via SFP+ sockets for both throughput and redundancy. It has 4 ports connecting up the fabric interconnect.

Cisco UCS 2208 has thirty-two 10GBASE-KR connections to the blade chassis midplane, with one connection per fabric extender for each of the chassis' eight half slots. This gives each half-slot blade server access to each of two 4x10-Gbps unified fabric-based networks via SFP+ sockets for both throughput and redundancy. It has 8 ports connecting up the fabric interconnect.

Note Cisco UCS 2208 fabric extenders were utilized in this study.

Cisco UCS Chassis

The Cisco UCS 5108 Series Blade Server Chassis is a 6 RU blade chassis that will accept up to eight half-width Cisco UCS B-Series Blade Servers or up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The Cisco UCS 5108 Series Blade Server Chassis can accept four redundant power supplies with automatic load-sharing and failover and two Cisco UCS (either 2100 or 2200 series) Fabric Extenders. The chassis is managed by Cisco UCS Chassis Management Controllers, which are mounted in the Cisco UCS Fabric Extenders and work in conjunction with the Cisco UCS Manager to control the chassis and its components.

A single Cisco UCS managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time Cisco supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows the Cisco UCS to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear, and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

Cisco UCS B200 M3 Blade Server

Cisco UCS B200 M3 is a third generation half-slot, two-socket Blade Server. The Cisco UCS B200 M3 harnesses the power of the latest Intel® Xeon® processor E5-2600 product family, with up to 384 GB of RAM (using 16-GB DIMMs), two optional SAS/SATA/SSD disk drives, and up to dual 4x 10 Gigabit Ethernet throughput, utilizing our VIC 1240 LAN on motherboard (LOM) design. The Cisco UCS B200 M3 further extends the capabilities of Cisco UCS by delivering new levels of manageability, performance, energy efficiency, reliability, security, and I/O bandwidth for enterprise-class virtualization and other mainstream data center workloads.

Cisco UCS VIC1240 Converged Network Adapter

A Cisco innovation, the Cisco UCS Virtual Interface Card (VIC) 1240 (Figure 1) is a 4-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional Port Expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

The Cisco UCS VIC 1240 enables a policy-based, stateless, agile server infrastructure that can present up to 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 1240 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.

Figure 3 Cisco UCS VIC1240 Converged Network Adapter

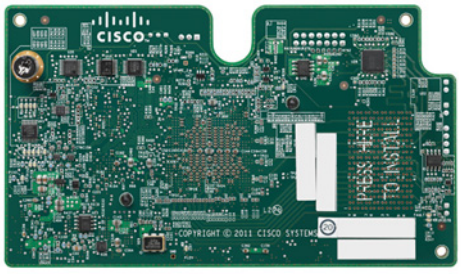

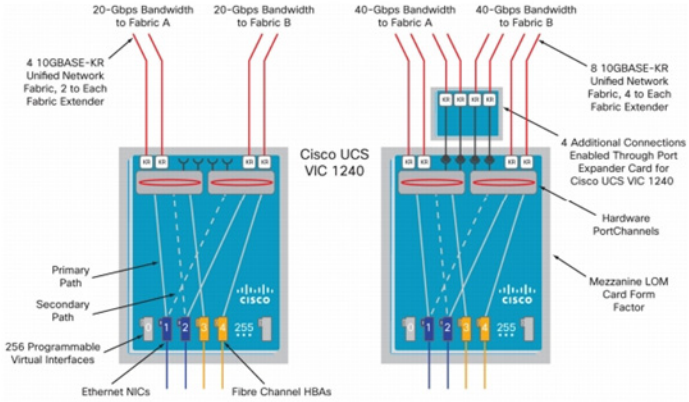

Note The Cisco UCS VIC1240 virtual interface cards are deployed in the Cisco UCS B-Series B200 M3 blade servers.

VMware Horizon View

VMware Horizon View (formerly known as VMware View) simplifies desktop and application management while increasing security and control. Horizon View delivers a personalized high fidelity experience for end-users across sessions and devices. It enables higher availability and agility of desktop services unmatched by traditional PCs while reducing the total cost of desktop ownership up to 50 percent. End-users can enjoy new levels of productivity and the freedom to access desktops from more devices and locations while giving IT greater policy control.

Horizon View 5.2 Features

Horizon View delivers rich, personalized virtual desktops as a managed service from a virtualization platform built to deliver the entire desktop, including the operating system, applications and data. With VMware Horizon View, desktop administrators virtualize the operating system, applications, and user data and deliver modern desktops to end-users. Get centralized automated management of these components for increased control and cost savings. Improve business agility while providing a flexible high performance desktop experience for end-users, across a variety of network conditions.

Automated Desktop Provisioning

Horizon View Manager provides a single management tool for greater IT efficiency to provision new desktops or groups of desktops, and an easy interface for setting desktop policies. Using a template, you can customize virtual pools of desktops and easily set policies, such as how many virtual machines can be in a pool, or logoff parameters.

Streamlined Application Management

VMware ThinApp application virtualization separates applications from underlying operating systems and reduces conflict between the OS and other applications for increased compatibility and streamlined management. Applications packaged with ThinApp can be run centrally from the datacenter, deployed locally to physical or virtual desktops or on USB drives for deployment flexibility.

Advanced Virtual Desktop Image Management

Horizon View Composer enables the rapid creation of desktop images from a golden image. Updates are instant and guaranteed across any number of virtual desktops. When combined with ThinApp, IT administrators can reduce the number of total images, storage requirements and operational costs.

Automate Desktop Operations Management

VMware vCenter Operations Manager for View allows administrators to gain insight into desktop and infrastructure performance, quickly pinpoint and troubleshoot issues. Administrators can optimize resource utilization, and proactively manage the desktop environment through the management dashboards. vCenter Operations Manager for View is an optional add-on for Horizon View customers. You can also leverage PCoIP Extension Services to collect Horizon View statistics into your existing WMI tool.

Efficient Resource Utilization

Horizon View Storage Accelerator optimizes storage load by caching common image blocks when reading virtual desktop images. Space Efficient Disks continuously reduce the storage needed per desktop. Both these technologies improve storage capacity and utilization, thereby reducing costs of additional hardware.

Built-in Security

Maintain control over data and intellectual property by keeping it secure in the datacenter. Encrypted protocol traffic provides secure end-users access virtual desktops inside or outside of the corporate network. Integration with vShield Endpoint enables offloaded and centralized anti-virus and anti-malware (AV) solutions. To eliminate agent sprawl and AV storm issues, risk of malware infection, and simplify AV administration. Horizon View also supports integration with Radius 2-factor authentication requirements.

Enhancements in Horizon View 5.2

VMware Horizon View 5.2 continues to build upon the advancements released in Horizon View 5.2. TCO was further reduced by optimizing storage reads, improved desktop migration and large scale management, and further enhanced the user-experience with lower bandwidth and client diversity.

Lower Total Cost of Ownership

Space-efficient disks, native in vSphere, reduce storage costs and administrative overhead by efficiently using and reclaiming storage space to minimize Horizon View Composer image size. This lowers storage capacity requirements for persistent desktops and decreases the need to continuously recompose and restore images.

Simplified Management

Improved large-scale management allows customers with large Horizon View deployments to efficiently and logically manage their virtual desktop infrastructure. Overall desktop architecture is simplified with a single VMware vCenter Server™ supporting up to 10,000 desktops in a pod. With support for 32 hosts per pool on VMFS along with NFS and pools spanning multiple VLANs, larger desktop pools can be created to decrease operational costs. Furthermore, View admin UI responsiveness increases and accelerates performance of operations such as provisioning, and rebalance improves the efficiency of the desktop administration team.

VMware vCenter Server virtual appliance support enables greater flexibility in Horizon View infrastructure deployment.

Seamless User Experience

Media services for rich 3D graphics add support for hardware accelerated 3D graphics for the most demanding 3D applications. By virtualizing the graphics processing unit (GPU), you can dedicate or share physical GPU resources across multiple users, providing a rich 3D experience from the data center. Using a combination of software and hardware-accelerated graphics, VMware Horizon View™ provides the greatest flexibility for delivering 3D graphics for virtual desktops and workstation use cases. 3D graphics acceleration is built upon the VMware vSphere® platform. Only Horizon View is designed to fully leverage vSphere, expanding the value of combined solutions.

Horizon View media services already support unified communications for Cisco Unified Communications.

Horizon View HTML access enables users to access desktops based on Horizon View from HTML5-capable browsers to securely access their data and applications. Without requiring the installation of any software or plug-ins, end users conveniently can access their desktops on any device. With Horizon Workspace™ integration, that same desktop convenience is expanded to provide end users access to their desktops, data and apps, all from a single location. Note: Horizon View HTML access is available in the Horizon View feature pack.

VMware Horizon View Clients for iOS and Android with Unity make it easier than ever to access Windows applications on your iPhone, iPad or Android device. Remove the frustration of working with Windows on mobile devices with a new mobile native user interface. With Unity users can easily browse, search, and open Windows applications and files, set applications and files as favorites, and easily switch between running applications.

Windows 8 support gives users ability to use the latest OS inside their virtual desktops. Horizon View Client has also been updated to run on the latest Windows 8 devices.

VMware Horizon View 5.2 Hosted Virtual Desktop Overview

Hosted Virtual Desktop (HVD) uses a hypervisor to host all the desktops in the data center.

Three types of HVD pools are available with Horizon View 5.2: Automated, Manual, and Terminal Services Pools. These pool types are discussed below.

- Automated HVD pools use Horizon View Composer to create some number of HVDs. HVD users can be assigned as floating or dedicated users. Floating users will be assigned randomly to HVDs as they log on. Once the user logs off, the HVD is available for any other user. Dedicated user assignments insure that a user is provided the same HVD each time he or she connects to the Horizon View Connection server. Automated pools can utilize the PCoIP protocol and View Persona Management.
- Automated HVD pools can create two types of HVDs: Full virtual machines created from a vCenter template or Horizon View Composer linked clones which share the same base image and use less storage.
- Manual HVD pools provide access to an existing set of HVDs. Any type of machine that can install the Horizon View Agent is supported. Examples could include vCenter virtual machines, physical machines, or blade PCs. Manual pools support the PCoIP protocol, View Persona Management, and Local Mode.
- Microsoft Terminal Services Pools provide Terminal Services sessions as desktops to Horizon View users. The Horizon View Connection Server manages these sessions in the same way it does for Automated or Manual HVD pools. Terminal Services Pools support View Persona Management.

For this study, we utilized Automated HVD pools with floating user assignments and Horizon View Composer linked clones over the PCoIP protocol.

Note View Persona Manager was not deployed.

The following figure shows the logical architecture for a Horizon View 5.2 deployment, including the optional related product; Thin App. Thin App provides application streaming capability and is not included in this study.

Figure 4 VMware Horizon View 5.2 Architecture

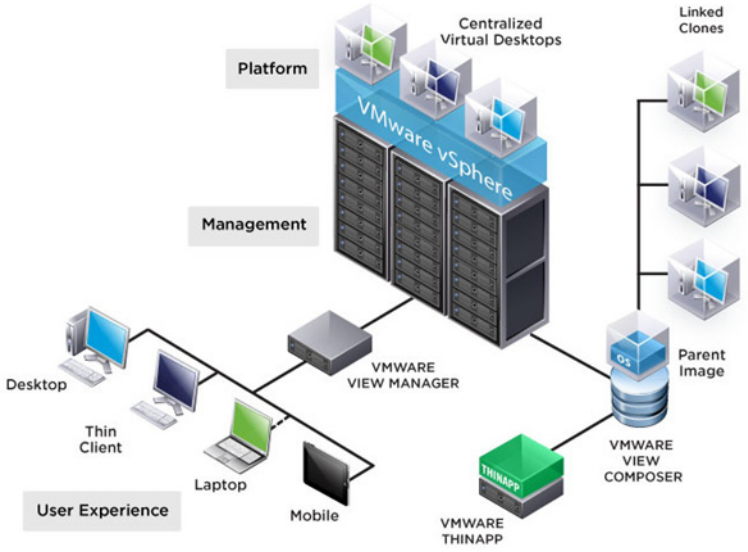

EMC VNX Series

The VNX series delivers uncompromising scalability and flexibility for the mid-tier while providing market-leading simplicity and efficiency to minimize total cost of ownership. Customers can benefit from VNX features such as:

- Next-generation unified storage, optimized for virtualized applications.
- Extended cache by using Flash drives with Fully Automated Storage Tiering for Virtual Pools (FAST VP) and FAST Cache that can be optimized for the highest system performance and lowest storage cost simultaneously on both block and file.
- Multiprotocol supports for file, block, and object with object access through EMC Atmos™ Virtual Edition (Atmos VE).
- Simplified management with EMC Unisphere™ for a single management framework for all NAS, SAN, and replication needs.
- Up to three times improvement in performance with the latest Intel Xeon multicore processor technology, optimized for Flash.
- 6 Gb/s SAS back end with the latest drive technologies supported:
  - 3.5" 100 GB and 200 GB Flash, 3.5" 300 GB, and 600 GB 15k or 10k rpm SAS, and 3.5" 1 TB, 2 TB and 3 TB 7.2k rpm NL-SAS
  - 2.5" 100 GB and 200 GB Flash, 300 GB, 600 GB and 900 GB 10k rpm SAS

- Expanded EMC UltraFlex™ I/O connectivity-Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Common Internet File System (CIFS), network file system (NFS) including parallel NFS (pNFS), Multi-Path File System (MPFS), and Fibre Channel over Ethernet (FCoE) connectivity for converged networking over Ethernet.

The VNX series includes five software suites and three software packs that make it easier and simpler to attain the maximum overall benefits.

- Software suites available:
  - VNX FAST Suite-Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (FAST VP is not part of the FAST Suite for VNX5100™).
  - VNX Local Protection Suite-Practices safe data protection and repurposing.
  - VNX Remote Protection Suite-Protects data against localized failures, outages, and disasters.
  - VNX Application Protection Suite-Automates application copies and proves compliance.
  - VNX Security and Compliance Suite-Keeps data safe from changes, deletions, and malicious activity.
- Software packs available:
  - VNX Total Efficiency Pack-Includes all five software suites (not available for VNX5100).
  - VNX Total Protection Pack-Includes local, remote, and application protection suites.
  - VNX Total Value Pack-Includes all three protection software suites and the Security and Compliance Suite (VNX5100 exclusively supports this package).

EMC VNX5500 Used in Testing

EMC VNX 5500 is a unified storage platform for multi-protocol file, block and object storage. It is powered by Intel quad-core Xeon 5500 series processors and delivers five 9's availability. It is designed to deliver maximum performance and scalability for enterprise and mid-tier companies, enabling them to dramatically grow, share, and cost-effectively manage multi-protocol file and block systems. It supports up to 250 drives and three X-Blades (also known as Data Movers) for file protocol support.This solution was validated Fibre Channel for hypervisor SAN boot, data storage of virtual desktops, SQL database, and infrastructure virtual machines such Horizon View Connection Servers, Horizon View Composer Servers, VMware vCenter Servers, and other supporting services. An NFS or iSCSI variant could be deployed on the VNX5500 using NFS or iSCSI for data storage of virtual desktops.

VMware ESXi 5.1

VMware, Inc. provides virtualization software. VMware's enterprise software hypervisors for servers-VMware vSphere ESX, VMware vSphere ESXi, and VSphere-are bare-metal hypervisors that run directly on server hardware without requiring an additional underlying operating system.

VMware on ESXi 5.1 Hypervisor

ESXi 5.1 is a "bare-metal" hypervisor, so it installs directly on top of the physical server and partitions it into multiple virtual machines that can run simultaneously, sharing the physical resources of the underlying server. VMware introduced ESXi in 2007 to deliver industry-leading performance and scalability while setting a new bar for reliability, security and hypervisor management efficiency.

Due to its ultra-thin architecture with less than 100MB of code-base disk footprint, ESXi delivers industry-leading performance and scalability plus:

- Improved Reliability and Security—with fewer lines of code and independence from general purpose OS, ESXi drastically reduces the risk of bugs or security vulnerabilities and makes it easier to secure your hypervisor layer.
- Streamlined Deployment and Configuration—ESXi has far fewer configuration items than ESX, greatly simplifying deployment and configuration and making it easier to maintain consistency.
- Higher Management Efficiency—the API-based, partner integration model of ESXi eliminates the need to install and manage third party management agents. You can automate routine tasks by leveraging remote command line scripting environments such as vCLI or PowerCLI.
- Simplified Hypervisor Patching and Updating—due to its smaller size and fewer components, ESXi requires far fewer patches than ESX, shortening service windows and reducing security vulnerabilities.

Modular Virtual Desktop Infrastructure Technical Overview

Modular Architecture

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

An increasingly mobile 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 to prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 5). 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 7.

Figure 5 The Evolving Workplace Landscape

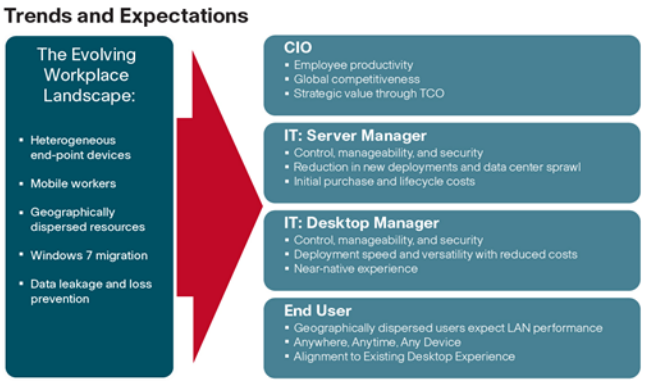

Some of the key drivers for desktop virtualization are increased data security and reduced TCO through increased control and reduced management costs.

Cisco Data Center Infrastructure for Desktop Virtualization

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 (Figure 6).

Figure 6 VMware Horizon View 5.2 on Cisco Unified Computing System

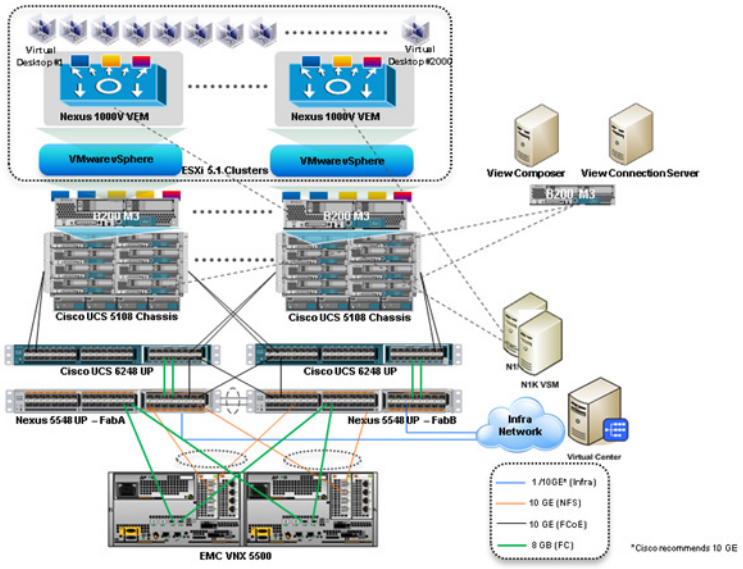

Simplified

Cisco Unified Computing System provides a radical new approach to industry standard computing and provides the heart of the data center infrastructure for desktop virtualization and the Cisco Virtualization Experience (VXI). Among the many features and benefits of Cisco Unified Computing System are the drastic reductions in the number of servers needed and number of cables per server and the ability to very quickly deploy or re-provision servers through Cisco UCS Service Profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco Service Profiles and Cisco storage partners' storage-based cloning. This speeds time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

IT tasks are further simplified through reduced management complexity, provided by the highly integrated Cisco UCS Manager, along with fewer servers, interfaces, and cables to manage and maintain. This is possible due to the industry-leading, highest virtual desktop density per blade of Cisco Unified Computing System along with the reduced cabling and port count due to the unified fabric and unified ports of Cisco Unified Computing System and desktop virtualization data center infrastructure.

Simplification also leads to improved and more rapid success of a desktop virtualization implementation. Cisco and its partners -VMware and EMC - have developed integrated, validated architectures, including available pre-defined, validated infrastructure packages, known as Cisco Solutions for EMC VSPEX End User Computing.

Secure

While virtual desktops are inherently more secure than their physical world predecessors, they introduce new security considerations. Desktop virtualization 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 Unified Computing System and Nexus data center infrastructure for desktop virtualization provides stronger 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 and it is critical to have a solution that can scale predictably with that growth. The Cisco solution supports more virtual 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 Service Profiles allow for on-demand desktop provisioning, making it easy to deploy dozens or thousands of additional desktops.

Each additional Cisco UCS server provides near linear performance and utilizes Cisco's dense memory servers and unified fabric to avoid desktop virtualization bottlenecks. The high performance, low latency network supports high volumes of virtual desktop traffic, including high resolution video and communications.

Cisco Unified Computing System and Nexus data center infrastructure is an ideal platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization.

Savings and Success

As demonstrated above, the simplified, secure, scalable Cisco data center infrastructure solution for desktop virtualization will save time and cost. There will be faster payback, better ROI, and lower TCO with the industry's highest virtual desktop density per server resulting in fewer servers required, reducing both capital expenditures (CapEx) and operating expenditures (OpEx). There will also be much lower network infrastructure costs, with fewer cables per server and fewer ports required, via the Cisco UCS architecture and unified fabric.

The simplified deployment of Cisco Unified Computing System for desktop virtualization speeds up time to productivity and enhances business agility. IT staff and end users are more productive more quickly and the business can react to new opportunities by simply 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, anywhere.

Understanding Desktop User Groups

There must be a considerable effort within the enterprise to identify desktop user groups and their memberships. The most broadly recognized, high-level user groups are as follows:

- Task Workers—Groups of users working in highly specialized environments where the number of tasks performed by each worker is essentially identical. These users are typically located at a corporate facility (for example, call center employees).

- Knowledge/Office Workers—Groups of users who use a relatively diverse set of applications that are Web-based and installed and whose data is regularly accessed. They typically have several applications running simultaneously throughout their workday and a requirement to utilize Flash video for business purposes. This is not a singular group within an organization. These workers are typically located at a corporate office (for example, workers in accounting groups).
- Power Users—Groups of users who run high-end, memory, processor, disk IO, and/or graphic-intensive applications, often simultaneously. These users have high requirements for reliability, speed, and real-time data access (for example, design engineers).
- Mobile Workers—Groups of users who may share common traits with Knowledge/Office Workers, with the added complexity of needing to access applications and data from wherever they are—whether at a remote corporate facility, customer location, at the airport, at a coffee shop, or at home—all in the same day (for example, a company's outbound sales force).
- Remote Workers—Groups of users who could fall into the Task Worker or Knowledge/Office Worker groups but whose experience is from a remote site that is not corporate owned, most often from the user's home. This scenario introduces several challenges in terms of type, available bandwidth, and latency and reliability of the user's connectivity to the data center (for example, a work-from-home accounts payable representative).
- Guest/Contract Workers—Groups of users who need access to a limited number of carefully controlled enterprise applications and data and resources for short periods of time. These workers may need access from the corporate LAN or remote access (for example, a medical data transcriptionist).

There is good reason to search for and identify multiple sub-groups of the major groups listed above in the enterprise. Typically, each sub-group has different application and data requirements.

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?

Provided 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 or Windows XP?
- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 7?
- 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 Thin App be used for streamed applications or will all applications be installed in the image?
- Will you use floating assignment or assigned user desktops?
- Will you use linked clone or full copy desktops?
- How will you manage user persona?
- 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?

Proof of Concept Pilot Program

To validate what you have learned during your analysis, create an isolated Proof of Concept environment to test the various workloads and validate your sizing calculations.

Then create a Pilot environment and get users from each user group who can exercise all of the organizations key applications. Use the pilot user feedback to further refine you design in preparation for production roll out.

Failure to follow these key steps will make a successful virtual desktop deployment project nearly impossible.

Cisco Services

Cisco offers assistance for customers in the analysis, planning, implementation, and support phases of the VDI lifecycle. These services are provided by the Cisco Advanced Services group. Some examples of Cisco services include:

- Cisco VXI Unified Solution Support
- Cisco VXI Desktop Virtualization Strategy Service
- Cisco VXI Desktop Virtualization Planning and Design Service

The Solution: A Unified, Pre-Tested and Validated Infrastructure

To meet the challenges of designing and implementing a modular desktop infrastructure, Cisco, EMC and VMware have collaborated to create the data center solution for virtual desktops outlined in this document.

Key elements of the solution include:

- A shared infrastructure that can scale easily
- A shared infrastructure that can accommodate a variety of virtual desktop workloads

Cisco Networking Infrastructure

This section describes the Cisco networking infrastructure components used in the configuration.

Cisco Nexus 5548 Switch

The Cisco Nexus 5548UP is a 1-RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports that accept modules and cables meeting the Small Form-Factor Pluggable Plus (SFP+) standard and one expansion slot. Expansion slot options include:

- Ethernet module that provides sixteen 1/10 Gigabit Ethernet and FCoE ports using the SFP+ interface.
- Fibre Channel plus Ethernet module that provides eight 1/10 Gigabit Ethernet and FCoE ports using the SFP+ interface, and eight ports of 8/4/2/1-Gbps native Fibre Channel connectivity using the SFP+/SFP interface.
- Unified port module that provides up to sixteen 1/10 Gigabit Ethernet and FCoE ports using the SFP+ interface or up to sixteen ports of 8/4/2/1-Gbps native Fibre Channel connectivity using the SFP+ and SFP interfaces; the use of 1/10 Gigabit Ethernet or 8/4/2/1-Gbps Fibre Channel on a port is mutually exclusive but can be selected for any of the 16 physical ports per module.
- Four port QSFP Ethernet module that provides 4 40 Gigabit Ethernet ports using QSFP interface.

The switch has a single serial console port and a single out-of-band 10/100/1000-Mbps Ethernet management port. Two N+1 redundant, hot-pluggable power supplies and five N+1 redundant, hot-pluggable fan modules provide highly reliable front-to-back cooling.

Cisco Nexus 5500 Series Feature Highlights

The switch family's rich feature set makes the series ideal for rack-level, access-layer applications. It protects investments in data center racks with standards-based Ethernet and FCoE features that allow IT departments to consolidate networks based on their own requirements and timing.

- The combination of high port density, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10 Gigabit Ethernet at the rack level. The switch family has sufficient port density to support single or multiple racks fully populated with blade and rack-mount servers.
- Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible. Hot-swappable power and cooling modules can be accessed from the front panel, where status lights offer an at-a-glance view of switch operation. Front-to-back cooling is

consistent with server designs, supporting efficient data center hot-aisle and cold-aisle designs. Serviceability is enhanced with all customer replaceable units accessible from the front panel. The use of SFP+ ports offers increased flexibility to use a range of interconnect solutions, including copper for short runs and fibre for long runs.

- FCoE and IEEE data center bridging features support I/O consolidation, ease management of multiple traffic flows, and optimize performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5500 Series switches provide additional features that create an even more easily managed, high-performance, unified network fabric.

Features and Benefits

This sections details the specific features and benefits provided by the Cisco Nexus 5500 Series.

10GB Ethernet, FCoE, and Unified Fabric Features

The switch series, using cut-through architecture, supports line-rate 10 Gigabit Ethernet on all ports while maintaining consistently low latency independent of packet size and services enabled. It supports a set of network technologies known collectively as Data Center Bridging (DCB) that increases the reliability, efficiency, and scalability of Ethernet networks. These features allow the switches to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN, and cluster environments. Its ability to connect Fibre Channel over Ethernet (FCoE) to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling.

Low Latency

The cut-through switching technology used in the Cisco Nexus 5500 Series ASICs enables the product to offer a low latency of 3.2 microseconds, which remains constant regardless of the size of the packet being switched. This latency was measured on fully configured interfaces, with access control lists (ACLs), QoS, and all other data path features turned on. The low latency on the Cisco Nexus 5500 Series enables application-to-application latency on the order of 10 microseconds (depending on the NIC). These numbers, together with the congestion management features described in the next section, make the Cisco Nexus 5500 Series a great choice for latency-sensitive environments.

Other Features

Other Nexus 5548UP features include: Nonblocking Line-Rate Performance, Single-Stage Fabric, Congestion Management, Virtual Output Queues, Lossless Ethernet (Priority Flow Control), Delayed Drop FC over Ethernet, Hardware-Level I/O Consolidation, and End-Port Virtualization.

Cisco Nexus 1000V Feature Highlight

Cisco Nexus 1000V Series Switches are virtual machine access switches that are an intelligent software switch implementation based on IEEE 802.1Q standard for VMware vSphere environments running the Cisco® NX-OS Software operating system. Operating inside the VMware ESX hypervisor, the Cisco Nexus 1000V Series supports Cisco VN-Link server virtualization technology to provide:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Non-disruptive operational model for server virtualization and networking teams

With the Cisco Nexus 1000V Series, you can have a consistent networking feature set and provisioning process all the way from the virtual machine access layer to the core of the data center network infrastructure. Virtual servers can now use the same network configuration, security policy, diagnostic tools, and operational models as their physical server counterparts attached to dedicated physical network ports. Virtualization administrators can access pre-defined network policy that follows mobile virtual machines to help ensure proper connectivity, saving valuable time for virtual machine administration.

Developed in close collaboration with VMware, the Cisco Nexus 1000V Series is certified by VMware to be compatible with VMware vSphere, vCenter, ESX, and ESXi, and with many other vSphere features. You can use the Cisco Nexus 1000V Series to manage your virtual machine connectivity with confidence in the integrity of the server virtualization infrastructure.

The Cisco Nexus 1000V Release 2.1 software is being offered in two editions:

- Cisco Nexus 1000V Essential Edition: This is available at no cost and provides most of the comprehensive Layer 2 networking features of the Cisco Nexus 1000V Series, including VXLAN, Cisco vPath for service insertion and chaining, and VMware vCloud Director integration.
- Cisco Nexus 1000V Advanced Edition: This version offers value-added security features such as Domain Host Control Protocol (DHCP) snooping, IP source guard, Dynamic Address Resolution Protocol (ARP) Inspection, and Cisco TrustSec® Secure Group Access (SGA) support (a new feature in Release 2.1). The Cisco VSG zone-based virtual firewall is also included in the Advanced Edition.

Cisco Nexus 1000V Product Architecture

Cisco Nexus 1000V Series Switches have two major components: the Virtual Ethernet Module (VEM), which runs inside the hypervisor, and the external Virtual Supervisor Module (VSM), which manages the VEMs.

Virtual Ethernet Module (VEM)

The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware virtual switch (vSwitch). This level of integration helps ensure that the Cisco Nexus 1000V Series is fully aware of all server virtualization events, such as VMware vMotion and Distributed Resource Scheduler (DRS). The VEM takes configuration information from the VSM and provides advanced networking functions: quality of service (QoS), security features, and monitoring features.

Virtual Supervisor Module (VSM)

The Cisco Nexus 1000V Series VSM controls multiple VEMs as one logical modular switch. Configuration is performed through the VSM and is automatically propagated to the VEMs. Instead of configuring soft switches inside the hypervisor on a host-by-host basis administrators can define configurations for immediate use on all VEMs being managed by the VSM from a single interface.

Cisco Nexus 1000V Features and Benefits

The Cisco Nexus 1000V Series provides a common management model for both physical and virtual network infrastructures through Cisco VN-Link technology, which includes policy-based virtual machine connectivity, mobility of virtual machine security and network properties, and a non-disruptive operational model.

Policy-Based Virtual Machine Connectivity

To facilitate easy creation and provisioning of virtual machines, the Cisco Nexus 1000V Series includes port profiles. Port profiles enable you to define virtual machine network policies for different types or classes of virtual machines and then apply the profiles through the VMware vCenter. Port profiles are a scalable mechanism for configuring networks with large numbers of virtual machines. When the Port Profiles include QoS and security policies, they formulate a complete service-level agreement (SLA) for the virtual machine's traffic.

Mobility of Virtual Machine Security and Network Properties

Network and security policies defined in the port profile follow the virtual machine throughout its lifecycle, whether it is being migrated from one server to another, suspended, hibernated, or restarted. In addition to migrating the policy, the Cisco Nexus 1000V Series VSM moves the virtual machine's network state. Virtual machines participating in traffic-monitoring activities can continue these activities uninterrupted by VMware vMotion operations. When a specific port profile is updated, the Cisco Nexus 1000V Series automatically provides live updates to all the virtual ports using that same port profile. The capability to migrate network and security policies through VMware vMotion makes regulatory compliance much easier to enforce with the Cisco Nexus 1000V Series because the security policy is defined in the same way as for physical servers and is constantly enforced by the switch.

Besides traditional switching capability, the Cisco Nexus 1000V Series offers the Cisco vPath architecture to support virtualized network services with:

- Intelligent Traffic Steering: This feature redirects packets in a network flow to a virtual service virtual machine called a Virtual Service Node (VSN), which can be on a different server. Thus, a VSN is not required on every server, providing flexible and consolidated deployment.
- Performance Acceleration: VEM caches the VSN's decision for a flow, implements the service in all subsequent packets of the flow, and accelerates virtualized network service in the hypervisor kernel.

Cisco Virtual Service Gateway (VSG) is the first VSN to leverage the Cisco vPath architecture and provides multi-tenant, scalable, security services for virtual machines on the Cisco Nexus 1000V Series Switches.

Non-disruptive Operational Model

Because of its close integration with VMware vCenter, the Cisco Nexus 1000V Series allows virtualization administrators to continue using VMware tools to provision virtual machines. At the same time, network administrators can provision and operate the virtual machine network the same way they do the physical network. While both teams work independently, the Cisco Nexus 1000V Series enforces consistent configuration and policy throughout the server virtualization environment. This level of integration lowers the cost of ownership while supporting organizational boundaries among server, network, security, and storage teams.

Inside VMware vCenter, virtual machines are configured as before. For network configuration, port profiles defined on the Cisco Nexus 1000V Series VSM are displayed by VMware vCenter as port groups. Virtualization administrators can take advantage of preconfigured port groups and focus on virtual machine management, and network administrators can use port profiles to apply policy for a large number of ports at the same time. Together, both teams can deploy server virtualization more efficiently and with lower operating costs.

Enhanced Deployment Scenarios

- Optimized server bandwidth for I/O-intensive applications: Today, network interfaces are often dedicated to a particular type of traffic, such as VMware Console or vMotion. With the Cisco Nexus 1000V Series, all network interface cards (NICs) can be treated as a single logical channel with QoS attached to each type of traffic. Consequently, the bandwidth to the server can be more efficiently utilized, with network-intensive applica-tions virtualized.

- Easier security audits with consistent security policy: Security audits on virtual machines are usually more difficult to perform because virtual machines are secured differently than physical servers. As the Cisco Nexus 1000V Series provides persistent security policy to mobile virtual machines, security audits are similar to those for physical servers.
- Virtual machine as basic building block of data center: With the Cisco Nexus 1000V Series, virtual machines are treated the same way as physical servers in security policy, monitoring and troubleshooting, and the operational model between network and server administrators, enabling virtual machines to be true basic building blocks of the data center. These operational efficiencies lead to greater scaling of server virtualization deployments with lower operating expenses.

VMware Product Compatibility

The Cisco Nexus 1000V Series is compatible with VMware vSphere as a VMware vNetwork Distributed Switch (vDS) with support for VMware ESX and ESXi hypervisors and integration with VMware vCenter Server. Cisco Nexus 1000V Series Switches are compatible with the various VMware vSphere features.

Architecture and Design of Horizon View 5.2 on Cisco Unified Computing System and EMC VNX Storage

Design Fundamentals

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 Computer (BYOC) 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? constituted 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 2008 R2, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- Hosted Virtual Desktop: A hosted virtual desktop is a virtual desktop running either on virtualization layer (ESX) or on bare metal hardware. The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- 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 only hosted virtual desktops were validated. Each of the sections provides some fundamental design decisions for this environment.

Hosted Virtual Desktop Design Fundamentals

This section details how the VMware Horizon View 5.2 can be used to deliver a variety of virtual desktop configurations.

Choosing a Display Protocol

A display protocol provides end users with a graphical interface to a View desktop that resides in the datacenter. You can use PCoIP (PC-over-IP), which VMware provides, or Microsoft RDP (Remote Desktop Protocol.)

You can set policies to control which protocol is used or to allow end users to choose the protocol when they login to a desktop.

Note For this study, the PCoIP protocol was used.

VMware View with PCoIP

PCoIP provides an optimized desktop experience for the delivery of the entire desktop environment, including applications, images, audio, and video content for a wide range of users on the LAN or across the WAN. PCoIP can compensate for an increase in latency or a reduction in bandwidth, to make sure that end users can remain productive regardless of network conditions.

PCoIP is supported as the display protocol for View desktops with virtual machines and with physical machines that contain Teradici host cards.

PCoIP Features

Key features of PCoIP include the following:

- For users outside the corporate firewall, you can use this protocol with your company's virtual private network or with View security servers.
- Advanced Encryption Standard (AES) 128-bit encryption is supported and is turned on by default.
- Connections from all types of View clients. For more information, go to:

  http://pubs.vmware.com/view-52/index.jsp?topic=/com.vmware.view.planning.doc/GUID-43E5EEF5-72D9-4CCA-8439-66D6FD6D1A1F.html
- USB redirection is supported.
- Audio redirection with dynamic audio quality adjustment for LAN and WAN is supported.
- Optimization controls for reducing bandwidth usage on the LAN and WAN.
- Multiple monitors are supported. You can use up to four monitors and adjust the resolution for each monitor separately, with a resolution of up to 2560x1600 per display. Pivot display and autofit are also supported. When the 3D feature is enabled, up to 2 monitors are supported with a resolution of up to 1920x1200.
- 32-bit color is supported for virtual displays.
- ClearType fonts are supported.
- Copy and paste of text and images between a local Windows client system and the desktop is supported, up to 1MB. Supported file formats include text, images, and RTF (Rich Text Format). You cannot copy and paste system objects such as folders and files between systems.

Video Quality

- 480p-formatted video You can play video at 480p or lower at native resolutions when the View desktop has a single virtual CPU. If the operating system is Windows 7 and you want to play the video in high-definition Flash or in full screen mode, the desktop requires a dual virtual CPU.
- 720p-formatted video You can play video at 720p at native resolutions if the View desktop has a dual virtual CPU. Performance might be affected if you play videos at 720p in high definition or in full screen mode.
- 1080p-formatted video If the View desktop has a dual virtual CPU, you can play 1080p formatted video, although the media player might need to be adjusted to a smaller window size.
- 3D If you plan to use 3D applications such as Windows Aero themes or Google Earth, the Windows 7 View desktop must have virtual hardware version 8, available with vSphere 5 and later. You must also turn on the pool setting called Windows 7 3D Rendering. Up to 2 monitors are supported, and the maximum screen resolution is 1920 x 1200. This non-hardware accelerated graphics feature enables you to run DirectX 9 and OpenGL 2.1 applications without requiring a physical graphics processing unit (GPU).

Recommended Guest Settings

Recommended guest operating system settings include the following settings:

- For Windows XP desktops: 768MB RAM or more and a single CPU
- For Windows 7 desktops: 1GB of RAM and a dual CPU

Microsoft RDP

Remote Desktop Protocol is the same multichannel protocol many people already use to access their work computer from their home computer. Microsoft Remote Desktop Connection (RDC) uses RDP to transmit data.

Microsoft RDP provides the following features:

- With RDP 6, you can use multiple monitors in span mode. RDP 7 has true multiple monitor support, for up to 16 monitors.
- You can copy and paste text and system objects such as folders and files between the local system and the View desktop.
- RDP supports 32-bit color.
- RDP supports 128-bit encryption.
- You can use this protocol for making secure, encrypted connections to a View security server in the corporate DMZ.
- The following are the RDP-related requirements and considerations for different Windows operating systems and features:For Windows XP and Windows XP Embedded systems, you should use Microsoft RDC 6.x.
- Windows Vista comes with RDC 6.x installed, though RDC 7 is recommended.
- Windows 7 comes with RDC 7 installed. Windows 7 SP1 comes with RDC 7.1 installed.
- You must have RDC 6.0 or later to use multiple monitors.
- For Windows XP desktop virtual machines, you must install the RDP patches listed in Microsoft Knowledge Base (KB) articles 323497 and 884020. If you do not install the RDP patches, a Windows Sockets failed error message might appear on the client.
- The View Agent installer configures the local firewall rule for inbound RDP connections to match the current RDP port of the host operating system, which is typically 3389. If you change the RDP port number, you must change the associated firewall rules.

Note You can download RDC versions from the Microsoft website.

Recommended Guest Settings

Client hardware requirements include the following:

- x86-based processor with SSE2 extensions, with a 800MHz or higher processor speed.
- ARM processor with NEON (preferred) or WMMX2 extensions, with a 600MHz or higher processor speed.
- 128 MB RAM

Choose a User Profile Management System

There are a number of options for managing user profiles for HVDs. The two methods we considered for this study were Microsoft Roaming User Profiles and View Persona Manager. It is important to select and deploy a method so that user settings for software applications and user preferences are maintained, particularly for floating desktops. Both methods are discussed briefly below. (We used Microsoft Roaming User Profiles in the study.)

Microsoft Roaming User Profiles and Folder Redirection

This technology has been around for more than a dozen years. It was significantly enhanced with the introduction of Windows Vista and updated again with Windows 7. Version two (v2) roaming profiles were introduced, adding 8 additional folders that can be redirected. This greatly reduces the time it takes to load the user's profile during logon.Using Roaming User Profiles and Folder redirection require a network shares that all users have access to during the virtual desktop session. The user must have read and write access to their profile folder and folder redirection folder, which get created on first login after Roaming User Profiles is configured.

Utilizing Microsoft Active Directory Group Policy is the recommended method for providing Roaming User Profiles and Folder Redirection to your users. See the article titled Managing Roaming User Data Deployment Guide at the following url for details on how to configure both Roaming User Profiles and Folder Redirection:

http://technet.microsoft.com/en-us/library/cc766489%28WS.10%29.aspx

Note The significant changes to Roaming User Profiles and Folder redirection applies to Microsoft Windows 7 and Microsoft Windows Vista.

VMware Persona Management

You can use View Persona Management with View desktops on physical computers and virtual machines that are not managed by View. View Persona Management retains changes that users make to their profiles. User profiles comprise a variety of user-generated information.

- User-specific data and desktop settings, which allow the desktop appearance to be the same regard less of which desktop a user logs in to.
- Application data and settings. For example, these settings allow applications to remember toolbar positions and preferences.
- Windows registry entries configured by user applications.

To facilitate these abilities, View Persona Management requires storage on a CIFS share equal or greater than the size of the user's local profile.

Minimizing Logon and Logoff Times

View Persona Management minimizes the time it takes to log on to and off of desktops.

- View takes recent changes in the profile on the View desktop and copies them to the remote repository at regular intervals. The default is every 10 minutes. In contrast, Windows roaming profiles wait until logoff time and copy all changes to the server at logoff.
- During logon, View downloads only the files that Windows requires, such as user registry files. Other files are copied to the View desktop when the user or an application opens them from the profile folder in the View desktop.
- With View Persona Management, during logoff, only files that were updated since the last replication are copied to the remote repository.

With View Persona Management, you can avoid making any changes to Active Directory in order to have a managed profile. To configure Persona Management, you specify a central repository, without changing the user's properties in Active Directory. With this central repository, you can manage a user's profile in one environment without affecting the physical machines that users might also log on to.

With View Persona Management, if you provision desktops with VMware ThinApp applications, the ThinApp sandbox data can also be stored in the user profile. This data can roam with the user but does not significantly affect logon times. This strategy provides better protection against data loss or corruption.

Configuration Options

You can configure View personas at several levels: a single View desktop, a desktop pool, an OU, or all View desktops in your deployment. You can also use a standalone version of View Persona Management on physical computers and virtual machines that are not managed by View.

By setting group policies (GPOs), you have granular control of the files and folders to include in a persona:

- Specify whether to include the local settings folder. For Windows 7 or Windows Vista, this policy affects the AppData\Local folder. For Windows XP, this policy affects the Local Settings folder.
- Specify which files and folders to load at login time. For example: Application Data\Microsoft\Certificates. Within a folder, you can also specify files to exclude.
- Specify which files and folders to download in the background after a user logs in to the desktop. Within a folder, you can also specify files to exclude.
- Specify which files and folders within a user's persona to manage with Windows roaming profiles functionality instead of View Persona Management. Within a folder, you can also specify files to exclude.

As with Windows roaming profiles, you can configure folder redirection. You can redirect the same folders that support redirection with Windows Roaming User Profiles.

Accessing USB Devices Connected to the End Point

Administrators can configure the ability to use USB devices, such as thumb flash drives, VoIP (voice-over-IP) devices, and printers, from a View desktop. This feature is called USB redirection. (It was not used in this study.)

When you use this feature, most USB devices that are attached to the local client system become available from a menu in View Client. You use the menu to connect and disconnect the devices.

You can specify which types of USB devices end users are allowed to connect to. For composite devices that contain multiple types of devices, such as a video input device and a storage device, you can split the device so that one device (for example, the video input device) is allowed but the other device (for example, the storage device) is not.

USB devices that do not appear in the menu, but are available in a View desktop, include smart card readers and human interface devices such as keyboards and pointing devices. The View desktop and the local computer use these devices at the same time.

This feature has the following limitations:

- When you access a USB device from a menu in View Client and use the device in a View desktop, you cannot access the device on the local computer.
- USB redirection is not supported on Windows 2000 systems or for View desktops sourced from Microsoft Terminal Servers.

Printing from a View Desktop

The virtual printing feature allows end users with View Client on Windows systems to use local or network printers from a View desktop without requiring that additional print drivers be installed in the View desktop.

The location-based printing feature allows you to map View desktops to the printer that is closest to the endpoint client device.

With virtual printing, after a printer is added on a local Windows computer, View adds that printer to the list of available printers on the View desktop. No further configuration is required. For each printer available through this feature, you can set preferences for data compression, print quality, double-sided printing, color, and so on. Users who have administrator privileges can still install printer drivers on the View desktop without creating a conflict with the virtual printing component. To send print jobs to a USB printer, you can either use the USB redirection feature or use the virtual printing feature.

The location-based printing feature is available for both Windows and non-Windows client systems. Location based printing allows IT organizations to map View desktops to the printer that is closest to the endpoint client device. Using this feature does require that the correct printer drivers be installed in the View desktop.

We did not use virtual printing in this study. Our workload generator, Login VSI, installs a pdf printer into the master image which is utilized for printing during the test.

Other Features to Consider

Horizon View 5.2 supports these additional features that were not deployed in this study:

- Streaming Multimedia with Wyse MMR. (Only used for Windows XP environments.)
- Single Sign-On for Logging In (Workload generator initiates multiple sessions from a single workstation.)
- Multiple Monitor Support (Workload generator supports single monitor.)

Designing a VMware Horizon View 5.2 Deployment

There are several elements that go into the design of a successful Horizon View 5.2 environment. This section covers those topics at a high level. Readers should consult the VMware View Architecture Planning guide for Horizon View 5.2 at the following URL for more details: https://pubs.vmware.com/view-52/topic/com.vmware.ICbase/PDF/horizon-view-52-architecture-planning.pdf

Determine Desktop Pools Required

Based on the analysis performed on user groups and the applications identified that will be supported by the Hosted Virtual Desktop (HVD) environment, a strategy for laying out your desktop pool structure should be create.

For this study, a single user group (knowledge workers) will be tested and the application workload that this group will run will be identified, which is based on the Login VSI 3.6 medium workload (with flash.) Virtual machines will be used as the desktop sourc.

If you use a vSphere virtual machine as a desktop source, you can automate the process of making as many identical virtual desktops as you need. You can set a minimum and maximum number of virtual desktops to be generated for the pool. Setting these parameters ensures that you always have enough View desktops available for immediate use but not so many that you overuse available resources.

Using pools to manage desktops allows you to apply settings or deploy applications to all virtual desktops in a pool. The following examples show some of the settings available:

- Specify which remote display protocol to use as the default for the View desktop and whether to let end users override the default.
- Configure the display quality and bandwidth throttling of Adobe Flash animations.
- If using a virtual machine, specify whether to power off the virtual machine when it is not in use and whether to delete it altogether.
- If using vSphere 4.1 or later, specify whether to use a Microsoft Sysprep customization specification or QuickPrep from VMware. Sysprep generates a unique SID and GUID for each virtual machine in the pool.
- Specify whether the View desktop can or must be downloaded and run on a local client system.

In addition, using desktop pools provides many conveniences.

- Dedicated-assignment pools: Each user is assigned a particular View desktop and returns to the same virtual desktop at each login. Users can personalize their desktops, install applications, and store data.
- Floating-assignment pools: The virtual desktop is optionally deleted and re-created after each use, offering a highly controlled environment. A floating-assignment desktop is like a computer lab or kiosk environment where each desktop is loaded with the necessary applications and all desktops have access to necessary data.

Using floating-assignment pools also allows you to create a pool of desktops that can be used by shifts of users. For example, a pool of 100 desktops could be used by 300 users if they worked in shifts of 100 users at a time.

Note For this study, Automated Pools with Floating Assignments in conjunction with View Composer linked clones was used.

Managing Storage Requirements

VMware vSphere lets you virtualize disk volumes and file systems so that you can manage and configure storage without having to consider where the data is physically stored.

Fibre Channel SAN arrays, iSCSI SAN arrays, and NAS arrays are widely used storage technologies supported by VMware vSphere to meet different datacenter storage needs. The storage arrays are connected to and shared between groups of servers through storage area networks. This arrangement allows aggregation of the storage resources and provides more flexibility in provisioning them to virtual machines.

With View 4.5 and later and vSphere 4.1 and later, you can now also use the following features:

- vStorage thin provisioning, which lets you start out with as little disk space as necessary and grow the disk to add space later
- Tiered storage, which allows you to distribute virtual disks in the View environment across high performance storage and lower-cost storage tiers, to maximize performance and cost savings
- Local storage on the ESX/ESXi host for the virtual machine swap files in the guest operating system.

With Horizon View 5.2 and later and vSphere 5.0 and later, you can now also use the following features:

- With the View storage accelerator feature, you can configure ESXi hosts to cache virtual machine disk data.

  Using this content-based read cache (CBRC) can reduce IOPS and improve performance during boot storms, when many desktops start up and run anti-virus scans at the same time. Instead of reading the entire OS from the storage system over and over, a host can read common data blocks from cache.

- You can deploy a desktop pool on a cluster that contains up to 32 ESXi hosts, but you must store the replica disks on NFS datastores.

  Although replica disks must be stored on NFS datastores, OS disks and persistent disks can be stored on NFS or VMFS datastores.

View Composer

Because View Composer creates desktop images that share virtual disks with a base image, you can reduce the required storage capacity by 50 to 90 percent.

View Composer uses a base image, or parent virtual machine, and creates a pool of up to 1,000 linked-clone virtual machines. Each linked clone acts like an independent desktop, with a unique host name and IP address, yet the linked clone requires significantly less storage.

