



Technical report: IBM System Storage N series and Microsoft Exchange Server 2007

Best practices guide

• • • • • • • • •

Document NS3578-0

February 3, 2008



Table of contents

Abstract	4
Introduction	4
Purpose and scope	4
Intended audience	4
Caveats.....	4
IBM N series for Microsoft Exchange Server 2007	5
Data management suite.....	5
SnapDrive.....	5
SnapManager for Microsoft Exchange	5
Single mailbox restore	6
Enterprise storage.....	6
Disk types	6
Fibre channel.....	6
ATA	7
RAID	7
RAID level comparison	7
Advantages of RAID-DP	7
Protecting Exchange data	7
RAID-DP performance.....	8
Rebuild times	8
IBM N series aggregates, volumes, and LUNs	8
Examples using traditional versus flexible volumes.....	9
Recommended layouts.....	11
Aggregate recommendations	11
Volume recommendations	12
LUN recommendations.....	12
Two LUNs per storage group.....	12
Two LUNs per backup set	12
NTFS hard links.....	13
Number of snapshot copies to keep online	14
Mapping Exchange users to storage groups.....	15
Sizing Exchange Server 2007 deployments	16
Planning storage capacity	17
Key Exchange Server factors	17
Mailbox size/quota.....	17
Database dumpster	17
Content indexing.....	17
Maintenance and moving mailboxes	17
Maintenance	17
Moving mailboxes.....	18
Capacity planning: Exchange Server 2007 Server roles	18
Edge Transport	18



Hub Transport	19
Transport Database.....	19
Planning for I/O	20
Transactional I/O.....	20
Non-transactional I/O	20
Server considerations	23
Effect of host-side memory on storage	23
Mailbox role	24
Minimum memory configuration for mailbox based on number of storage groups.....	24
Memory recommendations for log shipping	25
Sizing additional processing capacity for local continuous replication (LCR)	25
Storage monitoring	25
Conclusion.....	25
Appendix A: Best-practices summary	26
SnapDrive	26
SnapManager for Exchange	26
Disk types	26
RAID-DP	26
Aggregate recommendations.....	27
Volume recommendations	27
LUN recommendations	27
NTFS hard links	27
Mapping Exchange users to storage groups	27
Sizing Exchange Server 2007 deployments.....	28
Maintenance	28
Capacity planning for Exchange 2007 Server roles	28
I/O planning best practices	29
Processor configuration best practices.....	30
Mailbox role memory configuration.....	30
Trademarks and special notices.....	31



Abstract

This paper is designed to offer guidance in planning and sizing Microsoft Exchange Server 2007 environments. The best practices and recommendations set forth in this guide will ensure a highly available, easy-to-manage Exchange environment that will meet SLA requirements. Consult with a local IBM System Storage N series Exchange expert when planning and deploying Exchange Server environments when using N series storage solutions.

Introduction

Corporate messaging has become mission-critical in nearly all organizations, bringing additional constraints for information technology organizations. Not only do they have to deal with requests for increased user mailbox sizes and additional users, they need to ensure high availability for their messaging infrastructure while providing increased uptime percentages. This, in turn, reduces the amount of time available for operational maintenance.

In order to meet the demand of larger mailboxes as well as provide increased performance and availability, proper planning and deployment are critical in corporate messaging environments. Success or failure of the Exchange messaging deployment hinges on proper planning and configuration choices.

Purpose and scope

The following areas will be covered:

- IBM® System Storage® N series solutions for Microsoft® Exchange Server 2007
- Sizing Exchange Server 2007 deployments
- Capacity storage planning
- Planning for I/O
- Server considerations

NOTE: This document highlights specific best practices in boxed sections throughout the document. A complete list of best practices is also available in Appendix A.

Intended audience

This paper is intended for experienced Exchange server and/or storage administrators. Readers should have a solid understanding of Exchange server storage architecture as well as Exchange backup and restore concepts for Microsoft Exchange Server 2007.

Caveats

The best practices for Exchange Server 2007 storage architecture presented herein are focused exclusively on the N series storage operating system, IBM System Storage N series with Data ONTAP® 7.1 or later.

Guidance for Microsoft Unified Messaging Server is not included in this document.



IBM N series for Microsoft Exchange Server 2007

When deploying Exchange Server 2007 onto an IBM N series storage appliance, it is important to understand the key components involved and how to properly utilize them. There are both hardware and software components that make up the entire N series solution for Microsoft Exchange Server 2007.

Data management suite

SnapDrive

IBM System Storage N series with SnapDrive® for Windows® is an enterprise-class storage and data management solution for Windows environments. SnapDrive simplifies and automates application storage provisioning, provides consistent data copies through the use of IBM System Storage N series with Snapshot™, rapid application recovery, and overall data storage management for Windows servers.

When provisioning N series logical unit numbers (LUNs) for use with Exchange (or other Windows applications), SnapDrive will handle all of the necessary tasks to create, align, and present a LUN to the Windows operating system. No manual steps need to be taken by a storage administrator. Simply perform a “LUN create” operation through SnapDrive. The same ease-of-use principles are enjoyed on other SnapDrive operations.

SnapDrive for Windows 4.2 and later fully support volume mount points (VMPs). This support allows an Exchange administrator to surpass the drive-letter restrictions within Windows and mount more LUNs for a given Exchange server. This is key when following Microsoft best practices, which require more LUNs to be mounted per server than there are drive letters available. With volume mount points, SnapDrive provides the flexibility needed to handle very complex Exchange layouts and not restrict administrators to drive letter-mounted LUNs.

Best practice

Use SnapDrive to create and manage N series LUNs. SnapDrive handles the necessary tasks (partition, align, format) to properly create and mount a LUN on a given Exchange server.

Note: At the time of this writing, SnapDrive supports the creation of up to 120 mount points.

For more information on SnapDrive for Windows, visit the SnapDrive for Windows page on the IBM N series website.

SnapManager for Microsoft Exchange

IBM System Storage N series with SnapManager® 4.0 for Microsoft Exchange (SME), the current version at the time of this writing, supports Microsoft Exchange Server 2003 and 2007. SME has achieved a “[Certified for Windows](#)” logo for Windows Server 2003 Standard and Enterprise Editions. SME is also Microsoft “[SimpleSAN](#)”-designated and a Windows Server 2003 certified backup and recovery solution for Exchange Server.

SME is tightly integrated with Microsoft Exchange, which allows for consistent online backups of Exchange environments while leveraging N series Snapshot copy technology. SnapManager 4.0 for Exchange supports Microsoft Volume Shadow Copy Services (VSS), Virtual Disk Services (VDS), and other Microsoft storage technologies. (For details on VSS, see [Microsoft knowledge base article](#))



[822896](#).) SME is a VSS (snapshot copy) requestor, which means that it uses the Microsoft-supported VSS subsystem to initiate backups. SME provides a complementary feature set for new Microsoft Exchange Server 2007 data replication features. SME works with Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR) replica databases and provides a rich feature set to leverage these new technologies.

Best practice

Use SnapManager for Exchange when deploying Exchange Server 2007 on N series storage. SME will perform the data migration from local disks to N series LUNs. It will also manage that data, handling all backup, restore, and verification tasks.

Single mailbox restore

IBM System Storage N series with Single Mailbox Recovery (SMBR) allows users to search through and quickly locate mail items that may have been inadvertently deleted from a user's mailbox. SMBR enables users to access e-mail items up to and beyond the 14-day default limit of the deleted items' retention cache. This provides greater flexibility to Exchange administrators who need fast and easy user mailbox access to recover necessary e-mail items.

