



Technical Report

Storage Efficiency and Best Practices for Microsoft Exchange Server 2010

Brad Garvey, NetApp
April 2010 | TR-3824

TABLE OF CONTENTS

1	INTRODUCTION	3
1.1	SCOPE	3
1.2	INTENDED AUDIENCE	3
1.3	CAVEATS	3
2	BACKGROUND	4
2.1	WHAT IS STORAGE EFFICIENCY?	4
2.2	NETAPP TECHNOLOGY FOR STORAGE EFFICIENCY	4
3	ACHIEVING STORAGE EFFICIENCY FOR EXCHANGE SERVER 2010	5
4	DESIGNING STORAGE EFFICIENCY FOR EXCHANGE SERVER 2010	5
4.1	DEPLOYMENT AND PROVISIONING OF STORAGE: DISK DRIVES	7
4.2	DEPLOYMENT AND PROVISIONING OF STORAGE: RAID-DP	7
5	STORAGE LAYOUT PLANNING	8
5.1	AGGREGATE RECOMMENDATIONS	8
5.2	VOLUME PLANNING	8
5.3	LUN LAYOUT RECOMMENDATIONS	9
5.4	VOLUME SIZING	10
5.5	SNAPSHOT COPY SPACE MANAGEMENT	11
6	SIZING AND CAPACITY PLANNING	12
6.1	DATABASE CACHE RECOMMENDATIONS	12
6.2	CAPACITY PLANNING	14
7	PLANNING STORAGE CAPACITY	15
7.1	FACTORS AFFECTING DISK AND STORAGE CAPACITY	15
7.2	MAINTENANCE AND MOVING MAILBOXES	16
7.3	DATA PROTECTION	17
8	HIGH AVAILABILITY	17
8.1	DESIGN CONSIDERATIONS AND BEST PRACTICES	18
8.2	DEPLOYMENT SCENARIOS	19
9	VIRTUALIZATION	20
10	CONCLUSION	20
	APPENDIX: BEST PRACTICES	20

1 INTRODUCTION

This document provides best practices and insights related to the various design considerations that affect storage efficiency when implementing NetApp® storage solutions for the Microsoft® Exchange Server 2010 environment.

1.1 SCOPE

The scope of this guide is limited to technical discussions of design guidelines based on the principles, policies, and preferred standards recommended by NetApp for storage infrastructures that host Microsoft Exchange Server 2010. This document does not describe an end-to-end reference implementation.

1.2 INTENDED AUDIENCE

This guide offers guidance in planning and deploying Microsoft Exchange Server 2010 on NetApp storage. The best practices and recommendations presented in this guide enable a highly available, easy-to-manage Exchange environment that meets SLA requirements. Consult with a local NetApp Exchange expert for additional guidance when planning, designing, and deploying Exchange Server environments using NetApp storage solutions.

This guide covers:

- Best practices for deploying Exchange Server 2010 on NetApp storage
- Sizing considerations for Exchange Server 2010 deployments
- Storage layout considerations

This document highlights specific best practices in boxed sections. A complete list of storage efficiency solutions and best practices as they pertain to Exchange Server 2010 is also available in the appendix.

1.3 CAVEATS

The best practices for Exchange Server 2010 storage architecture presented in this guide focus exclusively on the latest NetApp storage operating system, Data ONTAP® 7.3.x.

This document does not include guidance for unified messaging in Exchange Server 2010.

2 BACKGROUND

2.1 WHAT IS STORAGE EFFICIENCY?

Most simply, NetApp storage efficiency is defined as increased storage use and decreased storage costs.

2.2 NETAPP TECHNOLOGY FOR STORAGE EFFICIENCY

Figure 1 illustrates the technologies that NetApp offers to implement storage efficiency and realize cost-saving benefits. This is accomplished by optimizing existing storage in the infrastructure and deferring or avoiding future storage needs. Savings increase as more and more of these technologies are used in conjunction with one another.

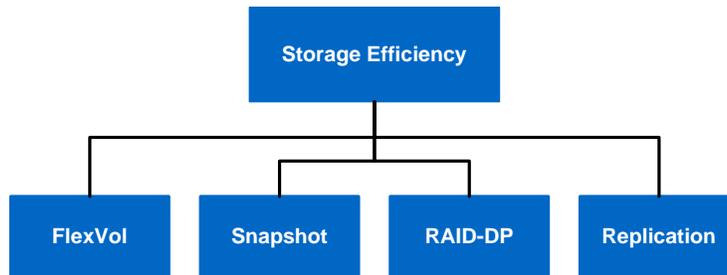


Figure 1) Storage efficiency.

RAID-DP

RAID-DP[®] is NetApp's implementation of double-parity RAID 6, which is an extension of NetApp's original Data ONTAP WAFL[®] RAID 4 design.

