



Reference Architecture-Based Design for 4000 Seat Virtual Desktop Infrastructure

**Citrix XenDesktop 5.6 and Citrix XenApp 6.5 Built on
Cisco Unified Computing System, Nexus 5500, EMC
VNX7500, VMware ESXi 5.1**

A Cisco Validated Design

May 2013



Data Center of the Future





1 Overview	7
1.1 Solution Component Benefits	7
1.1.1 Benefits of Cisco Unified Computing System	7
1.1.2 Benefits of Nexus 5548UP	8
1.1.3 Benefits of EMC VNX Family of Storage Controllers	9
1.1.4 Benefits of VMware ESXi 5.1	9
1.1.5 Benefits of Citrix Desktop Virtualization	10
1.2 Audience	11
2 Summary of Main Findings	11
3 Architecture	13
3.1 Hardware Deployed	13
3.2 Software Revisions	15
3.3 Configuration Guidelines.....	15
3.3.1 VLANs	15
3.3.2 VMware Clusters	16
4 Infrastructure Components.....	16
4.1 Cisco Unified Computing System (UCS)	16
4.1.1 Cisco Unified Computing System Components.....	17
4.1.2 Fabric Interconnect	18
4.1.3 Cisco UCS Manager	18
4.1.4 Cisco UCS 2200 Series IO Module	18
4.1.5 Cisco UCS Chassis	19
4.1.6 Cisco UCS B200 M3 Blade Server	19
4.1.7 Cisco UCS VIC1240 Converged Network adapter.....	19
4.2 Citrix Desktop Virtualization	20
4.2.1 Citrix XenDesktop	20
4.2.1.1 Enhancements in Citrix XenDesktop 5.6 Feature Pack 1	21
4.2.2 Citrix FlexCast Technology.....	21
4.2.3 High-Definition User Experience Technology.....	22
4.2.4 Citrix XenDesktop Hosted VM Overview.....	22
4.2.5 Citrix XenApp Hosted Shared Desktop Overview.....	25
4.2.6 Citrix Provisioning Services.....	25
4.3 EMC VNX Series.....	25
4.3.1 EMC on VNX 7500 Used in Testing.....	26
4.4 VMware ESXi 5.1	26
4.4.1 VMware on ESXi 5.1 Hypervisor.....	26
4.5 Modular Desktop Virtualization Technical Overview	27
4.5.1 Modular Architecture.....	27
4.5.2 Understanding Desktop User Groups	29
4.5.3 Understanding Applications and Data	30
4.5.4 Project Planning and Solution Sizing Sample Questions	30



4.5.5 Cisco Services	31
4.5.6 The Solution: A Unified, Pre-Tested and Validated Infrastructure	31
4.6 Cisco Networking Infrastructure	31
4.6.1 Cisco Nexus 5548 Switch	31
4.6.2 Cisco Nexus 1000V	33
5 Architecture and Design of Citrix Desktop Virtualization on Cisco Unified Computing System and EMC VNX Storage ...	36
5.1 Design Fundamentals	36
5.2 Hosted Virtual Desktop Design Fundamentals	37
5.2.1 Hypervisor Selection	37
5.2.2 Provisioning Services	38
5.2.3 Citrix XenDesktop Broker	39
5.2.4 Citrix XenDesktop with Personal vDisk	40
5.2.5 Citrix XenApp	40
5.3 Designing a Mixed Citrix XenDesktop 5.6 and XenApp 6.5 Deployment	40
5.3.1 Hosted Virtual Desktops	40
5.3.2 Hosted Shared Desktops	41
5.4 Storage Architecture Design	42
6 Solution Validation	43
6.1 Configuration Topology for Scalable Citrix Mixed Workload Desktop Virtualization Infrastructure on Cisco Unified Computing System and EMC Storage	43
6.2 Cisco Unified Computing System Configuration	44
6.2.1 Base Cisco UCS System Configuration	45
6.2.2 QoS and CoS in Cisco Unified Computing System	78
6.2.3 System Class Configuration	79
6.2.4 Cisco UCS System Class Configuration	79
6.2.5 Steps to Enable QOS on the Cisco Unified Computing System	80
6.3 LAN Configuration	82
6.3.1 UCS Connectivity	82
6.3.2 EMC VNX7500 LAN Connectivity	82
6.3.3 Nexus 1000V Configuration in L3 Mode	83
6.3.4 Configuring Cisco UCS VM-FEX	102
6.4 SAN Configuration	113
6.4.1 Boot from SAN Benefits	113
6.4.2 Configuring Boot from SAN Overview	114
6.4.3 SAN Configuration on Cisco Nexus 5548UP	114
6.4.4 Configuring Boot from SAN on EMC VNX	117
6.4.5 SAN Configuration on Cisco UCS Manager	121
6.5 EMC VNX7500 Storage Configuration	125
6.5.1 Example EMC Volume Configuration for PVS Write Cache	127
6.5.2 EMC Storage Configuration for PVS vDisks	128
6.5.3 EMC Storage Configuration for VMware ESXi 5.0 Infrastructure and VDA Clusters	128
6.5.4 Example EMC Boot LUN Configuration	128
6.5.5 Example EMC FC LUN Configuration for SQL Clustering	129



6.5.6 EMC FAST Cache in Practice	129
6.5.7 EMC Additional Configuration Information	131
6.6 Cisco UCS Manager Configuration for VMware ESXi 5.1	133
6.6.1 Service Profile Templates	133
6.6.2 VLAN Configuration	133
6.7 Installing and Configuring ESXi 5.1	134
6.7.2 Install and Configure vCenter 5.1	134
6.7.4 ESXi 5.1 Cluster Configuration	136
6.8 Installing and Configuring Citrix Provisioning Server 6.1	141
6.8.1 Pre-requisites	141
6.8.2 Create a Highly Available CIFS Share for PVS vDisks	143
6.8.3 Install Citrix Licensing and Licenses	143
6.8.3.1 Pre-requisites	143
6.8.4 Install Provisioning Services 6.1	149
6.8.5 Install Required PVS Hotfixes	158
6.8.6 Adding PVS Servers to the Farm	158
6.9 Installing and Configuring Citrix XenDesktop 5.6 FP1	160
6.9.1 Pre-requisites	161
6.9.2 Install XenDesktop, XenDesktop Studio, and Optional Components	162
6.9.3 Create SQL database for XenDesktop	163
6.9.4 Configure the XenDesktop Site Hosts and Storage	165
6.9.5 Configure XenDesktop HDX Policies	168
6.9.6 Configure the XenDesktop Desktop Group and Options	176
6.10 Installing and Configuring Citrix XenApp 6.5	179
6.10.1 Pre-requisites	179
6.10.2 Install Citrix XenApp 6.5	181
6.10.3 Configure XenApp License Information	195
6.10.4 Configure XenApp Farm	196
6.10.5 Install Required XenApp Hot Fixes	208
6.10.6 Configure the Hosted Shared Desktop (Citrix)	209
6.10.7 Configure XenApp Policies	218
6.11 Installing and Configuring Citrix NetScaler Virtual Appliance on VMware ESX for Use with Citrix StoreFront	228
6.12 Installing and Configuring Citrix StoreFront for XenDesktop and XenApp	240
6.12.1 Install	240
6.12.2 Create StoreFront Database	243
6.12.3 Create Multiple Server Group	243
6.12.4 Join an Existing Server Group	253
7 Desktop Delivery Infrastructure and Golden Image Creation	257
7.1 Overview of Solution Components	257
7.2 Citrix User Profile Management Servers	261
7.2.1 Setting up a Highly Available Profile Share	261
7.2.2 Setting Up a Two Node Citrix User Profile Server Cluster	261
7.2.3 Install User Profile Manager	262



7.2.4 Create a GPO for UPM.....	263
7.3 Microsoft Windows 7 Golden Image Creation.....	265
7.3.1 Create Base Windows 7 SP1 32bit Virtual Machine.....	265
7.3.2 Add Provisioning Services Target Device Software.....	265
7.3.3 Add a XenDesktop 5.6 Virtual Desktop Agent.....	270
7.3.4 Add Login VSI Target Software.....	275
7.3.5 Perform Additional PVS and XenDesktop Optimizations.....	277
7.3.6 Convert Golden Image Virtual Machine to PVS vDisk.....	278
7.3.7 Add Write-Cache Drives to Virtual Machine Templates in vCenter.....	281
7.4 Microsoft Windows Server 2008 R2 XenApp Golden Image Creation.....	282
7.4.1 Create Base Windows Server 2008 R2 64bit Virtual Machine.....	283
7.4.2 Add Provisioning Services Target Device Software.....	283
7.4.3 Prepare XenApp Server for Imaging and Provisioning.....	283
7.4.4 Add Login VSI 3.7 Target Software.....	289
7.4.5 Perform Additional PVS Optimizations.....	291
7.4.6 Convert Golden Image Virtual Machine to PVS vDisk.....	292
7.4.7 Add Write-Cache Drives to Virtual Machine Templates in vCenter.....	296
7.5 Citrix Provisioning Server (PVS) 6.1 Services.....	297
7.5.1 Storage Configuration for PVS.....	297
7.5.2 PVS for Use with Standard Mode Desktops.....	298
7.5.3 Process to Create Hosted VM Desktops with Tier 0 Storage Using XenDesktop Setup Wizard in PVS.....	299
7.5.4 Process to Create Virtual Desktops with Personal vDisk in PVS.....	306
7.5.5 Process to Create XenApp 6.5 Virtual Machines Using the Streamed VM Setup Wizard in PVS.....	311
8 Test Setup and Configurations.....	317
8.1 Cisco UCS Test Configuration for Single Blade Scalability.....	318
8.2 Cisco UCS Configuration for Cluster Tests.....	321
8.3 Cisco UCS Configuration for Four Chassis – Twenty Five 4000 Mixed Workload Blade Test.....	324
8.4 Testing Methodology and Success Criteria.....	325
8.4.1 Load Generation.....	325
8.4.2 User Workload Simulation – LoginVSI From Login VSI Inc.....	325
8.4.3 Testing Procedure.....	327
8.4.4 Success Criteria.....	328
8.4.4.1 Login VSImax.....	328
9 Citrix XenDesktop 5.6 Hosted VM and Citrix XenApp 6.5 Shared Hosted Desktop Mixed Workload on Cisco UCS B200 M3 Blades, EMC VNX7500 and VMware ESXi 5.1 Test Results.....	332
9.1 Cisco UCS Test Configuration for Single-Server Scalability Test Results.....	334
9.1.1 Pooled XenDesktop 5.6 Hosted VM with Tier 0 (SSD) Storage.....	335
9.1.2 Citrix XenDesktop 5.6 Hosted VM with Personal vDisk.....	340
9.1.3 Citrix XenApp 6.5 Shared Hosted Desktop.....	346
9.1.4 Citrix XenApp 6.5 Shared Hosted Desktop (VM-FEX).....	351
9.2 Cisco UCS Test Configuration for Single Cluster Scalability Test Results.....	356
9.2.1 Pooled XenDesktop 5.6 Hosted VM with Tier 0 (SSD) Storage.....	357
9.2.3 Citrix XenApp 6.5 Shared Hosted Desktop.....	370



9.3 Cisco UCS Test Configuration for 4100 Desktop Mixed Workload Scalability Test Results	376
10 Scalability Considerations and Guidelines	391
10.1 Cisco UCS System Configuration	391
10.2 Citrix XenDesktop 5.6 Hosted VDI.....	392
10.3 EMC VNX Storage Guidelines for Citrix XenDesktop Provisioned Virtual Machines.....	392
10.4 VMware ESXi 5 Guidelines for Virtual Desktop Infrastructure	392
11 References.....	393
11.1 Cisco Reference Documents	393
11.2 Citrix Reference Documents	393
11.2.1 XenDesktop 5.6	393
11.2.3 XenApp 6.5	394
11.2.4 Provisioning Services 6.1	394
11.2.5 Citrix User Profile Manager	394
11.3 EMC Reference Documents	394
11.4 VMware Reference Documents	395
Appendix A Nexus 5548UP Configurations	395
N5548UP-A Configuration.....	395
N5548UP-B Configuration.....	420
Appendix B Sample Nexus 1000V VSM Configuration.....	461
Citrix-VSM-01	461
Citrix-VSM-02	461
Appendix C Sample VM-FEX VSM Configuration	461
Citrix-VMFEX-01	461
Appendix D ESXtop Performance Charts for 4100 Session Run.....	461



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 a mixed workload environment including Citrix XenDesktop 5.6 Hosted Virtual Desktops with PVS write-cache on SSDs on the blades, XenDesktop 5.6 Hosted Virtual Desktop 5.6 with Personal vDisk, and XenApp 6.5 Hosted Shared Desktop environment, utilizing Citrix Provisioning Server 6.1, on Cisco UCS B-Series B200 M3 Blade Servers running VMware ESXi 5.1 hypervisor software connected to an EMC VNX 7500 Storage Array. We utilize second and third generation Unified Computing System hardware and software. We provide best practice recommendations and sizing guidelines for large scale customer deployments of the mixed workload 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 include the following:

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 20 Chassis and up to 160 blades in a single Cisco UCS management domain
- Tight integration with VMware vCenter for Cisco Unified Computing System



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, which includes the following:

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 EMC VNX Family of Storage Controllers

The EMC VNX Family

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 \$8K, 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 as follows:

- 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

1.1.4 Benefits of VMware ESXi 5.1



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

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

1. [Built on a robust, reliable foundation](#)
2. [Delivers a complete virtualization platform from desktop through the datacenter out to the public cloud](#)
3. [Provides the most comprehensive virtualization and cloud management](#)
4. [Integrates with your overall IT infrastructure](#)
5. [Proven with over 350,000 customers](#)

And best of all, VMware delivers while providing

6. [Low total-cost-of-ownership \(TCO\)](#)

1.1.5 Benefits of Citrix Desktop Virtualization

Citrix offers three key technologies that enable high definition end user desktops:

Citrix XenDesktop

XenDesktop is a comprehensive desktop virtualization solution that includes all the capabilities required to deliver desktops, applications 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.

XenDesktop helps businesses with the following:

- Enable virtual work styles 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 offshoring, 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.

Citrix XenApp

Citrix XenApp empowers users with on-demand self-service to enterprise hosted shared desktops and applications. Used by more than 100 million people worldwide, XenApp is an application delivery solution that enables any Windows® application to be virtualized, centralized and managed in the datacenter and instantly delivered as a service to users anywhere on any device. [Virtual application](#) delivery with XenApp enables organizations to improve [application management](#) by:



- Centralizing hosted shared desktop and applications in the datacenter to reduce complexity and lower the cost of [desktop management](#) by up to 50 percent
- Controlling and encrypting access to data and applications to improve security
- Delivering applications instantly to users anywhere on any device
- Simplifying and automating the process of delivering or updating applications, enabling IT to focus on strategic initiatives

XenApp reduces the cost of [desktop management](#) up to 50 percent by simplifying the management and delivery of all Windows applications. Centralizing hosted shared desktops and applications in the datacenter reduces costs and increases efficiency by enabling IT to manage a single instance of each application in an application hub, rather than using manual processes to install applications on every PC.

Citrix Provisioning Server

Citrix Provisioning Server provides the capability to create, maintain and start hundreds of virtual desktop and virtual servers from a single virtual disk image respectively. The benefits are greatly reduced disk storage requirements compared to other provisioning techniques and the ability to maintain all of the desktops or servers by changing their respective single virtual disk image.

Citrix Provisioning Server 6.1 was used to create and maintain XenDesktop 5.6 virtual desktops and XenApp 6.5 virtual servers for hosted shared desktops.

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 4000 seat mixed Desktop Virtualization Infrastructure delivering outstanding end-user experience with the following concurrently running workloads:

- 2000 Citrix XenApp 6.5 Hosted Shared Desktop sessions
- 1000 Citrix XenDesktop 5.6 Pooled Hosted Virtual Desktops with PVS Write Cache on Tier O Storage
- 1000 Citrix XenDesktop 5.6 Hosted Virtual Desktops with Personal vDisk

The combined power of the Cisco Unified Computing System, Nexus switching and EMC storage hardware, VMware ESXi 5.1, Citrix Provisioning Server 6.1, Citrix XenApp 6.5 and Citrix XenDesktop 5.6 software produces a high-density per blade and per chassis mixed workload Virtual Desktop delivery system with the following:

- Cisco UCS B200 M3 half-width blade with dual 8-core processors, 256GB of 1600 MHz memory and two 300GB SSDs for PVS write cache supports 40.9% more pooled hosted virtual desktop workloads than the previously studied full width blade using a new Login VSI medium workload with flash.
- Cisco UCS B200 M3 half-width blade with dual 8-core processors and 256GB of 1600 MHz memory supports 31.8% more hosted virtual desktop with Personal vDisk workloads than the previously studied full width blade using a new Login VSI medium workload with flash.
- Cisco UCS B200 M3 half-width blade with dual 8-core processors and 256GB of 1600 MHz memory supports 11.1% more hosted shared desktop sessions than the previously studied full width blade using a new Login VSI medium workload with flash.



- The study design based on 25 Cisco Unified Computing System B200 M3 desktop virtualization blades, each with dual 8-core processors, 256GB of memory and a Cisco UCS VIC1240 converged network adapter supports 4000 virtual desktop workloads running the new medium workload with flash, more than 3 times the density of previously studied chassis with full width blades.
- The 4000 seat system booted to a login prompt in less than 30 minutes without exhausting server CPU, memory, network or storage subsystems.
- We were able to ramp up (log in and start workloads) to steady state (with all 4000 users logged in and working) 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. 1. All of the Windows 7 SP1 virtual desktops and supporting infrastructure components, including Active Directory, Profile Servers, Provisioning Servers, SQL Servers, XenDesktop delivery controllers and XenApp servers were hosted as virtual servers.
- We maintain our industry leadership with our new Cisco UCS Manager 2.1(1a) software that makes scaling simple, consistency guaranteed and maintenance simple. Combined with UCS Central, our Cisco UCS management scope extends to over 100,000 virtual desktops.
- Our 10G unified fabric story gets additional validation on second generation Cisco UCS 6200 Series Fabric Interconnects and second generation Nexus 5500 Series access switches as we run more challenging workload testing, maintaining unsurpassed user response times.
- EMC's VNX 7500 system provides storage consolidation and outstanding efficiency. Both block and NFS storage resources were provided by a single system.
- EMC's Fast Cache technology delivers predictable performance and continuous availability for end user computing environment ensuring a rich end user experience while enforcing compliance, data security, high-availability, and increasing IT productivity.
- EMC delivers simplified management to an End User Computing infrastructure through Unisphere for ease of configuration and management, plug-in technologies for simplified desktop provisioning, and integrations that deliver rich metrics for monitoring your VNX storage platform.
- Citrix HDX technology, extended in XenDesktop 5.6 Feature Pack 1 software, provides excellent performance with host-rendered flash video and other demanding applications.
- Citrix XenApp 6.5 extends the flexibility of the solution design by adding hosted shared server desktops and published applications to the solution array.
- Citrix FlexCast technology provides the right desktop resources to the right device and insures a consistent user experience based on the device being used.



3 Architecture

3.1 Hardware Deployed

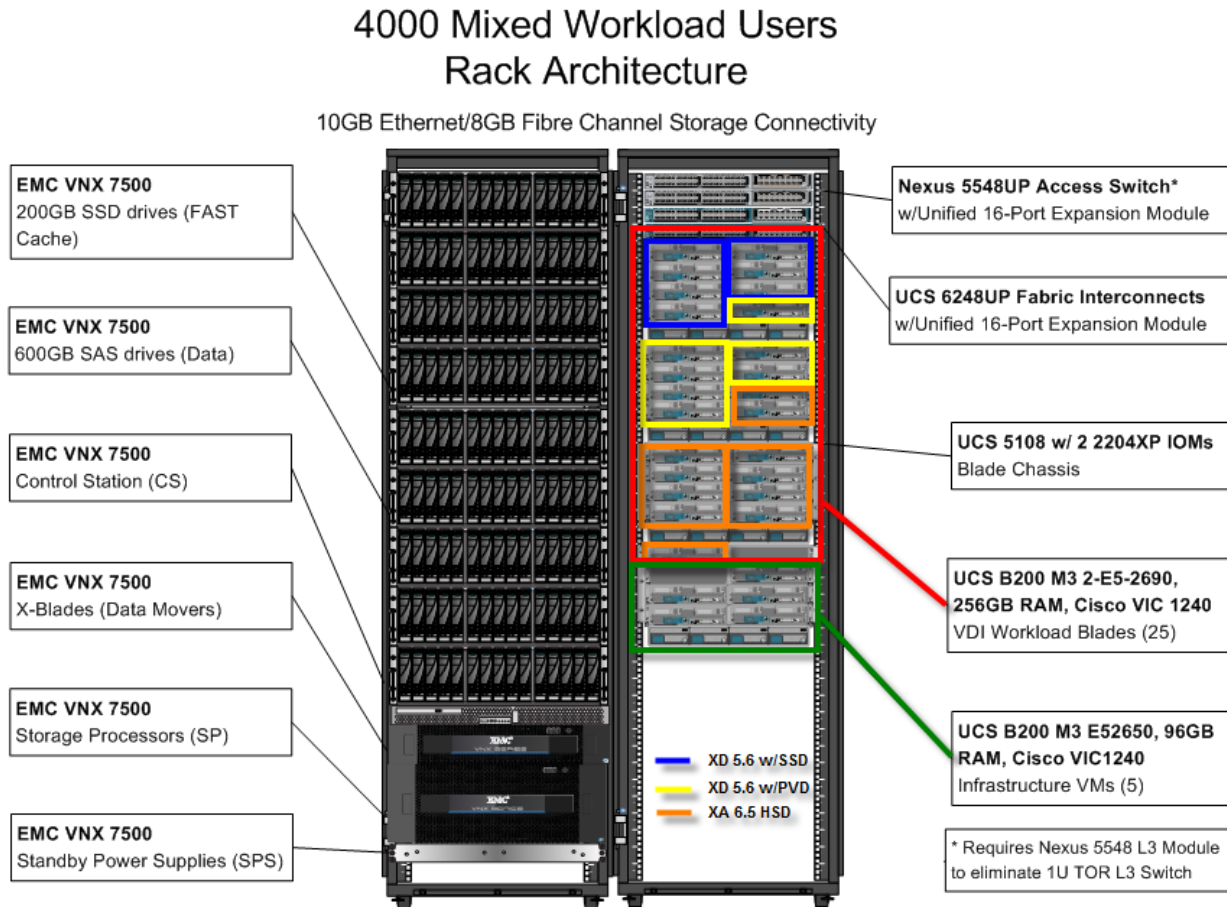
The architecture deployed is highly modular. While each customer's environment might vary in its exact configuration, when the reference architecture contained in this document is built, it can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources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

The 4000 User XenDesktop 5.6 and XenApp 6.5 solution includes Cisco networking, Cisco UCS and EMC storage which the computing and storage that fit in two data center racks, including the access layer network switches in the Cisco UCS rack.

This document details the deployment of a 4000 seat mixed workload desktop virtualization solution featuring:

- Citrix XenDesktop 5.6 Pooled Hosted Virtual Desktops with PVS Write Cache on Tier O Storage
- Citrix XenDesktop 5.6 Hosted Virtual Desktops with Personal vDisks
- Citrix XenApp 6.5 Hosted Shared Desktops
- Citrix Provisioning Server 6.1
- Citrix User Profile Manager
- Citrix StoreFront
- Citrix NetScaler VPX
- Cisco Nexus 1000V Distributed Virtual Switch
- Cisco Virtual Machine Fabric Extender (VM-FEX)
- VMware ESXi 5.1 Hypervisor

Figure 1. Citrix Mixed Workload Users 4000 User Hardware Components



The reference configuration includes:

- Two Cisco Nexus 5548UP switches with 16-universal port Expansion Modules
- Two Cisco UCS 6248 Series Fabric Interconnects with UCS 6200 16-universal port Expansion Modules
- Four 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
- Twenty-five Cisco UCS B200 M3 Blade servers with Intel E5-2690 processors, 256 GB RAM, and VIC1240 mezzanine cards for the mixed desktop virtualization workloads
- One EMC VNX7500 dual controller storage system for HA, 4 Datamovers, 600GB SAS Drives and 200GB SSD Fast Cache Drives

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 Interconnects	2.1 (1a)	Embedded Management
	Cisco UCS B200 M3	2.1 (1a)	Hardware BIOS
Network	Nexus 5548UP Switch	5.2(1)N1(1)	Operating System Version
Storage	EMC VNX7500	ONTAP 8.1.0 RC2	Operating System Version
Software	Cisco UCS Blade Hosts	B200: VMware ESXi 5.1 B230: VMware ESXi 5.1	Operating System Version
	Cisco Nexus 1000V	4.2(1)SV1(5.2)	Virtual Switch appliance version

3.3 Configuration Guidelines

The 4000 User Mixed Workload Desktop Virtualization solution described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided and refer to which redundant component is being configured with each step. Redundant components are designated as A or B. For example, SPA and SPB 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 Mixed Workload Desktop Virtualization customer environment as stand-alone solution.

3.3.1 VLANs

For the 4000 User Mixed Workload Desktop Virtualization 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_PKT	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) (One XenDesktop 5.6 with PVS write cache on local SSDs and Nexus 1000V, One XenDesktop 5.6 with Personal vDisk and Nexus 1000V, and One XenApp 6.5 Hosted Shared Desktop with VM-FEX.)
- 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 via the Nexus 5000 switches in each domain).

4 Infrastructure Components

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

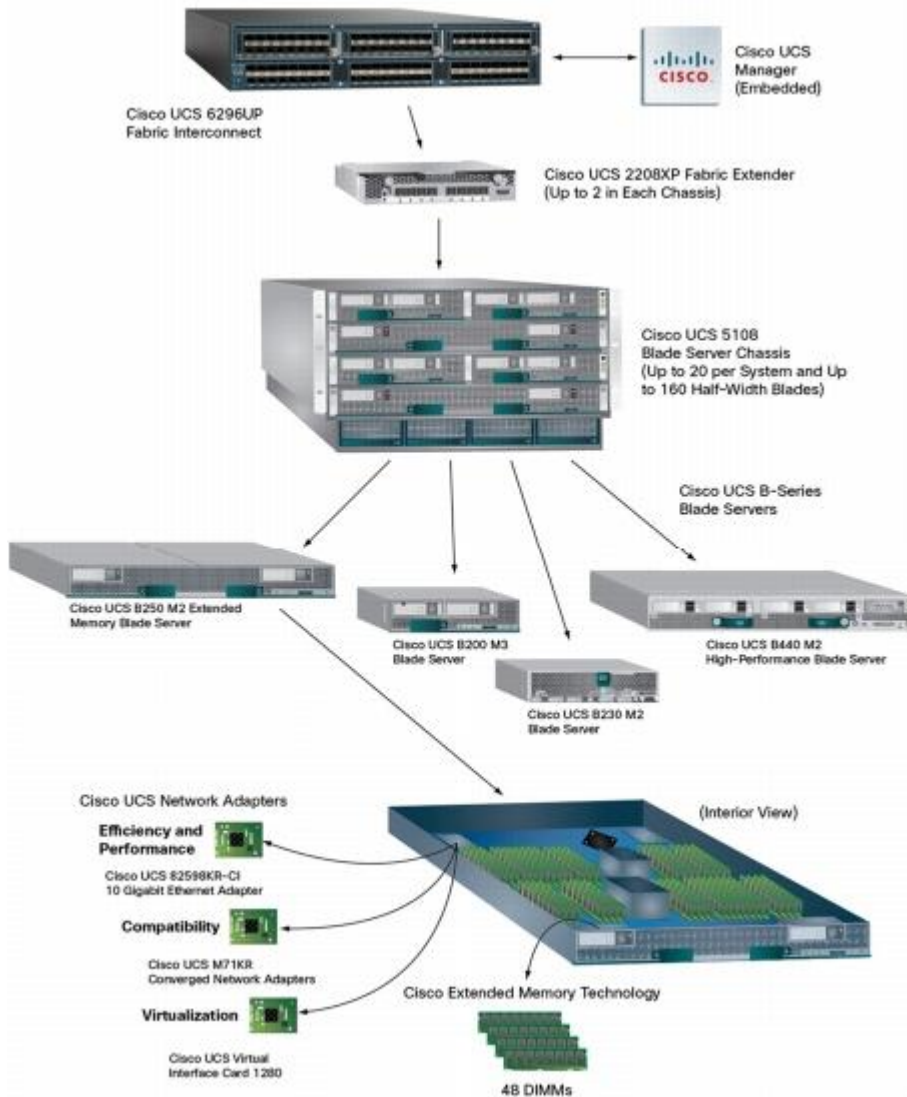
4.1 Cisco Unified Computing System (UCS)

The Cisco Unified Computing System™ (Cisco UCS™) is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

4.1.1 Cisco Unified Computing System Components

Cisco UCS components are shown in Figure 2.

Figure 2. Cisco Unified Computing System Components



The Cisco UCS 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

The Cisco UCS 6200 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system (Figure 2). The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

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

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, switching capacity of 2 terabits (Tb), and 320-Gbps bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco® low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated. 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 Manager

The Cisco UCS 6200 Series hosts and runs Cisco UCS Manager in a highly available configuration, enabling the fabric interconnects to fully manage all Cisco UCS elements. Connectivity to the Cisco UCS 5100 Series blade chassis is maintained through the Cisco UCS 2100 or 2200 Series Fabric Extenders in each blade chassis. The Cisco UCS 6200 Series interconnects support out-of-band management through a dedicated 10/100/1000-Mbps ethernet management port as well as in-band management. Cisco UCS Manager typically is deployed in a clustered active-passive configuration on redundant fabric interconnects connected through dual 10/100/1000 ethernet clustering ports.

4.1.4 Cisco UCS 2200 Series IO Module

The Cisco UCS 2200 Series IO Module 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.

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



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

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

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

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

4.1.6 Cisco UCS B200 M3 Blade Server

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

4.1.7 Cisco UCS VIC1240 Converged Network adapter

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

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

Figure 3. The Cisco UCS VIC 1240 Converged Network Adapter





Figure 4. The Cisco UCS VIC 1240 Logical Diagram



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

4.2 Citrix Desktop Virtualization

4.2.1 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. The Citrix XenDesktop open architecture enables customers to easily adopt desktop virtualization using any hypervisor, storage, or management infrastructure.



4.2.1.1 Enhancements in Citrix XenDesktop 5.6 Feature Pack 1

XenDesktop 5.6 Feature Pack 1, builds upon 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 Citrix 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.

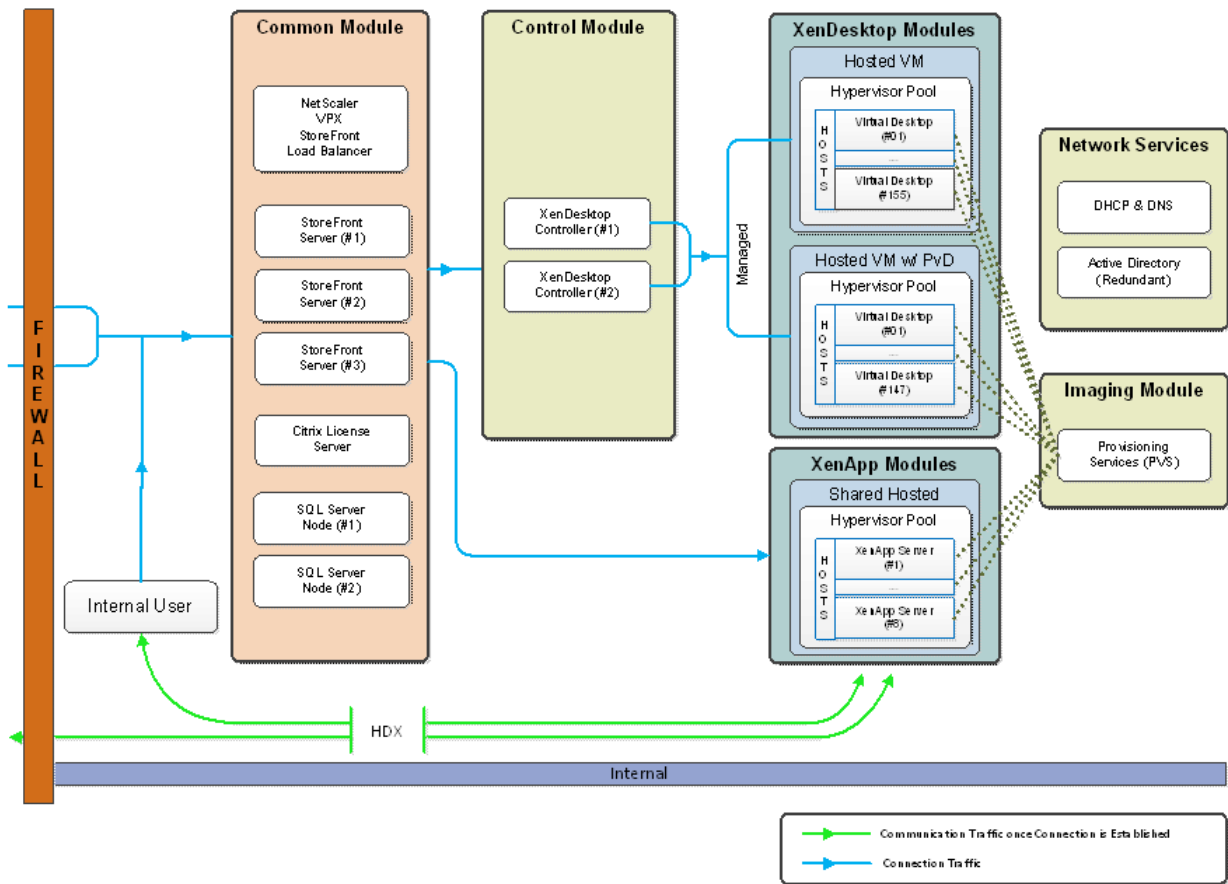
4.2.3 High-Definition User Experience Technology

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 upon 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 Figure 5.

Figure 5. Citrix XenDesktop and XenApp 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® ICA® 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 brokers user authentication, enumerates the available desktops and, upon 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 XenApp 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 not 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.



Available Software Suites

- 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.

Available Software Packs

- 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 5.1

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

4.4.1 VMware on ESXi 5.1 Hypervisor

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

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

- **Improved Reliability and Security** — with fewer lines of code and independence from general purpose OS, ESXi drastically reduces the risk of bugs or security vulnerabilities and makes it easier to secure your hypervisor layer.
- **Streamlined Deployment and Configuration** — ESXi has far fewer configuration items than ESX, greatly simplifying deployment and configuration and making it easier to maintain consistency.



- **Higher Management Efficiency** — The API-based, partner integration model of ESXi eliminates the need to install and manage third party management agents. You can automate routine tasks by leveraging remote command line scripting environments such as vCLI or PowerCLI.
- **Simplified Hypervisor Patching and Updating** — Due to its smaller size and fewer components, ESXi requires far fewer patches than ESX, shortening service windows and reducing security vulnerabilities.

4.5 Modular Desktop Virtualization Technical Overview

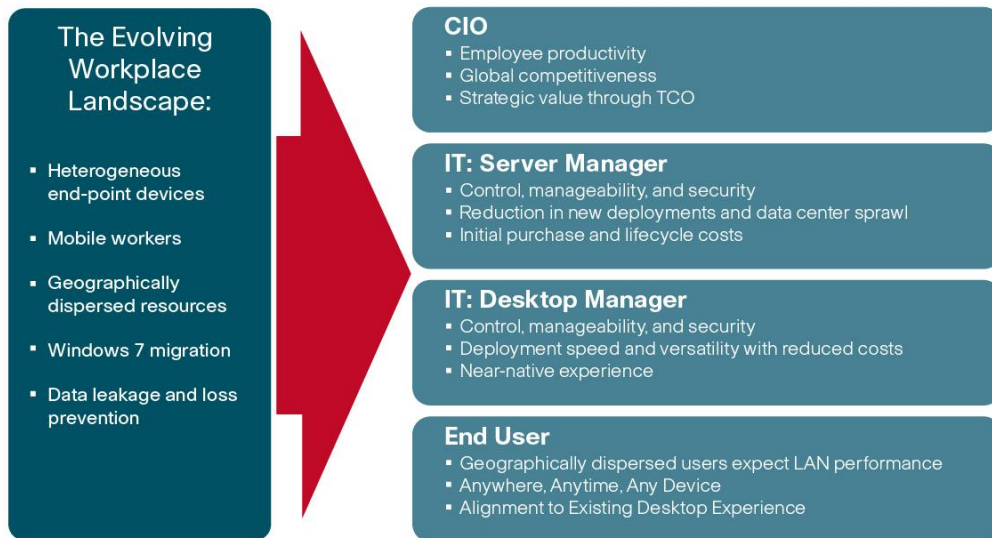
4.5.1 Modular Architecture

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

Figure 6. **The Evolving Workplace Landscape**

Trends and Expectations



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.

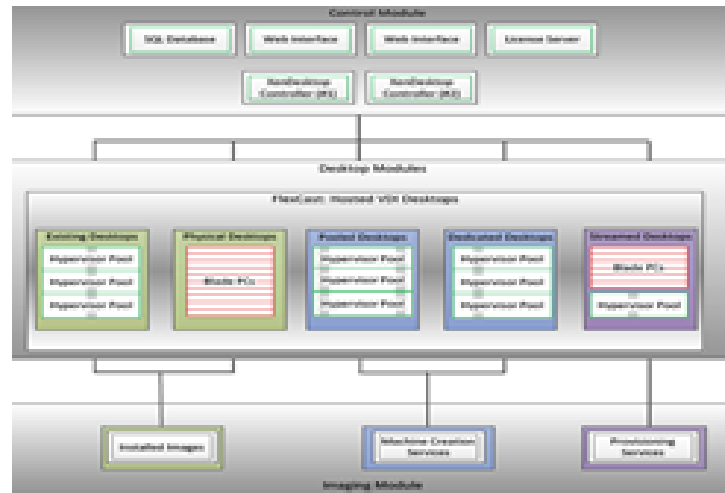


Figure 7. Citrix XenDesktop on Cisco UCS


Citrix XenDesktop on Cisco UCS: Modular Design Enables Best Scalability



Solution of the Year



Modular Software



Modular Infrastructure



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 NetApp – have developed integrated, validated architectures, including available pre-defined, validated infrastructure packages, known as FlexPod.



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 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 UCS 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, via the Cisco UCS architecture and unified fabric.

The simplified deployment of Cisco UCS 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 (e.g., 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 (e.g., 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 Desktop Virtualization 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

4.5.6 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 desktop virtualization 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.

4.6.1.1 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.

4.6.1.2 Features and Benefits

Specific features and benefits provided by the Cisco Nexus 5500 Series follow.

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.

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:](#)

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.

4.6.2 Cisco Nexus 1000V

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

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

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

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

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

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

4.6.2.1 Nexus 1000V Product Architecture

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

Virtual Ethernet Module (VEM)

The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware virtual switch (vSwitch). This level of integration helps ensure that the Cisco Nexus 1000V Series is fully aware of all server virtualization events, such as VMware vMotion and Distributed Resource Scheduler (DRS). The VEM takes



configuration information from the VSM and provides advanced networking functions: quality of service (QoS), security features, and monitoring features.

Virtual Supervisor Module (VSM)

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

4.6.2.2 Nexus 1000V Features and Benefits

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

Policy-Based Virtual Machine Connectivity

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

Mobility of Virtual Machine Security and Network Properties

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

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

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

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

Non-disruptive Operational Model

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



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



Enhanced Deployment Scenarios

- Optimized server bandwidth for I/O-intensive applications: Today, network interfaces are often dedicated to a particular type of traffic, such as VMware Console or vMotion. With the Cisco Nexus 1000V Series, all network interface cards (NICs) can be treated as a single logical channel with QoS attached to each type of traffic. Consequently, the bandwidth to the server can be more efficiently utilized, with network-intensive applications virtualized.
- Easier security audits with consistent security policy: Security audits on virtual machines are usually more difficult to perform because virtual machines are secured differently than physical servers. As the Cisco Nexus 1000V Series provides persistent security policy to mobile virtual machines, security audits are similar to those for physical servers.
- Virtual machine as basic building block of data center: With the Cisco Nexus 1000V Series, virtual machines are treated the same way as physical servers in security policy, monitoring and troubleshooting, and the operational model between network and server administrators, enabling virtual machines to be true basic building blocks of the data center. These operational efficiencies lead to greater scaling of server virtualization deployments with lower operating expenses.

VMware Product Compatibility

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

5 Architecture and Design of Citrix Desktop Virtualization on Cisco Unified Computing System and EMC VNX Storage

5.1 Design Fundamentals

There are many reasons to consider desktop virtualization solutions 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) programs. The first step in designing a desktop virtualization solution is to understand the user community and the type of tasks that are required to successfully execute their role. The following user classifications are provided:

- **Knowledge Workers** today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- **External Contractors** are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- **Task Workers** perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.



- **Mobile Workers** need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- **Shared Workstation** users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

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

Citrix XenDesktop 5.6 can be used to deliver a variety of virtual desktop configurations. The following sections describe the deployment considerations.

5.2.1 Hypervisor Selection

Citrix XenDesktop and Citrix XenApp are hypervisor agnostic, so any of the following three hypervisors can be used for the desktop virtualization solutions:



- 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 5.1 and vCenter 5.1.

5.2.2 Provisioning Services

XenDesktop Hosted Virtual Desktops and XenApp virtual servers for Hosted Shared Desktops can be deployed with and without Citrix Provisioning Services, but Citrix Provisioning Services enables you to stream a single desktop image to create multiple virtual desktops or XenApp servers on one or more hosts in the data center quickly and efficiently. This facility greatly reduces the amount of storage required compared to other methods of creating virtual desktops and virtual servers.

For XenDesktop Hosted Virtual Desktops, Citrix Provisioning Services created desktops can be deployed as 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.
- 2.
- **PVS Standard Mode 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 are 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 or virtual server 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 local HD is stored in a file on a secondary local hard drive of the 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:** Server Cache is stored in a file on the 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 upon shutdown.

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

Hosted Shared Desktops provided by Citrix XenApp can be deployed on bare metal hardware or as virtual machines. PVS is extremely useful for creating multiple XenApp virtual machines to support Hosted Shared Desktops. A single vDisk image supports all of the XenApp virtual machines in the deployment. Citrix XenApp servers were deployed as virtual machines in this exercise.

In this study, we used Provisioning Server 6.1 for the creation and management of Pooled Desktops with PVS Write Cache on device HD of each virtual machine. 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.

Similarly, we used Provisioning Server 6.1 for deployment and management of XenApp virtual servers with PVS write-cache on device HD of each virtual machine.

5.2.3 Citrix XenDesktop Broker

XenDesktop Controllers (XDCs) or brokers are one of the cornerstones of any XenDesktop deployment. They provide authentication for users coming into the XD farm; they enumerate resources available for all of those users, and they also direct launch requests from the users to the appropriate VDIs. In addition to providing user services, they also maintain desktop startups and shutdowns. The broker systems communicate constantly with the SQL database, ensuring that a single broker can fail without issue.

General best practice for XenDesktop Controllers is to have at least two per XenDesktop Site for redundancy. As additional users or resources come online, it may be necessary to either scale to larger XDCs or scale to a larger number of XDCs (scaling up or out, as the case may be). Additional XDCs can be deployed based upon geographic layout or other factors that may come up – including division of users or division of organizations. Additional information on best practices for XenDesktop Controllers can be found on the Citrix website at the following link:

<http://support.citrix.com/article/CTX132799>

In this particular environment, we deployed two XenDesktop Controllers, as that was the best number for the number of users that we deployed.



5.2.4 Citrix XenDesktop with Personal vDisk

In every environment, there exists a subset of users that wants complete personalization control over their desktop. In prior releases of XenDesktop, this was accomplished through the use of Dedicated Desktops. While providing some of the benefits of the VDI environment, this brought about additional complexities, such as the requirement to manage all of these additional images, at one image per user instead of a single image per a group of users. There is also a significant storage consumption add-on, as each of these images now consumes enterprise-class storage instead of just occupying someone's desktop.

With XenDesktop 5.6, Citrix introduced the concept of the Personal vDisk, or PvD. PvD allows consumers of the normal pooled random image to extensively personalize their desktop, eliminating the requirement for Dedicated Desktop images, while still allowing the IT group to only manage that smaller number of pooled images.

One important note when considering the use of PvD in an environment is that deployment of PvD creates a connection between a user and a particular virtual desktop. If that virtual desktop is unavailable for any reason, the user cannot log on.

Additional information about planning a Personal vDisk environment can be found in the PvD Planning Guide, available on the Citrix website at the following link:

<http://support.citrix.com/article/CTX133227>

The environment described in this document deployed half of the Hosted Virtual desktops with Personal vDisk technology.

5.2.5 Citrix XenApp

Citrix XenApp 6.5 provides the capability of efficiently sharing Server 2008 R2 server desktops. Additionally, applications can be delivered on the shared server desktops by installing them in XenApp 6.5 virtual disk (vDisk) or by hosting the applications on separate XenApp 6.5 virtual servers.

User sessions provided by XenApp 6.5 virtual servers are referred to as Hosted Virtual Desktops (HVDs.)

In this environment, HVDs were deployed for task worker workloads, with the required applications installed on the XenApp 6.5 server's vDisk.

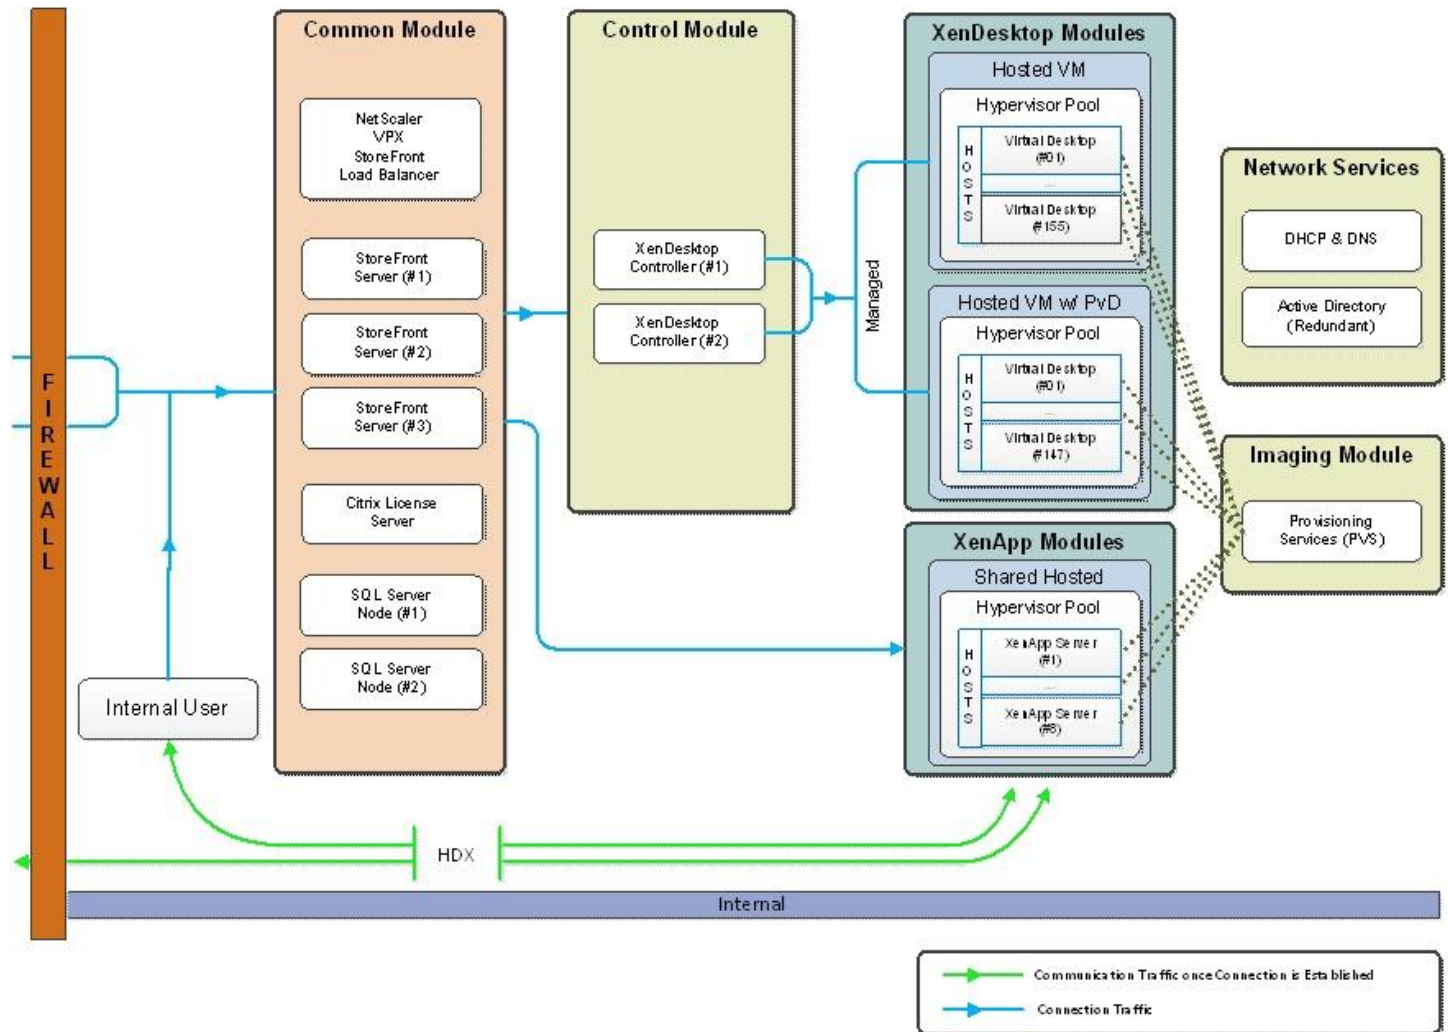
5.3 Designing a Mixed Citrix XenDesktop 5.6 and XenApp 6.5 Deployment

The design of this environment presented a combined use case – XenApp Hosted Shared Desktops (HSDs), XenDesktop Pooled Random Desktops and XenDesktop Pooled Random with PvD – in order to more closely simulate a customer environment with several different types of users.

5.3.1 Hosted Virtual Desktops

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. Citrix Pooled Desktop Infrastructure

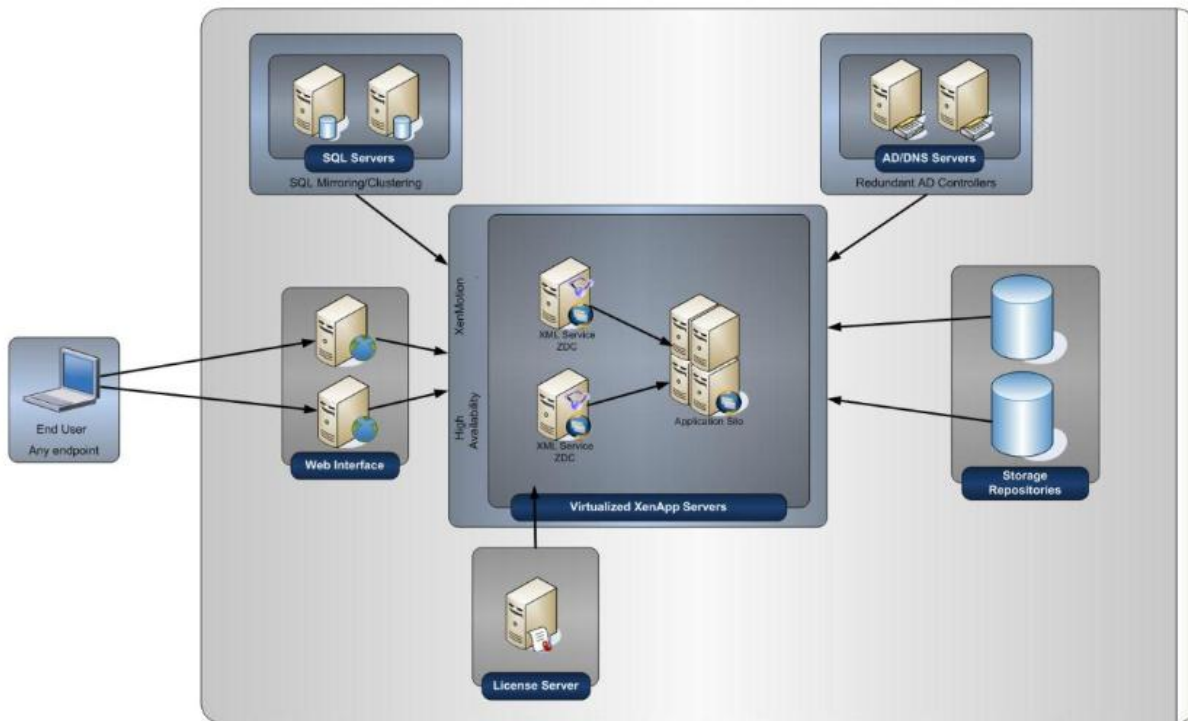


We elected to use Provisioning Services 6.1 in place of Machine Creation Services for this study so that our design would scale to many thousands of desktops.

5.3.2 Hosted Shared Desktops

To implement our shared server desktop delivery model for this study, known as Hosted Shared Desktops, we followed the Citrix Reference Architecture for highly available desktop virtualization.

Figure 9. Citrix Hosted Shared Desktop Infrastructure



Note: The infrastructure servers show in Figure X and Figure Y above were used for both XenDesktop and XenApp. There is no requirement to maintain separate infrastructures for the delivery methods.

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

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

Learn more about XenDesktop 5.6 Planning and Design at the following location:

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

To read about Citrix’s XenApp Reference Architecture, go to the following link:

<http://support.citrix.com/article/ctx131762>

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, the storage pools used to store PVS write cache are made up of 16 or 48 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 Mixed Workload Desktop Virtualization Infrastructure on Cisco Unified Computing System and EMC Storage

Figure 10. Architecture Block Diagram

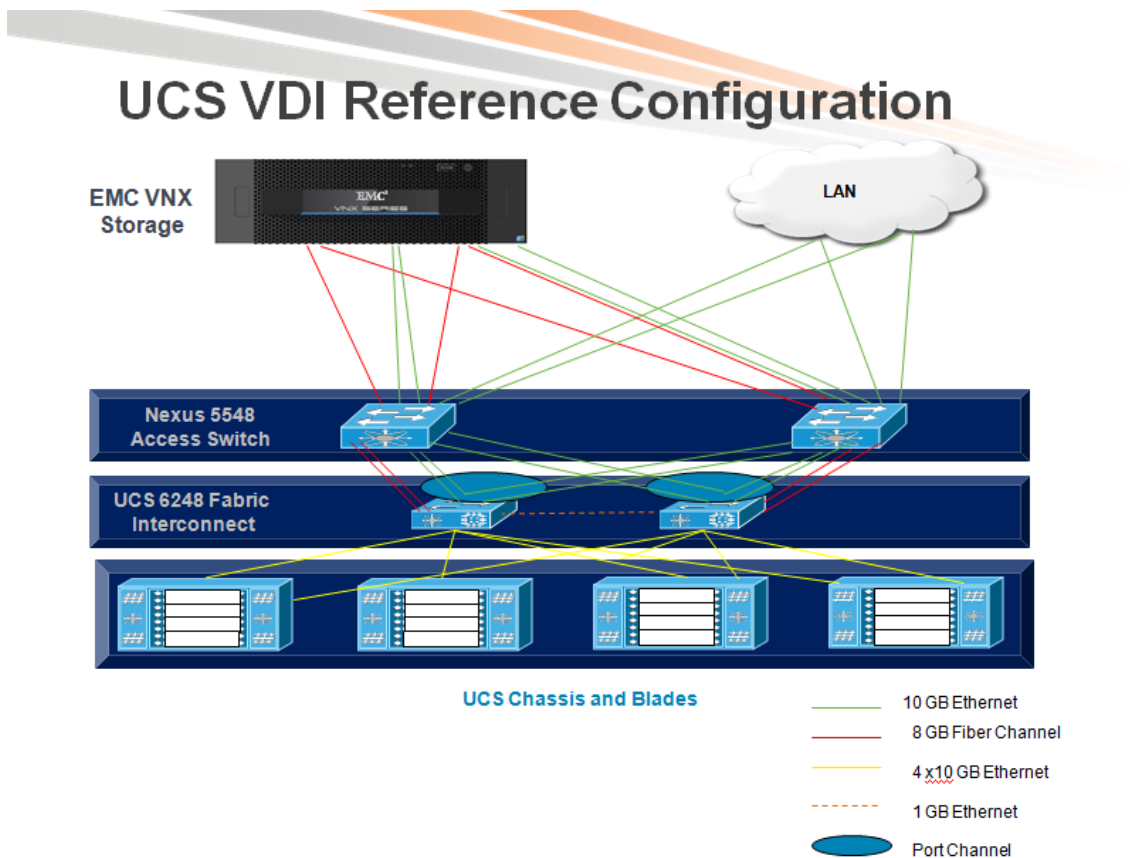
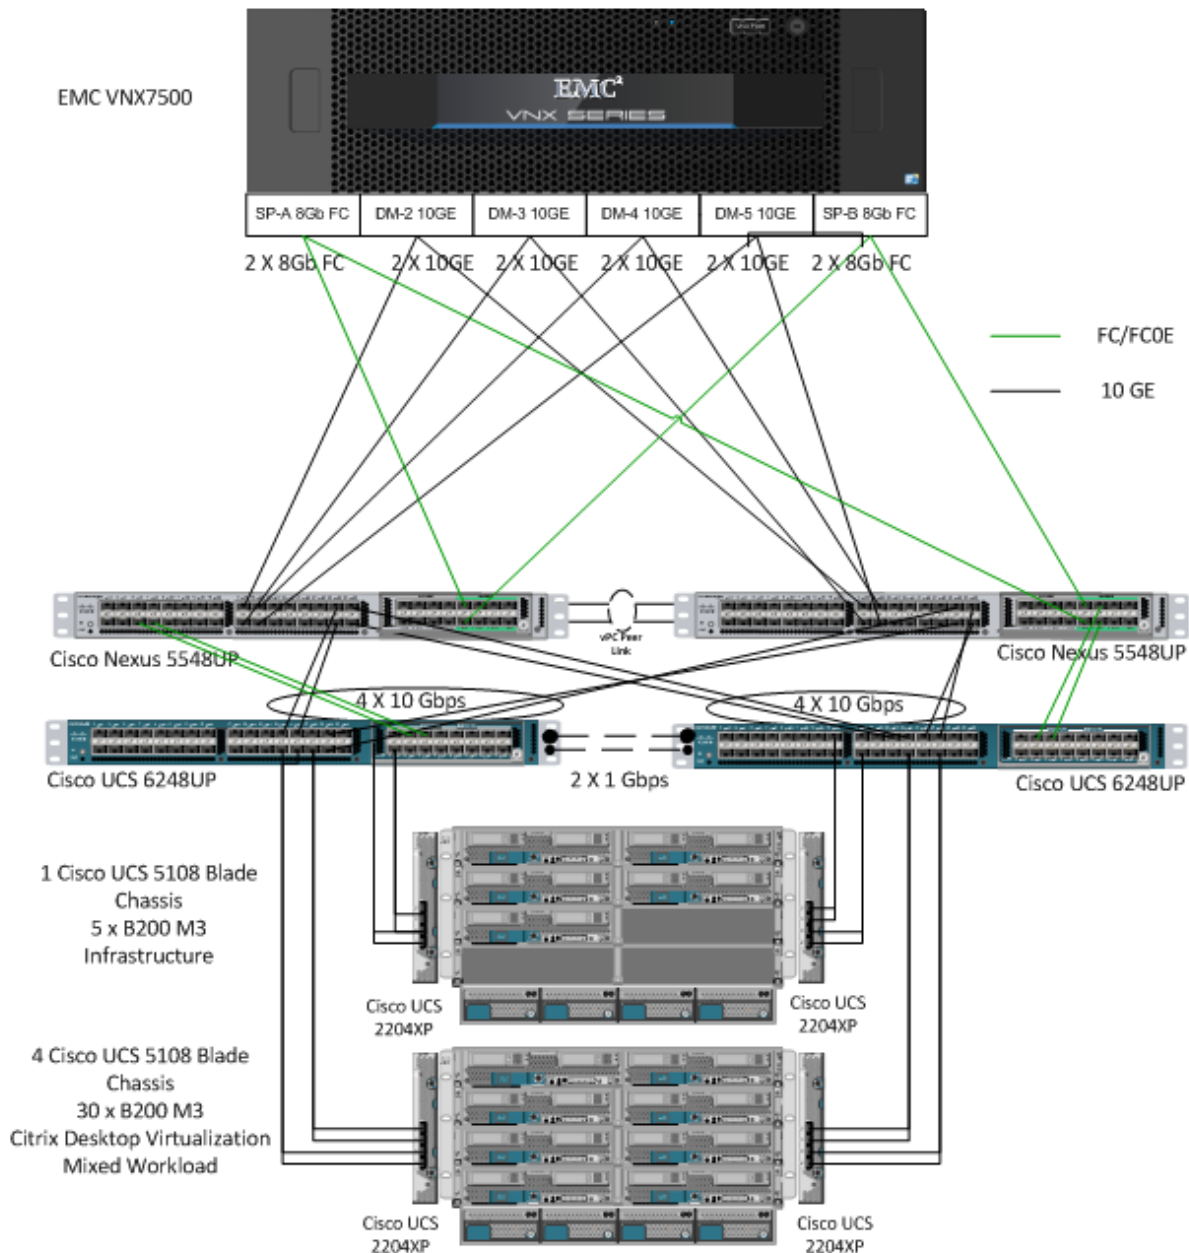


Figure 10 above captures the architectural diagram for the purpose of this study. The architecture is divided into four distinct layers:

- 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 11. Detailed Configuration Architecture



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) and it is beyond the scope of this document. More details on each step can be found in the following documents:

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

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

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

Name	Type	State	Version
ucs-k9-bundle-b-series.2.0.1s.B.bin	B Series Bundle	Active	2.0(1s)B
ucs-k9-bundle-b-series.2.0.3a.B.bin	B Series Bundle	Active	2.0(3a)B
ucs-k9-bundle-b-series.2.0.3c.B.bin	B Series Bundle	Active	2.0(3c)B
ucs-k9-bundle-b-series.2.0.4a.B.bin	B Series Bundle	Active	2.0(4a)B
ucs-k9-bundle-b-series.2.1.1a.B.bin	B Series Bundle	Active	2.1(1a)B
ucs-k9-bundle-c-series.2.0.1s.C.bin	C Series Bundle	Active	2.0(1s)C
ucs-k9-bundle-infra.2.0.1s.A.bin	Infrastructure Bundle	Active	2.0(1s)A
ucs-k9-bundle-infra.2.0.3a.A.bin	Infrastructure Bundle	Active	2.0(3a)A
ucs-k9-bundle-infra.2.0.3c.A.bin	Infrastructure Bundle	Active	2.0(3c)A
ucs-k9-bundle-infra.2.0.4a.A.bin	Infrastructure Bundle	Active	2.0(4a)A
ucs-k9-bundle-infra.2.1.1a.A.bin	Infrastructure Bundle	Active	2.1(1a)A

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

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

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

Equipment Servers LAN SAN VM Admin

Filter: Fabric Interconne...

Fabric Interconnects

- Fabric Interconnect A (primary)
 - Fixed Module
 - Ethernet Ports
 - Port 1
 - Port 2
 - Port 3
 - Port 4
 - Port 5
 - Port 6
 - Port 7
 - Port 8
 - Port 9
 - Port 10
 - Port 11
 - Port 12
 - Port 13
 - Port 14
 - Port 15
 - Port 16
 - Port 17
 - Port 18
 - Port 19
 - Port 20
 - Port 21
 - Port 22
 - Port 23

General Faults Events FSM Statistics

Fault Summary

0 0 0 0

Status

Overall Status: **Up**

Additional Info:

Admin State: **Enabled** ←

Actions

- Enable Port
- Disable Port
- Reconfigure
- Unconfigure
- Show Interface

Physical Display

Up Admin Down Fail Link Down

Properties

ID: 1 Slot ID: 1

User Label:

MAC: 54:7F:EE:76:D9:08

Mode: Fabric

Port Type: Physical **Role: Server**

Transceiver

Type: H10GB CU5M

Model: 74752-9047

Vendor: CISCO-MOLEX

Serial: MOC14322707

License Details

License State: License Ok

License Grace Period: 0

Equipment Servers LAN SAN VM Admin

Filter: Fabric Interconne...

Fabric Interconnects

- Fabric Interconnect A (primary)
 - Fixed Module
 - Ethernet Ports

Filter: Export Print Role: All Unconfigured Network Server FCoE Uplink Unfiled Uplink Appliance Storage FCoE Storage Unfiled Storage Monitor

Slot	Port ID	MAC	Role	Type	Overall Status	Administrative State
1	1	54:7F:EE:76:D9:08	Server	Physical	Up	Enabled
1	2	54:7F:EE:76:D9:09	Server	Physical	Up	Enabled
1	3	54:7F:EE:76:D9:0A	Server	Physical	Up	Enabled
1	4	54:7F:EE:76:D9:0B	Server	Physical	Up	Enabled
1	5	54:7F:EE:76:D9:0C	Server	Physical	Up	Enabled
1	6	54:7F:EE:76:D9:0D	Server	Physical	Up	Enabled
1	7	54:7F:EE:76:D9:0E	Server	Physical	Up	Enabled
1	8	54:7F:EE:76:D9:0F	Server	Physical	Up	Enabled
1	9	54:7F:EE:76:D9:10	Server	Physical	Up	Enabled
1	10	54:7F:EE:76:D9:11	Server	Physical	Up	Enabled
1	11	54:7F:EE:76:D9:12	Server	Physical	Up	Enabled
1	12	54:7F:EE:76:D9:13	Server	Physical	Up	Enabled
1	13	54:7F:EE:76:D9:14	Server	Physical	Up	Enabled
1	14	54:7F:EE:76:D9:15	Server	Physical	Up	Enabled
1	15	54:7F:EE:76:D9:16	Server	Physical	Up	Enabled
1	16	54:7F:EE:76:D9:17	Server	Physical	Up	Enabled
1	17	54:7F:EE:76:D9:18	Server	Physical	Up	Enabled
1	18	54:7F:EE:76:D9:19	Server	Physical	Up	Enabled
1	19	54:7F:EE:76:D9:1A	Server	Physical	Up	Enabled
1	20	54:7F:EE:76:D9:1B	Server	Physical	Up	Enabled
1	21	54:7F:EE:76:D9:1C	Server	Physical	Up	Enabled
1	22	54:7F:EE:76:D9:1D	Server	Physical	Up	Enabled
1	23	54:7F:EE:76:D9:1E	Server	Physical	Up	Enabled
1	24	54:7F:EE:76:D9:1F	Server	Physical	Up	Enabled
1	25	54:7F:EE:76:D9:20	Server	Physical	Up	Enabled
1	26	54:7F:EE:76:D9:21	Server	Physical	Up	Enabled
1	27	54:7F:EE:76:D9:22	Server	Physical	Up	Enabled
1	28	54:7F:EE:76:D9:23	Server	Physical	Up	Enabled
1	29	54:7F:EE:76:D9:24	Server	Physical	Up	Enabled
1	30	54:7F:EE:76:D9:25	Server	Physical	Up	Enabled
1	31	54:7F:EE:76:D9:26	Server	Physical	Up	Enabled
1	32	54:7F:EE:76:D9:27	Server	Physical	Up	Enabled

5 Configure and enable uplink Ethernet ports:

Equipment Servers LAN SAN VM Admin

Filter: Fabric Interconne...

Fabric Interconnects

- Fabric Interconnect A (primary)
 - Fixed Module
 - Expansion Module 2
 - Ethernet Ports
 - Port 1
 - Port 2
 - Port 3
 - Port 4
 - Port 5
 - Port 6
 - Port 7
 - Port 8
 - Port 9
 - Port 10
 - FC Ports
 - Fans
 - PSUs
 - Fabric Interconnect B (subordinate)

General Faults Events FSM Statistics

Fault Summary

0 0 0 0

Status

Overall Status: **Up**

Additional Info:

Admin State: **Enabled** ←

Actions

- Enable Port
- Disable Port
- Reconfigure
- Unconfigure
- Show Interface

Physical Display

Up Admin Down Fail Link Down

Properties

ID: 5 Slot ID: 2

User Label:

MAC: 54:7F:EE:3E:F0:B4

Mode: Trunk

Port Type: Physical **Role: Network**

Transceiver

Type: H10GB CU1M

Model: 74752-9044

Vendor: CISCO-MOLEX

Serial: MOC14295773

Slot	Port ID	MAC	IF Role	IF Type	Overall Status	Administrative State
2	1	54:7F:EE:3E:FD-B0	Server	Physical	Sfp Not Present	Enabled
2	2	54:7F:EE:3E:FD-B1	Server	Physical	Sfp Not Present	Enabled
2	3	54:7F:EE:3E:FD-B2	Server	Physical	Sfp Not Present	Enabled
2	4	54:7F:EE:3E:FD-B3	Server	Physical	Sfp Not Present	Enabled
2	5	54:7F:EE:3E:FD-B4	Network	Physical	Up	Enabled
2	6	54:7F:EE:3E:FD-B5	Network	Physical	Up	Enabled
2	7	54:7F:EE:3E:FD-B6	Network	Physical	Up	Enabled
2	8	54:7F:EE:3E:FD-B7	Network	Physical	Up	Enabled
2	9	54:7F:EE:3E:FD-B8	Network	Physical	Sfp Not Present	Enabled
2	10	54:7F:EE:3E:FD-B9	Network	Physical	Sfp Not Present	Enabled

and FC uplink ports:

General | **Faults** | **Events** | **FSM** | **Statistics**

Fault Summary
 0 0 0 0

Status
 Overall Status: **Up**
 Additional Info:
 Admin State: **Enabled**

Actions
 Enable Port
 Disable Port
 Configure as Uplink Port
 Configure as FC Storage Port
 Show Interface

Physical Display
 Legend: Up (Green), Admin Down (Yellow), Fail (Red), Link Down (Orange)

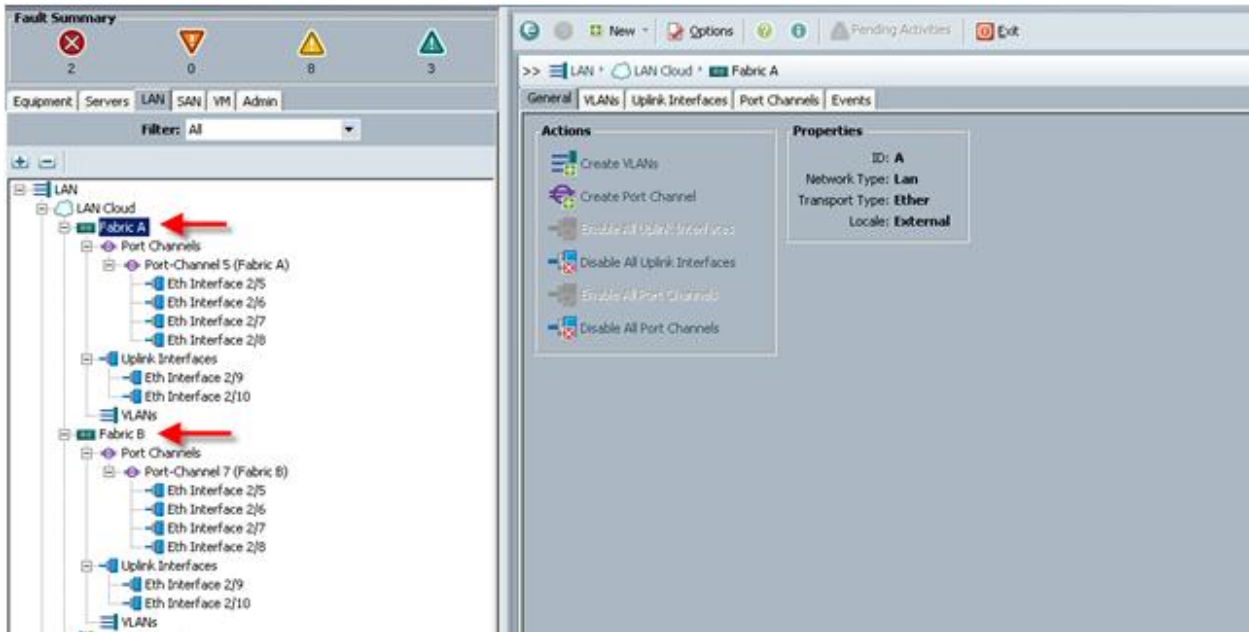
Properties
 ID: 13 Slot ID: 2
 User Label:
 WWPN: 20:4D:54:7F:EE:76:D9:00 Mode: **N Proxy**
 Port Type: **Physical** Negotiated Speed: **8 Gbps**
 VSAN: Fabric dual/vsan default ...

License Details
 License State: **License Ok**
 License Grace Period: 12636000

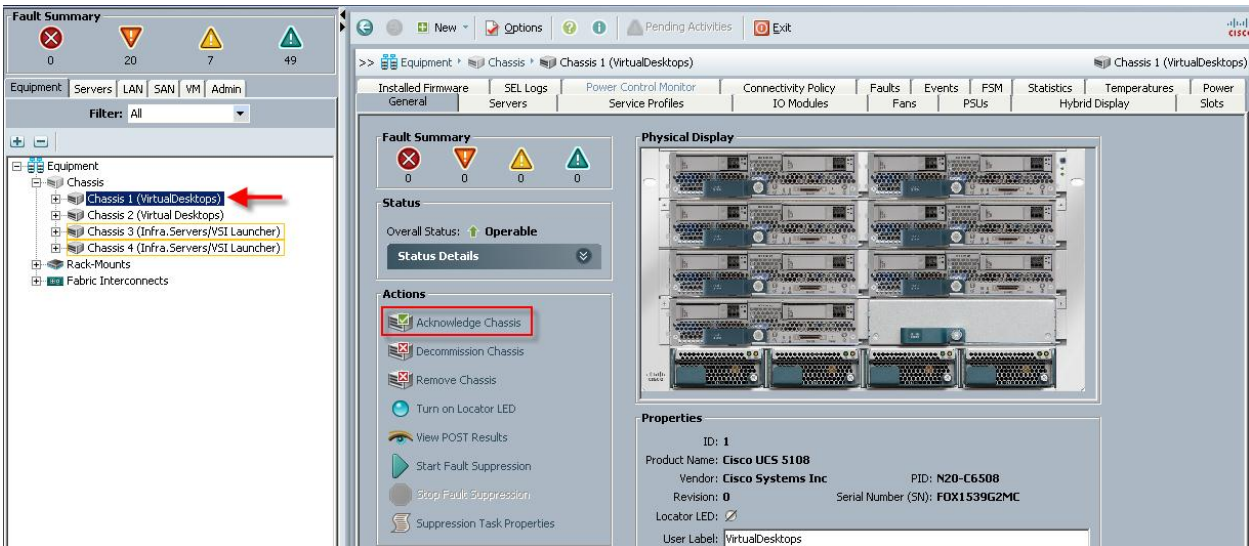
Slot	Port ID	WWPN	IF Role	IF Type	Overall Status	Administrative State
2	11	20:4B:54:7F:EE:76:D9:00	Network	Physical	Sfp Not Present	Enabled
2	12	20:4C:54:7F:EE:76:D9:00	Network	Physical	Sfp Not Present	Enabled
2	13	20:4D:54:7F:EE:76:D9:00	Network	Physical	Up	Enabled
2	14	20:4E:54:7F:EE:76:D9:00	Network	Physical	Up	Enabled
2	15	20:4F:54:7F:EE:76:D9:00	Network	Physical	Sfp Not Present	Enabled
2	16	20:50:54:7F:EE:76:D9:00	Network	Physical	Sfp Not Present	Enabled

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.

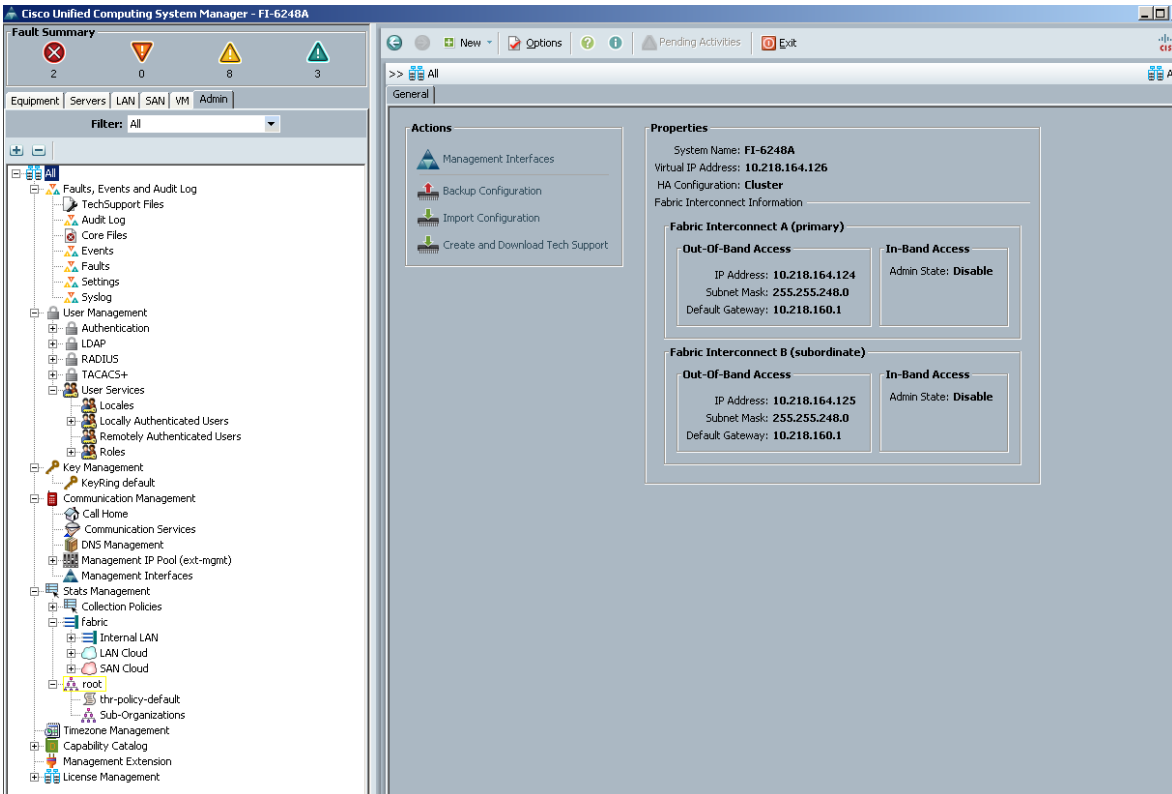
- 5a On the LAN tab in the Navigator pane, configure the required Port Channels and Uplink Interfaces on both Fabric Interconnects:



- 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.

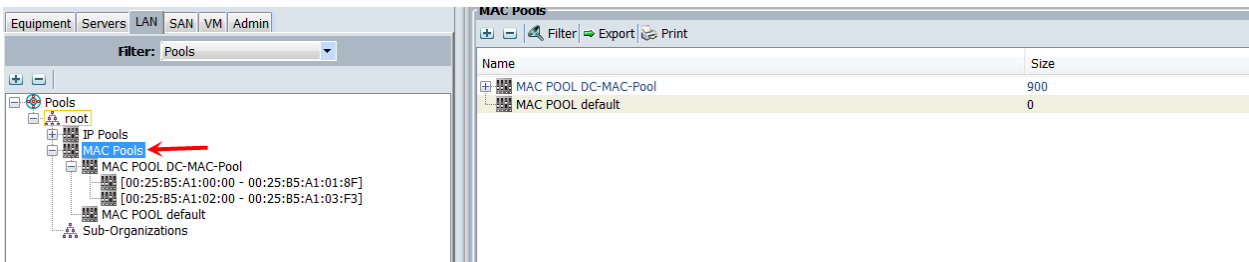


- Use the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and Licensing. Configuring your Management IP Pool (which provides IP based access to the KVM of each UCS Blade Server,) Time zone Management (including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

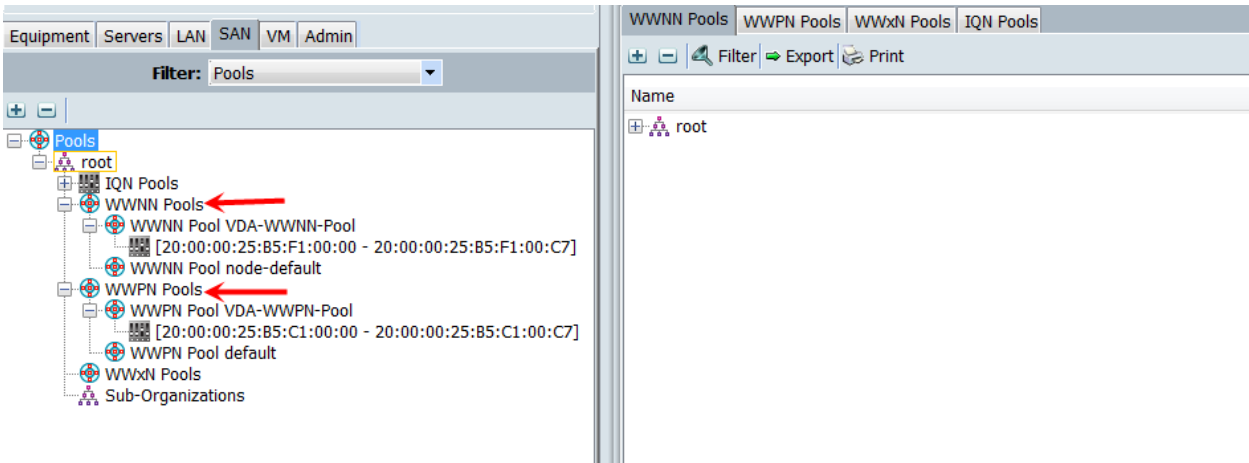


8 Create all the pools: MAC pool, WWPN pool, WWNN pool, UUID pool, Server pool.

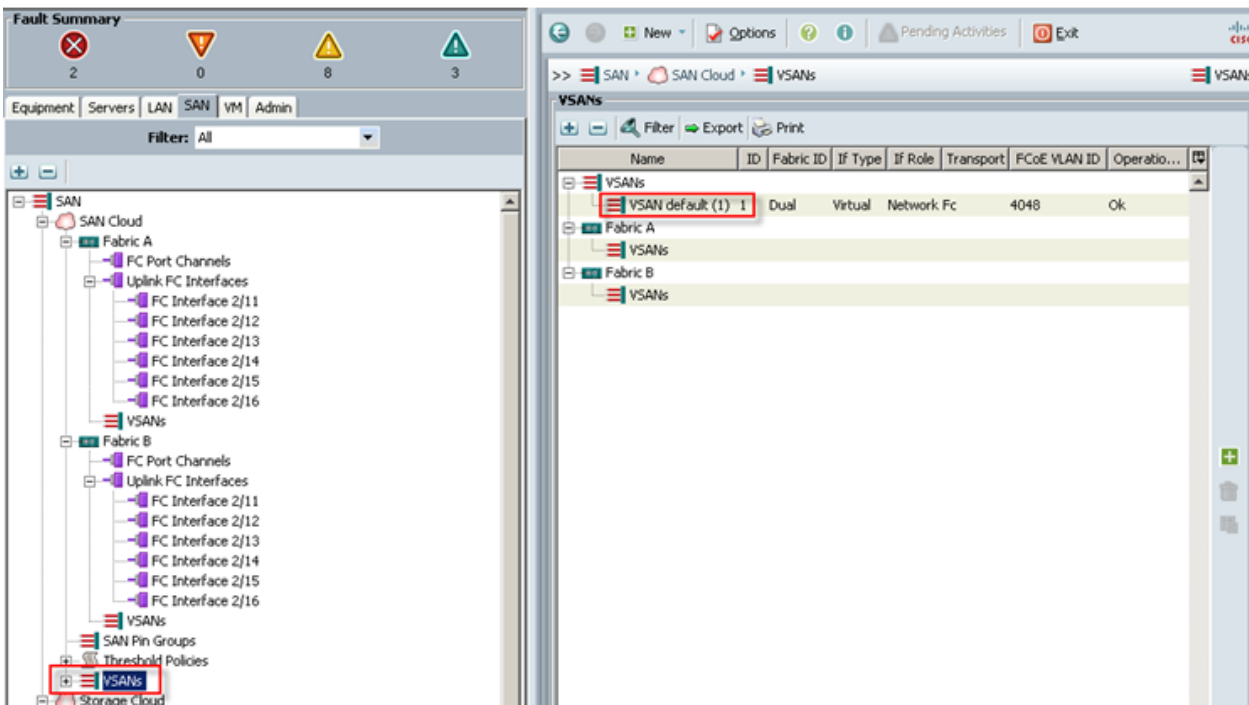
8.1 From the LAN tab in the navigator, under the Pools node, we created a MAC address pool of sufficient size for the environment. In this project, we created a single pool with two address ranges for expandability.



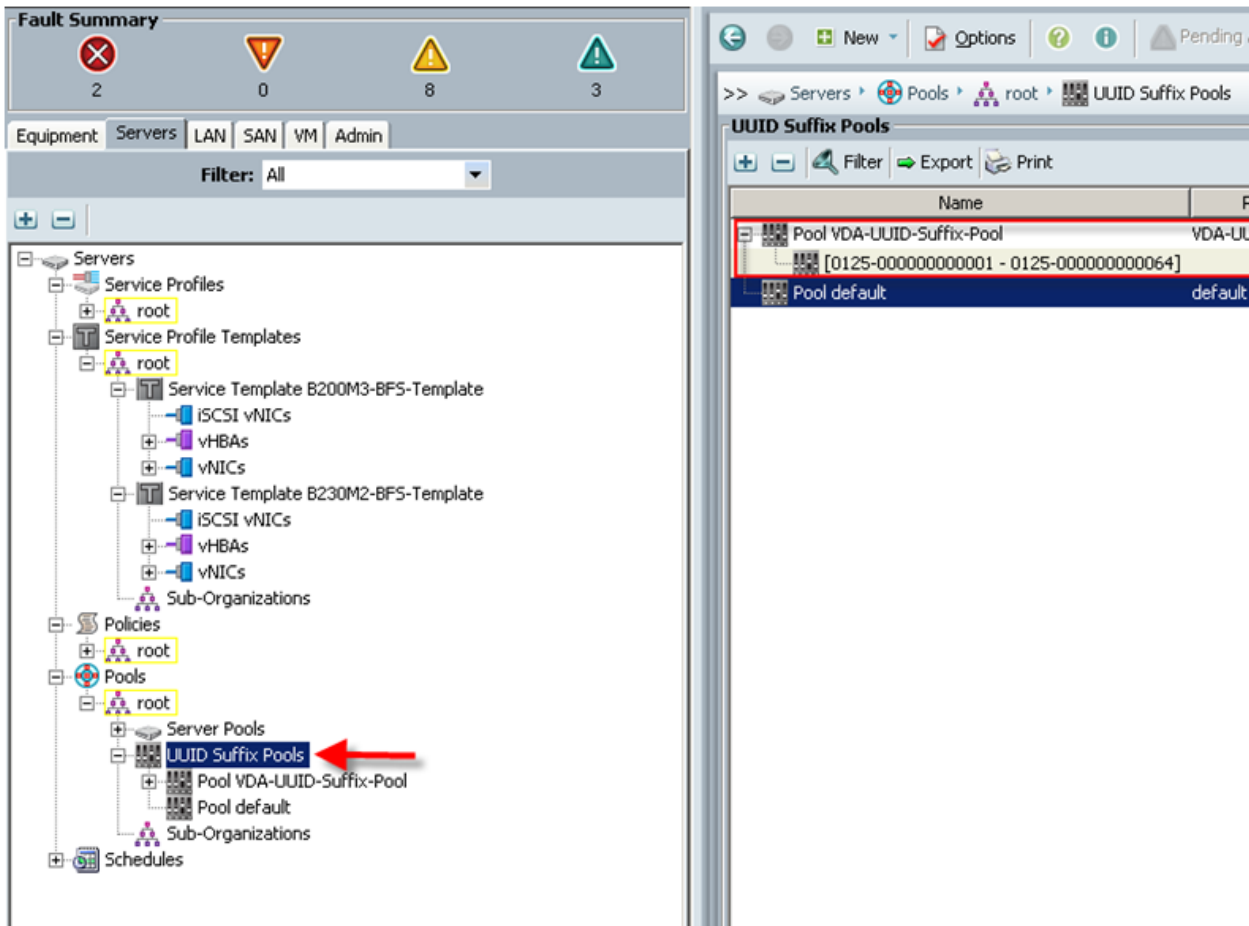
8.2 For Fiber Channel connectivity, WWNN and WWPN pools must be created from the SAN tab in the navigator pane, in the Pools node:




8.3 For this project, we used a single VSAN, the default VSAN with ID 1:



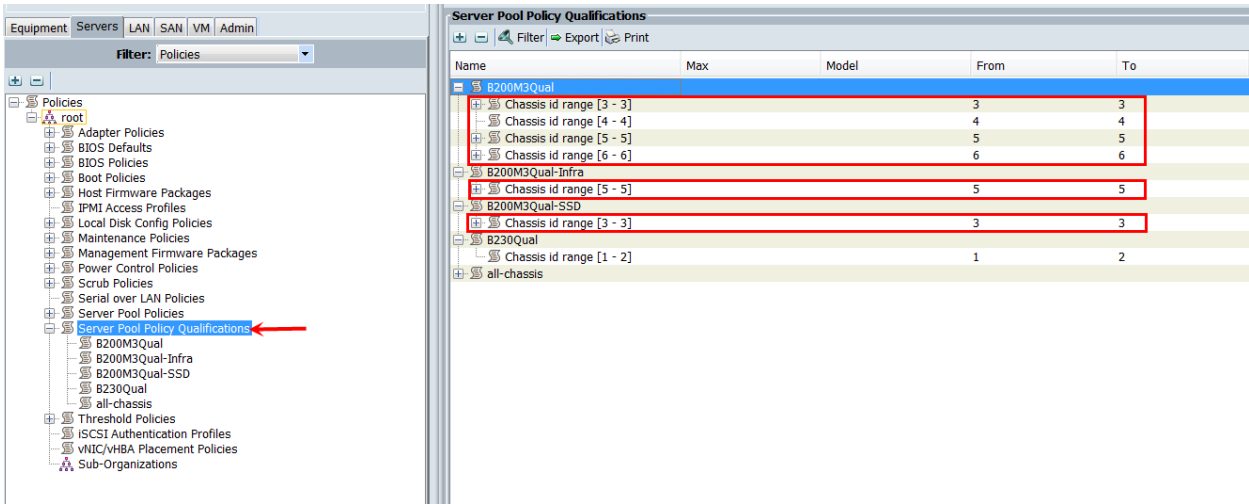
8.4 The next pool we created is the Server UUID pool. On the Servers tab in the Navigator page under the Pools node we created a single UUID Pool for the test environment. Each UCS Blade Server requires a unique UUID to be assigned by its Service profile.



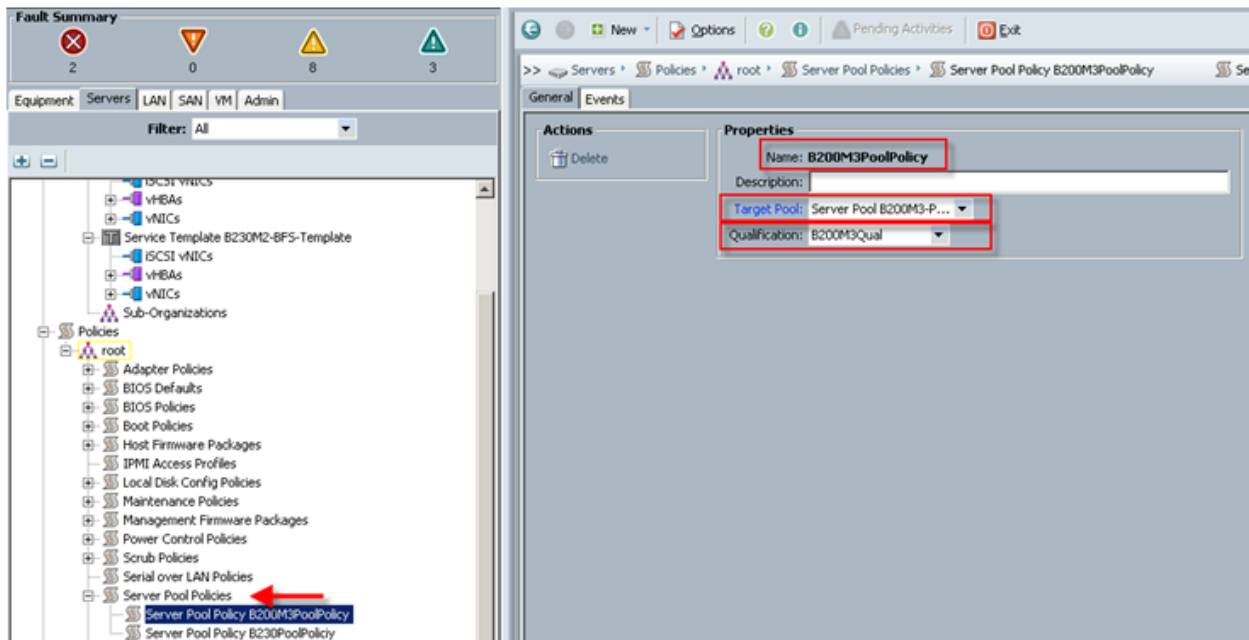
8.5 We created two Server Pools for use in our Service Profile Templates as selection criteria for automated profile association. Server Pools were created on the Servers tab in the navigation page under the Pools node. Only the pool name was created, no servers were added:



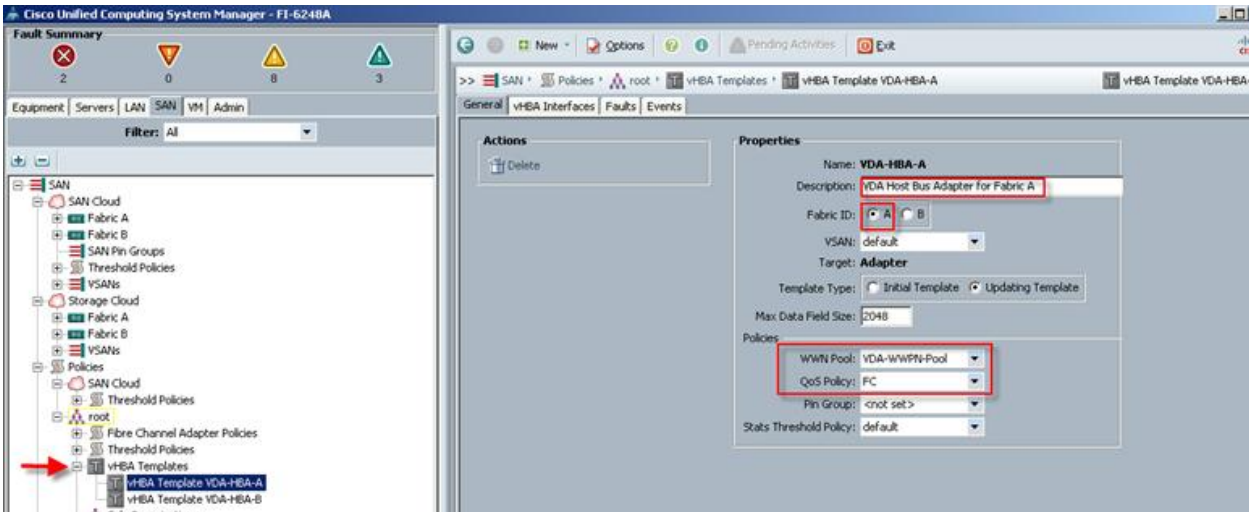
8.6 We created two Server Pool Policy Qualifications to identify the blade server model for placement into the correct pool using the Service Profile Template. In this case we used Chassis ids to select the servers. (We could have used slots or server models to make the selection.)



8.7 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.

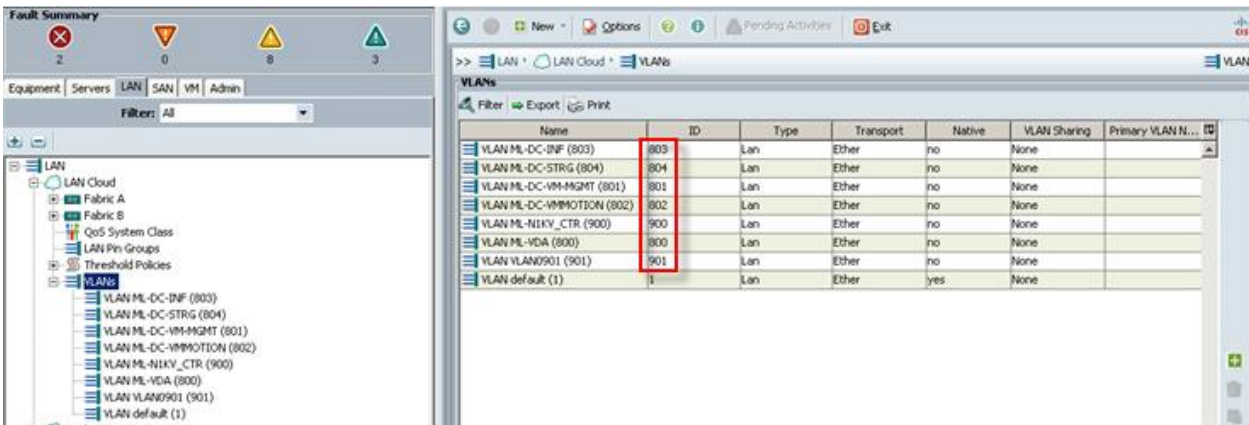


9 Virtual Host Bus Adapter templates were created for FC SAN connectivity from the SAN tab under the Policies node, one template for each fabric:



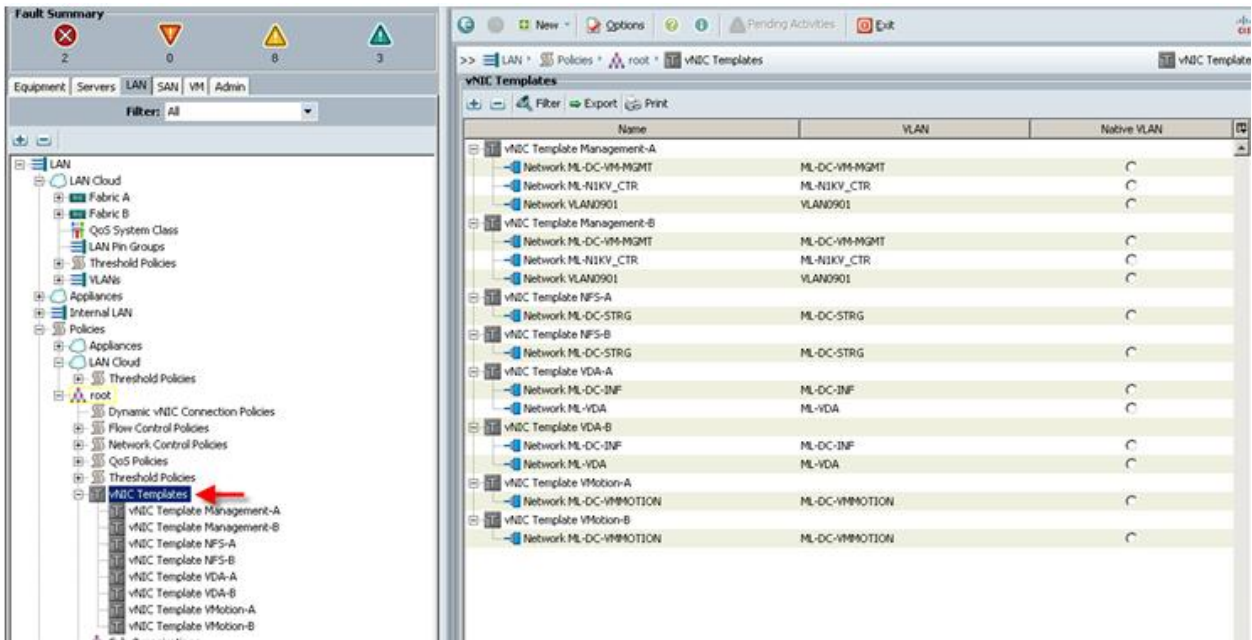
Create at least one HBA template for each Fabric Interconnect if block storage will be used. We used the WWPN pool created earlier and the QoS Policy created in section 5.2.4 below.

- 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. (N1KV management and VMware Management shared VLAN 801.

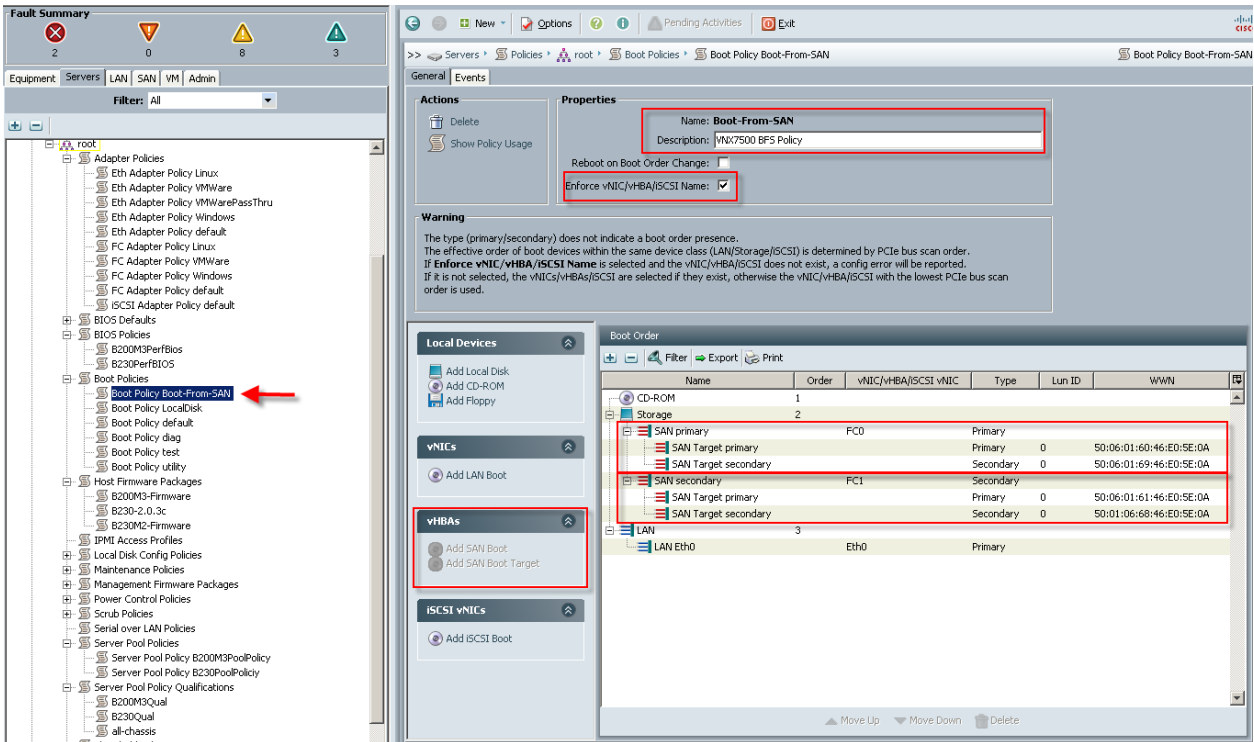
- 11 On the LAN tab in the navigator pane, under the policies node configure the vNIC templates that will be used in the Service Profiles. In this project, we utilize eight virtual NICs per host, four pairs, with each pair connected to both Fabric Interconnects for resiliency.



