



Deploying Microsoft® Exchange Server® 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Reference Architecture Guide

By Jeff Chen

May 2011



Feedback

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

Table of Contents

Solution Overview	3
Key Solution Components.....	5
Hitachi Virtual Storage Platform	5
Hitachi Dynamic Tiering Software	5
Hitachi Compute Blade 2000.....	7
Hitachi Dynamic Link Manager Software	10
Microsoft Exchange Server 2010	10
Solution Design	12
Hitachi Compute Blade 2000 Chassis Configuration	13
Hitachi Compute Blade 2000 Blade Configuration.....	13
Designing Exchange Server Roles	16
Mailbox High Availability.....	16
SAN Architecture	18
Network Architecture	19
Storage Architecture.....	21
Engineering Validation.....	26
Microsoft Exchange Jetstress 2010 Test	26
Microsoft Exchange Load Generator 2010 Test	30
Conclusion	31

Deploying Microsoft® Exchange Server® 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Reference Architecture Guide

This reference architecture deploys 20,000 users using Microsoft Exchange Server 2010 on a Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering. All Exchange servers were installed on logical partitions (LPAR) on a Hitachi Compute Blade 2000. Hitachi Dynamic Tiering software allows combining different tiers of storage media into a single storage pool. Using Hitachi Dynamic Tiering reduces the cost of an Exchange 2010 deployment by reducing the total amount of drives required while maintaining the performance requirements for a given environment.

Use of this reference architecture is for IT administrators involved in planning and deploying Microsoft Exchange Server 2010 and storage. There is a focus on planning and deploying the Exchange Server 2010 mailbox role. It assumes some familiarity with Hitachi Storage Navigator software, Microsoft Windows® 2008 R2 and Exchange Server 2010.

For information about deploying other Exchange server roles, see the Microsoft TechNet library [“Exchange Server 2010.”](#)

IT administrators need email solutions that provide data protection and simple management in environmentally friendly data centers. Using Microsoft® Exchange Server® 2010 with storage area network (SAN) storage like the Hitachi Virtual Storage Platform accomplishes those business-critical objectives. This reference architecture uses changes in Exchange Server 2010 and the advantages of the Virtual Storage Platform to lower administration costs, improve system efficiency, and enable virtualization.

Solution Overview

Deploying all Exchange components, including Active Directory, mailbox, hub transport and client access servers, in an environment using LPARs offers the greatest overall benefit. Fewer physical servers housing multiple Exchange server roles with better resource utilization increases the agility and availability of Exchange while lowering costs.

This reference architecture uses:

- An Hitachi Compute Blade 2000 with five server blades
- Microsoft Windows 2008 R2 Enterprise edition installed on logical partitions on the blades

The servers installed on the logical partitions include:

- Four Microsoft Exchange 2010 mailbox servers, with each server supporting three active and three passive mailbox databases.
- Four Exchange 2010 servers with the client access and hub transport roles installed.
- One Microsoft Windows 2008 domain controller.

Figure 1 shows the layout of the logical partitions (LPAR) on the Hitachi Compute Blade 2000 and the connections to the Hitachi Virtual Storage Platform.




Figure 1

This reference architecture is designed for 20,000 users with 1GB mailboxes. Based on the Exchange 2010 Mailbox Server Role Requirements Calculator results and Microsoft recommendations that the size of mailbox databases in DAG environments should not exceed 2TB, Hitachi Data Systems used 12 active databases and 12 passive databases (1,667 users per database) to house 20,000 users. For more information about the calculator, see the [“Exchange 2010 Mailbox Server Role Requirements Calculator”](#) entry on the Microsoft Exchange Team Blog.

Key Solution Components

The following sub-sections describe the key hardware and software components used to deploy this solution.

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform helps you use your information, important in today's data-driven economy. Information exists in many forms. It must be protected and readily accessible to ensure business survival and success. The Virtual Storage Platform maximizes cost efficiency and return on investment by creating a nimble storage infrastructure that reduces costs and increases performance, availability, scalability, and reliability.

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

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

Microsoft Exchange Server 2010 can be deployed for many types and numbers of users in a wide variety of infrastructure topologies. Implement Exchange Server 2010 with the Virtual Storage Platform storage system, to scale an environment from a few thousand users to hundreds of thousands of users effectively.

For more information about the Hitachi Virtual Storage Platform, see the Hitachi Data Systems [web site](#).

Hitachi Dynamic Tiering Software

Hitachi Dynamic Tiering eliminates the time-consuming manual processes of data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance. It is available only on the Hitachi Virtual Storage Platform.

Most data is accessed rarely after creation. Data that is accessed infrequently should not be stored on your most expensive storage tier. It should be moved to a lower, less expensive storage tier. Defining where and how long data should reside at any point in its life cycle can be complex and problematic.

Many organizations use data tiers to manage application performance, manually provisioning space from several storage technologies with different performance and cost characteristics. Data specialists look to past usage patterns to determine how to configure tiers manually. The storage infrastructure does not respond effectively to dynamic application and data use. When usage patterns change rapidly, manually tiered storage systems produce less than optimal results.

Using existing features of Hitachi Dynamic Provisioning software, Hitachi Dynamic Tiering automates the management of multiple storage tiers as a single entity. First, Hitachi Dynamic Provisioning software breaks the volume into 42MB pages. Then, after an initial setup process, Hitachi Dynamic Tiering monitors data access in real time to make decisions on moving pages of data between the available storage tiers based on actual use. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost storage tiers. When there is frequent access to a page on a lower tier, Hitachi Dynamic Tiering automatically moves the page to a higher tier. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

This maximizes service levels while it minimizes total cost of storage ownership. Additionally, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Previously, the creation of each dynamic provisioning pool had to use one RAID level and one disk type. Hitachi Dynamic Tiering on the Hitachi Virtual Storage Platform allows a single pool to contain tiers made up of multiple types of RAID groups and any type of disk, managing the various tiers within a dynamic provisioning pool automatically. This eliminates most user management of storage tiers within a storage system to maintain peak performance under dynamic conditions without storage administrator intervention.

Hitachi Dynamic Tiering inherits the Hitachi Dynamic Provisioning software program's thin provisioning and wide striping functionalities. This provides virtual storage capacity to eliminate application service interruptions, reduce costs, and simplify administration by:

- Optimizing storage performance and capacity based on business or application requirements.
- Supporting deferring storage capacity upgrades to align with actual business usage.
- Simplifying the storage administration process.
- Providing performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.
- Eliminating hot spots across the different RAID groups by smoothing the combined workload.
- Improving capacity utilization significantly.

Use Hitachi Dynamic Tiering Software to ensure that the right data is in the right place at the right time. This maximizes service levels and minimizes total cost of storage ownership. For more information, see the Hitachi Dynamic Tiering software datasheet.

Hitachi Compute Blade 2000

The Hitachi Compute Blade 2000 is an enterprise-class platform that offers the following:

- Balanced system architecture to eliminate bottlenecks in performance and throughput
- Embedded Hitachi logical partitioning (LPAR) virtualization
- Unprecedented configuration flexibility
- Eco-friendly power saving features and capabilities
- Fast recovery from server failures due to N+1 cold standby design to allow replacing failed servers within minutes

With its unique combination of power, efficiency, and flexibility, extend the benefits of virtualization to new areas of the enterprise data center — including mission-critical application servers and database servers — with minimal cost and maximum simplicity.

