

# Deploying VMware® View™ 4.6 with VMware vSphere™ 4.1 on the Hitachi Virtual Storage Platform

Reference Architecture

*By Bob Ingram and Michael Nakamura*

*August 2011*

## Feedback

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

# Table of Contents

<b>Solution Overview .....</b>	<b>3</b>
<b>Key Solution Components.....</b>	<b>6</b>
Hitachi Virtual Storage Platform .....	6
Hitachi Dynamic Provisioning.....	6
Hitachi Compute Blade 2000.....	7
VMware vSphere 4 .....	7
VMware View 4.6 .....	7
<b>Solution Design .....</b>	<b>8</b>
Infrastructure Block .....	8
Desktop Block .....	9
Server Design.....	10
Storage Design.....	11
SAN Design .....	12
VMware View Design .....	13
Scalability .....	15
<b>Engineering Validation.....</b>	<b>16</b>
512 Virtual Machine Baseline Environment .....	16
Workload Tests .....	17
Test Results Summary .....	17
<b>Conclusion .....</b>	<b>23</b>
<b>Appendix — VMware RAWC.....</b>	<b>24</b>

# Deploying VMware® View™ 4.6 with VMware vSphere™ 4.1 on the Hitachi Virtual Storage Platform

## Reference Architecture

The purpose of this reference architecture is to lower administration costs, improve system efficiency, and maintain performance. Using a building block design, this reference architecture helps IT organizations deploy VMware View desktop in a virtualized environment using VMware vSphere, Hitachi Dynamic Provisioning, and Hitachi Virtual Storage Platform.

This reference architecture assumes an average end user workload and provides guidance for sizing and deploying a Microsoft Windows 7 virtual desktop infrastructure (VDI) supporting 1,024 users per building block. This 1,024 user building block is the foundation to scale up to many thousands of users by adding additional building blocks.

This paper's intended use is by IT administrators responsible for desktops and storage administration. It assumes some familiarity with Hitachi Storage Navigator, VMware vSphere, and VMware View Composer.

---

**Note**—All testing was done in a lab environment. In production environments, results can be affected by many factors that cannot be predicted or duplicated in a lab. Conduct proof-of-concept testing using your target applications in a non-production, isolated test environment that is identical to your production environment. Following this best practice allows you to obtain results closest to what you can expect to experience in your deployment.

---

## Solution Overview

One of the most important factors to consider when sizing a VDI environment for performance is defining the type of end user workload. Hitachi Data Systems consulted with VMware to determine various end user workload levels. Three main types of workers and the expected type of applications utilized are categorized in Table 1.

**Table 1. Workloads**

<i>Workloads</i>	<i>User Workload</i>
Task-Based User	Microsoft Outlook Microsoft Word Microsoft Internet Explorer
Knowledge-Based User	Microsoft Outlook Microsoft Word Microsoft Excel Microsoft PowerPoint Microsoft Internet Explorer

<i>Workloads</i>	<i>User Workload</i>
Power User	Microsoft Outlook Microsoft Word Microsoft Excel Microsoft PowerPoint Microsoft Internet Explorer Adobe Reader 7-ZIP Windows Media Player