When creating a linked-clone desktop pool, a full clone is first made from the parent virtual machine. The full clone, or replica, and the clones linked to it can be placed in a variety of locations. The options are:

- Replica and linked clones on same datastore
- Replica and lined clones on different datastores

  As an example, you could place the replicas on low capacity read optimized drives with IOPS and place the linked clones on traditional spinning media

- Disposable Disks for Paging and Temp Files

  Guest OS page files and temp files are placed here. When the HVD is powered off, this disk is deleted

- Persistent disks for dedicated desktops

  End user's application data and profiles are stored here. The data survives refresh, recompose and rebalance operations.

- Local datastores for floating or stateless desktops

  Host local drives store linked clone files, presenting some advantages and several disadvantages. Use this option with care after considering your requirements.

Note For this study, the technique was to have the replicas and linked clones on different datastores.

Hosted Virtual Desktop Infrastructure Design

To implement the automated pool floating desktop delivery model for this study, the VMware View Reference Architecture for virtual desktop delivery was followed.

Figure 7 View Desktop Infrastructure

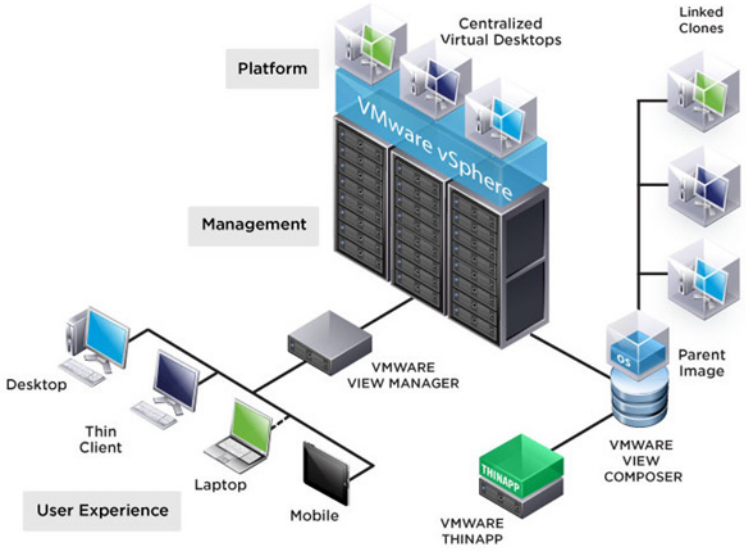

Learn more about VMware Horizon View planning and design at the following location:

https://pubs.vmware.com/view-52/topic/com.vmware.ICbase/PDF/horizon-view-52-architecture-planning.pdf

Solution Validation

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

Configuration Topology for Scalable VMware View 5.2 Virtual Desktop Infrastructure on Cisco Unified Computing System and EMC Storage

Figure 8 illustrates the Cisco Unified Computing System VDI configuration.

Figure 8 Architecture Block

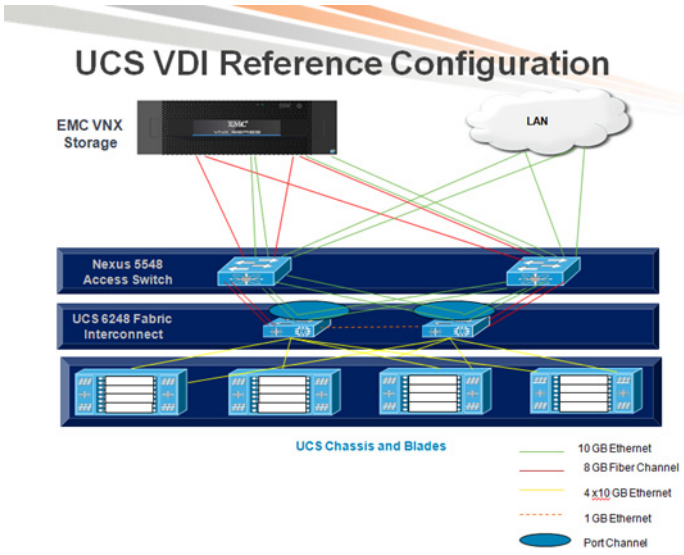

Figure 8 illustrates the architectural topology for the purpose of this study. The architecture is divided into four distinct layers:

- Cisco UCS Compute Platform
- The Virtual Desktop Infrastructure that runs on Cisco UCS blade hypervisor hosts
- Network Access layer and LAN
- Storage Access Network (SAN) and EMC VNX Storage array

Figure 9 details the physical configuration of the 2000 seat View 5.2 environment.

Figure 9 Detailed Architecture Configuration

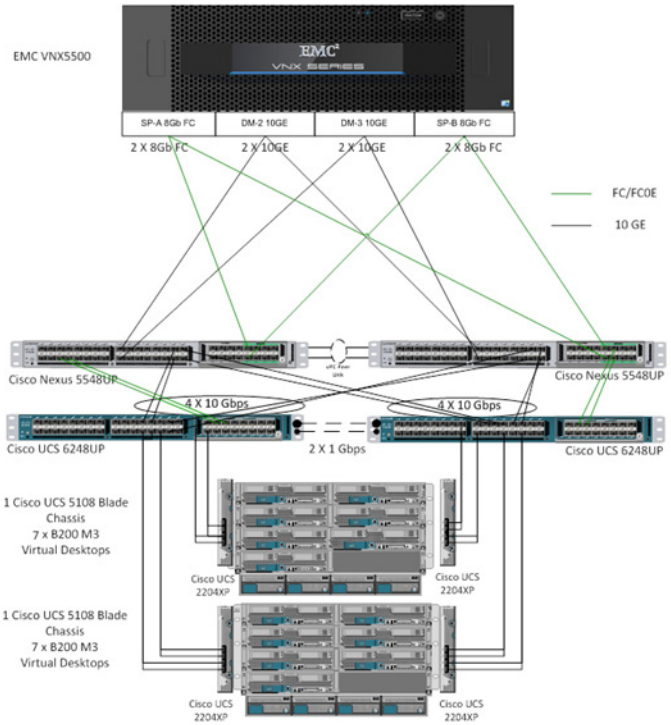

Cisco Unified Computing System Configuration

This section talks about 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 http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html) and it is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco UCS CLI Configuration guide http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.1/b_UCSM_CLI_Configuration_Guide_2_1.pdf
- Cisco UCS-M GUI Configuration guide http://www.cisco.com/en/US/partner/docs/unified_computing/ucs/sw/gui/config/guide/2.1/b_UCSM_GUI_Configuration_Guide_2_1.html

Base Cisco Unified Computing System Configuration

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

1. Bring up the Fabric interconnect and from a Serial Console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. Now bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary and ask if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname. When this is done all access to the FI can be done remotely. You will also configure the virtual IP address to connect to the FI, you need a total of three IP address to

bring it online. You can also wire up the chassis to the FI, using either 1, 2 or 4 links per IO Module, depending on your application bandwidth requirement. We connected all the four links to each module.

2. Now connect using your favorite browser to the Virtual IP and launch the Cisco UCS Manager. The Java based Cisco UCS Manager will let you do everything that you could do from the CLI. We will highlight the GUI methodology here.

3. First check the firmware on the system and see if it is current. Visit http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.0(4d)&relind=AVAILABLE&rellifecycle=&reltype=latest to download the most current Cisco UCS Infrastructure and Cisco UCS Manager software. Use the Cisco UCS Manager Equipment tab in the left pane, then the Firmware Management tab in the right pane and Packages sub-tab to view the packages on the system. Use the Download Tasks tab to download needed software to the FI. The firmware release used in this paper is 2.1(1a).

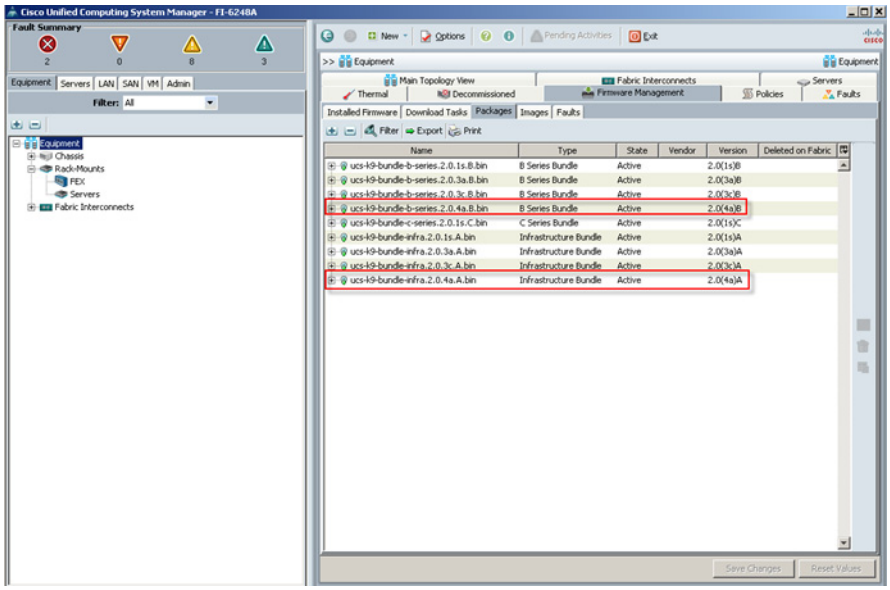

If the firmware is not current, follow the installation and upgrade guide to upgrade the Cisco UCS Manager firmware. We will use Cisco UCS Policy in Service Profiles later in this document to update all Cisco UCS components in the solution.

Note The Bios and Board Controller version numbers do not track the IO Module, Adapter, nor CIMC controller version numbers in the packages.

4. Configure and enable the server ports on the FI. These are the ports that will connect the chassis to the FIs.

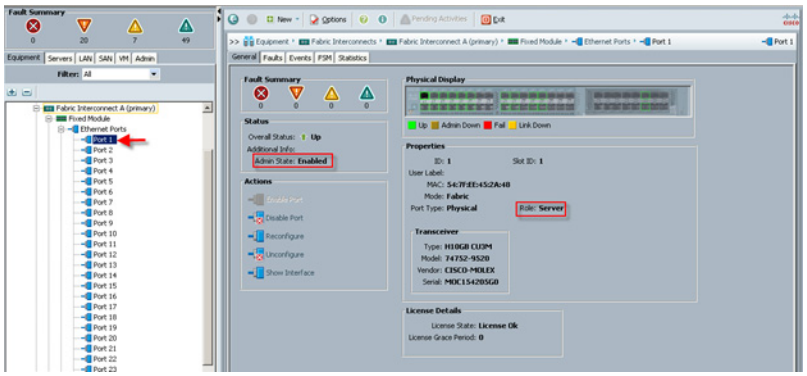

5. Configure and enable uplink Ethernet ports.

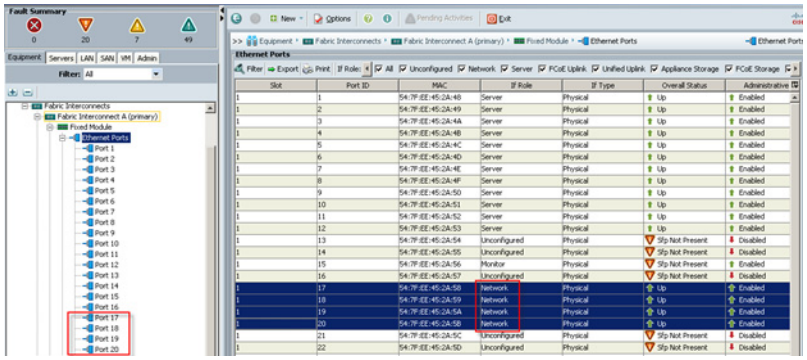

6. Configure and enable FC uplink ports.

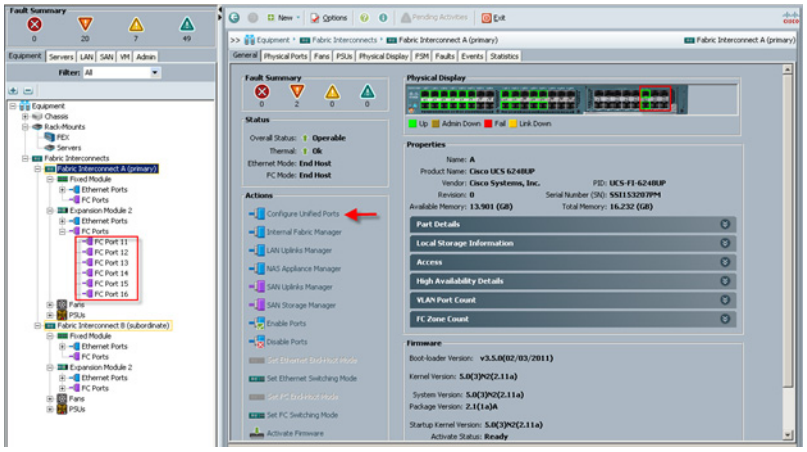

Use the Configure Unified Ports, Configure Expansion Module Ports to configure FC uplinks.

Note In this example, six FC ports are configured, two of which are in use.

7. From the LAN tab in the Navigator pane, configure the required Port Channels and Uplink Interfaces on both Fabric Interconnects.

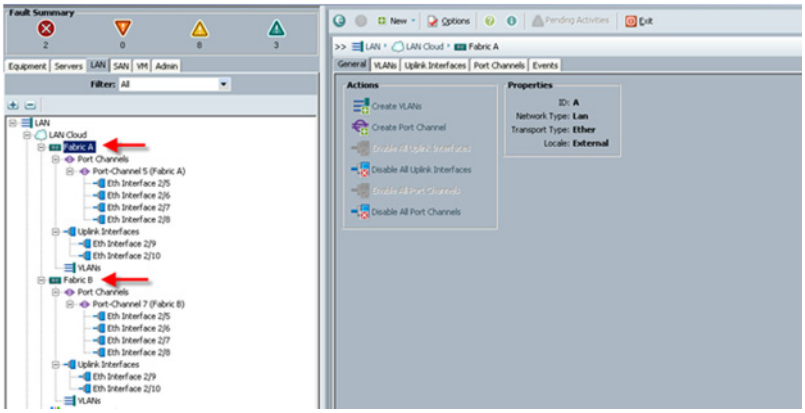

8. Expand the Chassis node in the left pane; click each chassis in the left pane, then click Acknowledge Chassis in the right pane to bring the chassis online and enable blade discovery.

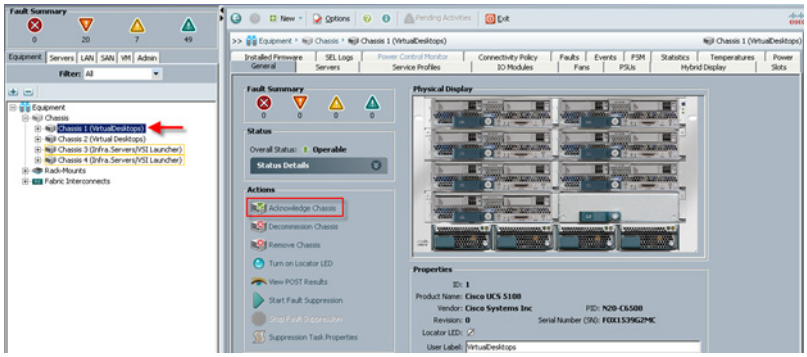

9. Use the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and Licensing. Configuring your Management IP Pool (which provides IP based access to the KVM of each Cisco UCS Blade Server,) Time zone Management (including NTP time source(s)) and uploading your license files are critical steps in the process.

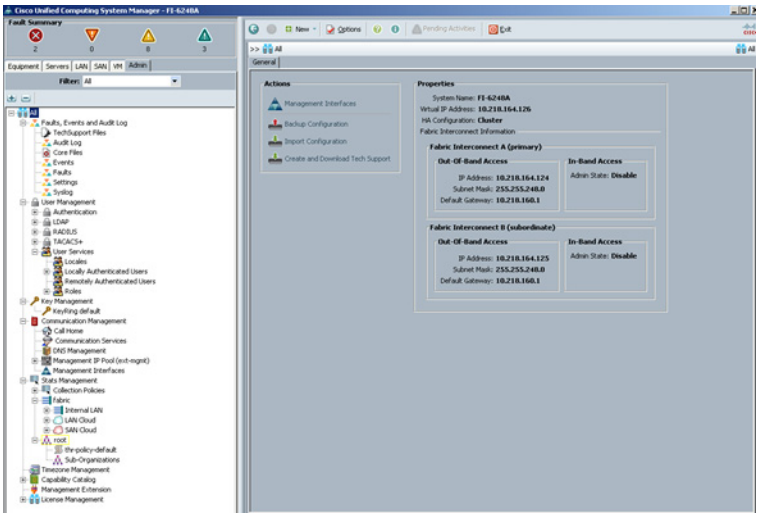

10. Create all the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, Server pool.
11. From the LAN tab in the navigator, under the Pools node, a MAC address pool was created of sufficient size for the environment. In this project, a single pool with two address ranges for expandability was created.

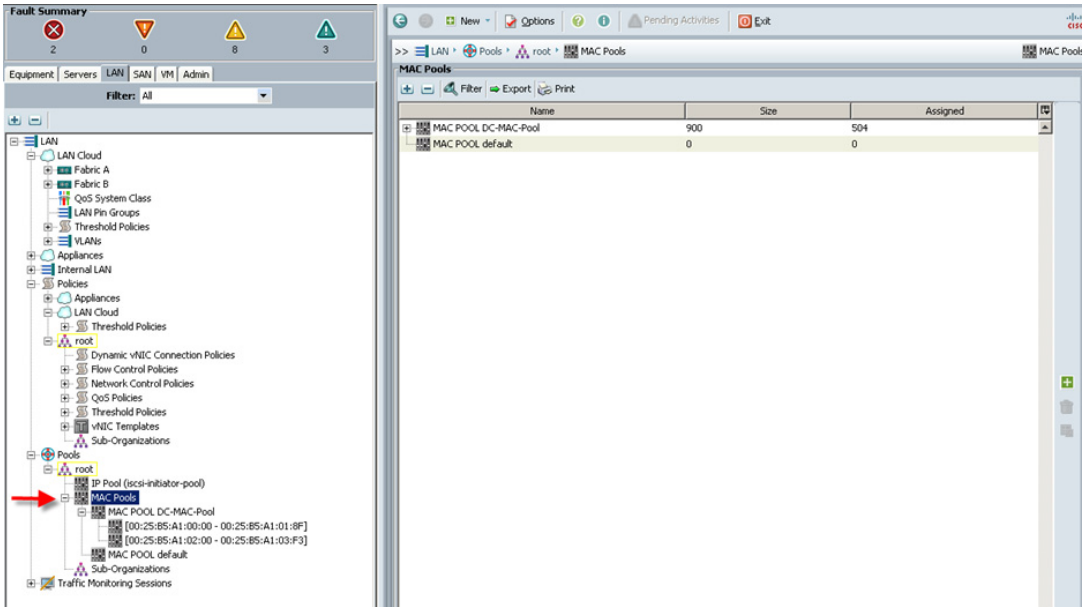

12. For Fiber Channel connectivity, WWNN and WWPN pools must be created from the SAN tab in the navigator pane, in the Pools node.

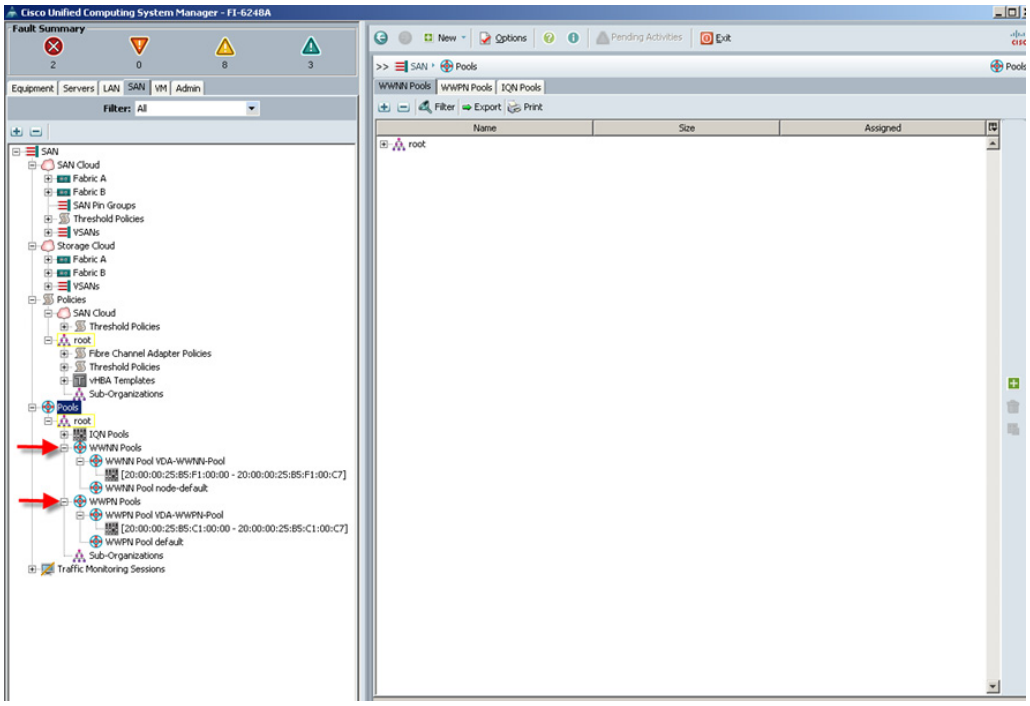

13. For this project, a single VSAN was used, with the default VSAN with ID 1.

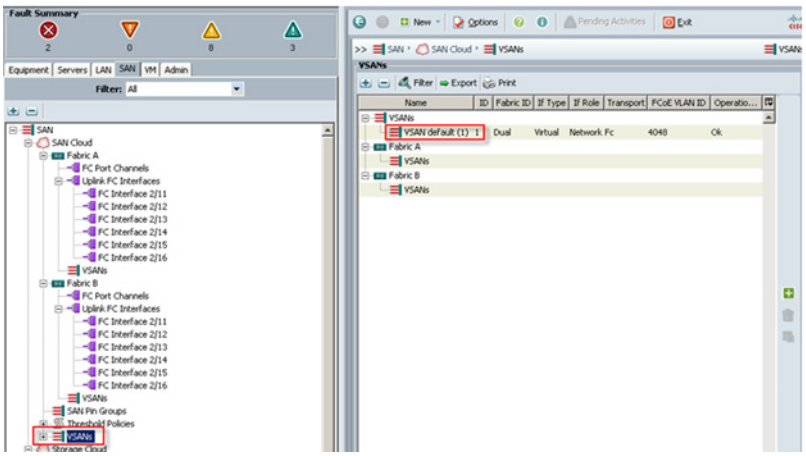

14. The next pool to create is the Server UUID pool. On the Servers tab in the Navigator page under the Pools node create a single UUID Pool for the test environment. Each Cisco UCS Blade Server requires a unique UUID to be assigned by its Service profile.

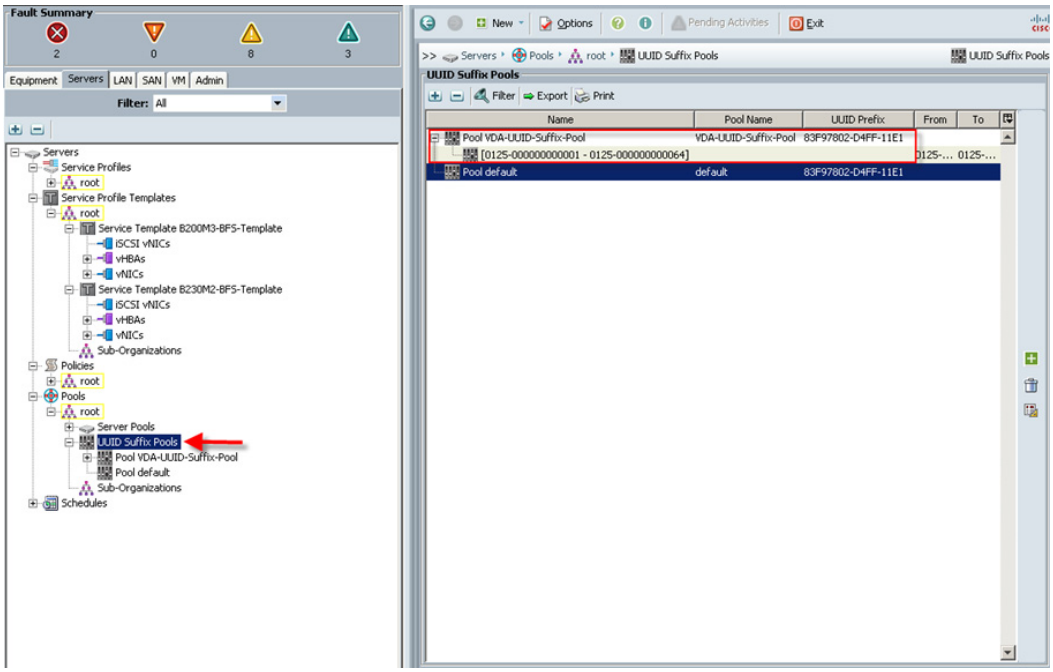

15. Create two Server Pools for use in the Service Profile Templates as selection criteria for automated profile association. Server Pools were created on the Servers tab in the navigation page under the Pools node. Only the pool name was created, no servers were added.

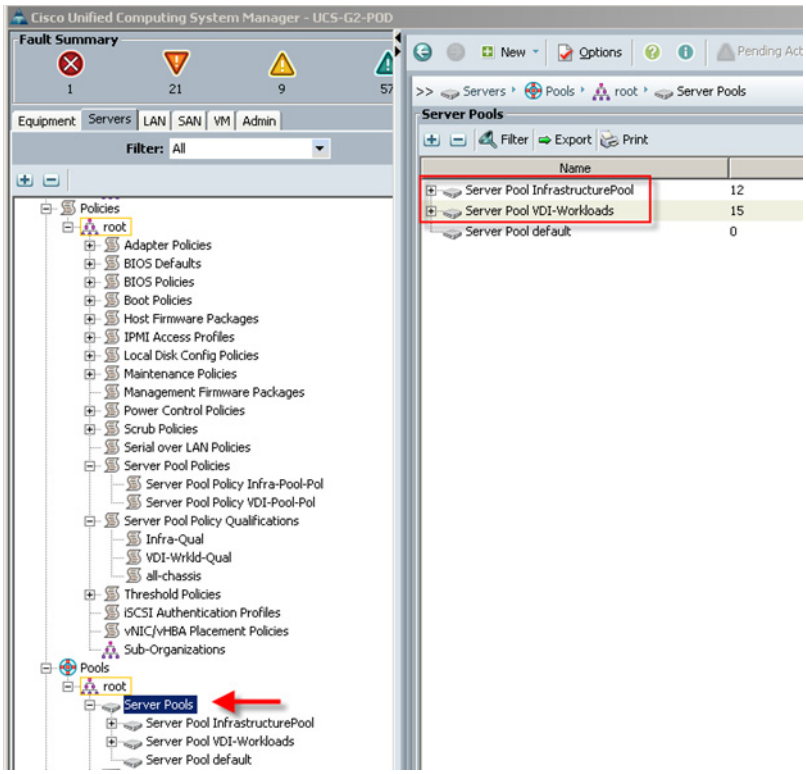

16. Create two Server Pool Policy Qualifications to identify the blade server model for placement into the correct pool using the Service Profile Template. In this case we used Chassis ids to select the servers. (You can use slots or server models to make the selection.)

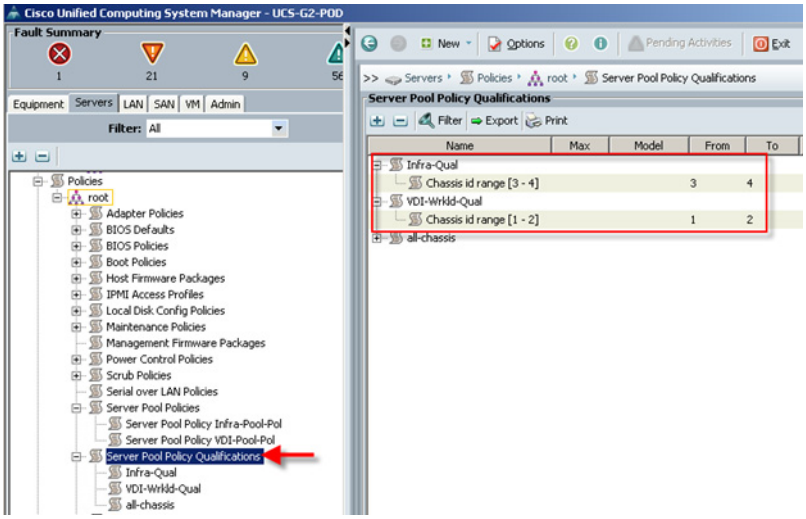

17. The next step in automating the server selection process is to create corresponding Server Pool Policies for each Cisco UCS Blade Server model, utilizing the Server Pool and Server Pool Policy Qualifications created earlier.

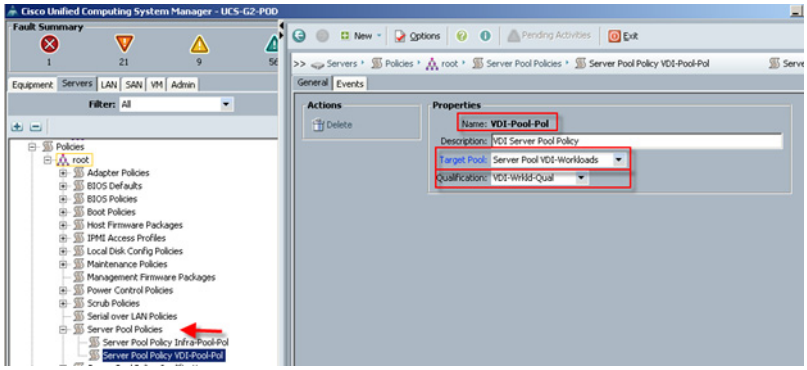

18. Virtual Host Bus Adapter templates were created for FC SAN connectivity from the SAN tab under the Polices node, one template for each fabric.

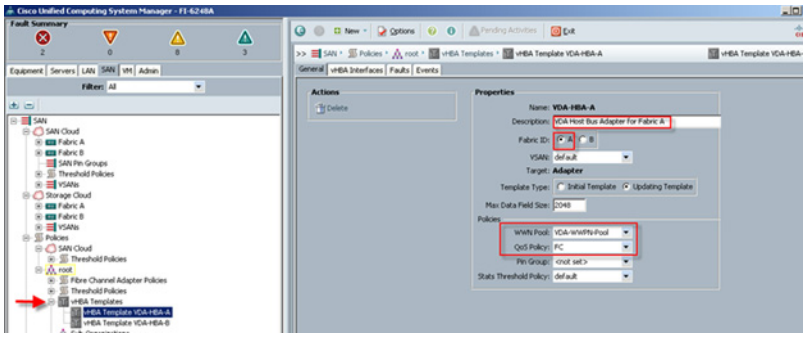

19. Create at least one HBA template for each Fabric Interconnect if block storage will be used. We used the WWPN pool created earlier and the QoS Policy.

20. From the LAN tab in the navigator pane, configure the VLANs for the environment.

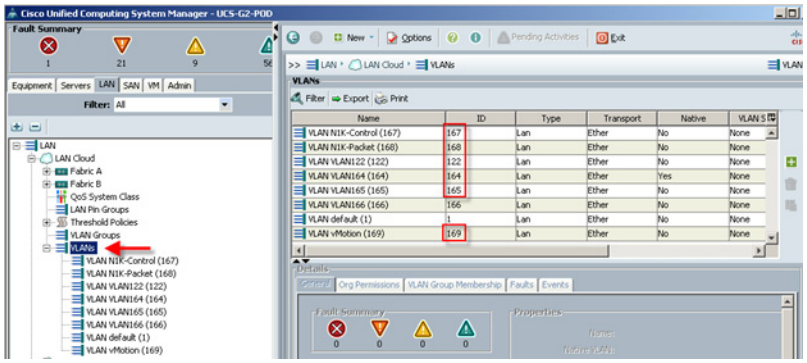

Note In this project we utilized six VLANs to accommodate our four ethernet system classes, a separate VLAN for infrastructure services, and two VLANs for Nexus 1000V packet and control functions. (N1KV management and VMware Management shared VLAN 164.) We did not use VLAN 166 in the FC variant we deployed in this study. However, if the NFS or iSCSI protocols were used, that VLAN is in place.

21. From the LAN tab in the navigator pane, under the policies node configure the vNIC templates that will be used in the Service Profiles. In this project, we utilize eight virtual NICs per host, four pairs, with each pair connected to both Fabric Interconnects for resiliency.

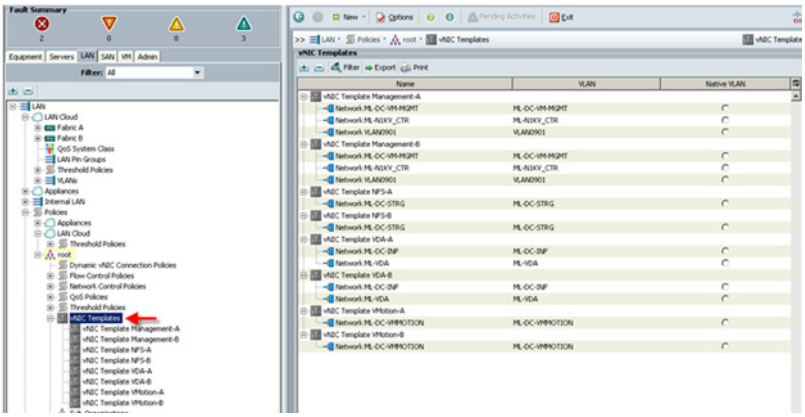

22. Create vNIC templates for both fabrics, check Enable Failover, select VLANs supported on adapter (optional,) set the MTU size, select the MAC Pool and QoS Policy, then click OK.

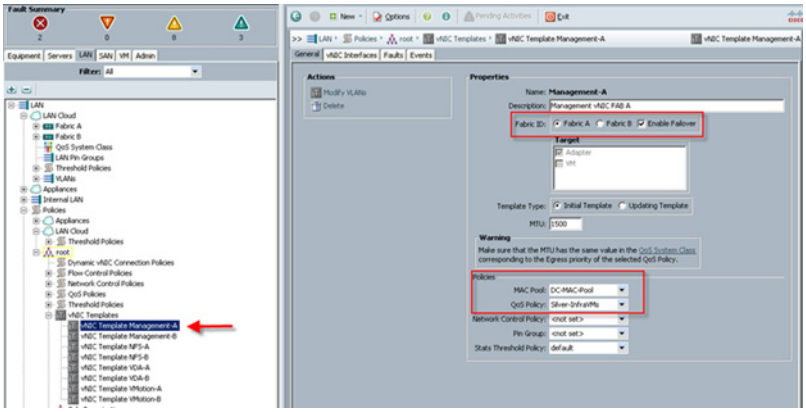

23. Create a boot from SAN policy to use for both Cisco UCS B250 M2 and Cisco UCS B200 M3 blades, using the WWNs from the VNX5500 storage system as SAN targets.

24. Create performance BIOS Policies for each blade type to insure optimal performance. The following screen captures show the settings for the Cisco UCS B200 M3 blades used in this study.

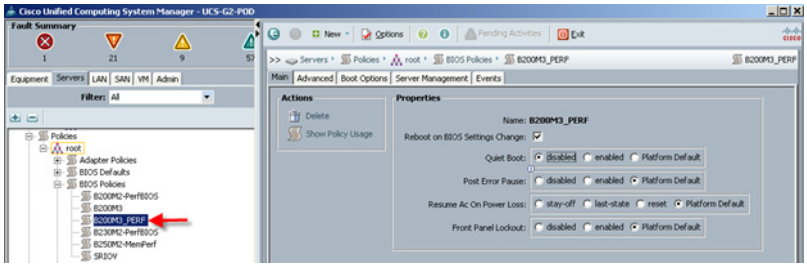

The Advanced Tab Settings:

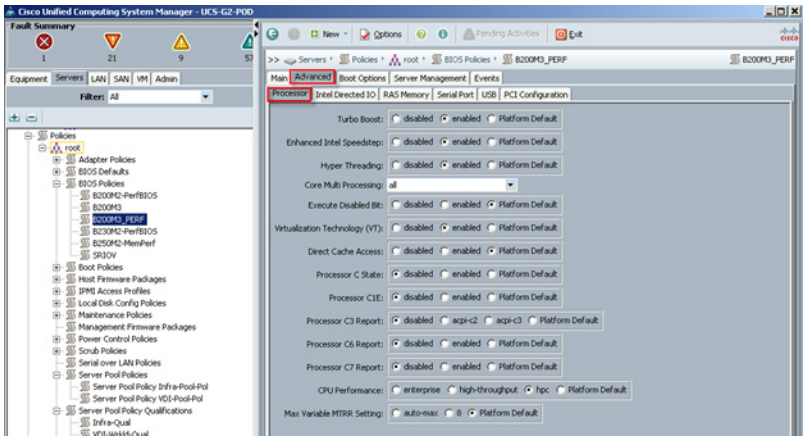

The remaining Advanced tab settings are at the platform default or not configured. Similarly, the Boot Options and Server Management tabs' settings are set at the defaults.

Note Be sure to Save Changes at the bottom of the page to preserve this setting. Also, be sure to add this policy to your blade service profile template.

25. New in Cisco UCS Manager 2.1(1a) is a way Host Firmware Package polices can be set by package version across the Cisco UCS domain rather than by server model .

Note You can still create specific packages for different models or for specific purposes.

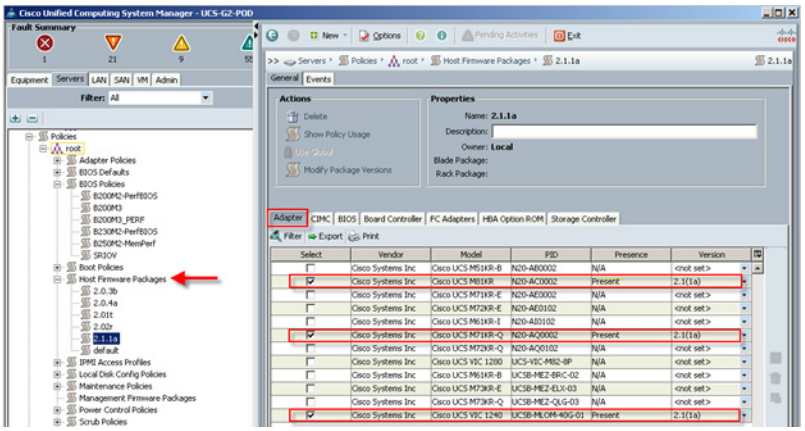

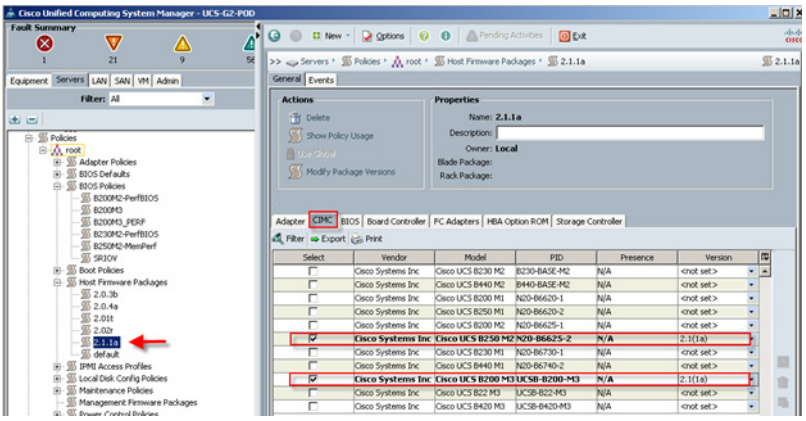

Note The process was continued across the remaining tabs choosing the appropriate model and versions.

26. Management Firmware Packages are no longer supported. Host Firmware packages replaced this functionality in Cisco UCS Manager 2.1.

27. Create a service profile template using the pools, templates, and policies configured above.

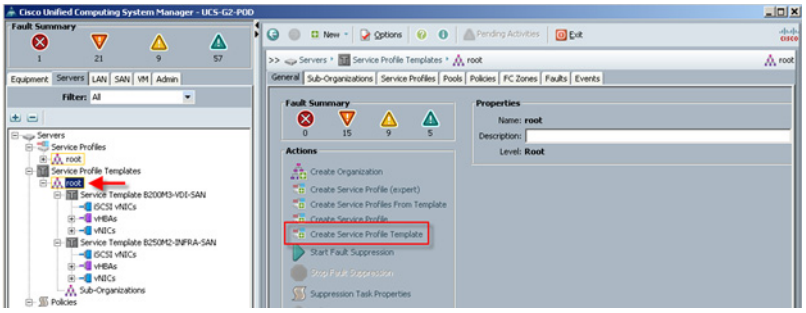

Note In this project, we created two templates, one for each of the Cisco UCS Blade Server models used.

Follow through each section, utilizing the policies and objects you created earlier, then click Finish.

Note On the Operational Policies screen, select the appropriate performance BIOS policy you created earlier to help ensure the maximum LV DIMM performance.

Note For automatic deployment of service profiles from your template(s), you must associate a server pool that contains blades with the template.

28. From the Create Service Profile Template wizard, enter a unique name, select the type as updating, and select the VDA-UUID-Suffix_Pool created earlier, then click Next.

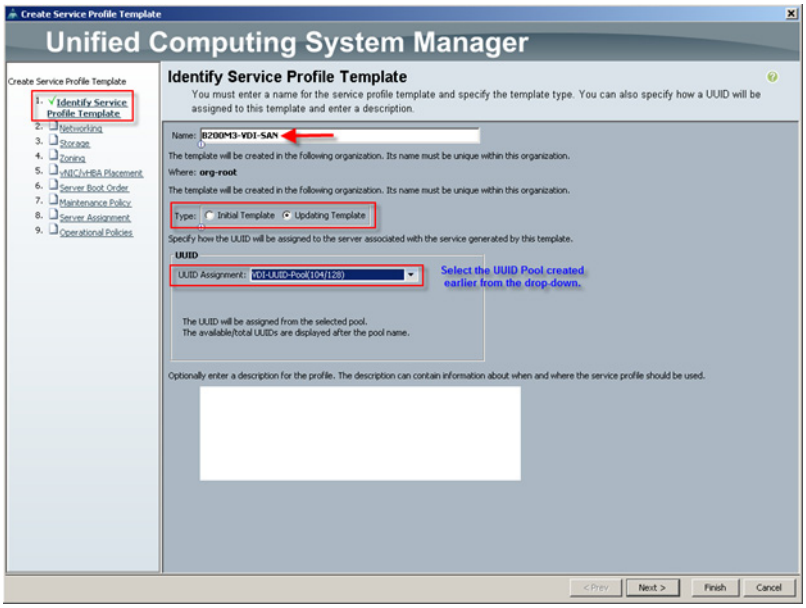

29. From the Storage page, select the Expert mode and select the WWNN Pool created earlier from the drop-down list and then click Add.

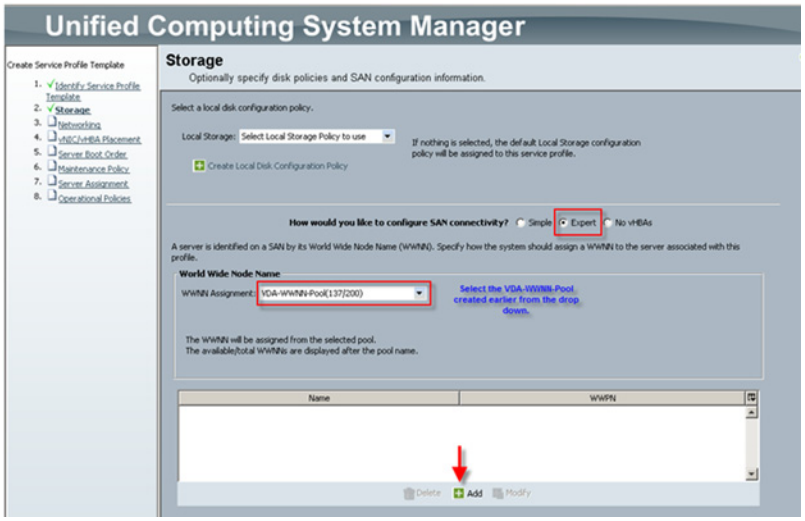

Note We used the default Local Storage configuration in this project. Local drives on the blades were not used.

30. From the Create HBA page, enter a name (FCO) and check Use SAN Connectivity Template, which will change the display as follows.

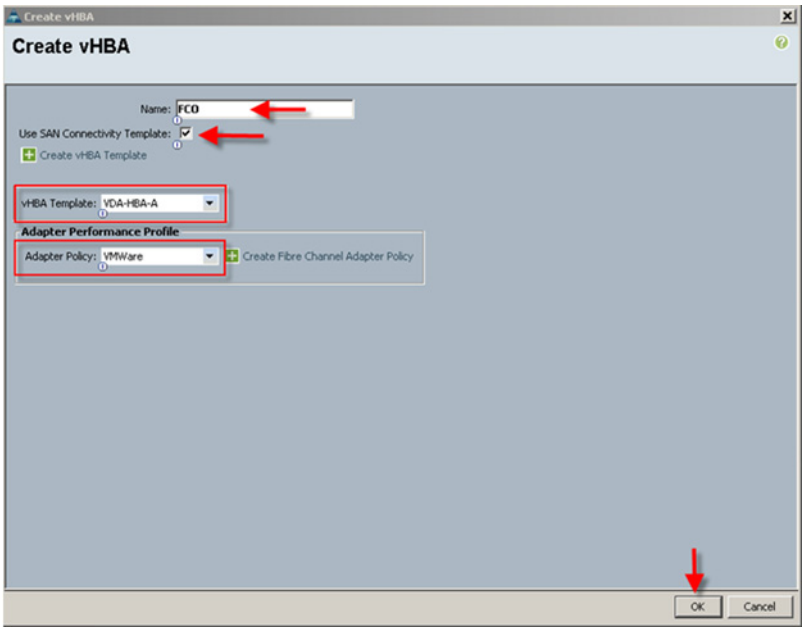

31. Select the vHBA template for Fabric Interconnect A and the VMware Adapter Policy from the drop-downs, then click OK.
32. Repeat this process for FC1, choosing VDA-HBA-B for Fabric Interconnect B. The result is shown in the Storage page that appears as follows:
33. Click Next.

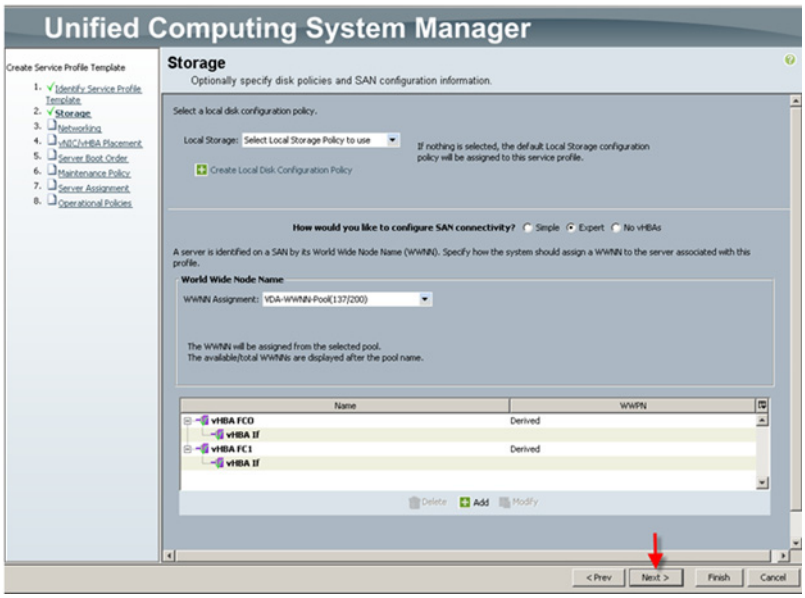

34. Click Next.
35. Select the Expert configuration option and click Add.

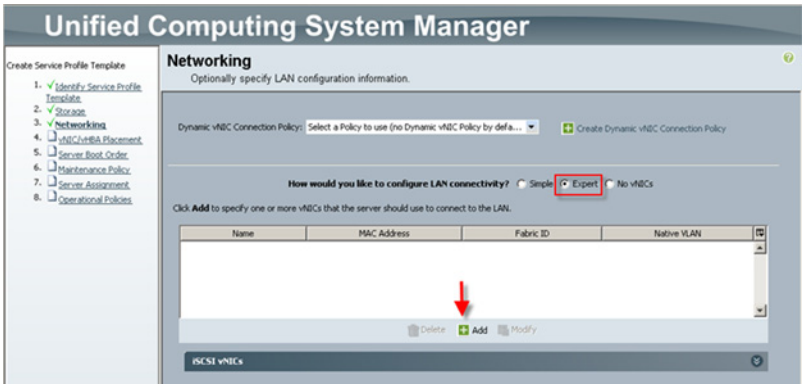

36. From the Create vNIC window, enter a unique Name, check the Use LAN Connectivity Template checkbox, select the vNIC Template from the drop-down, and repeat this process for the Adapter Policy.

