Technical Report

Microsoft Exchange Server 2010 and SnapManager for Exchange Best Practices Guide

Niyaz Mohamed & Robert Quimbey

March 2013 | TR-4033

TABLE OF CONTENTS

LIST OF FIGURES

1 Introduction

1.1 Purpose and Scope

This guide provides® NetApp best practices for deploying Microsoft® Exchange Server 2010 and using SnapManager® for Microsoft Exchange (SME).

1.2 Intended Audience

This paper is a best practice guide for experienced Microsoft Exchange administrators who have read the following documents:

- “SnapManager for Exchange Installation and Administration Guide”
- “SnapDrive for Windows Installation and Administration Guide”
- “Data ONTAP System Administrators Guide”

Readers of this best practice guide should have a solid understanding of the Exchange storage architecture and Exchange administration as well as Exchange backup and restore concepts. The recommendations in this document are best practices to assist with the design, implementation, and configuration of SnapManager for Exchange in Windows Server® 2008 and Windows Server 2008 R2 environments with Microsoft Exchange Server 2010.

2 Exchange Server 2010 Architecture

Exchange Server 2010 includes the following server roles:

- **Client access servers.** Support traditional components such as Post Office Protocol 3 (POP3) and Internet Message Access Protocol 4 (IMAP4), Exchange ActiveSync, Microsoft Outlook Web app, Outlook Anywhere, and several new features, including the RPC client access service and the Exchange control panel.
- **Edge transport servers.** Handle message traffic to and from the Internet and run spam filters.
- **Hub transport servers.** Perform internal message transfer, distribution list expansions, and message conversions between Internet mail and Exchange Server message formats.
- **Mailbox servers.** Maintain mailbox store databases, provide client access servers with access to the data, and support access to public folders for Outlook clients.
- **Unified messaging servers.** Integrate voice and fax with e-mail messaging and run Outlook voice access.

2.1 Database Availability Groups

Exchange Server 2007 introduced a built-in log shipping feature called continuous replication. Continuous replication, which was available in three forms—LCR, CCR, and SCR—significantly reduced the cost of deploying a highly available Exchange infrastructure and provided a much improved deployment and management experience over previous versions of Exchange.

Although the introduction of continuous replication in Exchange Server 2007 did provide high availability, it was still a challenge due to the integration between Exchange and Windows® failover clustering.

Exchange Server 2010 combines on-site data replication (CCR) and off-site data replication (SCR) into a single framework called a database availability group (DAG). A database availability group is a cluster of up to 16 nodes that provide automatic database-level failover.

DAGs use continuous replication and a subset of Windows failover clustering to provide continuous mailbox availability.

2.2 Personal Archive Mailbox

A personal archive is an additional mailbox associated with a user's primary mailbox. This new mailbox is known as an archive mailbox and is provisioned automatically for the user when the administrator enables the personal archive feature.

After the archive mailbox has been associated with the user account, mail can be moved by the user into the personal archive by dragging and dropping PST files or automatically through retention policies. Exchange Server 2010 SP1 allows the archive mailbox to be placed in a different database than the primary mailbox.

3 Exchange Server 2010 Planning Considerations

3.1 System Requirements

In this section we discuss the system requirements for Exchange Server 2010 on NetApp storage systems.

- Windows Server 2008 and Windows Server 2008 R2 64-bit edition
- Minimum and maximum page file size set to physical RAM plus 10MB
- Memory requirements vary depending on Exchange features that are installed; for detailed information about memory requirements for Exchange 2010, see "Understanding Memory Configurations and Exchange Performance"
- Disk space depends upon the requirements; at least 1.2GB of free space is required on the drive on which you will install Exchange Server
- Disk partitions formatted as NTFS file systems

For more detailed requirements, visit the Microsoft TechNet article Exchange 2010 System Requirements.

4 NetApp Storage Efficiencies

NetApp offers the following technologies that increase storage efficiency.

4.1 RAID-DP

RAID-DP® technology prevents data loss when up to two drives fail per RAID group.

RAID-DP is integrated with the WAFL® file system so that the dedicated parity drives don't become a performance bottleneck. RAID-DP makes SATA disks an option for your enterprise storage. Exchange administrators can use less-expensive SATA without worrying about data loss and also lower their storage acquisition costs.

Note: SyncMirror® can be used along with RAID-DP to provide a second layer of mirrored protection for a more robust disk protection strategy.

4.2 Snapshot

NetApp Snapshot™ technology provides low-cost, fast-backup, point-in-time copies of the file system (volume) or LUN by preserving Data ONTAP® architecture WAFL consistency points.

There is no performance penalty for creating Snapshot copies, because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not 100% for each backup as with mirror copies. It can result in savings in storage costs for backup and restore purposes and opens up a number of efficient data management possibilities.

Refer to the Backup and Recovery section for more information on how to leverage NetApp Snapshot technology for data protection requirements for Microsoft Exchange Server 2010 environments.

4.3 Thin Provisioning

Thin provisioning, in a shared storage environment, is a method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data versus the traditional method of allocating all of the blocks up front. This methodology eliminates almost all white space, which helps avoid poor utilization rates. Flexible volumes (FlexVol® volumes) are the enabling technology behind NetApp thin provisioning, which can be thought of as the virtualization layer of Data ONTAP. When a LUN is created, it does not dedicate specific blocks out of the NetApp volume for the LUN or for Snapshot copies of the LUN. Instead, it allocates the blocks from the NetApp aggregate when the data is actually written. This allows the administrator to provision more storage space, as seen from the connected servers, than is actually physically present in the storage system.