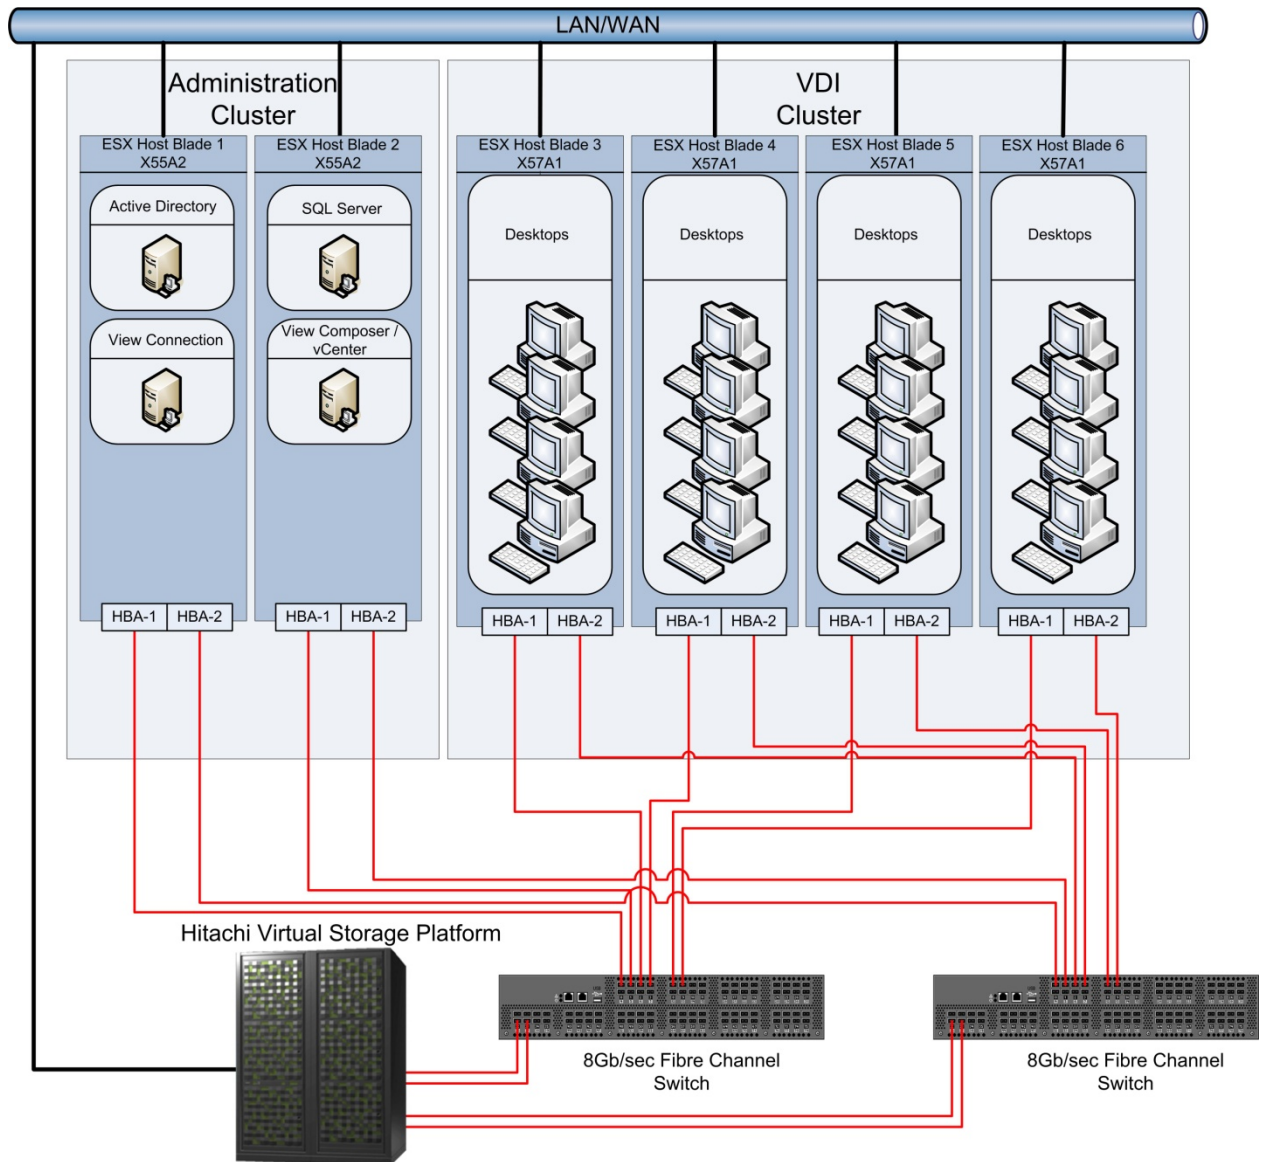
This reference architecture focuses on one of the three types of users described in Table 1, the knowledge-based user, to build out a scalable and converged server and storage VDI solution using Hitachi components.

The following are general characteristics of a knowledge-based user, as defined by the VMware Desktop RAWC tool:

- **Knowledge-based User Virtual Machines**—This level places an average of 16 virtual machines per CPU core. These users do the following:
  - Use Microsoft Word for editing medium document files
  - Use Microsoft PowerPoint and Microsoft Excel
  - Send and receive an average number of email messages per day
  - Use Adobe Reader for viewing simple PDF files with graphics
  - Browse simple Web sites

The end user workload is a key factor in determining the number of users that can be supported in a building block.

Figure 1 is an overview of the VMware vSphere environment used in the Hitachi Data Systems lab to validate the building block used in this reference architecture.



**Figure 1**

To achieve high availability, this reference architecture used the following:

- Redundant physical Fibre Channel connections using two host bus adapters (HBAs) per ESX host
- vSphere clustering using Distributed Resource Scheduler (DRS) and high availability (HA)

## Key Solution Components

The following are descriptions of the components used in this reference architecture.

### Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is a 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 performance, capacity, and connectivity by adding cache, processors, connections, and disks to the base system.
- **Scale Out**—Combine multiple chassis into a single logical system with shared resources.
- **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. Consolidating more virtual machines on a physical host requires storage systems to be able to dynamically add more storage resources to keep pace with I/O demand. The 3D scaling capability of the Virtual Storage Platform meets that requirement.

For more information, see the [Hitachi Virtual Storage Platform](#) on the Hitachi Data Systems website.

### Hitachi Dynamic Provisioning

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

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

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

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

Hitachi Dynamic Provisioning in a VMware vSphere View environment allows a leveling of the workload across a large number of disks. This allows sizing the disks for an average workload instead of a peak burst workload.

The advantages of using Hitachi Dynamic Provisioning in VDI are:

- **Performance Balance**—Hitachi Dynamic Provisioning distributes data across all disks in the pool.
- **Over Provisioning**—Hitachi Dynamic Provisioning allows over provisioning of the storage.

For more information, see the [Hitachi Dynamic Provisioning datasheet](#) and [Hitachi Dynamic Provisioning](#) on the Hitachi Data Systems website.

## Hitachi Compute Blade 2000

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

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

## VMware vSphere 4

VMware vSphere 4 is a highly efficient virtualization platform that provides a robust, scalable, and reliable infrastructure for the data center. It features Distributed Resource Scheduler (DRS), high availability (HA), and fault tolerance (FT).

Use of vSphere 4's round robin multipathing policy distributes load across multiple host bus adapters (HBAs) and multiple storage ports. Use of DRS with Hitachi Dynamic Provisioning automatically distributes loads on the ESX host and across the back end of the storage system.

VMware vSphere 4 has the following components:

- **ESX 4.1**—ESX 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**—vCenter Server allows management of the vSphere environment through a single user interface. With vCenter, advanced vSphere features are available, such as vMotion, Storage vMotion, Distributed Resource Scheduler, high availability, and fault tolerance.

