

Citrix XenDesktop 5.5 Built on Cisco UCS B-Series Blade Servers, Cisco Nexus 5000 Series Switches, and VMware ESXi 5.0: Reference Architecture and Deployment Guide



Contents

Overview	3
What You Will Learn	3
Audience	3
Objectives	3
Summary of Main Findings	3
Modular Virtual Desktop Infrastructure Overview	4
Modular Architecture	4
Cisco Data Center Infrastructure for Desktop Virtualization	5
Solution Components: Cisco and Citrix Reference Architecture	7
Understanding Desktop User Groups	8
Understanding Applications and Data	8
Project Planning and Solution Sizing Sample Questions	9
Cisco Services	10
The Solution: A Unified, Pretested, and Validated Infrastructure	10
Microsoft Windows 7 SP1 Image Creation and Provisioning	10
Architecture and Design of Citrix XenDesktop 5.5 on Cisco UCS and Storage Solution	25
Design Fundamentals	25
Hosted VDI Design Fundamentals	26
Designing a Citrix XenDesktop 5.5 Deployment	27
Solution Validation	27
Configuration Topology for Scalable Citrix XenDesktop 5.5 VDI on Cisco UCS and Storage Solution	27
Cisco Unified Computing System Configuration	29
Citrix XenDesktop 5.5 and Provisioning Server 5.6 Configuration	42
LAN Configuration	49
SAN Configuration	50
Boot from SAN Configuration	51
Hypervisor Operating System Installation and Configuration	57
Test Setup and Configuration	64
Cisco UCS Test Configuration for Single-Blade Scalability	65
Cisco UCS Configuration for Single-Chassis Test	66
Testing Methodology and Success Criteria	67
Load Generation	67
User Workload Simulation: Login VSI from Login Consultants	67
Success Criteria	69
Test Results	73
Citrix XenDesktop 5.5 Hosted VDI Standard Mode VDI Test Results	73
Eight Cisco UCS B230 M2 Blade Servers: Single-Chassis Validation	79
Scalability Considerations and Guidelines	83
Cisco UCS System Configuration	83
Storage Sizing Best Practices	84
For More Information	84

Overview

What You Will Learn

This deployment guide reports the results of a study evaluating the scalability of a Citrix XenDesktop 5.5 environment using Citrix Provisioning Server 5.6 on the Cisco UCS® B230 M2 Blade Server running VMware ESXi 5 hypervisor software connected to a storage array. The second-generation Cisco Unified Computing System™ (Cisco UCS) hardware and software are used. Best practice recommendations and sizing guidelines are provided for large-scale customer deployments of Citrix XenDesktop 5.5 running two workload scenarios on Cisco UCS.

Audience

This document is designed to assist solution architects, sales engineers, field engineers, and consultants with evaluation, planning, design, and deployment of Citrix XenDesktop 5.5 hosted shared desktop virtualization solutions on Cisco UCS. The reader should have an architectural understanding of Cisco UCS, Cisco Nexus® 5000 Series Switches, Citrix XenDesktop and Provisioning Server software, shared storage, and related software.

Objectives

The main objectives of this guide include articulation of the design considerations and validation efforts required to design and successfully deploy Citrix XenDesktop 5.5 on the reference Cisco UCS architecture with shared storage running on a VMware ESXi 5 hypervisor.

Summary of Main Findings

- The combination of Cisco UCS, Cisco Nexus switching, and storage hardware with VMware ESXi 5, Citrix Provisioning Server 5.6, and Citrix XenDesktop 5.5 software produces a virtual desktop delivery system with a high density per blade and chassis.
- The Cisco UCS B230 M2 half-width blade with dual 10-core processors and 256 GB of memory supports 22.7 percent more virtual desktop workloads than the previously studied full-width blade (Cisco UCS B250 M2 Blade Server) using a medium workload **with** flash video.
- The Cisco UCS B230 M2 half-width blade with dual 10-core processors and 256 GB of memory supports 36.4 percent more virtual desktop workloads than the previously studied full-width blade (Cisco UCS B250 M2) using a medium workload **without** flash video.
- A single Cisco UCS chassis with eight Cisco UCS B230 M2 blades with dual 10-core processors and 256 GB of memory and a Cisco UCS M81KR Virtual Interface Card (VIC) supports 1080 virtual desktop workloads running a medium workload **with** flash video, almost 2.5 times the density of the previously studied chassis with full-width blades (Cisco UCS B250 M2).
- A single Cisco UCS chassis with eight Cisco UCS B230 M2 blades with dual 10-core processors and 256 GB of memory and a Cisco UCS M81KR VIC supports 1200 virtual desktop workloads, more than 2.7 times the density of the previously studied chassis with full-width blades (Cisco UCS B250 M2).
- Compared to the previously studied full-width blades, this solution booted the Microsoft Windows 7 SP1 virtual workloads 50 percent faster without pegging the processor or exhausting memory or storage connections.
- Compared to the previously studied full-width blades, this solution ramped up to a steady state 50 percent faster without pegging the processor or exhausting memory or storage subsystems.

- Compared to the previously studied full-width blades, with this solution the rack space required to support 1000 users decreased from 18 rack units (18RU) to 6RU, for 2000 users it decreased from 30RU to 12RU, and for 5000 users it decreased from 72RU to 30RU.
- This solution offers a validated design that is 100 percent virtualized on VMware ESXi 5. All the Microsoft Windows 7 SP1 virtual desktops and supporting infrastructure components, including Microsoft Active Directory, profile servers, provisioning servers, Microsoft SQL Server, and Citrix XenDesktop delivery controllers, were hosted on VMware vSphere 5.
- Cisco maintains industry leadership with the new Cisco UCS Manager 2.0 software, which makes scaling easy, helps ensure consistency, and simplifies maintenance.
- The Cisco[®] 10-Gbps unified fabric is additionally validated on second-generation Cisco UCS 6200 Series Fabric Interconnects and second-generation Cisco Nexus 5500 platform access switches with tests running more challenging workloads, maintaining exceptional user response times.
- The storage system provides storage consolidation and efficiency. Both block and file storage resources are provided by a single system.
- Citrix HDX technology, extended in Citrix XenDesktop 5.5 software, provides excellent performance with host-rendered Adobe Flash video and other demanding applications.

Modular Virtual Desktop Infrastructure Overview

The main challenges businesses face when considering virtual desktop infrastructure (VDI) include understanding the business factors that influence a move to VDI, understanding VDI desktop user candidate groups, understanding the virtual desktop workloads used by each group, and perhaps most important, understanding where the data resides for a user group and workload pair.

Understanding all these factors allows the enterprise to analyze its VDI use cases, select those that have the best potential for good return on investment (ROI), and plan for successful VDI deployment.

The analysis must identify the acceptable end-user experience that will define a successful project.

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 centers, knowledge and task workers, partners, consultants, and executives connecting from locations around the world 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 help ensure protection of corporate data and to prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 1). These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 7.

Figure 1. The Evolving Workplace Landscape

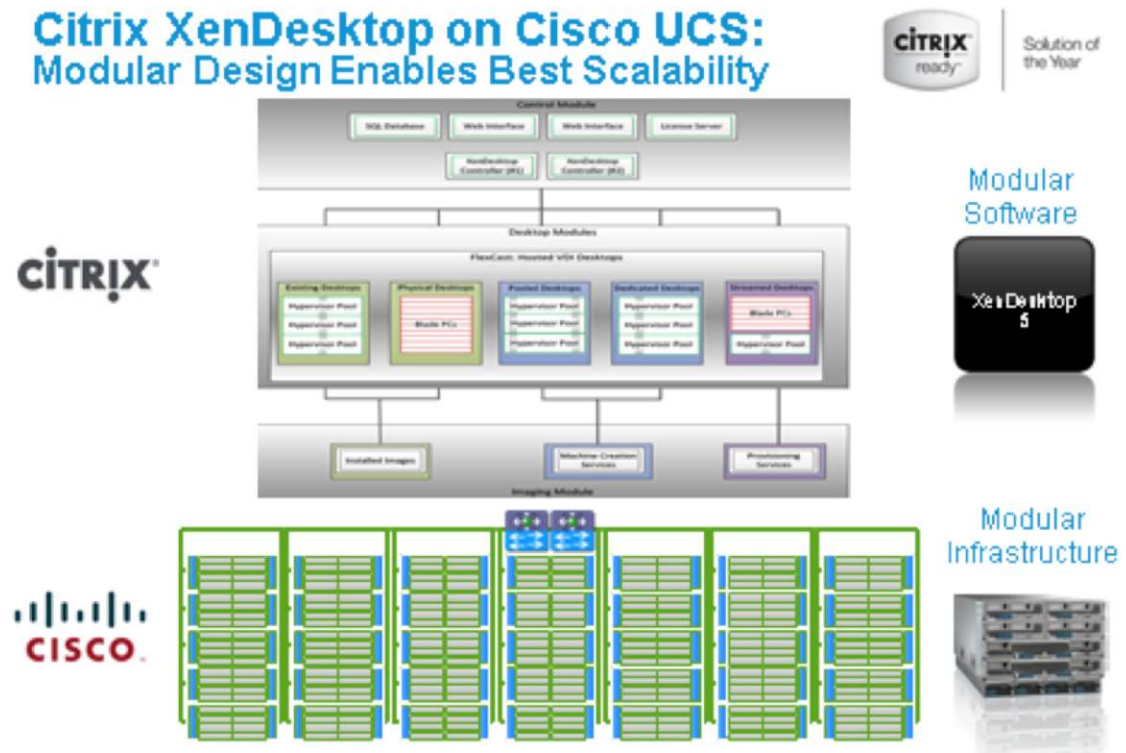


Some of the main reasons for the move to desktop virtualization are increased data security and reduced total cost of ownership (TCO) through increased control and reduced management costs.

Cisco Data Center Infrastructure for Desktop Virtualization

Cisco focuses on three main elements to deliver the best desktop virtualization data center infrastructure: simplification, security, and scalability. The solution software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform (Figure 2).

Figure 2. Citrix XenDesktop on Cisco UCS



Simplification

Cisco UCS provides a radical new approach to industry-standard computing and provides the core of the data center infrastructure for desktop virtualization and Cisco Virtualization Experience Infrastructure (VXI™). Among the many features and benefits of Cisco UCS are the drastic reductions in the number of servers needed and the number of cables per server, and the capability to quickly deploy or reprovision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and virtual desktop provisioning, operations are significantly simplified. Thousands of desktops can be provisioned in minutes with Cisco service profiles and Cisco storage partners' storage-based cloning. This accelerates the 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 reduction is possible because of the industry-leading, highest virtual desktop density per blade of Cisco UCS along with the reduced cabling and port count resulting from the unified fabric and unified ports of Cisco UCS and desktop virtualization data center infrastructure.

Simplification also leads to the improved and more rapid success of a desktop virtualization implementation. Cisco and its partners have developed integrated, validated architectures, including predefined and validated infrastructure packages.

Security

Although virtual desktops are inherently more secure than their physical counterparts, 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 data center infrastructure (Cisco UCS and Cisco Nexus platforms) 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.

Scalability

Growth of a desktop virtualization solution is almost inevitable, and it is critical to have a solution that can scale, and scale predictably, with that growth. The Cisco solution supports more virtual desktops per server, and additional servers scale with nearly linear performance. Cisco data center infrastructure provides a flexible platform for growth and improves business agility. Cisco UCS service profiles allow on-demand desktop provisioning, and it is as easy to deploy dozens or to deploy thousands of additional desktops.

Each additional Cisco UCS server provides nearly linear performance and uses 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 Cisco Nexus data center infrastructure is an excellent platform for growth, with transparent scaling of server, network, and storage resources to support desktop virtualization.

Savings and Success

As demonstrated in the preceding discussions, the simplified, secure, scalable Cisco data center infrastructure solution for desktop virtualization saves time and costs. It offers faster payback and ongoing savings (better ROI and lower TCO) with the industry's highest virtual desktop density per server, so fewer servers are needed, reducing both capital expenditures (CapEx) and operating expenses (OpEx). It also greatly reduces network infrastructure costs, with fewer cables per server and fewer ports required, through the use of the Cisco UCS architecture and unified fabric.

The simplified deployment of Cisco UCS for desktop virtualization accelerates up time to productivity and enhances business agility. IT staff and end users are more productive more quickly and the business can respond to new opportunities by simply deploying virtual desktops whenever and wherever needed. The high-performance Cisco systems and network deliver a near-native end-user experience, allowing users to be productive any time and anywhere.

Solution Components: Cisco and Citrix Reference Architecture

Cisco's desktop virtualization solution is the first in the industry to bind together the three critical elements of an end-to-end deployment: the end user, the network, and the data center. It builds on Cisco's architectural advantage to provide a solution that supports a diversity of endpoint devices, extends pervasive security and policy management to each virtual desktop, and uses a new and innovative virtualization-optimized stateless server computing model (Cisco UCS).

Base Components

- Cisco UCS computing platform includes:
 - Cisco UCS 6200 Series Fabric Interconnects
 - Cisco UCS 2200 Series Fabric Extenders
 - Cisco UCS 5108 Blade Server Chassis
 - Cisco UCS B230 M2 Blade Servers for virtual desktop hosting
 - Cisco UCS B200 M2 Blade Servers for infrastructure
- Cisco Nexus 5500 platform switches
- Hypervisor: VMware ESXi 5
- Virtual desktop connection broker: Citrix XenDesktop 5.5 with Provisioning Server 5.6 SP1

Understanding Desktop User Groups

The enterprise must take considerable care to identify desktop user groups and their memberships. The most broadly recognized high-level user groups are:

- **Task workers:** These groups of users work in highly specialized environments in which the number of tasks performed by each worker is essentially identical. These users typically are located at a corporate facility (for example, call center employees).
- **Knowledge workers and office workers:** These groups of users use a relatively diverse set of applications that are web based and installed and that access data. These workers typically have several applications running simultaneously throughout the workday and require the use of Adobe Flash video for business purposes. These workers do not exist in a single group within an organization. These workers are typically located at a corporate office (for example, workers in the accounting department).
- **Power users:** These groups of users run high-end, memory-, processor-, disk I/O-, and graphics-intensive applications, often simultaneously. These users have high requirements for reliability, speed, and real-time data access (for example, a design engineer).
- **Mobile workers:** These groups of users may share common traits with knowledge and office workers, with the added complexity of needing to access applications and data from wherever they are—at a remote corporate facility, at a customer location, at an airport, at a coffee shop, or at home—all in the same day (for example, a company's outbound sales force).
- **Remote workers:** These groups of users could be part of the task worker or knowledge and office worker groups, but their experience is from a remote site that is not corporate owned and most often is the user's home. This scenario introduces several challenges regarding the type, available bandwidth, and latency and reliability of the user's connectivity to the data center (for example, a work-from-home accounts payable representative).
- **Guest and contract workers:** These groups of users 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 remotely (for example, a medical data transcriptionist).

Organizations should search for and identify any subgroups of the major groups listed here. Typically, each subgroup has different application and data requirements.

Understanding Applications and Data

After the desktop user groups and subgroups have been identified, the next task is to catalog group application and data requirements. This task can be one of the most time-consuming processes in the VDI planning exercise, but it 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 of cloud applications, such as Salesforce.com.

This application and data analysis is beyond the scope of this Cisco Validated Design document, but it should not be omitted from the planning process. A variety of third-party tools are available to assist your organization with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

Now that you understand user groups and their applications and data requirements, you should consider some important project and solution sizing questions.

You should address these general project questions at the outset:

- Has a VDI pilot plan been created based on the business analysis of the desktop groups, applications, and data?
- Is the infrastructure and budget in place to run the pilot program?
- Are the required skill sets to run the VDI project available? Can you hire or contract for them?
- Do you have end-user experience performance metrics identified for each desktop subgroup?
- How will you measure success or failure?
- What are the future implications of success or failure?

The following is a nonexhaustive list of sizing questions that you should address for each user subgroup:

- What is the desktop OS planned? Microsoft Windows 7 or Windows XP?
- Will you use a 32-bit or 64-bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? Will all use Microsoft Windows 7?
- How much memory per target desktop group desktop will you need?
- Do you have any multimedia-, Adobe Flash-, or graphics-intensive workloads?
- What is the endpoint graphics processing capability you need?
- Will Citrix XenApp be used for hosted shared server desktops or exclusively for Citrix XenDesktop?
- Are any Citrix XenApp hosted applications planned? Are they packaged or installed?
- Will Citrix 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 sufficient I/O operations per second (IOPS) available for the write-intensive VDI workload?
- Will you have storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?

- Is antivirus software a part of the image?
- Is user profile management (nonroaming profile based) part of the solution?
- What are the fault-tolerance, failover, and disaster-recovery plans?
- Do you have any additional desktop subgroup-specific questions?

Cisco Services

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

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

The Solution: A Unified, Pretested, and Validated Infrastructure

To meet the challenges of designing and implementing a modular desktop infrastructure, Cisco, Citrix, and VMware have collaborated to create the data center solution described in this document.

Main elements of the solution include:

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

Microsoft Windows 7 SP1 Image Creation and Provisioning

The Microsoft Windows 7 SP1 master, or golden, image with additional software was initially installed and prepared as a standard virtual machine on VMware vSphere ESXi 5 prior to being converted into a separate Citrix Provisioning Server vDisk file. The vDisk file is used in conjunction with Citrix Provisioning Server and the Citrix XenDesktop 5.5 controller to create more than 1000 new desktop virtual machines on the VMware ESXi 5 hosts.

With Citrix XenDesktop 5.5 and Provisioning Server 5.6 SP1, the Citrix XenDesktop Setup Wizard integration is not operational. The Citrix Provisioning Server's Streamed Virtual Machine (VM) Setup Wizard was used to create the Microsoft Windows 7 virtual machines on VMware ESXi 5. The Citrix XenDesktop 5.5 Desktop Studio was then used to create streamed machine catalogs and import the virtual machines created by Citrix Provisioning Server.

The Citrix Provisioning Server effectively creates virtual machine objects; configures processors, memory, network adapters, and VLAN assignments; and creates and assigns a 3-GB virtual write-cache virtual disk hosted on a data store mounted on the hypervisor through the Network File System (NFS) from a storage solution. It also creates and configures the relevant Citrix Provisioning Server and Active Directory objects associated with the new machines.

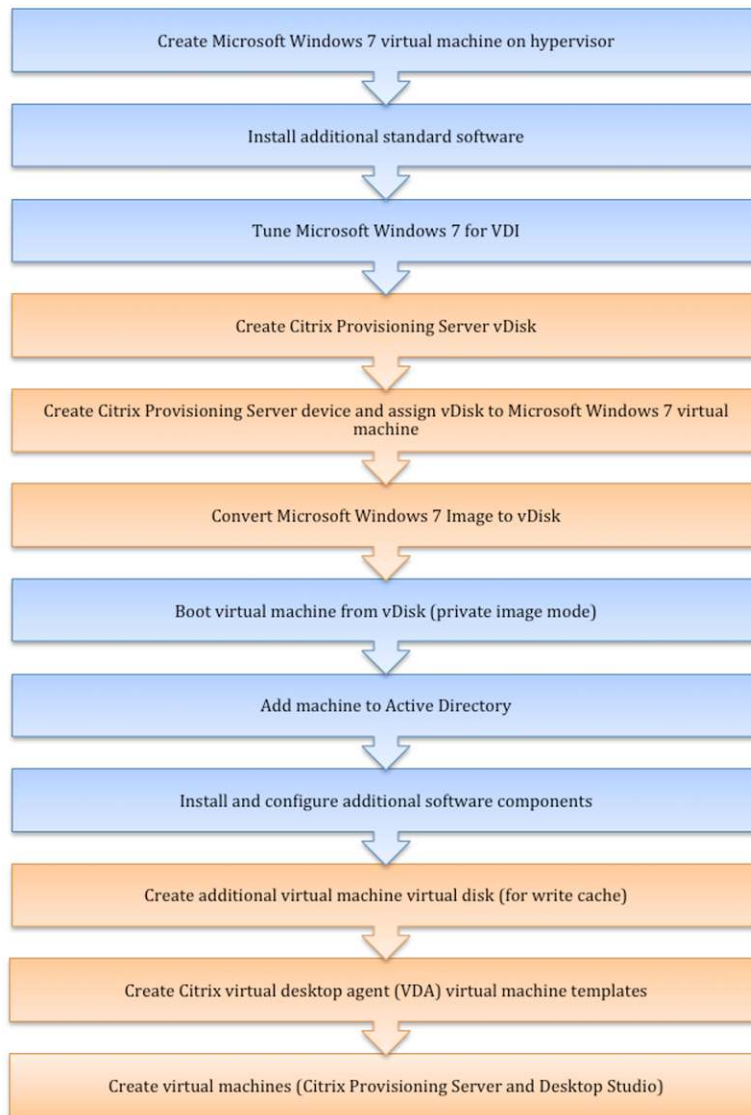
More information about why the additional virtual disks are needed can be found in the section [“Configuration Topology for Scalable Citrix XenDesktop 5.5 VDI on Cisco UCS and Storage Solution.”](#)

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

Create Microsoft Windows 7 SP1 Virtual Machine and Install Standard Software

Figure 3 provides a process overview for Microsoft Windows 7 image and vDisk provisioning.

Figure 3. Microsoft Windows 7 Image and vDisk Citrix Provisioning Process Overview



Create Microsoft Windows 7 Virtual Machine and Install Standard Software

The virtual machine configurations and software in Table 1 were used to create the initial Microsoft Windows 7 virtual machine on the hypervisor, which is then extracted to create a Citrix Provisioning Server vDisk in .vhd format.

Table 1. Citrix XenDesktop Virtual Desktop Image

OS	Microsoft Windows 7 Enterprise 32-bit	Service Pack	SP1
CPU	1 x virtual CPU (vCPU)	RAM	1536 MB
Disk		Network	1 x 10 Gigabit Ethernet VMXNET3
• C:\	1 x 16-GB (Citrix Provisioning Server vDisk)		
• E:\	1 x 3-GB virtual disk (Citrix Provisioning Server write cache)		

- Software installed prior to cloning to vDisk
 - VMware tools
 - Citrix Provisioning Server 5.6.1 target device
 - Microsoft Office 2007 SP2 Enterprise Edition
 - Internet Explorer 8.0.7601.17514

Tune Microsoft Windows 7 Image for VDI

When many Microsoft Windows desktops run on a hypervisor, you need to try to reduce unnecessary CPU cycles and disk I/O to improve system performance and stability. Turning off unnecessary processes and other unwanted desktop services helps achieve this.

The following tuning and configuration settings were applied to the standard image:

- Configure fixed 1.5-GB page file.
- Configure networking and firewall:
 - Turn off firewall.
 - Set the Domain Name Server (DNS) IP addresses for the domain.
 - Turn off IPV6.
- Apply the Microsoft Windows 7 optimization recommendations from the following Citrix article: Microsoft Windows 7 Optimization Guide for Desktop Virtualization (<http://support.citrix.com/article/CTX127050>).
- Apply the best practices for configuring the Citrix Provisioning Server on a network from the following Citrix article: Best Practices for Configuring Citrix Provisioning Server on a Network (<http://support.citrix.com/article/CTX117374>).
- Apply the fix for the VMXNET3 driver for virtual machines: Hotfix CPVS56SP1E011 for Citrix Provisioning Services 5.1, 5.1 SP1, 5.1 SP2, 5.6, and 5.6 SP1 (<http://support.citrix.com/article/CTX128160>).
- Using the Microsoft Management Console on the master (or golden) image, choose Local Computer Policy > User Configuration > Administrative Templates > System > Scripts and enable "Run logon scripts synchronously."
- Using the Microsoft Management Console on the master (or golden) image, choose Local Computer Policy > User Configuration > Administrative Templates > Start Menu and Taskbar and enable "Turn off user tracking."
- Apply the Citrix Provisioning Services 5.6 SP1 Hotfix 18672 target shutdown fix: Hotfix CPVS56SP1E033 for Citrix Provisioning Services 5.6 Service Pack 1 (<http://support.citrix.com/article/CTX130661>).
- Apply the Citrix Provisioning Services 5.6 SP1 and SP2 BNlstack update: Hotfix CPVS56SP2E002 for Citrix Provisioning Services 5.6 Service Pack 1 and Service Pack 2 (<http://support.citrix.com/article/CTX131544>).


Citrix Provisioning Services vDisk and Virtual Machine Creation

After the Microsoft Windows 7 master (or golden) image virtual machine has initially been created with the required software, it must be extracted into a Citrix Provisioning Server vDisk file. To do this, you use the Citrix Provisioning Services Imaging Wizard, which is part of the Citrix Provisioning Services target device installation.

Prior to installing the Citrix Provisioning Services target device software on the master (or golden) image virtual machine, perform the following steps:

1. Create a Citrix Provisioning Server vDisk:

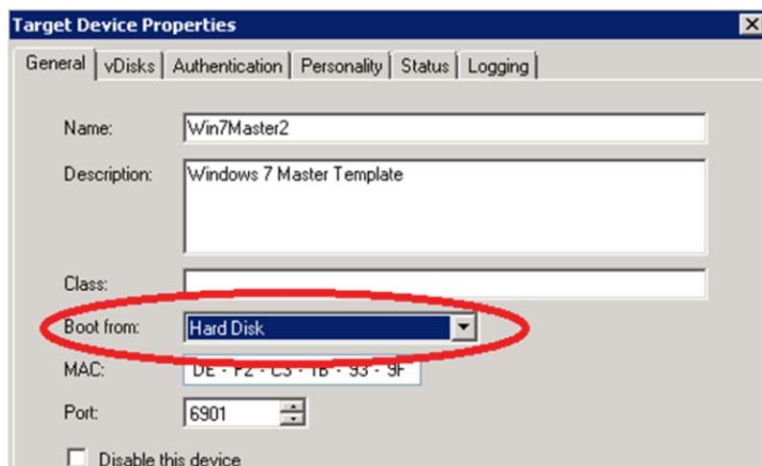
- a. Open the Citrix Provisioning Services Console on a Citrix Provisioning Server (Release 5.6 SP1 or later).
- b. Create a new, fixed-size vDisk (16 GB, but size may vary depending on requirements).
- c. Configure the Private Image Access Mode option on the new vDisk.



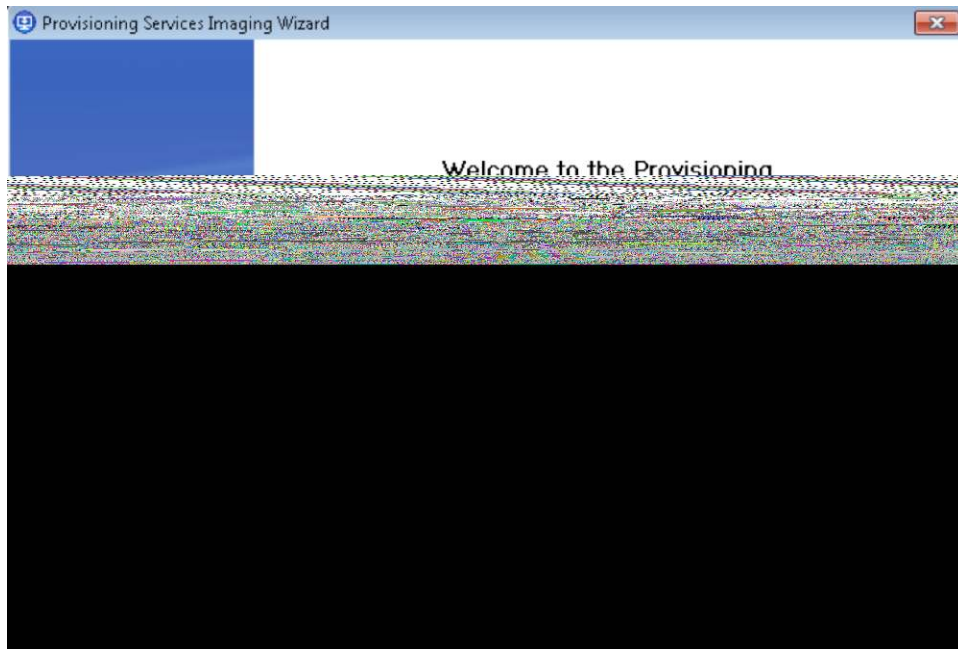
- d. Configure the “Active Directory machine password management” option.



