



**NetApp™**  
Go further, faster

NETAPP TECHNICAL REPORT

# Microsoft Exchange Server 2007 Best Practices Guide

Brad Garvey, NetApp

**March, 2009 | TR-3578**

**Microsoft®**  
**GOLD CERTIFIED**  

---

*Partner*

## TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION</b>	<b>4</b>
1.1	PURPOSE AND SCOPE	4
1.2	INTENDED AUDIENCE	4
1.3	CAVEATS	4
<b>2</b>	<b>NETAPP SOLUTIONS FOR MICROSOFT EXCHANGE SERVER 2007</b>	<b>4</b>
2.1	STORAGE AND DATA MANAGEMENT SOFTWARE	5
2.1.1	SnapDrive® for Windows	5
2.1.2	SnapManager® for Microsoft Exchange	5
2.1.3	Single Mailbox Recovery	6
2.2	ENTERPRISE STORAGE SYSTEM COMPONENTS	6
2.2.1	Disk Types	6
2.2.2	RAID	7
2.2.3	NetApp Aggregates, Volumes, and LUNs	8
<b>3</b>	<b>STORAGE SYSTEM VOLUME SIZING</b>	<b>15</b>
<b>4</b>	<b>SIZING EXCHANGE SERVER 2007 DEPLOYMENTS</b>	<b>18</b>
<b>5</b>	<b>PLANNING STORAGE CAPACITY</b>	<b>19</b>
5.1	KEY EXCHANGE SERVER FACTORS	19
5.1.1	Mailbox Size/Quota	19
5.1.2	Database Dumpster	19
5.1.3	Context Indexing	19
5.2	MAINTENANCE AND MOVING MAILBOXES	19
5.2.1	Maintenance	19
5.2.2	Moving Mailboxes	20
<b>6</b>	<b>CAPACITY PLANNING FOR EXCHANGE SERVER 2007 SERVER ROLES</b>	<b>20</b>
6.1	EDGE TRANSPORT	20
6.2	HUB TRANSPORT	21
<b>7</b>	<b>PLANNING FOR I/O</b>	<b>22</b>
7.1	TRANSACTIONAL I/O	22
7.2	NON-TRANSACTIONAL I/O	22
<b>8</b>	<b>SERVER CONSIDERATIONS</b>	<b>24</b>
8.1	EFFECT OF HOST SIDE MEMORY ON STORAGE	25
<b>9</b>	<b>CONTINUOUS REPLICATION</b>	<b>26</b>
9.1	OVERVIEW	26
9.2	STORAGE BEST PRACTICES FOR CONTINUOUS REPLICATION	27
9.3	BACKUP CONSIDERATIONS FOR CONTINUOUS REPLICATION	27
<b>10</b>	<b>SERVICE PACK 1</b>	<b>28</b>

10.1	EXCHANGE SP1 TRANSPORT FEATURES .....	28
10.2	MAILBOX SERVER ROLE .....	29
10.3	HIGH AVAILABILITY FEATURES.....	30
11	WINDOWS 2008 .....	30
12	STORAGE MONITORING .....	31
13	CONCLUSION .....	31
14	APPENDIX A – BEST PRACTICES SUMMARY .....	32

## 1 INTRODUCTION

The availability and performance of enterprise messaging can have a significant impact on productivity across the organization. For this reason, enterprise messaging is viewed as a mission critical application for many organizations today. And for most, Microsoft® Exchange Server has emerged as the enterprise messaging platform of choice.

The critical importance of enterprise messaging has led to some notable challenges for information technology (IT) managers. 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.

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.

### 1.1 PURPOSE AND SCOPE

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 NetApp Exchange expert when planning and deploying Exchange Server environments when using NetApp storage solutions.

The following areas will be covered:

- NetApp 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.

### 1.2 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.

### 1.3 CAVEATS

The best practices for Exchange Server 2007 storage architecture presented are focused exclusively on the latest NetApp storage operating system, Data ONTAP® 7G.

Guidance for Unified Messaging in Exchange Server 2007 is not included in this document.

## 2 NETAPP SOLUTIONS FOR MICROSOFT EXCHANGE SERVER 2007

When deploying Exchange Server 2007 onto a NetApp storage appliance, it is important to understand the key components involved and how to properly utilize them. The hardware and software components incorporated into NetApp solutions for Microsoft Exchange Server 2007

include enterprise storage systems and data management software. Specific components are described below.

## 2.1 STORAGE AND DATA MANAGEMENT SOFTWARE

### 2.1.1 SnapDrive® for Windows

NetApp 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 Snapshot™ copies, rapid application recovery, and overall data storage management for Windows servers.

When provisioning NetApp 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 NetApp LUNs. SnapDrive handles the necessary tasks (partition, align, format) to properly create and mount a LUN on a given Exchange server.**

**Note:** SnapDrive supports the creation of up to 120 mount points.

For more information on SnapDrive for Windows, please visit the [SnapDrive for Windows](#) page on the NetApp Website.

### 2.1.2 SnapManager® for Microsoft Exchange

NetApp SnapManager 4.0 for Microsoft Exchange (SME) 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 a 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 NetApp 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 KB 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 NetApp storage. SME will perform the data migration from local disks to NetApp LUNs. It will also**

**manage that data, handling all backup, restore, and verification tasks.**

For more information on SnapManager for Microsoft Exchange, please visit the [SnapManager for Exchange](#) page on the NetApp Website. Also, further best practices for SME can be found at: [TR-3541: SnapManager 4.0 for Microsoft Exchange: Best Practices Guide](#).

