

Planning for Microsoft® Exchange Server 2007 Deployments on the Hitachi Adaptable Modular Storage 2000 Family

Reference Architecture and Best Practices Guide

By Rick Andersen, Patricia Brailey, Steven Burns (MCSE 2003 Messaging and Security, MCITP Exchange 2007) and Larry Meese

December 2008



Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The authors would like to recognize and sincerely thank the following contributors and reviewers of this document:

- Mark Adams, Product Marketing
- Heidi Biggar, Application Solutions
- Steve Burr, Global Solution Services
- Dave Campbell, Product Management
- Jeff Maaks, PragmaticPS
- John Montgomery, Application Solutions
- Lisa Pampuch, Application Solutions
- Gil Rangel, Technical Operations

For best results use Acrobat Reader 8.0.



Executive Summary

Proper deployment of Microsoft® Exchange Server 2007 can be difficult. Without appropriate planning, organizations can run into a variety of performance, availability, scalability and data protection issues. Storage plays a key role in any Exchange deployment. Particular attention must be paid to the type of storage that is implemented to support the environment and, importantly, how that storage is ultimately leveraged and deployed to ensure optimal efficiency, protection and scale.

This Reference Architecture and Best Practices Guide details how organizations can deploy Hitachi Adaptable Modular Storage 2000 family systems in Exchange Server 2007 environments. The 2000 family's enterprise-class features are well-matched to Exchange environments. The 2000 family simplifies the disk provisioning and management processes of these environments as well as delivers the necessary performance to support even the most demanding transactional and non-transactional I/O workloads.


This guide is intended for Exchange administrators, storage administrators and IT managers tasked with rolling out or supporting Exchange 2007 deployments. The solutions described in this guide are based on testing performed using both Windows Server 2003 and Windows Server 2008, Exchange Server 2007 Service Pack 1 and Adaptable Modular Storage 2000 family storage systems.

Contents

Planning for Exchange Solution Deployment	2
Storage Design for Exchange Solution Deployment	2
Storage Solution Planning	3
Determining Exchange Server 2007 IOPS	4
Database Capacity Considerations	8
Transaction Log Capacity Considerations	9
Mailbox Server Processor Considerations	10
Mailbox Server Memory Considerations	10
Planning for Scalability	12
Adding Exchange Servers	12
Adding Storage Groups	12
Scaling Mailbox Quotas	13
Planning for Protection	14
Traditional Host-based Backup Software	14
Host-based Replication Software	14
Storage-based Backup and Replication	15
Designing for Performance and Protection	16
Host Considerations	16
Virtualizing Exchange 2007 with Hyper-V	17
Multipath Considerations	19
Storage Platform Guidelines	19
Storage Group Design	19
Checkpoint Depth	19
Transaction Log I/O	19
Exchange Database I/O	20
Edge and Hub Transport Server Database Design	20
Solution Architecture and Components	20
Hardware Components	23
Software and Service Components	24
Hitachi Data Systems Best Practices	26
Storage Availability Recommendations	26
Storage Design Recommendations	26
Microsoft Best Practices	27
LU Partition Allocation Size	27
Aligning Disk Tracks in Windows Server 2003 Environments	27
Summary	28



Appendix A: Example Deployment Planning Considerations.....	29
Planning Prerequisites	29
User Classification	30
Storage Design	30
Storage Configuration	31
Performance Monitoring	31
Appendix B: Additional Reference Material	32
Hitachi Data Systems.....	32
Microsoft.....	32
Appendix C: Exchange Server 2007 Tools	33
Exchange Server 2007 Storage Calculator	33
Exchange Server Jetstress Tool	33
Exchange Load Generator Tool	33
Exchange Server Stress and Performance Tool	34
Exchange Server Profile Analyzer.....	34
Exchange Server 2007 Best Practices Analyzer	34
Exchange Server Toolbox	34



Planning for Microsoft® Exchange Server 2007 on the Hitachi Adaptable Modular Storage 2000 Family

Reference Architecture and Best Practices Guide

By Rick Andersen, Patricia Brailey, Steven Burns (MCSE 2003 Messaging and Security, MCITP Exchange 2007) and Larry Meese

Proper deployment of Microsoft® Exchange Server 2007 can be difficult. Without appropriate planning, organizations can run into a variety of performance, availability, scalability and data protection issues. At the extreme, the lack of planning upfront can affect an organization's overall business goals. At minimum, it can affect an organization's ability to meet established service levels, recover data in an outage situation or scale to meet future capacity requirements.

Storage plays a key role in any Exchange deployment. Particular attention must be paid to the type of storage that is implemented to support the environment and, importantly, how that storage is ultimately leveraged and deployed to ensure optimal efficiency, protection and scale.

This Reference Architecture and Best Practices Guide is designed to help eliminate some of the guesswork around deploying storage for Exchange 2007. Specifically, the guide details how organizations can deploy Hitachi Adaptable Modular Storage 2000 family systems in Exchange Server 2007 environments. The 2000 family's enterprise-class features are well-matched to Exchange environments. The 2000 family simplifies the disk provisioning and management processes of these environments as well as delivers the necessary performance to support even the most demanding transactional and non-transactional I/O workloads. Further, the 2000 family is priced competitively, making it not only an ideal solution from a technology standpoint to support Exchange 2007 but also a cost-effective alternative to using direct-attached storage (DAS).

This guide provides a wealth of valuable information, including key considerations and design tradeoffs that all organizations should consider when planning an Exchange 2007 roll-out, detailed architectural guidelines for deploying 2000 family storage systems in support of Exchange 2007 environments as well as advice on how to avoid common pitfalls when deploying storage in support of Exchange.

This guide is intended for Exchange administrators, storage administrators and IT managers tasked with rolling out or supporting Exchange 2007 deployments. The solutions described in this guide are based on testing performed using both Windows Server 2003 and Windows Server 2008, Exchange Server 2007 Service Pack 1 and Adaptable Modular Storage 2000 family storage systems.

Planning for Exchange Solution Deployment

Exchange Server 2007 introduces new server hardware requirements, server roles, terminologies and architectural constraints that significantly change deployment and upgrade processes. An Exchange solution deployment is also influenced by the type of deployment, the number of users and whether the storage system is dedicated to the Exchange environment or shared with other applications. The three deployment types are a new environment, a transition from Exchange Server 2000 or 2003, or a migration from Exchange 5.5 or another e-mail product such as Lotus Notes.

Microsoft uses the term *transitioning* to describe the upgrade of an existing organization's Exchange Server 2000 or 2003 to Exchange Server 2007. Microsoft's thorough documentation of the transition process is available online at Microsoft TechNet.

Migrating to Exchange Server 2007 entails implementing a completely new Exchange Server environment and then migrating mailbox data from Exchange version 5.5 or earlier or from a non-Exchange messaging system such as Lotus Notes into the Exchange 2007 environment using various tools.

In addition to the thorough Exchange Server 2007 transition and migration documentation found on TechNet, Microsoft maintains an online list of qualified partners who can assist organizations in the process.

Storage Design for Exchange Solution Deployment


Storage is a critical component of overall Exchange solution architecture. All of the solutions in this guide meet Microsoft's required response time for database read and write operations of less than 20ms and are based on Hitachi Data Systems testing for Microsoft's Exchange Solution Review Program (ESRP). At the time of this writing, the Hitachi Adaptable Modular Storage 2000 family submissions are being reviewed by Microsoft and will be published on both the Microsoft and Hitachi Data Systems Web sites.

This document defines four 2000 family reference architectures that support various numbers of users on dedicated or shared storage, as follows:

- Up to 2500 mailboxes on a shared storage system:
 - Exchange shares storage system with one or more other primary applications or workloads.
 - I/O profile and capacity utilization dictate what resources remain to support other workloads.
- 2500 to 5000 mailboxes on a shared storage system:
 - Exchange is the primary application, but additional medium and light workloads are supported.
 - I/O profile and capacity utilization dictate what resources remain to support other workloads.
- 5000 to 10,000 mailboxes on a largely dedicated storage system:
 - Exchange is the only primary application, but additional light workloads are supported.
 - I/O profile and capacity utilization dictate what resources remain to support other workloads.
- 10,000 to 40,000 mailboxes on a dedicated storage system

In all cases, the storage designs for the solutions in this guide follow Microsoft best practice recommendations. Each Exchange 2007 storage group consists of a single LU created on a dedicated RAID-1+0 2D+2D group. Each storage group contains one mailbox store database. The LUs that support Exchange logs are created on RAID-1 1D+1D groups and two log LUs are placed in each RAID group.

While the testing validates the ultimate scalability and performance of the 2000 family, it does not fully account for the additional requirements that can influence response times in a full production environment.



Many other variables that can directly or indirectly affect the storage design must be taken into account when planning an Exchange 2007 environment and deployment. Consider IOPS (performance) requirements first and then storage requirements. Other variables to consider include the following:

- **Active Directory (AD) design** — Exchange 2007 is tightly integrated with AD and uses AD sites to define e-mail routing domains. In an environment consisting of multiple AD sites, the total required storage might be divided between multiple physical locations.
- **Dispersed versus centralized environments** — A dispersed environment might have different disaster recovery requirements from those of a centralized environment. Multiple copies of data might be stored in different locations.
- **Distribution of user profiles** — Different groups of users within an environment can widely vary in their usage profiles. It is important to understand this distribution prior to designing an Exchange environment.
- **Required storage per user** — In certain environments, different mailbox quotas might be defined for different class of users. It is best to group users with identical mailbox quotas into a storage group to reduce management complexity. This can necessitate the creation of different classes of storage groups with varying number of users. The use of large mailboxes can also increase non-transactional I/O.
- **Impact of mobile users (BlackBerry, Windows Mobile, and so on)** — Depending on the configuration of these clients, the IOPS required to support them can be much greater than for standard MAPI clients. It is important to keep this in mind when assigning users to storage groups.
- **User concurrency rate** — The percentage of total users that are connected to the Exchange environment concurrently can impact the storage design. If testing shows that your environment can handle 1000 users concurrently and your concurrency rate is 50 percent, your environment can handle 2000 users. Expected and unexpected spikes or changes in the average concurrency rate also need to be taken into account.
- **Government requirements for data retention** — Archiving requirements affect the performance and storage requirements for your environment.
- **Use of cache mode versus online mode** — The use of cache mode moves the processing demands from the server to the client. If a high percentage of your users use online mode, size your servers and storage accordingly to compensate for the higher IOPS requirements.
- **Backup, restore and replication methods** — The design of an Exchange environment must take into account the methodology used and the time allowed for backup and restore procedures. Do not allow the size of your databases to exceed the amount of data that your backup can handle within the allowed backup window.

Many other items need to be taken into account when designing an Exchange 2007 environment. This guide focuses on defining the storage performance and capacity required for an environment and configuring a 2300 to meet those requirements. For more information about planning and implementing an Exchange 2007 environment, see Appendix B.

Storage Solution Planning

Capacity, performance and availability requirements are often at odds with budget constraints and with each other when designing deployment architectures for Exchange Server 2007. Even with performance improvements, Exchange Server 2007 places significant demands on servers and on the underlying storage infrastructure.

