



## Reference Architecture-Based Design for 5000 Seat Virtual Desktop Infrastructure

Citrix XenDesktop 5.6 Feature Pack 1 Built on Cisco Unified Computing System and EMC VNX

Cisco Unified Computing System, Nexus 5500, Nexus 1000V, EMC VNX7500, and VMware ESXi 5

Last Updated: February 7, 2013




Cisco  
Validated  
Design



Building Architectures to Solve Business Problems





## About the Authors

### **Mike Brennan, Sr. Technical Marketing Engineer, 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, 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.

### **Hardik Patel, Virtualization Systems Engineer, Cisco Systems**

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

### **Erick Arteaga, Sr. Software Test Engineer, Citrix Systems**

Erick Arteaga, a Senior Software Test Engineer with the Solutions Lab at Citrix Systems, Inc., has 7 years of experience in the IT industry, 5 years in Consulting services including Server and Desktop Virtualization deployments and maintenance.

### **Vadim Lebedev, Sr. Software Test Engineer, Citrix Systems**

Vadim Lebedev, a Senior Software Test engineer with the Solutions Lab at Citrix Systems, Inc., has 13 years of computer industry experience and has 5 years of experience on server and desktop virtualization. He began his tenure at Citrix as a member of the XenServer Escalation team handling complex and highly sensitive customer issues relating to XenServer and several other Citrix products.

### **Alexander Lyublinski, Test Architect and Staff Software Test Engineer, Citrix Systems**

Alexander Lyublinski, a Test Architect and Staff Software Test Engineer with the Solutions Lab at Citrix Systems, Inc., is a 20 year veteran of the computer industry and has over 10 years of experience in the virtualization area specializing in VDI solution architecture planning, deployment and testing. Alexander has written multiple knowledge base and white paper articles, as well as participated in the creation of several technical books.

### **Ka-kit Wong, Solutions Engineer, Strategic Solutions Engineering, EMC**

Ka-Kit Wong is a solutions engineer for desktop virtualization in EMC's Strategic Solutions Engineering group, where he focuses on developing End User Computing (EUC) validated solutions. He has been at EMC for more than 13 years, and his roles have included systems, performance and solutions testing in various positions. He holds a master of science degree in computer science from Vanderbilt University.



# 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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2013 Cisco Systems, Inc. All rights reserved.

# 1 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 Citrix XenDesktop 5.6 environment, utilizing Citrix Provisioning Server 6.1, on Cisco UCS B-Series B230 M2 Blade Servers running VMware ESXi 5.0 Update 1 hypervisor software connected to an EMC VNX7500 Storage Array. We utilize second generation Unified Computing System hardware and software. We provide best practice recommendations and sizing guidelines for large scale customer deployments of XenDesktop 5.6 on the Cisco Unified Computing System.

## 1.1 Solution Component Benefits

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

### 1.1.1 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 Cisco Unified Computing System

##### 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 flexible, agile, high performance, self-integrating information technology with faster ROI
- Fabric Extender technology 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
- Leverage Cisco UCS Management Packs for System Center 2012 for integrated management

### 1.1.2 Benefits of Nexus 5548UP

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

#### Architectural Flexibility

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

### 1.1.3 Benefits of Nexus 1000v Distributed Virtual Switch

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

#### **Policy-Based Virtual Machine Connectivity**

To complement the ease of creating and provisioning virtual machines, the Cisco Nexus 1000V Series includes the Port profile feature to address the dynamic nature of server virtualization from the network's perspective (Figure 2). Port profiles enable you to define virtual machine network policies for different types or classes of virtual machines from the VSM and then apply the profiles to individual virtual machine virtual NICs (vNICs) through the VMware vCenter GUI for transparent provisioning of network resources. Port profiles are a scalable mechanism for configuring networks with large numbers of virtual machines.

#### **Mobility of Virtual Machine Security and Network Properties**

Network and security policies defined in the port profile follow the virtual machine throughout its lifecycle, whether it is being migrated from one server to another (Figure 3), suspended, hibernated, or restarted. In addition to migrating the policy, the VSM also moves the virtual machine's network state, such as the port counters and flow statistics. Virtual machines participating in traffic monitoring activities, such as Cisco NetFlow or ERSPAN, can continue these activities uninterrupted by VMware vMotion operations. When a specific port profile is updated, the Cisco Nexus 1000V Series automatically provides live updates to all of the virtual ports using that same port profile. With the ability to migrate network and security policies through VMware vMotion, regulatory compliance is much easier to enforce with the Cisco Nexus 1000V Series, because the security policy is defined in the same way as physical servers and constantly enforced by the switch.

#### **Non-disruptive Operational Model**

Because of its close integration with VMware vCenter Server, the Cisco Nexus 1000V Series allows virtualization administrators to continue using VMware tools to provision virtual machines. At the same time, network administrators can provision and operate the virtual machine network the same way they do the physical network using Cisco CLI and SNMP along with tools such as ERSPAN and NetFlow (Figure 4). While both teams work independently, using familiar tools, the Cisco Nexus 1000V Series enforces consistent configuration and policy throughout the server virtualization environment. This level of integration lowers the cost of ownership while supporting various organizational boundaries among server, network, security, and storage teams.

#### **Differentiated Quality of Service**

Today, network interfaces are often dedicated to a particular type of traffic, such as VMware Console or vMotion. With the Cisco Nexus 1000V Series, all the network interface cards (NICs) on the server can be treated as a single logical channel with QoS attached to each type of traffic. With VMware vSphere Version 4.1, the Cisco Nexus 1000V Series can even provide different service-level agreements (SLAs) for production virtual machines. Consequently, the bandwidth to the server can be more efficiently utilized with virtualization of network-intensive applications.

### 1.1.4 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.

#### **VNXe Series—Simple, Efficient, Affordable**

The VNXe Series was designed with the IT generalist in mind and provides an integrated storage system for small-to-medium businesses as well as remote offices, and departments in larger enterprise businesses. Starting at less than \$8,000, the VNXe series provides true storage consolidation with a unique application-driven approach that eliminates the boundaries between applications and their storage.

This simple application-driven approach to managing shared storage makes the VNXe series ideal for IT generalists/managers and application administrators who may have limited storage expertise. EMC Unisphere for the VNXe series enables easy, wizard-based provisioning of storage for Microsoft, Exchange, file shares, iSCSI volumes, VMware, and Hyper-V. VNXe supports tight integration with VMware to further facilitate efficient management of virtualized environments. Complemented by Unisphere Remote, the VNXe is also ideal for remote office-branch office (ROBO) deployments. Built-in efficiency capabilities, such as file de-duplication with compression and thin provisioning result in streamlined operations and can save up to 50 percent in upfront storage costs. Software packs aimed at facilitating backup, remote data protection, and disaster recovery include features such as easy-to-configure application snapshots.

The VNXe series supports high availability by using redundant components – power supplies, fans, and storage processors – as well as dynamic failover and failback. Additionally, the VNXe series supports the ability to upgrade system software or hardware while the VNXe system is running. It also delivers single click access to a world of resources such as comprehensive online documentation, training, and how-to-videos to expand your knowledge and answer questions.

#### **VNX Series—Simple, Efficient, Powerful**

A robust platform for consolidation of legacy block storage, file-servers, and direct-attached application storage, the VNX series enables organizations to dynamically grow, share, and cost-effectively manage multi-protocol file systems and multi-protocol block storage access. The VNX Operating environment enables Microsoft Windows and Linux/UNIX clients to share files in multi-protocol (NFS and CIFS) environments. At the same time it supports iSCSI, Fiber Channel, and FCoE access for high bandwidth and latency-sensitive block applications. The combination of EMC Atmos Virtual Edition software and VNX storage supports object-based storage and enables customers to manage web applications from EMC Unisphere. The VNX series next generation storage platform is powered by Intel quad-core Xeon 5600 series with a 6 –Gb/s SAS drive back-end and delivers demonstrable performance improvements over the previous generation mid-tier storage:

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

### **1.1.5 Benefits of VMware ESXi 5.0**

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)

### 1.1.6 Benefits of Citrix XenDesktop and Provisioning Server

Citrix XenDesktop is a comprehensive desktop virtualization solution that includes all the capabilities required to deliver desktops, apps and data securely to every user in an enterprise. Trusted by the world's largest organizations, XenDesktop has won numerous awards for its leading-edge technology and strategic approach to desktop virtualization.

Citrix XenDesktop helps businesses:

- Enable virtual workstyles to increase workforce productivity from anywhere
- Leverage the latest mobile devices to drive innovation throughout the business
- Rapidly adapt to change with fast, flexible desktop and app delivery for off-shoring, M&A, branch expansion and other initiatives
- Transform desktop computing with centralized delivery, management and security

A complete line of XenDesktop editions lets you choose the ideal solution for your business needs and IT strategy. XenDesktop VDI edition, a scalable solution for delivering [virtual desktops](#) in a VDI scenario, includes [Citrix HDX technology](#), provisioning services, and profile management. XenDesktop Enterprise edition is an enterprise-class desktop virtualization solution with [FlexCast delivery technology](#) that delivers the right type of virtual desktop with on-demand applications to any user, anywhere. The comprehensive Platinum edition includes advanced management, monitoring and security capabilities.

## 1.2 Audience

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

## 2 Summary of Main Findings

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

- The combined power of the Unified Computing System, Nexus switching and EMC storage hardware with VMware ESXi 5.0 Update 1, Citrix Provisioning Server 6.1, and Citrix XenDesktop 5.6 software produces a high density per blade and per chassis Virtual Desktop delivery system.
- B230 M2 half-width blade with dual 10-core processors and 256GB of memory supports 22.7% more virtual desktop workloads than the previously studied full width blade using a new medium workload with flash.
- The study design based on five Unified Computing System chassis, four each with 8 - B230 M2 blades and one with 7-B230 M2 blades, each with dual 10-core processors and 256GB of memory and a M81KR (Palo) converged network adapter supports 5000 virtual desktop workloads running the new medium workload with flash, more than 2.25 times the density of previously studied chassis with full width blades.
- The 39 B230 M2 Servers were organized in three VMware ESXi Clusters, each with N+1 Server fault tolerance. This highly available configuration is a first for our Cisco Validated Design series.
- We were able to ramp (log in and start workloads) up to steady state in 30 minutes without pegging the processor, exhausting memory or storage subsystems.
- Compared to previous studies with full width blades, the rack space required to support 5000 users was reduced from 72 Rack Units to 30 Rack units.
- Pure Virtualization: We continue to present a validated design that is 100% virtualized on ESXi 5.0 Update 1. All of the Windows 7 SP1 virtual desktops and supporting infrastructure components, including Active Directory, Profile Servers, Provisioning Servers, SQL Servers, and XenDesktop delivery controllers were hosted as virtual servers.
- We maintain our industry leadership with our new Cisco UCS Manager 2.0(4a) software that makes scaling simple, consistency guaranteed and maintenance simple.
- Our 10G unified fabric story gets additional validation on second generation 6200 Series Fabric Interconnects and second generation Nexus 5500 Series access switches as we run more challenging workload testing, maintaining unsurpassed user response times.
- For the first time in a Cisco VDI Cisco Validated Design, we incorporate our Nexus 1000V virtual distributed switch. A redundant pair of Nexus 1000V Virtual Supervisor Module appliances were deployed for each ESXi cluster to extend our Cisco end-to-end Quality of Service leadership, guaranteeing high priority virtual desktop related traffic to produce industry leading end user experience
- EMC's VNX 7500 system provides storage consolidation and outstanding efficiency. Both block and NFS storage resources were provided by a single system, utilizing EMC Fast Cache technology.
- Citrix HDX technology, extended in XenDesktop 5.6 Feature Pack 1 software, provides excellent performance with host-rendered flash video and other demanding applications.

## 3 Architecture

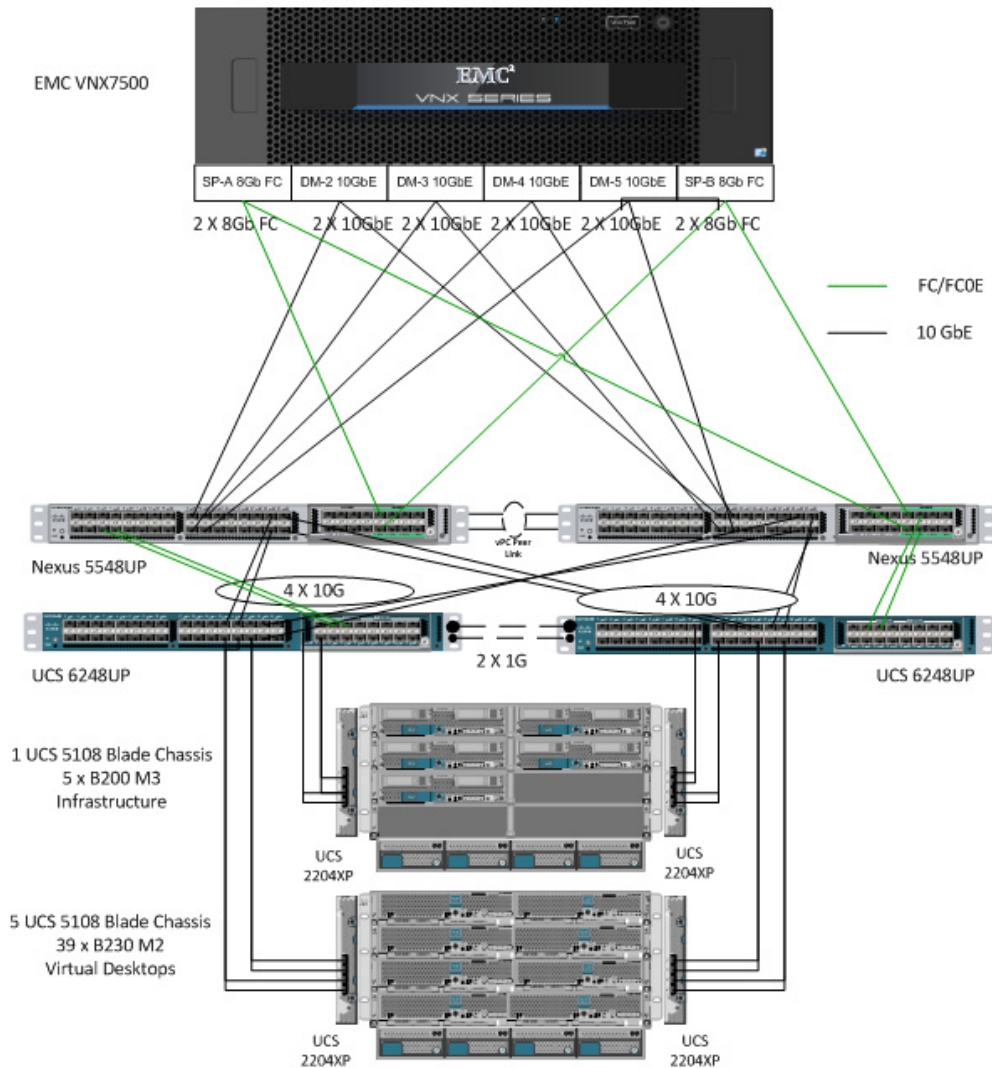
### 3.1 Hardware Deployed

The architecture deployed is highly modular. While each customer's environment might vary in its exact configuration, when 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 5000 User XenDesktop 5.6 solution includes Cisco networking, Cisco Unified Computing System and EMC storage which the computing and storage that fit in two data center racks, including the networking residing in the Cisco Unified Computing System rack.

This document details the deployment of Citrix XenDesktop 5.6 with Provisioning Server 6.1 on VMware ESXi 5.0 Update 1. Cisco Nexus 1000V distributed switch manages the three VMware Clusters hosting the virtual desktops, insuring end to end Quality of Service and ease of management by the network team.

**Figure 1 Citrix XenDesktop 5.6 5000 User Hardware Components**



The reference configuration includes:

- Two Cisco Nexus 5548UP switches with 16-universal port Expansion Modules
- Two Cisco UCS 6248UP Series Fabric Interconnects with Cisco UCS 6200 16-universal port Expansion Modules
- Six Cisco UCS 5108 Blade Server Chassis with two 2204XP IO Modules per chassis
- Five Cisco UCS B200 M3 Blade servers with Intel E5-2650 processors, 96 GB RAM, and VIC1240 mezzanine cards for infrastructure services

- Thirty nine Cisco UCS B230 M2 Blade servers with Intel E7-2780 processors, 256 GB RAM, and M81KR mezzanine cards for VDI workloads
- One EMC VNX7500 dual controller storage system for HA

The EMC VNX7500 disk shelf, disk and Fast Cache configurations are detailed in Section 5.4 Storage Architecture Design later in this document.

## 3.2 Software Revisions

**Table 1** Software Used in this Deployment

Layer	Compute	Version or Release	Details
Compute	Cisco UCS Fabric Interconnect	2.0(4a)	Embedded Management
	Cisco UCS B200 M2	2.0(4a)	Hardware BIOS
	Cisco UCS B230 M2	2.0(4a)	Hardware BIOS
Network	Nexus Fabric Switch	5.2(1)N1(1)	Operating System Version
Storage	EMC VNX7500	File: 7.1.47-5 Block: 05.32.000.5.006	Operating System Version
Software	Cisco UCS Blade Hosts	B200: VMware ESXi 5.0 Update 1 B230: VMware ESXi 5.0 Update 1	Operating System Version
	Cisco Nexus 1000V	4.2(1)SV1(5.2)	Virtual Switch appliance version

## 3.3 Configuration Guidelines

The 5000 User XenDesktop 5.6 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 Citrix XenDesktop 5.6 with Provisioning Server 6.1 customer environment as stand-alone solution.

### 3.3.1 VLANs

For the 5000 User XenDesktop 5.6 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
ML-VDA	800	Virtual Desktops	No
ML_DC-VM-MGMT	801	ESXi, N1KV Management	Yes
ML_DC-VMMOTION	802	vMotion	No
ML_DC-INF	803	Infrastructure VMs	No
ML_DC-STRG	804	NFS Storage	No
ML_Launcher-Inf	851	Login VSI Launchers	No
ML-N1KV_CTRL	900	N1KV Control	No
ML-N1KV_PK	901	N1KV Packet	No

### 3.3.2 VMware Clusters

We utilized five VMware Clusters to support the solution and testing environment:

- Infrastructure (Active Directory, DNS, DHCP, SQL Clusters, Citrix User Profile Manager clustered shares, PVS 6.1 virtual machines, XenDesktop controllers, Nexus 1000V Virtual Switch Manager appliances, etc.)
- VDA Clusters (3) (Windows 7 SP1 32-bit pooled virtual desktops; approximately 1700 per cluster per Nexus 1000V best practices recommended VSM density.)
- Launcher Cluster (The Login Consultants Login VSI launcher infrastructure was hosted on a completely separate Cisco UCS Domain using dedicated switching and storage. It was connected to the solution Cisco UCS Domain through the Nexus 5000 switches in each domain.)

## 4 Infrastructure Components

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

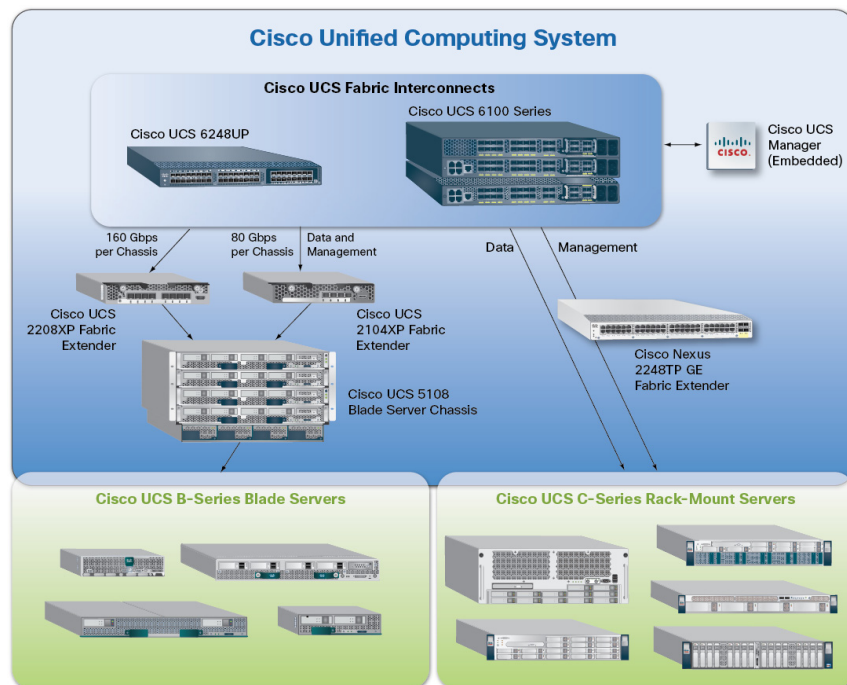
### 4.1 Cisco Unified Computing System (UCS)

[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.

#### 4.1.1 Cisco Unified Computing System Components

Cisco UCS components are shown in Cisco Unified Computing System Components.



**Figure 2 Cisco Unified Computing System Components**

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

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

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

### 4.1.2 Fabric Interconnect

Cisco UCS Fabric Interconnects create a unified network 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 UCS.

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

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

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

#### Cisco UCS 6248UP 48-Port Fabric Interconnect

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

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

### 4.1.3 Cisco UCS 2200 Series IO Module

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



**Note**

---

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

---

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

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

### 4.1.4 Cisco UCS Chassis

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

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

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

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear, and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel
- 4.1.5 Cisco UCS B200 M3 Blade Server

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

#### 4.1.5 Cisco UCS B200 M3 Blade Server

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

#### 4.1.6 Cisco UCS B230 M2 Blade Server

The UCS B230 M2 Blade Server is a full-slot, two-socket blade server featuring the performance and reliability of Intel Xeon Processor E7-2800 product family and up to 32 DIMM slots which support up to 512 GB of memory. The Cisco UCS B230 M2 supports two SSD drives and one CNA mezzanine slots for up to 20 Gbps of I/O throughput. The Cisco UCS B230 M2 Blade Server platform delivers outstanding performance, memory and I/O capacity to meet the diverse needs of a virtualized environment with advanced reliability and exceptional scalability for the most demanding applications.


#### 4.1.7 Cisco UCS Converged Network Adapter (CNA)

##### Cisco UCS M81KR Virtual Interface Card

A Cisco® innovation, the Cisco UCS M81KR (Palo) Virtual Interface Card (VIC) is a virtualization-optimized Fibre Channel over Ethernet (FCoE) mezzanine card designed for use with Cisco UCS B-Series Blade Servers (Figure 1). The VIC is a dual-port 10 Gigabit Ethernet mezzanine card that supports up to 128 Peripheral Component Interconnect Express (PCIe) standards-compliant

virtual interfaces that can be dynamically configured so that both their interface type (network interface card [NIC] or host bus adapter [HBA]) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. In addition, the Cisco UCS M81KR supports Cisco VN-Link technology, which adds server-virtualization intelligence to the network.

**Figure 3 Cisco UCS VIC M81KR Converged Network Adapter**





**Note**

The Cisco UCS M81KR virtual interface cards are deployed in the Cisco UCS B-Series B230 M2 Blade Servers.

**Cisco UCS VIC1240 Converged Network Adapter**

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





The Cisco UCS VIC 1240 enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1240 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.



**Note**

The Cisco UCS VIC1240 virtual interface cards are deployed in the Cisco UCS B-Series B200 M3 Blade Servers.

## 4.2 Citrix XenDesktop

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to any user, anywhere. With FlexCast™ delivery technology, XenDesktop can quickly and securely deliver individual applications or complete desktops to the entire enterprise, whether users are task workers, knowledge workers or mobile workers. Users now have the flexibility to access their desktop on any device, anytime, with a high definition user experience. With XenDesktop, IT can manage single instances of each OS, application, and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management. XenDesktop's open architecture enables customers to easily adopt desktop virtualization using any hypervisor, storage, or management infrastructure.

### 4.2.1 Enhancements in Citrix XenDesktop 5.6 Feature Pack 1

Citrix XenDesktop 5.6 Feature Pack 1, builds on the themes of the last release which are about reducing cost and making it easier to do desktop virtualization. Below, is an overview of new or updated technologies and capabilities contained in Feature Pack 1:



- **Remote PC**—Extends the FlexCast physical delivery model to include secure remote connections to office-based PCs with a high-definition user experience leveraging Receiver and HDX technologies. Simple auto-assignment setup is included so Remote PC can be easily provisioned to thousands of users. With this new FlexCast delivery feature, Citrix is simplifying desktop transformation by creating an easy on-ramp for desktop virtualization. [View the Remote PC video](#)
- **Universal Print Server**—Combined with the previously available Universal Print Driver, administrators may now install a single driver in the virtual desktop image or application server to permit local or network printing from any device, including thin clients and tablets.
- **Optimized Unified Communications**—A new connector from Citrix enables Microsoft Lync 2010 clients to create peer-to-peer connections for the ultimate user experience, while taking the load off datacenter processing and bandwidth resources. The Cisco Virtualization Experience Client (VXC), announced October 2011, was the first in the industry to provide peer-to-peer connection capability to deliver uncompromised user experience benefits to Citrix customers. [Download Cisco VXC](#). Webcam Video Compression adds support for WebEx (in addition to Office Communicator, GoToMeeting HDFaces, Skype and Adobe Connect).
- **Mobility Pack for VDI**—With the new Mobility Pack, XenDesktop dynamically transforms the user interfaces of Windows desktops and applications to look and feel like the native user interface of smartphones and tablets. Now, your existing Windows applications adapt to the way users interact with applications on smaller devices without any source code changes. Previously, this technology was available only for XenApp.
- **HDX 3D Pro**—This HDX update provides breakthrough visual performance of high-end graphics intensive applications obtained by producing much faster frame rates using NVIDIA's latest API and leveraging a new, ultra-efficient, deep compression codec.
- **XenClient Enterprise**—XenClient 4.1 now supports 9x more PCs, has wider graphics support with NVIDIA graphics & has broader server hypervisor support. Its backend management can now run on XenServer, vSphere & Hyper-V. The release brings robust policy controls to the platform & role based administration. XenClient 4.1 delivers enterprise level scalability with support of up to 10,000 endpoints.
- **Simple License Service**—This new license service automatically allocates and installs your XenDesktop and/or XenApp licenses directly from your license server, eliminating the need to go to My Citrix to fully allocate your licenses. For more details, reference [Citrix edocs](#). Version 11.6.1 or higher of the License Server is required.

## 4.2.2 FlexCast Technology

Citrix XenDesktop with FlexCast is an intelligent delivery technology that recognizes the user, device, and network, and delivers the correct virtual desktop and applications specifically tailored to meet the performance, security, and flexibility requirements of the user scenario. FlexCast technology delivers any type of virtual desktop to any device and can change this mix at any time. FlexCast also includes on-demand applications to deliver any type of virtual applications to any desktop, physical or virtual.

The FlexCast delivery technologies can be broken down into the following categories:

- **Hosted Shared Desktops** provide a locked down, streamlined and standardized environment with a core set of applications, ideally suited for task workers running a few lower-intensity applications with light personalization requirements
- **Hosted VM Desktops** offer a personalized Windows desktop experience, typically needed by knowledge workers with higher application performance needs and high personalization requirements

- **Streamed Virtual Hard Disk (VHD) Desktops** use the local processing power of rich clients while providing centralized single image management of the desktop. These types of desktops are often used in computer labs and training facilities and when users require local processing for certain applications or peripherals.
- **Local VM Desktops** utilize XenClient to extend the benefits of centralized, single-instance management to mobile workers that need to use their laptops offline. When they are able to connect to a suitable network, changes to the OS, applications, and user data are automatically synchronized with the data center.
- **Physical Desktops** utilize the Remote PC feature in XenDesktop to create secure remote connections to physical PCs on a LAN without having to build out a large scale XenDesktop infrastructure in the data center.
- **On-demand Applications** allows any Windows® application to be centralized and managed in the data center, hosted either on multi-user terminal servers or VMs and instantly delivered as a service to physical and virtual desktops. Optimized for each user device, network, and location, applications are delivered through a high speed protocol for use while connected or streamed through Citrix application virtualization or Microsoft App-V directly to the endpoint for use when offline.

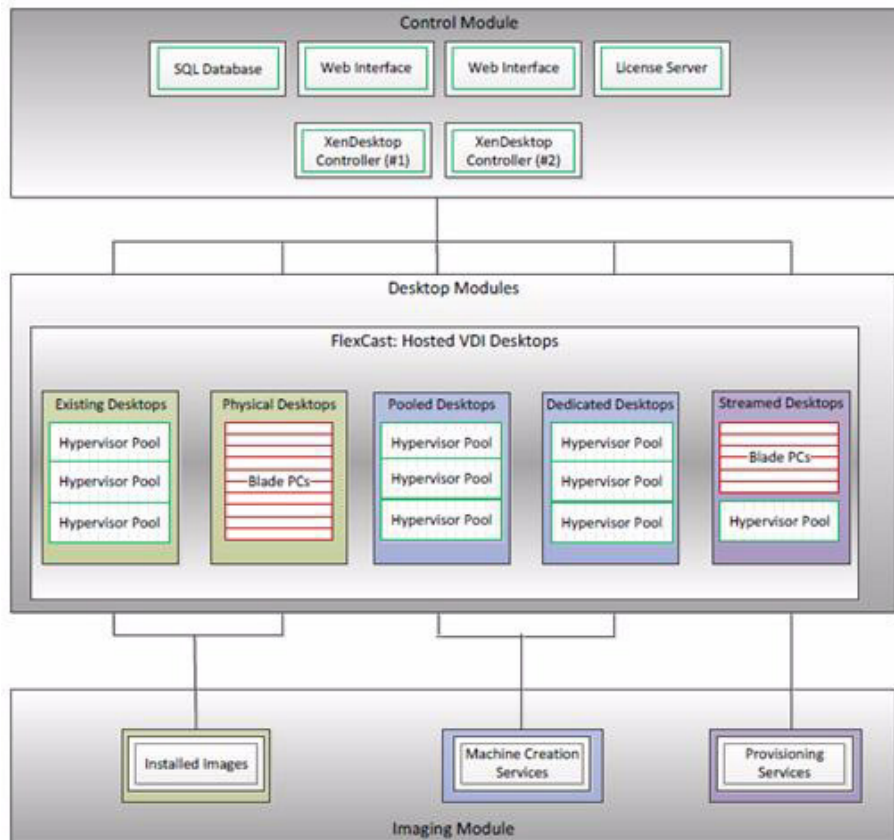
### 4.2.3 High-Definition User Experience Technology

Citrix High-Definition User Experience (HDX) technology is a set of capabilities that delivers a high definition desktop virtualization user experience to end users for any application, device, or network. These user experience enhancements balance performance with low bandwidth, whether it be plugging in a USB device, printing from a network printer or rendering real time video and audio. Citrix HDX technology provides network and application performance optimizations for a “like local PC” experience over LANs and a very usable experience over low bandwidth and high latency WAN connections.

### 4.2.4 Citrix XenDesktop Hosted VM Overview

Hosted VM uses a hypervisor to host all the desktops in the data center. Hosted VM desktops can either be pooled or assigned. Pooled virtual desktops use Citrix Provisioning Services to stream a standard desktop image to each desktop instance on boot-up. Therefore, the desktop is always returned to its clean, original state. Citrix Provisioning Services enables the streaming of a single desktop image to create multiple virtual desktops on one or more hypervisors in a data center. This feature greatly reduces the amount of storage required compared to other methods of creating virtual desktops. The high-level components of a Citrix XenDesktop architecture utilizing the Hosted VM model for desktop delivery are shown in Figure4.

Figure 4 Citrix XenDesktop on VMware vSphere



Components of a Citrix XenDesktop architecture using Hosted VM include:

- **Virtual Desktop Agent:** The Virtual Desktop Agent (VDA) is installed on the virtual desktops and enables direct Independent Computing Architecture (ICA) connections between the virtual desktop and user devices with the Citrix online plug-in.
- **Desktop Delivery Controller:** The XenDesktop controllers are responsible for maintaining the proper level of idle desktops to allow for instantaneous connections, monitoring the state of online and connected virtual desktops and shutting down virtual desktops as needed. The primary XD controller is configured as the farm master server. The farm master is able to focus on its role of managing the farm when an additional XenDesktop Controller acts as a dedicated XML server. The XML server is responsible for user authentication, resource enumeration, and desktop launching process. A failure in the XML broker service will result in users being unable to start their desktops. This is why multiple controllers per farm are recommended.
- **Citrix Receiver:** Installed on user devices, Citrix Receiver enables direct HDX connections from user devices to virtual desktops. Receiver is a mobile workspace available on a range of platforms so users can connect to their Windows applications and desktops from devices of their choice. Receiver for Web is also available for devices that don't support a native Receiver. Receiver incorporates the Citrix® HDX client engine and other technologies needed to communicate directly with backend resources, such as StoreFront.
- **Citrix XenApp:** Citrix XenApp is an on-demand application delivery solution that enables any Windows application to be virtualized, centralized, managed in the data center, and instantly delivered as a service to users anywhere on any device. XenApp can be used to deliver both virtualized applications and virtualized desktops. In the Hosted VM model, XenApp is typically used for on-demand access to streamed and hosted applications.

- **Provisioning Services:** PVS creates and provisions virtual desktops from a single desktop image (vDisk) on demand, optimizing storage utilization and providing a pristine virtual desktop to each user every time they log on. Desktop provisioning also simplifies desktop images, provides the best flexibility, and offers fewer points of desktop management for both applications and desktops. The Trivial File Transfer Protocol (TFTP) and Pre-boot eXecution Environment (PXE) services are required for the virtual desktop to boot off the network and download the bootstrap file which instructs the virtual desktop to connect to the PVS server for registration and vDisk access instructions.
- **Personal vDisk:** Personal vDisk technology is a powerful new tool that provides the persistence and customization users want with the management flexibility IT needs in pooled VDI deployments. Personal vDisk technology gives these users the ability to have a personalized experience of their virtual desktop. Personal apps, data and settings are easily accessible each time they log on. This enables broader enterprise-wide deployments of pooled virtual desktops by storing a single copy of Windows centrally, and combining it with a personal vDisk for each employee, enhancing user personalization and reducing storage costs.
- **Hypervisor:** XenDesktop has an open architecture that supports the use of XenServer, Microsoft Hyper-V, or VMware vSphere. For the purposes of the testing documented in this paper, VMware vSphere was the hypervisor of choice.
- **Storefront:** Storefront is the next-generation of Web Interface and provides the user interface to the XenDesktop environment. Storefront broker user authentication, enumerates the available desktops and, on launch, delivers an.ica file to Citrix Receiver on the user's local device to initiate a connection. Because StoreFront is a critical component, redundant servers must be available to provide fault tolerance.
- **License Server:** The Citrix License Server is responsible for managing the licenses for all of the components of XenDesktop. XenDesktop has a 90 day grace period which allows the system to function normally for 90 days if the license server becomes unavailable. This grace period offsets the complexity involved with building redundancy into the license server.
- **Data Store:** Each XenDesktop farm requires a database called the data store. Citrix XenDesktops use the data store to centralize configuration information for a farm in one location. The data store maintains all the static information about the XenDesktop environment.
- **Domain Controller:** The Domain Controller hosts Active Directory, Dynamic Host Configuration Protocol (DHCP), and Domain Name System (DNS). Active Directory provides a common namespace and secure method of communication between all the servers and desktops in the environment. DNS provides IP Host name resolution for the core XenDesktop infrastructure components. DHCP is used by the virtual desktop to request and obtain an IP address from the DHCP service. DHCP uses Option 66 and 67 to specify the bootstrap file location and file name to a virtual desktop. The DHCP service receives requests on UDP port 67 and sends data to UDP port 68 on a virtual desktop. The virtual desktops then have the operating system streamed over the network utilizing Citrix Provisioning Services (PVS).
- All of the aforementioned components interact to provide a virtual desktop to an end user based on the FlexCast Hosted VM desktop delivery model leveraging the Provisioning Services feature of XenDesktop. This architecture provides the end user with a pristine desktop at each logon based on a centralized desktop image that is owned and managed by IT.

#### 4.2.5 Citrix XenDesktop Hosted Shared Desktop Overview

In a typical large enterprise environment, IT will implement a mixture of Flexcast technologies to meet various workstyle needs. Like the test in this document, hosted shared desktops can be deployed alongside hosted VM desktops.

Host shared desktops has been a proven Citrix offering over many years and is deployed in some of the largest enterprises today due to its ease of deployment, reliability and scalability. Hosted shared desktops are appropriate for environments that have a standardized set of applications that do not deviate from one user to another. All users share the same desktop interface hosted on a Windows server in the backend datacenter. Hence, the level of desktop customization is limited compared to a Hosted VM desktop model.

If VM isolation is required and the ability to allocate resources to one user over another is important, the Hosted VM desktop should be the model of choice.

## 4.2.6 Citrix Provisioning Services

Citrix Provisioning Server provides images to physical and virtual desktops. Desktops utilize network booting to obtain the image and only portions of the desktop images are streamed across the network as needed. Provisioning Server does require additional server resources but these can be either physical or virtual servers depending on the capacity requirements and hardware configuration. Also, Provisioning Server does not require the desktop to be virtualized as Provisioning Server can deliver desktop images to physical desktops.

## 4.3 EMC VNX Series

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

- Next-generation unified storage, optimized for virtualized applications.
- Extended cache by using Flash drives with Fully Automated Storage Tiering for Virtual Pools (FAST VP) and FAST Cache that can be optimized for the highest system performance and lowest storage cost simultaneously on both block and file.
- Multiprotocol supports for file, block, and object with object access through EMC Atmos™ Virtual Edition (Atmos VE).
- Simplified management with EMC Unisphere™ for a single management framework for all NAS, SAN, and replication needs.
- Up to three times improvement in performance with the latest Intel Xeon multicore processor technology, optimized for Flash.
- 6 Gb/s SAS back end with the latest drive technologies supported:
  - 3.5" 100 GB and 200 GB Flash, 3.5" 300 GB, and 600 GB 15k or 10k rpm SAS, and 3.5" 1 TB, 2 TB and 3 TB 7.2k rpm NL-SAS
  - 2.5" 100 GB and 200 GB Flash, 300 GB, 600 GB and 900 GB 10k rpm SAS
- Expanded EMC UltraFlex™ I/O connectivity—Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Common Internet File System (CIFS), network file system (NFS) including parallel NFS (pNFS), Multi-Path File System (MPFS), and Fibre Channel over Ethernet (FCoE) connectivity for converged networking over Ethernet.

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

Software suites available:

- VNX FAST Suite—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (FAST VP is not part of the FAST Suite for VNX5100™).



- VNX Local Protection Suite—Practices safe data protection and repurposing.
- VNX Remote Protection Suite—Protects data against localized failures, outages, and disasters.
- VNX Application Protection Suite—Automates application copies and proves compliance.
- VNX Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

Software packs available:

- VNX Total Efficiency Pack—Includes all five software suites (not available for VNX5100).
- VNX Total Protection Pack—Includes local, remote, and application protection suites.
- VNX Total Value Pack—Includes all three protection software suites and the Security and Compliance Suite (VNX5100 exclusively supports this package).

### 4.3.1 EMC on VNX 7500 Used in Testing

EMC VNX 7500 unified storage platform is at the top of the line of the VNX series. It is powered by Intel quad-core Xeon 5600 series processors and delivers five 9's availability. It is designed to deliver maximum performance and scalability for enterprises, enabling them to dramatically grow, share, and cost-effectively manage multi-protocol file and block systems. It supports up to 1000 drives and eight X-Blades (also known as Data Movers) for file protocol support. This solution was validated using NFS for data storage of virtual desktops, and Fibre Channel for hypervisor SAN boot, SQL database, and infrastructure virtual machines such as Citrix XenDesktop controllers, VMware vCenter Servers, and other supporting services.

## 4.4 VMware ESXi 5.0

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.

### 4.4.1 VMware on ESXi 5.0 Hypervisor

ESXi 5.0 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.

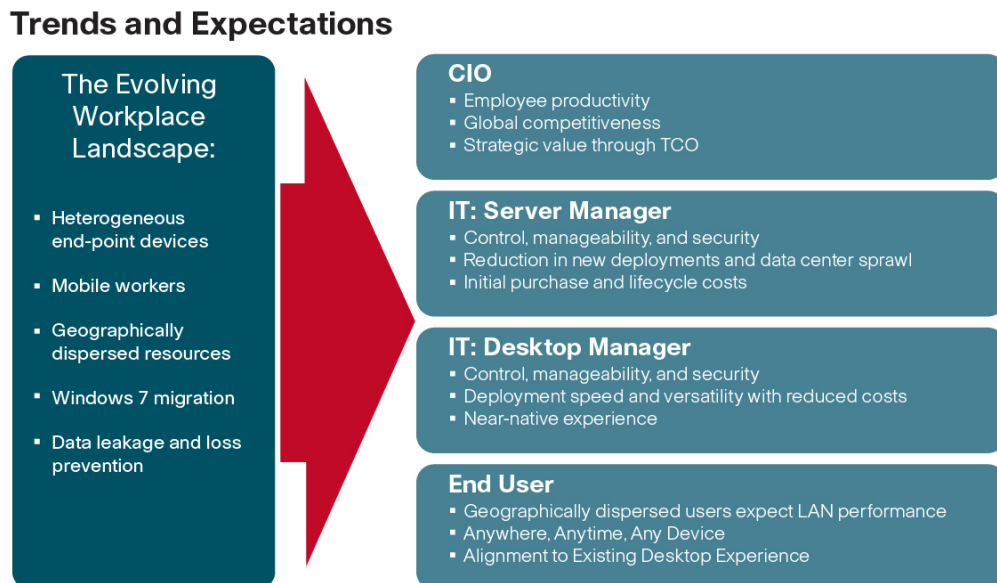
## 4.5 Modular Virtual Desktop Infrastructure Technical Overview

### 4.5.1 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 make sure protection of corporate data and to prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 10). These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 7.

**Figure 5 The Evolving Workplace Landscape**

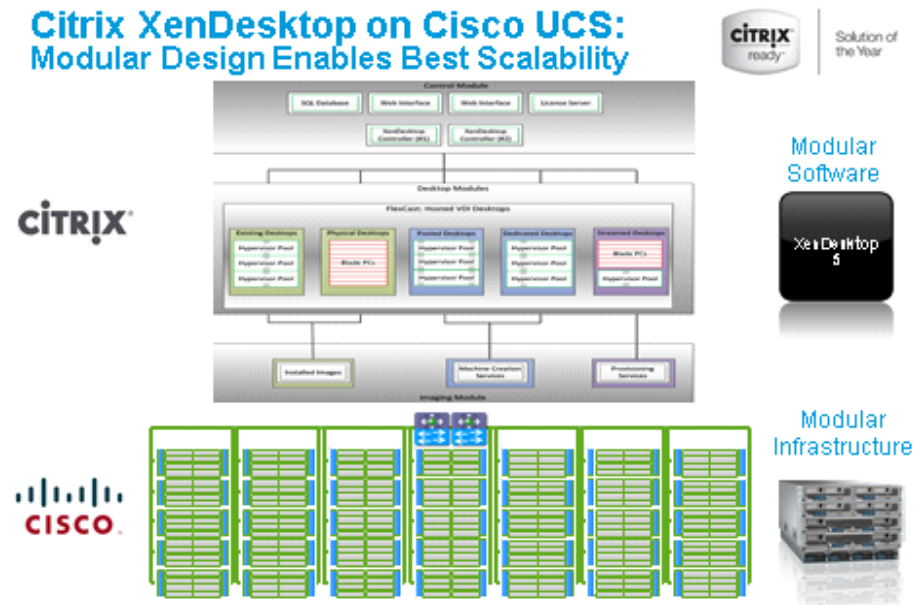


Some of the key drivers for desktop virtualization are increased data security and reduced TCO through increased control and reduced management costs.

#### 4.5.1.1 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 11).

Figure 6 Citrix XenDesktop on Cisco UCS



#### 4.5.1.2 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 –Citrix (XenDesktop and Provisioning Server) and EMC – have developed integrated, validated architectures, including available pre-defined, validated infrastructure packages, known as VSPEX.

#### 4.5.1.3 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.

#### 4.5.1.4 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.

#### 4.5.1.5 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, through 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.

### 4.5.2 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.

### 4.5.3 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.

### 4.5.4 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 XenApp be used for Hosted Shared Server Desktops or exclusively XenDesktop?
- Are there XenApp hosted applications planned? Are they packaged or installed?
- Will Provisioning Server or Machine Creation Services be used for virtual desktop deployment?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- Is user profile management (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?

### 4.5.5 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

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

To meet the challenges of designing and implementing a modular desktop infrastructure, Cisco, Citrix, 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

## 4.6 Cisco Networking Infrastructure

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

### 4.6.1 Cisco Nexus 5548 Switch

The Cisco Nexus 5548 Switch is a 1RU, 10 Gigabit Ethernet, FCoE access-layer switch built to provide more than 500 Gbps throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet and FCoE ports that accept modules and cables meeting the Small Form-Factor Pluggable Plus (SFP+) form factor.

One expansion module slot can be configured to support up to six additional 10 Gigabit Ethernet and FCoE ports, up to eight FC ports, or a combination of both. 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.

## 4.6.2 Cisco Nexus 5500 Series Feature Highlights

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

- The combination of high port density, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10 Gigabit Ethernet at the rack level. The switch family has sufficient port density to support single or multiple racks fully populated with blade and rack-mount servers.
- Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible. Hot-swappable power and cooling modules can be accessed from the front panel, where status lights offer an at-a-glance view of switch operation. Front-to-back cooling is consistent with server designs, supporting efficient data center hot-aisle and cold-aisle designs. Serviceability is enhanced with all customer replaceable units accessible from the front panel. The use of SFP+ ports offers increased flexibility to use a range of interconnect solutions, including copper for short runs and fibre for long runs.
- FCoE and IEEE data center bridging features support I/O consolidation, ease management of multiple traffic flows, and optimize performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5500 Series switches provide additional features that create an even more easily managed, high-performance, unified network fabric.

### 4.6.2.1 Features and Benefits

The following sections detail the specific features and benefits provided by the Cisco Nexus 5500 Series.

#### 4.6.2.2 10GB Ethernet, FCoE, and Unified Fabric Features

The Cisco Nexus 5500 Series is first and foremost a family of outstanding access switches for 10 Gigabit Ethernet connectivity. Most of the features on the switches are designed for high performance with 10 Gigabit Ethernet. The Cisco Nexus 5500 Series also supports FCoE on each 10 Gigabit Ethernet port that can be used to implement a unified data center fabric, consolidating LAN, SAN, and server clustering traffic.

#### 4.6.2.3 Low Latency

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

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

## 4.6.3 Nexus 1000V Virtual Distributed Switch

### 4.6.3.1 Product Overview

Cisco Nexus® 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking. The switches are designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner. Integrated into the VMware vSphere hypervisor and fully compatible with VMware vCloud Director, the Cisco® Nexus 1000V Series provides:

- Advanced virtual machine networking based on Cisco NX-OS operating system and IEEE 802.1Q switching technology
- Cisco vPath technology for efficient and optimized integration of virtual network services
- Virtual Extensible Local Area Network (VXLAN), supporting cloud networking

These capabilities help ensure that the virtual machine is a basic building block of the data center, with full switching capabilities and a variety of Layer 4 through 7 services in both dedicated and multitenant cloud environments. With the introduction of VXLAN on the Nexus 1000V Series, network isolation among virtual machines can scale beyond traditional VLANs for cloud-scale networking.

### 4.6.3.2 Advanced Virtual Machine Networking using Cisco Nexus 1000V Series

The Cisco Nexus 1000V Series Switches are virtual machine access switches for the VMware vSphere environments running the Cisco NX-OS operating system. Operating inside the VMware ESX or ESXi hypervisors, the Cisco Nexus 1000V Series provides:

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

When server virtualization is deployed in the data center, virtual servers typically are not managed the same way as physical servers. Server virtualization is treated as a special deployment, leading to longer deployment time, with a greater degree of coordination among server, network, storage, and security administrators. With the Cisco Nexus 1000V Series, you can have a consistent networking feature set and provisioning process all the way from the virtual machine access layer to the core of the data center network infrastructure. Virtual servers can now use the same network configuration, security policy, diagnostic tools, and operational models as their physical server counterparts attached to dedicated physical network ports. Virtualization administrators can access predefined network policies that follow mobile virtual machines to make sure proper connectivity saving valuable time to focus on virtual machine administration. This comprehensive set of capabilities helps you to deploy server virtualization faster and realize its benefits sooner.

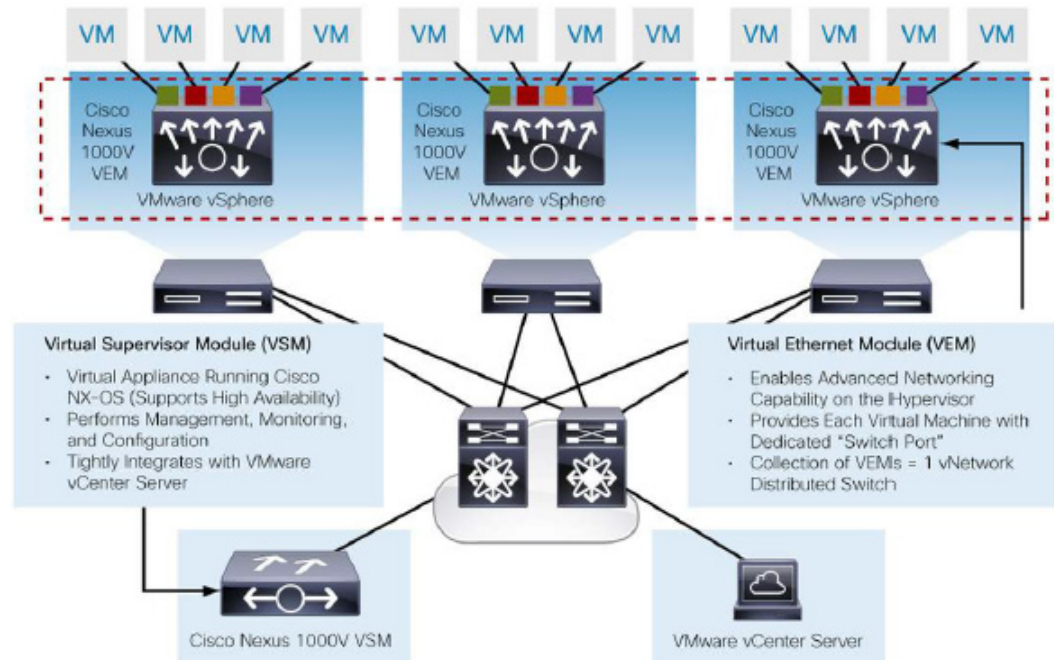
Developed in close collaboration with VMware, the Cisco Nexus 1000V Series is certified by VMware to be compatible with VMware vSphere, vCenter, vCloud Director, ESX, and ESXi, and many other VMware vSphere features. You can use the Cisco Nexus 1000V Series to manage your virtual machine connectivity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



### 4.6.3.3 Product Architecture

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

**Figure 7 Cisco Nexus 1000V Series Architecture**



#### Virtual Ethernet Module

The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware Virtual Switch functionality. The VEM uses the VMware vNetwork Distributed Switch (vDS) API, which was developed jointly by Cisco and VMware, to provide advanced networking capability to virtual machines. This level of integration helps ensure that the Cisco Nexus 1000V Series is fully aware of all server virtualization events, such as VMware VMotion and Distributed Resource Scheduler (DRS). The VEM takes configuration information from the VSM and performs Layer 2 switching and advanced networking functions:

- Port Channels
- Quality of service (QoS)
- Security: Private VLAN, access control lists (ACLs), and port security
- Monitoring: NetFlow, Switch Port Analyzer (SPAN), and Encapsulated Remote SPAN (ERSPAN)

In the event of loss of communication with the VSM, the VEM has Nonstop Forwarding (NSF) capability to continue to switch traffic based on the last known configuration. Thus, the VEM provides advanced switching with data center reliability for the server virtualization environment.

#### Virtual Supervisor Module

The Cisco Nexus 1000V Series VSM controls multiple VEMs as one logical modular switch. Instead of physical line-card modules, the VSM supports multiple VEMs running in software inside the physical servers. Configuration is performed through the VSM and is automatically propagated to the VEMs.

Instead of configuring soft switches inside the hypervisor on a host-by-host basis, administrators can define configurations for immediate use on all VEMs being managed by the VSM, from a single interface.

**By using the capabilities of Cisco NX-OS, the Cisco Nexus 1000V Series provides:**

- **Flexibility and scalability:** port profiles, a new Cisco NX-OS feature, provides configuration of ports by category, enabling the solution to scale to a large number of ports. Common software can run all areas of the data center network, including the LAN and SAN.
- **High availability:** Synchronized, redundant VSMs enable rapid, stateful failover and help ensure an always available virtual machine network.
- **Manageability:** The Cisco Nexus 1000V Series can be accessed through the Cisco command-line interface (CLI), Simple Network Management Protocol (SNMP), XML API, and Cisco Works LAN Management Solution (LMS).

The VSM is also integrated with VMware vCenter Server so that the virtualization administrator can take advantage of the network configuration in the Cisco Nexus 1000V Series.

## 5 Architecture and Design of Citrix XenDesktop 5.6 on Cisco Unified Computing System and EMC VNX Storage

### 5.1 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 desktop environments for each user:

- **Traditional PC:** A traditional PC is what —typically constituted a desktop environment: physical device with a locally installed operating system.
- **Hosted Shared Desktop:** A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2008 R2, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- **Hosted Virtual Desktop:** A hosted virtual desktop is a virtual desktop running either on virtualization layer (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.

## 5.2 Hosted VDI Design Fundamentals

Citrix XenDesktop 5.6 can be used to deliver a variety of virtual desktop configurations. When evaluating a Hosted VDI deployment, consider the following:

### 5.2.1 Hypervisor Selection

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

- **Hyper-V:** Microsoft Windows Server 2008 R2 Hyper-V builds on the architecture and functions of Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility. Hyper-V is available in a Standard, Server Core and free Hyper-V Server 2008 R2 versions. More information on Hyper-V can be obtained at the company web site.
- **vSphere:** VMware vSphere consists of the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. It offers features like Distributed resource scheduler, vMotion, HA, Storage vMotion, VMFS, and a multipathing storage layer. More information on vSphere can be obtained at the company website.
- **XenServer:** Citrix® XenServer® is a complete, managed [server virtualization](#) platform built on the powerful Xen® hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows® and Linux® [virtual servers](#) and delivers cost-effective server consolidation and business continuity. More information on Hyper-V can be obtained at the company website.

For this study, we utilized VMware ESXi 5.0 Update 1 and vCenter 5.0 Update 1.

## 5.2.2 Provisioning Services

Hosted-VDI desktops can be deployed with or without Citrix Provisioning Services (PVS.) However, PVS enables you to stream a single desktop image to create multiple virtual desktops on one or more servers in a data center. This facility greatly reduces the amount of storage required compared to other methods of creating virtual desktops. Virtual desktops streamed through Citrix Provisioning Services can be deployed as XenDesktop Pooled or Dedicated using the following types of virtual disks (vDisks):

- **PVS Private Mode vDisk:** A dedicated desktop is a desktop assigned to one distinct user. It uses a PVS vDisk that is set to private mode, enabling all changes to be captured when the virtual machine is restarted.
- **PVS Standard vDisk:** A pooled virtual desktop uses Citrix Provisioning Services to stream a standard or shared desktop image to multiple desktop instances on boot-up. Changes made to this type of PVS virtual disk is not preserved when the virtual desktop is restarted.

When considering a Provisioning Services deployment, there are some design decisions that need to be made regarding the write-cache for the virtual desktop device leveraging provisioning. The write-cache is a cache of all data that the target device has written. If data is written to the Provisioning Server vDisk in a caching mode, the data is not written back to the base vDisk. Instead it is written to a write-cache file in one of the locations specified below. The following options exist for the Provisioning Services write cache:

- **Cache on device HD:** Cache on device HD is stored in a file on a secondary local hard drive of the virtual device. It gets created as an invisible file in the root folder of the local HD. The Cache file size grows as needed, but never gets larger than the original vDisk, and frequently not larger than the free space on the original vDisk.
- **Cache in device RAM:** Cache is stored in client RAM (Memory), The Cache maximum size is fixed by a setting in vDisk properties. All written data can be read from local RAM instead of going back to server. RAM Cache is faster than server cache and works in a high availability environment.
- **Cache on server:** Device Cache is stored in a file on the PVS server, or on a share, SAN, or other. The file size grows as needed, but never gets larger than the original vDisk, and frequently not larger than the free space on the original vDisk. It is slower than RAM cache because all reads/writes have to go to the server and be read from a file. Cache gets deleted when the device reboots, in other words, on every boot the device reverts to the base image. Changes remain only during a single boot session.
- **Cache on device hard drive persisted:** (Experimental Phase) The same as Cache on device hard drive, except cache persists. At this time, this write cache method is an experimental feature only, and is only supported for NT6.1 or later (Windows 7 and Windows 2008 R2 and later). This method also requires a different bootstrap.
- **Cache on server persisted:** This cache option allows for the saving of changes between reboots. Using this option, after rebooting, a target device is able to retrieve changes made from previous sessions that differ from the read only vDisk image. If a vDisk is set to Cache on server persistent, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted on shutdown.

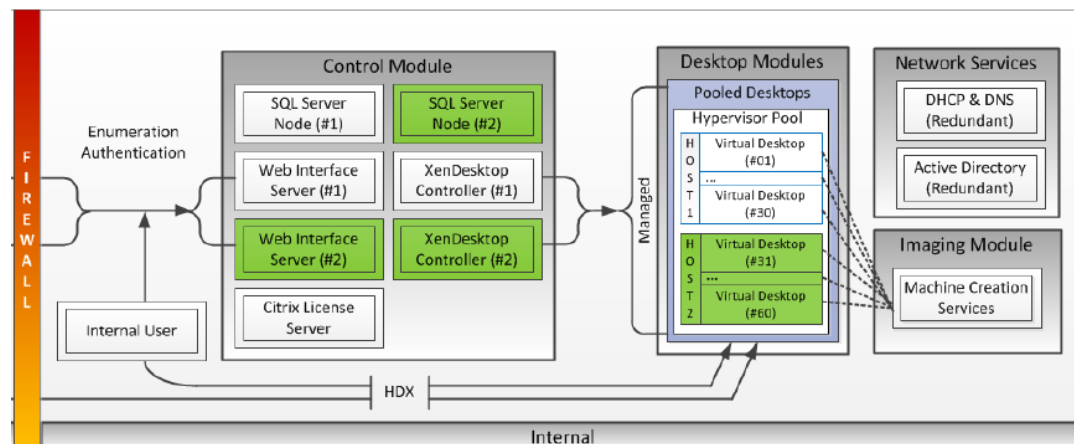
The alternative to Citrix Provisioning Services for pooled desktop deployments is Citrix Machine Creation Services, which is integrated directly with the XenDesktop Studio console.

For this study Provisioning Server 6.1 standard mode vDisks were used to stream 5000 XenDesktop Pooled Desktops with Cache on device HD. Provisioning Server 6.1 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

## 5.3 Designing a Citrix XenDesktop 5.6 Deployment

To implement our pooled desktop delivery model for this study, known as Hosted VDI Pooled Desktops, we followed the Citrix Reference Architecture for local desktop delivery.

**Figure 8 Pooled Desktop Infrastructure**



We used Provisioning Services 6.1 in place of Machine Creation Services shown in the figure above for this study.

To read about Citrix's XenDesktop Reference Architecture – Pooled Desktops (Local and Remote) go to the following link:

<http://support.citrix.com/article/CTX131049>

To learn more about XenDesktop 5.6 Planning and Design go to the following link:

<http://support.citrix.com/product/xd/v5.5/consulting/>

## 5.4 Storage Architecture Design

In a large scale PVS deployment, the option typically chosen for the PVS write cache destination is “Cache on device hard drive” (see 5.2.2 for other destination options) to achieve higher scalability, allow ease of manageability and agility when the write cache area resides in the EMC VNX unified storage system. In a virtualized environment, this cache area resides in the virtual hard disks that are attached to the virtual machines, and it accounts for the majority of IOPS requirements as the PVS servers absorb most of the read IOPS for the virtual desktops while all writes are redirected to the write cache area. Since the write Cache area is write intensive by nature, it is recommended to designate RAID 10 storage pools on the VNX for PVS write cache to minimize RAID penalty that are likely incurred by RAID 5 or similar RAID types. Because EMC VNX supports RAID 10 pools with multiples of eight drives, it is recommended to use that drive count unit as the building block while the number of deployed desktops continues to scale. In this solution, each building block storage pool is made up of 16 SAS drives.

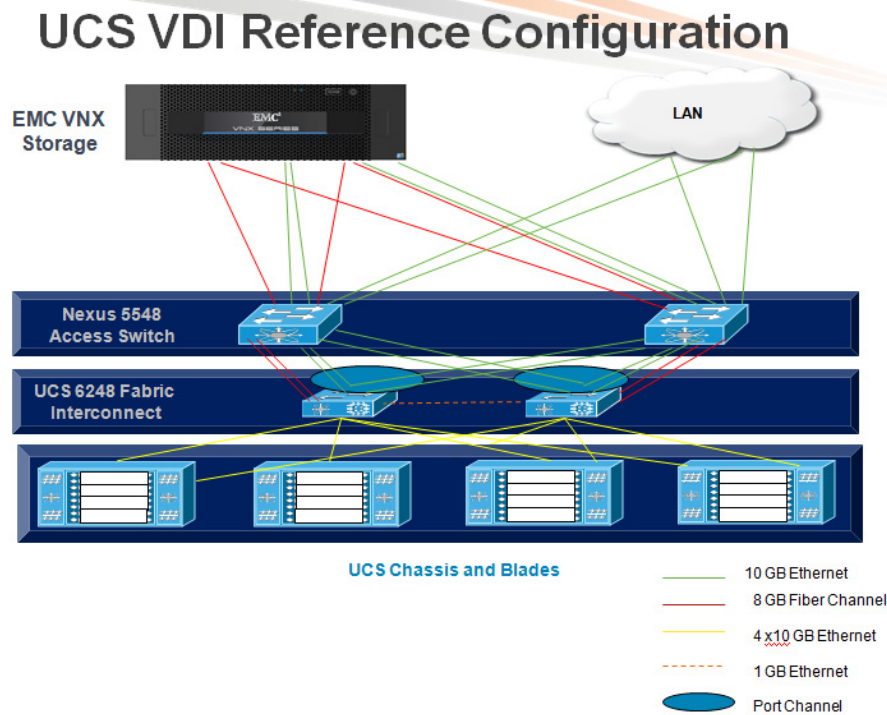
## 6 Solution Validation

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

## 6.1 Configuration Topology for Scalable Citrix XenDesktop 5.6 Virtual Desktop Infrastructure on Cisco Unified Computing System and EMC Storage

Figure 9

Architecture Block Diagram

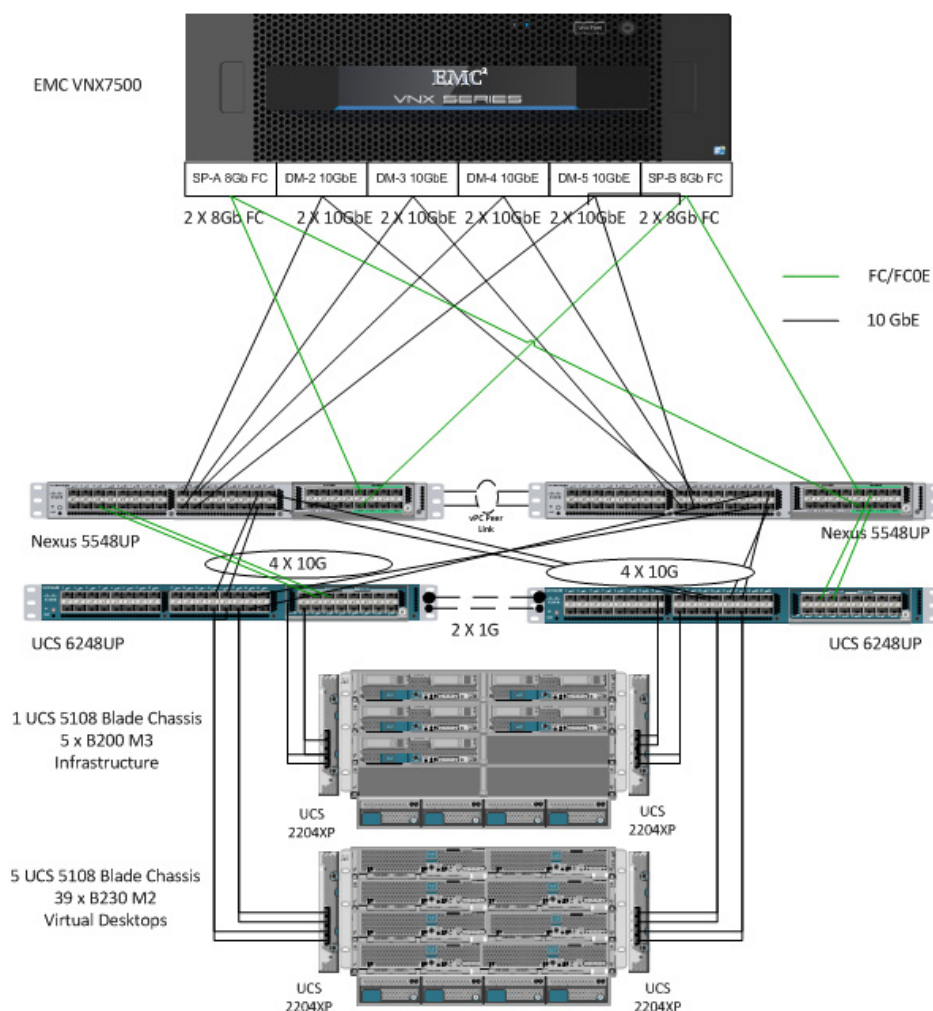


Architecture Block Diagram above captures the architectural diagram for the purpose of this study. The architecture is divided into four distinct layers:

- 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 5000 seat XenDesktop 5.6 environment.