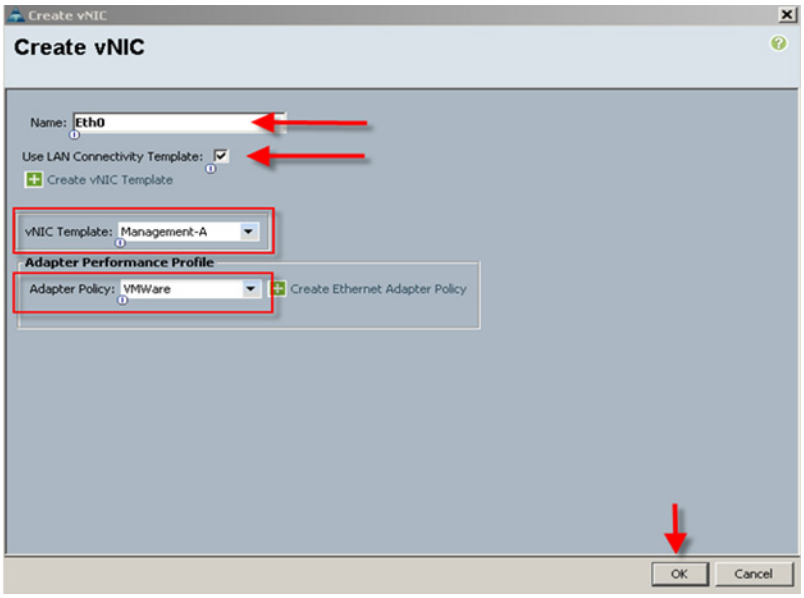

37. Repeat the process for the remaining seven vNICs, which will result in the following: (Eth5, 6, and 7 not shown).

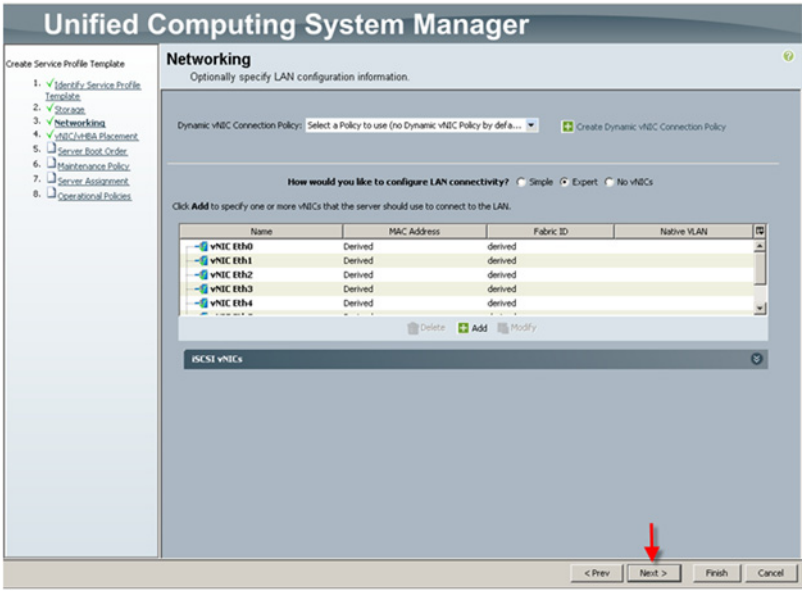

38. Click Next.
39. Accept the default placement and click Next.

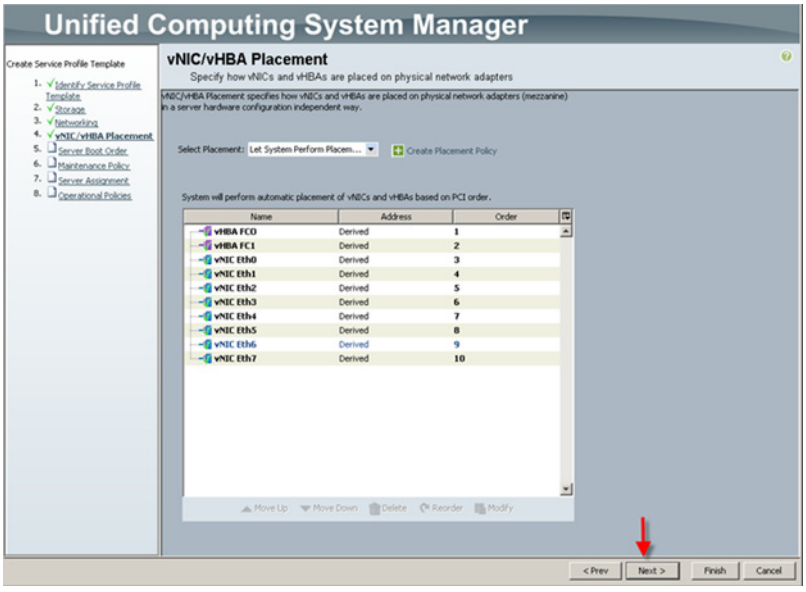

40. Select the Boot from SAN policy created in Section 5.4.5 from the drop-down.

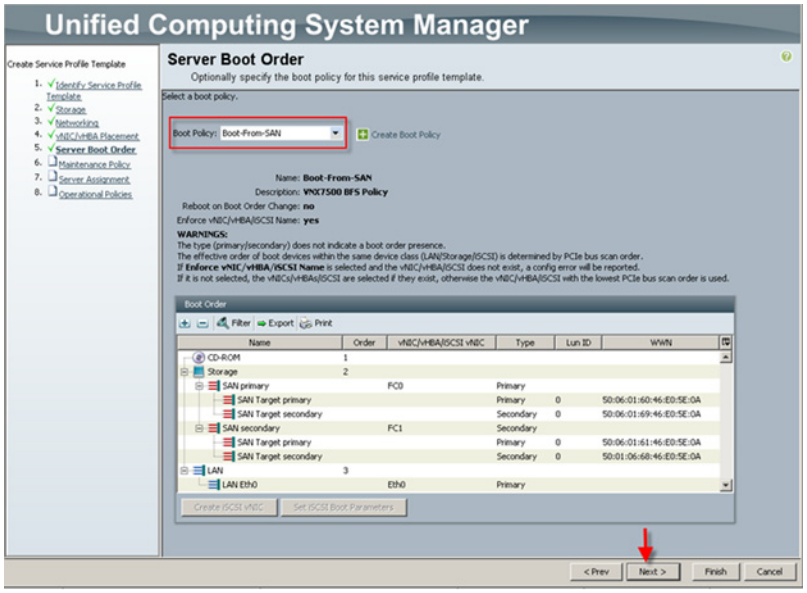

41. A Maintenance Policy for the project was not created, so click Next.

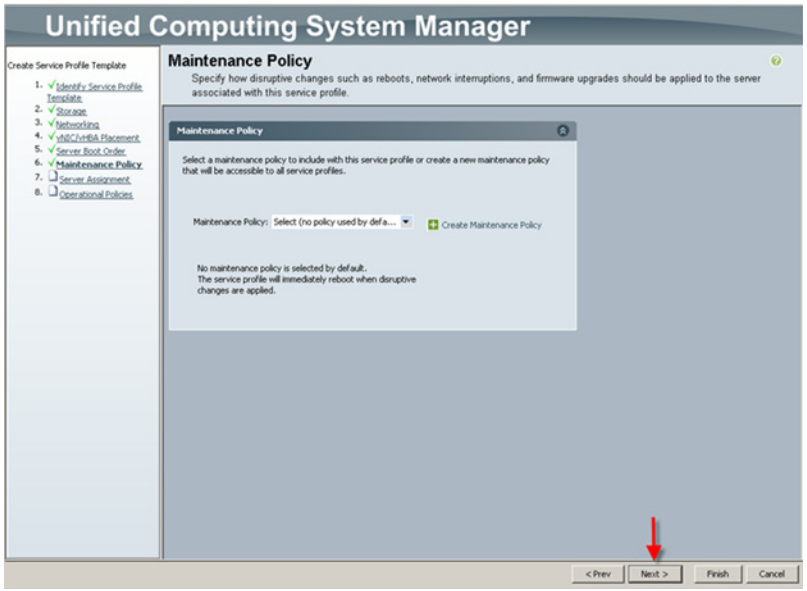

42. Make the following selections from the drop-downs as shown below, then click Next.

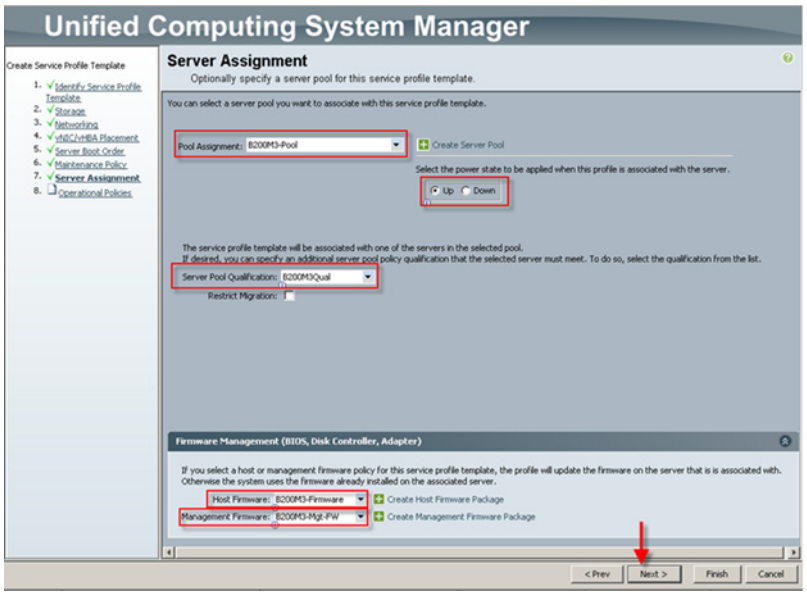

43. From the Operational Policies page, expand the BIOS Configuration section and select the BIOS Policy for the Cisco B200 M3 created earlier, then click Finish to complete the Service Profile Template.

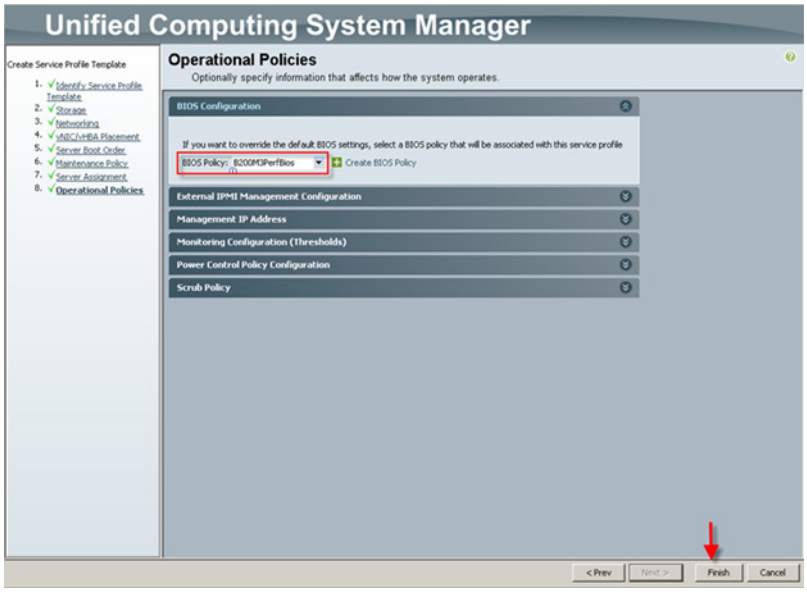

Note The result is a Service Profile Template for the Cisco UCS Blade Server B200 M3. We repeated the procedure to create a Service Profile Template for the Cisco UCS Blade Server B200 M3 used in the study.

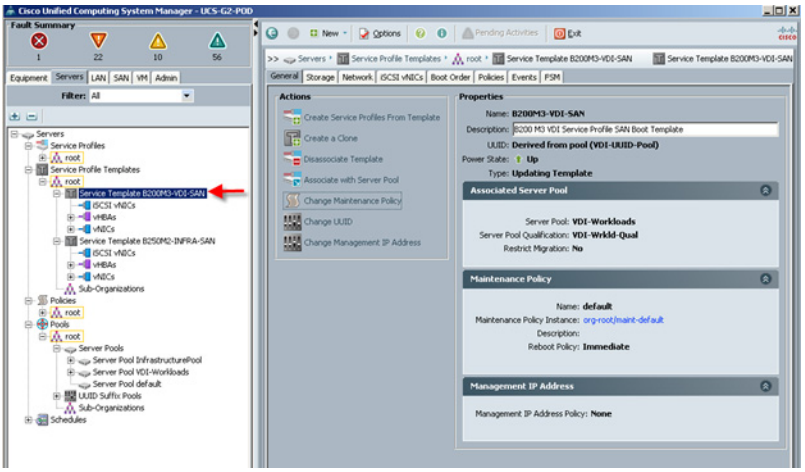

44. Use the newly created Service Profile Templates for each Cisco UCS Blade Server model for this project and create the appropriate number of Service Profiles. From the Servers tab in the navigation page, from the Service Profile Templates node, expand the root and select Service Template B200 M3, then click Create Service Profiles from Template in the right pane, Actions area.

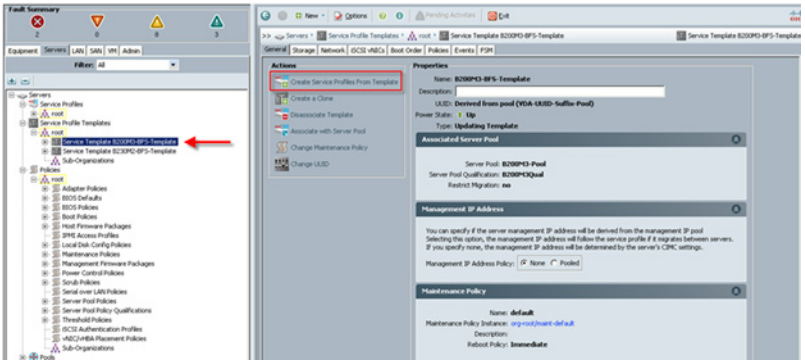

45. Provide the naming prefix and the number of Service Profiles to create and click OK.

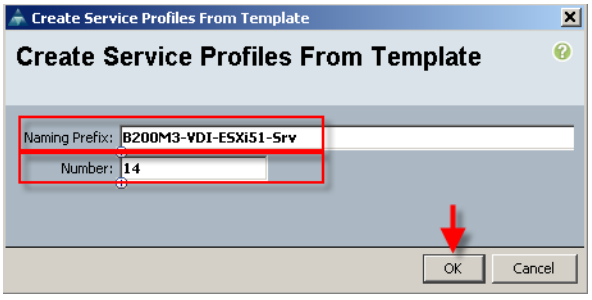

46. Cisco UCS Manager creates the requisite number of profiles and because of the Associated Server Pool and Server Pool Qualification policy, the Cisco UCS B200 M3 blades in the test environment automatically associate with the proper Service Profile.

Note This process was repeated for the Cisco UCS B250M2-INFRA-SAN template and the same result was achieved.

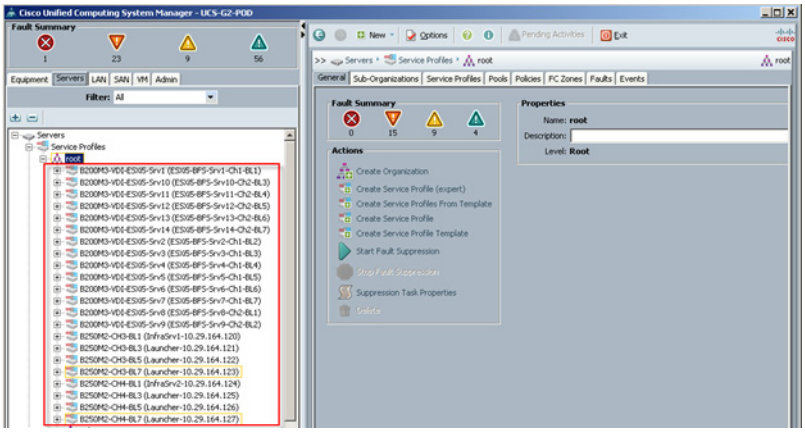

47. Verify that each server has a profile and that it receives the correct profile.

The image below is a Cisco UCS B200 M3 Sample:

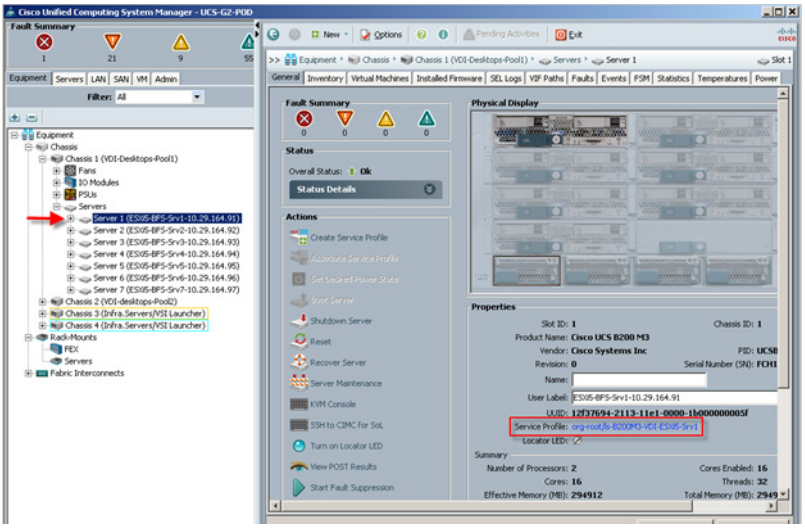

The Cisco UCS Blade Servers are ready for the hypervisor installation.

QoS and CoS in Cisco Unified Computing System

Cisco Unified Computing System provides different system class of service to implement quality of service including:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs

- Flow control policies that determine how uplink Ethernet ports handle pause frames.

Applications like the Cisco Unified Computing System and other time sensitive applications have to adhere to a strict QOS for optimal performance.

System Class Configuration

Systems Class is the global operation where entire system interfaces are with defined QoS rules.

- By default system has Best Effort Class and FCoE Class.

  Best effort is equivalent in MQC terminology as "match any"

  - FCoE is special Class define for FCoE traffic. In MQC terminology "match cos 3"
- System class allowed with 4 more users define class with following configurable rules.
  - CoS to Class Map
  - Weight: Bandwidth
  - Per class MTU
  - Property of Class (Drop v/s no drop)
- Max MTU per Class allowed is 9217.
- Through Cisco Unified Computing System you can map one CoS value to particular class.
- Apart from FcoE class there can be only one more class can be configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally system will calculate the bandwidth based on following equation (there will be rounding off the number).

$$\%\ \text{b/w shared of given Class} = \frac{(\text{Weight of the given priority} * 100)}{\text{Sum of weights of all priority}}$$

Cisco UCS Class Configuration

- Platinum
- Gold
- Silver
- Bronze

Table 3 Name Table Map Between Cisco UCS and the NXOS

Cisco UCS Names	NXOS Names
Best effort	Class-default
FC	Class-fc
Platinum	Class-Platinum
Gold	Class-Gold
Silver	Class-Silver
Bronze	Class-Bronze

Table 4 Class to CoS Map by default in Cisco Unified Computing System

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Fc	3
Platinum	5
Gold	4
Silver	2
Bronze	1

Table 5 Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
Fc	5

Steps to Enable QOS on Cisco Unified Computing System

For this study, we utilized four Cisco UCS QoS System Classes to priorities four types of traffic in the infrastructure.

Table 6 QoS Priority to vNIC and VLAN Mapping

Cisco UCS Qos Priority	vNIC Assignment	VLAN Supported
Platinum	eth2, eth3	166 (Storage – Not used in FC variant)
Gold	eth4, eth5	122 (VDA)
Silver	eth0, eth1	164 (Management)
Bronze	eth6, eth7	169 (vMotion)

Configure Platinum, Gold, Silver and Bronze policies by checking the enabled box. The Platinum Policy, used for NFS storage, was configured for Jumbo Frames in the MTU column. Notice the option to set no packet drop policy during this configuration.

Figure 10 Cisco UCS QoS System Class Configuration

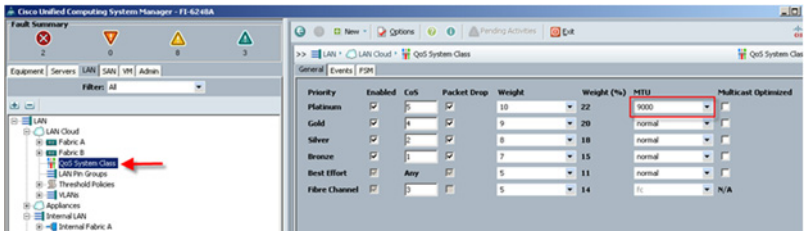

Next, from the LAN tab under Policies, Root, QoS Polices, verify QoS Policies Platinum, Gold, Silver and Bronze exist, with each QoS policy mapped to its corresponding Priority.

Figure 11 Cisco UCS QoS Policy Configuration

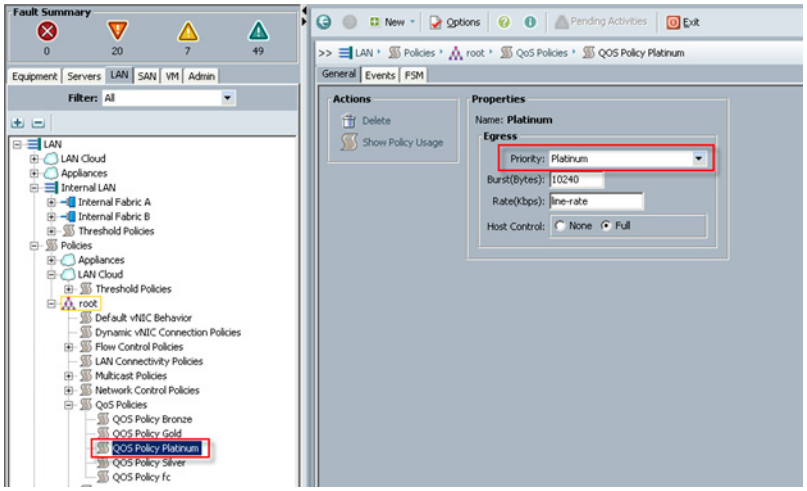

Finally, include the corresponding QoS Policy into each vNIC template using the QoS policy drop down, using the QoS Priority to vNIC and VLAN Mapping table above.

Figure 12 Utilize QoS Policy in vNIC Template

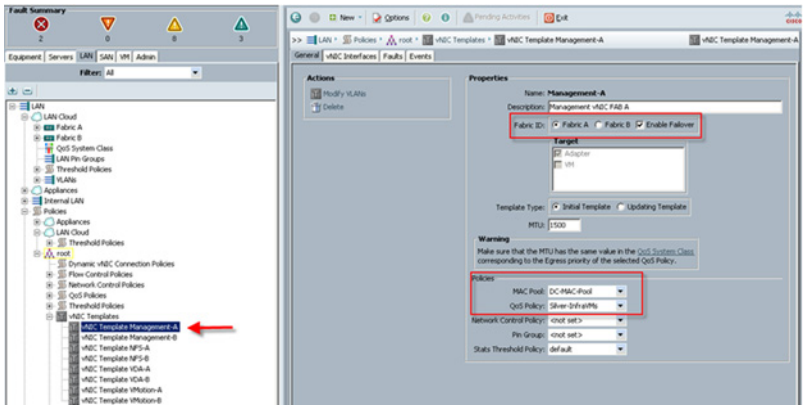

This is a unique value proposition for Cisco Unified Computing System with respect to end-to-end QOS. For example, there is a VLAN for the EMC storage; configure the Platinum policy with Jumbo frames and receive an end-to-end QOS and performance guarantees from the Blade Servers to the Nexus 1000V virtual distributed switches running in vCenter through the Nexus 5548UP access layer switches.

LAN Configuration

The access layer LAN configuration consists of a pair of Cisco Nexus 5548s (N5Ks,) a family member of our low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches for our VDI deployment.

Cisco UCS Connectivity

Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS 6248 fabric interconnects, and they are connected to the Cisco Nexus 5548 pair in a bow tie manner as shown below in a port channel.

The 6248 Fabric Interconnect is in End host mode, as we are doing both Fiber Channel as well as Ethernet (NAS) data access as per the recommended best practice of the Cisco Unified Computing System. We built this out for scale and have provisioned more than 40 G per Fabric interconnect (Figure 13).

Note The upstream configuration is beyond the scope of this document; there are some good reference document [4] that talks about best practices of using the Cisco Nexus 5000 and 7000 Series Switches. New with the Nexus 5500 series is an available Layer 3 module that was not used in these tests and that will not be covered in this document.

Figure 13 Ethernet Network Configuration with Upstream Cisco Nexus 5500 Series from the Cisco Unified Computing System 6200 Series Fabric Interconnects

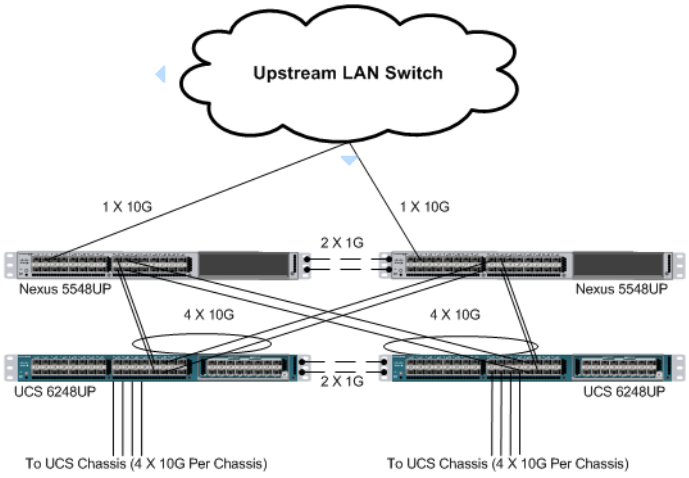

EMC VNX5500 LAN Connectivity

The Cisco Nexus 5548UP is used to connect to the EMC VNX5500 storage system for Fiber Channel and file-based access.

The VNX5500 is equipped with dual-port 8GB FC modules on each controller. These are connected to the pair of Nexus 5548 unified ports to provide block storage access to the environment. (See section SAN Configuration)

The VNX5500 supports two dual-port 10G Data Movers which are connected to the pair of N5Ks downstream. One of the Data Movers is set to Active, with the second providing failover capability. This allows end-to-end 10G access for file-based storage traffic. We have implemented jumbo frames on the ports and have priority flow control on, with Platinum CoS and QoS assigned to the vNICs carrying the storage data access on the Fabric Interconnects. (Note: This configuration was not used in this study, but is shown as a supported option.)

The EMC ethernet connectivity diagram is shown below. There is a total of 40 Gbps bandwidth available for the servers.

Figure 14 EMC VNX Ethernet Connectivity

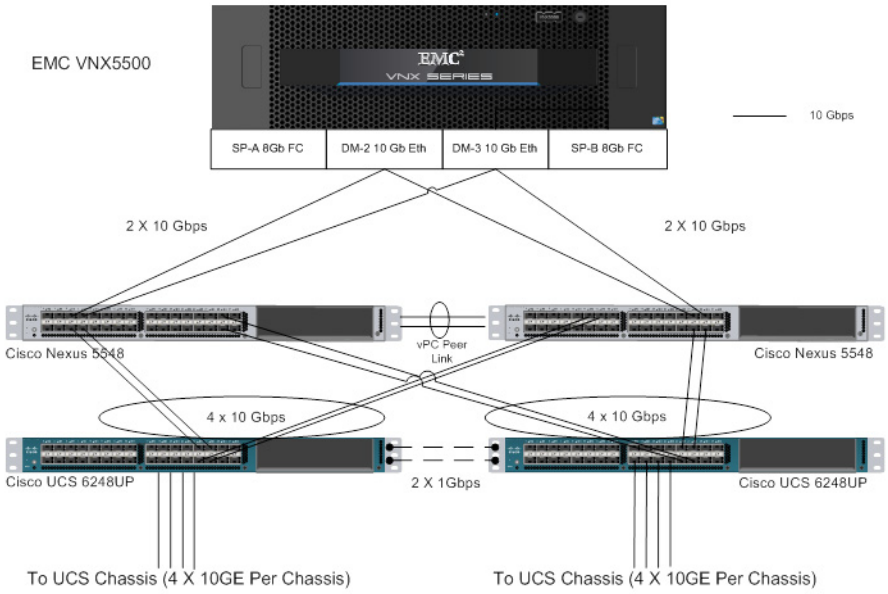

Figure 15 EMC VNX Fibre Channel Connectivity

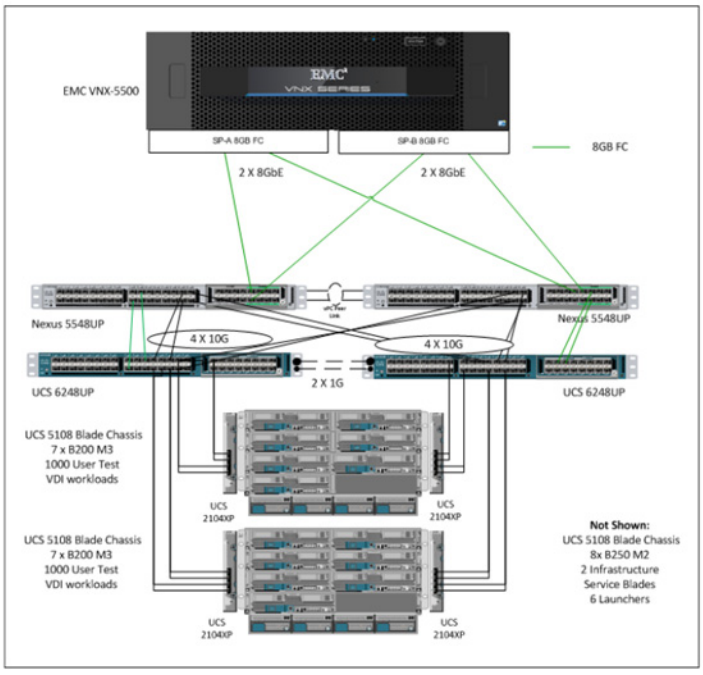

For information on configuring ethernet connectivity on a EMC VNX5500 Storage System, refer to the EMC website: www.emc.com.

Nexus 1000V Configuration in L3 Mode

1. Download the Nexus1000 V 4.2(1) SV1 (5.2): http://www.cisco.com/cisco/software/release.html?mdfid=282646785&flowid=3090&softwareid=282088129&release=4.2(1)SV1(5.2)&relind=AVAILABLE&rellifecycle=&reltype=latest
2. Extract the downloaded N1000V .zip file on the Windows host.
3. To start the N1000V installation, run the command below from the command prompt. (Make sure the Windows host has the latest Java version installed)

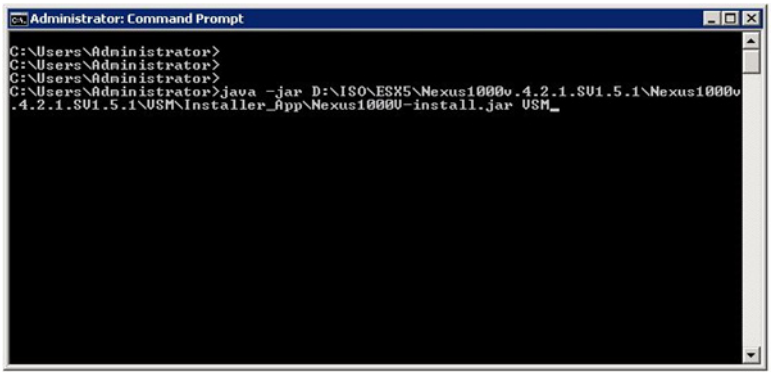

4. After running the installation command, you will see the "Nexus 1000V Installation Management Center."

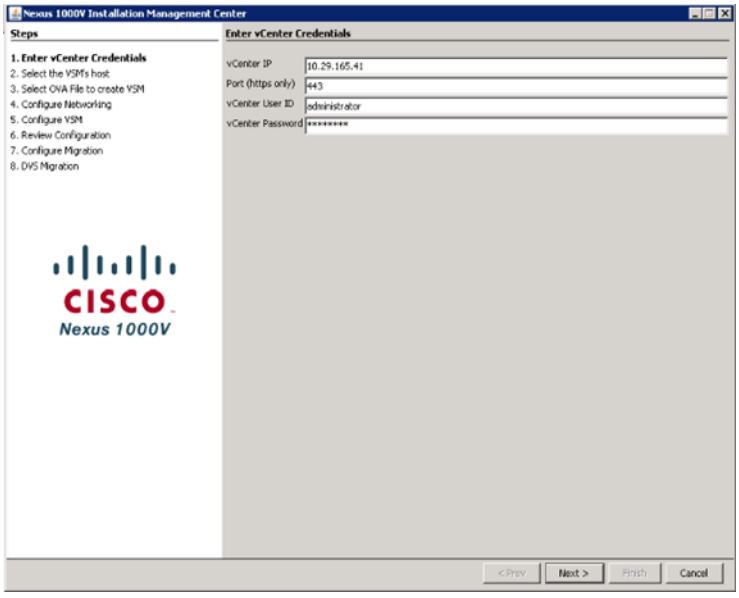

5. Enter the vCenter IP and the logon credentials.

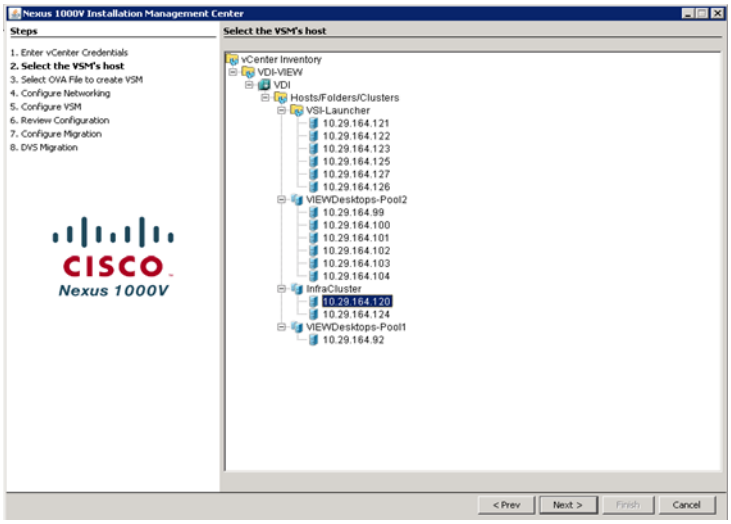

6. Select the ESX host on which to install N1KV Virtual Switch Manager.
7. Select the OVA file from the extracted N1KV location to create the VSM.

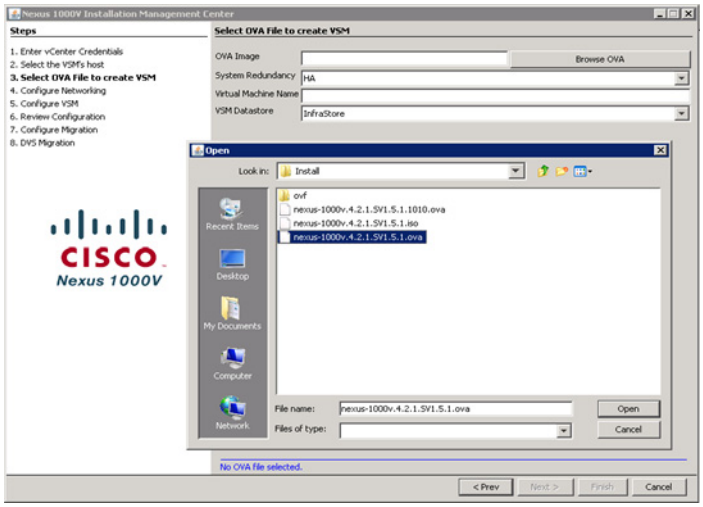

8. Select the System Redundancy type as "HA" and enter the virtual machine name for the N1KV VSM and choose the Datastore for the VSM.

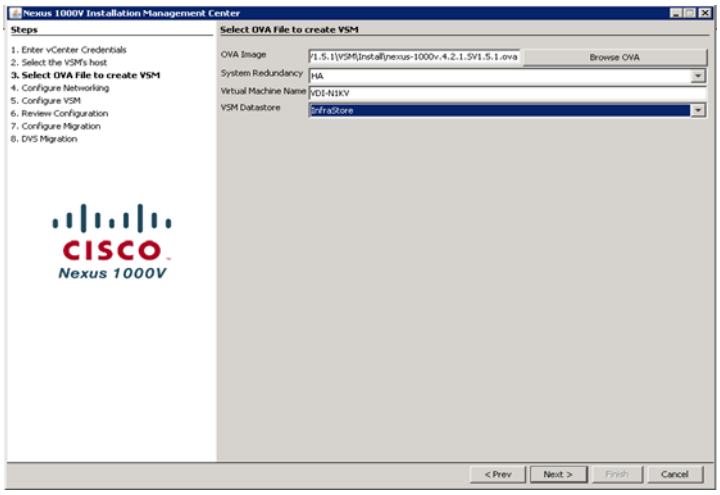

9. To configure L3 mode of installation, select the "L3 : Configure port groups for L3"
   a. Create Port-group as Control and specify the VLAN ID and select the corresponding vSwitch.
   b. Select the existing port group "VM Network" for N1K Mgmt and choose mgmt0 with the VLAN ID for the SVS connection between vCenter and VSM.
   c. In the option for L3 mgmt0 interface port-profile enter the vlan that was pre-defined for ESXi mgmt and accordingly it will create a port-group which will have L3 capability. In this case it is n1kv-L3 port-group as shown in the screenshot below.

```
VDI-N1KV-DVS# sh run port-profile n1kv-L3

!Command: show running-config port-profile n1kv-L3
!Time: Fri Oct 26 16:53:34 2012

version 4.2(1)SV1(5.2)
port-profile type vethernet n1kv-L3
  capability l3control
  vmware port-group
  switchport mode access
  switchport access vlan 164
  no shutdown
  system vlan 164
  state enabled
```

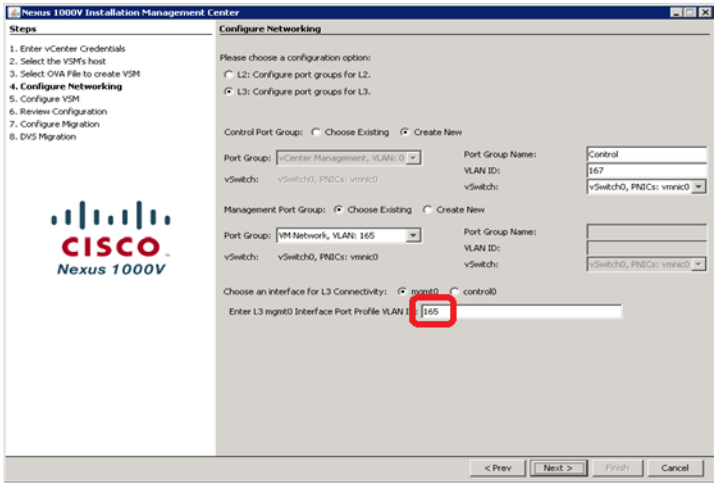

10. To Configure VSM, Type the Switch Name and enter the admin password for the VSM. Type the IP address, subnet mask, Gateway, Domain ID (Note: If there are multiple instance of N1KV VSM need to be install, make sure that each are configured with different Domain ID) and select the SVS datacenter Name and Type the vSwitch0 Native vlan ID. (Make sure the Native VLAN ID specified should match the Native VLAN ID of Cisco UCS and the Nexus 5k)

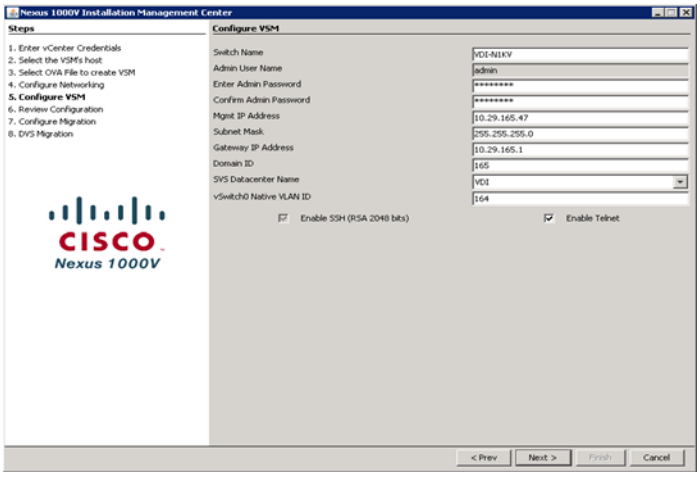

11. Review the configuration and click Next to proceed with the installation.

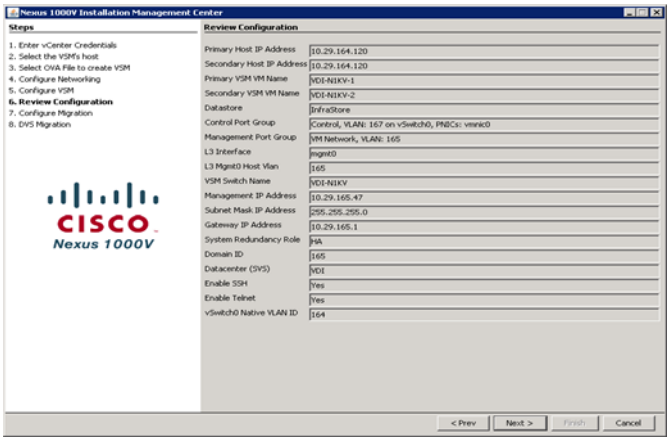

12. Wait for the Completion of Nexus 1000V VSM installation.

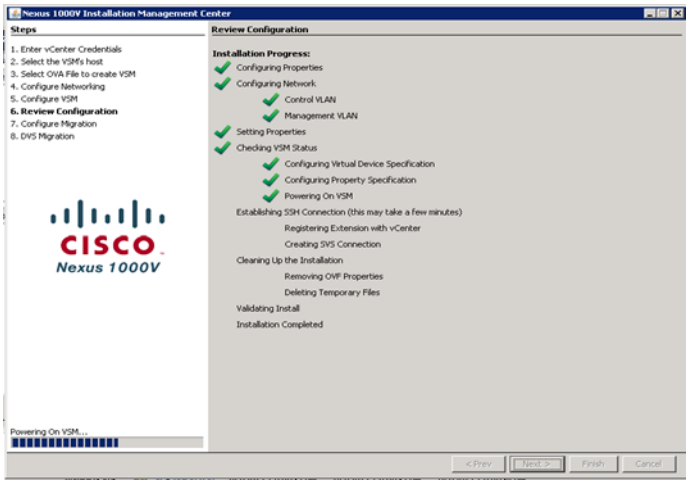

13. Click Finish to complete the VSM installation.

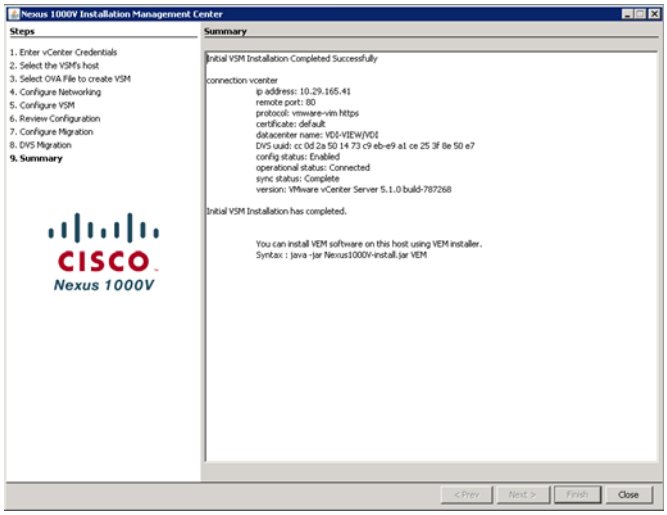

14. Log into (ssh or telnet) to the N1KV VSM with the IP address and configure VLAN for ESX Mgmt, Control, N1K Mgmt and also for Storage and vMotion purposes as mentioned below (VLAN ID differs based on your Network). No need to create VLANs for N1KV packet and control with version 4.2(1)SV2(1.1) of N1KV installer.

```
VDI-N1KV# conf t
```

15. Enter the following configuration commands, one per line, an d end with CNTL/Z:

```
VDI-N1KV(config)# vlan  122
VDI-N1KV(config-vlan)# name  VDA
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan  164
VDI-N1KV(config-vlan)# name  ESXi-Mgmt
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan  165
VDI-N1KV(config-vlan)# name  Infra-Mgmt
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan  166
VDI-N1KV(config-vlan)# name  Storage
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan  167VDI-N1KV(config-vlan)# name  N1K-Control
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan  169
VDI-N1KV(config-vlan)# name  vMotion
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)#<CNTRL/Z>
```

```
vlan 122
  name VDI-desktops
vlan 164
  name ESX_Mgmt
vlan 165
  name Infra_Mgmt
vlan 166
  name Storage
vlan 167
  name Control
vlan 169
  name vMotion
```

16. Run the following configuration command to configure jumbo mtu and qos polices:

```
VDI-N1KV# conf t
```

```
VDI-N1KV(config)# policy-map type qos jumbo-mtu
VDI-N1KV(config-pmap-qos)# policy-map type qos platinum_Cos_5
VDI-N1KV(config-pmap-qos)# class class-default
VDI-N1KV(config-pmap-c-qos)# set cos 5
VDI-N1KV(config-pmap-c-qos)# end
VDI-N1KV# copy running-config startup-config
```

```
policy-map type qos jumbo-mtu
policy-map type qos platinum_Cos_5
  class class-default
    set cos 5
```

17. To migrate and manage all the ESXi you will need to network using Nexus 1000V VSM, Configure Port Profiles and port groups as shown below.

```
port-profile type ethernet Unused_Or_Quarantine_Uplink
  vmware port-group
  shutdown
  description Port-group created for Nexus1000V internal usage.
Do not use.
  state enabled
port-profile type vethernet Unused_Or_Quarantine_Veth
  vmware port-group
  shutdown
  description Port-group created for Nexus1000V internal usage.
Do not use.
  state enabled
```

Note Do not make any changes to the port profiles; they are created by default.

18. Create the DC System Uplink for ESXi and Nexus 1000V Management:

```
VDI-N1KV(config)# port-profile type ethernet Mgmt-uplink
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode trunk
VDI-N1KV(config-port-prof)# switchport trunk allowed vlan  164-165,167
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no shutdown
VDI-N1KV(config-port-prof)# system vlan  164,167
VDI-N1KV(config-port-prof)#state enabled
```

19. Create the DC Storage Uplink port profile for NFS traffic:

```
VDI-N1KV(config)# port-profile type ethernet  storage-uplink
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  166
VDI-N1KV(config-port-prof)# mtu 9000
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no shutdown
VDI-N1KV(config-port-prof)# system vlan  166
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet storage-uplink
  vmware port-group
  mtu 9000
  switchport mode access
  switchport access vlan 166
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 166
  state enabled
```

20. Create the Storage virtual ethernet communications port profile:

```
VDI-N1KV(config)# port-profile type vethernet storage
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  166
VDI-N1KV(config-port-prof)# service -policy type qos input platinum_Cos_5
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan  166
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type vethernet storage
  vmware port-group
  switchport mode access
  switchport access vlan 166
  service-policy type qos input platinum_Cos_5
  no shutdown
  system vlan 166
  state enabled
```

21. Create the DC vMotion Uplink port profile:

```
VDI-N1KV(config)# port-profile type ethernet  vmotion-uplink
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  169
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan  169
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet vmotion-uplink
  vmware port-group
  switchport mode access
  switchport access vlan 169
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 169
  state enabled
```

22. Create the virtual ethernet port profile for vMotion:

```
VDI-N1KV(config)# port-profile type vethernet vmotion
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  169
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan  169
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type vethernet vmotion
  vmware port-group
  switchport mode access
  switchport access vlan 169
  no shutdown
  system vlan 169
  state enabled
```

23. Create the DC VDA Uplink port profile:

```
VDI-N1KV(config)# port-profile type ethernet  vdi-uplink
VDI-N1KV(config-port-prof)# vmware port-group
    VDI-N1KV(config-port-prof)#switchport mode access
    VDI-N1KV(config-port-prof)#switchport access vlan  122,164-165
VDI-N1KV(config-port-prof)#channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)#no shutdown
VDI-N1KV(config-port-prof)# system vlan 122,165
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet vdi-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 122,164-165
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 122,165
  state enabled
```

24. Create the virtual ethernet port profile for VDA traffic:

```
VDI-N1KV(config)# port-profile type vethernet  vdi-pool1
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# port-binding static auto expand
VDI-N1KV(config-port-prof)# max-ports 1024
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  122
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan  122
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type vethernet vdi-pool1
  vmware port-group
  port-binding static auto expand
  switchport mode access
  switchport access vlan 122
  no shutdown
  system vlan 122
  max-ports 1024
  state enabled
```

25. Create the virtual ethernet port profile for VDA1 traffic:

```
VDI-N1KV(config)# port-profile type vethernet  vdi-pool2
VDI-N1KV (config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# port-binding static auto expandVDI-N1KV
(config-port-prof)# max-ports 1024
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan  122
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV (config-port-prof)# system vlan  122
VDI-N1KV(config-port-prof)#state enable
```

```
port-profile type vethernet vdi-pool2
  vmware port-group
  switchport mode access
  switchport access vlan 122
  no shutdown
  system vlan 122
  max-ports 1024
  state enabled
```

26. After creating port profiles, make sure vCenter shows all the port profiles and port groups under the respective N1KV VSM. Then, Add the ESXi host to the VSM.
27. Go to Inventory> networking > select DVS for N1KV> click Hosts.

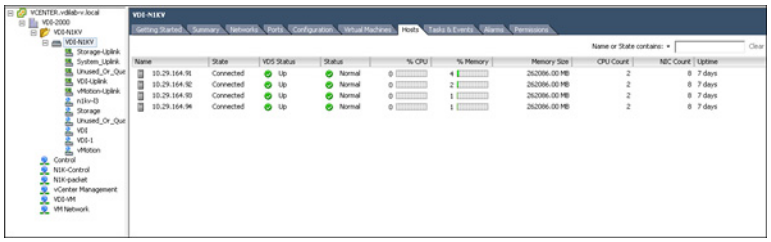

28. Right-click and select Add host to vSphere Distributed Switch.

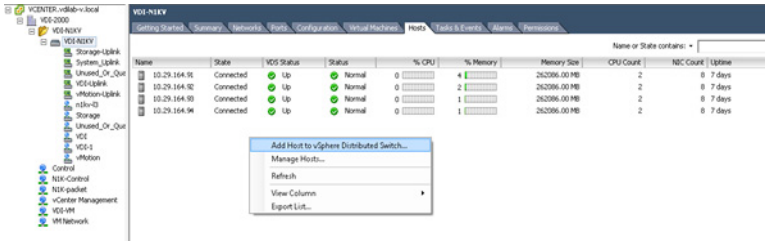

Note This brings up the ESXi hosts which are not part of the existing configuration.

29. Select ESXi hosts to add in N1KV.

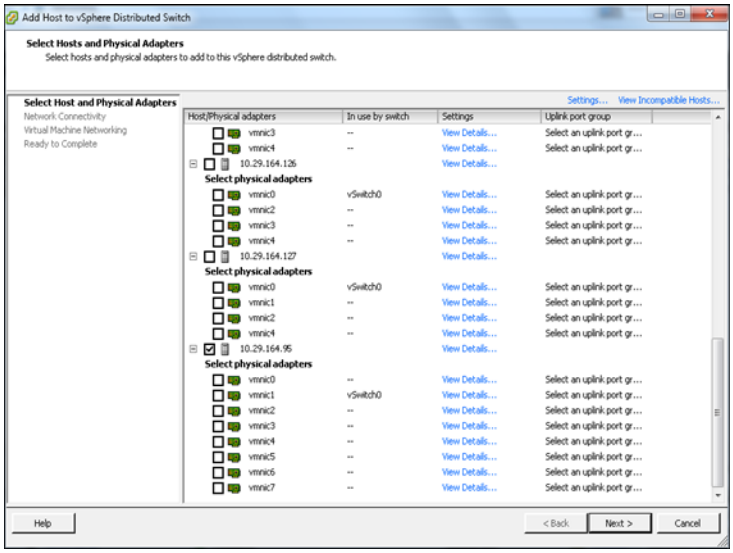

30. Click Select an uplink port-group and from the drop-down menu select appropriate Uplink that is allowed for corresponding vmnic as per the configuration on Cisco UCS Manager vNICS.

    For example, consider vmnic0 and vmnic1 for use as the System-Uplink. As per the best practices here we have 8 vmnics (4 pairs) and each pair of vmnics will be associated with one uplink; system/mgmt uplink, storage uplink, VM/VDI traffic uplink, vMotion Uplink.

