Cisco Desktop Virtualization Solutions for EMC VSPEX with VMware View 5.1.2 for 500 Desktops

Built on Cisco Unified Computing System, Nexus 2232PP, EMC VNX Storage, and VMware ESXi 5.1

Last Updated: May 7, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Hardik Patel, Support Engineer, Cisco Systems, Inc

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

Ramesh Guduru, Support Engineer, Cisco Systems, Inc

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

Mike Brennan, Sr. Technical Marketing Engineer, VDI Performance and Solutions Team Lead, Cisco Systems

Mike Brennan is a Cisco Unified Computing System architect, focusing on Virtual Desktop Infrastructure solutions with extensive experience with EMC VNX, VMware ESX/ESXi, XenDesktop and Provisioning Services. He has expert product knowledge in application and desktop virtualization across all three major hypervisor platforms, both major desktop brokers, Microsoft Windows Active Directory, User Profile Management, DNS, DHCP and Cisco networking technologies.

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.

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

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

Cisco Desktop Virtualization Solutions for EMC VSPEX with VMware View 5.1.2 for 500 Desktops

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 document reports the results of a study evaluating the scalability of a VMware View 5.1.2 environment, utilizing View Linked Clones, on managed Cisco UCS C-Series C220 M3 Rack-Mount Servers running VMware ESXi 5.1 hypervisor software connected to an EMC VNX5300 Storage Array. We utilize second generation Unified Computing System hardware and software. We provide best practice recommendations and sizing guidelines for a 500-600 virtual desktop customer deployment of View 5.1.2 on the Cisco Unified Computing System.

Five Cisco UCS C220 M3 Rack Servers were utilized in the design to provide N+1 fault tolerance for 500 Virtual Windows 7 desktops at the server level, guaranteeing the same end-user experience if just 4 C220 M3 servers are operational. In fact, the five server architecture can comfortably support 600 desktops with N+1 server fault tolerance. For that reason, the document architecture will refer to supporting a 600 desktop capacity with five Cisco UCS C220 M3 servers.

Alternatively, with just four Cisco UCS C220 M3 Rack Servers, we can effectively host 500 users with all servers online or 450 Users with 3 Cisco UCS C220 M3 servers running.

This study was performed in conjunction with EMC's VSPEX program and is aligned with the VSPEX View 500 Reference Architecture and Design Guide.

Audience

This document describes the architecture and deployment procedures of an infrastructure comprised of Cisco, EMC, and VMware virtualization. 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.

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

Benefits of the Solution Components

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

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

- Cisco UCS B-Series blade servers for infrastructure and virtual workload hosting
- 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 achieving a flexible, agile, high-performance, self-integrating information technology with faster ROI. With Cisco UCS Manager 2.1 introduced on November 21, 2012, Cisco UCS C-Series Servers and Fibre Channel (FC) SAN storage can be managed end-to-end by Cisco UCS 6200 Series Fabric Interconnects
- Fabric Extender technology, particularly the Nexus 2232PP Fabric Extender used in the FC variant, reduces the number of system components to purchase, configure and maintain
- 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
- With Cisco UCS Manger 2.1 and Cisco UCS Central 1.0, the scope of management extends to many Cisco UCS Domains
- Leverage Cisco UCS Management Packs for System Center 2012 for integrated management

Benefits of Nexus Switching

Cisco Nexus 5548 (NFS Variant)

The Cisco Nexus 5548UP Switch, used exclusively in the NFS variant or the EMC VSPEX C500 proven architecture, 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:
- 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

Note Cisco Nexus 5548UPs were utilized in the NFS variant of the study only.

Cisco Nexus 2232PP Fabric Extender (Fibre Channel Variant)

The Cisco Nexus 2232PP 10GE Fabric Extender provides 32 10 Gb Ethernet and Fibre Channel Over Ethernet (FCoE) Small Form-Factor Pluggable Plus (SFP+) server ports and eight 10 Gb Ethernet and FCoE SFP+ uplink ports in a compact 1 rack unit (1RU) form factor.

The Nexus 2232PP in conjunction with VIC1225 converged network adapters in the Cisco UCS C220 M3 rack servers provide fault-tolerant single wire management of the rack servers through up to 8 uplink ports to Cisco Fabric Interconnects.

Reduce TCO

- The innovative Fabric Extender approach reduces data center cabling costs and footprint with optimized inter-rack cabling
- Unified fabric and FCoE at the server access layer reduce capital expenditure and operating expenses Simplify Operation
- Cisco UCS 6248UP or 6296UP Fabric Interconnects provide a single point of management and policy enforcement
- Plug-and-play management includes auto-configuration

Note Cisco Nexus 2232PPs were utilized in the FC variant of the study only.

Benefits of EMC VNX Family of Storage Controllers

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, Fibre 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

As 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 only VMware delivers on the core requirements for a business virtualization infrastructure solution.

- Is 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
- Is proven over 350,000 customers

And best of all, VMware delivers while providing

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

Benefits of VMware View 5.1.2

Simplify desktop and application management while increasing security and control with VMware View. Deliver a personalized high fidelity experience for end-users across sessions and devices. Enable higher availability and agility of desktop services unmatched by traditional PCs while reducing the total cost of desktop ownership up to 50%. 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.

Automated Desktop Provisioning

VMware 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 log off 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

VMware 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 View customers.

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, called Cisco Solutions for EMC VSPEX End User Computing, produced a highly efficient, robust and scalable Virtual Desktop Infrastructure (VDI) for a hosted virtual desktop deployment. Key findings 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 View 5.1.2 software produces a high density per rack-server Virtual Desktop delivery system.
- Cisco UCS C-220 M3 Rack-Mount Servers support 500-600 virtual desktops in N+1 server fault tolerance configuration based on the number of rack servers deployed.
- The 500 seat design providing N+1 server fault tolerance for 450 users is based on four Cisco UCS C220 M3 1U rack-mount servers, each with dual 8-core processors, 256GB of 1600 MHz memory and a Cisco VIC1225 converged network adapter
- The 600 seat design providing N+1 server fault tolerance for 600 users is based on five Cisco UCS C220 M3 1U rack-mount servers, each with dual 8-core processors, 256GB of 1600 MHz memory and a Cisco VIC1225 converged network adapter.
- We were able to boot the full complement of desktops in under 25 minutes without pegging the processor, exhausting memory or storage subsystems.
- We were able to ramp (log in and exercise workloads) up to steady state in thirty minutes without pegging the processor, exhausting memory or storage subsystems.
- We maintain our industry leadership with our new Cisco UCS Manager 2.1(1a) 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 and Nexus 2200 Series fabric extenders as we run more challenging workload testing, maintaining unsurpassed user response times.
- For the Managed Cisco UCS C-Series FC variant, utilizing UCS Manager 2.1 Service Profile Templates and Service Profiles in conjunction with Nexus 2232PP Fabric Extenders, we were able fully configure all five Cisco UCS C220 M3 servers from cold start to ready to deploy VMware ESXi 5 boot from SAN in 30 minutes.
- For the Managed Cisco UCS C-Series FC variant of the study, utilizing UCS Manager 2.1, we were able to connect and integrate the EMC VNX5300 via FC on our UCS 6248UP Fabric Interconnects, including FC zoning, eliminating the requirement for upstream access layer switching or fiber channel switches for that purpose.
- For the Unmanaged Cisco UCS C-Series NFS variant of the study, we use a pair of Nexus 5548UP access layer switches to directly attach the unmanaged Cisco UCS C220 M3 servers and the EMC VNX5300.

- 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 Active Directory, Profile Servers, SQL Servers, View Connection Server, View Replica Server and View Composer Server were hosted as virtual servers.
- EMC's VNX5300 system provides storage consolidation and outstanding efficiency. Both block and file storage resources were provided by a single system, utilizing EMC Fast Cache technology.
- Whether using the Managed Cisco UCS C-Series FC or the Unmanaged C-Series NFS variant and the EMC VNX storage layout prescribed in this document, the same outstanding end user experience is achieved as measured by Login VSI 3.7 testing.
- VMware View 5.1.2 software utilizing floating assignment linked-clone technology provides excellent performance with host-rendered flash video and other demanding applications.

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 500 user View 5.1.2 solution with N+1 fault tolerance for 450 users includes Cisco networking, four Cisco UCS C220 M3 Rack-Mount Servers and an EMC VNX5300 storage system.

The 600 User View 5.1.2 solution with N+1 fault tolerance for 600 users includes Cisco networking, five Cisco UCS C220 M3 Rack-Mount Servers and the same EMC VNX5300 storage system. The study will illustrate this configuration of the solution.

The same VNX5300 configuration is used with the 500 and 600 user examples.

Two variants to the design are offered:

- Managed C-Series with Fibre Channel (FC) Storage
- Unmanaged C-Series with NFS Storage

This document details the deployment of VMware View 5.1.2 with floating assignment linked clones on VMware ESXi 5.1 using the five Cisco UCS C220 M3 rack-mount server case.

Figure 1 VMware View 5.1.2 600 User Hardware Components- Managed Cisco UCS C220 M3 Rack-Mount Servers 8 Gb Fibre Channel Storage to Cisco UCS Connectivity

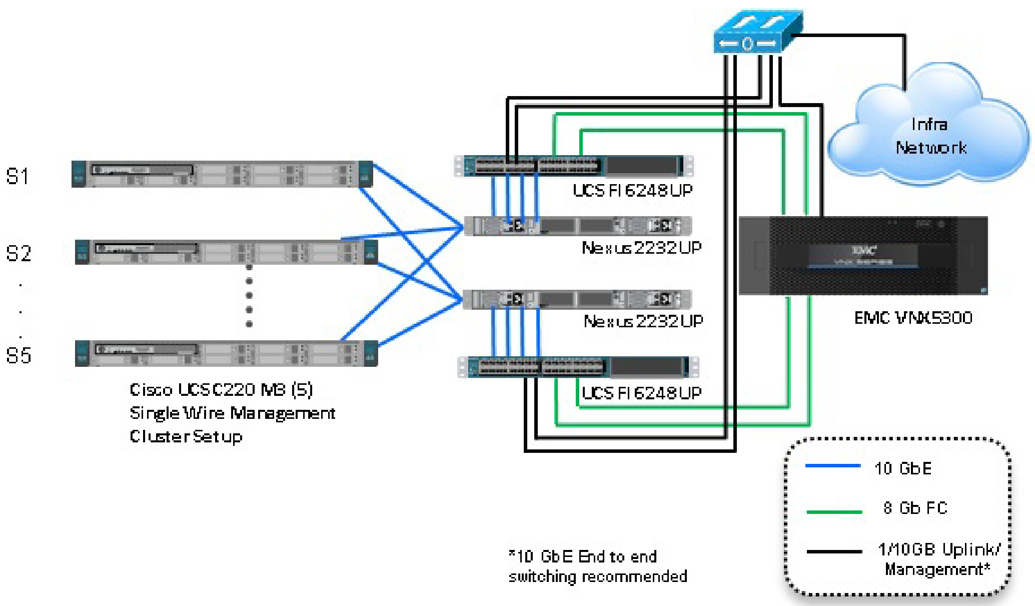

Figure 2 VMware View 5.1.2 600 User Hardware Components- Unmanaged Cisco UCS C220 M3 Rack-Mount Servers 10 Gb Ethernet NFS Storage to Cisco UCS Connectivity

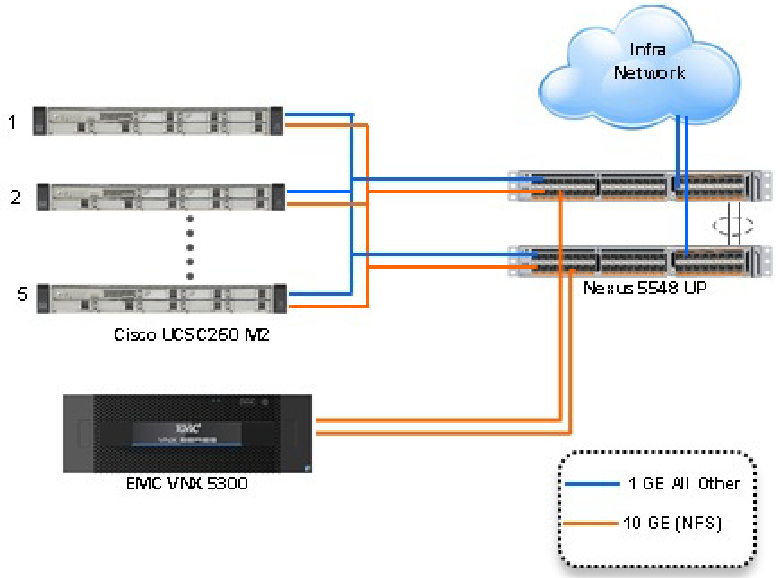

The reference configuration includes:

- Two Cisco Nexus 5548UP switches (Unmanaged C-Series NFS Variant only)

- Two Cisco UCS 6248UP Series Fabric Interconnects (Managed C-Series FC Variant Only)
- Two Cisco Nexus 2232PP Fabric Extenders (Managed C-Series FC Variant Only)
- Five Cisco UCS C220 M3 Blade servers with Intel E5-2690 processors, 256 GB RAM, and VIC1225 CNAs for 600 VDI workloads with N+1 Server fault tolerance for 600 desktops.
- Four Cisco UCS C220 M3 Blade servers with Intel E5-2690 processors, 256 GB RAM, and VIC1225 CNAs for 500 VDI workloads with N+1 Server fault tolerance for 450 desktops
- One EMC VNX5300 dual controller storage system for HA, 2 Datamovers, 600GB SAS Drives and 200GB SSD Fast Cache Drives

The EMC VNX5300 disk shelf, disk and Fast Cache configurations are detailed in Section 5.4 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 (FC Variant) Cisco UCS C220 M3	2.1 (1a) 1.4 (7b)	Embedded Management Hardware BIOS
Network	Nexus Fabric Switch	5.2(1)N1(1)	Operating System Version
Storage	EMC VNX5300	Block: 5.31.000.5.704 File: 7.0.50-2	Operating System Version
Software	Cisco UCS C220 M3 Hosts	VMware ESXi 5.0 Update 1	Operating System Version

Configuration Guidelines

The 500-600 User View 5.1.2 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 View 5.1.2 with Machine Configuration Server customer environment as stand-alone solution.

VLANs

For the 500-600 User View 5.1.2 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
ESXi_Management	132	ESXi Management	Yes
VDI	100	ESXi, N1KV Management	No
vMotion	51	vMotion	No

VMware Clusters

We utilized two VMware Clusters to support the solution and testing environment in both variants:

- Infrastructure Cluster (Active Directory, DNS, DHCP, SQL Server, File Shares for user profiles, vCenter, etc.) This would likely be the Customer's existing infrastructure cluster, adding a pair of virtual machines for View Connection Server and View Replica Server and a SQL database to an existing SQL server to support them.
- VDA Cluster (Windows 7 SP1 32-bit pooled virtual desktops)

Infrastructure Components

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

Cisco Unified Computing System (UCS)

Cisco Unified Computing System 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 3.

Figure 3 Cisco Unified Computing System Components

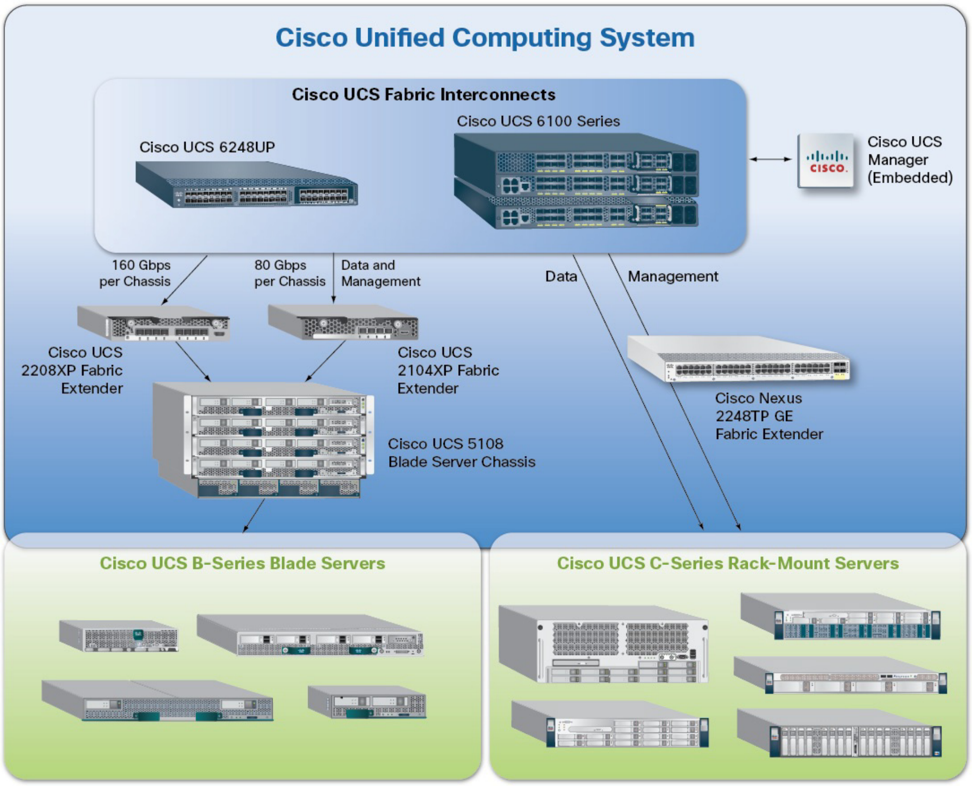
Cisco Unified Computing System
Cisco UCS Fabric Interconnects
Cisco UCS 6100 Series
Cisco UCS 6248UP
Cisco UCS Manager (Embedded)
160 Gbps per Chassis
80 Gbps per Chassis
Data and Management
Data
Management
Cisco UCS 2208XP Fabric Extender
Cisco UCS 2104XP Fabric Extender
Cisco Nexus 2248TP GE Fabric Extender
Cisco UCS 5108 Blade Server Chassis
Cisco UCS B-Series Blade Servers
Cisco UCS C-Series Rack-Mount Servers

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.

Note Only the Cisco UCS C-Series Rack-Mount Servers, specifically the Cisco UCS C220 M3, were used in both the Managed FC variant and the Unmanaged NFS variant for this study. For the Managed FC variant, Cisco UCS 6248UP Fabric Interconnects and Nexus 2232PPs were used in conjunction with the Cisco UCS C220 M3s and the EMC VNX5300. For the Unmanaged NFS variant, Nexus 5548UPs were used in conjunction with the Cisco UCS C220 M3s and the EMC VNX5300.

The components of the Unified Computing System and Nexus switches that were used in the study are discussed below.

Cisco Fabric Interconnects (Managed FC Variant Only)

Cisco UCS Fabric Interconnects create a unified network, storage and management fabric throughout the Cisco UCS. 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 1 Tbps 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+. The expansion module was not required for this study.

Cisco UCS C220 M3 Rack-Mount Server

Cisco Unified Computing System is the first truly unified data center platform that combines industry-standard, x86-architecture blade and rack servers with networking and storage access into a single system. Key innovations in the platform include a standards-based unified network fabric, Cisco Virtualized Interface Card (VIC) support, and Cisco UCS Manager Service Profile and Direct Storage Connection support. The system uses a wire- once architecture with a self-aware, self-integrating, intelligent infrastructure that eliminates the time-consuming, manual, error-prone assembly of components into systems.

Managed Cisco UCS C-Series Rack-Mount Servers reduce total cost of ownership (TCO) and increase business agility by extending Cisco Unified Computing System™ innovations to a rack-mount form factor. These servers:

- Can be managed and provisioned centrally using Cisco UCS Service Profiles with Cisco UCS Manager 2.1(1a,) Cisco UCS Fabric Interconnects and Nexus 2232PP Fabric Extenders
- Offer a form-factor-agnostic entry point into the Cisco Unified Computing System, which is a single converged system with configuration automated through integrated, model-based management

- Simplify and speed deployment of applications
- Increase customer choice with unique benefits in a familiar rack package
- Offer investment protection through the capability to deploy them either as standalone servers or as part of the Cisco Unified Computing System

Note This study highlights the use of Managed Cisco UCS C-Series Rack-Mount servers in the FC variant. The alternative NFS variant utilizes the Cisco UCS C220 M3 servers in stand-alone mode.

Cisco UCS Virtual Interface Card (VIC) Converged Network Adapter

A Cisco® innovation, the Cisco UCS Virtual Interface Card (VIC) 1225 (Figure 1) is a dual-port Enhanced Small Form-Factor Pluggable (SFP+) 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) card designed exclusively for Cisco UCS C-Series Rack Servers. With its half-height design, the card preserves full-height slots in servers for third-party adapters certified by Cisco. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases.

Cisco UCS Virtual Interface Card 1225

The card 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 1225 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 4 *Cisco UCS VIC M81KR Converged Network Adapter*

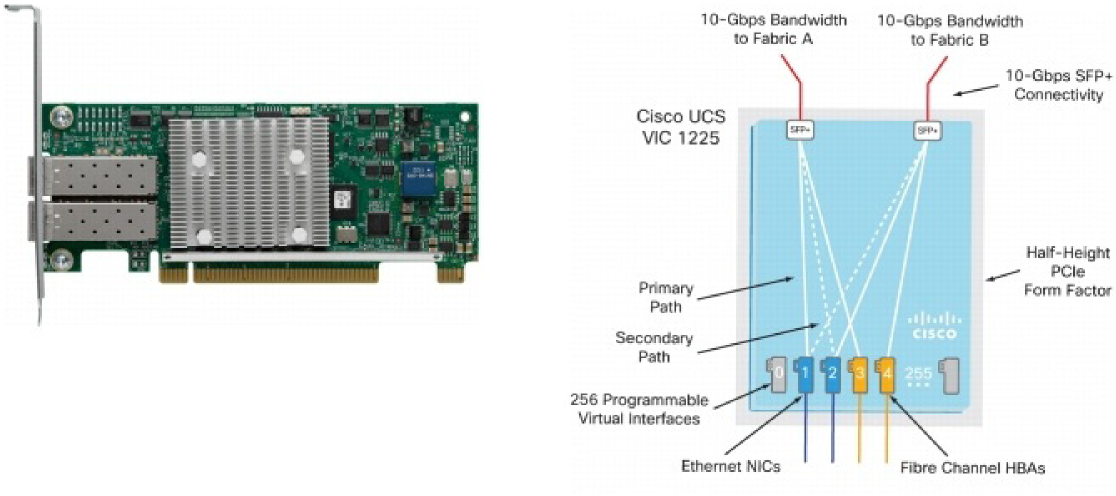
10-Gbps Bandwidth to Fabric A
10-Gbps Bandwidth to Fabric B
10-Gbps SFP+ Connectivity
Cisco UCS VIC 1225
SFP+
SFP+
Half-Height PCIe Form Factor
Primary Path
Secondary Path
cisco
256 Programmable Virtual Interfaces
0
1
2
3
4
255
Ethernet NICs
Fibre Channel HBAs

Note The Cisco UCS VIC 1225 virtual interface cards are deployed in the Cisco UCS C-Series C220 M3 rack-mount servers.

VMware View 5.1

VMware View simplifies desktop and application management while increasing security and control. 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.

VMware View 5 Features

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 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.

The following features of VMware View 5 provide measurable benefits for desktop virtualization.

- View Media Services for 3D Graphics enables basic 3D applications like Windows Aero, Office 2010 or those requiring OpenGL or DirectX in View desktops without the need for specialized graphics cards or client devices. (Requires vSphere 5 or later as the platform.)
- View Media Services for Integrated Unified Communications enables an integrated VOIP and View desktop experience for the end-user with an architecture to provide optimized performance for both the desktop and unified communication.
- View Persona Management (View Premier editions only) dynamically associates a user persona to stateless floating desktops. IT administrators can streamline migration from physical to stateless virtual desktops while preserving user settings
- View Storage Accelerator optimizes storage load by caching common image blocks when reading virtual desktop images to decrease storage load during boot storms.
- vSphere 5 support ensures that View desktop services run on the industry's most complete and robust cloud infrastructure platform for flexible, reliable IT services.
- PCoIP Extension Services allow WMI based tools to collect over 20 session statistics for monitoring, trending and troubleshooting end-user support issues.
- PCoIP Optimization Controls deliver protocol efficiency and enables IT administrators to configure the bandwidth settings by use case, user or network requirements to consume up to 75% less bandwidth.
- PCoIP Continuity Services deliver a seamless end-user experience regardless of network reliability by detecting interruptions and automatically reconnecting the session.

Enhancements in View 5.1

VMware View 5.1 continues to build upon the advancements released in View 5. 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

Storage Optimization—New in View 5.1, View Storage Accelerator is a technology that reduces storage loads generated by peak VDI storage reads caching the common blocks of desktop images into local host memory. The Accelerator leverages a VMware vSphere (version 5.0 or later) platform feature called Content Based Read Cache (CBRC) implemented inside the ESX/ESXi hypervisor. When enabled for specific VMs, the host hypervisor scans the storage disk blocks to generate digests of the block contents. When these blocks are read into the hypervisor, they are cached in the host based CBRC. Subsequent reads of blocks with the same digest will be served from the in-memory cache directly. This significantly improves the desktop performance, especially during boot storms or anti-virus scanning storms when a large number of blocks with identical contents are read.

Simplify Desktop Management and Deployment

View Persona Management—To help with physical to virtual desktop migrations, we've extended View Persona Management to physical desktops. This new feature also enables Windows XP to Windows 7 migration. The View Persona Management agent can be installed without the VMware View agent on physical desktops belonging to the same licensed VMware View desktop end-users.

During a physical to virtual migration, an administrator can first install View Persona Management on the physical desktop. When the same user uses a virtual desktop with Persona Management enabled, user data and settings are automatically synchronized. We also extend Persona Management to support a one-time Windows XP to Windows 7 migration.

VMware vCenter Operations for View—New for VMware View 5.1, we have integrated with our management products to give you VMware vCenter Operations (vCOps) Manager for View. Optimized for virtual desktop deployments, VMware vCenter Operations Manager for View provides end-to-end monitoring of desktops and users, displayed with user friendly dashboards, to help identify, troubleshoot, and trend potential issues.

View Administrator Enhancements—Some customers deploy VMware View in a restrictive environment in which write access to the Active Directory is prohibited. In this new version of View, an administrator can set a configuration option to reuse existing machine accounts in AD during the provisioning process.

As the numbers of VMware View deployments grow, our customers are expanding the scale of their View virtual desktop programs. Enhancements made in VMware View 5.1, make management at scale easier. VMware View Composer server in View 5.1 can be installed in a standalone server. An administrator can also configure VMware View Connection Server (via vdmadmin command line tool) to log events in syslog rather than a database.

Last but not least, View Admin UI response time has been greatly improved in a large-scale environment.

View Administrator Language Support—To serve a growing international customer base, the View Admin UI is localized to five major non-English languages: French, German, Japanese, Korean, and Simplified Chinese.

Create the Best End-User Experience

USB Enhancements—The USB redirect feature for the Windows client has been reworked. The new USB feature no longer requires device driver to be installed on the client side. A generic USB arbitrator is implemented on the client side, while a proper USB hub is implemented in the agent. This allows VMware View to support a much broader range of USB devices while supporting fine-grained remote device policy (e.g. enable/disable mass storage file copy) even on multi-function USB devices.

RADIUS Support —Based on customer feedback, the security authentication support in VMware View has been extended to other two-factor authentication vendors leveraging a RADIUS client in the View 5.1 Connection Server. This gives you more choice when implementing single sign-on or security tokens into your virtual desktops.

Continued PCoIP Enhancements—Continuous enhancements the PCoIP remote protocol have been made following the significant progress made in version 5.0. The optimal remote protocol performance cannot be achieved with code improvement alone. To help customers make the right choice in protocol with proper performance tuning, we published a white paper comparing the tuning and test results of all state-of-the-art remote protocols:
http://www.vmware.com/files/pdf/techpaper/PCoIPvHDXsinglesession03-05-12.pdf

VMware View 5.1 Hosted Virtual Desktop (HVD) 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 View 5.1: Automated, Manual, and Terminal Services Pools. These pool types are discussed below.

Automated HVD pools use 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 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 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 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 View users. The 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 View Composer linked clones over the PCoIP protocol.

View Persona Manager was not deployed.

The following figure shows the logical architecture for a View 5.1 deployment, including the optional related product; Thin App. Thin App provides application streaming capability and is not included in this study.

Figure 5 VMware View 5 Architecture Diagram

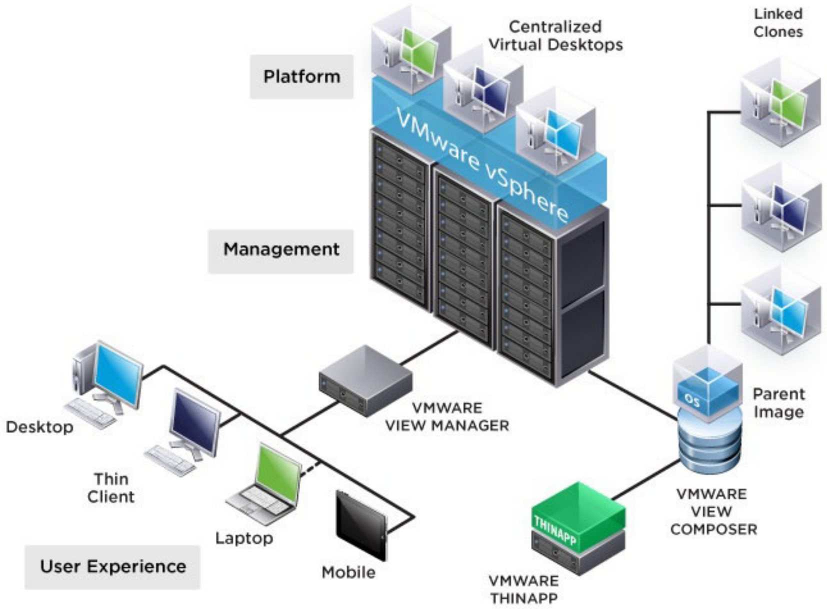

EMC VNX Series

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 re-purposing.
- 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 VNX 5300 Used in Testing

EMC VNX 5300 provides storage by using FC (SAN) or IP (NAS) connections for virtual desktops, and infrastructure virtual machines such as VMware View controllers, VMware vCenter Servers, Microsoft SQL Server databases, and other supporting services. Optionally, user profiles and home directories are redirected to CIFS network shares on the VNX5300.

VMware ESXi 5.1

VMware, Inc. provides virtualization software. VMware's enterprise software hypervisors for servers-VMware ESX, Vmware ESXi, and VSphere-are bare-metal embedded 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 6). 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 6 The Evolving Workplace Landscape

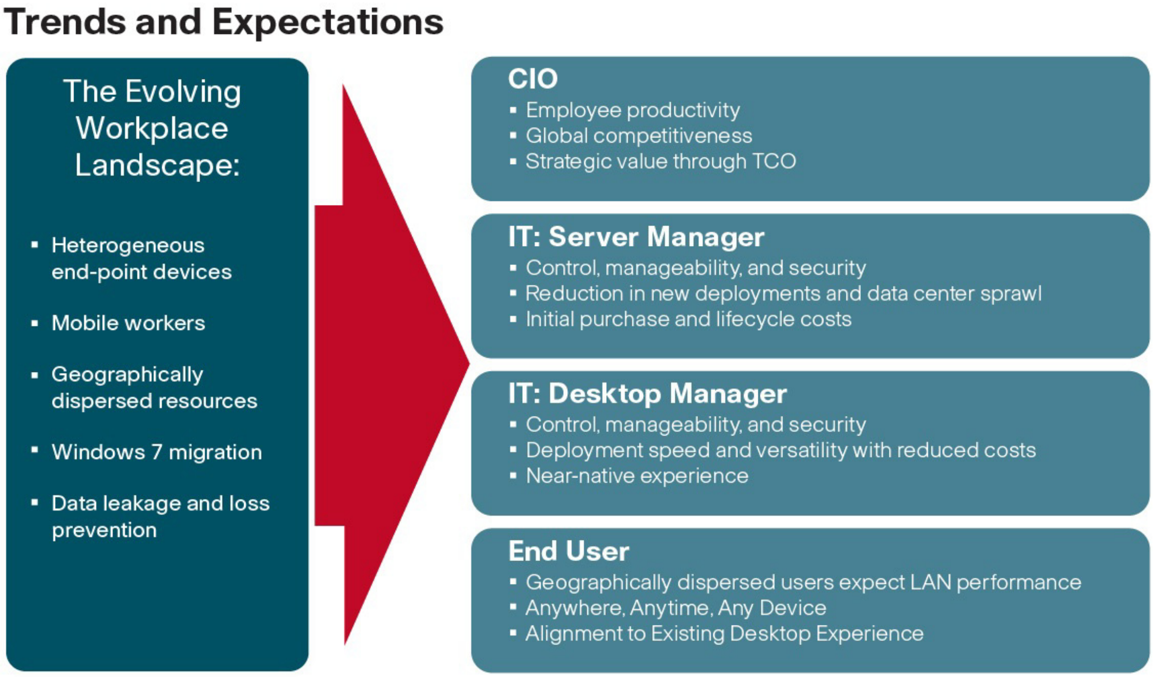
Trends and Expectations
The Evolving Workplace Landscape:
▪ Heterogeneous end-point devices
▪ Mobile workers
▪ Geographically dispersed resources
▪ Windows 7 migration
▪ Data leakage and loss prevention
CIO
▪ Employee productivity
▪ Global competitiveness
▪ Strategic value through TCO
IT: Server Manager
▪ Control, manageability, and security
▪ Reduction in new deployments and data center sprawl
▪ Initial purchase and lifecycle costs
IT: Desktop Manager
▪ Control, manageability, and security
▪ Deployment speed and versatility with reduced costs
▪ Near-native experience
End User
▪ Geographically dispersed users expect LAN performance
▪ Anywhere, Anytime, Any Device
▪ Alignment to Existing Desktop Experience

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 7).

Figure 7 VMware View 5.1.2 on Single-Wire Managed Cisco UCS C220 M3 Fibre Channel Variant

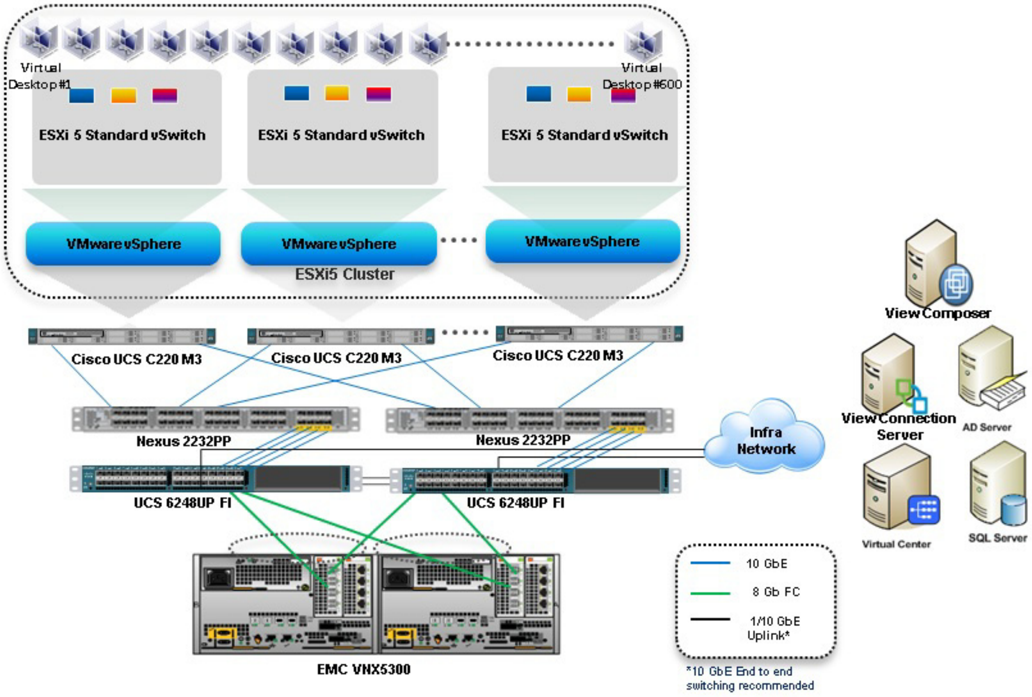

Figure 8 VMware View 5.1.2 on Unmanaged Cisco UCS C220 M3 NFS Variant

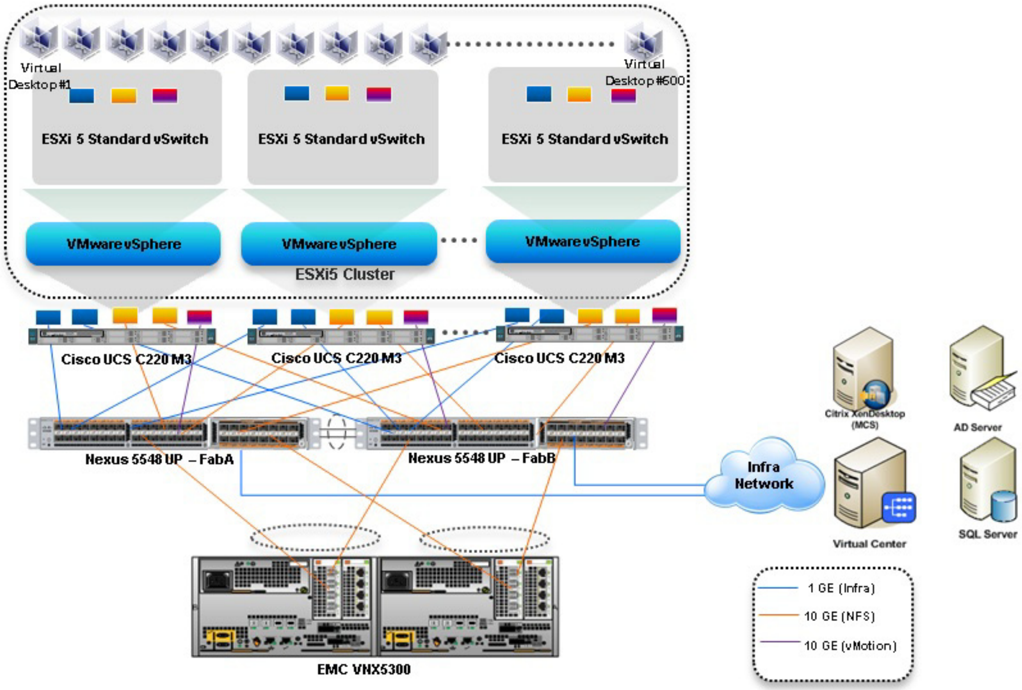

Simplified

Cisco UCS 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 UCS 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 UCS along with the reduced cabling and port count due to the unified fabric and unified ports of Cisco UCS 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 (View 5.1) 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 UCS 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 blade 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, meaning there will be fewer servers needed, 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:

- 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 (e.g., 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 (e.g. 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 (e.g., 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 (e.g., 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 RDP be used for Hosted Shared Server Desktops or exclusively View Hosted Virtual Desktops?
- Are there ThinApp streamed applications planned?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will SSD drives or Cisco Storage Accelerator be used for VDI IOPS?
- Will there be external 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?

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 5548UP Switch (Unmanaged FC Variant Only)

Two Cisco Nexus 5548UP access switches are 1RU 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 and one expansion slot. The Cisco Nexus 5500 platform can be equipped with an expansion module that can be used to increase the number of 10 Gigabit Ethernet and FCoE ports or to connect to Fibre Channel SANs with 8/4/2/1-Gbps Fibre Channel switch ports, or both. (Not required for this study.)

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 2232PP Fabric Extender (Managed FC Variant Only)

The Cisco Nexus 2232PP 10G provides 32 10 Gb Ethernet and Fibre Channel Over Ethernet (FCoE) Small Form-Factor Pluggable Plus (SFP+) server ports and eight 10 Gb Ethernet and FCoE SFP+ uplink ports in a compact 1 rack unit (1RU) form factor.

Two Nexus 2232PP 10GE Fabric Extenders were deployed to provide cluster-mode single wire management to the Cisco UCS C220 M3 rack servers.

Four of eight available 10 GbE uplinks from each Nexus 2232 were utilized to provide 40 Gb of bandwidth between the UCS 6248UP Fabric Interconnects and the Cisco UCS C220 M3 rack servers.

Architecture and Design of View 5.1 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: a physical device with a locally installed operating system.
- **Hosted Shared Desktop** A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 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 (XenServer, Hyper-V or 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 (HVD) Design Fundamentals

VMware View 5.1 can be used to deliver a variety of virtual desktop configurations. When evaluating a Hosted HVD deployment, consider the following:

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.

For this study, we used the PCoIP protocol.

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 ensure 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 https://www.vmware.com/support/viewclients/doc/viewclients_pubs.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.

Following are 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.

You can download RDC versions from the Microsoft Web site.

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 Even though the article talks about Windows Vista, the product where the significant changes to Roaming User Profiles and Folder Redirection were first implemented, it applies to Windows 7 as well.

VMware Persona Management

You can use View Persona Management with View desktops and with 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 Log Off 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 log off time and copy all changes to the server at log off.
- 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 log off, 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

View 5.1 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 View 5.1 Deployment

There are several elements that go into the design of a successful View 5.1 environment. This section covers those topics at a high level. Readers should consult the VMware View Architecture Planning guide for View 5.1 at the following URL for more details:

http://pubs.vmware.com/view-51/topic/com.vmware.ICbase/PDF/view-51-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, we will test a single user group (knowledge workers) and have identified the application workload this group will run, which is based on the Login VSI 3.7 medium workload (with flash.) We will also use virtual machines as our desktop source.

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.

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

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 View 5.1 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.

For this study, we utilized the replicas and linked clones on different datastores technique.

Hosted Virtual Desktop Infrastructure

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

Figure 9 View Desktop Infrastructure

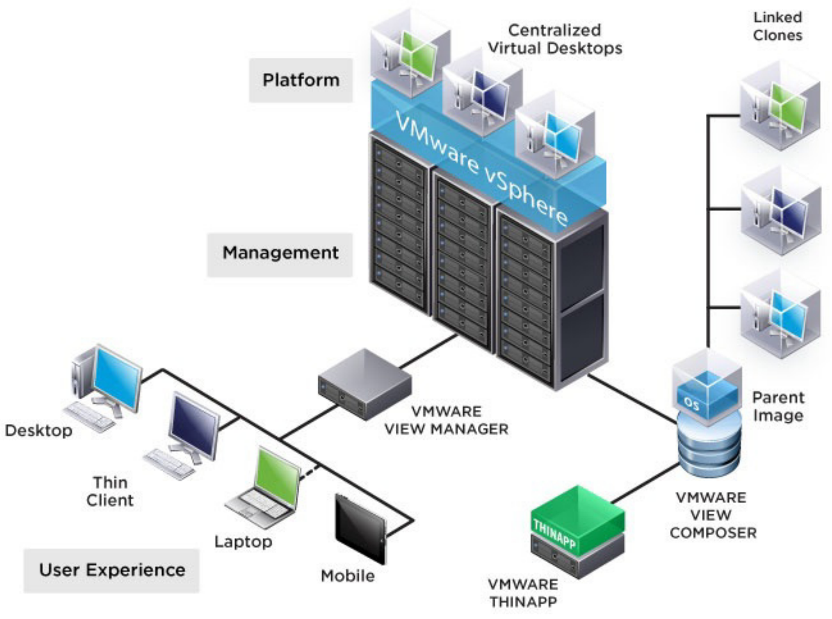

Learn more about VMware View 5.1 planning and design at the following location:

http://pubs.vmware.com/view-51/topic/com.vmware.ICbase/PDF/view-51-architecture-planning.pdf

Desktop Delivery Base Image Creation

Microsoft Windows 7 Golden Image Creation

Create a 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. We used windows 7 32 bit OS for our testing.

For the virtual machine, the following parameters were used.

- Memory: 1536Mb
- Processor: 1 vCPU
- Hard Disk: 18 GB
- Network Adapter: Attached to Standard Switch

Note The floppy drive was removed from the virtual machine.

2. Right-click the Windows 7 Golden Image properties and select Hardware Tab to attach the Windows -7 SP 1 ISO image
3. Click OK.

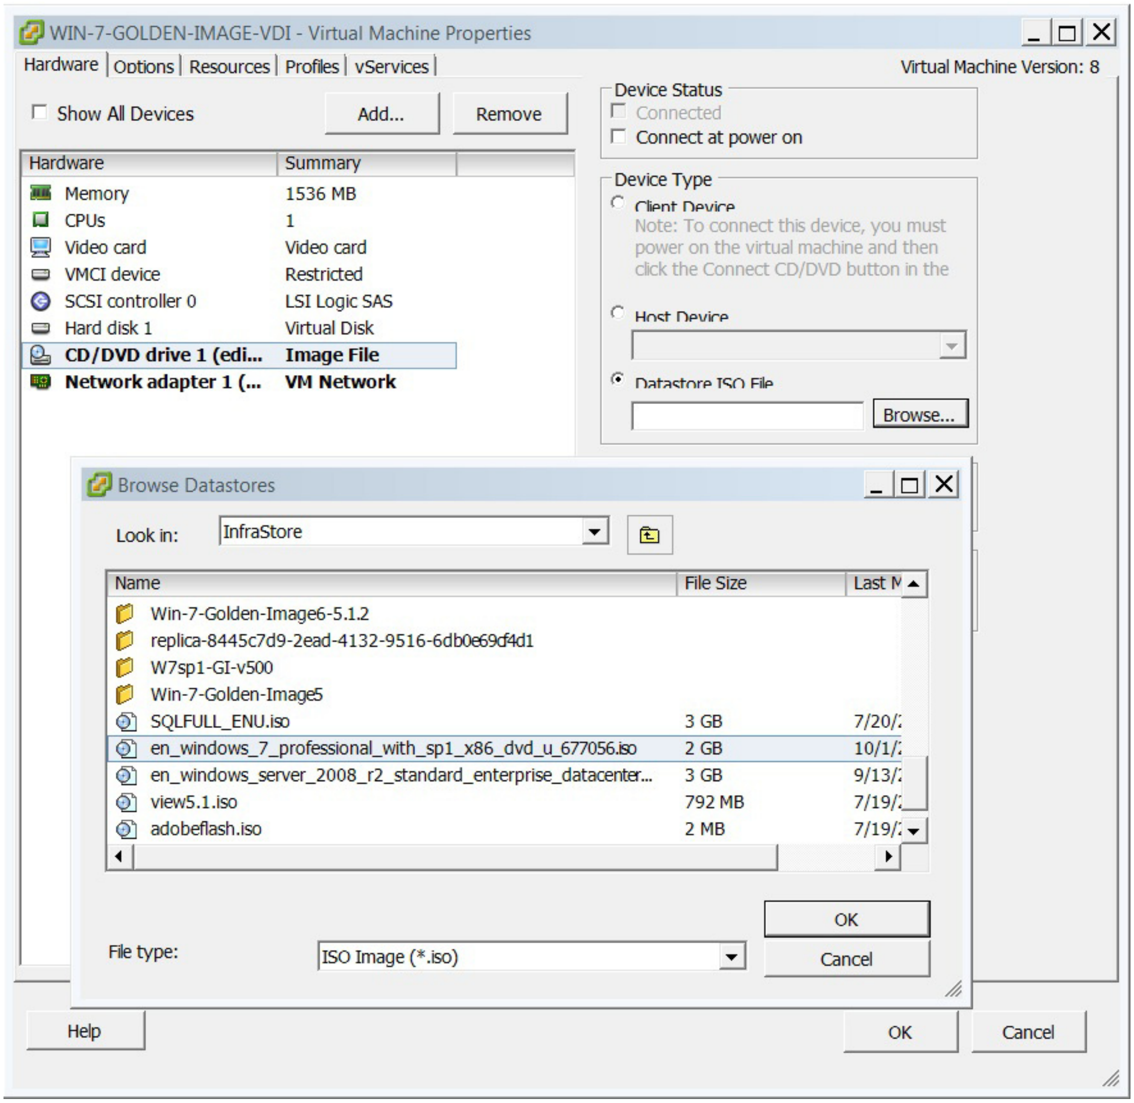

☐ Show All Devices

Hardware	Summary
Memory	1536 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic SAS
Hard disk 1	Virtual Disk
CD/DVD drive 1 (edi...	**Image File**
Network adapter 1 (...	**VM Network**

Device Status

☐ Connected

☐ Connect at power on

Device Type

○ Client Device

Note: To connect this device, you must power on the virtual machine and then click the Connect CD/DVD button in the

○ Host Device

● Datastore ISO File

Look in: InfraStore

Name	File Size	Last M
Win-7-Golden-Image6-5.1.2		
replica-8445c7d9-2ead-4132-9516-6db0e69df4d1		
W7sp1-GI-v500		
Win-7-Golden-Image5		
SQLFULL_ENU.iso	3 GB	7/20/
en_windows_7_professional_with_sp1_x86_dvd_u_677056.iso	2 GB	10/1/
en_windows_server_2008_r2_standard_enterprise_datacenter...	3 GB	9/13/
view5.1.iso	792 MB	7/19/
adobeflash.iso	2 MB	7/19/

File type: ISO Image (*.iso)

4. Right-click on 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 created and configure IP Address, join the domain and Restart the Virtual Machine.
6. Shutdown the Windows 7 golden image virtual machine; this completes the process of creating the virtual machine.

Optimization of Base Windows 7 SP1 Virtual Machine

1. Refer to the following link about how to optimize Windows 7 SP1 32 bit virtual machine for View 5.1.

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

Install View 5.1 Virtual Desktop Agent software.

1. Download the software from the link below:

   https://my.vmware.com/web/vmware/details?productId=268&downloadGroup=VIEW-512-PREMIERE

2. Open installer VMware-viewagent-5.1.2-928164.exe for 32bit OS or VMware-viewagent-x86_64-5.1.2-928164.exe 64bit OS.
3. Click Next in the installer wizard.

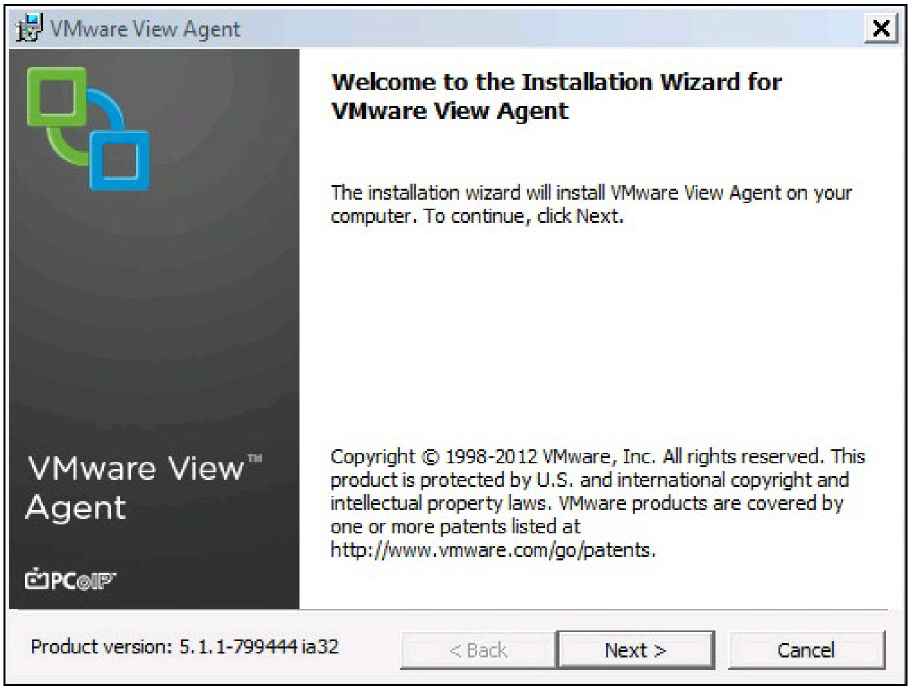

Welcome to the Installation Wizard for VMware View Agent

The installation wizard will install VMware View Agent on your computer. To continue, click Next.

Copyright © 1998-2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 5.1.1-799444 ia32

4. Read the License agreement and click Next.

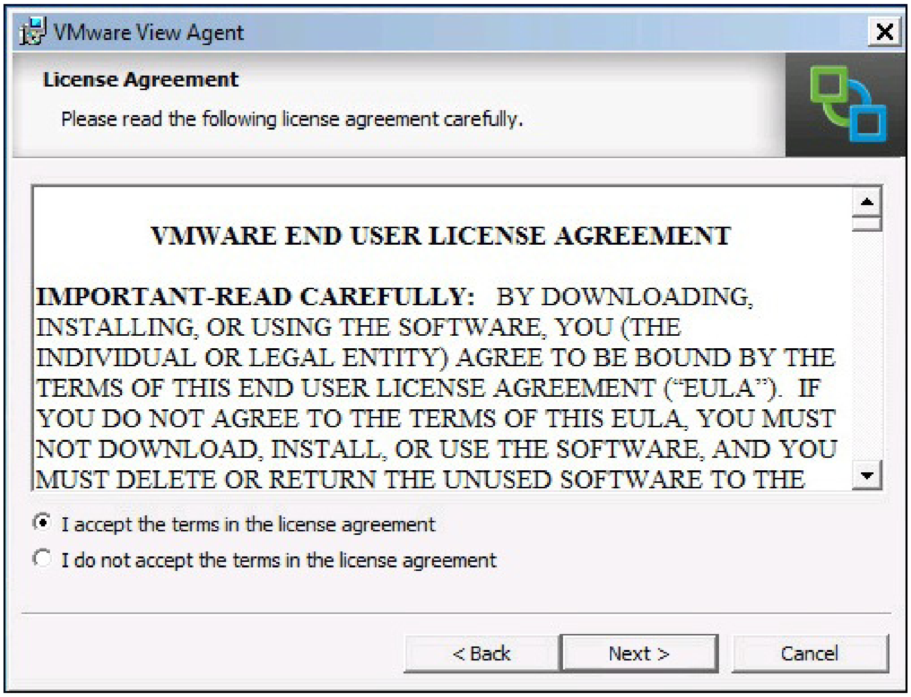

VMware View Agent

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

5. Select default setup or make necessary and click Next.

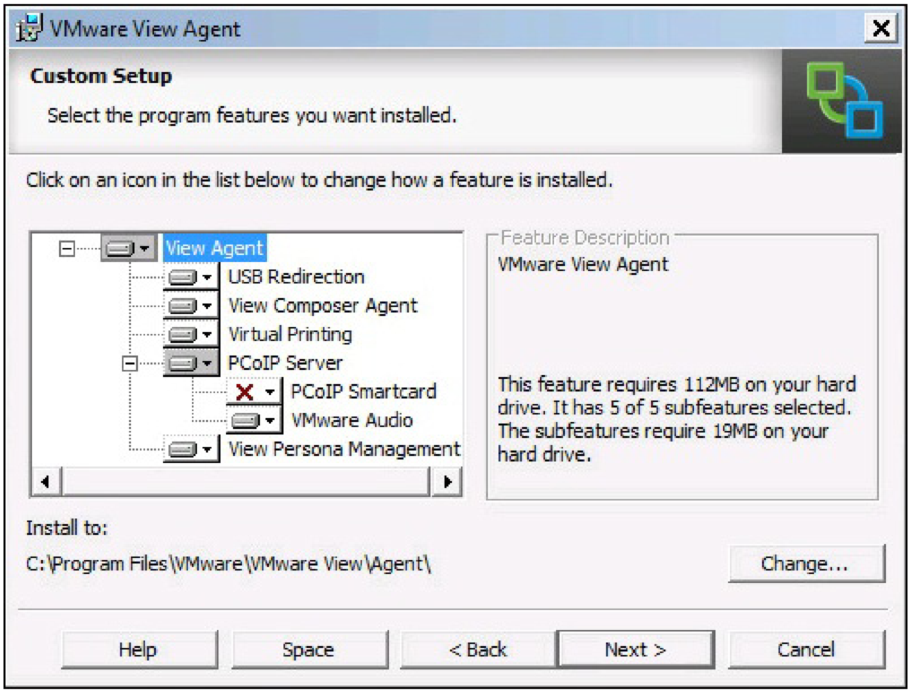

VMware View Agent

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- View Agent
  - USB Redirection
  - View Composer Agent
  - Virtual Printing
  - PCoIP Server
    - ✗ PCoIP Smartcard
    - VMware Audio
  - View Persona Management

Feature Description

VMware View Agent

This feature requires 112MB on your hard drive. It has 5 of 5 subfeatures selected. The subfeatures require 19MB on your hard drive.

Install to:

C:\Program Files\VMware\VMware View\Agent\

6. Click Install for ready to install program wizard.

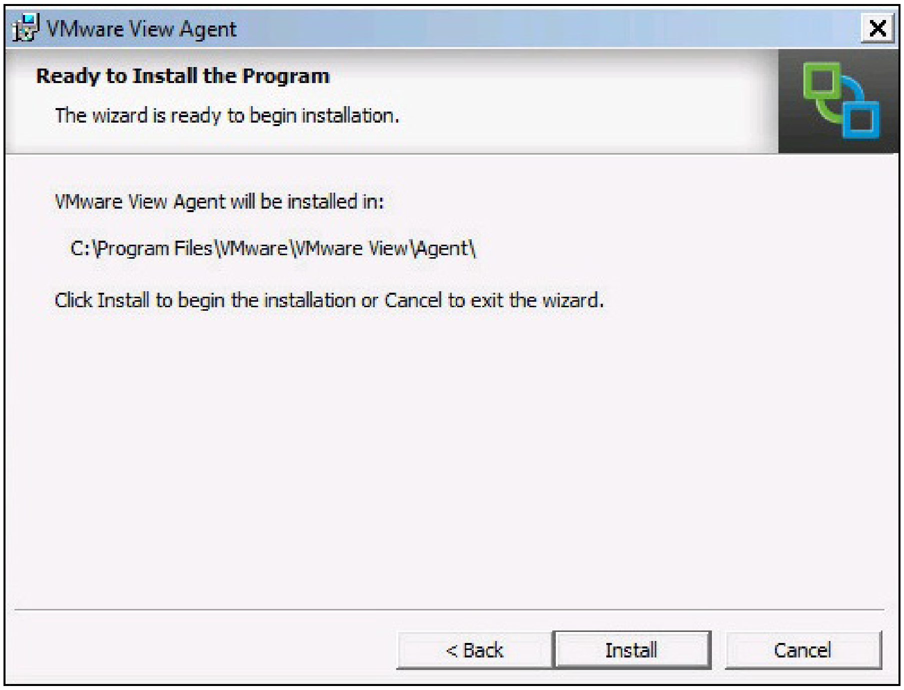

VMware View Agent

Ready to Install the Program

The wizard is ready to begin installation.

VMware View Agent will be installed in:

C:\Program Files\VMware\VMware View\Agent\

Click Install to begin the installation or Cancel to exit the wizard.

Install Additional Software

1. Install additional software required in your Windows 7 Golden Image.
   a. For our testing, we installed Microsoft Office 2010
   b. Login VSI Target software package to facilitate workload testing.
2. Reboot the Virtual Machine.
3. Install service packs and hot fixes required for the additional software components that were added.
4. Reboot the Virtual Machine.

Perform Additional View 5.1 Configuration

Create a Snapshot for the Virtual Machine

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

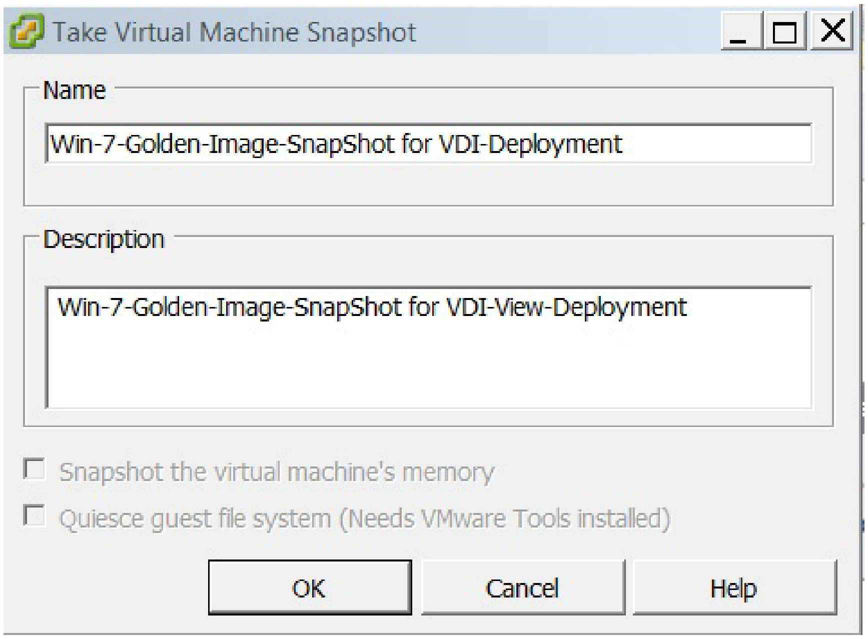

Name

Win-7-Golden-Image-SnapShot for VDI-Deployment

Description

Win-7-Golden-Image-SnapShot for VDI-View-Deployment

□ Snapshot the virtual machine's memory

□ Quiesce guest file system (Needs VMware Tools installed)

Create Customization Specifications or Virtual Desktops

1. Right-click on the powered off virtual machine after taking a snapshot and select Template and click on convert to template.
2. Provide a name to the template and provide the host /cluster, datastore details.
3. Select Guest Customization and check the radial button for Customize using the Customization wizard. click Next.

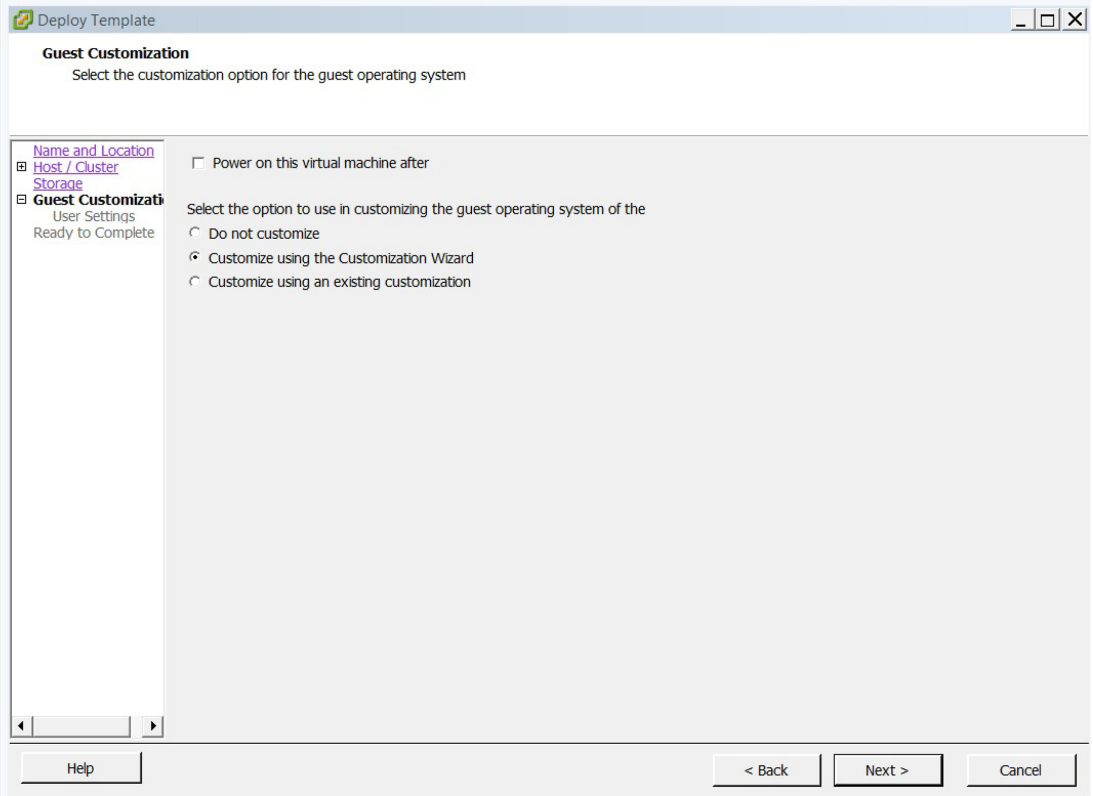
Deploy Template

Guest Customization

Select the customization option for the guest operating system

□ Power on this virtual machine after

Select the option to use in customizing the guest operating system of the

○ Do not customize

● Customize using the Customization Wizard

○ Customize using an existing customization

4. Select appropriate name and organization. Click Next.

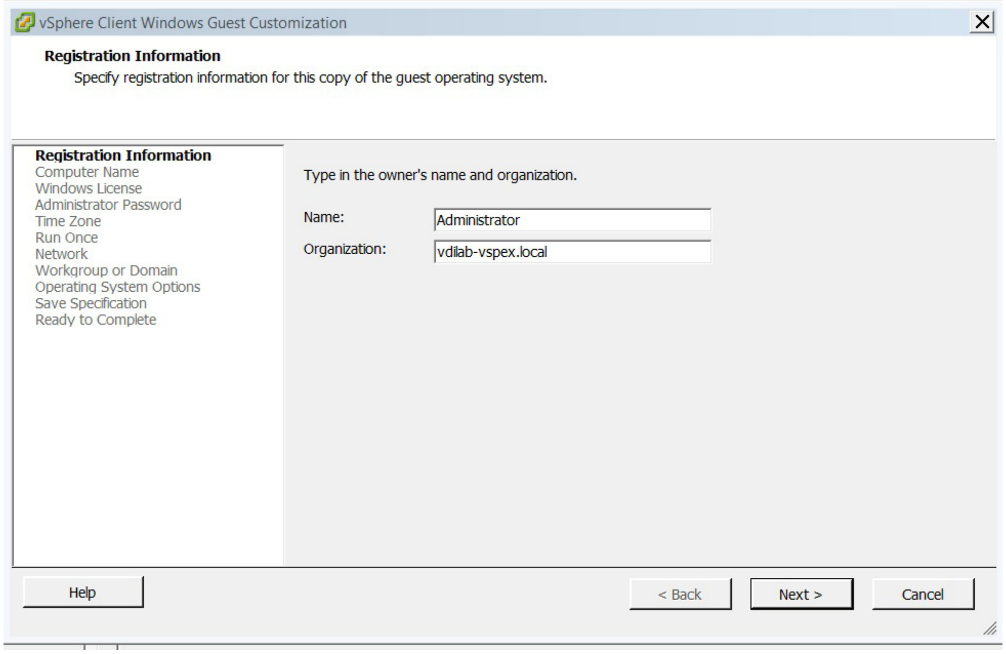
vSphere Client Windows Guest Customization

Registration Information

Specify registration information for this copy of the guest operating system.

Type in the owner's name and organization.

Name: Administrator

Organization: vdilab-vspex.local

5. Select the radio button for use virtual machine name. Click Next.

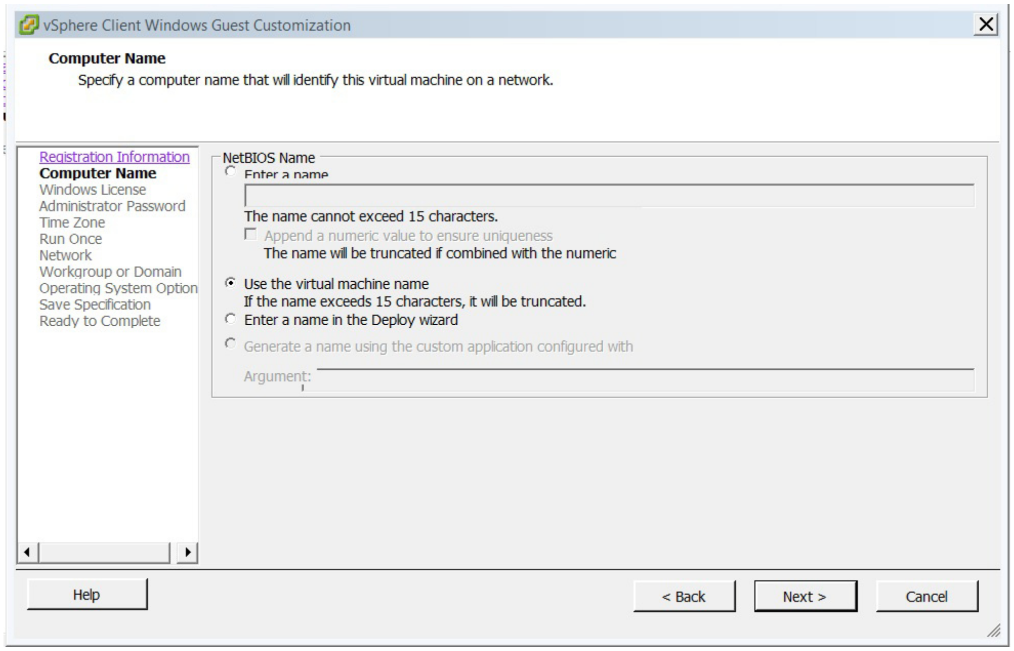

vSphere Client Windows Guest Customization

Computer Name

Specify a computer name that will identify this virtual machine on a network.

NetBIOS Name

○ Enter a name

The name cannot exceed 15 characters.

□ Append a numeric value to ensure uniqueness

The name will be truncated if combined with the numeric

● Use the virtual machine name

If the name exceeds 15 characters, it will be truncated.

○ Enter a name in the Deploy wizard

○ Generate a name using the custom application configured with

Argument:

6. Specify Volume License Key for Windows 7 and select per seat or per server maximum option. Click Next.

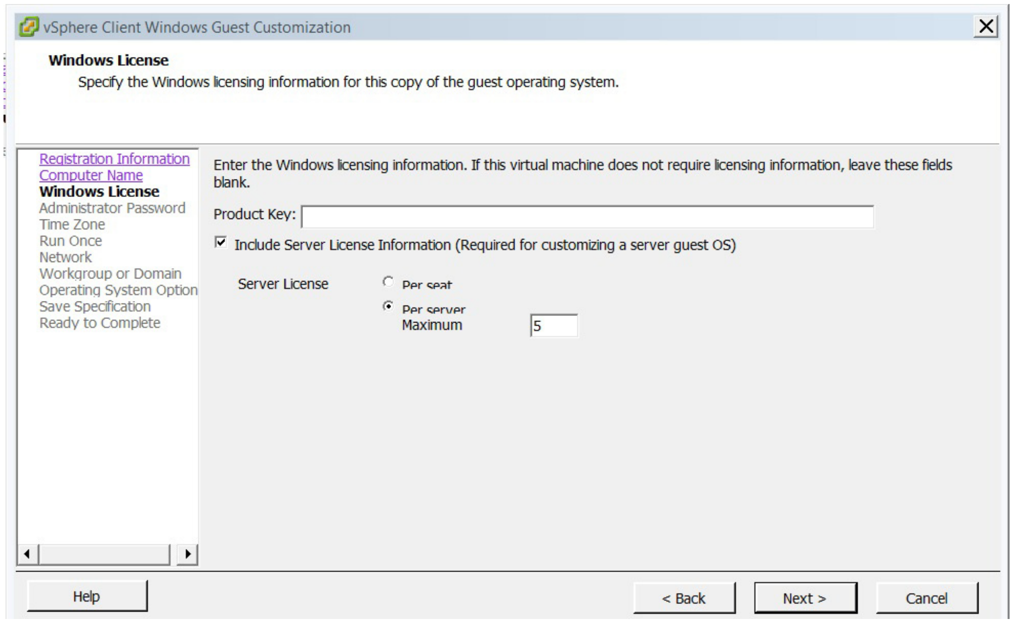

vSphere Client Windows Guest Customization

Windows License

Specify the Windows licensing information for this copy of the guest operating system.

Enter the Windows licensing information. If this virtual machine does not require licensing information, leave these fields blank.

Product Key:

☑ Include Server License Information (Required for customizing a server guest OS)

Server License ○ Per seat

● Per server

Maximum 5

7. Enter credential for administrator account. Click Next.

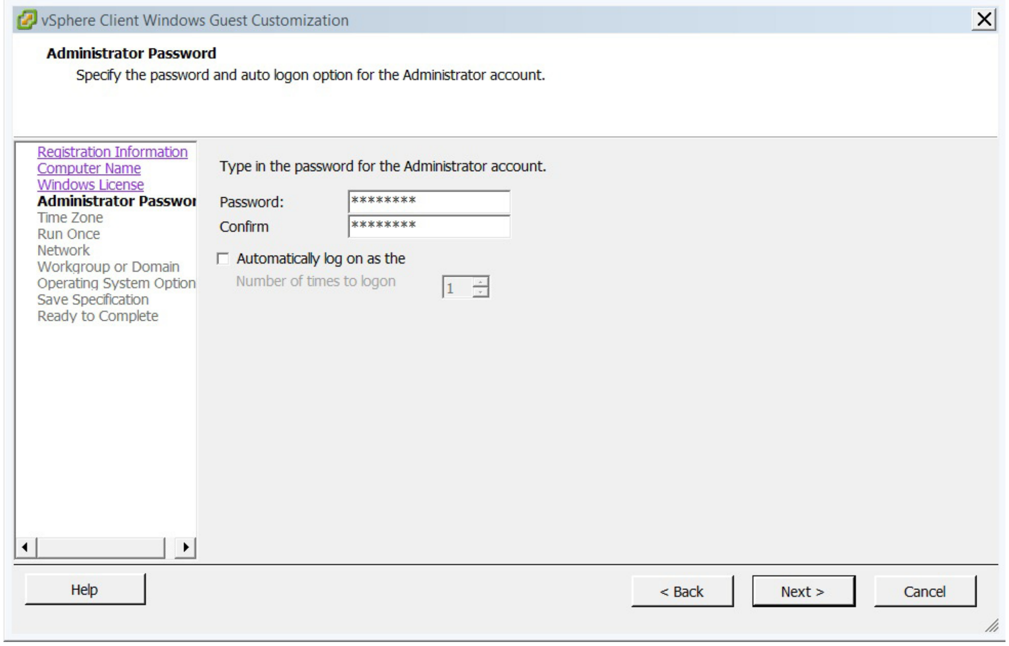

vSphere Client Windows Guest Customization

Administrator Password

Specify the password and auto logon option for the Administrator account.

Type in the password for the Administrator account.

Password: ********

Confirm ********

☐ Automatically log on as the

Number of times to logon 1

8. Select appropriate time zone, then click Next.

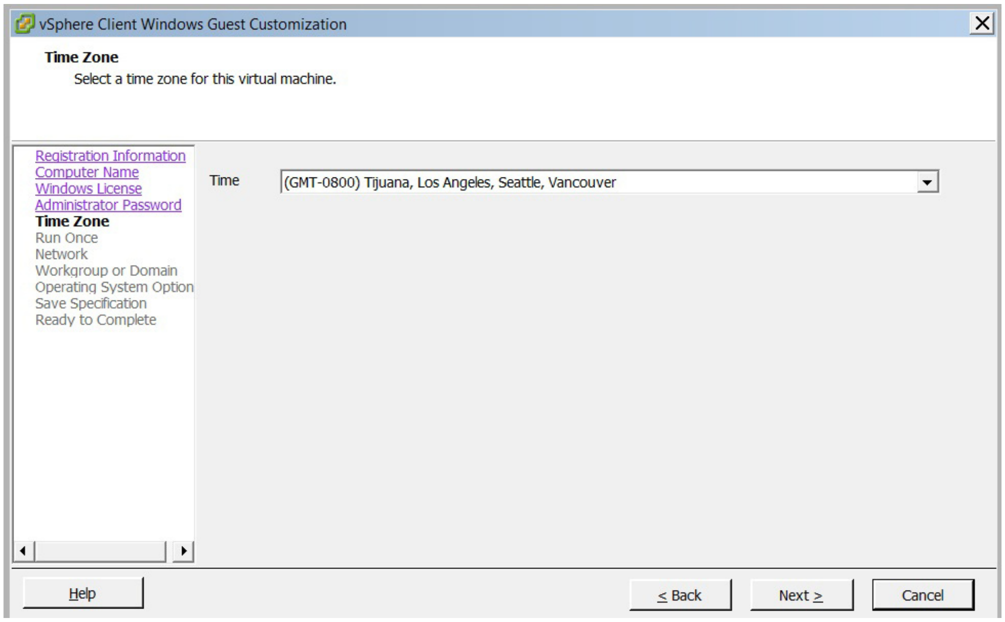

vSphere Client Windows Guest Customization

Time Zone

Select a time zone for this virtual machine.

Time (GMT-0800) Tijuana, Los Angeles, Seattle, Vancouver

9. For network select typical settings for virtual desktop networking. Click Next.

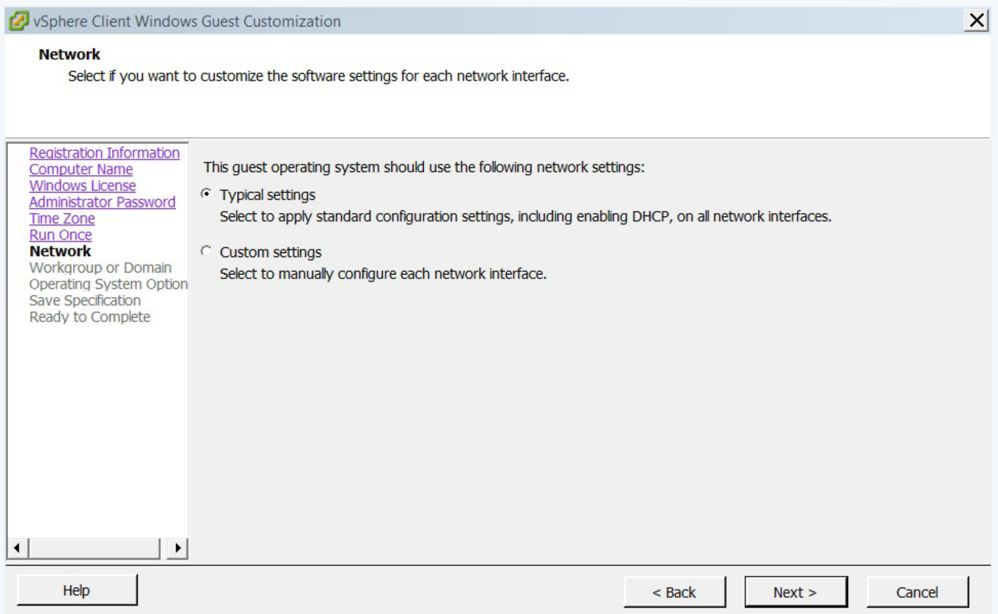

vSphere Client Windows Guest Customization

Network

Select if you want to customize the software settings for each network interface.

This guest operating system should use the following network settings:

● Typical settings
Select to apply standard configuration settings, including enabling DHCP, on all network interfaces.

○ Custom settings
Select to manually configure each network interface.