SMBR can be run on any machine that has access to the Exchange data from where the mail items will be retrieved. By mounting a LUN, either backed by a snapshot copy or from an IBM System Storage N series with FlexClone® copy, and locating the appropriate backup set, SMBR can quickly sort through the archived Exchange database and locate any mail items that an Exchange administrator is looking for. With its advanced search capabilities, keywords or other search criteria can be used and SMBR will search through the entire database, returning the matching criteria.

SMBR also includes extract wizards that are integrated with most popular tape environments. These wizards communicate with backup software to find and locate mail items.

SMBR does not require any additional storage to be available in order to recover mail items from a backup set or archive. Accessing the database within a snapshot copy, or using the extract wizards to locate items on a tape media, SMBR eliminates the storage required to restore the entire database to an Exchange server and then find the necessary mail items.

Enterprise storage

Disk types

Exchange Server 2007 is still an I/O-intensive application, even with the new enhancements to reduce the amount of I/O it generates. Operations like database writes, transaction log writes, and verifications place a large amount of I/O load against any storage platform hosting the Exchange data.

Fibre channel

Fibre channel (FC) disk drives have been on the market for a very long time. They are a proven reliable storage device, with high read/write speeds, and have the ability to handle high I/O loads. FC disks are ideally suited to handle the I/O requirements of Exchange 2007 and are an integral factor in meeting and maintaining low latency times required by Exchange.



ATA

ATA disk drives are becoming increasingly popular as a storage medium. Although ATA disks are a low-cost, high-capacity storage solution, they are not well suited for high I/O requirements for larger Exchange environments. ATA drives are better suited to serve as an archival destination, where data can be stored and protected for relatively low cost. Using IBM System Storage N series with SnapVault® technology, users can safely store Exchange backup sets onto ATA drives, where they will be kept safe and protected and still be readily and quickly accessible when needed.

Best practice

FC disks are the recommended choice to meet the needs for Exchange 2007. However, with SnapVault, ATA disks can also meet the needs for Exchange 2007 in less I/O intensive deployments. To determine the I/O requirement for your Exchange environment and the best disk-type match, please consult with your local IBM N series Exchange expert.

RAID

RAID level comparison

The following table, from the Microsoft TechNet article, "[Planning Disk Storage](#)," compares the three most common types of RAID readily available.

RAID type	Speed	Capacity utilization	Rebuild performance	Disk failure performance	I/O performance
RAID 10	Best	Poor	Best	Best	Best
RAID 5	Good	Best	Poor	Poor	Poor
RAID 6	Poor	Good	Poor	Poor	Poor

Table 1) Comparison of RAID solutions.

The table states that RAID 6 is not an acceptable RAID solution for Exchange Server 2007. Poor speed and poor performance in general far outweigh the data protection benefits of double parity in RAID 6. However, the N series double-parity RAID solution, IBM System Storage N series with RAID-DP™, has the double-disk protection of RAID 6 without incurring the performance penalties that are typical with RAID 6 implementations. This makes RAID-DP an optimal RAID solution for Exchange environments.

Advantages of RAID-DP

N series RAID-DP is often categorized as another type of RAID 6. This is not accurate. While RAID-DP and RAID 6 share common benefits (near zero data loss from a double disk failure), RAID-DP does not suffer from the same limitations of RAID 6. RAID-DP has optimizations over RAID 6 that offer benefits far superior to other RAID implementations.

Protecting Exchange data

Although double disk failures are a rare event, they do indeed happen. If you have a RAID 5 storage array for Exchange data, and there is a double disk failure, there is a 100% chance that Exchange data will be lost.



RAID 6 and RAID-DP are protected against a double disk failure because of double-parity protection. In the event that a second disk fails during a RAID reconstruction, the second parity disk will take over and continue to reconstruct the data on the failed disk.

RAID-DP performance

Again, RAID-DP does not suffer from the same performance issues as typical RAID 6 configurations. RAID-DP is streamlined and optimized to provide high performance while providing the high data availability that Exchange users and administrators demand.

All IBM N series [Exchange ESRP submissions](#) used RAID-DP. This is proof that RAID-DP provides excellent performance that meets or exceeds Exchange Server 2007 recommendations.

Rebuild times

RAID-DP will automatically start to rebuild a broken disk upon failure. There is no user intervention required. RAID-DP is optimized to shorten the time a given RAID group is in degraded mode.

Best practice

Always use RAID-DP when deploying Exchange server data onto an N series storage appliance. RAID-DP offers the best combination of performance and data protection for Exchange server environments.

IBM N series aggregates, volumes, and LUNs

Data ONTAP 7.1 and later provides functionality that enables the creation of flexible volumes for managing data without the need to assign physical disks to the volumes. Instead, the IBM System Storage N series with FlexVol® volumes enjoy performance benefits from a larger pool of physical disks, called an “aggregate.”

This results in the following additional benefits for Microsoft Exchange environments:

- A large number of volumes can be created, all with independent snapshot copy schedules, mirroring policies, etc.
- All volumes can be managed independently, while receiving the maximum I/O benefit of a much larger pool of disks.

Examples using traditional versus flexible volumes

The volume structure shown in Figure 1 illustrates a scenario using traditional volumes that yields the greatest operational flexibility. Using traditional volumes, each disk is associated with a specific volume and is for the exclusive use of that volume only.