31. After selecting appropriate uplinks click Next.

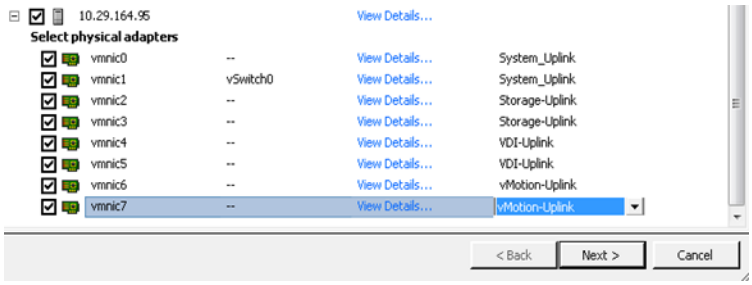

32. From the Network Connectivity tab, select Destination port group for vmk0.

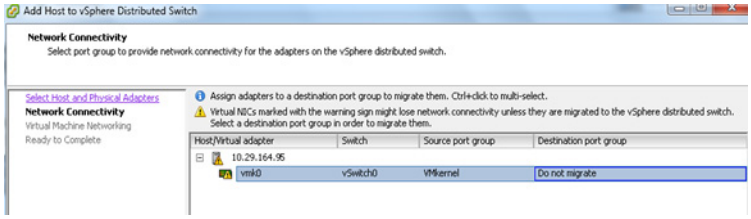

33. From the drop-down menu select a port group that was configured for L3 capability and for ESXi host management communication. In this case it is n1kv-l3 and click Next.

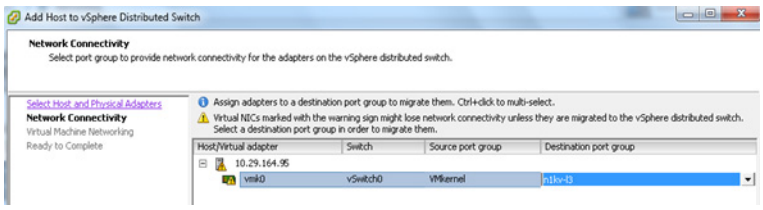

34. On the tab for virtual machine networking, select VMs and assign them to a destination port-group if there are any. Otherwise click Next.

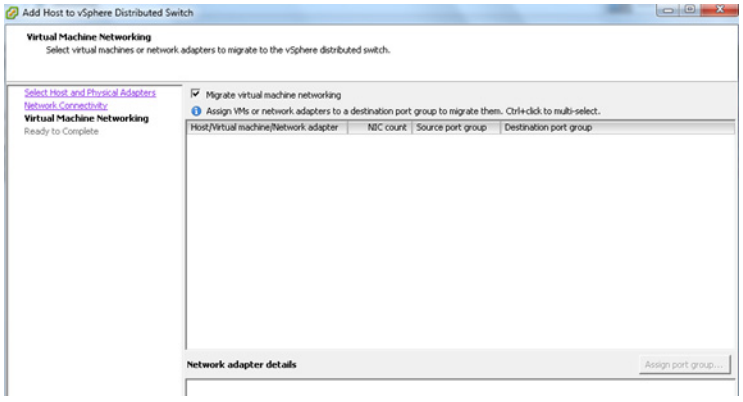

35. Verify the Settings and click Finish to add the ESXi host part of N1KV DVS.

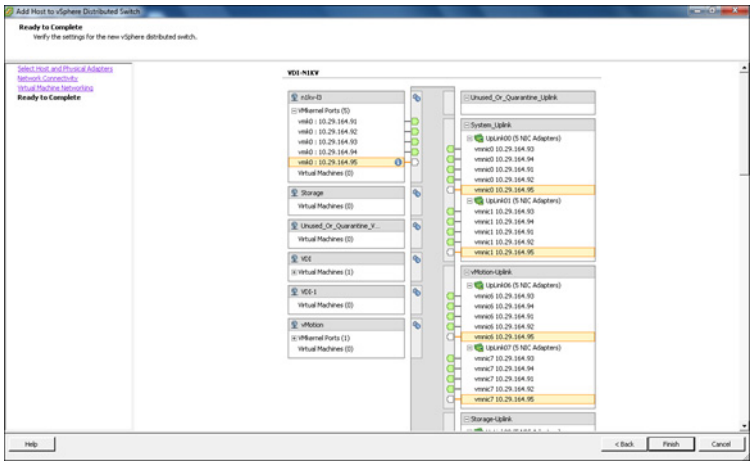

Note This will invoke VMware update manger (VUM) to automatically push the VEM installation for the selected ESXi hosts. After successful staging, install and remediation process , now the ESXi host will be added to N1KV VSM. From the vCenter task manager, quickly check the process of VEM installation.

In the absence of Update manager:

1. Upload vib file cross_cisco-vem-v144-4.2.1.1.5.2.0-3.0.1.vib for VEM installation to local or remote datastore which can be obtained by browsing to the management IP address for N1KV VSM.
2. Login to ESXi host using ESXi shell or SSH session.
3. Run the following command:

```
esxcli software vib install -v /vmfs/volumes/
datastore/ cross_cisco-vem-v144-4.2.1.1.5.2.0-3.0.1.vib
```

4. Verify the successful installation of ESXi VEM and the status of ESXi host.

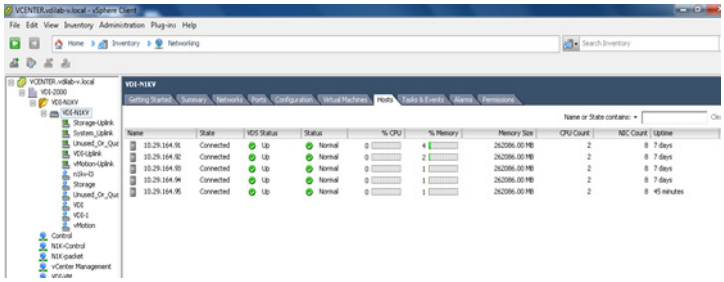

5. Verify putty into N1KV VSM. Run sh module command which will show all the ESXi hosts attached to that VSM.

```
VDI-N1KV(config)# sh module
```

```
VDI-N1KV(config)# sh module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------  -----------------  ------------
1    0      Virtual Supervisor Module        Nexus1000V         active *
2    0      Virtual Supervisor Module        Nexus1000V         ha-standby
16   248    Virtual Ethernet Module          NA                 ok
17   248    Virtual Ethernet Module          NA                 ok
18   248    Virtual Ethernet Module          NA                 ok
19   248    Virtual Ethernet Module          NA                 ok
20   248    Virtual Ethernet Module          NA                 ok

Mod  Sw                Hw
---  ----------------  ------------------------------------------------
1    4.2(1)SV1(5.2)    0.0
2    4.2(1)SV1(5.2)    0.0
16   4.2(1)SV1(5.2)    VMware ESXi 5.1.0 Releasebuild-799733 (3.1)
17   4.2(1)SV1(5.2)    VMware ESXi 5.1.0 Releasebuild-799733 (3.1)
18   4.2(1)SV1(5.2)    VMware ESXi 5.1.0 Releasebuild-799733 (3.1)
19   4.2(1)SV1(5.2)    VMware ESXi 5.1.0 Releasebuild-799733 (3.1)
20   4.2(1)SV1(5.2)    VMware ESXi 5.1.0 Releasebuild-799733 (3.1)
```

```
Mod  Server-IP        Server-UUID                           Server-Name
---  ---------------  ------------------------------------  -------------------
1    10.29.165.47     NA                                    NA
2    10.29.165.47     NA                                    NA
16   10.29.164.93     9476f312-1321-e111-0000-1b000000006e  10.29.164.93
17   10.29.164.91     9476f312-1321-e111-0000-1b000000005f  10.29.164.91
18   10.29.164.94     9476f312-1321-e111-0000-1b000000007e  10.29.164.94
19   10.29.164.92     9476f312-1321-e111-0000-1b000000001f  10.29.164.92
20   10.29.164.95     9476f312-1321-e111-0000-1b000000004e  ESXi5-BFS-Srv5
```

6. Repeat the procedure to configure additional VSM pairs for each ESX cluster.

SAN Configuration

The same pair of Nexus 5548UP switches were used in the configuration to connect between the FC ports on the EMC VNX5500 and the FC ports of the UCS 6248 Fabric Interconnects.

Boot from SAN Benefits

Booting from SAN is another key feature which helps in moving towards stateless computing in which there is no static binding between a physical server and the OS / applications it is tasked to run. The OS is installed on a SAN LUN and boot from SAN policy is applied to the service profile template or the service profile. If the service profile were to be moved to another server, the pwwn of the HBAs and the Boot from SAN (BFS) policy also moves along with it. The new server now takes the same exact character of the old server, providing the true unique stateless nature of the UCS Blade Server.

The key benefits of booting from the network:

- Reduce Server Footprints: Boot from SAN alleviates the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- Disaster and Server Failure Recovery: All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys functionality of the servers at the primary site, the remote site can take over with minimal downtime.
- Recovery from server failures is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of its image. As a result, boot from SAN can greatly reduce the time required for server recovery.

- High Availability: A typical data center is highly redundant in nature - redundant paths, redundant disks and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for mechanical failure of a local disk.
- Rapid Redeployment: Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage a cost effective endeavor.
- Centralized Image Management: When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

With Boot from SAN, the image resides on a SAN LUN and the server communicates with the SAN through a host bus adapter (HBA). The HBAs BIOS contain the instructions that enable the server to find the boot disk. All of the FC-capable Converged Network Adapter (CNA) cards supported on Cisco UCS B-series blade servers support Boot from SAN.

After power on self-test (POST), the server hardware component fetches the boot device that is designated as the boot device in the hardware BOIS settings. Once the hardware detects the boot device, it follows the regular boot process.

Configuring Boot from SAN Overview

There are three distinct phases during the configuration of Boot from SAN. The high-level procedures are:

1. SAN zone configuration on the Nexus 5548UPs
2. Storage array host initiator configuration
3. Cisco UCS configuration of Boot from SAN policy in the service profile

In each of the following sections, each high-level phase will be discussed.

SAN Configuration on Nexus 5548UP

The FCoE and NPIV feature has to be turned on in the Nexus 5500 series switch. Make sure you have 8 GB SFP+ modules connected to the Nexus 5548UP ports. The port mode is set to AUTO as well as the speed is set to AUTO. Rate mode is "dedicated" and when everything is configured correctly you should see something like the output below on a Nexus 5500 series switch for a given port (for example, Fc1/17).

Note A Nexus 5500 series switch supports multiple VSAN configurations. A single VSAN was deployed in this study.

Cisco Fabric Manager can also be used to get a overall picture of the SAN configuration and zoning information. As discussed earlier, the SAN zoning is done upfront for all the pwwns of the initiators with the EMC VNX5500 target pwwns.

```
VDI-N5548-A# show feature | grep npiv
npiv                  1         enabled

 VDI-N5548-A# show interface brief
--------------------------------------------------------------------------------
```

```
Interface  Vsan   Admin  Admin    Status          SFP  Oper  Oper    Port
                          Mode     Trunk                              Mode Speed
Channel
                                      Mode
(Gbps)
--------------------------------------------------------------------------------
fc1/17     1      auto   on       up              swl    F       8    --
fc1/18     1      auto   on       up              swl    F       8    --
```

The FC connection was used for configuring boot from SAN for all of server blades. In addition, a general purpose 1TB infrastructure LUN for infrastructure virtual machine storage and 14 write-cache LUNs for each VDI host were provisioned.

Single vSAN zoning was set up on the Nexus 5548's to make those VNX5500 LUNs visible to the infrastructure and test servers.

An example SAN zone configuration is shown below on the Fabric A side:

```
VDI-N5548-A# sh zone name B200M3-CH1-SERVER1-FC0 vsan 1
zone name B200M3-CH1-SERVER1-FC0 vsan 1
    member pwwn 20:00:00:25:b5:c1:00:af
!               [B200M3-CH1-SERVER1-fc0]
    member pwwn 50:06:01:60:46:e0:5e:0a
!               [VNX5500-A0]
    member pwwn 50:06:01:69:46:e0:5e:0a
!               [VNX5500-B1]

VDI-N5548-A# sh zone name B200M3-CH1-SERVER2-FC0 vsan 1
zone name B200M3-CH1-SERVER2-FC0 vsan 1
    member pwwn 20:00:00:25:b5:c1:00:9f
!               [B200M3-CH1-SERVER2-fc0]
    member pwwn 50:06:01:60:46:e0:5e:0a
!               [VNX5500-A0]
    member pwwn 50:06:01:69:46:e0:5e:0a
!               [VNX5500-B1]
```

Where 20:00:00:25:b5:c1:00:af /20:00:00:25:b5:c1:00:9f are blade servers pwwn's of their respective Converged Network Adapters (CNAs) that are part of the Fabric A side.

The EMC FC target ports are 50:06:01:60:46:e0:5e:0a /50:06:01:69:46:e0:5e:0a and belong to one port on the FC modules on SP-A and SP-B.

Similar zoning is done on the second Nexus 5548 in the pair to take care of the Fabric B side as shown below:

```
VDI-N5548-B# sh zone name B200M3-CH1-SERVER1-FC1 vsan 1    zone name
B200M3-CH1-SERVER1-FC1 vsan 1
    member pwwn 20:00:00:25:b5:c1:00:bf
               [B200M3-CH1-SERVER1-fc1]
    member pwwn 50:06:01:61:46:e0:5e:0a
               [VNX5500-A1]
    member pwwn 50:06:01:68:46:e0:5e:0a
               [VNX5500-B0]
 VDI-N5548-B# sh zone name B200M3-CH1-SERVER2-FC1 vsan 1
zone name B200M3-CH1-SERVER2-FC1 vsan 1
    member pwwn 20:00:00:25:b5:c1:00:8f
                [B200M3-CH1-SERVER2-fc1]
    member pwwn 50:06:01:61:46:e0:5e:0a
               [VNX5500-A1]
    member pwwn 50:06:01:68:46:e0:5e:0a
               [VNX5500-B0]
```

Where 20:00:00:25:b5:c1:00:bf /20:00:00:25:b5:c1:00:8f are blade servers pwwn's of their respective Converged Network Adapters (CNAs) that are part of the Fabric B side.

The EMC FC target ports are 50:06:01:61:46:e0:5e:0a /50:06:01:68:46:e0:5e:0a and belong to the other port on the FC modules on SP-A and SP-B. They were spread across the two controllers for redundancy as shown in Figure 16.

Figure 16 VNX5500 FC Target Ports

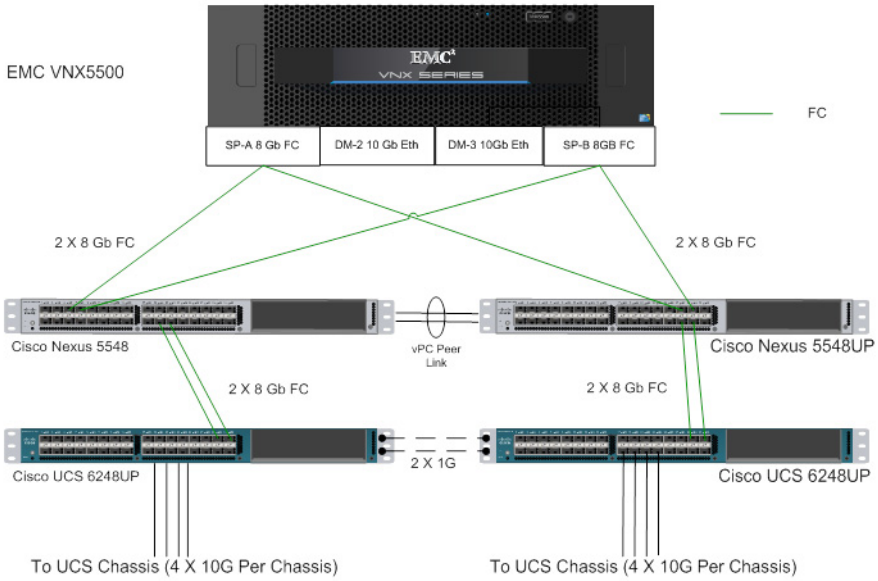

For detailed Nexus 5500 series switch configuration, refer to Cisco Nexus 5500 Series NX-OS SAN Switching Configuration Guide. (See the Reference Section of this document for the link.)

Configuring Boot from SAN on EMC VNX

The steps required to configure boot from SAN LUNs on EMC VNX are as follows:

1. Create a storage pool from which LUNs will be provisioned. RAID type, drive number and type are specified in the dialogue box below. Five 600GB SAS drives are used in this example to create a RAID 5 pool. Uncheck "Schedule Auto-Tiering" to disable automatic tiering.

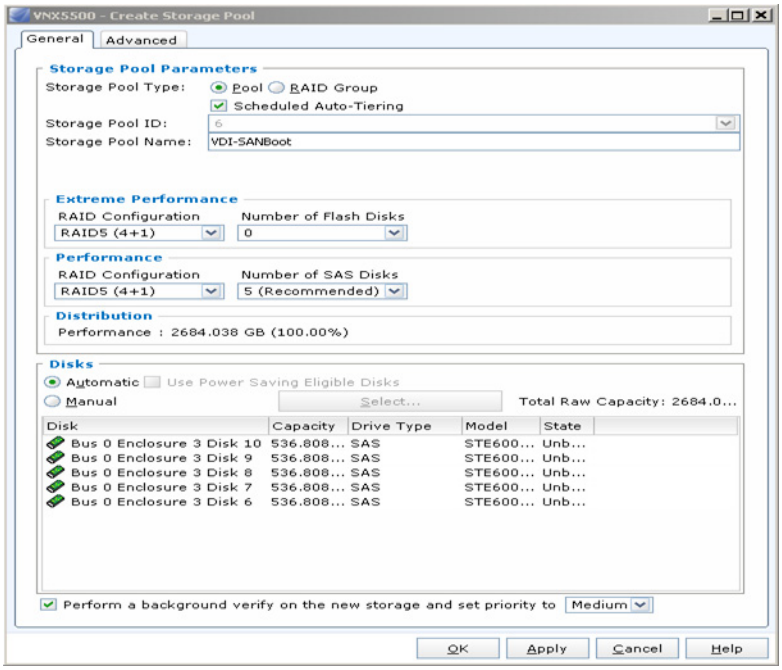

2. Provision LUNs from the storage pool created in step 1. Each LUN is 12GB in size to store the ESXi hypervisor OS.

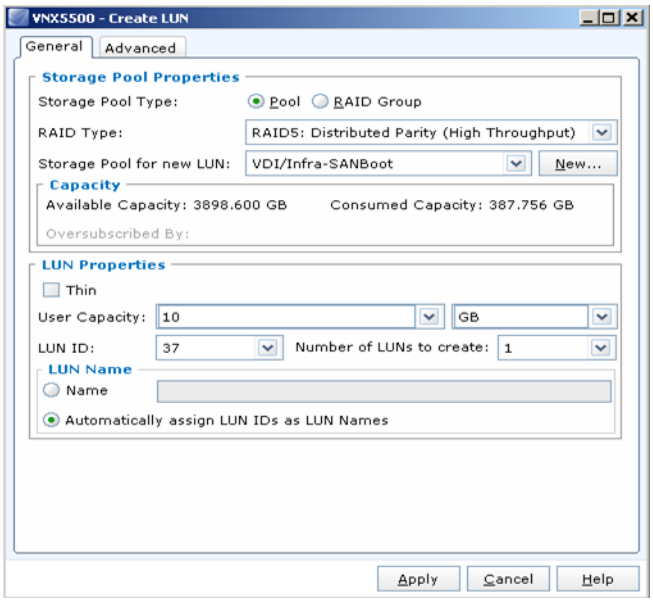

3. Create a storage group, the container used for host to LUN mapping, for each of the ESXi hosts.

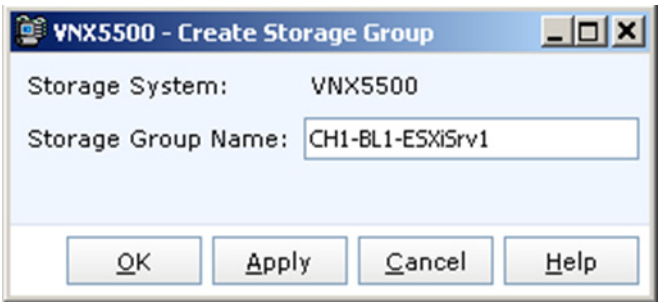

4. Register host initiators with the storage array to associate a set of initiators with a given host. The registered host will be mapped to a specific boot LUN in the following step.

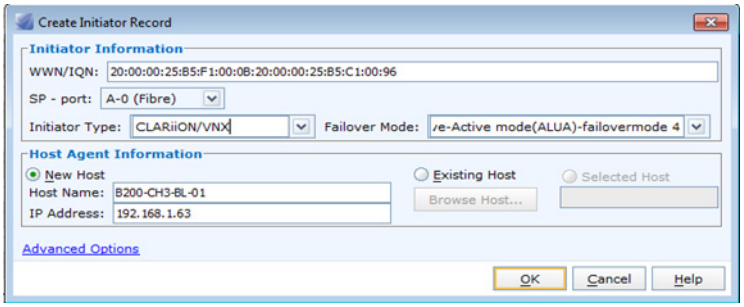

5. Assign each registered host to a separate storage group

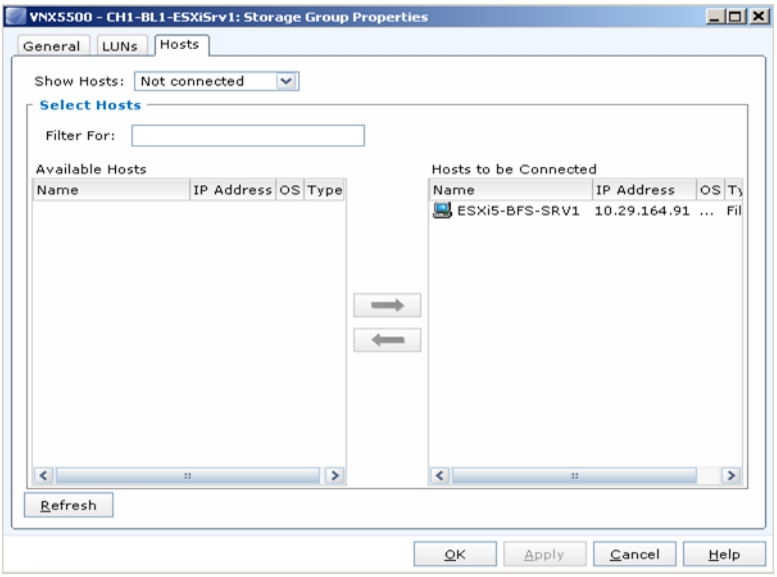

6. Assign a boot LUN to each of the storage groups. A host LUN ID is chosen to make visible to the host. It does not need to match the array LUN ID. All boot LUNs created for the testing are assigned host LUN ID 0.

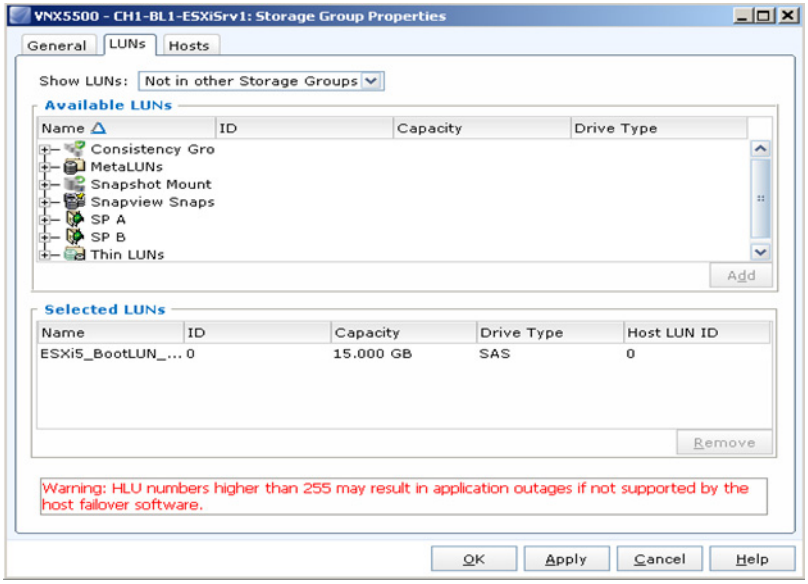

When the Cisco UCS Blade Server boots up, its vHBAs will connect to the provisioned EMC Boot LUNs and the hypervisor operating system can be installed.

SAN Configuration on Cisco UCS Manager

To enable Boot from SAN on the Cisco UCS Manager 2.1 (UCS-M) series, do the following:

1. Add SAN Boot for primary to the new policy. The vHBA name is optional, it can be left blank and it is unnecessary to enforce the vHBA name. Click OK.

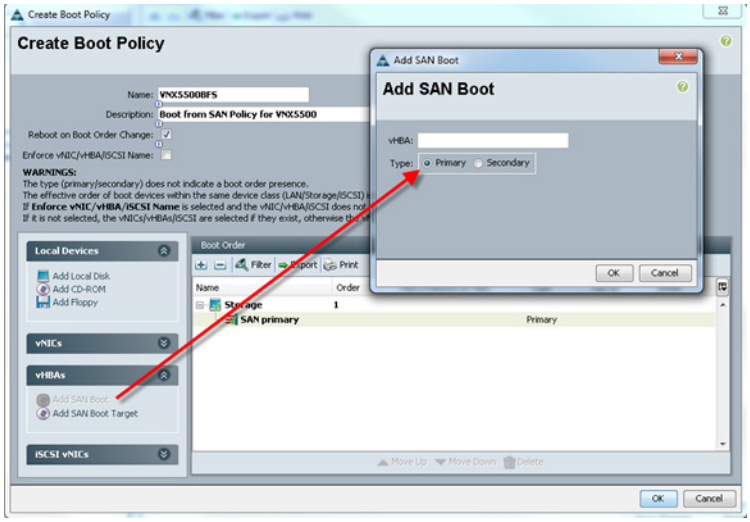

2. Add SAN Boot for primary to the new policy. The vHBA name is optional, it can be left blank and it is unnecessary to enforce the vHBA name. Click OK.

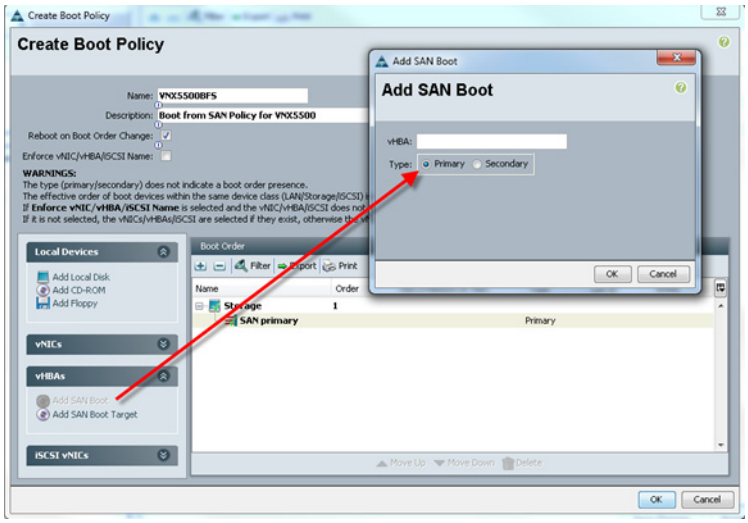

3. Add SAN boot for SAN Secondary, click OK. Leave the optional vHBA name blank.

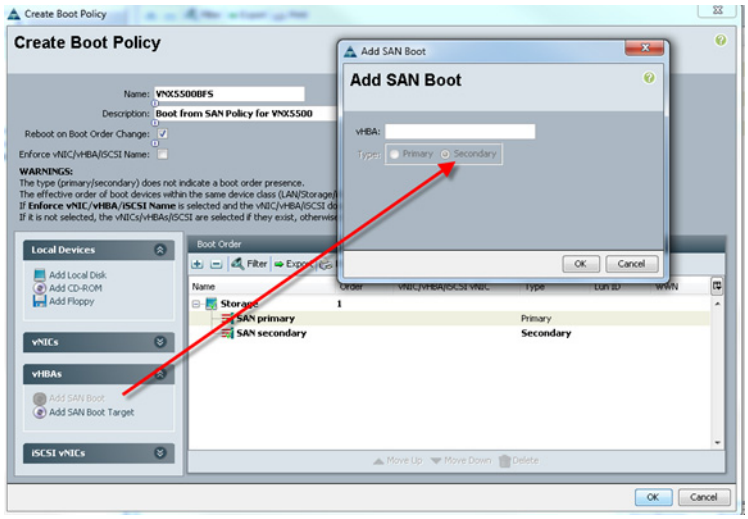

4. Add the Boot target WWPN to the SAN Primary; make sure this matches the EMC VNX pwwn. To avoid any mistakes, copy and paste from the Nexus 5500 Series command, for each switch, as shown below:

```
VDI-N5548-A# show fcns database vsan 1
0xe300ef     N      50:06:01:60:46:e0:5e:0a (Clariion)      scsi-fcp:both
0xe301ef     N      50:06:01:69:46:e0:5e:0a (Clariion)      scsi-fcp:both
VDI-N5548-B #  show fcns database vsan 1
0x470400     N      50:06:01:61:46:e0:5e:0a (Clariion)        scsi-fcp
0x470500     N      50:06:01:68:46:e0:5e:0a (Clariion)        scsi-fcp
```

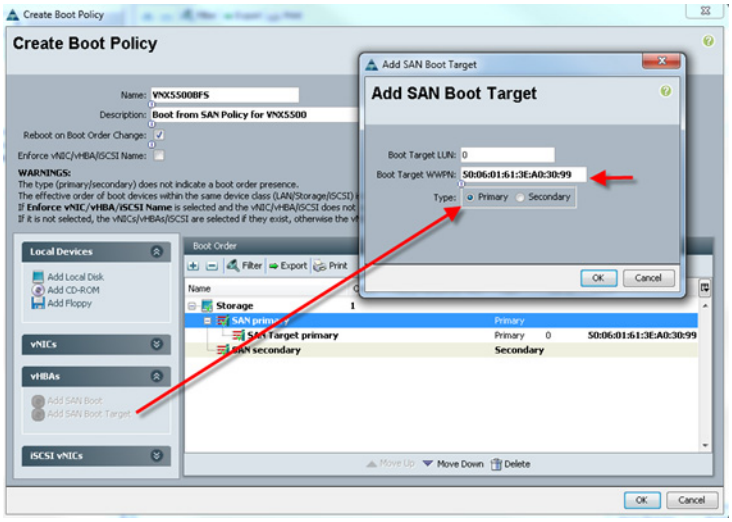

5. Repeat step 4 for SAN Primary's—SAN Target Secondary.
6. Repeat step 4 for SAN Secondary's—SAN Target Primary.
7. Repeat step 4 for SAN Secondary's—SAN Target Secondary.

At the end your Boot from SAN policy will look like the following:

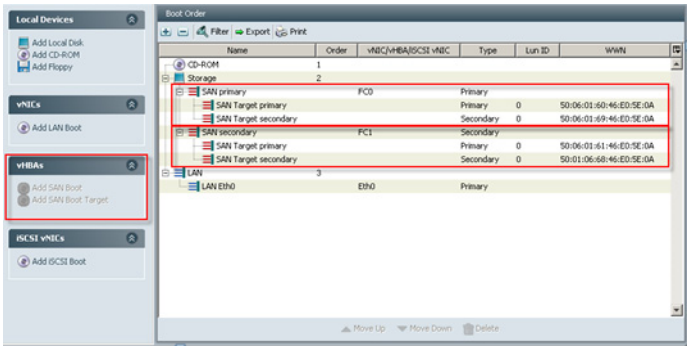

8. Make the association of the service profile template to the Boot from SAN policy during the service profile template configuration which was described earlier in this document.

EMC VNX5500 Storage Configuration

Figure 17 illustrates the layout of the disks that are required to store 2,000 desktop virtual machines. This layout does not include space for user profile data. Refer to "VNX shared file systems" on page 23 section for more information.

Figure 17 Core Storage Layout

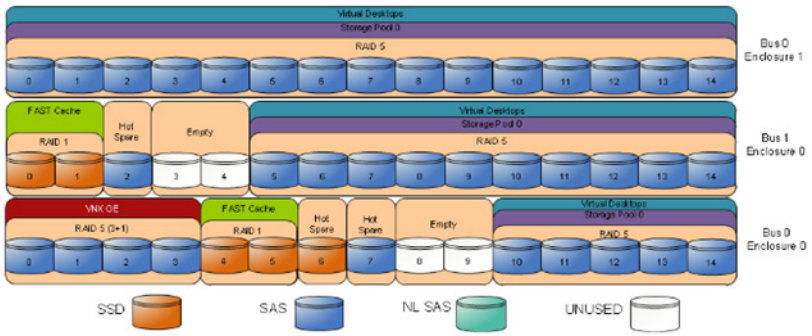

Core Storage Layout Overview

The following core configuration is used in the reference architecture:

- Four SAS disks (0_0_0 to 0_0_3) are used for the VNX OE
- Disks 0_0_6, 0_0_7, and 1_0_2 are hot spares. These disks are marked as hot spare in the storage
- layout diagram.
- Thirty SAS disks (0_0_10 to 0_0_14, 1_0_5 to 1_0_14, and 0_1_0 to 0_1_14) in the RAID 5 storage
- pool 0 are used to store virtual desktops. FAST Cache is enabled for the entire pool.
- For NAS, thirty LUNs of 200 GB each are carved out of the pool to provide the storage required to create fourteen 410 GB NFS file systems and two 50 GB file systems. The file systems are presented to the vSphere servers as NFS datastores.
- For FC, sixteen LUNs of 375 GB each and two LUNs of 50 GB each are carved out of the pool to present to the vSphere servers as eighteen VMFS datastores.
- Four Flash drives (0_0_4 to 0_0_5 and 1_0_0 to 1_0_1) are used for EMC VNX FAST Cache. There are no user-configurable LUNs on these drives.
- Disks 0_0_8 to 0_0_9 and 1_0_3 to 1_0_4 were not used for testing this solution

In solution validation testing, storage space for user data and infrastructure was allocated on the VNX array as shown in the following figure. This storage is in addition to the core storage shown above. If storage for user data exists elsewhere in the production environment, this storage is not required.

EMC Storage Configuration for VMware View

EMC Storage Configuration for VMware ESXi 5.1 Infrastructure Clusters

The steps required to configure LUNs for the VDI Datastores are as follows:

1. Create a storage pool using EMC Unisphere; select Storage > Storage Configuration > Storage Pools and click Create.

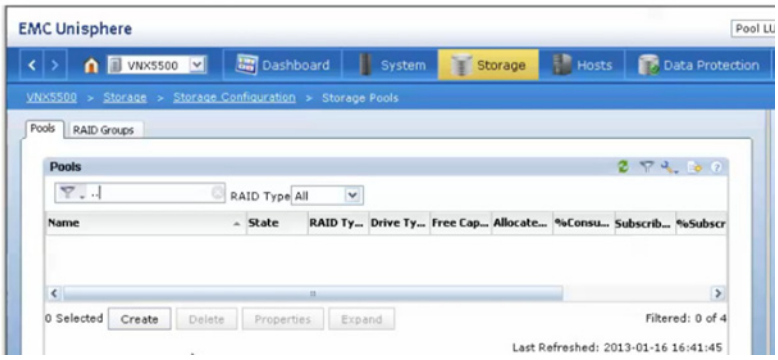

2. Create a storage pool from LUNs from the pool created in step 1. RAID type, drive number and type are specified in the dialog box. Select 30 x 600GB SAS drives from manual Selection to create RAID 5 Storage pool: Uncheck Schedule Auto-Tiering to disable automatic tiering.

Two LUNs of 2.08TB are carved out of the RAID 5 storage pool configured with 10 SAS drives. The LUNs are used to store infrastructure virtual machines such as VMware Horizon View 5.2 Connection Servers, View Composer servers, and VMware vCenter server.

Example EMC Boot LUN Configuration

Each ESXi server requires a boot LUN from SAN for the hypervisor OS. A total of 43 LUNs are carved out of the 5-disk RAID 5 pool. Each LUN is 5GB in size.

EMC FAST Cache in Practice

EMC FAST Cache uses Flash drives to add an extra layer of cache between the dynamic random access memory (DRAM) cache and rotating disk drives, thereby creating a faster medium for storing frequently accessed data. FAST Cache is an extendable Read/Write cache. It boosts application performance by ensuring that the most active data is served from high-performing Flash drives and can reside on this faster medium for as long as is needed.

FAST Cache tracks data activity at a granularity of 64KB and promotes hot data in to FAST Cache by copying from the hard disk drives (HDDs) to the Flash drives assigned to FAST Cache. Subsequent IO access to that data is handled by the Flash drives and is serviced at Flash drive response times-this ensures very low latency for the data. As data ages and becomes less active, it is flushed from FAST Cache to be replaced by more active data.

Only a small number of Flash drives are needed enabling FAST Cache to provide greater performance increases than implementing a large number of short-stroked HDDs. This results in cost savings in data center space, power, and cooling requirements that lowers overall TCO for the business.

FAST Cache is particularly suited to applications that randomly access storage with high frequency, such as Oracle and SQL OLTP databases. OLTP databases have inherent locality of reference with varied IO.

Cisco UCS Manager Configuration for VMware ESXi 5.1

This section addresses creation of the service profiles and VLANs to support the project.

Service Profile Templates

Two types of service profiles were required to support two different blade server types (Table 7).

Table 7 Role/Server/OS Deployment

Role	Blade Server Used	Operating System Deployed
Infrastructure	UCS B250 M2	ESXi 5.1
VDI Hosts	UCS B200 M3	ESXi 5.1

To support those different hardware platforms, service profile templates were created, utilizing various policies created earlier.

The service profile templates were then used to quickly deploy service profiles for each blade server in the Cisco UCS system. When each blade server booted for the first time, the service profile was deployed automatically, providing the perfect configuration for the VMware ESXi 5.1 installation.

VLAN Configuration

In addition, to control network traffic in the infrastructure and assure priority to high value traffic, virtual LANs (VLANs) were created on the Nexus 5548s, on the Cisco UCS Manager (Fabric Interconnects,) and on the Nexus 1000V Virtual Switch Modules in each vCenter Cluster. The virtual machines in the environment used the VLANs depending on their role in the system.

A total of seven Virtual LANs, VLANs, were utilized for the project. Table 8 identifies them and describes their use.

Table 8 VLAN Naming and Use

VLAN Name	VLAN ID Use	
VDA	122	VDI Virtual Machine Traffic
MGMT	164	VMware ESXi Management
Infra-Mgmt	165	Infrastructure Management Traffic (vCenter, SQL, AD, 1000Vetc)
STRG	166	VNX5500 NFS Traffic (Optional)
N1K-Control	167	Nexus 1000V Control Traffic
VMOTION	169	VMware vMotion Traffic

VLANs are configured using Cisco UCS Manager from the LAN tab; LAN\VLANs node in the left pane of Cisco UCS Manager. They were set up earlier in section 7.2.1 Base Cisco UCS System Configuration.

Installing and Configuring ESXi 5.1

In this study, we used Fibre Channel storage to boot the hosts from LUNs on the VNX5500 storage system. Prior to installing the operating system, storage groups are created, assigning specific boot LUNs to individual hosts. (See section Configuring Boot from SAN on EMC VNX for details.)

VMware ESXi 5.1 can be installed in boot-from-SAN mode using standard hypervisor deployment techniques including:

- Mounting a Cisco Customized ESXi 5.1 ISO image from the KVM of the blade
- Using automated deployment tools from third party sources (Optional)

Install VMware ESXi 5.1

ESXi was installed using the UCS Manager KVM console with the Cisco Customized ESXi 5.1 ISO image mounted. The Cisco UCS Manager boot policy deployed to each blade was set to boot from CD then SAN to accommodate hypervisor installs or updates.

The IP address, hostname, and NTP server were configured using Direct Console ESXi Interface accessed from Cisco UCS Manager KVM console; http://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc_50%2FGUID-26F3BC88-DAD8-43E7-9EA0-160054954507.html

http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-51-command-line-interface-getting-started-guide.pdf.

Install and Configure vCenter

To manage hypervisors and virtual machines a dedicated vCenter server instance was installed on Windows 2008R2 SP1 based virtual machine.

Vmware vCenter Server			
OS:	Windows 2008 R2	**Service Pack:**	
CPU:	4vCPUs	**RAM:**	16GB
Disk:	80GB	**Network:**	1x10Gbps

To support vCenter instance a Microsoft SQL Server 2008 R2 server was created to host vCenter database. Refer to the Microsoft documentation about configuring SQL Server and SQL Server clusters: http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx and http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx)

Install and Configure vCenter

1. Install the Microsoft® SQL Server® 2008 R2 Native Client for ODBC connections (http://www.microsoft.com/en-us/download/details.aspx?id=16978 look for Native Client for your architecture)
2. Create a System DSN (control panel, administrative tools, Data Sources ODBC) and connect to your vCenter-SQL server. Note: Make sure to use FQDN's for everything.
3. Create Active Directory user account and call it vcenter. (This user account will be used to connect to vCenter, you will have to follow a VMware specific procedure and assign specific permissions on vCenter for View components to connect to vCenter. (http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2034833).
4. Install vCenter server package, connect to the database.
5. Connect your vSphere client to vCenter and create a datacenter.
6. Create self-signed certificate (http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514).

Install VMware Licenses

1. Connect to vCenter using vSphere client.
2. Go to Home > Administration > Licensing.
3. Click Manage vSphere Licenses.

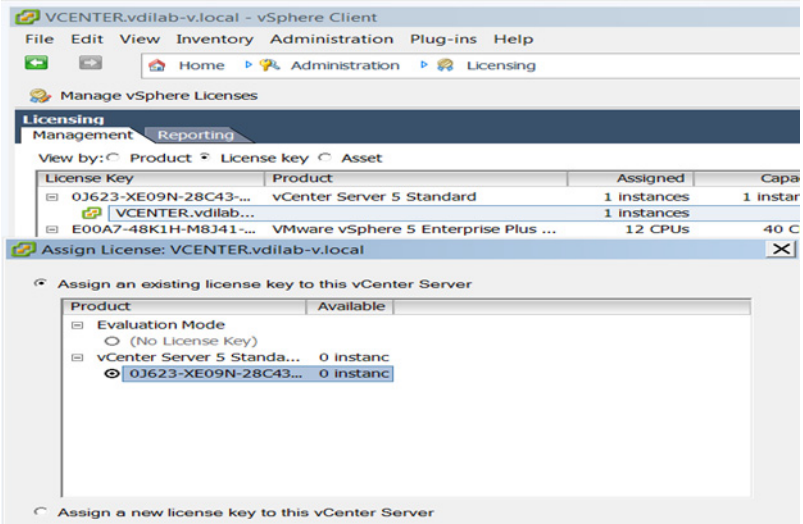

ESXi 5.1 Cluster Configuration

To accommodate maximum recommendations for View 5.2 on ESXi 5.1 we created four ESXi 5.1 clusters described below. For each of the two VDI clusters using seven Cisco UCS B200 M3 Blade Servers and 1000 virtual machines being hosted.

The 14 Cisco UCS B200 M3s and six Cisco UCS B250 M2 ESX hosts were configured into four Clusters:

- Infra-Cluster
- Launcher-CL
- VDI-Desktops-Pool1
- VDI-Desktops-Pool2

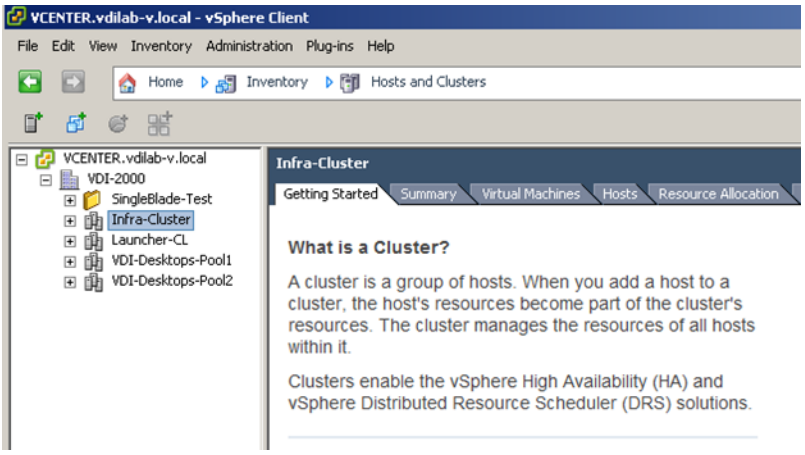

Infra-Cluster Infrastructure Cluster

The Infra-Cluster cluster was used to host all of the virtualized servers within the VDA Infrastructure, including two pairs of Nexus 1000V Virtual Switch Manager (VSM) appliances, one for each virtual desktop cluster.

Two physical Cisco UCS B250-M2 hosts were used in this cluster.

One standard switch to manage VMware Management, VDA, vMotion, and Storage traffic were configured on DC-INF cluster hosts. Three pairs of fault tolerant VSMs introduced the N1KV Management, Control and Packet VLANs to the environment.

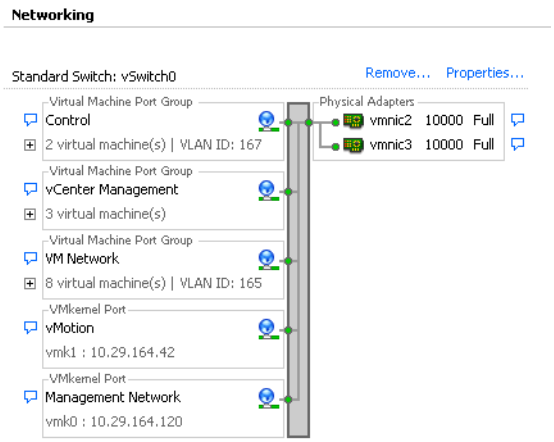

Virtual Desktop Clusters

The following clusters were used to host 2000 desktops:

VDI-CL1-N1KVVDI-CL2-N1KV

Each of these desktop clusters were configured identically with a Nexus 1000V high availability distributed virtual switch providing the required network connectivity.

The Nexus 1000V switches were configured to manage networking for all three ESX Clusters hosting virtual desktops, working in concert with the Cisco UCS Fabric Interconnects and Nexus 5548UP access layer switches to provide end to end Quality of Service for network communications, insuring the highest quality virtual desktop end user experience.

The Nexus 1000V configuration is described in detail in section Nexus 1000V Configuration in L3 Mode.

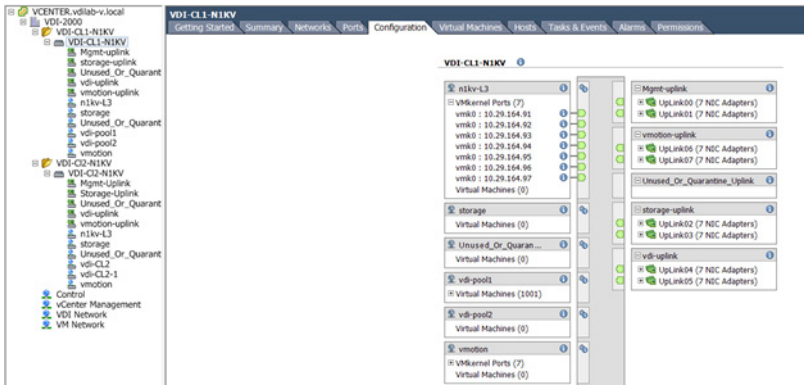

Login VSI Launcher Cluster

The separate Launcher-CL cluster was used to host Login Consultants' LoginVSI launcher VMs and a LoginVSI console VM. It was hosted on the same Cisco UCS Domain with dedicated storage.

The Laucher-CL cluster utilized the Standard vSwitch configured as shown below:

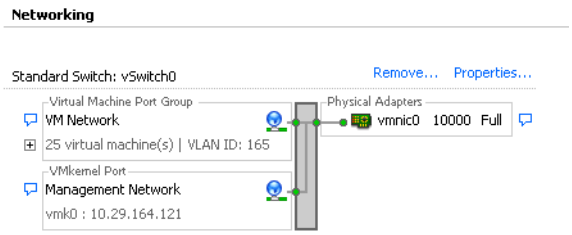

Installing and configuring VMware View 5.2

Building out a VMware View 5.2 environment requires the installation of the following components:

- View Connection Server
- View Replica Server
- View Administrator
- View Composer
- View Transfer Server

This section outlines the tasks required to build the View 5.2 environment used in this study. Refer to the VMware View Installation guide for View 5.2 for more details:

https://pubs.vmware.com/view-52/topic/com.vmware.ICbase/PDF/horizon-view-52-installation.pdf.

Pre-requisites

The following is a list of pre-requisites that are required with installing View 5.2 components. They are as follows:

- One of the following operating systems:
  - Windows Server 2008 R2, Standard or Enterprise Edition
  - Windows Server 2008 R2, Standard or Enterprise Edition SP1

Note You can mix operating systems within a site.

- vCenter 5.0 Update 1 or later
- A supported Microsoft SQL or Oracle database for vCenter and View Composer databases
- A supported vSphere hypervisor host operating system
- Physical or virtual hardware meeting the following recommended requirements
  - For View Connection Server: Pentium IV 2.0 Ghz or higher, 4 CPUs/vCPUs; 10GB+ RAM; 1GB NIC
  - For View Administrator: IE 8 or 9; Firefox 6 or 7; Adobe Flash 10 or later
  - For View Composer: 2.0 GHz or faster, 4 CPUs/vCPUs; 8GB+RAM; 1GB NIC; 60GB+ Disk Space
  - For View Transfer Server: Can co-exist on the same VM with any other View Manager component

Create SQL Databases for View 5.2

View Manager Installer needs a separate database for View Composer Server and View Server events.

Create Database for View Composer Server

1. Create a Database for View Composer server and create a user with server authentication.
2. On the VM where View Composer will be installed, go to Start' Administrative Tools ' ODBC.
3. Create a system DSN using DB server and user with SA authentication.

Create Event Database for View Administrator.

Create a Database for View Administrator Events and user with SA authentication.

Install View Manager and components

Download View Manager software from the link given below.

https://my.vmware.com/web/vmware/details?downloadGroup=VIEW-520-PREMIER&productId=320&rPId=3908

Install View Connection server

Log into View Connection server with Domain Administrator credentials.