10. For the Workgroup or Domain, select the radial button for windows server. Enter domain for the environment.

11. Specify the user account name password. Click Next.

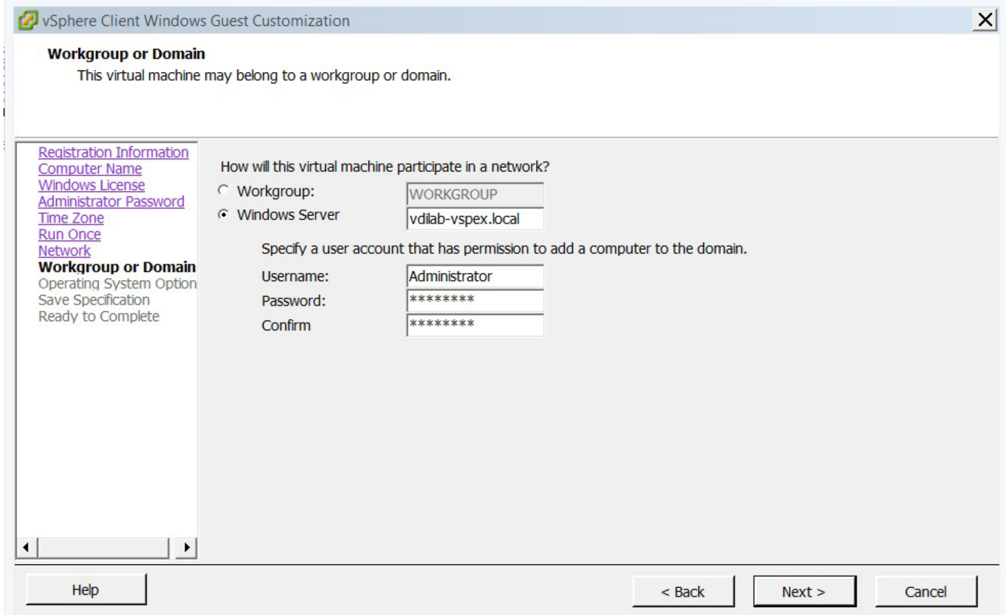

vSphere Client Windows Guest Customization

Workgroup or Domain

This virtual machine may belong to a workgroup or domain.

How will this virtual machine participate in a network?

○ Workgroup: WORKGROUP
● Windows Server vdilab-vspex.local

Specify a user account that has permission to add a computer to the domain.

Username: Administrator
Password: ********
Confirm ********

12. Check box to generate new security ID click Next.

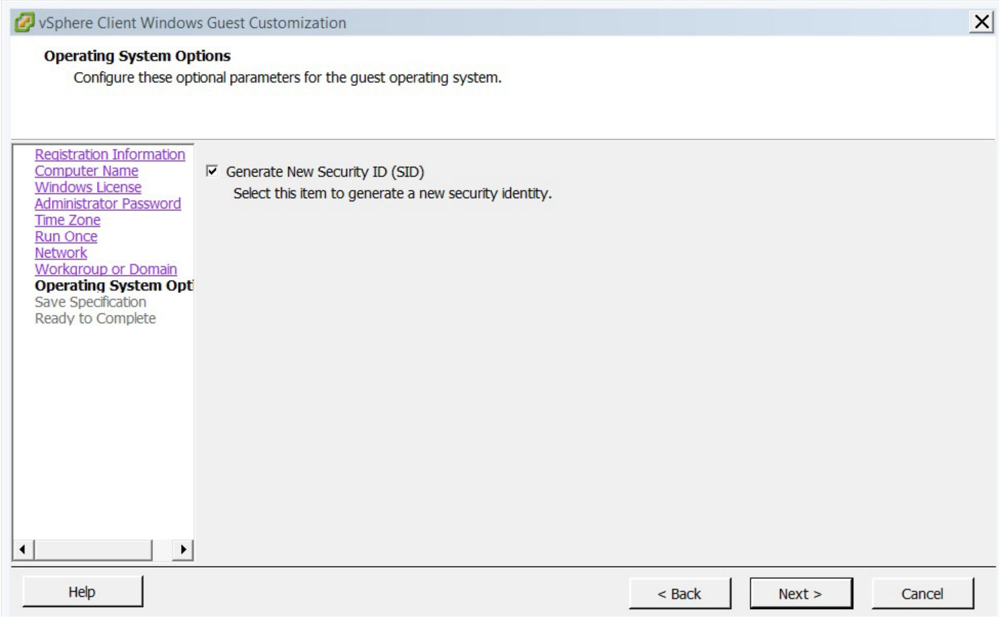

Operating System Options

Configure these optional parameters for the guest operating system.

☑ Generate New Security ID (SID)

Select this item to generate a new security identity.

13. Save the customization specification.

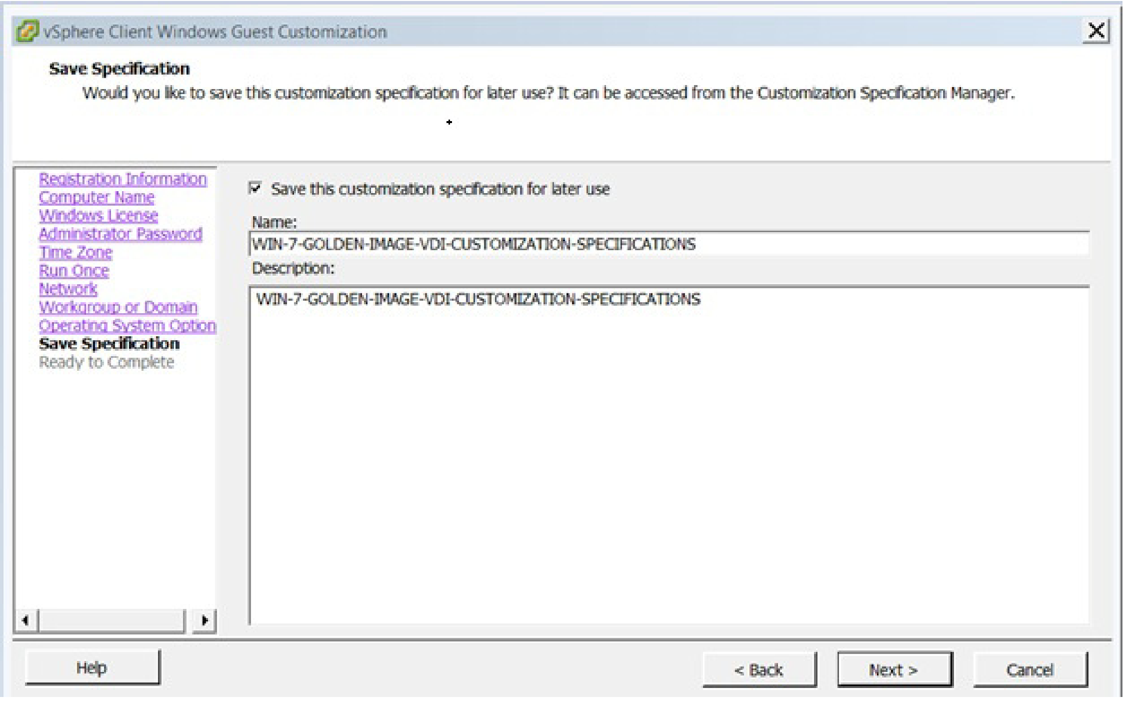

Save Specification

Would you like to save this customization specification for later use? It can be accessed from the Customization Specification Manager.

☑ Save this customization specification for later use

Name:
WIN-7-GOLDEN-IMAGE-VDI-CUSTOMIZATION-SPECIFICATIONS

Description:
WIN-7-GOLDEN-IMAGE-VDI-CUSTOMIZATION-SPECIFICATIONS

14. Verify and click Finish.

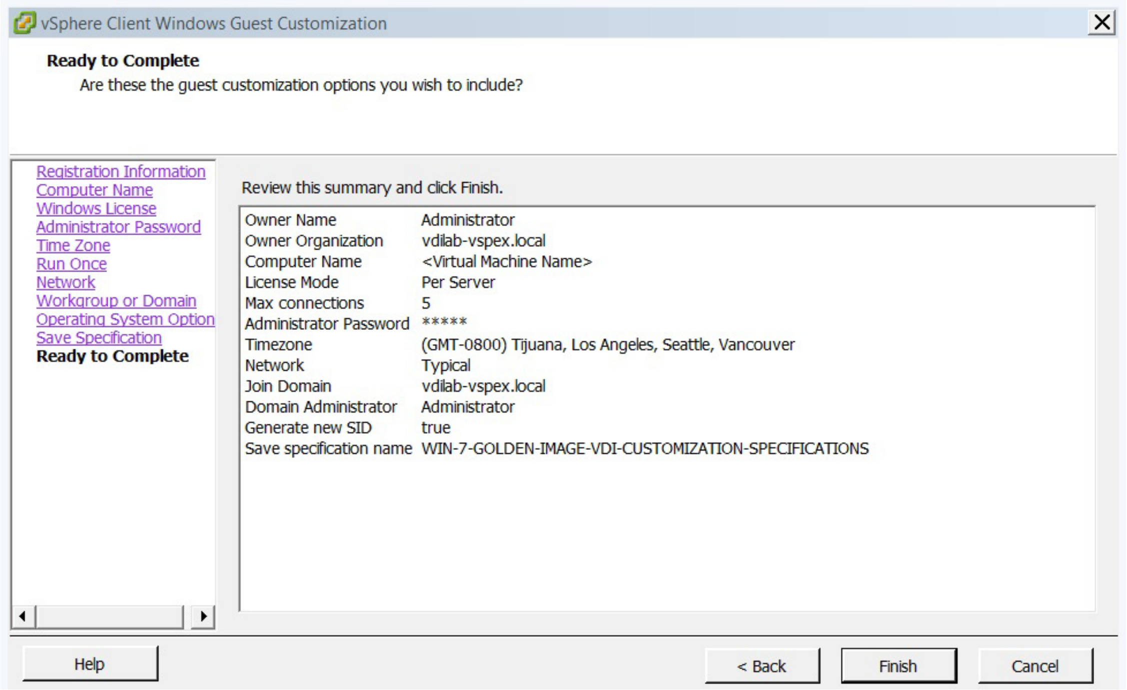

Review this summary and click Finish.

Owner Name	Administrator
Owner Organization	vdilab-vspex.local
Computer Name	<Virtual Machine Name>
License Mode	Per Server
Max connections	5
Administrator Password	*****
Timezone	(GMT-0800) Tijuana, Los Angeles, Seattle, Vancouver
Network	Typical
Join Domain	vdilab-vspex.local
Domain Administrator	Administrator
Generate new SID	true
Save specification name	WIN-7-GOLDEN-IMAGE-VDI-CUSTOMIZATION-SPECIFICATIONS

To edit or modify customization specification log on to vCenter Client with vCenter server IP and credentials. Go to home screen and select customization specification manager. Select saved customization, right-click and select Edit.

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 5.1.2 Virtual Desktop Infrastructure on Cisco Unified Computing System and EMC Storage

There are two variants of the configuration topology for this solution:

- Managed Fibre Channel Variant
- Unmanaged NFS Variant

Diagrams for each variant are shown below.

Figure 10 Architecture Block Diagram- Single Wire Managed Fibre Channel Variant

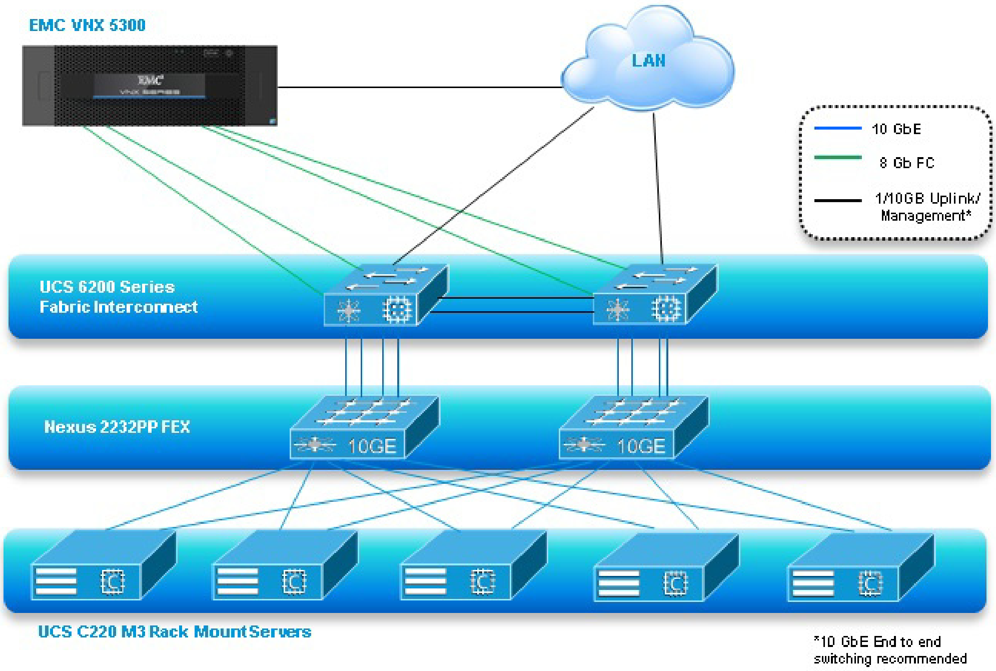
EMC VNX 5300
LAN
10 GbE
8 Gb FC
1/10GB Uplink/ Management*
UCS 6200 Series Fabric Interconnect
Nexus 2232PP FEX
10GE
10GE
UCS C220 M3 Rack Mount Servers
*10 GbE End to end switching recommended

Figure 11 Architecture Block Diagram- Unmanaged NFS Variant

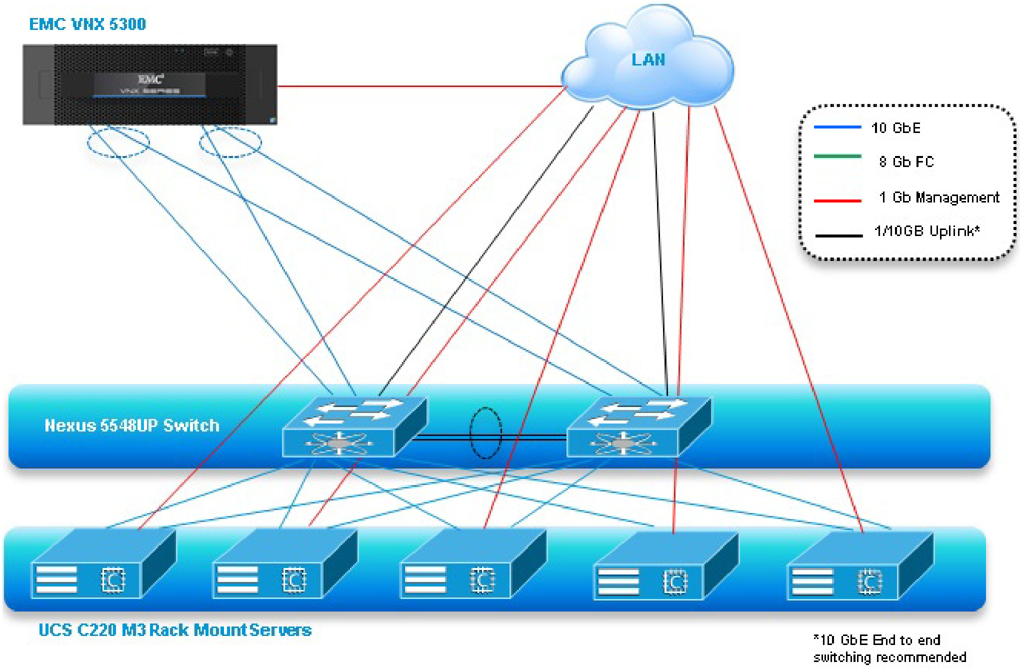

The figures above capture the architectural diagram for the purpose of this study. The architecture is divided into four distinct layers:

- Cisco UCS Compute Platform

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

The following figure details the physical configuration of the 500-600 seat VMware View 5.1.2 environment.

Figure 12 Detailed Architecture of the Configuration for the Single Wire Managed Fibre Channel Variant

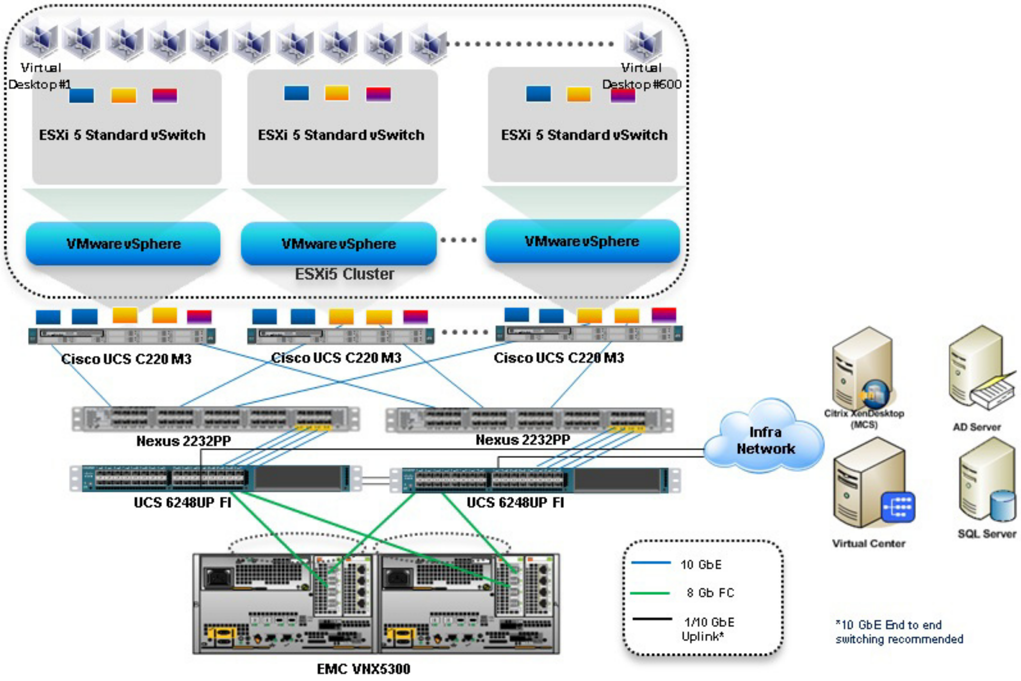

Figure 13 Detailed Architecture of the Configuration for the Unmanaged NFS Variant

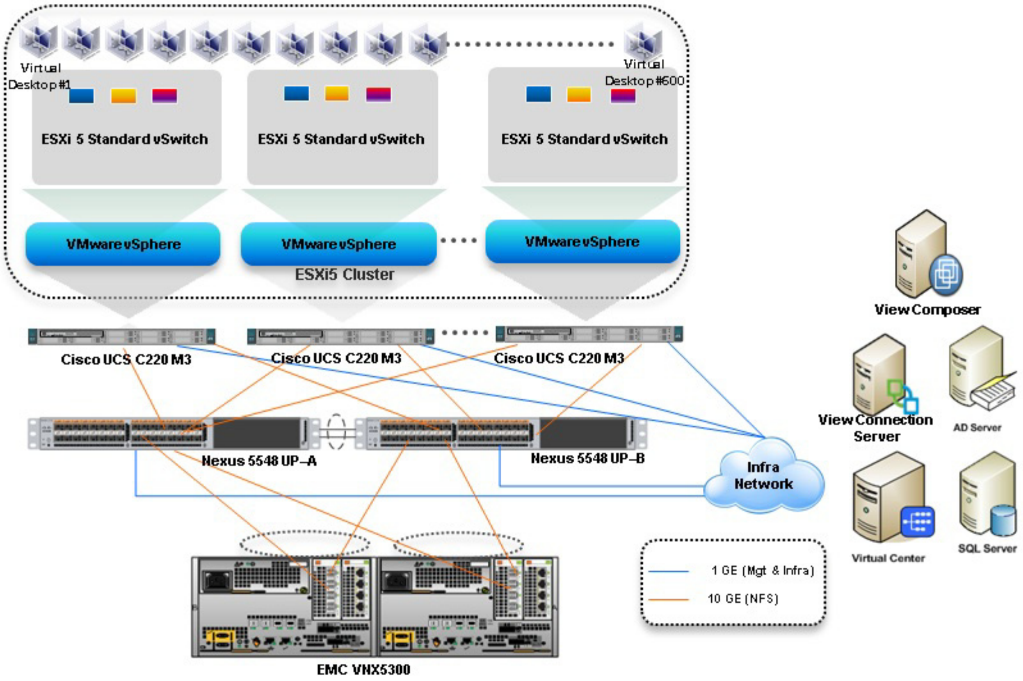

Note The figure above does not include the 1 x 1Gb management connection from each unmanaged Cisco UCS C220 M3 to the Infrastructure Network.

Cisco Unified Computing System Configuration (Managed FC Variant Only)

This section talks about the 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

Note The linke above requires a cisco.com login.

Base Cisco UCS System Configuration

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

1. Bring up the Fabric Interconnect (FI) and from a serial console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. 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. 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.
3. Check the firmware on the system and see if it is current. Go to 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).

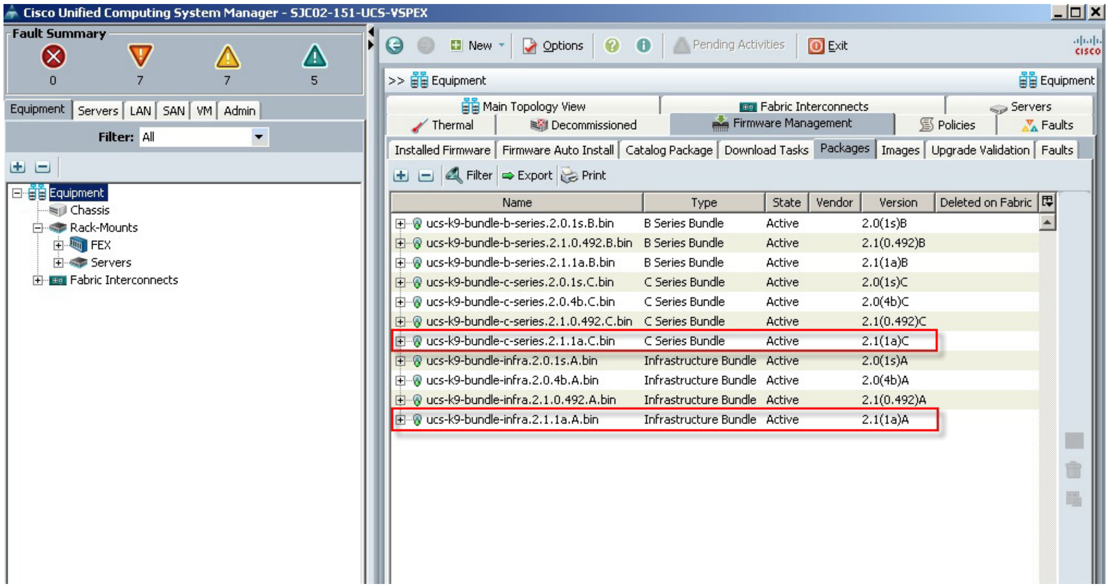

Name	Type	State	Vendor	Version	Deleted on Fabric
ucs-k9-bundle-b-series.2.0.1s.B.bin	B Series Bundle	Active		2.0(1s)B	
ucs-k9-bundle-b-series.2.1.0.492.B.bin	B Series Bundle	Active		2.1(0.492)B	
ucs-k9-bundle-b-series.2.1.1a.B.bin	B Series Bundle	Active		2.1(1a)B	
ucs-k9-bundle-c-series.2.0.1s.C.bin	C Series Bundle	Active		2.0(1s)C	
ucs-k9-bundle-c-series.2.0.4b.C.bin	C Series Bundle	Active		2.0(4b)C	
ucs-k9-bundle-c-series.2.1.0.492.C.bin	C Series Bundle	Active		2.1(0.492)C	
ucs-k9-bundle-c-series.2.1.1a.C.bin	C Series Bundle	Active		2.1(1a)C	
ucs-k9-bundle-infra.2.0.1s.A.bin	Infrastructure Bundle	Active		2.0(1s)A	
ucs-k9-bundle-infra.2.0.4b.A.bin	Infrastructure Bundle	Active		2.0(4b)A	
ucs-k9-bundle-infra.2.1.0.492.A.bin	Infrastructure Bundle	Active		2.1(0.492)A	
ucs-k9-bundle-infra.2.1.1a.A.bin	Infrastructure Bundle	Active		2.1(1a)A	

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 Nexus 2232PP Fabric Extenders to the FIs.

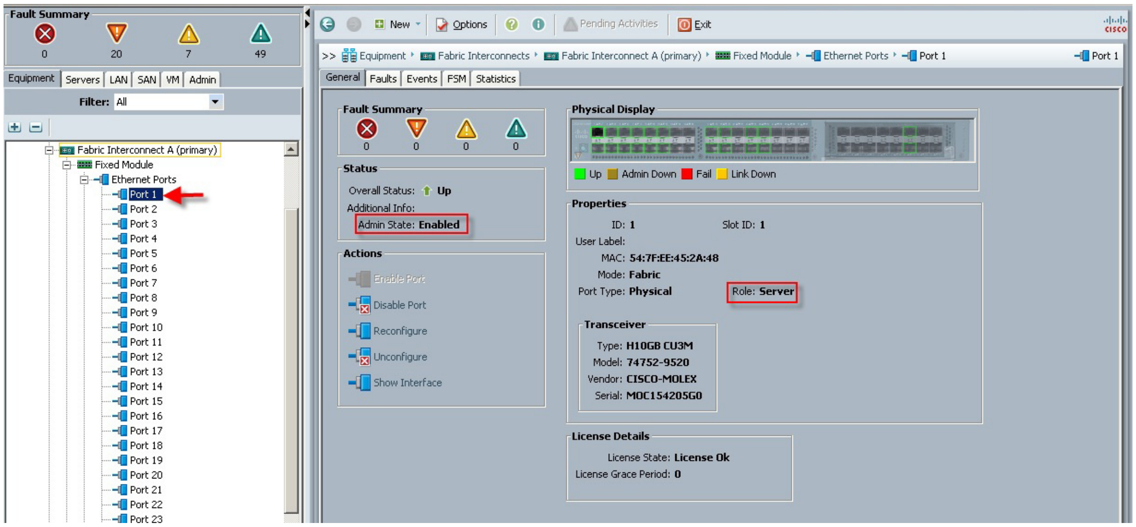

Status

Overall Status: Up

Additional Info:

Admin State: Enabled

Properties

ID: 1 Slot ID: 1

User Label:

MAC: 54:7F:EE:45:2A:48

Mode: Fabric

Port Type: Physical Role: Server

Transceiver

Type: H10GB CU3M

Model: 74752-9520

Vendor: CISCO-MOLEX

Serial: MOC154205G0

License Details

License State: License Ok

License Grace Period: 0

5. Configure and enable at least one uplink Ethernet ports to connect the Cisco UCS system to your LAN.

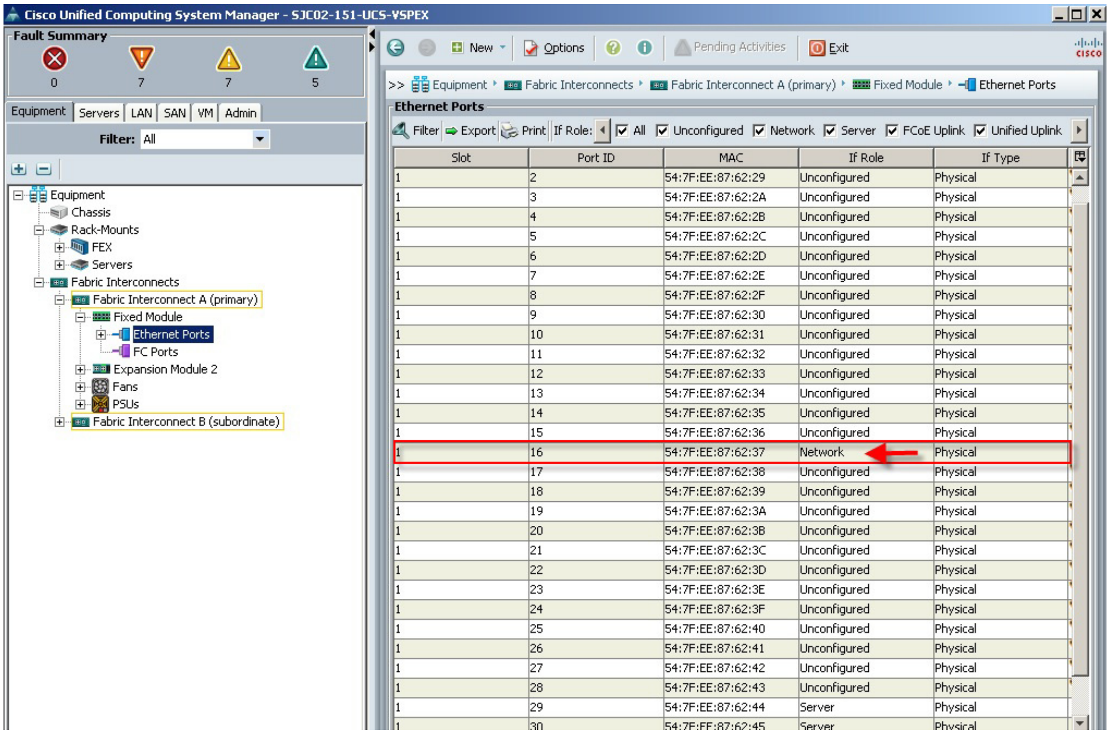
Cisco Unified Computing System Manager - SJC02-151-UCS-VSPEX
Fault Summary
0
7
7
5
Equipment
Servers
LAN
SAN
VM
Admin
Filter: All
Equipment
Chassis
Rack-Mounts
FEX
Servers
Fabric Interconnects
Fabric Interconnect A (primary)
Fixed Module
Ethernet Ports
FC Ports
Expansion Module 2
Fans
PSUs
Fabric Interconnect B (subordinate)
New
Options
Pending Activities
Exit
>> Equipment ▸ Fabric Interconnects ▸ Fabric Interconnect A (primary) ▸ Fixed Module ▸ Ethernet Ports
Filter
Export
Print

Ethernet Ports

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

Slot	Port ID	MAC	If Role	If Type
1	2	54:7F:EE:87:62:29	Unconfigured	Physical
1	3	54:7F:EE:87:62:2A	Unconfigured	Physical
1	4	54:7F:EE:87:62:2B	Unconfigured	Physical
1	5	54:7F:EE:87:62:2C	Unconfigured	Physical
1	6	54:7F:EE:87:62:2D	Unconfigured	Physical
1	7	54:7F:EE:87:62:2E	Unconfigured	Physical
1	8	54:7F:EE:87:62:2F	Unconfigured	Physical
1	9	54:7F:EE:87:62:30	Unconfigured	Physical
1	10	54:7F:EE:87:62:31	Unconfigured	Physical
1	11	54:7F:EE:87:62:32	Unconfigured	Physical
1	12	54:7F:EE:87:62:33	Unconfigured	Physical
1	13	54:7F:EE:87:62:34	Unconfigured	Physical
1	14	54:7F:EE:87:62:35	Unconfigured	Physical
1	15	54:7F:EE:87:62:36	Unconfigured	Physical
1	16	54:7F:EE:87:62:37	Network	Physical
1	17	54:7F:EE:87:62:38	Unconfigured	Physical
1	18	54:7F:EE:87:62:39	Unconfigured	Physical
1	19	54:7F:EE:87:62:3A	Unconfigured	Physical
1	20	54:7F:EE:87:62:3B	Unconfigured	Physical
1	21	54:7F:EE:87:62:3C	Unconfigured	Physical
1	22	54:7F:EE:87:62:3D	Unconfigured	Physical
1	23	54:7F:EE:87:62:3E	Unconfigured	Physical
1	24	54:7F:EE:87:62:3F	Unconfigured	Physical
1	25	54:7F:EE:87:62:40	Unconfigured	Physical
1	26	54:7F:EE:87:62:41	Unconfigured	Physical
1	27	54:7F:EE:87:62:42	Unconfigured	Physical
1	28	54:7F:EE:87:62:43	Unconfigured	Physical
1	29	54:7F:EE:87:62:44	Server	Physical
1	30	54:7F:EE:87:62:45	Server	Physical

6. From the Equipment tab with one of the two Fabric Interconnects highlighted in the navigation pane, use the Configure Unified Ports in the Actions panel on the General tab, to configure FC Storage ports.

Note In this example, we configured four FC Storage ports, two of which are in use. Ports to the left of the slider shown above are Ethernet ports. Ports to the right of the slider are Fibre Channel ports.

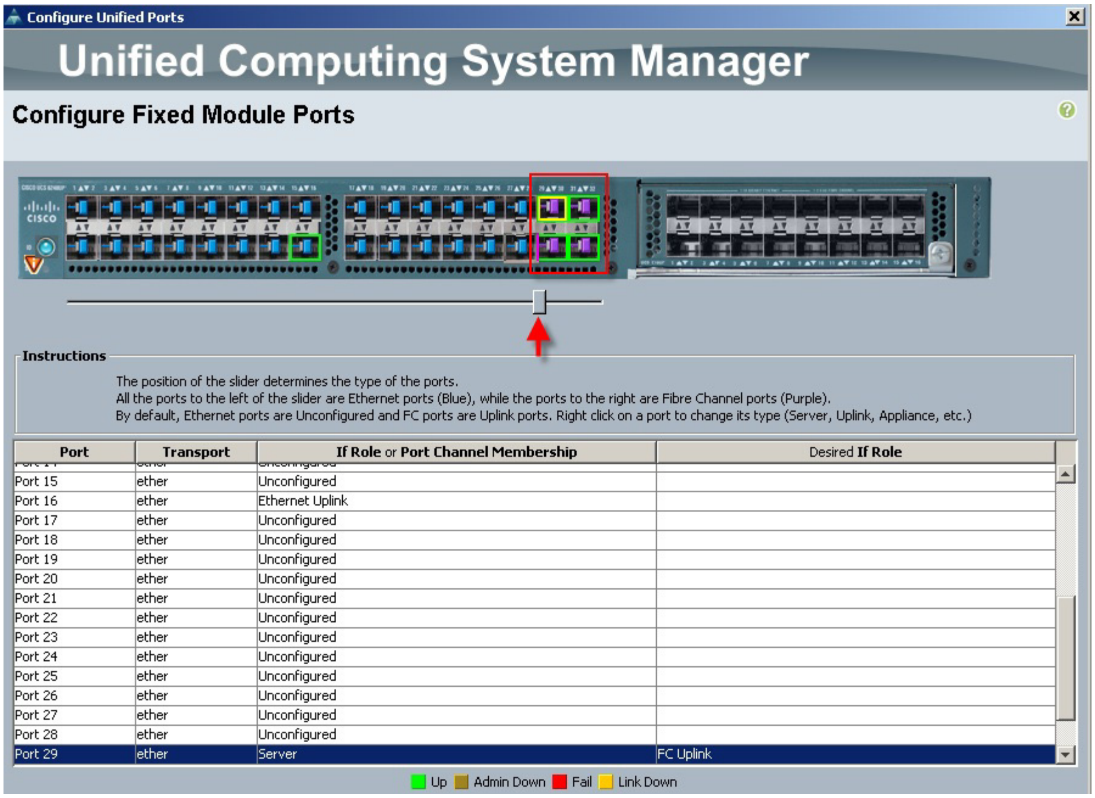

Configure Unified Ports

Unified Computing System Manager

Configure Fixed Module Ports

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Ethernet ports (Blue), while the ports to the right are Fibre Channel ports (Purple).
By default, Ethernet ports are Unconfigured and FC ports are Uplink ports. Right click on a port to change its type (Server, Uplink, Appliance, etc.)

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Ethernet Uplink	
Port 17	ether	Unconfigured	
Port 18	ether	Unconfigured	
Port 19	ether	Unconfigured	
Port 20	ether	Unconfigured	
Port 21	ether	Unconfigured	
Port 22	ether	Unconfigured	
Port 23	ether	Unconfigured	
Port 24	ether	Unconfigured	
Port 25	ether	Unconfigured	
Port 26	ether	Unconfigured	
Port 27	ether	Unconfigured	
Port 28	ether	Unconfigured	
Port 29	ether	Server	FC Uplink

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down

7. Connect four uplink ports on each of the Nexus 2232PP Fabric Extenders to four server ports on one of the Cisco UCS Fabric Interconnects. (All four to one Fabric Interconnect.) Power on the Nexus 2232PPs.

8. Expand the Rack-Mounts, FEX node in the left pane of the Cisco UCS Manager console, the click each FEX in the left pane, then click Acknowledge FEX in the right pane to bring the FEX online and enable server discovery.

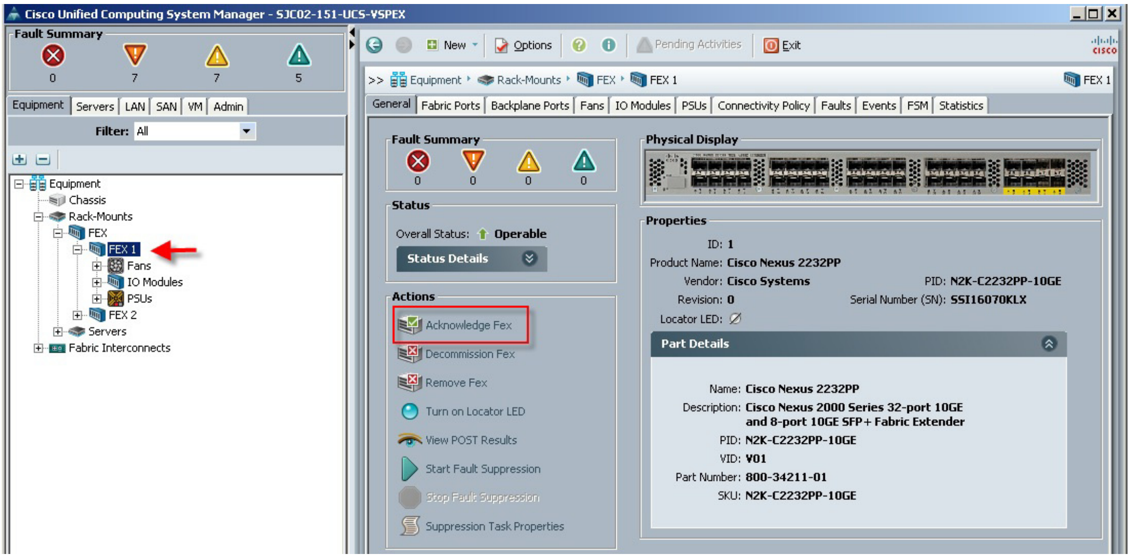

Overall Status: Operable

Status Details

ID: 1
Product Name: Cisco Nexus 2232PP
Vendor: Cisco Systems
PID: N2K-C2232PP-10GE
Revision: 0
Serial Number (SN): SSI16070KLX
Locator LED:

Name: Cisco Nexus 2232PP
Description: Cisco Nexus 2000 Series 32-port 10GE and 8-port 10GE SFP+ Fabric Extender
PID: N2K-C2232PP-10GE
VID: V01
Part Number: 800-34211-01
SKU: N2K-C2232PP-10GE

9. Repeat the procedure for FEX2.

10. Download the UCS C220 M3 Rack Server Software version 1.4.(7a) or later:

    http://software.cisco.com/download/release.html?mdfid=284296253&flowid=31742&softwareid=283850974&release=1.4(7a)1&relind=null&rellifecycle=null&reltype=null&i=rb

11. Burn the software to a CD.

12. Insert the CD into each Cisco UCS C220 M3, restart the server, press F6 during the post process to view the boot menu, select the CD Rom for the boot device and update all components via the menu.

13. After the server firmware and CIMC have been updated, during the restart, press F8 to access the server CIMC.

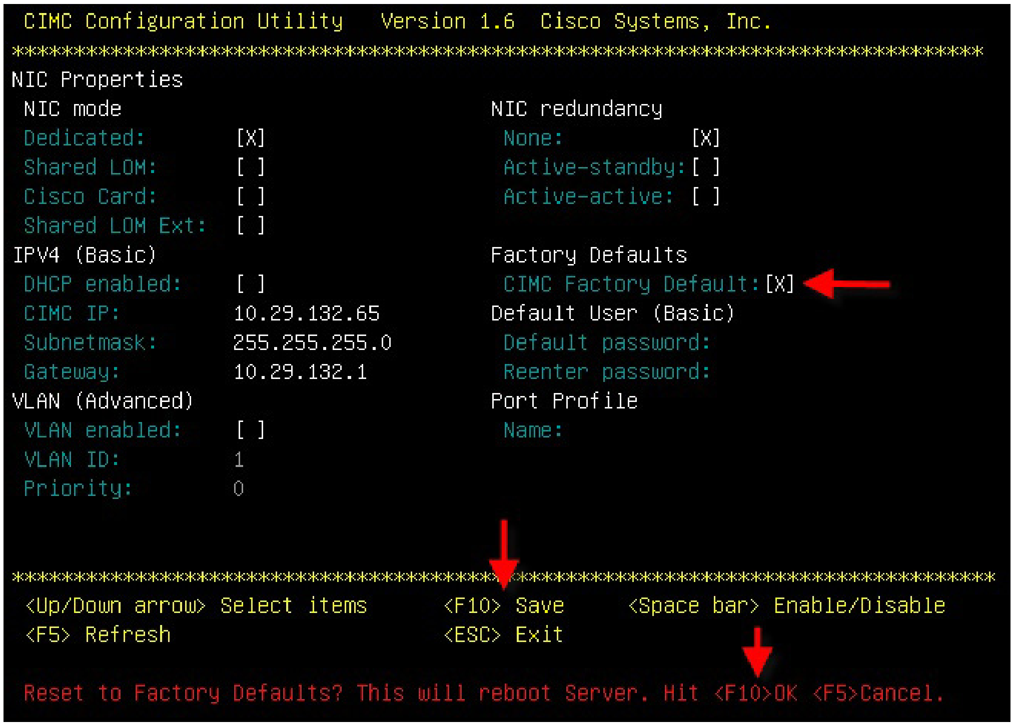

```
 CIMC Configuration Utility   Version 1.6  Cisco Systems, Inc.
********************************************************************************
NIC Properties
 NIC mode                             NIC redundancy
 Dedicated:        [X]                 None:          [X]
 Shared LOM:       [ ]                 Active-standby:[ ]
 Cisco Card:       [ ]                 Active-active: [ ]
 Shared LOM Ext:   [ ]
IPV4 (Basic)                          Factory Defaults
 DHCP enabled:     [ ]                 CIMC Factory Default:[X]
 CIMC IP:          10.29.132.65       Default User (Basic)
 Subnetmask:       255.255.255.0       Default password:
 Gateway:          10.29.132.1         Reenter password:
VLAN (Advanced)                       Port Profile
 VLAN enabled:     [ ]                 Name:
 VLAN ID:          1
 Priority:         0


*********************************************************************************
 <Up/Down arrow> Select items      <F10> Save     <Space bar> Enable/Disable
 <F5> Refresh                      <ESC> Exit

 Reset to Factory Defaults? This will reboot Server. Hit <F10>OK <F5>Cancel.
```

14. Select CIMC Factory Default with spacebar, press F10, then F10 again to confirm factory defaults.

15. This step prepares the server for management through the Nexus 2232PP FEX by Cisco UCS Manager on the Fabric Interconnects.

16. Connect one port of the VIC1225 from each Cisco UCS C220 M3 server to one Nexus 2232PP Fabric Extender. Connect the second VIC1225 port to the other Nexus 2232PP Fabric Extender. Repeat for all servers.

17. Reboot all of the Cisco UCS C220 M3 servers.

18. From the Equipment tab in Cisco UCS Manager, navigate to Rack-Mounts, Servers. The five Cisco UCS C220 M3 servers should be visible.

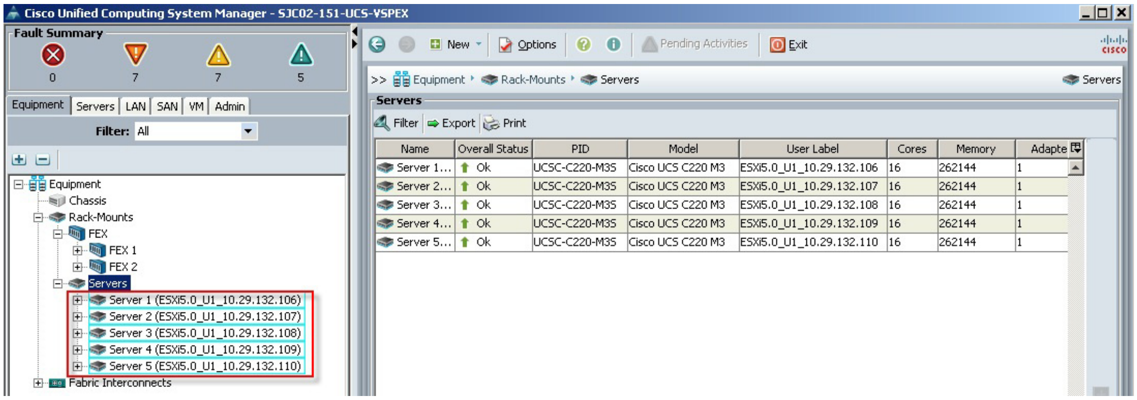

- Equipment
  - Chassis
  - Rack-Mounts
    - FEX
      - FEX 1
      - FEX 2
    - Servers
      - Server 1 (ESXi5.0_U1_10.29.132.106)
      - Server 2 (ESXi5.0_U1_10.29.132.107)
      - Server 3 (ESXi5.0_U1_10.29.132.108)
      - Server 4 (ESXi5.0_U1_10.29.132.109)
      - Server 5 (ESXi5.0_U1_10.29.132.110)
  - Fabric Interconnects

Name	Overall Status	PID	Model	User Label	Cores	Memory	Adapte
Server 1...	Ok	UCSC-C220-M3S	Cisco UCS C220 M3	ESXi5.0_U1_10.29.132.106	16	262144	1
Server 2...	Ok	UCSC-C220-M3S	Cisco UCS C220 M3	ESXi5.0_U1_10.29.132.107	16	262144	1
Server 3...	Ok	UCSC-C220-M3S	Cisco UCS C220 M3	ESXi5.0_U1_10.29.132.108	16	262144	1
Server 4...	Ok	UCSC-C220-M3S	Cisco UCS C220 M3	ESXi5.0_U1_10.29.132.109	16	262144	1
Server 5...	Ok	UCSC-C220-M3S	Cisco UCS C220 M3	ESXi5.0_U1_10.29.132.110	16	262144	1

Note If the servers do not appear under the server node, repeat step 6, Acknowledge FEX for each FEX.

19. 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.

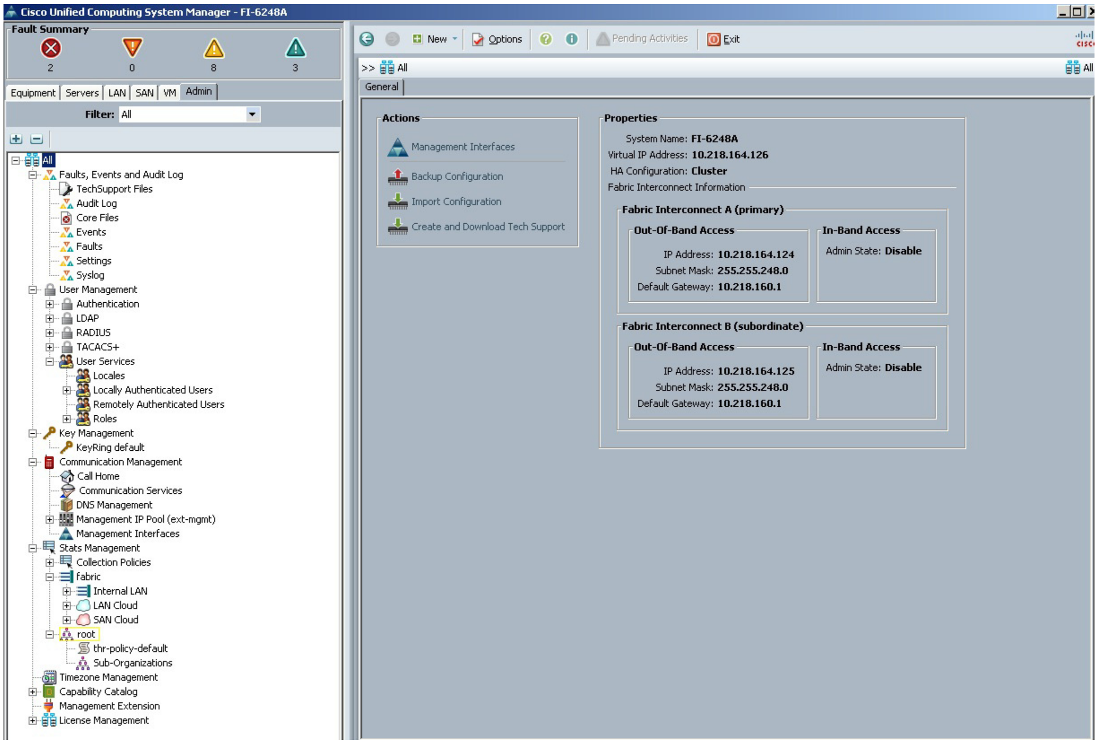

Properties

System Name: FI-6248A
Virtual IP Address: 10.218.164.126
HA Configuration: Cluster
Fabric Interconnect Information

Fabric Interconnect A (primary)

Out-Of-Band Access
IP Address: 10.218.164.124
Subnet Mask: 255.255.248.0
Default Gateway: 10.218.160.1

In-Band Access
Admin State: Disable

Fabric Interconnect B (subordinate)

Out-Of-Band Access
IP Address: 10.218.164.125
Subnet Mask: 255.255.248.0
Default Gateway: 10.218.160.1

In-Band Access
Admin State: Disable

20. Create all the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, Server pool.

21. From the LAN tab in the navigator, under the Pools node, we created a MAC address pool of sufficient size for the environment. In this project, we created a single pool with two address ranges for expandability.

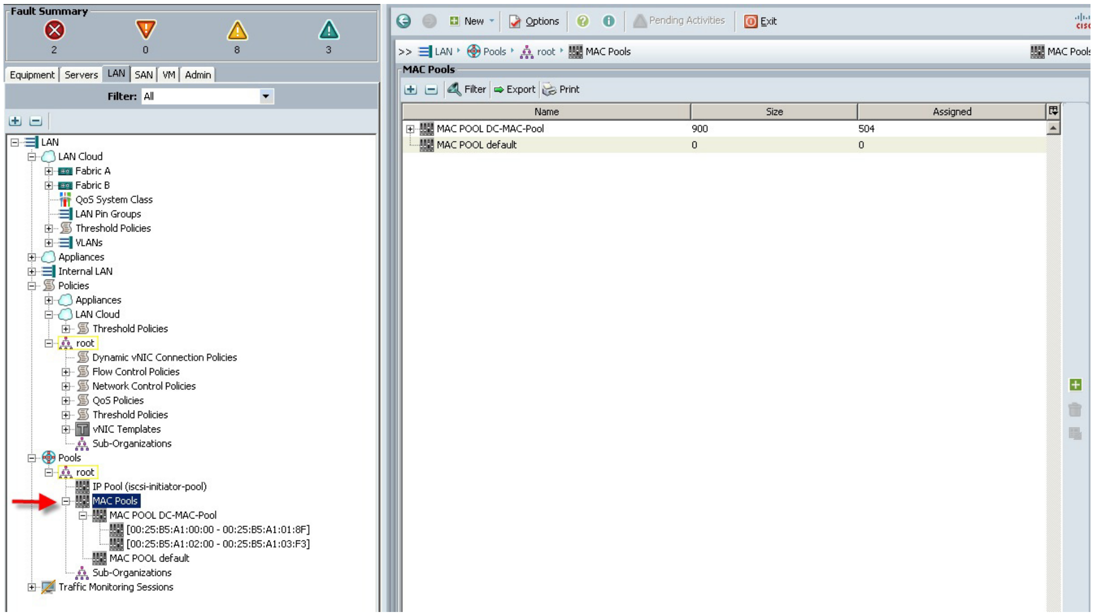

Name	Size	Assigned
MAC POOL DC-MAC-Pool	900	504
MAC POOL default	0	0

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

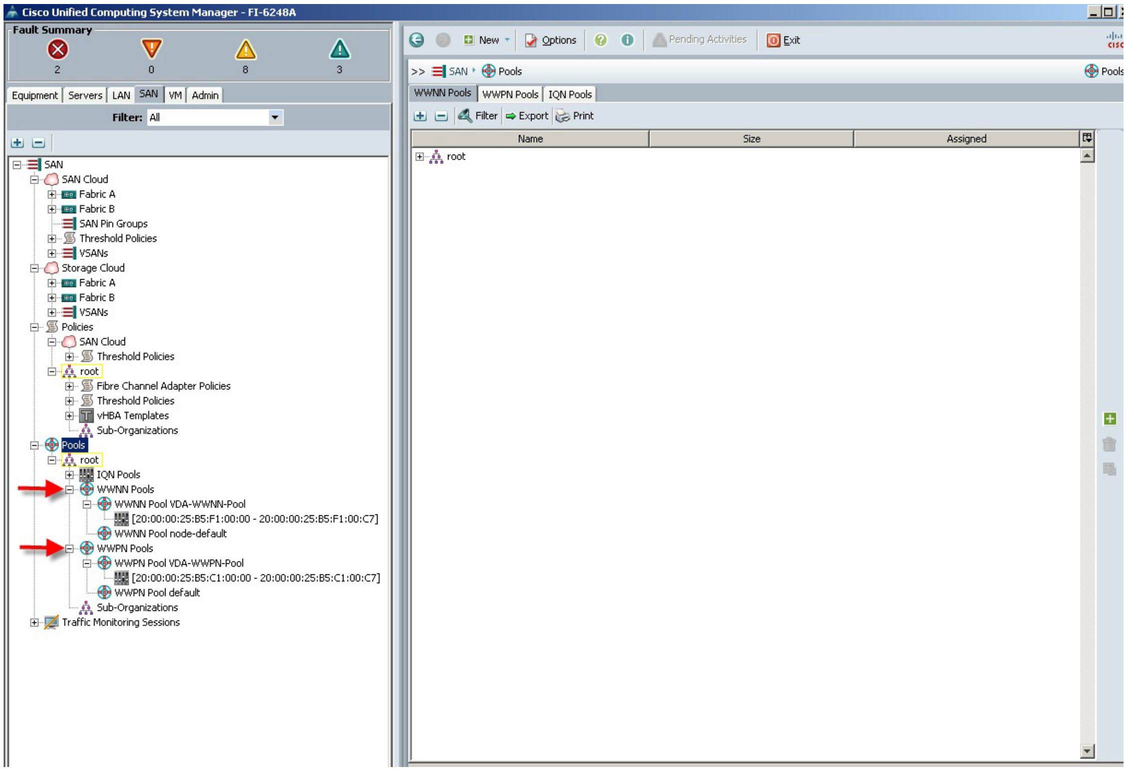

Name	Size	Assigned
root		

23. For this project, we used a single VSAN, the default VSAN with ID 1.

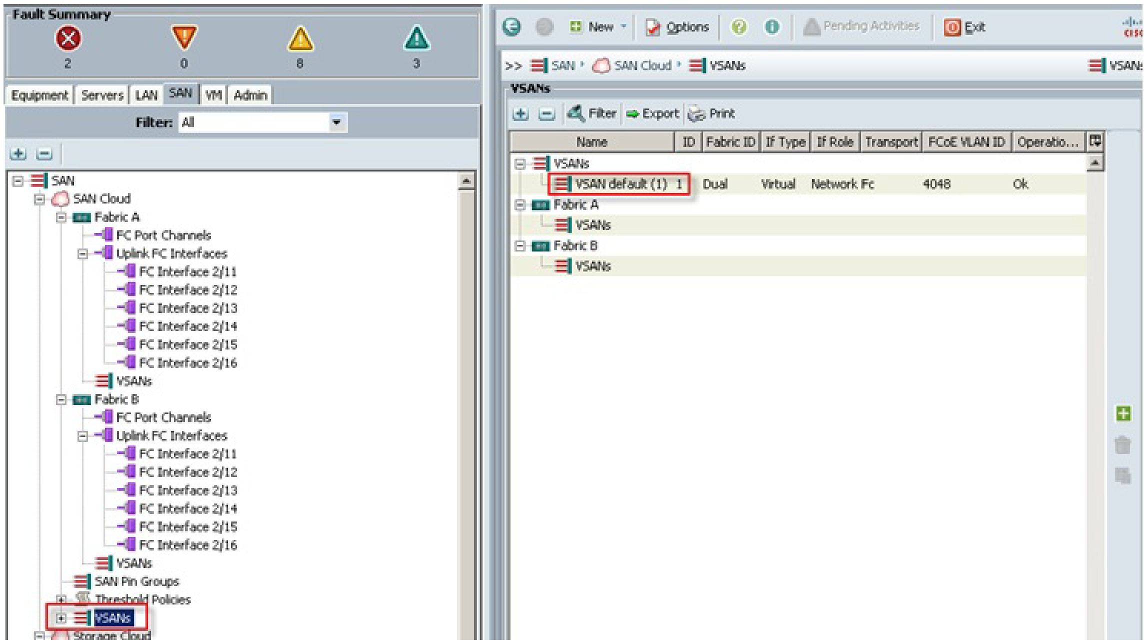

VSANs

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operatio...
VSANs							
VSAN default (1)	1	Dual		Virtual	Network Fc	4048	Ok
Fabric A							
VSANs							
Fabric B							
VSANs							

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

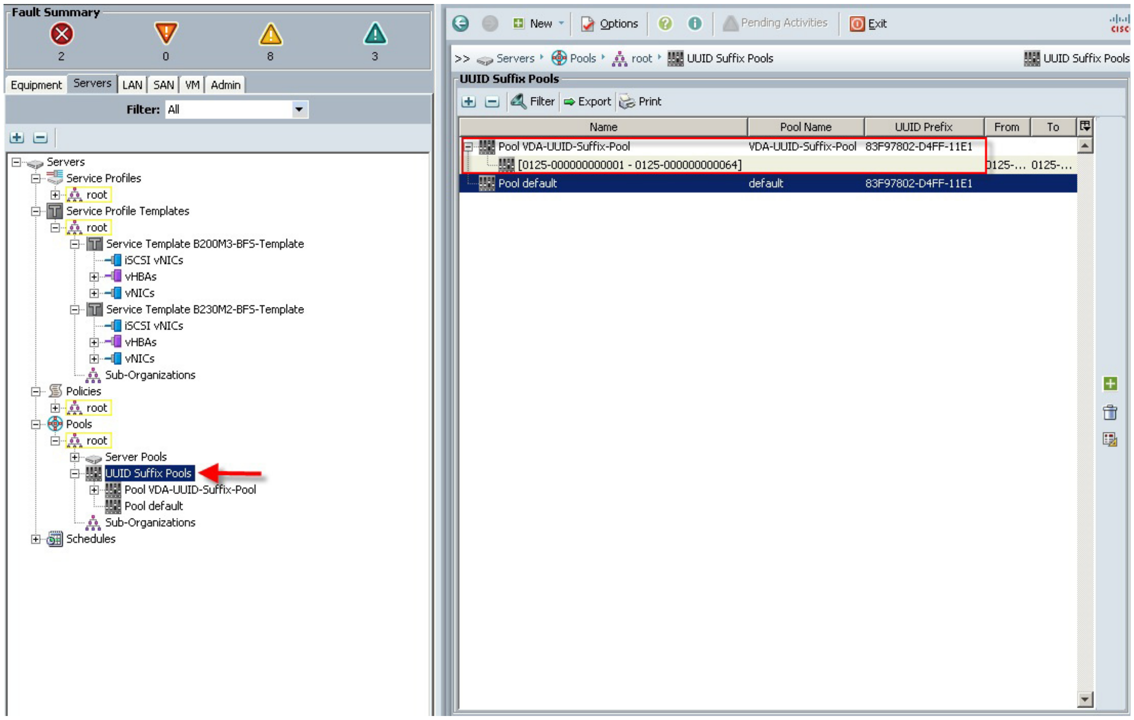

UUID Suffix Pools

Name	Pool Name	UUID Prefix	From	To
Pool VDA-UUID-Suffix-Pool	VDA-UUID-Suffix-Pool	83F97802-D4FF-11E1		
[0125-000000000001 - 0125-000000000064]			0125-...	0125-...
Pool default	default	83F97802-D4FF-11E1		

Note We created one Server Pool for use in our Service Profile Templates as the selection criteria for automated profile association. The Server Pool was created in the Servers tab in the navigation page under the Pools node. Only the pool name was created, no servers were added.

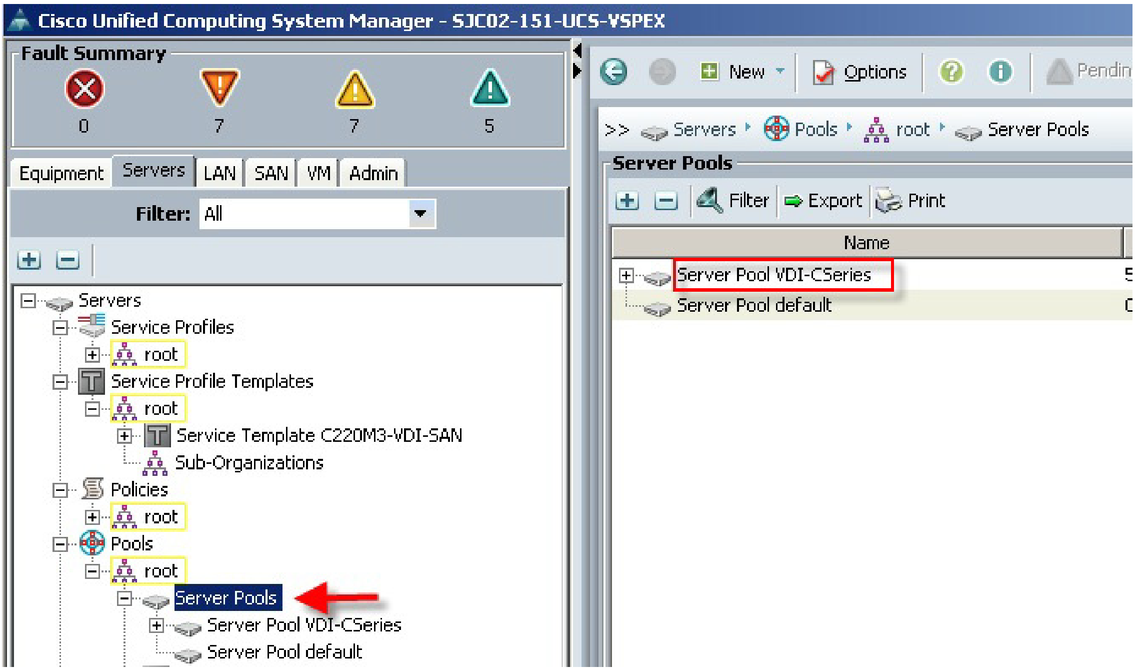

Note We created one Server Pool Policy Qualification to identify the Cisco UCS C220 M3 rack-mount server model for placement into the correct Server pool using the Service Profile Template. In this case we used the Cisco UCS C220 M3 Product ID (PID) to select the servers. (This is helpful if the deployment grew and different Cisco UCS C-Series Rack-Mount models were incorporated later.)

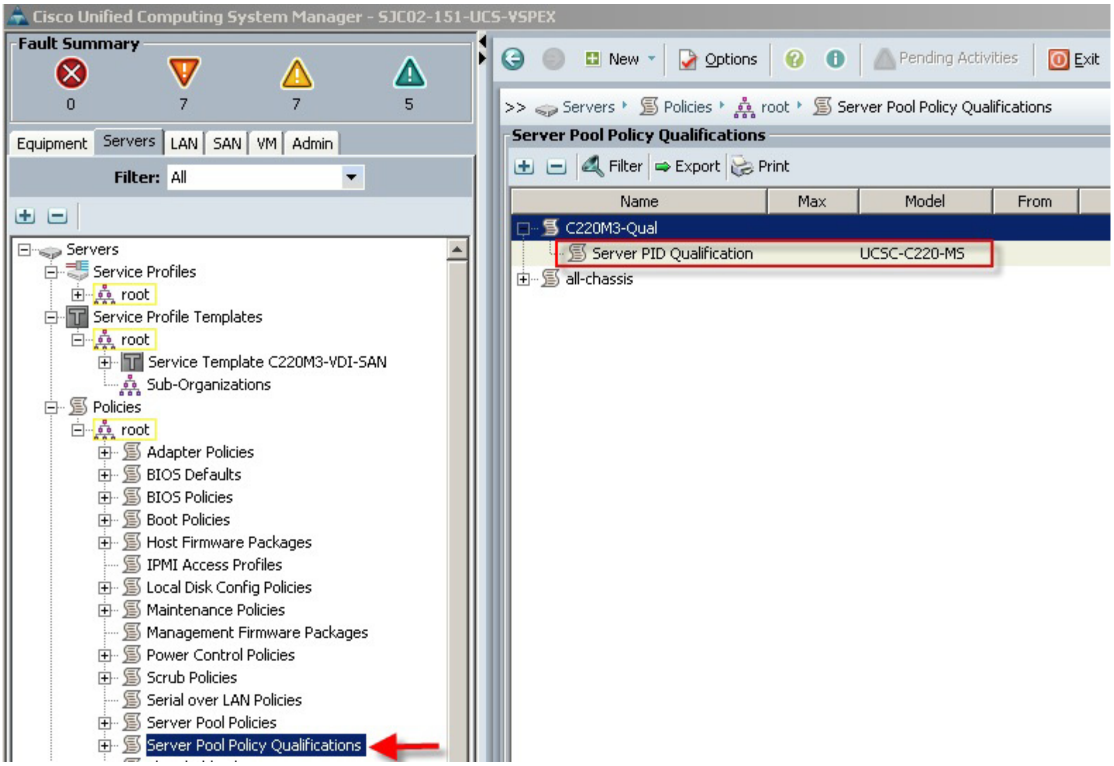

Server Pool Policy Qualifications

Name	Max	Model	From
C220M3-Qual			
Server PID Qualification		UCSC-C220-MS	
all-chassis			

25. The next step in automating the server selection process is to create corresponding Server Pool Policy for each Cisco UCS C-Series server model, utilizing the Server Pool and Server Pool Policy Qualification created earlier.

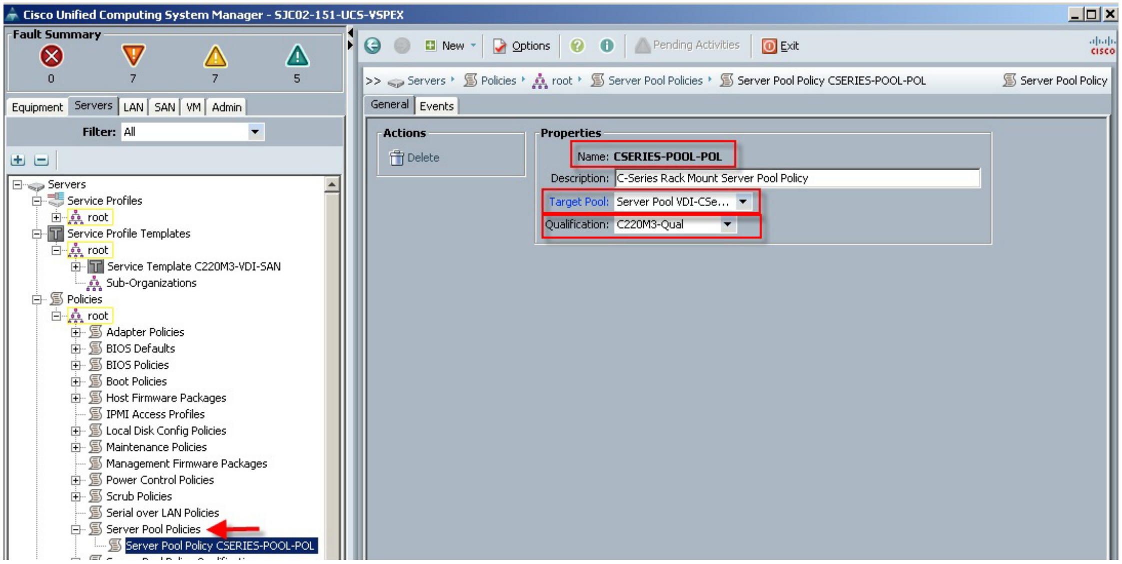

Name: CSERIES-POOL-POL

Description: C-Series Rack Mount Server Pool Policy

Target Pool: Server Pool VDI-CSe...

Qualification: C220M3-Qual

26. The Virtual Host Bus Adapter updating templates were created for FC SAN connectivity from the SAN tab under the Polices node, utilizing the WWPN pool created earlier and the default FC QoS policy, one template for each fabric.

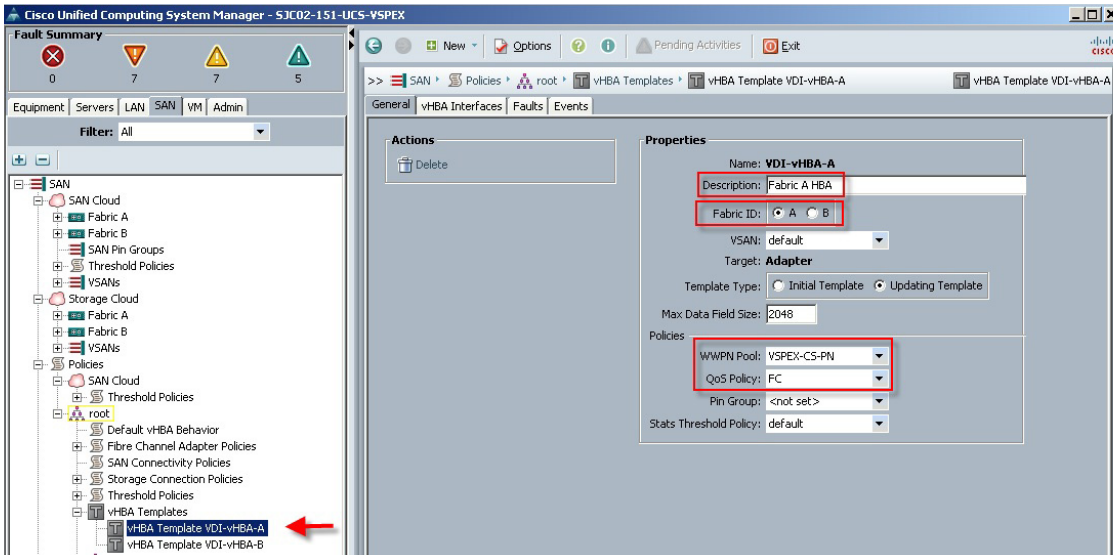

Name: VDI-vHBA-A

Description: Fabric A HBA

Fabric ID: ● A ○ B

VSAN: default

Target: Adapter

Template Type: ○ Initial Template ● Updating Template

Max Data Field Size: 2048

Policies

WWPN Pool: VSPEX-CS-PN

QoS Policy: FC

Pin Group: <not set>

Stats Threshold Policy: default

27. On the LAN tab in the navigator pane, configure the VLANs for the environment.

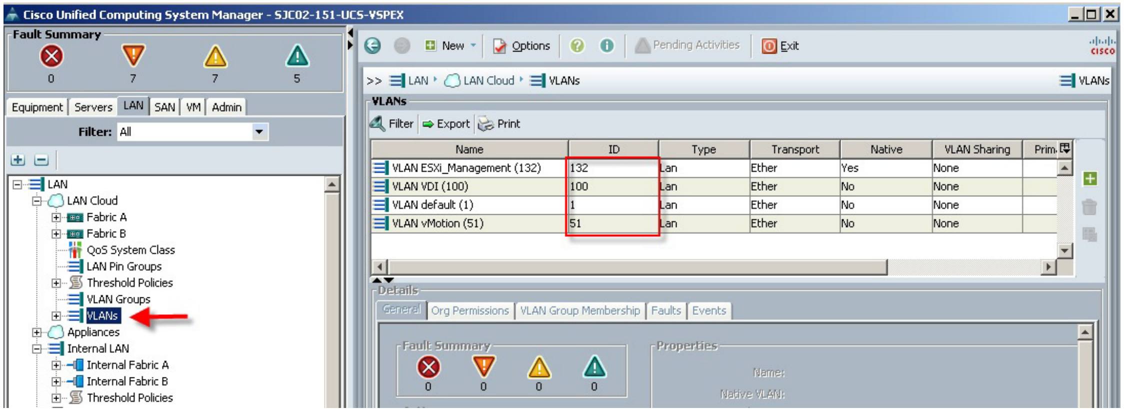

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN ESXi_Management (132)	132	Lan	Ether	Yes	None
VLAN VDI (100)	100	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	No	None
VLAN vMotion (51)	51	Lan	Ether	No	None

Note In this project we utilized three VLANs for the Managed FC variant to accommodate our three ethernet system classes. Infrastructure services shared VLAN 132.

28. On 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 six virtual NICs per host, three pairs, with one member of each pair connected to one of the two Fabric Interconnects for resiliency.

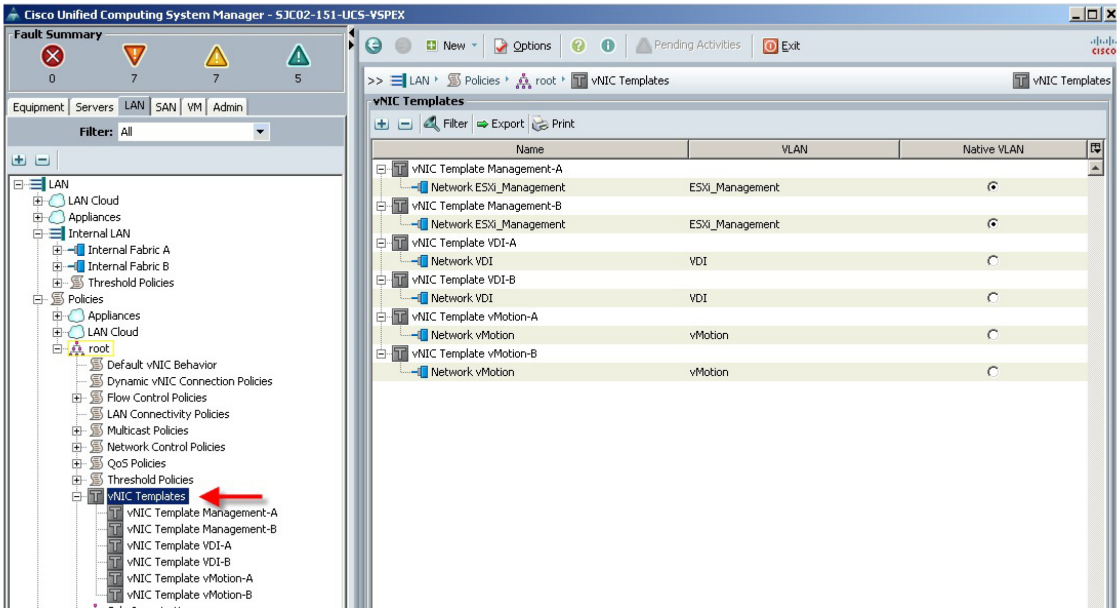
Cisco Unified Computing System Manager - SJC02-151-UCS-VSPEX

>> LAN ▸ Policies ▸ root ▸ vNIC Templates

vNIC Templates

Name	VLAN	Native VLAN
vNIC Template Management-A		
Network ESXi_Management	ESXi_Management	●
vNIC Template Management-B		
Network ESXi_Management	ESXi_Management	●
vNIC Template VDI-A		
Network VDI	VDI	○
vNIC Template VDI-B		
Network VDI	VDI	○
vNIC Template vMotion-A		
Network vMotion	vMotion	○
vNIC Template vMotion-B		
Network vMotion	vMotion	○

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

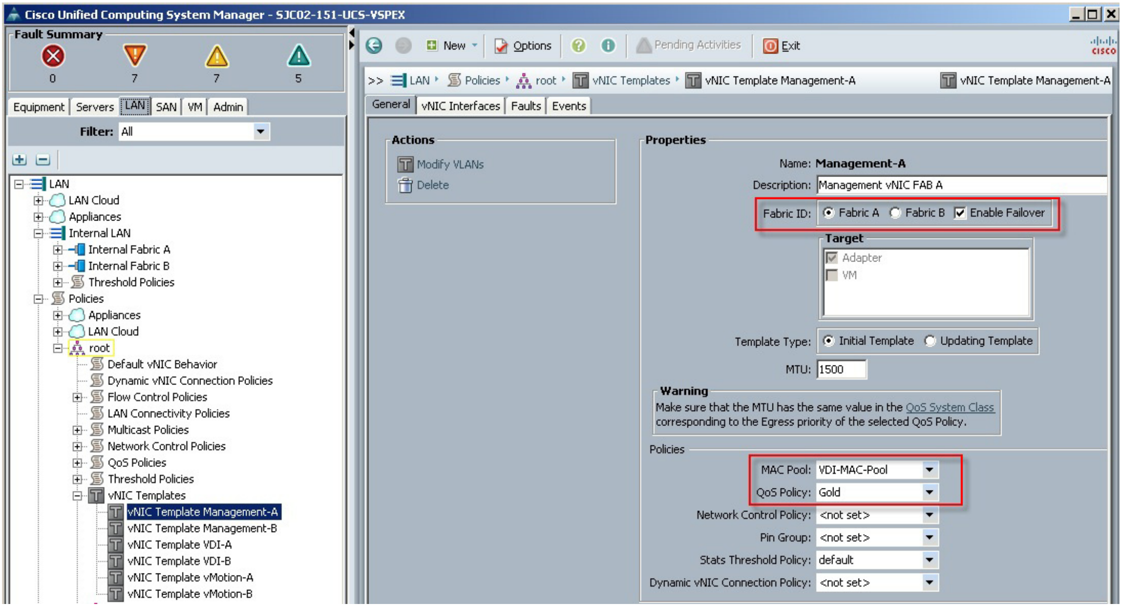
Cisco Unified Computing System Manager - SJC02-151-UCS-VSPEX

>> LAN ▸ Policies ▸ root ▸ vNIC Templates ▸ vNIC Template Management-A

Properties

Name: Management-A

Description: Management vNIC FAB A

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

Target: ☑ Adapter ☐ VM

Template Type: ● Initial Template ○ Updating Template

MTU: 1500

Warning
Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

Policies

MAC Pool: VDI-MAC-Pool

QoS Policy: Gold

Network Control Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

Dynamic vNIC Connection Policy: <not set>

30. Create a performance BIOS Policy for the Cisco UCS C220 M3 server to insure optimal performance. The following screen captures show the settings for the Cisco UCS C220 M3 servers used in this study.

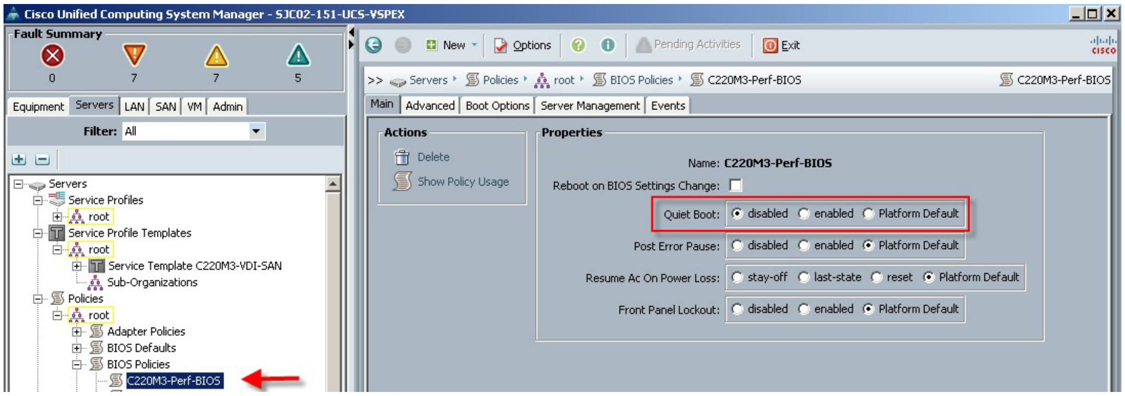

Name: **C220M3-Perf-BIOS**

Reboot on BIOS Settings Change: □

Quiet Boot: ● disabled ○ enabled ○ Platform Default

Post Error Pause: ○ disabled ○ enabled ● Platform Default

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

Front Panel Lockout: ○ disabled ○ enabled ● Platform Default

Advanced tab, Processor settings.