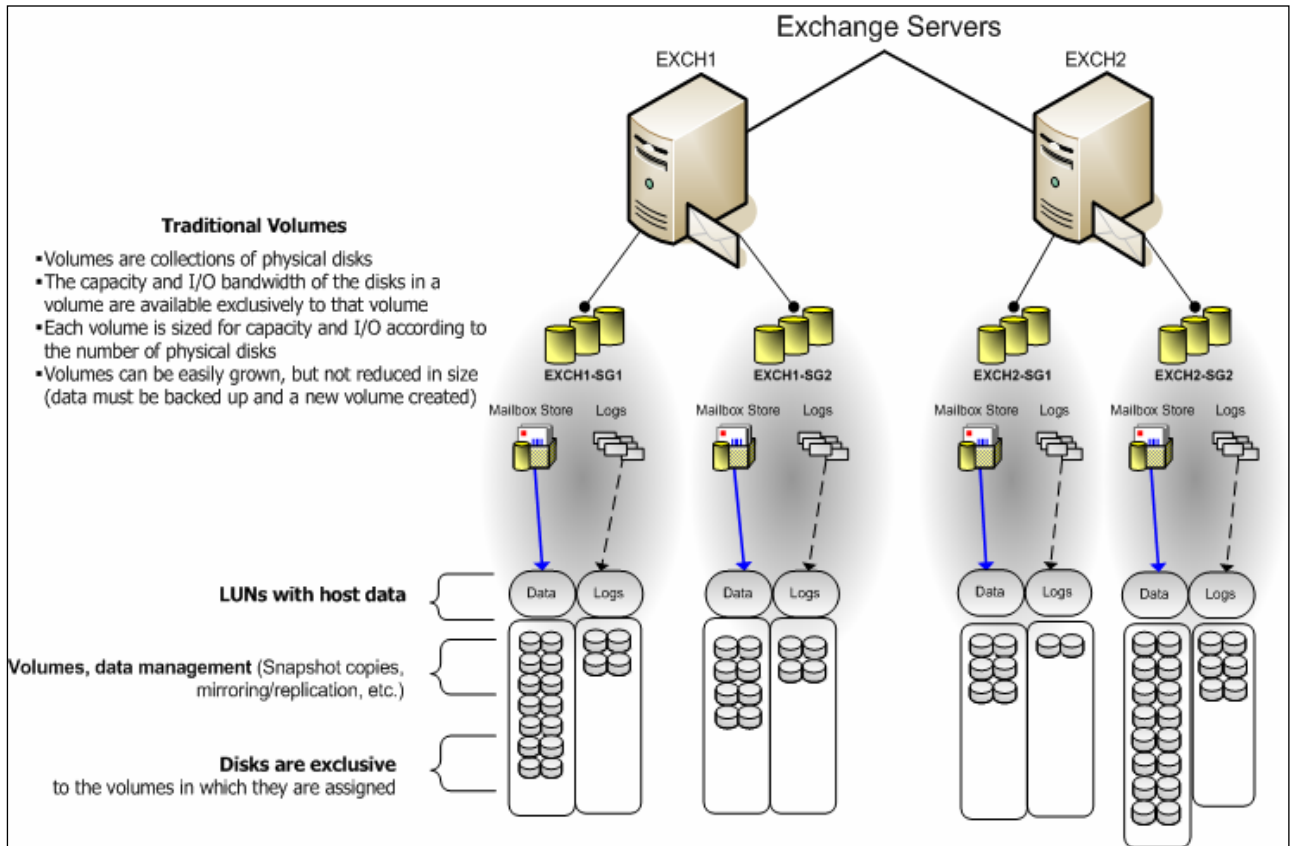


Figure 1) Using Data ONTAP 7G (7.1 or later) with traditional volumes.

Figure 2 shows a simplified example of how the same volume structure that existed previously can be used in a configuration with flexible volumes. This configuration yields the same operational flexibility for Exchange, but provides significant advantages. When using flexible volumes, the capacity and performance bandwidth of a large collection of fast disks is made available to all volumes, no matter their size. Even small volumes enjoy the benefit of a very large number of spindles. Thus, volumes can be better leveraged for managing data and still realize the performance benefit of the aggregate's total spindle count.

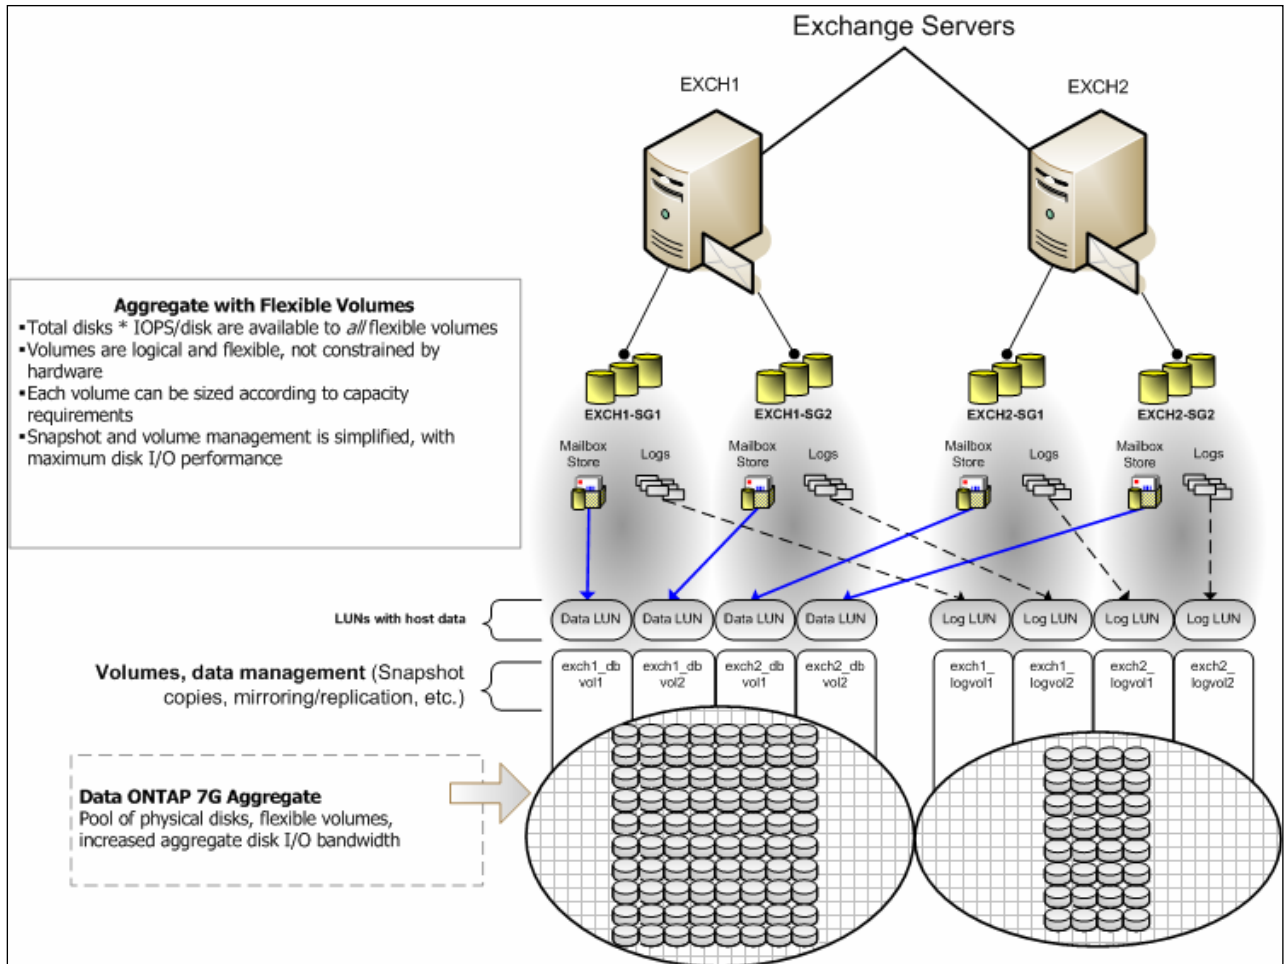


Figure 2) Using Data ONTAP 7G (7.1 or later) with flexible volumes.

Data ONTAP 7.1 or later has flexible options for Exchange administrators to efficiently use all available disks attached to N series platforms. The benefits are easily realized for all Exchange servers that store data on N series systems. Traditional volumes or flexible volumes can be used, depending on the operational and performance needs of small to very large enterprise deployments.

For accurate sizing and deployment of Exchange storage, consult a local N series Exchange expert. Experts are the best resource to ensure a successful deployment that is sized correctly.



Recommended layouts

Aggregate recommendations

Pooling all of the available disks into a single, large aggregate may maximize performance; however it might not meet the data availability requirements set forth in the service level agreement (SLA) agreement.

Best practice

When using aggregates, it is recommended to move the database flexible volumes and LUNs from the transaction log flexible volumes and LUNs onto different aggregates.

Creating separate aggregates for Exchange database volumes and Exchange transaction log/SnapInfo volumes can still meet the performance requirements of Exchange Server 2007 while still providing the data availability required by most typical SLA agreements. If, in the unlikely event an aggregate is lost, part of the Exchange data is still available. An Exchange administrator can potentially recover data from the available aggregate.

The table below details the benefits and tradeoffs for two possible aggregate layouts.

Number of aggregates	Benefits	Tradeoffs	Recommended configuration
One single aggregate for all Exchange data	<ul style="list-style-type: none">Minimize the number of parity disks required	<ul style="list-style-type: none">In the unlikely event of losing an aggregate, Exchange data is completely lost from the affected aggregate	<ul style="list-style-type: none">Low I/O Exchange deployments
Multiple aggregates for Exchange data (database and transaction logs on separate aggregates)	<ul style="list-style-type: none">Dedicated aggregate handling the I/O for a database or for transaction logsIf a single aggregate is lost (very unlikely), there is not a total loss of Exchange data	<ul style="list-style-type: none">More parity disks required because there are more aggregates	<ul style="list-style-type: none">I/O intensive Exchange deployments

Table 2) Aggregate layouts.