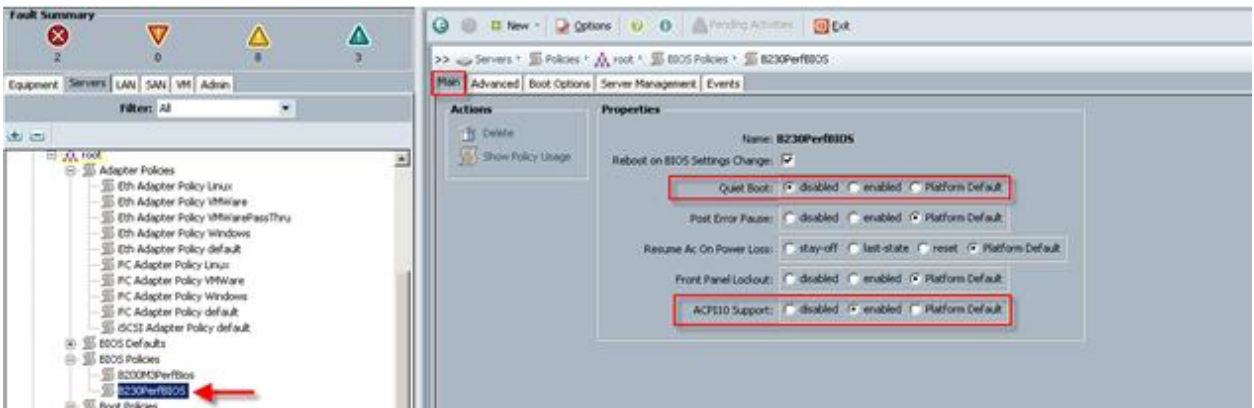
11a Create vNIC templates for both fabrics, check Enable Failover, select VLANs supported on adapter (optional,) set the MTU size, select the MAC Pool and QoS Policy, click OK.




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:



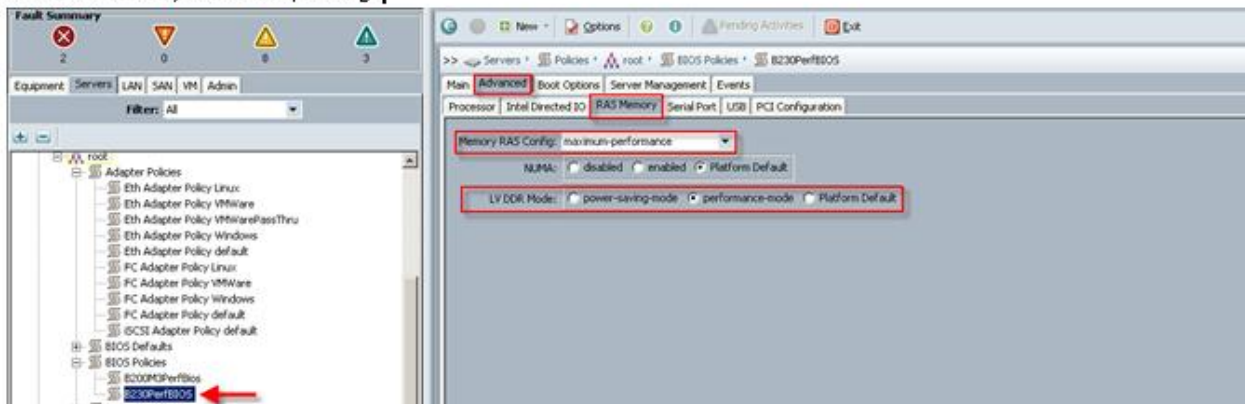
The Advanced tab, Processor settings:



The Advanced tab, Intel Directed IO tab settings:



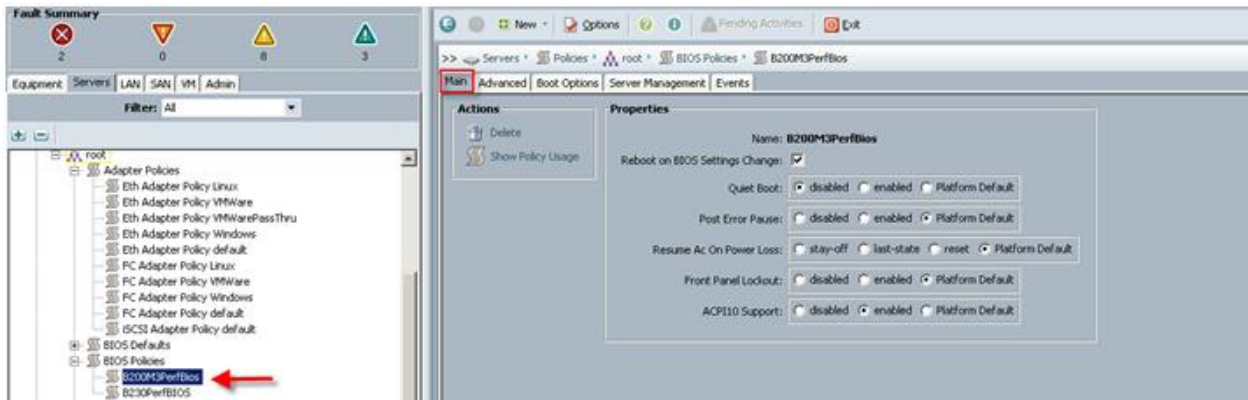
The Advanced tab, RAS Memory settings:



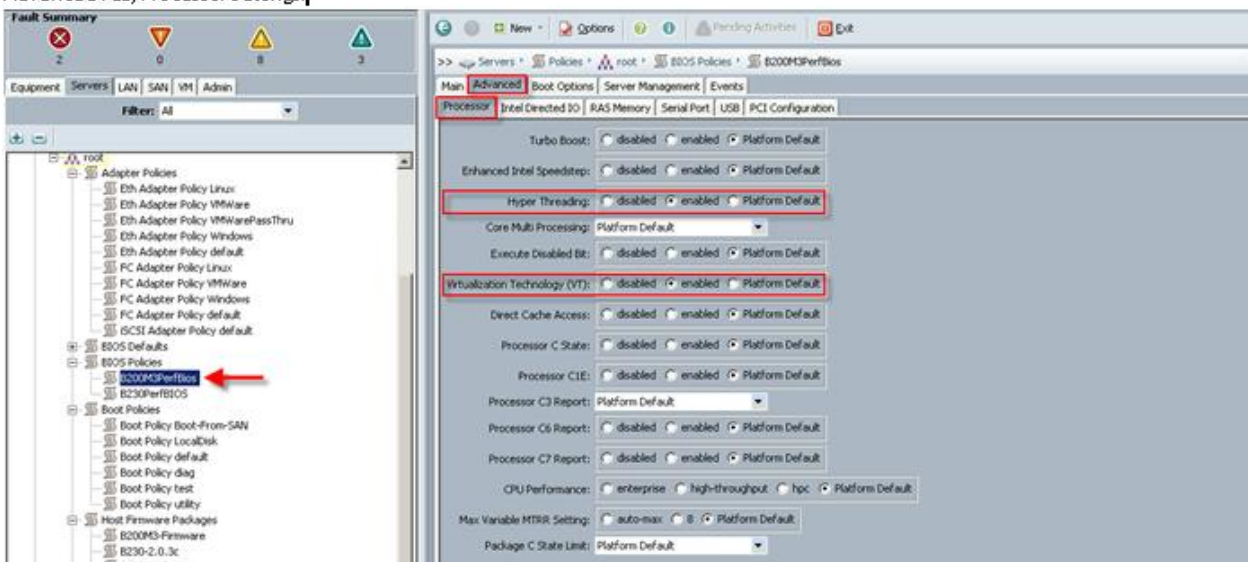
The remaining Advanced tab settings are at platform default or not configured. Similarly, the Boot Options and Server Management tab settings are at defaults.

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.

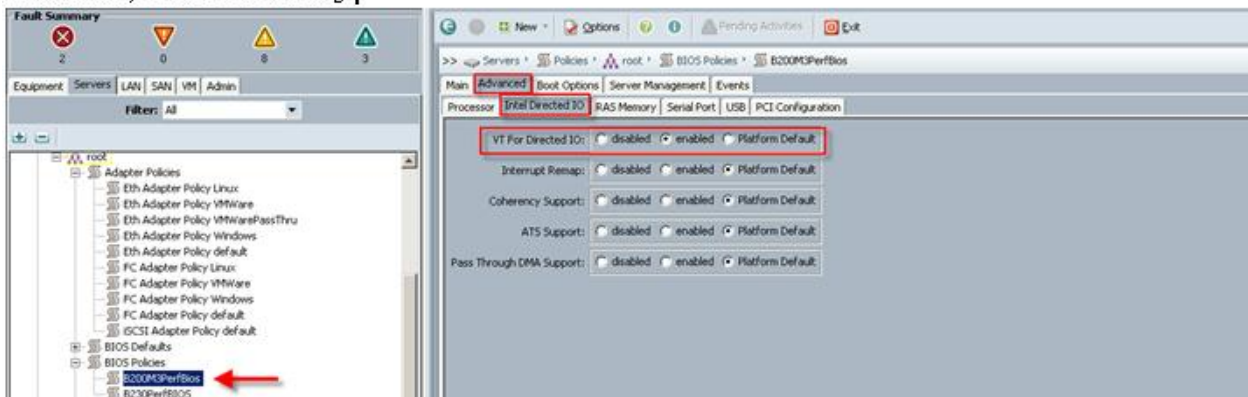
13a The following screen captures show the settings for the Cisco UCS B200 M3 blades used in this study:



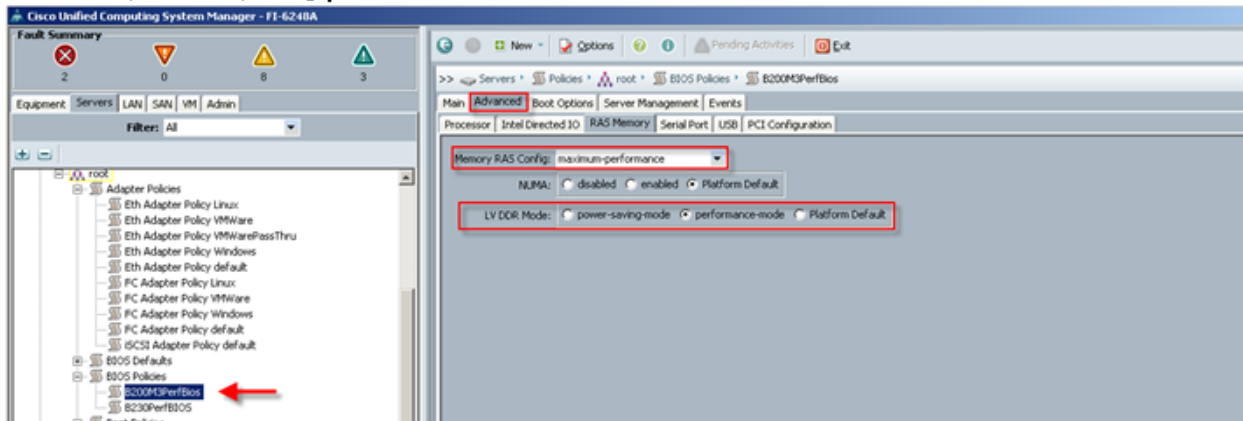
AdvancedTab, Processor settings:



AdvancedTab, Intel Directed IO settings:



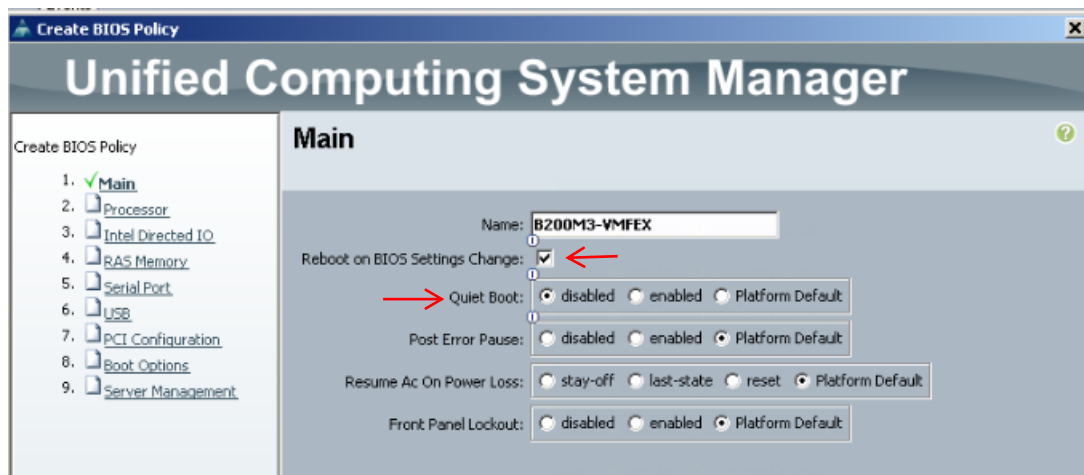
AdvancedTab, RAS Memory Settings:



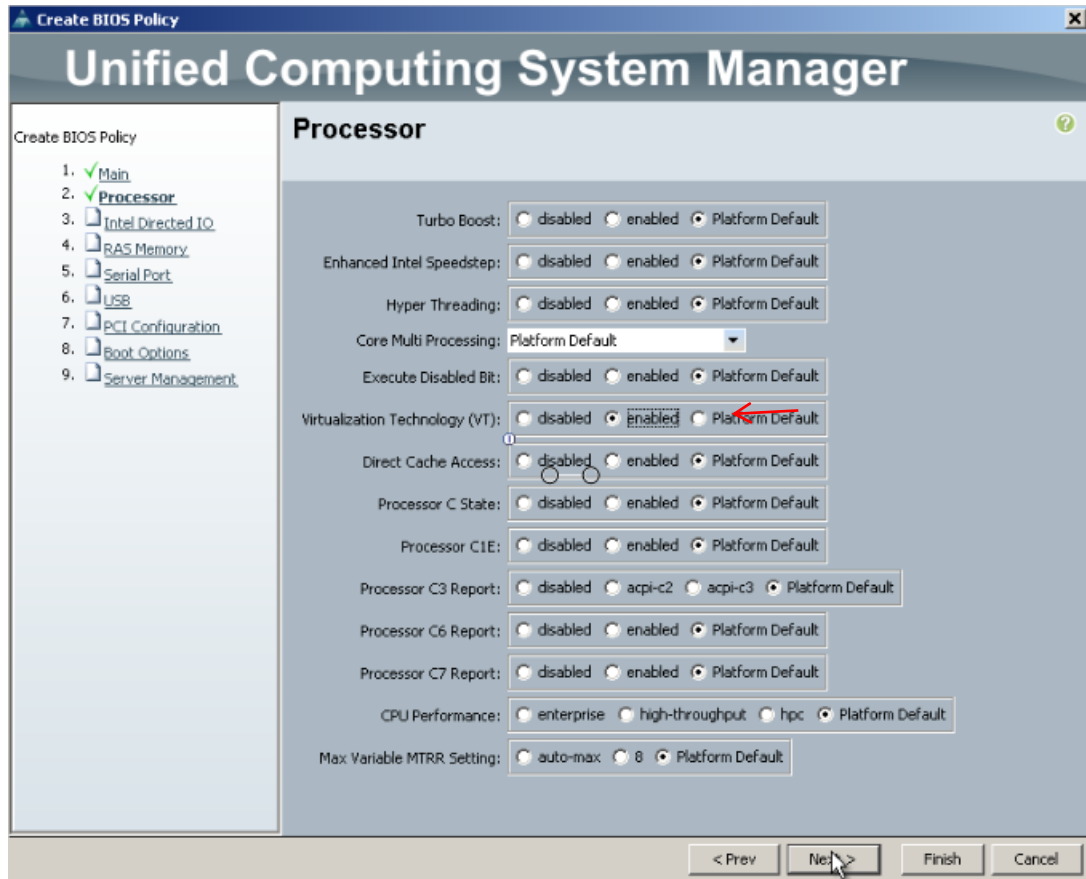
The remaining Advanced tab settings are at platform default or not configured. Similarly, the Boot Options and Server Management tab settings are at defaults.

13b. Following BIOS Policy for B200 M3 used with VM-FEX configuration:

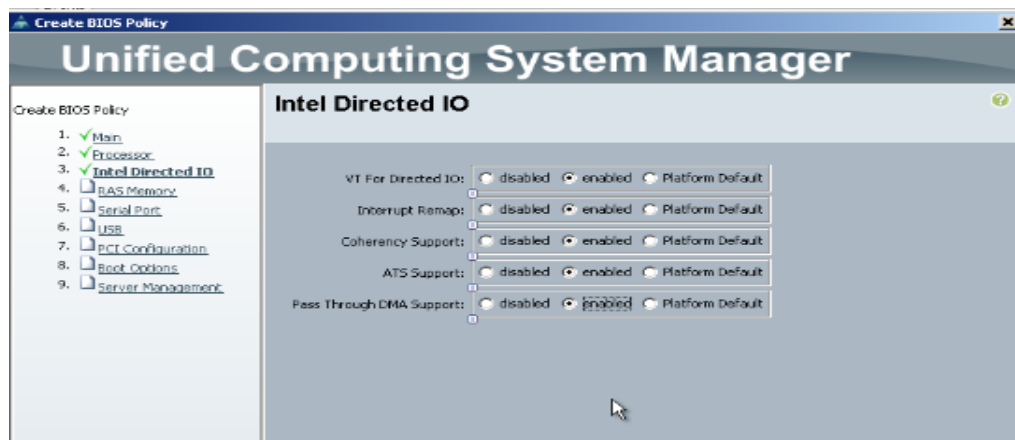
Main; Check box for Reboot on BIOS Settings change, Disable Quiet Boot. Click Next.



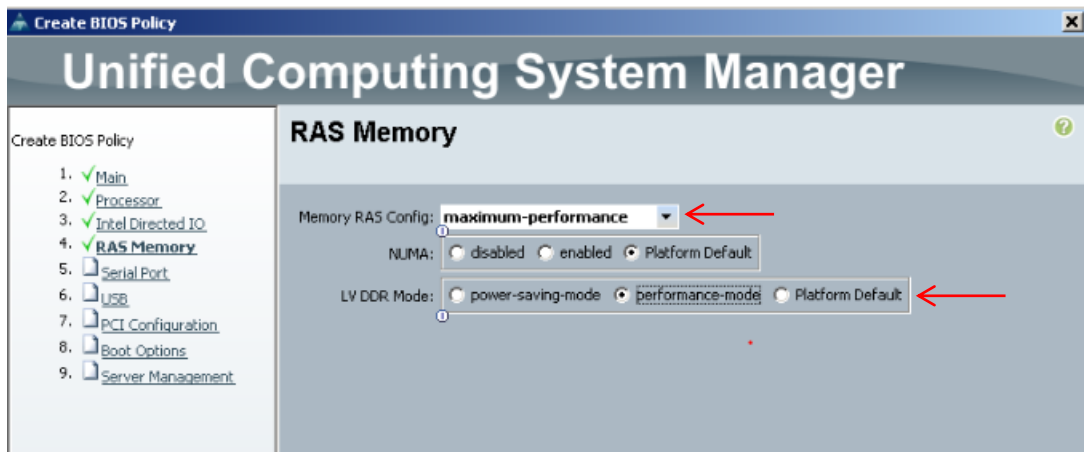
Processor; Enable Virtualization Technology (VT). Click Next.



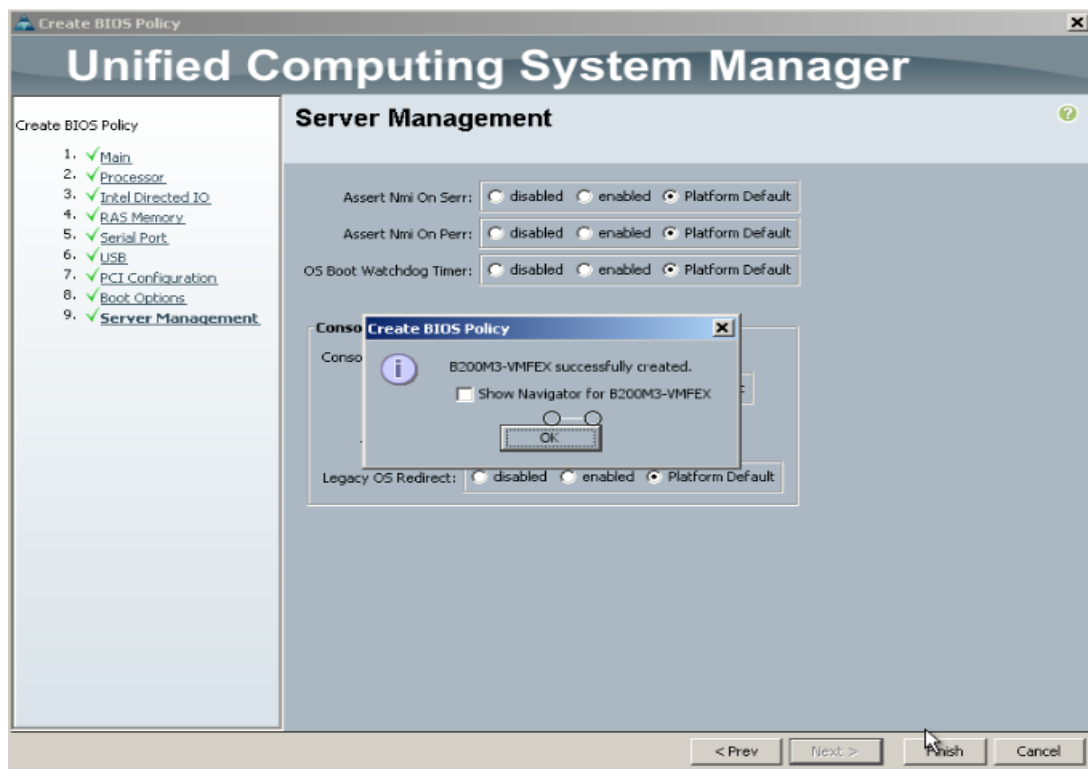
Intel Directed IO; enable VT for Directed IO, Interrupt Remap, Coherency Support, ATS Support, Pass Through DMA Support. Click Next.



RAS Memory; Memory RAS Config select maximum-performance
LV DDR Mode: enable performance-mode. Click Next.



Remaining sections from 5. **Serial Port** to 9. **Server management** set to perform-default and no changes were made. Click Finish. Click OK.



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 Host Firmware Package policies were set for Adapter and BIOS:

Properties
Name: B200-M3-2.1.1b
Description: Network/storage adapter and BIOS policy for 2.1 Delinear
Owner: Local
Blade Package:
Rack Package:

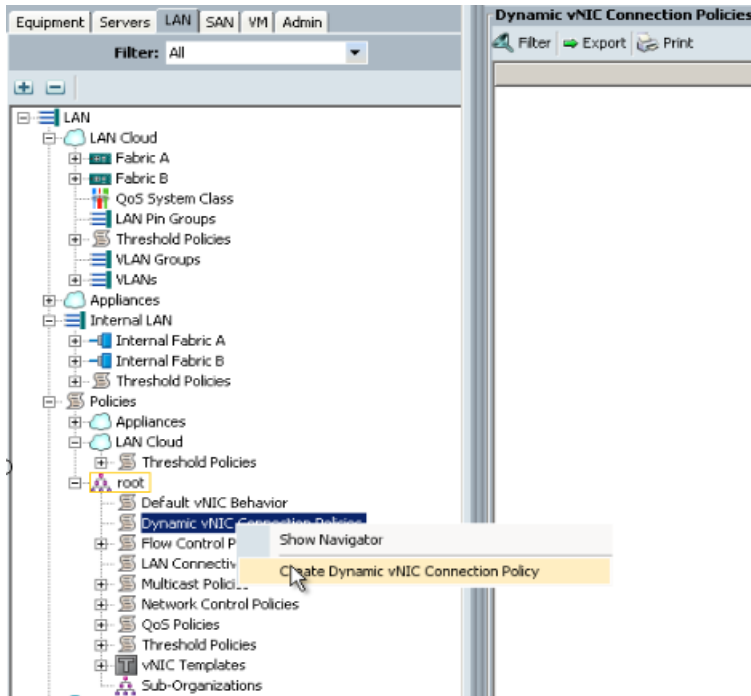
Select	Vendor	Model	PID	Presence	Version
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M51KR-B	N20-AB0002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M8KR	N20-A00002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M71KR-E	N20-AE0002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M72KR-E	N20-AE1002	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M61KR-1	N20-AD0102	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-AQ1002	N/A	<not set>
<input type="checkbox"/>	Broadcom Corp.	Broadcom 10GbE Onboard Adapter	N2XX-ABPC01	N/A	<not set>
<input type="checkbox"/>	Broadcom Corp.	Broadcom 10GbE Onboard Adapter	N2XX-ABPC02	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS P8E	N2XX-ACPC01	N/A	<not set>
<input type="checkbox"/>	Emulex Corp.	Emulex OCe10102-F	N2XX-AEFC01	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Intel 10GbE Adapter	N2XX-APCB01	N/A	<not set>
<input type="checkbox"/>	Qlogic Corp.	QLogic QLE8132	N2XX-AQPC01	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS VIC 1280	UCS-VIC-M82-8P	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M61KR-B	UCSB-MEZ-8RC-02	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M73KR-E	UCSB-MEZ-ELX-03	N/A	<not set>
<input type="checkbox"/>	Cisco Systems Inc	Cisco UCS M73KR-Q	UCSB-MEZ-QLG-03	N/A	<not set>
<input checked="" type="checkbox"/>	Cisco Systems Inc	Cisco UCS VIC 1240	UCSB-MEOM-40G-01	Present	2.1(1a)

Properties
Name: B200-M3-2.1.1b
Description: Network/storage adapter and BIOS policy for 2.1 Delinear
Owner: Local
Blade Package:
Rack Package:

Select	Vendor	Model	PID	Presence	Version
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B230 M2	B230-BA5E-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B440 M2	B440-BA5E-M2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B200 M1	N20-86620-1	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Cisco UCS B200 M1	N20-86620-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B250 M1	N20-86620-2	N/A	<not set>
<input type="checkbox"/>	Intel Corp.	Cisco UCS B250 M1	N20-86620-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B200 M2	N20-86625-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B250 M2	N20-86625-2	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B230 M1	N20-86730-1	N/A	<not set>
<input type="checkbox"/>	Cisco Systems, Inc.	Cisco UCS B440 M1	N20-86740-2	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	UCSB-B200-M3	Present	B200M3-2.1.1.0.10052012419
<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-B5E-SFF-C200	N/A	<not set>

15. Create Dynamic vNIC Connection Policy:

To create a dynamic vNIC connection policy go to LAN tab on Cisco UCS Manager. Select Policies → root → select Dynamic vNIC Connection Policy. Right click and select create Dynamic vNIC Connection Policy.

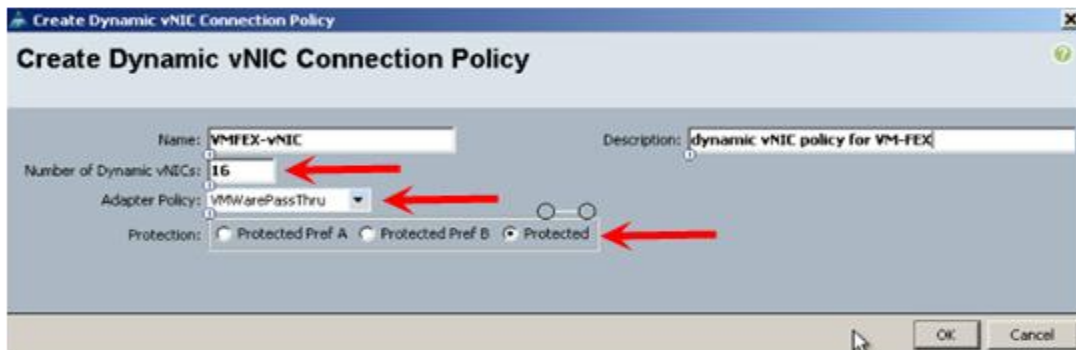


Select name for new dynamic vNIC connection policy, Set number of Dynamic vNICs as per the requirement.

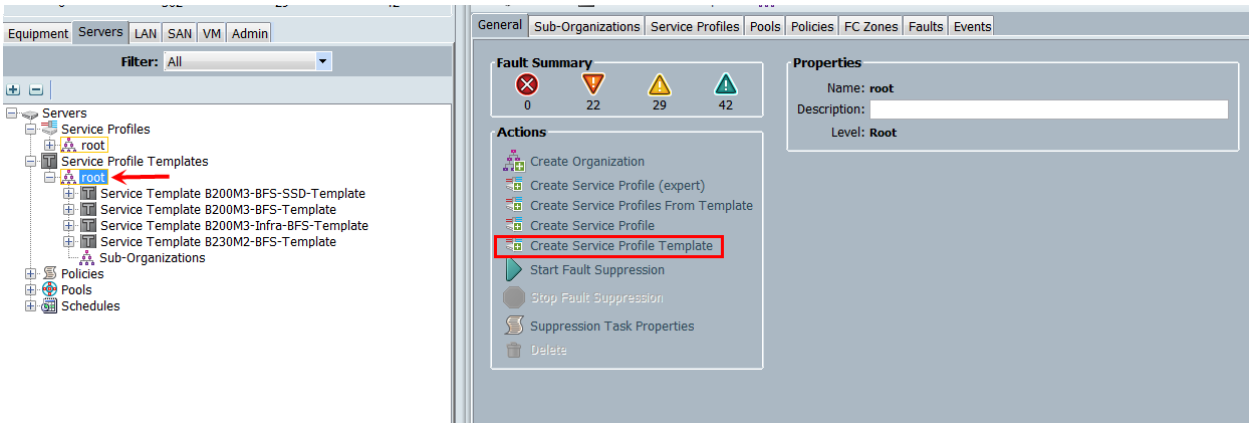
Adapter Policy: VMWarePassThru

Protection: Protected

Click OK. Click OK on the confirmation message.



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



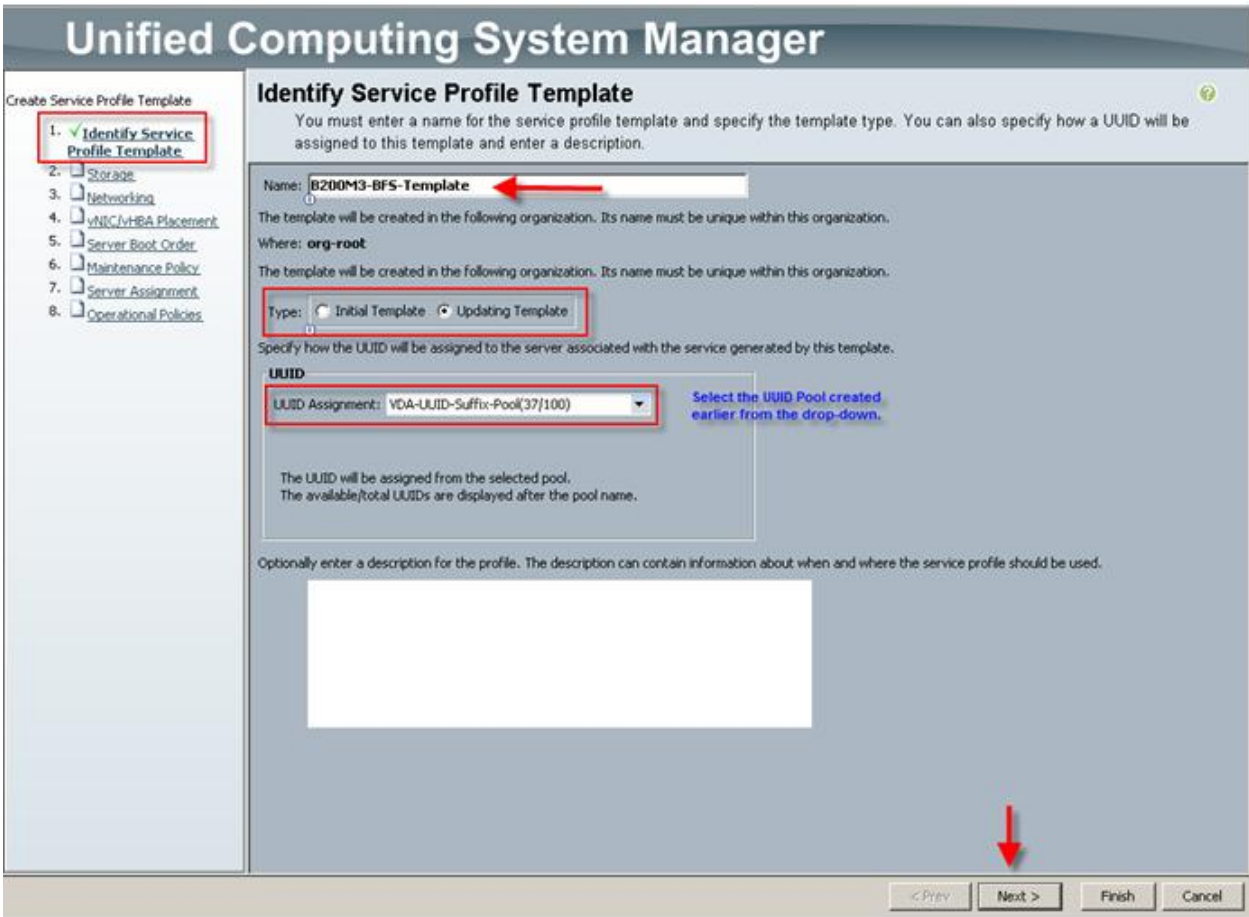
In this project, we created one template for the UCS B200 M3 Blade Server models used.

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.

- 16a 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.

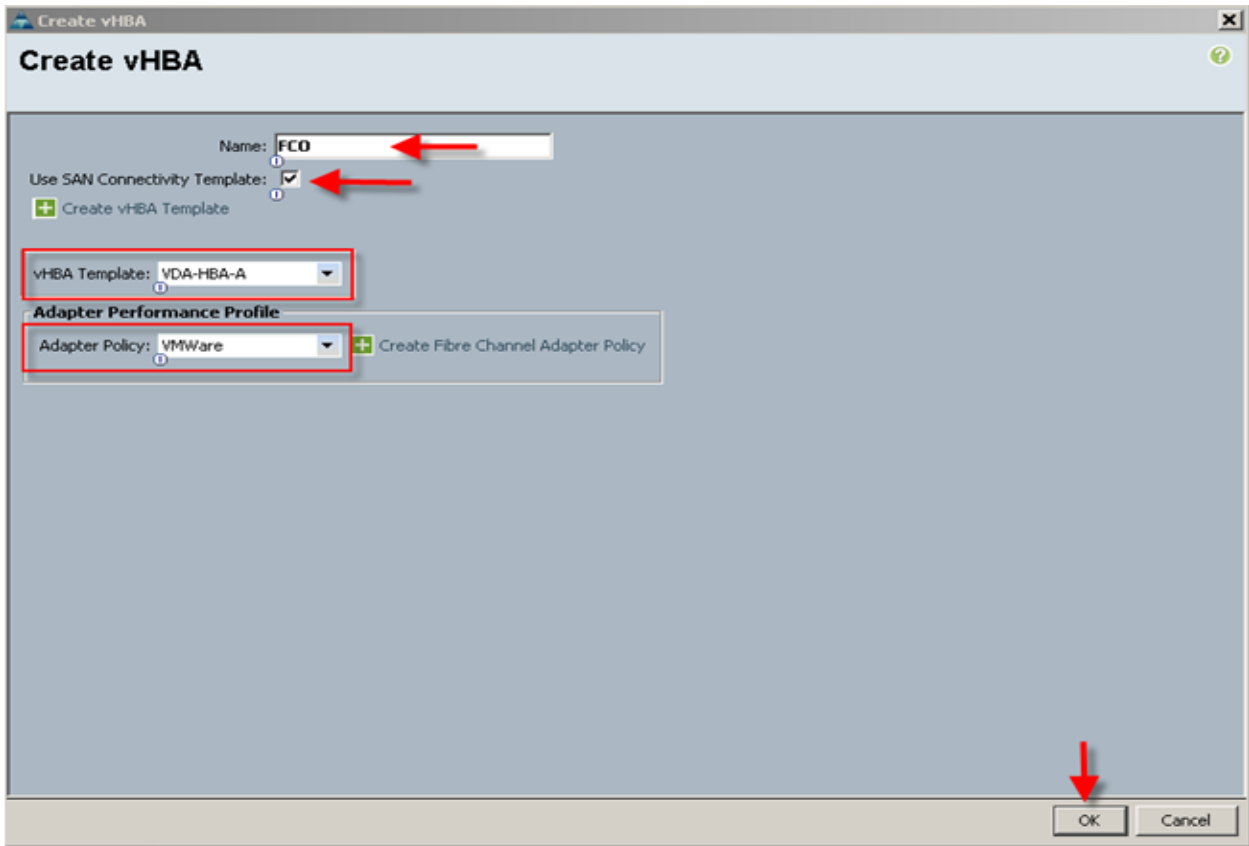


16b On the Storage page, we selected the Expert mode, we selected the WWNN Pool we created earlier from the drop down list and then click Add



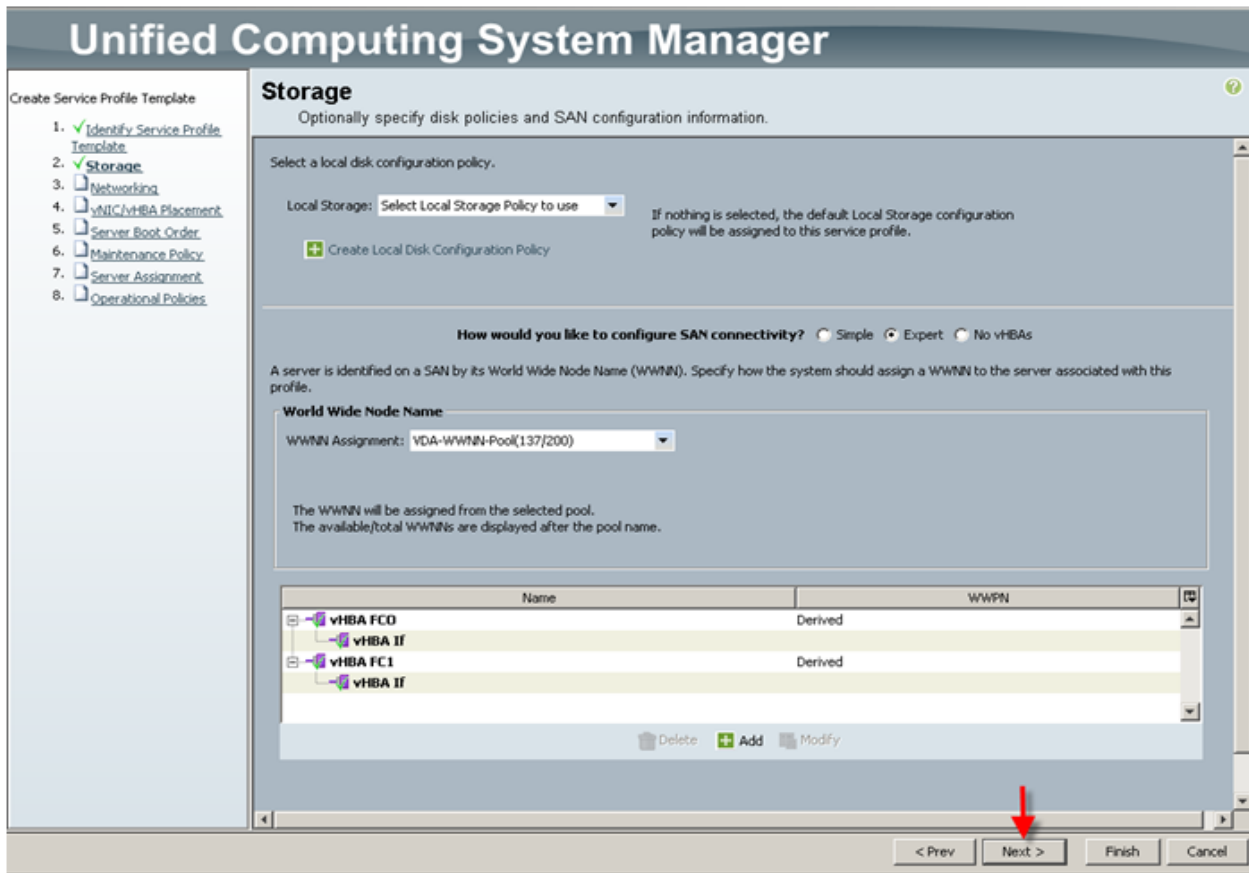
Note that we used the default Local Storage configuration in this project. Local drives on the blades were not used.

16c On the Create HBA page, we entered a name (FCO) and checked Use SAN Connectivity Template, which changed the display to the following:

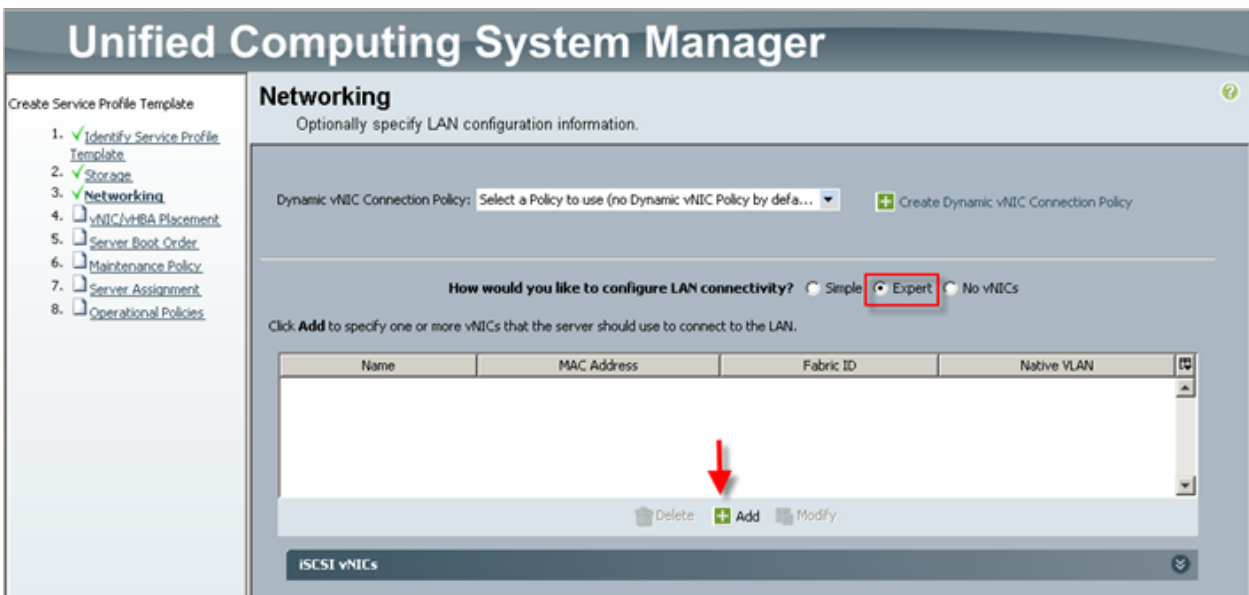


We selected the vHBA template for Fabric Interconnect A and the VMWare Adapter Policy from the drop downs, then clicked OK.

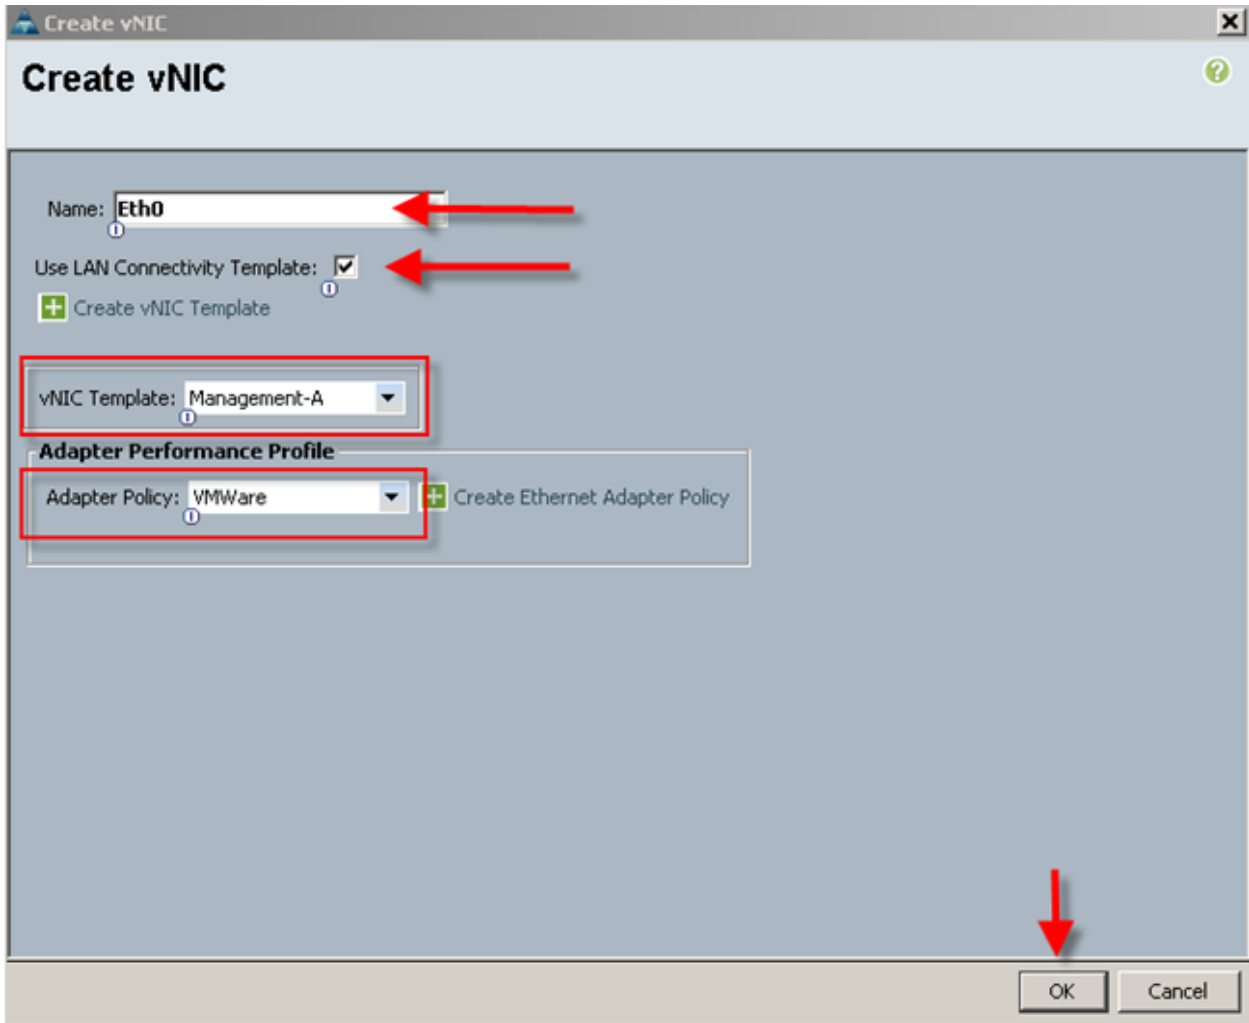
We repeated the process for FC1, choosing VDA-HBA-B for Fabric Interconnect B. The result is the Storage page that appears as follows:



16d Click Next to continue. We selected the Expert configuration option and clicked Add in the adapters window:



16e 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.



16f 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

- Identify Service Profile Template
- Storage
- Networking**
- vNIC/vHBA Placement
- Server Boot Order
- Maintenance Policy
- Server Assignment
- Operational Policies

Networking

Optionally specify LAN configuration information.

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	
vNIC Eth1	Derived	derived	
vNIC Eth2	Derived	derived	
vNIC Eth3	Derived	derived	
vNIC Eth4	Derived	derived	

ISCSI vNICs

< Prev **Next >** Finish Cancel

16g Click Next. We accepted the default placement and 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

vNIC/vHBA Placement ?

Specify how vNICs and vHBAs are placed on physical network adapters

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

Select Placement: Let System Perform Placem... + Create Placement Policy

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

Name	Address	Order	
vHBA FCD	Derived	1	
vHBA FC1	Derived	2	
vNIC Eth0	Derived	3	
vNIC Eth1	Derived	4	
vNIC Eth2	Derived	5	
vNIC Eth3	Derived	6	
vNIC Eth4	Derived	7	
vNIC Eth5	Derived	8	
vNIC Eth6	Derived	9	
vNIC Eth7	Derived	10	

▲ Move Up ▼ Move Down 🗑 Delete ↻ Reorder ✎ Modify

< Prev
Next >
Finish
Cancel

16h 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/iSCSI 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/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

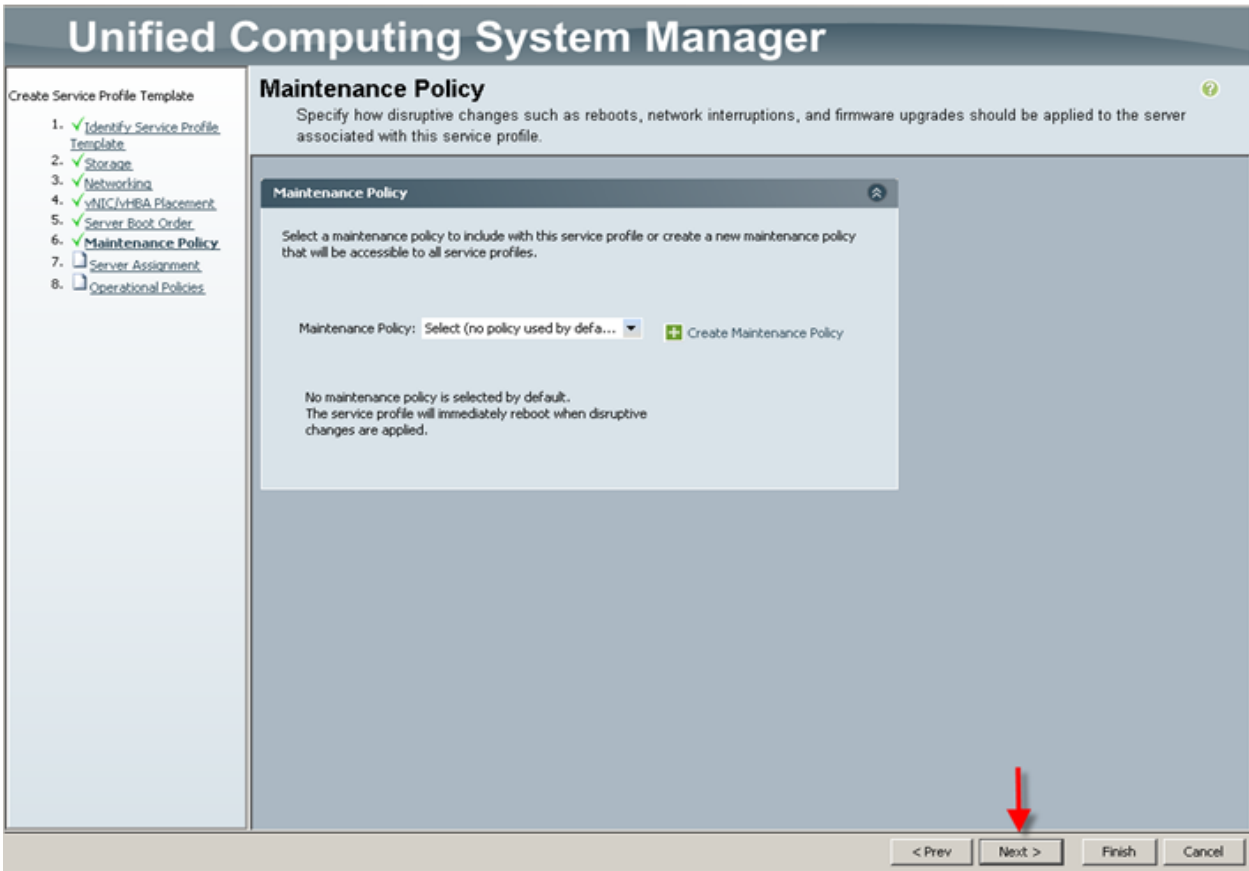
Boot Order

Filter Export Print

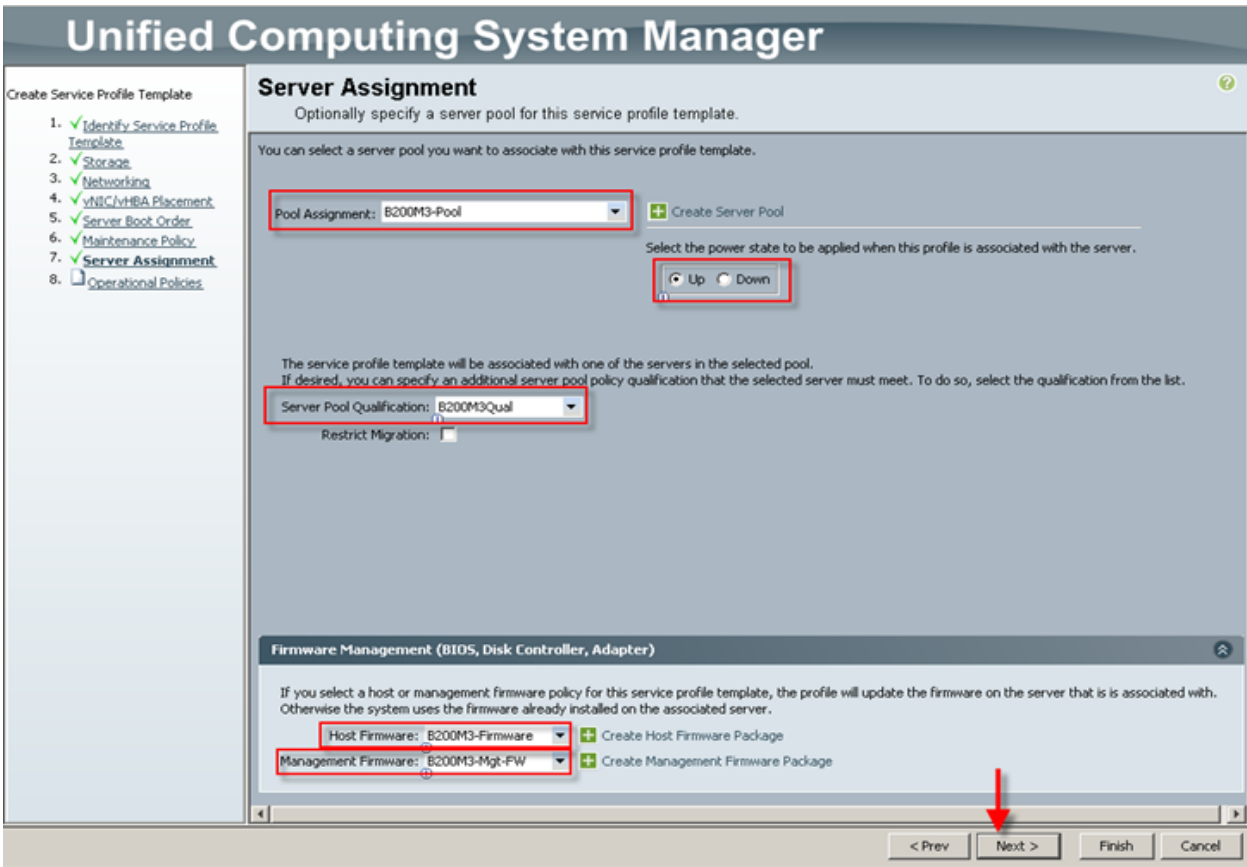
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
CD-ROM	1				
Storage	2				
SAN primary		FC0	Primary		
SAN Target primary			Primary	0	50:06:01:60:46:E0:5E:0A
SAN Target secondary			Secondary	0	50:06:01:69:46:E0:5E:0A
SAN secondary		FC1	Secondary		
SAN Target primary			Primary	0	50:06:01:61:46:E0:5E:0A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

< Prev
Next >
Finish
Cancel

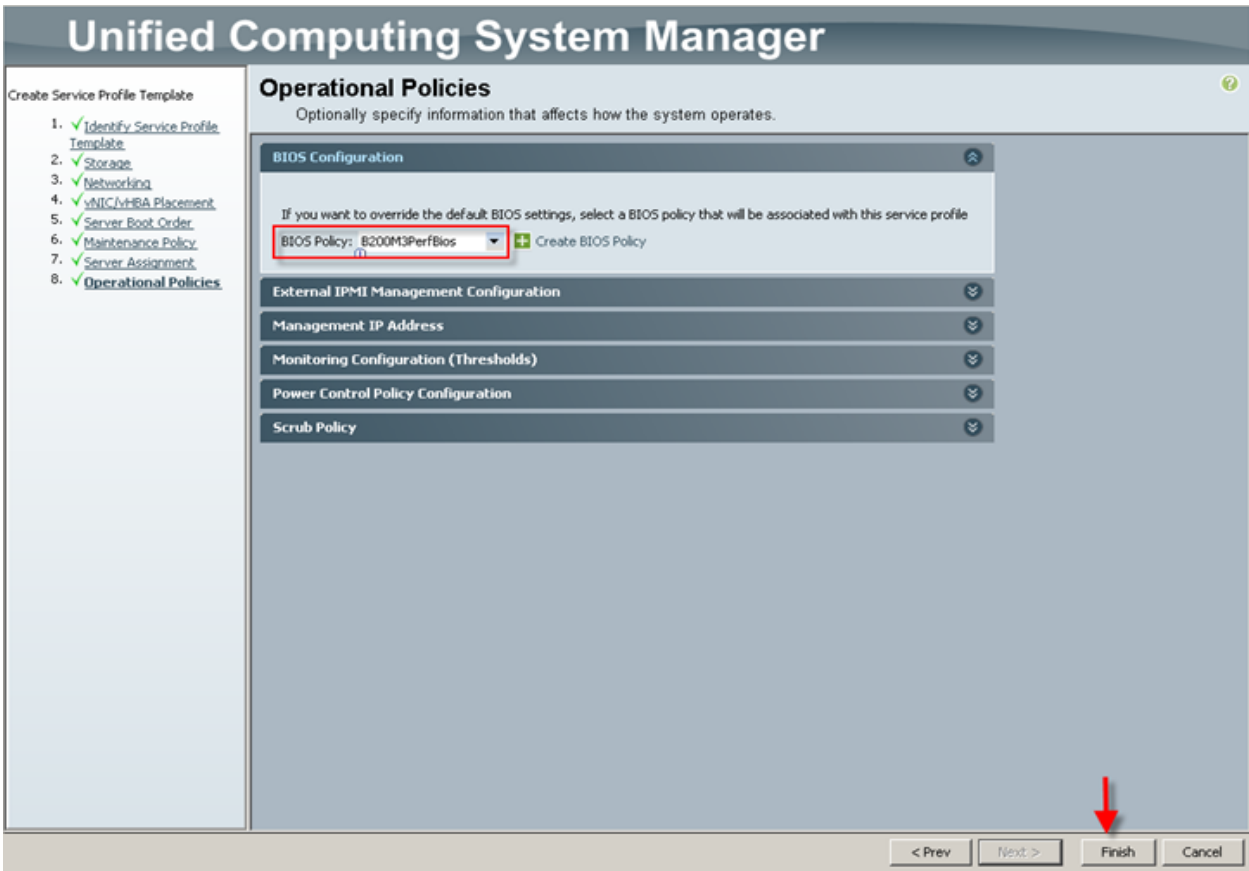
16i We did not create a Maintenance Policy for the project, so we clicked Next to continue:



16j We made the following selections from the drop downs as shown, then clicked Next to continue:

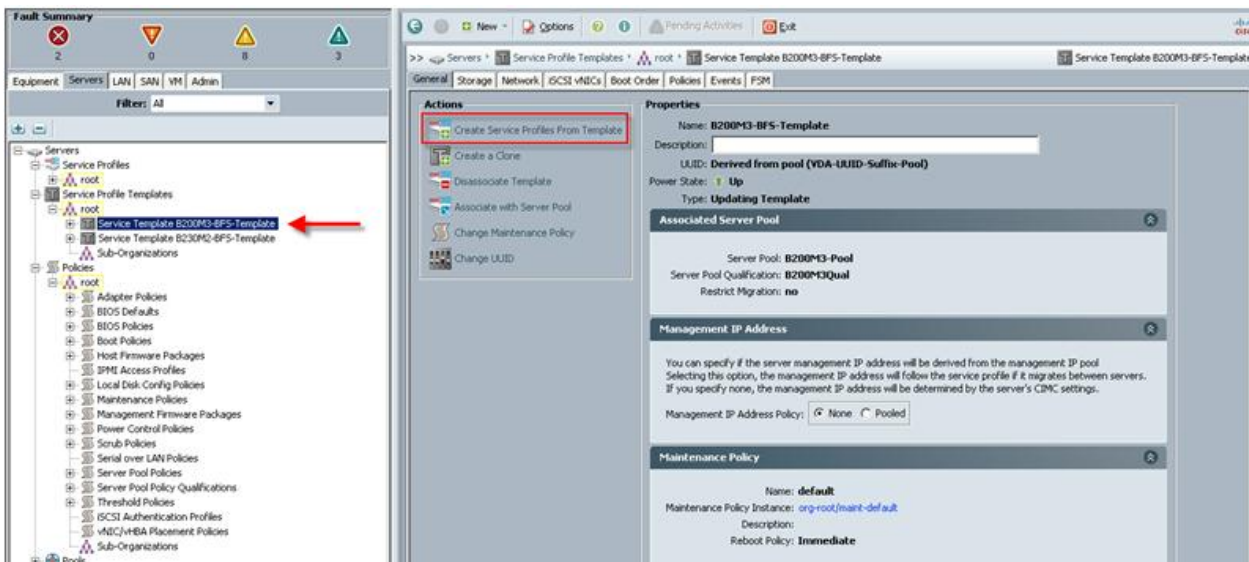


16k On the Operational Policies page, we expanded the BIOS Configuration section and selected the BIOS Policy for the Cisco UCS B200 M3 created earlier, then clicked Finish to complete the Service Profile Template:

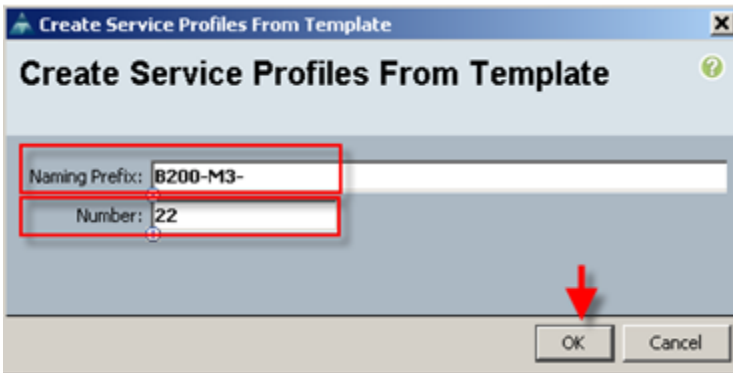


The result is a Service Profile Template for the Cisco UCS Blade Server B200 M3.


- 17 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 Cisco UCS B200 M3, then clicked on Create Service Profiles from Template in the right pane, Actions area:



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



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

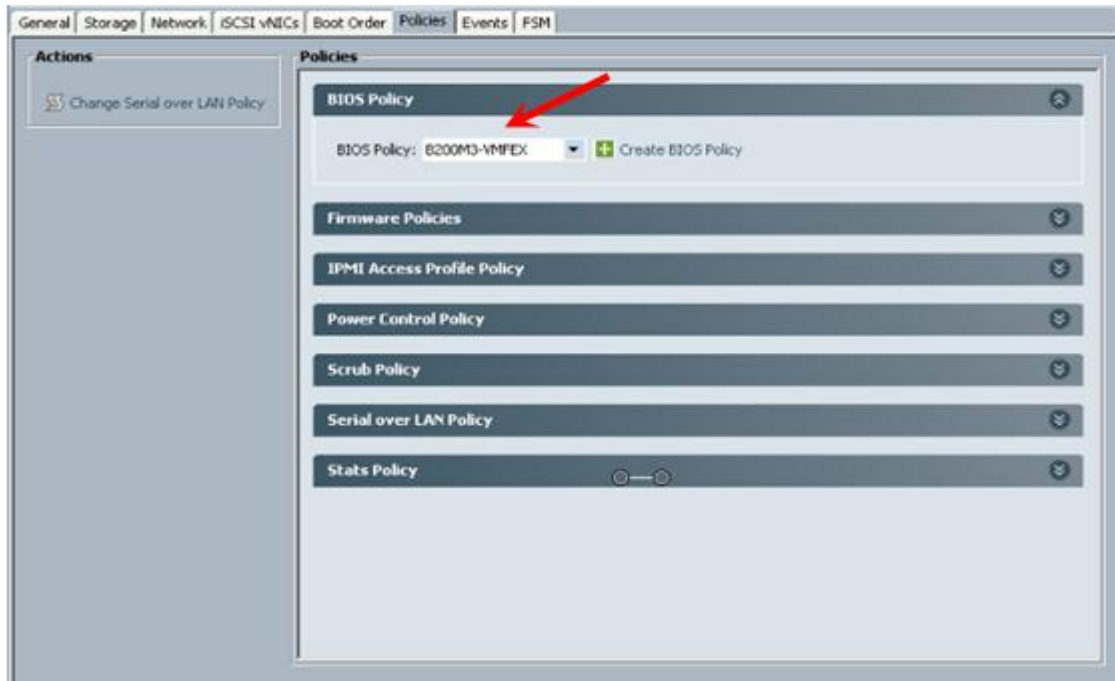


- 20 We verified that each server had a profile and that it received the correct profile.

Name	Overall S...	PID	Model	Serial	Profile	User Label	Co...	Cores E...	Thre...	Me...	Adap...	NICs	HB...	Opera...	Powe...	Assoc State	Fault Sup.	
Server 1 (192.168.1.62)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16207...	org-root/ls-B200-M3-22	192.168.1.62...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 2 (192.168.1.63)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16237...	org-root/ls-B200-M3-19	192.168.1.63...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 3 (192.168.1.64)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16217...	org-root/ls-B200-M3-16	192.168.1.64...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 4 (192.168.1.65)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16207...	org-root/ls-B200-M3-13	192.168.1.65...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 5 (192.168.1.66)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16237...	org-root/ls-B200-M3-2	192.168.1.66...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 6 (192.168.1.67)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16237...	org-root/ls-B200-M3-7	192.168.1.67...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 7 (192.168.1.68)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16237...	org-root/ls-B200-M3-4	192.168.1.68...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A
Server 8 (192.168.1.69)	Ok		UCSB-B200...	Cisco UCS B2...	FCH16207...	org-root/ls-B200-M3-10	192.168.1.69...16	16	32	262144	1	8	2	↑	↑	↑	↑	Associa.../N/A

21. To configure service profiles for the blades to be used for VM-FEX configuration follow steps given above with following changes:

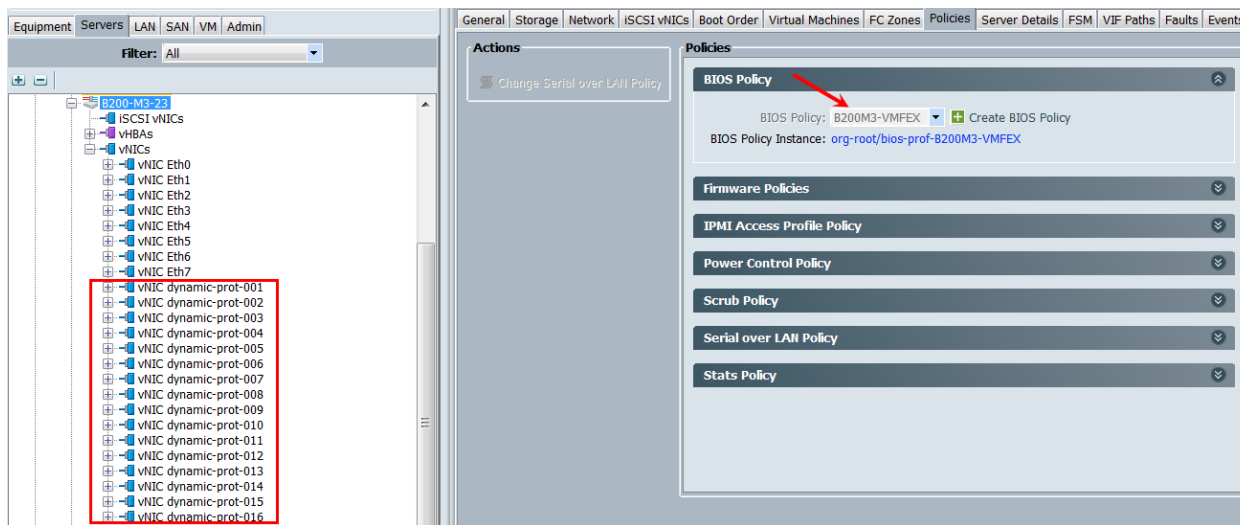
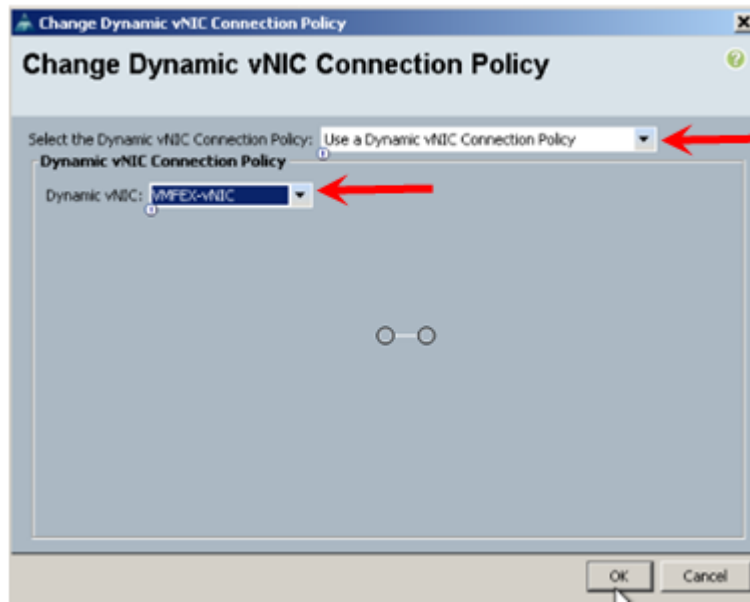
- Select existing service profile or service profile template which need to modified to use with VM-FEX or create new service profile or service profile template. All the steps remain the same except the one showed below.
 - Go to Policies tab and Select BIOS Policy.
 - Select BIOS policy created to use with VM-FEX as shown in Section 6.2.13b



21d. Go to Network tab and click on Change Dynamic vNIC Connection Policy.



21e. Select Dynamic vNIC Connection Policy: Use a Dynamic vNIC Connection Policy. Which was created in section 6.2.21a. Dynamic vNIC: VMFEX-vNIC. Click OK.



As shown in the screenshot above, after applying changes on the service profile or template there will be X number of dynamic vNIC port, where X is selected number of vNIC during creation of dynamic vNIC connection policy. At this point, the UCS Blade Servers are ready for hypervisor installation.

Note: If you are using existing service profile or service profile template and it is associated with blades. Applying these changes will make them all associated servers to Reboot. It is advised to take necessary precautions before applying above mention changes; Ideally all blades should be in maintenance mode.

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.
- Via UCS we can map one CoS value to particular class.
- Apart from FCoE class there can be only one more class can be configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally system will calculate the bandwidth based on following equation (there will be rounding off the number).

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

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

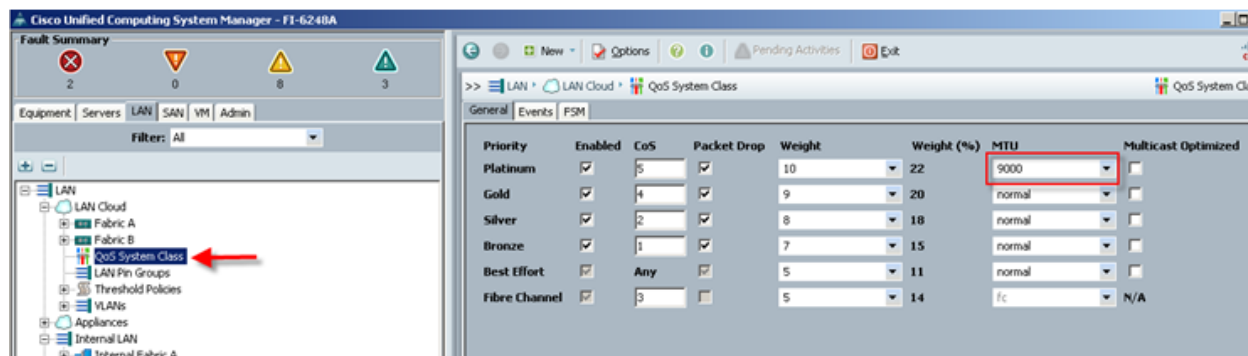
Table 5. Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
Fc	5

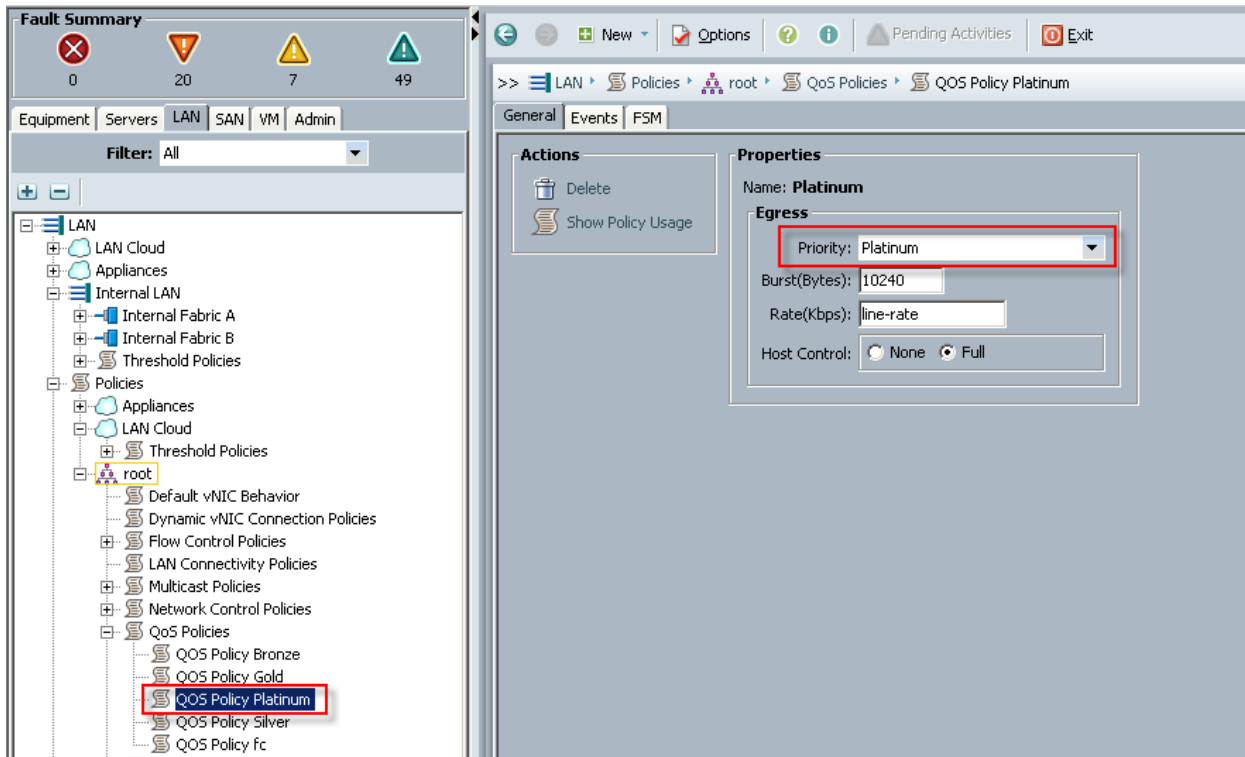
6.2.5 Steps to Enable QOS on the Cisco Unified Computing System

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

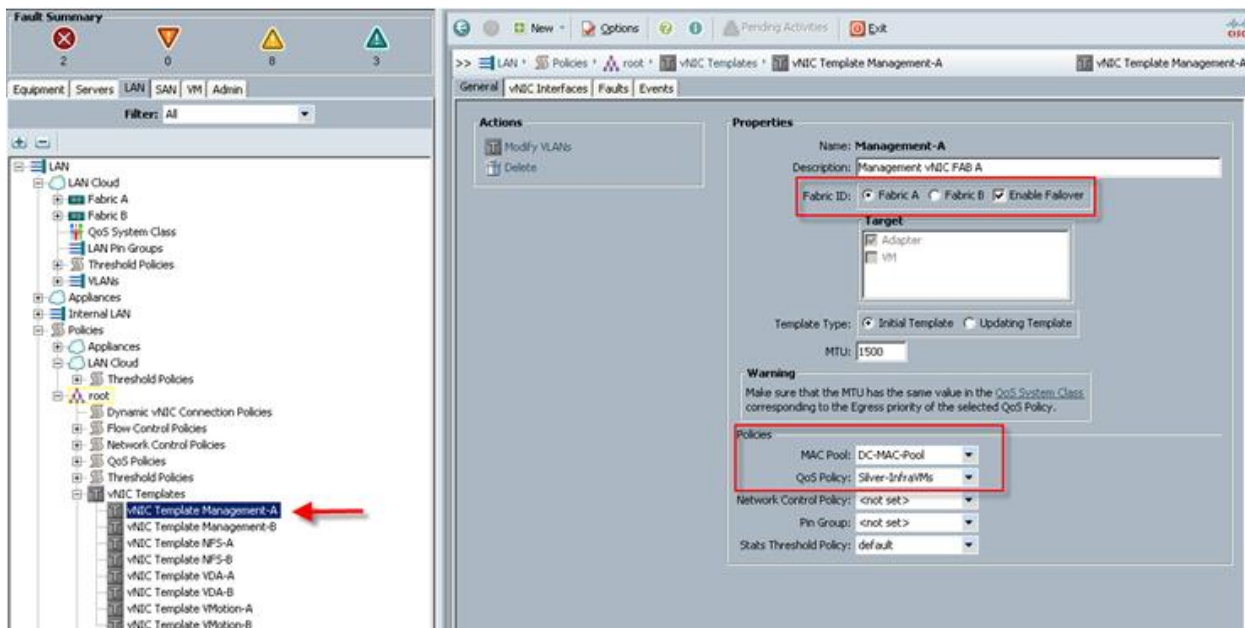
Configure Platinum, Gold, Silver and Bronze policies by checking the enabled box. For the Platinum Policy, used for NFS storage was configured for Jumbo Frames in the MTU column. Notice the option to set no packet drop policy during this configuration.



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



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.



This is a unique value proposition for the Cisco Unified Computing System 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 UCS Connectivity

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

The 6248 Fabric Interconnect is in End host mode, as we are doing both Fiber Channel as well as Ethernet (NAS) data access as per the recommended best practice of the Cisco Unified Computing System. We built this out for scale and have provisioned more than 40 G per Fabric Interconnect (Figure 32).

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

Figure 12. **Ethernet Network Configuration with Upstream Cisco Nexus 5500 Series from the Cisco Unified Computing System 6200 Series Fabric Interconnects**



6.3.2 EMC VNX7500 LAN Connectivity

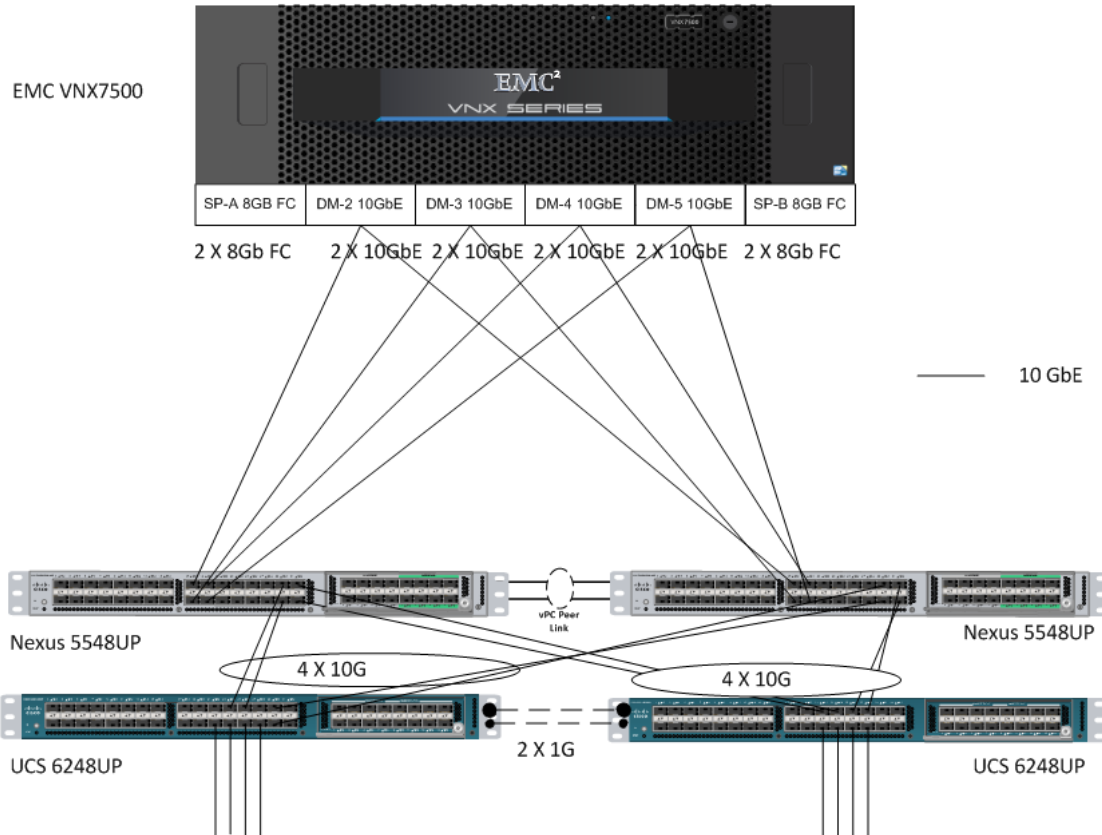
The Cisco Nexus 5548 is used to connect to the EMC VNX7500 storage system for Fiber Channel and file-based access.

The EMC 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 EMC 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.

Figure 13. **EMC VNX Ethernet Connectivity**



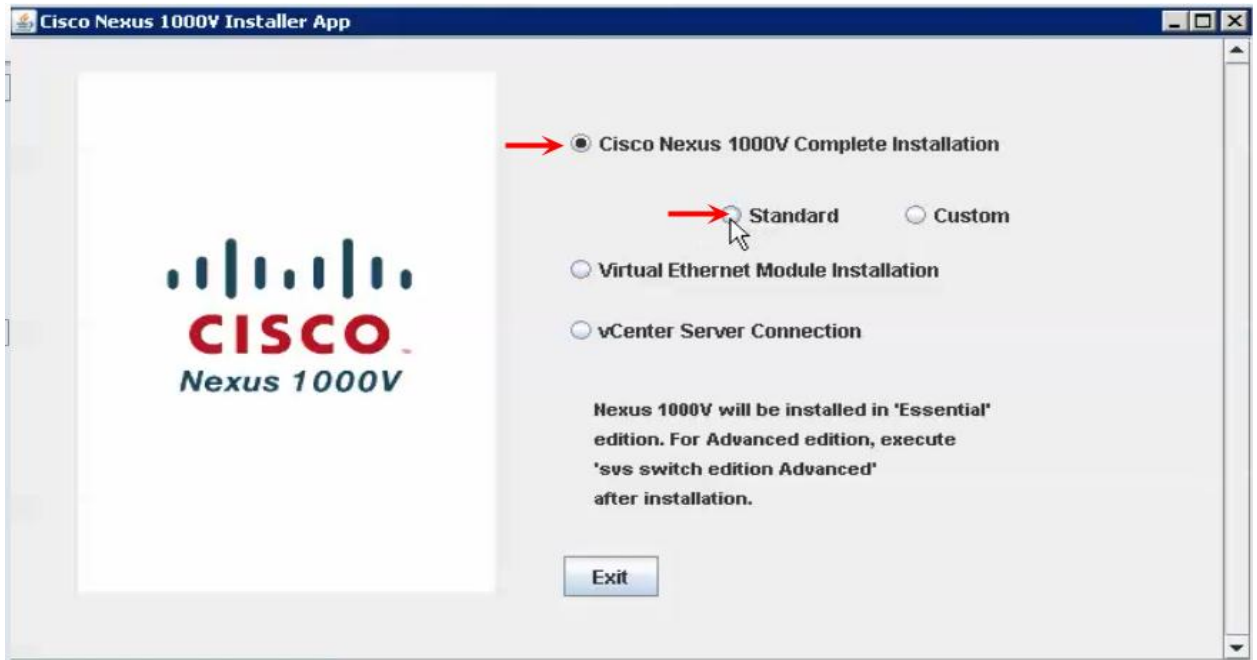
For information about configuring ethernet connectivity on an EMC VNX7500 Storage System, refer to the EMC website.

6.3.3 Nexus 1000V Configuration in L3 Mode

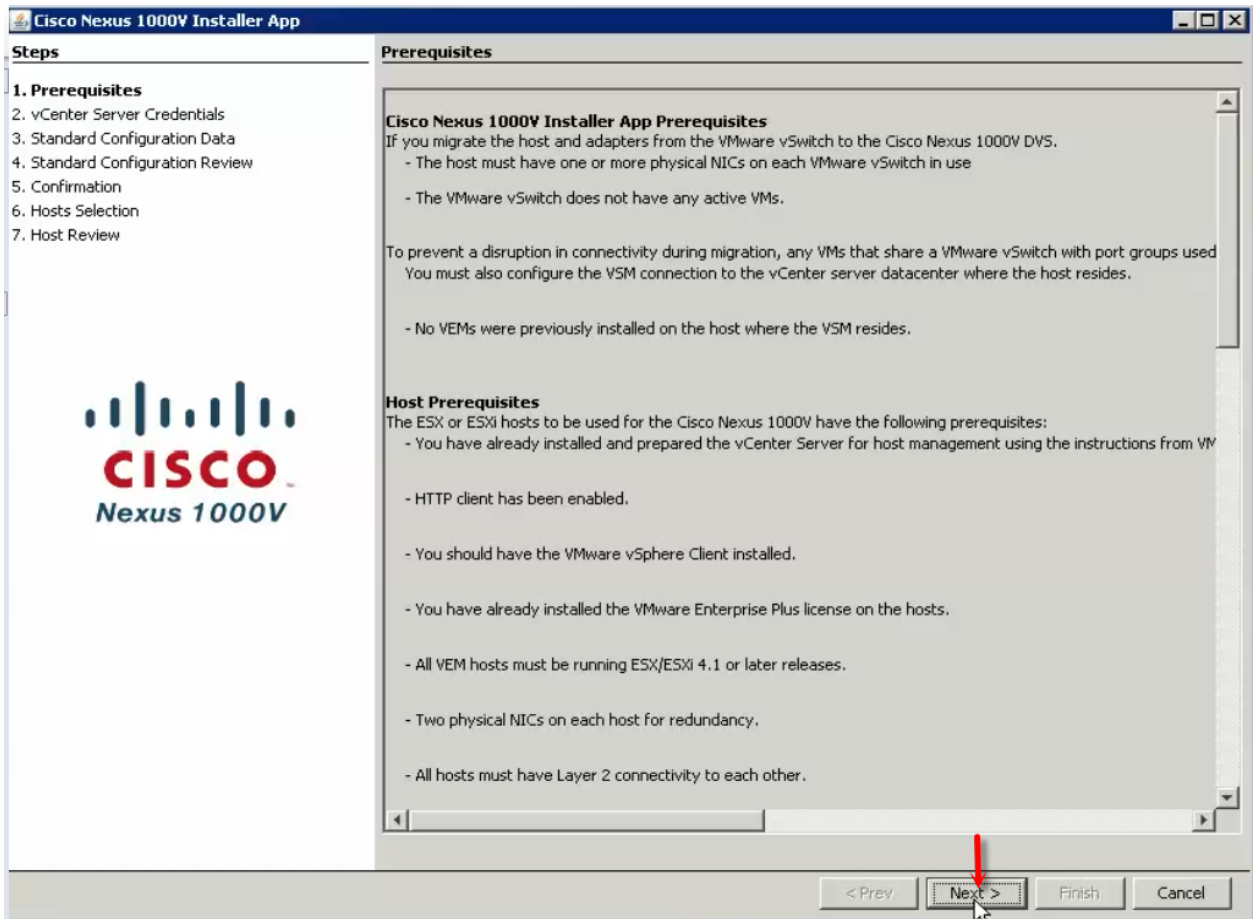
1. To download the Nexus1000 V 4.2(1) SV2 (1.1), Click the link below. <http://software.cisco.com/download/release.html?mdfid=282646785&flowid=3090&softwareid=282088129&release=4.2%281%29SV2%281.1a%29&reind=AVAILABLE&rellifecycle=&reltype=latest>
- 3.
2. Extract the downloaded N1000V .zip file on the Windows host.
- 4.
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).

```
C:\Program Files\Java\jre7\bin>java -jar c:\Nexus1000v.4.2.1.SU2.1.1\Nexus1000v.4.2.1.SU2.1.1\USM\Installer_app\Nexus1000v-install_CNX.jar
```

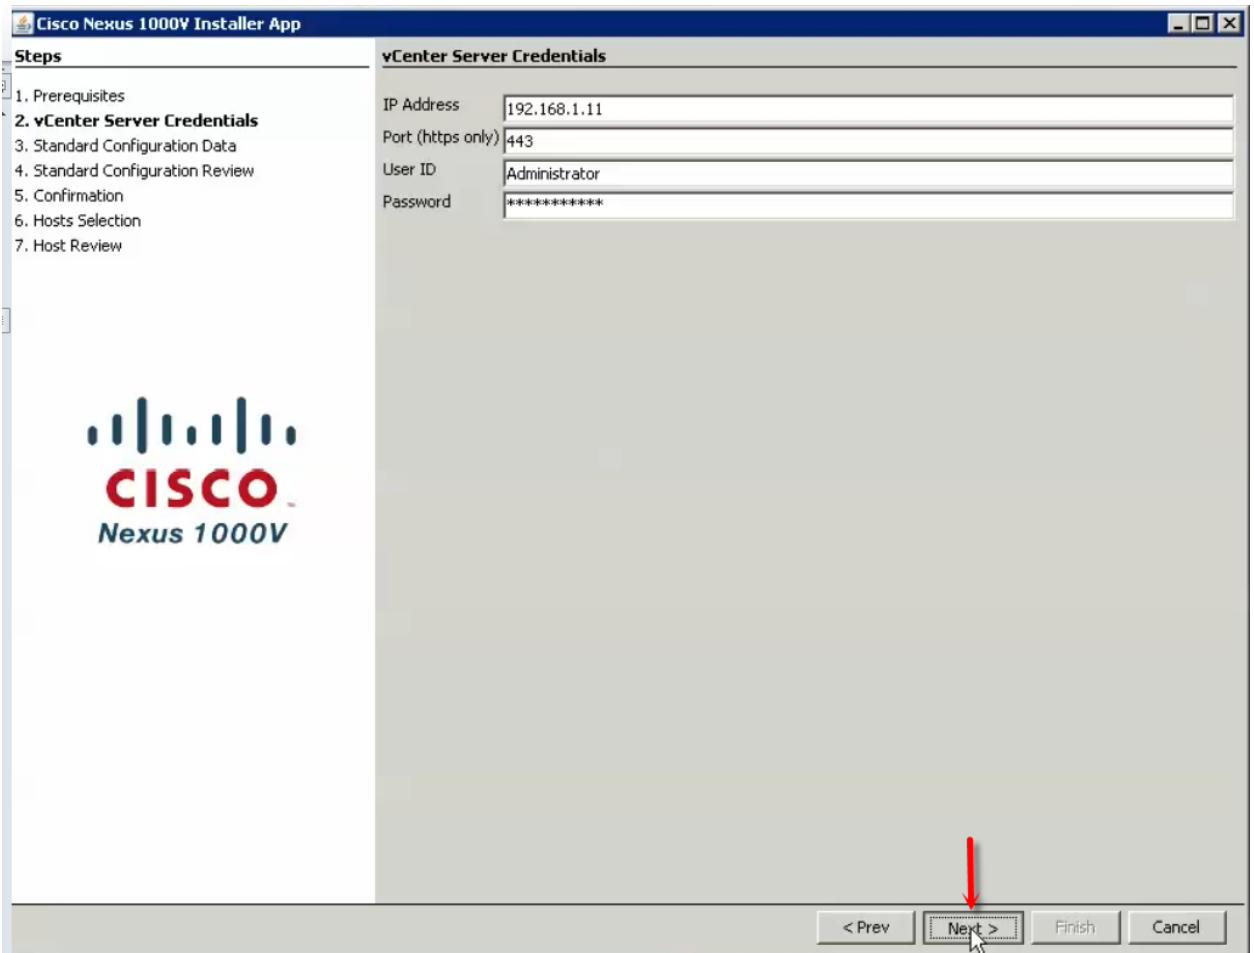
4. After running the installation command, you will see the "Nexus 1000V Installation Management Center".
5. Select the radio button for Cisco Nexus 1000V Complete Installation then select radio button for Standard or Custom installation.
5. For this study we used Standard installation.



- 6.
- 7.
6. Go through the prerequisite for Cisco Nexus 1000V installer, ESXi hosts click Next.



- 8.
- 9.
7. Enter the vCenter IP and the logon credentials.



10.
11.

8. On the standard configuration Data enter or browse intended ESXi host and datastore for primary and secondary instance of VSM. It will install HA pair of VSM.
 9. Enter name for Virtual machine to be created.
 10. Browse to the location of OVA image from the installer package which can be found from extracted file of Nexus 1000V installer:
 12. \Nexus1000v.4.2.1.SV2.1.1\Nexus1000v.4.2.1.SV2.1.1\VSM\Install\nexus1000v.4.2.1.SV2.1.1.ova
 11. Enter IP address, subnet mask, gateway and domain ID for VSM.
 12. Enter Management VLAN; the installer will create vmkernel port-group for ESXi management and uplink port-profile.
 13. Select No for select for Migrate Hosts to DVS.
 14. Click Save configuration for further installation or future reference.
 15. Installer app will verify all the data input and summaries configuration to review.
- 13.

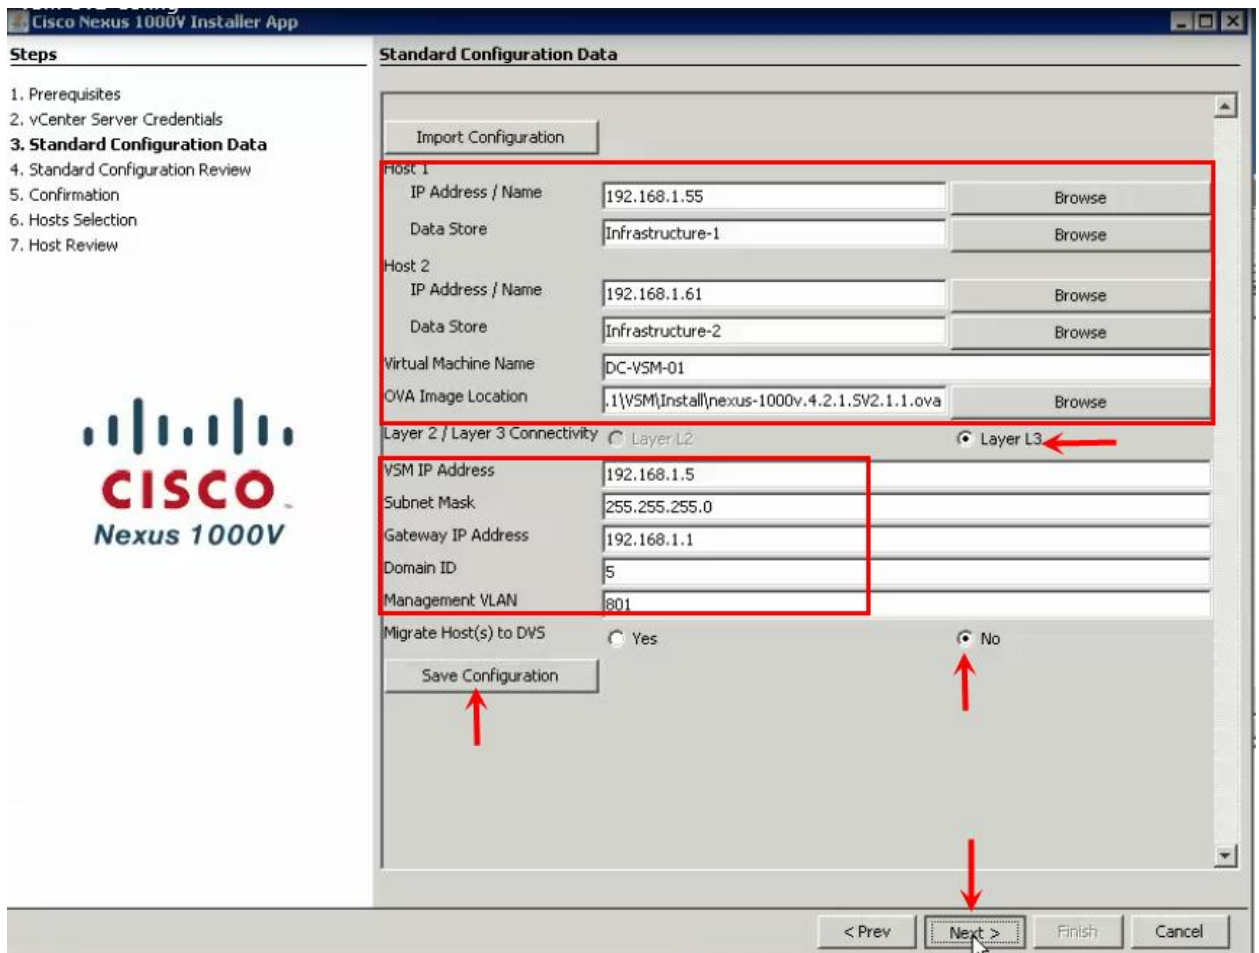
```
port-profile type vethernet n1kv-veth-vlan-801-13
  capability l3control
  vmware port-group
  port-binding static auto
  switchport mode access
  switchport access vlan 801
  no shutdown
  system vlan 801
  max-ports 256
  min-ports 16
  state enabled
```

14.

```
port-profile type ethernet n1kv-eth-2
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 801
  switchport trunk native vlan 1
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 801
  state enabled
```

15.

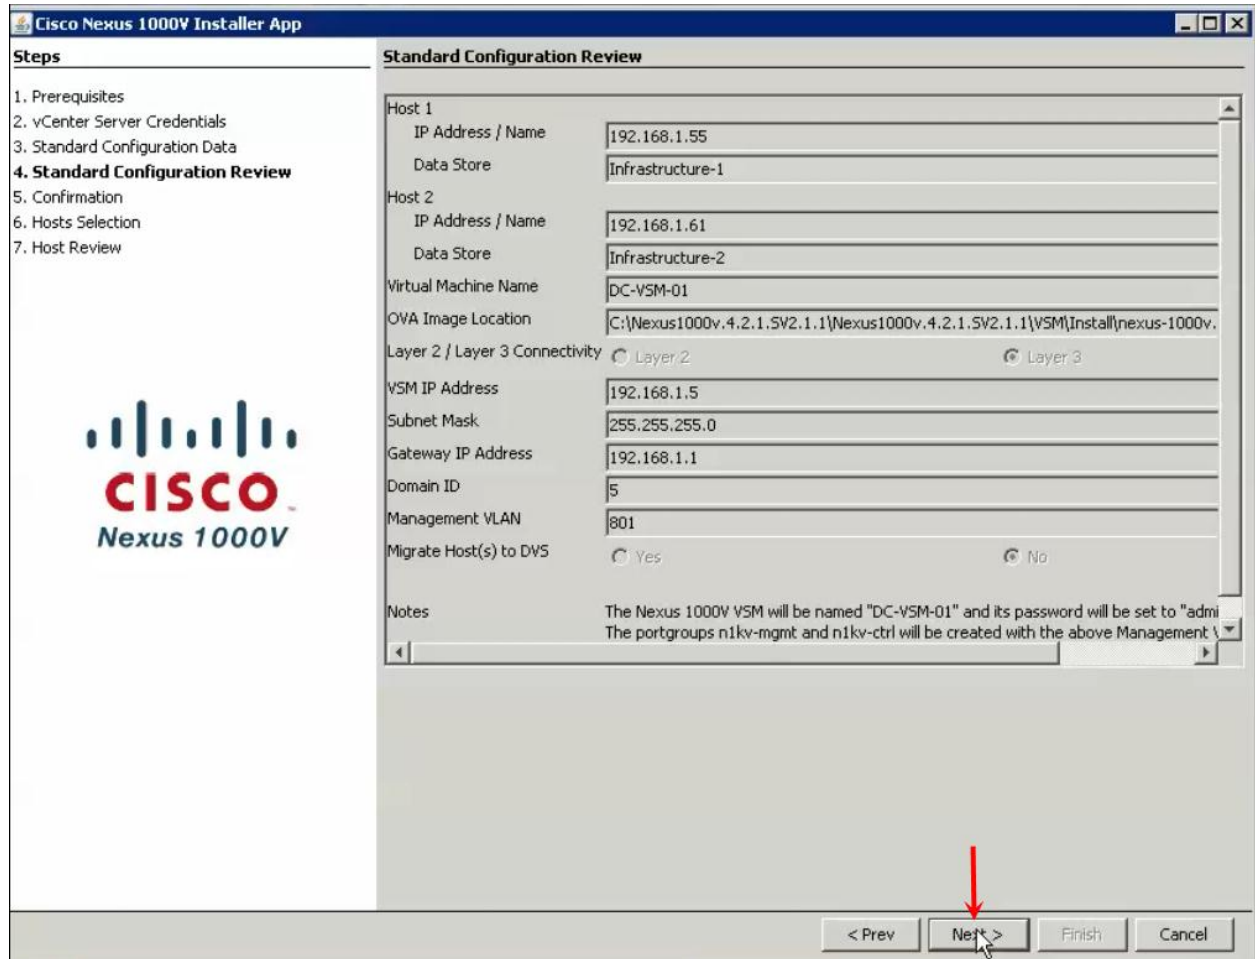
16.



17.

18.

