HITACHI
Inspire the Next

Deploying Citrix XenDesktop 5 with Citrix XenServer 5.6 SP2 on Hitachi Virtual Storage Platform

Reference Architecture Guide

By Roger Clark

September 2011

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. Be sure to include the title of this white paper in your email message.

Table of Contents

Deploying Citrix XenDesktop 5 with Citrix XenServer 5.6 SP2 on Hitachi Virtual Storage Platform

Reference Architecture Guide

This reference architecture provides planning information to deploy Citrix XenDesktop 5 on Citrix XenServer 5.6 SP2 using Hitachi Compute Blade 2000 and Hitachi Virtual Storage Platform utilizing a building block approach. It contains details on deploying a non-persistent Microsoft Windows 7 virtual desktop infrastructure in a provisioning services (PVS) environment for typical users. The design objectives of this reference architecture are to:

- Lower administration costs
- Improve system efficiency
- Maintain performance

Testing in the Hitachi Data Systems lab shows that Hitachi Virtual Storage Platform and Hitachi Compute Blade 2000 make an ideal platform for a scalable virtual desktop infrastructure (VDI) implementation using Citrix XenDesktop running on Citrix XenServer.

The building block design used for this paper provides the scalability and performance needed to deliver a VDI solution. It scales from 205 to 3000 user desktop environment using the shared infrastructure components built into the infrastructure foundation block. With this design, adding additional user desktops to this environment means adding additional blocks of storage and server blades.

This paper’s intended use is by IT administrators responsible for desktops and storage administration. It assumes some familiarity with Hitachi Storage Navigator, Citrix XenDesktop, and Citrix XenServer.

Solution Overview

This describes deploying a XenDesktop environment that can scale up to 3000 desktops by using a two step building block approach.

The first step consists of an infrastructure foundation server and storage block that hosts desktops and all of the infrastructure components for the virtual desktop environment. The infrastructure foundation block has storage for desktops and infrastructure virtual machines.

- Two of the server blades in the infrastructure foundation block are dedicated to hosting the XenDesktop infrastructure components.
- The remaining six server blades are dedicated to hosting desktops.

Figure 1 shows the topology of the infrastructure foundation block.

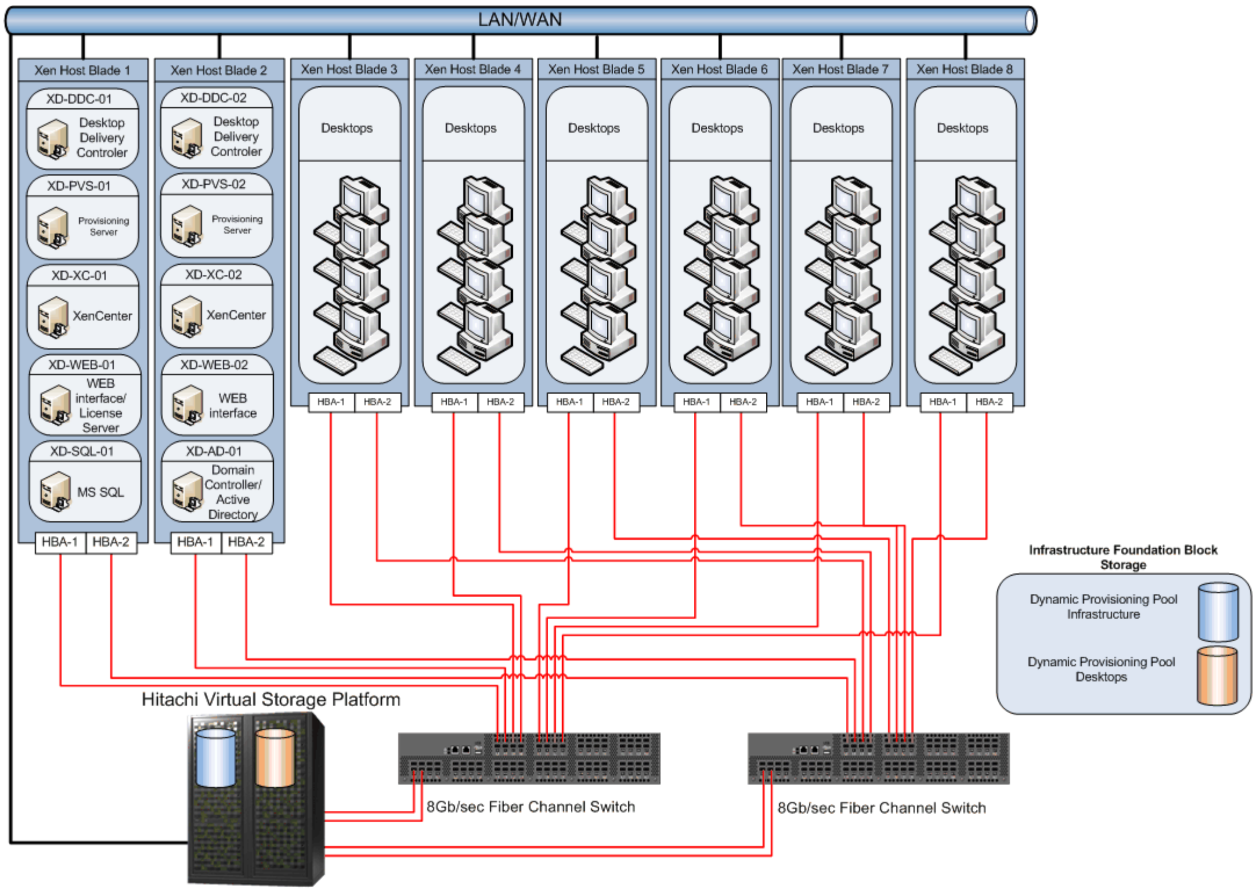

Figure 1

The second step adds server and storage blocks to scale out the environment. The storage blocks consist of storage for the additional desktops. The server blocks used to scale out the solution use all eight blades for desktops. They also use the infrastructure components in the infrastructure foundation block.

Figure 2 shows the topology of the server block.

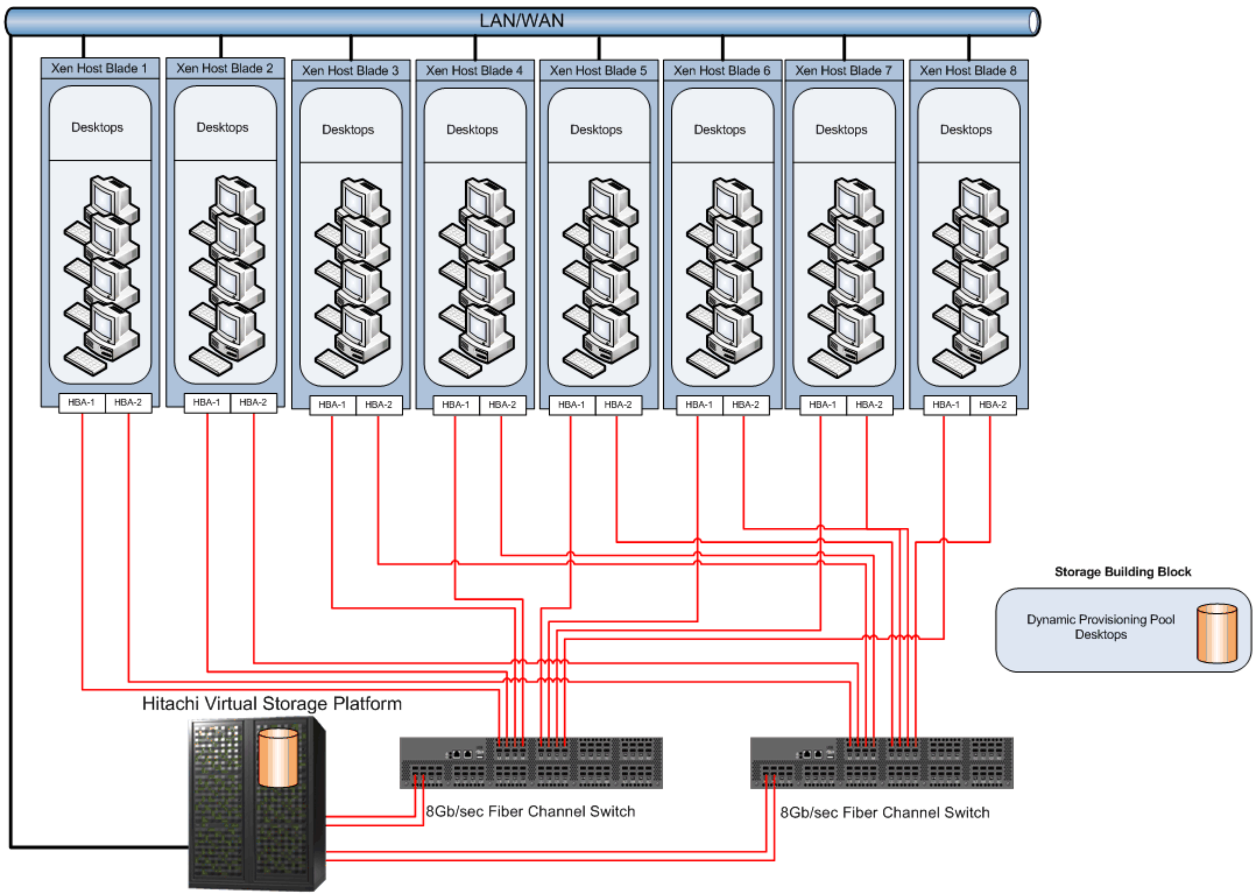

Figure 2

The infrastructure foundation block contains a storage pool for the infrastructure as well as for the desktops. The server building block only has storage for the desktops.