The Hitachi Compute Blade 2000 features a modular architecture to deliver unprecedented configuration flexibility, as shown in Figure 2.




Figure 2

The Hitachi Compute Blade 2000 combines all the benefits of virtualization with all the advantages of the blade server format: simplicity, flexibility, high compute density, and power efficiency. Take advantage of the following:

- Consolidate more resources.
- Extend the benefits of virtualization solutions, whether Hitachi logical partitioning, VMware vSphere, Microsoft Hyper-V, or all three.
- Cut costs without sacrificing performance.

The Hitachi Compute Blade 2000 enables you to consolidate application and database servers for backbone systems using virtualization, areas where effective consolidation was difficult in the past. By removing performance and I/O bottlenecks, Hitachi Compute Blade 2000 opens new opportunities for increasing efficiency and utilization rates to reduce the administrative burden in your data center.

Configure and administer the Hitachi Compute Blade 2000 using an HTML browser that supports secure encrypted communications. Or use the optional management suite to manage multiple chassis using a unified GUI-based interface.

Chassis

The Hitachi Compute Blade 2000 is a 19-inch rack compatible, 10U-high chassis with a high degree of configuration flexibility. The front of the chassis has:

- Up to 8 server blades
- Up to 4 power supply modules

The back of the chassis has:

- 6 bays for I/O switch modules
- 8 fan modules
- Up to 2 management modules
- 16 half-height PCIe slots
- 2 AC power input modules

The Hitachi Compute Blade 2000 can be configured with mirrored power supplies, providing backup on each side of the chassis for higher reliability.

Efficient, variable-speed, redundant fan modules provide cooling. Each fan module includes 3 fans to tolerate a fan failure within a module. If an entire module fails, the other fan modules continue to cool the chassis.

To view the specific server configuration used in this reference architecture, see “Hitachi Compute Blade 2000 Chassis Configuration.”

Server Blades

The Hitachi Compute Blade 2000 supports 2 blade server options. These can be combined within the same chassis. Table 1 lists the specifications for each option.

Table 1. Server Blade Specifications

<i>Feature</i>	<i>X55A2</i>	<i>X57A1</i>
Processors (up to two per blade)	Intel Xeon 5600 — 4 or 6 cores	Intel Xeon 7500 — 6 or 8 core
Processor cores	4,6,8, or 12	6, 8, 12, or 16
Memory slots	18	32
Maximum memory	144GB (with 8GB DIMMS)	256GB (with 8GB DIMMS)
Hard drives	Up to 4	N/A
Network interface cards (on-board)	Up to 2 1Gb/sec Ethernet	Up to 2 1Gb/sec Ethernet
Other interfaces	2 USB 2.0 ports and 1 serial port	2 USB 2.0 ports and 1 serial port
Mezzanine slots	2	2
PCIe 2.0 (8x) expansion slots	2	2

Up to 4 X57A1 blades can be connected using the SMP interface connector to create a single 8-socket SMP system with up to 64 cores and 1024GB of memory.

I/O Options

The connections from the server blades through the chassis' mid-plane to the bays or slots on the back of the chassis consist of the following:

- Two on-board NICs to connect to switch bays one and two.
- Optional mezzanine card in mezzanine slot 1 to connect to switch bays 3 and 4
- Optional mezzanine card in mezzanine slot 2 to connect to switch bays 5 and 6
- Two connections to dedicated PCIe slots

The I/O options supported by the optional mezzanine cards and the switch modules are either 1Gb/sec Ethernet or 8Gb/sec Fibre Channel connectivity.

For more information about currently supported mezzanine cards, switch modules, and PCIe cards, see the Hitachi Compute Blade 2000 datasheet.

Logical Partitioning

The firmware of the Hitachi Compute Blade 2000 server blades has the Hitachi Compute Blade 2000 logical partitioning (LPAR) feature embedded. This proven, mainframe-class technology combines Hitachi logical partitioning expertise with Intel VT technologies to improve performance, reliability, and security.

Unlike emulation solutions, this embedded logical partitioning virtualization feature does not degrade application performance. Unlike third-party virtualization solutions, it does not require the purchase and installation of additional components. This keeps the total cost of ownership low.

A blade can operate in one of two modes, basic or Hitachi Virtualization Manager (HVM). There are 2 licensing options available in the HVM mode:

- **Basic** — The blade operates as a standard server without logical partitioning support
- **HVM with essential license** — The blade supports 2 logical partitions. No additional purchase is required for this mode.
- **HVM with enterprise license** — The blade supports up to 16 logical partitions. The enterprise license is an additional cost.

Additional flexibility is provided by using the embedded logical partitioning feature alone or combined with Microsoft Hyper-V, VMware vSphere, or both in a single system.

Hitachi Blade Server Management Modules

The Hitachi Compute Blade 2000 supports up to two management modules to provide redundancy. Each module may be hot-swapped. They support live firmware updates without the need for shutting down the blades. Each module supports an independent management LAN interface from the data network for remote and secure management of the chassis and all blades. Each module supports a serial command line interface and a web interface. SNMP and email alerts are supported.

N+1 or N+M Cold Standby Failover

The Hitachi Compute Blade 2000 maintains high uptime levels through sophisticated failover mechanisms. The N+1 cold standby function enables multiple servers to share 1 standby server. This increases system availability while it decreases the need for multiple standby servers or costly software-based high-availability servers. The system can detect a fault in a server blade and switch to the standby server, manually or automatically. Hardware switching in the absence of an administrator enables the system to return to normal operations within a short time.

The N+M cold standby function has “M” backup server blades for every “N” active server blade, so failover is cascading. In the event of multiple hardware failures, the system automatically detects the fault and identifies the problem by indicating the faulty server blade, allowing immediate failure recovery. This approach can reduce total downtime by sharing of the application workload among the working servers.

Hitachi Dynamic Link Manager Software

The Hitachi Dynamic Link Manager software was used for SAN multipathing, configured with the round-robin multipathing policy. The round-robin load balancing algorithm of Hitachi Dynamic Link Manager software automatically selects a path by rotating through all available paths. This balances the load across all available paths and optimizing IOPS and response time.

Microsoft Exchange Server 2010

Database Availability Groups

To support database mobility and site resiliency, Microsoft Exchange Server 2010 introduced Database Availability Groups (DAGs). A DAG is an object in Active Directory that can include up to 16 mailbox servers that host a set of databases. Any server within a DAG has the ability to host a copy of a mailbox database from any other server within the DAG. This supports mailbox database replication as well as database and server switchovers and failovers.

Setting up a Windows failover cluster is no longer necessary for high availability. However, the prerequisites for setting up a DAG are similar to that of a failover cluster. Hitachi Data Systems recommends using DAGs for high availability and mailbox resiliency.

shows 12 active databases and 12 passive databases placed in a single DAG. The databases are stored in a Hitachi Dynamic Tiering pool on a Hitachi Virtual Storage Platform system.