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

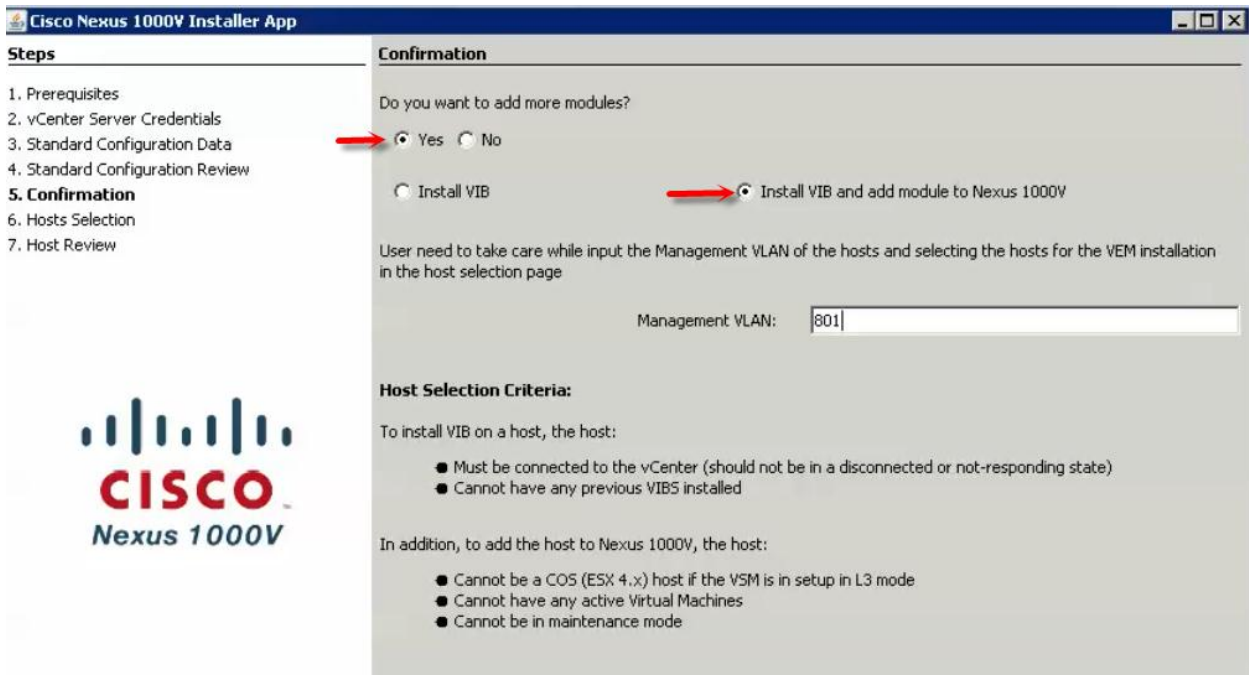


17. Wait for the Completion of Nexus 1000V VSM installation.

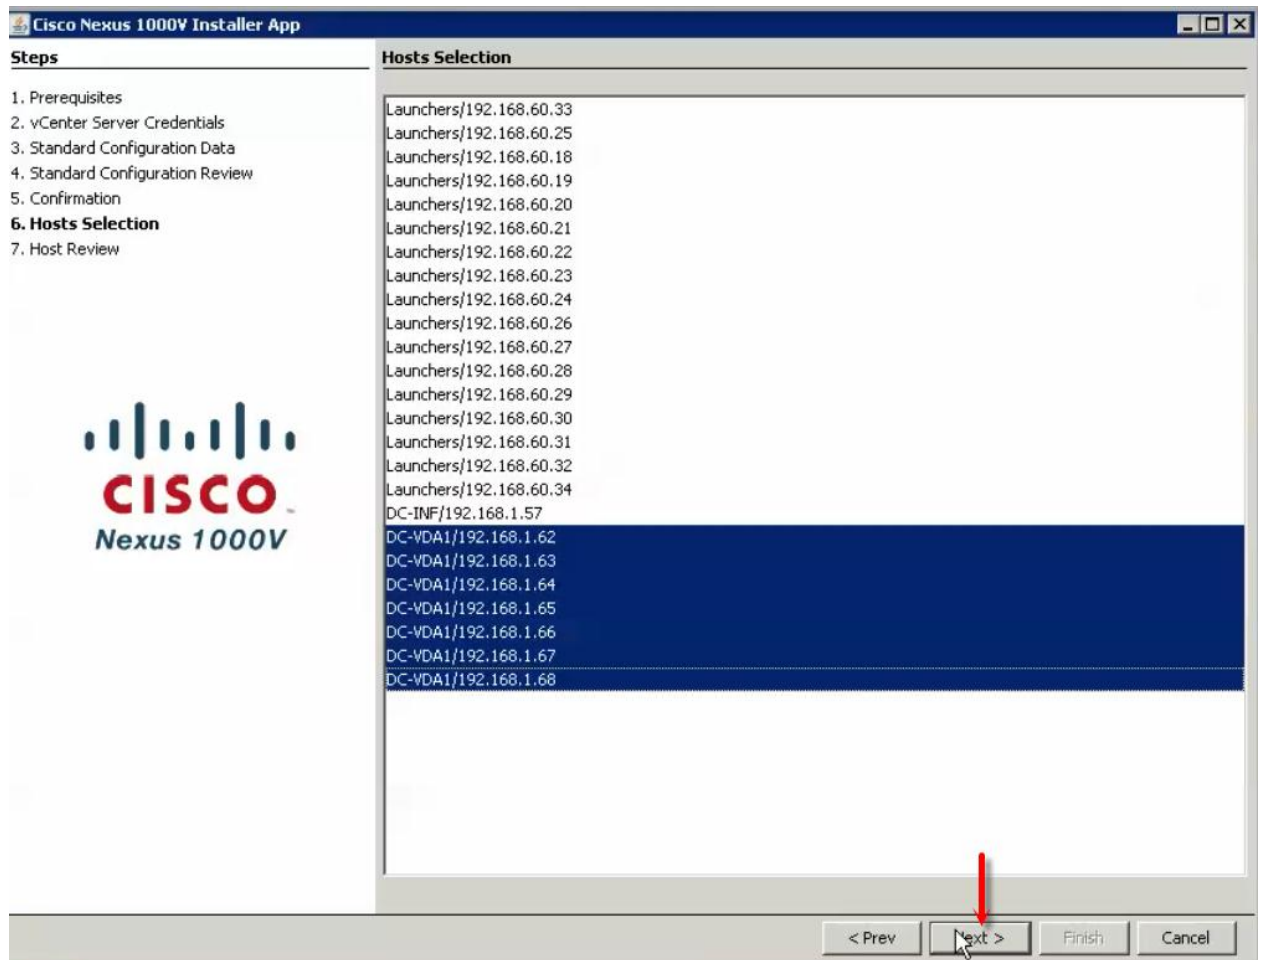


19.
20.

18. Select Yes to add modules. Select Install VIB and add modules to Nexus 1000V.
19. Enter Management VLAN to add ESXi host for VEM installation.
20. Go through the prerequisites to add Host for VIB installation and add it to Nexus 1000V.
21. Click Next.

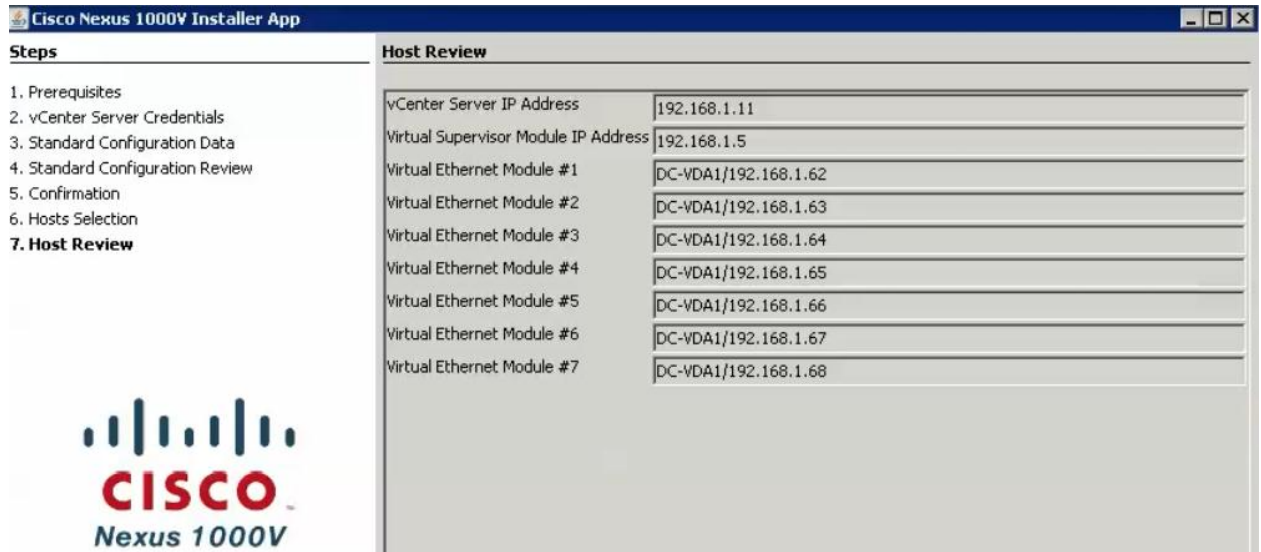


- 22. From the list of hosts, select the ESXi hosts to add to the N1KV DVS.
- 23. Click Next.



- 21.
- 22.


24. Review Host selection to install VIB and add to Nexus 1000V. Click Finish.



23.

24.

25. Review the Summary and click Close.



25.

26.

26. Logon (ssh or telnet) to the N1KV VSM with the IP address and configure VLAN for ESX Mgmt, 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.

```
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 803
VDI-N1KV(config-vlan)# name ML_DC-INF
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)# copy running-config startup-config
```

```
vrf context management
  ip route 0.0.0.0/0 192.168.1.1
vlan 1,800-804
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
```

27. 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# copy running-config startup-config
```

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

28. To migrate and manage 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
```

27. **Note:** These port-profiles are created by default and do not make any changes.

28.

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

29.

```
VDI-N1KV(config)# port-profile type ethernet n1kv-eth-2
VDI-N1KV(config)# vmware port-group
VDI-N1KV(config-port-prof)# switchport mode trunk
VDI-N1KV(config-port-prof)# switchport trunk allowed vlan 801
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
VDI-N1KV(config-port-prof)#state enabled
```

```
port-profile type ethernet n1kv-eth-2
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 801
  switchport trunk native vlan 1
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 801
  state enabled
```

30. 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
```

31. 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
```

32. 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
```

33. 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
```

34. 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 trunk
VDI-N1KV(config-port-prof)# switchport trunk allowed 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
```

35. Create the virtual ethernet port profile for VDA2 traffic.

```
VDI-N1KV(config)# port-profile type vethernet VDA2
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 VDA2
vmware port-group
switchport mode access
switchport access vlan 800
no shutdown
system vlan 800
max-ports 1024
state enabled
```

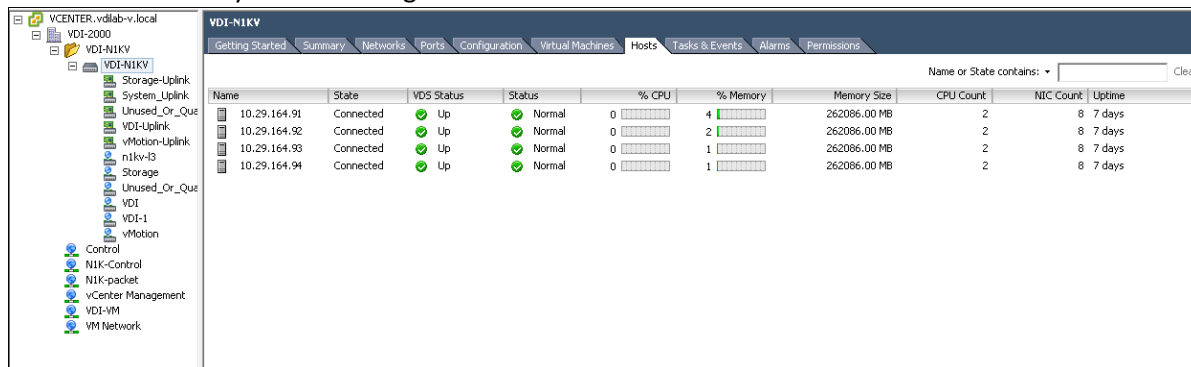
36. Create the virtual ethernet port profile for VDA3 traffic.

```
VDI-N1KV(config)# port-profile type vethernet VDA3
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 VDA3
vmware port-group
switchport mode access
switchport access vlan 800
no shutdown
system vlan 800
max-ports 1024
state enabled
```

37. 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.

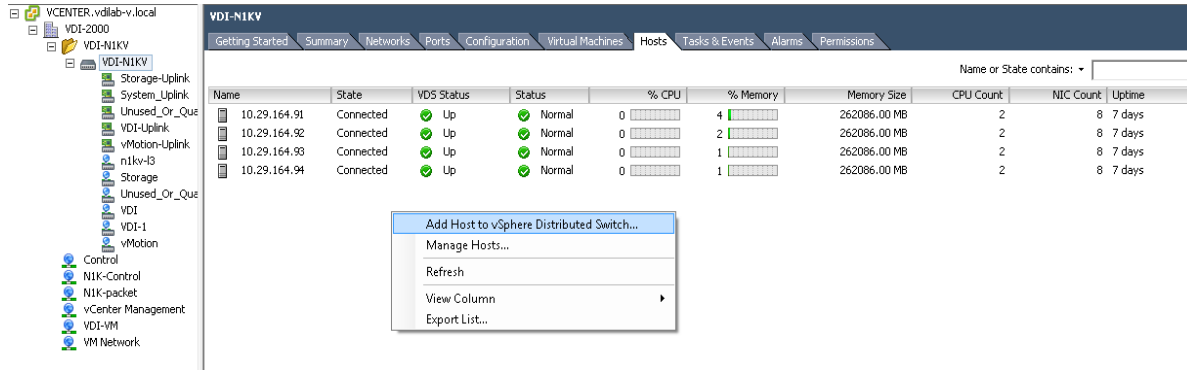
38. Go to Inventory →networking →select DVS for N1KV →click on tab for hosts.



30.

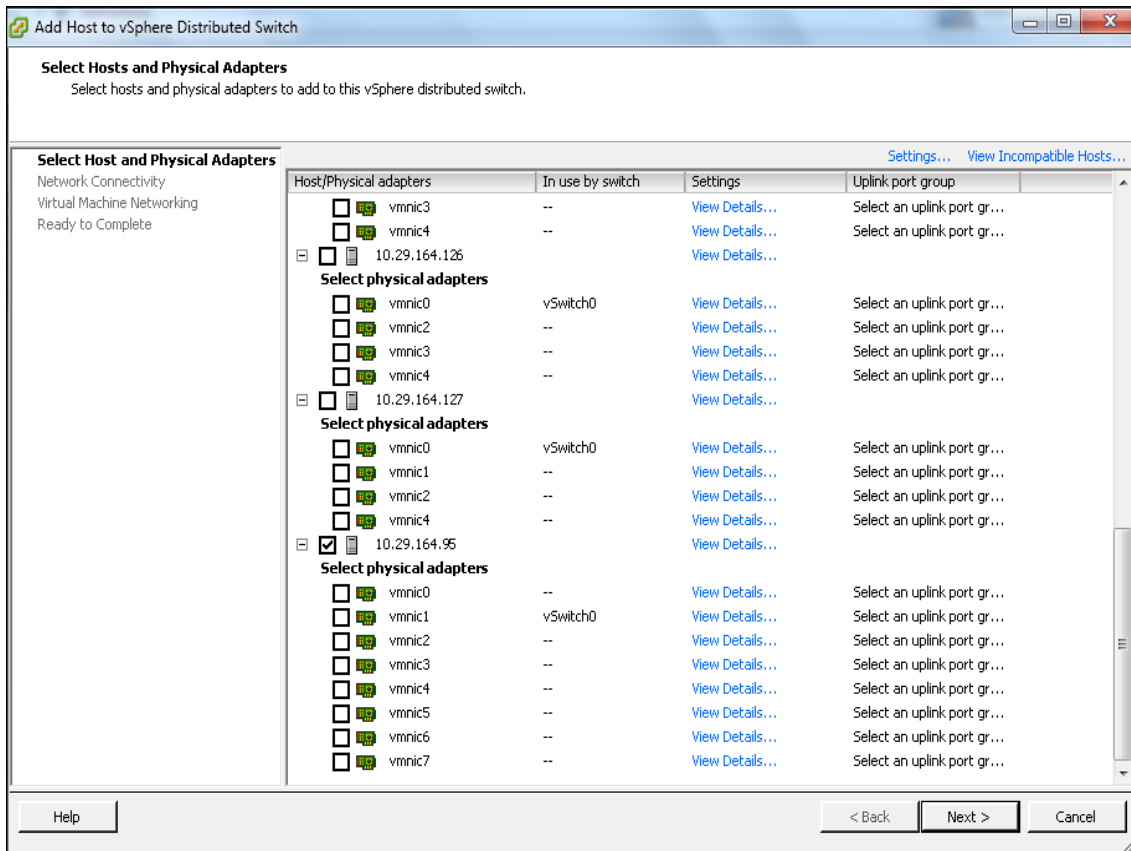
39. Right-click and select add host to vSphere Distributed Switch.

31.



Note: ESXi hosts are not part of the existing configuration.

40. Select ESXi hosts to add in N1KV.



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

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

10.29.164.95 [View Details...](#)

Select physical adapters

<input checked="" type="checkbox"/>		vmnic0	--	View Details...	Select an uplink port group
<input checked="" type="checkbox"/>		vmnic1	vSwitch0	View Details...	Select an uplink port group Storage-Uplink
<input checked="" type="checkbox"/>		vmnic2	--	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic3	--	View Details...	Unused_Or_Quarantine_Upli
<input checked="" type="checkbox"/>		vmnic4	--	View Details...	VDI-Uplink
<input checked="" type="checkbox"/>		vmnic5	--	View Details...	vMotion-Uplink
<input checked="" type="checkbox"/>		vmnic6	--	View Details...	Select an uplink port gr...
<input checked="" type="checkbox"/>		vmnic7	--	View Details...	Select an uplink port gr...

10.29.164.95 [View Details...](#)

Select physical adapters

<input checked="" type="checkbox"/>		vmnic0	--	View Details...	Select an uplink port gr...
<input checked="" type="checkbox"/>		vmnic1	vSwitch0	View Details...	Select an uplink port gr...
<input checked="" type="checkbox"/>		vmnic2	--	View Details...	Select an uplink port group
<input checked="" type="checkbox"/>		vmnic3	--	View Details...	Select an uplink port group Storage-Uplink
<input checked="" type="checkbox"/>		vmnic4	--	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic5	--	View Details...	Unused_Or_Quarantine_Upli
<input checked="" type="checkbox"/>		vmnic6	--	View Details...	VDI-Uplink
<input checked="" type="checkbox"/>		vmnic7	--	View Details...	vMotion-Uplink

10.29.164.95 [View Details...](#)

Select physical adapters

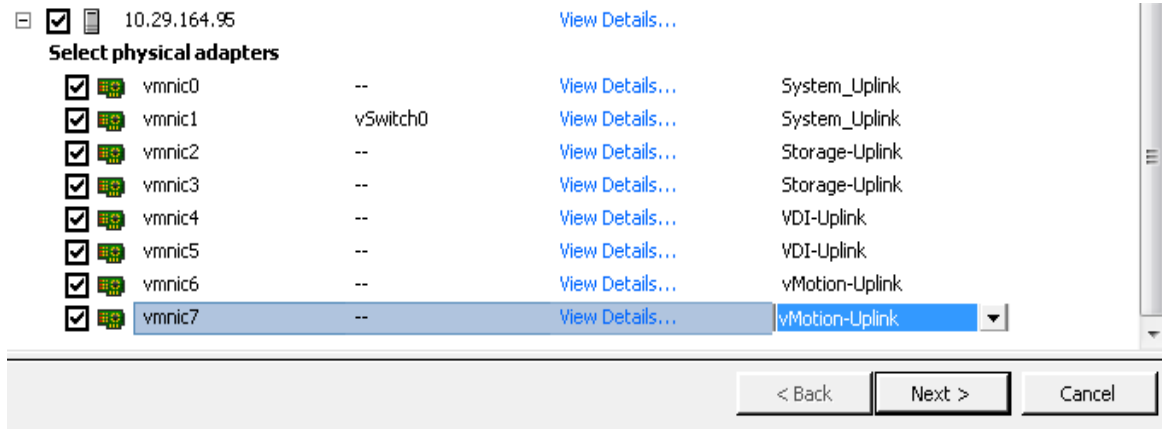
<input checked="" type="checkbox"/>		vmnic0	--	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic1	vSwitch0	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic2	--	View Details...	Storage-Uplink
<input checked="" type="checkbox"/>		vmnic3	--	View Details...	Storage-Uplink
<input checked="" type="checkbox"/>		vmnic4	--	View Details...	Select an uplink port group
<input checked="" type="checkbox"/>		vmnic5	--	View Details...	Select an uplink port group Storage-Uplink
<input checked="" type="checkbox"/>		vmnic6	--	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic7	--	View Details...	Unused_Or_Quarantine_Upli

10.29.164.95 [View Details...](#)

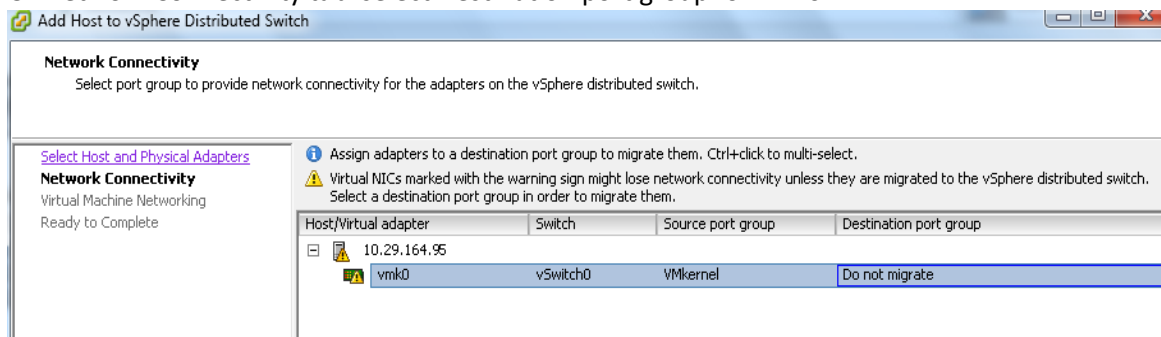
Select physical adapters

<input checked="" type="checkbox"/>		vmnic0	--	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic1	vSwitch0	View Details...	System_Uplink
<input checked="" type="checkbox"/>		vmnic2	--	View Details...	Storage-Uplink
<input checked="" type="checkbox"/>		vmnic3	--	View Details...	Storage-Uplink
<input checked="" type="checkbox"/>		vmnic4	--	View Details...	VDI-Uplink
<input checked="" type="checkbox"/>		vmnic5	--	View Details...	VDI-Uplink
<input checked="" type="checkbox"/>		vmnic6	--	View Details...	Select an uplink port group
<input checked="" type="checkbox"/>		vmnic7	--	View Details...	Select an uplink port group Storage-Uplink

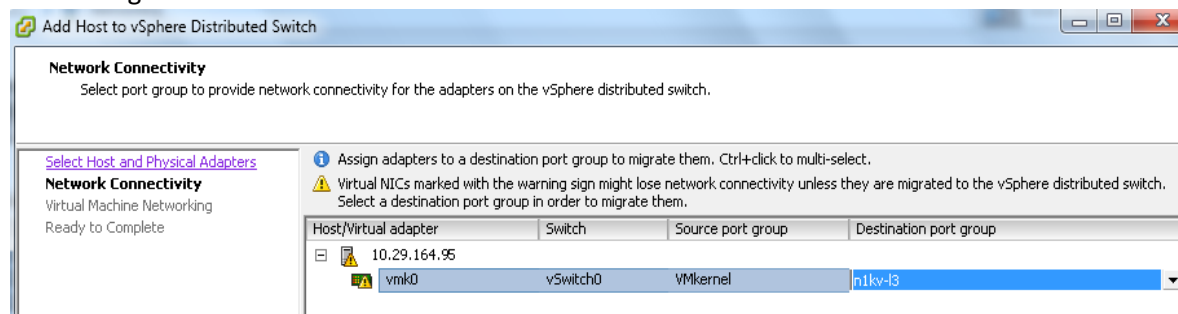
42. After selecting appropriate uplinks click Next.



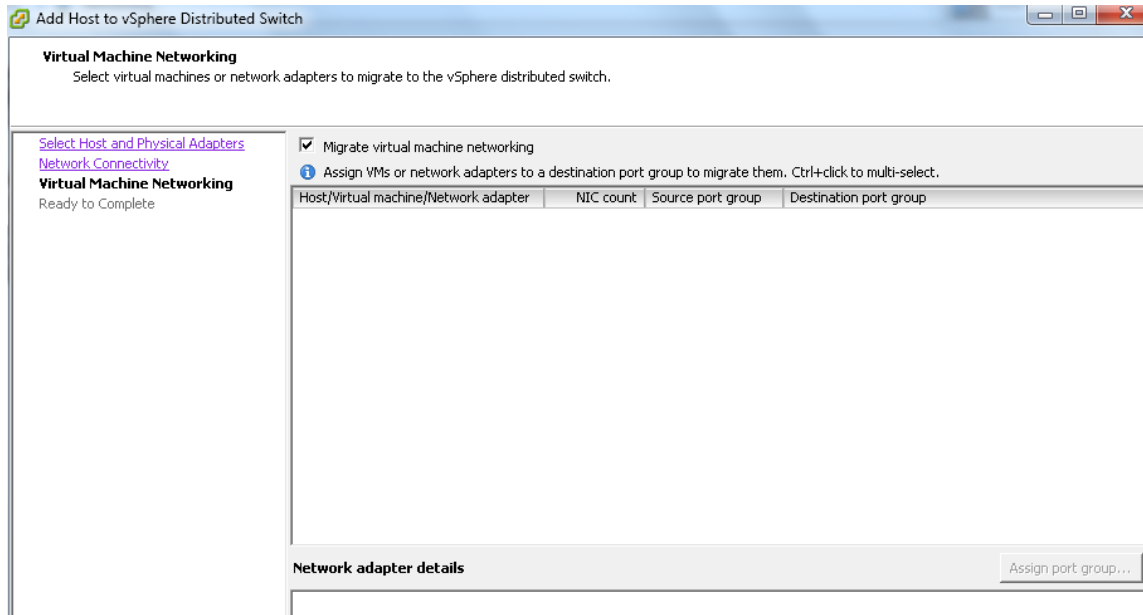
43. Network Connectivity tab select Destination port group for vmk0.



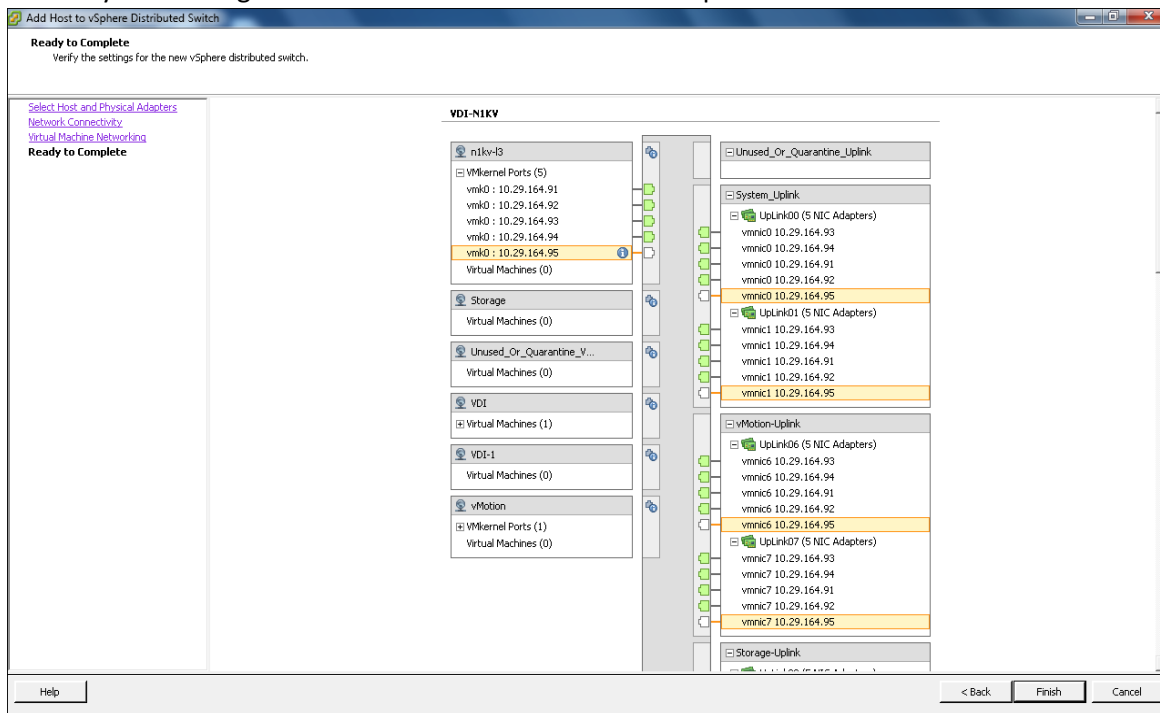
44. From the drop down menu select a port group which was configured for L3 capability and for ESXi host management communication. In this case it is n1kv-l3 and Click Next.



45. On the tab for virtual machine networking, select VMs and assign them to a destination port-group. Click Next.



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



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



Name	State	VDS Status	Status	% CPU	% Memory	Memory Size	CPU Count	NIC Count	Uptime
192.168.1.69	Connected	Up	No...	0	1	262085.80 MB	2	2	8 47 days
192.168.1.70	Connected	Up	No...	0	1	262085.80 MB	2	2	8 47 days
192.168.1.71	Connected	Up	No...	0	1	262085.80 MB	2	2	8 37 days
192.168.1.72	Connected	Up	No...	0	1	262085.80 MB	2	2	8 37 days
192.168.1.73	Connected	Up	No...	0	1	262085.80 MB	2	2	8 47 days
192.168.1.74	Connected	Up	No...	0	0	262085.80 MB	2	2	8 0 second
192.168.1.75	Connected	Up	No...	0	1	262085.80 MB	2	2	8 47 days

48. Run sh module command which will show all the ESXi hosts attached to that VSM.

VDI-N1KV(config)# sh module

```

DC-VSM-02# sh module
Mod  Ports  Module-Type                Model                Status
-----
1    0      Virtual Supervisor Module  Nexus1000V          active *
2    0      Virtual Supervisor Module  Nexus1000V          ha-standby
3    248    Virtual Ethernet Module    NA                   ok
4    248    Virtual Ethernet Module    NA                   ok
5    248    Virtual Ethernet Module    NA                   ok
6    248    Virtual Ethernet Module    NA                   ok
9    248    Virtual Ethernet Module    NA                   ok
10   248    Virtual Ethernet Module    NA                   ok
11   248    Virtual Ethernet Module    NA                   ok

Mod  Sw                Hw
-----
1    4.2(1)SV2(1.1)    0.0
2    4.2(1)SV2(1.1)    0.0
3    4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
4    4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
5    4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
6    4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
9    4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
10   4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)
11   4.2(1)SV2(1.1)    VMware ESXi 5.1.0 Releasebuild-838463 (3.1)

```

```

Mod  Server-IP          Server-UUID          Server-Name
-----
1    192.168.1.6        NA                   NA
2    192.168.1.6        NA                   NA
3    192.168.1.72      0278f983-ffd4-e111-0125-000000000047  192.168.1.72
4    192.168.1.73      0278f983-ffd4-e111-0125-000000000018  192.168.1.73
5    192.168.1.74      0278f983-ffd4-e111-0125-000000000048  192.168.1.74
6    192.168.1.75      0278f983-ffd4-e111-0125-000000000019  192.168.1.75
9    192.168.1.70      0278f983-ffd4-e111-0125-00000000003e  192.168.1.70
10   192.168.1.69      0278f983-ffd4-e111-0125-000000000038  192.168.1.69
11   192.168.1.71      0278f983-ffd4-e111-0125-000000000017  192.168.1.71

```

49. Repeat the procedure to configure additional VSM pairs the second XenDesktop VDA ESX cluster.

6.3.4 Configuring Cisco UCS VM-FEX

1. Click the download link given below to install latest VEM software installer.
<http://software.cisco.com/download/release.html?mdfid=283853163&flowid=25821&softwareid=283853158&release=2.0%285b%29&relind=AVAILABLE&rellifecycle=&reltype=latest>
2. From the retrieved ISO file unzip it and browse to `\ucs-bxxx-drivers.2.1.1a\VMware\VM-FEX\Cisco\MLOM\ESXi_5.1\ cisco-vem-v151-5.1-1.1.1.1.vib`

Note: Select accordingly your Network Interface card and ESXi hypervisor version. Current Supported NICs are 1280, M81KR, MLOM and ESX/ESXi 4.1 U2, ESX/ESXi4.1 U3, ESXi 5.0 U1, ESXi 5.1.

6.3.4.1 Installing Cisco VEM Software Bundle Using SSH or ESXi Shell cmdline:

1. Upload VEM installer vib file to datastore on ESXi host. Preferably on the shared storage
2. Login into ESXi host using ssh or ESXi shell. Run the command shown below.

`esxcli software vib install -v /vmfs/volumes/xxxxx/cisco-vem-v151-5.1-1.1.1.1.vib`

3. To further verify run following command:

`Vmkload_mod -l |grep pts`

```

~ # esxcli software vib install -v /vmfs/volumes/XA_WC1_FS1/cisco-vem-v151-5.1-1.1.1.1.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: Cisco_bootbank_cisco-vem-v151-esx_5.1-1.1.1.1
  VIBs Removed: Cisco_bootbank_cisco-vem-v150-esx_4.2.1.2.1.1.0-3.1.1
  VIBs Skipped:
~ #
~ #
~ # vmkload_mod -l |grep pts
vem-v151-vmfex-pts      0      148
~ #
  
```


You can also install using the VMware Update Manager.

Please refer to link given below for different method of installation or upgrading VEM software on ESX/ESXi host.

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/GUI_VMware_VM-FEX_UCSM_Configuration_Guide_chapter3.html

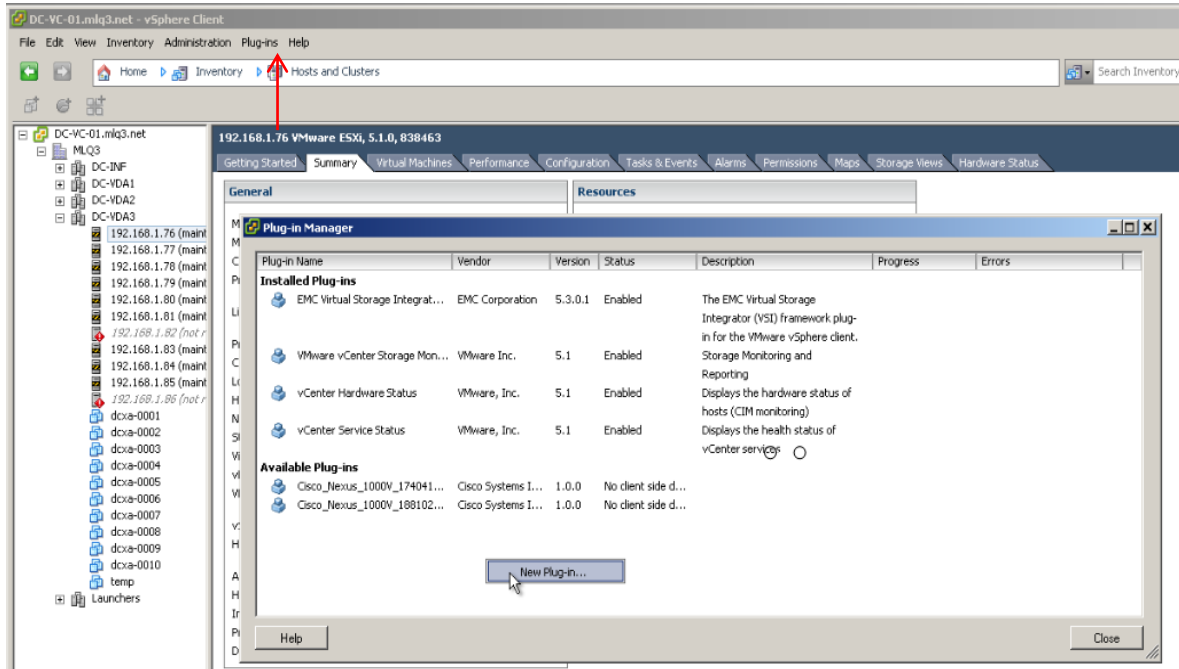
6.3.4.2 Configuring VM-FEX on Cisco UCS Manager and Integration with vCenter

1. Add Cisco UCSM extension for VM-FEX on vCenter.
2. Export vCenter Extension.
3. On the Cisco UCS Manager go to VM tab. Click VMware and click Export vCenter Extension. Save the extension file and click OK.

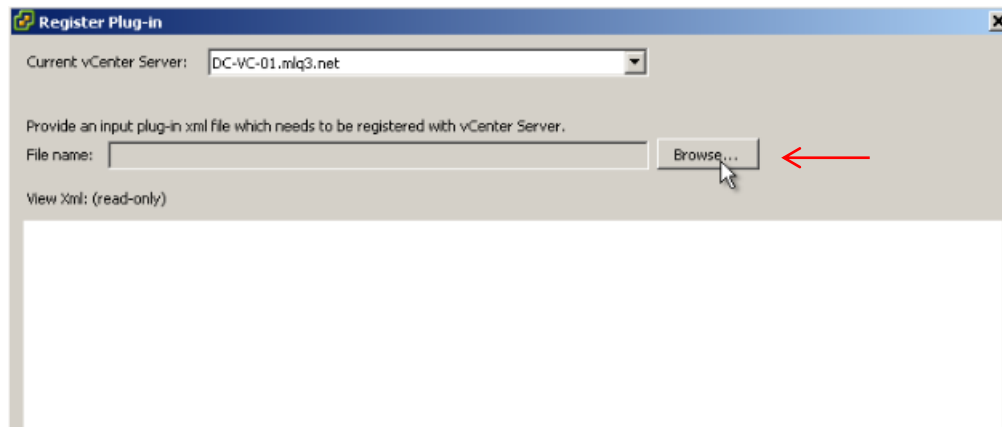


4. Login to vCenter Client. Select Plug-ins and select Plug-in Manager.

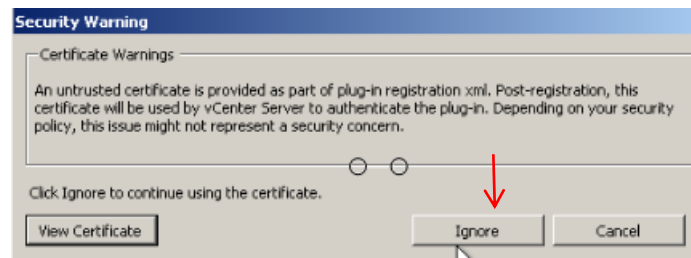
- Right-click in the empty pane for Plug-in Manager and select New Plug-in.



- Click Browse and select saved vCenter Extension file.



- Click Ignore.



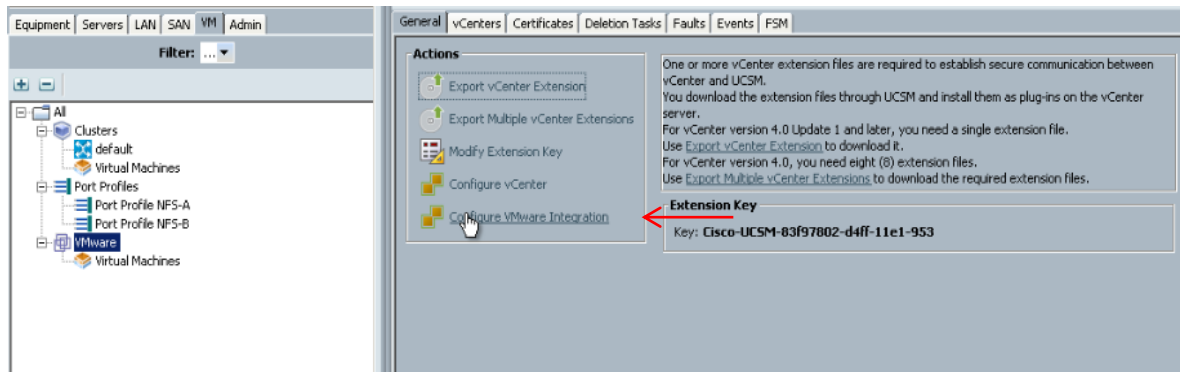
- Verify that Plug-in registered and shows in Available Plug-ins.



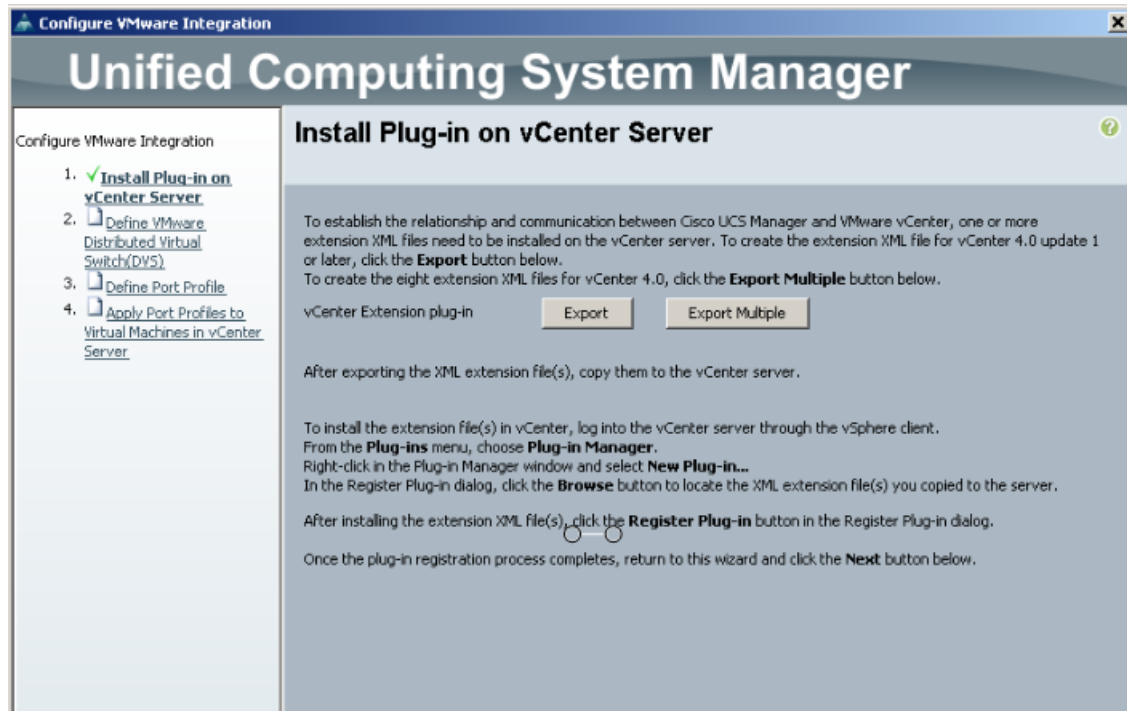
Available Plug-ins

	Cisco_Nexus_1000V_174041...	Cisco Systems I...	1.0.0	No client side d...
	Cisco_Nexus_1000V_188102...	Cisco Systems I...	1.0.0	No client side d...
	Cisco-UCSM-83f97802-d4ff-1...	Cisco Systems, I...	1.0.0	No client side d...

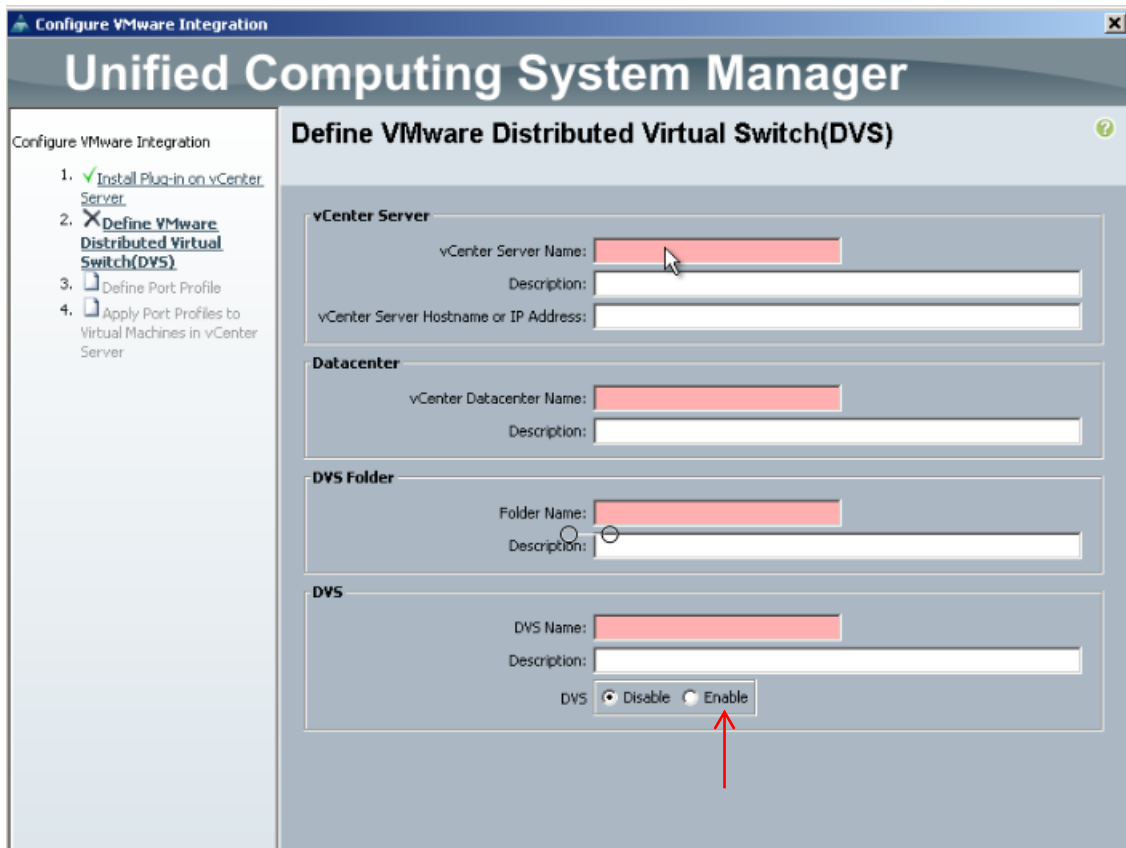
9. Configure VMware Integration.
10. From the Cisco UCS Manger, go to VM tab, click VMware and on the right side pane click Configure VMware Intergation.
11. Login to vCenter Client, go to networking and create a Folder to use with VM-FEX distributed virtual switch.



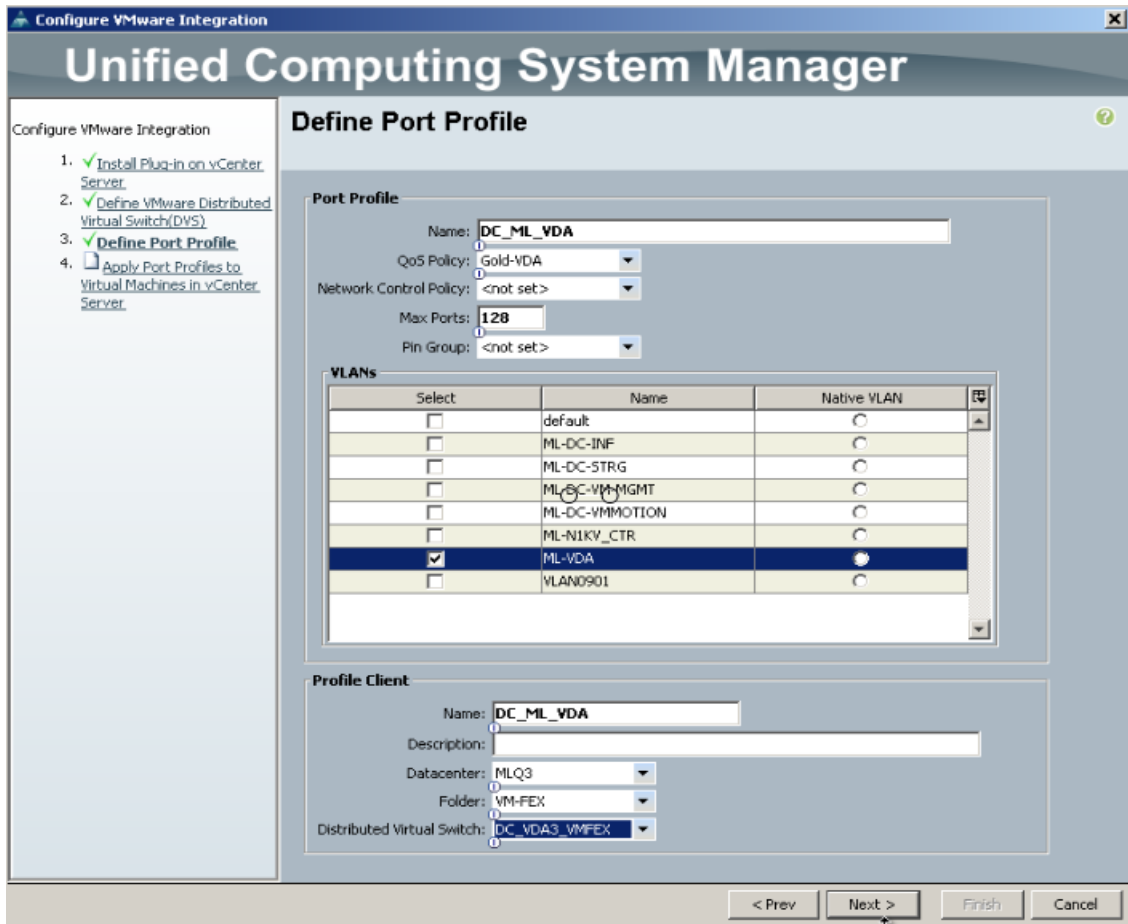
12. Install Plug-in on vCenter Server; which was completed in steps above section x.1.
13. Click Next.



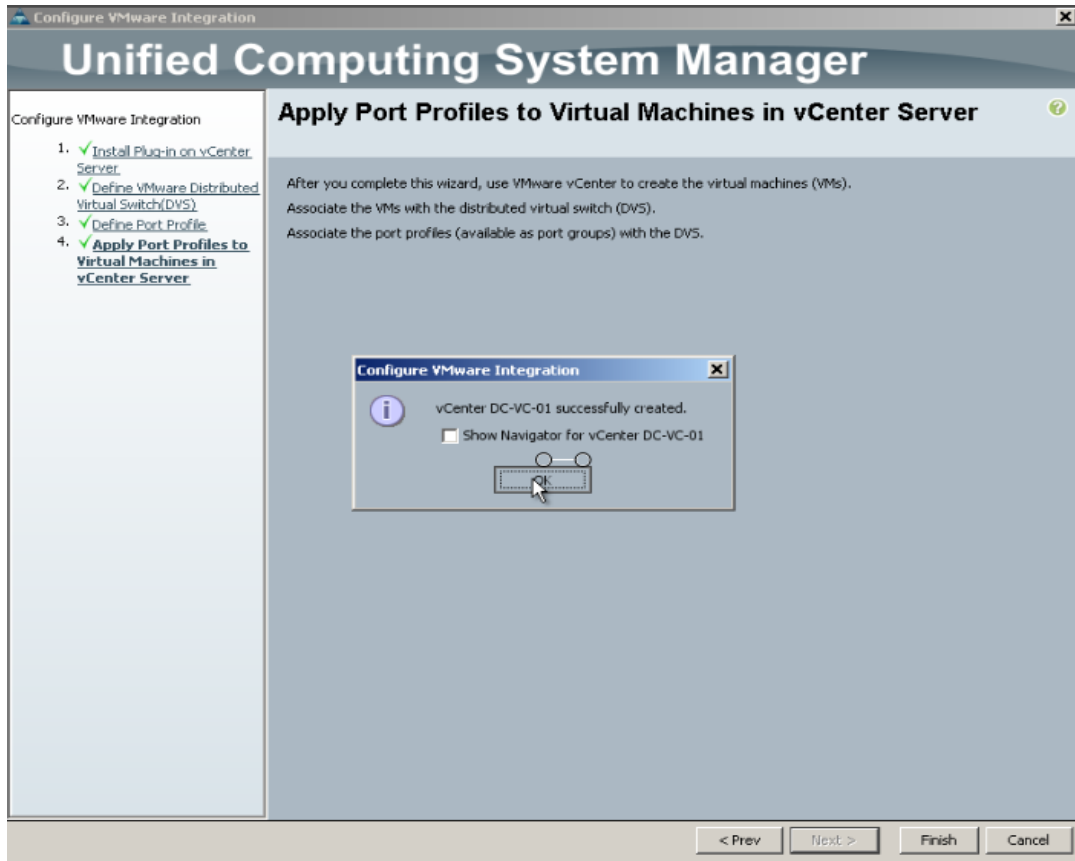
14. Fill out the section for Define VMware Distributed Virtual Switch.
15. Click Enable in the DVS section. Click Next.



16. Define Port Profile.
17. Select a name for port-profile, QoS Policy, Max Ports, VLANs.
18. Select Datacenter, Folder, Distributed Virtual Switch.
19. Click Next.

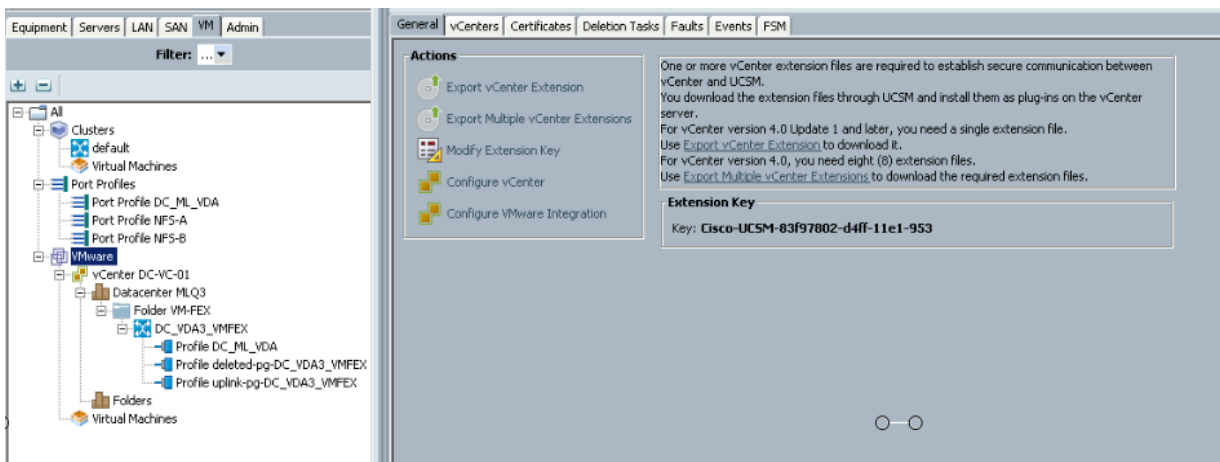



20. Apply Port Profile to Virtual Machines in vCenter Server.
21. Click Finish.
22. Click Ok.



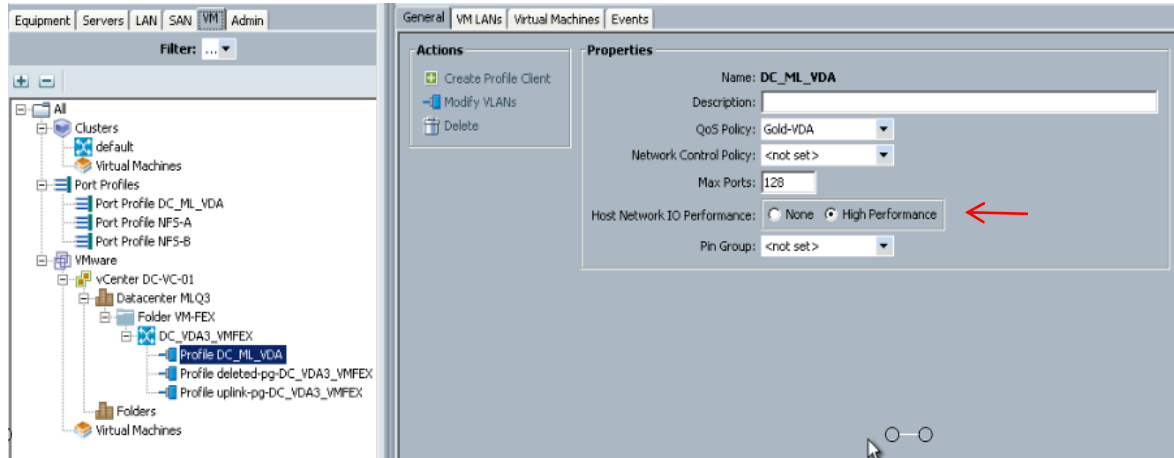
23. Above completed configuration is shown in the screenshot below.

24. Under VMware node vCenter → Datacenter → Folder → VM-FEX switch → port-profile.





25. Select port-profile created and set the Host Network IO performance to High Performance.

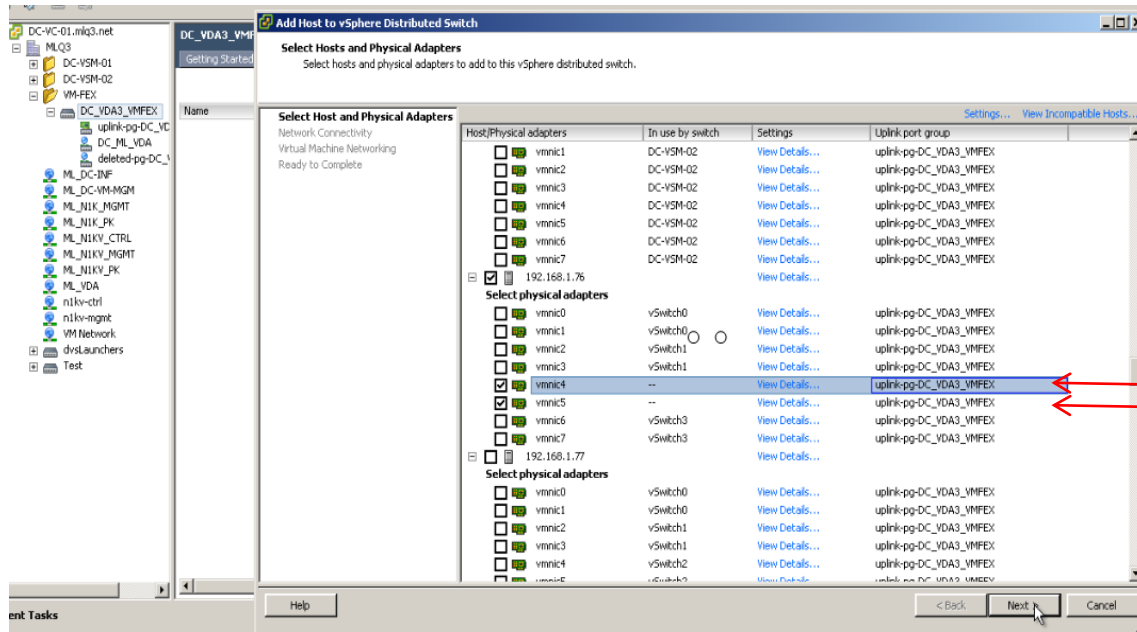


26. On vCenter console, go to networking tab. Select Distributed virtual switch for VM-FEX.

27. Click Hosts tab.

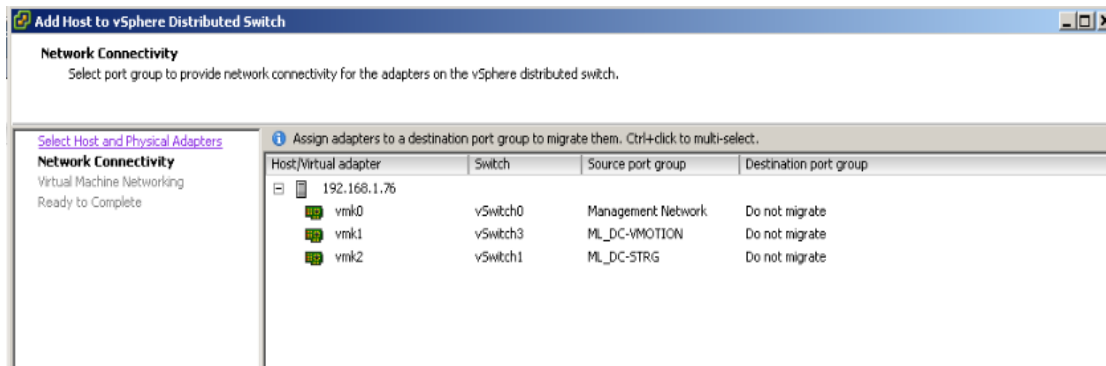
28. Right-click in the empty field and select Add host to vSphere Distributed Switch.

29. Check the box next to the ESXi host and vmnic needed to be add on VM-FEX vDS.



In our study we configured vmnic 4 and vmnic 5 for VDA to add on the VM-FEX configuration as shown above.

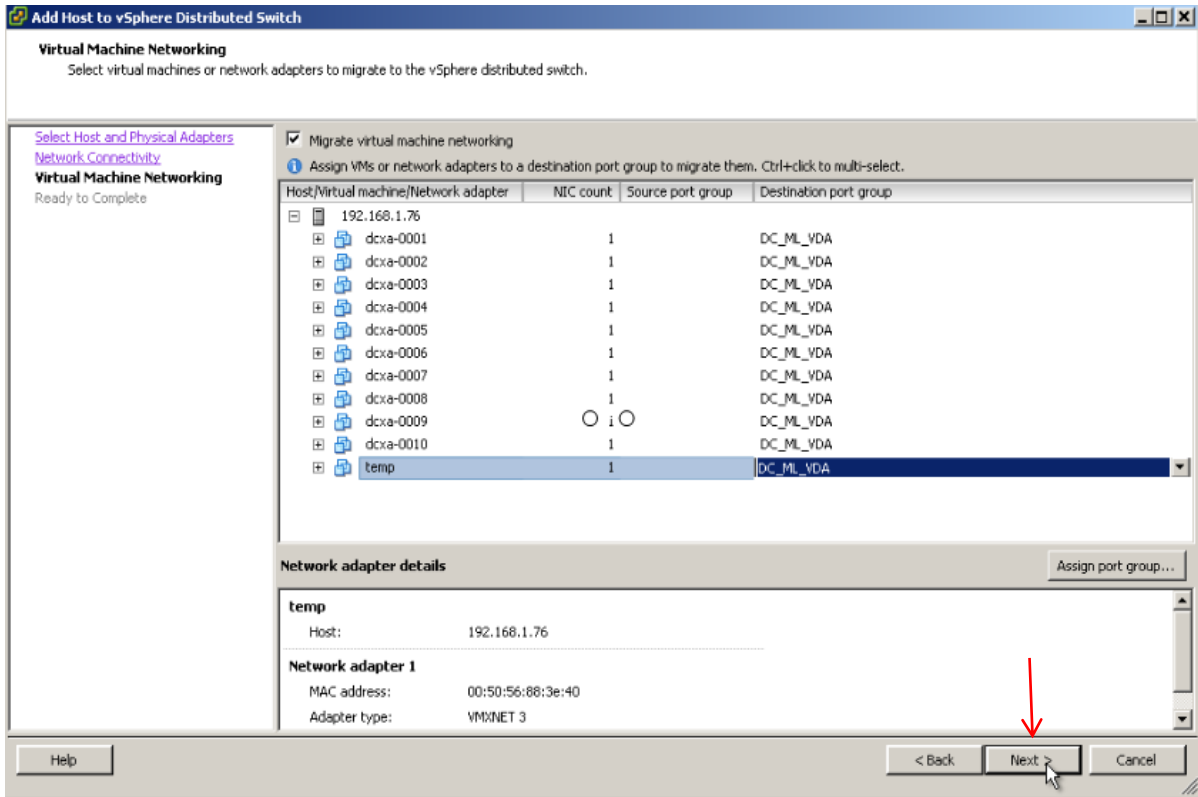
30. Click Next.



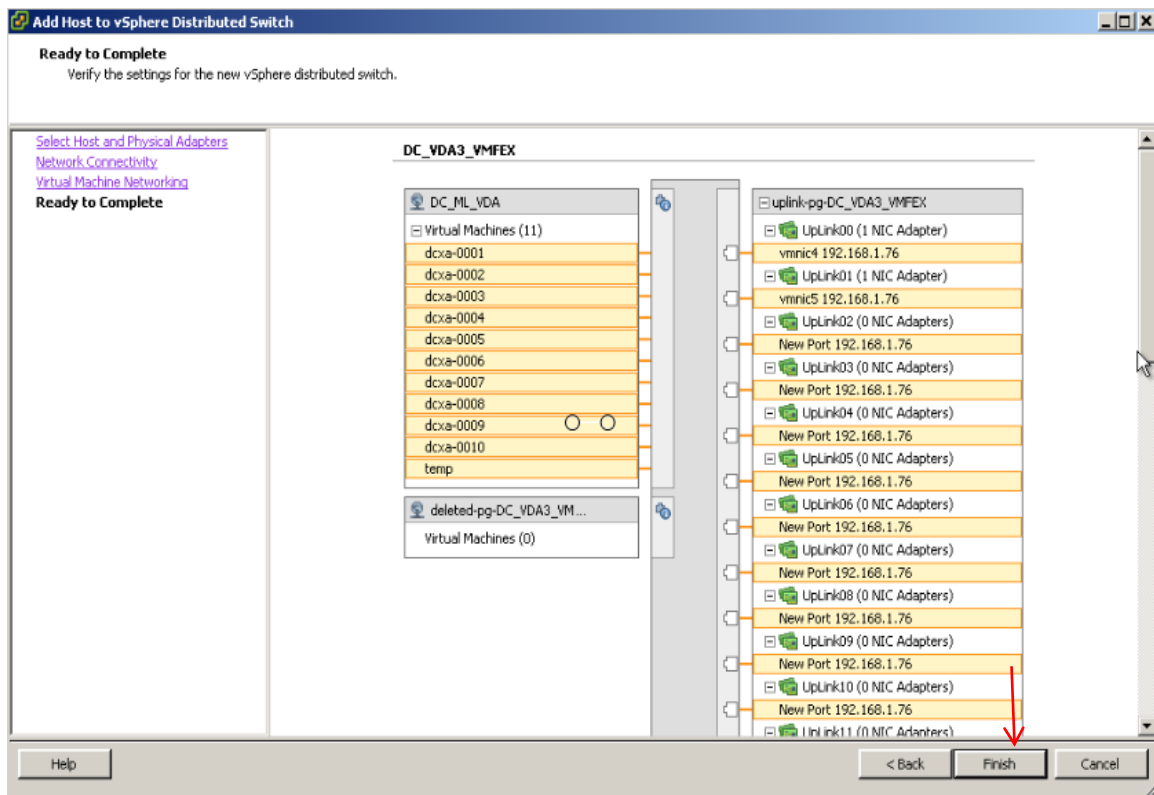
31. On the tab for the Virtual machine networking, click the check box.

32. Migrate the destination port-group for all the desired VMs to port-group created for the VM networking.

33. Click Next.

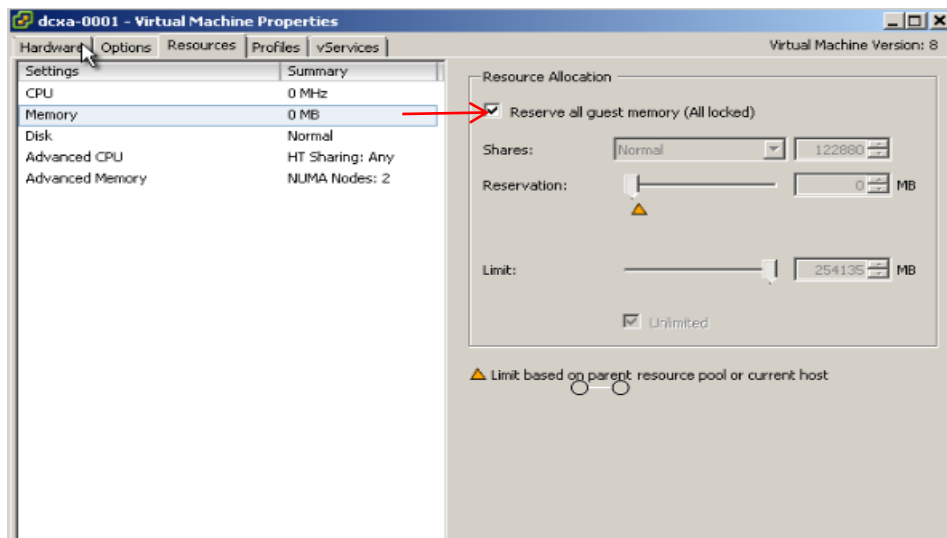


34. Click Finish.

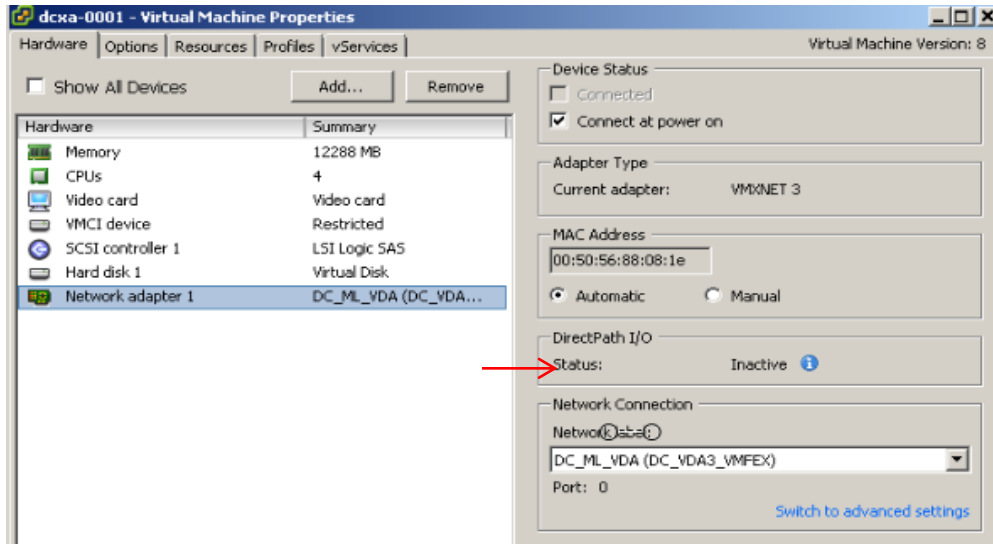





35. Select the virtual machines that are part of the VM-FEX configuration.
36. Right-click and select edit settings.
37. Click the Resources tab and select Memory.
38. Check the box to reserve all guest memory.
39. Select the ESXi host that is part of the VM-FEX vDS configuration.
40. Select the appropriate uplink port-profile as per the vmnic assignment and associated VLAN.



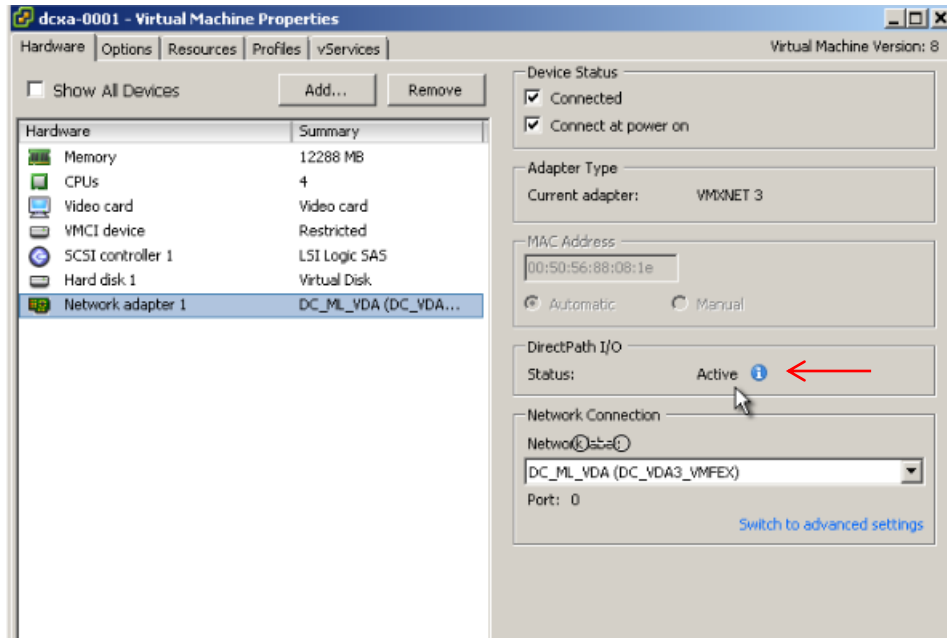
41. Verify the label for the network adapter is connected to the desired port-group on VM-FEX vDS.
42. Power on the virtual machine.



43. Go to the Cisco UCS Manager, VM tab and under the Virtual Machines node expand the Host Blade. It will show the Host Blade Chassis/Blade. Expand the virtual machines node and verify the virtual machines are part of the VM-FEX vDS configuration on that blade with its associated VLAN.



44. Edit the setting for the powered on virtual machine that was added to VM-FEX vDS. Select Network adapter.
45. Check the status for DirectPath I/O: it shows as Active.



Follow this procedure to add the remaining ESXi hosts and virtual machines as part of the VM-FEX configuration.

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 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 pwwn of the HBAs and the Boot from SAN (BFS) policy also moves along with it. The new server now takes the same exact character of the old server, providing the true unique stateless nature of the Cisco UCS Blade Server.

The key benefits of booting from the network are:

- **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 Cisco 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 (for example, Fc1/17).

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

The 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 Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Oper Channel	Port
-----------	------	------------	------------------	--------	-----	-----------	-------------------	--------------	------

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 pwwn 20:00:00:25:b5:c1:00:af
```

```
! [B230M2-CH1-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]
```

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

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

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

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

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]
```

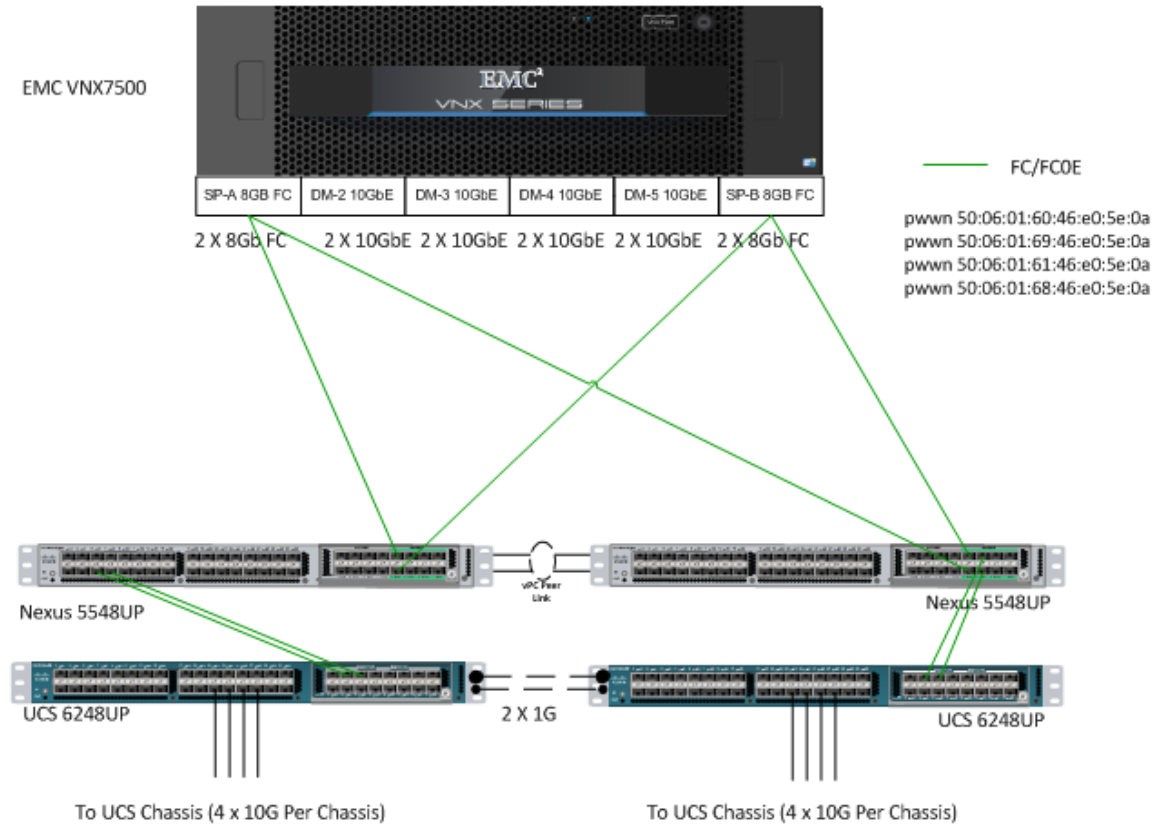
```
member pwwn 50:06:01:68:46:e0:5e:0a
```

```
[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 14. VNX7500 FC Target Ports

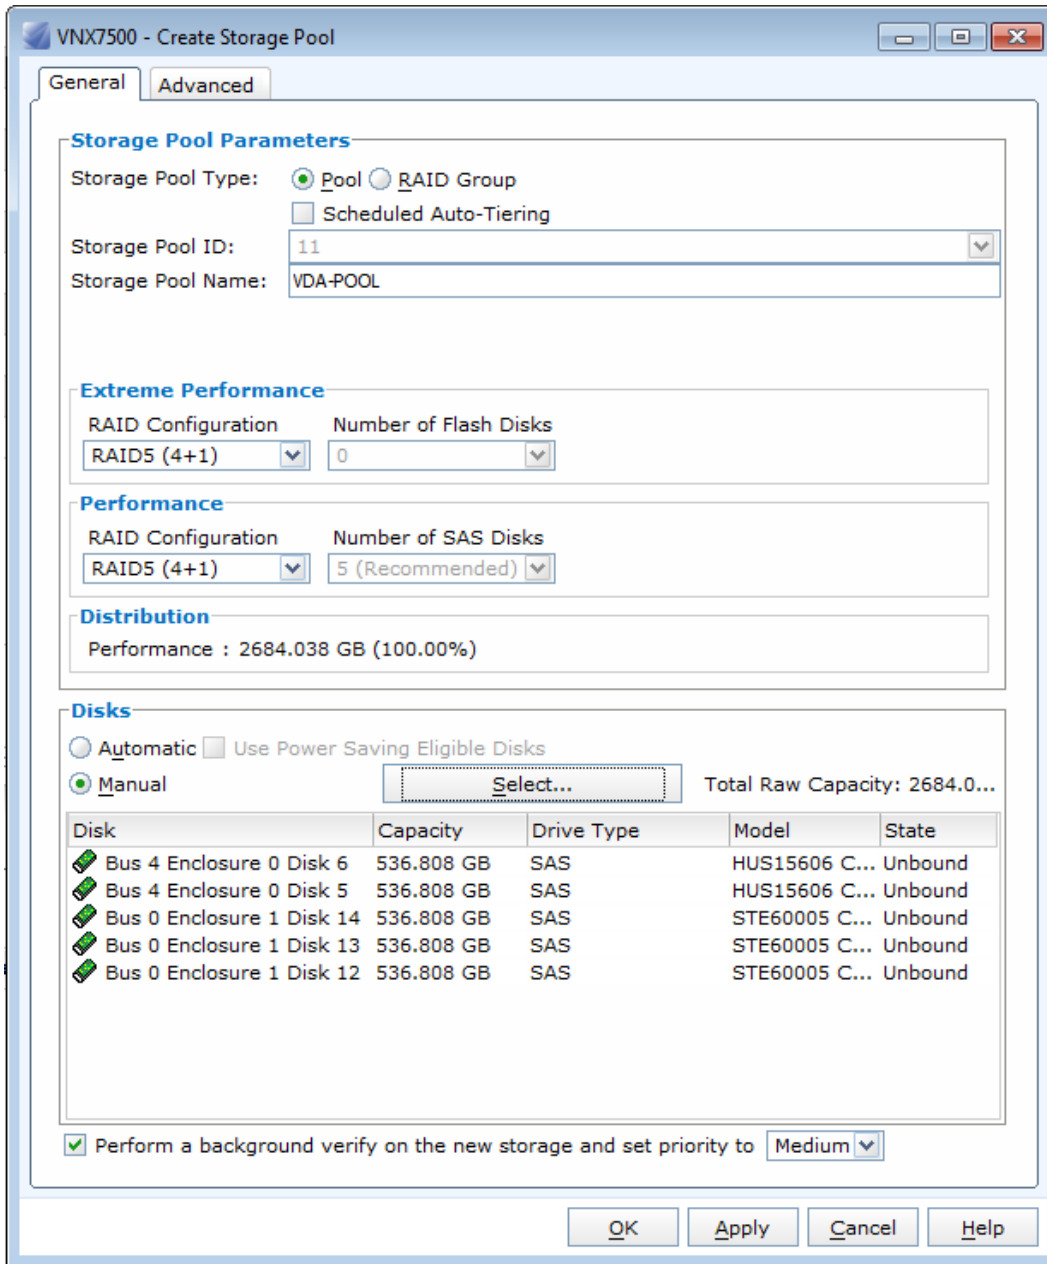


For detailed information about the Nexus 5500 series switch configuration, refer to the Cisco Nexus 5500 Series NX-OS SAN Switching Configuration Guide. (See the Reference Section of this document for a link.)

6.4.4 Configuring Boot from SAN on EMC VNX

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

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



- 33.
- 34.
2. Provision LUNs from the storage pool created in step 1. Each LUN is 12GB in size to store the ESXi hypervisor OS.
- 35.

36.
37.

3. Create a storage group, the container used for host to LUN mapping, for each of the ESXi hosts.
- 38.

39.
40.

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.
- 41.

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

42.
43.

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

44.

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

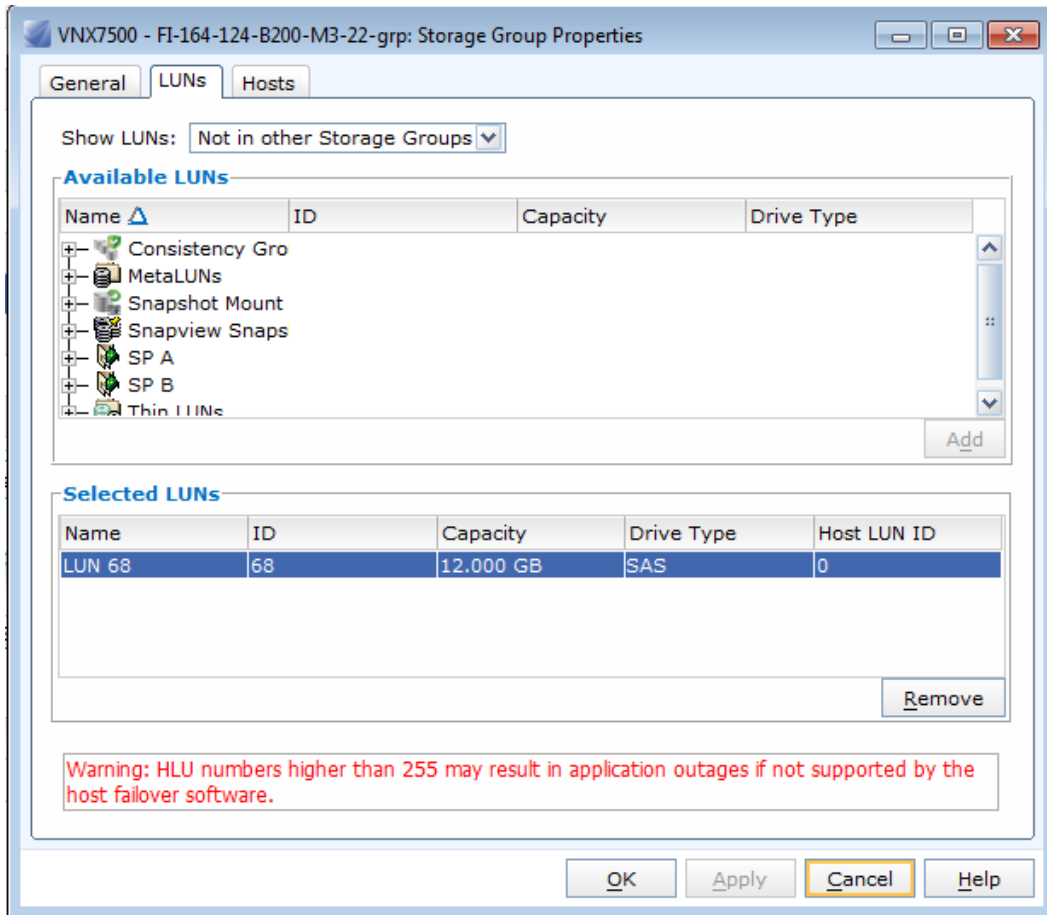
Refresh

OK Apply Cancel Help

45.
46.

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

47.
48.



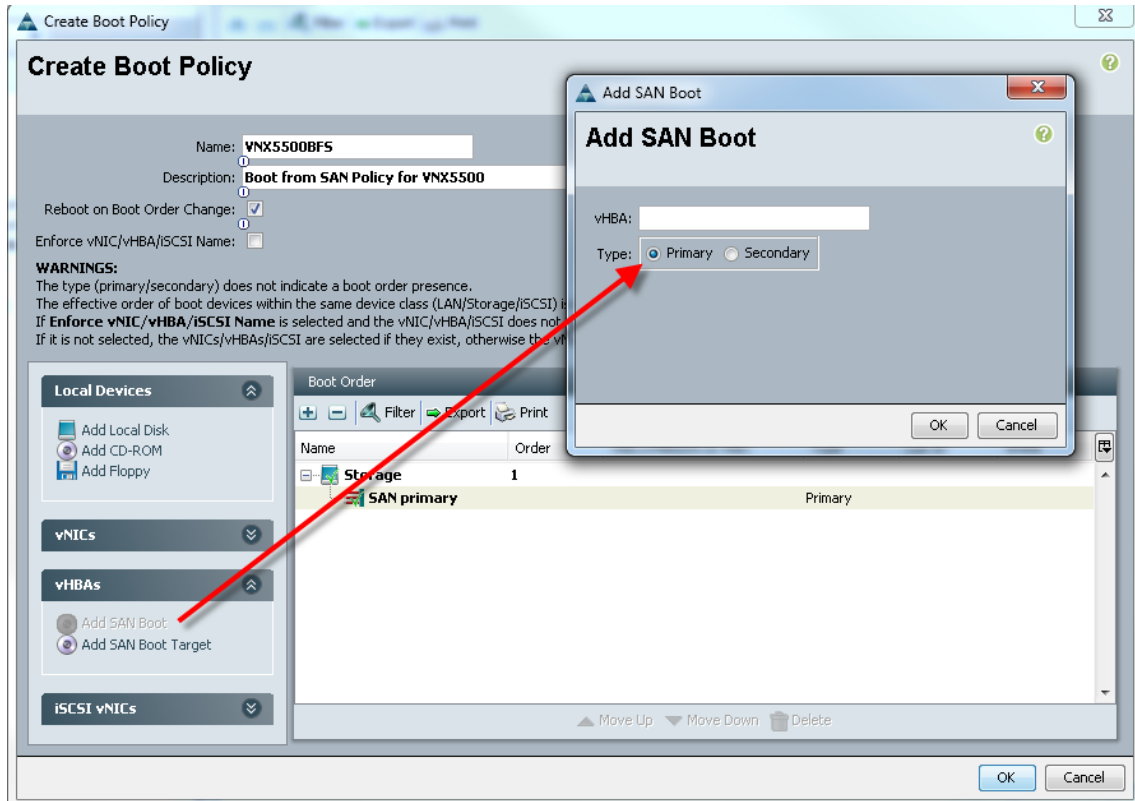
49.

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.

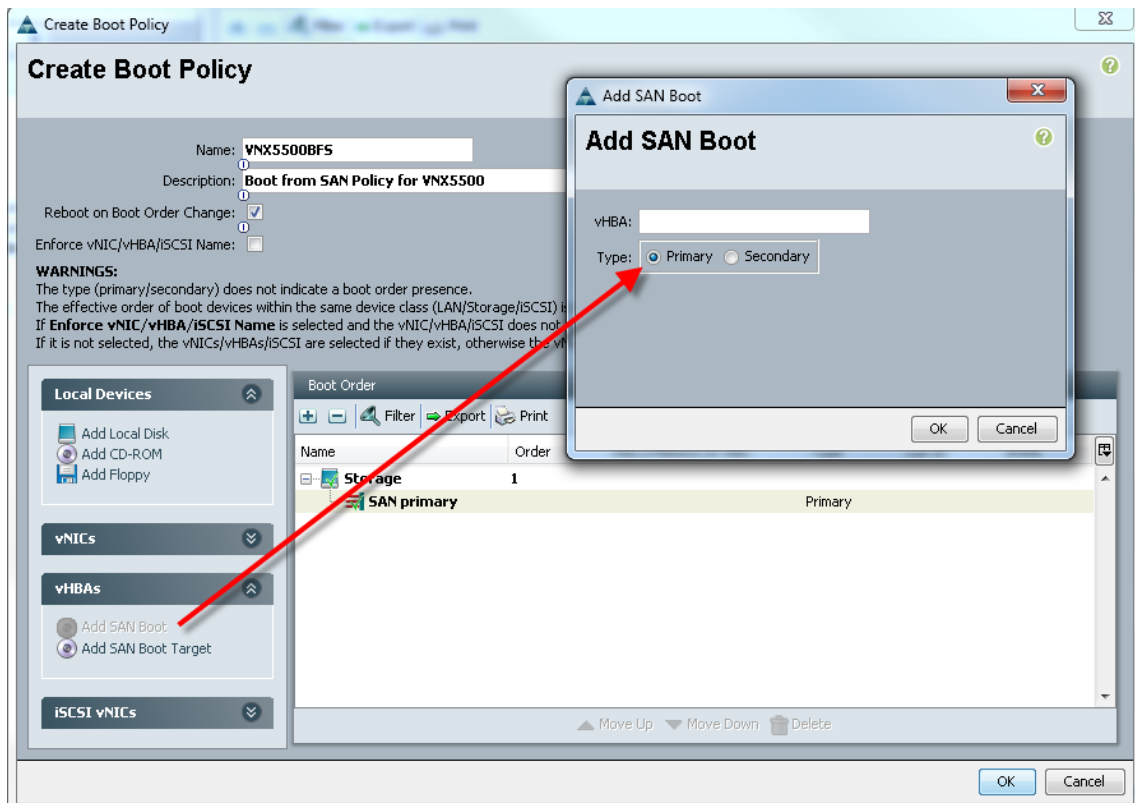
6.4.5 SAN Configuration on Cisco UCS Manager

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

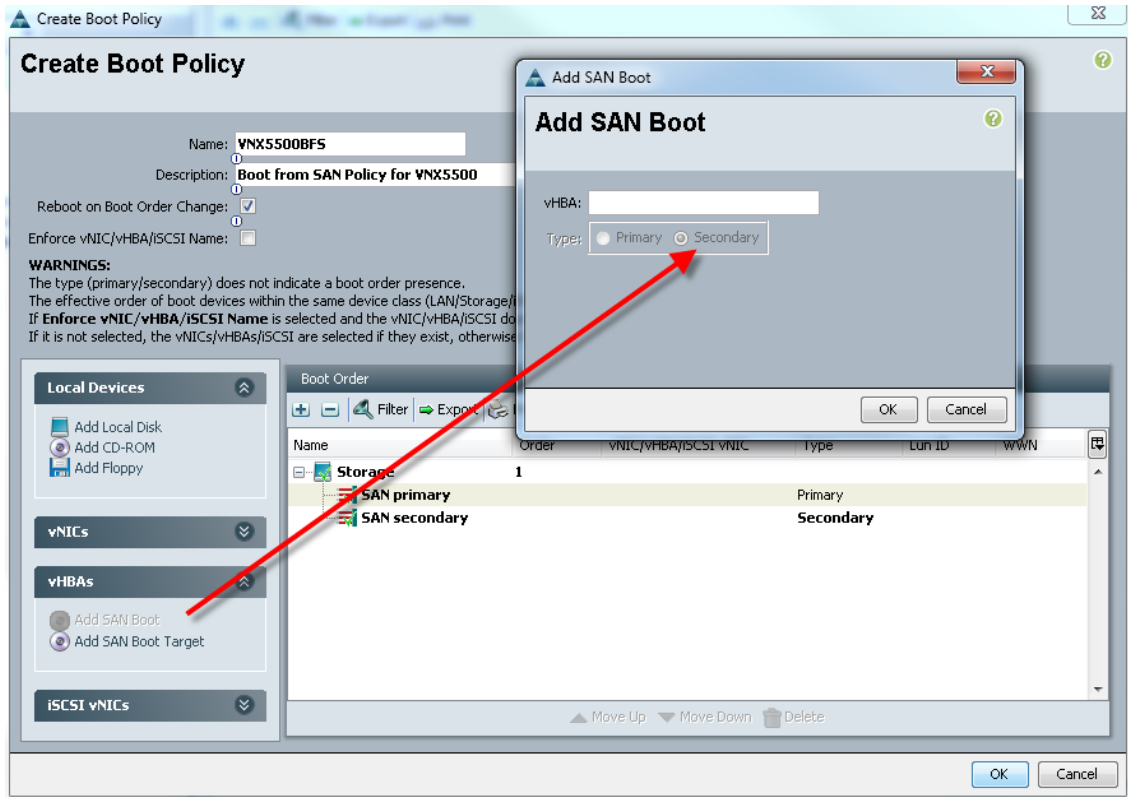
1. Add SAN Boot for primary to the new policy. The vHBA name is optional, it could be left blank and it does not have to enforce the vHBA name. Click OK.



- 50.
2. Add SAN Boot for primary to the new policy. The vHBA name is optional, it could be left blank and it does not have to enforce the vHBA name. Click OK



- 51.
3. Add SAN boot for SAN Secondary, Click OK. Leave the optional vHBA name blank.



- 52.
4. Now add Boot target WWPN to the SAN Primary, make sure this matches the EMC VNX pwwn. To avoid any typos, copy and paste from Nexus 5500 Series command as follows from each switch:

53.

54. VDI-N5548-A# show fcns database vsan 1

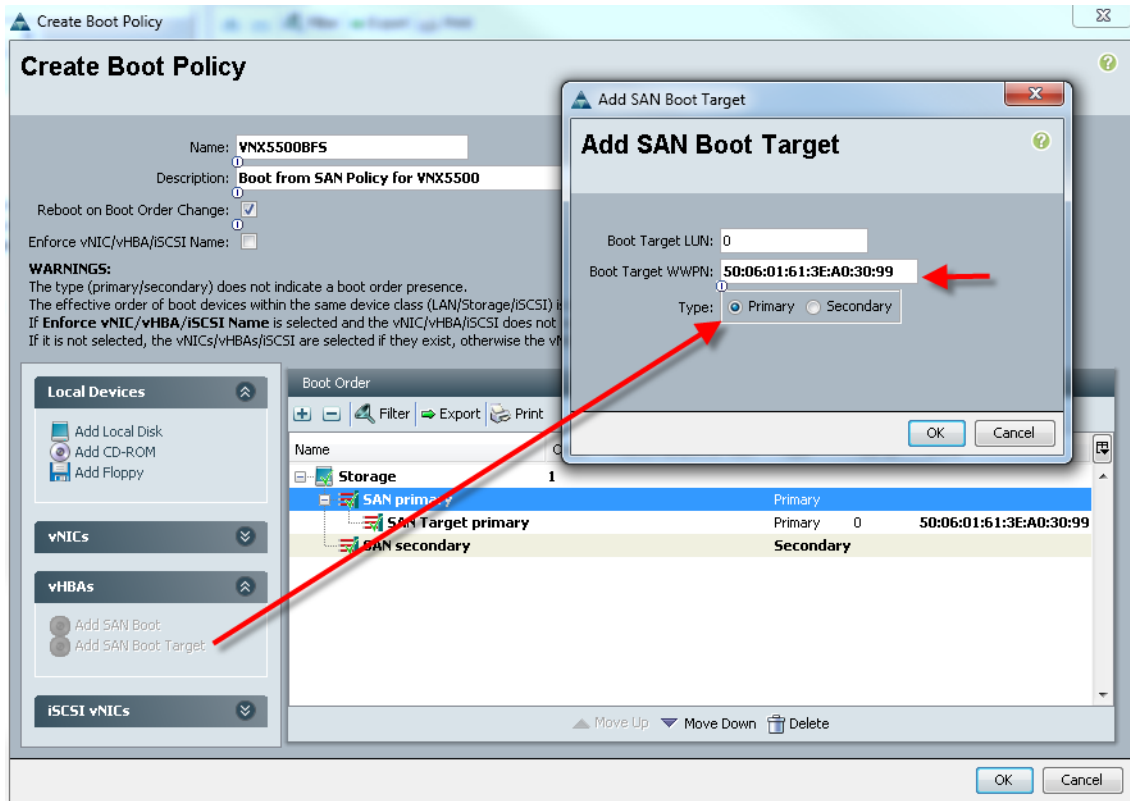
55. 0xe300ef N 50:06:01:60:46:e0:5e:0a (Clariion) scsi-fcp:both

56. 0xe301ef N 50:06:01:69:46:e0:5e:0a (Clariion) scsi-fcp:both

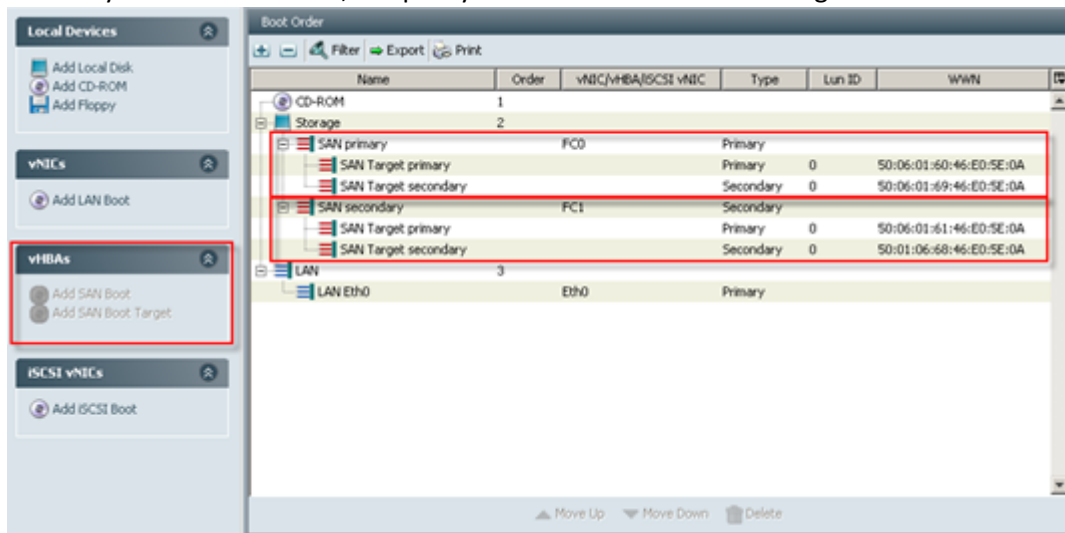
VDI-N5548-B # show fcns database vsan 1

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, the policy should look like the following:




- 57.
- 58.
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 a 4000 user Citrix mixed workload desktop virtualization environment, hypervisor boot LUNs, SQL database, and infrastructure VMs.

Figure 15. **Physical Storage Disks Layout**



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.
- Four disks (0_0_8 to 0_0_9 and 2_0_3 to 2_0_4) 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.

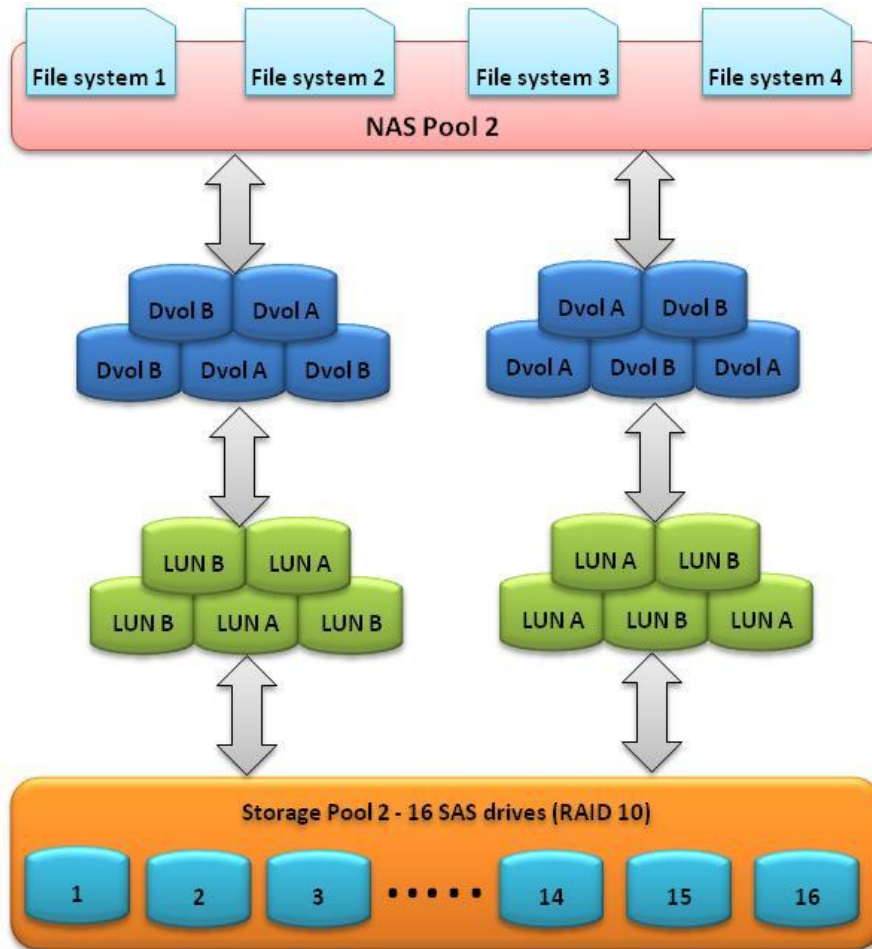


- 48 SAS disks (spread across enclosures 0_0 through 4_0) on the RAID 10 storage pool 1 are used to store PVS write cache allocated for the XenApp based virtual desktops. FAST Cache is enabled on this pool.
- 16 SAS disks (4_0_5 to 4_0_14 and 5_0_0 to 5_0_5) on the RAID 10 storage pool 2 are used to store PVS write cache allocated for the XenDesktop based virtual desktops. FAST Cache is enabled on this pool.
- 15 SAS disks (5_0_6 to 5_0_14 and 6_0_2 to 6_0_7) on the RAID 5 storage pool 3 are used to store the Personal vDisks. FAST Cache is enabled on this pool.
- Five SAS disks (6_0_8 to 6_0_12) on the RAID 5 storage pool 4 are used to store the PVS vDisks. FAST Cache is enabled on this pool.
- Ten SAS disks (7_0_0 to 7_0_9) on the RAID 5 storage pool 5 are used to store the infrastructure virtual machines.
- Five SAS disks (7_0_10 to 7_0_14) on the RAID 5 storage pool 6 are used to store the ESXi boot LUNs.
- Four SAS disks (0_1_0 to 0_1_3) on the RAID 10 storage pool 7 are used to store the SQL databases and quorum disk.
- Disks 6_0_13 to 6_0_14 and 0_1_4 to 0_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

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

Figure 16. NFS File System Layout



60.
61.