To achieve high availability for this solution, this architecture used the following:

- Redundant physical paths enabled via multiple host bus adapters (HBAs) on the servers
- Two Fibre Channel switches
- Multiple Fibre Channel paths to a Hitachi Virtual Storage Platform
- Citrix native multipathing for continued operation in the event of a hardware component failure

Key Solution Components

These are the key solution components.

Hitachi Virtual Storage Platform

Hitachi Virtual Storage Platform is the first 3D scaling storage platform designed for all data types. Its storage architecture flexibly adapts for performance, capacity, and multi-vendor storage. Combined with the unique Hitachi Command Suite management platform, it transforms the data center.

- **Scale Up** — Meet increasing demands by dynamically adding processors, connectivity, and capacity in a single unit. Provide the highest performance for both open and mainframe environments.
- **Scale Out** — Meet multiple demands by dynamically combining multiple units into a single logical system with shared resources. Support increased demand in virtualized server environments. Ensure safe multi-tenancy and quality of service through partitioning of cache and ports.
- **Scale Deep** — Extend storage value by virtualizing new and existing external storage systems dynamically. Extend the advanced functions of Hitachi Virtual Storage Platform to multivendor storage. Offload less demanding data to external tiers to save costs and to optimize the availability of tier one resources.

The trend in desktop virtualization is to consolidate the I/O workload of many desktops onto a single storage system. Consolidating virtual machines onto a physical host requires storage systems to dynamically add more storage resources to keep up with I/O demand. The 3D scaling capability of a Virtual Storage Platform meets that requirement.

For more information, see Hitachi Virtual Storage Platform on the Hitachi Data Systems website.

Hitachi Dynamic Provisioning

On Hitachi Virtual Storage Platform, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Hitachi Dynamic Provisioning is similar to using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups within a Hitachi Virtual Storage Platform. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB created against it (with no initial physical space allocated).

Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. This distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the DP-VOL happens as needed using 42MB pages, up to the logical size specified for each DP-VOL. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. An expanded pool can be rebalanced across the current and newly added RAID groups for an even striping of the data and the workload.

For more information, see the Hitachi Dynamic Provisioning datasheet and Hitachi Dynamic Provisioning on the Hitachi Data Systems website.

Hitachi Compute Blade 2000

Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

Hitachi embeds logical partitioning virtualization in the firmware of Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines the logical partitioning expertise of Hitachi with Intel VT technologies to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

For more information, see Hitachi Compute Blade Family on the Hitachi Data Systems website.

Citrix XenDesktop 5

The Citrix website describes XenDesktop 5 as follows:

> By transforming complex, distributed desktops into a simple, on-demand service, XenDesktop frees you from the costs and constraints of traditional computing architectures. Centralized delivery, management and control of virtual desktops bring new levels of efficiency to your IT organization while streamlining security and compliance. Self-service application provisioning, simplified helpdesk support and support for mobile and virtual work styles give you a foundation to leverage a new generation of IT models and strategies.
>
> Because it utilizes network streaming, a single image can be used to deliver a desktop to physical or virtual machines. However, this requires a fast and stable network. A network connection of less than 1Gb/sec to the client may result in slower desktop experiences by users.
>
> Find out more how Citrix XenDesktop helps your business.

This reference architecture describes a proven approach to help simplify the design and planning of the most commonly used form of virtual desktops, Provisioning Services (PVS).

Citrix XenServer 5.6 SP2

The Citrix website describes XenServer as follows:

> Citrix XenServer is a complete, managed server virtualization platform built on the powerful Xen hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows® and Linux® servers and delivers cost-effective server consolidation and business continuity.

Read more about XenServer on the Citrix website.

Solution Design

This reference architecture uses Citrix XenDesktop with the Citrix Provisioning Services model to deliver non-persistent desktops. This robust desktop delivery method is used commonly in large VDI implementations. Your solution may require a different delivery model.

This is detailed information on the Citrix XenDesktop environment design used for this reference architecture. Although it includes application and hardware design information, it focuses on the storage design required to build the basic infrastructure for the XenDesktop environment. Reference to this white paper is one part of planning your complete Citrix implementation.

Note — Validation of this design was done in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before solution implementation in your production environment. This means to test applications in a non-production, isolated test environment that otherwise matches your production environment.

This solution design discussion includes the following:

- Infrastructure Foundation Building Block
- Building Blocks
- Environment Scaling
- Fibre Channel Design
- Tested Hardware Components
- Hardware Sizing
- Storage Architecture
- Environment Storage Architecture

Infrastructure Foundation Building Block

This describes the server and storage design of the infrastructure foundation block. See “Building Blocks” for the difference between an infrastructure foundation building block and a server block.

Servers

The server portion of the infrastructure foundation block consists of one Hitachi Compute Blade 2000 chassis with eight Hitachi X55A2 server blades running Citrix XenServer 5.6 SP2.

The first two server blades of the chassis host the XenDesktop Servers supporting the virtual desktops running in the environment. Two server blades are used to host the XenDesktop Servers for high availability. Virtual desktop requests are load balanced between all XenDesktop servers using the load balancing functionality built into the Citrix products.

Also, the infrastructure foundation block is the location of the additional Citrix Provisioning Services servers that are added as the environment is scaled out. Citrix recommends having one Provisioning Services server for every 500 desktops. For this reference architecture solution, that is roughly one Provisioning Services server added for every two server blocks.

Table 1 lists the virtual machines hosted on Blade 0 and Blade 1. Table 2 lists the desktops hosted on Blade 2 through Blade 7.

Table 1. Virtual Machines Hosted on Blade 0 and Blade 1

Server	*Number of Servers*	*Configuration*
Citrix XenDesktop Desktop Delivery Controller (DDC)	2	2 × vCPU 4GB RAM 1 × 24GB VHD
Citrix Provisioning Service servers (PVS)*	2	2 × vCPU 4GB RAM 1 × 24GB VHD 1 × 100GB VHD (Image storage)
Citrix Web Interface and License servers† (WI)	2	2 × vCPU 4GB RAM 1 × 24GB VHD
Citrix XenCenter servers (XC)	2	2 × vCPU 4GB RAM 1 × 24GB VHD
Microsoft SQL Server 2008	1	2 × vCPU 4GB RAM 1 × 24GB VHD 1 × 20GB VHD (SQL logs) 1 × 100GB VHD (SQL data)

*Two PVS servers are part of the infrastructure foundation block. Add additional PVS servers as more virtual desktops are added to the environment.

†License Server only runs on one of the web servers

Table 2. Desktops Hosted on Blade 2 through Blade 7

Desktop	*Configuration*
Microsoft Windows 7 64bit	1 × vCPU 2GB RAM 1 × 5GB VHD (write cache and pagefile)

Storage

The storage for the infrastructure foundation block consists of the following:

- **Infrastructure pool**—12 × 300GB 10k RPM SAS drives in a RAID-5 (3D+1P) Hitachi Dynamic Provisioning pool with a single 1.5TB LUN.
- **Desktop pool**—36 × 146GB 15k RPM SAS drives in a RAID-1+0 (2D+2D) Hitachi Dynamic Provisioning pool with two 1TB LUNs.

Add additional LUNs to the infrastructure pool to support environments with larger user profile requirements or to support more golden images for the PVS servers.

Building Blocks

Once the infrastructure foundation block is in place, scale out the environment by adding server building block and storage building block. A server block and a storage block are separate blocks to add to the capacity of the environment. Use Citrix XenDesktop components in the infrastructure foundation block to add additional desktops to the environment.

Server Building Block

A server building block consists of the following:

- One Hitachi Compute Blade 2000 chassis
- Eight Hitachi X55A2 server blades running Citrix XenServer 5.6 SP2. All eight server blades host desktops.
    - Of the eight server blades in an infrastructure foundation building block, two blades host infrastructure servers and six blades host desktops. See “Infrastructure Foundation Building Block.”

Figure 3 shows the difference between the infrastructure foundation block and a server building block.

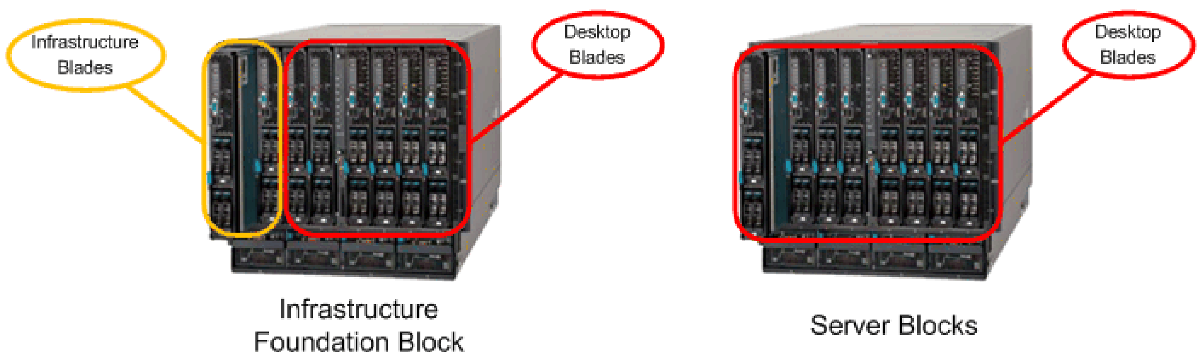

Figure 3

Storage Building Block

A storage block supports two server blocks. Each storage block consists of one dynamic provisioning pool with the following:

- 72 × 146GB 15k RPM SAS drives configured in a RAID-1+0 (2D+2D) array
- 4 × 1TB LUNs

It is based on the following:

- A non-persistent desktop with an I/O profile of 10 to 15 IOPS per desktop
- A cache of 5GB per user

Figure 4 shows the difference between the first building block and additional building blocks.