Consult a local N series Exchange expert to provide guidance on whether to deploy a single aggregate or multiple aggregates for an Exchange environment.



Volume recommendations

Volume layout is critical in creating and sustaining a highly available Exchange environment. Careful consideration for various backup groups, disaster recovery scenarios, and even archiving solutions will help determine the placement of volumes onto aggregates, and the corresponding LUNs onto those volumes.

Best practices

When using aggregates, it is recommended that the database flexible volumes and LUNs are separate from the transaction log flexible volumes and LUNs. In the unlikely event that an aggregate is lost, Exchange data is still available for disaster recovery.

It is recommended that there are different database and transaction log volumes for different Exchange servers, to prevent a potential “busy” snapshot copy problem. Because there are separate volumes for each server, there is no need for concern regarding snapshot schedules overlapping each other for different servers.

Place database LUNs and transaction log/SnapInfo LUNs in separate volumes.

If there are separate LUNs for the Exchange transaction log files and the SnapInfo directory, place those LUNs in the same volume. Both of these LUNs will have a similar I/O profile, allowing them to share the same volume. And for disaster recovery scenarios, having the entire log set for Exchange on the same volume will help achieve SLAs.

LUN recommendations

Best practice

Separate Exchange database and transaction log files onto separate LUNs and separate volumes. This allows for greater flexibility for backup and recovery procedures and data protection strategies.

Current Microsoft best practices state that all databases in a storage group should be on a single LUN. Also, it is a best practice to place no more than one database in each storage group. Exchange 2007 has increased the storage group limit from 4 to 50. To reduce complexity, Microsoft recommends the two following designs:

- Two LUNs per storage group
- Two LUNs per backup set.

Two LUNs per storage group

This recommendation gives you the greatest flexibility for storage group management, but has great complexity from a LUN management standpoint. With 50 storage groups as a maximum for Exchange 2007, the number of LUNs required may exceed current N series limitations.

Two LUNs per backup set

This recommendation provides the greatest LUN management simplicity, but has the least amount of flexibility when it comes to storage group management. Also, with fewer LUNs, there is more data stored on each one. With a current-version 2-TB limit on master boot record (MBR) partitions, scaling the storage to meet the needs of a particular storage group will be challenging.

N series provides Exchange administrators the flexibility to deploy Exchange data in a layout that best suits their requirements. Identifying a key balance between the number of storage groups required for an Exchange environment and the number of LUNs that have to be managed is a key component of a successful Exchange deployment.



When planning how Exchange data will be deployed onto N series LUNs, remember the following:

- RTO/RPO times
 - With many storage groups and a high number of LUNs, recovery may take longer to achieve, as there are more physical objects to handle
- Backup schedules
 - With an increased number of LUNs and storage groups, backup schedules can become quite complex
- Replication/mirroring
 - An increased number of LUNs typically means smaller dataset sizes; when replicating key storage groups, smaller dataset sizes may help shorten replication times.

As always, consult with a local N series Exchange expert for guidance on deploying an Exchange environment.

NTFS hard links

The number of transaction logs generated for a storage group directly affects the size of the SME SnapInfo directory. When the transaction log directory resides on a different NTFS volume from the SnapInfo directory, SME will perform a file copy for each transaction log, copying the logs into the SnapInfo directory. A copy operation is performed instead of a move operation because it eliminates the potential for file corruption if something interrupts the move operation. Once the transaction logs are copied, Exchange will truncate any logs that were committed to the database if SME is configured to request a log truncation.

When the transaction log directory resides on the same NTFS volume as the SnapInfo directory, SME will utilize [NTFS hard links](#) to avoid the file copy operation and maximize storage utilization. During the backup process, SME will edit the file pointers for a given transaction log. SME inserts a pointer to the SnapInfo directory and removes the pointer to the Exchange transaction log directory. As a result, it will appear as if a transaction log was moved from one directory to the other. However, a physical move operation was not performed, and no I/O overhead was incurred on the Exchange server. Only the pointers for a transaction log file were updated, thus saving the additional space required to copy the transaction log files into the SnapInfo directory.

Best practice

Utilize NTFS hard links whenever possible to increase storage utilization, eliminate the physical copy overhead on the Exchange server, and increase backup performance. SME will automatically utilize NTFS hard links when it detects the transaction log file directory and the SnapInfo directory on the same NTFS volume.