62. 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. Four 1TB file systems are then carved out of the NAS pool and are presented to the ESXi servers as four NFS datastores.

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 ensures very low latency for the data. As data ages and becomes less active, it is flushed from FAST Cache to be replaced by more active data.

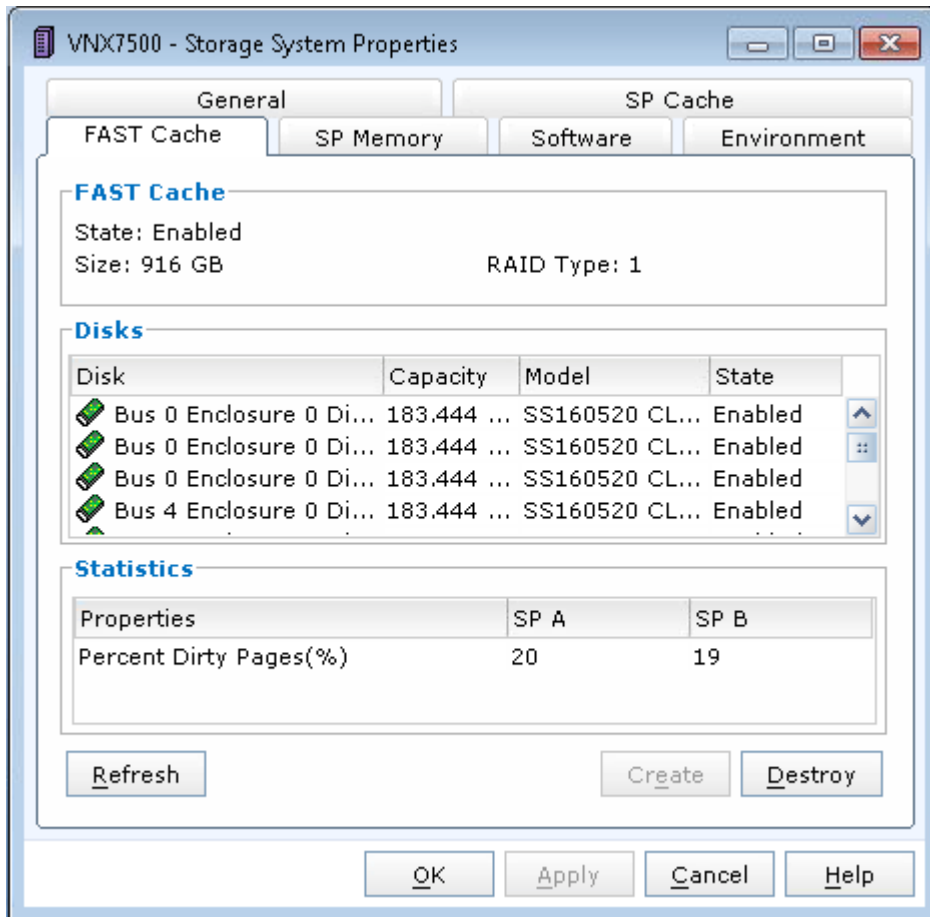
Only a small number of Flash drives are needed enabling FAST Cache to provide greater performance increases than implementing a large number of short-stroked HDDs. This results in cost savings in data center space, power, and cooling requirements that lowers overall TCO for the business.

FAST Cache is particularly suited to applications that randomly access storage with high frequency, such as Oracle and SQL OLTP databases. OLTP databases have inherent locality of reference with varied IO 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 ensure disks 1 and 2 reside in the same enclosure (likewise for disks 3 and 4) to avoid drive mirroring across enclosures.

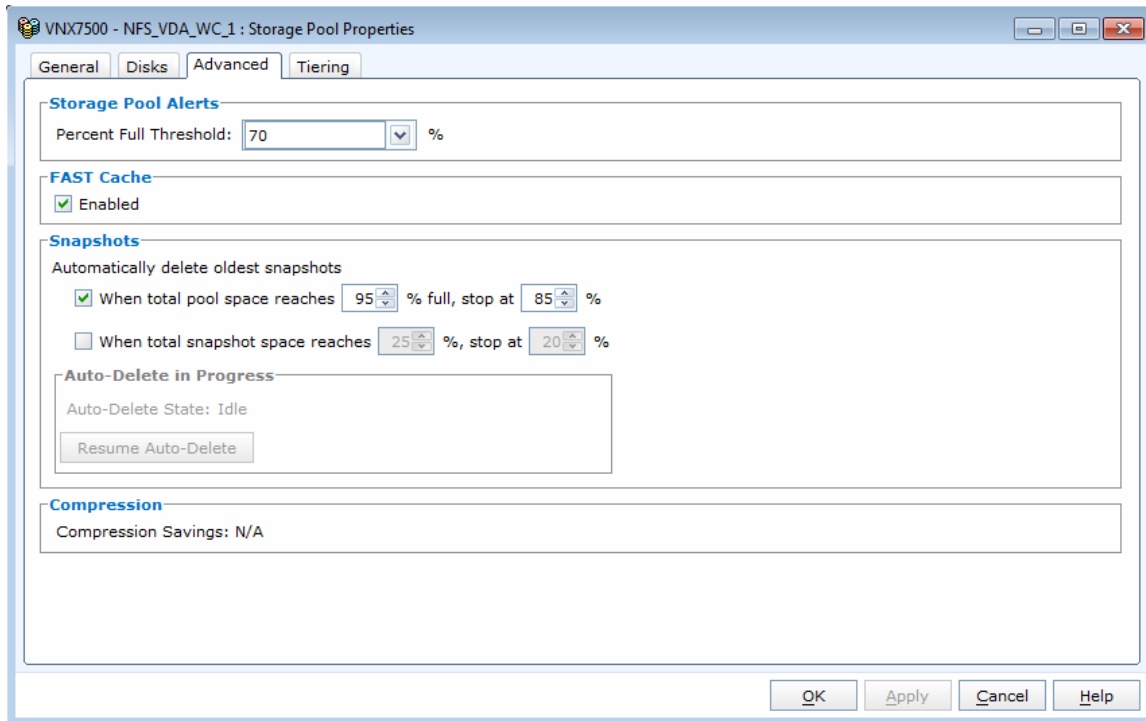
In this solution, ten 200GB SSD drives were used for FAST Cache. The ten drives were spread across the even number buses. Figure 21 shows the FAST Cache settings for VNX7500 array used in this solution.

Figure 17. VNX7500-FAST Cache tab



To enable FAST Cache for any LUN in a pool, do the following:

1. Navigate to the Storage Pool Properties page in Unisphere.
2. Click the Advanced tab.
3. Select Enabled to enable FAST Cache as shown below.



6.5.7 EMC Additional Configuration Information

The following sections detail the tuning configurations to 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
```

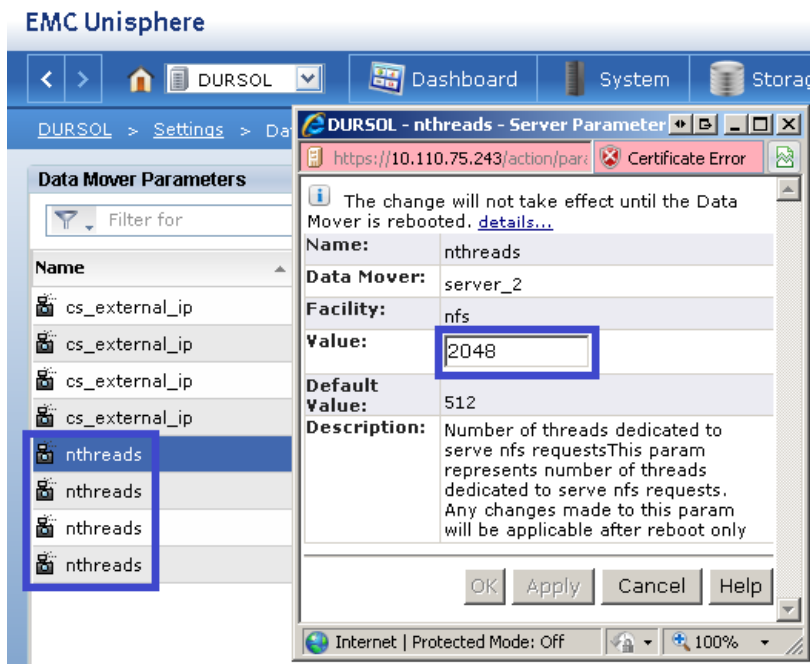
```
change_effective = reboot DataMover
```

```
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. 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**. Perform

this procedure for each of the **nthreads** Data Mover parameters listed menu. Reboot the Data Movers for the change to take effect.



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 uxfs,perm,rw,uncached
fs6 on /fs6 uxfs,perm,rw,uncached
fs7 on /fs7 uxfs,perm,rw,uncached
fs8 on /fs8 uxfs,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. 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.



Path:	/FS1
Data Mover:	server_5
File System Name:	FS1
Read Only:	<input checked="" type="radio"/> Read/Write <input type="radio"/> Read Only
Access-Checking Policy:	<input type="radio"/> NT - CIFS client rights checked against ACLs; NFS client rights checked against ACLs and permission bits <input type="radio"/> UNIX - NFS client rights checked against permission bits; CIFS client rights checked against permission bits AND ACLs <input type="radio"/> SECURE - Both NFS and CIFS client rights checked against both permission bits AND ACLs <input checked="" type="radio"/> NATIVE - NFS client rights checked against permission bits; CIFS client rights checked against ACLs <input type="radio"/> MIXED - Both NFS and CIFS client rights checked against ACL; Only a single set of security attributes maintained <input type="radio"/> 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:	<input type="checkbox"/>
Cifs Oplocks Enabled:	<input type="checkbox"/>
Set Advanced Options:	<input checked="" type="checkbox"/>
Use NT Credential:	<input type="checkbox"/>
Direct Writes Enabled:	<input checked="" type="checkbox"/>
Prefetch Enabled:	<input checked="" type="checkbox"/>
Multi-Protocol Locking Policy:	<input checked="" type="radio"/> nolock <input type="radio"/> writelock <input type="radio"/> rmllock
CIFS Sync Writes Enabled:	<input type="checkbox"/>
CIFS Notify Enabled:	<input type="checkbox"/>
CIFS Notify Trigger Level:	512
CIFS Notify On Access Enabled:	<input type="checkbox"/>
CIFS Notify On Write Enabled:	<input type="checkbox"/>

6.6 Cisco UCS Manager Configuration for VMware ESXi 5.1

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

<u>Role</u>	<u>Blade Server Used</u>	<u>Operating System Deployed</u>
Infrastructure	Cisco UCS B200 M3	ESXi 5.1
VDI Hosts	Cisco UCS B200 M3	ESXi 5.1

To support this hardware platform, a single service profile template was created, utilizing various policies created earlier.

The service profile template was 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 5.1 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 UCS Manager (Fabric Interconnects,) on the Nexus 1000V Virtual Switch Modules in the vCenter XenDesktop Clusters, and on the VM-FEX on the XenApp 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 list identifies them and describes their use:

<u>VLAN Name</u>	<u>VLAN ID</u>	<u>Use</u>
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-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 5.1

Install VMware ESXi 5.0 Update 1

ESXi was installed using an interactive Installation method.

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

http://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc_50%2FGUID-26F3BC88-DAD8-43E7-9EA0-160054954506.html

Additionally the VMware ESXi 5.1 Patch ESXi510-201210001 was installed on ESXi servers hosting the XenDesktop Hosted VMs and XenApp Hosted Shared Desktops
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2034548

6.7.2 Install and Configure vCenter 5.1

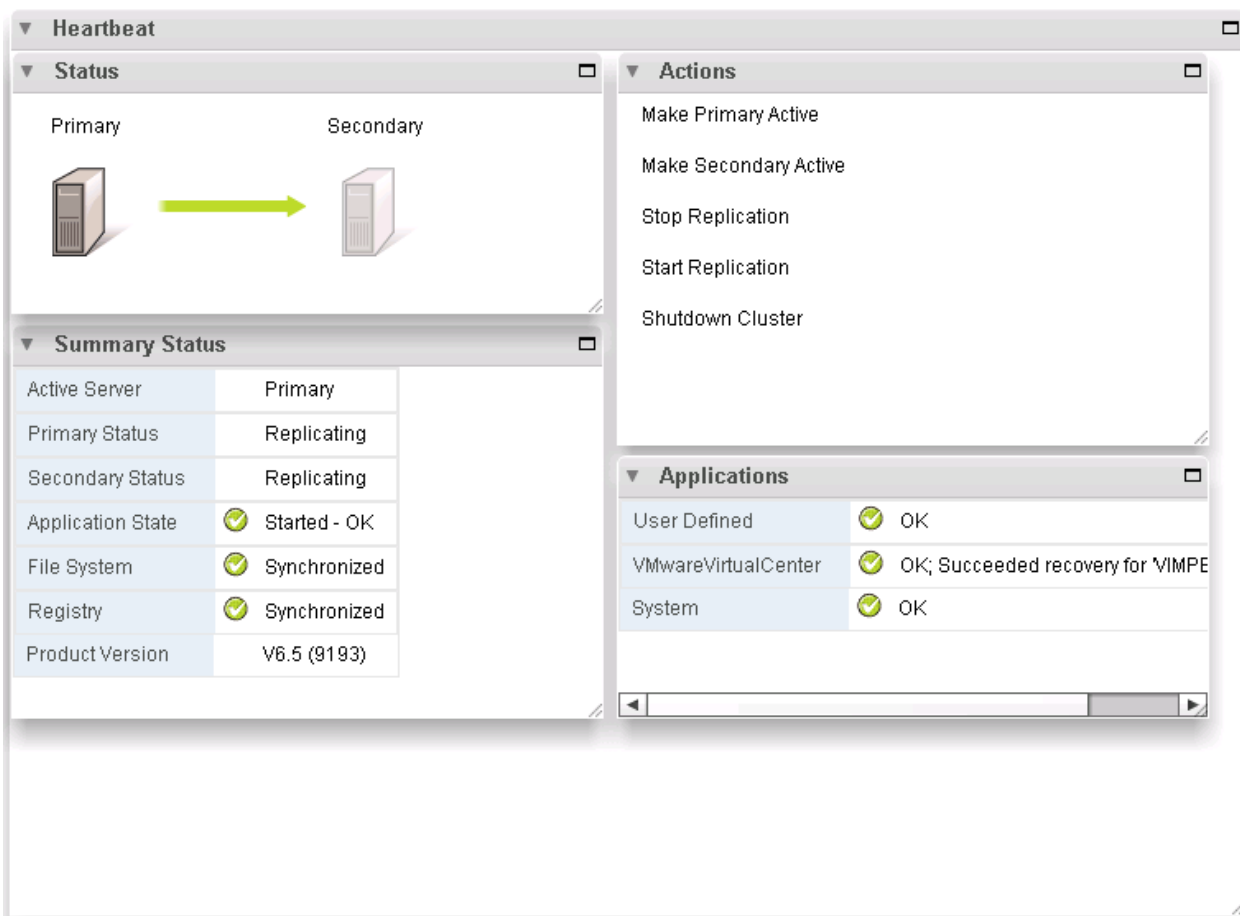
To manage hypervisors and virtual machines a dedicated vCenter server instance was installed on Windows 2008 R2 virtual machine.

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))

Install and configure vCenter


1. Install the Microsoft® SQL Server® 2008 R2 Native Client for ODBC connections (<http://www.microsoft.com/en-us/download/details.aspx?id=16978> look for Native Client for your architecture)
2. Create a System DSN (control panel, administrative tools, Data Sources ODBC) and connect to your vCenter-SQL server. **Note:** Ensure 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).
7. Heartbeat 6.5 was installed to support High Availability requirement for vCenter <http://www.vmware.com/support/heartbeat/doc/vcenter-server-heartbeat-65-release-notes.html>



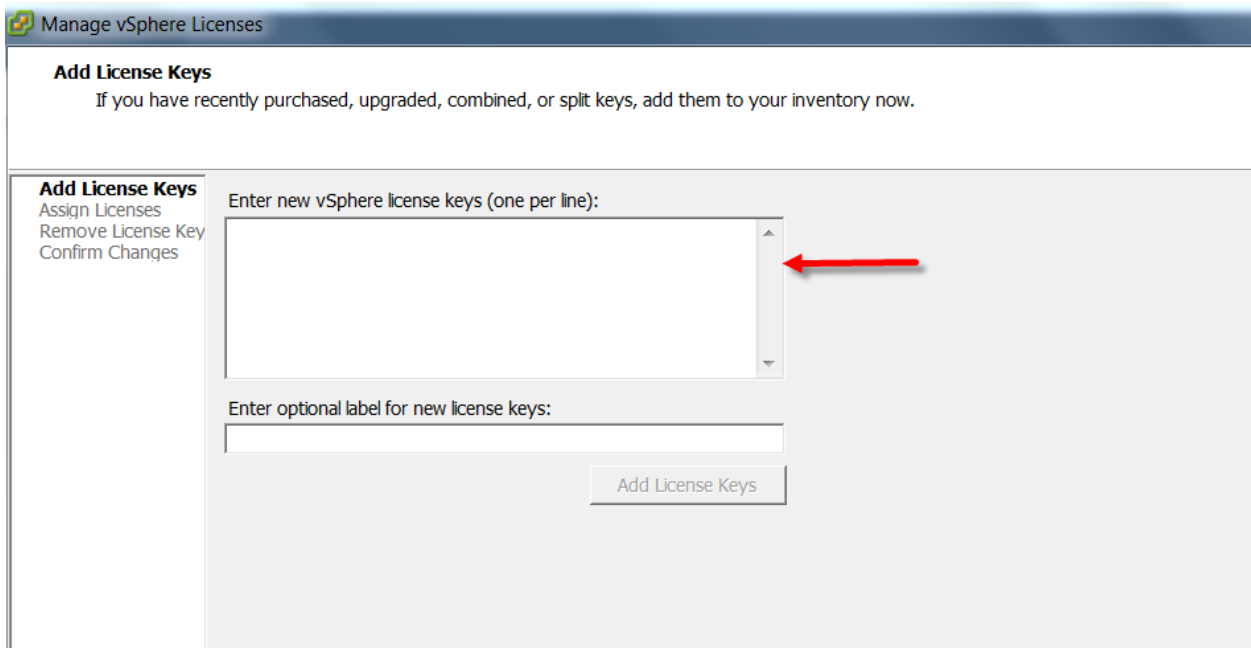


Install Licenses

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



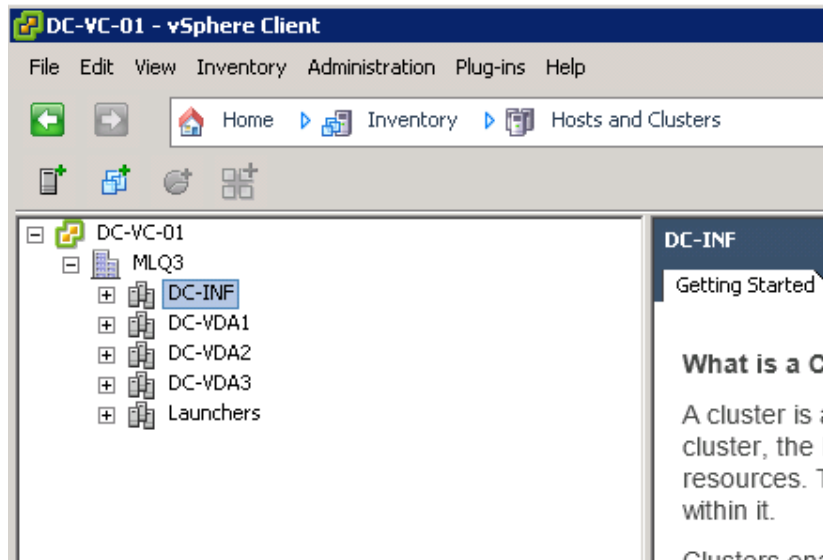
4. Add **License keys** for **vCenter** and **Hosts**



6.7.4 ESXi 5.1 Cluster Configuration

The 30 Cisco UCS B200 M3s, and 20 Cisco UCS B250 M2s ESX hosts were configured into five Clusters:

- DC-INF (5 Cisco UCS B200 M3 blade servers)
- DC-VDA1 (7 Cisco UCS B200 M3 blade servers)
- DC-VDA1 (7 Cisco UCS B200 M3 blade servers)
- DC-VDA1 (11 Cisco UCS B200 M3 blade servers)
- Launchers (20 Cisco UCS B250 M2 blade servers)



The DC-INF cluster was used to host all of the virtualized servers within the VDA Infrastructure. The following clusters were used to host 4100 desktops:

- DC-VDA1: 1085 Virtual Desktops with Tier 0 (SSD) Storage
- DC-VDA2: 1015 Virtual Desktops with Personal vDisk
- DC-VDA3: 2000 XenApp 6.5 Remote Desktops

Launchers cluster was used to host VSI Client Launcher Virtual Machines, Launcher PVS, and VSI console.

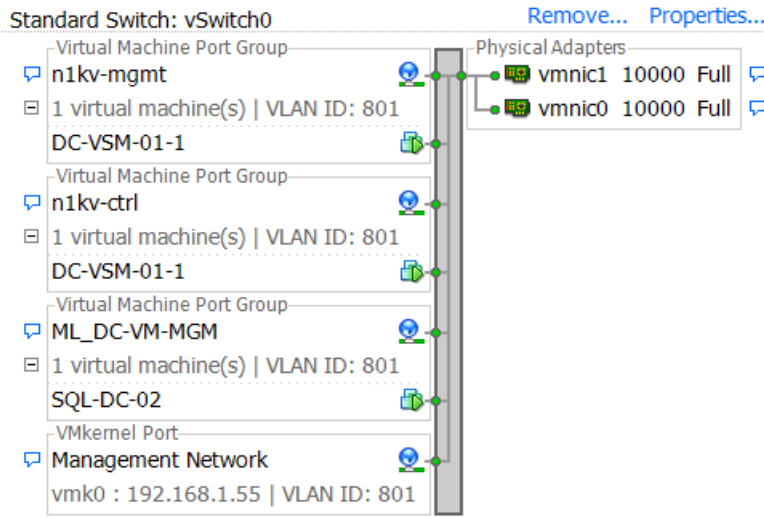
6.7.4.1 DC-INF Infrastructure Cluster

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 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 XenDesktop Hosted Virtual Desktop Clusters

The following clusters were used to host 2100 desktops:

- DC-VDA1 (1085 XenDesktop Pooled Desktops with Tier 0 Storage)
- DC-VDA2 (1015 XenDesktop Personal vDisk Desktops)

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

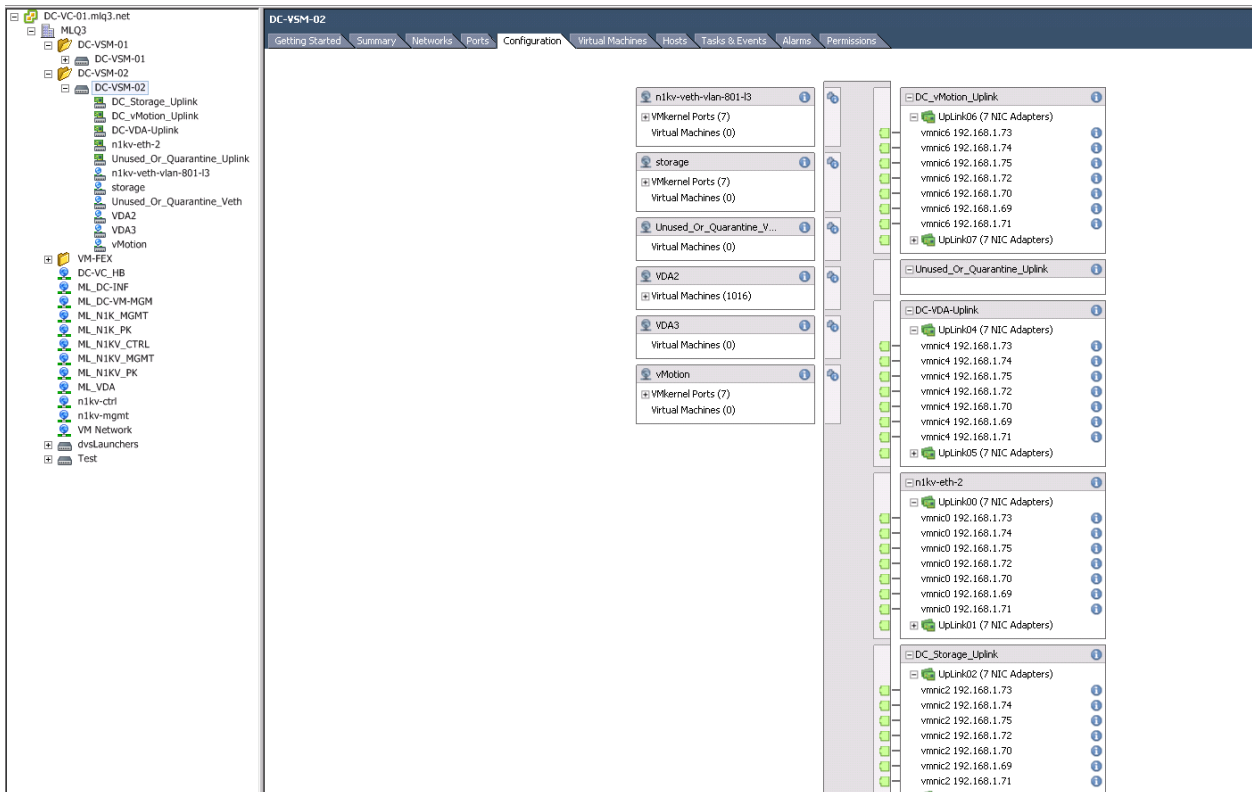


The Nexus 1000V switches were configured to manage networking for both ESX Clusters hosting virtual desktops, working in concert with the UCS Fabric Interconnects and Nexus 5548UP access layer switches to provide end to end Quality of Service for network communications, insuring the highest quality virtual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

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

The screenshot displays the vCenter configuration interface for a host named DC-VSM-01. The interface is divided into several sections:

- Left Panel (Inventory):** A tree view showing the host's configuration hierarchy, including folders like DC-VC-01, MLQ3, DC-VSM-01, and various network and storage components.
- Configuration Tab:** The main configuration area, currently showing the 'Configuration' sub-tab. It lists several components with their associated virtual machines (VMs):
 - n1kv-veth-vlan-801-13:** 0 VMs
 - Storage:** 7 VMkernel Ports, 0 VMs
 - Unused_Or_Quarantine_V...:** 0 VMs
 - VDA:** 937 VMs
 - VDA1:** 156 VMs
 - vMotion:** 7 VMkernel Ports, 0 VMs
- Right Panel (Network and Storage Adapters):** A detailed view of network and storage adapters, showing their status and associated VMkernel ports:
 - DC_vMotion_Uplink:** 7 NIC Adapters (UpLink06-07), each with 7 VMkernel ports (vmnic6).
 - Unused_Or_Quarantine_Uplink:** 0 adapters.
 - DC-VDA-Uplink:** 7 NIC Adapters (UpLink04-05), each with 7 VMkernel ports (vmnic4).
 - n1kv-eth-2:** 7 NIC Adapters (UpLink00-01), each with 7 VMkernel ports (vmnic0).
 - DC_Storage_Uplink:** 7 NIC Adapters (UpLink02-03), each with 7 VMkernel ports (vmnic2).



6.7.4.3 XenApp Hosted Shared Desktop Cluster


The following cluster was used to host 2000 Hosted Shared Desktop sessions:

- DC-VDA3

This cluster was configured with Cisco VM-FEX technology.

The VM-FEX switches were configured to manage networking for the ESX Cluster hosting shared server desktops, working in concert with the Cisco UCS Fabric Interconnects and Nexus 5548UP access layer switches to provide end to end Quality of Service for network communications, insuring the highest quality virtual desktop end user experience.

The study has proven that the VM-FEX technology offloads network related CPU cycles that allows approximately 8% more hosted shared desktop sessions on a Cisco UCS B200 M3 blade server.



6.7.4.4 Login VSI Launcher Cluster

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

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

The Launcher cluster was connected to the Data Center clusters via 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. Thousands of target devices can use a single vDisk shared across multiple PVS servers in the same farm, making virtual desktop management easier

Disk storage management is very important because 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\BNTFTP\Parameters Port

The TFTP port defaults to UDP 69.

Additionally, Netscaler VPX was configured for TFTP Load Balancing in One-Arm Mode ([CTX131954](#))



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)

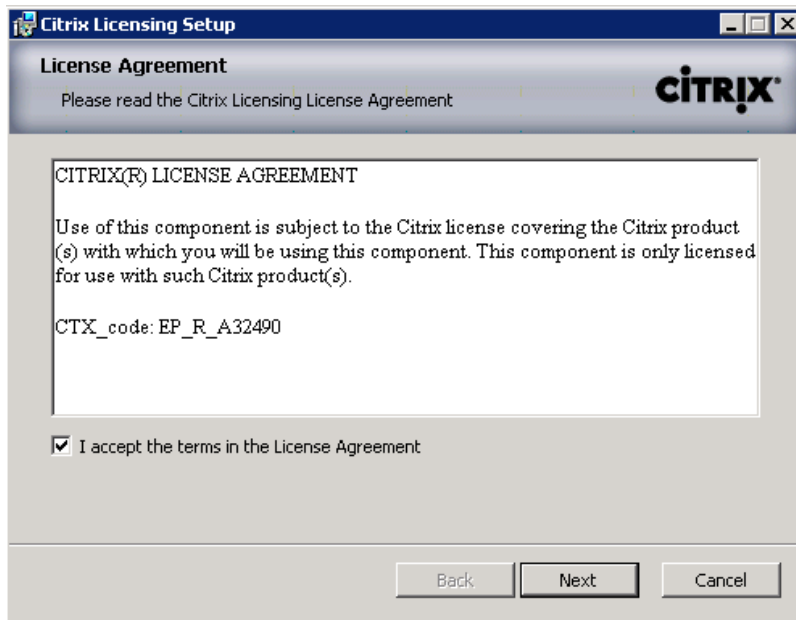
1. Open Computer Management.
2. Click on Add Role.
3. Select Web Server (IIS).
4. Click Next.
5. Click Next.
6. Under Application Development select ASP.NET.



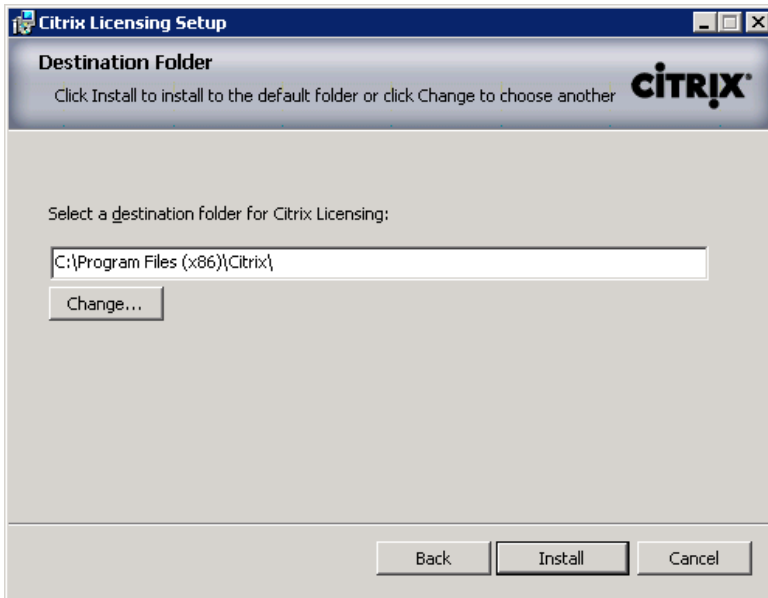
7. Click Add Required Role Services.
8. Click Next.
9. Click Install.
10. Click Close to complete the installation process.

6.8.3.2 Install License Server

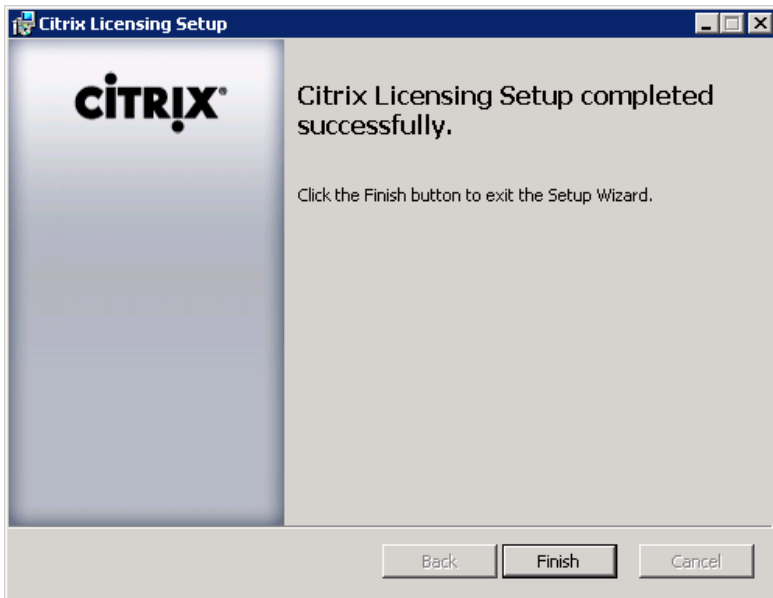
1. Locate and launch CTX_Licensing.msi.
2. Select I accept the terms in the License Agreement.



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

Specify the numbers for the following ports:

License Server Port (default 27000): 26994

Vendor Daemon Port (default 7279): 7279

Management Console Web Port (default 8082): 8082

Set a password for the License Administration Console:

Username: admin

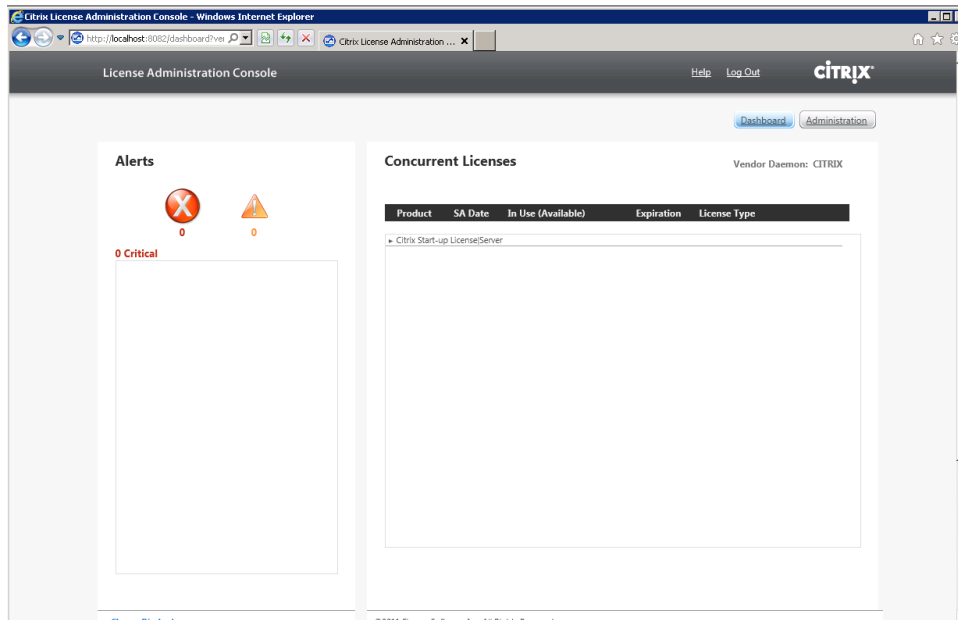
Password: XXXXXXXXXX

Confirm Password: XXXXXXXXXX

OK Cancel

3. Click OK.

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



5. Click the Administration button.

6. Enter the Admin credentials.



Log On

User Name:

Password:

Display Language:

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

License Administration Console User Name: admin Help Log Out **CITRIX**

[Dashboard](#) [Administration](#)

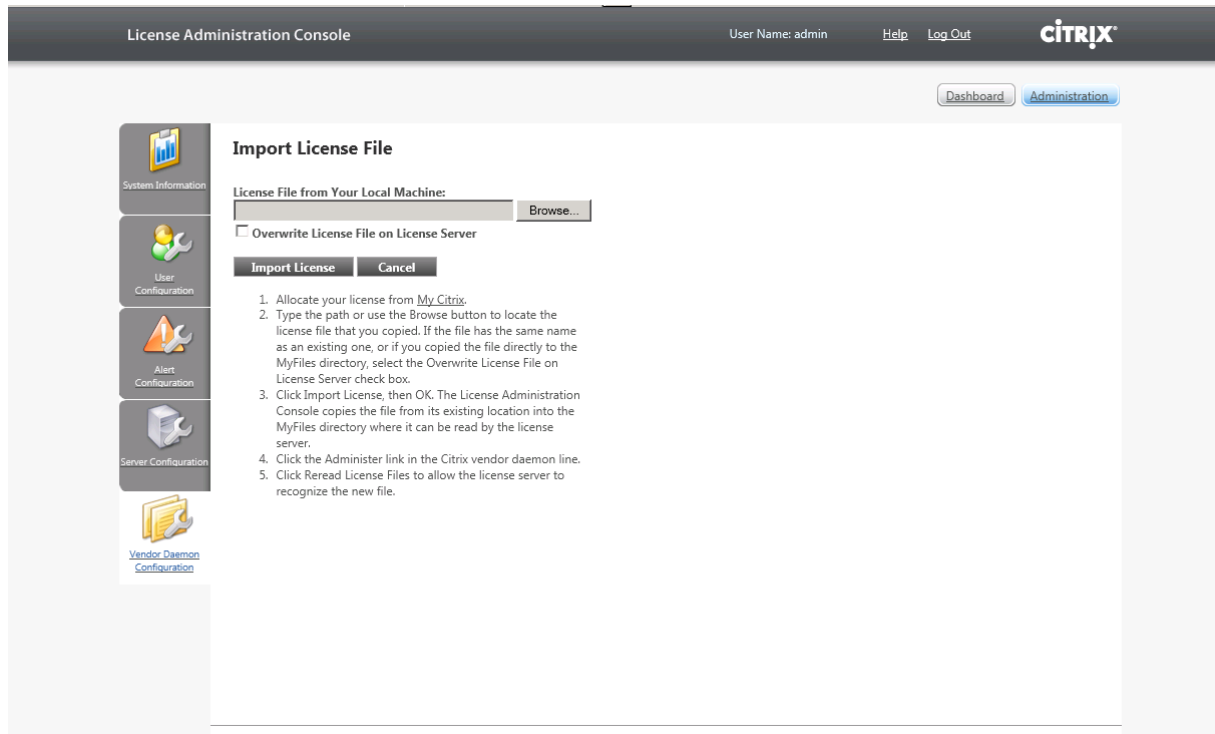
Vendor Daemon Configuration

Import License

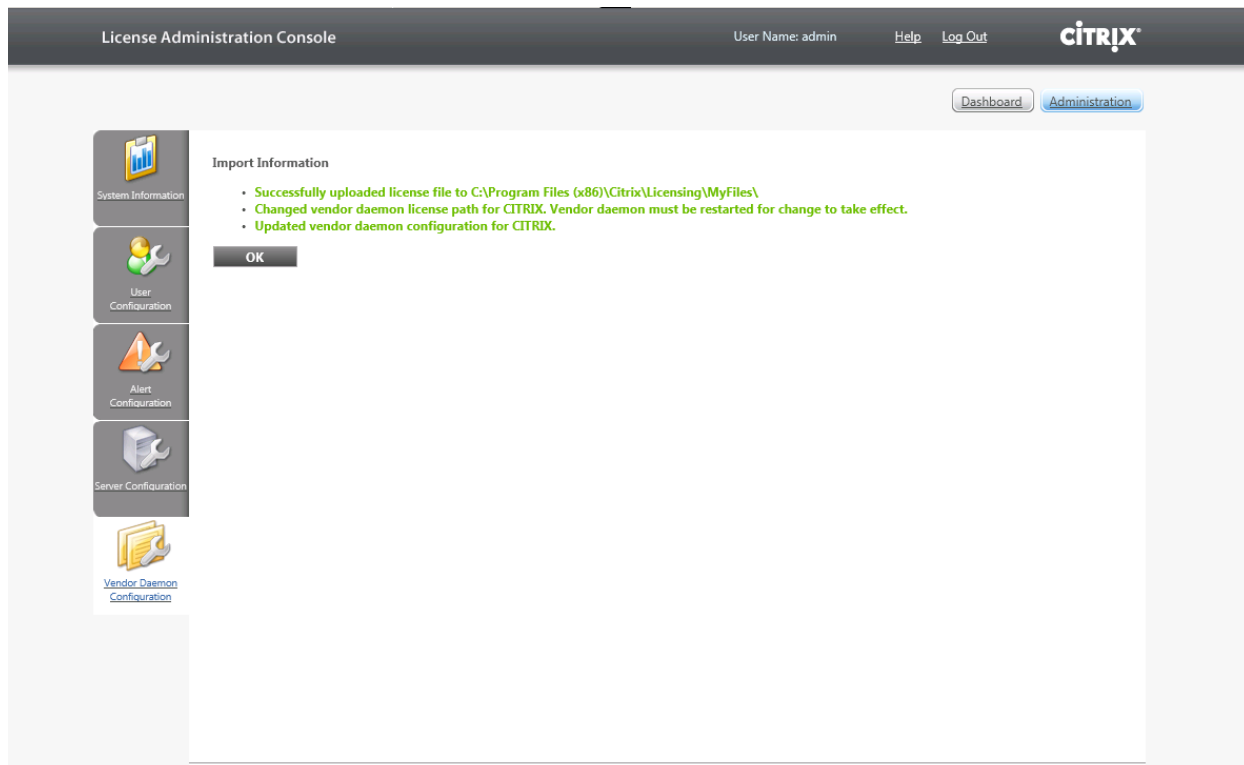
Name	Status	Version	Port		
CITRIX	RUNNING	11.10	7279	Administer	Delete

- System Information
- User Configuration
- Alert Configuration
- Server Configuration
- Vendor Daemon Configuration**

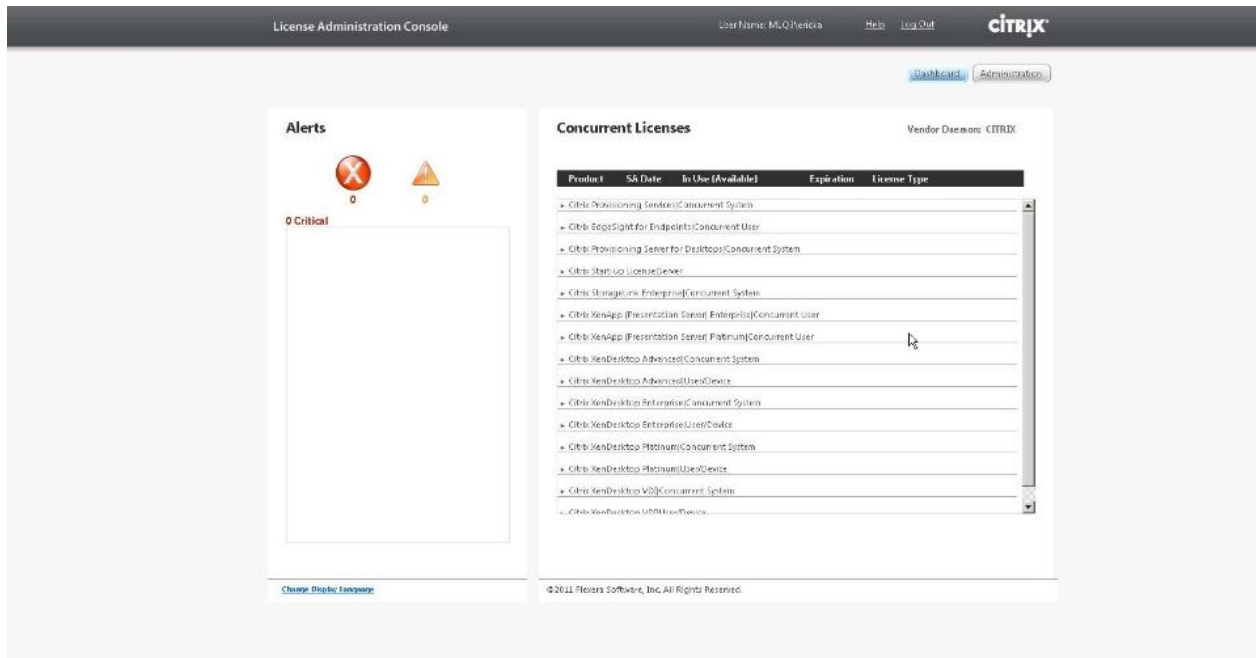
9. Click Import License.
10. Click Browse to locate the license file you are applying to the server.



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



14. Click OK.
15. Click the Dashboard button.
16. Validate that the necessary licenses have been installed.

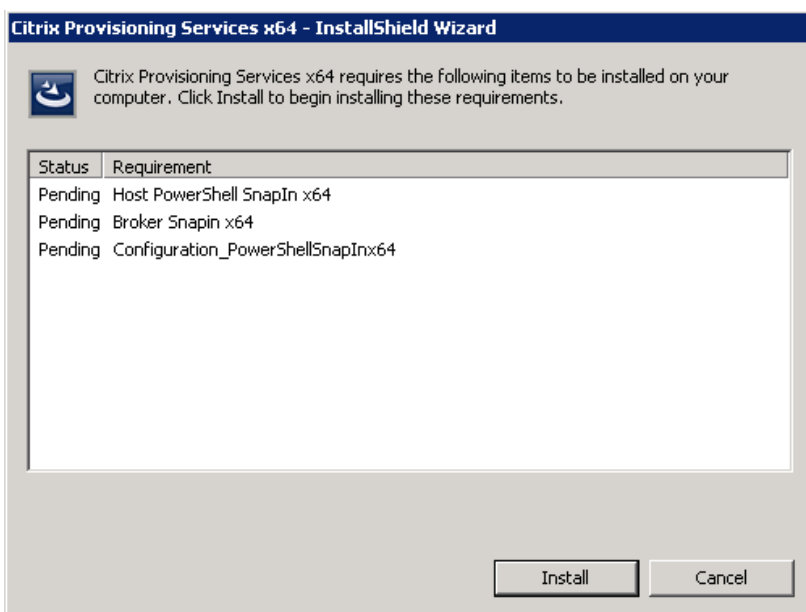


6.8.4 Install Provisioning Services 6.1

6.8.4.1 Base Install

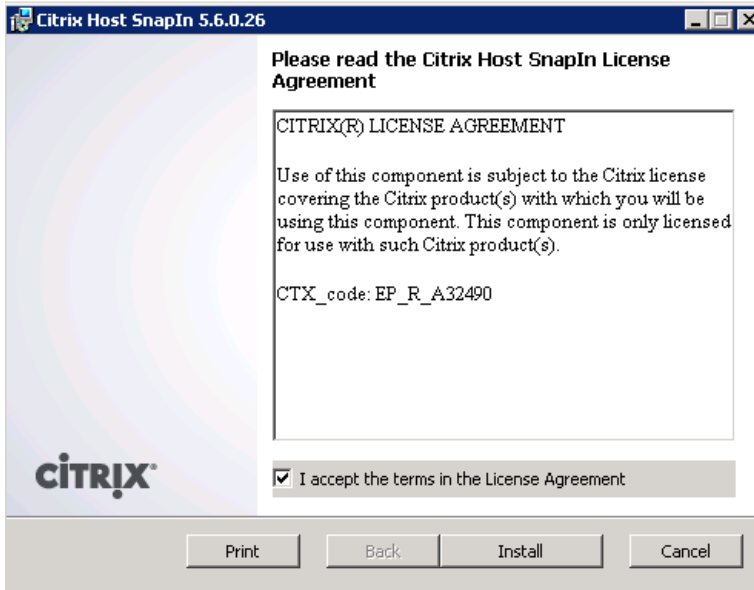
The following steps outline the process taken to install Provisioning Services 6.1

1. Locate the PVS_Server_x64.exe and run the executable.
 2. If prompted to installed required software click Install.
- 63.

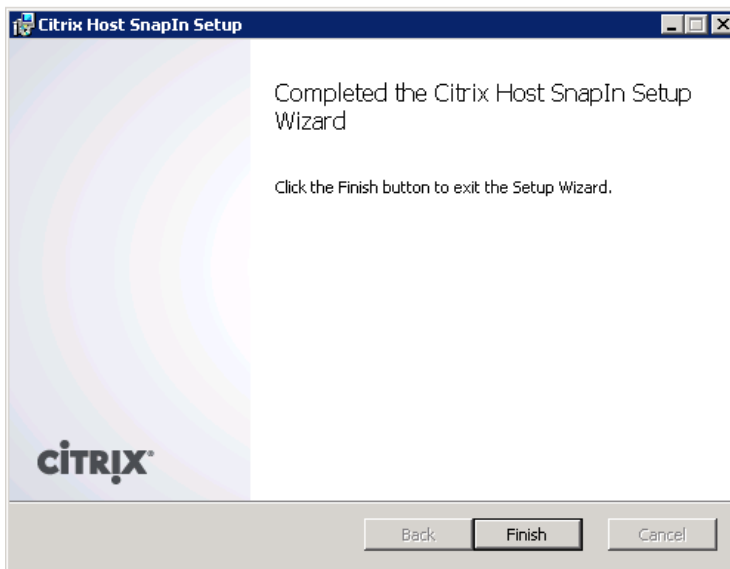


64.

- 65.
- 3. Select "I accept the terms in the License Agreement" and Click Install.
- 66.



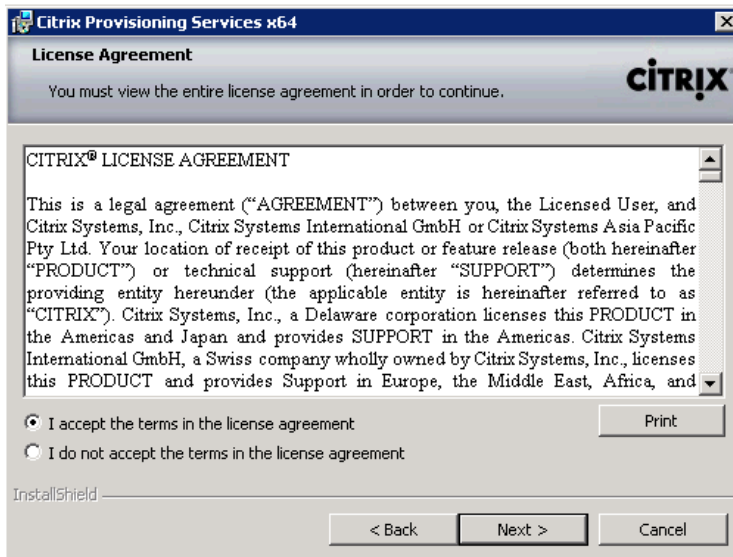
- 67.
- 68.
- 4. Click Finish.



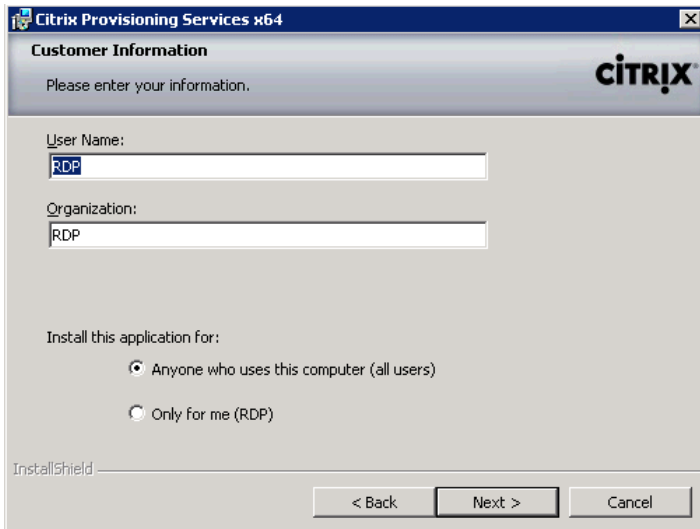
- 69.
- 70.
- 5. Click Next.
- 71.



6. Select I accept the terms in the license agreement.
- 72.

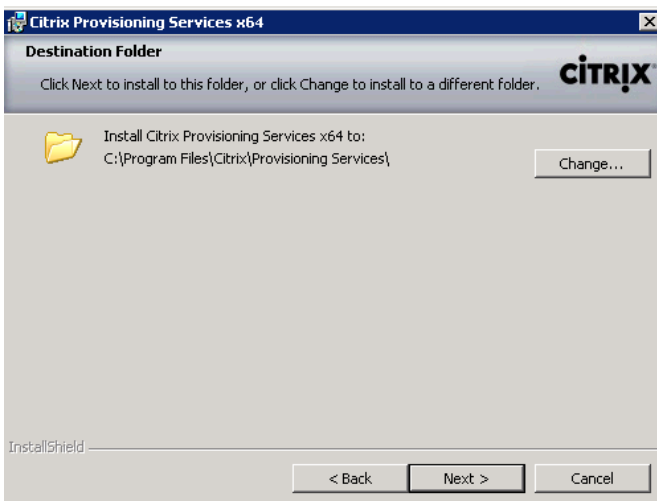


- 73.
- 74.
7. Enter the User name and Organization specific for your environment.
 8. Select Anyone who uses this computer (all users).
- 75.



9. Click Next.

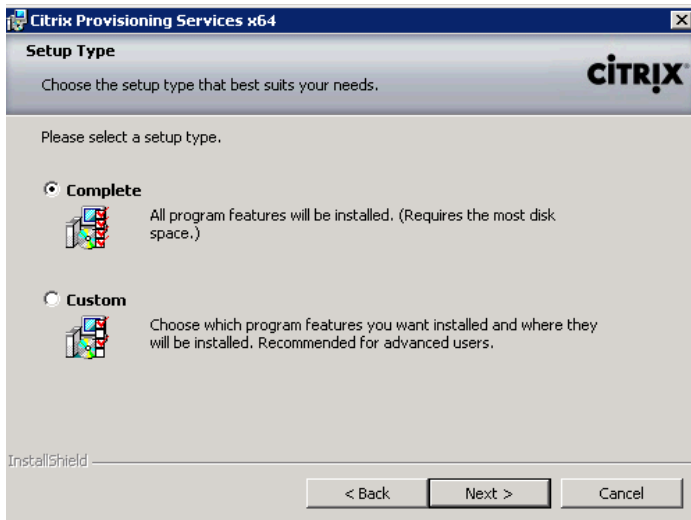
10. Leave the Destination Folder as the default location.



11. Click Next.

12. Select Complete.

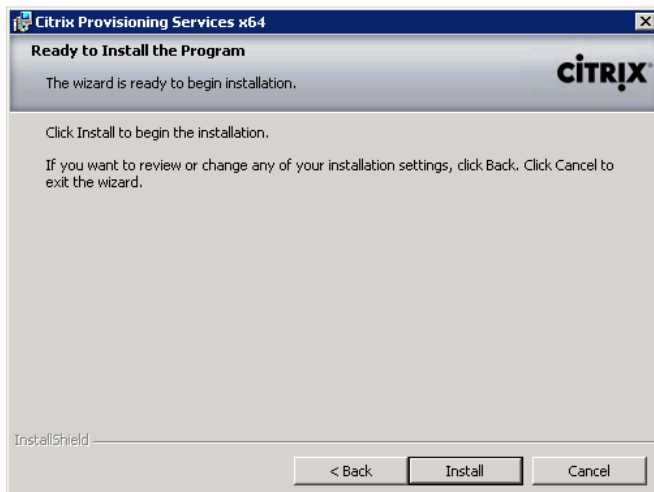
76.



13. Click Next.

14. Click Install to begin the PVS installation process.

77.



15. Click Finish to complete the installation.

78.



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.

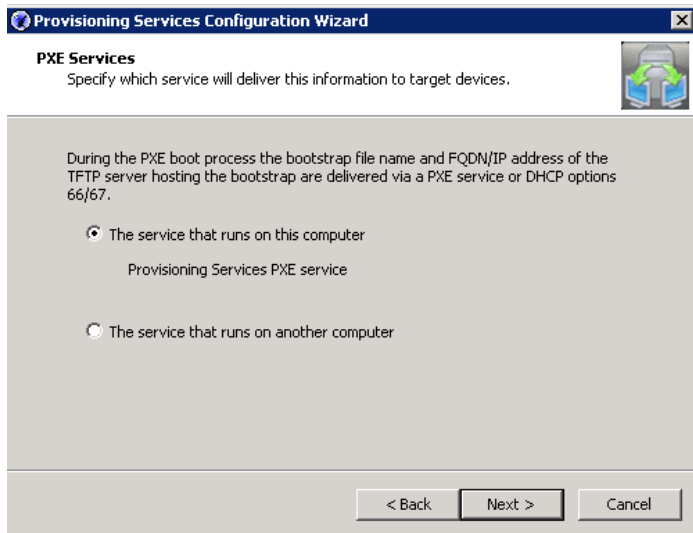


2. Click Next.
3. In the DHCP services window, select the service that runs on another computer.



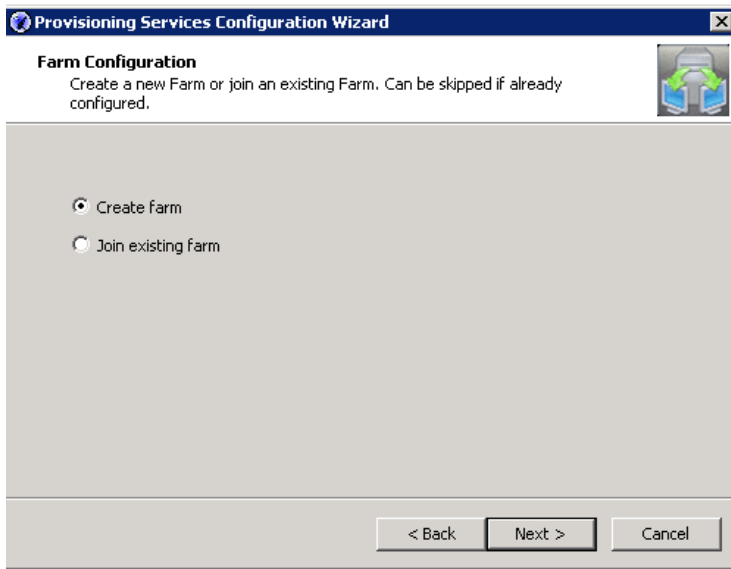
4. Click Next.

5. For PXE services, select the service that runs on this computer.

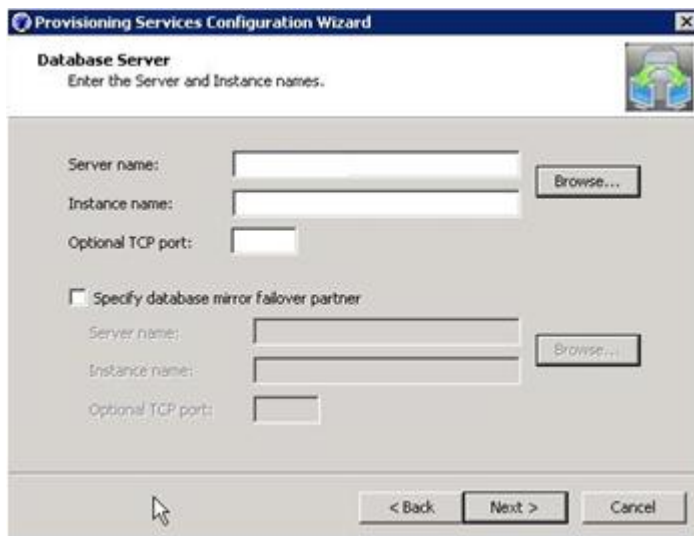


6. Click Next.

7. Select Create Farm.



8. Click Next.
9. Within the Database Server window, enter the DB Server name or IP address.



10. Click Next.
11. In the New Farm window, enter the environment specific information for the Farm Name, Site Name, and Collection Name. Additionally, choose the appropriate Active Directory group that will be identified as the Farm Administrators .



The screenshot shows the 'New Farm' step of the Provisioning Services Configuration Wizard. The title bar reads 'Provisioning Services Configuration Wizard'. The main heading is 'New Farm' with the instruction 'Enter the new Database and Farm names.' Below this, there are four input fields: 'Database name:' with a dropdown menu showing 'ProvisioningServices', 'Farm name:' with a text box containing 'Farm', 'Site name:' with a text box containing 'Site', and 'Collection name:' with a text box containing 'Collection'. There are two radio button options: 'Use Active Directory groups for security' (selected) and 'Use Windows groups for security'. Below these is a 'Farm Administrator group:' dropdown menu showing 'mlq3.net/Users/Domain Admins'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

12. Click Next.

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

14. Click Next.

The screenshot shows the 'Network Communications' step of the Provisioning Services Configuration Wizard. The title bar reads 'Provisioning Services Configuration Wizard'. The main heading is 'Network Communications' with the instruction 'Specify networks settings.' Below this, there is a section titled 'Available network cards for stream services:' with a list of network cards, each with a checkbox. One card is selected. Below this, there is a text box for '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.' Below this is a note: 'Note: All servers must have the same port configurations.' There are two input fields: 'First communications port:' with a text box containing '6890' and 'Console port:' with a text box containing '54321'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

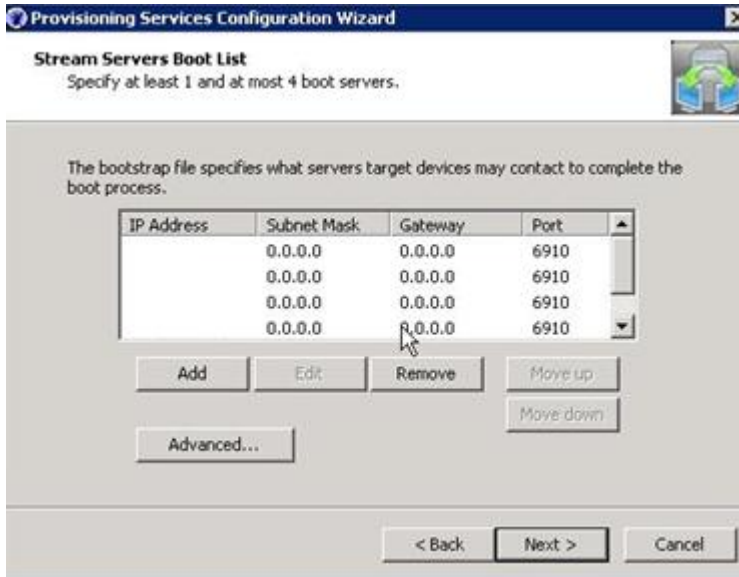
15. Select Use the Provisioning Services TFTP service.

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

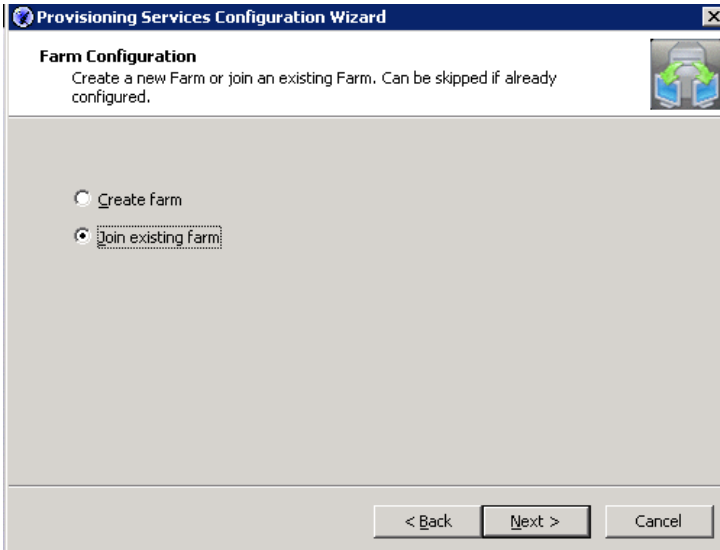
There are several recommended Hotfixes available for Provisioning Services 6.1. The Hotfixes that are listed below are specific to either the PVS server or the PVS target devices. The hotfixes applied to this environment are as follows:

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

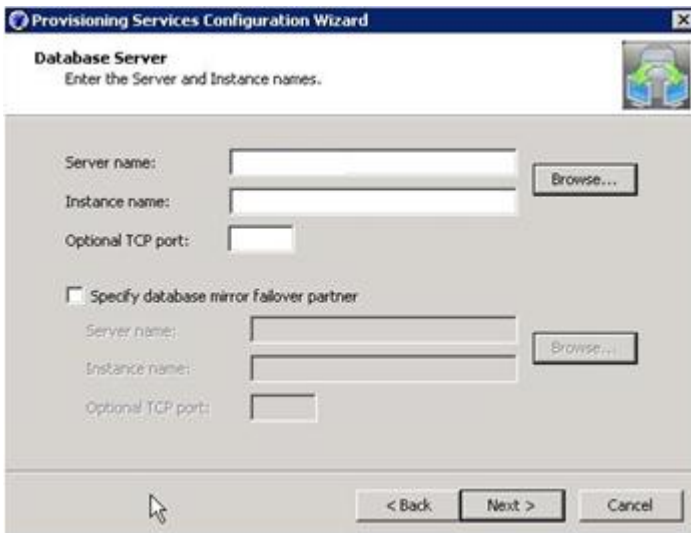
Each of the listed hotfixes has unique installation steps. Please refer to the Reference section of this document to view the specific installation steps for each hotfix.

6.8.6 Adding PVS Servers to the Farm

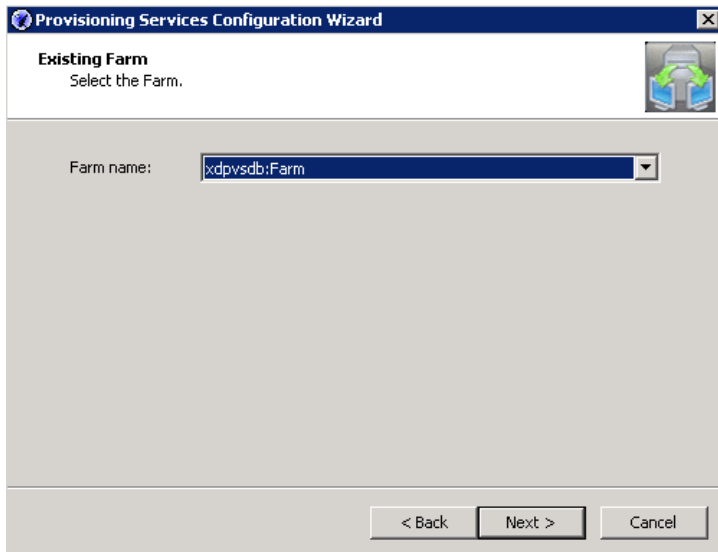
1. After installing the PVS 6.1 software on additional servers, launch the Provisioning Services Configuration Wizard.
2. Select Join existing Farm.
3. Click Next.



4. Enter Database information.



5. Click Next.
6. Select the existing PVS Farm from the drop-down box.



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

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

17. Click Next.
18. Configure Boot strap Boot list.
19. List the first four PVS servers in your farm.
20. Click Next.
21. 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 UCS 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: 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 redistrib\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 ensure 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

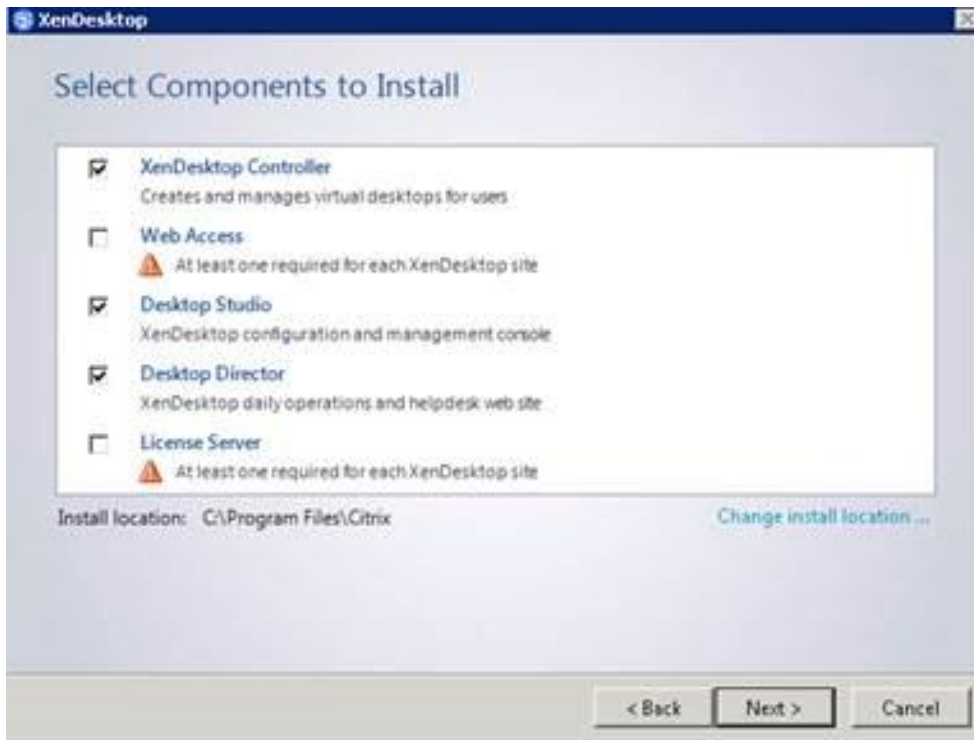
1. Start the XenDesktop installation wizard.
2. Click Install XenDesktop.



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.

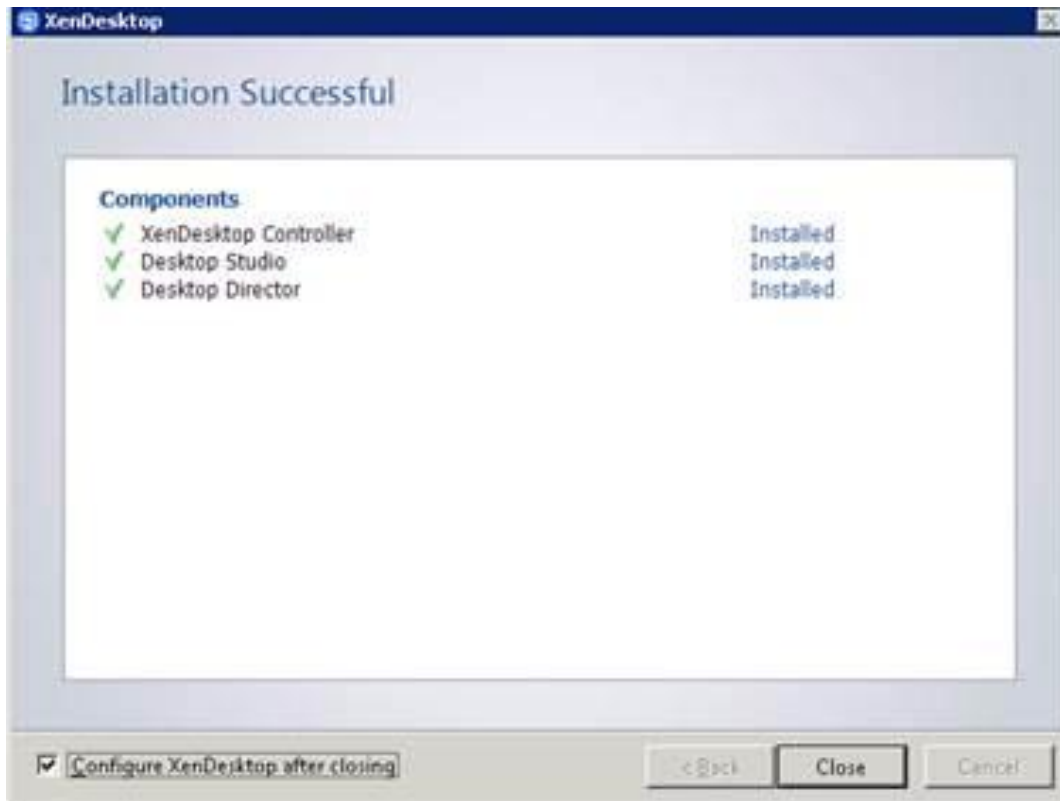
5. Verify that Install SQL Server Express is NOT selected.



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

6.9.3 Create SQL database for XenDesktop

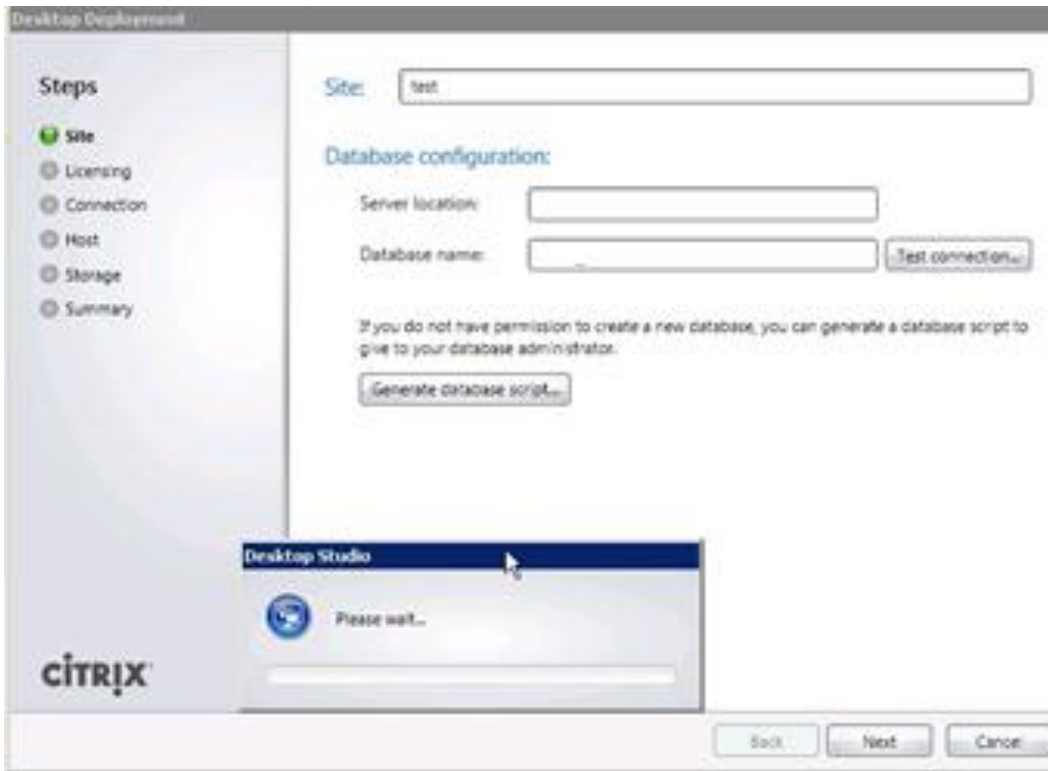
1. After successfully installing XenDesktop, select the checkbox Configure XenDesktop after Closing.
2. Click Close.



3. Open Desktop Studio: Start > All Programs > Citrix > Desktop Studio.
4. Click on 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 3rd 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

Back Next Cancel

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

Desktop Deployment

Steps

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

Host name

FS1_VDA1

Cluster

Select a cluster for the new virtual machines.

DC-VDA1 Browse...

Network

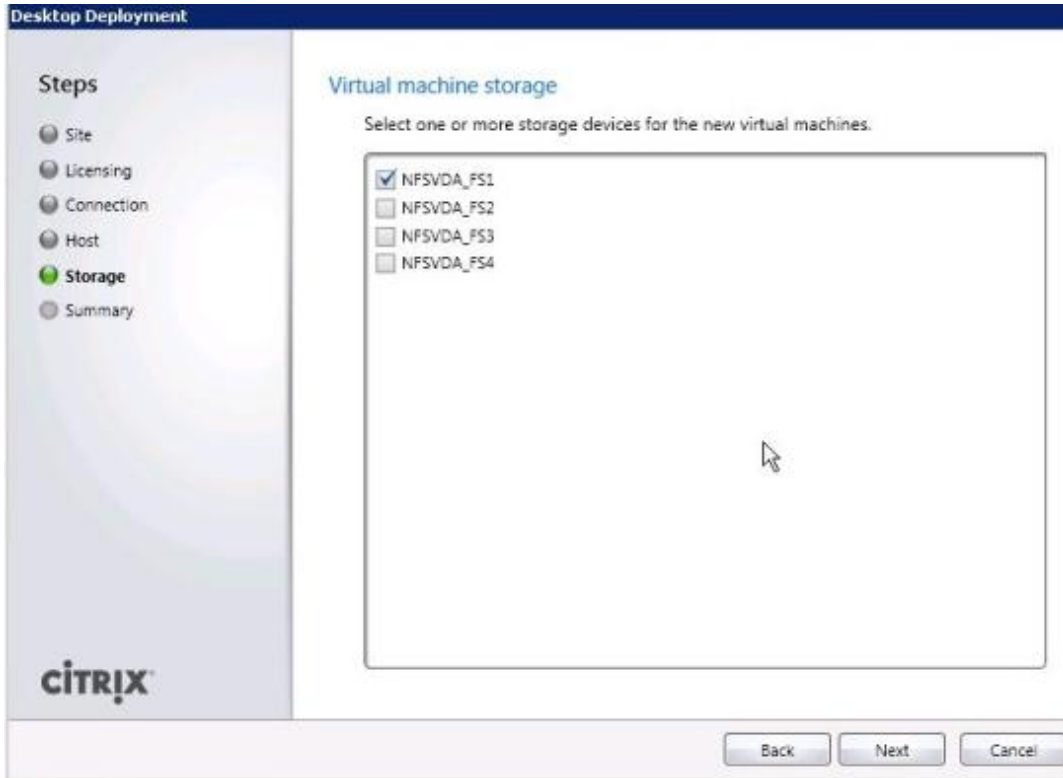
Select a network for the virtual machines to use.

- N1K3ctrn
- Storage
- UnusedOr_Quarantine_Veth
- VDA
- VDA1
- VM Network
- vMotion

Back Next Cancel

7. Click Next.

8. Select storage - this correlates to your vCenter's Datastore.



9. Click Next.
10. Click Finish.
11. Create additional hosts per datastore and network:
 - Right-click on your existing VCenter Storage Host connection.
 - Select Add Storage.
 - Select Use an Existing Host connection.

Add Host

Steps

- Connection
- Host
- Storage
- Summary

Use an existing Host Connection

dc-vc-01

Connect to a new Host

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:

Host type:

Address:

Username:

Password:

HA Servers: **None selected**

Virtual machines:

Use XenDesktop to create desktops

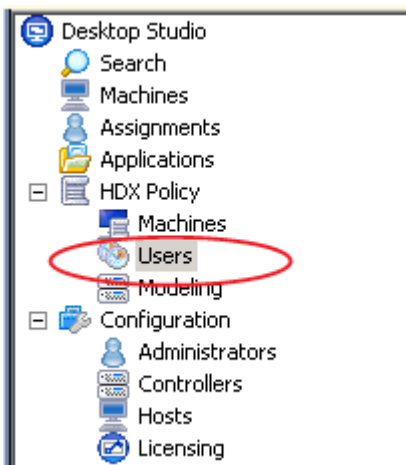
Manually create virtual machines

CITRIX

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.

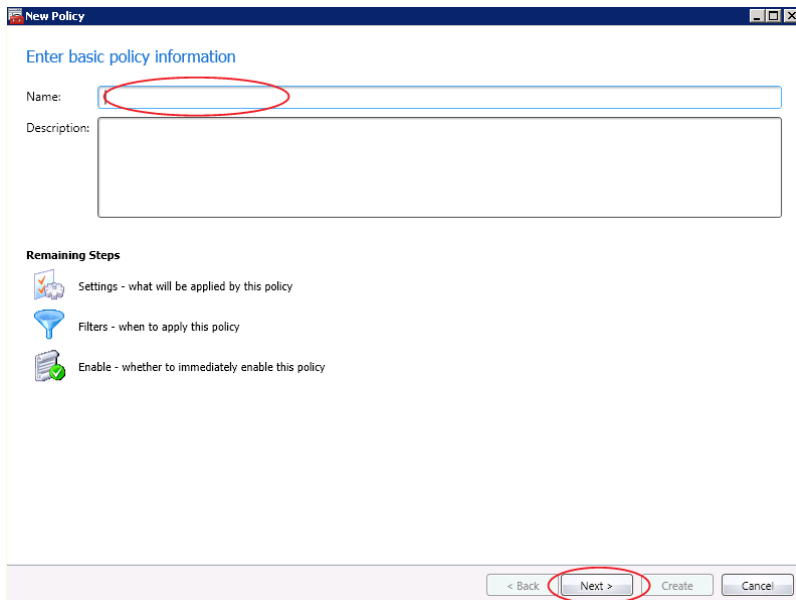




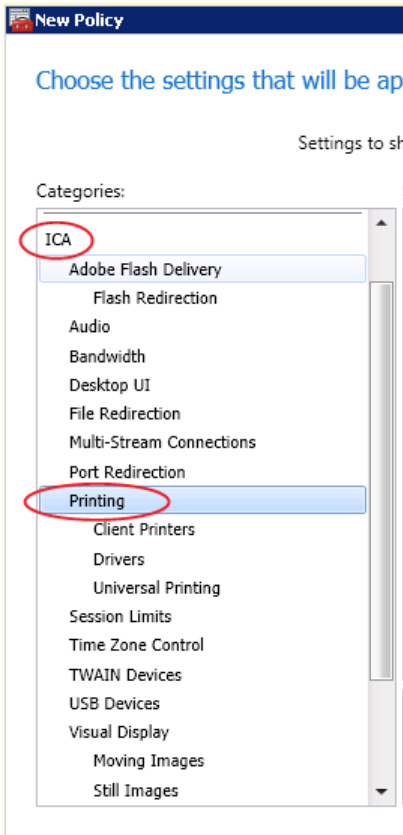
3. Click New to start the policy creation process.



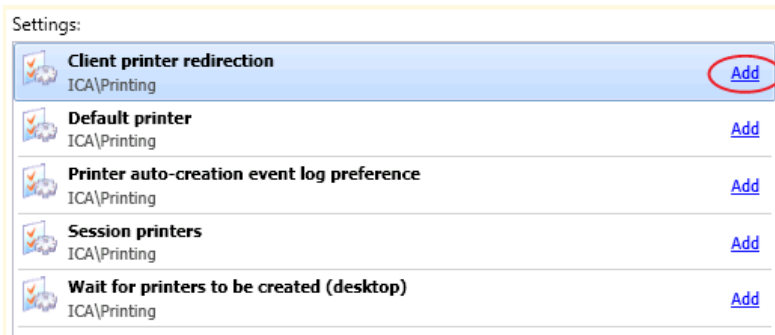
4. Enter a name for your policy and click Next.



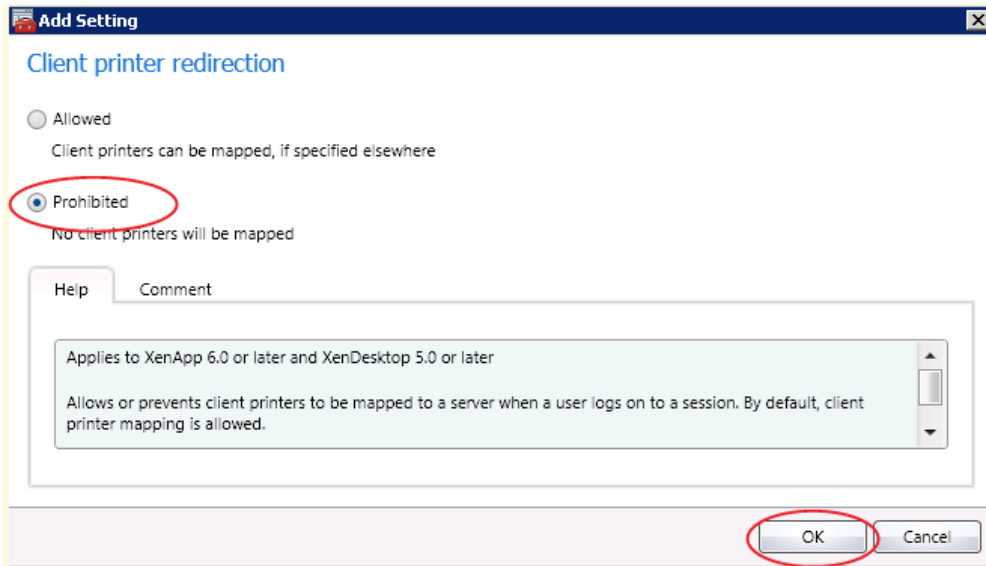
5. Select from Categories ICA → Printing.



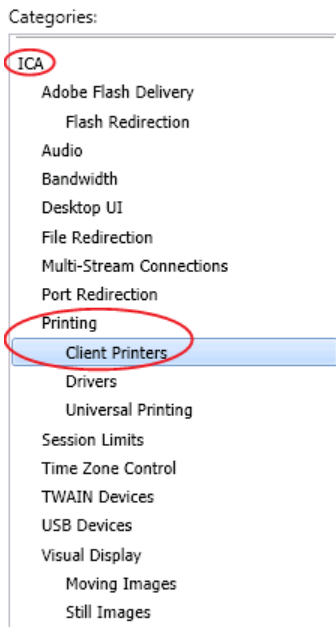
6. Select Client printer redirection and click Add.



7. Click Prohibited and then click OK.



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

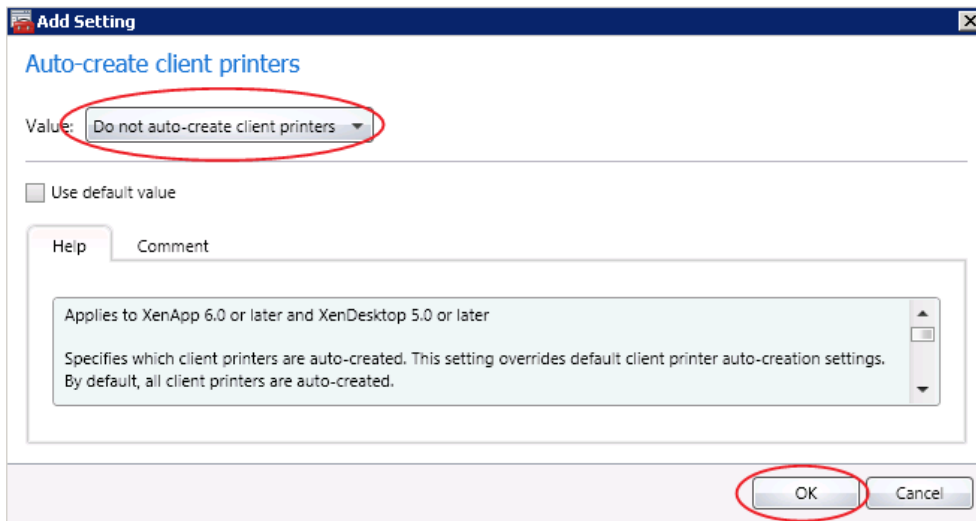


9. Select Auto-create client printers and click Add.

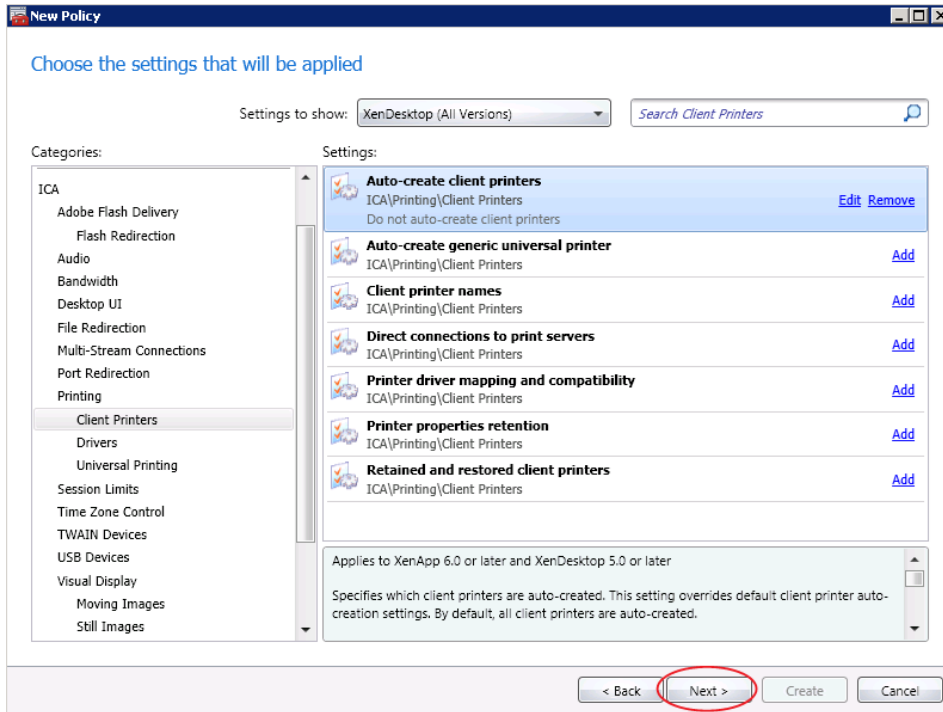
Settings:

Auto-create client printers ICA\Printing\Client Printers	Add
Auto-create generic universal printer ICA\Printing\Client Printers	Add
Client printer names ICA\Printing\Client Printers	Add
Direct connections to print servers ICA\Printing\Client Printers	Add
Printer driver mapping and compatibility ICA\Printing\Client Printers	Add
Printer properties retention ICA\Printing\Client Printers	Add
Retained and restored client printers ICA\Printing\Client Printers	Add

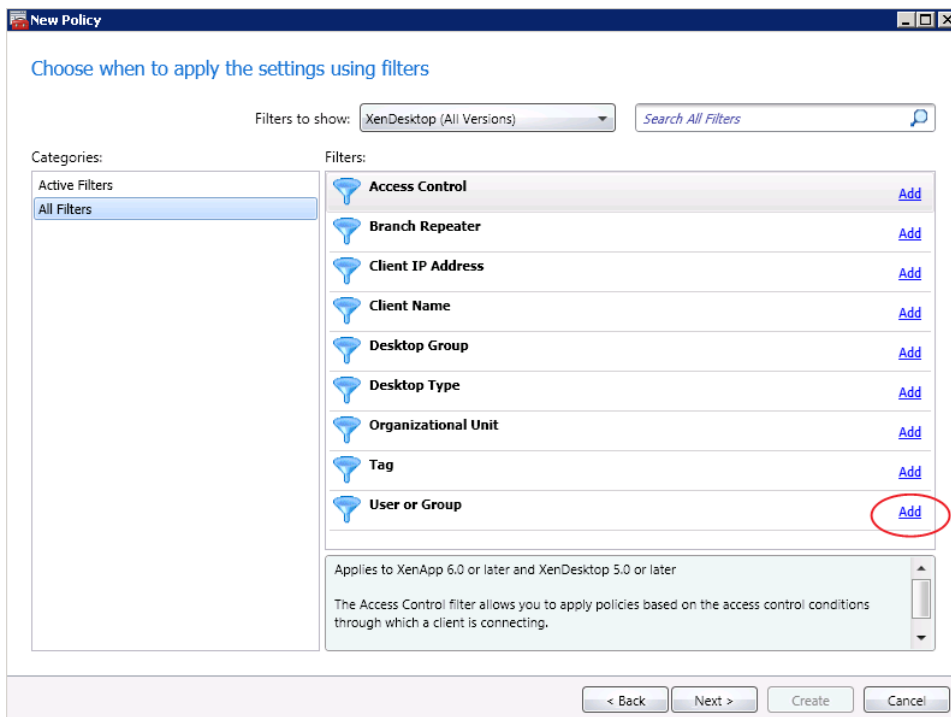
10. From the drop-down list, pick Do not auto-create client printers and click OK.



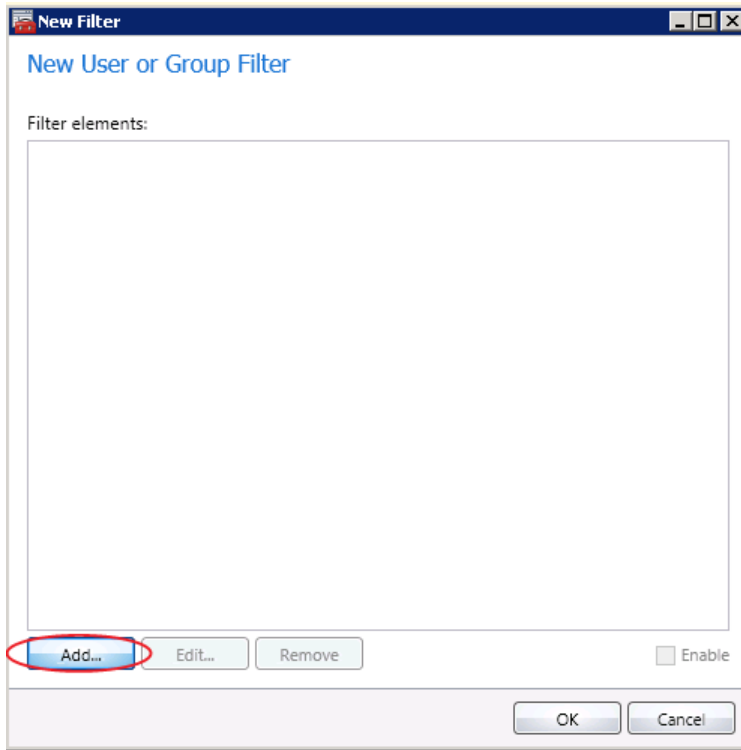
11. Click Next.



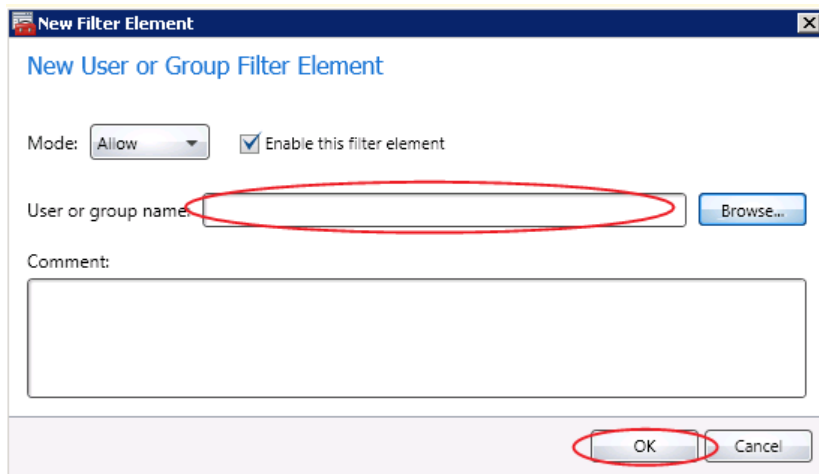
12. Select All Filters → User or Group and click Add Create policy filter.




13. Click Add.



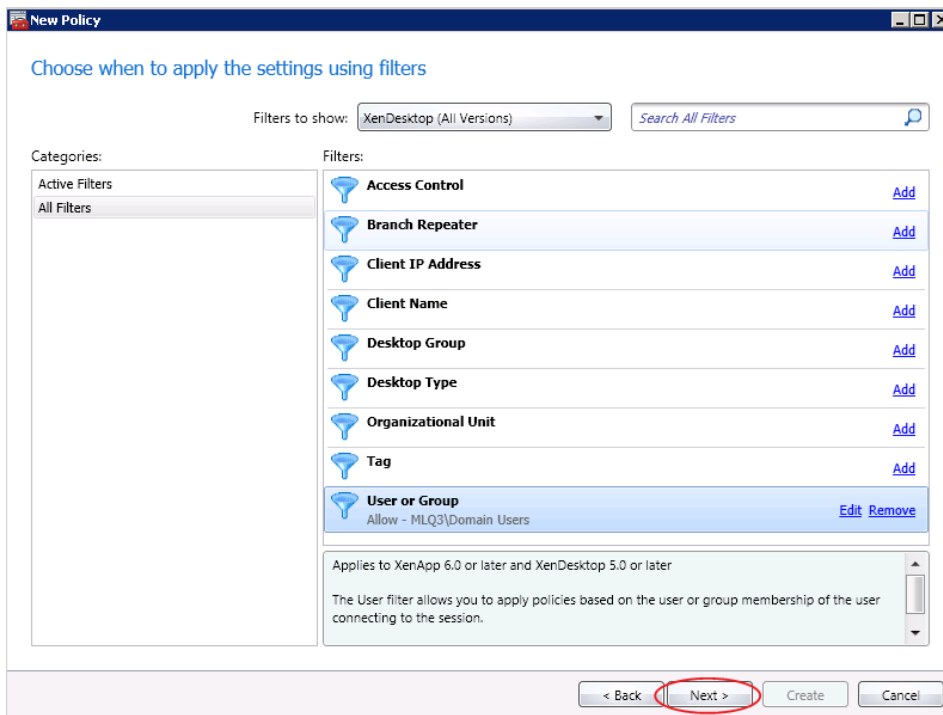
14. Select the User Group for this policy.



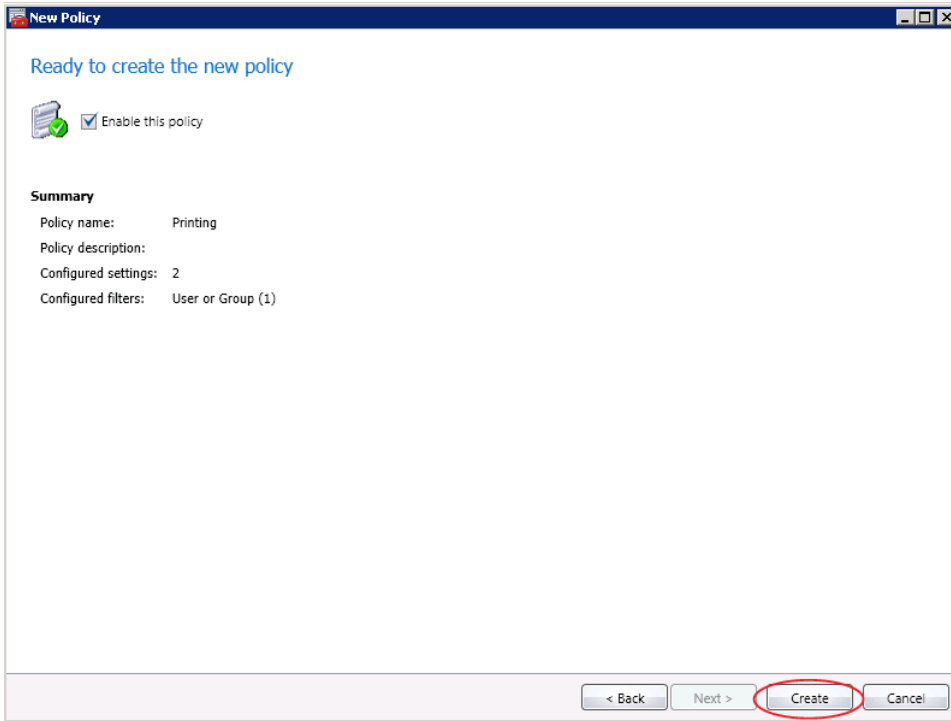
15. Click OK.



16. Click Next.




17. Click Create.



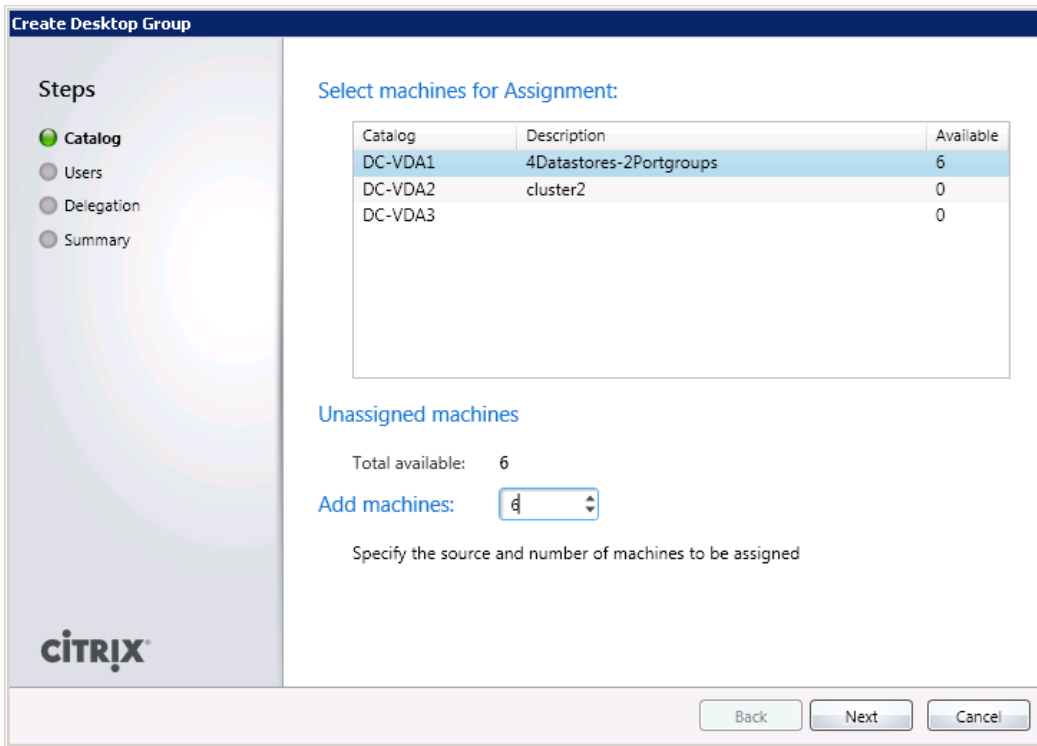
6.9.6 Configure the XenDesktop Desktop Group and Options

To configure the Citrix XenDesktop, do the following:

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.