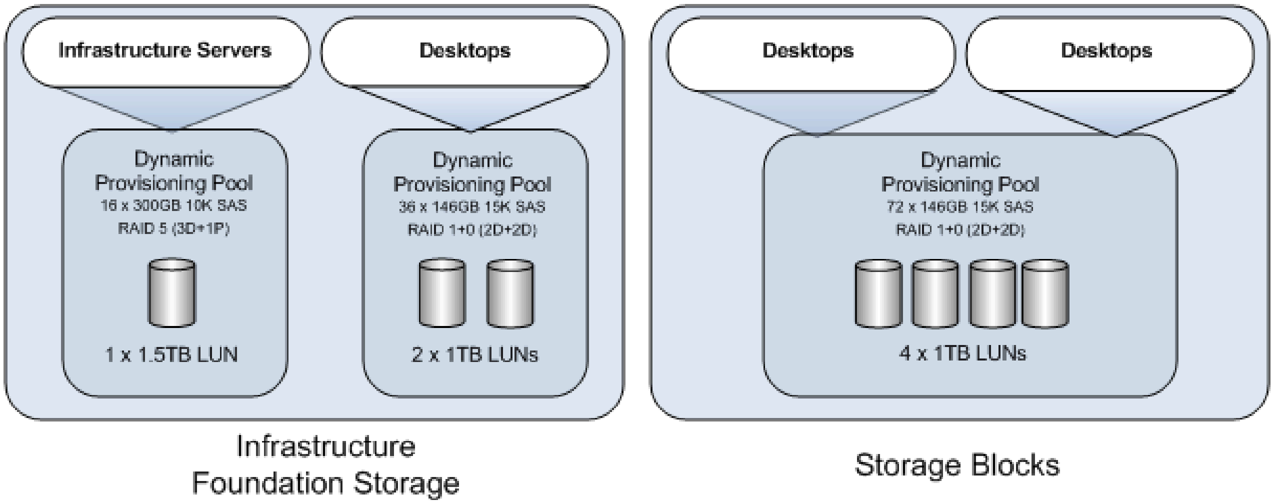

Figure 4

Environment Scaling

To increase the number of desktops in this solution, additional server blocks and storage blocks are added.

- The infrastructure foundation block in this reference architecture supports a Citrix XenDesktop infrastructure and 205 desktops.
- Each server block supports 271 desktops.
- Each storage block supports two server blocks.

This shared infrastructure supports up to 3000 desktops, according to the maximum desktops recommended by Citrix per XenDesktop farm. For more information, see *Best Practices for Citrix XenDesktop with Provisioning Server*.

To scale the environment out to 3,000 desktops requires the following:

- 11 server blocks
- 6 storage blocks
- 6 PVS servers in the infrastructure block

Figure 5 shows how the solution can be scaled out to 3,000 desktops by adding additional server and storage building blocks to the environment.

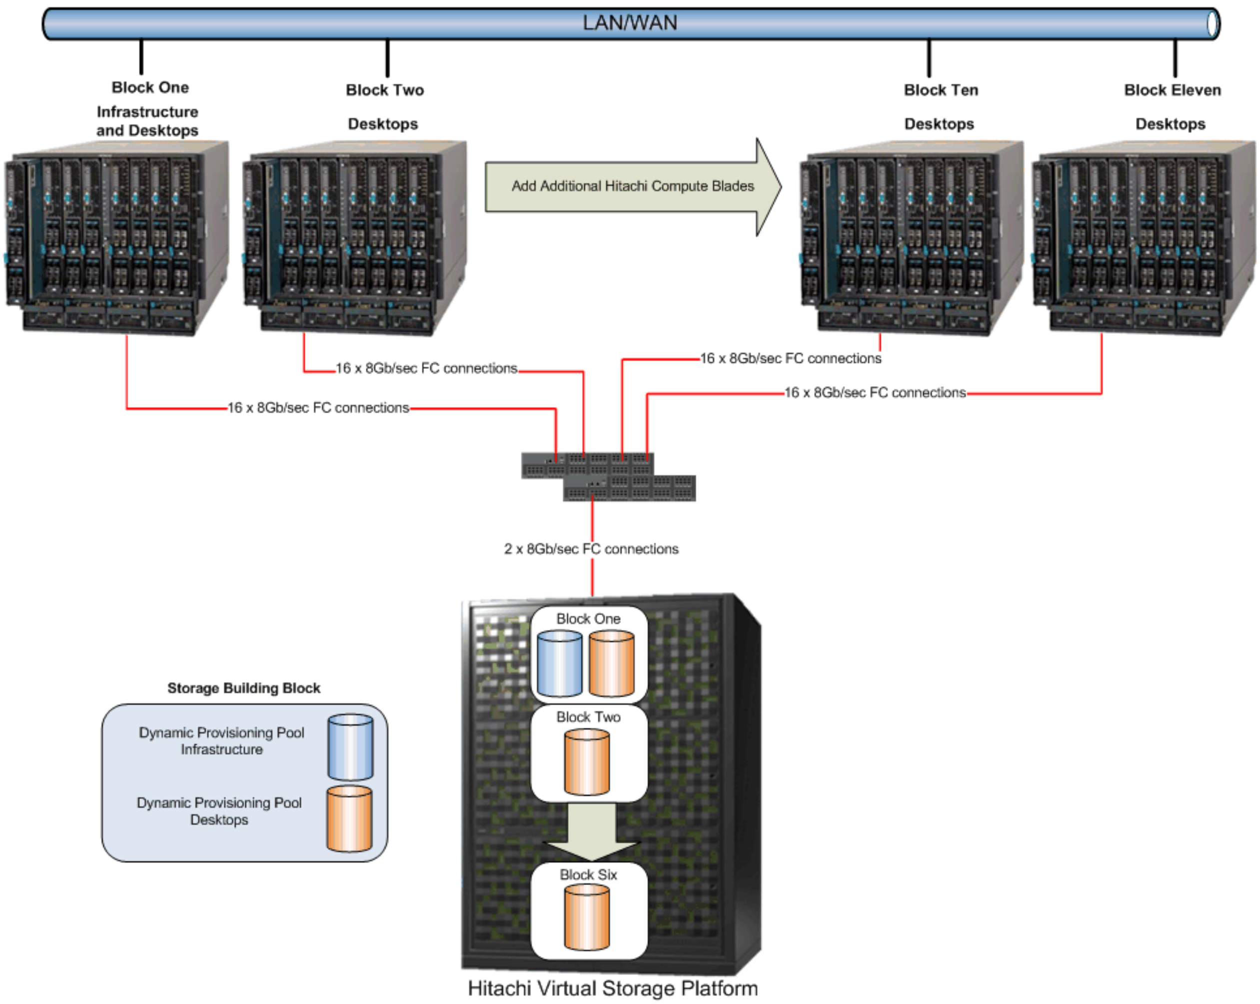

LAN/WAN
Block One
Infrastructure and Desktops
Block Two
Desktops
Add Additional Hitachi Compute Blades
Block Ten
Desktops
Block Eleven
Desktops
16 x 8Gb/sec FC connections
16 x 8Gb/sec FC connections
16 x 8Gb/sec FC connections
16 x 8Gb/sec FC connections
2 x 8Gb/sec FC connections
Storage Building Block
Dynamic Provisioning Pool Infrastructure
Dynamic Provisioning Pool Desktops
Block One
Block Two
Block Six
Hitachi Virtual Storage Platform

Figure 5

Fibre Channel Design

A Hitachi Compute Blade 2000 chassis and a Hitachi Virtual Storage Platform are connected using two Fibre Channel switches for redundancy and high availability. Figure 6 shows the Fibre Channel connections for this solution.