Unlike other RAID technologies, RAID-DP provides the ability to achieve a higher level of data protection without any performance effect while consuming a minimal amount of storage.

SATA

The performance acceleration provided by WAFL and the double-disk protection provided by RAID-DP make economical and large-capacity SATA drives practical for production application use. In addition, to negate the read latencies associated with large drives, SATA drives can be used with the NetApp Performance Acceleration Module (PAM) card, which significantly increases performance with large working set sizes.

SATA drives are more susceptible to unrecoverable BIT errors and operational failures. Without high-performance double-disk RAID protection such as RAID-DP, large SATA disk drives are more prone to failure, especially during their long RAID reconstruction time for a single drive failure. Table 1 compares RAID types and the probability of data loss in a five-year period.

Table 1) Data loss probability comparisons.

RAID Type	Probability of Data Loss in Five-Year Period	Relative Risk of Data Loss Compared to RAID-DP
RAID 10 (1 data disk)	0.33%	163
RAID 5 (7 data disks)	6.0%	3,955
RAID 6 (7 data disks)	0.002%	1.0
RAID-DP (7 data disks)	0.002%	1.0

For additional information on Exchange Server and RAID-DP, see <http://media.netapp.com/documents/tr-3574.pdf>.

THIN PROVISIONING, FLEXVOL

Thin provisioning is a function of NetApp FlexVol®, which allows storage to be provisioned just like traditional storage. However, it is not consumed until the data is written (just-in-time storage).

DEDUPLICATION

NetApp deduplication technology leverages NetApp WAFL block sharing to perform protocol-agnostic data-in-place deduplication as a property of the storage itself.

SNAPSHOT

NetApp Snapshot™ technology provides zero-cost, near-instantaneous backup, point-in-time copies of the volume or LUN by preserving Data ONTAP WAFL consistency points (CPs).

Creating Snapshot copies incurs minimal performance effect because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is at the rate of block-level changes, not 100% for each backup as it is with mirror copies. Using Snapshot can result in savings in storage cost for backup and restore purposes and opens up a number of efficient data management possibilities.

NETAPP STRATEGY FOR STORAGE EFFICIENCY

As seen in the previous section on technologies for storage efficiency, NetApp's strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core Data ONTAP operating system and the WAFL file system. Unlike its competitors' technologies, NetApp's technologies surrounding its FAS and V-Series product line have storage efficiency built into their core.

Customers who already have other vendors' storage systems and disk shelves can still leverage all the storage saving features that come with the NetApp FAS system simply by using the NetApp V-Series product line. This is again in alignment with NetApp's philosophy of storage efficiency (helping increase storage use and decrease storage cost) because customers can continue to use their existing third-party storage infrastructure and disk shelves, yet save more by leveraging NetApp's storage-efficient technologies.

3 ACHIEVING STORAGE EFFICIENCY FOR EXCHANGE SERVER 2010

In general, NetApp's storage efficiency strategy and corresponding technologies are relevant for any kind of enterprise application. As with any environment, however, proper application of the principles is key. The remaining sections of this document specifically deal with applying these storage efficiency technologies to a Microsoft Exchange Server 2010 environment as well as best practices for implementing Exchange 2010 on NetApp storage.

4 DESIGNING STORAGE EFFICIENCY FOR EXCHANGE SERVER 2010

Table 2 presents the principles for designing storage efficiency for Exchange Server 2010.

Table 2) Storage efficiency principles.

Principle 1	
STATEMENT	STANDARDIZE AROUND RAID-DP AS THE RAID TECHNOLOGY
Rationale	RAID-DP provides double disk failure protection in a RAID group with minimal storage needs.
Implications	RAID-DP provides a high level of protection with no performance effect and storage efficiency.
Principle 2	
STATEMENT	USE SATA DISKS WHEREVER APPROPRIATE

Rationale	SATA disks can be a good fit in environments with high capacity requirements. Higher read latency associated with SATA drives can be reduced by using the NetApp PAM card.
Implications	High-capacity, economical disks have a better mapping to Exchange environments with larger mailbox sizes and high-capacity demands.
Principle 3	
STATEMENT	THIN PROVISION THE STORAGE USING NETAPP FLEXVOL
Rationale	Thin provisioning a dedicated restore LUN can provide the ability to use the dedicated restore LUN when needed without permanently allocating volume and LUN space.
Implications	Storage silos can be removed.
Principle 4	
STATEMENT	WHEREVER APPROPRIATE, USE SERVER VIRTUALIZATION TECHNOLOGIES WITH NETAPP STORAGE TO GAIN ADDITIONAL SAVINGS
Rationale	NetApp's deduplication technology is extremely effective in virtual infrastructure and hence provides more storage savings.
Implications	NetApp's deduplication technology makes possible a more dynamic and storage-efficient infrastructure.
Principle 5	
STATEMENT	USING DEDUPLICATION IN EXCHANGE 2010 ENVIRONMENTS CAN REDUCE THE AMOUNT OF STORAGE USED
Rationale	NetApp provides an in-place deduplication strategy, applicable for both primary and secondary storage.
Implications	Depending on the message profile, using deduplication can save anywhere from 10% to 30% of storage.

