

Deploying Citrix XenDesktop 5 with VMware vSphere 4.1 on the Hitachi Virtual Storage Platform

Reference Architecture Guide

By Roger Clark

June 2011

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

Solution Overview	3
Key Solution Components.....	4
Hitachi Virtual Storage Platform	4
Hitachi Dynamic Provisioning Software	5
Hitachi Compute Blade 2000.....	5
Citrix XenDesktop 5.....	6
VMware vSphere 4.1	6
Solution Design	6
Server Building Block	7
Storage Building Block	8
Fibre Channel Design.....	10
Hardware Components	10
Hardware Sizing	11
Storage Architecture.....	13
IOPS Planning.....	16
Environment Storage Architecture	17
Engineering Validation.....	19
Conclusion	23

Deploying Citrix XenDesktop 5 with VMware vSphere 4.1 on the Hitachi Virtual Storage Platform

Reference Architecture Guide

This reference architecture provides planning information to assist in deploying Citrix XenDesktop 5 on VMware vSphere 4.1 using a Hitachi Compute Blade 2000 server and the Hitachi Virtual Storage Platform using a building block approach. It has the details of a deployment of a non-persistent Microsoft Windows 7 virtual desktop infrastructure in the two most common XenDesktop delivery methods, provisioning services (PVS) and machine creation services (MCS). Its purpose is to lower administration costs, improve system efficiency, and maintain performance.

Testing in the Hitachi Data Systems lab shows that the Hitachi Virtual Storage Platform and the Hitachi Compute Blade 2000 make an ideal platform for a scalable virtual desktop infrastructure (VDI) implementation using Citrix XenDesktop running on VMware vSphere. A building block design provides the scalability and performance needed to deliver a VDI solution, scaling from a 384 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 only meant 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 software, Citrix XenDesktop, and VMware vSphere.

Solution Overview

This describes deploying a XenDesktop environment that can scale up to 3000 desktops by using a building block approach. Each new building block can support up to 512 more desktops. Hitachi Data Systems testing used a heavy workload that simulates a scenario of 10 to 15 IOPS per desktop to exercise the building block architecture.

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 the Hitachi Virtual Storage Platform

Proper zoning within the SAN fabric and the use of multipathing software allows for continued operation in the event of a hardware component failure.

Figure 1 shows the topology of this reference architecture.

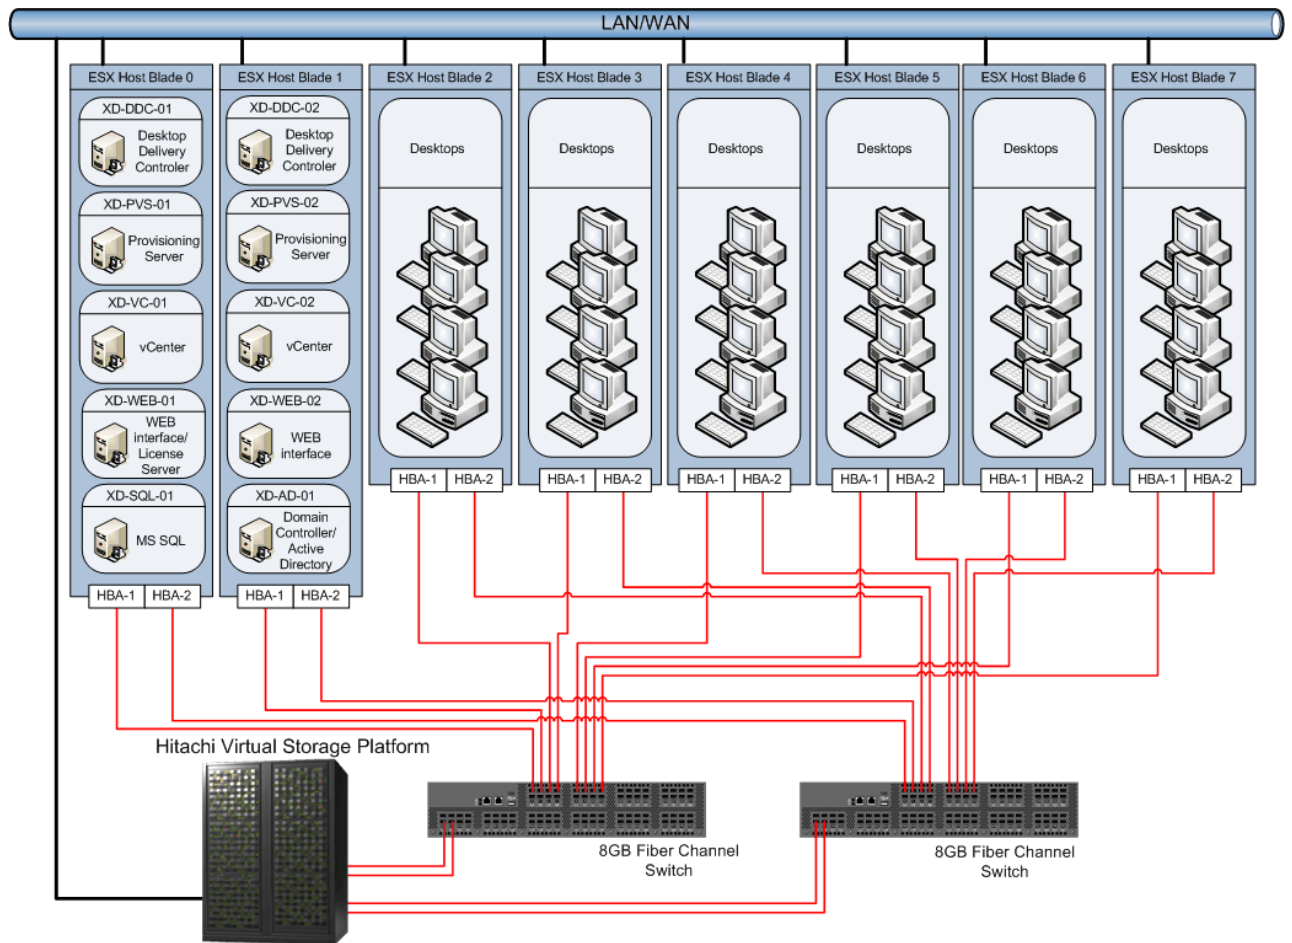


Figure 1

Key Solution Components

Following are the key solution components.

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to scale up, scale out, and scale deep at the same time in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity, and virtualization.