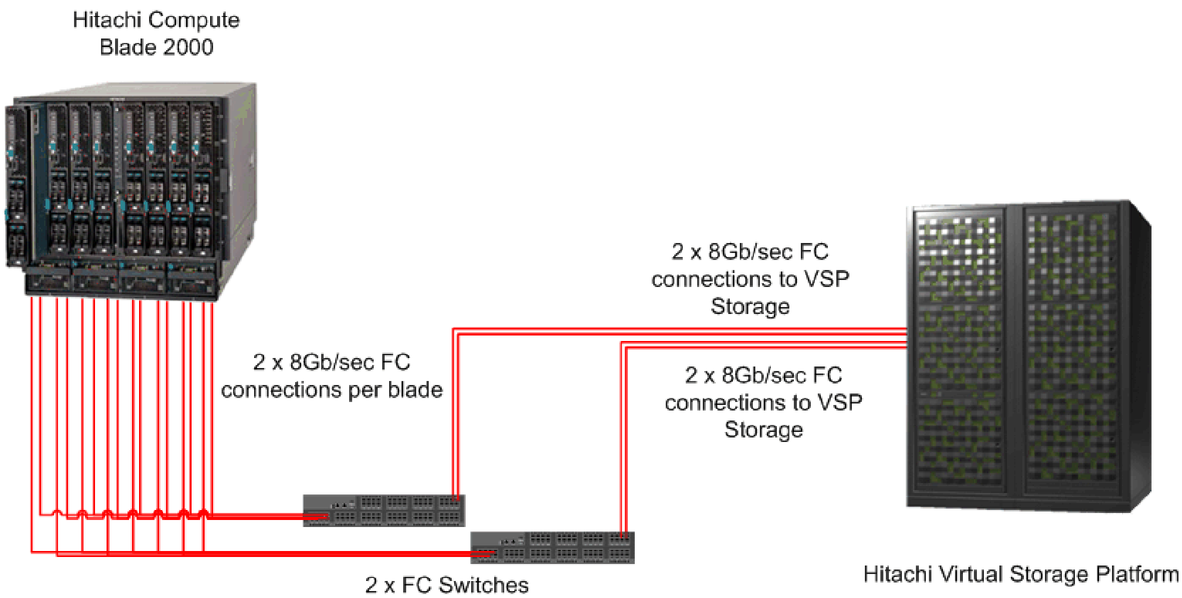

Figure 6

Another option is to deploy the environment using an enterprise-level director that contains multiple blades that can support high availability and redundancy.

Tested Hardware Components

The hardware for this solution architecture consists of the following:

- Two Hitachi Compute Blade 2000 chassis
- Ten X55A2 server blades
- One Hitachi Virtual Storage Platform
- Two Brocade 5000 Fibre Channel switches

Table 3 describes the hardware components used in this reference architecture.

Table 3. Hardware Components

Hardware	*Detail Description*	*Version*	*Qty.*
Hitachi Virtual Storage Platform	80GB cache memory 32 × 8Gb/sec ports 72 × 146GB, 15k RPM, SAS disks 16 × 300GB 10k RPM, SAS disks	70-02-00-00/7	1
Hitachi Compute Blade 2000 Chassis	4 × 1Gb/sec Ethernet switch modules 10 × Emulex 8GFC 2P PCIe Card	A0160-G-5666	2
Hitachi X55A2 Server Blades	2 × Six-core Intel CPU 72GB RAM 2 × 600GB 10k SAS disks 1 × Hitachi GbE 4 Port Mezzanine Card	03-79	10
Brocade 5000 Switch	SAN Switch with 8Gb/sec Fibre Channel ports	FOS 5.3.1a	2

Hardware Sizing

Determining the number of desktops your hardware supports is dictated by a number of factors, such as the following:

- User workload
- Desktop OS
- Applications running on the desktops
- The amount of random access memory required for each desktop

Most Citrix XenDesktop implementations have a number of different user categories:

- Light
- Medium (most users)
- Heavy

Use these categories to define the resources needed for the design and planning of your environment. For the purpose of this reference architecture, the hardware was sized around a medium workload profile.

The formulas used in this reference architecture to determine the recommended number of desktops that could be run on this hardware configuration were taken from the Citrix XenDesktop E-Docs document library.

Based on the breakdown of the number of users in each category, and following the CPU and memory allocation guidelines, Table 4 has recommendations to calculate resource allocation for desktops using Microsoft Windows XP and Microsoft Windows 7.

Table 4. Microsoft Windows XP and Windows 7 Resource Recommendations

User Group	Operating System	vCPU Allocation	Memory Allocation per User	Average Steady State IOPS*	Estimated Users/Core
Light	Windows XP	1	768MB to 1GB	3-5	10 to 12
	Windows 7	1	1GB to 1.5GB	4-6	8 to 10
Medium	Windows XP	1	1GB to 1.5GB	6-10	8 to 10
	Windows 7	1	1.5GB to 2GB	8-12	6 to 8
Heavy	Windows XP	1	2GB	20-40	4 to 6
	Windows 7	2	4GB	25-50	2 to 4

*The IOPS estimates are an average of when the user is logged onto the virtual desktop and working. It is not a peak average, which also takes into account start ups, logons, and logoffs.

Use the guidelines in Table 4 with the following two formulas to estimate the scalability of a particular physical server.

- **Memory Perspective**

$$Number\ of\ Users\ = \frac{(Server\ RAM - Hypervisor\ Overhead\ RAM)}{(Average\ Memory\ Allocation\ /\ User)}$$

- **Core Perspective**

$$Number\ of\ Users = (Cores - 1\ core) \times (Average\ Users/Core)$$

The lower result from the two formulas is the *recommended* maximum number of users.

Note — Each hypervisor has memory and CPU overhead, as shown in both equations. On the average, the hypervisor needs 6GB to 8GB of RAM. Subtract one core from the number of available cores to account for the CPU overhead.

For the test environment this is calculated as the following:

- All users are medium users on Microsoft Windows 7
- The server configuration has 72GB of memory with 12 cores

The values used from Table 2 are:

- 2GB of RAM per user
- Average of 8 users per core

The number of users from a memory perspective is 33 users.

$$33\ Users = \frac{(72GB - 6GB)}{(2\ GB/User)}$$

The number of users from a core perspective is 88 users.

$$88\ Users = (12\ cores - 1\ core) \times (8\ Users/Core)$$

In the Hitachi Data System test environment, each server blade was configured with 72GB of memory. This limits the *recommended* maximum number of desktops per server blade to only 33 desktops per blade.

To increase the number of users on this reference architecture, increase the memory so it supports roughly the same number of users as the cores can support. This means to reconfigure the server with 192GB of memory.

$$93\ Users = \frac{(192GB - 6GB)}{(2\ GB\ /\ User)}$$

Storage Architecture

The storage architecture used to validate this solution takes into consideration Citrix best practices for the deployment of a XenDesktop environment. This reference architecture uses industry-standard formulas to determine the number of IOPS required to support the number of desktops in our test environment.

Storage Sizing

The many variables that dictate capacity requirements for a Citrix XenDesktop environment include the following:

- Number of users and their workload
- Type of delivery model for the desktop
- Persistent or non-persistent desktops

The storage in this reference architecture is sized around a medium workload profile. To validate that the environment can handle spikes in user load and other unforeseen bursts of I/O without users suffering any performance degradation, testing used a medium workload. This was done to validate that this environment.

Provisioning Services Storage

When using Citrix Provisioning Services (PVS), which is based on software streaming technology, do the following:

1. Prepare a master target device to be imaged by installing an operating system and software.
2. Create a virtual disk (vDisk) image from the hard drive of the master target device and save it to the provisioning server (PVS).
3. Stream the contents of the vDisk to the target device on demand, in real time, using the software-streaming technology of PVS.

Once the vDisk image is available from the network, a virtual machine on a target device no longer needs a local hard drive to operate. A target device boots directly from the network. It behaves as if it is running from a local drive on the target device.

Unlike thin-client technology, processing takes place on the target device. Each target machine contains a volatile write cache file. Each restart cycle deletes this file.

The size of the cache file for each virtual machine depends on several factors, including:

- The types of applications used
- User workloads
- Reboot frequency

The general estimate for the cache size of a provisioned workstation which is rebooted daily that only runs applications such as Microsoft Word and Outlook is 300MB to 500MB. If a workstation is rebooted less often, or when the workstation uses applications that are graphic intensive such as Microsoft PowerPoint, Microsoft Visual Studio, or a CAD/CAM application, the cache size can grow much larger.

As an application workload can vary for each environment, perform a detailed analysis on your environment to determine an expected cache file size.

A common practice is to place the Windows page file on the same disk as the write cache. If you plan to do this, it is recommended practice to at least double the size of your write memory cache.

There are several options for storing the cache file for provisioned desktops. Two common locations with their benefits and limitations are the following:

- **Physical Local Storage.** If using physical desktops or blade servers, Citrix recommends using the memory on the target device or the hard disk to store cache files. Consider the following when using this method:
    - Accessing memory is considerably faster than accessing a hard disk. Placing the cache file in memory often provides the best performance for a single application environment.
    - Storing the cache in the memory of the target machine limits its size to the available amount of physical memory in the device. If the cache file is expected to grow larger than that, use the hard drive or an enterprise storage device. Filling a RAM-based cache file may cause an undesired result or errors.
    - The hard drive on a target machine usually has more space available for the cache file than the memory. If the target machine does not contain hard drives or enough memory for the required size of the cache file, then store the cache file on a shared enterprise storage device and use the Provisioning Server for Desktops as a proxy.
- **Client Disk (Virtual Machine vDisks).** A virtual machine infrastructure hosting the desktops also can use for the location of the cache the enterprise storage associated with each virtual machine or the allocated memory for the virtual machine. For example, created each virtual machine with an additional disk to store the cache file on each the vDisk of target device. Consider the following when using virtual machine vDisk-based cache locations:
    - Storing the cache file as part of the virtual machine requires space for the allocated size of each cache file in addition to the size required for each unique virtual machine.
    - This option provides added resiliency. If a cache file for a single virtual machine runs out of space, it affects only that virtual machine.
    - Utilizing a client disk for the cache file lowers the overall network traffic, when compared to utilizing an enterprise server-based option. All caching is done within the virtual machine.

Figure 7 shows the PVS file layout for the write cache and vDisk.