4.1 DEPLOYMENT AND PROVISIONING OF STORAGE: DISK DRIVES

Enhancements to Exchange Server 2010 have dramatically decreased the amount of I/O as compared to previous versions of Exchange. With the reduction in I/O and the increase in mailbox sizes, some drive types might provide advantages in capacity and I/O.

FIBRE CHANNEL

Fibre Channel (FC) disk drives are proven, reliable storage devices with high read/write speeds that have the ability to handle high I/O loads. FC disks are ideally suited to handle the I/O requirements of Exchange Server deployments.

SAS

SAS disks are newer to the market than FC and SATA disks. SAS disk drives provide capacity as well as throughput speeds that make them an excellent choice for Exchange Server workloads.

SATA

The decreased I/O profile in Exchange Server 2010 makes SATA a viable solution in many Exchange environments. Although the SATA disks might be a good fit for performance and capacity, NetApp recommends that when using SATA disks for Exchange deployments, they be used in a DS4243 disk shelf. A PAM card should be used in solutions involving SATA drives in order to account for periods of increased I/O.

When choosing drive types for Exchange 2010, NetApp recommends consulting a local NetApp Exchange expert. The expert can provide deployment guidance to make sure that items such as I/O latencies and SLAs are in line with the requirements of Exchange Server 2010 and Exchange administration staff.

4.2 DEPLOYMENT AND PROVISIONING OF STORAGE: RAID-DP

RAID-DP is NetApp's implementation of double-parity RAID 6, an extension of NetApp's original Data ONTAP WAFL RAID 4 design.

Unlike other RAID and enhanced RAID technologies such as RAID 10, RAID 5, RAID 6, and RAID 50, what makes RAID-DP special is its ability to provide a higher level of data protection with minimal performance effect while consuming a minimal amount of storage.

5 STORAGE LAYOUT PLANNING

5.1 AGGREGATE RECOMMENDATIONS

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

Creating separate aggregates for Exchange database volumes and Exchange transaction log/SnapInfo volumes can meet the performance requirements of Exchange Server 2010 while providing the data availability required by most typical SLA agreements. 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.

In Exchange Server 2010 environments that use mailbox resiliency, Microsoft has relaxed the requirements for isolating database and log files to separate sets of disks. This means that database and log volumes can be placed in the same LUN. Table 3 details the benefits and trade-offs for two possible aggregate layouts.

Table 3) Aggregate layouts.

Number of Aggregates	Benefits	Trade-Offs	Recommended Configuration
One single aggregate for all Exchange data	Minimize the number of parity disks required	In the unlikely event of losing an aggregate, Exchange data is completely lost from the affected aggregate	Low I/O Exchange deployments Smaller storage controller deployments
Multiple aggregates for Exchange data (database and transaction logs on separate aggregates)	Dedicated aggregate handling the I/O for a database or for transaction logs If a single aggregate is lost (very unlikely), there is not a total loss of Exchange data	More parity disks required because there are more aggregates	I/O-intensive Exchange deployments

Best Practice

NetApp recommends having at least 10% free space available in an aggregate hosting Exchange data. This allows the storage system to perform optimally.

5.2 VOLUME PLANNING

Data ONTAP 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, and so on.
- All volumes can be managed independently while receiving the maximum I/O benefit of a much larger pool of disks.

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 helps determine the placement of volumes onto aggregates and the corresponding LUNs onto those volumes.

Best Practice

- NetApp recommends separating database and transaction logs from different servers into separate volumes to prevent a potential “busy” Snapshot copy problem. Because there are separate volumes for each server, there is no concern about Snapshot schedules overlapping 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 LUNs have a similar I/O profile, allowing them to share the same volume. For disaster recovery scenarios, having the entire log set for Exchange on the same volume helps achieve SLAs.
- NetApp recommends having at least 10% free space available in a volume hosting Exchange data.

5.3 LUN LAYOUT RECOMMENDATIONS

When considering Exchange 2010 LUN configuration, the number of LUNs you provision largely depends on the recovery point objectives (RPOs) and the recovery time objectives (RTOs).

