TECHNICAL PAPER

Deploy Hitachi Unified Compute Platform Select for VMware vSphere using Hitachi Unified Storage VM in a Scalable Environment

Reference Architecture Guide

By Daniel Worden, Michael Nakamura, Henry Chu

January 9, 2013

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

Solution Overview.......... 3

Key Solution Components...5

Hardware Components.......... 5
Software Components.......... 7

Solution Design.......... 9

Infrastructure Cell for Compute Resources.......... 11
Infrastructure Cell for Storage Resources.......... 16
Application Cell for Hitachi Unified Compute Platform Select Management.......... 18
Application Cell for VMware vSphere.......... 22
Scaling Using Expansion Cell for Compute Resources.......... 28
Scaling Using Expansion Cell for Storage Resources.......... 30

Engineering Validation.......... 31

Test Methodology.......... 31
Test Results.......... 37

Conclusion.......... 46

Deploy Hitachi Unified Compute Platform Select for VMware vSphere using Hitachi Unified Storage VM in a Scalable Environment

Reference Architecture Guide

This is a reference architecture guide for Hitachi Unified Compute Platform Select with VMware vSphere using Hitachi Unified Storage VM. It contains advice on how to build a virtual infrastructure that meets the unique requirements of your organization, providing the flexibility for to scale out as organizational needs grow.

The benefits of this solution include the following:

- Faster deployment
- Reduced risk
- Predictability
- Ability to scale out
- Lower cost of ownership

Hitachi Unified Compute Platform is a family of completely integrated and flexible solutions. Each solution is configured for immediate deployment to run top-tier infrastructure applications without over-purchasing or provisioning unnecessary equipment. Each custom-built-solution has its entire solution stack-certified. There are no compatibility issues.

This reference architecture guide focuses on designing a virtual infrastructure capable of hosting virtual machines running general server application workloads. It is strongly recommended to run a server capacity-planning pilot to gather sizing and IOPS information before designing your environment.

You need familiarity with the use and configuration of the following to use this reference architecture guide:

- Hitachi Unified Storage
- Hitachi Compute Blade 500
- Hitachi Dynamic Provisioning
- VMware vSphere 5

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 a VMware infrastructure supported by Hitachi and Brocade hardware to create a flexible and pre-validated end-to-end converged stack solution.

This converged solution validates the integration of the hardware stack (compute, storage and networking) with the software stack (hypervisor and management for both software and hardware components).

The following components create this Unified Compute Platform Select for VMware vSphere environment:

- **Hitachi Compute Blade 500** — Enterprise-class server platform providing dense compute resources and high I/O throughput
- **Hitachi Unified Storage VM** — Hitachi Unified Storage VM storage virtualization system is designed for organizations that need to manage their storage assets more efficiently.
- **Hitachi Dynamic Provisioning** — Provides wide striping and thin provisioning functionalities for greater operational and storage efficiency
- **VMware vSphere 5** — Virtualization technology providing the infrastructure for the data center
- **Brocade VDX 6720 fabric switch** — 24-port 10 Gb/sec Ethernet fabric switch provides connectivity to the data center network
- **Brocade 6510 Enterprise-class fabric switch** — Provides SAN connectivity for Hitachi Unified Storage VM

Figure 1 on page 4 illustrates the high-level logical design of this reference architecture on Hitachi Unified Storage VM and Hitachi Compute Blade 500.

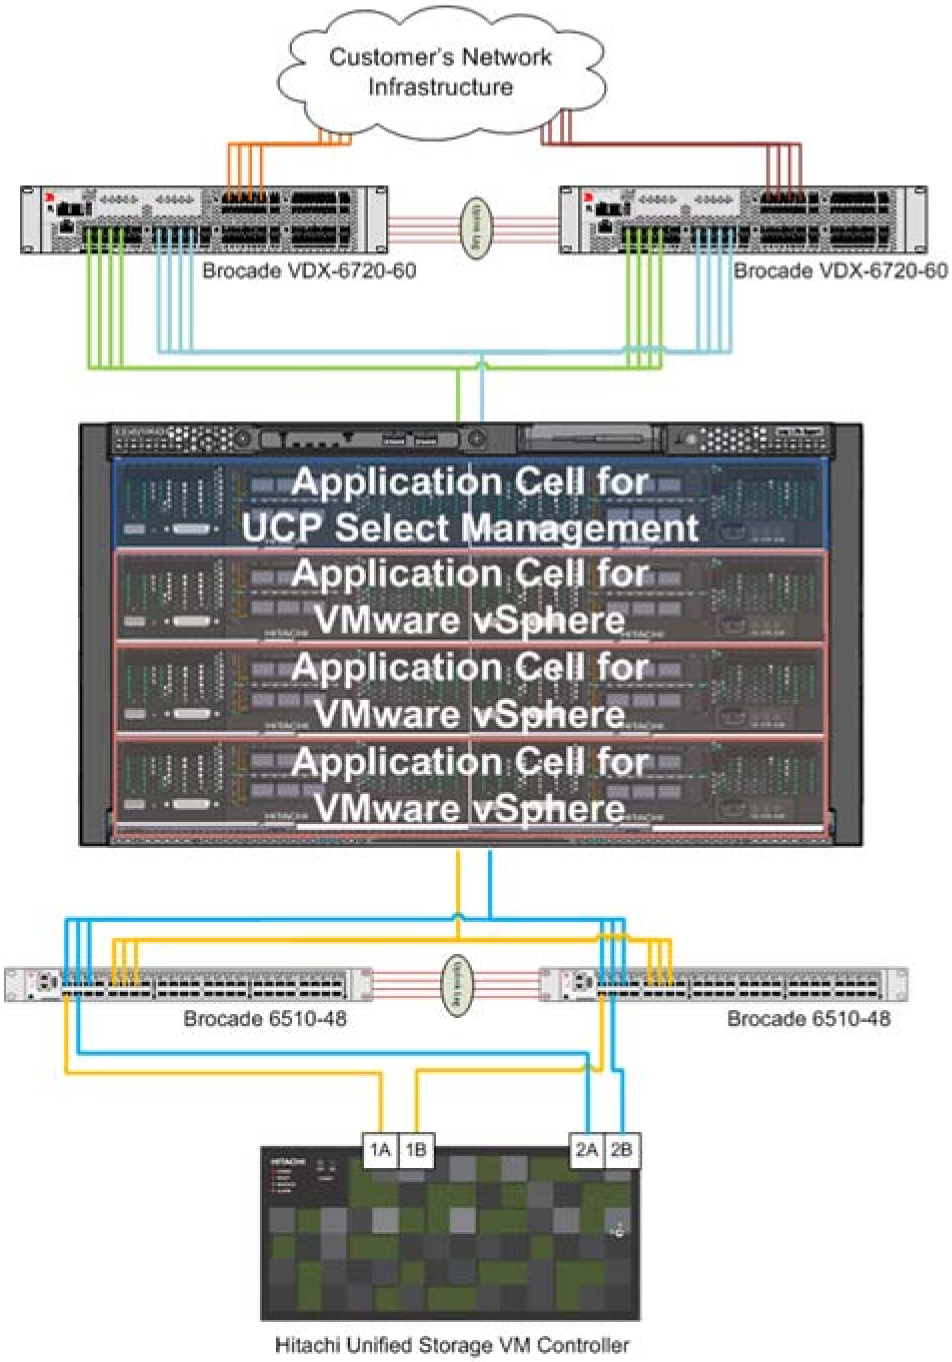
Customer's Network Infrastructure
Brocade VDX-6720-60
Brocade VDX-6720-60
Application Cell for UCP Select Management
Application Cell for VMware vSphere
Application Cell for VMware vSphere
Application Cell for VMware vSphere
Brocade 6510-48
Brocade 6510-48
1A
1B
2A
2B
Hitachi Unified Storage VM Controller

Figure 1

Key Solution Components

These are descriptions of the key hardware and software components used to deploy this Hitachi Unified Compute Platform for VMware vSphere reference solution with an entry level enterprise storage platform.

Hardware Components

Table 1 lists the detailed information about the hardware components used in the Hitachi Data Systems lab to validate this solution.

Table 1. Hardware Components

Hardware	Description	Version	Quantity
Hitachi Unified Storage VM	▪ Dual controller ▪ 16 × 8 Gb/sec Fibre Channel ports ▪ 64 GB cache memory ▪ 72 SAS 600 GB 10k RPM disks, 2.5 inch SFF	73-01-02-00/00	1
Hitachi Compute Blade 500 chassis	▪ 8-blade chassis ▪ 2 Brocade 5460 FC Switch Modules each with 6 × 8 Gb/sec uplink ports ▪ 2 Brocade VDX 6746 Ethernet Switch Modules each with 8 × 10 Gb/sec uplink ports ▪ 2 Management modules ▪ 6 Cooling fan modules ▪ 4 Power supply modules	SVP: A0108-B-5923 5460: FOS 6.3.2d VDX6746: NOS 2.0.1_kat4	1
Hitachi Compute Blade 520BH1 blades	▪ Half blade ▪ 2 × 8-core Intel Xeon E5-2680 processor, 2.70 GHz ▪ 96 GB RAM ▪ 6 × 16 GB DIMMs	BMC/EFI: 01-27	4
Brocade 6510	▪ SAN switch with 48 × 8 Gb Fibre Channel ports	FOS 7.0.1a	2
Brocade VDX 6720	▪ Ethernet switch with 24 × 10 Gb/sec ports	NOS 2.0.1b	2

Hitachi Unified Storage VM

Hitachi Unified Storage VM is an entry-level enterprise storage platform. It combines storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a storage virtualization system to provide central storage services to existing storage assets.

Unified management delivers end-to-end central storage management of all virtualized internal and external storage on Unified Storage VM. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, as well as virtual NAS functionality, without compromising performance or scalability.

The benefits of Unified Storage VM are the following:

- Enables the move to a new storage platform with less effort and cost when compared to the industry average
- Increases performance and lowers operating cost with automated data placement
- Supports scalable management for growing and complex storage environment while using fewer resources
- Achieves better power efficiency and with more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions
- Consolidates management with end-to-end virtualization to prevent virtual server sprawl

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.

The Hitachi Compute Blade 500 chassis contains internal Fibre Channel and network switches for the high availability requirements of Hitachi Unified Compute Platform for VMware vSphere.

Brocade Storage Area Network Switches

Brocade and Hitachi Data Systems have collaborated 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 reference architecture uses the following Brocade products:

- Brocade 6510 Switch
- Brocade VDX 6720 Data Center Switch

Software Components

This describes the software components deployed for this reference architecture.

Table 2 describes the software used in this reference architecture.

Table 2. Software Components

Software	Version
Hitachi Storage Navigator Modular 2	Microcode Dependent
Hitachi Dynamic Provisioning	Microcode Dependent
VMware vCenter server	5.1.0
VMware Virtual Infrastructure Client	5.1.0
VMware ESXi	5.1.0
Microsoft Windows Server 2008	Enterprise edition, R2
Microsoft SQL Server 2008	Enterprise edition, R2

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 VM, 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 re-balance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

VMware vSphere 5

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** — 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 5** — 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.

Solution Design

This is detailed information on this Hitachi Unified Compute Platform Select for VMware vSphere reference solution. It includes software and hardware design information required to build the basic infrastructure for the virtualized data center environment.

To provide you with options for scaling out your environment in modular increments, this solution uses a cell architecture. This design defines compute and storage resource groups to support a specific usage scenario. You can add additional cells to scale out the environment to meet your organization's requirements.

Figure 2 illustrates a high-level concept of the cell architecture.