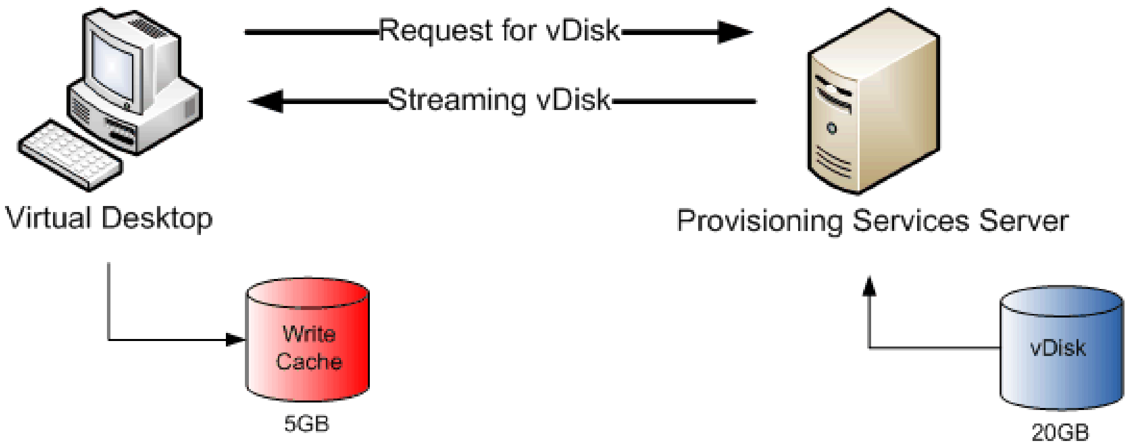

Figure 7

To estimate the storage required for the PVS model, use this formula:

$$Storage\ Required = (Number\ of\ Users\ \times Write\ Cache\ Size) + \ (Number\ of\ Desktop\ Images\ \times Desktop\ Image\ Size)$$

For example, to calculate the space required for 271 non-persistent PVS desktops using one 20GB master image and a 5GB vDisk for write cache and the Windows pagefile:

$$1375\ GB = (271\ users\ \times 5120 MB/user) + \ (1\ image\ \times 20480 MB/image)$$

1375GB is approximately 1.5TB, the size that as was used for the storage block in this architecture.

Each of these options has benefits and limitations. Evaluate your specific requirements before determining the location of the cache file.

For this reference architecture, the cache file and Windows page file was stored on a 5GB client vDisk.

IOPS Planning

No precise method exists to establish VDI IOPS requirements. Follow the recommended practice of doing a proof-of-concept test in a VDI non-production environment to establish your IOPS requirements for your VDI production environment.

Although not precise, when determining your IOPS requirements, the following helps to calculate an approximate number of spindles needed to support a given workload:

- Determine an accurate I/O profile of the current desktops in your environment with an I/O monitoring tool.
- Consider the times that users will log on to their desktops and start their desktops to reduce the effects of boot storms.
- Consider using a server-based antivirus solution, such as VMware V-Shield, to reduce excess I/O on a user desktop.

Use the following formulas to calculate the approximate functional IOPS available for a VDI workload.

$$Total\ Raw\ IOPS = (Disk\ Speed\ IOPS/Drive) \times (Number\ of\ Disk\ Drives)$$

$$Functional\ IOPS = \left\{\frac{[(Total\ Raw\ IOPS) \times (Percentage\ of\ Write\ Operations\ written\ as\ a\ decimal)]}{RAID\ Penalty}\right\} + \{(Total\ Raw\ IOPS) \times (Percentage\ of\ Write\ Operations\ written\ as\ a\ decimal)\}$$

Assumptions for this reference architecture were for a standard of 80% write operations and 20% read operations to simulate a user running common office applications.

For example, to calculate the functional IOPS for this test environment:

- 146GB 15k RPM SAS drives capable of 150 IOPS per drive
- 72 total drives
- RAID-1+0 configuration

$$Total\ Raw\ IOPS = (150\ IOPS/drive) \times (72\ Drives)$$

$$Total\ Raw\ IOPS = 10800\ IOPS$$

$$Functional\ IOPS = \left\{\frac{[(10800\ IOPS) \times (0.8)]}{2}\right\} + \{(10800\ IOPS) \times (0.2)\}$$

$$Functional\ IOPS = 6480\ IOPS$$

Using this approach, storage for this reference architecture was built to handle 6480 IOPS under a sustained workload.

Boot Storm Consideration

A boot storm occurs when a large number of desktops start in a relatively short amount of time. This adversely affects the user experience because of a spike in IOPS exceeds the number of calculated IOPS for sustained workloads.

Test your environment to determine the effect of starting up a large numbers of desktops. Determine how many desktops need pre-booting in your environment and schedule the desktop boot times to meet those needs. Then, use the XenDesktop power management features to start up user desktops automatically prior to users needing to log on to reduce the affect of a boot storm in your environment.

Environment Storage Architecture

For the purpose of testing, the environment was built in two pieces.

- A Hitachi Compute Blade 2000 chassis with two blades running the environments infrastructure.
- A second Hitachi Compute Blade 2000 with eight blades running the desktops.

This was done to isolate each environment and to simulate a full server chassis and storage block.

The storage design used for this environment consists of:

- 72 × 146GB 15k RPM SAS drives configured as RAID-1+0 (2D+2D) hosting all of the desktop virtual machines, the desktop base images, and the users write cache in the PVS environment
- 12 × 300GB 10k RPM SAS drives configured as RAID-5 (3D+1P) assigned to two dedicated dynamic provisioning pools hosting all of the user profiles and infrastructure virtual machines

Table 5 describes the dynamic provisioning pools used in this white paper

Table 5. Dynamic Provisioning Pools

Pool Name (ID)	RAID Group Configuration	Drive Type	Number of RAID Groups	Pool Capacity (TB)
HDP-XD-D001	RAID-1+0 (2D+2D)	72 × 146GB 15k RPM SAS	18	4.74TB
HDP-XD-U001	RAID-5 (3D+1P)	12 × 300GB 10k RPM SAS	4	3.14TB

Infrastructure Storage Overview

The storage block for the infrastructure consists of a 1.5TB LUN carved out of a 3.14TB RAID-5 dynamic provisioning pool. This additional infrastructure storage space allows for the future growth of the environment to handle the additional user profiles if the environment is scaled up to 3,000 desktop users and any additional provisioning services servers needed in the PVS environment.

Table 6 describes the LDEV created for the infrastructure virtual machines and user profiles.

Table 6. Infrastructure LDEV configuration

LDEV Name (ID)	LDEV	Size (TB)	Purpose	Storage Port
USP-LDEV-01	00:00:12	1.5	Infrastructure Virtual Machines and User Profiles	3D/4D

Provisioning Services (PVS) Storage Overview

The storage block for the PVS environment consists of a 4.74TB RAID-1+0 dynamic provisioning pool carved into two 1TB LUNs that store all Citrix XenDesktop virtual machines.

Table 7 describes the LDEVs created for the PVS storage blocks in the test environment.

Table 7. PVS LDEV Configuration

LDEV Name (ID)	LDEV	Size (TB)	Purpose	Storage Port
PVS-LDEV-01	00:00:33	1TB	Desktop Virtual Machines User Write Cache	3D/4D
PVS-LDEV-02	00:00:34	1TB	Desktop Virtual Machines User Write Cache	3D/4D

Engineering Validation

After building the reference architecture in the Hitachi Data Systems lab, a series of tests were run to validate the designs. The testing focused on the entire virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon, user workload execution, and user logoff. Test metrics were gathered from the hypervisor, virtual desktops, storage, and load generation software to assess the overall success of an individual test cycle.

User Workload Simulation – Login VSI from Login Consultants

To validate a Citrix XenDesktop deployment accurately requires using a real-world user workload. To accurately validate this architecture, third-party tools from Login Consultants were used throughout virtual desktop infrastructure testing. These tools have the benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for an individual desktop throughout large scale testing, including logon storms.

Login Virtual Session Indexer (Login Consultants VSI 3.1) methodology designed for benchmarking server-based computing (SBC) and virtual desktop infrastructure (VDI) environments is completely platform and protocol independent. This allows you to test your specific environment. Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.

Login VSI simulates a light to heavy workload user running generic applications such as Microsoft Office 2007, Microsoft Internet Explorer including Adobe Flash applets, and Adobe Reader.

For this test, applications are installed locally, not streamed or hosted on Citrix XenApp.

Like real users, the scripted session leave multiple applications open at the same time. Every session averages about 20% minimal user activity, similar to real world usage. During each 12 minute loop users open and close files a couple of time per minute, which is probably more intensive that most users.

The following outline the automated Login VSI simulated user workflows that were used for this validation testing:

- Use Microsoft Office 2007 applications, Microsoft Internet Explorer, 7-Zip, Bullzip PDF Printer, and Adobe Reader to open up to five applications at the same time with a type rate of 160ms for each character. The workload observes approximately 2 minutes of idle time, which closely simulates real-world users.
- After starting a session, the workload repeats every 12 minutes. During each loop, measure the response time every 2 minutes.
- Each loop consists of the following operations:
    - Browse and compose Microsoft Outlook 2007 messages.
    - Open multiple instances of Microsoft Internet Explorer-based browsing sessions, including heavy multimedia websites.
    - Open multiple instances of Microsoft Word 2007 to perform open, close, and edit operations.
    - Print and review PDF documents using Bullzip PDF Printer and Acrobat Reader.
    - Open, edit, and close a randomized large Microsoft Excel 2007 sheet.
    - Review and edit a Microsoft PowerPoint 2007 presentation.
    - Perform file compression operations using 7-Zip.