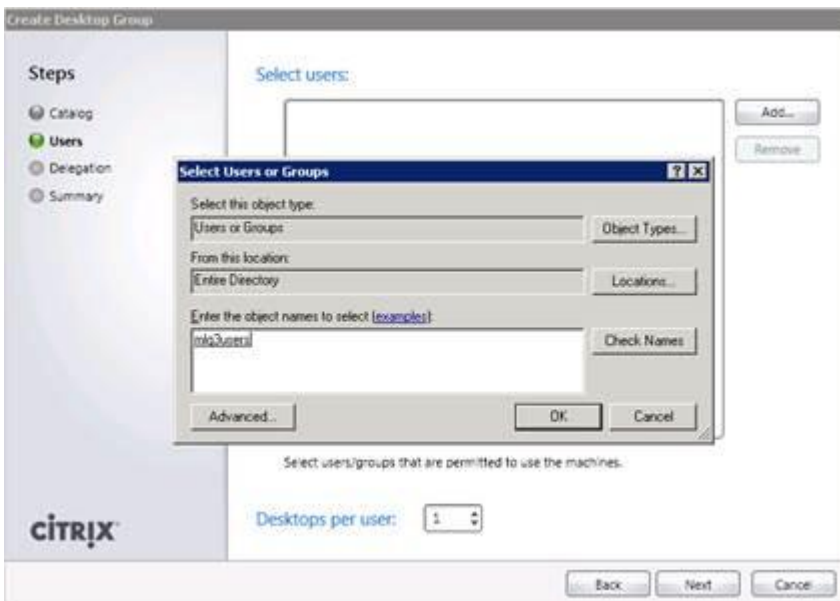


4. Select a catalog with available machines.
5. Add machines from the Catalog.



6. Click Next.

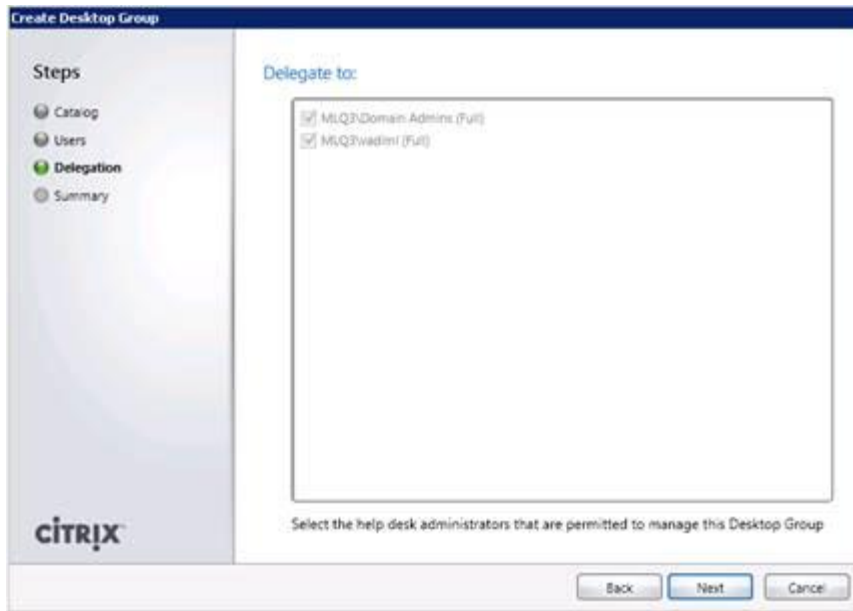
7. Click Add and select a user or user group for the Desktop Group.



8. Select one desktop per user.

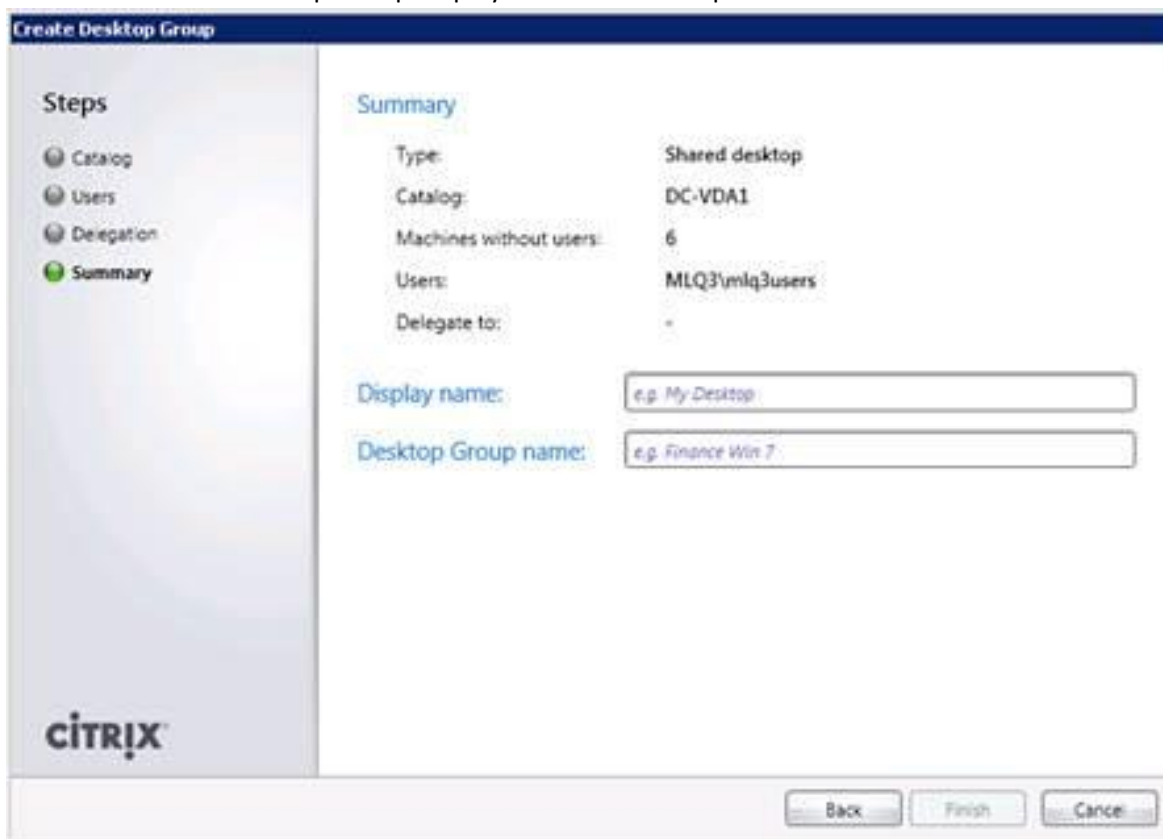
9. Click Next.

10. Select Administrators delegated to manage the Desktop Group.



11. Click Next.

12. Enter a Desktop Group Display Name and Description.



13. Click Finish.



6.10 Installing and Configuring Citrix XenApp 6.5

6.10.1 Pre-requisites

During a wizard-based installation, the XenApp Server Role Manager (using the Server Role Installer) automatically installs XenApp prerequisites.

For command-line installations, you must install the prerequisite software and Windows roles before installing XenApp (except as noted). You can deploy prerequisites with PowerShell cmdlets, the Microsoft ServerManagerCmd.exe command, or the Microsoft Deployment Image Servicing and Management (DISM) tool.

If installation of a required Windows role or other software requires a restart (reboot), restart the server before starting the XenApp server role installation.

XenApp Server Role

Supported operating systems: Windows Server 2008 R2 and Windows Server 2008 R2 SP1 (Enterprise, Standard, Datacenter, and Foundation).

Most servers running the supported operating systems meet the hardware requirements for XenApp with ample processing power to host user sessions accessing the published resources. However, additional research may be needed to determine if current hardware meets the requirements.

- CPU:
 - 64-bit architecture with Intel Pentium
 - Xeon family with Intel Extended Memory 64 Technology
 - AMD Opteron family
 - AMD Athlon 64 family
 - Compatible processor
- Memory: 512MB RAM (minimum)
- Disk space: up to 3.2GB

The XenApp Server Role Manager deploys the following software (except as noted), if it is not already installed:

- .NET Framework 3.5 SP1 (this is a prerequisite for the XenApp Server Role Manager; it is deployed automatically when you choose to add the XenApp server role from the Autorun menu)
- Windows Server Remote Desktop Services role (if you do not have this prerequisite installed, the Server Role Manager installs it and enables the RDP client connection option; you will be asked to restart the server and resume the installation when you log on again)
- Windows Application Server role
- Microsoft Visual C++ 2005 SP1 Redistributable (x64)
- Microsoft Visual C++ 2008 SP1 Redistributable (x64)

When you install the XenApp server role, XML and Internet Integration Service (IIS) integration is an optional component. When this component is installed, the Citrix XML Service and IIS share a port (default = 80). When this component is not installed, the Citrix XML Service defaults to standalone mode with its own port settings. You can change the port during or after XenApp configuration. The Server Role Installer checks for installed IIS role services and whether the component is selected or specified. For complete information, see [Before Installing XenApp](#). The IIS role services are listed below.



- Web Server (IIS) > Common HTTP Features > Default Document (selecting this automatically selects Web Server (IIS) > Management Tools > Management Console, which is not required or checked for XenApp installation)
- Web Server (IIS) > Application Development > ASP.NET (selecting this automatically selects Web Server (IIS) > Application Development > .NET Extensibility; although not checked for XenApp installation, ASP.NET requires .NET Extensibility)
- Web Server (IIS) > Application Development > ISAPI Extensions
- Web Server (IIS) > Application Development > ISAPI Filters
- Web Server (IIS) > Security > Windows Authentication
- Web Server (IIS) > Security > Request Filtering
- Web Server (IIS) > Management Tools > IIS 6 Management Compatibility (includes IIS 6 Metabase Compatibility, IIS 6 WMI Compatibility, IIS 6 Scripting Tools, and IIS 6 Management Console)

If you plan to use Philips SpeechMike devices with XenApp, you may need to install drivers on the servers hosting sessions that record audio before installing XenApp. For more information, see Citrix information on the Philips web site.

AppCenter

XenApp Management includes the AppCenter. By default, the AppCenter is installed on the same server where you install the XenApp server role; however, you can install and run the AppCenter on a separate computer. To install the AppCenter on a workstation, from the XenApp Autorun menu, select **Manually Install Components > Common Components > Management Consoles**.

Supported operating systems:

- Windows Server 2008 R2, 64-bit edition, SP1
- Windows Server 2008 R2, 64-bit edition
- Windows Server 2008 Enterprise, 32-bit edition, SP2
- Windows Server 2003 R2, 32-bit and 64-bit editions
- Windows Server 2003, 32-bit and 64-bit editions, SP2
- Windows 7 Enterprise, 32-bit and 64-bit editions, SP1
- Windows Vista Enterprise, 32-bit and 64-bit editions, SP2
- Windows XP Professional, 32-bit edition, SP3
- Windows XP Professional, 64-bit edition, SP2

Requirements:

- Disk space: 25MB
- Microsoft Management Console (MMC):
 - For Windows Vista, Windows 7, Windows Server 2008 R2, and Windows Server 2008 R2 SP1: MMC 3.0 (installed by default)
 - For other supported Windows operating systems: MMC 2.0 or 3.0

The XenApp Server Role Manager deploys the following software, if it is not already installed:

- Microsoft .NET Framework 3.5 SP1
- Microsoft Windows Installer (MSI) 3.0
- Microsoft Windows Group Policy Management Console



- Microsoft Visual C++ 2005 SP1 Redistributable (x64)
- Microsoft Visual C++ 2008 SP1 Redistributable (x64)
- Microsoft Visual C++ 2008 SP1 Redistributable
- Microsoft Visual C++ 2005 SP1 Redistributable
- Microsoft Primary Interoperability Assemblies 2005

If you install the AppCenter on a computer that previously contained the Microsoft Group Policy Management Console (GPMC) and a Citrix Delivery Services Console earlier than the version delivered with XenApp 6.0, you may also need to uninstall and reinstall the Citrix XenApp Group Policy Management Experience (x64) program in order to use the GPMC to configure Citrix policies.

Data Store Database

The following databases are supported for the XenApp data store:

- Microsoft SQL Server 2008 Express R2
- Microsoft SQL Server 2008 Express SP3
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008 SP2
- Microsoft SQL Server 2005 SP4
- Oracle 11g R2 32-bit Enterprise Edition

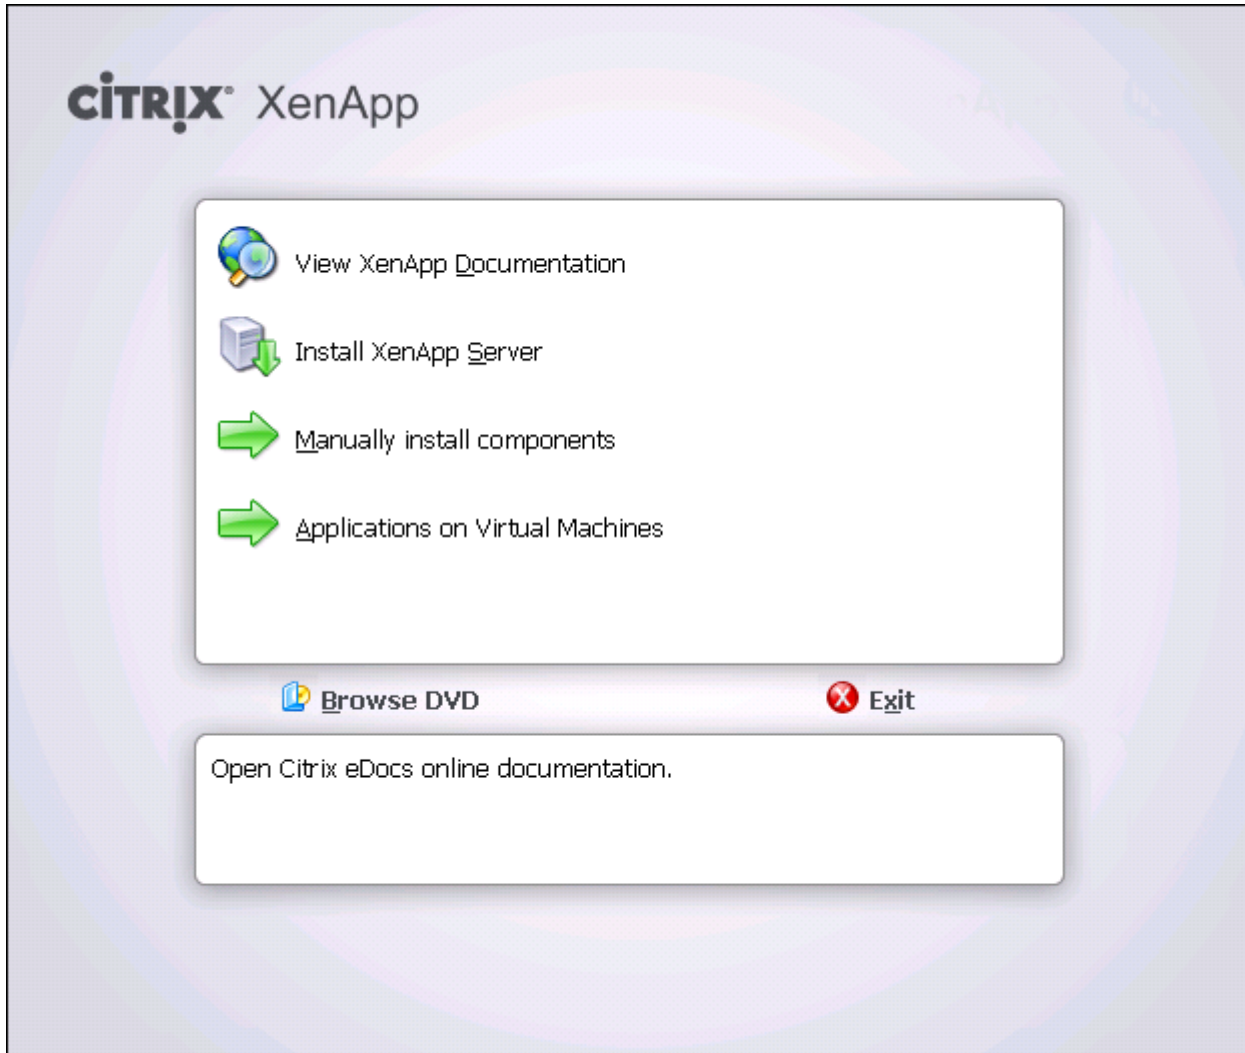
Microsoft SQL Server 2008 Express can be deployed for you by the XenApp Server Configuration Tool when creating a XenApp farm.

For information about the latest supported database versions, see [CTX114501](#). For information about requirements, see [Data Store Database Reference](#).


6.10.2 Install Citrix XenApp 6.5

To install Citrix XenApp 6.5, do the following:

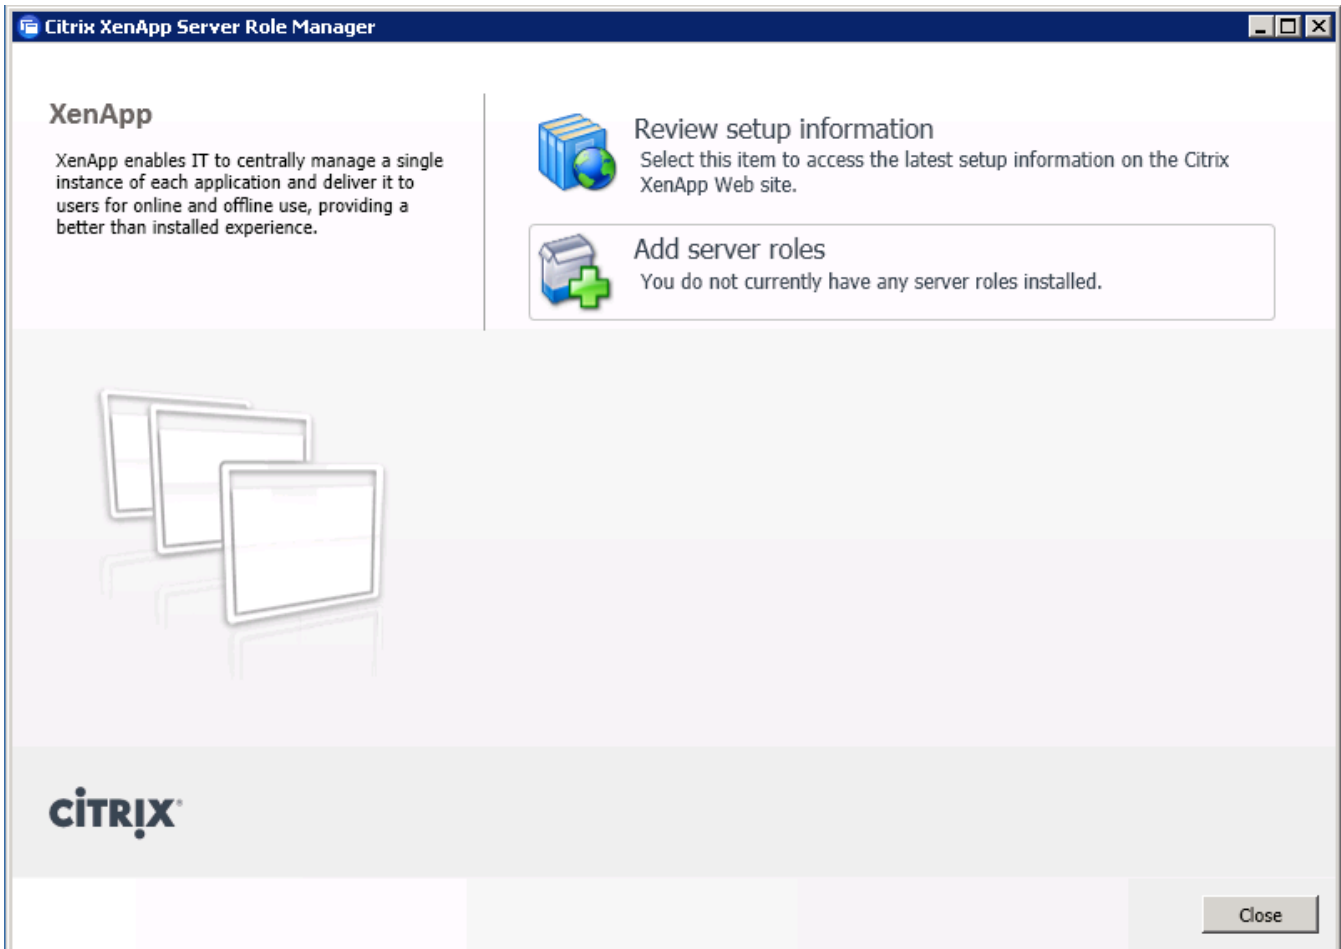
1. On the installation media, double-click autorun.exe. The Autorun menu launches.
2. Select Install XenApp Server. The Server Role Manager launches and checks if any roles are already installed.



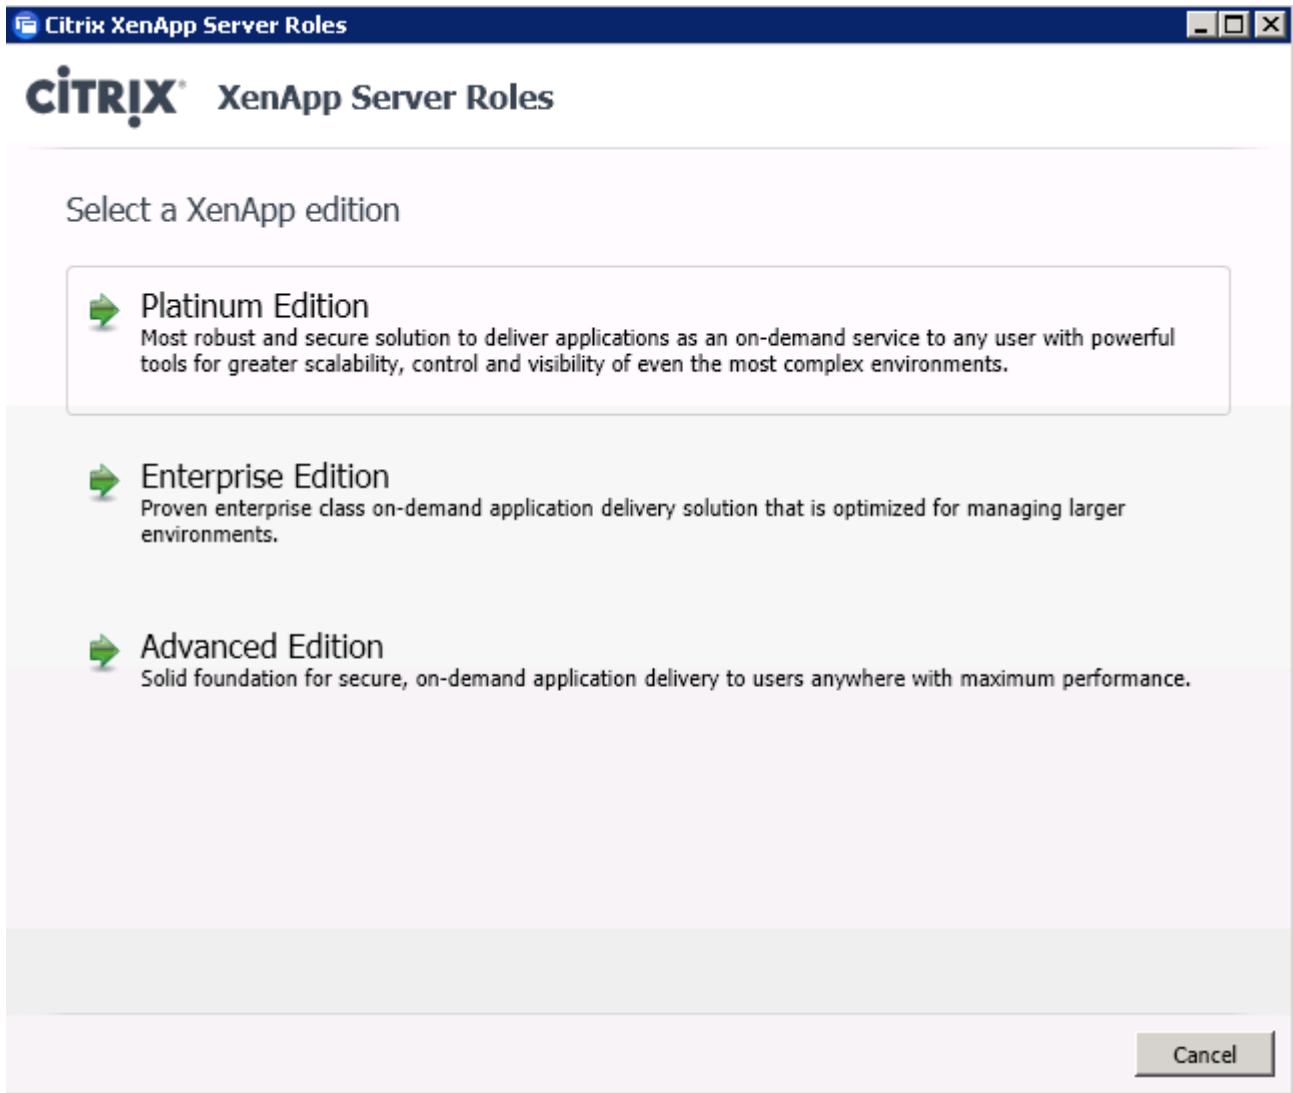
Note: If .Net hasn't been installed click OK to install .Net.



3. Select Add server roles.



4. Select your XenApp edition.



5. Accept the End User License Agreement.



The screenshot shows a window titled "Citrix XenApp Server Roles" with a blue header bar. Below the header, the Citrix logo and "XenApp Server Roles" are displayed. The main content area is titled "License agreement" and contains the following text:

You must accept the terms of the license agreement to continue. Use the scroll bar or Page Down key to read the entire agreement.

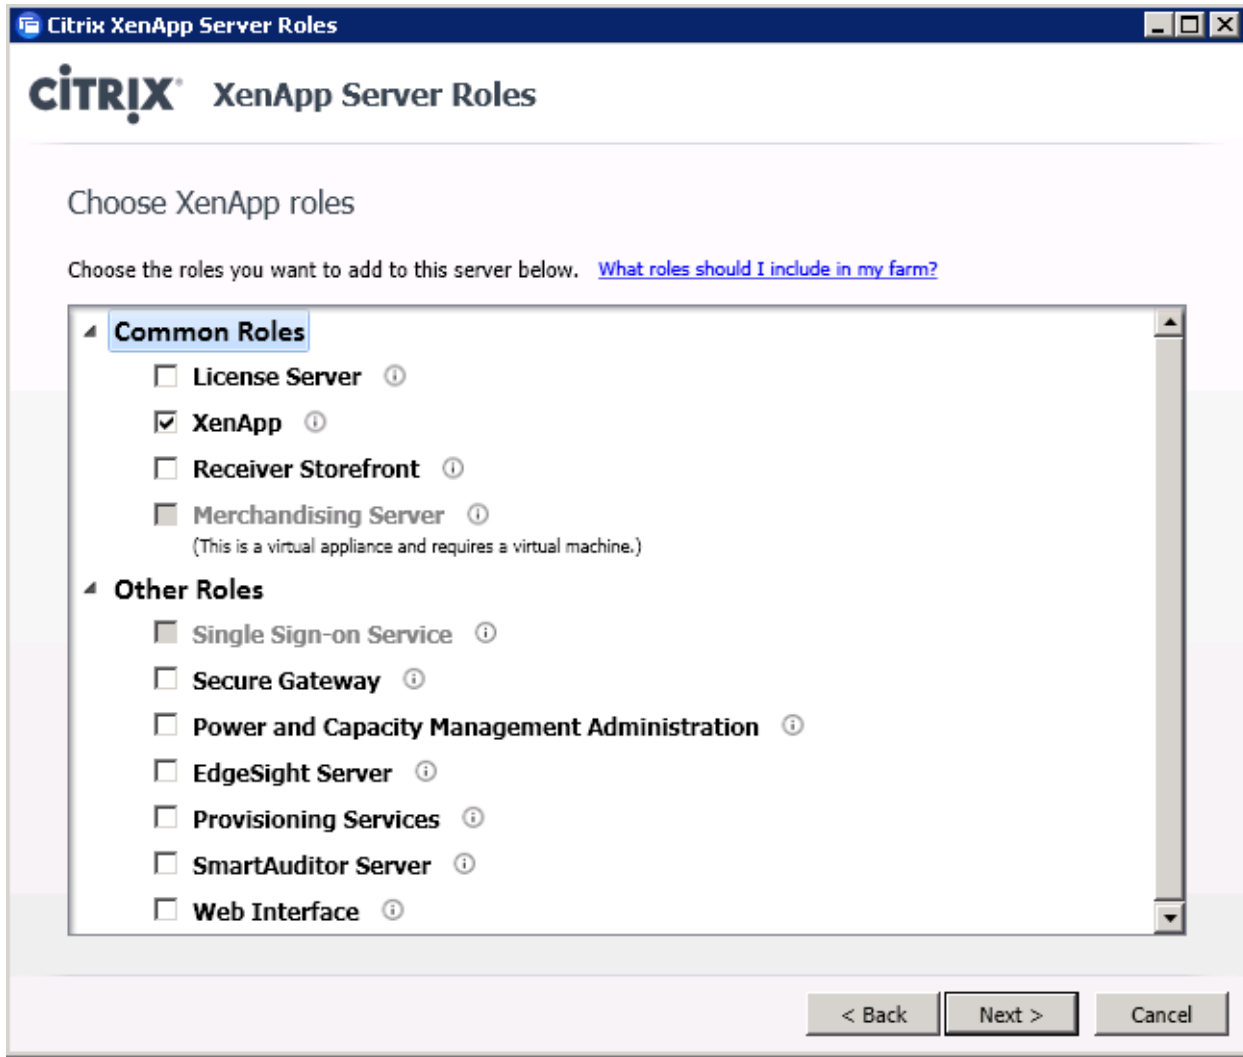
CITRIX® LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between you, the Licensed User, and Citrix Systems, Inc., Citrix Systems International GmbH, or Citrix Systems Asia Pacific Pty Ltd. Your location of receipt of this product or feature release (both hereinafter "PRODUCT") or technical support (hereinafter "SUPPORT") determines the providing entity hereunder (the applicable entity is hereinafter referred to as "CITRIX"). Citrix Systems, Inc., a Delaware corporation, licenses this PRODUCT in the Americas and Japan and provides SUPPORT in the Americas. Citrix Systems International GmbH, a Swiss company wholly owned by Citrix Systems, Inc., licenses this PRODUCT and provides SUPPORT in Europe, the Middle East, and Africa, and licenses the PRODUCT in Asia and the Pacific (excluding Japan). Citrix Systems Asia Pacific Pty Ltd. provides SUPPORT in Asia and the Pacific (excluding Japan). Citrix Systems Japan KK provides SUPPORT in Japan. BY INSTALLING AND/OR USING THE PRODUCT, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT.

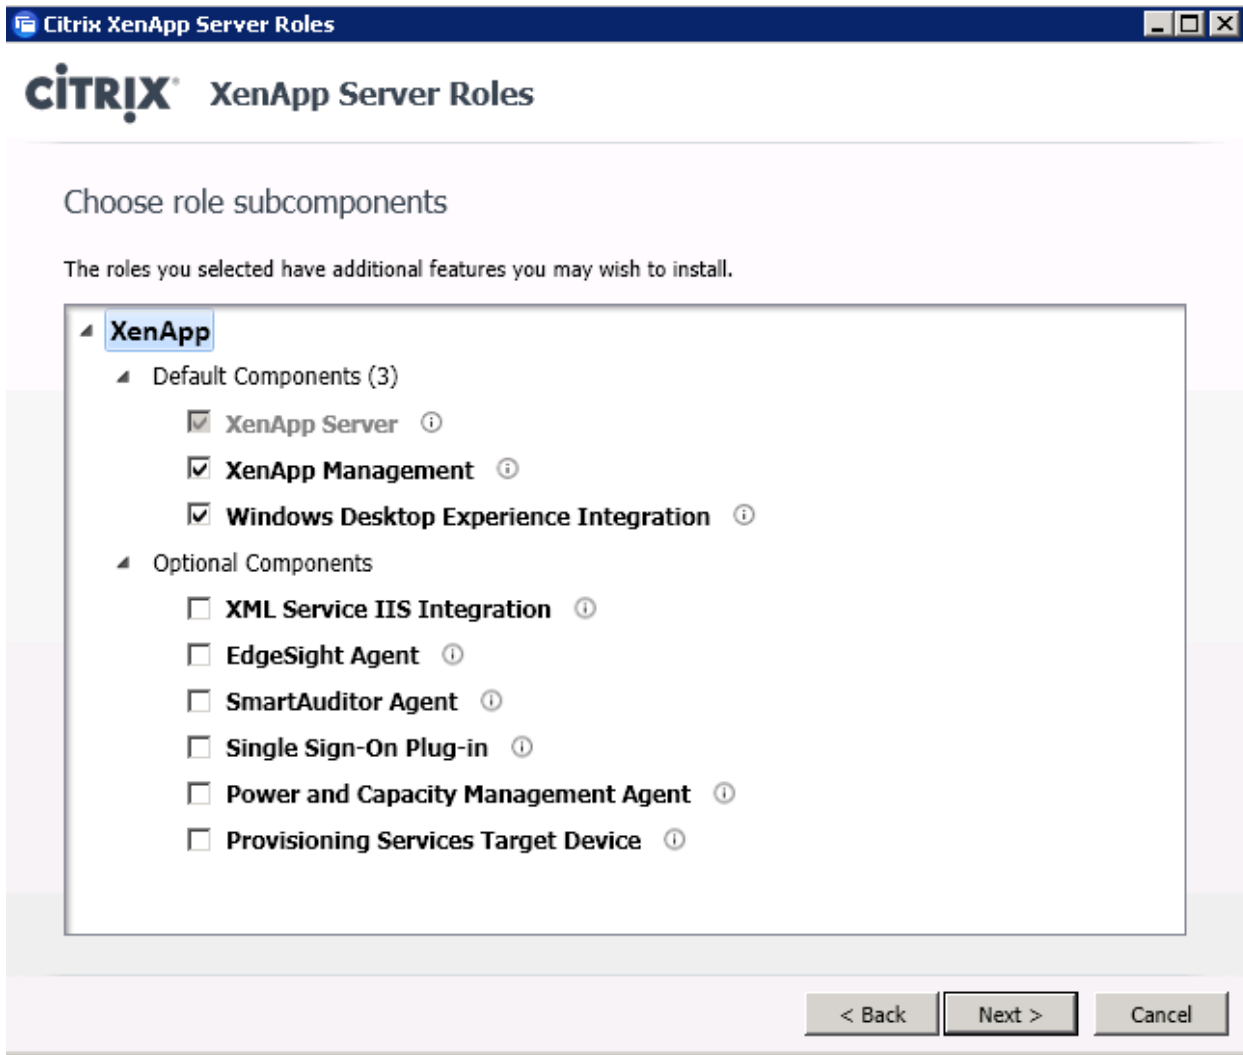
1. GRANT OF LICENSE. This PRODUCT contains software that provides services on a computer called a server ("Server Software") and contains software that allows a computer to

At the bottom of the window, there is a checkbox labeled "I accept the terms of this license agreement" which is checked. Below this are three buttons: "< Back", "Next >", and "Cancel".

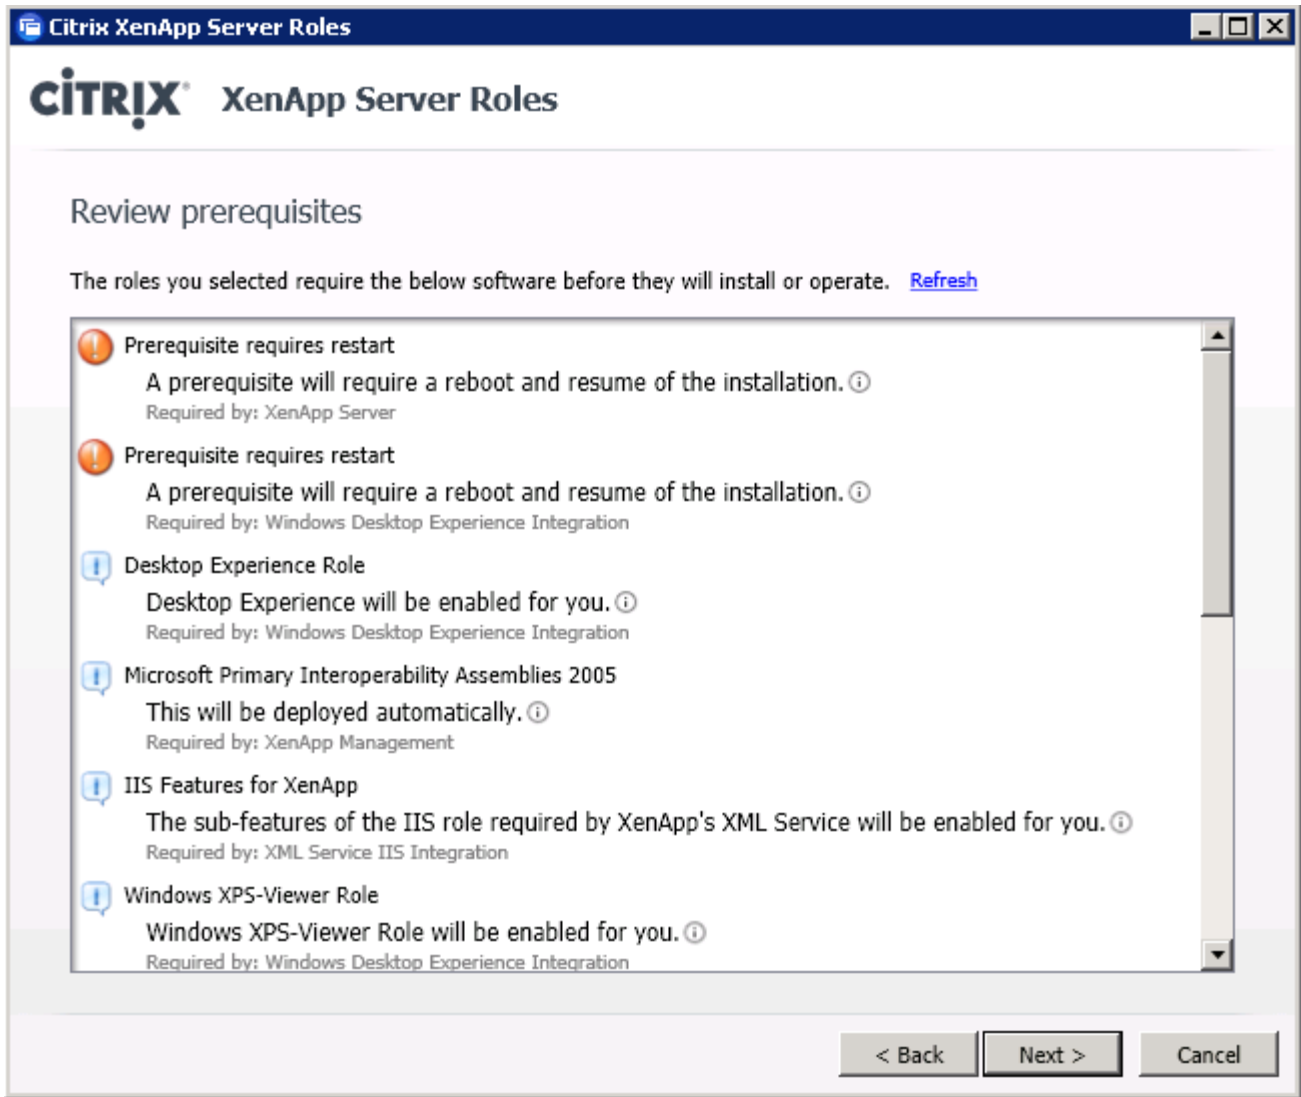
6. Select the roles you want to add. (For our environment we only used XenApp role).



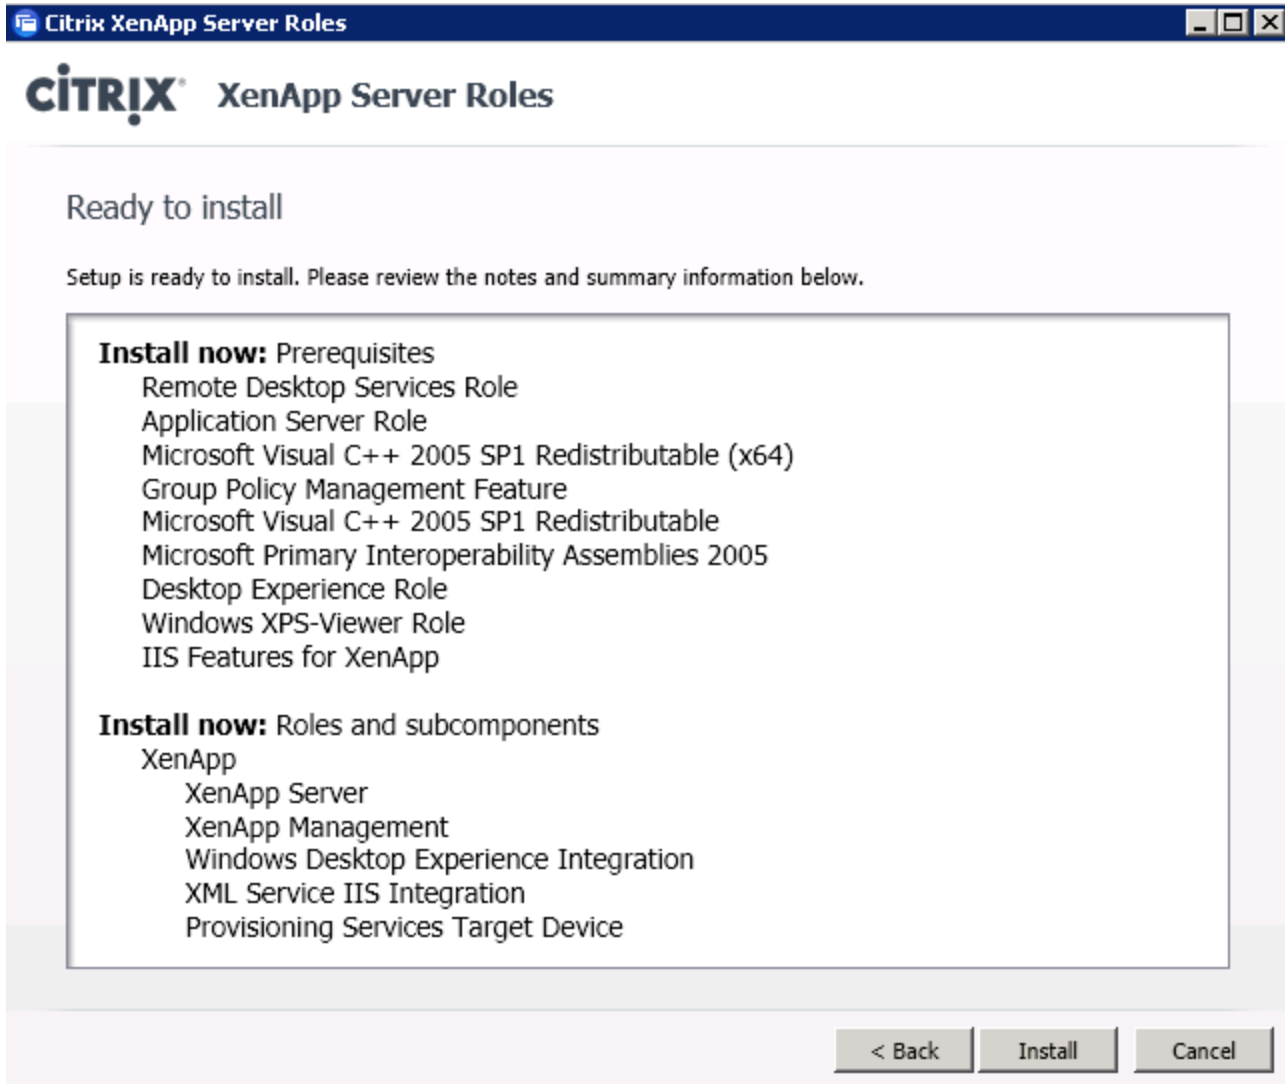
7. Select role components. Roles may have default and optional components.



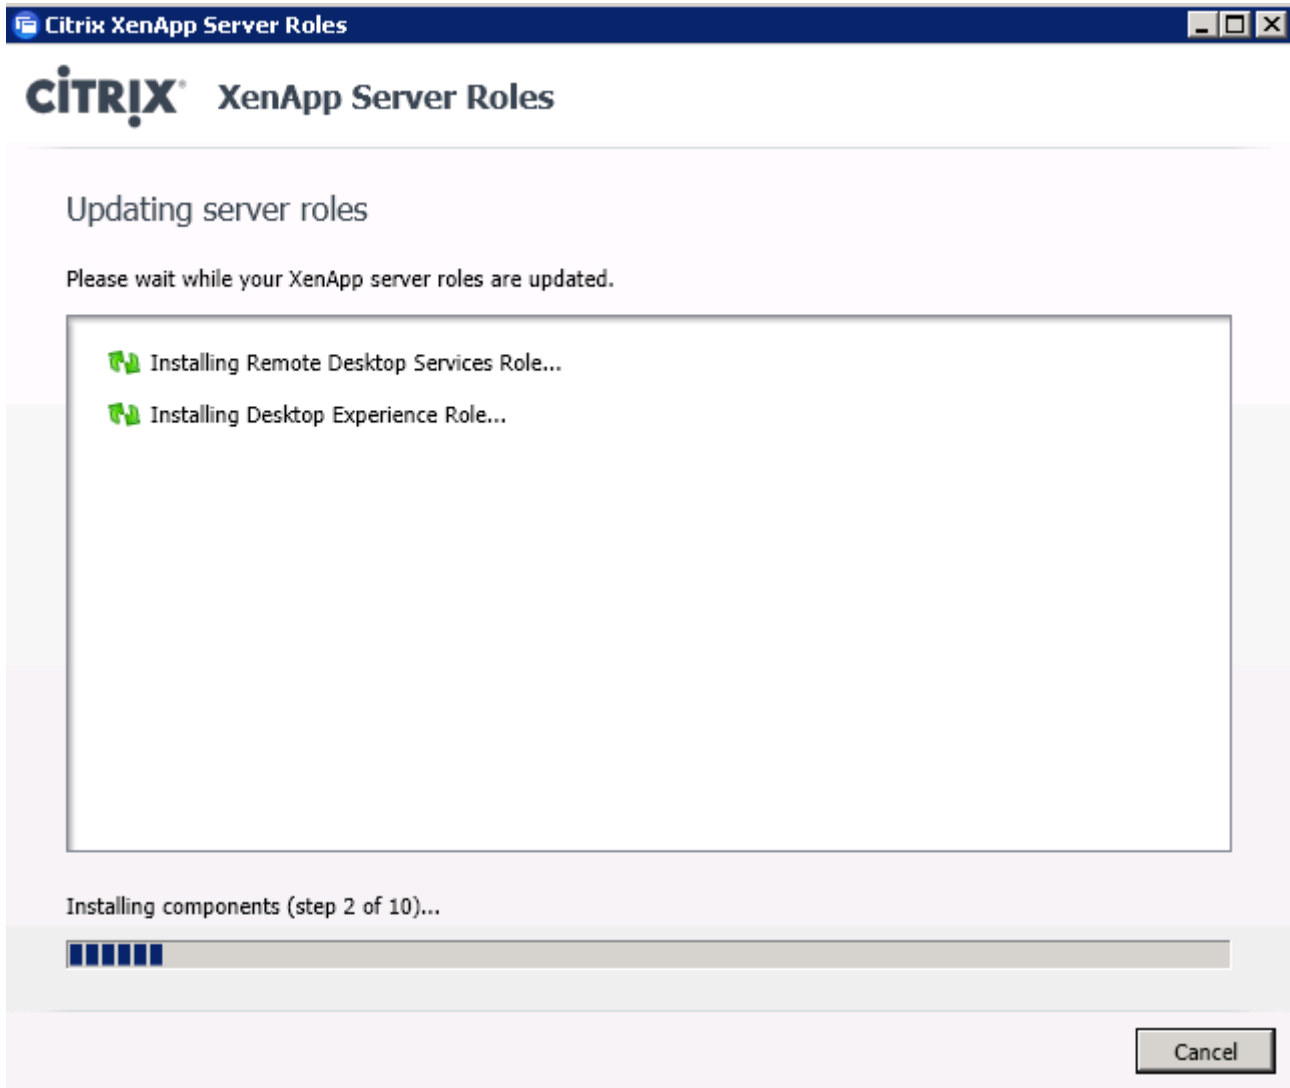
8. Review the prerequisites summary, which indicates which role or component needs the prerequisite, and whether the Server Role Installer installs it or you must install it. For software you must install, the display indicates whether the XenApp installation media contains the software or you must obtain it elsewhere.



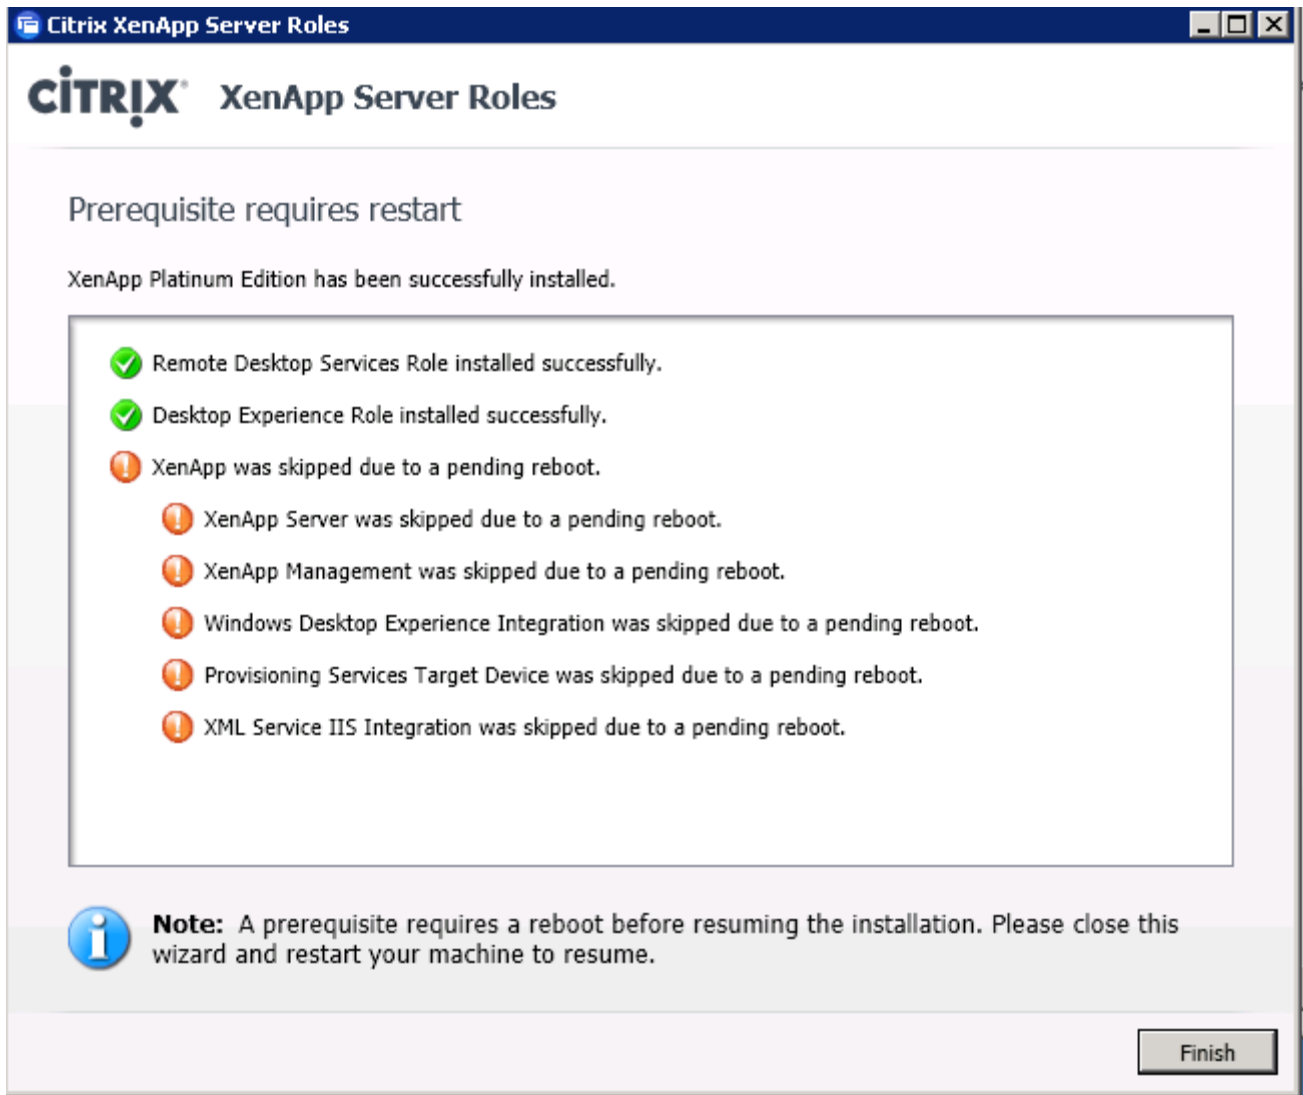
- Review the summary, which lists the selected roles and components to be installed or prepared. It also lists prerequisites which will be automatically deployed for all selected roles.



10. Click Install.



11. Click Finish.



The screenshot shows a Windows-style window titled "Citrix XenApp Server Roles". The window content includes the Citrix logo and the title "XenApp Server Roles". Below the title, the text reads "Prerequisite requires restart" and "XenApp Platinum Edition has been successfully installed." A list of installation results is shown in a box, with two green checkmarks for successful roles and several red exclamation marks for skipped components. A note at the bottom explains that a restart is required. A "Finish" button is located in the bottom right corner.

Citrix XenApp Server Roles

CITRIX XenApp Server Roles

Prerequisite requires restart

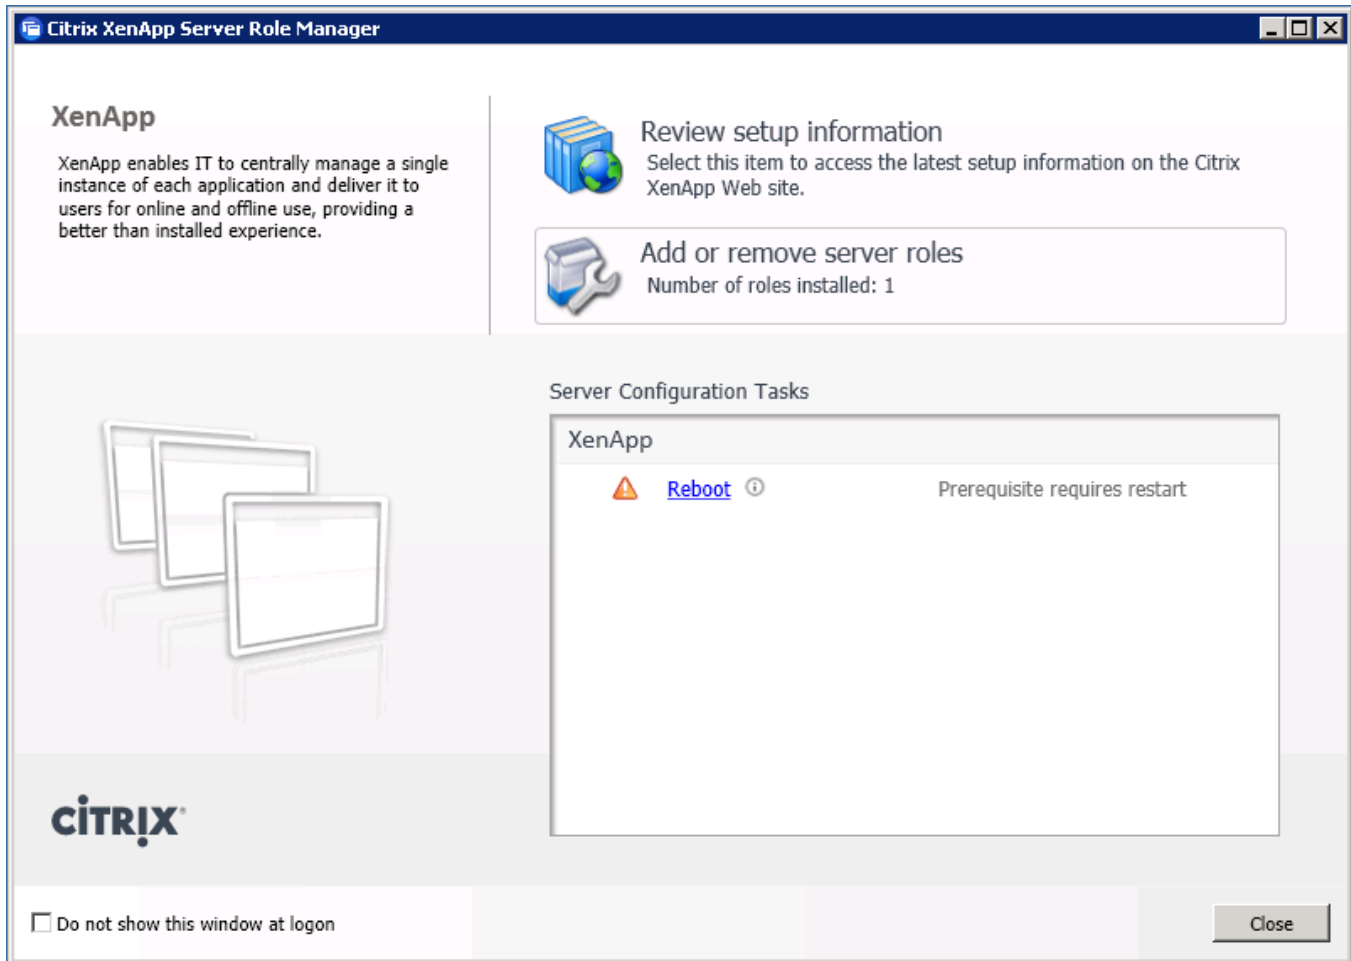
XenApp Platinum Edition has been successfully installed.

- ✓ Remote Desktop Services Role installed successfully.
- ✓ Desktop Experience Role installed successfully.
- ✗ XenApp was skipped due to a pending reboot.
 - ✗ XenApp Server was skipped due to a pending reboot.
 - ✗ XenApp Management was skipped due to a pending reboot.
 - ✗ Windows Desktop Experience Integration was skipped due to a pending reboot.
 - ✗ Provisioning Services Target Device was skipped due to a pending reboot.
 - ✗ XML Service IIS Integration was skipped due to a pending reboot.


Note: A prerequisite requires a reboot before resuming the installation. Please close this wizard and restart your machine to resume.

Finish

12. Reboot the server.



13. Click Yes to reboot.



14. Resume Install upon reboot.

Citrix XenApp Server Role Manager

XenApp

XenApp enables IT to centrally manage a single instance of each application and deliver it to users for online and offline use, providing a better than installed experience.

Review setup information
Select this item to access the latest setup information on the Citrix XenApp Web site.

Add or remove server roles
Number of roles installed: 1

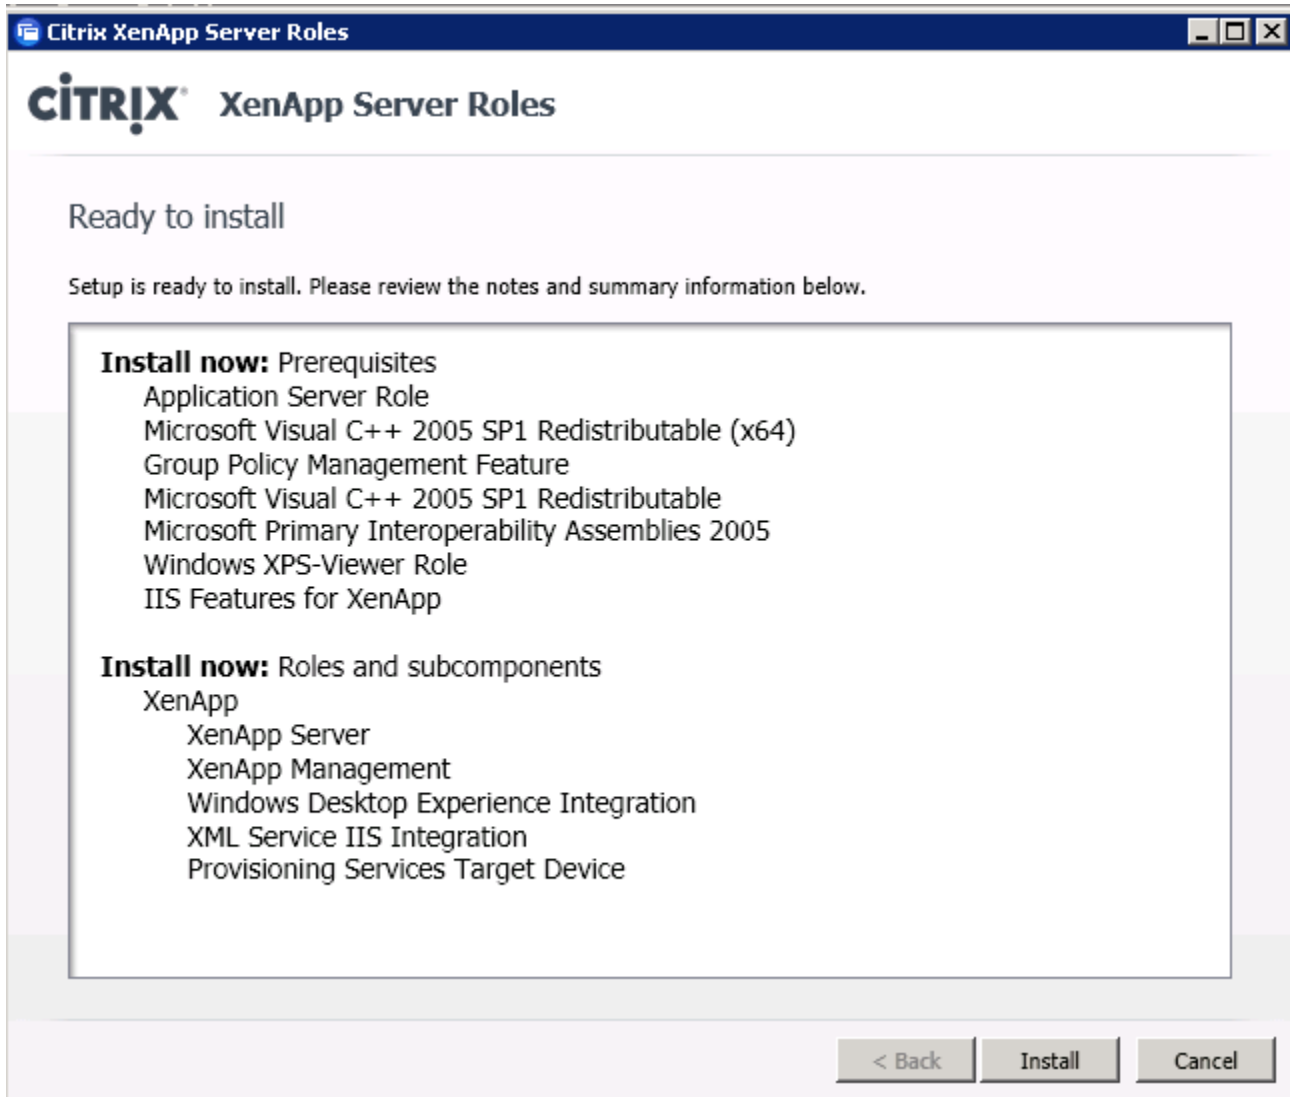
Server Configuration Tasks

XenApp	
Resume Install ⓘ	Not Installed

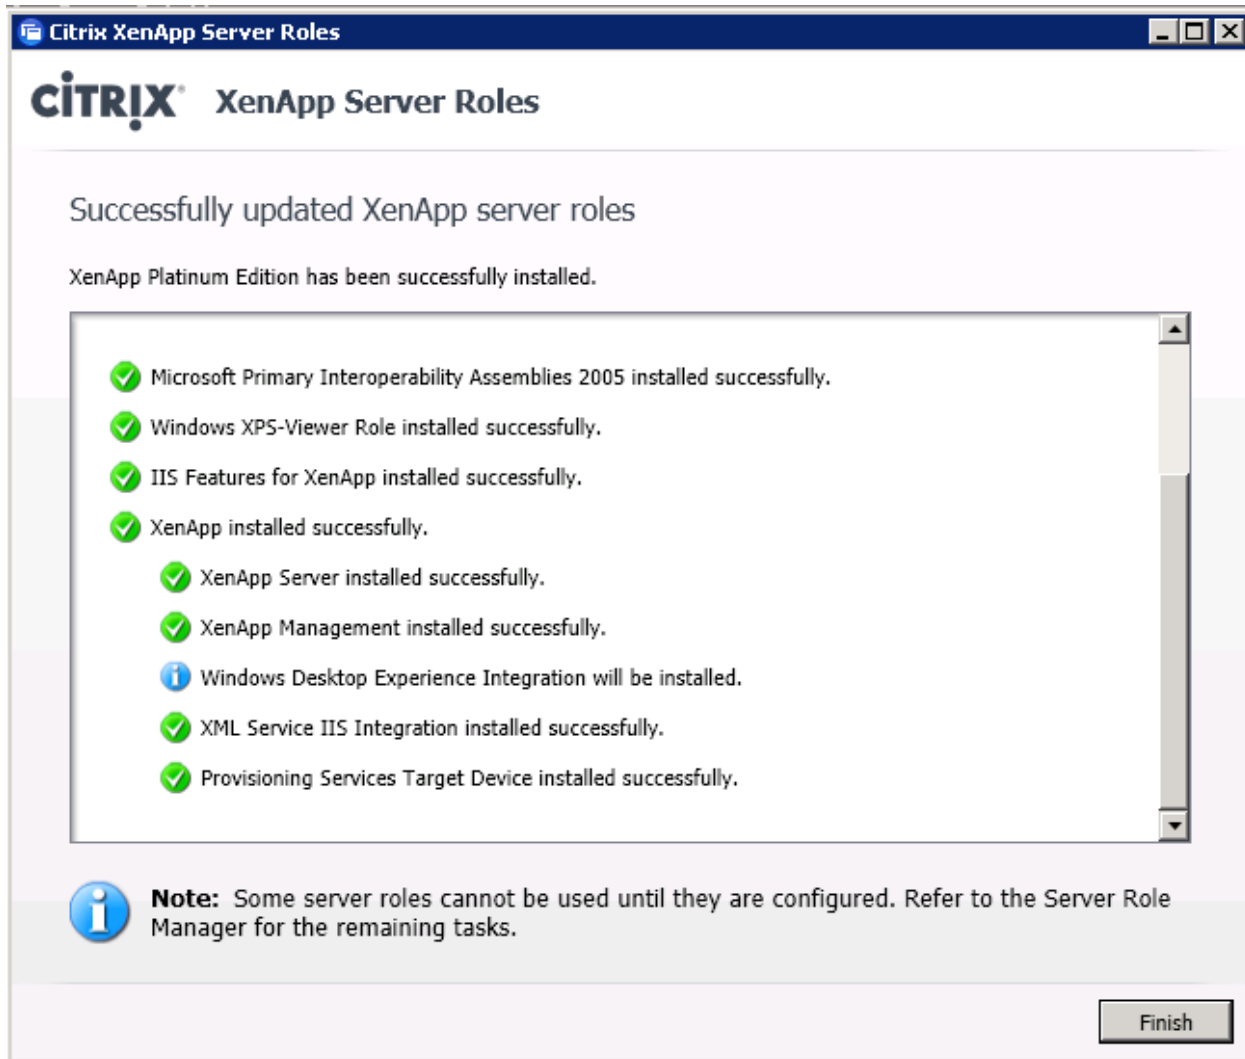
Do not show this window at logon

Close

15. Click Install.



16. Click Finish.



6.10.3 Configure XenApp License Information

1. Access the XenApp Server Role Manager.
2. Click Specify licensing. The Licensing Configuration Tool launches.
3. On the Enter License Server Information page, select one of the following:
 - Connect to existing license server. Specify the case-sensitive license server name. If you do not change the license server port value, the default value 27000 is used.
 - Click Test Connection to verify that the specified license server is running and using a compatible software version, and to check if the license server has any licenses.

4. On the Select Licensing Model page, you can select a licensing model option or defer the selection to a later time.

Note: Select the licensing model best suited to your planned deployment, which may differ from the recommendation, which is based on the licenses currently on the license server.

- **XenApp.** Select this model if you plan to use only XenApp licenses. This option is recommended if the Test Connection operation discovered no licenses, only XenApp licenses, or a mixture of unique XenApp and XenDesktop licenses on the license server.
- **XenDesktop concurrent system.** Select this model if you plan to use XenDesktop concurrent user licenses. This option is recommended if the Test Connection operation discovered only XenDesktop concurrent licenses on the license server.
- **XenDesktop user/device.** Select this model if you plan to use XenDesktop user or device licenses. This option is recommended if the Test Connection operation discovered XenDesktop user/device licenses or both XenDesktop user/device and XenDesktop concurrent licenses.

6.10.4 Configure XenApp Farm

1. Access the Server Role Manager.
2. Click Configure under XenApp. The Server Configuration Tool launches.

Citrix XenApp Server Role Manager

XenApp

XenApp enables IT to centrally manage a single instance of each application and deliver it to users for online and offline use, providing a better than installed experience.

Review setup information
Select this item to access the latest setup information on the Citrix XenApp Web site.

Add or remove server roles
Number of roles installed: 1

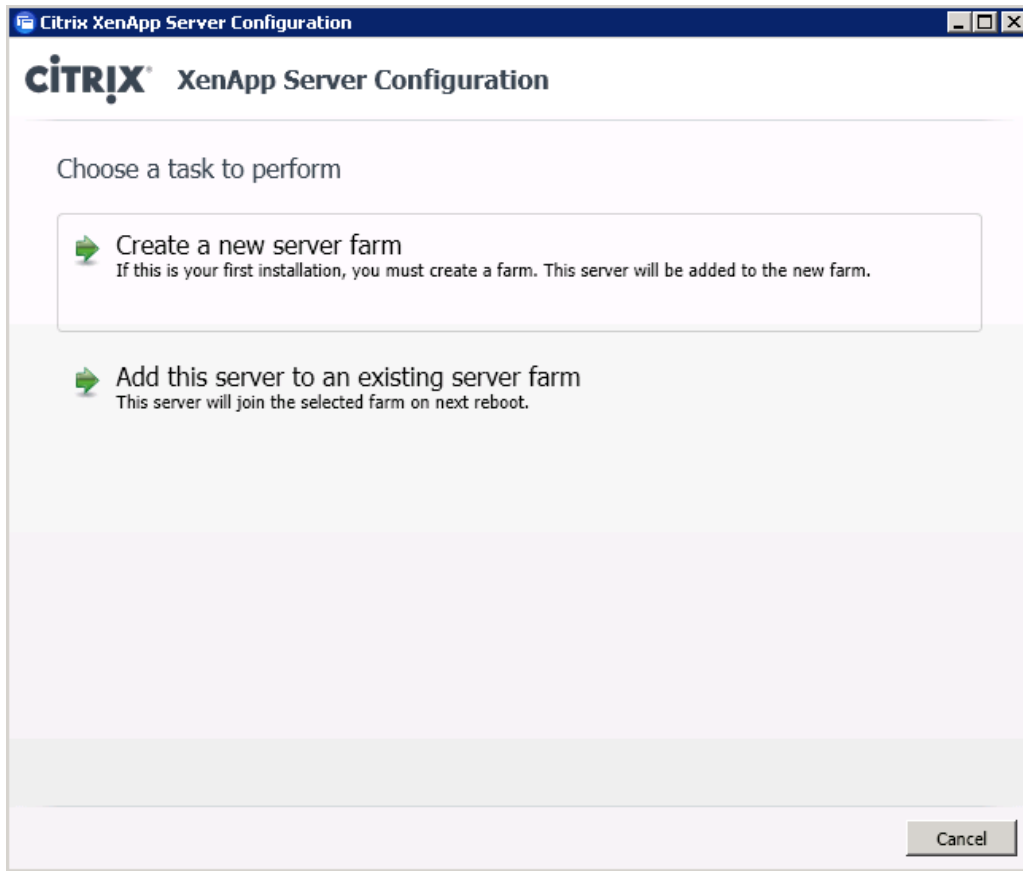
Server Configuration Tasks

XenApp		
✓	Licensing specified	Edit Licensing ⓘ
⚠	Configure ⓘ	Not configured

Do not show this window at logon

Close

3. Select Create a new server farm.



4. When creating a farm, on the Enter basic information page:

- Enter a farm name, up to 32 characters (can include spaces).
- Specify the domain and username for a user who will be the first Citrix administrator. The administrator has full permissions to the farm and can create additional administrator accounts.



Citrix XenApp Server Configuration

CITRIX XenApp Server Configuration

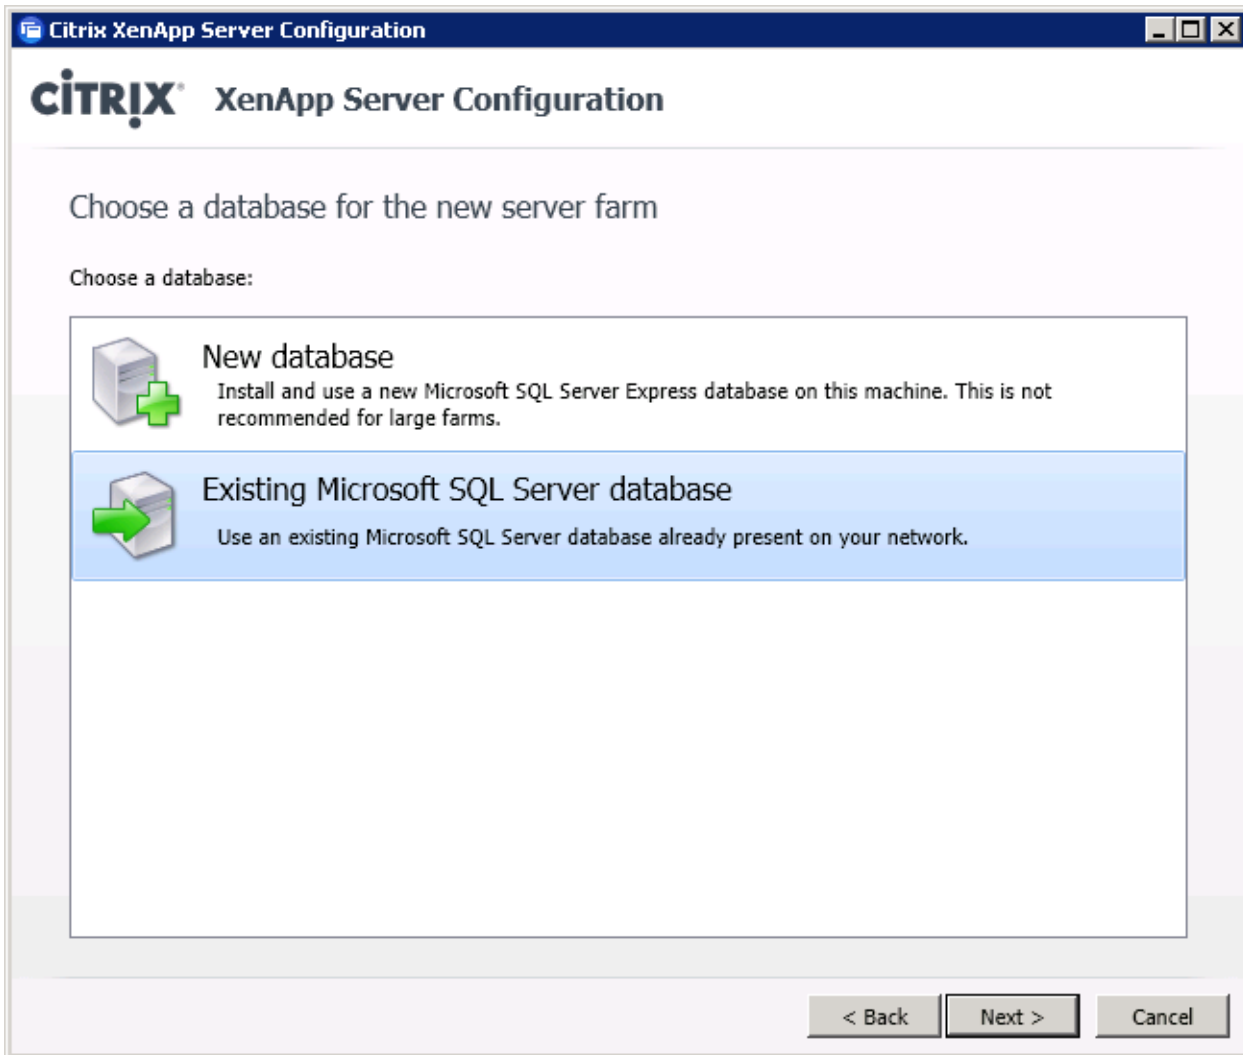
Enter basic information about the new server farm

New XenApp Server farm name:

First Citrix administrator account:

< Back Next > Cancel

5. Select New database.



6. Specify the database credentials. Specify the user name in the form <DBMACHINE>\<USER> or <DOMAIN>\<USER>.

Citrix XenApp Server Configuration

CITRIX XenApp Server Configuration

Enter connection information for the new server farm

Database server name:
XD-SQL-CL1
Examples: TestServerA or TestServerB\Sq\Express or ActServer\Accounting,1444 or 10.1.1.0 or 192.168.1.1,1433


Database name:
dcxadb

Integrated Windows authentication
Select this option to use your Windows identity to access the database.

SQL Server authentication using a login ID and password
Select this option to use a SQL login ID and password to access the database.

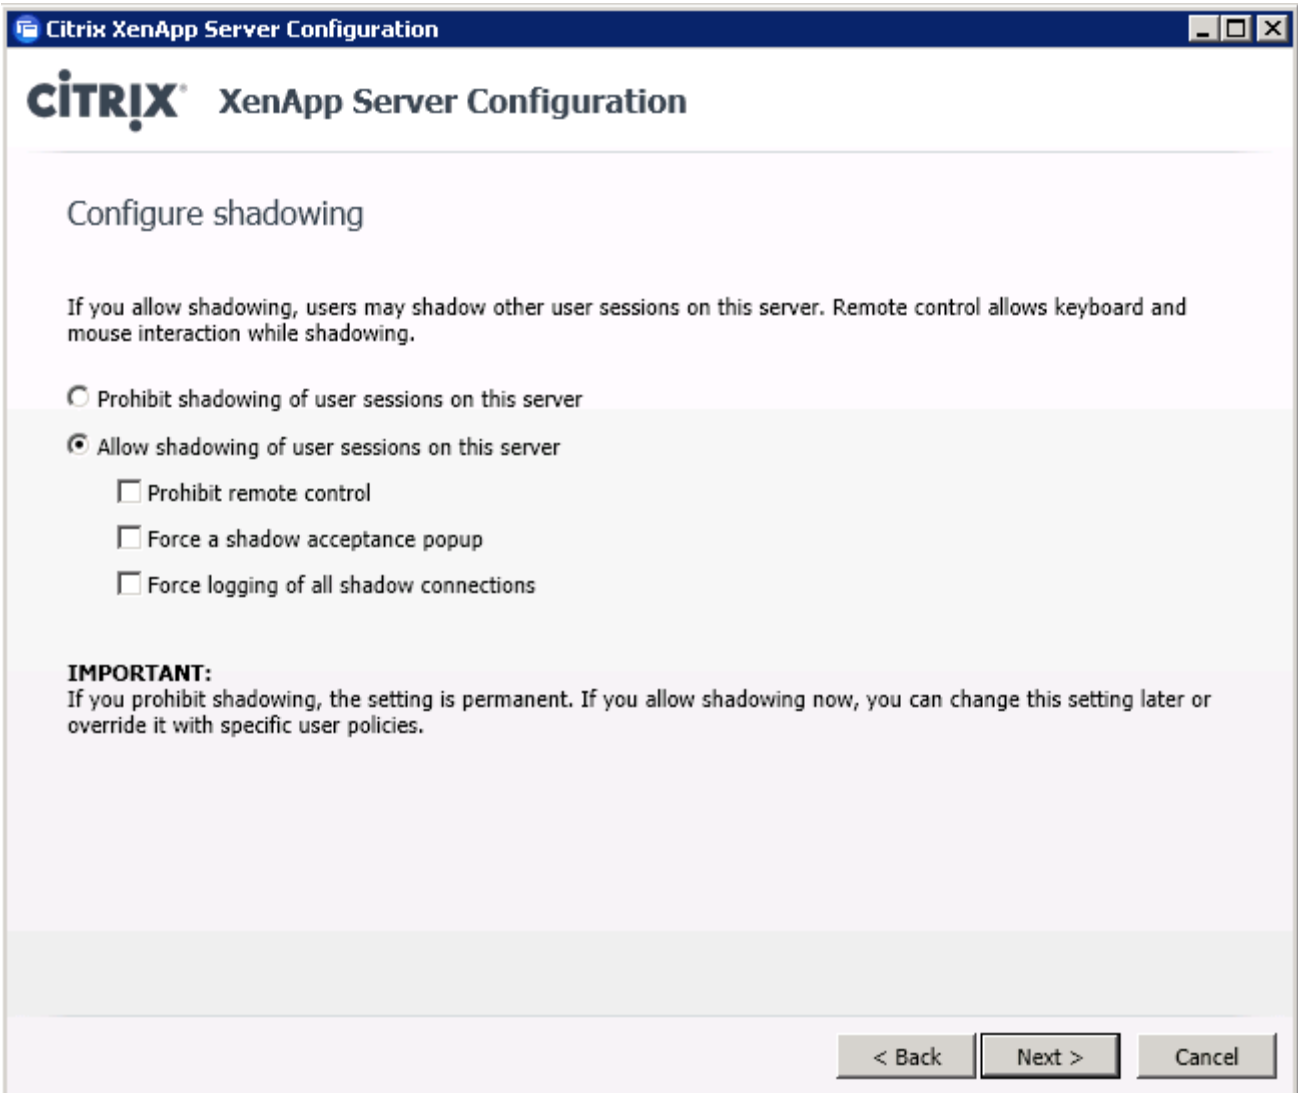
< Back Next > Cancel

7. Test database connection.



8. Choose default session shadowing settings (which allow shadowing).

Note: This is recommended for most farms.



The screenshot shows a window titled "Citrix XenApp Server Configuration". The main heading is "Configure shadowing". Below this, there is explanatory text: "If you allow shadowing, users may shadow other user sessions on this server. Remote control allows keyboard and mouse interaction while shadowing." There are two radio button options: "Prohibit shadowing of user sessions on this server" (unselected) and "Allow shadowing of user sessions on this server" (selected). Under the selected option, there are three checkboxes: "Prohibit remote control" (unchecked), "Force a shadow acceptance popup" (unchecked), and "Force logging of all shadow connections" (unchecked). An "IMPORTANT:" section follows, stating: "If you prohibit shadowing, the setting is permanent. If you allow shadowing now, you can change this setting later or override it with specific user policies." At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

9. Click Next.

Citrix XenApp Server Configuration

CITRIX XenApp Server Configuration

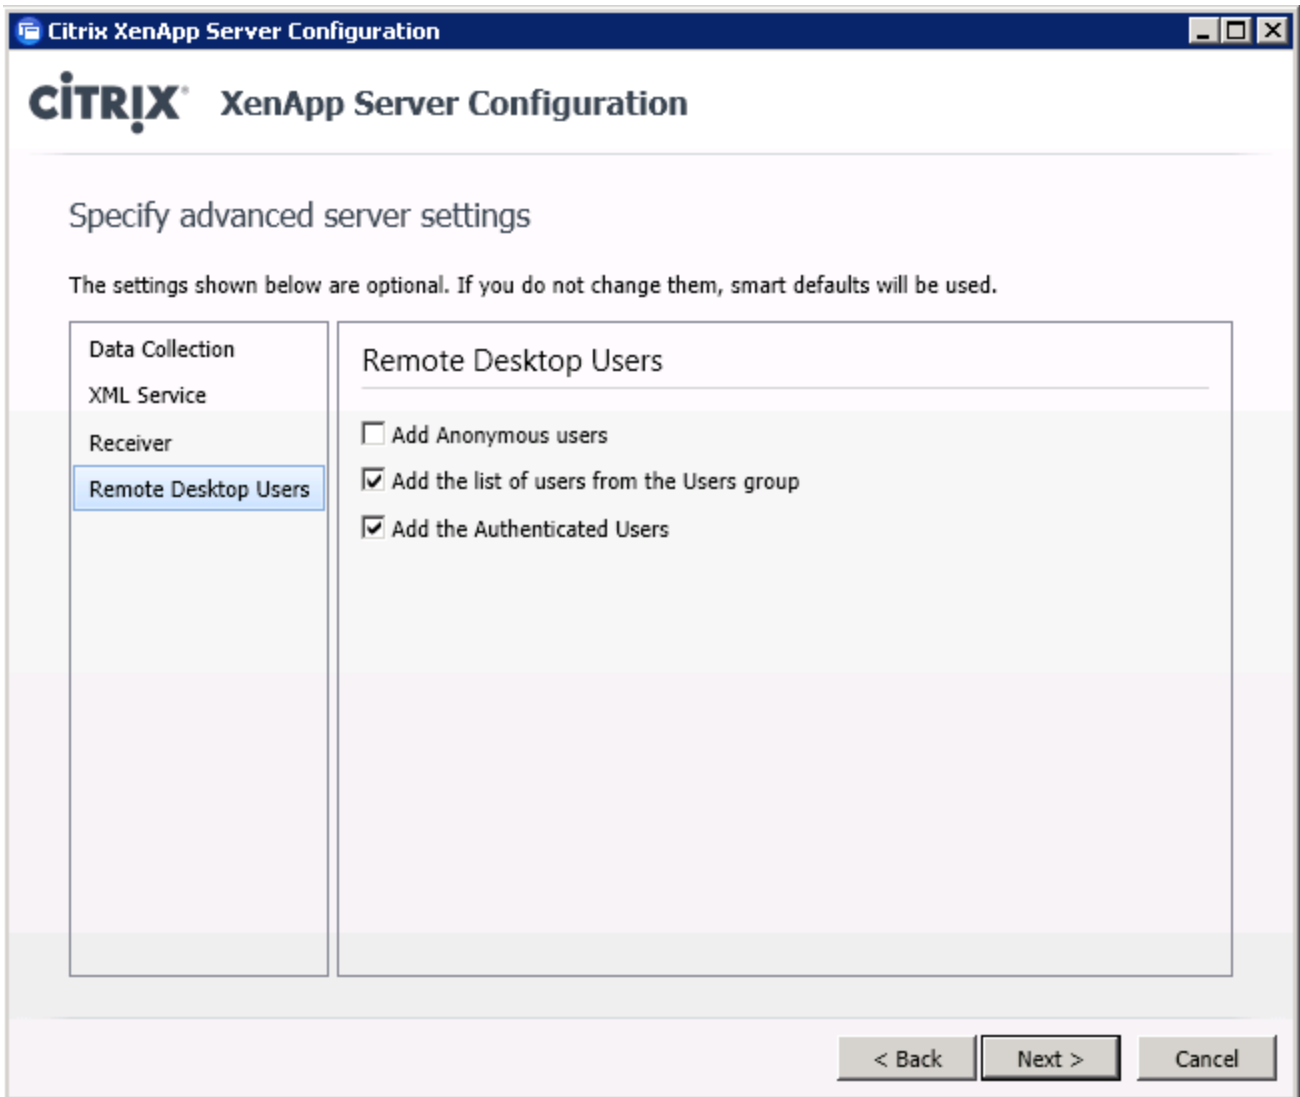
Specify advanced server settings

The settings shown below are optional. If you do not change them, smart defaults will be used.

Data Collection	Data Collection Options
XML Service	<input checked="" type="radio"/> Enable Controller and Session-host modes
Receiver	<input checked="" type="radio"/> Enable Session-host mode only ⓘ
Remote Desktop Users	<input type="checkbox"/> Use a custom zone name
	Zone name: <input type="text" value="Default Zone"/>

< Back Next > Cancel

10. Check Add the Authenticated Users.



Citrix XenApp Server Configuration

CITRIX XenApp Server Configuration

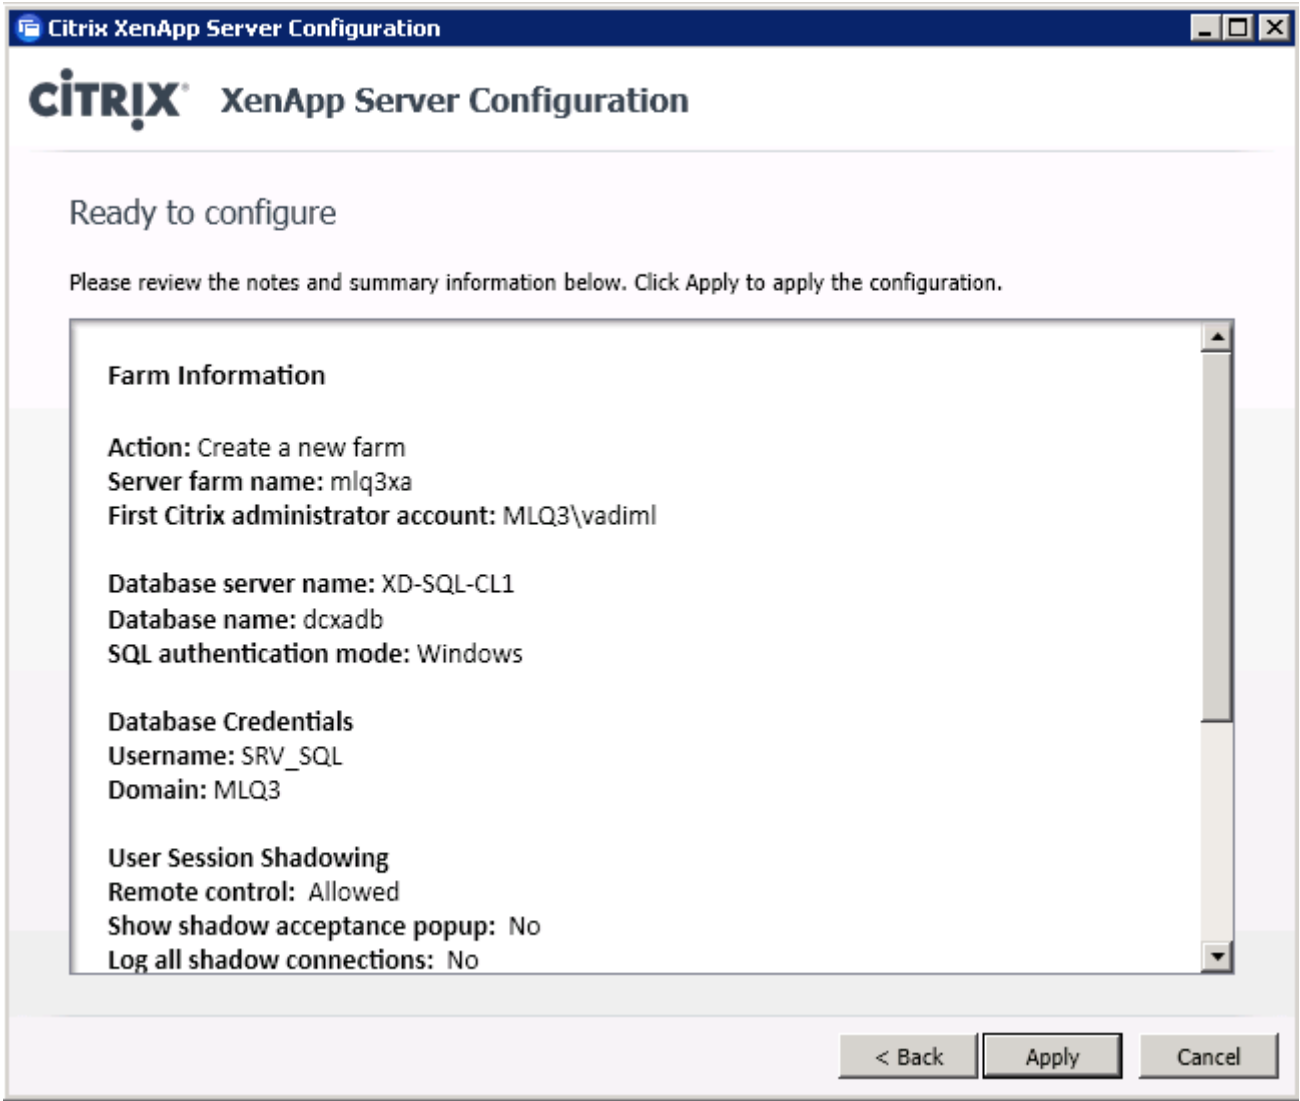
Specify advanced server settings

The settings shown below are optional. If you do not change them, smart defaults will be used.

Data Collection	<h4>Remote Desktop Users</h4> <ul style="list-style-type: none"><input type="checkbox"/> Add Anonymous users<input checked="" type="checkbox"/> Add the list of users from the Users group<input checked="" type="checkbox"/> Add the Authenticated Users
XML Service	
Receiver	
Remote Desktop Users	

< Back Next > Cancel

11. Review the summary page and click Apply.



12. Click Finish upon successful completion.




Citrix XenApp Server Configuration

CITRIX XenApp Server Configuration

Configuring server

- ✔ Saving configuration settings...
- ✔ Creating Citrix data store...
- ✔ Initializing Citrix data store...
- ✔ Starting Citrix IMA Services...
- ✔ Registering Citrix Administrators...
- ✔ Configuring Commands Remoting Service...

Server configured successfully!

 These settings will be applied when you clone or reboot this server.

XenApp

XenApp enables IT to centrally manage a single instance of each application and deliver it to users for online and offline use, providing a better than installed experience.

Review setup information
Select this item to access the latest setup information on the Citrix XenApp Web site.

Add or remove server roles
Number of roles installed: 1

Server Configuration Tasks

XenApp		
✓	Licensing specified	Edit Licensing ⓘ
✓	Configured	Edit Configuration ⓘ

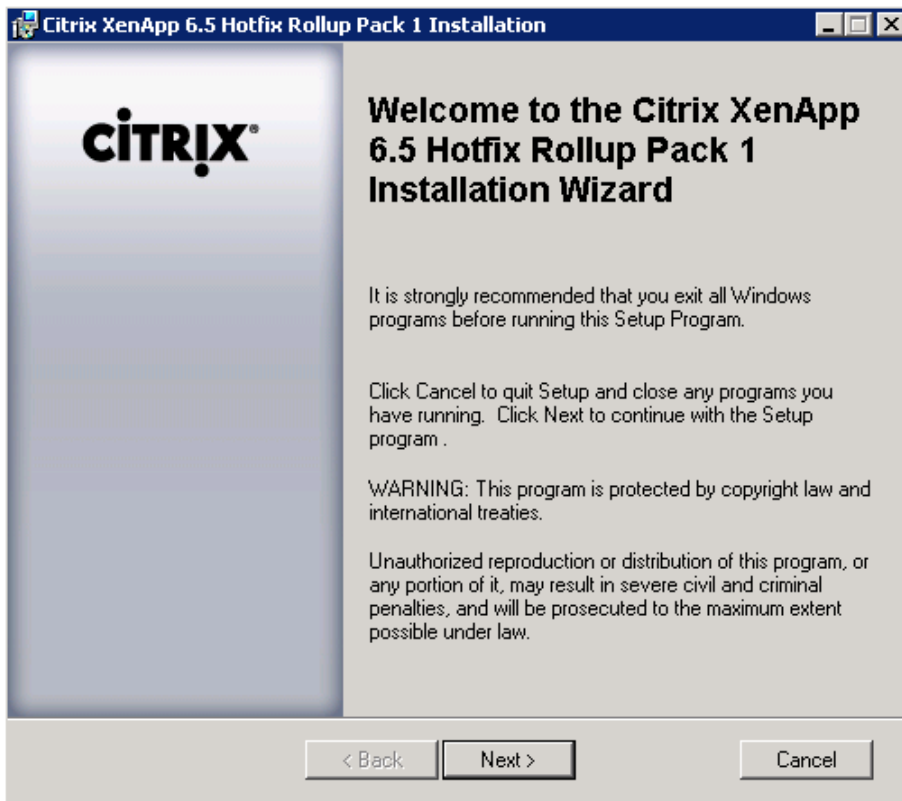
Do not show this window at logon

Close

6.10.5 Install Required XenApp Hot Fixes

Hotfix Rollup Pack 1 for Citrix XenApp 6.5 for Microsoft Windows Server 2008 R2 is required for proper system operation. It contains previously released Hot Fixes as well as introduces a number of new ones. For detailed information about issues fixed in the Hotfix Rollup Pack 1, see [CTX132122](#).

1. Run XA650W2K8R2X64R01.msp.
2. Click Next.



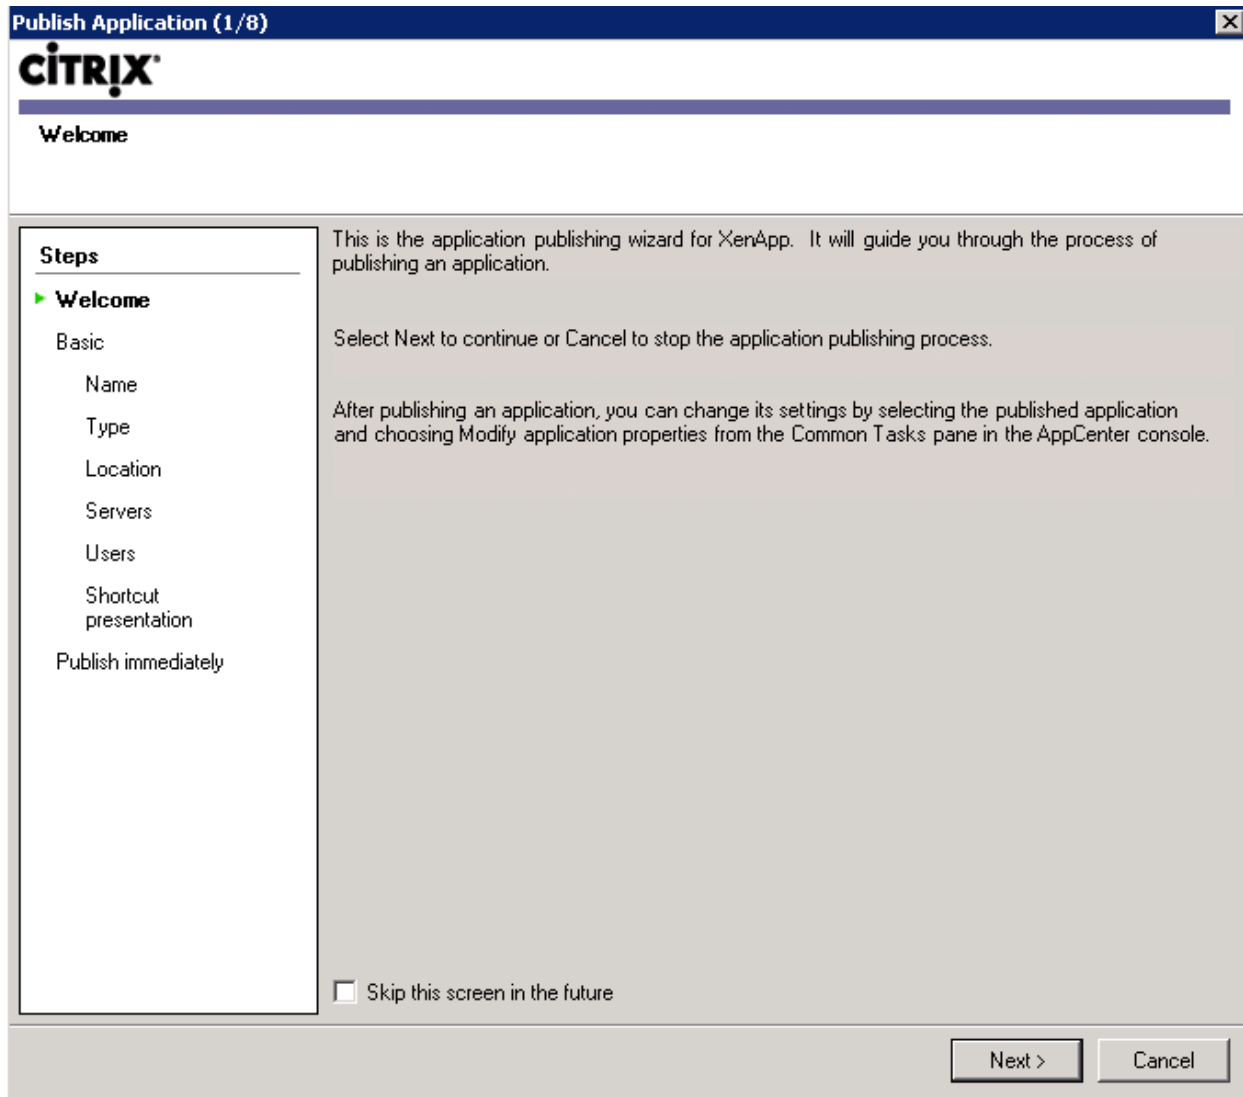
3. Click Close.



6.10.6 Configure the Hosted Shared Desktop (Citrix)

1. Open Citrix AppCenter management console (Administrative Tools | Citrix | Management Consoles).

2. Under the XenApp node, expand the farm.
3. Select the Applications node and from the Actions pane choose Publish application to lunch Publish Application wizard.
4. Click Next.



5. Provide a display name (maximum 256 characters) and application description. Then click Next.

Note: The name appears on user devices when users access the application and on the AppCenter for the farm applications. XenApp supports application names that use Latin-1 and Unicode character sets, including characters used in Japanese, Chinese, and Korean.

HostedSharedDesktop - Publish Application

CITRIX

Name

Enter the name and description that you want to be displayed to clients for this application.

Display name:

Application description:

Steps

- ✓ Welcome
- Basic**
 - ▶ **Name**
 - Type
 - Location
 - Servers
 - Users
 - Shortcut presentation
 - Publish immediately

< Back Next > Cancel

6. Select Server desktop and click Next.

HostedSharedDesktop - Publish Application (3/7) x

CITRIX

Type
View the application type. To change the type, use the Change application type task.

Steps

- ✓ Welcome
- Basic**
- ✓ Name
- ▶ **Type**
- Servers
- Users
- Shortcut presentation
- Publish immediately

Choose the type of application to publish.

Server desktop
 Content
 Application

Application type

Accessed from a server
 Streamed if possible, otherwise accessed from a server

Server application type:

Streamed to client

Note: To change the application type after publishing it, you must use the Change Application Type task.

Quick Help

Grants users access to an entire server desktop including resources and applications available from desktop menus and shortcuts.

7. Add the individual servers or worker groups on which the published application runs when accessed by an ICA connection.

HostedSharedDesktop - Publish Application (4/7)

CITRIX

Servers
Configure which servers will host the application.

Steps

- ✓ Welcome
- Basic**
- ✓ Name
- ✓ Type
- ▶ **Servers**
- Users
- Shortcut presentation
- Publish immediately

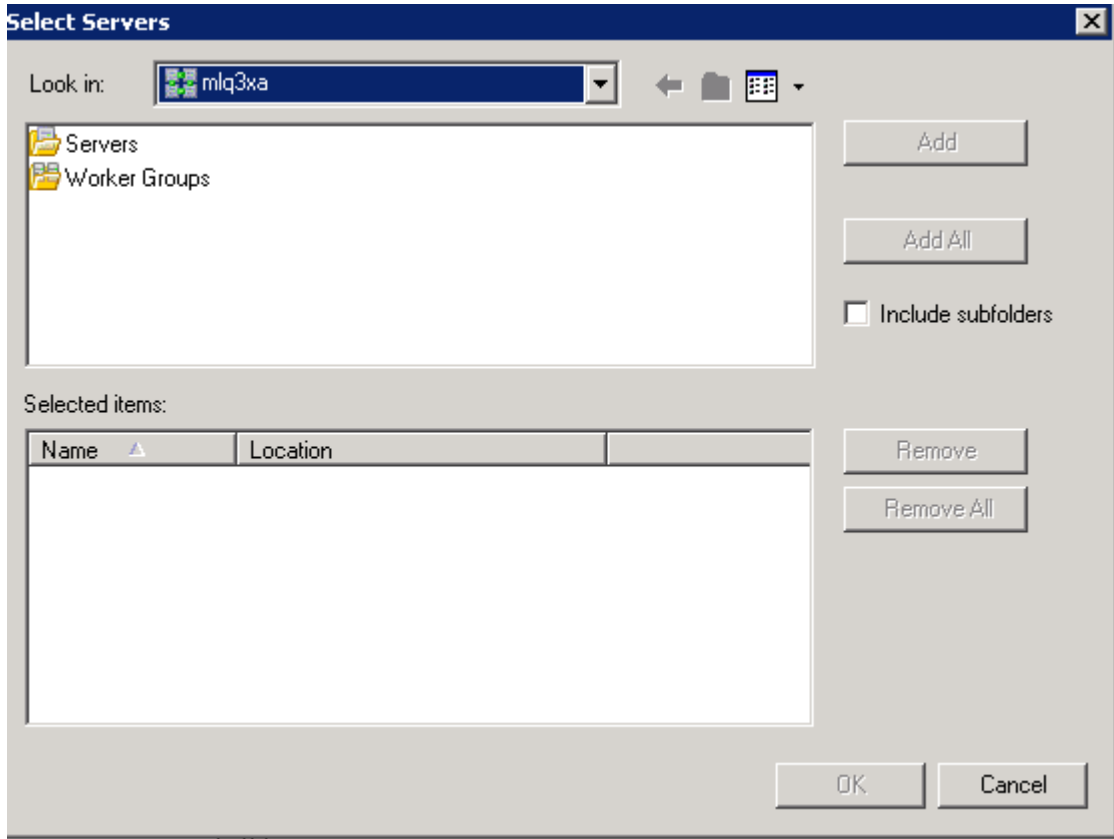
Choose the servers on which this published application will run when being delivered via ICA.