In general, planning considerations for Exchange Server 2007 involve the following factors:

- Ensuring that the storage solution provides acceptable response times
- Ensuring that Exchange Server has enough disk capacity to store the data
- Ensuring that Exchange Server 2007 is properly protected to provide for resumption of service after a failure or corruption

Part of the storage planning process for Exchange is to determine how specific Exchange related changes to your environment affect planning for both I/O requirements and capacity requirements. Table 1 outlines the reasons for change and the effects on I/O and capacity.

Table 1. Effects of Exchange Related Changes on I/O and Capacity

Reason	I/O Impact	Capacity Requirements Impact
Increase in number of users	Yes. Use standard formulas to calculate.	Yes. More mailboxes are required. Use standard formulas to calculate.
Change in concurrency rate (percent of users connected at same time)	Yes. Multiply concurrency rate percentage times total I/O load calculated.	No.
Change in usage profile for all or a portion of total users (includes mobile devices)	Calculate I/O requirements using formulas below.	No.
Change in mailbox quota (maximum allowable size) for all or a portion of users	Yes. See non-transactional I/O below.	Yes. Calculate required change using standard formulas.
Non-transactional I/Os caused by background operations such as indexing and online maintenance	Yes. This can increase or decrease depending on total amount of data stored.	No. Note, however, that changes to mailbox sizes (quotas) can affect the amount of non-transactional I/O.

Determining Exchange Server 2007 IOPS

Microsoft uses non-standard terminology when describing Exchange workloads. Microsoft uses the term *IOPS* to refer to database I/O transactions per second (TPS). The physical access count to the disks varies by the RAID level in use and the expected read-write I/O mixture. For example, when reading from RAID-5 or RAID-1+0 based logical volumes, one physical disk I/O occurs per application read operation (although many reads might be issued to satisfy one apparent “user read” due to housekeeping and overhead operations in a database). However, for writes, this is a function of the RAID level in use. For RAID-1 and RAID-1+0, two physical IOPS occur per application write, but four IOPS are required for RAID-5 and six IOPS are required for RAID-6. Thus, when Microsoft reports “X number of IOPS achieved” (meaning TPS) for a configuration, far more are actually achieved due to the write component (hardware IOPS) and database housekeeping operations.

The following sections describe host- and availability-related IOPS requirements that must be calculated before determining the performance requirements for the underlying storage platform.

Database IOPS

According to Microsoft, the two largest predictors of Exchange Server 2007 database IOPS are the amount of database cache per user and the number of messages each user sends and receives per day. The following formula is based on a standard worker who uses Microsoft Office Outlook 2007 in cached Exchange mode. The predictions are only valid for user database cache sizes between 2MB and 5MB. This formula is not

validated with users sending and receiving over 150 messages per day. The average message size for formula validation is 50KB, but message size is not a primary factor for IOPS.

Database cache and estimated IOPS per user are based on user profile and message activity. Table 2 provides estimated values for IOPS per user that you can use to predict your baseline Exchange 2007 IOPS requirements.

Table 2. Estimated IOPS per User by Usage Profile

Usage Profile	Send/Receive per Day 50KB Message Size	Database Cache per User	Estimated IOPS per User
Light	5 sent/20 received	2MB	0.08
Average	10 sent/40 received	3MB	0.16
Heavy	20 sent/80 received	5MB	0.32
Very Heavy	30 sent/120 received	5MB	0.48

After you determine the total number of users and their associated usage profiles, use the following formula to calculate the host IOPS or TPS requirement:

$$\text{(# of users)} \times \text{(estimated IOPS per user from table)} = \text{required host IOPS (TPS)}$$

For this example, 480 IOPS are required to support 1000 users with Very Heavy user profiles, as illustrated by the following calculation:

$$\text{(1000 users)} \times \text{(0.48 IOPS per user)} = \text{480 IOPS}$$

This calculation provides the number of application IOPS required by the host to service this environment, but it does not calculate the exact number of physical IOPS required on the storage side. Factor in the read-write ratio used by Exchange and the penalty incurred with writes by the various types of RAID levels.

Exchange 2007 normally generates a 1:1 (or 50 percent to 50 percent) read-write ratio for the database. Sort the estimated host (application) IOPS calculated from the above formula into database read operations and database write operations, as follows:

$$\text{(host IOPS)} \times \text{(read \%)} = \text{host read operations}$$

$$\text{(host IOPS)} \times \text{(write \%)} = \text{host write operations}$$

$$\text{(480)} \times \text{(0.50)} = \text{240 host read IOPS}$$

$$\text{(480)} \times \text{(0.50)} = \text{240 host write IOPS}$$

RAID-1+0 has no penalty for reads, but each host write operation incurs two back-end storage operations. Therefore, calculate the effective number of disk operations for a RAID-1+0 solution as follows:

$$\text{240 host read IOPS} = \text{240 storage read IOPS}$$

$$\text{(240 host write IOPS)} \times \text{(2)} = \text{480 storage write IOPS}$$

When using RAID-5, where the write penalty is a factor of four, these calculations apply:

$$\text{240 host read IOPS} = \text{240 storage read IOPS}$$

$$\text{(240 host write IOPS)} \times \text{(4)} = \text{960 storage write IOPS}$$

Note: RAID-5 can be an appropriate choice when capacity requirements are high in comparison to the I/O demand. This might be the case with environments with very large mailbox quotas. If you are considering using RAID-5 for Exchange databases, test the storage environment with Jetstress before implementing into production to verify acceptable performance levels.

So, to sustain 480 host IOPS, the underlying disk group must be able to sustain 720 total random disk IOPS if used in a RAID-1+0 configuration, and 1200 random disk IOPS if used with a RAID-5 configuration, as described in Table 3.

Table 3. RAID Configurations

Configuration	Host Read IOPS	Host Write IOPS	Write Penalty	Total Disk IOPS Required
RAID-1+0	240	240	240 (X2)	720
RAID-5	240	240	240 (X4)	1200

The required storage IOPS now allow you to calculate the minimum number of disks required for a given configuration. Using the assumption that a 146GB, 15K RPM SAS drive can support a maximum of 180 random 8KB IOPS (based on average seek, which is about 50 percent capacity utilization), the minimum number of disks required for this configuration uses the following calculation:

$$(720 \text{ storage IOPS}) / (180 \text{ IOPS per disk}) = 4^\dagger \text{ disks for RAID-1+0}$$

$$(1200 \text{ storage IOPS}) / (180 \text{ IOPS per disk}) = 7^\dagger \text{ disks for RAID-5}$$

[†] Round up to the next whole number

An important conclusion to draw from this is that a solution sized with fewer than this number of disks probably does not meet the performance requirements of the environment. For this reason, Hitachi Data Systems strongly recommends that Exchange solutions be sized based primarily on performance requirements, with a secondary focus on storage capacity.

Transaction Log IOPS

The number of database reads is lower in Exchange Server 2007 than in Exchange Server 2003, due to server caching. In Exchange Server 2003, a transaction log LU for a storage group required approximately 10 percent as many I/Os as the databases in the storage group. For example, if the database LU is using 500 I/Os, the log LU uses approximately 50 I/Os.

In Exchange Server 2007, server caching, combined with the smaller log file size and the ability to have more storage groups, increases the log-to-database write ratio to approximately 3:4. This means that if the database LU is consuming 500 write I/Os, expect your log LU to consume approximately 250-375 write I/Os. After measuring or predicting the transactional log I/O, apply a 20 percent I/O overhead factor to ensure adequate headroom for busier than normal periods.

When using built-in Exchange Server 2007 replication technologies such as local continuous replication (LCR), cluster continuous replication (CCR) and standby continuous replication (SCR), the primary transaction logs must be read and then written to the passive copy. Continuous replication adds read overhead approximately equivalent to 10 percent of the write activity.

Mobile Devices

A large number of BlackBerry and Windows Mobile e-mail users can dramatically affect the Exchange I/O workload, which directly affects the IOPS requirement. For purposes of calculating IOPS, use these guidelines to calculate mobile devices' effect on I/O workloads:

- Equate one BlackBerry user to four Exchange MAPI clients.
- Equate one Windows Mobile or ActiveSync device that synchronizes frequently to the Exchange Server to four typical Exchange MAPI clients.

Non-transactional I/O

Retrieving, receiving, sending and deleting items are examples of transactional I/O. The reduction in transactional I/O makes non-transactional I/O more important. With the trend toward larger mailbox sizes and the resulting increase in the amount of data on disk, you must now consider and plan for non-transaction I/O when designing your storage configuration. Following are examples of non-transactional I/O:

- Message record management (MRM) allows users to manage their mailboxes through the use of policies. These policies normally address compliance requirements for e-mail retention and as a result e-mail might be moved or deleted. These actions require a synchronous crawl of the Exchange databases. Do not execute MRM at the same time as backup or online maintenance operations.
- Content indexing with Exchange 2007 increases IOPS and size of the Exchange database by approximately five percent. This index is placed on the same LU as the database. Messages are indexed as they are received.
- Online maintenance is the term applied to a series of operations performed by the Exchange Information Store to ensure logical consistency in the databases. Generally speaking, the operations ensure that items no longer needed by the system are removed from the database. Online maintenance activities include online defragmentation, which frees pages in the database by compacting records onto the fewest number of pages possible, thus reducing the amount of I/O necessary. By default, Exchange Server database maintenance occurs daily between 1:00 a.m. and 5:00 a.m., although these time periods can be modified. These processes increase the I/O load on Exchange databases.
- Streaming backup and restore operations can affect the I/O requirements for the storage system. The key metric for backup and restore is throughput, or the number of megabytes per second that can be copied from and to the database LUs. If these backup or restore operations are scheduled for an active production period, be sure to take this into account when calculating I/O requirements, paying particular attention to the effect this can have on the storage system.
- Streaming backup operations increase the load on mailbox servers. Schedule these backups during off-peak hours. If backups are required during peak periods, ensure that they are differential or incremental and perform full backups during off-peak hours. Consider running VSS backups from a passive copy, thus minimizing the effect of streaming backups on the mailbox servers.

- VSS verification processes require the verification of databases in a backup set. This verification causes substantial I/O to the storage system. Consider the following techniques when conducting VSS verifications:
 - Run verification during non-peak hours.
 - Offload verification processes onto separate storage systems.
- Anti-virus and anti-spam applications can increase the I/O profile on mailbox servers, especially if these are real time scans. If possible, schedule these applications to run in off peak hours.

Note: To increase the efficiency of defragmentation and backup processes, schedule your maintenance processes and backup operations to run at different times.

Data Protection Requirements

Backup and replication processes place demands Exchange Server 2007 components infrastructure and require additional storage resources regardless of the deployed technology. You can allow for the performance effect of backup and replication processes in terms of IOPS and the effect on response times in one of two ways:

- Decrease the number of mailboxes per database to accommodate the IOPS and response time requirements for the backup and replication process.
- Increase the number of physical disk spindles that comprise the RAID groups for the database LUs for a storage group to account for backup and replication overhead.