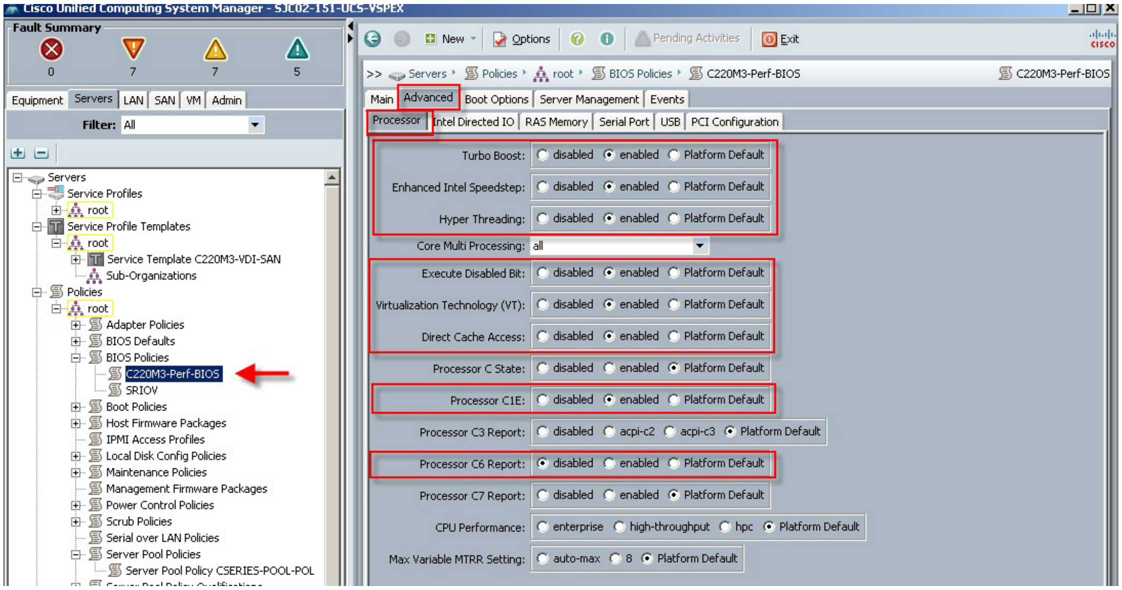

Turbo Boost: ○ disabled ● enabled ○ Platform Default

Enhanced Intel Speedstep: ○ disabled ● enabled ○ Platform Default

Hyper Threading: ○ disabled ● enabled ○ Platform Default

Core Multi Processing: all

Execute Disabled Bit: ○ disabled ● enabled ○ Platform Default

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

Direct Cache Access: ○ disabled ● enabled ○ Platform Default

Processor C State: ○ disabled ○ enabled ● Platform Default

Processor C1E: ○ disabled ● enabled ○ Platform Default

Processor C3 Report: ○ disabled ○ acpi-c2 ○ acpi-c3 ● Platform Default

Processor C6 Report: ● disabled ○ enabled ○ Platform Default

Processor C7 Report: ○ disabled ○ enabled ● Platform Default

CPU Performance: ○ enterprise ○ high-throughput ○ hpc ● Platform Default

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

Advanced Tab, Intel Directed IO settings.

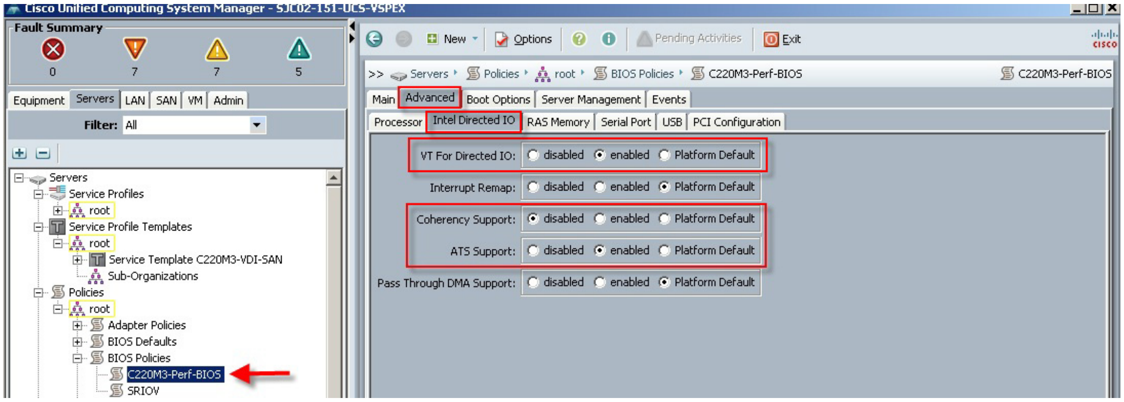

VT For Directed IO: ○ disabled ● enabled ○ Platform Default

Interrupt Remap: ○ disabled ○ enabled ● Platform Default

Coherency Support: ● disabled ○ enabled ○ Platform Default

ATS Support: ○ disabled ● enabled ○ Platform Default

Pass Through DMA Support: ○ disabled ○ enabled ● Platform Default

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

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

31. To enable Boot from SAN on the Cisco UCS Manager 2.0 (UCS-M) series, create a Boot from SAN policy.

32. Add SAN Boot for primary to the new policy. The vHBA name is optional, it could be left blank and you do not have to enforce the vHBA name. Click OK.

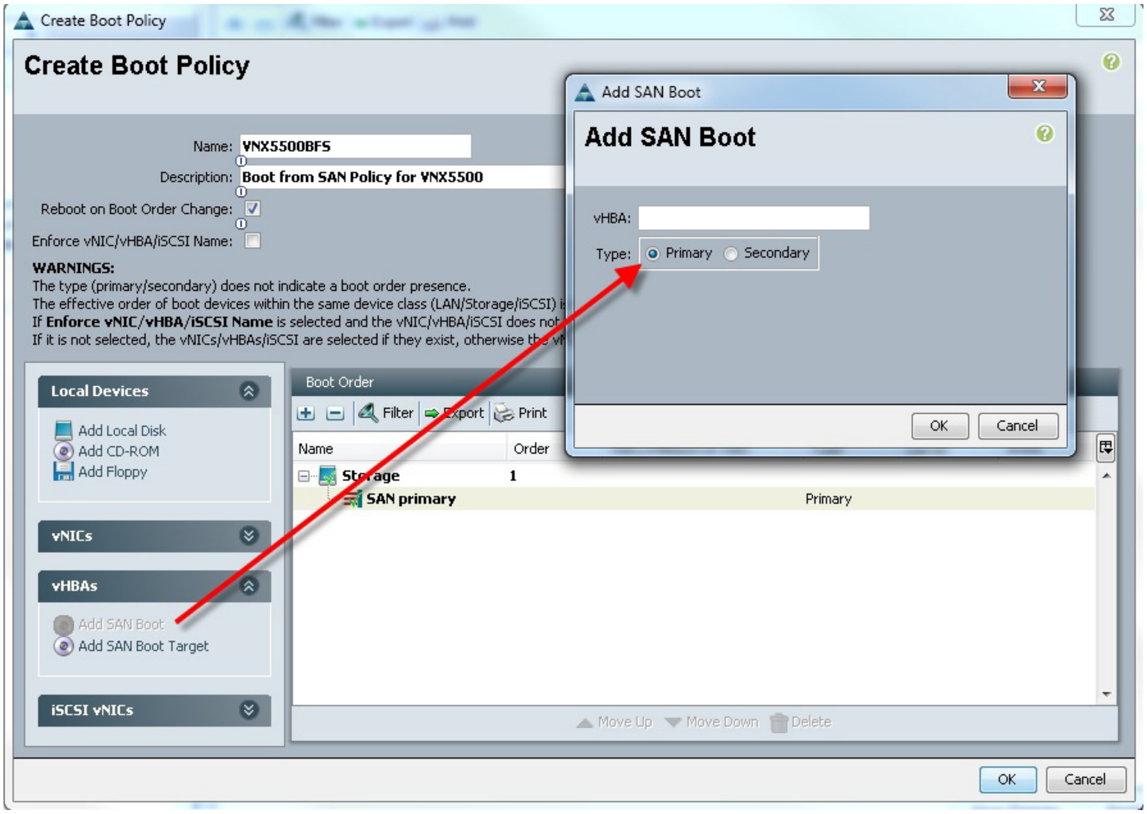

Create Boot Policy

Name: VNX5500BFS

Description: Boot from SAN Policy for VNX5500

Reboot on Boot Order Change: ☑

Enforce vNIC/vHBA/iSCSI Name: □

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) i
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the v

Add SAN Boot

vHBA:

Type: ● Primary ○ Secondary

Name	Order	
Storage	1	
SAN primary		Primary

33. Add SAN boot for SAN Secondary, Click OK. Again, we left the optional vHBA name blank.

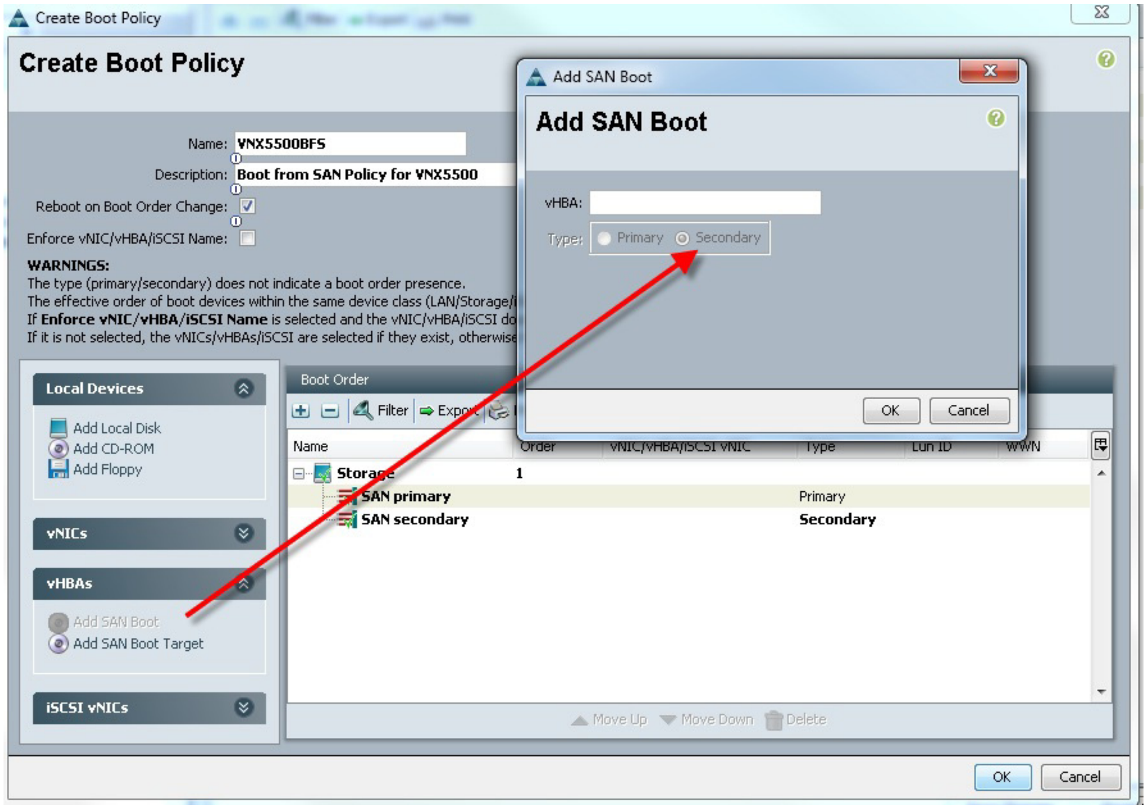

Create Boot Policy

Name: VNX5500BFS

Description: Boot from SAN Policy for VNX5500

Reboot on Boot Order Change: ☑

Enforce vNIC/vHBA/iSCSI Name: □

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/i
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI do
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise

Add SAN Boot

vHBA:

Type: ○ Primary ● Secondary

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Storage	1				
SAN primary			Primary		
SAN secondary			Secondary		

34. Add Boot target WWPN to the SAN Primary, make sure this matches the EMC VNX pwwn. To avoid any typos, copy and paste from Cisco UCS 6248UP command as follows from each Fabric Interconnect:

```
SJC02-151-UCS-VSPEX-A(nxos)# show fcns database vsan 1
0x5600ef    N     50:06:01:6f:3e:a0:64:c8 (Clariion)      scsi-fcp:both
0x5601ef    N     50:06:01:67:3e:a0:64:c8 (Clariion)      scsi-fcp:both
```

```
SJC02-151-UCS-VSPEX-B(nxos)#  show fcns database vsan 1
0xbf00ef    N     50:06:01:66:3e:a0:64:c8 (Clariion)      scsi-fcp:both
0xbf01ef    N     50:06:01:6e:3e:a0:64:c8 (Clariion)      scsi-fcp:both
```

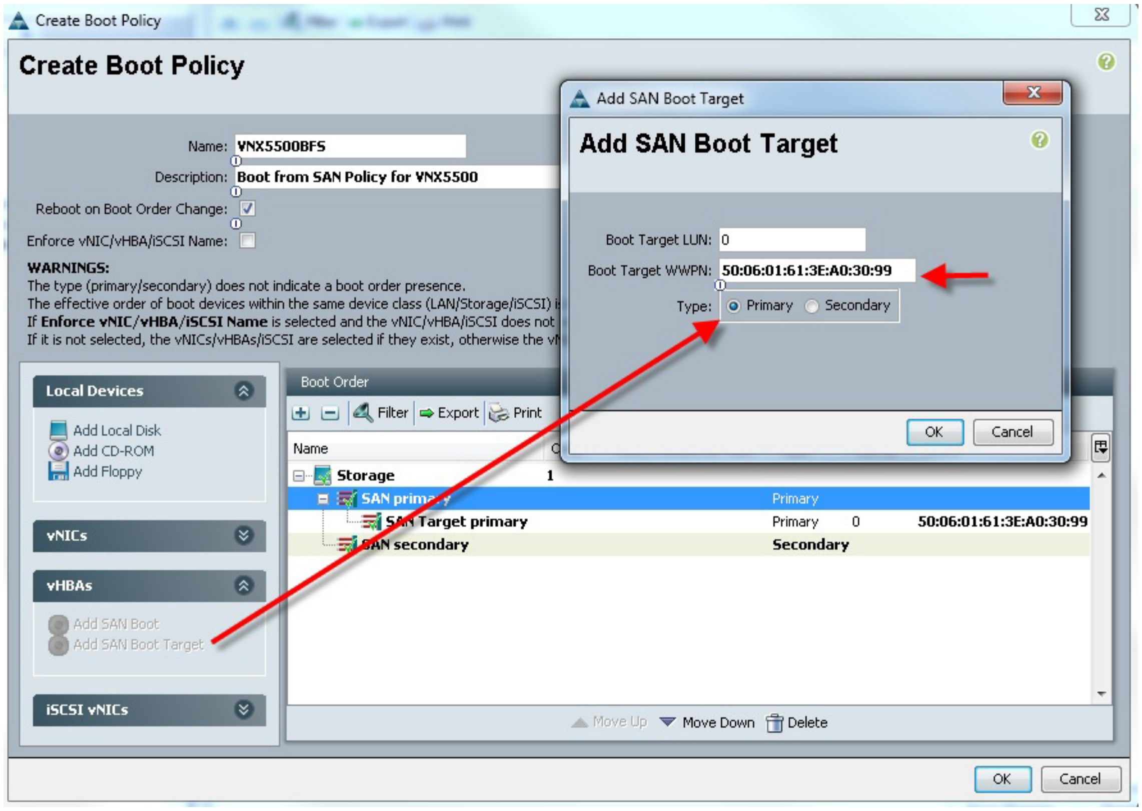

Create Boot Policy

Name: VNX5500BFS

Description: Boot from SAN Policy for VNX5500

Reboot on Boot Order Change: ☑

Enforce vNIC/vHBA/iSCSI Name: ☐

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) i
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vN

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 50:06:01:61:3E:A0:30:99

Type: ● Primary ○ Secondary

Name				
Storage	1			
SAN primary		Primary		
SAN Target primary		Primary	0	50:06:01:61:3E:A0:30:99
SAN secondary		Secondary		

Repeat step 4 for SAN primary's SAN Target Secondary

Repeat step 4 for SAN Secondary's—SAN Target Primary

Repeat step 4 for SAN Secondary's—SAN Target Secondary

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

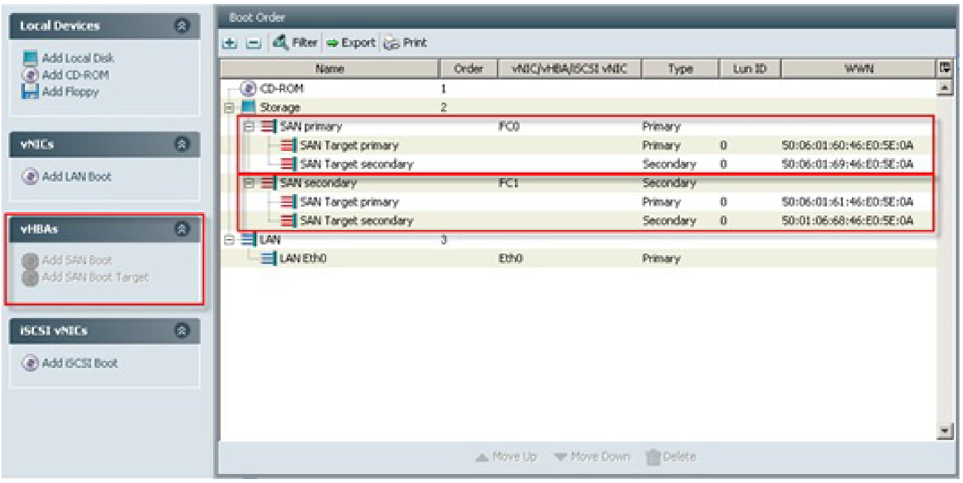

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		FC0	Primary		
SAN Target primary			Primary	0	50:06:01:60:46:E0:5E:0A
SAN Target secondary			Secondary	0	50:06:01:69:46:E0:5E:0A
SAN secondary		FC1	Secondary		
SAN Target primary			Primary	0	50:06:01:61:46:E0:5E:0A
SAN Target secondary			Secondary	0	50:01:06:68:46:E0:5E:0A
LAN	3				
LAN Eth0		Eth0	Primary		

35. The Cisco UCS C220 M3 Host Firmware Package polices were set for Adapter, CIMC and BIOS.

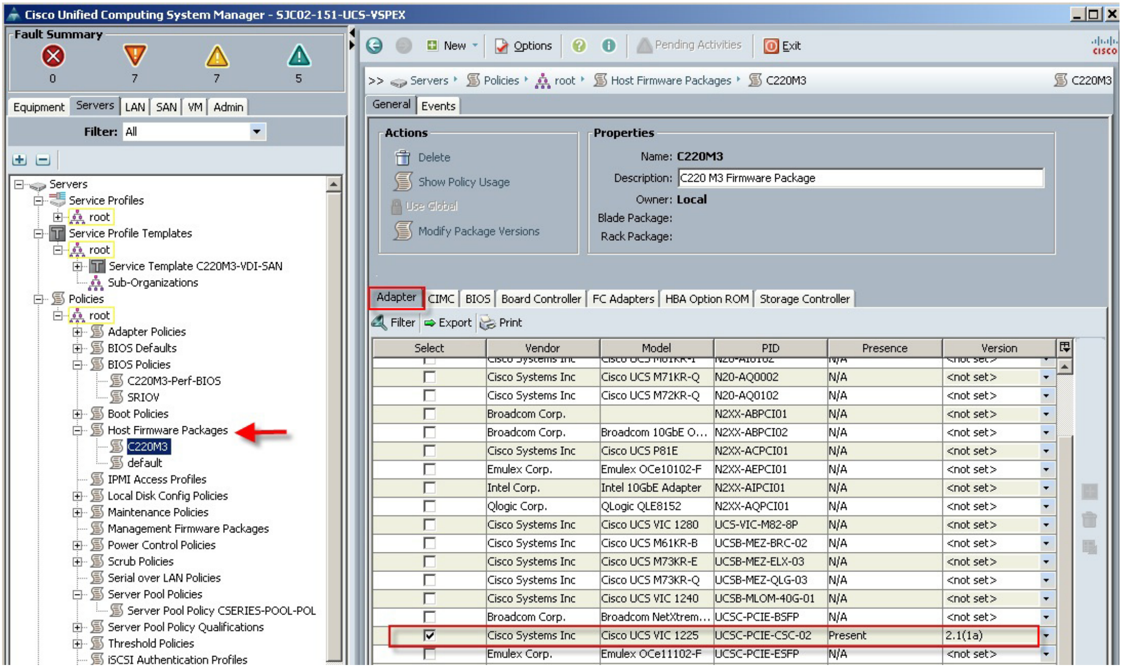

Properties

Name: **C220M3**

Description: C220 M3 Firmware Package

Owner: **Local**

Blade Package:

Rack Package:

Select	Vendor	Model	PID	Presence	Version
□	Cisco Systems Inc	Cisco UCS M81KR-I	N20-AI0102	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS M71KR-Q	N20-AQ0002	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS M72KR-Q	N20-AQ0102	N/A	<not set>
□	Broadcom Corp.		N2XX-ABPCI01	N/A	<not set>
□	Broadcom Corp.	Broadcom 10GbE O...	N2XX-ABPCI02	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS P81E	N2XX-ACPCI01	N/A	<not set>
□	Emulex Corp.	Emulex OCe10102-F	N2XX-AEPCI01	N/A	<not set>
□	Intel Corp.	Intel 10GbE Adapter	N2XX-AIPCI01	N/A	<not set>
□	Qlogic Corp.	QLogic QLE8152	N2XX-AQPCI01	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS VIC 1280	UCS-VIC-M82-8P	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS M61KR-B	UCSB-MEZ-BRC-02	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS M73KR-E	UCSB-MEZ-ELX-03	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS M73KR-Q	UCSB-MEZ-QLG-03	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS VIC 1240	UCSB-MLOM-40G-01	N/A	<not set>
□	Broadcom Corp.	Broadcom NetXtrem...	UCSC-PCIE-BSFP	N/A	<not set>
☑	Cisco Systems Inc	Cisco UCS VIC 1225	UCSC-PCIE-CSC-02	Present	2.1(1a)
□	Emulex Corp.	Emulex OCe11102-F	UCSC-PCIE-ESFP	N/A	<not set>

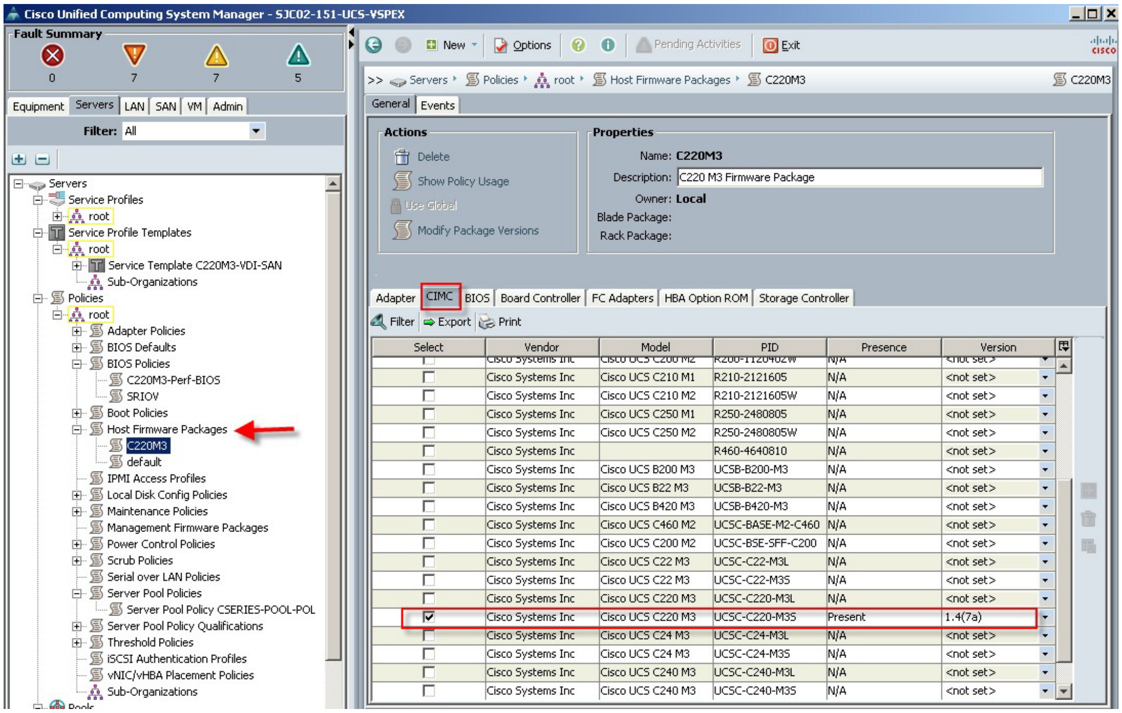

Properties

Name: **C220M3**

Description: C220 M3 Firmware Package

Owner: **Local**

Blade Package:

Rack Package:

Select	Vendor	Model	PID	Presence	Version
□	Cisco Systems Inc	Cisco UCS C200 M2	R200-1120402W	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C210 M1	R210-2121605	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C210 M2	R210-2121605W	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C250 M1	R250-2480805	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C250 M2	R250-2480805W	N/A	<not set>
□	Cisco Systems Inc		R460-4640810	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS B200 M3	UCSB-B200-M3	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS B22 M3	UCSB-B22-M3	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS B420 M3	UCSB-B420-M3	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C460 M2	UCSC-BASE-M2-C460	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C200 M2	UCSC-BSE-SFF-C200	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C22 M3	UCSC-C22-M3L	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C22 M3	UCSC-C22-M3S	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C220 M3	UCSC-C220-M3L	N/A	<not set>
☑	Cisco Systems Inc	Cisco UCS C220 M3	UCSC-C220-M3S	Present	1.4(7a)
□	Cisco Systems Inc	Cisco UCS C24 M3	UCSC-C24-M3L	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C24 M3	UCSC-C24-M3S	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C240 M3	UCSC-C240-M3L	N/A	<not set>
□	Cisco Systems Inc	Cisco UCS C240 M3	UCSC-C240-M3S	N/A	<not set>

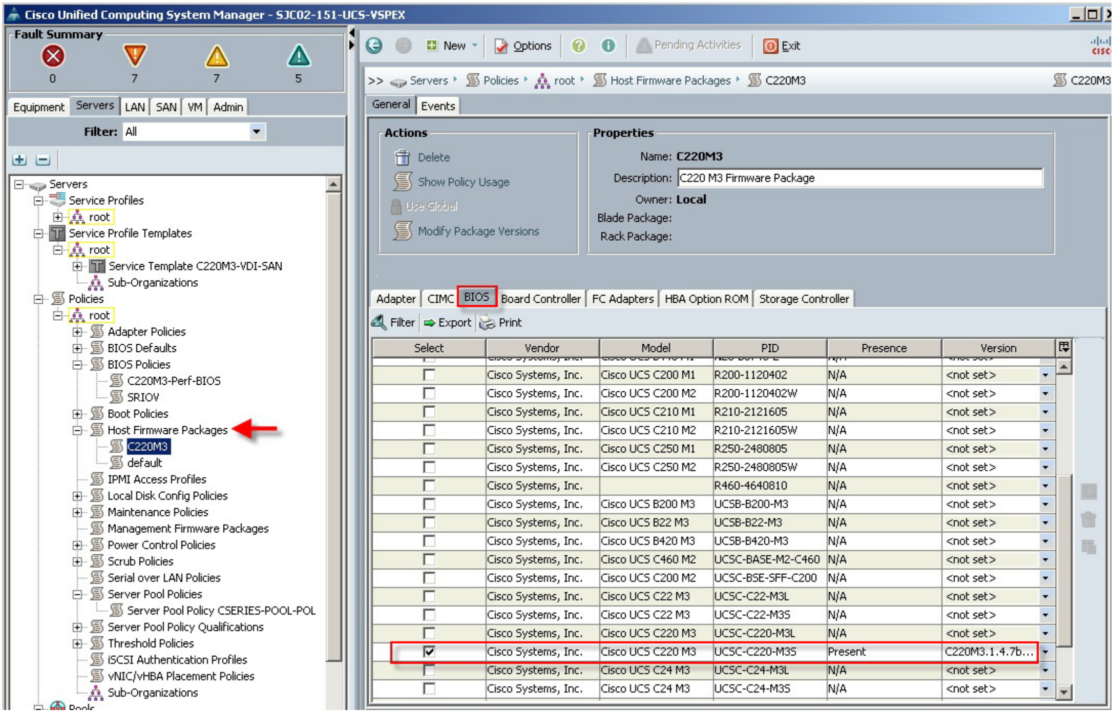

Properties

Name: **C220M3**

Description: C220 M3 Firmware Package

Owner: **Local**

Blade Package:

Rack Package:

Select	Vendor	Model	PID	Presence	Version
☐	Cisco Systems, Inc.	Cisco UCS C200 M1	R200-1120402	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C200 M2	R200-1120402W	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C210 M1	R210-2121605	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C210 M2	R210-2121605W	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C250 M1	R250-2480805	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C250 M2	R250-2480805W	N/A	<not set>
☐	Cisco Systems, Inc.		R460-4640810	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS B200 M3	UCSB-B200-M3	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS B22 M3	UCSB-B22-M3	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS B420 M3	UCSB-B420-M3	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C460 M2	UCSC-BASE-M2-C460	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C200 M2	UCSC-BSE-SFF-C200	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C22 M3	UCSC-C22-M3L	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C22 M3	UCSC-C22-M3S	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C220 M3	UCSC-C220-M3L	N/A	<not set>
☑	Cisco Systems, Inc.	Cisco UCS C220 M3	UCSC-C220-M3S	Present	C220M3.1.4.7b...
☐	Cisco Systems, Inc.	Cisco UCS C24 M3	UCSC-C24-M3L	N/A	<not set>
☐	Cisco Systems, Inc.	Cisco UCS C24 M3	UCSC-C24-M3S	N/A	<not set>

36. Set FC Switching Mode on the Fabric Interconnects to enable Cisco UCS Local Zoning feature.

37. In Cisco UCS Manager, navigate to the "SAN Tab" in the navigation pane, Select top level "SAN" tab in the navigation tree. In the Main window, select the "SAN Uplinks Tab" which will display the "Port and Port Channels" and "SAN Pin Groups" windows.

38. Click the "SAN Uplinks Manager" in the Main window.

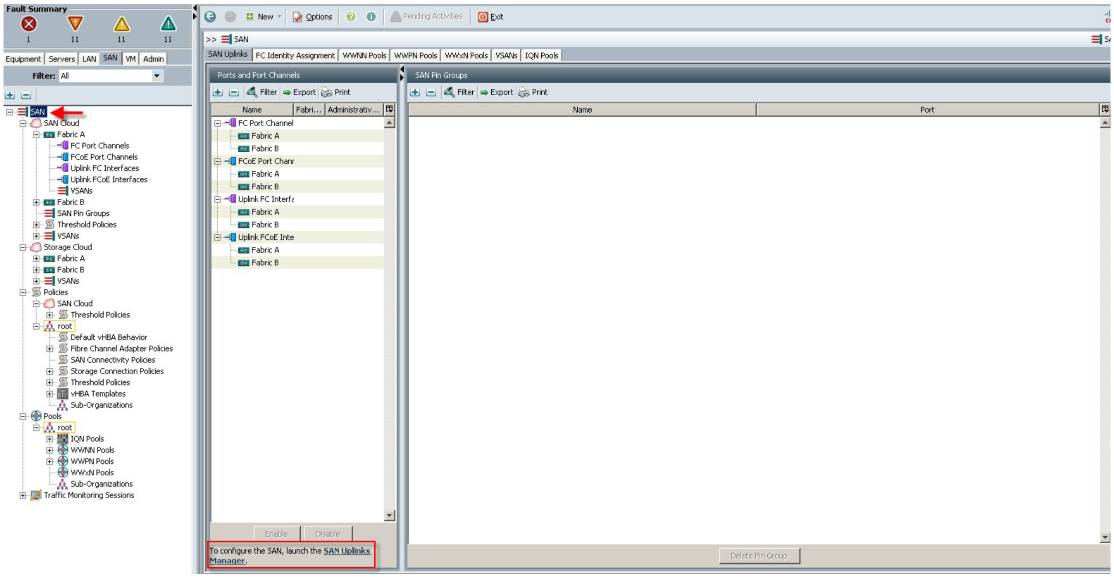

To configure the SAN, launch the SAN Uplinks Manager.

The SAN Uplinks Manager windows displays.

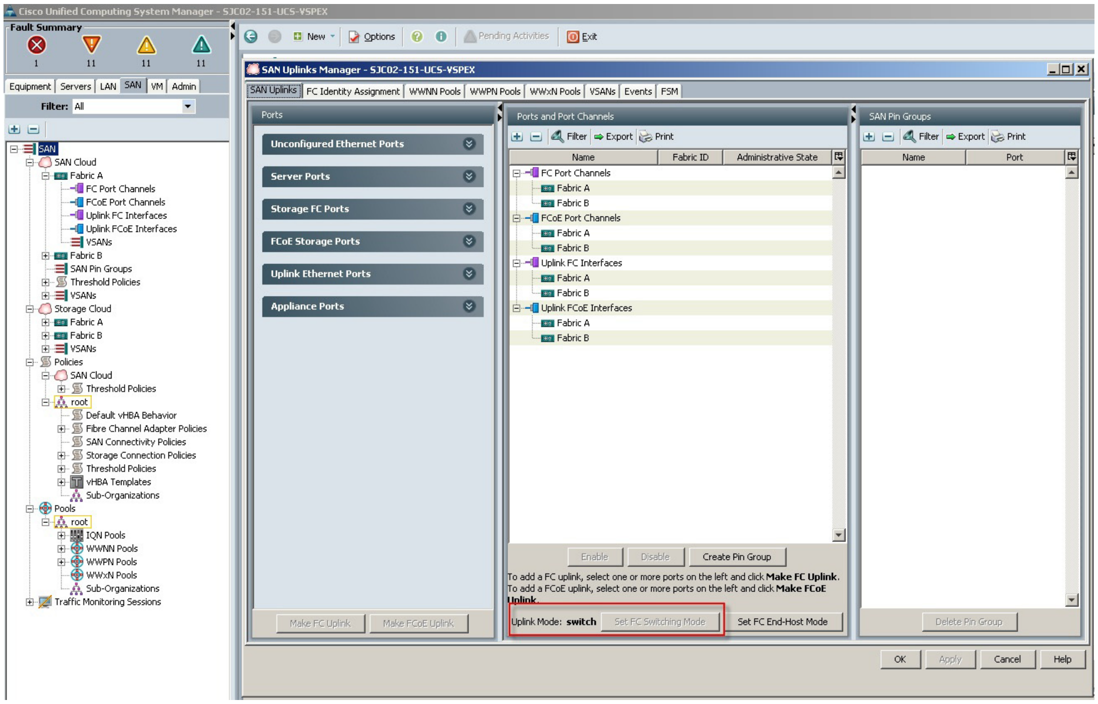

To add a FC uplink, select one or more ports on the left and click **Make FC Uplink**.
To add a FCoE uplink, select one or more ports on the left and click **Make FCoE Uplink**.

Uplink Mode: **switch**

39. Click the Set FC Switching Mode button.

Note Changing the FC Uplink Mode will result in both Fabric Interconnects' expansion modules immediately rebooting resulting in a 10-15 minute outage. Changing Uplink Modes should only occur during a planned maintenance window.

40. The Uplink Mode will show whether Cisco Unified Computing System is currently in "End-Host" mode or "Switching" mode. If Cisco Unified Computing System is already in "FC Switching" mode click the "Cancel" Button and proceed to the next steps. If Cisco Unified Computing System is in "End-Host" mode, click the button "Set FC Switching Mode" to change the Uplink Mode to FC Switching Mode.

41. If you are using the default VSAN 1, from the SAN tab under the Storage Cloud node, right-click the VSAN default(1) object and choose Show Navigator. In the FC Zoning Settings area, click the FC Zoning: Enabled radio button, then click OK.

Note Use the default VSAN 1, a dual fabric VSAN, with caution. In situations where you will change the VLANs down the road, migrating from the default VSAN will cause both sides of your storage connectivity to go down simultaneously. Best practice is to configure single fabric VSANs for this reason.

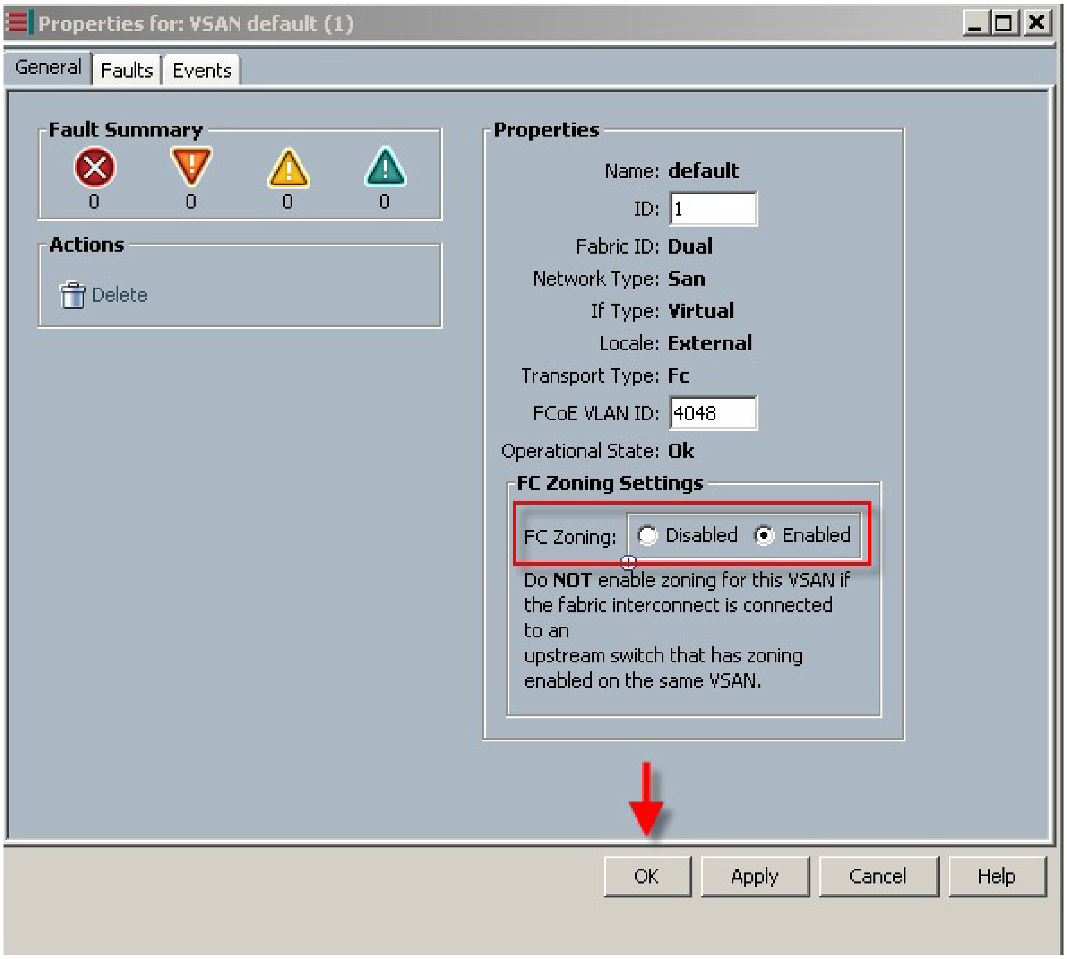

Properties

Name: **default**

ID: 1

Fabric ID: **Dual**

Network Type: **San**

If Type: **Virtual**

Locale: **External**

Transport Type: **Fc**

FCoE VLAN ID: 4048

Operational State: **Ok**

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do **NOT** enable zoning for this VSAN if the fabric interconnect is connected to an
upstream switch that has zoning enabled on the same VSAN.

42. Click OK to complete the FC Zoning enablement.
43. Create a service profile template using the pools, templates, and policies configured above.

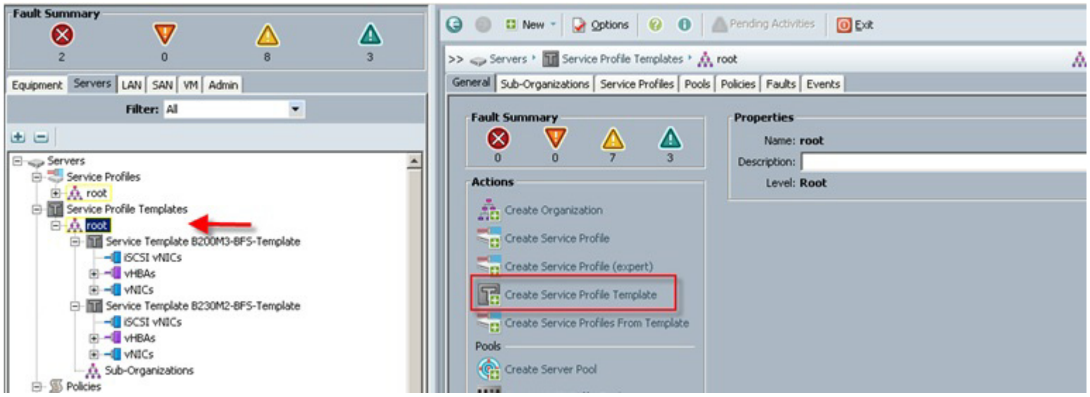

Properties

Name: **root**

Description:

Level: **Root**

In this project, we created one template for the Cisco UCS C220 M3 Rack-Mount server model used.

44. Follow 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 insure maximum LV DIMM performance.

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

45. On the Create Service Profile Template wizard, enter a unique name, select the type as Updating Template, and selected the VDI-UUID-Pool created earlier, then click Next.

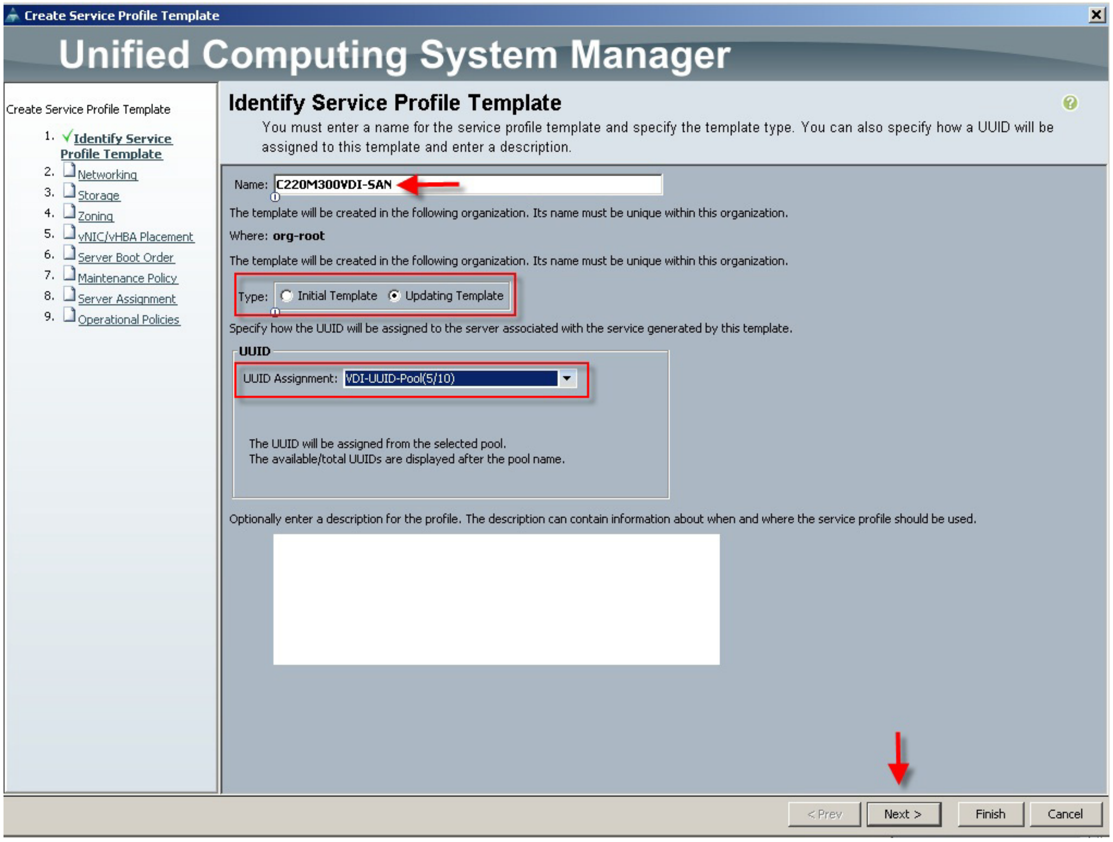

Identify Service Profile Template

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

Name: C220M300VDI-SAN

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

Where: **org-root**

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

Type: ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: VDI-UUID-Pool(5/10)

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

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

46. On the Network page, we select Expert mode, we click Add to create virtual NICs.

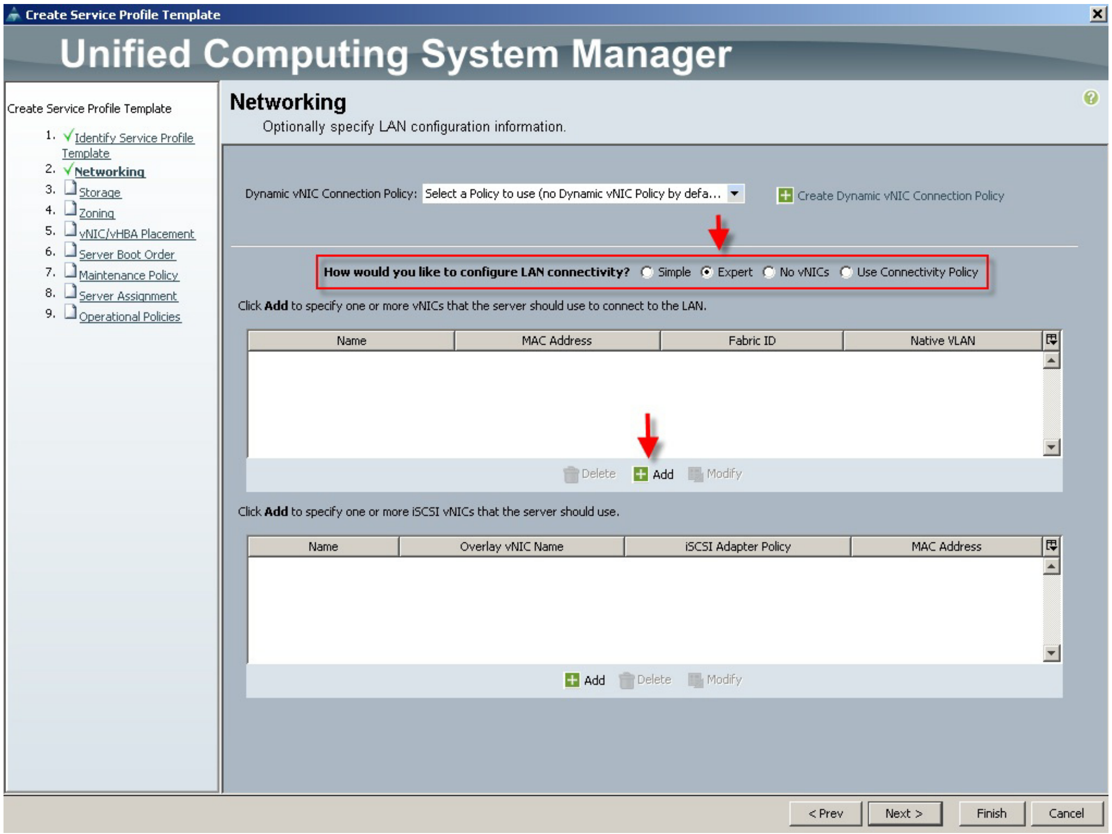

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ **Networking**
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

Networking

Optionally specify LAN configuration information.

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

How would you like to configure LAN connectivity? ○ Simple ● Expert ○ No vNICs ○ Use Connectivity Policy

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

Name	MAC Address	Fabric ID	Native VLAN

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

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address

47. On the Create vNIC page, provide a name, typically eth0, eth1, eth2, etc. Check the Use vNIC Template checkbox. Use the drop-down to select one of the vNIC templates created earlier. Use the drop-down list to choose the VMware Adapter Policy

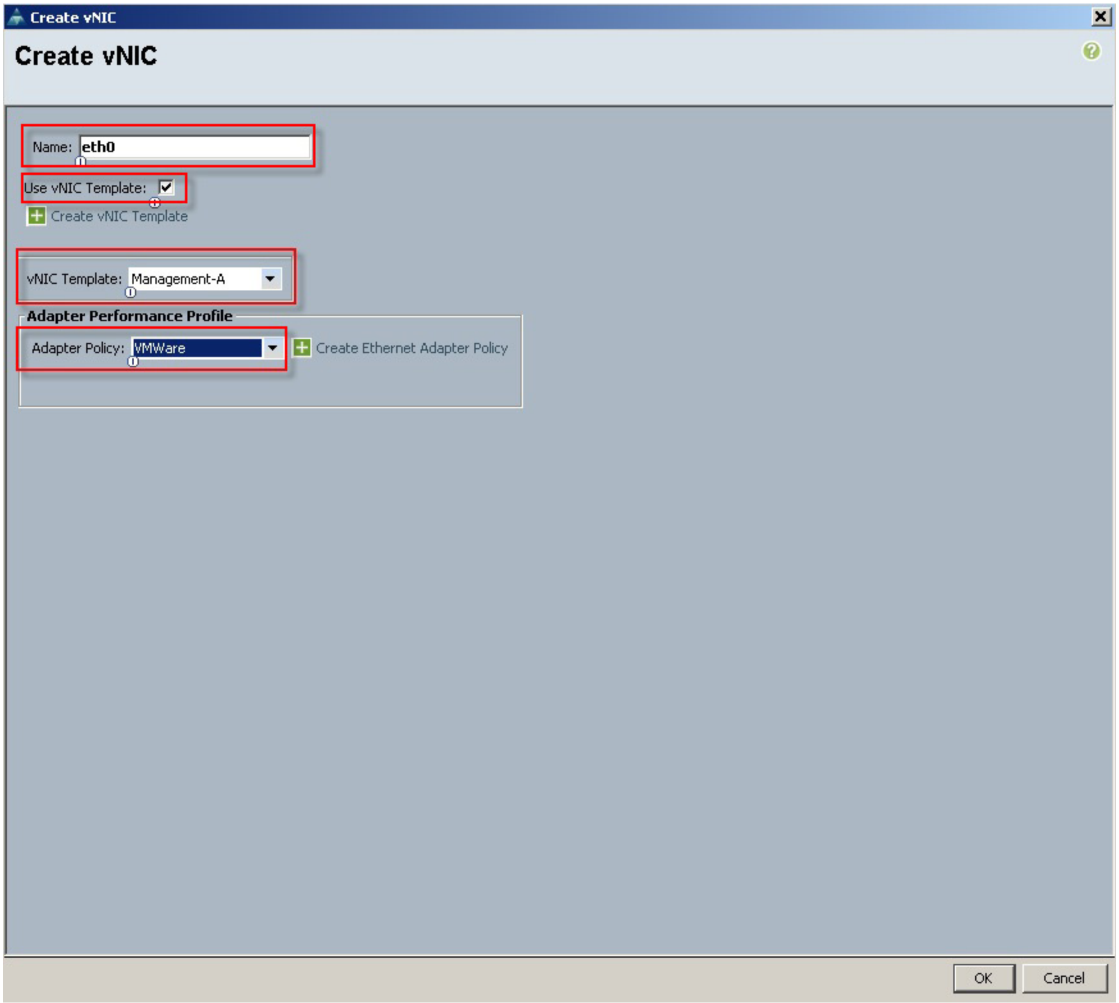

Create vNIC

Name: eth0

Use vNIC Template: ☑

vNIC Template: Management-A

Adapter Performance Profile

Adapter Policy: VMWare

48. We repeat this process to create eth1, eth2, eth3 eth4 and eth5 using a different vNIC template for each vNIC. Then click Next.

Note eth5 is not shown in the graphic.

49. On the Storage page, we selected the Expert mode, we selected the WWNN Pool we created earlier from the drop down list and then click Add.

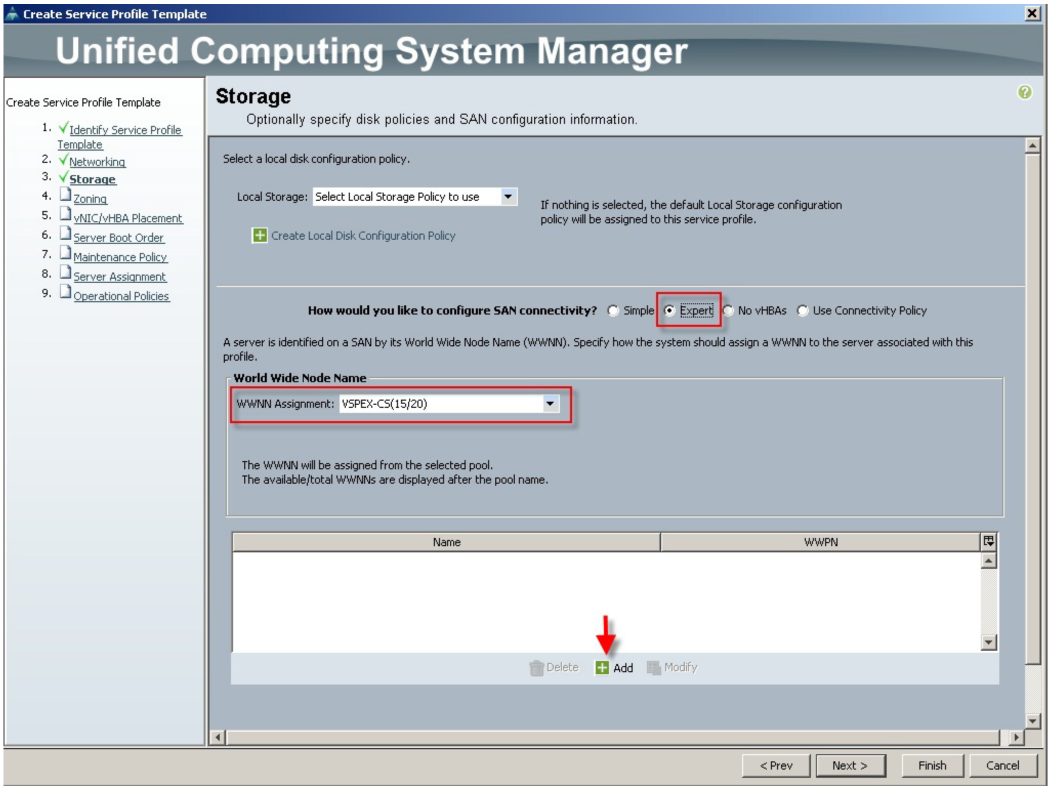

Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: Select Local Storage Policy to use

If nothing is selected, the default Local Storage configuration policy will be assigned to this service profile.

How would you like to configure SAN connectivity? ○ Simple ● Expert ○ No vHBAs ○ Use Connectivity Policy

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

World Wide Node Name

WWNN Assignment: VSPEX-CS(15/20)

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

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

50. On the Create HBA page, enter a name (FC0) and check Use SAN Connectivity Template, which changed the display to the following:

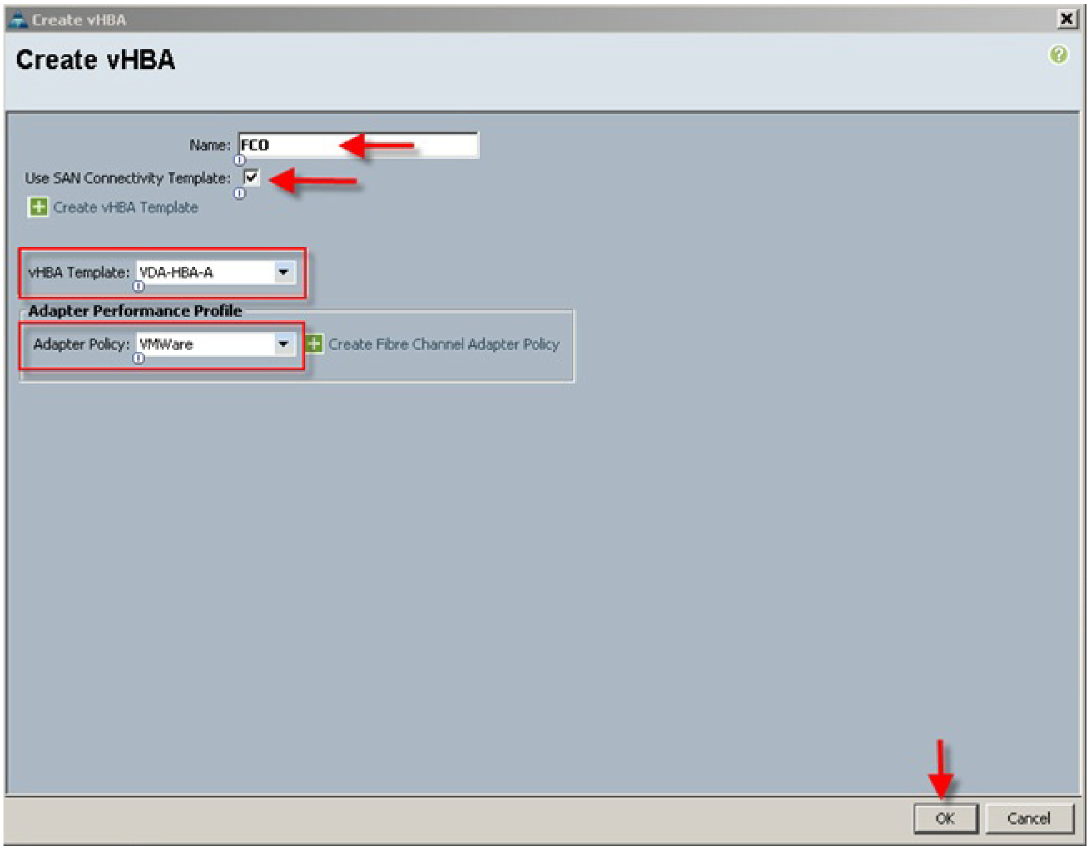
Create vHBA
Create vHBA
Create vHBA Template
Adapter Performance Profile
Create Fibre Channel Adapter Policy
OK
Cancel

Name: FC0

Use SAN Connectivity Template: ☑

vHBA Template: VDA-HBA-A

Adapter Policy: VMWare

51. Select the vHBA template for Fabric Interconnect A and the VMware Adapter Policy from the drop downs, then click OK.

52. Repeat the process for FC1, choosing VDA-HBA-B for Fabric Interconnect B. The Storage page displays as follows:

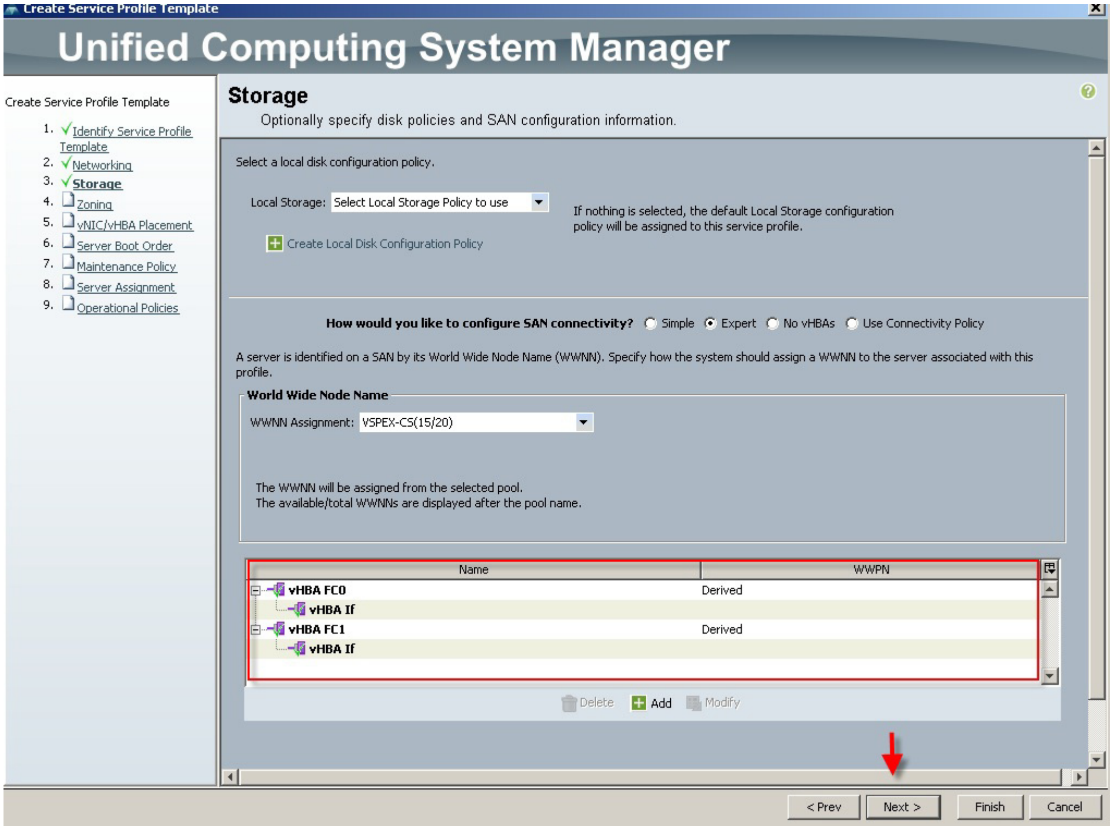

Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: Select Local Storage Policy to use

If nothing is selected, the default Local Storage configuration policy will be assigned to this service profile.

Create Local Disk Configuration Policy

How would you like to configure SAN connectivity? ○ Simple ● Expert ○ No vHBAs ○ Use Connectivity Policy

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

World Wide Node Name

WWNN Assignment: VSPEX-CS(15/20)

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

Name	WWPN
vHBA FC0	Derived
vHBA If	
vHBA FC1	Derived
vHBA If	

53. Click Next to continue.

54. Part of the process for creating a Service Profile Template in Cisco UCS Manager 2.1 is to perform FC Zoning tasks. When you reach this task in the wizard, the vHBA Initiators added in step 3 appear in the Select vHBA Initiators section of the right pane. Before you can proceed, you must add vHBA Initiator Groups created previously or you can create them here.

55. Click the Add option in the Select vHBA Initiator Groups section in the right pane.

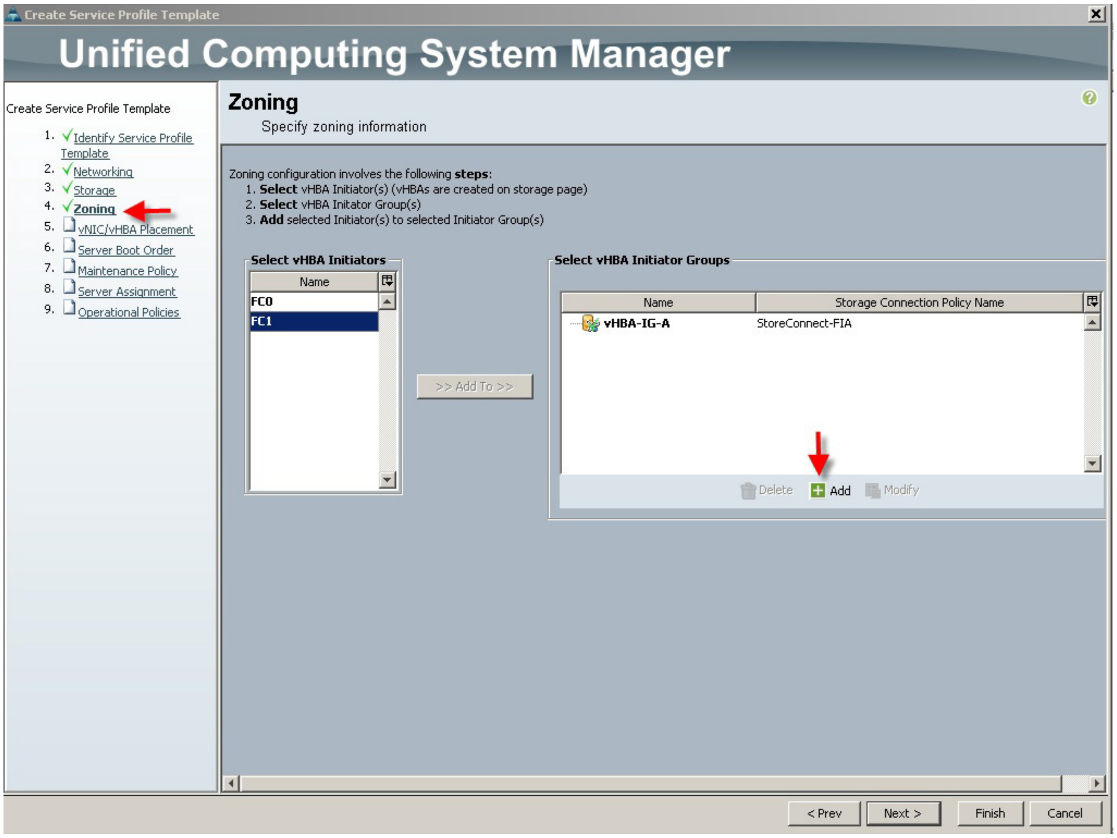

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

Zoning

Specify zoning information

Zoning configuration involves the following **steps**:

1. **Select** vHBA Initiator(s) (vHBAs are created on storage page)
2. **Select** vHBA Initator Group(s)
3. **Add** selected Initiator(s) to selected Initiator Group(s)

Select vHBA Initiators

Name
FC0
FC1

Select vHBA Initiator Groups

Name	Storage Connection Policy Name
vHBA-IG-A	StoreConnect-FIA

56. On the Create vHBA Initiator Group window, provide a unique name, optional description and then select a previously created Storage Connection Policy or click the +Create Storage Connection Policy control.

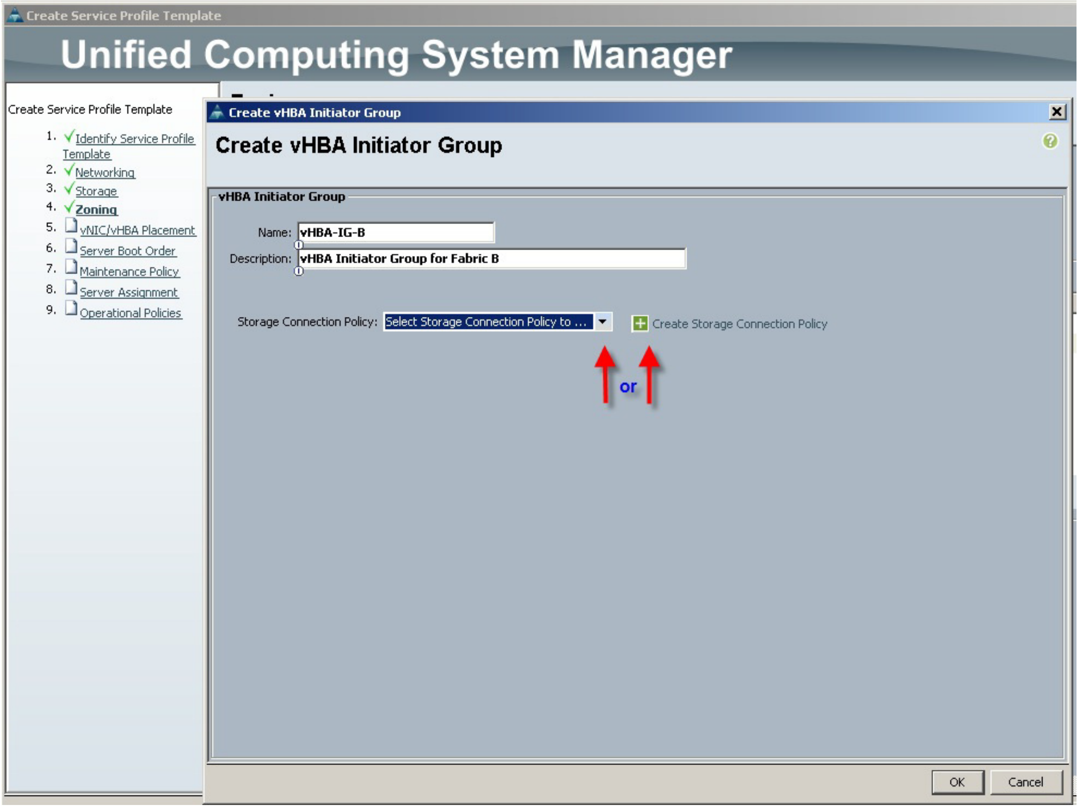
Create Service Profile Template
Unified Computing System Manager
Create Service Profile Template
1. Identify Service Profile Template
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies
Create vHBA Initiator Group
Create Storage Connection Policy
or
OK
Cancel

Create vHBA Initiator Group

vHBA Initiator Group

Name: vHBA-IG-B

Description: vHBA Initiator Group for Fabric B

Storage Connection Policy: Select Storage Connection Policy to ...

57. In this case, we click the + Create Storage Connection Policy to demonstrate creating a new policy for Fabric B:

58. In the Create Storage Connection Policy window, provide a unique policy name, an optional description, select the Single Initiator Multiple Targets radio button when you have physically provisioned paths to multiple FC targets on each Fabric Interconnect (recommended,) and then click the + (Add) control on the right edge of the FC Target Endpoints area.

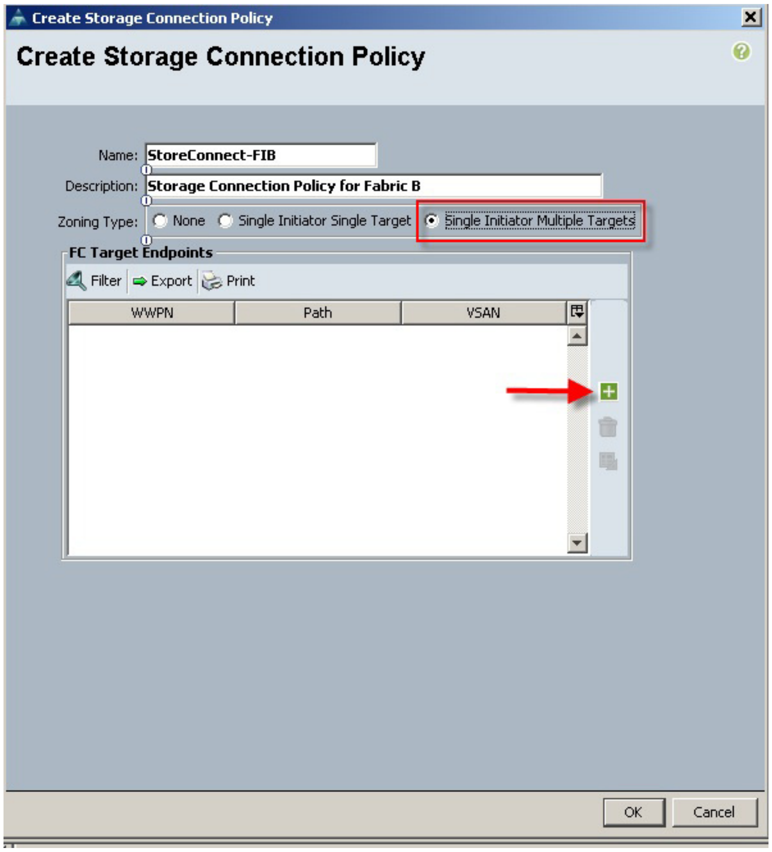

Create Storage Connection Policy

Name: StoreConnect-FIB

Description: Storage Connection Policy for Fabric B

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN

59. Input the World Wide Port Name value from the FC storage controllers that are connected to fabric B in our example case below, provide and optional description, select the B radio button in the path section, and select an existing VSAN or create a new one corresponding to a VSAN on your SAN controller.

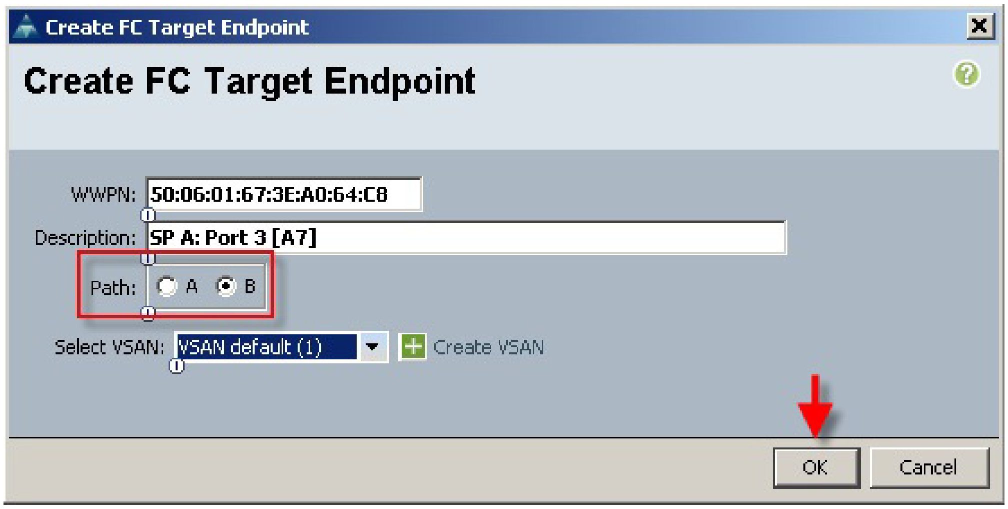

Create FC Target Endpoint

WWPN: 50:06:01:67:3E:A0:64:C8

Description: SP A: Port 3 [A7]

Path: ○ A ● B

Select VSAN: VSAN default (1)

60. Repeat the process to add the second port to the SAN Controller.

61. Click OK to add the Storage Connection Policy.

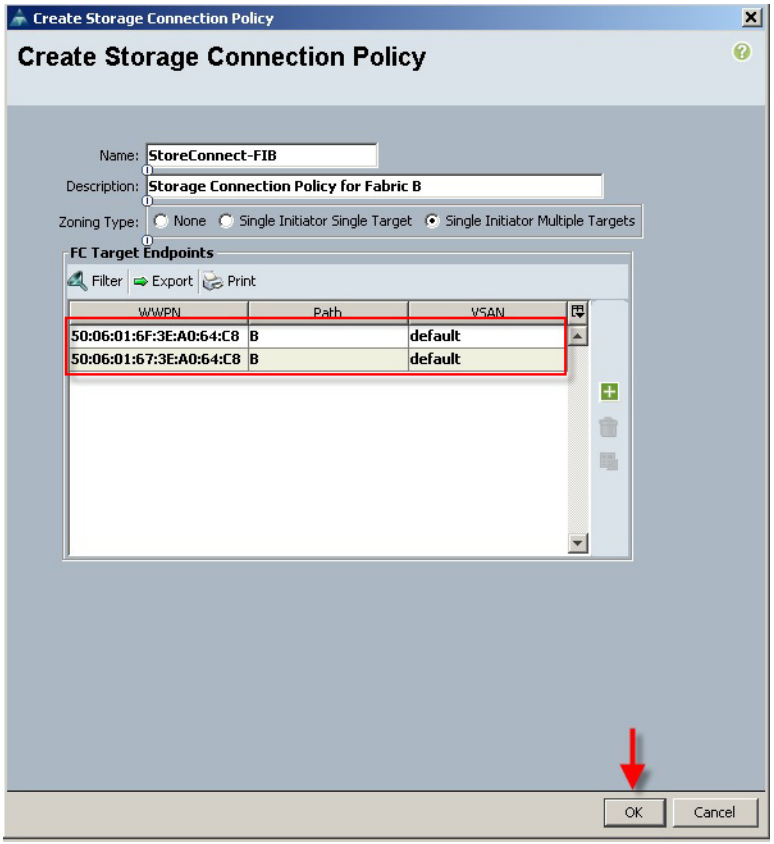
Create Storage Connection Policy

Name: StoreConnect-FIB

Description: Storage Connection Policy for Fabric B

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN
50:06:01:6F:3E:A0:64:C8	B	default
50:06:01:67:3E:A0:64:C8	B	default

62. Click OK to acknowledge successful creation of the Storage Connection Policy. Repeat the process to create a Storage Connection Policy for Fabric A

63. From the drop-down menu, add the Storage Connection Policy created above to the vHBA Initiator Group. Notice that the Connection Policy details are added to the dialogue. Click OK to create the vHBA Initiator Group.

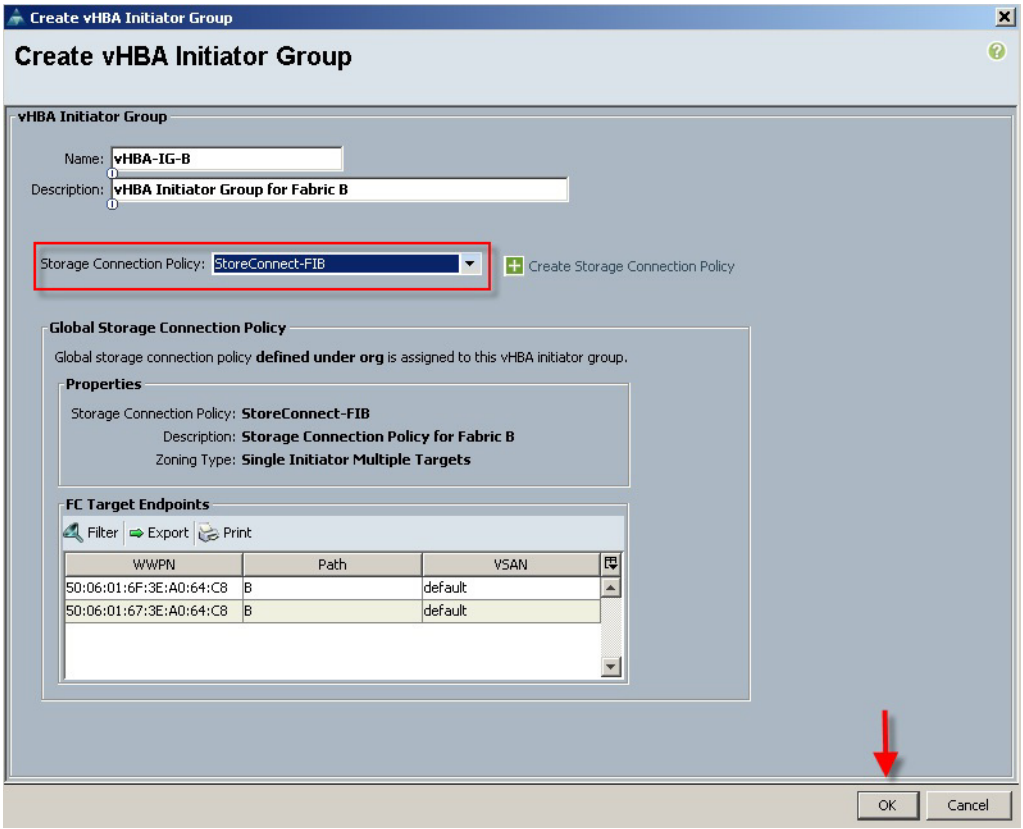

Create vHBA Initiator Group

vHBA Initiator Group

Name: vHBA-IG-B

Description: vHBA Initiator Group for Fabric B

Storage Connection Policy: StoreConnect-FIB

Global Storage Connection Policy

Global storage connection policy **defined under org** is assigned to this vHBA initiator group.