### **2.1.3 Single Mailbox Recovery**

[NetApp 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 a 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.

## **2.2 ENTERPRISE STORAGE SYSTEM COMPONENTS**

### **2.2.1 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. Fibre Channel 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 NetApp 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.

If users have a requirement to deploy Exchange 2007 on a NetApp storage appliance using ATA drives, please consult a local NetApp Exchange expert. The expert will provide guidance on the deployment to ensure that items like I/O latencies and SLAs are in line with the requirements of Exchange Server 2007 and Exchange administration staff.

#### Best Practice

**FC disks are the recommended choice to meet the needs for Exchange 2007. However, SATA disks can also meet the needs for Exchange 2007 in certain less I/O intensive deployments. To determine the I/O requirement for an Exchange environment, please consult with your local NetApp Exchange expert.**

## 2.2.2 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 NetApp double parity RAID solution, 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

NetApp [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. A detailed paper clearly identifying those benefits can be located on the Technical Reports section of the [NetApp Library – TR3574 – Using RAID-DP in Exchange Server 2007 Storage Designs](#).

### 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

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 NetApp [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 a NetApp storage appliance. RAID-DP offers the best combination of performance and data protection for Exchange server environments.**

### **2.2.3 NetApp Aggregates, Volumes, and LUNs**

Data ONTAP 7G provides functionality that enables the creation of flexible volumes for managing data without the need to assign physical disks to the volumes. Instead, the 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.