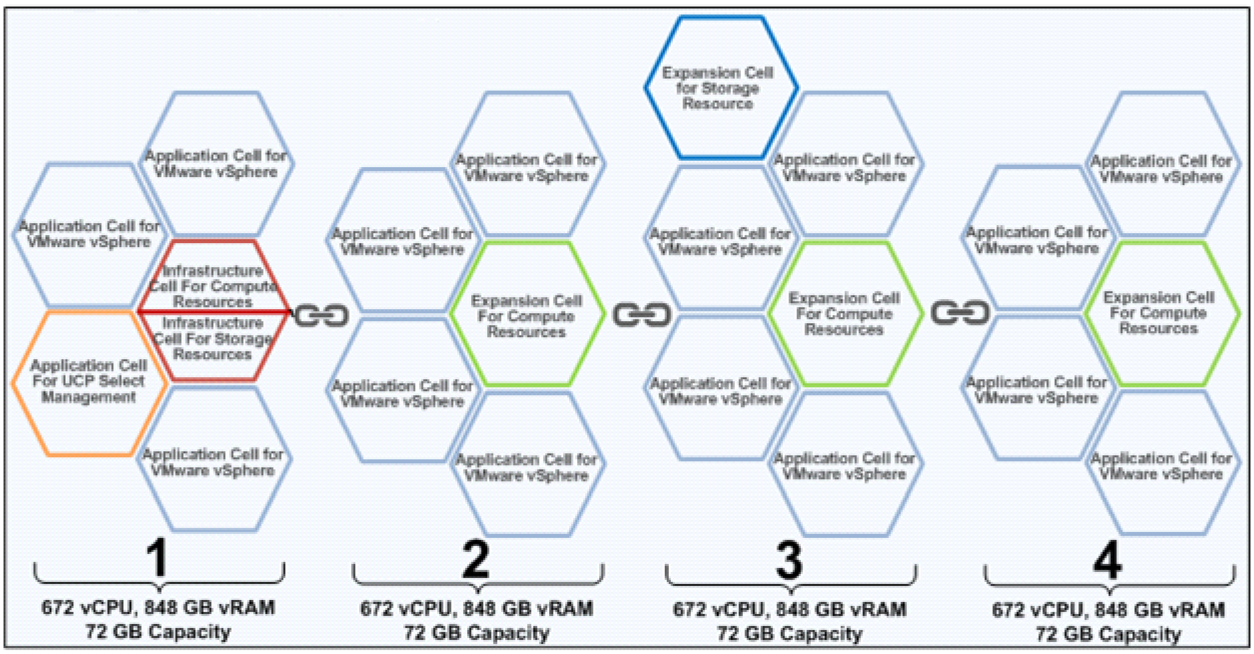

Figure 2

The architecture consists of preconfigured cells designed to support general server workload. These cells provide the following:

- **Infrastructure cell for compute resources** — Foundation for compute components
- **Infrastructure cell for storage resources** — Foundation for storage components
- **Application cell for Hitachi Unified Compute Platform Select management** — Resource to manage this environment
  - This cell is required only if an existing configuration for managing a VMware vSphere environment does not exist.
- **Application cell for VMware vSphere** — Provides the resource for hosting virtual machines running general server application workloads.
- **Expansion cell for compute resources** — Provides the compute resources for scaling out the Unified Compute Platform Select for VMware vSphere environment.
- **Expansion cell for storage resources** — Provides the storage resources for scaling out the Unified Compute Platform Select for VMware vSphere environment.

These cells provide the compute and storage hardware needed to build this scalable Hitachi Unified Compute Platform Select for VMware vSphere solution.

Infrastructure Cell for Compute Resources

The infrastructure cell for compute resources provides the foundation for the compute components needed to start building this solution.

Figure 3 shows the infrastructure cell for compute resources.

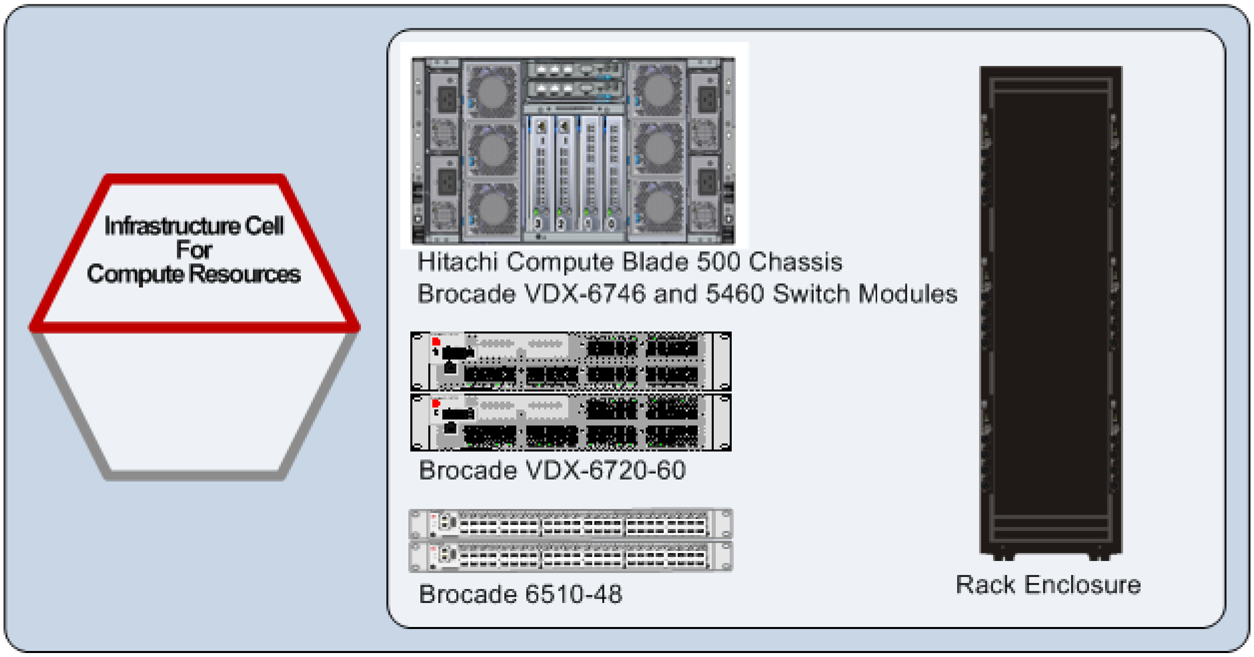

Figure 3

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

- Infrastructure cell for storage resources
- Application cell for Hitachi Unified Compute Platform Select management
- Application cell for VMware vSphere
- 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 solution. Both infrastructure cells 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.

Table 3 shows the components of the infrastructure cell for compute.

Table 3. Hardware Components for the Infrastructure Cell for Compute Resources

Hardware	Detail Description	Version	Quantity
Hitachi Compute Blade 500 Chassis	■ 2 Brocade VDX6746 DCB switch modules ■ 2 Brocade 5460 6-port 8 Gb/sec Fibre Channel switch modules ■ 2 Chassis management modules ■ 6 Cooling fan modules ■ 4 Power supply modules	SVP: A0108-B-5923 5460: FOS 6.3.2d VDX6746: NOS 2.0.1_kat4	1
Brocade Ethernet Switch	■ Brocade VDX6720-60 10 Gb/sec 60 port Ethernet Switch	NOS 2.0.1b	2
Brocade Fibre Channel Switch	■ Brocade 6510-48 8 Gb/sec 48 port Fibre Channel Switch	FOS 7.0.1a	2

The hardware in the infrastructure cell for compute resources makes up the core compute hardware in this Hitachi Unified Compute Platform Select for VMware vSphere solution.

Chassis Components

The Hitachi Compute Blade 500 chassis has 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:

- A fully populated infrastructure cell for compute resources,
- An infrastructure cell for storage resources,
- Up to three expansion cells for compute resources

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

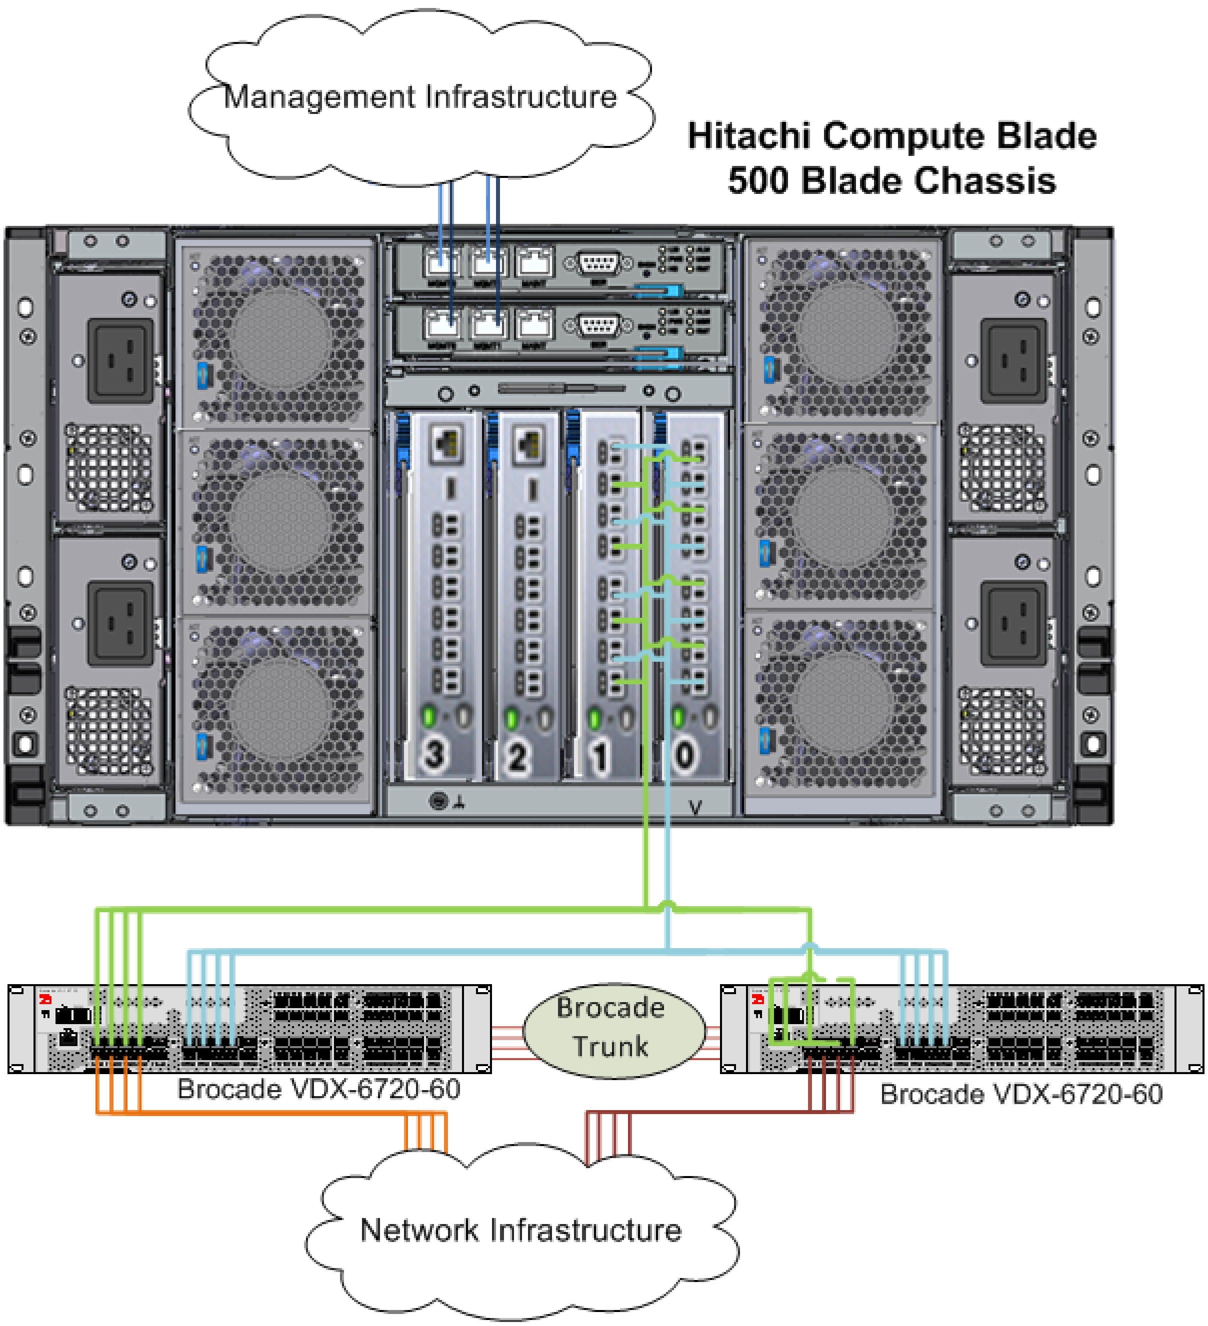

Figure 4

The network design also allows for the utilization of advanced features inherent in the Brocade VDX switch family such as Brocade's VCS Fabric Technology. This helps provide:

- Non-stop networking
- Simplified, automated networks
- An evolutionary approach that protects existing IT investments

See the Brocade website for more information about Brocade VCS Fabric Technology.