ONE LUN PER DATABASE

- Single LUN per database architecture means that both the database and its corresponding log files are placed on the same LUN. In order to use this LUN configuration, mailbox databases must be part of a database availability group (DAG) that has two or more copies and does not use a hardware-based Volume Shadow Copy Service (VSS) solution.
- Deploying LUNs in this configuration simplifies storage administration because there are fewer LUNs to manage; however, it also limits the ability to perform hardware-based VSS backup and restore procedures.

TWO LUNS PER DATABASE

- In this solution, each mailbox database and transaction log set is placed on a separate LUN. This solution provides greater flexibility in terms of recovery options but increases the total number of LUNs required.
- In environments with high LUN counts, transaction logs for multiple mailbox databases can be placed on a single LUN. In this situation, each transaction log LUN should be placed on a separate volume to provide the best RPO/RTO times. When deploying LUNs in this fashion, NetApp recommends limiting the number of log streams per LUN to from five through 10.

Best Practice

- When creating LUNs, use volume mountpoints. This alleviates drive letter constraints when a large number of LUNs are required.
- NetApp recommends separating Exchange database and transaction log files onto separate LUNs and separate volumes whenever possible. This allows greater flexibility for backup and recovery procedures as well as data protection strategies.
- 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. It also avoids an out-of-space condition for Exchange, thus taking the affected storage group offline.

5.4 VOLUME SIZING

Volume sizing is broken into two different parts:

- Database volume sizing
- Transaction log volume sizing

NetApp recommends that when sizing Exchange volumes, the Microsoft sizing spreadsheet for Exchange first be used to determine the 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: the total size of all transaction log LUNs that will be stored in one transaction log volume
- Snapshot copy space: the space consumed by transaction logs generated during a 24-hour period

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

Transaction log volume size = Total transaction log LUN size + (Snapshot copy space * online backup retention duration)

Provided that an accurate change rate is known, the following formula can be used to calculate total volume size. To calculate the Exchange DB volume size, use the following variables:

- Database LUN size: the size of the LUN used to store the 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 of the database size
- Online backup retention duration: the number of days that backups are kept online; a day is measured by 24 hours
- Fault tolerance window: the number of days that backup failures can be tolerated before running out of Snapshot copy space; a typical value for this is two through four

Database volume size can be calculated using the following formula:

Database volume size = (Sum of the database LUN sizes that will share the database volume) + ((fault tolerance window + online backup retention duration] * database daily change rate)

DATABASE DAILY RATE OF CHANGE

The database rate of change can be estimated by the following formula, which takes into account the number of users, the number of messages sent and received per day, and the average message size.

A = average message size

R = number of messages received per user per day

S = number of messages sent per user per day

U = number of users

With tuning, the log change delta = (R+S) * A * U * 1.2.

The database change delta is log change delta * 1.4.

Example:

Average message size (A) = 100K

Number of messages received per user/day (R) = 80

Number of messages sent per user/day (S) = 20

Number of users (U) = 25K

Log Change Delta $(80+20) * 100K * 25K * 1.2 = 300GB/day$

Database Change Delta = Log change delta * 1.4 = 420GB

Total Change Delta = 720GB

5.5 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 allows sufficient space for proper Exchange Server database operation and online backup retention.

NetApp recommends the following storage system volume and Snapshot copy space management parameters:

- 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

Note: If `volume autogrow` is used, `autodelete` must also be enabled to allow adequate volume space.

Snapshot autodelete settings are:

- 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 Data ONTAP frees by deleting Snapshot copies with the autodelete function. By default, this value is set to 20%. This means that if autodelete is invoked by a low-space condition, autodelete deletes Snapshot copies until the volume has 20% free space. NetApp recommends that this value be adjusted to 2% greater than the autodelete threshold.

The autodelete threshold setting is based on the volume size shown in Table 4.

Table 4) Autodelete threshold settings.

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 making sure of adequate volume free space. The amount of volume free space and the daily change rate should be monitored on a regular basis because these might 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 Data ONTAP command line interface:

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

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

To size storage and use space reservation correctly, NetApp recommends that you work with your local NetApp professional. All factors are taken into consideration, and a successful fractional space reservation policy is deployed.

6 SIZING AND CAPACITY PLANNING

6.1 DATABASE CACHE RECOMMENDATIONS

In previous versions of Exchange Server, one of the key metrics needed for sizing storage is the amount of database I/O per second (IOPS) consumed by each user. The best predictors for Exchange Server 2010 mailbox IOPS are:

- Amount of database cache per mailbox
- Number of messages each user sends and receives per day

These estimates are only valid for database cache sizes between 3MB and 30MB per mailbox. The average message size used for the estimates is 75KB, but message size is not a primary factor for IOPS. Note that other client types and usage scenarios might yield different results.

Table 5 provides estimated values for IOPS per mailbox based on usage profile and database cache, which can be used to predict baseline Exchange 2010 mailbox I/O requirements.

