TECHNICAL PAPER

HITACHI
Inspire the Next

# Deploy Hitachi Unified Compute Platform Select for Citrix XenDesktop with VMware vSphere

Reference Architecture Guide

*By Roger Clark, Henry Chu*

*January 11, 2013*

Hitachi Data Systems

# Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

# Table of Contents

# Deploy Hitachi Unified Compute Platform Select for Citrix XenDesktop with VMware vSphere

## Reference Architecture Guide

This reference architecture guide describes a solution using Hitachi Unified Compute Platform Select for Citrix XenDesktop with VMware vSphere 5.0. This solution scales in a repeatable, building block design to deploy desktops using Citrix XenDesktop and Citrix Provisioning Services. This solution provides information to plan and deploy a Citrix XenDesktop 5.5 environment using the following:

- Hitachi Unified Storage 150
- Hitachi Compute Blade 500
- Brocade Ethernet and Fibre Channel networking components
- VMware vSphere 5.0 as the hypervisor

Hitachi Unified Compute Platform is a family of integrated and flexible reference solutions. Each Unified Compute Platform solution, configured for immediate deployment, runs top tier infrastructure applications without over-purchasing or provisioning unnecessary equipment. The entire solution is stack certified and compatible.

You need a working familiarity with techniques and practices used for the products listed in this guide.

**Note** — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

# Solution Overview

This reference architecture uses the following:

- Hitachi Unified Storage 150
- Hitachi Compute Blade 500
- Hitachi Dynamic Provisioning
- Citrix XenDesktop 5.5
- Citrix Provisioning Services 5.6.2
- VMware vSphere 5.0
- Brocade 6510 enterprise fabric switches
- Brocade VDX-6720 Ethernet switches

## Determining User Workload

An important factor when sizing a Citrix XenDesktop environment is defining the typical workload profile for the end-users that will be using the environment. The workload used for sizing this solution is based on industry standards and Login VSI.

The light workload definition used in this solution represents an average task-based user running a common set of applications.

Table 1 shows the applications used to define the workload types.

**Table 1. User Application Definitions Used for Testing**

| *User Workload* | **Light** | **Medium** | **Heavy** |
|---|---|---|---|
| | ■ Used in this solution | ■ Not used in this solution | ■ Not used in this solution |
| *Workload* | ■ Microsoft Internet Explorer<br>■ Microsoft Word<br>■ Microsoft Excel<br>■ Microsoft Outlook | ■ Microsoft Internet Explorer<br>■ Microsoft Word<br>■ Microsoft Outlook<br>■ Microsoft PowerPoint<br>■ Adobe Acrobat Reader and PDF printer<br>■ 7-Zip | ■ Microsoft Internet Explorer<br>■ Microsoft Excel<br>■ Microsoft PowerPoint<br>■ Adobe Acrobat Reader and PDF printer<br>■ 7-Zip<br>■ Media Player |

The desktop configuration used for this solution followed the recommendations in Table 2 for a light user workload. See Table 2 for a description of the different workloads.

**Table 2. User Desktop Configurations**

| *User Workload* | **Light** | **Medium** | **Heavy** |
|---|---|---|---|
| *Operating System* | Microsoft Windows 7, 64-bit | Microsoft Windows 7, 64-bit | Microsoft Windows 7, 64-bit |
| *vCPU Allocation* | 1 | 1 | 2 |
| *Memory Allocation per User* | 1 GB to 2 GB | 1.5 GB to 2 GB | 4 GB |
| *Average Steday State IOPS** | 4 to 7 | 8 to 12 | 12 to 25 |
| *Estimated Users per Core* | 8 to 10 | 6 to 8 | 4 to 6 |

*The IOPS estimates are an average of when the user is logged on to the virtual desktop and working. It does not take into account start-ups, users logging on, and users logging off.

**Note** — Hitachi Data Systems recommends that you do in-depth testing to determine the correct workloads for the end-users using your environment.

## Logical Design

Figure 1 on page 5 illustrates the high-level logical design of this solution with the following:

- One infrastructure cell for Hitachi Unified Compute Platform Select management
- One application cell for Citrix provisioning services
- One application cell for Citrix XenDesktop

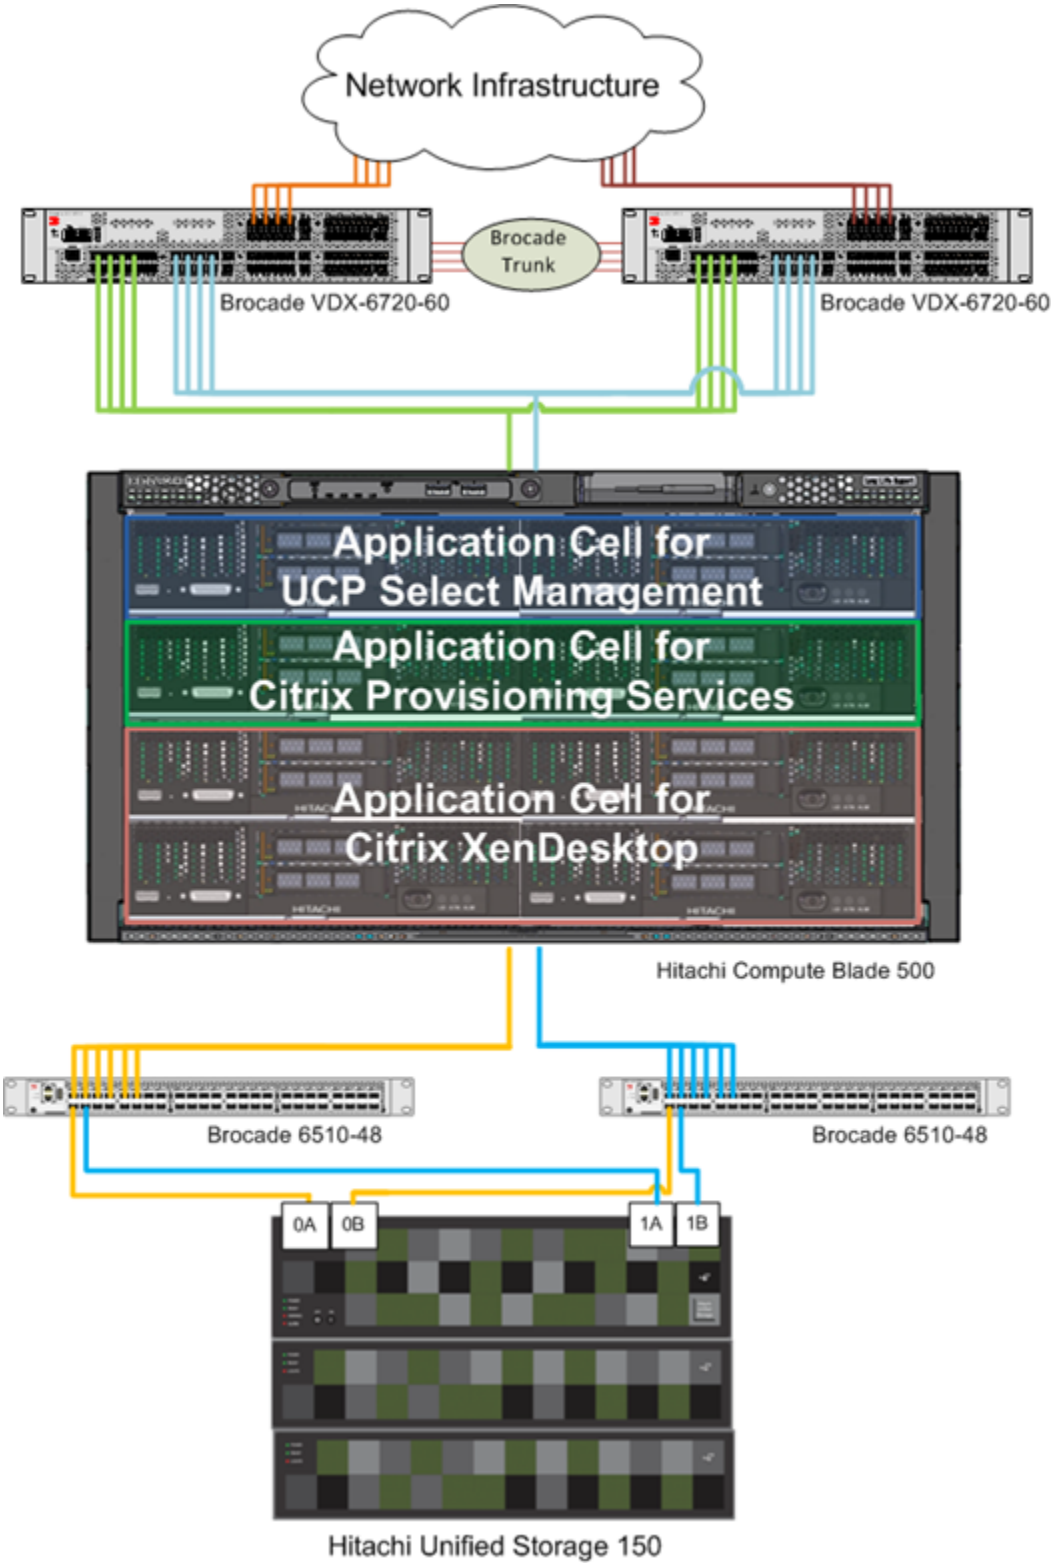
Network Infrastructure
Brocade Trunk
Brocade VDX-6720-60
Brocade VDX-6720-60
Application Cell for UCP Select Management
Application Cell for Citrix Provisioning Services
Application Cell for Citrix XenDesktop
Hitachi Compute Blade 500
Brocade 6510-48
Brocade 6510-48
0A
0B
1A
1B
Hitachi Unified Storage 150

**Figure 1**

This solution supports up to 500-streamed desktops designed with the following:

- Microsoft Windows 7, 64-bit
- 1 vCPU
- 2 GB of RAM
- Light user workload

Using the Hitachi Data Systems cell design, this solution can be scaled out to support thousands of desktops.

# Solution Components

These are the major components used in this solution.

## Hitachi Compute Blade 500

Hitachi Compute Blade 500 combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

## Hitachi Unified Storage

Hitachi Unified Storage is a midrange storage platform for all data. It helps businesses meet their service level agreements for availability, performance, and data protection.

The performance provided by Hitachi Unified Storage is reliable, scalable, and available for block and file data. Unified Storage is simple to manage, optimized for critical business applications, and efficient.

Using Unified Storage requires a smaller capital investment. Deploy this storage, which grows to meet expanding requirements and service level agreements, for critical business applications. Simplify your operations with integrated set-up and management for a quicker time to value.

Unified Storage enables extensive cost savings through file and block consolidation. Build a cloud infrastructure at your own pace to deliver your services.

Hitachi Unified Storage 150 provides reliable, flexible, scalable, and cost-effective modular storage. Its symmetric active-active controllers provide input-output load balancing that is integrated, automated, and hardware-based.

Both controllers in Unified Storage 150 dynamically and automatically assign the access paths from the controller to a logical unit (LU). All LUs are accessible, regardless of the physical port or the server that requests access.

## Hitachi Dynamic Provisioning

On Hitachi storage systems, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like 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. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a logical size you specify of up to 60 TB created against it without allocating any physical space initially.

Deploying 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. Dynamic provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

When used with Hitachi Unified Storage, Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 1 GB chunks, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

## Citrix XenDesktop 5.5

Citrix XenDesktop is a desktop virtualization solution that transforms Microsoft Windows desktops and applications into an on-demand service available to any user, anywhere, on any device. With XenDesktop, you can securely deliver individual Windows, web and SaaS applications, or full virtual desktops, to PCs, Macs, tablets, smart phones, laptops, and thin clients — all with a high-definition user experience.

The approach used in this solution simplifies the design and planning of the most commonly used form of virtual desktops, Citrix Provisioning Services.

## Citrix Provisioning Services 5.6.2

Use Citrix Provisioning Services to stream a single desktop image to create multiple virtual desktops on one or more servers in your data center. Provisioning Services greatly reduces the amount of storage required when compared to other methods of creating virtual desktops.

## VMware vSphere 5.0

VMware vSphere 5 is a virtualization platform that provides a data center infrastructure. It features vSphere Distributed Resource Scheduler (DRS), high availability, and fault tolerance.

VMware vSphere 5 has the following components:

- **ESXi 5.0** — this is a hypervisor that loads directly on a physical server. It partitions one physical machine into many virtual machines that share hardware resources.
- **vCenter Server** — this allows management of the vSphere environment through a single user interface. With vCenter, there are features available such as vMotion, Storage vMotion, Storage Distributed Resource Scheduler, High Availability, and Fault Tolerance.

## Brocade Switches

Brocade and Hitachi Data Systems partner to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

This solution uses the following Brocade products:

- Brocade 5460 8 Gb SAN switch for Hitachi Compute Blade 500
- Brocade 6510
- Brocade VDX 6720
- Brocade VDX 6746

# Solution Design

This solution uses a cell architecture. A cell provides the type of a specific type of component necessary to build a solution. The cell architecture defines the compute and storage resources needed to support a defined workload. The cell types in this solution are the following:

- **Infrastructure Cells**
  - **Compute resources** — Foundation for compute components
  - **Storage resources** — Foundation for storage components
- **Application Cells**
  - **Hitachi Unified Compute Platform Select management** — Required for resources for hosting management services for VMware vSphere, Citrix XenDesktop, or other vendor management software if no management environment currently exists or if specific application management services need isolation
  - **Citrix Provisioning Services** — Resources for hosting the Citrix Provisioning Services environments
  - **Citrix XenDesktop** — Resources for hosting Citrix XenDesktop virtual desktops
- **Expansion Cell**
  - **Compute Resources** — Resources for scaling out an application cell for two Citrix XenDesktop application cells

The starting design of this Citrix XenDesktop solution supports 500 light workload users. When fully scaled out with the maximum number of cells, the solution supports 3,500 light workload user desktops.

Figure 2 shows the cell architecture scale-out configuration that supports 3,500 desktops.