SAN Infrastructure

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

For this reference architecture, zone the infrastructure cell for compute resources to four ports on the Hitachi Unified Storage VM controller, two ports per cluster. When adding expansion cells for compute resources to the solution, zone four new open storage ports on the cluster.

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

Figure 5 on page 15 illustrates the physical SAN architecture of the infrastructure cell for compute.

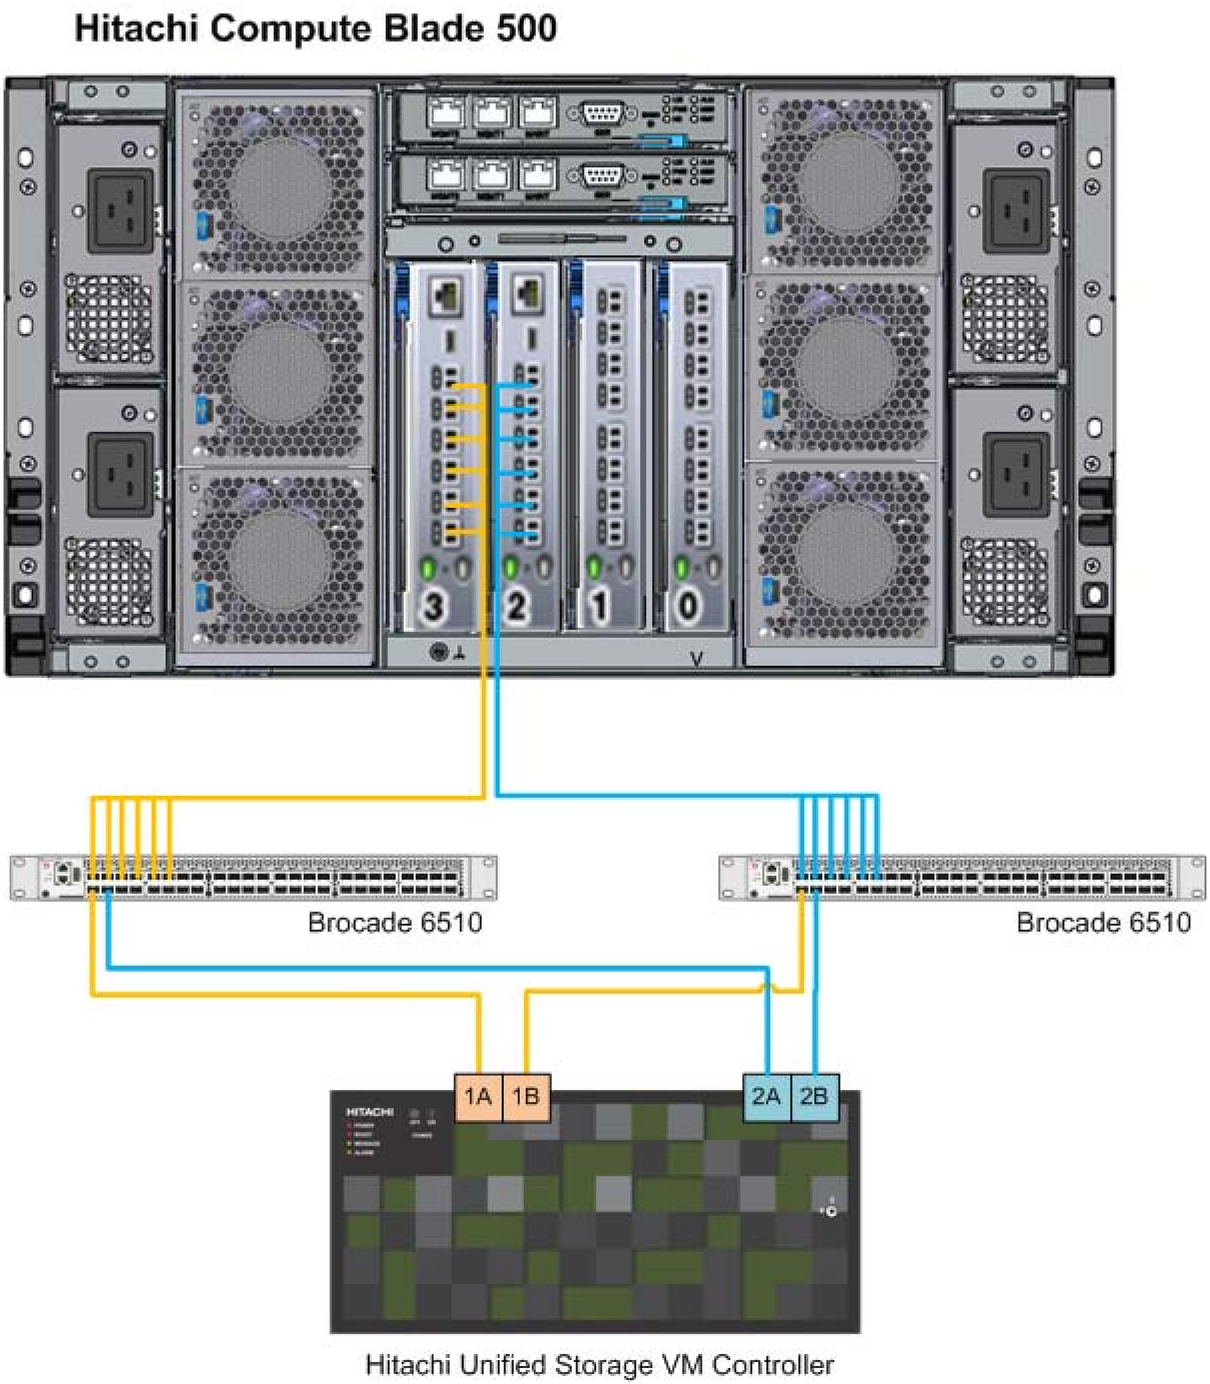
Hitachi Compute Blade 500
3
2
1
0
Brocade 6510
Brocade 6510
1A
1B
2A
2B
Hitachi Unified Storage VM Controller

Figure 5

Infrastructure Cell for Storage Resources

The infrastructure cell for storage resources contains all of the base storage hardware required to start building this solution.

Figure 6 shows the infrastructure cell for storage resources.

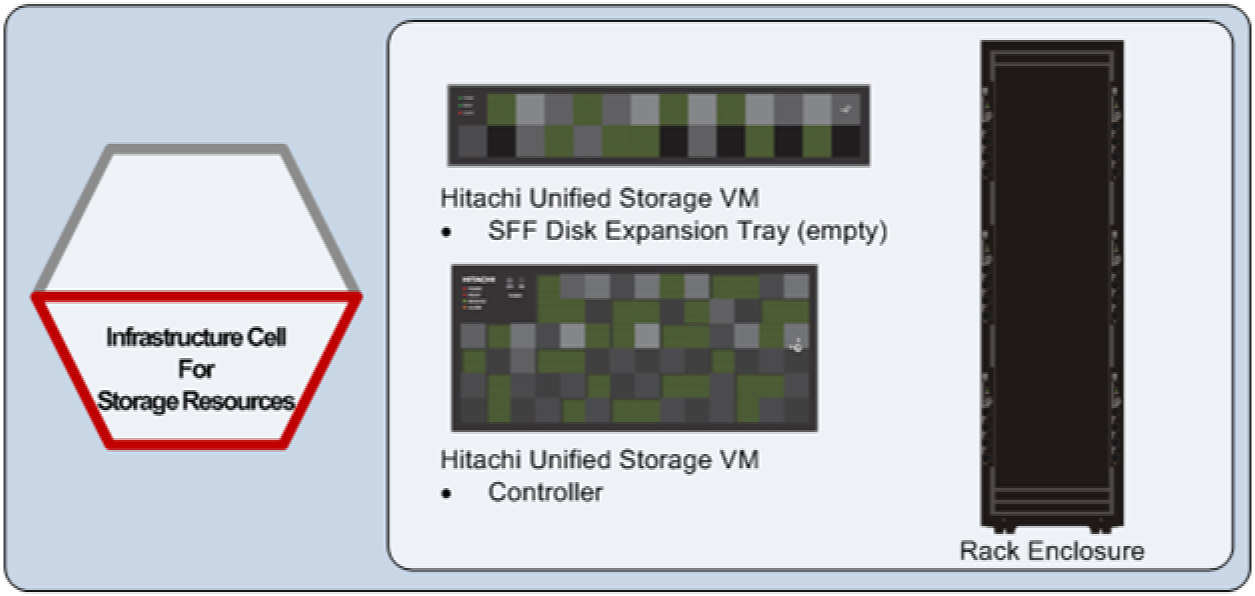

Figure 6

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 VMware vSphere

The infrastructure cell for storage resources provides the storage infrastructure for the other cells in the solution. Once an infrastructure cell for storage resources is fully populated, add additional infrastructure cells for storage resources to scale out the solution.

Table 4 shows the components of the infrastructure cell for storage.

Table 4. Infrastructure Cell for Storage Resources Hardware

Hardware	Detail Description	Version	Quantity
Hitachi Unified Storage VM	▪ Dual controllers and Fibre Channel Modules ▪ 16 × 8Gb/sec Fibre Channel ports ▪ 32 GB cache	73-01-02-00/00	1
SFF disk expansion tray for Hitachi Unified Storage VM	▪ Contains disks for other cells		1

The infrastructure cell for storage resources contains a Hitachi Unified Storage VM controller and a disk expansion tray.

This disk expansion tray holds disks for this infrastructure cell. Add storage disks to this cell for the following:

- Application cell for Hitachi Unified Compute Platform Select management
- Hot spares (optional)

Each infrastructure cell for storage resources can physically support up to 11 application cells for VMware vSphere.

Note — Scalability limits depend on application workloads running on this infrastructure.

Application Cell for Hitachi Unified Compute Platform Select Management

The application cell for Hitachi Unified Compute Platform Select management contains the compute and storage components for hosting the VMware vSphere infrastructure services.

Figure 7 shows the application cell for Unified Compute Platform Select management.

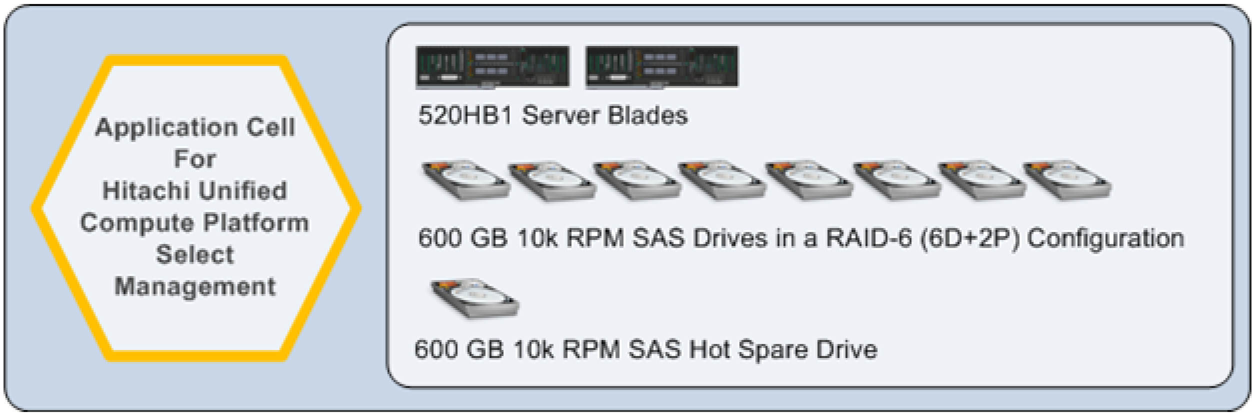

Figure 7

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

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources
- Application cell for VMware vSphere

Use an application cell for Hitachi Unified Compute Platform Select management when a VMware vSphere environment does not already exist.