Success Criteria

Multiple metrics were captured during each test run. The success criteria for considering a single test run as pass or fail was based mainly on the following main metrics:

- Login VSI Max evaluates the user response time during increasing user load.
- Login VSI Correct Optimal Performance Index (COPI) assess the successful start-to-finish execution of all the initiated virtual desktop sessions.

These main metrics are important on the raw data that they provide and in their ability to align the test results between the hosted shared and hosted VDI models.

Login VSI Corrected Optimal Performance Index (COPI)

The Login VSI Corrected Optimal Performance Index (COPI) is calculated from specific measurements during each test run to determine how many desktops can run simultaneously without excessively affecting user experience.

COPI is based on these measurements:

- **The Uncorrected Optimal Performance Index (UOPI)**—This is based on the first 5 consecutive sessions that hit the "Optimal Performance Max Reached" threshold. This is calculated on the response time average of four sessions higher than 2000ms, with a 4 session average response time greater than 8000ms.
- **The Stuck Session Count (SSC)**—This represents sessions which have become stuck before UOPI and must be accounted for in COPI.
- **The Lost Session Count (LSC)**—This is a count of completely missing log files; these tests are discarded completely in the corrected index.

Incorporating SSC and LSC into COPI helps ensure that the test results are fair and comparable. COPI is calculated: as:

$$COPI = UOPI - \left(\frac{SSC}{2}\right) - LSC$$

Login VSI Max

VSI Max is the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users, as indicated during the workload execution. Response time is generally an indicator of host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.

The user response time has a threshold of 2000ms. All user response times are expected to be less than 2000ms in order to assume that the user interaction with the virtual desktop is at a functional level.

VSI Max is reached when the response times reaches or exceeds 2000ms for 6 consecutive occurrences. If VSI Max is reached, then the test run is considered a failure because user experience has significantly degraded.

EdgeSite works with VSI to manage the tests and collect performance data. To collect performance data from the Citrix XenServer environment, dstat collected data from all of the XenServer servers.

The testing consisted of two test scenarios, as follows:

- **Desktop Start-up Test**—This simulates a start of business day or shift change scenario. All or some of the desktops are pre-started prior to users log on. This test uses the power management features of Citrix XenDesktop. This tries to prevent boot storms that can adversely affect a user's experience.
- **Workload Performance Test**—This simulates the load created in an environment when users log on and start running a measured workload. This test uses Login Consultants VSI. This validates the calculations used to build the environment.

This testing provides the data needed to validate Citrix XenDesktop with Citrix Provisioning Servers with Citrix XenServer 5.6 SP2 that virtualizes Microsoft Windows 7 desktops on Hitachi Compute Blade server blades using a Hitachi Virtual Storage Platform storage array. This information provides data points that you may reference when designing your own implementation and environment.

These validation results are an example of what is possible under the specific environment conditions outlined in this white paper. These results do not represent the full characterization of Citrix XenDesktop with XenServer scalability.

Desktop Start-up Test Results

A boot storm occurs when a large number of desktops start in a relatively short amount of time. This adversely affects the user experience because of a spike in IOPS exceeds the number of calculated IOPS for sustained workloads.

Citrix XenDesktop has power management features in the Desktop Group properties to help prevent boot storms. This schedules the booting of desktops before they are needed. VDI administrators may set a predetermined number of desktops to be booted and available for use at a specific time.

This test determined the time required to pre-boot 271 desktops in the test environment.

The desktop booting test followed these steps:

1. The desktop group is powered down and placed in maintenance mode.
2. The desktop group is taken out of maintenance mode.
3. XenDesktop sends power up commands to the hypervisor to start each desktop until the designated number of desktops have booted into a ready state.

This test determined the amount of time needed to pre-boot a given number of desktops in a VDI environment. In the test environment, all 271 desktop booted and entered a ready state in roughly 9 minutes. The results for your environment may be different.

Another key benefit of using the XenDesktop power management feature is how this spreads out the IOPS evenly over the boot up sequence. This helps to prevent large IOPS spikes that could degrade user performance.

With the Citrix XenDesktop power management feature set to boot all 271 desktops, Figure 8 and Figure 9 charted the IOPS during the boot process for the desktop and infrastructure environments. The system was at steady state in 9 minutes.

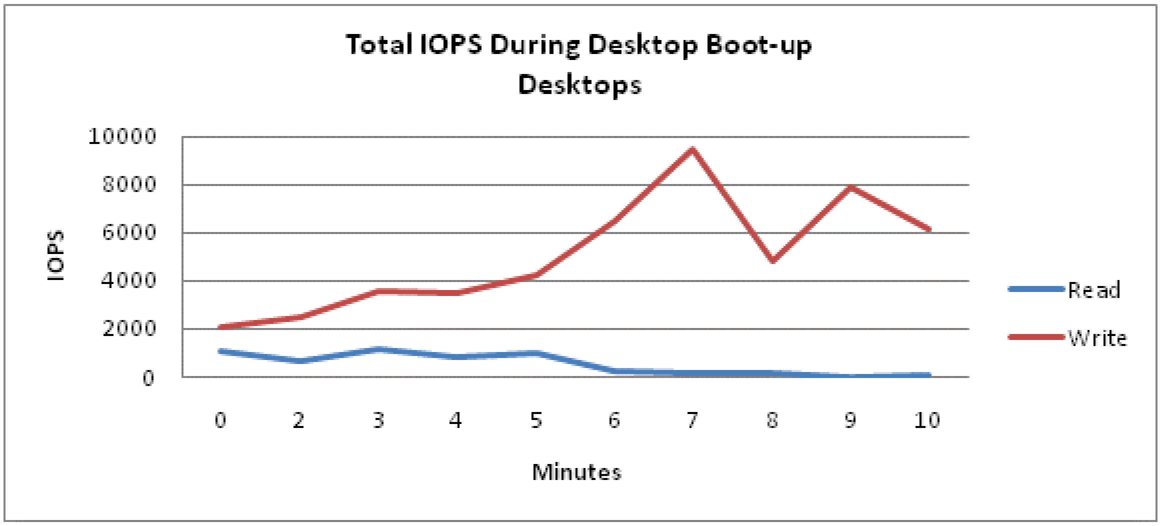

Figure 8

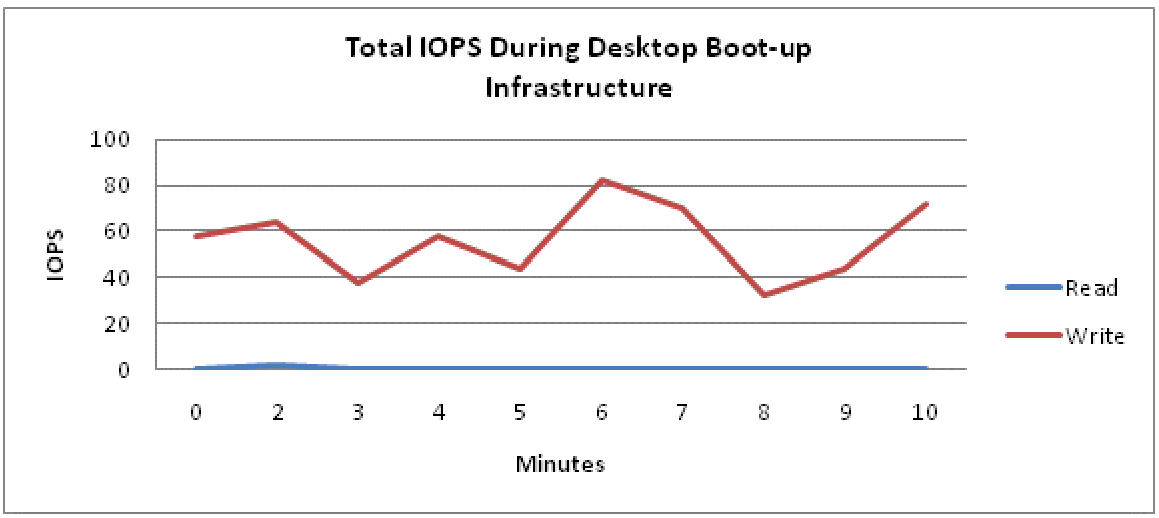

Figure 9

Figure 8 and Figure 9 show that using the power management feature in Citrix XenDesktop evenly spreads out the IOPS. No major spikes from power up adversely affected the performance of any user already logged on to the environment.

Average CPU utilization shows a slight increase during boot-ups. This should be considered whenever scheduling automated boot-ups. Figure 10 show the CPU usage for the desktop environment.