**Figure 10 Detailed Architecture of the Configuration**



## 6.2 Cisco Unified Computing System Configuration

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](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:

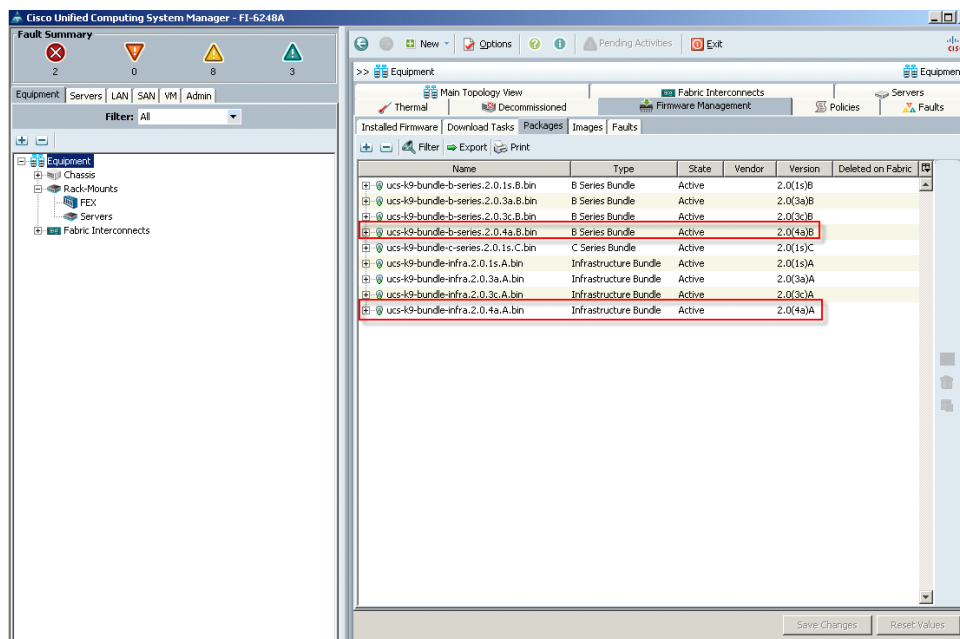
- Cisco UCS CLI Configuration guide  
[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/sw/cli/config/guide/2.1/b\\_UCSM\\_CLI\\_Configuration\\_Guide\\_2\\_1.pdf](http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.1/b_UCSM_CLI_Configuration_Guide_2_1.pdf)
- Cisco UCS-M GUI Configuration guide  
[http://www.cisco.com/en/US/partner/docs/unified\\_computing/ucs/sw/gui/config/guide/2.1/b\\_UCS\\_M\\_GUI\\_Configuration\\_Guide\\_2\\_1.html](http://www.cisco.com/en/US/partner/docs/unified_computing/ucs/sw/gui/config/guide/2.1/b_UCS_M_GUI_Configuration_Guide_2_1.html)



## 6.2.1 Base Cisco UCS System Configuration

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

1. Bring up the Fabric interconnect and from a Serial Console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. Now bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary and ask if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname. Once this is done all access to the FI can be done remotely. You will also configure the virtual IP address to connect to the FI, you need a total of three IP address to bring it online. You can also wire up the chassis to the FI, using either 1, 2 or 4 links per IO Module, depending on your application bandwidth requirement. We connected all the four links to each module.
2. Now connect using your favorite browser to the Virtual IP and launch the Cisco UCS-Manager. The Java based Cisco UCSM will let you do everything that you could do from the CLI. We will highlight the GUI methodology here.
3. First check the firmware on the system and see if it is current. Visit [http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.0\(4d\)&relind=AVAILABLE&rellifecycle=&reltype=latest](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 UCS Infrastructure and 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 2.0(4a).



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.



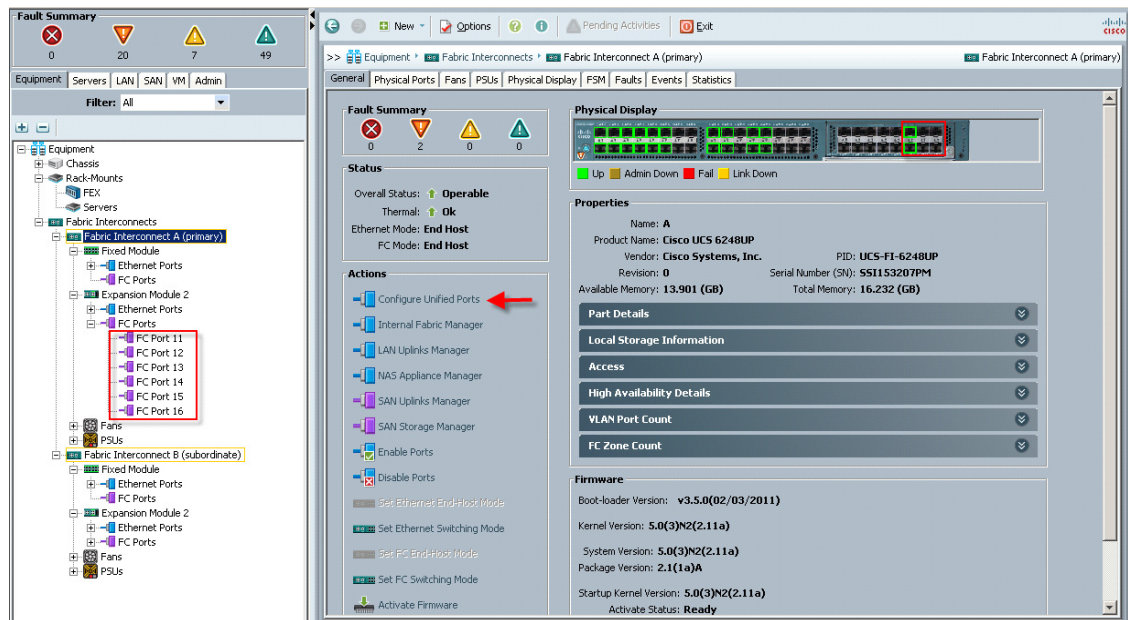
- Configure and enable the server ports on the FI. These are the ports that will connect the chassis to the FIs.

The screenshot shows the Cisco UCS Manager interface. On the left, the 'Fabric Interconnect A (primary)' tree is expanded, and 'Port 1' is selected. On the right, the 'Port 1' configuration page is displayed. The 'Status' section shows 'Overall Status: Up' and 'Admin State: Enabled'. The 'Properties' section shows 'ID: 1', 'Slot ID: 1', 'User Label', 'MAC: 54:7F:EE:45:2A:48', 'Mode: Fabric', 'Port Type: Physical', and 'Role: Server'. The 'Transceiver' section shows 'Type: H10GB CU3M', 'Model: 74752-9520', 'Vendor: CISCO-MOLEX', and 'Serial: MOC154205G0'. The 'License Details' section shows 'License State: License Ok' and 'License Grace Period: 0'.

- Configure and enable uplink Ethernet ports and FC uplink ports.

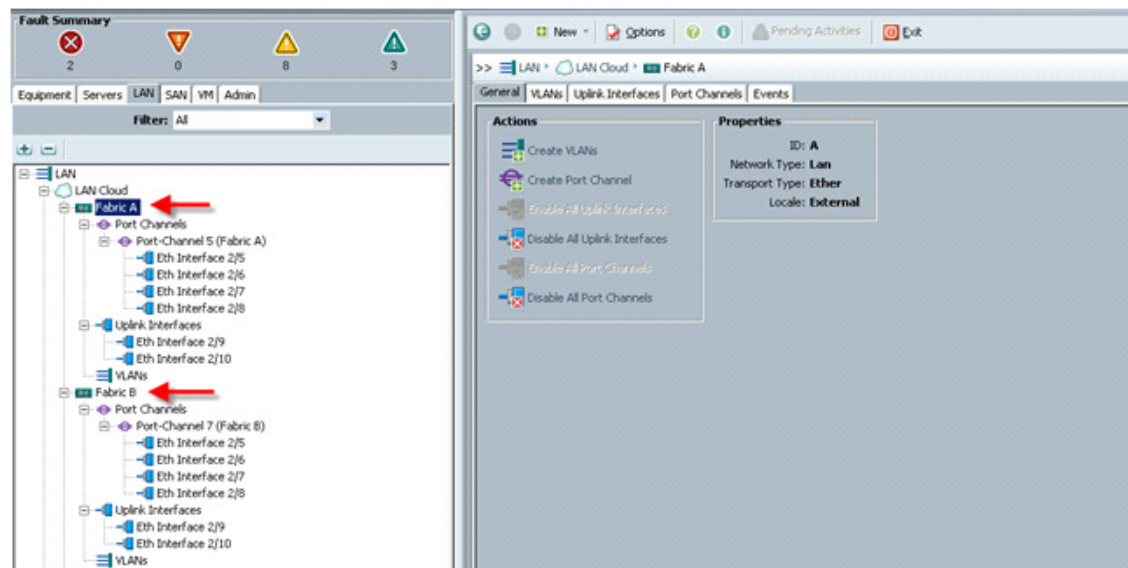
The screenshot shows the Cisco UCS Manager interface. On the left, the 'Fabric Interconnects' tree is expanded, and 'Ethernet Ports' is selected. On the right, the 'Ethernet Ports' configuration page is displayed. The table below shows the configuration for Ethernet ports:

Slot	Port ID	MAC	If Role	If Type	Overall Status	Administrative State
1	1	54:7F:EE:45:2A:48	Server	Physical	Up	Enabled
1	2	54:7F:EE:45:2A:49	Server	Physical	Up	Enabled
1	3	54:7F:EE:45:2A:4A	Server	Physical	Up	Enabled
1	4	54:7F:EE:45:2A:4B	Server	Physical	Up	Enabled
1	5	54:7F:EE:45:2A:4C	Server	Physical	Up	Enabled
1	6	54:7F:EE:45:2A:4D	Server	Physical	Up	Enabled
1	7	54:7F:EE:45:2A:4E	Server	Physical	Up	Enabled
1	8	54:7F:EE:45:2A:4F	Server	Physical	Up	Enabled
1	9	54:7F:EE:45:2A:50	Server	Physical	Up	Enabled
1	10	54:7F:EE:45:2A:51	Server	Physical	Up	Enabled
1	11	54:7F:EE:45:2A:52	Server	Physical	Up	Enabled
1	12	54:7F:EE:45:2A:53	Server	Physical	Up	Enabled
1	13	54:7F:EE:45:2A:54	Unconfigured	Physical	Sfp Not Present	Disabled
1	14	54:7F:EE:45:2A:55	Unconfigured	Physical	Sfp Not Present	Disabled
1	15	54:7F:EE:45:2A:56	Monitor	Physical	Sfp Not Present	Enabled
1	16	54:7F:EE:45:2A:57	Unconfigured	Physical	Sfp Not Present	Disabled
1	17	54:7F:EE:45:2A:58	Network	Physical	Up	Enabled
1	18	54:7F:EE:45:2A:59	Network	Physical	Up	Enabled
1	19	54:7F:EE:45:2A:5A	Network	Physical	Up	Enabled
1	20	54:7F:EE:45:2A:5B	Network	Physical	Up	Enabled
1	21	54:7F:EE:45:2A:5C	Unconfigured	Physical	Sfp Not Present	Disabled
1	22	54:7F:EE:45:2A:5D	Unconfigured	Physical	Sfp Not Present	Disabled

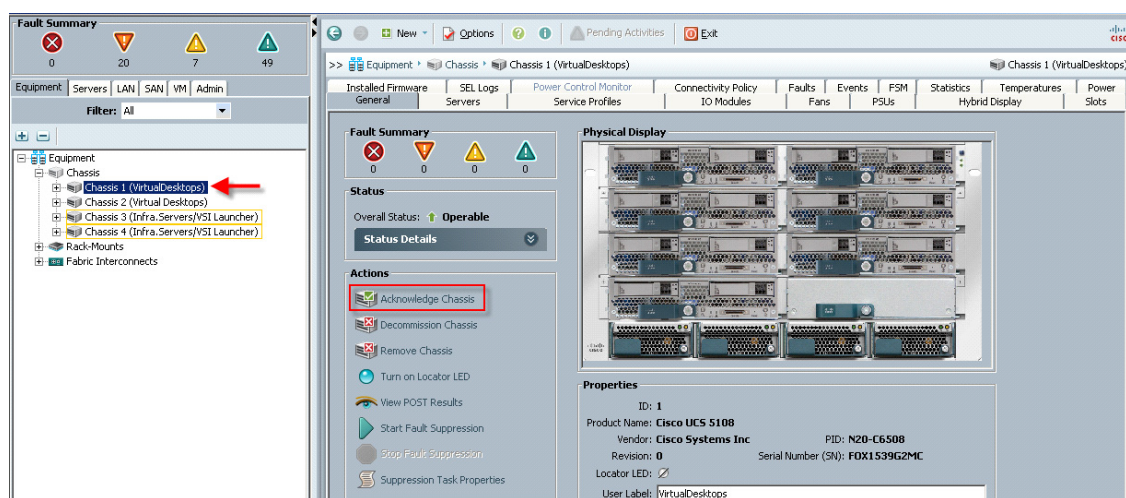


Use the Configure Unified Ports, Configure Expansion Module Ports to configure FC uplinks. Note: In this example, we configured six FC ports, two of which are in use.

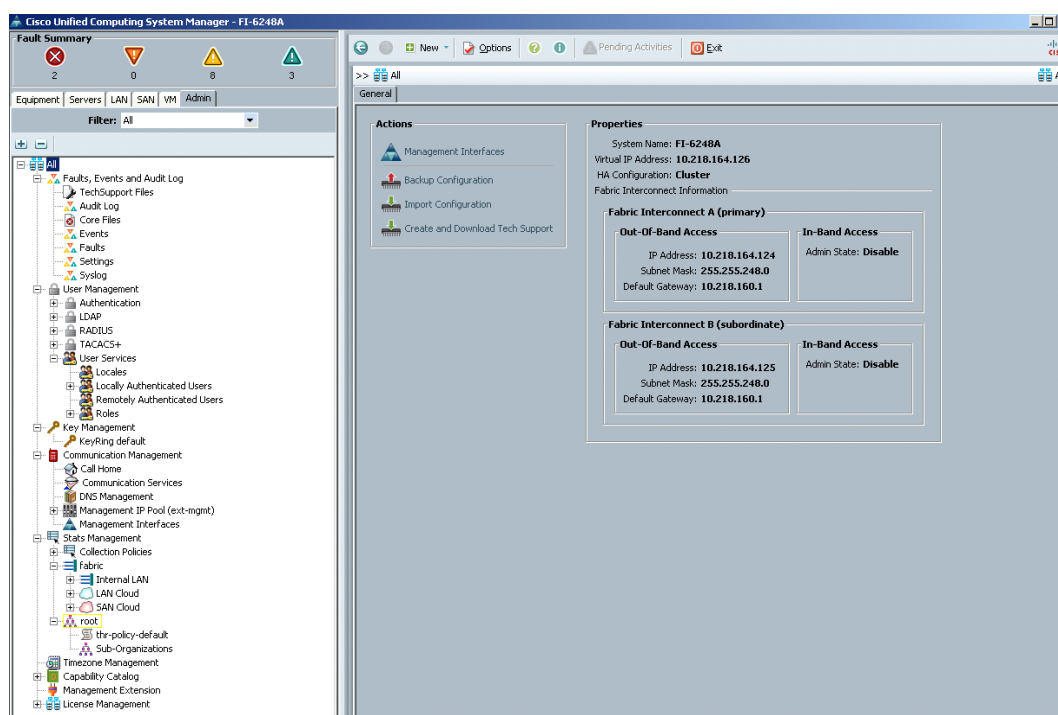
- a. On the LAN tab in the Navigator pane, configure the required Port Channels and Uplink Interfaces on both Fabric Interconnects.



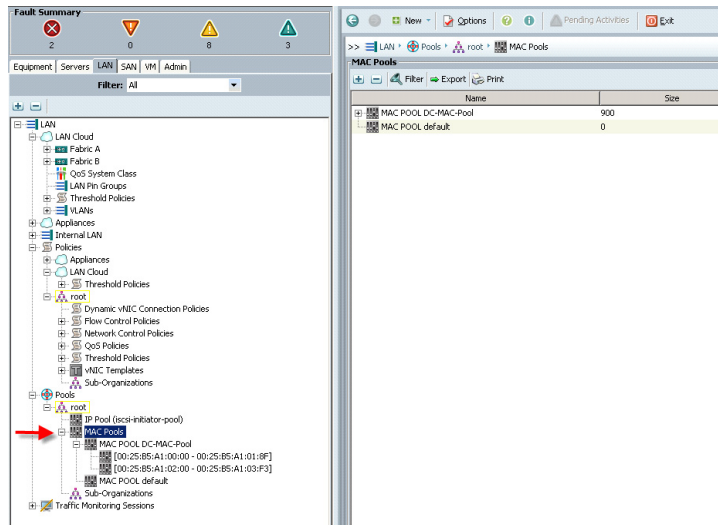
6. Expand the Chassis node in the left pane, then click on each chassis in the left pane, then click Acknowledge Chassis in the right pane to bring the chassis online and enable blade discovery.



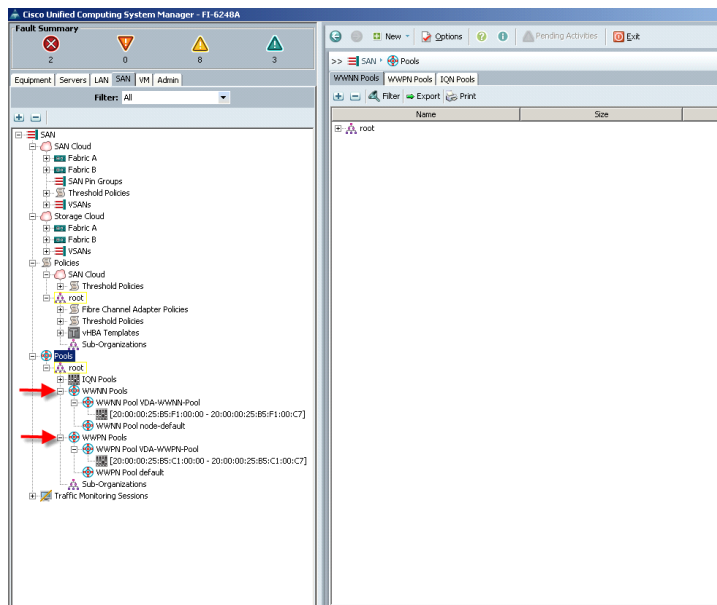
7. Use the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and Licensing. Configuring your Management IP Pool (which provides IP based access to the KVM of each Cisco UCS Blade Server,) Time zone Management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



8. Create all the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, Server pool.
  - a. 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.

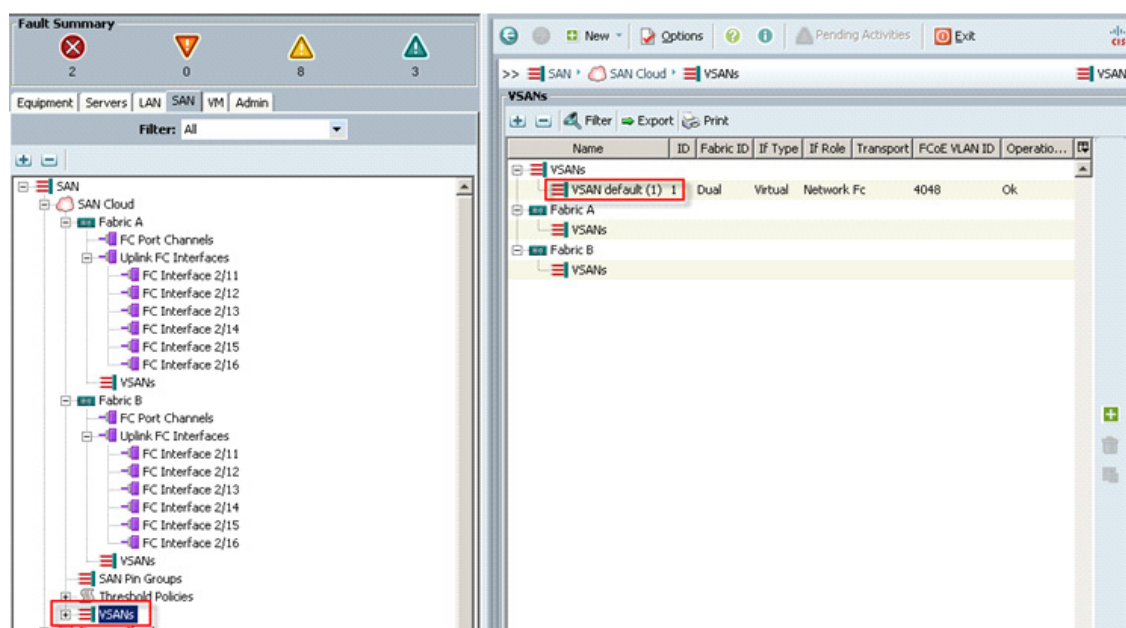


- b. For Fiber Channel connectivity, WWNN and WWPn pools must be created from the SAN tab in the navigator pane, in the Pools node:

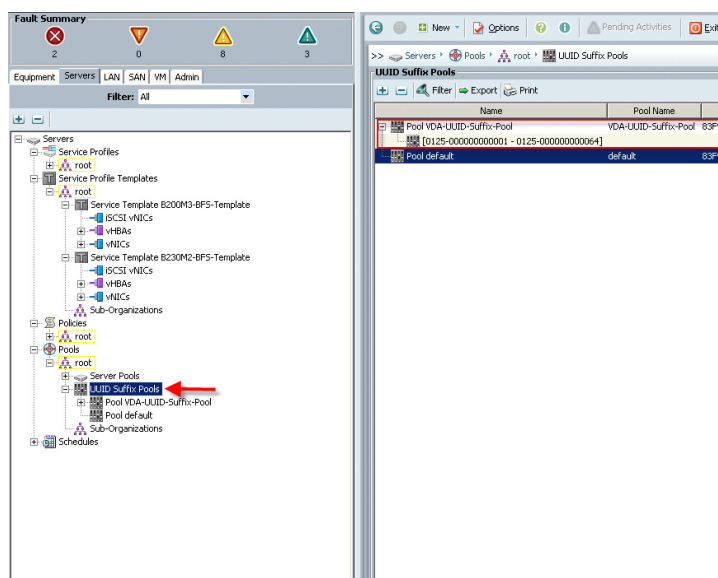


- c. For this project, we used a single VSAN, the default VSAN with ID 1:

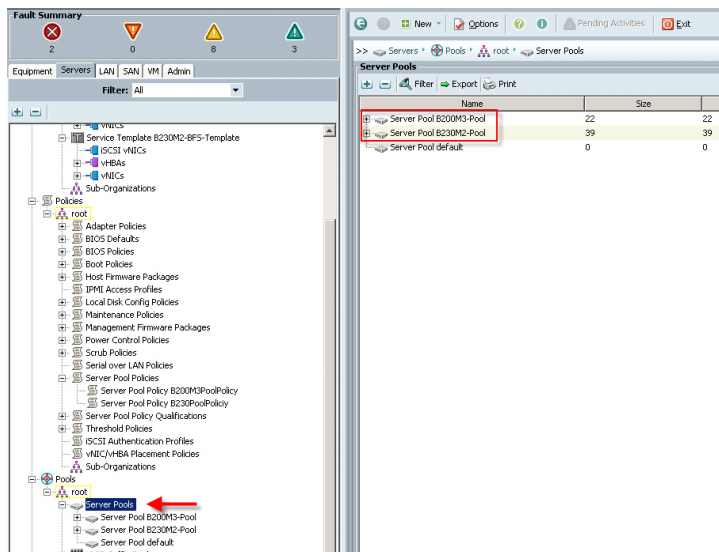




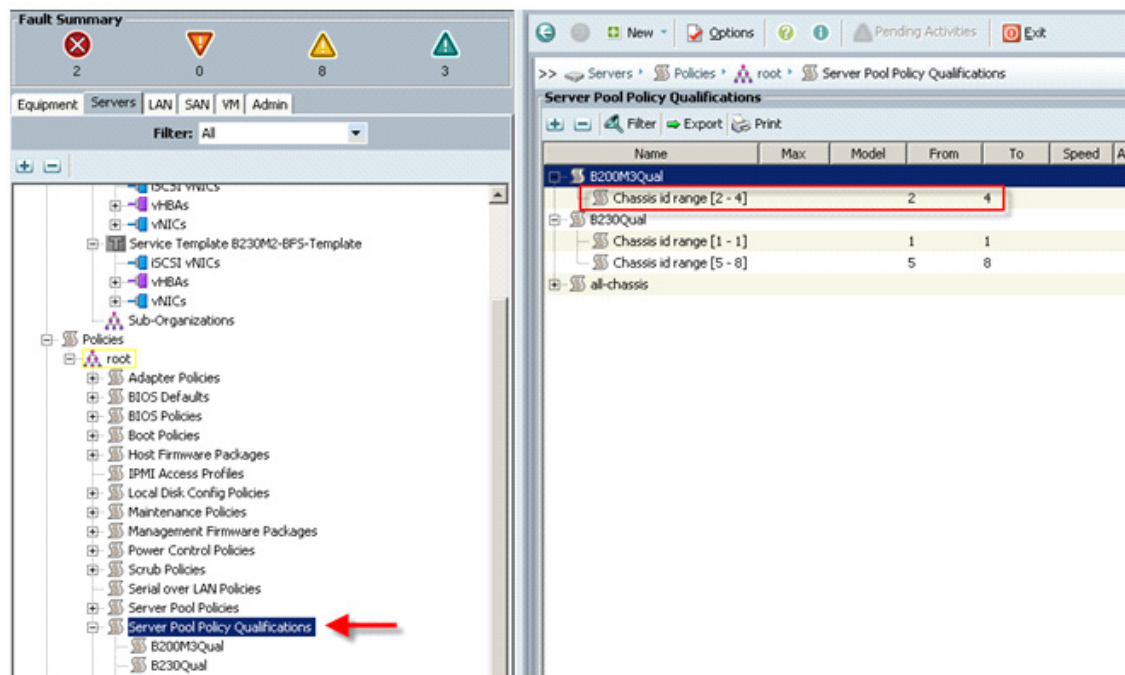
- d. The next pool we created is the Server UUID pool. On the Servers tab in the Navigator page under the Pools node we created a single UUID Pool for the test environment. Each Cisco UCS Blade Server requires a unique UUID to be assigned by its Service profile.



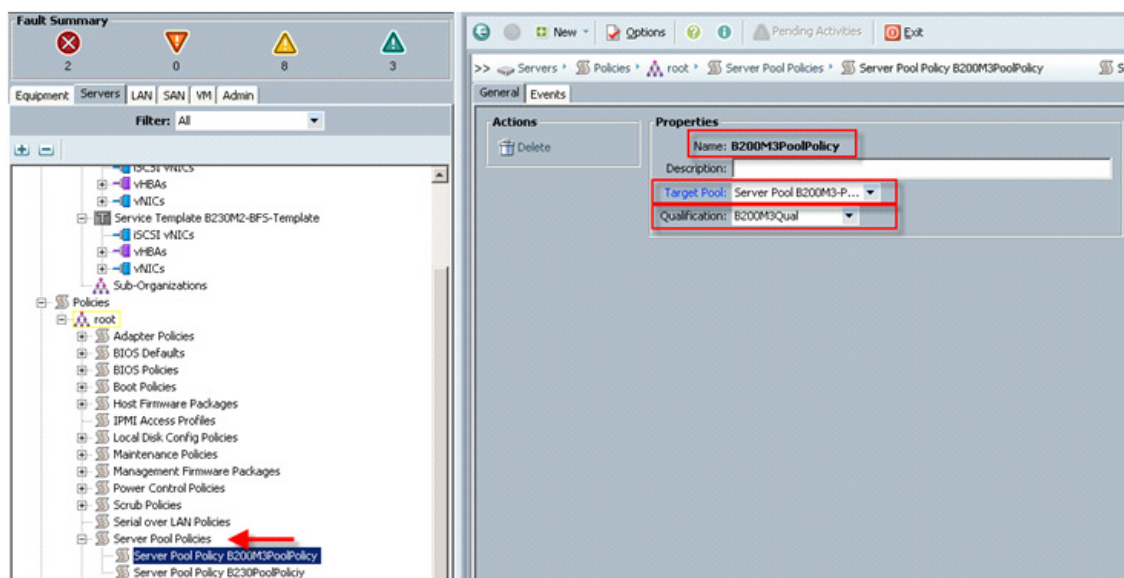
- e. We created two Server Pools for use in our Service Profile Templates as selection criteria for automated profile association. Server Pools were created on the Servers tab in the navigation page under the Pools node. Only the pool name was created, no servers were added:



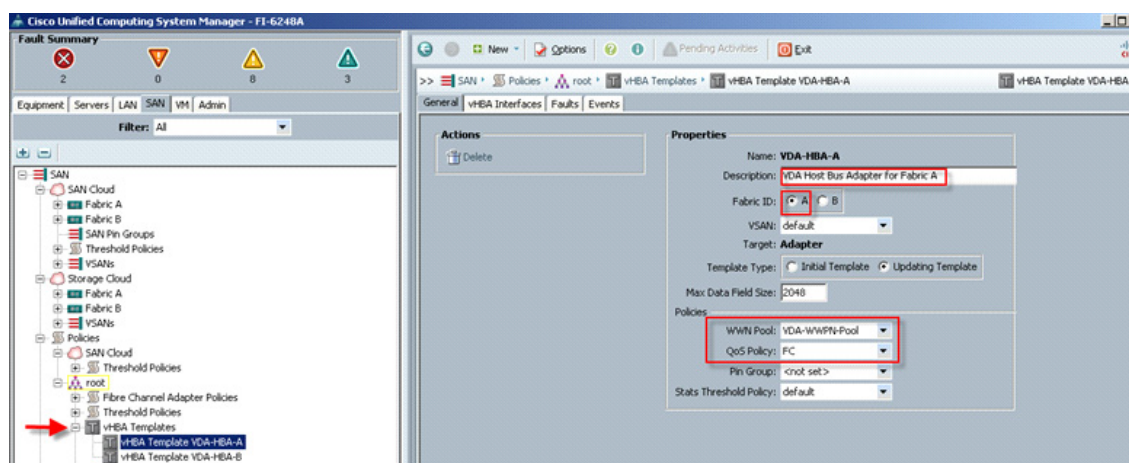
- f. We created two Server Pool Policy Qualifications to identify the blade server model for placement into the correct pool using the Service Profile Template. In this case we used Chassis ids to select the servers. (We could have used slots or server models to make the selection.)



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

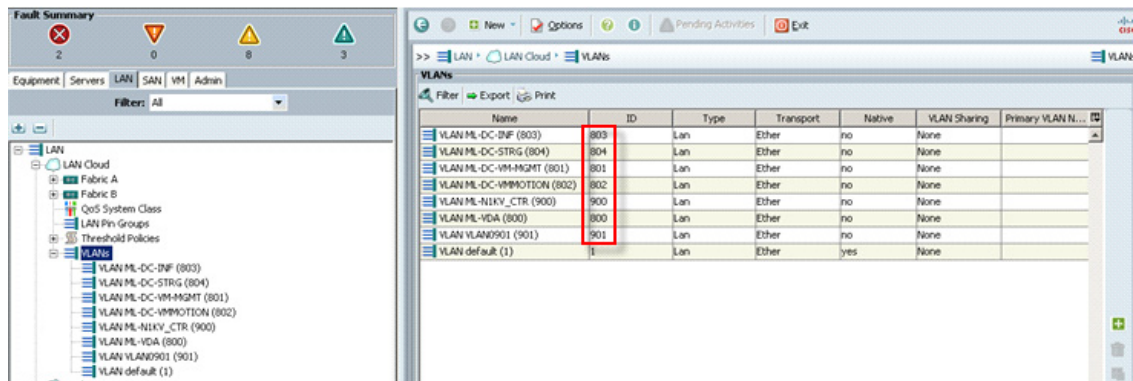


9. Virtual Host Bus Adapter templates were created for FC SAN connectivity from the SAN tab under the Policies node, one template for each fabric.



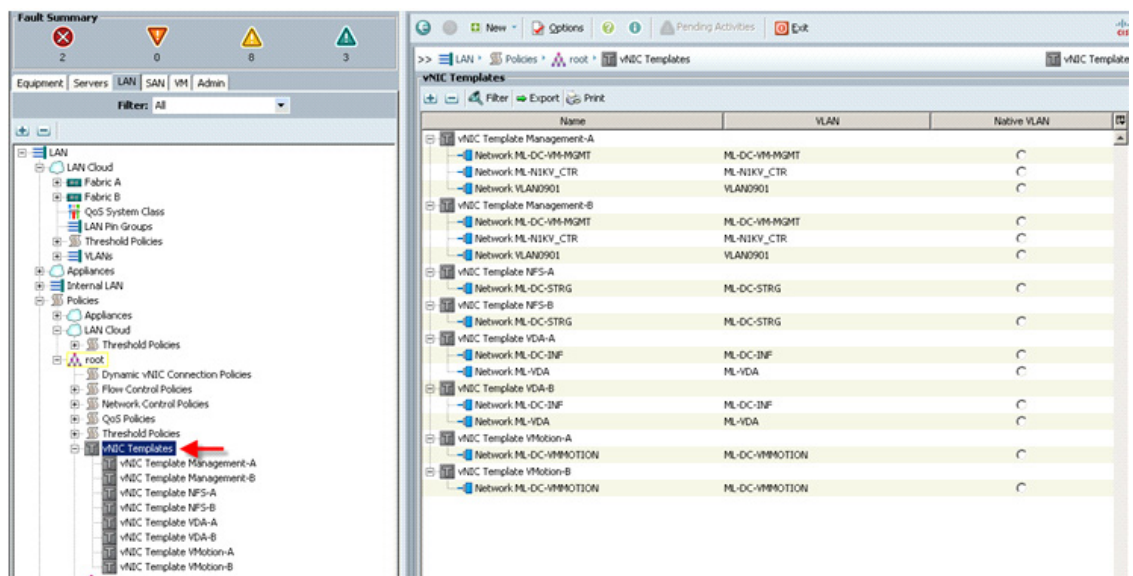
Create at least one HBA template for each Fabric Interconnect if block storage will be used. We used the WWPN pool created earlier and the QoS Policy created in section 5.2.4.

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



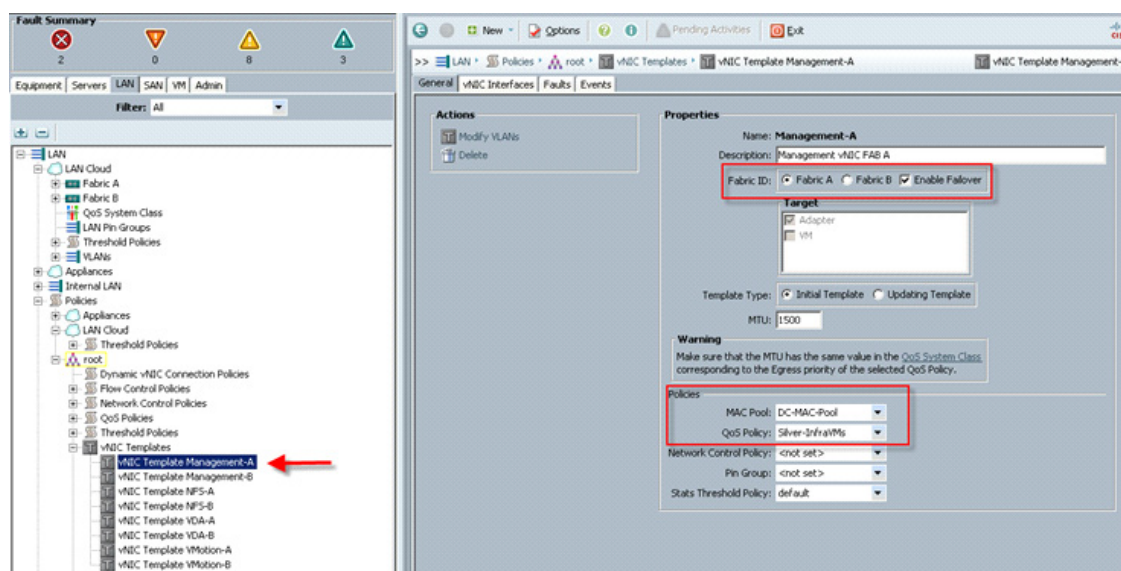
In this project we utilized seven VLANs to accommodate our four ethernet system classes, a separate VLAN for infrastructure services, and two VLANs for Nexus 1000V packet and control functions. (NIKV management and VMware Management shared VLAN 801).

11. 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 eight virtual NICs per host, four pairs, with each pair connected to both Fabric Interconnects for resiliency.

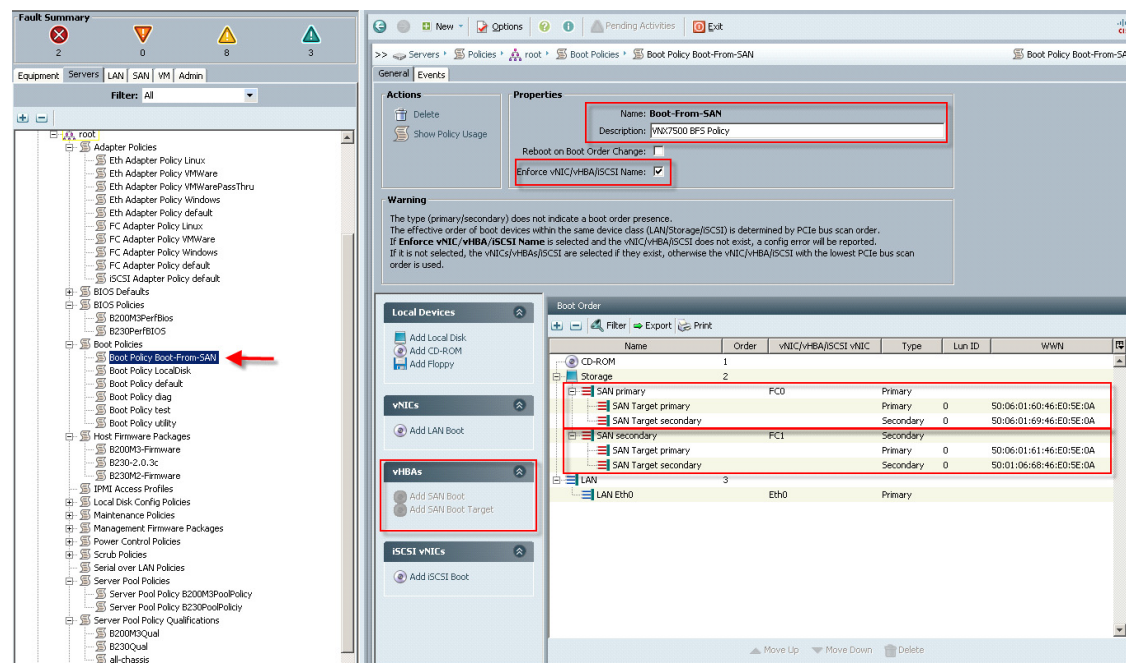


- a. Create vNIC templates for both fabrics, check Enable Failover, select VLANs supported on adapter (optional,) set the MTU size, select the MAC Pool and QoS Policy, then click OK.

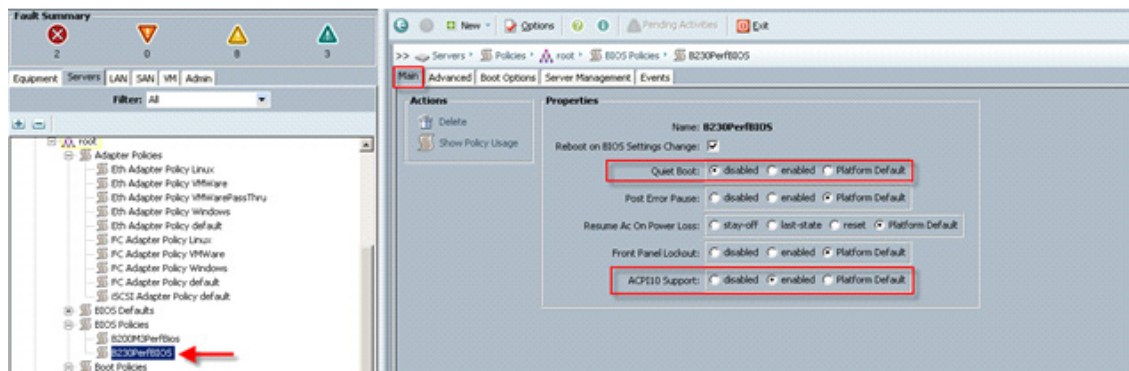




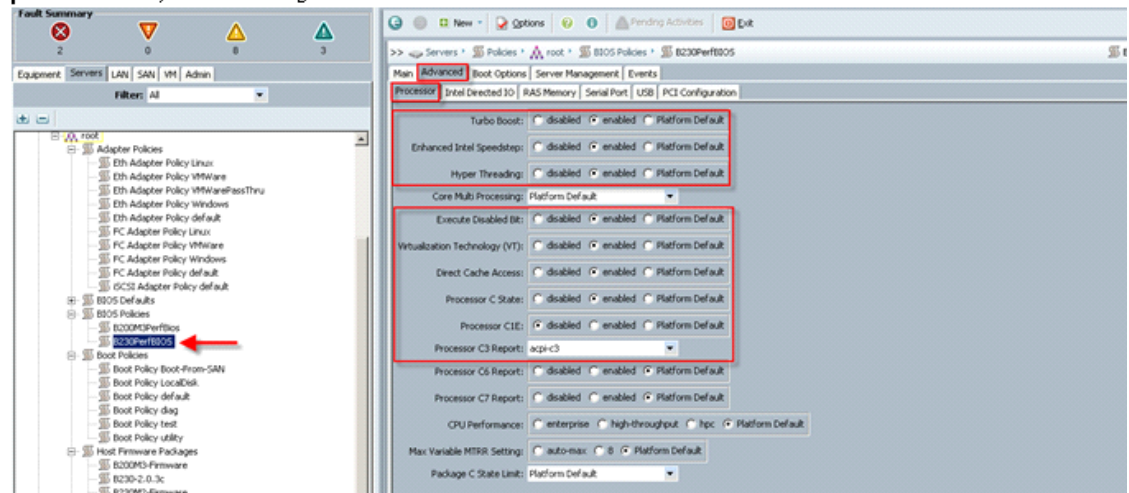
12. Create boot from SAN policy that was used for both B230 M2 and B200 M3 blades, using the WWNs from the VNX7500 storage system as SAN targets.



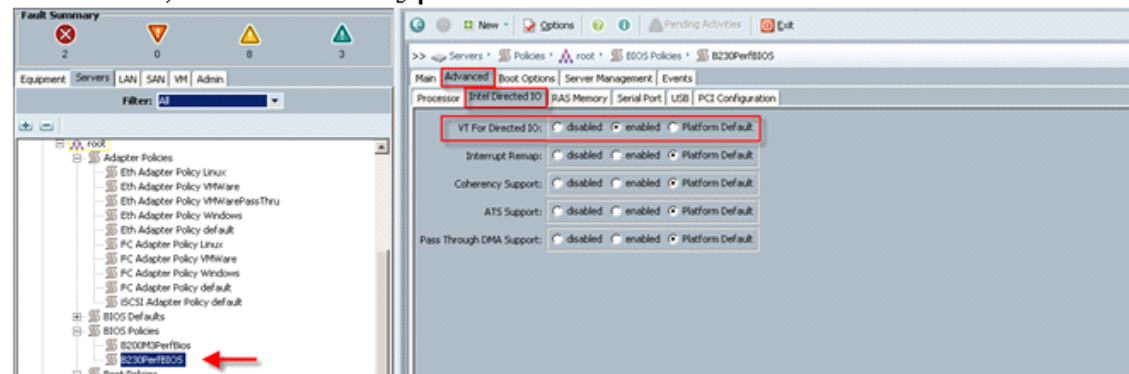
13. Create performance BIOS Policies for each blade type to insure optimal performance. The following screen captures show the settings for the Cisco UCS B230 M2 blades used in this study.



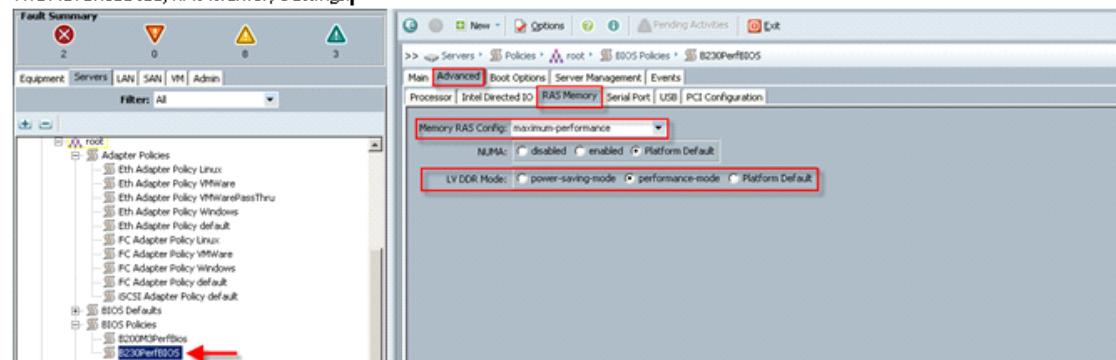
The Advanced tab, Processor settings:



The Advanced tab, Intel Directed I/O tab settings:



The Advanced tab, RAS Memory settings:



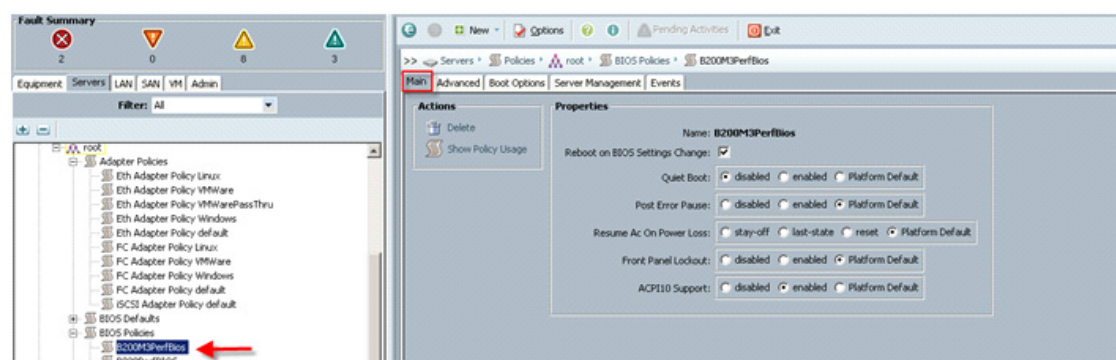
The remaining Advanced tab settings are at platform default or not configured. Similarly, the Boot Options and Server Management tabs' settings are at defaults.



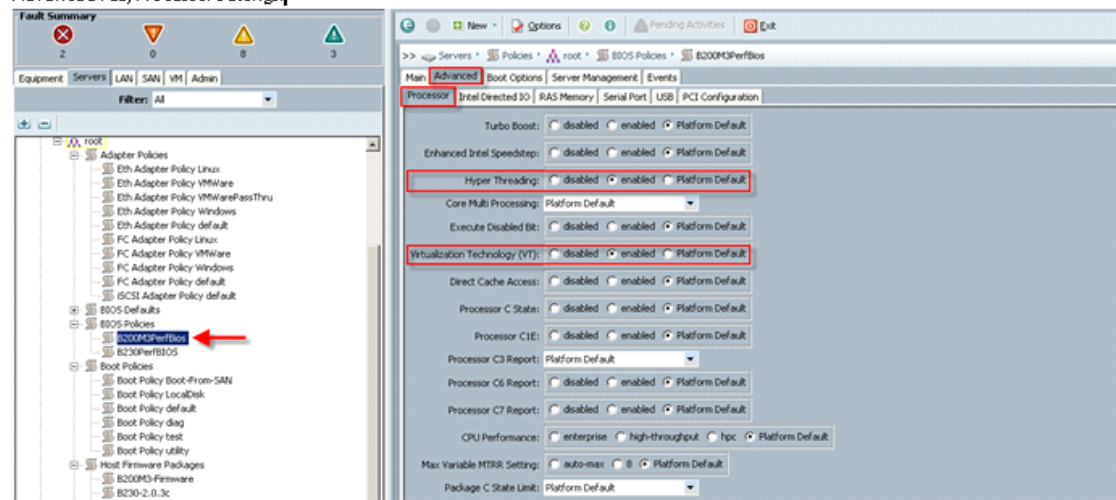
#### Note

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.

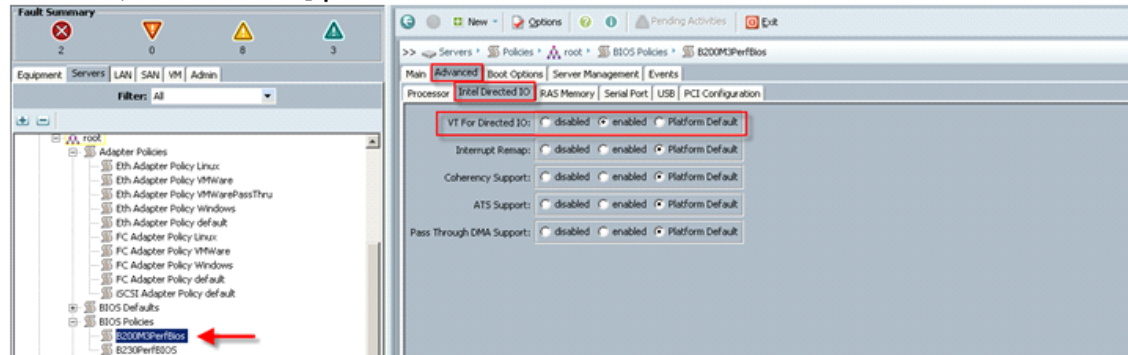
- a. The following screen captures show the settings for the B200 M3 blades used in this study:



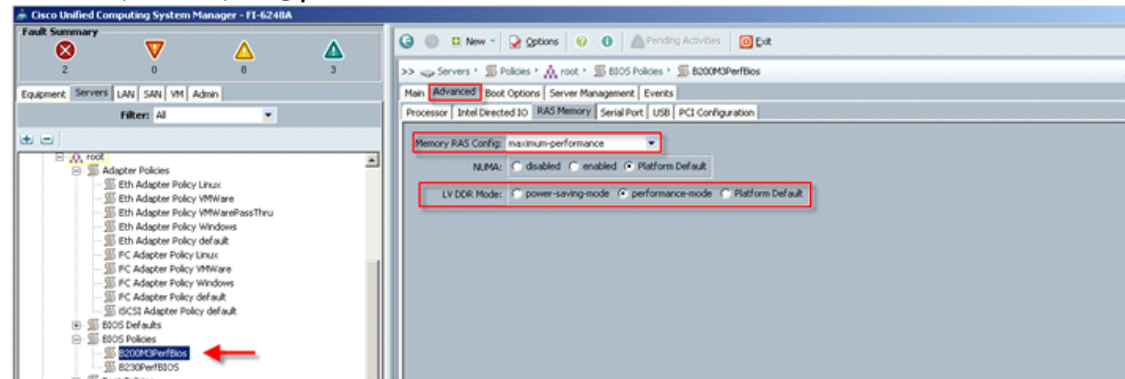
Advanced Tab, Processor settings:



AdvancedTab, Intel Directed IO settings:



AdvancedTab, RAS Memory Settings:



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

14. Cisco UCS B200 M3 and Cisco UCS B230 M2 Host Firmware Package policies were set for Adapter and BIOS.



## B200 M3 Adapter Firmware

The screenshot displays the Cisco Unified Computing System Manager interface. On the left, the 'Host Firmware Packages' folder is selected under the 'B200M3-Firmware' package. The right pane shows the 'Properties' tab for 'B200M3-Firmware'. Below the properties, a table lists the installed firmware packages:

Select	Vendor	Model	PID	Presence	Version
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M51KR-B	N20-AB0002	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS M61KR-E	N20-AC0002	Present	2.0(4a)
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M71KR-E	N20-AE0002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M72KR-E	N20-AE0102	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M61KR-I	N20-AI0102	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M71KR-Q	N20-AQ0002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M72KR-Q	N20-AQ0102	N/A	<not set>
<input type="checkbox"/>	Broadcom Corp.	Broadcom 10GbE Onboard ...	N20-ABPC101	N/A	<not set>
<input type="checkbox"/>	Broadcom Corp.	Broadcom 10GbE Onboard ...	N20-ABPC102	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS P81E	N20-ACPC101	N/A	<not set>
<input type="checkbox"/>	Emulex Corp.	Emulex 0C610102-F	N20-ABPC101	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Intel 10GbE Adapter	N20-ABPC101	N/A	<not set>
<input type="checkbox"/>	QLogic Corp.	QLogic QLE8152	N20-AQPC101	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS VIC 1280	UCS-VIC-M82-8P	Present	2.0(4a)
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M61KR-B	UCSB-MEZ-BRC-02	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M71KR-E	UCSB-MEZ-ELX-03	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M72KR-Q	UCSB-MEZ-QLG-03	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS VIC 1240	UCSB-PLOM-40G-01	Present	2.0(4a)

## B200 M3 BIOS Firmware

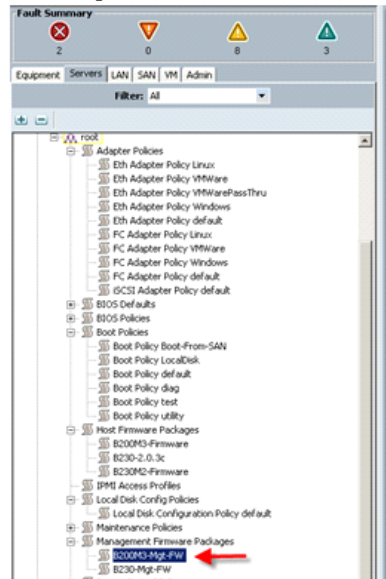
The screenshot displays the Cisco Unified Computing System Manager interface. On the left, the 'Host Firmware Packages' folder is selected under the 'B200M3-Firmware' package. The right pane shows the 'Properties' tab for 'B200M3-Firmware'. Below the properties, a table lists the installed BIOS firmware packages:

Select	Vendor	Model	PID	Presence	Version
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B230 M2	B230-BASE-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B440 M2	B440-BASE-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B200 M1	N20-B6620-1	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Cisco UCS B200 M1	N20-B6620-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B250 M1	N20-B6620-2	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Cisco UCS B250 M1	N20-B6620-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B200 M2	N20-B6625-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B250 M2	N20-B6625-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B230 M1	N20-B6730-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B440 M1	N20-B6740-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C200 M1	R200-1120402	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C200 M2	R200-1120402W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C210 M1	R210-2121605	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C210 M2	R210-2121605W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C250 M1	R250-2460805	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C250 M2	R250-2460805W	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B200 M3	UCSB-B200-M3	Present	B200M3 2.0.4a.0.080...
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B22 M3	UCSB-B22-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B420 M3	UCSB-B420-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS C200 M2	UCSC-BSE-SFP-C200	N/A	<not set>

The process was repeated for the Cisco UCS B230 M2 Blade Servers, choosing the appropriate model and versions.

15. Create Management Firmware Packages for each Blade Type in the environment, assigning the latest Cisco UCS Infrastructure and Management software downloaded earlier.

## B200 M3-Mgt-FW



General Events

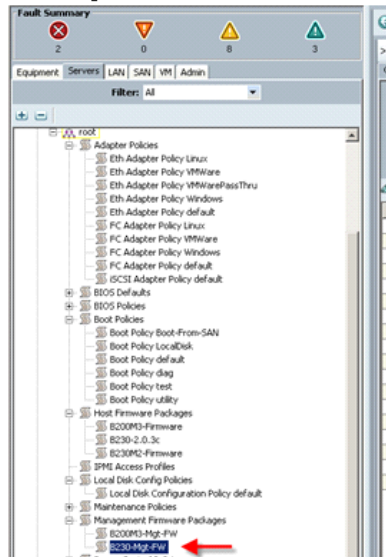
Actions: Delete, Show Policy Usage

Properties: Name: B200M3-Mgt-FW, Description:

Filter Export Print

Select	Vendor	Model	PID	Presence	Version
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B230 M2	B230-BASE-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B440 M2	B440-BASE-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M1	N20-B6620-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B250 M1	N20-B6620-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M2	N20-B6625-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B250 M2	N20-B6625-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B230 M1	N20-B6730-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B440 M1	N20-B6740-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M1	R200-1120402	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M2	R200-1120402W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C210 M1	R210-2121605	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C210 M2	R210-2121605W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C250 M1	R250-2480805	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C250 M2	R250-2480805W	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M3	UC3B-B200-M3	Present	2.0(4a)
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B22 M3	UC3B-B22-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B420 M3	UC3B-B420-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M2	UC3C-B5E-9FF-C200	N/A	<not set>

## B230 M2-Mgt-FW



General Events

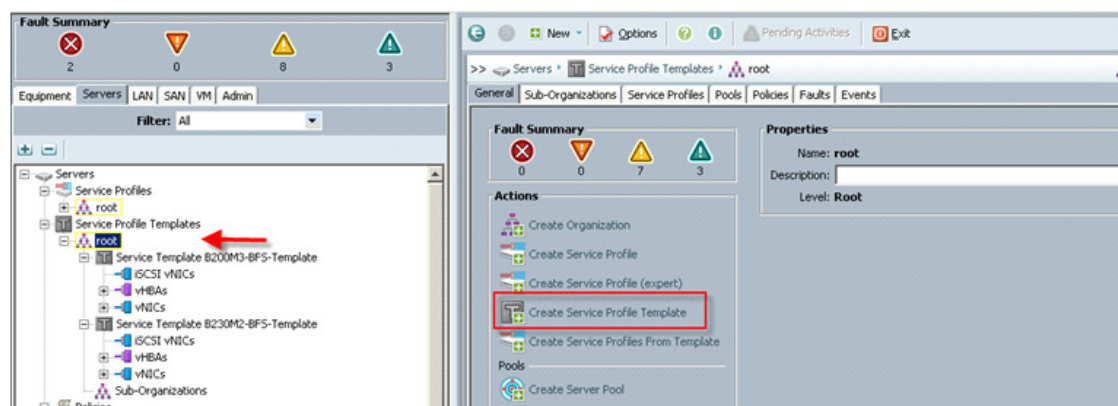
Actions: Delete, Show Policy Usage

Properties: Name: B230-Mgt-FW, Description:

Filter Export Print

Select	Vendor	Model	PID	Presence	Version
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS B230 M2	B230-BASE-M2	Present	2.0(4a)
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B440 M2	B440-BASE-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M1	N20-B6620-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B250 M1	N20-B6620-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M2	N20-B6625-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B250 M2	N20-B6625-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B230 M1	N20-B6730-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B440 M1	N20-B6740-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M1	R200-1120402	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M2	R200-1120402W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C210 M1	R210-2121605	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C210 M2	R210-2121605W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C250 M1	R250-2480805	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C250 M2	R250-2480805W	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B200 M3	UC3B-B200-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B22 M3	UC3B-B22-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS B420 M3	UC3B-B420-M3	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS C200 M2	UC3C-B5E-9FF-C200	N/A	<not set>

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



In this project, we created two templates, one for each of the Cisco UCS Blade Server models used. Follow through each section, utilizing the policies and objects you created earlier, then click Finish.



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



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

- a. On the Create Service Profile Template wizard, we entered a unique name, selected the type as updating, and selected the VDA-UUID-Suffix\_Pool created earlier, then clicked Next.

## Unified Computing System Manager

Create Service Profile Template

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

### Identify Service Profile Template

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

Name: B200M3-BFS-Template ←

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: YDA-UUID-Suffix-Pool(37/100) Select the UUID Pool created earlier from the drop-down.

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.

< Prev
Next >
Finish
Cancel

- b. 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.



**Unified Computing System Manager**

Create Service Profile Template

1. [Identify Service Profile Template](#)
2. [Storage](#)
3. [Networking](#)
4. [vNIC/vHBA Placement](#)
5. [Server Boot Order](#)
6. [Maintenance Policy](#)
7. [Server Assignment](#)
8. [Operational Policies](#)

**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

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: VDA-WWNN-Pool(137/200) Select the VDA-WWNN-Pool created earlier from the drop down.

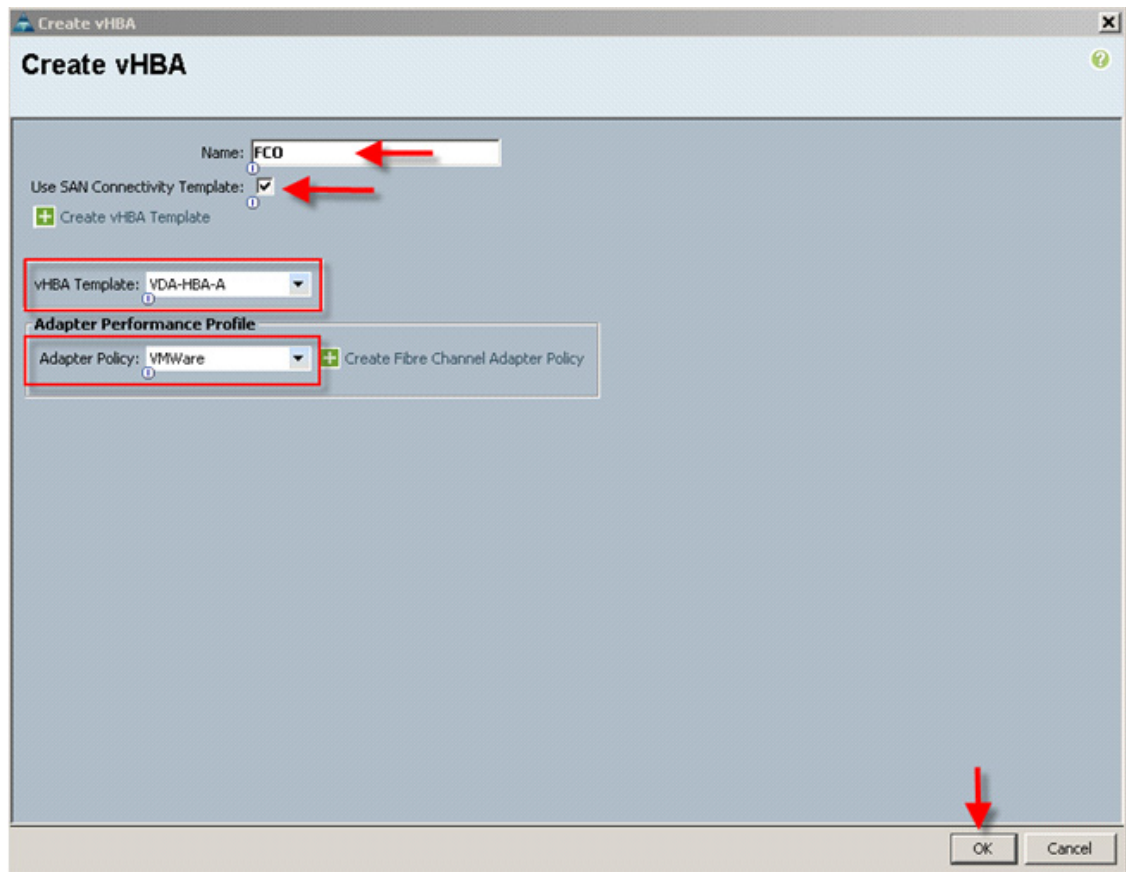
The WWNN will be assigned from the selected pool.  
The available/total WWNNs are displayed after the pool name.