Note — Scalability limits depend on application workloads running on this infrastructure.

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 Unified Compute Platform Select management to ensure virtual machine failover in the event of a hardware failure.

Table 5 shows the details of the hardware configuration in the application cell for Unified Compute Platform Select management.

Table 5. Application Cell for Hitachi Compute Platform Select Management Hardware

Hardware	Detail Description	Version	Quantity
520HB1 server blade	▪ 2 × 8-Core Intel Xeon E5-2680 processor, 2.7 GHz ▪ 128 GB Memory per blade ▪ 1 Emulex 2-port 10 GbE on-board CNA card ▪ 1 Hitachi FIVE-EX 2-port 8 Gb Fibre Channel mezzanine card	BMC/EFI: 01-27	2
SFF Disk Drives ▪ 600 GB 10k RPM SAS drives ▪ Installed in the infrastructure cell for storage resources disk tray	▪ RAID-6 (6D+2P)		8
	▪ Hot spare		1

The compute infrastructure of the application cell for Unified Compute Platform Select management supports all associated Microsoft SQL Server, Active Directory, and VMware vCenter requirements.

Manage your environment using the above resources or by connecting to a pre-existing VMware vSphere management environment.

Network Infrastructure

Configure each of the 520HB1 server blades with a single on-board two-channel 10 GbE CNA card for network traffic.

Split each CNA card into four logical NICs per channel, for eight NICs per server blade. This design only uses three NICs per channel. This allows maximum bandwidth for the virtual machine network.

Set bandwidth allocation for each NIC 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**

  Virtual machine 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 Hitachi Compute Blade 500.

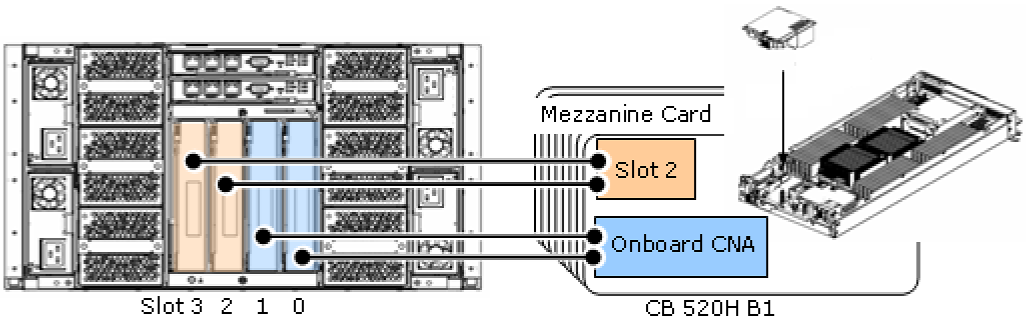

Figure 8

This solution uses the following VLANs to separate network traffic in the application cell for VMware vSphere:

- **Management-VLAN** — Chassis management connections and primary management of the ESXi hypervisors
- **vMotion-VLAN** — Configured for vMotion
- **VM-VLAN** — Configured for the virtual machine network

Following best practice, separate the management, vMotion, and virtual machine traffic to achieve greater security or better performance.

- Team the logical NICs to allow network path redundancy
- Perform maintenance upgrades with zero downtime of the Brocade VDX6746 switch modules while you keep the server blades online.

With enhancements to VMware vSphere 5, the VMkernel load balances vMotion traffic over all vmkernel ports configured for vMotion. This improves performance and reduces migration times.

Storage Infrastructure

The storage infrastructure of the application cell for Hitachi Unified Compute Platform Select management consists of eight 600 GB 10k RPM SAS drives housed in the disk expansion tray contained in the infrastructure cell for storage. Configure the storage into a single RAID-6 (6D+2P) dynamic provisioning pool dedicated to management servers. The pool provides an overall capacity of 3 TB.

Zone each server blade in the application cell for Unified Compute Platform Select management to Hitachi Unified Storage VM through the Brocade 5460 Fiber Chanel switch modules using single initiator to multi target zoning for each port on the 520HB1 server blades. Following best practice, the SAN environment was configured 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

The storage multipathing policy for each target in ESXi was set to **round robin**. This results in optimal load distribution during an all paths available situation.

Table 6 shows the zoning configuration used for the application cell for Unified Compute Platform Select management.

Table 6. Application Cell for Unified Compute Platform Select Management Zone Configuration

Host	Host HBA Number	Fabric	Zone Name	Storage Port
Blade0-ESX 0	HBA1_1	Fabric 1	ESX0_HBA1_1_HUS_VM_1A_2A	1A
				2A
	HBA1_2	Fabric 2	ESX0_HBA1_2_ HUS_VM_1B_2B	1B
				2B
Blade1-ESX 1	HBA1_1	Fabric 1	ESX 1_ HBA1_1_ HUS_VM_0A_0B	1A
				2A
	HBA1_2	Fabric 2	ESX1_ HBA1_2_ HUS_VM_1A_1B	1B
				2B

Server Configuration Sizing Guidelines

Apply the proper resource allocation for virtual machines used to manage the Hitachi Unified Compute Platform Select for VMware vSphere environment. If using a separate environment outside of this solution for management, use the virtual machine sizing recommendations in Table 7.

Table 7 lists the virtual machine configurations used for each component of the management infrastructure used in this reference architecture.

Table 7. Virtual Machine Sizing Recommendations

Virtual Machine	*Configuration*	*Count*
Microsoft Active Directory, DNS, DHCP	vCPU — 1 vMemory — 4 GB	1
VMware vCenter	vCPU — 2 vMemory — 10 GB	1
Microsoft SQL 2008 database for VMware vCenter	vCPU — 2 vMemory — 8 GB	1

Application Cell for VMware vSphere

The application cell for VMware vSphere contains all compute and storage components necessary to run general server application workloads consisting of the following:

- 168 virtual CPUs
- 212 GB of virtual machine memory
- 18 TB of storage capacity

Figure 9 on page 23 shows the application cell for VMware vSphere.

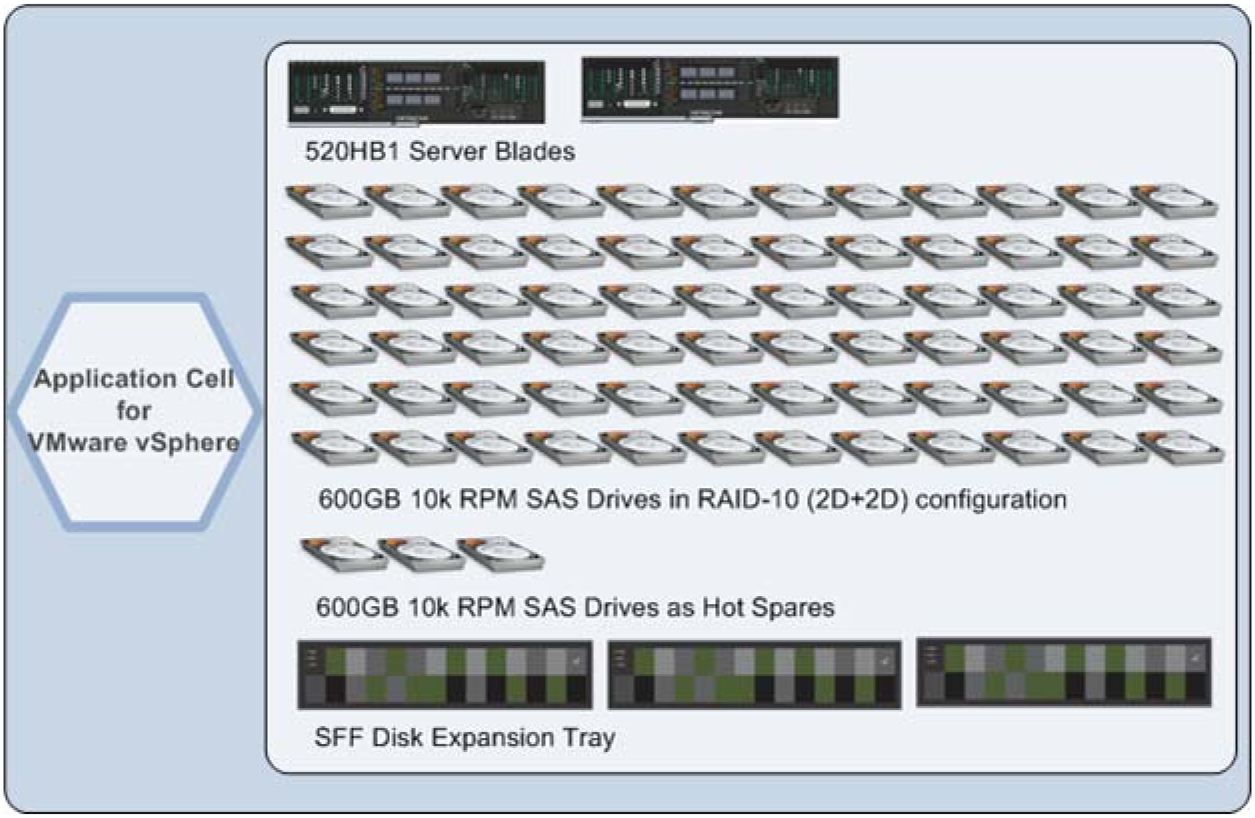

Figure 9

Use the application cell for VMware vSphere 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)

Add the compute components of the application cell for VMware vSphere to the infrastructure cell for compute and the storage components to the infrastructure cell for storage to start building a scalable Hitachi Unified Compute Platform Select for VMware vSphere environment.

To scale out the solution and increase capacity, add additional application cells for VMware vSphere to your infrastructure cells for compute resources or expansion cells for compute resources. A single infrastructure cell for compute resources and an infrastructure cell for storage resources physically supports up to 16 application cells for VMware vSphere before you require a new infrastructure cells.

Note — Scalability limits depend on application workloads running on this infrastructure.

Compute Infrastructure

The application cell for VMware vSphere supports a maximum density of 168 virtual CPUs and 212 GB of virtual machine memory.

In a maximum density configuration, a cell cannot support the failover of virtual machines in case of a server blade failure. To provide high availability, do the following:

- Reduce the number of virtual CPUs and virtual machine memory per host up to 50%.
- Configure a VMware High Availability and Distributed Resource Scheduler cluster dedicated to application cells for VMware vSphere.

Place additional hosts from each application cell for VMware vSphere into the cluster. When scaling the solution, increase the number of virtual machines per host as you add more resources to the cluster.

Based on VMware maximums, each High Availability and Distributed Resource Scheduler cluster can support up to 16 application cells for VMware vSphere (32 hosts).

Table 8 shows the details of the hardware used in the application cell for VMware vSphere.

Table 8. Application Cell for VMware vSphere Hardware

Hardware	Detail Description	Version	Quantity
520HB1 server blade	■ 2 × 8-Core Intel Xeon E5-2680 processors, 2.7 GHz ■ 96 GB RAM per server blade ■ 1 Emulex 2-port 10 GbE on-board CNA card ■ 1 Hitachi FIVE-EX 2-port 8 Gb Fibre Channel mezzanine card	BMC/EFI: 01-27	2
SFF disk drives ■ 600GB 10k RPM SAS drives	■ RAID-10 (2D+2D)		72
	■ Hot spare ■ Installed in infrastructure cell for storage resources disk tray		3
SFF disk expansion tray	■ Added to the infrastructure cell for storage resources		1

Network Infrastructure

The application cell for VMware vSphere uses the same networking configuration described in the application cell for Hitachi Unified Compute Platform Select management.

Storage Infrastructure

The storage infrastructure of the application cell for VMware vSphere consists of seventy-two 600 GB 10k RPM SAS drives in two dynamic provisioning pools with the following configuration:

- **Pool 0** — 24 drives (1 tray) consisting of 6 RAID-10 (2D+2D) parity groups.
- **Pool 1** — 48 drives 48 drives (2 trays) consisting of 12 RAID-10 (2D+2D) parity groups.