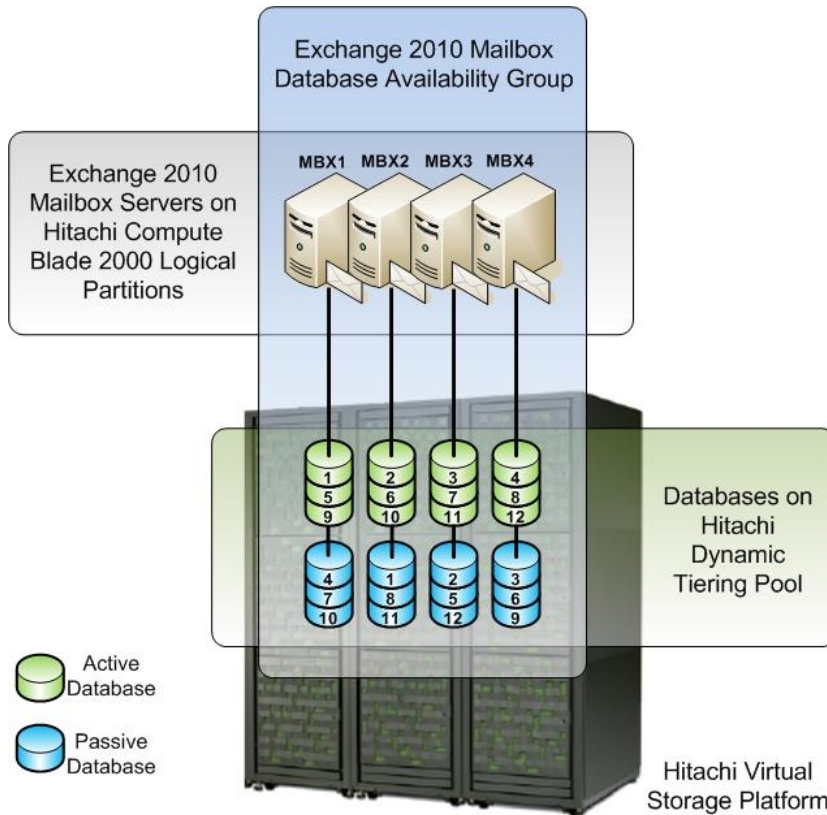


Figure 3

Databases

In Microsoft Exchange Server 2010, the changes to the Extensible Storage Engine (ESE) enable the use of large databases on larger, slower disks while maintaining adequate performance. The Exchange Store's database tables make better use of the underlying storage system and cache. The Store no longer relies on secondary indexing, making it less sensitive to performance issues.

Exchange 2010 supports databases up to approximately 16TB, but Microsoft recommends using databases of 2TB or less when using DAGs.

Solution Design

Testing performed to validate this 20,000-user reference architecture used a Hitachi Virtual Storage Platform. The user profile assumes each user sends and receives an average of 100 messages per day, which gives an I/O profile of 0.12 IOPS per mailbox. This is information regarding the following:

- The design of this solution
- The factors to be considered when designing the storage infrastructure
- The appropriate sizing of the solution

Table 2 details the information about the hardware components used in the Hitachi Data Systems lab.

Table 2. Hardware Components

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Virtual Storage Platform	6 × 8Gb/sec Fibre Channel ports used 128GB cache memory 96 × 300GB, 10K RPM, SAS disks 44 × 2TB 7.2K RPM, SATA disks	70-01-62-00/00	1
Hitachi Compute Blade 2000 Chassis	10 × 8Gb/sec dual-port HBAs 2 × 1Gb/sec LAN Switch Module 2 × management modules 8 × cooling fan modules 4 × power supply modules	A0154-E-5234	1
Hitachi Compute Blade 2000 — X55A2 Blade	2 × 6-Core Intel Xeon X5670 2.93GHz 72GB memory	58.22	4
Hitachi Compute Blade 2000 — X57A1 Blade	2 × 8-Core Intel Xeon X7560 2.26GHz 256GB memory	58.22	1
Brocade 5300 Fibre Channel Switch	8Gb/sec Fibre Channel ports	FOS 6.4.0E	2

Table 3 has the software components used in the Hitachi Data Systems lab.

Table 3. Software Components

<i>Hardware</i>	<i>Version</i>
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Tiering	Dependent on microcode version
Hitachi Dynamic Link Manager	6.5
Windows Server	2008 R2 Enterprise
Microsoft Exchange Server	2010 SP1
Microsoft Exchange Jetstress 2010	14.01.0225.017
Microsoft Exchange Load Generator 2010	14.01.0180.003

Hitachi Compute Blade 2000 Chassis Configuration

This reference architecture uses:

- 1 SMP X57A1 blade
- 4 standard X55A2 blades
- 2 1Gb/sec LAN switch modules
- 10 Hitachi dual port, 8Gb/sec HBA cards.

Each server blade has two on-board NICs. Each NIC is connected to a LAN switch module.

Each blade has two PCIe slots available. All slots have Hitachi dual port 8Gb/sec HBA cards installed. Hitachi HBAs are required to support logical partitioning. With Hitachi HBAs, when logical partitioning is enabled, it creates eight virtual WWNs for the logical partitions on each blade. Figure 4 shows the front and back view of Hitachi Compute Blade 2000 used in this solution.




Figure 4

Hitachi Compute Blade 2000 Blade Configuration

To meet the 20,000-user requirement and to use the hardware resources efficiently, the Microsoft Windows servers listed in Table 4 were deployed on the five physical server blades.

Table 4. Hitachi Compute Blade 2000 Configuration

<i>Blade Server</i>	<i>Logical Partition</i>	<i>Server Name</i>	<i>Role</i>	<i>Number of CPU Cores</i>	<i>Memory (GB)</i>
Blade 1	LPAR2	AD2	Active Directory and DNS	4	8
Blade 2	LPAR1	MBX1	Mailbox Server	6	64
	LPAR2	CAS-HT1	Client Access and Hub Transport Server	6	8
Blade 3	LPAR1	MBX2	Mailbox Server	6	64
	LPAR2	CAS-HT3	Client Access and Hub Transport Server	6	8
Blade 4	LPAR1	MBX3	Mailbox Server	6	64
	LPAR2	CAS-HT3	Client Access and Hub Transport Server	6	8
Blade 5	LPAR1	MBX4	Mailbox Server	6	64
	LPAR2	CAS-HT4	Client Access and Hub Transport Server	6	8

Processor Capacity

With the release of Microsoft Exchange 2010, Microsoft has new processor configuration recommendations for servers that host the mailbox role resulting from the implementation of mailbox resiliency. The recommendations are based on two factors:

- Whether the server hosts active and passive database copies
- The number of database copies

A passive database copy requires CPU resources to perform the following tasks:

- Check or validate replicated logs
- Replay replicated logs into the database
- Maintain the content index associated with the database copy

For this reference architecture, the following formulas were used to calculate the CPU requirement for the Microsoft Exchange roles, taking failover into consideration:

$$\text{Megacycles per mailbox} = \frac{(\text{average CPU usage} \times \text{speed of processors in megacycles}) \times (\text{number of processors} \div \text{number of mailboxes})}{\text{number of processors} \times \text{speed of processors in megacycles}}$$

$$\text{CPU usage} = \frac{(\text{number of users} \times \text{current megacycles per mailbox}) \div (\text{number of processors} \times \text{speed of processors in megacycles})}{\text{number of processors} \times \text{speed of processors in megacycles}}$$

Physical Memory

This reference architecture supports 20,000 users who send and receive 100 messages per day. Table 5 has the database cache guidelines.