Table 5) I/O profiles.

Messages Sent/Received per Mailbox per Day (~75KB Average Message Size)	Database Cache per Mailbox (MB)	Single Database Copy (Standalone): Estimated IOPS per Mailbox	Multiple Database Copies (Mailbox Resiliency): Estimated IOPS per Mailbox
50	3	.060	.050
100	6	.120	.100
150	9	.180	.150
200	12	.240	.200
250	15	.300	.250
300	18	.360	.300
350	21	.420	.350
400	24	.480	.400
450	27	.540	.450
500	30	.600	.500

After the database cache size requirements have been determined, the next step is to determine the minimum memory requirements per server. This verifies that the database cache size requirements can be met. The database cache size must be factored into the sizing process to make sure that the amount of physical memory per server is adequate to meet the needs of the mailbox count with a given user profile. Table 6 lists the default mailbox database cache sizes for both single-role mailbox servers as well as multirole servers.

Table 6) Default mailbox database cache sizes.

Server Physical Memory (RAM)	DB Cache Size: Mailbox Role Only	DB Cache Size: Multirole (for Example, Mailbox + Hub)
2GB	512MB	Not supported
4GB	1GB	Not supported
8GB	3.6GB	2GB
16GB	10.4GB	8GB
24GB	17.6GB	14GB
32GB	24.4GB	20GB
48GB	39.2GB	32GB
64GB	53.6GB	44GB
96GB	82.4GB	68GB
128GB	111.2GB	92GB

DETERMINE SERVER MEMORY REQUIREMENTS

Perform the following steps to determine server memory requirements.

1. Determine the amount of required database cache by multiplying the mailbox count times the memory requirements based on the user.
 Example: 2,500 150-messages/day mailboxes require 22.5GB of database cache (2,500 * 9MB = 22.5GB)
2. Determine the amount of required physical memory by determining which server configuration provides 22.5GB of database cache.
 Example: A single-role mailbox server with 32GB of physical RAM provides 24.4GB of database cache; therefore, 32GB of physical ram is the ideal memory configuration based on this mailbox count/user profile.
3. Determine the number of databases required for the design so that the memory requirement per database is met. The same simple process can be used to size multirole server configurations as well. The memory requirements of additional applications/workloads should then be added to the physical RAM resources required by Exchange Server 2010.

LOG CAPACITY CONSIDERATIONS

Exchange 2010 generates approximately 10% fewer transaction logs per day than Exchange 2007. Use Table 7 to estimate the number of transaction logs. It is important to note that the numbers in the following table represent the number of logs generated per mailbox per day, and that the numbers are based on an average message size of 75KB.

Table 7) Transaction logs generated.

Send/Receive Profile	Transaction Logs Generated (per Mailbox/Day)
10 sent/40 received	10
20 sent/80 received	20
30 sent/120 received	30
40 sent/160 received	40
50 sent/200 received	50
60 sent/240 received	60
70 sent/280 received	70
80 sent/320 received	80
90 sent/360 received	90
100 sent/400 received	100

For additional information on database cache and I/O recommendations, refer to <http://technet.microsoft.com/en-us/library/ee832793.aspx>.

6.2 CAPACITY PLANNING

A properly sized Exchange environment meets both Microsoft requirements for Exchange storage and customer requirements indicated in their SLAs. To get a properly sized environment, tools convert information about the customer environment into a physical storage recommendation.

You should use the following primary tools when planning an Exchange environment for a customer:

- Microsoft storage calculator
- 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 for planning an Exchange environment and provides a framework for storage group layout and LUN requirements. It is important to realize that the Microsoft storage calculator does not make recommendations on storage design (RAID parity, number of disks, and so on) because the storage design largely depends on the type of storage array being used. When sizing Exchange Server deployments using NetApp storage, it is important to use the NetApp Exchange Sizing Tool.

Best Practice

Use the NetApp Sizing Tool for Exchange to size all Exchange Server deployments using NetApp storage.

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

7 PLANNING STORAGE CAPACITY

7.1 FACTORS AFFECTING DISK AND STORAGE CAPACITY

Many factors are required to correctly size storage for Microsoft Exchange Server 2010. Following are three important factors directly affecting disk and storage capacity:

- Mailbox size
- Database dumpster
- Exchange search

MAILBOX SIZE/QUOTA

An important factor (if not the most important factor) is the mailbox size or quota. Understanding the mailbox size/quota is the starting point for calculating the amount of data storage required. It also indicates 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. Doing so will prevent future storage problems.