When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any.

Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.

If the time required to procure new storage is very long, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

The potential risk when configuring the Exchange environment for thin provisioning is a LUN going offline when there is not enough space to write further data.

4.4 Space Guarantee

The space guarantee is the enabler of thin provisioning. Space guarantees can be set at the volume or the LUN level, depending on the space guarantee requirements of the application. Typically, if the space guarantee at the volume level is set to “volume,” the amount of space required by the flexible volume or FlexVol volume is always available from its aggregate. This is the default setting for FlexVol volumes. When the space guarantee is set to “volume,” the space is reserved from the aggregate’s available space at volume creation time.

When the space guarantee is set to “none,” the volume reserves no space from the aggregate during volume creation. Space is first taken from the aggregate when data is actually written to the volume. Write operations to space-reserved LUNs in a volume with ‘guarantee=none’ will fail if the containing aggregate does not have enough available space.

LUN reservation enables the LUN to have space in the volume, but ‘guarantee=none’ does not enable the volume to have space in the aggregate. When the space guarantee for the volume is set to “File,” the aggregate enables space to be available to completely rewrite LUNs that have space reservation enabled.

4.5 Space Reclamation

Space reclamation must be initiated from time to time to recover the unused space in a LUN. Storage space can be reclaimed at the storage level using the SnapDrive -> Start Space Reclaimer option.

4.6 Fractional Reserve

Fractional reserve is a volume option that determines how much space Data ONTAP will reserve for Snapshot overwrite data for LUNs to be used after all other space in the volume is used. The default value for fractional_reserve is 100%. However, using the autodelete functionality, the fractional reserve can be set to “0”; through the command line interface (CLI), it can be set to anything from 0% to 100%.

4.7 Autodelete and Autosize

The autosize volume setting (available in Data ONTAP 7.1 and later) defines whether a volume should automatically grow to avoid filling up to capacity. It is possible to define how quickly the volume should grow with the “-i” option. The default growth increment is 5% of the volume size at creation. It is also possible to define how large the volume is allowed to grow with the “-m” option. If volume autosize is enabled, the default maximum size to which to grow is 120% of the original volume size.

Example:

```
vol autosize vol0 -m 1500g -i 1g on

vol status –v vol0

        Volume autosize settings:

                                state=on

                                maximum-size=1500GB

                                increment-size=1GB
```

Best Practices

NetApp recommends planning for additional buffer space when using thin provisioning for Microsoft Exchange Server 2010 environments.

NetApp recommends prioritizing autosize over autodelete because deletions occur at the Data ONTAP level, and it is possible to have a backup set of a transaction log and database where one of the Snapshot copies has been automatically deleted or orphaned.

When autodelete is used, NetApp recommends that the Snapshot autodelete functionality in SnapManager for Exchange be used, and not the Data ONTAP autodelete feature. If not, SnapManager will not continue deleting Snapshot copies as per the retention. Snapshot copies should only be deleted using SnapManager, either with the retention or with the delete backup wizard. Snap autodelete needs to be off on ALL volumes managed by a SnapManager product.

In such scenarios, the recommendation is to use autosize; however, it might fail due to space constraints in the aggregate and must be properly monitored using Operations Manager. For volume autosize to work, it is mandatory that the containing aggregate has enough space (at least 1.2 times the volume size).

Both autodelete and autosize work at the volume level and not on individual LUNs. This means that LUNs will not automatically grow and must be handled separately with different commands. NetApp SnapDrive® for Windows (SDW) can be used to make more space available for the LUN.

The autodelete volume setting (available in Data ONTAP 7.1 and later) allows Data ONTAP to delete Snapshot copies if a threshold is met. This threshold is called a “trigger” and can be set so that Snapshot copies will be automatically deleted when one of the following conditions is met.

- **Volume.** The volume is near full. This is reported in the first line reported for each volume in the `df` command. It should be noted that the volume can be full even though there might still be space in the snap_reserve areas.
- **Snap_reserve.** The snap reserve space is near full.
- **Space_reserve.** The “overwrite reserved” space is full. This is the space determined by the LUNs with space reservations enabled and the fractional_reserve option. The reserve space will never be filled until both the volume and the snap_reserve areas are full.

Note: The `df` command is available when accessing NetApp storage using the CLI.

```
6240b> df

File system          Kbytes       used        avail        capacity   Mounted on

/vol/vol0/           1407415772   12094808    1395320964   1%         /vol/vol0/

/vol/vol0/.snapshot  74074512     455588      73618924     1%         /vol/vol0/.snapshot
```

Best Practice

NetApp recommends using autogrow instead of autodelete. When using autodelete, set the autodelete trigger to volume.

The order in which Snapshot copies are deleted is determined by the following three options.

- **Delete_order.** This option determines whether the oldest or newest Snapshot copies should be deleted first.
- **Defer_deleted.** This option allows the user to define a group of Snapshot copies that should first be deleted when no other Snapshot copies are available. It is possible to defer the deletion of user-created Snapshot copies, scheduled Snapshot copies, or Snapshot copies beginning with a configurable prefix.
- **Commitment.** This option determines how Snapshot copies used for SnapMirror® and dump operations should be handled. If set to “try,” it will only delete these Snapshot copies if they are not locked. If set to “disrupt,” these Snapshot copies will be deleted even if they are locked.