Properties

Storage Connection Policy: **StoreConnect-FIB**

Description: **Storage Connection Policy for Fabric B**

Zoning Type: **Single Initiator Multiple Targets**

FC Target Endpoints

WWPN	Path	VSAN
50:06:01:6F:3E:A0:64:C8	B	default
50:06:01:67:3E:A0:64:C8	B	default

You are returned to the Zoning page.

64. In the Select vHBA Initiators area, click one of the initiators, then in the Select vHBA Initiator Groups area, click on the group that is corresponds to the vHBA initiator. In our case, we are connecting FC0, which is configured for Fabric A to the vHBA-IG-A which is also configured for Fabric A. Click the Add To control between the lists to complete the association.

65. Repeat the process for FC1 and vHBA-IG-B. Your configuration should look similar to the following:

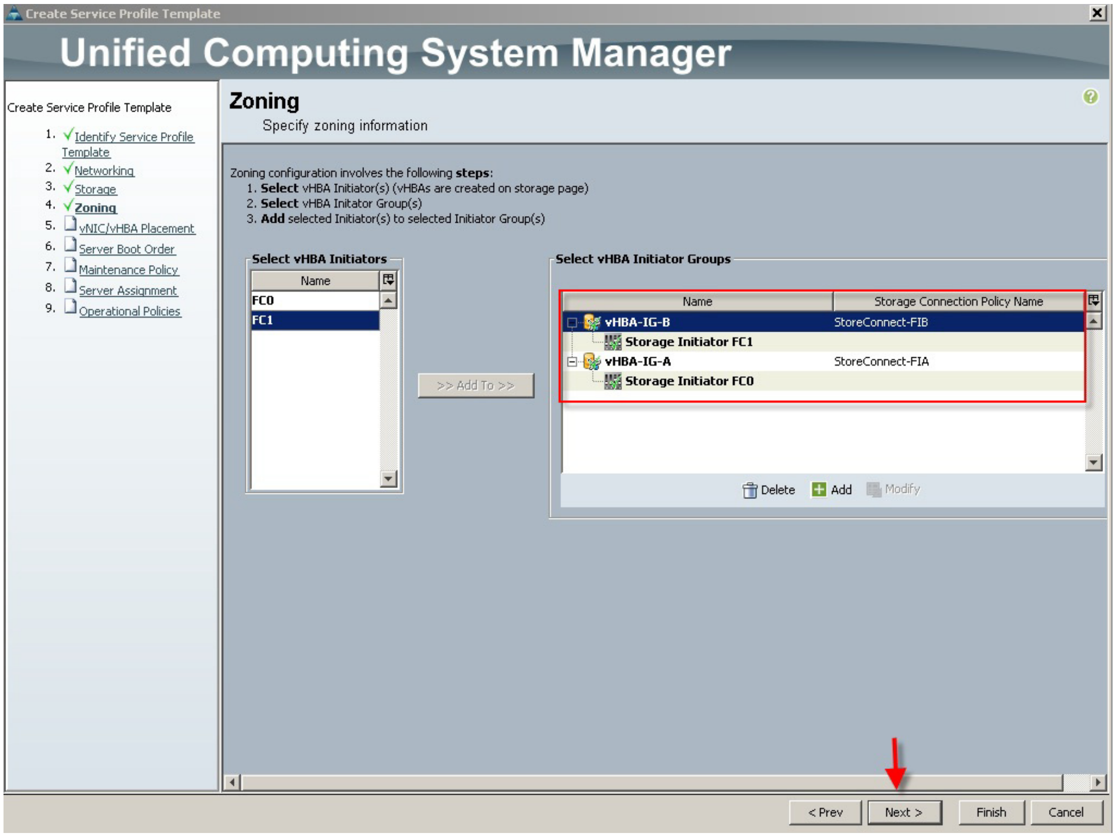

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ **Zoning**
5. □ vNIC/vHBA Placement
6. □ Server Boot Order
7. □ Maintenance Policy
8. □ Server Assignment
9. □ Operational Policies

Zoning

Specify zoning information

Zoning configuration involves the following **steps**:
1. **Select** vHBA Initiator(s) (vHBAs are created on storage page)
2. **Select** vHBA Initator Group(s)
3. **Add** selected Initiator(s) to selected Initiator Group(s)

Select vHBA Initiators

Name
FC0
FC1

Select vHBA Initiator Groups

Name	Storage Connection Policy Name
vHBA-IG-B	StoreConnect-FIB
Storage Initiator FC1	
vHBA-IG-A	StoreConnect-FIA
Storage Initiator FC0	

66. Click Next to continue.

67. You can use the vNIC/vHBA Placement step to set the placement order of the ethernet and fiber channel adapters for the template.

Note You can create a policy to set the placement order which is useful if you are creating multiple profiles.

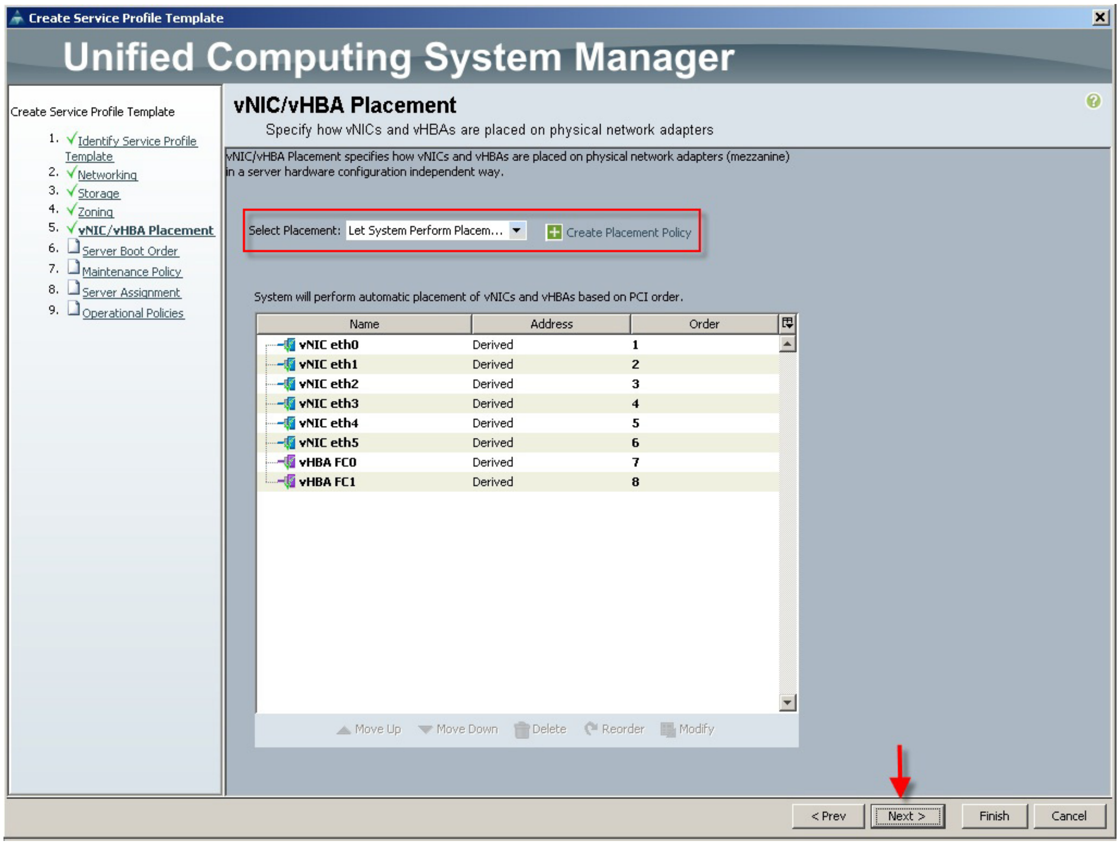
Create Service Profile Template
Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

vNIC/vHBA Placement

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: Let System Perform Placem...

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order
vNIC eth0	Derived	1
vNIC eth1	Derived	2
vNIC eth2	Derived	3
vNIC eth3	Derived	4
vNIC eth4	Derived	5
vNIC eth5	Derived	6
vHBA FC0	Derived	7
vHBA FC1	Derived	8

68. Accept the default placement and click Next.

69. Select the Boot Policy created earlier, use an existing default boot policy, or Create Boot Policy by clicking the + control. In our case, we selected the Boot from SAN policy created earlier. When selected, the details of the policy are displayed in the lower portion of the window.

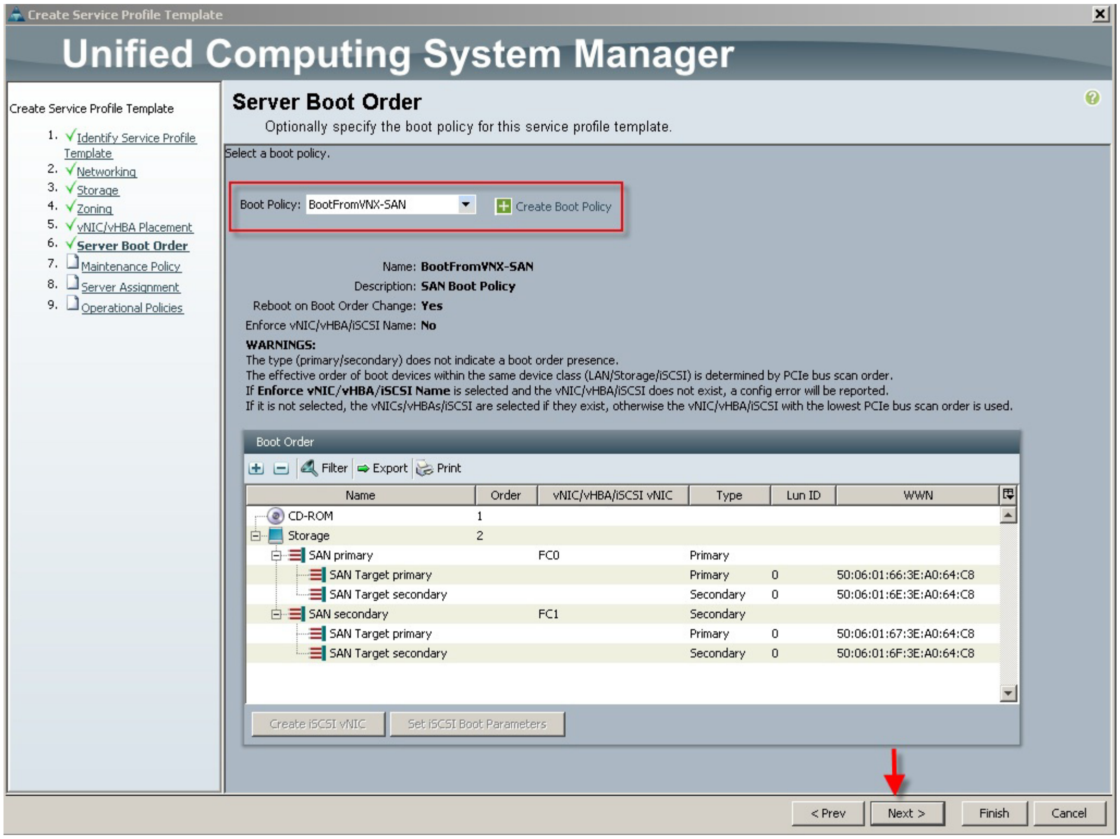

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: BootFromVNX-SAN

Name: **BootFromVNX-SAN**
Description: **SAN Boot Policy**
Reboot on Boot Order Change: **Yes**
Enforce vNIC/vHBA/iSCSI Name: **No**

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

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		FC0	Primary		
SAN Target primary			Primary	0	50:06:01:66:3E:A0:64:C8
SAN Target secondary			Secondary	0	50:06:01:6E:3E:A0:64:C8
SAN secondary		FC1	Secondary		
SAN Target primary			Primary	0	50:06:01:67:3E:A0:64:C8
SAN Target secondary			Secondary	0	50:06:01:6F:3E:A0:64:C8

70. Click Next to continue.

71. On the Maintenance Policy page, select the Maintenance Policy previously created, if any, or Create Maintenance Policy by clicking the + control. In this study, no maintenance policy was created nor used.

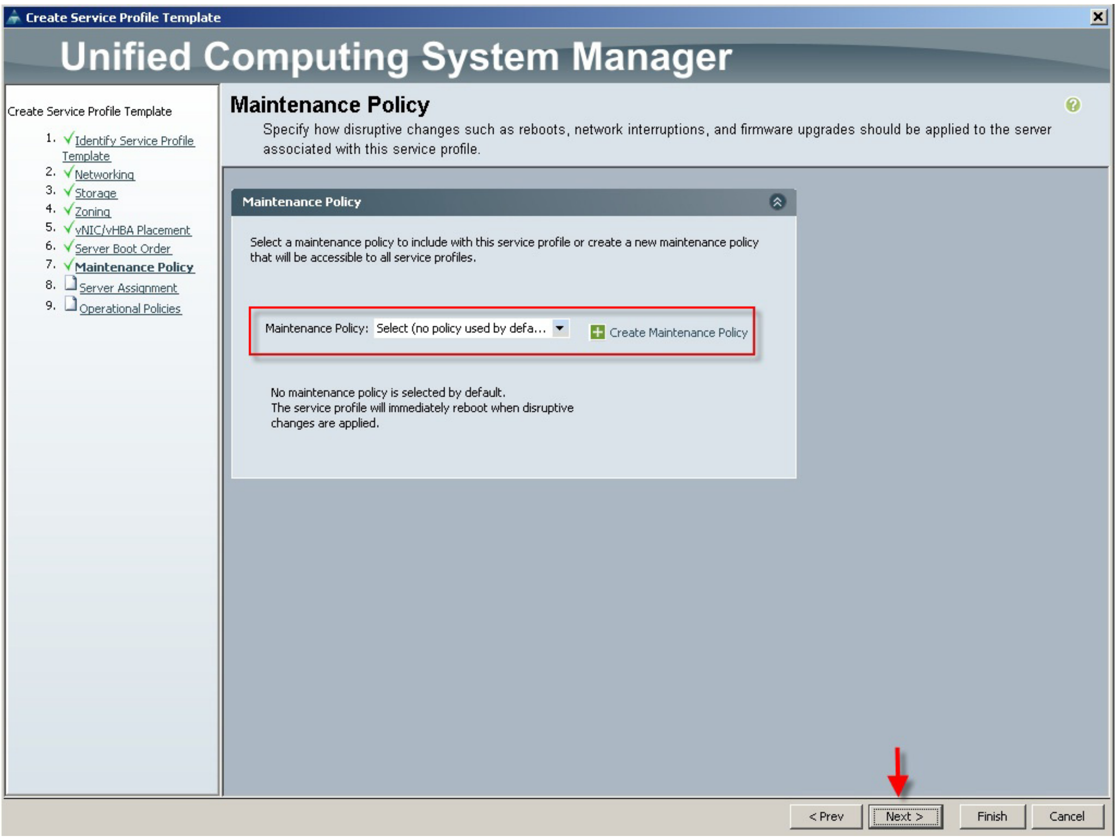

Create Service Profile Template

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. Server Assignment
9. Operational Policies

Maintenance Policy

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: Select (no policy used by defa... ▾ | + Create Maintenance Policy

No maintenance policy is selected by default.
The service profile will immediately reboot when disruptive changes are applied.

72. Click Next to continue.

73. On the Server Assignment page, utilize the Server Pool and Server Pool Qualification from the drop downs. You can create a Server Pool from the wizard, but you cannot create a Server Pool Qualification policy from here. In our study, we utilized the Server Pool and Server Pool Qualification policy we created earlier. We expanded the Firmware Management node near the bottom of the page and selected the policy we created earlier.

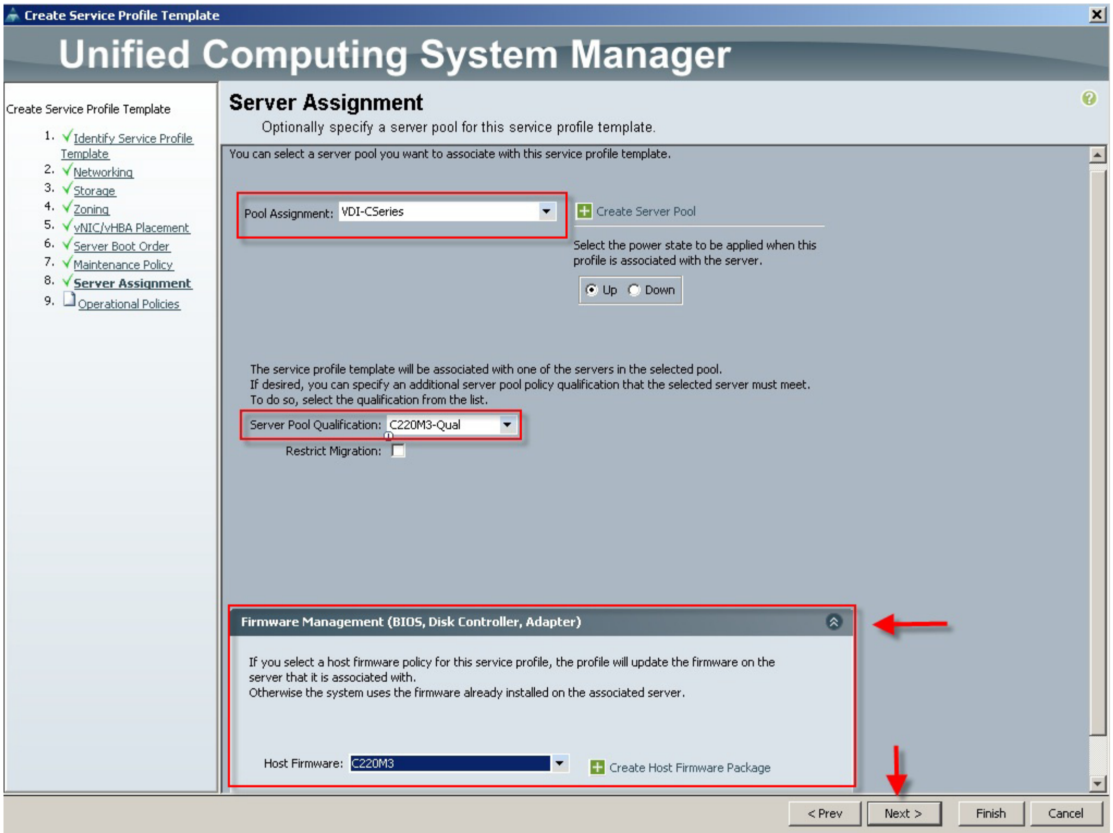

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. ✓ Server Assignment
9. □ Operational Policies

Server Assignment

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: VDI-CSeries

Create Server Pool

Select the power state to be applied when this profile is associated with the server.

● Up ○ Down

The service profile template will be associated with one of the servers in the selected pool.
If desired, you can specify an additional server pool policy qualification that the selected server must meet.
To do so, select the qualification from the list.

Server Pool Qualification: C220M3-Qual

Restrict Migration: □

Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with.
Otherwise the system uses the firmware already installed on the associated server.

Host Firmware: C220M3

Create Host Firmware Package

74. Click Next to continue.

75. In the Operational Policy page, you can configure BIOS, External IPMI Management, Management IP Address, Monitoring, Power Control and Scrub configuration and policies for the Cisco UCS Servers for which the template will be applied.

Note In our study, we configured a BIOS Policy and created a Management IP Address pool which we applied to the template. Note that you can create policies and pools from within this wizard.

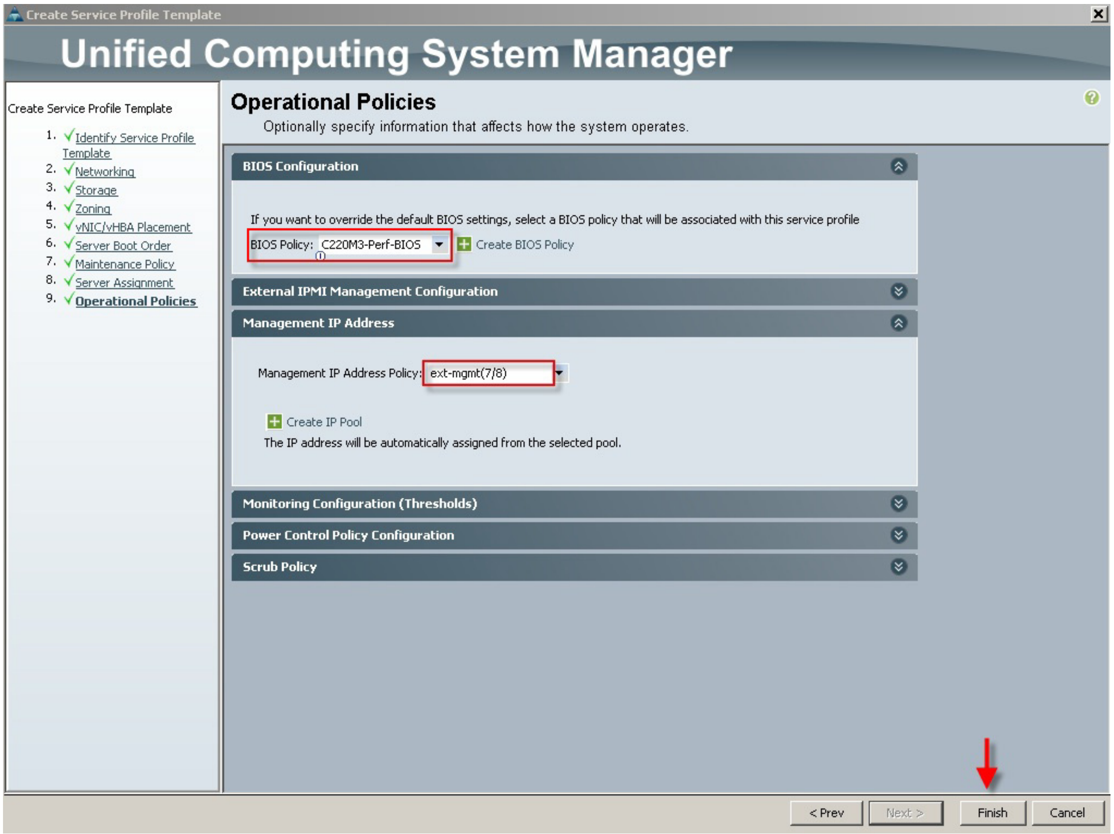

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. ✓ Server Assignment
9. ✓ Operational Policies

Operational Policies

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy: C220M3-Perf-BIOS

External IPMI Management Configuration

Management IP Address

Management IP Address Policy: ext-mgmt(7/8)

The IP address will be automatically assigned from the selected pool.

Create BIOS Policy
Create IP Pool
Monitoring Configuration (Thresholds)
Power Control Policy Configuration
Scrub Policy
< Prev
Next >
Finish
Cancel

76. Click Finish to complete the wizard.

77. Click OK to acknowledge successful creation.

78. To create Service Profiles, from the Servers tab in the navigation page, in the Service Profile Templates node, expand the root and select Service Template C220 M3, then click Create Service Profiles from Template in the right pane, Actions area:

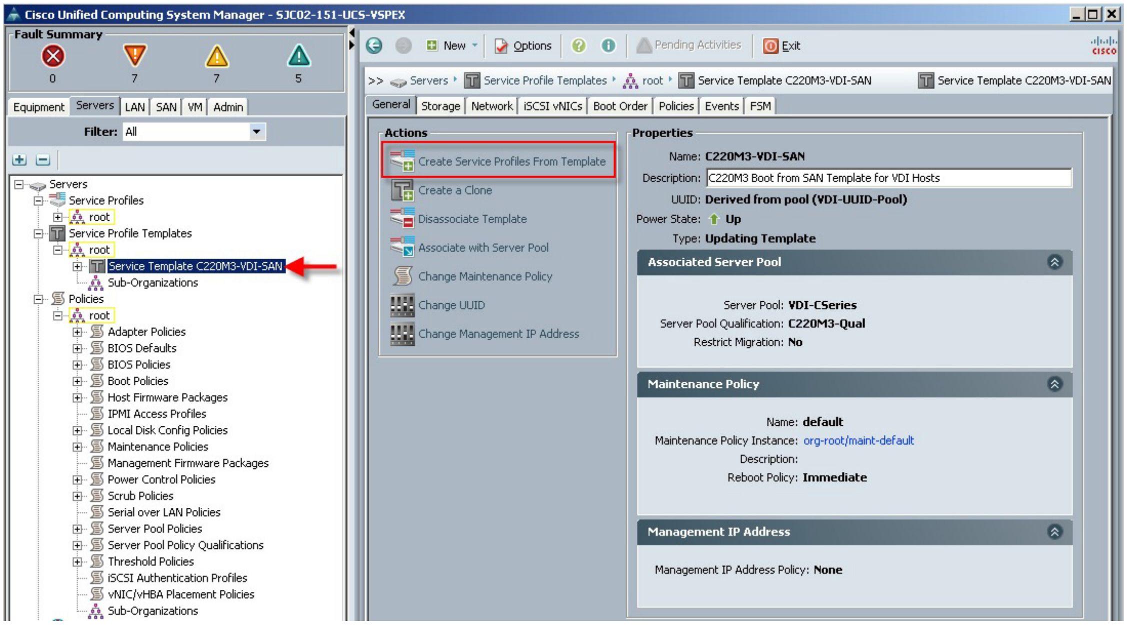
Cisco Unified Computing System Manager - SJC02-151-UCS-VSPEX
Fault Summary
0
7
7
5
Equipment
Servers
LAN
SAN
VM
Admin
New
Options
Pending Activities
Exit
>> Servers > Service Profile Templates > root > Service Template C220M3-VDI-SAN
Service Template C220M3-VDI-SAN
General
Storage
Network
iSCSI vNICs
Boot Order
Policies
Events
FSM
Servers
Service Profiles
root
Service Profile Templates
root
Service Template C220M3-VDI-SAN
Sub-Organizations
Policies
root
Adapter Policies
BIOS Defaults
BIOS Policies
Boot Policies
Host Firmware Packages
IPMI Access Profiles
Local Disk Config Policies
Maintenance Policies
Management Firmware Packages
Power Control Policies
Scrub Policies
Serial over LAN Policies
Server Pool Policies
Server Pool Policy Qualifications
Threshold Policies
iSCSI Authentication Profiles
vNIC/vHBA Placement Policies
Sub-Organizations
Actions
Create Service Profiles From Template
Create a Clone
Disassociate Template
Associate with Server Pool
Change Maintenance Policy
Change UUID
Change Management IP Address

Filter: All

Properties

Name: C220M3-VDI-SAN
Description: C220M3 Boot from SAN Template for VDI Hosts
UUID: Derived from pool (VDI-UUID-Pool)
Power State: Up
Type: Updating Template

Associated Server Pool

Server Pool: VDI-CSeries
Server Pool Qualification: C220M3-Qual
Restrict Migration: No

Maintenance Policy

Name: default
Maintenance Policy Instance: org-root/maint-default
Description:
Reboot Policy: Immediate

Management IP Address

Management IP Address Policy: None

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

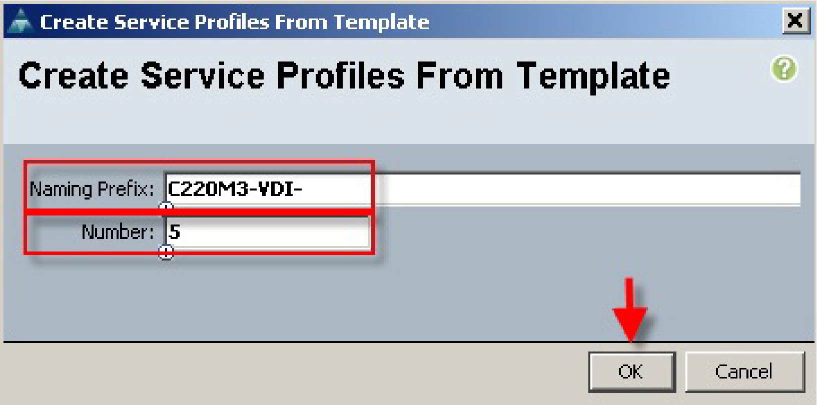
Create Service Profiles From Template
Create Service Profiles From Template
OK
Cancel

Naming Prefix: C220M3-VDI-
Number: 5

Cisco UCS Manager created the requisite number of profiles and because of the Associated Server Pool and Server Pool Qualification policy, the Cisco UCS C220 M3 blades in the test environment began automatically associating with the proper Service Profile.

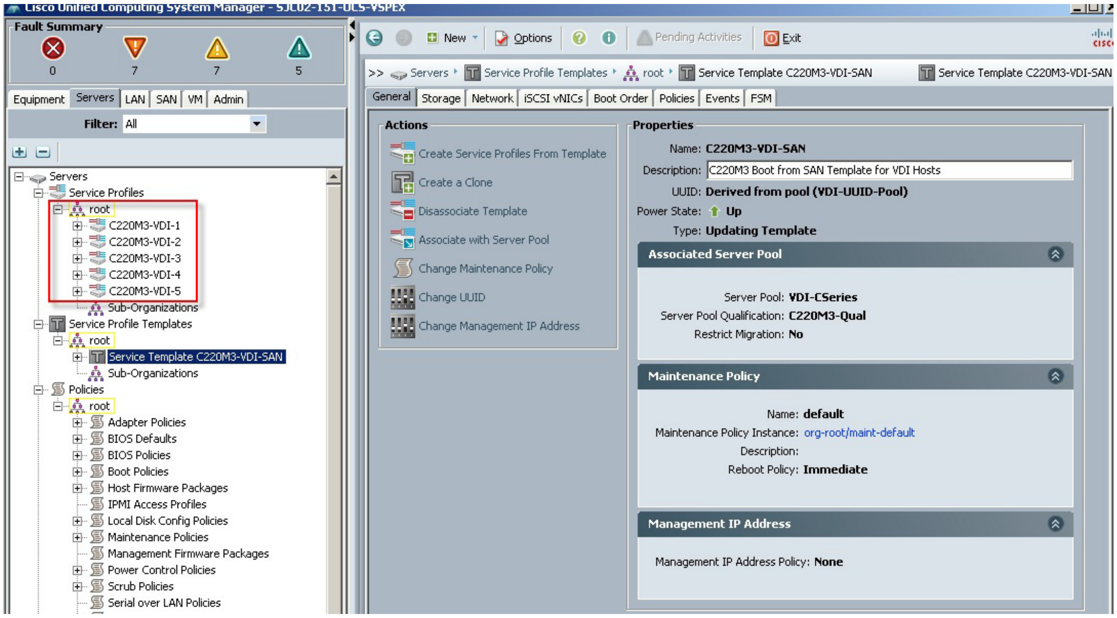

Properties

Name: **C220M3-VDI-SAN**
Description: C220M3 Boot from SAN Template for VDI Hosts
UUID: **Derived from pool (VDI-UUID-Pool)**
Power State: **Up**
Type: **Updating Template**

Associated Server Pool

Server Pool: **VDI-CSeries**
Server Pool Qualification: **C220M3-Qual**
Restrict Migration: **No**

Maintenance Policy

Name: **default**
Maintenance Policy Instance: org-root/maint-default
Description:
Reboot Policy: **Immediate**

Management IP Address

Management IP Address Policy: **None**

80. Each of the Cisco UCS C220 M3 servers were automatically assigned one of the service profiles based on its pool and pool qualification policy.

Note We verified that each server had a Service Profile.

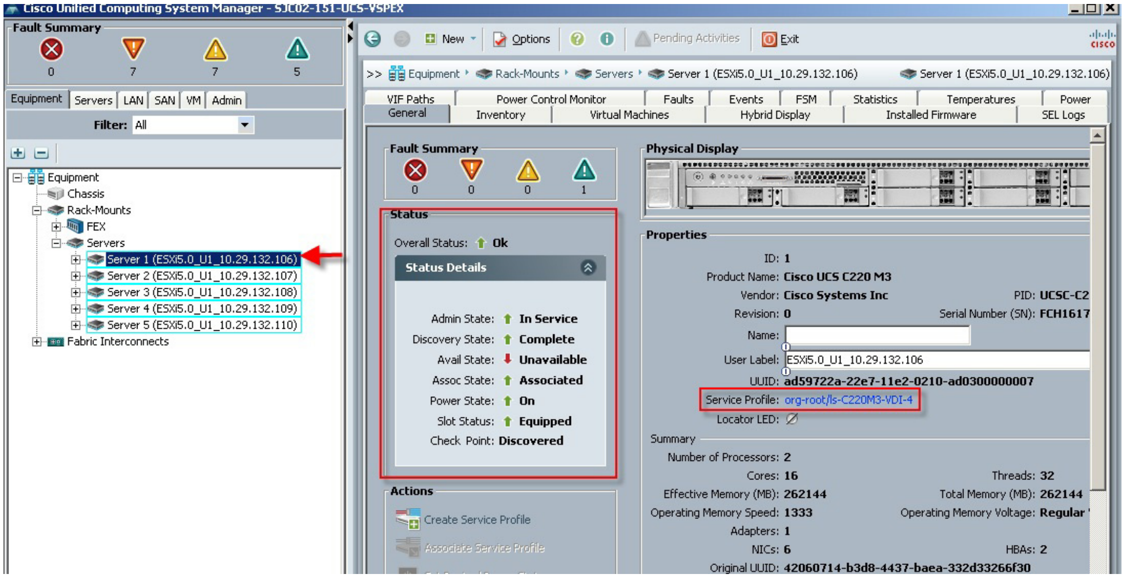

Status

Overall Status: **Ok**

Status Details

Admin State: **In Service**
Discovery State: **Complete**
Avail State: **Unavailable**
Assoc State: **Associated**
Power State: **On**
Slot Status: **Equipped**
Check Point: **Discovered**

Properties

ID: **1**
Product Name: **Cisco UCS C220 M3**
Vendor: **Cisco Systems Inc** PID: **UCSC-C2**
Revision: **0** Serial Number (SN): **FCH1617**
Name:
User Label: ESXi5.0_U1_10.29.132.106
UUID: **ad59722a-22e7-11e2-0210-ad0300000007**
Service Profile: org-root/ls-C220M3-VDI-4
Locator LED:
Summary
Number of Processors: **2**
Cores: **16** Threads: **32**
Effective Memory (MB): **262144** Total Memory (MB): **262144**
Operating Memory Speed: **1333** Operating Memory Voltage: **Regular**
Adapters: **1**
NICs: **6** HBAs: **2**
Original UUID: **42060714-b3d8-4437-baea-332d33266f30**

The Cisco UCS Blade Servers are ready for 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 9216.
- Via UCS we 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 System Class Configuration

Cisco UCS defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

Table 3 Name Table Map between Cisco Unified Computing System 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

In this study, we set FC to a weight of 6 and Best Effort to a weight of 4

Steps to Enable Qos on the Cisco Unified Computing System

For the Managed FC variant of this study, we utilized three Cisco UCS QoS System Classes to priorities three 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	100 (VDI)
Gold	Eth0, eth1	132 (ESXi_Management)
Bronze	Eth4, eth5	51 (vMotion)

1. Configure Platinum, Gold, and Bronze System Classes by checking the enabled box.

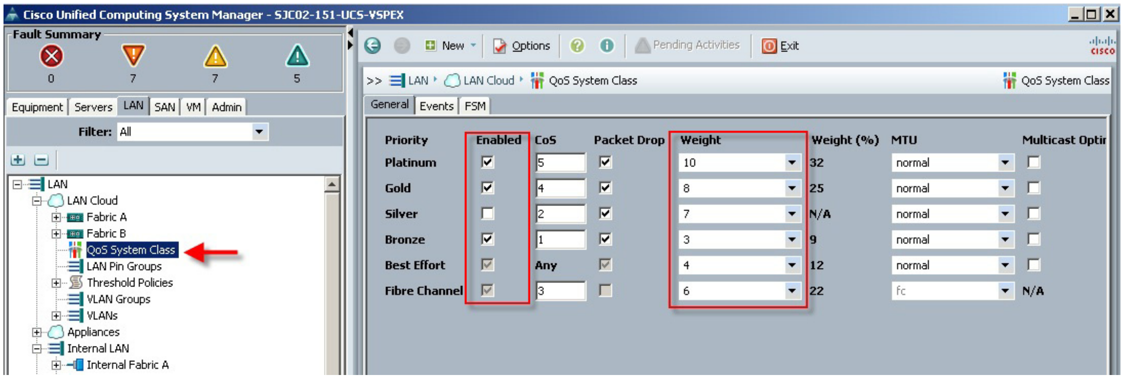

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optir
Platinum	☑	5	☑	10	32	normal	□
Gold	☑	4	☑	8	25	normal	□
Silver	□	2	☑	7	N/A	normal	□
Bronze	☑	1	☑	3	9	normal	□
Best Effort	☑	Any	☑	4	12	normal	□
Fibre Channel	☑	3	□	6	22	fc	N/A

2. In.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.

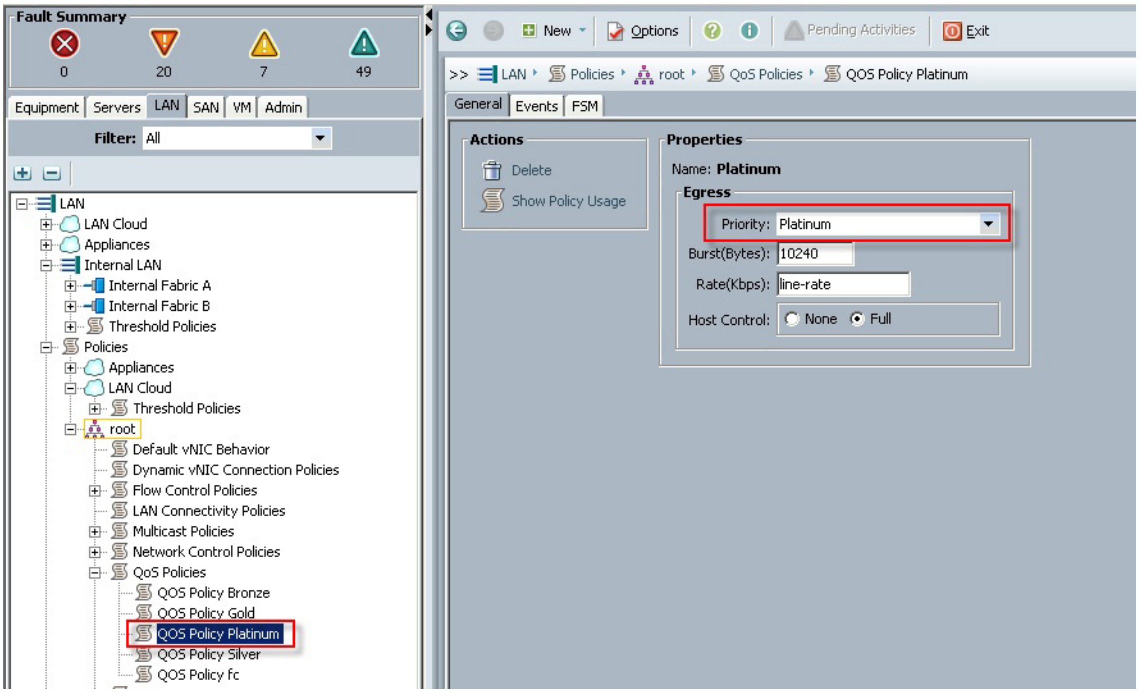

Properties

Name: **Platinum**

Egress

Priority: Platinum

Burst(Bytes): 10240

Rate(Kbps): line-rate

Host Control: ○ None ● Full

3. Include the corresponding QoS Policy into each vNIC template using the QoS policy drop-down list, using the QoS Priority to vNIC and VLAN Mapping table above.

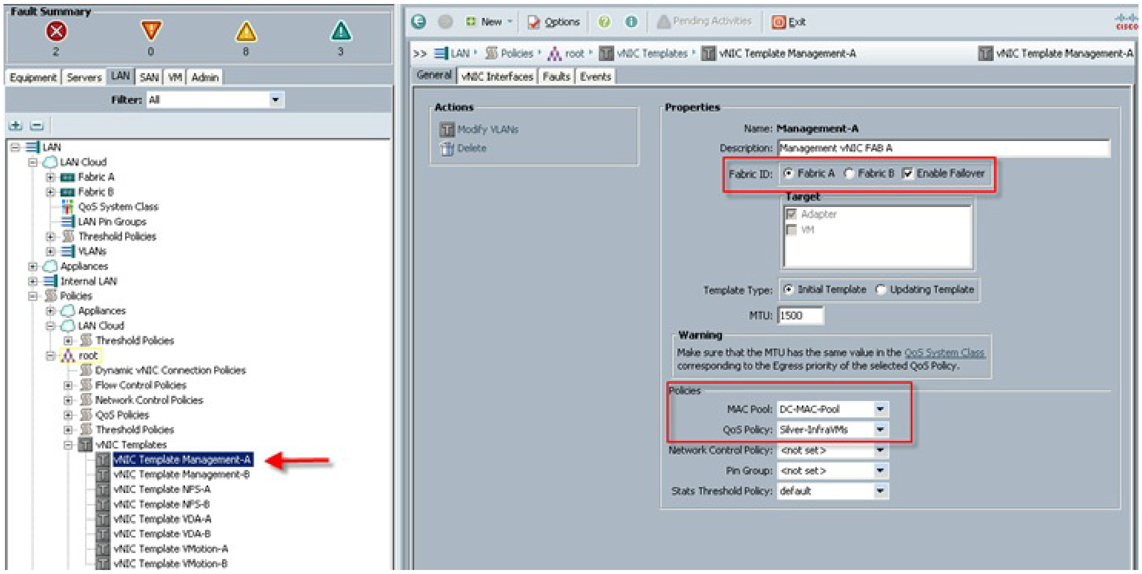

Properties

Name: Management-A

Description: Management vNIC FAB A

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

Target

☑ Adapter

□ VM

Template Type: ● Initial Template ○ Updating Template

MTU: 1500

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

Policies

MAC Pool: DC-MAC-Pool

QoS Policy: Silver-InfraVMs

Network Control Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

This is a unique value proposition for Cisco UCS with respect to end-to-end QOS.

LAN Configuration

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

Nexus 5548UP and VNX5300 Connectivity (Unmanaged NFS Variant)

The access layer LAN configuration consists of a pair of Cisco Nexus 5548UPs (N5Ks,) a family member of our low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches for our VDI deployment uplinked to the Customers existing L3 network.

In the Unmanaged NFS Variant, the Cisco UCS C220 M3 servers are managed in standalone mode, requiring a management/infrastructure network connection and a separate high speed data connection to the EMC VNX5300 storage system.

Five Cisco UCS C220 M3 Rack-Mount servers are connected via their VIC1225 to 10Gb ports on a pair of N5Ks to the EMC VNX NFS storage. One or both of each servers' integrated 1 Gb ethernet ports is/are connected to the upstream or top of rack L3 switch for management and all infrastructure communications.

Note The upstream configuration is beyond the scope of this document; there are some good reference documents that talk 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 14 Unmanaged NFS Variant Ethernet Network Configuration with Cisco Nexus 5500 Series Switches

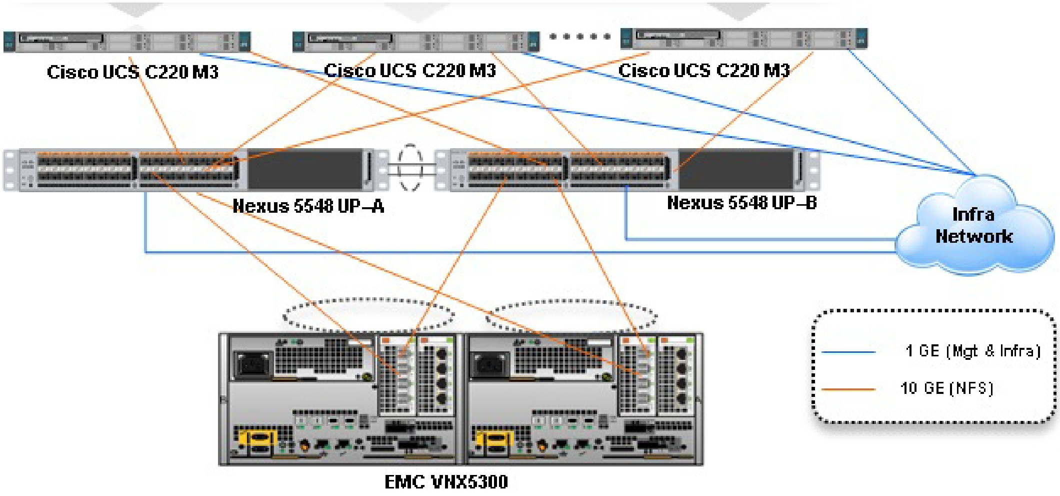

Cisco UCS Fabric Interconnect 6248UP and VNX5300 Connectivity (Managed FC Variant)

In the Managed FC Variant, the UCS C220 M3 servers are connected to a pair of Cisco UCS 6248UP Fabric Interconnects (FIs) via two Nexus 2232PP Fabric Extenders. The VNX5300 is connected to the FI FC Storage Ports directly via Fiber Channel. Fibre Channel zoning is done on the FIs. The FIs are uplinked to the Customer top of rack L3 switch.

In this configuration, called Managed Single Wire Cluster Setup, only the two 10 GbE ports on the VIC1225 CNA are used for all communications. In addition, Cisco UCS Manager 2.1 manages the configuration of the servers through Cisco UCS Service Profiles.

Figure 15 Managed FC Variant Network Configuration

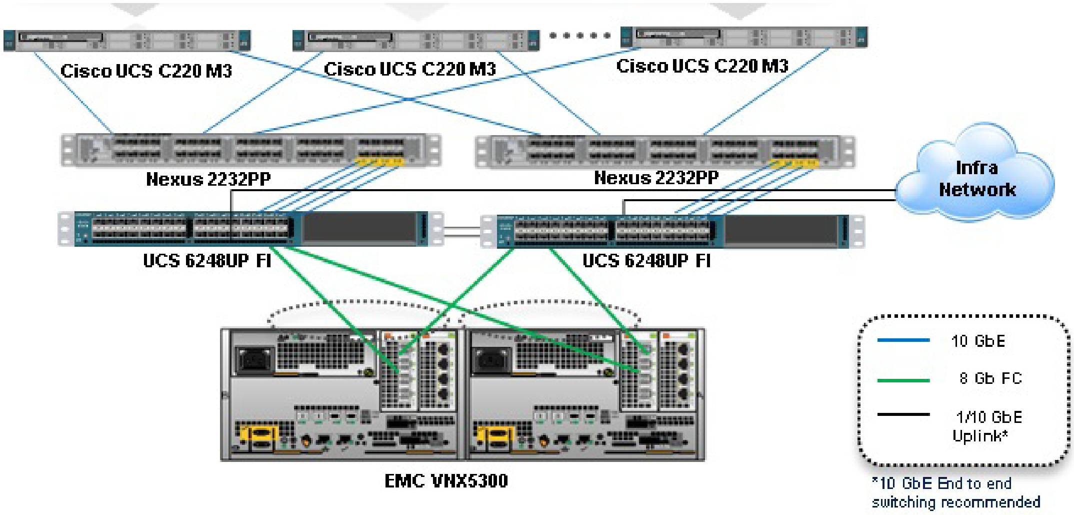

SAN Configuration

For the two variants of this study, different equipment was used to connect to the VNX5300 storage outlined in Section 6.3 above. Only the Managed FC Variant supports booting from the VNX5300. This section describes the SAN Configuration supporting boot from SAN.

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 Cisco 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 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:

- SAN zone configuration through the Cisco UCS Manager Service Profile Template creation process
- Storage array host initiator configuration
- 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 Zone Configuration on Cisco UCS Manager

The Cisco UCS Local Zoning feature requires that the UCS Fabric Interconnects be configured in FC Switching Mode rather than the default of FC End Host Mode. When that task is completed, the properties of the VSAN used in the deployment must have the FC Zoning Setting set to Enabled.

These tasks are outlined in section Base Cisco UCS System Configuration above in step 15.

When enabled, the Cisco UCS Manager Service Profile Template creation wizard will provide the steps necessary to complete the Fibre Channel Zoning for the project.

These tasks are outlined in section Base Cisco UCS System Configuration above in step 16.

It is possible to perform FC zoning on the Fabric Interconnects from the command line. Using that method is not covered in this paper. Please refer to the Cisco UCS Manager CLI Command Reference, Release 2.1 that can be found at:

http://www.cisco.com/en/US/partner/docs/unified_computing/ucs/sw/cli/command/reference/2.1/b_command_reference_2.1_chapter_0100.html

Configuring Boot from SAN on EMC VNX

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

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

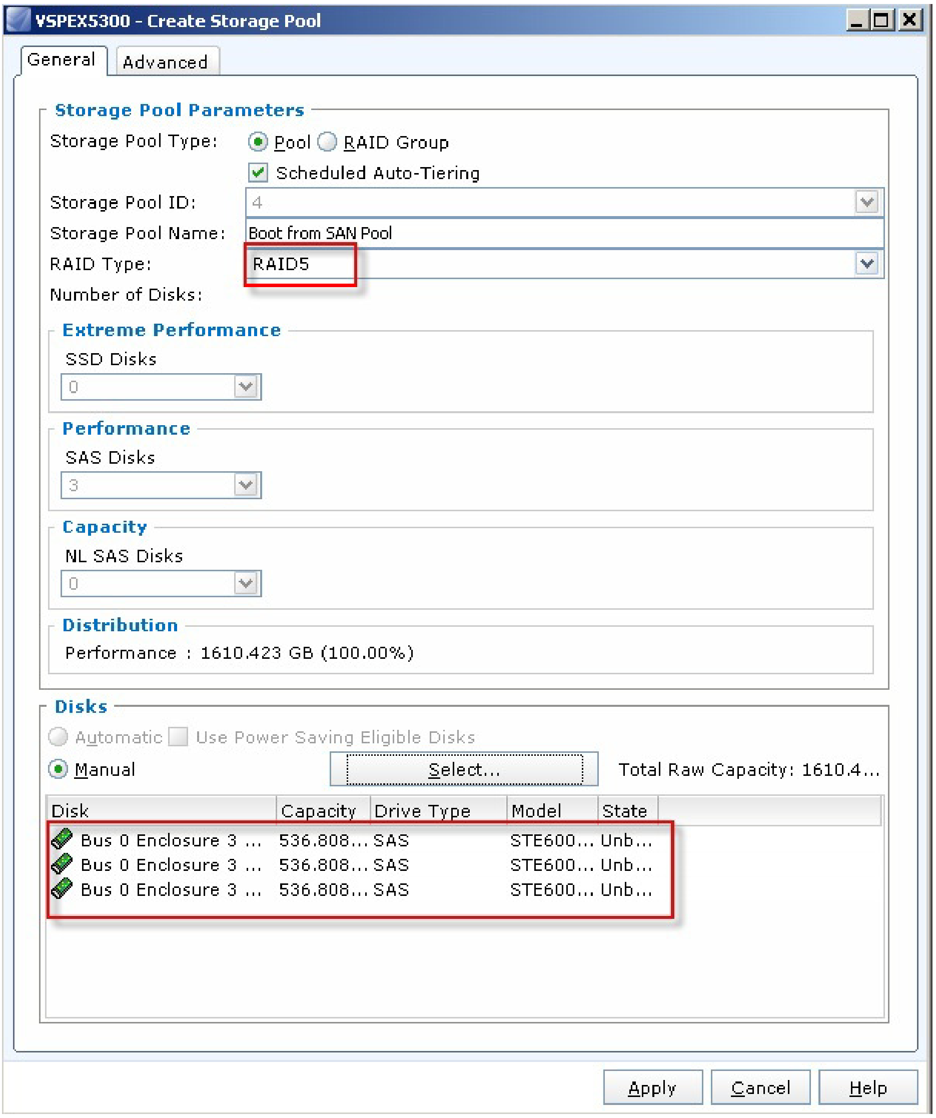

VSPEX5300 - Create Storage Pool

General | Advanced

Storage Pool Parameters

Storage Pool Type: ● Pool ○ RAID Group

☑ Scheduled Auto-Tiering

Storage Pool ID: 4

Storage Pool Name: Boot from SAN Pool

RAID Type: RAID5

Number of Disks:

Extreme Performance

SSD Disks: 0

Performance

SAS Disks: 3

Capacity

NL SAS Disks: 0

Distribution

Performance : 1610.423 GB (100.00%)

Disks

○ Automatic □ Use Power Saving Eligible Disks

● Manual

Total Raw Capacity: 1610.4...

Disk	Capacity	Drive Type	Model	State
Bus 0 Enclosure 3 ...	536.808...	SAS	STE600...	Unb...
Bus 0 Enclosure 3 ...	536.808...	SAS	STE600...	Unb...
Bus 0 Enclosure 3 ...	536.808...	SAS	STE600...	Unb...

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

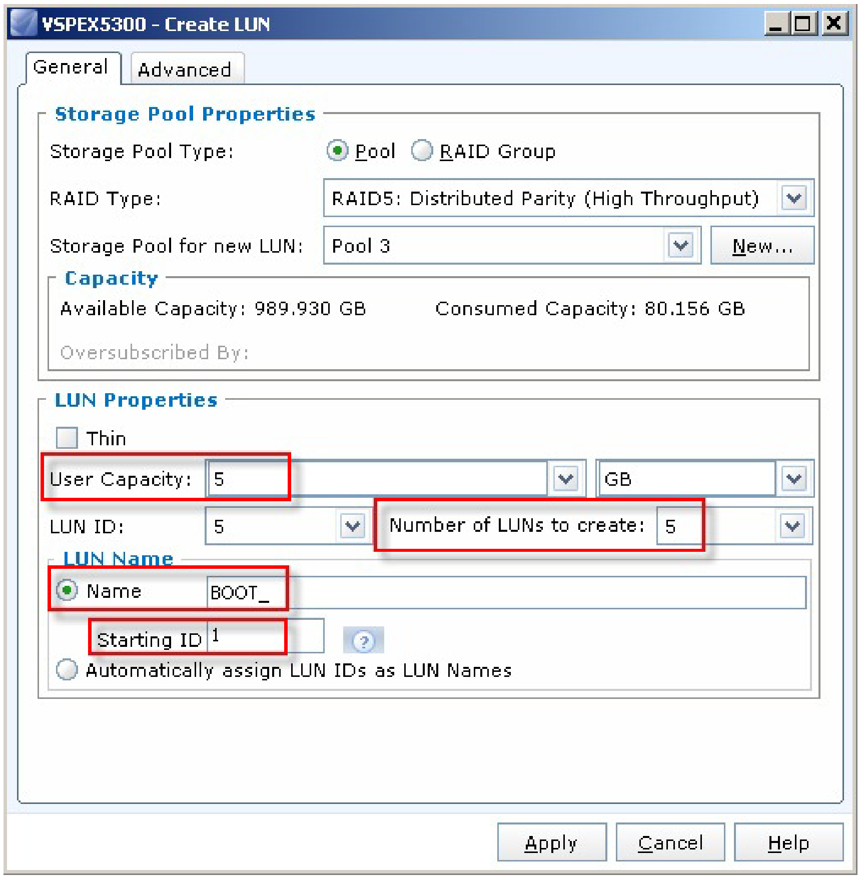

Storage Pool Properties

Storage Pool Type: ● Pool ○ RAID Group

RAID Type: RAID5: Distributed Parity (High Throughput)

Storage Pool for new LUN: Pool 3 New...

Capacity

Available Capacity: 989.930 GB Consumed Capacity: 80.156 GB

Oversubscribed By:

LUN Properties

□ Thin

User Capacity: 5 GB

LUN ID: 5 Number of LUNs to create: 5

LUN Name

● Name BOOT_

Starting ID 1

○ Automatically assign LUN IDs as LUN Names

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

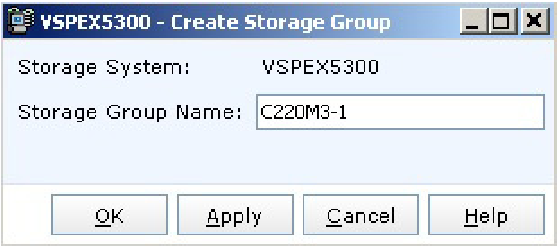

Storage System: VSPEX5300

Storage Group Name: C220M3-1

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.

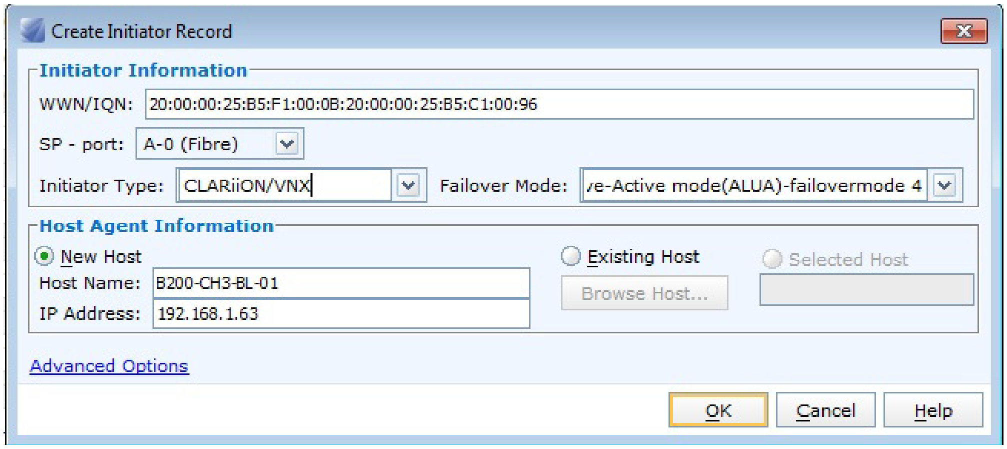

Create Initiator Record

Initiator Information

WWN/IQN: 20:00:00:25:B5:F1:00:0B:20:00:00:25:B5:C1:00:96

SP - port: A-0 (Fibre)

Initiator Type: CLARiiON/VNX

Failover Mode: ve-Active mode(ALUA)-failovermode 4

Host Agent Information

● New Host ○ Existing Host ○ Selected Host

Host Name: B200-CH3-BL-01

IP Address: 192.168.1.63

5. Assign each registered host to a separate storage group as shown below.

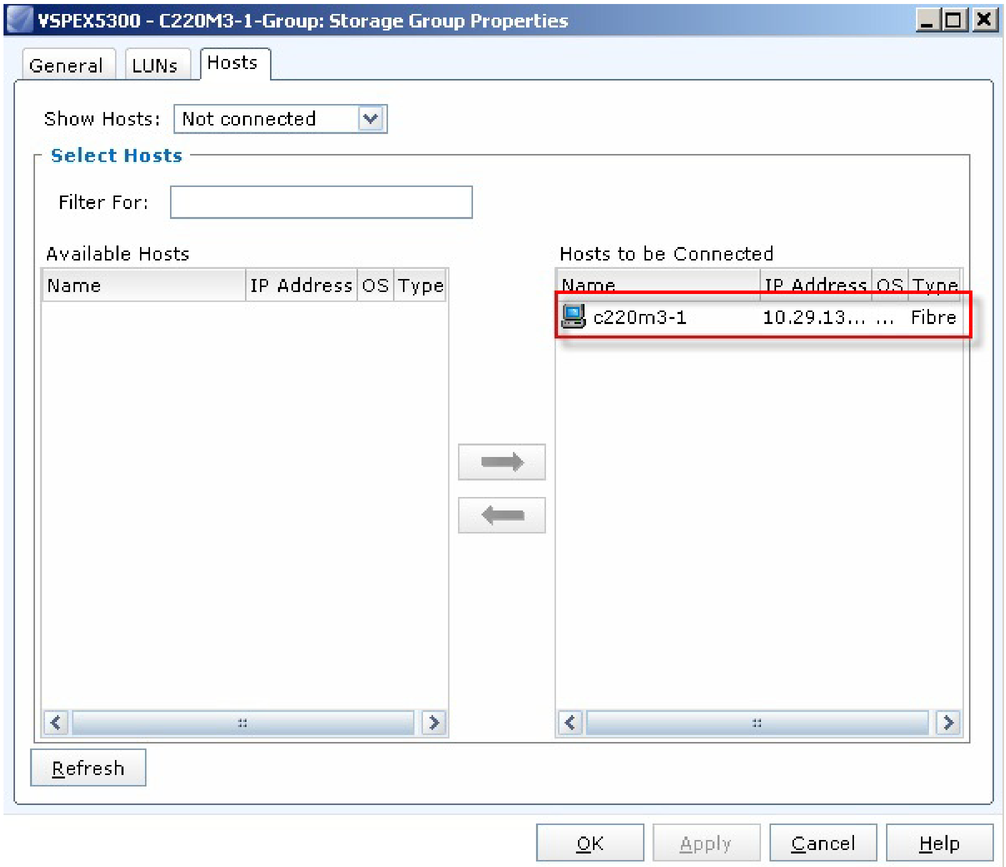

VSPEX5300 - C220M3-1-Group: Storage Group Properties

Show Hosts: Not connected

Select Hosts

Filter For:

Available Hosts

Name	IP Address	OS	Type

Hosts to be Connected

Name	IP Address	OS	Type
c220m3-1	10.29.13...	...	Fibre

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.

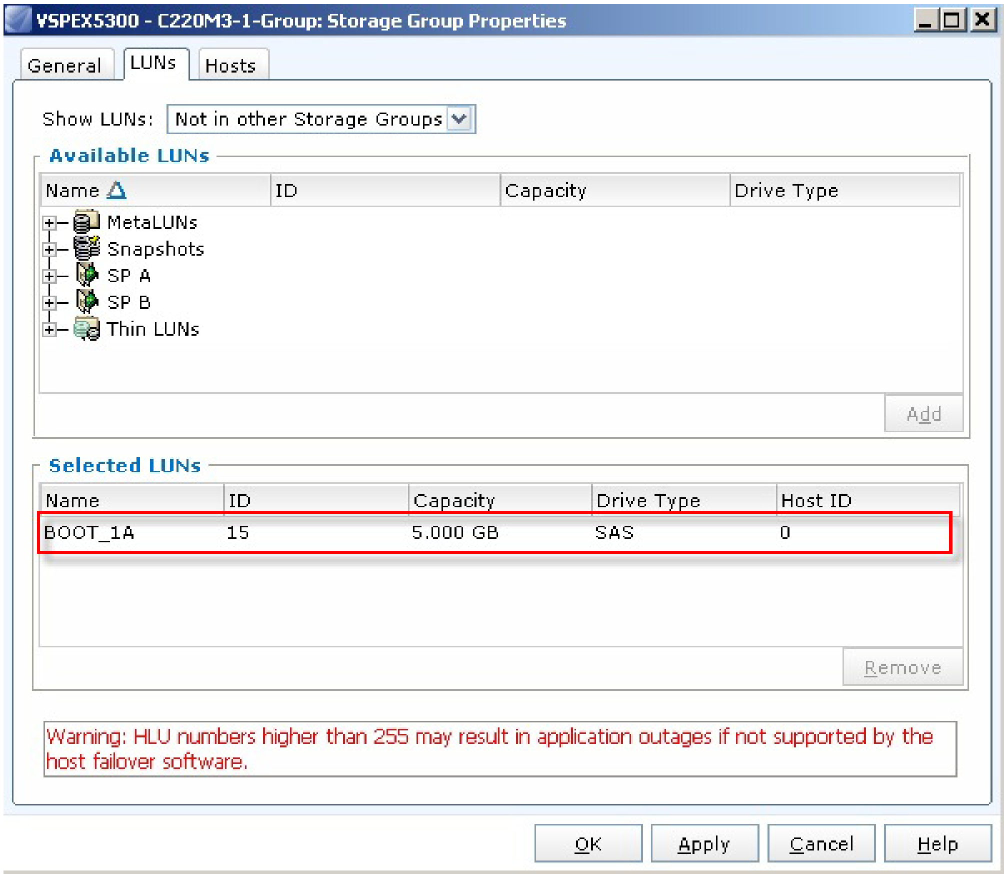

Show LUNs: Not in other Storage Groups

Name	ID	Capacity	Drive Type	Host ID
BOOT_1A	15	5.000 GB	SAS	0

Warning: HLU numbers higher than 255 may result in application outages if not supported by the host failover software.

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

The configuration of the Boot from SAN policy in Cisco UCS Manager is covered in section Base Cisco UCS System Configuration above in step 13.

The policy created is incorporate in the Service Profile Template in section Base Cisco UCS System Configuration above in step 16o.

EMC VNX5300 Storage Configuration

The Managed FC Variant and Unmanaged NFS Variant have the same configuration to the pool level on the VNX5300. The configuration varies from there as described below.

Physical and Logical Storage Layout for Managed FC and Unmanaged NFS Variants

The figure below shows the physical storage layout of the disks in the reference architecture. This configuration accommodates up to 600 virtual desktops.

Figure 16 Physical Storage Layout

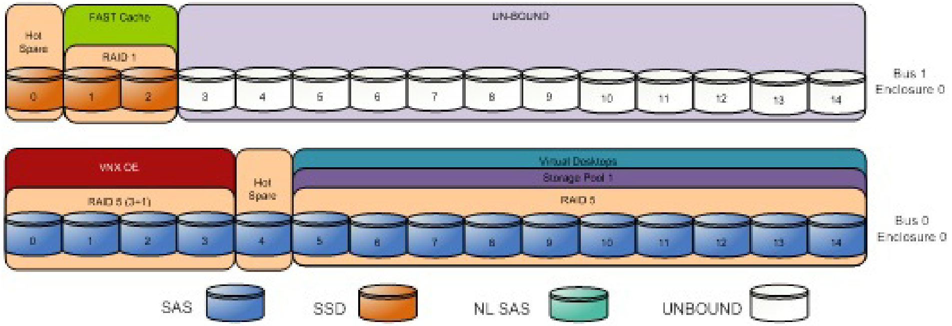

The above storage layout is used for the following configurations:

Managed Fibre Channel Variant

- Four SAS disks are used for the VNX OE.
- One SAS disk is a hot spare for SAS disks.
- One SSD Disk is hot spare for SSD drives.
- Two 100GB Flash drives are used for EMC VNX FAST Cache. See the "EMC FAST Cache in Practice" section below to follow the FAST Cache configuration best practices.
- 10 600GB SAS disks on a single RAID 5 storage pool with Fast Cache enabled for virtual desktop write cache drives.
- Four LUNs of 500GB each and one 50GB LUN are carve out to the block pool to present to the ESXi servers as four VMFS datastores for linked clones and one VMFS datastore for View replicas.
- The Managed Fibre Channel Variant was validated with 600 virtual desktops.

Unmanaged NFS Variant

- Four SAS disks are used for the VNX OE.
- One SAS disk is a hot spare for SAS disks.
- One SSD Disk is hot spare for SSD drives.
- Two 100GB Flash drives are used for EMC VNX FAST Cache. See the "EMC FAST Cache in Practice" section below to follow the FAST Cache configuration best practices.
- 10 600GB SAS disks on a single RAID 5 storage pool with Fast Cache enabled for virtual desktop write cache drives.
- Ten LUNs of 200GB each are carve out to the NAS pool to present to the Data Movers as dvols that belong to a system defined NAS pool.
- Four file systems of 500 GB each are carved out of the NAS pool to present to the ESXi servers as four NFS datastores

1. To enable an NFS performance fix for VNX File that significantly reduces NFS write latency, the file systems must be mounted on the Data Mover using the Direct Writes mode. The Set Advanced Options check box must be selected to enable the Direct Writes check box.

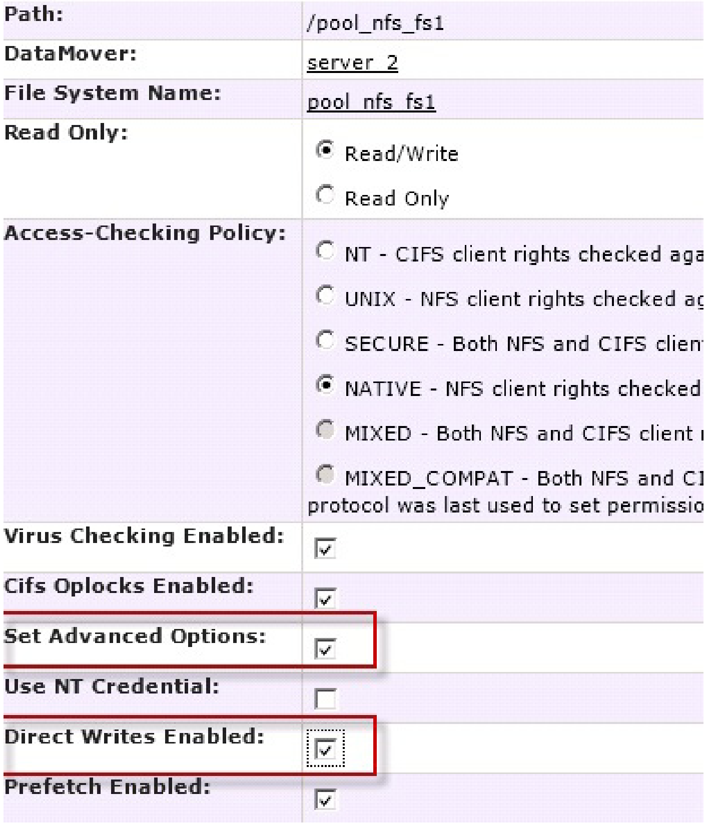

Path:	/pool_nfs_fs1
DataMover:	server_2
File System Name:	pool_nfs_fs1
Read Only:	● Read/Write ○ Read Only
Access-Checking Policy:	○ NT - CIFS client rights checked aga ○ UNIX - NFS client rights checked ag ○ SECURE - Both NFS and CIFS clien ● NATIVE - NFS client rights checked ○ MIXED - Both NFS and CIFS client r ○ MIXED_COMPAT - Both NFS and CI protocol was last used to set permissio
Virus Checking Enabled:	☑
Cifs Oplocks Enabled:	☑
Set Advanced Options:	☑
Use NT Credential:	☐
Direct Writes Enabled:	☑
Prefetch Enabled:	☑

2. Export the file systems using NFS, and give root access to ESXi servers.

3. In Unisphere, click Settings > Data Mover Parameters to make changes to the Data Mover configuration. Click the drop down menu to the right of Set Parameters and change the setting to "All Parameters". Scroll down to the nthreads parameter as and click Properties to update the setting. The default number of threads dedicated to serve NFS requests is 384 per Data Mover on VNX. Since up to 600 desktop connections are required in this solution, it is recommended to increase the number of active NFS threads to a maximum of 1024 on each Data Mover.

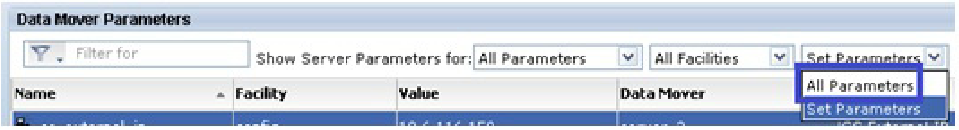

Data Mover Parameters

Filter for | Show Server Parameters for: All Parameters | All Facilities | Set Parameters

All Parameters
Set Parameters

Name	Facility	Value	Data Mover

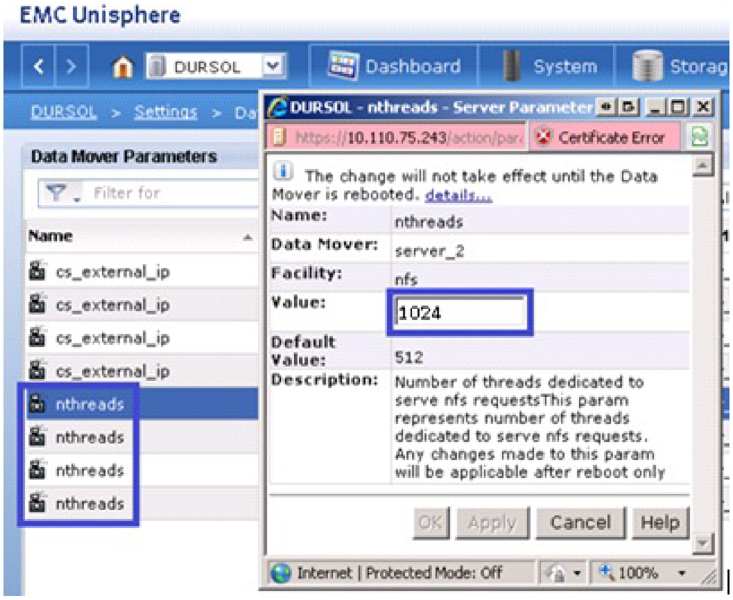

The change will not take effect until the Data Mover is rebooted. details...

Name: nthreads

Data Mover: server_2

Facility: nfs

Value: 1024

Default Value: 512

Description: Number of threads dedicated to serve nfs requestsThis param represents number of threads dedicated to serve nfs requests. Any changes made to this param will be applicable after reboot only

Note The Unmanaged NFS Variant was validated with 600 virtual desktops.

Configure SP Cache

To help ensure optimal performance of the solution the array SP cache settings must be modified. Complete the following steps in Unisphere to configure the SP cache:

1. Click the System tab.
2. From the System home page, click the Manage Cache link on the right hand side of the screen to open the Storage System Properties window.
3. In the Storage System Properties window select the SP Cache tab.
4. Set the Low Watermark setting to 70 and the High Watermark setting to 90 as shown in the red highlighted area.
5. Click OK to implement the changes and close the Storage System Properties Window.

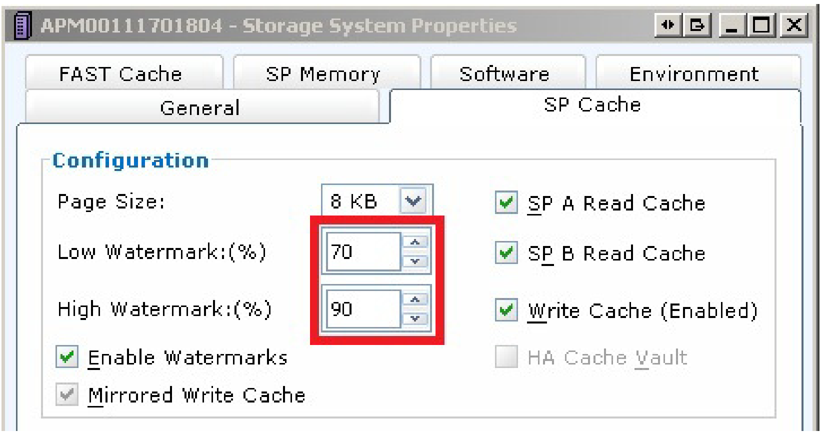

Page Size: 8 KB

Low Watermark:(%) 70

High Watermark:(%) 90

☑ Enable Watermarks

☑ Mirrored Write Cache

☑ SP A Read Cache

☑ SP B Read Cache

☑ Write Cache (Enabled)

☐ HA Cache Vault

EMC Storage Configuration for VMware ESXi 5.1 Infrastructure Servers

If storage required for infrastructure virtual machines (that is, SQL server, domain controller, vCenter server, and/or View controllers) does not exist in the production environment already and the optional user data disk pack has been purchased, configure a NFS file system or one or more FC LUNS on VNX to be used as a NFS datastore or VMFS datastore in which the infrastructure virtual machines reside.

In this study we used 5 600GB SAS drives for the Infrastructure Pool in a RAID 5 array to provision 2 500GB VMFS LUNS for our infrastructure virtual machines and files. Fast Cache was disabled on the Infrastructure Pool.

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

Installing and Configuring ESXi 5.1

In this study, we used Fibre Channel storage to boot the hosts from LUNs on the VNX5300 storage system. Prior to installing the operating system, storage groups are created, assigning specific boot LUNs to individual hosts. (See Section 7.4.4 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 from the UCS Manager (UCSM) KVM console using a ESXi 5.1 ISO image downloaded from the VMware site.

The IP address, hostname, and NTP server were configured using Direct Console ESXi Interface accessed from UCSM KVM console.

Refer to the following VMware documentation for configuring network settings help:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc_50/GUID-26F3BC88-DAD8-43E7-9EA0-160054954506.html

Install and Configure vCenter

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

Table 7 *VMware vCenter Server*

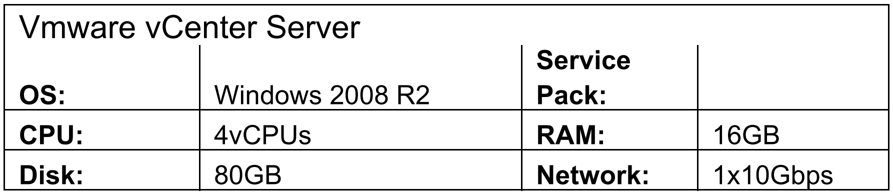

Vmware vCenter Server			
OS:	Windows 2008 R2	**Service Pack:**	
CPU:	4vCPUs	**RAM:**	16GB
Disk:	80GB	**Network:**	1x10Gbps

To support vCenter instance, one Microsoft SQL Server 2008 R2 was created to host vCenter database.

If the Customer wants to utilize fault tolerance at the SQL Server level, refer to Microsoft documentation on configuring SQL Server clusters.

http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx

http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx

To install and configure vCenter, use the following steps:

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 Ensure to use FQDN's for everything.

3. Create Active Directory user account and call it vpxuser. (This user account will be used for XD 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).
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 Licenses

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

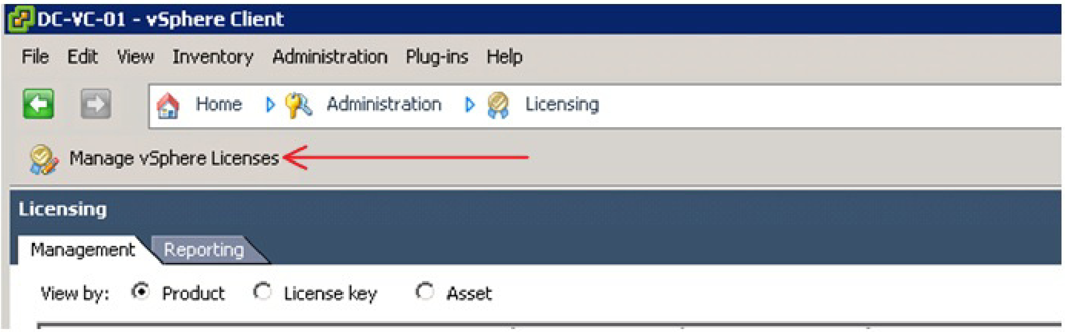
DC-VC-01 - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home
Administration
Licensing
Manage vSphere Licenses
Licensing
Management
Reporting
View by: ● Product ○ License key ○ Asset

4. Add License keys for vCenter and Hosts.

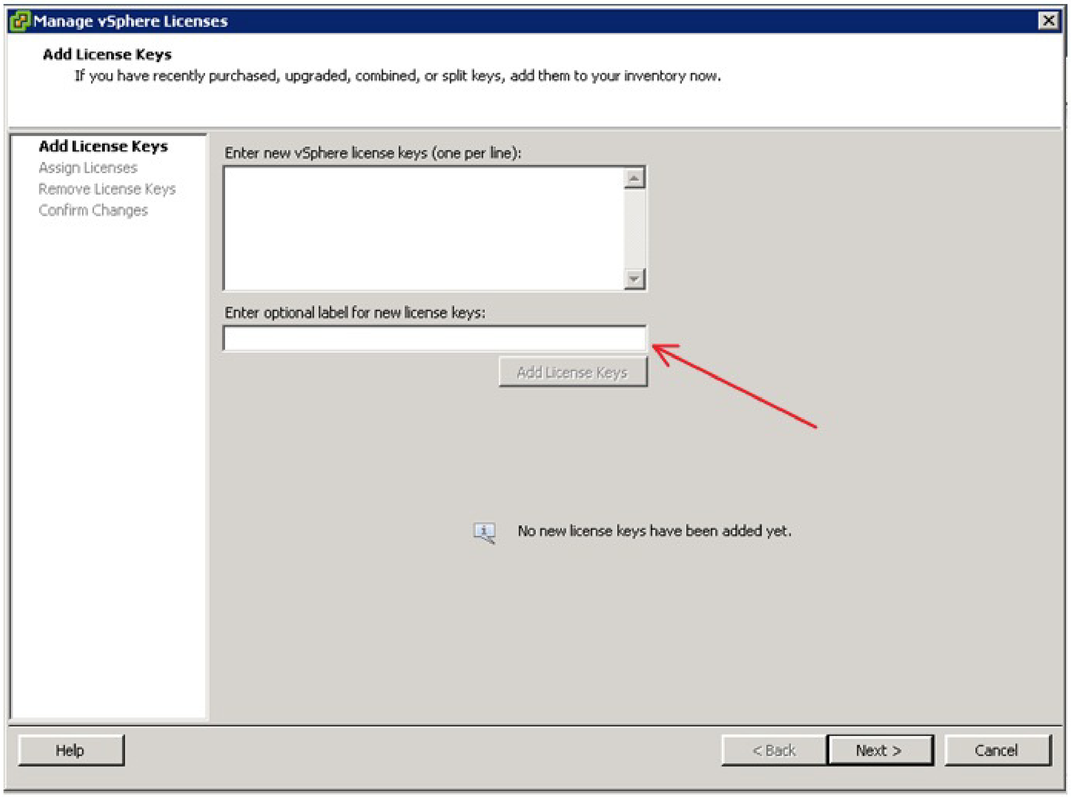

Add License Keys

If you have recently purchased, upgraded, combined, or split keys, add them to your inventory now.