1. Open installer file VMware-viewconnectionserver-x86_64-5.2.0-987719.exe with "Run as administrator."
2. Click Next.

3. Read the VMware End User License Agreement and click Next.

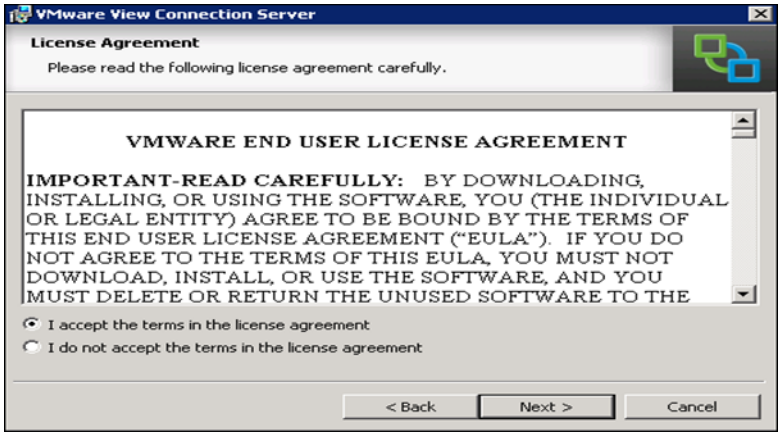

4. Select the location for the installer to install all the components and click Next.

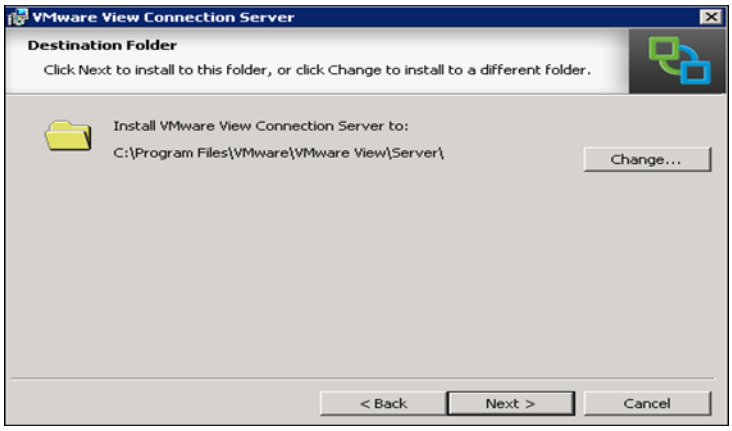

5. Select Standard server installation.

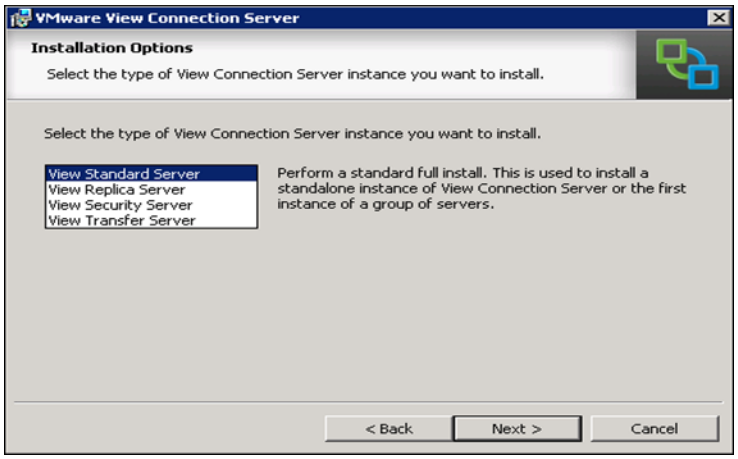

6. Enter a password and click Next.

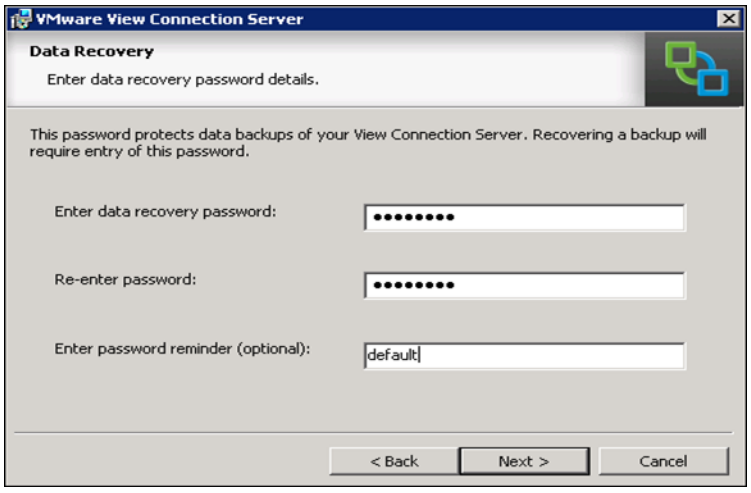

7. Select the radio button Configure Windows Firewall automatically. Click Next.

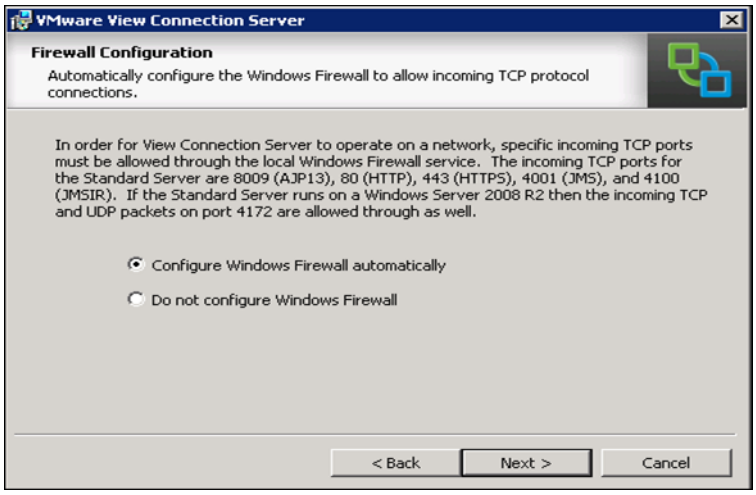

8. Select the radio button Authorize a specific domain user or domain group. Click Next.

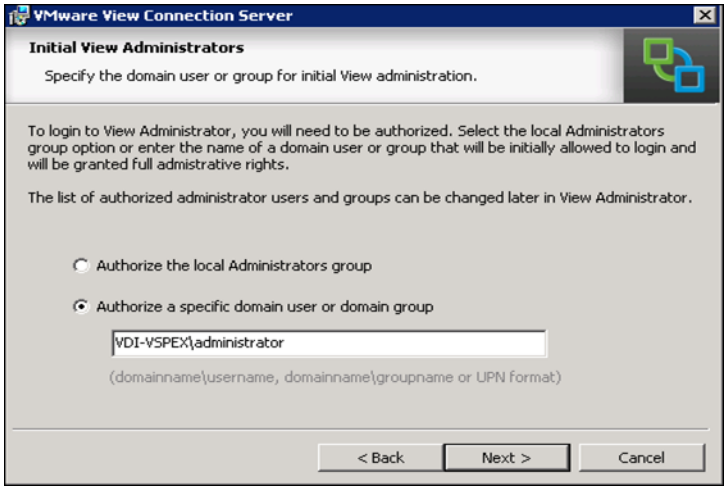

9. Uncheck the box Participate anonymously in the user experience improvement program. Click Next.
10. Click Install.

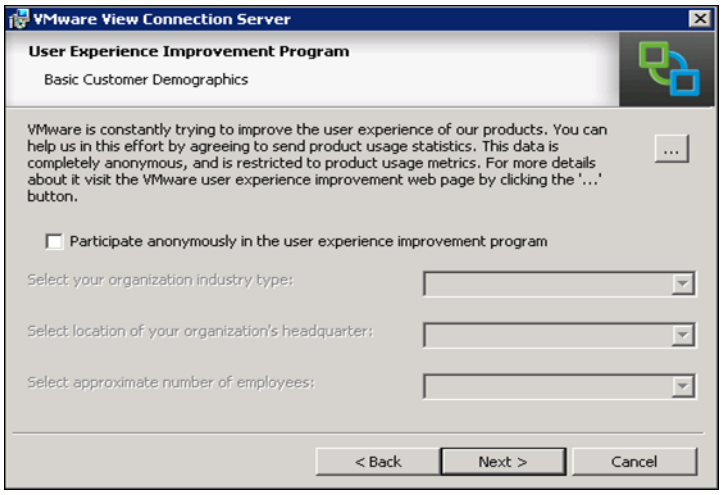

11. Double-click the View Administrator icon on the desktop; ignore the security warning on IE.

Note You need to install Flash player plugin v10.3 or higher to use Web Browser for Login to View Administrator.

12. Log into View Administrator GUI by entering username, password and Domain name.

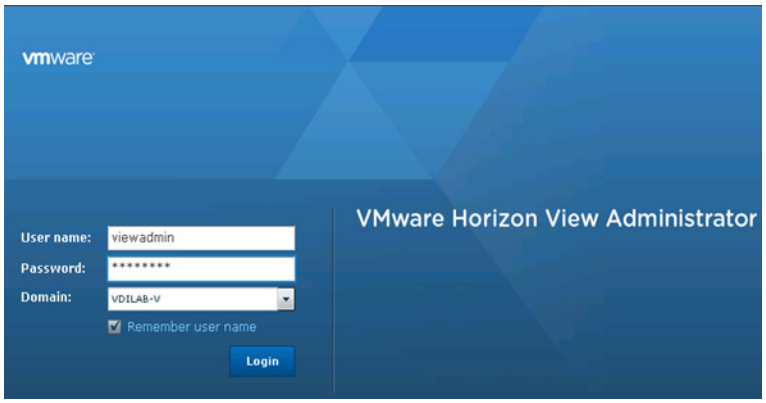

Install View Replica Server

1. Log into the replica server with domain administrator credentials
2. Open the installer file VMware-viewconnectionserver-x86_64-5.2.0-987719.exe with "Run as administrator."
3. Click Next.

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

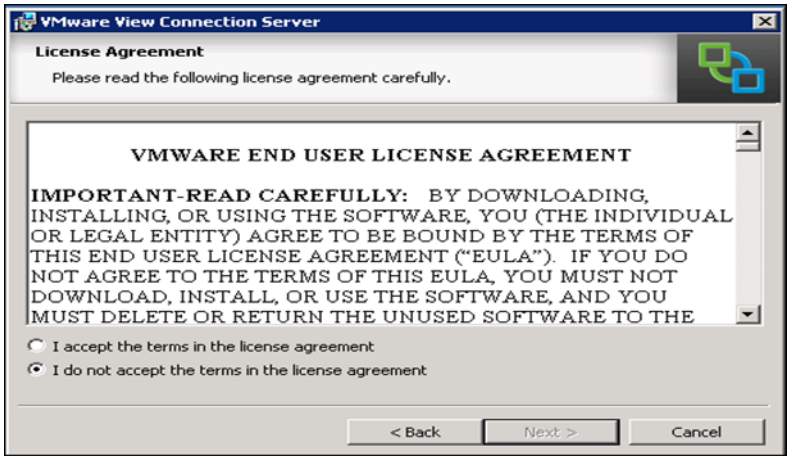

5. Select a destination for installation. Click Next.

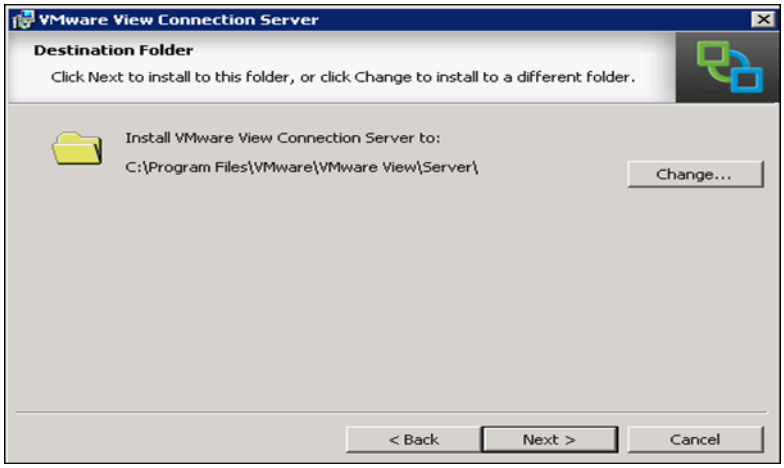

6. Select Replica server installation. Click Next.

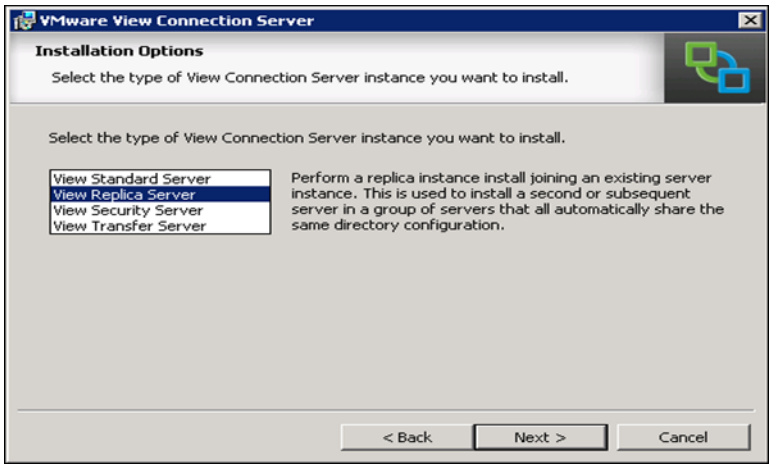

7. Select IP Address/Host name for view connection server primary instance to connect with replica server. (FQDN is recommended)

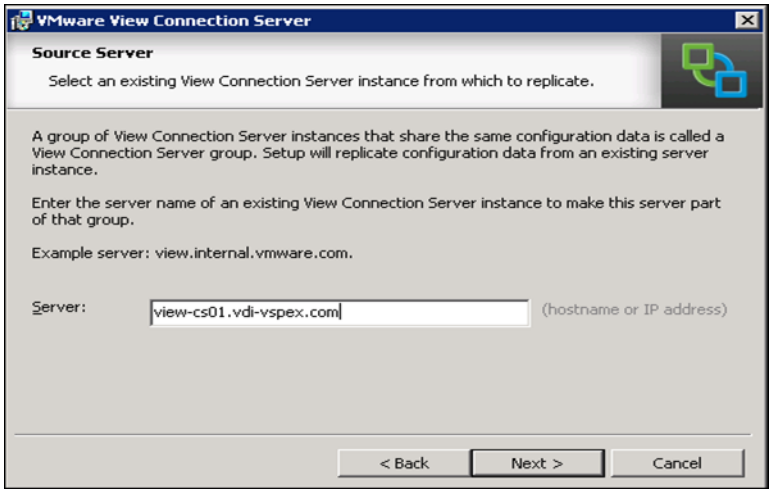

8. Click Next.

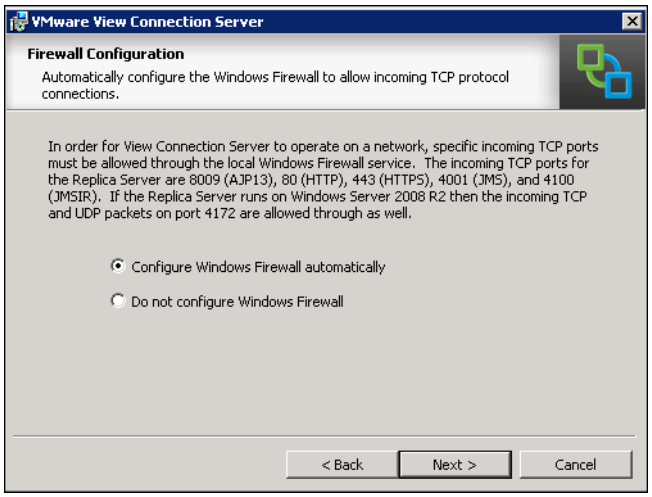

9. Click Next.

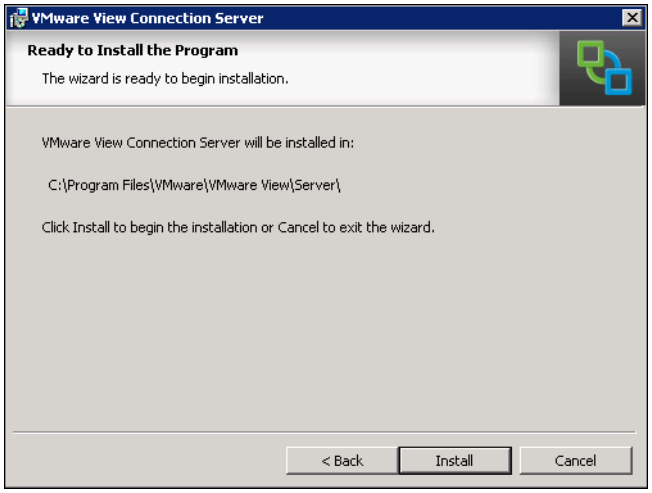

Install View Composer Server

View Composer server can be install on a separate stand alone server or on the same server that was used for vCenter server installer. For our test, we installed View Composer server on the same server we used for vCenter server.

1. Open View composer installer VMware-viewcomposer- 5.2.0-983460.
2. Create a database and ODBC connection for view composer installation. See section 7.8.2.1 for how to create database for view composer server.
3. Click Next.

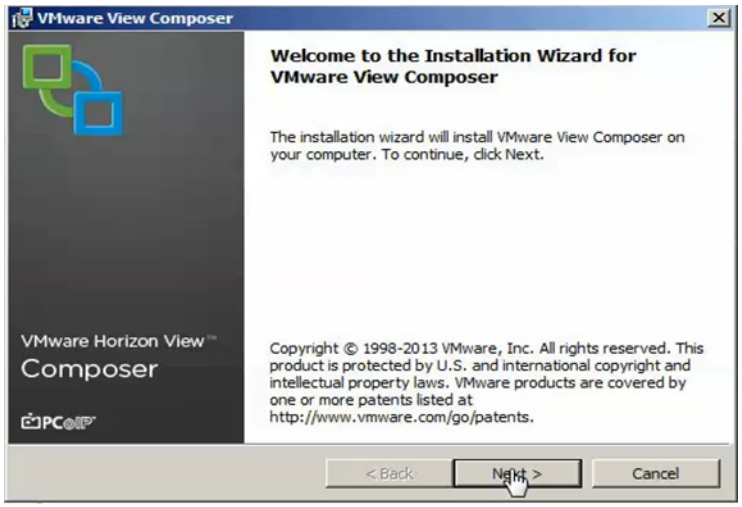

4. Read the VMware End User License and click Next.

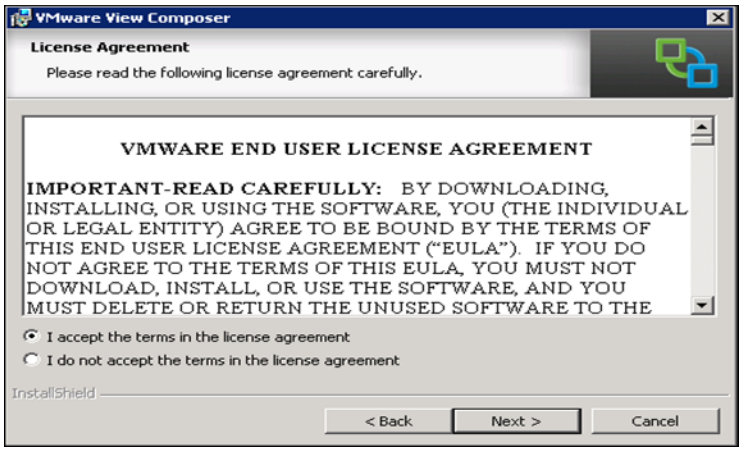

5. Select a location for the View Composer installation. Click Next.

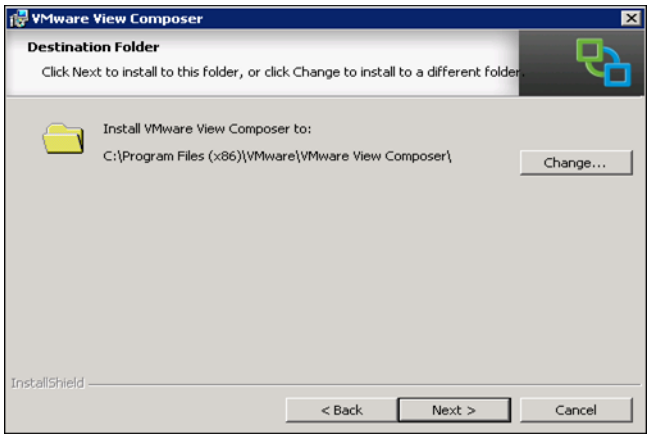

6. Enter the newly created Database and SA user Information for the View Composer installation.Click Next.

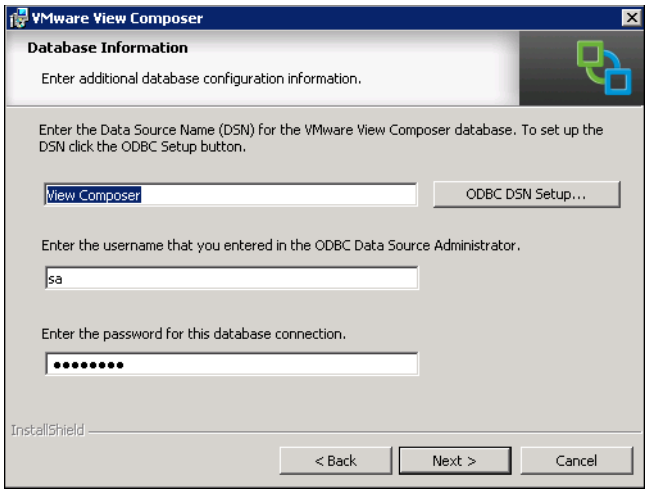

7. Accept the default port settings and click Next.

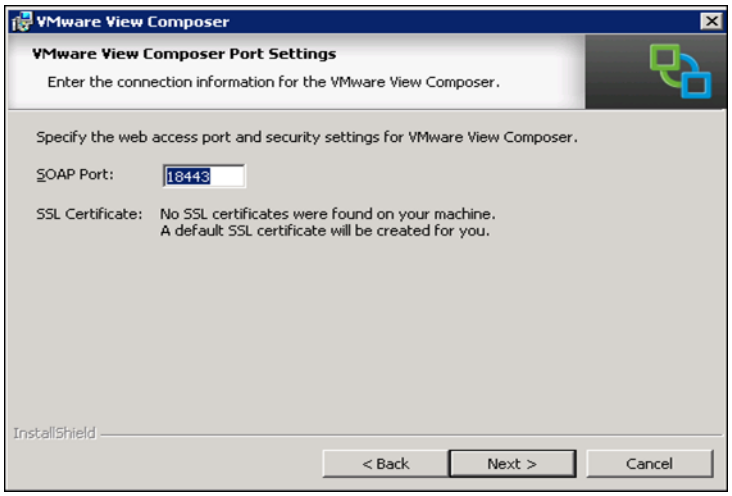

8. Click Install.

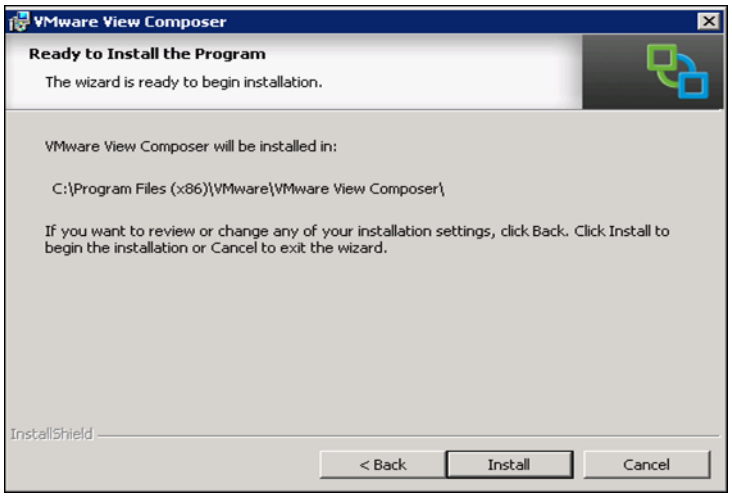

View Administrator Configuration

To configure the View 5.2 system, follow these steps:

1. Log into Vmware View Administrator using web browser.
2. Select View configuration.
3. From the drop-down menu select Product Licensing and Usage.
4. Click Edit settings and enter a valid License key for View Manager.

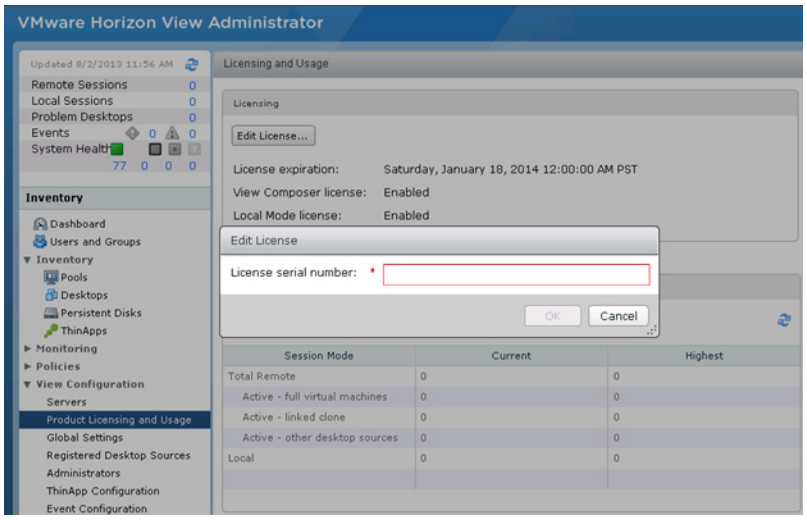

5. In View Configuration Click on servers. Select vCenter Servers tab. Click Add.

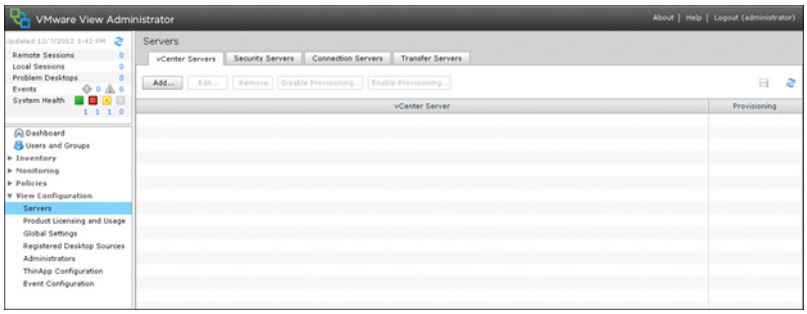

6. Enter FQDN for vCenter server and username/password. Make the necessary changes for Advanced settings. Click Next.

For this test case we used the following parameters.

Max concurrent vCenter Provisioning operations: 20

Max concurrent Power operations: 10

Max concurrent View Composer maintenance operations: 50

Max concurrent View Composer provisioning operations: 50

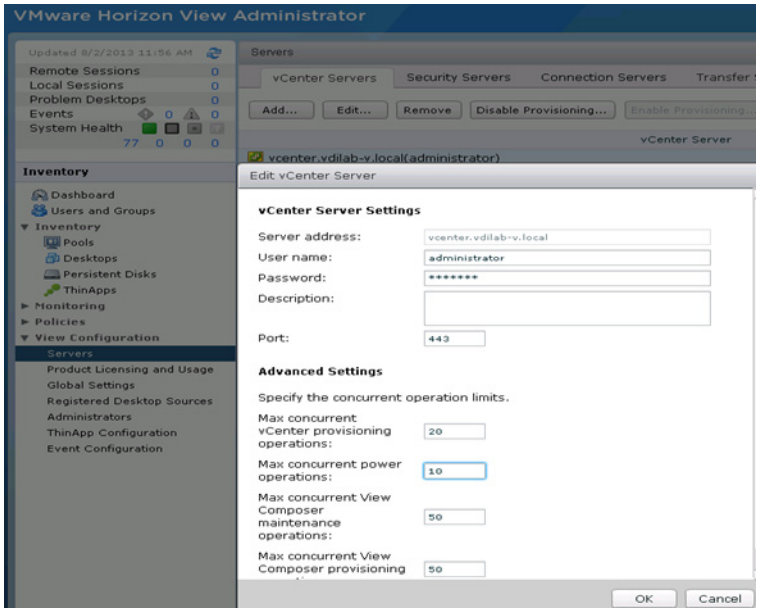

7. Click View Certificate and accept certificate warning.

8. Select View Composer settings. Select the radio button for View Composer Server; either co-installed with Vcenter or Standalone View Composer Server.

Note For the study, the View Composer Server was installed as a Standalone.

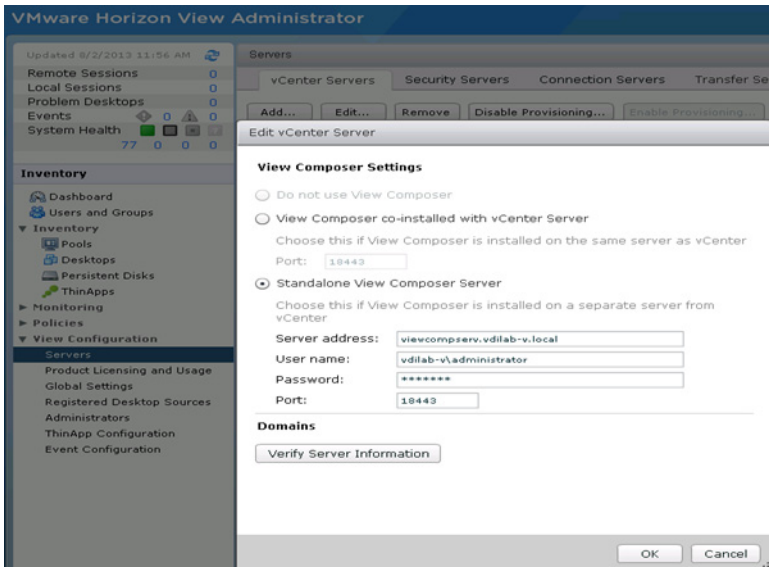

9. Click View Certificate and accept the certificate.

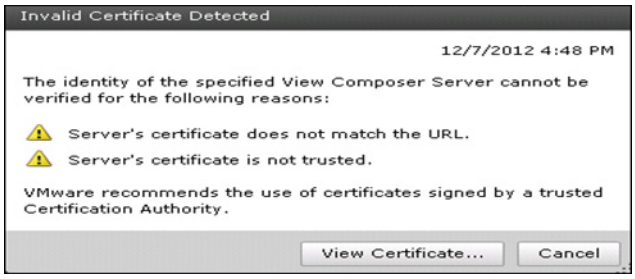

10. Click Add to add view composer domain.

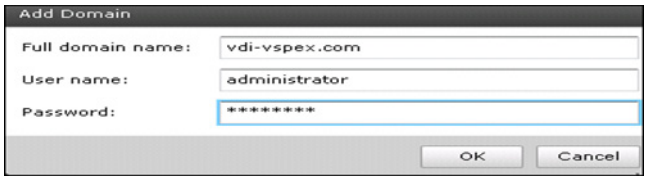

11. Click Next.

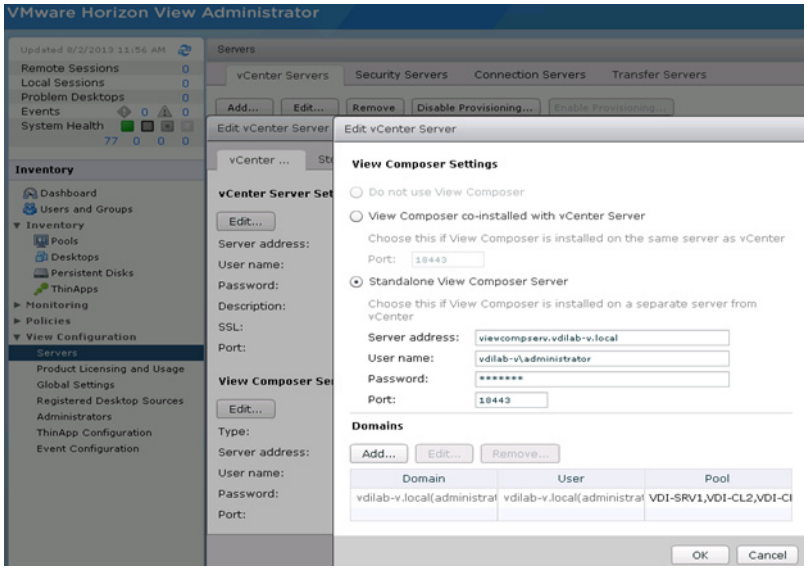

12. Select the check box to enable host caching. Set Default host cache siz. Click Next.

Note For this test case we used 2048MB cache size.

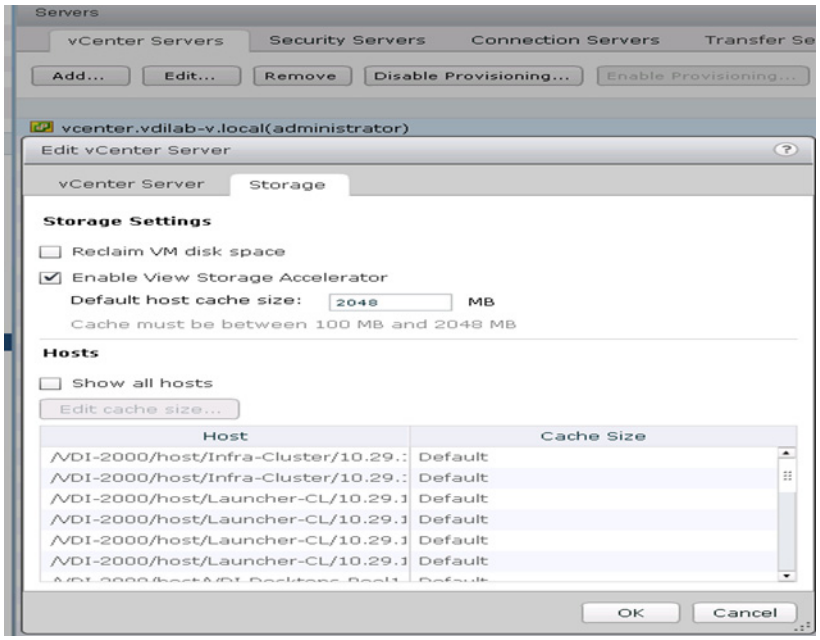

13. Create a new database for View Event Database in SQL server.
14. Click Event Database configuration.
15. Enter Database server information, database name, username/password. For the table prefix add VE_

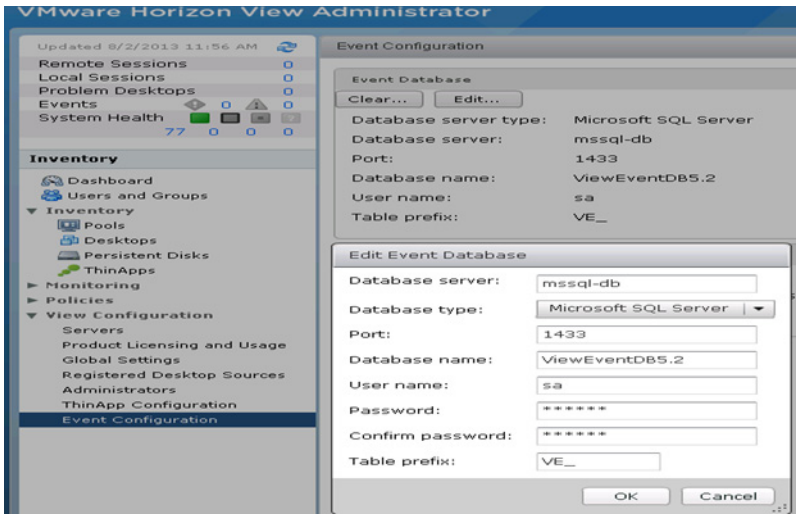

16. After completion of the View configuration go to Dashboard for View Administrator and check System Health and verify all components are shown as green.

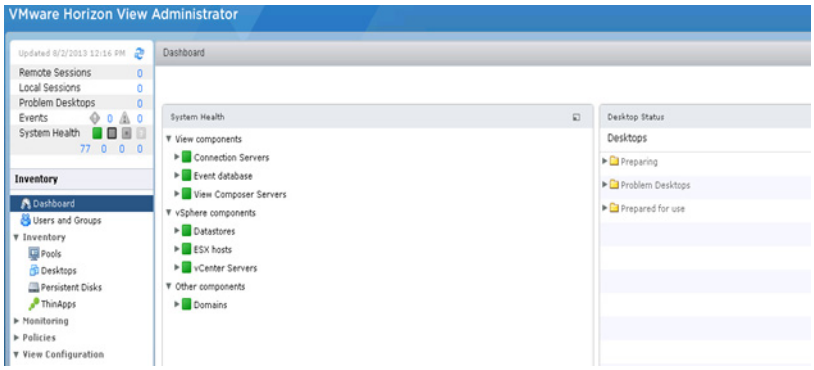

Install SSL Certificate for View Connection and Replica Server

1. Log into AD server and Add role for Active Directory Certifiacate services if does not exist.
2. Go to start Menu > Run > mmc.

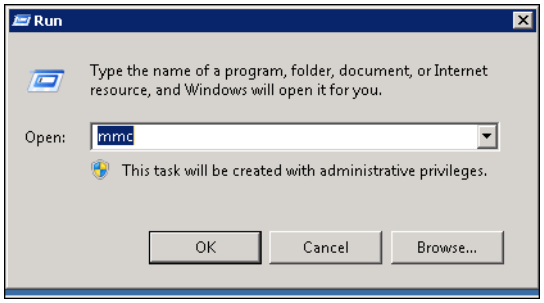

3. Click File and select Add/Remove Snap-in.

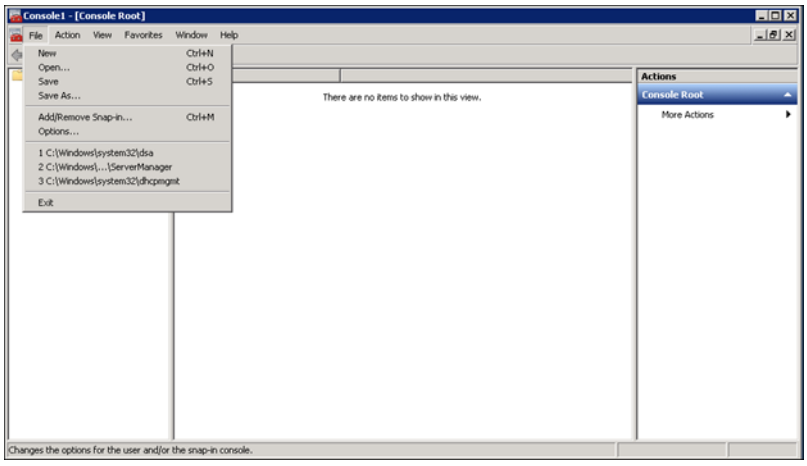

4. Select Certificates and click Add.

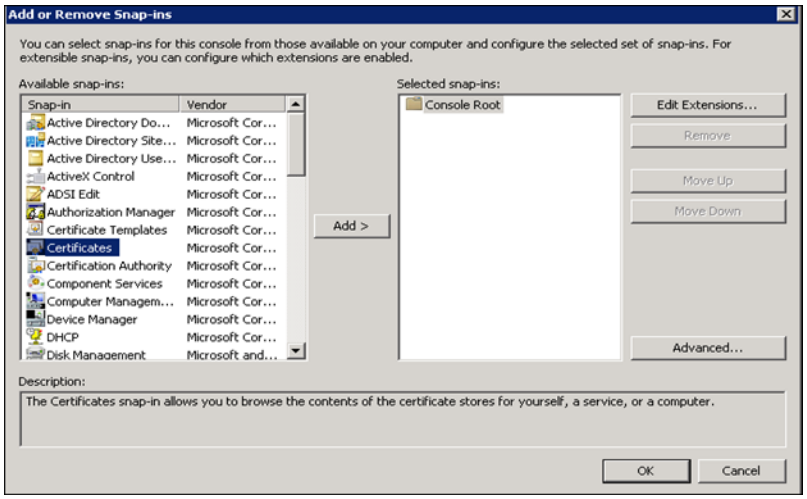

5. Select the radio button for Computer account.

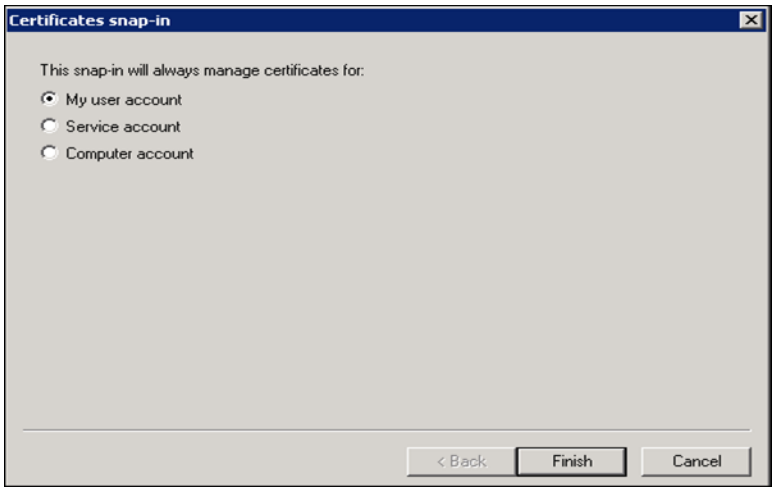

6. Select the radio button for Local Computer. Click Finish.

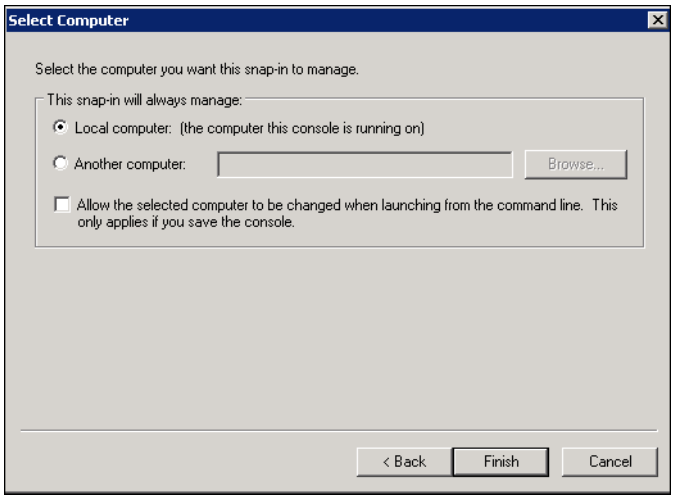

7. Add Certificate Templates and Certification Authority. Click OK.

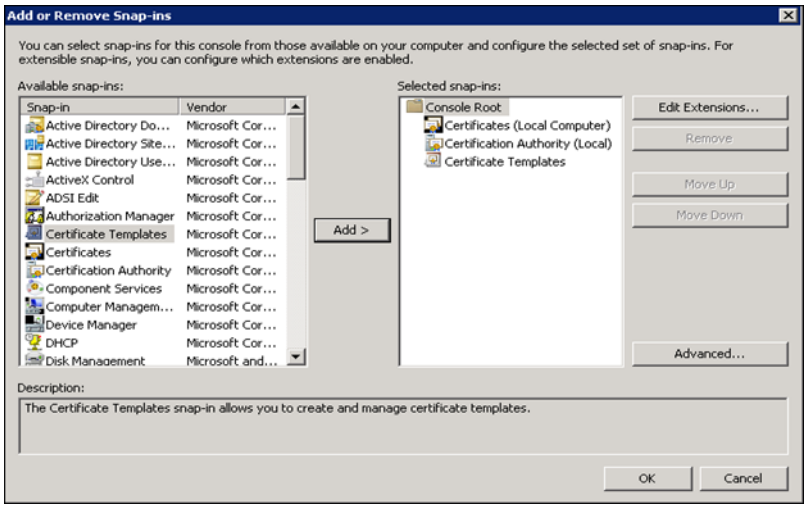

8. Click Certificate Template and from the list of template displayed on the right side select Web Server.
9. Right -click Web Server; select properties.

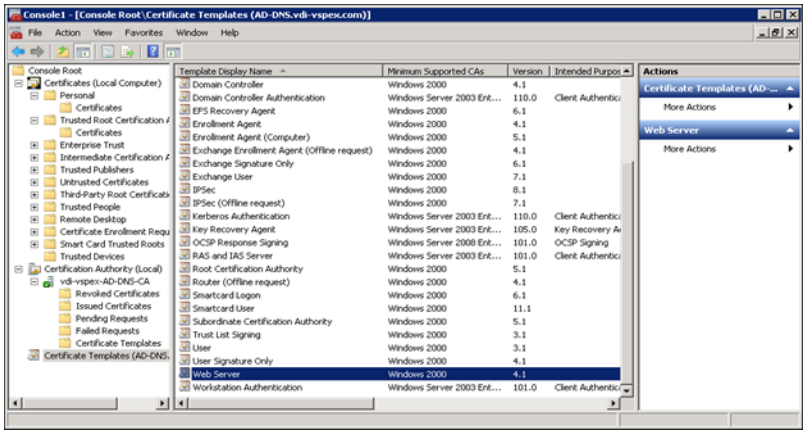

10. Select for the Security tab and add the computer name to assign for connection server, replica server. Allow full control to both servers.

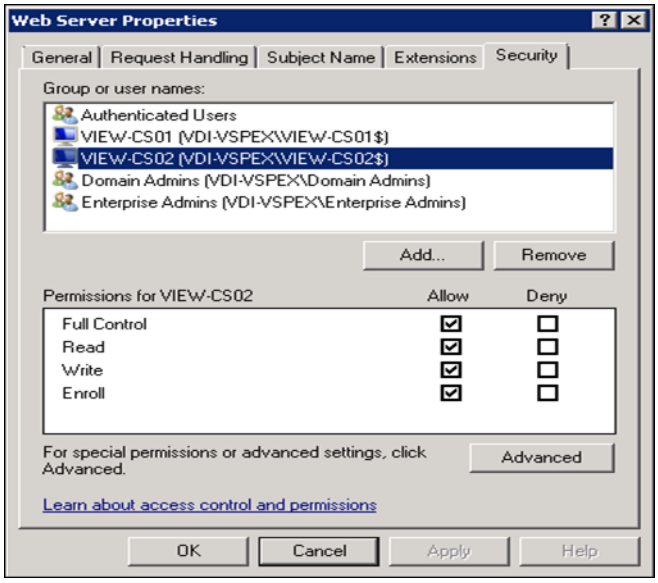

11. Select Certificates on the Console Root > Personal > Right-click Certificates.
12. Select Request New Certificate on All Tasks.

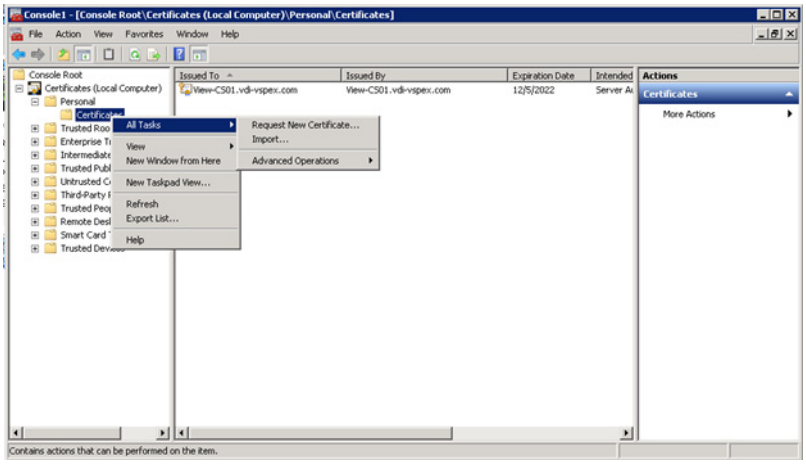

13. Click Next.

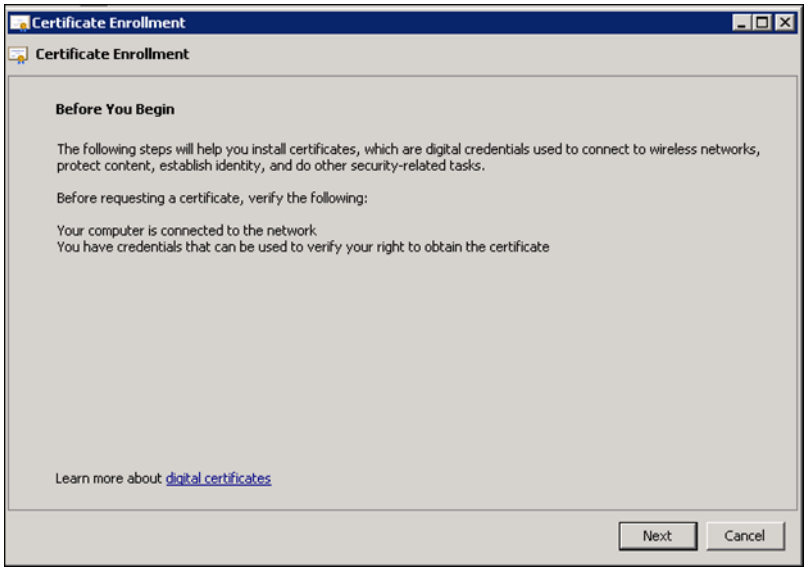

14. Check the box for Web Server. Click Details.
15. Click Properties.

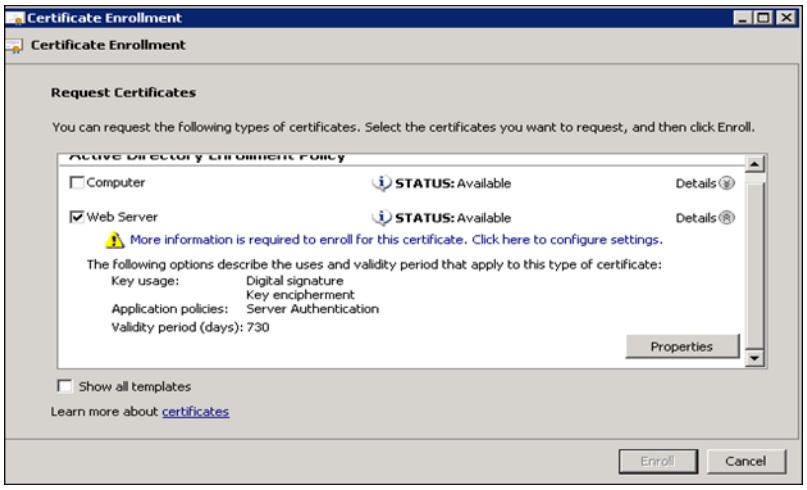

16. On the left side from the drop menu for Subject Type select Common Name, Organization, Country, Locale and add them with their appropriate value as shown in the screenshot below.
17. Alternative name: from the drop menu for type select DNS and add DNS name for view connection server. Do the same for view Replica server.
18. Click Apply.

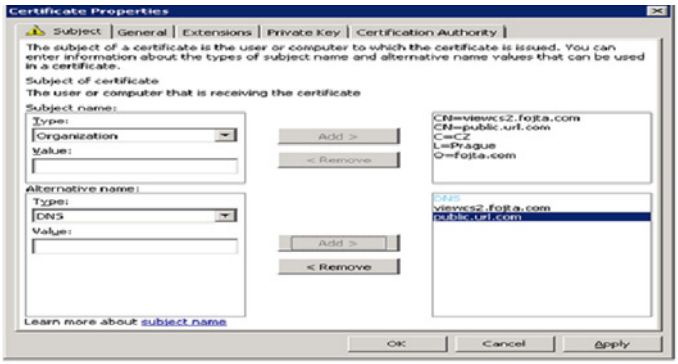

19. Click Apply > Click OK.

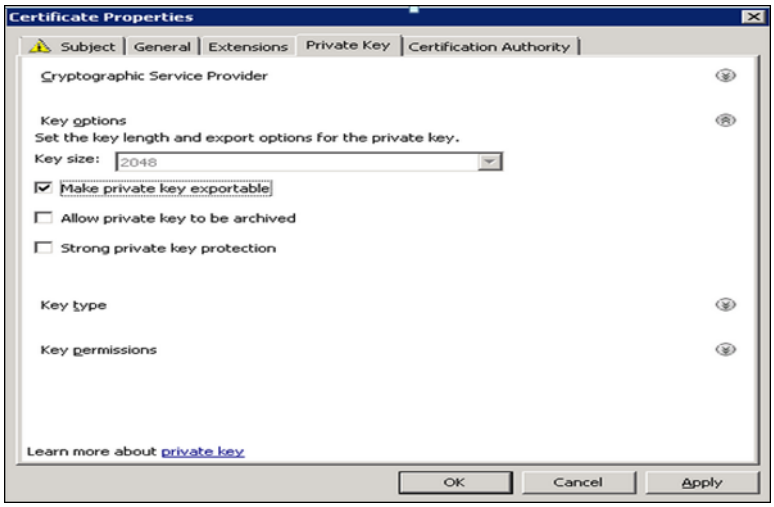

20. Export certificate created for view connection server and Replica server. Copy them to their corresponding server.
21. Go to View connection server/Replica server. Start Menu > Run > mmc.
22. Click File and Select Add/remove Snap-in.