## VMware View 4.6

VMware View 4.6 provides virtual desktops as a managed service. Using View, create images of approved desktops and then deploy them automatically as needed. Desktop users quickly access their personalized desktop, including data, applications, and settings from anywhere with network connectivity to the server. PCoIP, a high performance display protocol, provides enhanced end-user experience compared to traditional remote display protocols.

VMware View allows management of software upgrades and patches and software license usage. View lets the desktops take advantage of key vSphere features, such as vMotion, high availability and Dynamic Resource Scheduler, to provide robust and highly available desktops.

View can deploy desktop clones from a single OS image. Or, use View Composer to deploy virtual desktops linked to a master OS image.

VMware View has the following components:

- **View Connection Server**—The View Connection Server runs a broker service for client connections. Authenticating users through Windows Active Directory services, it directs them to the appropriate virtual desktop or terminal services client. It permits a single logon to give access to a secure environment.
- **View Agent**—Install View Agent on virtual machines, terminal services, clients and physical machines that will be used as View desktops.
- **View Client**—Use View Client for connection to View desktops. It runs on the Microsoft Windows and Apple Macintosh operating systems as an application or on a thin client.
- **View Administrator**—Administrators deploy and manage View desktops, control user authentication, and manage View using View Administrator as a web-based application.
- **View Composer**—Install View Composer on the vCenter server to allow administrators to create pools of linked clones virtual machines. Linked clones can provide significant storage savings by linking many virtual machines to a single virtual machine image. Because linked clones share a single image, managing patches and updates are easier. View Composer allows tiering of storage by allowing the persistent disks for end user data to be placed on a different VMFS than the OS image.

For more information see the [VMware View](#) website.

## Solution Design


Using the key solution components described in the previous section, this reference architecture provides the infrastructure details and methodology for calculating the compute and storage requirements necessary to build a 1,024 user environment while sustaining ideal performance for knowledge-based user workloads.

### Infrastructure Block

The Hitachi Compute Blade 2000 X55A2 was used for the View environment infrastructure block. Table 2 shows the hardware configuration with the two server blades in Chassis 1 as shown in Figure 2.

**Table 2. Chassis 1 Hitachi Compute Blade 2000 X55A2 Configuration**

<i>Hardware</i>	<i>Detail Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Compute Blade 2000 X55A2 blades	Full blade 2 × 6-Core Intel Xeon X7670 processors 64GB Memory per blade	04-21	2



**Figure 2**

The infrastructure block consists of server blades running VMware ESX that are dedicated to hosting administrative machines. Table 3 lists the major software components used.

**Table 3. Software Components**


<i>Software</i>	<i>Version</i>
VMware vCenter server	4.0.0 build 140742
VMware View	4.6 build 366101
VMware View Composer	2.6 build 349155
VMware ESX	4.1.0 build 260247
Microsoft Windows Server	2008 R2, 64 bit Enterprise
Microsoft SQL Server	2008 64 bit Enterprise

## Desktop Block

The Hitachi Compute Blade 2000 X57A1 was used for the View environment desktop block. The desktop block is the basis for a building block design from which multiple desktop blocks can be added to scale out your VDI environment. Table 4 shows the desktop block hardware configuration with the four server blades in Chassis 2, as shown in Figure 3.

**Table 4. Chassis 2 Hitachi Compute Blade 2000 X57A1 Configuration**

<i>Hardware</i>	<i>Detail Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Compute Blade 2000 X57A1 blades	Full blade 2 x 8-Core Intel Xeon X7560 processors 256GB Memory per blade	04-21	4



**Figure 3**

The desktop block consists of server blades running VMware ESX that are dedicated to hosting all virtual desktop machines which users will remotely access. Table 5 lists the major software components used.

**Table 5. Software Components**

<i>Software</i>	<i>Version</i>
VMware ESX	4.1.0 build 260247
VMware View Agent	4.6 build 366101
Microsoft Windows	7, 64 bit Enterprise

## Server Design

Calculate the scaling for a server block using the following recommendations, taken from “Calculating Maximum Users by Hardware” and “Calculating Maximum Users by Memory.” Use the smaller result for the recommended maximum number of virtual user machines in the environment.

### *Calculating Maximum Users by Hardware*

The formula to calculate the approximate maximum number of virtual user machines that can be hosted based on server hardware is the following:

$$\text{Maximum Number of Virtual User Machines based on Hardware} = (\text{Number of Host Server Blades}) \times \left( \text{Number of } \frac{\text{Processors}}{\text{Server Blade}} \right) \times \left( \text{Number of } \frac{\text{CPU Cores}}{\text{Processor}} \right) \times \left( \text{Number of } \frac{\text{Virtual Machines}}{\text{CPU Core}} \right)$$

VMware recommends 16 virtual machines per CPU core for knowledge-based users.

Using the four server blade configuration and the above formula, this is the maximum number of users from a CPU perspective:

$$1,024 \text{ users} = (4 \text{ Host Server Blades}) \times \left( 2 \frac{\text{Processors}}{\text{Blade}} \right) \times \left( 8 \frac{\text{CPU Cores}}{\text{Processor}} \right) \times \left( 16 \frac{\text{Virtual Machines}}{\text{CPU Core}} \right)$$

## Calculating Maximum Users by Memory

The formula to calculate the approximate maximum number of virtual user machines that can be hosted based on server memory is the following:

$$\text{Maximum Number of Virtual User Machines based on Memory} = \frac{(\text{Number of Host Server Blades}) \times (\text{Memory in Host Server Blades}) + [(\text{Memory in Host Server Blades}) \times (\text{Over Commitment Percentage})]}{\left(\text{Amount of } \frac{\text{Memory}}{\text{Virtual Machine}}\right)}$$

VMware recommends 1GB per virtual machine for knowledge-based users.