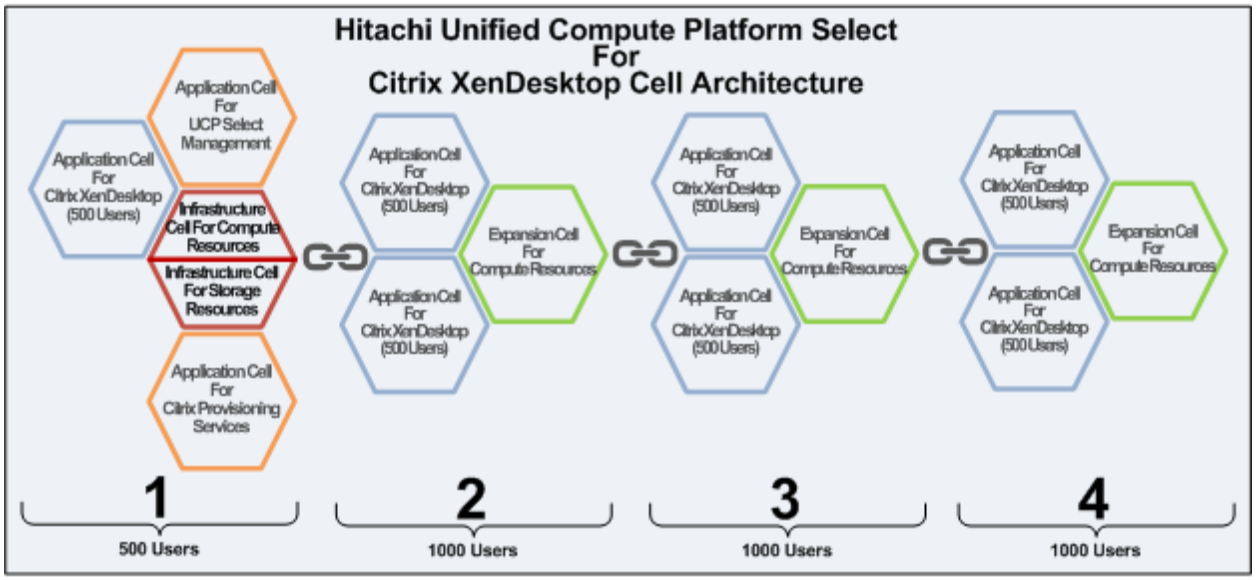


**Figure 2**

## Infrastructure Cell for Compute Resources

The infrastructure cell for compute resources (Figure 3 on page 13) provides the foundation for the compute components needed to start building a scalable Citrix XenDesktop solution. It consists of these components:

- 42U rack enclosure
- Hitachi Compute Blade 500 Chassis
  - 6 cooling fan modules
  - 4 power supply modules
  - Server blades housed in this chassis are for other cells
  - 2 Brocade VDX 6746 DCB switch modules
  - 2 Brocade 5460 Fibre Channel switch modules
    - 6 ports
    - 8 Gb/sec
- 2 Brocade 6720-60 Ethernet switches
  - 60 ports
  - 10 Gb/sec
- Brocade 6510-48 Fibre Channel switches
  - 48 ports
  - 8 Gb/sec

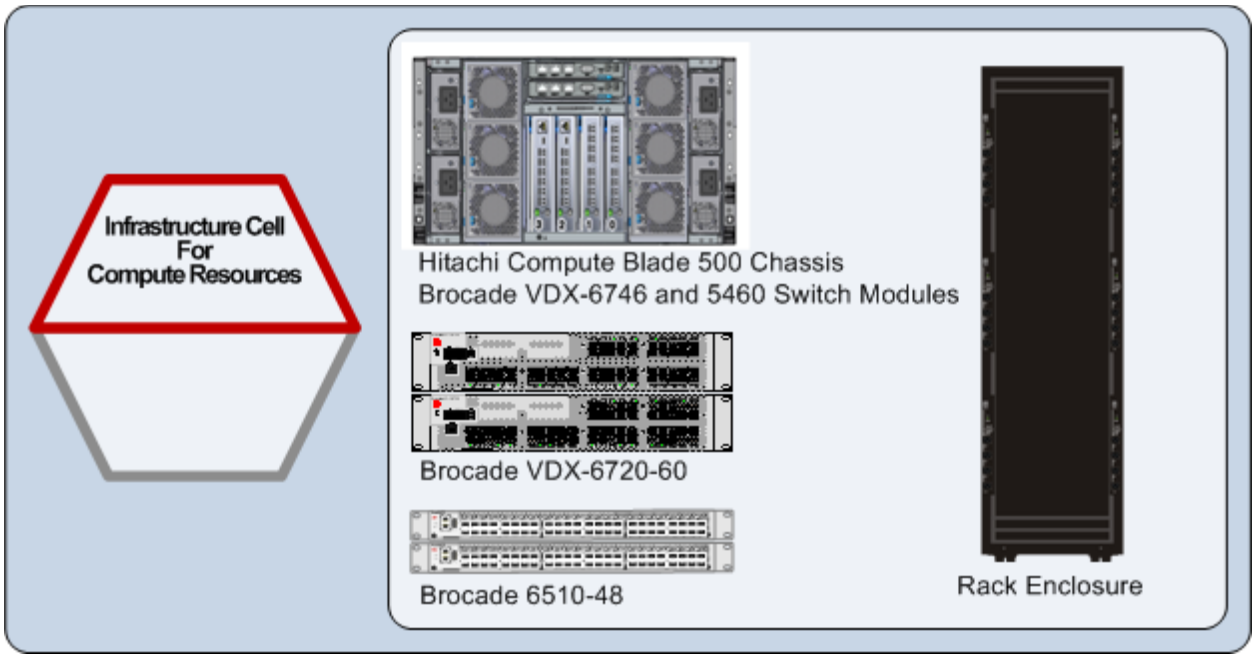


**Figure 3**

Use the infrastructure cell for compute resources for the following cells:

- Infrastructure cell for storage resources
- Application cell for Hitachi Unified Compute Platform Select management
- Application cell for Citrix Provisioning Services
- Application cell for Citrix XenDesktop
- Expansion cell for compute resources

The infrastructure cell for compute resources and the infrastructure cell for storage resources are the core infrastructure cells required to build a scalable Citrix XenDesktop solution. The Hitachi Compute Blade 500 chassis in the infrastructure cell for compute resources supports the two configurations in Table 3.

**Table 3. Infrastructure Cell for Compute Resources Configurations for the Hitachi Compute Blade 500 Chassis**

| *Configuration Options* | *Server Blades 0-1* | *Server Blades 2-3* | *Server Blades 4-5* | *Server Blades 6-7* |
|---|---|---|---|---|
| XenDesktop infrastructure part of the Unified Compute Platform solution | Application Cell for Hitachi Unified Compute Platform Select management | Application Cell for Citrix Provisioning Services | Application Cell for Citrix XenDesktop | |
| XenDesktop infrastructure outside of the Unified Compute Platform solution | Application Cell for Citrix XenDesktop | | Application Cell for Citrix XenDesktop | |

The infrastructure cell for compute resources and the infrastructure cell for storage resources support up to three expansion cells for Hitachi Compute Blade 500 before requiring new infrastructure cells. Every infrastructure cell for compute resources requires one infrastructure cell for storage resources.

## Chassis Components

The Hitachi Compute Blade 500 chassis is equipped with redundant management modules to provide high availability access to manage and monitor the chassis, switch modules, and server blades.

The chassis contains redundant switch modules for high availability and maximum throughput.

Hot swappable power and fan modules allow for non-disruptive maintenance.

## Network Infrastructure

The network design used in this solution provides ample bandwidth and redundancy for the following:

- One fully populated infrastructure cell for compute resources
- One infrastructure cell for storage resources
- Up to three expansion cells for Hitachi Compute Blade 500

Figure 4 illustrates the physical network configuration of the infrastructure cell for compute.

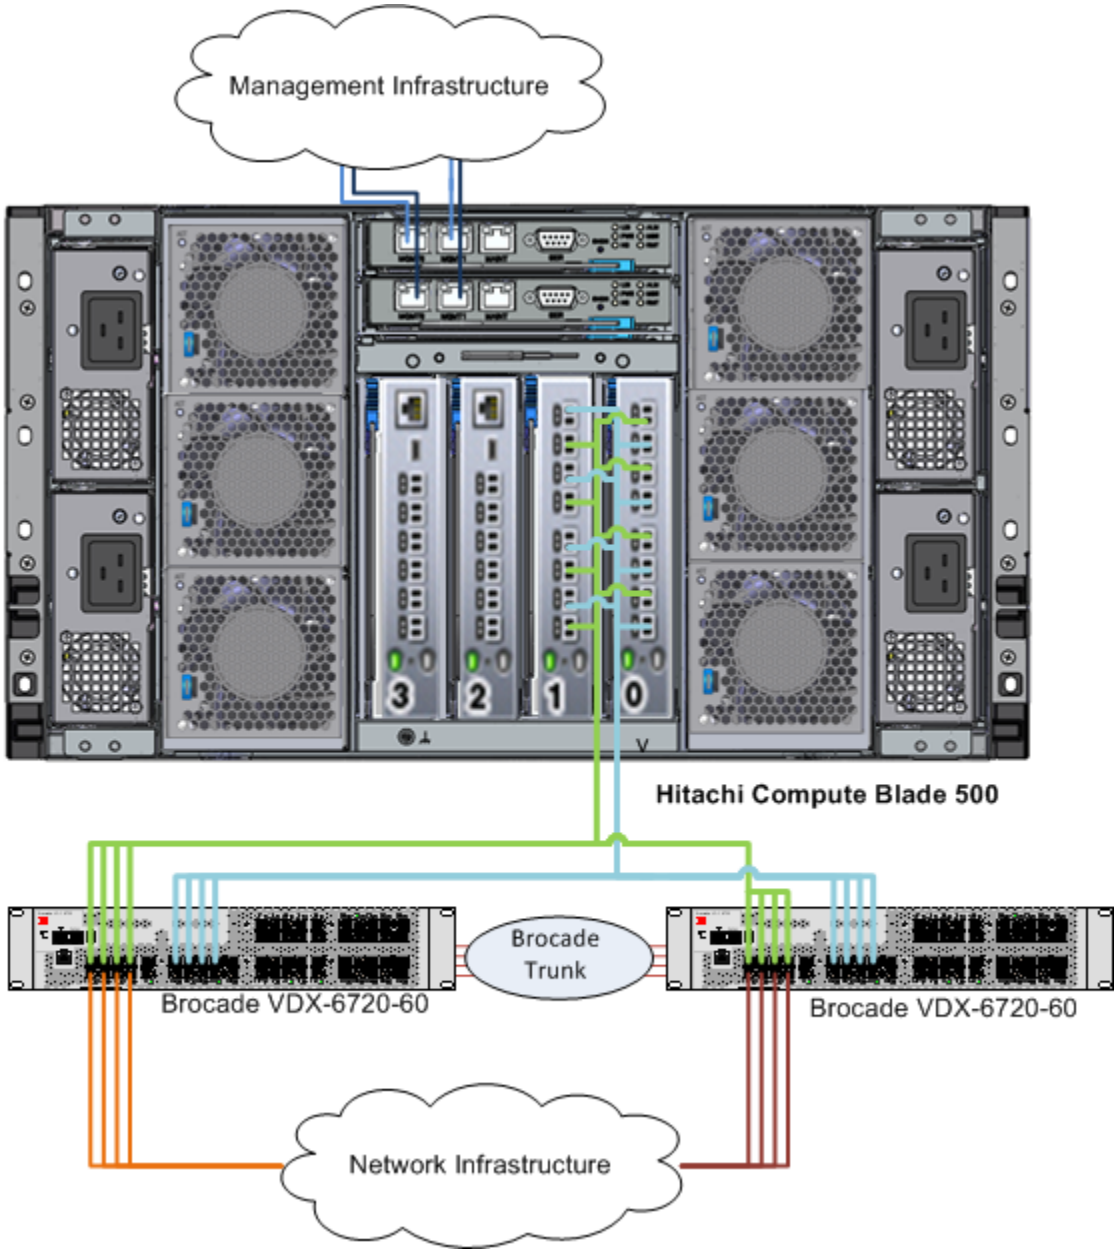


**Figure 4**

The network design allows for the utilization of advanced features inherent in the Brocade VDX switch family, such as VCS Fabric Technology. This provides:

- Non-stop networking
- Simplified, automated networks

## SAN Infrastructure

The Hitachi Unified Storage 150 controller used for this solution has 16 ports for connections to the Brocade 6510 enterprise fabric switches.

For this solution, zone the infrastructure cell for compute resources to four ports, with two ports per processor. When adding expansion cells for Hitachi Compute Blade 500 to the solution, zone four new open ports on the controller to each new cell.

Dedicating four ports on Hitachi Unified Storage 150 to each Hitachi Compute Blade 500 chassis ensures bandwidth between each chassis and Hitachi Unified Storage 150.

Figure 5 on page 17 illustrates the physical SAN configuration of the infrastructure cell for compute.