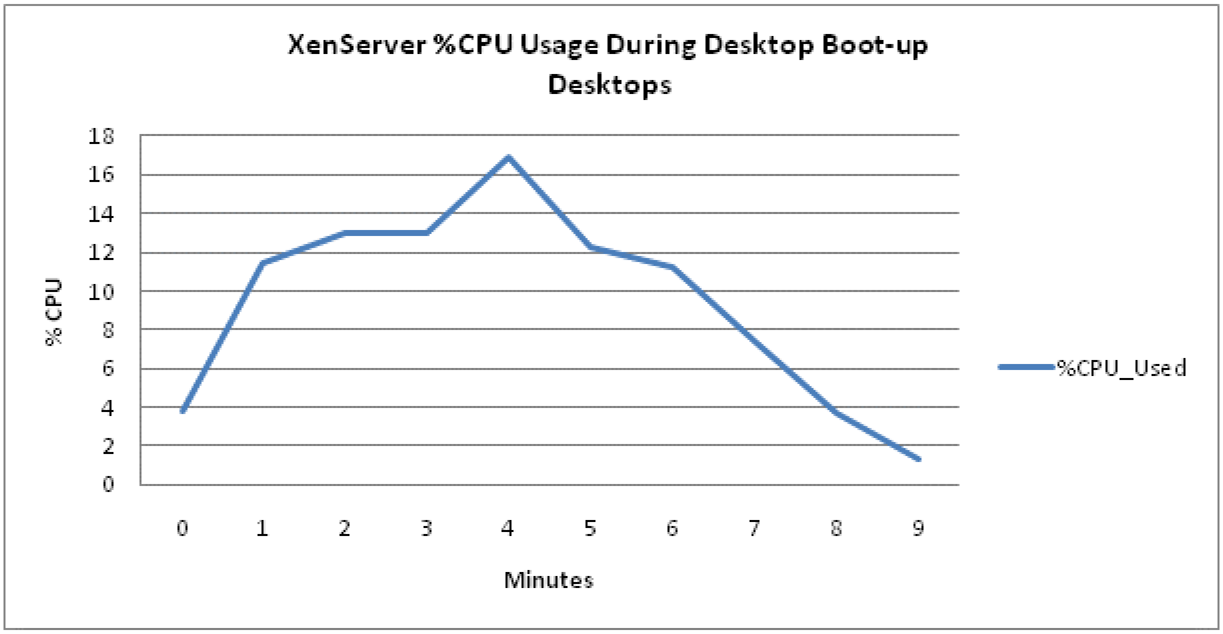

Figure 10

Figure 11 shows the CPU usage for the infrastructure environment.

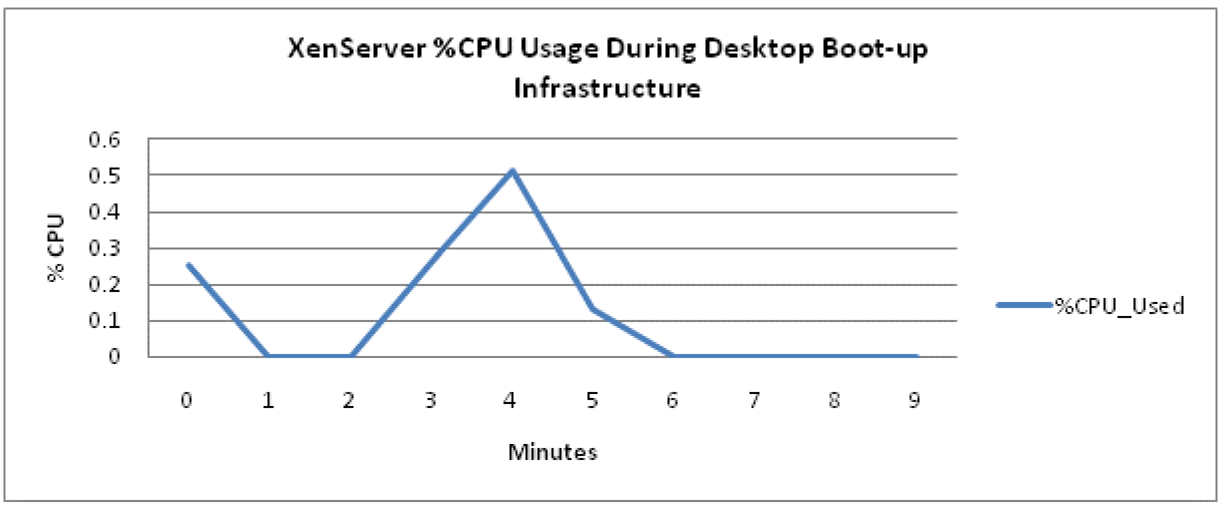

Figure 11

Workload Performance Test Results

This test validates that a Hitachi Compute Blade 2000 server blade was able to support the desired load of 271 virtual desktops under a medium workload.

Table 8 provides the VSI COPI score for the overall Hitachi Compute Blade 2000 server environment and shows that 100 percent of all the 271 virtual desktop sessions executed without issue.

Table 8. Workload Performance Test Metrics

Workload Performance Test Metrics	Number
Total Sessions Launched	271
Uncorrected VSI Max (UOPI)	271
Stuck Session Count before UVM (SSC)	0
Lost Session Count before UVM (LSC)	0
Correct Optimal Performance Index (COPI)	**271**

The Login VSI test measures the user response time and evaluates the user experience based on workload response time. Figure 12 shows that under a medium workload the user response time was within the 2000ms threshold.

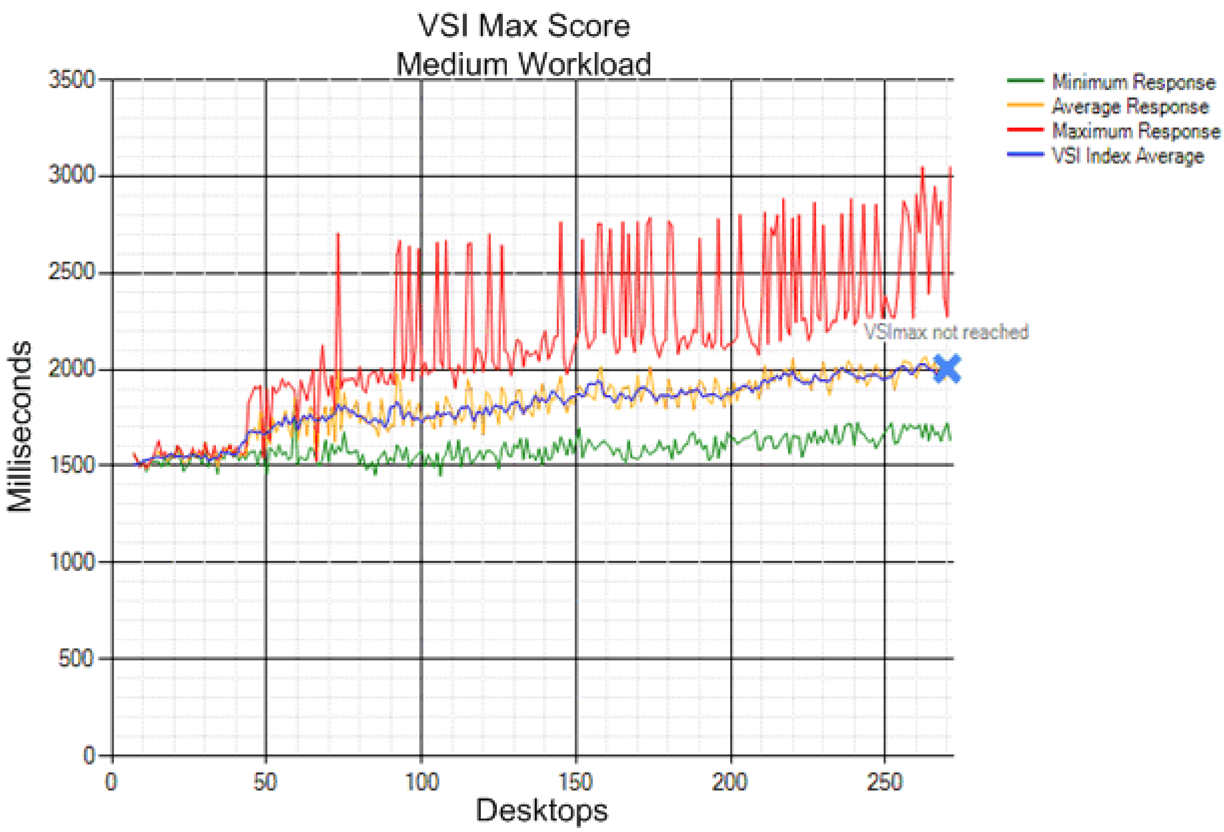

Figure 12

Each storage block is designed for a steady workload of 6480 IOPS and to support two server blocks. Since this test was performed with only one Hitachi Compute Blade 2000 chassis (one server block) the test results should be roughly half the total IOPS the storage block was designed to support.

Figure 13 shows the total IOPS for the desktop environments during the medium VSI test run. The highest usage occurs during the workload portion of the testing, averaging 3767 IOPS.