- **Scale Up**—Increase virtual server consolidation, improve utilization of resources, and reduce costs.
- **Scale Out**—Add new physical or virtual servers to your environment to meet business demands.
- **Scale Deep**—Extend the advanced functions of the Virtual Storage Platform to external multivendor storage.

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

Hitachi Dynamic Provisioning

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

Using Hitachi Dynamic Provisioning is similar to the use of a host-based logical volume manager (LVM) without incurring host processing overhead. It uses one or more wide-striping pools across many RAID groups within a Virtual Storage Platform. Each pool has one or more virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB created against it.

Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These come from 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](#).

Hitachi Compute Blade 2000

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

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines Hitachi's logical partitioning expertise 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.

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 workstyles 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 forms of virtual desktops, Machine Creation Services (MCS) and Provisioning Services (PVS).

VMware vSphere 4.1

This reference architecture uses VMware vSphere 4.1 for the hypervisor. VMware describes it this way:

With this release, the VMware virtual datacenter operating system continues to transform x86 IT infrastructure into the most efficient, shared, on-demand utility, with built-in availability, scalability, and security services for all applications and simple, proactive automated management.

VMware vSphere offers features like Distributed Resource Scheduler, vMotion, HA, Storage vMotion, VMFS, and a multipathing storage layer.

Find out [what is new in vSphere 4.1](#).

Solution Design

The reference architecture that follows contains two of the XenDesktop delivery models for non-persistent desktops; PVS and MCS. These are the most common implementations. Your solution may require a different delivery model.

In this solution, the MCS and PVS models were built as two separate environments. If appropriate for your environment, these models can be combined into a single environment or used with other XenDesktop delivery models.

This section provides detailed information on the Citrix XenDesktop environment design used for this reference architecture. Although it includes both application and hardware design information, it focuses primarily on the storage design required to build the basic infrastructure for the XenDesktop environment.

Server Building Block

For ease of management, scalability and predictable performance, this solution uses a building block approach. Each building block in this reference architecture consists of a server block and a storage block.

In the first building block, all eight blades in the Hitachi Compute Blade 2000 chassis have VMware ESX 4.1 running on them. All the blades are using the underlying storage from the Hitachi Virtual Storage Platform to make up this first building block.

Blade 0 and Blade 1 host the following infrastructure VMs.

- Two XenDesktop Desktop Delivery Controllers (DDC)
- Two Citrix Provisioning Service servers (PVS)
- Two Citrix Web Interface servers (WI)
- Two VMware vCenter servers (VC)
- One Microsoft SQL Server 2008 instance

The remaining six blades are dedicated to hosting the desktops.

The first building block hosts 384 desktops. It contains the infrastructure components needed to scale out to 3,000 desktops, based on recommendations from Citrix and VMware. To scale beyond 3,000 desktops, add additional DDCs and VCs. Depending on the load in your environment, additional DDCs and VCs can be added to your first building block or a second infrastructure block can be added. Figure 2 shows the difference between the first building block and additional building blocks.

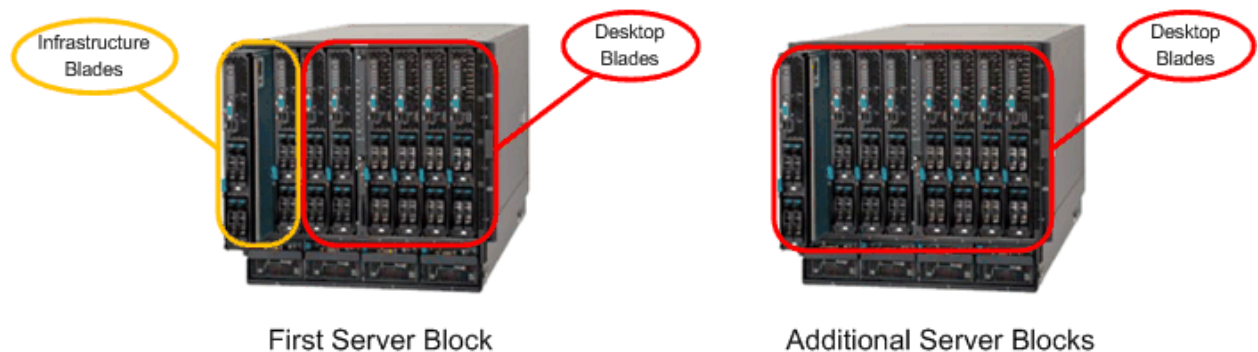


Figure 2

All building blocks added later consist of one storage block and eight compute blades. Each added building block uses the infrastructure in the first building block and can support up to 512 additional desktops, based on the compute blade configuration used in the first building block.

Note—For PVS environments, each subsequent building block requires adding an additional PVS server to blade 0 and blade 1 of the infrastructure on the first building block.

Storage Building Block

The first storage block for this solution consists of two separate Hitachi Dynamic Provisioning pools.

- **Infrastructure pool**—The infrastructure storage block consists of 12 300GB 10K RPM SAS drives in a RAID-5 (3D+1P) Hitachi Dynamic Provisioning pool.
- **Desktop pool**—The desktop storage block consists of 32 146GB 15K RPM SAS drives in a RAID-1+0 (2D+2D) Hitachi Dynamic Provisioning pool. It is based on the following:
 - A non persistent desktop with an I/O profile of 10 to 15 IOPS per desktop
 - A per user cache of 5GB

To scale out the solution, add additional storage to the desktop pool. Figure 3 shows the difference between the first building block and additional building blocks.