Table 5. Database Cache Guidelines

<i>Messages Sent and Received per Mailbox per Day</i>	<i>Database Cache per Mailbox (MB)</i>	<i>Standalone Estimated IOPS per Mailbox</i>	<i>Mailbox Resiliency Estimated IOPS per Mailbox</i>
50	3	0.06	0.05
100	6	0.12	0.10
150	9	0.18	0.15
200	12	0.24	0.20
250	15	0.30	0.25
300	18	0.36	0.30
350	21	0.42	0.35
400	24	0.48	0.40
450	27	0.54	0.45
500	30	0.60	0.50

Based on Table 5:

- The mailbox resiliency estimate per mailbox was 0.10 IOPS. Adding a 20 percent overhead gives a total of 0.12 IOPS
- The database cache size per mailbox is 6MB

This reference architecture design is for 20,000 users spread over 4 servers, or 5,000 users per server. Each server also requires cache for 1,667 additional users in the case of a server failure where one of the passive databases becomes active. At 6MB cache per user mailbox, Table 6 shows the calculation for the total cache size needed on each server.

Table 6. Total Cache Calculation

<i>Type</i>	<i>Number of Users</i>	<i>Cache Size Needed</i>
Active	5,000	30,000MB
Passive	1,667	About 10,000MB
Total	6,667	About 40,000MB or about 40GB

After determining the 40GB database cache size, Hitachi Data Systems used Table 7 to determine the amount of physical memory. Using 64GB is the ideal memory configuration, based on this mailbox count and user profile. Table 7 lists Microsoft’s guidelines for determining physical memory capacity.

Table 7. Physical Memory Guidelines

<i>Server Physical Memory (GB)</i>	<i>Database Cache Size (GB)</i>
2	0.5
4	1.0
8	3.6
16	10.4
24	17.6
32	24.4
48	39.2
64	53.6
96	82.4

Designing Exchange Server Roles

In this reference architecture, the client access server and hub transport server roles are combined, while the mailbox role was installed on a separate logical partition. The primary reasons are:

- To minimize the number of servers, operating system instances, and Exchange servers to manage
- To optimize performance for planned or unplanned failover scenarios

When deploying roles in isolation, Microsoft recommends:

- One hub transport server logical processor for every 5 mailbox server logical processors
- Three client access server logical processors for every 4 mailbox server logical processors

This can become confusing, especially when you factor in providing sufficient client access and hub transport servers during multiple VM server failures or maintenance scenarios. Simplify the configuration by deploying one combined client access and hub transport server for every one mailbox server, with each server having the same number of virtual CPUs.

For more information, see the Microsoft TechNet article "[Understanding Client Access and Hub Transport Combined Role Configurations in Capacity Planning.](#)"

Mailbox High Availability

Because the Microsoft Exchange mailbox databases reside on intelligent, RAID-protected disks this reference architecture only deploys two copies of the databases: 1 active and 1 passive. The decision to use a single DAG was based on Microsoft's recommendation to minimize the number of DAGs. Consider using more than one DAG only if one of the following conditions applies to your environment:

- More than 16 mailbox servers
- Active mailbox exist users in multiple sites
- Requirement of separate DAG-level administrative boundaries
- Mailbox servers exist in separate domains

Each mailbox server hosts 3 active databases and 3 passive databases for 5,000 users. The passive databases are distributed evenly across the mailbox servers, so that if one mailbox server fails, one passive mailbox database on each of the other 3 servers becomes active. Each server is able to handle the load of all 6667 users (5000 active users, plus the 1667 passive users that could become active with a server failure).

For more information, see the Microsoft TechNet article "[Database Copy Layout Design.](#)"

Table 8 shows the DAG database distribution used in the Hitachi Data Systems lab.

Table 8. DAG Database Distribution

<i>Database</i>	<i>MBX1</i>	<i>MBX2</i>	<i>MBX3</i>	<i>MBX4</i>
DB1	Active	Passive		
DB2		Active	Passive	
DB3			Active	Passive
DB4	Passive			Active
DB5	Active		Passive	
DB6		Active		Passive
DB7	Passive		Active	
DB8		Passive		Active
DB9	Active			Passive
DB10	Passive	Active		
DB11		Passive	Active	
DB12			Passive	Active

Microsoft recommends having at least three database copies (one active and two passive) when deploying with direct attached storage (DAS) or with just a bunch of disks (JBOD). This takes in consideration server failure and storage (hard drive) failure.

However, in this reference architecture, only two copies of the databases are deployed (one active and one passive) because the Microsoft Exchange mailbox databases reside on the Hitachi Virtual Storage Platform. Providing high performance and the most intelligent RAID-protected storage system in the industry, the Virtual Storage Platform reduces the possibility of storage failure. The consideration in this case is protection against server failure.

Reducing database copies to two instead of three provides a number of benefits, including:

- Uses less storage
- Requires less server resources
- Consumes less network traffic to replicate passive databases
- Reduces the number of databases to manage

The number of database copies required in a production environment also depends on factors such as the use of lagged database copies, or backup and recovery methodologies.

SAN Architecture

To provide high availability, the storage area network (SAN) configuration for this reference architecture uses two Fibre Channel switches. SAN boot volumes for the OS connects to two HBA ports. Microsoft Exchange volumes connect to 4 HBA ports. There are 4 redundant paths from the switches to the Virtual Storage Platform configured for Microsoft Exchange volumes. The SAN OS boot volumes use the remaining 2 redundant paths.

Exchange databases and logs are mapped to ports 1A, 2A, 5A and 6A on the Virtual Storage Platform.. The SAN OS boot volumes are mapped to ports 3B and 8B. Hitachi Dynamic Link Manager software is used for multipathing with the round-robin load balancing algorithm.

Figure 5 illustrates the SAN design used for this reference architecture.