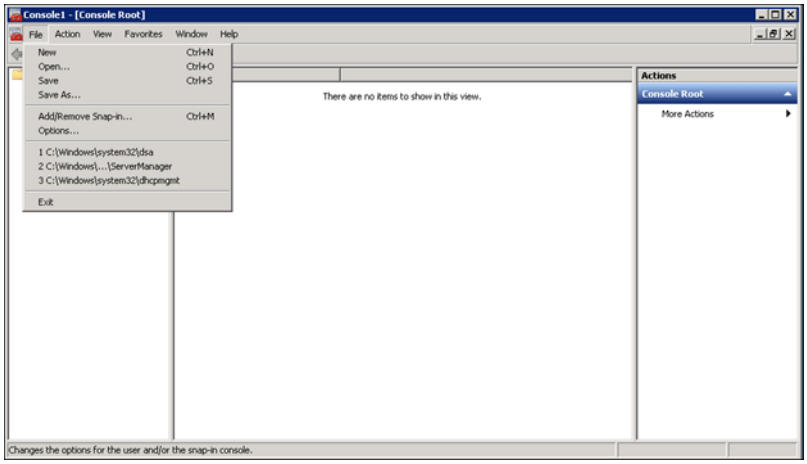

23. Select Certificate from the Available snap-ins on the left side and click Add.

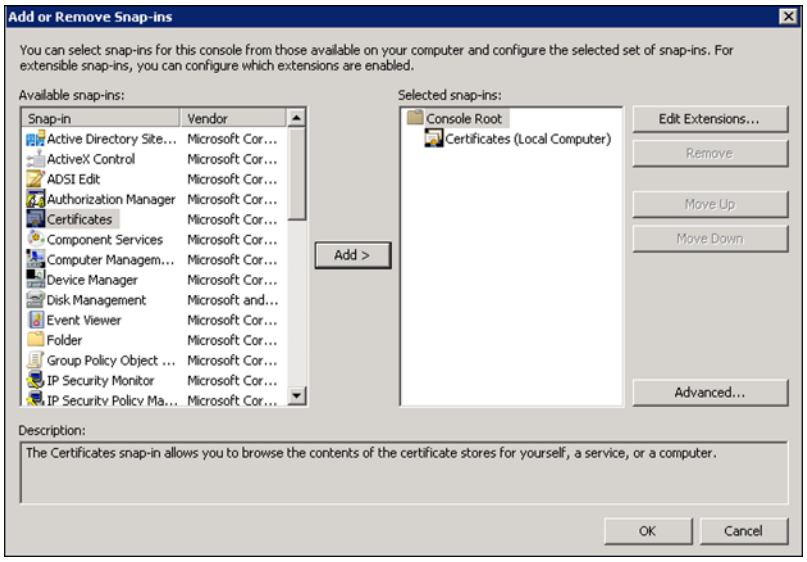

24. Select the radio button for Computer account.

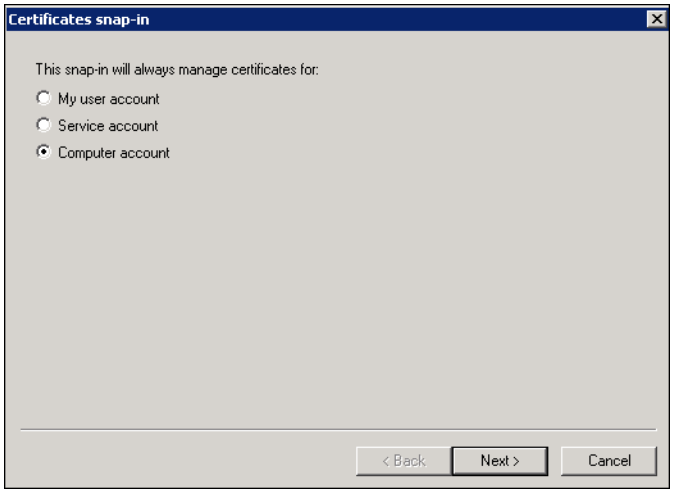

25. Select the radio button for Local computer. Click Finish.

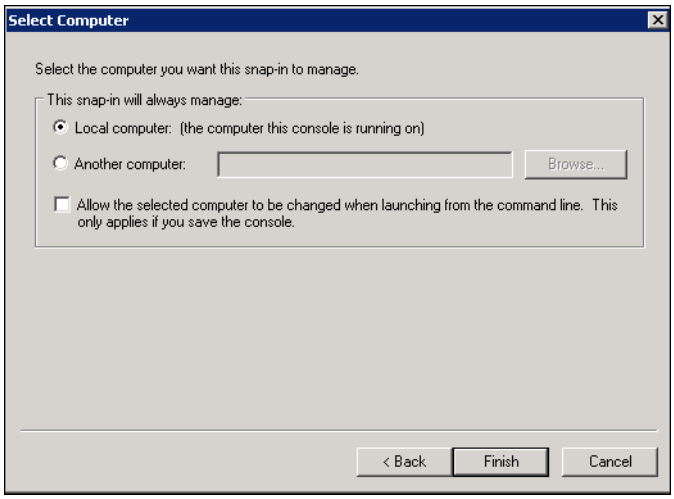

26. Select Certificates on the Console Root; Select Personal > Certificates > All Tasks > Import.

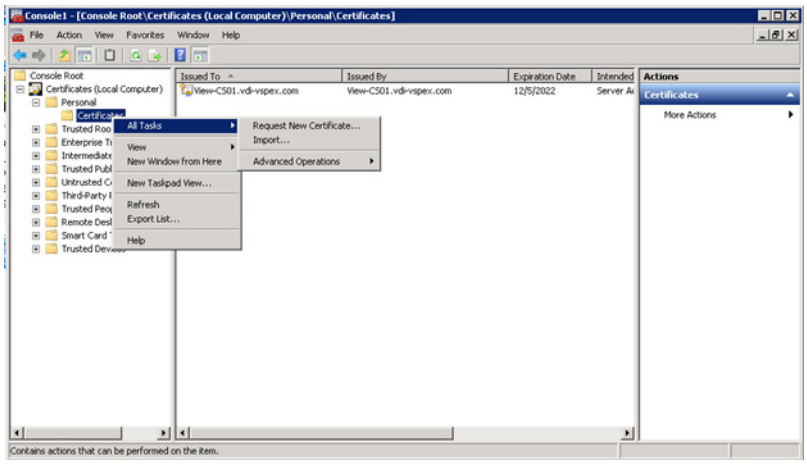

27. Browse and select copied certificate for view connection server and follow the same for view Replica server.

28. Select the previous installed certificate and change the friendly name. Replace the newly created certificate with vdm as the friendly name.

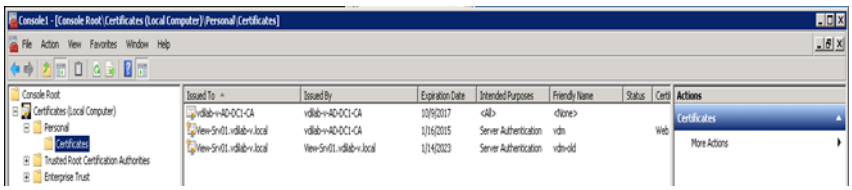

For more information about VMare's best practices to obtain the self-signed certificate, go to: http://pubs.vmware.com/view-52/topic/com.vmware.ICbase/PDF/horizon-view-52-obtaining-certificates.pdf

Install View Client on End Points

1. Download the installer file from the link given below:

https://my.vmware.com/web/vmware/info/slug/desktop_end_user_computing/vmware_horizon_view/5_2?rct=j&q=cmware%20view%205.2%20horizon%20download&source=web&cd=1&ved=0CEIQFjAA&url=http://www.vmware.com/go/downloadview

2. Open the installer file for 32-bit or 64-bit OS and click Next.

3. Read the VMWare EndUser Lincednse Agreement and click Next.

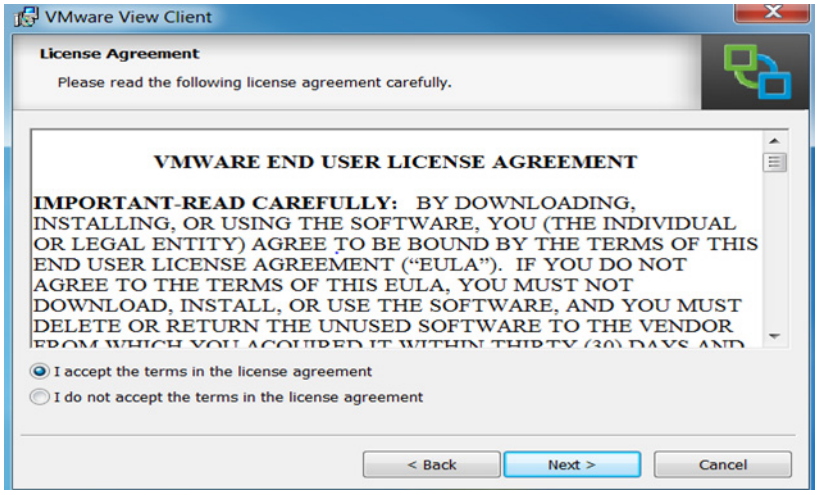

4. Click Next.

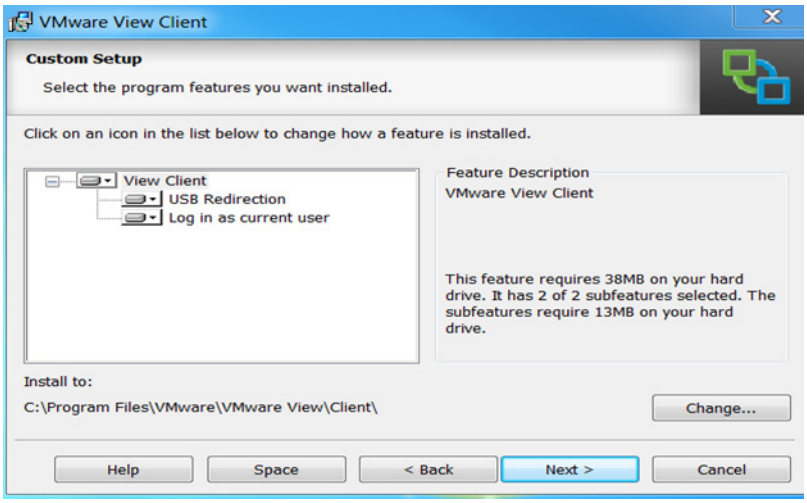

5. Enter FQDN for View Connection server ad click Next.

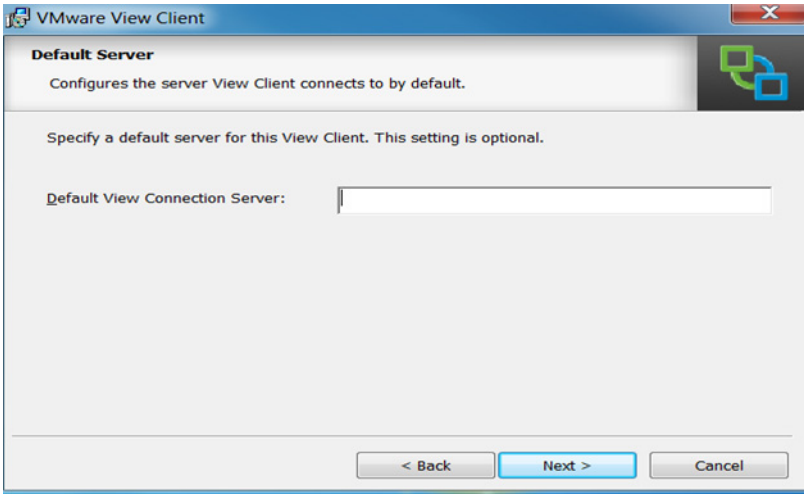

6. Accept the default or add the fqdn of your View Connection Server and click Next.

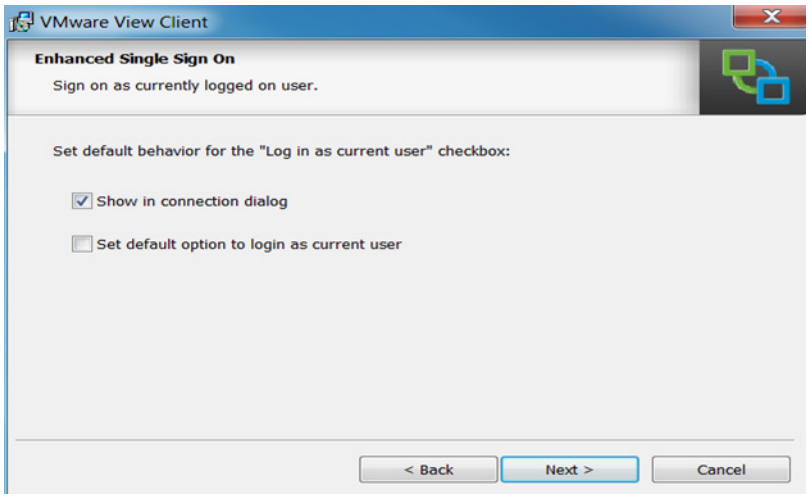

7. Click Next.

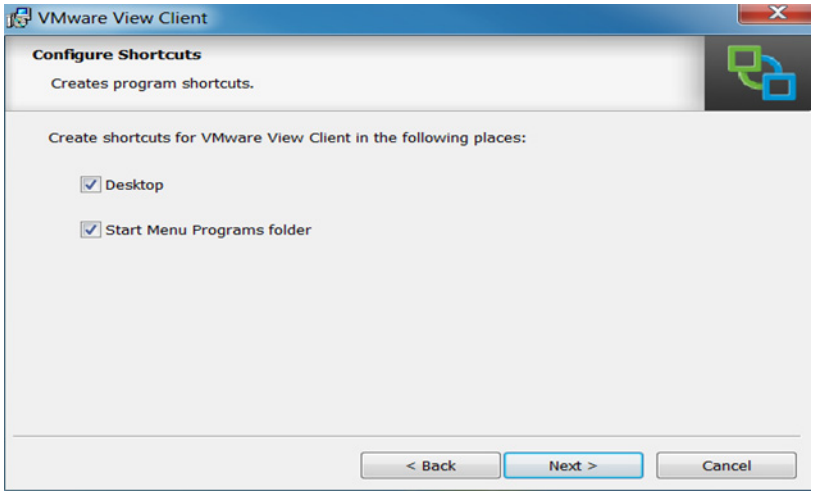

8. Click Install.

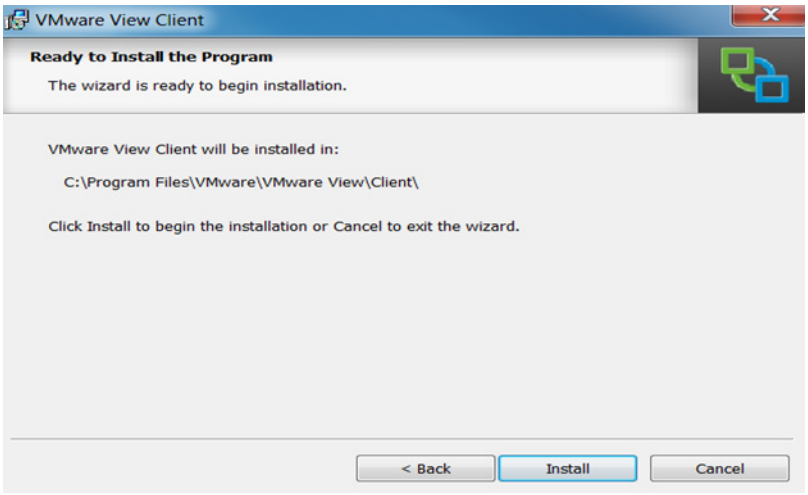

9. Reboot is required after completing the installation.

Configure the View 5.2 Hosts and Storage

Configure Content Based Read Cache (CBRC) on View 5.2 Hosts

CBRC was introduced as a feature of vSphere 5. It is a read cache that is particularly useful during boot storms. It becomes an essential configuration for floating assignment View 5.2 Linked Clones.

The CBRC feature provides a per-host RAM-based solution for View desktops. This considerably reduces the read I/O requests that are issued to the storage layer, and also addresses boot storm snags.

CBRC is configured in vCenter by highlighting the host; access the Configuration Tab, Software, and Advanced Settings.

Each ESXi host used for View Desktops we enabled CBRC and increased the CBRC.DCacheMemReserved to 2048.

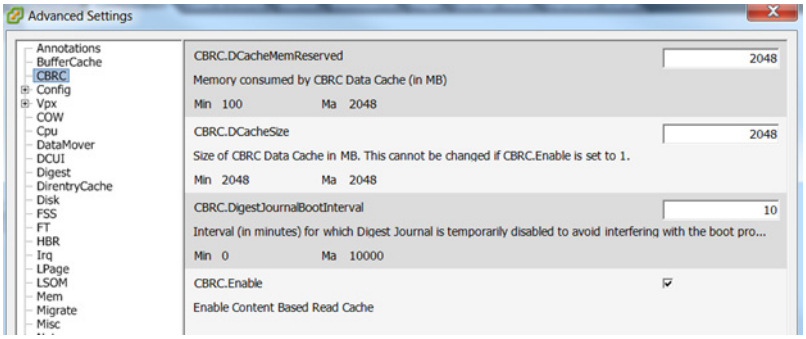

These CBRC settings are used in conjunction with the View 5.2 Administrator, View Configuration, Servers, vCenter Server Properties, Host Cashing tab.

In our test environment, we enabled 2GB of CBRD and correspondingly, 2GB of Host Cache in View Administrator. This combination enables the View Storage Accelerator feature.

Storage Configuration for View 5.2 Hosts

On VNX 5500 30 SAS disks with 300 GB capacity were used to create 16 LUNs, each with a capacity of 370 GB. 2 LUNs with capacities of 50 GB each were created to store replica disks.

Each ESXi host in each cluster was assigned 8 LUNs as VMFS5 datastores for linked clones and one 50Gb VMFS5 datastore to hold the Replica disk.

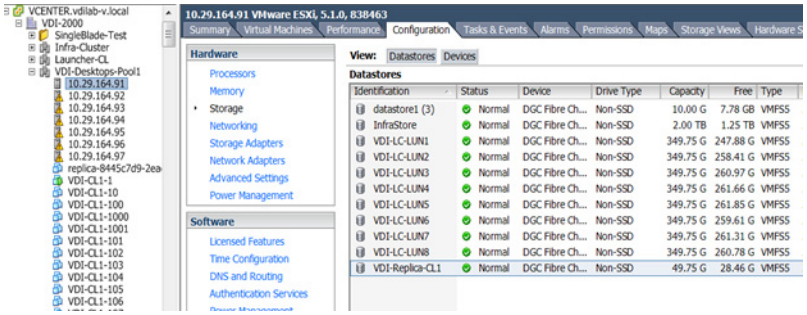

Note The same configuration was done for second Cluster.

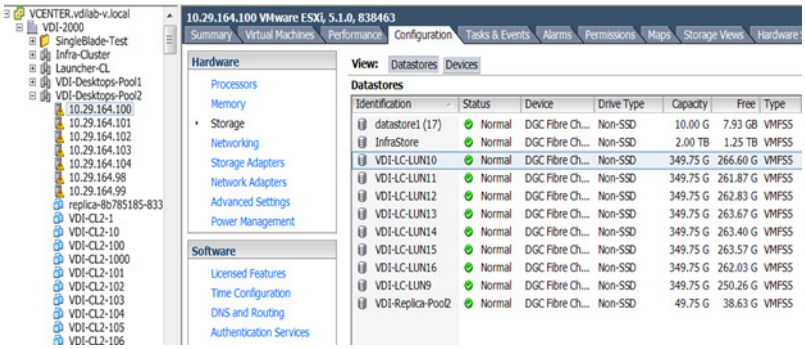

Desktop Delivery Base Image Creation and Desktop Deployment

Microsoft Windows 7 Golden Image Creation

Create base Windows SP1 Virtual Machine

1. Select ESXi host in Infrastructure cluster and create a virtual machine to use as Golden Image with windows 7 OS. MS Windows 7 32 bit OS was used for testing.

   For the virtual machine following parameters were used:

   Memory: 1536Mb

   Processor: 1vCPU

   Hard Disk: 18Gb

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

2. Right-click Windows 7 Golden Image properties and select Hardware TAB to attach the Windows -7 SP 1 ISO.

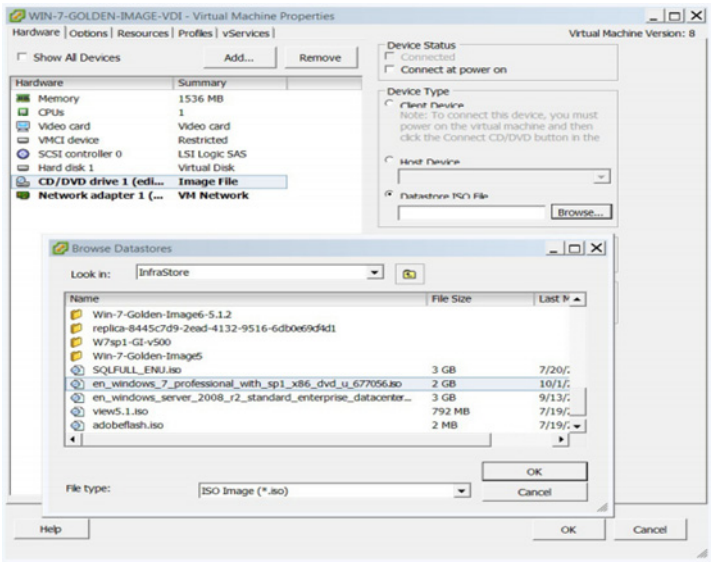

3. Click OK.
4. Right-click Windows 7 Golden Image Properties and click Edit Setting. Click the Options tab
   a. Go to the Options tab.
   b. Select Boot Options and check box for Force BIOS Setup
   c. Click OK and complete installation.
5. After the installation, log in to Windows 7 Golden Image virtual machine and configure IP Address, join the domain and Restart the Virtual Machine.
6. Install Windows Upates, then disable the Windows Update service on the Golden Image machine.

This complete the process of creating the Golden Image virtual machine.

Optimization of Base Windows 7 SP1 Virtual Machine

1. Click the link below for instructions about how to optimize MS Windows 7 SP1 32 bit virtual machine.

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

Install View 5.2 Virtual Desktop Agent Software

1. Download software from the following link:

   https://my.vmware.com/web/vmware/details?downloadGroup=VIEW-520-PREMIER&productId=320&rPId=4175

2. Open the installer VMware-viewagent-5.2.0-987719.exe for 32bit OS or VMware-viewagent-x86_64-5.2.0-987719.exe 64bit OS.
3. Click Next.

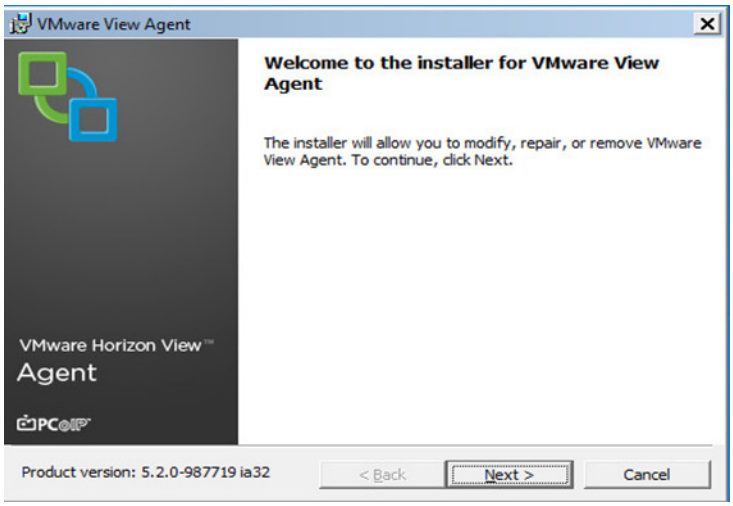

4. Read the VMware End User License Agreement, click I accept and click Next.

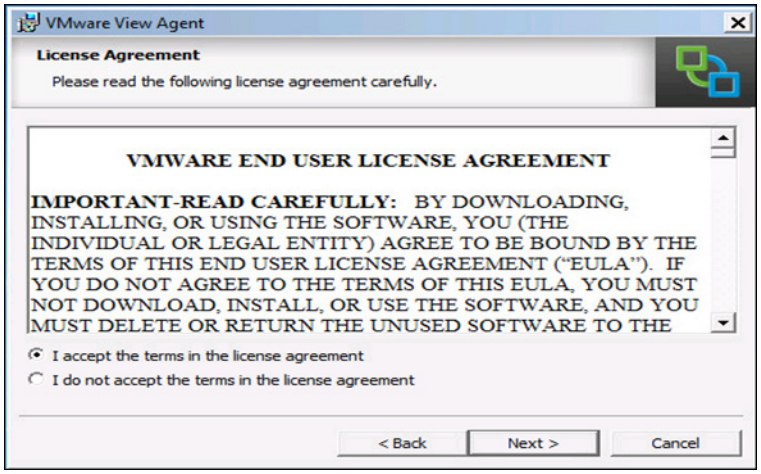

5. Click Next to accept the default setup.

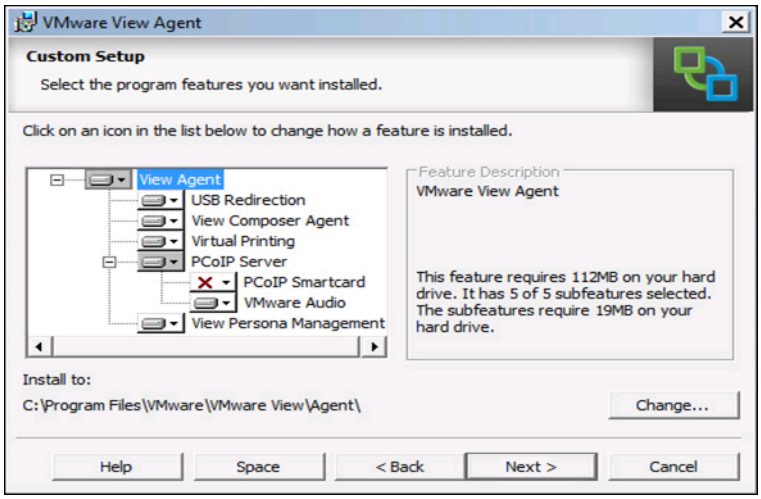

6. Click Install.

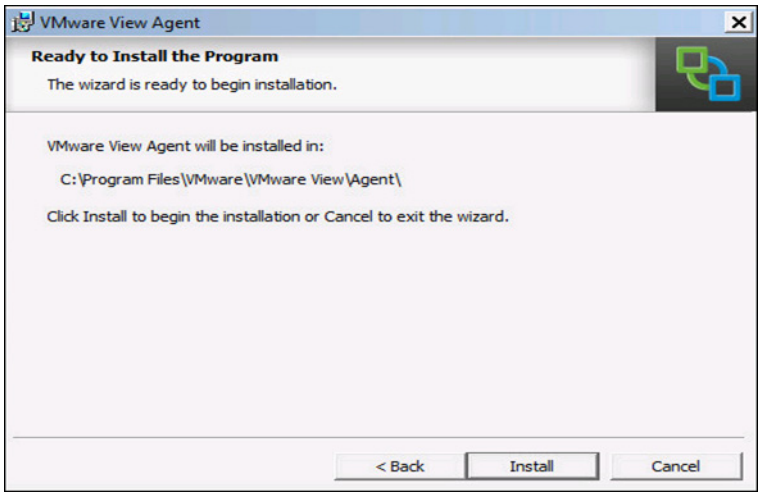

Install Additional Software

1. Install additional software required for your base windows image. We installed the following software:

- MS Office 2010 was installed for test environment.
- The VSI Target software package was installed to facilitate workload testing. (Optional)

2. Reboot the virtual machine.
3. Install the service packs and hot fixes required for the additional software components that were added.
4. Shut down the virtual machine.

Perform Additional View 5.2 Configuration

Create a Snapshot for the Virtual Machine

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

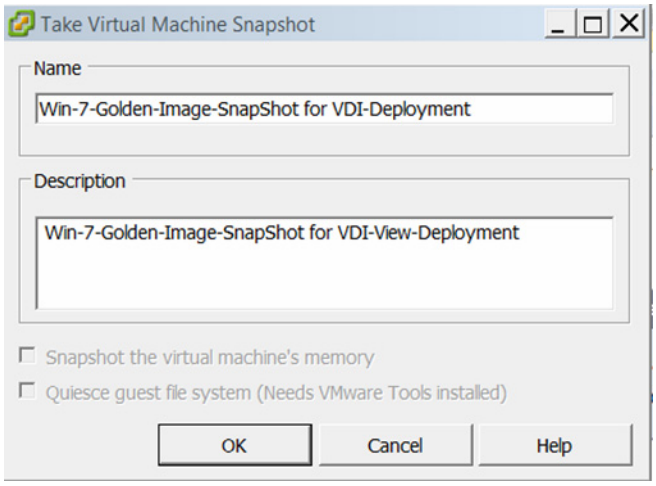

Create Customization Specification for Virtual Desktops

1. From vCenter, Right-click the powered off virtual machine after taking a snapshot and select Convert to Template.
2. Provide a name for the template and provide the host /cluster, data store details.
3. Convert the template back to a virtual machine or deploy a virtual machine from the template.
4. Provide a name and select the data center, click Next.
5. Select the cluster for the virtual machine.
6. Select the virtual machine host.
7. Select the datastore.
8. Select Guest Customization and check the radio button for Customize using the Customization Wizard.
9. Click Next.

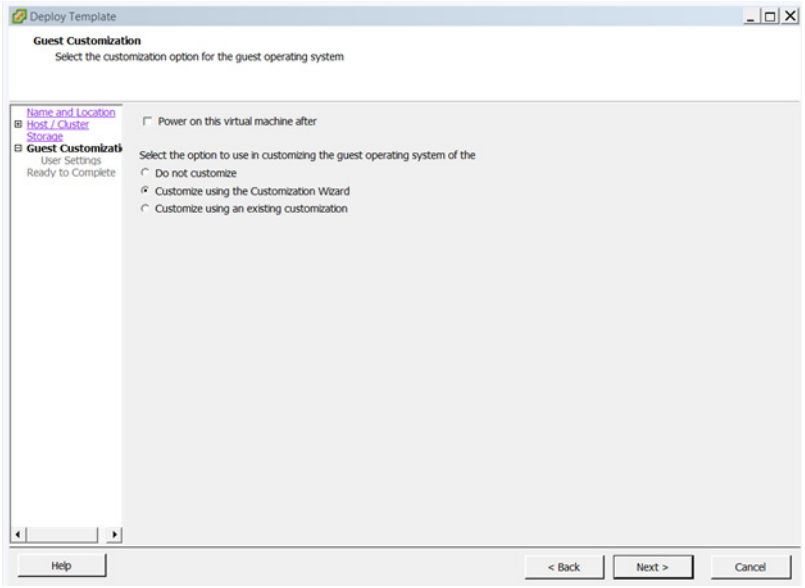

10. Select the appropriate name and organization. Click Next.

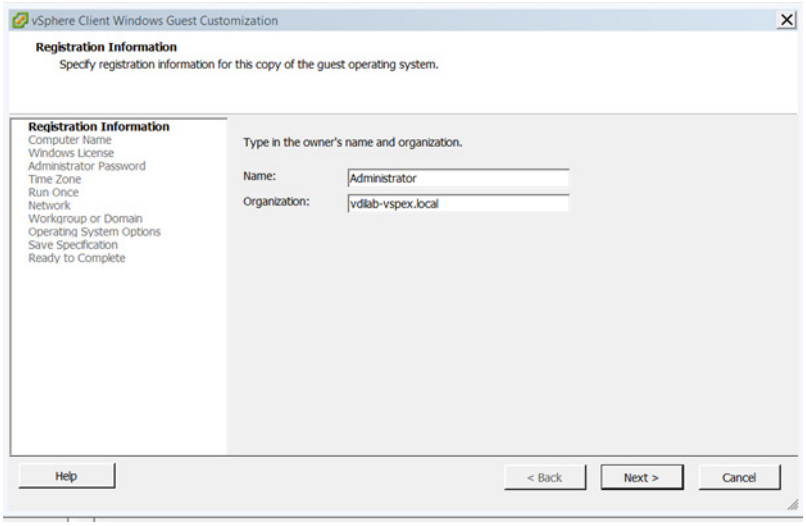

11. Select the radio button for Use the virtual machine name. Click Next.

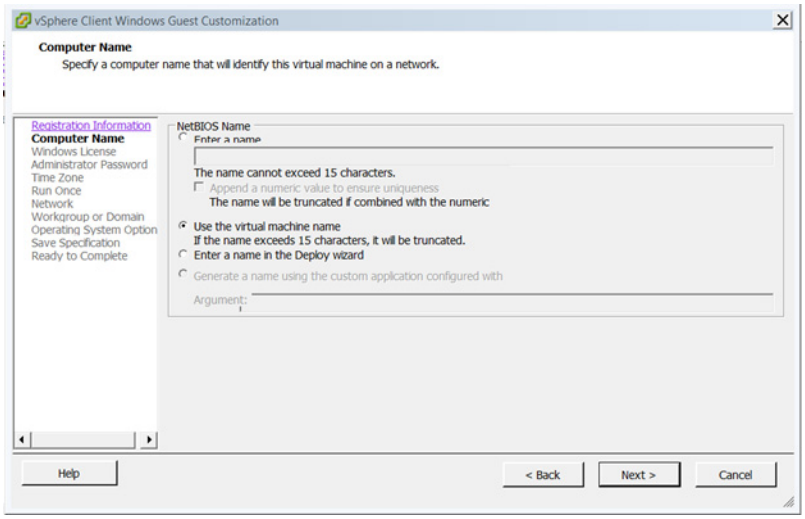

12. Specify the Volume License Key for Windows 7 and select Per seat or Per server Maximum. Click Next.

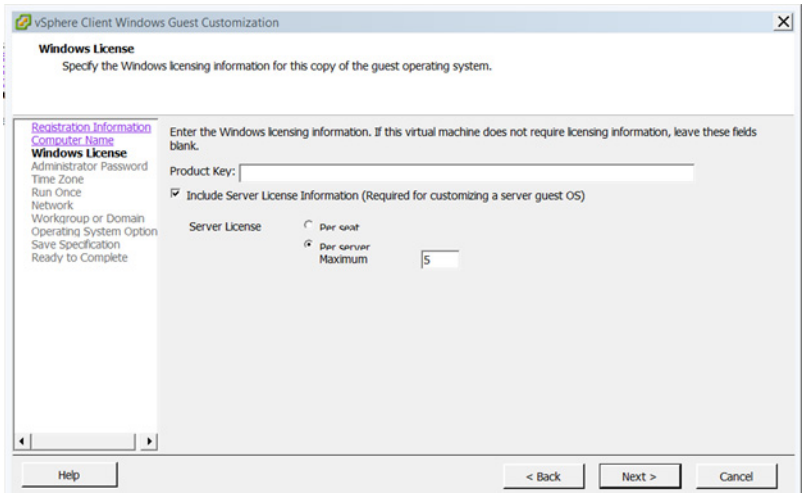

13. Enter the credential for the administrator account. Click Next.

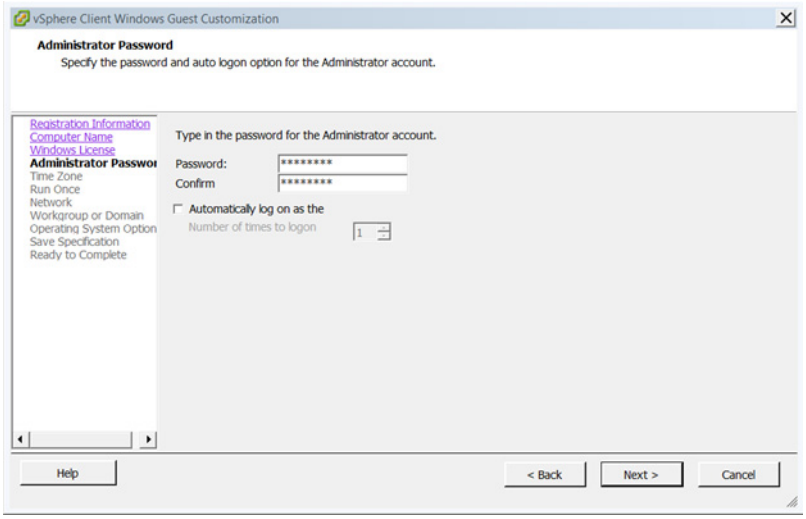

14. Select the appropriate time zone, Click Next.

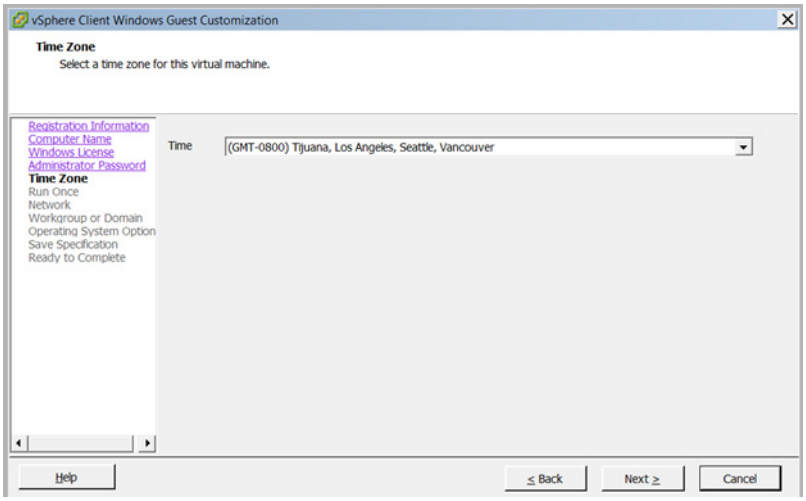

15. For Network, select Typical Setting for the virtual desktop networking. Click Next.

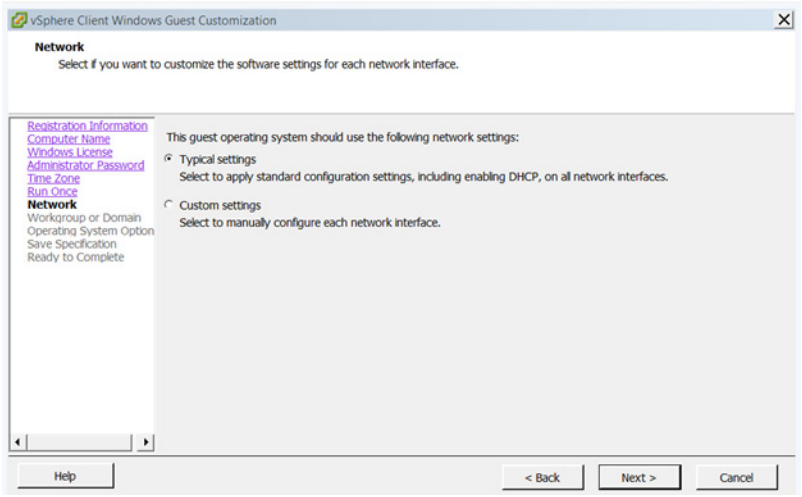

16. For Workgroup or Domain, select the radio button for Windows Server. Enter the domain for environment.
17. Specify the user account name password. Click Next.

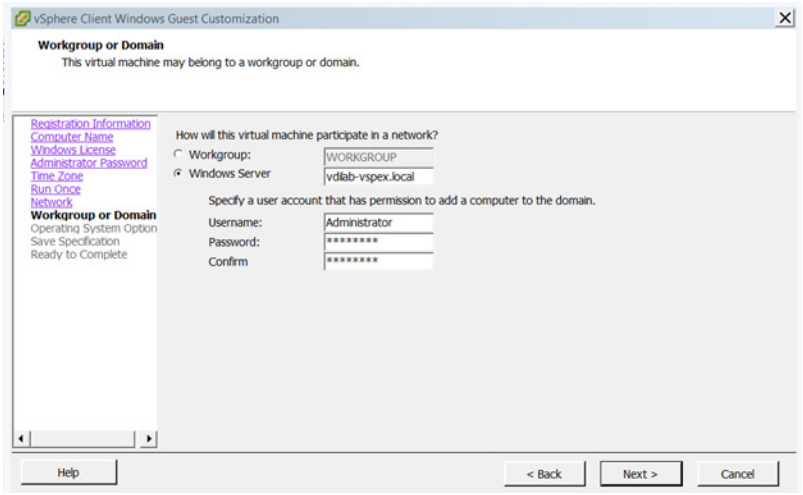

18. Select the checkbox to generate the new security ID. Click Next.

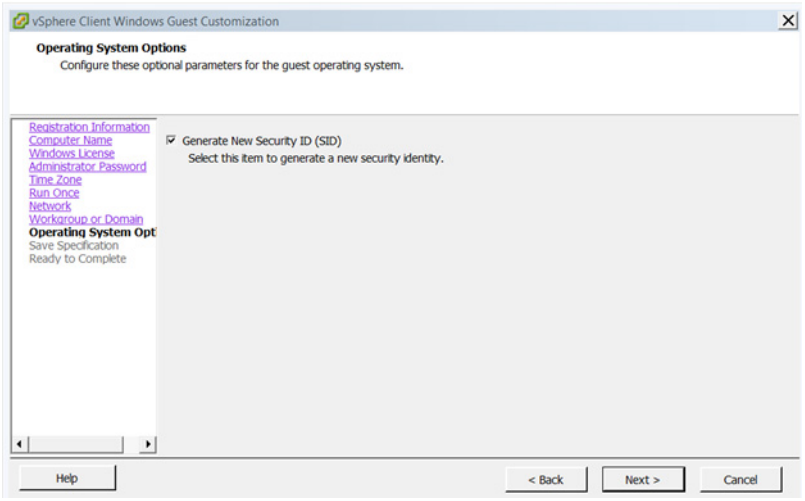

19. Select the checkbox to Save the customization specification for later use and provide a name.

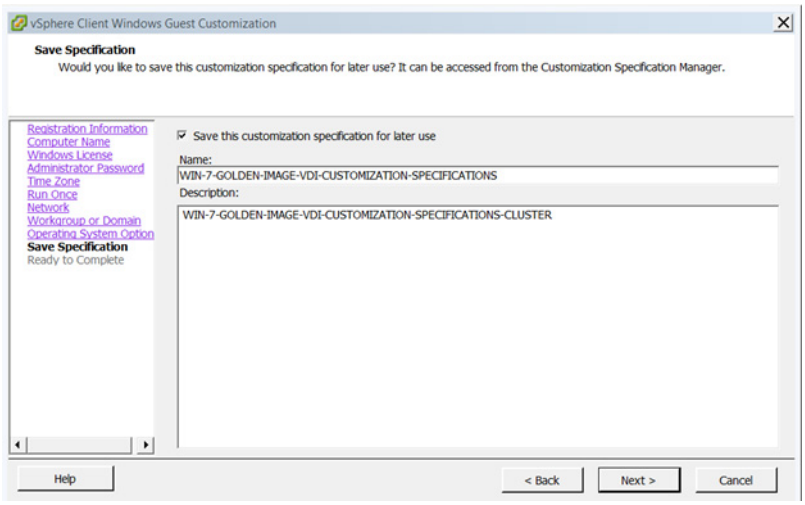

20. Verify and click Finish.

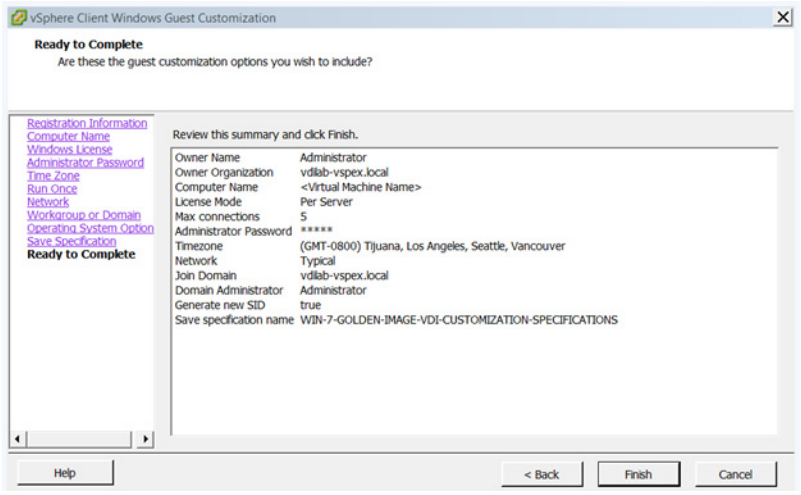

21. To edit or modify the customization specification, log into vCenter Client with vCenter server IP and credentials. Go to the Home screen and select Customization Specification Manager. Select Saved customization, right -click and select edit.

Configure the View Desktop Pools and Options.

Desktop Pools are the containment object in View 5.2 Administrator that hold the configuration and the provisioned linked clones in the View environment.

The maximum recommended number of virtual machines in a VMware ESXi cluster is 1000. Therefore, we created two View 5.2 pools with identical settings to match up with our two ESXi 5.1 clusters for VDI described earlier in this document.

The following sections describe how we configured our View 5.2 environment.

Create the Desktop Pools

1. Log into View Administrator console. on the left side; from drop menu for Inventory select Pools. Click Add to create a new desktop pool.

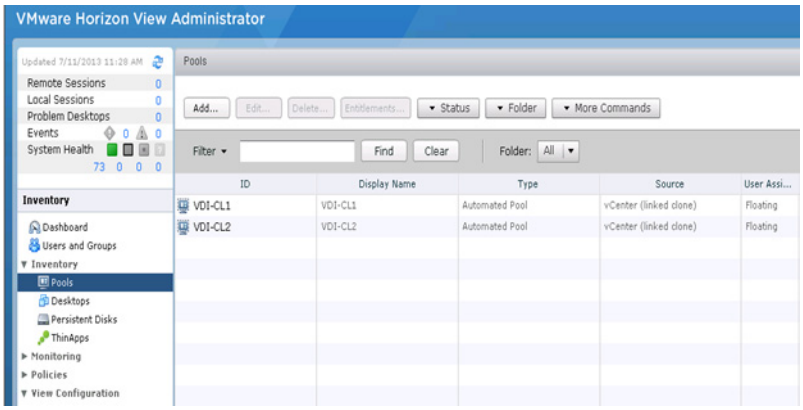

There are three types of Desktop Pools you can create and descriptions for each type is given on the right side of the screen.

2. Select Automated Pool and click Next.

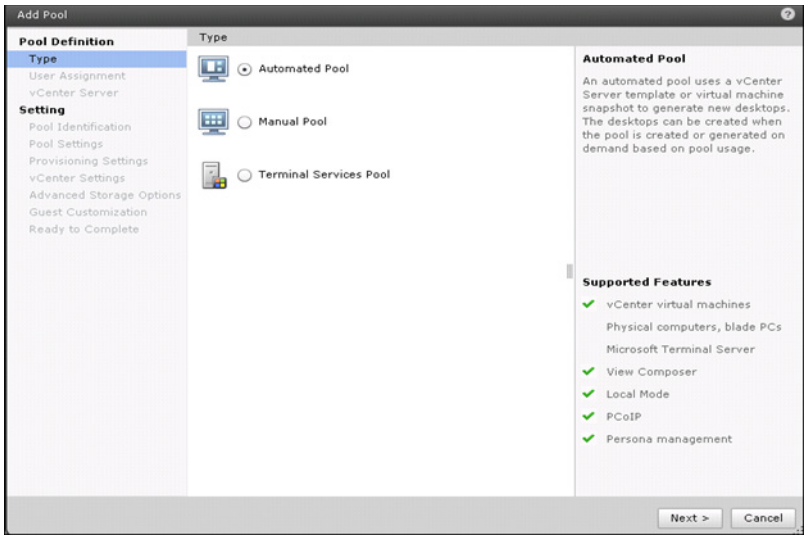

There are two types of User Assignment options available:

- Dedicated (desktops that are manually or automatically assigned to users)
- Floating (desktops that are randomly assigned to users from the pool.

Note For the purpose of testing, the Floating user assignment is used.

3. Select Floating. Click Next.

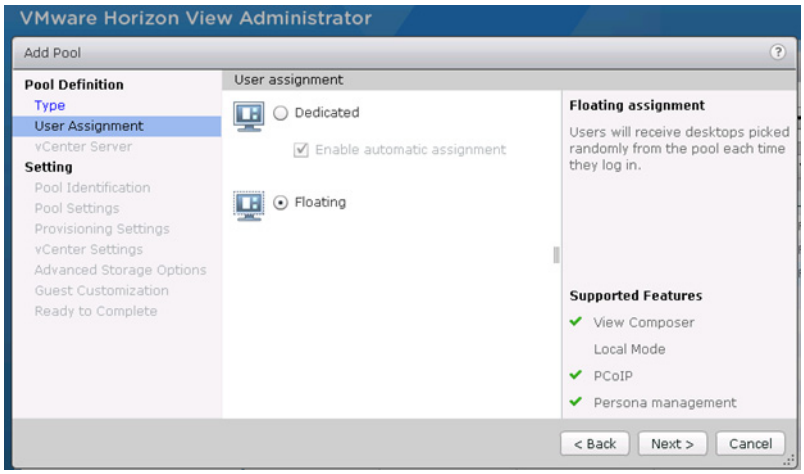

4. Select the radio button for either Full Vitual Machine or view composer linked clones. Click Next.

Note For this study, we chose View Composer linked clones.