Using the Hitachi Compute Blade 2000 4 server blade configuration and the above formula, the following sizing numbers can be reached from a memory perspective:

$$1,024 \text{ users} = (4 \text{ Host Server Blades}) \times \frac{(256\text{GB Memory in Host Server Blades})}{\left(\frac{1\text{GB Memory}}{\text{Virtual Machine}}\right)}$$

The Hitachi Compute Blade 2000 is available with a variety of hardware configurations. When selecting your hardware configuration, keep the above formulas in mind to ensure that you are able to host the maximum number of virtual machines relative to the CPU and memory available.

Calculations from the CPU and memory sizing formula above show that four Hitachi Compute Blade 2000 server blades, each with two 8-core processors and 256GB of RAM, is an ideal configuration for 1,024 users.

## Storage Design

For this reference architecture, Hitachi Dynamic Provisioning was used on the Hitachi Virtual Storage Platform to define the storage block. This simplifies the management of the storage through the use of over provisioning, wide striping and on-line expansion of dynamic provisioning pools.

A storage block is comprised of a single dynamic provisioning pool. The sizing of the pool was based on assumptions made about the disks used and VDI environments.

- **Disk**—Disk speed, RAID level, and the required IOPS
- **VDI**—Estimated IOPS for knowledge user workloads

The initial calculations for disk I/O are based on the estimated IOPS for 1,024 knowledge-based users. Industry standard numbers were used to estimate this number of IOPS. Table 6 shows the initial calculations for IOPS.

**Table 6. Peak IOPS for 1,024 Virtual Machines**

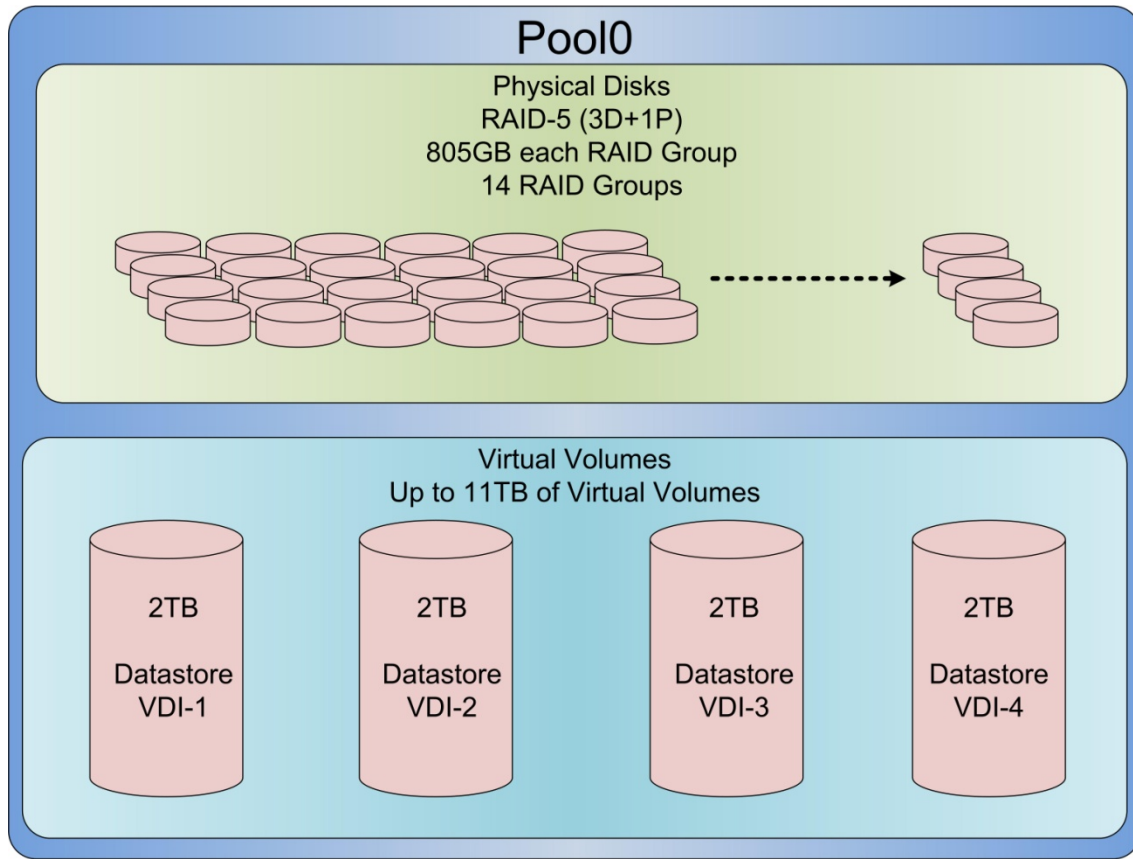
<i>User Type</i>	<i>Peak IOPS Estimates for 1 Virtual Machine</i>	<i>Peak IOPS Estimates for 1,024 Virtual Machines</i>
Knowledge-Based Users	8	8,192

Taking the knowledge user workload as an example, a dynamic provisioning pool with 14 RAID groups consisting of 56 300GB 10k RPM SAS drives in a RAID-5 (3D+1P) configuration was built out to support a building block of 1,024 knowledge users. This provides up to 11TB of space and supports the average IOPS required.

This reference architecture used dynamic provisioning volumes of 2TB without oversubscription. While the aggregate of the dynamic provisioning volume can be oversubscribed, ESX 4.1 only supports a maximum LUN size of 2TB.