Local Cache Mode Versus Online Mode

The use of Outlook 2003 and 2007 in local cached mode can decrease the workload of the mailbox servers because the users are normally working on a local copy of their mailbox. This becomes more critical when deploying larger mailboxes with a large number of items in users' mailboxes. Cached mode lowers the read I/O burden on the server and storage system by transferring this to the client.

Database Capacity Considerations

To determine total database storage capacity required, you must first determine the mailbox size. Storage limits or quotas define the amount of data a user is allowed to store in their mailboxes, which determines the number of users that can be hosted on a particular Exchange Server 2007 mailbox server. Most organizations impose a limit on the size of a user's mailbox through quotas, which allow for more accurate and effective capacity requirements and cost predictions. If your environment has multiple mailbox sizes, use this formula for each mailbox size and add the results together. The examples in this section assume 2500 users. Use the following formula to calculate the total capacity required for the Exchange databases:

$$\text{(# of users) x (mailbox size) = required storage capacity}$$

$$\text{(2500 users) x (250MB per user) = 625GB}$$

You must consider other factors when calculating the total database size. The database always has free pages or white space spread throughout, adding about 10 percent overhead. During online maintenance, items marked for removal from the database are removed, freeing these pages. The percentage of white space is constantly changing, with the highest percentage immediately after online maintenance, and the lowest percentage right before online maintenance. Ten percent is a good average to use when planning capacity, as this white space is used by the database before extending the file size.

Each database has a dumpster that stores hard-deleted items. By default, items are stored for seven days in Exchange 2003, and 14 days in Exchange Server 2007. The dumpster includes items that are purged from

deleted items. Exchange Server 2007 increases the overhead consumed by the database dumpster because deleted items are stored for twice as long by default.

Over time, user mailboxes reach quotas requiring those users to delete, move or archive messages to free space for their mailboxes. This means that the database dumpster increases to a maximum size that is equivalent to two weeks of incoming mail. If the majority of users do not reach the mailbox quota, only some of the incoming mail is deleted, so the growth is split between the dumpster and the increase in mailbox size.

Table 4 summarizes the database size requirements factoring in the additional capacity requirements of the database dumpster plus white space.

Table 4. Mailbox Database Size Requirements

<i>Consideration</i>	<i>Required Size</i>
Mailbox quota	250MB
White space	25MB
Weekly inbox mail	10MB x 2 weeks = 20MB
Total size	295MB

The following formula calculates the total amount of capacity required for the storage groups:

$$\begin{aligned} & \text{(# of users) x (total mailbox size) = required storage capacity} \\ & \text{(2500 users) x (295MB/user) = 737GB} \end{aligned}$$

From an availability and recovery time objective perspective, smaller databases are always better, but specific sizing must be balanced with other factors, especially capacity and complexity. Larger databases take longer to back up and restore. Accordingly, these operations might not successfully complete within the allotted windows. However, deploying with many smaller databases adds complexity. With Exchange Server 2007, the maximum number of databases per server is increased from 20 to 50. If you are not using CCR, Microsoft recommends limiting the database size to 100GB. On servers that do use CCR, Microsoft recommends limiting the database size to 200GB. However, if you use Hitachi Data Protection software or storage system-based hardware replication, larger databases might be acceptable. This is due to the ability of both Hitachi Data Protection software and storage system-based hardware replication to back up and restore large databases quickly.

Additional physical disk resources might be needed, beyond those required for raw mailbox capacity, to provide acceptable response times and to conform to the required levels of availability (for example, use a higher spindle count to maximize storage performance).

Transaction Log Capacity Considerations

The transaction log files maintain a record of every transaction and operation performed by the Exchange Server 2007 database engine. All transactions are synchronously written to the log first and then written to the database using a checkpoint process. Exchange Server 2007 reduces the size of the transaction logs from 5MB to 1MB. This design change facilitates the support of features such as cluster continuous replication and minimizes the amount of data loss if a failure or corruption occurs. Consider these factors when determining transaction log capacity:

- **Number of logs generated each day per Exchange user** — A good rule of thumb for a mailbox server that hosts medium to heavy users is about 1GB of logs per 100 users or about 10GB of transaction logs per day for a 1000 users.

- **Number of days transaction logs are retained** — Full or incremental backups normally run on a daily basis and truncate several transaction logs, while differential backup types do not perform the log truncation process.
- **Storage required for mailbox moves** — When a mailbox is moved from one storage group to another, logs are generated that match the size of the mailbox being moved.

For example, use the following assumptions about each Exchange Server 2007 storage group:

- 1000 users, each with a 200MB mailbox
- 10 logs generated per day per user
- Three-day log-retention requirement
- Ability move at least 10 mailboxes in parallel

Use the following formula to calculate the capacity required for the Exchange transaction logs:

(# of users) x (logs generated per day per user) x (retention days) = required storage capacity

(1000 users) x (10MB logs per day) x (3 retention days) = 30GB

Using this example, multiply 10 mailboxes by the mailbox size of 200MB for an additional 2GB transaction log size increase to factor in mailbox moves. Add 20 percent capacity to ensure adequate headroom and to mitigate peak time.

Mailbox Server Processor Considerations

The recommended processor configuration for servers that host the mailbox role is based on mailbox count and user profile. A server with four processor cores provides a good balance between price and performance and can host several thousand mailboxes. Rule-of-thumb sizing for the mailbox role requires an understanding of the average client user profile.

Mailbox server utilization depends on the following factors:

- User activity profile (for example, 2000 average users)
- Access mode (cache versus online, Outlook, Outlook Web Access)
- Use of data replication features (for example, LCR, CCR and SCR)
- Mobile devices
- Third-party applications

A good rule of thumb is 1000 average users per CPU core.

Mailbox Server Memory Considerations

Memory sizing for the Exchange Mailbox server role is critical to reducing the I/O workload presented by the server to the storage platform. In general, increasing the amount of memory on the mailbox server results in fewer physical disk I/Os being generated by the Exchange databases. However, a point of diminishing returns exists where adding memory to the server is not justified based on the price-performance ratio. Accordingly, Microsoft recommends up to 32GB of memory on a mailbox server, based on its internal testing that showed diminishing returns above 32GB. Table 5 provides estimates of minimum memory requirements based on the number of hosted mailboxes and the profile type.

Table 5. Mailbox Server Minimum Memory Recommendations

User Type	Minimum Memory Recommendation
Light	2GB + 2MB/mailbox
Average	2GB + 3MB/mailbox
Heavy	2GB + 5MB/mailbox
Very Heavy	2GB + 5MB/mailbox

Please note that 2GB is the minimum memory requirement. Memory requirements are also dictated by the number of storage groups allocated to a mailbox server as outlined in Table 6.

Storage Group Memory Requirements

The number of storage groups defined affects the database cache used by the Extensible Storage Engine because server cache is used for both read and write I/O for each storage group. Adding a storage group increases the amount of database cache required to support the workload. This has the positive effect of reducing database I/O to storage. However, if too many storage groups are defined on a server with insufficient memory, Exchange server performance can degrade. Microsoft publishes specific guidelines for the number of storage groups and corresponding minimum memory requirements. Table 6 outlines these requirements.

Table 6. Memory Recommendation by Number of Storage Groups

Storage Group Count	Minimum Required Physical Memory
1-4	2GB
5-8	4GB
9-12	5GB
13-16	6GB
17-20	7GB
21-24	8GB
25-28	9GB
29-32	10GB
33-36	11GB
37-40	12GB
41-44	13GB
45-48	14GB
49-50	15GB

Planning for Scalability

Exchange 2007 environments can grow to accommodate demand for more mailboxes or larger mailboxes. The 2000 family has many features that make it easy to scale an Exchange 2007 environment.

Adding Exchange Servers

Additional Exchange servers are required if the number of users increases or the average user profile changes significantly.

IOPS Increase

The addition of users not only produces a need for more storage capacity, it also produces more IOPS. Depending on the specifications of the existing servers, more servers might need to be implemented to handle the increase. An increase in user workload profile only increases the number of IOPS and does not require an increase in capacity. More Exchange servers and more disks might need to be added, however, to handle the increase in load. Consider the use of the 2000 family feature that allows for online RAID group and LU expansion. With this feature on the storage system, disks can be added to a RAID group to provide additional performance and capacity.

Fibre Channel Connections

Adding more servers to the Exchange environment requires additional Fiber Channel connections. Models in the 2000 family have up to 16 Fibre Channel host ports and can handle up to 2048 host connections. Any new logical units assigned to the new servers are accessible using either controller due to the symmetric, active-active design. Any open port can be assigned to the new server. Remember to assign at least two paths, one on each controller, if multipathing is implemented.

Server Standards


Creating an organizational standard for servers that host the Mailbox role can make adding servers easier. Consider standardizing the following:

- Number of processor cores and the amount of memory per server
- Number of mailboxes per server
- Number of storage group LUs and transaction log LUs per server
- Mailbox quota

Adding Storage Groups

A standard server configuration does not mean servers can be added to the Exchange environment without regard to the effect they have on the storage system. It is important to remember that every host I/O is associated with disk I/Os. When adding servers, it is very likely new storage groups with their associated RAID groups and logical units must also be added. In most cases, it is better to add storage groups rather than enlarge existing storage groups, because enlarging storage groups entails reconfiguring the RAID group and restriping the storage group data across the enlarged LU. While this can be done on the 2000 family, it is simpler to create new RAID groups and logical units and assign them to the new servers. It might even be possible to use pre-allocated, standard-sized logical units that can be quickly assigned to the new servers using either Hitachi Storage Navigator Modular 2 or Hitachi Device Manager 6.0 software.

It is not always possible to use preallocated logical units, however, and creating new ones on an ad hoc basis is completely acceptable. Due to the switched, 3Gb/sec SAS back-end architecture and the presence of active-active controllers on the 2000 family, configuration of additional storage is simplified and more efficient.



Administrators can let Storage Navigator Modular 2 automatically place the new RAID groups and automatically assign the new logical units to alternating controllers. The 2000 family handles balancing the load to the new storage groups on the back-end and controller level after workloads are applied.

When enlarging storage groups, you must also consider the amount of time it takes to complete backup and replication processes. Larger storage groups take longer to back up or replicate, placing recovery point objectives and recovery time objectives in danger of being missed. A standard server configuration is helpful in this situation because the administrator already knows how long a backup for a given server takes. As always, test backups and replication processes before placing the environment into production.

When adding mailboxes to an Exchange environment, consider I/O requirements before capacity requirements. Use the standard formulas outlined earlier in this paper to calculate I/O effect and capacity requirements. It can be tempting to add mailboxes to existing storage groups if they have excess capacity; however, the excess capacity does not mean the storage group can handle the additional I/O load. A large increase in mailboxes likely means more storage groups and Exchange servers are needed to meet the increase in I/O.

Scaling Mailbox Quotas

Increasing mailbox quotas affects capacity utilization and non-transactional I/O and requires an overall increase in capacity allocated to Exchange 2007.