One of the new features of Exchange 2010 is the archive mailbox. If archiving is enabled, an additional archive mailbox for the user is created within the same mailbox database. Users can manually move items from their active mailbox to their archive mailbox. They also have the ability to import PSTs into the archive mailbox. Retention policies control when mail is automatically moved from the active mailbox to the archive mailbox. Quotas can be set on archive mailboxes just as they are set on active mailboxes.

DATABASE DUMPSTER

In Exchange 2007, the database dumpster was essentially a view. Items deleted from the folder remained in that folder and were flagged within the database that hides them from normal views.

The database dumpster has undergone architectural changes in Exchange 2010. The new version of database dumpster, termed Dumpster 2.0, is implemented as a folder called Recoverable Items and is located within the non-IPM subtree of the user's mailbox.

The folder has three subfolders:

- Deletions
- Versions
- Purges

The Deletions folder replaces the view that is displayed when a user accessed the Recover Deleted Items tool. When a user soft deletes or performs an Outlook hard delete against an item, the item is moved to the Recoverable Items\Deletions folder. When the user accesses Outlook/OWA Recover Deleted Items, the RPC Client Access service translates the request and returns the Recoverable Items\Deletions folder view.

Single Item Recovery prevents purging of data and provides versioning capability (the ability to retain the unaltered form of the item). By default, this data is retained until the age of the deleted item has exceeded the deleted item retention window. In addition, Exchange 2010 enables long-term preservation of data for litigation hold scenarios by preventing the purging of data all together.

Table 8 summarizes the behavior in Exchange 2010.

Table 8) Single item recovery.

Feature State	Soft-Deleted Items Kept in Dumpster	Modified (Versions) and Store Hard-Deleted Items Kept in Dumpster	User Can Purge Items from Dumpster	MRM Automatically Purges Items from Dumpster
Single item recovery disabled	Yes	No	Yes	Yes, 14 days by default and 120 days for calendar items
Single item recovery enabled	Yes	Yes	No	Yes, 14 days by default and 120 days for calendar items
Litigation hold enabled	Yes	Yes	No	No

It is important to note that when litigation hold is enabled, dumpster quotas are not enforced, preventing that data from being purged or preventing it from being stored in the dumpster. This means that additional capacity might need to be allocated when enabling litigation hold.

EXCHANGE SEARCH

Exchange Server 2010 creates an index of an Exchange database that is approximately 10% of the total database size. This index is placed on the same LUN as the indexed database. Sizing both volume and the hosting LUN to account for the index is very important. Administrators may turn off content indexing for a given database, but it is enabled by default in Exchange 2010.

7.2 MAINTENANCE AND MOVING MAILBOXES

MAINTENANCE

Online maintenance has changed drastically in Exchange 2010. By default, online maintenance is performed at runtime and is set to run in the background.

Maintenance changes in Exchange 2010 include:

- Cleanup is performed at run time (when hard delete occurs).
- Online maintenance happens during store dumpster cleanup; page zero is the default.
- Online defragmentation occurs and space is reclaimed at run time and is autothrottled.
- Database checksum by default is set to Always On, resulting in a more sequential I/O pattern.

Best Practice

Enable online maintenance to run 24x7.

MOVING MAILBOXES

Moving user mailboxes from previous versions of Exchange to Exchange 2010 generates 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 helps prevent the transaction log LUN from running out of space and causing downtime. Understanding the migration process for an Exchange Server helps determine how much additional storage is 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.

7.3 DATA PROTECTION

NetApp SnapManager[®] for Microsoft Exchange (SME) supports Microsoft Exchange Server 2007 and 2010.

SME is tightly integrated with Microsoft Exchange, which allows for consistent online backups of Exchange environments while leveraging NetApp Snapshot copy technology. SME is a VSS (Snapshot copy) requestor, which means that it uses the VSS subsystem supported by Microsoft to initiate backups. SME provides a complementary feature set for new Microsoft Exchange Server 2010 high-availability features. SME works with a DAG, providing the ability to back up and restore data from either the active database copy or one or more passive database copies.

Best Practice

Use SnapManager for Exchange when deploying Exchange Server 2010 on NetApp storage. SME performs the data migration from local disks to NetApp LUNs. It also manages that data, handling all backup, restore, and verification tasks.

For more information on SnapManager for Microsoft Exchange, visit the [SnapManager for Exchange](#) page on the NetApp Web site. Also, further best practices for SME can be found in the *SnapManager for Exchange Best Practice Guide*.

8 HIGH AVAILABILITY

Microsoft made significant changes to the high-availability option available in Exchange 2010. Some of the options and features in previous versions were removed. LCR, CCR, SCR, and SCC as known in Exchange 2007 are no longer available.