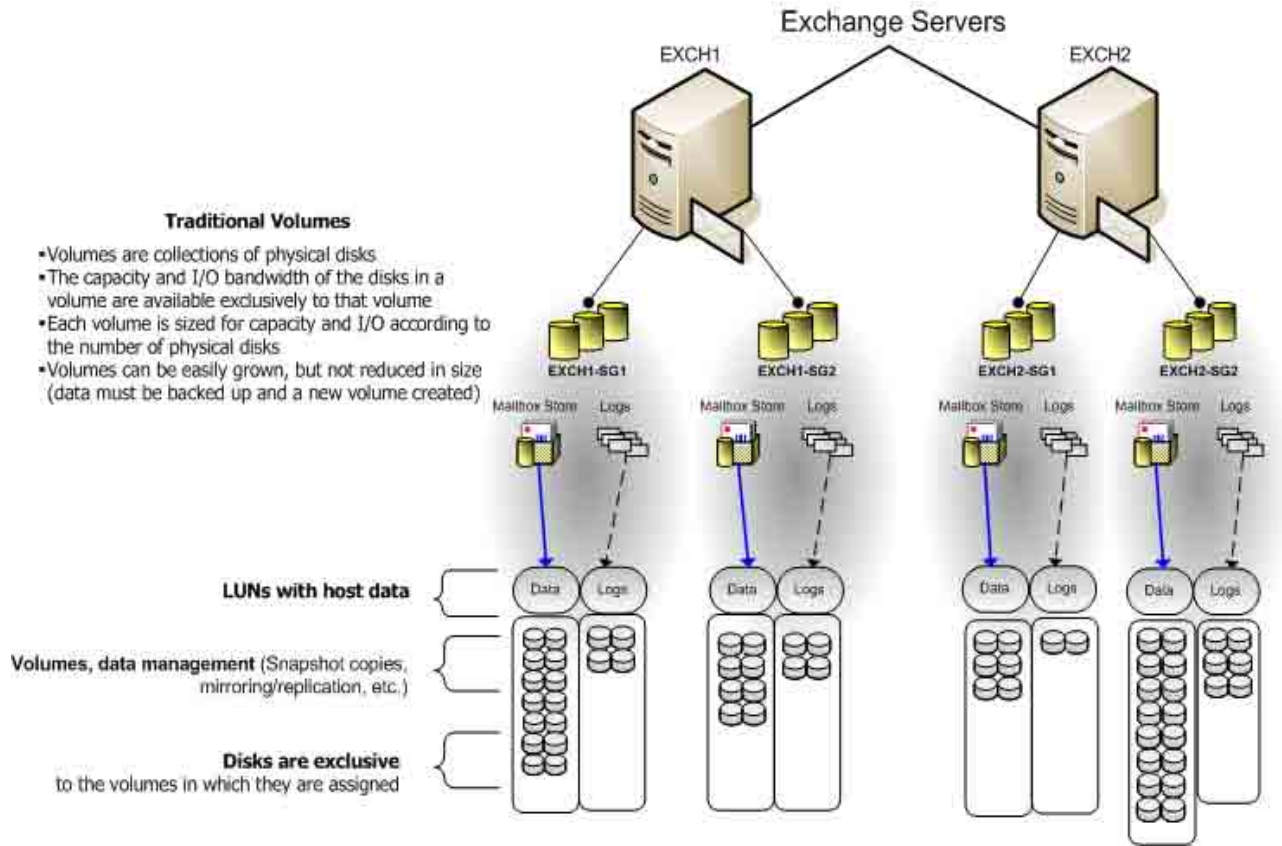
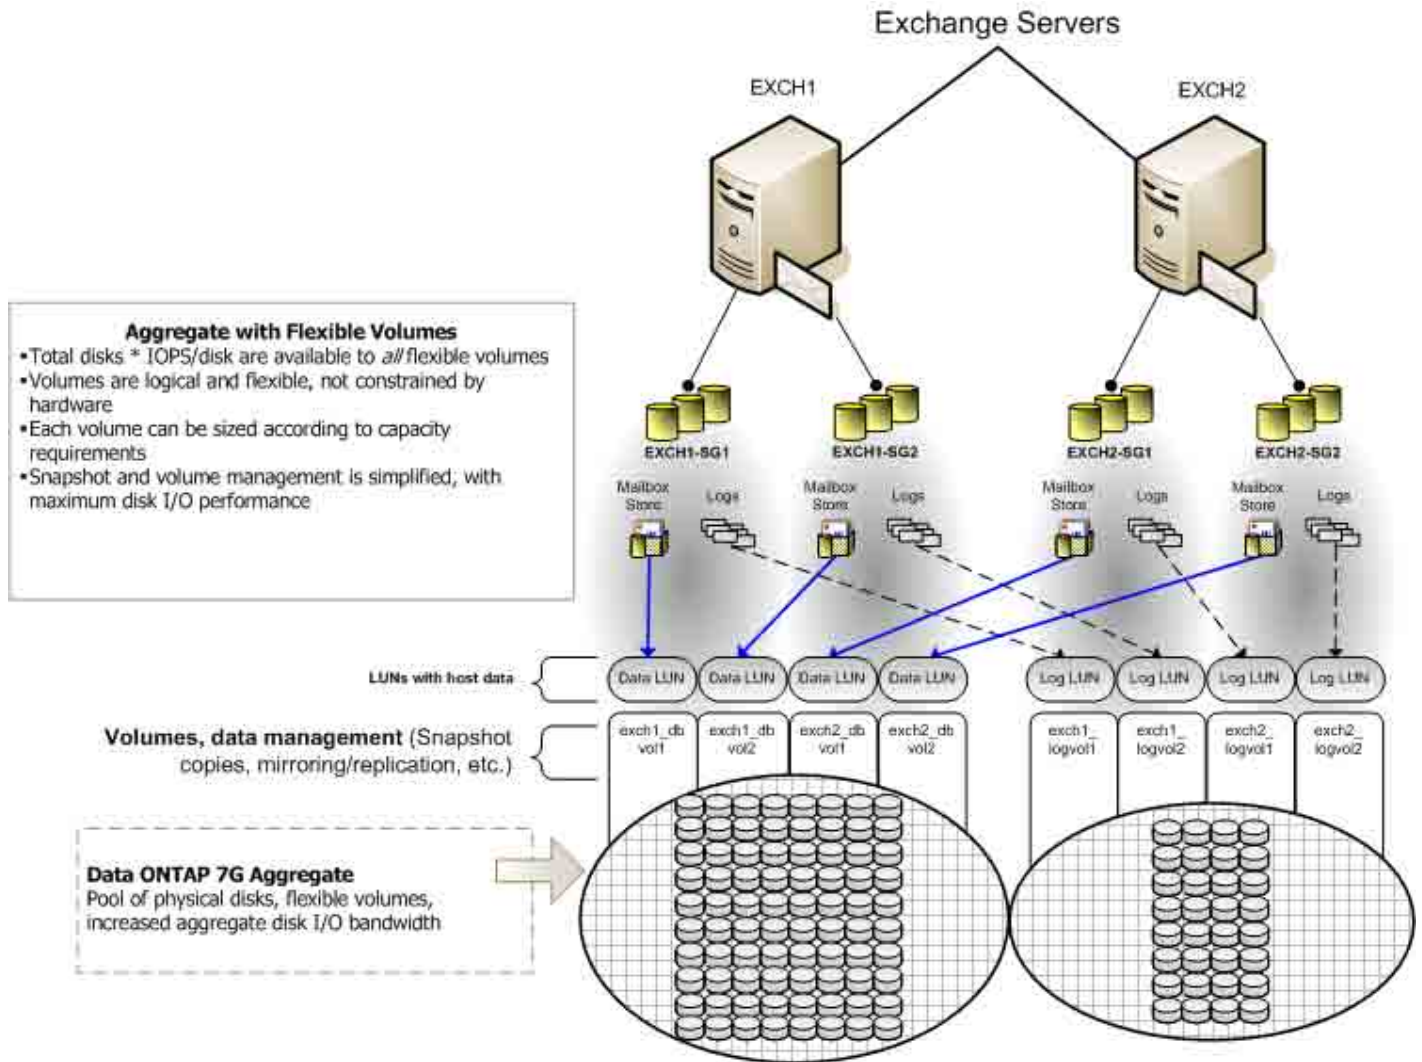


Figure 1) Using Data ONTAP 7G with traditional volumes.

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 with flexible volumes.**

Data ONTAP 7G provides flexible options that enable Exchange administrators to efficiently utilize all available disks attached to NetApp platforms. The benefits are easily realized for all Exchange servers that are storing data on one or more NetApp systems. Traditional volumes or flexible volumes can be used, depending on the operational requirements and performance needs of small to very large enterprise deployments.

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

For detailed calculations on volume and LUN sizing, please see the [Creating Storage System Volumes and LUNs chapter in the SME Installation and Administration Guide](#).

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

Table 2 Aggregate Layouts

#### Best Practice

**It is recommended to have at least 10% free space available in an aggregate hosting Exchange data. This allows for optimal performance of the storage system.**

Consult a local NetApp 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 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. For more information about busy Snapshot copy issues, see [Chapter 12 of the SME Installation and Administrator Guide](#).

#### 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.

#### Best Practice

It is recommended to have at least 10% free space available in a volume hosting Exchange data. This allows for optimal performance of the storage system.

### LUN Recommendations

#### Best Practice

When possible, it is recommended to 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.

#### Best Practice

Microsoft recommends approximately 20% free disk space for optimal Windows performance. This 20% free disk space can help prevent an Exchange outage by providing additional storage space in the event additional Exchange data is written to a LUN and avoiding an out of space condition for Exchange, thus taking the affected storage group offline. Please refer to the Microsoft article [Deploying a Large Exchange Server 2007 Organization](#).