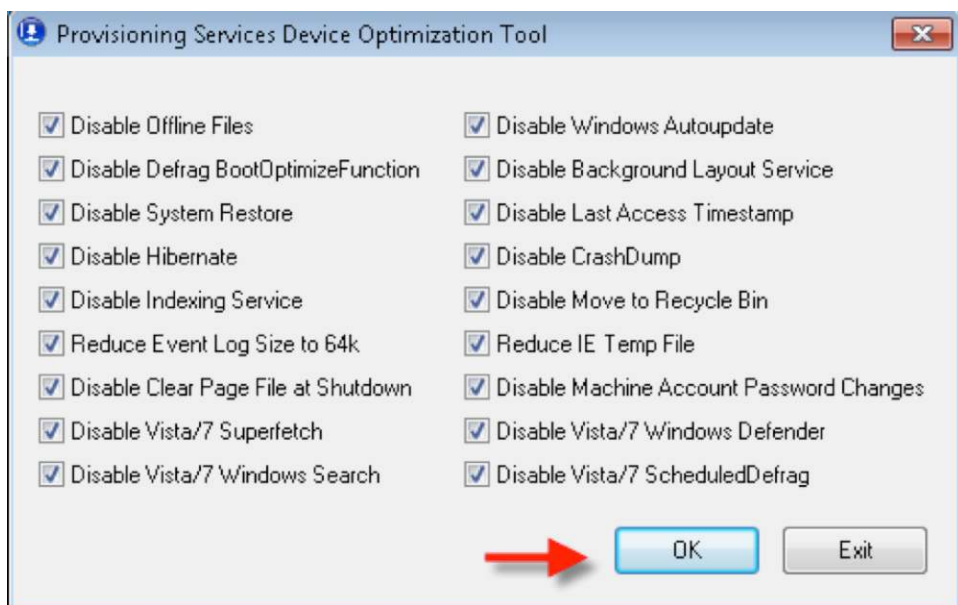
2. Create a new device in a new or existing Citrix Provisioning Services collection and assign the vDisk.
 - a. Assign the MAC address of the master (or golden) image virtual machine to this Citrix Provisioning Services object.
 - b. Set the device to boot from the hard disk.



- c. Assign the vDisk created in step 1 on the vDisks tab.
3. Boot the Microsoft Windows 7 golden image virtual machine from its VMware ESXi 5 host and verify that the vDisk is attached.
4. Clone the Microsoft Windows 7 Image to the vDisk:
 - a. Run the Citrix Provisioning Services Imaging Wizard.

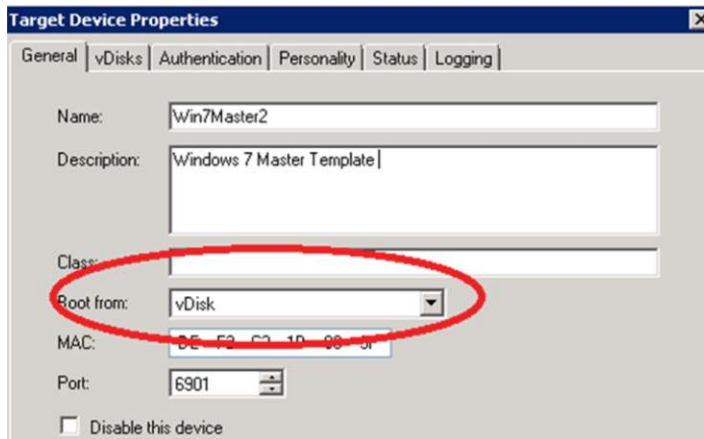


- b. Run the Citrix Provisioning Services Device Optimization Tool by clicking OK.



The image is to assigned the vDisk (E:\).


5. After the Imaging process is complete, shut down the virtual machine.
6. Set the virtual machine to boot from the Citrix Provisioning Server vDisk (rather than the hard disk) and start the virtual machine:
 - a. Using the Citrix Provisioning Services Console, change the target device options to "Boot from vDisk."



7. Using VMware vCenter, start the virtual machine.
 - a. Add the host to the domain.
 - b. Restart the guest OS.
8. Install and configure additional software components. Software Installed after vDisk cloning:
 - a. Citrix XenDesktop VDA 4.0.5010
 - b. Login VSI 3.0 Target Software (tools used for benchmarking)
9. Add a 3-GB Writecache.VMDK/.VHD file to the master (or golden) image virtual machine (so the disk signature matches).

You need to create and format an additional virtual disk on the Microsoft Windows 7 master (or golden) image virtual machine. This virtual disk will later be detached and used as the default virtual machine template for the cloning process, so that each provisioned virtual machine will have a unique 3-GB virtual disk (E:\ drive). This drive is where the provisioned virtual machine's Citrix Provisioning Server write cache will be placed and where all write I/O processing will subsequently be performed.

 - a. Create a new 3-GB hard disk using VMware vSphere standard procedures for the Microsoft Windows 7 SP1 master (or golden) image virtual machine.
 - b. Activate the new 3-GB hard disk on the Microsoft Windows 7 master (or golden) image virtual machine using the disk manager. (Bring the disk online and use the Standard mode. Do not use the Dynamic mode.)
 - c. Format the new volume as NTFS.
 - d. Shut down the virtual machine.
10. Using VMware vCenter, detach the new 3-GB virtual disk from the master (or golden) image virtual machine, but **do not** delete it (note where it is stored for the next stage).
11. On the Citrix Provisioning Services Console, change the vDisk mode to "Standard Image" and change the cache location to "Cache on device's HD."



12. Create a VDA virtual machine template (for provisioning).

Create the virtual machine templates on the relevant NFS data stores hosted on the storage solution. If you are creating large numbers of clones, you should mount several NFS volumes on the hypervisors balanced between at least two storage controllers.

- a. After the NFS volumes have been mounted on the hypervisors, use the VMware vSphere client to create a Microsoft Windows virtual machine in the usual way, but do not start it:
 - i. Create a new virtual machine (Win7_PVS_Temp) with a standard hard disk for Microsoft Windows 7.
 - ii. Allocate 1.5 GB of RAM.
 - iii. Assign the virtual machine to correct virtual machine network.
 - iv. Change boot order to use network boot.
 - v. Delete the assigned standard hard disk for the Microsoft Windows 7 virtual disk.
 - vi. Attach the virtual disk created in step 10.
- b. Convert the virtual machine to a template using VMware vCenter.

13. Create a copy of the template for each NFS volume.

- a. Create and apply a full copy of the template to NFS volume 1 and name it (for example, Win7PVSTemp (1)).
- b. Create and apply a full copy of the template to NFS volume 2 and name it (for example, Win7PVSTemp (2)).
- c. Continue creating templates until you have applied a template to each target NFS volume that you want to use.

14. Provision VDI Microsoft Windows 7 virtual machines.

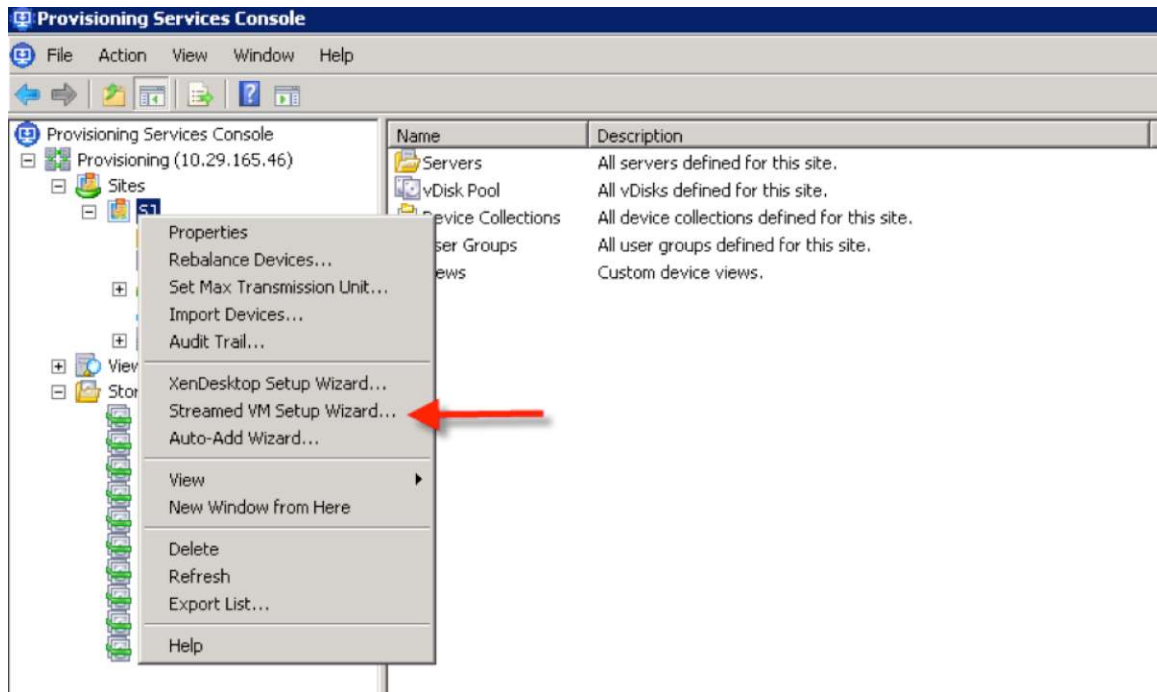
Large-scale provisioning can be achieved easily by using the Citrix XenDesktop Setup Wizard, which is invoked from the Citrix Provisioning Server. Alternatively, you can use the Streamed VM Setup Wizard on Citrix Provisioning Server 5.6 SP1 to provision the virtual machines in conjunction with the Citrix XenDesktop 5.5 Desktop Studio to import the streamed machines. This is the approach used here.

Note: The entire Citrix XenDesktop infrastructure should be set up and tested prior to creating clones registered or configured on each of the components, including Active Directory, with this tool.

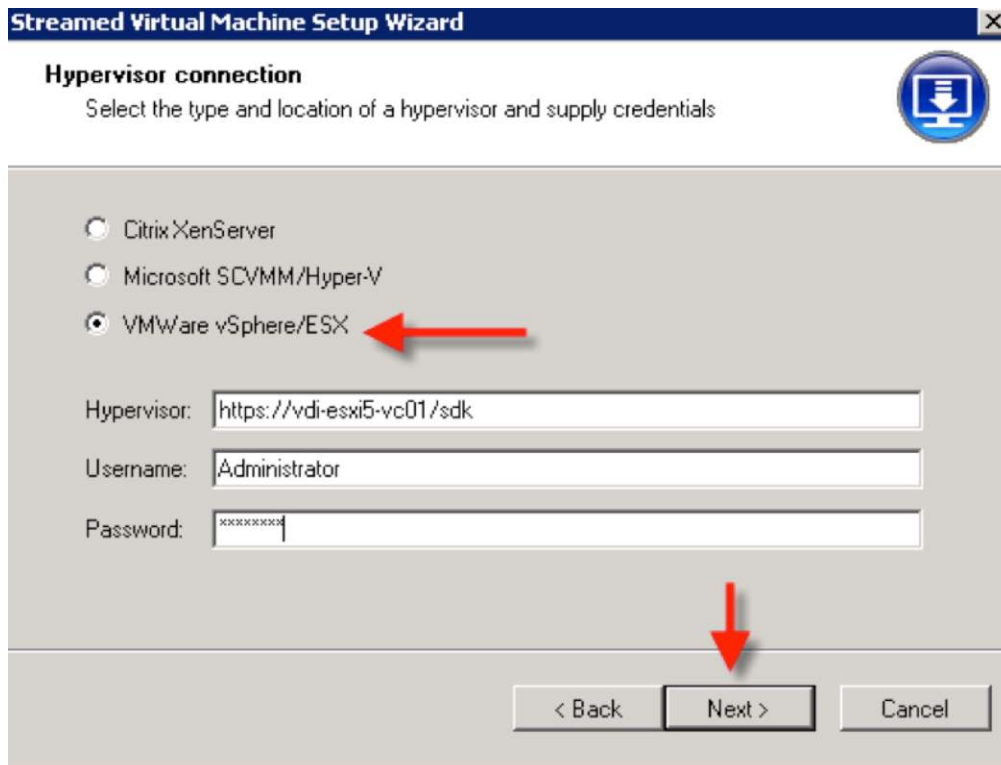
The aim of this step is to create VDI virtual machines evenly distributed across all available mounted NFS data stores, so calculate how many virtual machines you will be creating on each data store and then run the Citrix XenDesktop Setup Wizard or the Streamed VM Setup Wizard.

The Citrix XenDesktop Setup Wizard or the Streamed VM Setup Wizard should be installed and run on the Citrix Provisioning Server from the Citrix Provisioning Services Console.

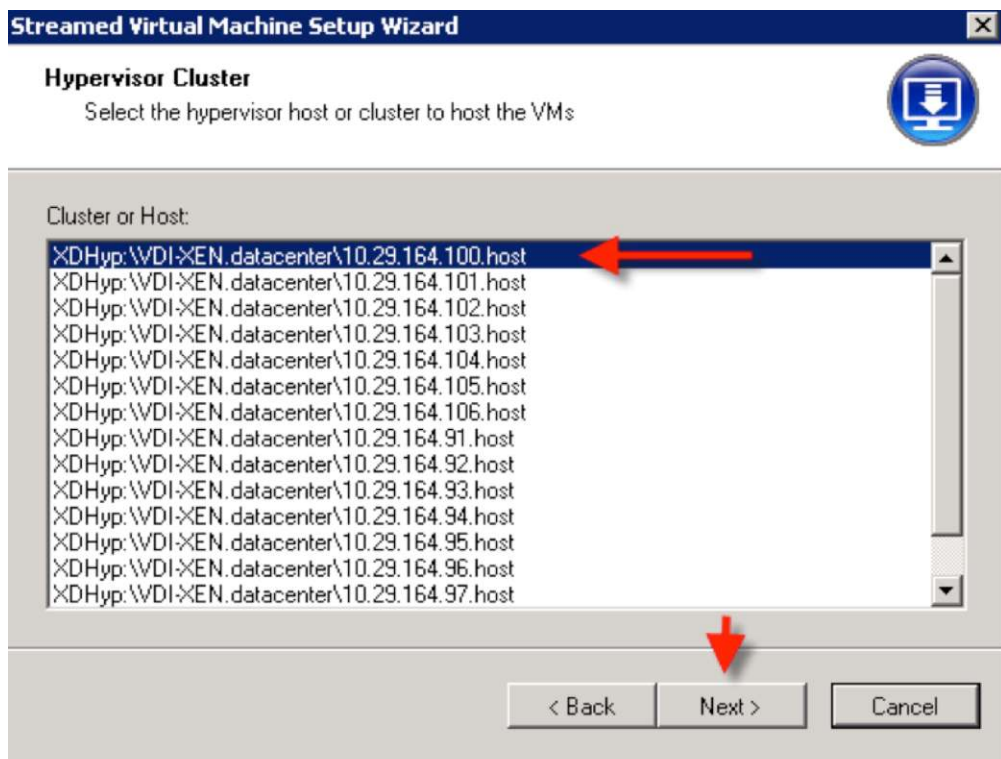
- a. Choose the site and choose “Streamed VM Setup Wizard.”



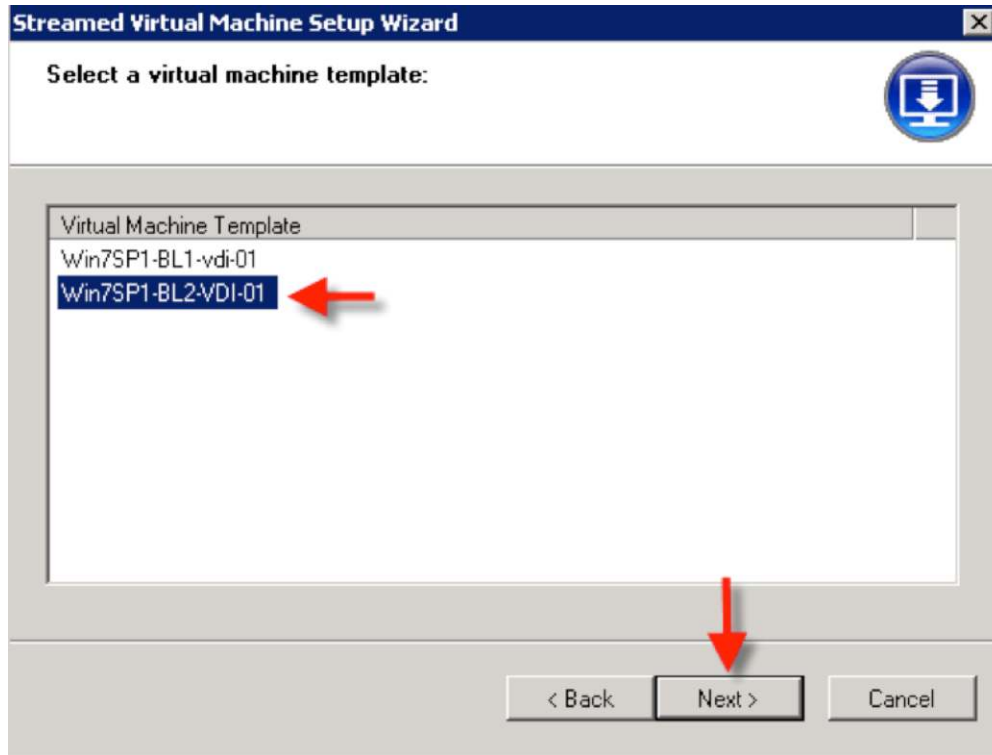
- b. Choose “VMware vSphere/ESX,” provide the hypervisor location and credentials, and click Next.



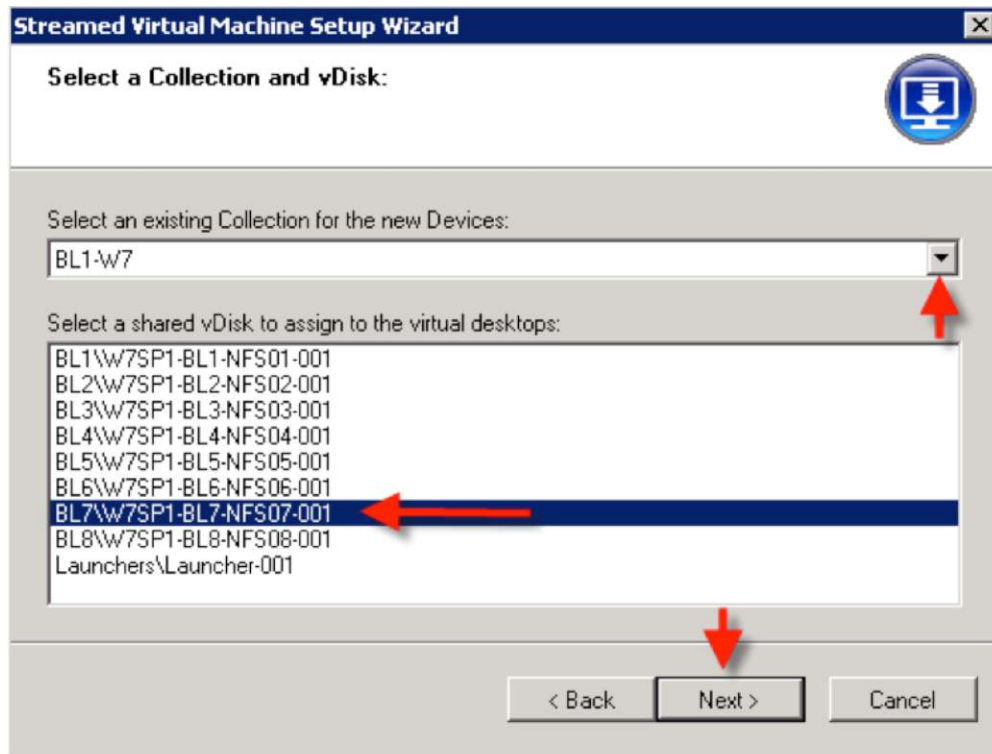
- c. Select the host or cluster to which the virtual machines will be provisioned; then click Next.



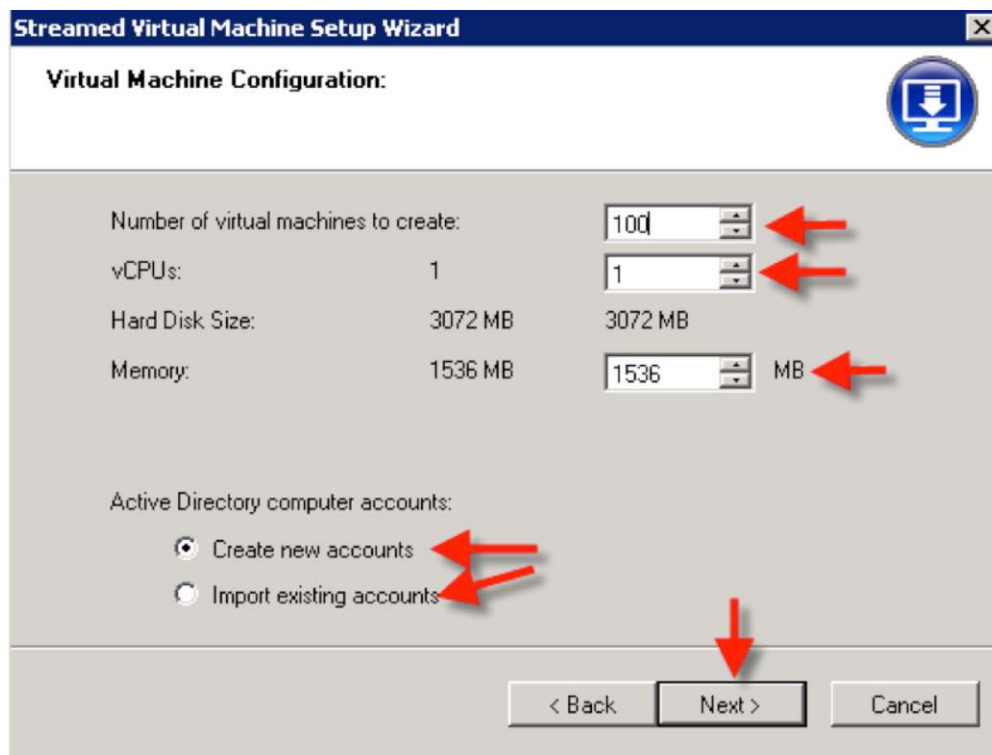
- d. Select the template associated with the host or cluster and NFS volume to which you want to add virtual machine instances; then Click Next.



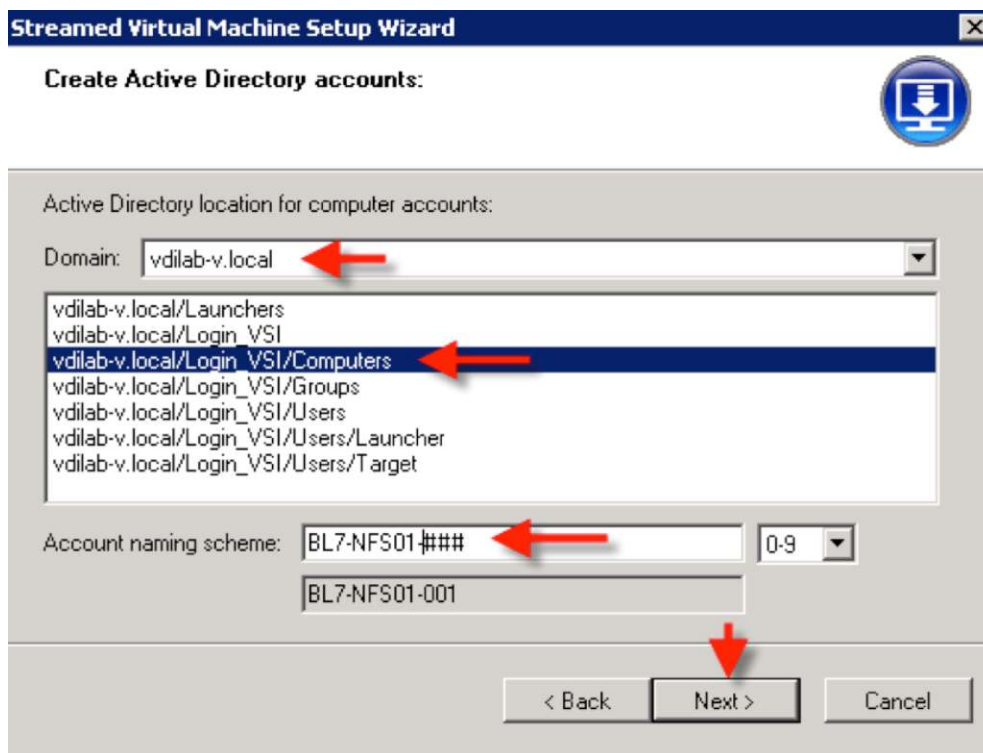
- e. Select an existing Citrix Provisioning Services device collection and the vDisk to assign to the virtual desktops; then click Next.



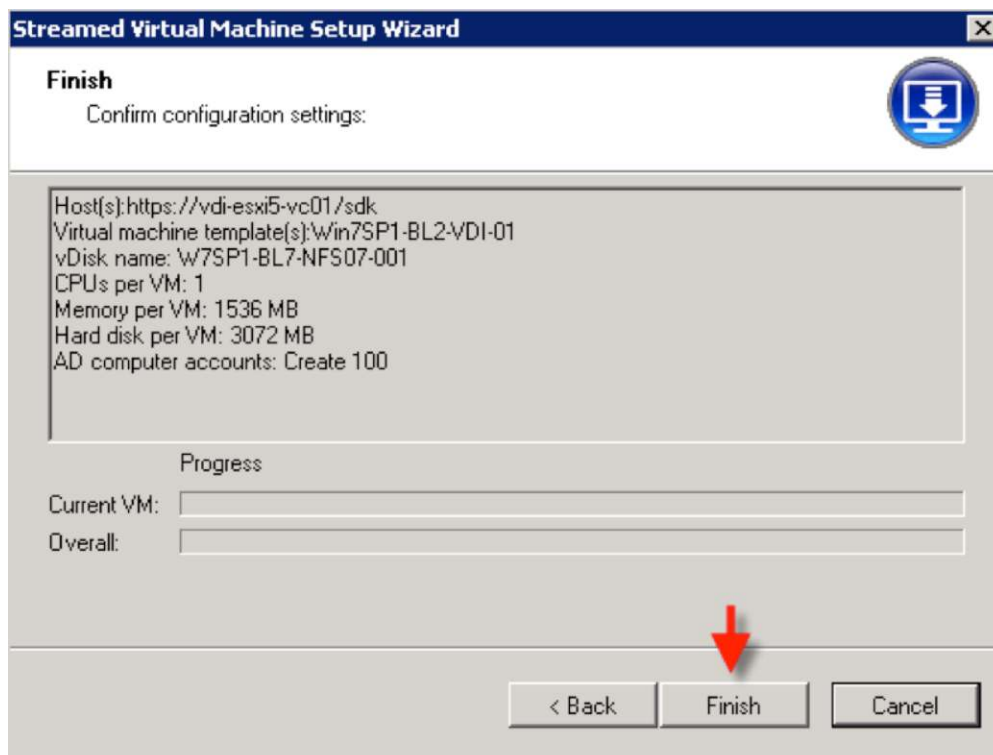
- f. Complete the Virtual Machine Configuration dialog box; then click Next.



- g. Select the desired organization unit in which the virtual machines will be created in Active Directory and create the naming scheme; then click Next.



- h. Click Finish to start the provisioning process.



- i. After the process is complete, run the Streamed VM Setup Wizard again using the same process except this time select a different template with a different virtual desktop numbering scheme from the next available host and NFS volume until you have virtual machines for all NFS volumes.