Capacity Utilization

Capacity utilization is calculated by dividing available disk capacity by used disk capacity within a RAID group. Monitor capacity utilization closely, as it can cause performance to gradually decrease as the ratio increases. Increasing the quota without changing the number of mailboxes per storage group increases capacity utilization. Adding more storage groups to handle larger quotas allows for more capacity without degrading performance.

Non-transactional I/O

Non-transactional I/O includes functions like indexing and archiving of mailboxes. Increasing mailbox quotas influences the time it takes to perform these functions Exchange needs to crawl through more data. Consider this extra load that is placed on Exchange storage groups when configuring for larger quotas. More storage groups with fewer users keep the I/O load at a manageable level. It is very important to monitor your Exchange environment regularly because non-transactional I/O is different for every environment.

Capacity Requirements

Increasing quotas has a different effect on the Exchange 2007 architecture than other components. Adding Exchange servers, storage groups and mailboxes has a more linear effect in that you can add more of the same components to accommodate the higher I/O and capacity requirements. A new quota requirement likely changes the number of mailboxes per storage group in addition to needing more storage. New calculations must be made to determine the optimal configuration. Any standard server configuration used to simplify the scaling of Exchange servers, storage groups and mailboxes might need to be updated to reflect the new mailbox quota.

Planning for Protection

Data protection operations — specifically, backup and replication — place demands on one or more components of an Exchange Server 2007 infrastructure and require additional storage resources regardless of the deployed technology. Some of the associated costs, benefits and requirements that accompany host- and storage-based backup and replication solutions are summarized in Table 7.

Table 7. Host- and Storage-based Backup and Replication Considerations

Method	Host RAM and CPU Requirements	Added Workload to Production Exchange Storage	Required Backup Window	Recovery Point Objective (RPO)	Recovery Time Objective (RTO)
Host-based backup	Medium to high	Medium to high	Hours	Hours to days	Hours to days
Host-based replication	Medium to high	Low to high	N/A*	Minutes to days	Minutes to days**
Storage-based backup and replication	Low	Low to medium	Seconds to minutes	Hours to days	Minutes to hours

* Backup not performed, but can offload backup process on the passive replica

** If restore is required

E-mail is a critical resource for conducting business and is thus a high priority for any business continuity solution.

Traditional Host-based Backup Software

Traditional host-based backup applications for Exchange Server 2007 are perfectly acceptable protection solutions for some organizations, especially for small environments with less stringent data protection requirements and smaller budgets. By nature, the streaming backup process for Exchange Server 2007 demands increased host resources during normal backup operations. Specifically, additional CPU cycles are required to perform the page checksum process. Memory and disk resources must be sufficient to support the intensive sequential workloads associated with the streaming backup.


The amount of time required to perform this type of backup (commonly known as the *backup window*) is much greater than the time required for storage-based backup methods, which prohibit this type of solution from delivering attractive recovery point objective (RPO) and recovery time objective (RTO) levels. However, mostly due to cost, tape backups are a critical business requirement and provide enhanced data protection when combined with storage-based replication.

Host-based Replication Software

Exchange Server 2007 introduces new replication and availability features that can affect overall storage design. LCR, CCR and SCR enhance the recoverability of Exchange 2007 by complementing backup solutions, not replacing them. Each technology provides varying levels of protection by providing a passive copy for rapid recovery.

Local Continuous Replication

LCR provides functionality that maintains a second copy of a storage group locally on a standalone Exchange 2007 server, which can facilitate recovery when the second copy is manually activated. In essence, when deployed using modern, high-availability storage systems from Hitachi, the value of LCR is limited to offloading the backup process to the second copy and potential recovery capabilities from some logical corruptions. Like CCR and SCR, LCR with Exchange 2007 SP1 uses the Exchange Replication Service, which requires



additional server hardware resources, specifically processor cycles and memory address space. Follow Microsoft's recommendations for additional resources for CPU, RAM and other server and storage hardware requirements.

Adding LCR on the 2000 family entails allocating additional storage group RAID groups for the passive copies. This satisfies two key storage recommendations for LCR enabled storage groups: It isolates the physical disks for the active and passive copies and ensures that the disks allocated to the passive copies have the same performance capabilities as the disks used to support the active copy. Use a separate path to the storage for the passive copy to eliminate a single point of failure.

Cluster Continuous Replication

CCR also leverages Exchange Server 2007 asynchronous log shipping technology to replicate storage groups between the active and passive nodes of a cluster. In addition, CCR is integrated into the Microsoft Cluster Service (MSCS) to facilitate Exchange service management and failover. While still resource intensive, many of the key processes and their associated performance effects occur on the passive Exchange node. The primary Exchange server does incur a resource penalty due to the overhead of managing the replication service, a network share and other components. Because log shipping technologies inherently operate at a point in time that is behind the production system, the RPO and RTO levels delivered through host-based replication, in most cases, cannot match those of a storage-based, synchronous replication solution.

Implementing CCR on the 2000 family requires allocating separate RAID groups for the passive copy and sometimes allocating RAID groups for the passive copy on an additional, often remotely located, storage system. For Exchange 2007 SP1, ensure that the CCR passive copies have the same performance and capacity capabilities as the disks on which the active copy resides. For non-SP1 environments, see Microsoft documentation because disk performance requirements are significantly different for the released-to-market (RTM) and SP1 versions of Exchange 2007.

Standby Continuous Replication with Service Pack 1


SCR is an extension of Microsoft's other two replication products, LCR and CCR, and it is only available with SP1. SCR is not meant to be used for backups. It is instead meant to be used in standby environments. In SCR, passive copies must be activated before they can be used for recovery. SCR uses the same log-shipping technology as LCR and CCR, but allows multiple targets per storage group rather than just one. It also automatically delays replay activity to intercept corrupt data before it is replicated.

Deploying SCR involves creating RAID groups for database and log LUs on one or more storage systems at remote sites. As in LCR and CCR, the passive copies must have the same performance and capacity capabilities as the disks on which the active copy resides in the event that a failover to the remote site occurs.

Storage-based Backup and Replication

Many storage-based backup and replication features are offered with the 2000 family. These offerings include but are not limited to Hitachi ShadowImage® Replication (in-system), Hitachi TrueCopy® Synchronous, Hitachi Protection Manager and Hitachi Storage Cluster. The performance effect associated with using storage-based replication technologies varies due to any number of factors. Expect to incorporate a performance allowance into the design for the increased data protection workload.

The performance allowance might need to be increased for larger, higher IOPS Exchange environments that use advanced replication architectures. Similarly, performance effects might be negligible for smaller environments using basic replication architectures. Detailed planning and design are critical for Exchange Server 2007 environments with high IOPS, aggressive RPO and RTO, demanding business continuity requirements or that otherwise possess advanced levels of complexity. Performance and business continuity



consultants from the Hitachi Data Systems Global Solution Services division can assist in implementing these complex backup and replication solutions.

Designing for Performance and Protection

Designing a high-performing and resilient solution requires allowing for the performance impact on IOPS and response time associated with backup and replication processes. Regardless of the backup or replication technology, additional workload requirements are placed on the server and storage resources housing the Exchange Server 2007 infrastructure. To ensure application performance does not degrade during times of high activity, you must incorporate the increased resource demands of backup and replication functions into the storage design.

From a storage perspective, two possible methods exist for adapting the Adaptable Modular Storage 2000 family for backup and replication:

- **Decrease the number of mailboxes allocated to each storage group for the Exchange server.** This decreases the total number of mailboxes hosted by the Exchange server, thus reducing the I/O workload to the physical storage resources associated with those storage groups.
- **Increase the number of RAID-1+0 disk spindles allocated to the database RAID groups.** Ultimately, this allocates the mailbox store file over an increased number of physical disks from the larger RAID-1+0 group, which provides the ability to support an increased number of IOPS within an acceptable average response time. In other words, increasing the number of physical disk spindles allows for backup or replication overhead. Keep in mind that the addition of physical disks to the RAID group and the restriping of data across the new RAID group represent significant configuration changes.

Because these are significant configuration changes, it is important to test them before putting them into production. Use Jetstress to collect benchmark statistics first for IOPS and response time, then make comparisons after the backup or replication process is in place:

- If you decrease the number of mailboxes per storage group, the benchmark IOPS and response time metrics are the targeted maximum levels when the effect of backup or replication processes is included.
- If you increase the number of physical spindles per storage group, the benchmark IOPS and response time become the expected minimum levels.

In both cases, ensure that a 10 to 15 percent factor is built into the calculations to account for growth and spikes in latency.

Host Considerations

Keep these host considerations in mind when planning and deploying storage for your Exchange environment:

- **Use of mount points** — Because Exchange 2007 supports 50 storage groups, it is very easy to consume all available host drive letters. Therefore, it is best to use mount points in any Exchange 2007 environment to simplify management and to allow for additional storage groups.
- **Host bus adapter (HBA) queue depth** — It is important to configure the queue depths on HBAs so that the combined queue depths for HBAs connected to a specific port on the storage system does not exceed the maximum that a single port can handle.
- **Disk alignment** — This applies to Windows Server 2003 environments only. It does not apply to Windows Server 2008 environments.

Virtualizing Exchange 2007 with Hyper-V

Hosting Exchange Server 2007 in a virtualized environment using Microsoft Hyper-V does not change the design criteria from a storage perspective. Exchange application sizing guidance for the underlying storage infrastructure is the same whether the Exchange server is being hosted on a single physical server or as a guest on a Hyper-V server. When allocating storage to an Exchange virtual guest in a Hyper-V environment, Microsoft requires the use of fixed virtual hard disks (VHD), SCSI pass-through (raw disks) or iSCSI disks.