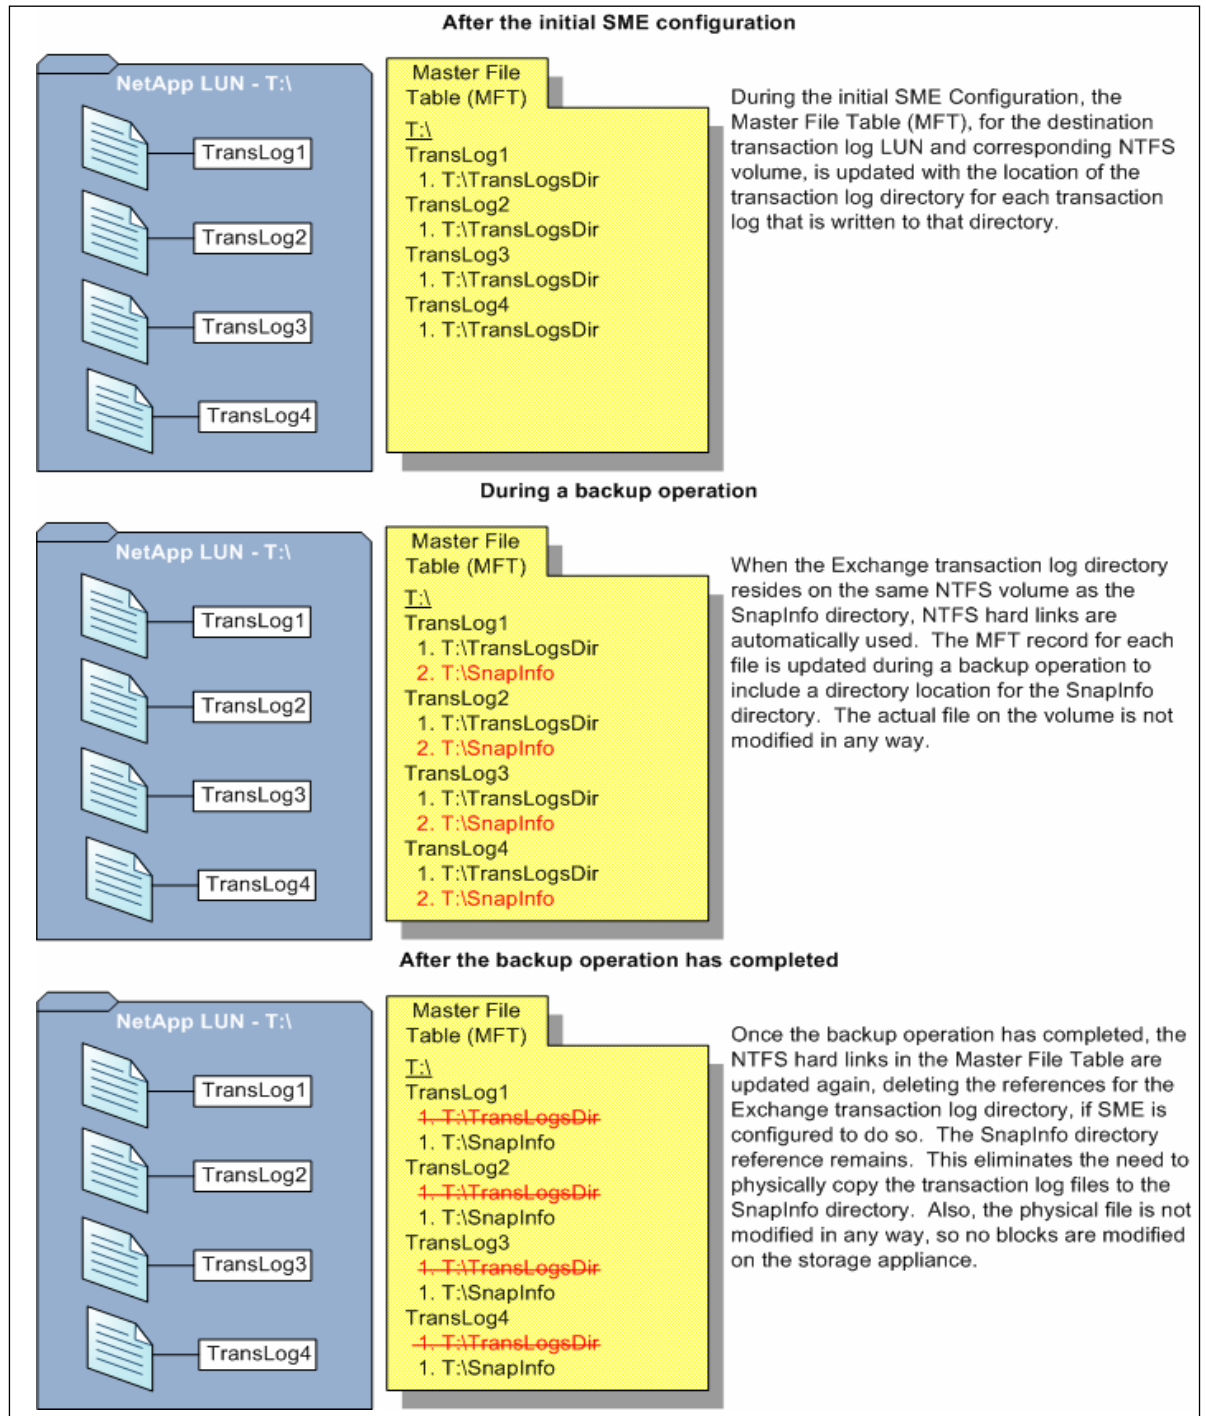


Figure 3) NTFS hard links.

Number of snapshot copies to keep online

When a snapshot copy is kept online, it is readily accessible by the Exchange administrator to perform recovery procedures on. The number of snapshot copies to keep online will be determined by the requirements in the SLA agreement. More production storage capacity is



required if a larger number of snapshot copies are kept online. However, offloading older snapshot copies onto less expensive storage can be a less expensive solution, while still providing fast access to those snapshot copies. IBM System Storage N series with NearStore® storage appliances provide an ideal, cost-efficient storage solution for archiving older snapshot copies. Using N series SnapVault technology, you can quickly and easily replicate the snapshot copies onto a NearStore appliance where they will be stored and protected until they are needed.

Mapping Exchange users to storage groups

All users on a given Exchange server may not have the same SLA requirements. Typically, executive-type user mailboxes are more important than other user mailboxes. Careful consideration when laying out the storage and the user mailboxes onto that storage is an important part of storage design.

Best practice

Place user mailboxes that have stringent SLA requirements in a separate storage group. That storage group should have its own dedicated set of LUNs residing on a separate volume. Configuring the storage in such a way enables much more flexibility in managing those user mailboxes and protecting them as required by the SLA requirements.

Managing important user mailboxes so that are separated from other mailboxes on the same Exchange server is much easier when they are located on their own dedicated storage. For example, a storage group that contains executive users for a company would be backed up more often than a regular user storage group for that same company. Having those executive user mailboxes on a LUN that is on a dedicated storage volume enables scheduling of more frequent backups for that group of users. Also, the size of the Exchange database may be smaller (depending on the I/O profile of the users) than the other users on the same Exchange server. This means faster database verification times, faster recovery times, and faster mirroring times to an alternate storage facility for data protection.



Sizing Exchange Server 2007 deployments

A properly sized Exchange environment will meet or beat both Microsoft requirements for Exchange storage and customer requirements indicated in their SLAs. To get a properly sized environment, information about the customer environment is collected and tools are used to convert that information into a physical storage recommendation.

There are two primary tools that should be utilized when planning an Exchange environment for a customer:

- The Microsoft storage calculator
- The N series Exchange Sizing Tool.

The Microsoft storage calculator provides sizing information based on the following parameters:

- Exchange server configuration
- Mailbox configuration
- I/O requirements
- Exchange data and backup configuration
- Client requirements.

The sizing information provided by these tools is an important component and provides a framework for storage group layout and LUN requirements. It is important to realize that the Microsoft storage calculator does not make any recommendations on storage design (RAID parity, number of disks, etc), as the storage design is largely dependent upon the type of storage array being utilized. When sizing Exchange server deployments using N series storage it is important the N series Exchange Sizing Tool be used.

Best practice

The N series Sizing Tool for Exchange should be used to size all Exchange server deployments utilizing N series storage.

Planning storage capacity

Key Exchange Server factors

There are many factors required for correctly sized storage for Microsoft Exchange Server 2007. Listed below are three important factors directly affecting disk and storage capacity:

- Mailbox size/quota
- Database dumpster
- Content indexing.

Mailbox size/quota

An important factor (if not the most important factor) is mailbox size or quota. Understanding the mailbox size/quota is the starting point for the amount of data storage required. It will also indicate how many users can be hosted per Exchange server.

When determining the mailbox size per user for a given Exchange environment, remember to think about future storage growth needs. Keeping potential future storage needs in mind when sizing the current Exchange environment will prevent storage headaches down the road.

Database dumpster

For each Exchange database, there is a database dumpster that stores hard-deleted items. These items will be retained, by default under current versions, for 14 days in Exchange Server 2007 (twice as long as Exchange Server 2003). The increased retention times also increase the amount of storage required for the database dumpster. The amount of storage and I/O overhead depends on the mailbox sizes and the message delivery rates for those mailboxes.

When sizing an Exchange 2007 environment, always take into consideration database dumpster retention policies and size accordingly. Depending on the needs of a customer, the default retention period may be adjusted, which directly affects the overhead it has on the storage.

Content indexing

Exchange Server 2007 creates an index of an Exchange database that is approximately 5% of the total database size. This index is placed on the same LUN as the indexed database. Sizing both the volume and the hosting LUN to account for the index is very important. Users may turn off Content Indexing for a given database, but by default in Exchange 2007 it is turned on.

Maintenance and moving mailboxes

Maintenance

Planning storage for database maintenance during the initial phase of deploying an Exchange environment can save headaches when the time comes to perform those tasks. A repair operation or a compact database operation will require storage that is equal to the size of the database plus 10%.

Best practice

When planning maintenance of Exchange databases, keep in mind the following:

- Always plan maintenance windows during off-peak or low-activity times.
- Never schedule a maintenance window during an SME backup operation.
- Never schedule a maintenance window during a verification operation if both operations are performed on the same server.

Moving mailboxes

Moving user mailboxes from previous versions of Exchange to Exchange Server 2007 will generate a substantial number of transaction logs on the destination Exchange server. During the mailbox move process, all message data is first written to the transaction logs, then committed to the database. Allocating more storage to handle user mailbox moves helps prevent the transaction log LUN from running out of space and causing downtime. Understanding the migration process for Exchange Server helps determine how much additional storage will be required for the user move operations.