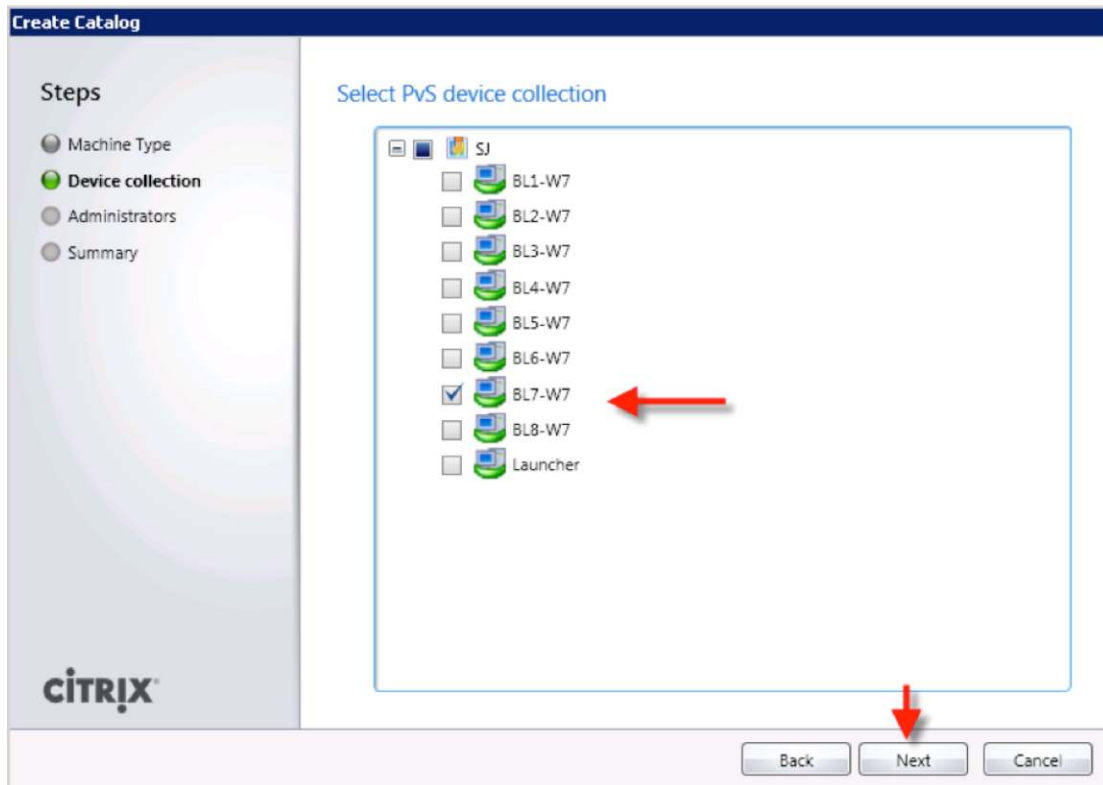
Integrate Citrix XenDesktop 5.5

After all of the Microsoft Windows 7 virtual machines are created with the Streamed VM Setup Wizard, use Citrix XenDesktop 5.5 Desktop Studio to create machine catalogs and add the provisioned machines.

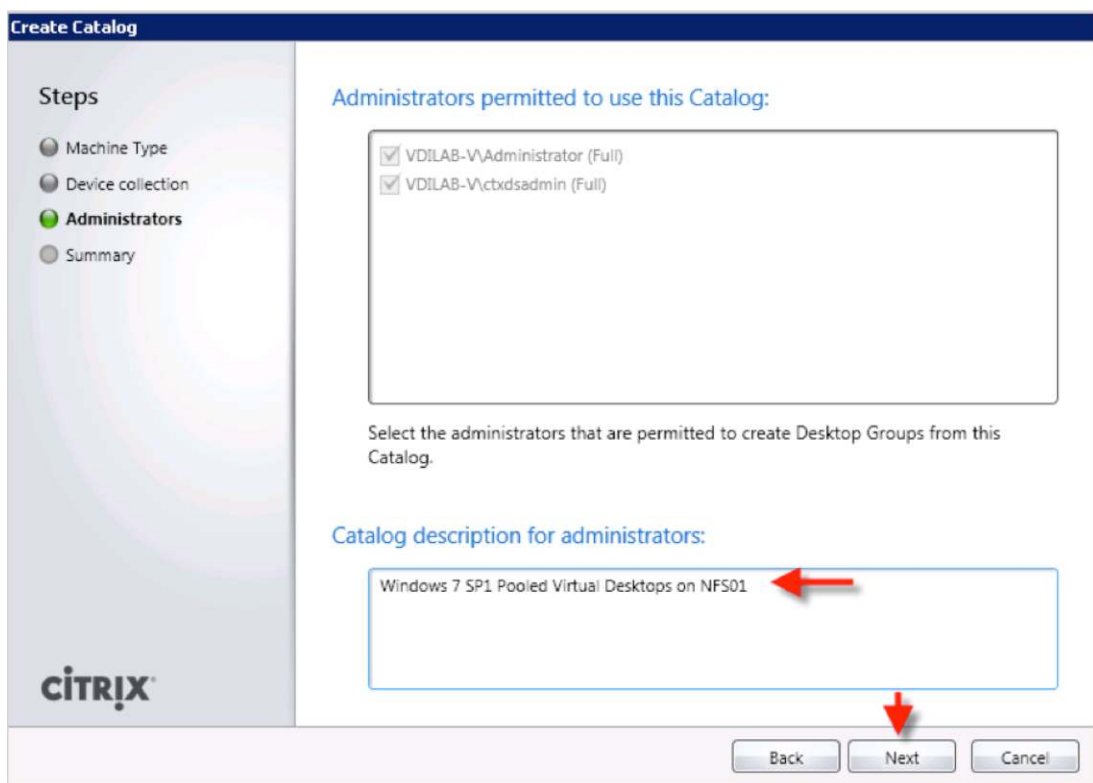
1. Open the Citrix Desktop Studio on a Citrix XenDesktop 5.5 server, select the Machine node in the left pane, and click Create Catalog in the right Action pane. Choose Streamed as the machine type, provide the IP address of the Citrix Provisioning Server 5.6 SP1 hosting the device collection and vDisk, select the device collection domain, and select "Virtual" as the target device type; then click Next.

The screenshot shows the 'Create Catalog' wizard in Citrix Desktop Studio. The left pane shows the 'Steps' section with 'Machine Type' selected. The main area displays the configuration for the 'Streamed' machine type. The 'Machine type' dropdown is set to 'Streamed'. The 'Provisioning Services address' text box contains '10.29.165.46'. The 'Device collection domain' dropdown is set to 'vdi1ab-v.local'. The 'Target device type' radio buttons have 'Virtual' selected. At the bottom, the 'Next' button is highlighted. Red arrows point to the 'Machine type' dropdown, the 'Provisioning Services address' text box, the 'Device collection domain' dropdown, the 'Virtual' radio button, and the 'Next' button.

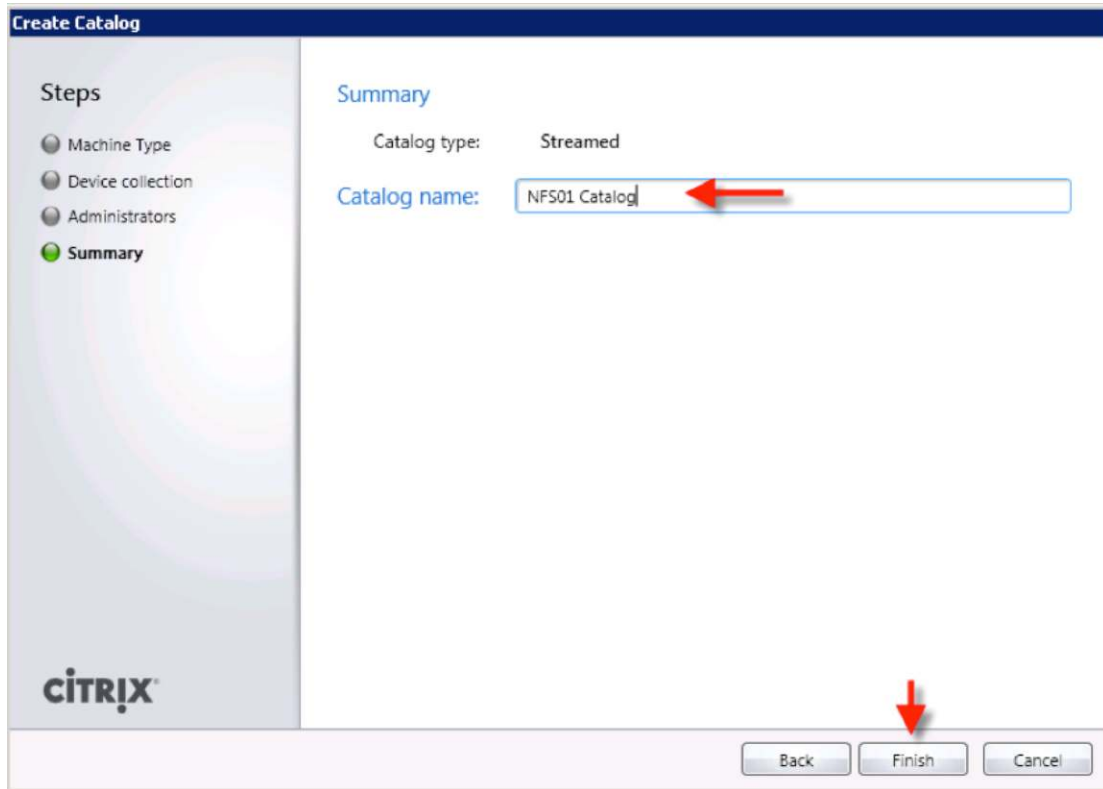
2. Expand the site in the "Select PvS device collection" list and choose one or more Citrix Provisioning Services device collections to be used for the catalog; then click Next.



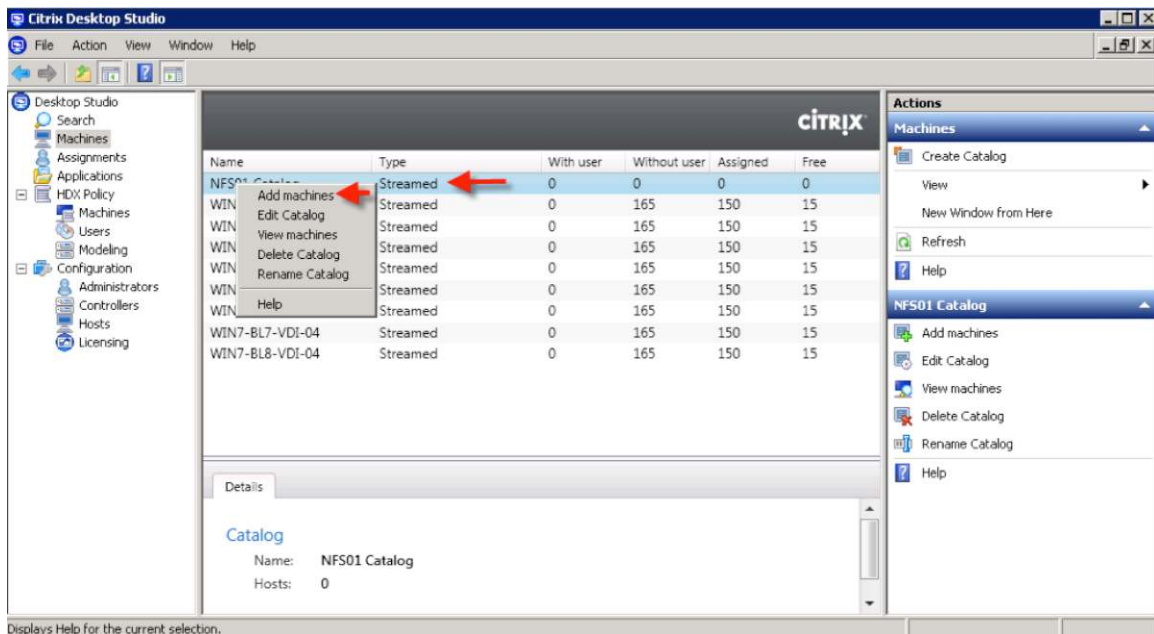
3. Add a description for the catalog; then click Next.



4. Enter a catalog name; then click Finish.



5. After the catalog is created, right-click it and choose "Add machines."



6. The dialog box appears as in step 2. Select the Citrix Provisioning Services device collections containing the virtual machines to be added to the catalog you created; then click Finish.

Architecture and Design of Citrix XenDesktop 5.5 on Cisco UCS and Storage Solution

Design Fundamentals

There are many reasons to consider a virtual desktop solution: for instance, a growing and diverse base of user devices, complexity in management of traditional desktops, security, and bring your own device (BYOD) computing in the work environment. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully perform the users' roles. The following user classifications is provided:

- **Knowledge workers** today do not just work in their offices all day; they attend meetings, visit branch offices, work from home, and even work from coffee shops. These “anywhere” workers expect access to all their applications and data wherever they are located.
- **External contractors** are increasingly part of everyday business. They need access to all your applications and data, yet administrators still have little control over the devices they use and the locations from which they work from. Consequently, IT has to make trade-offs, weighing the cost of providing these workers with a device compared to 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 desktops from everywhere, regardless of their capability to connect to a network. In addition, these workers expect the capability 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 reprovision desktops with the latest operating systems and applications as the needs of the organization change.

After you have identified the user classifications and defined the business requirements for each user classification, you need to evaluate the types of virtual desktops that are available based on user requirements. There are essentially five potential desktop environments for each user:

- **Traditional PC:** A traditional PC is what typically constituted a desktop environment: a physical device with a locally installed operating system.
- **Hosted, server-based desktop:** A hosted, server-based desktop is a desktop with which 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 may affect the other users.
- **Hosted virtual desktop:** A hosted virtual desktop is a virtual desktop running either on a virtualization layer (Citrix XenServer, Microsoft Hyper-V, or VMware ESXi) 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 desktop:** A streamed desktop is a desktop running entirely on the user's local client device. The user interacts with the desktop directly, but the desktop is available only while the user is 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 the user is disconnected from the network.

For the purposes of the validation represented in this document, the hosted virtual desktop was used and is called the hosted VDI.

Hosted VDI Design Fundamentals

Citrix XenDesktop can be used to deliver a variety of virtual desktop configurations. When evaluating a hosted VDI deployment, consider the hypervisor selection and whether to use Citrix Provisioning Services.

Hypervisor Selection

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

- **Citrix XenServer:** Citrix XenServer is a complete, managed server virtualization platform built on the powerful Citrix Xen hypervisor. Citrix Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. Citrix XenServer is designed for efficient management of Microsoft Windows and Linux virtual servers and delivers cost-effective server consolidation and business continuity. More information about Citrix XenServer can be obtained at the company website.
- **VMware 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 such as VMware Distributed Resource Scheduler, vMotion, High Availability, and Storage vMotion; VMware File System (VMFS); and a multipathing storage layer. More information about VMware vSphere can be obtained at the company website.
- **Microsoft Hyper-V:** Microsoft Windows Server 2008 R2 Hyper-V builds on the architecture and functions of Microsoft Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility. Microsoft Hyper-V is available in standard, server core, and free Microsoft Hyper-V Server 2008 R2 versions. More information about Microsoft Hyper-V can be obtained at the company web site.

Citrix Provisioning Services

Hosted-VDI desktops can be deployed with or without Citrix Provisioning Services, but Citrix Provisioning Services enables you to stream a single desktop image to create multiple virtual desktops on one or more servers in a data center. This facility greatly reduces the amount of storage required compared to other methods of creating virtual desktops. Citrix Provisioning Services desktops can be deployed as pooled or private:

- **Private desktop:** A private desktop is a single private desktop assigned to one distinct user.
- **Pooled desktop:** A pooled virtual desktop uses Citrix Provisioning Services to stream a standard desktop image to multiple desktop instances upon bootup.

When considering a Citrix Provisioning Services deployment, you need to make some design decisions regarding the write cache for the virtual desktop device that is using provisioning. The write cache is a cache of all data that the target device has written. If data is written to the Citrix Provisioning Server vDisk in caching mode, the data is not written to the base vDisk. Instead, it is written to a write-cache file in one of the locations specified here:

- **Cache on local HD:** The cache on the local hard drive is stored in a file on a secondary local hard drive of the device. It is created as an invisible file in the root folder of the local hard drive. The cache file size

grows as needed, but it never gets larger than the original vDisk, and it is frequently not larger than the free space on the original vDisk.

- **RAM cache:** The RAM cache is stored in the client's RAM (memory). The cache's maximum size is fixed by a setting in the vDisk properties. All written data can be read from local RAM instead of going back to the server. The RAM cache is faster than the server cache and works in a high-availability environment.
- **Server cache:** The server cache is stored in a file on the server, or on a share, SAN, or other device. The file size grows as needed, but it never becomes larger than the original vDisk, and it frequently is not larger than the free space on the original vDisk. It is slower than the RAM cache because all read and write operations have to go to the server and be read from a file. The cache is 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.
- **Difference cache:** The difference cache is in a file on the server, or on a share, SAN, or other device. The cache file size grows as needed, but it never becomes larger than the original vDisk, and it frequently is not larger than the free space on the original vDisk. It is slower than the RAM cache and the server cache.

Designing a Citrix XenDesktop 5.5 Deployment

For a complete overview of configurations, architecture, and design recommendations for delivering virtual desktops with Citrix XenDesktop, visit <http://support.citrix.com/proddocs/index.jsp?topic=/xendesktop-bdx/cds-admin-deploy-planwrapper-bdx.html>.

Solution Validation

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

Configuration Topology for Scalable Citrix XenDesktop 5.5 VDI on Cisco UCS and Storage Solution

Figure 4 shows the reference configuration for the Citrix XenDesktop 5.5 VDI on Cisco UCS and Storage solution.

Figure 4. Cisco UCS VDI Reference Configuration

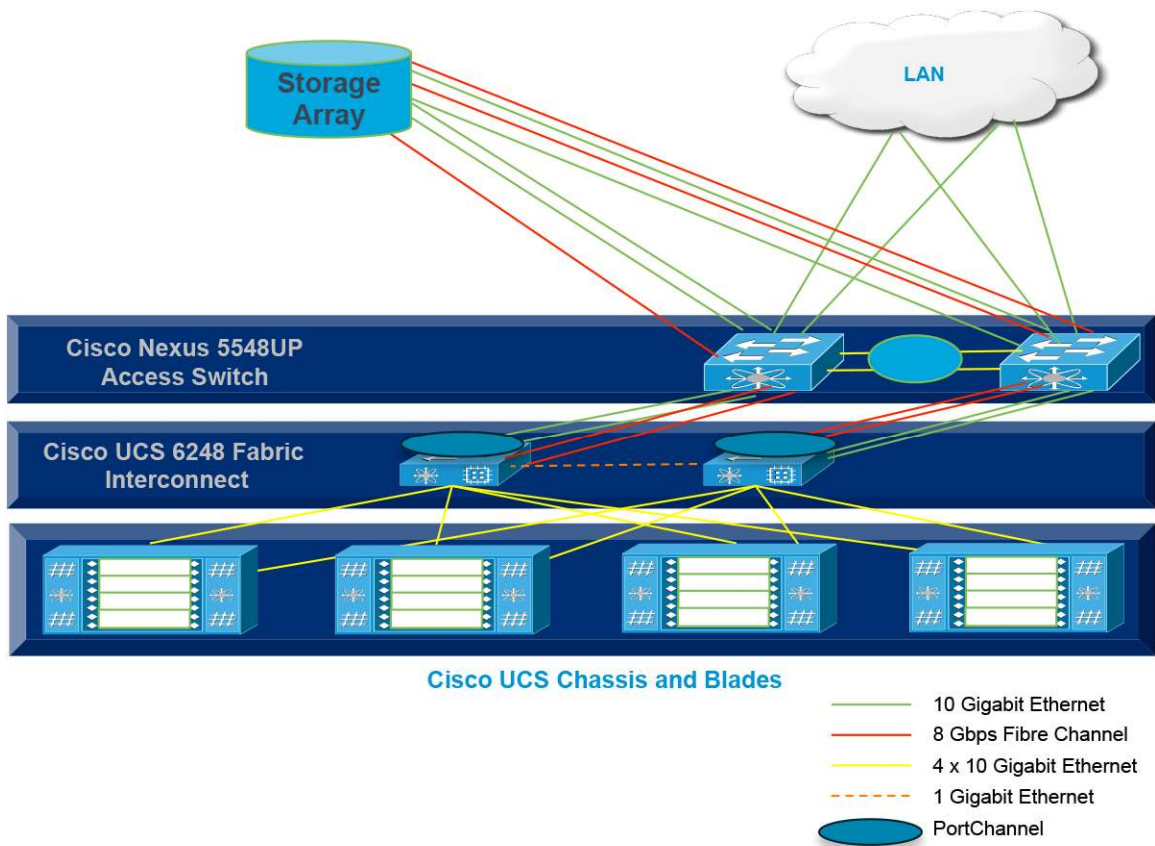



Figure 5 shows the architecture used in this study. The architecture is divided into four layers:

- Cisco UCS computing platform
- VDI that runs on Cisco UCS blade hypervisor hosts
- Network access layer and LAN
- SAN and storage array

Figure 5. Detailed Architecture of the Configuration



Cisco Unified Computing System Configuration

This section discusses the Cisco UCS configuration performed as part of the infrastructure buildout. The racking, power, and installation of the chassis are described in the installation guide (see http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html) and are beyond the scope of this document. More details about each step can be found in the following documents:

- Cisco UCS command-line interface (CLI) configuration guide:
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.0/b_UCSM_CLI_Configuration_Guide_2_0.html
- Cisco UCS Manager GUI configuration guide:
http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/2.0/b_UCSM_GUI_Configuration_Guide_2_0.html

Base Cisco UCS System Configuration

To configure Cisco UCS, perform the steps described here.

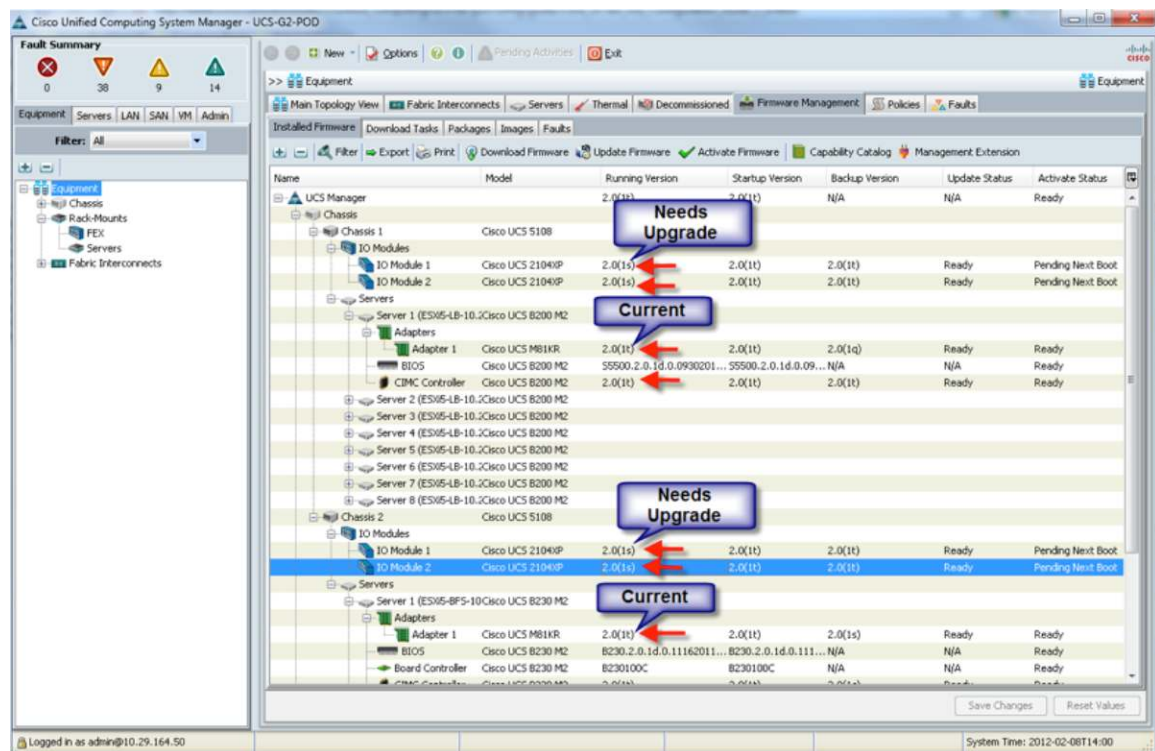
- Step 1. Bring up the fabric interconnect and from a serial console connection, set the IP address, gateway, and hostname of the primary fabric interconnect. Then bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary interconnect and a dialog box asks you if you want to be part of the cluster. Answer Yes and set the IP address, gateway, and hostname. After this is done, all access to the fabric interconnect can be

performed remotely. You will also configure the virtual IP address to connect to the fabric interconnect. You need a total of three IP addresses to bring the fabric interconnect online. You can also connect the chassis to the fabric interconnect, using one, two, or four links per I/O module, depending on your application's bandwidth requirements. In the configuration discussed in this document, all four links were connected to each module.

Step 2. Next connect using your browser of choice to the virtual IP address and launch Cisco UCS Manager. The Java-based Cisco UCS Manager will let you do everything that you can do from the CLI; this document focuses on the GUI methodology.

Step 3. First check the firmware on the system and to see whether it is current (Figure 6). The firmware release used in this document is Cisco UCS Firmware Release 2.0(1t).

Figure 6. Checking the Firmware Release

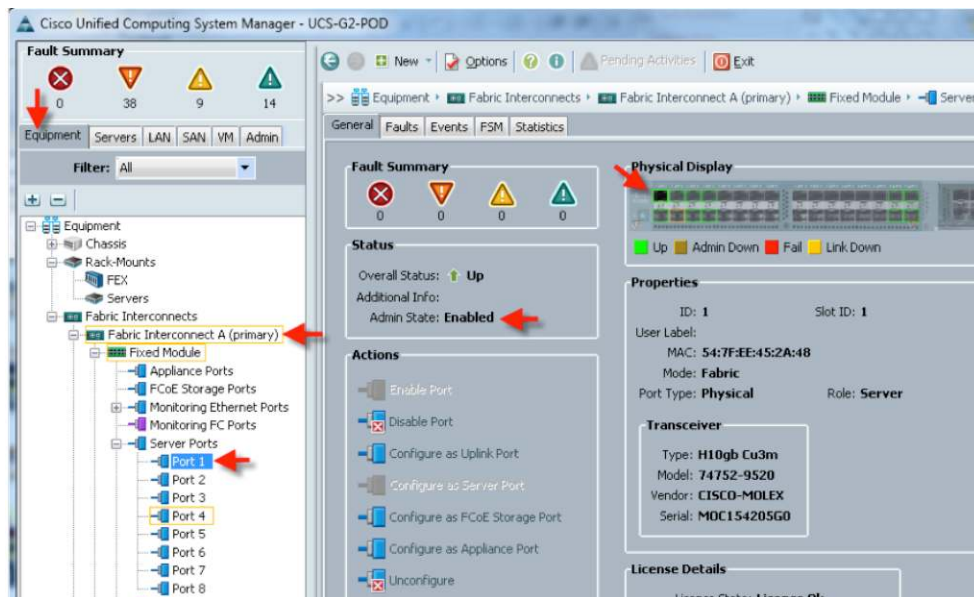


If the firmware is not current, refer to the installation and upgrade guide to upgrade the Cisco UCS firmware. Also do not forget to upgrade the BIOS to the latest version and associate it with all the blades.

Note: The BIOS and board controller version numbers do not match the I/O module, adapter, or Cisco Integrated Management Controller version number.

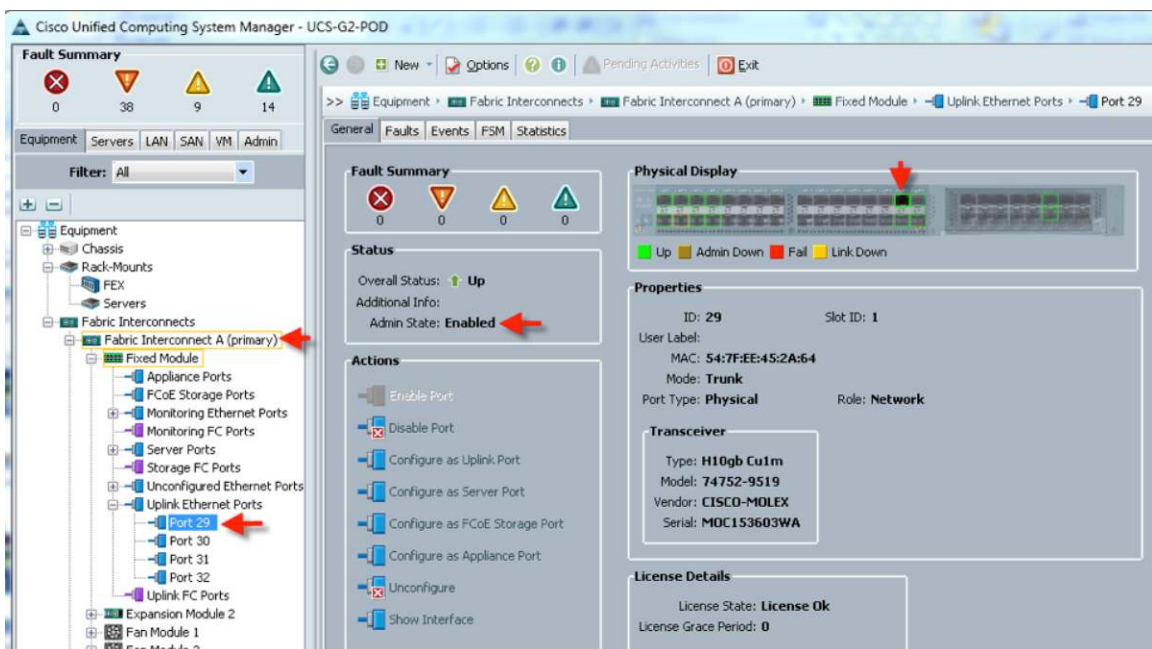
Configure and enable the server port on the fabric interconnect (Figure 7). To bring the chassis online, acknowledge the chassis.