Best Practice

When using SnapMirror products or SnapVault® hardware for replicating Microsoft Exchange Server 2010 databases, NetApp recommends not using the “disrupt**”** option for commitment. This is because SnapMirror baseline Snapshot copies can be destroyed by autodelete even though they will always be the last Snapshot copies deleted. In many configurations, deleting the last SnapMirror Snapshot copy is not desired because a new full baseline copy is required to resume mirroring operations. If, for example, the source and destination are at different sites, recreating this baseline can be a time-consuming and costly process.

4.8 Best Practice Configurations When Using Thin Provisioning for Microsoft Exchange Server 2010 Environments

There are many ways to configure the NetApp storage appliance for LUN thin provisioning; each has advantages and disadvantages. It should be noted that it is possible to have thinly provisioned volumes and non–thinly provisioned volumes on the same storage system or even the same aggregate. The following are considered to be best practice configurations when using thin provisioning for Microsoft Exchange Server 2010.

Option 1: Volume Guarantee Set to ‘None’

Volume guarantee	= none
LUN reservation	= enabled
fractional_reserve	= 0%
snap_reserve	= 0%
autodelete	= volume / oldest_first
autosize	= off
try_first	= snap_delete

This configuration has the advantage of the free space in the aggregate being used as a shared pool of free space. The disadvantages of this configuration are the high level of dependency between volumes and that the level of thin provisioning cannot easily be tuned on an individual volume basis. When using this configuration, the total size of the volumes is greater than the actual storage available in the host aggregate. With this configuration storage administrators can generally size the volume so that they only need to manage and monitor the used space in the aggregate. This option does not affect the space for hosting the live data, but rather allows the backup space to dynamically change.

Option 2: Using Autogrow/Autodelete

Volume guarantee	= volume
LUN reservation	= disabled
fractional_reserve	= 0%
snap_reserve	= 0%
autodelete	= volume / oldest_first
autosize	= on
try_first	= autogrow

This configuration allows the administrator to finely tune the level of thin provisioning for Microsoft Exchange Server 2010 environments. With this configuration the volume size defines or guarantees an amount of space that is only available to LUNs within that volume. The aggregate provides a shared storage pool of available space for all the volumes contained within it. If the LUNs or Snapshot copies require more space than available in the volume, the volumes will automatically grow, taking more space from the containing aggregate. Additionally, the advantage of having the LUN space reservation disabled in that case is that Snapshot copies can then use the space that is not needed by the LUNs. The LUNs themselves are also not in danger of running out of space because the autodelete feature will remove the Snapshot copies consuming space. This is an ideal setting for many migrations in which Snapshot copy space will be high during the initial mailbox moves, but will taper off in the months and years to come when more space is required within the database to store mail.

Note: Snapshot copies used to create FlexClone® volumes will not be deleted by the autodelete option.

Best Practice
Using autogrow is the most common deployment configuration.

4.9 Monitoring

When using NetApp efficiency features, the volumes should be appropriately sized so that autosize and/or autodelete policies are not triggered unless there is an abnormal rate of change or a problem with Snapshot copy retention. NetApp OnCommand® management software that includes Operations Manager is the recommended tool to monitor Exchange volumes for these events and to send notifications to the storage administration team to follow up further with the Exchange administration team. SNMP can also be used to monitor these events. See the appendix for SNMP examples.

After a notification for a volume autogrow or Snapshot autodelete event has been received by the storage administration team, the recommended action is for the storage administration team to examine the affected storage controllers and then follow up with the Exchange administration team for further administrative actions.

A typical cause of volume autosize events is that the rate of change greatly surpassed the rate of change assumption used in sizing the volume. Adding additional Exchange mailboxes beyond the original database design parameters or e-mail storms can cause increased data change rates. Another cause for volume autosize events is that older Snapshot copies created by SnapManager for Exchange are not being deleted. As Snapshot copies age, they can grow in size and consume more capacity than originally allocated in the volume. A typical cause of SnapManager for Exchange not deleting backups is that SnapManager for Exchange backups are failing. By default, SnapManager for Exchange does not delete Snapshot copies of older SnapManager for Exchange backup sets if the backup fails. Another cause for SnapManager for Exchange not deleting backups is that the SnapManager for Exchange backup retention policies are not being enforced correctly because Snapshot copies were manually removed outside of SnapManager for Exchange on the controller itself.

Monitoring the health of SnapManager for Exchange can be done by monitoring for SnapManager for Exchange event IDs and the enhanced enterprise monitoring functionality in SME 6.0.2R1. To monitor the health of SnapManager for Exchange retention external to SnapManager for Exchange, the number of Snapshot copies and SnapInfo directories should be calculated for a specific Exchange Server. For example, if the SnapManager for Exchange retention policy for a particular server is 10 backups online (-RetainBackups 10 parameter in the `new-backup` command), there will be 10 SnapManager for Exchange Snapshot copies in each Exchange database volume (with a prefix of exchsnap__) and 20 SnapManager for Exchange Snapshot copies in each Exchange transaction log volume (10 with a prefix of exchsnap__ and 10 with a prefix of eloginfo__). If the SnapManager for Exchange retention policy for a particular server is 10 days of backups online (-RetainDays 10 parameter in the `new-backup` command), there will be SnapManager for Exchange Snapshot copies in the Exchange database and transaction log volumes no older than 10 days. An alternate way of calculating SnapManager for Exchange retention when using the -RetainDays backup parameter is to multiply the number of days you keep backups online by the number of backups taken each day. If 1 backup per day is taken, then there would be 10 SnapManager for Exchange Snapshot copies in each Exchange database volume and 20 SnapManager for Exchange Snapshot copies in each Exchange transaction log volume.