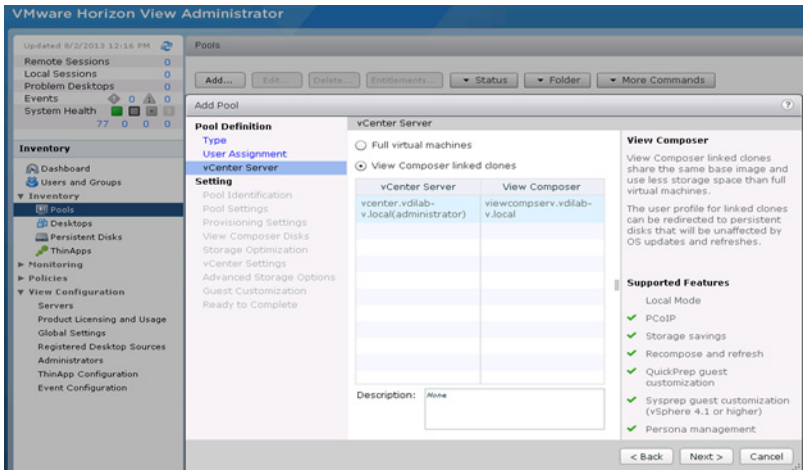

5. Enter a unique pool ID and Display name optionally, select a folder for the Desktops. Click Next.

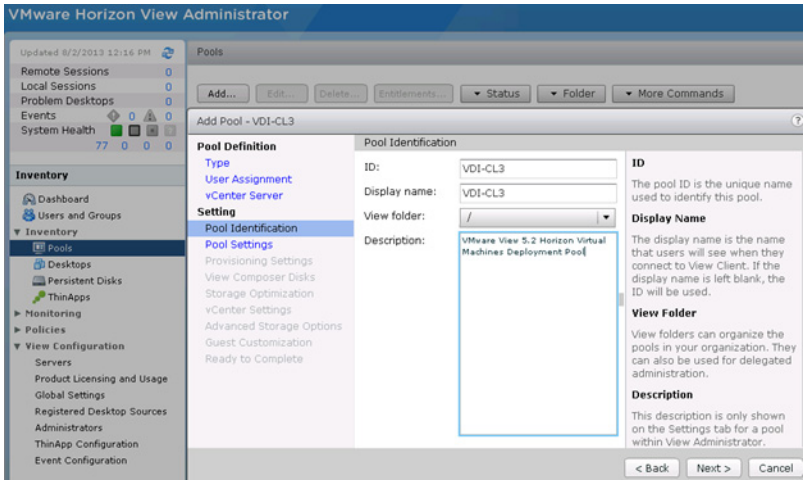

6. Configure the Pool Settings as needed. Select all default settings except the Remote Desktop Power Policy. Select Ensure deskstops are always powered on. Click Next.

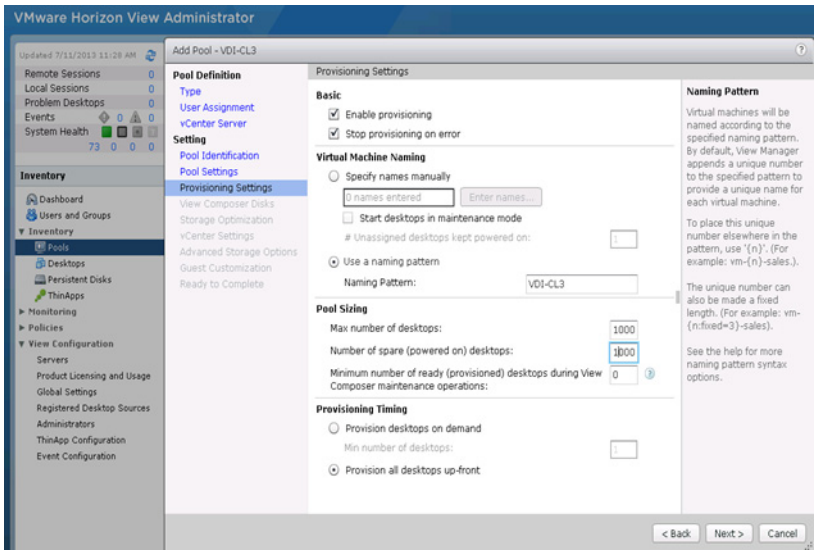

7. On the Provisioning Settings page, set the following options:

- Basic: Enable provisioning and Stop provisioning at error.
- Virtual Machine Naming: Use naming pattern.

Note Use {n} to deploy multiple desktops with same naming pattern. In case of name used VM-{n} deployed desktops will be VM-1, VM-2 ….. VM-10

- Pool Sizing: Select maximum number of desktops, number of powered on desktops and how to provision the desktops

Note For this study, we provisioned the entire desktops up front.

8. Select the radio button to Redirect disposable files to a non-persistent disk, set the Drive size for the disk, and select a Drive. Click Next.

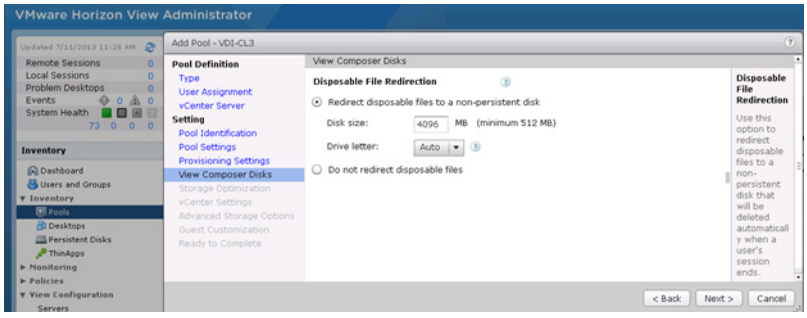

9. Check the box for Select separate datastore for replica disk and OS disk. Click Next.

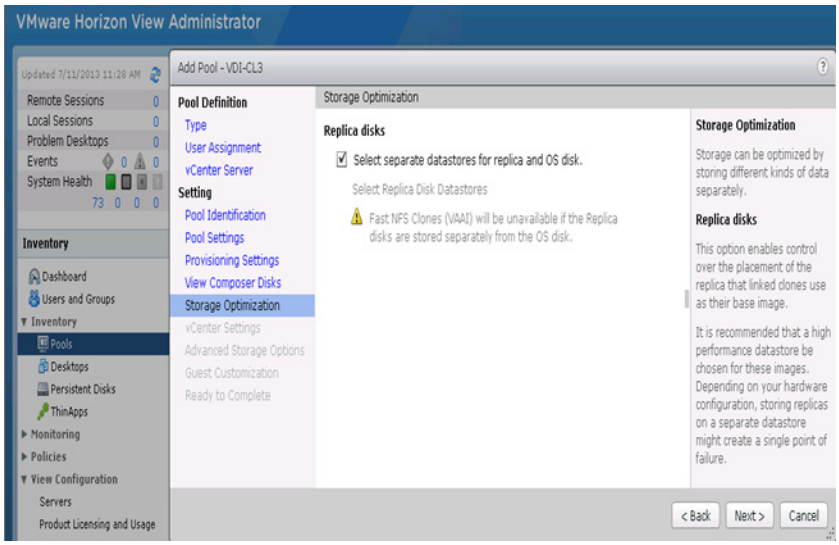

10. Select parent image (Golden Image), associated snapshot with GI image, location for VM if any specific folder was created, Host or Cluster where desktops are going to provision, Resource Pool, Linked Clone datastore, Replica disk datastores.

Note For our testing we created 2 Pools with 1000 desktops in each. One Pool was created with Cluster1 as host Resource Pool. 8 VMFS5 datastores for Linked Clones and one Replica disk datastore were selected. The second Pool was configured similarly with the remaining 8 VMFS5 datastores for Linked Clones and the remaining Replica VMFS5 datastore for the Replica disk.

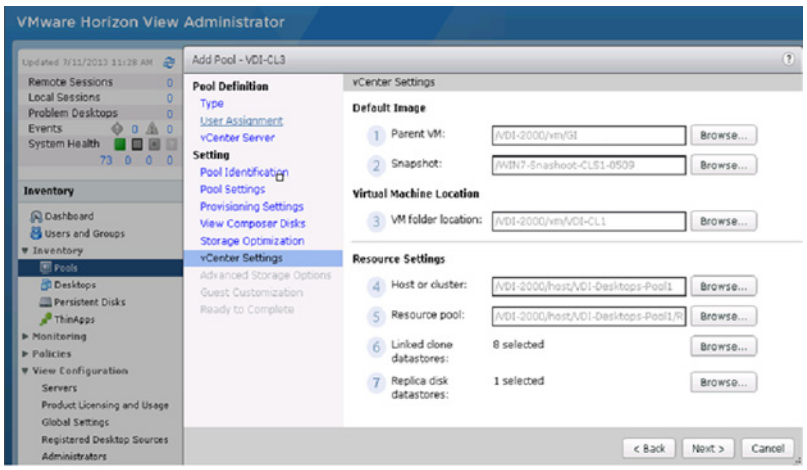

11. On the Advanced Storage Options page, check box to enable Use Storage Accelerator and add Blackout time if necessary. Click Next.

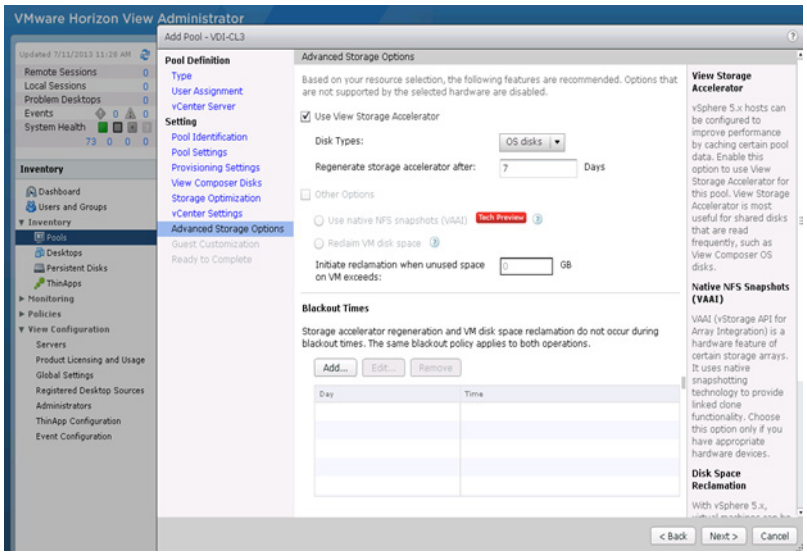

12. Browse to select the AD container to be used for the provisioned machines. Select the radio button for Use a customization specification and select customization created from the Parent Windows 7 image VM.

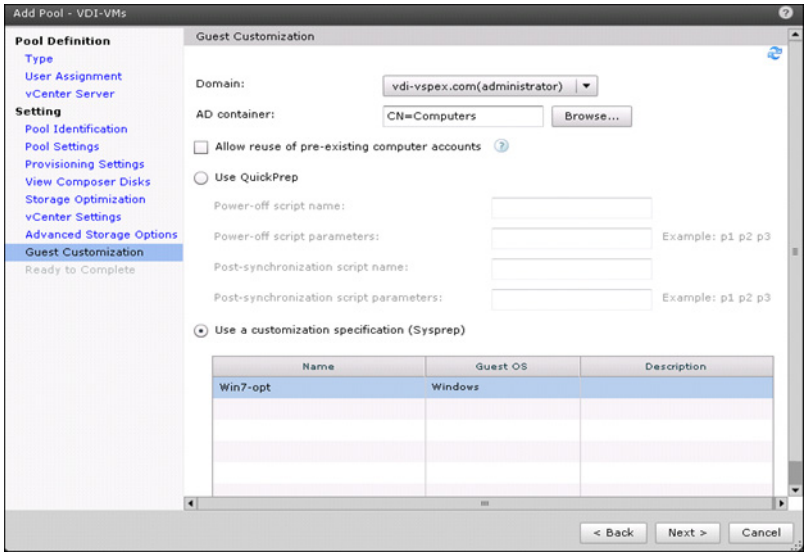

13. Verify all the details provided for the pool settings and check box to entitle specific users and groups to provide access to the desktops in the Pool and click Finish.

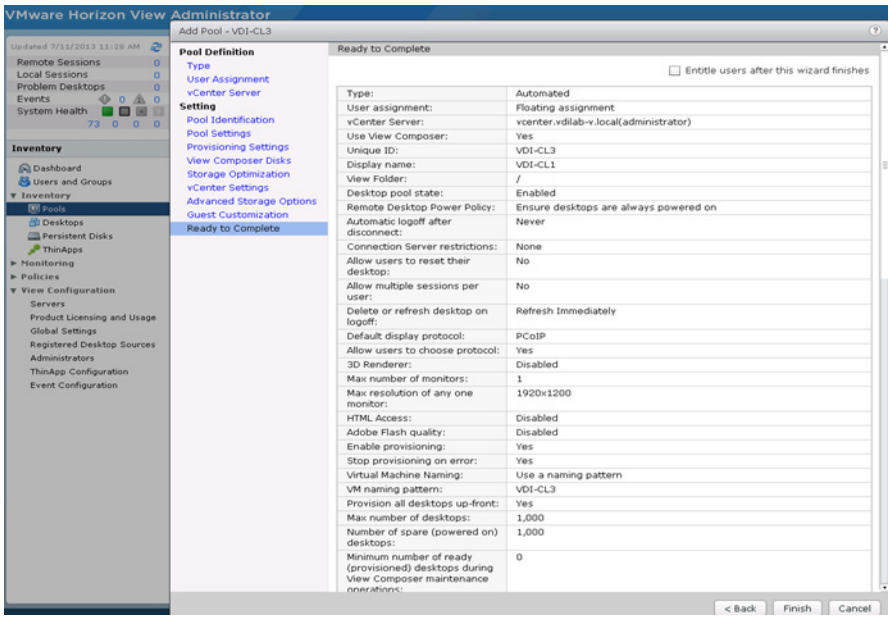

14. On the Entitlements page, click Add.

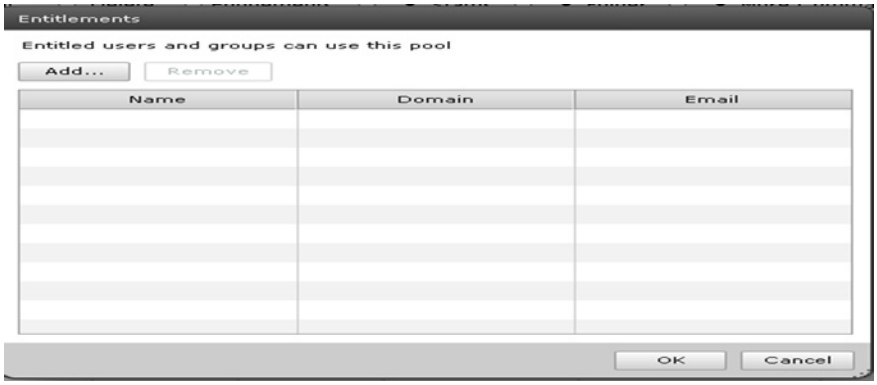

15. Enter a name for the users or groups who will be authorized to use View desktops in the pool. Click Find.
16. Select the appropriate users and group from the list. Click OK.

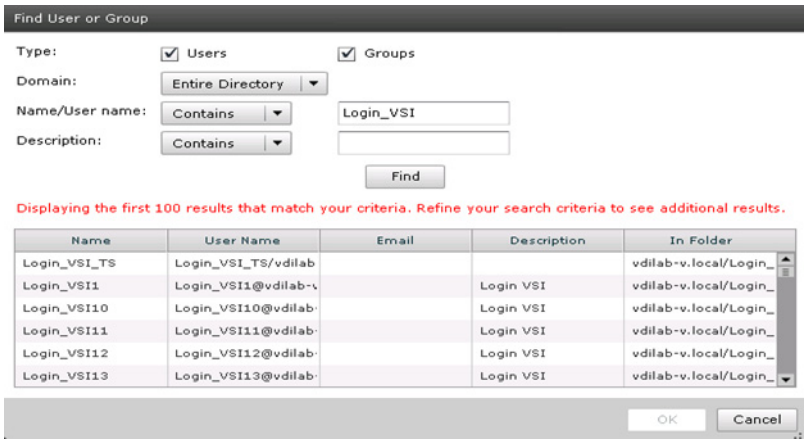

17. When the pool is Enabled and has Entitlements, both the columns will turn green.

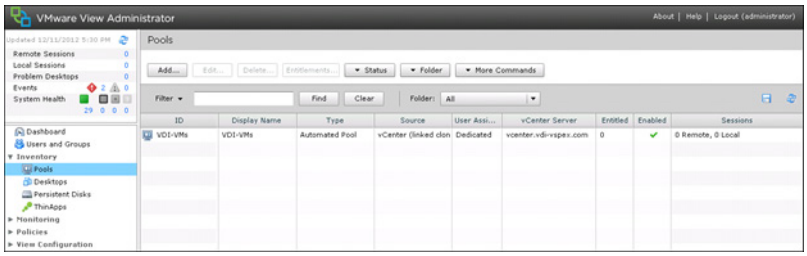

After the completion of the pool setting, a replica from the parent VM is created and the provisioning of the desktops as per the pool settings begins.

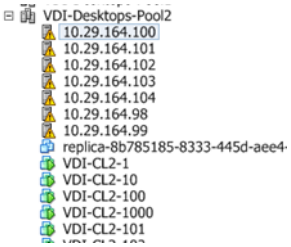

Test Setup and Configurations

In this project, we tested a single Cisco UCS B200 M3 blade in a single chassis and fourteen Cisco UCS B200 M3 blades in two chassis to illustrate linear scalability.

Cisco UCS Test Configuration for Single Blade Scalability

Figure 18 illustrates the Cisco UCS test configuration single-blade scalability.

Figure 18 Cisco UCS B200 M3 Blade Server for Single Server Scalability

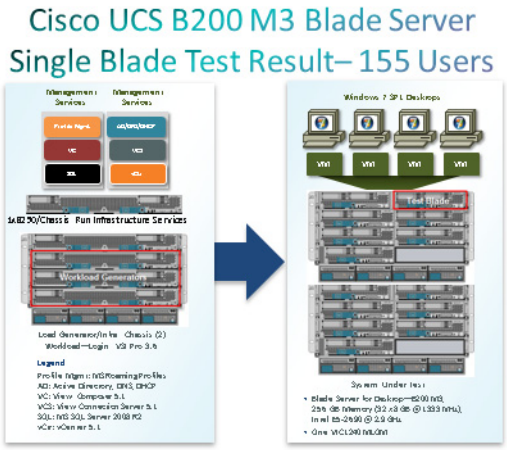

Hardware components

- 1 X Cisco UCS B200-M3 (2 X E5-2690 @ 2.9 GHz) blade server with 256GB of memory (16 GB X 16 DIMMS @ 1666 MHz) Windows 7 SP1 Virtual Desktop hosts
- 2 X Cisco UCS B250-M2 (2 X 5680 @ 3.333 GHz) blade servers with 96 GB of memory (4 GB X 24 DIMMS @ 1333 MHz) Infrastructure Servers
- 4 X Cisco UCS B250-M2 (2 X 5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) Load Generators
- 2 X M81KR (Palo) Converged Network Adapter/Blade (B250 M2)
- 1X VIC1240 Converged Network Adapter/Blade (B200 M3)
- 2 X Cisco Fabric Interconnect 6248UPs
- 2 X Cisco Nexus 5548UP Access Switches
- 1 X EMC VNX System storage array, two controllers, two Datamovers, 2 x dual port 8GB FC cards, 2 x dual port 10 GbE cards, 4 x 200GB Flash Drives for EMC Fast Cache, 30 x 300GB SAS drives for VMFS datastores, 8 x 600GB SAS Drives for Infrastructure and Boot LUNs and 2 x 300GB SAS drives and 1 200GB Flash Drive for hot spares

Software components

- Cisco UCS firmware 2.1(1a)
- Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.1 for VDI Hosts
- Horizon View 5.2
- Windows 7 SP1 32 bit, 1vCPU, 1 GB of memory, 18 GB/VM

Cisco UCS Configuration for Two Chassis—Fourteen Blade Test

Figure 19 illustrates the two chassis test configuration.

Figure 19 Two Chassis Test Configuration-14 x B200 M3 Blade Servers

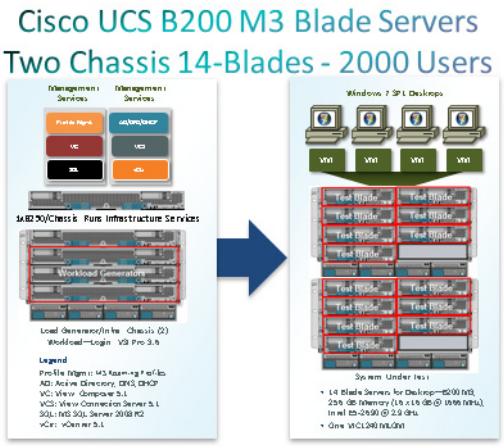

Hardware components

- 14 X Cisco UCS B200-M3 (2 X E5-2690 @ 2.9 GHz) blade server with 256GB of memory (16 GB X 16 DIMMS @ 1666 MHz) Windows 7 SP1 Virtual Desktop hosts
- 2 X Cisco UCS B250-M2 (2 X 5680 @ 3.333 GHz) blade servers with 96 GB of memory (4 GB X 24 DIMMS @ 1333 MHz) Infrastructure Servers
- 6 X Cisco UCS B250-M2 (2 X 5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) Load Generators
- 2 X M81KR (Palo) Converged Network Adapter/Blade (B250 M2)
- 1X VIC1240 Converged Network Adapter/Blade (B200 M3)
- 2 X Cisco Fabric Interconnect 6248UPs
- 2 X Cisco Nexus 5548UP Access Switches
- 1 X EMC VNX System storage array, two controllers, two Datamovers, 2 x dual port 8GB FC cards, 2 x dual port 10 GbE cards, 4 x 200GB Flash Drives for EMC Fast Cache, 30 x 300GB SAS drives for VMFS datastores, 8 x 600GB SAS Drives for Infrastructure and Boot LUNs and 2 x 600GB SAS drives and 1 200GB Flash Drive for hot spares

Software components

- Cisco UCS firmware 2.1(1a)
- Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.1 for VDI Hosts
- Horizon View 5.2
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 18 GB/VM

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco Labs with joint support from both Cisco and EMC resources.

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 Hosted VDI model under test.

Test metrics were gathered from the hypervisor, virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the Horizon View 5.2 environment and executing a typical end-user workflow. To generate load within the environment, an auxiliary software application was required to generate the end user connection to the Horizon View environment, to provide unique user credentials, to initiate the workload, and to evaluate the end user experience.

In the Hosted VDI test environment, sessions launchers were used simulate multiple users making a direct connection to Horizon View 5.2 via a VMware Horizon View PCoIP protocol connection.

User Workload Simulation—LoginVSI From Login Consultants

One of the most critical factors of validating a Horizon View deployment is identifying a real-world user workload that is easy for customers to replicate and standardized across platforms to allow customers to realistically test the impact of a variety of worker tasks. To accurately represent a real-world user workload, a third-party tool from Login Consultants was used throughout the Hosted VDI testing.

The tool has the benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for individual desktop throughout large scale testing, including login storms.

The Virtual Session Indexer (Login Consultants' Login VSI 3.6) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and allows customers to easily replicate the testing results in their environment.

Note In this testing, we utilized the tool to benchmark our VDI environment only.

- Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.
- Login VSI simulates a medium workload user (also known as knowledge worker) running generic applications such as: Microsoft Office 2007 or 2010, Internet Explorer 8 including a Flash video applet and Adobe Acrobat Reader (Note: For the purposes of this test, applications were installed locally, not streamed nor hosted on XenApp).

Like real users, the scripted Login VSI session will leave multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. This workload emulated a medium knowledge working using Office, IE, printing and PDF viewing.

- When a session has been started the medium workload will repeat every 12 minutes.

- During each loop the response time is measured every two minutes.
- The medium workload opens up to five apps simultaneously.
- The type rate is 160ms for each character.
- Approximately two minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007/2010, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and heavy
- 480 p Flash application gettheglass.com.
- Word 2007/2010, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007/2010, a very large randomized sheet is opened.
- PowerPoint 2007/2010, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

A graphical representation of the medium workload is shown below.

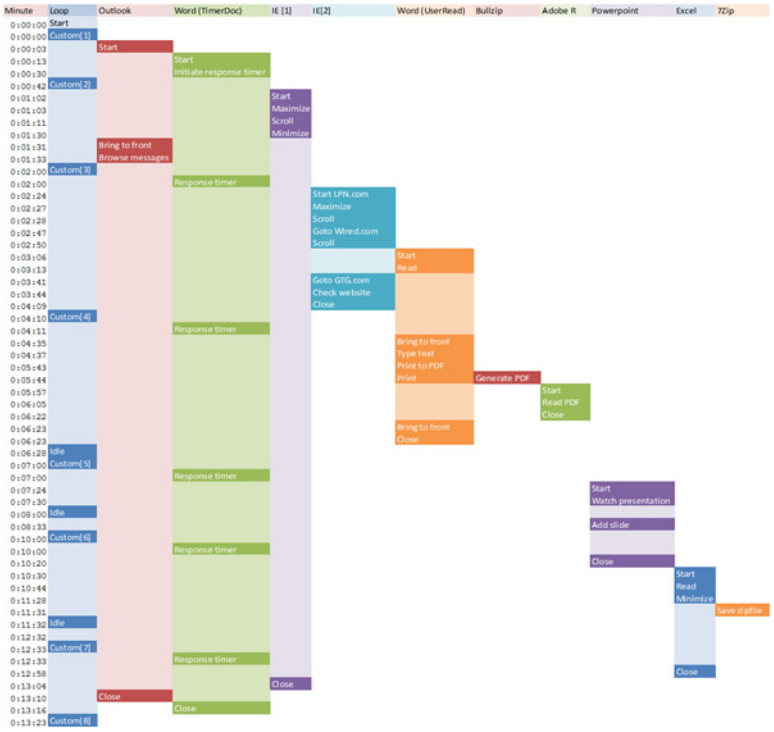

You can obtain additional information on Login VSI from http://www.loginvsi.com.

Testing Procedure

The following protocol was used for each test cycle in this study to help ensure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

- All virtual machines were shut down utilizing the vCenter.
- All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a "waiting for test to start" state.
- All VMware ESXi 5.1 VDI host blades to be tested were restarted prior to each test cycle.

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 30 minutes. Additionally, we require all sessions started, whether 155 single server users or 2000 full scale test users to become active within 2 minutes after the session is launched.

For each of the three consecutive runs on single blade (155 User) and 14-blade (2000 User) tests, the same process was followed:

1. Time 0:00:00 Started ESXtop Logging on the following systems:
   - VDI Host Blades used in test run
   - Profile Servers used in test run
   - SQL Servers used in test run
   - 7 or 80 Launcher VMs
2. Time 0:00:10 Started EMC Basic Performance Logging on SPs
3. Time 0:00:15 Started EMC NFS Performance Logging on Datamovers
4. Time 0:05 Take 155 or 2000 desktops out of maintenance mode on Horizon View Admin Console
5. Time 0:06 First machines boot
6. Time 0:33 155 or 2000 desktops booted on 1 or 14 blades
7. Time 0:35 155 or 2000 desktops available on 1 or 14 blades
8. Time 0:50 Start Login VSI 3.6 Test with 155 or 2000 desktops utilizing 7 or 80 Launchers
9. Time 1:20 155 or 2000 desktops launched
10. Time 1:22 155 or 2000 desktops active
11. Time 1:35 Login VSI Test Ends
12. Time 1:50 155 or 2000 desktops logged off
13. Time 2:00 All logging terminated

Success Criteria

There were multiple metrics that were captured during each test run, but the success criteria for considering a single test run as pass or fail was based on the key metric, VSI Max. The Login VSI Max evaluates the user response time during increasing user load and assesses the successful start-to-finish execution of all the initiated virtual desktop sessions.

Login VSI Max

VSI Max represents the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users as indicated during the workload execution. The user response time has a threshold of 4000ms and all users response times are expected to be less than 4000ms in order to assume that the user interaction with the virtual desktop is at a functional level. VSI Max is reached when the response times reaches or exceeds 4000ms for 6 consecutive occurrences. If VSI Max is reached, that indicates the point at which the user experience has significantly degraded. The response time is generally an indicator of the host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.

Note In the prior version of Login VSI, the threshold for response time was 2000ms. The workloads and the analysis have been upgraded in Login VSI 3 to make the testing more aligned to real-world use. In the medium workload in Login VSI 3.0, a CPU intensive 480p flash movie is incorporated in each test loop. In general, the redesigned workload would result in an approximate 20 percent decrease in the number of users passing the test versus Login VSI 2.0 on the same server and storage hardware.

Calculating VSIMax

Typically the desktop workload is scripted in a 12-14 minute loop when a simulated Login VSI user is logged on. After the loop is finished it will restart automatically. Within each loop the response times of seven specific operations is measured in a regular interval: six times in within each loop. The response times if these seven operations are used to establish VSImax.

The seven operations from which the response times are measured are:

- Copy new document from the document pool in the home drive
  - This operation will refresh a new document to be used for measuring the response time. This activity is mostly a file-system operation.
- Starting Microsoft Word with a document
  - This operation will measure the responsiveness of the Operating System and the file system. Microsoft Word is started and loaded into memory, also the new document is automatically loaded into Microsoft Word. When the disk I/O is extensive or even saturated, this will impact the file open dialogue considerably.
- Starting the "File Open" dialogue
  - This operation is handled for small part by Word and a large part by the operating system. The file open dialogue uses generic subsystems and interface components of the OS. The OS provides the contents of this dialogue.
- Starting "Notepad"
  - This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe itself through execution. This operation seems instant from an end-user's point of view.
- Starting the "Print" dialogue
  - This operation is handled for a large part by the OS subsystems, as the print dialogue is provided by the OS. This dialogue loads the print-subsystem and the drivers of the selected printer. As a result, this dialogue is also dependent on disk performance.
- Starting the "Search and Replace" dialogue
  - This operation is handled within the application completely; the presentation of the dialogue is almost instant. Serious bottlenecks on application level will impact the speed of this dialogue.

- Compress the document into a zip file with 7-zip command line
  - This operation is handled by the command line version of 7-zip. The compression will very briefly spike CPU and disk I/O.

These measured operations with 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 are 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. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

With Login VSI 3.0 and later it is now possible to choose between 'VSImax Classic' and 'VSImax Dynamic' results analysis. For these tests, we utilized VSImax Dynamic analysis.

VSIMax Dynamic

VSImax Dynamic is calculated when the response times are consistently above a certain threshold. However, this threshold is now dynamically calculated on the baseline response time of the test.

Five individual measurements are weighted to better support this approach:

- Copy new doc from the document pool in the home drive: 100%
- Microsoft Word with a document: 33.3%
- Starting the "File Open" dialogue: 100%
- Starting "Notepad": 300%
- Starting the "Print" dialogue: 200%
- Starting the "Search and Replace" dialogue: 400%
- Compress the document into a zip file with 7-zip command line 200%

A sample of the VSImax Dynamic response time calculation is displayed below:

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	33.3%	467
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	300%	150
Print Dialogue (PRINT)	220	200%	440
Replace Dialogue (FIND)	10	400%	40
Zip documents (ZIP)	130	200%	230
	VSImax Dynamic Response Time		**1837**

Then the average VSImax response time is calculated based on the amount of active Login VSI users logged on to the system. For this the average VSImax response times need to consistently higher than a dynamically calculated threshold.

To determine this dynamic threshold, first the average baseline response time is calculated. This is done by averaging the baseline response time of the first 15 Login VSI users on the system.

The formula for the dynamic threshold is: Avg. Baseline Response Time x 125% + 3000. As a result, when the baseline response time is 1800, the VSImax threshold will now be 1800 x 125% + 3000 = 5250ms.

Especially when application virtualization is used, the baseline response time can wildly vary per vendor and streaming strategy. Therefore it is recommend to use VSImax Dynamic when comparisons are made with application virtualization or anti-virus agents. The resulting VSImax Dynamic scores are aligned again with saturation on a CPU, Memory or Disk level, also when the baseline response time are relatively high.

Determining VSIMax

The Login VSI analyzer will automatically identify the "VSImax". In the example below the VSImax is 99. The analyzer will automatically determine "stuck sessions" and correct the final VSImax score.

- Vertical axis: Response Time in milliseconds
- Horizontal axis: Total Active Sessions

Figure 20 Sample Login VSI Analyzer Graphic Output

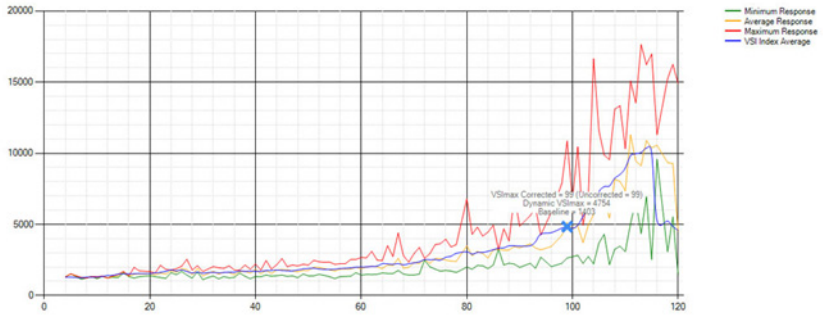

- Red line: Maximum Response (worst response time of an individual measurement within a single session)
- Orange line: Average Response Time within for each level of active sessions
- Blue line: the VSImax average.
- Green line: Minimum Response (best response time of an individual measurement within a single session)

Note We discovered a technical issue with the VSIMax dynamic calculation in our testing on Cisco B200 M3 blades where the VSIMax Dynamic was not reached during extreme conditions. Working with Login Consultants, we devised a methodology to validate the testing without reaching VSIMax Dynamic until such time as a new calculation is available.

Our Login VSI "pass" criteria, accepted by Login Consultants for this testing is as follows:

- Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, Memory utilization, Storage utilization and Network utilization.
- We will use Login VSI to launch version 3.6 medium workloads, including flash.
- Number of Launched Sessions must equal Active Sessions within two minutes of the last session launched in a test.
- The VMware Horizon View Desktop Administrator will be monitored throughout the steady state to make sure that:
  - All running sessions report In Use throughout the steady state

  - No sessions move to Unregistered 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.
- We will publish our CVD with our recommendation following the process above and will note that we did not reach a VSIMax dynamic in our testing due to a technical issue with the analyzer formula that calculates VSIMax.

The purpose of this testing is to provide the data needed to validate VMware View 5.2 automated pool, floating assignment linked clone virtual desktops using ESXi 5.1 and vCenter 5.1 to virtualize Microsoft Windows 7 SP1 desktops on Cisco UCS B200 M3 blade servers using a EMC VNX5500 storage system.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of View 5.2 with VMware vSphere.

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

One additional series of stress tests on a single blade server was conducted to establish the official Login VSI Max Score. To reach the Login VSI Max, we ran 195 Medium Workload (with flash) Windows 7 SP1 sessions on a single server. The Login VSI score was achieved on three consecutive runs and is shown in the next section of the document.

Cisco UCS Test Configuration for Single-Server Scalability Test Results

This section details the results from the View 5.2 Hosted VDI single blade server validation testing. The primary success criteria used to validate the overall success of the test cycle is an output chart from Login Consultants' VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)

Note We did not reach a VSIMax Dynamic in our testing due to a technical issue with the analyzer formula that calculates VSIMax. See Section 8.3.4.5 Determining VSIMax for a discussion of this issue.

We ran the single server test at approximately 10% lower user density than prescribed by the Login VSI Max to achieve a successful pass of the test with server hardware performance in a realistic range.

Additionally, graphs detailing the CPU, Memory utilization and network throughput during peak session load are also presented. Given adequate storage capability, the CPU utilization determined the maximum VM density per blade.

The charts below present our recommended maximum Login VSI Medium workload loading on a single blade server.

Figure 21 Login VSI Max graph- 155 Sessions Average Response Time below 2000 m seconds

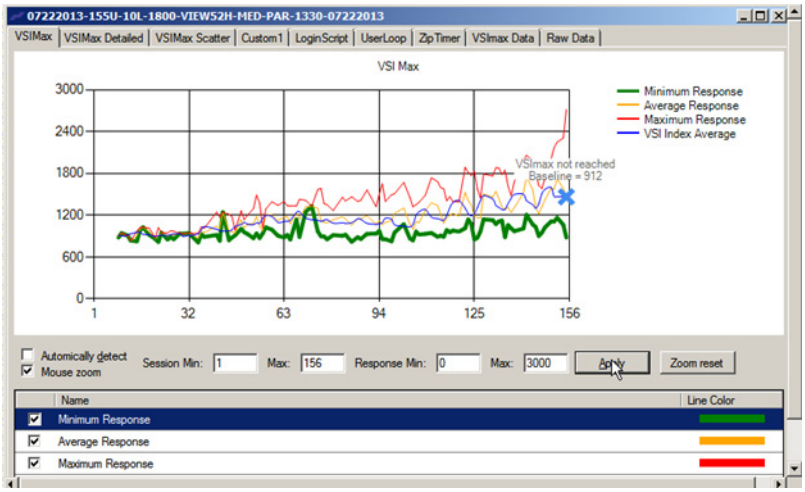

The following graphs detail CPU, Memory, Disk and Network performance on the Single Cisco UCS B200-M3 Blades

Figure 22 155 User Single B200 M3 CPU Utilization Boot Phase

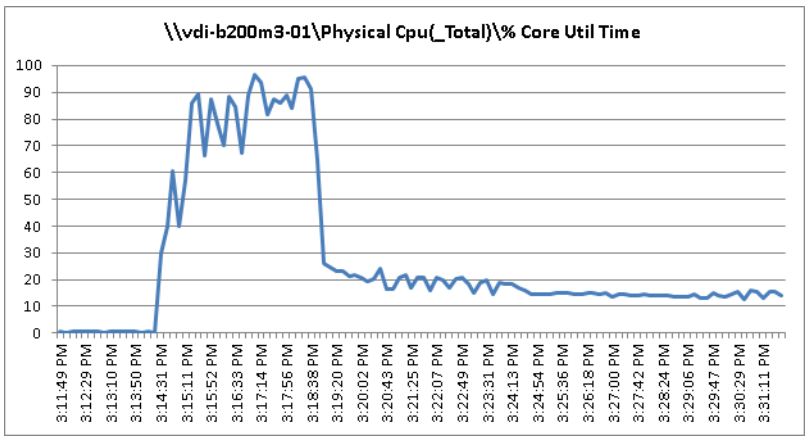

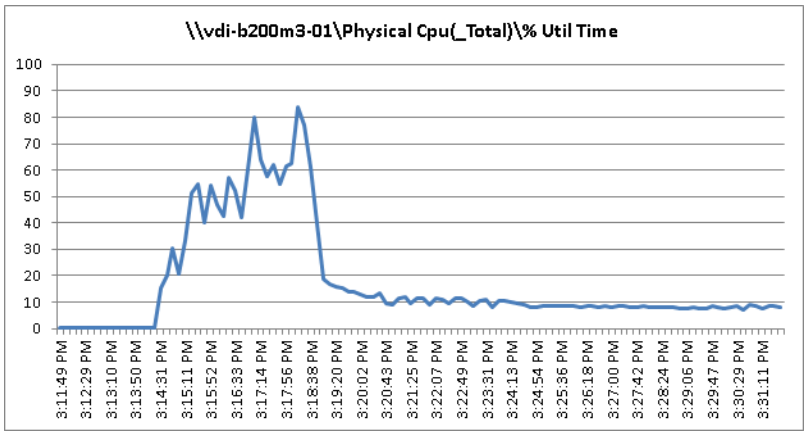

Figure 23 155 User Single B200 M3 Available Memory Boot Phase

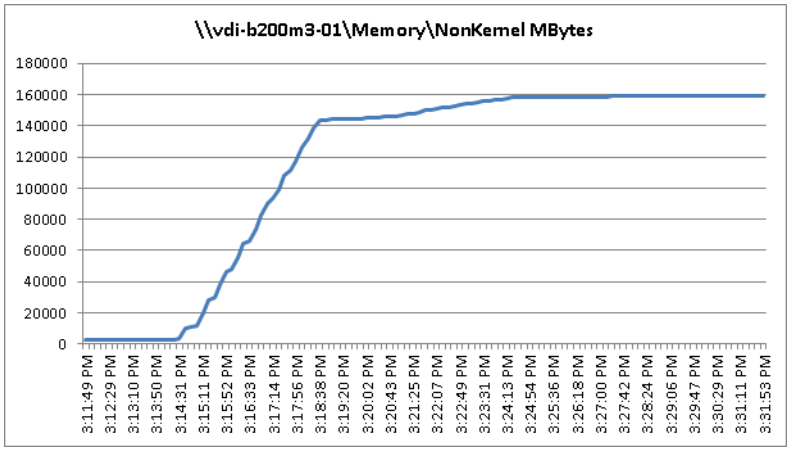

Figure 24 155 User Single B200 M3 Cisco VIC1240 MLOM Mbps Received/TransmittedBoot Phase

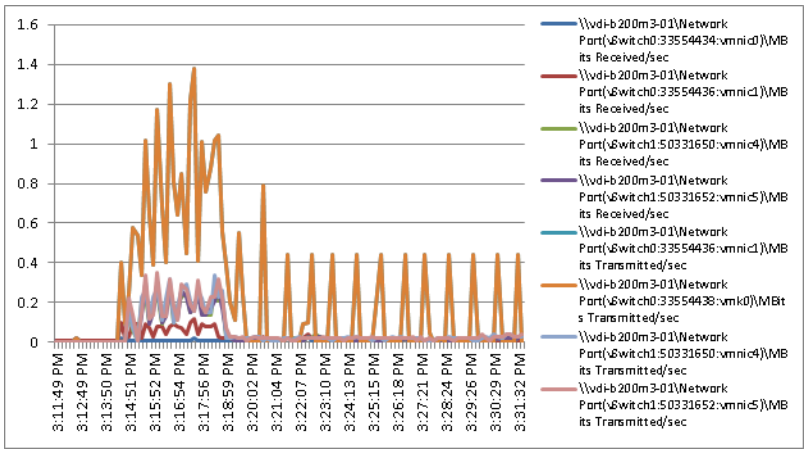

Figure 25 155 User Single B200 M3 Cisco VIC1240 MLOM Physical Disk Adapter MBps Read/Write Boot Phase

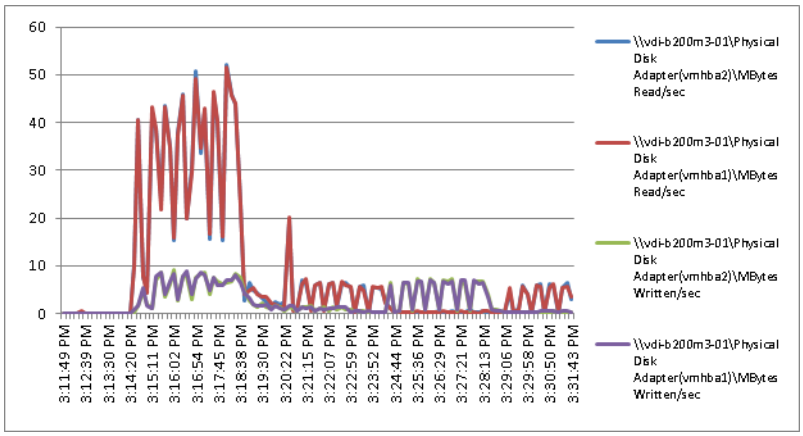

Figure 26 155 User Single B200 M3 CPU Core Utilization - Test Phase

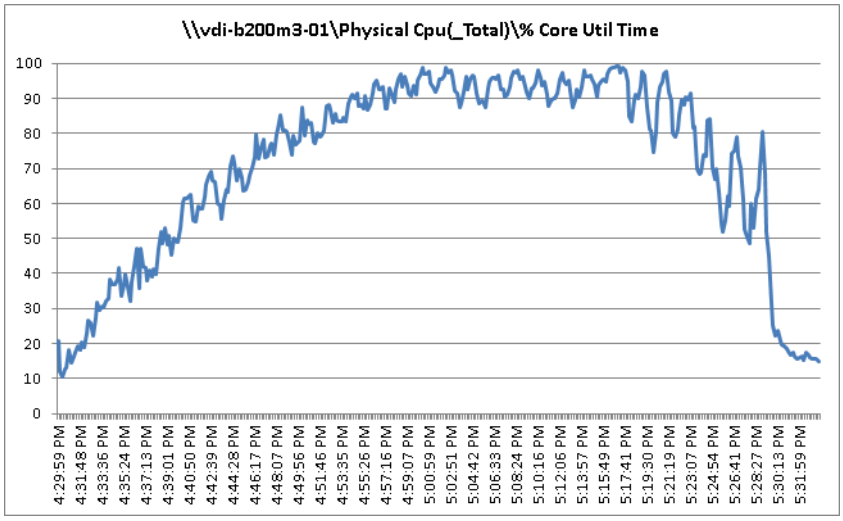

Figure 27 155 User Single B200 M3 CPU Processor Total Utilization Time - Test Phase

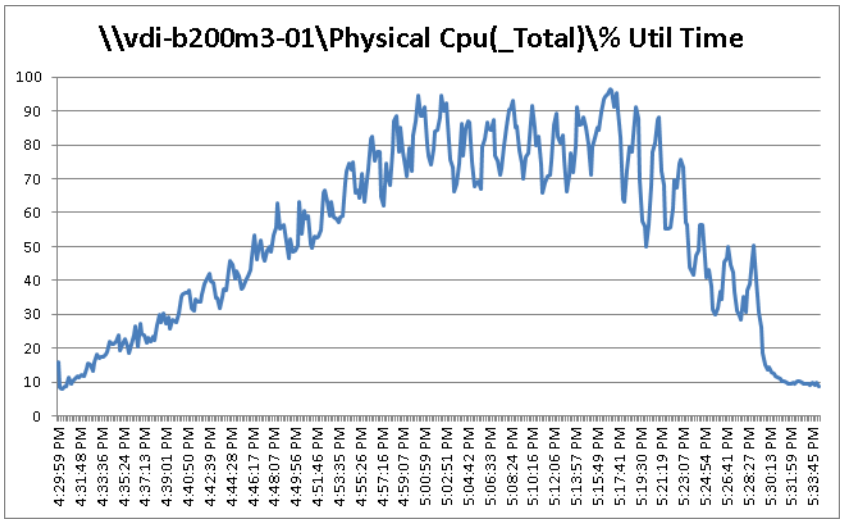

Figure 28 155 User Single B200 M3 Available Memory Test Phase

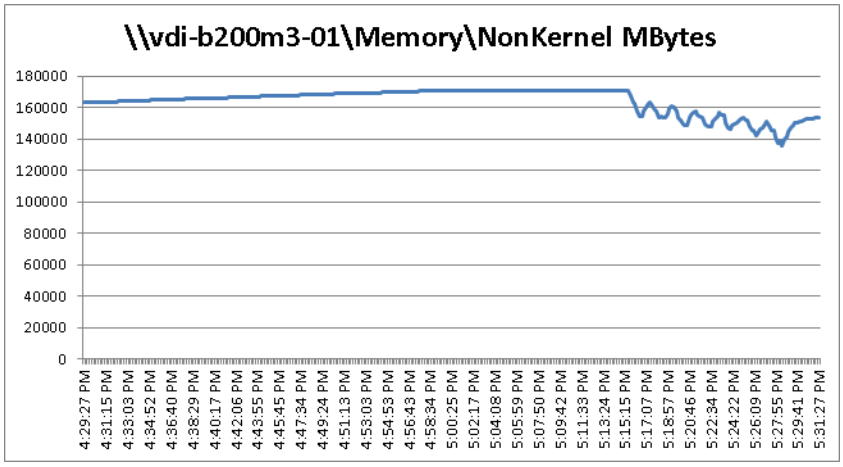

Figure 29 155 User Single B200 M3 Cisco VIC1240 MLOM Mbps Received/Transmitted Test Phase

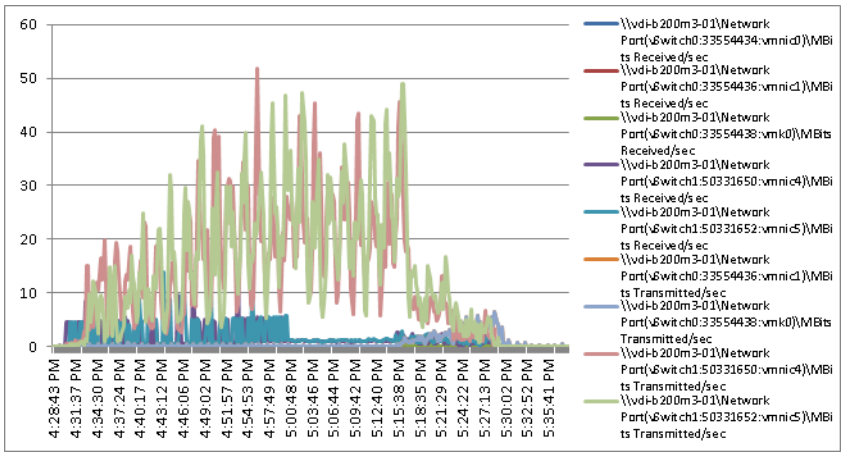

Figure 30 155 User Single B200 M3 Cisco VIC1240 MLOM Mbps Read/Written Test Phase

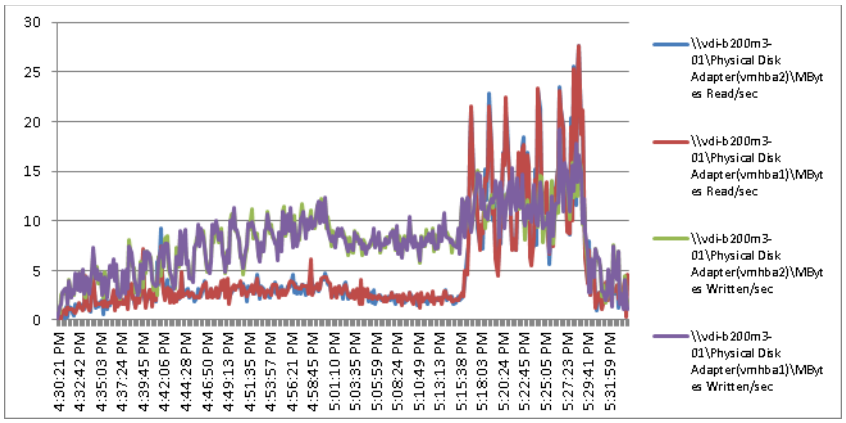

The following charts details for VIEW Connection Server performance during the single blade, 155 User test: Boot Phase