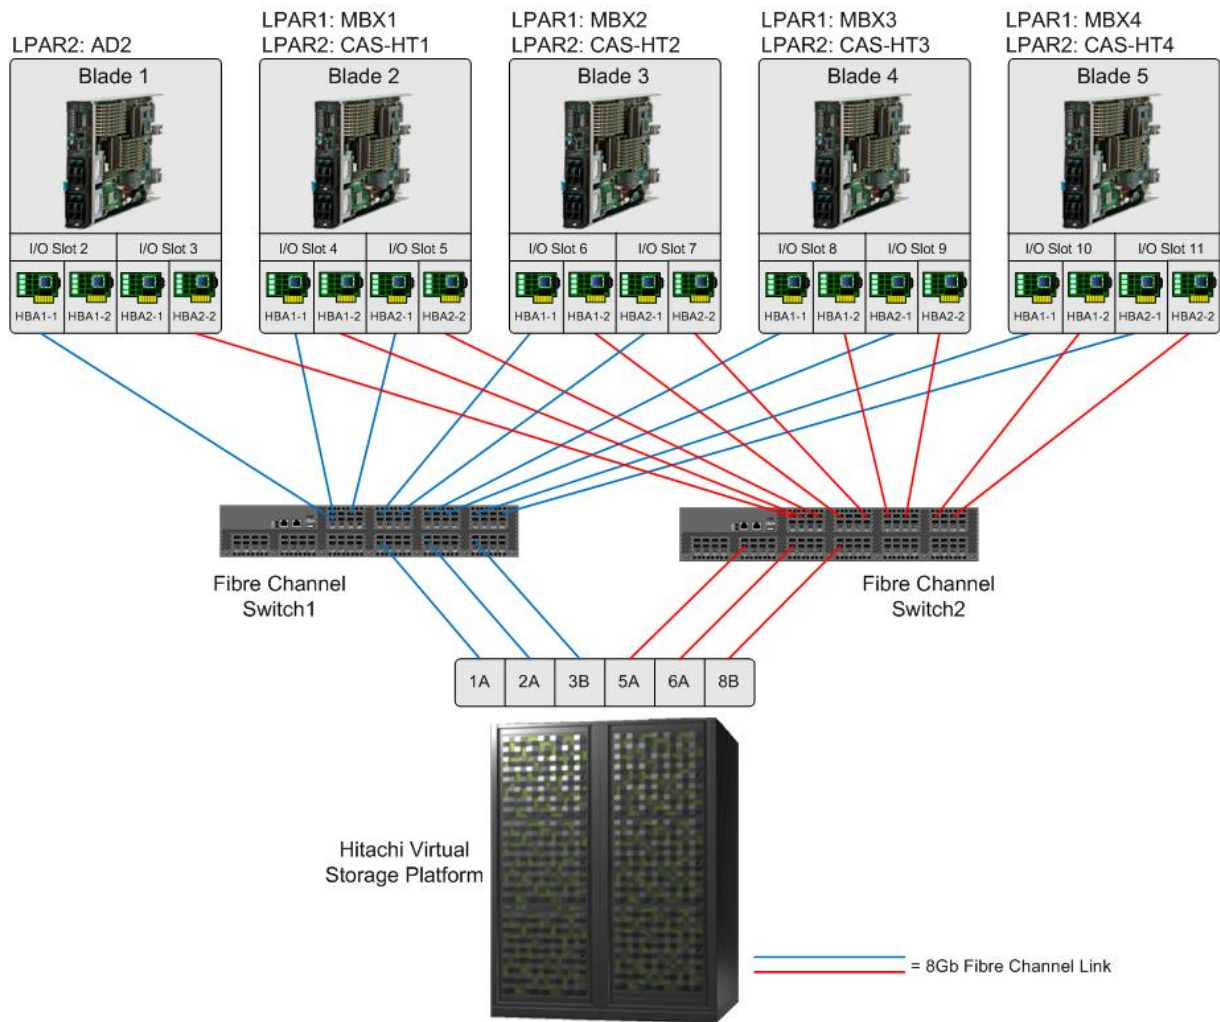


Figure 5

Network Architecture

When deploying Exchange mailbox servers in a DAG, Hitachi Data Systems recommends having separate local area network subnets available to the members for:

- Client access
- Replication

The configuration is similar to the public, mixed, and private networks used in previous versions of Microsoft Exchange. In Exchange 2010 the two networks are referred to as the MAPI network, and the replication network which have the following uses:

- MAPI network - dedicated to communication among the DAG members and other Exchange servers
- Replication network - dedicated to log shipping and seeding

While Microsoft supports using a single network configuration, Hitachi Data Systems does not recommend it. Having at least two networks connected to two separate network adapters in each server:

- Provides redundancy
- Enables Microsoft Exchange to distinguish between a server failure and a network failure

Each DAG member must have the same number of networks and at least one MAPI network.

If your environment has members of the same DAG located in different geographic sites, follow Microsoft's guidance regarding network latency for DAG members.

For more information about network planning for Exchange Server 2010, see the Microsoft TechNet article "[Planning for High Availability and Site Resilience.](#)"

Figure 6 shows the network configuration used for this reference architecture.

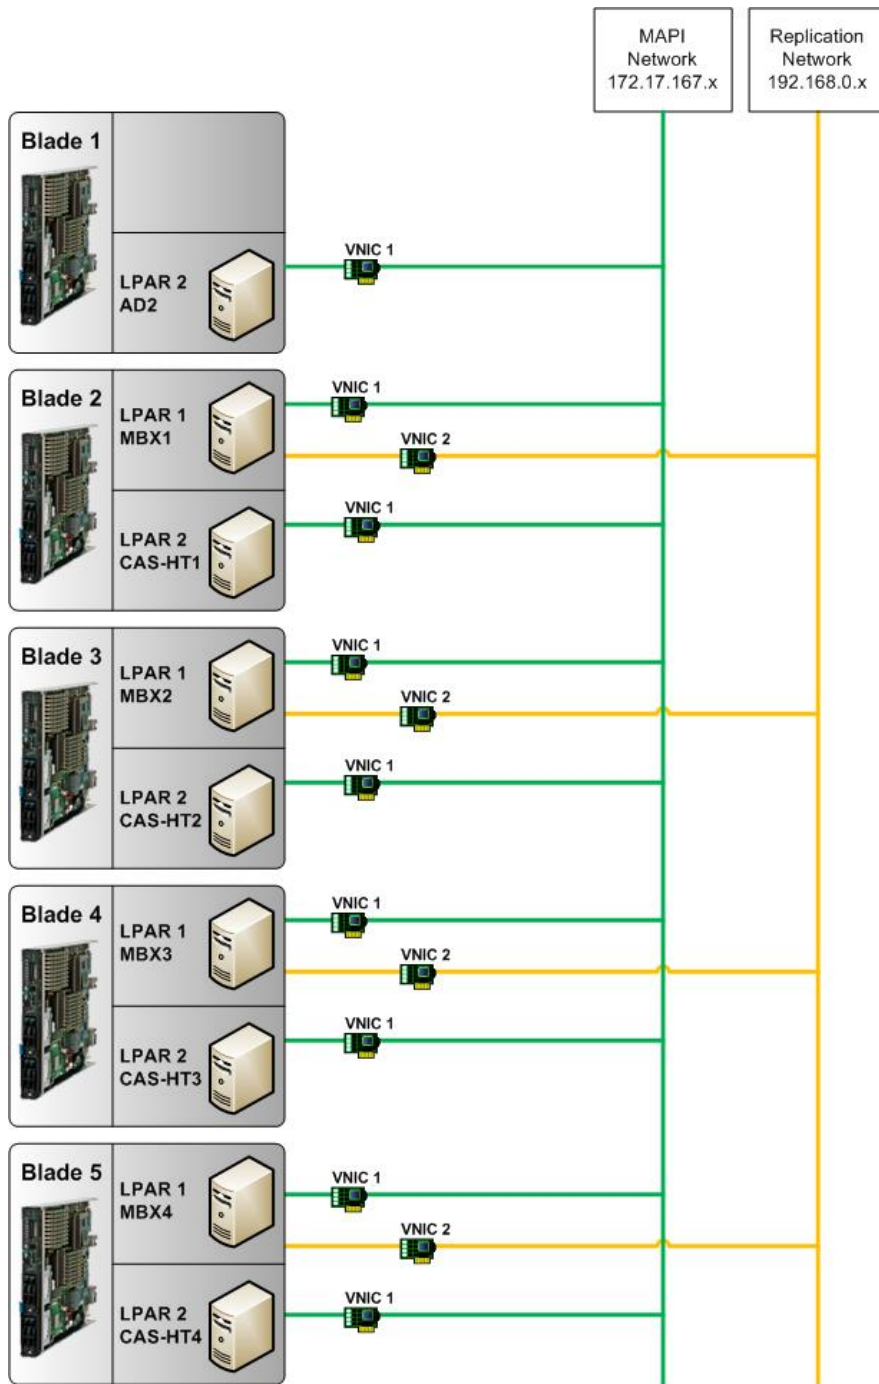


Figure 6

Storage Architecture

To satisfy 20,000 users needing 1GB of mailbox capacity and an I/O profile of 0.12 IOPS, this reference architecture uses one dynamic tiering pool and two dynamic provisioning pools.