Figure 10 on page 26 shows the storage configuration for the application cell for VMware vSphere.

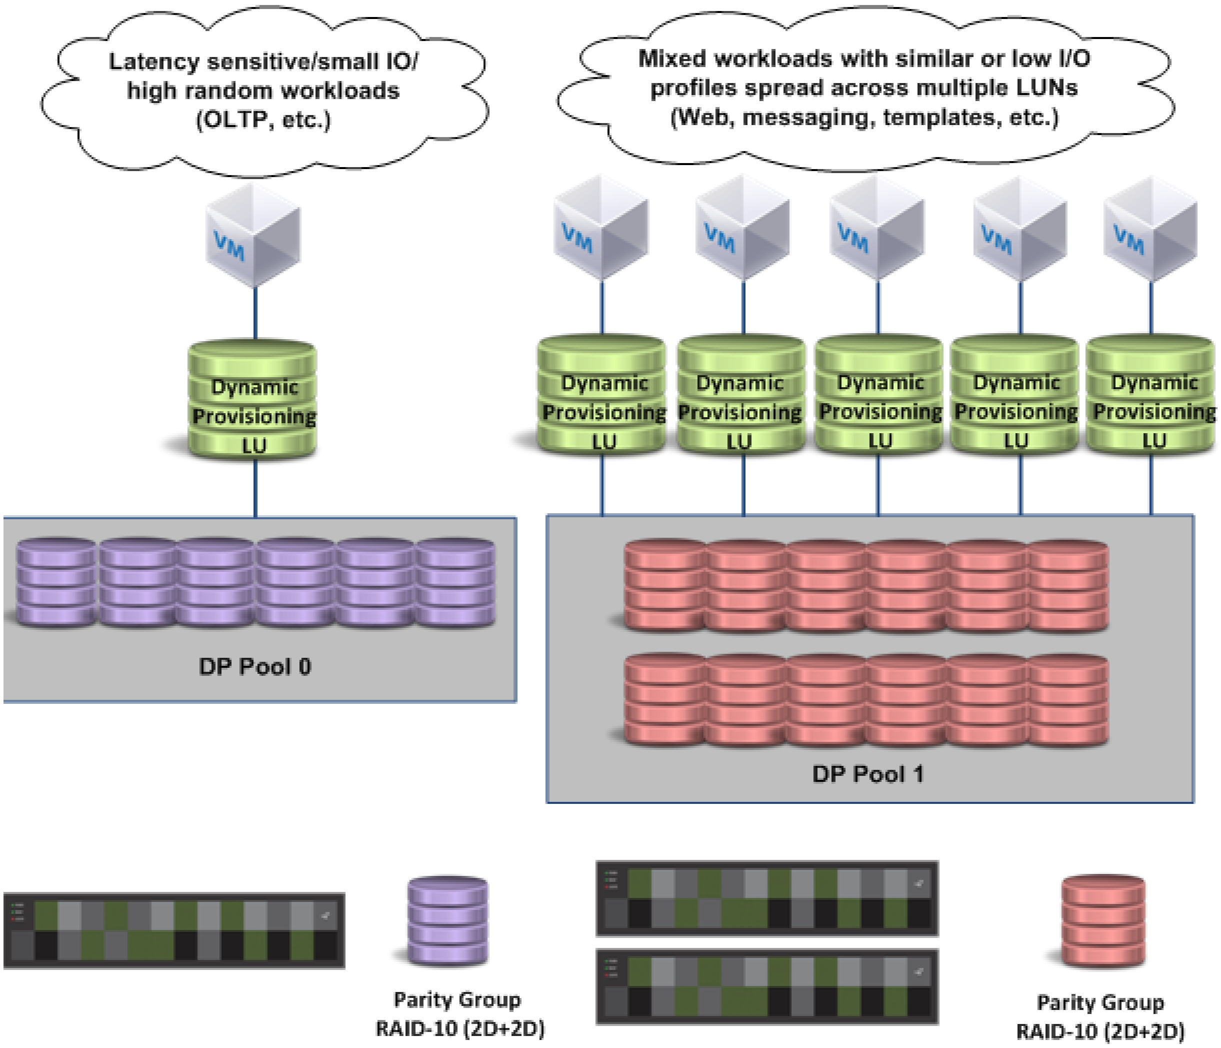

Figure 10

Use RAID-10 to maximize performance for random workloads, which is common with virtualized environments.

Create two pools to separate virtual machine workloads with different performance characteristics.

Because of its wide striping capability, Hitachi Dynamic Provisioning can balance the I/O load in pools of RAID groups. Mixing workloads in a single dynamic provisioning pool is possible to obtain certain levels of performance. However, grouping virtual machines with similar I/O profiles optimizes storage performance and results in a more efficient use of disk resources. Within a pool, create additional LUN's as necessary to spread the workload and avoid possible queue depth issues.

When scaling out with additional application cells for VMware vSphere, add RAID groups to grow the existing pools. Increasing spindle count allows the pool to support the increasing IOPS requirement dynamically. As stated before, create additional LUNs to prevent virtual machine workloads from saturating the LUN.

SAN Infrastructure

The 520HB1 server blades used in this reference architecture use dual-port 8 Gb/sec Fibre Channel mezzanine cards with redundant connections to the Brocade 6510 enterprise fabric switches.

The environment uses single initiator to multi-target zoning for each 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

The storage multipathing policy for each target in ESXi was set to **round robin**. This results in optimal load distribution during an all paths available situation.

Table 9 shows the zone configuration used for the application cell for VMware vSphere.

Table 9. Application Cell for VMware vSphere Zone Configuration

Host	*Host HBA Number*	*Fabric*	*Zone Name*	*Storage Port**
Blade2-ESX 2	HBA1_1	Fabric 1	ESX2_HBA1_1_HUS_VM_1A_2A	1A
				2A
	HBA1_2	Fabric 2	ESX2_ HBA1_2_ HUS_VM _1B_2B	1B
				2B
Blade3-ESX 3	HBA1_1	Fabric 1	ESX3_ HBA1_1_ HUS_VM_1A_2A	1A
				2A
	HBA1_2	Fabric 2	ESX3_ HBA1_2_ HUS_VM _1B_2B	1B
				2B

*The storage target ports for each cell depend on the Hitachi Compute Blade 500 chassis in which they are hosted. See “Scaling Using Expansion Cell for Compute Resources” on page 28 for details.

Scaling Using Expansion Cell for Compute Resources

Use an expansion cell for compute resources to scale out this solution beyond the first infrastructure cell for compute resources.

Figure 11 shows the expansion cell for compute resources.

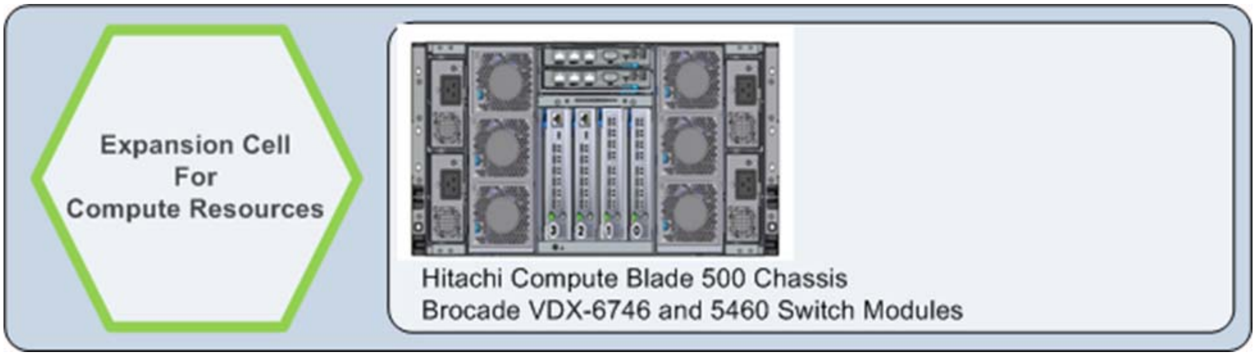

Figure 11

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

- Infrastructure cell for compute resources
- Application cell for VMware vSphere

Once the chassis in the infrastructure cell for compute resources becomes fully populated, use an expansion cell for compute resources to provide additional resource capacity. This 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 rack enclosure of the infrastructure cell for compute resources.

You can physically add up to three expansion cells for compute resources to an infrastructure cell for compute resources and an infrastructure cell for storage resources before you need to add new infrastructure.

One infrastructure cell for compute resources and two expansion cells for compute resources support a maximum of 11 application cells for VMware vSphere (22 server blades and 33 storage trays).

Note — Scalability limits depend on application workloads running on this infrastructure.

Chassis Components

The expansion cell for compute resources uses the same chassis components contained in the infrastructure cell for compute resources.

Networking Infrastructure

The networking for the expansion cell for compute resources uses the same networking configurations as the infrastructure cell for compute resources.

Storage Infrastructure

Use four of the open storage target ports on Hitachi Unified Storage VM in the infrastructure cell for storage resources. Follow the same storage configuration described for the infrastructure cell for compute resources to use the newly provisioned storage target ports in the zoning configuration.

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

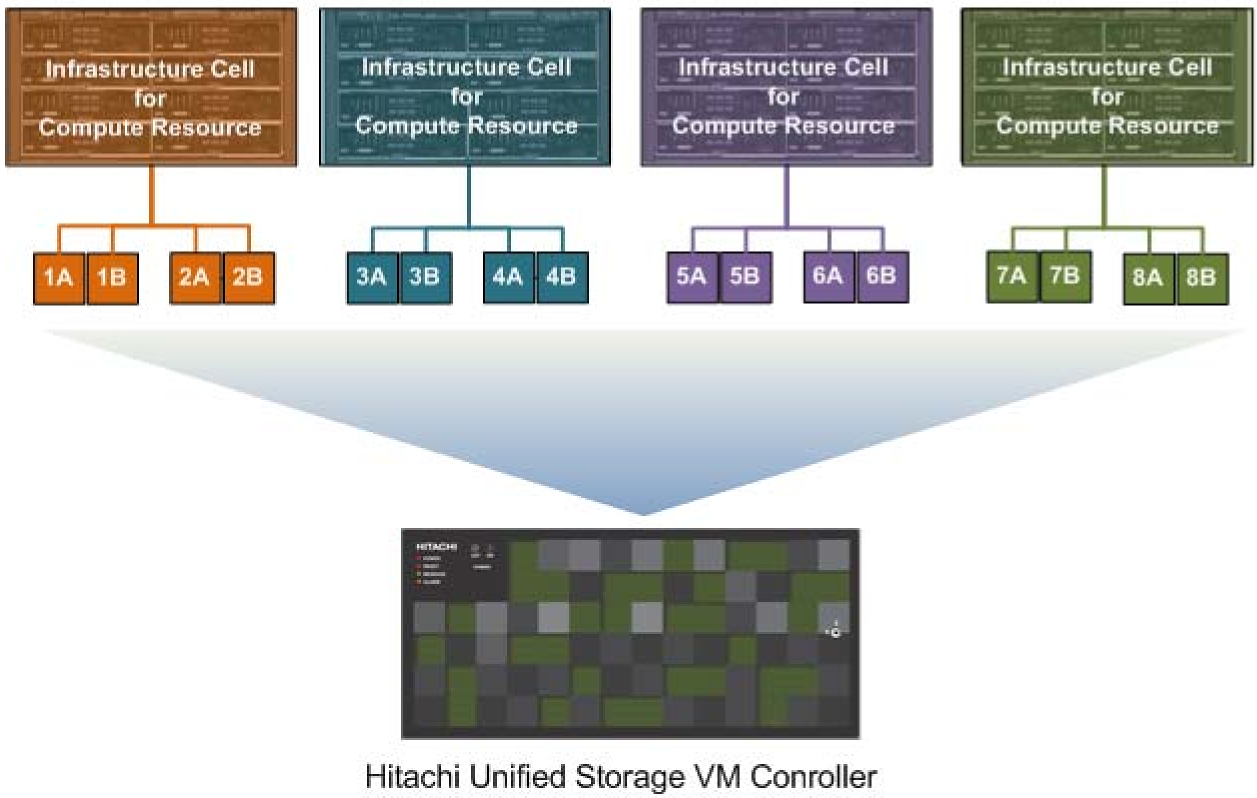

Figure 12

Scaling Using Expansion Cell for Storage Resources

Use an expansion cell for storage resources to scale out the VMware vSphere solution beyond the first infrastructure cell for storage resources.