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 NetApp 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 2TB limit on master boot record (MBR) partitions, scaling the storage to meet the needs of a particular storage group will be challenging.

NetApp 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/volumes that have to be managed is a key component of a successful Exchange deployment.

When planning how Exchange data will be deployed onto NetApp 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 NetApp Exchange expert for guidance on deploying an Exchange environment.

### **SnapManager for Exchange Server Metadata Repository Location**

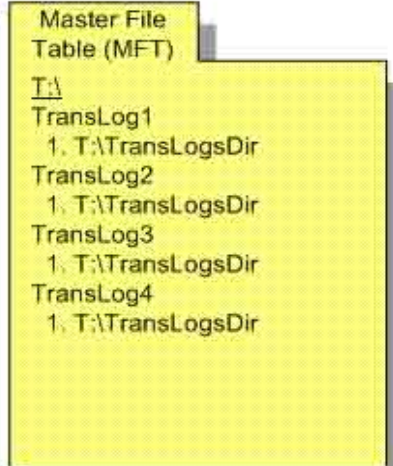
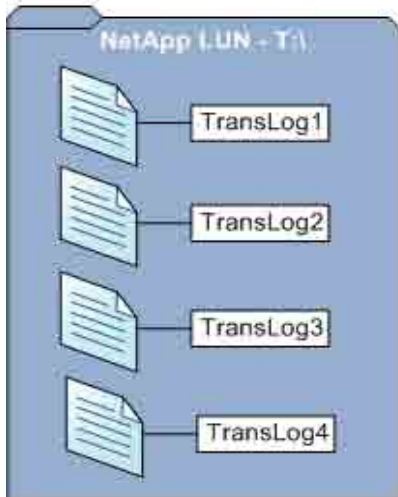
The number of transaction logs generated for a storage group directly affects the size of the SME SnapInfo directory. When the transaction log directory resides on a different NTFS volume from the SnapInfo directory, SME will perform a file copy for each transaction log, copying the logs into the SnapInfo 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. 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**

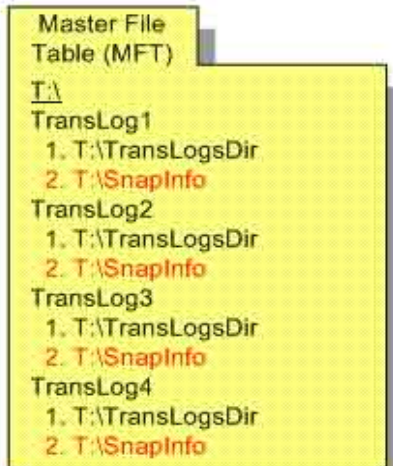

**Utilize NTFS hard links whenever possible to increase storage utilization, eliminate the physical copy overhead on the Exchange server, and increase backup performance. SME will automatically utilize NTFS hard links when it detects the transaction log file directory and the SnapInfo directory on the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

### After the initial SME configuration



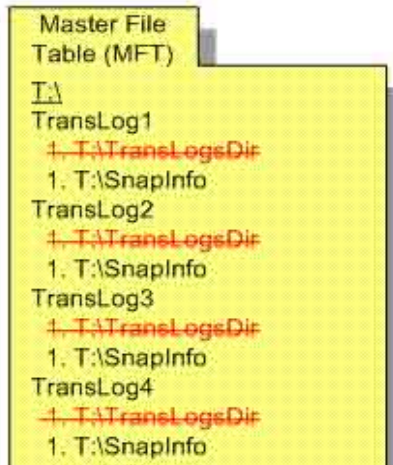
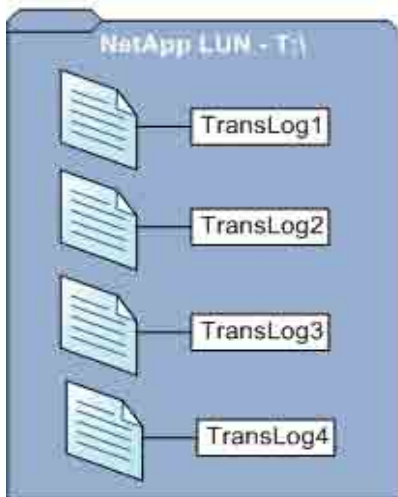
During the initial SME Configuration, the Master File Table (MFT), for the destination transaction log LUN and corresponding NTFS volume, is updated with the location of the transaction log directory for each transaction log that is written to that directory.

### During a backup operation



When the Exchange transaction log directory resides on the same NTFS volume as the SnapInfo directory, NTFS hard links are automatically used. The MFT record for each file is updated during a backup operation to include a directory location for the SnapInfo directory. The actual file on the volume is not modified in any way.

### After the backup operation has completed



Once the backup operation has completed, the NTFS hard links in the Master File Table are updated again, deleting the references for the Exchange transaction log directory, if SME is configured to do so. The SnapInfo directory reference remains. This eliminates the need to physically copy the transaction log files to the SnapInfo directory. Also, the physical file is not modified in any way, so no blocks are modified on the storage appliance.

Figure 3) NTFS hard links.

### 3 STORAGE SYSTEM VOLUME SIZING

Volume sizing is broken into two different parts, database volume sizing and transaction log volume sizing. NetApp recommends that when sizing Exchange volumes the Microsoft sizing spreadsheet for Exchange first be used in order to determine appropriate amount of space needed for each LUN. Total LUN Requirements are used as the basis for calculating volume space requirements.