Figure 7. Configuring and Enabling the Server Port on the Fabric Interconnect



Configure and enable upstream Ethernet links and Fibre Channel links (Figure 8).

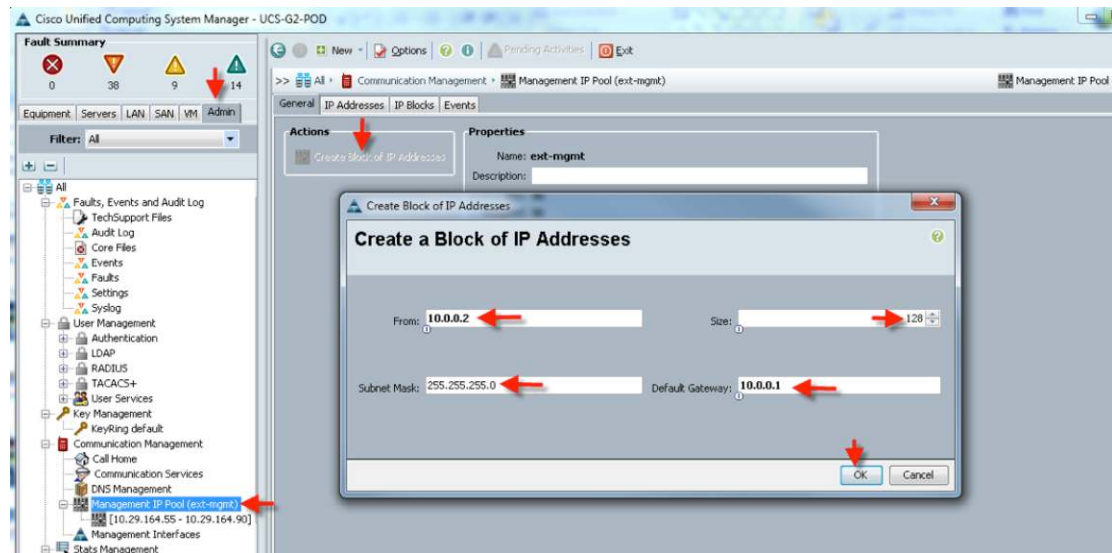
Figure 8. Configuring and Enabling Upstream Ethernet Links and Fibre Channel Links



After the blades are discovered, set the kernel-based virtual machine (KVM) IP addresses for each blade. To do this, select the Admin tab and choose Communication Management > Management IP Address Pool (Figure 9).

Note: Make sure that you have a sufficient number of IP addresses for all the blades, and make sure that the gateway and subnet mask are set correctly.

Figure 9. Setting the KVM IP Addresses for Each Blade



Create all the pools: MAC address pool, worldwide port name (WWPN) pool, worldwide node name (WWNN) pool, universal user ID (UUID) pool, and server pool (Figure 10 through 14).

Figure 10. Creating the MAC Address Pool

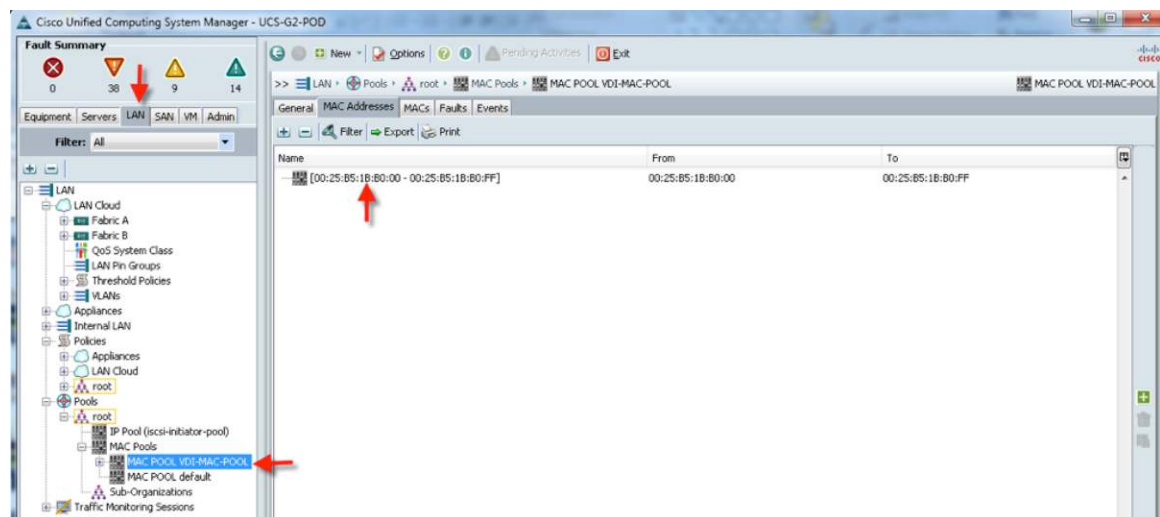


Figure 11. Creating the WWPN Pool

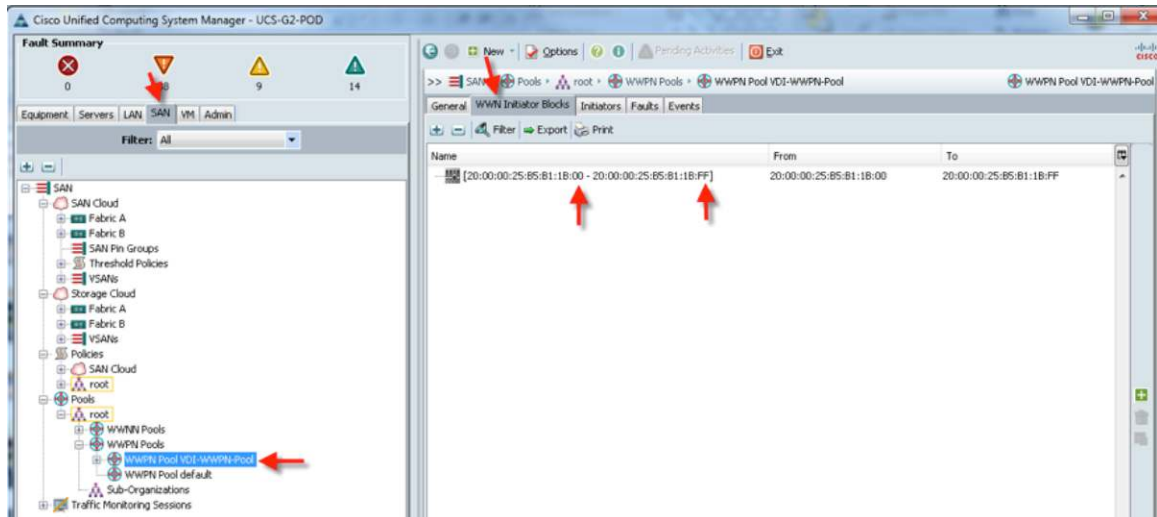


Figure 12. Creating the WWNN Pool

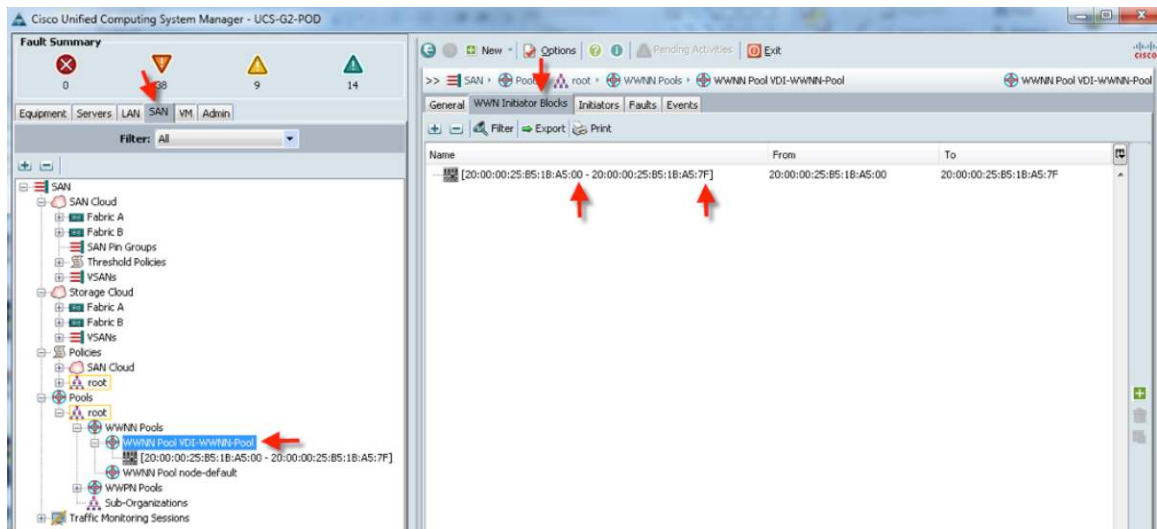


Figure 13. Creating the UUID Pool

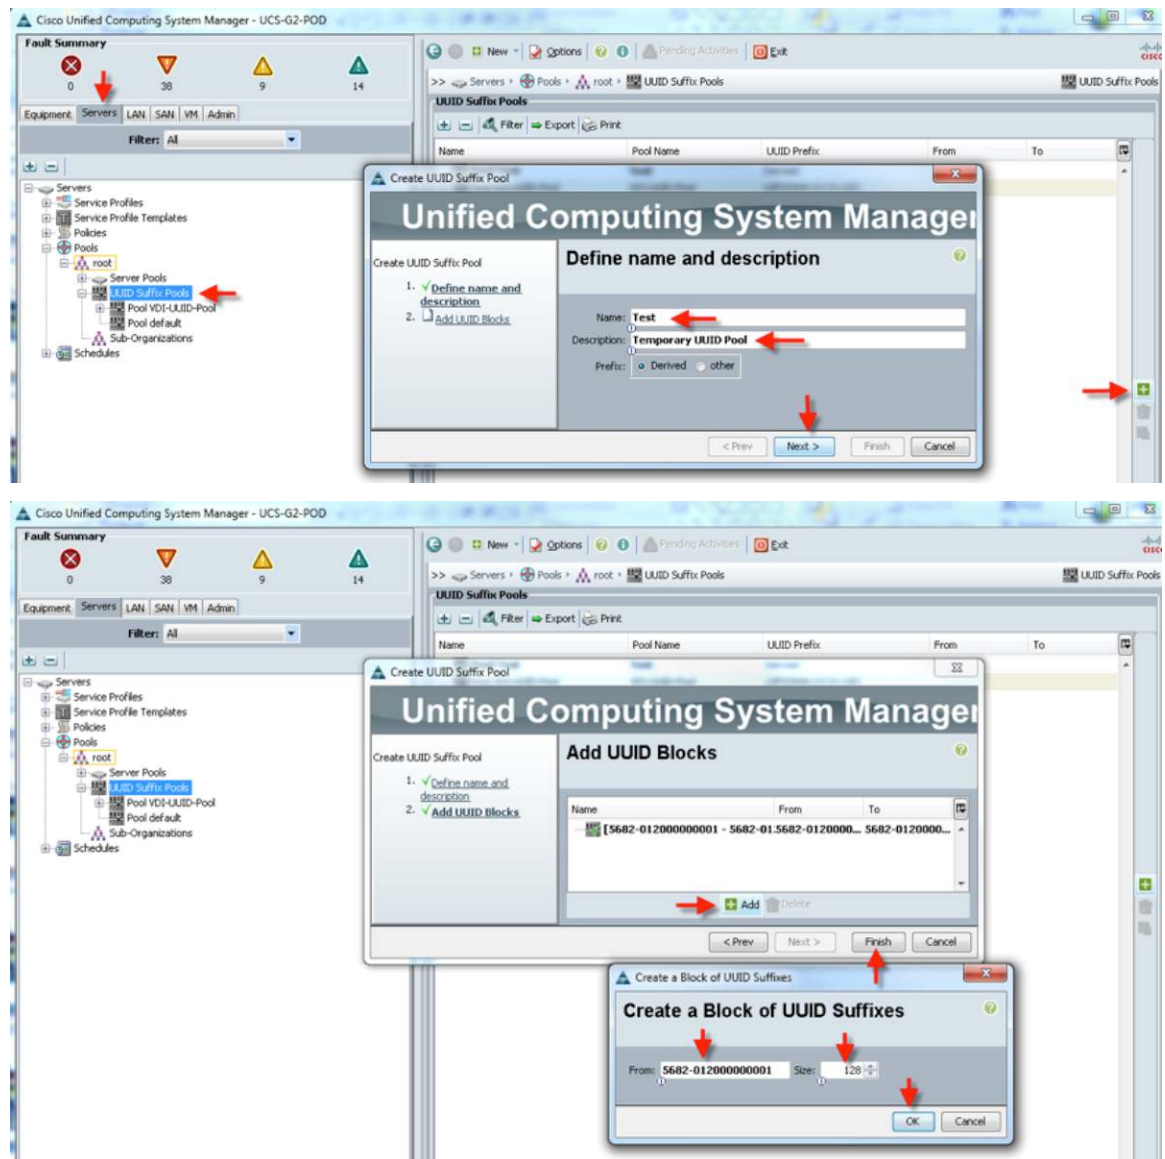
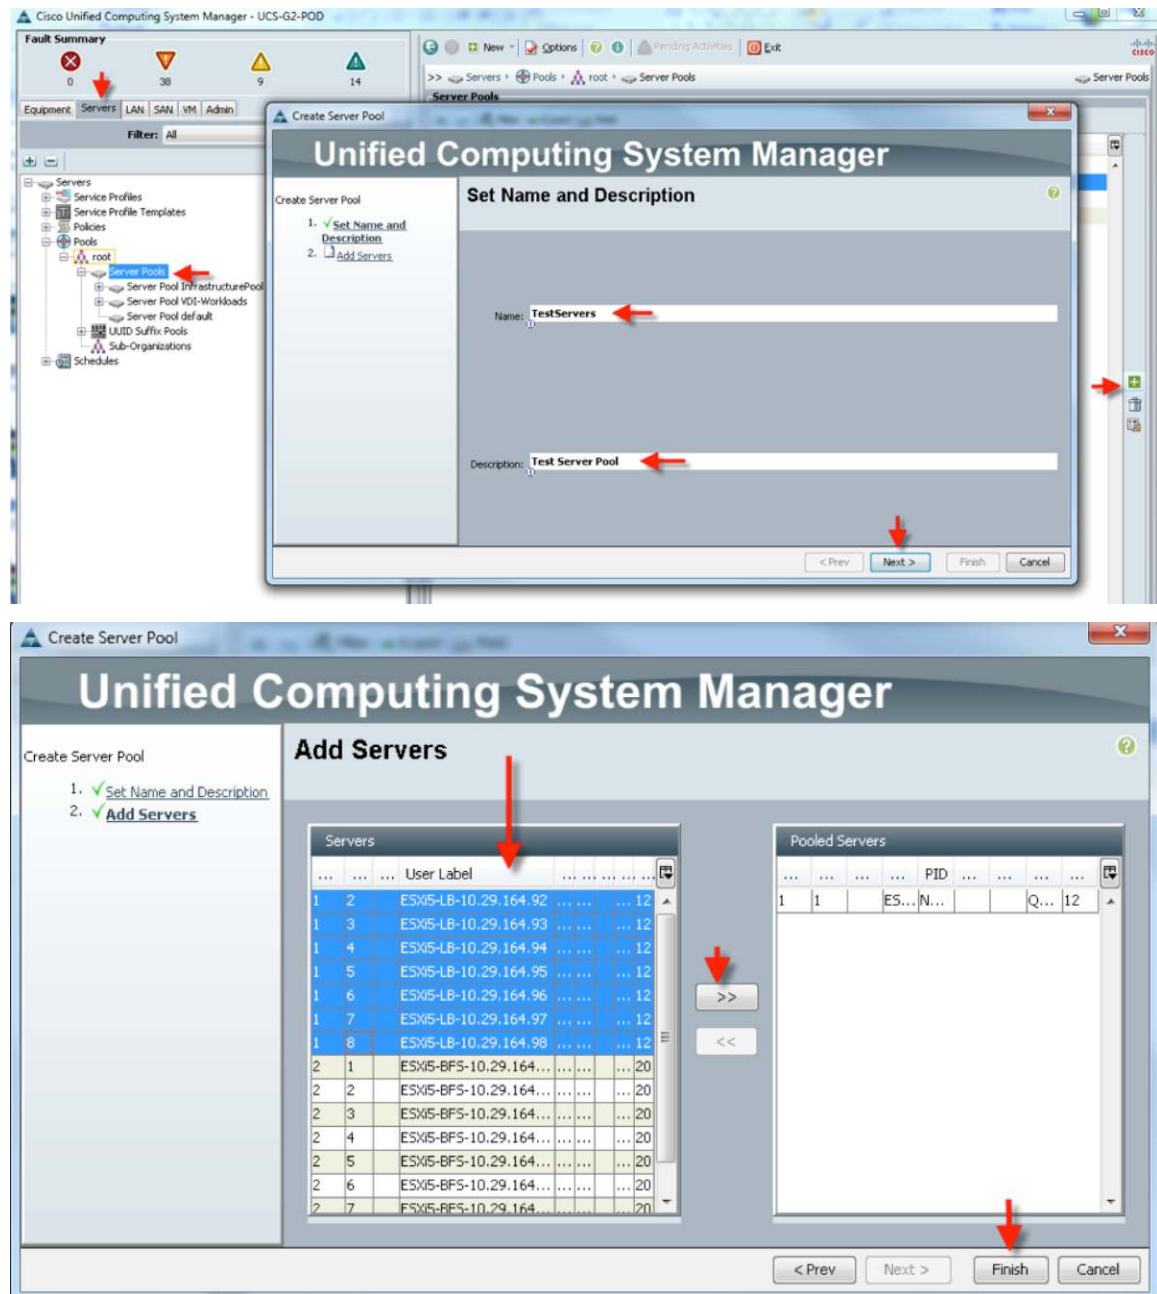
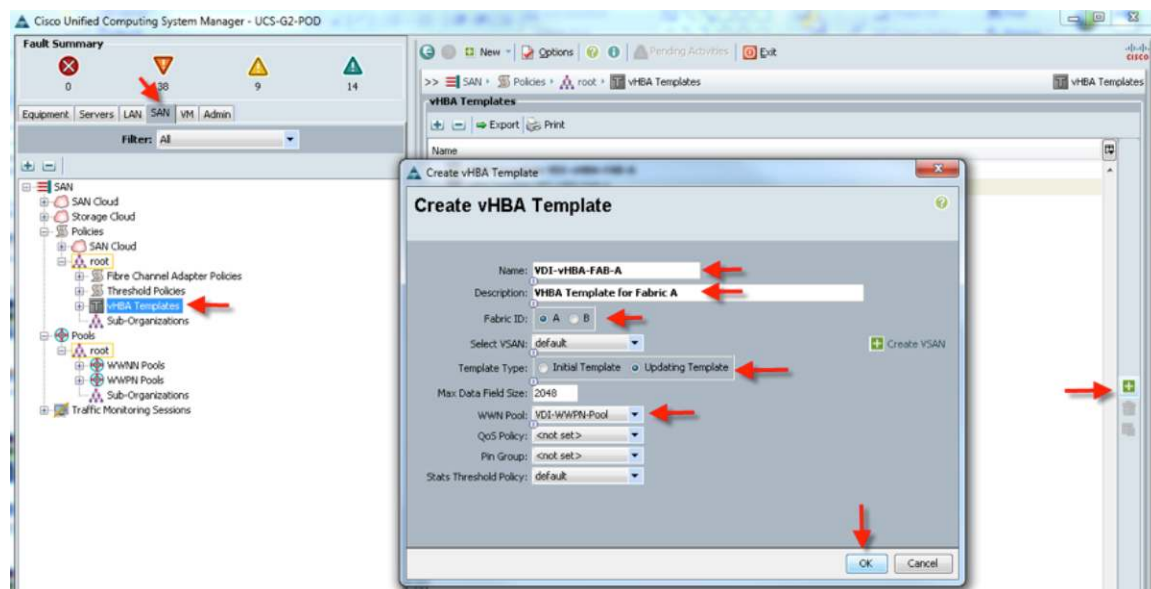


Figure 14. Creating the Server Pool



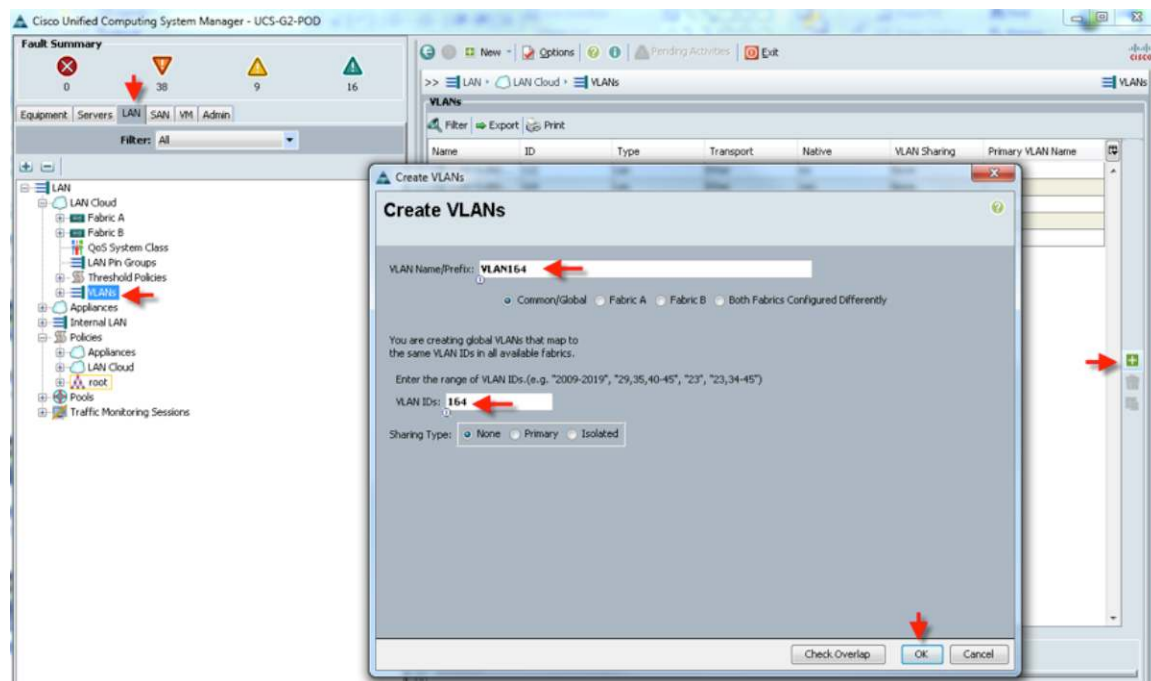
Create the virtual host bus adapter (vHBA) template (Figure 15). Create at least one vHBA template for each fabric interconnect if block storage will be used.

Figure 15. Creating the vHBA Template



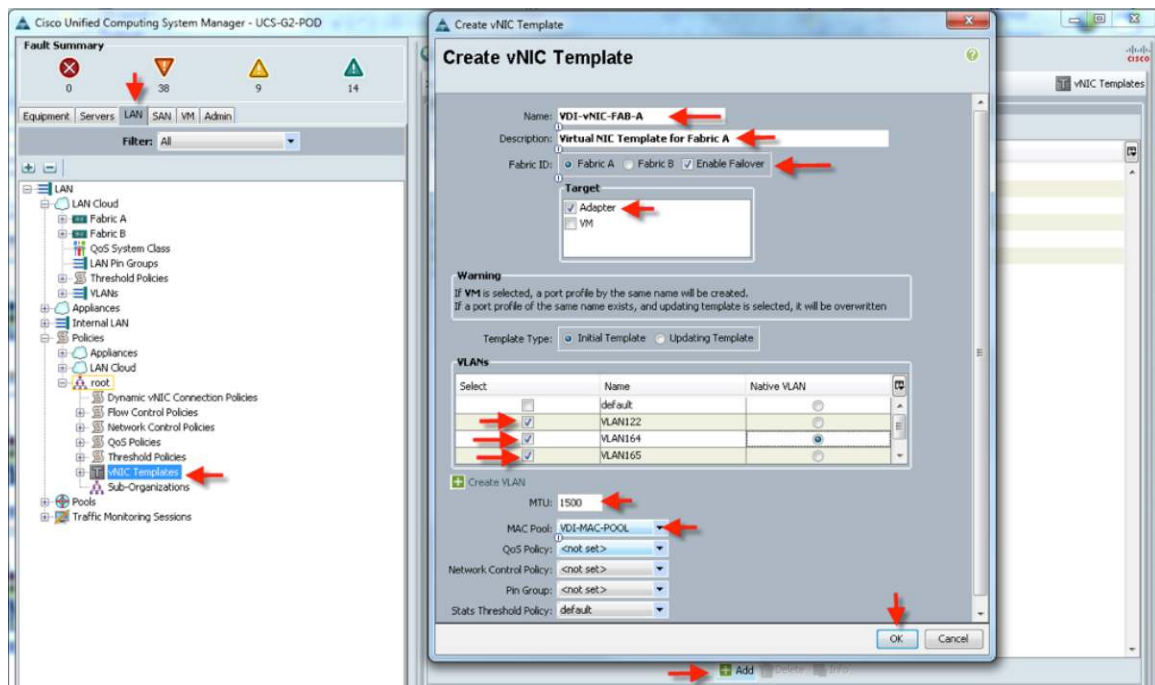
Optionally, create VLANs (Figure 16).

Figure 16. Creating VLANs



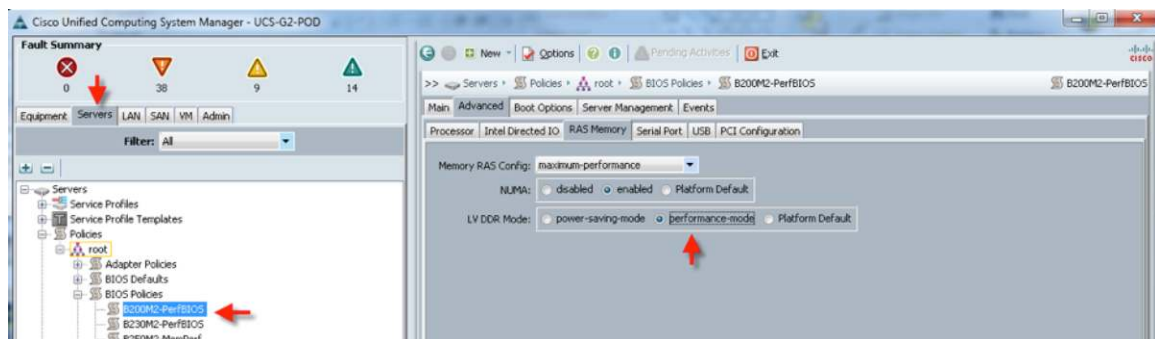
Create virtual network interface card (vNIC) templates for both fabrics, select Enable Failover, select the VLANs supported on the adapter (optional,) set the maximum transmission unit (MTU) size, and select the MAC address pool; then click OK (Figure 17).

Figure 17. Creating a vNIC Template



Create boot-from-SAN policies, adapter policies, and local boot policy if desired (for VMware ESXi 5, the installation was configured as RAID mirrored to avoid any catastrophic disk failures). Create performance BIOS policies for each blade type to help ensure that your low-voltage 8-GB 1333-MHz DIMMs will operate at top speed (Figure 18).