Enter new vSphere license keys (one per line):

Enter optional label for new license keys:

No new license keys have been added yet.

5. Use the top box to enter the license key and lower box to add an optional label.

ESXi 5.1 Cluster Configuration

To accommodate maximum recommendations for ESXi 5 clustering, we created two ESXi 5 clusters described below:

- VDI-v500
- Infra-CL

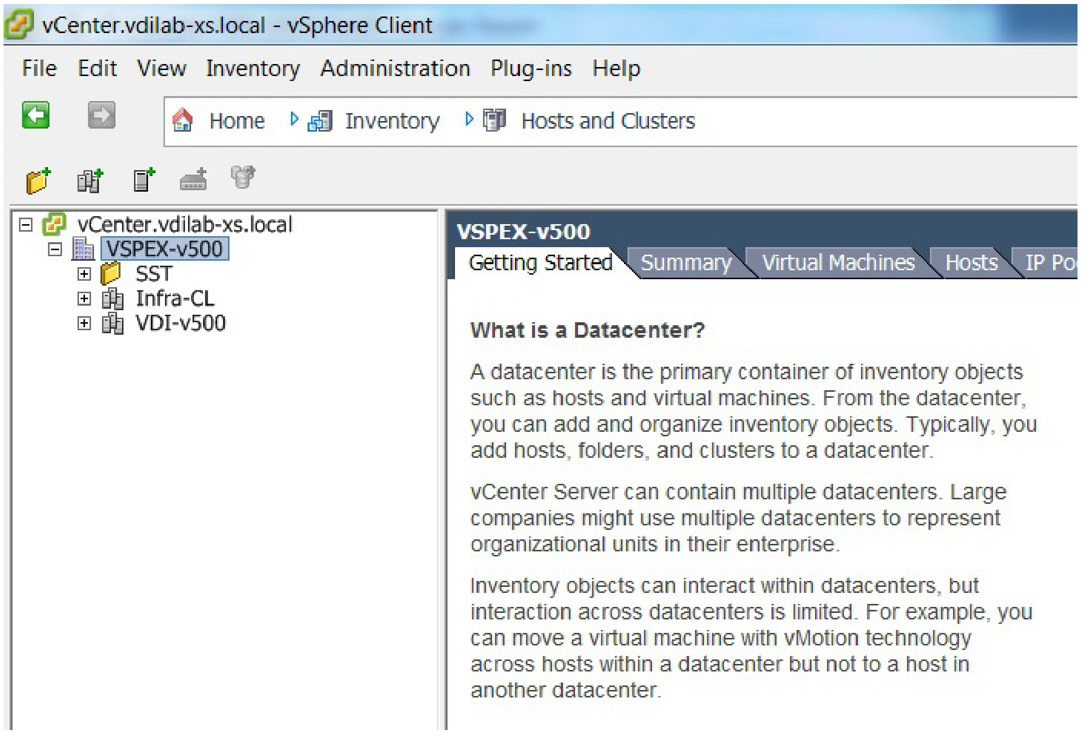

What is a Datacenter?

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

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

Inventory objects can interact within datacenters, but interaction across datacenters is limited. For example, you can move a virtual machine with vMotion technology across hosts within a datacenter but not to a host in another datacenter.

Infra-CL Infrastructure Cluster

The Infra-CL cluster was used to host all of the virtualized servers within the VDA Infrastructure, including Active Directory servers providing authentication, DNS, DHCP and group policy services, a MS SQL 2008 R2 server for vCenter View Composer and View Connection server databases, a pair of XenDesktop 5.6 FP1 servers, a file server for MS Roaming Profiles, a virtual file server for Login VSI files, and twenty Server 2008 R2 SP1 virtual machines as Login VSI Launchers.

Note Two physical Cisco UCS C220 M3 hosts were used in this cluster.

Standard ESXi vSwitches were used and configure as follows:

- Managed FC Variant (All 10 Gb)
  - Default vmkernel vSwitch for VMware Management
  - Vmkernel vSwitch for VMotion
  - Standard vSwitch for VDI traffic
- Unmanaged NFS Variant
  - Default vmkernel vSwitch for VMware Management and VDI traffic (1 Gb)
  - Vmkernel vSwitch for vMotion (10 Gb)
  - Vmkernel vSwitch for NFS storage (10 Gb)

Figure 17 Infra-CL ESXi host network configuration

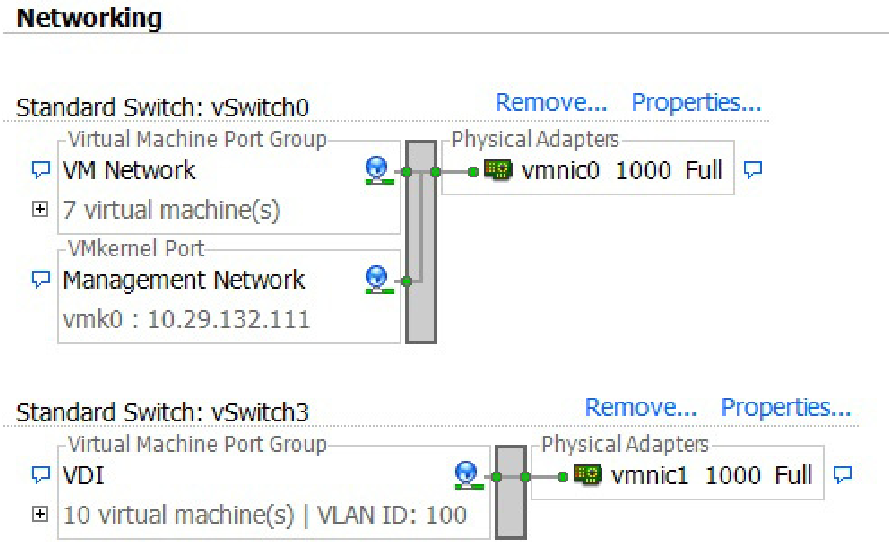

Networking

Standard Switch: vSwitch0

Virtual Machine Port Group

VM Network

7 virtual machine(s)

VMkernel Port

Management Network

vmk0 : 10.29.132.111

Physical Adapters

vmnic0 1000 Full

Standard Switch: vSwitch3

Virtual Machine Port Group

VDI

10 virtual machine(s) | VLAN ID: 100

Physical Adapters

vmnic1 1000 Full

It is not necessary to create an Infrastructure Cluster. If the existing Customer ESXi deployment has capacity, the required virtual machines can be added there.

Note Separate LUNS or NFS file systems were created for the Infrastructure Cluster hosts.

VDI-v500 Virtual Desktop Cluster

The VDI-v500 desktop cluster was used to host 600 of the Windows 7 SP1 32-bit virtual desktops.

Note Five physical Cisco UCS C220 M3 hosts were used in this cluster.

Standard ESXi vSwitches were used and configure as follows:

- Managed FC Variant (All 10 Gb)
  - Default vmkernel vSwitch for VMware Management
  - Vmkernel vSwitch for VMotion
  - Standard vSwitch for VDI traffic
- Unmanaged NFS Variant
  - Default vmkernel vSwitch for VMware Management and VDI traffic (1 Gb)
  - Vmkernel vSwitch for vMotion (10 Gb)
  - Vmkernel vSwitch for NFS storage (10 Gb)

Note It is recommended that you create a separate cluster for VDI workloads. It will provide a management container that will allow growth for the VDI use.

Figure 18 VDI-v500 ESXi Host Network Configuration

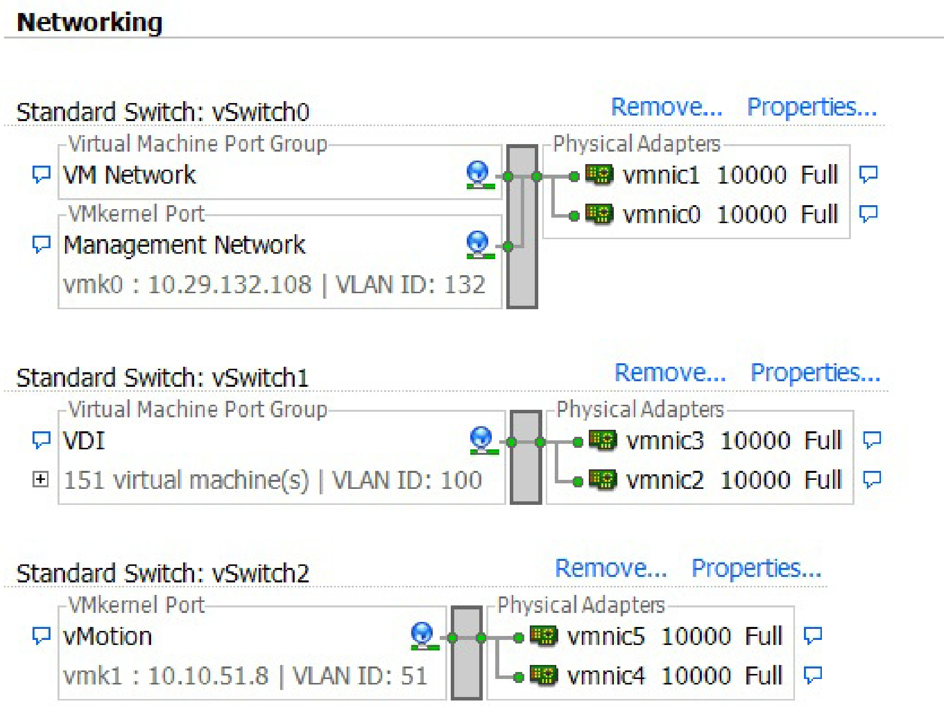

Standard Switch: vSwitch0

Virtual Machine Port Group
VM Network

VMkernel Port
Management Network
vmk0 : 10.29.132.108 | VLAN ID: 132

Physical Adapters
vmnic1 10000 Full
vmnic0 10000 Full

Standard Switch: vSwitch1

Virtual Machine Port Group
VDI
151 virtual machine(s) | VLAN ID: 100

Physical Adapters
vmnic3 10000 Full
vmnic2 10000 Full

Standard Switch: vSwitch2

VMkernel Port
vMotion
vmk1 : 10.10.51.8 | VLAN ID: 51

Physical Adapters
vmnic5 10000 Full
vmnic4 10000 Full

Installing and configuring VMware View 5.1

Building Out a VMware View 5.1 Environment Requires the Installation of the Following Components:

- View Connection Server
- View Replica Server
- View Administrator
- View Composer
- View Security Server (Optional)
- View Transfer Server (Optional)

This section outlines the tasks required to build the View 5.1 environment used in this study. Refer to the VMware View Installation guide for View 5.1 for more details.

Prerequisites

The following is a list of pre-requisites that are required with installing View 5.1 components:

- 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 5.1 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 5.1

The View Manager installer requires 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 below:

https://my.vmware.com/web/vmware/details?productId=268&downloadGroup=VIEW-512-PREMIERE

Install View Connection server

1. Log into the View Connection server with Domain Administrator credentials.
2. Open installer file VMware-viewconnectionserver-x86_64-5.1.2-928164.exe with "Run as administrator"
3. Click Next to installation wizard.

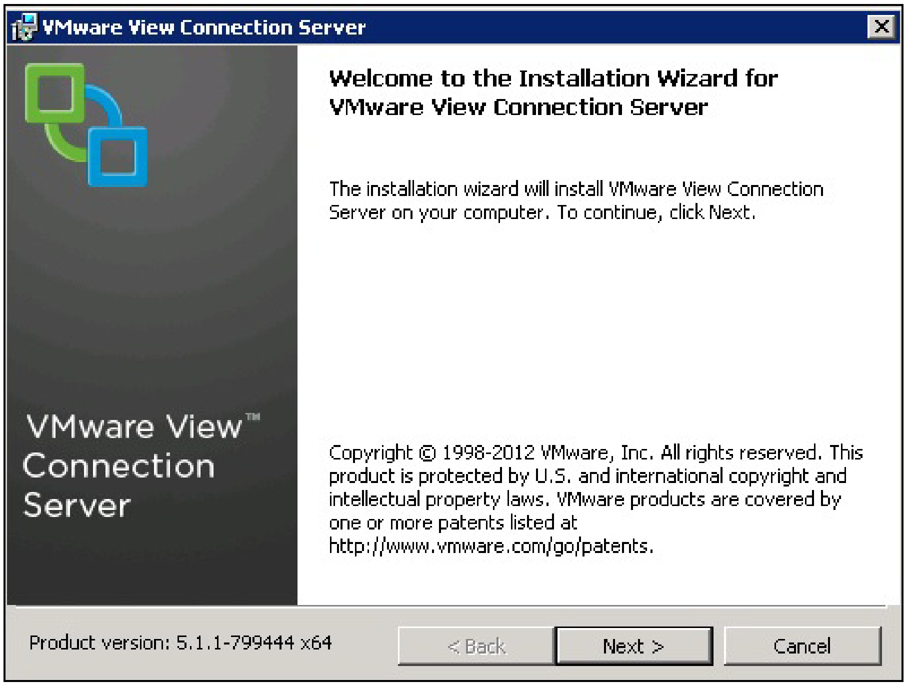

Welcome to the Installation Wizard for VMware View Connection Server

The installation wizard will install VMware View Connection Server on your computer. To continue, click Next.

Copyright © 1998-2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 5.1.1-799444 x64

4. Accept End User License agreement and Click Next.

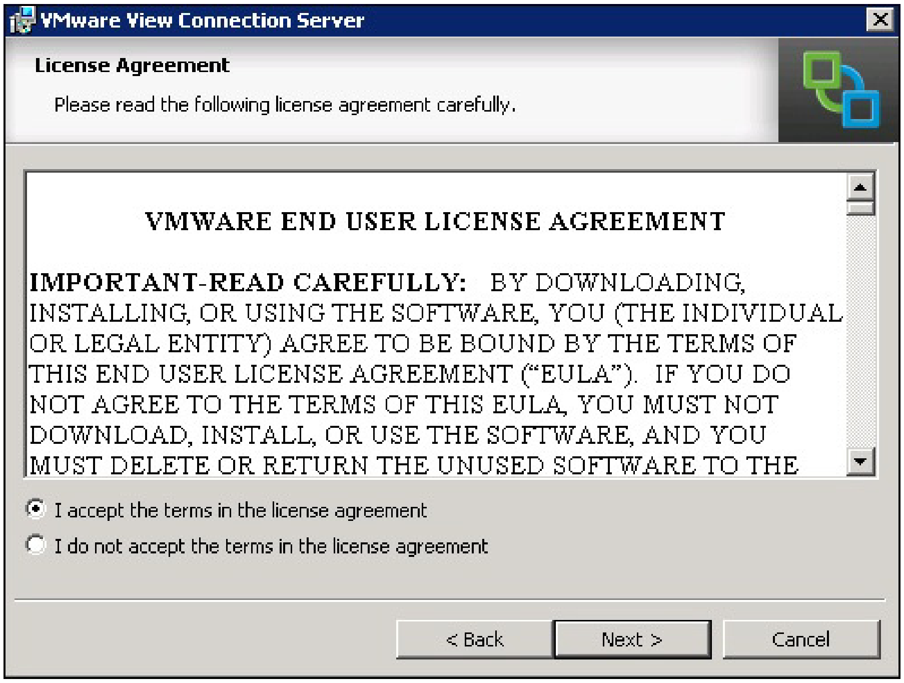

License Agreement

Please read the following license agreement carefully.

VMWARE END USER LICENSE AGREEMENT

IMPORTANT-READ CAREFULLY: BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU (THE INDIVIDUAL OR LEGAL ENTITY) AGREE TO BE BOUND BY THE TERMS OF THIS END USER LICENSE AGREEMENT ("EULA"). IF YOU DO NOT AGREE TO THE TERMS OF THIS EULA, YOU MUST NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND YOU MUST DELETE OR RETURN THE UNUSED SOFTWARE TO THE

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

5. Select desired location for installer to install all the components and Click Next.

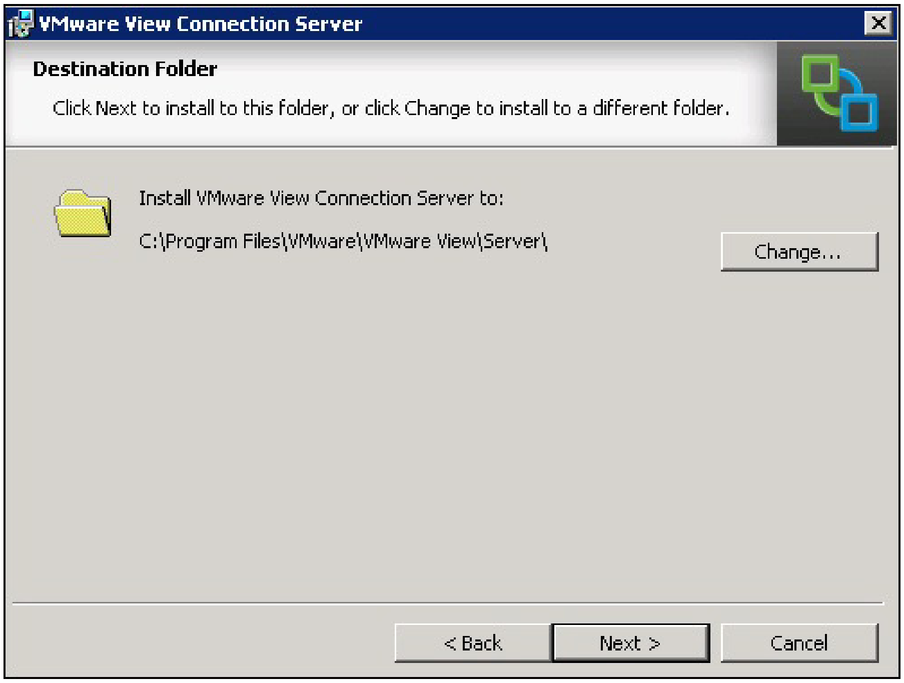

VMware View Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware View Connection Server to:

C:\Program Files\VMware\VMware View\Server\

6. Select Standard server installation.

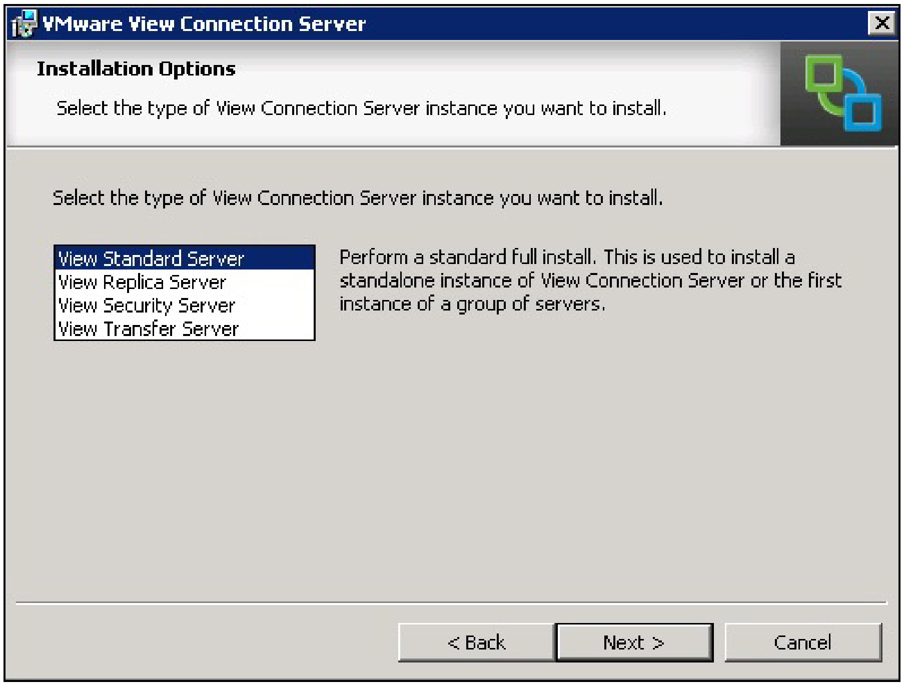

VMware View Connection Server

Installation Options

Select the type of View Connection Server instance you want to install.

Select the type of View Connection Server instance you want to install.

- View Standard Server
- View Replica Server
- View Security Server
- View Transfer Server

Perform a standard full install. This is used to install a standalone instance of View Connection Server or the first instance of a group of servers.

7. Enter password and Click Next.

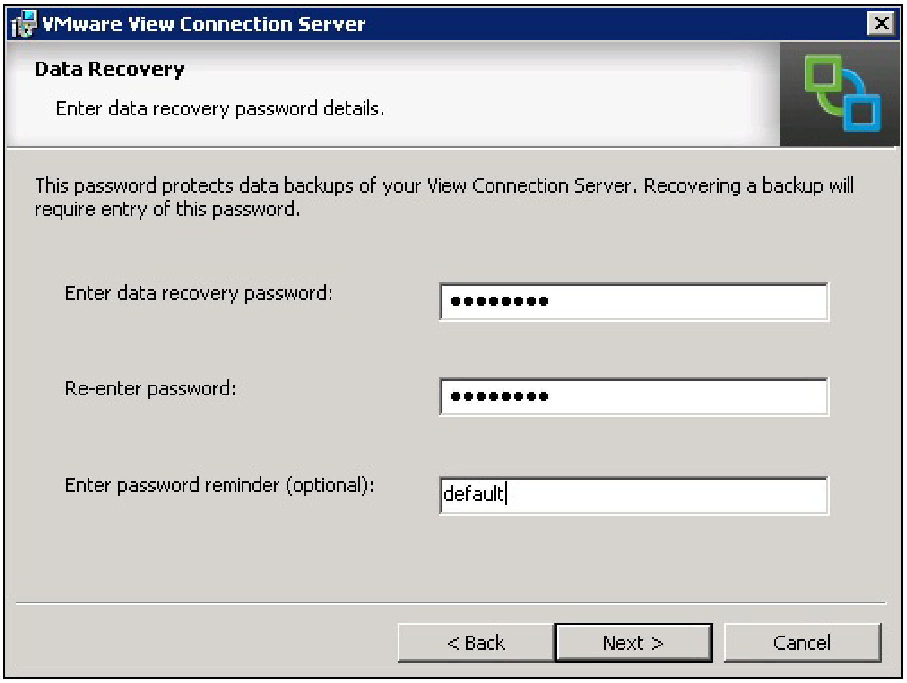

Data Recovery

Enter data recovery password details.

This password protects data backups of your View Connection Server. Recovering a backup will require entry of this password.

Enter data recovery password: ●●●●●●●●●

Re-enter password: ●●●●●●●●●

Enter password reminder (optional): default

8. Select the radio button Configure Windows Firewall Automatically. Click Next.

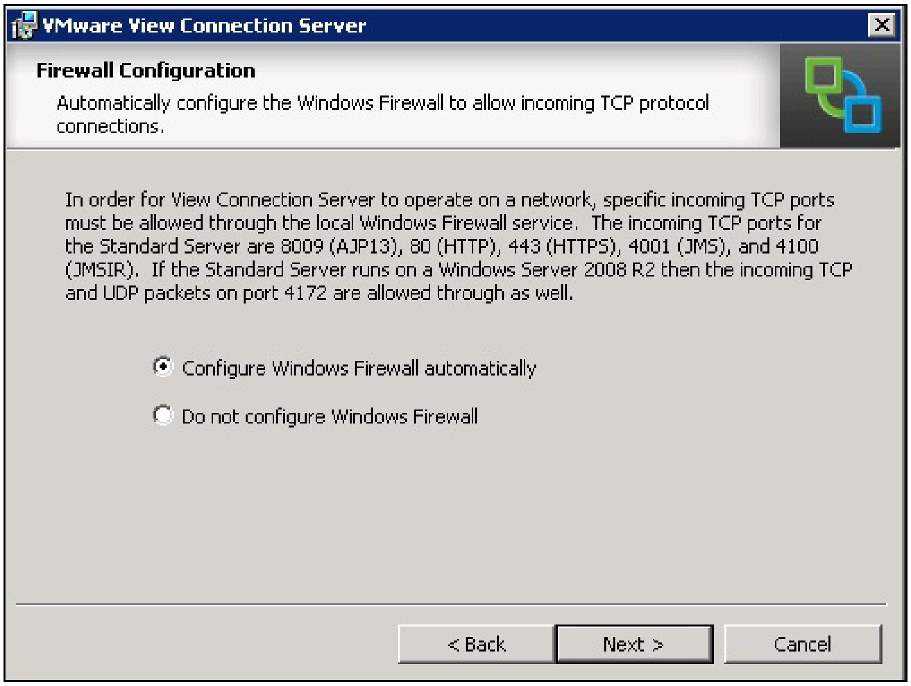

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for View Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Standard Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), and 4100 (JMSIR). If the Standard Server runs on a Windows Server 2008 R2 then the incoming TCP and UDP packets on port 4172 are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

9. Select the radio button Authorize A Specific Domain User Or Domain Group. Click Next.

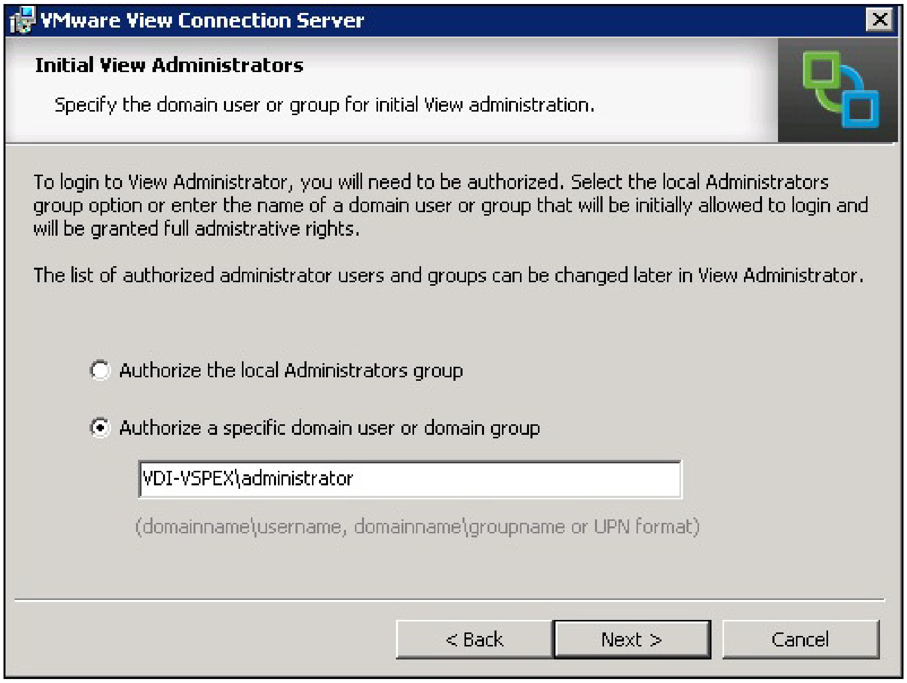

VMware View Connection Server

Initial View Administrators

Specify the domain user or group for initial View administration.

To login to View Administrator, you will need to be authorized. Select the local Administrators group option or enter the name of a domain user or group that will be initially allowed to login and will be granted full admistrative rights.

The list of authorized administrator users and groups can be changed later in View Administrator.

○ Authorize the local Administrators group

● Authorize a specific domain user or domain group

VDI-VSPEX\administrator

(domainname\username, domainname\groupname or UPN format)

10. Uncheck box for participate anonymously in the user experience improvement program. Click Next and click Install.

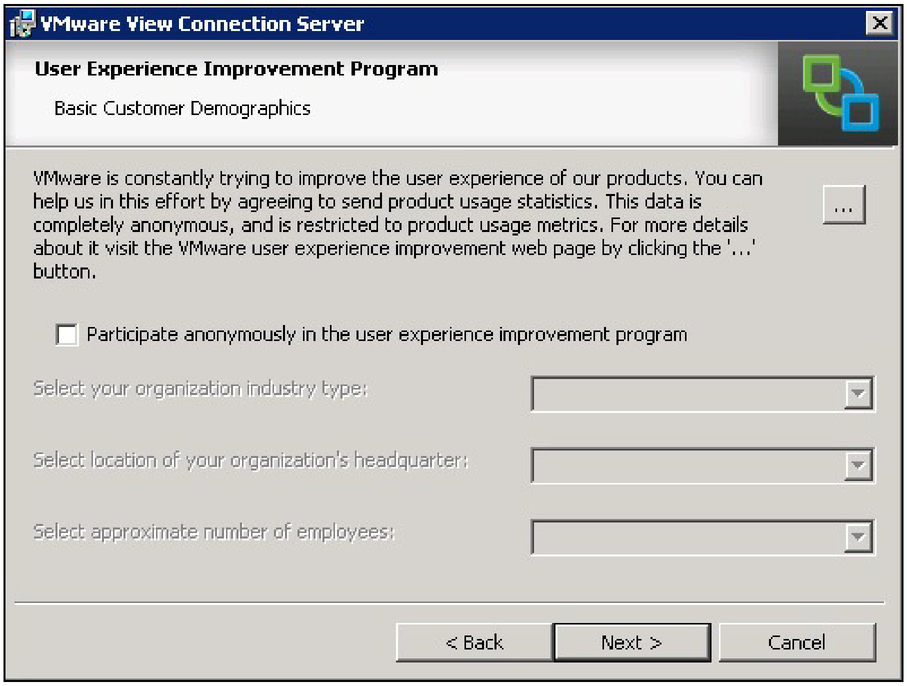

VMware View Connection Server

User Experience Improvement Program

Basic Customer Demographics

VMware is constantly trying to improve the user experience of our products. You can help us in this effort by agreeing to send product usage statistics. This data is completely anonymous, and is restricted to product usage metrics. For more details about it visit the VMware user experience improvement web page by clicking the '...' button.

□ Participate anonymously in the user experience improvement program

Select your organization industry type:

Select location of your organization's headquarter:

Select approximate number of employees:

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

Note You will need to install Flash player plugin v10.3 or higher to use the web browser to log into View Administrator.

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

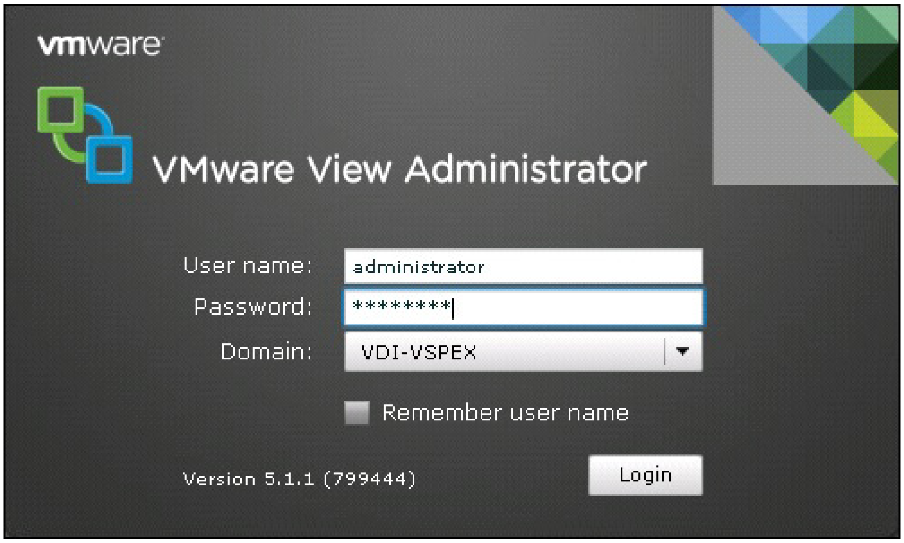

User name: administrator

Password: ********

Domain: VDI-VSPEX

☐ Remember user name

Install View Replica Server

1. Log into the replica server with your domain administrator credentials.
2. Open the installer file VMware-viewconnectionserver-x86_64-5.1.2-928164.exe with "Run as administrator."
3. Click Next.

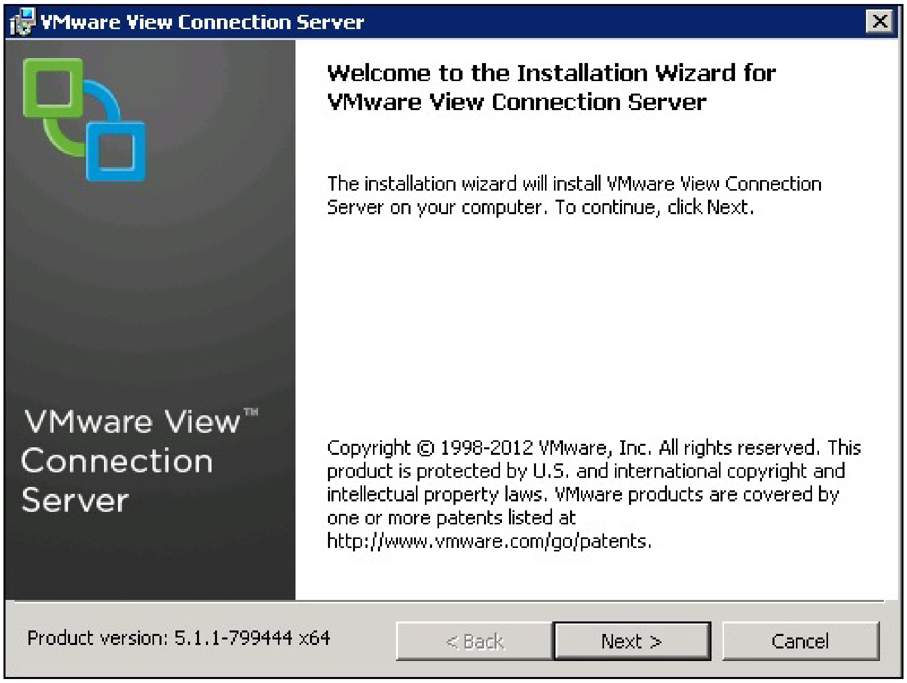

VMware View Connection Server

Welcome to the Installation Wizard for VMware View Connection Server

The installation wizard will install VMware View Connection Server on your computer. To continue, click Next.

Copyright © 1998-2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 5.1.1-799444 x64

4. Accept the License agreement. Click Next.
5. Select the destination for installation. Click Next.

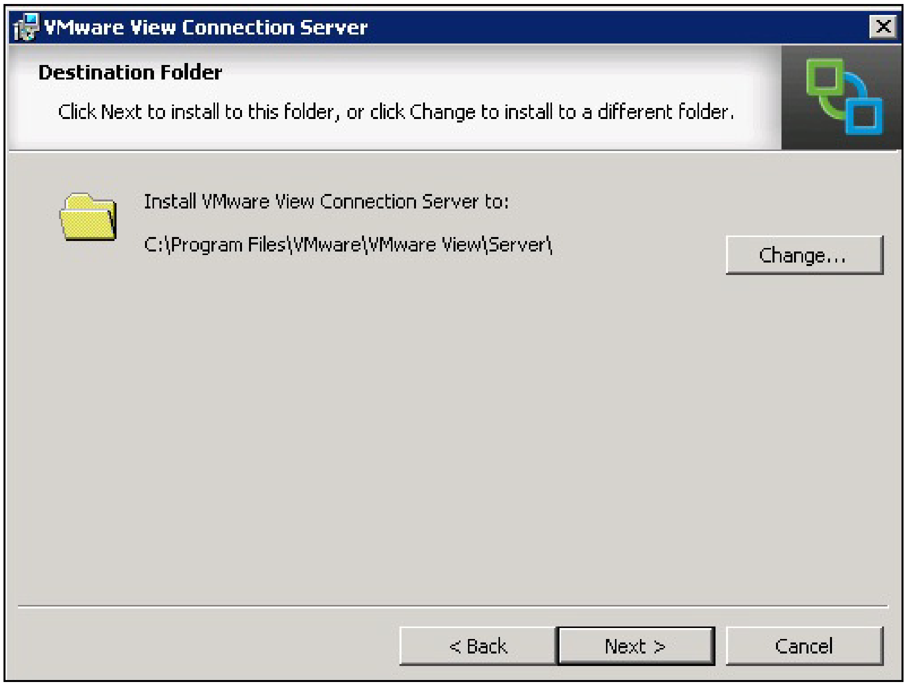

VMware View Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware View Connection Server to:

C:\Program Files\VMware\VMware View\Server\

6. Select Replica server installation. Click Next.

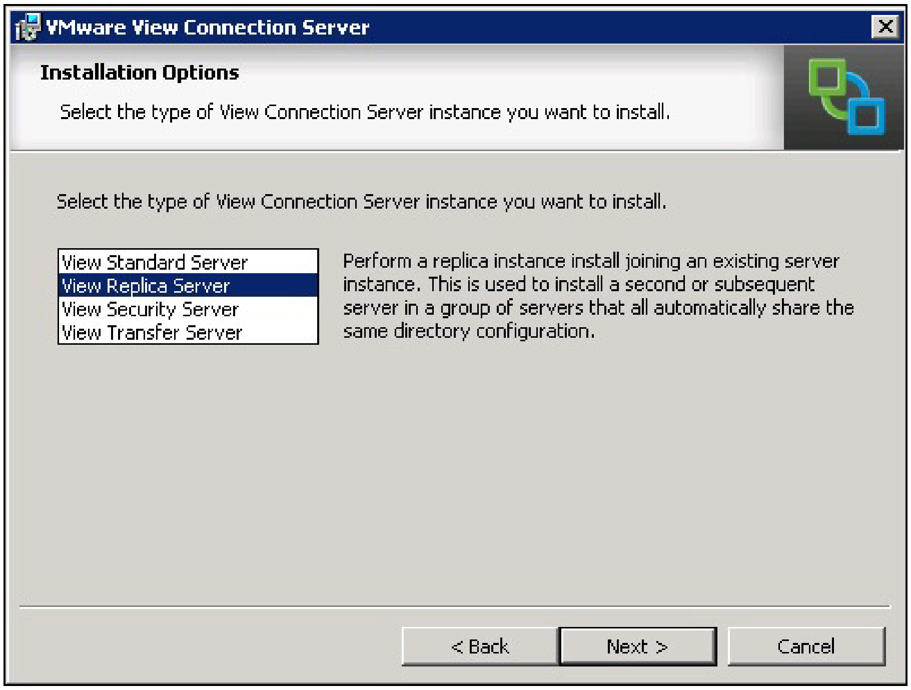

VMware View Connection Server

Installation Options

Select the type of View Connection Server instance you want to install.

Select the type of View Connection Server instance you want to install.

- View Standard Server
- View Replica Server
- View Security Server
- View Transfer Server

Perform a replica instance install joining an existing server instance. This is used to install a second or subsequent server in a group of servers that all automatically share the same directory configuration.

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

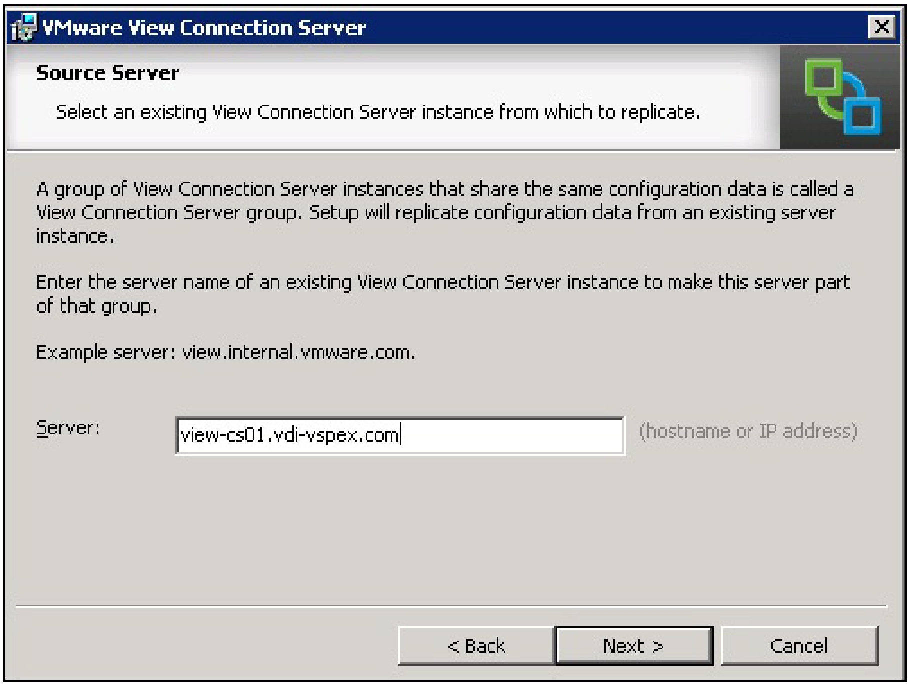

VMware View Connection Server

Source Server

Select an existing View Connection Server instance from which to replicate.

A group of View Connection Server instances that share the same configuration data is called a View Connection Server group. Setup will replicate configuration data from an existing server instance.

Enter the server name of an existing View Connection Server instance to make this server part of that group.

Example server: view.internal.vmware.com.

Server: view-cs01.vdi-vspex.com (hostname or IP address)

8. Click Next.

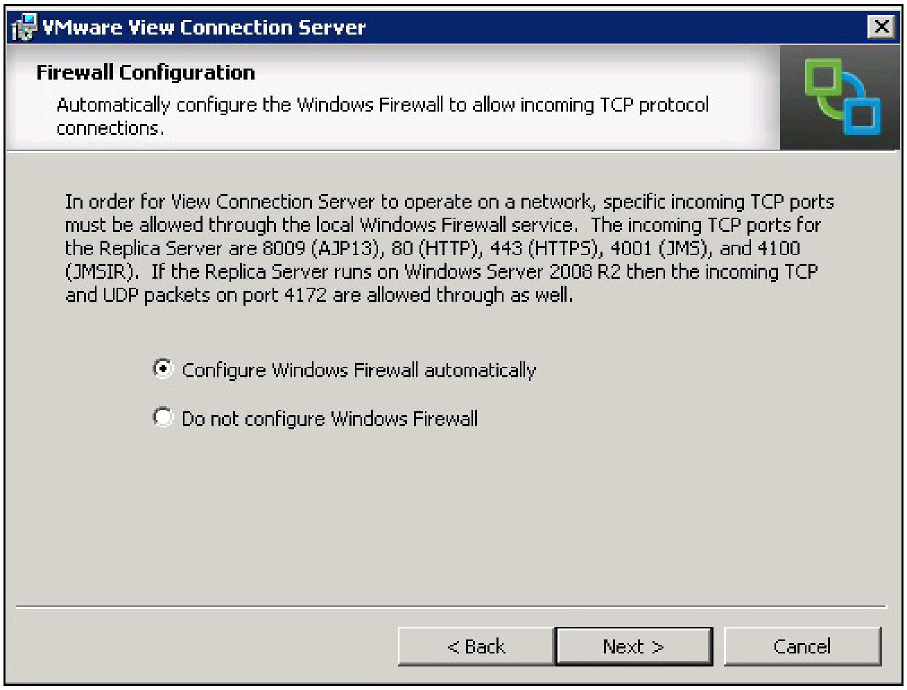

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for View Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Replica Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), and 4100 (JMSIR). If the Replica Server runs on Windows Server 2008 R2 then the incoming TCP and UDP packets on port 4172 are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

9. Click Next.

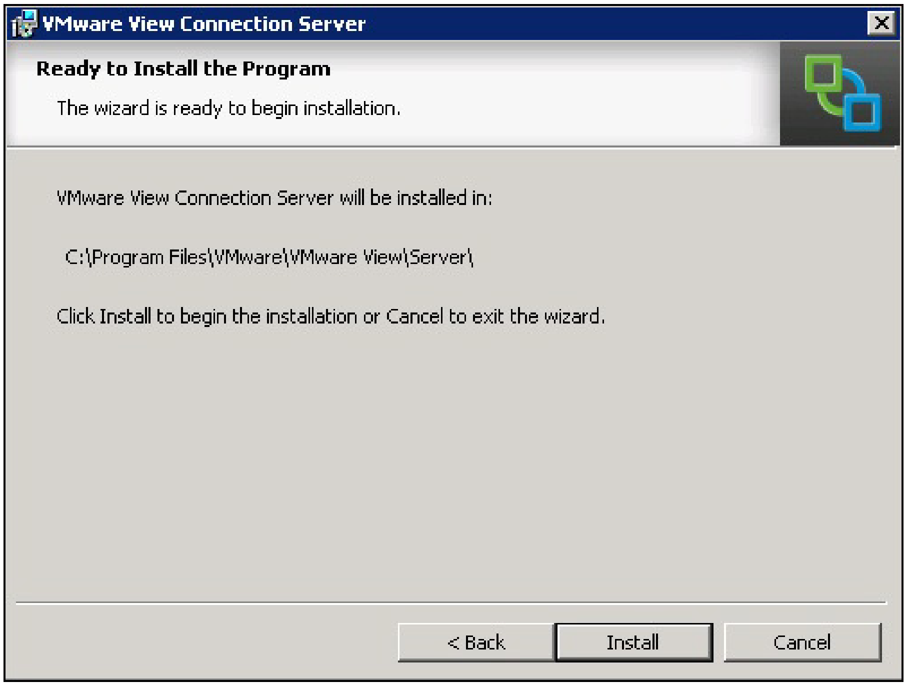

Ready to Install the Program

The wizard is ready to begin installation.

VMware View Connection Server will be installed in:

C:\Program Files\VMware\VMware View\Server\

Click Install to begin the installation or Cancel to exit the wizard.

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 the View composer installer VMware-viewcomposer-3.0.0-691993.exe
2. Create a database and ODBC connection for view composer installation. See section Create Database for View Composer Server about how to create database for view composer server.
3. Click Next on the installation wizard.

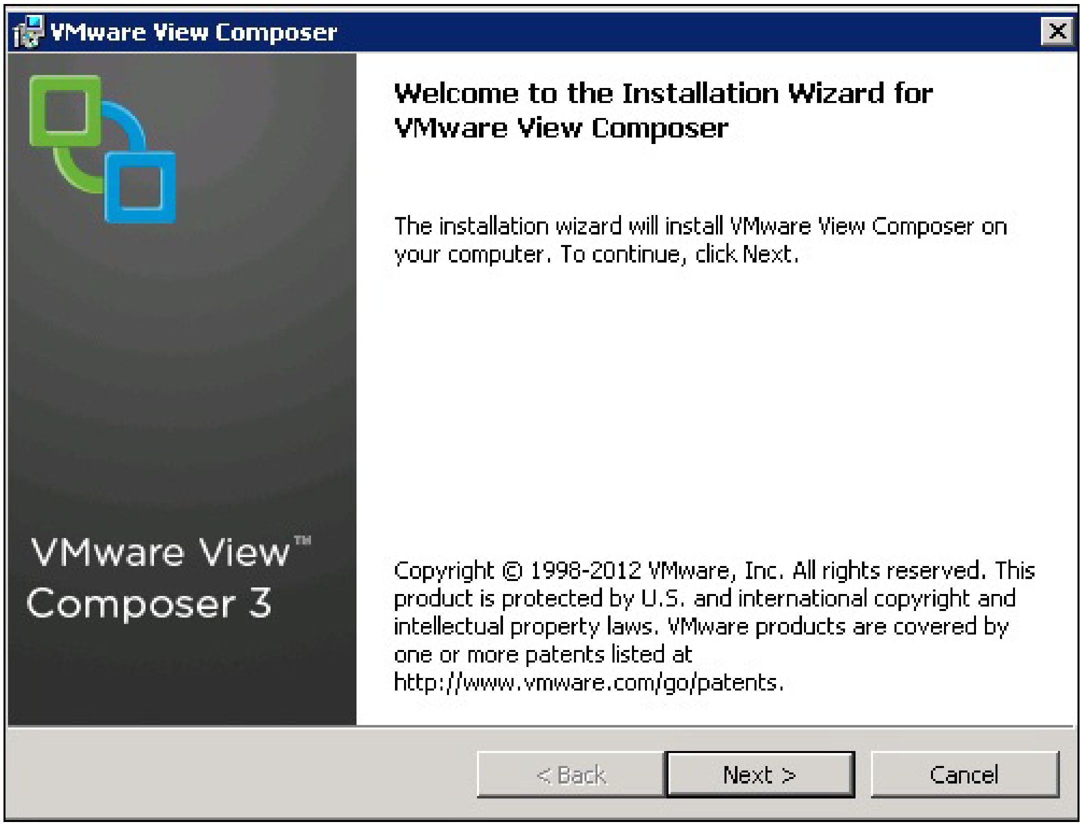

VMware View Composer

Welcome to the Installation Wizard for VMware View Composer

The installation wizard will install VMware View Composer on your computer. To continue, click Next.

Copyright © 1998-2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

4. Accept the License agreement. Click Next.
5. Select the location for the View Composer installation. Click Next.

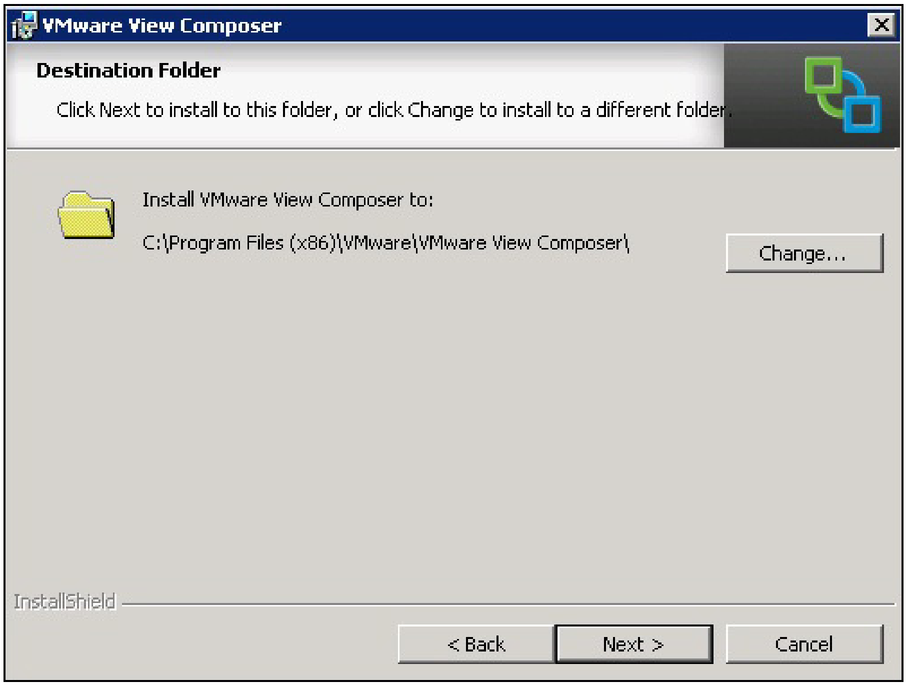

VMware View Composer

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware View Composer to:

C:\Program Files (x86)\VMware\VMware View Composer\

6. Enter the newly created Database and SA user Information for the View Composer installation by going to section Create Database for View Composer Server. Click Next.

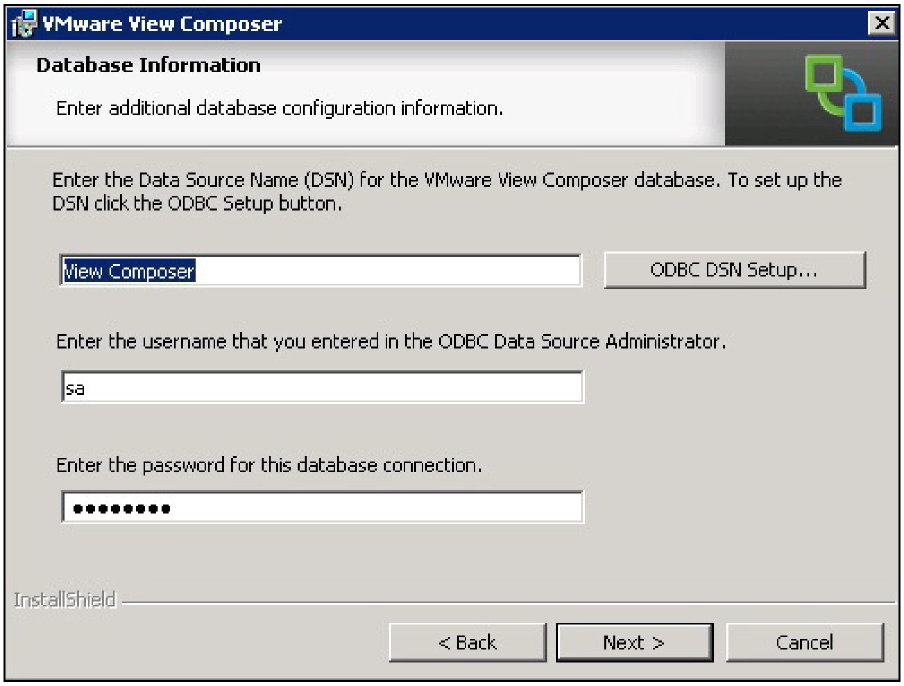

VMware View Composer

Database Information

Enter additional database configuration information.

Enter the Data Source Name (DSN) for the VMware View Composer database. To set up the DSN click the ODBC Setup button.

View Composer

Enter the username that you entered in the ODBC Data Source Administrator.

sa

Enter the password for this database connection.

●●●●●●●●●

7. Accept the default port settings and click Next.

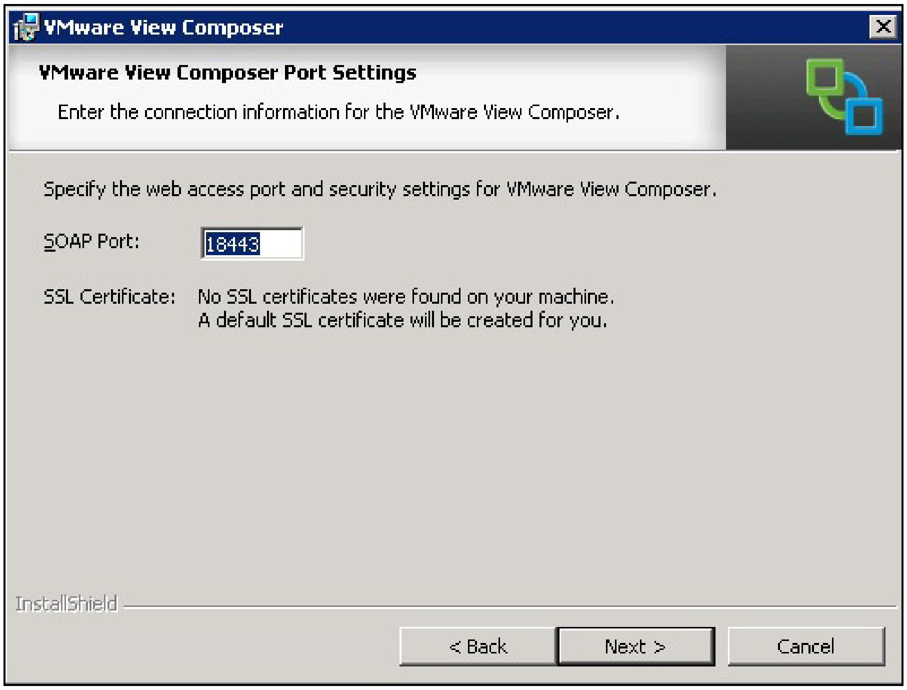

VMware View Composer

VMware View Composer Port Settings

Enter the connection information for the VMware View Composer.

Specify the web access port and security settings for VMware View Composer.

SOAP Port: 18443

SSL Certificate: No SSL certificates were found on your machine.
A default SSL certificate will be created for you.

8. Click Install.

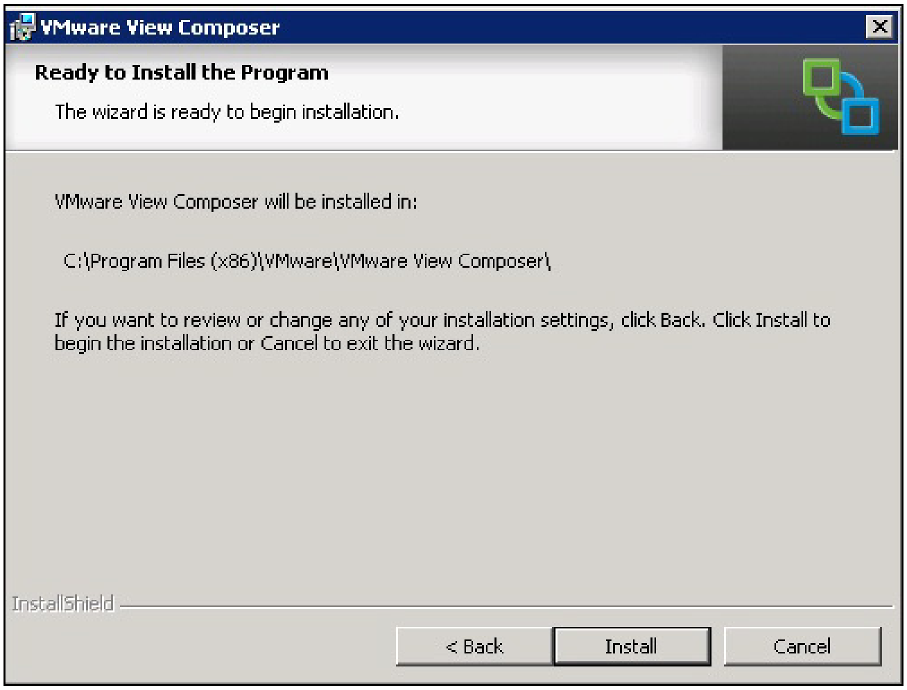

VMware View Composer

Ready to Install the Program

The wizard is ready to begin installation.

VMware View Composer will be installed in:

C:\Program Files (x86)\VMware\VMware View Composer\

If you want to review or change any of your installation settings, click Back. Click Install to begin the installation or Cancel to exit the wizard.

View Administrator Configuration

To configure the View 5.1 system, follow these steps:

1. Log into VMware View Administrator using a web browser.

2. Select View configuration.
3. From the drop-down menu select Product Licensing and Usage.
4. Click n Edit settings and enter a valid license key for View Manager.

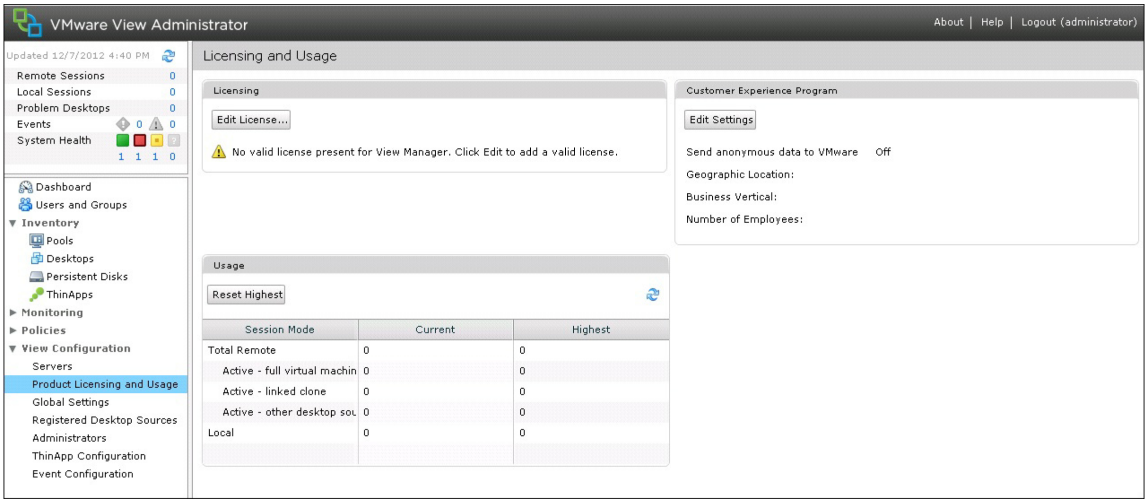

No valid license present for View Manager. Click Edit to add a valid license.

Send anonymous data to VMware Off

Geographic Location:

Business Vertical:

Number of Employees:

Session Mode	Current	Highest
Total Remote	0	0
Active - full virtual machin	0	0
Active - linked clone	0	0
Active - other desktop sou	0	0
Local	0	0

5. In the View Configuration, click Servers. Select vCenter Servers tab. Click Add.

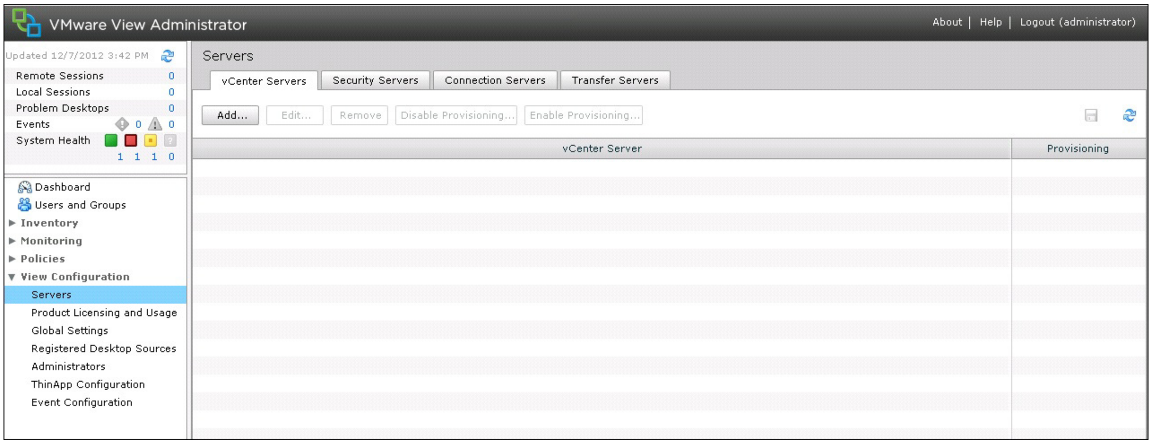

vCenter Server	Provisioning

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

- For this test case we used 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

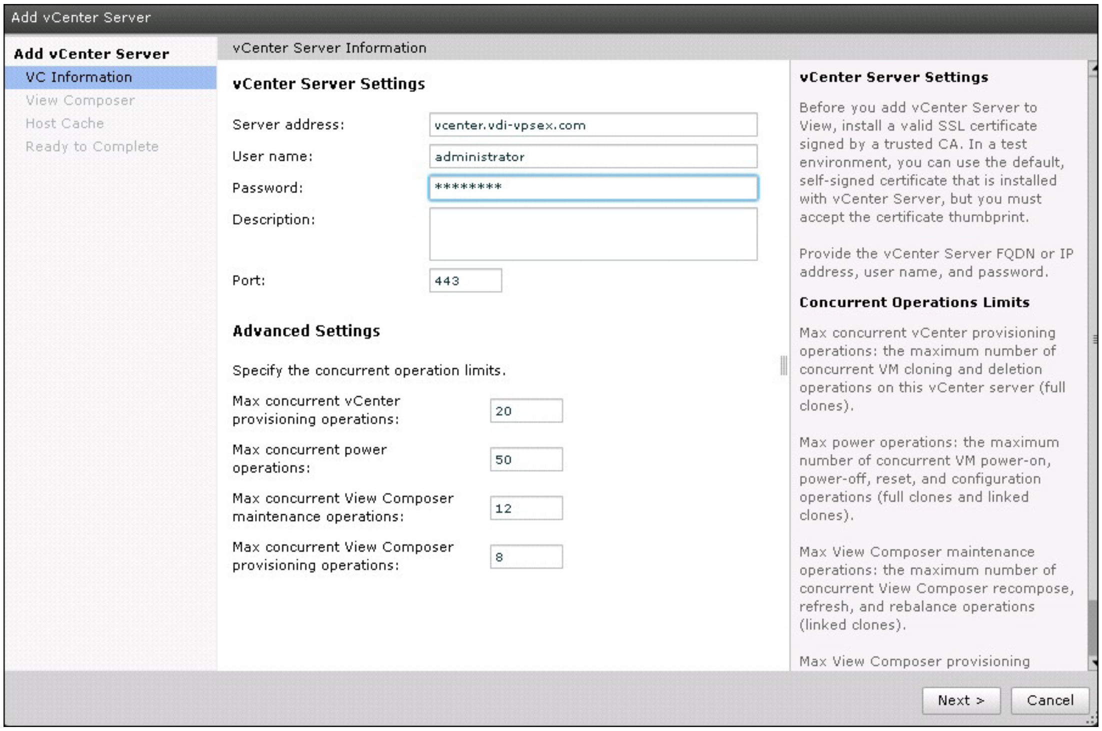
Add vCenter Server

vCenter Server Information

vCenter Server Settings

Server address: vcenter.vdi-vpsex.com

User name: administrator

Password: ********

Description:

Port: 443

Advanced Settings

Specify the concurrent operation limits.

Max concurrent vCenter provisioning operations: 20

Max concurrent power operations: 50

Max concurrent View Composer maintenance operations: 12

Max concurrent View Composer provisioning operations: 8

vCenter Server Settings

Before you add vCenter Server to View, install a valid SSL certificate signed by a trusted CA. In a test environment, you can use the default, self-signed certificate that is installed with vCenter Server, but you must accept the certificate thumbprint.

Provide the vCenter Server FQDN or IP address, user name, and password.

Concurrent Operations Limits

Max concurrent vCenter provisioning operations: the maximum number of concurrent VM cloning and deletion operations on this vCenter server (full clones).

Max power operations: the maximum number of concurrent VM power-on, power-off, reset, and configuration operations (full clones and linked clones).

Max View Composer maintenance operations: the maximum number of concurrent View Composer recompose, refresh, and rebalance operations (linked clones).

Max View Composer provisioning

7. Click View Certificate and accept the certificate warning.

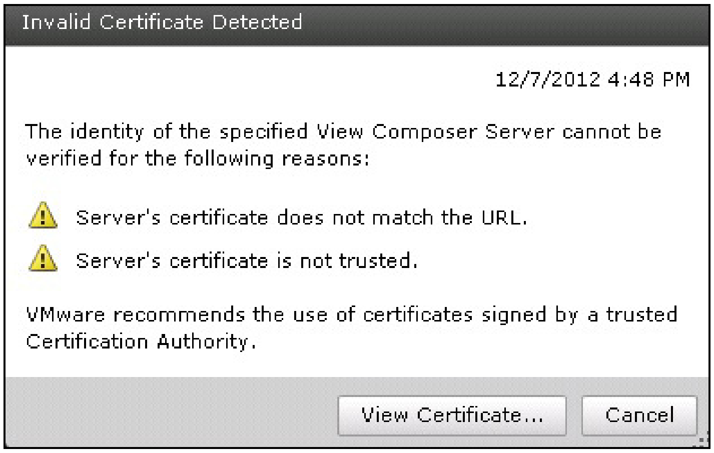
Invalid Certificate Detected

12/7/2012 4:48 PM

The identity of the specified View Composer Server cannot be verified for the following reasons:

⚠ Server's certificate does not match the URL.

⚠ Server's certificate is not trusted.

VMware recommends the use of certificates signed by a trusted Certification Authority.

8. Select the View Composer settings. Select the radio button for View Composer server installed with either vCenter server or as a standalone server. In the case of a standalone server, enter the server address, username, and password. Click Next.

9. Click View Certificate and accept the certificate.

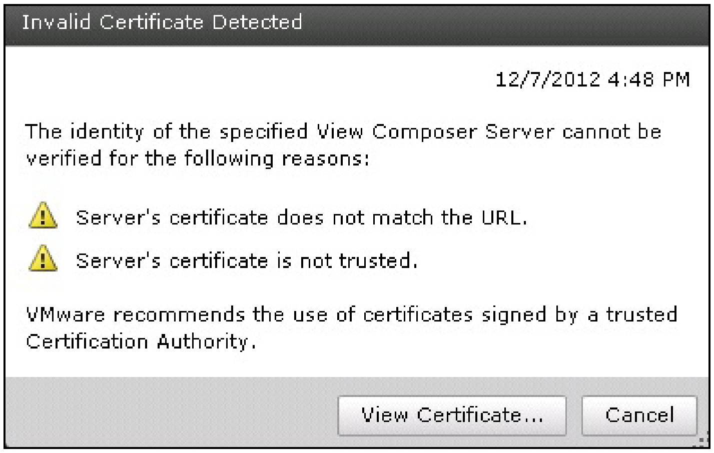

Invalid Certificate Detected

12/7/2012 4:48 PM

The identity of the specified View Composer Server cannot be verified for the following reasons:

⚠ Server's certificate does not match the URL.

⚠ Server's certificate is not trusted.

VMware recommends the use of certificates signed by a trusted Certification Authority.

10. Click Add to add a view composer domain.

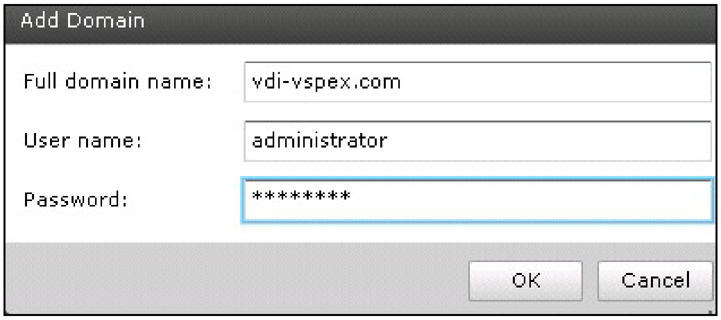

Add Domain

Full domain name: vdi-vspex.com

User name: administrator

Password: ********

11. Click Next.

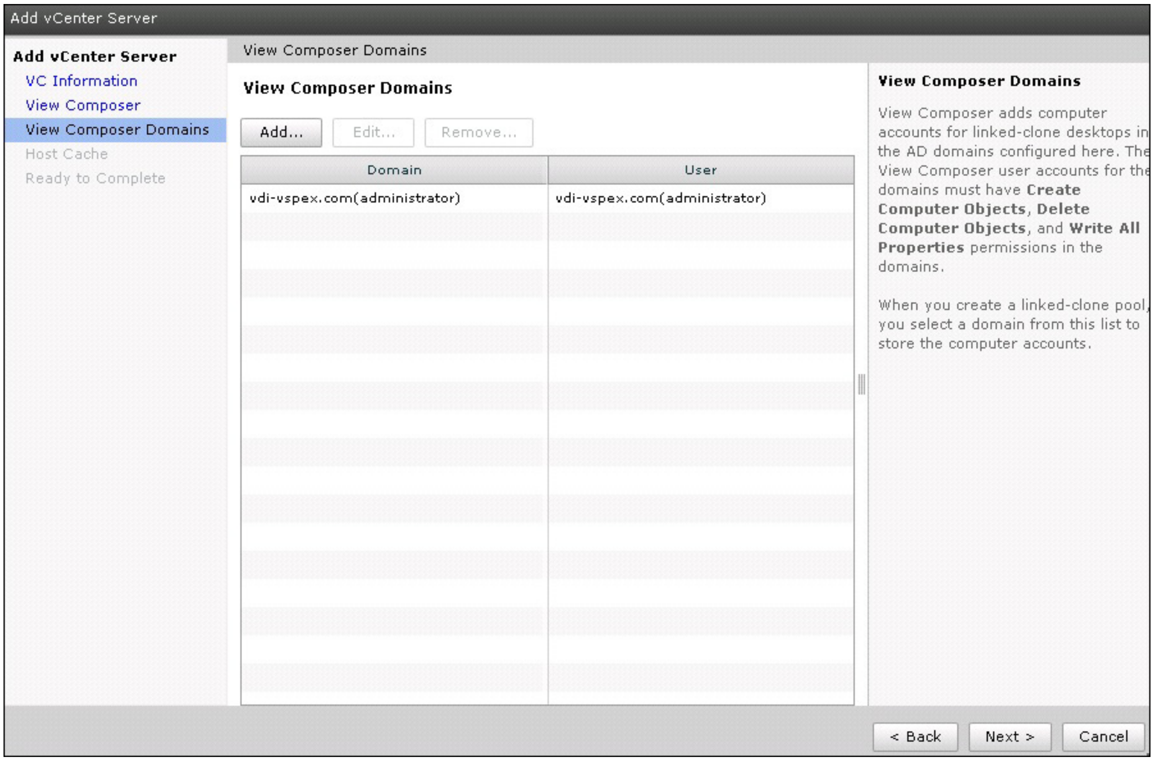

View Composer Domains

Domain	User
vdi-vspex.com(administrator)	vdi-vspex.com(administrator)

View Composer Domains

View Composer adds computer accounts for linked-clone desktops in the AD domains configured here. The View Composer user accounts for the domains must have **Create Computer Objects**, **Delete Computer Objects**, and **Write All Properties** permissions in the domains.

When you create a linked-clone pool, you select a domain from this list to store the computer accounts.

12. Click the check box to enable host caching. Set Default host cache size and Click Next.

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

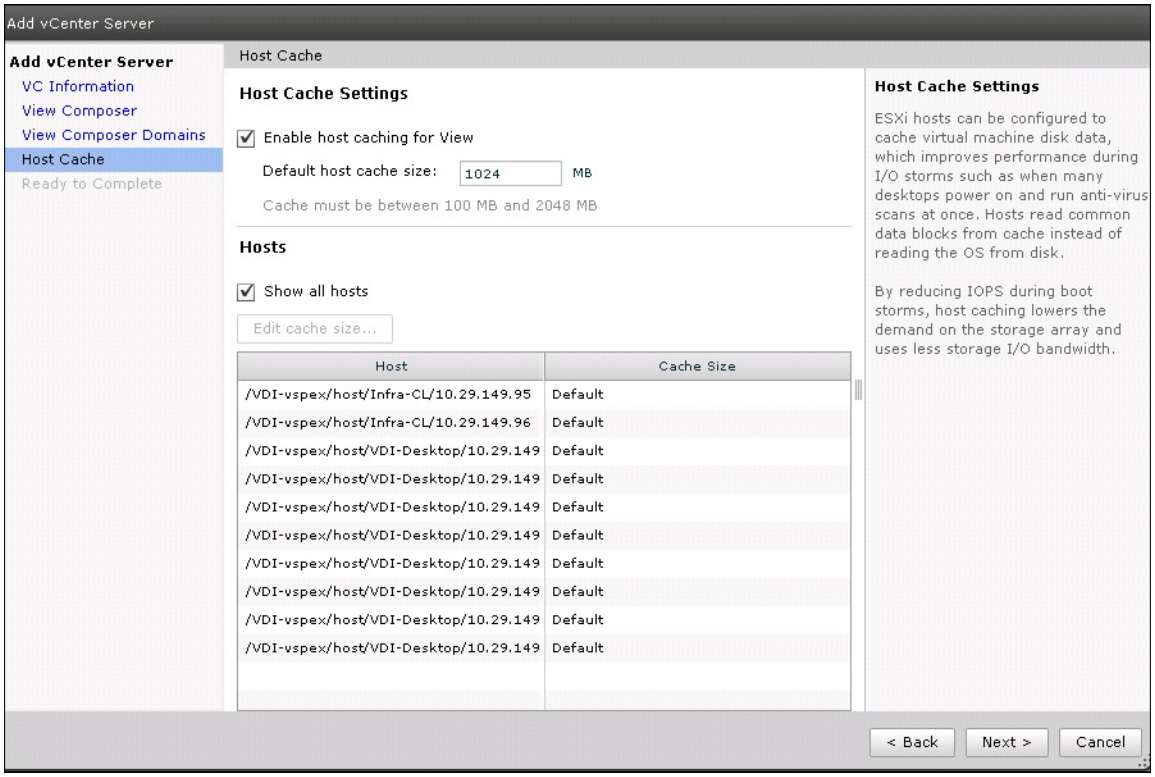

Add vCenter Server

Host Cache

Host Cache Settings

☑ Enable host caching for View

Default host cache size: 1024 MB

Cache must be between 100 MB and 2048 MB

Hosts

☑ Show all hosts

Host	Cache Size
/VDI-vspex/host/Infra-CL/10.29.149.95	Default
/VDI-vspex/host/Infra-CL/10.29.149.96	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default
/VDI-vspex/host/VDI-Desktop/10.29.149	Default

Host Cache Settings

ESXi hosts can be configured to cache virtual machine disk data, which improves performance during I/O storms such as when many desktops power on and run anti-virus scans at once. Hosts read common data blocks from cache instead of reading the OS from disk.

By reducing IOPS during boot storms, host caching lowers the demand on the storage array and uses less storage I/O bandwidth.