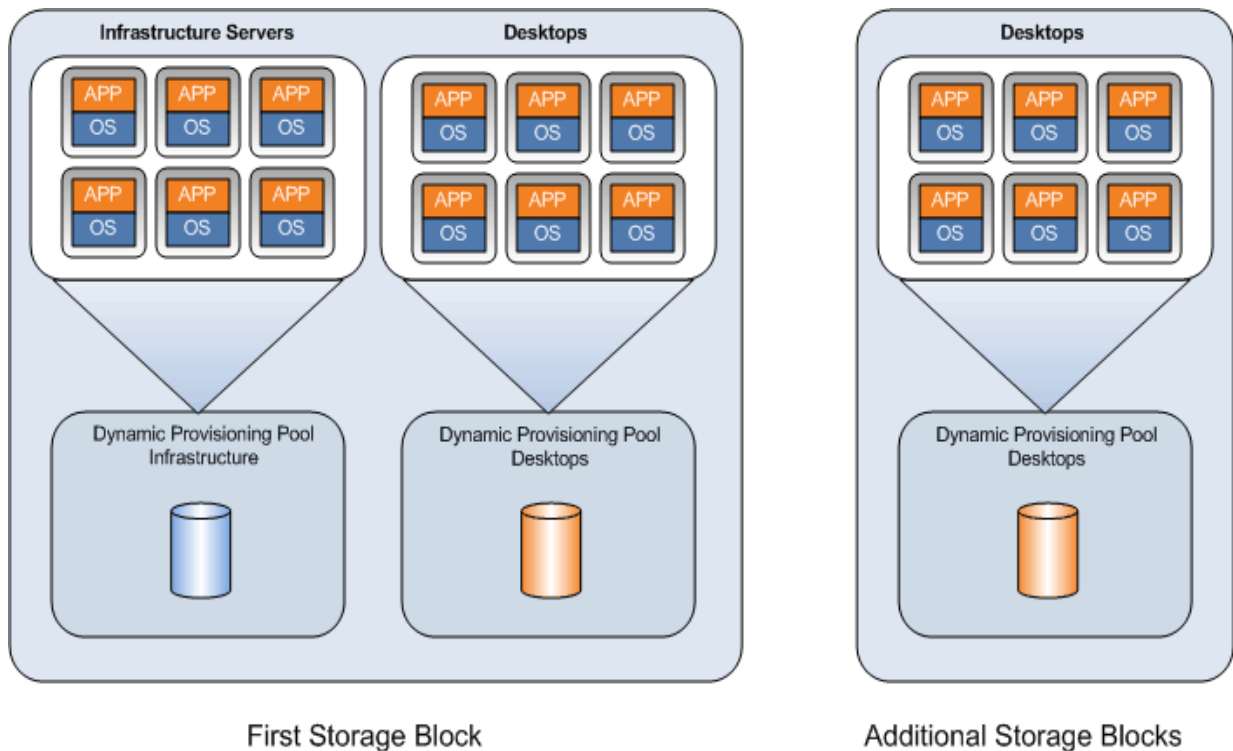


Figure 3

The first server and storage block in this solution is designed to support 384 concurrent desktops. To increase the number of desktops in this solution, add additional server and storage blocks.

Figure 4 shows how the solution can be scaled out to 3,000 desktops by adding additional server and storage building blocks to the environment.

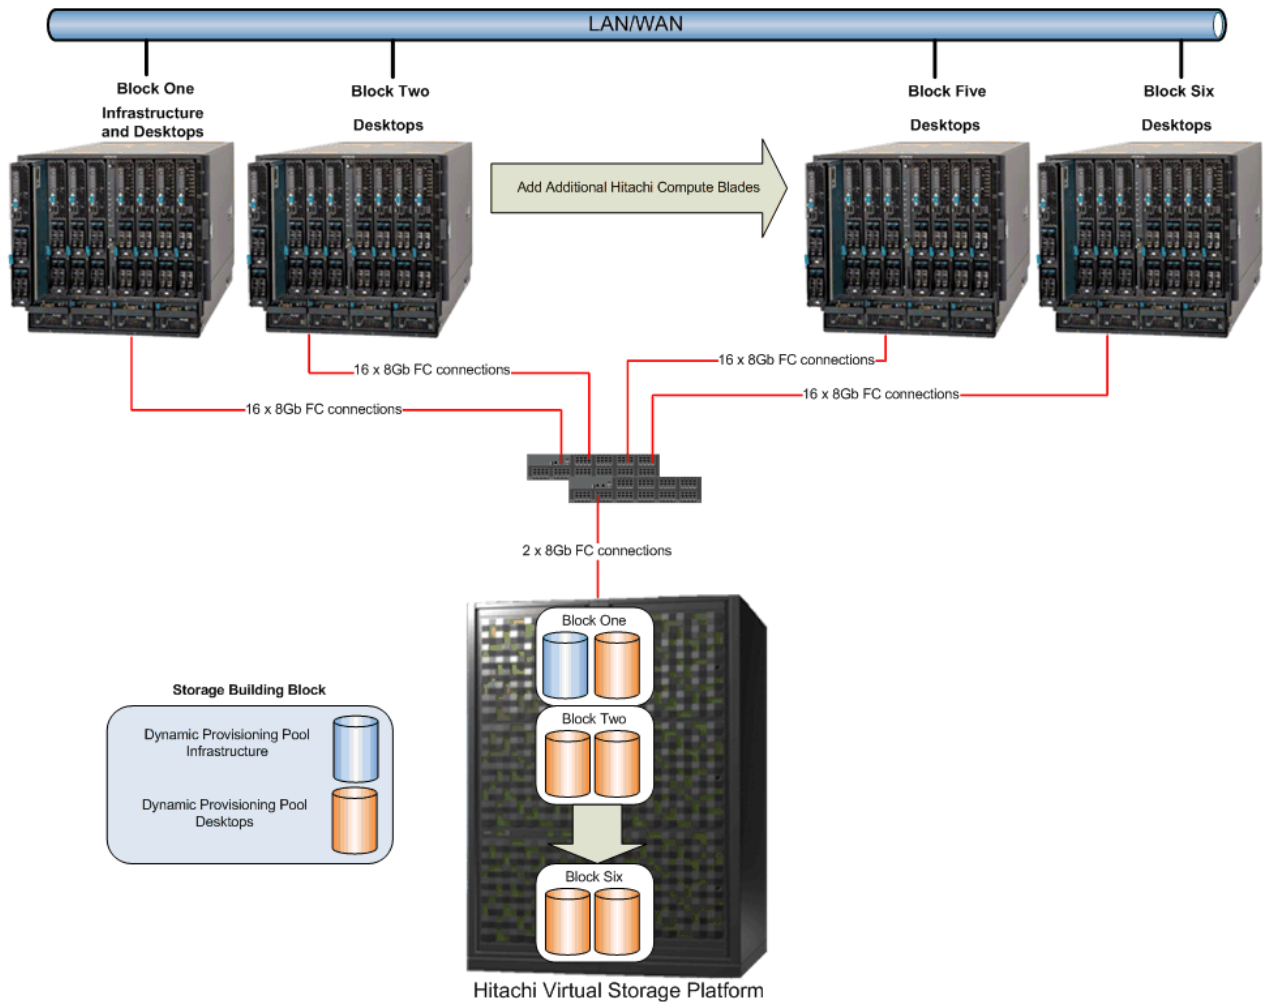


Figure 4

The first storage block of the environment consists of one 1.5TB LUN for the infrastructure VMs and user profiles. This LUN contains extra space to support more user profiles while adding desktops and additional PVS servers for the environment.