To monitor the health of SnapManager retention, use Windows PowerShell™ commands from the Data ONTAP PowerShell Toolkit and Windows native Windows PowerShell commands. See some sample Windows PowerShell scripts in the appendix.

4.10 NetApp FlexClone

A FlexClone volume is a writable point-in-time Snapshot copy of a FlexVol volume or another FlexClone volume. FlexClone uses space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and the clone. FlexClone volumes are great for any situation in which testing or development occurs, any situation in which progress is made by locking in incremental improvements, and any situation in which there is a desire to distribute data in changeable form without endangering the integrity of the original. A common scenario is to use FlexClone in an environment before committing a Microsoft Exchange Server 2010 rollup or hotfix into production.

FlexClone technology can be leveraged both at the primary storage system and at the SnapMirror destinations for effective utilization of resources. FlexClone can also be used for disaster recovery testing without affecting the operational continuity of the Microsoft Exchange Server 2010 environment.

FlexClone can be created from the storage controller console using the following command:

```
vol clone create cl_vol_name [-s {volume | file | none}] –b f_p_vol_name [parent_snap]
```

For example:

```
vol clone create Exchange_clone -s file –b Exchange SME_snap
```

The preceding example creates a new read-write FlexClone volume named Exchange_clone, based on an existing volume named Exchange from the point-in-time Snapshot copy name SME_snap.

Refer to FlexClone documentation in the “Data ONTAP Administration Guide” for more detailed information on how FlexClone works and on command line references.

Best Practice
Use SnapDrive for Windows to create FlexClone volumes. This automates the creation of the FlexClone volumes and connects the LUNs within the clone to the test and development host.

4.11 NetApp Deduplication

The deduplication process only stores unique blocks of data in the volume and creates additional metadata in this process.

Each 4KB block in the storage system has a digital fingerprint, which will be compared to other fingerprints on the volume. If two fingerprints are found to be the same, a byte-for-byte comparison is done of all bytes in the block. If they are an exact match, the duplicate block is discarded, and the space is reclaimed.

The core enabling technology of deduplication is fingerprints. When deduplication runs for the first time on a FlexVol volume, it scans the blocks and creates a fingerprint database that contains a sorted list of all fingerprints for used blocks in the flexible volume.

Deduplication consumes system resources and can alter the data layout on disk. Due to the application I/O pattern and the effect of deduplication on the data layout, the read and write I/O performance can vary.

Note: Setting read_realloc to 'on' for a volume that has enabled deduplication will not affect performance, nor will it reduce storage efficiency.

Note: Deduplication is transparent to Exchange, and the block changes are not recognized by Exchange. Therefore the Exchange database remains unchanged in size from the host's perspective, even though there are capacity savings at the volume level.

Note: Tests have shown space savings on Exchange Server 2010 databases in the 15–35% range.

Best Practices
• Deduplication rates cannot be predicted and shouldn't be used when sizing capacity. Deduplication can provide overhead for user growth and/or Snapshot retention. • NetApp recommends deduplication for database volumes, not for transaction log volumes. • Turn scheduled deduplication on and schedule it for nonpeak hours (late at night). • Replication of a deduplicated volume is supported by using SnapMirror. However, NetApp does not recommend using deduplication with synchronous SnapMirror, since that could add substantial overhead to the storage subsystem and introduce performance overhead to Exchange Server 2010 databases.

Refer to the Storage Management Guide for more detail on configuring deduplication.

5 NetApp Solution for Microsoft Exchange Server 2010

5.1 NetApp Storage Software and Tools

- **NetApp Windows host utilities kit.** This kit should be used in both physical and virtual environments; it configures Windows Server to access virtual disks on a NetApp storage system through the Fibre Channel, iSCSI, or FCoE protocol. It also helps to align the master boot record for the Microsoft VHD file layout, preventing it from getting out of alignment with the underlying NetApp LUN. This is very important for optimal I/O performance.

- **MPIO.** The NetApp Windows host utilities kit uses the Microsoft framework for MPIO, and it helps storage providers develop multiple paths to optimize connectivity with the storage arrays.

- **MPIO load balancing.** This type of load balancing, supported by MPIO, uses multiple data paths between server and storage to provide greater throughput of data than could be achieved with only one connection.

- **MPIO-based fault-tolerant failover.** In this scenario, multiple data paths to the storage are configured. If one path fails, the HBA or NIC fails over to the other path and resends any outstanding I/O.

  - For a server that has one or more HBAs or NICs, MPIO offers support for redundant switch fabrics or connections from the switch to the storage array.

  - For a server that has more than one HBA or NIC, MPIO also offers protection against the failure of one of those adapters directly within the server.

- **NetApp OnCommand host agent.** The OnCommand host agent gathers host-specific data and sends it to NetApp OnCommand. It reports on files, folders, drive size, HBA info, and average

CPU and memory usage. It helps applications like Exchange by showing how the server is running and tracks the hourly average of many data points.