Name	WWPN

[Delete](#) [Add](#) [Modify](#)

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

- On the Create HBA page, we entered a name (FCO) and checked Use SAN Connectivity Template, which changed the display to the following:



17. Select the vHBA template for Fabric Interconnect A and the VMware Adapter Policy from the drop down lists, then click OK.
18. Repeat the process for FC1, choosing VDA-HBA-B for Fabric Interconnect B. The result is the Storage page that appears as follows:

**Unified Computing System Manager**

Create Service Profile Template

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

**Storage**

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage:  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

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:

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

Name	WWNN
vHBA FCO	Derived
vHBA IF	Derived
vHBA FC1	Derived
vHBA IF	Derived

Add

< Prev **Next >** Finish Cancel

19. Click Next to continue.

20. We selected the Expert configuration option and clicked Add in the adapters window:

**Unified Computing System Manager**

Create Service Profile Template

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

**Networking**

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:   Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity? ☐ Simple ☒ **Expert** ☐ No vNICs

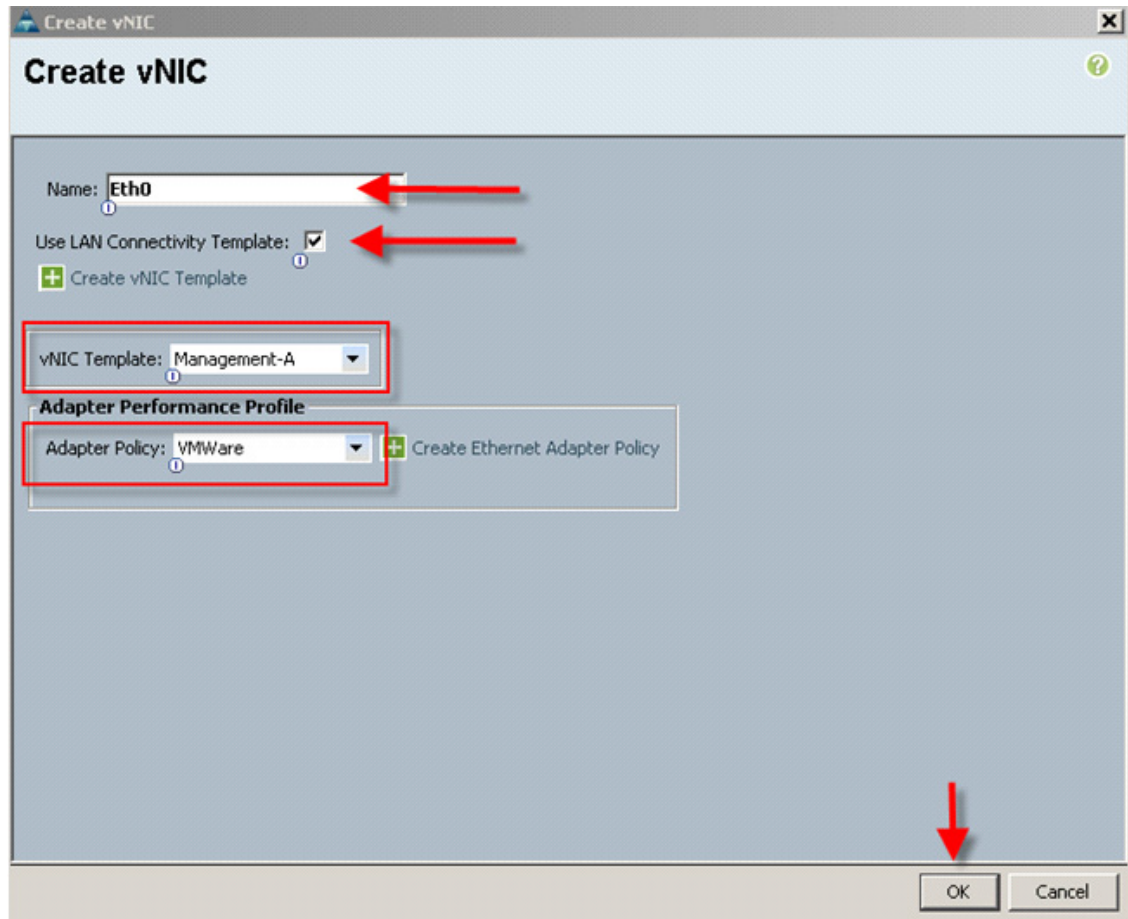
Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN
------	-------------	-----------	-------------

Add

ISCSI vNICs

21. In the Create vNIC window, we entered a unique Name, checked the Use LAN Connectivity Template checkbox, selected the vNIC Template from the drop down, and the Adapter Policy the same way.



22. We repeated the process for the remaining seven vNICs, resulting in the following: (Eth5, 6, and 7 not shown)

**Unified Computing System Manager**

Create Service Profile Template

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

### Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by defa... [+ Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity? ☒ Simple ☐ Expert ☐ No vNICs

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

Name	MAC Address	Fabric ID	Native VLAN
vNIC Eth0	Derived	derived	derived
vNIC Eth1	Derived	derived	derived
vNIC Eth2	Derived	derived	derived
vNIC Eth3	Derived	derived	derived
vNIC Eth4	Derived	derived	derived

[Delete](#) [+ Add](#) [Modify](#)

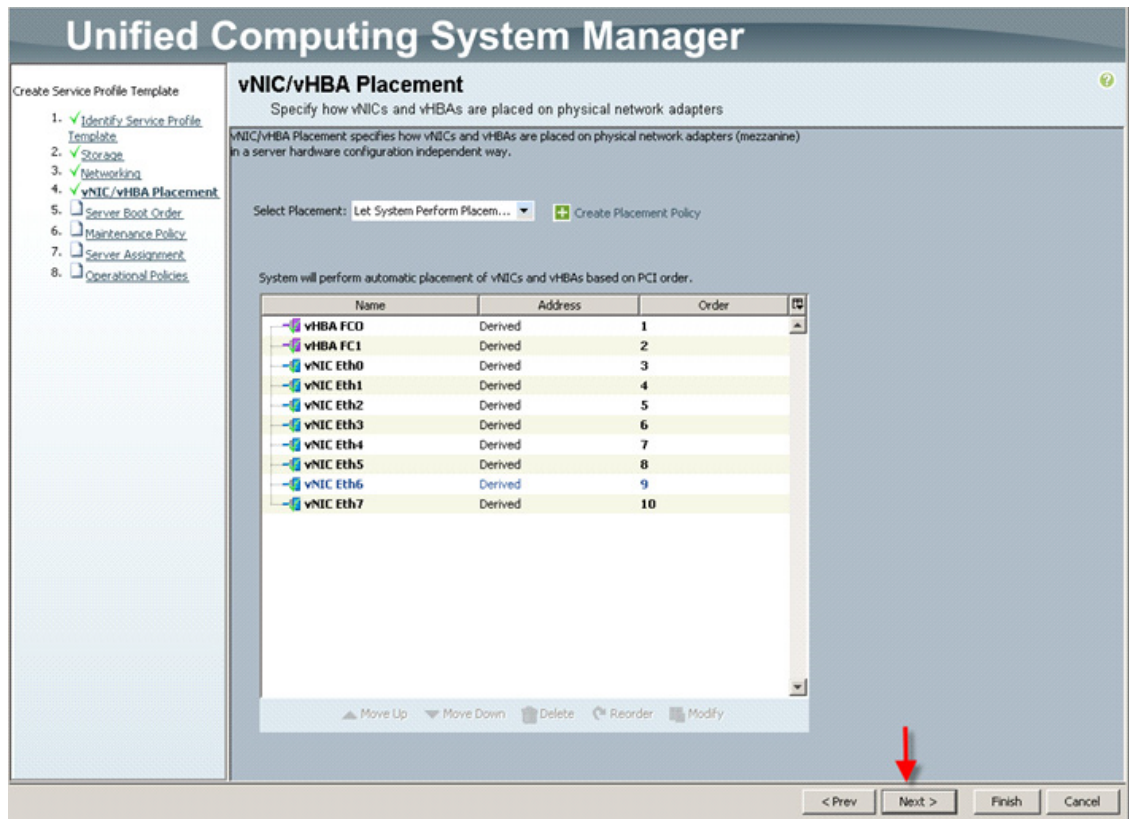
iscsi vNICs

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

23. Click Next to continue.

24. We accepted the default placement and clicked Next.





25. We selected the Boot from SAN policy created in Section 6.4.5 from the drop down, then proceeded.

## Unified Computing System Manager

Create Service Profile Template

1. [Identify Service Profile Template](#)
2. [Storage](#)
3. [Networking](#)
4. [vNIC/vHBA Placement](#)
5. **[Server Boot Order](#)**
6. [Maintenance Policy](#)
7. [Server Assignment](#)
8. [Operational Policies](#)

### Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Boot-From-SAN + Create Boot Policy

Name: **Boot-From-SAN**  
 Description: **VNX7500 BFS Policy**

Reboot on Boot Order Change: **no**  
 Enforce vNIC/vHBA/SCSI Name: **yes**

**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/SCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/SCSI Name** is selected and the vNIC/vHBA/SCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/SCSI are selected if they exist, otherwise the vNIC/vHBA/SCSI with the lowest PCIe bus scan order is used.

Boot Order

+ Add Filter Export Print

Name	Order	vNIC/vHBA/SCSI 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		

Create iSCSI vNIC Set iSCSI Boot Parameters

< Prev Next > Finish Cancel

**Note**

We did not create a Maintenance Policy for the project, so we clicked Next to continue.

The screenshot shows the 'Unified Computing System Manager' interface. On the left, a sidebar titled 'Create Service Profile Template' lists eight steps: 1. Identify Service Profile Template (checked), 2. Storage (checked), 3. Networking (checked), 4. vNIC/vHBA Placement (checked), 5. Server Boot Order (checked), 6. Maintenance Policy (checked), 7. Server Assignment (unchecked), and 8. Operational Policies (unchecked). The main panel is titled 'Maintenance Policy' and contains a sub-panel with the same title. The sub-panel text reads: 'Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.' Below this, there is a dropdown menu labeled 'Maintenance Policy:' with the text 'Select (no policy used by default)' and a '+ Create Maintenance Policy' button. A note below the dropdown states: 'No maintenance policy is selected by default. The service profile will immediately reboot when disruptive changes are applied.' At the bottom right of the main panel, a red arrow points to the 'Next >' button in a navigation bar that also includes '< Prev', 'Finish', and 'Cancel' buttons.

26. We made the following selections from the drop down lists as shown, then clicked Next to continue:



**Unified Computing System Manager**

**Create Service Profile Template**

1. ☒ Identify Service Profile Template
2. ☒ Storage
3. ☒ Networking
4. ☒ vNIC/vHBA Placement
5. ☒ Server Boot Order
6. ☒ Maintenance Policy
7. ☒ **Server Assignment**
8. ☐ 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: B200M3-Pool + 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: B200M3Qual

Restrict Migration: ☐

**Firmware Management (BIOS, Disk Controller, Adapter)**

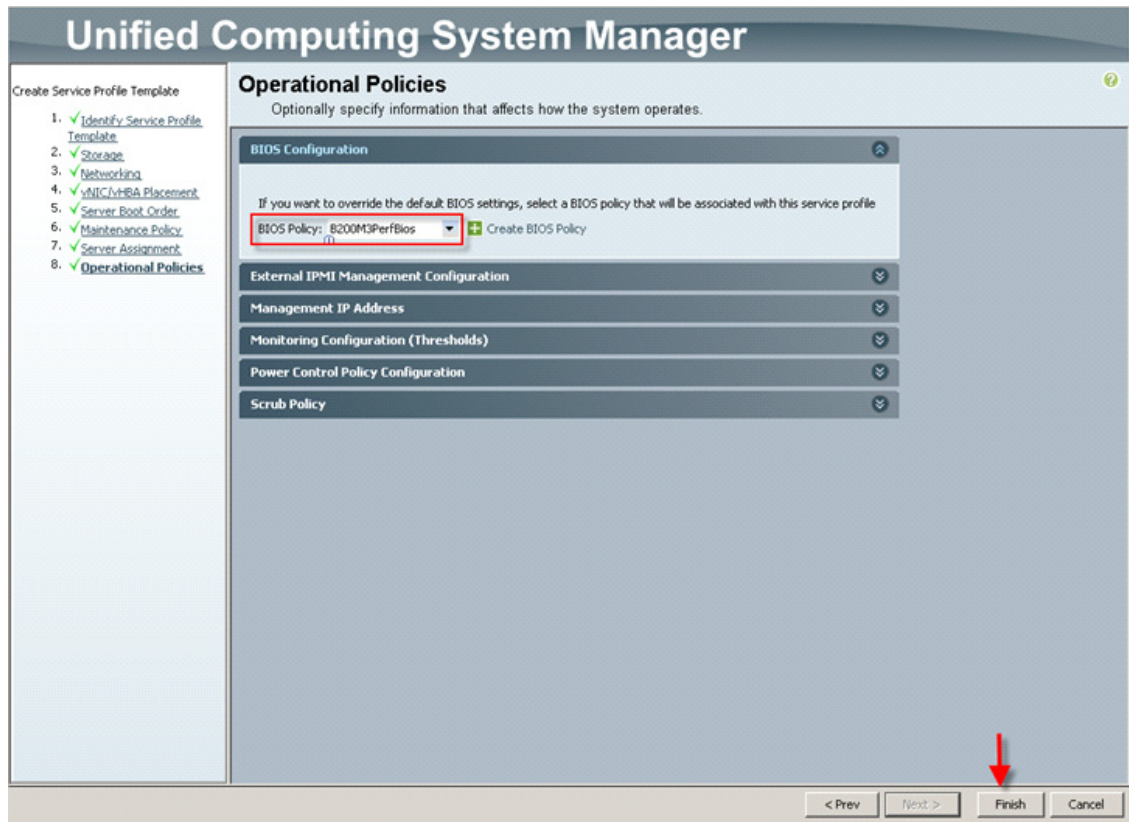
If you select a host or management firmware policy for this service profile template, the profile will update the firmware on the server that is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware: B200M3-Firmware + Create Host Firmware Package

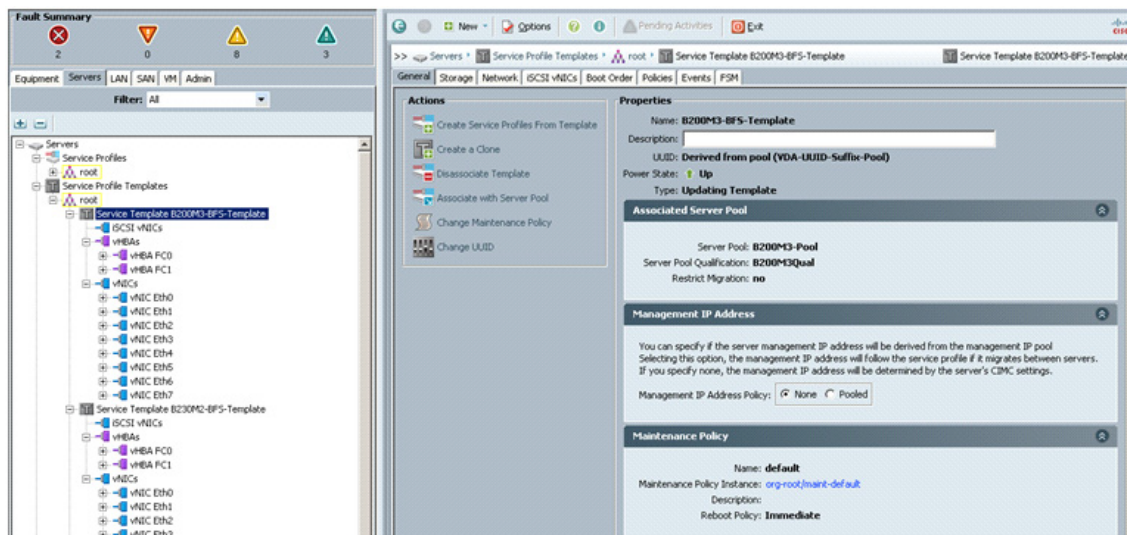
Management Firmware: B200M3-Mgt-FW + Create Management Firmware Package

< Prev **Next >** Finish Cancel

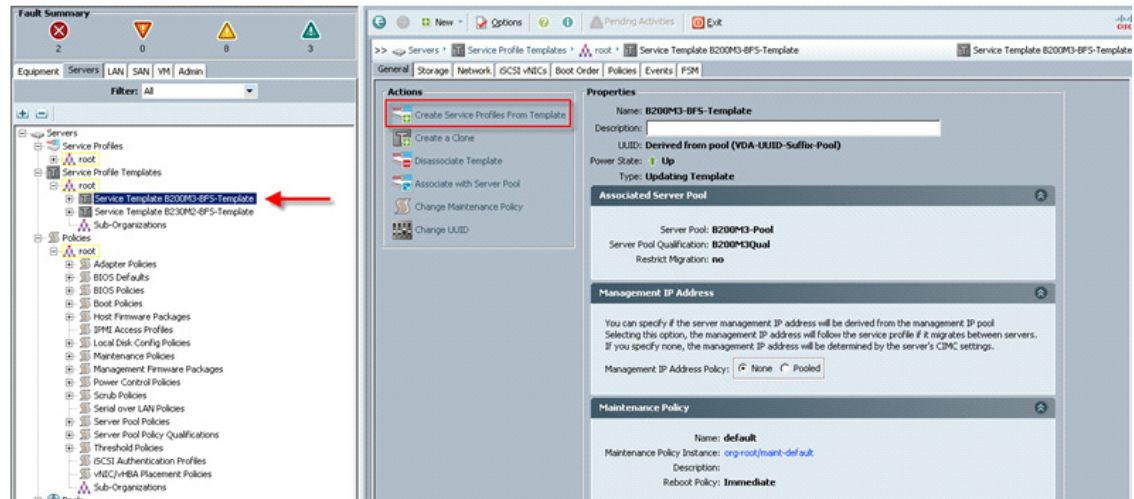
27. On the Operational Policies page, we expanded the BIOS Configuration section and selected the BIOS Policy for the Cisco UCS B200 M3 blade servers created earlier, then clicked Finish to complete the Service Profile Template.



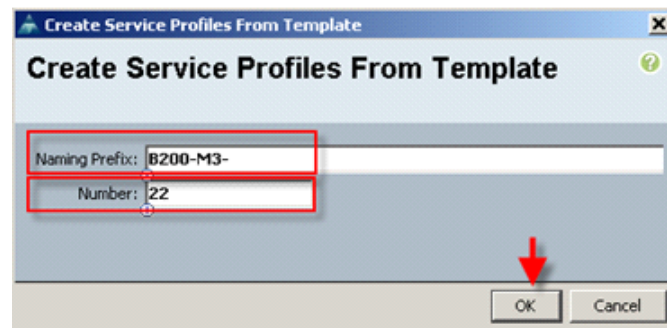
The result is a Service Profile Template for the Cisco UCS Blade Server B200 M. We repeated the procedure to create a Service Profile Template for the Cisco UCS Blade Server B230 M2 used in the study.



28. Now that we had created the Service Profile Templates for each Cisco UCS Blade Server model used in the project, we used them to create the appropriate number of Service Profiles. To do so, in the Servers tab in the navigation page, in the Service Profile Templates node, we expanded the root and selected Service Template B200M3, then clicked on Create Service Profiles from Template in the right pane, Actions area.



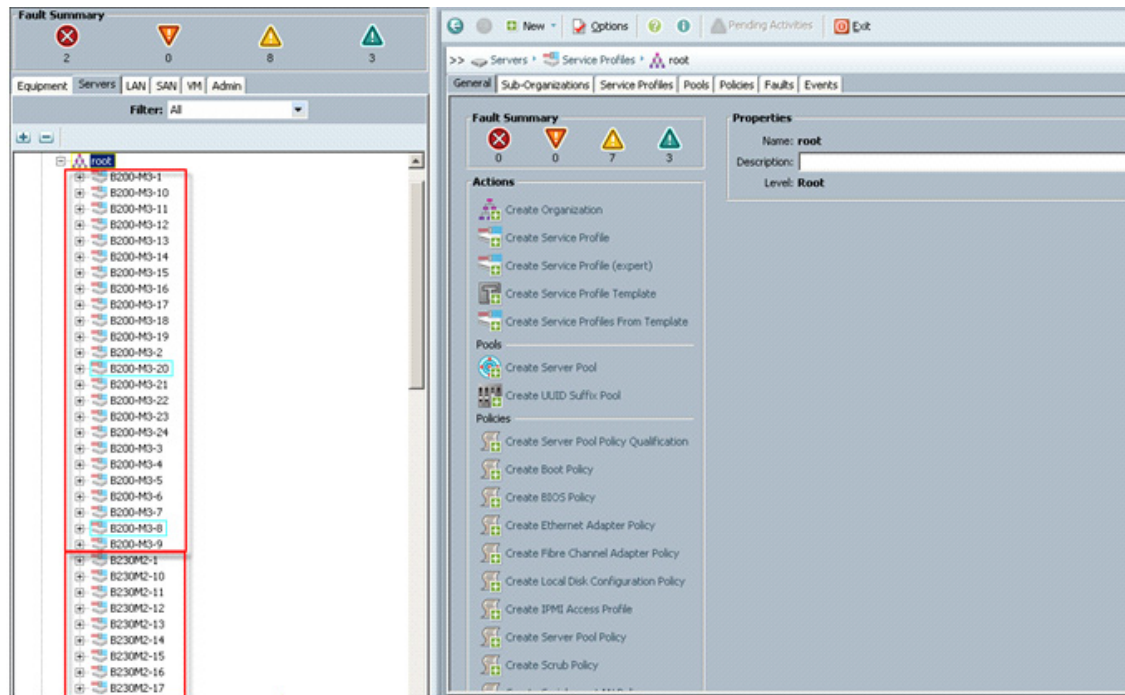
29. We provided the naming prefix and the number of Service Profiles to create and clicked OK.



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

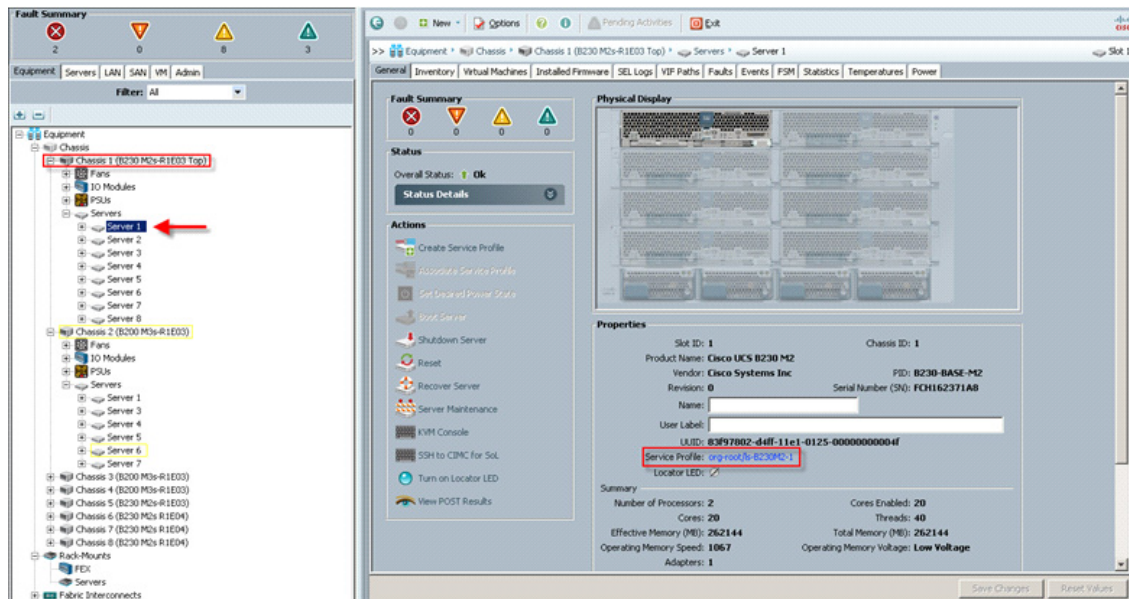
The process was repeated for the Cisco UCS B230 M2-BFS-Template and the same result was achieved.





31. We verified that each server had a profile and that it received the correct profile.

B230 M2 Sample:



At this point, the Cisco UCS Blade Servers are ready for hypervisor installation.

## 6.2.2 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.

### 6.2.3 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.
- Through Cisco Unified Computing System 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).

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

### 6.2.4 Cisco UCS System Class Configuration

Cisco Unified Computing System 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

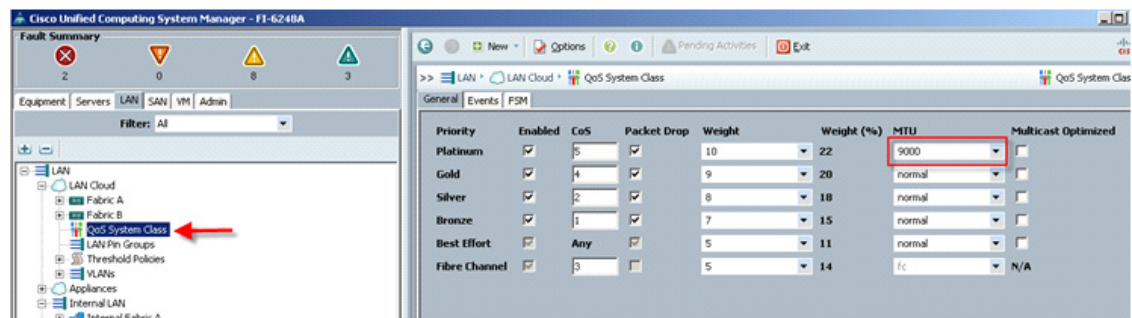
## 6.2.5 Steps to Enable QOS on the Cisco Unified Computing System

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

**Table 6** QoS Priority to vNIC and VLAN Mapping

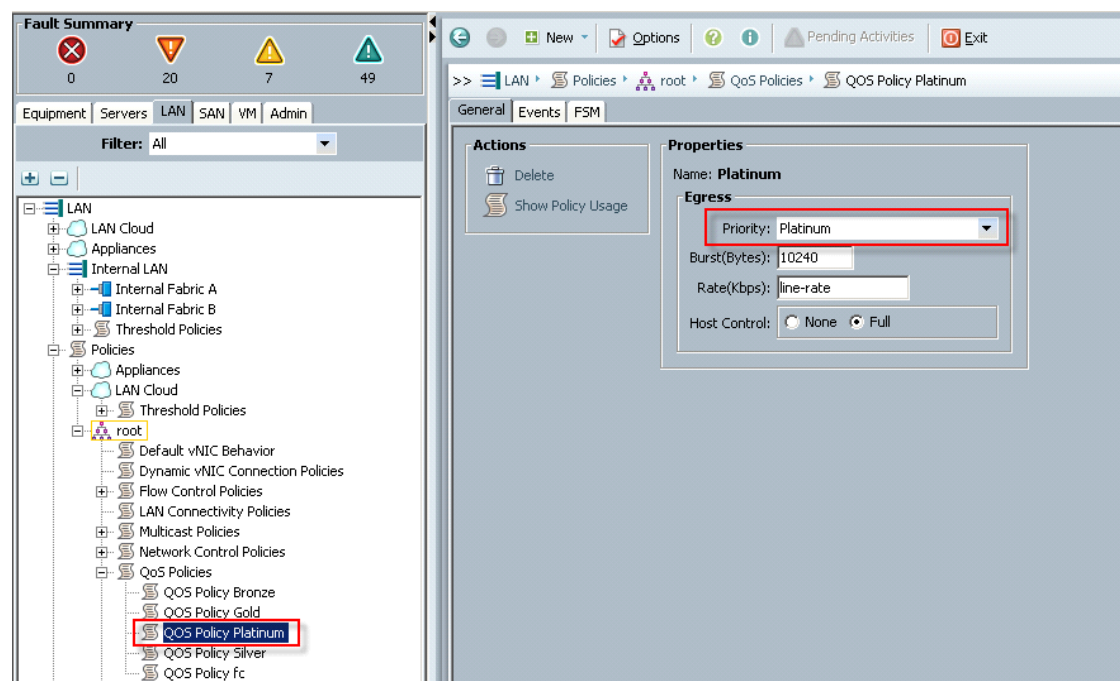
Cisco UCS Qos Priority	vNIC Assignment	VLAN Supported
Platinum	eth2, eth3	804 (Storage)
Gold	eth4, eth5	800 (VDA)
Silver	eth0, eth1	801 (Management)
Bronze	eth6, eth7	802 (vMotion)

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

**Figure 11** Cisco UCS QoS System Class Configuration

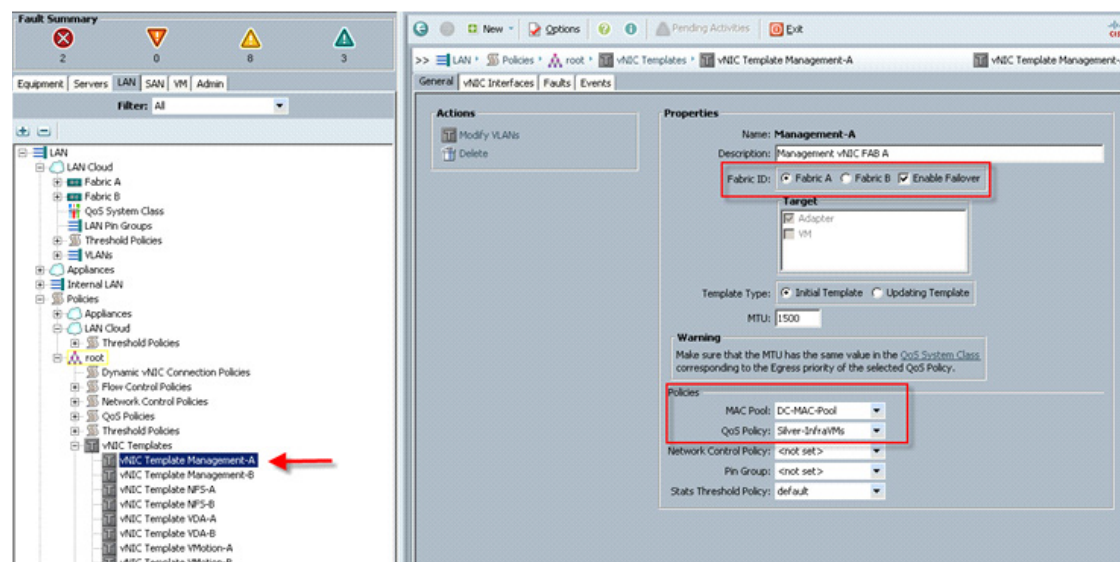
Next, in the LAN tab under Policies, Root, QoS Policies, verify QoS Policies Platinum, Gold, Silver and Bronze exist, with each QoS policy mapped to its corresponding Priority.

Figure 12 Cisco UCS QoS Policy Configuration



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

Figure 13 Utilize QoS Policy in vNIC Template



This is a unique value proposition for Cisco UCS with respect to end-to-end QoS. For example, we have a VLAN for the EMC storage, configure Platinum policy with Jumbo frames and get an end-to-end QoS and performance guarantees from the Blade Servers to the Nexus 1000V virtual distributed switches running in vCenter through the Nexus 5548UP access layer switches.




## 6.3 LAN Configuration

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

### 6.3.1 Cisco UCS Connectivity

Figure 14 Unified Computing System 6200 Series Fabric Interconnects




### 6.3.2 EMC VNX7500 LAN Connectivity

The Cisco Nexus 5548 is used to connect to the EMC VNX7500 storage system for Fiber Channel and file-based access.

The VNX7500 is equipped with dual-port 8GB FC modules on each controller. These are connected to the pair of Nexus 5548s to provide block storage access to the environment. (See Section 6.4 SAN configuration below.)

The VNX7500 supports four dual-port 10G Data Movers which are connected to the pair of N5Ks downstream. Three of the Data Movers were set to Active, with the fourth providing failover capability. This allows end-to-end 10G access for file-based storage traffic. We have implemented jumbo frames on the ports and have priority flow control on, with Platinum CoS and QoS assigned to the vNICs carrying the storage data access on the Fabric Interconnects.


The EMC ethernet connectivity diagram is shown below. Here again, we have a total of 40G bandwidth available for the servers.

**Figure 15 EMC VNX Ethernet Connectivity**

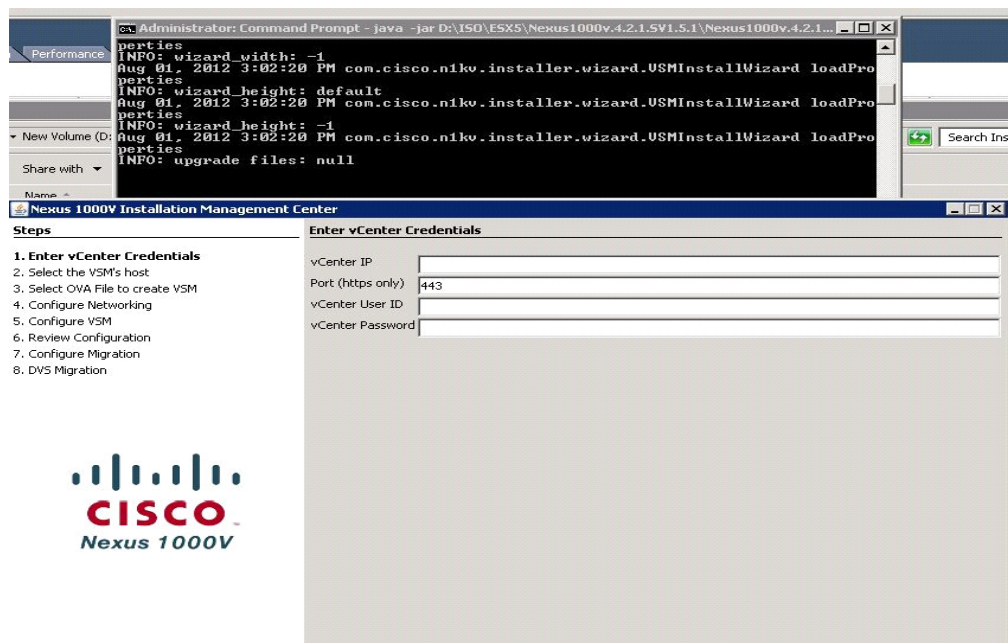
For information on configuring ethernet connectivity on a EMC VNX7500 Storage System, refer to the EMC website.

### 6.3.3 Nexus 1000V Configuration in L3 Mode

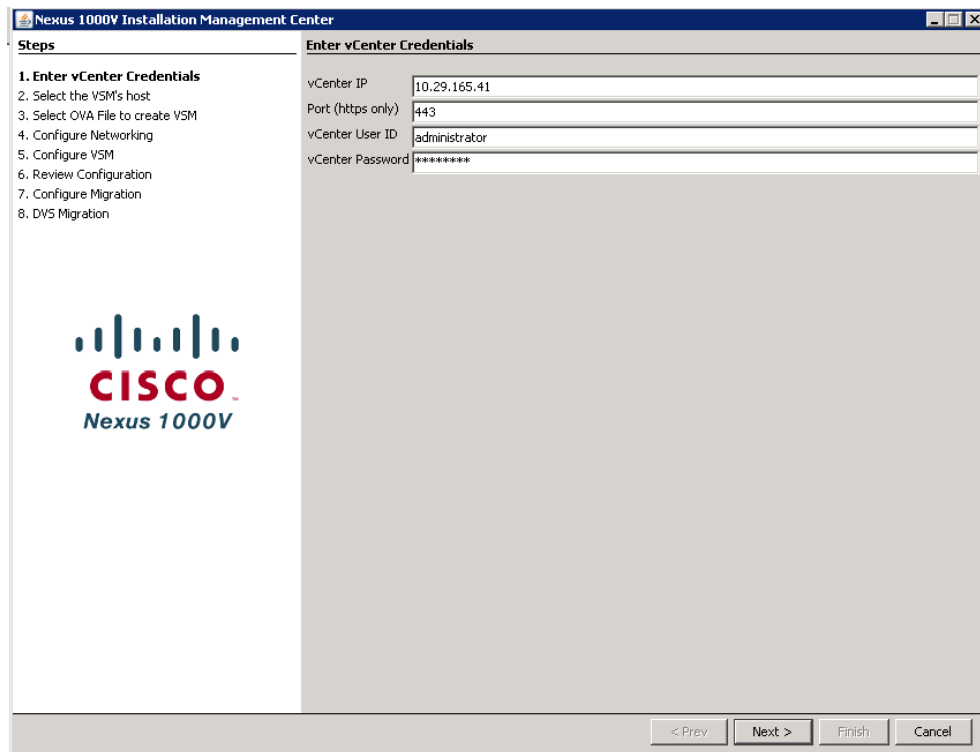
1. To download the Nexus1000 V 4.2(1) SV1 (5.2), Click the link below.  
[http://www.cisco.com/cisco/software/release.html?mdfid=282646785&flowid=3090&softwareid=282088129&release=4.2\(1\)SV1\(5.2\)&relind=AVAILABLE&rellifecycle=&reltype=latest](http://www.cisco.com/cisco/software/release.html?mdfid=282646785&flowid=3090&softwareid=282088129&release=4.2(1)SV1(5.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



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



5. Type the vCenter IP and the logon credentials.



**Nexus 1000V Installation Management Center**

**Steps**

1. Enter vCenter Credentials
2. Select the VSM's host
3. Select OVA File to create VSM
4. Configure Networking
5. Configure VSM
6. Review Configuration
7. Configure Migration
8. DVS Migration

**Enter vCenter Credentials**

vCenter IP: 10.29.165.41

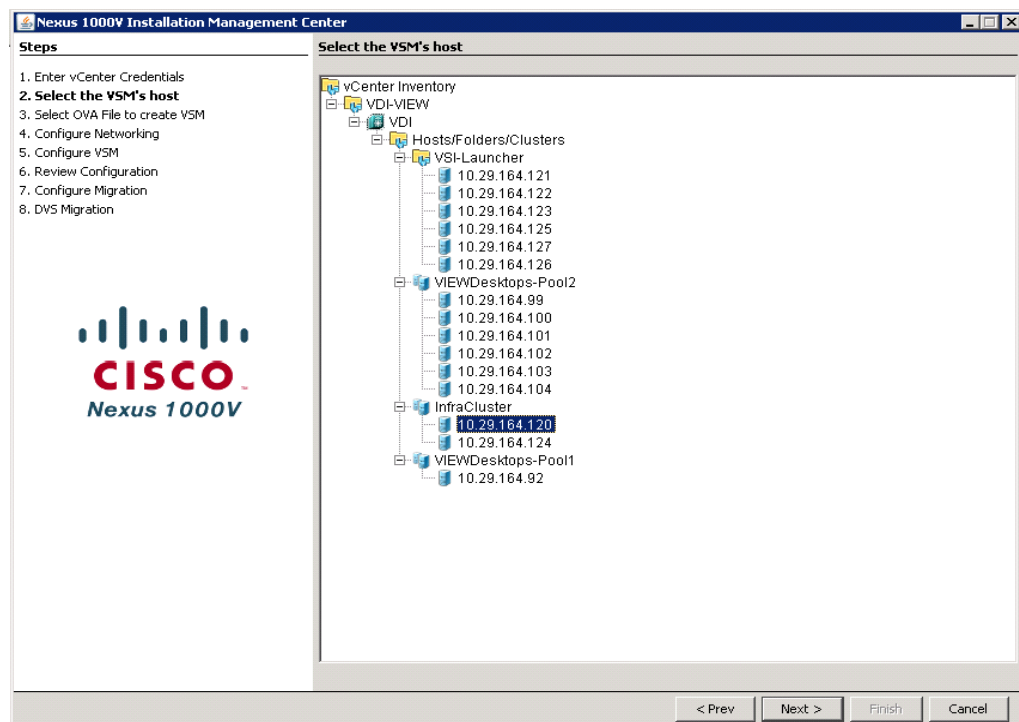
Port (https only): 443

vCenter User ID: administrator

vCenter Password: \*\*\*\*\*

< Prev Next > Finish Cancel

6. Select the ESX host on which to install N1KV Virtual Switch Manager.



**Nexus 1000V Installation Management Center**

**Steps**

1. Enter vCenter Credentials
2. Select the VSM's host
3. Select OVA File to create VSM
4. Configure Networking
5. Configure VSM
6. Review Configuration
7. Configure Migration
8. DVS Migration

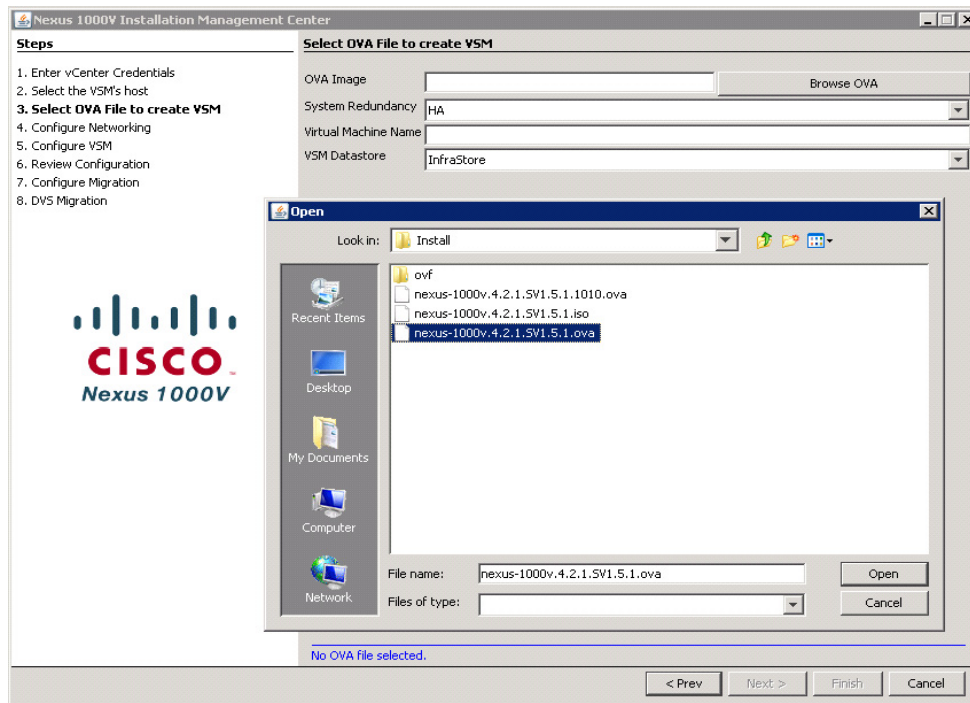
**Select the VSM's host**

vCenter Inventory

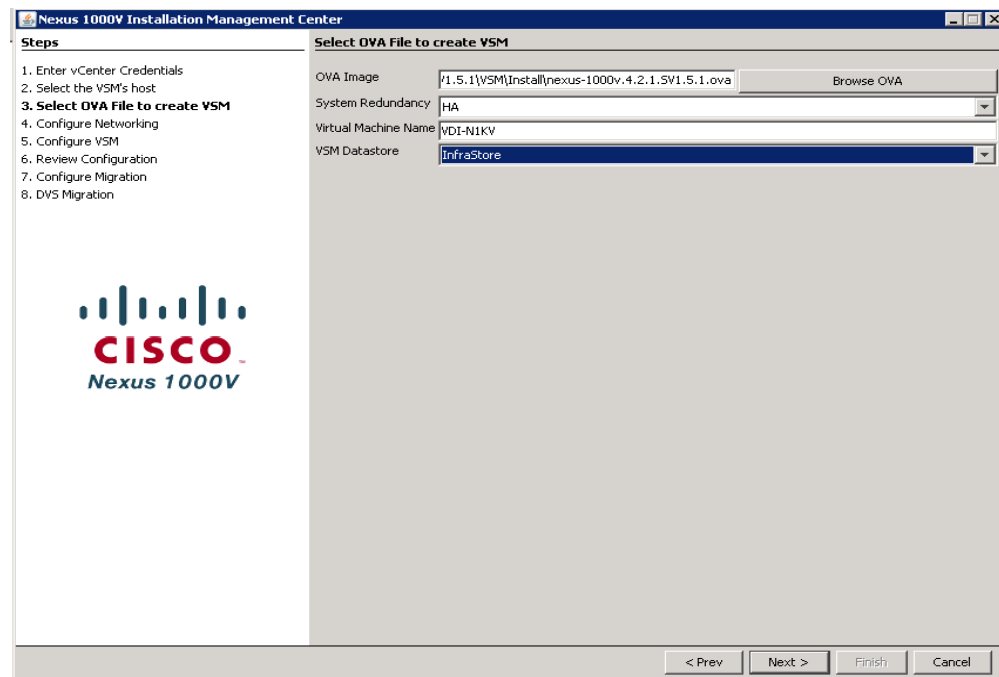
- VDI-VIEW
  - VDI
    - Hosts/Folders/Clusters
      - VSI-Launcher
        - 10.29.164.121
        - 10.29.164.122
        - 10.29.164.123
        - 10.29.164.125
        - 10.29.164.127
        - 10.29.164.126
      - VIEWDesktops-Pool2
        - 10.29.164.99
        - 10.29.164.100
        - 10.29.164.101
        - 10.29.164.102
        - 10.29.164.103
        - 10.29.164.104
      - InfraCluster
        - 10.29.164.120
        - 10.29.164.124
      - VIEWDesktops-Pool1
        - 10.29.164.92

< Prev Next > Finish Cancel

7. Select the OVA file from the extracted N1KV location to create the VSM.



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



9. To configure L3 mode of installation, choose the “L3: Configure port groups for L3”.
  - a. Create Port-group as Control and specify the VLAN ID and select the corresponding vSwitch.

- b. Select the existing port group “VM Network” for N1K Mgmt and choose mgmt0 with the VLAN ID for the SVS connection between vCenter and VSM.
- c. In the option for L3 mgmt0 interface port-profile enter the vlan that was pre-defined for ESXi mgmt and accordingly it will create a port-group which will have L3 capability. In this case it is n1kv-L3 port-group as shown in the screenshot below.

```
UDI-N1KV-DUS# sh run port-profile n1kv-L3
!Command: show running-config port-profile n1kv-L3
!Time: Fri Oct 26 16:53:34 2012

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

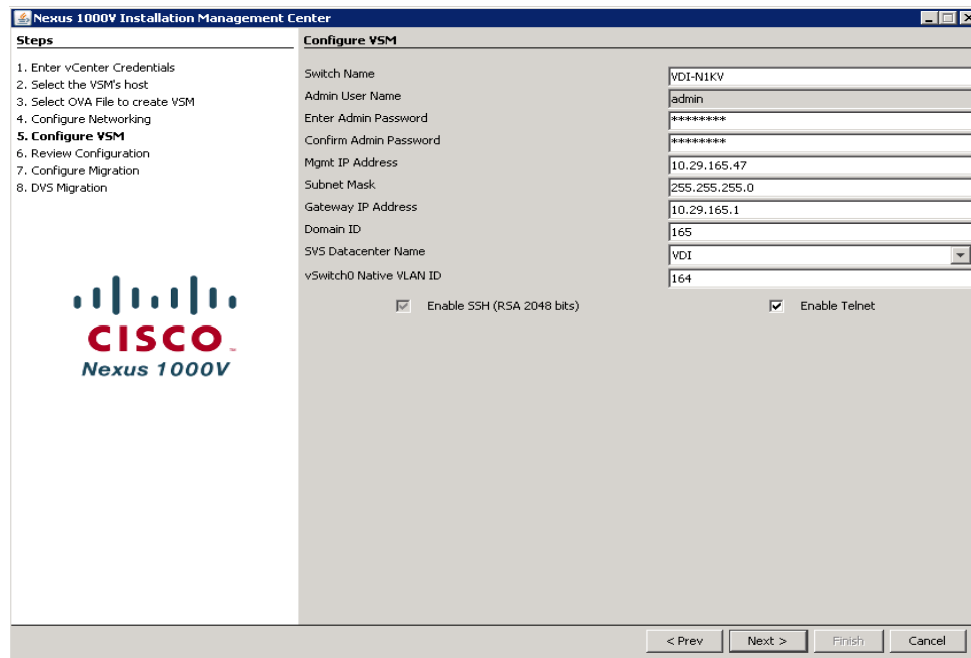
10. To Configure VSM, Type the Switch Name and Enter the admin password for the VSM. Type the IP address, subnet mask, Gateway, Domain ID.



#### Note

If there are multiple instance of N1KV VSM need to be install, make sure they each configured with different Domain ID and select the SVS datacenter Name and Type the vSwitch0 Native vlan ID. Make sure the Native VLAN ID specified should match the Native VLAN ID of Cisco UCS and the Nexus 5k.





**Nexus 1000V Installation Management Center**

**Steps**

1. Enter vCenter Credentials
2. Select the VSM's host
3. Select OVA File to create VSM
4. Configure Networking
- 5. Configure VSM**
6. Review Configuration
7. Configure Migration
8. DVS Migration

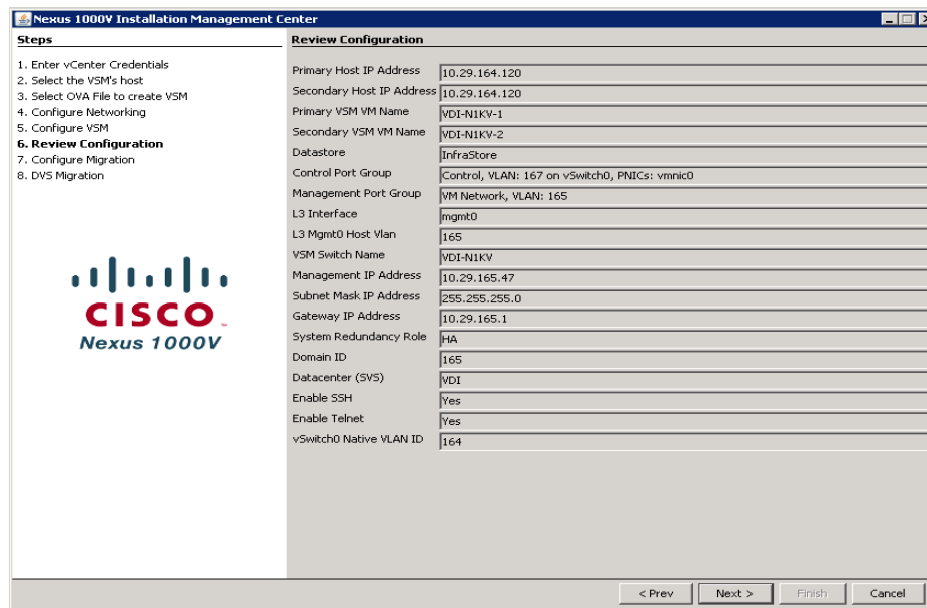
**Configure VSM**

Switch Name	VDI-N1KV
Admin User Name	admin
Enter Admin Password	*****
Confirm Admin Password	*****
Mgmt IP Address	10.29.165.47
Subnet Mask	255.255.255.0
Gateway IP Address	10.29.165.1
Domain ID	165
SVS Datacenter Name	VDI
vSwitch0 Native VLAN ID	164

☒ Enable SSH (RSA 2048 bits) ☒ Enable Telnet

< Prev Next > Finish Cancel

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



**Nexus 1000V Installation Management Center**

**Steps**

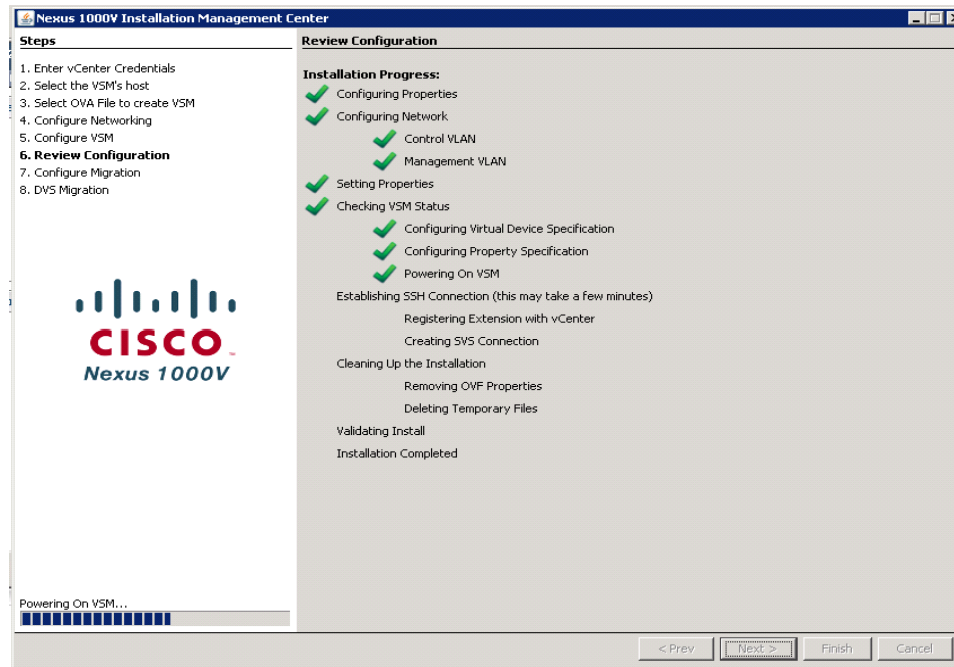
1. Enter vCenter Credentials
2. Select the VSM's host
3. Select OVA File to create VSM
4. Configure Networking
5. Configure VSM
- 6. Review Configuration**
7. Configure Migration
8. DVS Migration

**Review Configuration**

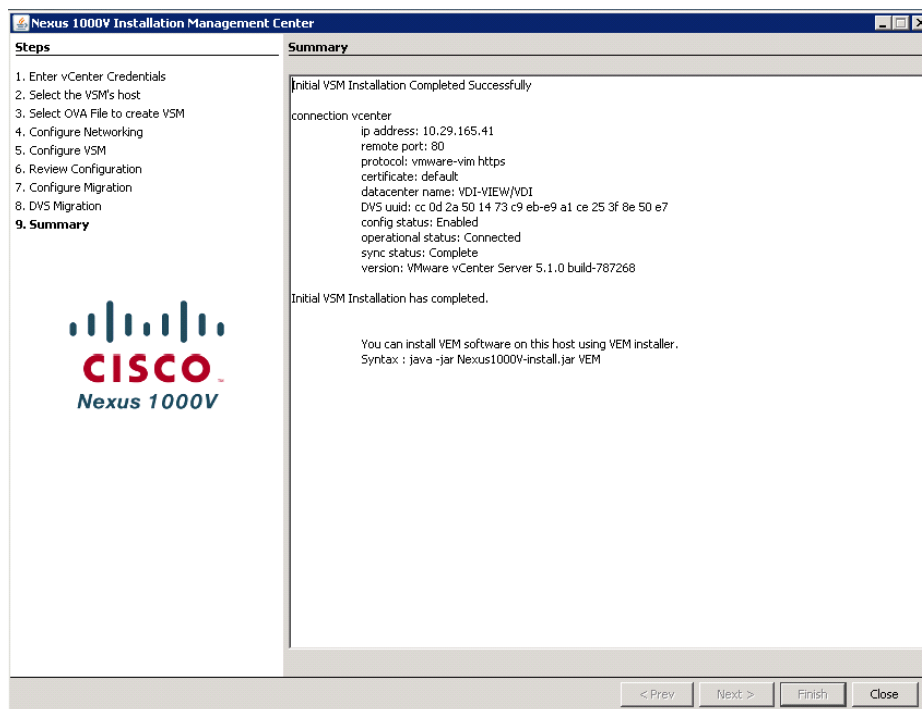
Primary Host IP Address	10.29.164.120
Secondary Host IP Address	10.29.164.120
Primary VSM VM Name	VDI-N1KV-1
Secondary VSM VM Name	VDI-N1KV-2
Datastore	InfraStore
Control Port Group	Control, VLAN: 167 on vSwitch0, PNICs: vmnic0
Management Port Group	VM Network, VLAN: 165
L3 Interface	mgmt0
L3 Mgmt0 Host Vlan	165
VSM Switch Name	VDI-N1KV
Management IP Address	10.29.165.47
Subnet Mask IP Address	255.255.255.0
Gateway IP Address	10.29.165.1
System Redundancy Role	HA
Domain ID	165
Datacenter (SVS)	VDI
Enable SSH	Yes
Enable Telnet	Yes
vSwitch0 Native VLAN ID	164

< Prev Next > Finish Cancel

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



13. Click Finish to complete the VSM installation.



14. Logon (ssh or telnet) to the N1KV VSM with the IP address and configure VLAN for ESX Mgmt, Control, N1K Mgmt and also for Storage and vMotion purposes as mentioned below (VLAN ID differs based on your Network).

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

Enter the following configuration commands, one per line. End with CNTL/Z.