The expansion cell for storage contains only a 2.5 inch SFF disk tray for Hitachi Unified Storage VM.

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

- Infrastructure cell for storage resources
- Application cell for VMware vSphere

Once the original infrastructure cell for storage drive chassis becomes fully populated, use an expansion cell for storage resources to provide additional capacity.

Put hot spares for the first application cells in the disk tray for the infrastructure cell for storage resources. When the tray in the infrastructure cell fills, use the expansion cell for storage resources.

Engineering Validation

This describes the test methodology used to validate this reference architecture and the results of the validation testing. These tests demonstrated the maximum utilization of the reference architecture.

The purpose of the tests was to determine maximum loads that the solution could support and still maintain an acceptable application performance.

Test Methodology

This reference architecture tested the core components of this Hitachi Unified Compute Platform Select for VMware vSphere solution to validate its performance and design. For validation purposes, testing a mixed workload of the following:

- Email messages
- Web pages
- Online transaction processing (OLTP)

The workload grouping was into a tile-based system to measure application performance and scalability. Each tile contained mixed workloads that stress critical compute and storage resources. These workloads represent a general purpose environment for VMware vSphere.

Each tile consists of the following virtual machines listed in Table 10.

Table 10. Virtual Machines for Each Testing Tile

	Microsoft Exchange 2007	*Olio Web Server*	*Olio Database Server*	*DVD Store 2 Database Server*	*DVD Store 2 Web Server*	*Standby*
Quantity	1	1	1	1	3	1
CPU	4 vCPUs	4 vCPUs	2 vCPUs	4 vCPUs	2 vCPUs	1 vCPUs
Memory	8192 MB	6144 MB	2048 MB	4096 MB	2048 MB	512 MB

Each tile represented a simulation of the following types of workloads:

- Microsoft Exchange 2007 mail servers for general email workloads
- Olio web and database servers for Web 2.0 workloads
- DVD Store 2 web and database servers for OLTP workloads
- Standby servers for idle general infrastructure workload

Testing involved these cells:

- Infrastructure cell for compute resources
- Infrastructure cell for storage resources
- Application cell for Unified Compute Platform Select management
- Application cell for VMware vSphere
    - 168 vCPUs
    - 212 GB vRAM
    - 18 TB capacity

Figure 13 shows those cells used to validate this reference architecture.

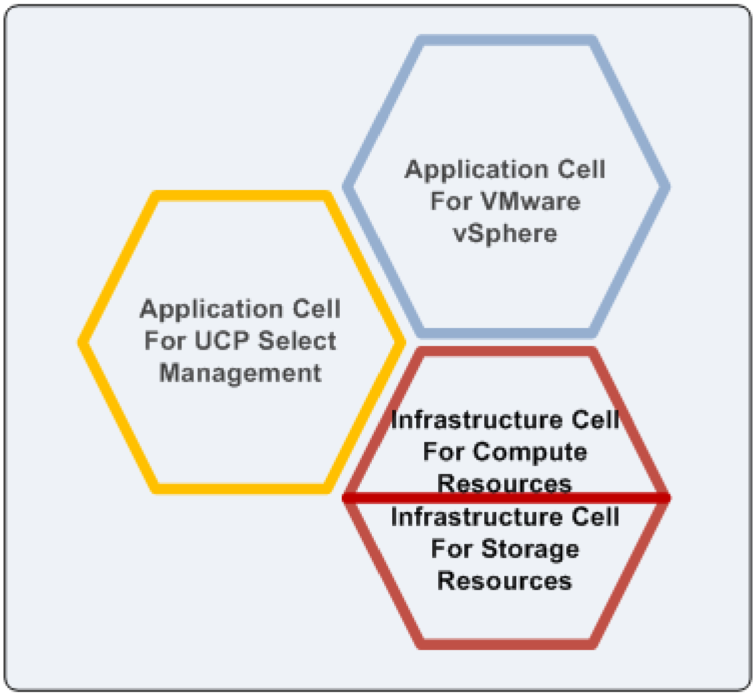

Figure 13

Testing used eight tiles between two ESXi hosts in the application cell for VMware vSphere. There were a total of the following:

- 64 virtual machines
- 168 virtual CPUs
- 212 GB of configured virtual machine memory

A single client controls each tile. A primary client controls each tile client. The clients ran in other hosts, outside of the ESXi hosts for the workload virtual machines.

Table 11 shows how the tiles were distributed.

Table 11. Tile distribution Across Compute and Storage Resources

	Tile 1	Tile 2	Tile 3	Tile 4	Tile 5	Tile 6	Tile 7	Tile 8
ESXi Host	2	3	2	3	2	3	2	3
HDP Pool 0	LUN 1: DVD Store 2 Database Servers							
HDP Pool 1	LUN 2: DVD Store 2 Web Servers, Olio Web Servers, Olio Database Servers							
	LUN 3: Mail Servers 1-3			LUN 4: Mail Servers 4-6			LUN 5: Mail Servers 7-8	
	LUN 6: Standby Servers							

Due to their higher random read and I/O intensive workload characterization, the DVD Store 2 database servers were placed on Dynamic Provisioning Pool 0.

All other servers had similar higher random write workload characterization and therefore placed on dynamic provisioning Pool 1.

Table 12 shows the tile workload definitions.

Table 12. Tile Workload Definitions

Workloads	Applications	Virtual Machine Platform	Simulated Load per Tile
Mail Server	▪ Microsoft Exchange 2007 ▪ Microsoft Exchange LoadGen	▪ Microsoft Windows 2008 R2 (64 Bit) ▪ 4 vCPU ▪ 8 GB RAM ▪ >100 GB boot and data disks	1000 users with a heavy workload profile
Standby	▪ None	▪ Microsoft Windows 2003 (32 Bit) ▪ 1 vCPU ▪ 512 MB RAM ▪ 4 GB boot disk	Non-load based functional test to activate idle resources for on-demand usage.

Table 12. Tile Workload Definitions (Continued)

Workloads	Applications	Virtual Machine Platform	Simulated Load per Tile

Web 2.0 load simulation	▪ Olio DB
	▪ Web application servers
	▪ 2-tier Java-based implementation of the Olio workload, including the following operations:
	  ▪ HomePage
	  ▪ Login
	  ▪ TagSearch
	  ▪ EventDetail
	  ▪ PersonDetail
	  ▪ AddPerson
	  ▪ AddEvent
		Database:
		▪ SUSE Linux 11 64-bit
		▪ 2 vCPU
		▪ 2 GB RAM
		▪ 10 GB boot and 4 GB data disks
		Web-Server:
		▪ SLES 11 64-bit
		▪ 4 vCPU
		▪ 6 GB RAM
		▪ 10 GB boot and 70 GB data disks
			400 concurrent users

E-commerce simulation	DVD Store 2	Database: ▪ SUSE Linux 11 (64-bit) ▪ 4 vCPU ▪ 4 GB RAM ▪ 10 GB boot and 35 GB data disks Front end (×3): ▪ SLES 11 64-bit ▪ 2 vCPU ▪ 2 GB RAM ▪ 10 GB disk	10 constant driver thread loads from one web server 20 burst based driver threads from two web servers Performance Metric is transactions per minute (TPM)

Table 12. Tile Workload Definitions (Continued)

Workloads	*Applications*	*Virtual Machine Platform*	*Simulated Load per Tile*
Virtual machine cloning and deployment	▪ vSphere Cloning ▪ vSphere virtual machine customization ▪ Updates VMTools operation ▪ Destroy virtual machine, wait 5 minutes, and restarts	▪ Uses Standby Server as base (see above)	Concurrent deployments increase per number of tiles. Metric is measured in Deployments Per Hour (DPH)
Dynamic virtual machine relocation	▪ VMware vMotion ▪ Relocation and then wait 3 minutes to repeat	▪ Olio database selected at random in round robin fashion (see specifications above)	Concurrent relocations Increase per number of tiles Metric is measured in Relocations Per Hour (RPH)
Dynamic storage relocation (Storage vMotion)	▪ Automated storage relocation to temporary location ▪ Automated storage relocation back, and then wait 5 minutes to repeat.	▪ Standby virtual machine chosen at random (see specifications above)	Concurrent relocations increase per number of tiles Metric is measured in relocations per hour (RPH)
Automated load balancing (vMotion)	▪ VMware Distributed Resource Scheduler ▪ Set to aggressive	▪ All	Infrastructure functionality to load balance tile workload.

Compute Infrastructure

Multiple performance metrics were collected from the ESXi hypervisor during the test.

With hyper-threading enabled on the 2 × 8-Core Intel Xeon E5-2680 processors, 32 physical CPUs are available for each host. There were 32 virtual machines configured with 84 virtual CPUs ran on each blade.

Each 520HB1 server blade contained 96 GB of RAM. The 32 virtual machines were configured with a total of 106 GB of vRAM on each blade. Over commitment of memory allocation was 10 GB.

There was the collection of guest operating system metrics for each type of server during the workload run.

Storage Infrastructure

There was the collection of multiple performance metrics from Hitachi Unified Storage VM during the test. The analysis of the metrics were analyzed from the Hitachi Unified Storage VM controllers was to verify the following:

- Physical disks were not saturated
- Storage processor and cache were not overwhelmed
- Hitachi Dynamic Provisioning pools performed well

Table 13 shows the MPU to LDEV mapping.

Table 13. MPU to LDEV Mapping

MPU	LDEV
MPU-10	▪ Mail3
MPU-11	▪ Standby
MPU-20	▪ Deploy ▪ Mail2 ▪ Storage
MPU-21	▪ Mail2 ▪ vMotion

Application Performance

This analyzes the application performance of this reference architecture. The capture of the performance of each application was with the configuration described in "Storage Infrastructure."

Test Results

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

Compute Infrastructure

Figure 14 on page 37 and Figure 15 on page 38 illustrate the physical CPU metrics collected on both ESXi hypervisors while running the 8-tile general server application workload.

- The CPU usage between each host is similar, but slightly off balance due to DRS balancing the workloads during the heavy start up. The difference in load during testing was not enough for DRS to warrant redistributing the load a second time.
- With the core utilization averaging 90% during steady state, the blades are running with high CPU utilization but have some buffer room to handle random CPU spikes.
- The test was used to determine maximum loads per host. For high availability, reduce the number of servers or workload by up to fifty percent and place them in a cluster with VMware High Availability and Dynamic Resource Scheduler enabled.