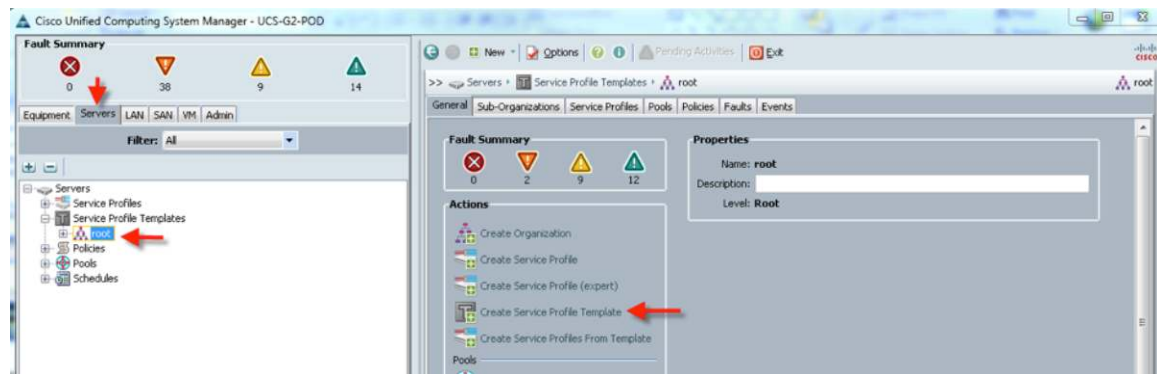
Figure 18. Creating Performance Policies



Note: Be sure to save your changes by clicking the button at the bottom of the page to preserve these settings. Be sure to add this policy to your blade service profile template.

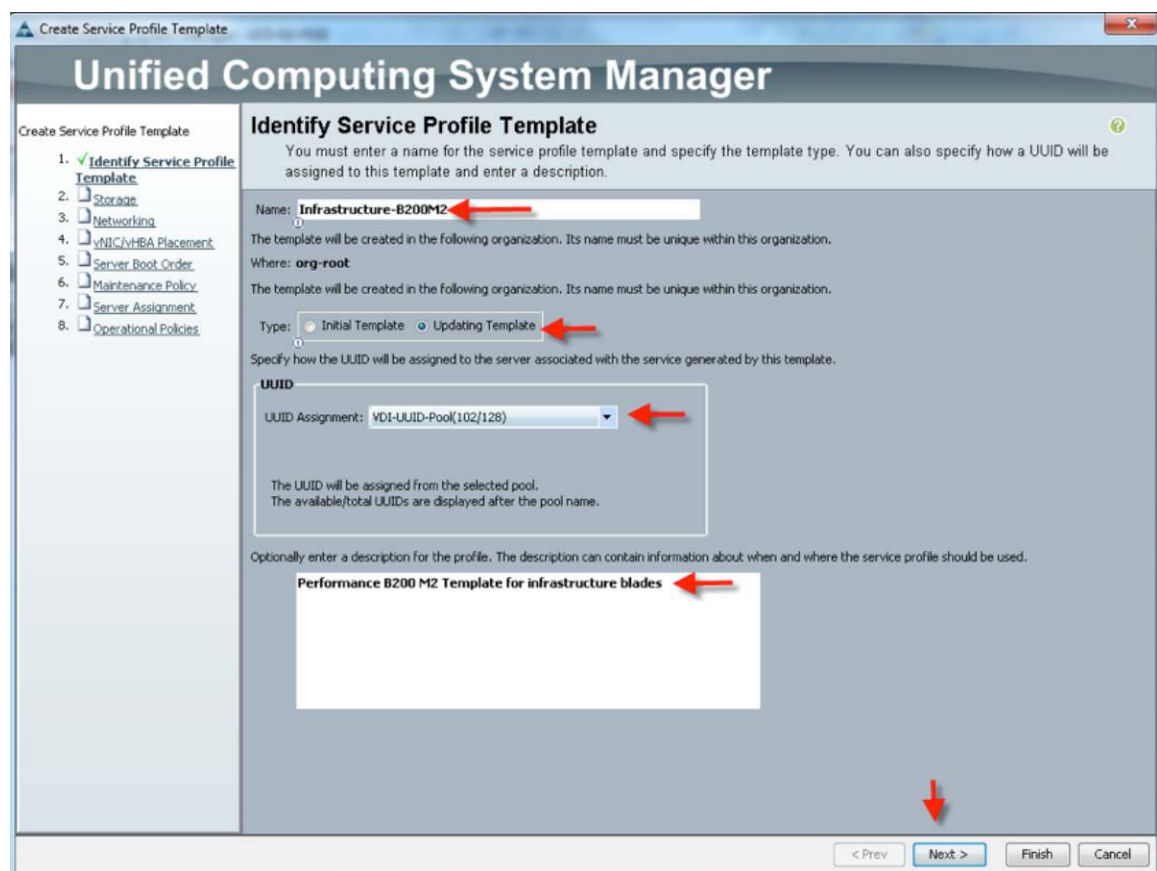
Create a service profile template using the pools, templates, and policies you have configured (Figure 19).

Figure 19. Creating a Service Profile Template



Complete each section of the dialog box that appears, using the policies and objects that you created earlier; then click Finish (Figure 20).

Figure 20. Completing the Service Profile Template



Note: On the Operational Policies page, select the appropriate performance BIOS policy you created earlier to help ensure maximum low-voltage DIMM performance.

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

Right-click a service profile template to deploy as many service profiles as you need, and Cisco UCS Manager will automatically start configuring these new service profile templates on the selected blade servers.

At this point, the servers are ready for OS provisioning. You should set up a preboot execution environment (PXE) server for the OS installation. Virtual media CD-based OS installation is also possible

Quality of Service and Class of Service in Cisco UCS

Cisco UCS provides different system classes of service (CoSs) to implement quality of service (QoS), 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 Cisco UCS and other time-sensitive applications have to adhere to a strict QoS policy for optimal performance.

System Class Configuration

System class is the global operation where entire system interfaces are with defined QoS rules.

- By default, the system has a best-effort class and a Fibre Channel over Ethernet (FCoE) class.
 - Best effort is equivalent to “match any” in modular QoS CLI (MQC) terminology.
 - FCoE is a special class defined for FCoE traffic. The equivalent in MQC terminology is “match CoS 3.”
- System class allowed with four more users define class with following configurable rules:
 - CoS-to-class map
 - Weight: Bandwidth
 - Per-class MTU
 - Class property: drop or no drop
- The highest MTU per class allowed is 9216.
- With Cisco UCS, one CoS value can be mapped to one particular class.
- In addition to the FCoE class, only one more class can be configured with the no-drop property.
- Weight can be configured as a value from 0 to 10. Internally, the system will calculate the bandwidth based on following equation (the number will be rounded):

Percentage of bandwidth shared by given class = (Weight of the given priority x 100) divided by (Sum of weights of all priorities)

Cisco UCS System Class Configuration

Cisco UCS defines user class names as follows:

- Platinum
- Gold
- Silver

- Bronze

Table 2 shows the relationship between Cisco UCS class names and Cisco NX-OS Software class names. Table 3 shows the default class-to-CoS mapping for Cisco UCS. Table 4 shows the default weight of classes in Cisco UCS.

Table 2. Mapping of Cisco UCS and Cisco NX-OS Class Names

Cisco UCS Class Names	Cisco NX-OS Class Names
Best effort	Class-default
Fibre Channel (FC)	Class-fc
Platinum	Class-Platinum
Gold	Class-Gold
Silver	Class-Silver
Bronze	Class-Bronze

Table 3. Default Class-to-CoS Map for Cisco UCS

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Fibre Channel (FC)	3
Platinum	5
Gold	4
Silver	2
Bronze	1

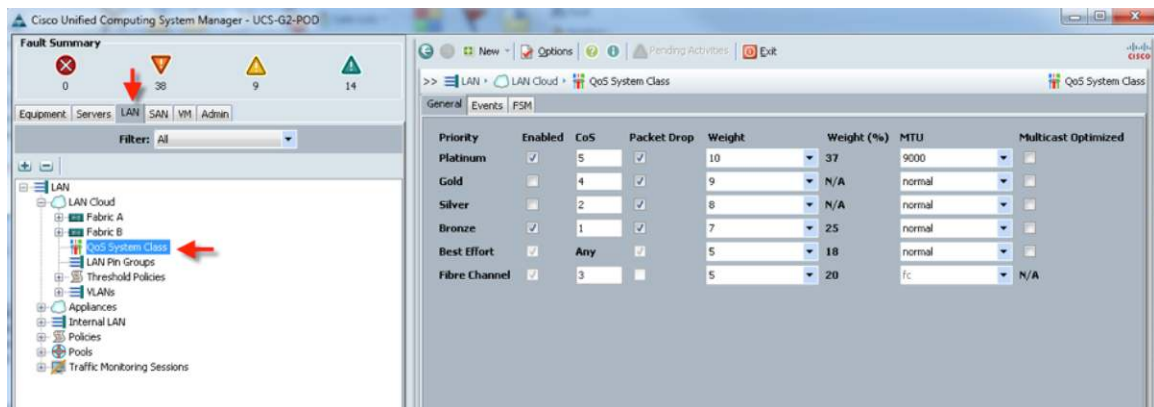
Table 4. Default Class-to-CoS Map for Cisco UCS

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	5
Fibre Channel (FC)	5

Steps to Enable QoS on Cisco UCS

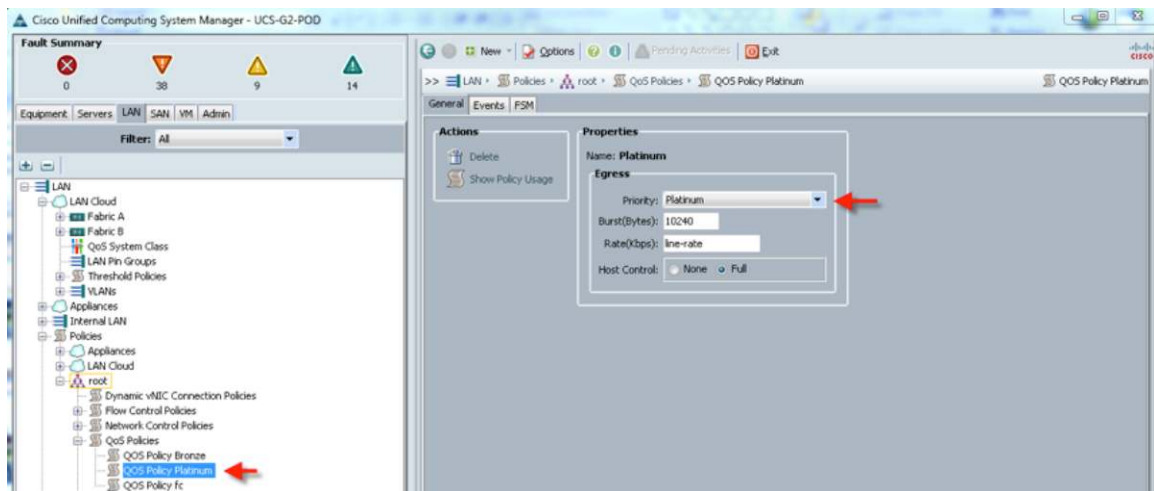
Configure the platinum policy by selecting the Platinum policy box. If you want jumbo frames enabled, change the MTU value from normal to 9000. Notice the option to set no-packet-drop policy for this configuration (Figure 21).

Figure 21. Cisco UCS QoS System Class Configuration



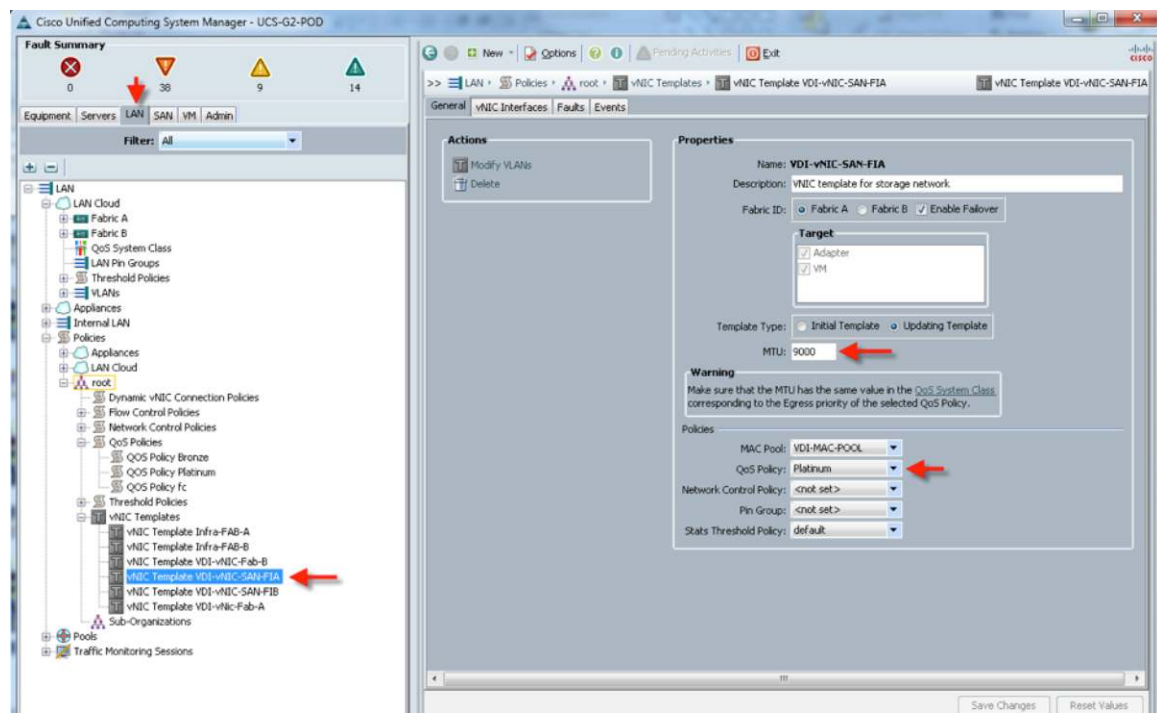
Next, on the LAN tab under Policies > Root > QoS Policies, verify that QoS Policy Platinum exists and that Platinum is set as the priority (Figure 22).

Figure 22. Cisco UCS QoS Policy Configuration



Finally, include the QoS Policy Platinum policy in the vNIC template under the QoS policy (Figure 23).

Figure 23. Use QoS Policy in the vNIC Template

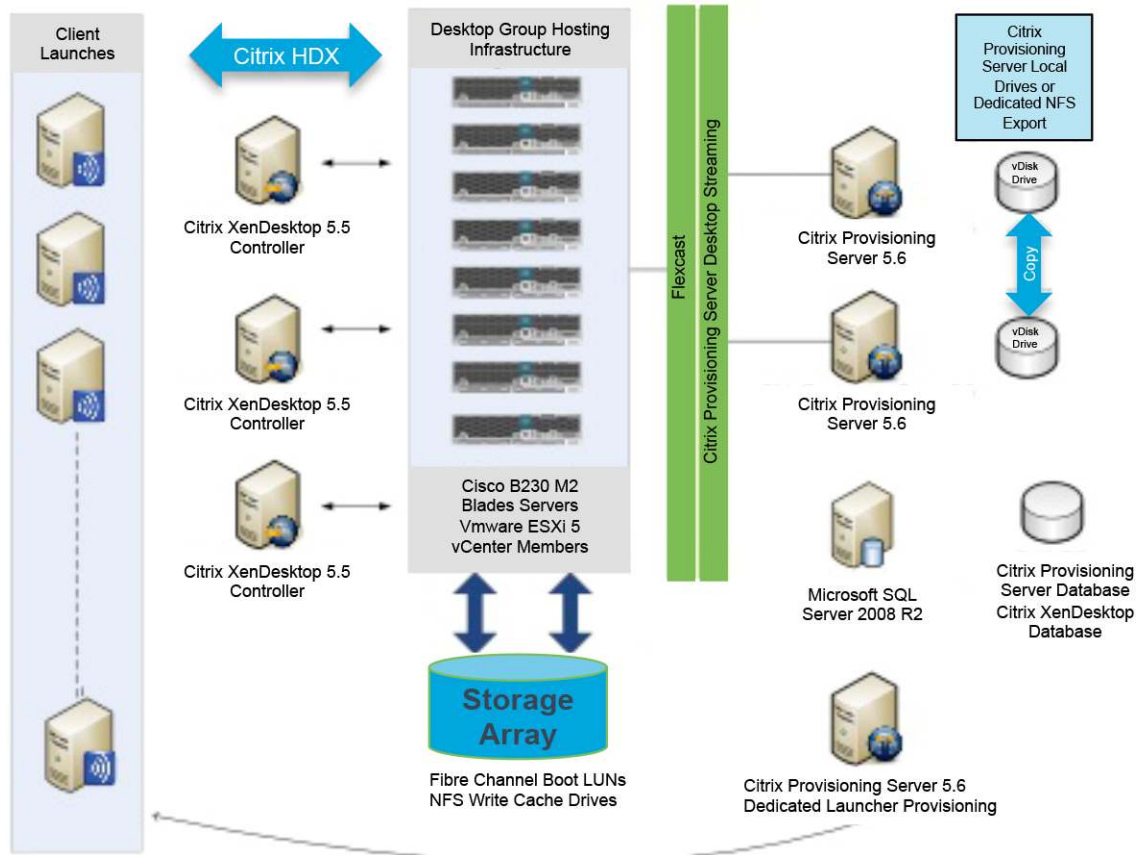


This end-to-end QoS policy is a unique value proposition for Cisco UCS. For example, if you have a VLAN for storage, configure the platinum policy with jumbo frames and get end-to-end QoS and performance guarantees.

Citrix XenDesktop 5.5 and Provisioning Server 5.6 Configuration

Figure 24 shows the Citrix XenDesktop high-level configuration.

Figure 24. Citrix XenDesktop Configuration



Here is a summary of the environment:

- 3 XenCitrix XenDesktop 5.5 delivery controllers
- 2 Citrix Provisioning Server 5.6 SP1 servers for virtual desktops
- 1 Citrix Provisioning Server 5.6 SP1 for virtual desktops
- 1 VMware vCenter with 2 infrastructure Cisco UCS B200 M2 blades, 8 Cisco UCS B230 M2 blades hosting Microsoft Windows 7 virtual machines, and 2 launcher blades for Login VSI client launchers
- 1080 virtual desktops
- 1 Citrix Licensing Server
- 2 file servers: 1 for Microsoft roaming profiles, and 1 for Login VSI
- 1 Microsoft SQL Server 2008 R2 for Citrix Provisioning Services and Citrix XenDesktop
- 1 storage solution includes 6 NFS volumes, and 12 Fibre Channel boot logical unit numbers (LUNs)
- 40 Login VSI launchers

Tables 5 through 9 provide additional details about the environment.

Table 5. VMware ESXi 5

VMware vSphere ESXi 5.0.0 Build 469512 hosts			
Hardware	Cisco B-Series blade servers	Model	Cisco UCS B230 M2
OS	VMware vSphere 5.0	Service Pack	–
CPU	2 x 10-core Intel Xeon processor E7-2870, at 2.4 GHz (40 logical cores total)	RAM	256 GB
Disk	Boot from SAN	Network	4 x 10 Gigabit Ethernet
<ul style="list-style-type: none"> • Cisco UCS M81KR NIC driver updated to enic_driver_2.1.2.22-564611 • Cisco UCS M81KR HBA driver updated to fnic_driver_1.5.0.7-563432 			

Table 6. Citrix Provisioning Server 5.6 SP1

Citrix Provisioning Server 5.6			
OS	Microsoft Windows 2008 R2 Enterprise 64-bit	Service Pack	SP1
CPU	2 x vCPU	RAM	8192 MB
Disk	1 x 24-GB boot virtual disk and 1 x 350-GB vDisk drive (hosted on NFS target volume)	Network	1 x 10 Gigabit Ethernet
<ul style="list-style-type: none"> • Database for Citrix Provisioning Services hosted on separate Microsoft SQL Server 2008 R2 SP1 64-bit • CPVS56SP1E029, CPVS56SP1E043, CPVS56SP1E046, and CPVS56SP2E003 applied to Citrix Provisioning Servers • CPVS56SP1E033 and CPVS56SP2E002 applied to target vDisks 			

Table 7. Citrix XenDesktop 5.5

Citrix XenDesktop 5.5 desktop delivery controllers			
OS	Microsoft Windows 2008 R2 Enterprise 64-bit	Service Pack	SP1
CPU	1 x vCPU, Intel Xeon processor 5690, 3.46 GHz	RAM	2048 MB
Disk	1 x 24-GB virtual disk (hosted on NFS infrastructure volume on storage solution)	Network	1 x 10 Gigabit Ethernet
<ul style="list-style-type: none"> • Database for DDC hosted on separate Microsoft SQL Server 2008 R2 SP1 64-bit 			

Table 8. Citrix License Server

Citrix License Server			
OS	Microsoft Windows 2008 R2 Enterprise 64-bit	Service Pack	SP1
CPU	1 x vCPU	RAM	2048 MB
Disk	1 x 24-GB virtual disk (hosted on NFS infrastructure volume on storage solution)	Network	1 x 10 Gigabit Ethernet

Table 9. Login VSI 3.0 Launchers

ICA client hosts (VSI Launchers)			
OS	Microsoft Windows 2008 R2 Enterprise 64-bit	Service Pack	SP1
CPU	1 x vCPU	RAM	4096 MB
Disk	1 x 24-GB virtual disk (hosted on NFS infrastructure volume on storage solution)	Network	1 x 10 Gigabit Ethernet

Other dedicated Infrastructure Virtual Machines, all running Server 2008 R2 SP1:

- 2 Microsoft Active Directory servers (directory services, DNS, and Dynamic Host Configuration Protocol [DHCP])
- User profile server (hosts Microsoft roaming profiles)

- File server (collects Login VSI data)
- VMware vCenter server

Citrix XenDesktop 5.5 Desktop Delivery Controller

The Citrix XenDesktop 5.5 DDCs were virtualized on VMware ESXi 5 running on Cisco UCS B200 M2 infrastructure blades.

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

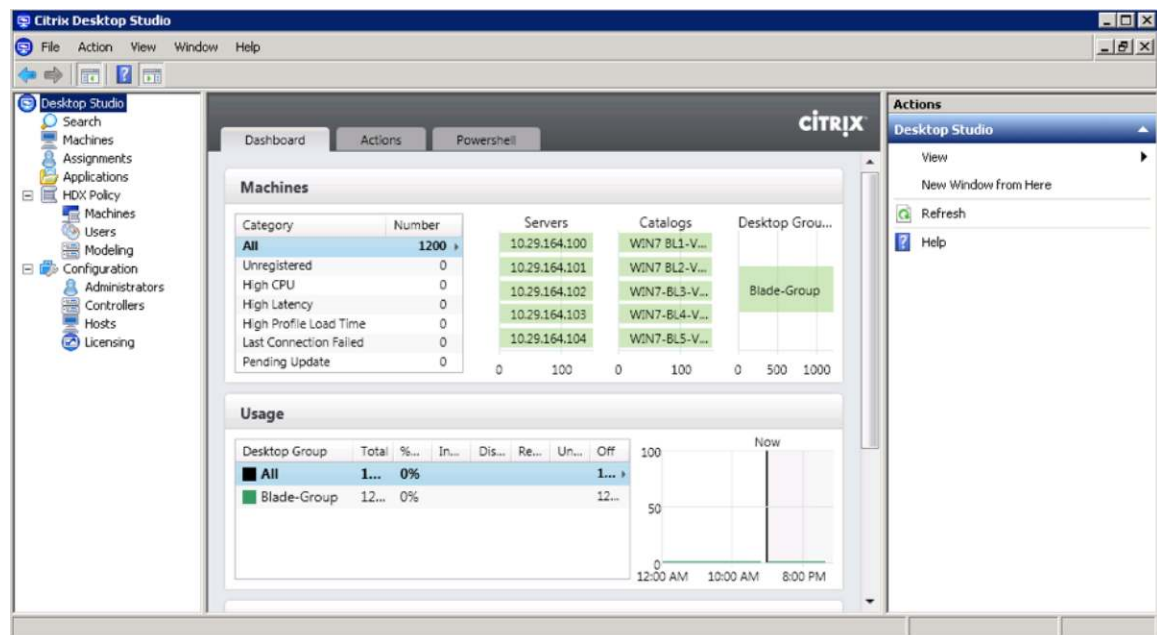
There is no longer a “master” DDC. All DDCs communicate directly with the SQL database continuously to update status and use, providing transparent fault tolerance. Citrix recommends at least two DDCs per site.

For this study, three DDCs were created and used. There were no public hot fixes or service packs for Citrix XenDesktop 5.5 posted during the study. For management, Citrix introduced two new management consoles beginning with Citrix XenDesktop 5:

- Citrix Desktop Studio
- Citrix Desktop Director

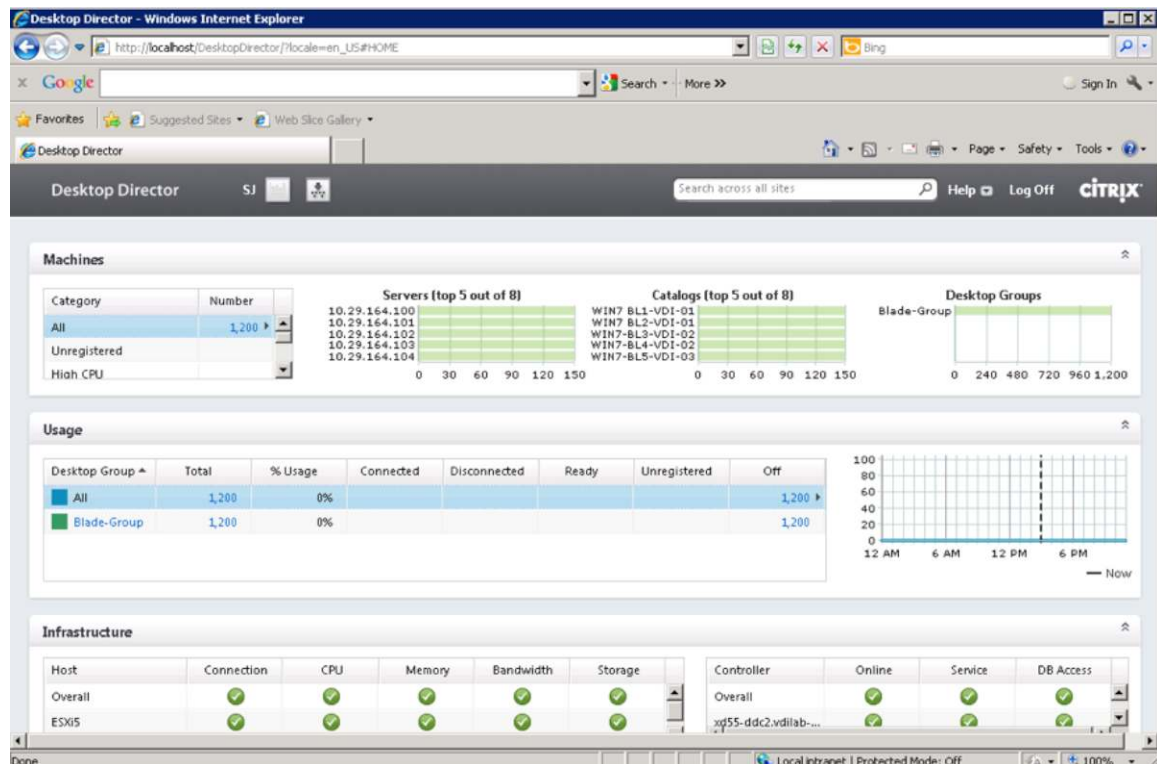
Citrix Desktop Studio is the main administrative console at which hosts, machine catalogs, desktop groups, and applications are created and managed. Citrix Desktop Studio is where Citrix HDX policy is configured and applied to the site. Citrix Desktop Studio is a Microsoft Management Console plug-in and fully supports Microsoft Windows PowerShell (Figure 25).

Figure 25. Citrix XenDesktop 5.5 Desktop Studio



Citrix Desktop Director is designed for use by level-1 and level-2 help-desk personnel. It provides real-time status and limited management capabilities for the running Citrix XenDesktop infrastructure. Help-desk personnel can provide real-time status to end users, shadow their sessions, and restart their desktops. Citrix Desktop Director uses Microsoft Internet Explorer and Adobe Flash to present data (Figure 26).