- Pool 10 is a dynamic tiering pool that contains 12 active Exchange databases and 12 passive databases.
- Pool 11 is a dynamic provisioning pool that contains 12 active Exchange logs and 12 passive logs.
- Pool 00 is a dynamic provisioning pool that contains 12 LUNs for the OS boot volumes for the logical partitions.

Hitachi Data Systems recommends keeping the database and log on separate pools for performance reasons.

Table 9 shows configurations for each of the dynamic tiering or dynamic provisioning pools used in the Hitachi Data Systems lab.

Table 9. Pool Configuration

<i>Dynamic Provisioning Pool</i>	<i>Multi-Tier Pool (Dynamic Tiering)</i>	<i>Drive Type</i>	<i>RAID Configuration</i>	<i>Tier</i>	<i>Number of Drives</i>	<i>Drive Capacity (GB)</i>	<i>Usable Pool Capacity (TB)</i>
Pool 10	Enabled	SAS	RAID-5 (7D+1P)	1	96	300	22.01
		SATA	RAID-5 (7D+1P)	2	40	2,000	62.67
Pool 11	Disabled	SAS	RAID-5 (7D+1P)	N/A	24	300	5.50
Pool 00	Disabled	SATA	RAID-5 (3D+1P)	N/A	4	2,000	5.37

Table 10 shows the LUN allocation for each of the dynamic tiering and dynamic provisioning pools used in the Hitachi Data Systems lab.

Table 10. LUN Allocation

<i>LUN Allocation</i>	<i>Pool</i>	<i>LDEV Size (GB)</i>	<i>LDEVs</i>	<i>Storage Ports</i>
Database LUNs	Pool 10	2,048	10:01,10:02,10:03,10:04,10:05,10:06,10:07,10:08,10:09,10:0A,10:0B,10:0C 12:01,12:02,12:03,12:04,12:05,12:06,12:07,12:08,12:09,12:0A,12:0B,12:0C	CL1A,CL2A,CL5A,CL6A
Log LUNs	Pool 11	200	11:01,11:02,11:03,11:04,11:05,11:06,11:07,11:08,11:09,11:0A,11:0B,11:0C 13:01,13:02,13:03,13:04,13:05,13:06,13:07,13:08,13:09,13:0A,13:0B,13:0C	CL1A,CL2A,CL5A,CL6A
OS boot LUNs	Pool 00	127	05:06,05:07,05:08,05:09,05:10,05:11,05:12,05:13	CL3B,CL8B

Determining I/O Requirements

When designing the storage architecture for Microsoft Exchange 2010, always start by calculating the I/O requirements. Determining how many IOPS each mailbox needs is known also as determining the I/O profile. Microsoft has guidelines and tools available to help you determine this number.

Two factors are used to estimate the I/O profile:

- The number of messages a user sends and receives per day.
- The total amount of database cache available to the mailbox.

The Extensible Storage Engine (ESE) uses the database cache (located on the mailbox server) to reduce I/O operations. Generally, more cache means less I/O operations eventually hitting the storage system. Table 11 lists Microsoft's guidelines.

Table 11. Estimated IOPS per Mailbox

<i>Messages Sent and Received per Mailbox per Day</i>	<i>Database Cache per Mailbox (MB)</i>	<i>Standalone Estimated IOPS per Mailbox</i>	<i>Mailbox Resiliency Estimated IOPS per Mailbox</i>
50	3	0.06	0.05
100	6	0.12	0.10
150	9	0.18	0.15
200	12	0.24	0.20
250	15	0.30	0.25
300	18	0.36	0.30
350	21	0.42	0.35
400	24	0.48	0.40
450	27	0.54	0.45
500	30	0.60	0.50

This reference architecture uses an I/O profile of 100 messages a day or 0.1 IOPS per mailbox. To ensure that the architecture can provide sufficient overhead for periods of extremely high workload, Hitachi adds 20 percent overhead for testing scenarios, for a total of 0.12 IOPS.

To calculate the total number of IOPS or transactions per second (TPS) for an Exchange environment use the following formula:

$$\text{Required host IOPS (TPS)} = (\# \text{ of users}) \times (\text{estimated IOPS per mailbox})$$

For example:

$$2,400 \text{ IOPS} = 20,000 \text{ users} \times 0.12 \text{ IOPS}$$

This means that each of the four mailbox servers in this reference architecture must be able to support 600 IOPS.

This calculation provides the number of application IOPS required by the host to service the environment, but it does not calculate the exact number of physical IOPS required on the storage side. Perform additional calculations to factor in the read-write ratio used by Microsoft Exchange Server 2010 and the write penalty incurred by the various types of RAID levels. To calculate how many read and write IOPS are required, see the Microsoft TechNet article "[Understanding Database and Log Performance Factors.](#)"

The transaction logs in Microsoft Exchange Server 2010 require approximately 10 percent as many I/O operations as the databases. For mailbox databases in a resiliency scenario, the log write I/O is 50 percent of the amount of database write I/O. For log read I/O in a mailbox resiliency situation, apply a 10 percent overhead to account for the use of continuous replication. For example, if the databases on a server require 600 I/Os, the logs require 60 I/Os. Of those 60 I/Os, 30 are write I/O and 30 plus 10 percent (or 33) are read I/O. After you calculate the transactional log I/O, Microsoft recommends adding another 20 percent overhead to ensure adequate capacity for busier-than-normal periods.

Determining Capacity Requirements

In addition to the mailbox quota requirement, also consider the size of the database dumpster and the amount of white space the database is likely to have. The database always has free pages or white space spread throughout. During online maintenance, removing items marked for removal from the database frees those pages. Estimate the amount of white space in the database by knowing the amount of megabytes of mail sent and received by those users with mailboxes inside the database.

Each database also has a dumpster that stores items deleted from a user's mailbox. By default, most items are stored for 14 days, with calendar items stored for 120 days. In addition to the dumpster, Microsoft Exchange Server 2010 includes the single item recovery feature. Disabled by default, this feature prevents the purging of data before the deleted item retention window has passed. If enabled, this feature increases the size of the mailbox for a two-week period, which must be considered when determining the capacity requirements.

Consider what content indexing adds to the total overhead. This allows making a quick and easy search of mail items. This contributes about another 10 percent to the total database size.

The Microsoft Exchange personal archive feature was not included in this reference architecture. Using personal archives eliminate the need for .pst files by keeping archive data within the Exchange database rather than on their local computers. This allows users and administrators to do these things:

- Manage this data with features like retention policies and legal holds
- Recall the data easily with multi-mailbox search
- Protect this data by native Microsoft Exchange processes

If you plan to use personal archives, increase the storage quota needed by each mailbox for the additional data. The creation of personal archive mailboxes can be done at any time. Before SP1, the archive mailbox resided on the same mailbox database as the user's primary mailbox. SP1 allows storage of the archive mailbox in a different database. By default, the archive quotas are set to unlimited. Configure special quotas for the archive at creation to manage the growth of the archive data.

If you plan to use a recovery database in your environment, allocate sufficient capacity for all the databases you plan to recover per server.

The transaction log files maintain a record of every transaction and operation performed by the Microsoft Exchange 2010 database engine. It writes transactions first to the log and then to the database. The message size and I/O profile (based on the number of messages per mailbox per day) can help estimate the number of transaction logs generated per day. Table 12 provides guidelines for estimating how many transaction logs are generated for a 75KB average message size.