Figure 31 155 User View Connection Server 5.2 CPU Utilization

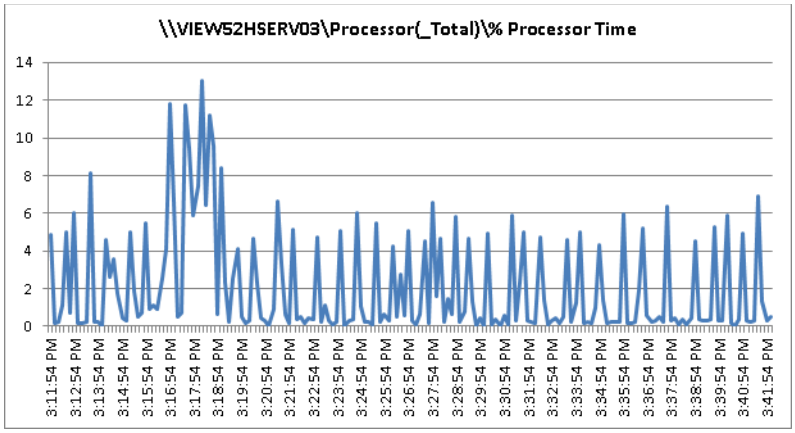

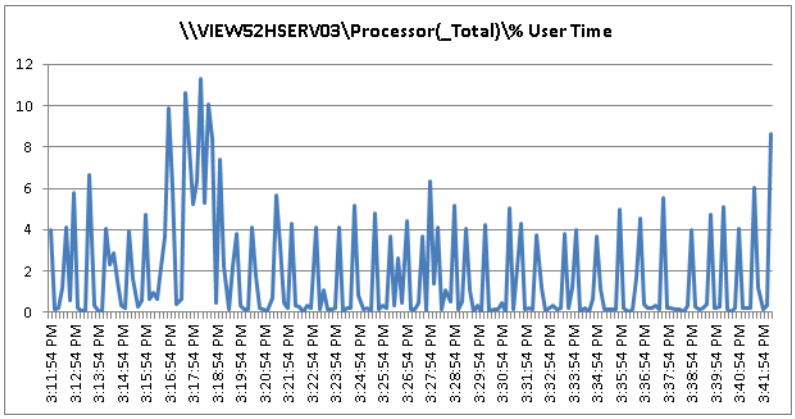

Figure 32 155 User View Connection Server 5.2 Available Memory

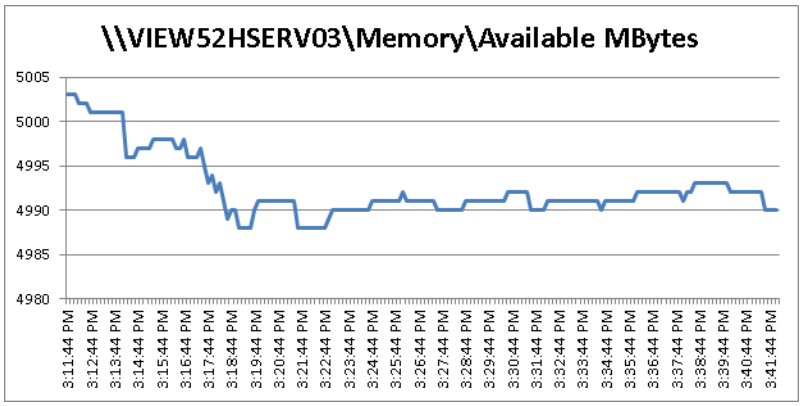

Figure 33 155 User View Connection Server 5.2 Bytes Received /Second

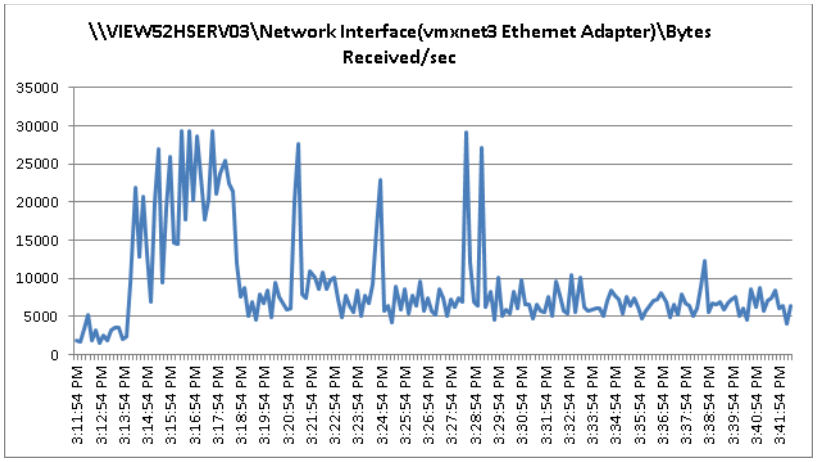

Figure 34 155 User View Connection Server 5.2 Bytes Sent/Second

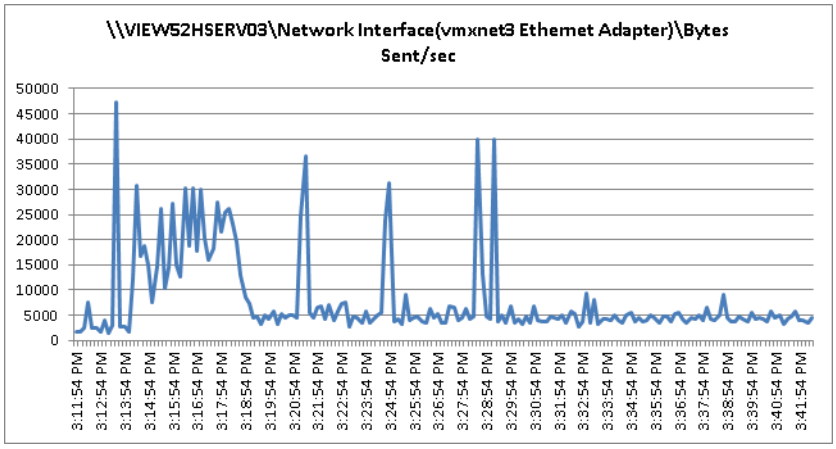

Figure 35 155 User View 5.2 - View Connection Server CPU Utilization-Test Phase

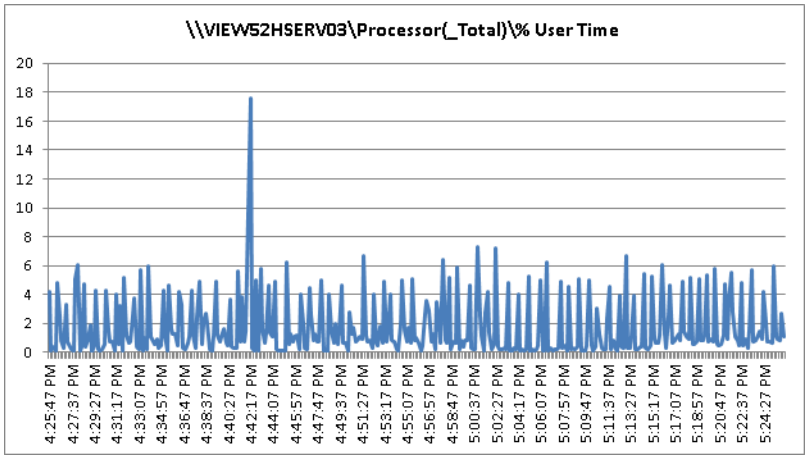

Figure 36 155 User View 5.2 - View Connection Total User Time- TEST Phase

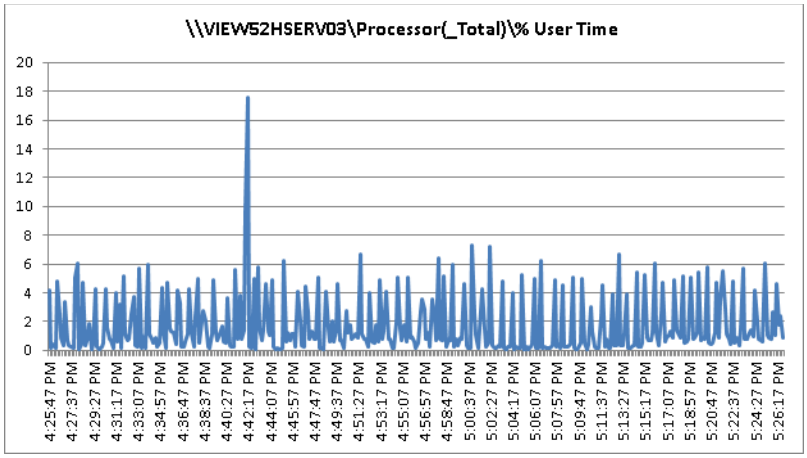

Figure 37 155 User View 5.2 - View Connection Server Available Memory

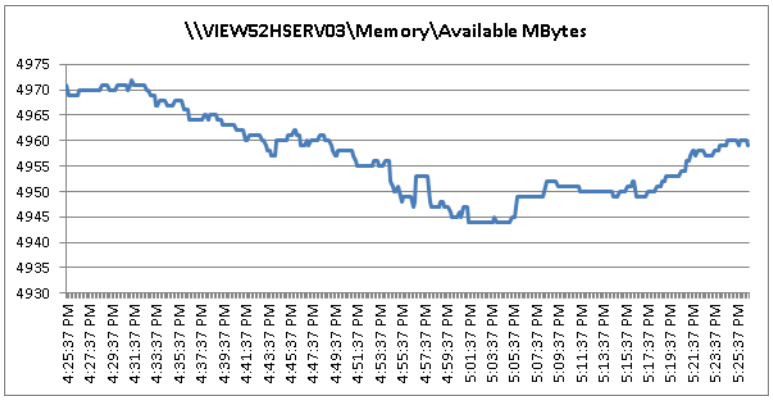

Figure 38 155 User View 5.2 View Connection Server Bytes Received

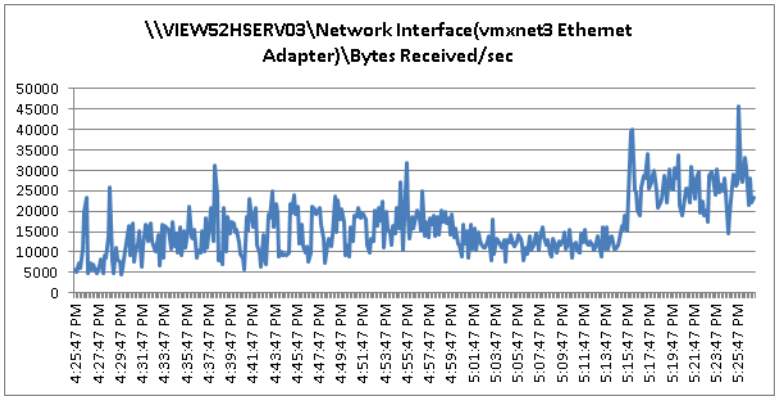

Figure 39 155 User View 5.2 View Connection Server Bytes Sent

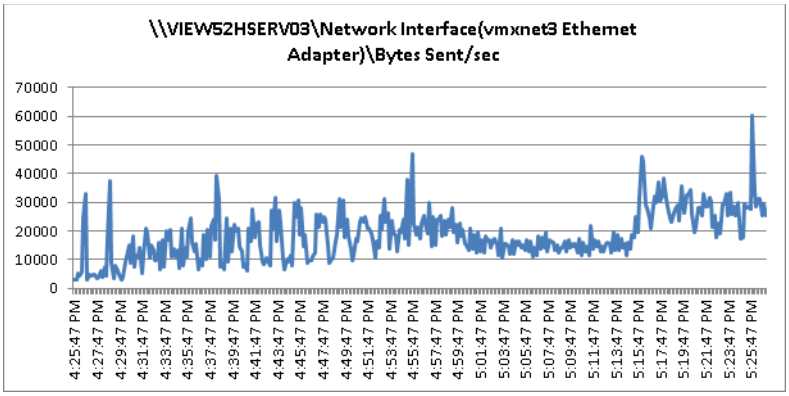

Cisco UCS Test Configuration for 2000 Desktop Two-Cluster Scalability Test Results

This section details the results from the View 5.2 Hosted VDI seven blade server, two-cluster, 2000 user validation testing. It demonstrates linear scalability for the system. The primary success criteria used to validate the overall success of the test cycle is an output chart from Login Consultants' VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)

Note We did not reach a VSIMax Dynamic in our testing due to a technical issue with the analyzer formula that calculates VSIMax. See Section 8.3.4.5 Determining VSIMax for a discussion of this issue.

We ran the multi-server test at an average user density slightly below 143 users per blade across the system. Two ESX Clusters, each containing seven B200 M3s ran the entire workload. In fact the fourteen blade test harness provides N+1 server fault tolerance on a system basis to achieve a successful pass of the test with server hardware performance in a realistic range.

Additionally, graphs detailing the CPU, Memory utilization and network throughput during peak session load are also presented for a representative blade running 143 user sessions below. We have provided performance charts for all 14 blades in Appendix B to illustrate this point.

Given adequate storage capability, the blade CPU utilization determined the maximum recommended VM density per blade for the 2000 user environment.

We also present performance information on key infrastructure virtual machines with the tested blade data.

For the large scale test, we are including the EMC VNX5500 performance metrics as well.

Figure 40 2000 Desktop Sessions on VMware ESXi 5.1 below 4000 ms

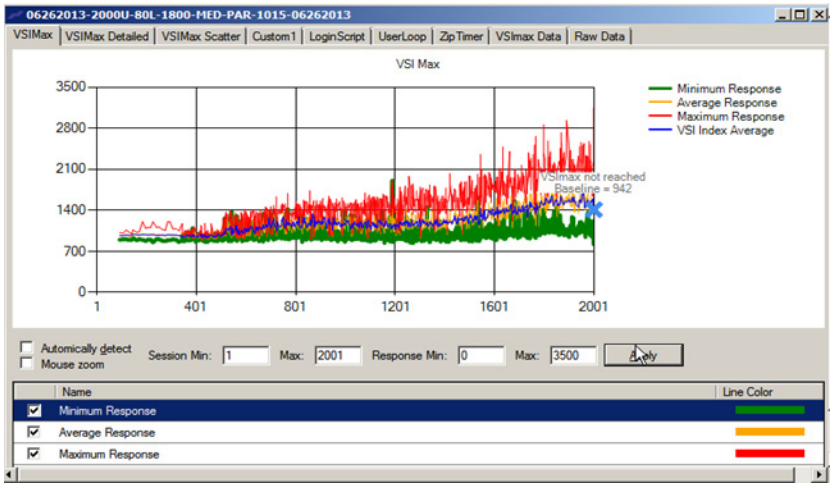

The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B200 M3 Blade during the fourteen blade, 2000 User test. (Representative results for all fourteen blades in one of the vCenter clusters can be found in Appendix B.)

Figure 41 2000 User Single B200 M3 Core CPU Utilization Boot Phase

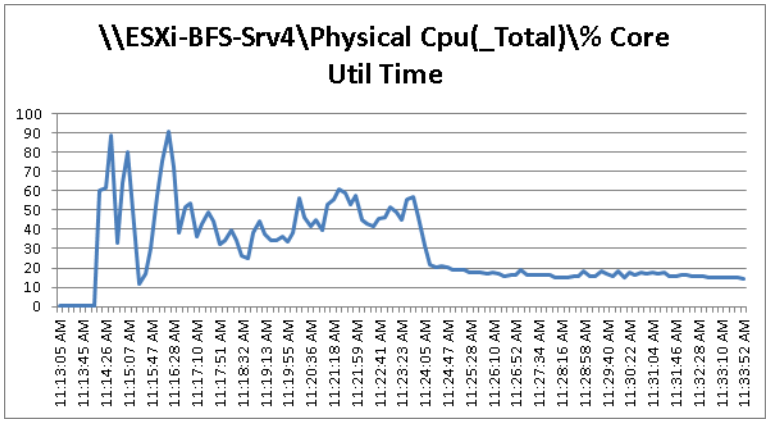

Figure 42 2000 User Single B200 M3 CPU Utilization Boot Phase

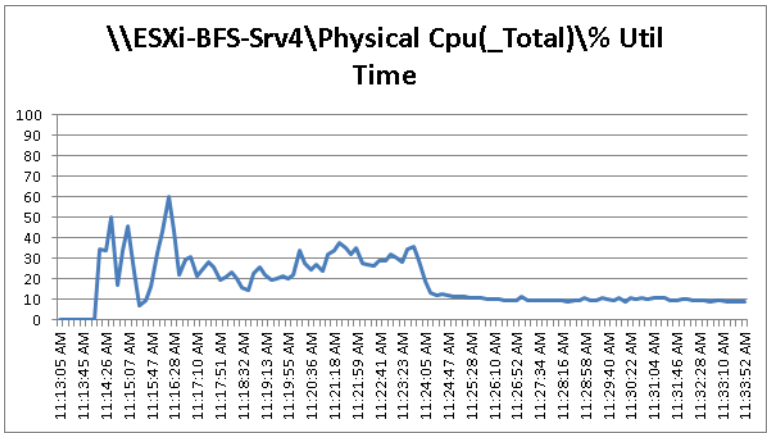

Figure 43 2000 User Single Cisco UCS B200 M3 Cisco NonKernel Mbytes Available Boot Phase

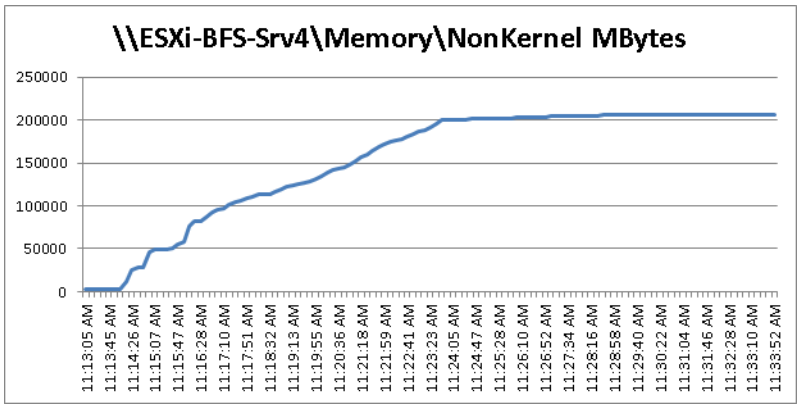

Figure 44 2000 User Single Cisco B200 M3 Cisco VIC1240 MLOM Physical Disk Adapter

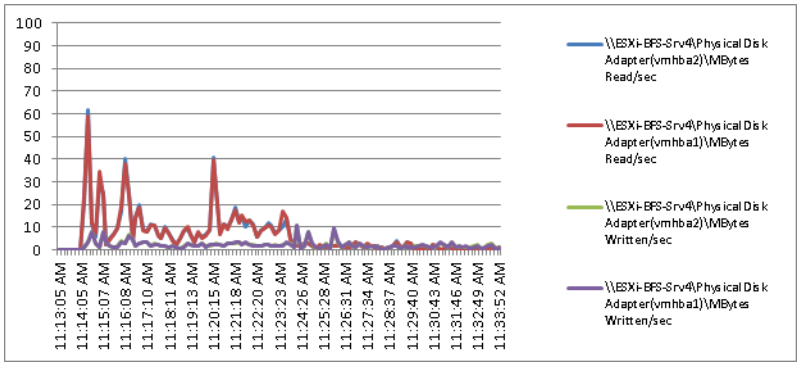

Figure 45 2000 User Single Cisco UCS B200 M3 Cisco VIC1240 MLOM VIC Mbps Receive/Transmit Boot Phase

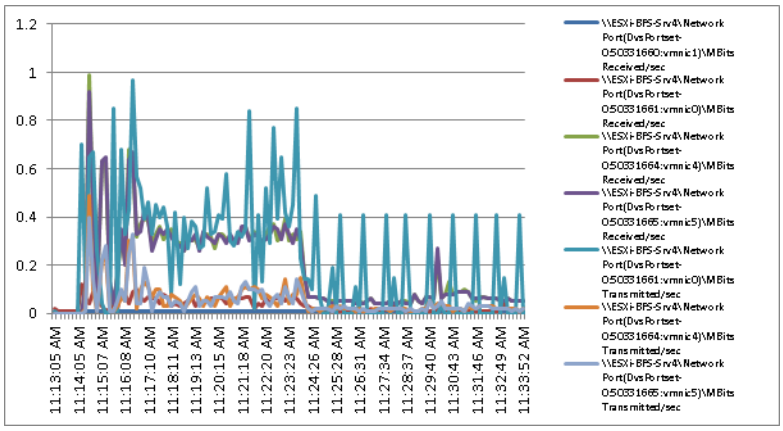

Figure 46 2000 User Single Cisco UCS B200 M3 CPU Utilization Test Phase

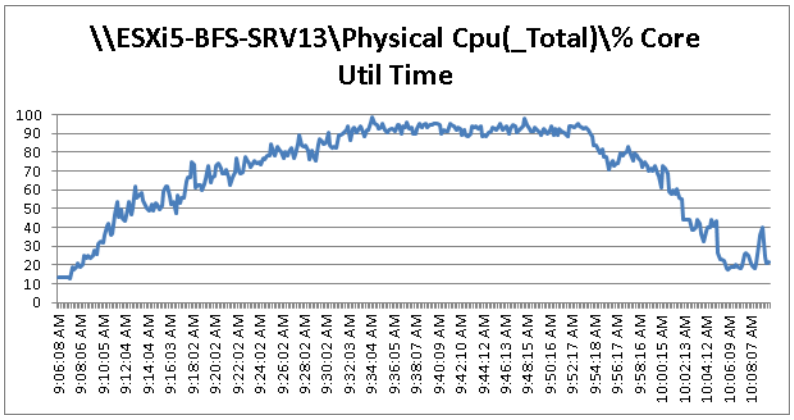

Figure 47 2000 User Single Cisco UCS B200 M3 Cisco VIC1240 MLOM Network Adapter Mbps Receive/Transmit Test Phase

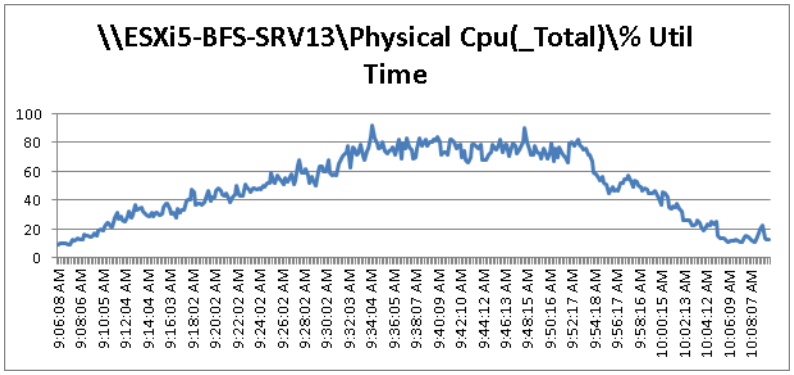

Figure 48 2000 User Single Cisco UCS B200 M3 Cisco VIC1240 MLOM Physical Disk Adapter MBps Read/Write Test Phase

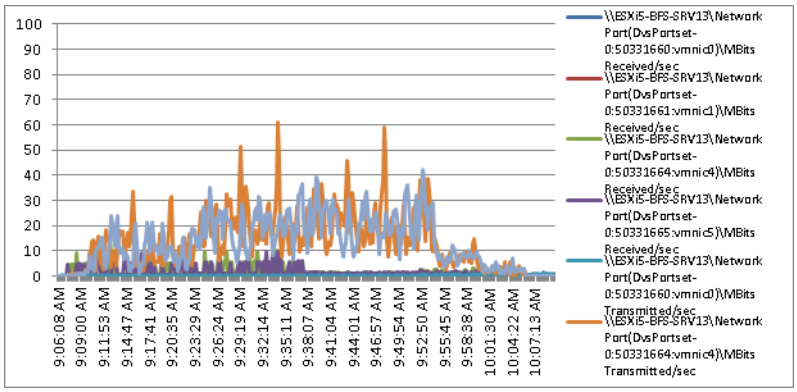

Figure 49 2000 User Single Cisco UCS B200 M3 Cisco VIC1240 MLOM VIC Mbps Receive/Transmit Test Phase

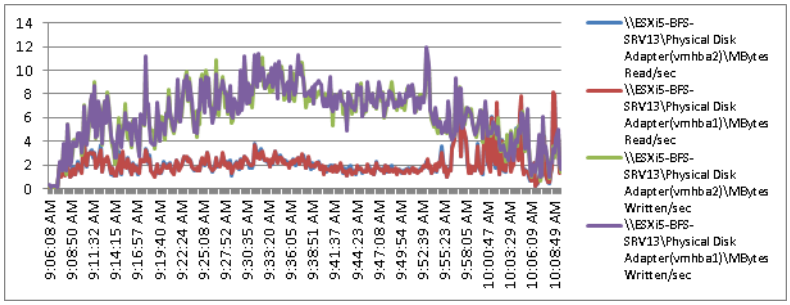

Figure 50 2000 User Single Cisco UCS B200 M3 Cisco Non-Kernel Mbytes Avaaible

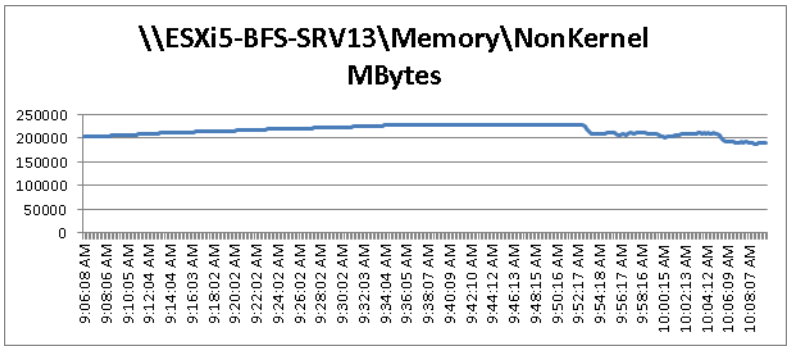

The following charts detail the VNX5500 performance during the fourteen blade, 2000 User test:

Figure 51 2000 Users EMC VNX5500 SP Utilization Boot Phase

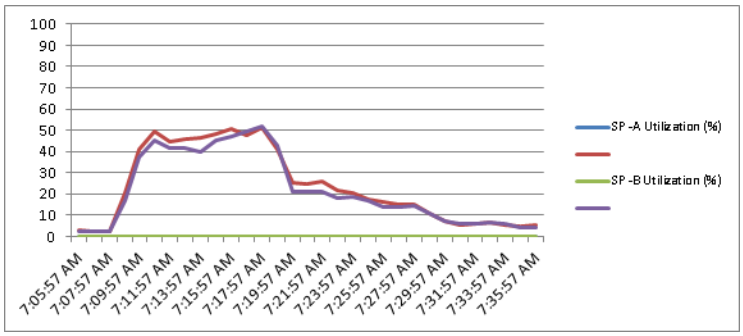

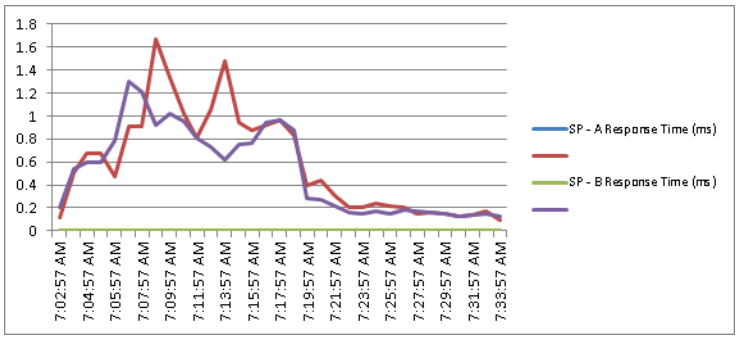

Figure 52 2000 Users EMC VNX5500 SP Queue Lengths Boot Phase

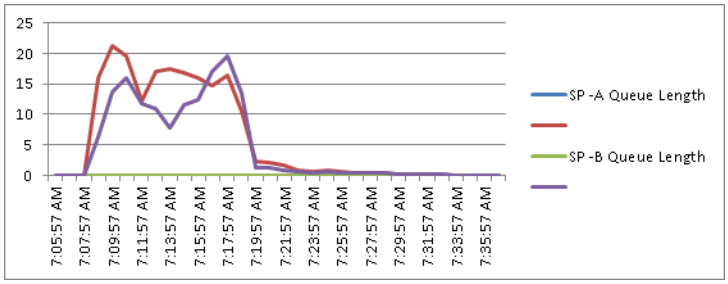

Figure 53 2000 Users EMC VNX5500 SP Total Throughput Boot Phase

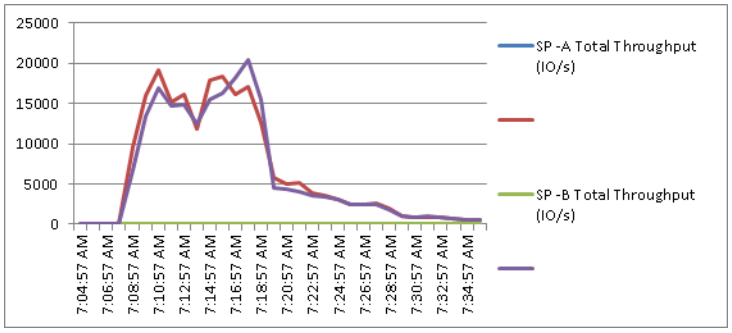

Figure 54 2000 Users EMC VNX5500 SP Utilization Test Phase

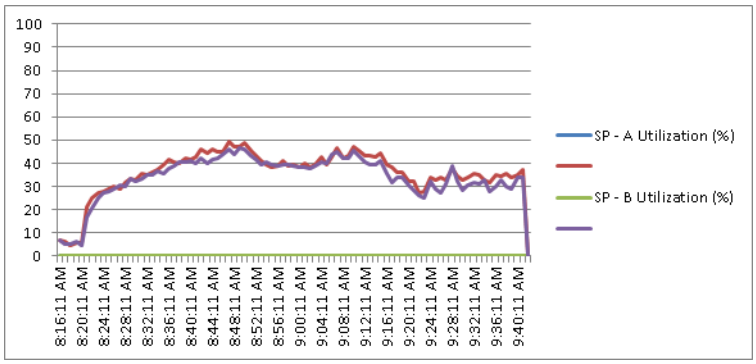

Figure 55 2000 Users EMC VNX5500 SP Queue Lengths Test Phase

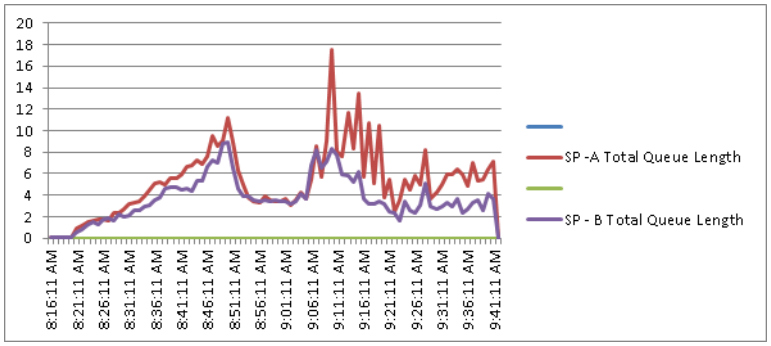

Figure 56 2000 Users EMC VNX5500 SP Total Throughput Test Phase

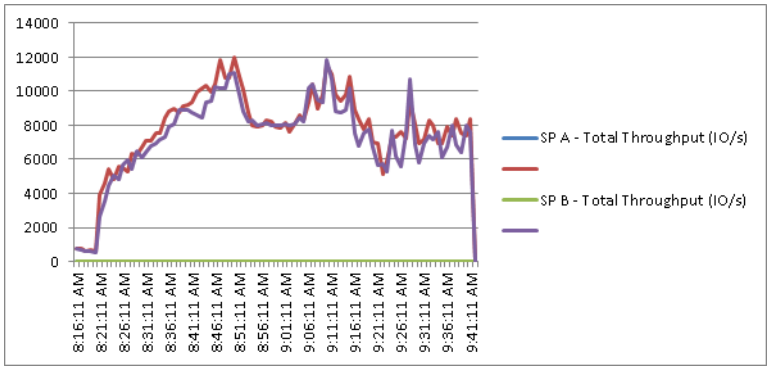

Figure 57 2000 Users EMC VNX5500 SP Total Response Time Test Phase

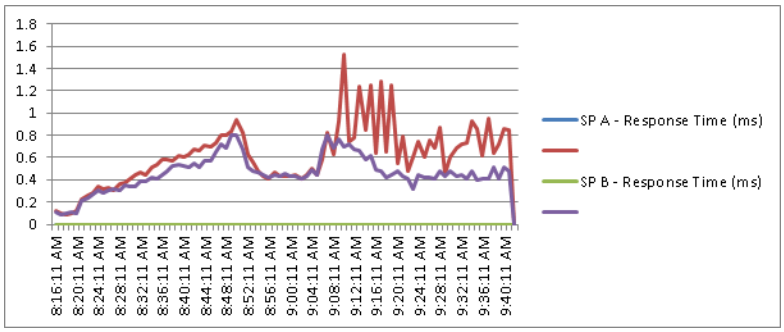

The following charts detail infrastructure server performance during the fourteen blades, 2000 User test:

Figure 58 2000 User View Connection Server 5.2 CPU Utilization Test Phase

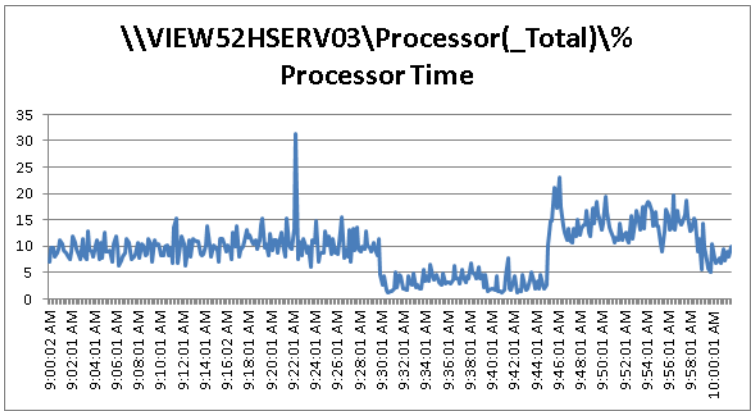

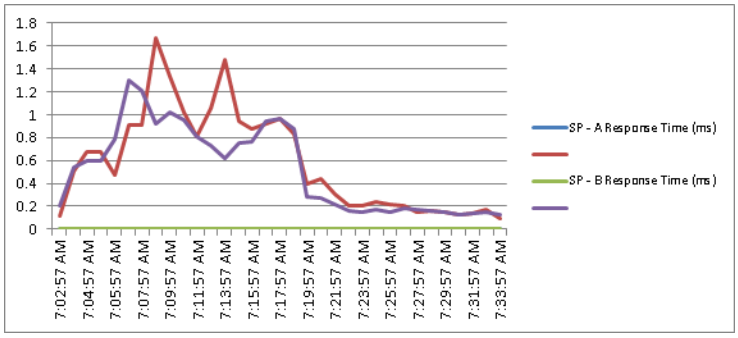

Figure 59 2000 Users EMC VNX5500 SP Queue Lengths Boot Phase

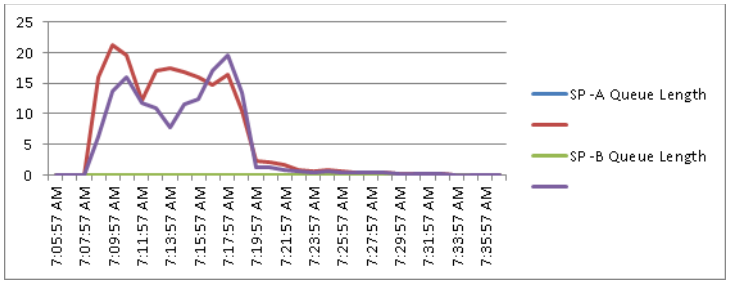

Figure 60 2000 Users EMC VNX5500 SP Total Throughput Boot Phase

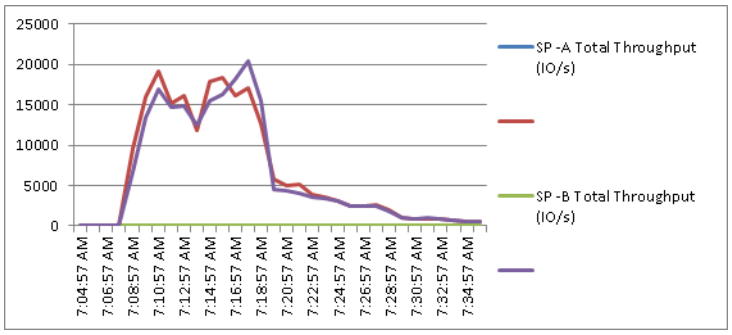

Figure 61 2000 Users EMC VNX5500 SP Utilization Test Phase

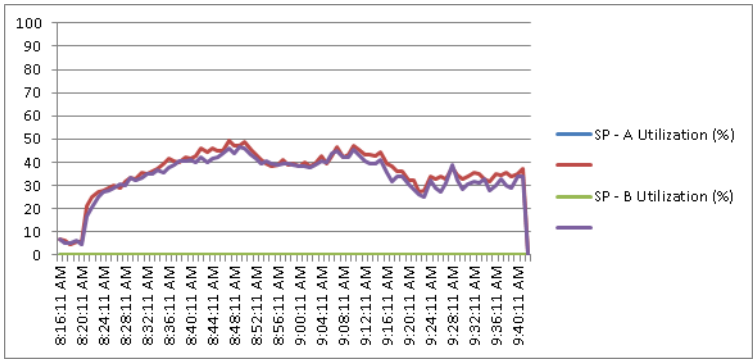

Figure 62 2000 Users EMC VNX5500 SP Queue Lengths Test Phase

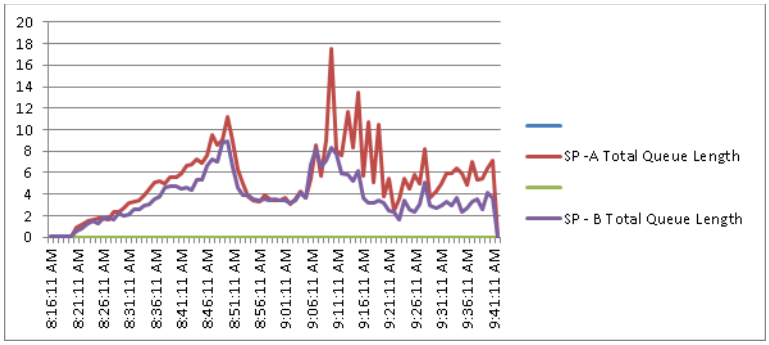

Figure 63 2000 Users EMC VNX5500 SP Total Throughput Test Phase

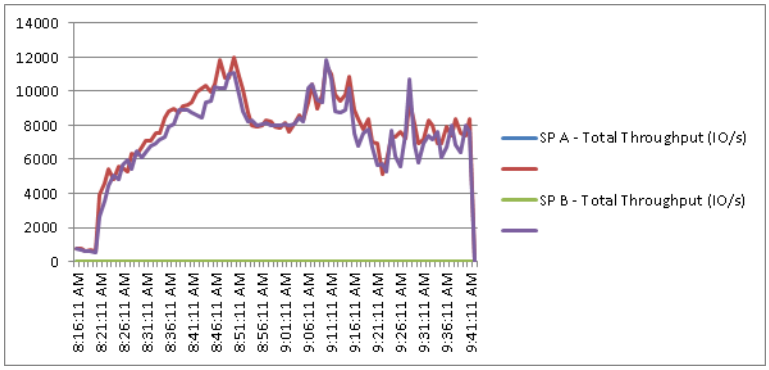

Figure 64 2000 Users EMC VNX5500 SP Total Response Time Test Phase

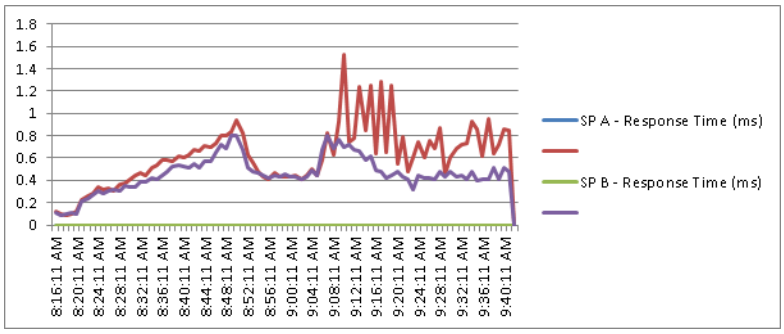

The following charts detail infrastructure server performance during the fourteen blades, 2000 User test:

Figure 65 2000 User View Connection Server 5.2 CPU Utilization Test Phase

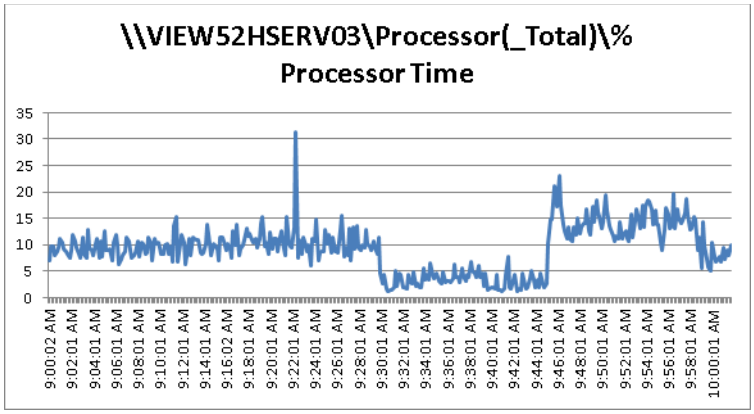

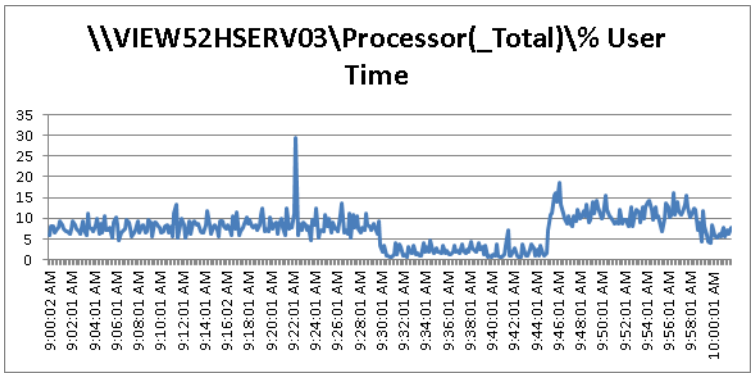

Figure 66 2000 User View Connection Server 5.2 Available Memory Test Phase

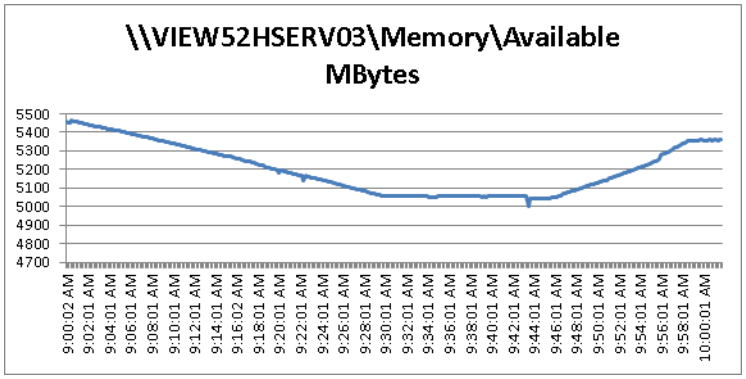

Figure 67 2000 User View Connection Server 5.2 Bytes Received/Second Test Phase

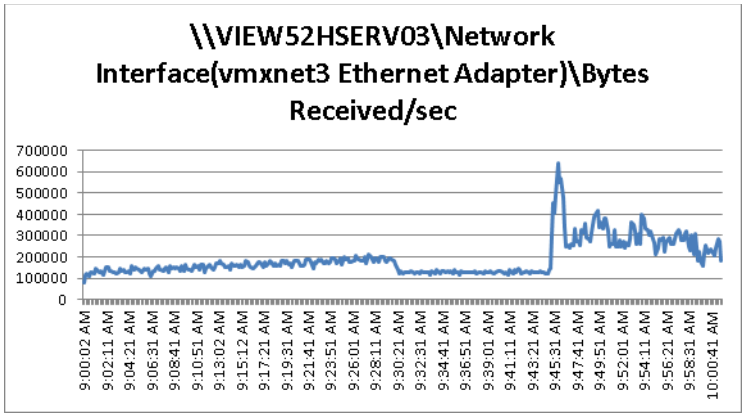

Figure 68 2000 User View Connection Server 5.2 Bytes Sent/Second Test Phase

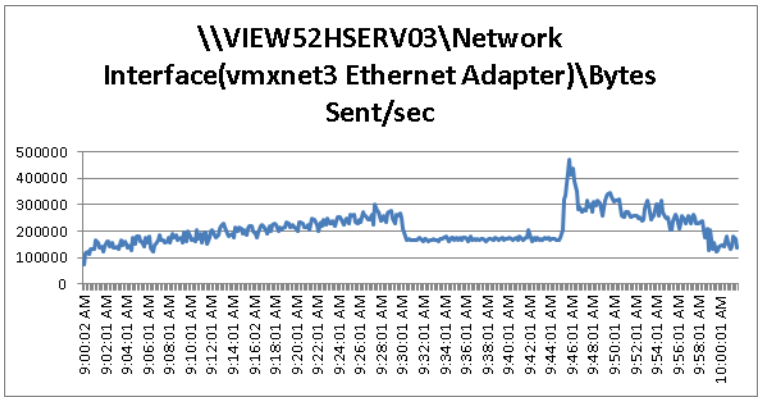

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 2000 User, six chassis, 14 VDI host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 2000 user system.

Cisco UCS System Configuration Considerations

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested.

- Cisco UCS 2.1 management software supports up to 20 chassis within a single Cisco UCS domain on our second generation Cisco UCS Fabric Interconnect 6248 and 6296 models. Our single UCS domain can grow to 160 half-width blades.
- With Cisco UCS 2.1 management software, released late in November 2012, each Cisco UCS 2.1 Management domain is ostensibly manageable by Cisco UCS Central, our new manager of managers, vastly increasing the reach of the Cisco UCS system.

- As scale grows, the value of the combined Cisco UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100% of the time.
- To accommodate the Cisco Nexus 5500 upstream connectivity in the way we describe in the LAN and SAN Configuration section, we need four Ethernet uplinks and two Fibre Channel uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we can calculate number of desktops can be hosted in a single Cisco UCS domain. Assuming eight links per chassis, four to each 6248, scaling beyond 10 chassis would require a pair of Cisco UCS 6296 fabric interconnects. A 20,000 virtual desktop building block, with its support infrastructure services can be built out of the RA described in this study with eight links per chassis and 20 Cisco UCS chassis comprised of seven Cisco UCS B200 M3 VDI blade server and one Cisco UCS B200 M3 Infrastructure blades servers in each chassis.

Of course, the backend storage has to be scaled accordingly, based on the IOPS considerations as described in the EMC scaling section. Please refer the EMC section that follows this one for scalability guidelines.

VMware View 5.2 Considerations

VMware View Composer can create and provision up to 1000 desktops per pool when deployed on vSphere 4.1 or later. View Composer can also perform a recompose operation on up to 1,000 desktops at a time. Desktop pool size is limited by the following factors:

- Each desktop pool can contain only one ESX/ESXi cluster.
- With View 5.2 and later and vSphere 5.0 and later, an ESXi cluster can contain more than 8 ESXi hosts (up to 32), but you must store the linked-clone replica disks on NFS datastores.
- Each CPU core has compute capacity for 8 to 10 virtual desktops.

A single VMware View Connection server can host up to 2000 simultaneous connections over any supported connection type. Seven View Connection Servers (5 active plus 2 spares) can host up to 10000 direct, RDP or PCoIP connections simultaneously. The sever View Connection Server cluster configuration should not be clustered across WAN links.

VMware View deployments can use VMware HA clusters to guard against physical server failures. With View 5.2 and later and vSphere 5 and later, if you use View Composer and store replica disks on NFS datastores, the cluster can contain up to 32 servers, or nodes.

With vCenter 4.1 and 5.0, each vCenter Server can support up to 10,000 virtual machines.

For more information on VMware View 5.2 configuration and guidelines, see Chapter 11 References.

EMC VNX Storage Guidelines for Horizon View 5.2 Provisioned Virtual Machines

Sizing VNX storage system to meet virtual desktop IOPS requirement is a complicated process. When an I/O reaches the VNX storage, it is served by several components such as Data Mover (NFS), backend dynamic random access memory (DRAM) cache, FAST Cache, and disks. To reduce the complexity, EMC recommends using a building block approach to scale to thousands of virtual desktops.

For more information on storage sizing guidelines to implement virtual desktop infrastructure in VNX unified storage systems, refer to the EMC white paper "Sizing EMC VNX Series for VDI workload - An Architectural Guideline".

VMware ESXi 5.1 Guidelines for Virtual Desktop Infrastructure

In our test environment two adjustments were performed to support our scale:

- The amount of memory configured for the Tomcat Maximum memory pool was increased to 3072.
- The cost threshold for parallelism was increased to 15.

For further explanations on a basis for these adjustments and details on how to perform them, refer to the VMware documentation sited in the References section of this document.

References

This section provides links to additional information for each partner's solution component of this document.

Cisco Reference Documents

Third-Generation Fabric Computing: The Power of Unification webcast replay

http://tools.cisco.com/gems/cust/customerSite.do?METHOD=W&LANGUAGE_ID=E&PRIORITY_CODE=215011_15&SEMINAR_CODE=S15897&CAMPAIGN=UCS+Momentum&COUNTRY_SITE=us&POSITION=banner&REFERRING_SITE=go+unified+computing&CREATIVE=carousel+banner+event+replay

Cisco Unified Computing System Manager Home Page

http://www.cisco.com/en/US/products/ps10281/index.html

Cisco UCS B200 M3 Blade Server Resources

http://www.cisco.com/en/US/partner/products/ps12288/index.html

Cisco UCS 6200 Series Fabric Interconnects

http://www.cisco.com/en/US/partner/products/ps11544/index.html

Cisco Nexus 1000V Series Switches Resources

http://www.cisco.com/en/US/partner/products/ps9902/index.html

Cisco Nexus 5500 Series Switches Resources

http://www.cisco.com/en/US/products/ps9670/index.html

Download Driver Software for Cisco UCS B200 M3 Blade Server

http://software.cisco.com/download/release.html?mdfid=283612660&flowid=22121&softwareid=283655658&release=2.1%281a%29&relind=AVAILABLE&rellifecycle=&reltype=latest

Download Cisco UCS Manager and Blade Software Version 2.1(1b)

http://software.cisco.com/download/release.html?mdfid=283612660&flowid=22121&softwareid=283655658&release=2.1%281a%29&relind=AVAILABLE&rellifecycle=&reltype=latest

Download Cisco UCS Central Software Version 1.0(1a)

http://software.cisco.com/download/cart.html?imageGuId=8CAAAD77B3A1DB35B157BE84ED109A4703849F53&i=rs

VMware View Reference Documents

View 5 Documents

Performance and Best Practices

http://www.vmware.com/files/pdf/view/vmware-horizon-view-best-practices-performance-study.pdf

View 5.2 Architecture and Planning
https://pubs.vmware.com/view-52/topic/com.vmware.ICbase/PDF/horizon-view-52-architecture-planning.pdf

View 5 with PCoIP Network Optimization Guide

http://www.vmware.com/files/pdf/view/VMware-View-5-PCoIP-Network-Optimization-Guide.pdf

Virtual Desktop

Windows 7 Optimization Guide
http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf

EMC Reference Documents

- Sizing EMC VNX Series for VDI Workload - An Architectural Guideline

VMware Reference Documents

- Accessing a vCenter Server using Web access or vSphere Client fails with an SSL certificate error: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514
- VMware vSphere ESXi and vCenter Server 5 Documentation: http://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.upgrade.doc%2FGUID-200B9E03-D46B-44A9-9B0E-4863D067CFFF.html
- VMware vCenter Management Webservices features do not function properly: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1039180
- VMware® vCenter Server™ 5.1 Database Performance Improvements and Best Practices for Large-Scale Environments: http://www.vmware.com/files/pdf/techpaper/VMware-vCenter-DBPerfBestPractices.pdf
- Performance Best Practices for VMware vSphere™ 5.0: http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf

Appendix

Click the following link to open the Appendix:
http://www.cisco.com/en/US/docs/unified_computing/ucs/UCS_CVDs/ucs_vspex_vview5.2_2000_appendix.pdf.