#### Transaction Log Volume Sizing

Providing accurate sizing for transaction log volumes depends on the following factors.

- Total Transaction log LUN size = Total size of all transaction log LUNs that will be stored in one transaction log volume
- Snapshot copy space = Space consumed by transaction logs generated during a 24 hour period
- Online Backup Retention Duration = Number of day that backups are kept online. A day is measure by 24 hours

Transaction log volume sizing can be calculated using the following formula.

$$\text{Transaction Log Volume size} = \text{Total Transaction Log LUN size} + (\text{snapshot copy space} * \text{Online Backup Retention Duration})$$

#### Database Volume Sizing

Providing that an accurate change rate is known the following formula would be used in calculating total volume size. The formula below is based on 2 different key variables.

To calculate the Exchange DB volume size a number of variable are used:

- Database LUN Size = The size the LUN used to store Exchange Mailbox or Public Folder Database
- Database daily change rate = The amount the Exchange Mailbox or Public Folder Database changes in a day expressed as a percentage value of the database size.
- Online Backup Retention Duration = Number of day that backups are kept online. A day is measure by 24 hours
- Fault Tolerance Window = This is the number of days that backup failures can be tolerated before running out of Snapshot copy space. A typical value for this is 2 – 4.

Database volume size can be calculated using the following formula

$$\text{Database Volume Size} = (\text{Sum of the database LUN sizes that will share the database volume}) + ((\text{Fault Tolerance Window} + \text{Online Backup Retention Duration}) * \text{database daily change rate})$$

## Database Daily Change Rate

An accurate rate of change is a critical part of sizing. Rate of change will differ between environments and specific numbers should be used whenever possible. The rate of change can not only vary depending on the number of messages sent and received but other factors such as online maintenance.

Online maintenance process can have a large impact on the amount of space consumed by Snapshot copies. The online defrag portion of the online maintenance process changes pages in the database. The more databases pages that are changed by this process, the more space needed to store the snapshot copies, thus affecting the overall rate of change.

Microsoft recommends that 1 online defragmentation pass per mailbox database be completed in 14 days, upon checking metrics and finding that defragmentation is completing on a more aggressive schedule, you can shorten the length of the daily online maintenance window to reduce the amount of daily database churn, which in turn reduces Snapshot space consumption.

The following performance counters can be tracked to determine your online defragmentation trend:

- MSEXchange Database ==> Instances\Online Defrag Pages Freed/Sec
- MSEXchange Database ==> Instances\Online Defrag Pages Read/sec.

If the ratio of reads to freed pages is greater than 100 to 1, reduce the online maintenance window. If the read-to-freed ratio is less than 50 to 1, increase the online maintenance window.

## Snapshot Copy Space Management

When sizing storage system volumes for Exchange server use NetApp recommends setting proper snapshot copy space and retention management policies for volumes hosting LUNs containing Exchange data. This guarantees sufficient space for proper Exchange Server databases operation and online backup retention.

The following storage system volume and snapshot copy space management parameters are recommended.

- Guarantee = volume
- LUN reservation = on
- Fractional\_reserve = 0%
- Snap\_reserve = 0%
- Auto\_delete = volume
- Auto\_grow = off (This can also be set to on depending on environment)
- Try\_first = snap\_delete

Snapshot auto delete settings

- state =on
- Trigger =volume
- Delete\_order =oldest\_first
- Defer\_delete =prefix
- Prefix =exchsnap



- Target free space = the target free space setting is the amount of space that ONTAP will free by deleting snapshots with the autodelete function. By default this value is set to 20%, this means that if auto delete is invoked by a low space condition, that autodelete will delete snapshots until the volume has 20% free space. It is recommended that this value be adjusted to 2% greater than the auto delete threshold.

The autodelete threshold setting is based on the volume size as follows.

Volume Size	Trigger
< 20GB	85%
20 < 100GB	90%
100 < 500GB	92%
500 < 1TB	95%
= >1TB	98%

Example:

For a volume size of 500GB the target\_free\_space setting would be adjusted to 93%.

### Monitoring Volume Free Space and Rate of Change

Ongoing space monitoring is key to ensure enough volume free space is available. The amount of volume free space and the daily change rate should be monitored on a regular basis since these change as use of the environment changes over time.

The amount of free space in a volume can be viewed by using the following command from the ONTAP command line interface.

```
df -g <vol_name>
```

To accurately monitor the rate of change, the `snap delta` command can be used to get the size of each Snapshot copy. The sizes of each Snapshot copy will aid in estimating how much data is changing and if adjustments need to be made in volume sizing to accommodate the rate of change.

In order to correctly size storage and correctly use space reservation as always it is recommended to work with your local NetApp professional to ensure your storage is sized correctly, all factors are taken into consideration, and a successful fractional space reservation policy is deployed

### 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. [NetApp NearStore® storage appliances](#) provide an ideal, cost-efficient storage solution for archiving older Snapshot copies. Using [NetApp 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.

## 4 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:

1. The Microsoft storage calculator
2. The NetApp 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 NetApp storage it is important the NetApp Exchange Sizing Tool be used.

### Best Practice

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

Please consult a local NetApp Exchange expert to provide accurate sizing for Exchange environments.

## 5 PLANNING STORAGE CAPACITY

### 5.1 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.

#### 5.1.1 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.

#### 5.1.2 Database Dumpster

For each Exchange database, there is a database dumpster that stores hard-deleted items. These items will be retained, by default, 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.

#### 5.1.3 Context 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 Context Indexing for a given database, but by default in Exchange 2007 it is turned on.