13. Verify the details and click Finish.

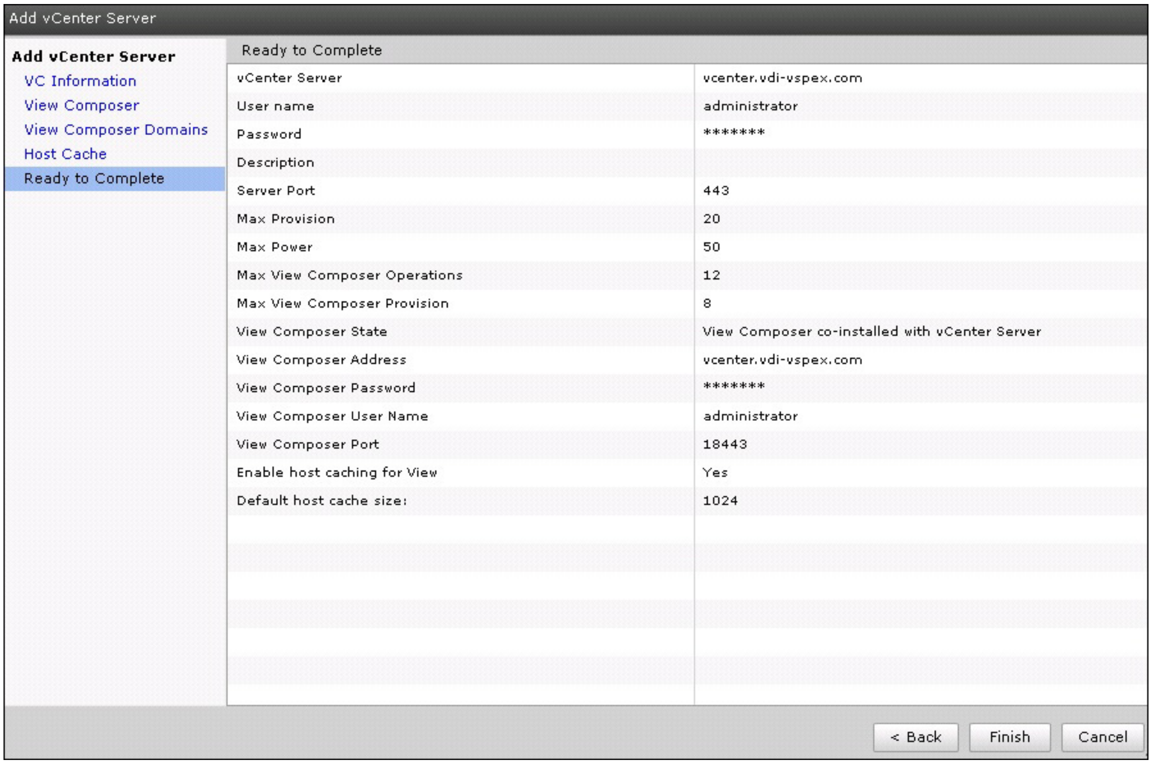

Ready to Complete

vCenter Server	vcenter.vdi-vspex.com
User name	administrator
Password	*******
Description	
Server Port	443
Max Provision	20
Max Power	50
Max View Composer Operations	12
Max View Composer Provision	8
View Composer State	View Composer co-installed with vCenter Server
View Composer Address	vcenter.vdi-vspex.com
View Composer Password	*******
View Composer User Name	administrator
View Composer Port	18443
Enable host caching for View	Yes
Default host cache size:	1024

14. Create a new database for the View Event Database in SQL server. (See section Create Event Database for View Administrator. for how to create Event Database for View administrator.)

15. Click Event Database configuration.

16. Enter Database server information, database name and username/password. For the table prefix add VE_

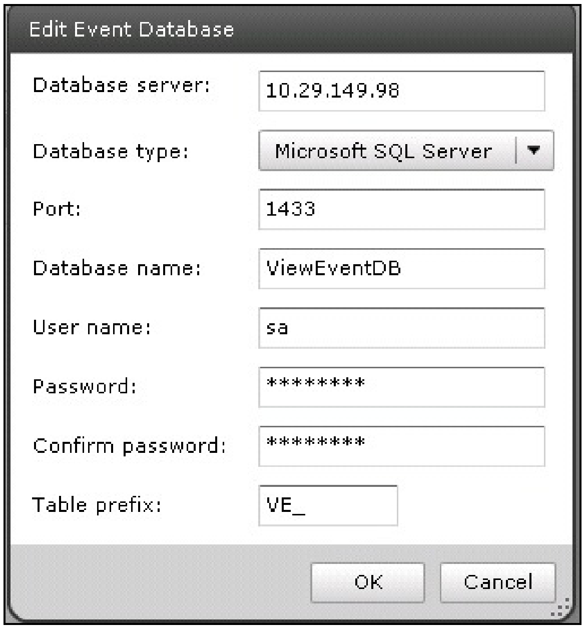

Database server:	10.29.149.98
Database type:	Microsoft SQL Server
Port:	1433
Database name:	ViewEventDB
User name:	sa
Password:	********
Confirm password:	********
Table prefix:	VE_

17. After completing the view configuration go to the Dashboard for View Administrator and check System Health and verify all components are shown as green.

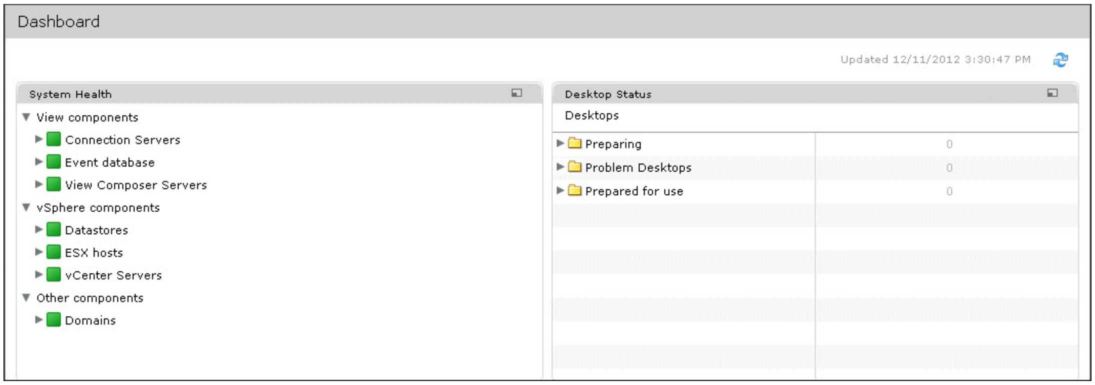

Install SSL Certificate for View Connection and Replica Server

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

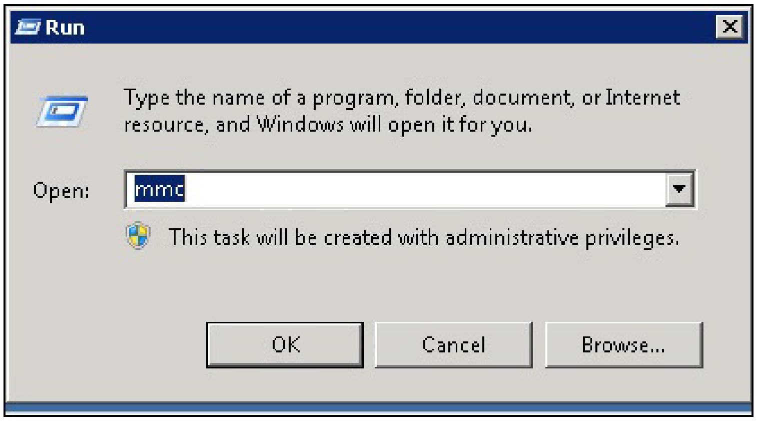

Type the name of a program, folder, document, or Internet resource, and Windows will open it for you.

Open: mmc

This task will be created with administrative privileges.

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

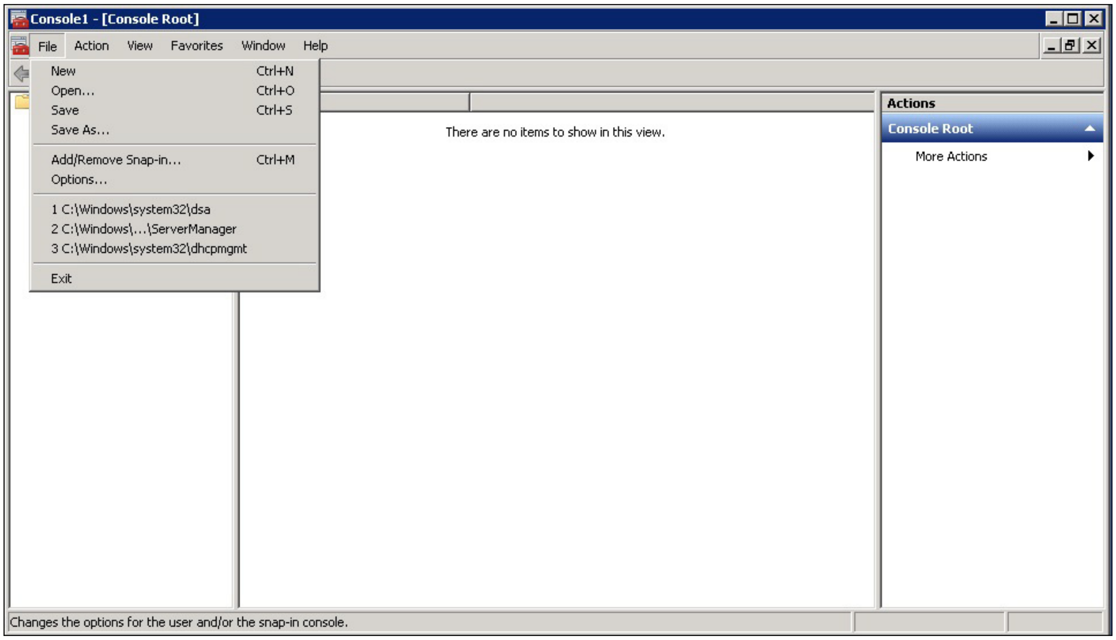
Console1 - [Console Root]
File
Action
View
Favorites
Window
Help
New Ctrl+N
Open... Ctrl+O
Save Ctrl+S
Save As...
Add/Remove Snap-in... Ctrl+M
Options...
1 C:\Windows\system32\dsa
2 C:\Windows\...\ServerManager
3 C:\Windows\system32\dhcpmgmt
Exit
Actions
Console Root
More Actions

There are no items to show in this view.

Changes the options for the user and/or the snap-in console.

4. Select Certificates and click Add.

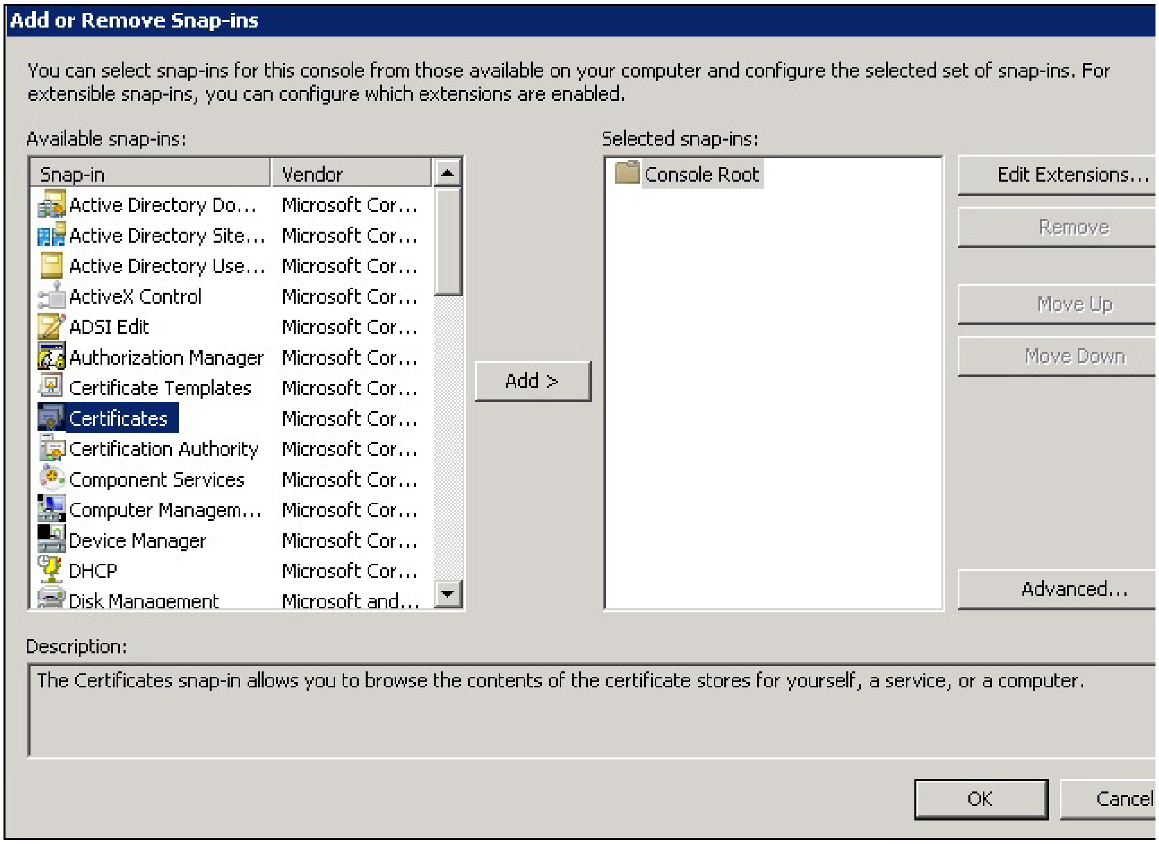

You can select snap-ins for this console from those available on your computer and configure the selected set of snap-ins. For extensible snap-ins, you can configure which extensions are enabled.

Available snap-ins:

Snap-in	Vendor
Active Directory Do...	Microsoft Cor...
Active Directory Site...	Microsoft Cor...
Active Directory Use...	Microsoft Cor...
ActiveX Control	Microsoft Cor...
ADSI Edit	Microsoft Cor...
Authorization Manager	Microsoft Cor...
Certificate Templates	Microsoft Cor...
Certificates	Microsoft Cor...
Certification Authority	Microsoft Cor...
Component Services	Microsoft Cor...
Computer Managem...	Microsoft Cor...
Device Manager	Microsoft Cor...
DHCP	Microsoft Cor...
Disk Management	Microsoft and...

Selected snap-ins:

Console Root

Description:

The Certificates snap-in allows you to browse the contents of the certificate stores for yourself, a service, or a computer.

5. Select the radio button for the computer account.

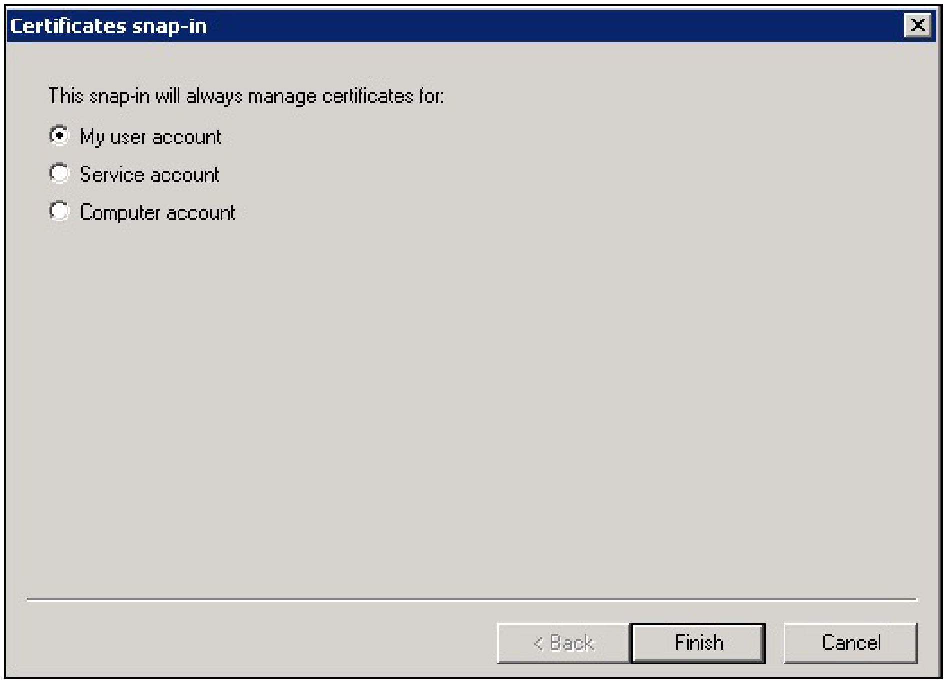

Certificates snap-in

This snap-in will always manage certificates for:

● My user account

○ Service account

○ Computer account

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

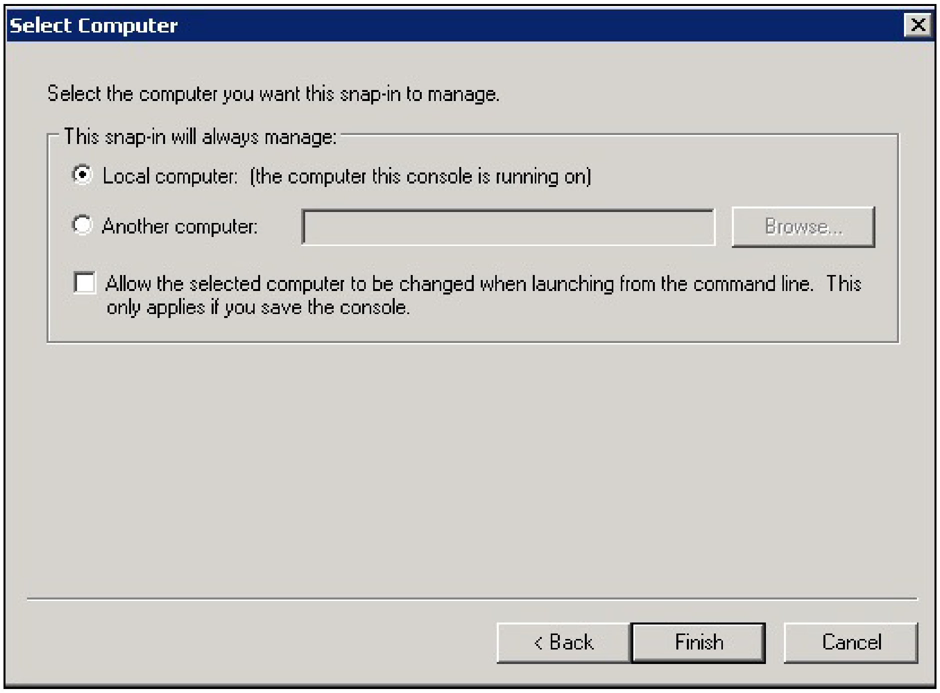

Select Computer

Select the computer you want this snap-in to manage.

This snap-in will always manage:

● Local computer: (the computer this console is running on)

○ Another computer:

□ Allow the selected computer to be changed when launching from the command line. This only applies if you save the console.

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

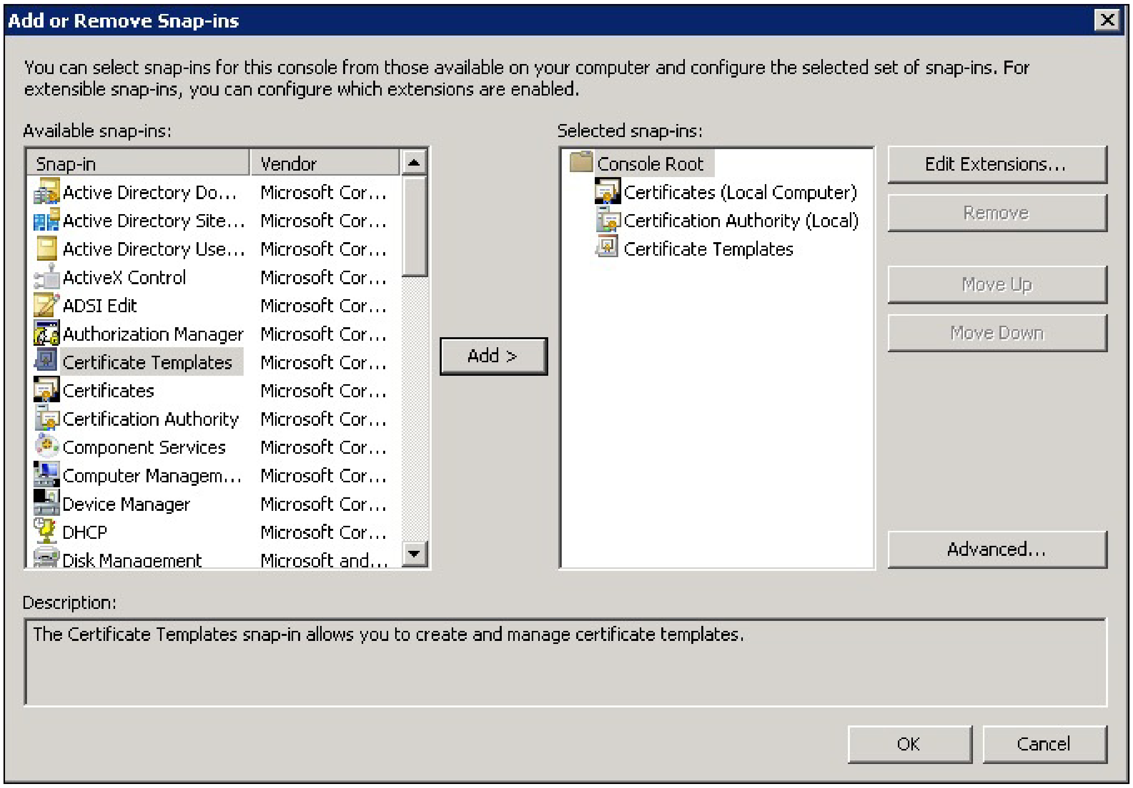

Add or Remove Snap-ins

You can select snap-ins for this console from those available on your computer and configure the selected set of snap-ins. For extensible snap-ins, you can configure which extensions are enabled.

Available snap-ins:

Snap-in	Vendor
Active Directory Do...	Microsoft Cor...
Active Directory Site...	Microsoft Cor...
Active Directory Use...	Microsoft Cor...
ActiveX Control	Microsoft Cor...
ADSI Edit	Microsoft Cor...
Authorization Manager	Microsoft Cor...
Certificate Templates	Microsoft Cor...
Certificates	Microsoft Cor...
Certification Authority	Microsoft Cor...
Component Services	Microsoft Cor...
Computer Managem...	Microsoft Cor...
Device Manager	Microsoft Cor...
DHCP	Microsoft Cor...
Disk Management	Microsoft and...

Selected snap-ins:

- Console Root
  - Certificates (Local Computer)
  - Certification Authority (Local)
  - Certificate Templates

Description:

The Certificate Templates snap-in allows you to create and manage certificate templates.

8. Click Certificate Template and from the list of template displayed on the right side select Web Server.

9. Right-click on Web Server; select properties.

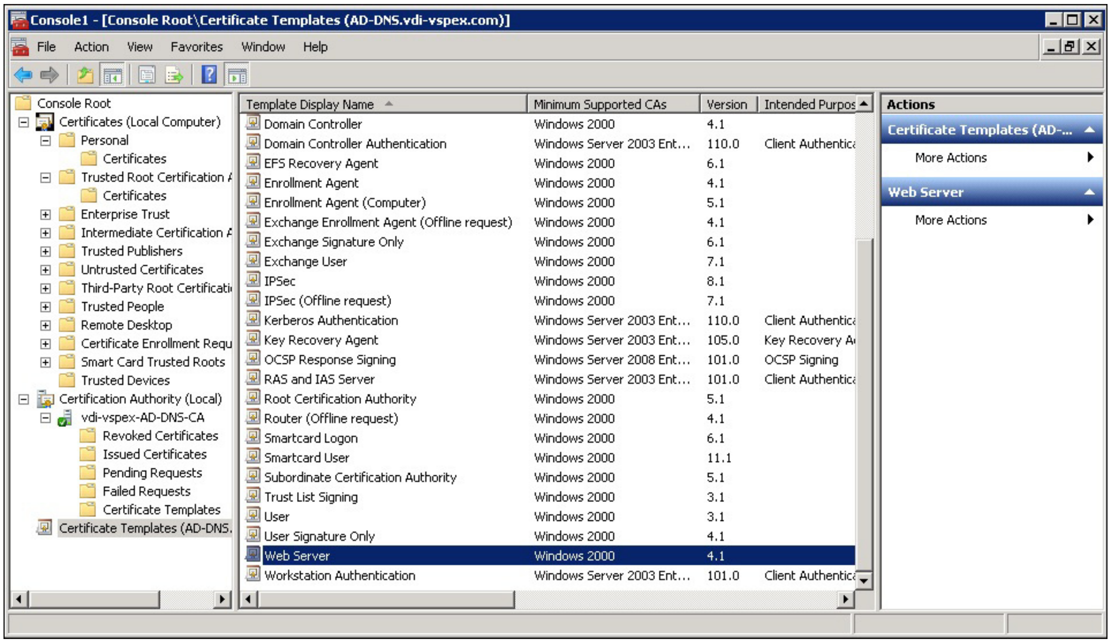

Console1 - [Console Root\Certificate Templates (AD-DNS.vdi-vspex.com)]

Template Display Name	Minimum Supported CAs	Version	Intended Purpos
Domain Controller	Windows 2000	4.1	
Domain Controller Authentication	Windows Server 2003 Ent...	110.0	Client Authentic
EFS Recovery Agent	Windows 2000	6.1	
Enrollment Agent	Windows 2000	4.1	
Enrollment Agent (Computer)	Windows 2000	5.1	
Exchange Enrollment Agent (Offline request)	Windows 2000	4.1	
Exchange Signature Only	Windows 2000	6.1	
Exchange User	Windows 2000	7.1	
IPSec	Windows 2000	8.1	
IPSec (Offline request)	Windows 2000	7.1	
Kerberos Authentication	Windows Server 2003 Ent...	110.0	Client Authentic
Key Recovery Agent	Windows Server 2003 Ent...	105.0	Key Recovery A
OCSP Response Signing	Windows Server 2008 Ent...	101.0	OCSP Signing
RAS and IAS Server	Windows Server 2003 Ent...	101.0	Client Authentic
Root Certification Authority	Windows 2000	5.1	
Router (Offline request)	Windows 2000	4.1	
Smartcard Logon	Windows 2000	6.1	
Smartcard User	Windows 2000	11.1	
Subordinate Certification Authority	Windows 2000	5.1	
Trust List Signing	Windows 2000	3.1	
User	Windows 2000	3.1	
User Signature Only	Windows 2000	4.1	
Web Server	Windows 2000	4.1	
Workstation Authentication	Windows Server 2003 Ent...	101.0	Client Authentic

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

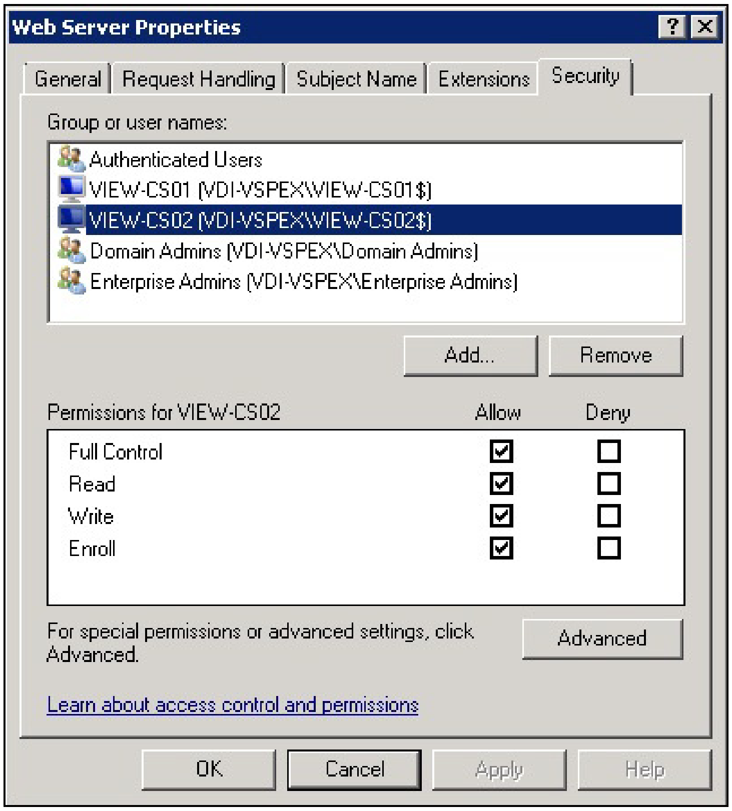

Group or user names:

- Authenticated Users
- VIEW-CS01 (VDI-VSPEX\VIEW-CS01$)
- VIEW-CS02 (VDI-VSPEX\VIEW-CS02$)
- Domain Admins (VDI-VSPEX\Domain Admins)
- Enterprise Admins (VDI-VSPEX\Enterprise Admins)

Permissions for VIEW-CS02	Allow	Deny
Full Control	☑	□
Read	☑	□
Write	☑	□
Enroll	☑	□

For special permissions or advanced settings, click Advanced.

11. Select Certificates on the Console Root ' Personal ' Right Click on Certificates.
12. Select Request New Certificate on All Tasks.

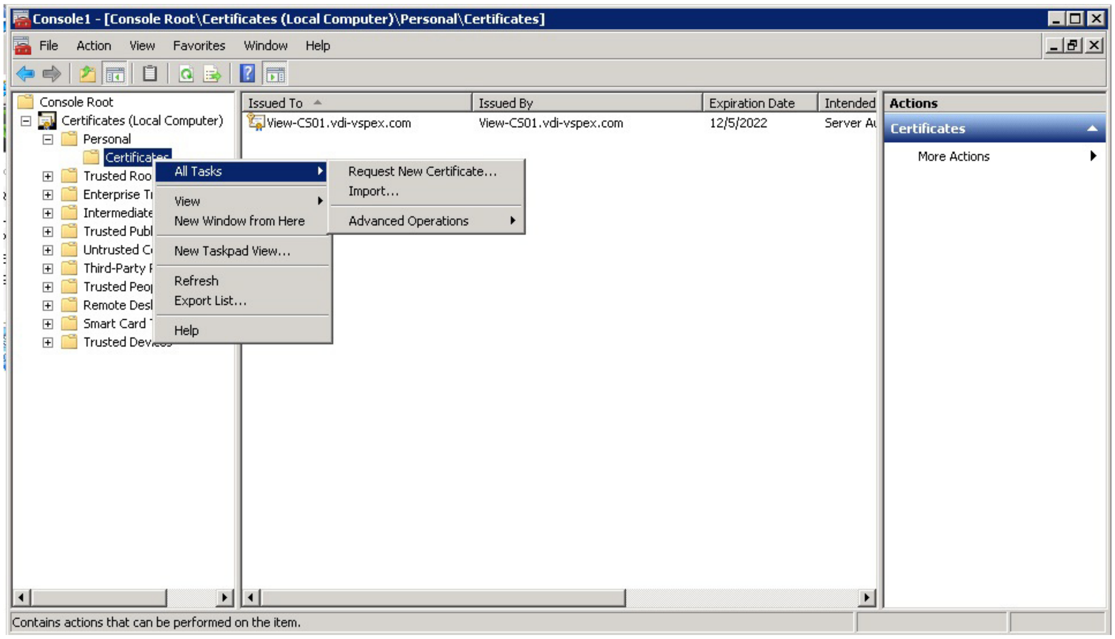

Issued To	Issued By	Expiration Date	Intended
View-CS01.vdi-vspex.com	View-CS01.vdi-vspex.com	12/5/2022	Server Au

13. Click Next.

14. Select the checkbox for Web Server. Click Details.

15. Click Properties.

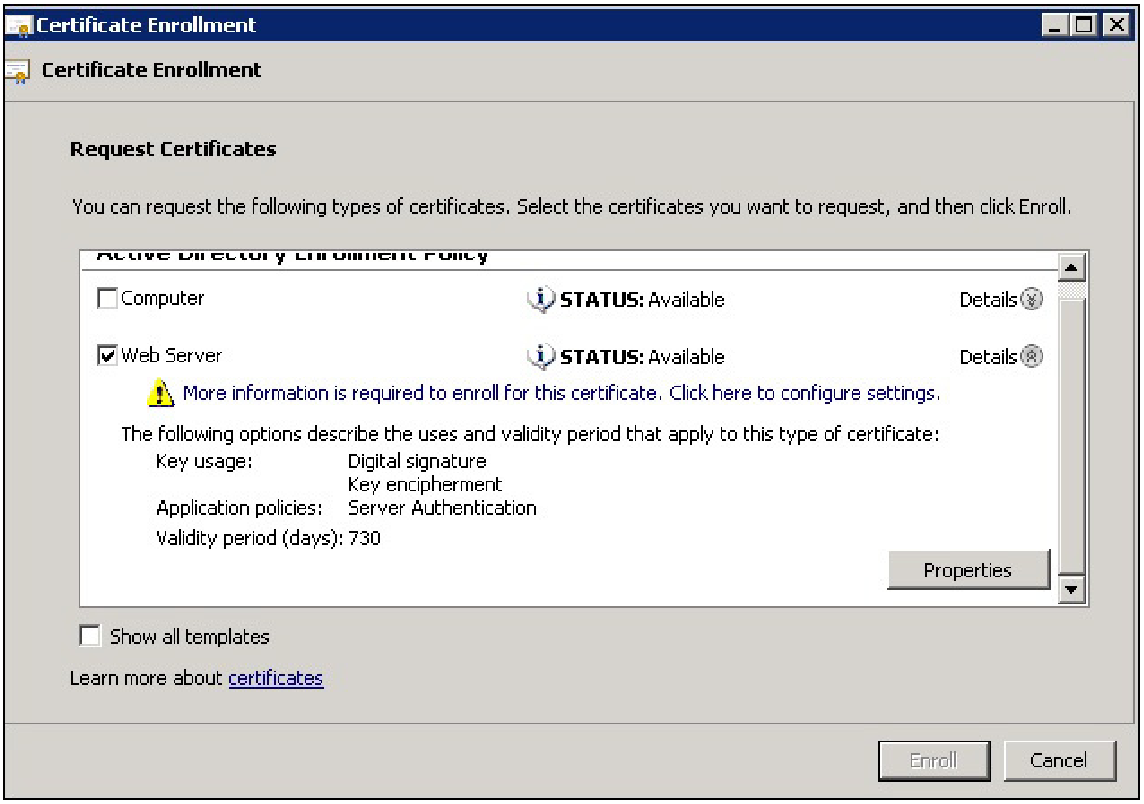

Certificate Enrollment

Certificate Enrollment

Request Certificates

You can request the following types of certificates. Select the certificates you want to request, and then click Enroll.

Active Directory Enrollment Policy

☐ Computer STATUS: Available Details

☑ Web Server STATUS: Available Details

More information is required to enroll for this certificate. Click here to configure settings.

The following options describe the uses and validity period that apply to this type of certificate:

Key usage:	Digital signature Key encipherment
Application policies:	Server Authentication
Validity period (days):	730

☐ Show all templates

Learn more about certificates

16. On the left side of the drop-down 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.

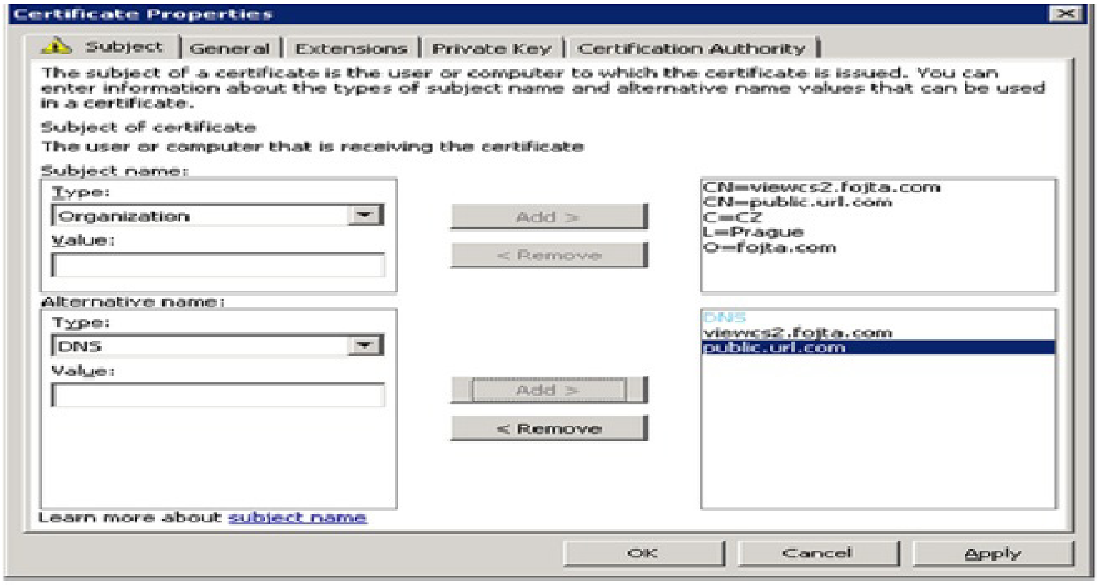

The subject of a certificate is the user or computer to which the certificate is issued. You can enter information about the types of subject name and alternative name values that can be used in a certificate.

Subject of certificate

The user or computer that is receiving the certificate

Subject name:

Type: Organization

Value:

CN=viewcs2.fojta.com
CN=public.url.com
C=CZ
L=Prague
O=fojta.com

Alternative name:

Type: DNS

Value:

DNS
viewcs2.fojta.com
public.url.com

Learn more about subject name

19. Click Apply, then click OK.

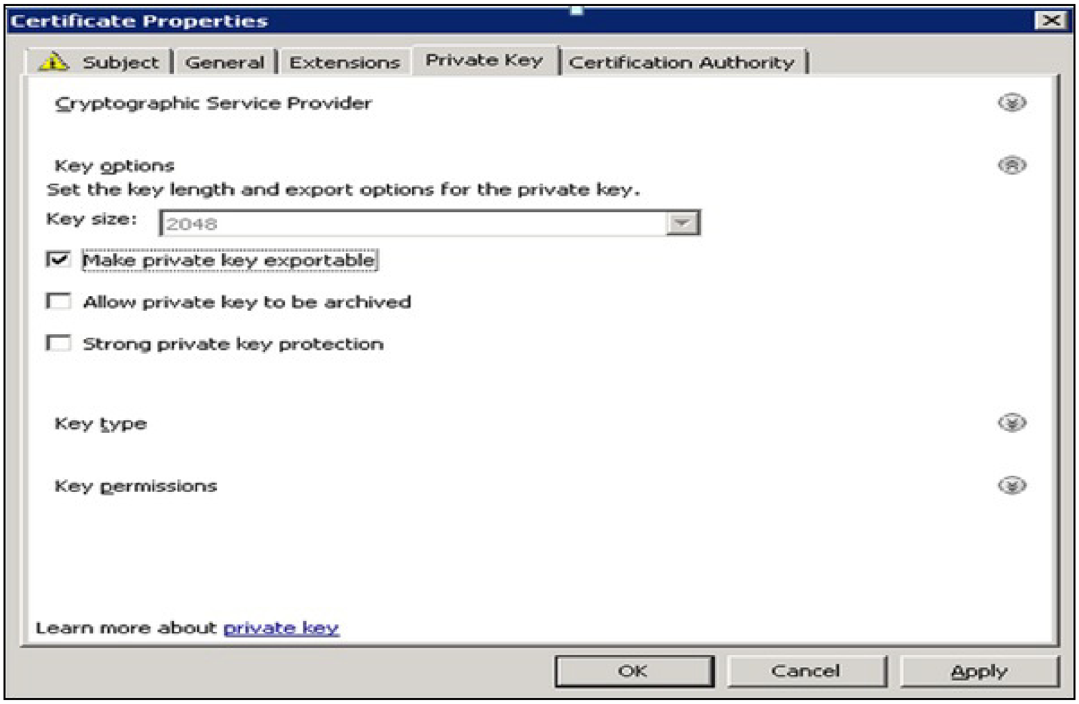

Cryptographic Service Provider

Key options

Set the key length and export options for the private key.

Key size: 2048

☑ Make private key exportable

□ Allow private key to be archived

□ Strong private key protection

Key type

Key permissions

Learn more about private key

20. Export the certificate created for the View Connection Server and Replica Server. Copy them to their corresponding server.

21. Go to the View connection server/Replica server. Start Menu > Run > mmc

22. Click File and select Add/remove Snap-in.

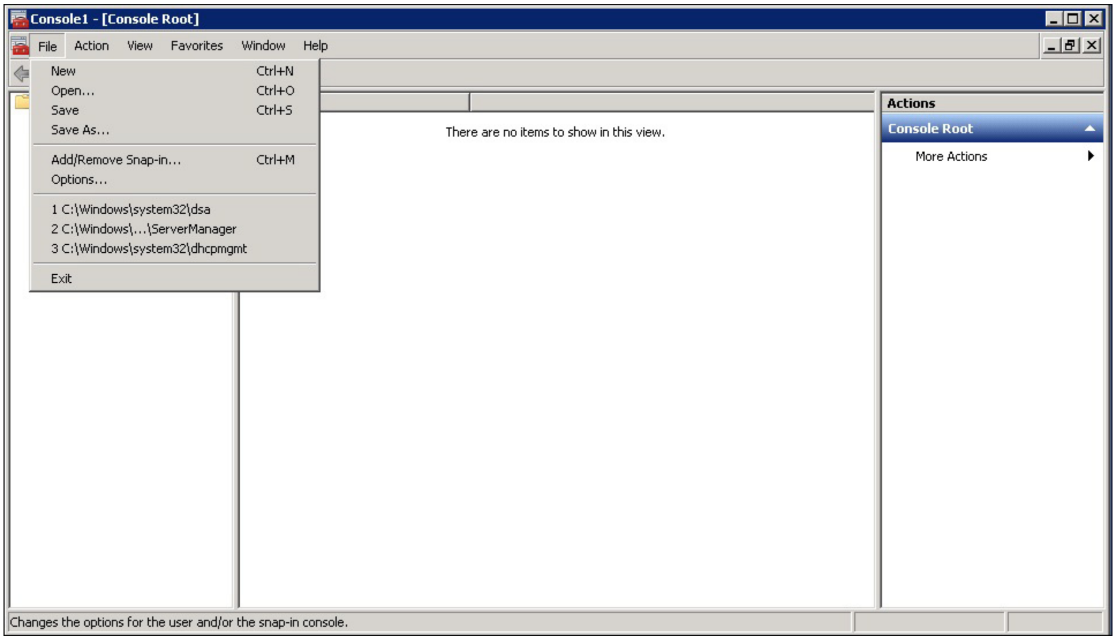
Console1 - [Console Root]
File
Action
View
Favorites
Window
Help
New Ctrl+N
Open... Ctrl+O
Save Ctrl+S
Save As...
Add/Remove Snap-in... Ctrl+M
Options...
1 C:\Windows\system32\dsa
2 C:\Windows\...\ServerManager
3 C:\Windows\system32\dhcpmgmt
Exit
Actions
Console Root
More Actions

There are no items to show in this view.

Changes the options for the user and/or the snap-in console.

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

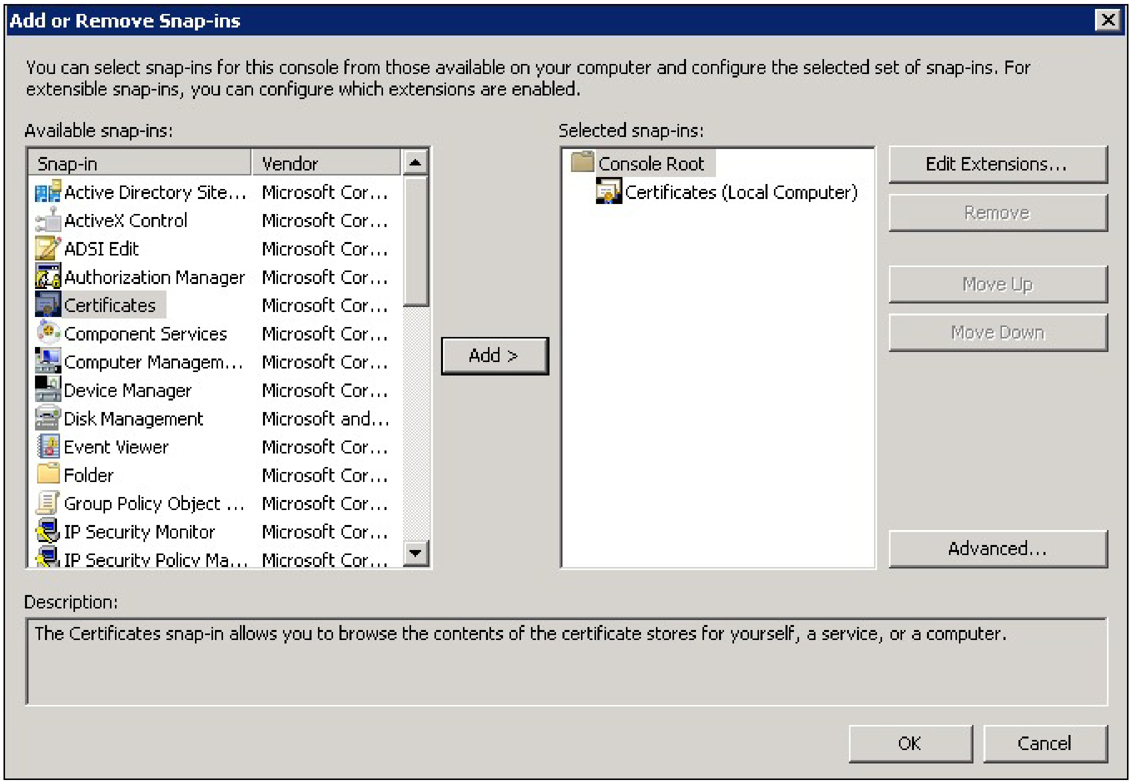

You can select snap-ins for this console from those available on your computer and configure the selected set of snap-ins. For extensible snap-ins, you can configure which extensions are enabled.

Available snap-ins:

Snap-in	Vendor
Active Directory Site...	Microsoft Cor...
ActiveX Control	Microsoft Cor...
ADSI Edit	Microsoft Cor...
Authorization Manager	Microsoft Cor...
Certificates	Microsoft Cor...
Component Services	Microsoft Cor...
Computer Managem...	Microsoft Cor...
Device Manager	Microsoft Cor...
Disk Management	Microsoft and...
Event Viewer	Microsoft Cor...
Folder	Microsoft Cor...
Group Policy Object ...	Microsoft Cor...
IP Security Monitor	Microsoft Cor...
IP Security Policy Ma...	Microsoft Cor...

Selected snap-ins:

- Console Root
  - Certificates (Local Computer)

Description:

The Certificates snap-in allows you to browse the contents of the certificate stores for yourself, a service, or a computer.

24. Select the radio button for the Computer account.

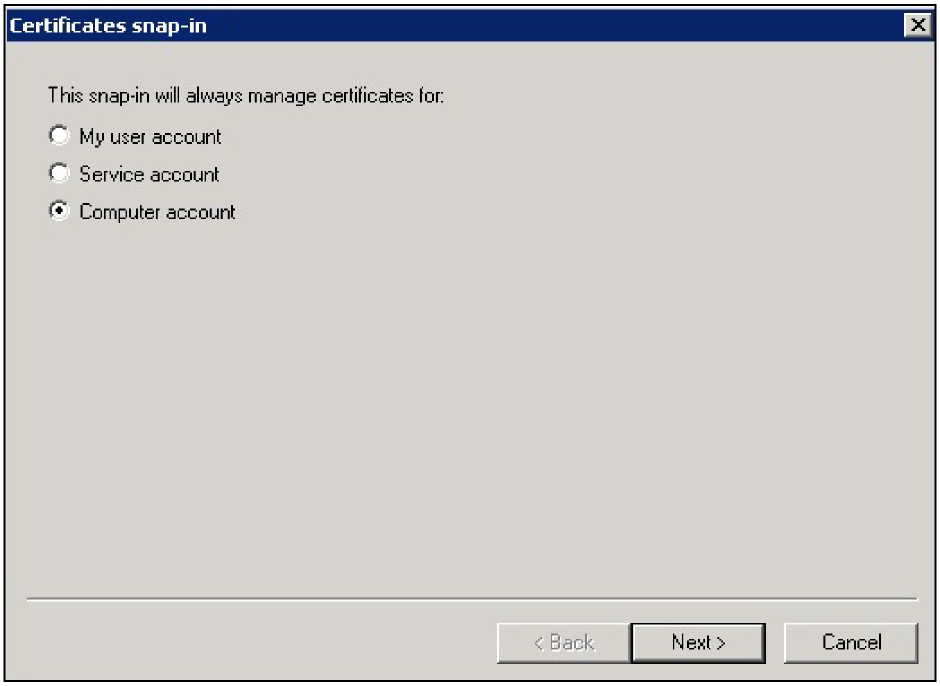

Certificates snap-in

This snap-in will always manage certificates for:

○ My user account

○ Service account

● Computer account

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

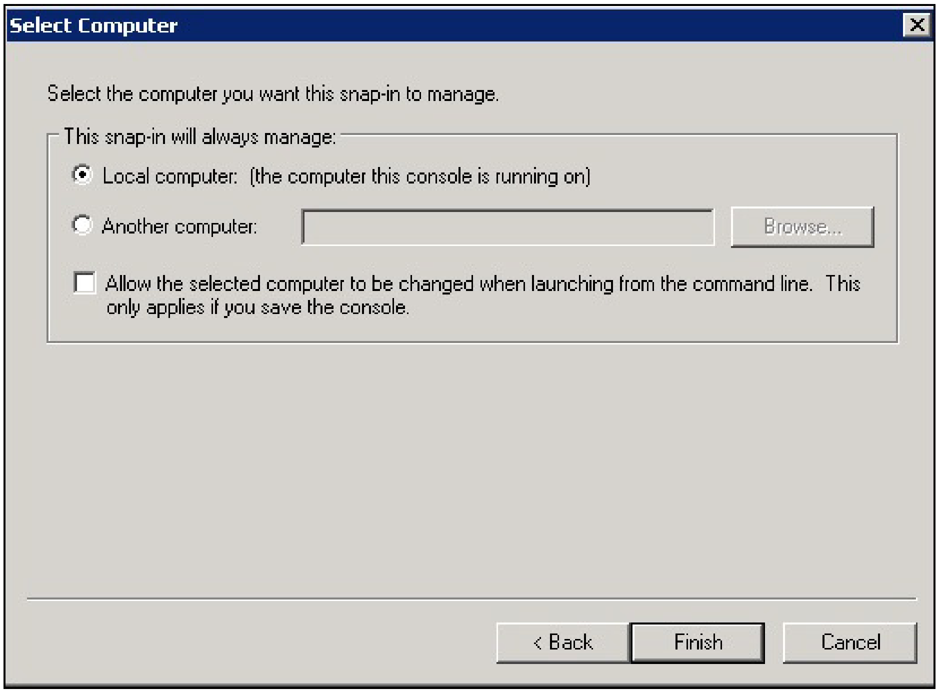

Select Computer

Select the computer you want this snap-in to manage.

This snap-in will always manage:

● Local computer: (the computer this console is running on)

○ Another computer:

□ Allow the selected computer to be changed when launching from the command line. This only applies if you save the console.

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

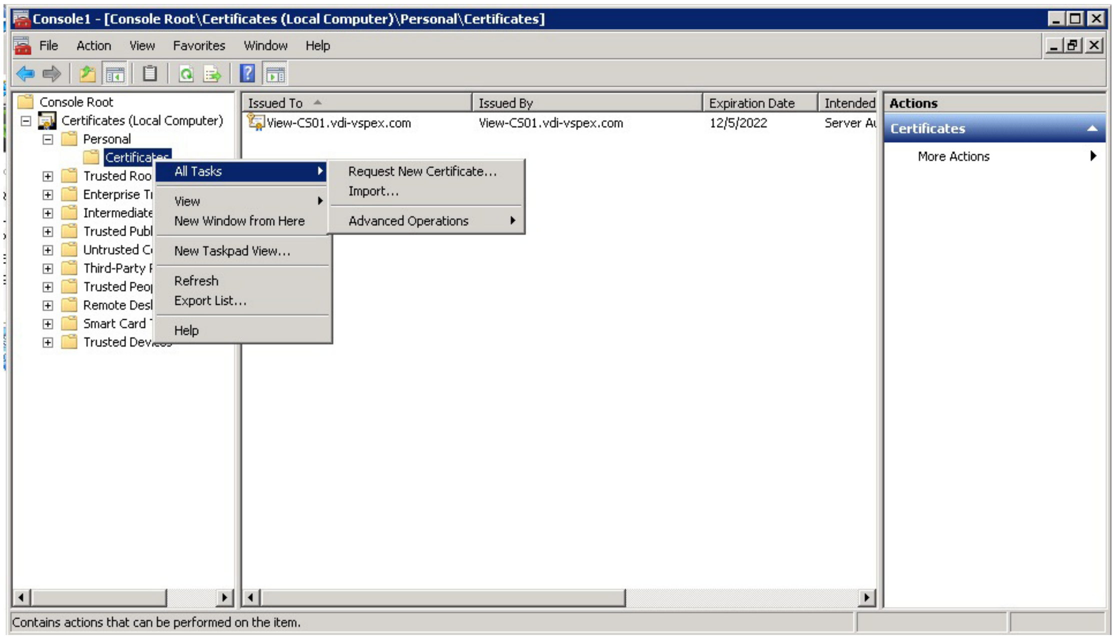

Issued To	Issued By	Expiration Date	Intended
View-CS01.vdi-vspex.com	View-CS01.vdi-vspex.com	12/5/2022	Server Au

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 friendly name. Replace the newly created certificate with vdm as the friendly name.

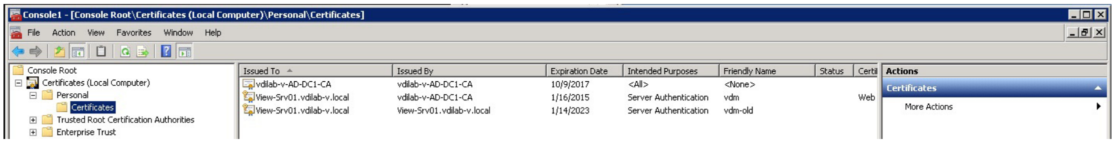
Console1 - [Console Root\Certificates (Local Computer)\Personal\Certificates]
File
Action
View
Favorites
Window
Help
Console Root
Certificates (Local Computer)
Personal
Certificates
Trusted Root Certification Authorities
Enterprise Trust
Actions
Certificates
More Actions

Issued To	Issued By	Expiration Date	Intended Purposes	Friendly Name	Status	Certi
vdilab-v-AD-DC1-CA	vdilab-v-AD-DC1-CA	10/9/2017	<All>	<None>		
View-Srv01.vdilab-v.local	vdilab-v-AD-DC1-CA	1/16/2015	Server Authentication	vdm		Web
View-Srv01.vdilab-v.local	View-Srv01.vdilab-v.local	1/14/2023	Server Authentication	vdm-old		

Install View Client on End Points.

1. Download installer file from the link below:

   https://my.vmware.com/web/vmware/info/slug/desktop_end_user_computing/vmware_view/5_1

2. Open installer file for 32-bit or 64-bit OS, Click Next on Installation wizard.

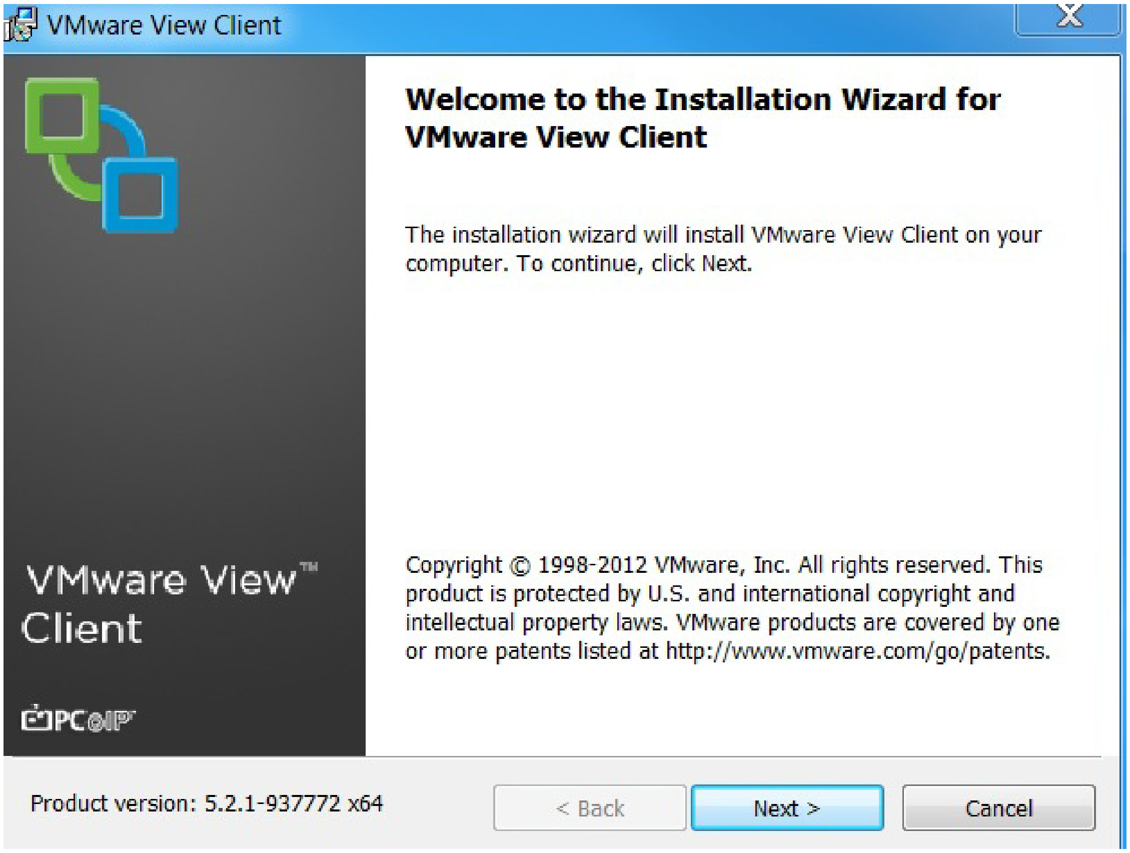

Welcome to the Installation Wizard for VMware View Client

The installation wizard will install VMware View Client on your computer. To continue, click Next.

Copyright © 1998-2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 5.2.1-937772 x64

3. Accept the License agreement. Click Next.
4. Accept the default features and click Next.

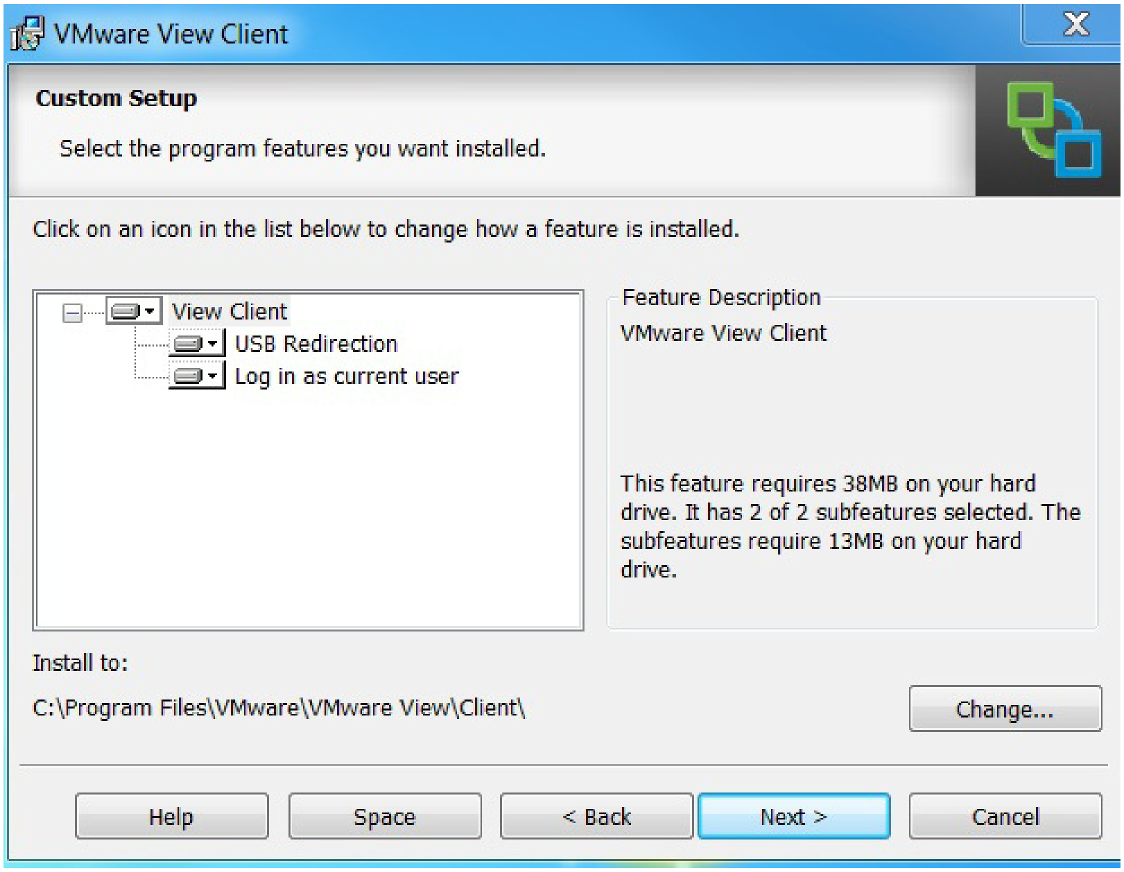

VMware View Client

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- View Client
  - USB Redirection
  - Log in as current user

Feature Description

VMware View Client

This feature requires 38MB on your hard drive. It has 2 of 2 subfeatures selected. The subfeatures require 13MB on your hard drive.

Install to:

C:\Program Files\VMware\VMware View\Client\

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

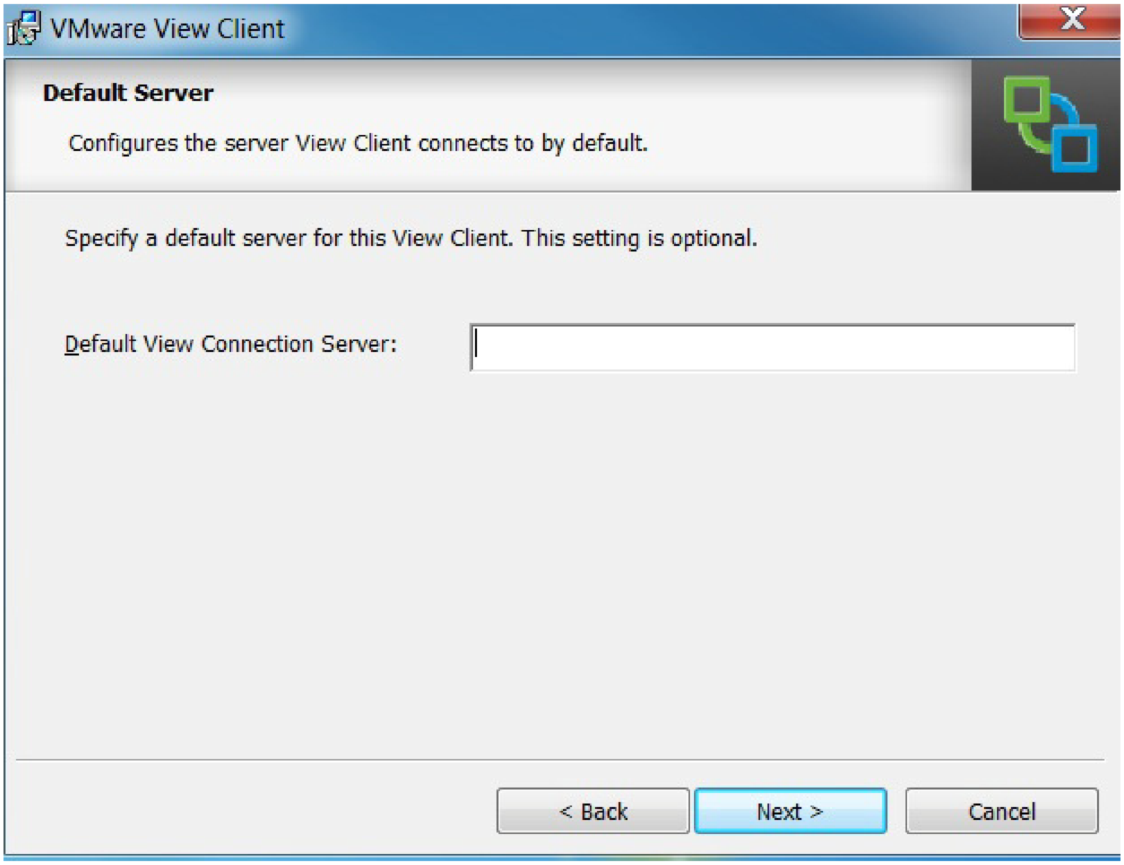

VMware View Client

Default Server

Configures the server View Client connects to by default.

Specify a default server for this View Client. This setting is optional.

Default View Connection Server:

6. Click Next.

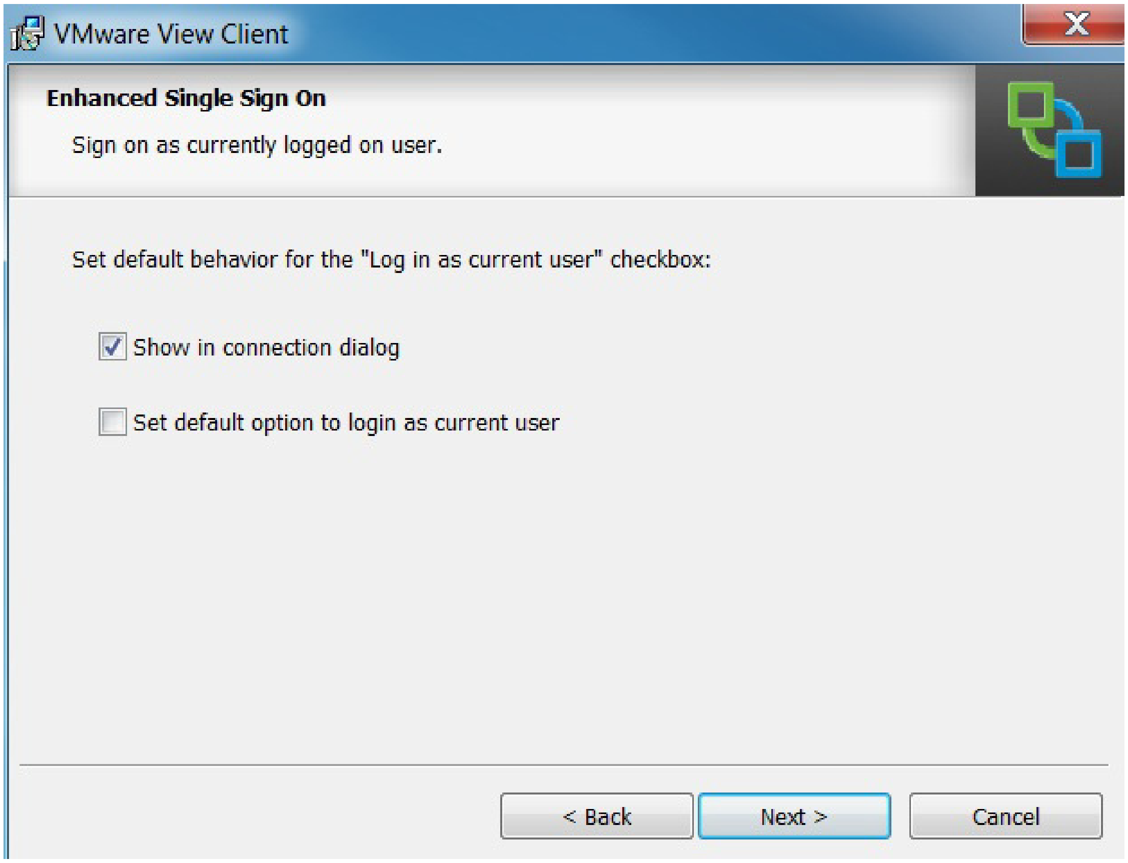

VMware View Client

Enhanced Single Sign On

Sign on as currently logged on user.

Set default behavior for the "Log in as current user" checkbox:

☑ Show in connection dialog

□ Set default option to login as current user

7. Click Next.

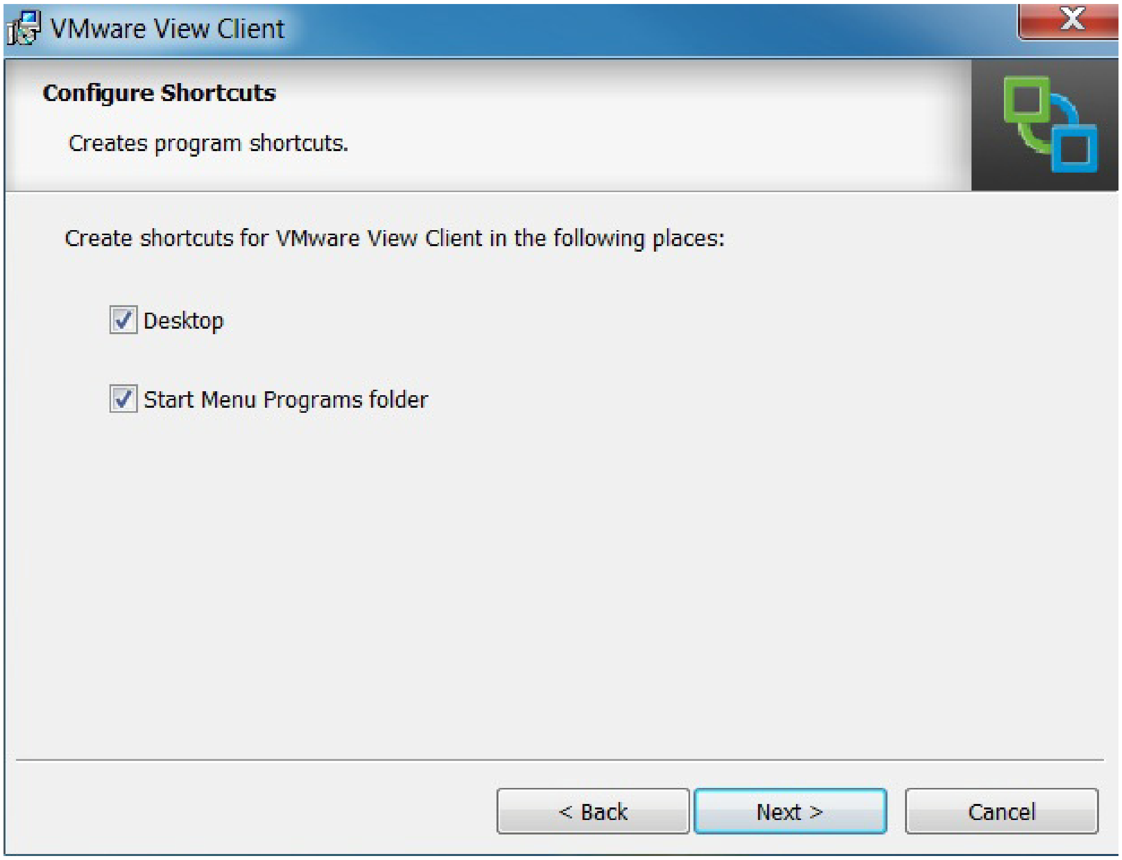

VMware View Client

Configure Shortcuts

Creates program shortcuts.

Create shortcuts for VMware View Client in the following places:

☑ Desktop

☑ Start Menu Programs folder

8. Click Install.

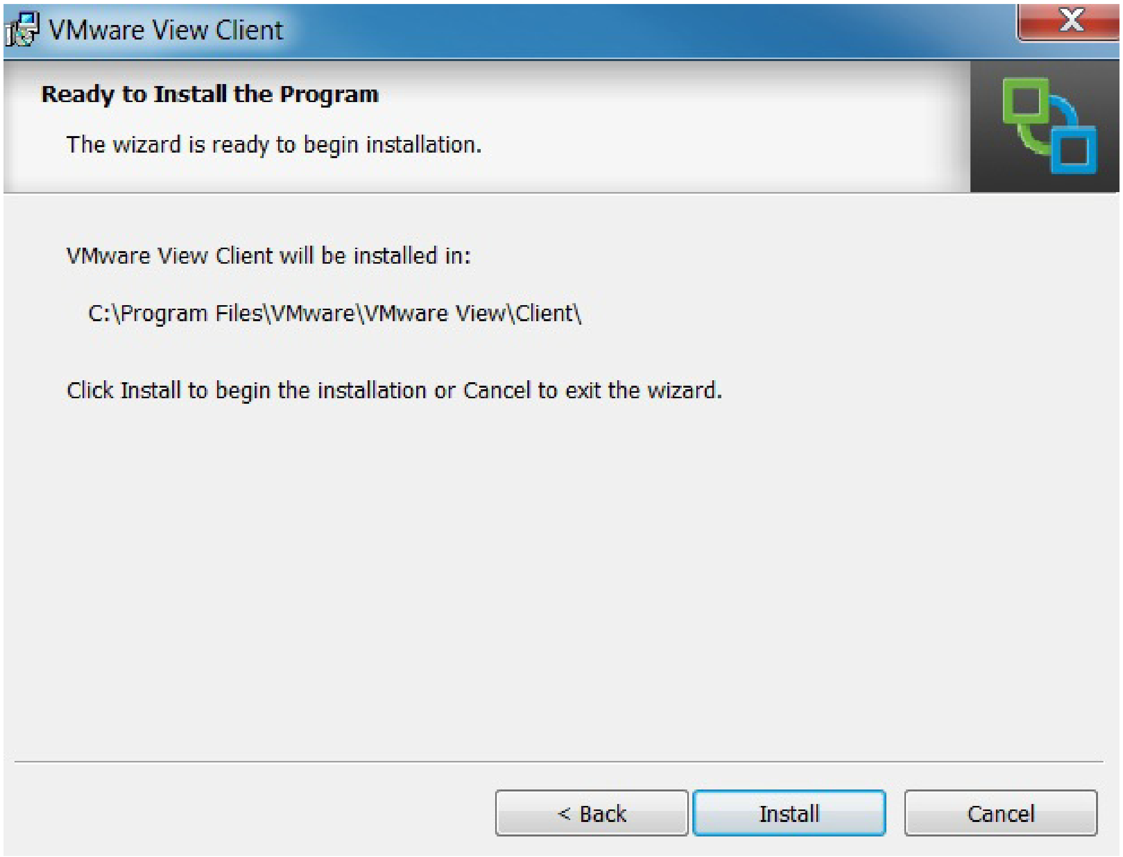

VMware View Client

Ready to Install the Program

The wizard is ready to begin installation.

VMware View Client will be installed in:

C:\Program Files\VMware\VMware View\Client\

Click Install to begin the installation or Cancel to exit the wizard.

9. Reboot.

Configure the View 5.1 Hosts and Storage

Configure Content Based Read Cache (CBRC) on View 5.1 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 5.1 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, Advanced Settings.

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

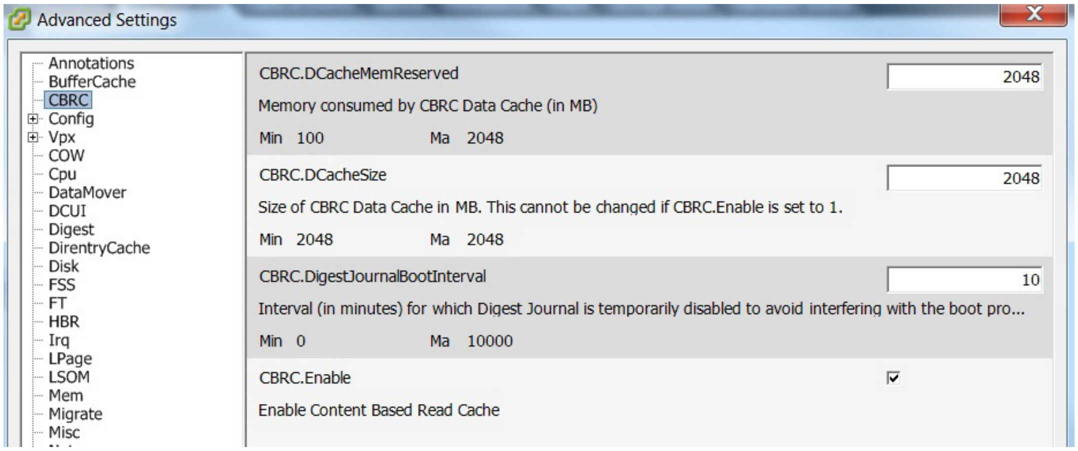

CBRC.DCacheMemReserved 2048

Memory consumed by CBRC Data Cache (in MB)

Min 100 Ma 2048

CBRC.DCacheSize 2048

Size of CBRC Data Cache in MB. This cannot be changed if CBRC.Enable is set to 1.

Min 2048 Ma 2048

CBRC.DigestJournalBootInterval 10

Interval (in minutes) for which Digest Journal is temporarily disabled to avoid interfering with the boot pro...

Min 0 Ma 10000

CBRC.Enable ☑

Enable Content Based Read Cache

These CBRC settings are used in conjunction with the View 5.1 Administrator, View Configuration, Servers, vCenter Server Properties, Host Cashing tab.

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

Storage Configuration for View 5.1 Hosts

On VNX 5300 10 SAS disks with 600Gb capacity were used to create 4 LUNs, each with a capacity of 485 GB. One LUN with capacity of 50 GB was created to store replica disks.

Each ESXi host in cluster was assigned 4 LUNs as VMFS5 datastores for linked clones and one 50Gb VMFS5 datastore to hold the Replica disk.

Configure the View Desktop Pools and Options

Desktop Pools are the containment object in View 5.1 Administrator that hold the configuration and the provisioned linked clones in the View environment.

The following sections describe how we configured our View 5.1 environment.

Create the Desktop Pools

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

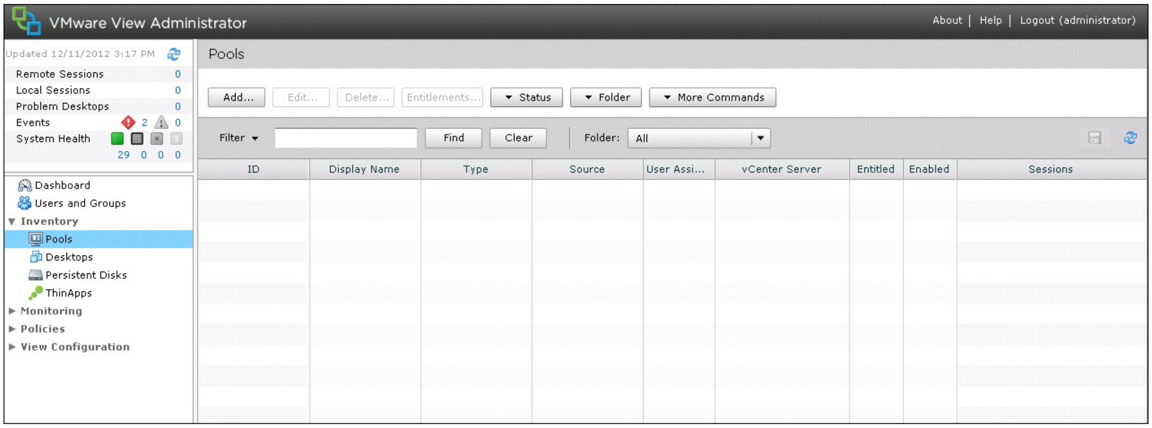

There are three types of Desktop Pools that we can create and the description for each type is located on the right side of the screen.