Servers:

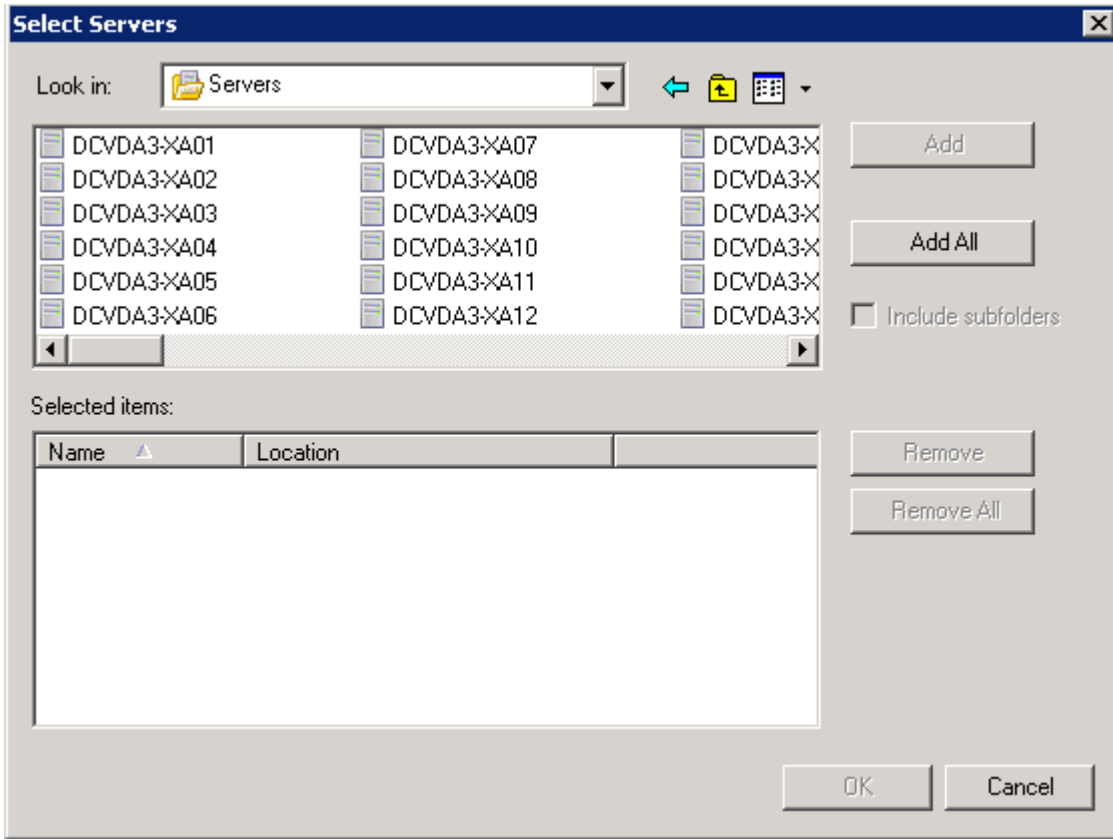
Name	Relative location
------	-------------------

0 items

8. Click Servers.



9. Click Add All.



10. On the Users page, create the Configured users list for users or groups who have access to the application.

HostedSharedDesktop - Publish Application (5/7)

CITRIX

Users

Configure the users who may access the application.

Steps

- ✓ Welcome
- Basic**
- ✓ Name
- ✓ Type
- ✓ Servers
- ▶ **Users**
- Shortcut presentation
- Publish immediately

Specify the users who can access this application.
To add users, choose a directory type at the bottom and select Add. [More...](#)

Allow anonymous users
 Allow only configured users

Configured users:

Select directory type:

< Back Next > Cancel

11. Click Next.

HostedSharedDesktop - Publish Application (6/7)


CITRIX

Shortcut presentation

Configure the appearance and location of the application shortcut.

These settings function differently on different clients. [More...](#)

Application icon

Icon: 

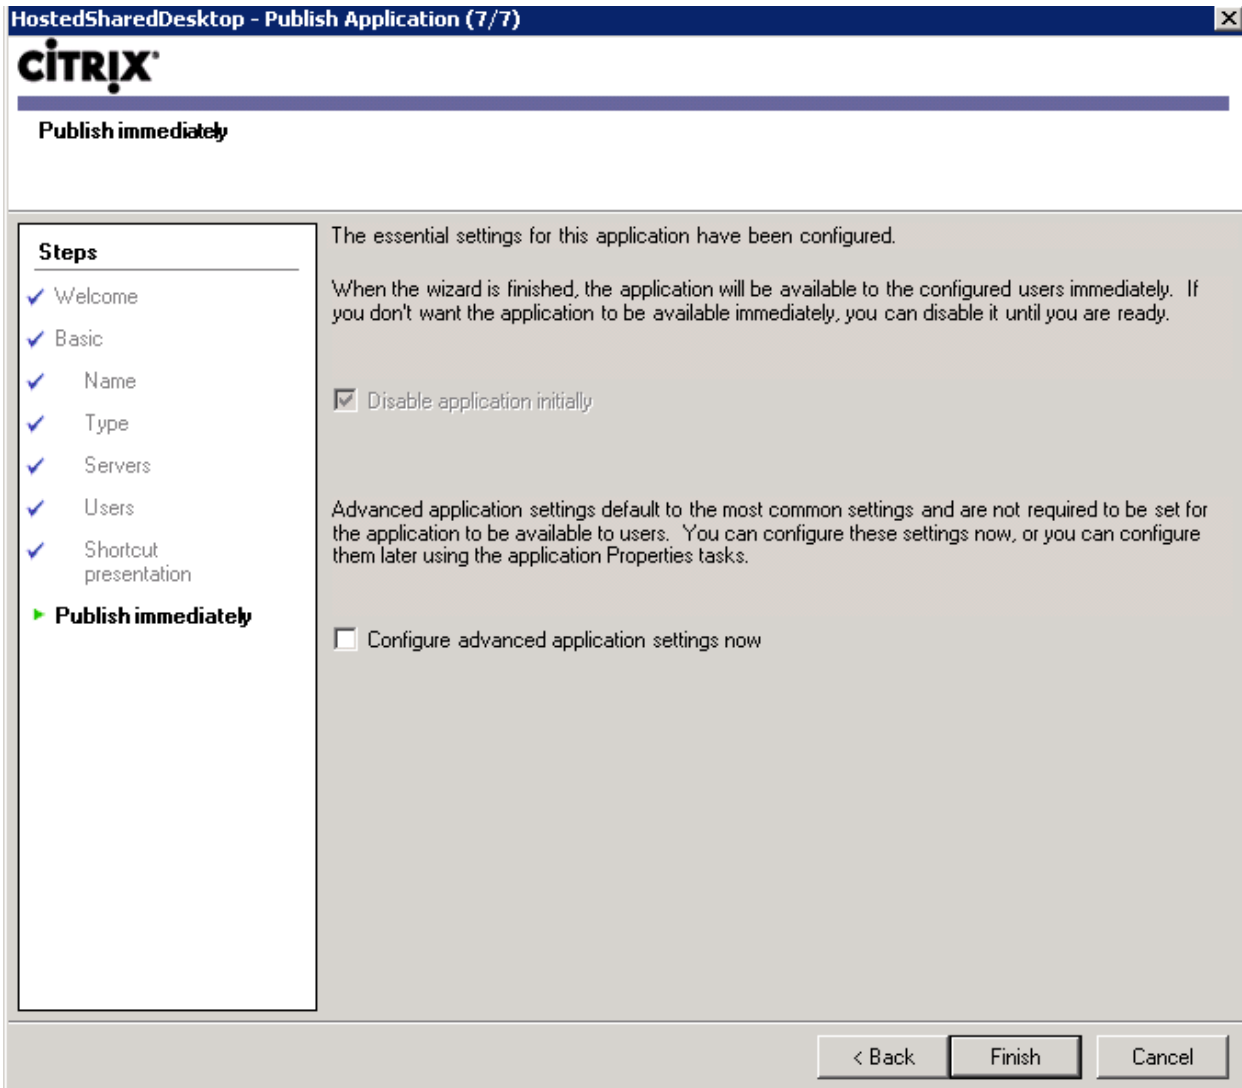
Client application folder:

Application shortcut placement

Add to the client's Start menu
Start menu folder (Citrix XenApp plugin only):

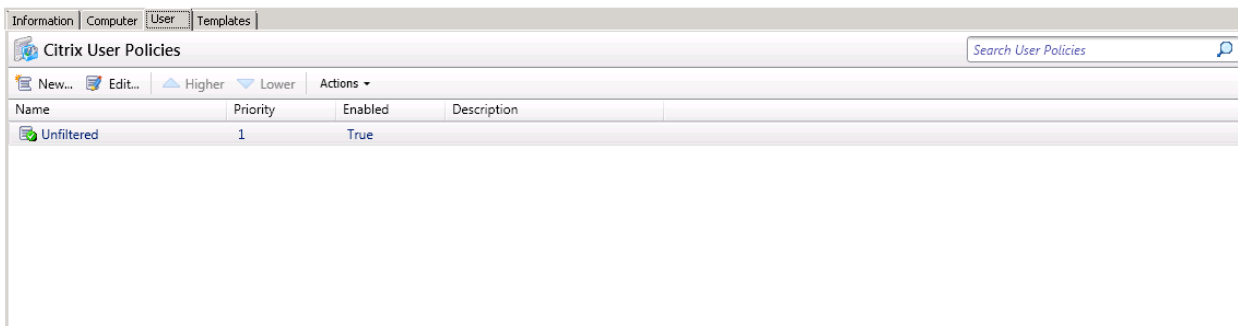
Add shortcut to the client's desktop

12. Click Finish to the publish application.



6.10.7 Configure XenApp Policies

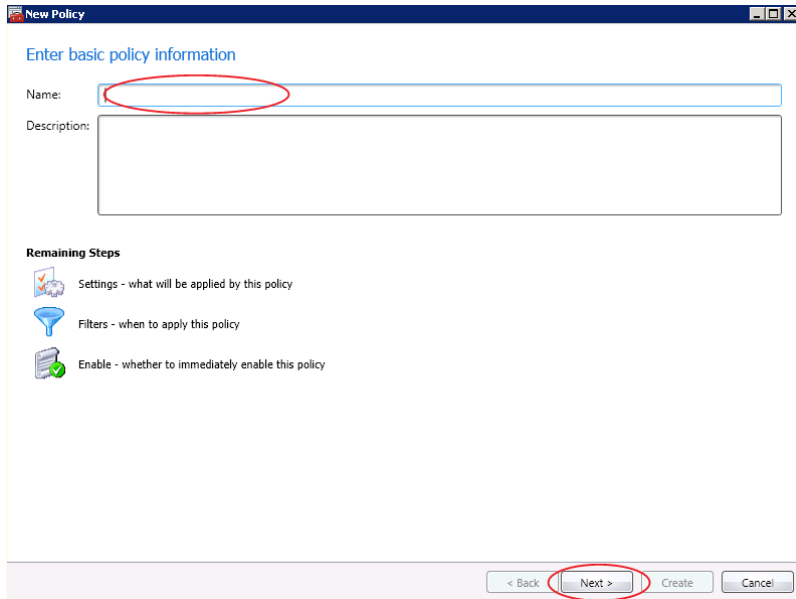
1. From the AppCenter, select the Policies node in the left pane and then select the User tab.



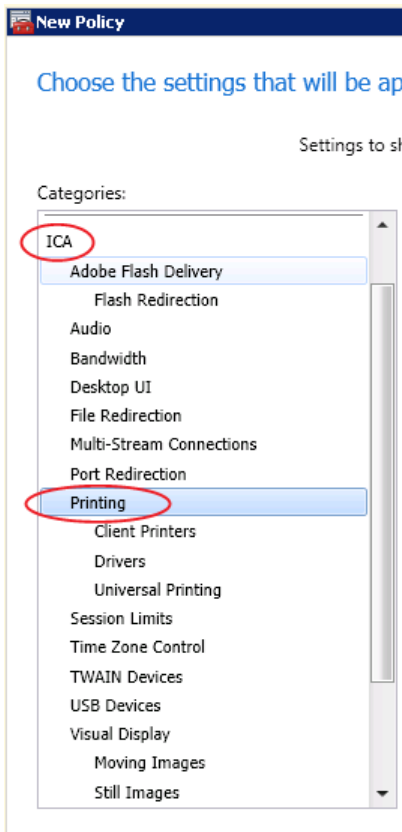
2. Click New to start the policy creation process.



3. Enter a name for your policy and click Next.



4. Select from Categories ICA → Printing.



5. Select Client printer redirection and click Add.

Settings:

	Client printer redirection ICA\Printing	Add
	Default printer ICA\Printing	Add
	Printer auto-creation event log preference ICA\Printing	Add
	Session printers ICA\Printing	Add
	Wait for printers to be created (desktop) ICA\Printing	Add

6. Click Prohibited and then click OK.

Add Setting

Client printer redirection

Allowed
Client printers can be mapped, if specified elsewhere

Prohibited
No client printers will be mapped

Help Comment

Applies to XenApp 6.0 or later and XenDesktop 5.0 or later

Allows or prevents client printers to be mapped to a server when a user logs on to a session. By default, client printer mapping is allowed.

OK Cancel

7. Select from Categories ICA → Printing → Client Printers.

Categories:

- ICA**
- Adobe Flash Delivery
 - Flash Redirection
- Audio
- Bandwidth
- Desktop UI
- File Redirection
- Multi-Stream Connections
- Port Redirection
- Printing**
 - Client Printers**
 - Drivers
 - Universal Printing
- Session Limits
- Time Zone Control
- TWAIN Devices
- USB Devices
- Visual Display
 - Moving Images
 - Still Images

8. Select Auto-create client printers and click Add.



Settings:

	Auto-create client printers ICA\Printing\Client Printers	Add
	Auto-create generic universal printer ICA\Printing\Client Printers	Add
	Client printer names ICA\Printing\Client Printers	Add
	Direct connections to print servers ICA\Printing\Client Printers	Add
	Printer driver mapping and compatibility ICA\Printing\Client Printers	Add
	Printer properties retention ICA\Printing\Client Printers	Add
	Retained and restored client printers ICA\Printing\Client Printers	Add

9. From the drop-down list, pick Do not auto-create client printers and click OK.

Add Setting

Auto-create client printers

Value: **Do not auto-create client printers**

Use default value

Help | Comment

Applies to XenApp 6.0 or later and XenDesktop 5.0 or later

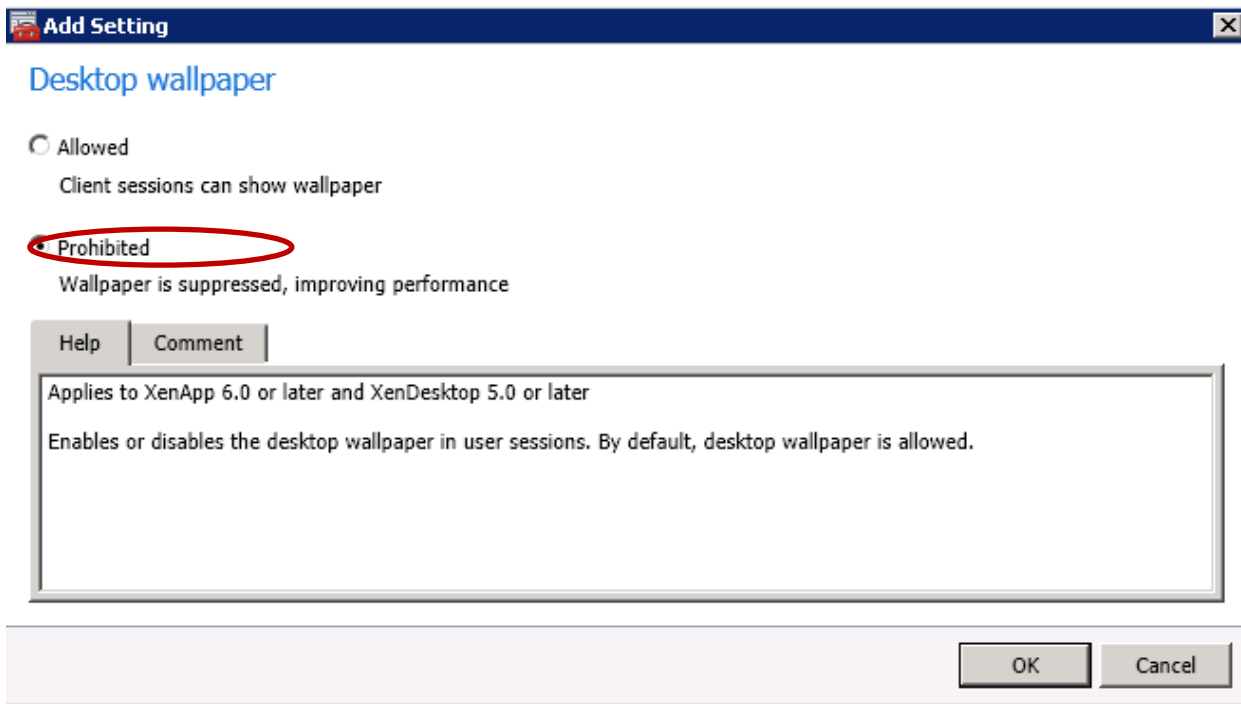
Specifies which client printers are auto-created. This setting overrides default client printer auto-creation settings. By default, all client printers are auto-created.

OK Cancel

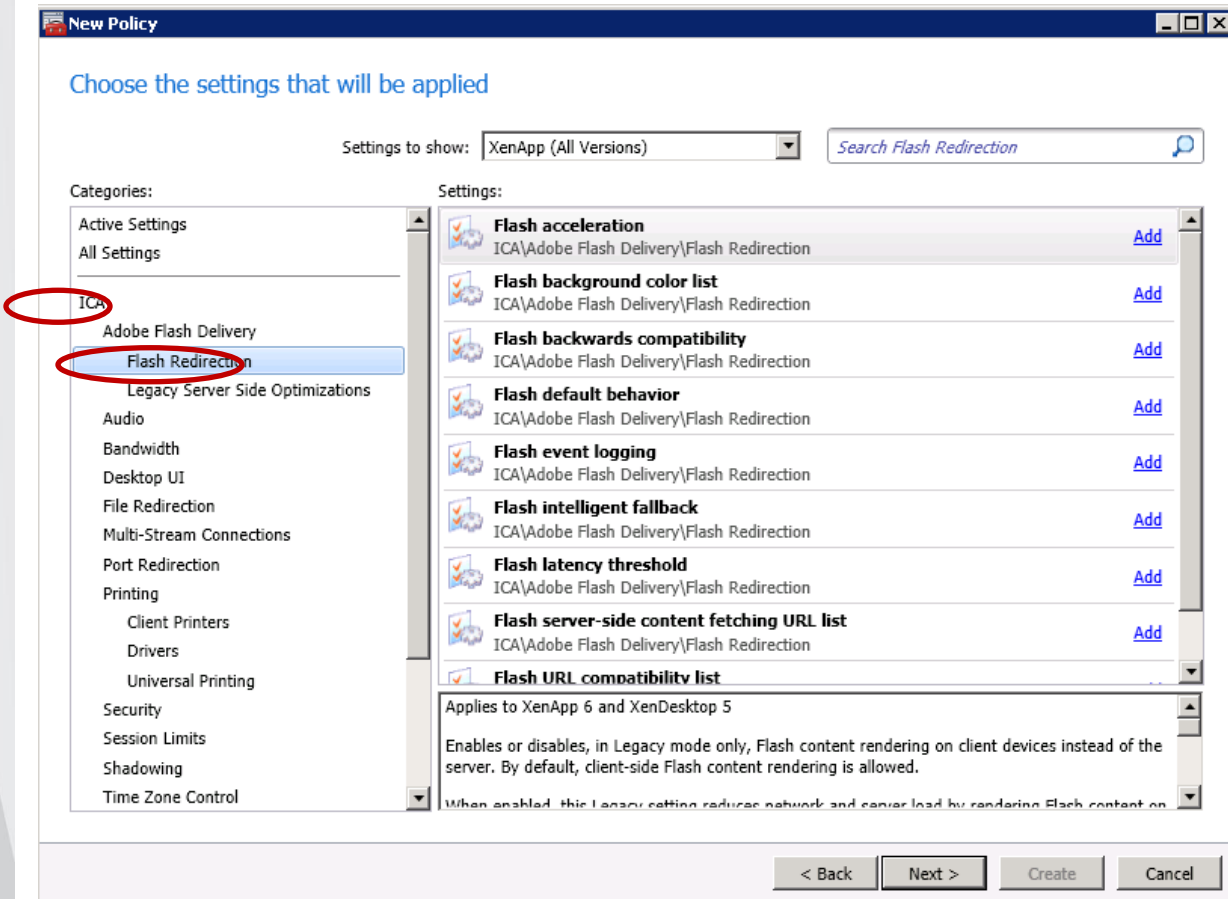
10. Select from Categories ICA → Desktop UI.

11. Select Desktop Wallpaper and click Add.

12. Select Prohibited and click OK.



13. Select from Categories ICA → Flash Redirection.



13. Select Flash acceleration and click Add.

Settings:

	Flash acceleration ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash background color list ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash backwards compatibility ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash default behavior ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash event logging ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash intelligent fallback ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash latency threshold ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash server-side content fetching URL list ICA\Adobe Flash Delivery\Flash Redirection	Add
	Flash URL compatibility list	Add

14. Select Disabled and click OK.

Add Setting ✕

Flash acceleration

Enabled
Unless prevented by the Flash URL Compatibility list, Flash content is rendered on the user device.

Disabled
All web sites are rendered on the server

Help

Comment

Applies to XenApp 6 and XenDesktop 5

Enables or disables, in Legacy mode only, Flash content rendering on client devices instead of the server. By default, client-side Flash content rendering is allowed.

When enabled, this Legacy setting reduces network and server load by rendering Flash content on the client device. Additionally, the Flash URL Compatibility list setting forces Flash content from specific Web sites to be rendered on the server.

15. Click Next.

16. Select All Filters → User or Group and click Add Create policy filter.

New Policy

Choose when to apply the settings using filters

Filters to show: **XenApp (All Versions)**

Categories:

- Active Filters
- All Filters

Filters:

	Access Control	Add
	Branch Repeater	Add
	Client IP Address	Add
	Client Name	Add
	Organizational Unit	Add
	User or Group	Add
	Worker Group	Add

Applies to XenApp 6.0 or later and XenDesktop 5.0 or later

The Access Control filter allows you to apply policies based on the access control conditions through which a client is connecting.

If no filters are configured for an enabled policy then all of the configured settings will be applied.

< Back Next > Create Cancel

17. Click Add.

New Filter

New User or Group Filter

Filter elements:

[Add...](#) [Edit...](#) [Remove](#) Enable

OK Cancel



18. Pick the User Group for the policy.

New Filter Element

New User or Group Filter Element

Mode: Allow Enable this filter element

User or group name: Browse...

Comment:

OK Cancel

19. Click OK.

New Filter

New User or Group Filter

Filter elements:

Allow - MLQ3\Domain Users	Enabled
---------------------------	---------

Add... Edit... Remove Enable

OK Cancel

20. Click Next.

New Policy

Choose when to apply the settings using filters

Filters to show: XenApp (All Versions) [Search All Filters](#)

Categories:

- Active Filters
- All Filters

Filters:

- Access Control** [Add](#)
- Branch Repeater** [Add](#)
- Client IP Address** [Add](#)
- Client Name** [Add](#)
- Organizational Unit** [Add](#)
- User or Group** [Edit](#) [Remove](#)
Allow - MLQ3\Domain Users
- Worker Group** [Add](#)


Applies to XenApp 6.0 or later and XenDesktop 5.0 or later

The User filter allows you to apply policies based on the user or group membership of the user connecting to the session.

If no filters are configured for an enabled policy then all of the configured settings will be applied.

< Back Next > Create Cancel

21. Click Create.



6.11 Installing and Configuring Citrix NetScaler Virtual Appliance on VMware ESX for Use with Citrix StoreFront

To install NetScaler you must follow the procedure here for importing and deploying a NetScaler template: <http://support.citrix.com/proddocs/topic/netscaler-vpx-10/ns-vpx-install-on-esx-wrapper-con.html>

When the VM is up, continue with these steps:

1. Enter an IP address, Subnet Mask and Gateway for NetScaler.

```

The key fingerprint is:
ed:f2:d2:bc:6a:0f:0b:a9:9d:ba:fd:2e:5a:36:a4:ca root@ns
.
machdep.cpu_idle_hlt: 0 -> 1
Start daemons: syslogd Jan 21 19:07:25 <kern.info> ns syslogd: kernel boot file
is /flash/ns-10.0-71.6
inetd cron httpd monit sshd vmware_guestd .

!There is no ns.conf in the /nsconfig!

Start Netscaler software
tput: no terminal type specified and no TERM environmental variable.
No machine id

Enter NetScaler's IPv4 address []: 10.218.241.71
Enter Netmask []: 255.255.255.0
Enter Gateway IPv4 address []: 10.218.241.1

```

2. Enter option 4 to save and quit.

```

NSVPX-DC-02 on 192.168.1.59
File View VM
[Icons: Stop, Run, Refresh, Home, Back, Forward, Print, Help, Copy]

Enter NetScaler's IPv4 address []: 10.218.241.71
Enter Netmask []: 255.255.255.0
Enter Gateway IPv4 address []: 10.218.241.1

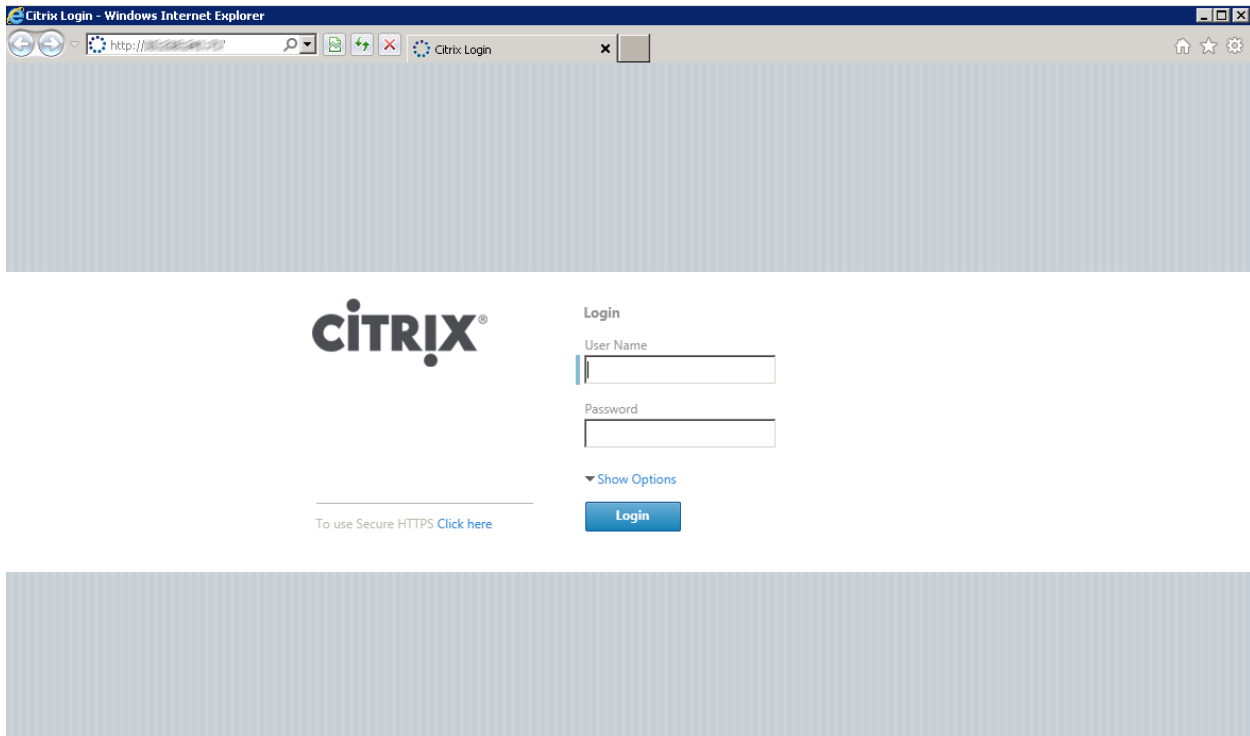
-----
Netscaler Virtual Appliance Initial Network Address Configuration.
This menu allows you to set and modify the initial IPv4 network addresses.
The current value is displayed in brackets ([]).
Selecting the listed number allows the address to be changed.

After the network changes are saved, you may either login as nsroot and
use the Netscaler command line interface, or use a web browser to
http://10.218.241.71 to complete or change the Netscaler configuration.
-----

1. NetScaler's IPv4 address [10.218.241.71]
2. Netmask [255.255.255.0]
3. Gateway IPv4 address [10.218.241.1]
4. Save and quit
Select item (1-4) [4]: █

```

3. Use a browser to access the NetScaler web interface and login.



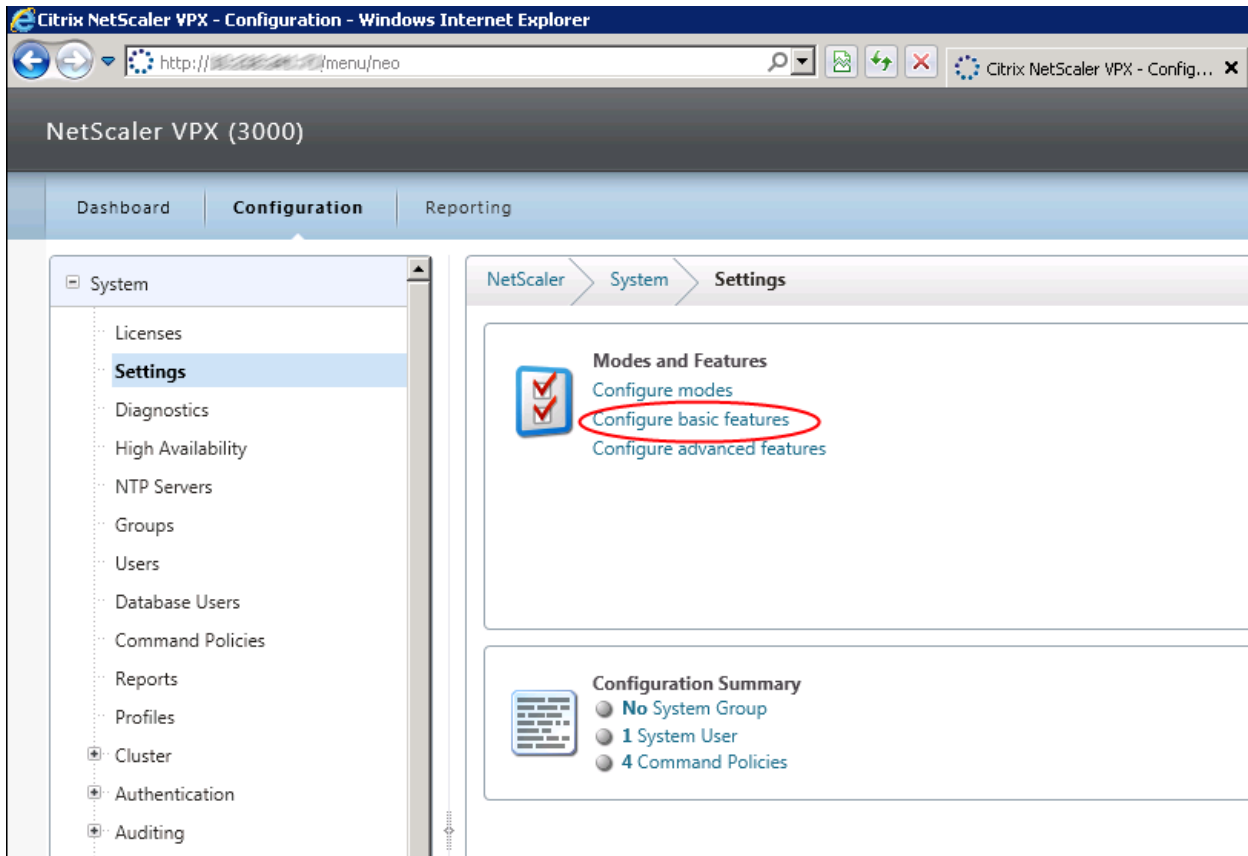
4. Under Configuration, System, Licenses click Manage Licenses.

NetScaler > System > Licenses

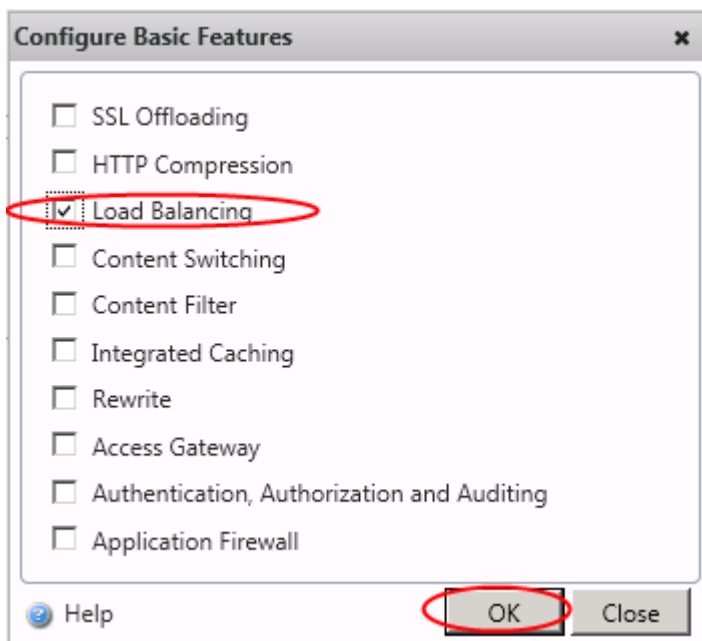
License Name	License Key	License Type	License Status
		S	
		Co	
		Maximum IC	
		Authentication, Authorizati	
		F	
		Global Server	
		D	

Manage Licenses...

5. Click **Add** and import your license. Make sure to use an appropriate license to accommodate the size of your environment. A platinum license was used for this exercise. Learn more about NetScaler VPX licensing here: <http://support.citrix.com/article/CTX122426>
6. From Settings click Configure basic features.



7. Check the checkbox for Load Balancing and click "OK".



8. From Network click IPs.

The screenshot shows the NetScaler VPX (3000) Configuration page. The left sidebar contains a 'Network' menu with 'IPs' highlighted. The main content area shows the 'Network' configuration page with a 'Monitor Connections' section and a 'Configuration Summary' section.

Configuration Summary

- 3 System IPs
- 1 System IPv6
- 1 VLAN
- No IPv4 Tunnel
- No IPv6 Tunnel

9. Click Add.

NetScaler > Network > IPs > IPV4s Refresh Help Save

IPV4s IPV6s

25 Per Page | 1 - 1 of 1 | 1

IP Address	State	Type	Mode	ARP	ICMP	Virtual Server
	Enabled	Netscaler IP	Active	ENABLED	ENABLED	-N/A-

Add... Open... Remove Enable Disable Add Range... Statistics

10. In the Create IP menu, enter a new IP address for server-side connections, select Mapped IP and click Create.

Create IP [X]

IP Address* [] Netmask* []

Virtual Router ID [] ICMP Response [NONE]

ARP Response [NONE]

State

IP Type

Subnet IP Virtual IP Mapped IP GSLB Site IP Cluster IP

Options

ARP ICMP Virtual Server Dynamic Routing

Host Route

Enable Gateway IP [. . .] Metric []

V Server RHI Level

NONE ONE_VSERVER ALL_VSERVERS

OSPF LSA Type

TYPE5 TYPE1 Area []

Application Access Controls

Enable Management Access control to support the below listed applications.

Applications

Telnet FTP SSH SNMP

GUI

Secure access only

Allow access only to management applications

Help [Create] [Close]

11. From Servers. click Add.



NetScaler > Load Balancing > Servers Refresh | Help | Save

25 Per Page | 0 - 0 of 0 | Search

Name	State	IPAddress / Domain
------	-------	--------------------

Add... | Open... | Enable | Disable | Remove | Add Range... | Show Bindings... | Restart DBS Monitors | Rename

12. Enter your StoreFront Server's name in the field, and enter its IP address. Repeat this step for your 3 StoreFront servers. Click Create.



Create Server [X]

Server Name* [Redacted]

IP Address Domain Name

IP Address* [Redacted] IPv6

Translation IP Address [. . .]

Translation Mask [. . .]

Resolve Retry (secs) [5]

IPv6 Domain

Enable after Creating

Comments []

[Help](#) **Create** **Close**

13. Click Virtual Servers and click Add.

Load Balancing

- Virtual Servers
- Services
- Service Groups
- Monitors
- Metric Tables
- Servers
- Persistency Groups

+	AAA - Application Traffic	!
+	Content Switching	!
+	Cache Redirection	!
+	GSLB	!
+	Rewrite	!
+	Responder	!
+	Access Gateway	!
+	Web Interface	
+	Application Firewall	!

< Add... Open.. Enable Disable Remove

14. Enter a name for the VIP, give it an IP address, leave the default settings and click “Add” to add a Service.

Create Virtual Server (Load Balancing)

Name* IP Address Based IP Pattern Based

Protocol* HTTP IP Address* IPv6

Network VServer Range Port*

Directly Addressable State AppFlow Logging

Services | Service Groups | Policies | Method and Persistence | Advanced | Profiles | SSL Settings

Activate All Deactivate All

Active	Service Name	IP Address	Port	Protocol	State	Weight	Dynamic Weight
--------	--------------	------------	------	----------	-------	--------	----------------

Comments

15. Select your StoreFront server under Server. Select Protocol HTTP, port 80 and click Create. Repeat this step for your 3 StoreFront Servers.

Create Service [x]

Service Name* dc-store-01 Server* dc-store-01 ()

Protocol* HTTP Port* 80

Enable Service Enable Health Monitoring AppFlow Logging

Monitors Policies Profiles Advanced SSL Settings

Available

Monitors
arp
nd6
ping
tcp
tcp-ecv
http-ecv
udp-ecv
dns
ftp
tcps
https

Add >

< Remove

Configured

Monitors	Weight	State

Comments

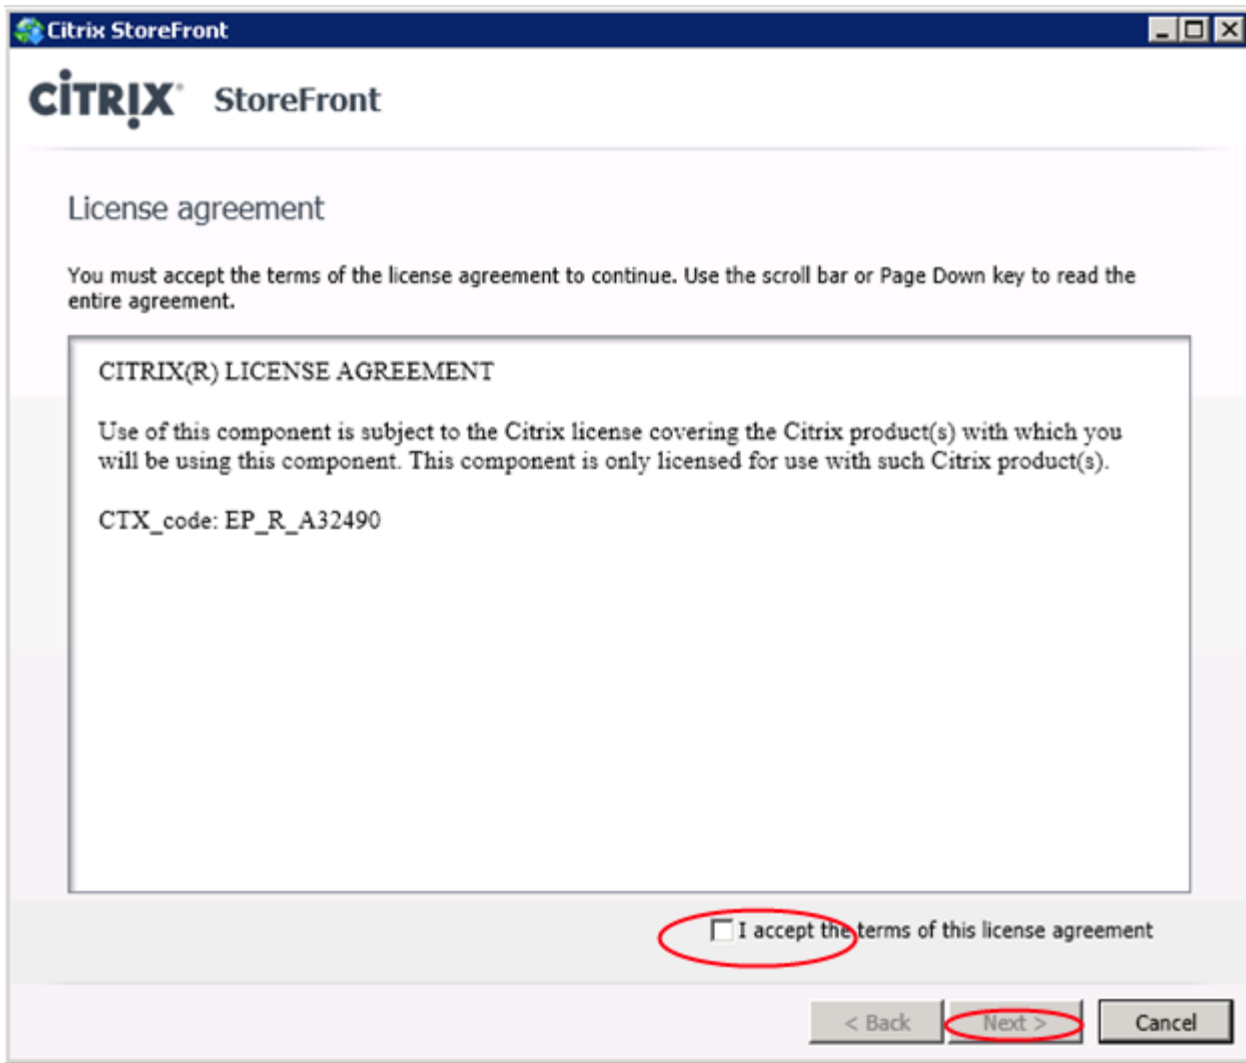
Help Create Close

16. Click the Method and Persistence tab, select least Connection Method, COOKIEINSERT Persistence, and SOURCEIP Backup Persistence. Click OK.

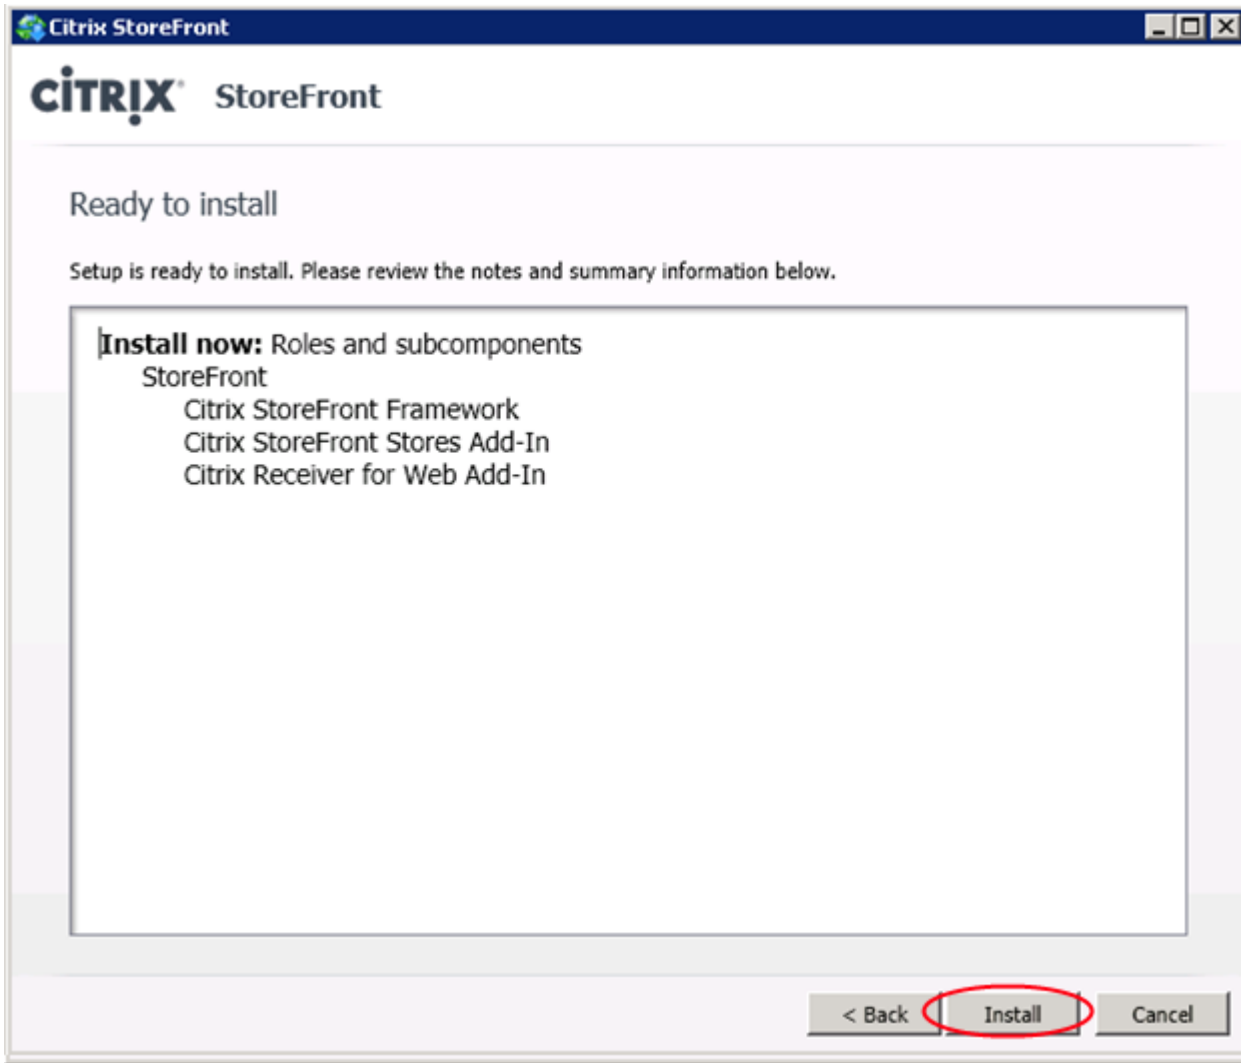
6.12 Installing and Configuring Citrix StoreFront for XenDesktop and XenApp

6.12.1 Install

1. Log on to the StoreFront server using an account with local administrator permissions.
2. Browse your installation media or download package, locate CitrixStoreFront-x64.exe, and run the file as an administrator. If a message appears indicating that Microsoft .NET Framework 3.5 with Service Pack 1 will be enabled, click Yes.
3. Read and accept the license agreement, and click Next.



4. If the Review prerequisites page appears, click Next.
5. On the Ready to install page, check that all three StoreFront components are listed for installation and click Install.



Note: Before the components are installed, the .NET Framework 3.5.1 Features > .NET Framework 3.5.1 feature and the Web Server (IIS) role are deployed, and the following role services are enabled if they are not already configured on the server.

- Web Server > Common HTTP Features > Static Content, Default Document, HTTP Errors, HTTP Redirection
 - Web Server > Application Development > ASP.NET, .NET Extensibility, ISAPI Extensions, ISAPI Filters
 - Web Server > Health and Diagnostics > HTTP Logging
 - Web Server > Security > Windows Authentication, Request Filtering
 - Management Tools > IIS Management Console, IIS Management Scripts and Tools
 - Management Tools > IIS 6 Management Compatibility > IIS 6 Metabase Compatibility, IIS 6 WMI Compatibility, IIS 6 Scripting Tools
6. When the installation is complete, click Finish.
 7. Go to Start, Control Panel, Administrative Tools and Click on Services.
 8. Change Citrix Services to start in Automatic (Delayed Start)



- Citrix Configuration Replication
- Citrix Credential Wallet
- Citrix Peer Resolution Service

Provides a ... Started Automatic (Delayed Start)
Provides a ... Started Automatic (Delayed Start)
Resolves p... Started Automatic (Delayed Start)

Network S...
Network S...
Network S...

6.12.2 Create StoreFront Database

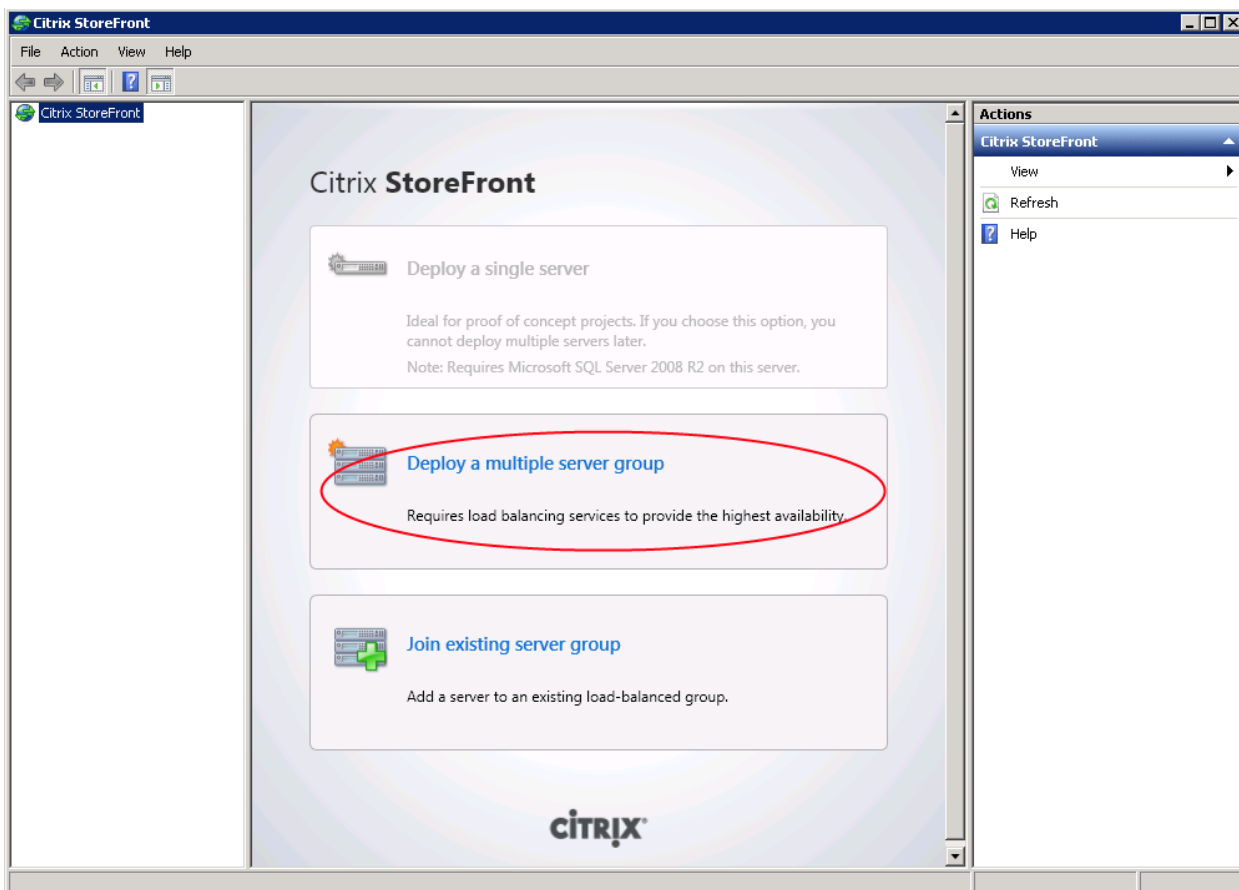
If you plan to use a remote database with a multiple server StoreFront deployment, follow the Procedure to properly create the database, users, and set permissions:

<http://support.citrix.com/proddocs/topic/dws-storefront-12/dws-deploy-multi-database.html>

Note: Create and use AD security group for the Storefront servers.

6.12.3 Create Multiple Server Group

1. If the Citrix StoreFront management console is not already open after installation of StoreFront, click Start > All Programs > Citrix > Citrix StoreFront.
2. In the results pane of the Citrix StoreFront management console, click Deploy a multiple server group.



3. Specify the URL of the load balancing environment hosting the StoreFront server in the Hostname (load balancer) box.



Note: In order to configure a multiple server deployment, the StoreFront servers must be part of an existing load balancing environment.

Hostname
(load balancer):

4. Provide details of the SQL Server instance to be used to record details of users' application subscriptions for your first store. Enter the fully qualified domain name of the database server and the name of the database.

Database server:

Database name:

5. Click Test Connection to ensure that StoreFront can access the specified database. If the database details you provide cannot be verified, select the Specify connection string check box to manually compose the database connection string.

Note: The credentials with which you log on to the StoreFront server are used to test the database connection. Ensure that this user account has permissions to access the database to enable StoreFront to validate the connection details.


6. Click Create to set up the authentication service, which authenticates users to XenDesktop sites, XenApp farms, and AppController.

Deploy Multiple Server Group ✕

CITRIX

Deploy Multiple Server Group

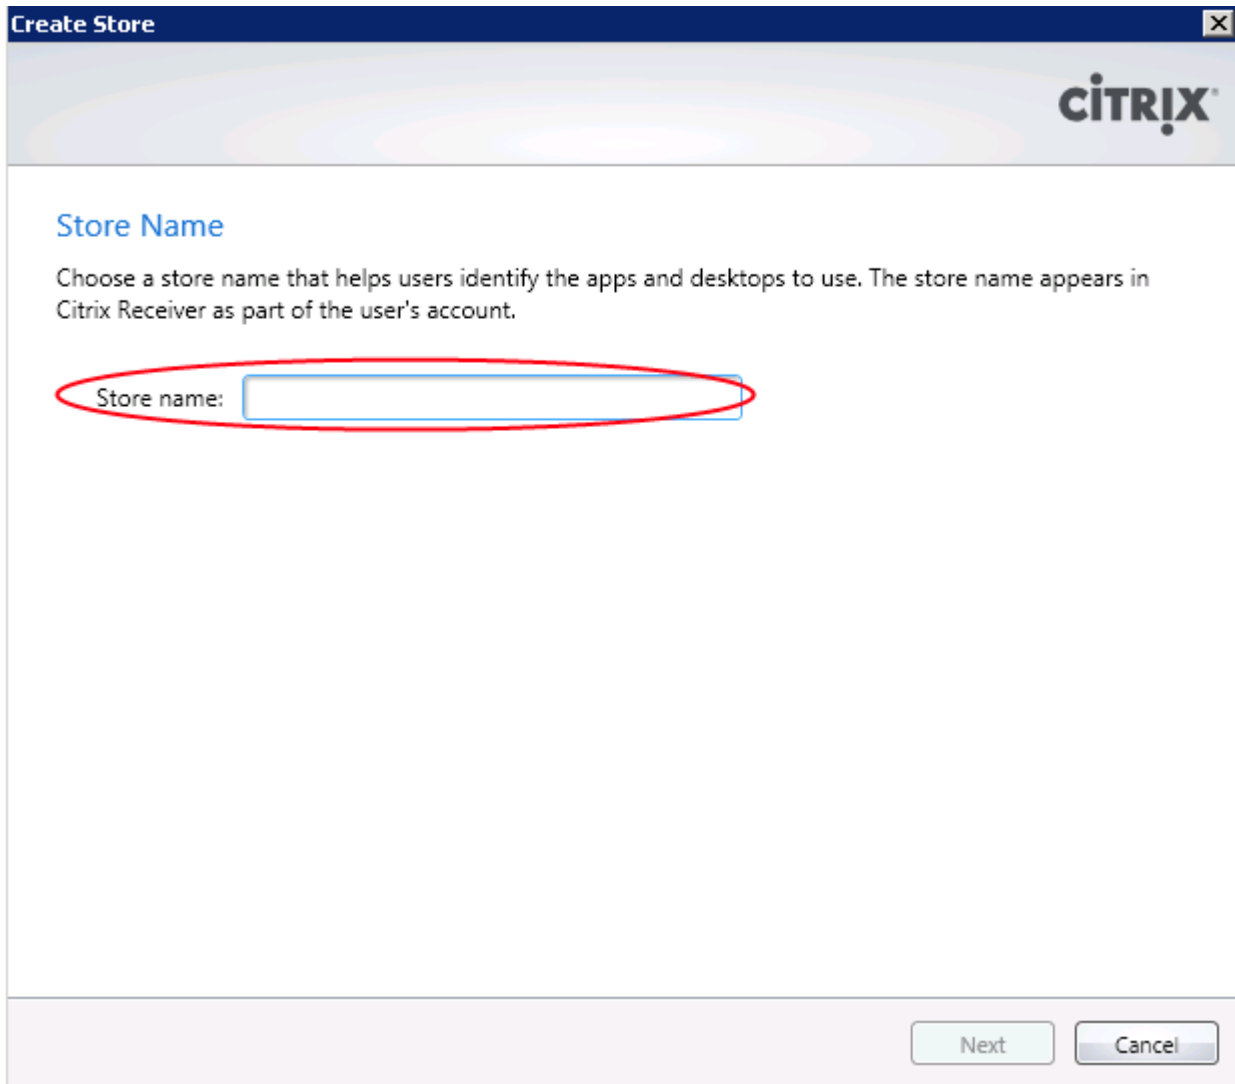
This server must be part of an established load-balancing system. To ensure that users have the same apps on each of their devices, a database server is also required for subscription services.

Hostname (load balancer): 

Database server:

Database name:

7. On the Store Name page, specify a name for your store and click Next.



Create Store

CITRIX

Store Name

Choose a store name that helps users identify the apps and desktops to use. The store name appears in Citrix Receiver as part of the user's account.

Store name:

Next Cancel

StoreFront stores enumerate and aggregate desktops and applications from XenDesktop sites, XenApp farms, and AppController, making these resources available to users.

8. On the Delivery Controllers page, list the XenDesktop, and XenApp Server names providing the resources that you want to make available in the store. Click Add.
9. Use this section to add your XenDesktop Delivery Controllers by name. Select transport type as HTTP port 80. Repeat the same step for XenApp.
10. Click OK.

Add Delivery Controller [X]

Display name:

Type:

- CloudGateway Enterprise
- XenApp
- XenDesktop

Servers:
(in failover order)

XA-01	▲
-------	---

Transport type: ⓘ SSL Relay port:

Port:

11. With both of your Delivery Controller resources added, click Next.

Create Store ✕

CITRIX

Delivery Controllers

Specify the delivery controllers and servers for this store.

Delivery Controllers:

Name	Type	Servers
XA	XenApp	XA-01
XD	XenDesktop	XD-01

12. Select None as Remote access type. Click Create.

Create Store ✕

CITRIX

Remote Access

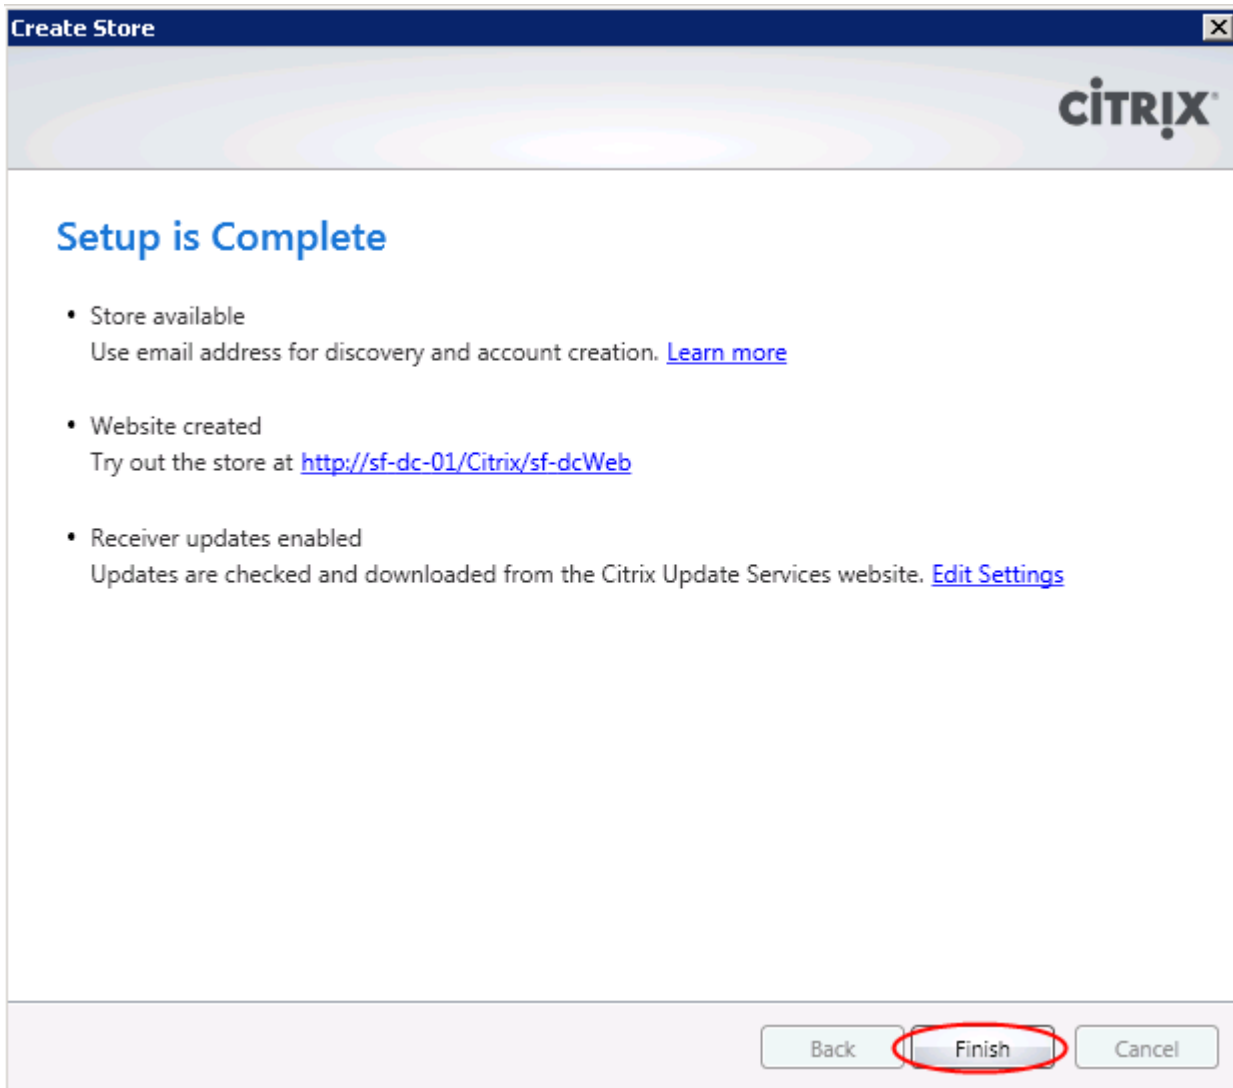
If you have users connecting from external networks, add gateway servers to provide remote access.

Remote access: None No VPN tunnel Full VPN tunnel

Access Gateways:

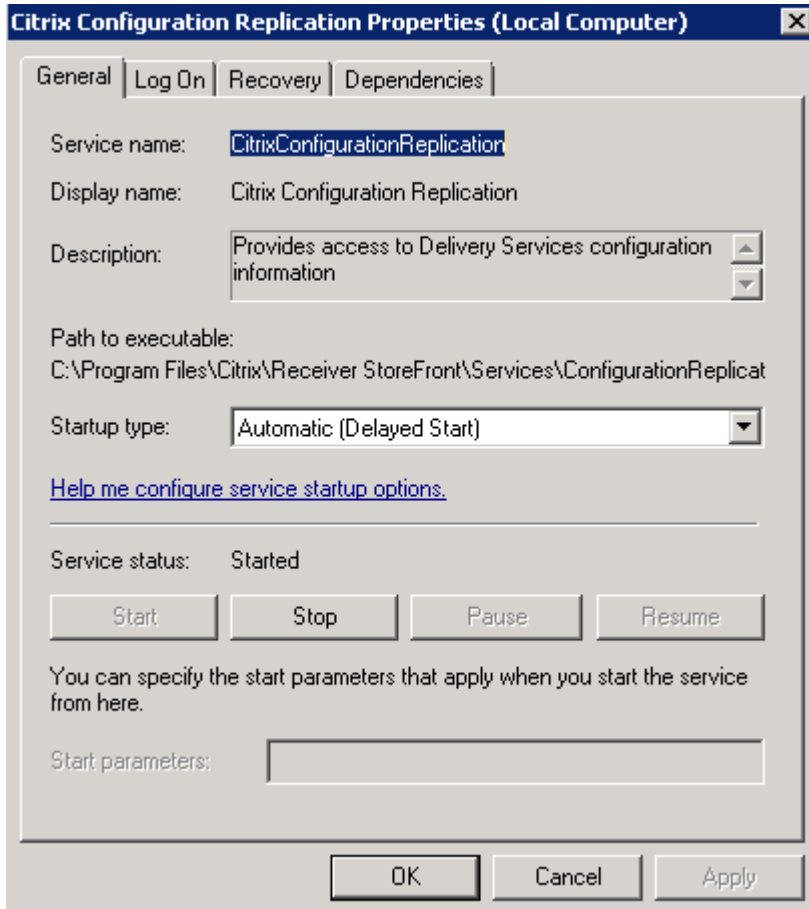
Default gateway:

13. Click Finish to complete the process.



14. Configure the following services for Automatic (Delayed Start)

- Citrix Configuration Replication



- Citrix Credential Wallet

Citrix Credential Wallet Properties (Local Computer) [X]

General | Log On | Recovery | Dependencies

Service name:

Display name: Citrix Credential Wallet

Description: [Up] [Down]

Path to executable:
C:\Program Files\Citrix\Receiver StoreFront\Services\CredentialWallet\Citri:

Startup type: [Down]

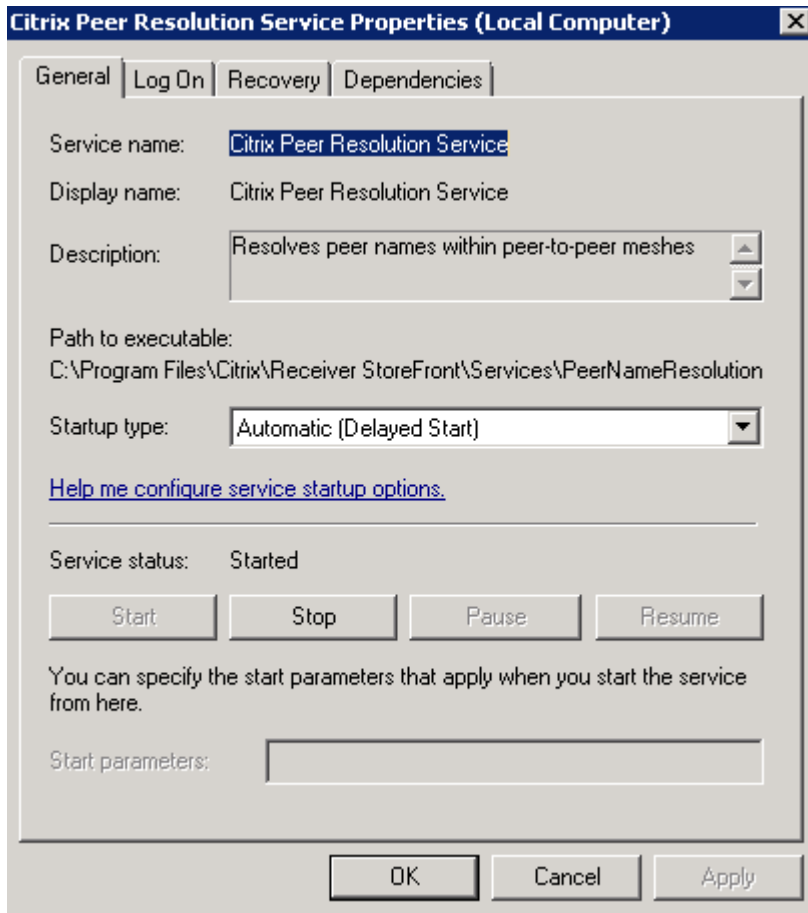
[Help me configure service startup options.](#)

Service status: Started

You can specify the start parameters that apply when you start the service from here.

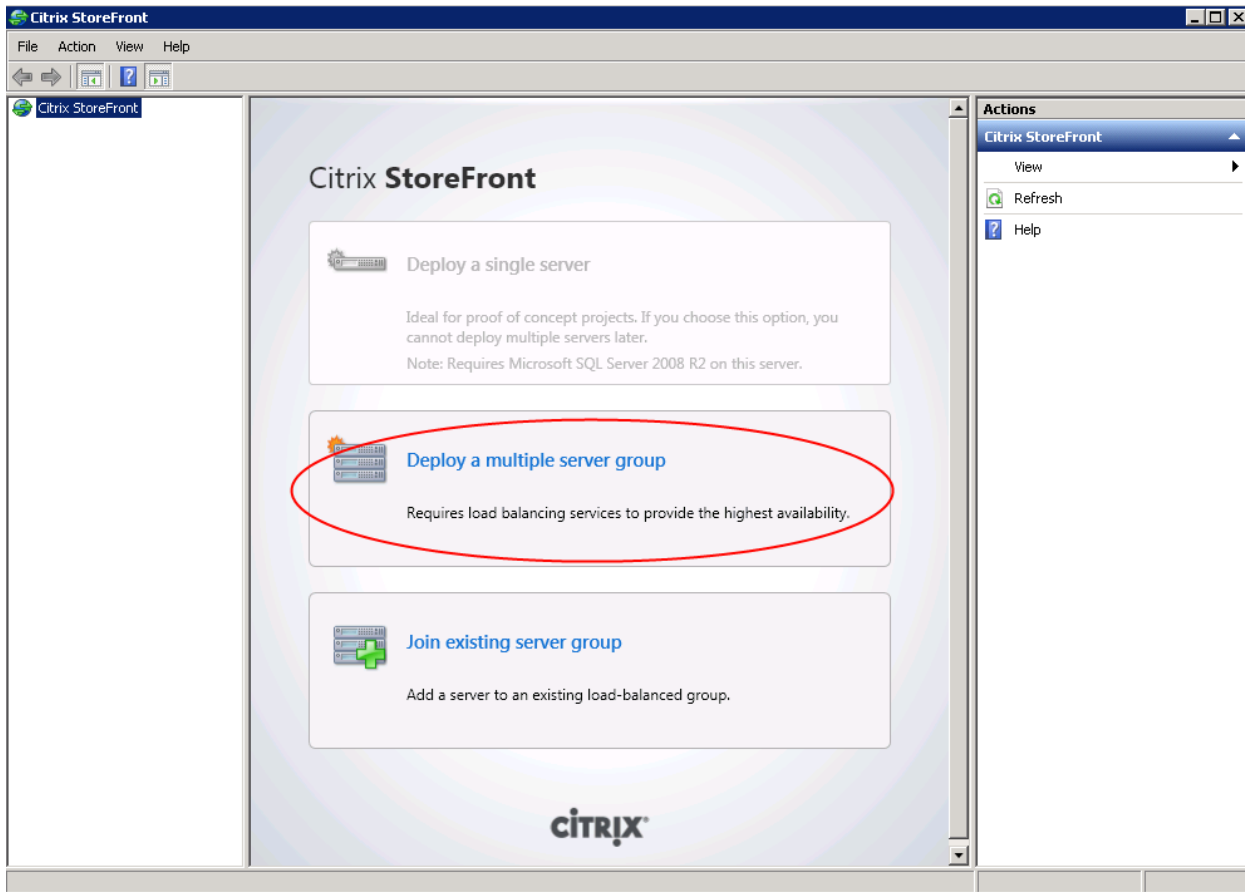
Start parameters:

- Citrix Peed Resolution Service

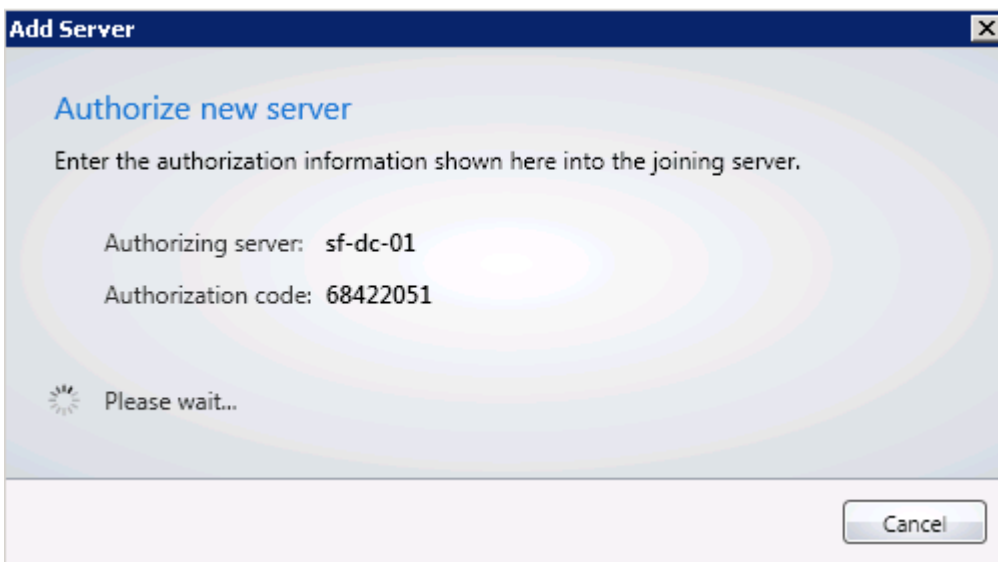
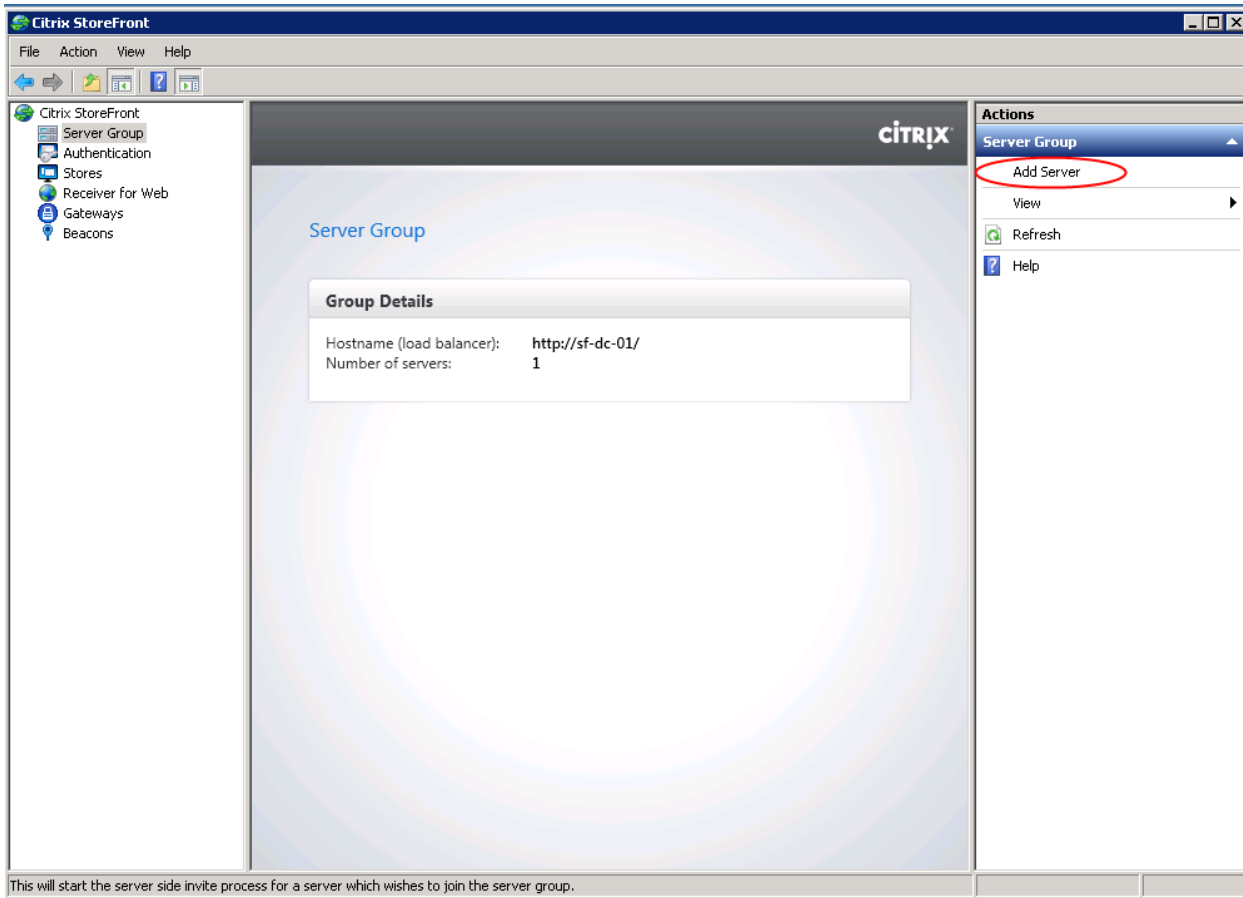


6.12.4 Join an Existing Server Group

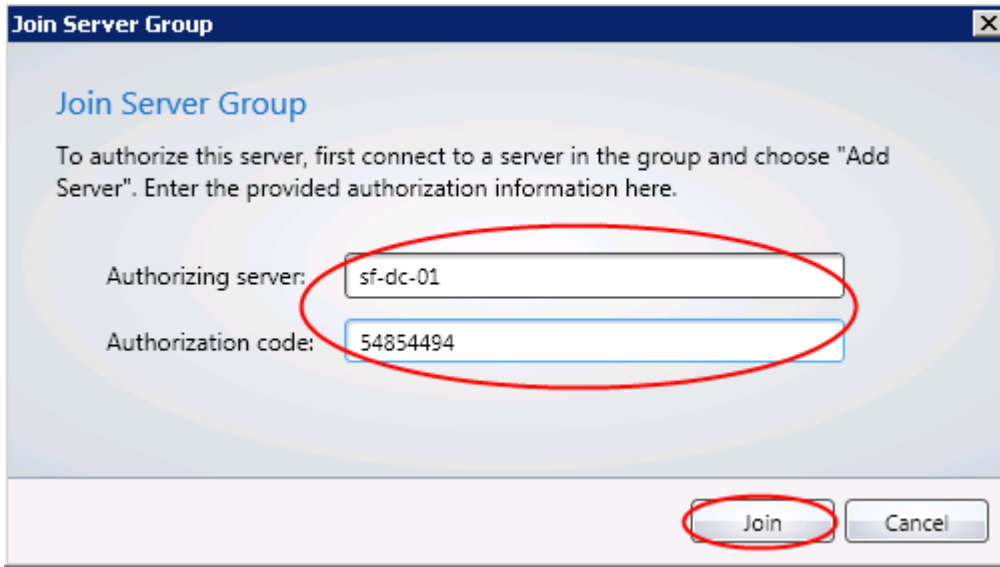
1. If the Citrix StoreFront management console is not already open after installation of StoreFront, click Start > All Programs > Citrix > Citrix StoreFront.
2. In the results pane of the Citrix StoreFront management console, click Join existing server group.



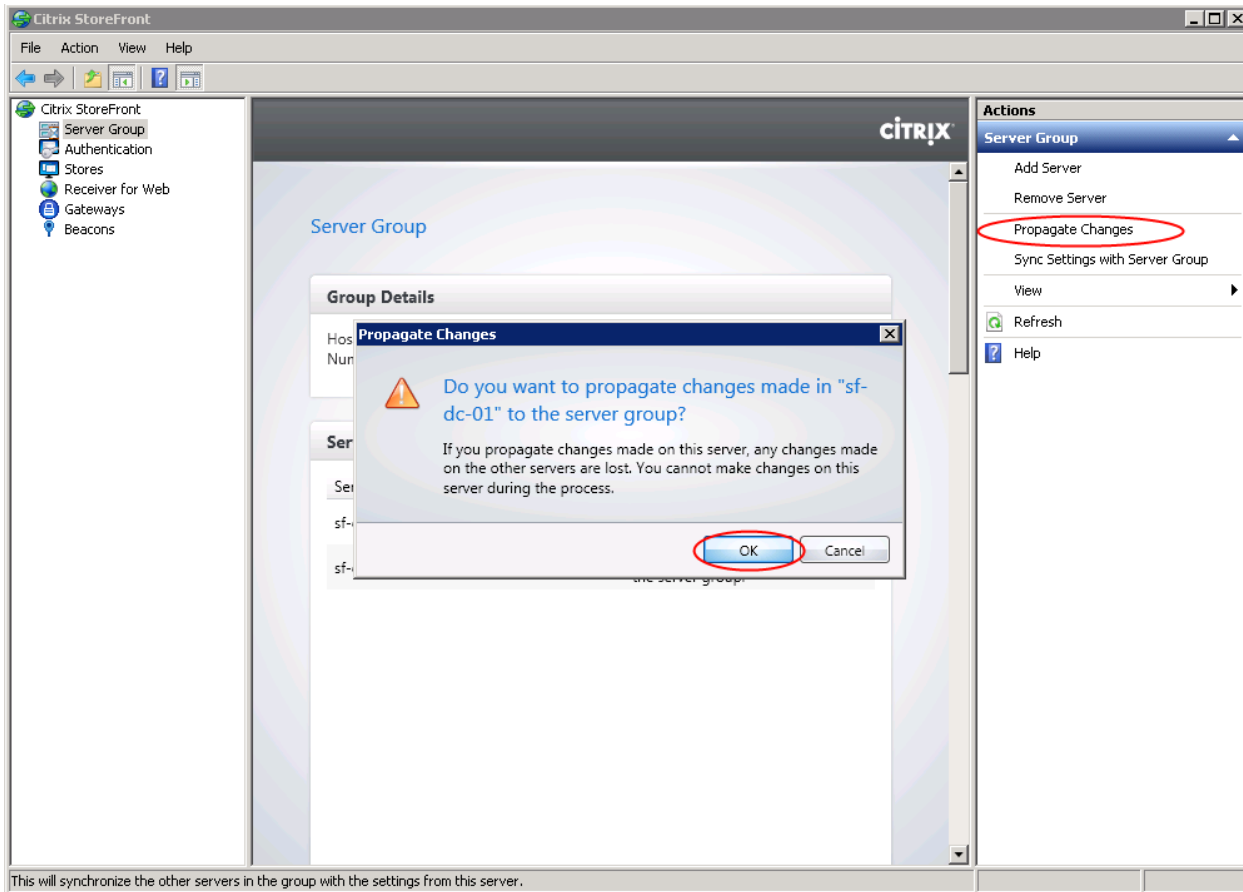
3. Log on to the primary server in the StoreFront deployment that you wish to join and open the Citrix StoreFront management console. Select the Server Group node in the left pane of the console and, in the Actions pane, click Add Server. Make a note of the authorization code that is displayed.



- Return to the secondary server and, in the Join Server Group dialog box, specify the name of the primary server in the Authorizing server box. Enter the authorization code obtained from that server and click Join.



5. When the new server has joined the deployment, return to the primary server and, in the left pane of the Citrix StoreFront management console, select the Server Group node. In the Actions pane, click Propagate Changes.
6. In the Propagate Changes dialog box, click OK.





The configurations of all the secondary servers in the deployment, including the newly added server, are updated to match the configuration of the primary server.

The new secondary server is added to your deployment and all servers in the group are updated with details of the new server.

7. Repeat these steps to add the third StoreFront Server.

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, XenDesktop 5.6 FP1 in combination with Provisioning Services 6.1 managed the VDI environment in this project.


The next 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 18. Citrix XenDesktop 5.6, Citrix XenApp 6.5 and Provisioning Server 6.1 Logical Diagram



Summary of the Environment:

- 30 ESXi 5.1.0 (838463-custom-Cisco-2.1.0.3) on B200 M3 divided into 4 clusters
- 20 ESXi-5.0.0-on B250 (build 469512-custom-Cisco-2.0.1d) into 1 cluster for VSI launchers
- B200 M3 ESXi 5.1.0 clusters were divided into the following:
 - 5 Hosts assigned to Common Infrastructure
 - 7 Hosts assigned to stream 1000 Hosted VM Desktops on Local Tier 0 storage
 - 7 Hosts assigned to stream 1000 Hosted VM Desktops on Shared NFS storage using Virtual Private Disks (PvD)
 - 11 Hosts assigned to stream 88 Virtual XenApp 6.5 Machines capable of 2000 Hosted Shared Desktop sessions
- 4 XenDesktop 5.6 FP1 Delivery Controllers
- 88 XenApp 6.5 Streamed Virtual Machines
- 10 Provisioning Server 6.1 Hotfix 61E015 Server for Virtual Desktops broken between 2 Sites for XenDesktop and XenApp
- Citrix UPM profile Management 4.1.1
- 2 VMware vCenter 5.1.0 build 838463 with Heartbeat 6.5
- 220 VSI Launchers
- 1 Citrix Licensing Server
- 3 Citrix StoreFront Servers 12.12.17357
- 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
- 1 NetScaler VPX (virtual) NS10.0 build 71.6nc configured with One-Arm Single Subnet Topology



Storage on EMC VNX 7500:

- 64 12 GB Fibre Channel Boot LUN
- 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
- 1 250 GB NFS mount for centralized logging
- 4 1.5 TB NFS mounts for Hosted VM Desktop Write Cache
- 4 1 TB NFS mounts for Hosted VM Desktop Personal vDisks Cache
- 2 1.5 TB NFS mounts for XenApp 6.5 Virtual Machines Write Cache

The following tables provide details on components configuration.

Desktop Virtualization			
VMware ESXi 5.1 Hosts			
Hardware:	Cisco B-Series Blade Servers	Model:	B200 - M3
OS:	VMware ESXi 5.1	RAM:	256GB
CPU:	2X Intel Xeon E5-2680 @2.7GHz 8-Core CPUs	Network:	Cisco UCS VIC 1240
Disk:	(Boot LUNs)	Disk	2 X 300 GB SSDs

Desktop Infrastructure			
VMware ESXi 5.1 Hosts			
Hardware:	Cisco B-Series Blade Servers	Model:	B200 - M3
OS:	VMware ESXi 5.1	RAM:	128 GB
CPU:	2X Intel Xeon E5-2665 @2.4GHz 8-Core CPUs	Network:	Cisco UCS VIC 1240
Disk:	(Boot LUNs)		

Citrix Provisioning Server 6.1			
Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	16GB
CPU:	4vCPUs	Network:	1xE1000e NIC for Management;1xVMXNET



			3 NIC for streaming
Disk:	40GB		

Citrix XenDesktop 5.6 Delivery Controllers

Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	4GB
CPU:	1vCPU	Network:	1xVMXNET 3 NIC
Disk:	40GB		

Vmware vCenter Servers

Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	16GB
CPU:	4vCPUs	Network:	1xVMXNET 3 NIC
Disk:	40GB		

Microsoft SQL Server 2008 R2 for XenDesktop

Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	4GB
CPU:	2vCPU	Network:	1xVMXNET 3 NIC
Disk:	40GB		

Microsoft SQL Server 2008 R2 for vCenter

Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	80GB
CPU:	8vCPU	Network:	1xVMXNET 3 NIC
Disk:	40GB		



StoreFront Server			
Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	8GB
CPU:	8vCPU	Network:	1xVMXNET 3 NIC
Disk:	40GB		

Citrix XenApp 6.5			
Hardware:	Virtual Machine	Model:	
OS:	Windows 2008 R2	RAM:	12GB
CPU:	4vCPU	Network:	1xVMXNET 3 NIC
Disk:	40GB		

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

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

7.2 Citrix User Profile Management Servers

This section explains the installation and configuration of the profile cluster. The installation has the following:

- Two clustered virtual machines
- A highly available file share

7.2.1 Setting up a Highly Available Profile Share

Installing File Services

To set up a highly available profile share:

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.

7.2.2 Setting Up a Two Node Citrix User Profile Server Cluster

Two LUNs will be used for Quorum and file server clustered disk. Please follow the steps below to complete the cluster 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.

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.

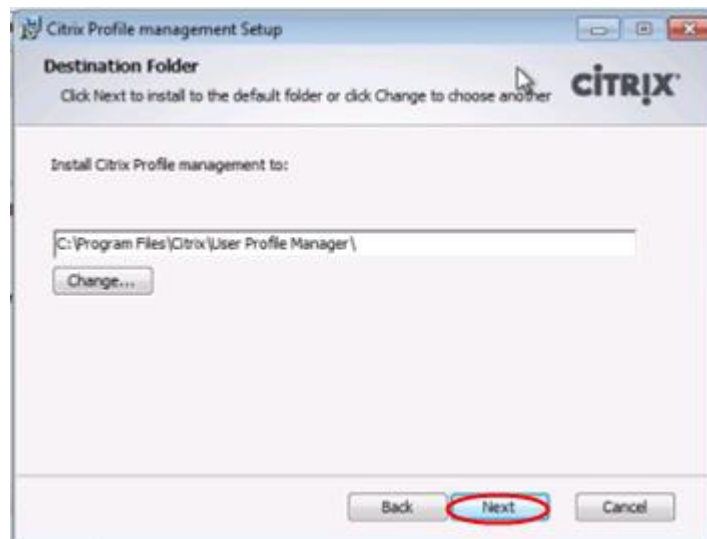
7.2.3 Install User Profile Manager

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.

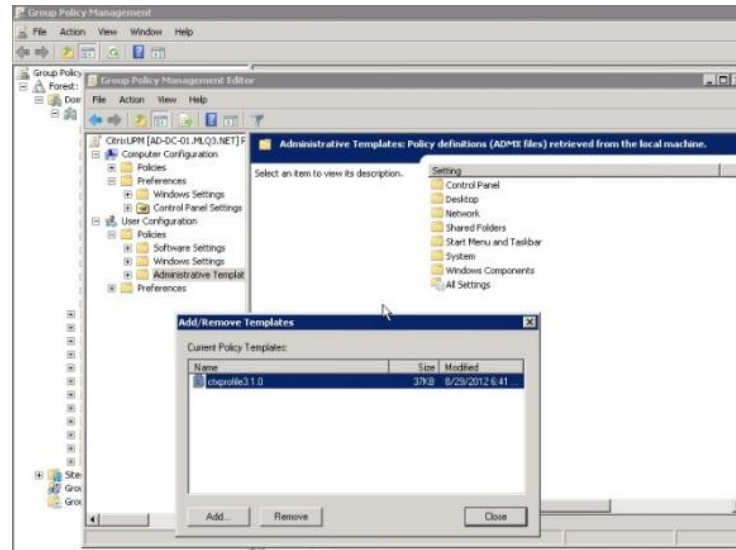


3. Click Next.
4. Click Finish after installation is complete.

7.2.4 Create a GPO for UPM

1. Create a GPO linked to users OU (organizational Unit).
2. Add the Citrix UPM administrative template.
3. Edit the new GPO, browse to User Configuration > Policies > Administrative Template

4. Right-click on “Administrative Template” and select Add/Remove Template.
5. Click Add.
6. Browse to the location of the template file provided with UPM installation files.



7. Configure the following settings under Administrative templates | Citrix | Profile Management:
 - Enable Active write back
 - Enable Profile Management
 - Enter the absolute path for the location where the profiles will be stored. (An example of the syntax would be [\\upmshare\profiles\%username%](#))
8. Select Enable for Process logons of local administrators
9. Select Enable for the File system | Exclusion list – directories and enter the following information:
 - AppData\LocalLow
 - AppData\Roaming
 - \$Recycle.Bin
 - AppData\Local
10. Click Log Settings | Enable Logging and select Enable.
11. Click Profile handling | Delete locally cached profiles on logoff and select Disabled.
12. Click Local profile conflict handling.
13. Select If both local windows profile and Citrix Profile exist.



14. Select Delete local profile.
15. Click Streamed User Profiles.
16. 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

Create base Windows 7 SP1 32bit Virtual Machine

The Microsoft Windows 7 Enterprise 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 the new desktop virtual machines on the ESXi hosts.

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

Each Hosted VM desktop virtual machine was created with 1vCPU, 1.5GB of Memory, a 3.0 GB write cache disk and Page file set to 1.5GB

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

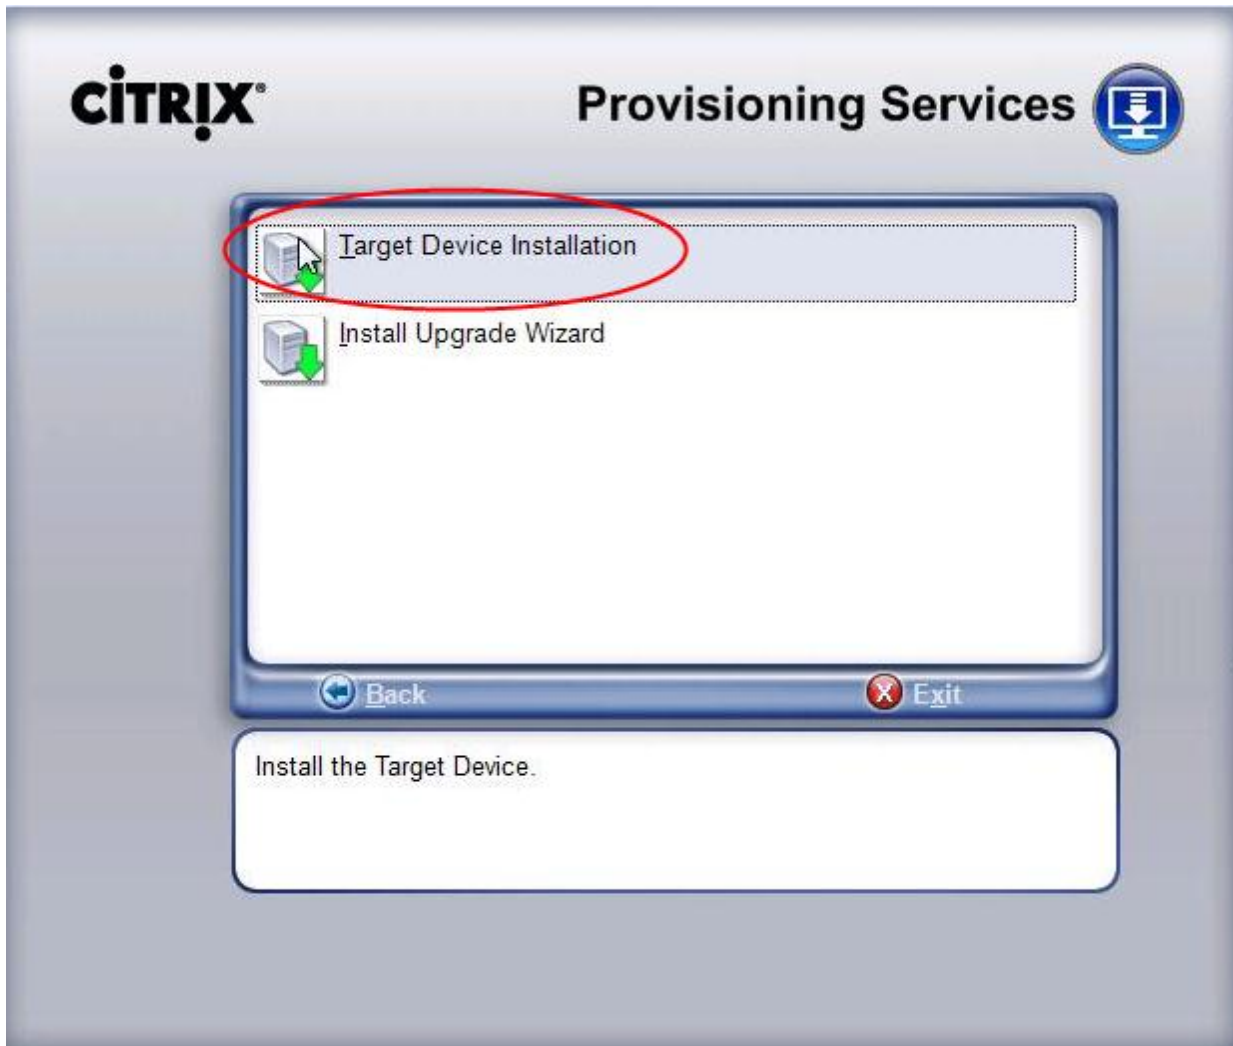
7.3.1 Create Base Windows 7 SP1 32bit Virtual Machine

1. Install Win7 32bit SP1 Enterprise
2. Install Office 2010 Professional with Run All From My Computer
3. Install Office 2010 Service Pack (most recent)
4. 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

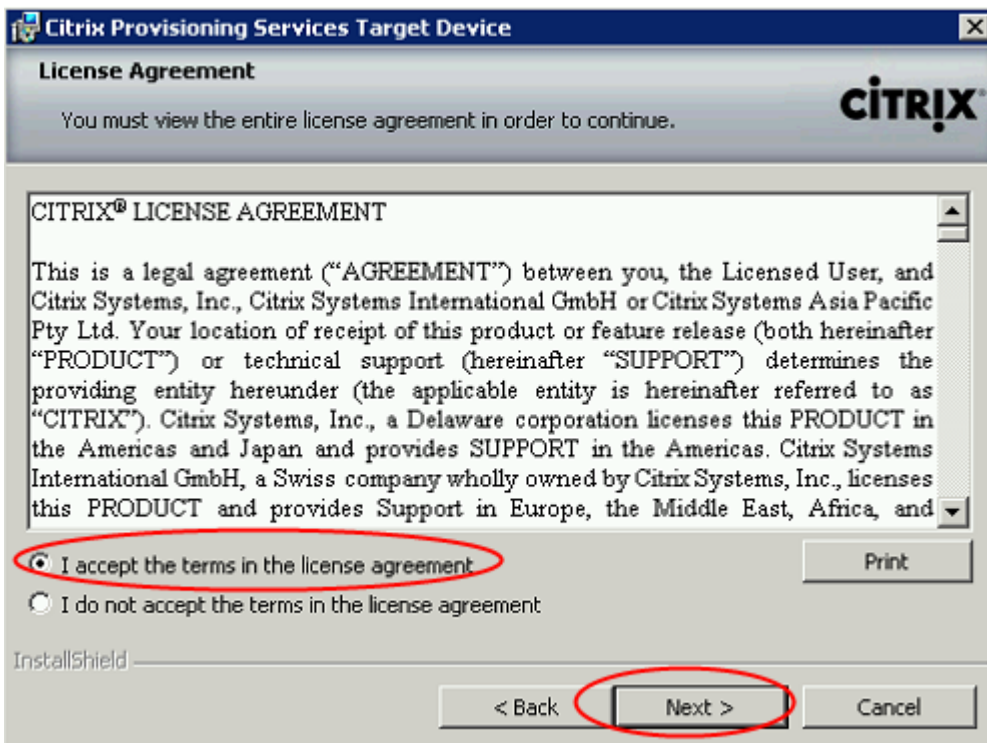
1. Launch the PVS Device executable, select Target Device Installation.



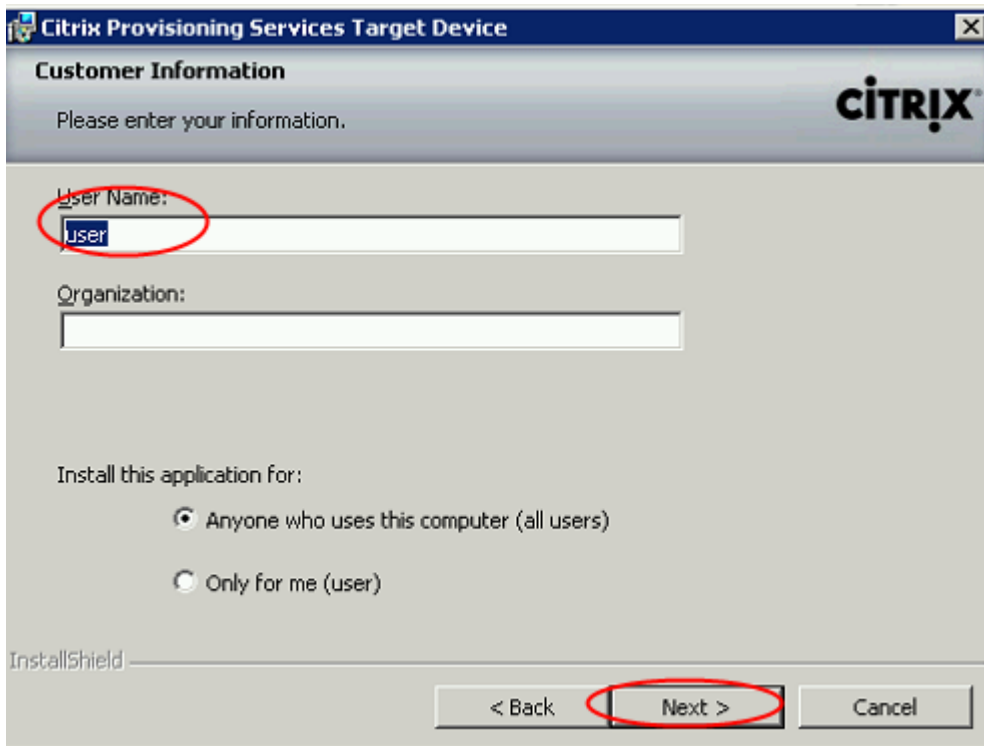
2. Click Next.