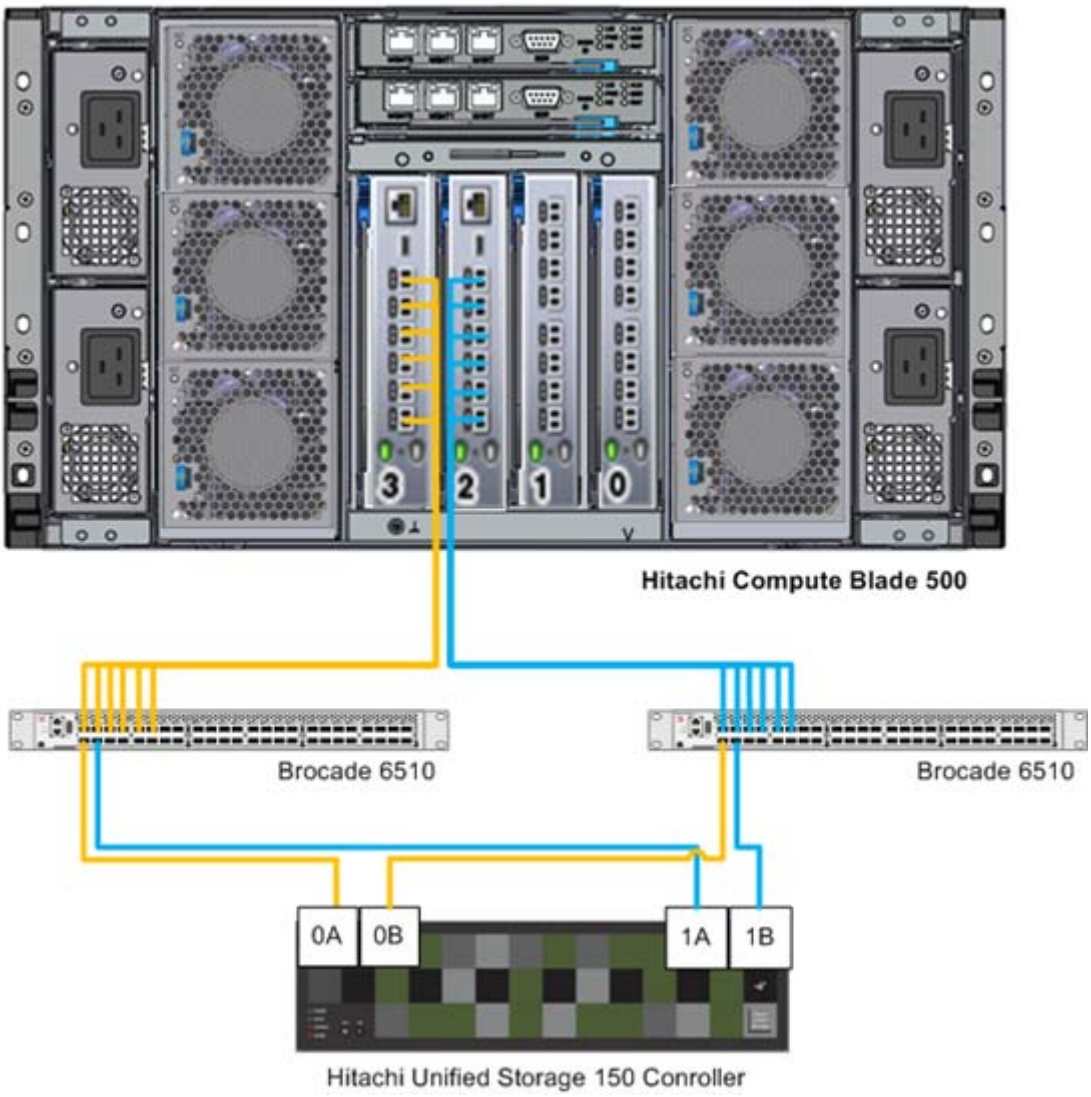
3
2
1
0
Hitachi Compute Blade 500
Brocade 6510
Brocade 6510
0A
0B
1A
1B
Hitachi Unified Storage 150 Conroller

**Figure 5**

## Infrastructure Cell for Storage Resources

The infrastructure cell for storage resources (Figure 6) contains all of the base storage hardware required to start building a Citrix XenDesktop solution. It consists of these components:

- 42U rack enclosure
- 1 Hitachi Unified Storage 150
  - 2 controller modules
  - 2 Fibre Channel modules
    - 16 × 8 Gb/sec Fibre Channel ports
  - 1 SFF disk expansion tray
    - Disk drives housed in this tray are for other cells
  - 32 GB cache
- Hitachi Unified Storage controller

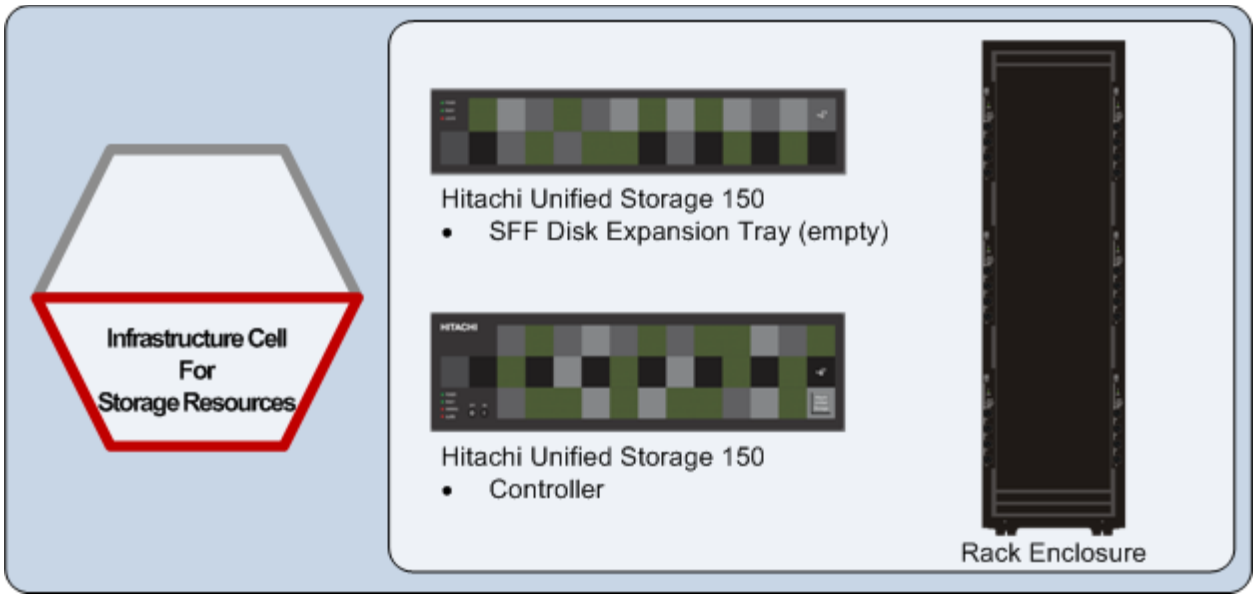


**Figure 6**

Use an infrastructure cell for storage resources in conjunction with the following cells:

- Infrastructure cell for compute resources
- Application cell for Hitachi Unified Compute Platform Select management
- Application cell for Citrix Provisioning Services
- Application cell for Citrix XenDesktop

The infrastructure cell for storage provides the storage infrastructure for all the cells in the solution. After fully populating an infrastructure cell for storage resources, add infrastructure cells for storage resources to scale out the solution.

The Hitachi Unified Storage 150 controller in each infrastructure cell for storage resources can support up to eight application cells for Citrix XenDesktop.

The Hitachi Unified Storage 1500 SFF disk expansion tray in the infrastructure cell for storage resources can house the following:

- An application cell for Hitachi Unified Compute Platform Select management
- Application cells for Citrix Provisioning Services
- Hot spare drives

## Application Cell for Hitachi Unified Compute Platform Select Management

The application cell for Hitachi Unified Compute Platform Select management (Figure 7) contains the compute and storage components for hosting the VMware vSphere and Citrix XenDesktop infrastructure services. It consists of these components:

- 2 × 520HB1 server blades

  Each server blade has the following:

  - 2 × 6-core Intel Xeon E5-2640 processors, 2.5 GHz
  - 128 GB RAM
  - 1 Emulex 2-port 10 GbE on-board CNA card
  - 1 Emulex 2-port 8 Gb Fibre Channel mezzanine card
- 8 × 600 GB 10k RPM SAS drives
  - RAID-6 (6D+2P) configuration
  - Installed in the disk tray of the infrastructure cell for storage resources
- 1 × 600 GB 10k RPM SAS drive
  - Hot spare drive
  - Installed in the disk tray of the infrastructure cell for storage resources

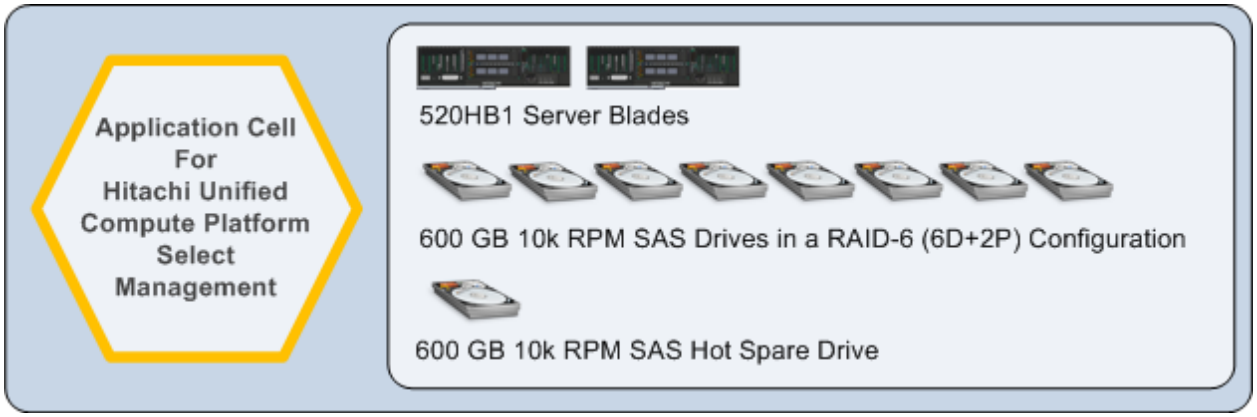


**Figure 7**

Use an application cell for Hitachi Unified Compute Platform Select management in conjunction with the following cells:

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources

Use an application cell for Hitachi Unified Compute Platform Select management if either of the following exists:

- A VMware vSphere environment does not exist already
- The Citrix XenDesktop infrastructure and management components need to be isolated

The application cell for Hitachi Unified Compute Platform Select management can be added to either of the following:

- The infrastructure cell for compute resources and infrastructure cell for storage resources
- A separate environment outside of the Citrix XenDesktop solution
  - If installed outside of the Citrix XenDesktop solution, follow the recommended virtual machine requirements outlined in this solution to size the compute hardware used to host the infrastructure and management virtual machines.

Table 4 lists the software components in the application cell for Hitachi Unified Compute Platform Select management.

**Table 4. Software Components for the Application Cell for Hitachi Unified Compute Platform Select Management**

| *Software* | *Version* |
|---|---|
| Citrix XenDesktop | 5 SP1 |
| VMware vCenter | 5 |
| VMware ESXi | 5.0.0 U1 Build 623860 |
| Microsoft Windows Server | 2008 R2 Enterprise |
| Microsoft SQL Server | 2008 R2 Standard |

## Compute Infrastructure

The application cell for Hitachi Unified Compute Platform Select management provides enough capacity to support an emergency high availability event if a single server blade fails. Use VMware High Availability and Distributed Resource Scheduler to configure a cluster dedicated to the application cell for Hitachi Unified Compute Platform Select management to ensure virtual machine failover in the event of a hardware failure.

The compute infrastructure of the application cell for Hitachi Unified Compute Platform Select management supports all associated requirements for the following:

- Citrix desktop delivery controller
- Microsoft Active Directory
- DHCP
- VMware vCenter

## Network Infrastructure

Each 520HB1 server blade is configured with a single onboard two-channel 10 GbE CNA card for network traffic. Split each CNA card into four logical NICs per channel, for a total of eight NICs per blade. This solution only uses three NICs per channel. Set bandwidth and NIC allocation as follows:

- **Channel 0 and 1 NIC 0**

  Virtual Machine Management Network

  - VMKernel management network vSwitch
  - 1 GbE per NIC, for a total of 2 GbE

- **Channel 0 and 1 NIC 1**

  vMotion network

  - VMKernel vMotion network vSwitch
  - 2 GbE per NIC, for a total of 4 GbE

- **Channel 0 and 1 NIC 2**

  Desktop Network

  - Virtual machine network vSwitch
  - 7 GbE per NIC, for a total of 14 GbE

Figure 8 illustrates the CNA and Fibre Channel to switch module mapping for the Hitachi Compute Blade 500.

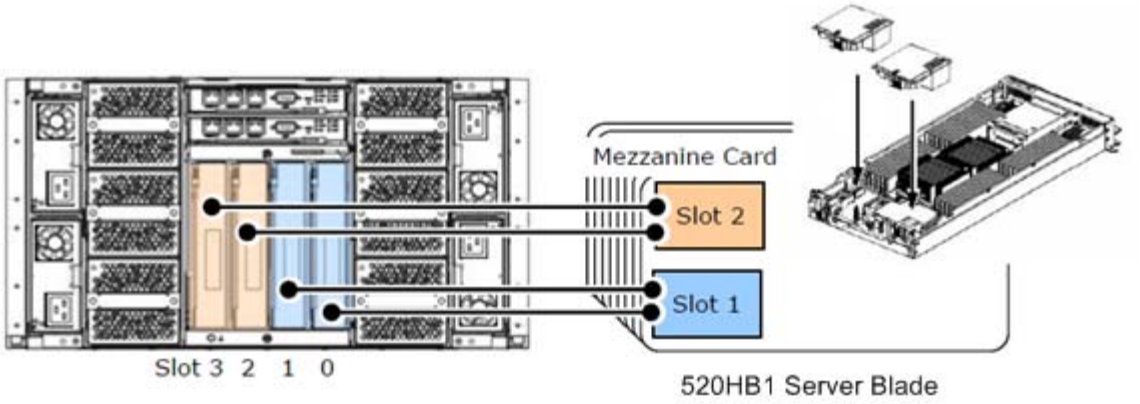


**Figure 8**

The following VLANs separate network traffic in the application cell for Hitachi Unified Compute Platform Select management:

- **Management-VLAN** — Chassis management connections and primary management of the ESXi hypervisors
- **vMotion-VLAN** — Configured for vMotion
- **Desktop-VLAN** — Configured for the desktop network

## Storage Infrastructure

The storage infrastructure of the application cell for Hitachi Unified Compute Platform Select management consists of nine 600 GB 10k RPM SAS drives housed in the disk expansion tray contained in the infrastructure cell for storage resources.

Using Hitachi Dynamic Provisioning, configure eight of the disks as follows:

- A single RAID-6 (6D+2P) dynamic provisioning pool dedicated to the following:
  - Citrix XenDesktop infrastructure
  - Management servers
- One disk used as a hot spare

Each deployed application cell for Hitachi Unified Compute Platform Select management uses a separate, dynamic provisioning pool configured with one RAID-6 (6D+2P) group dedicated to a Citrix XenDesktop infrastructure and management virtual machines.