When considering which Exchange 2007 roles to virtualize, keep in mind those roles that make good candidates for virtualization. Except for the Unified Messaging server, any role can be virtualized. Those roles that do not store any persistent data, such as the Client Access servers and the Hub Transport servers, are ideal candidates for virtualization. This means that they can be quickly deployed or replaced on another host. Take care when considering a decision to virtualize the Mailbox server role to ensure that enough resources are available to support the guest machine on the Hyper-V physical server. Appropriate planning for the Mailbox server role resources considers the following requirements:

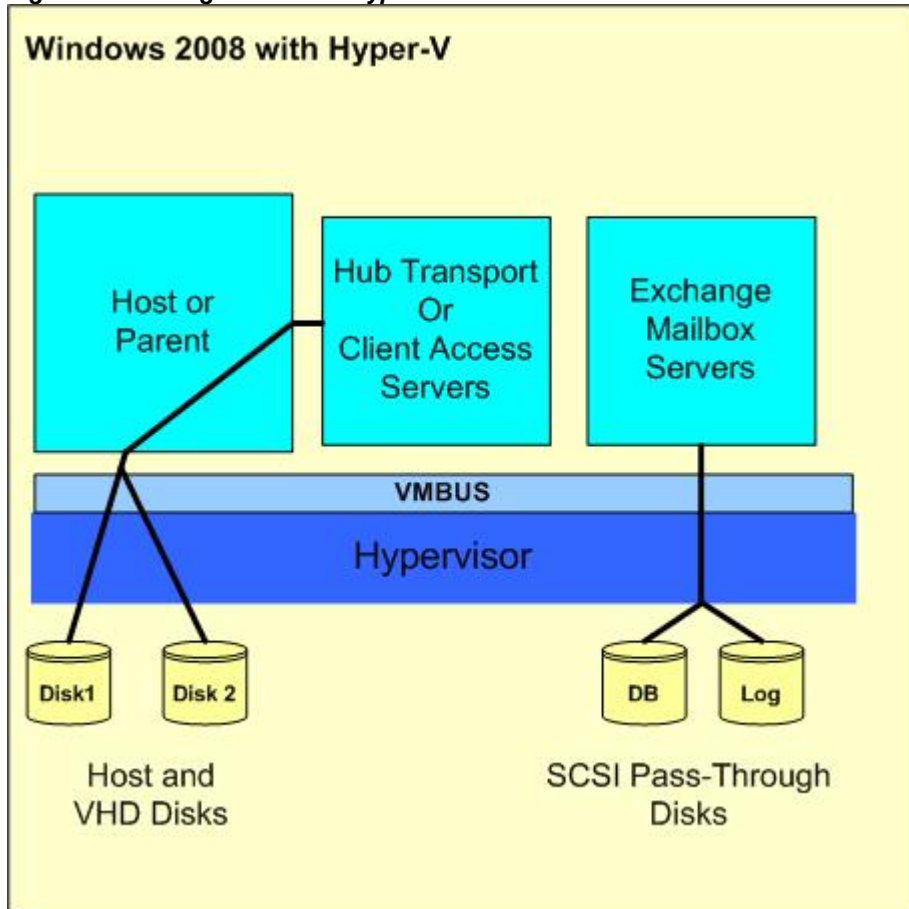
- Processor
- Memory
- I/O and storage

For roles such as Client Access and Hub Transport, either SCSI pass-through or fixed VHD disks can be deployed. When hosting the Exchange Mailbox server role in a guest virtual machine, Microsoft recommends that SCSI pass-through disks be utilized for the database and transaction logs. Although this limits the portability of the Exchange Mailbox guest virtual machine, it generally provides the best performance.

Figure 1 details one scenario for running Exchange 2007 in a Hyper-V configuration:

- Disk 1 is connected to the Hyper-V host server that runs the Hyper-V operating system.
- Disk 2 contains the VHD file that is directly connected to the host server and contains the Hub or Client Access server.
- The database and log disks utilize pass-through disks that directly expose the storage to the Exchange Mailbox server guest machine. This configuration bypasses the host's file system, allowing for improved performance.

Figure 1. Exchange 2007 with Hyper-V



Keep these considerations in mind when running Exchange 2007 as a virtual guest:

- Place storage for the virtual guest Exchange database and transaction logs on storage spindles that are separate from the virtual machines' OS VHD storage.
- Fibre Channel HBAs must be configured to the host OS and the LUs presented to the Exchange virtual guest need to be configured as either pass-through or fixed VHD disks.
- Use of dynamic disks is not supported.
- Dynamic expansion of virtual disks is not supported.
- Hypervisor snapshots of an Exchange virtual machine are not supported,
- The use of differencing disks is not supported with Exchange 2007 as a virtual guest.
- VSS-based backups cannot be taken from the Hyper-V OS, only from within the guest virtual machine.

For more information about the Microsoft support policy for running Exchange in a virtualization environment, see Microsoft TechNet article cc794548, "Microsoft Support Policies and Recommendations for Exchange Servers in Hardware Virtualization Environments."

Multipath Considerations

The 2000 family supports active-active multipath connectivity. To obtain maximum availability, design and implement your host-storage connections so that at least two unique paths exist from the host to the storage system. Hitachi recommends the use of dual SAN fabrics, multiple HBAs and host-based multipathing software when deploying Exchange Server.

Multipathing software such as Hitachi Dynamic Link Manager and Windows Server 2008 native MPIO are critical components of a highly available system. Multipathing software allows the Windows operating system to see and access multiple paths to the same LU, enabling data to travel down any available path for increased performance or continued access to data in the event of a failed path. While multiple load-balancing settings exist in both Hitachi Dynamic Link Manager and Windows Server 2008 native MPIO, the symmetrical active-active controller feature of the 2000 family enables either controller to respond to I/O regardless of the originating HBA port, without having to select a host load-balancing option. However, if the workload is large enough to consume more bandwidth than a single HBA port can handle, Hitachi recommends using the round robin load-balancing algorithm in both Hitachi Dynamic Link Manager and Windows Server 2008 native MPIO to distribute load evenly over all available HBAs.

Storage Platform Guidelines

Proper storage group and mailbox design can improve reliability and performance. The design of the storage groups and mailbox directly affects the checkpoint depth, database and transaction log I/O performance.

Storage Group Design

Exchange Server 2007 now supports a larger number of storage groups and databases per server, with up to 50 storage groups allowed per server with the Enterprise Edition. Although a storage group can contain as many as five databases, Microsoft recommends having only one database per storage group.

The previous recommendation for Exchange Server 2003 was to fill a storage group (to a maximum of five databases) before creating a new storage group. Now, with Exchange Server 2007, mailbox data can be distributed across more databases and mailbox databases can be distributed across more storage groups. By limiting a storage group to a single database, each database uses a dedicated set of log files, which can improve performance, recovery, availability and manageability.


Checkpoint Depth

The default checkpoint depth for Exchange Server 2007 is 20MB per storage group. The checkpoint depth is the maximum amount of data that can be written to the transaction logs before Exchange commits this data to the database file. This means that the larger the checkpoint depth, the less frequently Exchange is required to physically commit data from cache to the database file.

By spreading the databases across multiple storage groups, checkpoint depth can be effectively increased. For example, if you have a single storage group with four databases in the storage group, each database is sending data to one transaction log, filling the checkpoint depth of 20MB four times faster than if just one database sends log data. Hitting the checkpoint depth more frequently leads to more frequent and less optimized writes to the database, which can reduce Exchange performance.

Transaction Log I/O

It is important to design for low write latency and high availability for the storage volumes supporting transaction log files. Place log files on separate RAID-1 or RAID-1+0 groups that do not hold database files. Do not allow transaction log files (which generate sequential reads and writes) to share the same physical spindles as database files (which generate random reads and writes). This allows the 2000 family storage controller to



optimize timing of the I/O from the battery-protected cache to the disks by detecting the sequential nature of the log writes and reads. For the deployments described in this paper, two transaction log LUs were allocated on the same RAID-1 group.

Exchange Database I/O

Earlier versions of Exchange required additional physical disks simply to meet the I/O-intensive nature of the Exchange 2003 workload. With Exchange Server 2007, the larger database server cache decreases the number of reads to the database on disk, resulting in a reduced number of physical reads as a percentage of total I/O. In Exchange Server 2003, a typical database read to write ratio was about 2:1, meaning two reads occurred for each write.

For environments that follow Microsoft-recommended Exchange 2007 memory requirements, and that use Outlook running in cached mode, the read-to-write ratio is closer to 1:1, or 50 percent fewer physical reads. RAID-1+0 is used in the solution design for enhanced availability and optimized performance for random reads and writes. In some specific environments, RAID-5 can be an appropriate choice when application capacity requirements are high in comparison to the I/O demand. If you are considering using RAID-5 for Exchange databases, test the storage environment with Jetstress before implementing into production to verify acceptable performance levels.

When Outlook is running in online mode, or when using desktop search engines that do not use the Exchange Server 2007 Content Indexing Service, the read-to-write ratio increases depending on the mailbox size, typically resulting in a workload with an increased level of physical reads.

Edge and Hub Transport Server Database Design

Edge and Hub Transport server roles are new with Exchange Server 2007. These roles maintain databases that are often hosted on internal drives. However, to improve availability, place these databases on the 2000 family.

Hub Transport Server

Microsoft recommends a minimum allocation of 7GB of disk space per server per 1000 users. This provides headroom in case of an outage where the Hub Transport server cannot deliver e-mail. The Transport dumpster cannot currently be clustered or replicated. Consider booting the Hub Transport server from a LU that can be replicated on the storage system to a remote site to allow for transport dumpster recovery. This can provide recovery from a disaster on the passive node and allows for the reclaim of lost e-mail messages.

Edge Transport Server

Microsoft recommends reserving 8GB of Edge Transport database space for every 1000 users that the Edge server supports. This typically provides enough headroom for a 48-hour outage if the Edge Transport server cannot deliver e-mail.

Solution Architecture and Components


The 2000 family has new features that make it easy to configure and implement an Exchange 2007 environment. The placement of RAID groups is no longer dependent on a back-end loop structure. The switched, 3Gb/sec SAS back-end architecture in the 2000 family allows bandwidth to all disks in a tray by utilizing 16 wide links. The switched architecture also provides fault tolerance in that if one of the wide links fails the SAS controller can detect the failure and continue to use the remaining links. The ability to use either SAS or SATA drive technology in the 2000 family expansion tray offers additional flexibility.

The sample architecture described in this section allows for 5,000 Exchange users on a 2300 supporting a 250MB mailbox quota and an average workload of 0.48 IOPS. Each Exchange 2007 storage group and its

associated 260GB database LU are deployed on four dedicated RAID-1+0 2D+2D RAID groups. The transaction logs for this deployment are hosted on two shared RAID-1 1D+1D RAID groups configured with two 20GB logical units each. Because of the SAS back-end architecture, these RAID groups can technically be placed anywhere on the array. To make configuration simple, all of the RAID-10 groups were configured allowing Storage Navigator Modular 2 to automatically assign the disks, then all of the RAID-1 groups were assigned also allowing Storage Navigator Modular 2 to automatically assign the disks.

Figure 2 illustrates a configuration with four database RAID groups and two log RAID groups.

Figure 2. RAID Group Configuration



One host with 64GB of RAM and four dual core AMD Opteron processors is attached to the 2300 with the mailbox server role installed. On the 2300, the four storage group logical units are spread evenly across two units. The transaction log logical units are also spread evenly across two ports. See Table 8 for configuration details.

Table 8. Detailed Storage Configuration

RAID Group	Port Assignment	LUN	Size (GB)	RAID Level	RAID Type	Disk Spec	Description
1	0A, 1A	1	260	RAID 1+0	2+2	146GB 15K	Storage Group 1
2	0A, 1A	2	260	RAID 1+0	2+2	146GB 15K	Storage Group 2
3	0C, 1C	3	260	RAID 1+0	2+2	146GB 15K	Storage Group 3
4	0C, 1C	4	260	RAID 1+0	2+2	146GB 15K	Storage Group 4
41	0A, 1A	41	20	RAID 1	1+1	146GB 15K	SG 1 Log
41	0A, 1A	42	20	RAID 1	1+1	146GB 15K	SG 2 Log
42	0C, 1C	43	20	RAID 1	1+1	146GB 15K	SG 3 Log
42	0C, 1C	44	20	RAID 1	1+1	146GB 15K	SG 4 Log

Remember to consider factors such as the disk utilization ratio when designing the storage environment for Exchange 2007. The disk space utilization ratio is calculated by dividing available disk capacity by used disk

capacity within a RAID group. Disk space utilization shows how much headroom is available in the environment. Table 9 shows the difference in deliverable IOPS based on the space capacity used by an eight LU Exchange database configuration.