- **SnapDrive for Windows.** This application helps with storage provisioning and managing disks in both physical and virtual environments. SnapDrive for Windows manages the LUNs on the storage system, making them available as local disks on Windows hosts. Here are the key features of SnapDrive for Windows:
  - Enhances online storage configuration, LUN expansion, and shrinking; provides streamlined management
  - Works in conjunction with NetApp SnapMirror software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes
  - Enables NetApp SnapVault updates of qtrees to a SnapVault destination
  - Enables management of SnapDrive for Windows on multiple hosts
  - Enhances support on Microsoft cluster configurations
  - Simplifies iSCSI session management
  - Enables technology for SnapManager for Exchange products

5.2 SnapManager for Exchange Server Overview

SnapManager for Exchange provides an integrated data management solution for Microsoft Exchange Server 2010 that enhances the availability, scalability, and reliability of Exchange databases. SnapManager for Exchange provides rapid online backup and restoration of databases, along with local or remote backup set mirroring for disaster recovery.

SnapManager for Exchange uses online Snapshot technologies that are part of Data ONTAP. It integrates with Exchange backup and restore~~s~~ APIs and the Volume Shadow Copy Service (VSS). SnapManager for Exchange uses SnapMirror to support disaster recovery.

SnapManager for Exchange provides the following data management capabilities:

- Migrating Exchange databases and transaction logs to NetApp LUNs
- Backing up Exchange databases and transaction logs from NetApp LUNs
- Verifying Exchange databases and transaction logs in backup sets
- Managing backup sets
- Archiving backup sets
- Restoring Exchange databases and transaction logs from previously created backup sets

Some of the new features released in SnapManager for Exchange 6.0.2 include:

- Improved backup performance
- Backup retention management enhancements
- Gapless database availability group (DAG) backup feature
- Copy backup support
- Enterprise monitoring and reporting enhancements
- Business continuance model (BCM) enhancements for Exchange 2007

5.3 SnapManager for Exchange Server Architecture

SnapManager for Microsoft Exchange supports both Microsoft Exchange Server 2007 and Microsoft Exchange Server 2010. SnapManager for Exchange is tightly integrated with Microsoft Exchange, which allows consistent online backups of Microsoft Exchange environments while leveraging NetApp Snapshot copy technology. SnapManager for Exchange is a VSS requestor, which means that it uses the VSS subsystem supported by Microsoft to initiate backups. SnapManager for Exchange works with a DAG, providing the ability to back up and restore data from both active database copies and passive database copies.

For more information about VSS, refer to Microsoft's Volume Shadow Copy Service Overview.

Figure 1) SnapManager for Exchange Server architecture.

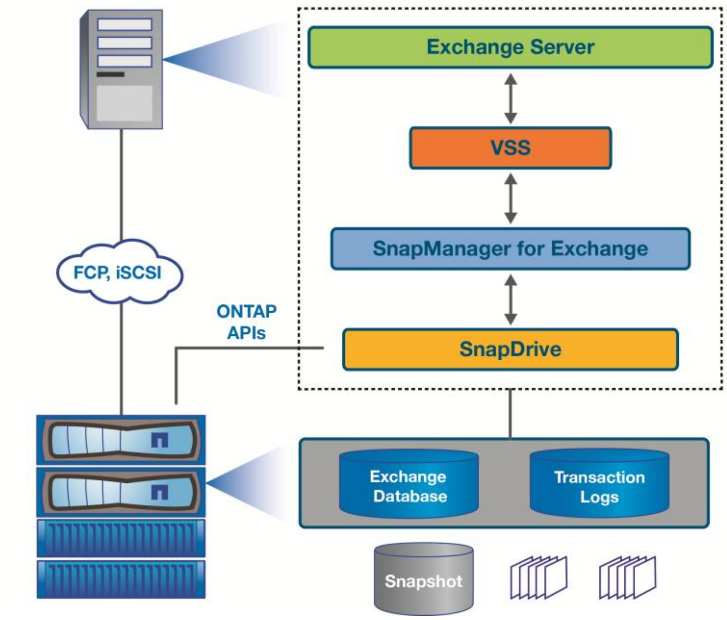

5.4 SnapManager for Exchange Server Installation and Upgrade Considerations

For information about compatible versions of SnapManager for Exchange, SnapDrive for Windows, and Data ONTAP, see the SnapManager and SnapDrive Compatibility Matrix.

Before upgrading SnapManager for Exchange, consider the following steps.

- Back up the operating system installation on the Exchange server. This includes backing up all of the server system state, which consists of the registry, the boot files, and the COM+ class registry.
- Back up the data on the local drives on the Exchange server.
- Back up the boot and system drives.
- Use your backup utility to create and maintain a current emergency repair disk (ERD).

Best Practice
It is a NetApp best practice to install SnapManager for Exchange and SnapDrive for Windows on all member servers of the DAG. SnapDrive for Windows must be installed on all member servers of the DAG when using a heterogeneous environment.

5.5 SnapManager for Exchange 6.0.2: Gapless Backup

Configuration: Less than 10 total databases per node in a 2-node configuration.

The gapless backup feature is designed to make sure that a Snapshot copy that is older than the most recent full backup, which truncates the transaction logs, can utilize up-to-the-minute restore (roll-forward recovery).

Figure 2) SnapManager for Exchange example.