Zone each 520HB1 server blade in the application cell for Hitachi Unified Compute Platform Select management to Hitachi Unified Storage 150 through the Brocade 5460 Fibre Chanel switch modules. Use a single initiator to multi target zoning for each HBA port on the 520HB1 server blades. Following best practice, configure the SAN environment in a dual fabric topology for redundancy and high availability. This results in four paths available to each ESXi host, providing the following:

- Resiliency to failure
- Redundant paths to the storage subsystem

Table 5 has the zoning configuration for the application cell for Unified Compute Platform Select management.

**Table 5. Application Cell for Hitachi Unified Compute Platform Select Management Zone Configuration**

| *Host* | *Host HBA Number* | *Fabric* | *Zone Name* | *Storage Port* |
|---|---|---|---|---|
| Blade 0-ESX 0 | HBA1_1 | Fabric 1 | esx0_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx0_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |
| Blade 1-ESX 1 | HBA1_1 | Fabric 1 | esx1_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx1_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |

## Recommended Practices for Scaling

These are guidelines and recommended practices to maintain adequate end user performance when scaling out the application cell for Hitachi Unified Compute Platform Select management:

- Always create a new, dedicated application cell for Unified Compute Platform Select management for every eight application cells for Citrix XenDesktop. (See “Application Cell for Citrix XenDesktop” on page 29.)
- Design for high availability compliance and rolling host maintenance if end users require uninterrupted service.
- Set the multipathing policy for each target to **round robin** in the ESXi configuration.

## Server Configuration Sizing Guidelines

It is critical to apply proper resource allocation for the management infrastructure virtual machines used for the Citrix XenDesktop environment. Even with proper sizing of end user resources, end user experience can suffer if the management infrastructure is resource starved or under sized.

If a management infrastructure already exists and the infrastructure cell for Hitachi Unified Compute Platform Select management is not necessary, use the virtual machine sizing recommendations in Table 6 to size the hosting hardware listed in “Application Cell for Hitachi Unified Compute Platform Select Management” on page 20.

Table 6 lists the virtual machine configurations for each component of the management infrastructure in this reference architecture.

**Table 6. Virtual Machine Sizing Recommendations**

| *Virtual Machine* | *Configuration* | *Count* |
|---|---|---|
| Microsoft Active Directory, DNS, and DHCP | ■ vCPU — 1<br>■ vRAM— 8GB<br>■ Operating system drive — 20 GB | 1 |
| VMware vCenter | ■ vCPU — 2<br>■ vRAM — 8GB<br>■ Operating system drive — 20 GB | 1 |
| Microsoft SQL Server 2008 database for the following:<br>■ VMware vCenter<br>■ Citrix XenDesktop<br>■ Citrix Provisioning Services | ■ vCPU — 2<br>■ vRAM — 8GB<br>■ Operating system drive — 20 GB<br>■ Data drive — 20 GB | 1 |
| Citrix XenDesktop Delivery Controllers | ■ vCPU — 4<br>■ vRAM — 8GB<br>■ Operating system drive — 20GB | 2 |
| Citrix XenDesktop Web Interface Servers | ■ vCPU — 2<br>■ vRAM — 4GB<br>■ Operating System Drive — 20GB | 2 |

## Application Cell for Citrix Provisioning Services

The application cell for Citrix Provisioning Services (Figure 9) contains all compute and storage components necessary to host the following:

- Virtual machines for Citrix Provisioning Services supporting the Citrix XenDesktop environment
- Desktop master images

It consists of these components:

- 2 × 520HB1 server blades

  Each server blade has the following:

  - 2 × 6-core Intel Xeon E5-2640 processors, 2.5 GHz
  - 128 GB RAM
  - 1 Emulex 2-port 10 GbE on-board CNA card
  - 1 Emulex 2-port 8 Gb Fibre Channel mezzanine card
- 8 × 600 GB 10k RPM SAS drives
  - RAID-6 (6D+2P) configuration
  - Installed in the disk tray of the infrastructure cell for storage resources


**Figure 9**

Use the application cell for Citrix Provisioning Services in conjunction with the following cells:

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources

Table 7 has the software components in the application cell for Citrix Provisioning Services for this architecture.

**Table 7. Application Cell for Citrix Provisioning Services Software Components**

| *Software* | *Version* |
|---|---|
| Citrix Provisioning Services | 5.6 SP2 |
| VMware ESXi | 5.0.0 U1 Build 623860 |
| Microsoft Windows Server | 2008 R2 Enterprise |

### Compute Infrastructure

Based on the Citrix recommendation of 500 desktop connections per provisioning server, the compute components in the cell are sized to host nine provisioning servers. This is to support either of the following:

- Eight application cells for Citrix XenDesktop (“Application Cell for Citrix XenDesktop” on page 29)
- 4000 desktops in an N+1 configuration

The application cell for Citrix Provisioning Service provides enough capacity to support an emergency high availability event if a single server blade fails. Use a VMware High Availability and Distributed Resource Scheduler cluster dedicated to the application cell for Citrix Provisioning Services to ensure virtual machine failover in the event of a hardware failure.

### Network Infrastructure

The application cell for Citrix Provisioning Services uses the same networking configuration described in the application cell for Hitachi Unified Compute Platform Select management. See “Application Cell for Hitachi Unified Compute Platform Select Management” on page 20.

### Storage Infrastructure

The storage in the application cell for Citrix Provisioning Services provides 3.6 TB of space to host the virtual disk files for up to nine Citrix Provisioning Services servers and up to twenty 20 GB desktop images per Citrix Provisioning Services server.

Using Hitachi Dynamic Provisioning, configure the storage into a single RAID-6 (6D+2P) dynamic provisioning pool dedicated to Citrix Provisioning Services servers.

- Each application cell for Citrix Provisioning Services deployed uses a separate dynamic provisioning pool with one RAID-6 (6D+2P) group dedicated to Citrix Provisioning Services servers.
- The application cell for Citrix Provisioning Services uses the hot spare drive contained in the application cell for Hitachi Unified Compute Platform Select.

Zone each 520HB1 server blade in the application cell for Citrix Provisioning Services to Hitachi Unified Storage 150 through the Brocade 5460 Fiber Chanel switch modules. Use single initiator to multi-target zoning for each port on the server blades. Following best practice, configure the SAN environment in a dual fabric topology for redundancy and high availability. This results in four paths available to each ESXi host, providing the following:

- Resiliency to failure
- Redundant paths to the storage subsystem

Table 8 has the zoning configuration for the application cell for Citrix Provisioning Services.

**Table 8. Application cell for Citrix Provisioning Service Zone Configuration**

| *Host* | *Host HBA Number* | *Fibre* | *Zone Name* | *Storage Port* |
|---|---|---|---|---|
| Blade2-ESX 2 | HBA1_1 | Fabric 1 | Esx2_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | Esx2_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |
| Blade3-ESX 3 | HBA1_1 | Fabric 1 | Esx3_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | Esx3_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |

## Server Configuration Sizing Guidelines

The application cell for Citrix Provisioning Services supports up to both of the following:

- Eight application cells for Citrix XenDesktop
- Seven additional Citrix Provisioning Services virtual machines

Add a new Citrix Provisioning Services virtual machine to the environment for either of the following:

- Each application cell for Citrix XenDesktop added
- Every 500 desktops

Use the virtual machine sizing recommendations in this reference architecture guide when adding Citrix Provisioning Services virtual machines to the environment.

Table 9 has the recommended virtual machine configuration for a Citrix Provisioning Services server supporting 500 desktops.

**Table 9. Virtual Machine Sizing Recommendations for 500 Desktops**

| *Virtual Machine* | *Configuration* |
|---|---|
| Citrix Provisioning Services | ■ **vCPU** — 4<br>■ **vRAM** — 8GB<br>■ **Operating system drive** — 20 GB |

## Application Cell for Citrix XenDesktop

The application cell for Citrix XenDesktop (Figure 10 on page 30) contains all compute and storage components necessary to support up to 500 streamed desktops for light workload users. It consists of these components:

- 4 × 520HB1 server blades

  Each server blade has the following:

  - 2 × 6-core Intel Xeon E5-2640 processors, 2.5 GHz
  - 144 GB RAM
  - 1 Emulex 2-port 10 GbE onboard CNA card
  - 1 Emulex 2-port 8 Gb Fibre Channel mezzanine card
- 24 × 600 GB 10k RPM SAS drives
  - RAID-10 (2D+2P) configuration
  - Installed in the disk tray of the infrastructure cell for storage resources
- 1 × 600 GB 10k RPM SAS drive
  - Hot spare drive
  - Installed in the disk tray of the infrastructure cell for storage resources
- 1 SFF disk expansion tray
  - Added to the infrastructure cell for storage resources

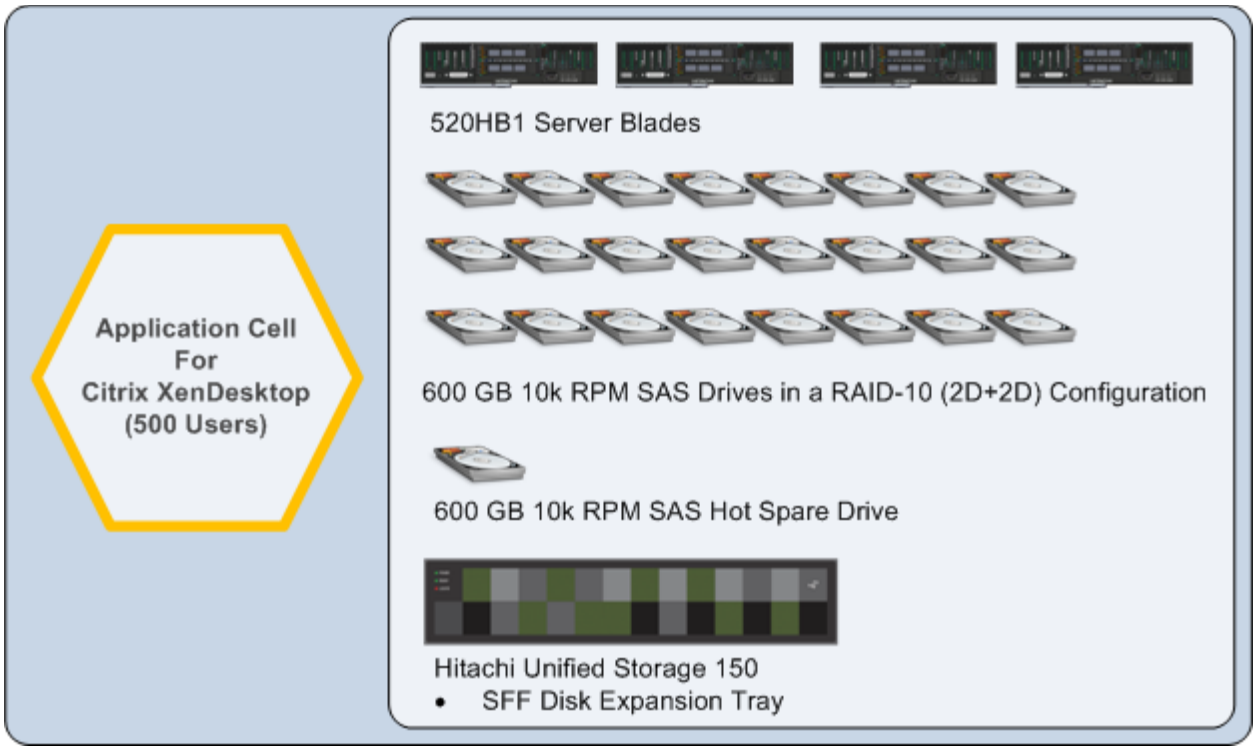


**Figure 10**

Use an application cell for Citrix XenDesktop in conjunction with the following cells:

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources
- Expansion cell for compute resources (used for scale-out)

To start building a scalable Citrix XenDesktop environment:

- Add the compute components of the application cell for Citrix XenDesktop to the Hitachi Compute Blade 500 chassis included with the infrastructure cell for compute resources
- Add the storage components to the infrastructure cell for storage resources

Each application cell for Citrix XenDesktop supports up to 500 light workload users.

To scale out the solution, add one of the following to your environment:

- Additional application cells for Citrix XenDesktop to the infrastructure cell for compute resources
- An expansion cell for compute resources

Up to eight application cells for Citrix XenDesktop can be supported by a single infrastructure cell for compute resources and infrastructure cell for storage resources before you require new infrastructure cells.

### Compute Infrastructure

The application cell for Citrix XenDesktop supports a maximum density of 500 light user desktops per cell. In a maximum density configuration, a cell cannot support the failover of desktops in the case of a server blade failure.

To provide high availability, reduce the total number of desktops in units of 125 users. For every 125 users you reduce from the total, the environment can sustain one additional server blade failure.

For high availability, create a dedicated VMware High Availability and Distributed Resource Scheduler cluster. Then, place the hosts from each application cell for Citrix XenDesktop into the cluster. Based on VMware maximums each cluster can support up to eight application cells for Citrix XenDesktop (32 hosts).

Table 10 has the configuration of the Microsoft Windows 7 virtual machine for the streamed desktops.

**Table 10. Microsoft Windows 7 Virtual Machine Configuration**

| *Virtual Machine Configuration Option* | *Value* |
|---|---|
| Operating system | Microsoft Windows 7, 64-bit |
| vCPU | 1 |
| Memory | 2 GB |
| Operating system disk size | 20 GB |
| Write cache disk size | 5 GB |

Table 11 lists the software components in the application cell for Citrix XenDesktop.

**Table 11. Application Cell for Citrix XenDesktop Software Components**

| *Software* | *Version* |
|---|---|
| Citrix Virtual Desktop Agent | 5.5 |
| Citrix Provisioning Services Agent | 5.6.2 |
| VMware Tools | 8.6.0.6261 |
| VMware ESXi | 5.0.0 U1 Build 623860 |
| Microsoft Windows 7 | 64-bit Enterprise |

### Network Infrastructure

The application cell for Citrix XenDesktop uses the same networking configuration described in the application cell for Hitachi Unified Compute Platform Select management. See “Application Cell for Hitachi Unified Compute Platform Select Management” on page 20.