Table 9. Capacity Usage Ratio for an 8 LU Configuration

Storage Configuration	Capacity Utilization (Percent)	Exchange Total Database Size (GB)	Percent of IOPS Measured Within 20ms
8 LU	22	500	100
8 LU	44	1000	84.5
8 LU	100*	2304	77.5

* Estimated values based on physical disk performance characteristics.

Used disk capacity can grow by either increasing the mailbox quota or by increasing the number of users. Table 9 shows that an eight LU configuration can support growth in the Exchange environment while sustaining similar performance.

Figure 3 shows a simple topology for an Exchange 2007 environment of 5000 users.

Figure 3. Mailbox Server Topology

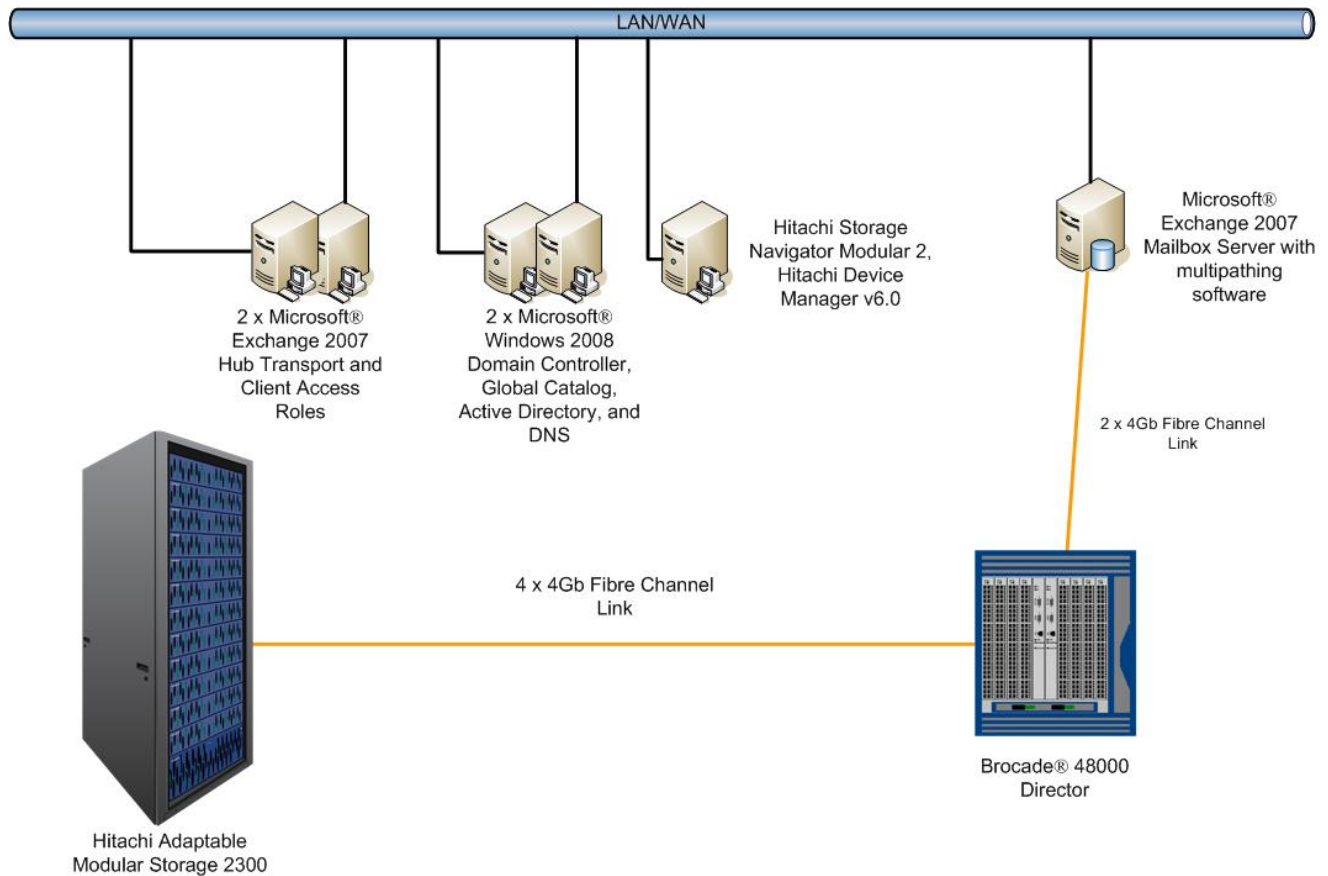
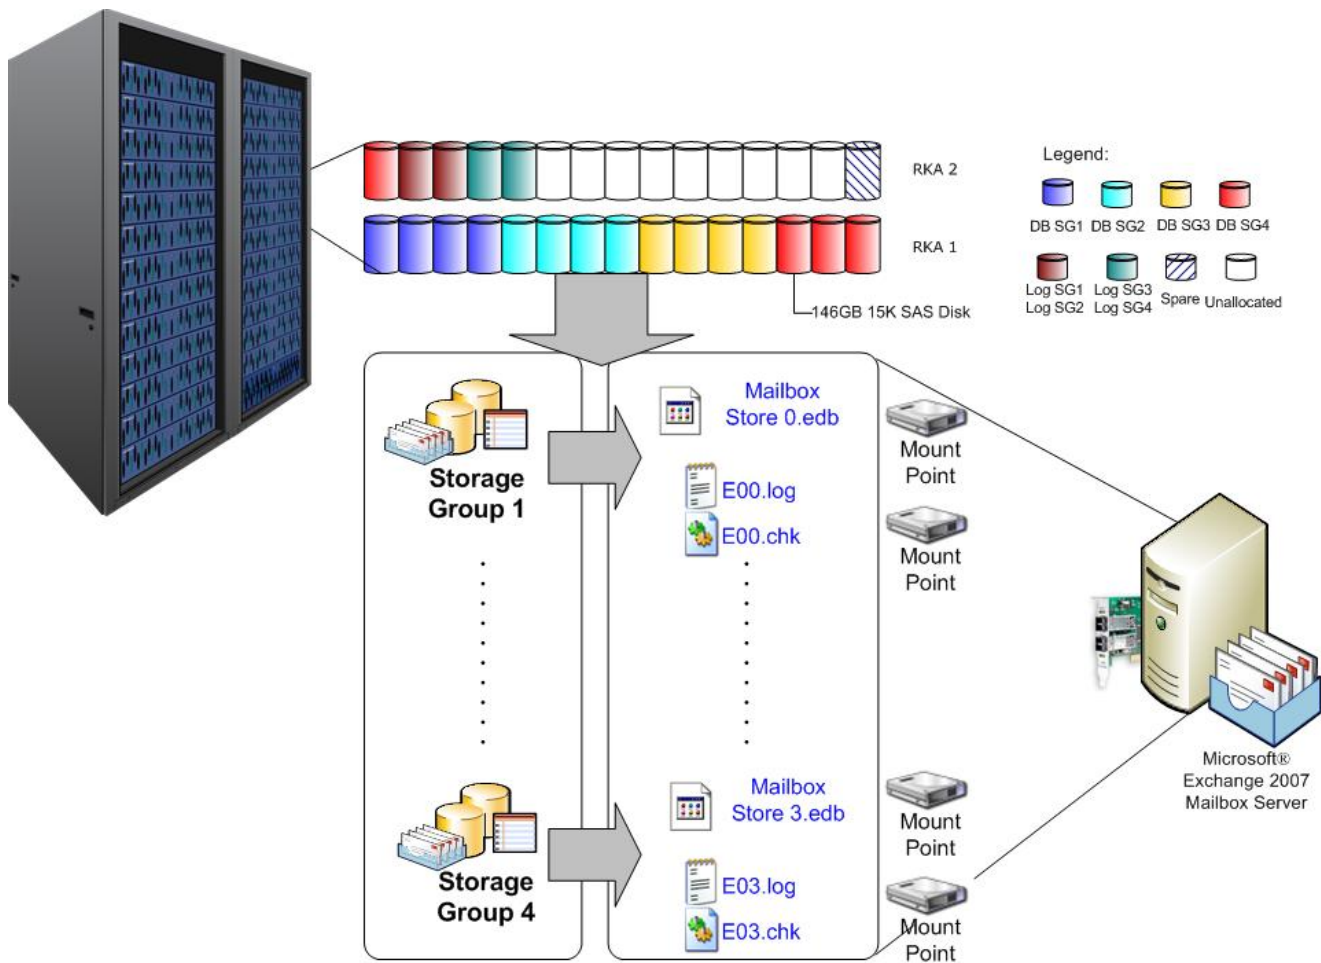


Figure 4 provides an overview of the storage configuration for a 5000 user environment.

Figure 4. Mailbox Server Storage Overview



Hardware Components

The following hardware components for the Microsoft Exchange Server 2007 environment are referenced in this section:

- One Hitachi Adaptable Modular Storage 2300 disk storage system:
 - 8GB cache
 - Two Fibre Channel 4Gb ports
 - 21–146GB 15,000RPM SAS hard disk drives
- One Brocade® 48000 Fibre Channel Director with 6 x 4Gb ports
- One Sun Microsystems SunFire x4600 M2 for the Exchange 2007 Mailbox server:
 - Four dual core 2.6GHz AMD Opteron 2220 processors
 - 32GB RAM
- Two QLogic 2462 HBAs with QLogic driver version 9.1.7.15

- Two Hewlett-Packard DL385 G2 x64 servers for Exchange 2007 Hub Transport and Client Access roles:
 - Two dual core 2.8 GHz AMD Opteron 2220 processors
 - 8GB RAM
- Three Hewlett-Packard DL385 G2 x64 servers for the Hitachi management server and two Windows 2008 servers running Active Directory, Global Catalog and DNS servers:
 - Two dual core 2.8 GHz AMD Opteron 2220 processors
 - 8GB RAM

These components are only meant to serve as an example of the class of server and storage system configuration that are needed for an Exchange 2007 environment.

Software and Service Components

Hitachi offers a number of software applications to manage and configure the 2000 family. Storage Navigator Modular 2 is redesigned to look and function more like Hitachi Device Manager 6.0. Use Storage Navigator Modular 2 to create RAID groups and logical units, and to assign those logical units to the mailbox servers. Alternatively, Hitachi Device Manager 6.0 can be used to allocate the storage. All Hitachi software components can reside on a single management server. The deployment of Microsoft software products follows Microsoft's best practice recommendations.