The following components were removed from Exchange 2010 in terms of clustering:

- There are no more EVS/CMS.
- The database is no longer associated with a server, but is an organization-level resource.
- There is no longer a requirement to choose cluster or noncluster at installation; an Exchange 2010 Server can move in and out of a database availability group (DAG) as needed.
- There is no longer the limitation of only hosting the mailbox role on a clustered Exchange Server.
- Failover and failback functions are no longer handled as part of Windows Cluster Services.

To replace the server and data resiliency options available in earlier versions of Exchange, Microsoft implemented the DAG feature. The DAG uses the same log shipping concept used in CCR. A DAG consists of two or more mailbox servers; however, there can be no more than 16 servers in the DAG. Each mailbox server can hold one or more active or passive copies of the database. Each database maintains separate status, so one server can host copies of multiple databases and have just some of those copies active at one time.

The DAG uses a new component in Exchange called Active Manager. The Active Manager facilitates failover and failback. In the event of a failure, Exchange 2010 promotes one of the copies of the database to active status, and the mailbox role then takes up the task of serving up the mailboxes on that database.

The following are some of the features of the DAG:

- Passive copy backup (for example, CCR) versus the HA+Hold policy (no backup)
- Storage failure detection and automatic failover
- No subnet or DNS requirements
- Fast failovers (<30s)

Figure 2 shows an example of a DAG.

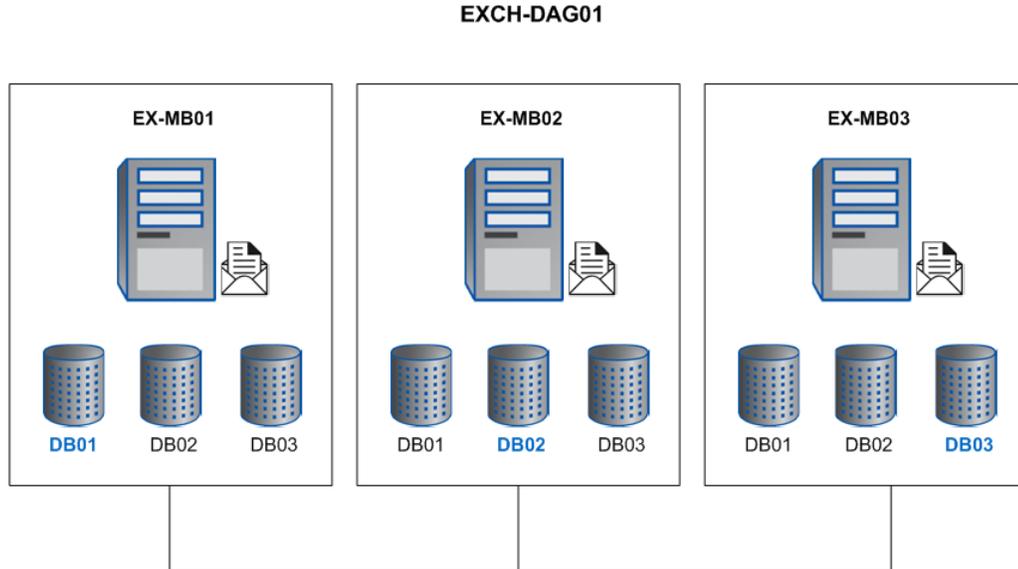


Figure 2) Database availability group.

In Figure 2, there are three servers and three copies of each database, one on each server. The active database copy is the one with the blue font. The passive copies of each database are updated by way of log shipping.

8.1 DESIGN CONSIDERATIONS AND BEST PRACTICES

The Exchange 2010 high-availability models address storage-related outages, plus server- and application-related outages, by adding additional copies of the mailbox databases. NetApp recommends the following practices when designing highly available solutions for Exchange 2010.

STORAGE-RELATED FAILURES

To limit exposure to storage-related failure, especially with JBOD deployments, Microsoft recommends a minimum of three copies of each mailbox database. This provides the ability to sustain a double-disk failure in a RAID-less JBOD solution without RAID. When deploying on NetApp storage, RAID-DP can provide fault tolerance for double-disk failure with minimal performance effect.

POINT-IN-TIME RESTORES

In a DAG configuration, each copy of the database is the current up-to-the-minute copy. This makes restoring to points of time in the past more difficult. To minimize this difficulty, an additional copy, called a LAG copy, can be added. With Exchange 2010 DAG, you can specify a truncation lag time of up to 14 days. This allows point-in-time restores of up to two weeks. For point-in-time restore functionality, SnapManager for Exchange can be used to provide space-efficient Snapshot copies, providing the flexibility to quickly and easily restore to a point in time without incorporating a LAG copy into the design.

8.2 DEPLOYMENT SCENARIOS

SINGLE-SITE SCENARIO

Deploying a two-node DAG with a minimum of two copies of each mailbox database in a single site is best suited for companies that want to achieve server and application level redundancy. In this situation, deploying a two-node DAG using RAID-DP provides not only server and application level redundancy but double disk failure as well. Adding SnapManager for Exchange in a single-site scenario gives the ability to do point-in-time restores without the added capacity requirements and complexity of a LAG copy.