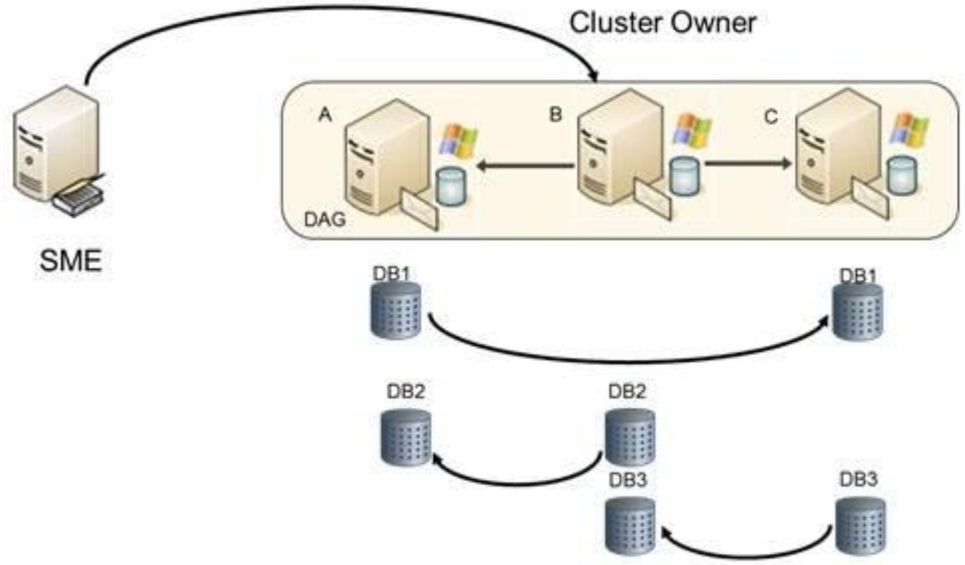

SnapManager for Exchange (SME) Example

- Three-node DAG; Node-B is the owner node of the DAG.
- Each database has two copies. The arrow indicates the log shipping direction, pointing to the passive database copy.
- SME is shown running on a client machine, though it can be run from any of the DAG nodes.

The administrator selects a “full backup” on all of the databases. Additional remote copy backup needs to be selected to perform a copy backup on all passive database copies.

Sequence of Operation

When SME connects to Node-B, which is the DAG owner, it gets the list of databases to back up and then starts the backup operation.

Figure 3) Sequence of operation.

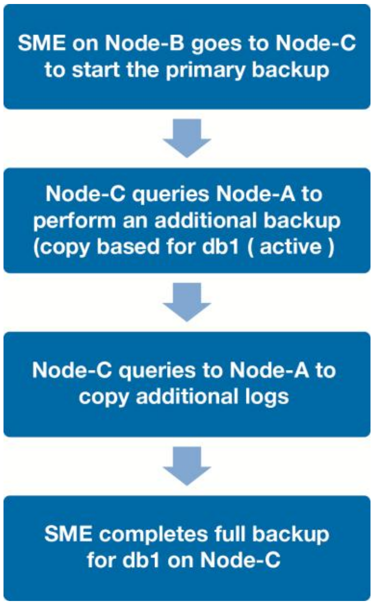
SME on Node-B goes to Node-C to start the primary backup
Node-C queries Node-A to perform an additional backup (copy based for db1 (active)
Node-C queries to Node-A to copy additional logs
SME completes full backup for db1 on Node-C

In the end you will have a backup for each database copy (6), and all backups can perform an up-to-the-minute restore. Since the job was initiated on Node-B, the overall operation resides on Node-B.

5.6 SnapManager for Exchange: Server Backup and Frequent Recovery Point

Configuration: More than 10 databases per node and more than 2 nodes.

A smaller recovery point objective can be achieved in larger environments with many databases and server nodes in a DAG. This is a sample configuration that was successfully implemented at a couple of large customer environments. The strategy employed utilizes a “full backup” on one copy of the database that does not need to be the active database copy. The remaining database copies run a “copy backup,” which does not truncate logs. All databases participate in a frequent recovery point, or transaction log backup every 30 minutes. Should one copy of the database fail, the rapid reseed process can be used to perform a nondisruptive restore.

00:30–02:00 database copy 4 full backup
02:00–03:00 database copy 3 copy backup
03:00–04:00 database copy 1 and database copy 2 copy backup
03:00–05:00 database copy 3 and database copy 4 FRP backup
05:00–00:00 all database copies FRP backup every 30 minutes

Figure 4) Backup schedule.

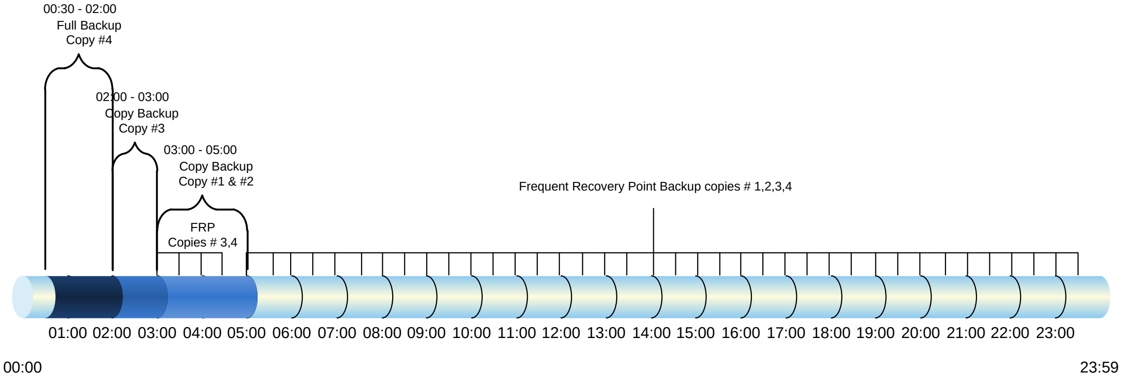

Note: It is important that a copy backup is scheduled at least 10 minutes after the full backup. This is to make sure that the log truncation activity from the full backup has successfully replicated to each database copy.