```
VDI-N1KV(config)# vlan 800
VDI-N1KV(config-vlan)# name ML-VDA
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan 801
VDI-N1KV(config-vlan)# name ML-DC-VM-MGMT
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan 802
VDI-N1KV(config-vlan)# name ML-DC-VMOTION
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan 804
VDI-N1KV(config-vlan)# name ML-DC-STRG
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan 900
VDI-N1KV(config-vlan)# name ML-N1KV_CTRL
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# vlan 901
VDI-N1KV(config-vlan)# name ML-N1KV_PK
VDI-N1KV(config-vlan)# no sh
VDI-N1KV(config)# <CNTRL/Z>
```

```
vrf context management
  ip route 0.0.0.0/0 192.168.1.1
vlan 1,800-804,900-901
vlan 1
vlan 800
  name ML_VDA
vlan 801
  name ML_DC-VM-MGMT
vlan 802
  name ML_DC-VMOTION
vlan 803
  name ML_DC-INF
vlan 804
  name ML_DC-STRG
vlan 900
  name ML-N1KV_CTRL
vlan 901
  name ML-N1KV_PK
```

**15.** Run following configuration command to configure jumbo mtu and qos polices.

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

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

16. To Migrate and Manage all the ESXi host network using Nexus 1000V VSM, Configure Port Profiles and port groups as mentioned below.

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

**Note**

These port-profiles are created by default and do not make any changes.

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

```
VDI-N1KV(config)# port-profile type ethernet DC_System_Uplink
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode trunk
VDI-N1KV(config-port-prof)# switchport trunk allowed vlan 801,900-901
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no shutdown
VDI-N1KV(config-port-prof)# system vlan 801,900
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet DC_System_Uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 801,900-901
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 801,900
```

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

```
VDI-N1KV(config)# port-profile type ethernet DC_Storage_Uplink
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 804
```

```

VDI-N1KV(config-port-prof)# mtu 9000
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no shutdown
VDI-N1KV(config-port-prof)# system vlan 804
VDI-N1KV(config-port-prof)#state enabled

```

```

port-profile type ethernet DC_Storage_Uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 804
  mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 804
  state enabled

```

**19. Create the Storage virtual ethernet communications port profile:**

```

VDI-N1KV(config)# port-profile type vethernet Storage
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 804
VDI-N1KV(config-port-prof)# service -policy type qos input platinum_Cos_5
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan 804
VDI-N1KV(config-port-prof)#state enabled

```

```

port-profile type vethernet Storage
  vmware port-group
  switchport mode access
  switchport access vlan 804
  service-policy type qos input platinum_Cos_5
  no shutdown
  system vlan 804
  state enabled

```

**20. Create the DC vMotion Uplink port profile:**

```

VDI-N1KV(config)# port-profile type ethernet DC_vMotion_Uplink
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 802
VDI-N1KV(config-port-prof)# channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan 802
VDI-N1KV(config-port-prof)#state enabled

```

```
port-profile type ethernet DC_vMotion_Uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 802
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 802
  state enabled
```

**21. Create the virtual ethernet port profile for vMotion:**

```
VDI-N1KV(config)# port-profile type vethernet vMotion
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 802
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan 802
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type vethernet vMotion
  vmware port-group
  switchport mode access
  switchport access vlan 802
  no shutdown
  system vlan 802
  state enabled
```

**22. Create the DC VDA Uplink port profile:**

```
VDI-N1KV(config)# port-profile type ethernet DC_VDA_Uplink
VDI-N1KV(config-port-prof)# vmware port-group
VDI-N1KV(config-port-prof)#switchport mode access
VDI-N1KV(config-port-prof)#switchport access vlan 800,803
VDI-N1KV(config-port-prof)#channel-group auto mode on mac-pinning
VDI-N1KV(config-port-prof)#no shutdown
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet DC-VDA-Uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 800,803
  channel-group auto mode on mac-pinning
  no shutdown
  state enabled
```

**23. Create the virtual ethernet port profile for VDA traffic.**

```
VDI-N1KV(config)# port-profile type vethernet VDA
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# max-ports 1024
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 800
```



```

VDI-N1KV(config-port-prof)# no sh
VDI-N1KV(config-port-prof)# system vlan 800
VDI-N1KV(config-port-prof)#state enabled

```

```

port-profile type vethernet VDA
  vmware port-group
  switchport mode access
  switchport access vlan 800
  no shutdown
  system vlan 800
  max-ports 1024
  state enabled

```

**24.** Create the virtual ethernet port profile for VDA1 traffic:

```

VDI-N1KV(config)# port-profile type vethernet VDA1
VDI-N1KV (config-port-prof)# vmware port-group
VDI-N1KV (config-port-prof)# max-ports 1024
VDI-N1KV(config-port-prof)# switchport mode access
VDI-N1KV(config-port-prof)# switchport access vlan 800
VDI-N1KV(config-port-prof)# no sh
VDI-N1KV (config-port-prof)# system vlan 800
VDI-N1KV(config-port-prof)#state enable

```

```

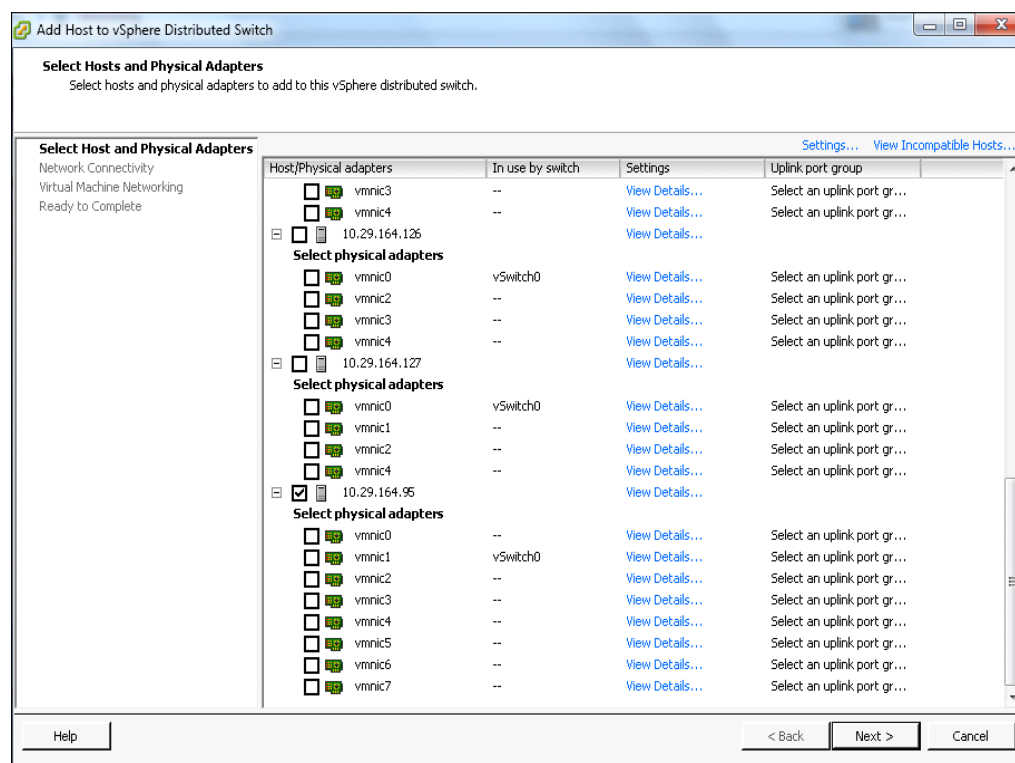
port-profile type vethernet VDA1
  vmware port-group
  switchport mode access
  switchport access vlan 800
  no shutdown
  system vlan 800
  max-ports 1024
  state enabled

```

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

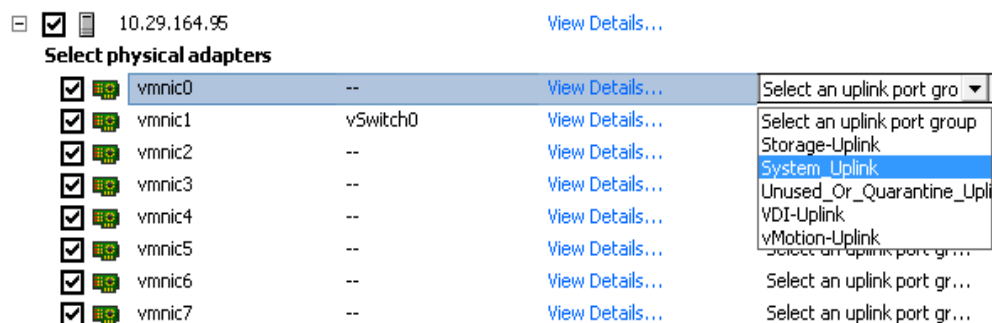
Name	State	VDS Status	Status	% CPU	% Memory	Memory Size	CPU Count	NIC Count	Uptime
10.29.164.91	Connected	Up	Normal	0	4	262086.00 MB	2	8	7 days
10.29.164.92	Connected	Up	Normal	0	2	262086.00 MB	2	8	7 days
10.29.164.93	Connected	Up	Normal	0	1	262086.00 MB	2	8	7 days
10.29.164.94	Connected	Up	Normal	0	1	262086.00 MB	2	8	7 days

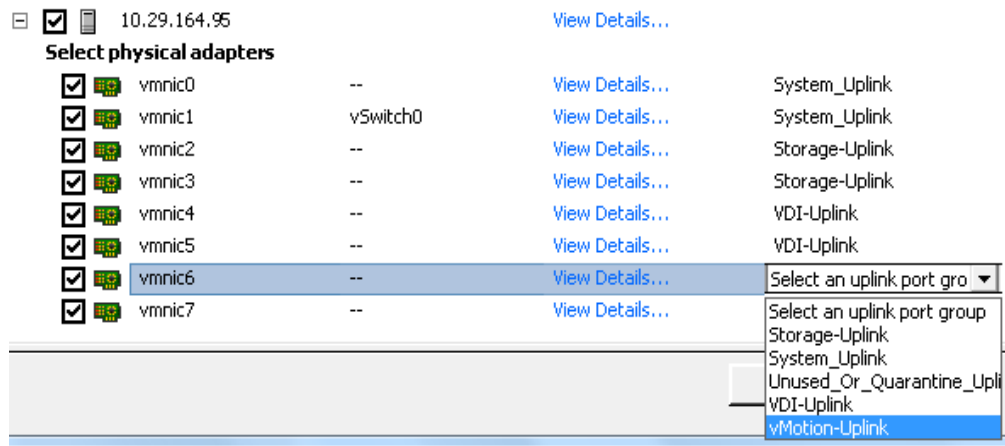

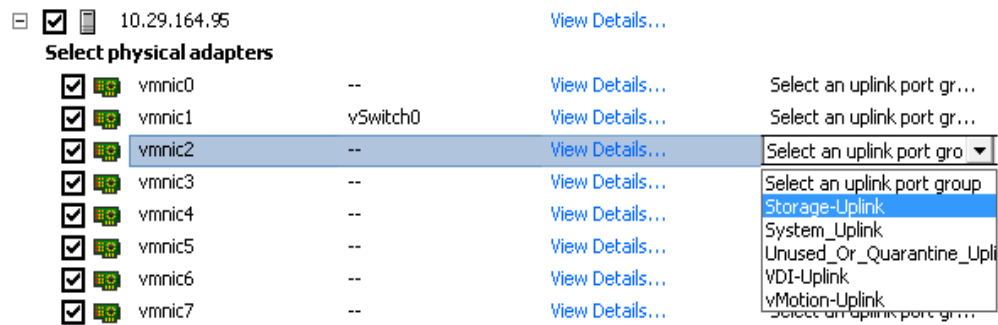
- 27.** Right-click and select add host to vSphere Distributed Switch.
- It will bring up ESXi hosts which are not part of existing configuration.
- 28.** Select ESXi hosts to add in N1KV.



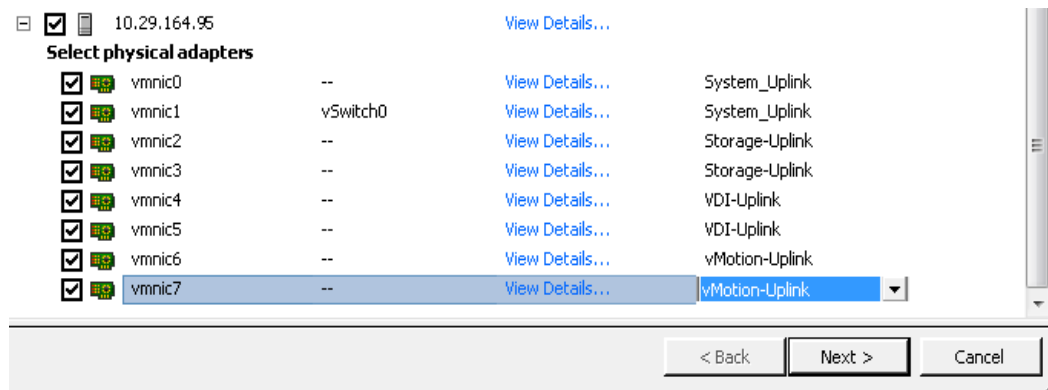
29. Click Select an uplink port-group and from the drop down menu select appropriate Uplink that is allowed for corresponding vmnic as per the configuration on Cisco UCSM vNICs.

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

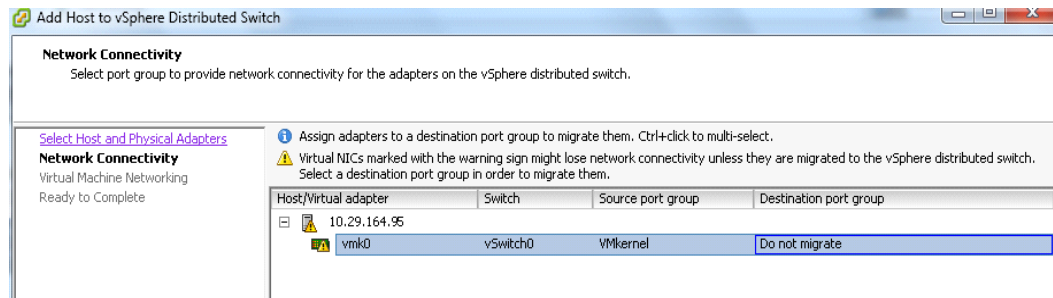




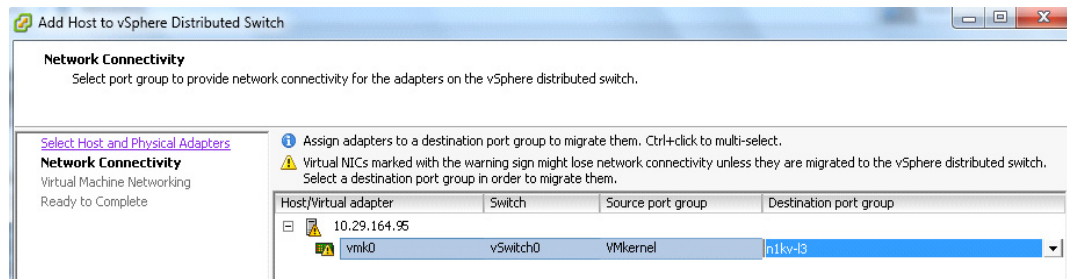
30. After selecting appropriate uplinks click Next.



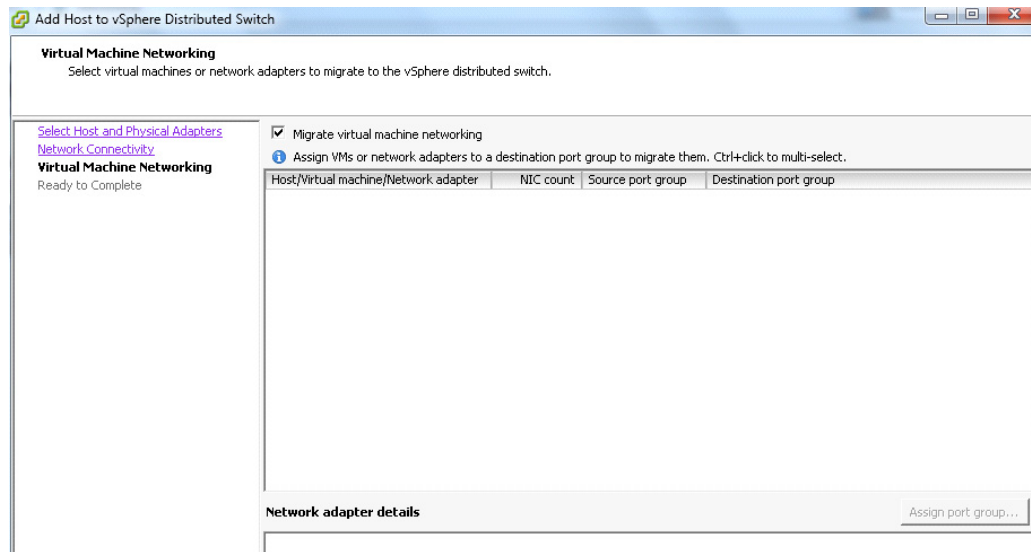
31. Network Connectivity tab select Destination port group for vmk0.



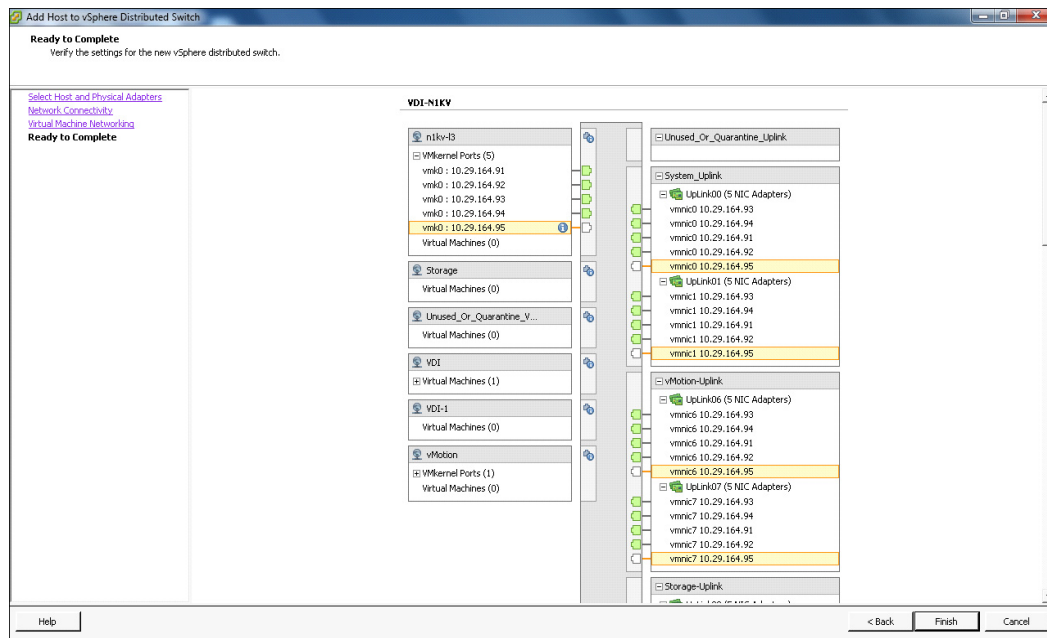
32. From the drop down menu select a port group which was configured for L3 capability and for ESXi host management communication. In this case it is n1kv-l3 and Click Next.



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



**34. Verify the Settings and Click Finish to add the ESXi host part of N1KV DVS.**

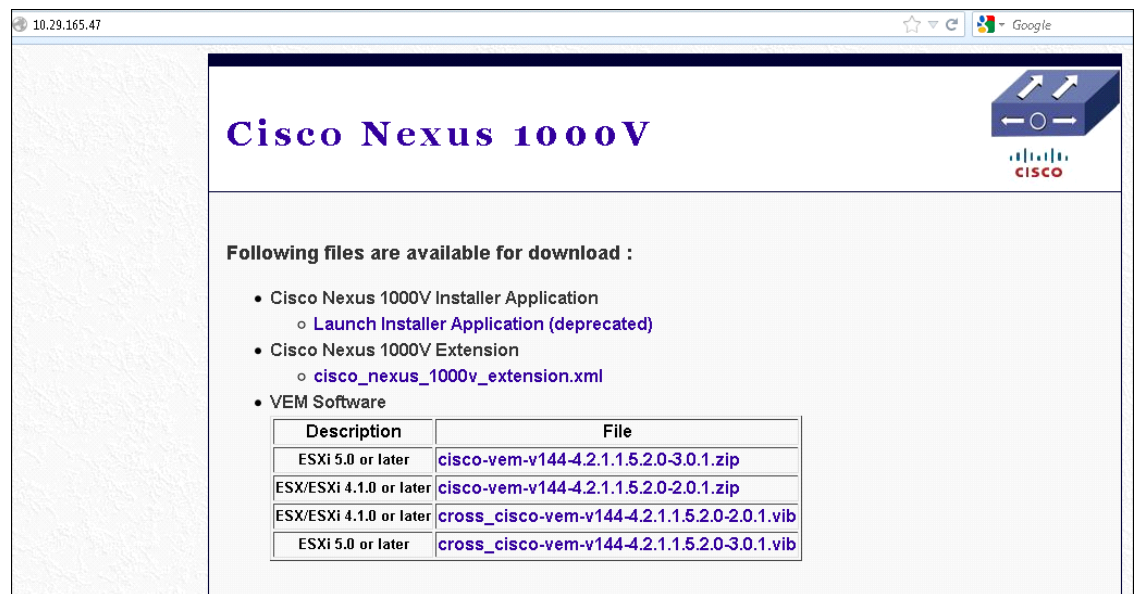


### Note

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

In the absence of Update manager:

35. Upload vib file [cross\\_cisco-vem-v144-4.2.1.1.5.2.0-3.0.1.vib](#) for VEM installation to local or remote datastore which can be obtained by browsing to the management IP address for N1KV VSM.

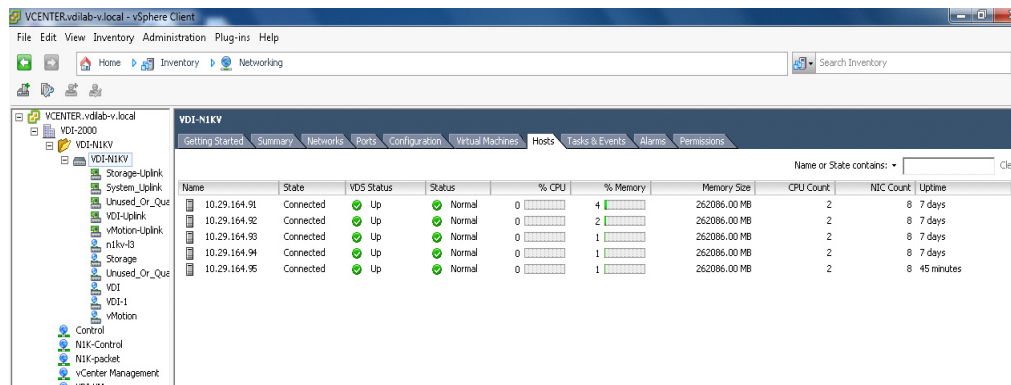


36. Login to ESXi host using ESXi shell or SSH session.

37. Run following command:

esxcli software vib install -v /vmfs/volumes/ datastore/ [cross\\_cisco-vem-v144-4.2.1.1.5.2.0-3.0.1.vib](#)

38. To verify the successful installation of ESXi VEM and the status of ESXi host.



39. To verify putty into N1KV VSM. Run sh module command which will show all the ESXi hosts attached to that VSM.

VDI-N1KV(config)# sh module



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

Mod	Ports	Module-Type	Model	Status
1	0	Virtual Supervisor Module	Nexus1000V	active *
2	0	Virtual Supervisor Module	Nexus1000V	ha-standby
16	248	Virtual Ethernet Module	NA	ok
17	248	Virtual Ethernet Module	NA	ok
18	248	Virtual Ethernet Module	NA	ok
19	248	Virtual Ethernet Module	NA	ok
20	248	Virtual Ethernet Module	NA	ok

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

Mod	Server-IP	Server-UUID	Server-Name
1	10.29.165.47	NA	NA
2	10.29.165.47	NA	NA
16	10.29.164.93	9476f312-1321-e111-0000-1b0000000006e	10.29.164.93
17	10.29.164.91	9476f312-1321-e111-0000-1b0000000005f	10.29.164.91
18	10.29.164.94	9476f312-1321-e111-0000-1b0000000007e	10.29.164.94
19	10.29.164.92	9476f312-1321-e111-0000-1b0000000001f	10.29.164.92
20	10.29.164.95	9476f312-1321-e111-0000-1b0000000004e	ESXi5-BFS-Srv5

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

## 6.4 SAN Configuration

The same pair of Nexus 5548UP switches were used in the configuration to connect between the FC ports on the EMC VNX7500 and the FC ports of the Cisco UCS 6248 Fabric Interconnects' expansion module.

### 6.4.1 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 pwn 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 BIOS settings. Once the hardware detects the boot device, it follows the regular boot process.

## 6.4.2 Configuring Boot from SAN Overview

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

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

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

## 6.4.3 SAN Configuration on Nexus 5548UP

The FCoE and NPIV feature has to be turned on in the Nexus 5500 series switch. Make sure you have 8 GB SPF+ modules connected to the Cisco UCS 6200UP series Fabric Interconnect expansion ports. The port mode is set to AUTO as well as the speed is set to AUTO. Rate mode is “dedicated” and when everything is configured correctly you should see something like the output below on a Nexus 5500 series switch for a given port (eg. Fc1/17).



### Note

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

Cisco Fabric Manager can also be used to get a overall picture of the SAN configuration and zoning information. As discussed earlier, the SAN zoning is done upfront for all the pwwns of the initiators with the EMC VNX7500 target pwwns.

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

```
npiv          1      enabled
```

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

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

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

Single vSAN zoning was set up on the Nexus 5548's to make those FAS3240 LUNs visible to the infrastructure and test servers.

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

```
VDI-N5548-A# sh zone name B230M2-CH1-SERVER1-FC0 vsan 1
```

```
zone name B230M2-CH1-SERVER1-FC0 vsan 1
```

```
member pwn 20:00:00:25:b5:c1:00:af
```

```
! [B230M2-CH1-SERVER1-fc0]
```

```
member pwn 50:06:01:60:46:e0:5e:0a
```

```
! [VNX7500-A0]
```

```
member pwn 50:06:01:69:46:e0:5e:0a
```

```
! [VNX7500-B1]
```

```
VDI-N5548-A# sh zone name B230M2-CH1-SERVER2-FC0 vsan 1
```

```
zone name B230M2-CH1-SERVER2-FC0 vsan 1
```

```
member pwn 20:00:00:25:b5:c1:00:9f
```

```
! [B230M2-CH1-SERVER2-fc0]
```

```
member pwn 50:06:01:60:46:e0:5e:0a
```

```
! [VNX7500-A0]
```

```
member pwn 50:06:01:69:46:e0:5e:0a
```

! [VNX7500-B1]

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

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

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

```
VDI-N5548-B# sh zone name B230M2-CH1-SERVER1-FC1 vsan 1
```

```
zone name B230M2-CH1-SERVER1-FC1 vsan 1
```

```
member pwwn 20:00:00:25:b5:c1:00:bf
```

```
[B230M2-CH1-SERVER1-fc1]
```

```
member pwwn 50:06:01:61:46:e0:5e:0a
```

```
[VNX7500-A1]
```

```
member pwwn 50:06:01:68:46:e0:5e:0a
```

```
[VNX7500-B0]
```

```
VDI-N5548-B# sh zone name B230M2-CH1-SERVER2-FC1 vsan 1
```

```
zone name B230M2-CH1-SERVER2-FC1 vsan 1
```

```
member pwwn 20:00:00:25:b5:c1:00:8f
```

```
[B230M2-CH1-SERVER2-fc1]
```

```
member pwwn 50:06:01:61:46:e0:5e:0a
```

```
[VNX7500-A1]
```

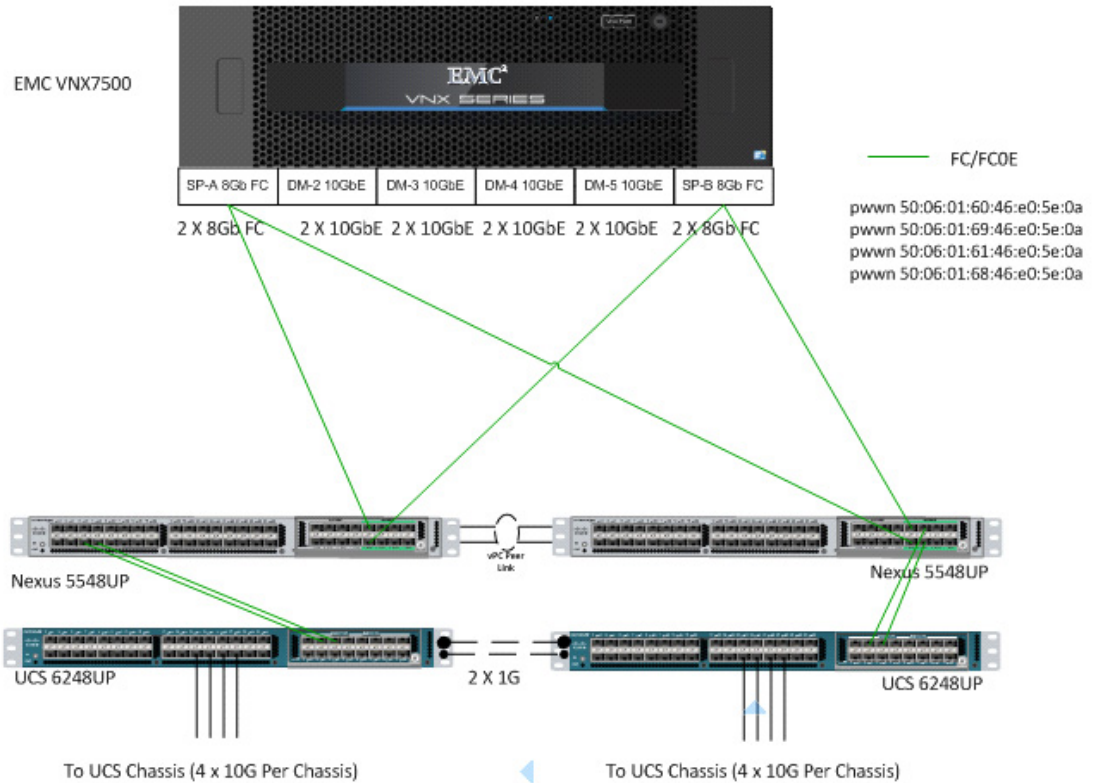
```
member pwwn 50:06:01:68:46:e0:5e:0a
```

```
[VNX7500-B0]
```

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

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

Figure 16 VNX7500 FC Target Ports



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

## 6.4.4 Configuring Boot from SAN on EMC VNX

The steps required 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. Five 600GB SAS drives are used in this example to create a RAID 5 pool. Uncheck "Schedule Auto-Tiering" to disable automatic tiering.

VNX7500 - Create Storage Pool

General Advanced

**Storage Pool Parameters**

Storage Pool Type: ☒ Pool ☐ RAID Group

☐ Scheduled Auto-Tiering

Storage Pool ID: 11

Storage Pool Name: VDA-POOL

**Extreme Performance**

RAID Configuration: RAID5 (4+1) Number of Flash Disks: 0

**Performance**

RAID Configuration: RAID5 (4+1) Number of SAS Disks: 5 (Recommended)

**Distribution**

Performance : 2684.038 GB (100.00%)

**Disks**

☐ Automatic ☐ Use Power Saving Eligible Disks

☒ Manual  Total Raw Capacity: 2684.0...


Disk	Capacity	Drive Type	Model	State
Bus 4 Enclosure 0 Disk 6	536.808 GB	SAS	HUS15606 C...	Unbound
Bus 4 Enclosure 0 Disk 5	536.808 GB	SAS	HUS15606 C...	Unbound
Bus 0 Enclosure 1 Disk 14	536.808 GB	SAS	STE60005 C...	Unbound
Bus 0 Enclosure 1 Disk 13	536.808 GB	SAS	STE60005 C...	Unbound
Bus 0 Enclosure 1 Disk 12	536.808 GB	SAS	STE60005 C...	Unbound

☒ Perform a background verify on the new storage and set priority to Medium


OK Apply Cancel Help

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

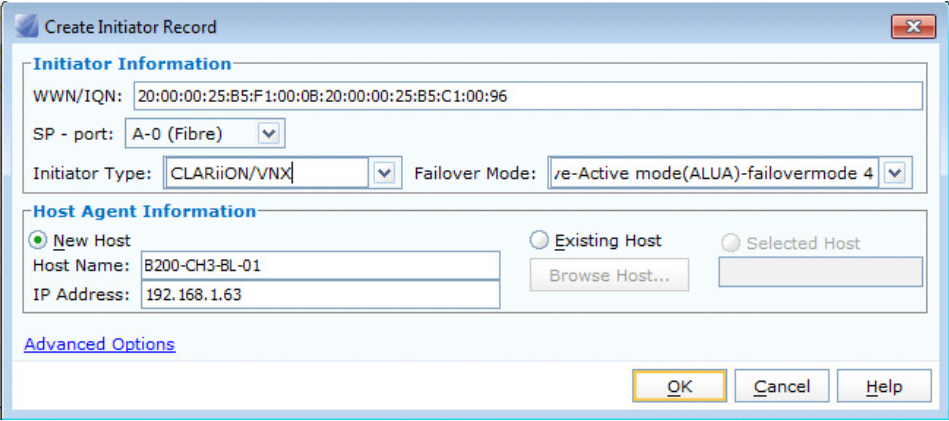




3. Create a storage group, the container used for host to LUN mapping, for each of the ESXi hosts.



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



**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: /e-Active mode(ALUA)-failovermode 4

**Host Agent Information**

☒ New Host ☐ Existing Host ☐ Selected Host

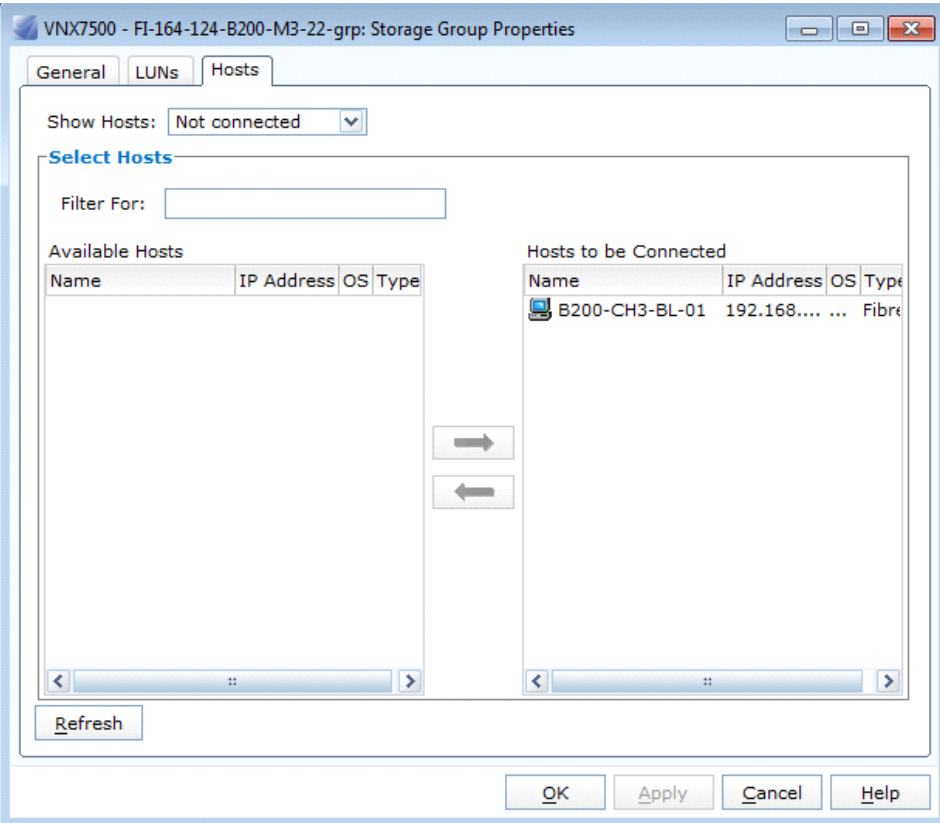
Host Name: B200-CH3-BL-01

IP Address: 192.168.1.63

[Advanced Options](#)

OK Cancel Help

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



**VNX7500 - FI-164-124-B200-M3-22-grp: Storage Group Properties**

General LUNs Hosts

Show Hosts: Not connected

**Select Hosts**


Filter For:

Available Hosts				Hosts to be Connected			
Name	IP Address	OS	Type	Name	IP Address	OS	Type
				B200-CH3-BL-01	192.168....	...	Fibre

Refresh

OK Apply Cancel Help

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

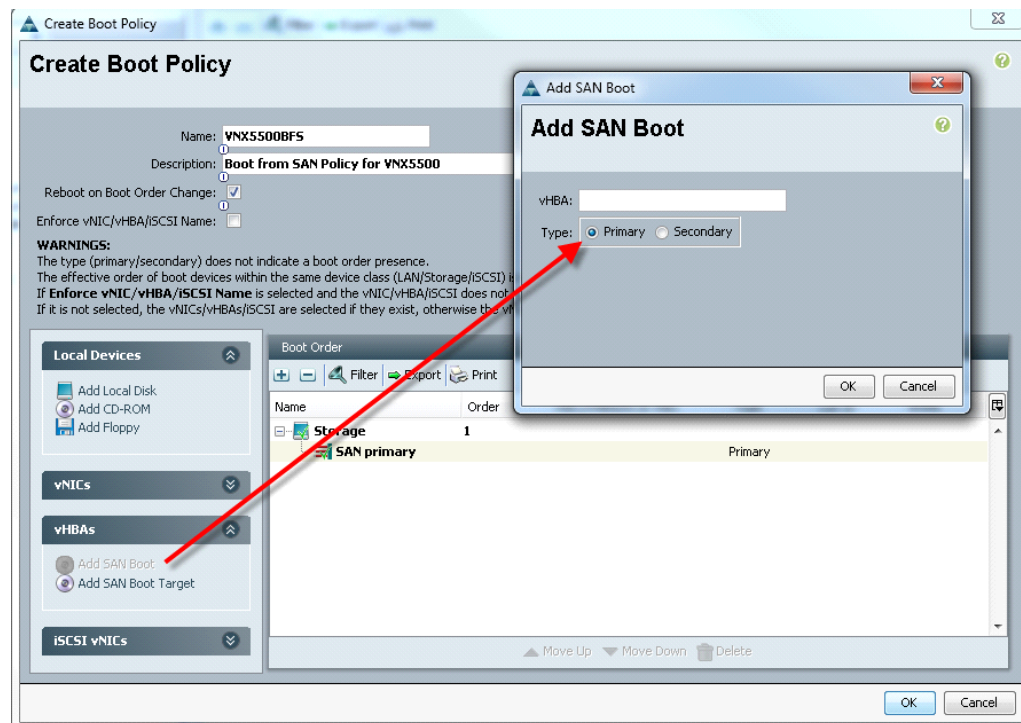


When the Cisco UCS Blade Server boots up, its vHBAs will connect to the provisioned EMC Boot LUNs and the hypervisor operating system can be installed.

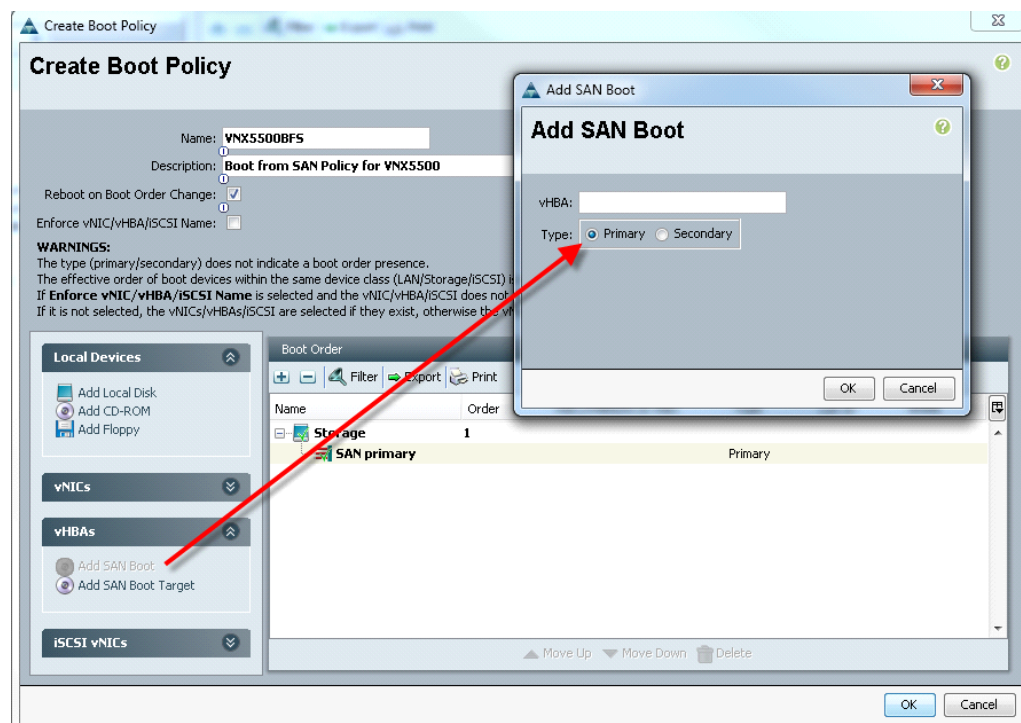
### 6.4.5 SAN Configuration on Cisco UCS Manager

To enable Boot from SAN on the Cisco UCS Manager 2.0 (UCS-M) series, do the following:

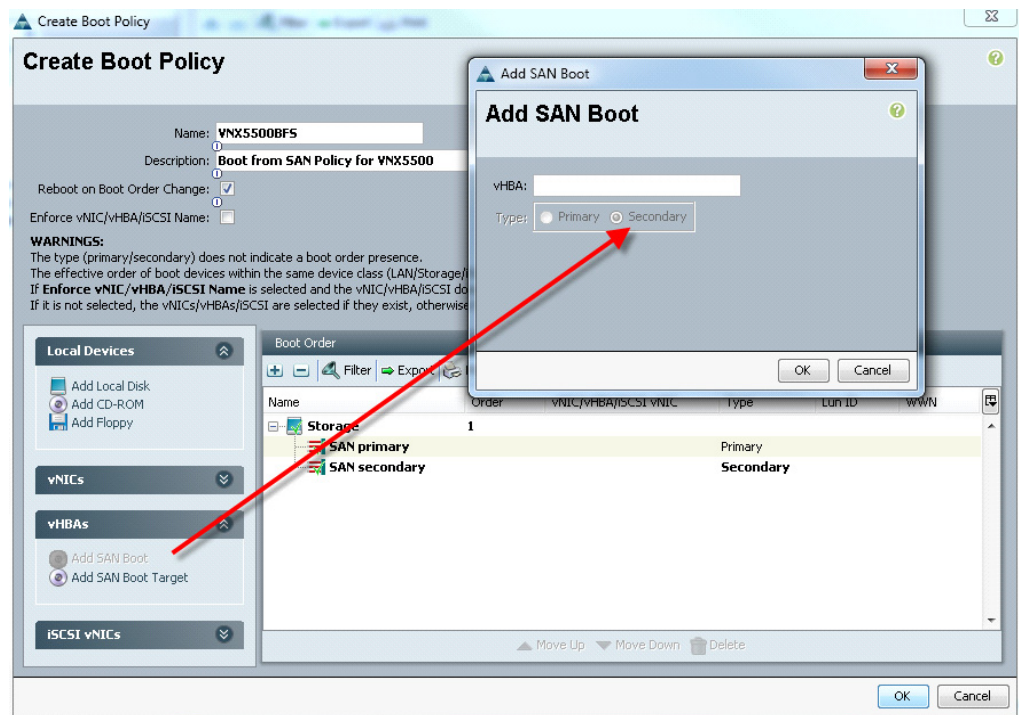
1. Add SAN Boot for primary to the new policy. The vHBA name is optional, it could be left blank and we don't have to enforce the vHBA name. Click OK.



2. Add SAN Boot for primary to the new policy. The vHBA name is optional, it could be left blank and we don't have to enforce the vHBA name. Click OK.



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



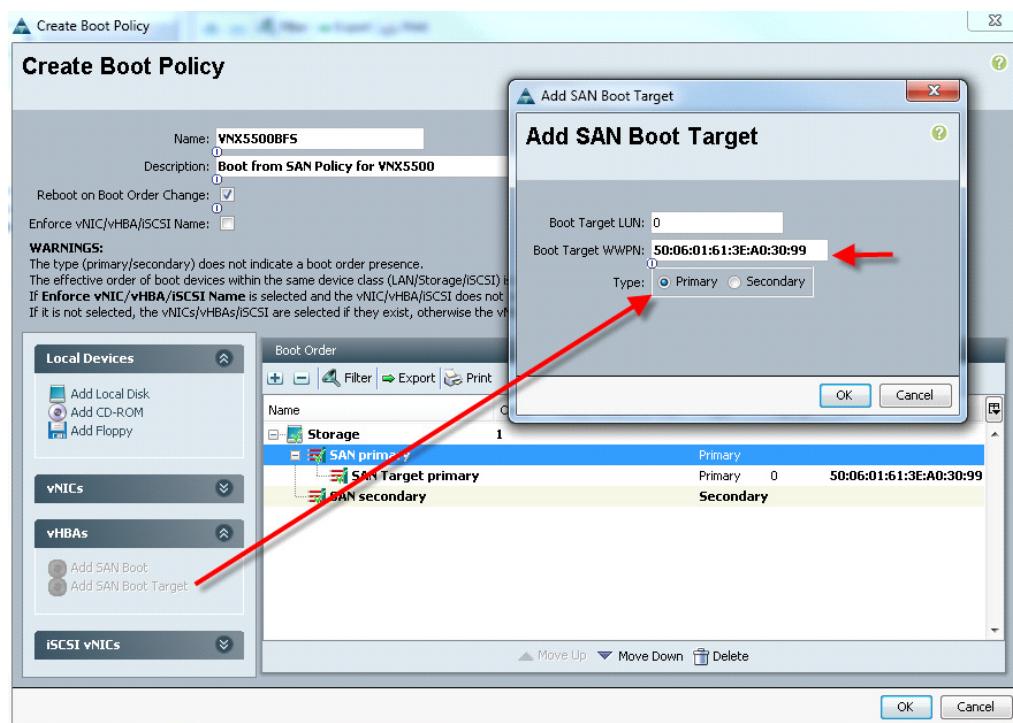
- Now add Boot target WWPN to the SAN Primary, make sure this is exactly matches the EMC VNX pwwn. To avoid any typos, copy and paste from Nexus 5500 Series command as follows from each switch:

**VDI-N5548-A# show fcns database vsan 1**

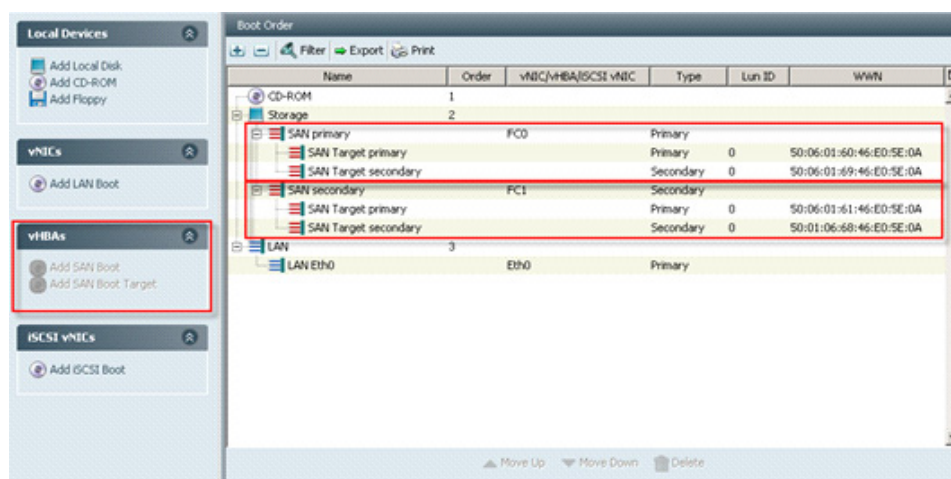
```
0xe300ef N 50:06:01:60:46:e0:5e:0a (Clariion) scsi-fcp:both
0xe301ef N 50:06:01:69:46:e0:5e:0a (Clariion) scsi-fcp:both
```

**VDI-N5548-B # show fcns database vsan 1**

```
0x470400 N 50:06:01:61:46:e0:5e:0a (Clariion) scsi-fcp
0x470500 N 50:06:01:68:46:e0:5e:0a (Clariion) scsi-fcp
```



5. Repeat step 4 for SAN primary's SAN Target Secondary.
6. Repeat step 4 for SAN Secondary's – SAN Target Primary.
7. Repeat step 4 for SAN Secondary's – SAN Target Secondary.
8. At the end your Boot from SAN policy should look like:

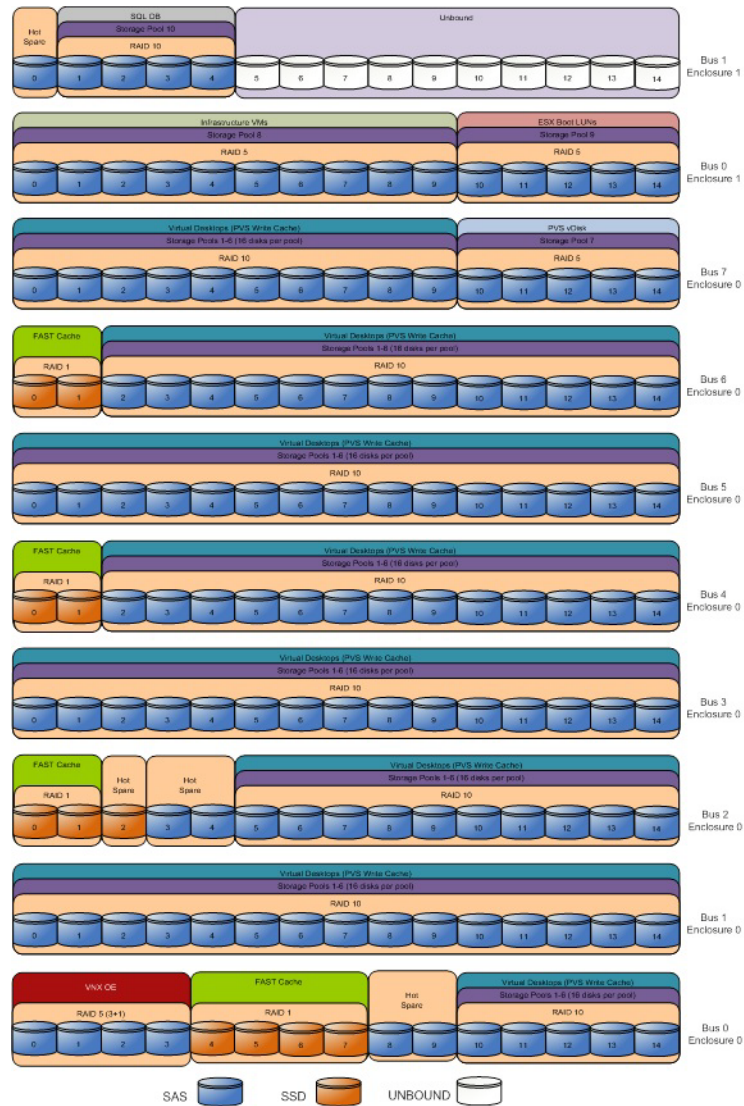


9. The last step is to make the association of the service profile template to the Boot from SAN policy during the service profile template configuration which we covered earlier in this document.



## 6.5 EMC VNX7500 Storage Configuration

The figure below shows the physical storage layout of the disks in the reference architecture. This configuration accommodates 5000 virtual desktops, hypervisor boot LUNs, SQL database, and infrastructure virtual machines.



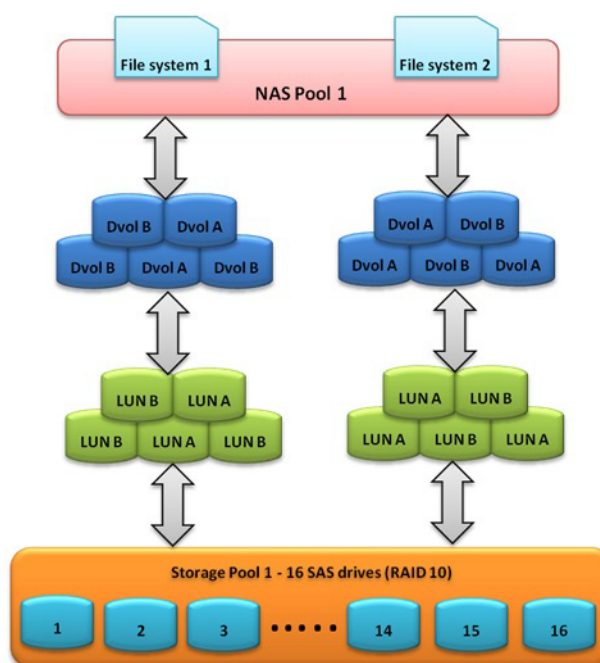
The above storage layout is used for the following configurations:

- Four SAS disks (0\_0\_0 to 0\_0\_3) are used for the VNX OE.
- Five disks (0\_0\_8 to 0\_0\_9, 2\_0\_3 to 2\_0\_4, 1\_1\_0) are hot spares for SAS disks. Disk 2\_0\_2 is hot spare for SSD drives. These disks are marked as hot spare in the storage layout diagram.
- Ten 200GB Flash drives (on the even number buses) are used for EMC VNX FAST Cache. See the “EMC FAST Cache in Practice” section below to follow the FAST Cache configuration best practices.
- 96 SAS disks (spread across enclosures 0\_0 through 7\_0) on the RAID 10 storage pools 1-6 are used to store PVS write cache allocated for the virtual desktops. Each of the six storage pools consists of 16 drives (8+8). FAST Cache is enabled on these pools.

- Five SAS disks (7\_0\_10 to 7\_0\_14) on the RAID 5 storage pool 7 are used to store the PVS vDisks. FAST Cache is enabled on this pool.
- Ten SAS disks (0\_1\_0 to 0\_1\_9) on the RAID 5 storage pool 8 are used to store the infrastructure virtual machines.
- Five SAS disks (0\_1\_10 to 0\_1\_14) on the RAID 5 storage pool 9 are used to store the ESXi boot LUNs.
- Four SAS disks (1\_1\_1 to 1\_1\_4) on the RAID 10 storage pool 10 are used to store the SQL databases and quorum disk.
- Disks 1\_1\_5 to 1\_1\_14 are unbound. They are not used for testing this solution.
- All SAS disks used for this solution are 600GB.

### 6.5.1 Example EMC Volume Configuration for PVS Write Cache

The figure below shows the layout of the NFS file systems used to store the PVS write cache for the virtual desktops:



Ten LUNs of 411GB each are carved out of a RAID 10 storage pool configured with 16 SAS drives. The LUNs are presented to VNX File as dvols that belong to a system defined NAS pool. Two 2.1TB file systems are then carved out of the NAS pool and are presented to the ESXi servers as two NFS datastores. Because there are six storage pools to support PVS write cache, a total of twelve 2.1TB NFS file systems are used.

### 6.5.2 EMC Storage Configuration for PVS vDisks

Similar to the PVS write cache storage, ten LUNs of 205GB each are carved out of the RAID 5 storage pool configured with 5 SAS drives to support a 500GB NFS file system that is designated to store PVS vDisks for the desktops.

### 6.5.3 EMC Storage Configuration for VMware ESXi 5.0 Infrastructure and VDA Clusters

Two LUNs of 2.08TB are carved out of the RAID 5 storage pool configured with 10 SAS drives. The LUNs are used to store infrastructure virtual machines such as XenDesktop controllers, PVS servers, and VMware vCenter server.

### 6.5.4 Example EMC Boot LUN Configuration

Each ESXi server requires a boot LUN from SAN for the hypervisor OS. A total of 43 LUNs are carved out of the 5-disk RAID 5 pool. Each LUN is 12GB in size.

### 6.5.5 Example EMC FC LUN Configuration for SQL Clustering

Three LUNs are provisioned from the 4-disk RAID 10 storage pool to support SQL Server clustering. The LUN configurations are as follow:

1. Data – 100GB
2. Logs – 25GB
3. Quorum – 2GB

The SQL server cluster is used to provide the required database services for XenDesktop controllers, Provisioning services, and VMware vCenter server.

### 6.5.6 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 helps ensure 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.

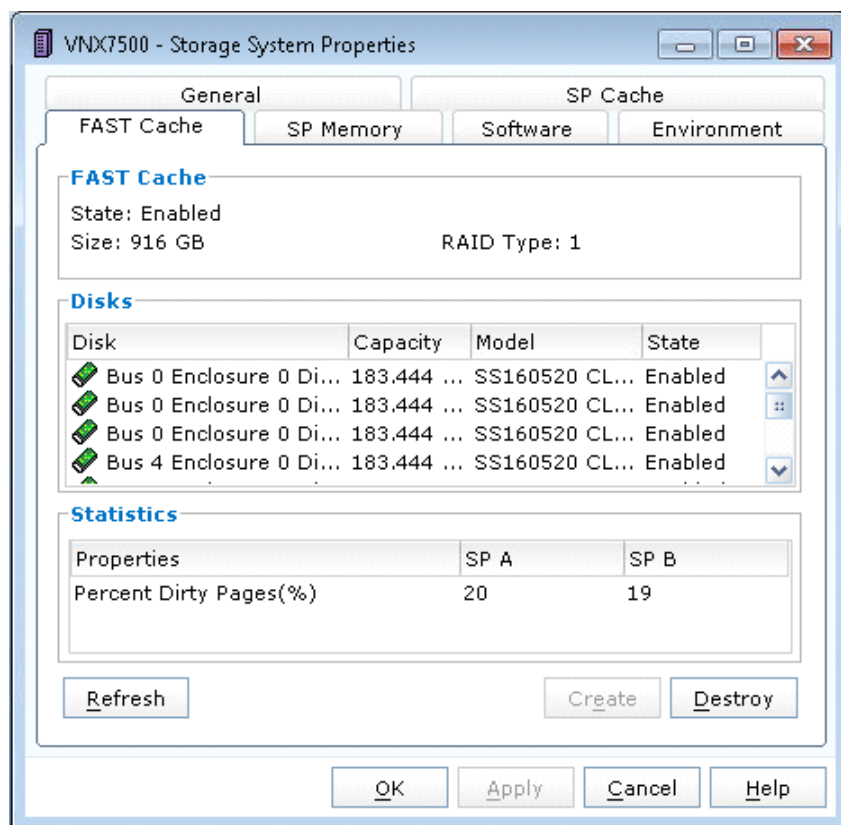
Only a small number of Flash drives are needed enabling FAST Cache to provide greater performance increases than implementing a large number of short-stroked HDDs. This results in cost savings in data center space, power, and cooling requirements that lowers overall TCO for the business.

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 patterns. Applications with these characteristics benefit most from deploying FAST Cache. The optimal use of FAST Cache is achieved when the working data set can fit within the FAST Cache.

FAST Cache is enabled as an array-wide feature in the system properties of the array in EMC Unisphere. Click the **FAST Cache** tab, then click **Create** and select the Flash drives to create the FAST Cache. RAID 1 is the only RAID type allowed. There are no user-configurable parameters for FAST Cache. However, it is important to select the drives in the right order such that each RAID 1 pairs are not mirrored across buses/enclosures. For example, if disks 1, 2, 3, and 4 are selected in that order to create a 4-disk FAST Cache, disks 1 and 2 will belong to one RAID-1 group, and disks 3 and 4 will belong to another group. You have to make sure disks 1 and 2 reside in the same enclosure (likewise for disks 3 and 4) to avoid drive mirroring across enclosures.

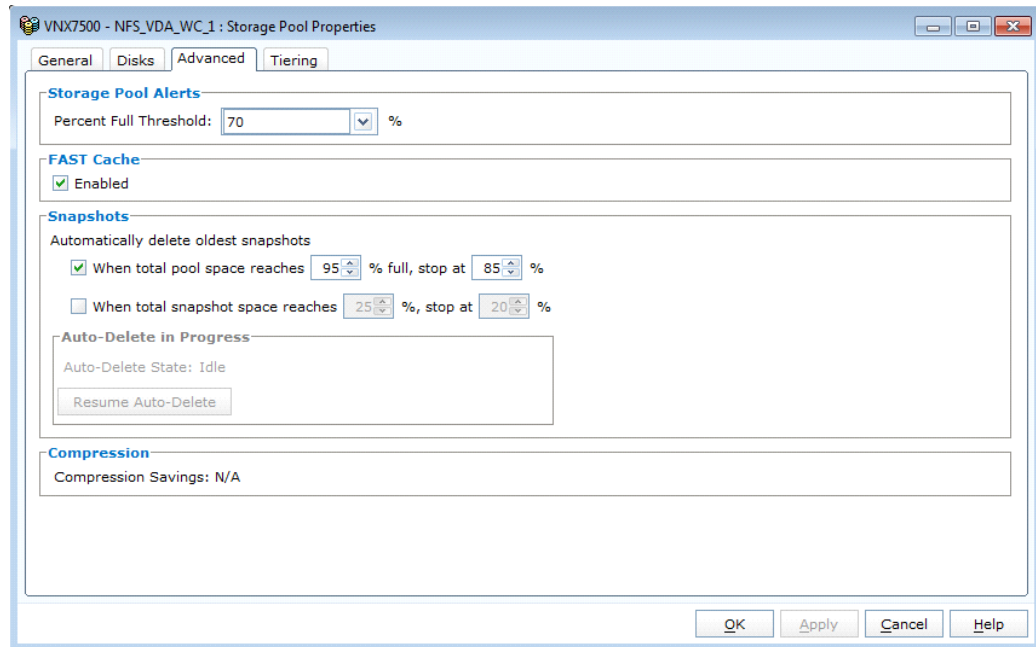
In this solution, ten 200GB SSD drives were used for FAST Cache. The ten drives were spread across the even number buses. Figure 16 shows the FAST Cache settings for VNX7500 array used in this solution.

**Figure 17** VNX7500–FAST Cache tab



To enable FAST Cache for any LUN in a pool, navigate to the **Storage Pool Properties** page in Unisphere, and then click the **Advanced** tab. Select **Enabled** to enable FAST Cache as shown in Figure 17.

Figure 18 VNX7500–Enable FAST Cache



## 6.5.7 EMC Additional Configuration Information

The following tuning configurations optimize the NFS performance on the VNX 7500 Data Movers:

### 6.5.7.1 NFS Active Threads Per Data Mover

The default number of threads dedicated to serve NFS requests is 384 per Data Mover on the VNX. Some use cases such as the scanning of desktops might require more number of NFS active threads. It is recommended to increase the number of active NFS threads to the maximum of 2048 on each Data Mover. The **nthreads** parameter can be set by using the following command:

```
# server_param server_2 -facility nfs -modify nthreads -value 2048
```

Reboot the Data Mover for the change to take effect.

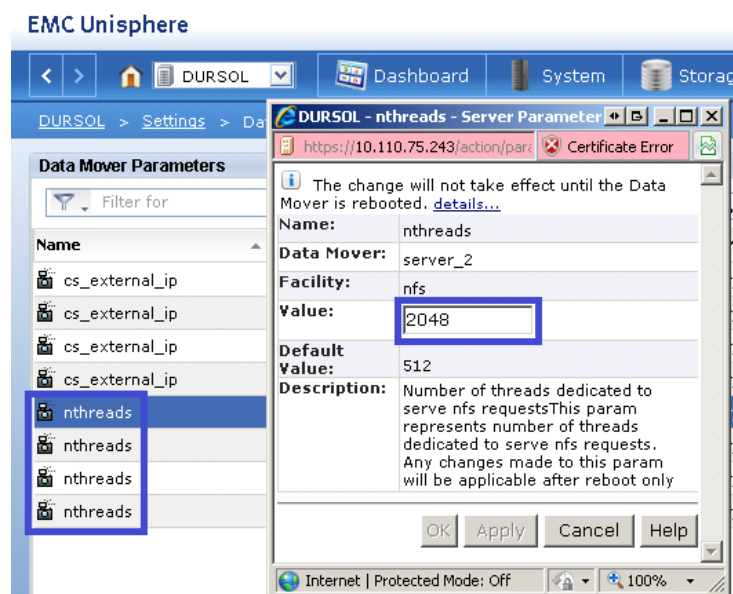
Type the following command to confirm the value of the parameter:

```
# server_param server_2 -facility nfs -info nthreads
server_2 :
name                = nthreads
facility_name        = nfs
default_value        = 384
current_value        = 2048
configured_value     = 2048
user_action          = reboot DataMover
change_effective     = reboot DataMover
range                = (32,2048)
```

description = Number of threads dedicated to serve nfs requests This param represents number of threads dedicated to serve nfs requests. Any changes made to this param will be applicable after reboot only.

The NFS active threads value can also be configured by editing the properties of the **nthreads** Data Mover parameter in **Settings–Data Mover Parameters** menu in Unisphere, as shown in VNX7500–nThreads properties. Highlight the **nthreads** value you want to edit and select **Properties** to open the nthreads properties window. Update the **Value** field with the new value and click **OK** as shown in VNX7500–nThreads properties. Perform this procedure for each of the **nthreads** Data Mover parameters listed menu. Reboot the Data Movers for the change to take effect.

Figure 19 VNX7500–nThreads properties



### 6.5.7.2 NFS Performance Fix

VNX file software contains a performance fix that significantly reduces NFS write latency. The minimum software patch required for the fix is 7.0.13.0. In addition to the patch upgrade, the performance fix only takes effect when the NFS file system is mounted by using the **uncached** option as shown below:

```
# server_mount server_2 -option uncached fs1 /fs1
```

The **uncached** option can be verified by using the following command:

```
# server_mount server_2
server_2 :
root_fs_2 on / uxfs,perm,rw
root_fs_common on /.etc_common uxfs,perm,ro
fs1 on /fs1 uxfs,perm,rw,uncached
fs2 on /fs2 uxfs,perm,rw,uncached
fs3 on /fs3 uxfs,perm,rw,uncached
fs4 on /fs4 uxfs,perm,rw,uncached
```



```
fs5 on /fs5 ufs,perm,rw,uncached
fs6 on /fs6 ufs,perm,rw,uncached
fs7 on /fs7 ufs,perm,rw,uncached
fs8 on /fs8 ufs,perm,rw,uncached
```

The uncached option can also be configured by editing the properties of the file system mount in **Storage–Storage Configuration–File Systems–Mounts** menu in Unisphere. Highlight the file system mount you want to edit and select **Properties** to open the **Mount Properties** window as shown in **VNX7500–File System Mount Properties**. Select the **Set Advanced Options** checkbox to display the advanced menu options, and then select the **Direct Writes Enabled** checkbox and click **OK**. The uncached option is now enabled for the selected file system.

**Figure 20 VNX7500–File System Mount Properties**

Path: /FS1

Data Mover: server\_5

File System Name: FS1

Read Only: ☒ Read/Write  
☐ Read Only

Access-Checking Policy: ☐ NT - CIFS client rights checked against ACLs; NFS client rights checked against ACLs and permission bits  
☐ UNIX - NFS client rights checked against permission bits; CIFS client rights checked against permission bits AND ACLs  
☐ SECURE - Both NFS and CIFS client rights checked against both permission bits AND ACLs  
☒ NATIVE - NFS client rights checked against permission bits; CIFS client rights checked against ACLs  
☐ MIXED - Both NFS and CIFS client rights checked against ACL; Only a single set of security attributes maintained  
☐ MIXED\_COMPAT - Both NFS and CIFS client rights checked against either permission bits or ACL depending on which protocol was last used to set permissions

Virus Checking Enabled: ☐

Cifs Oplocks Enabled: ☐

Set Advanced Options: ☒

Use NT Credential: ☐

Direct Writes Enabled: ☒

Prefetch Enabled: ☒

Multi-Protocol Locking Policy: ☒ nolock  
☐ writelock  
☐ relock

CIFS Sync Writes Enabled: ☐

CIFS Notify Enabled: ☐

CIFS Notify Trigger Level: 512

CIFS Notify On Access Enabled: ☐

CIFS Notify On Write Enabled: ☐

OK Apply Cancel Help

## 6.6 Cisco UCS Manager Configuration for VMware ESXi 5.0 Update 1

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

### 6.6.1 Service Profile Templates

Two types of service profiles were required to support two different blade server types:

**Table 7 Role/Server/OS Deployment**

Role	Blade Server Used	Operating System Deployed
Infrastructure	Cisco UCS B200 M3	ESXi 5.0 Update 1

VDI Hosts                      Cisco UCS B230 M2                      ESXi 5.0 Update 1

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

The service profile templates were then used to quickly deploy service profiles for each blade server in the Cisco Unified Computing System. When each blade server booted for the first time, the service profile was deployed automatically, providing the perfect configuration for the VMware ESXi 5.0 Update 1 installation.

## 6.6.2 VLAN Configuration

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

A total of seven Virtual LANs, VLANs, were utilized for the project. The following table identifies them and describes their use:

**Table 8**

**VLAN Naming and Use**

VLAN Name	VLAN ID	Use
ML-VDA	800	VDI Virtual Machine Traffic
ML-DC-VM-MGMT	801	VMware Management and Nexus 1000V Management Traffic
ML-DC-VMOTION	802	VMware vMotion Traffic
ML_DC-INF	803	VDI Infrastructure
ML-DC-STRG	804	VNX7500 NFS Traffic
ML-N1KV_CTR	900	Nexus 1000V Control Traffic
VLAN0901	901	Nexus 1000V Packet Traffic

VLANs are configured in UCS Manager on the LAN tab, LAN\VLANs node in the left pane of Cisco UCS Manager. They were set up earlier in section 6.2.1 Base Cisco UCS System Configuration.

## 6.7 Installing and Configuring ESXi 5.0 Update 1

In this study, we used Fibre Channel storage to boot the hosts from LUNs on the VNX7500 storage system. Prior to installing the operating system, storage groups are created, assigning specific boot LUNs to individual hosts. (See Section 6.4.4 Configuring Boot from SAN on EMC VNX for details.)

VMware ESXi 5.0 Update 1 can be installed in boot-from-SAN mode using standard hypervisor deployment techniques including:

1. Mounting a Cisco Customized ESXi 5.0 Update 1 ISO image from the KVM of the blade
2. Using automated deployment tools from third party sources (Optional)

### 6.7.1 Install VMware ESXi 5.0 Update 1

ESXi was installed using Symantec Altiris Deployment Solution 7.1.

ESXi installation script that performs a standard installation to the first detected disk, and enables SSH and ESXi Shell was used to automate installation process

([http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc\\_50/GUID-C3F32E0F-297B-4B75-8B3E-C28BD08680C8.html](http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc_50/GUID-C3F32E0F-297B-4B75-8B3E-C28BD08680C8.html)).

Example of ks.cfg file used in install:

```
accepteula
install --firstdisk --overwritevmfs
rootpw Password
network --bootproto=dhcp --device=vmnic0 --addvmportgroup=0
reboot
%firstboot --interpreter=busybox
# enable & start remote ESXi Shell (SSH)
vim-cmd hostsvc/enable_ssh
vim-cmd hostsvc/start_ssh
# enable & start ESXi Shell (TSM)
vim-cmd hostsvc/enable_esx_shell
vim-cmd hostsvc/start_esx_shell
reboot
```

The IP address, hostname, and NTP server were configured using Direct Console ESXi Interface accessed from Cisco UCS Manager KVM console

([http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc\\_50/GUID-26F3BC88-DA D8-43E7-9EA0-160054954506.html](http://pubs.vmware.com/vsphere-50/topic/com.vmware.vsphere.install.doc_50/GUID-26F3BC88-DA D8-43E7-9EA0-160054954506.html)).

## 6.7.2 Install and Configure vCenter

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

Vmware vCenter Server			
OS:	Windows 2008 R2	Service Pack:	
CPU:	4vCPUs	RAM:	16GB
Disk:	40GB	Network:	1x10Gbps

To support vCenter instance two node Microsoft SQL Server 2008 R2 cluster was created to host vCenter database. 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) and [http://msdn.microsoft.com/en-us/library/ms189134\(v=sql.105\).aspx](http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx))

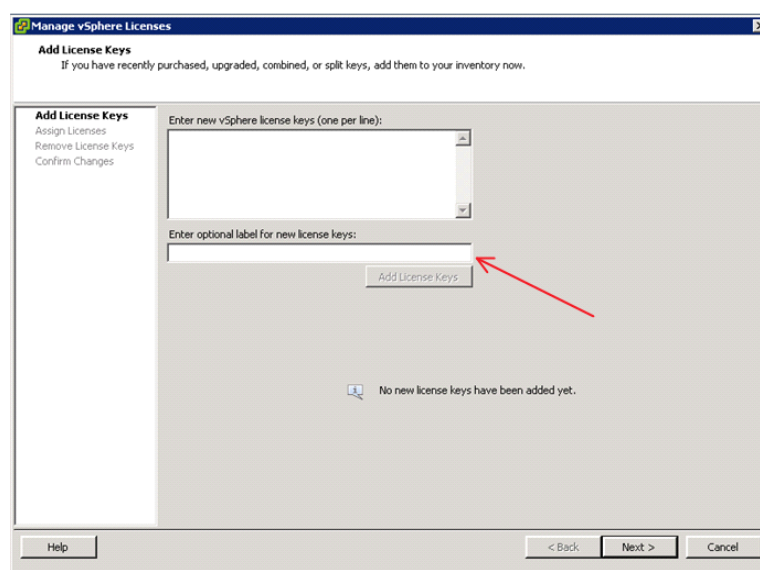
To install and configure vCenter, do the following:

1. Install the Microsoft® SQL Server® 2008 R2 Native Client for ODBC connections (<http://www.microsoft.com/en-us/download/details.aspx?id=16978> look for Native Client for your architecture)
2. Create a System DSN (control panel, administrative tools, Data Sources ODBC) and connect to your vCenter-SQL server. Note: Make sure to use FQDN's for everything.

3. Create Active Directory user account and call it vcenter. (This user account will be used for XD to connect to vCenter, you will have to follow a Citrix specific procedure and assign specific permissions on vCenter for XD to connect to vCenter <http://support.citrix.com/proddocs/topic/xendesktop-rho/cds-vmware-rho.html>).
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](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514)).

### 6.7.3 Install Licenses

1. Connect to vCenter using vSphere client.
2. Go to Home → Administration → Licensing.
3. Click Manage vSphereLicenses.
4. Add License keys for vCenter and Hosts.

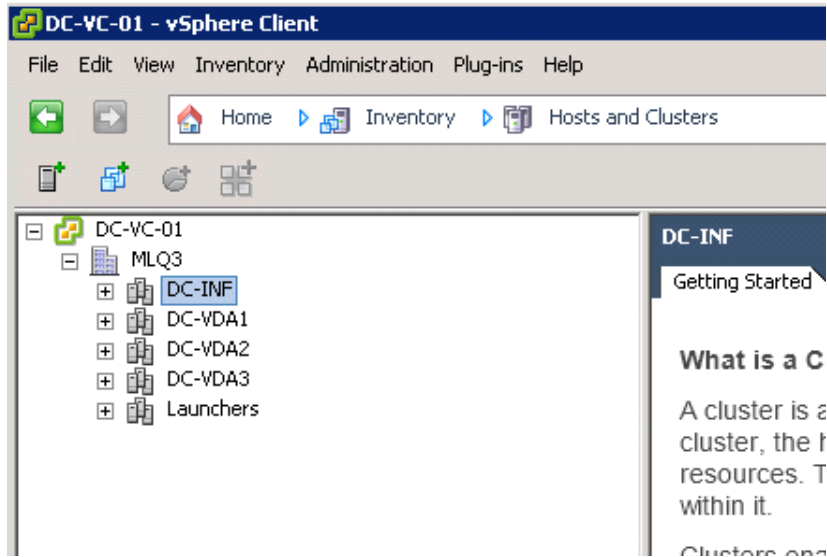


### 6.7.4 ESXi 5.0 Update 1 Cluster Configuration

To accommodate maximum recommendations for ESXi 5 clustering and to support the Nexus 1000V virtual switching best practices, we created five ESXi 5 clusters described below. For each of the three VDA clusters, 13 Cisco B230 M2 blade servers and approximately 1700 virtual machines were hosted. The 5 B200 M3s, 39 B230 M2s, and 20 B250 M2 ESX hosts were configured into 5 Clusters:

- DC-INF
- DC-VDA1
- DC-VDA2
- DC-VDA3

- Launchers



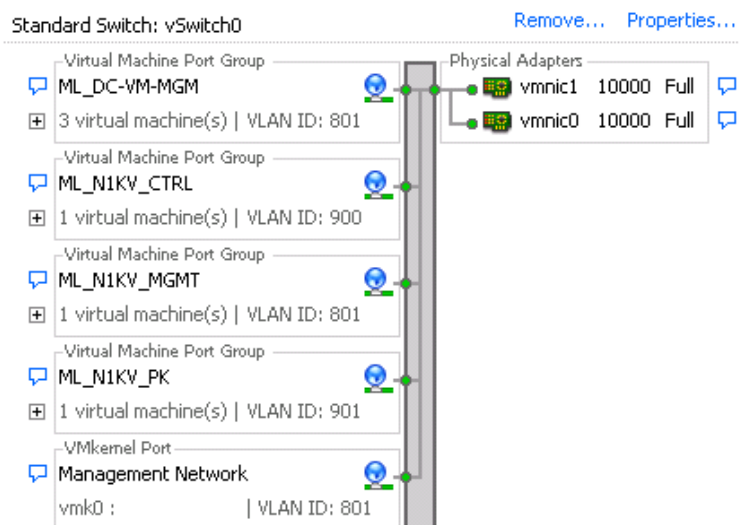
#### 6.7.4.1 DC-INF Infrastructure Cluster

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


Five physical Cisco UCS B200-M3 hosts were used in this cluster.

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


##### 1. Management



## 2. VDA



## 3. vMotion



## 4. Storage



### 6.7.4.2 Virtual Desktop Clusters

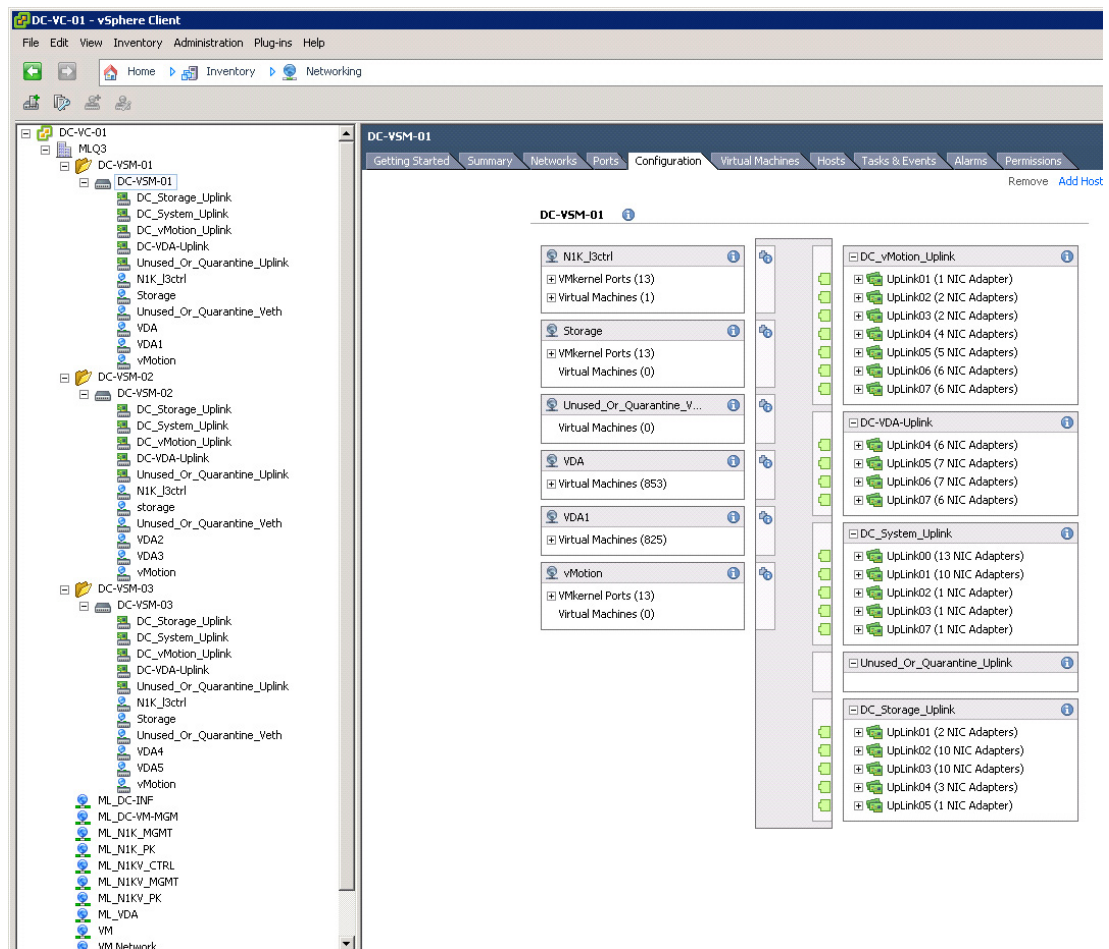
The following clusters were used to host 5004 desktops:

- DC-VDA1
- DC-VDA2
- DC-VDA3

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

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

The Nexus 1000V configuration is described in detail in Section 6.3.3 Nexus 1000V Configuration earlier in this document.



#### 6.7.4.3 Login VSI Launcher Cluster

The separate **Launchers** cluster was used to host Login Consultants' LoginVSI launchers, Launcher PVS, and a LoginVSI console.

It was hosted on a separate Cisco UCS Domain with dedicated storage.

The Launcher cluster was connected to the Data Center clusters through a Nexus 5000 layer 2 switch. Standard ESXi 5 vSwitches were used on the Launchers cluster.

## 6.8 Installing and configuring Citrix Provisioning Server 6.1

### 6.8.1 Pre-requisites

In most implementations, there is a single vDisk providing the standard image for multiple target devices. The more target devices using the same vDisk image, the less vDisks need to be created; making vDisk management easier. In order to have a single vDisk, all target devices must have certain similarities to make sure that the OS has all of the drivers it requires to run properly. The three key components that should be consistent are the motherboard, network card, or video card.



Disk storage management is very important because a Provisioning Server can have many vDisks stored on it, and each disk can be several gigabytes in size. Your streaming performance can be improved using a RAID array, SAN, or NAS.

Software and hardware requirements are available at <http://support.citrix.com/proddocs/topic/provisioning-61/pvs-install-task1-plan-6-0.html>.

#### Provisioning Server to Provisioning Server Communication

Each Provisioning Server must be configured to use the same ports (UDP) in order to communicate with each other (uses the Messaging Manager). At least five ports must exist in the port range selected. The port range is configured on the Stream Services dialog when the Configuration Wizard is run.



#### Note

If configuring for a high availability (HA), all Provisioning Servers selected as failover servers must reside within the same site. HA is not intended to cross between sites.

The first port in the default range is UDP 6890 and the last port is 6909.

#### Provisioning Servers to Target Device Communication

Each Provisioning Server must be configured to use the same ports (UDP) in order to communicate with target devices (uses the StreamProcess). The port range is configured using the Console's Network tab on the Server Properties dialog.

The default ports include:

UDP 6910 6930

#### Target Device to Provisioning Services Communication

Target devices communicate with Provisioning Services using the following ports:

UDP 6901, 6902, 6905



#### Note

Unlike Provisioning Servers to target device ports numbers, target device to Provisioning Services cannot be configured.

#### Login Server Communication

Each Provisioning Server that will be used as a login server must be configured on the Stream Servers Boot List dialog when the Configuration Wizard is run.

The default port for login servers to use is UDP 6910.

#### Console Communication

The Soap Server is used when accessing the Console. The ports (TCP) are configured on the Stream Services dialog when the Configuration Wizard is run.

The default ports are TCP 54321 and 54322 (Provisioning Services automatically sets a second port by incrementing the port number entered by 1; 54321 + 1).

If this value is modified, the following command must be run.

For PowerShell: MCLI-Run SetupConnection

For MCLI: MCLI Run SetupConnection



#### Note

Refer to the Provisioning Server Programmers Guides for details.

#### TFTP Communication

The TFTP port value is stored in the registry:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\BNFTFTP\Parameters Port

The TFTP port defaults to UDP 69.

#### TSB Communication

The TSB port value is stored in the registry:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\PVSTSB\Parameters Port

The TSB port defaults to UDP 6969.

#### Port Fast

Port Fast must be enabled.

#### Network Card

PXE 0.99j, PXE 2.1 or later.

#### Network Addressing

#### DHCP

## 6.8.2 Create a Highly Available CIFS Share for PVS vDisks

The following steps outline the process taken to create Highly Available CIFS share for vDisk hosting.



#### Note

A two-node Microsoft cluster was setup prior to this process. Procedure to set up the cluster described in detail in section “Setting Up a Two Node Citrix User Profile Server Cluster”.

1. Open Failover Cluster Manager.
2. Click Services and Applications node.
3. Select your File Server and click “**Add a shared folder**” in the Actions pane.
4. Click Browse and set the folder intended for the vDisk store, then click “**Next**”.
5. Leave the NTFS permission and click “**Next**”.
6. Validate SMB is checked and input Share name and then click “**Next**”.
7. Accept defaults and click “**Next**”.
8. Check Users and groups have custom share permission.
9. Click Permissions and set permissions to Everyone at Full Control; then click “**OK**”.
10. Click “**Next**”.
11. Accept Defaults then click “**Next**”.
12. Review summary and click “**Create**”.

## 6.8.3 Install Citrix Licensing and Licenses

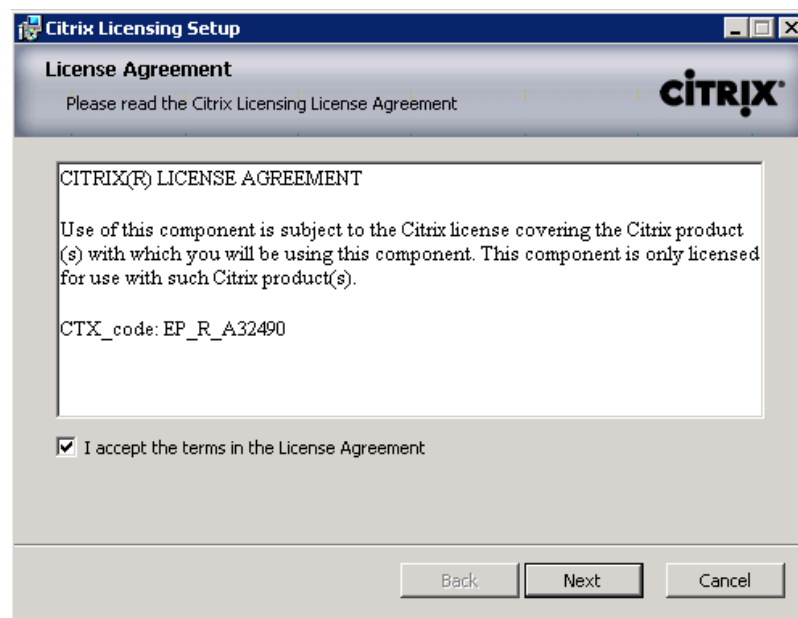
The steps that are outlined below describe the process of installing the Citrix License Server. Additionally, there are steps provided for configuring the license server as well as installing licenses.

### 6.8.3.1 Pre-requisites (Web Server)

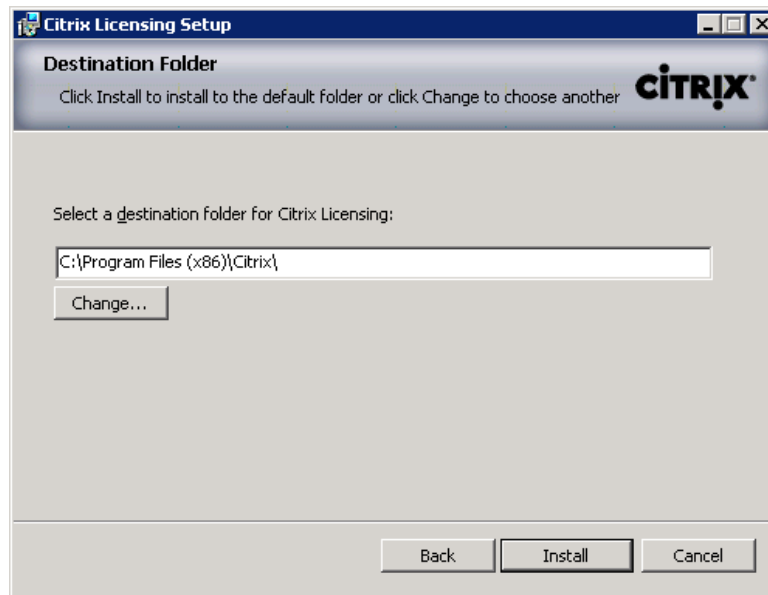
1. Open Computer Management.
2. Click on Add Role.
3. Select Web Server (IIS).
4. Click Next.
5. Click Next.
6. Under Application Development select ASP.NET.
7. Click Add Required Role Services.
8. Click Next.
9. Click Install.
10. Click Close to complete the installation process.

### 6.8.3.2 Install License Server

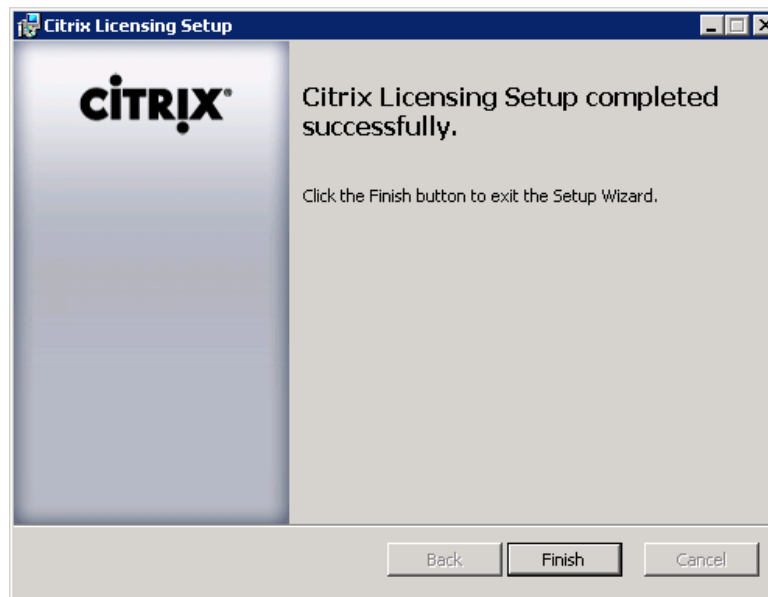
1. Locate and launch CTX\_Licensing.msi.
2. Select "I accept the terms in the License Agreement".



3. Click Next.
4. Accept the default destination folder.

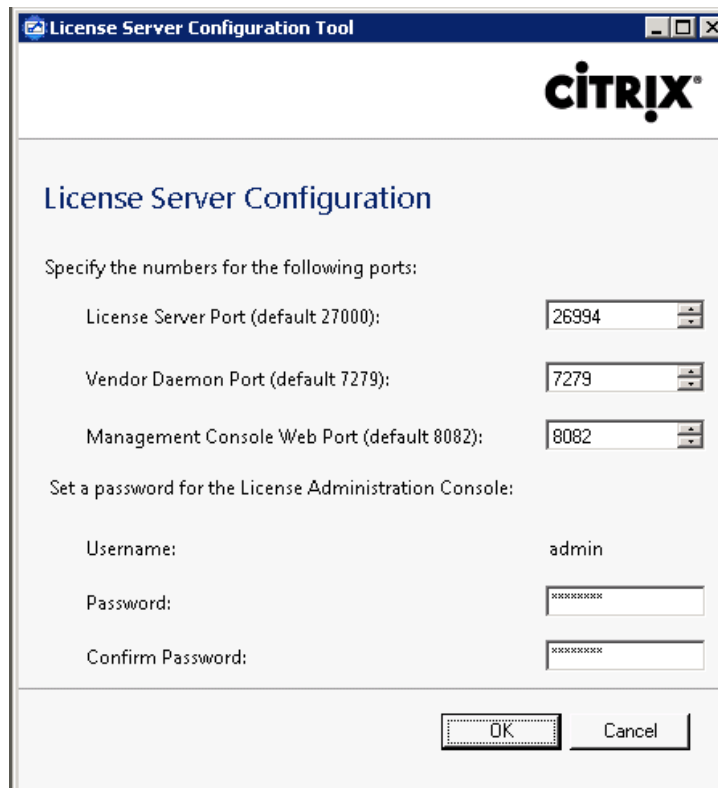


5. Click Install.
6. Click Finish to complete the installation process.



### 6.8.3.3 Configuring the License Server

1. Open the License Server Configuration Tool.
2. Accept the Default ports and provide the password for the Admin account.



**License Server Configuration Tool**

**CITRIX®**

### License Server Configuration

Specify the numbers for the following ports:

License Server Port (default 27000):

Vendor Daemon Port (default 7279):

Management Console Web Port (default 8082):

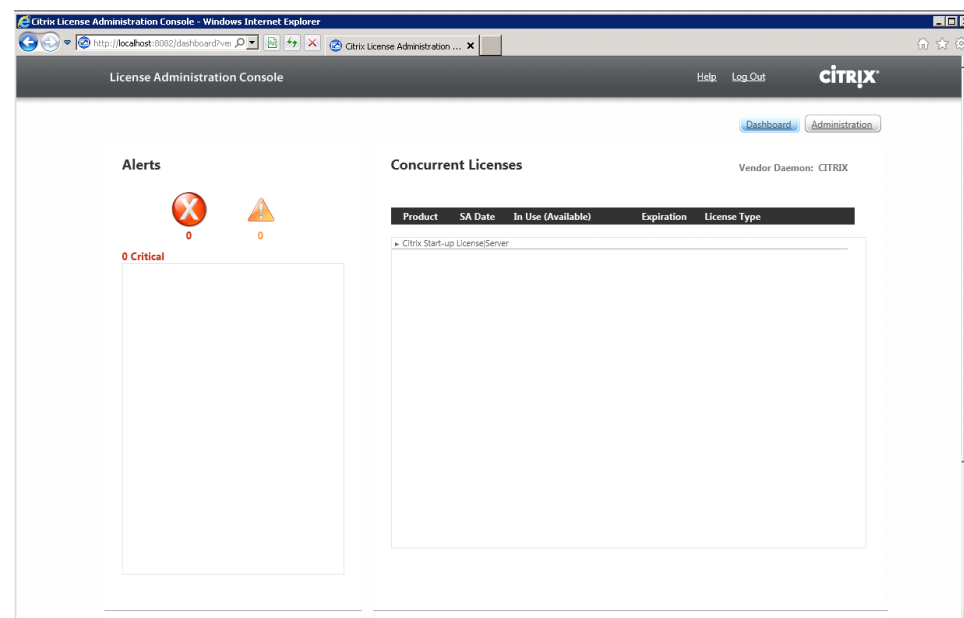
Set a password for the License Administration Console:

Username:

Password:

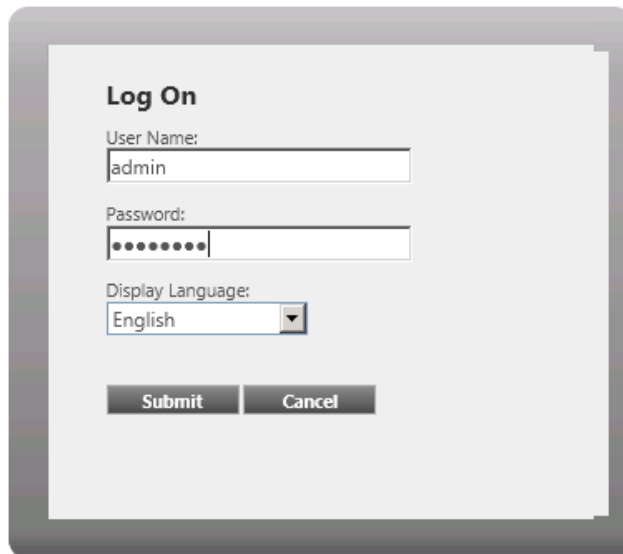
Confirm Password:

3. Click OK.
4. Go to Start | All Programs | Citrix | Management Consoles and click on License Administration Console.



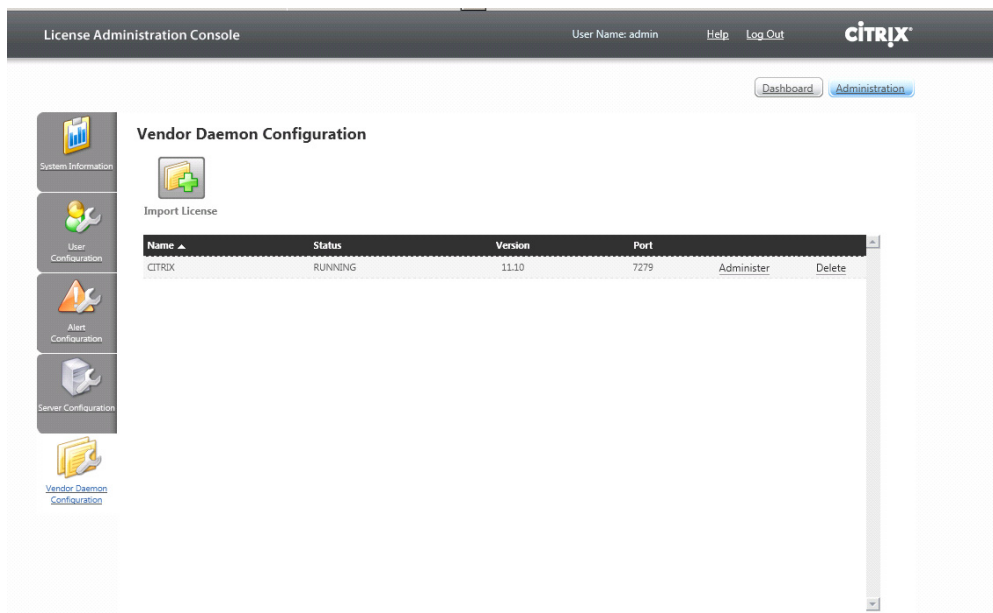
5. Click on the Administration button.

6. Enter the Admin credentials.



A 'Log On' dialog box with a light gray background and a dark gray border. It contains three input fields: 'User Name' with 'admin' entered, 'Password' with masked characters (dots), and 'Display Language' with a dropdown menu showing 'English'. At the bottom are two buttons: 'Submit' and 'Cancel'.

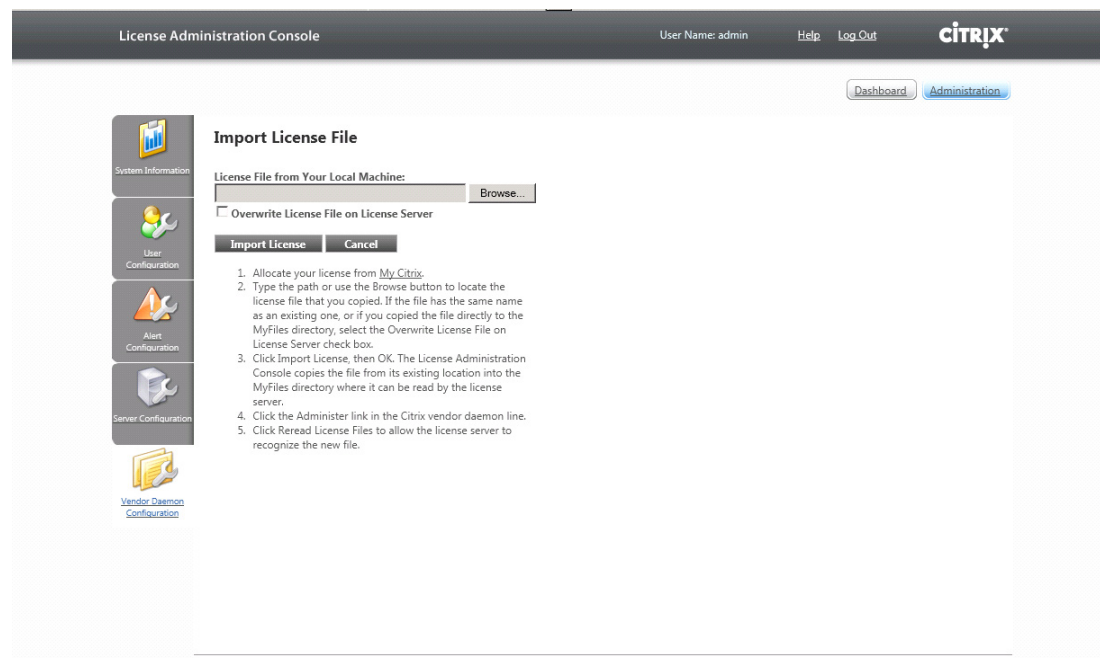
7. Click Submit.
8. Click on the Vendor Daemon Configuration tab on the left part of the screen.



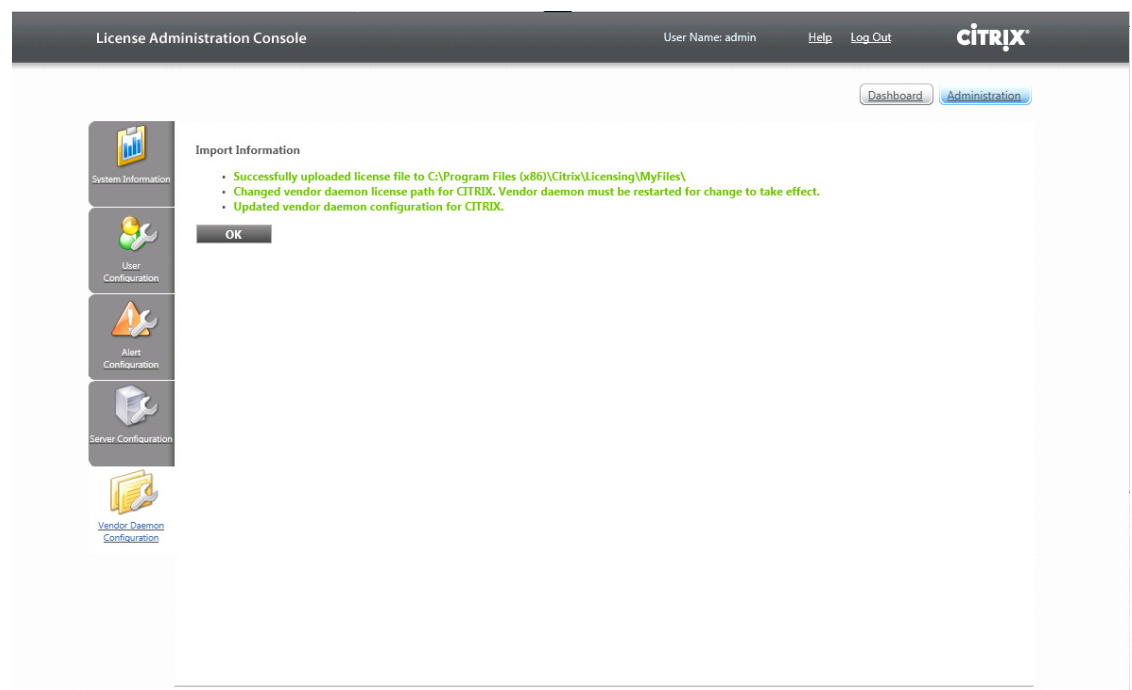
The screenshot shows the 'License Administration Console' interface. The top header bar includes 'License Administration Console', 'User Name: admin', 'Help', 'Log Out', and the Citrix logo. Below the header, there are two tabs: 'Dashboard' and 'Administration'. The left sidebar contains five icons with labels: 'Systems Information', 'User Configuration', 'Alert Configuration', 'Server Configuration', and 'Vendor Daemon Configuration'. The 'Vendor Daemon Configuration' tab is selected, showing an 'Import License' section with a table. The table has columns for Name, Status, Version, Port, and actions (Administer, Delete). The table contains one row for 'CITRIX' with status 'RUNNING', version '11.10', and port '7279'.

Name	Status	Version	Port	
CITRIX	RUNNING	11.10	7279	<a href="#">Administer</a> <a href="#">Delete</a>

9. Click on Import License.
10. Click Browse to locate the license file you are applying to the server.



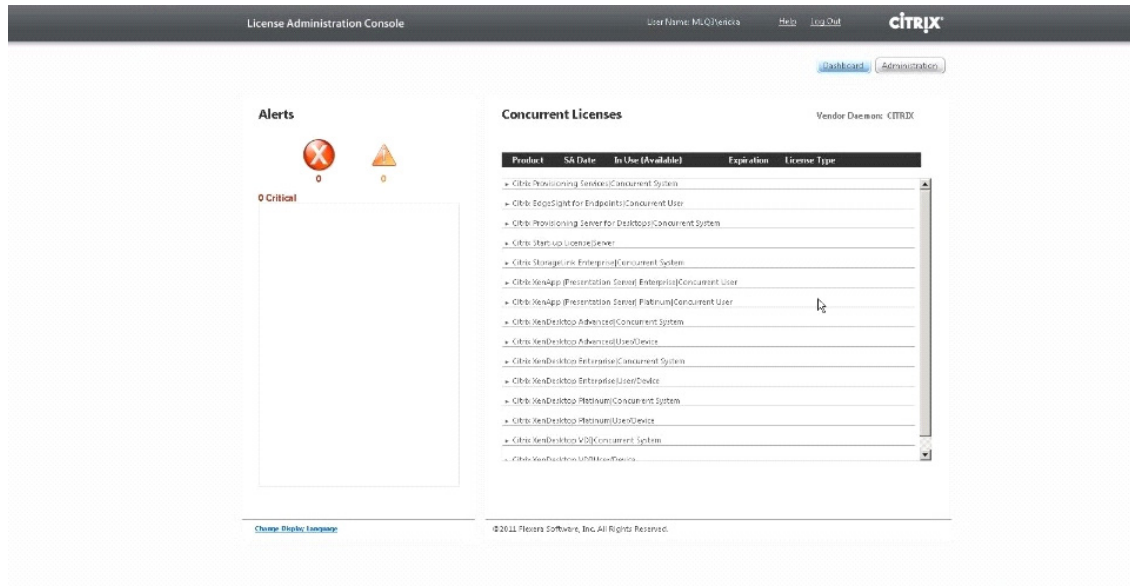
11. Select the file and Click Open.
12. Click on Import License.
13. Validate that the import was successful.



14. Click OK.
15. Click on the Dashboard button.



16. Validate that the necessary licenses have been installed.

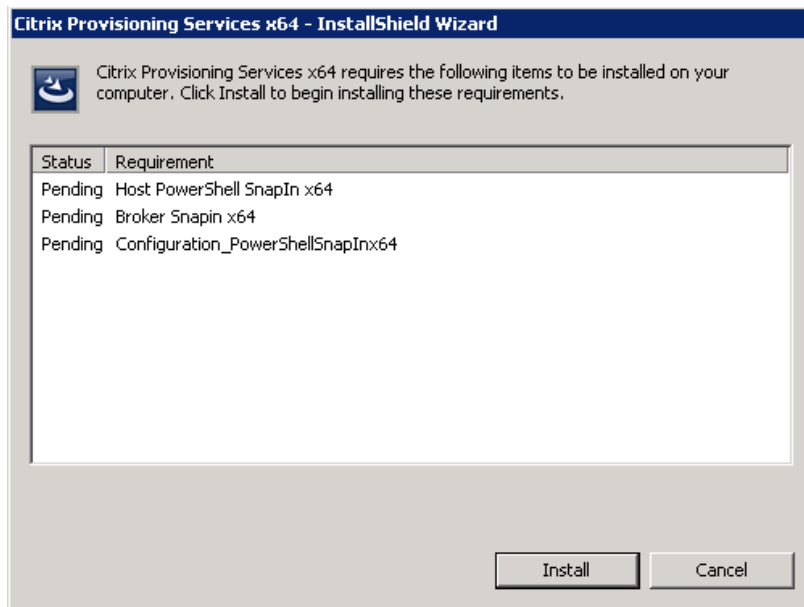


## 6.8.4 Install Provisioning Services 6.1

### 6.8.4.1 Base Install

The following steps outline the process taken to install Provisioning Services 6.1


1. Locate the PVS\_Server\_x64.exe and run the executable.
2. If prompted to installed required software click **Install**.



3. Select "I accept the terms in the License Agreement" and click **Install**.



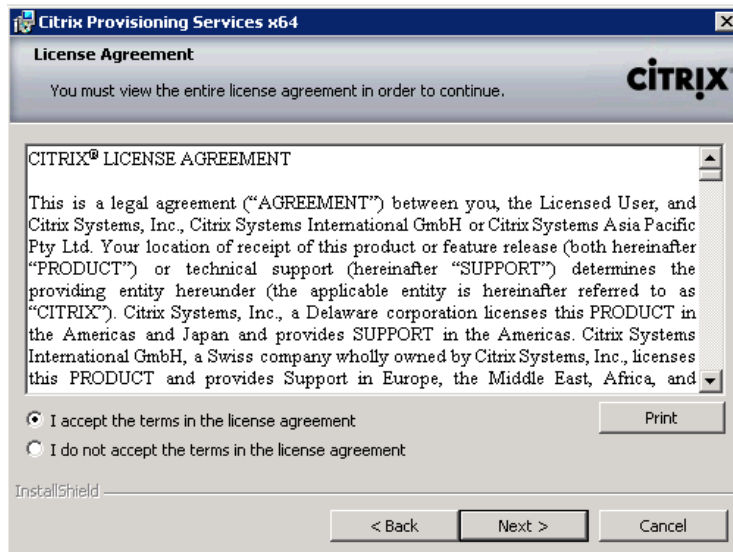
4. Click **Finish**.



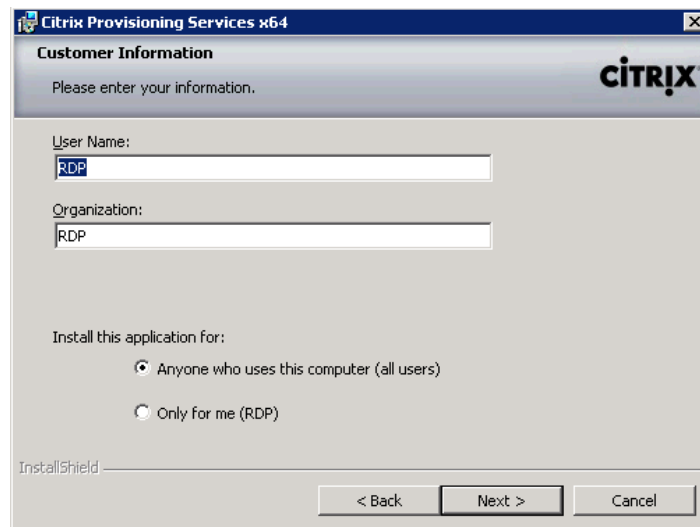
5. Click **Next**.




6. Select **"I accept the terms in the license agreement"**.



7. Enter the User name and Organization specific for your environment.
8. Select **Anyone who uses this computer (all users)**.



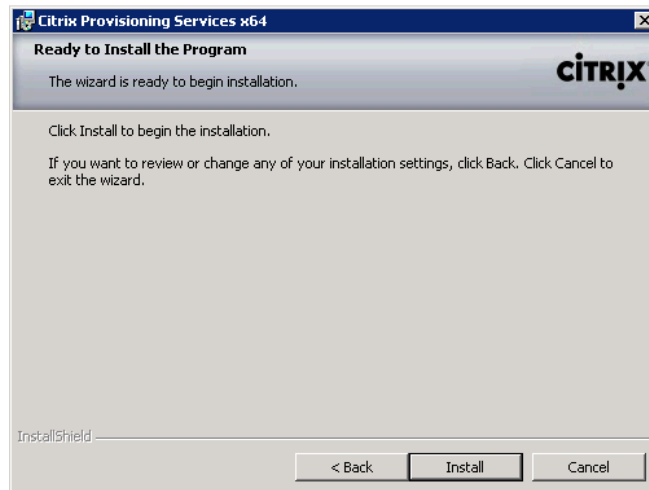
9. Click **Next**.
10. Leave the Destination Folder as the default location.



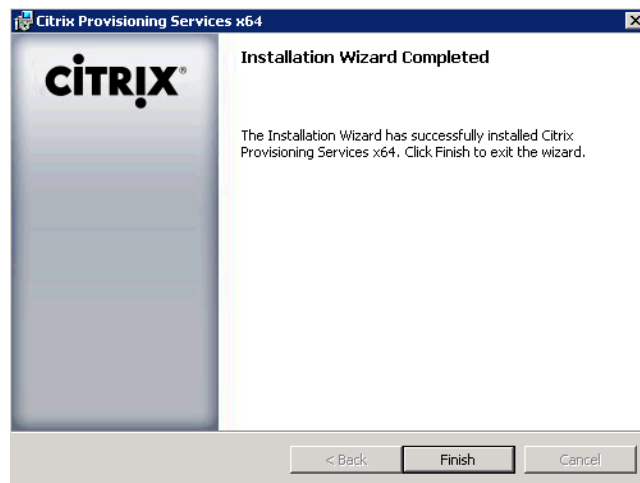
11. Click **Next**.
12. Select **Complete**.



13. Click **Next**.
14. Click **Install** to begin the PVS installation process.



15. Click **Finish** to complete the installation.



16. Click **Finish**.

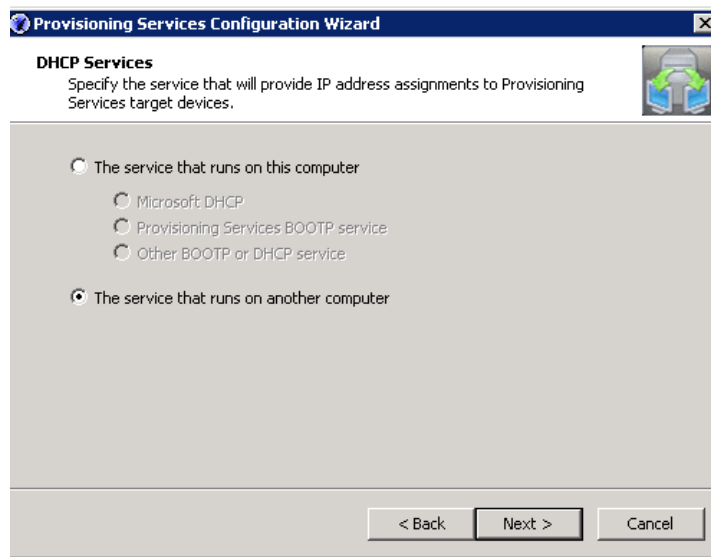
#### 6.8.4.2 Configure PVS Using the Provisioning Services Configuration Wizard

The steps that are identified below provide the steps taken to configure PVS using the Provisioning Services Configuration Wizard.

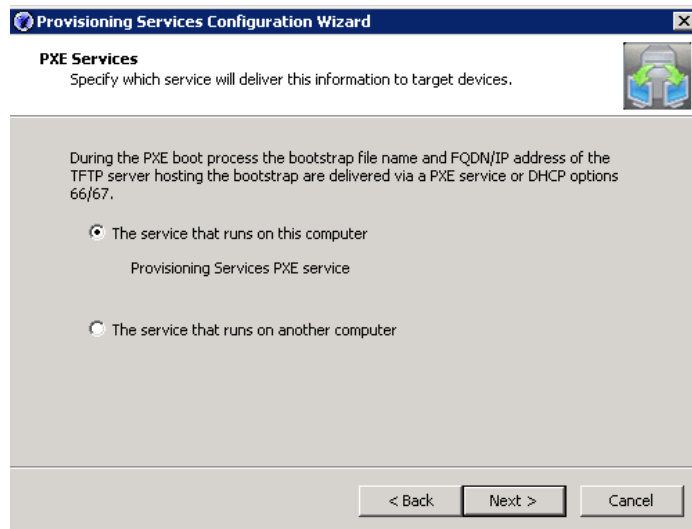
1. Start the PVS Configuration wizard.



2. Click **Next**.
3. In the DHCP services window, select **"The service that runs on another computer"**.

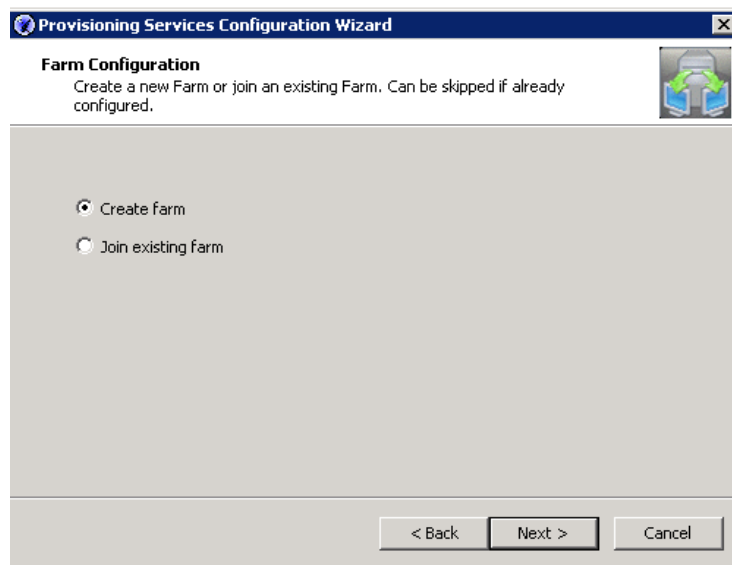


4. Click **“Next”**.
5. PXE services select **“the service that runs on this computer”**.

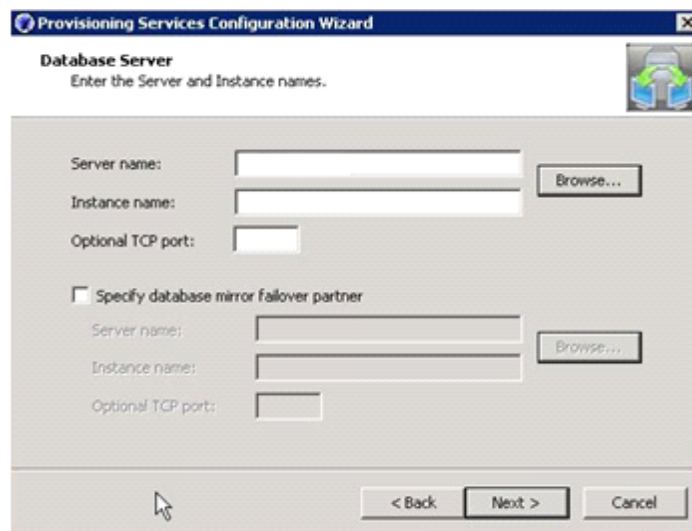


6. Click **“Next”**.
7. Select **Create Farm**.

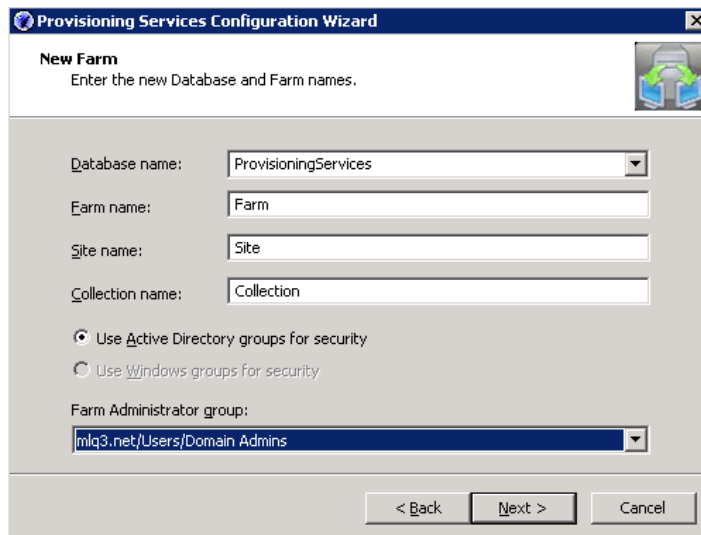




8. Collection Name. Additionally, choose the appropriate Active Directory group that will be identified as the Farm Administrators.
9. Click **Next**.



10. Configure Boot strap Boot list.
11. List the first four PVS servers in your farm.



**Provisioning Services Configuration Wizard**

**New Farm**  
Enter the new Database and Farm names.

Database name: ProvisioningServices

Farm name: Farm

Site name: Site

Collection name: Collection

☒ Use Active Directory groups for security  
☐ Use Windows groups for security

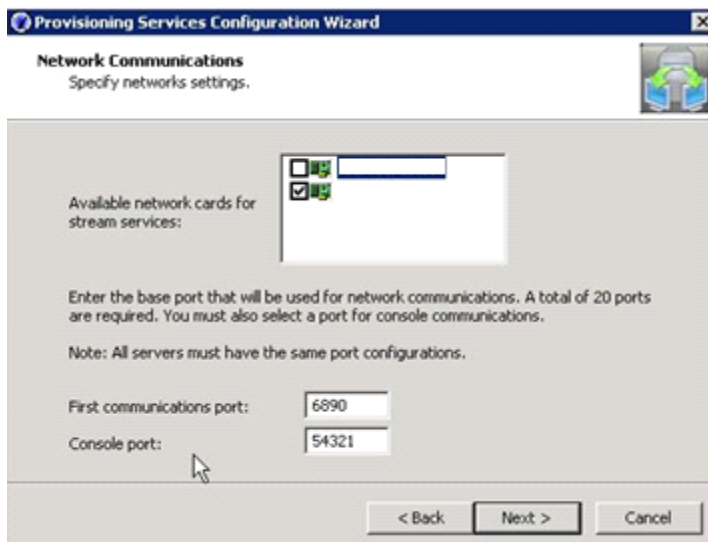
Farm Administrator group: mlq3.net/Users/Domain Admins

< Back Next > Cancel

12. Click **Next**.

13. Configure PVS streaming service NIC. Select your corresponding 10Gbps NIC.

14. Click **Next**.



**Provisioning Services Configuration Wizard**

**Network Communications**  
Specify network settings.

Available network cards for stream services:

Enter the base port that will be used for network communications. A total of 20 ports are required. You must also select a port for console communications.

Note: All servers must have the same port configurations.

First communications port: 6890

Console port: 54321

< Back Next > Cancel

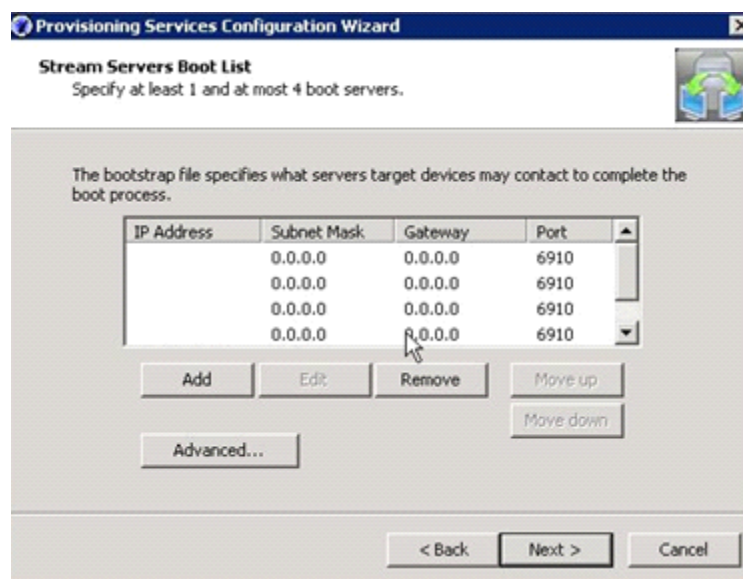
15. Select "Use the Provisioning Services TFTP service".

16. Click **Next**.

17. Configure Boot strap Boot list.

18. List the first four PVS servers in your farm.

19. Click **Next**.



20. Click **Finish**.

## 6.8.5 Install Required PVS Hotfixes

There are several recommended Hotfixes available for Provisioning Services 6.1. The Hotfixes that are listed below are specific to either the PVS server or the PVS target devices. The hotfixes applied to this environment were as follows:

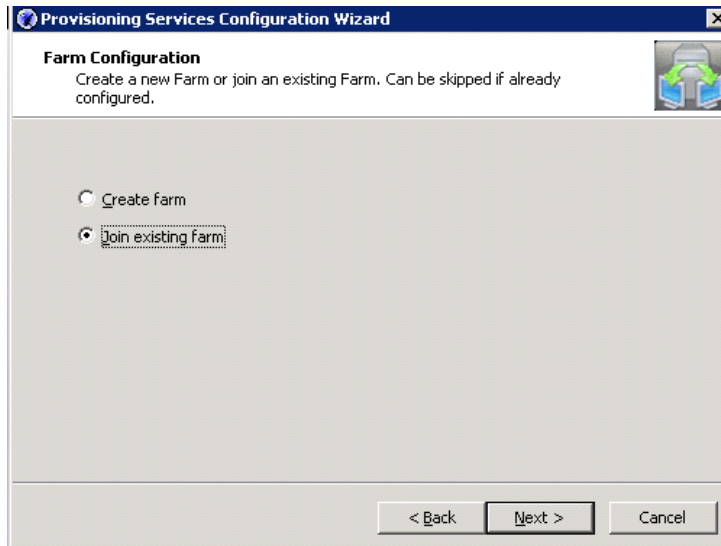
- CPVS61E001
- CPVS61E002
- CPVS61E003
- CPVS61E004
- CPVS61E005
- CPVS61E006
- CPVS61E007
- CPVS61E008
- CPVS61E009
- CPVS61E010
- CPVS61E011
- CPVS61E014
- CPVS61E015

Each of the listed hotfixes has unique installation steps. Please refer to the Reference section of this document to view the specific installation steps for each hotfix.

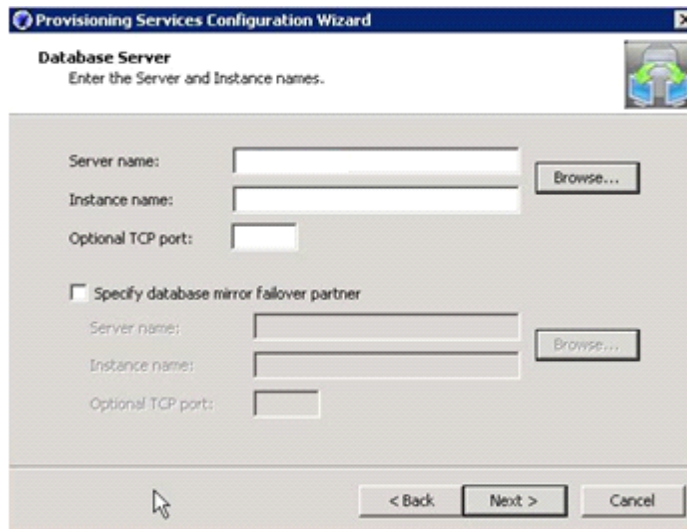
## 6.8.6 Adding PVS Servers to the Farm

After installing the PVS 6.1 software on additional servers, launch the Provisioning Services Configuration Wizard.

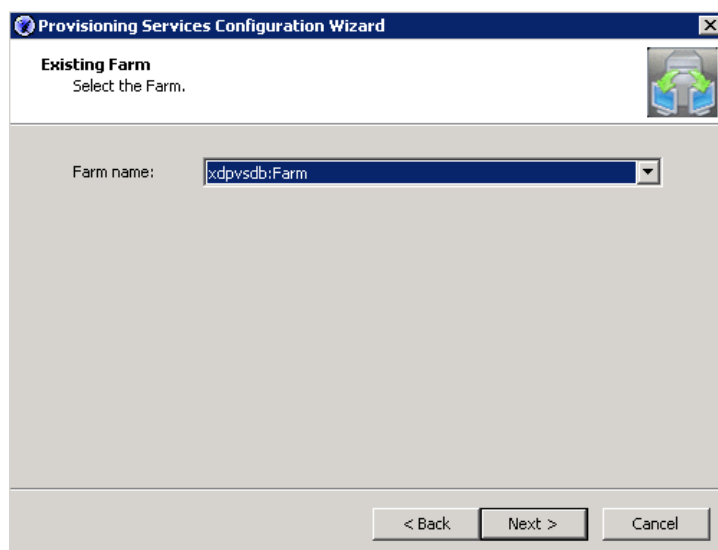
1. Select **Join existing Farm**.



2. Click **Next**.
3. Enter Database information.



4. Click **Next**.
5. Select the existing PVS Farm from the drop down box.



6. Click **Next**.
7. Select **Existing site Name**.
8. Click **Next**.
9. Select **Existing Store**.
10. Click **Next**.
11. Select **Network Service Account**.
12. Click **Next**.
13. Select your corresponding NIC that is configured for the farm PVS streaming service NIC.
14. Click **Next**.
15. Select **“Use the Provisioning Services TFTP service”**.

**Note**


---

This is only selected if the server is going to be a Login server.

---

16. Click **Next**.
  - Configure Boot strap Boot list.  
List the first four PVS servers in your farm. Click **Next**.
17. Click **Finish**.

## 6.9 Installing and Configuring Citrix XenDesktop 5.6 FP1

Four XenDesktop 5.6 Delivery Controllers were virtualized on VMware ESXi 5.0 hosted on Cisco B200 M3 infrastructure blades.

Beginning with XenDesktop 5, Citrix replaced the proprietary IMA protocol encrypted data store in favor of Microsoft SQL Server databases. Concurrently the concept of XenDesktop Farms (used in XenDesktop 4 and earlier) was eliminated in favor of the concept of Sites.

From a management standpoint, Citrix introduced two new management consoles beginning with XenDesktop 5.

- Desktop Studio
- Desktop Director

The Desktop Studio is the main administration console where hosts, machine catalogs, desktop groups and applications are created and managed. The Desktop Studio is where HDX policy is configured and applied to the site. The Desktop Studio is a Microsoft Management Console snap in and fully supports PowerShell.

## 6.9.1 Pre-Requisites

The following is a list of pre-requisites that are required with installing XenDesktop 5.6. They are as follows:

- One of the following operating systems:
  - Windows Server 2008, Standard or Enterprise Edition (32- or 64-bit), with Service Pack 2
  - Windows Server 2008 R2, Standard or Enterprise Edition (64-bit only)

Note that you can mix operating systems within a site.

- Microsoft .NET Framework 3.5 with Service Pack 1.
  - If you do not have this on your server, it is installed automatically for you. The XenDesktop installation media also contain this installer in the Support\DotNet35SP1 folder.
- Internet Information Services (IIS) and ASP.NET 2.0. IIS is required only if you are installing the Web Interface or Desktop Director:
  - For Windows Server 2008, IIS Version 7.0
  - For Windows Server 2008 R2, IIS Version 7.5
  - If you do not have these on your server, you may be prompted for the Windows Server installation media, and they are installed for you.
- Visual J# 2.0 Redistributable Package, Second Edition.
- This is required only if the Web Interface is installed on the server. If you do not have this on your server, it is installed automatically for you. The XenDesktop installation media also contain this installer in the Support\JSharp20SE folder.
- Visual C++ 2008 with Service Pack 1 Redistributable Package.
  - If you do not have this on your server, it is installed automatically for you. The XenDesktop installation media also contain this installer in the Support\vc\_redist\2008\_SP1 folder.
- Windows PowerShell version 2.0.
  - If you are using Windows Server 2008 (not Windows Server 2008 R2), Windows Management Framework is installed automatically if it is not already present on the server; it includes Windows PowerShell 2.0.



### Note

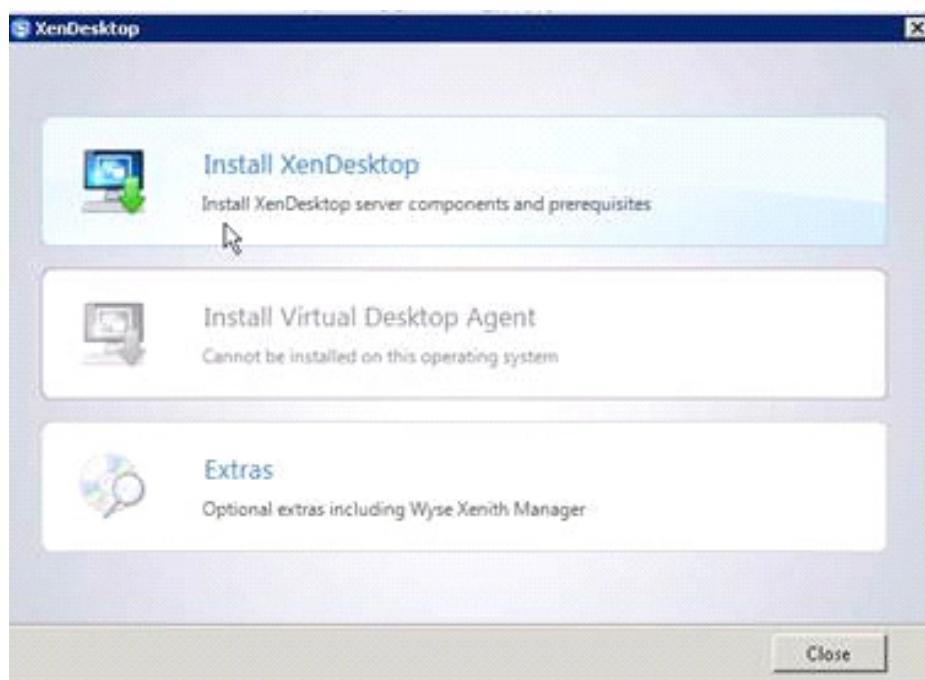
Windows Management Framework must be downloaded, so either make sure an Internet connection is available or pre-install Windows Management Framework.

- One of the following browsers if you are running the License Administration Console on the controller:
  - Internet Explorer 8 or 9
  - Firefox 3 to 8.x
  - Google Chrome
- Disk space requirements:
  - 100 MB for the Controller and SDKs
  - 50 MB for Desktop Studio
  - 50 MB for Desktop Director
  - 40 MB for Citrix Licensing
  - 100 MB for the Web Interface (and client software included in the installation)

## 6.9.2 Install XenDesktop, XenDesktop Studio, and Optional Components

The steps identified below show the process used when installing XenDesktop, XenDesktop Studio and optional components

1. Start the XenDesktop installation wizard.
2. Click **Install XenDesktop**.



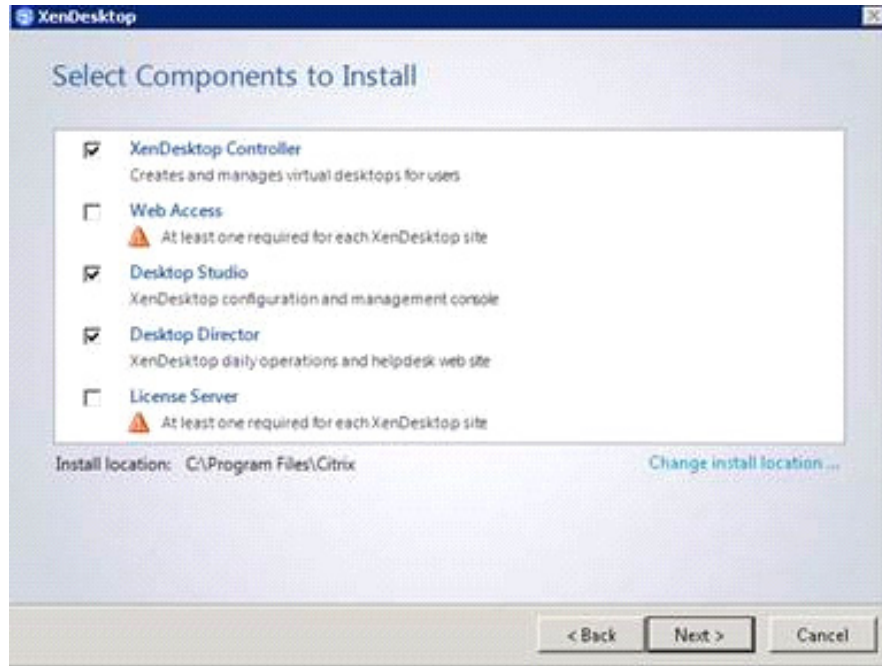
3. Select Components to install.
4. Verify that XenDesktop Controller, Desktop Studio and Desktop Directory are selected (License Server and Web Access components were not installed).





**Note** Desktop Director was installed on only the first XenDesktop Controller.

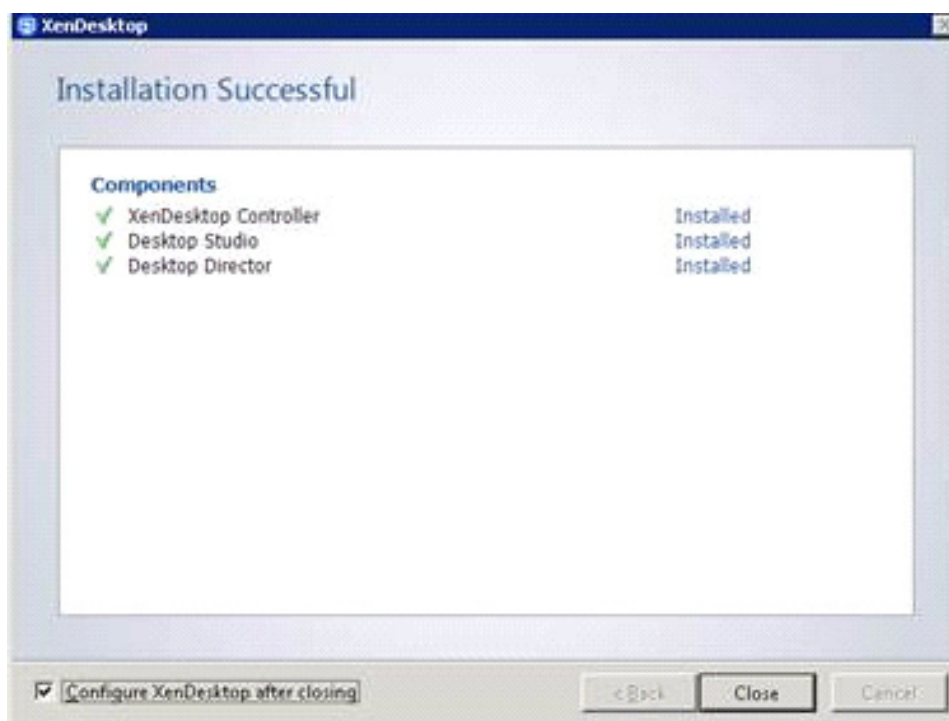
5. Verify that “**Install SQL Server Express**” is NOT selected.



6. Click **Next**.
7. Click **Install** in the Summary page to continue installation.
8. Click **Finish** to finalize your installation.

### 6.9.3 Create SQL Database for XenDesktop

1. After successfully installing XenDesktop, select the checkbox “**Configure XenDesktop after Closing**”.
2. Click **Close** on the Installation Successful dialogue.

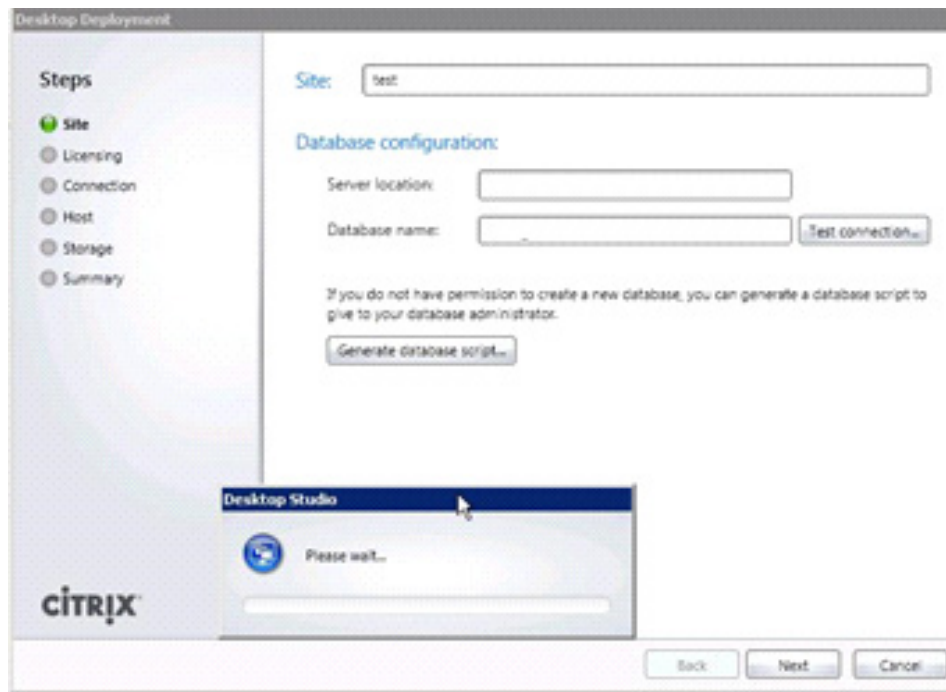


3. Open **Desktop Studio**: **Start > All Programs > Citrix > Desktop Studio**.
4. Click **XenDesktop deployment** in the XenDesktop Wizard.
5. Name the Site.
6. Database Configuration:
  - Enter the name of the SQL server installed earlier.
  - In the Database name field, enter the name of the database.
  - In this scenario, leave the prepopulated default Database name as it is, so that the wizard can create the database.

**Note**

To validate connectivity to the SQL Server, use the **Test Connection** button. If the database does not exist then an exception will be shown. However, connectivity to the SQL database validates successfully.

7. Click **“Next”**.



When your first XenDesktop server is installed and configured, you can repeat the procedure in Section 6.9.2 Install XenDesktop above to create load balancing and a fault tolerant XenDesktop environment.

### 6.9.4 Configure the XenDesktop Site Hosts and Storage

1. Enter licensing server name and select a license. Click **Next**.
2. Select **VMware Virtualization** from the Host Type dropdown.
3. Enter VCenter URL information.



**Note** Your vCenter server or appliance must have a trusted 3<sup>rd</sup> Party SSL certificate to use https.

4. Enter the vCenter administrator user name and password.
5. Click **Next**.

**Desktop Deployment**

**Steps**

- Site
- Licensing
- Connection**
- Host
- Storage
- Summary

**Host type:** VMware virtualization

**Address:** https://dc-vc-01.miq3.net/sdk

**Username:** vcenter

**Password:** .....

The Connection name will be displayed in Desktop Studio. Consider using a name that will help administrators to identify the host type and address of the deployment to which the connection relates.

**Connection name:** vcenter

**Virtual machines:**

- ☒ Use XenDesktop to create virtual machines
- ☐ Manually create virtual machines

**CITRIX**

Back Next Cancel

- Configure the hostname and select cluster and Network (port profile on VMware).

**Desktop Deployment**

**Steps**

- Site
- Licensing
- Connection
- Host**
- Storage
- Summary

**Host name**

FS1\_VDA

**Cluster**

Select a cluster for the new virtual machines.

DC-VDA1 Browse...

**Network**

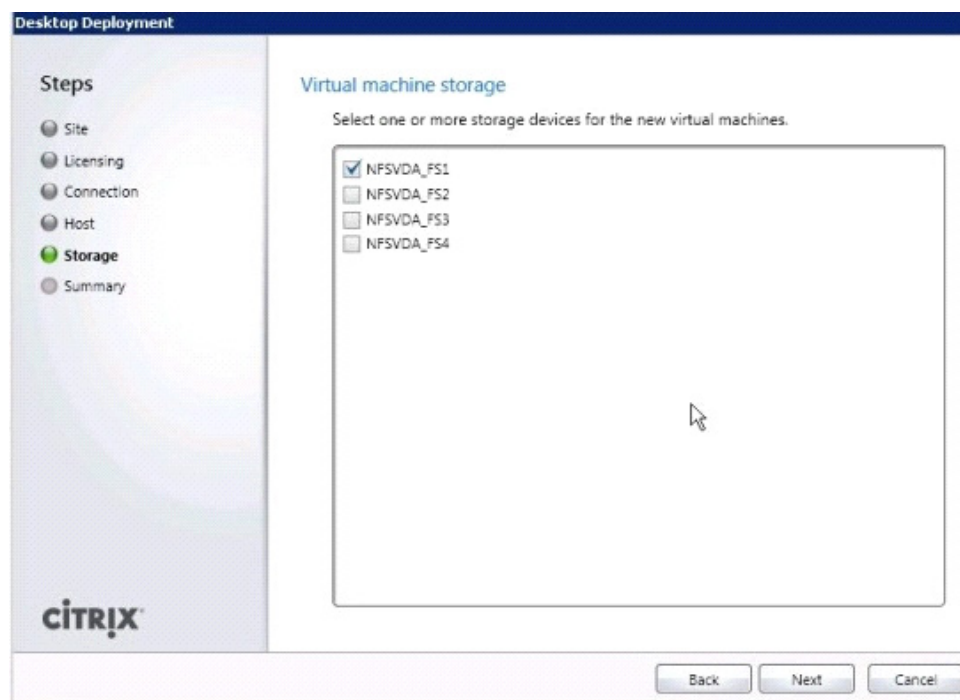
Select a network for the virtual machines to use.

- ☐ N1K3ctrl
- ☐ Storage
- ☐ UnusedOr\_Quarantine\_Veth
- ☒ VDA
- ☐ VDA1
- ☐ VM Network
- ☐ vMotion

**CITRIX**

Back Next Cancel

- Click **Next**.
- Select **storage** - this correlates to your vCenter's Datastore.




9. Click **Next**.
10. Click **Finish**. Create additional hosts per datastore & network.
  - Right click your existing VCenter Storage “Host” connection
  - Select **Add Storage**
  - Select **Use an Existing Host connection**

### 6.9.5 Configure XenDesktop HDX Policies

When testing with VSI a XenDesktop policy should be created to disable client printer mapping which is enabled by default. HDX policies configured and applied in Citrix Desktop Studio.

1. Open Desktop Studio.
2. Proceed to HDX Policy | Users.



3. Click **New** to start the policy creation process.

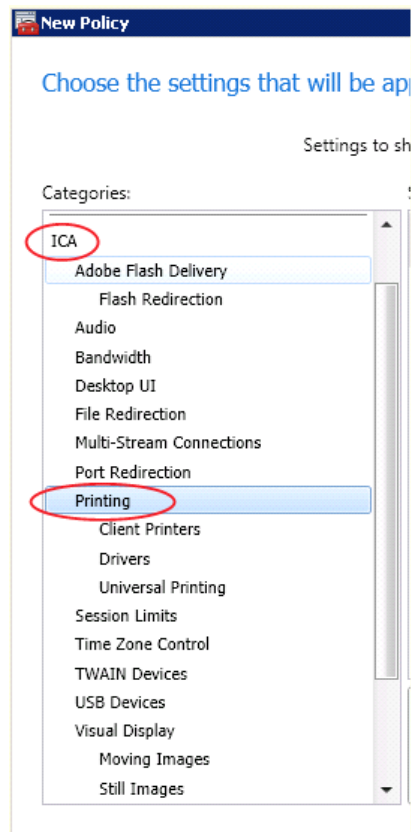


4. Enter a name for your policy and click **Next**.

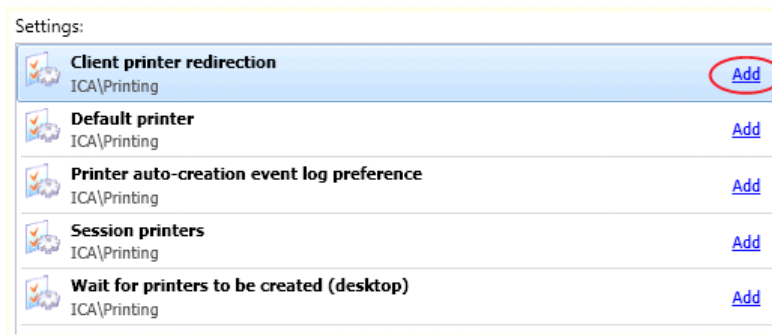
A screenshot of a 'New Policy' dialog box. The dialog has a title bar and a main area. The main area is divided into two sections. The top section is titled 'Enter basic policy information' and contains a 'Name:' label followed by a text input field (circled in red) and a 'Description:' label followed by a larger text area. The bottom section is titled 'Remaining Steps' and lists three steps: 'Settings - what will be applied by this policy', 'Filters - when to apply this policy', and 'Enable - whether to immediately enable this policy'. At the bottom of the dialog, there are four buttons: '< Back', 'Next >' (circled in red), 'Create', and 'Cancel'.