The management server was deployed as follows:

- Microsoft Windows Server 2008 Enterprise x64 Edition, Service Pack 1
- Hitachi Basic Operating System M software
- Hitachi Device Manager 6.0 software
- Java Runtime Environment 1.6.0

The Mailbox server required the following software:

- Microsoft Windows Server 2008 Enterprise x64 Edition, Service Pack 1
- Microsoft Exchange Server 2007 Enterprise Edition, Service Pack 1
- Microsoft Exchange Server 2007 Management Tools, including Windows PowerShell
- Microsoft .NET Framework 2.0, Service Pack 1, or version 3.0
- Microsoft Management Console (MMC) 3.0
- Multipathing software
- Supported HBA driver

The Hub Transport and Client Access servers required the following software:

- Microsoft Windows Server 2008 Enterprise x64 Edition, Service Pack 1
- Microsoft Exchange Server 2007 Enterprise Edition, Service Pack 1
- Microsoft .NET Framework 2.0, Service Pack 1, or version 3.0

The Active Directory and DNS servers required the following software:

- Microsoft Windows Server 2008 Enterprise x64 Edition, Service Pack 1
- Microsoft .NET Framework 2.0, Service Pack 1, or version 3.0

To scale up or down from 5000 users, see Table 10. For all of the architectures, a 0.48 IOPS workload and 250MB mailboxes were used, as well as RAID-1+0 2D+2D and RAID-1 1D+1D RAID groups for the Exchange storage groups and logs. These architectures are also using 146GB 15K rpm SAS drives. Using 300GB 15K-rpm SAS drives allows for an increase in mailbox quota, but not the number of mailboxes. For example, 300GB drives at the same capacity utilization ratio can approximately double the mailbox quota for the 5000 user example while delivering similar performance. Keep in mind that an increase in mailbox size might result in an increase in non-transactional I/O. Also remember to use all eight Fibre Channel ports on the 2300 or 2500 when scaling up from the 20,000 user configuration.

Table 10. Reference Architectures Based on the Tested Environment

Number of Mailboxes	2500	5000	10,000	20,000	40,000
Total number of storage group LUs	4	4	8	24	48
Total number of transaction log LUs	4	4	8	24	48
Total number of paths	4	4	4	4	8
Number of servers	1	1	1	1	2
Total number of processor cores	8	8	8	8	8
Amount of memory per server (GB)	16	16	16	32	32

The various tests validate that the 2300 can scale to support more than 40,000 users because the Jetstress results for all of these architectures meet the Microsoft criteria of delivering the expected IOPS based on the I/O profile within 20ms response times. While the smaller configurations were tested with a server that far exceeds the Microsoft recommendations, the larger configurations did not meet Microsoft's recommendations for number of processor cores but still delivered the required IOPS within 20ms. The architectures for the larger configurations outlined in Table 10 match the criteria for organizations considering Exchange 2007 consolidation initiatives.

In most environments, mailboxes are distributed over multiple smaller servers. Organizational requirements, preferences or standards likely dictate the ultimate server hardware configuration. Table 11 shows possible configurations for maintaining up to 5000 users per server in a practical scenario. Memory is determined by number of storage groups, number of mailboxes and other factors. Keep in mind that Microsoft identifies 32GB of memory per server as the point of diminishing returns.

Table 11. Small Server Configuration Examples

Number of Mailboxes	2500	5000	10,000	15,000	30,000	40,000
Number of servers	1	1	2	3	6	8
Total number of processor cores per server	4	6	6	6	6	6

Hitachi Data Systems Best Practices

Following these guidelines provides an environment that can supply the required availability, performance and storage capacity. Not all recommendations are applicable to all environments.

Storage Availability Recommendations

Data path redundancy is enhanced on the Adaptable Modular Storage 2000 family due to the active-active controller. To ensure the highest availability of the data path from an Exchange server to its corresponding storage port on the storage system, follow these best practices:

- Use dual HBAs on each Exchange Server 2007 mailbox server.
- Deploy dual dedicated fabrics with multiple switches or directors.
- Use Hitachi Dynamic Link Manager multipathing software for automatic failover on each mailbox server.

Storage Design Recommendations

- Keep the Exchange Server 2007 workload isolated to its own RAID groups. Mixing intensive I/O from another application whose I/O profile is different can degrade Exchange Server 2007 performance.
- Due to the difference in I/O characteristics, isolate the Exchange Server 2007 storage group RAID groups from the transaction log RAID groups to take advantage of the intelligent storage controller's optimization routines.
- Use RAID-1+0 for database file RAID groups.
- Use RAID-1 or RAID-1+0 for transaction log files.
- Do not use LU concatenation because the extra disk spindles do not generally contribute to increased I/O rates or increased queue depth for the LU.
- Size storage solutions for Exchange based primarily on performance criteria. The number of disks, RAID level and percent utilization of each disk directly affect the level of achievable performance. Consider capacity requirements only after performance is addressed.
- The number of drives, coupled with the RAID level, determines the physical IOPS capacity of a RAID group and all of its LUs. If the storage design has an insufficient number of drives per RAID group, response times grow to large values very quickly.
- Validate storage design by testing with Jetstress and LoadGen.

Microsoft Best Practices

Following are some of Microsoft's most critical recommendations for Exchange Server 2007 regarding storage:

LU Partition Allocation Size

With Exchange Server 2007, Microsoft recommends that when defining the partitions that host the storage group databases, that an NTFS allocation unit of 64KB be defined. This recommendation is based on measured performance improvements seen with large sequential read operations. This type of profile is typically seen with streaming backup and Exchange Server Database Utilities (Eseutil) tasks. Performance benefits might result when running a 64KB allocation unit with Eseutil.

Microsoft testing validates that changing the NTFS allocation size from 4KB to 64KB does not result in any increase in transaction log sequential throughput. Therefore, you can use the default 4KB NTFS allocation size for NTFS volumes hosting transaction log files.

Aligning Disk Tracks in Windows Server 2003 Environments

Due to changes implemented in Windows Server 2008, it is no longer necessary to use `Diskpart.exe` to align disk tracks. This section is only applicable to Windows Server 2003 environments.

`Diskpart.exe` is a Microsoft utility that ensures that the disk tracks are sector and cache aligned. Windows implements a hidden boot record area before the start of the logical volume, which means that block 0 does not start on a true physical RAID stripe or chunk boundary. Disk operations that complete with a single access use fewer resources than those that overlap two physical disks. Also, I/Os that can be processed with a single cache slot use fewer cache and CPU resources than I/Os that overlap into a second cache slot.

An offset setting of 64KB aligns a logical disk for an Adaptable Modular Storage RAID disk track. Using the Windows `Diskpart.exe` utility sets the hidden sectors to equal 128 for all LUs that host any Exchange database or transaction log files on a given server. This aligns the disk records on a 64KB boundary and provides optimal disk and cache access. Performance improvements seen by aligning with `Diskpart.exe`, as measured in the Hitachi Data Systems labs, can be up to six percent for Adaptable Modular Storage systems.



Summary

Rapid adoption of Exchange Server 2007 is increasing the demand for resilient storage infrastructures, such as the Hitachi Adaptable Modular Storage 2000 family, to support those deployments. Planning, designing and implementing the underlying storage architecture are critical factors that an organization must consider when transitioning or migrating to Exchange Server 2007. With its new Dynamic Load Balancing Controller architecture, symmetrical active-active controllers and simplified Hitachi Storage Navigator Modular 2, the Hitachi Adaptable Modular Storage 2000 family delivers leading performance and makes deploying and managing storage for Exchange 2007 fast and easy.

A high-performing, durable storage infrastructure is a critical success factor for supporting a diverse set of Exchange Server 2007 deployments. The reference configurations outlined in this guide provide a solid foundation upon which solution architectures that can support varying levels of scalability, availability and performance requirements can be built. The solution architectures introduced in this guide provide guidelines necessary to architect high-performing and highly available Exchange Server 2007 deployments using the Adaptable Modular Storage 2300. These solutions are critical components in meeting business needs while leveraging and maximizing investments in Microsoft and Hitachi technology.


Appendix A: Example Deployment Planning Considerations

This appendix can be used as a high-level checklist by technical staff involved in the planning and implementation of a Microsoft® Exchange environment.

Planning Prerequisites

The first step in a successful Exchange 2007 deployment is to collect the data that enables you to fully understand your e-mail environment and the performance and capacity requirements you need to meet. As part of the planning stage of your deployment, consider these factors:

- **Active Directory topology** — An Exchange environment is either centralized or distributed. Because Exchange 2007 uses the Active Directory site topology for routing, understanding the existing Active Directory design is imperative. Certain requirements and restrictions exist relating to placement of Exchange 2007 roles within the Active Directory topology. If the environment is physically dispersed, multiple storage systems might be required.
- **Servers** — The amount of memory the servers have and the maximum amount that can be installed needs to be considered. Memory requirements increase as storage groups and users are added. A point exists where the addition of memory is not cost-effective. Microsoft considers this to be at 32GB. Consider the number of processor cores that are installed on the servers and the maximum number of users that can be supported. A good rule of thumb is that each processor core supports up to 1000 users. Verify that HBAs and their drivers are supported for the environment. Plan to install at least two HBAs per server to support multipathing. The 2000 family has active-active controllers that facilitate multipathing.
- **User I/O requirements** — User I/O requirements vary depending on multiple factors. The primary factor is the user profile. Microsoft defines four user profiles ranging from Light to Very Heavy. It is possible that some, if not all, of the users in a given environment fall outside of the Microsoft's defined profiles. One reason for this is the additional I/O requirements of BlackBerry or Windows Mobile users. These users can impose I/O requirements four or more times greater than a standard MAPI user. Another factor is the use of online mode instead of cache mode. Online mode puts a much heavier I/O load on the server, while cache mode moves a large percentage of the processing to the client. The ideal condition is to have accurate numbers for I/O requirements per user. If transitioning from Exchange 2003, keep in mind that the I/O load decreases substantially due to the improved memory usage in Exchange 2007. Also plan for future changes and growth. Take into account any changes that are expected in user access patterns, such as an increased number of BlackBerry users and any planned increase in the number of users due to expansion or acquisition. If possible, define classes of users with similar I/O requirements.
- **User capacity requirements** — User capacity requirements vary depending on multiple factors. These include requirements that are government or corporate mandated. Government related requirements usually address data retention. This capacity will be in addition to the user mailbox capacity and can take the form of journaling or archiving of e-mails for some or all users. Corporate mandates take the form of mailbox quotas for all e-mail users. Quotas are required to avoid uncontrolled growth of the mailbox databases. Companies might define one quota size for all users or might define different quotas for different departments, positions or roles. Certain types of businesses or certain types of employees within a company might require larger amounts of storage for a high number of attachments or large attachments.
- **Replication** — Replication includes any duplication of data, such as copies of data for backup and recovery purposes and disaster recovery or rapid recovery procedures. Data replication has a two-fold effect on storage infrastructure. You must plan for the capacity required to hold the copies and you must plan for the impact to backup procedures and backup and recovery windows. Certain types of replication, such as CCR and SCR, might require additional storage systems. Hitachi also has hardware-based replication



technologies that provide rapid backup and recovery of large databases while minimizing the impact on the backup window and the production environment in general.

Note that the factors you need to consider for your specific environment can vary and might be a subset of this list or might include factors not listed.