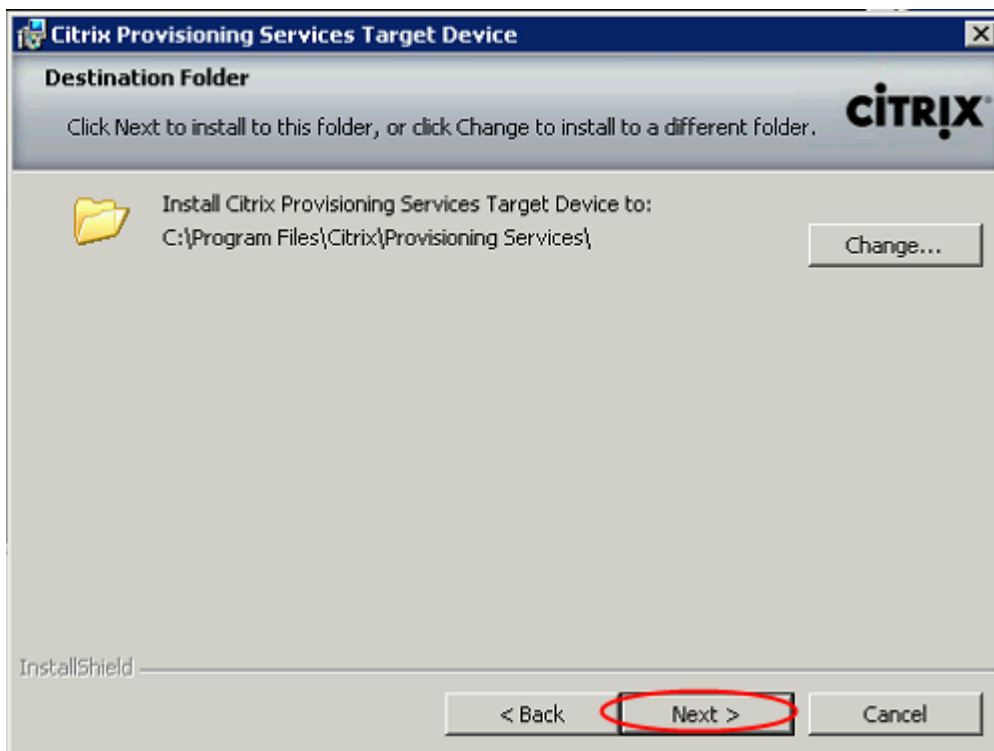
3. Accept the license agreement.
4. Click Next.



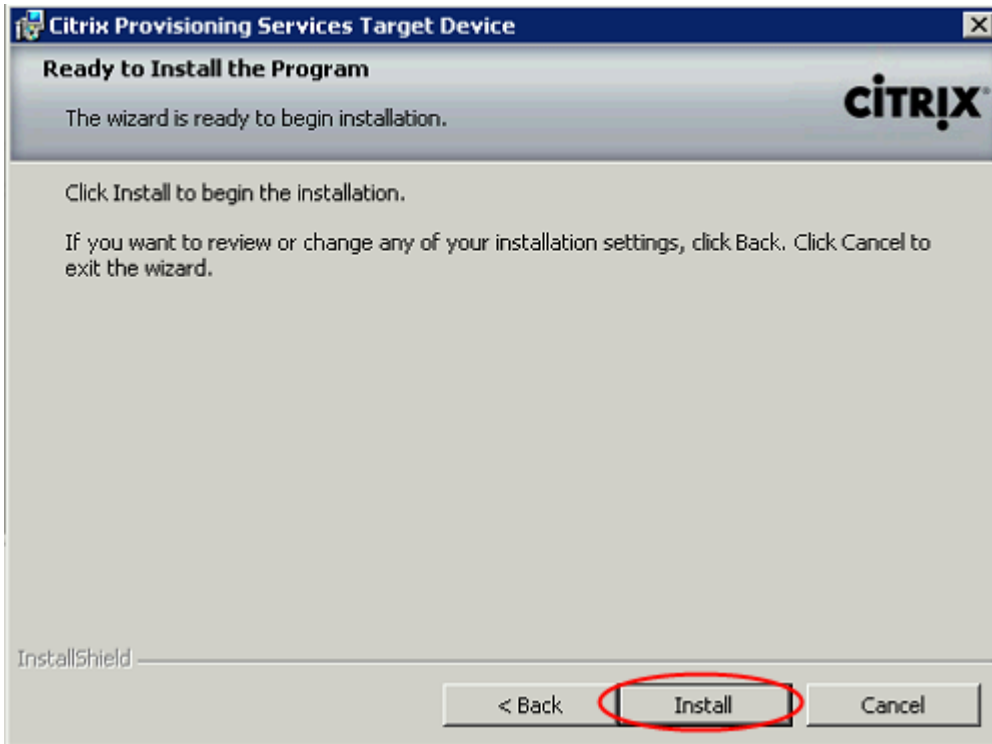
5. Enter in the customer information.
6. Click Next.



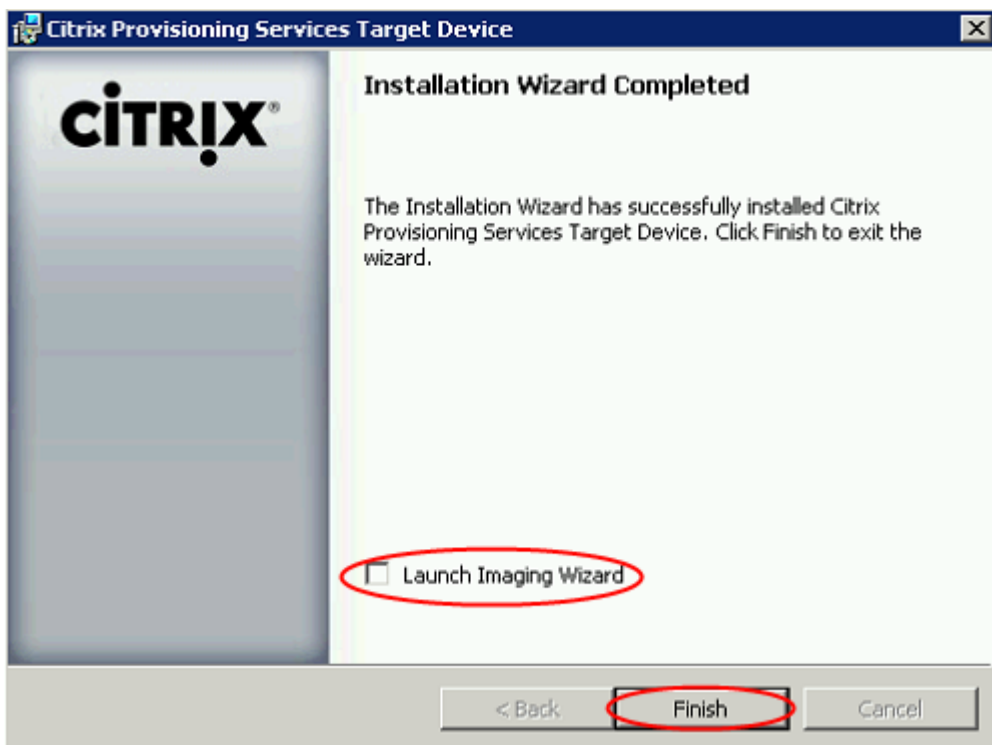
7. Choose the default installation location.
8. Click Next.



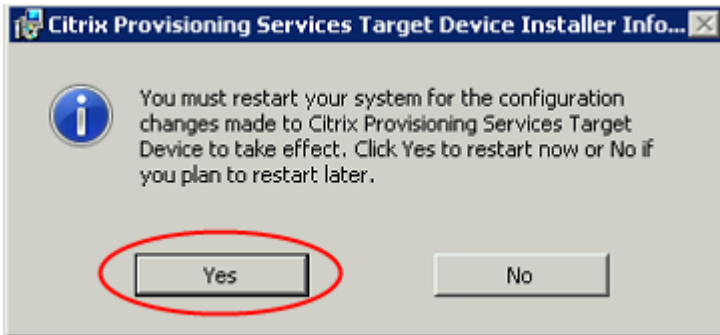
9. Click Install to begin the PVS Client installation process.



- 10. Uncheck Launch Imaging Wizard (This process will take place at a later point in the conversion process).
- 11. Click Finish.

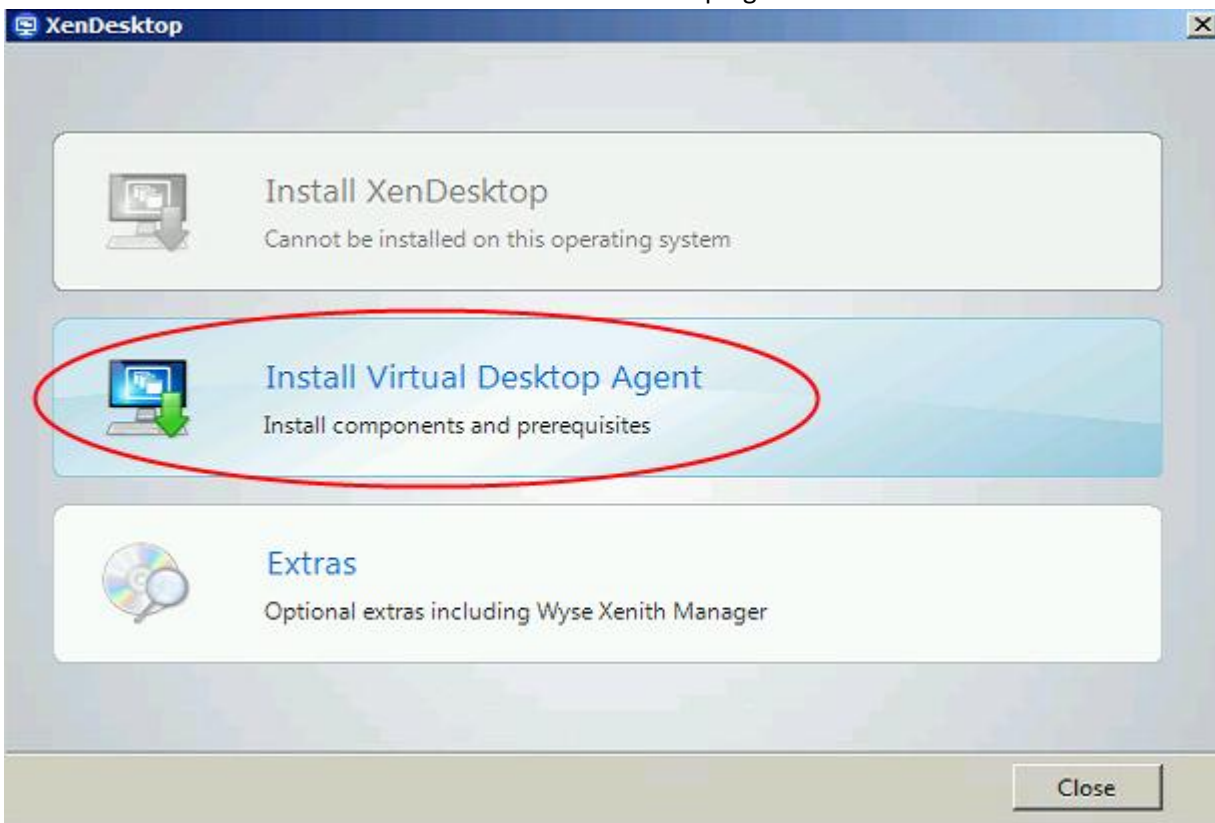


- 12. Click Yes to restart the virtual machine.

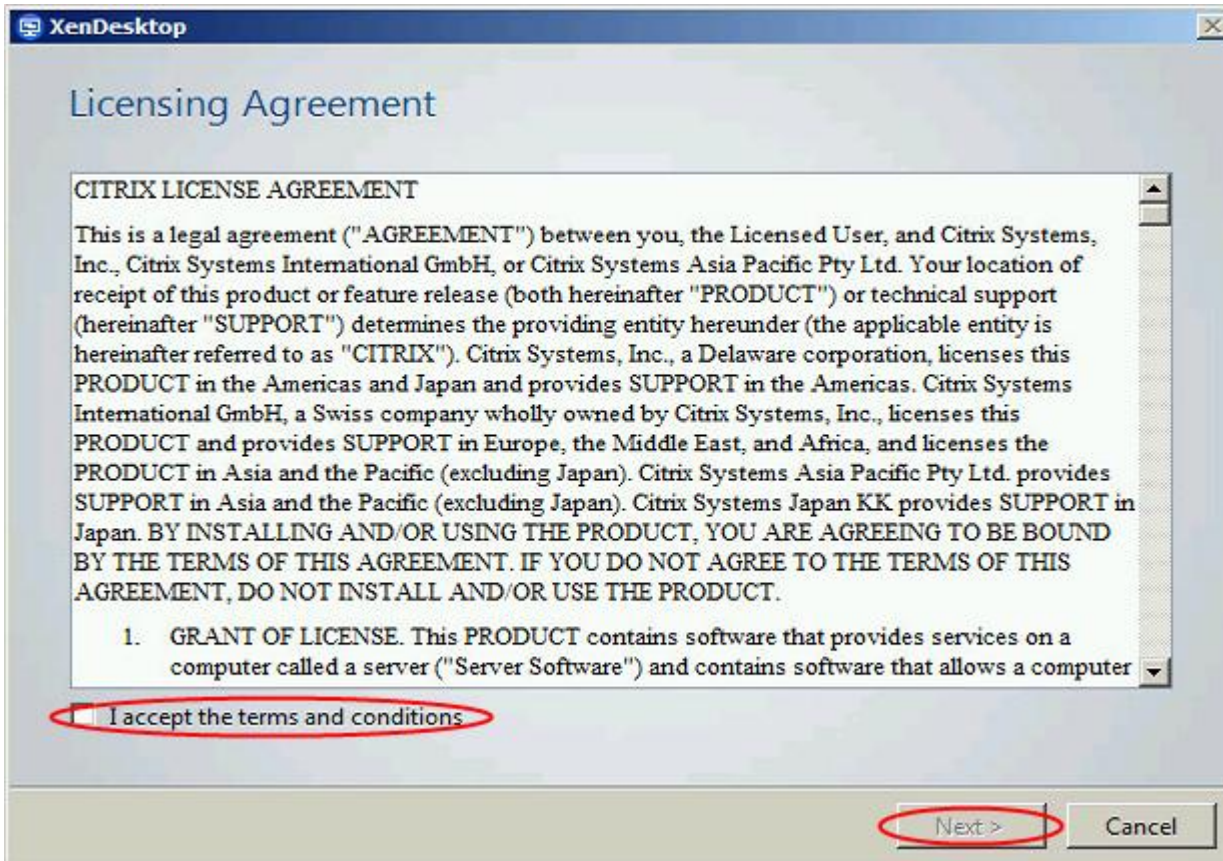


7.3.3 Add a XenDesktop 5.6 Virtual Desktop Agent

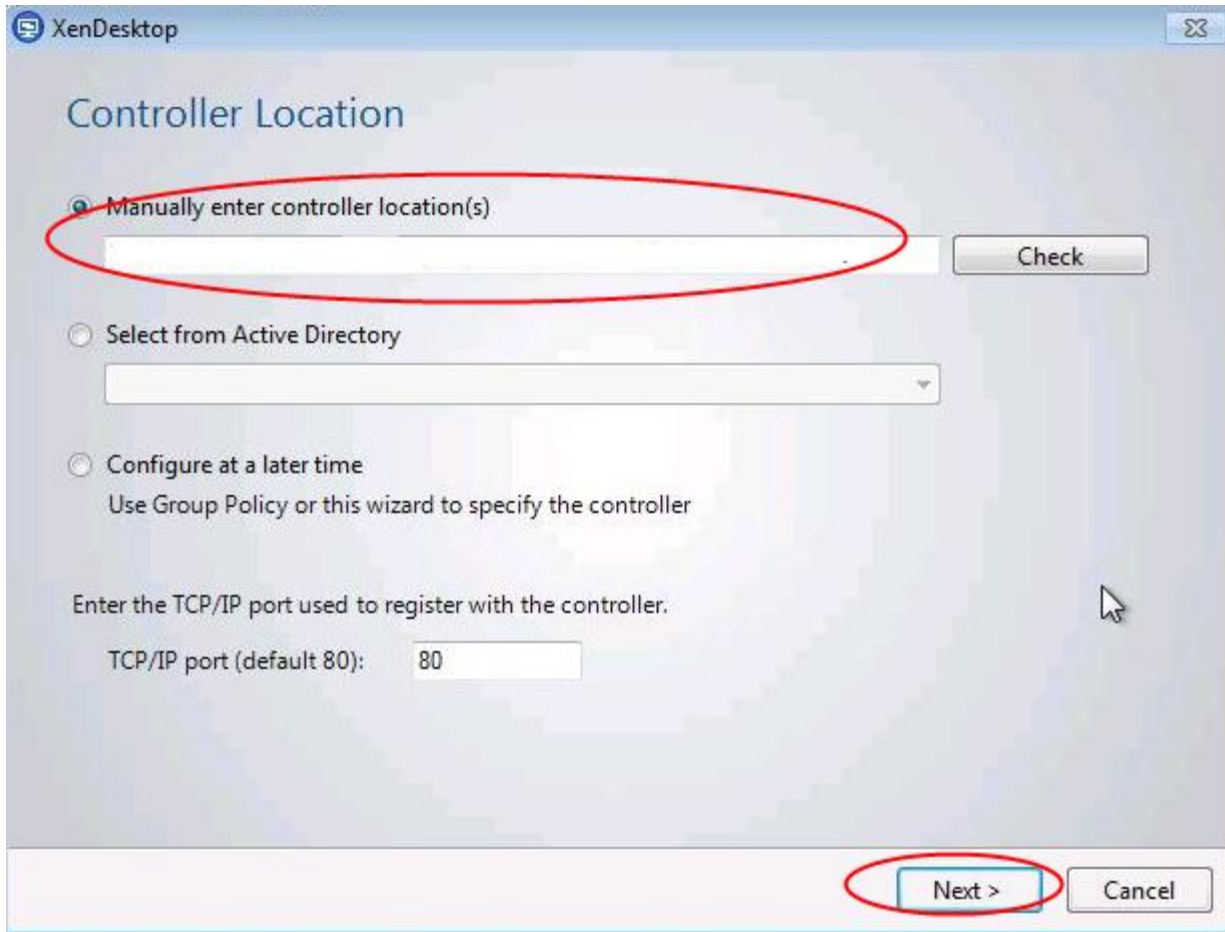
1. Copy the VDA executable to the local machine.
2. Launch executable and select Install Virtual Desktop Agent.



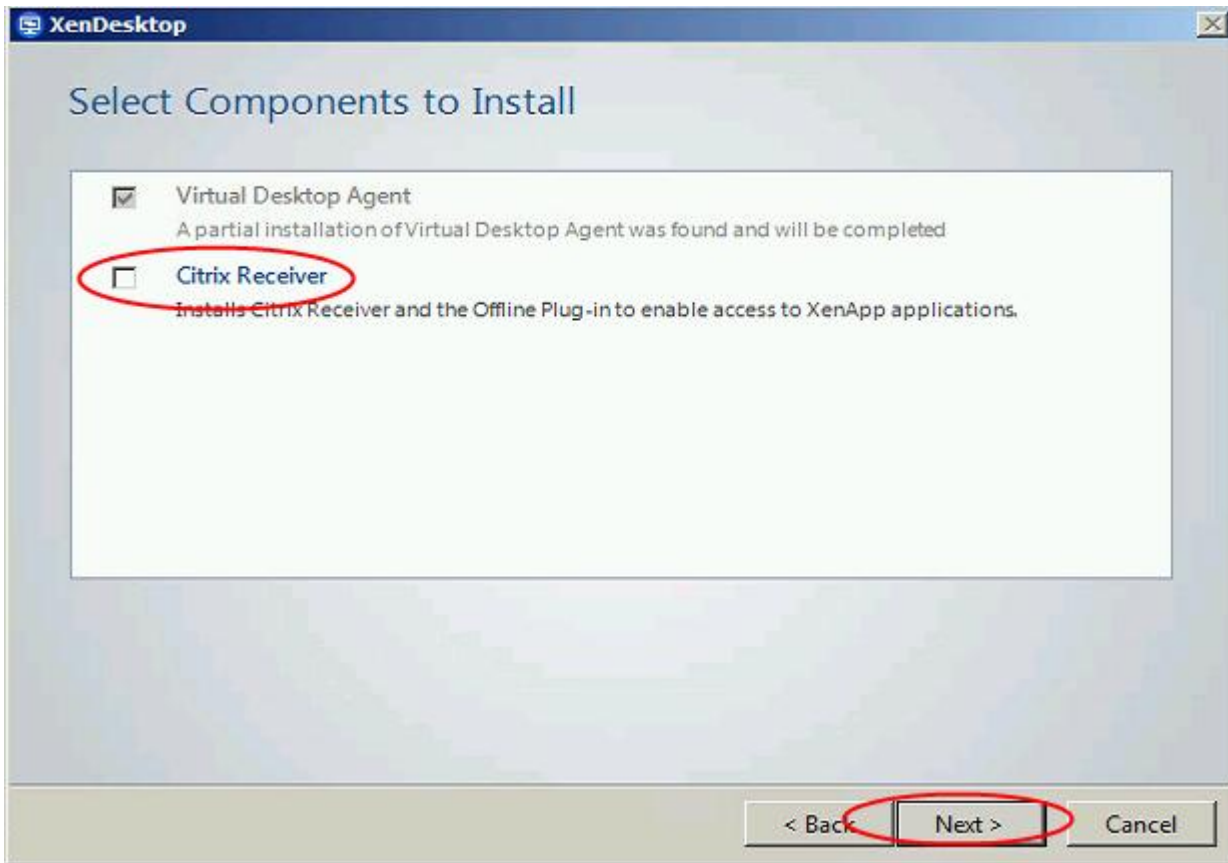
3. Accept terms and conditions of Licensing Agreement.



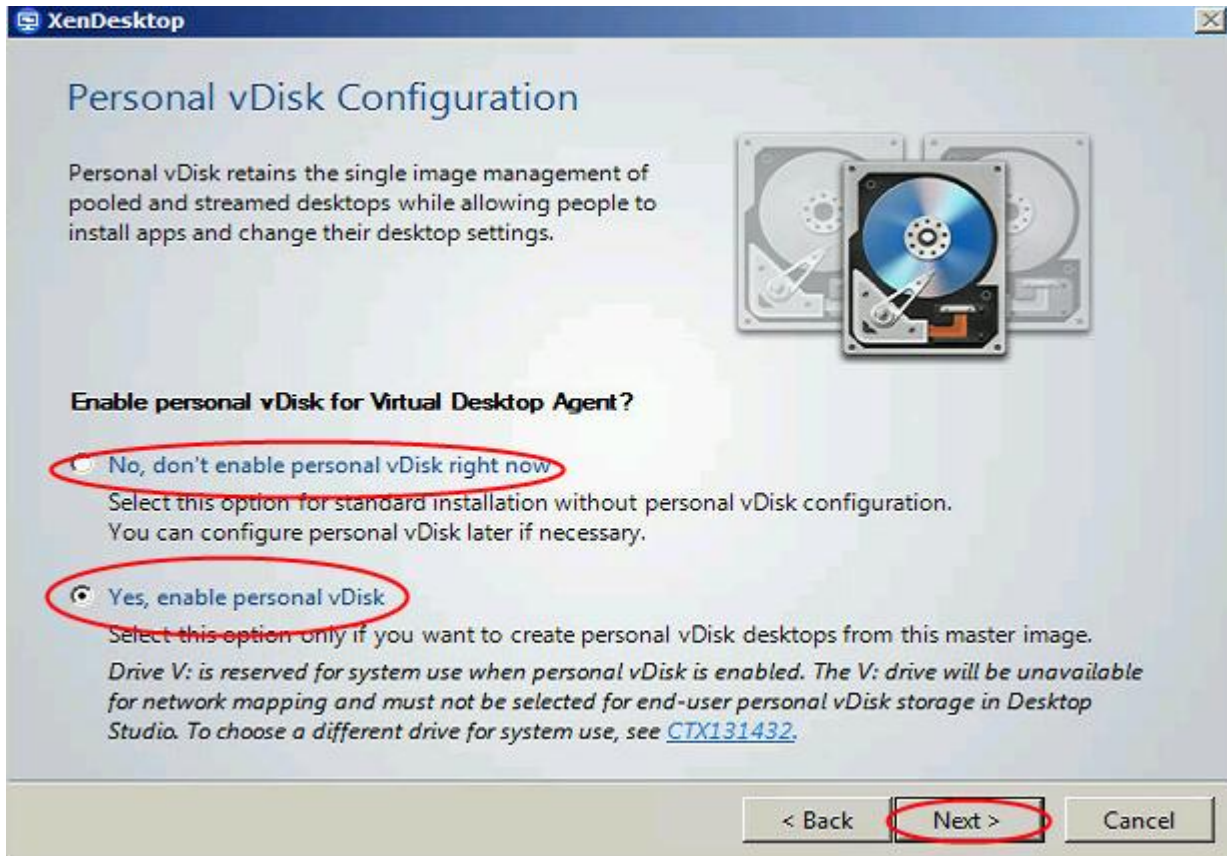
4. Select Advanced Install.
5. Use default installation paths.
6. Enter all XenDesktop Delivery Controller's FQDNs.



7. Click Next.
8. Do not install HDX or Citrix Receiver.



9. For PvD, select Yes, enable personal vDisk.
10. For non-PvD select No, don't enable Personal vDisk right now.




11. Click Next.
12. Use default settings. Leave checked: Optimize XenDesktop Performance, User Desktop Shadowing, Real Time Monitoring.



13. Click Next.
14. Click Finish.
15. Remove VDA Welcome Screen program from the Windows Startup folder.
16. Restart VM.
17. Log in and check the event log to ensure 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.

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.

The screenshot shows the 'Provisioning Services Imaging Wizard' window. The title bar reads 'Provisioning Services Imaging Wizard'. Below the title bar, the main heading is 'Connect to Farm' with the instruction 'Enter the name or address of a server in the farm to connect to.' The form is divided into two sections: 'Server information' and 'Credentials'. In the 'Server information' section, the 'Server' field contains 'xd-pvs-10.miq3.net' and the 'Port' field contains '54321'. In the 'Credentials' section, the radio button for 'Use my Windows credentials' is selected. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is circled in red.

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.

Provisioning Services Imaging Wizard

New vDisk

Enter the details for the new vDisk.

vDisk name:

Store:
Accessible by server: XD-PVS-10

VHD type:

VHD block size:

< Back **Next >** Cancel

13. Select KMS for Licensing Management.

Provisioning Services Imaging Wizard

Microsoft Volume Licensing

Choose if the vDisk is to be configured for Microsoft KMS or MAK volume license management.

None

Key Management Service (KMS)

Multiple Activation Key (MAK)

< Back **Next >** Cancel

14. Click Next.
15. Use default size image volumes.

Provisioning Services Imaging Wizard
Configure Image Volumes
 Define the size of each volume.

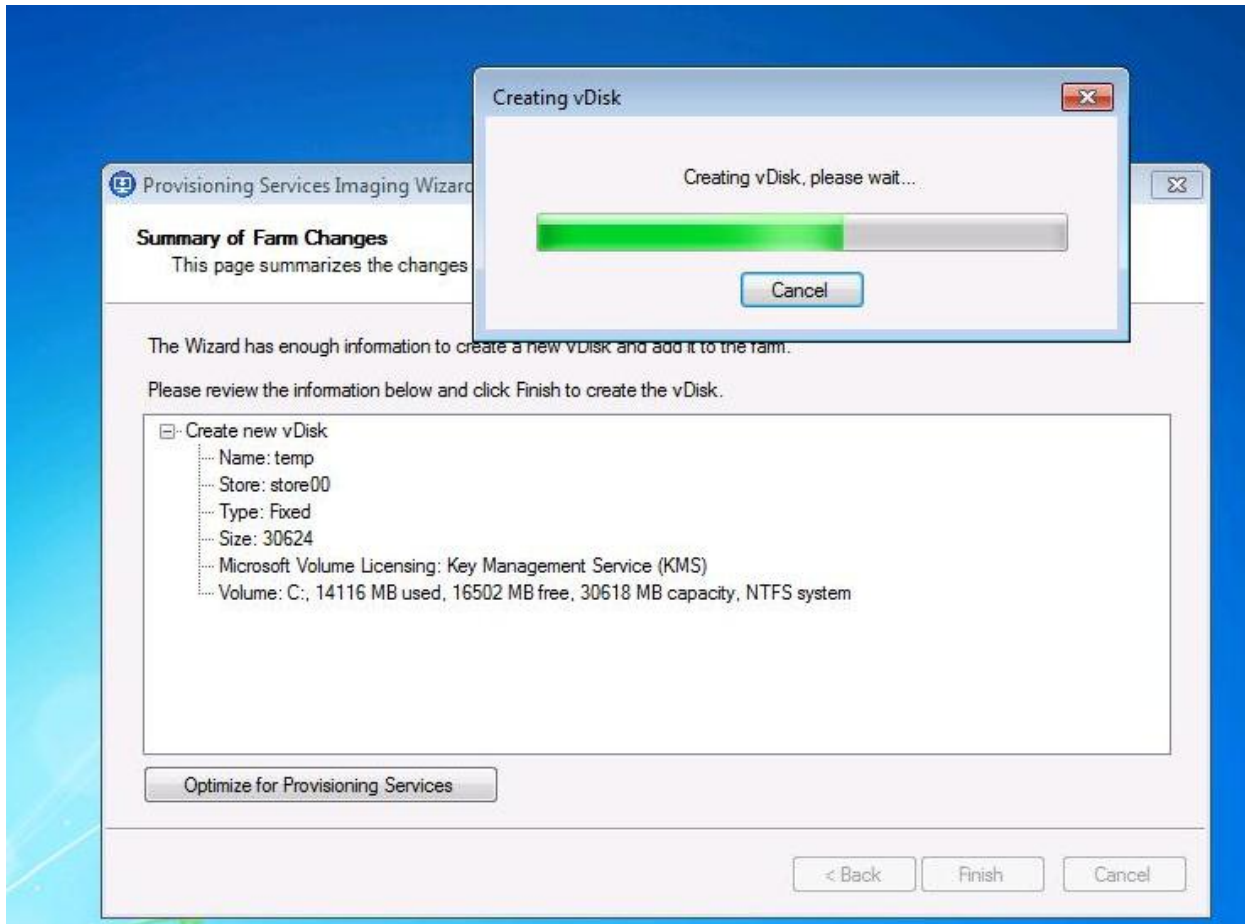
Source Volume	Used Space	Free Space	Capacity	File System
1 C: Boot	14116 MB 46 %	16502 MB 54 %	30618 MB	NTFS
2 None				
3 None				
4 None				

↓

Destination Volume	Used Space	Free Space	Capacity	File System
C: Boot	14116 MB 46 %	16502 MB 54 %	30618 MB	NTFS

vDisk	Allocated Space	Unallocated Space	Capacity
Summary	30618 MB 100 %	6 MB 0 %	30624 MB

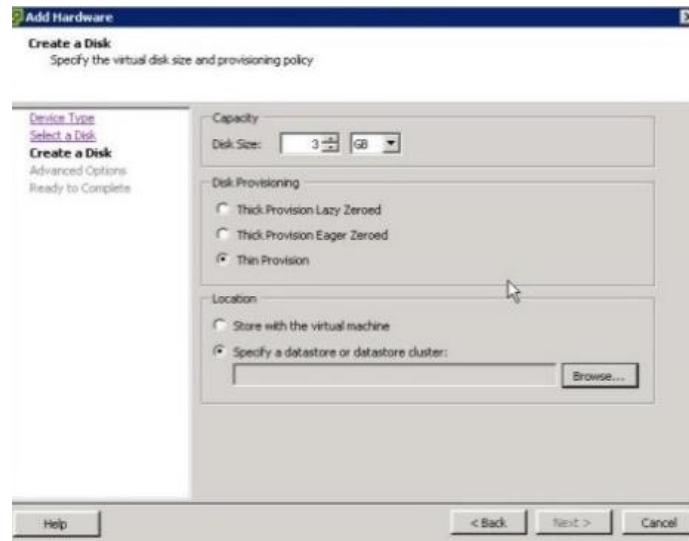
16. Click Next.
17. Click Finish to begin the vDisk creation.



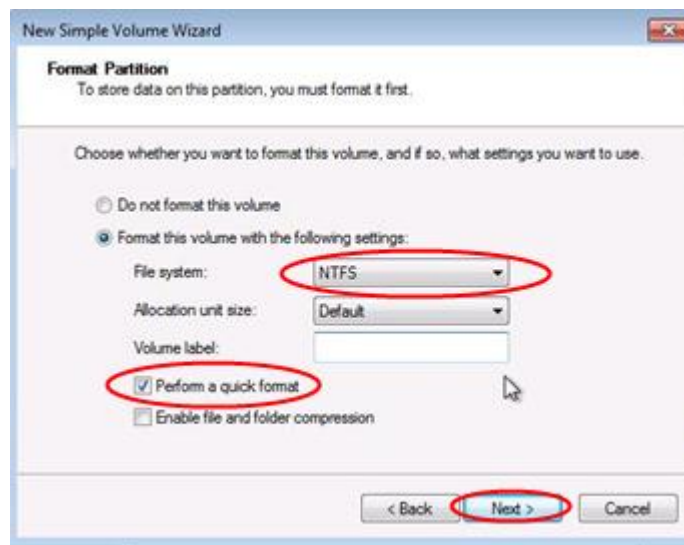
18. You will be prompted to reboot the source virtual machine. Prior to rebooting, go to the properties of the source virtual machine and change the boot options so that it performs a Network boot.
19. Click Yes to reboot the source virtual machine.
20. Logon as the same user account that was used at the beginning of this process.
21. When 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 Microsoft Windows Server 2008 R2 XenApp Golden Image Creation

Create a base Citrix XenApp Server Virtual Machine

The Microsoft Windows Server 2008 R2 64 bit Enterprise master or golden image with Citrix XenApp and 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 to provision 88 new streamed XenApp Server virtual machines on the ESXi hosts.



With Provisioning Server 6.1, the Streamed VM Setup Wizard was utilized.

Each virtual XenApp Server virtual machine was created with 4vCPUs, 12GB of Memory, a 12.0 GB write cache disk and configured with 8GB of static page file.

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

7.4.1 Create Base Windows Server 2008 R2 64bit Virtual Machine

1. Install Windows 2008 R2 64 bit
2. Install and configure XenApp (procedure in XenApp installation Section)
3. Install Office 2010 Professional with Run All From My Computer
4. Install Office 2010 Service Pack (most recent)
5. Windows Updates (Be sure not to install IE9. Use IE8)

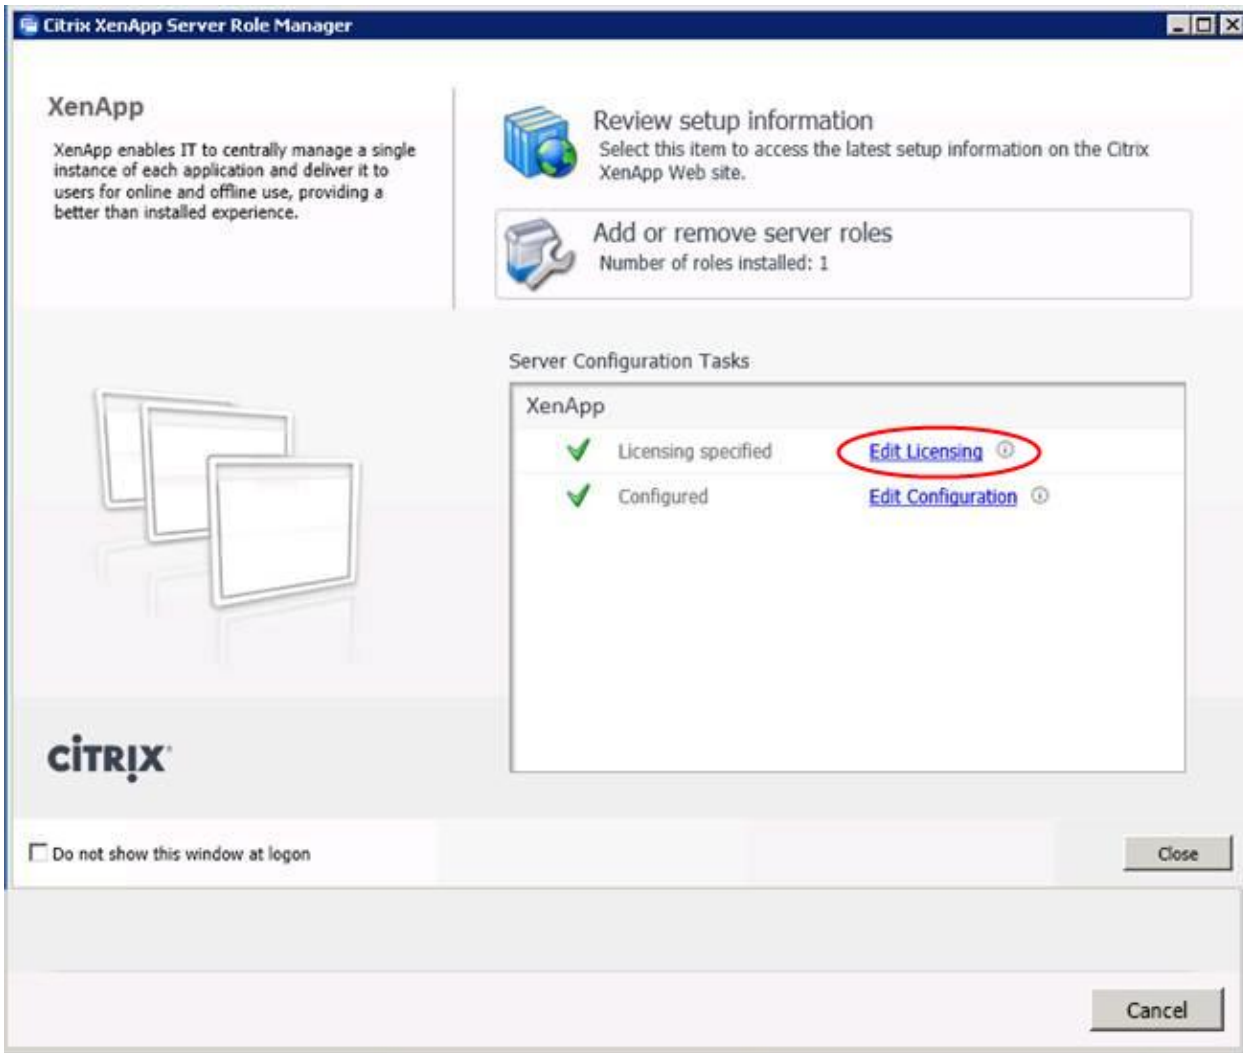
7.4.2 Add Provisioning Services Target Device Software

Note: The latest version of the Target Device Software can be installed directly by installing Hotfix CPVS61E015.

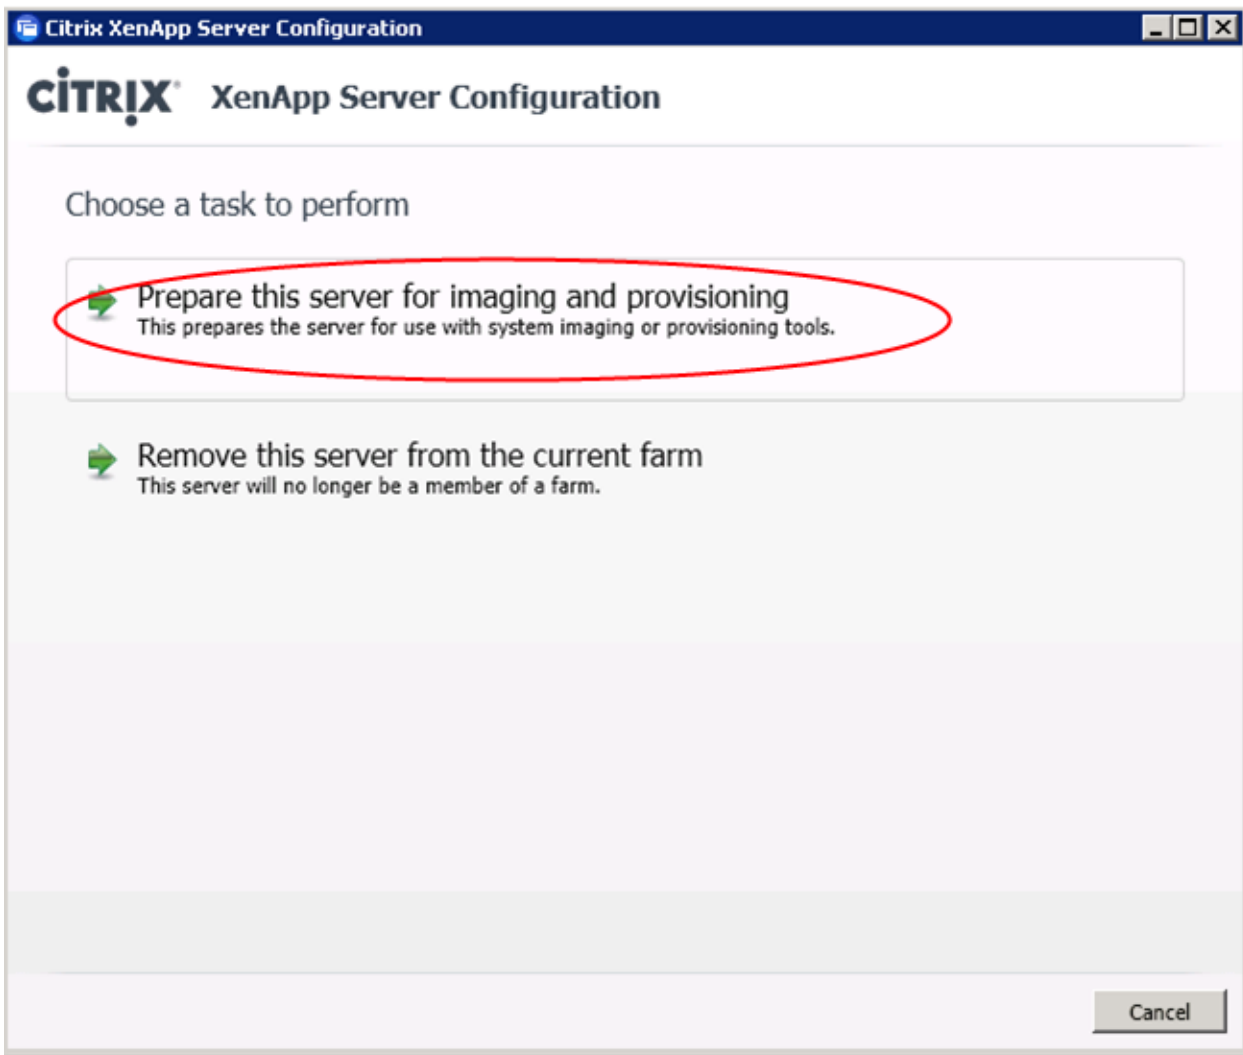
1. Launch the PVS Device executable and select Target Device Installation.
2. Click Next.
3. Accept the license agreement.
4. Click Next.
5. Enter in the customer information.
6. Click Next.
7. Choose the default installation location.
8. Click Next.
9. Click Install to begin the PVS Client installation process.
10. Uncheck Launch Imaging Wizard (This process will take place at a later point in the conversion process).
11. Click Finish.
12. Click Yes to restart the virtual machine.

7.4.3 Prepare XenApp Server for Imaging and Provisioning

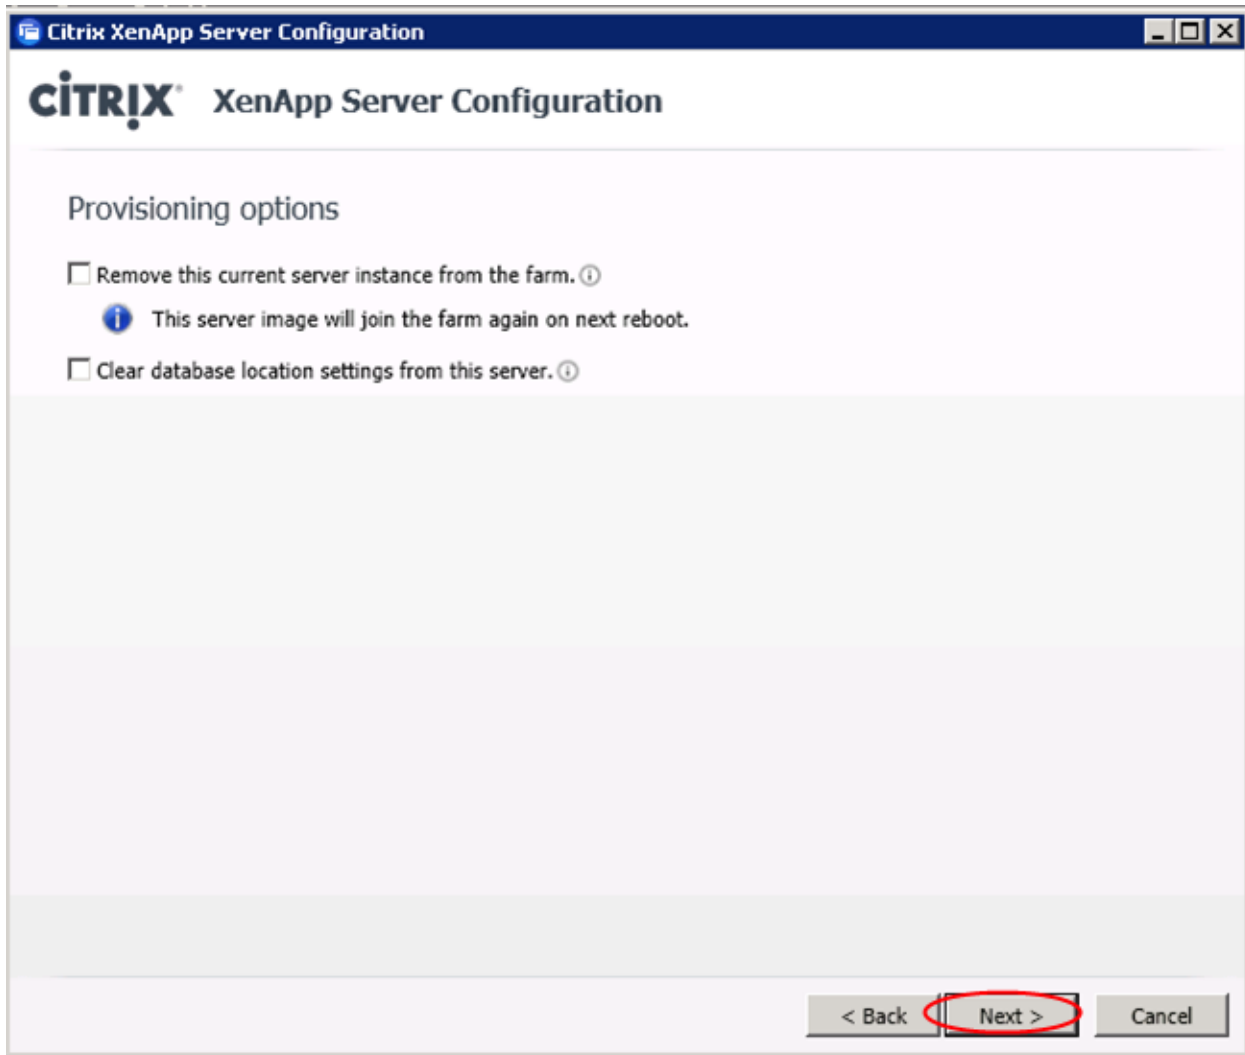
1. Open Citrix XenApp Server Role Manager.
2. Click Edit Configuration.



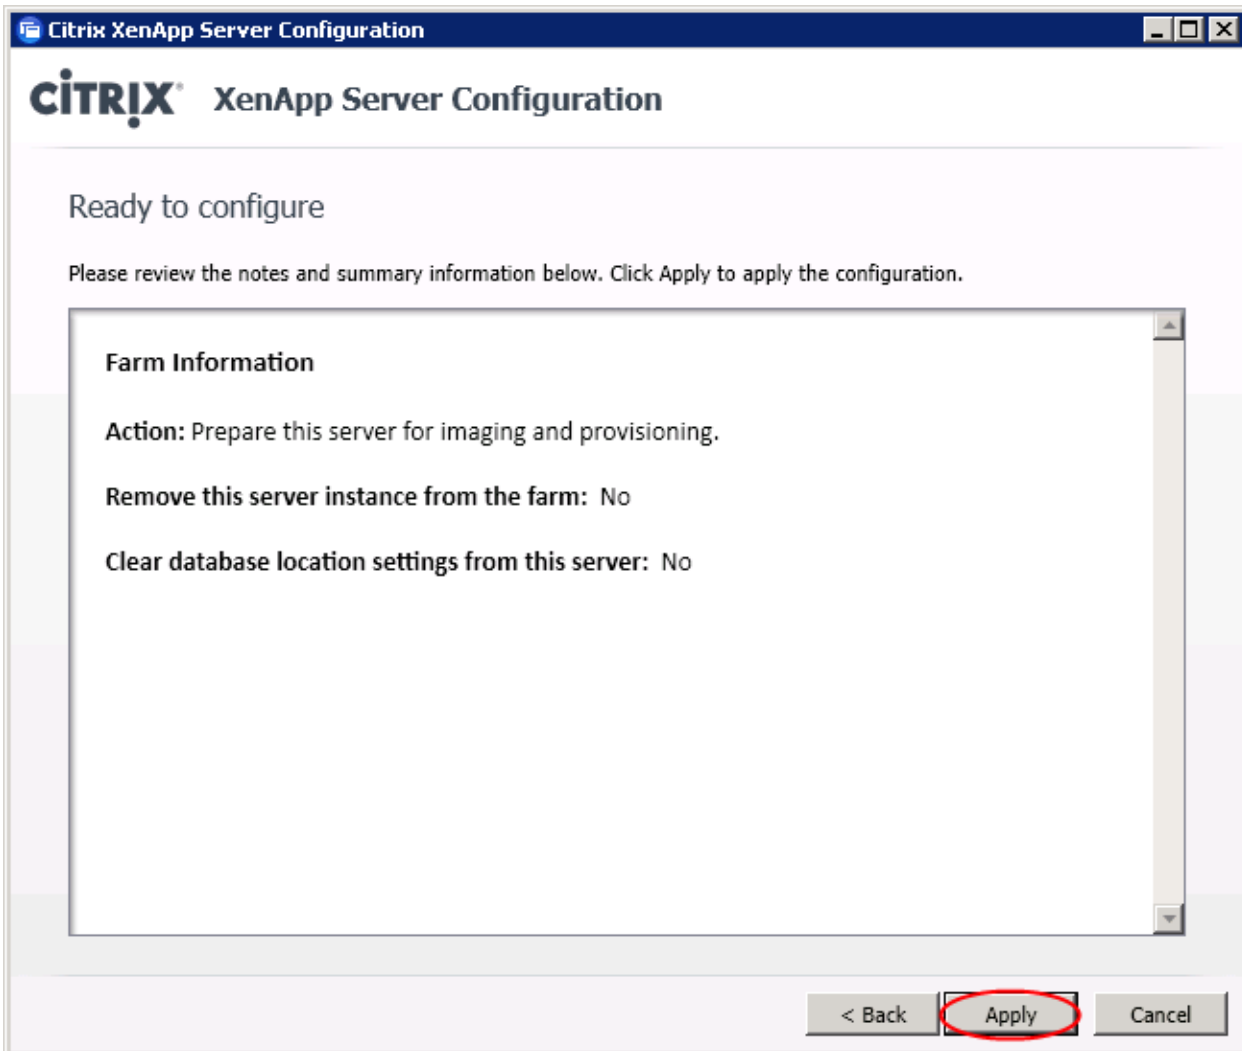
3. Select Prepare this Server for Imaging and Provisioning.



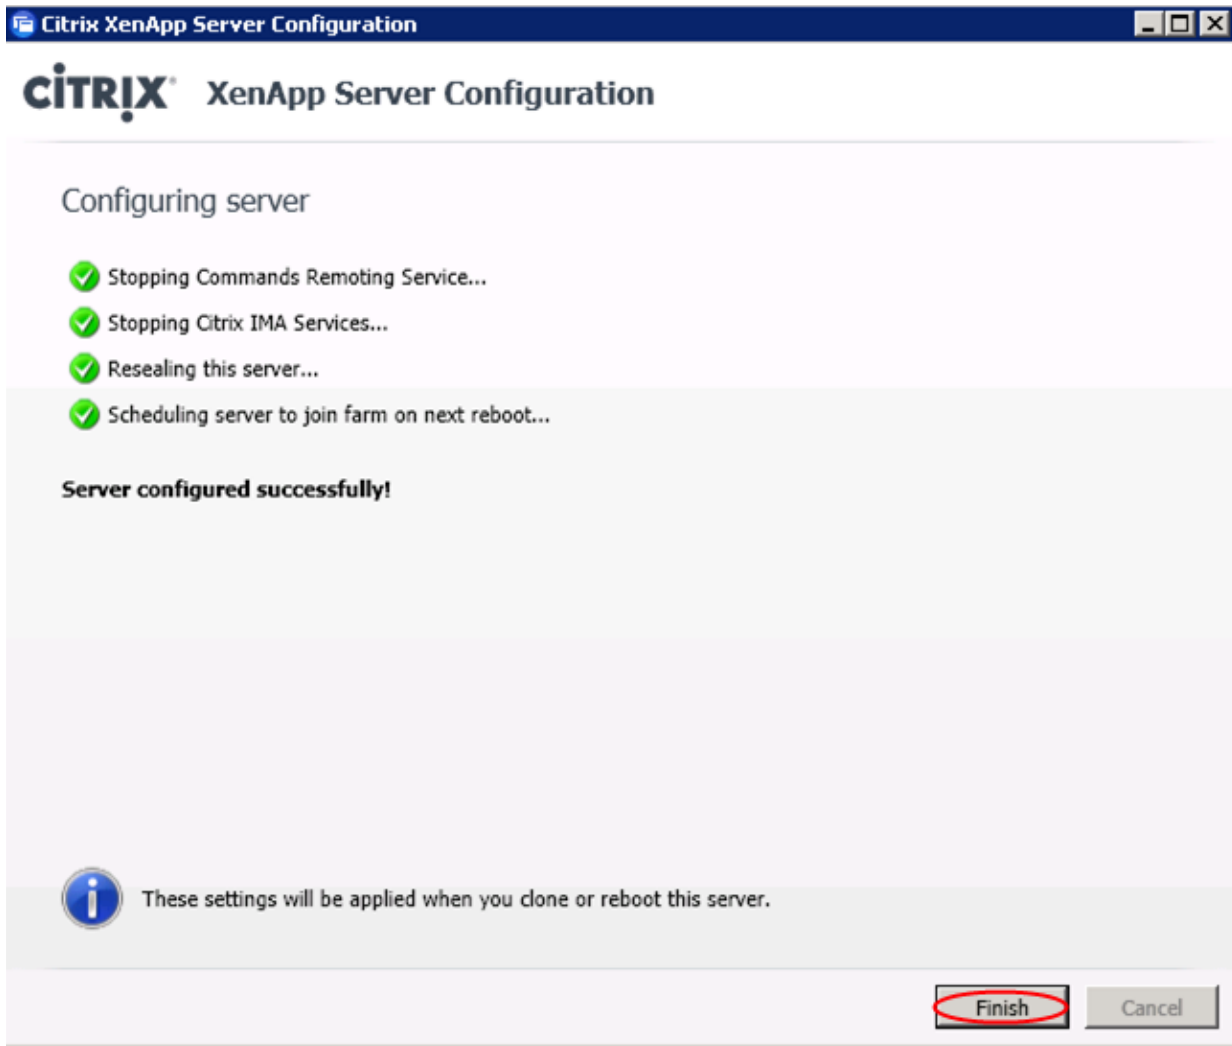
4. Click Next Leaving the checkboxes clear.



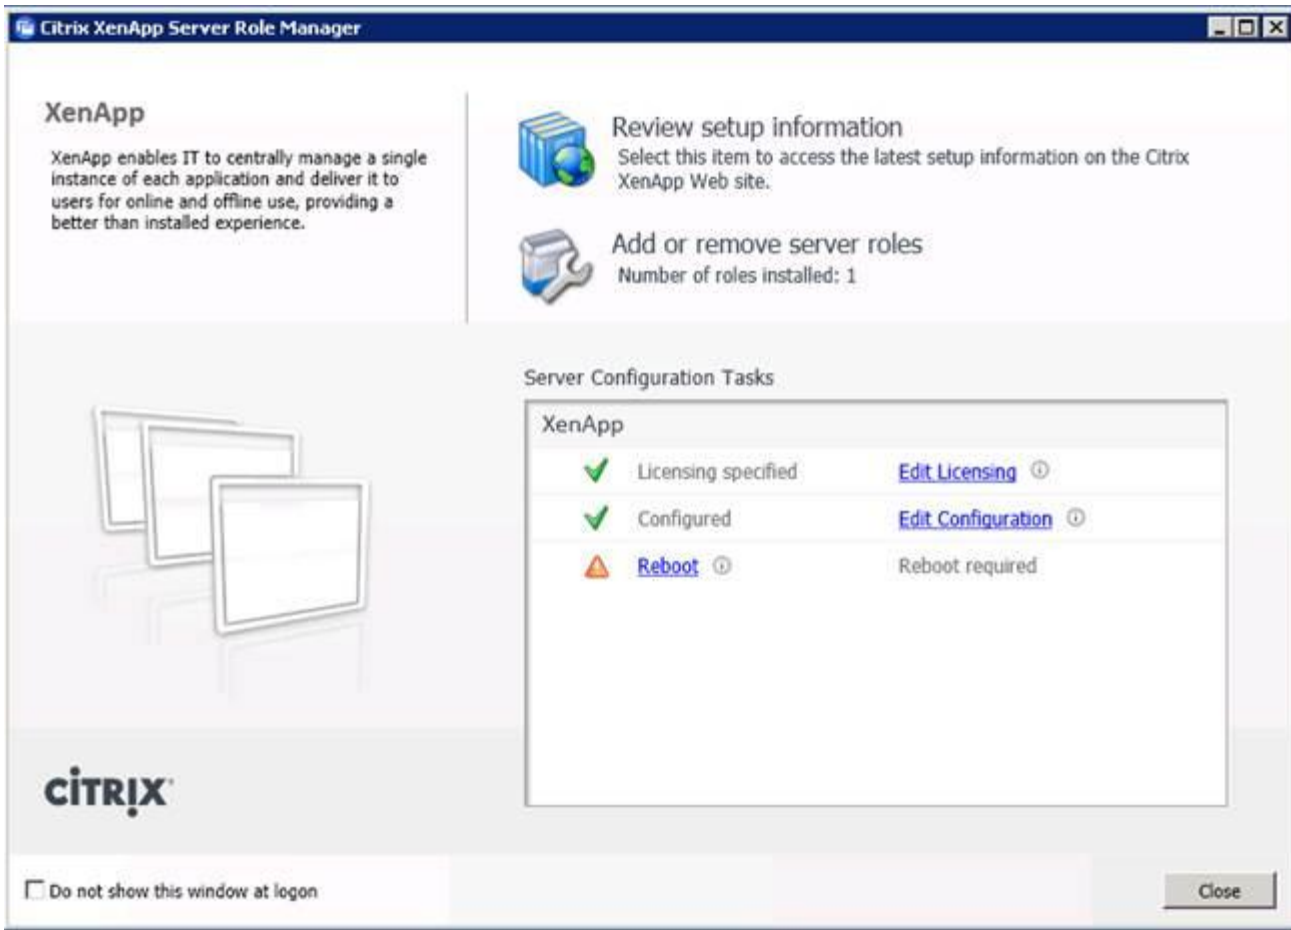
5. Click Apply.



6. Click Finish.

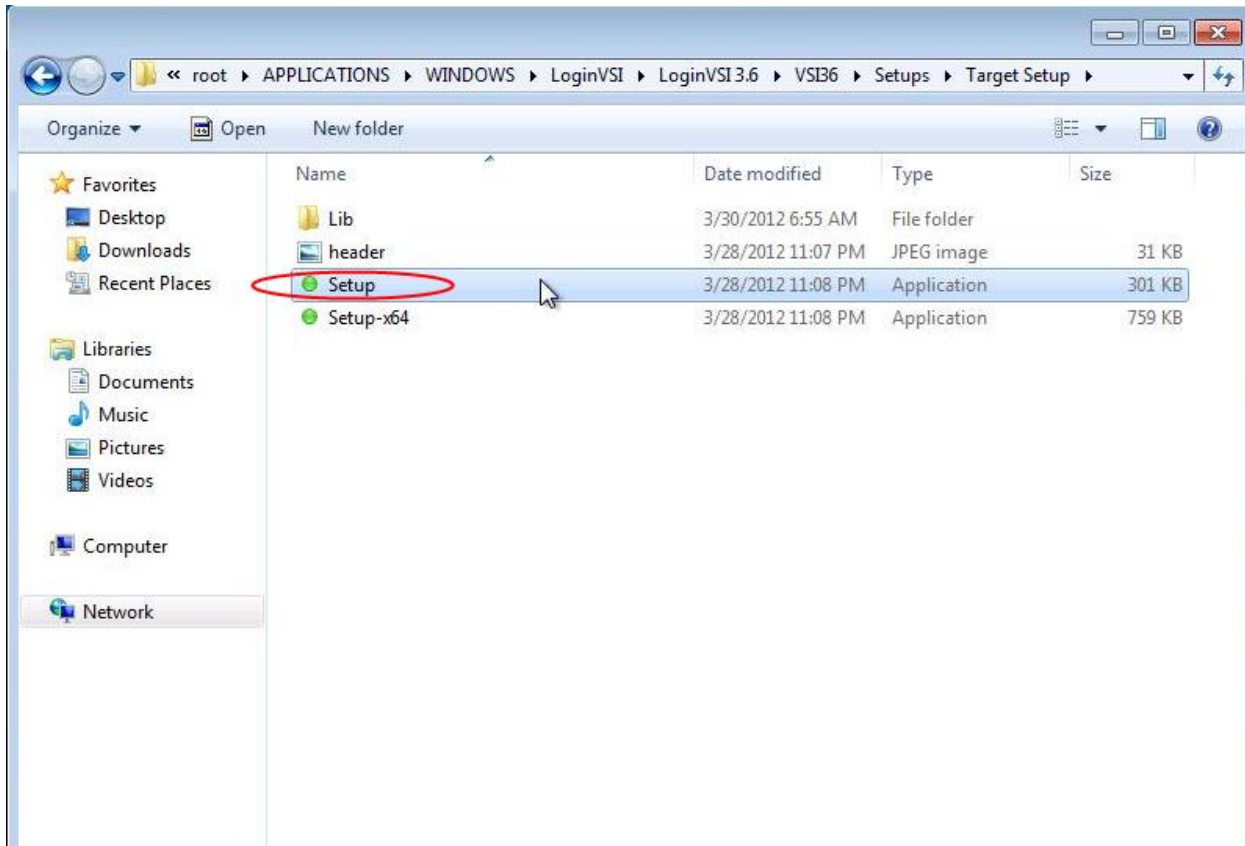


7. Shutdown the server and create PVS image to be used for streaming.



7.4.4 Add Login VSI 3.7 Target Software

1. Launch setup wizard for Login VSI 3.7 using run as Administrator.
2. Launch setup wizard using run as Administrator.



3. Enter VSI share path.



4. Use default installation paths.

7.5.5 Perform Additional PVS 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 8192 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.

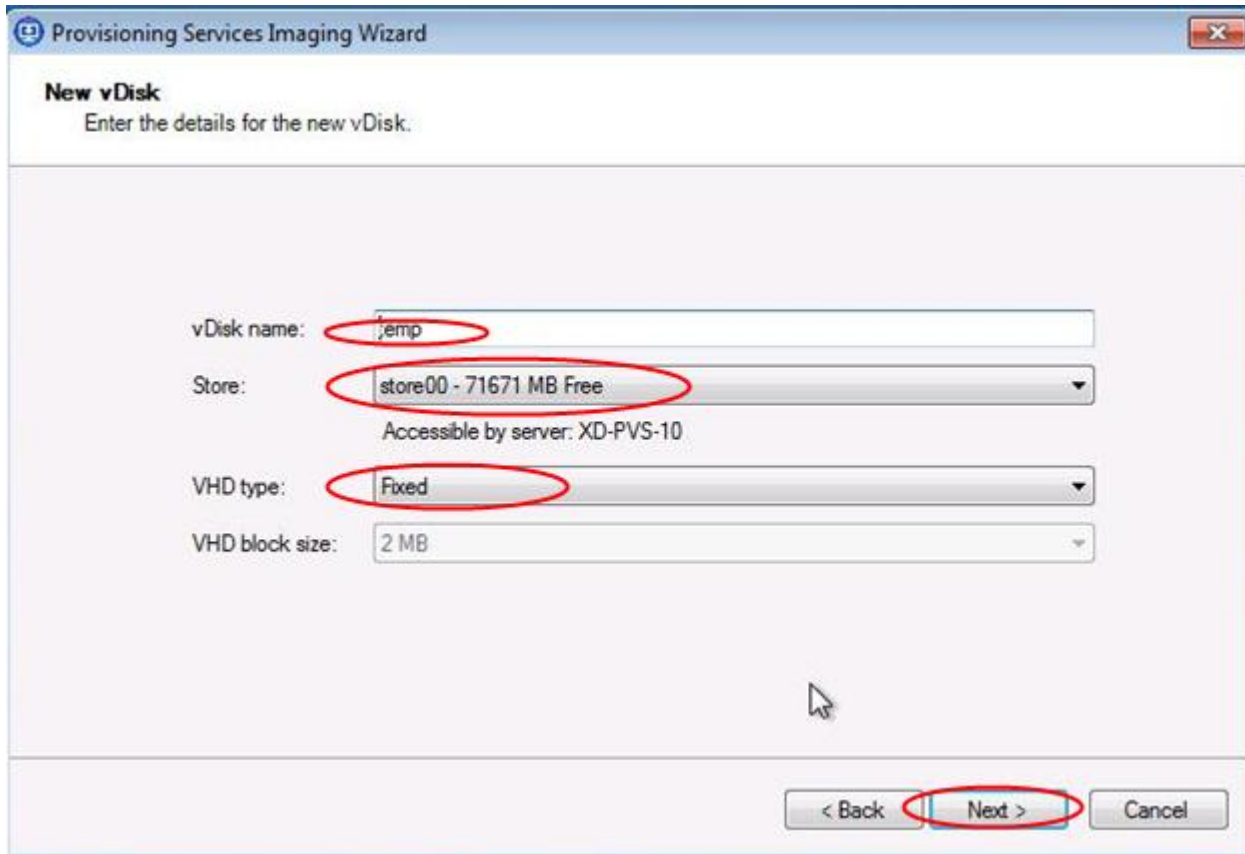
7.4.6 Convert Golden Image Virtual Machine to PVS vDisk

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:

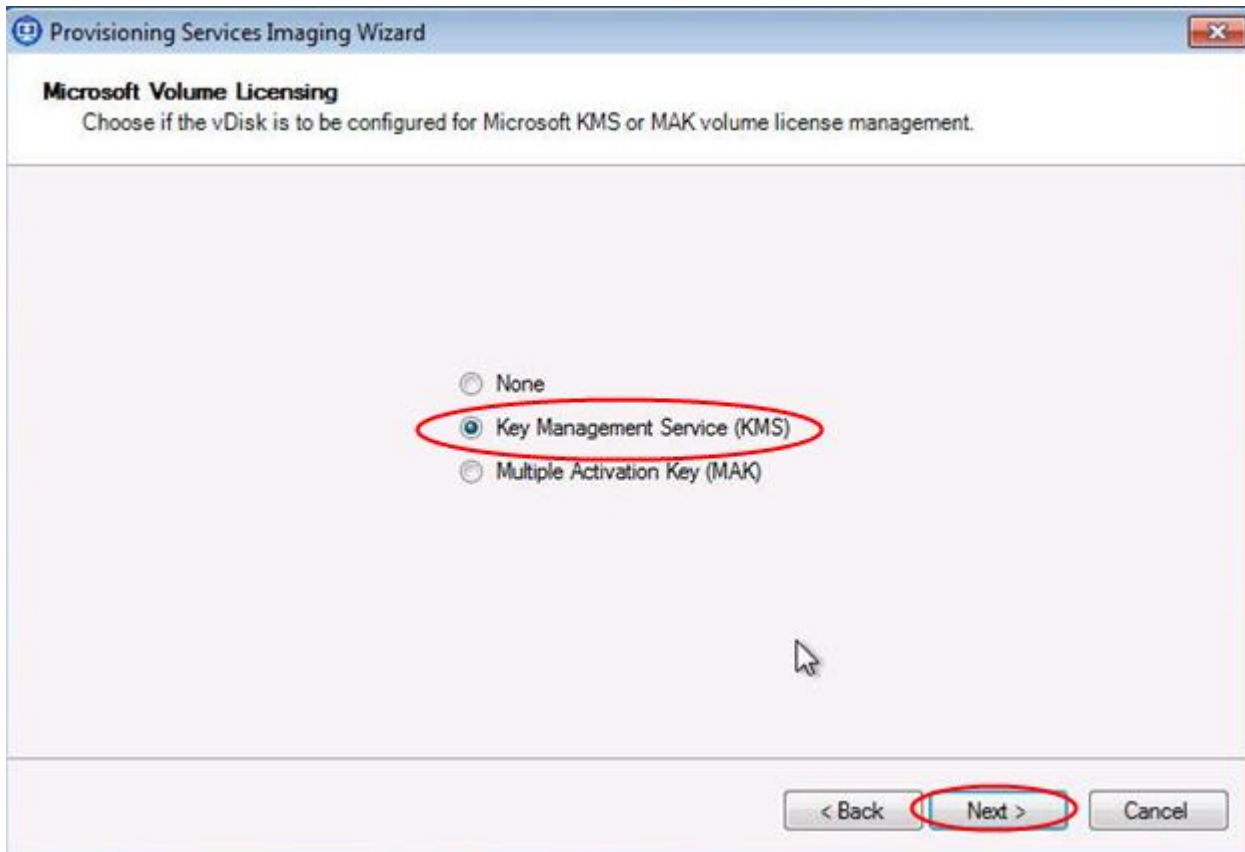
1. Reboot the source virtual machine.
2. Log in to the virtual machine using an account that has administrative privileges.
3. Go to Start | All Programs | Citrix | Provisioning Services.
4. Launch the PVS Imaging Wizard.
5. Click Next at the Welcome Screen.
6. Enter the Server Name or IP address of the PVS server you will be connecting to in order to create the new vDisk.

The screenshot shows the 'Provisioning Services Imaging Wizard' window with the title 'Connect to Farm'. Below the title is the instruction 'Enter the name or address of a server in the farm to connect to.' The form is divided into two sections: 'Server information' and 'Credentials'. In the 'Server information' section, the 'Server' field contains 'xd-pvs-10.mlq3.net' and the 'Port' field contains '54321'. In the 'Credentials' section, the radio button for 'Use my Windows credentials' is selected. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is circled in red.

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. Click Next.
12. Select VHD type Fixed.
13. Click Next.
14. Re-arm windows licensing, following the procedure:
<http://support.citrix.com/proddocs/topic/provisioning-61/pvs-collections-kms-licensing.html>.
15. Select KMS for Licensing Management.



16. Click Next.

17. Use default size image volumes.

Provisioning Services Imaging Wizard

Configure Image Volumes
Define the size of each volume.

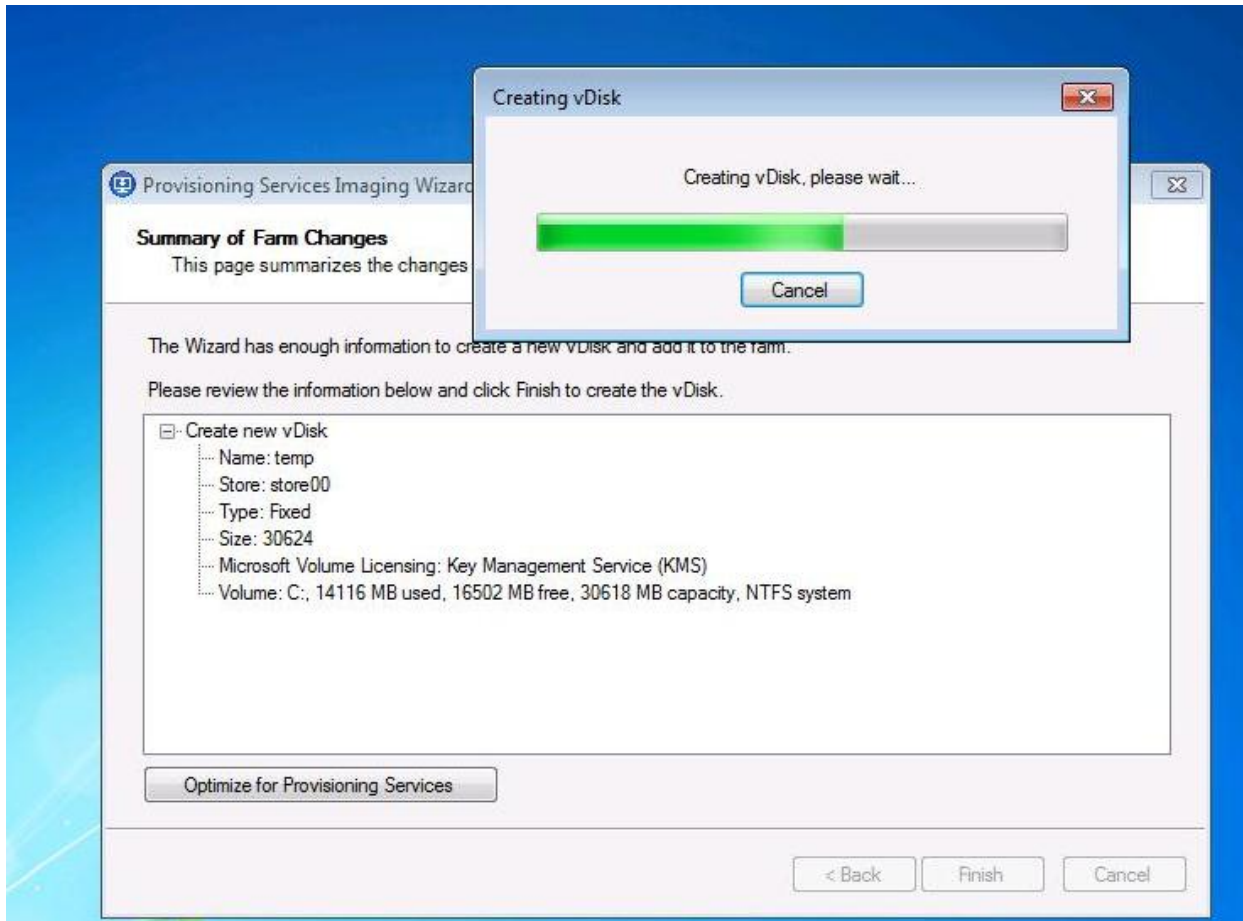
Source Volume	Used Space	Free Space	Capacity	File System
1 C: Boot	14116 MB 46 %	16502 MB 54 %	30618 MB	NTFS
2 None				
3 None				
4 None				

↓

Destination Volume	Used Space	Free Space	Capacity	File System
C: Boot	14116 MB 46 %	16502 MB 54 %	30618 MB	NTFS

vDisk	Allocated Space	Unallocated Space	Capacity
Summary	30618 MB 100 %	6 MB 0 %	30624 MB

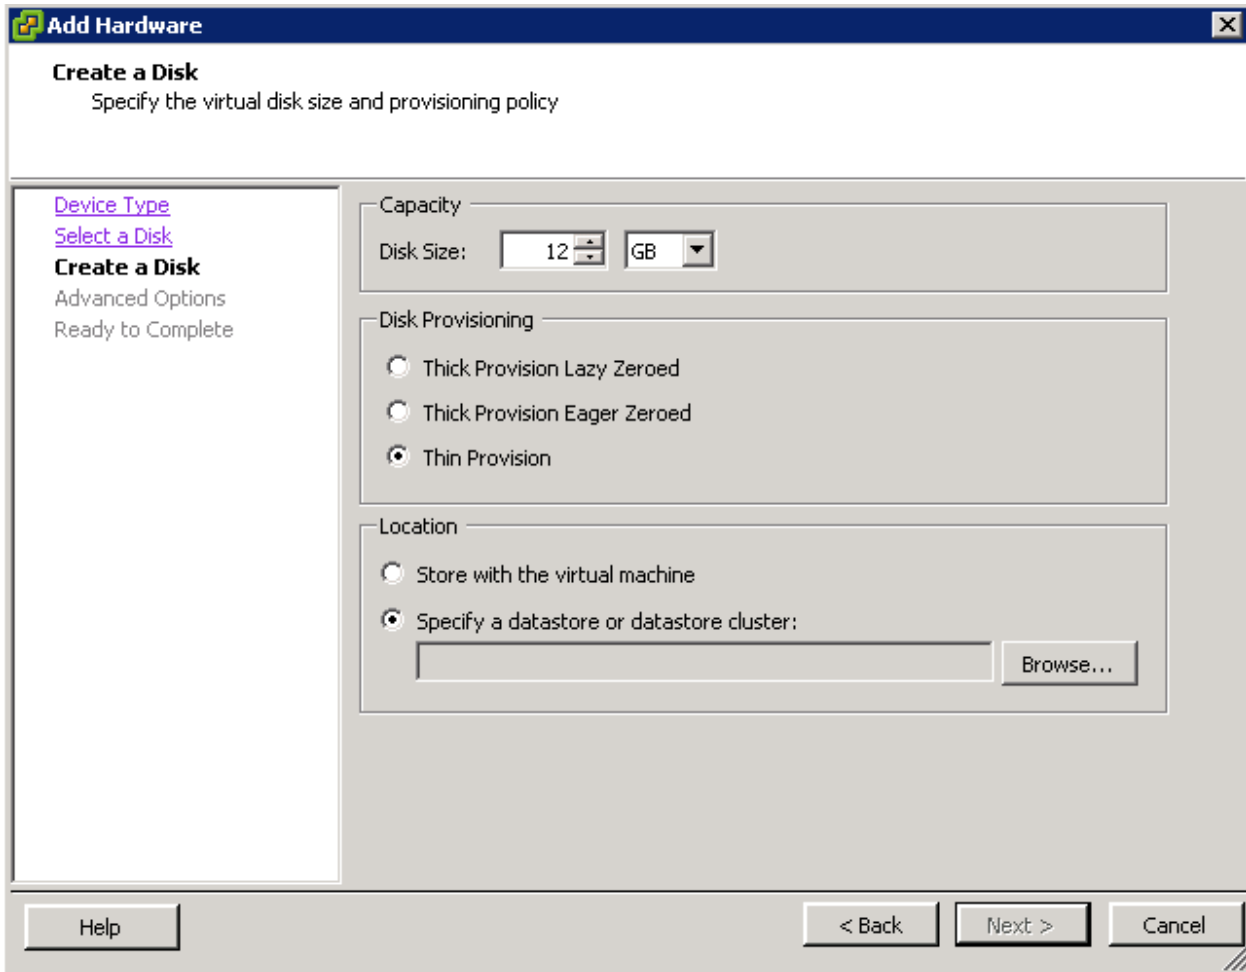
18. Click Next.
19. Click Finish to begin the vDisk creation.



20. 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.
21. Click Yes to reboot the source virtual machine.
22. Logon as the same user account that was used at the beginning of this process.
23. When 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.
24. Shutdown the source virtual machine.
25. Ensure that the VM boots to Network.
26. In PVS, switch the collection account to boot to vDisk.

7.4.7 Add Write-Cache Drives to Virtual Machine Templates in vCenter

1. Add a new 12GB 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

7.5 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.

7.5.1 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 - 6 NFS datastores
- Windows 7 SP1 vDisks storage accessed by the Provisioning Servers (NFS Share) via Windows File Server Share
- XenApp servers (windows 2008 R2 64) vDisk storage accessed by the Provisioning Servers (NFS Share) via Windows File Server Share
- The Launcher vDisks were stored on Fibre Channel LUNs



7.5.2 PVS for Use with Standard Mode Desktops

The Windows 7 SP1 desktop image and XenApp Server image are 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) and Streamed VM Wizard (XenApp).
- 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 Hosted VM 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.
- PVS streamed VM's also feature Personal vDisks. Personal vDisks allow users to retain personal configuration and data. Pvd Hosted VM Desktops are configured with a 5GB Pvd disk.

Ten PVS servers were configured in a farm with a two sites to provide streaming services for 2000 Hosted VM Desktop Virtual Machines and 88 Virtual XenApp 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.

Seven PVS servers are assigned to a site designated to 2000 Hosted VM Desktop Machines. (nonPvd and with Pvd).

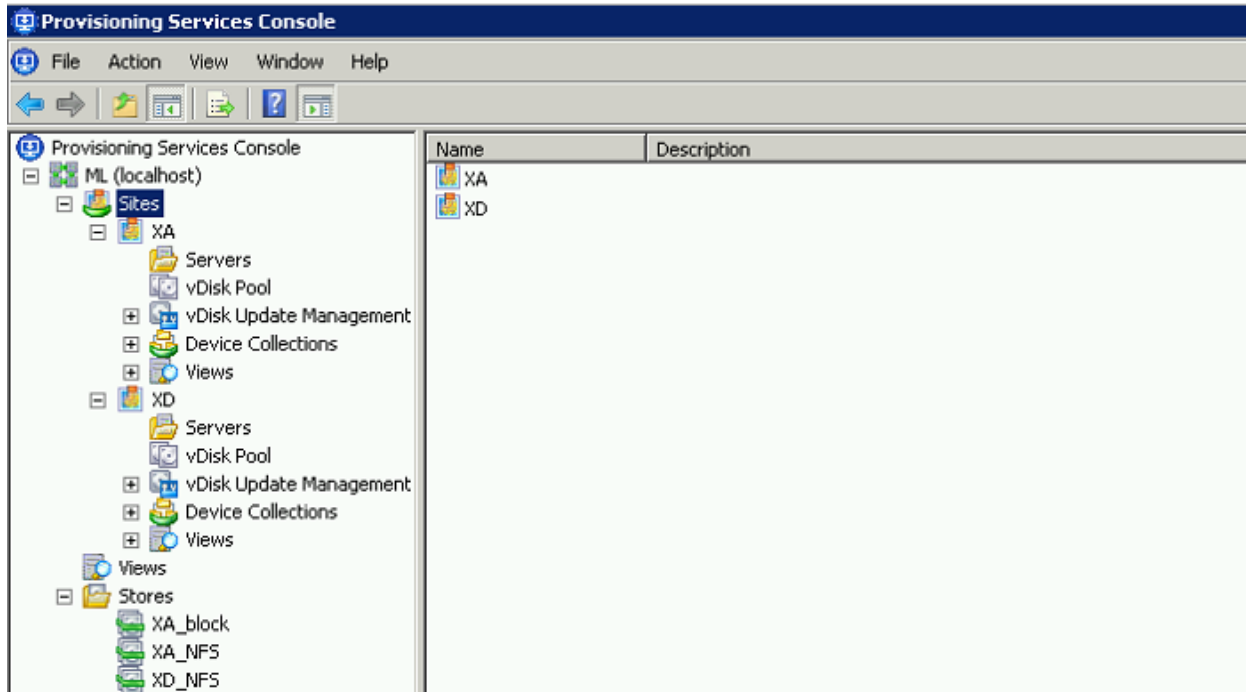
Three PVS servers are assigned to a separate site for streaming 88 XenApp Virtual Machines capable of 2000 Hosted Shared Desktop sessions.

The vDisks are hosted on a clustered file server share and were 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 19. 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.

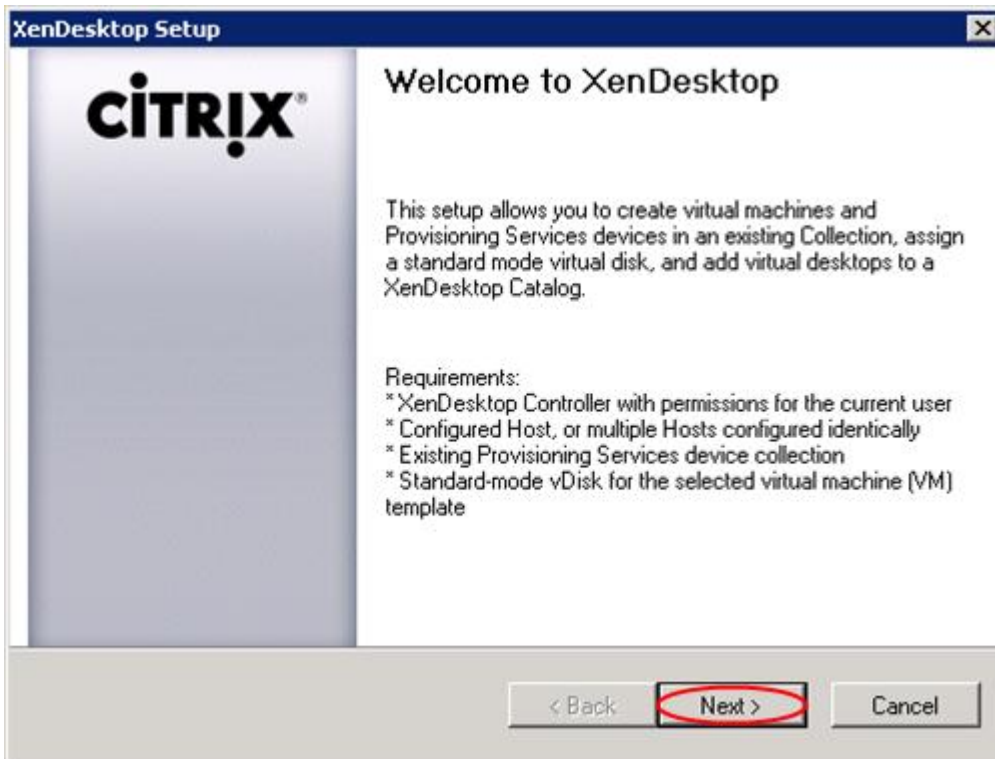
The PVS Target device agent installed in the Windows 7 gold image and XenApp Server gold image automatically places the Windows swap file on the same drive used by the PVS Write Cache when this mode is enabled.

7.5.3 Process to Create Hosted VM Desktops with Tier 0 Storage Using XenDesktop Setup Wizard in PVS

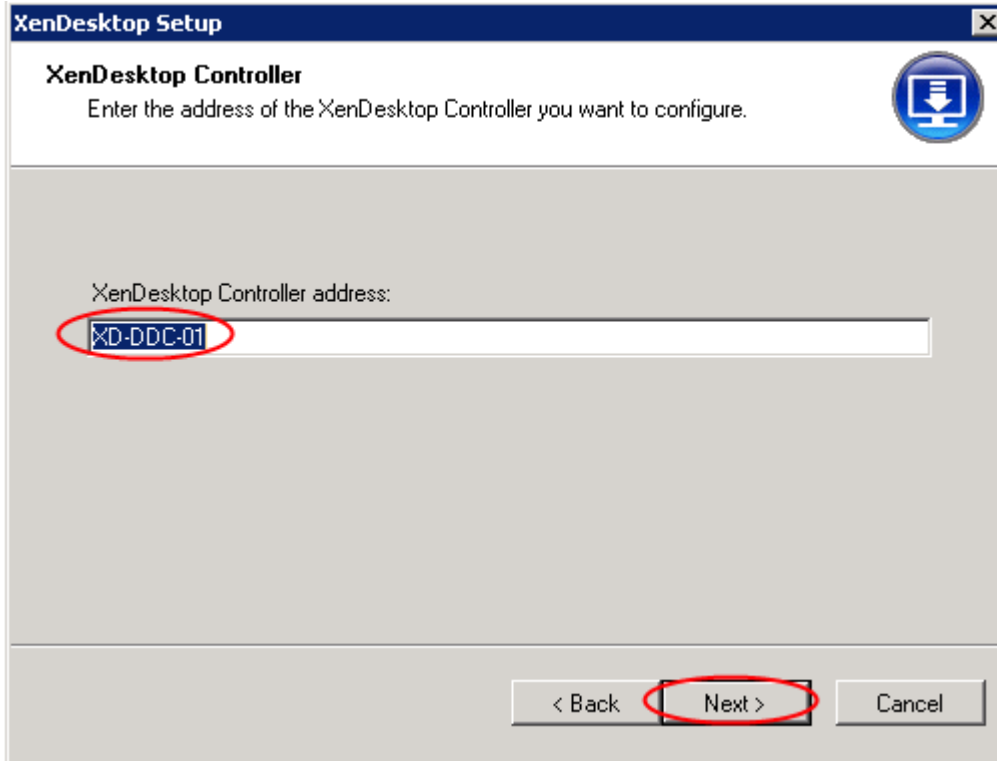
1. Start XenDesktop Setup Wizard.



2. Click Next.

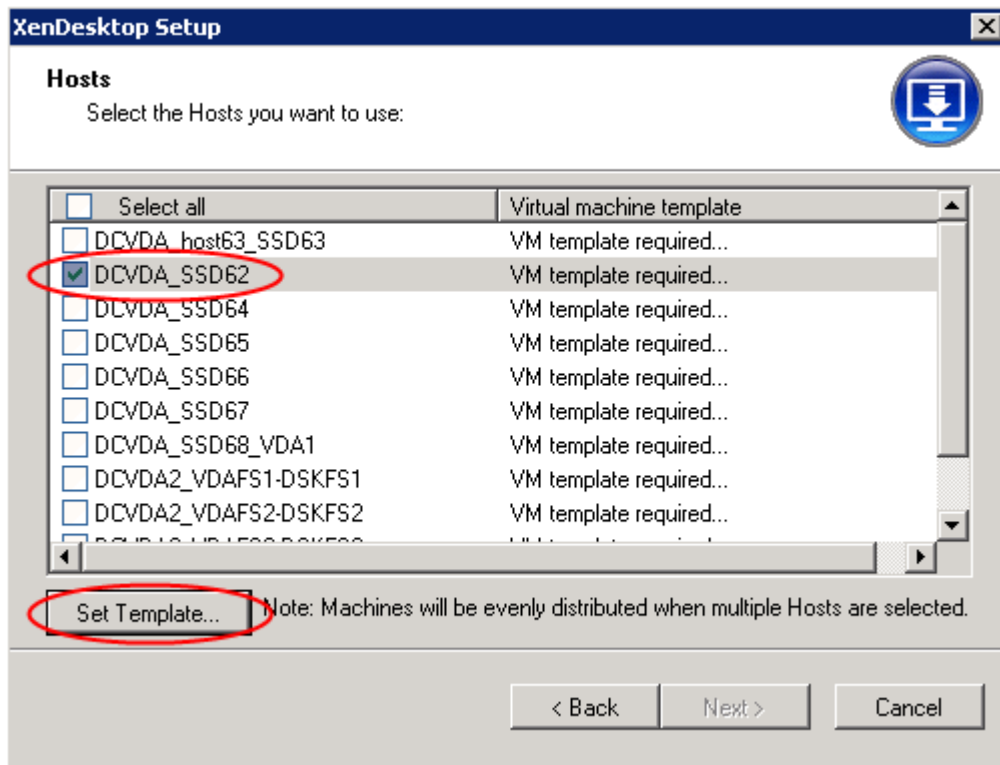


3. Connect to XenDesktop Controller and click Next.



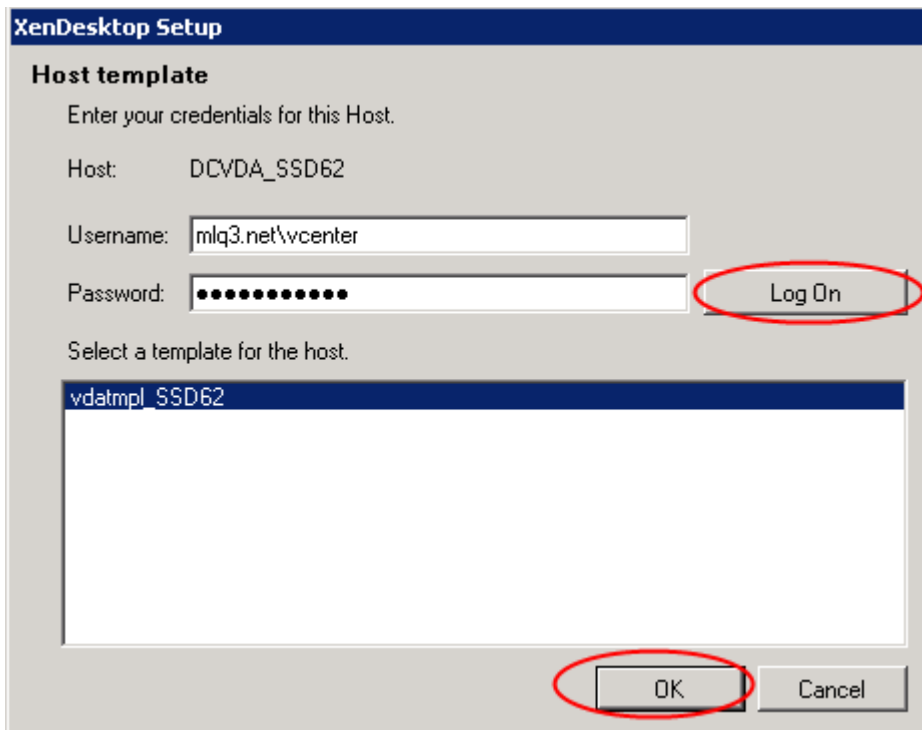
4. Configure VM template.

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.

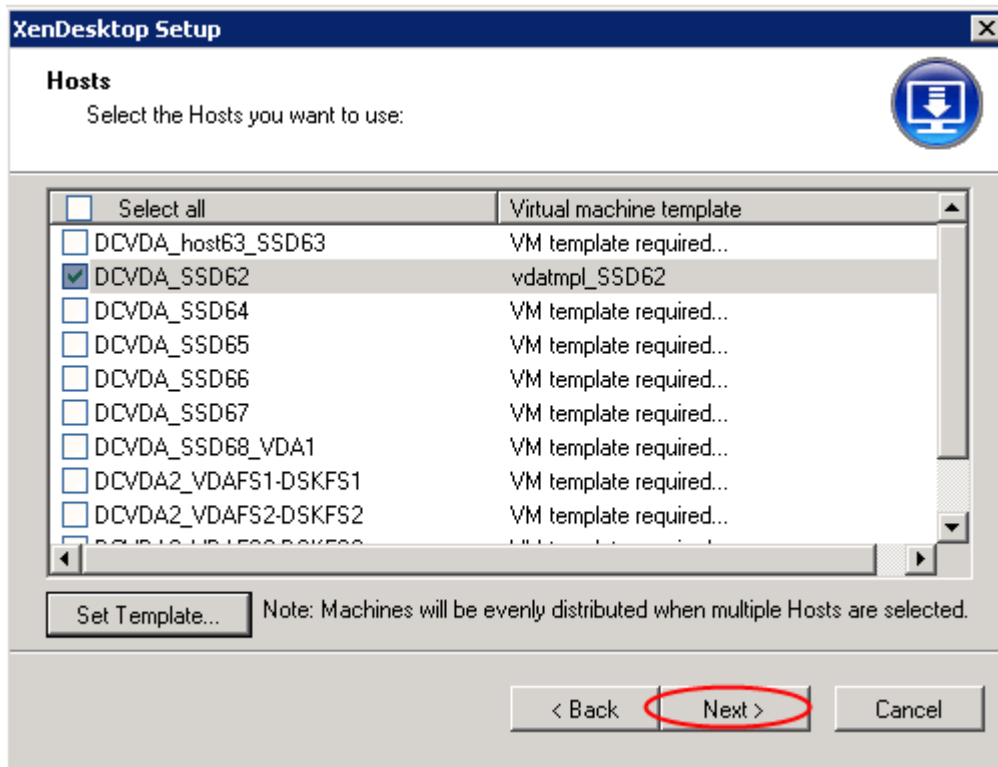


5. Enter the password to connect to vCenter and Select VM template you are going to use.

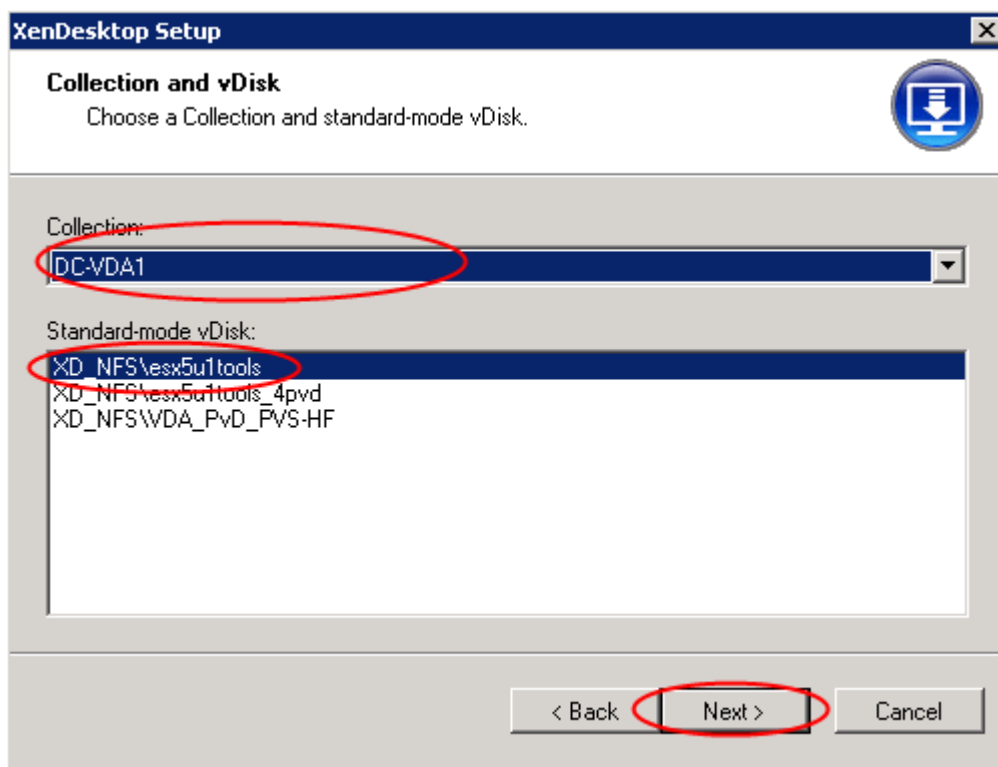
Note: Prior to this step, the VM templates need to be configured on each VMware datastore that will contain drives for the streamed desktops.



6. Click Next.



7. Select PVS device Collection for non-PVD Collection and vDisk.



8. Select XenDesktop Catalog preferences.

XenDesktop Setup
Catalog
Select your Catalog preferences.

Create a new catalog
 Use an existing catalog

Machine type: Streamed

Catalog name: DC-VDA1

Description: nonPvD

Select Administrators:

<input checked="" type="checkbox"/>	Administrator name	Description
<input checked="" type="checkbox"/>	MLQ3\Domain Admins	Full
<input checked="" type="checkbox"/>	MLQ3\vadiml	Full

< Back Next > Cancel

9. Select Virtual Machine preferences .

XenDesktop Setup
Virtual machines
Select your virtual machine preferences.

Number of virtual machines to create: 155

vCPUs: 1

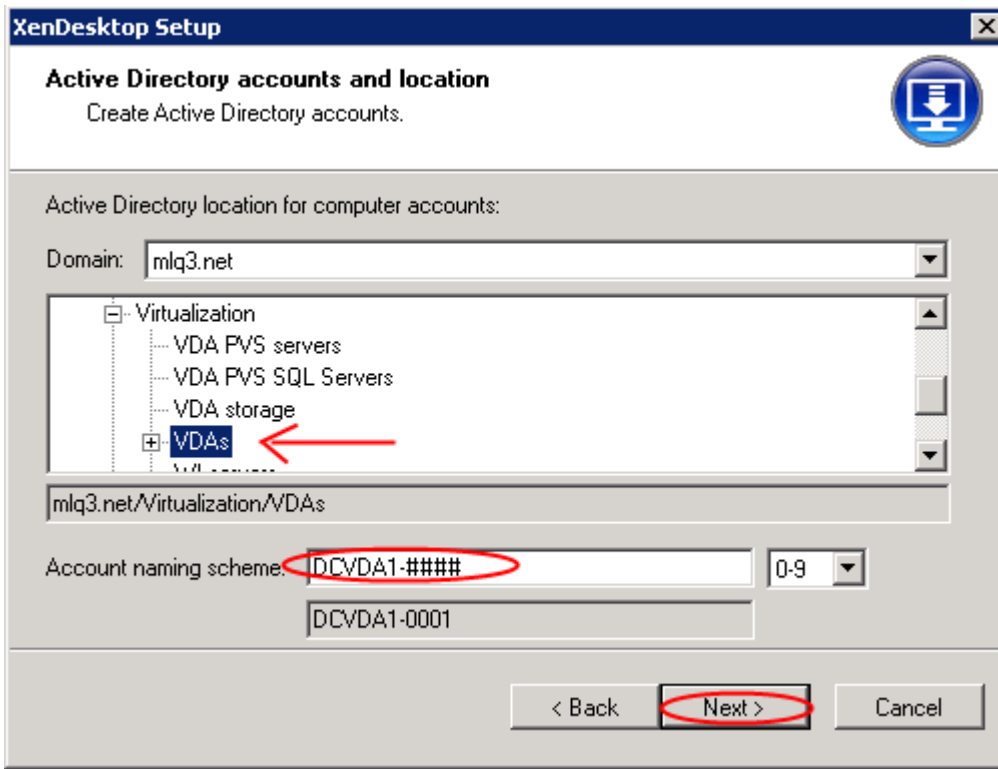
Memory: 1536 MB

Local write cache disk: 3072 MB 3072 MB

Active Directory computer accounts:
 Create new accounts
 Import existing accounts

< Back Next > Cancel

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.
- Create XenDesktop machines in the assigned existing catalog or a new XenDesktop catalog.