Monitor the Hitachi Dynamic Provisioning pools for any performance degradation as you scale out your environment. Depending on the user load, some environments may benefit from multiple LUNs for the infrastructure storage to prevent I/O queuing.

Because the solution uses non-persistent desktops that will be rebooted on a regular basis, there is over-provisioning of the LUNs hosting the virtual desktops in the MCS environment. This lowers the initial cost of storage while still providing the required performance.

During testing, the amount of storage used by the virtual desktops in the MCS environment never exceeded more than 33% of the 800GB capacity of the LUNs because of the following:

- The thin provisioning capabilities of vSphere
- The user cache is not saved when rebooting the desktop

Rebooting a pooled non persistent desktop destroys the write cache. This results in the target device starting with the same image each time.

If you use over provisioning, set up proper threshold alerting and to check your storage resources on a regular basis. To prevent outages, take immediate actions upon reaching a threshold, such as adding more disks to your resource pools or reducing the number of virtual desktops for the pool.

Do not use over provisioning when using persistent desktops.

Base your decision whether or not to over provision your environment on the type of desktops in use, workstation reboot schedules, and your company's IT policies.

Fibre Channel Design

This solution connected the Hitachi Compute Blade 2000 chassis and the Hitachi Virtual Storage Platform with two Fibre Channel switches for redundancy and high availability purposes. Figure 5 shows the Fibre Channel connections for this solution.

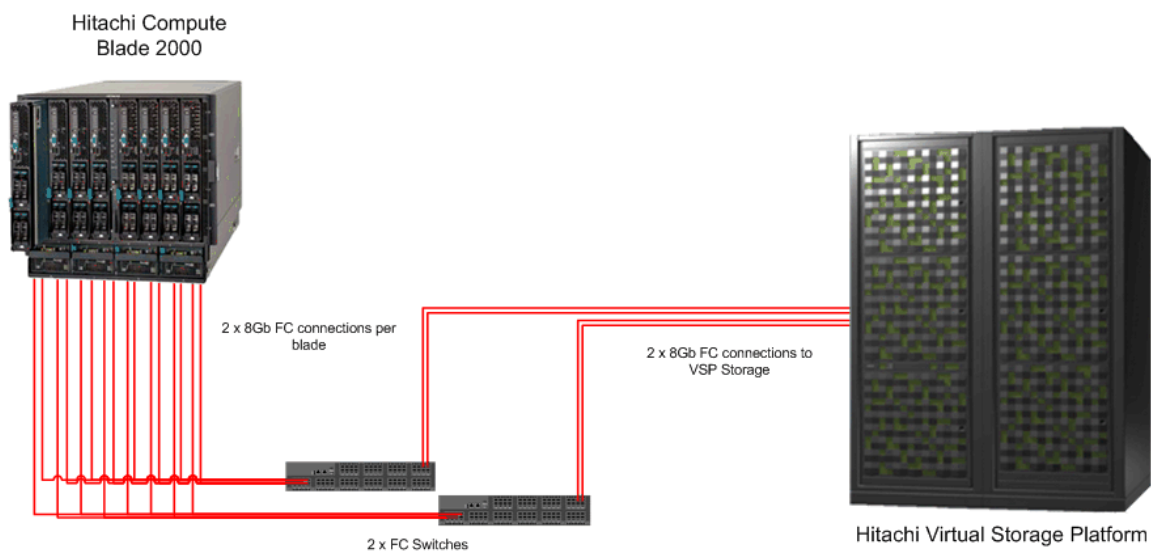


Figure 5

Another option not used in this reference architecture is to deploy the environment using an enterprise-level director that contains multiple blades that can support high availability and redundancy.

Hardware Components

The hardware for this solution infrastructure consists of the following:

- A single Hitachi Compute Blade 2000 chassis with eight X55A2 blades
- A single Hitachi Virtual Storage Platform
- Two Brocade 5000 Fibre Channel switches

Table 1 describes the hardware components used in this reference architecture.

Table 1. Hardware Components in this Reference Architecture

Hardware	Detail Description	Version	Quantity
Hitachi Virtual Storage Platform storage system	80GB cache memory 32 x 8Gb/sec ports 72 x 146GB, 15K RPM, SAS disks 184 x 300GB 10K RPM, SAS disks	70-02-00-00/7	1
Hitachi Compute Blade 2000 chassis	2 x 8Gb/sec Fibre Channel switch modules 4 x 1Gb/sec Ethernet switch modules	A0160-G-5666	1
Hitachi X55A2 server blades	2 x six core Intel CPU 72GB RAM 2 x 36GB 15K SAS disks	03-79	8
Brocade 5000 switch	SAN Switch with 8Gb/sec Fibre Channel ports	FOS 5.3.1a	2

Hardware Sizing

Determining the number of desktops that can be supported by your specific hardware is dictated by a number of factors, such as:

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

Most XenDesktop implementations have a number of different user categories:

- Light
- Normal (most users)
- Power
- Heavy

These categories can be used to help 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 to heavy workload profile. During the testing phase, a heavy test load was used to stress the desktops and the environment.

This reference architecture uses industry-standard formulas to help determine the recommended number of desktops that could be run on the hardware configuration in the Hitachi Data Systems lab.

Based on the breakdown of the number of users in each category, and following the CPU and memory allocation guidelines, Table 2 provides recommendations used to calculate resource allocation for desktops using Windows XP and Windows 7.

Table 2. 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
Normal	Windows XP	1	1GB to 1.5GB	6-10	8 to 10
	Windows 7	1	1.5GB to 2GB	8-12	6 to 8
Power	Windows XP	1	1.5GB to 2GB	12-16	6 to 8
	Windows 7	1 to 2	2GB to 3GB	15-25	4 to 6
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 take into account start ups, logons, and logoffs.

Use the guidelines in Table 2 with the following two formulas to estimate the scalability of a particular physical server.