## Storage Infrastructure

The storage infrastructure of the application cell for Citrix XenDesktop consists of twenty-four 600 GB 10k RPM SAS drives. Following best practice, the SAN environment was configured in a dual fabric topology for redundancy and high availability. These drives are in two dynamic provisioning pools, each consisting of three RAID-10 (2D+2D) groups created with Hitachi Dynamic Provisioning. Each pool is dedicated for the following:

- Streamed desktops virtual machine disk files
- Write cache

This supports sustained light workloads and situations such as boot storms and log on storms.

The 520HB1 server blades used in this reference architecture use dual-port 8 Gb/sec Fibre Channel mezzanine cards. There are redundant connections to the Brocade 5460 fabric switch modules.

The environment uses single initiator to multi target zoning for each port on the 520HB1 server blades. This results in four paths available to each ESXi host, providing the following:

- Resiliency to failure
- Redundant paths to the storage subsystem

Set the multipathing policy for each target to **round robin** in the ESXi configuration. Following best practice, configure the SAN environment in a dual fabric topology for redundancy and high availability. This results in optimal load distribution during an all paths available situation.

Table 12 on page 33 shows the zone configuration for the application cell for Citrix XenDesktop.

**Table 12. Application cell for Citrix XenDesktop Zone Configuration**

| *Host and Cell* | *Host HBA Number* | *Fabric* | *Zone Name* | *Storage Port** |
|---|---|---|---|---|
| Blade4-ESX 4<br>▪ Application cell for Citrix XenDesktop | HBA1_1 | Fabric 1 | esx4_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx4_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |
| Blade5-ESX 5<br>▪ Application cell for Citrix XenDesktop | HBA1_1 | Fabric 1 | esx5_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx5_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |
| Blade6-ESX 6<br>▪ Application cell for Citrix XenDesktop | HBA1_1 | Fabric 1 | esx6_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx6_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |
| Blade7-ESX 7<br>▪ Application cell for Citrix XenDesktop | HBA1_1 | Fabric 1 | esx7_hba1_1_hus_0A_1A | 0A |
| | | | | 1A |
| | HBA1_2 | Fabric 2 | esx7_hba1_2_hus_0B_1B | 0B |
| | | | | 1B |

*The storage target ports for each cell are dependent on the Hitachi Compute Blade 500 chassis in which each is hosted. See “Expansion Cell for Compute Resources” on page 34 for details.

## Recommended Practices for Scaling

Follow this recommended practice to maintain adequate end user performance when scaling out using application cells for Citrix XenDesktop:

- Use dedicated dynamic provisioning pools for each defined user workload in the environment.

## Expansion Cell for Compute Resources

Use the expansion cell for compute resources (Figure 11) to scale out the Citrix XenDesktop solution beyond the first infrastructure cell for compute resources. It consists of these components:

- Housed in the 42U rack enclosure used by the infrastructure cell for compute resources
- Hitachi Compute Blade 500 Chassis
  - 6 cooling fan modules
  - 4 power supply modules
  - Holds server blades for application cells
  - 2 Brocade VDX 6746 DCB switch modules
  - 2 Brocade 5460 Fibre Channel switch modules
    - 6 ports
    - 8 Gb/sec

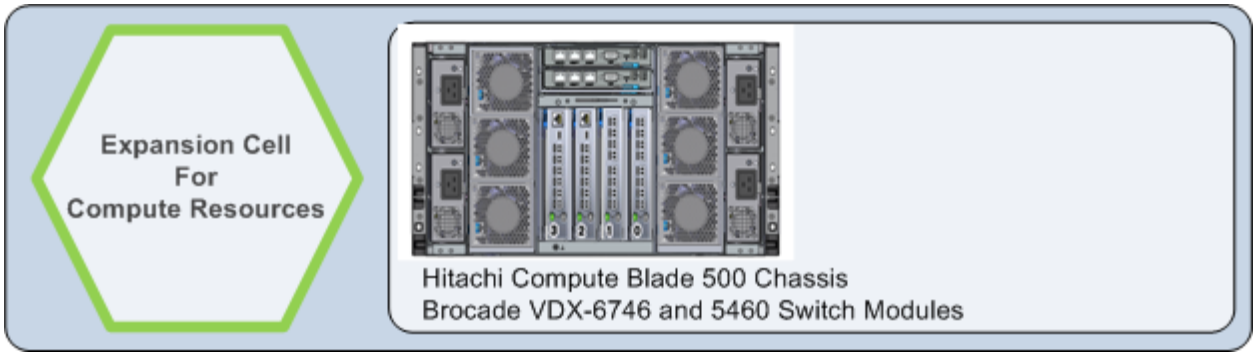


**Figure 11**

Use an expansion cell for compute resources in conjunction with the following cells:

- Infrastructure cell for compute resources
- Application cell for Citrix XenDesktop

Once the chassis in the infrastructure cell for compute resources is full, use an expansion cell for compute resources to provide additional desktop capacity.

The expansion cell for compute resources uses the storage and networking infrastructure provided in the infrastructure cells for compute resources and storage resources. House this cell in the infrastructure cell for compute resources rack enclosure.

You can add up to three expansion cells for Hitachi Compute Blade 500 to infrastructure cells for compute resources and storage resources before requiring new infrastructure cells.

### Chassis Components

The expansion cell for compute resources uses the same chassis components contained in the infrastructure cell for compute resources. See “Infrastructure Cell for Compute Resources” on page 12.

### Network Infrastructure

The expansion cell for compute resources uses the same networking configurations as the infrastructure cell for compute resources chassis. See “Infrastructure Cell for Compute Resources” on page 12.

### Storage Infrastructure

Utilizing four of the open storage target ports on Hitachi Unified Storage 150 in the infrastructure cell for storage resources, follow the same storage configuration described for the infrastructure cell for compute resources. Utilize the newly provisioned storage target ports in the zoning configuration.

Figure 12 shows the storage target ports of a fully scaled out solution.

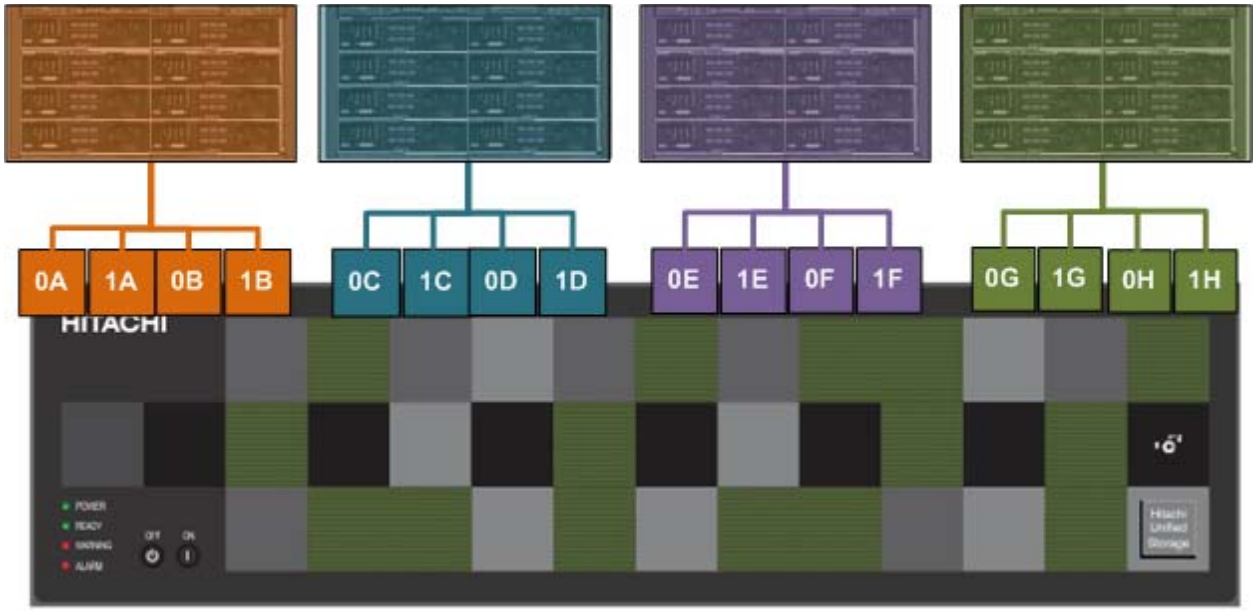


Hitachi Unified Storage 150 Controller

**Figure 12**

# Engineering Validation

This describes the test methodology used to validate the reference architecture with the results of the validation testing.

This reference architecture tested the core components of the Hitachi Unified Compute Platform Select for Citrix XenDesktop solution to determine maximum loads per application cell that the solution could support maintaining an acceptable end-user experience.

The tested components were validated to support streaming up to 500 desktops with a light user workload per application cell. The actual number of desktops in an environment will vary, depending on workload and high availability requirements.

One of the following of each cell was tested in this 500-user reference architecture (Figure 13 on page 37):

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources
- Application cell for Citrix VDI (500 desktops for light workload users)
- Application cell for Hitachi Unified Compute Platform Select
- Application cell for Citrix Provisioning Services
- Application cell for Citrix XenDesktop

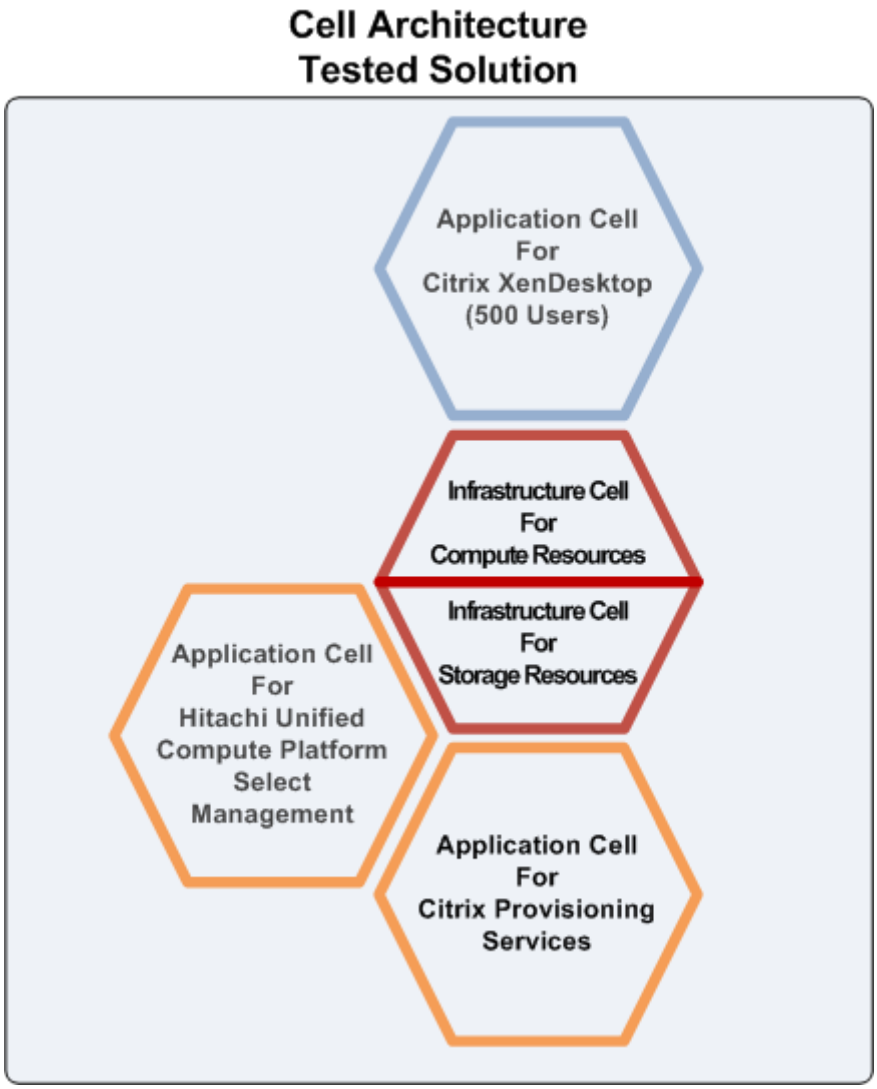


**Figure 13**

## Test Methodology

Testing used a single chassis to test a 500 seat streamed desktop configuration.

Login Consultants Login VSI generated the workload for the lab validation testing.

Testing was with a light user workload profile that generated an average of seven IOPS per desktop during steady state.

### Compute Infrastructure, Storage Infrastructure, and Application Experience

Table 13 defines the testing configuration used on the Login VSI controller to generate the workload.

**Table 13. Test Configuration**

| *Login VSI Configuration Option* | *Value* |
|---|---|
| Application launch delay | 10 seconds |
| Total number of virtual machines under test | 500 |
| Test mode | Light workload |
| Log off user(s) after test completion | Disabled |
| Workload | Microsoft Outlook, random Microsoft Word document, Microsoft Internet Explorer, Microsoft Excel |

### Logon and Steady State

Performance metrics were gathered from Hitachi Unified Storage 150 and ESXi hosts during logon and steady state tests to ensure that the Citrix XenDesktop environment could maintain performance under maximum loads.

### Boot Storm

Citrix XenDesktop Power Management controlled the desktop power state during the boot storm tests. The default settings were used:

- Maximum active actions — 40
- Max New actions per minute — 10
- Max power actions — 20%

Cache write pending and core utilization of the storage processor on Hitachi Unified Storage 150 was monitored throughout the entirety of the boot storm test.

## Test Results

These are the test results for the environment operating in a logon and steady state condition.

### Compute Infrastructure

During the logon and steady state test, the CPU and memory statistics were captured using esxtop to determine the maximum density that the compute hardware could support.

*Hypervisor CPU Performance*

Figure 14 illustrates the physical CPU metrics collected on the ESXi hypervisors while running the 500-user logon and steady-state workload.