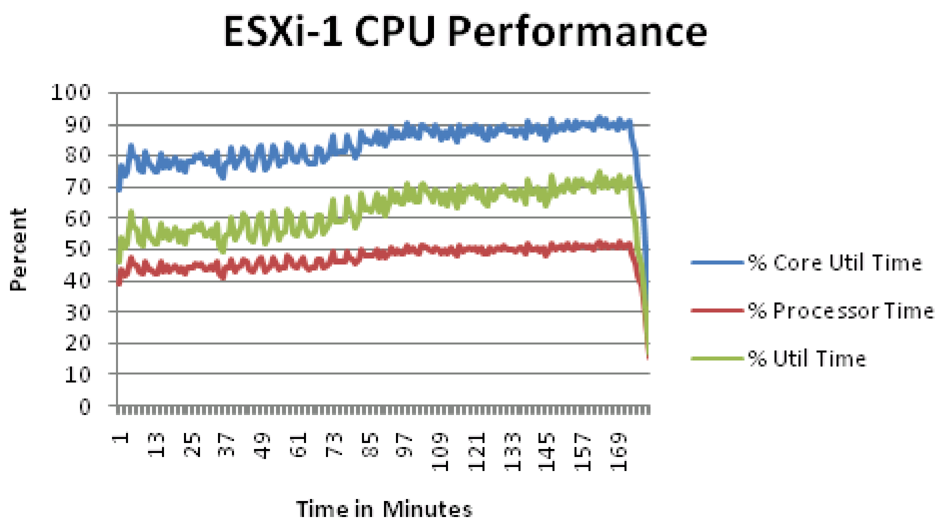

Figure 14

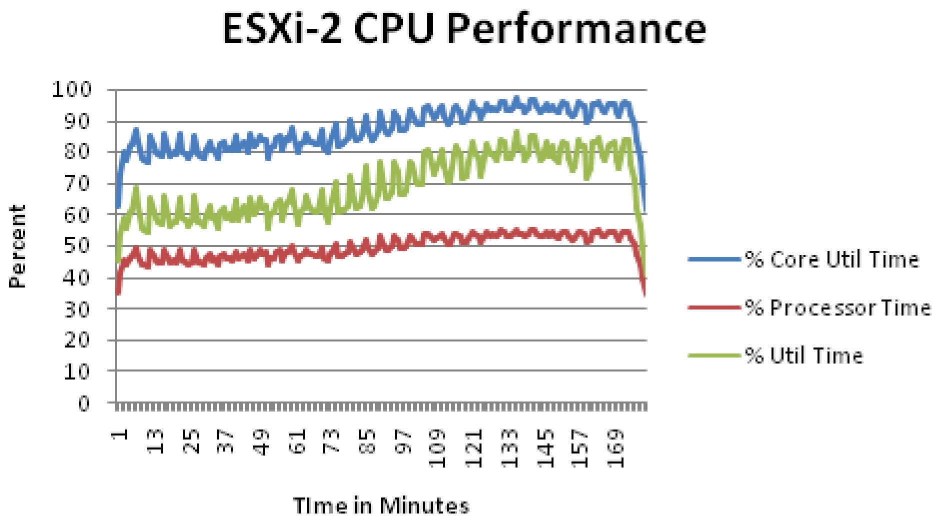

Figure 15

Figure 16 and Figure 17 on page 39 illustrate the physical memory usage collected on both ESXi hypervisors while running the 8-tile general server application workload.

- The memory usage between each host is similar, but slightly off balance due to DRS balancing the workloads during the heavy start up. The difference in load during testing was not enough for DRS to warrant redistributing the load a second time.
- The difference in Used Physical RAM (with Shared VM RAM) and Used Physical RAM (without Shared VM RAM) indicates the amount of memory saved from transparent page sharing. Due to the variation in application workloads, benefit from transparent page sharing was minimal.
- The overall used memory throughout the test was relatively low. The 520H B1 blade has adequate memory headroom for more memory intensive workloads.

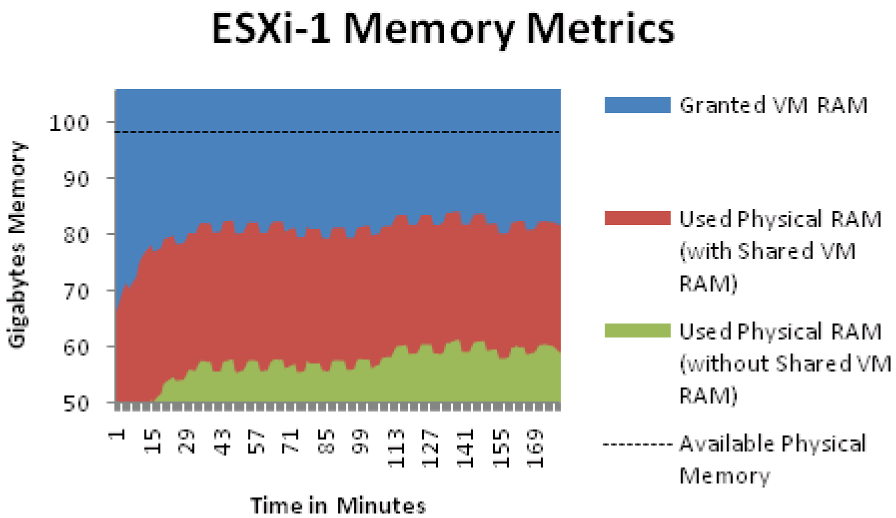

Figure 16

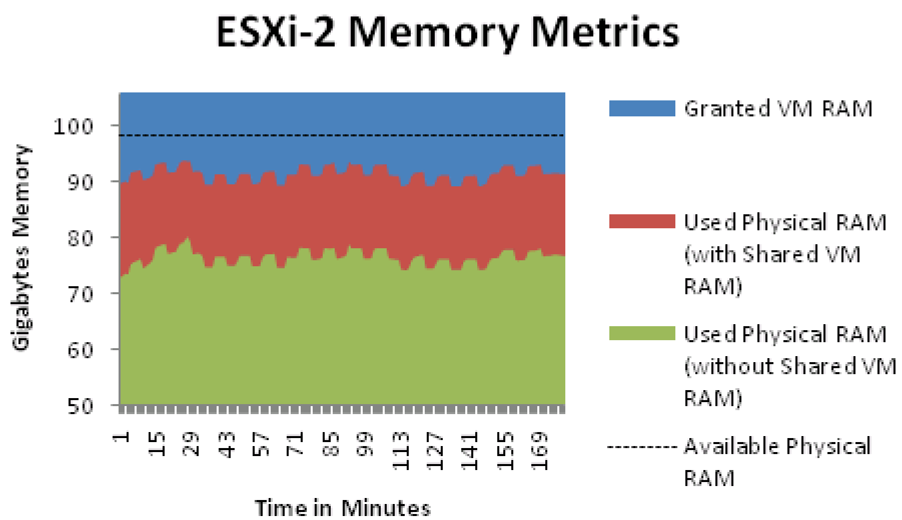

Figure 17

Figure 18 illustrates a sample of a mail server and DVD Store 2 database server, showing the combined VMDK IOPS statistics for those specific virtual machines.

- All servers, with the exception of the mail servers and DVD Store 2 database servers, had relatively low I/O with none peaking over 70 IOPS. Therefore, they were omitted from Figure 18.
- Due to their high IO profile, no more than three mail server virtual machines were stored on a LUN to avoid LUN queue depth issues.

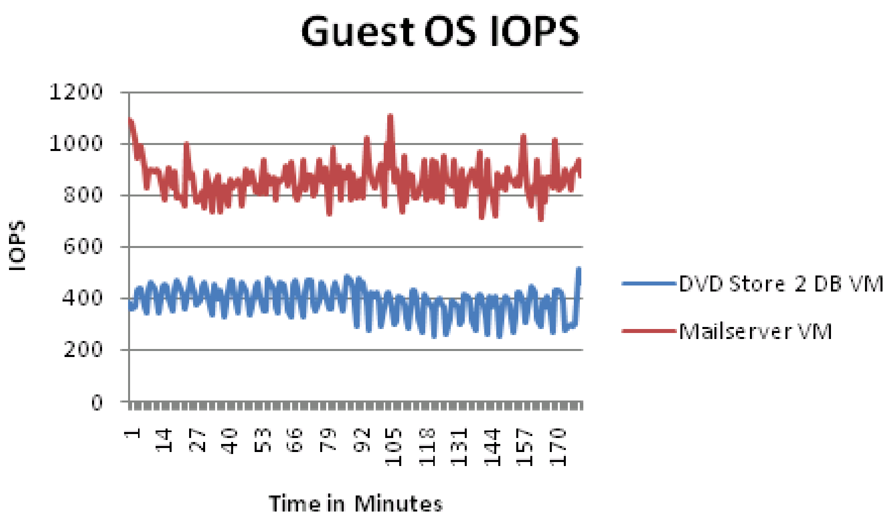

Figure 18

Storage Infrastructure

Physical disk performance was acceptable during the general server application workloads.

Figure 19 on page 41 illustrates the average performance for physical drives from both dynamic provisioning pools. The percent busy statistic indicates the workloads are generating medium to high disk utilization but well within an acceptable range. These disk performance graphs show the following:

- Disk utilization peaked at 47% for dynamic provisioning Pool 0 disk and 62% for dynamic provisioning Pool 1 disk.
- There was between 40% to 50% headroom during steady state peaks.

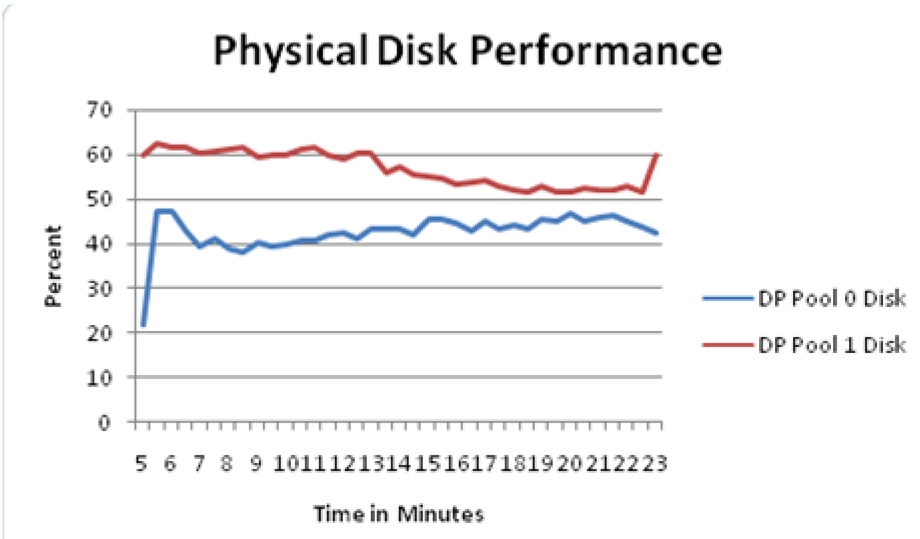

Figure 19

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

Figure 20 and Figure 21 on page 42 illustrate the processor and cache performance metrics.

- There are eight processor cores per controller, totaling sixteen cores. Only storage processor 20-00 saw regular fluctuation in utilization, which can be common for server workloads. However, steady state was generally below thirty percent for the other three storage processors.
    - Only 4 of the 16 cores saw any regular usage and so only those 4 were diagrammed in Figure 20.
    - The other 12 storage processor cores usage was minimal and not significant enough to show in the figure.
- Write pending percentage peaked at sixteen percent indicating plenty of headroom for future growth.

Table 13 on page 36 shows why the utilization of some MPUs is higher or lower than others.

This data indicates the Hitachi Unified Storage VM controllers are capable of handling various server applications with minimal fluctuation in workload.

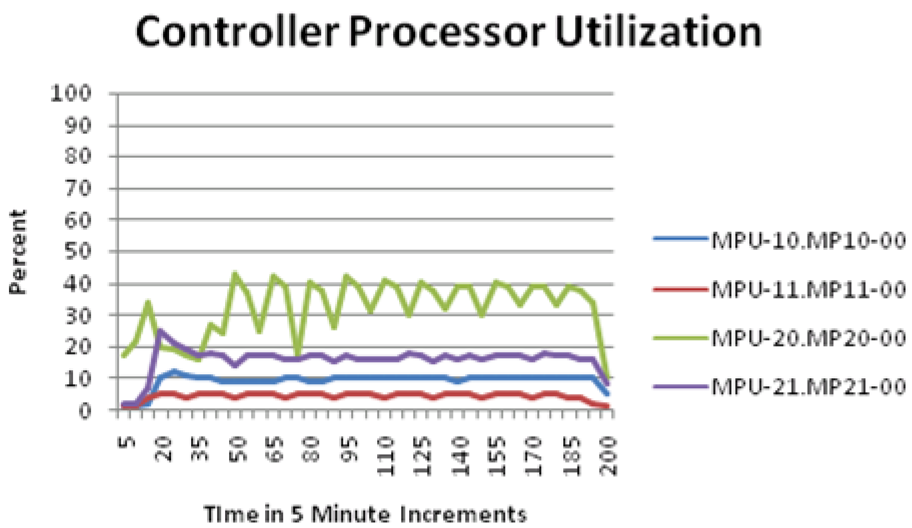

Figure 20

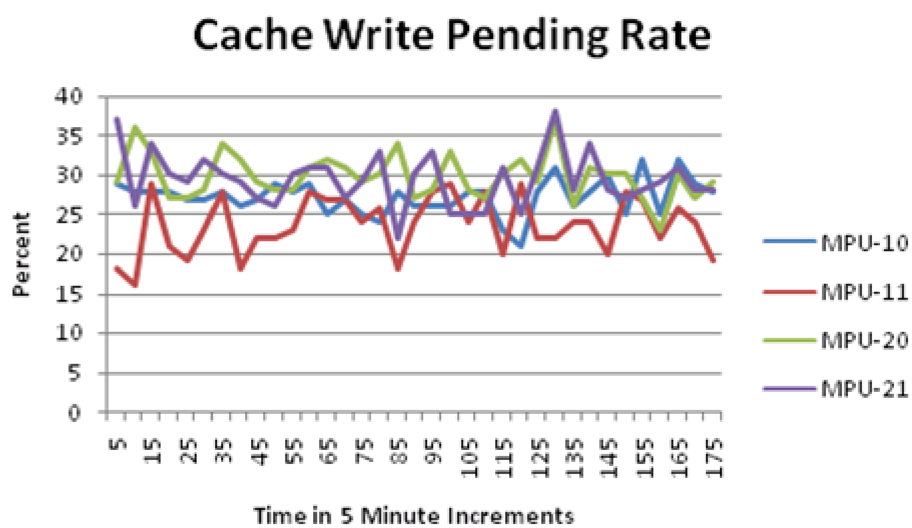

Figure 21

Figure 22 illustrates the aggregated performance for dynamic provisioning Pool 0 dedicated to OLTP workloads.