Table 12. Number of Transaction Logs Generated per I/O Profile for 75K Average Message

<i>I/O Profile</i>	<i>Transaction Logs Generated per Day</i>
50	10
100	20
150	30
200	40
250	50
300	60
350	70
400	80
450	90
500	100

As message size increases, the number of logs generated per day grows. Read [“Understanding Mailbox Database and Log Capacity Factors”](#) on Microsoft *TechNet* to find out how to prepare for this.

Consider these additional factors when determining transaction log capacity:

- Backup and restore factors
- Move mailbox operations
- Log growth overhead
- High availability factors

If you plan to include lag copies in your Microsoft Exchange environment, determine the capacity for both the database copy and the logs. The log capacity requirements depend on the delay, and usually require more capacity than the non-lagged copy.

The amount of space required for logs also depends on your backup methodology and how often logs are truncated.

For more information about calculating capacity requirements, see the Microsoft TechNet article [“Understanding Mailbox Database and Log Capacity Factors.”](#)

Download the [Microsoft Exchange 2010 Mailbox Server Role Requirements Calculator](#) from the Microsoft Exchange Team Blog.

Dynamic Tiering Pool Design

Hitachi Dynamic Tiering software allows you to combine different types of drives into a single storage pool to achieve same level of performance while reducing cost by using fewer amounts of drives. This reference architecture shows you how to combine SAS and SATA drives into one single dynamic tiering pool with capacity of 48TB to host 12 active and 12 passive databases for 20,000 users.

Selecting Drive Types

A Hitachi Virtual Storage Platform can contain SSD, SAS and SATA drives. When creating a dynamic tiering pool one of the first things to consider is the selection of appropriate drive types.

For Exchange 2010 environment, Hitachi Data System recommends not using SSD for the following reasons:

- **Required IOPS can be achieved with SAS** — For a user profile with 0.12 IOPS and 1GB mailbox size, Hitachi Data System recommends using SAS drives in a Hitachi Dynamic Provisioning environment. The IOPS achieved on Microsoft Exchange Jetstress 2010 test result from this environment was more than twice the target IOPS. The additional performance of SSD drives is not needed.
- **SSD drives might not be cost effective** — Exchange environments need a large amount of capacity for databases. It is not cost effective to add SSD drives into a data pool in most Microsoft Exchange environments.

Consider following benefits that SATA drives can provide in Microsoft Exchange environment:

- **SATA drives are relatively inexpensive** — When compared to SSD and SAS, the cost of SATA drives is lower.
- **SATA drives provide high capacity** — SATA drives have a very high capacity (up to 2TB), when compared to SSD (up to 400GB) and SAS (up to 600GB).

With more than enough IOPS achieved from a SAS-only Hitachi Dynamic Provisioning pool, adding SATA drives into a dynamic tiering pool can be a cost effective solution.

Determine Capacity for each Tier

On a Hitachi Dynamic Provisioning pool, you can enable the Multi-Tier Pool feature and then expand the pool using different types of drives. For this reference architecture, the drive count of 96 SAS drives (tier 1) and 40 SATA drives (tier 2) was used on a single pool for the following reasons:

- The IOPS achieved from SAS-only pool was substantially more than two times the target IOPS. The target IOPS can be achieved with half the number of SAS drives.
- Half of the database LUNs were for passive databases. Passive databases have very few disk I/Os compared to the active databases. The active databases can reside on tier 1 (SAS) and passive databases can be on tier 2 (SATA). There were 96 SAS drives used to accommodate the active databases.

- When there is a mailbox server failure, a DAG failover occurs, and the passive databases associated with the active databases on the failed server will become active temporarily. To meet the 600 IOPS requirement for a single mailbox server failover, 40 SATA drives were added to the pool. With 40 SATA drives on RAID 5 (7D+1P), over 800 IOPS were achieved for the Jetstress workload.
- When a DAG database failover occurs, the disk I/O of the failed over databases increases and the Hitachi Dynamic Tiering software starts to move this data into the SAS tier.

The decision about the tier drive ratio depends on the I/O and capacity requirements of your Exchange environment. In general, Hitachi Data System recommends the following:

- **Migrate active databases first** — Hitachi Dynamic Tiering fills higher tiers first. If you are migrating databases into a dynamic tiering pool, migrate active databases first to ensure that they reside on the tier 1 (SAS tier) of the pool.
- **Set enough Cycle Time** — It takes some time for the Hitachi Dynamic Tiering software to migrate data into the appropriate tier. If you have a large pool with a large amount of SATA capacity, as the pool used for this reference architecture, set the cycle time to at least 4 hours. This allows enough time for the monitoring and reallocation cycles. Manual monitoring and reallocation is recommended for the initial pool migration or a major pool expansion, also.

Compare Dynamic Tiering with Dynamic Provisioning

In this reference architecture, the total drive count was reduced by 37 percent from 216 (SAS) to 136 (SAS+SATA) by utilizing Hitachi Dynamic Tiering and still met the required IOPS. Table 12 shows the drive allocation used on the dynamic tiering pool and the dynamic provisioning pool.

Table 12. Dynamic Tiering Pool vs. Dynamic Provisioning Pool

<i>Pool</i>	<i>Drive Type</i>	<i>Drive Capacity (GB)</i>	<i>RAID Configuration</i>	<i>Tier</i>	<i>Number of Drives</i>	
Dynamic Tiering	SAS	300	RAID-5 (7D+1P)	1	96	136
	SATA	2,000	RAID-5 (7D+1P)	2	40	
Dynamic Provisioning	SAS	300	RAID-5 (7D+1P)	N/A	216	

Engineering Validation

The following describes the testing performed to validate the Microsoft Exchange solution documented in this reference architecture.

Microsoft Exchange Jetstress 2010 Test

Microsoft Exchange Jetstress 2010 is used to verify the performance and stability of a disk subsystem prior to putting a Microsoft Exchange 2010 server into production. It helps verify disk performance by simulating an Exchange database and log file I/O loads. It uses Performance Monitor, Event Viewer, and ESEUTIL to ensure that the storage system meets or exceeds your performance criteria. Jetstress generates I/O based on Microsoft's estimated IOPS per mailbox user profiles.

The test was performed on mailbox servers installed on four logical partitions on a Hitachi Compute Blade 2000 with three active databases on each server, for a total of 12 databases for a 24-hour period.

The goal was to verify that the storage was able to handle the I/O load for a long period. Table 13 has the Jetstress parameters used in the testing.

Table 13. Jetstress Test Parameters

<i>Parameter</i>	<i>Value</i>
Number of databases	12
User profile	0.12
Number of users per database	1667
Total number of users	20,000
Mailbox size	1GB

Two sets of tests were performed against a single Hitachi Dynamic Tiering pool:

- **Test Case 1** — This test was against 12 active databases. For this test, most of the databases resided on the SAS tier.
- **Test Case 2** — This test simulated one mailbox server failure (MBX1) and a DAG database failover scenario. On this test, three additional passive databases were created in the pool. The majority of these databases resided on the SATA tier. The test was run on three mailbox servers with an additional passive database attached to each server.

A 24-hour test was run concurrently against the 12 database instances. All latency and achieved IOPS results met Microsoft requirements and all tests passed without errors.