### 5.2 MAINTENANCE AND MOVING MAILBOXES

#### 5.2.1 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.**

## 5.2.2 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 additional storage to handle user mailbox moves will help prevent the transaction log LUN from running out of space and causing downtime. Understanding the migration process for an Exchange server will help determine how much additional storage will be required to handle 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.

# 6 CAPACITY PLANNING FOR EXCHANGE SERVER 2007 SERVER ROLES

## 6.1 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.

A minimum of 4GB 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 the following factors:

- Message Tracking Logs
- Protocol Logs
- Mail Database
- 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, and the capacity reserved should be the average message size multiplied by the maximum queue, in the case where the queue is at maximum and the server is shut down. A 500,000-item queue with an average message size of 50KB is approximately 25GB of data in the database.

Edge Transport servers that run anti-virus scan 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. 1GB is a good starting point, although each organization's actual needs will be different.

## 6.2 HUB TRANSPORT

The Hub Transport server routes mail to the proper mailbox server and it also includes a transport dumpster for those storage groups in a cluster continuous replication (CCR) environment. For more 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
- 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 additional information 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
- MaxDumpsterTime

By default the value for the MaxdumpsterSizeperStorageGroup is set to 18. To size the transport dumpster properly, take your largest acceptable message size, and increase that number by 50%. For example, if the message quota is 10MB, the MaxdumpsterSizeperStorageGroup value should be set to 15MB.

For organizations that have not implemented message size limits, it is a best practice to set the 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, it cannot be guaranteed that all messages will be recovered in the event of a node failure.

It is recommended that the MaxDumpster Time be set to the default value of seven 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 15MB, and the Hub Transport server services 100 storage groups in a CCR environment, then 1.5GB 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 multiplied by maximum dumpster size

## 7 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.

### 7.1 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. The database I/O is 8KB in size and random, though it can be a multiple of 8KB 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).

### 7.2 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 2GB mailbox, many enterprises are not doubling the user capacity, but in some cases increasing it tenfold. One such example would be moving from 200MB to 2GB. 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 utilizes additional physical RAM by caching more data and thus reducing the database I/O footprint for transactional I/O (I/O that is generated by send/receive/client processing of e-mail). There are several sources of non-transactional I/O 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 32GB.

#### 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.

<b>Message records management (MRM)</b>	<b>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.</b>
<b>Backup operations</b>	<b>Backup operations should not interfere with scheduled maintenance or scheduled MRM operations.</b>
<b>Database verification</b>	<p><b>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.</b></p> <p><b>The following techniques should be used when planning VSS verifications:</b></p> <ul style="list-style-type: none"> <li>• <b>Run verification during off-peak hours</b></li> <li>• <b>Use a verification server</b></li> <li>• <b>Offload verification process onto separate storage</b></li> <li>• <b>Do not delete transaction logs from non-verified backups</b></li> </ul> <p><b>For additional information on database verification options refer to the chapter on <a href="#">Selecting the Database Verification Server in the SME Installation and Administration Guide</a>.</b></p>

## 8 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.

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

**Table 3 Processor Configuration Table for Exchange 2007**



**NOTE:** Minimum hardware requirements must be met to receive Microsoft Product Support Services.

## 8.1 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 900MB of database cache. With additional memory available in Exchange Server 2007, this amount can increase from 900MB 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 8GB of server memory, the database cache is limited to 1.48MB/user [ ( 8GB – 2048 ) / 4000 ]. Increasing the amount of server memory to 16GB increases the amount of database cache to 3.48MB/user [ ( 16GB – 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	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

Table 4 Memory requirements per mailbox type

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

### Required Minimum Memory Configuration for Mailbox Based on the 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	2GB
5-8	4GB
9-12	6GB
13-16	8GB
17-20	10GB
21-24	12GB
25-28	14GB
29-32	16GB
33-36	18GB
37-40	20GB
41-44	22GB
45-48	24GB
49-50	26GB

**Table 5 Minimum memory requirements**

**Example 1:** A 4,000-user Mailbox server with a heavy user profile would calculate out to 22GB of RAM ( 2048MB + ( 4000 \* 5MB ) ). 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 ( 2048MB + ( 1000 \* 2MB ) ). 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 1GB 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.

## **9 CONTINUOUS REPLICATION**

### **9.1 OVERVIEW**

With Exchange 2007 the concept of log shipping or continuous replication was introduced. Continuous replication is an availability feature which comes in two forms:

1. Local Continuous Replication (LCR)
2. Clustered Continuous Replication (CCR)

Local Continuous Replication provides storage level resiliency while CCR provides server and storage level resiliency. Another benefit of continuous replication is the Exchange database can be backed up from the target node via a Volume Shadow Copy Service (VSS) agent or requester. This offloads disk I/O incurred during the backup and integrity check processes from the active spindles.

Log shipping consists of two storage locations, an active and target location. Log shipping works by initially seeding or copying the database to the target location. When a log file is full, it is renamed from the active Exx.log to the next log file in the sequence, such as Exx00000002.log. The target then pulls or copies the closed log file from the source, performs a transaction log verification, and then replays the log file into the target database. In the event of a storage failure on the active storage, the target storage becomes the active, thus returning exchange to a functional state very quickly incurring minimal data loss.

## **9.2 STORAGE BEST PRACTICES FOR CONTINUOUS REPLICATION**

There are several considerations to keep in mind when designing storage solutions for continuous replication.

### **Storage Design**

Designing storage for continuous replication environments differs slightly from that of traditional Exchange Server storage design. As best practice the following rules should be followed when designing storage.