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

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

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 an example, consider the following:

- All users are split evenly between normal and power users on Windows 7
- The server configuration has 64GB of memory and has 8 cores

The values used from Table 2 are:

- Average of 2.5 GB of RAM per user
- Average of 7 users per core

The most number of users possible from a memory perspective is 23 users.

$$23 \text{ Users} = \frac{(64\text{GB} - 6\text{GB})}{(2.5\text{GB}/\text{User})}$$

The most number of users possible from a core perspective is 49 users.

$$49 \text{ Users} = (8 \text{ cores} - 1 \text{ core}) \times (7 \text{ Users/Core})$$

With the given configuration, 23 is the recommended maximum number of users.

One way to make the number of users from a memory and core roughly equal is to reconfigure the server with 126GB of memory.

$$48 \text{ Users} = \frac{(126\text{GB} - 6\text{GB})}{(2.5\text{GB} / \text{User})}$$

In the Hitachi Data System test environment, each Hitachi Compute Blade 2000 server blade was configured with 72GB of memory. This limited the number of desktops per server blade to 64.

Storage Architecture

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

Storage Sizing

The capacity requirements for a XenDesktop environment are dictated by many variables such as the following:

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

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

Provisioning Services Storage

When using 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. The PVS streams the contents of the vDisk to the target device on demand, in real time, using software-streaming technology.

Once the vDisk image is available from the network, a virtual machine (VM) on a target device no longer needs its 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. This file is deleted upon each restart cycle.

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

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

This reference architecture uses a cache file of 5GB to represent a heavy user that may not reboot the desktop every day.

There are several options for storing the cache file for provisioned desktops. Two common locations are target-based physical local storage and client disk (virtual machine vDisks). The benefits and limitations of each option are as follows:

- **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 RAM. If the cache file is expected to grow larger than the amount of available physical RAM in the device, use the hard drive or an enterprise storage device. Filling a RAM-based cache file, undesired result or errors may occur.
 - 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 the enterprise storage associated with each VM or the allocated RAM for the VM for the location of the cache. For example, created each VM 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 VM requires space equal to the allocated size of each cache file. This is in addition to the size required for each unique VM.
 - This option provides added resiliency. If a cache file for a single virtual machine runs out of space, only that virtual machine is affected.
 - 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 VM.

Figure 6 shows the file layout for the PVS model.

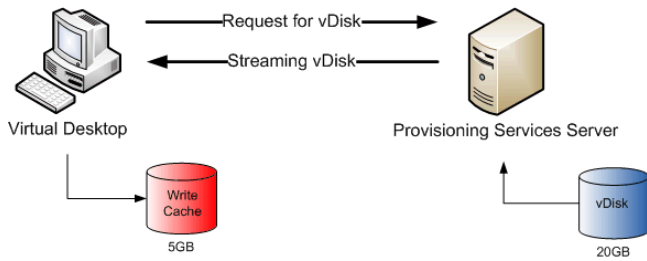


Figure 6

To estimate the storage required for a PVS model:

Storage Required

$$= (\text{Number of Users} \times \text{Write Cache Size}) + (\text{Number of Desktop Images} \times \text{Desktop Image Size})$$

For example, to calculate the space required for 384 non-persistent PVS desktops using one 20GB master image and a 5GB write cache:

$$1986560\text{MB} = (384 \text{ users} \times 5120\text{MB}/\text{user}) + (1 \text{ image} \times 20480\text{MB}/\text{image})$$

1986560MB is approximately 2TB.

Each of these options has benefits and limitations. It is important to evaluate your specific requirements before determining the location of the cache file.

For this reference architecture, the cache file was stored on the client vDisk.

Machine Creation Services Storage

The storage in the MCS environment works much differently than in the PVS environment. Understand how the MCS creates desktops and how the user cache works to properly size the storage for a MCS environment before using it. It is beyond the scope of this document to give this instruction.

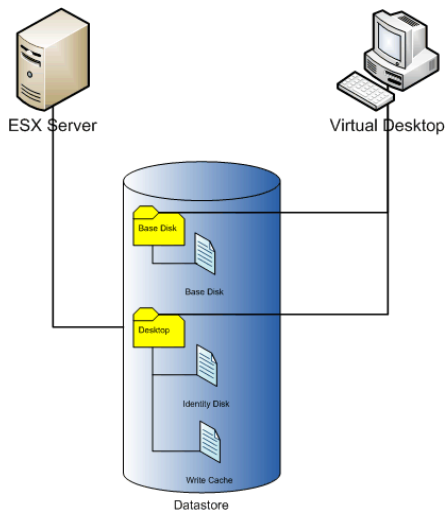


Figure 7

The base disk is a thin-provisioned copy of the master VM used to create the desktop image. If you provision a master VM with 20GB for the OS but only use 16MB of the disk, then your base disk provisions 20GB with only 16MB actually stored.

To calculate the amount of storage needed for a non-persistent desktop MCS environment use the following formula:

Desktop Pool Space Needed

$$= \{ [(Identity\ Disk\ Size)/User] \times (Number\ of\ Users) + [(Write\ Cache\ Size)/User] \times (Number\ of\ Users) \} + \{ [(Master\ Image\ Provisioned\ Size)/(Master\ Images \cdot Datastores)] \times (Number\ of\ Master\ Images) \times (Number\ of\ Datastores) \}$$

For example, here is how to calculate the space required for the following:

- 384 non-persistent MCS desktops for users
- One 20GB provisioned master image per master image per datastore (20480MB)
- 16MB identity disk size per user
- 5GB write cache per user (5120MB)
- Three datastores

Desktop Pool Space Needed

$$= \left[\left(\frac{16MB}{User} \right) \times (384\ Users) + \left(\frac{5,120MB}{User} \right) \times (384\ Users) \right] + \left[\left(\frac{20,480MB}{[Master\ Images \cdot Datastore]} \right) \times (1\ Master\ Image) \times (3\ Datastores) \right]$$

$$Desktop\ Pool\ Space\ Needed = 2,033,664MB$$

The result of 2,033,664MB means approximately 2TB of desktop pool space needed.

You might also need to consider additional storage requirements depending on the features you plan to install in your XenDesktop environment. This reference architecture guide focuses on the components that an enterprise-class XenDesktop environment must have to support a standard environment.

IOPS Planning

No precise method exists to establish VDI IOPS requirements. Like other production server environments, the best practice is to deploy a test environment first to establish whether your VDI environment can support your IOPS requirements.

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

- Determine an accurate IO profile of the current desktops in your environment. Use an I/O monitoring tool to assess your current environment.
- Consider the times that users will be logging onto their desktops and start up their desktops. This is to reduce the effects of boot storms.
- Consider using a server-based antivirus solution, such as VMware V-Shield. This reduces excess I/O on a user desktop.