User Classification


Users can be classified in many different ways depending upon specific corporate requirements. These classifications can be used to segregate like users into specific storage groups. This practice can simplify management and aid in backup and recovery procedures. Consider which of the following user classifications might be most appropriate for the environment:

- **Organizational** — Users can be divided based on one or more organizational factors such as department, job title, role and so on. This type of division is probably the most common and can help simplify management of the environment. This type of division can be beneficial if backup procedures vary depending on department or if mailbox quotas vary by department.
- **Retention requirements** — Users might need to meet requirements for retention of data. In some cases, certain levels of management or users in particular roles might be required by law to retain data for specific amounts of time. Storage planning needs to take this into account.
- **Quotas** — To simplify management, institute consistent quotas for all mailboxes in a mailbox store. Mailbox stores containing mailboxes with different quotas can make it difficult to know if the underlying storage is overallocated.
- **Location** — Depending on your Active Directory design it may be necessary to divide users by physical location. Be sure that a user is assigned to a mailbox server and associated storage in the location that will give them the best user experience.

Storage Design

Preferred storage design varies widely depending on the IT standards for the corporation. Many IT departments implement one or more standard LU sizes to simplify storage management. Keep the following considerations in mind when designing the storage for your Exchange deployment.

- **Drive specifications** — Drive specifications and RAID types can be used to calculate the maximum number of IOPS that can be supported by a specific RAID layout. Whenever sizing environments always leave enough buffer both with I/O and capacity to handle any unexpected spikes or growth.
- **Standard LU size** — Many IT departments implement standard LU sizes to simplify management. For Exchange 2007 deployments using the 2000 family, use RAID-1+0 for the mailbox databases and RAID-1 for the logs. Never combine database and log LUs on the same set of disk spindles. In certain situations, RAID-5 can give acceptable performance for mailbox databases where the I/O load is low compared to the capacity requirements. Hitachi Data Systems recommends RAID-1+0 over RAID-5, but if you decide to use RAID-5 make sure you understand all of the possible results and make sure the environment is tested before putting it into production. RAID-5 has an I/O penalty during write operations. The use of standard-size LUs allows you to assign users to mailbox stores based on the I/O and capacity requirements of the user and the maximum IOPS and capacity of the LU. An LU might support a certain number of users based on I/O requirements or fewer users with larger mailboxes based on the capacity requirements. Always consider I/O requirements before capacity requirements.
- **Database size** — As a general best practice, a higher number of smaller databases is a better choice than fewer large databases. Because Exchange 2007 supports a maximum of 50 mailbox stores and 50 storage



groups per server, this practice is easy to follow. Make sure that the server has enough memory to support the number of storage groups configured. Microsoft recommends a maximum database size of 200GB when using CCR and 100GB without CCR. Larger mailbox stores can safely be used when using hardware-based replication technologies supported by the 2000 family. Backup and recovery procedures and requirements must be taken into account when defining the size of databases for an Exchange environment. If larger disks are used, and larger LUs are configured, make sure that your backup and recovery methods and schedule will not be affected.

- **Number of LUs per RAID group** — Configure only one LU per RAID group for mailbox store. For logs, you can configure two log LUs per RAID group.
- **User distribution** — Make sure that the RAID group can handle the number of users planned. The number of users per RAID group might be much lower due to capacity considerations.

Storage Configuration

Storage configuration requires multiple steps, including creation and mapping of LUs, creation and formatting of volumes on the host and configuration of databases and logs:

- **LU creation** — The 2000 family greatly simplifies this process. Use either Hitachi Storage Navigator or Hitachi Device Manager to quickly create LUs. Due to the architecture of the SAS backplane on the storage system, placement of the RAID groups is no longer as important as it was in earlier generations of the Hitachi Adaptable Modular Storage family. Create all database LUs before creating log LUs. This ensures balanced distribution of the LUs across the active-active controllers of the 2000 family.
- **Configure SAN** — Connect SAN cables and zone switches (if applicable) so that each server has at least two unique paths to the storage system.
- **Map LUs** — Map each of the LUs to one port on each controller. While any of the ports can communicate with either controller, this design allows for the best use of the active-active controllers and is the best design to take advantage of MPIO technology on the host.
- **Create and format volumes** — Discover LUs. Create volumes and mount these volumes to mount points. Create a mount point structure that is intuitive and indicates the use of each volume. Mount points simplify management and allow for growth.
- **Configure databases and logs** — Delete default databases and logs and create new ones on the mount points.

Performance Monitoring

After you complete the storage configuration, use Jetstress to verify performance of the environment before placing it into production. Add users and create mailboxes only after testing verifies the required performance levels. Add users in stages while monitoring the environment continuously. Focus ongoing monitoring on performance and storage capacity utilization.

Appendix B: Additional Reference Material

For more information about related topics, see the following documents.

Hitachi Data Systems

Additional material from Hitachi is available on the [Hitachi Data Systems support portal](#).

Training courses on the Hitachi Adaptable Modular Storage 2000 family are available from the [Hitachi Data Systems Learning Center](#).

For additional information and contacts, see <http://www.hds.com/>.

Microsoft

The Microsoft® [TechNet](#) site is a rich resource for Exchange content. For specific Exchange Server 2007 content, see these resources:

- [Exchange Server Library](#)
- [High Availability](#)
- [Exchange Server Tools Documentation](#)

Contact members of the Microsoft Exchange team at the official [Microsoft Exchange Team Blog](#).

Microsoft Exchange Server 2007 Books

The following books are good references for Exchange 2007 implementations:

- *Designing Storage for Microsoft Exchange 2007 SP1*, by Pierre Bijaoui and Juergen Hasslauer, ISBN-13: 978-1555583088
- *Microsoft Exchange 2007 with SP1: Tony Redmond's Guide to Successful Implementation*, by Tony Redmond, ISBN-13: 978-1555583552
- *Microsoft Exchange Server 2007: Implementation and Administration*, by Jim Mcbee, ISBN-13: 978-0470187418
- *Microsoft Exchange Server 2007 Administrator's Companion*, by Walter Glenn, Scott Lowe and Joshua Maher. ISBN-13: 978-0735623507

PowerShell Books

PowerShell, the scripting language that Microsoft is using to replace cmd. exe, is an important topic for Exchange administrators. For Exchange 2007, PowerShell is the basis for Exchange Management Shell, which is the CLI for Exchange 2007. The Exchange functionality within PowerShell is supported by an add-in to the base PowerShell install. The authors recommend the following PowerShell books:

- *Professional Windows PowerShell*, by Andrew Watt, ISBN-13: 978-0471946939.
- *Professional Windows PowerShell for Exchange Server 2007 Service Pack 1*, by Joezer Cookey-Gam, Brendan Keane, Jeffrey Rosen and Jonathan Runyon, ISBN-13: 978-0470226445.

Appendix C: Exchange Server 2007 Tools

Microsoft offers several tools that can aid in the validation of a particular Exchange 2007 Server and storage design prior to implementation in a production environment. Links to the specific download page for the 64-bit version of each tool are given within each subsection. The 32-bit version of any of these tools can be accessed from the main tools page. The following tools can all be accessed through the [Tools for Exchange Server 2007](#) page on Microsoft TechNet.

Exchange Server 2007 Storage Calculator

The Microsoft Exchange team provides an [Exchange 2007 Mailbox Server Role Storage Requirements Calculator](#) to assist organizations in designing their storage layout for Exchange Server 2007. This calculator assists in the determination of the storage requirements (I/O performance and capacity) and the optimal LU layout based on a set of user-defined and pre-configured values.

Exchange Server Jetstress Tool

The [Microsoft Exchange Server Jetstress Tool](#) is used to verify the performance and stability of a disk storage system prior to putting an Exchange server into production. Jetstress verifies disk performance by simulating Exchange disk I/O load. You use Performance Monitor, Event Viewer, and ESEUTIL in conjunction with Jetstress to verify that your disk storage system meets or exceeds the performance criteria you establish. Microsoft's criteria for success are outlined in Table 12.

Table 12. Jetstress Success Criteria

Parameter	Success Criteria
Achieved I/O	➡ = Estimated IOPS
Database Avg. Disk sec/Read	Avg. < 20ms, Max < 50ms
Database Avg. Disk sec/Write	Avg. < 20ms
Log Avg. Disk sec/Read	Avg. < 20ms, Max < 50ms
Log Avg. Disk sec/Write	Avg. < 10ms, Max < 50ms

After a successful completion of the Jetstress disk performance and stress tests in a non-production environment, you can be confident that your Exchange 2007 disk storage system is adequately sized in terms of the supplied performance criteria for the user count and user profiles you establish.

Exchange Load Generator Tool

Perform benchmarking, pre-deployment validation and stress testing tasks that introduce various types of workloads into a test Exchange messaging system. Simulate the delivery of multiple MAPI client messaging requests to an Exchange server.

Use [Microsoft Exchange Load Generator](#) (LoadGen) as a simulation tool to measure the impact of MAPI clients on Exchange servers. LoadGen allows you to test how a server running Exchange responds to e-mail loads. LoadGen is a useful tool for administrators who are sizing servers and validating a deployment plan. Specifically, LoadGen helps you determine whether each of the servers can handle the load that they are intended to carry. LoadGen can also help validate the complete solution, such as the network, along with server and storage performance. LoadGen is not normally used to validate a storage infrastructure. Instead use Jetstress for this purpose. However, Jetstress does not enable you to predict or test how the I/O footprint changes on an Exchange 2007 Mailbox server as the memory configuration is changed. LoadGen is required to quantify how memory affects Exchange I/O characteristics.



Exchange Server Stress and Performance Tool

The [Exchange Server Stress and Performance](#) tool simulates large numbers of client sessions by concurrently accessing one or more protocol servers. Exchange Server Stress and Performance includes multiple modules that you can use to simulate several protocols and loads.

Exchange Server Profile Analyzer

The [Exchange Server Profile Analyzer](#) collects estimated statistical information from a single mailbox store or from across an entire Exchange Server organization. You can then use the collected data for tasks such as analyzing the performance and health of a server that has mailboxes.

Exchange Server 2007 Best Practices Analyzer

The [Microsoft Exchange Best Practices Analyzer](#) programmatically collects settings and values from various data repositories. After the data is collected, a set of comprehensive best practice rules is applied to the topology. Administrators running this tool receive a detailed report listing the recommendations that can be made to the environment to achieve greater performance, scalability and uptime.

Exchange Server Toolbox

The [Exchange Server 2007 Toolbox](#) includes various management and troubleshooting tools. These include tools for configuration management, disaster recovery, mail flow and performance. The Exchange Server 2007 Toolbox is installed by default when the Exchange management console is installed, and is available as a node from that application.



Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom

Contact Information: + 44 (0) 1753 618000 www.hds.com / info.emea@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., and/or its affiliates 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.

TrueCopy and ShadowImage are registered trademarks of Hitachi Data Systems Corporation.

Microsoft is a registered trademark of Microsoft Corporation.

All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to <http://www.hds.com/corporate/legal/index.html> or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

© Hitachi Data Systems Corporation 2008. All Rights Reserved.

WHP-303-01 DG December 2008