#### **Size for Performance of Target database LUNS**

Because of the intensity of the transaction log verification and replay disk I/O on the target database LUNS can be 2 to 3 times that of the active database LUNS. It is recommended that both active and target storage be sized to accommodate for the increased workload of the target database LUNS. This accounts for increased I/O on database LUNS in a fail over situation.

With the introduction of Exchange 2007 SP1 the target database I/O is much lower than that of the active node. This improvement means that target databases from multiple CCR clusters can be combined in the same aggregate. Prior to SP1 target storage for each target CCR node needed to be isolated in a separate aggregate due to the heavy I/O load of the target database.

#### **Separate Volumes and Aggregates**

When designing storage for continuous replication environments storage for active and target database and log LUNS should be separated, Separate aggregates should be created for both the active and target transaction log LUNS. Active and target database LUNS should also be split into separate aggregates.

#### **Storage Isolation**

In order to provide maximum continuity in the event of a failure, active and target storage should be located on separate storage controllers. Placing storage for active and target nodes on the same storage controller makes the storage controller a single point of failure. Clustering storage controller will minimize this risk, but best practice is to utilize completely separate storage.

## **9.3 BACKUP CONSIDERATIONS FOR CONTINUOUS REPLICATION**

When implementing CCR environments it is best practice to use SME 4.0 to backup Exchange Server databases from the target node. This will minimize increased CPU and storage load caused by the integrity check process.

For additional information on continuous replication refer to the [Microsoft Exchange 2007 Continuous Replication Best Practice Guide](#).

Best Practice	
<b>Size active and target storage equally</b>	<b>Both active and target storage should be designed equally</b>
<b>Separate Aggregates and Volumes</b>	<b>Transaction log LUNS should be placed in separate aggregates for active and target storage.</b>
<b>Storage Isolation</b>	<b>Active and target storage should be placed on separate storage controllers.</b>
<b>Backing up CCR deployments</b>	<b>Mailbox database backups should be done from the passive node.</b>

## 10 SERVICE PACK 1

Exchange Server 2007 SP 1 includes not only enhancements to the existing features but new features as well. New features that increase availability, decrease I/O and improve transport performance. This section outlines some of the new features as well as important enhancements affecting storage planning.

### 10.1 EXCHANGE SP1 TRANSPORT FEATURES

#### BACK PRESSURE

Back pressure is a system resource monitoring feature on transport servers that monitors important system resources. When default thresholds are exceeded, new connections are rejected until system resource utilization falls below those thresholds. Algorithm improvements were made in SP1, and the free disk space requirement was reduced from 4GB to 500MB, which is important for transport sizing. For more information, see Microsoft's document [Understanding Back Pressure](#).

#### PRIORITY QUEUING

The categorizer will now act on the message priority that users set, whereas in previous Exchange versions it did not. Messages that are flagged with high importance will be queued before those with normal importance, which will be queued before those with low importance. High-importance messages that exceed the `MaxHighPriorityMessageSize` threshold (250kB) will be set to normal importance. For more information, see Microsoft's document [Understanding Priority Queuing in Exchange 2007 SP1](#).

#### TRANSPORT DUMPSTER

The hub transport servers in an organization can send any mail lost during a continuous replication failure from the hub transport dumpster. In Exchange 2007 SP1, LCR-enabled storage groups join CCR-enabled storage groups in being able to utilize this dumpster. In most cases, the only data at risk is client activity that does not go through transport, such as changing a message property, tasks, or many calendaring tasks. The hub transport servers cannot be clustered, so in a site disaster, hub transport servers that are lost are not available to recover mail from, and that mail might be lost. For more information, see the Microsoft document [Working with the Queue Database on Transport Servers](#).

#### BACKUP AND RESTORE

Exchange does not provide a mechanism to back up a hub transport server. One strategy is to utilize VSS to back up the hub transport database and transaction logs. The mailbox server role must be installed on the hub transport server, and the default database must be placed on the same NetApp LUN as the hub transport database. Next, use SnapManager® 4.0 for Exchange to perform a VSS backup of the mailbox database. The VSS writer will quiesce the mailbox databases and will create a Snapshot™ copy of the entire LUN. Proper freezing and quiescing of the hub transport database cannot occur. The best solution is to stop the transport service and then create the backup.

## 10.2 MAILBOX SERVER ROLE

### ONLINE DEFRAGMENTATION

Online defragmentation was seldom monitored because it was difficult to do so prior to Exchange 2007 SP1. In SP1, detailed information is presented in the event viewer as a single event: 703. This event shows how long the defragmentation pass took and how often the database has been fully defragmented since the installation of SP1.

This is very important information that can be used to adjust the online maintenance window. Microsoft recommends that defragmentation be completed in 14 days, so if the defragmentation is completing much sooner, shortening the online maintenance window can reduce the amount of disk churn and reduce the Snapshot space consumption on NetApp storage. For more precision, two performance counters should be logged to measure the online defragmentation trend:

- MExchange Database == > Instances\Online Defrag Pages Freed/Sec
- MExchange Database == > Instances\Online Defrag Pages Read/sec.

Reduce the online maintenance window if the read-to-freed ratio is greater than 100 to 1 and increase the online maintenance window if the read-to-freed ratio is less than 50 to 1. For example, if the read counter is averaging 600, and the freed counter is averaging 5, the ratio would be 120 to 1, and the recommendation would be to reduce the online maintenance window. These counters can be enabled with the registry key "[Show Advanced Counters.](#)"

### ONLINE MAINTENANCE DATABASE SCANNING, AND PAGE ZEROING

Prior to Exchange 2007 SP1, the only way to ensure a live database was not corrupt, and to zero deleted pages in the database without taking the database offline, is to perform an online streaming backup. During the streaming online backup, each page in the database is read, and the checksum integrity is verified. Page zeroing only occurs during streaming online backup, as the setting is ignored for a VSS backup.