MULTISITE SCENARIO

Extending a DAG across multiple data centers provides high availability of servers and storage components and also adds site resiliency. When planning a multisite scenario, NetApp recommends at least three mailbox servers as well as three copies of each mailbox database, two in the primary site and one in the secondary site. Adding at least two copies in both primary and secondary sites provides site resiliency and also provides high availability in each site.

For additional information on DAG layout planning, see <http://technet.microsoft.com/en-us/library/dd979781.aspx>.

When designing the storage layout for a DAG scenario, use the following design considerations and best practices.

Deployment

Best Practice

- In a single-site scenario, deploy a minimum of a two-node DAG with at least two copies of each mailbox database.
- In a multisite scenario, deploy at least three mailbox servers as well as three copies of each mailbox database, two in the primary site and one in the secondary site. Adding at least two copies in both primary and secondary sites provides site resiliency and also provides high availability in each site.

Storage Design

Best Practice

- Design identical storage for active and passive copies of the mailboxes regarding capacity and performance.
- Provision the active and passive LUNs identically regarding capacity and performance.

Volume Separation

Best Practice

- Place active and passive copies of the database in separate volumes.
- Place active and passive copies of the mailbox database in separate volumes.

Backup

Best Practice

Perform VSS backups on one of the passive nodes.

9 VIRTUALIZATION

Virtualizing Exchange environments deliver significant benefits, including:

- Reduced server hardware costs
- Power and space savings
- Improved server use
- Rapid server provisioning

When virtualizing Exchange 2010 roles, NetApp recommends that the roles be separated onto different servers. In the event of a host server failure, this prevents failure of any particular role. For example, deploying one CAS, one HUB, and two mailbox servers per host server provides a good mix in terms of disbursement of roles.

For additional information and recommendations for virtualizing Exchange 2010 services, refer to <http://technet.microsoft.com/en-us/library/aa996719.aspx>.

10 CONCLUSION

Microsoft Exchange Server 2010 is not a one-size-fits-all application. Multiple configuration options are available to suit most of the needs of 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 2010 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 provides a highly available, easy-to-manage Exchange environment that meets SLAs.

Consult with a local NetApp Exchange expert when planning and deploying Exchange environments onto NetApp storage. NetApp Exchange experts can quickly identify the needs and demands of any Exchange environment and adjust the storage solution accordingly.

APPENDIX: BEST PRACTICES

The following list contains the storage efficiency solutions and best practices as they pertain to Exchange Server 2010:

- **Storage design and layout**
 - NetApp recommends having at least 10% free space available in an aggregate hosting Exchange data. This allows the storage system to perform optimally.
 - NetApp recommends separating database and transaction logs from different servers into separate volumes 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 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 LUNs have a similar I/O profile, allowing them to share the same volume. For disaster recovery scenarios, having the entire log set for Exchange on the same volume helps achieve SLAs.
 - NetApp recommends having at least 10% free space available in a volume hosting Exchange data.
 - When creating LUNs, use volume mountpoints. This alleviates drive letter constraints when a large number of LUNs are required.

- NetApp recommends separating Exchange database and transaction log files onto separate LUNs and separate volumes whenever possible. This allows greater flexibility for backup and recovery procedures as well as data protection strategies.
- 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. It also avoids an out of space condition for Exchange, thus taking the affected storage group offline.
- **Sizing and capacity planning**
Use the NetApp Sizing Tool for Exchange to size all Exchange Server deployments using NetApp storage.
- **Database maintenance**
Enable online maintenance to run 24x7.
- **Data protection**
Use SnapManager for Exchange when deploying Exchange Server 2010 on NetApp storage. SME performs the data migration from local disks to NetApp LUNs. It also manages that data, handling all backup, restore, and verification tasks.
- **High availability (deployment)**
 - In a single-site scenario, deploy a minimum of a two-node DAG with at least two copies of each mailbox database.
 - In a multisite scenario, deploy at least three mailbox servers as well as three copies of each mailbox database, two in the primary site and one in the secondary site. Adding at least two copies in both primary and secondary sites provides site resiliency and also provides high availability in each site.
- **High availability (storage design)**
 - Design identical storage for active and passive copies of the mailboxes regarding capacity and performance.
 - Provision the active and passive LUNs identically regarding capacity and performance.
- **High availability (volume separation)**
 - Place active and passive copies of the database in separate volumes.
 - Place active and passive copies of the mailbox database in separate volumes.
- **High availability (backup)**
Perform VSS backups on one of the passive nodes.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.