Figure 26. Citrix XenDesktop 5.5 Desktop Director

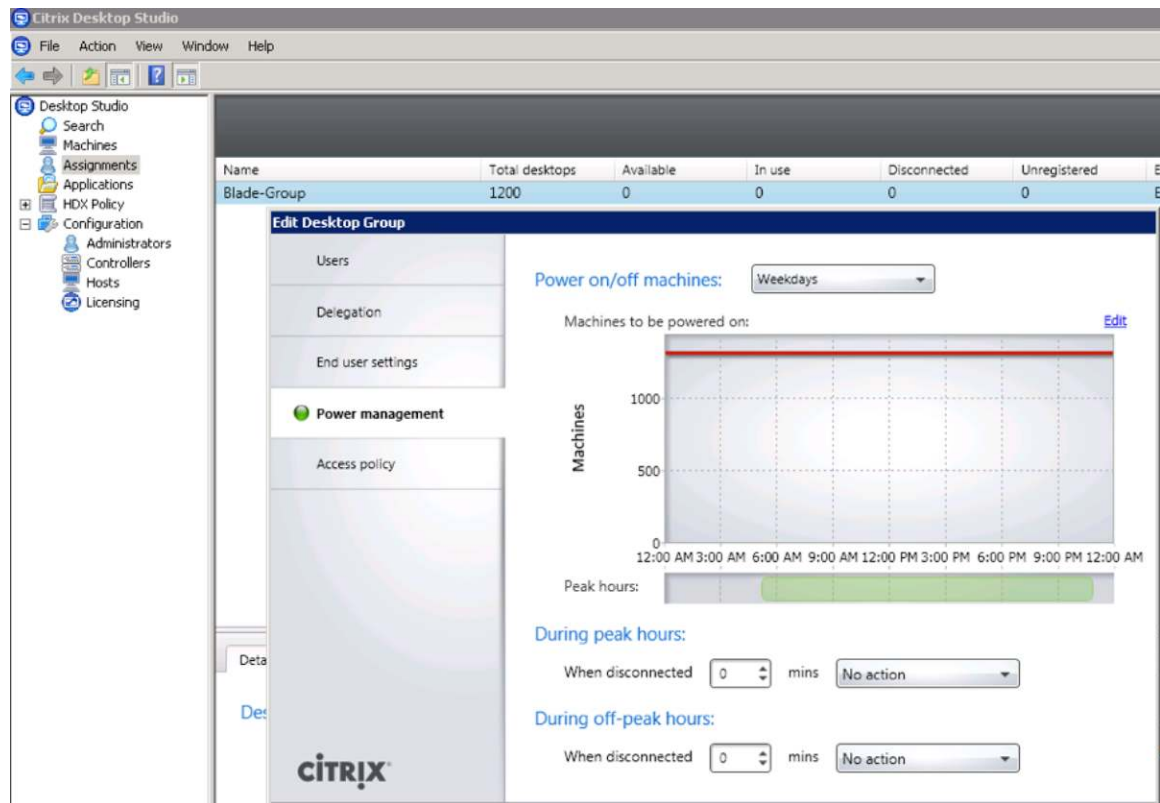


Site Configuration

In addition to the standard Citrix XenDesktop site installation, the following additional items were configured:

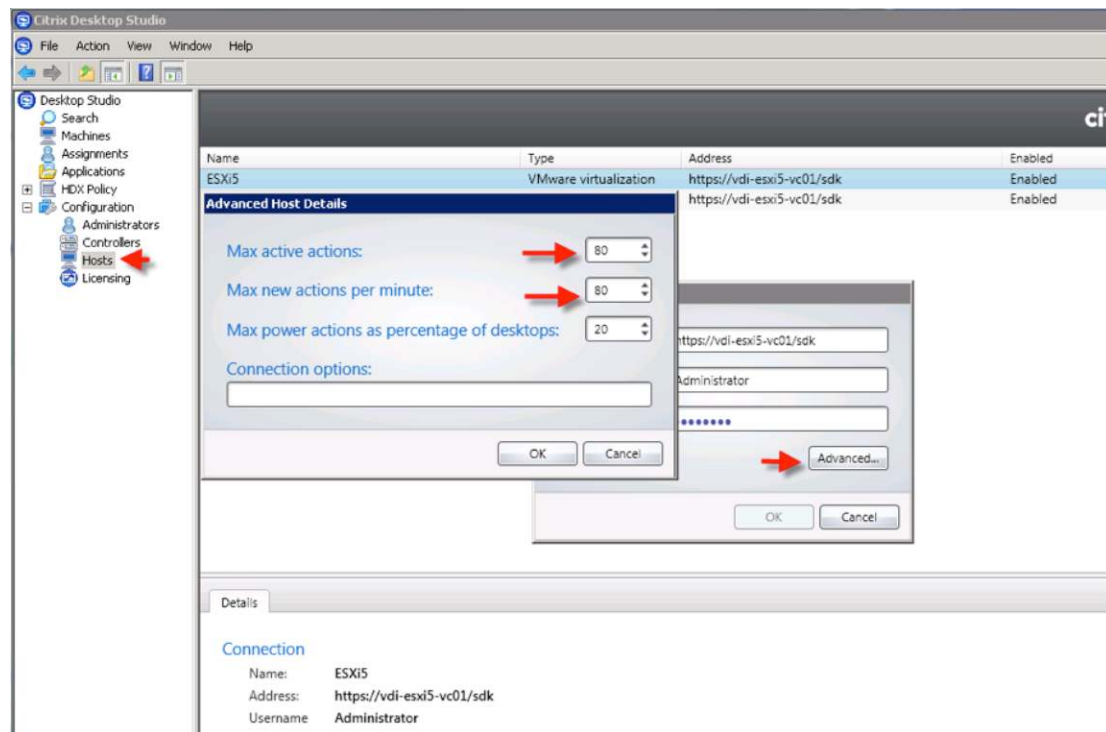
- Eight machine catalogs corresponding to each of the eight blades in the test environment
- One desktop group containing all the desktops from each catalog
- Citrix XenDesktop HDX user policies to disable autoconnect client drives, client printer mappings, autocreation of a generic universal printer, and automatic installation of in-box printer drivers
- Power settings for the desktop group to start the desired number of virtual desktops when the group is removed from maintenance mode on the Assignments node of Citrix Desktop Studio; action upon disconnect set to “No action” (Figure 27)

Figure 27. Citrix XenDesktop Power Settings



- Maximum active actions and maximum new actions per minute changed from 10 (the default value) to 10 x Y, where Y represented the number of Cisco UCS B230 M2 blades participating in a test in the Advanced Host Details section of the Configuration > Hosts node of Citrix Desktop Studio (Figure 28)

Figure 28. Citrix XenDesktop Maximum Active and New Actions



Citrix Provisioning Services

Citrix Provisioning Services is part of the Citrix XenDesktop Enterprise and Platinum suites and was used in all tested scenarios. Provisioning provides the capability to create thousands of virtual machines hosted on hypervisor servers with identical virtual machine configurations. It allows those virtual machines to boot using PXE from a single Microsoft Windows 7 golden image.

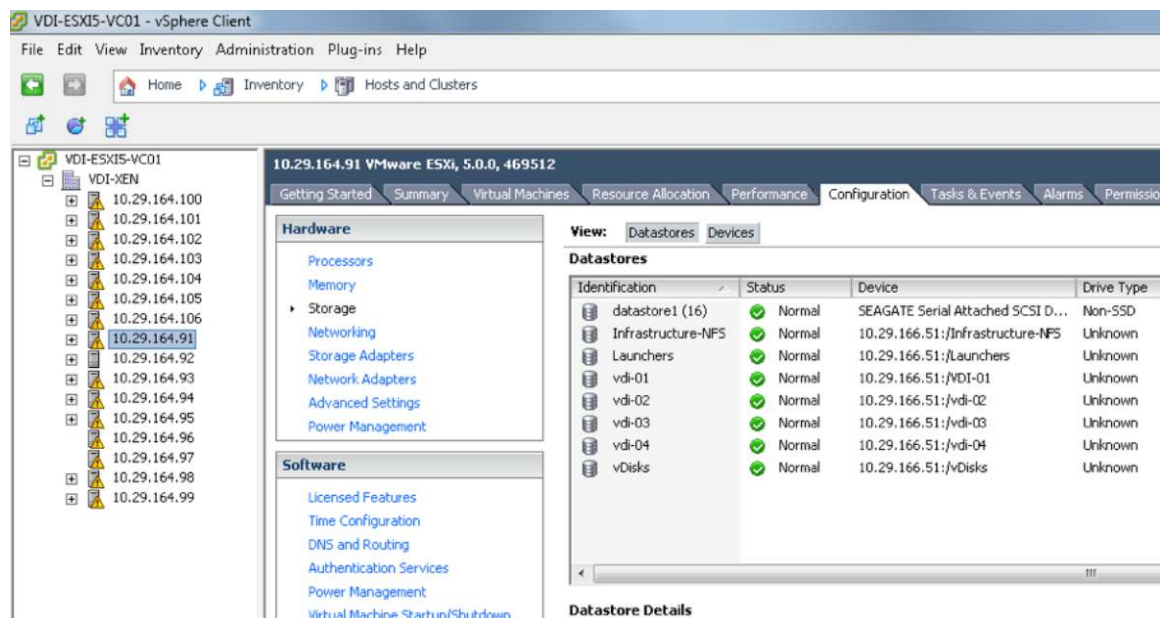
Storage Configuration for VMware ESXi 5 Server Hosting Citrix Provisioning Server and Virtual Desktop vDisks

The test environment used a single storage system to provide both boot-from-SAN LUNs for the Cisco UCS B230 M2 blades hosting the Microsoft Windows 7 SP1 virtual machines and NFS volumes that were used for:

- Infrastructure virtual machine hard drives (one NFS volume)
- Write-cache drives provisioned with each Microsoft Windows 7 SP1 virtual machine (four NFS volumes)
- Microsoft Windows 7 SP1 vDisks storage accessed by the two of the provisioning servers (one NFS volume)
- Launcher vDisks accessed by a dedicated Citrix Provisioning Service launcher (one NFS volume)

Figure 29 shows the configuration.

Figure 29. Data Store Configuration on VMware ESXi 5



Citrix Provisioning Server 5.6 SP1 for Standard-Mode Desktops

The windows desktop image is converted to a vDisk (.vhd) image. The vDisk is then locked in a shared (read-only) mode and hosted on the Citrix Provisioning Server's local disk or in a shared file location.

- Citrix Provisioning Server is used to create the desired number of virtual machines based on parameters specified during setup in the setup wizard, and the virtual machines are associated with a machine template you create in advance on a hypervisor or in a cluster and a vDisk. Citrix Provisioning Server creates machine accounts in Active Directory as it creates the virtual machines.
- Virtual desktops are then configured to boot using PXE on the hypervisor server.
- Citrix Provisioning Server streams the vDisk image upon startup to the hypervisor, and the image is loaded into RAM.
- Citrix Provisioning Server injects a security identifier (SID) and hostname associated with the virtual machine as each desktop boots to maintain uniqueness in Active Directory. These object mappings are maintained and managed in the Citrix Provisioning Server and are visible in the Citrix Provisioning Services Console in the Collections view.
- Each virtual desktop is assigned a write cache (temporary file) in which any delta changes (write operations) to the default image are recorded and used by the virtual Microsoft Windows operating system throughout its lifecycle. The write cache is written to a dedicated 3-GB hard drive created by Citrix Provisioning Server with each new virtual desktop when it is properly configured.

Note: For best performance, a copy of the vDisk was hosted and maintained on each Citrix Provisioning Server's local disk to provide high availability and load balancing by all servers in the farm. The drive assigned by the hypervisor to each Citrix Provisioning Server for vDisk storage was on a dedicated NFS mount. Citrix

Provisioning Servers were assigned 8 GB of RAM to help make the images persistent and capable of being serviced by RAM after they are read for the first time by each server.

Two Citrix Provisioning Servers were configured in a farm to provide high availability and resilience for virtual desktop provisioning. If a failure occurs, connections automatically fail over to working servers in the farm without interruption of the desktop.

A separate Citrix Provisioning Server and NFS export were used to provision Login VSI launcher machines for workload testing.

Four NFS exports on the storage system were used to create and store each virtual machine's write-cache drive.

The write cache is a temporary file in which any delta changes (write operations) to the default image are recorded and used by the virtual Microsoft Windows operating system throughout its lifecycle. You should carefully consider where the write cache is placed when scaling virtual desktops using Citrix Provisioning Server. There are several options as to where the write cache can be placed:

- Citrix Provisioning Server
- Hypervisor RAM
- Device local disk (an additional virtual disk for VDI instances)

For optimal performance and scalability, the "Cache on devices HD" option is used. A 3-GB virtual disk is assigned to the virtual machine templates used in the clone creation process (described in the section ["Citrix XenDesktop 5.5 Integration"](#)).

The Citrix Provisioning Services target device agent installed on the Microsoft Windows 7 golden image automatically places the Microsoft Windows swap file on the same drive used by the Citrix Provisioning Server write cache when this mode is enabled.

To further increase scalability, load balancing across multiple NFS volumes was achieved by using four virtual machine templates (each created on different data stores or storage repositories) and running the Citrix Provisioning Services Streamed VM Setup Wizard tool four times, using a different virtual machine template for each process.

Figure 29 shows multiple virtual machine instances hosted on a hypervisor server booting from a Citrix Provisioning Server single master image. Each instance has a virtual disk hosted on different NFS mounts on which the Citrix Provisioning Server cache is placed. This approach helps ensure that all write I/O takes place over NFS using high-performance storage.

LAN Configuration


The access-layer LAN configuration consists of a pair of Cisco Nexus 5548UP Switches, one of the low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches used in the VDI deployment.

Cisco UCS Connectivity

Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS 6248UP Fabric Interconnects, and they are connected to the Cisco Nexus 5548UP pair in a PortChannel as shown in Figure 30. The Cisco 6248UP is in end-host mode, because both Fibre Channel and Ethernet NAS data access are deployed following the recommended best practices for Cisco UCS. This configuration was built out for scale, and more than 40 Gbps per fabric interconnect are provisioned.

Note: The upstream configuration is beyond the scope of this document. There are some good reference documents that cover best practices for Cisco Nexus 5000 and 7000 Series Switches. New with the Cisco Nexus 5500 platform is a Layer 3 module that was not used in these tests and that will not be discussed in this document.

Figure 30. Network Configuration with Upstream Cisco Nexus 5000 Series Switches from Cisco UCS



Storage LAN Connectivity

The Cisco Nexus 5548UP is used to connect to the storage system for Fibre Channel and NAS access.

The storage system is equipped with dual-port 8-Gbps Fibre Channel modules on each service processor on the storage system, and these are connected to the pair of Cisco Nexus 5548UP Switches to provide block storage access to the environment.

The storage system supports dual-port 10 Gigabit Ethernet modules on each service processor, and these are configured in a PortChannel and connected to the pair of Cisco Nexus 5000 Series Switches downstream. This configuration allows end-to-end 10-Gbps access to file-based storage. In this test, jumbo frames are implemented on the ports and have priority flow control (PFC) enabled, with platinum CoS and QoS assigned to the vNICs carrying the storage data access on the fabric interconnects.

SAN Configuration

The same pair of Cisco Nexus 5548UP Switches was used in the configuration to connect the Fibre Channel ports on the storage system to the Fibre Channel ports of the Cisco UCS 6248UP Fabric Interconnects' expansion modules. The Fibre Channel connection was exclusively used for configuring boot-from-SAN for the VMware ESXi 5 server blades. For more information, see the following section, "Boot from SAN."

Single virtual SAN (vSAN) zoning was set up on the Cisco Nexus 5548UP Switches to make these storage system LUNs visible to the infrastructure and test servers.

Boot from SAN Configuration

Booting from the SAN is another feature that facilitates the move toward stateless computing, in which there is no static binding between a physical server and the OS and applications that 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 WWPN of the HBAs and the boot-from-SAN policy would move along with it. The new server now takes the same exact view of the old server, demonstrating the true stateless nature of the blade server.

The main benefits of booting from the network include:

- Smaller server footprint: 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.
- Simplified 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 the functionality of the servers at the primary site, the remote site can take over with little downtime. Recovery from server failure 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 the server image. As a result, boot-from-SAN capability can greatly reduce the time required for server recovery.
- High availability: A typical data center is highly redundant: redundant paths, redundant disks, and redundant storage controllers. Storage of operating system images on the SAN 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 need to be in production for only hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server use 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.

Configuring Boot from SAN on the Storage System

With boot from SAN, the image resides on a SAN LUN, and the server communicates with the SAN through an HBA. The HBA BIOS contains the instructions that enable the server to find the boot disk. All Fibre Channel–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. After the hardware detects the boot device, it follows the regular boot process.

Note: The two SAN fabrics are separate from a data perspective, and the dual-port HBAs and storage controller provide redundancy.

The boot-from-SAN process is divided into three main steps:

- Storage array configuration
- SAN zone configuration
- Cisco UCS service profile configuration

Storage Array Configuration

The storage array administrator needs to provision LUNs of the required size for installing the OS and to enable boot from SAN. The boot-from-SAN LUN is usually LUN 0. The SAN administrator also needs to identify the WWPN of the adapter to implement the necessary LUN masking. The LUN masking is also a critical step in the SAN LUN configuration.

SAN Configuration on Cisco Nexus 5500 Platform Switches

You must turn on FCoE and the N-port ID virtualization (NPIV) feature on the Cisco Nexus 5500 platform switch. Make sure you have 8-GB Enhanced Small Form-Factor Pluggable (SFP+) modules connected to the Cisco UCS 6200 Series Fabric Interconnect. Be sure that the port mode and the speed are set to Auto and that the rate mode is set to Dedicated. When everything is configured correctly, you should see something like the output shown here on a Cisco Nexus 5500 platform switch for a given port (for example, Fc1/17).

Note: A Cisco Nexus 5500 platform switch supports a single VSAN configuration. If more VSANs are required, you should use a SAN switch such as the Cisco MDS 9100 Series Multilayer Fabric Switches.

Cisco Fabric Manager can also be used to get an overall view of the SAN configuration and zoning information. As discussed earlier, SAN zoning is performed upfront for all WWPNs of the initiators with the storage system target WWPN.

```
VDI-N5548-A# show feature | grep npiv
npiv                1                enabled
VDI-N5548-A# show interface brief
-----
Interface  Vsan   Admin  Admin  Status      SFP  Oper  Oper  Port
Speed  Channel                                Mode  Trunk                                Mode
(Gbps)
-----
fc1/17     1      auto   on     up           swl   F     8     --
fc1/18     1      auto   on     up           swl   F     8     --

VDI-N5548-A# show interface fc1/17 brief
-----
Interface  Vsan   Admin  Admin  Status      SFP  Oper  Oper  Port
Mode  Speed  Channel                                Mode  Trunk                                Mode
(Gbps)
-----
fc1/17     1      auto   on     up           swl   F     8     --
```

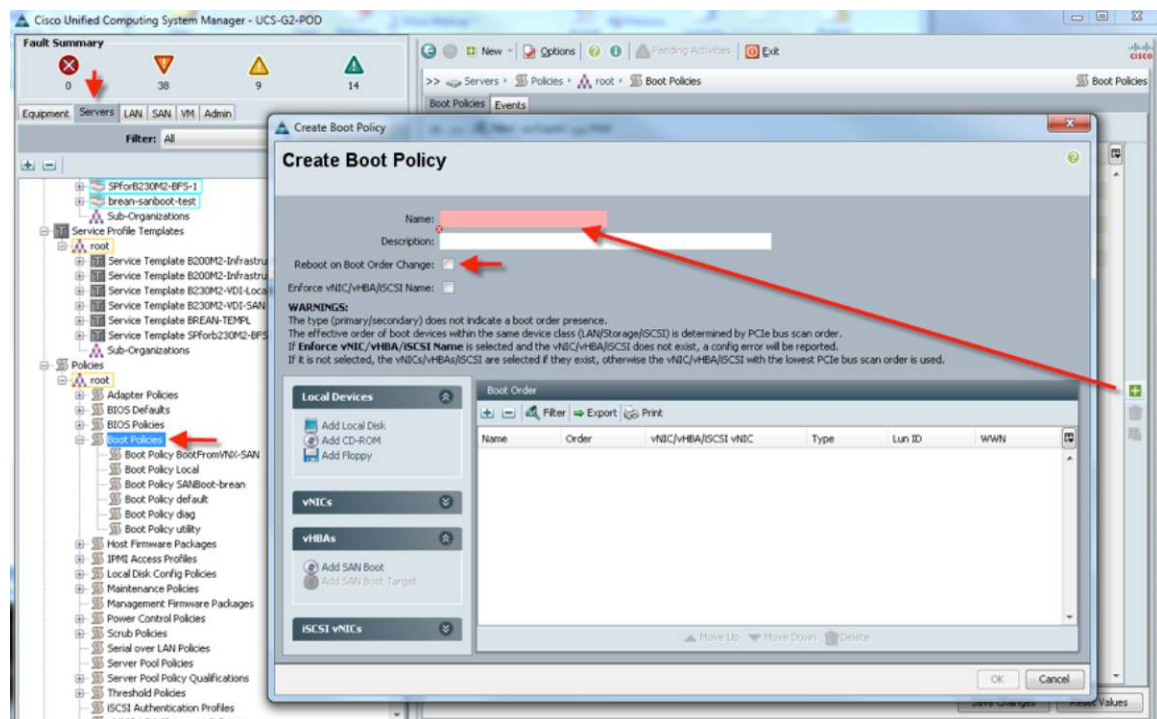

For detailed information about Cisco Nexus 5500 platform switch configuration, refer to the Cisco Nexus 5500 platform and Cisco NX-OS SAN switching configuration guide (see the [“For More Information”](#) section of this document).

Cisco UCS Manager Configuration for Boot from SAN

To enable boot from SAN on Cisco UCS Manager 2.0, follow these steps:

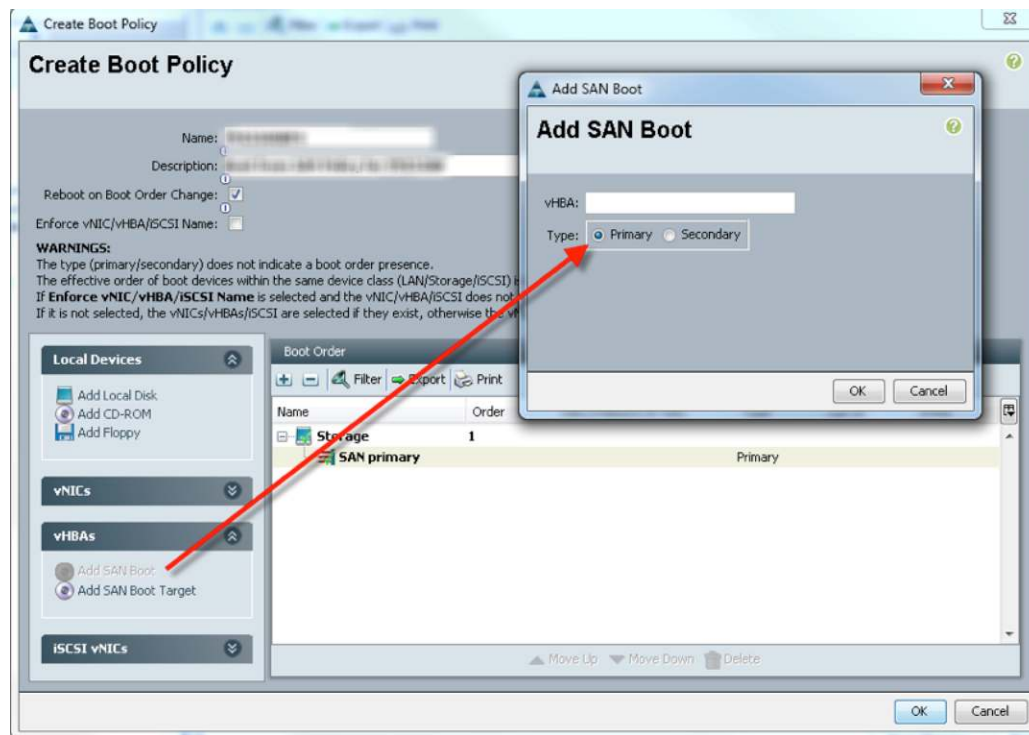
- Step 1. Create a boot policy on the Servers tab. Select the policies and select Boot Policies; in the right pane, click Add (+) button. Enter a name and select Reboot on Boot Order Change. Be sure that Enforce vNIC/vHBA/iSCSI Name is not selected (Figure 31).

Figure 31. Creating Boot Policy



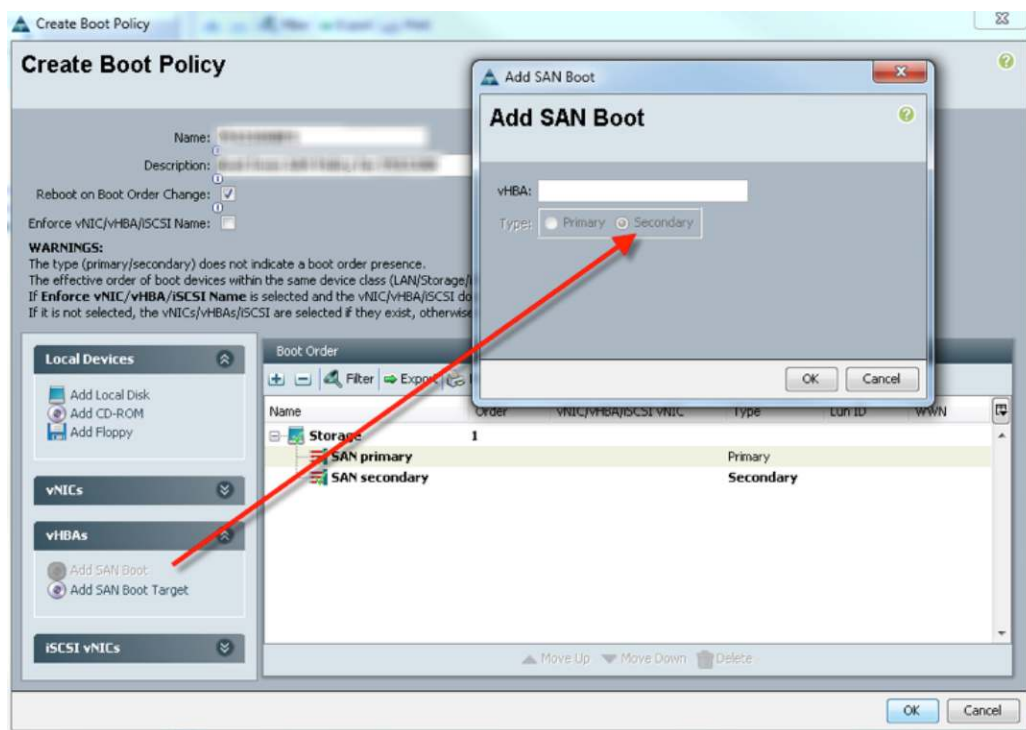
- Step 2. Add SAN boot for the SAN primary to the new policy. The vHBA name is optional; this name does not need to be enforced. Click OK (Figure 32).

Figure 32. Adding SAN Boot for SAN Primary to New Policy



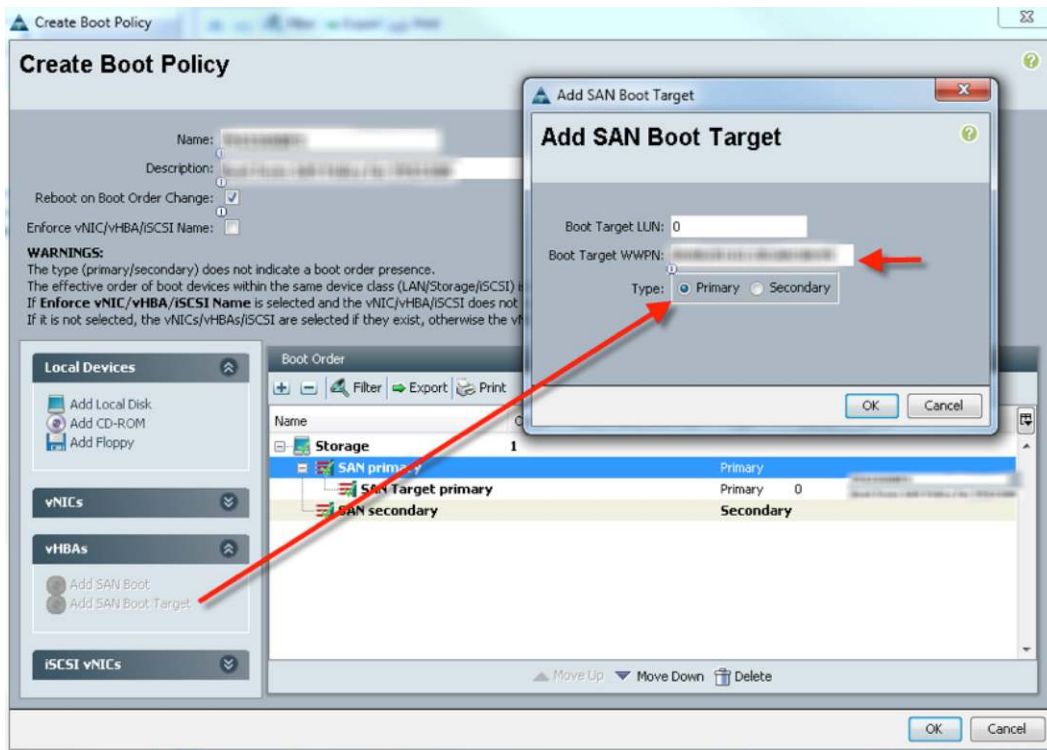
Step 3. Add SAN boot for the SAN secondary. Click OK. Again, the vHBA name is optional, and this field can be left blank (Figure 33).