- With the core utilization averaging at 96% during steady state, the server blades ran at maximum CPU utilization.
- The test determined maximum density per host.

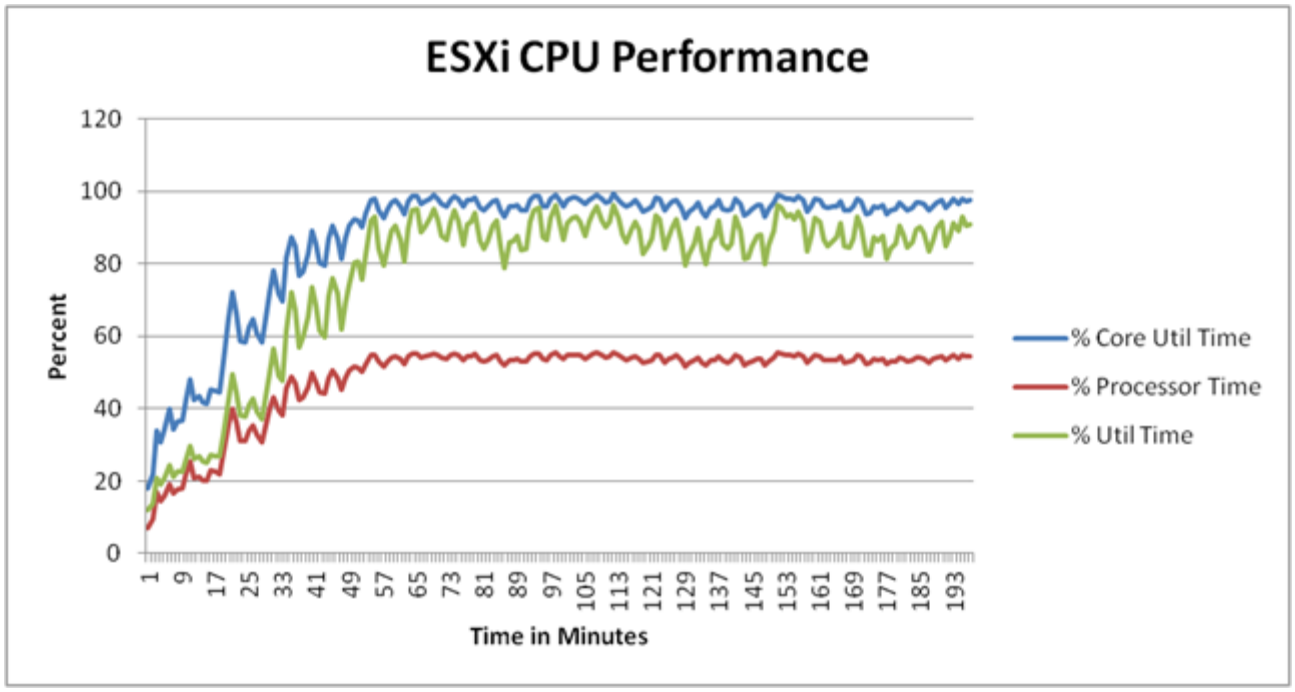


**Figure 14**

*Hypervisor Memory Performance*

Each 520HB1 server blade contained 144 GB of RAM. Splitting 500 desktops between four server blades, there was approximately 250 GB of vRAM configured on each server blade (125 virtual machines × 2 GB) under maximum density.

Figure 15 illustrates the benefits of transparent page sharing in VMware vSphere 5, with sharing about 45% of the granted 256 GB of guest RAM. This allows for over-allocation of RAM on each server blade for increased density of streamed desktops.

- Swap used on the hypervisor throughout the workload did not rise above 274 MB.
- The swap in and swap out rates were very low.
- The overall used memory throughout the 500-user test was close to the available memory. The 520HB1 server blade has adequate memory for light workloads.

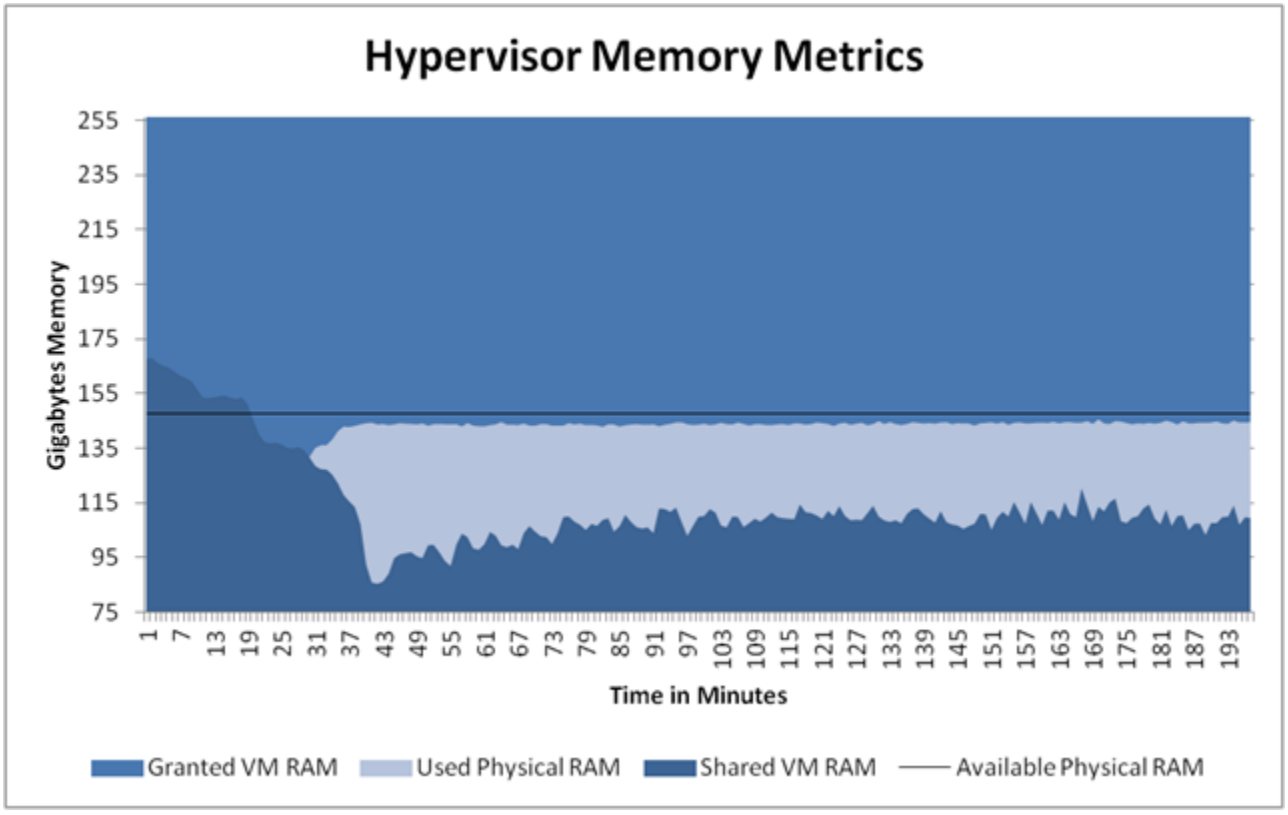


**Figure 15**

With the memory savings provided by the transparent page sharing, the 520HB1 server blades had adequate memory to support 500 desktops running a light workload.

## Storage Infrastructure

To understand the I/O profile of the typical light workload user, multiple performance metrics were collected from the Hitachi Unified Storage 150 controllers to test the following:

- Physical disks were saturated
- Cache was overwhelmed
- RAID groups performed well

Physical disk performance was acceptable during the 500-user test.

### *Physical Disk Performance*

Figure 16 on page 41 and Figure 17 on page 42 illustrate busy rate and tag count statistics. These are for a representative physical drive from a dynamic provisioning pool. The results for both metrics are acceptable, showing that there is room to support bursts of heavier workloads, if necessary.

For example, this physical drive shows the following:

- Peak at 72%
- Up to 30% headroom during steady state peaks

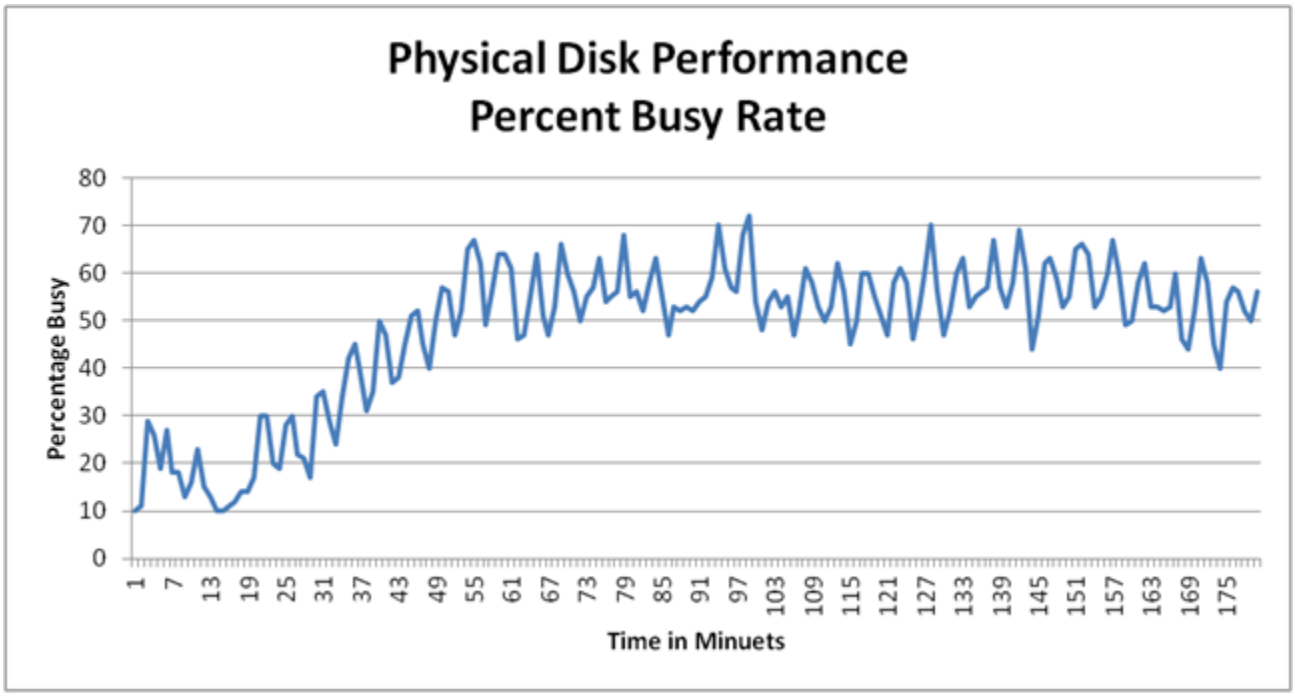


**Figure 16**

The low tag count statistics indicate the physical disks queue was not saturated.

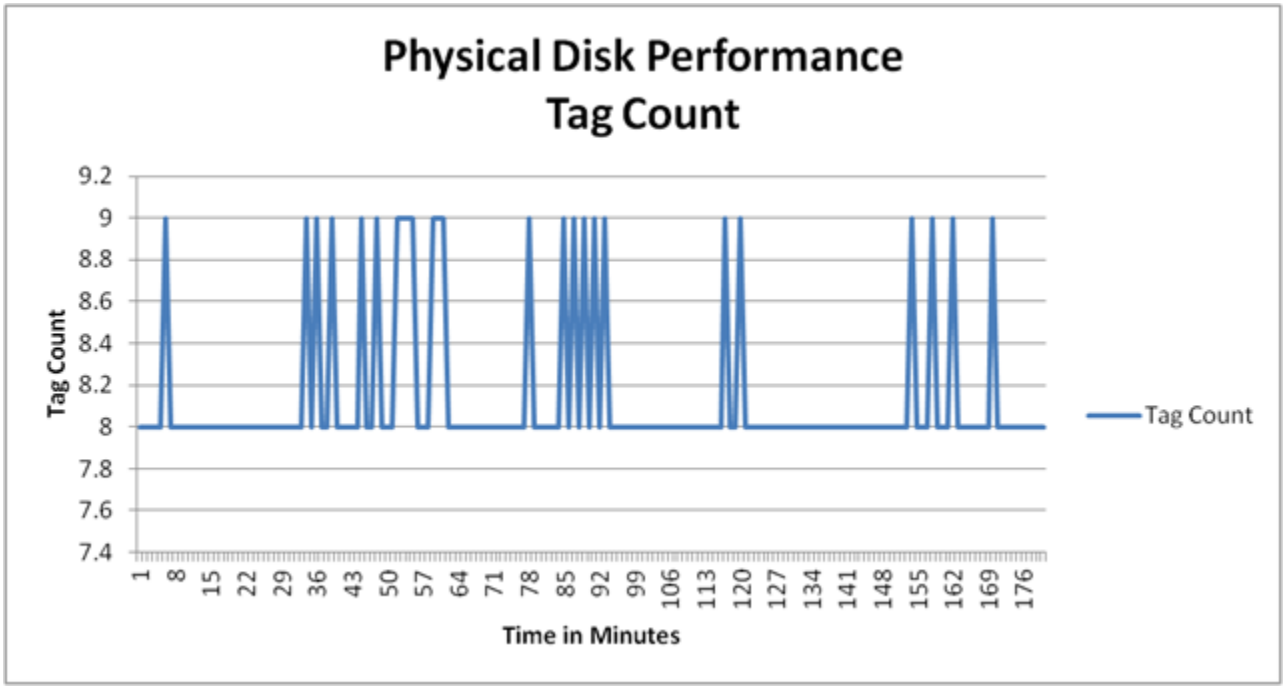


**Figure 17**

### *Storage Processor and Cache Performance*

Write cache saturation and storage processor core utilization on Hitachi Unified Storage 150 was monitored throughout the entirety of the test.

- Write pending percentage averaged below 4%.
- Core utilization peaked around 18%.

This data indicates the Hitachi Unified Storage 150 controllers are adequate for 500 light workload streamed desktops, with plenty of headroom for future growth.