Four 2TB VMFS LUNs were provisioned to the four blade servers with each LUN capable of hosting 256 virtual machines.

Figure 4 shows the configuration for the dynamic provisioning pool.



**Figure 4**

The VAAI hardware-assisted locking mechanism was used for this reference architecture.

## SAN Design

The Hitachi Virtual Storage Platform and two Hitachi Compute Blade 2000s are connected to the SAN through two Brocade 5400 switches.

Each of the two server blades in the infrastructure block and each of the four server blades in the desktop block have two Fibre Channel connections with a connection to each Brocade switch for redundancy and load balancing. VMware native multipathing round-robin balances the I/O across both connections.

The Hitachi Virtual Storage Platform has four Fibre Channel connections, with two connections to each Brocade switch. The Brocade SAN uses single initiator zoning.

Figure 5 shows the SAN connectivity.

Hitachi Compute Blade 2000

Hitachi Virtual Storage Platform





Figure 5

## VMware View Design

Table 7 shows each virtual machine's configuration that was installed on the VMware vSphere infrastructure cluster consisting of two Hitachi Compute Blade 2000 server blades. This configuration is shown in Figure 6.

Table 7. VMware View Virtual Machines

<i>Virtual Machine</i>	<i>Configuration</i>
Microsoft Active Directory	vCPU—1 vMemory—4GB
VMware vCenter	vCPU—1 vMemory—4GB
Microsoft SQL 2008 database virtual machine for vCenter and View	vCPU—2 vMemory—4GB
VMware View Connection	vCPU—4 vMemory—16GB



**Figure 6**

A single virtual machine, configured with 1 vCPU and 1GB of RAM, was used as a master image template for the virtual desktop. Microsoft Windows 7 Enterprise and the following applications were installed:

- VMware View 4.6 agent
- Microsoft Office 2007
- Microsoft Internet Explorer 8
- Microsoft Media Player 12
- Adobe Reader 9
- McAfee Virus Scan Enterprise 8.7.0i
- Sun Java JDK 6 Update 16

This reference architecture uses VMware View Composer linked clone desktops. Using View Composer linked clones reduces the storage footprint of the virtual desktops.

View Composer allows the creation of persistent virtual disks so that end user data can be stored on a separate datastore. This allows for storage tiering. For example, you can store the data for power users on high speed SAS disks, while storing data for low usage users on low cost SATA disks within the same Hitachi Virtual Storage Platform.

Four pools of VMware View Composer linked clone desktops were used, one for each blade server hosting 256 virtual machines. Each pool was based on the master image template described above.

## Scalability

This reference architecture has defined a 1,024 knowledge user building block with the following:

- **Desktop Block** — 4 × Hitachi Compute Blade 2000 X57A1 blades each with 2 × 8-Core Intel Xeon X7560 processors and 256GB memory per blade
- **Storage Block** — A dynamic provisioning pool consisting of 56 300GB 10k RPM SAS drives in 14 RAID-5 (3D+1P) groups
- **Virtual Machines** — 1 vCPU and 1GB RAM

In order to scale out the environment in increments of 1,024 users each, each building block requires an additional desktop block of four server blades plus a storage block with a dynamic provisioning pool with 14 RAID groups, as defined in “Infrastructure Block,” “Storage Design,” and “SAN Design.”

Scale to larger number of users by adding one or more additional building blocks, as shown in Figure 7.