The following formulas can be used to calculate the approximate functional IOPS available for a VDI workload.

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

Functional IOPS

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

This reference architecture used a standard of 80% write operations and 20% read operations. This user I/O profile simulates a user running common office applications.

For example, to calculate the functional IOPS for the following:

- 146GB 15K RPM SAS drives capable of 150 IOPS per drive
- A total of 36 of these drives
- RAID-1+0 configuration

$$\text{Total Raw IOPS} = (150 \text{ IOPS/drive}) \times (36 \text{ Drives})$$

$$\text{Total Raw IOPS} = 5400 \text{ IOPS}$$

$$\text{Functional IOPS} = \left\{ \frac{[(5400 \text{ IOPS}) \times (0.8)]}{2} \right\} + \{(5400 \text{ IOPS}) \times (0.2)\}$$

$$\text{Functional IOPS} = 3240 \text{ IOPS}$$

Using this approach, the test environment for this reference architecture was built to handle 3240 IOPS under a sustained workload.

Environment Storage Architecture

The storage design used for this environment consists of:

- 36 × 146GB 15K RPM SAS drives configured as RAID-1+0 (2D+2D) hosting all of the desktop VMs, the desktop base images, and the users write cache in the PVS environment
- 36 × 146GB 15K RPM SAS drives configured as RAID-1+0 (2D+2D) hosting all of the desktop VMs, the desktop base images, and the users write cache in the MCS 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 VMs

Table 3 shows the dynamic provisioning pools used in this white paper

Table 3. Dynamic Provisioning Pools

HDP Pool Name (ID)	RAID Group Configuration	Drive Type	Number of RAID Groups	Pool Capacity (TB)
HDP-XD-D001	RAID-1+0 (2D+2D)	36 × 146GB 15K RPM SAS	9	2.34TB
HDP-XD-D002	RAID-1+0 (2D+2D)	36 × 146GB 15K RPM SAS	9	2.34TB
HDP-XD-U001	RAID-5 (3D+1P)	12 × 300GB 10K RPM SAS	3	2.35TB

Infrastructure Storage Overview

The storage block for the infrastructure consists of a 1.5TB LUN carved out of a 2.35TB RAID-5 dynamic provisioning pool. The purpose of the additional infrastructure storage space is to allow for the future growth of the environment. This infrastructure storage handles the additional user profiles when the environment is scaled up to 3,000 desktop users and any additional provision services servers needed in the PVS environment.

Table 4 shows the LDEV created for the infrastructure virtual machines and user profiles.

Table 4. Infrastructure LDEV configuration

LDEV Name (ID)	LDEV	Size (TB)	Purpose	Storage Port
USP-LDEV-01	00:00:12	1.5	Infrastructure VMs and user profiles	3D/4D

Machine Creation Services (MCS) Storage Overview

The storage block for the MCS environment consists of a 2.34TB RAID 1+0 dynamic provisioning pool carved into three 800GB LUNs that stores all XenDesktop VMs as well as the desktop base image.

Table 5 show the LDEVs created for the MCS storage blocks in the test environment.

Table 5. MCS LDEV Configuration

LDEV Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
MCS-LDEV-01	00:00:24	800	XenDesktop golden Image Desktop VMs	3D/4D
MCS-LDEV-02	00:00:37	800	XenDesktop golden Image Desktop VMs	3D/4D
MCS-LDEV-03	00:00:38	800	XenDesktop golden Image Desktop VMs	3D/4D

Provisioning Services (PVS) Storage Overview

The storage block for the PVS environment consists of a 2.34TB RAID-1+0 dynamic provisioning pool carved into three 800GB LUNs that store all XenDesktop VMs and a 220GB LUN that stores the master PVS images. The size of the LUN used for the master PVS images varies depending on the size and number of master PVS images each environment needs.

Table 6 show the LDEVs created for the PVS storage blocks in the test environment.

Table 6. PVS LDEV Configuration

LDEV Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
PVS-LDEV-01	00:00:33	800	Desktop VMs User write Cache	3D/4D
PVS-LDEV-02	00:00:34	800	Desktop VMs User write Cache	3D/4D
PVS-LDEV-03	00:00:35	800	Desktop VMs User write Cache	3D/4D
PVS-Images	00:00:10	220	PVS Golden Image	3D/4D

Engineering Validation

After building the reference architectures in the Hitachi Data Systems lab, a series of tests were run to validate the designs. The testing was done using Logon Consultants VSI software. VSI is third party load generation software designed to replicate a multiple user workload and measure session and application response time. To collect performance data on the infrastructure, VMware's esxtop was used to collect data from all of the ESX servers. The testing consisted of two test scenarios, desktop start ups and workload performance.

The desktop start up test simulates a start of business day or shift change scenario. All or some of the desktops are pre started prior to any workers arriving to log on using the XenDesktop power management feature. This helps prevent boot storms that can adversely affect the user's experience.

The workload performance test simulates the load created in an environment when users log on and start running a measured workload. This test was done using the Logon Consultants VSI software. This validates the calculations used to build the environment.

Boot-up Testing Results

A boot storm occurs when a large number of desktops are started in a relatively short amount of time. This causes a spike in IOPS, exceeding the number of calculated IOPS for sustained workloads.

XenDesktop has a feature to help prevent boot storms. This feature schedules the booting of desktops before they are needed. VDI administrators may set a pre determined number of desktops to be booted and available for use.

This test shows the time required to pre-boot all 384 desktop in the test environment.

Determine how many desktops that need pre-booting in your environment and schedule the desktop boot times to meet those needs.

The desktop booting test followed these steps:

1. The desktop group was powered down and placed in maintenance mode.
2. The desktop group was taken out of maintenance mode.
3. XenDesktop sent power up commands to the hypervisor to start the first batch of six desktops all at one time.
4. After the first batch of desktops was started, XenDesktop sent powers up commands to the next batch of six desktops.
5. XenDesktop repeated sending power up commands to the next batch of six desktops when the previous batch finished until the designated number of desktops were started.

With the XenDesktop power management feature set to boot all 384 user desktops at one time, Figure 8 shows the boot times for the MCS and PVS environments.