Exchange 2007 SP1 forces page zeroing to create log files, propagating the page zeroing to all continuous replication targets.

When enabled in the registry ([Online Maintenance Database Scanning in Exchange 2007 SP1](#)), database scanning and page zeroing will run during the online maintenance window at the server level. Online database scanning addresses the CCR scenario where the backup is performed on the passive, and the active database is never completely scanned. Online maintenance will split the maintenance window effectively in half, between online defragmentation and database scanning. Databases should have page zeroing enabled when created; otherwise, the first time it is run against the database, the disk I/O workload (sequential read) will severely impact performance. Throttling can be enabled to reduce the impact the first time page zeroing is run.

### 10.3 HIGH AVAILABILITY FEATURES

#### STANDBY CONTINUOUS REPLICATION

SCR is a new feature introduced in Exchange 2007 SP1, and an SCR target can be configured for any source mailbox server role except LCR. The server that is configured as an SCR target must have the mailbox server role installed and cannot have any LCR-enabled databases. The SCR server can be in a clustered configuration; however, no configured clustered mailbox servers can be on that cluster (only SCR targets). SCR storage groups (like LCR and CCR) cannot have more than one database.

All source mailbox server roles (except LCR) can have multiple SCR targets for each storage group, though the practical (and tested by Microsoft) limit is four SCR targets per source. An SCR server can be configured as the SCR target for storage groups from multiple (non-LCR) source servers.

Microsoft does not provide a backup mechanism for SCR. However, one solution to this is to stop the replication service and to create a NetApp VSS Snapshot copy of the volumes housing the Exchange data.

## 11 WINDOWS 2008

Windows 2008 includes a number of improvements which can directly benefit Exchange Server deployments. Some of the benefits to running Exchange on Windows 2008 are as follows;

**Support for multi-subnet failover clusters** – This improvement enhances support not only for CCR clusters in Exchange 2007 SP1 but also makes configuration of NetApp's MetroCluster solution less complex.

**Improved SMB protocol** – Windows 2008 includes SMB version 2 which provides performance improvements over previous versions. This improvement means improvement in log shipping performance.

**Virtualization** – Virtualization is another area of Window 2008 which been greatly improved. These improvements can add flexibility to Exchange Server deployments. One example of this is the case of disaster recovery.

In order to provide redundancy for a remote site, some organizations may require a Warm Site that contains a duplicate of the primary production Exchange 2007 infrastructure. The intent of this standby site is to provide as near to the same level of functionality as possible in the event of the loss of the primary site. However, keeping a duplicate infrastructure for standby purposes can be cost prohibitive. Hyper-V makes it possible to provide a virtual duplicate of the entire primary site.

## 12 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. Prior to making storage related changes, increasing the number of users or mailbox quotas, sizing and performance should be reviewed with a NetApp Exchange specialist.

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.

## 13 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. NetApp storage appliances and data management 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, NetApp offers the flexible storage and data management solutions to support Exchange Server 2007 enterprise messaging systems.

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.

As always, please consult with a local NetApp Exchange expert when planning and deploying Exchange environments onto NetApp storage. They are the experts who can quickly identify the needs and demands for any Exchange environment, and adjust the storage solution accordingly.

## 14 APPENDIX A – BEST PRACTICES SUMMARY

### SnapDrive

#### Best Practice

Use SnapDrive to create and manage NetApp 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 NetApp storage. SME will perform the data migration from local disks to NetApp 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, SATA disks can also meet the needs for Exchange 2007 in certain less I/O intensive deployments. To determine the I/O requirement for an Exchange environment, please consult with your local NetApp Exchange expert.

### RAID-DP

#### Best Practice

Always use RAID-DP when deploying Exchange server data onto a NetApp 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.

#### Best Practice

It is recommended to have at least 10% free space available in an aggregate hosting Exchange data. This allows for optimal performance of the storage system.

### 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. For more information about busy Snapshot copy issues, see [Chapter 12 of the \*SME Installation and Administrator Guide\*](#).

#### 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.

#### Best Practice

It is recommended to have at least 10% free space available in a volume hosting Exchange data. This allows for optimal performance of the storage system.

### 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.

#### Best Practice

Microsoft recommends approximately 20% free disk space for optimal Windows performance. This 20% free disk space can help prevent an Exchange outage by providing additional storage space in the event additional Exchange data is written to a LUN and avoiding an out of space condition for Exchange, thus taking the affected storage group offline. Please refer to the Microsoft article [Deploying a Large Exchange Server 2007 Organization](#).

### 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 NetApp Sizing Tool for Exchange should be used to size all Exchange server deployments utilizing NetApp 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	<p>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:</p> <ul style="list-style-type: none"> <li>• Run verification during off-peak hours</li> <li>• Use a verification server</li> <li>• Offload verification process onto separate storage</li> <li>• Do not delete transaction logs from non-verified backups</li> </ul> <p>For additional information on database verification options refer to the chapter on <a href="#">Selecting the Database Verification Server in the SME Installation and Administration Guide</a>.</p>

## 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

## Continuous Replication

Best Practice	
Size active and target storage equally	Both active and target storage should be designed equally
Separate Aggregates and Volumes	Transaction log LUNS should be placed in separate aggregates for active and target storage.
Storage Isolation	Active and target storage should be placed on separate storage controllers.
Backing up CCR deployments	Mailbox database backups should be done from the passive node.