2. For our testing purpose, we created Automated Pool. Click Next.

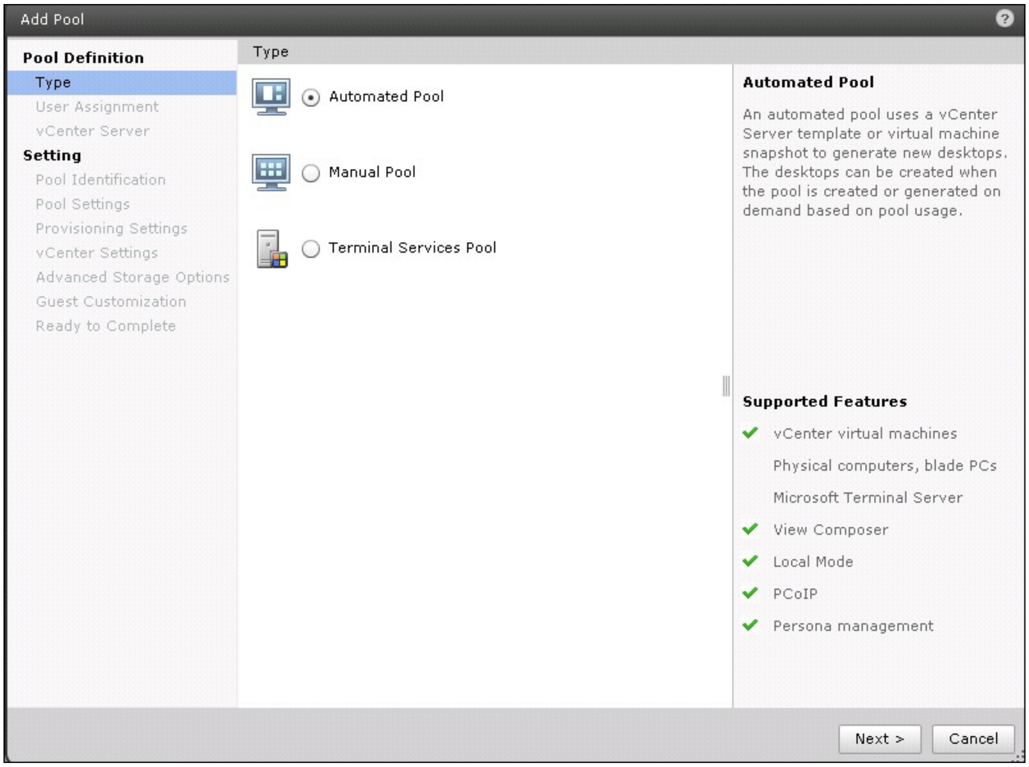

● Automated Pool

○ Manual Pool

○ Terminal Services Pool

Automated Pool

An automated pool uses a vCenter Server template or virtual machine snapshot to generate new desktops. The desktops can be created when the pool is created or generated on demand based on pool usage.

Supported Features

- ✓ vCenter virtual machines
- Physical computers, blade PCs
- Microsoft Terminal Server
- ✓ View Composer
- ✓ Local Mode
- ✓ PCoIP
- ✓ Persona management

User assignment options are available:

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

For our test we used Floating user assignments.

3. Click the Floating radio button, then click Next.

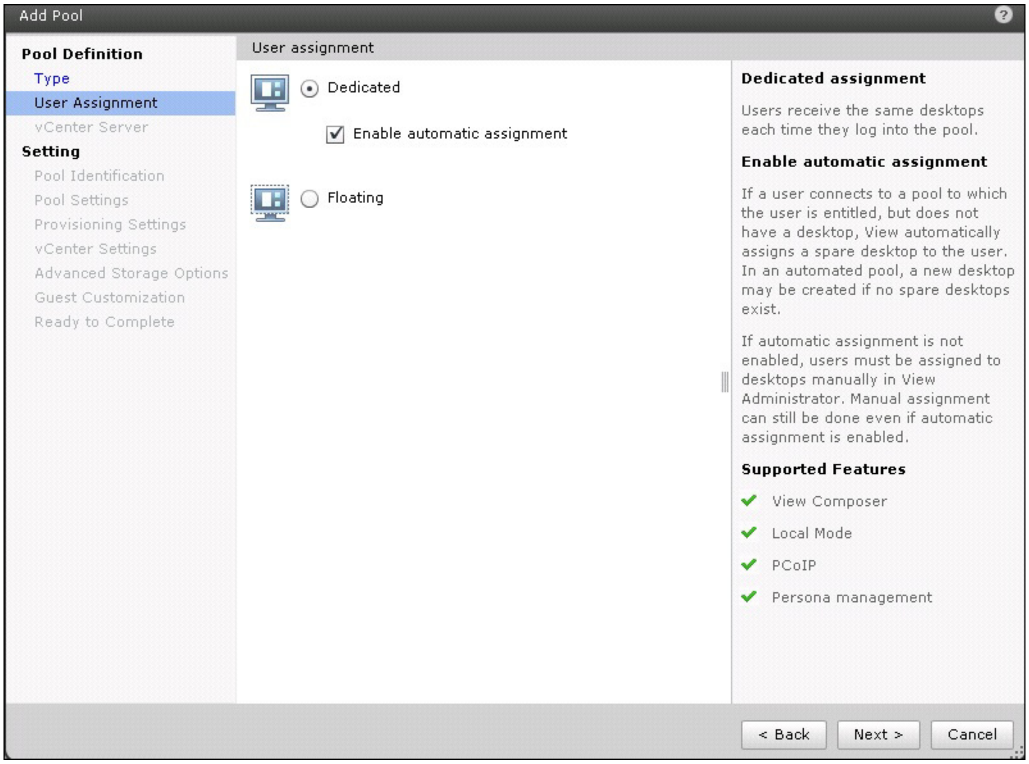

Add Pool

User assignment

● Dedicated

☑ Enable automatic assignment

○ Floating

Dedicated assignment

Users receive the same desktops each time they log into the pool.

Enable automatic assignment

If a user connects to a pool to which the user is entitled, but does not have a desktop, View automatically assigns a spare desktop to the user. In an automated pool, a new desktop may be created if no spare desktops exist.

If automatic assignment is not enabled, users must be assigned to desktops manually in View Administrator. Manual assignment can still be done even if automatic assignment is enabled.

Supported Features

✓ View Composer

✓ Local Mode

✓ PCoIP

✓ Persona management

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

Note For this study, we chose View Composer linked clones

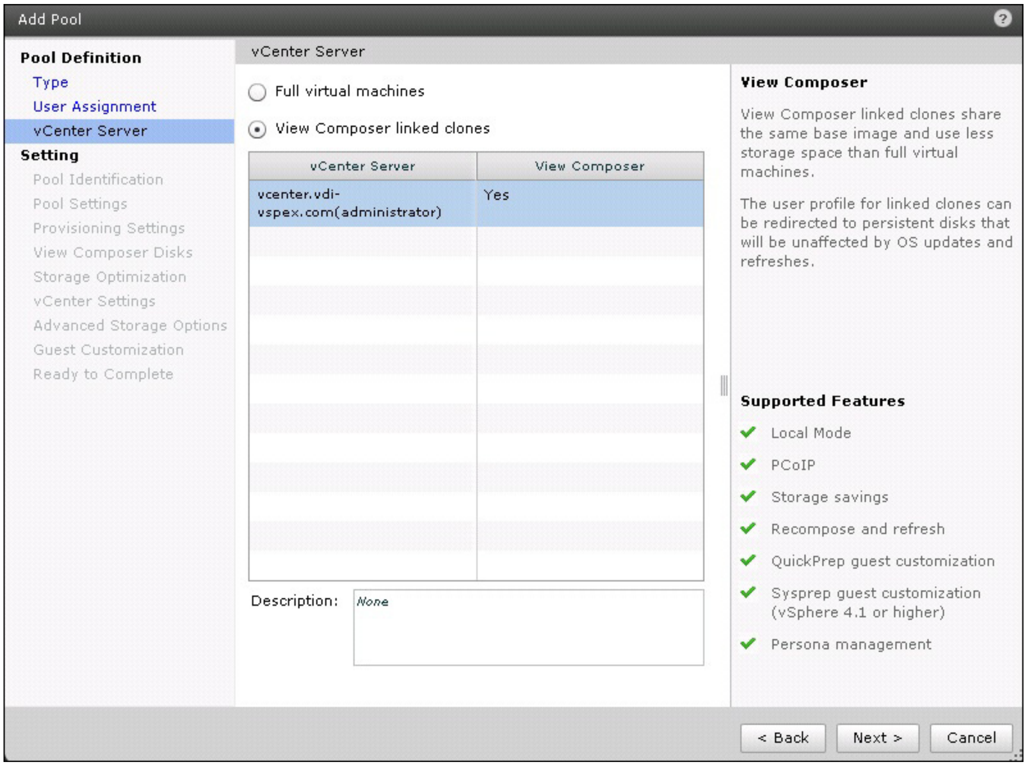

vCenter Server

○ Full virtual machines

● View Composer linked clones

vCenter Server	View Composer
vcenter.vdi-vspex.com(administrator)	Yes

Description: *None*

View Composer

View Composer linked clones share the same base image and use less storage space than full virtual machines.

The user profile for linked clones can be redirected to persistent disks that will be unaffected by OS updates and refreshes.

Supported Features

- ✓ Local Mode
- ✓ PCoIP
- ✓ Storage savings
- ✓ Recompose and refresh
- ✓ QuickPrep guest customization
- ✓ Sysprep guest customization (vSphere 4.1 or higher)
- ✓ Persona management

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

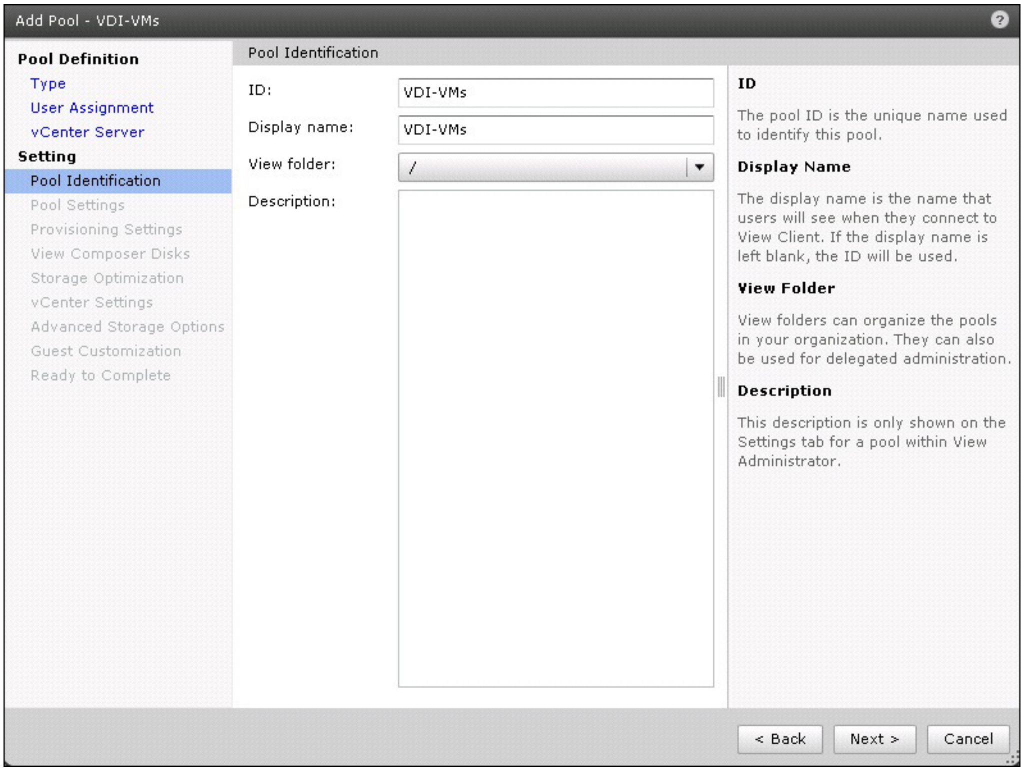

Pool Identification

ID: VDI-VMs

Display name: VDI-VMs

View folder: /

Description:

ID

The pool ID is the unique name used to identify this pool.

Display Name

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

View Folder

View folders can organize the pools in your organization. They can also be used for delegated administration.

Description

This description is only shown on the Settings tab for a pool within View Administrator.

6. Configure the Pool Settings as needed. We selected all default settings except the Remote Desktop Power Policy. We selected Ensure desktops are always powered on. Click next.

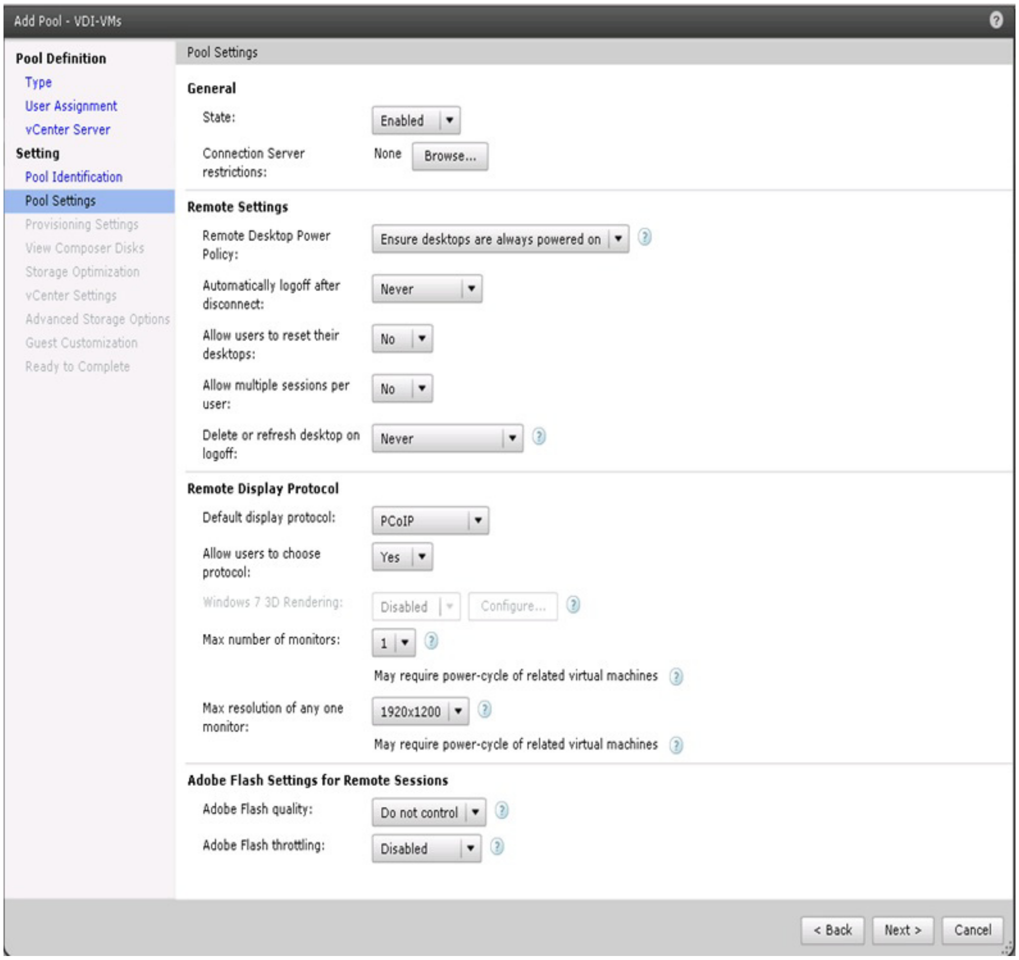

Add Pool - VDI-VMs

Pool Settings

General

State: Enabled

Connection Server restrictions: None Browse...

Remote Settings

Remote Desktop Power Policy: Ensure desktops are always powered on

Automatically logoff after disconnect: Never

Allow users to reset their desktops: No

Allow multiple sessions per user: No

Delete or refresh desktop on logoff: Never

Remote Display Protocol

Default display protocol: PCoIP

Allow users to choose protocol: Yes

Windows 7 3D Rendering: Disabled Configure...

Max number of monitors: 1

May require power-cycle of related virtual machines

Max resolution of any one monitor: 1920x1200

May require power-cycle of related virtual machines

Adobe Flash Settings for Remote Sessions

Adobe Flash quality: Do not control

Adobe Flash throttling: Disabled

7. On the Provisioning Settings page, we 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

For this study, we provisioned all of the desktops up-front.

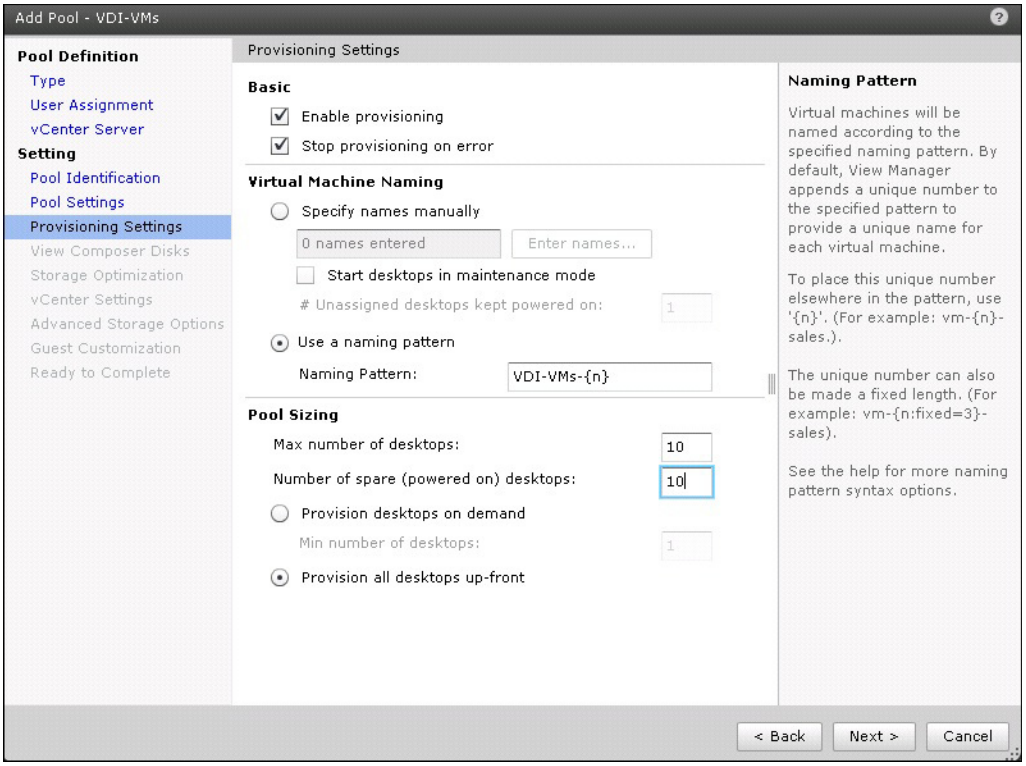

Add Pool - VDI-VMs

Pool Definition
- Type
- User Assignment
- vCenter Server

Setting
- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

Provisioning Settings

Basic

☑ Enable provisioning

☑ Stop provisioning on error

Virtual Machine Naming

○ Specify names manually

0 names entered | Enter names...

□ Start desktops in maintenance mode

# Unassigned desktops kept powered on: 1

● Use a naming pattern

Naming Pattern: VDI-VMs-{n}

Pool Sizing

Max number of desktops: 10

Number of spare (powered on) desktops: 10

○ Provision desktops on demand

Min number of desktops: 1

● Provision all desktops up-front

Naming Pattern

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.).

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales).

See the help for more naming pattern syntax options.

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.

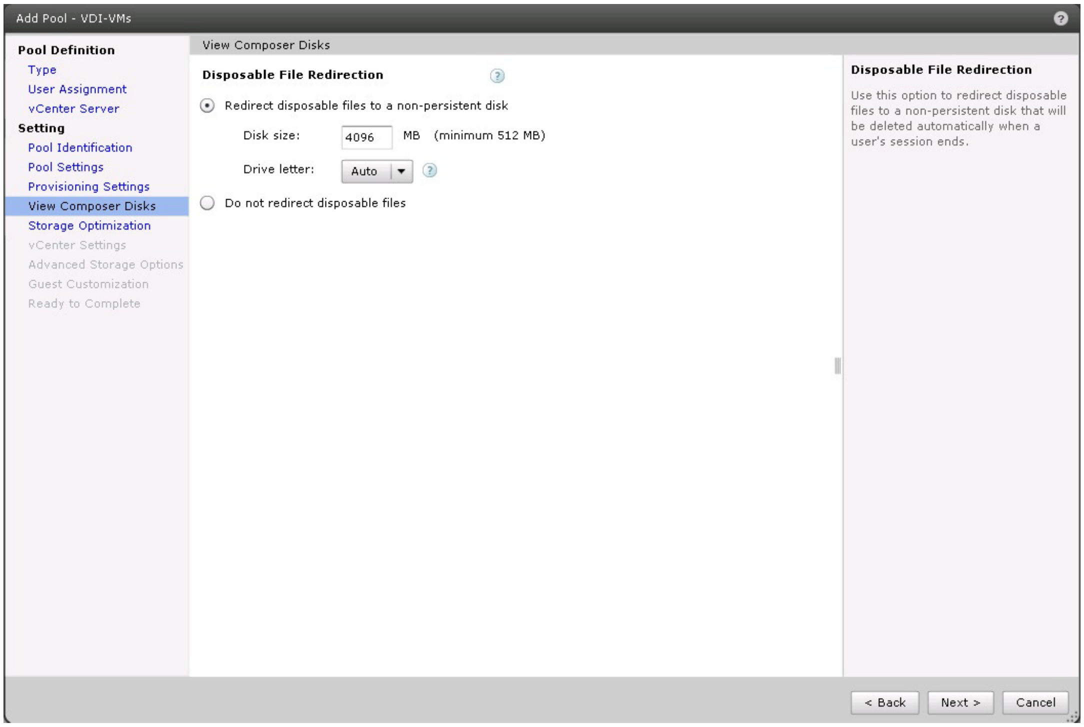

Add Pool - VDI-VMs

Pool Definition
- Type
- User Assignment
- vCenter Server

Setting
- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

View Composer Disks

Disposable File Redirection

● Redirect disposable files to a non-persistent disk

Disk size: 4096 MB (minimum 512 MB)

Drive letter: Auto

○ Do not redirect disposable files

Disposable File Redirection

Use this option to redirect disposable files to a non-persistent disk that will be deleted automatically when a user's session ends.

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

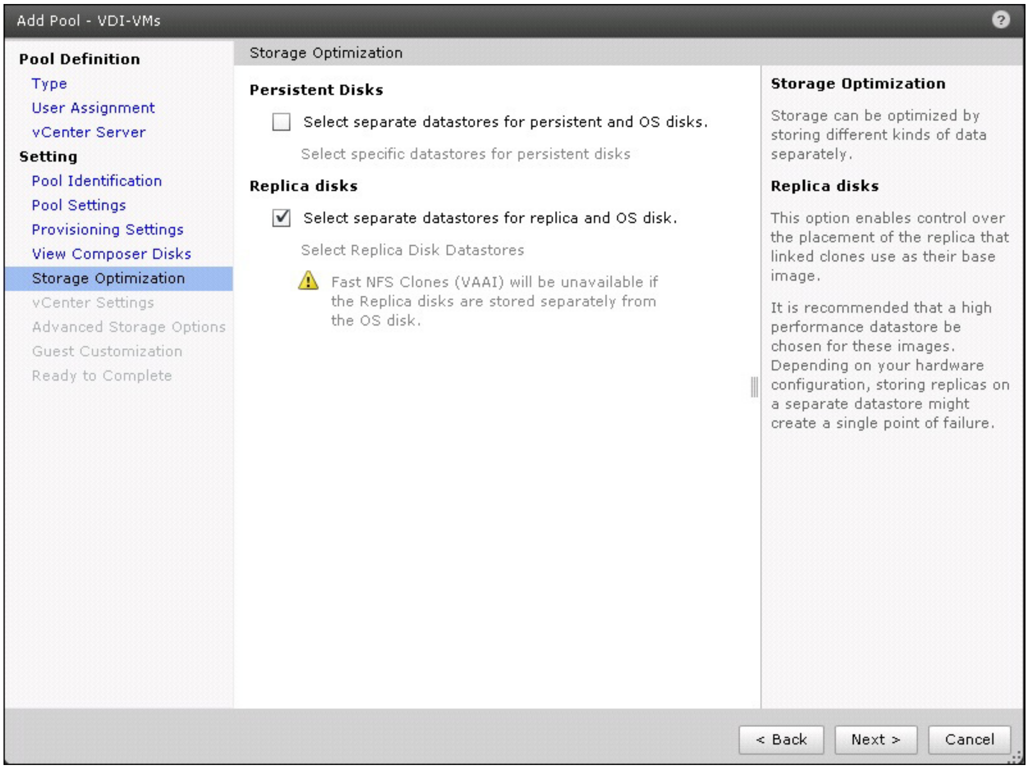

Storage Optimization

Persistent Disks

□ Select separate datastores for persistent and OS disks.

Select specific datastores for persistent disks

Replica disks

☑ Select separate datastores for replica and OS disk.

Select Replica Disk Datastores

⚠ Fast NFS Clones (VAAI) will be unavailable if the Replica disks are stored separately from the OS disk.

Storage Optimization

Storage can be optimized by storing different kinds of data separately.

Replica disks

This option enables control over the placement of the replica that linked clones use as their base image.

It is recommended that a high performance datastore be chosen for these images. Depending on your hardware configuration, storing replicas on a separate datastore might create a single point of failure.

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.

For our testing we created 1 Pool with 600 desktops. One Pool was created with VDI-v500 as host Resource Pool. 4 VMFS5 datastores for Linked Clones and one Replica disk datastore were selected

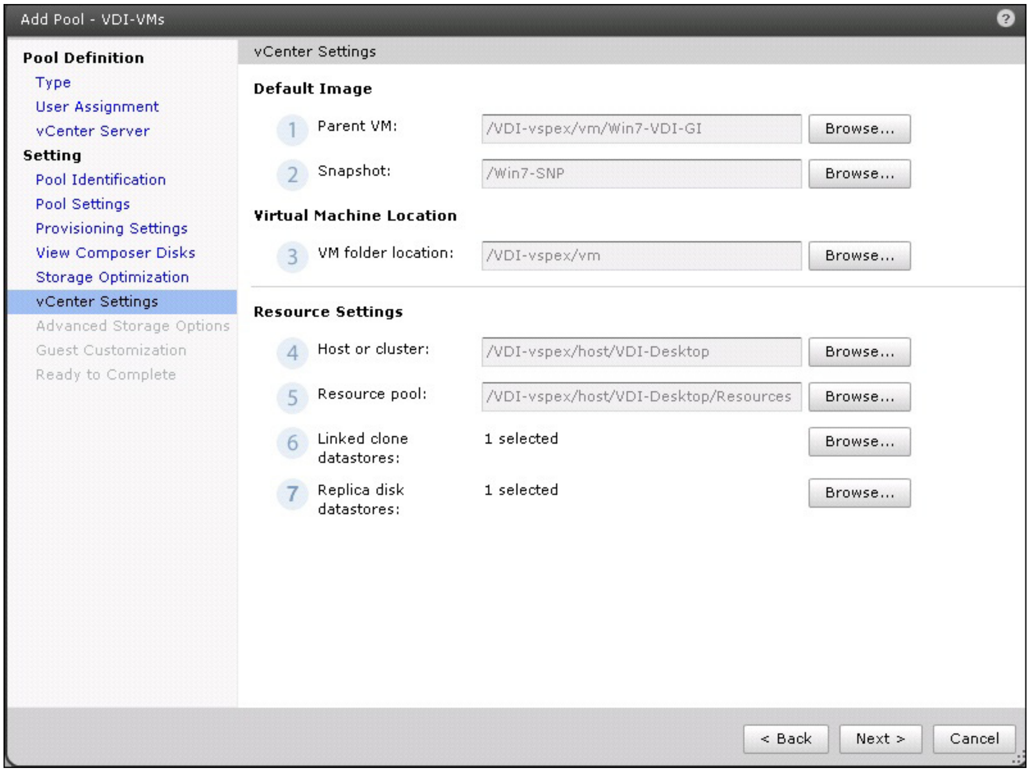

vCenter Settings

Default Image

1 Parent VM: /VDI-vspex/vm/Win7-VDI-GI

2 Snapshot: /Win7-SNP

Virtual Machine Location

3 VM folder location: /VDI-vspex/vm

Resource Settings

4 Host or cluster: /VDI-vspex/host/VDI-Desktop

5 Resource pool: /VDI-vspex/host/VDI-Desktop/Resources

6 Linked clone datastores: 1 selected

7 Replica disk datastores: 1 selected

11. On the Advanced Storage Options page, check box to enable Use Host caching. Add Blackout time if desired. Click Next.

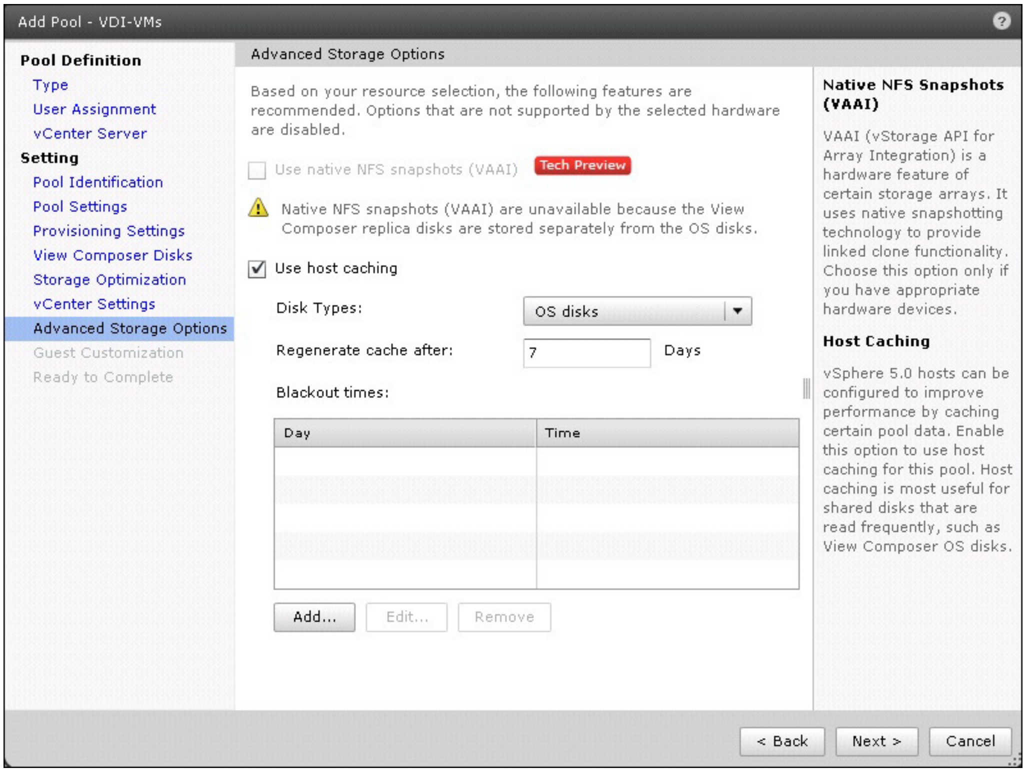

Advanced Storage Options

Based on your resource selection, the following features are recommended. Options that are not supported by the selected hardware are disabled.

☐ Use native NFS snapshots (VAAI)

⚠ Native NFS snapshots (VAAI) are unavailable because the View Composer replica disks are stored separately from the OS disks.

☑ Use host caching

Disk Types: OS disks

Regenerate cache after: 7 Days

Blackout times:

Day	Time

Native NFS Snapshots (VAAI)

VAAI (vStorage API for Array Integration) is a hardware feature of certain storage arrays. It uses native snapshotting technology to provide linked clone functionality. Choose this option only if you have appropriate hardware devices.

Host Caching

vSphere 5.0 hosts can be configured to improve performance by caching certain pool data. Enable this option to use host caching for this pool. Host caching is most useful for shared disks that are read frequently, such as View Composer OS disks.

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

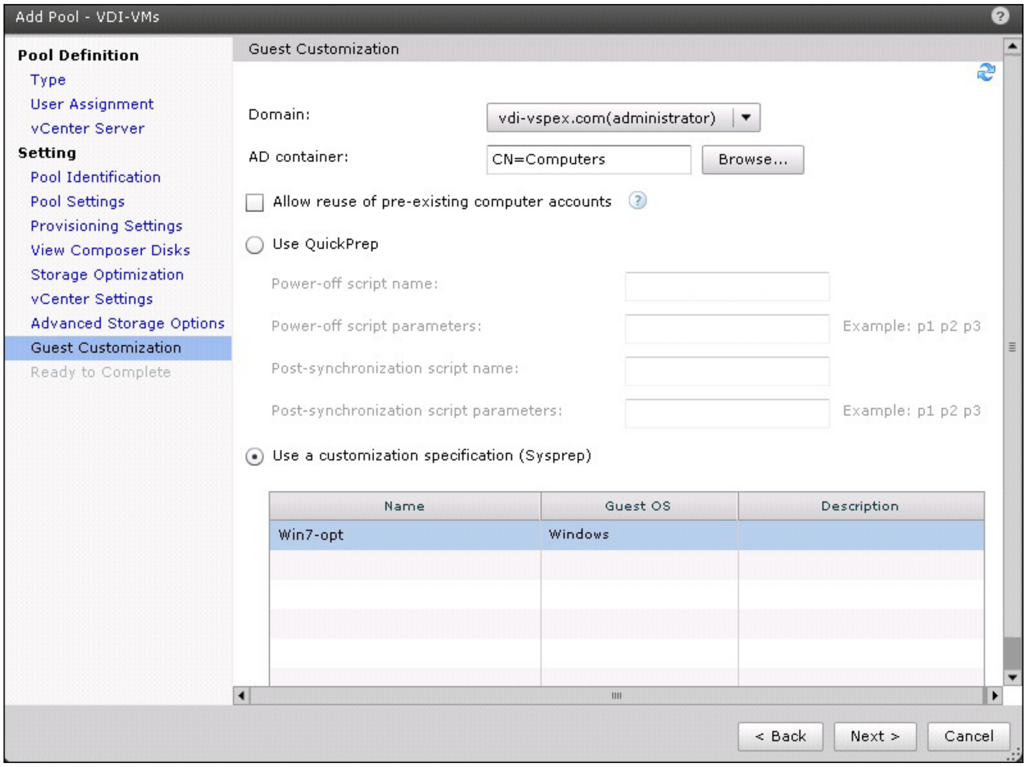

Domain: vdi-vspex.com(administrator)

AD container: CN=Computers Browse...

□ Allow reuse of pre-existing computer accounts

○ Use QuickPrep

Power-off script name:

Power-off script parameters: Example: p1 p2 p3

Post-synchronization script name:

Post-synchronization script parameters: Example: p1 p2 p3

● Use a customization specification (Sysprep)

Name	Guest OS	Description
Win7-opt	Windows	

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.

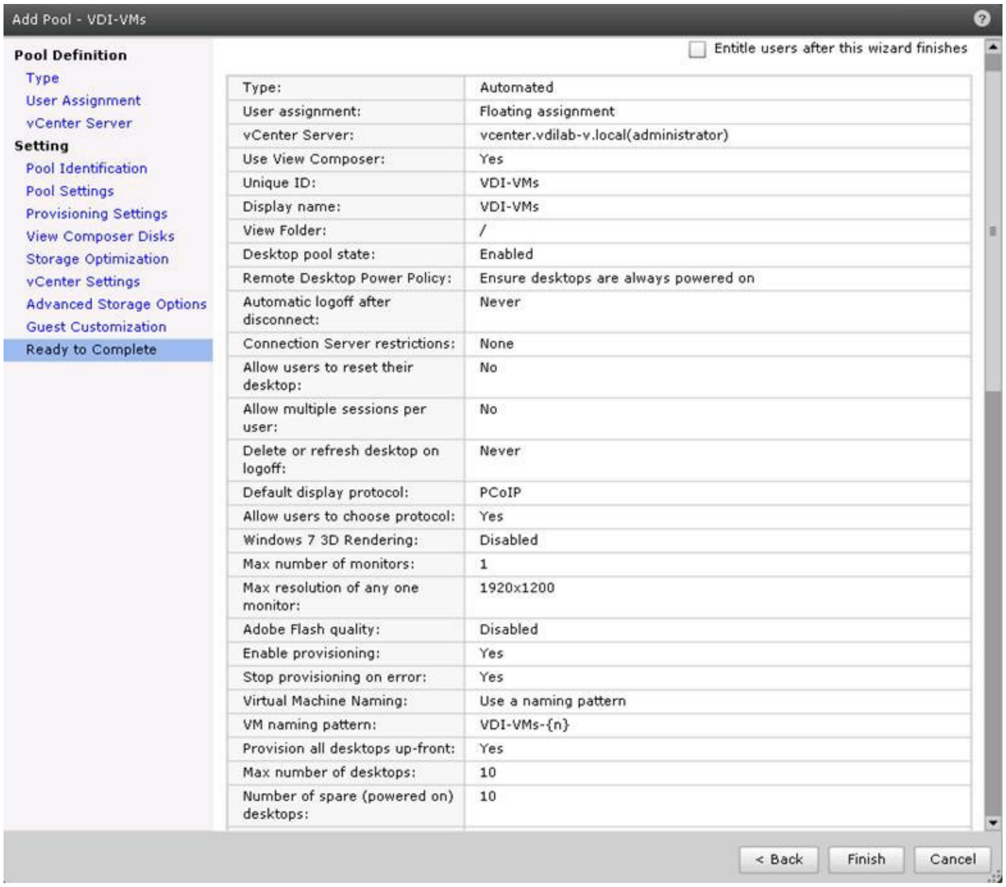

□ Entitle users after this wizard finishes

Type:	Automated
User assignment:	Floating assignment
vCenter Server:	vcenter.vdilab-v.local(administrator)
Use View Composer:	Yes
Unique ID:	VDI-VMs
Display name:	VDI-VMs
View Folder:	/
Desktop pool state:	Enabled
Remote Desktop Power Policy:	Ensure desktops are always powered on
Automatic logoff after disconnect:	Never
Connection Server restrictions:	None
Allow users to reset their desktop:	No
Allow multiple sessions per user:	No
Delete or refresh desktop on logoff:	Never
Default display protocol:	PCoIP
Allow users to choose protocol:	Yes
Windows 7 3D Rendering:	Disabled
Max number of monitors:	1
Max resolution of any one monitor:	1920x1200
Adobe Flash quality:	Disabled
Enable provisioning:	Yes
Stop provisioning on error:	Yes
Virtual Machine Naming:	Use a naming pattern
VM naming pattern:	VDI-VMs-{n}
Provision all desktops up-front:	Yes
Max number of desktops:	10
Number of spare (powered on) desktops:	10

14. On the Entitlements page, Click on Add.

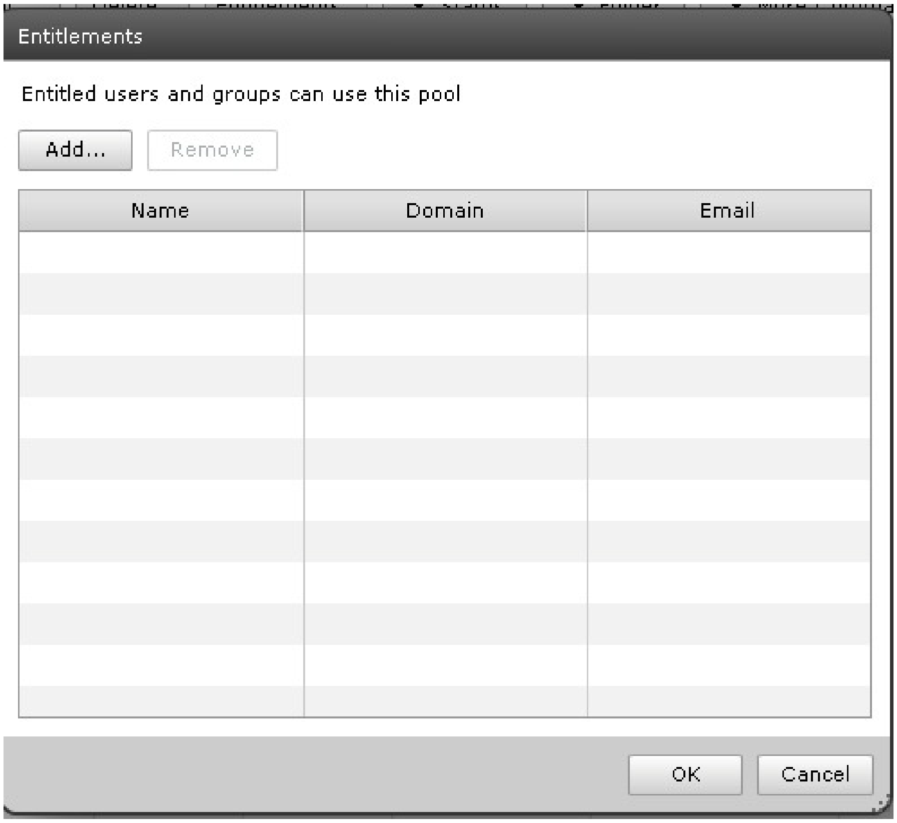

Entitlements

Entitled users and groups can use this pool

Name	Domain	Email

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

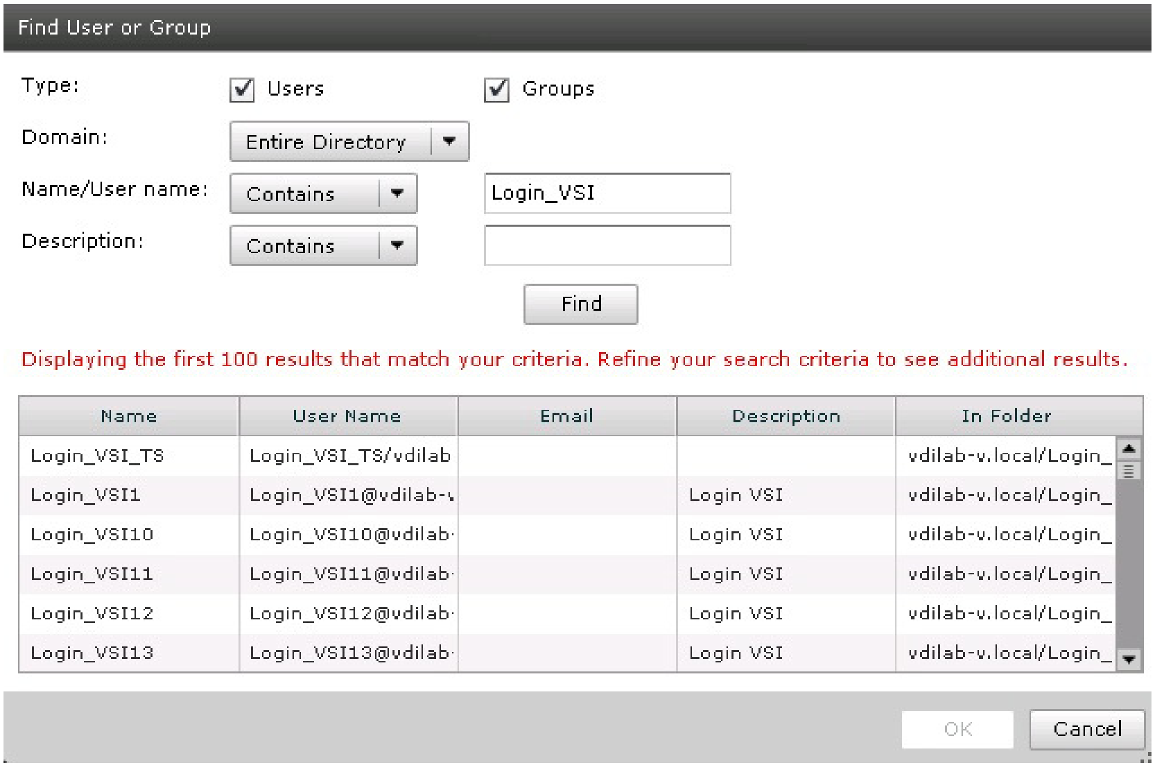

Find User or Group

Type: ☑ Users ☑ Groups

Domain: Entire Directory

Name/User name: Contains Login_VSI

Description: Contains

Displaying the first 100 results that match your criteria. Refine your search criteria to see additional results.

Name	User Name	Email	Description	In Folder
Login_VSI_TS	Login_VSI_TS/vdilab			vdilab-v.local/Login_
Login_VSI1	Login_VSI1@vdilab-v		Login VSI	vdilab-v.local/Login_
Login_VSI10	Login_VSI10@vdilab-		Login VSI	vdilab-v.local/Login_
Login_VSI11	Login_VSI11@vdilab-		Login VSI	vdilab-v.local/Login_
Login_VSI12	Login_VSI12@vdilab-		Login VSI	vdilab-v.local/Login_
Login_VSI13	Login_VSI13@vdilab-		Login VSI	vdilab-v.local/Login_

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

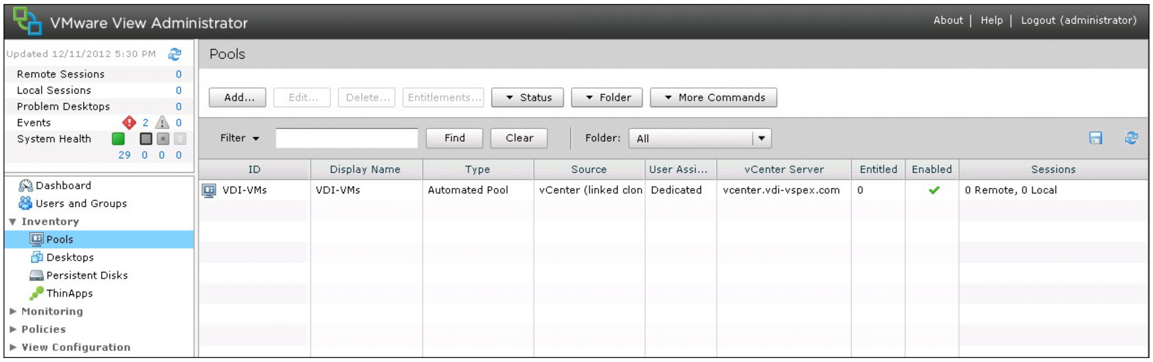

Remote Sessions	0
Local Sessions	0
Problem Desktops	0
Events	2 0
System Health	29 0 0 0

Pools

Folder: All

ID	Display Name	Type	Source	User Assi...	vCenter Server	Entitled	Enabled	Sessions
VDI-VMs	VDI-VMs	Automated Pool	vCenter (linked clon	Dedicated	vcenter.vdi-vspex.com	0	✓	0 Remote, 0 Local

18. After completing the pool setting, a replica from the parent VM is created and View Composer starts provisioning of the desktops per the pool settings.

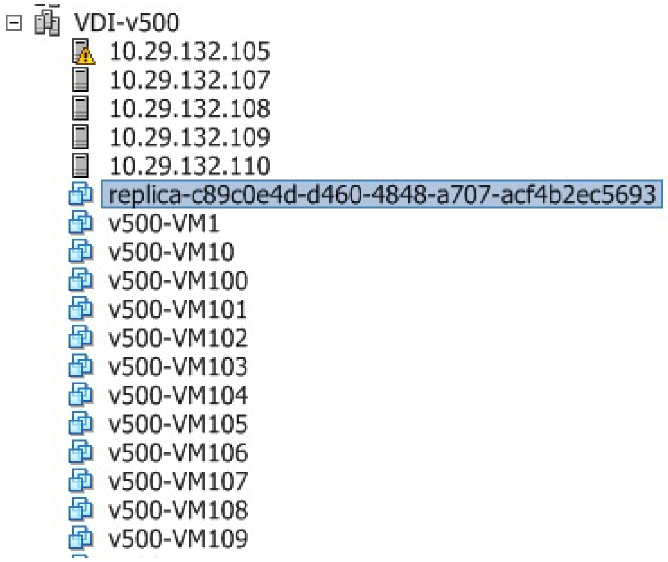

VDI-v500
10.29.132.105
10.29.132.107
10.29.132.108
10.29.132.109
10.29.132.110
replica-c89c0e4d-d460-4848-a707-acf4b2ec5693
v500-VM1
v500-VM10
v500-VM100
v500-VM101
v500-VM102
v500-VM103
v500-VM104
v500-VM105
v500-VM106
v500-VM107
v500-VM108
v500-VM109

Test Setup and Configurations

In this project, we tested a single Cisco UCS C220 M3 Rack-Mount server and five Cisco UCS C220 M3 Rack-Mount servers to illustrate linear scalability in a 600 User configuration with N+1 Server Fault Tolerance.

Cisco UCS Test Configuration for Single Server Scalability

Figure 19 Cisco UCS C220 M3 Rack-Mount Server for Single Server Scalability

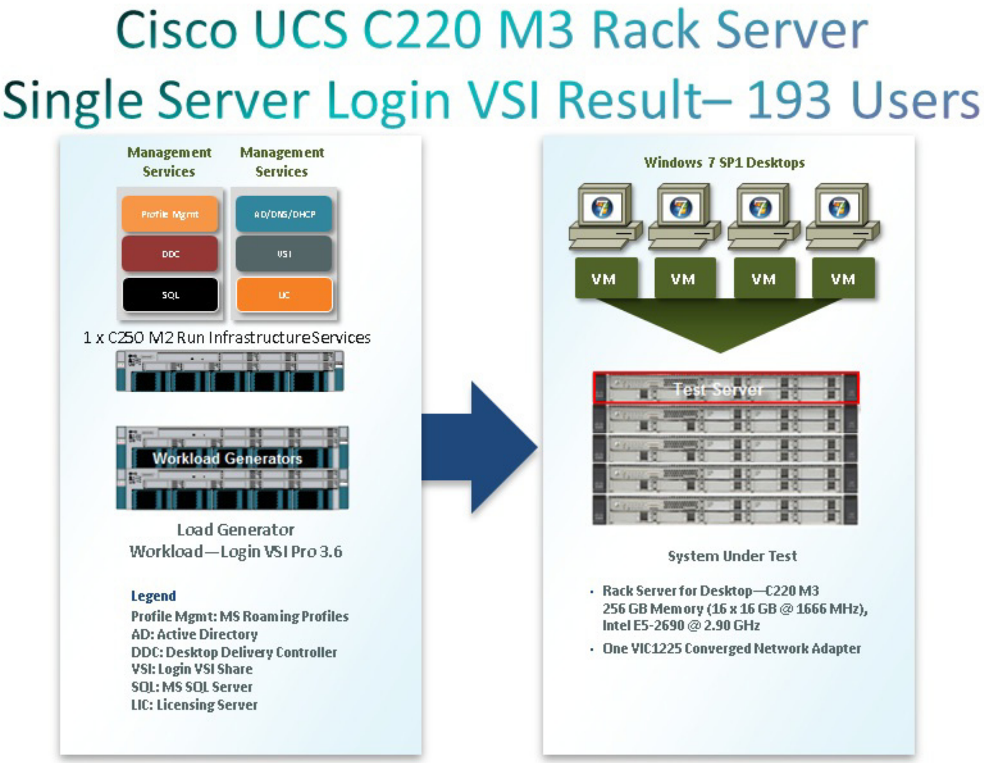

Cisco UCS C220 M3 Rack Server

Single Server Login VSI Result– 193 Users

Legend
Profile Mgmt: MS Roaming Profiles
AD: Active Directory
DDC: Desktop Delivery Controller
VSI: Login VSI Share
SQL: MS SQL Server
LIC: Licensing Server

- Rack Server for Desktop—C220 M3
256 GB Memory (16 x 16 GB @ 1666 MHz),
Intel E5-2690 @ 2.90 GHz
- One VIC1225 Converged Network Adapter

Hardware components

- 1 X Cisco UCS C220 M3 (2 x E5-2690 @ 2.90 GHz) blade server with 256GB of memory (16GB X 16 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktop hosts
- 1 X Cisco UCS C250 M2 (Xeon 5690 @ 3.47 GHz) blade servers with 96 GB of memory (4GB X 24 DIMMS @ 1333 MHz) Infrastructure Servers
- 2 X Cisco UCS C250 M2 (Xeon 5690 @ 3.47 GHz) blade servers with 96 GB of memory (4GB X 24 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Server (C250 M2)
- 1X VIC1225 Converged Network Adapter/Server (C220 M3)
- 2 X Cisco Fabric Interconnect 6248UPs(Managed FC Variant Only)
- 2 X Cisco Nexus 5548UP Access Switches (Unmanaged NFS Variant Only)
- 2 X Cisco Nexus 2232PP 10 GE Fabric Extenders (Managed FC Variant Only)
- 1 X EMC VNX5300 System storage array, two controllers, four Datamovers, 2 x dual port 8GB FC cards, 4 x dual port 10 GbE cards, 2 x 100GB Flash Drives for EMC Fast Cache, 10 x 600GB SAS drives for View 5.1 linked clone disks, 5 x 600GB SAS Drives for Infrastructure and Boot LUNs, 1 x 100GB Flash Drive for hot spare and 1 x 600GB SAS drives for hot spare

Software components

- Cisco UCS firmware 2.1(1a) (Managed FC Variant Only)

- Cisco UCS C220 M3 firmware 1.4.7b
- VMware ESXi 5.1 for VDI Hosts
- VMware View 5.1.2
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 18 GB/VM

Cisco UCS Configuration for Five Server

Figure 20 600 Desktop Test Configuration-5 x Cisco UCS C220 M3 Rack-Mount Servers

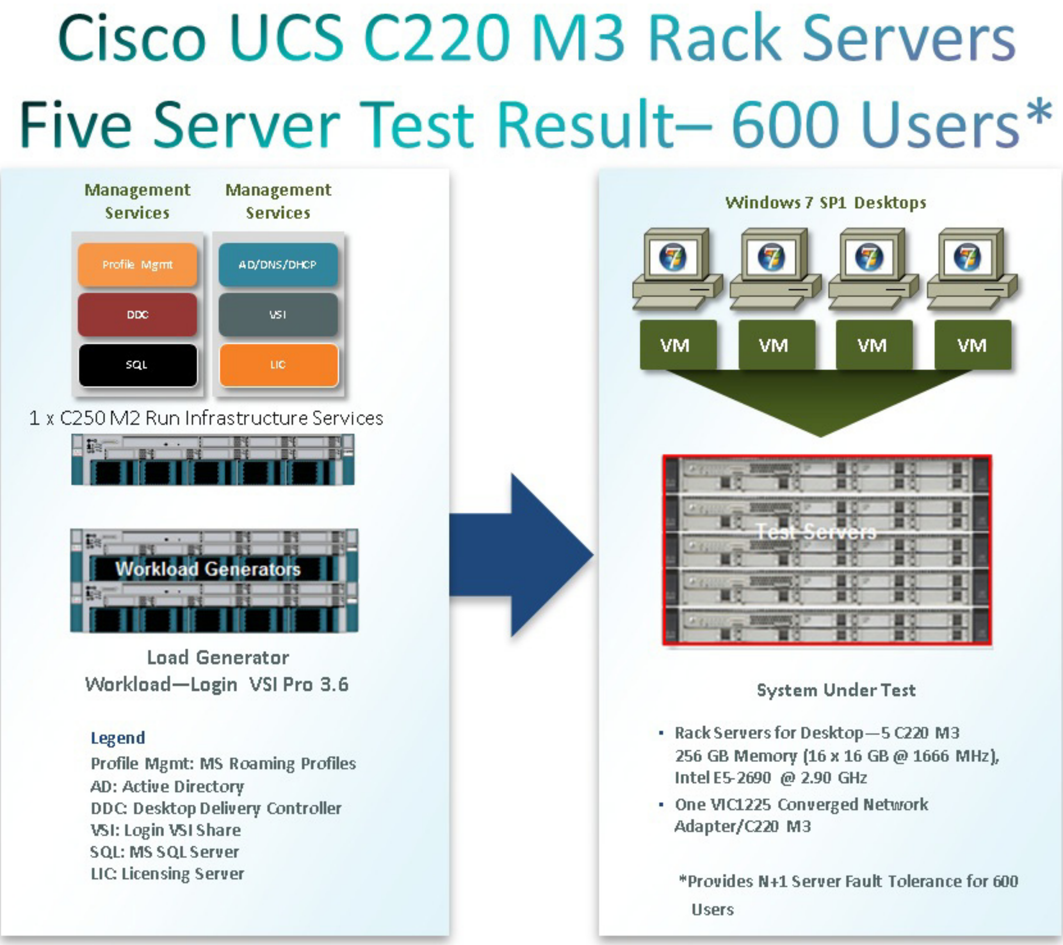
Cisco UCS C220 M3 Rack Servers
Five Server Test Result– 600 Users*
Management Services
Management Services
Profile Mgmt
AD/DNS/DHCP
DDC
VSI
SQL
LIC
1 x C250 M2 Run Infrastructure Services
Workload Generators
Load Generator
Workload—Login VSI Pro 3.6
Legend
Profile Mgmt: MS Roaming Profiles
AD: Active Directory
DDC: Desktop Delivery Controller
VSI: Login VSI Share
SQL: MS SQL Server
LIC: Licensing Server
Windows 7 SP1 Desktops
VM
VM
VM
VM
Test Servers
System Under Test
• Rack Servers for Desktop—5 C220 M3 256 GB Memory (16 x 16 GB @ 1666 MHz), Intel E5-2690 @ 2.90 GHz
• One VIC1225 Converged Network Adapter/C220 M3
*Provides N+1 Server Fault Tolerance for 600 Users

Hardware components

- 5 X Cisco UCS C220 M3 (2 x E5-2690 @ 2.90 GHz) blade server with 256GB of memory (16GB X 16 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktop hosts
- 1 X Cisco UCS C250 M2 (Xeon 5690 @ 3.47 GHz) blade servers with 96 GB of memory (4GB X 24 DIMMS @ 1333 MHz) Infrastructure Servers
- 2 X Cisco UCS C250 M2 (Xeon 5690 @ 3.47 GHz) blade servers with 96 GB of memory (4GB X 24 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Server (C250 M2)

- 1X VIC1225 Converged Network Adapter/Server (C220 M3)
- 2 X Cisco Fabric Interconnect 6248UPs (Managed FC Variant Only)
- 2 X Cisco Nexus 5548UP Access Switches (Unmanaged NFS Variant Only)
- 2 X Cisco Nexus 2232PP 10 GE Fabric Extenders (Managed FC Variant Only)
- 1 X EMC VNX5300 System storage array, two controllers, four Datamovers, 2 x dual port 8GB FC cards, 4 x dual port 10 GbE cards, 2 x 100GB Flash Drives for EMC Fast Cache, 10 x 600GB SAS drives for View 5.1 linked clone disks, 5 x 600GB SAS Drives for Infrastructure and Boot LUNs, 1 x 100GB Flash Drive for hot spare and 1 x 600GB SAS drives for hot spare

Software components

- Cisco UCS firmware 2.1(1a) (Managed FC Variant Only)
- Cisco UCS C220 M3 firmware 1.4.7b
- VMware ESXi 5.1 for VDI Hosts
- VMware View 5.1.2
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 17 GB/VM

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco Labs in San Jose, CA.

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 log off 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 View 5.1.2 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 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 View 5.1 via a View PCoIP protocol connection.

User Workload Simulation-Login VSI From Login VSI Inc.

One of the most critical factors of validating a desktop virtualization 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 VSI Inc 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 Login Virtual Session Indexer (Login VSI Inc' Login VSI 3.7) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and hence 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 via Thinapp

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 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The type rate is 160ms for each character.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Microsoft Outlook 2007/2010, browse 10 messages.
- Microsoft 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.
- Microsoft 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.
- Microsoft Excel 2007/2010, a very large randomized sheet is opened.
- Microsoft 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.

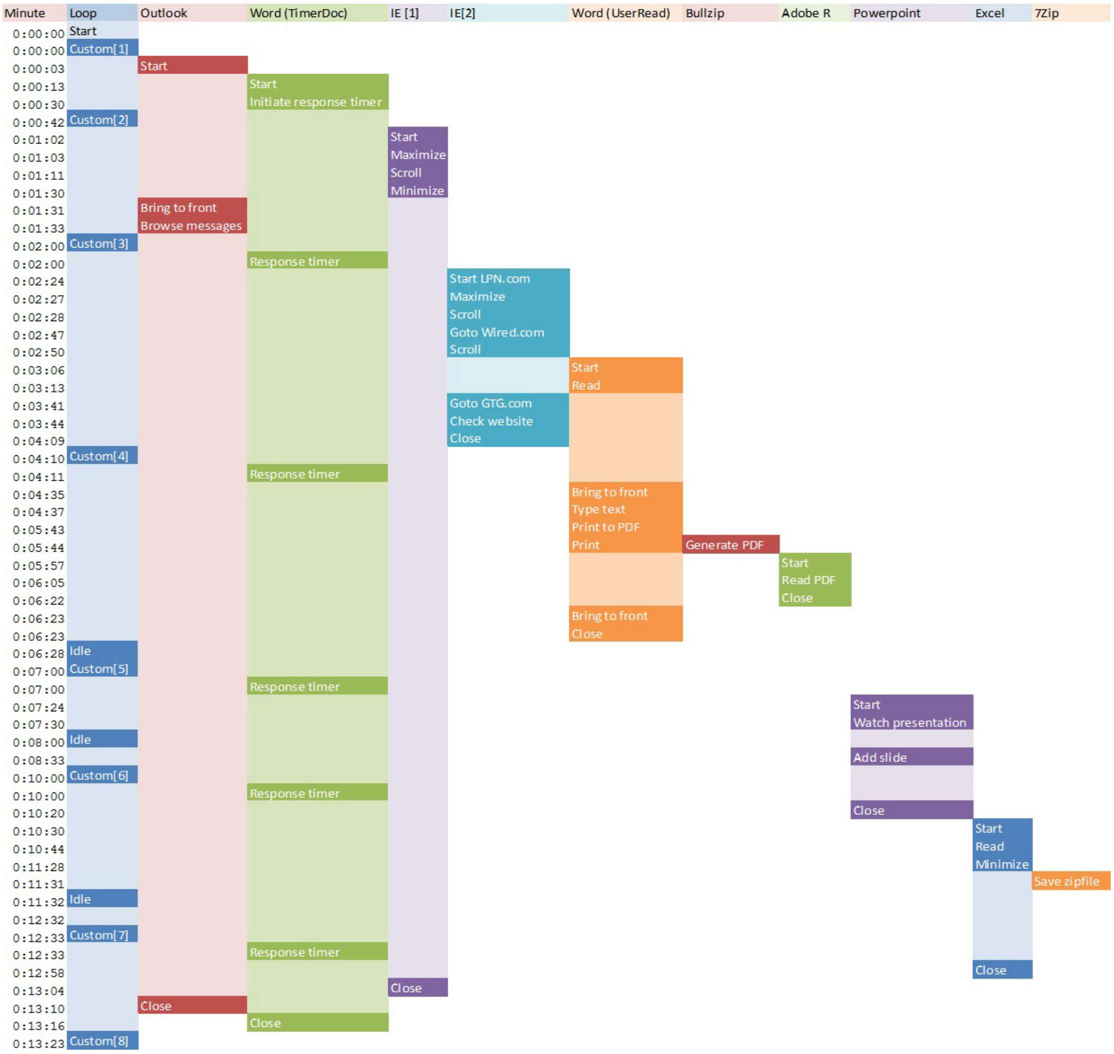

Minute	Loop	Outlook	Word (TimerDoc)	IE [1]	IE[2]	Word (UserRead)	Bullzip	Adobe R	Powerpoint	Excel	7Zip
0:00:00	Start										
0:00:00	Custom[1]										
0:00:03		Start									
0:00:13			Start								
0:00:30			Initiate response timer								
0:00:42	Custom[2]										
0:01:02				Start							
0:01:03				Maximize							
0:01:11				Scroll							
0:01:30				Minimize							
0:01:31		Bring to front									
0:01:33		Browse messages									
0:02:00	Custom[3]										
0:02:00			Response timer								
0:02:24					Start LPN.com						
0:02:27					Maximize						
0:02:28					Scroll						
0:02:47					Goto Wired.com						
0:02:50					Scroll						
0:03:06						Start					
0:03:13						Read					
0:03:41					Goto GTG.com						
0:03:44					Check website						
0:04:09					Close						
0:04:10	Custom[4]										
0:04:11			Response timer								
0:04:35						Bring to front					
0:04:37						Type text					
0:05:43						Print to PDF					
0:05:44						Print	Generate PDF				
0:05:57								Start			
0:06:05								Read PDF			
0:06:22								Close			
0:06:23						Bring to front					
0:06:23						Close					
0:06:28	Idle										
0:07:00	Custom[5]										
0:07:00			Response timer								
0:07:24									Start		
0:07:30									Watch presentation		
0:08:00	Idle										
0:08:33									Add slide		
0:10:00	Custom[6]										
0:10:00			Response timer								
0:10:20									Close		
0:10:30										Start	
0:10:44										Read	
0:11:28										Minimize	
0:11:31											Save zipfile
0:11:32	Idle										
0:12:32											
0:12:33	Custom[7]										
0:12:33			Response timer								
0:12:58										Close	
0:13:04				Close							
0:13:10		Close									
0:13:16			Close								
0:13:23	Custom[8]										

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 insure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

All virtual machines were shut down utilizing the VMware View Administrator and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a "waiting for test to start" state.

All VMware ESXi 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 195 single server users or 600 full scale test users to become active within 2 minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Parallel Launching method is used for all single server and scale testing. This assures that our tests represent real-world scenarios.

Note The Login VSI Sequential Launching method allows the CPU, storage and network components to rest between each logins. This does not produce results that are consistent with the real-world scenarios that our Customers run in.

For each of the three consecutive runs on single server (195 User) and 4 and 5 server (500 and 600 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
   - DDCs used in test run
   - Profile Server(s) used in test run
   - SQL Server(s) used in test run
   - 3 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 (Unmanaged NFS Variant Only)
4. Time 0:05 Take 195, 500 or 600 desktops out of maintenance mode on View Administrator
5. Time 0:06 First machines boot
6. Time 0:26 195, 500 or 600 desktops booted on 1 or 5 servers
7. Time 0:28 195, 500 or 600 desktops available on 1 or 5 servers
8. Time 1:28 Start Login VSI 3.7 Test with 195, 500 or 600 desktops utilizing 7, 17 or 20 Launchers
9. Time 1:58 195, 500 or 600 sessions launched
10. Time 2:00 195, 500 or 600 sessions active
11. Time 2:15 Login VSI Test Ends
12. Time 2:30 195, 500 or 600 sessions logged off
13. Time 2:35 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.

The following 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:

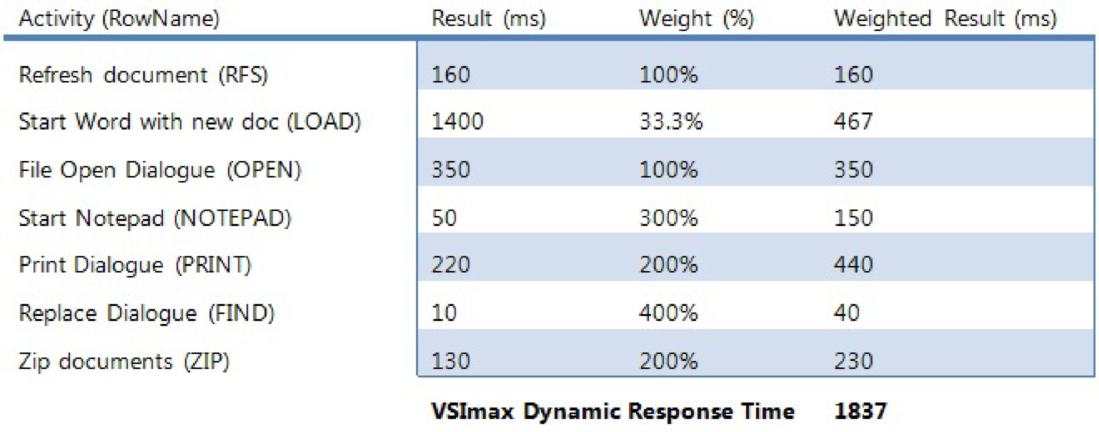

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	33.3%	467
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	300%	150
Print Dialogue (PRINT)	220	200%	440
Replace Dialogue (FIND)	10	400%	40
Zip documents (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 98. 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 21 ***Sample Login VSI Analyzer Graphic Output***

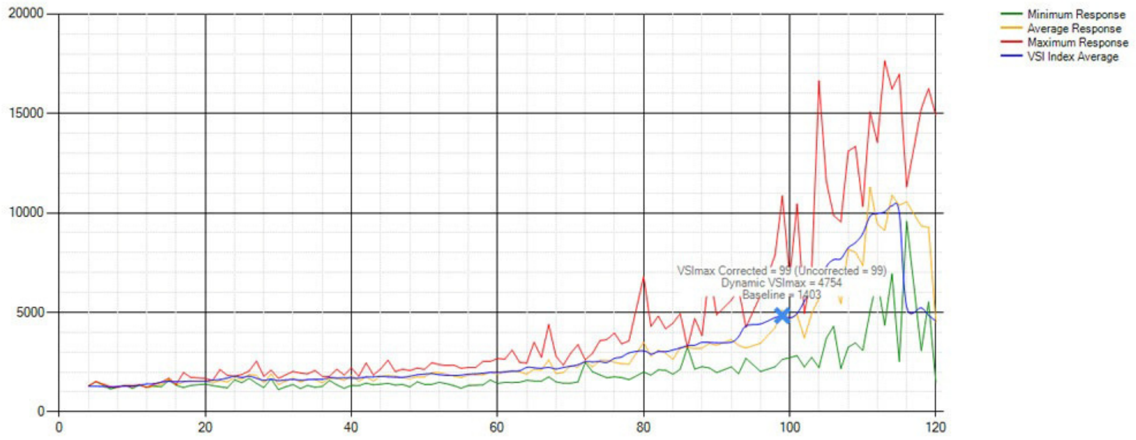

- Red line: Maximum Response (worst response time of an individual measurement within a single session)
- Orange line: Average Response Time within for each level of active sessions
- Blue line: the VSImax average.
- Green line: Minimum Response (best response time of an individual measurement within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

In our tests, the total number of users in the test run had to login, become active and run at least one test loop and log out automatically without reaching the VSI Max to be considered a success.

Note We discovered a technical issue with the VSIMax dynamic calculation in our testing on Cisco C220 M2 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 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 3.7 medium workloads, including flash.
- Number of Launched Sessions must equal Active Sessions within two minutes of the last session launched in a test.
- The View Administrator console will be monitored throughout the steady state to insure 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 recommendations 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.

VDI Test Results

The purpose of this testing is to provide the data needed to validate VMware View 5.1 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 C220 M3 rack mount servers using a EMC VNX5300 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 5.1 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 VSImax Score. To reach the Login VSImax, we ran 205 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 below.

Figure 22 Log In VSImax Reached: 193 Users

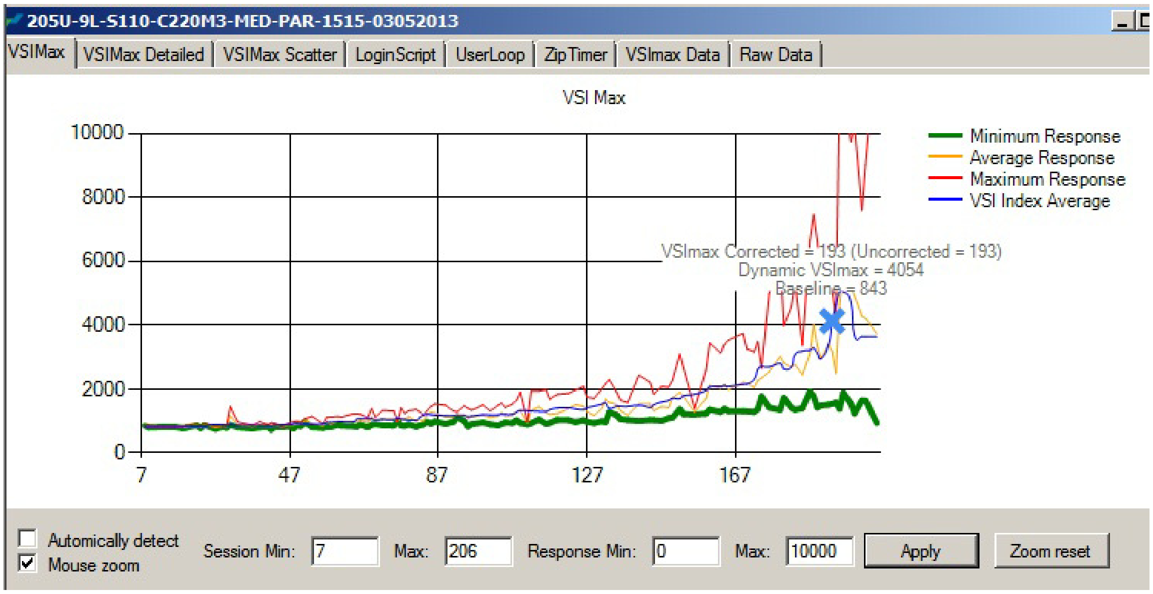

□ Automically detect
☑ Mouse zoom

Session Min: 7 Max: 206 Response Min: 0 Max: 10000

Cisco UCS Configuration for Single-Server Scalability Test Results

This section details the results from the View 5.1 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 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 20% lower user density than prescribed by the Login VSImax to achieve a successful pass of the test with server hardware performance in a realistic range.

The data below demonstrates that a single UCS C220 M3 is able to host 155 users sessions with average end user response times of under 2 seconds, well below the Login VSImax Dynamic threshold for unacceptable performance.

This provides evidence that 4 UCS C220 M3s can handle the full 600 user compliment in the event that one of the 5 UCS C220 M3s prescribed for the 600 Users workload fails.

Similarly, for the alternate 4 UCS C220 M3s 500 User alternate solution, this data supports 450 Users can run on 3 UCS C220 M3s in the event that one of the 4 UCS C220 M3s fails.

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.

Figure 23 155 View 5.1.2 Desktop Sessions on VMware ESXi 5.1 below 2500 ms

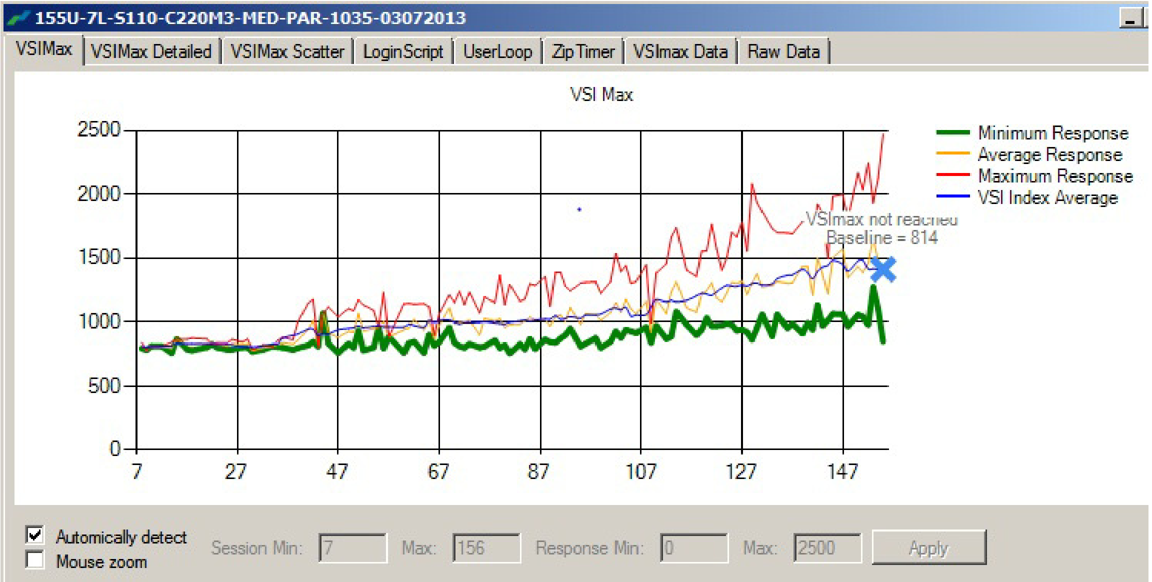

155U-7L-S110-C220M3-MED-PAR-1035-03072013

VSImax not reached
Baseline = 814

☑ Automically detect
□ Mouse zoom

Session Min: 7 Max: 156 Response Min: 0 Max: 2500

The following graphs detail CPU, Memory, Disk and Network performance on the Single Cisco UCS C220 M3 Blades.