Figure 18 on page 43 illustrates the processor and cache performance metrics.

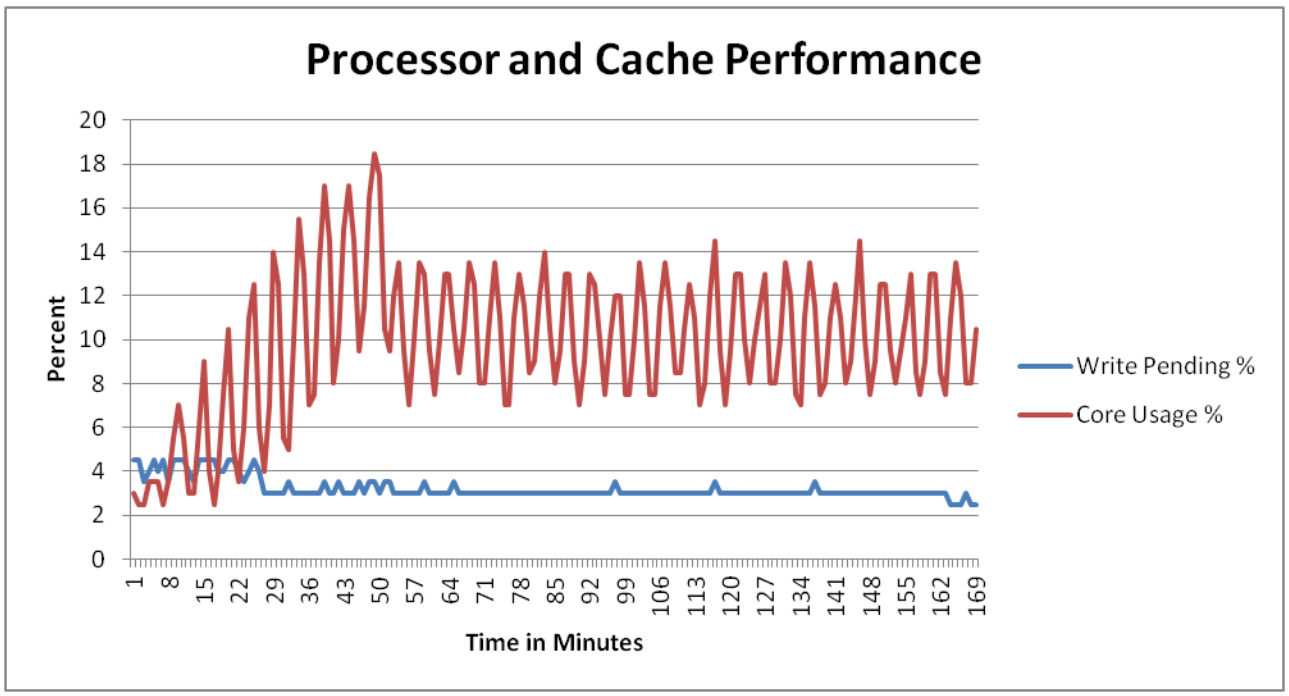


**Figure 18**

### *Datastore Performance*

Figure 19 and Figure 20 on page 44 illustrate the aggregated performance for the datastore dedicated to the XenDesktop virtual machine disk files and desktop write cache.

The following data points are well within acceptable ranges:

- The read IOPS peaked at approximately 440 during the logon storm.
- The write IOPS peaked at approximately 5217 during the logon storm.
- The read steady state IOPS peaked at approximately 61.
- The write steady state IOPS peaked at approximately 4579.
- Latency peaked at approximately 0.9 milliseconds.

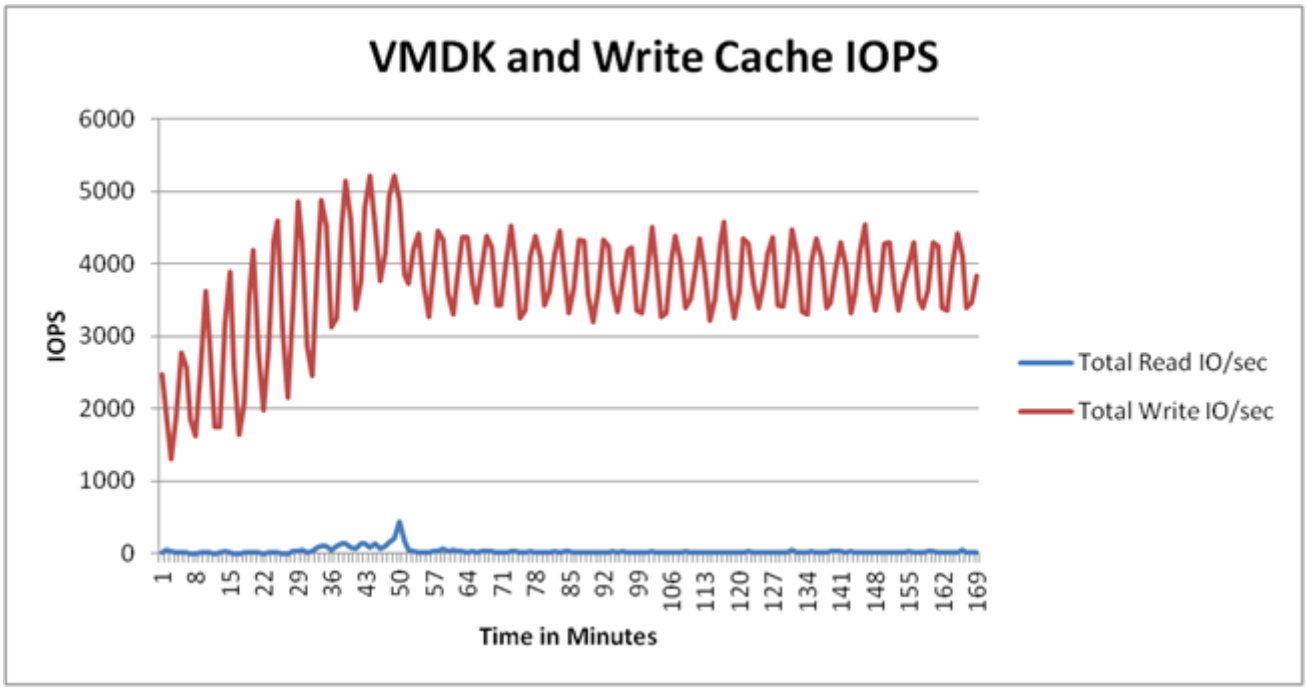


**Figure 19**

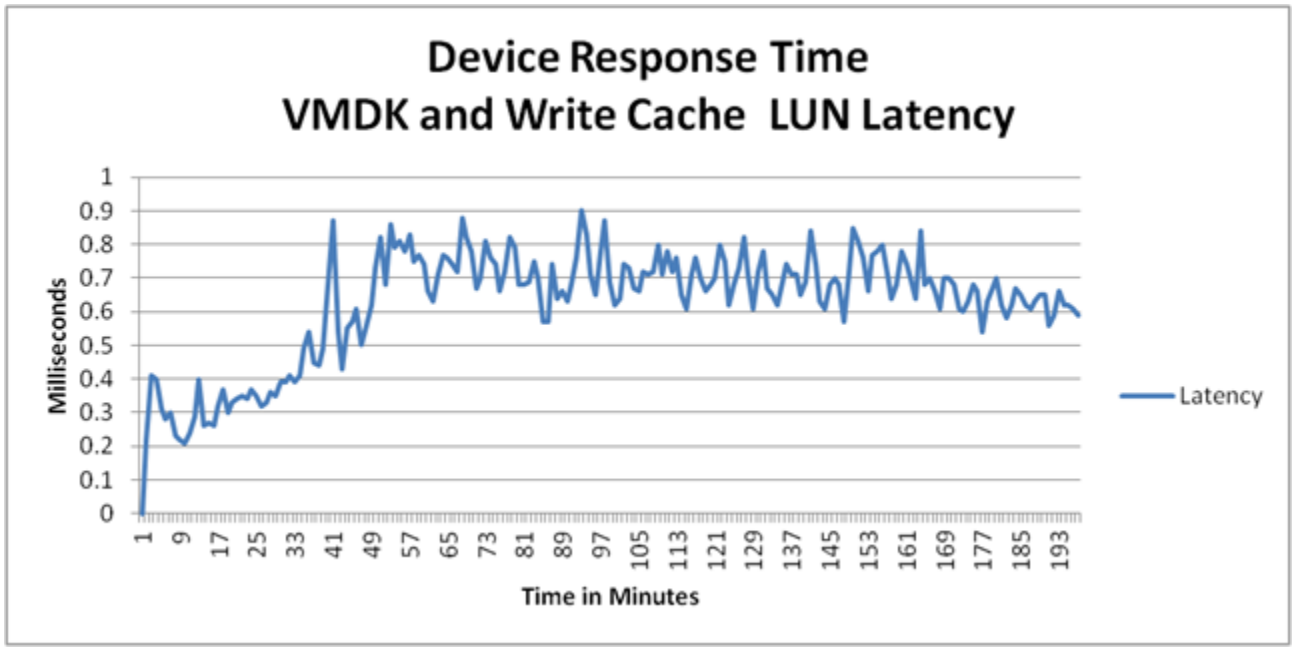


**Figure 20**

Metrics were collected for every virtual desktop during the 500-user run. Figure 21 illustrates a representative sample of five guests, showing the combined virtual machine disk files and write cache IOPS statistics for those specific guests.

- The IOPS peaked during logon storm for each desktop between one minute and sixty minutes (the first peak for each user).
- During the remainder of the steady state workload, the desktops averaged seven IOPS.

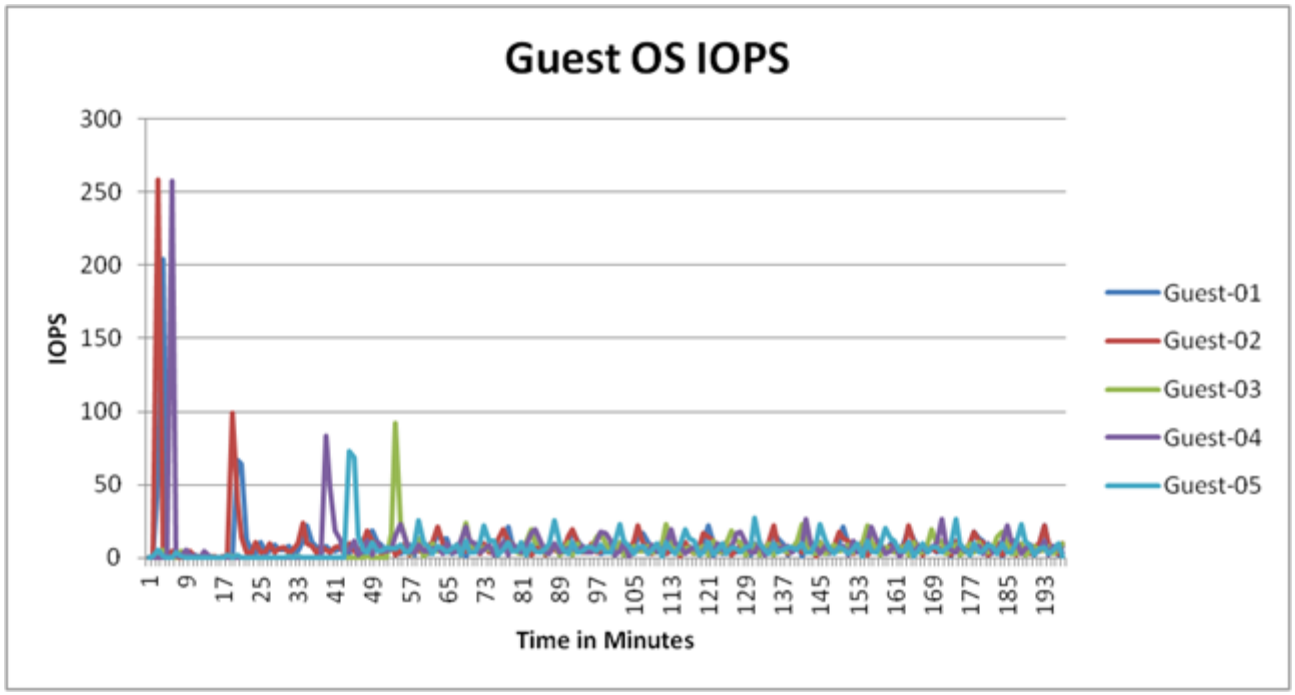


**Figure 21**

### *Desktop IO Profile*

Figure 22 illustrates the percentage of read and write I/O operations that were random versus sequential on the virtual machine disk files and write cache dedicated datastores. These measurements were taken at the back end storage level to show the I/O profile of the workload as seen by the Hitachi Unified Storage 150 controllers.

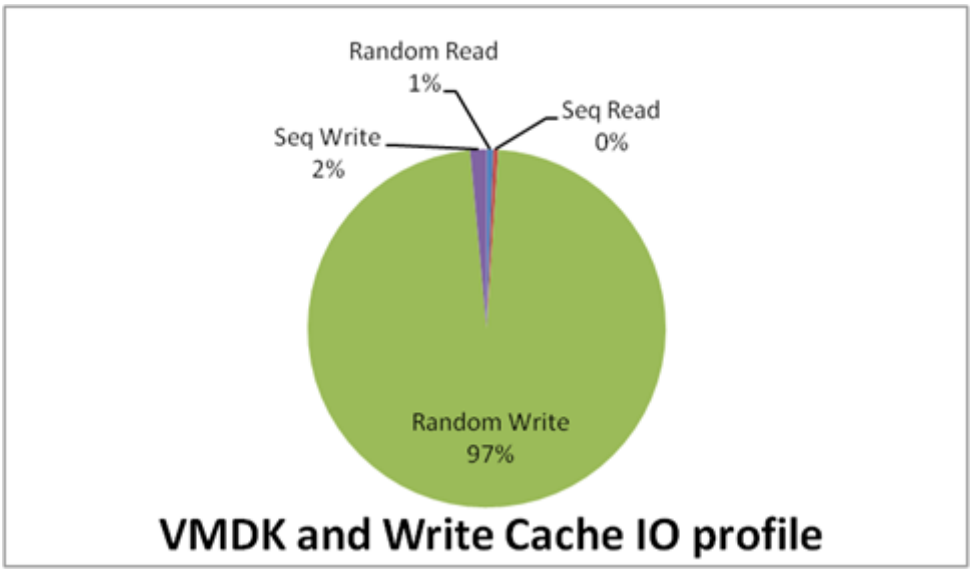


**Figure 22**

## Application Experience

The Login VSI testing tool reported the time required to open and close the simulated applications during the test.