- The average read IOPS was 1883 while average write IOPS was 1306 indicating a 3:2 read/write ratio.
- The workload was 99% random. RAID-10 was used to maximize performance for random workloads.

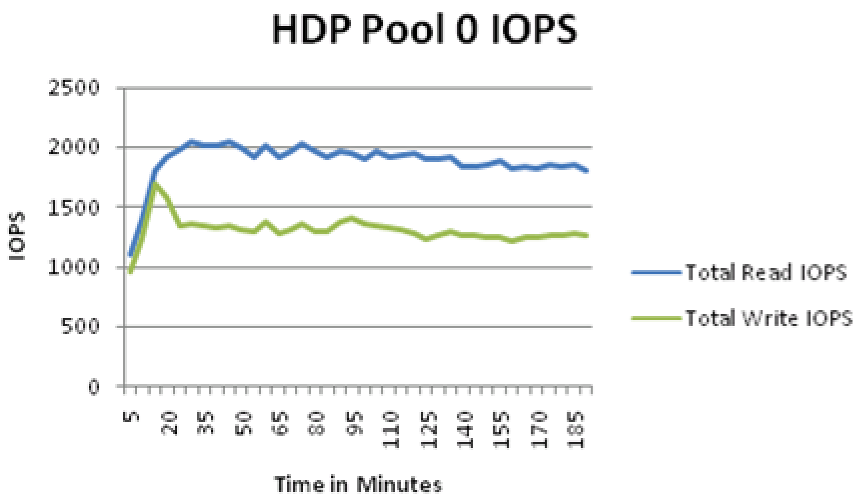

Figure 22

Figure 23 on page 44 illustrates the aggregated performance for Dynamic Provisioning Pool 1, which hosted all other workloads.

- The average read IOPS was 702 while average write IOPS was 6540 indicating a 1:9 read/write ratio.
- The workload was 99% random. RAID-10 was used to maximize performance for random workloads.

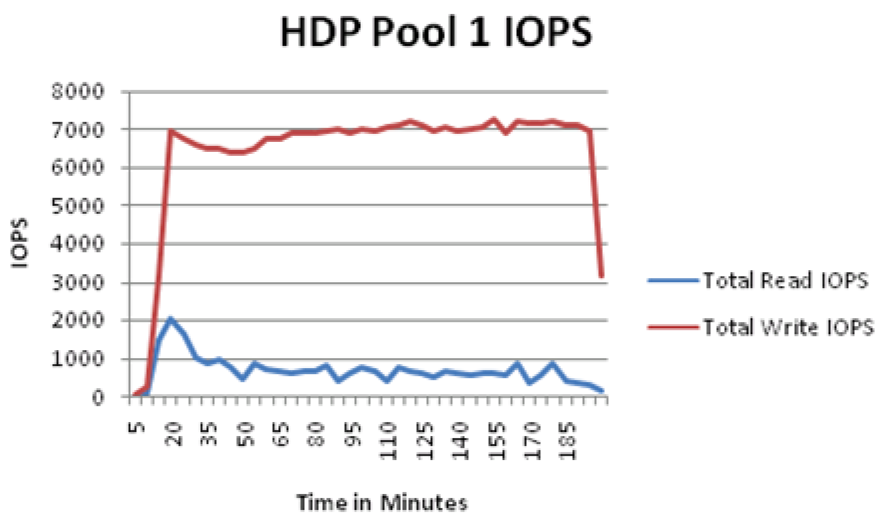

Figure 23

Figure 24 shows the LDEV latency from the storage device. The LDEV latency is far less than the application latency shown in Figure 26 on page 45. Unlike application latency, LDEV latency does not include the effect of things like CPU and memory latency.

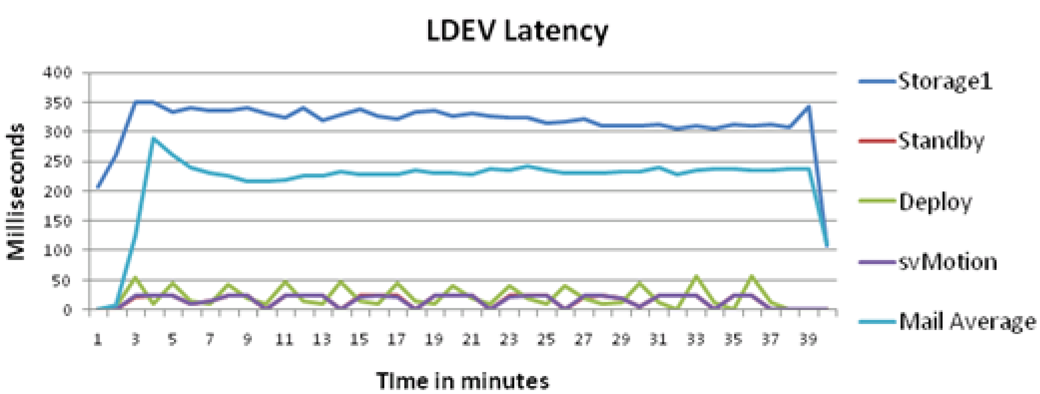

Figure 24

Note — The LDEV latency in Figure 24 is in microseconds. The application latency in Figure 26 on page 45 is in milliseconds.

Separating workloads into different pools based on I/O profile did the following:

- Decreased the disk latency reported by the ESXi hosts
- Provided a boost in storage performance

Application Performance

Figure 25 shows the application workload throughput.

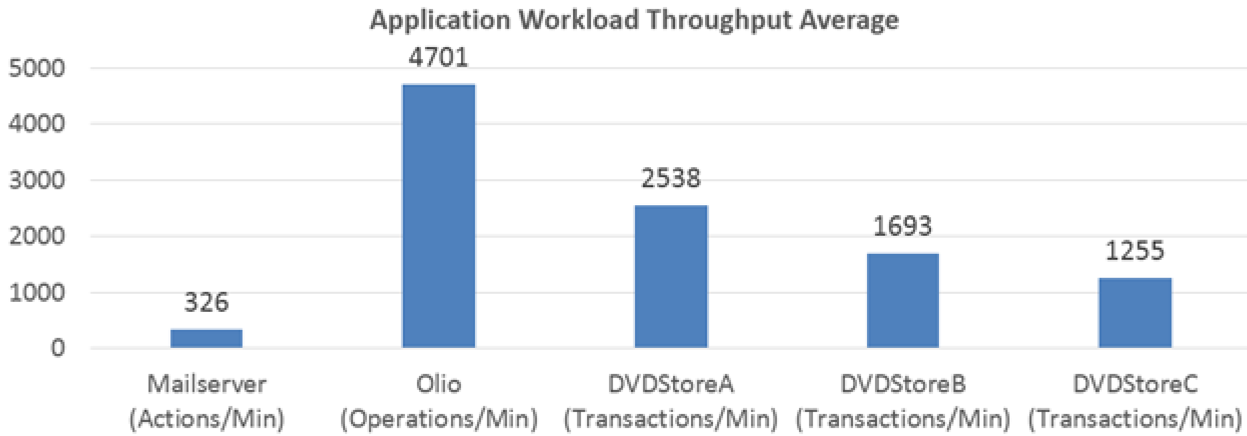

Figure 25

Figure 26 shows the application workload latency.

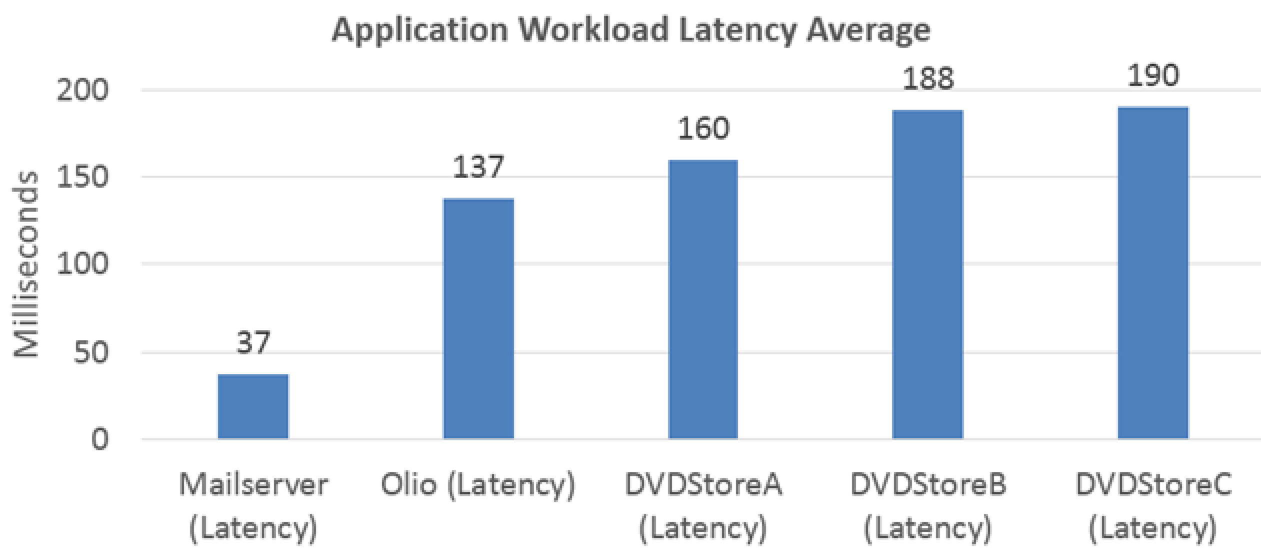

Figure 26

Conclusion

This reference architecture guide discusses how to design a Hitachi Unified Compute Platform Select for VMware vSphere solution with Hitachi Unified Storage VM. The purpose of the general server application workloads testing in the Hitachi Data Systems laboratory was to provide general guidance on the virtual resources available with this solution.

Each implementation has its own unique set of data center and application requirements. Design your implementation of this environment by understanding the I/O workload of the server applications in your environment. Creating an environment that meets your unique needs results in increased ROI from avoiding over or under provisioning resources.

Use Hitachi Dynamic Provisioning to reallocate I/O capabilities dynamically, as necessary. Having the capability to provision additional spindles to an already-provisioned datastore within vSphere allows for non-disruptive upgrades to the underlying storage infrastructure. This provides immediate benefits to your environment without confusing shuffling of virtual machines, datastores, or LUs.

This Unified Compute Platform Select design gives you a build-as-you-go model that uses performance-proven hardware resources, including Unified Storage VM and Hitachi Compute Blade 500. The modular design, using a cell architecture, permits implementing an environment for modest needs that gives you the flexibility to scale out as your IT needs grow.

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 95050-2639
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-187-00, January 2013