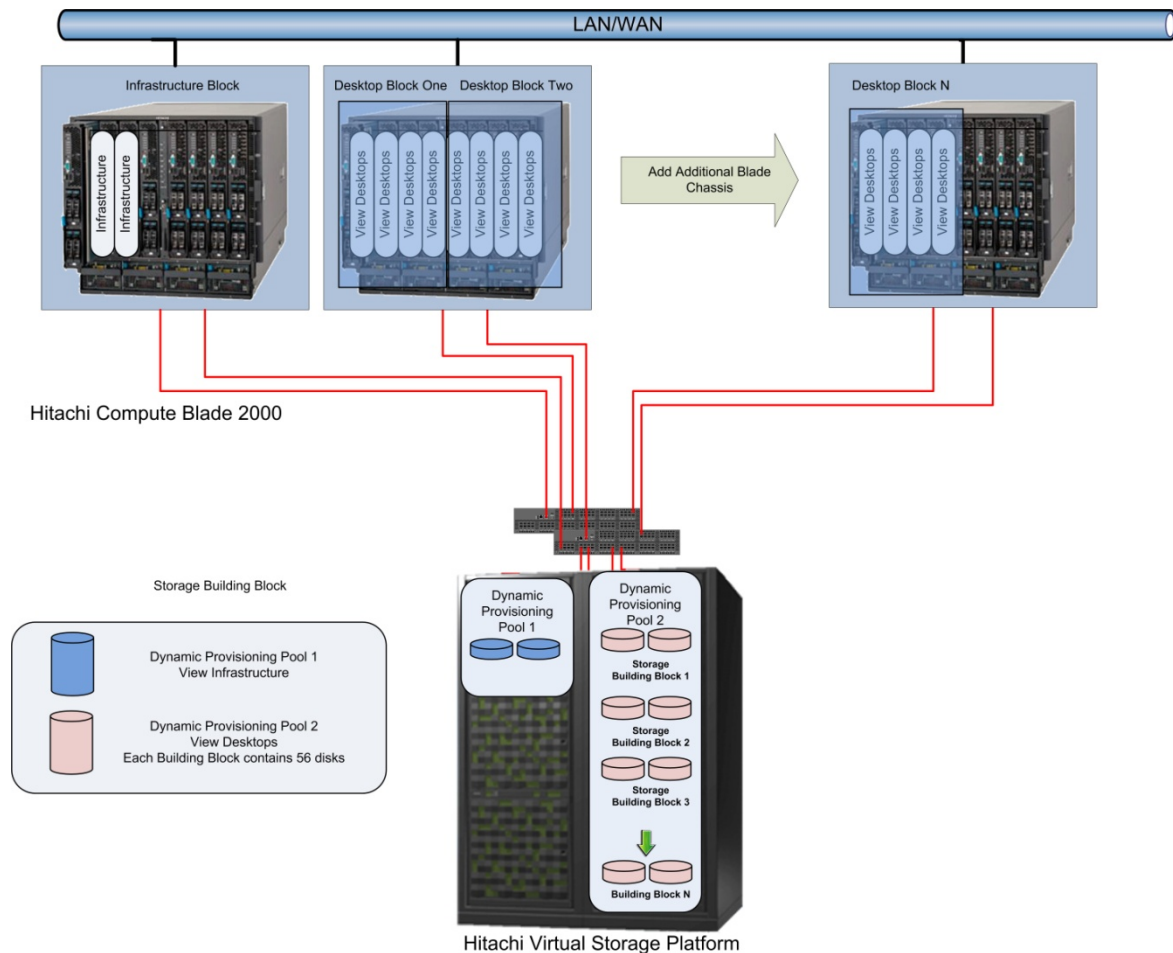


Figure 7

You can add additional server blades to the existing DRS cluster. When adding ESX hosts to the DRS cluster,

- Zone the SAN for connectivity to the Fibre Channel ports.
- Create host groups containing all dynamic provisioning volumes.

Execute DRS in automatic or manual mode to re-distribute the workload across the additional ESX hosts.

You can add more RAID-5 (3D+1P) groups to the dynamic provisioning pool. With dynamic provisioning, the data automatically re-distributes across the newly-added RAID groups. When scaling to more than a total of eight ESX hosts, Hitachi Data Systems recommends configuring additional Fibre Channel ports on the Hitachi Virtual Storage Platform for additional I/O bandwidth.

There are VDI maximums imposed at the vCenter, DRS cluster, and View server level when scaling out. See [Configuration Maximums for VMware vSphere 4.1 Guide](#) for details.

## Engineering Validation

While the design of the reference architecture is for a 1,024 virtual machine environment for knowledge-based users, a 512 virtual machine environment for knowledge-based users was initially built to provide a performance baseline for comparison. Then, tests were run on a 1,024 virtual machine environment.

### 512 Virtual Machine Baseline Environment

Concepts used in “Solution Design” were applied to create the 512 virtual machine baseline environment:

- The virtual machine density per CPU core. See “Calculating Maximum Users by Hardware”.

$$512 \text{ users} = (2 \text{ Host Server Blades}) \times \left(2 \frac{\text{Processors}}{\text{Blade}}\right) \times \left(8 \frac{\text{CPU Cores}}{\text{Processor}}\right) \times \left(16 \frac{\text{Virtual Machines}}{\text{CPU Core}}\right)$$

- The virtual machine memory requirements. See “Calculating Maximum Users by Memory”

$$512 \text{ users} = (2 \text{ Host Server Blades}) \times \frac{(256\text{GB Memory in Host Server Blades})}{\left(\frac{1\text{GB Memory}}{\text{Virtual Machine}}\right)}$$

- The IOPS requirement for the 512 virtual machine workload is in Table 8.

**Table 8. Peak IOPS for 512 Virtual Machines**

<i>User Type</i>	<i>Peak IOPS Estimates for 1 Virtual Machine</i>	<i>Peak IOPS Estimates for 512 Virtual Machines</i>
Knowledge-Based Users	8	4096

A dynamic provisioning pool consisting of 28 300GB 10k RPM SAS drives in 7 RAID-5 (3D+1P) groups was built to support a building block of 512 knowledge-based users. This provides up to 5.5TB of space and supports the average IOPS required.

As data capacity was not a factor and ESX 4.1 supports a maximum LUN size of 2TB, two 2TB LUN's were provisioned to the two blade servers, with each LUN capable of hosting 256 virtual machines.

- Determining the View Desktop Pool design

Two pools of VMware View Composer linked clone pools were created, one for each blade server supporting 256 virtual machines.

This design was used to keep a one-to-one host server blade to datastore architecture.

## Workload Tests

The workloads were driven using VMware Desktop Reference Architecture Workload Code (RAWC). This tool allows the tester to generate workloads using Microsoft Office, Adobe Viewer, Microsoft Media Player, Java, and other common desktop applications. There is more information on the RAWC configuration in "Appendix—VMware RAWC."

Data was collected from the following:

- *esxtop*
- Hitachi Virtual Storage Platform performance export tool

The key metrics analyzed were:

- ESX host CPU utilization
- Hardware device response time
- IOPS
- Disk utilization rate

Two test cases were run to capture performance metrics:


- **Baseline 512 Virtual Machines**—Two RAWC Controller virtual machines were configured with 20 session launcher virtual machines to run three iterations of the RAWC desktop workload simulating typical knowledge-based user behavior.
- **Complete 1,024 Virtual Machines**—Four RAWC controller virtual machines were configured with 40 session launcher virtual machines to run 3 iterations of the RAWC desktop workload simulating typical knowledge-based user behavior.

## Test Results Summary

The presentation of test results is by test stage.