Figure 33. Adding SAN Boot for SAN Secondary



Step 4. Add the boot target WWPN to the SAN primary. Make sure that this name is exactly the same as the storage system WWPN (Figure 34). To avoid any typos, copy and paste the Cisco Nexus 5500 platform command as shown here for each switch:

Figure 34. Creating Boot Policy



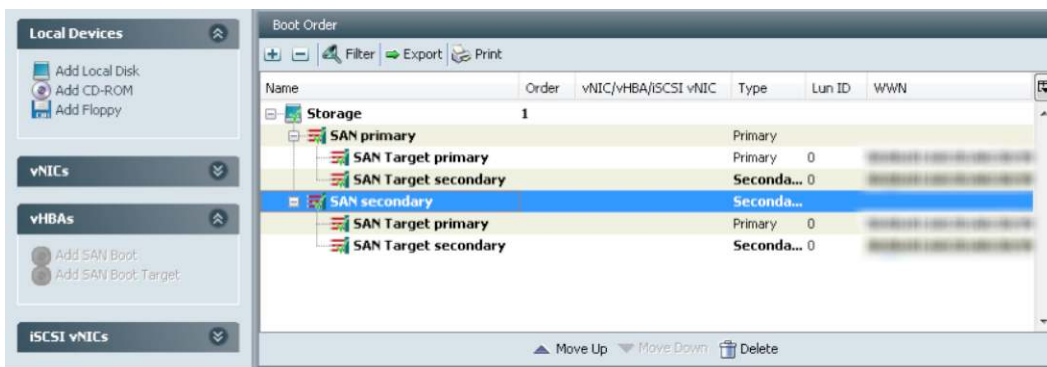
Step 5. Repeat step 4 for the SAN primary's SAN target secondary.

Step 6. Repeat step 4 for the SAN secondary's SAN target Primary.

Step 7. Repeat step 4 for the SAN secondary's SAN target Secondary.

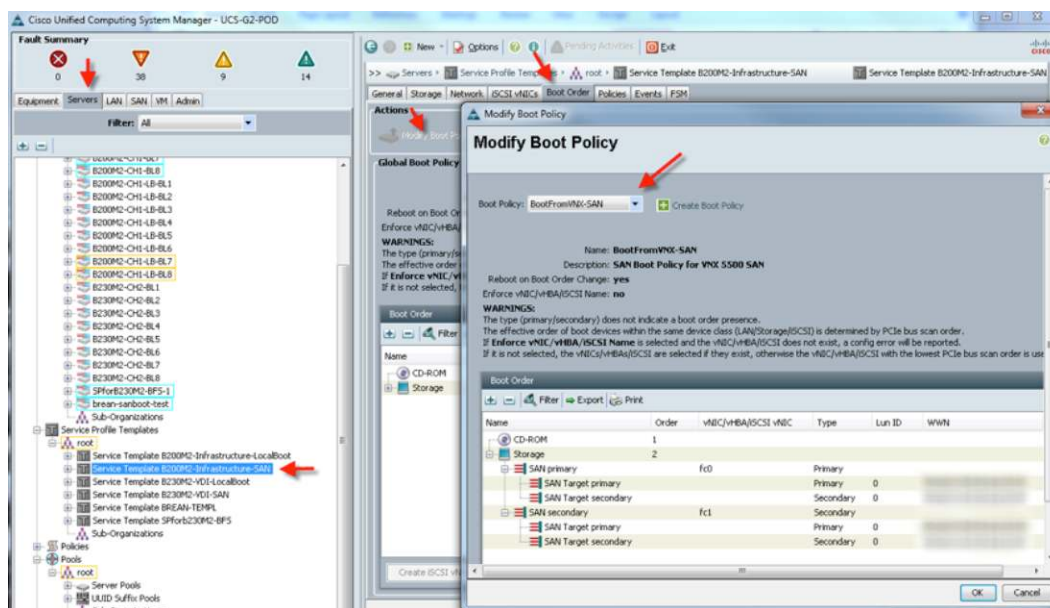
When you are finished, your boot-from-SAN policy should look like Figure 35.

Figure 35. Boot-from-SAN Policy



Step 8. Associate the service profile template with the boot-from-SAN policy during the service profile template configuration. You can also modify the boot policy in use in a service profile template as shown in Figure 36.

Figure 36. Modifying a Service Profile Template Boot Policy



This completes the boot-from-SAN configuration on Cisco UCS Manager. When the service profile is created from the template, each server will be ready to boot from the SAN assuming that the appropriate OS installation steps have been implemented.

Hypervisor Operating System Installation and Configuration

This section describes VMware ESXi 5 installation and configuration and any tuning that was performed in the testing. Tables 10, 11, and 12 provide details about the software components.

Table 10. Cisco UCS B230 M2 Blade and Microsoft Windows 7 SP1 Virtual Machine Test Blade

VMware vSphere ESXi 5 Host Build 469512			
Hardware	Cisco UCS B-Series Blade Servers	Model	Cisco UCS B230 M2
OS	VMware ESXi 5.0.0	Service Pack	—
CPU	2 x 10-core Intel Xeon processors E7-2870, 2.4 GHz (40 logical cores total)	RAM	256 GB
Disk	Boot from SAN NAS (NFS) volumes for cache and vDisk	Network	4 x 10 Gigabit Ethernet (2 x 10 Gigabit Ethernet storage)

Table 11. Cisco UCS B200 M2 Infrastructure Blade

VMware vSphere ESXi 5 Host Build 469512			
Hardware	Cisco UCS B-Series Blade Servers	Model	Cisco UCS B200 M2
OS	VMware ESXi 5.0.0	Service Pack	—
CPU	2 x 6-core Intel Xeon processor 5690, 3.466 GHz (24 logical cores total)	RAM	96 GB (8 GB, 1333 MHz)
Disk	Local disk installation	Network	4 x 10 Gigabit Ethernet (2 x 10 Gigabit Ethernet storage)

Table 12. Cisco UCS B200 M2 Infrastructure Blade

VMware vSphere vCenter 5.0.0.16608			
Hardware	Running in a virtual environment on top of Cisco UCS B200 M2	Model	Cisco UCS B200 M2
OS	Microsoft Windows Server 2008 R2 Enterprise Edition	Service Pack	SP1
CPU	4 vCPUs	RAM	8 GB
Disk	C:\ (50 GB)	Network	1 x 1 Gigabit Ethernet
Network Driver	E1000 on VMware vCenter Server	VMware Version	8

One goal of the test was to virtualize the entire infrastructure, including VMware vCenter and all the Citrix XenDesktop management services. This goal was accomplished with no bare-metal installation of any operating system. All infrastructure components were run as virtual machines.

Installing the OS

This test used local media using the Cisco UCS Manager KVM to install VMware ESXi 5 on the Cisco UCS blade servers. The UCS Cisco UCS B230 M2 Blade Servers were configured to boot from SAN, and the Cisco UCS B200 M2 infrastructure was installed on the local SAS drives.

Configuring VMware vCenter for Citrix XenDesktop DDC Communications

Citrix XenDesktop 5.5 requires communication with the VMware vCenter SDK over HTTP. You need to update the proxy.xml file (located on the VMware vCenter Server at C:\ProgramData\VMware\VMware VirtualCenter\proxy.xml) to allow the SDK over HTTP. Set the access mode for /sdk to **httpAndHttps**.

```
<config>
  <EndpointList>
    <_length>15</_length>
    <_type>vim.ProxyService.EndpointSpec[ ]</_type>
    <e id="0">
      <_type>vim.ProxyService.NamedPipeServiceSpec</_type>
      <accessMode>httpAndHttps</accessMode>
      <pipeName>\\.\pipe\vmware-vpxd-webserver-pipe</pipeName>
      <serverNamespace>/</serverNamespace>
    </e>
    .
    <e id="5">
      <_type>vim.ProxyService.LocalServiceSpec</_type>
      <accessMode>httpAndHttps</accessMode>
      <port>8085</port>
      <serverNamespace>/sdk</serverNamespace>
    </e>
  </EndpointList>
</config>
```

Restart VMware vCenter Service

Restart the VMware vCenter service from the command line or the services.msc console.


Note: Citrix recommends that you either import the default certificate from the VMware vCenter Server on your desktop controllers or upgrade the default certificates with trusted third-party certificates, which should be installed on the DDCs.

See the article at <http://support.citrix.com/article/CTX125578>.

Setting Up the Network Time Protocol Server in VMware vCenter

One important aspect of running a benchmark in a virtualized environment is configuring and setting up a Network Time Protocol (NTP) server and configuring it from VMware vCenter for each server. This setting helps ensure synchronization of performance data collected across various components (Figure 37).

Figure 37. NTP Server Setup



Tuning VMware ESXi

Apart from the Citrix and storage system best practices for using NFS storage on VMware vSphere for the virtual machine write cache and to host virtual disks used to provision Microsoft Windows 7 SP1 virtual machines and the Login Consultants' launcher virtual machines, no other tuning was performed on the VMware ESXi servers.

The virtual NIC (enic) and virtual HBA (fnic) drivers were updated to the following versions respectively:

- enic_driver_2.1.2.22-564611
- fnic_driver_1.5.0.7-563432

These drivers support the Cisco UCS M81KR, which is installed on each blade. Updated drivers for Cisco UCS components are available at

<http://www.cisco.com/cisco/software/release.html?mdfid=283853163&softwareid=283853158&release=2.0%281f%29&reind=AVAILABLE&rellifecycle=&reltype=latest>.

Configuring VMware vSphere

For single-server scalability testing, one VMware ESXi 5.0.0 server was configured with boot from SAN, and the driver updates specified in the preceding section were applied. One NFS storage volume was configured for this test and was used for Microsoft Windows 7 SP1 virtual machine write-cache volumes.

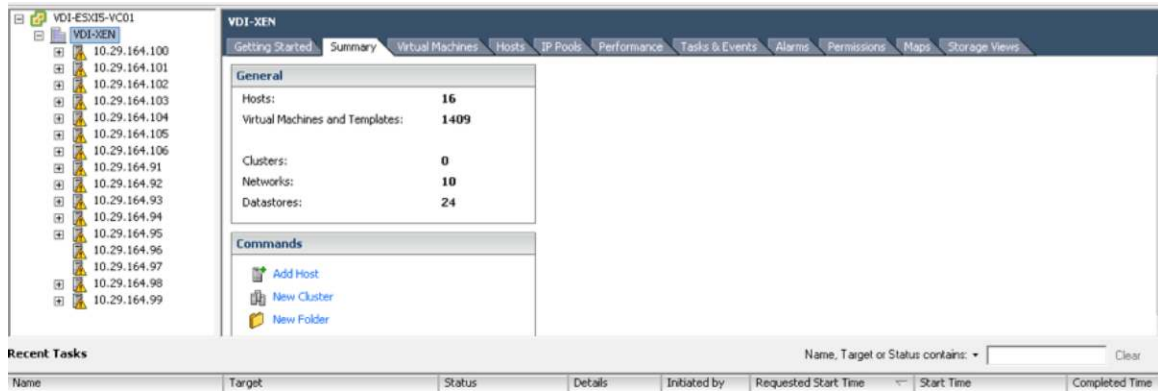
For eight-server scalability testing, eight VMware ESXi 5.0 servers were configured with boot from SAN, and the driver updates specified in the preceding section were applied. Four NFS storage volumes were configured for this test and were used for Microsoft Windows 7 SP1 virtual machine write-cache volumes. Two blade servers shared each NFS volume.

For the single-chassis testing, a VMware cluster was not created. A disaster recovery system therefore was not available by design. Four NAS devices were mounted on the eight servers as NFS mounts, and the launcher virtual machines were used to generate the load to the desktops on the eight servers. Each blade was assigned

135 desktops, for a total of 1080 desktops; 270 desktops were distributed on each of the four volumes deployed for that purpose.

Figure 38 shows the VMware data center under test.

Figure 38. VMware Data Center Under Test



Two Citrix Provisioning Server 5.6 SP1 virtual machines share the fifth NFS volume dedicated to hosting the standard-image virtual disks (vDisks) that were used to provision and launch the Microsoft Windows 7 SP1 virtual desktops.

One additional Citrix Provisioning Server 5.6 SP1 virtual machine used a sixth NFS volume dedicated to hosting the standard-image vDisks that were used to provision and launch the Login Consultants' launcher machines for workload testing.

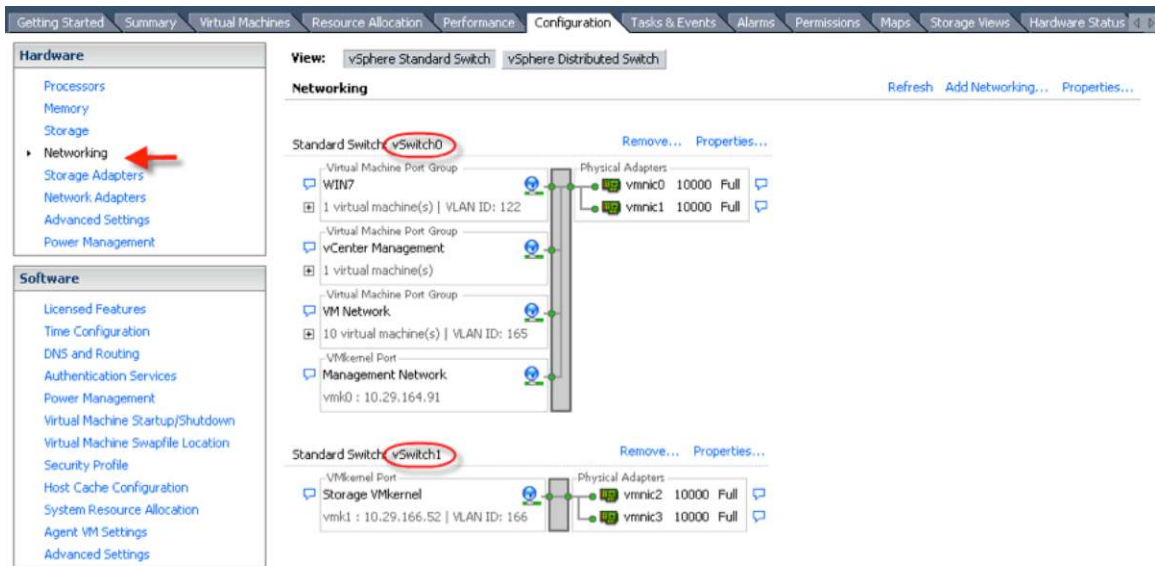
Configuring the Network

For the Cisco UCS M81KR (Palo), four static vNICs and two Fibre Channel vHBAs were created per blade.

Different VLANs were used to separate the desktop traffic from the management traffic and the storage traffic. Simple VMware vSwitch configurations for four vNICs, two per vSwitch, are performed with two active virtual machine NICs per switch at the vSwitch level.

All the infrastructure (Cisco UCS B200 M2 servers) and tests (Cisco UCS B230 M2 servers) running VMware ESXi 5, were configured as shown in Figure 39.

Figure 39. vSwitch Configuration



VLAN 164 was configured as the native VLAN on each host using Cisco UCS Manager. For vSwitch0 (Figure 40), the vmkernel Management Network vNIC settings for this VLAN were set to active and standby (vmnic0 and vmnic1) for VMware vMotion and vmkernel traffic (Figure 41).

Figure 40. VMware vSwitch Management Network Edit

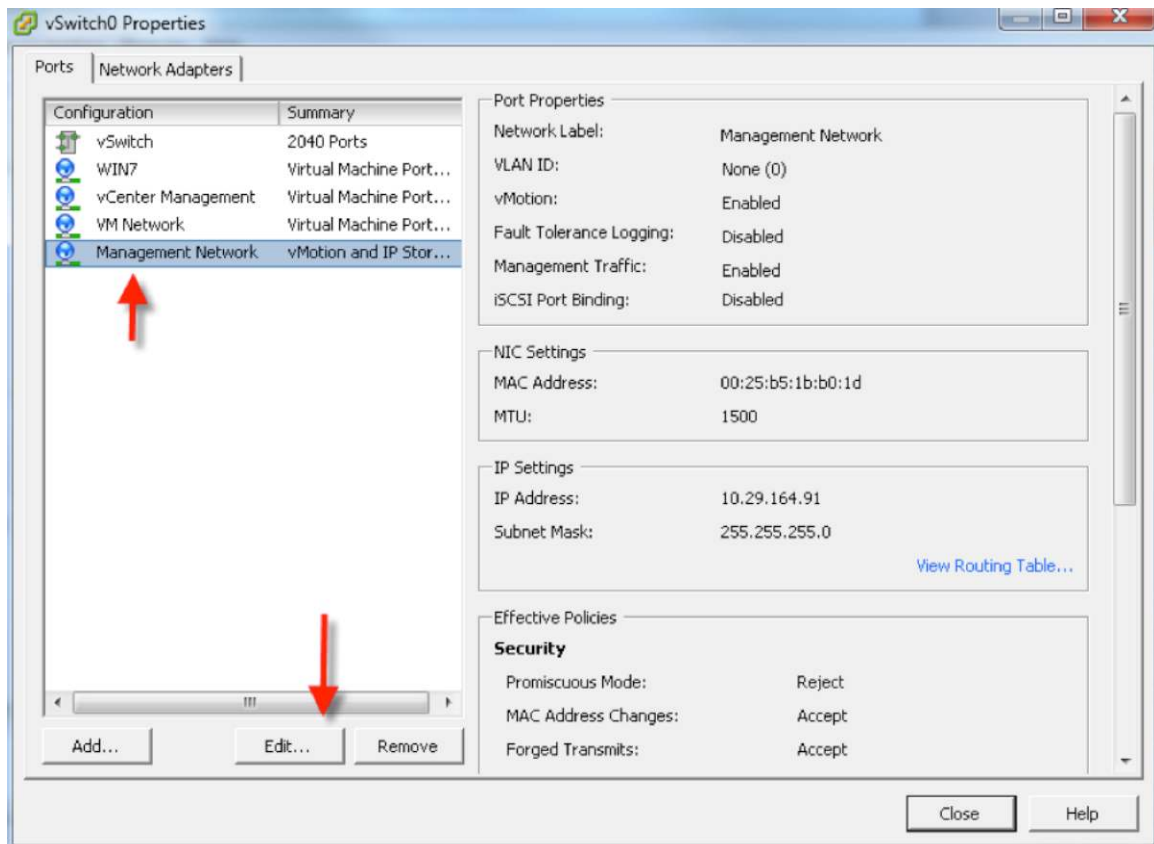
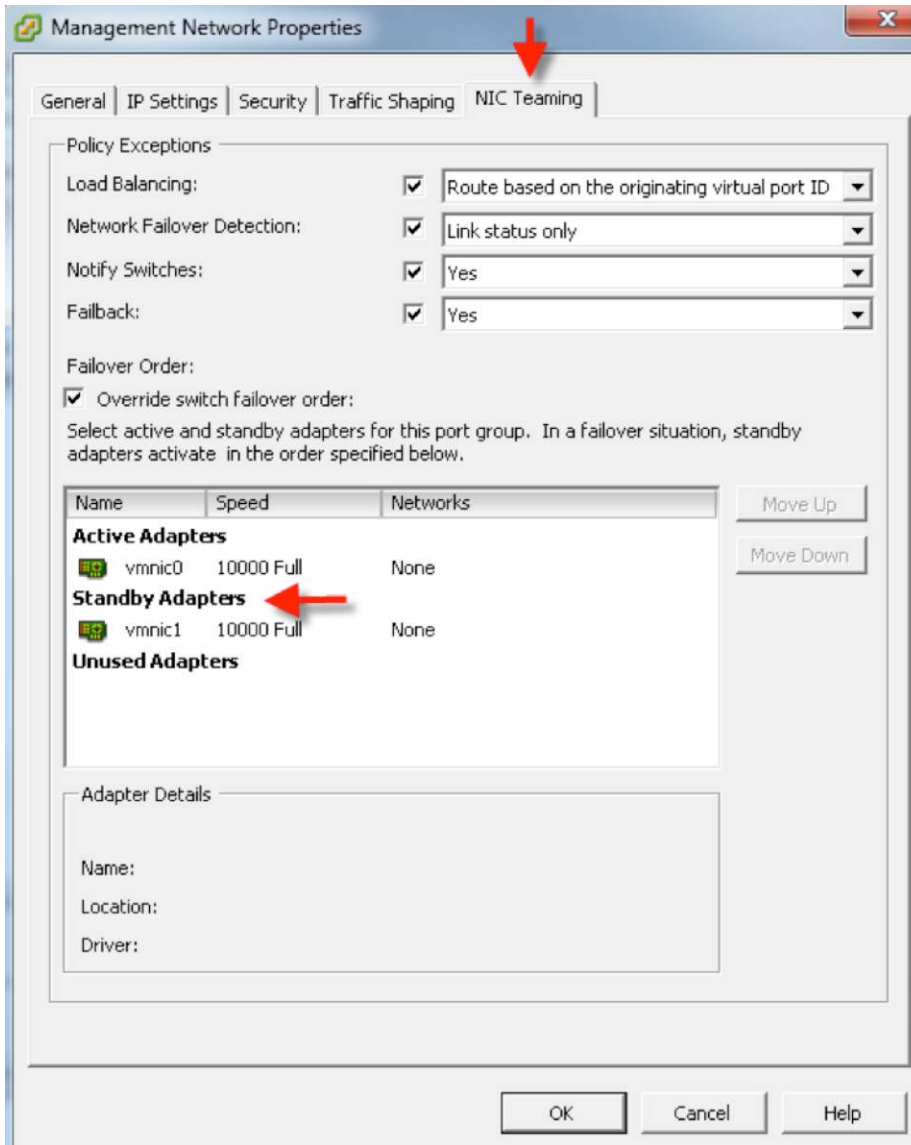



Figure 41. Vmnic1 Set as Standby Adapter for VMware vMotion and vmkernel Network Only on vSwitch0



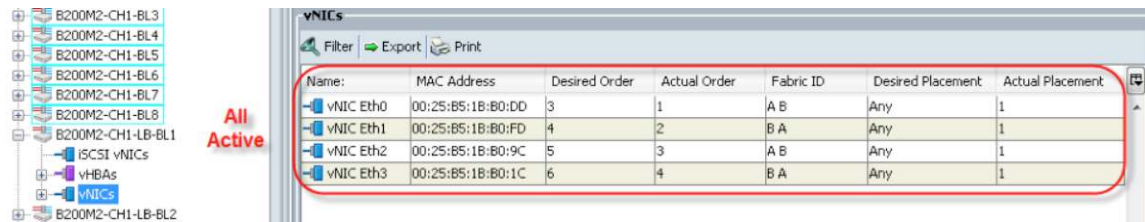
The VMware ESXi 5 console (Figure 42) shows the two virtual machine NICs, vmnic0 and vmnic1, used for management network communications, with their MAC addresses.

Figure 42. VMware ESXi 5 Management vNICs



The MAC addresses shown in Cisco UCS Manager (Figure 43) map to either Fabric A or Fabric B.

Figure 43. Active vNICs



Name	MAC Address	Desired Order	Actual Order	Fabric ID	Desired Placement	Actual Placement
vNIC Eth0	00:25:B5:1B:B0:DD	3	1	A B	Any	1
vNIC Eth1	00:25:B5:1B:B0:FD	4	2	B A	Any	1
vNIC Eth2	00:25:B5:1B:B0:9C	5	3	A B	Any	1
vNIC Eth3	00:25:B5:1B:B0:1C	6	4	B A	Any	1

The teaming is performed so that the vNICs on Fabric A and Fabric B are all active in Cisco UCS Manager and on the VMware vSwitches at the switch level.

Cisco recommends using Cisco Nexus 1000V Series Switches for network configuration because they provide advanced capabilities for DHCP snooping along with other smart network switch capabilities in the hypervisor itself. These features provide a substantial benefit in a virtual desktop environment in which a vSwitch would not be able to provide such features.

Note: Keep in mind when mounting the NFS data store that NFS data stores require a vmkernel port that can connect to the storage NFS server. Cisco uses a different vSwitch and VLAN to separate the traffic from other desktop and management network traffic and to provide the necessary vmkernel port.

After the storage is configured with the NFS volumes and the volumes are exported to all hosts, you need to add the NFS data store to the appropriate ESXi servers. This is easily accomplished using the plug-in for VMware vCenter. Right-click the cluster or VMware vCenter to add NFS exports to all its members, or select an individual VMware ESXi host in the left pane and then select the storage type: choose Unified Storage > Provision Storage > and follow the wizard, specifying Network File System as the storage type.

With the storage plug-in for VMware vCenter, you can mount existing NFS file exports or create a new export from within VMware vCenter. After the wizard completes, the NFS export is mounted by all the specified hosts in one simple action (Figure 44).

Figure 44. NFS Export Mounted on VMware ESXi 5 Host



Identification	Status	Device	Drive Type	Capacity	Free	Type	L
datastore1 (7)	Normal	DGC Fibre Channel...	Non-SSD	5.00 GB	4.29 GB	VMFS5	2
vd1-01	Normal	10.29.166.51:/VD...	Unknown	1,008.38 G	751.77 GB	NFS	2

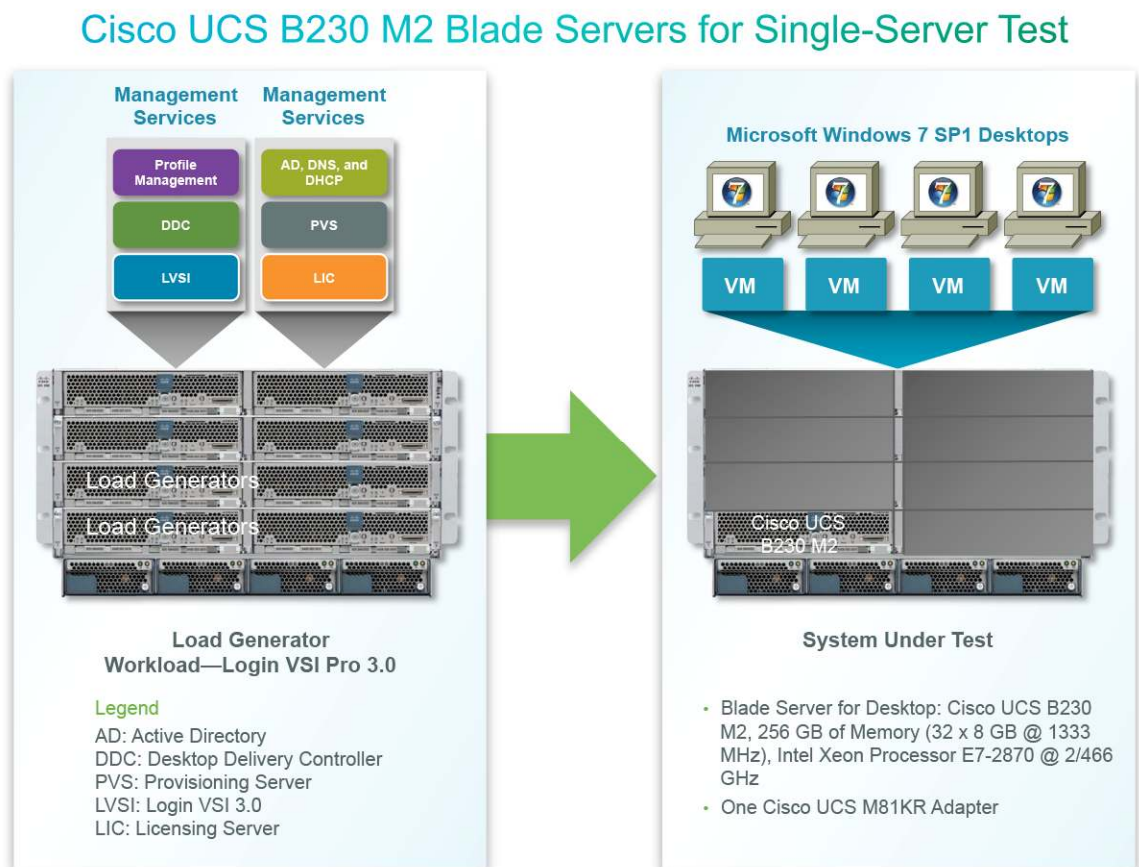
Test Setup and Configuration

A single Cisco UCS B230 M2 Blade Server and eight Cisco UCS B230 M2 Blade Servers were tested in a single chassis to illustrate linear scalability.