5. Select from Categories ICA | Printing.

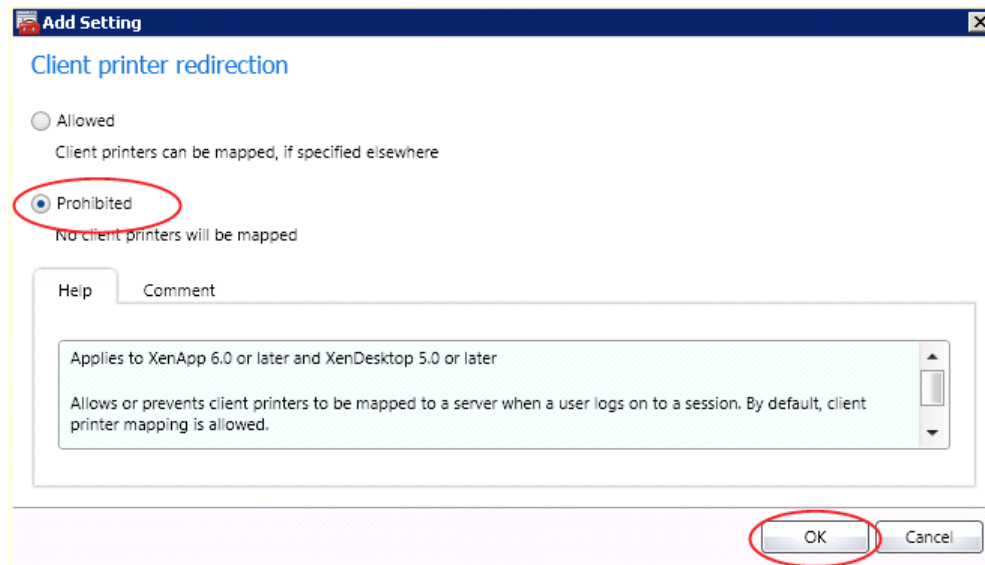




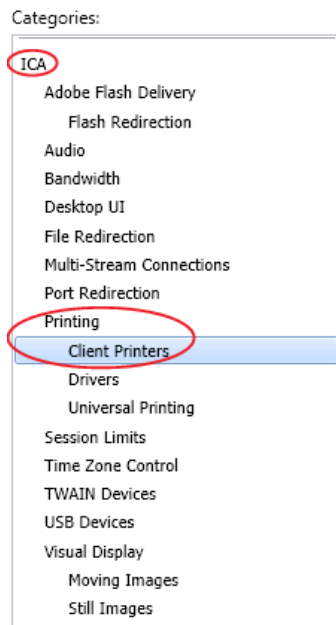
6. Select Client printer redirection and click **Add**.



7. Click **Prohibited** and then click **OK**.










8. Select from Categories ICA | Printing | Client Printers.

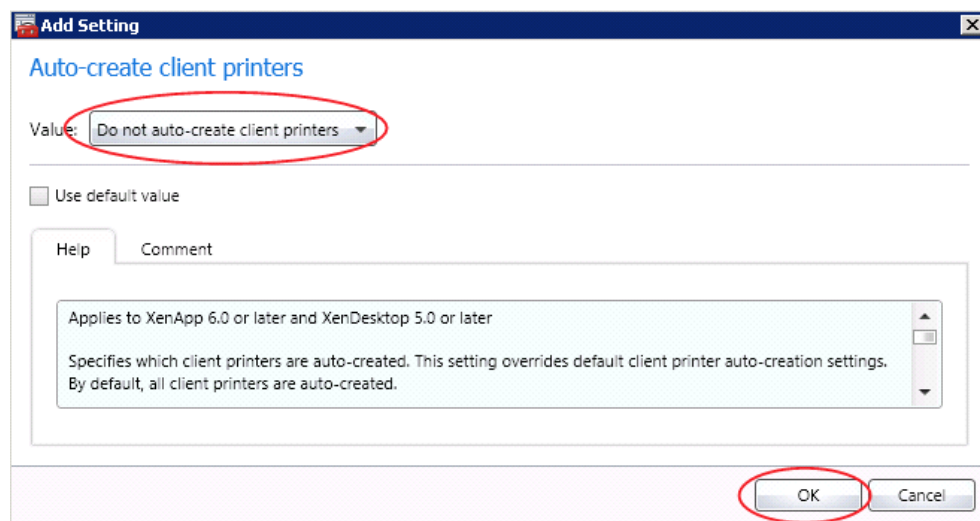


9. Select “Auto-create client printers” and click Add.

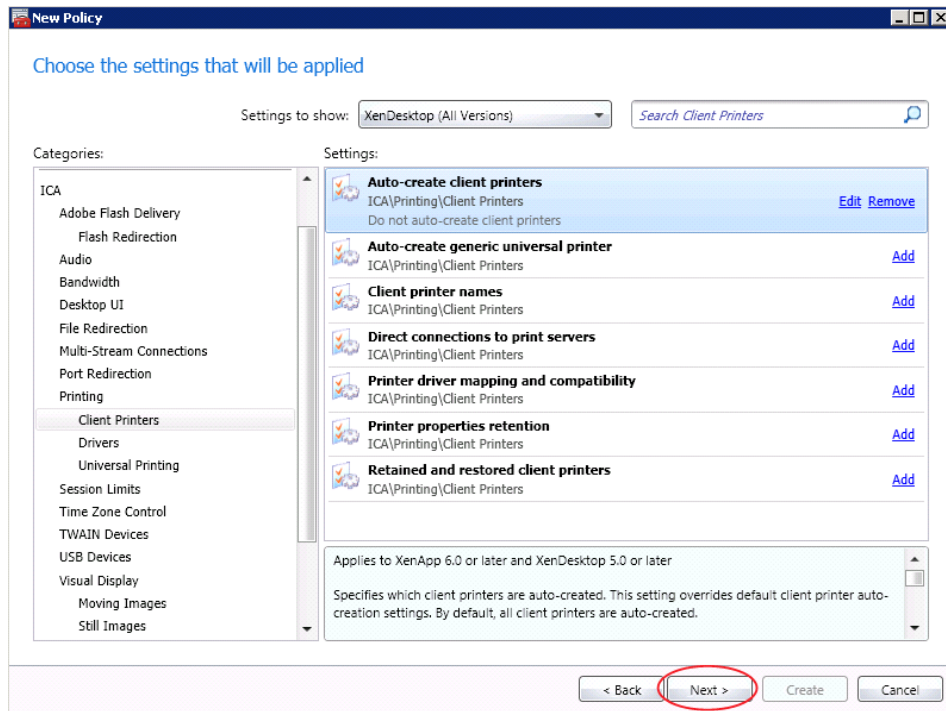
Settings:

	<b>Auto-create client printers</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Auto-create generic universal printer</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Client printer names</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Direct connections to print servers</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Printer driver mapping and compatibility</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Printer properties retention</b> ICA\Printing\Client Printers	<a href="#">Add</a>
	<b>Retained and restored client printers</b> ICA\Printing\Client Printers	<a href="#">Add</a>

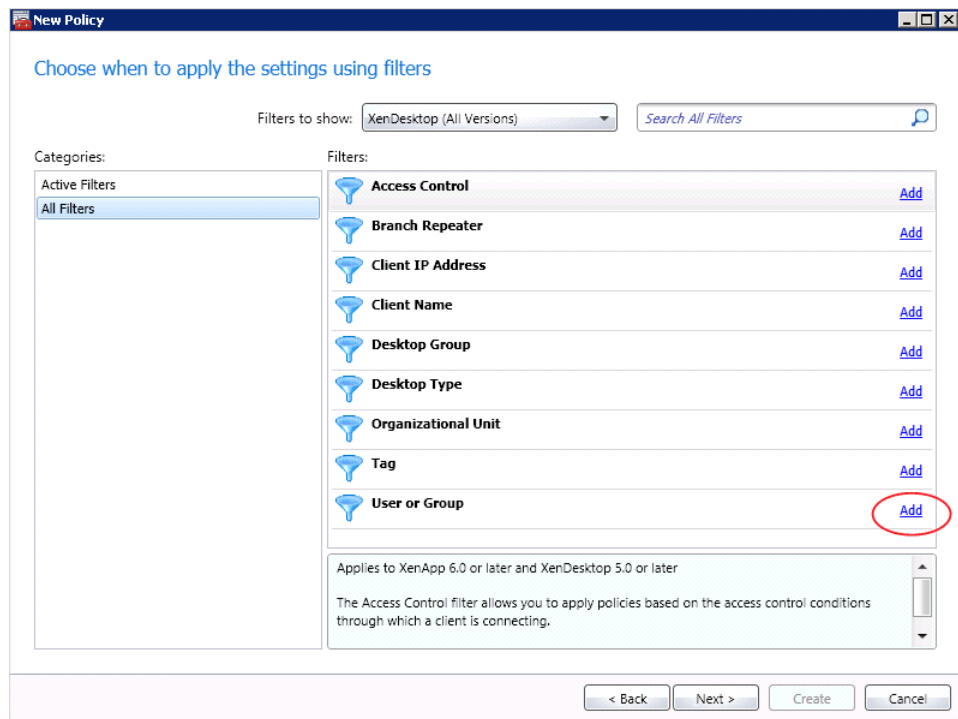
- From the drop-down list, select “Do not auto-create client printers” and click **OK**.



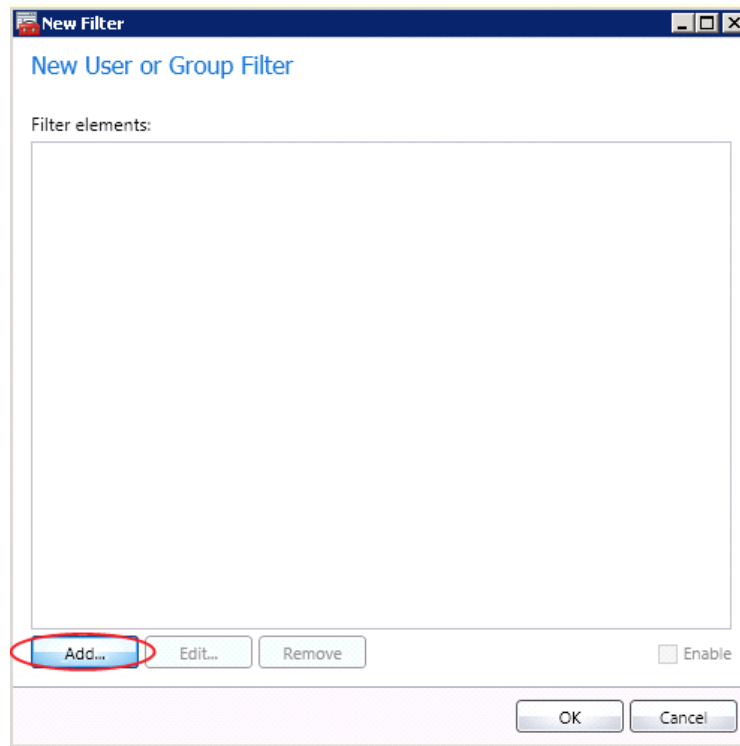
- Click **Next**.



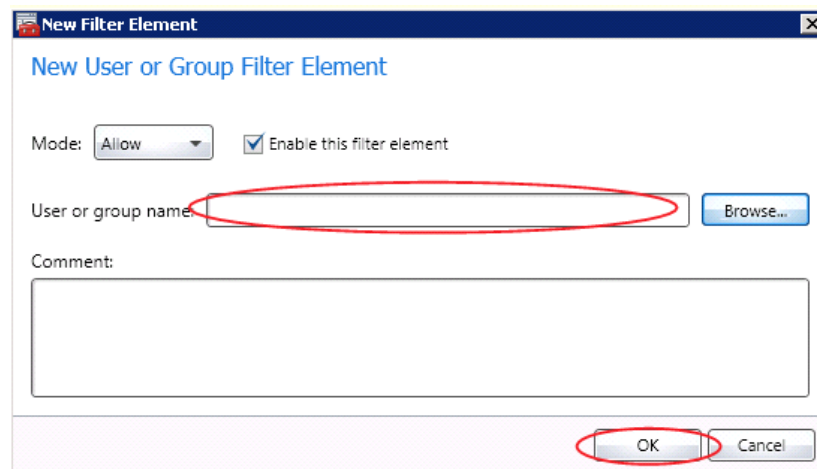
12. Select All Filters, “User or Group” and click Add Create policy filter.




13. Click Add.



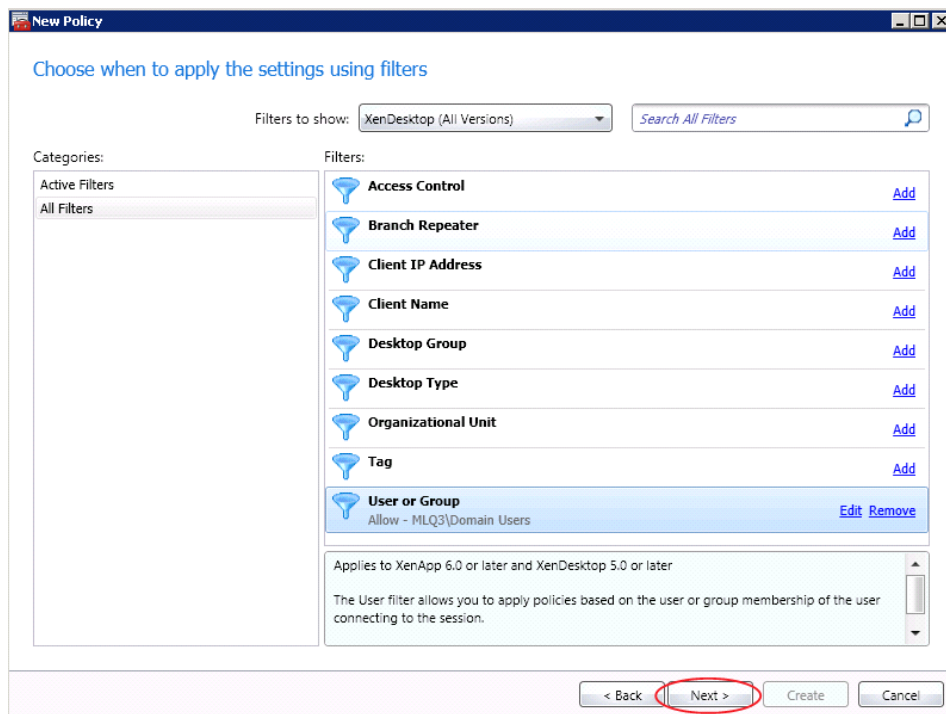
14. Pick the User Group you will be applying this policy.



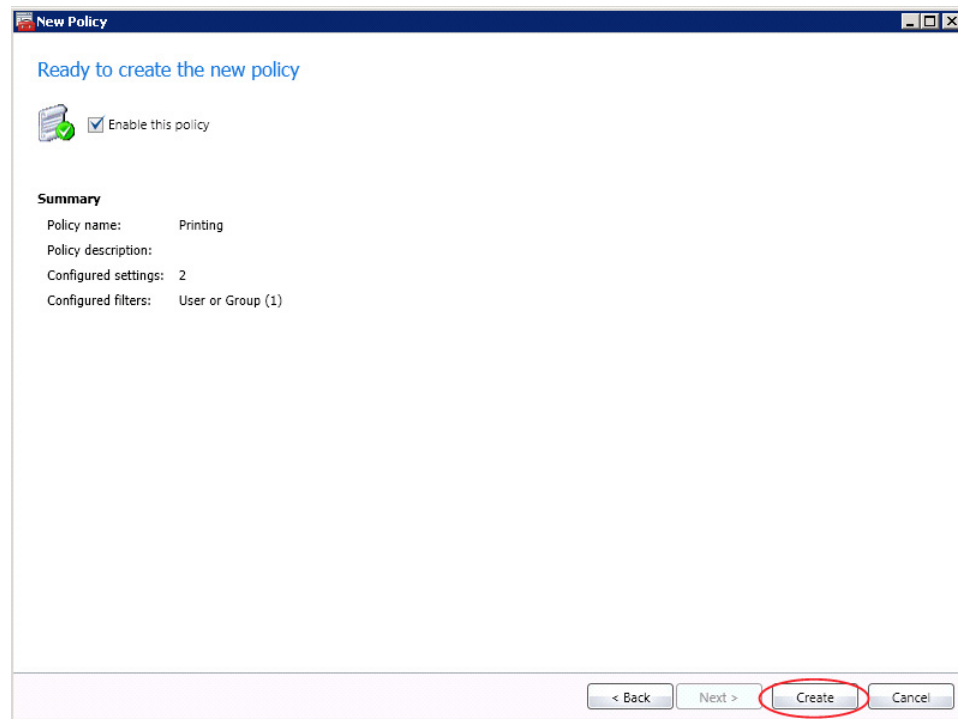
15. Click OK.



16. Click **Next**.




17. Click **Create**.



### 6.9.6 Configure the XenDesktop Desktop Group and Options

1. PVS XenDesktop Wizard is used for Catalog and VM creation.
2. From the XenDesktop Studio, click the **Assignments** node.
3. Click “**Create Desktop Group**”.



4. Select a catalog with available machines.
5. Add machines from the Catalog.



**Create Desktop Group**

**Steps**

- Catalog**
- Users
- Delegation
- Summary

**Select machines for Assignment:**

Catalog	Description	Available
DC-VDA1	4Datastores-2Portgroups	6
DC-VDA2	cluster2	0
DC-VDA3		0

**Unassigned machines**

Total available: 6

Add machines:

Specify the source and number of machines to be assigned

Back Next Cancel

6. Click **Next**.

7. Click **Add** and select a user or user group to use the Desktop Group.

**Create Desktop Group**

**Steps**

- Catalog
- Users**
- Delegation
- Summary

**Select users:**

Add... Remove

**Select Users or Groups**

Select this object type:  Object Types...

From this location:  Locations...

Enter the object names to select (examples):

Check Names


Advanced... OK Cancel

Select users/groups that are permitted to use the machines.

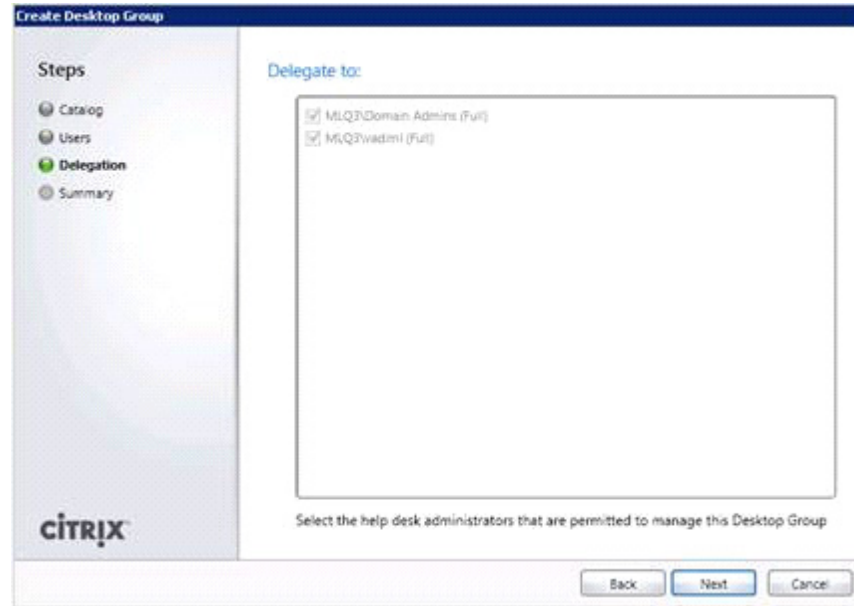
Desktops per user:

Back Next Cancel

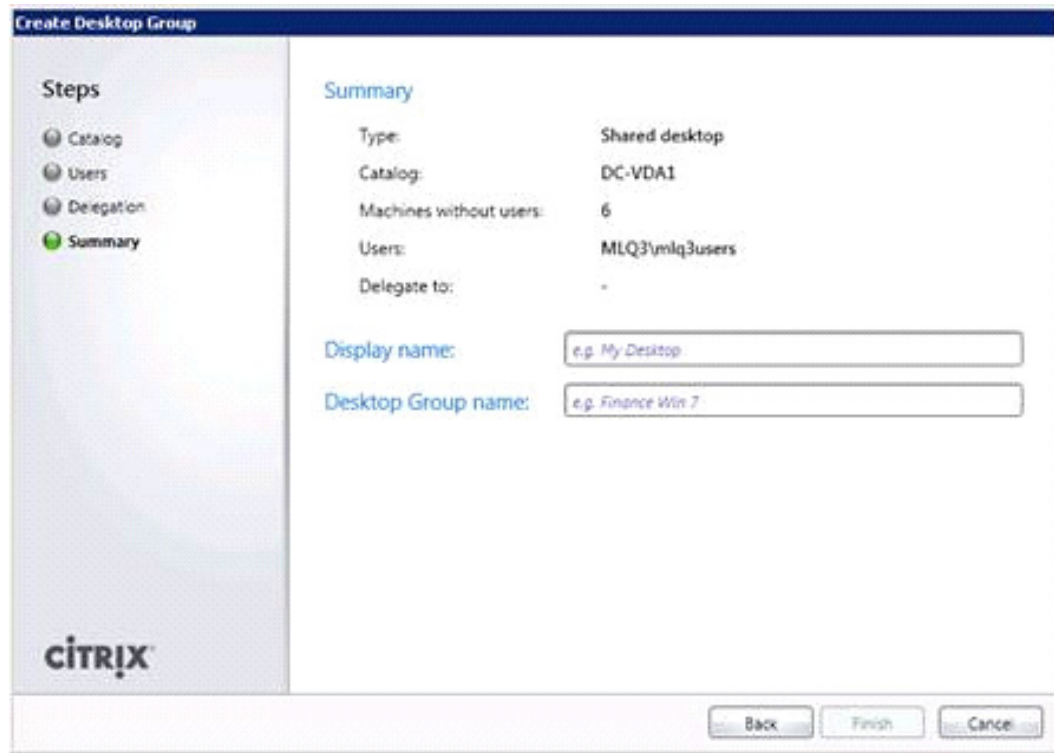
8. Select 1 Desktop Per user.



9. Click **Next**.
10. Select Administrators delegated to manage the Desktop Group.



11. Click **Next**.
12. Enter a Desktop Group Display Name and Description.



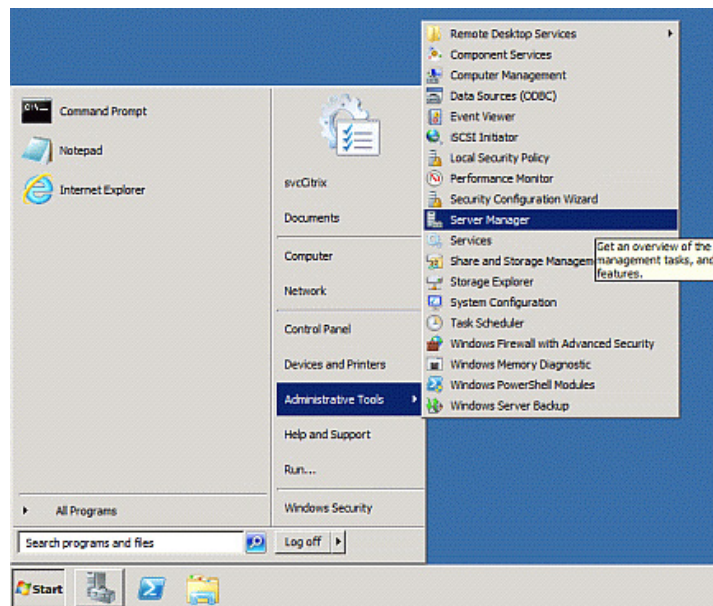
13. Click **Finish**.

## 6.9.7 Installing and Configuring Citrix Web Interface for XenDesktop

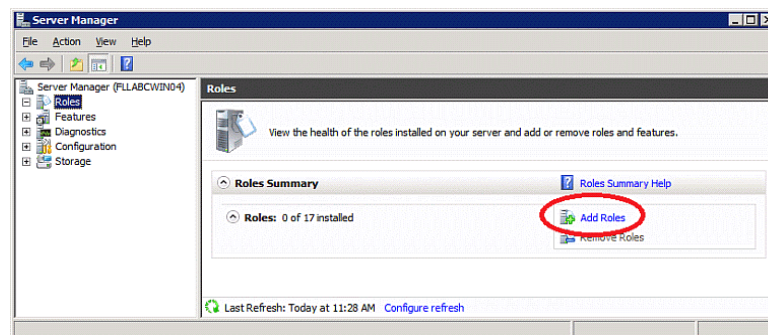
In XenDesktop Implementations, a key component is Web Interface.

### 6.9.7.1 Pre-requisites

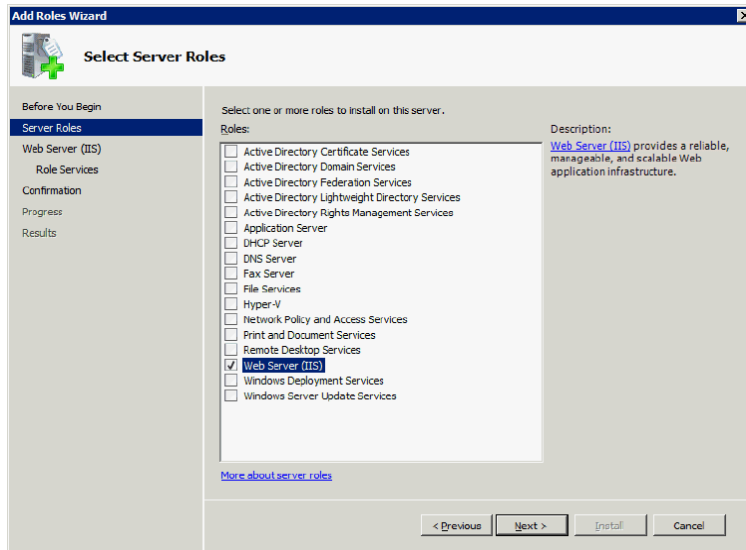
- Windows Server 2008 R2.
  - IIS 7
  - Microsoft Visual J# .NET 2.0.
  
  - Processor
    - Intel or AMD x86 or x64 compatible; 2 GHz minimum; 3 GHz preferred
  - Memory
    - Minimum of 2 GB RAM; 4 GB preferred
  - Hard Disk and Storage
    - Windows 2003 and 2003 x64; minimum of 250 MB on the application drive.
1. Launch **Server Manager** from the Start Menu. The path to Server Manager is **Start > Administrative Tools > Server Manager**.



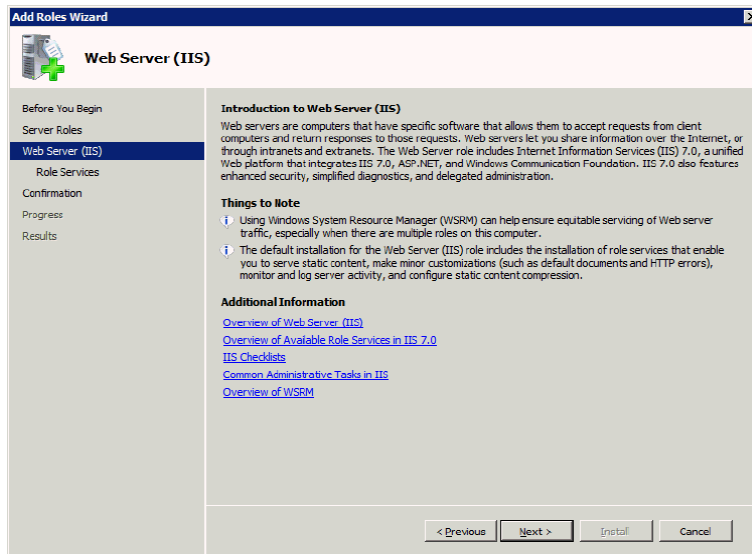
2. Select the **Roles** node in the left hand navigation pane.
3. Click **Add Roles**.



4. Agree to the Introductory Screen for the Add Roles Wizard by clicking **Next**.
5. Select **Web Server** role from the checkboxes available.



6. Agree to the screen that introduces IIS by selecting **Next**.



7. Make sure that all appropriate checkboxes are selected. Many of these are selected by default. Checkboxes include:

Web Server | Common HTTP Features:

- > Static Content
- > Default Document
- > Directory Browsing
- > HTTP Errors

Web Server | Application Development:

- > ASP.NET
- > .NET Extensibility
- > ISAPI Extensions

> ISAPI Filters

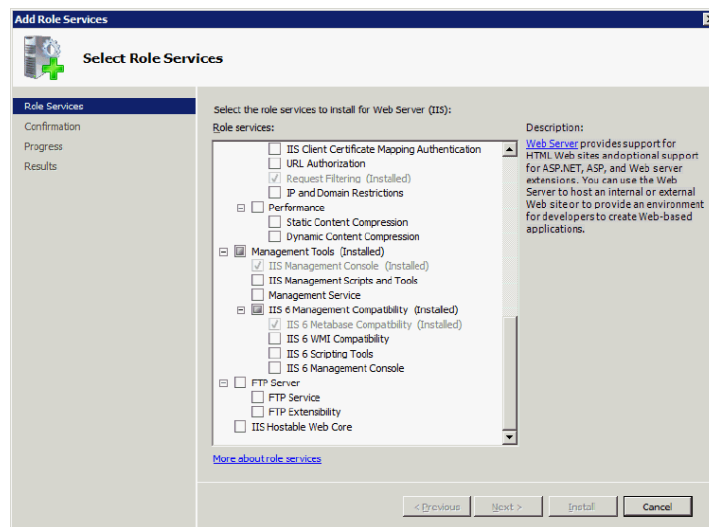
Web Server | Health and Diagnostics:

- > HTTP Logging
- > .Request Monitor

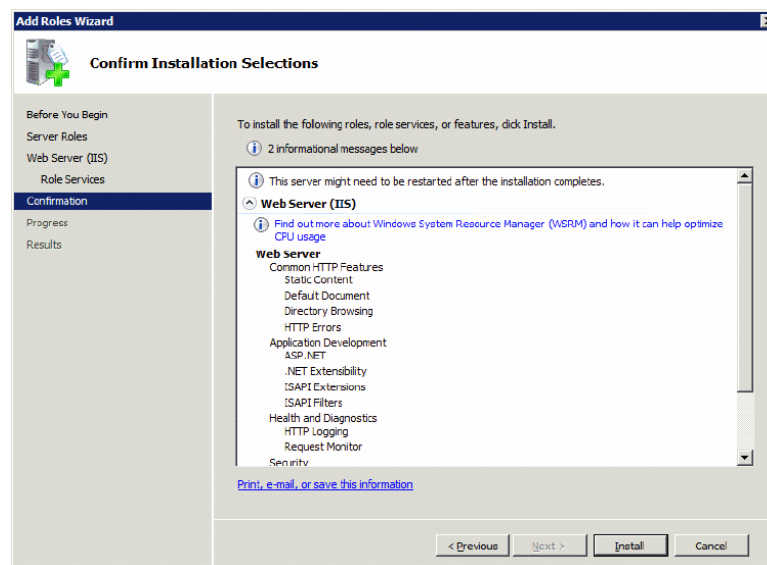
Web Server | Management Tools:

- > IIS Management Console

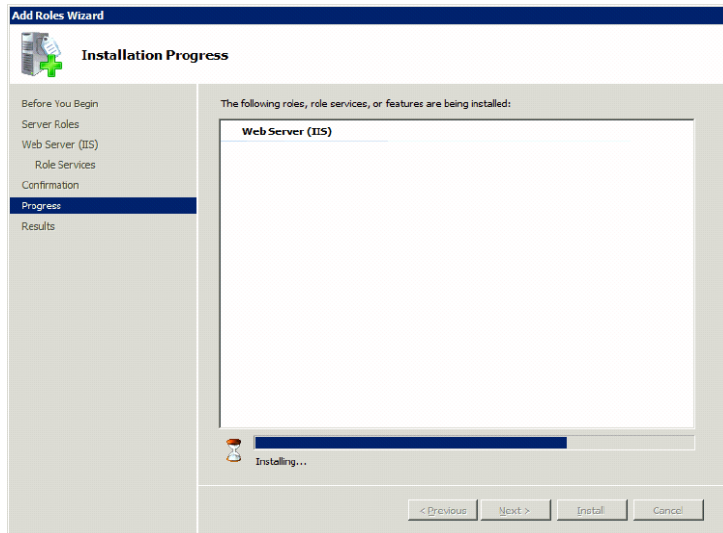
Web Server | Management Tools | IIS Management Compatibility  
| IIS 6 Metabase Compatibility



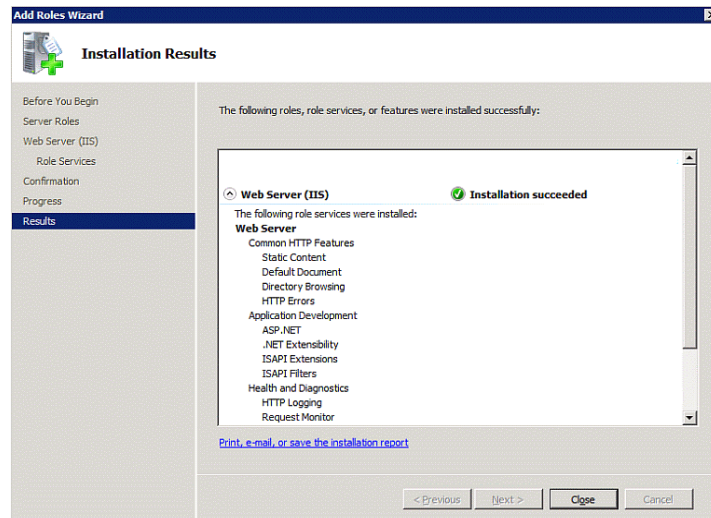
8. Confirm the installation selections by selecting **Install** at the prompt.



9. Allow the installation to proceed to completion.




10. Validate that the installation completed. If successful, click **Close** to complete the installation the IIS Server Role. Close out Server Manager and return to the system. It is now possible to proceed with the installation of secondary components required for Web Interface.

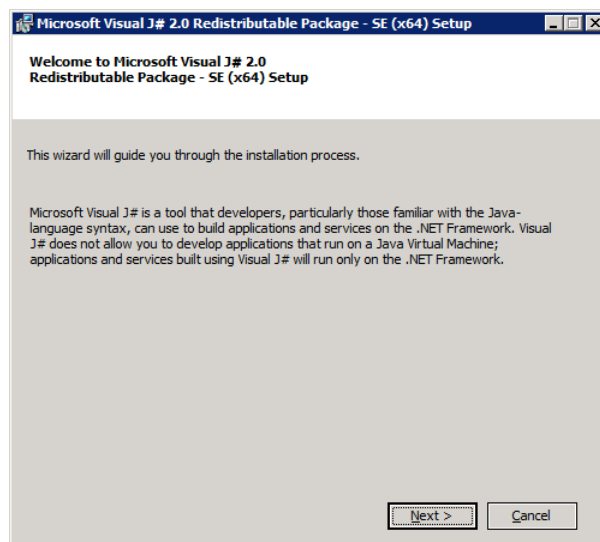


11. Navigate to the directory containing Microsoft Visual J# .NET 2.0. The installer is available in the Support folder of the XenDesktop 5 media.
12. Double-click the installer file. For 64-bit systems, the installer is vjredist64.exe.





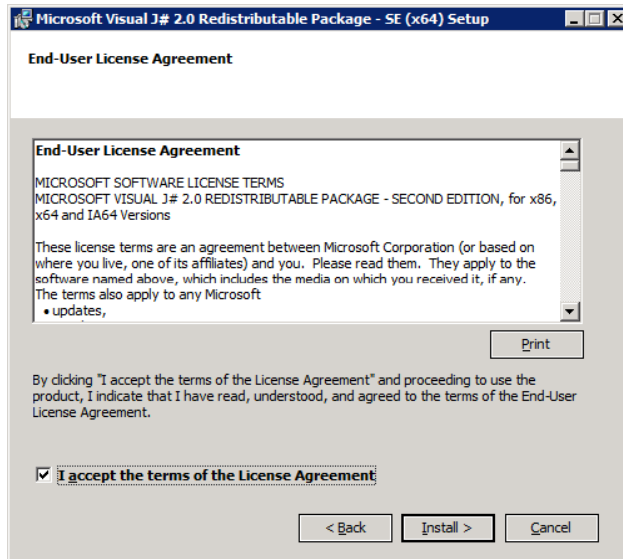
13. Read the overview and click **Next** to continue.



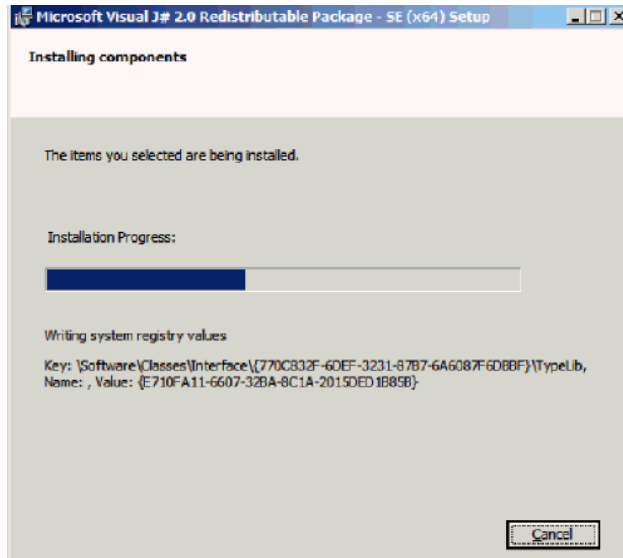
14. Read the entire End-User License Agreement.

15. Check the box marked **I accept the terms of the License Agreement**.


16. Click **Install** to proceed.



17. Allow the wizard to install files related to the software.



18. When the installer is completed, a Setup Complete summary screen appears.
19. Click **Finish** to complete the installation.




### 6.9.7.2 Installing Web Interface

1. Double-click the installer package to begin the setup process.

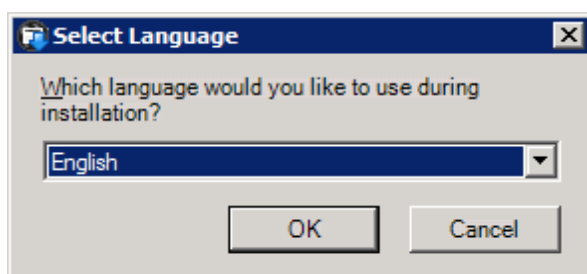


#### Note

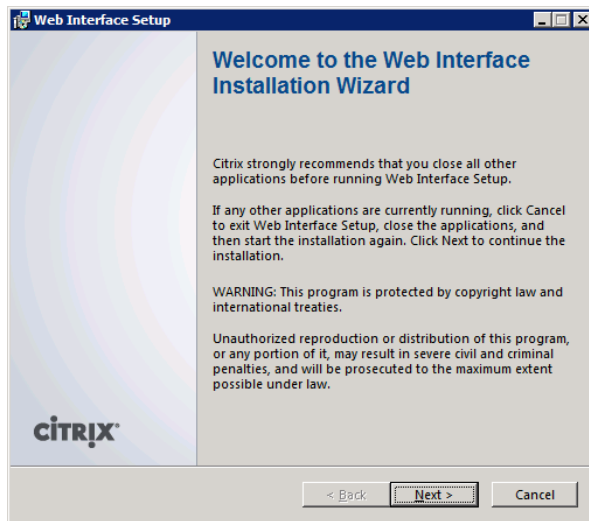
If using a service account to install Citrix Web Interface, make sure that the user account has elevated privileges. Consult with Active Directory Administrators to make sure compliance with specific enterprise account standards.



2. Select the language for the Web Interface installation. For the purposes of this scenario, select **English** and click **OK**.



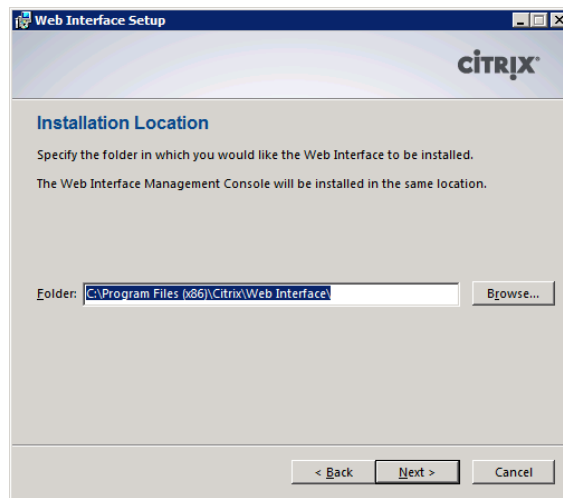
3. The installer presents a welcome screen. Read the entire welcome screen text and click **Next** to continue.



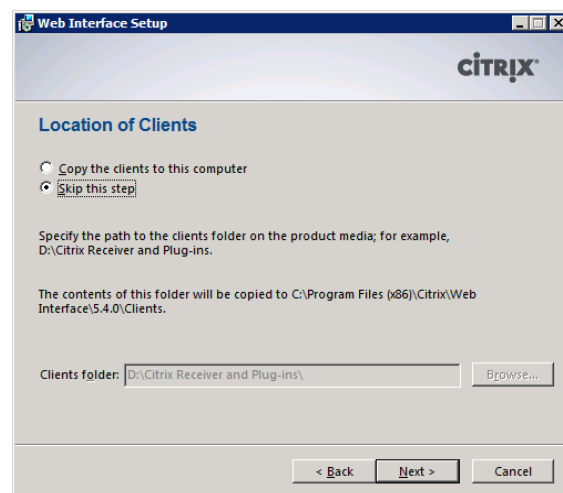
4. Read the entire Citrix License Agreement. Select the button marked **I accept the license agreement** and click **Next**.




5. Choose the default system path to install Citrix Web Interface and click **Next** to continue.



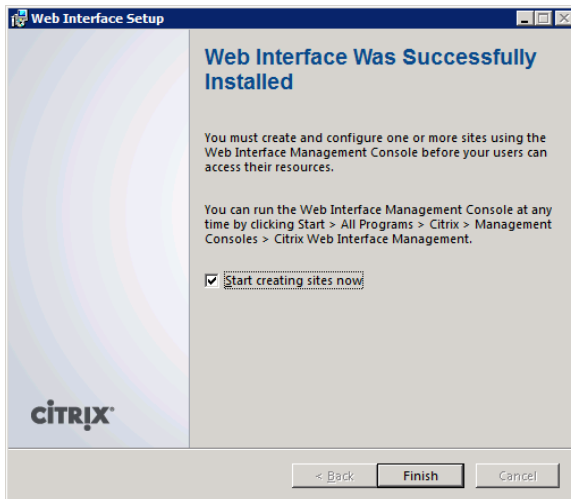
6. Select the button marked **Skip this step** and click **Next** to continue the installation.



7. Read the summary screen and click **Next** to begin the setup file installation.



8. On successful installation of Web Interface components, review the Summary page.
9. Ensure that the checkbox marked **Start creating sites** is checked.
10. Click **Finish** to complete the installation of Web Interface.



### 6.9.7.3 Configuring Web Interface

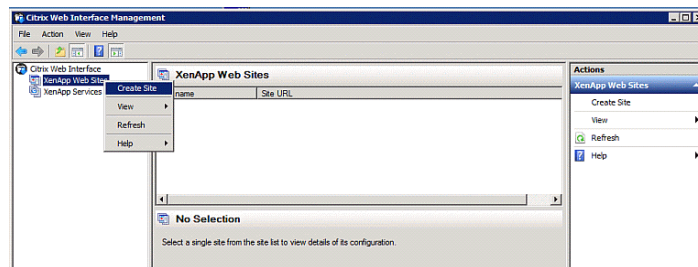
1. In the Web Interface Management Snap-In, right-click **XenApp Web Site** and select **Create Site**.



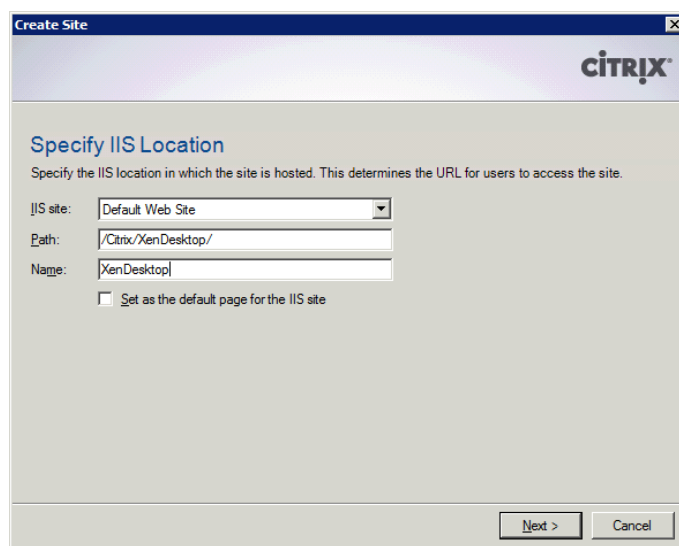
#### Note

Web Interface Management can be found by navigating to:

Start > All Programs > Citrix > Citrix Web Interface Management.



2. In the Create Site dialog, specify information to label the site.
3. Click **Next** to continue.



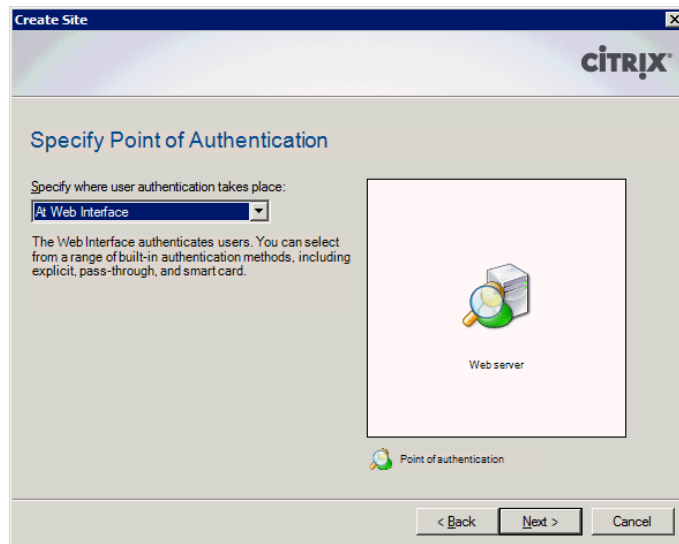
4. Select the Point of Authentication.



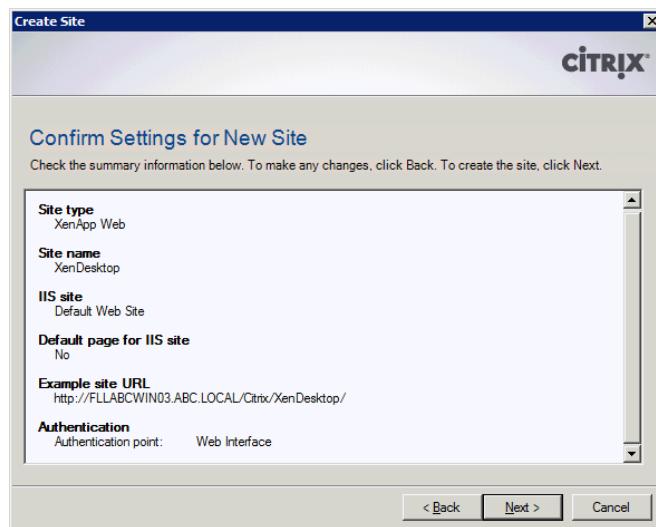
**Note** The Point of Authentication is at the Web Interface.

5. Click **Next** to continue.

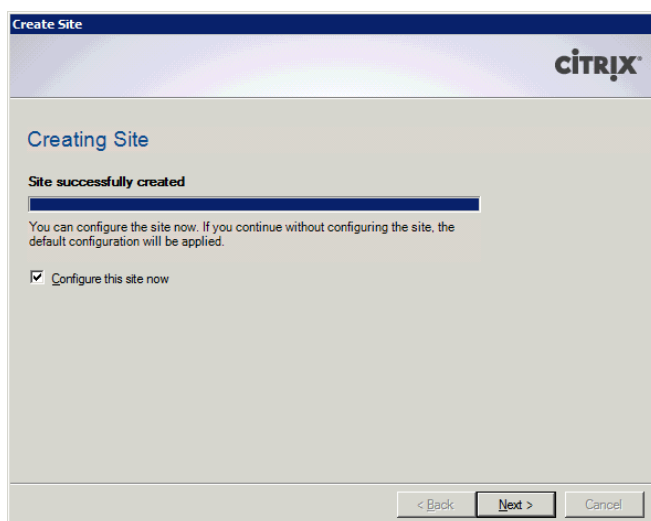




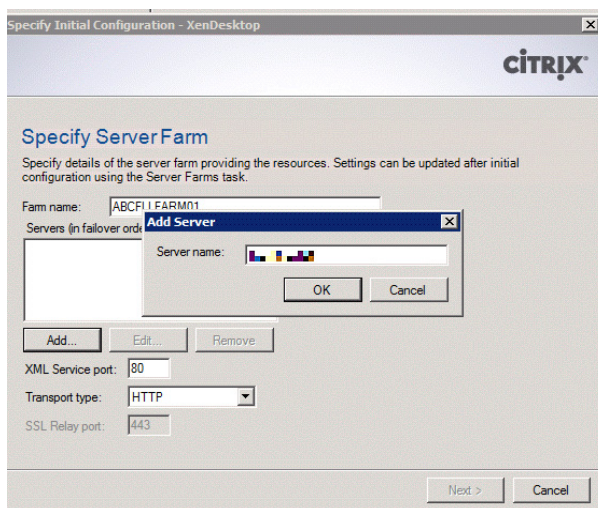
6. Click **Next** to begin configuration of this XenApp Web Site.



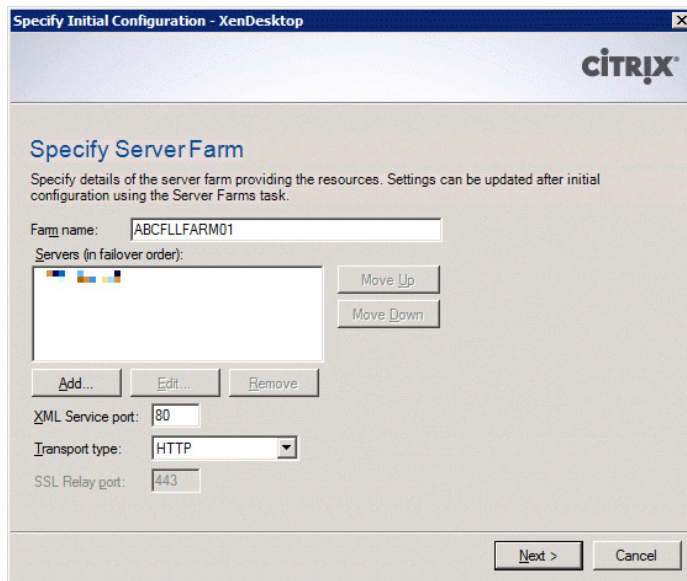
7. Allow the setup wizard to complete installation of the Web Interface Web Site.
8. When the setup wizard has completed, a site summary box appears.
9. Make sure that the checkbox marked **Configure the site now** is checked and click **Next** to continue.



10. The Farm Setup dialog box launches. Specify a Farm Name.
11. In the **Servers** area, click **Add**.
12. Specify the XenDesktop Controller address in the field marked server name. Specify either the fully qualified domain name (FQDN) or the IP address of the controller.
13. Click **OK** to add the server to the Farm.
14. Repeat this step for all of XenDesktop Controllers in the enterprise configuration.



15. When all servers in the XenDesktop Site have been added, click **Next** to continue.



**Specify Initial Configuration - XenDesktop**

**Specify Server Farm**

Specify details of the server farm providing the resources. Settings can be updated after initial configuration using the Server Farms task.

Farm name:

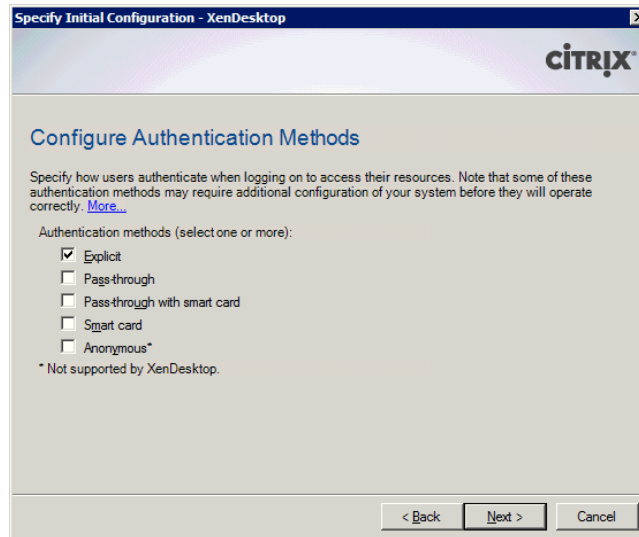
Servers (in failover order):

XML Service port:

Transport type:

SSL Relay port:

16. Configure the Authentication Method appropriate for the enterprise.
17. Click **Next** to continue.



**Specify Initial Configuration - XenDesktop**

**Configure Authentication Methods**

Specify how users authenticate when logging on to access their resources. Note that some of these authentication methods may require additional configuration of your system before they will operate correctly. [More...](#)

Authentication methods (select one or more):

- ☒ Explicit
- ☐ Pass-through
- ☐ Pass-through with smart card
- ☐ Smart card
- ☐ Anonymous\*

\* Not supported by XenDesktop.

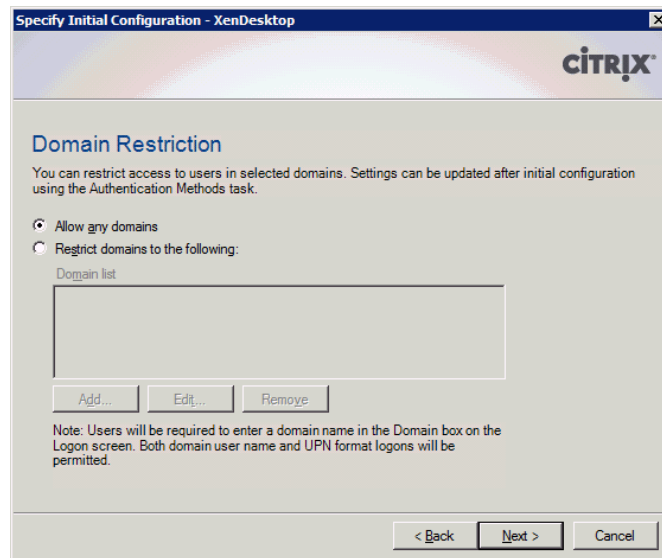
18. Configure the Domain Restrictions.



**Note**

Domain Restriction settings depend on enterprise security needs. Consult with Security Administrators to choose the model most appropriate. For the purposes of this scenario, Allow any domains will be chosen.

19. Click **Next** to continue.

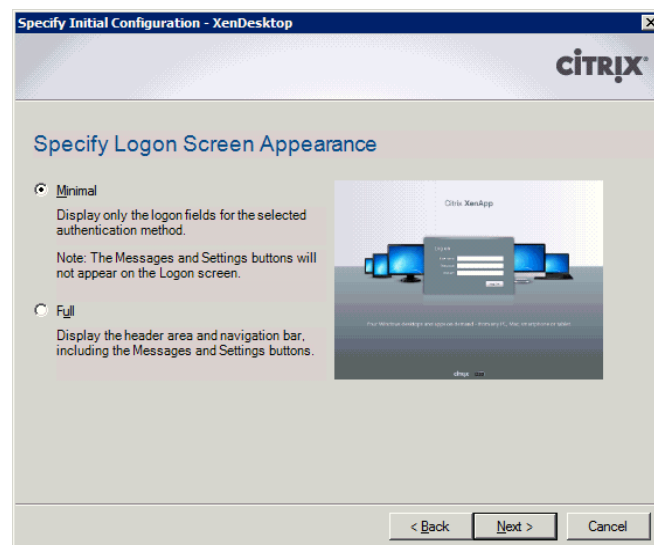


20. Choose the most appropriate Logon Screen appearance. For the purposes of this scenario, choose **Minimal**.

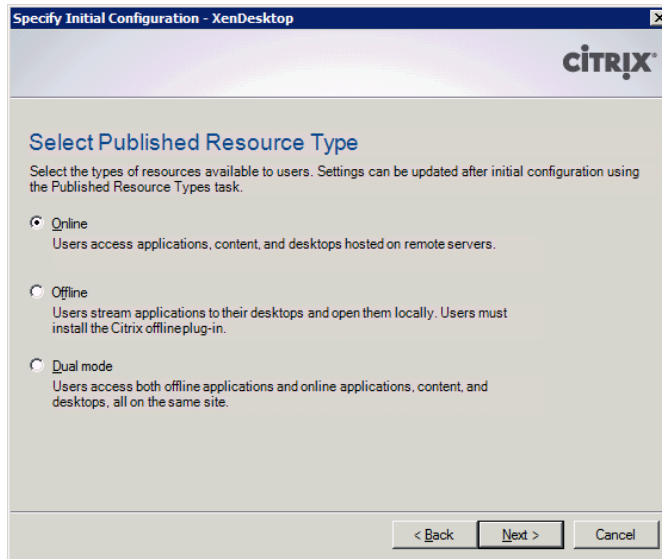


**Note** This logon screen appearance was configured for VSI launcher access.

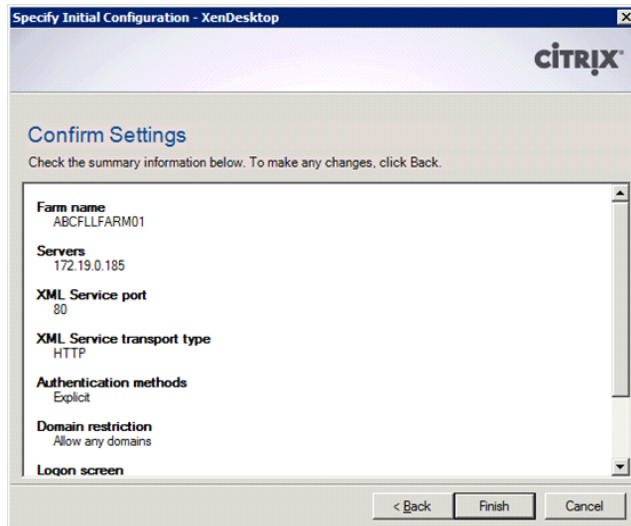
21. Click **Next** to continue.



22. Specify the Published Resource type.
23. Click **Next** to continue.

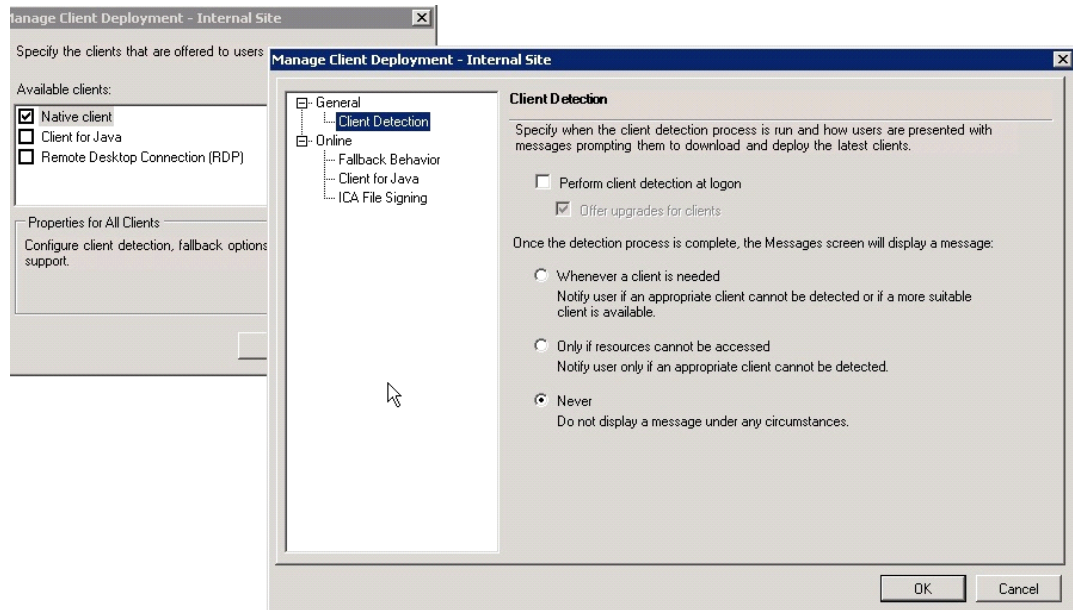


24. Review the Initial Configuration summary for accuracy.
25. Click **Finish** to complete the installation of XenDesktop Web Interface site.



26. On successful configuration, exit from Web Interface management Snap-In.
27. Disable client detection at login:
  - Right-click on your site name.
  - Select **Client Deployment**.
  - Select **Native Client**.
  - Click **Properties**.
  - Click **Client Detection**.
  - Un-check **Perform Client Detection at login**.
  - Select **Never Do not Display** a message under any circumstances.

- Exit with OK.



## 7 Desktop Delivery Infrastructure and Golden Image Creation

Many components are required to support the Virtual Desktop Infrastructure used in this project. This section details the key servers and their roles in the solution. Ultimately, Citrix XenDesktop 5.6 FP1 in combination with Provisioning Services 6.1 managed the VDI environment in this project.

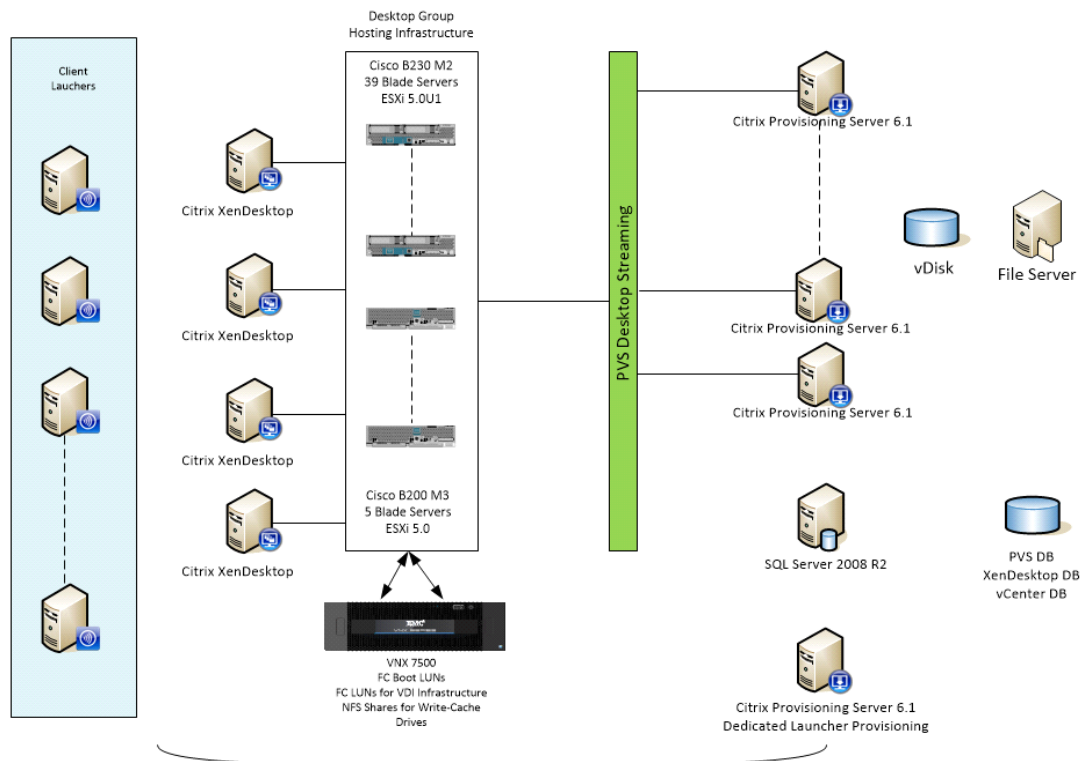
This section includes:

- An overview of the component servers in the solution
- Citrix User Profile Management
- Citrix Provisioning Services 6.1
- Creating the Windows 7 SP1 Golden Image and converting it to a Provisioning Services vDisk

### 7.1 Overview of Solution Components

The image below provides a logical overview of the solution components in the environment.

**Figure 21 Citrix XenDesktop 5.6 and Provisioning Server 6.1 Logical Diagram**



#### Summary of the Environment:

- 5 ESXi-5.0.0-on B200 M3(build 469512-custom-Cisco-2.0.1d)
- 39 ESXi-5.0.0 Update 1-on B230 M2 (build 623860-custom-Cisco-2.0.1.6)
- 20 ESXi-5.0.0-on B250 (build 469512-custom-Cisco-2.0.1d)
- 4 XenDesktop 5.6 FP1 Delivery Controllers
- 10 Provisioning Server 6.1 Server for Virtual Desktops
- 1 ESX vCenter 5.0.0 Update 1b Build 804277
- 210 VSI Launchers
- 5004 Virtual Desktops equally distributed on 3 Clusters
- 1 Citrix Licensing Server
- 1 Web Interface Server
- 2 node File Server Failover Cluster for vDisk and User Profiles
- 2 node Microsoft SQL Server 2008 R2 cluster for Provisioning Services, XenDesktop.
- 2 node Microsoft SQL Server 2008 R2 cluster for vCenter

#### Storage on EMC VNX 7500.

- 64 12 GB Fibre Channel Boot LUNs
- 2 2 TB Shared XenDesktop infrastructure Fibre Channel LUNs

- 2 1 TB Shared Login VSI Infrastructure Fibre Channel LUNs
- 1 200 GB vCenter SQL Fibre Channel LUN
- 12 2TB NFS mounts for Virtual Desktop Write Cache

The following tables provide details on the configuration of the solution components.

**Table 9 Solution Component Configurations**

Vmware ESXi 5.0U1 Hosts			
<b>Hardware :</b>	Cisco B-Series Blade Servers	<b>Model:</b>	B230 - M2
<b>OS:</b>	VMware ESXi 5.0U1	<b>RAM:</b>	256GB
<b>CPU:</b>	2X Deca-Core CPUs	<b>Network:</b>	UCS M81KR (Palo) Converged Network Adapter
<b>Disk:</b>	(Boot LUNs)		

Enterprise Infrastructure Hosts			
<b>Hardware :</b>	Cisco B-Series Blade Servers	<b>Model:</b>	B200 - M3
<b>OS:</b>	VMware ESXi 5.0	<b>RAM:</b>	256GB
<b>CPU:</b>	2X Octa-Core CPUs	<b>Network:</b>	UCS VIC1240 Converged Network Adapter
<b>Disk:</b>	(Boot LUNs)		

Citrix Provisioning Server 6.1			
<b>Hardware :</b>	Virtual Machine	<b>Model:</b>	
<b>OS:</b>	Windows 2008 R2	<b>RAM:</b>	16GB
<b>CPU:</b>	4vCPUs	<b>Network:</b>	1x1Gbps; 1x10Gbps
<b>Disk:</b>	40GB		

Citrix XenDesktop 5.6 Delivery Controllers			
<b>Hardware :</b>	Virtual Machine	<b>Model:</b>	
<b>OS:</b>	Windows 2008 R2	<b>RAM:</b>	4GB
<b>CPU:</b>	1vCPU	<b>Network:</b>	1x1Gbps
<b>Disk:</b>	40GB		

Vmware vCenter Server			
-----------------------	--	--	--



<b>Hardware</b>	Virtual Machine	<b>Model:</b>	
<b>OS:</b>	Windows 2008 R2	<b>RAM:</b>	16GB
<b>CPU:</b>	4vCPUs	<b>Network:</b>	1x10Gbps
<b>Disk:</b>	40GB		

Microsoft SQL Server 2008 R2 for XenDesktop			
<b>Hardware</b>	Virtual Machine	<b>Model:</b>	
<b>OS:</b>	Windows 2008 R2	<b>RAM:</b>	4GB
<b>CPU:</b>	2vCPU	<b>Network:</b>	1x10Gbps
<b>Disk:</b>	40GB		

Microsoft SQL Server 2008 R2 for vCenter			
<b>Hardware</b>	Virtual Machine	<b>Model:</b>	
<b>OS:</b>	Windows 2008 R2	<b>RAM:</b>	80GB
<b>CPU:</b>	8vCPU	<b>Network:</b>	1x10Gbps
<b>Disk:</b>	40GB		

The other dedicated Infrastructure Virtual Machines, all running Microsoft Windows Server 2008 R2 SP1:

- Two Active Directory Servers (Directory Services, DNS, and DHCP)

## 7.2 Citrix User Profile Management Servers

We used Microsoft Server 2008 R2 SP1 File Servers, clustered to provide File Share for Citrix UPM.

### 7.2.1 Setting Up a Highly Available Profile Share

This section explains the installation and configuration of the profile cluster required to support Citrix User Profile Manager. The installation did the following:

- Clustered the two virtual machines
- Create a highly available file share

#### 7.2.1.1 Installing File Services

To set up a highly available profile share, do the following:

1. Open the Server Manager.
2. Browse to the Role node.
3. Click **Add Roles**.
4. Check **File Services**.

5. Click **Next**.
6. Check **File Server**.
7. Click **Next** to install.

### 7.2.1.2 Clustering File Server Application

When that is completed on both Nodes, proceed to Failover Cluster Manager.

1. Right-click Services and applications node and click “**Configure a Service or Application...**”
2. Select File Server and click “**Next**”.
3. Input a File Server name and IP address; then click “**Next**”.
4. Select the cluster disk to be used then click “**Next**”.
5. When that is complete, click “**Add a shared folder**” in the Actions pane.
6. Click Browse and set the profile folder intended for the profiles, then click “**Next**”.
7. Leave the NTFS permission and click “**Next**”.
8. Validate SMB is checked and input Share name and then click “**Next**”.
9. Accept defaults and click “**Next**”.
10. Check Users and groups have custom share permission.
11. Click Permissions and set permissions to Everyone at Full Control; then click “**OK**”.
12. Click “**Next**”.
13. Accept Defaults then click “**Next**”.
14. Review summary and click “**Create**”.

### 7.2.2 Setting Up a Two-Node Citrix User Profile Server Cluster

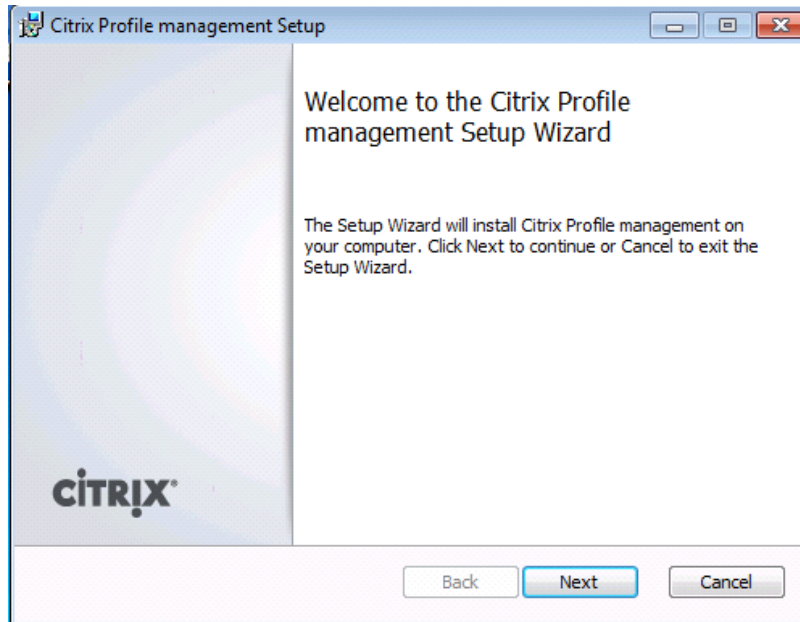
The two LUNs created in the section above, will be used for Quorum and file server clustered disk. Follow the steps below to complete the clustering setup.

1. Install Failover Cluster feature with both servers:
  - a. Open the Server manager.
  - b. Browse to the feature node.
  - c. Click **Add Feature**.
  - d. Check **Failover Clustering**.
  - e. Click **Install**.
2. When this is completed, proceed to the **Start Menu>Administrative Tools>Failover Cluster Manager**.
3. Click **Validate a Configuration** and follow the on screen instructions to validate the two nodes: Profile01 and Profile02. When that succeeds, proceed to the next step.
4. Click **Create a Cluster** and follow the on screen instructions to create a two node cluster.

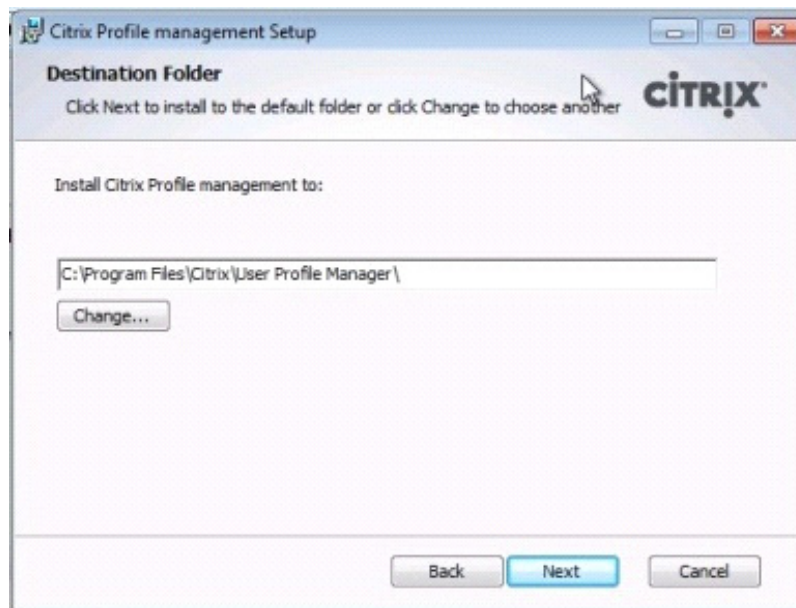
## 7.2.3 Install User Profile Manager

The following steps outline the steps taken to install and configure User Profile Manager In the Virtual Desktop Master Image.

1. Start the UPM installer.

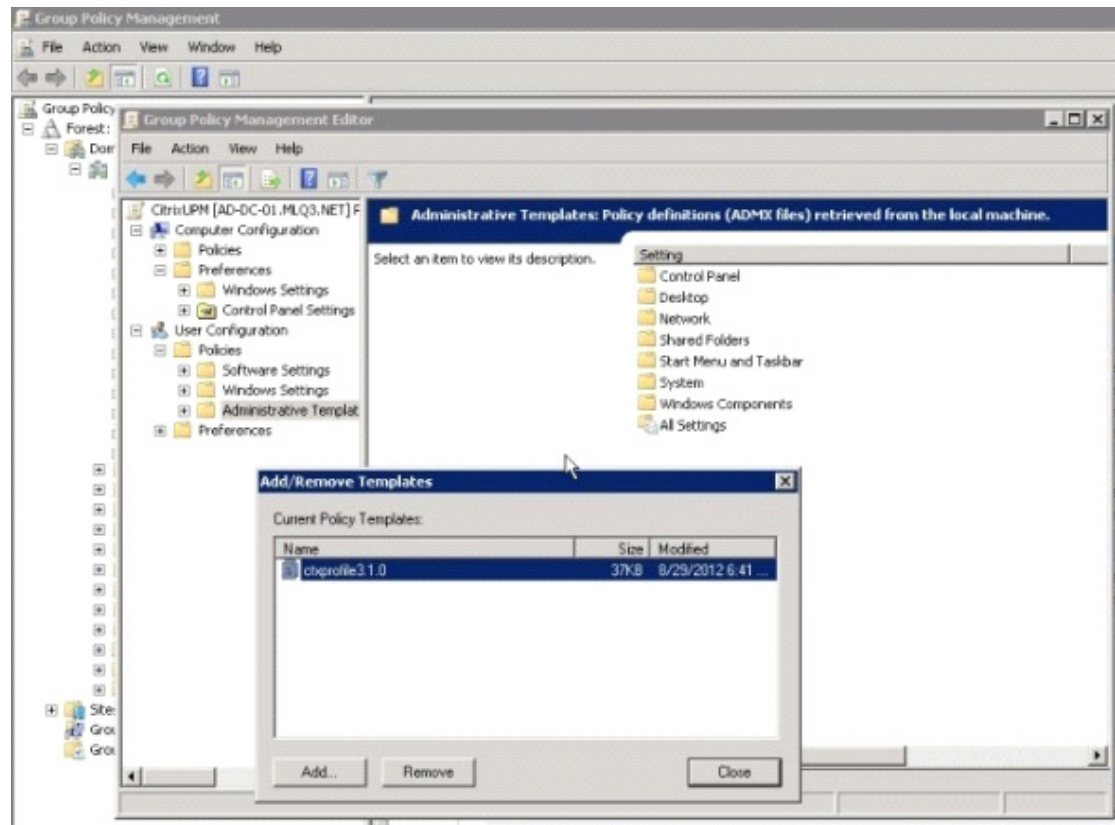


2. Use default installation paths.



3. Click **Next**.
4. Click **Finish** after installation is complete.
5. Create a GPO linked to users OU (organizational Unit).

6. Add Citrix UPM administrative template.
7. Edit the new GPO, browse to **User Configuration > Policies > Administrative Template**.
8. Right-click **Administrative Template** and select **Add/Remove Template**.
9. Click Add.
10. Browse to the location of the template file provided with UPM installation files. (ctxprofile4.1.1.adm)



11. Configure the following settings under **Administrative templates > Citrix>Profile Management**:
12. Enable **Active write back**.
13. Enable **Profile Management**.
14. Enter the absolute path for the location where the profiles will be stored. (An example of the syntax would be `\\upmshare\profiles\%username%`).
15. Select **Enable** for Process logins of local administrators
16. Select **Enable** for the File system |Exclusion list – directories and enter the following information:
  - AppData\LocalLow
  - AppData\Roaming
  - \$Recycle.Bin
  - AppData\Local
17. Click **Log Settings > Enable Logging** and select **Enable**.
18. Click **Profile handling>Delete locally cached profiles on logoff** and select **Disabled**.

19. Click **Local profile conflict handling**.
20. Select **If both local windows profile and Citrix Profile exist**.
21. Select **Delete local profile**.
22. Click **Streamed User Profiles**.
23. Enable **Profile Streaming**.



**Note** These settings were used based on Citrix documentation. Refer to the Reference section of this document for more information.

## 7.3 Microsoft Windows 7 Golden Image Creation

### 7.3.1 Create Base Windows 7 SP1 32bit Virtual Machine

The Microsoft Windows 7 SP1 master or golden image with additional software was initially installed and prepared as a standard virtual machine on VMware ESXi prior to being converted into a separate Citrix Provisioning Server vDisk file. The vDisk is used in conjunction with Provisioning Server 6.1 and the XenDesktop 5.6 controller to provision 5000 new desktop virtual machines on the ESXi hosts.

With XenDesktop 5.6 and Provisioning Server 6.1, the XenDesktop Setup Wizard was utilized.

Each virtual desktop virtual machine was created with a 3.0 GB write cache disk.

The following section describes the process used to create the master or golden image and centralized Windows 7 vDisk used by Provisioning Services.

Virtual Machine Base Software

- Install Win7 32-bit SP1 Enterprise
- Install Office 2010 Professional with Run All From My Computer
- Install Office 2010 Service Pack (most recent)
- Windows Updates (Be sure not to install IE9. Use IE8)

### 7.3.2 Add Provisioning Services Target Device Software



**Note**

Latest version of Target Device Software can be installed directly by installing Hotfix CPVS61E015

1. Launch the PVS Device executable.
2. Click **Next**.
3. Accept the license agreement.
4. Click **Next**.
5. Enter in the customer information.
6. Click **Next**.
7. Choose the default installation location.
8. Click **Next**.

9. Click **Install** to begin the PVS Client installation process.
10. Uncheck **Launch Imaging Wizard** (This process will take place at a later point in the conversion process).
11. Click **Finish**.
12. Click on **Yes** to restart the virtual machine.

### 7.3.3 Add XenDesktop 5.6 Virtual Desktop Agent

1. Copy the VDA executable to the local machine.
2. Launch executable.
3. Select Advanced Install.
4. Use default installation paths.
5. Enter XenDesktop Delivery Controller's FQDN's.
6. Click "**Next**".
7. Do not install HDX or Citrix Receiver.
8. Click "**Finish**".
9. Remove VDA Welcome Screen program from the Windows Startup folder.
10. Restart VM.
11. Log in and check the event log to make sure that the DDC registration has successfully completed.

### 7.3.4 Add Login VSI Target Software

1. Launch setup wizard using run as Administrator.
2. Enter VSI share path.
3. Use default installation paths.
4. Reboot after installation completes.

### 7.3.5 Perform additional PVS and XenDesktop Optimizations

1. Increased the ARP cache lifespan to 600 seconds for stream service bound NICs (article is located in Provisioning Services Reference documentation at the end of the document).
2. Delete XPS Printer.
3. Ensure that Bullzip PDF is the default printer.
4. Optimize:
  - Configure SWAP file to 1536 MB (Cisco requirement)
  - Disable Windows Restore and service
    - Delete the restore points
  - Perform a disk cleanup
  - Disable Windows Firewall Service
  - Disable Windows Backup scheduled jobs

- Open Computer Management > System Tools > Task Scheduler> Task Scheduler Library >Microsoft > Windows and disable the following:
  - Defrag
  - Offline files
  - Windows Backup
- Windows Performance Settings
  - Smooth Edges
  - Use Visual Styles
  - Show Translucent
  - Show Window contents when dragging
- 5. Modify Action Center settings (uncheck all warnings).
- 6. Ensure that the Shadow Copy service is running and set to auto.

### 7.3.6 Convert Golden Image Virtual Machine to PVS vDisk

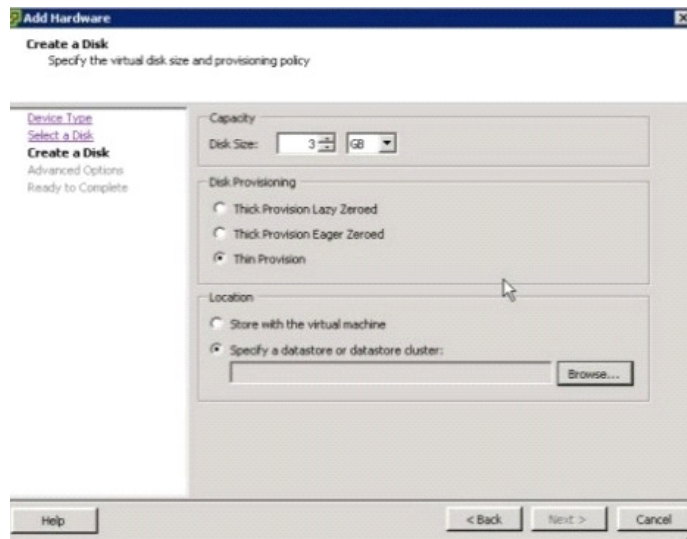
The following is a list of steps taken to convert a virtual machine to a vDisk that will be used to stream desktops through PVS:

1. Reboot the source virtual machine.
2. Log in to the virtual machine using an account that has administrative privileges.
3. Go to Start > All Programs > Citrix> Provisioning Services.
4. Launch the PVS Imaging Wizard.
5. Click **Next** at the Welcome Screen.
6. Enter the Server Name or IP address of the PVS server you will be connecting to in order to create the new vDisk.
7. Select Create A New vDisk.
8. Click “**Next**”.
9. Enter the vDisk Name.
10. Select the PVS Store where the new vDisk will reside.
11. Select VHD type “Fixed”.
12. Click “**Next**”.
13. Select “KMS” for Licensing Management.
14. Click “**Next**”.
15. Use default size image volumes.
16. Click “**Next**”.
17. Click Finish to begin the vDisk creation.
18. You will be prompted to reboot the source virtual machine. Prior to rebooting, go to the properties of the source virtual machine and change the boot options so that it performs a Network boot.
19. Click Yes to reboot the source virtual machine.
20. Logon as the same user account that was used at the beginning of this process.

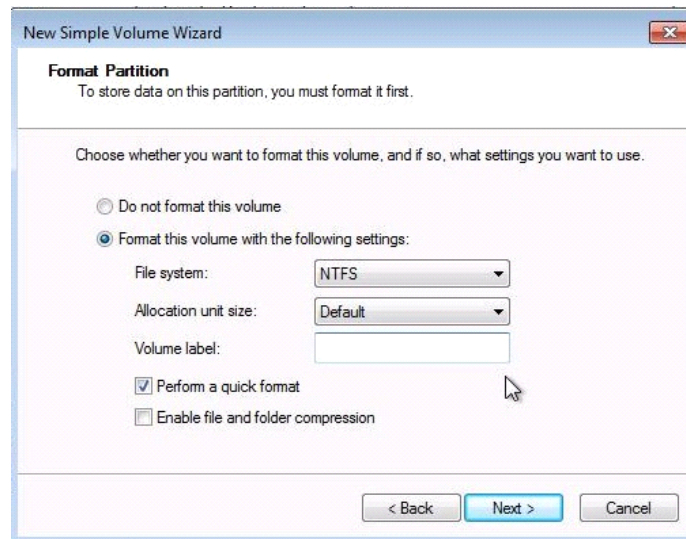
21. Once logged in the Imaging wizard will start the data conversion process. The time needed to complete this process is dependent on the size of the vDisk
22. Shutdown the source virtual machine.
23. Ensure that the VM to boot to Network.
24. In PVS, switch the collection account to boot to vDisk.

### 7.3.7 Add Write-Cache Drives To Virtual Machine Templates In vCenter

1. Add a new 3GB VMDK disk to VM.



2. In the VM's OS, Initiate the disk, and create a simple Volume.
3. Do not assign a Drive Letter.
4. Format as NTFS.