The 512 virtual machine tests were performed using two Hitachi Compute Blade blades and seven RAID-5 (3D+1P) groups.

Figure 8 shows average CPU Percent busy rate for the 512 virtual machine test.



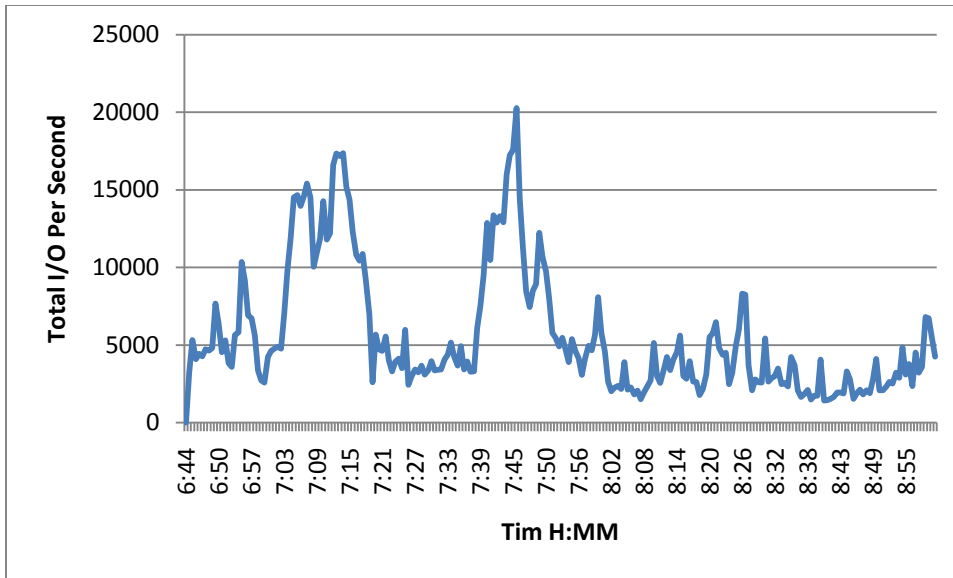
**Figure 8**

Figure 9 shows average free memory for the ESX hosts during the 512 virtual machine test.




**Figure 9**

Figure 10 shows the total IOPS during the 512 virtual machine test.




**Figure 10**

Figure 11 shows the disk percent busy rate during the 512 virtual machine test.



**Figure 11**

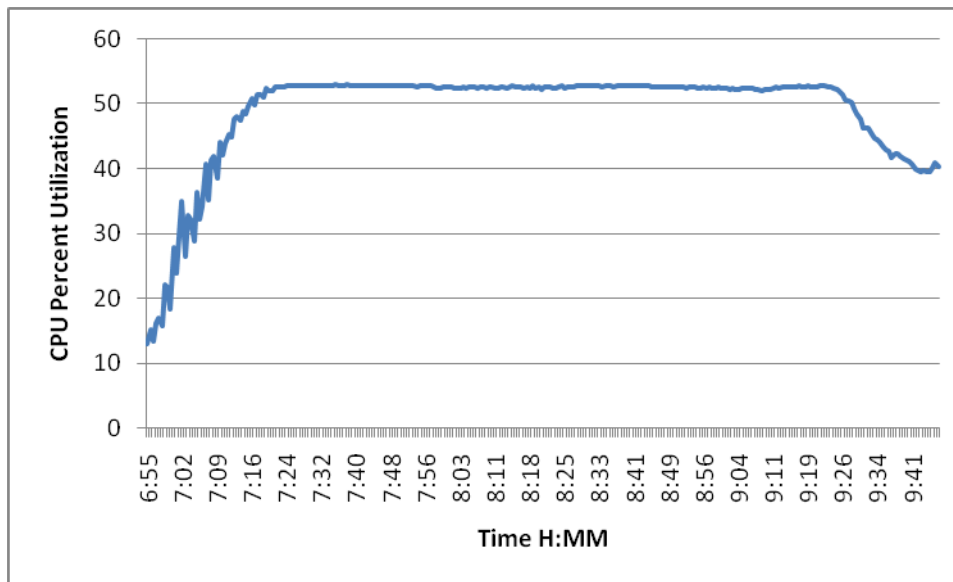
Figure 12 shows the hardware device response time during the 512 virtual machine test.



**Figure 12**


The 1,024 virtual machine tests used 4 Hitachi Compute Blade server blades and 14 RAID-5 (3D+1P) groups.

Figure 13 shows the average CPU percent busy rate for the 1,024 virtual machine test




**Figure 13**

Figure 14 shows average free memory for the ESX hosts during the 1,024 virtual machine test.