Cisco UCS Test Configuration for Single-Blade Scalability

The test configuration for single-blade scalability is shown in Figure 45.

Figure 45. Cisco UCS B200 M2 Infrastructure Services and Cisco UCS B230 M2 Blade Server with Microsoft Windows 7 Virtual Machine Host



Hardware components included:

- 1 Cisco UCS B230 M2 (Intel Xeon processor E7-2870 at 2.4 GHz) blade server with 256 GB of memory (8 GB x 32 DIMMs at 1333 MHz) Microsoft Windows 7 SP1 virtual desktops
- 2 Cisco UCS B200 M2 (Intel Xeon processor X5690 at 3.466 GHz) blade servers with 96 GB of memory (8 GB x 12 DIMMs at 1333 MHz) infrastructure servers
- 2 Cisco UCS B200 M2 (Intel Xeon processor X5690 at 3.466 GHz) blade servers with 96 GB of memory (8 GB x 12 DIMMs at 1333 MHz) load generators
- 1 Cisco UCS M81KR VIC per blade
- 2 Cisco UCS 6248UP Fabric Interconnects
- 2 Cisco Nexus 5548UP access switches
- 1 Storage Array, 2 Controllers, 2 dual-port 8-Gbps Fibre Channel cards, and 2 dual-port 10 Gigabit Ethernet cards with at least 26 600GB 15,000-rpm SAS drives and 2 100 GB SSDs for cache

Software components included:

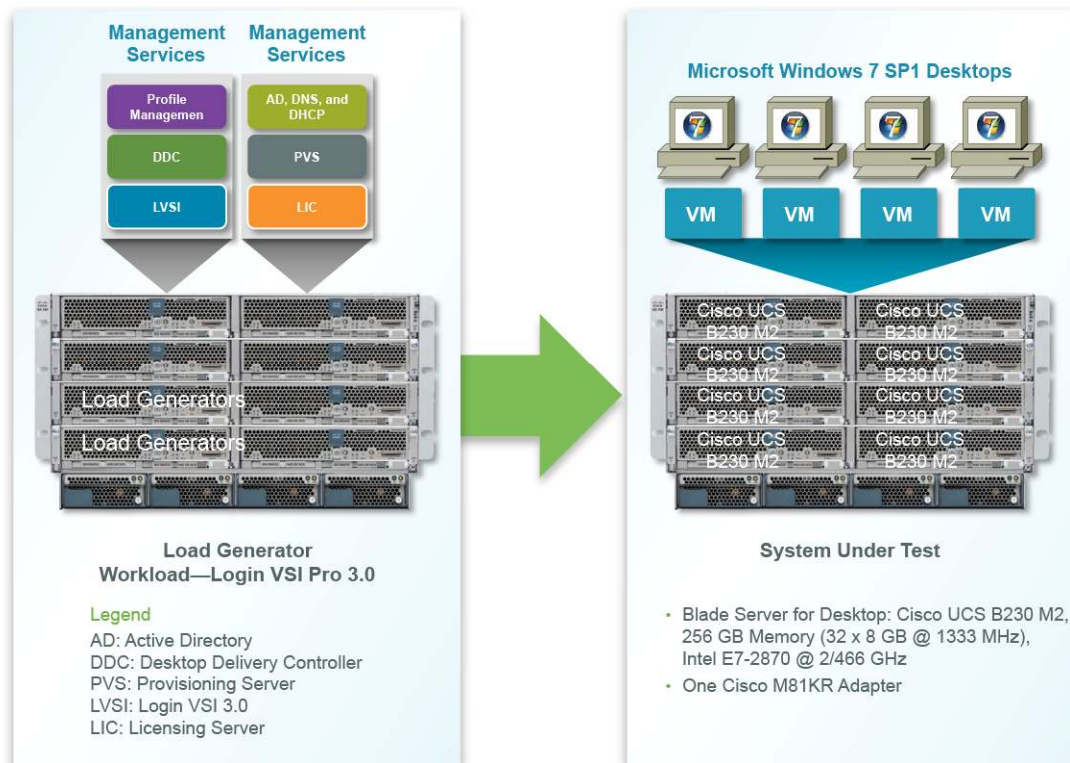
- Cisco UCS Firmware Release 2.0(1t)
- VMware ESXi 5.0.0 Build 469512 and VMware vCenter 5.0
- Citrix XenDesktop 5.5
- Citrix Provisioning Server 5.6 SP1
- Microsoft Windows 7 SP1 32-bit, 1 vCPU, 1.5 GB of memory, and 20 GB per virtual machine

Cisco UCS Configuration for Single-Chassis Test

The configuration for a single server chassis in the test environment is shown in Figure 46.

Figure 46. Single-Chassis Test Configuration: Eight Cisco UCS B230 M2 Blade Servers

Full Chassis Test Configuration: 8 Cisco UCS B230 M2 Blade Servers



Hardware components included:

- 8 Cisco UCS B230 M2 (Intel Xeon processor E7-2870 at 2.4 GHz) blade servers with 256 GB of memory (8 GB x 32 DIMMs at 1333 MHz) Microsoft Windows 7 SP1 virtual desktops
- 2 Cisco UCS B200 M2 (Intel Xeon processor X5690 at 3.466 GHz) blade servers with 96 GB of memory (8 GB x 12 DIMMs at 1333 MHz) infrastructure servers
- 2 Cisco UCS B200 M2 (Intel Xeon processor X5690 at 3.466 GHz) blade servers with 96 GB of memory (8 GB x 12 DIMMs at 1333 MHz) load generators

- 1 Cisco UCS M81KR VIC per blade
- 2 Cisco 6248UP Fabric Interconnects
- 2 Cisco Nexus 5548UP access switches
- 1 storage array, two service processors, 2 dual-port 8-GB Fibre Channel cards, and 2 dual-port 10 Gigabit Ethernet cards with 26 600-GB 10,000-rpm SAS drives and 3 100 GB SSDs for cache

Software components included:

- Cisco UCS Firmware Release 2.0(1t)
- VMware ESXi 5.0.0 Build 469512 and VMware vCenter 5.0
- Citrix XenDesktop 5.5
- Citrix Provisioning Server 5.6 SP1
- Microsoft Windows 7 SP1 32-bit, 1 vCPU, 1.5 GB of memory, and 20 GB per virtual machine

Testing Methodology and Success Criteria

All validation testing was conducted on site within the Cisco labs with joint support from both Citrix and storage resources. The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop bootup, user logon, virtual desktop acquisition (also referred to as ramp up), user workload processing (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 to have passed unless all the planned test users completed the ramp-up and steady state phases (described in the following sections) and unless all metrics were within the permissible thresholds and noted as success criteria. Three completed test cycles were conducted for each hardware configuration, and results were found to be relatively consistent from one test to the next.

Load Generation

Within each test environment, load generators were used to simulate multiple users accessing the Citrix XenDesktop 5.5 environment and running a typical end-user workflow. To generate load within the environment, an auxiliary software application was used to generate the end-user connection to the Citrix XenDesktop environment, to provide unique user credentials, to initiate the workload, and to evaluate the end-user experience.

In the hosted VDI test environment, session launchers were used to simulate multiple users making direct connections to Citrix XenDesktop 5.5 through a Citrix ICA protocol connection.

User Workload Simulation: Login VSI from Login Consultants

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

The tool measures the in-session response time, providing an objective means of measuring the expected user experience for individual desktops throughout large-scale use, including login storms.

The Virtual Session Indexer (Login Consultants Login VSI 3.0) methodology, designed for benchmarking server-based computing (SBC) and VDI environments, is completely platform and protocol independent and hence allows customers to easily replicate the test results in their environments.

Note: This test used the tool to benchmark the VDI environment only.

Login VSI calculates an index based on the number of simultaneous sessions that can be run on a single machine. It simulates a user with a medium workload (also known as a knowledge worker) running general applications such as Microsoft Office 2007 or 2010, Microsoft Internet Explorer 8 including Adobe Flash applets, and Adobe Acrobat Reader.

Note: For the purposes of this test, applications were installed locally, not streamed or hosted on Citrix XenApp.

To simulate real users, the scripted Login VSI session leaves multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. MediumNoFlash is a workload based on the medium workload with only the Adobe Flash components disabled.

The overall testing is summarized here:

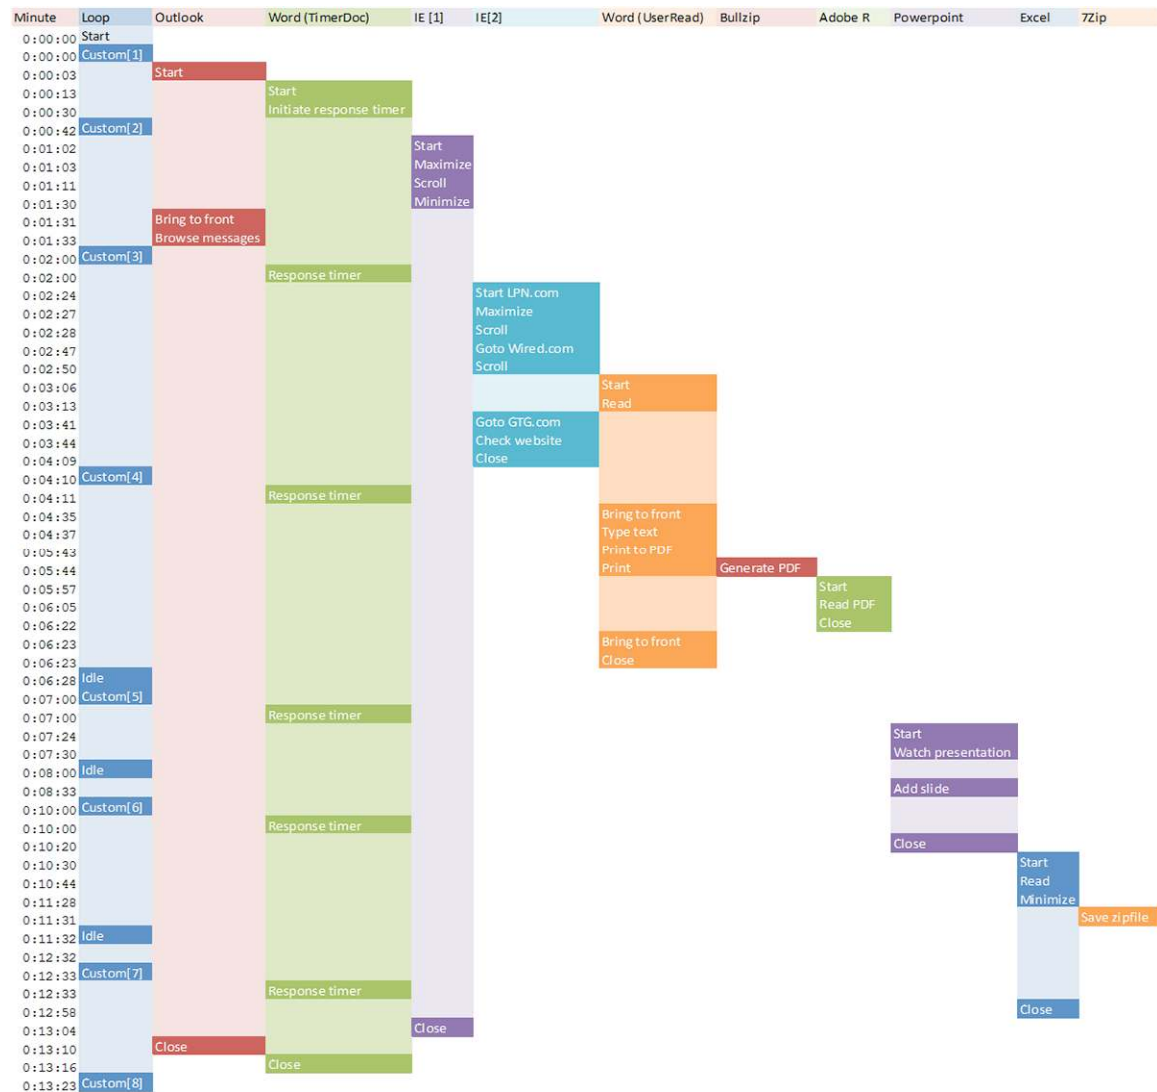
- The workload emulated the medium workload of a knowledge worker using Microsoft Office and Internet Explorer and PDF documents.
- After a session was started, the medium workload repeated every 12 minutes.
- During each loop, the response time was measured every two minutes.
- The medium workload opened up to five applications simultaneously.
- The typing rate was 160 milliseconds (ms) for each character.
- Approximately two minutes of idle time was included to simulate real-world users.

Each loop opened and used the following applications and processes:

- Microsoft Outlook 2007 or 2010: 10 messages were browsed.
- Microsoft Internet Explorer: One instance was left open (BBC.co.uk), and one instance was browsed to Wired.com, and Lonelyplanet.com.
- 480p Adobe Flash application gettheglass.com: This application was not used with the MediumNoFlash workload.
- Microsoft Word 2007 or 2010: One instance was used to measure response time, and one instance was used to review and edit a document.
- Bullzip PDF Printer and Adobe Acrobat Reader: The Microsoft Word document was printed and reviewed as a PDF file.
- Microsoft Excel 2007 or 2010: A very large randomized sheet was opened.
- Microsoft PowerPoint 2007 or 2010: A presentation was reviewed and edited.
- 7-zip: Using the command-line version, the output of the session was zipped.

Figure 47 shows a graphical representation of the medium workload.

Figure 47. Graphical Overview of Test Workload



Additional information about Login VSI is available at <http://www.loginvsi.com>.

Success Criteria

Multiple metrics were captured during each test run, but the success criteria for considering a single test run as passed or failed was based on one important metric: the Login VSImax value. VSImax evaluates the user response time during increasing user load and assesses the successful start-to-finish processing of all the initiated virtual desktop sessions.

Login VSImax Value

VSImax represents the maximum number of users that the environment can handle before serious performance degradation occurs. It is calculated based on the response times of individual users as indicated during the workload processing. The user response time has a threshold of 4000 ms, and all users' response times are expected to be less than 4000 ms to assume that the user interaction with the virtual desktop is at a functional level. VSImax is reached when the response time reaches or exceeds 4000 ms for six 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 means of comparison that can be aligned with host CPU performance.

Note: In the prior version of Login VSI, the threshold for the response time was 2000 ms. The workloads and the analyses have been upgraded in Login VSI 3 to make the testing better aligned with real-world use. In the medium workload in Login VSI 3.0, a CPU-intensive 480p Adobe Flash movie is incorporated in each test loop. In general, the redesigned workload results in an approximately 20 percent decrease in the number of users passing the test compared to Login VSI 2.0 on the same server and storage hardware.

Calculating the Login VSImax Value

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

The seven operations for which the response times are measured are listed here:

- Copying a new document from the document pool on the home drive: This operation refreshes 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 measures the responsiveness of the operating system and the file system. Microsoft Word is started and loaded into memory, and the new document is automatically loaded into Microsoft Word. Extensive or even saturated disk I/O will affect the File Open dialog box considerably.
- Starting the File Open dialog box: This operation is handled partially by Microsoft Word and mainly by the operating system. The File Open dialog box uses generic subsystems and interface components of the OS. The OS provides the contents of this dialog box.
- Starting Microsoft Notepad: This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe program itself through processing. This operation seems instant from an end user's point of view.
- Starting the Print dialog box: This operation is handled mainly by the OS subsystems, because the Print dialog box is provided by the OS. This dialog box loads the print subsystem and the drivers of the selected printer. As a result, this dialog box also depends on disk performance.
- Starting the Search and Replace dialog box: This operation is handled within the application completely; the presentation of the dialog box is almost instant. Serious bottlenecks at the application level affect the speed of this dialog box.
- Compressing the document into a Zip file with the 7-zip command line: This operation is handled by the command-line version of 7-zip. The compression very briefly spikes CPU and disk I/O.

These measured operations with Login VSI do address considerably different subsystems such as CPU (user and kernel), memory, disk, the OS in general, the application itself, print, graphics device interface, 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 then escalate,

and 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, you can choose between VSImax Classic and VSImax Dynamic results analysis. For these tests, VSImax Classic analysis was used.

VSImax Classic

VSImax Classic is based on the previous versions of Login VSI and is achieved when the average Login VSI response time is higher than a fixed threshold of 4000 ms. This method proves to be reliable when no antivirus or application virtualization is used.

To calculate the response times, the seven activities listed in the previous section were totaled. To balance these measurements, they were weighted before they were summed. Without weighting individual response times before they are totaled, one specific measurement (out of seven) could dominate the results.

Within VSImax Classic, two measurements were weighted before they were added to the total VSImax response time:

- The task “starting Microsoft Word with a document” is divided by two (50 percent).
- The task “starting the Search and Replace dialog box” is multiplied by 5 (500 percent).

A sample of the VSImax Classic response time calculation is displayed in Figure 48.

Figure 48. Sample of VSImax Classic Response Time Calculation

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	50%	700
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	100%	50
Print Dialogue (PRINT)	220	100%	220
Replace Dialogue (FIND)	10	500%	50
Zip documents (ZIP)	130	100%	130
VSImax Classic Response Time			1660

The average VSImax response time was then calculated based on the amount of active Login VSI users logged on to the system. When the average VSImax response times are consistently higher than the default threshold of 4000 ms, VSImax is achieved.

In practice, however, tests have shown a substantial increase in application response time when antivirus software or application virtualization is used. The baseline response time is typically approximately 1400 to 1800 ms without application virtualization or antivirus software. However, when antivirus software or application virtualization is used, the baseline response time varies between 2500 and 3500 ms.

When the baseline response time is already high, the VSImax threshold of 4000 ms is reached too easily. VSImax Classic will report a maximum value long before system resources such as CPU, memory, and disks are actually saturated. It was therefore decided to further optimize the VSImax calculation. Login VSI 3.0 thus introduced VSImax Dynamic, which supports widely varying baseline response times when antivirus software and application virtualization are used.

VSImax Dynamic

Similar to VSImax Classic, 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:

- Copying a new document from the document pool in the home drive: 100 percent
- Opening Microsoft Word with a document: 33.3 percent
- Starting the File Open dialog box: 100 percent
- Starting Microsoft Notepad: 300 percent
- Starting the Print dialog box: 200 percent
- Starting the Search and Replace dialog box: 400 percent
- Compressing the document into a Zip file using the 7-zip command line: 200 percent

A sample of the VSImax Dynamic response time calculation is shown in Figure 49.

Figure 49. Sample of VSImax Dynamic Response Time Calculation

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

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

To determine this dynamic threshold, first the average baseline response time was calculated by averaging the baseline response time of the first 15 Login VSI users on the system. The formula for the dynamic threshold is:

Average baseline response time x 125% + 3000

As a result, when the baseline response time is 1800, the VSImax threshold will now be $1800 \times 125\% + 3000 = 5250$ ms.


Especially when application virtualization is used, the baseline response time can vary widely depending on the vendor and streaming strategy. Therefore, VSImax Dynamic is recommended when making comparisons with application virtualization or antivirus agents. The resulting VSImax Dynamic scores are aligned again with saturation on the CPU, memory, or disk and also when the baseline response times are relatively high.

Determining the VSImax Threshold

The Login VSI analyzer will automatically identify the VSImax value. In the example in Figure 50, the VSImax value is 98. The analyzer automatically identifies “stuck sessions” and corrects the final VSImax score. The figure provides the following information:

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

Figure 50. Sample Login VSI Analyzer Graphic Output



For the tests described here to be considered successful, the total number of users in the test run had to log in, become active, and run at least one test loop and log out automatically without reaching the VSImax threshold.

Test Results

The purpose of this testing is to provide the data needed to validate a Citrix XenDesktop 5.5 hosted VDI FlexCast model and Citrix Provisioning Services 5.6 SP1 using VMware ESXi 5 to virtualize Microsoft Windows 7 SP1 desktops on Cisco UCS B230 M2 blade servers using a dual-controller storage system.

The information contained in this section provides data points that customers can reference in designing their own implementations. These validation results provide an example of what is possible under the specific environment conditions outlined in this document and do not represent a full characterization of Citrix XenDesktop with VMware ESXi 5.

Citrix XenDesktop 5.5 Hosted VDI Standard Mode VDI Test Results

Two test sequences were performed to establish single-blade performance and multiple-blade, linear scalability across the tested system.

Single Cisco UCS B230 M2 Blade Server Single-Chassis Validation

This section details the results from the Citrix XenDesktop hosted VDI single-blade server validation testing.

The primary success criteria to validate the overall success of the test cycle are provided as an output chart from Login Consultants' VSI Analyzer Professional Edition, VSImax Dynamic for the Medium Workload (with Adobe Flash).

Additional graphs detailing the CPU and memory use during peak session load are also presented. Given adequate storage capability, the limiting factor in the testing was CPU utilization.


The single-server graphs shown in this section are representative of a single VMware ESXi 5 host in the larger environment for validation purposes, but these graphs are representative of the behavior for all servers in the respective environment.

Performance information for the relevant infrastructure virtual machines is presented with the tested blade.

Single-Blade Login VSI Medium Workload with Adobe Flash: 135 Users

This section details testing results for the Login VSI Medium Workload (with Adobe Flash) on a single Cisco UCS B230 M2 blade. This test used 15-minute virtual machine bootup and 15-minute Login VSI ramp-up times. This testing demonstrates the capability of the Cisco UCS B230 M2 to start and run high-density workloads twice as fast as the current industry-tested standard (Figure 51).

Figure 51. 135 Desktop Sessions on VMware ESXi 5 Below 4000 ms



The graphs in Figures 52 through 56 detail CPU, memory, disk, and network performance on the single Cisco UCS B230 M2 blades during the bootup and testing phases.

Figure 52. 135-User Single Cisco UCS B230 M2 CPU Utilization Boot Phase: 15 Minutes




Figure 53. 135-User Single Cisco UCS B230 M2 Memory Utilization Boot Phase: 15 Minutes




Figure 54. 135-User Single Cisco UCS B230 M2 Network Utilization Boot Phase: 15 Minutes




Figure 55. 135-User Single Cisco UCS B230 M2 CPU Test Phase: 15-Minute Ramp Up (2 Times Faster)






Figure 56. 135-User Single Cisco UCS B230 M2 Network Rate Test Phase: 15-Minute Ramp Up (2 Times Faster)



Single-Blade Login VSI Medium No Flash Workload: 155 Users

This section details testing results for the Login VSI Medium No Flash workload on a single Cisco UCS B230 M2 blade. The same 15-minute bootup and 15-minute Login VSI ramp-up times were used (Figure 57).

Figure 57. 155 Desktop Sessions on VMware ESXi 5 Below 4000 ms



The graphs in Figures 58 through 60 detail CPU, memory, disk, and network performance on the single Cisco UCS B230 M2 blades during the testing phase of the Login VSI Medium No Flash testing.

Figure 58. 155-User Single Cisco UCS B230 M2 CPU Test Phase: 15-Minute Ramp Up (2 Times Faster)




Figure 59. 155-User Single Cisco UCS B230 M2 Active Compared to Allocated Memory Test Phase: 15-Minute Ramp Up (2 Times Faster)





Figure 60. 155-User Single Cisco UCS B230 M2 Network Rate Test Phase: 15-Minute Ramp Up (2 Times Faster)



Eight Cisco UCS B230 M2 Blade Servers: Single-Chassis Validation

This section details the results from the Citrix XenDesktop hosted VDI eight-blade server validation testing. It illustrates linear scalability from one blade with 135 users running the Login VSI Medium workload (with Adobe Flash) to eight blades with 1080 users.

The primary success criteria metrics are provided to validate the overall success of the test cycle as an output chart from Login Consultants' VSI Analyzer Professional Edition, VSImax for the Medium Workload (with Adobe Flash).

Additional graphs detailing the CPU and memory utilization during peak session load are also presented. Given adequate storage capability, the limiting factor in the testing was CPU utilization. The performance charts for each of the eight Cisco UCS B230 M2 blades are essentially identical for the multiple-blade runs. Data from one randomly chosen blade is presented in this document to represent all the blades in this portion of the study.

The single-server graphs shown in this section (Figures 61 through 67) are representative of a single VMware ESXi 5 host in the larger environment for validation purposes, but these graphs are representative of the behavior for all servers in the respective environment.

Performance information for relevant infrastructure virtual machines is presented with the randomly chosen tested blade information.

Figure 61. 1080 Citrix XenDesktop 5.5 Sessions on VMware ESXi 5 Below 4000 ms




Figure 62. 1080-User Blade 6-CPU Boot Phase: 15-Minute Ramp Up (2 Times Faster)




Figure 63. 1080-User Blade 6 Memory Granted Compared to Active Boot Phase: 15-Minute Boot Up (2 Times Faster)




Figure 64. 1080-User Blade 6 Network Rate Boot Phase: 15-Minute Bootup (2 Times Faster)




Figure 65. 1080-User Blade 6 Cisco UCS B230 M2 CPU Test Phase: 15-Minute Ramp Up (2 Times Faster)




Figure 66. 1080-User Blade 6 Memory Granted Compared to Active Test Phase: 15-Minute Ramp Up (2 Times Faster)





Figure 67. 1080-User Blade 6 Memory Granted Compared to Active Test Phase: 15-Minute Ramp Up (2 Times Faster)



Scalability Considerations and Guidelines

You need to consider many factors when you begin to scale beyond the Cisco FlexPod 2-chassis, 14-VDI-host-server configuration, which this reference architecture successfully tested. This section offers guidance about scaling beyond two Cisco UCS chassis.

Cisco UCS System Configuration

As the results discussed here indicate, the Cisco UCS reference architecture tested has proven linear scalability (Table 13).

Table 13. VMware ESXi 5 Hosting Citrix XenDesktop 5.5 with Citrix Provisioning Services 5.6.1

Number of Chasis Tested	Number of Cisco UCS B230 M2 Servers Tested	Number of Virtual Machines Hosted	Virtual Machines per Physical Core Ratio
1	1	135	6.75
1	8	1080	6.75

- Cisco UCS 2.0 management software supports up to 20 chassis within a single Cisco UCS domain on the second-generation Cisco UCS 6248UP and 6296UP Fabric Interconnects.
- Given adequate storage capability to support 20 chassis, you could easily scale to 20,000 users in a single Cisco UCS domain. You would need to add hypervisor virtual machine management virtual servers on each two-chassis deployment to manage the VDI environment.
- To accommodate the Cisco Nexus 5500 platform upstream connectivity as described in the [“LAN Configuration”](#) section, you need four Ethernet uplinks configured on the Cisco UCS fabric interconnect. Also, depending on the number of uplinks from each chassis, you could calculate the number of desktops that can be hosted in a single Cisco UCS domain. Assuming eight links per chassis, four to each Cisco

UCS 6248UP, scaling beyond 10 chassis would require a pair of Cisco UCS 6296UP Fabric Interconnects. A building block of 20,000 virtual desktops can be built from the Reference Architecture described in this study with eight links per chassis and 20 Cisco UCS chassis composed of seven Cisco UCS B230 M2 servers and one Cisco UCS B200 M2 server in each chassis.

Of course, the back-end storage has to be scaled accordingly, based on the IOPS considerations

Storage Sizing Best Practices

Storage estimation for deploying VDI solutions on enterprise storage includes the following steps:

- Gather essential solution requirements.
- Perform performance-based and capacity-based storage estimation.
- Get recommendations about storage system physical and logical configuration from the storage provider.

For More Information

Cisco Validated Design for Desktop Virtualization: www.cisco.com/vdidesigns



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)