The rapid reseed process leverages Windows PowerShell, Exchange Windows PowerShell commands, and SnapDrive commands to provide a rapid nondisruptive restore of a failed Exchange 2010 database copy. The process is detailed in the Microsoft SIG blog post: https://communities.netapp.com/community/netapp-blogs/msenviro/blog/2011/12/01/don-t-reseed-that-exchange-server-2010-database. The rapid reseed script is in the appendix.

The steps for an incremental reseed are:

1. Suspend failed database copy
2. Remove failed database copy
3. Add the database copy with seeding postponed
4. Restore the database LUN from recent Snapshot copy
5. Resume database copy

The rapid reseed restore process leverages the native Exchange incremental reseed using any healthy database copy as the source for the necessary transaction log files.

If there is a failure in the seeding, the application event log on the server where the database copy is being restored will have an event log that will show the logs that are needed to complete the incremental reseed. On the server where database copy 4 is located, manually copy the missing logs from the appropriate SnapInfo directory and into the transaction log directory on the source server. You can also run eseutil /mh on the restored database file to see which log files are needed for the reseed.

A database can also be restored from a full backup with FRP or a copy backup with FRP using SME.

Best Practice
NetApp recommends using server-based full and copy backups with frequent recovery points to achieve the smallest recovery point objective (RPO).

5.7 SnapManager for Exchange Management

The SnapManager for Exchange service must be a member of the Exchange server's local administrators group.

5.8 Sizing and Storage Layout for Exchange Server 2010

5.8.1 Aggregate Recommendations

Fewer, larger aggregates will maximize performance; however, they might not meet the data availability requirements set forth in the SLA. In Exchange Server 2010 environments with multiple database copies, Microsoft no longer requires separating database and transaction log files to separate sets of disks. This means that database and transaction log volumes can be placed in the same aggregate. Each database copy of the same database must be placed in a separate aggregate.

Best Practices

- NetApp recommends having at least 10% free space available in an aggregate hosting Exchange data for optimal storage performance.
- On controllers that are isolated to the Exchange server, setting the global wafl.optimize_write_once flag to off can optimize random workloads. The flag must be set before Exchange aggregates are created. If Exchange aggregates are already present, reallocate –A must be run on each Exchange aggregate. This is a time-consuming process that can affect performance during the reallocation scan.

5.8.2 Volume Planning and Layout

Data ONTAP enables the creation of flexible volumes for managing data without the need to assign physical disks to the volumes. Instead, the flexible volumes (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 Server 2010 environments:

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

Best Practices

- NetApp recommends **separating database and transaction logs** from different servers into separate volumes to prevent a potential “busy” Snapshot copy problem. Utilizing separate volumes for each server reduces complexity, since there is no concern about Snapshot copy schedules overlapping different servers.
- NetApp recommends having at least **10% free space** available in a volume hosting Exchange data.
- NetApp recommends placing each database in a separate volume with copies of the same database isolated in separate aggregates.

Volume sizing is different for transaction logs and database volumes. Transaction log sizing involves calculating the size of the transaction log LUN(s) in the volume, adding space for the Snapshot retention length, plus 10% free space in the volume.

Transaction Log Volume

A. Transaction log LUN size: 2561GB
    a. Holds 800 users sending 100 messages a day.
B. Messages per day 80,000 * 75KB = 5.72GB per day
C. Snapshot retention of 7 days with 3 days of fault tolerance is 10 days.
D. 5.72GB * 10 days = 57.2GB

E. Log volume (57.2GB + 2561GB / (1-.1)) or 2909GB
   a. 10 days of Snapshot copies + transaction log LUN + 10% free space in the volume.

Database Volume

A. Database LUN size: 2560GB (2TB database)
B. If you size the database LUN for quota, this includes maximum mailbox size, deleted items in the dumpster, calendar, 3 days of incoming mail, and whitespace in the database
C. Snapshot retention of 7 days with 3 days of fault tolerance is 10 days.
D. Daily change rate (2%) * 10 days
E. 2560 + (2048 * 20%) = 2970GB

5.8.3 LUN Planning and Layout

A database and its corresponding transaction log must be placed on separate LUNs for SnapManager for Exchange. In environments with high LUN counts, transaction logs for multiple mailbox databases can be consolidated on a single LUN. NetApp recommends limiting the number of transaction log streams per LUN to fewer than 10.

Best Practices
• When creating LUNs, use volume mount points. There are a finite number of drive letters, and in a DAG each database path must be the same on every server that has a copy of that database. • Place each database on a separate LUN in a separate volume. • Use larger databases. Microsoft supports up to 16TB databases with a best practice size of 2TB. Many customers, including NetApp IT, run Exchange Server 2010 with larger than 2TB databases on NetApp storage.