In addition to migration, many organizations move a percentage of mailboxes on an ongoing basis. Ongoing moves should be factored in when planning the needed amount of storage.

Capacity planning: Exchange Server 2007 Server roles

Edge Transport

The Edge Transport server is responsible for message hygiene, routing messages to appropriate hub servers, as well as handling inbound and outbound e-mail traffic.

Under current versions, a minimum of 4 GB of available space is required on the LUN containing the Message Queue database. If this condition is not satisfied, the system will activate a condition known as back pressure. For more information about back pressure, and the options to configure back pressure, see the Microsoft TechNet article on [Understanding Back Pressure](#).

Capacity requirements of the Edge Transport server are affected by message tracking logs, protocol logs, mail database, and anti-virus.

If message tracking logs are enabled, additional capacity will be required. Message tracking capacity requirements depend on the number of messages received by the transport server. You can model your Exchange 2003 log generation rate, and set a hard limit for the number of days to keep data.

Transaction logs do not require much disk capacity because normal log creation is limited by the use of circular logging. As a result, transaction logs can be placed on the LUN containing the operating system.

The database (mail.que) does not store items indefinitely. The capacity reserved must be the average message size times the maximum queue. When the queue is at maximum, the server is shut down. A 500,000-item queue with an average message size of 50 KB is about 25 GB of data in the database.

Edge Transport servers that run anti-virus scans on incoming mail will need enough space for the anti-virus quarantine. The disk I/O resource requirements depend on the percentage of incoming mail that is infected with viruses, which is typically very small. The quantity of infected messages and attachments and



how long they remain in quarantine dictate the amount of space that quarantine will require. 1 GB is a good starting point, although each organization's actual needs will be different.

Hub Transport

The Hub Transport server routes mail to the proper mailbox server. It also includes a transport dumpster for those storage groups in a cluster continuous replication (CCR) environment. For information about the Hub Transport server role, see the Microsoft TechNet article on [Planning for Hub Transport Servers](#). Capacity requirements for Hub Transport servers are affected by the following factors: message tracking logs, protocol logs, transport database, and transport dumpster.

Message tracking and protocol logs are enabled by default on Hub Transport servers. Logs are located in the following locations:

C:\Program Files\Microsoft\Exchange Server\TransportRoles\Logs.

The default log size is set to 10MB. Log files are kept until the log directory reaches the maximum directory size set or the number of days.

Best practice

Since capacity sizing for message tracking and protocol logs can be difficult, it is recommended that the maximum directory size setting be used. Size can be monitored and adjusted depending on the environment.

For details on managing transport logs, see the Microsoft TechNet article on [Managing Transport Logs](#).

Best practice

To best utilize storage, storage utilized for the Edge and Hub Transport server roles should not be placed on the same storage subsystem as the Mailbox server.

Transport Database

When deploying a CCR configuration, design of the transport dumpster should be taken into consideration. The storage capacity of the Hub Transport server should contain enough capacity to store mail long enough for all storage groups in its site, so that messages can be recovered to the passive CCR node in the event of an unscheduled outage. This feature is known as the Transport Dumpster. For more information on the Transport Dumpster and its functionality, see the Microsoft TechNet article on [Cluster Continuous Replication and the Transport Dumpster](#).

There are two settings that control how long a message remains in the Transport Dumpster: MaxdumpsterSizeperStorageGroup and MaxDumpsterTime.

Under current versions, by default, the value for the MaxdumpsterSizeperStorageGroup is set to 18. To size the Transport Dumpster, increase the largest acceptable message size by 50%. For example, if the message quota is 10MB, the MaxdumpsterSizeperStorageGroup value should be 15 MB.

For organizations with no message size limits, it is best is to set MaxdumpsterSizeperStorageGroup to a value that is 1.5 times the average message size sent within the organization.

NOTE: If message size limits are not set, there is no guarantee that all messages will be recoverable from a node failure. It is recommended that the MaxDumpster Time be set to the default value (7 days).



Capacity consumed by the dumpster should be sized at the number of storage groups with transport dumpster enabled multiplied by the maximum dumpster size. For example, if the maximum dumpster size is 15 MB, and the Hub Transport server services 100 storage groups in a CCR environment, then 1.5 GB should be allocated for the Transport Dumpster.

Transport Dumpster best practices	
Transport Dumpster settings	MaxDumpsterSizeperStorageGroup: Message size limit plus 50% MaxDumpsterTime: 7 days (default setting 0)
Capacity for dumpster database	Number of storage groups times the maximum dumpster size

Planning for I/O

When planning Exchange server environments, not only is it important to understand and properly size for storage utilization but the same must also be done for I/O. Exchange servers generate I/O through a number of different processes, which can be categorized as Transactional I/O and Non-Transactional I/O. It is important to understand the processes that generate each type of I/O and how to best configure the Exchange environment to handle those I/O requirements. Understanding the sources of Exchange disk I/O helps in planning and configuring storage layout in a way that maximizes performance. When considering the sources of Exchange disk I/O, focus most attention on the I/O behavior that is generated during log file and database file access.

Transactional I/O

Transactional I/O is caused by end users performing actions on the Exchange server. Retrieving, receiving, sending, and deleting items cause disk I/O. Under current versions, the database I/O is 8 KB in size and random, though it can be a multiple of 8 KB when the I/O can be coalesced. Outlook users that are not using Cached Mode Outlook are directly affected by poor server disk latency and this is one of the most important concerns in storage design. To prevent a poor user experience, Exchange storage has specific latency requirements for database and transaction log LUNs. The transaction log LUN should be placed on the fastest storage with a goal of less than 10 millisecond (<10ms) writes. The database LUN requires read and write response times of less than 20 milliseconds (<20ms).

Non-transactional I/O

Transactional I/O occurs in response to direct user action and message delivery. It usually has the highest priority, and is the focus for storage design. The reduction in transactional I/O makes non-transactional I/O more important. With large mailboxes, particularly in the case of the 2-GB mailbox, many enterprises are not doubling the user capacity, but in some cases increasing it tenfold. One such example would be moving from 200 MB to 2 GB. When you have such a significant increase in the amount of data on disk, you must now consider and account for non-transaction I/O when planning your storage design.

Non-transactional I/O: The Mailbox server uses more physical RAM by caching more data and thus reducing the database I/O footprint for transactional I/O (I/O generated by send/receive/client processing of e-mail). Several sources of non-transactional I/O are generated on the Mailbox server. These include:

- Online maintenance (e.g., online defrag)
- Offline maintenance (e.g., offline defrag, database repair operations)
- Backup/restore operations



- Mailbox management processes (e.g., messaging records management [MRM])
- Content indexing
- Online maintenance
- Backup and restore operations
- VSS verifications process.

All of these operations require database I/O to properly maintain and recover the server. Although Exchange Server 2007 has reduced transactional I/O significantly, adequate storage performance is still required for proper maintenance of the Mailbox server. For this reason, there comes a point of diminishing returns when adding memory to the server. In general, the purpose of adding memory to the Exchange Mailbox server is to reduce storage requirements (specifically storage subsystem performance) and thus storage cost. Due to non-transactional I/O requirements, the storage requirements of the system may not be significantly reduced by adding server memory beyond 32 GB.

Best practice
 After familiarizing yourself with the Exchange activities that generate disk I/O, you should organize your storage system in a way that maximizes performance. The following table provides a list of best practices for placement of each of your data files and planning Non- Transactional I/O.