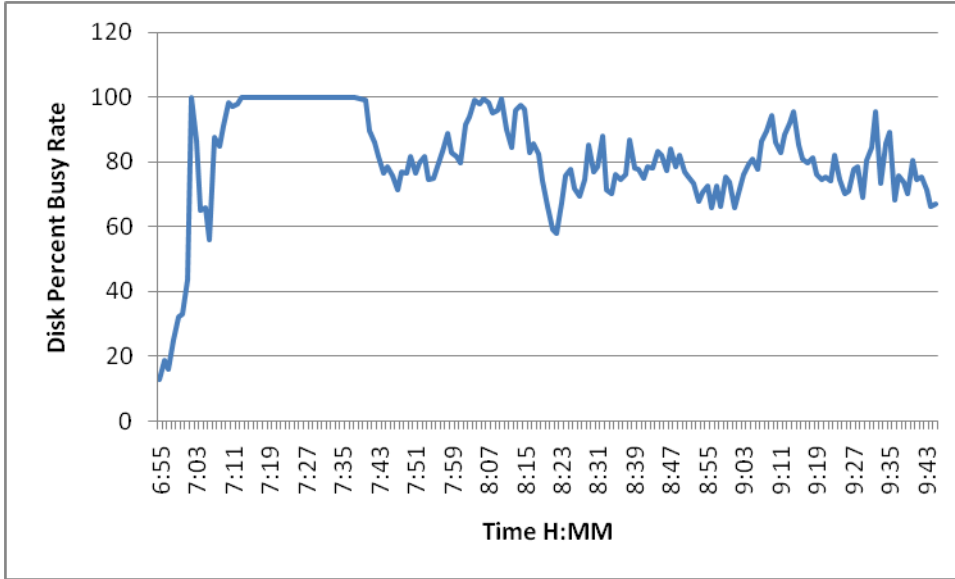
**Figure 14**

Figure 15 shows the total IOPS during the 1,024 virtual machine test.




**Figure 15**

Figure 16 shows the disk percent busy rate during the 1,024 virtual machine test.



**Figure 16**

Figure 17 shows the hardware device response time during the 1,024 virtual machine test.



**Figure 17**

***Scaling From 512 to 1,024 Virtual Machine Workload Test Analysis***

For the 512 virtual machine test, two of the Hitachi Compute Blade 2000 server blades were used with a dynamic provisioning pool containing seven RAID-5 (3D+1P) groups.

For the 1,024 user test, four server blades were used with a dynamic provisioning pool was expanded to 14 RAID-5 (3D+1P) groups.

Table 9 shows the results captured during a steady state of workloads to exclude skewing caused by logon storms.

**Table 9. Steady State Test Results Summary**

<i>Test</i>	<i>ESX Host Average CPU Utilization</i>	<i>ESX Host Average Free Memory</i>	<i>IOPS</i>	<i>IOPS/User</i>	<i>Response Time</i>	<i>Disk Utilization</i>
512 Virtual Machines 2 Server Blades 7 RAID Group Pool	49%	17GB	17,000 peak 3,900 average	7.6	2ms	100% peak 70% average
1,024 Virtual Machine 4 Server Blades 14 RAID Group Pool	53%	19GB	18,000 peak 7,100 average	6.9	3ms	100% peak 80% average

The results from Table 9 show the average CPU utilization on the ESX hosts and free memory indicates higher levels of workload could be supported from a compute perspective.

The average IOPS per user was around 7, indicating the generated RAWC desktop workload was close to our assumption of 8 IOPS per user for knowledge-based users.

Disk percent utilization was observed with peaks of 100%. However, this was not a concern as response time was below 10ms. Response time greater than 10ms to the vmkernel can result in poor user experience.

Based on the scaling test results in the Hitachi Data Systems lab, the View environment can scale by adding additional desktop and storage blocks supporting increments of 1,024 users.

## Conclusion

Testing shows that the Hitachi Virtual Storage Platform and the Hitachi Compute Blade 2000 make an ideal platform for hosting a scalable VDI implementation using VMware View running on VMware vSphere. This reference architecture supports up to 1,024 knowledge-based users.

Using a building block design consisting of desktop and storage blocks can simplify the process of architecting a VDI environment.

When sizing out a VDI environment, characterize the user workload before migrating to a VMware View environment. While this document provides proven methods to assist with the sizing and testing of your VDI environment, deployments vary greatly and the design implemented in the Hitachi Data Systems Lab may not reflect your implementation. The performance metrics provided are for guidance. A full desktop capacity planning assessment provides more precise sizing requirements for 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 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](#).

## Appendix — VMware RAWC

VMware RAWC was installed on two Hitachi Compute Blade 2000 server blades dedicated to administrative tasks.

- **RAWC Controller** — Defines the test configuration parameters and opens VMware RAWC Session Launcher virtual machines.
- **RAWC Session Launcher** — Initiates the user's logon to the VMware View Composer virtual machines. Each RAWC Session Launcher virtual machine can initialize approximately 20 View sessions.

Figure 18 shows the RAWC and View configuration.




Figure 18

Table 10 describes the VMware RAWC configuration.

**Table 10. VMware RAWC Configuration**

<i>Virtual Machine Role</i>	<i>Configuration</i>	<i>Quantity</i>
RAWC Session Launcher	vCPU—1 vMemory—4GB	40
RAWC Controller	vCPU—1 vMemory—4GB	4

VMware RAWC generates workloads based on the following applications.

- **Microsoft Word**—Open, minimize, maximize, close, insert text, save modifications
- **Microsoft Word (Random)**—Open, minimize, maximize, close, write random words/numbers, save modifications
- **Microsoft Excel**—Open, minimize, maximize, close, write random numbers, insert/delete columns/rows, copy/paste formulas, save modifications
- **Microsoft PowerPoint**—Open, minimize, maximize, close, conduct a slide show presentation
- **Microsoft Outlook**—Open, minimize, maximize, close, create/send emails
- **Internet Explorer**—Open, minimize, maximize, close, browse page
- **Windows Media Player**—Open, close, view a video
- **Java**—Compile a java project comprised of several hundred files using the JDK to create a software engineering type of workload
- **Adobe Reader**—Open, minimize, maximize, close, browse pages in PDF document
- **7-Zip**—Open, close, compress a large file

 **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 2011. All Rights Reserved. AS-102-00 August 2011

**Corporate Headquarters**

750 Central Expressway,  
Santa Clara, California 95050-2627 USA  
[www.hds.com](http://www.hds.com)

**Regional Contact Information**

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