7.5.4 Process to Create Virtual Desktops with Personal vDisk in PVS

1. Start XenDesktop Setup Wizard.

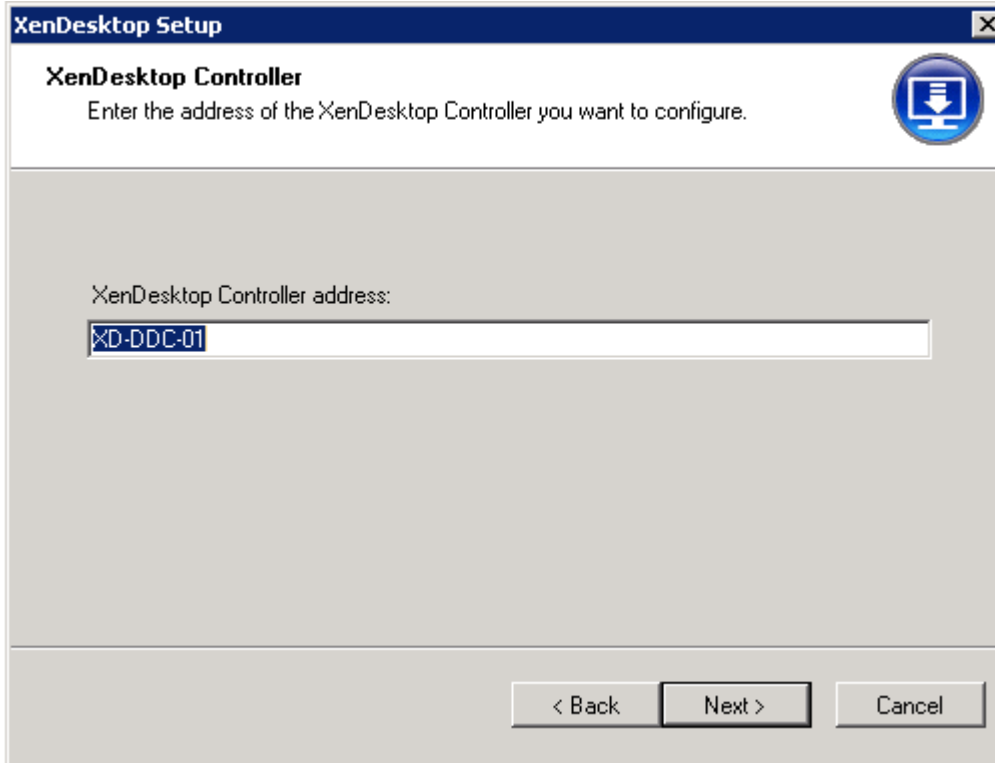


2. Click Next.



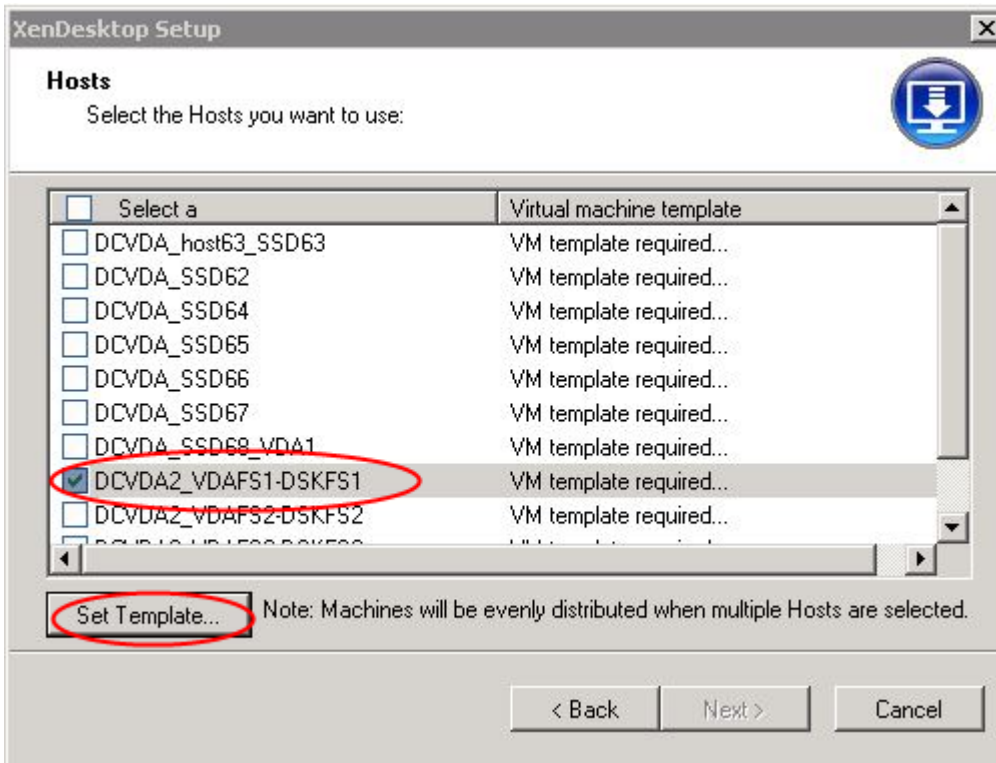


3. Connect to XenDesktop Controller.



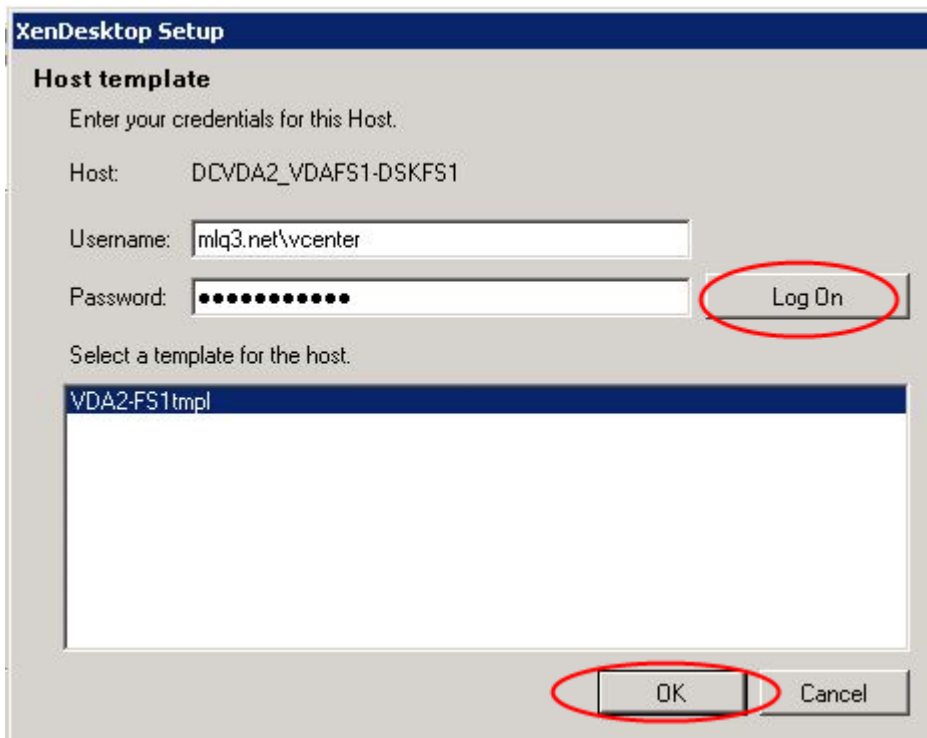
4. Configure VM template.

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.

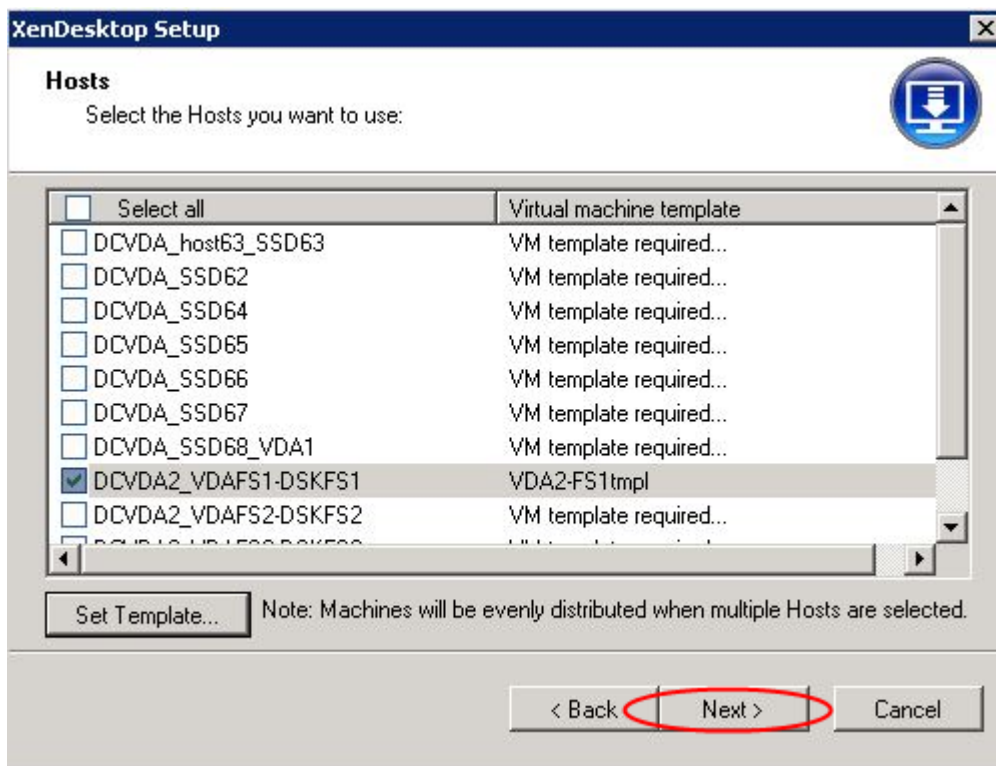


5. Input Password to connect to vCenter and Select VM template you are going to use.

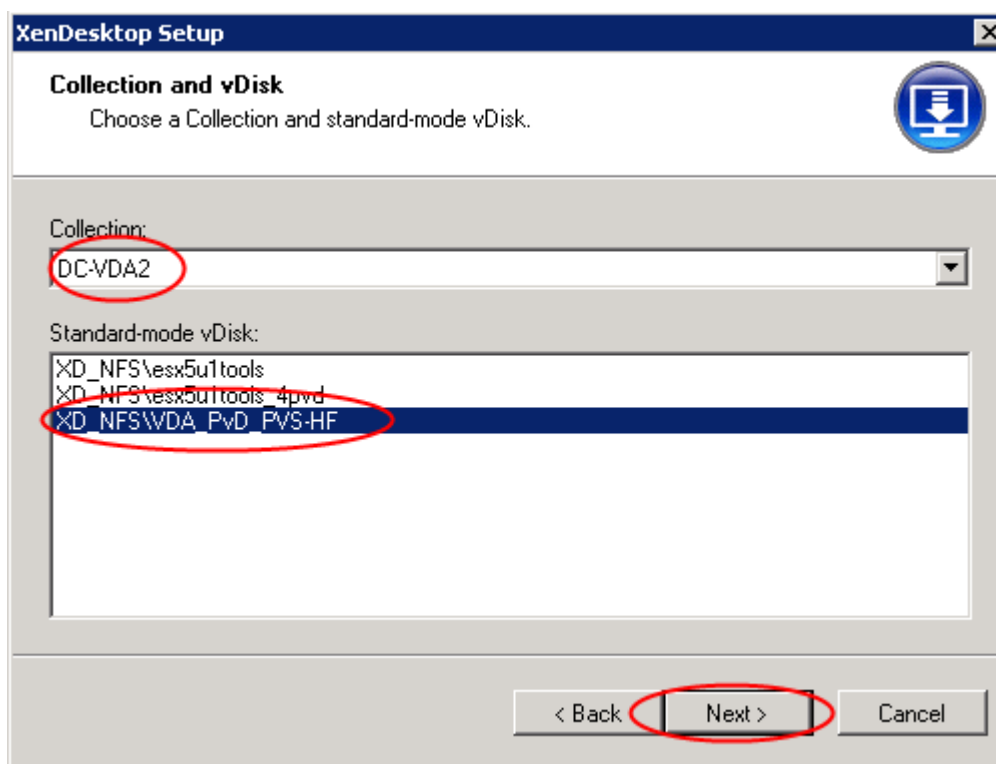
Note: Prior to this step, the VM templates need to be configured on each VMware datastore that will contain drives for the streamed desktops.



6. Click Next.



7. Select PVS device Collection.



8. Select XenDesktop Catalog preferences.

XenDesktop Setup

Catalog
Select your Catalog preferences.

Create a new catalog
 Use an existing catalog

Machine type: Streamed with personal vDisk

Catalog name: DC_VDA2_PVD
Description: pvd

Select Administrators:

<input checked="" type="checkbox"/>	Administrator name	Description
<input checked="" type="checkbox"/>	MLQ3\Domain Admins	Full
<input checked="" type="checkbox"/>	MLQ3\vadiml	Full

< Back **Next >** Cancel

9. Select Virtual Machine preferences Memory and Pvd size.

XenDesktop Setup

Virtual machines
Select your virtual machine preferences.

Number of virtual machines to create: 1015

vCPUs: 1

Memory: 1536 MB

Local write cache disk: 3072 MB 3072 MB

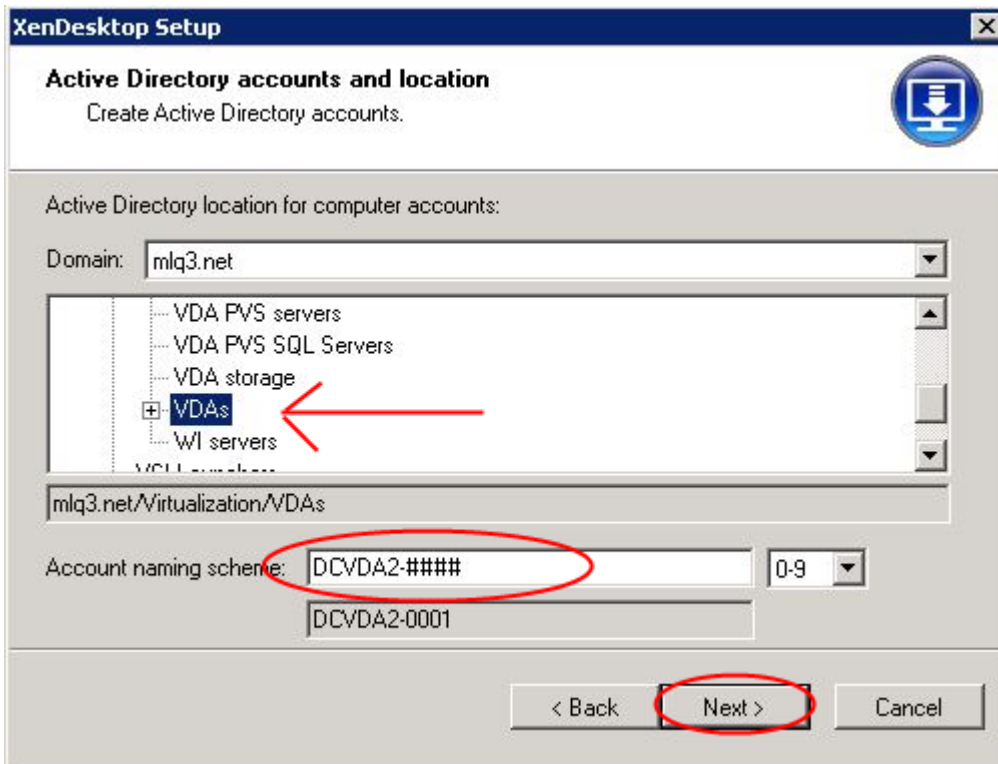
Personal vDisk size: 10 GB 5 GB

Personal vDisk drive letter: P: P:

Active Directory computer accounts:
 Create new accounts
 Import existing accounts

< Back **Next >** Cancel

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.
- Create XenDesktop machines in the assigned existing catalog or a new XenDesktop catalog.

7.5.5 Process to Create XenApp 6.5 Virtual Machines Using the Streamed VM Setup Wizard in PVS

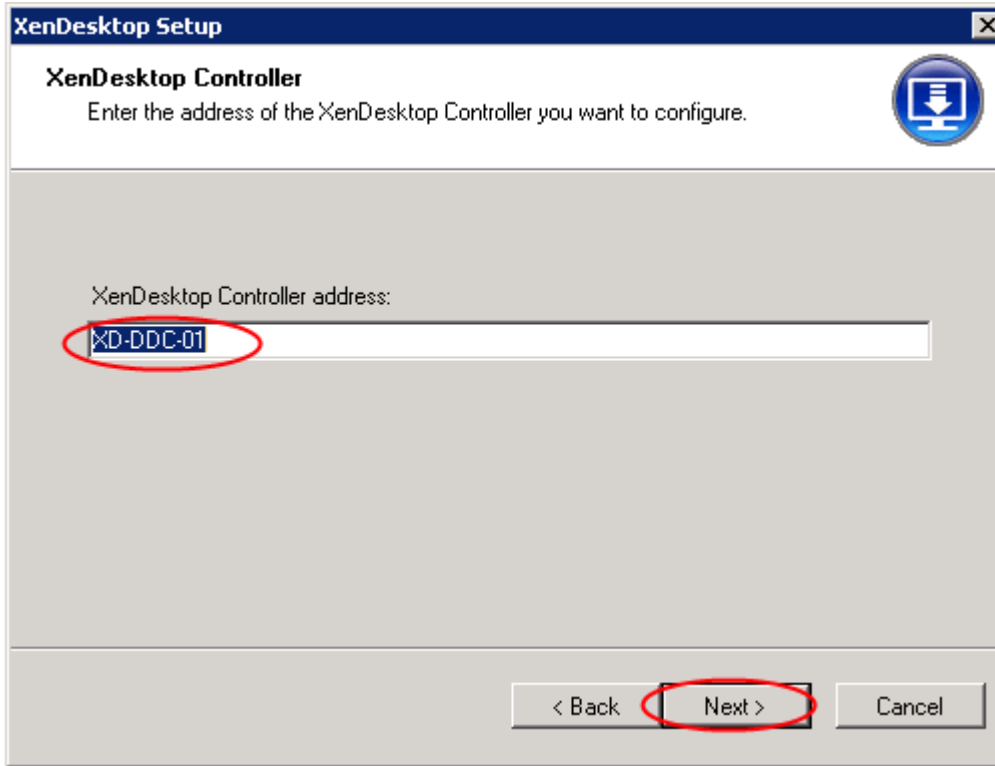
1. Start the XenDesktop Setup Wizard.



2. Click Next.

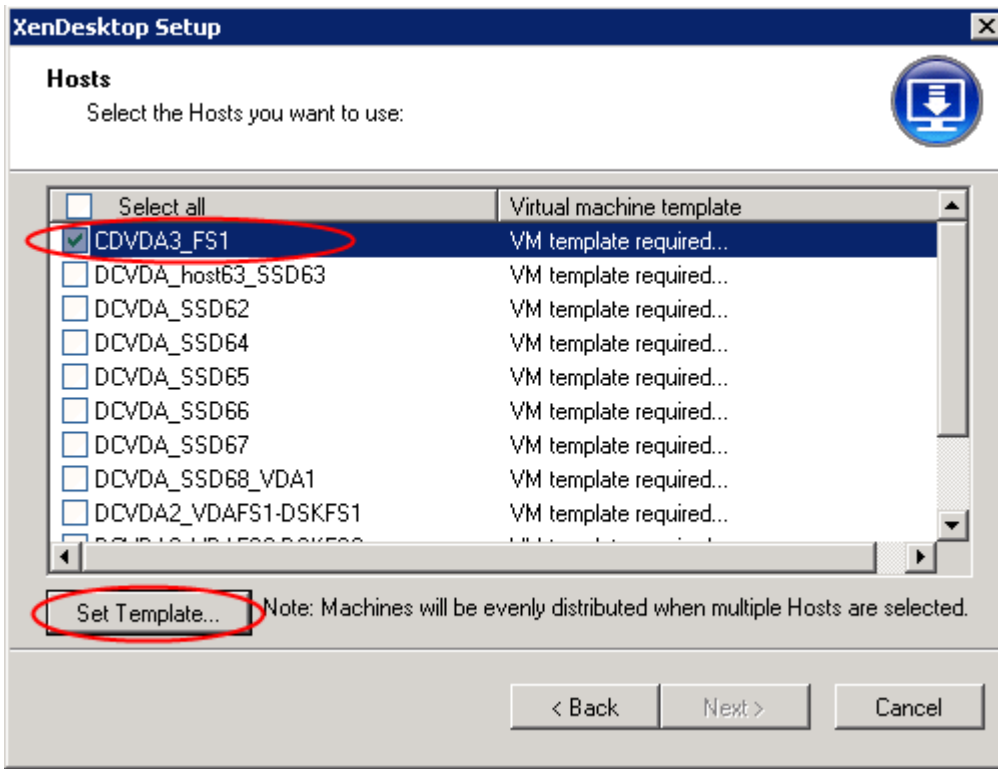


3. Connect to XenDesktop Controller.



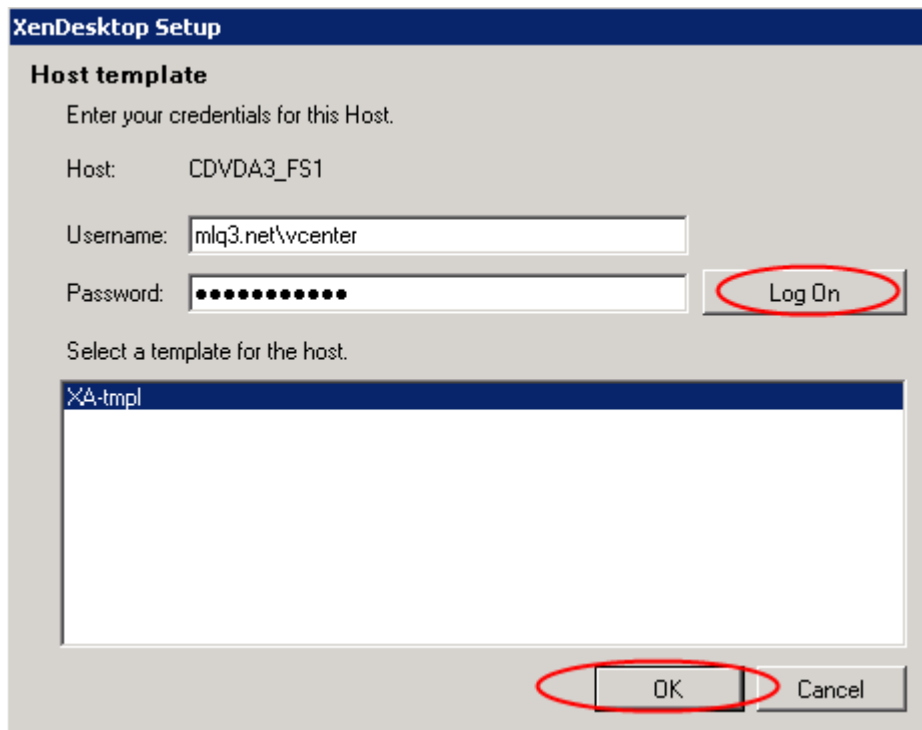
4. Configure VM template.

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.

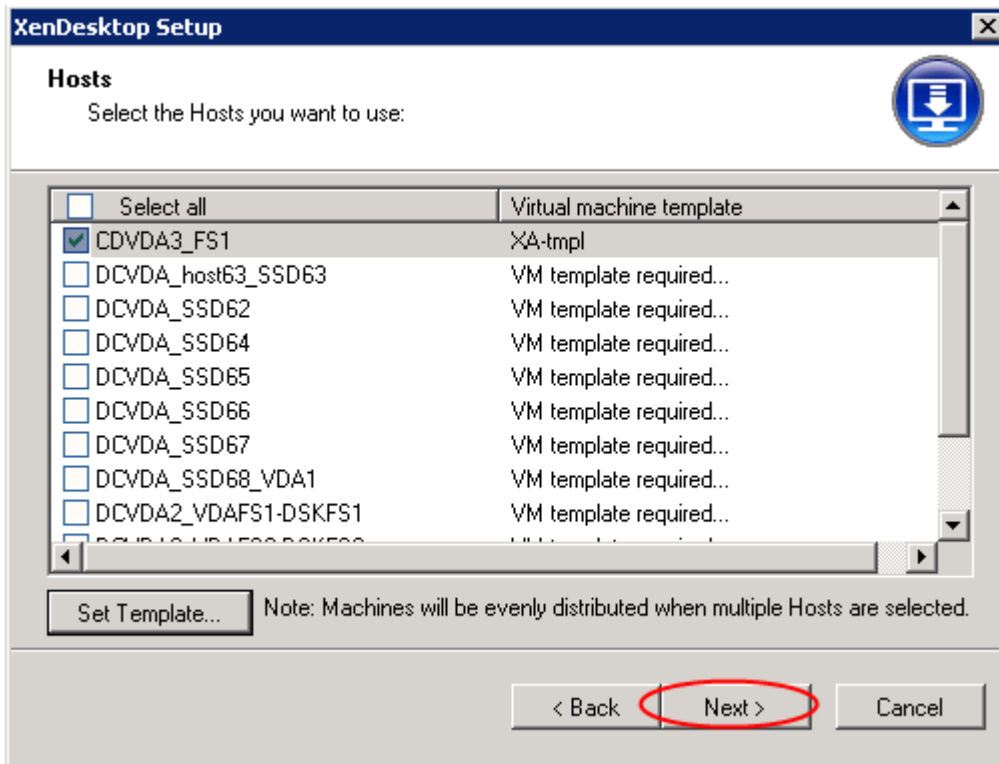


5. Input Password to connect to vCenter and Select VM template you are going to use.

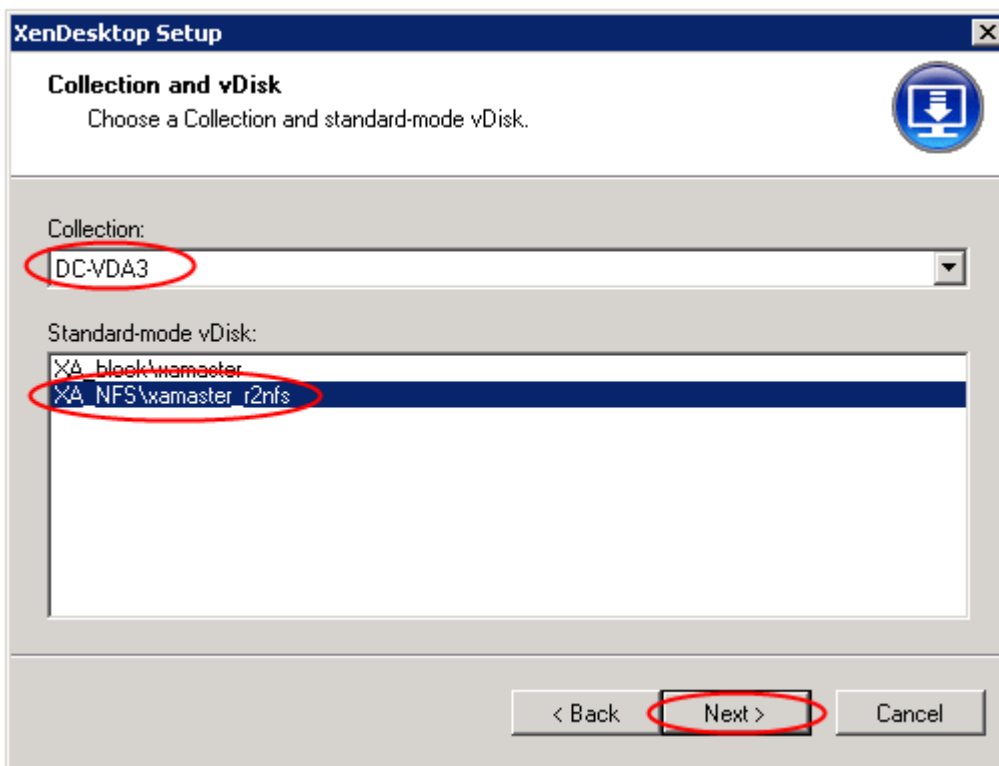
Note: Prior to this step, the VM templates need to be configured on each VMware datastore that will contain drives for the streamed desktops.



6. Click Next.



7. Select PVS device Collection for non-PVD Collection designated for XA and vDisk.



8. Select XenDesktop Catalog preferences.

XenDesktop Setup
Catalog
Select your Catalog preferences.

Create a new catalog
 Use an existing catalog

Machine type: Streamed

Catalog name: DC-VDA3

Description: XA

Select Administrators:

<input checked="" type="checkbox"/>	Administrator name	Description
<input checked="" type="checkbox"/>	MLQ3\Domain Admins	Full
<input checked="" type="checkbox"/>	MLQ3\vadiml	Full

< Back Next > Cancel

9. Select Virtual Machine preferences.

XenDesktop Setup
Virtual machines
Select your virtual machine preferences.

Number of virtual machines to create: 4

vCPU: 4

Memory: 12288 MB

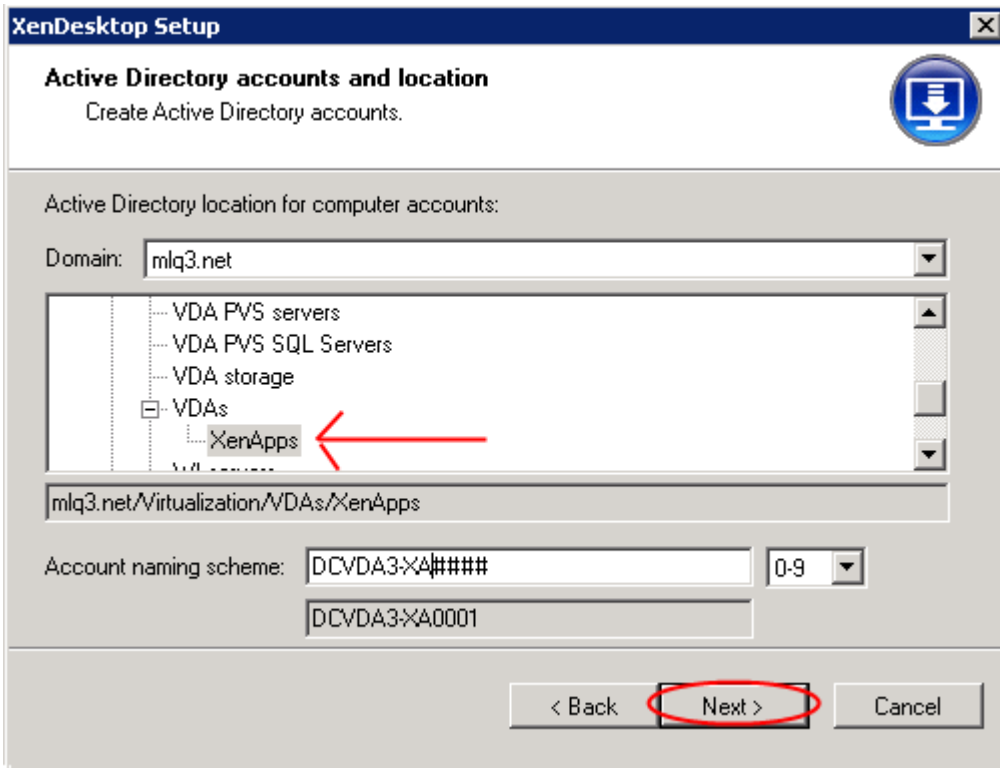
Local write cache disk: 12288 MB 12288 MB

Active Directory computer accounts:

Create new accounts
 Import existing accounts

< Back Next > Cancel

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.
- 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 B200 M3 blade in a single chassis and twenty five Cisco UCS B200 M3 blades in 4 chassis to illustrate linear scalability.

8.1 Cisco UCS Test Configuration for Single Blade Scalability

Figure 20. Cisco UCS B200 M3 Blade Server for Single Server Scalability XenDesktop 5.6 with PVS Write Cache on Tier 0 SSD.

Cisco UCS B200 M3 Blade Server Single Blade XD Test Result— 155 Users

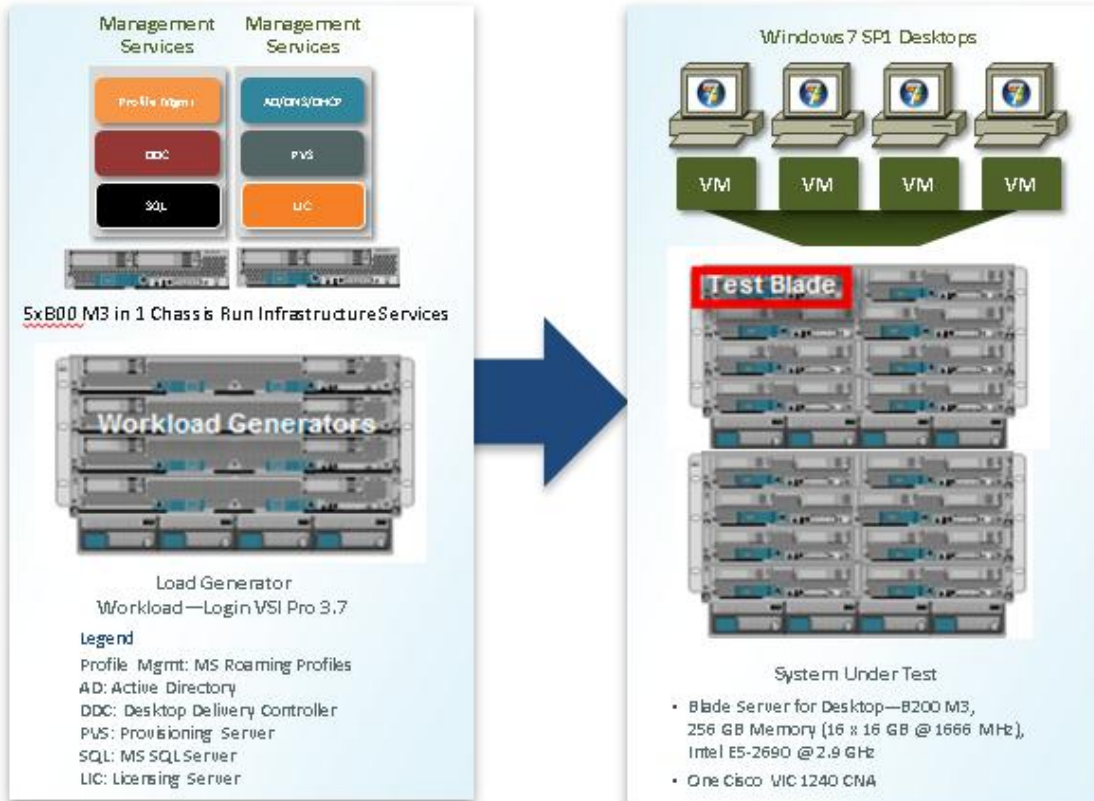


Figure 21. Cisco UCS B200 M3 Blade Server for Single Server Scalability XenDesktop 5.6 with Personal vDisk

Cisco UCS B200 M3 Blade Server

Single Blade XD PvD Test Result— 145 Users

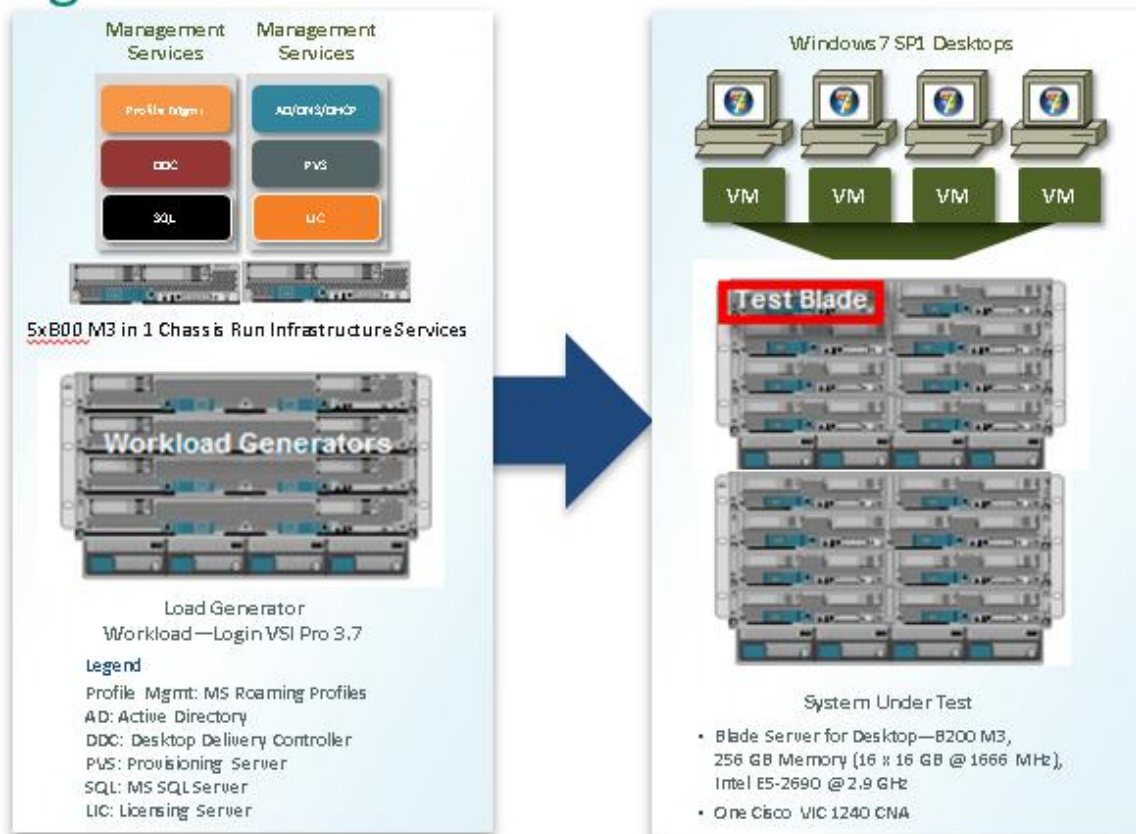
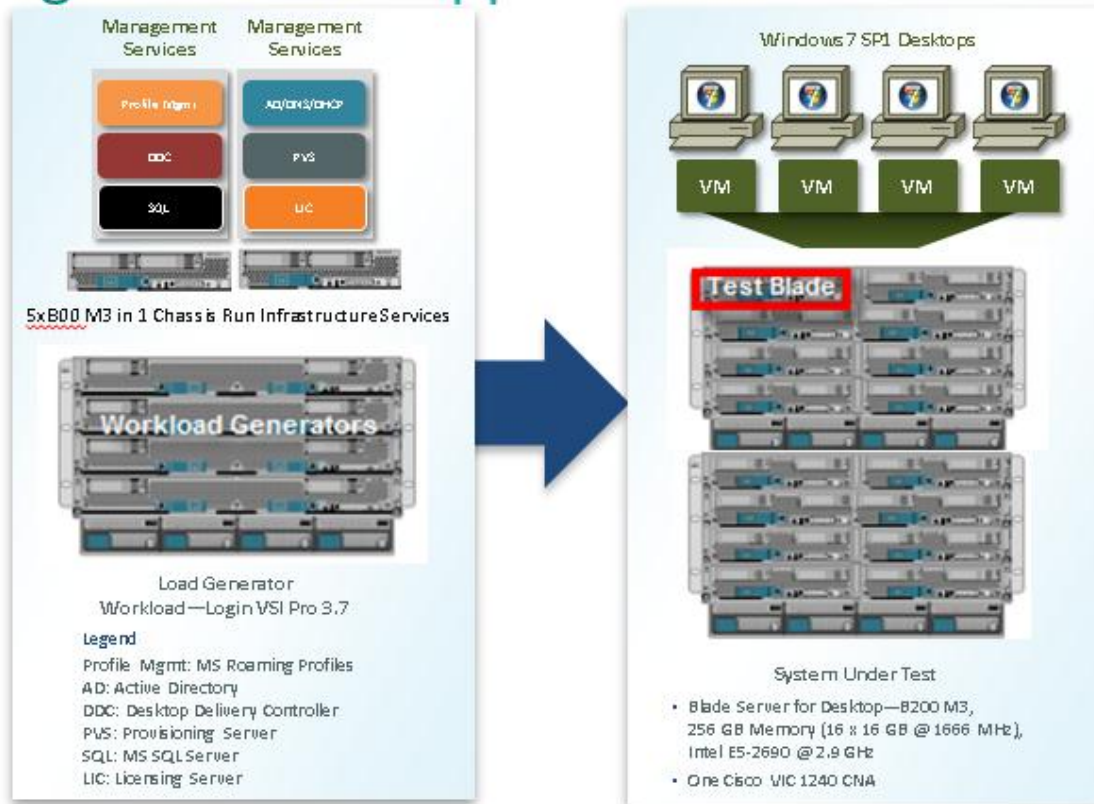


Figure 22. Cisco UCS B200 M3 Blade Server for Single Server Scalability XenApp 6.5 with VM-FEX

Cisco UCS B200 M3 Blade Server Single Blade XenApp Test Result– 200 Users



Hardware components

- 1 X Cisco UCS B200-M3 (E5-2690 @ 2.9 GHz) blade server with 256GB of memory (16 GB X 16 DIMMS @ 1600 MHz) Windows 7 SP1 Virtual Desktop hosts
- 5 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 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 B200 M3)
- 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, 48 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 2.1(1a)
- Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.1 VDI Hosts
- Citrix XenDesktop 5.6 Feature Pack 1
- Citrix XenApp 6.5
- Citrix Provisioning Server 6.1
- Citrix User Profile Manager
- Microsoft Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 17 GB/VM

8.2 Cisco UCS Configuration for Cluster Tests

Figure 23. Seven Blade Cluster XenDesktop 5.6 with Provisioning Server Write-Cache on Blade SSDs

Cisco UCS B200 M3 Blade Servers

7-Blades - 1085 XD 5.6 Users Tier 0 Write-Cache

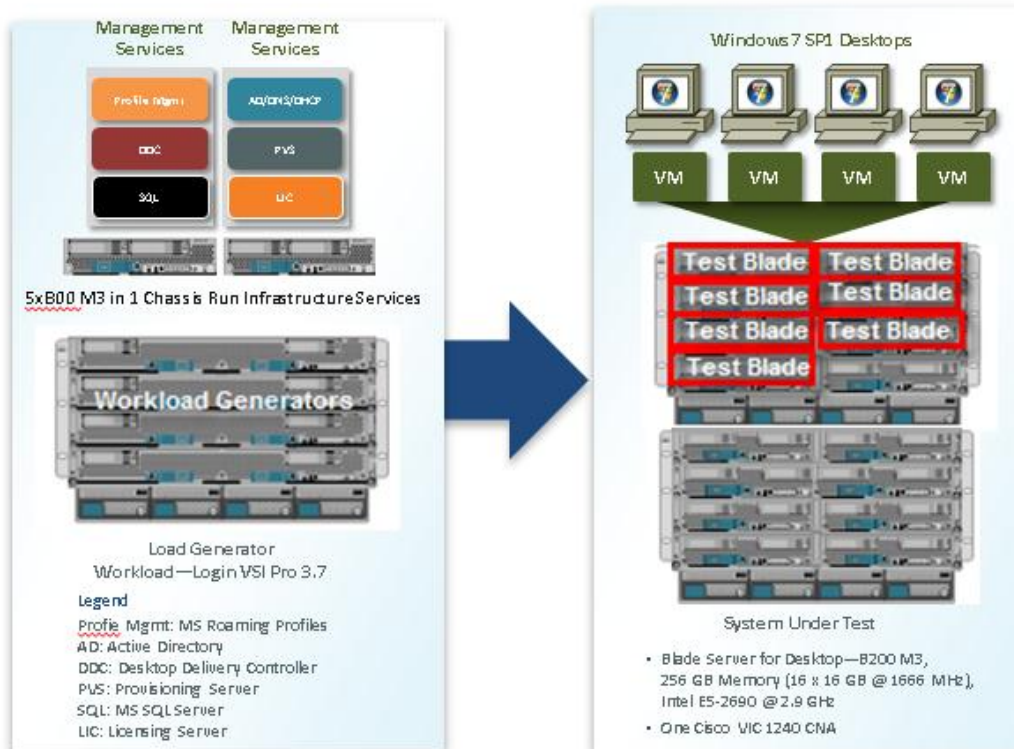


Figure 24. Seven Blade Cluster XenDesktop 5.6 with Personal vDisk

Cisco UCS B200 M3 Blade Servers 7-Blades - 1015 XD 5.6 with PvD Users

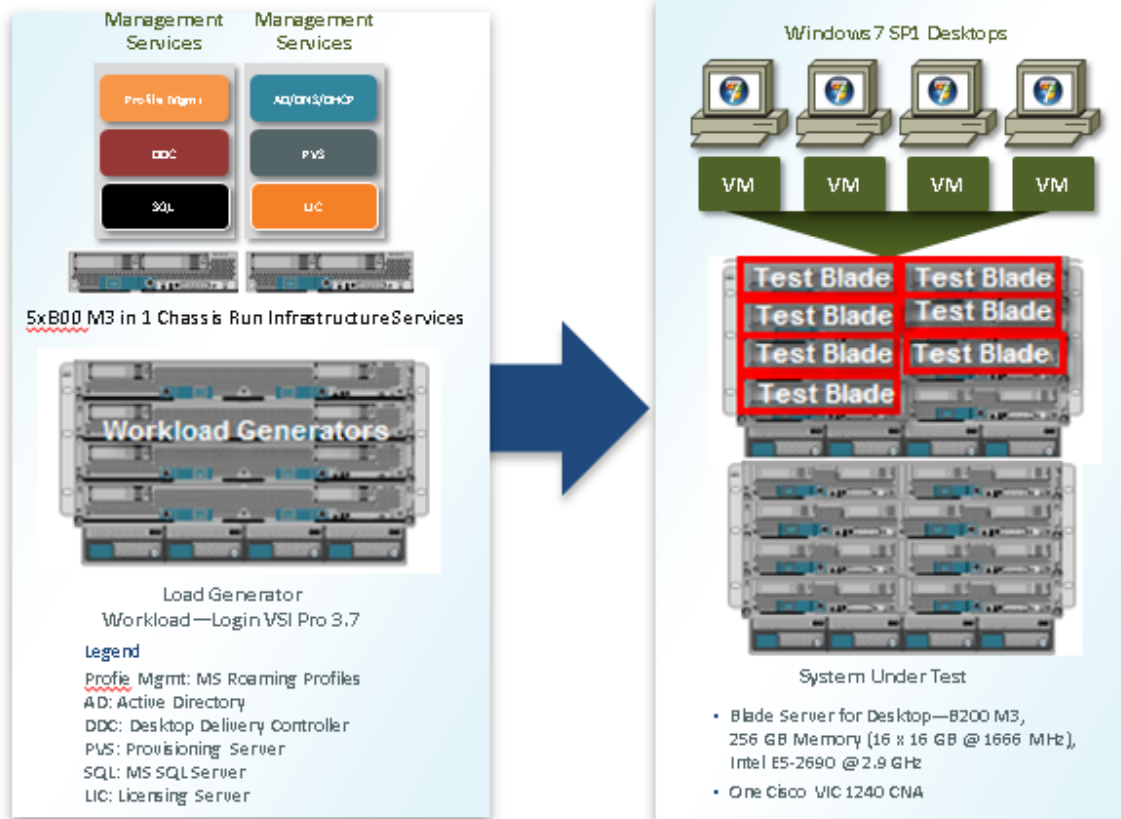
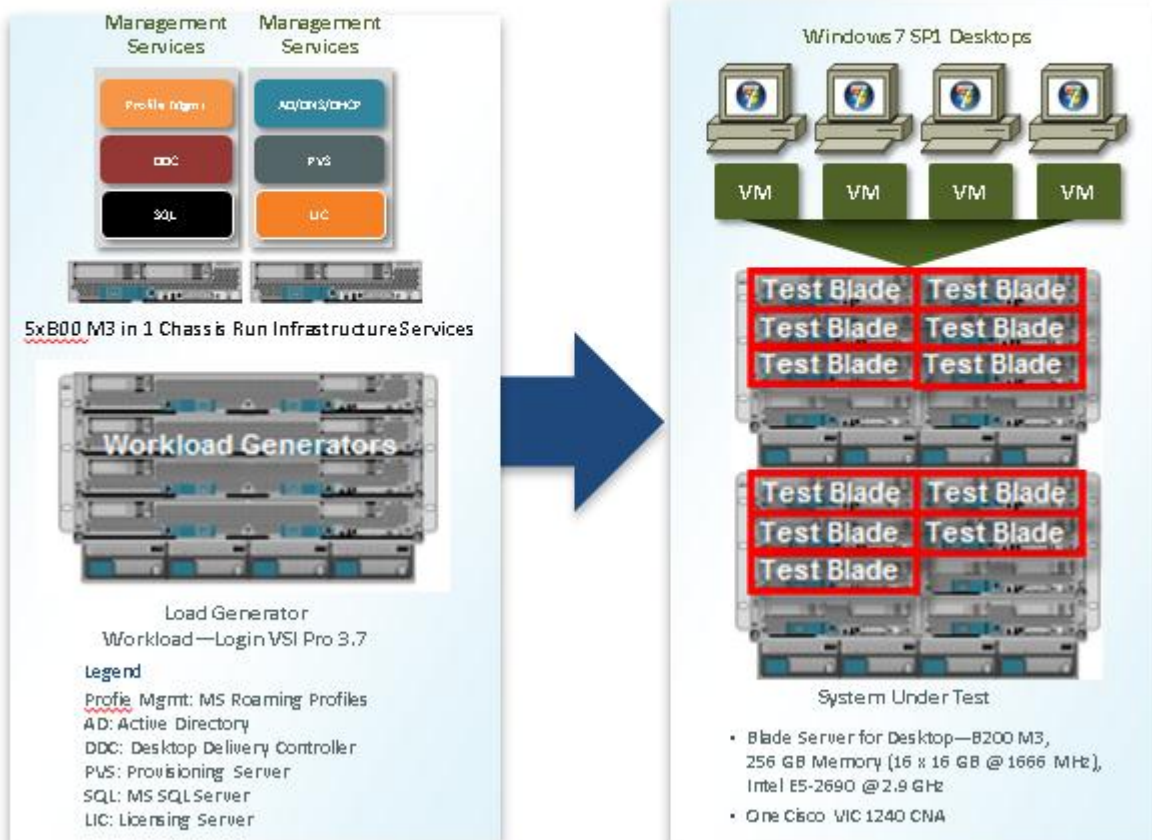


Figure 25. **Eleven Blade Cluster XenApp 6.5**

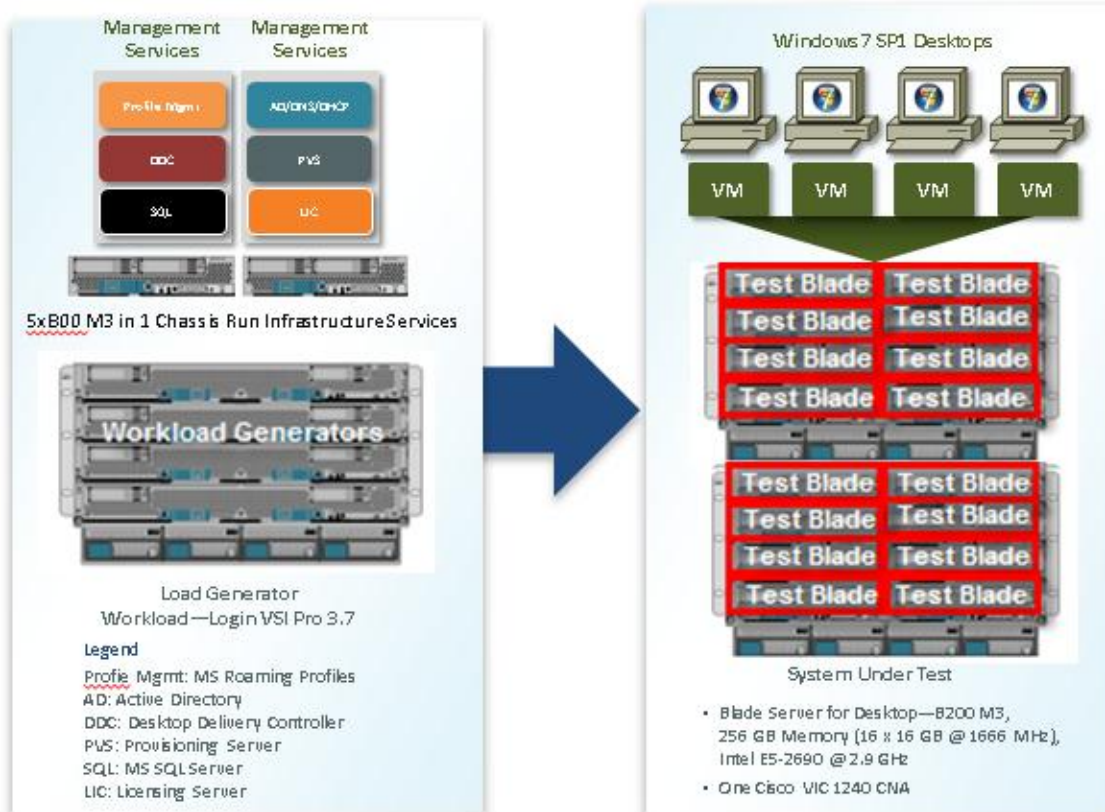
Cisco UCS B200 M3 Blade Servers 11-Blades - 2000 XA 6.5 Users



8.3 Cisco UCS Configuration for Four Chassis – Twenty Five 4000 Mixed Workload Blade Test

Figure 26. Four Chassis Test Configuration-25 B200 M3 Blade Servers

Cisco UCS B200 M3 Blade Servers 25-Blade - 4000 Mixed Workload Users



Hardware Components:

- 25 X Cisco UCS B200-M032 (E5-2690 @ 2.9 GHz) blade server with 256GB of memory (16 GB X 16 DIMMS @ 1600 MHz) Windows 7 SP1 Virtual Desktop hosts
- 5 X Cisco UCS B200-M3 (E5-2665) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz) Infrastructure Servers
- 24 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 B200 M3)
- 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, 48 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
- Citrix XenDesktop 5.6 Feature Pack 1
- Citrix XenApp 6.5
- Citrix Provisioning Server 6.1
- Citrix User Profile Manager
- Microsoft Windows 7 SP1 32 bit, 1vCPU, 1.5 GB of memory, 17 GB/VM

8.4 Testing Methodology and Success Criteria

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

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

8.4.1 Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the View 5.1 environment and executing a typical end-user workflow. To generate load within the environment, an auxiliary software application was required to generate the end user connection to the View environment, to provide unique user credentials, to initiate the workload, and to evaluate the end user experience.

In the Hosted VDI test environment, sessions launchers were used simulate multiple users making a direct connection to View 5.1 via a VMware PCOIP protocol connection.

8.4.2 User Workload Simulation - LoginVSI From Login VSI Inc.

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

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

The Login Virtual Session Indexer ([Login VSI Inc' Login VSI 3.7](#)) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and hence allows customers to easily replicate the testing results in their environment. **Note:** In this testing, we utilized the tool to benchmark our VDI environment only.



Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.

Login VSI simulates a medium workload user (also known as knowledge worker) running generic applications such as: Microsoft Office 2007 or 2010, Internet Explorer 8 including a Flash video applet and Adobe Acrobat Reader (Note: For the purposes of this test, applications were installed locally, not streamed by ThinApp).

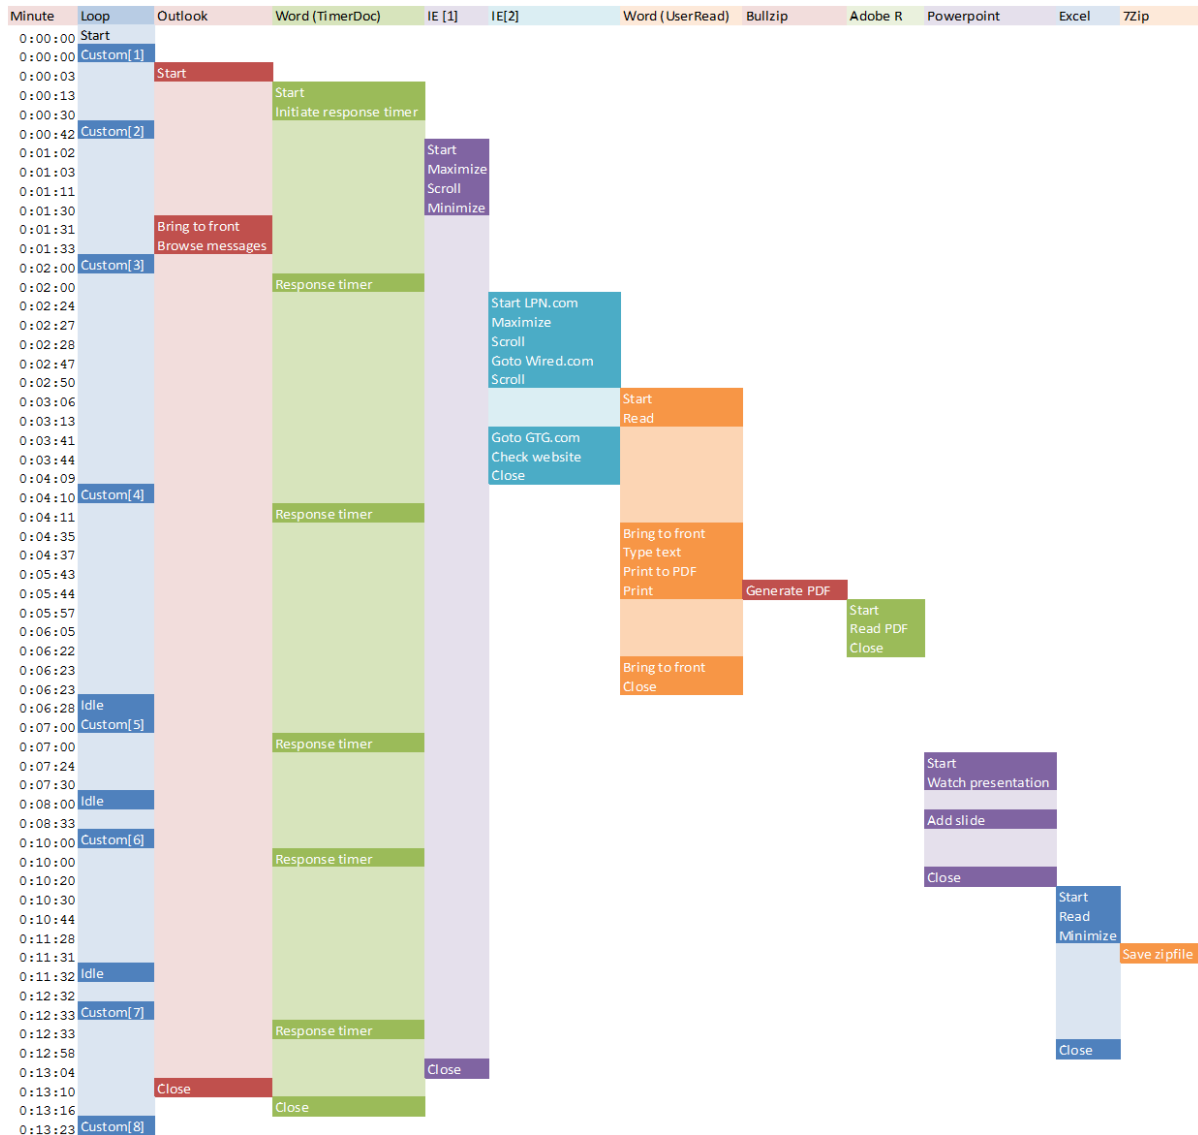
Like real users, the scripted Login VSI session will leave multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. This workload emulated a medium knowledge working using Office, IE, printing and PDF viewing.

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

Each loop will open and use:

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

A graphical representation of the medium workload is shown below.



You can obtain additional information and a free test license from <http://www.loginvsi.com>.

8.4.3 Testing Procedure

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

8.4.3.1 Pre-Test Setup for Single and Multi-Blade Testing

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

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

All VMware ESXi 5.0 VDI host blades to be tested were restarted prior to each test cycle.



8.4.3.2 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 195 single server users or 600 full scale test users to become active within 2 minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Parallel Launching method is used for all single server and scale testing. This assures that our tests represent real-world scenarios. (**Note:** The Login VSI Sequential Launching method allows the CPU, storage and network components to rest between each logins. This does not produce results that are consistent with the real-world scenarios that our Customers run in.)

For each of the three consecutive runs on single server (195 User) and 4 and 5 server (500 and 600 User) tests, the same process was followed:

1. Time 0:00:00 Started ESXTOP Logging on the following systems:
 - i. VDI Host Blades used in test run
 - ii. DDCs used in test run
 - iii. Profile Server(s) used in test run
 - iv. SQL Server(s) used in test run
 - v. 3 Launcher VMs
2. Time 0:00:10 Started EMC Basic Performance Logging on SPs
3. Time 0:00:15 Started EMC NFS Performance Logging on Datamovers (Unmanaged NFS Variant Only)
4. Time 0:05 Take 195, 500 or 600 desktops out of maintenance mode on View Administrator
5. Time 0:06 First machines boot
6. Time 0:26 195, 500 or 600 desktops booted on 1 or 5 servers
7. Time 0:28 195, 500 or 600 desktops available on 1 or 5 servers
8. Time 1:28 Start Login VSI 3.6 Test with 195, 500 or 600 desktops utilizing 7, 17 or 20 Launchers
9. Time 1:58 195, 500 or 600 sessions launched
10. Time 2:00 195, 500 or 600 sessions active
11. Time 2:15 Login VSI Test Ends
12. Time 2:30 195, 500 or 600 sessions logged off
13. Time 2:35 All logging terminated

8.4.4 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, VSImax. The Login VSImax evaluates the user response time during increasing user load and assesses the successful start-to-finish execution of all the initiated virtual desktop sessions.

8.4.4.1 Login VSImax

VSImax represents the maximum number of users the environment can handle before serious performance degradation occurs. VSImax 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. VSImax is reached when the response times reaches or exceeds 4000ms for 6 consecutive occurrences. If VSImax 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% decrease in the number of users passing the test versus Login VSI 2.0 on the same server and storage hardware.

8.4.4.2 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 \ul style="list-style-type: none;"> - 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.

8.4.4.3 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.

Five individual measurements are weighted to better support this approach:

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

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

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

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

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

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

Especially when application virtualization is used, the baseline response time can wildly vary per vendor and streaming strategy. Therefore it is recommend to use VSImax Dynamic when comparisons are made with application


virtualization or anti-virus agents. The resulting VSImax Dynamic scores are aligned again with saturation on a CPU, Memory or Disk level, also when the baseline response time are relatively high.

8.4.4.5 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 27. 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 VSImax 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 VSI Inc, 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 VSI Inc 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 View Administrator will be monitored throughout the steady state to insure that:



- 1) All running sessions report In Use throughout the steady state
 - 2) No sessions move to Agent unreachable or Disconnected state at any time during Steady State
- e) 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.
- 5) 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 Citrix XenDesktop 5.6 Hosted VM and Citrix XenApp 6.5 Shared Hosted Desktop Mixed Workload on Cisco UCS B200 M3 Blades, EMC VNX7500 and VMware ESXi 5.1 Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 5.6 FP1 Hosted VM, Hosted VM with Personal vDisk and Citrix XenApp 6.5 Shared Hosted Desktop models and Citrix Provisioning Services 6.1 using ESXi 5.1 and vCenter 5.1 to virtualize Microsoft Windows 7 SP1 desktops and Microsoft 2008R2 Windows Servers on Cisco UCS B200 M3 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 with XenDesktop 5.6 Hosted VM, we ran 200 Medium with flash Windows 7 SP1 sessions on a single blade. The consistent Login VSI score was achieved on three consecutive runs and is shown below.


Figure 28. **Login VSI Max Reached: 177 Users XenDesktop 5.6 Hosted VDI with PVS write-cache on SSD**






To reach the Login VSI Max with XenDesktop 5.6 Hosted VM with Personal vDisk, we ran 180 Medium with flash Windows 7 SP1 sessions on a single blade. The consistent Login VSI score was achieved on three consecutive runs and is shown below.

Figure 29. **Login VSI Max Reached: 158 Users XenDesktop 5.6 with Personal vDisk**




To reach the Login VSI Max with Citrix XenApp 6.5 Shared Hosted Desktop, we ran 200 Medium with flash Windows 7 SP1 sessions on a single blade hosting 8 XenApp Virtual Machines. The consistent Login VSI score was achieved on three consecutive runs and is shown below.

Figure 30. Login VSI Max Reached: 180 Users XenApp 6.5 without Cisco VM-FEX



To reach the Login VSI Max with Citrix XenApp 6.5 Shared Hosted Desktop (VM FEX), we ran 248 Medium with flash Windows 7 SP1 sessions on a single blade hosting 8 XenApp Virtual Machines. The consistent Login VSI score was achieved on three consecutive runs and is shown below.

Figure 31. Login VSI Max Reached: 220 Users XenApp 6.5 with Cisco VM-FEX



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

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.) 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.

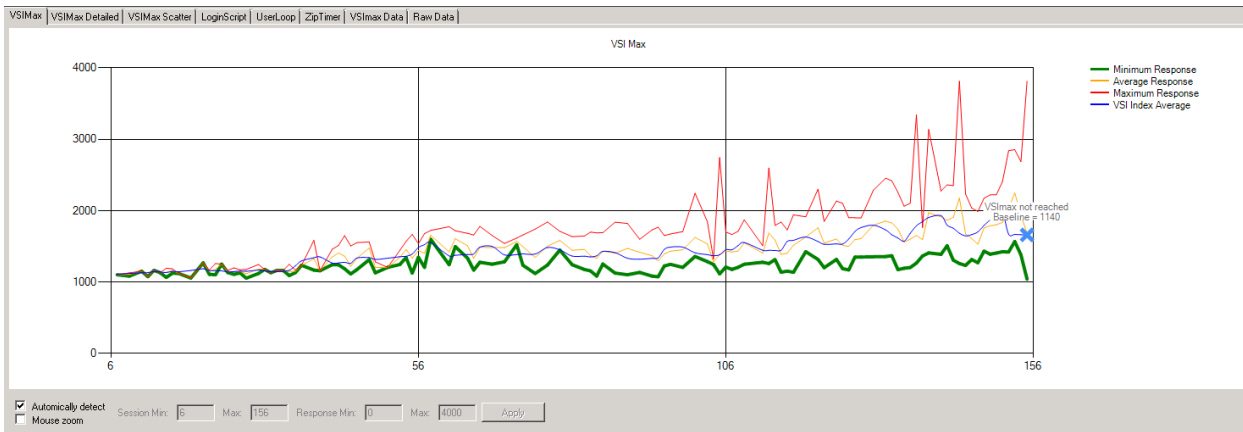


Given adequate storage capability, the CPU utilization determined the maximum recommended VM density per blade for the single server scalability testing.


9.1.1 Pooled XenDesktop 5.6 Hosted VM with Tier 0 (SSD) Storage

This section details the results from the XenDesktop 5.6 Hosted VM single blade server validation testing.

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



Test Phase	Boot storm Start	Boot storm End	Test Start	All Users Logged In	Log Off Start	All Users Logged Off
Time	5:07PM	5:37PM	5:58PM	6:29PM	6:47PM	7:01PM









Figure 32. XenDesktop Controller




Figure 33. XenDesktop Controller




Figure 34. XenDesktop Controller




Figure 35. XenDesktop Controller




Figure 36. Provisioning Services




Figure 37. Provisioning Services




Figure 38. Provisioning Services





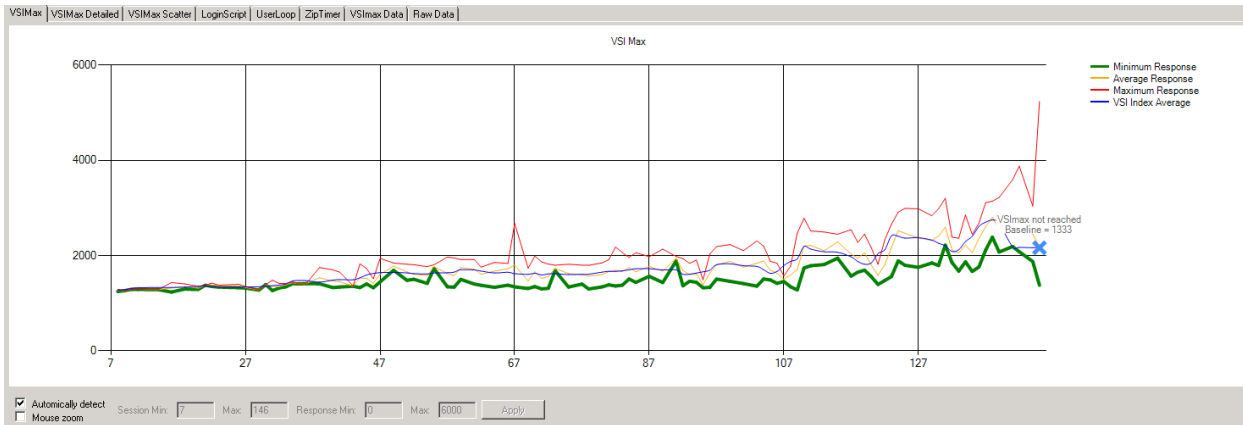
Figure 39. Provisioning Services




9.1.2 Citrix XenDesktop 5.6 Hosted VM with Personal vDisk

This section details the results from the XenDesktop 5.6 Hosted VM with Personal vDisk single blade server validation testing.

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



Test Phase	Boot storm Start	Boot storm End	Test Start	All Users Logged In	Log Off Start	All Users Logged Off
Time	4:07PM	4:34PM	5:35PM	6:05PM	6:23PM	6:37PM









Figure 40. XenDesktop Controller




Figure 41. XenDesktop Controller




Figure 42. XenDesktop Controller




Figure 43. XenDesktop Controller






Figure 44. Provisioning Services




Figure 45. Provisioning Services




Figure 46. Provisioning Services





Figure 47. Provisioning Services




9.1.3 Citrix XenApp 6.5 Shared Hosted Desktop


This section details the results from the XenApp 6.5 Shared Hosted Desktop single blade server hosting 8 XenApp Virtual Machines streamed by 3 Provisioning servers validation testing.

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

Note: Boot storm data is not provided due to the size of the Virtual Machines being hosted.



Test Phase	Boot storm Start	Boot storm End	Test Start	All Users Logged In	Log Off Start	All Users Logged Off
Time	10:00AM	10:05PM	11:00AM	11:31AM	11:51AM	12:02PM









Figure 48. Provisioning Services






Figure 49. Provisioning Services




Figure 50. Provisioning Services







Figure 51. Provisioning Services




9.1.4 Citrix XenApp 6.5 Shared Hosted Desktop (VM-FEX)

This section details the results from the XenApp 6.5 Shared Hosted Desktop single blade server with VM-FEX hosting 8 XenApp Virtual Machines streamed by 3 Provisioning servers validation testing.

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


Note: Boot storm data is not provided due to the size of the Virtual Machines being hosted.



Test Phase	Boot storm Start	Boot storm End	Test Start	All Users Logged In	Log Off Start	All Users Logged Off
Time	10:04AM	10:06AM	11:06AM	11:36AM	11:54AM	12:07PM



\\B200M3-CH4-BL08\Memory\NonKernel
MBytes



\\B200M3-CH4-BL08\Physical Cpu(_Total)\%
Core Util Time





Figure 52. Provisioning Services




Figure 53. Provisioning Services






Figure 54. Provisioning Services




Figure 55. Provisioning Services




Figure 56. XenApp Server




Figure 57. XenApp Server






Figure 58. XenApp Server





Figure 59. XenApp Server




9.2 Cisco UCS Test Configuration for Single Cluster Scalability Test Results



This section details the results from the XenDesktop 5.6 Hosted VM on Tier 0 (SSD) Storage, XenDesktop 5.6 Hosted VM with Personal vDisk, and XenApp 6.5 Shared Hosted Desktop with VM-FEX individual cluster 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).



9.2.1 Pooled XenDesktop 5.6 Hosted VM with Tier 0 (SSD) Storage



The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B200-M3 Blade during the seven blade, 1085 User test.






Figure 60. XenDesktop Controller




Figure 61. XenDesktop Controller




Figure 62. XenDesktop Controller






Figure 63. XenDesktop Controller




Figure 64. Provisioning Services

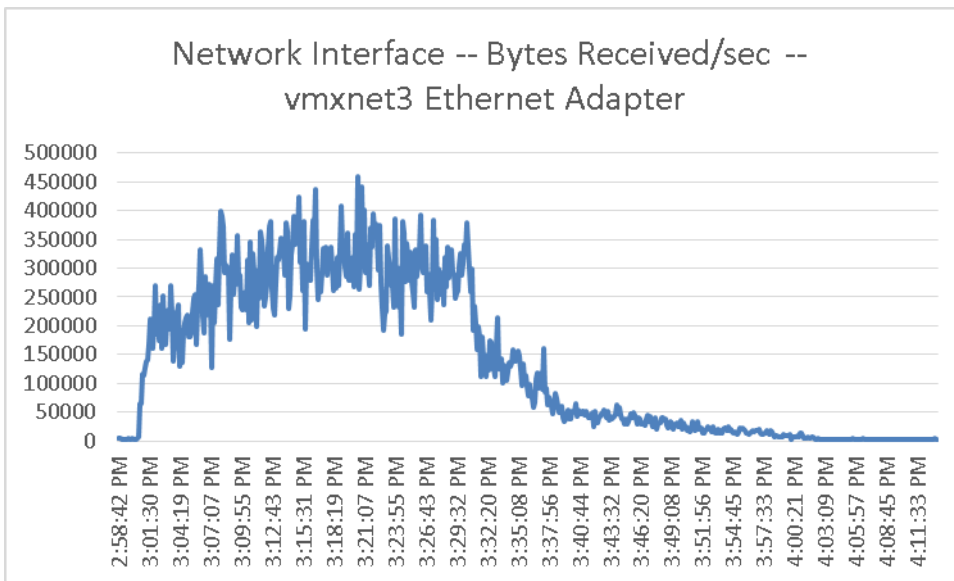


Figure 65. Provisioning Services




Figure 66. Provisioning Services








Figure 67. Provisioning Services








9.2.2 Citrix XenDesktop 5.6 Shared Hosted Desktop with Personal vDisk



The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B200-M3 Blade during the seven blade, 1015 User test.





The following graphs detail performance of the EMC VNX7500 during this test run.
Figure 68. **EMC VNX7500 NFS Read Operations Per Second**






Figure 69. EMC VNX7500 NFS Write Operations Per Second




Figure 70. EMC VNX7500 Storage Processor Total IOPS




Figure 71. XenDesktop Controller




Figure 72. XenDesktop Controller




Figure 73. XenDesktop Controller




Figure 74. XenDesktop Controller






Figure 75. Provisioning Services




Figure 76. Provisioning Services






Figure 77. Provisioning Services








Figure 78. Provisioning Services




9.2.3 Citrix XenApp 6.5 Shared Hosted Desktop



The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B200-M3 Blade during the eleven blade, 2000 User test.





The following graphs detail performance of the EMC VNX 7500 during the test run.

Figure 79. EMC VNX7500 Storage Processor Total IOPS






Figure 80. EMC VNX7500 NFS Read Operations Per Second




Figure 81. EMC VNX7500 NFS Write Operations Per Second




Figure 82. Provisioning Services




Figure 83. Provisioning Services






Figure 84. Provisioning Services




Figure 85. Provisioning Services




Figure 86. XenApp Server




Figure 87. XenApp Server




Figure 88. XenApp Server





Figure 89. XenApp Server



9.3 Cisco UCS Test Configuration for 4100 Desktop Mixed Workload Scalability Test Results

This section details the results from the XenDesktop 5.6 Hosted VM on Tier 0 (SSD) Storage, XenDesktop 5.6 Hosted VM with Personal vDisk, and XenApp 6.5 Shared Hosted Desktop with VM-FEX 4100 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).


We ran the multi-server test with 3 different user density based on based on VDI type N+1 fault tolerance on a cluster basis 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 4100 user environment.

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

Figure 90. **4100 Desktop Sessions on VMWare ESXi 5.1**



The following graphs detail CPU, Memory, Disk and Network performance on a representative Cisco UCS B200-M3 Blade during the twenty-five blade, 4100 User test. (Representative results for all blades in every of the vCenter clusters can be found in Appendix C).

Figure 91. **DC-VDA1 ESXi 5.1 server**




Figure 92. DC-VDA1 ESXi 5.1 Server




Figure 93. DC-VDA1 ESXi 5.1 server






Figure 94. DC-VDA2 ESXi 5.1 server




Figure 95. DC-VDA2 ESXi 5.1 server






Figure 96. DC-VDA2 ESXi 5.1 server




Figure 97. DC-VDA3 ESXi 5.1 server






Figure 98. DC-VDA3 ESXi 5.1 server





Figure 99. DC-VDA3 ESXi 5.1 server



The following graphs detail performance of the EMC VNX 7500 during the full scale test run.



Figure 100. EMC VNX7500 NFS Read Operations Per Second




Figure 101. EMC VNX7500 NFS Write Operations Per Second




Figure 102. EMC VNX7500 Storage Processor Total IOPS




Figure 103. XenDesktop Controller




Figure 104. XenDesktop Controller




Figure 105. XenDesktop Controller




Figure 106. XenDesktop Controller




Figure 107. Desktop Provisioning Services




Figure 108. Desktop Provisioning Services




Figure 109. Desktop Provisioning Services




Figure 110. Desktop Provisioning Services




Figure 111. XenApp Provisioning Services




Figure 112. XenApp Provisioning Services




Figure 113. XenApp Provisioning Services




Figure 114. XenApp Provisioning Services




Figure 115. XenApp Server

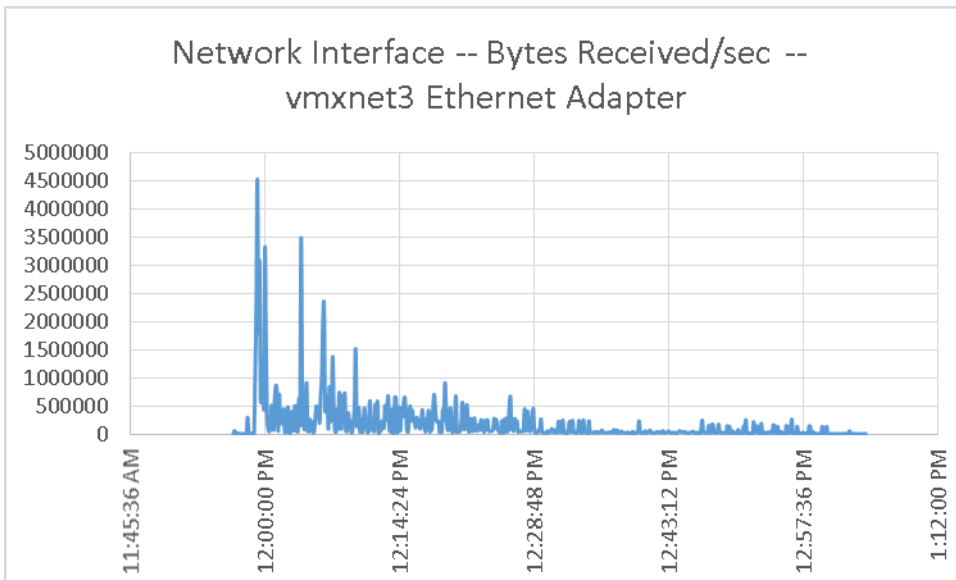


Figure 116. XenApp Server

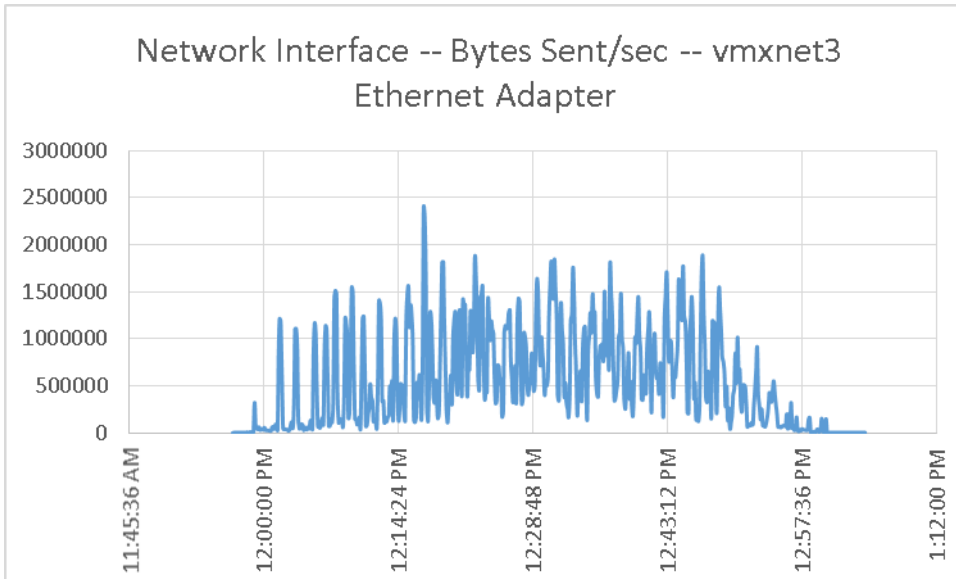


Figure 117. XenApp Server





Figure 118. XenApp Server



10 Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 5000 User, 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 UCS 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 Cisco UCS domain can grow to 160 blades.
- With Cisco UCS 2.1 management software, released late in November 2012, each Cisco UCS 2.1 Management domain is extensively manageable by Cisco UCS Central, our new manager of managers, vastly increasing the reach of the UCS system.
- As scale grows, the value of the combined Cisco UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100% of the time.
- To accommodate the Cisco Nexus 5500 upstream connectivity in the way we describe in the LAN and SAN Configuration section, we need four Ethernet uplinks and two Fibre Channel uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we can calculate number of desktops can be hosted in a single Cisco UCS domain. Assuming eight links per chassis, four to each 6248, scaling beyond 10 chassis would require a pair of Cisco UCS 6296 fabric interconnects. A 20,000 virtual desktop building block, with its support infrastructure services can be built out of the RA described in



this study with eight links per chassis and 20 Cisco UCS chassis comprised of seven Cisco UCS B230 M2 and one Cisco UCS 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 NetApp scaling section. Please refer the NetApp 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 in particular worth mentioning here.

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.

10.3 EMC VNX Storage Guidelines for Citrix 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 detailed information, refer to the VMware documentation sited 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

Cisco Unified Computing System Manager Home Page

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

Cisco UCS C220 M3 Rack Server Resources

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

Cisco UCS 6200 Series Fabric Interconnects

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

Cisco Nexus 2232PP 10GE Fabric Extender

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

Cisco Nexus 5500 Series Switches Resources

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

[Download Software for UCS C220 M3 Rack Server](#)

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

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

[http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=1.4\(4k\)&reliand=AVAILABLE&rellifecycle=&reltype=latest](http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=1.4(4k)&reliand=AVAILABLE&rellifecycle=&reltype=latest)

Download Cisco UCS Central Software Version 1.0(1a)

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

11.2 Citrix Reference Documents

Product Documentation - <http://support.citrix.com/proddocs/topic/infocenter/ic-how-to-use.html>

11.2.1 XenDesktop 5.6

Modular Reference Architecture - <http://support.citrix.com/article/CTX133162>



11.2.3 XenApp 6.5

Hotfix Rollup Pack 1 for Citrix XenApp 6.5 for Microsoft Windows Server 2008 R2 - <http://support.citrix.com/article/CTX132122>

11.2.4 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>

11.2.5 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 VSPEX End-User Computing, Citrix XenDesktop 5.6, VMware vSphere 5.1 for up to 2,000 Virtual Desktops Enabled by EMC VNX and EMC Next-Generation Backup—Proven Infrastructure
- EMC VSPEX End-User Computing, Citrix XenDesktop 5.6, VMware vSphere 5.1 for up to 250 Virtual Desktops Enabled by EMC VNXe and EMC Next-Generation Backup—Proven Infrastructure
- 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
- VMware vSphere ESXi and vCenter Server 5 Documentation: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021514
- VMware vCenter Management Webservices features do not function properly: - http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1039180
- VMware® vCenter Server™ 5.1 Database Performance Improvements and Best Practices for Large-Scale Environments: - <http://www.vmware.com/files/pdf/techpaper/VMware-vCenter-DBPerfBestPractices.pdf>
- Performance Best Practices for VMware vSphere™ 5.0: - http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf

Appendix A Nexus 5548UP Configurations

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 pwwn 50:06:01:60:46:e0:5e:0a
device-alias name VNX7500-B1 pwwn 50:06:01:69:46:e0:5e:0a
device-alias name B200M3-CH2-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:1a
device-alias name B200M3-CH2-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:57
device-alias name B200M3-CH2-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:b7
device-alias name B200M3-CH2-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:48
device-alias name B200M3-CH2-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:a8
device-alias name B200M3-CH2-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:49
device-alias name B200M3-CH2-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:c6
device-alias name B200M3-CH3-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:86
device-alias name B200M3-CH3-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:17
device-alias name B200M3-CH3-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:77
device-alias name B200M3-CH3-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:08
device-alias name B200M3-CH3-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:68
device-alias name B200M3-CH3-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:09
device-alias name B200M3-CH3-SERVER7-fc0 pwwn 20:00:00:25:b5:c1:00:69
device-alias name B200M3-CH3-SERVER8-fc0 pwwn 20:00:00:25:b5:c1:00:a9
device-alias name B200M3-CH4-SERVER1-fc0 pwwn 20:00:00:25:b5:c1:00:a6
device-alias name B200M3-CH4-SERVER2-fc0 pwwn 20:00:00:25:b5:c1:00:37
device-alias name B200M3-CH4-SERVER3-fc0 pwwn 20:00:00:25:b5:c1:00:97
device-alias name B200M3-CH4-SERVER4-fc0 pwwn 20:00:00:25:b5:c1:00:28
device-alias name B200M3-CH4-SERVER5-fc0 pwwn 20:00:00:25:b5:c1:00:88
device-alias name B200M3-CH4-SERVER6-fc0 pwwn 20:00:00:25:b5:c1:00:29
device-alias name B200M3-CH4-SERVER7-fc0 pwwn 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 pwwn 20:00:00:25:b5:c1:00:af
```

```
! [B230M2-CH1-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 B230M2-CH1-SERVER2-FC0 vsan 1
```

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

```
! [B230M2-CH1-SERVER2-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 B230M2-CH1-SERVER3-FC0 vsan 1
```

```
member pwwn 20:00:00:25:b5:c1:00:7f
```

```
! [B230M2-CH1-SERVER3-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 B230M2-CH1-SERVER4-FC0 vsan 1
```

```
member pwwn 20:00:00:25:b5:c1:00:3c
```

```
! [B230M2-CH1-SERVER4-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 B230M2-CH1-SERVER5-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:1d
!   [B230M2-CH1-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 B230M2-CH1-SERVER6-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:ad
!   [B230M2-CH1-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 B230M2-CH1-SERVER7-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:9e
!   [B230M2-CH1-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-CH2-SERVER1-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:1a
!   [B200M3-CH2-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-CH2-SERVER2-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:c6
!   [B200M3-CH2-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-CH2-SERVER3-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:57
```




```
! [B200M3-CH2-SERVER3-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-CH2-SERVER4-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:b7
! [B200M3-CH2-SERVER4-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-CH2-SERVER5-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:48
! [B200M3-CH2-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-CH2-SERVER6-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:b5
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-CH2-SERVER7-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:49
! [B200M3-CH2-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-SERVER1-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:86
! [B200M3-CH3-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-CH3-SERVER2-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:17

! [B200M3-CH3-SERVER2-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-SERVER3-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:77

! [B200M3-CH3-SERVER3-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-SERVER4-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:08

! [B200M3-CH3-SERVER4-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-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 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-SERVER3-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:97
! [B200M3-CH4-SERVER3-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-SERVER4-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:28
! [B200M3-CH4-SERVER4-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-SERVER5-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:88
! [B200M3-CH4-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-CH4-SERVER6-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:29
! [B200M3-CH4-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-CH4-SERVER7-FC0 vsan 1
member pwwn 20:00:00:25:b5:c1:00:89
! [B200M3-CH4-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 B230M2-CH5-SERVER1-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:7a
! [B200M3-CH5-SERVER1-fc0]

zone name B230M2-CH5-SERVER2-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:3b
! [B230M2-CH5-SERVER2-fc0]
```



```
zone name B230M2-CH5-SERVER3-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!    [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!    [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:ab
!    [B230M2-CH5-SERVER3-fc0]
```

```
zone name B230M2-CH5-SERVER4-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!    [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!    [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:9c
!    [B230M2-CH5-SERVER4-fc0]
```

```
zone name B230M2-CH5-SERVER5-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!    [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!    [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:7d
!    [B230M2-CH5-SERVER5-fc0]
```

```
zone name B230M2-CH5-SERVER6-FC0 vsan 1
  member pwwn 50:06:01:69:46:e0:5e:0a
!    [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:5e
!    [B230M2-CH5-SERVER6-fc0]
  member pwwn 50:06:01:60:46:e0:5e:0a
!    [VNX7500-A0]
```

```
zone name B230M2-CH5-SERVER7-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!    [VNX7500-A0]
  member pwwn 20:00:00:25:b5:c1:00:3f
!    [B230M2-CH5-SERVER7-fc0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!    [VNX7500-B1]
```

```
zone name B230M2-CH6-SERVER1-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
```



```
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:3a
! [B230M2-CH6-SERVER1-fc0]

zone name B230M2-CH6-SERVER2-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:aa
! [B230M2-CH6-SERVER2-fc0]

zone name B230M2-CH6-SERVER3-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:7b
! [B230M2-CH6-SERVER3-fc0]

zone name B230M2-CH6-SERVER4-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:5c
! [B230M2-CH6-SERVER4-fc0]

zone name B230M2-CH6-SERVER5-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:3d
! [B230M2-CH6-SERVER5-fc0]

zone name B230M2-CH6-SERVER6-FC0 vsan 1
member pwwn 50:06:01:60:46:e0:5e:0a
! [VNX7500-A0]
member pwwn 50:06:01:69:46:e0:5e:0a
! [VNX7500-B1]
```



```
member pwwn 20:00:00:25:b5:c1:00:1e
!      [B230M2-CH6-SERVER6-fc0]

zone name B230M2-CH6-SERVER7-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:ae
!      [B230M2-CH6-SERVER7-fc0]

zone name B230M2-CH7-SERVER1-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:5a
!      [B230M2-CH7-SERVER1-fc0]

zone name B230M2-CH7-SERVER2-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:1b
!      [B230M2-CH7-SERVER2-fc0]

zone name B230M2-CH7-SERVER3-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:9b
!      [B230M2-CH7-SERVER3-fc0]

zone name B230M2-CH7-SERVER4-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:7c
!      [B230M2-CH7-SERVER4-fc0]
```



```
zone name B230M2-CH7-SERVER5-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!   [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:5d
!   [B230M2-CH7-SERVER5-fc0]
```

```
zone name B230M2-CH7-SERVER6-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!   [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:3e
!   [B230M2-CH7-SERVER6-fc0]
```

```
zone name B230M2-CH7-SERVER7-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!   [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:1f
!   [B230M2-CH7-SERVER7-fc0]
```

```
zone name B230M2-CH8-SERVER1-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!   [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:9a
!   [B230M2-CH8-SERVER1-fc0]
```

```
zone name B230M2-CH8-SERVER2-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!   [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:5b
!   [B230M2-CH8-SERVER2-fc0]
```

```
zone name B230M2-CH8-SERVER3-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!   [VNX7500-A0]
```




```
member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
member pwwn 20:00:00:25:b5:c1:00:1c
!      [B230M2-CH8-SERVER3-fc0]

zone name B230M2-CH8-SERVER4-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:ac
!      [B230M2-CH8-SERVER4-fc0]

zone name B230M2-CH8-SERVER5-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:9d
!      [B230M2-CH8-SERVER5-fc0]

zone name B230M2-CH8-SERVER6-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:7e
!      [B230M2-CH8-SERVER6-fc0]

zone name B230M2-CH8-SERVER7-FC0 vsan 1
  member pwwn 50:06:01:60:46:e0:5e:0a
!      [VNX7500-A0]
  member pwwn 50:06:01:69:46:e0:5e:0a
!      [VNX7500-B1]
  member pwwn 20:00:00:25:b5:c1:00:5f
!      [B230M2-CH8-SERVER7-fc0]

zone name B230M2-CH8-SERVER8-FC0 vsan 1
  member pwwn 20:00:00:25:b5:c1:00:46
!      [B230M2-CH8-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 B230M2-CH1-SERVER8-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:66

! [B230M2-CH1-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 B230M2-CH6-SERVER8-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:06

! [B230M2-CH6-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 B230M2-CH7-SERVER8-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:26

! [B230M2-CH7-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-CH2-SERVER8-FC0 vsan 1

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

! [VNX7500-A0]

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

! [VNX7500-B1]

member pwwn 20:00:00:25:b5:c1:00:c6

! [B200M3-CH2-SERVER8-fc0]

zone name B200M3-CH4-SERVER8-FC0 vsan 1

member pwwn 20:00:00:25:b5:c1:00:95

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

! [VNX7500-A0]

member pwwn 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

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

device-alias name VNX7500-B0 pwwn 50:06:01:68:46:e0:5e:0a

device-alias name B200M3-CH2-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:0a

device-alias name B200M3-CH2-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:67

device-alias name B200M3-CH2-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:c7

device-alias name B200M3-CH2-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:58

device-alias name B200M3-CH2-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:b8

device-alias name B200M3-CH2-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:59

device-alias name B200M3-CH2-SERVER8-fc1 pwwn 20:00:00:25:b5:c1:00:07

device-alias name B200M3-CH3-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:96

device-alias name B200M3-CH3-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:27

device-alias name B200M3-CH3-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:87

device-alias name B200M3-CH3-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:18

device-alias name B200M3-CH3-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:78

device-alias name B200M3-CH3-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:19

device-alias name B200M3-CH3-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:79

device-alias name B200M3-CH3-SERVER8-fc1 pwwn 20:00:00:25:b5:c1:00:b9

device-alias name B200M3-CH4-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:b6

device-alias name B200M3-CH4-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:47

device-alias name B200M3-CH4-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:a7



device-alias name B200M3-CH4-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:38

device-alias name B200M3-CH4-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:98

device-alias name B200M3-CH4-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:39

device-alias name B200M3-CH4-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:99

device-alias name B230M2-CH1-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:bf

device-alias name B230M2-CH1-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:8f

device-alias name B230M2-CH1-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:6f

device-alias name B230M2-CH1-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:2c

device-alias name B230M2-CH1-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:0d

device-alias name B230M2-CH1-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:bd

device-alias name B230M2-CH1-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:8e

device-alias name B230M2-CH1-SERVER8-fc1 pwwn 20:00:00:25:b5:c1:00:76

device-alias name B230M2-CH5-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:6a

device-alias name B230M2-CH5-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:2b

device-alias name B230M2-CH5-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:bb

device-alias name B230M2-CH5-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:8c

device-alias name B230M2-CH5-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:6d

device-alias name B230M2-CH5-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:4e

device-alias name B230M2-CH5-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:2f

device-alias name B230M2-CH6-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:2a

device-alias name B230M2-CH6-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:ba

device-alias name B230M2-CH6-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:6b

device-alias name B230M2-CH6-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:4c

device-alias name B230M2-CH6-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:2d

device-alias name B230M2-CH6-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:0e

device-alias name B230M2-CH6-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:be



```
device-alias name B230M2-CH6-SERVER8-fc1 pwwn 20:00:00:25:b5:c1:00:16
device-alias name B230M2-CH7-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:4a
device-alias name B230M2-CH7-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:0b
device-alias name B230M2-CH7-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:8b
device-alias name B230M2-CH7-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:6c
device-alias name B230M2-CH7-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:4d
device-alias name B230M2-CH7-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:2e
device-alias name B230M2-CH7-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:0f
device-alias name B230M2-CH7-SERVER8-fc1 pwwn 20:00:00:25:b5:c1:00:36
device-alias name B230M2-CH8-SERVER1-fc1 pwwn 20:00:00:25:b5:c1:00:8a
device-alias name B230M2-CH8-SERVER2-fc1 pwwn 20:00:00:25:b5:c1:00:4b
device-alias name B230M2-CH8-SERVER3-fc1 pwwn 20:00:00:25:b5:c1:00:0c
device-alias name B230M2-CH8-SERVER4-fc1 pwwn 20:00:00:25:b5:c1:00:bc
device-alias name B230M2-CH8-SERVER5-fc1 pwwn 20:00:00:25:b5:c1:00:8d
device-alias name B230M2-CH8-SERVER6-fc1 pwwn 20:00:00:25:b5:c1:00:6e
device-alias name B230M2-CH8-SERVER7-fc1 pwwn 20:00:00:25:b5:c1:00:4f
device-alias name B230M2-CH8-SERVER8-fc1 pwwn 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
! [B200M3-CH4-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:87 fcid 0xad0018 dynamic
! [B200M3-CH3-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:99 fcid 0xad0038 dynamic
! [B200M3-CH4-SERVER7-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:98 fcid 0xad0019 dynamic
! [B200M3-CH4-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:19 fcid 0xad0039 dynamic
! [B200M3-CH3-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:96 fcid 0xad001a dynamic
! [B200M3-CH3-SERVER1-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:18 fcid 0xad003a dynamic
! [B200M3-CH3-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:38 fcid 0xad001b dynamic
```



```
! [B200M3-CH4-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:39 fcid 0xad003b dynamic
! [B200M3-CH4-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:a7 fcid 0xad001c dynamic
! [B200M3-CH4-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:c7 fcid 0xad003c dynamic
! [B200M3-CH2-SERVER4-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:b8 fcid 0xad003d dynamic
! [B200M3-CH2-SERVER6-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:67 fcid 0xad001d dynamic
! [B200M3-CH2-SERVER3-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:58 fcid 0xad001e dynamic
! [B200M3-CH2-SERVER5-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:76 fcid 0xad001f dynamic
! [B230M2-CH1-SERVER8-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:56 fcid 0xad003e dynamic
! [B230M2-CH8-SERVER8-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:36 fcid 0xad003f dynamic
! [B230M2-CH7-SERVER8-fc1]
vsan 1 wwn 20:00:00:25:b5:c1:00:16 fcid 0xad0040 dynamic
! [B230M2-CH6-SERVER8-fc1]
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 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]
```

```
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]

zone name B230M2-CH1-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:6f

! [B230M2-CH1-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-CH1-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:2c

! [B230M2-CH1-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-CH1-SERVER5-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0d

! [B230M2-CH1-SERVER5-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-CH1-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:bd

! [B230M2-CH1-SERVER6-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-CH1-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:8e

! [B230M2-CH1-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 B200M3-CH2-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0a

! [B200M3-CH2-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 B200M3-CH2-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:07

! [B200M3-CH2-SERVER8-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 B200M3-CH2-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:67

! [B200M3-CH2-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 B200M3-CH2-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:c7

! [B200M3-CH2-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 B200M3-CH2-SERVER5-FC1 vsan 1



member pwwn 20:00:00:25:b5:c1:00:58

! [B200M3-CH2-SERVER5-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 B200M3-CH2-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:c5

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 B200M3-CH2-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:59

! [B200M3-CH2-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 B200M3-CH3-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:96

! [B200M3-CH3-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 B200M3-CH3-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:27

! [B200M3-CH3-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 B200M3-CH3-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:87

! [B200M3-CH3-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 B200M3-CH3-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:18

! [B200M3-CH3-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 B200M3-CH3-SERVER5-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:78

! [B200M3-CH3-SERVER5-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 B200M3-CH3-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:19

! [B200M3-CH3-SERVER6-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 B200M3-CH3-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:79

! [B200M3-CH3-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 B200M3-CH3-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:b9

! [B200M3-CH3-SERVER8-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 B200M3-CH4-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:b6

! [B200M3-CH4-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 B200M3-CH4-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:47

! [B200M3-CH4-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 B200M3-CH4-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:a7



! [B200M3-CH4-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 B200M3-CH4-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:38

! [B200M3-CH4-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 B200M3-CH4-SERVER5-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:98

! [B200M3-CH4-SERVER5-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 B200M3-CH4-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:39

! [B200M3-CH4-SERVER6-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 B200M3-CH4-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:99

! [B200M3-CH4-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-CH5-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:6a

! [B230M2-CH5-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-CH5-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:2b

! [B230M2-CH5-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-CH5-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:bb

! [B230M2-CH5-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-CH5-SERVER4-FC1 vsan 1

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

! [VNX7500-A1]

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

! [VNX7500-B0]

member pwwn 20:00:00:25:b5:c1:00:8c

! [B230M2-CH5-SERVER4-fc1]

zone name B230M2-CH5-SERVER5-FC1 vsan 1

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

! [VNX7500-B0]

member pwwn 20:00:00:25:b5:c1:00:6d

! [B230M2-CH5-SERVER5-fc1]

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

! [VNX7500-A1]



zone name B230M2-CH5-SERVER6-FC1 vsan 1

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

! [VNX7500-A1]

member pwwn 20:00:00:25:b5:c1:00:4e

! [B230M2-CH5-SERVER6-fc1]

member pwwn 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 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-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0e

! [B230M2-CH6-SERVER6-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-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:be

! [B230M2-CH6-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-CH7-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:4a

! [B230M2-CH7-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-CH7-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0b

! [B230M2-CH7-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-CH7-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:8b

! [B230M2-CH7-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-CH7-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:6c

! [B230M2-CH7-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-CH7-SERVER5-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:4d

! [B230M2-CH7-SERVER5-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-CH7-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:2e

! [B230M2-CH7-SERVER6-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-CH7-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0f

! [B230M2-CH7-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-CH8-SERVER1-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:8a



! [B230M2-CH8-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-CH8-SERVER2-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:4b

! [B230M2-CH8-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-CH8-SERVER3-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:0c

! [B230M2-CH8-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-CH8-SERVER4-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:bc

! [B230M2-CH8-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-CH8-SERVER5-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:8d

! [B230M2-CH8-SERVER5-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-CH8-SERVER6-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:6e

! [B230M2-CH8-SERVER6-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-CH8-SERVER7-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:4f

! [B230M2-CH8-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-CH8-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:56

! [B230M2-CH8-SERVER8-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-CH1-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:76

! [B230M2-CH1-SERVER8-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-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:16

! [B230M2-CH6-SERVER8-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-CH7-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:36

! [B230M2-CH7-SERVER8-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 B200M3-CH2-SERVER8-FC1 vsan 1

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

! [VNX7500-A1]

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

! [VNX7500-B0]

member pwwn 20:00:00:25:b5:c1:00:07

! [B200M3-CH2-SERVER8-fc1]

zone name B200M3-CH4-SERVER8-FC1 vsan 1

member pwwn 20:00:00:25:b5:c1:00:a5

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

! [VNX7500-A1]

member pwwn 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

Citrix-VSM-01

Citrix-VSM-02

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 Sample VM-FEX VSM Configuration



Citrix-VMFEX-01



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 D ESXtop Performance Charts for 4100 Session Run



Twenty-five Blades' Data Shown Representing 4100 Sessions.



Run 179










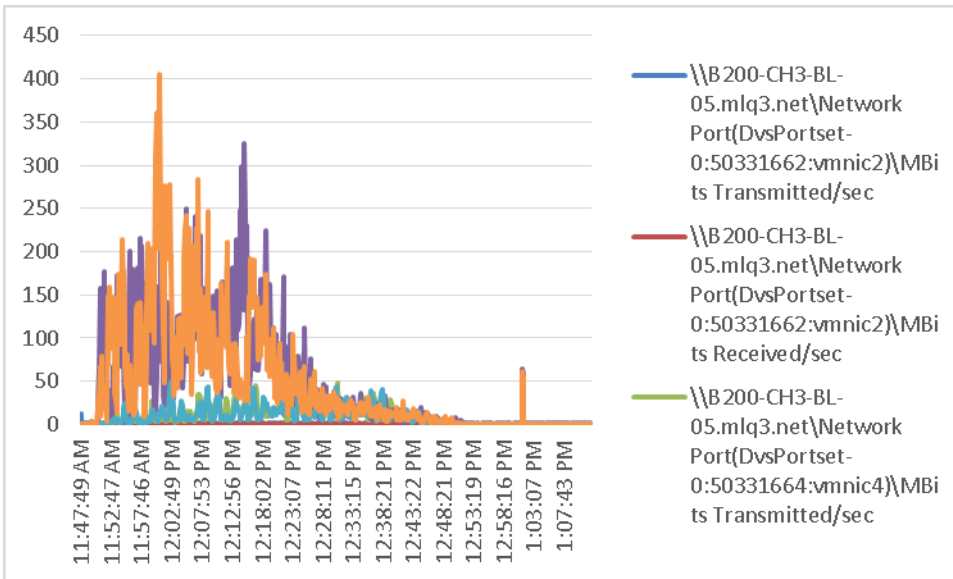












































\\B200-CH6-BL-02\Physical Cpu(_Total)\% Core Util Time















\\B230-CH6-BL-04\Physical Cpu(_Total)\% Core Util Time









\\B230-CH6-BL-06\Physical Cpu(_Total)\% Core Util Time