Best practices for optimizing disk I/O	
Source of Exchange I/O	Best practice to maximize performance
Database files	All database files within a storage group on a single volume that is dedicated to these databases. Disks that hold database files should have fast random access speeds.
Content indexing files	Because the content indexing file is a random-access file, it should be placed on the same volume as the databases.
Single instance storage	To maximize the benefit of single instance storage, mailboxes that belong to the same workgroups and distribution lists should be housed on the same database if possible.
Transaction log files	For optimal recoverability, it is recommended that you place the transaction logs and databases for a given storage group onto separate aggregates and FlexVol volumes, configured as RAID-DP. In addition, it is recommended that the separate LUNs hosting the logs and database files for the same storage group do not share physical disks within the array. This is to increase the level of recoverability, allowing you to recover from multiple disk failure scenarios with minimal data loss.
SMTP queue	Place SMTP queues on any LUN or disk that is not hosting any Exchange database files. For standalone systems, a dedicated LUN or the default installation location is typically fine. In an MSCS environment, the queues must be located on a shared LUN that does not contain any Exchange database files.
Online maintenance	By default, online maintenance is scheduled to be performed daily between 1 AM and 4 AM. Special care should be taken to ensure that online maintenance does not interfere with backup and MRM operations.
Message records management (MRM)	Scheduled MRM operations cause additional I/O to the database disks and should not be run at the same time as either backup or online maintenance.
Backup operations	Backup operations should not interfere with scheduled maintenance or scheduled MRM operations.
Database verification	When using VSS backup methods, such as an SME backup operation, it is necessary to verify the databases in the backup set. The verification operation is not only CPU intensive but also causes substantial I/O to the storage subsystem.

	<p>The following techniques should be used when planning VSS verifications:</p> <ul style="list-style-type: none">• Run verification during off-peak hours• Use a verification server• Offload verification process onto separate storage• Do not delete transaction logs from non-verified backups
--	--



Server considerations

Capacity sizing is not the only thing to take into consideration when sizing Exchange Server 2007 deployments. This section provides guidelines for processor and memory considerations that provide a strong platform for Exchange server deployments.

The primary hardware difference between Exchange Server 2007 and previous versions is the move to a 64-bit platform. This change requires a new approach when selecting hardware for Exchange Server 2007. This change needs to be factored in when choosing a hardware platform for each server role.

The following table provides guidelines for choosing processor and memory configurations for each server as set forth by Microsoft. For additional information on processor and memory configurations for specific server roles, see the [Planning Processor and Memory Configurations](#).

The goal of this chart is to provide minimum, recommended, and recommended maximum configurations.

Role	Minimum	Recommended	Maximum
Edge Transport	1 x Processor Core	2 x Processor Cores	4 x Processor Cores
Hub Transport	1 x Processor Core	4 x Processor Cores	8 x Processor Cores
Client Access Server (CAS)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores
Unified Messaging Server (UM)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores
Mailbox Server	1 x Processor Core	4 x Processor Cores	8 x Processor Cores
Multi Role (Hub, CAS, UM, Mailbox)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores

Table 3) Processor configuration table for Exchange 2007.

Effect of host-side memory on storage

As a result of larger amounts of memory available in the 64-bit platform, the Exchange Server 2007 database cache can be larger. This can reduce the number of reads and writes to disk. In Exchange Server 2003, the maximum was 900 MB of database cache. With additional memory available in Exchange Server 2007, this amount can increase from 900 MB to multiple GBs, depending on the amount of server system memory installed. As additional memory is added to the host server, the Exchange server database cache can increase, further decreasing the number of disk reads. The increased sizes of database cache make understanding the effects of host-side memory an important concept when planning storage and IOPS.

Increasing the amount of database cache also affects IOPS. When testing 4,000 users with 8 GB of server memory, the database cache is limited to 1.48 MB/user [(8 GB – 2048) / 4000]. Increasing server memory to 16 GB increases the amount of database cache to 3.48 MB/user [(16 GB – 2048) / 4000].

Mailbox role

Memory configuration for the Mailbox role depends largely upon the mailbox count and the client profile. Memory sizing for the Mailbox role is critical to reducing the I/O requirements of the server. The more memory you add to the Mailbox server, the larger the database cache size. Larger database cache sizes will result in less I/O being generated by the Exchange databases. Because I/O requirements are affected by memory, it is important to note that defining memory requirements of the Exchange 2007 Mailbox server is required prior to defining the storage requirements/configuration. The following table can be used to assist in estimating the memory requirements of a given mailbox.

User type	Send/receive per day	Mailbox memory recommendation
Light	5 sent/20 received	2 GB + 2 MB/Mailbox
Average	10 sent/40 received	2 GB + 3 ½ MB/Mailbox
Heavy	20 sent/80 received	2 GB + 5 MB/Mailbox
Very Heavy	30 sent/120 received	2 GB + 5 MB/Mailbox

Table 4) Memory requirements per mailbox type.

NOTE: The table above assumes an average message size of 50 KB.

Minimum memory configuration for mailbox based on number of storage groups

Each storage group requires a minimum amount of memory to function properly. The maximum number of Storage Groups configurable in Exchange 2007 has been increased to 50 in the Enterprise Edition. This increase provides greater flexibility in server/storage architecture, but the increase has a significant effect on the memory utilization of the Exchange 2007 Mailbox server. As the number of storage groups used increases, the amount of memory required also increases. The following table provides minimum memory requirements based on storage group counts.

Storage group count	Minimum required physical memory
1-4	2 GB
5-8	4 GB
9-12	6 GB
13-16	8 GB
17-20	10 GB
21-24	12 GB
25-28	14 GB
29-32	16 GB
33-36	18 GB
37-40	20 GB
41-44	22 GB
45-48	24 GB
49-50	26 GB

Table 5) Minimum memory requirements.



Example 1: A 4,000-user Mailbox server with a heavy user profile would calculate out to 22 GB of RAM (2048 MB + (4000 * 5 MB)). Based on the above support requirements, the administrator will have the flexibility to use up to 44 Storage Groups. Additional RAM would be required to deploy more than 44 Storage Groups based on the above supportability requirements.

Example 2: A 1,000-user Mailbox server with a light user profile would calculate out to 4GB of RAM (2048 MB + (1000 * 2 MB)). Based on the above support requirements the administrator will have the flexibility to use up to 8 Storage Groups. Additional RAM would be required to deploy more than 8 Storage Groups based on the above supportability requirements.

Memory recommendations for log shipping

When using Exchange 2007 log shipping capabilities, it is recommended that an additional 1 GB of physical memory be provisioned on each Mailbox server and Multi-Role servers.

Sizing additional processing capacity for local continuous replication (LCR)

Local Continuous Replication (LCR) has all of the Exchange 2007 Mailbox role services as well as the Exchange 2007 Replication Service running on the same server. On an LCR Mailbox server, there is additional processing overhead generated from the Replication Service copying and playing in logs to the target database copy. This additional processing cost is roughly 20% and should be factored in when sizing LCR Mailbox servers.

Storage monitoring

Once the storage plan has been implemented, ongoing configuration control is necessary to maintain a stable Exchange environment and ensure the deployment is meeting the capacity and I/O needs placed on the Exchange servers.

The [Exchange 2007 Management Pack for MOM 2005 SP1](#) combined with the [Performance Troubleshooter](#), built in to the Exchange 2007 Management Console, can be used to determine current capacity utilization and performance. It can also provide metrics as to when a given deployment requires additional Hub, Edge, CAS, or Mailbox server roles based on performance. This method can be used to fine-tune the server role ratios for a specific deployment.

Conclusion

Microsoft Exchange Server 2007 is not a one-size-fits-all application. There are multiple configuration options available to suit most needs for any customer. IBM N series storage appliances and software are built in a similar fashion, providing users with the flexibility to manage Exchange data in a manner that best meets their business requirements. With high performance, easy to manage storage appliances, and robust software offerings, N series offers the best end to end storage solutions for Exchange Server 2007.