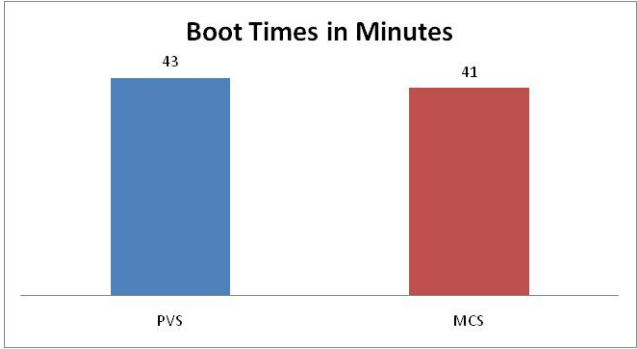


Figure 8

This test determines the amount of time needed to pre-boot a given number of desktops in your VDI environment.

Another key benefit of using the XenDesktop power management function is 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 XenDesktop power management feature set to boot all 384 desktops, Figure 9 and Figure 10 show the IOPS during the boot process for the MCS and PVS environments.

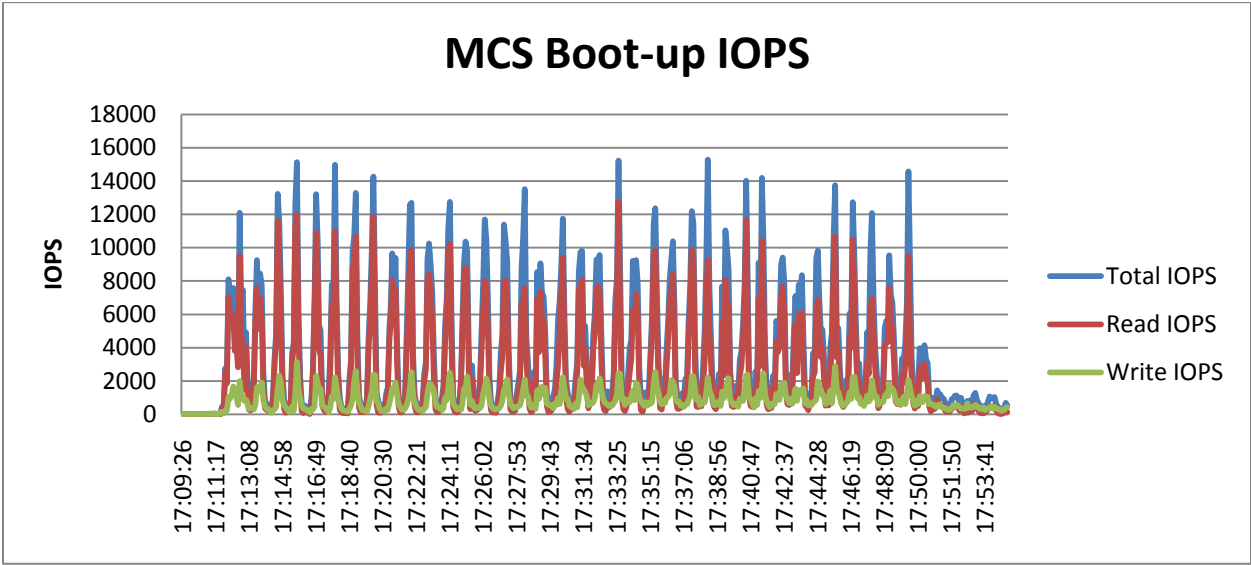


Figure 9

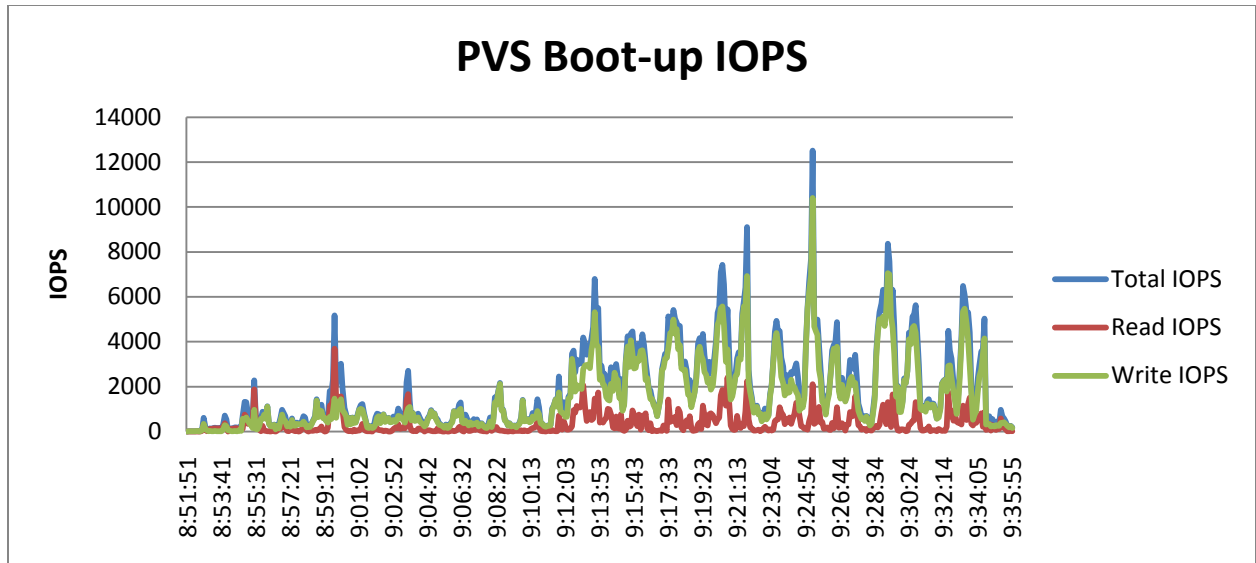


Figure 10

Figure 9 and Figure 10 show that using the power management function in XenDesktop spreads out the IOPS evenly. No major spikes from power up adversely affect the performance of any user already logged on to the environment.

Load Testing Results

The purpose of the load test is to validate that the reference architecture environment can exceed the calculated number of IOPS for the environment, yet continue to maintain a 2000 millisecond response time for the desktops.

For the load test:

1. All user desktops are started using twelve VSI launch agents running on VMs in a separate environment. All the desktops were started and in a ready state prior to starting the load generation test to simulate the desktops being pre-started prior to workers starting work.
2. A heavy work cycle is started on each desktop as it is logged on. Each heavy work cycle consists of creating and editing multiple office documents, running multimedia files, browsing Internet contents, and generating PDF files. Each user desktop continues to repeat this work cycle until all 384 desktops have logged on.
3. Once all the desktops have been logged on, each user desktop completes its current work cycle and then logs off.