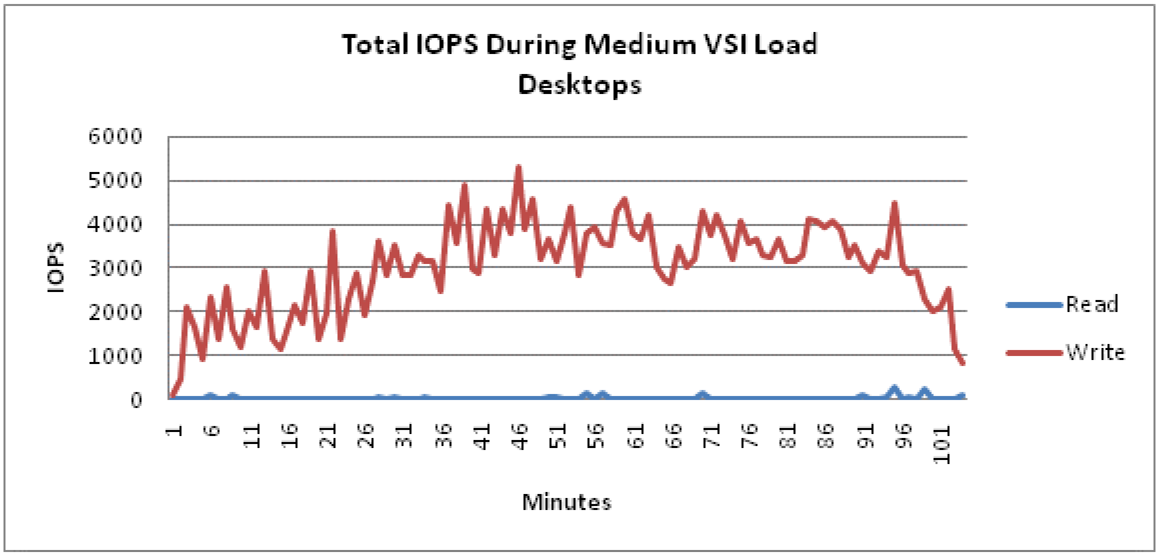

Figure 13

Figure 14 shows the IOPS for the infrastructure environment.

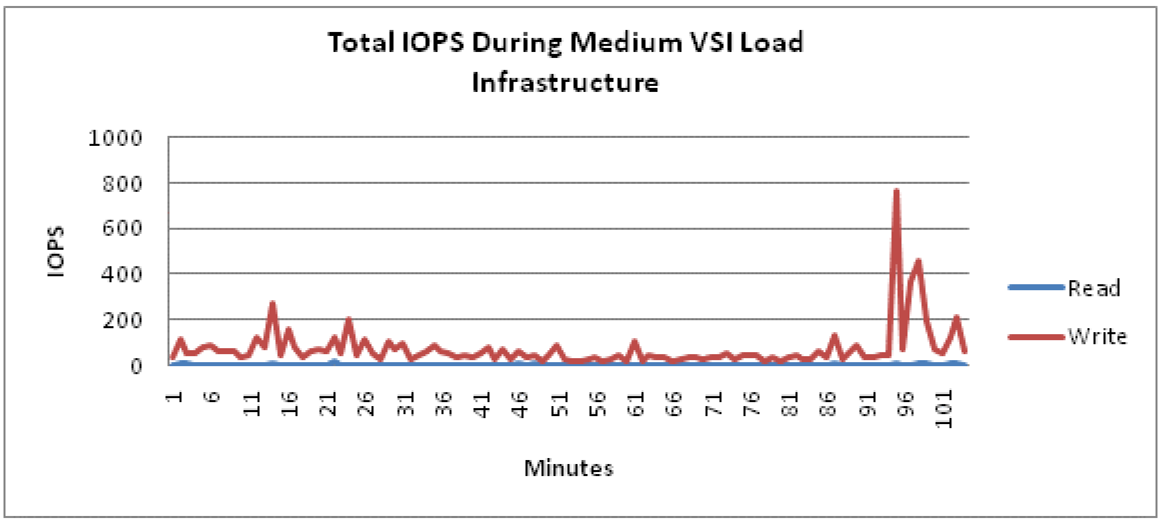

Figure 14

In the infrastructure environment, IOPS peaked slightly during the log off process. Even though the user is logged off, the workstation remains active.

To further evaluate the performance of the environment, CPU usage was monitored during the VSI load testing. Figure 15 shows the CPU usage for the desktop environment.

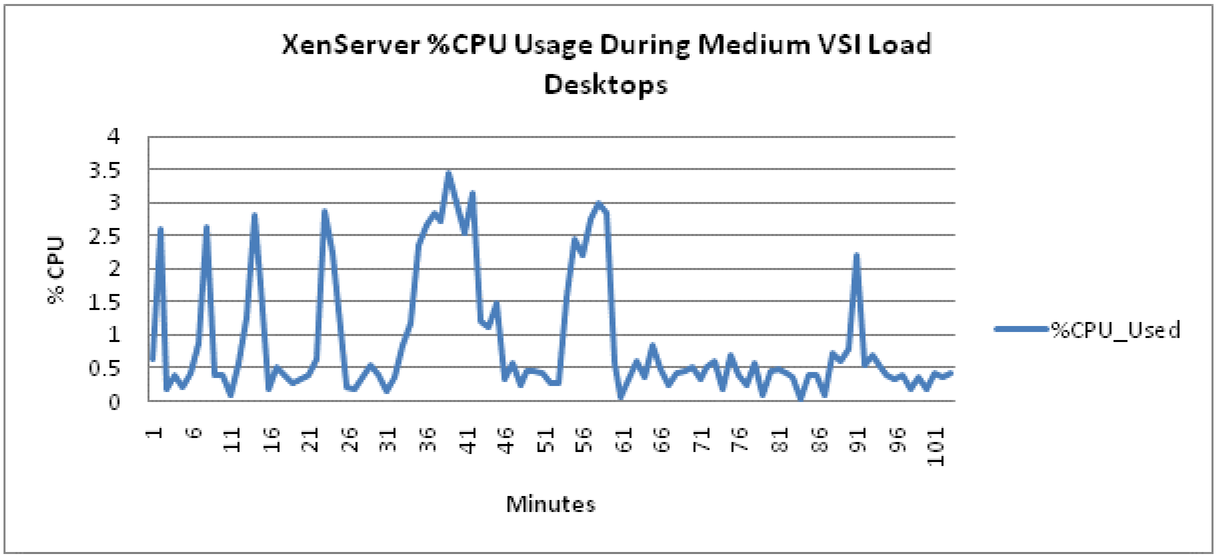

Figure 15

Figure 16 shows the CPU usage in the infrastructure environment during the VSI load testing.

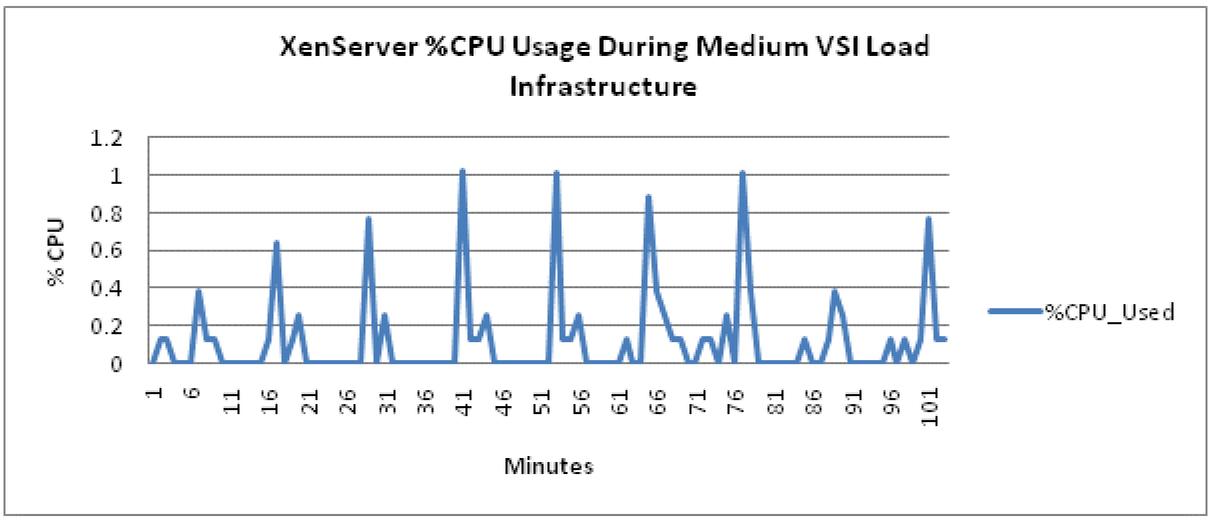

Figure 16

The CPU utilization remains relatively constant throughout the test run. For the desktop blades, the CPU utilization shows that Hitachi Compute Blade 2000 can support more desktops with memory added to each blade. The CPU utilization of the infrastructure blades also shows there are enough CPU resources to support an additional 10 server building blocks. Assume that CPU usage grows proportionally with each server building block added.

Conclusion

Testing shows that Hitachi Virtual Storage Platform and Hitachi Compute Blade 2000 make an ideal platform for a scalable VDI implementation using Citrix XenDesktop running on Citrix XenServer. This reference architecture provides the scalability and performance to deliver a VDI solution.

Following best practices from Citrix, this architecture scales from a 205 user desktop deployment to a 3000 user desktop environment using the shared infrastructure components built into the first design block. With this design, adding to this environment means adding additional blocks of storage, compute blades, and PVS servers.

This document provides proven methods for sizing and testing of your VDI environment to assure the success of VDI deployments. This information helps with selecting the delivery model that best suits your needs. This also provides the tools needed to design a high performance, scalable, and stable VDI environment. Always test your infrastructure in a non-production environment, including the scheduling of the pre-booting of desktops.

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies, and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi Data System products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best in class training on Hitachi products, technology, solutions, and certifications. Hitachi Data Systems Academy delivers on-demand web based training (WBT), classroom based instructor led training (ILT), and virtual instructor led training (vILT) courses. For more information, see the Hitachi Data Systems Academy website.

For more information about Hitachi products and services, contact your sales representative or channel partner, or visit the Hitachi Data Systems website.

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com