## 7.4 Citrix Provisioning Server (PVS) 6.1 Services

Citrix Provisioning Server (PVS) is part of the XenDesktop Enterprise and Platinum and was used in all tested scenarios. Provisioning provides the ability to create and manage 1000's of virtual machines hosted on hypervisor servers streamed from a single virtual machine vDisk Image.

### 7.4.1 Storage Configuration for PVS

The test environment utilized a single EMC VNX 7500 SAN system to provide storage for PVS 6.1 virtual machines and vDisks. EMC hosted LUNs and volumes were used for:

- PVS 6.1 virtual machines hard drives - Two common Infrastructure LUNs
- Write-cache drives provisioned with each Windows 7 SP1 virtual machine - Twelve NFS datastores
- Windows 7 SP1 vDisks storage accessed by the Provisioning Servers (NFS Share) through Windows File Server Share

The Launcher vDisks were stored on Fibre Channel LUNs on a separate storage pool.

### 7.4.2 PVS for Use With Standard Mode Desktops

The Windows 7 SP1 desktop image is converted into a vDisk (.vhd) image. The vDisk is then configured in a Shared (Read-only) mode and hosted within a shared file location.

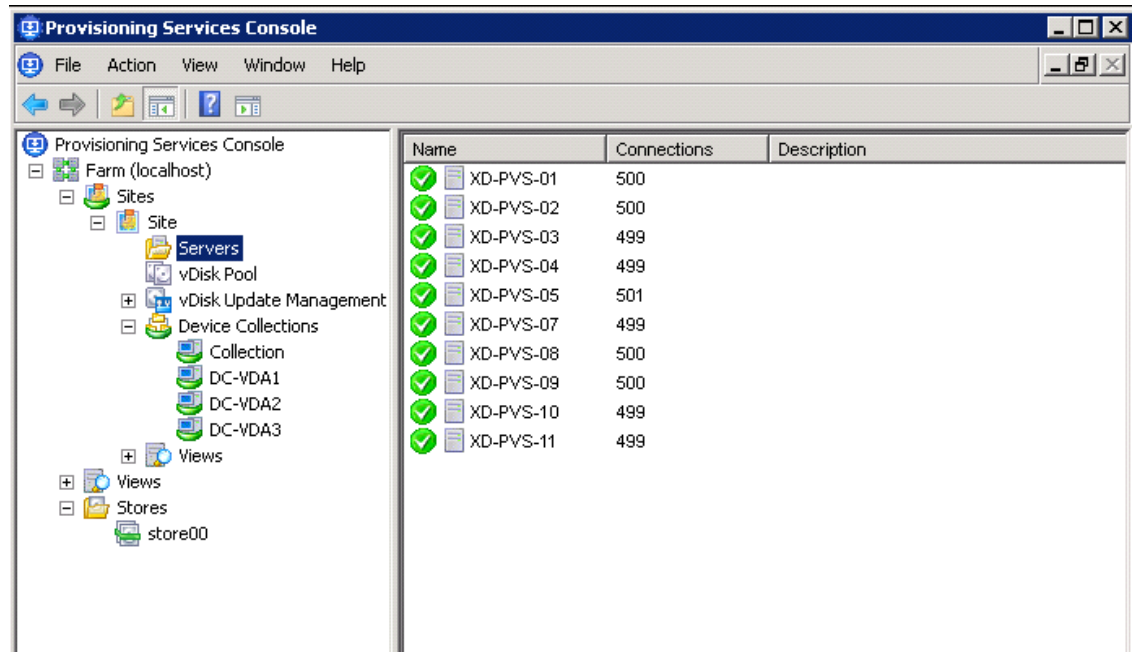
- PVS was used to create the desired number of virtual machines and machine accounts in Active Directory based on parameters specified using the built-in XenDesktop setup wizard (referenced in the next section).
- PVS streams the vDisk image on start-up of the Virtual Machine to the Hypervisor and is loaded into RAM.
- PVS injects a Security Identifier (SID) and host name associated with the virtual machine as each desktop boots to maintain uniqueness in AD. These object mappings are maintained and managed within the PVS server and are visible in the PVS Console under "Collections" view.
- Each virtual desktop is assigned a "Write Cache" (temporary file) where any delta changes (writes) to the default image are recorded and is used by the virtual windows operating system throughout its working life cycle. The Write Cache is written to a dedicated 3GB hard drive.

Ten PVS servers were configured in a farm with a single site to provide streaming services for 5000 Virtual Desktop Virtual Machines with high availability and resilience. Streaming connections are automatically failed over to a working server/s within the farm in the event of a failure without interruption to the desktop.

The vDisk was hosted on a file server cluster and was accessible by all servers in the farm for ease of management and to support high availability. The drive assigned by the hypervisor to the file server cluster for vDisk storage was on a datastore created on a dedicated NFS mount.

Three Device collections were created, one for each ESX cluster, to contain target device records for ease of management.

We assigned PVS servers with 4 vCPUs and 16GB RAM.

**Figure 22** Provisioning Services Farm Layout

A separate PVS server with local storage was used to provision Login VSI Launcher machines for test workload. We used two FC LUNs on the VNX 7500 to create and store each virtual machine's Write-Cache drive.

It is important to consider where the Write Cache is placed when scaling virtual desktops using PVS server. There are several options as to where the Write Cache can be placed:

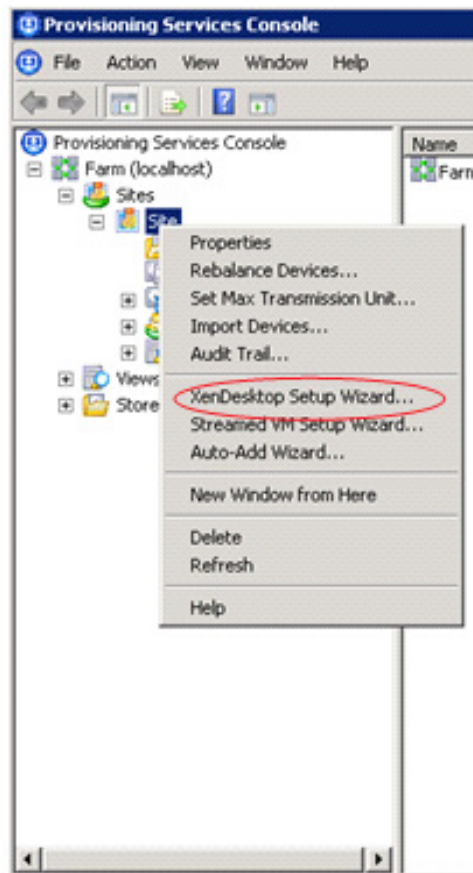
- PVS Server
- Hypervisor RAM
- Device Local Disk (an additional Virtual Disk for XenDesktop Virtual Machine)

For this project's optimal performance and scalability the Cache on device HD option is used. A 3GB virtual disk is assigned to the virtual machine templates used in the clone creation process.

The PVS Target device agent installed in the Windows 7 gold image automatically places the Windows swap file on the same drive used by the PVS Write Cache when this mode is enabled.

### 7.4.3 Process to Create Virtual Desktops Using XenDesktop Wizard in PVS

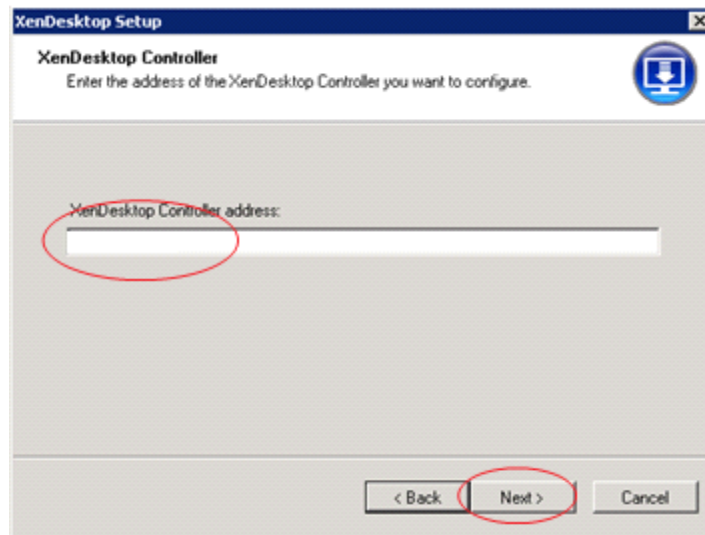
1. Start XenDesktop Setup Wizard.



2. Click Next.



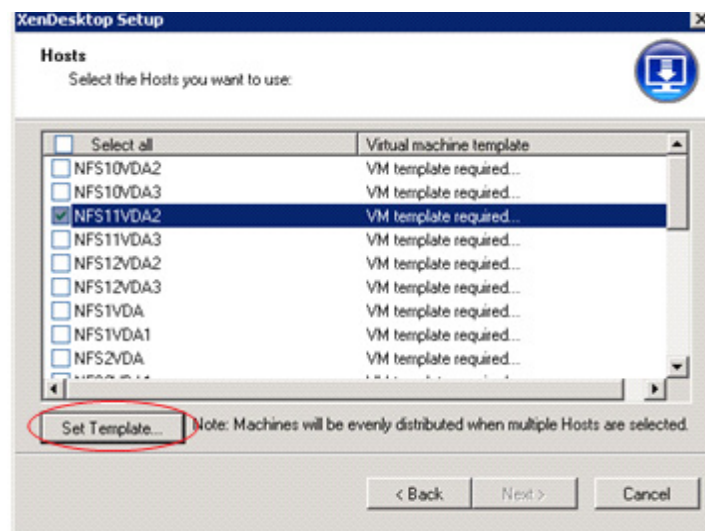
3. Connect to XenDesktop Controller.



#### 4. Configure VM template.



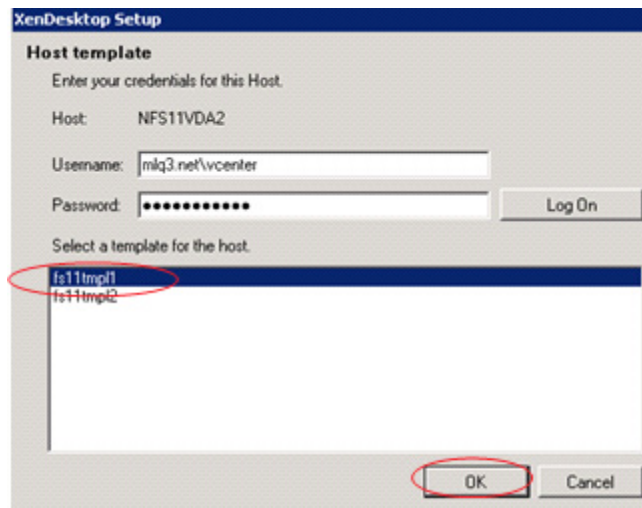
**Note** The connection to the type of host you are using with the credentials to use when accessing it have to be defined on Desktop Delivery Controller prior to this step.



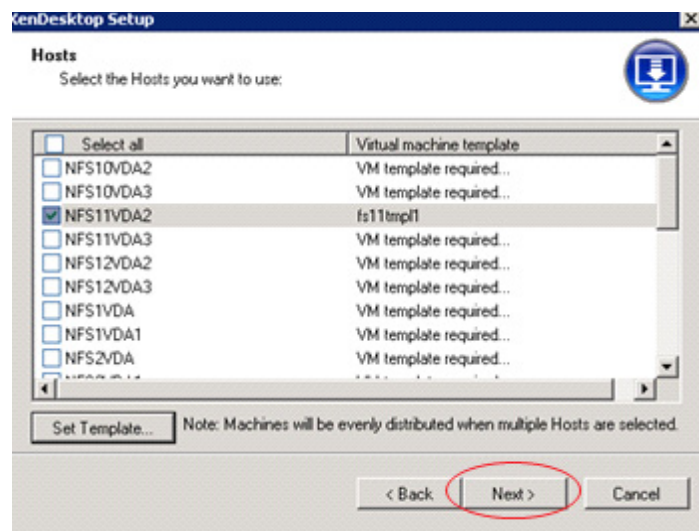
#### 5. Input Password to connect to vCenter and Select VM template you are going to use.



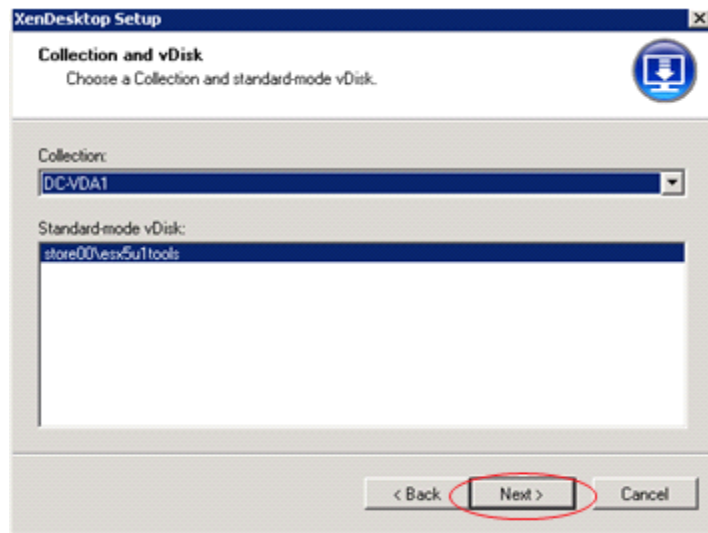
**Note** Prior to this step, the VM templates need to be configured on each VMware datastore that will contain drives for the streamed desktops.



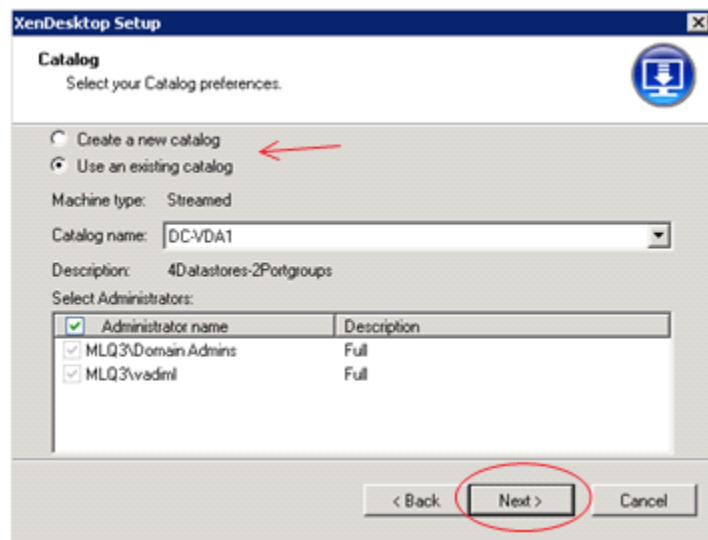
6. Click Next.



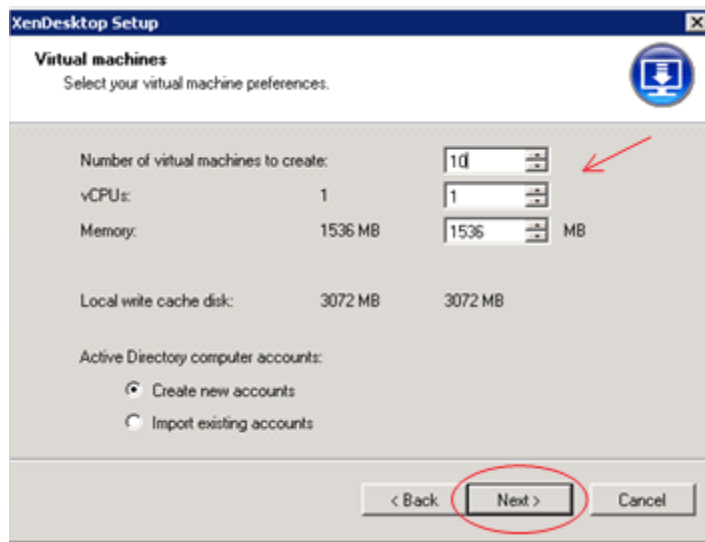
7. Select PVS device Collection.



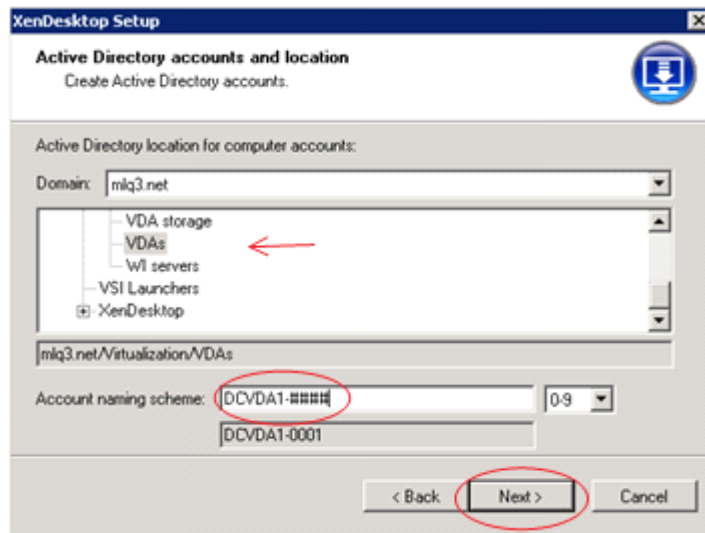
8. Select XenDesktop Catalog preferences.



9. Select Virtual Machine preferences.



10. Create machine accounts in Active Directory.



11. Click Finish to complete the following:
  - Create virtual machines on selected hypervisor hosts.
  - Create Provisioning Services target devices in the selected collection.
  - Create Active Directory computer accounts.
12. Create XenDesktop machines in the assigned existing catalog or a new XenDesktop catalog.

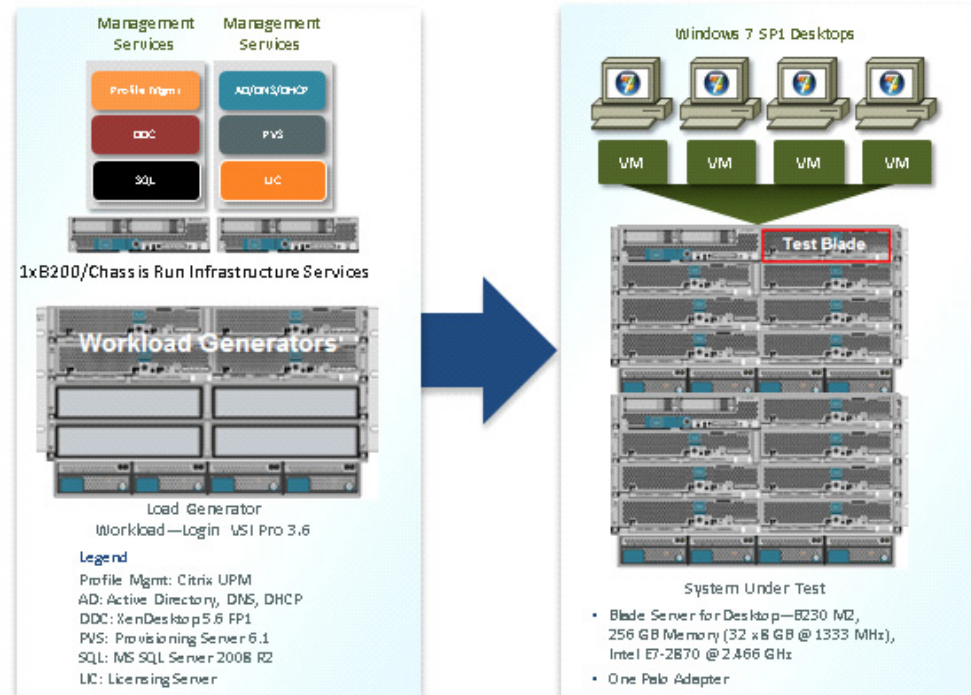
## 8 Test Setup and Configurations

In this project, we tested a single Cisco UCS B230 M2 blade in a single chassis and thirty nine B230 M2 blades in five chassis to illustrate linear scalability.

## 8.1 Cisco UCS Test Configuration for Single Blade Scalability

Figure 23 Cisco UCS B230 M2 Blade Server for Single Server Scalability

### Cisco UCS B230 M2 Blade Server Single Blade Test Result— 149 Users



#### Hardware Components

- 1 X Cisco UCS B230-M2 (E7-2870 @ 2.4 GHz) blade server with 256GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktop hosts
- 5 X Cisco UCS B200-M3 (E5-2650) blade servers with 96 GB of memory (8 GB X 12 DIMMS @ 1600 MHz) Infrastructure Servers
- 4 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Blade (B250 M2 and B230 M2)
- 1X VIC1240 Converged Network Adapter/Blade (B200 M3)
- 2 X Cisco Fabric Interconnect 6248UPs
- 2 X Cisco Nexus 5548UP Access Switches
- 1 X EMC VNX System storage array, two controllers, four Datamovers, 2 x dual port 8GB FC cards, 4 x dual port 10 GbE cards, 10 x 200GB Flash Drives for EMC Fast Cache, 96 x 600GB SAS drives for PVS Write Cache, 24 x 600GB SAS Drives for Infrastructure and Boot LUNs and 5 x 600GB SAS drives for hot spares

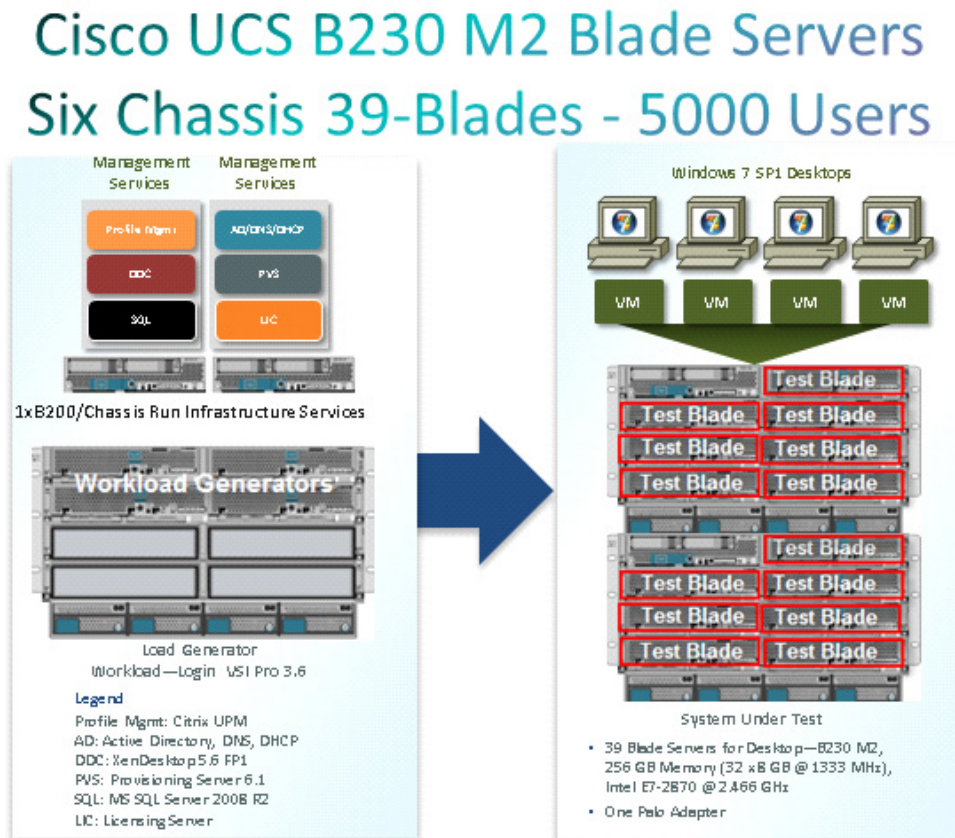


### Software Components

- Cisco UCS firmware 2.0(4a)
- Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.0 Update 1 for VDI Hosts
- XenDesktop 5.6 Feature Pack 1
- Provisioning Server 6.1
- Citrix User Profile Manager
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 17 GB/VM

## 8.2 Cisco UCS Configuration for Six Chassis Thirty-Nine Blade Test

Figure 24 Six Chassis Test Configuration-39 x B250 Blade Server



### Hardware Components

- 39 X Cisco UCS B230-M2 (E7-2870 @ 2.4 GHz) blade server with 256GB of memory (8 GB X 32 DIMMS @ 1333 MHz) Windows 7 SP1 Virtual Desktop hosts
- 5 X Cisco UCS B200-M3 (E5-2650) blade servers with 96 GB of memory (8 GB X 12 DIMMS @ 1600 MHz) Infrastructure Servers

- 4 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) Load Generators
- 1 X M81KR (Palo) Converged Network Adapter/Blade (B250 M2 and B230 M2)
- 1X VIC1240 Converged Network Adapter/Blade (B200 M3)
- 2 X Cisco Fabric Interconnect 6248UPs
- 2 X Cisco Nexus 5548UP Access Switches
- 1 X EMC VNX System storage array, two controllers, four Datamovers, 2 x dual port 8GB FC cards, 4 x dual port 10 GbE cards, 10 x 200GB Flash Drives for EMC Fast Cache, 96 x 600GB SAS drives for PVS Write Cache, 24 x 600GB SAS Drives for Infrastructure and Boot LUNs and 5 x 600GB SAS drives for hot spares

#### Software Components

- Cisco UCS firmware 2.0(4a)
- Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.0 Update 1 for VDI Hosts
- XenDesktop 5.6 Feature Pack 1
- Provisioning Server 6.1
- Citrix User Profile Manager
- Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 17 GB/VM

## 8.3 Testing Methodology and Success Criteria

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

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the Hosted VDI model under test.

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

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

### 8.3.1 Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the XenDesktop 5.6 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 XenDesktop 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 XenDesktop 5.6 through a Citrix HDX protocol connection.

### 8.3.2 User Workload Simulation–LoginVSI From Login Consultants

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

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

The Virtual Session Indexer 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 nor hosted on XenApp).

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

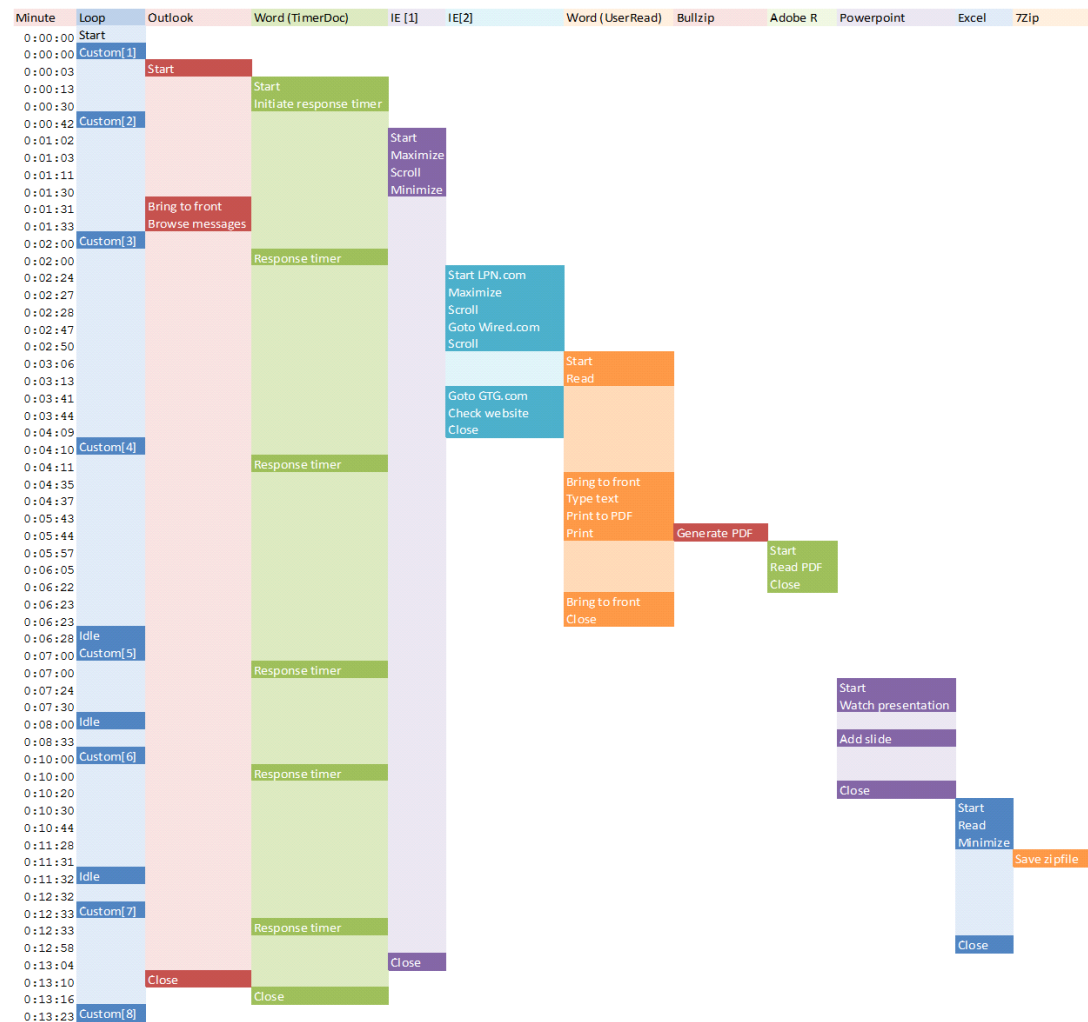
- When a session has been started the medium workload will repeat every 12 minutes.
- During each loop the response time is measured every 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The type rate is 160ms for each character.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

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

A graphical representation of the medium workload is shown below.

### Figure 25



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

### 8.3.3 Testing Procedure

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

#### 8.3.3.1 Pre-Test Setup for Single and Multi-Blade Testing

All virtual machines were shut down utilizing the Citrix XenDesktop 5.6 Desktop Studio.

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 Hyper-V Server 2008 R2 SP1 VDI host blades to be tested were restarted prior to each test cycle.

### 8.3.3.2 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 175 single server users or 5000 full scale test users to become active within 2 minutes after the session is launched.

For each of the three consecutive runs on single blade (175 User) and 39-blade (5000 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
  - PVS Server(s) used in test run
  - DDCs used in test run
  - Profile Servers used in test run
  - SQL Servers 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
4. Time 0:05 Take 175 or 5000 desktops out of maintenance mode on XenDesktop Studio
5. Time 0:06 First machines boot
6. Time 0:46 175 or 5000 desktops booted on 1 or 39 blades
7. Time 0:48 175 or 5000 desktops available on 1 or 39 blades
8. Time 1:50 Start Login VSI 3.6 Test with 175 or 5000 desktops utilizing 4 or 200 Launchers
9. Time 2:20 145 or 5000 desktops launched
10. Time 2:22 145 or 5000 desktops active
11. Time 12:37 Login VSI Test Ends
12. Time 2:52 145 or 5000 desktops logged off
13. Time 2:55 All logging terminated.

### 8.3.4 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.

#### 8.3.4.1 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.

### 8.3.4.2 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.

### 8.3.4.3 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 be 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  $1800 \times 125\% + 3000 = 5250\text{ms}$ .


Especially when application virtualization is used, the baseline response time can wildly vary per vendor and streaming strategy. Therefore it is recommended 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.

### 8.3.4.5 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 26 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 B230 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 3.6 medium workloads, including flash.
- Number of Launched Sessions must equal Active Sessions within two minutes of the last session launched in a test.
- The Citrix Desktop Studio 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 recommendation following the process above and will note that we did not reach a VSIMax dynamic in our testing due to a technical issue with the analyzer formula that calculates VSIMax.

## 9 VDI Test Results


The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 5.6 FP1 Hosted VDI FlexCast model and Citrix Provisioning Services 6.1 using ESXi 5.0 Update 1 and vCenter 5.0 Update 1B to virtualize Microsoft Windows 7 SP1 desktops on Cisco UCS B230 M2 Blade Servers using a EMC VNX 7500 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 XenDesktop with VMware vSphere.

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

One series of stress tests on a single blade server was conducted to establish the official Login VSI Max Score. To reach the Login VSI Max, we ran 175 Medium with flash Windows 7 SP1 sessions on a single blade. The Login VSI score was achieved on three consecutive runs and is shown below.

**Figure 27 Login VSIMax Reached: 149 Users**



## 9.1 Cisco UCS Test Configuration for Single-Server Scalability Test Results

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



### Note

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


We ran the single server test at approximately 9% lower user density than indicated by the Login VSIMax to achieve a successful pass of the test with server hardware performance in a realistic range.

Given adequate storage capability, the CPU utilization determined the maximum recommended VM density per blade for the single server scalability testing.



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.

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


**Figure 28** 139 Desktop Sessions on VMware ESXi 5.0U1 below 4000 ms



The following graphs detail CPU, Memory, Disk and Network performance on the Single Cisco UCS B230-M2 Blades.

**Figure 29** 139 User Single B230 M2 CPU Utilization Boot Phase**Figure 30** 139 User Single B230 M2 Available Memory


**Figure 31** 139 User Single B230 M2 Cisco M81KR VIC Mbps Receive/Transmit



The following graphs detail performance of the EMC VNX 7500 during the single blade, 139 user test. VNX Graphs (EMC).

The following charts detail infrastructure server performance during the single blade, 139 User test:

**Figure 32** 139 User Provisioning Services 6.1 XD-PVS-01 CPU Utilization





**Figure 33** 139 User Provisioning Services 6.1 XD-PVS-01 Available Memory**Figure 34** 139 User Provisioning Server XD-PVS-01 Bytes Received /Second

Figure 35 139 User Provisioning Server XD-PVS-01 Bytes Sent/Second




Figure 36 139 User XenDesktop 5.6 FP1 Controller XD-DDC01 CPU Utilization




Figure 37 139 User XenDesktop 5.6 FP1 Controller XD-DDC01 Available Memory






Figure 38 139 User XenDesktop 5.6 FP1 Controller XD-DDC01 Bytes Received



**Figure 39** 139 User XenDesktop 5.6 FP1 Controller XD-DDC01 Bytes Sent



## 9.2 Cisco UCS Test Configuration for 5000 Desktop Scalability Test Results

This section details the results from the XenDesktop Hosted VDI thirty-nine blade server 5000 user validation testing. It demonstrates linear scalability for the system. The primary success criteria used to validate the overall success of the test cycle is an output chart from Login Consultants' VSI Analyzer Professional Edition, VSIMax Dynamic for the Medium workload (with Flash.)



### Note

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

We ran the multi-server test at an average user density slightly higher than 128 users per blade across the system. Three ESX Clusters, each containing 13 B230 M2s ran the entire workload. In fact thirty-six of the thirty-nine blades, twelve in each cluster, ran at 130 users and three blades, one in each cluster ran at 108 users, providing 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 130 user sessions. The single server graphs for blades running 130 user sessions are essentially the same. We have provided one cluster's 13-blade 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 5000 user environment.

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

For the large scale test, we are including the EMC VNX7500 performance metrics as well.



**Figure 40** 5004 Desktop Sessions on VMware ESXi 5.0U1 below 4000 ms

The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B230-M2 Blade during the thirty-nine blade, 5000 User test. (Representative results for all thirteen blades in one of the vCenter clusters can be found in Appendix C.)


**Figure 41** 5000 User Single B230 M2 CPU Utilization Boot Phase

Figure 42 5000 User Single B230 M2 CPU Utilization Boot Phase




Figure 43 5000 User Single B230 M2 Cisco M81KR VIC Mbps Receive/Transmit Boot Phase




Figure 44 5000 User Single B230 M2 CPU Utilization Test Phase






Figure 45 5000 User Single B230 M2 CPU Utilization Test Phase



**Figure 46** 5000 User Single B230 M2 Cisco M81KR VIC Mbps Receive/Transmit Test Phase



The following charts detail infrastructure server performance during the thirty-nine blade, 5004 User test:

**Figure 47** 5000 User Provisioning Services 6.1 XD-PVS-01 CPU Utilization Test Phase




Figure 48 5000 User Provisioning Services 6.1 XD-PVS-01 Available Memory Test Phase




Figure 49 5000 User Provisioning Server XD-PVS-01 Bytes Received/Second Test Phase




Figure 50 5000 User Provisioning Server XD-PVS-01 Bytes Sent/Second Test Phase







Figure 51 5000 User XenDesktop 5.6 FP1 Controller XD-DDC-01 CPU Utilization Test Phase



**Figure 52** VNX7500 NFS Read Operations/Second**Figure 53** VNX7500 NFS Write Operations/Second



**Figure 54** VNX7500 Disk IOPS**Figure 55** 5000 User XenDesktop 5.6 FP1 Controller XD-DDC-01 Available Memory Test Phase



Figure 56 5000 User XenDesktop 5.6 FP1 Controller XD-DDC-01 Bytes Received/Second Test Phase





Figure 57 5000 User XenDesktop 5.6 FP1 Controller XD-DDC-01 Bytes Sent/Second Test Phase



## 10 Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 5000 users, six chassis, 39 VDI host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 5000 user system.

### 10.1 Cisco Unified Computing System Configuration

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

- Cisco UCS 2.0 management software supports up to 20 chassis within a single Cisco UCS domain on our second generation Cisco UCS Fabric Interconnect 6248 and 6296 models. Our single UCS domain can grow to 160 blades.
- With Cisco UCS 2.1 management software, released late in November 2012, each UCS 2.1 Management domain is extensibly manageable by UCS Central, our new manager of managers, vastly increasing the reach of the 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% of the time.
- To accommodate the Cisco Nexus 5500 upstream connectivity in the way we describe in the LAN and SAN Configuration section, we need four Ethernet uplinks and two Fibre Channel uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we can calculate number of desktops can be hosted in a single 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 B230 M2 and one B200 M3 blades servers in each chassis.

Of course, 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.

## 10.2 Citrix XenDesktop 5.6 Hosted VDI

XenDesktop environments can scale to large numbers. When implementing Citrix XenDesktop hosted VDI considerations include but not limited to:

- Types of Storage in your environment
- Types of desktops that will be deployed
- Data protection requirements
- For Citrix Provisioning Server pooled desktops write cache size and placement

These and other various aspects of scalability considerations described in greater detail in “XenDesktop - Modular Reference Architecture” document and should be a part of any VDI design.

Designing and deploying our test environment we followed best practices whenever possible.

The following are worth mentioning:

- Citrix always recommends using N+1 schema for VDI servers, to accommodate resiliency. In our test environment, this was applied to all infrastructure servers.
- All Provisioning Server Network Adapters were configured to have a static IP and management and streaming traffic was separated between different Network adapters.
- All the PVS services to start as: Automatic (Delayed Start).
- We used the XenDesktop Setup Wizard in PVS. Wizard does an excellent job of creating the desktops automatically and it's possible to run multiple instances of the wizard provided the deployed desktops are placed in different catalogs and have different naming conventions.
- To run wizard at a minimum you need to install the Provisioning Server, the XenDesktop Controller, and configure hosts, as well as create VM templates on all datastores where desktops will be deployed.

- At the 5000 desktop scale, we utilized 3 VMware Clusters, each with a dedicated Nexus 1000V VSM pair. Twelve NFS datastores were used to host the desktop write cache drives, four per ESX cluster. We had to deploy 2 templates per NFS datastore. Each of the two VM templates were attached to a different virtual ethernet port group, each of which was connected to the single VDA uplink port group. This was done due to a 1024 port limitation per virtual ethernet port profile in the Nexus 1000V.

## 10.3 EMC VNX Storage Guidelines for XenDesktop Provisioned Virtual Machines

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

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

## 10.4 VMware ESXi 5 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 cited in References section of this document.

# 11 References

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

## 11.1 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](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 B230 M2 Blade Server Resources

<http://www.cisco.com/en/US/products/ps11583/index.html>

Cisco UCS B200 M3 Blade Server Resources

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

Cisco Nexus 1000V Series Switches Resources

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

Cisco Nexus 5500 Series Switches Resources

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

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

<http://software.cisco.com/download/release.html?mdfid=283612660&flowid=22121&softwareid=283655658&release=2.0%284d%29&relind=AVAILABLE&rellifecycle=&reltype=latest>

Download Cisco UCS Central Software Version 1.0(1a)

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

## 11.2 Citrix Reference Documents

### XenDesktop 5.6

- Modular Reference Architecture - <http://support.citrix.com/article/CTX133162>

### Provisioning Services 6.1

- Hotfixes
  - CPVS61E001: <http://support.citrix.com/article/CTX133149>
  - CPVS61E002 <http://support.citrix.com/article/CTX133500>
  - CPVS61E003 <http://support.citrix.com/article/CTX133349>
  - CPVS61E004 <http://support.citrix.com/article/CTX133516>
  - CPVS61E005 <http://support.citrix.com/article/CTX133518>
  - CPVS61E006 <http://support.citrix.com/article/CTX133707>
  - CPVS61E007 <http://support.citrix.com/article/CTX133811>
  - CPVS61E008 <http://support.citrix.com/article/CTX135769>
  - CPVS61E009 <http://support.citrix.com/article/CTX133998>
  - CPVS61E010 <http://support.citrix.com/article/CTX134071>
  - CPVS61E011 <http://support.citrix.com/article/CTX134519>
  - CPVS61E014 <http://support.citrix.com/article/CTX134611>
  - CPVS61E015 <http://support.citrix.com/article/CTX134721>
- Description of Address Resolution Protocol (ARP) caching behavior in Windows Vista TCP/IP implementations: <http://support.microsoft.com/kb/949589>

## Virtual Desktop

- Windows 7 Optimization Guide: <http://support.citrix.com/article/CTX127050>
- Web Interface Best Practices for Avoiding Major Production Outages: <http://support.citrix.com/article/CTX125715>

## Citrix User Profile Manager

- UPM settings for XenDesktop: <http://support.citrix.com/proddocs/topic/user-profile-manager-sou/upm-plan-decide-wrapper.html>

## 11.3 EMC Reference Documents

- Sizing EMC VNX Series for VDI Workload – An Architectural Guideline
- EMC Infrastructure for Citrix XenDesktop 5.6, EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1—Reference Architecture
- EMC Infrastructure for Citrix XenDesktop 5.6, EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1—Proven Solutions Guide
- EMC Infrastructure for Citrix XenDesktop 5.5 (PVS), EMC VNX Series (NFS), Cisco UCS, Citrix XenDesktop 5.5 (PVS), Citrix XenApp 6.5, and XenServer 6—Reference Architecture
- EMC Infrastructure for Citrix XenDesktop 5.5 (PVS), EMC VNX Series (NFS), Cisco UCS, Citrix XenDesktop 5.5 (PVS), Citrix XenApp 6.5, and XenServer 6—Proven Solution Guide
- EMC Infrastructure for Citrix XenDesktop 5.5 , EMC VNX Series (NFS), Cisco UCS, Citrix XenDesktop 5.5, Citrix XenApp 6.5, and XenServer 6—Reference Architecture
- EMC Infrastructure for Citrix XenDesktop 5.5, EMC VNX Series (NFS), Cisco UCS, Citrix XenDesktop 5.5, Citrix XenApp 6.5, and XenServer 6—Proven Solution Guide

## 11.4 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](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](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](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](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf)

# Appendix A—Nexus 5548UP Configurations

## A.1 N5548UP-A Configuration