Figure 24 155 User Single UCS C220 M3 CPU Utilization Boot Phase

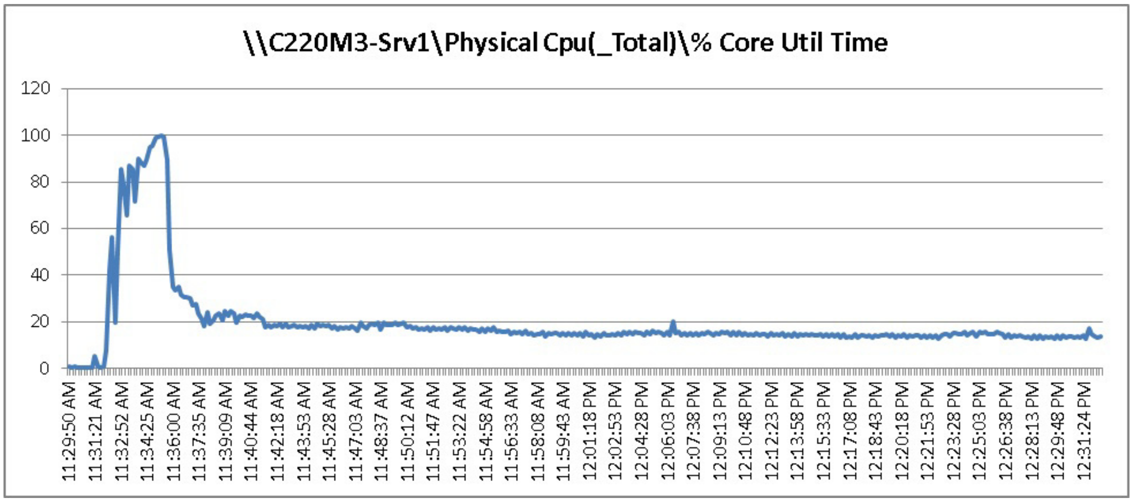

Figure 25 155 User Single UCS C220 M3 Available Memory Boot Phase

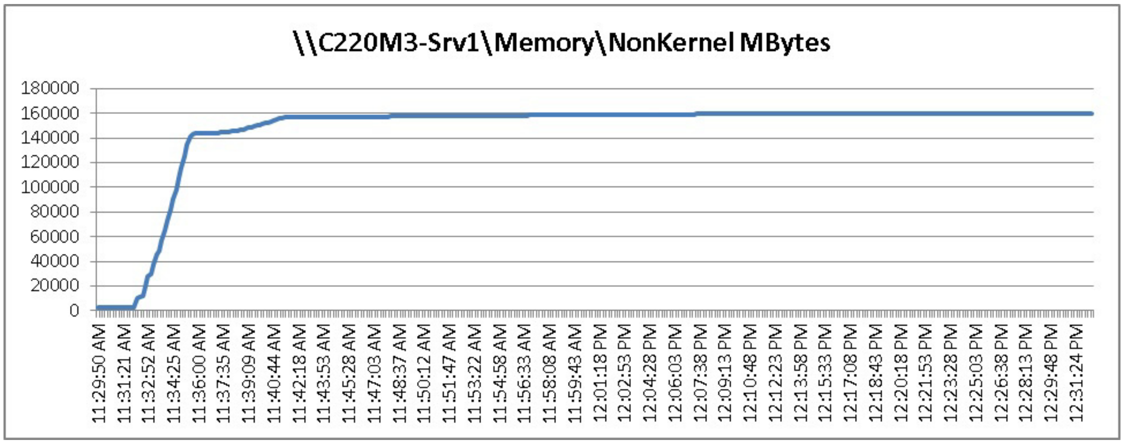

Figure 26 155 User Single UCS C220 M3 Cisco VIC1225 VIC Network Rates Boot Phase

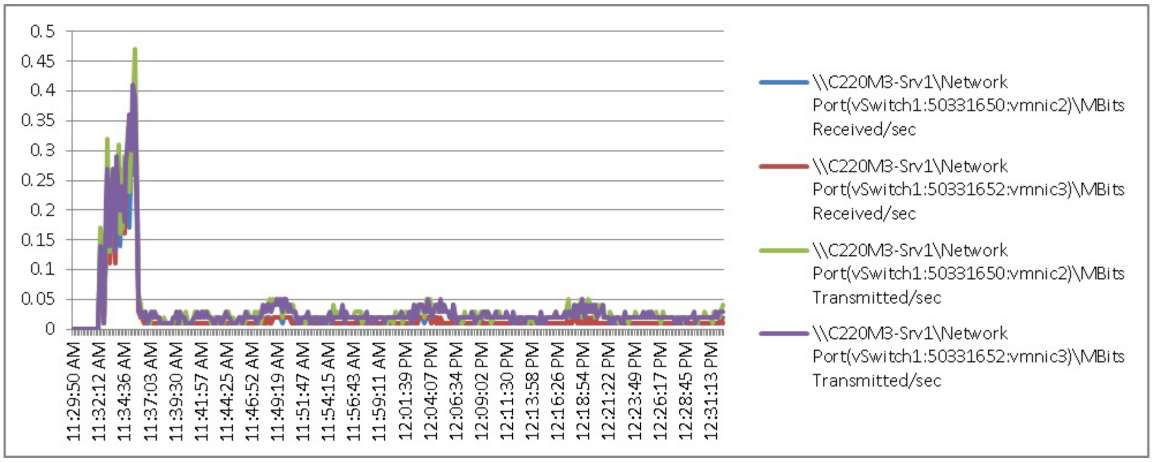

Figure 27 155 User Single UCS C220 M3 CPU Utilization Test Phase

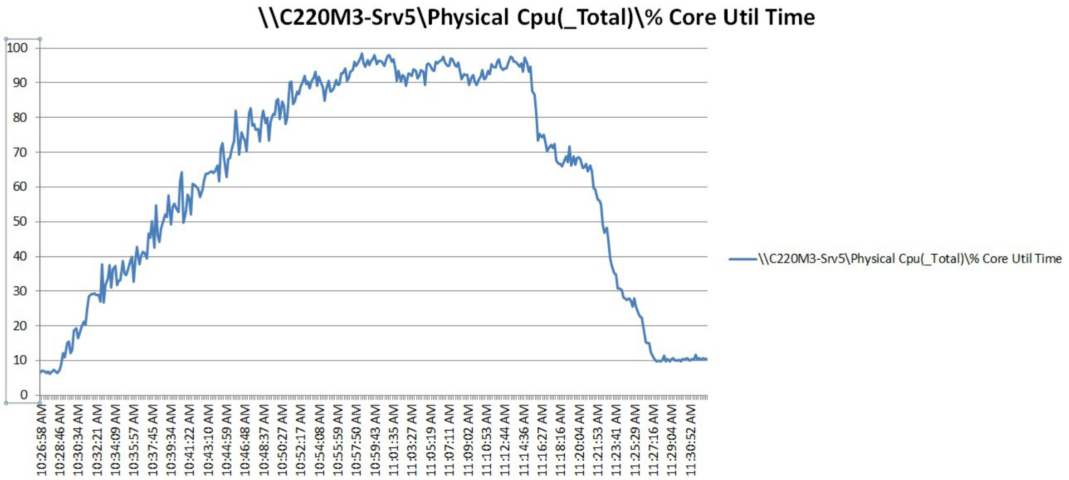

Figure 28 155 User Single UCS C220 M3 Available Memory Test Phase

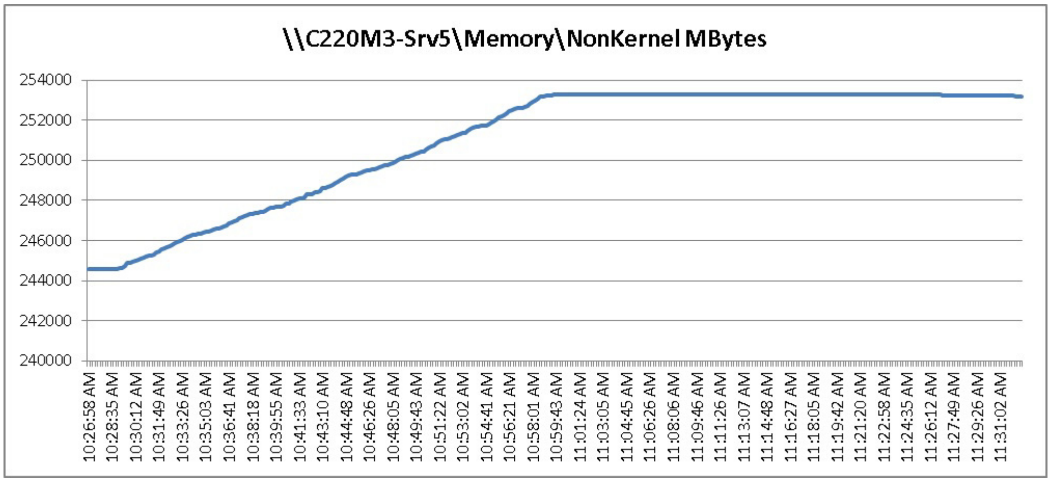

Figure 29 155 User Single UCS C220 M3 Cisco VIC1225 VIC Network Rates Test Phase

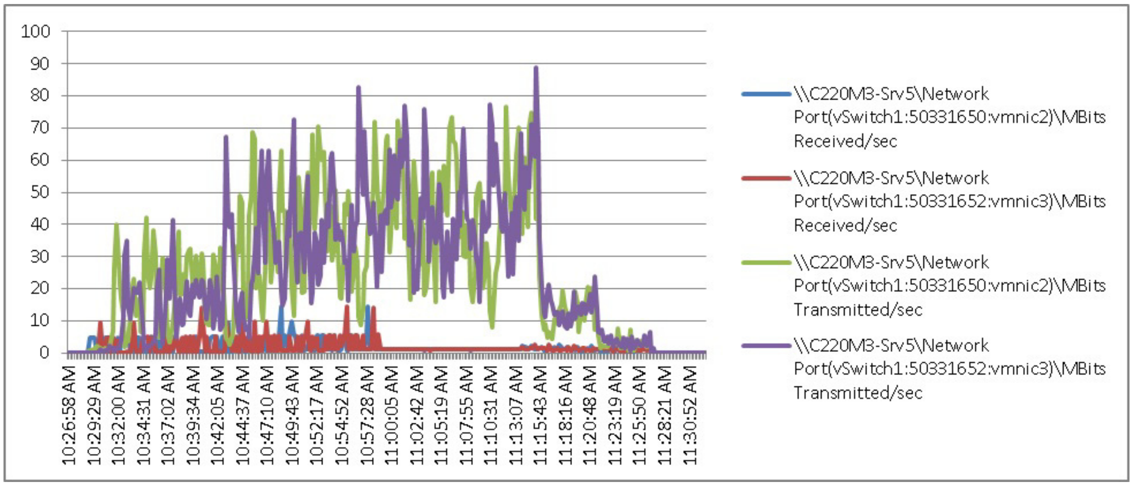

The following graphs detail performance of the EMC VNX5300 during the single blade, 155 user test phase.

Figure 30 155 User View 5.1.2 VNX5300 SP A and SP B Queue Length

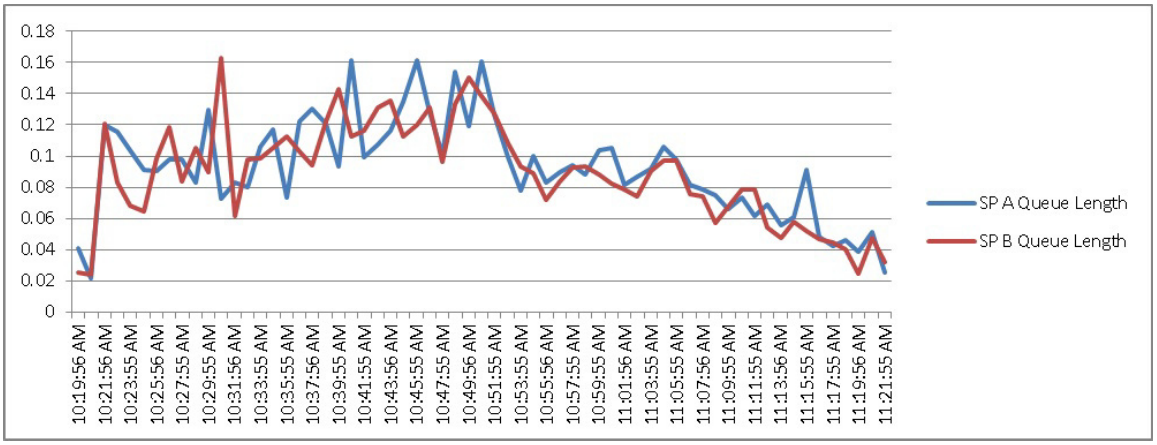

Figure 31 155 User View 5.1.2 VNX5300 SP A and SP B Utilization

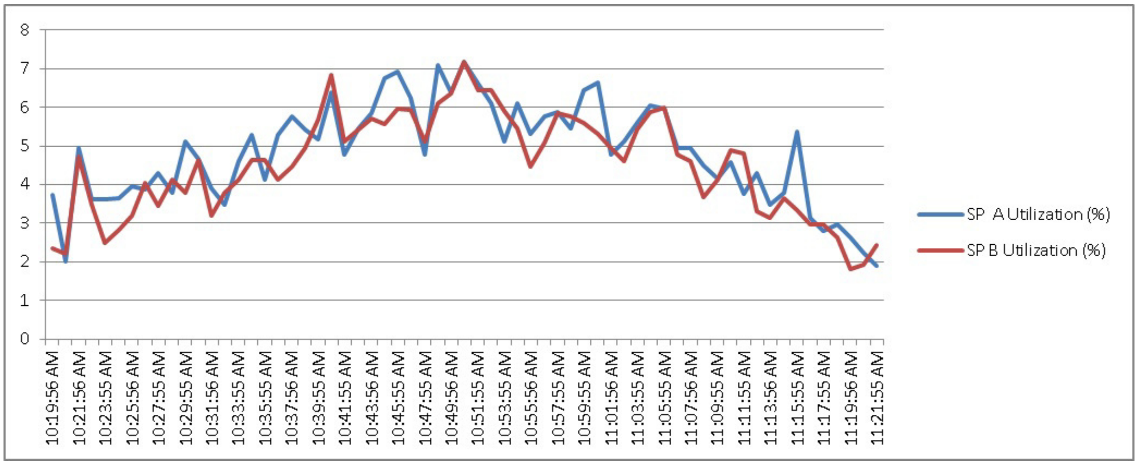

Figure 32 155 User View 5.1.2 VNX5300 SP A and SP B Total Throughput

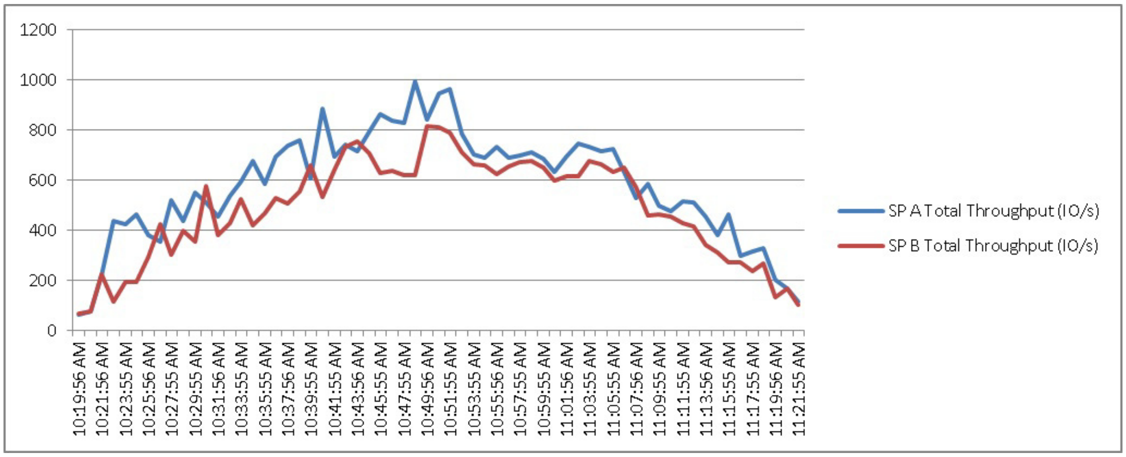

Cisco UCS Test Configuration for 600 Desktop Scalability Test Results

This section details the results from the View 5.1 Hosted VDI five rack server, 600 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 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 a user density of 120 users per server across the system. One ESX Cluster containing five Cisco UCS C220 M3 servers ran the entire workload. This configuration provides N+1 fault tolerance on a cluster 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 120 user sessions. The single server graphs for blades running 120 user sessions are essentially the same. We have provided the remaining four Cisco UCS C220 M3 servers' performance charts in Appendix D to illustrate this point.

Given adequate storage capability, the CPU utilization determined the maximum recommended VM density per blade for the 600 user environment.

For this scale test, we are including the EMC VNX5300 performance metrics as well.

Figure 33 600 User View 5.1.2 Desktop Sessions on 5 C220 M3s running VMware ESXi 5.1 below 2500 ms

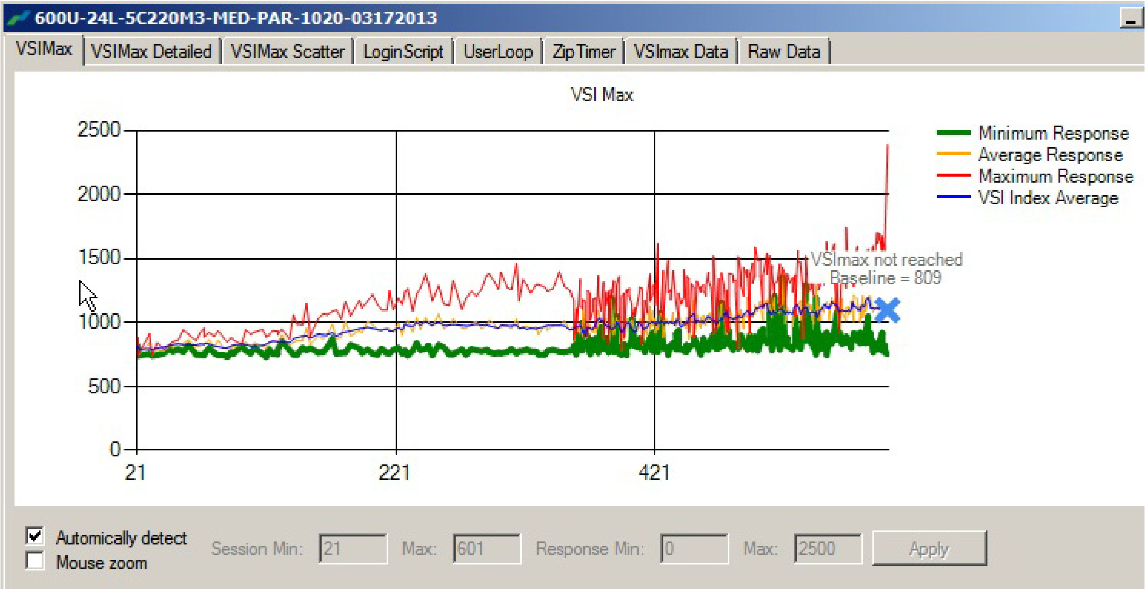

☑ Automicaly detect
□ Mouse zoom

Session Min: 21 Max: 601 Response Min: 0 Max: 2500

The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS C220 M3server during the five server, 600 User test. (Representative results for all five servers in one vCenter clusters can be found in Appendix C.)

Figure 34 600 Desktop Sessions on 5 C220 M3 CPU Utilization Boot Phase

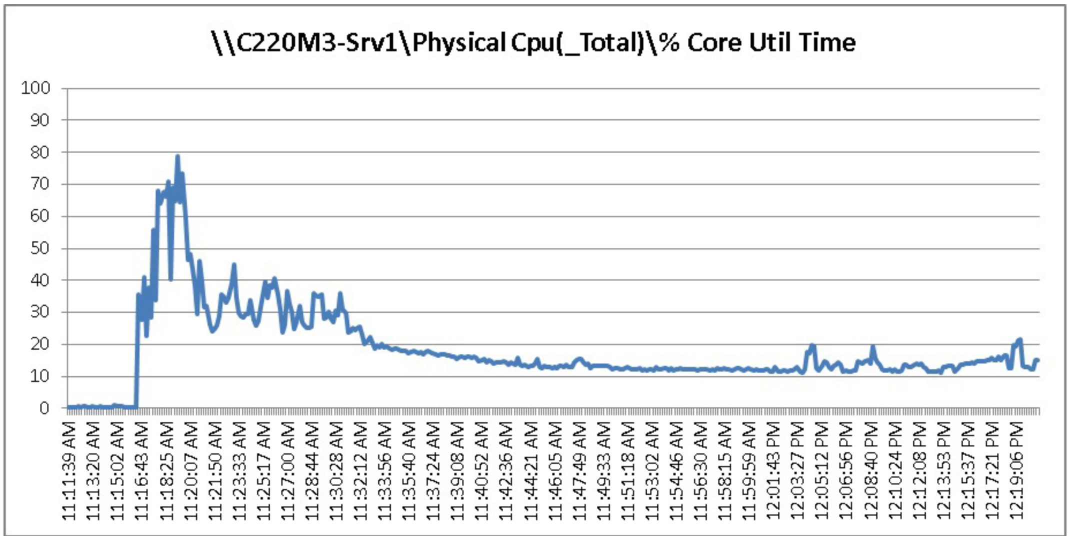

Figure 35 600 Desktop Sessions on 5 C220 M3 Memory Utilization Boot Phase

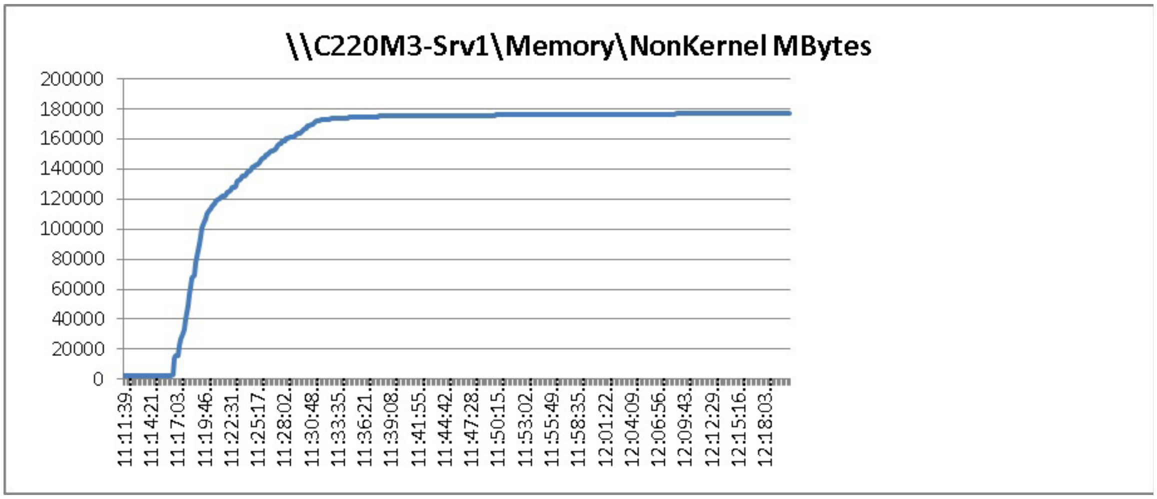

Figure 36 600 Desktop Sessions on 5 C220 M3 Cisco VIC1225 Mbps Receive/Transmit Boot Phase

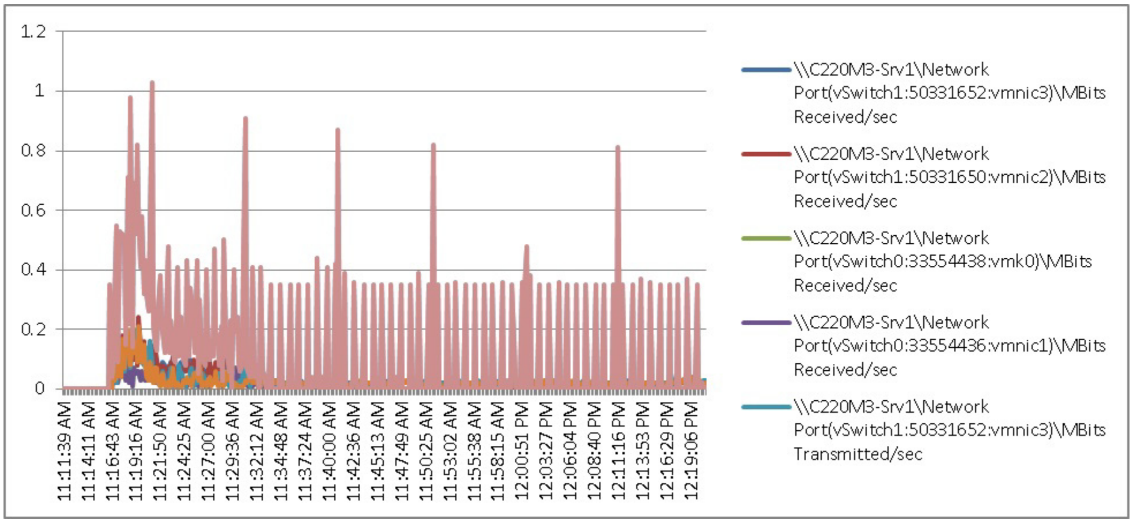

Figure 37 600 Desktop Sessions on 5 C220 M3 CPU Utilization Test Phase

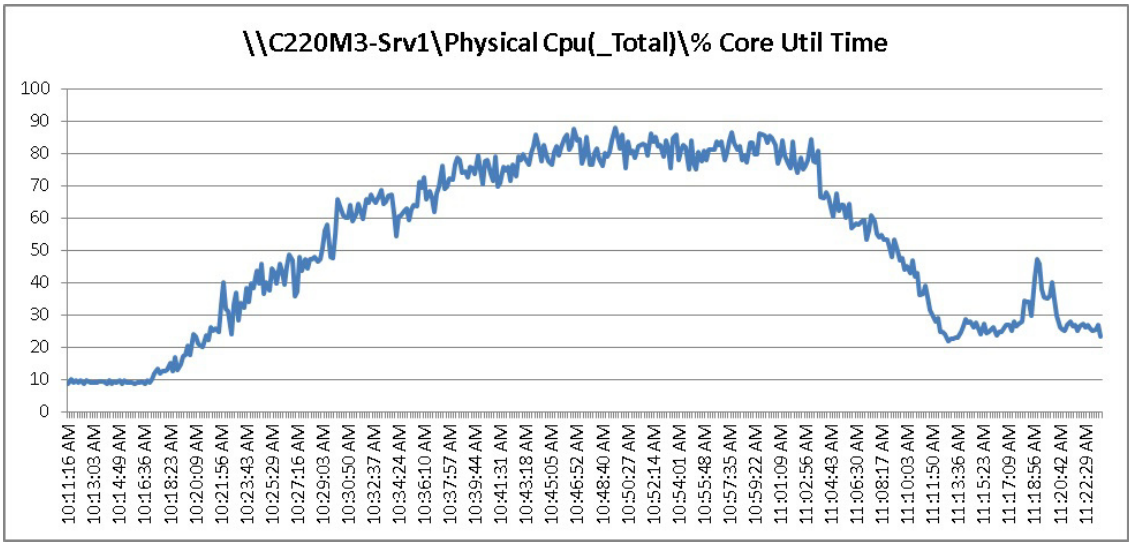

Figure 38 600 Desktop Sessions on 5 C220 M3 Memory Utilization Test Phase

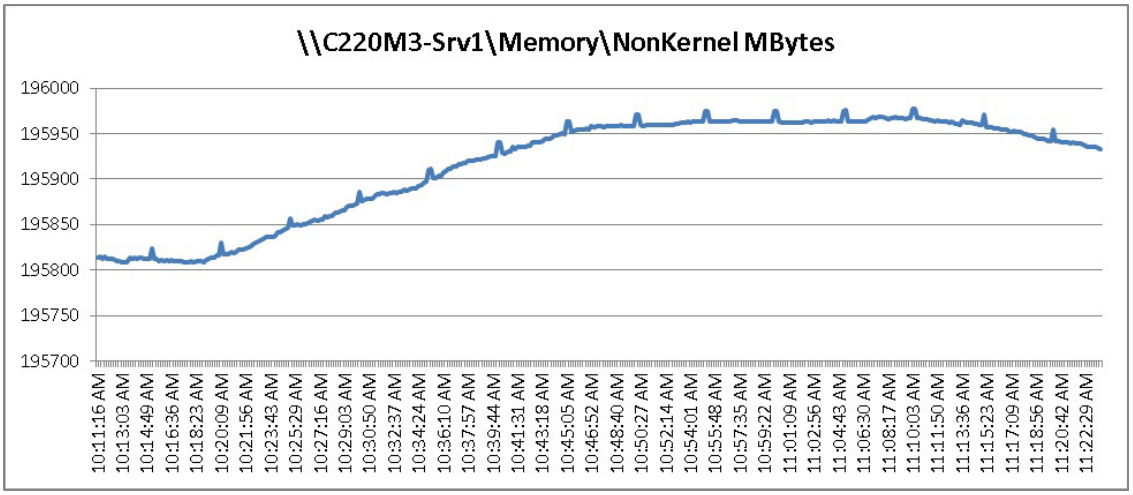

Figure 39 600 Desktop Sessions on 5 C220 M3 Cisco VIC1225 Mbps Receive/Transmit Test Phase

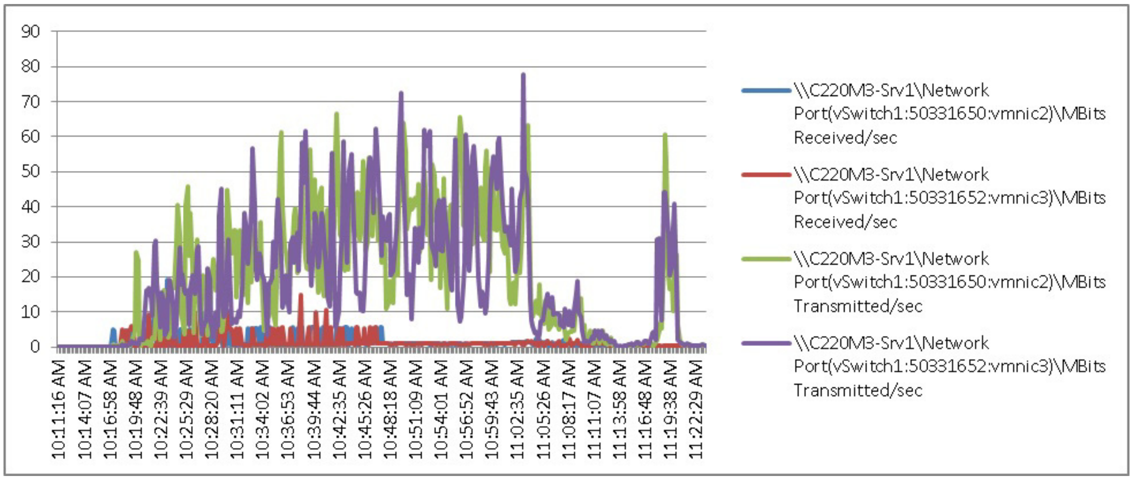

The following charts detail infrastructure server performance during the five server, 600 User test:

Figure 40 600 User View 5.1.2 Connection Server CPU Utilization Test Phase

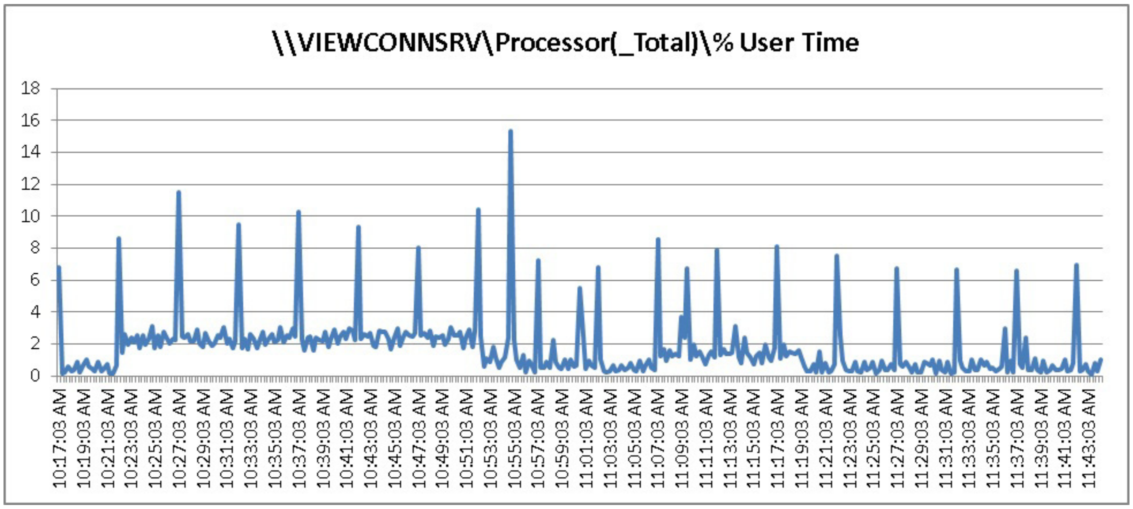

Figure 41 600 User View 5.1.2 Connection Server Available Memory Test Phase

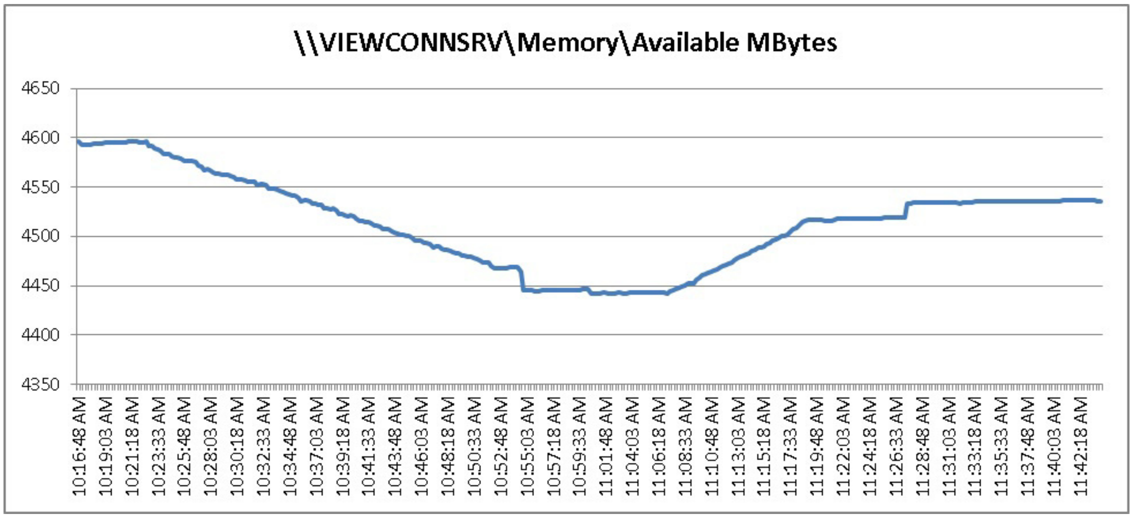

Figure 42 600 User View 5.1.2 Connection Server Bytes Received/Second Test Phase

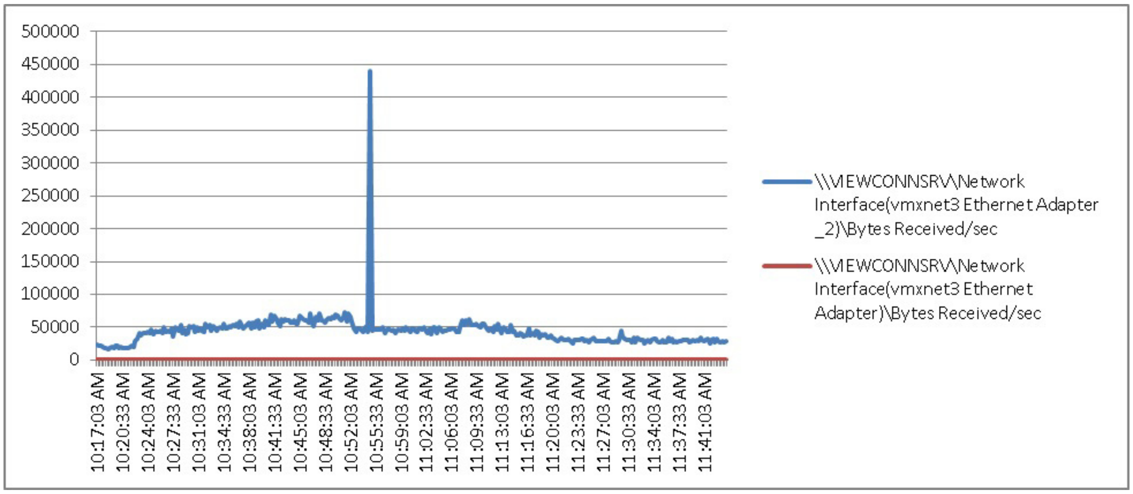

Figure 43 600 User View 5.1.2 Connection Server Bytes Sent/Second Test Phase

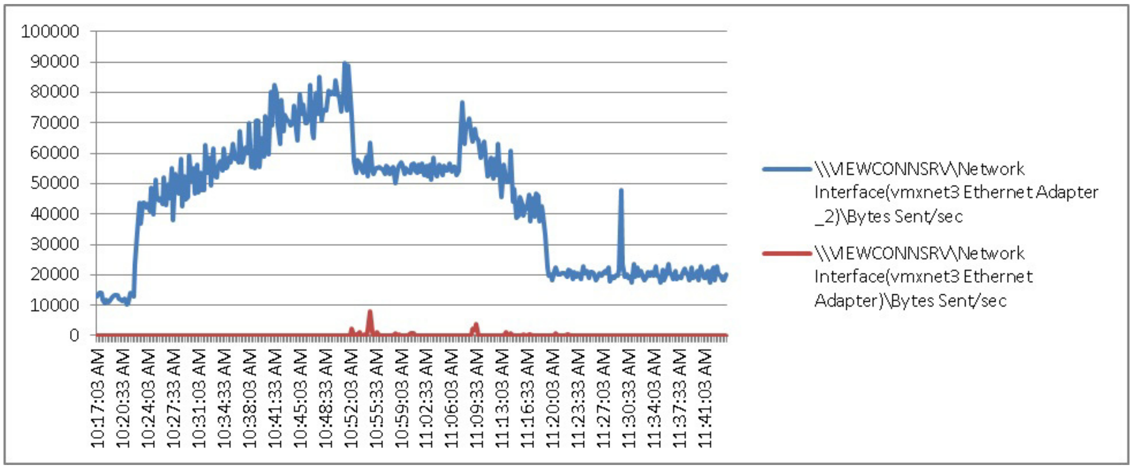

Figure 44 600 User View 5.1.2 VNX5300 SP A and SP B Queue Length

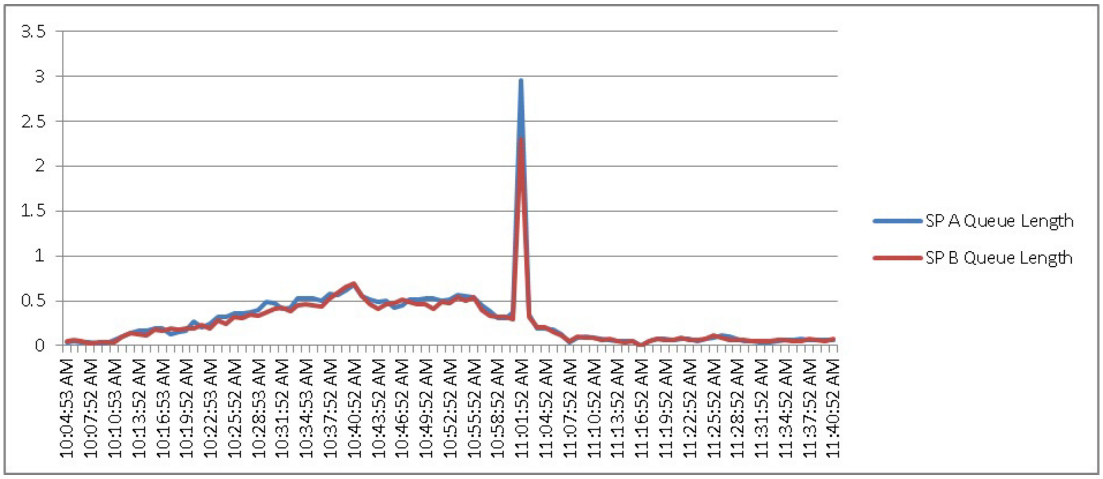

Figure 45 600 User View 5.1.2 VNX5300 SP A and SP B Utilization

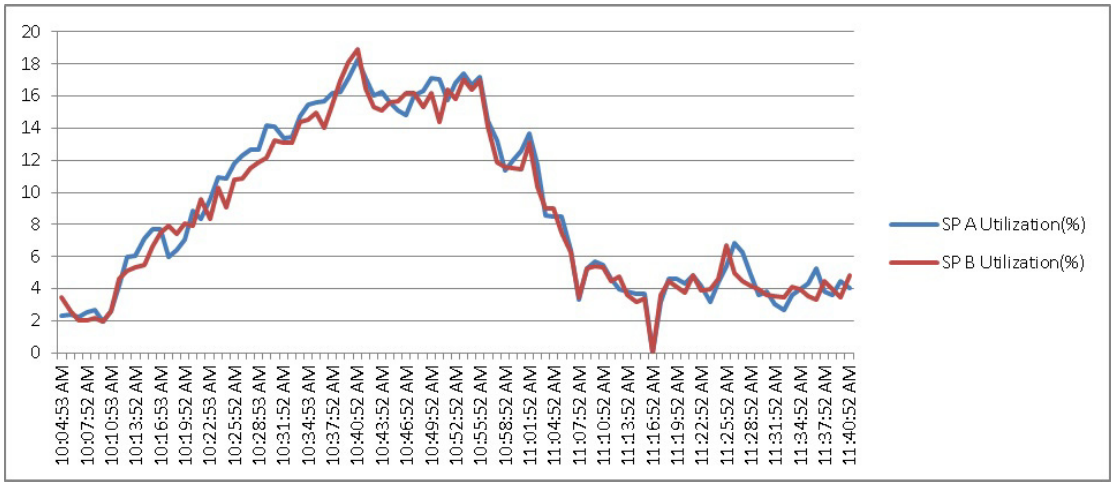

Figure 46 600 User View 5.1.2 VNX5300 SP A and SP B Total Throughput

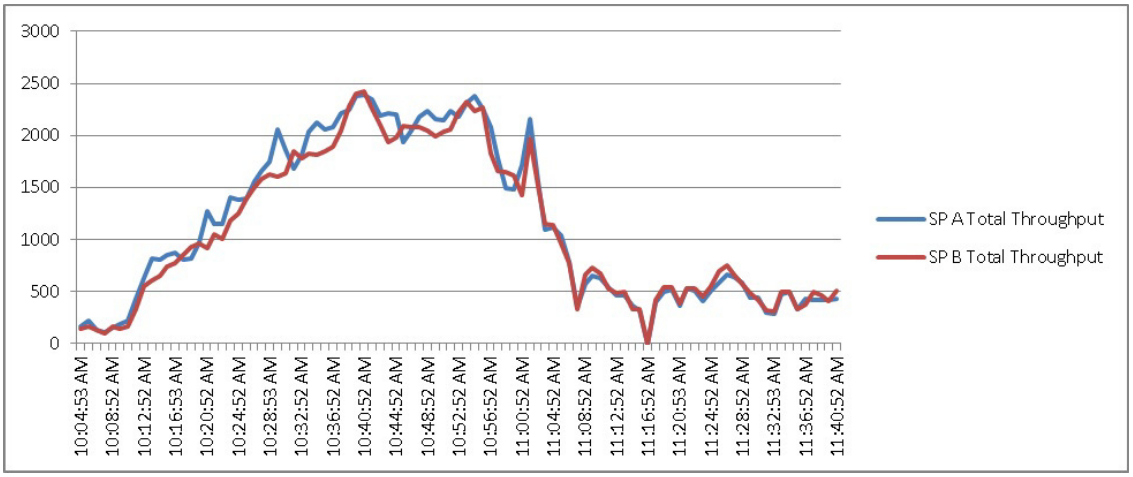

Cisco UCS Test Configuration for 500 Desktop Scalability Test Results

This section details the results from the View 5.1 Hosted VDI four rack servers 500 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 Determining VSIMax for a discussion of this issue.

We ran the multi-server test at an average user density of 125 users per server across the system. One ESX Cluster containing four Cisco UCS C220 M3 servers ran the entire workload. This configuration provides N+1 fault tolerance for 450 users 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 server running 125 user sessions. The single server graphs for blades running 125 user sessions are essentially the same. We have provided the remaining three Cisco UCS C220 M3 servers' performance charts in Appendix D to illustrate this point.

Given adequate storage capability, the CPU utilization determined the maximum recommended VM density per blade for the 500 user environment.

For this scale test, we are including the EMC VNX5300 performance metrics as well.

Figure 47 500 View 5.1 Desktop Sessions on 4 C220 M3s running VMware ESXi 5. 1 below 4000 ms

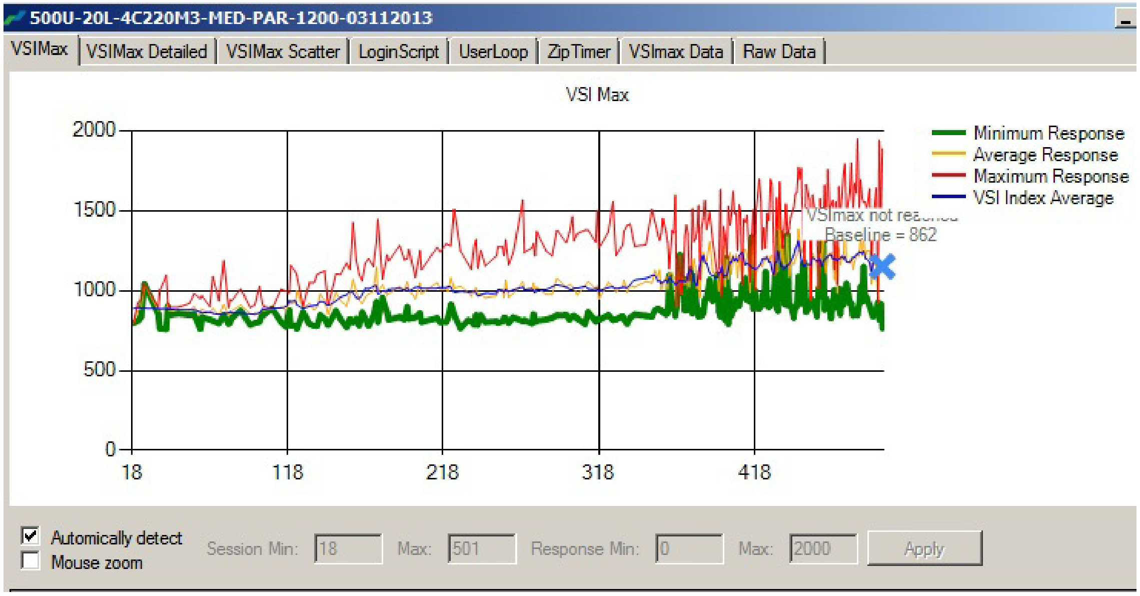

☑ Automically detect
□ Mouse zoom

Session Min: 18 Max: 501 Response Min: 0 Max: 2000

Figure 48 500 Desktop Sessions on 4 C220 M3 CPU Utilization Boot Phase

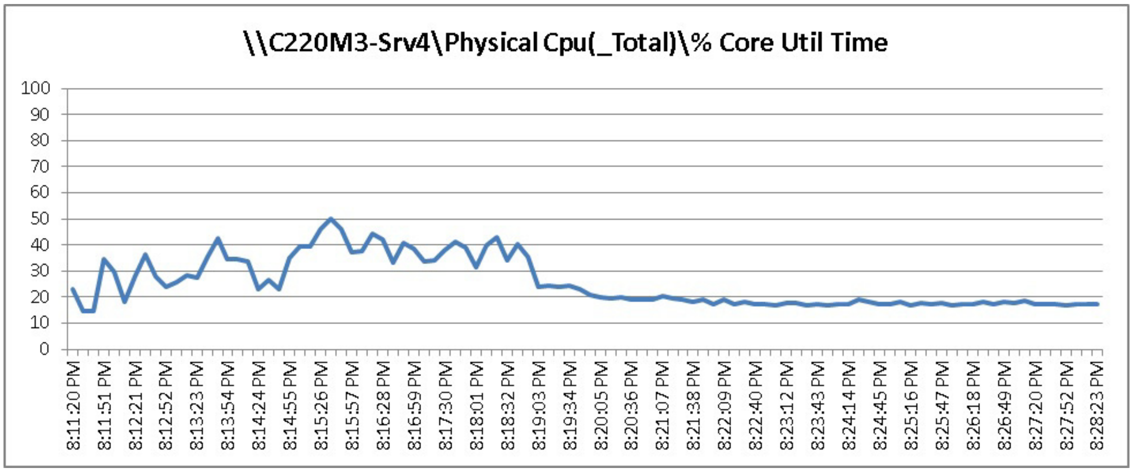

Figure 49 500 Desktop Sessions on 4 C220 M3 Memory Utilization Boot Phase

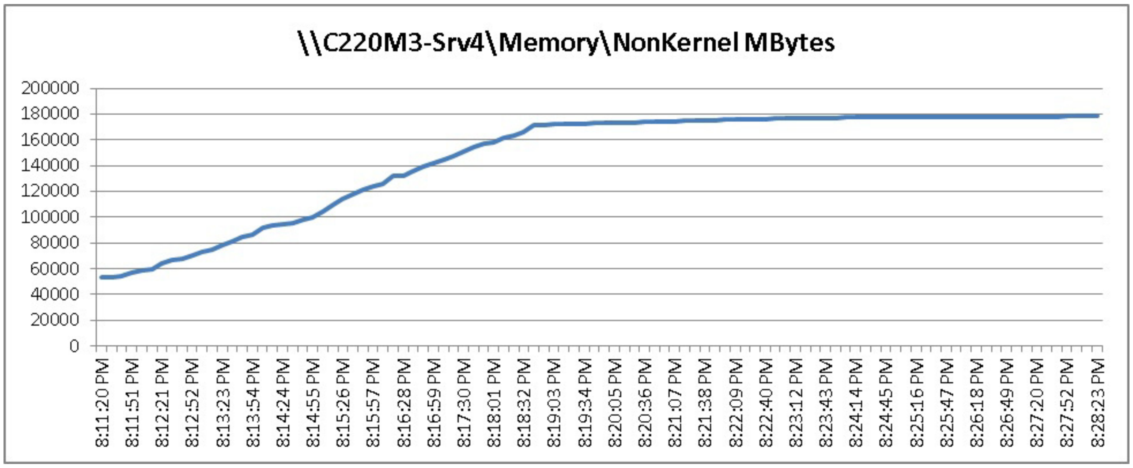

Figure 50 500 Desktop Sessions on 4 C220 M3 Cisco VIC1225 Mbps Receive/Transmit Boot Phase

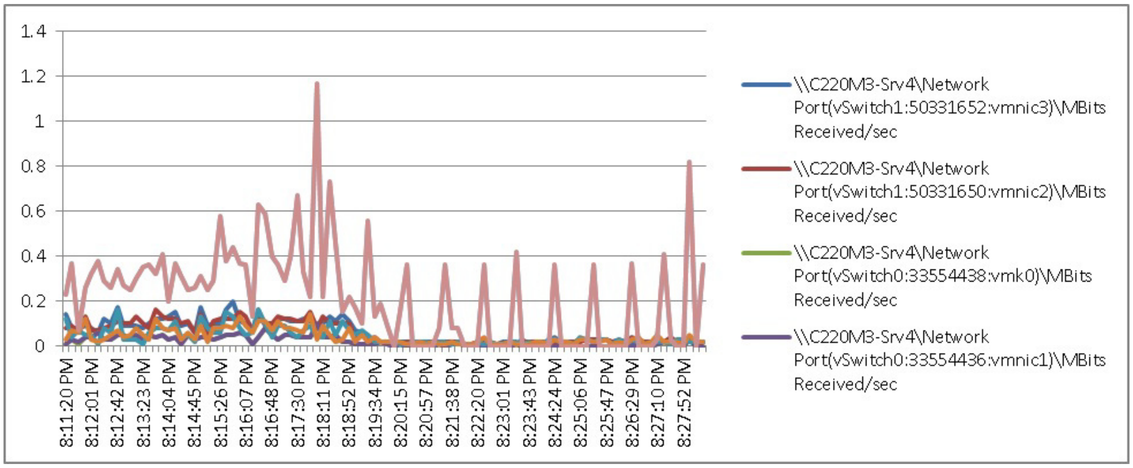

Figure 51 500 Desktop Sessions on 4 C220 M3 CPU Utilization Test Phase

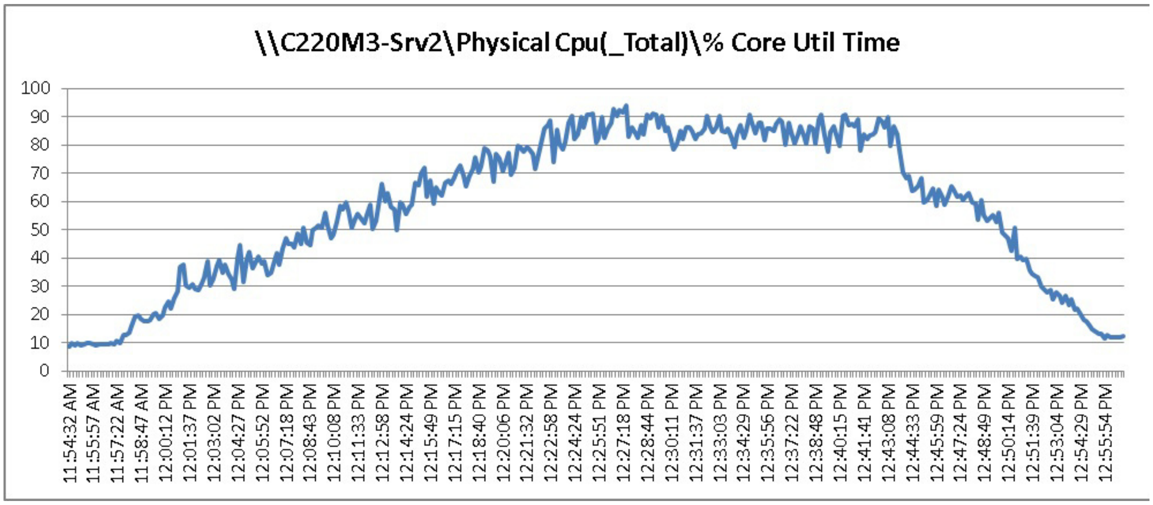

Figure 52 500 Desktop Sessions on 4 C220 M3 Memory Utilization Test Phase

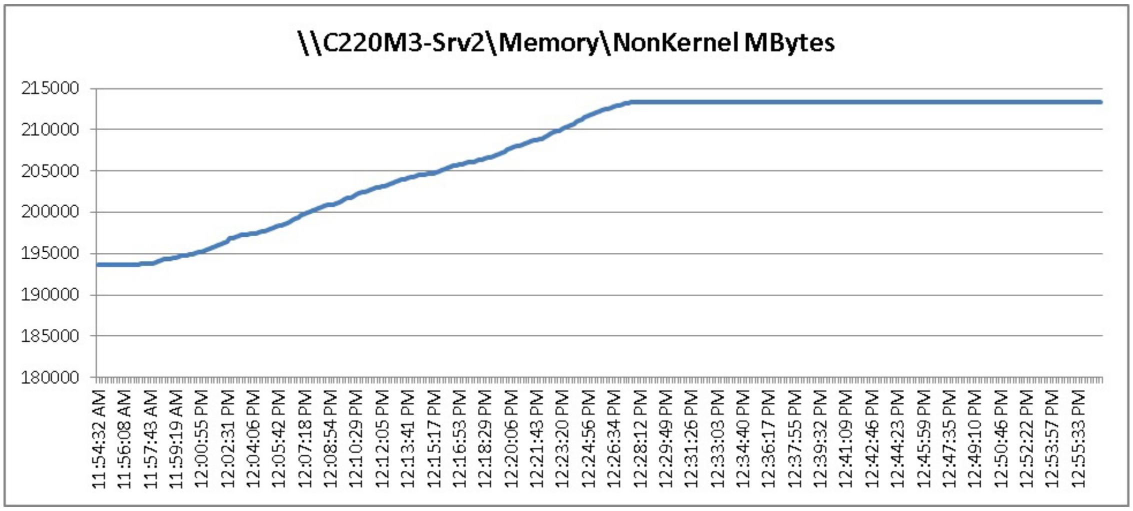

Figure 53 500 Desktop Sessions on 4 C220 M3 Cisco VIC1225 Mbps Receive/Transmit Test Phase

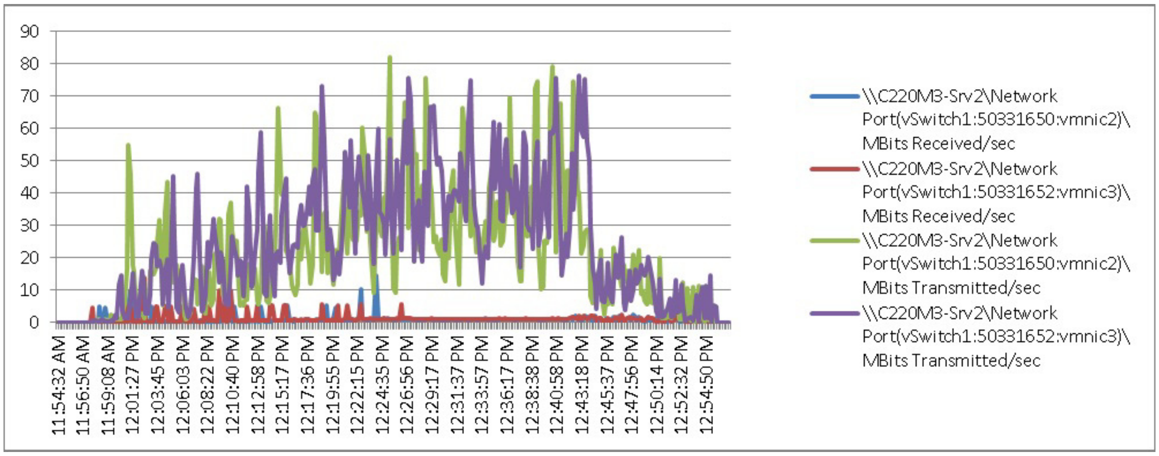

The following charts detail infrastructure server performance during the five server, 500 User test:

Figure 54 500 User View 5.1.2 Connection Server CPU Utilization Test Phase

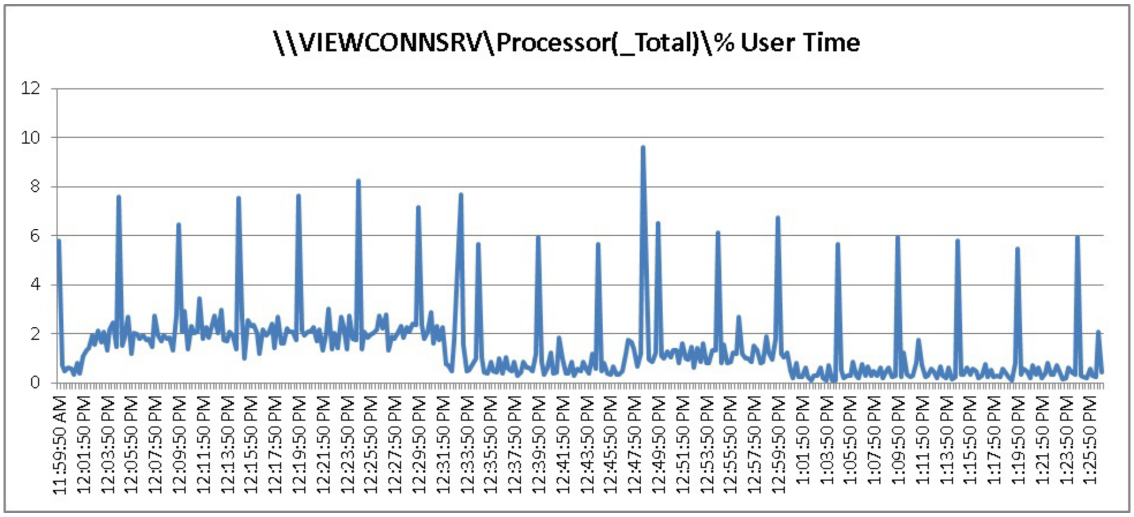

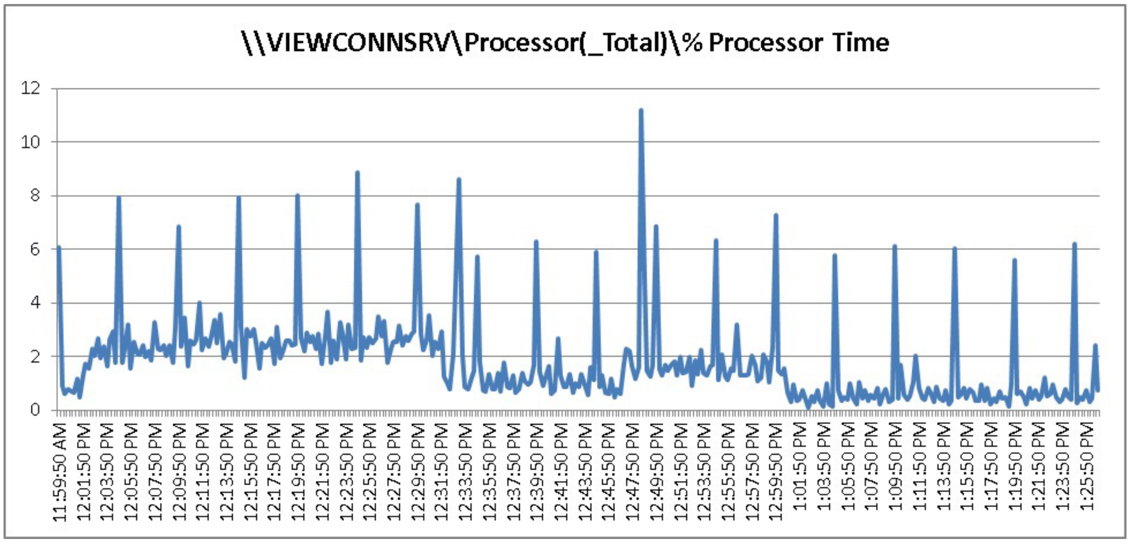

Figure 55 500 User View 5.1.2 Connection Server Available Memory Test Phase

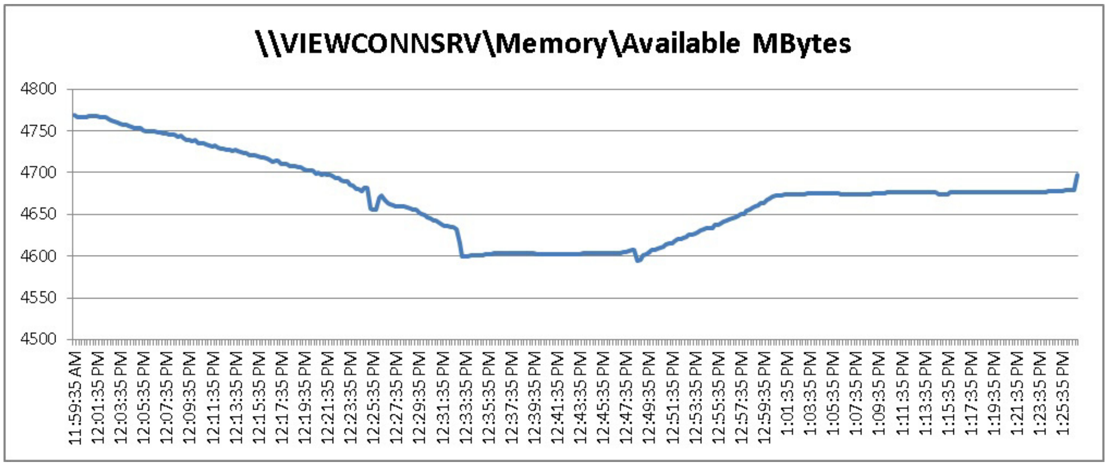

Figure 56 500 User View 5.1.2 Connection Server Bytes Received/Second Test Phase

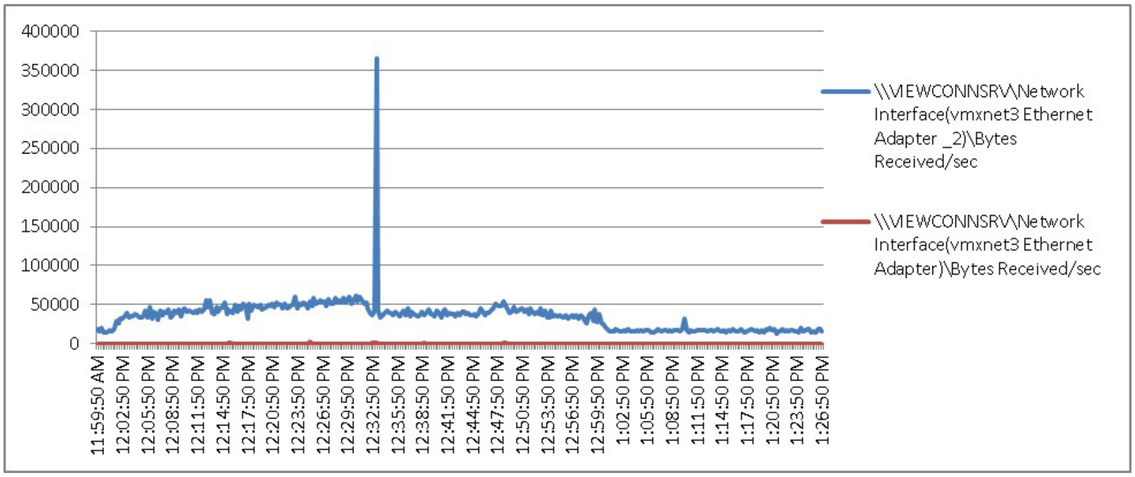

Figure 57 500 User View 5.1.2 Connection Server Bytes Sent/Second Test Phase

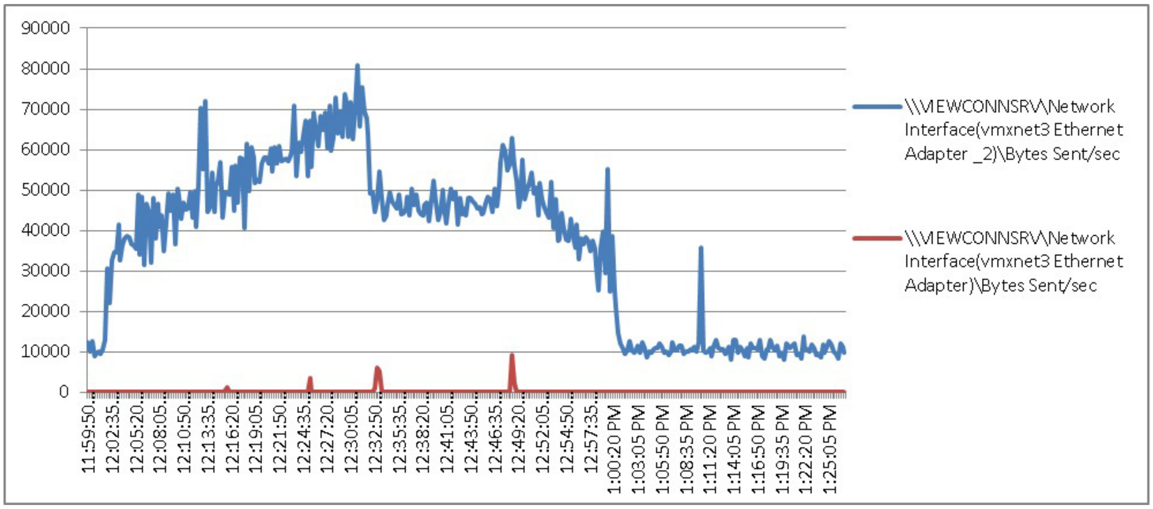

Figure 58 500 User View 5.1.2 VNX5300 SP A and SP B Queue Length

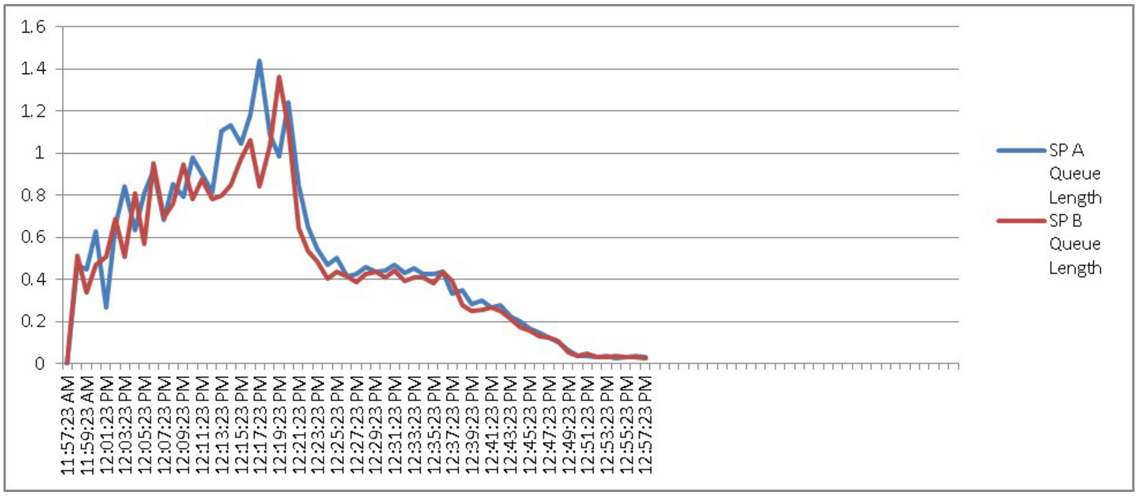

Figure 59 500 User View 5.1.2 VNX5300 SP A and SP B Utilization

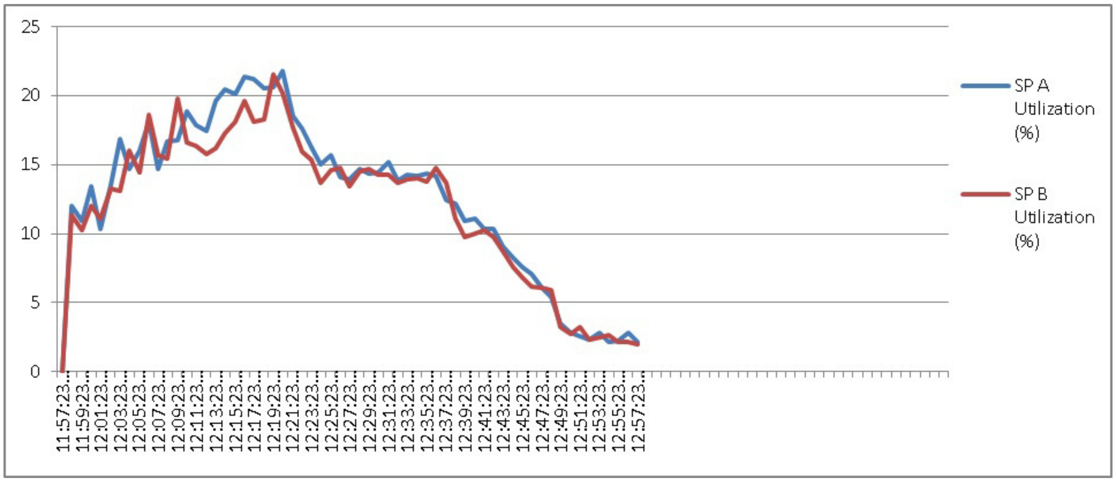

Figure 60 500 User View 5.1.2 VNX5300 SP A and SP B Total Throughput

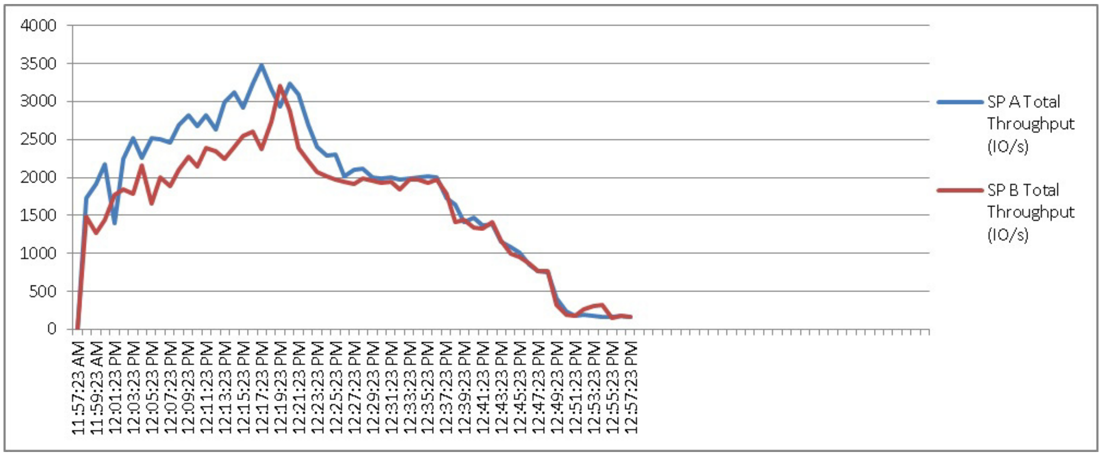

Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond a 500-600 User, four to five Cisco UCS C220 host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 500-600 user system.

Cisco UCS System Configuration

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 624UP8 and 6296UP models. Our single UCS domain can grow to 160 blades or rack servers or any combination of the two totaling 160.
- With Cisco UCS 2.1 management software, released late in November 2012, each Cisco UCS 2.1 Management domain is extensibly manageable by UCS Central, our new manager of managers, vastly increasing the reach of the Cisco UCS system.
- As scale grows, the value of the combined UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100 percent of the time.
- To accommodate the Cisco Nexus 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. Based on the number of uplinks from each chassis, we can calculate number of desktops can be hosted in a single 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 B230 M2 and one Cisco UCS B200 M3 blades servers in each chassis.

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

VMware View 5.1.2 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 5.1 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 5.1 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 5.1 configuration and guidelines, see Chapter 11 References.

EMC VNX Storage Guidelines for VMware View 5.1 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 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 C220 M3 Rack Server Resources

http://www.cisco.com/en/US/partner/products/ps12369/index.html

Cisco UCS 6200 Series Fabric Interconnects

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

Cisco Nexus 2232PP 10GE Fabric Extender

http://www.cisco.com/en/US/partner/products/ps10784/index.html

Cisco Nexus 5500 Series Switches Resources

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

Download Software for UCS C220 M3 Rack Server

http://software.cisco.com/download/type.html?mdfid=284296253&i=rs

Download Cisco UCS Manager and Blade Software Version 2.0(4d)

http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=1.4(4k)&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 5.1 Reference Documents

View 5 Documents

Performance and Best Practices

http://www.vmware.com/files/pdf/view/VMware-View-Performance-Study-Best-Practices-Technical-White-Paper.pdf

View 5.1 Architecture and Planning

http://pubs.vmware.com/view-51/topic/com.vmware.ICbase/PDF/view-51-architecture-planning.pdf

View Storage Accelerator in VMware View 5.1

http://www.vmware.com/files/pdf/techpaper/vmware-view-storage-accelerator-host-caching-content-based-read-cache.pdf

View 5 with PCoIP Network Optimization Guide

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

VMFS File Locking Impact in View 5.1

http://www.vmware.com/files/pdf/techpaper/vmware-view-vmfs-file-locking.pdf

Virtual Desktop

Windows 7 Optimization Guide

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

EMC Reference Documents

Sizing EMC VNX Series for VDI Workload - An Architectural Guideline:
http://www.emc.com/collateral/software/white-papers/h11096-vdi-sizing-wp.pdf

Deploying Microsoft Windows 7 Virtual Desktops with VMware View - Applied Best Practices Guide
http://www.emc.com/collateral/software/white-papers/h8043-windows-virtual-desktop-view-wp.pdf

Deploying Microsoft Windows 8 Virtual - Applied Best Practices Guide
http://powerlink.emc.com/km/live1/en_US/Offering_Technical/White_Paper/H11155-Windows_8_for_VDI_VM_ABGP.pdf

EMC Infrastructure for VMware View 5.1: VNX (NFS) vSphere 5.0, View Storage Accelerator, Persona Management, and Composer - Reference Architecture

http://emea.emc.com/collateral/software/technical-documentation/h10994-infrastructure-vmwareview-nfs-ra.pdf

EMC Infrastructure for VMware View 5.1: VNX (NFS) vSphere 5.0, View Storage Accelerator, Persona Management, and Composer - Proven Solution Guide

http://www.emc.com/collateral/software/technical-documentation/h10993-infrastructure-vmwareview-nfs-psg.pdf

EMC Infrastructure for VMware View 5.1: VNX (FC) vSphere 5.0, View Storage Accelerator, Persona Management, and Composer - Reference Architecture

http://www.emc.com/collateral/software/technical-documentation/h10996-infrastructure-vmwareview-fc-ra.pdf

EMC Infrastructure for VMware View 5.1: VNX (FC) vSphere 5.0, View Storage Accelerator, Persona Management, and Composer - Proven Solution Guide

http://www.emc.com/collateral/software/technical-documentation/h10993-infrastructure-vmwareview-nfs-psg.pdf

EMC Infrastructure for VMware View 5.0, EMC VNX Series (NFS), VMware vSphere 5.0, VMware View 5.0, and VMware View Composer 2.7 - Reference Architecture

http://www.emc.com/collateral/software/technical-documentation/h8305-infra-view5.0-vsphere5.0-ra.pdf

EMC Infrastructure for VMware View 5.0, EMC VNX Series (NFS), VMware vSphere 5.0, VMware View 5.0, and VMware View Composer 2.7 - Proven Solution Guide http://www.emc.com/collateral/hardware/technical-documentation/h8306-infra-view5.0-vsphere5.0-psg.pdf

EMC Infrastructure for User Virtualization with VMware View 5.0, VNX (NFS), and AppSense Environment Manager 8.1

http://www.emc.com/collateral/hardware/technical-documentation/h10654-view5-0-vnx(nfs)-vsphere5-0-appSense-ra.pdf

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://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514

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_vSphere5.0.pdf

Appendix A—Nexus 5548 Configuration (NFS Variant Only)

```
!Nexus 5548A NFS Variant VSPEX C500
!Command: show running-config
!Time: Sun Dec  6 23:15:29 2009

version 5.1(3)N1(1a)
hostname SJ2-B21-N5548-A

feature telnet
no feature http-server
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
feature fex

username admin password 5 $1$ZGA7bQiM$rD7QcH45.4ZOLIf2/J.ur1  role
network-adminno password strength-check

banner motd #Nexus 5000 Switch
```

```
#

ip domain-lookup
ip domain-name cisco.com
ip name-server 171.70.168.183 171.68.226.120
system jumbomtu 9000
logging event link-status default
class-map type qos class-fcoe
class-map type qos match-any class-platinum
  match cos 5
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing class-platinum
  match qos-group 2
class-map type queuing class-all-flood
  match qos-group 2
class-map type queuing class-ip-multicast
  match qos-group 2
policy-map type qos jumbo
  class class-default
    set qos-group 0
policy-map type qos system_qos_policy
  class class-platinum
    set qos-group 2
  class class-default
    set qos-group 0
policy-map type queuing system_q_in_policy
  class type queuing class-platinum
    bandwidth percent 50
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-default
    bandwidth percent 30
policy-map type queuing system_q_out_policy
  class type queuing class-platinum
    bandwidth percent 50
  class type queuing class-fcoe
    bandwidth percent 20
  class type queuing class-default
    bandwidth percent 30
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos class-platinum
  match qos-group 2
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos system_nq_policy
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos system_nq_policy
  class type network-qos class-platinum
    pause no-drop
    mtu 9000
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-default
```

```
    mtu 9000
    multicast-optimize
system qos
  service-policy type qos input system_qos_policy
  service-policy type queuing input system_q_in_policy
  service-policy type queuing output system_q_out_policy
  service-policy type network-qos system_nq_policy
fex 130
  pinning max-links 1
  description "FEX0130"
snmp-server user admin network-admin auth md5 0x0ae428b6495ff67f478fd90e941c15d7
priv
0x0ae428b6495ff67f478fd90e941c15d7 localizedkey

vrf context management
  ip route 0.0.0.0/0 10.29.132.1
vlan 1
vlan 47
  name N1K-Mgmt
vlan 48
  name N1K-Ctrl
vlan 49
  name N1K-Pckt
vlan 50
  name Storage
vlan 51
  name vMotion
vlan 52
  name VDIAB
vlan 100
  name VDI
vlan 132
  name ESXi_Management
spanning-tree vlan 1,10-20,50-52,100,132 priority 16384
vpc domain 30
  role priority 4000
  peer-keepalive destination 10.29.132.6


interface Vlan1

interface Vlan50
  no shutdown
  ip address 10.10.50.3/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1

interface Vlan51
  no shutdown
  ip address 10.10.51.3/24
  hsrp version 2
  hsrp 51
    preempt
    priority 110
    ip 10.10.51.1
```

```
interface Vlan52
  no shutdown
  ip address 10.10.52.3/24
  hsrp version 2
  hsrp 52
    preempt
    priority 110
    ip 10.10.52.1

interface Vlan100
  no shutdown
  ip address 10.10.1.3/22
  hsrp version 2
  hsrp 100
    preempt
    priority 110
    ip 10.10.1.1

interface port-channel2
  untagged cos 5
  switchport access vlan 50
  vpc 2

interface port-channel30
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132
  spanning-tree port type network
  vpc peer-link

interface port-channel47
  description VPCforN1KVUplinks
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132
  spanning-tree port type edge trunk
  speed 10000
  vpc 47

interface Ethernet1/1
  switchport access vlan 132
  speed 1000

interface Ethernet1/2

interface Ethernet1/3
  description N1K-UplinkforInfraServers
  switchport access vlan 100

interface Ethernet1/4
  switchport mode trunk
  switchport access vlan 51
  switchport trunk allowed vlan 47-52,100,132
  speed 1000

interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132
```

```
interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132

interface Ethernet1/7
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132

interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132

interface Ethernet1/9
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132

interface Ethernet1/10
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13
  switchport access vlan 100
  speed 1000

interface Ethernet1/14
  switchport access vlan 100
  speed 1000

interface Ethernet1/15
  switchport access vlan 100
  speed 1000

interface Ethernet1/16
  switchport access vlan 100
  speed 1000

interface Ethernet1/17
  switchport access vlan 100
  speed 1000

interface Ethernet1/18
  switchport access vlan 100
  speed 1000

interface Ethernet1/19
  switchport access vlan 100
  speed 1000

interface Ethernet1/20
  switchport access vlan 100
  speed 1000

interface Ethernet1/21
interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30
  shutdown
  switchport mode fex-fabric
  fex associate 130

interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132
  channel-group 30 mode active

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 47-52,100,132
  channel-group 30 mode active

interface Ethernet2/1
  description VNX-5300-10GB-UPLINK
  switchport access vlan 50
  channel-group 2 mode active

interface Ethernet2/2
  description VNX-5300-10GB-UPLINK
  switchport access vlan 50

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface mgmt0
  ip address 10.29.132.7/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N1.1a.bin
boot system bootflash:/n5000-uk9.5.1.3.N1.1a.bin
```