The log ons for the load generation test were spread out over an hour to simulate users arriving and logging on at different times. An average of 6 desktops logged on each minute over the hour.

When using the heavy workload, the reference architecture implementation handled a peak load of 8635 IOPS in the MCS environment and 4703 IOPS for the PVS environment while still maintaining the 2000 millisecond response time for the desktops.

In the MCS environment, the heavy work load generated a peak of 8635 IOPS, with a read peak of 5237 IOPS and write peak of 4620 IOPS. Figure 11 shows the IOPS for the MCS load test.

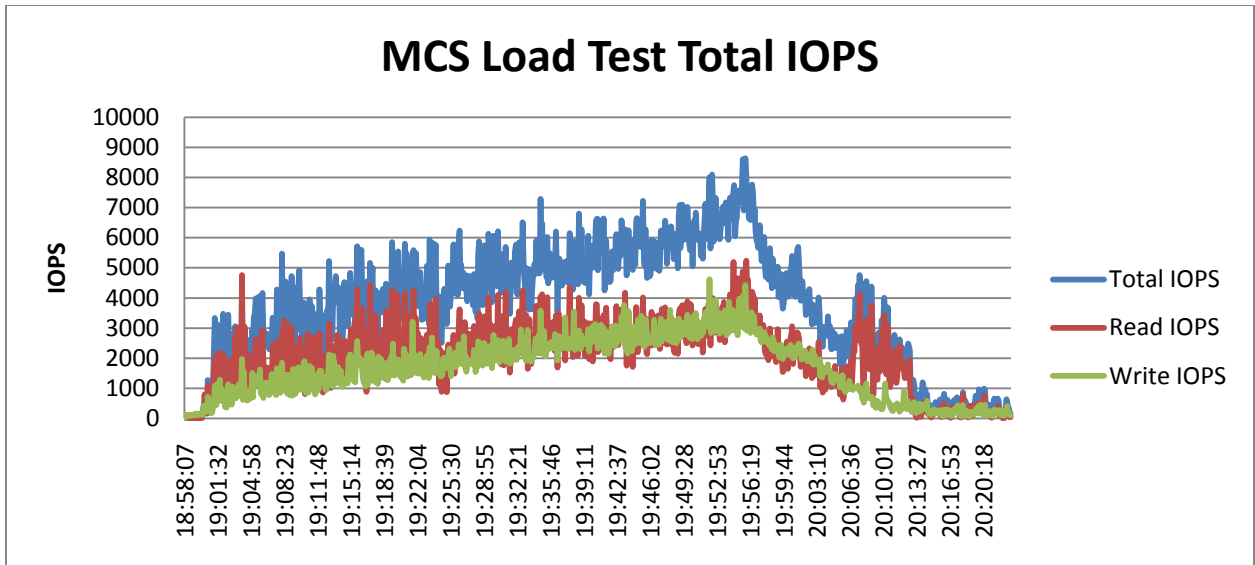


Figure 11

In the PVS environment, the heavy workload generated a peak of 4703 total IOPS, with a read peak of 3261 IOPS and write peak of 3857 IOPS. Figure 12 shows the IOPS for the PVS load test.

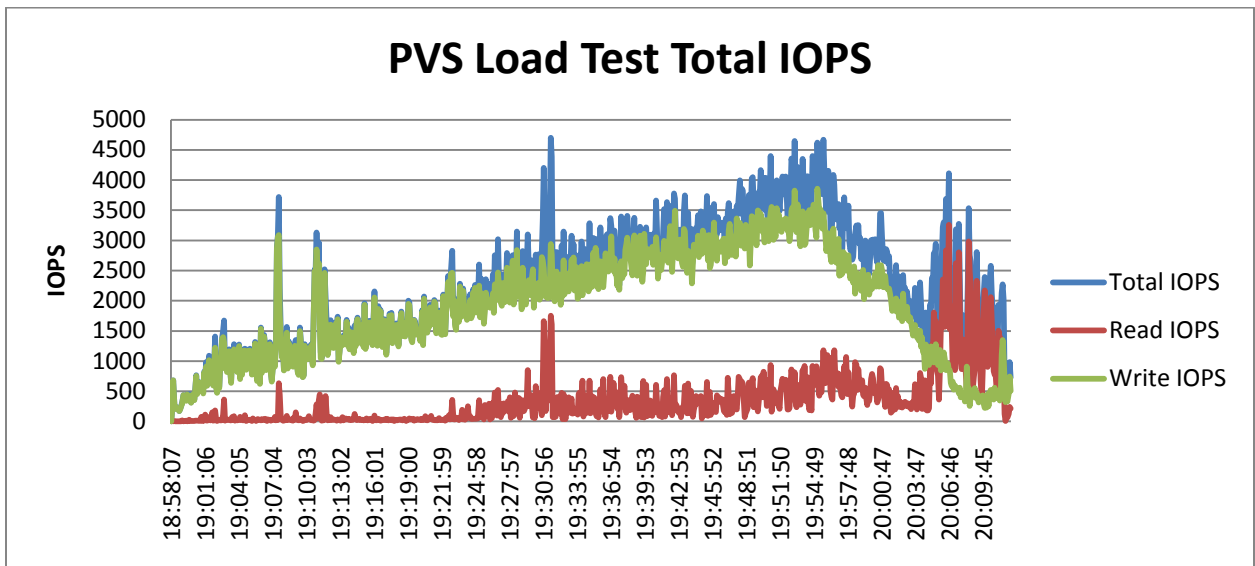


Figure 12

This test shows that the design handles a medium to heavy workload and still maintains a 2000 millisecond response time for the desktops. While this test showed that the design was able to handle the higher IOPS, do not over prescribe your VDI environment.

Conclusion

Testing shows that the Hitachi Virtual Storage Platform and the Hitachi Compute Blade 2000 make an ideal platform for a scalable VDI implementation using Citrix XenDesktop running on VMware vSphere. The design used in the paper provides the scalability and performance needed to deliver a VDI solution, scaling from a 384 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 only means adding additional blocks of storage and compute blades.

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 and provides the tools needed to design a high performance, scalable, and stable VDI environment.

Test your environment to determine the effect of large numbers of desktops starting up in your environment. 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.

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 web site](#).

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 web site](#). 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 web site](#).

For more information about Hitachi products and services, contact your sales representative or channel partner, or visit the [Hitachi Data Systems web site](#).

 **Hitachi Data Systems Corporation**

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-090-00 June 2011

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