```

version 5.2(1)N1(1)
feature fcoe
logging level feature-mgr 0
hostname N5548UP-A
feature npiv
feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp

username admin password 5 $1$puGfnNws$kvJHTyKpPn6bkDRDhA1Yy. role
network-adminno password strength-check

banner motd #Nexus 5000 Switch
#

ip domain-lookup
logging event link-status default
ip access-list acl-vnx
  10 permit ip any any
class-map type qos class-fcoe
class-map type qos match-all cm-qos-vnx
  match access-group name acl-vnx
class-map type qos match-all cm-qos-cos5
  match cos 5
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing cm-que-qosgrp5
  match qos-group 5
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 pm-qos-vnx
  class cm-qos-vnx
    set qos-group 5

```

```

class class-default
policy-map type qos pm-qos-global
  class cm-qos-cos5
    set qos-group 5
  class class-default
class-map type network-qos class-fcoe
  match qos-group 1
class-map type network-qos cm-nq-grp5
  match qos-group 5
class-map type network-qos class-all-flood
  match qos-group 2
class-map type network-qos class-ip-multicast
  match qos-group 2
policy-map type network-qos pm-nq-global
  class type network-qos cm-nq-grp5
    mtu 9216
    set cos 5
  class type network-qos class-default
    multicast-optimize
system qos
  service-policy type qos input pm-qos-global
  service-policy type network-qos pm-nq-global
slot 1
slot 2
  provision model N55-M8P8FP
snmp-server user admin network-admin auth md5
0xe9ffdef984e3778fa0237fc8b0a8faa1 priv

0xe9ffdef984e3778fa0237fc8b0a8faa1 localizedkey
vrf context management
  ip route 0.0.0.0/0 10.218.160.1
vlan 1
vlan 329
  name uplink2_R2E04-3750-Sw1EDGE
vlan 800
  name ML-VDA
vlan 801
  name ML-DC-VM-MGMT
vlan 802
  name ML-DC-VMMOTION
vlan 803
  name ML-DC-INF
vlan 804

```

```

        name ML-DC-STRG
vlan 850
        name ML_BR-MGMT
vlan 851
        name ML_Launcher-Inf
vlan 852
        name ML_LauncherA
vlan 853
        name ML_LauncherB
vlan 854
        name ML_LauncherC
vlan 900
        name ML-N1KV_CTRL
vlan 901
vpc domain 10
    role priority 1000
    peer-keepalive destination 10.218.164.121
port-profile default max-ports 512
device-alias database
    device-alias name VNX7500-A0 pwn 50:06:01:60:46:e0:5e:0a
    device-alias name VNX7500-B1 pwn 50:06:01:69:46:e0:5e:0a
    device-alias name B200M3-CH2-SERVER1-fc0 pwn 20:00:00:25:b5:c1:00:1a
    device-alias name B200M3-CH2-SERVER3-fc0 pwn 20:00:00:25:b5:c1:00:57
    device-alias name B200M3-CH2-SERVER4-fc0 pwn 20:00:00:25:b5:c1:00:b7
    device-alias name B200M3-CH2-SERVER5-fc0 pwn 20:00:00:25:b5:c1:00:48
    device-alias name B200M3-CH2-SERVER6-fc0 pwn 20:00:00:25:b5:c1:00:a8
    device-alias name B200M3-CH2-SERVER7-fc0 pwn 20:00:00:25:b5:c1:00:49
    device-alias name B200M3-CH2-SERVER8-fc0 pwn 20:00:00:25:b5:c1:00:c6
    device-alias name B200M3-CH3-SERVER1-fc0 pwn 20:00:00:25:b5:c1:00:86
    device-alias name B200M3-CH3-SERVER2-fc0 pwn 20:00:00:25:b5:c1:00:17
    device-alias name B200M3-CH3-SERVER3-fc0 pwn 20:00:00:25:b5:c1:00:77
    device-alias name B200M3-CH3-SERVER4-fc0 pwn 20:00:00:25:b5:c1:00:08
    device-alias name B200M3-CH3-SERVER5-fc0 pwn 20:00:00:25:b5:c1:00:68
    device-alias name B200M3-CH3-SERVER6-fc0 pwn 20:00:00:25:b5:c1:00:09
    device-alias name B200M3-CH3-SERVER7-fc0 pwn 20:00:00:25:b5:c1:00:69
    device-alias name B200M3-CH3-SERVER8-fc0 pwn 20:00:00:25:b5:c1:00:a9
    device-alias name B200M3-CH4-SERVER1-fc0 pwn 20:00:00:25:b5:c1:00:a6
    device-alias name B200M3-CH4-SERVER2-fc0 pwn 20:00:00:25:b5:c1:00:37
    device-alias name B200M3-CH4-SERVER3-fc0 pwn 20:00:00:25:b5:c1:00:97
    device-alias name B200M3-CH4-SERVER4-fc0 pwn 20:00:00:25:b5:c1:00:28
    device-alias name B200M3-CH4-SERVER5-fc0 pwn 20:00:00:25:b5:c1:00:88
    device-alias name B200M3-CH4-SERVER6-fc0 pwn 20:00:00:25:b5:c1:00:29
    device-alias name B200M3-CH4-SERVER7-fc0 pwn 20:00:00:25:b5:c1:00:89

```



```

device-alias name B200M3-CH5-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:7a
device-alias name B230M2-CH1-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:af
device-alias name B230M2-CH1-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:9f
device-alias name B230M2-CH1-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:7f
device-alias name B230M2-CH1-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:3c
device-alias name B230M2-CH1-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:1d
device-alias name B230M2-CH1-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:ad
device-alias name B230M2-CH1-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:9e
device-alias name B230M2-CH1-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:66
device-alias name B230M2-CH5-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:3b
device-alias name B230M2-CH5-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:ab
device-alias name B230M2-CH5-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:9c
device-alias name B230M2-CH5-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:7d
device-alias name B230M2-CH5-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:5e
device-alias name B230M2-CH5-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:3f
device-alias name B230M2-CH6-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:3a
device-alias name B230M2-CH6-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:aa
device-alias name B230M2-CH6-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:7b
device-alias name B230M2-CH6-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:5c
device-alias name B230M2-CH6-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:3d
device-alias name B230M2-CH6-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:1e
device-alias name B230M2-CH6-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:ae
device-alias name B230M2-CH6-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:06
device-alias name B230M2-CH7-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:5a
device-alias name B230M2-CH7-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:1b
device-alias name B230M2-CH7-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:9b
device-alias name B230M2-CH7-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:7c
device-alias name B230M2-CH7-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:5d
device-alias name B230M2-CH7-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:3e
device-alias name B230M2-CH7-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:1f
device-alias name B230M2-CH7-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:26
device-alias name B230M2-CH8-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:9a
device-alias name B230M2-CH8-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:5b
device-alias name B230M2-CH8-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:1c
device-alias name B230M2-CH8-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:ac
device-alias name B230M2-CH8-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:9d
device-alias name B230M2-CH8-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:7e
device-alias name B230M2-CH8-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:5f
device-alias name B230M2-CH8-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:46

```

```
device-alias commit
```

```
fcdomain fcid database
```

```

vsan 1 wwn 50:06:01:60:46:e0:5e:0a fcid 0xa300ef dynamic
!
[ VNX7500-A0]
vsan 1 wwn 50:06:01:69:46:e0:5e:0a fcid 0xa301ef dynamic
!
[ VNX7500-B1]
vsan 1 wwn 20:49:54:7f:ee:76:d9:00 fcid 0xa30000 dynamic
vsan 1 wwn 20:4a:54:7f:ee:76:d9:00 fcid 0xa30020 dynamic
vsan 1 wwn 20:4d:54:7f:ee:76:d9:00 fcid 0xa30001 dynamic
vsan 1 wwn 20:4e:54:7f:ee:76:d9:00 fcid 0xa30021 dynamic
vsan 1 wwn 20:00:00:25:b5:c1:00:af fcid 0xa30002 dynamic
!
[ B230M2-CH1-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9f fcid 0xa30022 dynamic
!
[ B230M2-CH1-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7f fcid 0xa30003 dynamic
!
[ B230M2-CH1-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5a fcid 0xa30004 dynamic
!
[ B230M2-CH7-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7c fcid 0xa30005 dynamic
!
[ B230M2-CH7-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9b fcid 0xa30023 dynamic
!
[ B230M2-CH7-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3a fcid 0xa30006 dynamic
!
[ B230M2-CH6-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9e fcid 0xa30024 dynamic
!
[ B230M2-CH1-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:aa fcid 0xa30025 dynamic
!
[ B230M2-CH6-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7d fcid 0xa30007 dynamic
!
[ B230M2-CH5-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3e fcid 0xa30026 dynamic
!
[ B230M2-CH7-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1d fcid 0xa30008 dynamic
!
[ B230M2-CH1-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3b fcid 0xa30027 dynamic
!
[ B230M2-CH5-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3f fcid 0xa30009 dynamic
!
[ B230M2-CH5-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9c fcid 0xa30028 dynamic
!
[ B230M2-CH5-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:ad fcid 0xa30029 dynamic
!
[ B230M2-CH1-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5d fcid 0xa3000a dynamic
!
[ B230M2-CH7-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1e fcid 0xa3000b dynamic

```

```

!           [B230M2-CH6-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5c fcid 0xa3000c dynamic
!           [B230M2-CH6-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3d fcid 0xa3002a dynamic
!           [B230M2-CH6-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7b fcid 0xa3000d dynamic
!           [B230M2-CH6-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:3c fcid 0xa3002b dynamic
!           [B230M2-CH1-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7a fcid 0xa3002c dynamic
!           [B200M3-CH5-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1f fcid 0xa3002d dynamic
!           [B230M2-CH7-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:ae fcid 0xa3000e dynamic
!           [B230M2-CH6-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:ab fcid 0xa3002e dynamic
!           [B230M2-CH5-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1b fcid 0xa3000f dynamic
!           [B230M2-CH7-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9a fcid 0xa30010 dynamic
!           [B230M2-CH8-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1c fcid 0xa3002f dynamic
!           [B230M2-CH8-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:ac fcid 0xa30030 dynamic
!           [B230M2-CH8-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:9d fcid 0xa30031 dynamic
!           [B230M2-CH8-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5b fcid 0xa30011 dynamic
!           [B230M2-CH8-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:7e fcid 0xa30012 dynamic
!           [B230M2-CH8-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5f fcid 0xa30032 dynamic
!           [B230M2-CH8-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:5e fcid 0xa30013 dynamic
!           [B230M2-CH5-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:1a fcid 0xa30033 dynamic
!           [B200M3-CH2-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:a9 fcid 0xa30034 dynamic
!           [B200M3-CH3-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:17 fcid 0xa30014 dynamic
!           [B200M3-CH3-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:49 fcid 0xa30015 dynamic
!           [B200M3-CH2-SERVER7-fc0]

```

```

vsan 1 wwn 20:00:00:25:b5:c1:00:69 fcid 0xa30035 dynamic
!      [B200M3-CH3-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:c6 fcid 0xa30016 dynamic
!      [B200M3-CH2-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:68 fcid 0xa30036 dynamic
!      [B200M3-CH3-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:37 fcid 0xa30017 dynamic
!      [B200M3-CH4-SERVER2-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:a6 fcid 0xa30037 dynamic
!      [B200M3-CH4-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:77 fcid 0xa30018 dynamic
!      [B200M3-CH3-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:89 fcid 0xa30038 dynamic
!      [B200M3-CH4-SERVER7-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:88 fcid 0xa30019 dynamic
!      [B200M3-CH4-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:09 fcid 0xa30039 dynamic
!      [B200M3-CH3-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:86 fcid 0xa3001a dynamic
!      [B200M3-CH3-SERVER1-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:08 fcid 0xa3003a dynamic
!      [B200M3-CH3-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:28 fcid 0xa3001b dynamic
!      [B200M3-CH4-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:29 fcid 0xa3001c dynamic
!      [B200M3-CH4-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:97 fcid 0xa3003b dynamic
!      [B200M3-CH4-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:b7 fcid 0xa3003c dynamic
!      [B200M3-CH2-SERVER4-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:a8 fcid 0xa3003d dynamic
!      [B200M3-CH2-SERVER6-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:57 fcid 0xa3001d dynamic
!      [B200M3-CH2-SERVER3-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:48 fcid 0xa3001e dynamic
!      [B200M3-CH2-SERVER5-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:66 fcid 0xa3001f dynamic
!      [B230M2-CH1-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:46 fcid 0xa3003e dynamic
!      [B230M2-CH8-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:26 fcid 0xa3003f dynamic
!      [B230M2-CH7-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:06 fcid 0xa30040 dynamic

```

```
! [B230M2-CH6-SERVER8-fc0]
vsan 1 wwn 20:00:00:25:b5:c1:00:95 fcid 0xa30060 dynamic
vsan 1 wwn 20:00:00:25:b5:c1:00:b5 fcid 0xa30041 dynamic

interface Vlan1

interface port-channel5
  description PC-TO-6248-A
  switchport mode trunk
  vpc 5

interface port-channel7
  description PC-TO-6248-B
  switchport mode trunk
  vpc 7

interface port-channel10
  description VPC-PEER-2-PORT-CHANNEL
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link

interface port-channel17
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
  service-policy type qos input pm-qos-vnx
  vpc 17

interface port-channel18
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
  service-policy type qos input pm-qos-vnx
  vpc 18

interface port-channel19
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
  service-policy type qos input pm-qos-vnx
  vpc 19

interface port-channel20
  switchport mode trunk
```

```
switchport trunk allowed vlan 1,800-804,850-854
service-policy type qos input pm-qos-vnx
vpc 20

interface fc2/1
no shutdown

interface fc2/2
no shutdown

interface fc2/3
no shutdown

interface fc2/4
no shutdown

interface fc2/5
no shutdown

interface fc2/6
no shutdown

interface fc2/7
no shutdown

interface fc2/8
no shutdown

interface Ethernet1/1
description uplink2 R1E05U42-5820x-SW1 P1
switchport mode trunk
switchport trunk allowed vlan 800-804,850-854

interface Ethernet1/2
switchport mode trunk
switchport trunk allowed vlan 329

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5
```

```
interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15
  description VPC-LINK-TO-N5548UP-A
  switchport mode trunk
  channel-group 10 mode active

interface Ethernet1/16
  description VPC-LINK-TO-N5548UP-A
  switchport mode trunk
  channel-group 10 mode active

interface Ethernet1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
  spanning-tree port type edge
  channel-group 17 mode active

interface Ethernet1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
  spanning-tree port type edge
  channel-group 18 mode active

interface Ethernet1/19
  switchport mode trunk
  switchport trunk allowed vlan 1,800-804,850-854
```

```
spanning-tree port type edge
channel-group 19 mode active

interface Ethernet1/20
switchport mode trunk
switchport trunk allowed vlan 1,800-804,850-854
spanning-tree port type edge
channel-group 20 mode active

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29
switchport mode trunk
channel-group 5 mode active

interface Ethernet1/30
switchport mode trunk
channel-group 5 mode active

interface Ethernet1/31
switchport mode trunk
channel-group 7 mode active

interface Ethernet1/32
switchport mode trunk
channel-group 7 mode active

interface Ethernet2/1
```



```

interface Ethernet2/2

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6

interface Ethernet2/7

interface Ethernet2/8

interface mgmt0
    ip address 10.218.164.120/21
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1.bin
ip route 0.0.0.0/0 10.218.255.73
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
interface fc2/7
interface fc2/8
!Full Zone Database Section for vsan 1
zone name B230M2-CH1-SERVER1-FC0 vsan 1
    member pwn 20:00:00:25:b5:c1:00:af
!
    [B230M2-CH1-SERVER1-fc0]
    member pwn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
    member pwn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]

zone name B230M2-CH1-SERVER2-FC0 vsan 1
    member pwn 20:00:00:25:b5:c1:00:9f
!
    [B230M2-CH1-SERVER2-fc0]
    member pwn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]

```

```

        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B230M2-CH1-SERVER3-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:7f
!
        [B230M2-CH1-SERVER3-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B230M2-CH1-SERVER4-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:3c
!
        [B230M2-CH1-SERVER4-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B230M2-CH1-SERVER5-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:1d
!
        [B230M2-CH1-SERVER5-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B230M2-CH1-SERVER6-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:ad
!
        [B230M2-CH1-SERVER6-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B230M2-CH1-SERVER7-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:9e
!
        [B230M2-CH1-SERVER7-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER1-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:1a
!
  [B200M3-CH2-SERVER1-fc0]
  member pwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER2-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:c6
!
  [B200M3-CH2-SERVER8-fc0]
  member pwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER3-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:57
!
  [B200M3-CH2-SERVER3-fc0]
  member pwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER4-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:b7
!
  [B200M3-CH2-SERVER4-fc0]
  member pwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER5-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:48
!
  [B200M3-CH2-SERVER5-fc0]
  member pwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER6-FC0 vsan 1
  member pwn 20:00:00:25:b5:c1:00:b5
  member pwn 50:06:01:60:46:e0:5e:0a

```

```

!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH2-SERVER7-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:49
!                               [B200M3-CH2-SERVER7-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH3-SERVER1-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:86
!                               [B200M3-CH3-SERVER1-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH3-SERVER2-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:17
!                               [B200M3-CH3-SERVER2-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH3-SERVER3-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:77
!                               [B200M3-CH3-SERVER3-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH3-SERVER4-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:08
!                               [B200M3-CH3-SERVER4-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

```

```

zone name B200M3-CH3-SERVER5-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:68
!
  [B200M3-CH3-SERVER5-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH3-SERVER6-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:09
!
  [B200M3-CH3-SERVER6-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH3-SERVER7-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:69
!
  [B200M3-CH3-SERVER7-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH3-SERVER8-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:a9
!
  [B200M3-CH3-SERVER8-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH4-SERVER1-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:a6
!
  [B200M3-CH4-SERVER1-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!
  [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!
  [VNX7500-B1]

```

```

zone name B200M3-CH4-SERVER2-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:37

```

```

!                               [B200M3-CH4-SERVER2-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH4-SERVER3-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:97
!                               [B200M3-CH4-SERVER3-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH4-SERVER4-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:28
!                               [B200M3-CH4-SERVER4-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH4-SERVER5-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:88
!                               [B200M3-CH4-SERVER5-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
                               [VNX7500-B1]

zone name B200M3-CH4-SERVER6-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:29
!                               [B200M3-CH4-SERVER6-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B200M3-CH4-SERVER7-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:89
!                               [B200M3-CH4-SERVER7-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]

```

```

        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]

zone name B230M2-CH5-SERVER1-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:7a
!
        [B200M3-CH5-SERVER1-fc0]

zone name B230M2-CH5-SERVER2-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:3b
!
        [B230M2-CH5-SERVER2-fc0]

zone name B230M2-CH5-SERVER3-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:ab
!
        [B230M2-CH5-SERVER3-fc0]

zone name B230M2-CH5-SERVER4-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:9c
!
        [B230M2-CH5-SERVER4-fc0]

zone name B230M2-CH5-SERVER5-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:7d
!
        [B230M2-CH5-SERVER5-fc0]

```

```

zone name B230M2-CH5-SERVER6-FC0 vsan 1
  member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:5e
!
    [B230M2-CH5-SERVER6-fc0]
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]

```

```

zone name B230M2-CH5-SERVER7-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
  member pwnn 20:00:00:25:b5:c1:00:3f
!
    [B230M2-CH5-SERVER7-fc0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]

```

```

zone name B230M2-CH6-SERVER1-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:3a
!
    [B230M2-CH6-SERVER1-fc0]

```

```

zone name B230M2-CH6-SERVER2-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:aa
!
    [B230M2-CH6-SERVER2-fc0]

```

```

zone name B230M2-CH6-SERVER3-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:7b
!
    [B230M2-CH6-SERVER3-fc0]

```

```

zone name B230M2-CH6-SERVER4-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]

```



```

        member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:5c
!           [B230M2-CH6-SERVER4-fc0]

zone name B230M2-CH6-SERVER5-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!           [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:3d
!           [B230M2-CH6-SERVER5-fc0]

zone name B230M2-CH6-SERVER6-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!           [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:1e
!           [B230M2-CH6-SERVER6-fc0]

zone name B230M2-CH6-SERVER7-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!           [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:ae
!           [B230M2-CH6-SERVER7-fc0]

zone name B230M2-CH7-SERVER1-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!           [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:5a
!           [B230M2-CH7-SERVER1-fc0]

zone name B230M2-CH7-SERVER2-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!           [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!           [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:1b

```

```

!                               [B230M2-CH7-SERVER2-fc0]

zone name B230M2-CH7-SERVER3-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:9b
!                               [B230M2-CH7-SERVER3-fc0]

zone name B230M2-CH7-SERVER4-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:7c
!                               [B230M2-CH7-SERVER4-fc0]

zone name B230M2-CH7-SERVER5-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:5d
!                               [B230M2-CH7-SERVER5-fc0]

zone name B230M2-CH7-SERVER6-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:3e
!                               [B230M2-CH7-SERVER6-fc0]

zone name B230M2-CH7-SERVER7-FC0 vsan 1
  member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
  member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
  member pwnn 20:00:00:25:b5:c1:00:1f
!                               [B230M2-CH7-SERVER7-fc0]

zone name B230M2-CH8-SERVER1-FC0 vsan 1

```

```

        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:9a
!
        [B230M2-CH8-SERVER1-fc0]

zone name B230M2-CH8-SERVER2-FC0 vsan 1
        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:5b
!
        [B230M2-CH8-SERVER2-fc0]

zone name B230M2-CH8-SERVER3-FC0 vsan 1
        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:1c
!
        [B230M2-CH8-SERVER3-fc0]

zone name B230M2-CH8-SERVER4-FC0 vsan 1
        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:ac
!
        [B230M2-CH8-SERVER4-fc0]

zone name B230M2-CH8-SERVER5-FC0 vsan 1
        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a
!
        [VNX7500-B1]
        member pwnn 20:00:00:25:b5:c1:00:9d
!
        [B230M2-CH8-SERVER5-fc0]

zone name B230M2-CH8-SERVER6-FC0 vsan 1
        member pwnn 50:06:01:60:46:e0:5e:0a
!
        [VNX7500-A0]
        member pwnn 50:06:01:69:46:e0:5e:0a

```

```

!                               [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:7e
!                               [B230M2-CH8-SERVER6-fc0]

zone name B230M2-CH8-SERVER7-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:5f
!                               [B230M2-CH8-SERVER7-fc0]

zone name B230M2-CH8-SERVER8-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:46
!                               [B230M2-CH8-SERVER8-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B230M2-CH1-SERVER8-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:66
!                               [B230M2-CH1-SERVER8-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B230M2-CH6-SERVER8-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:06
!                               [B230M2-CH6-SERVER8-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

zone name B230M2-CH7-SERVER8-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:26
!                               [B230M2-CH7-SERVER8-fc0]
    member pwnn 50:06:01:60:46:e0:5e:0a
!                               [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!                               [VNX7500-B1]

```

```

zone name B200M3-CH2-SERVER8-FC0 vsan 1
    member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]
    member pwnn 20:00:00:25:b5:c1:00:c6
!
    [B200M3-CH2-SERVER8-fc0]

```

```

zone name B200M3-CH4-SERVER8-FC0 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:95
    member pwnn 50:06:01:60:46:e0:5e:0a
!
    [VNX7500-A0]
    member pwnn 50:06:01:69:46:e0:5e:0a
!
    [VNX7500-B1]

```

```

zoneset name DC-UCS-POD-A vsan 1
    member B230M2-CH1-SERVER1-FC0
    member B230M2-CH1-SERVER2-FC0
    member B230M2-CH1-SERVER3-FC0
    member B230M2-CH1-SERVER4-FC0
    member B230M2-CH1-SERVER5-FC0
    member B230M2-CH1-SERVER6-FC0
    member B230M2-CH1-SERVER7-FC0
    member B200M3-CH2-SERVER1-FC0
    member B200M3-CH2-SERVER2-FC0
    member B200M3-CH2-SERVER3-FC0
    member B200M3-CH2-SERVER4-FC0
    member B200M3-CH2-SERVER5-FC0
    member B200M3-CH2-SERVER6-FC0
    member B200M3-CH2-SERVER7-FC0
    member B200M3-CH3-SERVER1-FC0
    member B200M3-CH3-SERVER2-FC0
    member B200M3-CH3-SERVER3-FC0
    member B200M3-CH3-SERVER4-FC0
    member B200M3-CH3-SERVER5-FC0
    member B200M3-CH3-SERVER6-FC0
    member B200M3-CH3-SERVER7-FC0
    member B200M3-CH3-SERVER8-FC0
    member B200M3-CH4-SERVER1-FC0
    member B200M3-CH4-SERVER2-FC0
    member B200M3-CH4-SERVER3-FC0
    member B200M3-CH4-SERVER4-FC0

```

```
member B200M3-CH4-SERVER5-FC0
member B200M3-CH4-SERVER6-FC0
member B200M3-CH4-SERVER7-FC0
member B230M2-CH5-SERVER1-FC0
member B230M2-CH5-SERVER2-FC0
member B230M2-CH5-SERVER3-FC0
member B230M2-CH5-SERVER4-FC0
member B230M2-CH5-SERVER5-FC0
member B230M2-CH5-SERVER6-FC0
member B230M2-CH5-SERVER7-FC0
member B230M2-CH6-SERVER1-FC0
member B230M2-CH6-SERVER2-FC0
member B230M2-CH6-SERVER3-FC0
member B230M2-CH6-SERVER4-FC0
member B230M2-CH6-SERVER5-FC0
member B230M2-CH6-SERVER6-FC0
member B230M2-CH6-SERVER7-FC0
member B230M2-CH7-SERVER1-FC0
member B230M2-CH7-SERVER2-FC0
member B230M2-CH7-SERVER3-FC0
member B230M2-CH7-SERVER4-FC0
member B230M2-CH7-SERVER5-FC0
member B230M2-CH7-SERVER6-FC0
member B230M2-CH7-SERVER7-FC0
member B230M2-CH8-SERVER1-FC0
member B230M2-CH8-SERVER2-FC0
member B230M2-CH8-SERVER3-FC0
member B230M2-CH8-SERVER4-FC0
member B230M2-CH8-SERVER5-FC0
member B230M2-CH8-SERVER6-FC0
member B230M2-CH8-SERVER7-FC0
member B230M2-CH8-SERVER8-FC0
member B230M2-CH1-SERVER8-FC0
member B230M2-CH6-SERVER8-FC0
member B230M2-CH7-SERVER8-FC0
member B200M3-CH2-SERVER8-FC0
member B200M3-CH4-SERVER8-FC0
```

```
zoneset name CD-UCS-POD-A vsan 1
```

```
zoneset activate name DC-UCS-POD-A vsan 1
```

## A.2 N5548UP-B Configuration

```
version 5.2(1)N1(1)
feature fcoe
logging level feature-mgr 0
hostname N5548UP-B
feature npiv
feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature vpc
feature lldp
username admin password 5 $1$Nj5oguWP$n4UckqZNp6GYRe8OQTaTP1 role
network-adminno password strength-check

banner motd #Nexus 5000 Switch
#

ip domain-lookup
logging event link-status default
ip access-list acl-vnx
  10 permit ip any any
class-map type qos class-fcoe
class-map type qos match-all cm-qos-vnx
  match access-group name acl-vnx
class-map type qos match-all cm-qos-cos5
  match cos 5
class-map type queuing class-fcoe
  match qos-group 1
class-map type queuing cm-que-qosgrp5
  match qos-group 5
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 pm-qos-vnx
  class cm-qos-vnx
    set qos-group 5
  class class-default
policy-map type qos pm-qos-global
  class cm-qos-cos5
```

```

        set qos-group 5
    class class-default
class-map type network-qos class-fcoe
    match qos-group 1
class-map type network-qos cm-nq-grp5
    match qos-group 5
class-map type network-qos class-all-flood
    match qos-group 2
class-map type network-qos class-ip-multicast
    match qos-group 2
policy-map type network-qos pm-nq-global
    class type network-qos cm-nq-grp5
        mtu 9216
        set cos 5
    class type network-qos class-default
        multicast-optimize
system qos
    service-policy type qos input pm-qos-global
    service-policy type network-qos pm-nq-global
snmp-server user admin network-admin auth md5
0x7e352b106a1735fa03b1e8d75e42c608 priv

0x7e352b106a1735fa03b1e8d75e42c608 localizedkey
vrf context management
    ip route 0.0.0.0/0 10.218.160.1
vlan 1
vlan 800
    name ML-VDA
vlan 801
    name ML-DC-VM-MGMT
vlan 802
    name ML-DC-VMMOTION
vlan 803
    name ML-DC-INF
vlan 804
    name ML-DC-STRG
vlan 850
    name ML_BR-MGMT
vlan 851
    name ML_Launcher-Inf
vlan 852
    name ML_LauncherA
vlan 853

```



```

    name ML_LauncherB
vlan 854
    name ML_LauncherC
vlan 900
    name ML-N1KV_CTRL
vlan 901
vpc domain 10
    role priority 1000
    peer-keepalive destination 10.218.164.120
port-profile default max-ports 512
device-alias database
    device-alias name VNX7500-A1 pwnn 50:06:01:61:46:e0:5e:0a
    device-alias name VNX7500-B0 pwnn 50:06:01:68:46:e0:5e:0a
    device-alias name B200M3-CH2-SERVER1-fc1 pwnn 20:00:00:25:b5:c1:00:0a
    device-alias name B200M3-CH2-SERVER3-fc1 pwnn 20:00:00:25:b5:c1:00:67
    device-alias name B200M3-CH2-SERVER4-fc1 pwnn 20:00:00:25:b5:c1:00:c7
    device-alias name B200M3-CH2-SERVER5-fc1 pwnn 20:00:00:25:b5:c1:00:58
    device-alias name B200M3-CH2-SERVER6-fc1 pwnn 20:00:00:25:b5:c1:00:b8
    device-alias name B200M3-CH2-SERVER7-fc1 pwnn 20:00:00:25:b5:c1:00:59
    device-alias name B200M3-CH2-SERVER8-fc1 pwnn 20:00:00:25:b5:c1:00:07
    device-alias name B200M3-CH3-SERVER1-fc1 pwnn 20:00:00:25:b5:c1:00:96
    device-alias name B200M3-CH3-SERVER2-fc1 pwnn 20:00:00:25:b5:c1:00:27
    device-alias name B200M3-CH3-SERVER3-fc1 pwnn 20:00:00:25:b5:c1:00:87
    device-alias name B200M3-CH3-SERVER4-fc1 pwnn 20:00:00:25:b5:c1:00:18
    device-alias name B200M3-CH3-SERVER5-fc1 pwnn 20:00:00:25:b5:c1:00:78
    device-alias name B200M3-CH3-SERVER6-fc1 pwnn 20:00:00:25:b5:c1:00:19
    device-alias name B200M3-CH3-SERVER7-fc1 pwnn 20:00:00:25:b5:c1:00:79
    device-alias name B200M3-CH3-SERVER8-fc1 pwnn 20:00:00:25:b5:c1:00:b9
    device-alias name B200M3-CH4-SERVER1-fc1 pwnn 20:00:00:25:b5:c1:00:b6
    device-alias name B200M3-CH4-SERVER2-fc1 pwnn 20:00:00:25:b5:c1:00:47
    device-alias name B200M3-CH4-SERVER3-fc1 pwnn 20:00:00:25:b5:c1:00:a7
    device-alias name B200M3-CH4-SERVER4-fc1 pwnn 20:00:00:25:b5:c1:00:38
    device-alias name B200M3-CH4-SERVER5-fc1 pwnn 20:00:00:25:b5:c1:00:98
    device-alias name B200M3-CH4-SERVER6-fc1 pwnn 20:00:00:25:b5:c1:00:39
    device-alias name B200M3-CH4-SERVER7-fc1 pwnn 20:00:00:25:b5:c1:00:99
    device-alias name B230M2-CH1-SERVER1-fc1 pwnn 20:00:00:25:b5:c1:00:bf
    device-alias name B230M2-CH1-SERVER2-fc1 pwnn 20:00:00:25:b5:c1:00:8f
    device-alias name B230M2-CH1-SERVER3-fc1 pwnn 20:00:00:25:b5:c1:00:6f
    device-alias name B230M2-CH1-SERVER4-fc1 pwnn 20:00:00:25:b5:c1:00:2c
    device-alias name B230M2-CH1-SERVER5-fc1 pwnn 20:00:00:25:b5:c1:00:0d
    device-alias name B230M2-CH1-SERVER6-fc1 pwnn 20:00:00:25:b5:c1:00:bd
    device-alias name B230M2-CH1-SERVER7-fc1 pwnn 20:00:00:25:b5:c1:00:8e
    device-alias name B230M2-CH1-SERVER8-fc1 pwnn 20:00:00:25:b5:c1:00:76

```

```

device-alias name B230M2-CH5-SERVER1-fc1 pwn 20:00:00:25:b5:c1:00:6a
device-alias name B230M2-CH5-SERVER2-fc1 pwn 20:00:00:25:b5:c1:00:2b
device-alias name B230M2-CH5-SERVER3-fc1 pwn 20:00:00:25:b5:c1:00:bb
device-alias name B230M2-CH5-SERVER4-fc1 pwn 20:00:00:25:b5:c1:00:8c
device-alias name B230M2-CH5-SERVER5-fc1 pwn 20:00:00:25:b5:c1:00:6d
device-alias name B230M2-CH5-SERVER6-fc1 pwn 20:00:00:25:b5:c1:00:4e
device-alias name B230M2-CH5-SERVER7-fc1 pwn 20:00:00:25:b5:c1:00:2f
device-alias name B230M2-CH6-SERVER1-fc1 pwn 20:00:00:25:b5:c1:00:2a
device-alias name B230M2-CH6-SERVER2-fc1 pwn 20:00:00:25:b5:c1:00:ba
device-alias name B230M2-CH6-SERVER3-fc1 pwn 20:00:00:25:b5:c1:00:6b
device-alias name B230M2-CH6-SERVER4-fc1 pwn 20:00:00:25:b5:c1:00:4c
device-alias name B230M2-CH6-SERVER5-fc1 pwn 20:00:00:25:b5:c1:00:2d
device-alias name B230M2-CH6-SERVER6-fc1 pwn 20:00:00:25:b5:c1:00:0e
device-alias name B230M2-CH6-SERVER7-fc1 pwn 20:00:00:25:b5:c1:00:be
device-alias name B230M2-CH6-SERVER8-fc1 pwn 20:00:00:25:b5:c1:00:16
device-alias name B230M2-CH7-SERVER1-fc1 pwn 20:00:00:25:b5:c1:00:4a
device-alias name B230M2-CH7-SERVER2-fc1 pwn 20:00:00:25:b5:c1:00:0b
device-alias name B230M2-CH7-SERVER3-fc1 pwn 20:00:00:25:b5:c1:00:8b
device-alias name B230M2-CH7-SERVER4-fc1 pwn 20:00:00:25:b5:c1:00:6c
device-alias name B230M2-CH7-SERVER5-fc1 pwn 20:00:00:25:b5:c1:00:4d
device-alias name B230M2-CH7-SERVER6-fc1 pwn 20:00:00:25:b5:c1:00:2e
device-alias name B230M2-CH7-SERVER7-fc1 pwn 20:00:00:25:b5:c1:00:0f
device-alias name B230M2-CH7-SERVER8-fc1 pwn 20:00:00:25:b5:c1:00:36
device-alias name B230M2-CH8-SERVER1-fc1 pwn 20:00:00:25:b5:c1:00:8a
device-alias name B230M2-CH8-SERVER2-fc1 pwn 20:00:00:25:b5:c1:00:4b
device-alias name B230M2-CH8-SERVER3-fc1 pwn 20:00:00:25:b5:c1:00:0c
device-alias name B230M2-CH8-SERVER4-fc1 pwn 20:00:00:25:b5:c1:00:bc
device-alias name B230M2-CH8-SERVER5-fc1 pwn 20:00:00:25:b5:c1:00:8d
device-alias name B230M2-CH8-SERVER6-fc1 pwn 20:00:00:25:b5:c1:00:6e
device-alias name B230M2-CH8-SERVER7-fc1 pwn 20:00:00:25:b5:c1:00:4f
device-alias name B230M2-CH8-SERVER8-fc1 pwn 20:00:00:25:b5:c1:00:56

device-alias commit

fcdomain fcid database
    vsan 1 wwn 50:06:01:61:46:e0:5e:0a fcid 0xad00ef dynamic
    !
    [VNX7500-A1]
    vsan 1 wwn 50:06:01:68:46:e0:5e:0a fcid 0xad01ef dynamic
    !
    [VNX7500-B0]
    vsan 1 wwn 20:49:54:7f:ee:76:ce:40 fcid 0xad0000 dynamic
    vsan 1 wwn 20:4a:54:7f:ee:76:ce:40 fcid 0xad0020 dynamic
    vsan 1 wwn 20:4d:54:7f:ee:76:ce:40 fcid 0xad0001 dynamic
    vsan 1 wwn 20:4e:54:7f:ee:76:ce:40 fcid 0xad0021 dynamic

```

```

vsan 1 wwn 20:00:00:25:b5:c1:00:bf fcid 0xad0002 dynamic
!
[B230M2-CH1-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8f fcid 0xad0022 dynamic
!
[B230M2-CH1-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6f fcid 0xad0003 dynamic
!
[B230M2-CH1-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4a fcid 0xad0004 dynamic
!
[B230M2-CH7-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6c fcid 0xad0005 dynamic
!
[B230M2-CH7-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8b fcid 0xad0023 dynamic
!
[B230M2-CH7-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2a fcid 0xad0024 dynamic
!
[B230M2-CH6-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8e fcid 0xad0025 dynamic
!
[B230M2-CH1-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:ba fcid 0xad0026 dynamic
!
[B230M2-CH6-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6d fcid 0xad0027 dynamic
!
[B230M2-CH5-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2e fcid 0xad0006 dynamic
!
[B230M2-CH7-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0d fcid 0xad0007 dynamic
!
[B230M2-CH1-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2b fcid 0xad0008 dynamic
!
[B230M2-CH5-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2f fcid 0xad0028 dynamic
!
[B230M2-CH5-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8c fcid 0xad0009 dynamic
!
[B230M2-CH5-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:bd fcid 0xad0029 dynamic
!
[B230M2-CH1-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4d fcid 0xad002a dynamic
!
[B230M2-CH7-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0e fcid 0xad002b dynamic
!
[B230M2-CH6-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4c fcid 0xad000a dynamic
!
[B230M2-CH6-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2d fcid 0xad000b dynamic
!
[B230M2-CH6-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6b fcid 0xad002c dynamic
!
[B230M2-CH6-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:2c fcid 0xad000c dynamic

```

```

!           [B230M2-CH1-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6a fcid 0xad000d dynamic
!           [B230M2-CH5-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0f fcid 0xad002d dynamic
!           [B230M2-CH7-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:be fcid 0xad002e dynamic
!           [B230M2-CH6-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:bb fcid 0xad000e dynamic
!           [B230M2-CH5-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0b fcid 0xad000f dynamic
!           [B230M2-CH7-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4e fcid 0xad002f dynamic
!           [B230M2-CH5-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0a fcid 0xad0010 dynamic
!           [B200M3-CH2-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8a fcid 0xad0030 dynamic
!           [B230M2-CH8-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:bc fcid 0xad0011 dynamic
!           [B230M2-CH8-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:0c fcid 0xad0031 dynamic
!           [B230M2-CH8-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4b fcid 0xad0012 dynamic
!           [B230M2-CH8-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:4f fcid 0xad0032 dynamic
!           [B230M2-CH8-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:8d fcid 0xad0013 dynamic
!           [B230M2-CH8-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:6e fcid 0xad0033 dynamic
!           [B230M2-CH8-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:b9 fcid 0xad0034 dynamic
!           [B200M3-CH3-SERVER8-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:27 fcid 0xad0035 dynamic
!           [B200M3-CH3-SERVER2-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:59 fcid 0xad0014 dynamic
!           [B200M3-CH2-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:79 fcid 0xad0015 dynamic
!           [B200M3-CH3-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:07 fcid 0xad0016 dynamic
!           [B200M3-CH2-SERVER8-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:78 fcid 0xad0036 dynamic
!           [B200M3-CH3-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:47 fcid 0xad0017 dynamic
!           [B200M3-CH4-SERVER2-fc1]

```

```

    vsan 1 wwn 20:00:00:25:b5:c1:00:b6 fcid 0xad0037 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:87 fcid 0xad0018 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:99 fcid 0xad0038 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:98 fcid 0xad0019 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:19 fcid 0xad0039 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:96 fcid 0xad001a dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:18 fcid 0xad003a dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:38 fcid 0xad001b dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:39 fcid 0xad003b dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:a7 fcid 0xad001c dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:c7 fcid 0xad003c dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:b8 fcid 0xad003d dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:67 fcid 0xad001d dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:58 fcid 0xad001e dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:76 fcid 0xad001f dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:56 fcid 0xad003e dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:36 fcid 0xad003f dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:16 fcid 0xad0040 dynamic
    !
    vsan 1 wwn 20:00:00:25:b5:c1:00:a5 fcid 0xad0041 dynamic
    vsan 1 wwn 20:00:00:25:b5:c1:00:c5 fcid 0xad0042 dynamic

interface Vlan1

interface port-channel5

```

```
description PC-TO-6248-A
switchport mode trunk
vpc 5

interface port-channel7
description PC-TO-6248-B
switchport mode trunk
vpc 7

interface port-channel10
description VPC-PEER-2-PORT-CHANNEL
switchport mode trunk
spanning-tree port type network
vpc peer-link

interface port-channel17
switchport mode trunk
switchport trunk allowed vlan 1,800-804,850-854
service-policy type qos input pm-qos-vnx
vpc 17

interface port-channel18
switchport mode trunk
switchport trunk allowed vlan 1,800-804,850-854
service-policy type qos input pm-qos-vnx
vpc 18

interface port-channel19
switchport mode trunk
switchport trunk allowed vlan 1,800-804,850-854
service-policy type qos input pm-qos-vnx
vpc 19

interface port-channel20
switchport mode trunk
switchport trunk allowed vlan 1,800-804,850-854
service-policy type qos input pm-qos-vnx
vpc 20

interface fc2/1
no shutdown

interface fc2/2
```

```
no shutdown

interface fc2/3
no shutdown

interface fc2/4
no shutdown

interface fc2/5
no shutdown

interface fc2/6
no shutdown

interface fc2/7
no shutdown

interface fc2/8
no shutdown

interface Ethernet1/1
description uplink2 R1E05U42-5820x-SW1 P2
switchport trunk allowed vlan 800-804,850-854

interface Ethernet1/2

interface Ethernet1/3
description R1E10U06-Altiris
switchport access vlan 801

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10
```

```
interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15
    description VPC-LINK-TO-N5548UP-B
    switchport mode trunk
    channel-group 10 mode active

interface Ethernet1/16
    description VPC-LINK-TO-N5548UP-B
    switchport mode trunk
    channel-group 10 mode active

interface Ethernet1/17
    switchport mode trunk
    switchport trunk allowed vlan 1,800-804,850-854
    spanning-tree port type edge
    channel-group 17 mode active

interface Ethernet1/18
    switchport mode trunk
    switchport trunk allowed vlan 1,800-804,850-854
    spanning-tree port type edge
    channel-group 18 mode active

interface Ethernet1/19
    switchport mode trunk
    switchport trunk allowed vlan 1,800-804,850-854
    spanning-tree port type edge
    channel-group 19 mode active

interface Ethernet1/20
    switchport mode trunk
    switchport trunk allowed vlan 1,800-804,850-854
    spanning-tree port type edge
    channel-group 20 mode active
```



```
interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29
    switchport mode trunk
    channel-group 7 mode active

interface Ethernet1/30
    switchport mode trunk
    channel-group 7 mode active

interface Ethernet1/31
    switchport mode trunk
    channel-group 5 mode active

interface Ethernet1/32
    switchport mode trunk
    channel-group 5 mode active

interface Ethernet2/1

interface Ethernet2/2

interface Ethernet2/3

interface Ethernet2/4

interface Ethernet2/5

interface Ethernet2/6
```

```

interface Ethernet2/7
    switchport mode trunk

interface Ethernet2/8
    switchport mode trunk

interface mgmt0
    ip address 10.218.164.121/21
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.2.1.N1.1.bin
boot system bootflash:/n5000-uk9.5.2.1.N1.1.bin
ip route 0.0.0.0/0 10.218.255.73
interface fc2/1
interface fc2/2
interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
interface fc2/7
interface fc2/8
!Full Zone Database Section for vsan 1
zone name B230M2-CH1-SERVER1-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:bf
!
    [B230M2-CH1-SERVER1-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B230M2-CH1-SERVER2-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:8f
!
    [B230M2-CH1-SERVER2-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B230M2-CH1-SERVER3-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:6f
!
    [B230M2-CH1-SERVER3-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a

```

```

!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH1-SERVER4-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:2c
!                               [B230M2-CH1-SERVER4-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH1-SERVER5-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:0d
!                               [B230M2-CH1-SERVER5-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH1-SERVER6-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:bd
!                               [B230M2-CH1-SERVER6-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH1-SERVER7-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:8e
!                               [B230M2-CH1-SERVER7-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH2-SERVER1-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:0a
!                               [B200M3-CH2-SERVER1-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER2-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:07
!
    [B200M3-CH2-SERVER8-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER3-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:67
!
    [B200M3-CH2-SERVER3-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER4-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:c7
!
    [B200M3-CH2-SERVER4-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER5-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:58
!
    [B200M3-CH2-SERVER5-fc1]
    member pwn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER6-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:c5
    member pwn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

```

```

zone name B200M3-CH2-SERVER7-FC1 vsan 1
    member pwn 20:00:00:25:b5:c1:00:59
!
    [B200M3-CH2-SERVER7-fc1]

```

```

        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH3-SERVER1-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:96
!           [B200M3-CH3-SERVER1-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH3-SERVER2-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:27
!           [B200M3-CH3-SERVER2-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH3-SERVER3-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:87
!           [B200M3-CH3-SERVER3-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH3-SERVER4-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:18
!           [B200M3-CH3-SERVER4-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH3-SERVER5-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:78
!           [B200M3-CH3-SERVER5-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a

```

```

!                               [VNX7500-B0]

zone name B200M3-CH3-SERVER6-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:19
!                               [B200M3-CH3-SERVER6-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH3-SERVER7-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:79
!                               [B200M3-CH3-SERVER7-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH3-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:b9
!                               [B200M3-CH3-SERVER8-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH4-SERVER1-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:b6
!                               [B200M3-CH4-SERVER1-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH4-SERVER2-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:47
!                               [B200M3-CH4-SERVER2-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B200M3-CH4-SERVER3-FC1 vsan 1

```

```

        member pwnn 20:00:00:25:b5:c1:00:a7
!           [B200M3-CH4-SERVER3-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH4-SERVER4-FC1 vsan 1
        member pwnn 20:00:00:25:b5:c1:00:38
!           [B200M3-CH4-SERVER4-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH4-SERVER5-FC1 vsan 1
        member pwnn 20:00:00:25:b5:c1:00:98
!           [B200M3-CH4-SERVER5-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH4-SERVER6-FC1 vsan 1
        member pwnn 20:00:00:25:b5:c1:00:39
!           [B200M3-CH4-SERVER6-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B200M3-CH4-SERVER7-FC1 vsan 1
        member pwnn 20:00:00:25:b5:c1:00:99
!           [B200M3-CH4-SERVER7-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a
!           [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!           [VNX7500-B0]

zone name B230M2-CH5-SERVER1-FC1 vsan 1
        member pwnn 20:00:00:25:b5:c1:00:6a
!           [B230M2-CH5-SERVER1-fc1]
        member pwnn 50:06:01:61:46:e0:5e:0a

```

```

!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH5-SERVER2-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:2b
!                               [B230M2-CH5-SERVER2-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH5-SERVER3-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:bb
!                               [B230M2-CH5-SERVER3-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH5-SERVER4-FC1 vsan 1
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]
    member pwnn 20:00:00:25:b5:c1:00:8c
!                               [B230M2-CH5-SERVER4-fc1]

zone name B230M2-CH5-SERVER5-FC1 vsan 1
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]
    member pwnn 20:00:00:25:b5:c1:00:6d
!                               [B230M2-CH5-SERVER5-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]

zone name B230M2-CH5-SERVER6-FC1 vsan 1
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 20:00:00:25:b5:c1:00:4e
!                               [B230M2-CH5-SERVER6-fc1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

```



```

zone name B230M2-CH5-SERVER7-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:2f
!
  [B230M2-CH5-SERVER7-fc1]
  member pwwn 50:06:01:61:46:e0:5e:0a
!
  [VNX7500-A1]
  member pwwn 50:06:01:68:46:e0:5e:0a
!
  [VNX7500-B0]

```

```

zone name B230M2-CH6-SERVER1-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:2a
!
  [B230M2-CH6-SERVER1-fc1]
  member pwwn 50:06:01:61:46:e0:5e:0a
!
  [VNX7500-A1]
  member pwwn 50:06:01:68:46:e0:5e:0a
!
  [VNX7500-B0]

```

```

zone name B230M2-CH6-SERVER2-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:ba
!
  [B230M2-CH6-SERVER2-fc1]
  member pwwn 50:06:01:61:46:e0:5e:0a
!
  [VNX7500-A1]
  member pwwn 50:06:01:68:46:e0:5e:0a
!
  [VNX7500-B0]

```

```

zone name B230M2-CH6-SERVER3-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:6b
!
  [B230M2-CH6-SERVER3-fc1]
  member pwwn 50:06:01:61:46:e0:5e:0a
!
  [VNX7500-A1]
  member pwwn 50:06:01:68:46:e0:5e:0a
!
  [VNX7500-B0]

```

```

zone name B230M2-CH6-SERVER4-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:4c
!
  [B230M2-CH6-SERVER4-fc1]
  member pwwn 50:06:01:61:46:e0:5e:0a
!
  [VNX7500-A1]
  member pwwn 50:06:01:68:46:e0:5e:0a
!
  [VNX7500-B0]

```

```

zone name B230M2-CH6-SERVER5-FC1 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:2d

```

```

!                               [B230M2-CH6-SERVER5-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH6-SERVER6-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:0e
!                               [B230M2-CH6-SERVER6-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH6-SERVER7-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:be
!                               [B230M2-CH6-SERVER7-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH7-SERVER1-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:4a
!                               [B230M2-CH7-SERVER1-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH7-SERVER2-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:0b
!                               [B230M2-CH7-SERVER2-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zone name B230M2-CH7-SERVER3-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:8b
!                               [B230M2-CH7-SERVER3-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]

```

```

        member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH7-SERVER4-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:6c
!
        [B230M2-CH7-SERVER4-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH7-SERVER5-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:4d
!
        [B230M2-CH7-SERVER5-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH7-SERVER6-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:2e
!
        [B230M2-CH7-SERVER6-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH7-SERVER7-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:0f
!
        [B230M2-CH7-SERVER7-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH8-SERVER1-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:8a
!
        [B230M2-CH8-SERVER1-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER2-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:4b
!          [B230M2-CH8-SERVER2-fc1]
  member pwn 50:06:01:61:46:e0:5e:0a
!          [VNX7500-A1]
  member pwn 50:06:01:68:46:e0:5e:0a
!          [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER3-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:0c
!          [B230M2-CH8-SERVER3-fc1]
  member pwn 50:06:01:61:46:e0:5e:0a
!          [VNX7500-A1]
  member pwn 50:06:01:68:46:e0:5e:0a
!          [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER4-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:bc
!          [B230M2-CH8-SERVER4-fc1]
  member pwn 50:06:01:61:46:e0:5e:0a
!          [VNX7500-A1]
  member pwn 50:06:01:68:46:e0:5e:0a
!          [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER5-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:8d
!          [B230M2-CH8-SERVER5-fc1]
  member pwn 50:06:01:61:46:e0:5e:0a
!          [VNX7500-A1]
  member pwn 50:06:01:68:46:e0:5e:0a
!          [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER6-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:6e
!          [B230M2-CH8-SERVER6-fc1]
  member pwn 50:06:01:61:46:e0:5e:0a
!          [VNX7500-A1]
  member pwn 50:06:01:68:46:e0:5e:0a
!          [VNX7500-B0]

```

```

zone name B230M2-CH8-SERVER7-FC1 vsan 1
  member pwn 20:00:00:25:b5:c1:00:4f
!          [B230M2-CH8-SERVER7-fc1]

```

```

        member pwnn 50:06:01:61:46:e0:5e:0a
!
        [VNX7500-A1]
        member pwnn 50:06:01:68:46:e0:5e:0a
!
        [VNX7500-B0]

zone name B230M2-CH8-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:56
!
    [B230M2-CH8-SERVER8-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B230M2-CH1-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:76
!
    [B230M2-CH1-SERVER8-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B230M2-CH6-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:16
!
    [B230M2-CH6-SERVER8-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B230M2-CH7-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:36
!
    [B230M2-CH7-SERVER8-fc1]
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]

zone name B200M3-CH2-SERVER8-FC1 vsan 1
    member pwnn 50:06:01:61:46:e0:5e:0a
!
    [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!
    [VNX7500-B0]
    member pwnn 20:00:00:25:b5:c1:00:07

```

```

!                               [B200M3-CH2-SERVER8-fc1]

zone name B200M3-CH4-SERVER8-FC1 vsan 1
    member pwnn 20:00:00:25:b5:c1:00:a5
    member pwnn 50:06:01:61:46:e0:5e:0a
!                               [VNX7500-A1]
    member pwnn 50:06:01:68:46:e0:5e:0a
!                               [VNX7500-B0]

zoneset name DC-UCS-POD-B vsan 1
    member B230M2-CH1-SERVER1-FC1
    member B230M2-CH1-SERVER2-FC1
    member B230M2-CH1-SERVER3-FC1
    member B230M2-CH1-SERVER4-FC1
    member B230M2-CH1-SERVER5-FC1
    member B230M2-CH1-SERVER6-FC1
    member B230M2-CH1-SERVER7-FC1
    member B200M3-CH2-SERVER1-FC1
    member B200M3-CH2-SERVER2-FC1
    member B200M3-CH2-SERVER3-FC1
    member B200M3-CH2-SERVER4-FC1
    member B200M3-CH2-SERVER5-FC1
    member B200M3-CH2-SERVER6-FC1
    member B200M3-CH2-SERVER7-FC1
    member B200M3-CH3-SERVER1-FC1
    member B200M3-CH3-SERVER2-FC1
    member B200M3-CH3-SERVER3-FC1
    member B200M3-CH3-SERVER4-FC1
    member B200M3-CH3-SERVER5-FC1
    member B200M3-CH3-SERVER6-FC1
    member B200M3-CH3-SERVER7-FC1
    member B200M3-CH3-SERVER8-FC1
    member B200M3-CH4-SERVER1-FC1
    member B200M3-CH4-SERVER2-FC1
    member B200M3-CH4-SERVER3-FC1
    member B200M3-CH4-SERVER4-FC1
    member B200M3-CH4-SERVER5-FC1
    member B200M3-CH4-SERVER6-FC1
    member B200M3-CH4-SERVER7-FC1
    member B230M2-CH5-SERVER1-FC1
    member B230M2-CH5-SERVER2-FC1
    member B230M2-CH5-SERVER3-FC1
    member B230M2-CH5-SERVER4-FC1

```

```

member B230M2-CH5-SERVER5-FC1
member B230M2-CH5-SERVER6-FC1
member B230M2-CH5-SERVER7-FC1
member B230M2-CH6-SERVER1-FC1
member B230M2-CH6-SERVER2-FC1
member B230M2-CH6-SERVER3-FC1
member B230M2-CH6-SERVER4-FC1
member B230M2-CH6-SERVER5-FC1
member B230M2-CH6-SERVER6-FC1
member B230M2-CH6-SERVER7-FC1
member B230M2-CH7-SERVER1-FC1
member B230M2-CH7-SERVER2-FC1
member B230M2-CH7-SERVER3-FC1
member B230M2-CH7-SERVER4-FC1
member B230M2-CH7-SERVER5-FC1
member B230M2-CH7-SERVER6-FC1
member B230M2-CH7-SERVER7-FC1
member B230M2-CH8-SERVER1-FC1
member B230M2-CH8-SERVER2-FC1
member B230M2-CH8-SERVER3-FC1
member B230M2-CH8-SERVER4-FC1
member B230M2-CH8-SERVER5-FC1
member B230M2-CH8-SERVER6-FC1
member B230M2-CH8-SERVER7-FC1
member B230M2-CH8-SERVER8-FC1
member B230M2-CH1-SERVER8-FC1
member B230M2-CH6-SERVER8-FC1
member B230M2-CH7-SERVER8-FC1
member B200M3-CH2-SERVER8-FC1
member B200M3-CH4-SERVER8-FC1

```

```

zoneset activate name DC-UCS-POD-B vsan 1

```



## Appendix B—Sample Nexus 1000V VSM Configuration

These files are extremely large. If the reader is interested in a sample configuration file from one of the VSMs from this study, please contact your Cisco Account Manager or Sales Engineer to request it on your behalf.



# Appendix C—ESXtop Performance Charts for 5000 Virtual Machine Run



Single Cluster of 13 Blades' Data Shown Representing 1668 Virtual Desktops.



Run 100



























































































Run 106























### Cpu(\_Total)\% Core Util Time































































## Run 108








































































# Acknowledgments

- Vijayakumar D, Cisco Systems
- Usha Ramachandran, Cisco Systems
- TJ Singh, Cisco Systems
- Hector Jhong, Citrix Systems
- Seth Roth, Citrix Systems
- Randy Loveless, Citrix Systems
- Bahram Daleki, Citrix Systems
- Kevin Phillips, EMC