- Load times for all tested applications were less than 1.5 seconds.
- File open times for all tested applications were less than 0.5 seconds.

These performance metrics are very close to physical desktop performance. This proves that this reference architecture provides adequate performance for 500 users with a light workload profile.

Figure 23 on page 47 illustrates the application experience metrics reported by Login VSI.

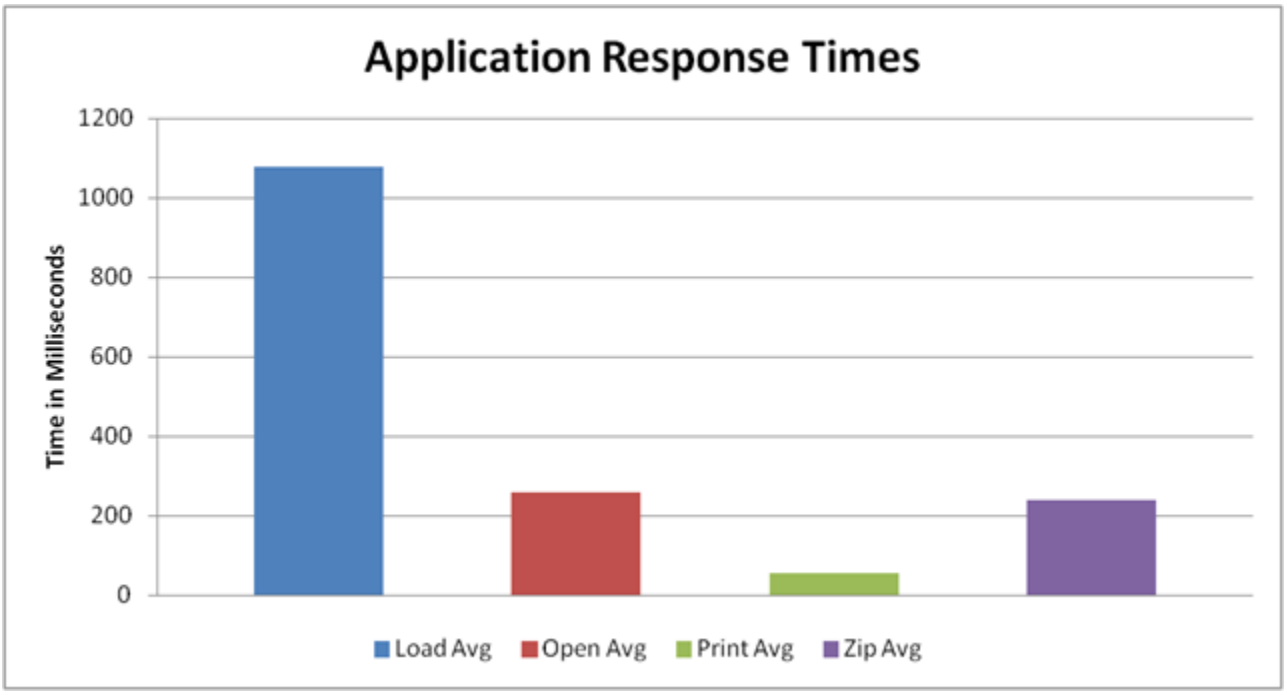


**Figure 23**

## Boot Storm

Using Citrix XenDesktop power management to power on the XenDesktop environment, it took approximately 60 minutes until all desktops were registered as ready in Desktop Studio.

Figure 24 on page 48 illustrates the aggregated performance of the RAID groups dedicated to the streamed desktops virtual machine disk files and write cache datastores.

- Read IOPS peaked at approximately 434.
- Write IOPs peaked at approximately 4016.

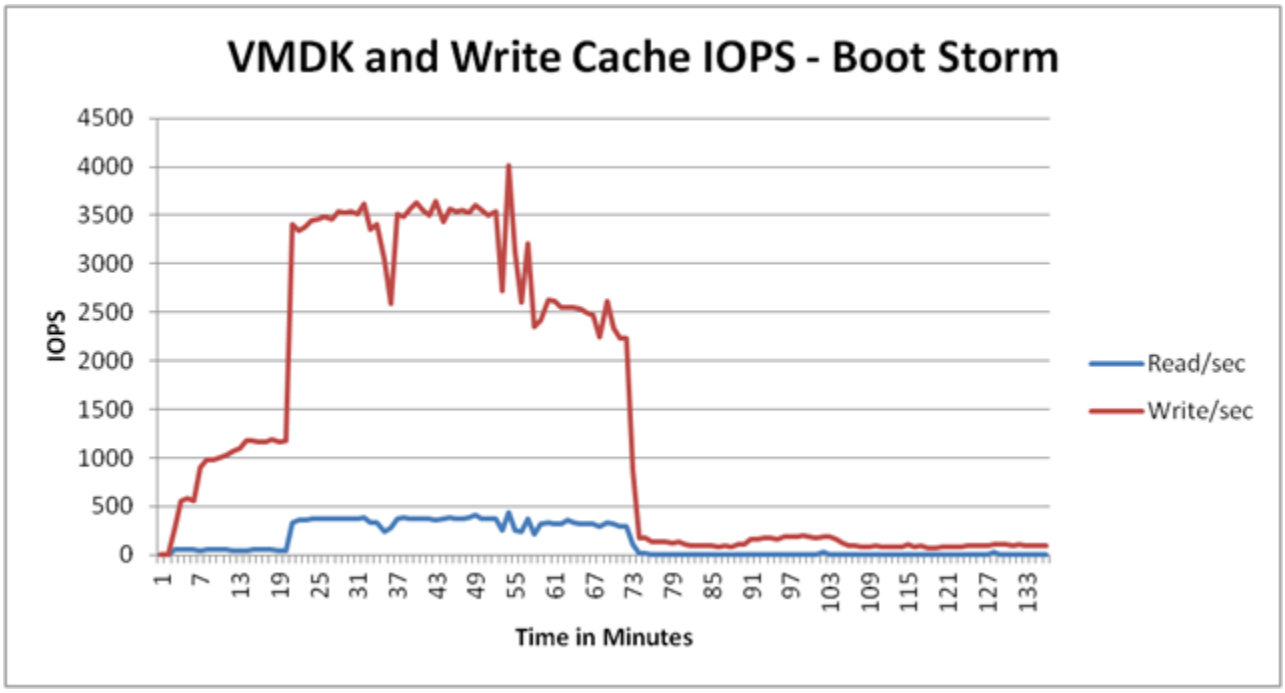


**Figure 24**

The data in Figure 24 proves the underlying storage supports a managed power-on of the 500 desktops to have them ready for use in an hour.

Figure 25 on page 49 illustrates the percent busy statistics for a representative drive from a dynamic provisioning pool for a virtual machine disk files and write cache during the boot storm test.

This physical drive showed the following:

- The drive peaked at 20% during the boot storm test.

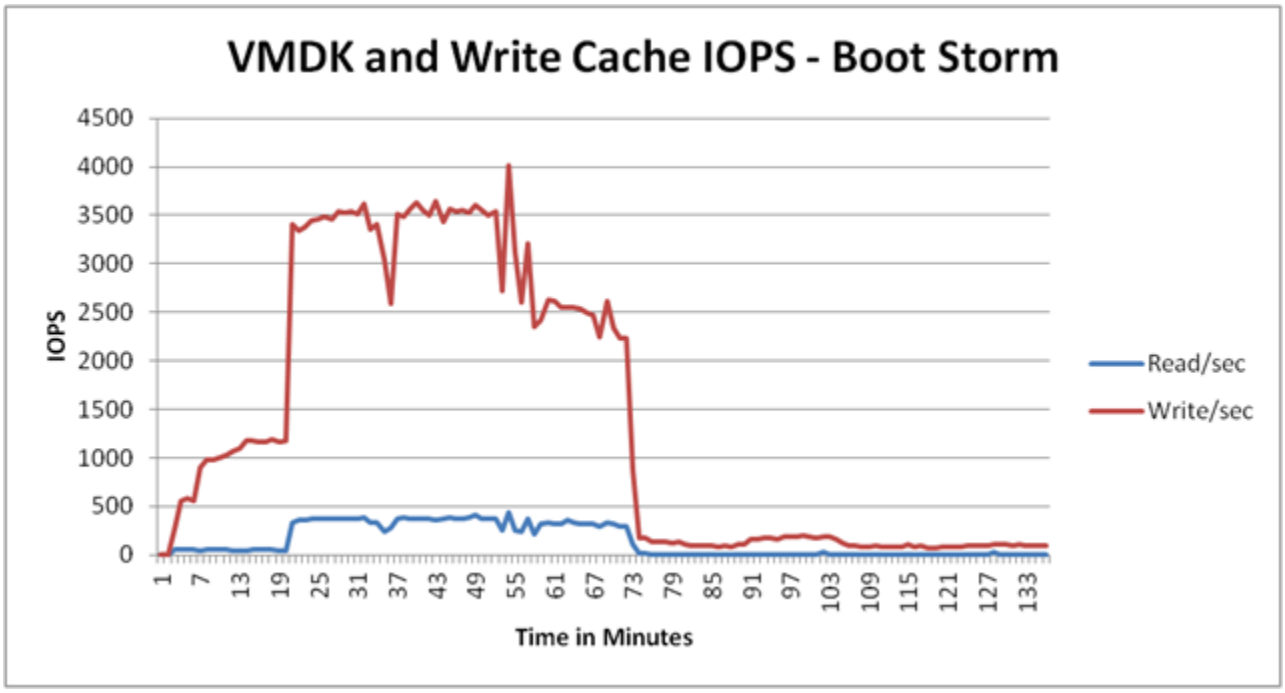


**Figure 25**

The percent busy rate seen in Figure 25 remained low during the boot storm test, showing that there is room to support bursts.

- Write pending percentage averaged below four percent.
- Core utilization peaked around eight percent.

Figure 26 on page 50 illustrates the processor and cache performance metrics.

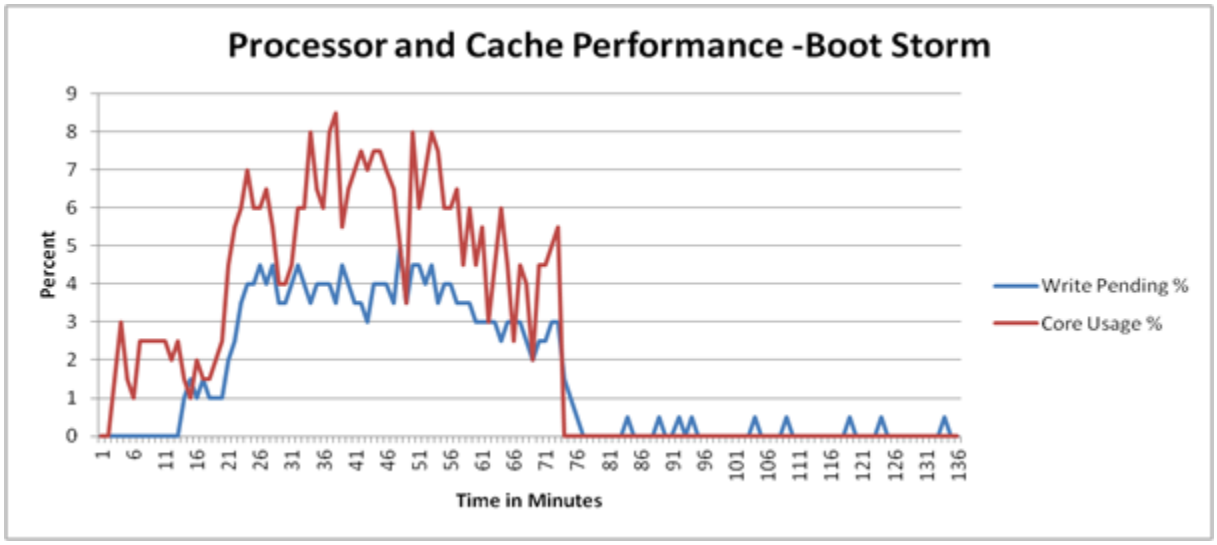


**Figure 26**

The default Citrix XenDesktop Power Management settings should be sufficient in most environments. If not sufficient, configure the settings under advanced settings of the host details in Desktop Studio. If you decide to increase these settings, make small incremental changes and monitor the affects on the hypervisor and storage to make sure you do not overwhelm the systems.

# Conclusion

This reference architecture guide shows how to design a Hitachi Unified Compute Platform Select for Citrix XenDesktop solution using Hitachi Compute Blade 500 and Hitachi Unified Storage 150. The cell design validated in the Hitachi Data Systems laboratory allows a build-as-you-go model with performance-proven sets of hardware resources.

By using this modular approach, you can scale an environment from 500 light workload users to the VMware vSphere configuration maximums. Create a right-sized design that allows purchasing flexibility to meet changing business or project needs.

Understand the I/O workload of the desktops in your Citrix XenDesktop environment so you properly design the architecture. This reduces costs and increases ROI.

Using Hitachi Dynamic Provisioning to create dynamic provisioning pools in this reference architecture allows increasing or decreasing I/O requirements dynamically, as necessary. Having the capability to provision additional spindles to an already-provisioned datastore within VMware vSphere allows for non-disruptive upgrades to the underlying storage infrastructure. This provides immediate benefits to the environment with no confusing shuffling of virtual machines, datastores, or LUs.

Further use of Hitachi Dynamic Provisioning to create dedicated dynamic provisioning pools for different types of XenDesktop users could drive utilization and ROI even higher.

# For More Information

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 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 Services Education 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**
2845 Lafayette Street, 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



AS-185-01, January 2013