The best practices and recommendations set forth in this guide are also not a one-size-fits-all solution. This document contains a collection of best practices and recommendations that provide a guideline to plan, deploy, and manage Exchange data. Following these guidelines will ensure a highly available, easy to manage Exchange environment that will meet SLAs. Please consult with your local IBM N series Exchange expert to maximize best-practice benefit realization.



Appendix A: Best-practices summary

SnapDrive

Best practice

Use SnapDrive to create and manage N series LUNs. SnapDrive handles the necessary tasks (partition, align, format) to properly create and mount a LUN on a given Exchange server.

SnapManager for Exchange

Best practice

Use SnapManager for Exchange when deploying Exchange Server 2007 into N series storage. SME will perform the data migration from local disks to N series LUNs. It will also manage that data, handling the backup, restore, and verification tasks.

Disk types

Best practice

FC disks are the recommended choice to meet the needs for Exchange 2007. However, with SnapVault, ATA disks can also meet the needs for Exchange 2007 in less I/O intensive deployments. To determine the I/O requirement for your Exchange environment and the best disk-type match, please consult with your local IBM N series Exchange expert.

RAID-DP

Best practice

Always use RAID-DP when deploying Exchange server data onto an N series storage appliance. RAID-DP offers the best combination of performance and data protection for Exchange server environments.



Aggregate recommendations

Best practice

When using aggregates, it is recommended to separate the database flexible volumes and LUNs from the transaction log flexible volumes and LUNs onto different aggregates.

Volume recommendations

Best practice

When using aggregates, it is recommended that the database flexible volumes and LUNs are separate from the transaction log flexible volumes and LUNs. In the unlikely event that an aggregate is lost, Exchange data is still available for disaster recovery.

Best practice

It is recommended that there are different database and transaction log volumes for different Exchange servers, to prevent a potential “busy” snapshot copy problem. Because there are separate volumes for each server, there is no need for concern regarding snapshot schedules overlapping each other for different servers.

Best practice

Place database LUNs and transaction log/SnapInfo LUNs in separate volumes.

Best practice

If there are separate LUNs for the Exchange transaction log files and the SnapInfo directory, place those LUNs in the same volume. Both of these LUNs will have a similar I/O profile, allowing them to share the same volume. And for disaster recovery scenarios, having the entire log set for Exchange on the same volume will help achieve SLAs.

LUN recommendations

Best practice

Separate Exchange database and transaction log files onto separate LUNs and separate volumes. This allows for greater flexibility for backup and recovery procedures and data protection strategies.

NTFS hard links

Best practice

Utilize NTFS hard links whenever possible to increase storage utilization, eliminate the physical copy overhead on the Exchange server, and increase backup performance. SME will automatically utilize NTFS hard links when it detects the transaction log file directory and the SnapInfo directory are on the same NTFS volume.

Mapping Exchange users to storage groups

Best practice

Place user mailboxes that have stringent SLA requirements in a separate storage group. That storage group should have its own dedicated set of LUNs residing on a separate volume. Configuring the storage in such a way enables much more flexibility in managing those user mailboxes and protecting them as required by the SLA requirements.



Sizing Exchange Server 2007 deployments

Best practice

The IBM N series Sizing Tool for Exchange should be used to size all Exchange server deployments utilizing N series storage.

Maintenance

Best practice

When planning maintenance of Exchange databases, keep in mind the following:

- Always plan maintenance windows during off-peak or low-activity times.
- Never schedule a maintenance window during an SME backup operation.
- Never schedule a maintenance window during a verification operation if both operations are performed on the same server.

Capacity planning for Exchange 2007 Server roles

Best practice

Since capacity sizing for message tracking and protocol logs can be difficult it is recommended that the maximum directory size setting be used. Size can be monitored and adjusted depending on the environment.

Best practice

To best utilize storage, storage utilized for the Edge and Hub Transport server roles should not be placed on **the same storage subsystem as the Mailbox server.**

Transport Dumpster Best practices

Transport Dumpster settings	MaxDumpsterSizeperStorageGroup: Message size limit plus 50% MaxDumpsterTime: 7 days (default setting 0)
Capacity for dumpster database	Number of storage groups multiplied by maximum dumpster size



I/O planning best practices

Best practice

After familiarizing yourself with the Exchange activities that generate disk I/O, you should organize your storage system in a way that maximizes performance. The following table provides a list of best practices for placement of each of your data files and planning Non- Transactional I/O.

Best practices for optimizing Disk I/O	
Source of Exchange I/O	Best practice to maximize performance
Database files	All database files within a storage group on a single volume that is dedicated to these databases. Disks that hold database files should have fast random access speeds.
Content indexing files	Because the content indexing file is a random-access file, it should be placed on the same volume as the databases.
Single instance storage	To maximize the benefit of single instance storage, mailboxes that belong to the same workgroups and distribution lists should be housed on the same database if possible.
Transaction log files	For optimal recoverability, it is recommended that you place the transaction logs and databases for a given storage group onto separate aggregates and FlexVol volumes, configured as RAID-DP. In addition, it is recommended that the separate LUNs hosting the logs and database files for the same storage group do not share physical disks within the array. This is to increase the level of recoverability, allowing you to recover from multiple disk failure scenarios with minimal data loss.
SMTP queue	Place SMTP queues on any LUN or disk that is not hosting any Exchange database files. For standalone systems, a dedicated LUN or the default installation location is typically fine. In an MSCS environment, the queues must be located on a shared LUN that does not contain any Exchange database files. You can also place the SMTP and MTA queues on a transaction log or on the SnapInfo LUN.
Online maintenance	By default, online maintenance is scheduled to be performed daily between 1 AM and 4 AM. Special care should be taken to ensure that online maintenance does not interfere with backup and MRM operations.
Message records management (MRM)	Scheduled MRM operations cause additional I/O to the database disks and should not be run at the same time as either backup or online maintenance.
Backup operations	Backup operations should not interfere with scheduled maintenance or scheduled MRM operations.
Database verification	When using VSS backup methods, such as an SME backup operation, it is necessary to verify the databases in the backup set. The verification operation is not only CPU intensive but also causes substantial I/O to the storage subsystem. The following techniques should be used when planning VSS verifications: <ul style="list-style-type: none"> • Run verification during off-peak hours • Use a verification server • Offload verification process onto separate storage • Do not delete transaction logs from non-verified backups



Processor configuration best practices

Role	Minimum	Recommended	Maximum
Edge Transport	1 x Processor Core	2 x Processor Cores	4 x Processor Cores
Hub Transport	1 x Processor Core	4 x Processor Cores	8 x Processor Cores
Client Access Server (CAS)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores
Unified Messaging Server (UM)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores
Mailbox Server	1 x Processor Core	4 x Processor Cores	8 x Processor Cores
Multi Role (Hub, CAS, UM, Mailbox)	1 x Processor Core	4 x Processor Cores	4 x Processor Cores

Mailbox role memory configuration

User Type	Send/Receive per Day	Mailbox Memory Recommendation
Light	5 sent/20 received	2GB + 2MB/Mailbox
Average	10 sent/40 received	2GB + 3 ½MB/Mailbox
Heavy	20 sent/80 received	2GB + 5MB/Mailbox
Very Heavy	30 sent/120 received	2GB + 5MB/Mailbox



Trademarks and special notices

© International Business Machines 1994-2008. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Data ONTAP, FlexClone, NearStore, Network Appliance, the Network Appliance logo, RAID-DP, SnapDrive, SnapManager, Snapshot and SnapVault are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.