Test Case 1 — Individual Server Metrics

The following tables give the results for Test Case 1.

Table 14. Transactional I/O Performance (MBX1)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
1	7.978	0.958	322.762	190.718	0.386	111.172
5	6.570	1.355	321.304	190.696	0.394	110.168
9	8.442	1.765	322.326	190.500	0.386	111.345

Table 15. Transactional I/O Performance (MBX2)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
2	7.782	1.001	339.242	198.553	0.397	114.833
6	6.888	1.446	339.204	199.053	0.400	114.598
10	7.559	1.906	339.667	199.012	0.394	114.751

Table 16. Transactional I/O Performance (MBX3)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
3	9.601	0.956	299.292	176.014	0.398	103.437
7	7.329	1.333	299.029	177.028	0.395	102.486
11	7.290	1.745	298.933	177.073	0.398	102.529

Table 17. Transactional I/O Performance (MBX4)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
4	7.697	0.968	314.428	185.032	0.399	107.442
8	7.044	1.326	314.302	185.338	0.402	107.123
12	8.781	1.712	315.875	185.491	0.396	108.113

Table 18. Database Sizing and Throughput for 3 Database Files (MBX1)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	1538.306
Target transactional I/O per second	600
Initial database size (bytes)	5268491665408
Final database size (bytes)	5303765762048

Table 19. Database Sizing and Throughput for 3 Database files (MBX2)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	1614.731
Target transactional I/O per second	600
Initial database size (bytes)	5310854135808
Final database size (bytes)	5347327803392

Table 20. Database Sizing and Throughput for 3 Database files (MBX3)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	1427.37
Target transactional I/O per second	600
Initial database size (bytes)	5307221868544
Final database size (bytes)	5339585118208

Table 21. Database Sizing and Throughput for 3 Database files (MBX4)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	1500.466
Target transactional I/O per second	600
Initial database size (bytes)	5311609110528
Final database size (bytes)	5345574584320

Test Case 2 — Individual Server Metrics

For this test, MBX1 was failed and the passive databases for the active databases on MBX1 became active on the other three servers. The Hitachi Dynamic Tiering monitoring and reallocation was stopped during the test to keep the passive databases from migrating from the SATA tier to the SAS tier when they became the active databases.

If the Cycle Time was set to 4 hours, the data on the failed over databases that initially resided on the SATA tier would be gradually moved over to SAS tier. Allowing the migration to take place would have improved the result of the test from the results shown below.

Table 22. Transactional I/O Performance (MBX2)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
2	9.500	0.898	151.271	88.895	0.350	73.460
6	8.521	1.346	151.408	89.019	0.353	73.168
10	8.302	1.752	151.355	88.973	0.385	72.892
1	17.675	2.100	151.044	88.596	0.334	74.894

Table 23. Transactional I/O Performance (MBX3)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
3	10.515	0.853	141.786	83.482	0.355	69.320
7	8.055	1.206	141.917	83.617	0.386	68.661
11	8.558	1.624	142.170	83.763	0.346	69.023
5	18.530	1.920	142.183	83.579	0.336	70.822

Table 24. Transactional I/O Performance (MBX4)

<i>Database</i>	<i>I/O Database Reads Average Latency (ms)</i>	<i>I/O Database Writes Average Latency (ms)</i>	<i>I/O Database Reads/sec</i>	<i>I/O Database Writes/sec</i>	<i>I/O Log Writes Average Latency (ms)</i>	<i>I/O Log Writes/sec</i>
4	6.814	1.049	277.006	168.064	0.377	97.833
8	6.444	1.444	276.573	168.025	0.381	97.485
12	7.507	1.776	275.842	167.227	0.386	97.852
9	10.121	2.132	277.533	165.562	0.384	98.824

Table 25. Database Sizing and Throughput for 3 Database files (MBX2)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	960.561
Target transactional I/O per second	800.04
Initial database size (bytes)	7043977314304
Final database size (bytes)	7076525113344

Table 26. Database Sizing and Throughput for 3 Database files (MBX3)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	902.497
Target transactional I/O per second	800.04
Initial database size (bytes)	7041393623040
Final database size (bytes)	7072028819456

Table 27. Database Sizing and Throughput for 3 Database files (MBX4)

<i>Metric</i>	<i>Result</i>
Achieved transactional I/O per second	1775.83
Target transactional I/O per second	800.04
Initial database size (bytes)	7035437711360
Final database size (bytes)	7076399284224

Microsoft Exchange Load Generator 2010 Test

Microsoft Exchange Load Generator 2010 (Exchange Loadgen) is a pre-deployment validation and stress-testing tool that introduces various types of workloads into a test (non-production) Microsoft Exchange messaging system. This lets you simulate the delivery of multiple MAPI client messaging requests to a Microsoft Exchange server.

To simulate the delivery of these messaging requests, run Exchange Loadgen tests on client computers. These tests send multiple messaging requests to the Microsoft Exchange server, which causes a mail load.

The simulation was set for an eight-hour test run with 12 databases, 4 mailbox servers, and an Outlook 2007 Online Mode profile. The performance counters shown in Table 28 were monitored on each mailbox server during the test. All tests passed without errors.

Table 28 lists the average values of performance counters monitored during the test.

Table 28. Mailbox Servers Loadgen Test Results

<i>Performance Monitor Counter</i>	<i>Criteria</i>	<i>MBX1</i>	<i>MBX2</i>	<i>MBX3</i>	<i>MBX4</i>
Processor (Total) Percentage Processor Time	Avarage < 75%	36	19	48	26
MS Exchange Database: I/O Database Reads (Attached) Average Latency	Avarage < 20 ms.	7	6	7	7
MS Exchange Database: I/O Database Writes (Attached) Average Latency	Avarage < 20 ms.	0.766	0.703	0.671	0.660
MS Exchange Database: I/O Log Writes Average Latency	Avarage < 10 ms.	0.463	0.311	0.351	0.329

Conclusion

Testing confirms that the reference architecture delivers the IOPS and capacity requirements needed to support the active and replicated databases for 20,000 Microsoft Exchange mailboxes configured with the specified user profile, while maintaining additional headroom to support peak throughput.

This reference architecture guide documents a solution using the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering to reduce cost by using fewer drives while still achieving the required IOPS. The Microsoft Exchange Server 2010 resiliency solution presented in this paper supports 20,000 users with a 0.12 IOPS per user profile and a user mailbox size of 1GB. There was a single DAG configured with four mailbox servers installed on the logical partitions of a Hitachi Compute Blade 2000.

The solution outlined in this document does not include data protection components, such as VSS snapshot, or clone backups. To provide high availability and protection from logical and physical failures, the architecture relies on the built-in mailbox resiliency features of Microsoft Exchange Server 2010 coupled with the RAID technology of the Hitachi Virtual Storage Platform. Adding protection requirements can affect performance and capacity requirements of the underlying storage configuration. When designing your solution, factor these needs into your storage design accordingly.

Many factors affect the storage design for a Microsoft Exchange 2010 environment. It is critical that you validate the storage design for your environment before deployment.

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

Live and recorded product demonstrations are available for many Hitachi Data Services products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the [Hitachi Data Systems Corporate Resources web site](#). Click the **Product Demos** tab for a list of available recorded demonstrations.

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

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

 **Hitachi Data Systems Corporation**

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

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

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-087-00 May 2011

Corporate Headquarters

750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

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