Do not create mount points for additional LUNs on another LUN that holds an Exchange Server 2010 database. If you have to complete a restore of a database residing on a LUN with volume mount points, the restore operation removes any mount points that were created after the backup, disrupting access to the data on the mounted volumes referenced by these volume mount points.

It is a NetApp best practice to place the transaction logs and database files on separate LUNs. These calculations are for the primary active database and its corresponding transaction log files. Each additional copy of the database would require a multiple of the sizing. The same calculations can be used to estimate the size of the archive database and its corresponding transaction log files.

Database

The database LUN houses the 20% free disk space, the database itself, and the content index files.

1. The MBXSize is the MBXLimit plus the dumpster.

2. The MBXLimit is the stated maximum mailbox size; in our example case, that is 5GB.

3. Calculate the space consumed in the dumpster, which also includes space consumed by enabling both single-item recovery and calendar version storage.

   Single-item recovery = MBXLimit * 0.02 (2%)

   Calendar version storage = MBXLimit * 0.03 (3%)

   Dumpster = SingleItemRecovery + CalendarVersionStore + (#Messages * MessageSize * (DeletedItemRetention+1 [today]))

   Example: Assuming a 5GB mailbox and default 1- day retention

   Dumpster = 102 + 154 + (150 * 75KB * (14+1))

= 102 + 154 + 165.04MB = 420.8MB

4. The database size is the MBXSize multiplied by the number of users.

5. The “DB size + overhead” adds 20% to the database size.

6. DB LUN is calculated by adding the DBS size + overhead to the content index, while padding in 20% free disk space.

   Example:

   a. DB size of 1024GB
   b. DB size + overhead = 1228.8GB
   c. ContentIndex = 1024 * 10% = 102.4GB
   d. (1228.8 + 102.4) / (1-.2) = 1664GB

Transaction Logs

The transaction logs are 1MB in size and must include transaction logs generated by the users from move mailbox requests and the backup fault tolerance window. Microsoft by default assumes a 3-day backup fault tolerance window.

A) User logs (calcnumusertlogs): LogGen * users * Datagrowth

   a. LogGen = 10 for each 50 messages per day in a user profile
   b. Data growth by default pads in an extra 20% to capacity
   c. Users: An example 100 message per day, 1000-user configuration would be:
      i. 20 * 1000 * 20% = 24,000 logs

B) Move mailbox (logdiskspacereqmove)= (users * 1%) * (MBXSize * 1.2)

      i. (1000 * 1%) * (5120 * 1.2)
      ii. 10 * 6144
      iii. 61,440GB

C) Log backup (logdiskspacereqbackup)

   a. UserLogs * BackupFailTol (default 3)
   b. 24,000 * 3 = 72,000 logs

D) Total log disk space (totlogdiskspace)

   a. Log backup + move mailbox
   b. (72,000 /1024) + 61440GB
   c. 61,510GB (60.04TB)

E) Add 20% free disk space

   a. 60.04TB / (1-.2) = 75.05TB

F) Divide by the number of databases, in this case, 30

   a. 75.05TB / 30 = 2.501TB per transaction log LUN
   b. Notice how the majority of this space is because of move mailbox and large mailbox sizes.

5.8.4 SnapInfo Data and LUN

The SnapInfo directory is the central repository for all SnapManager for Exchange–related activities. This directory contains the backup metadata and reports as well as truncated transaction log files.

In SnapManager for Exchange, if the SnapInfo directory is placed in the same LUN as its corresponding transaction log, then SnapManager for Exchange stores NTFS hard links to transaction log files in the SnapInfo directory during backup. This saves space and decreases transaction log backup time.

Best Practices
• Place the transaction log files and the SnapInfo directory on the same LUN. • If a database's transaction log files and the SnapInfo directory are placed on separate LUNs, place them both in the same volume.

5.8.5 Exchange 2010 Server Database Cache

The best predictors for Exchange Server 2010 mailbox user transactional IOPS are the amount of database cache per mailbox and the number of messages each user sends and receives per day. Microsoft has guidance on the amount of database cache required for a particular user profile, which can be used to accurately size the RAM on the mailbox server. For more details, see the Microsoft TechNet article Understanding the Mailbox Database Cache.

Transaction Log Capacity Considerations

Microsoft has created guidance on user profiles that are in 50 messages per day increments utilizing a 75KB average message size. Each 50 messages per day will cause approximately 10 transaction logs to be generated. For more details, see the Microsoft TechNet article Understanding Mailbox Database and Log Capacity Factors.

5.9 Capacity Planning

A properly sized Exchange environment will meet or exceed the customer service-level agreement (SLA). To properly size an environment, information from the customer environment is collected, and tools are used to convert that information into a physical storage recommendation.

Two primary tools should be utilized when planning an Exchange environment for a customer:

- The Microsoft Exchange 2010 Mailbox Server Role Requirements Calculator
- The NetApp Exchange Sizing Tool (SPM); work with your local NetApp partner or your NetApp representative to enable proper sizing

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 cannot accurately make recommendations on proprietary storage technology because the storage design largely depends on the type of storage array being utilized. When sizing Exchange server deployments using NetApp storage, it is important to use the NetApp Exchange Sizing Tool with the data from the Microsoft Exchange 2010 Mailbox Server Role Requirements Calculator.

Best Practice
Consult a local NetApp Exchange expert or your NetApp partner to assist in accurately sizing Exchange Server 2010. Use the NetApp Sizing Tool for Exchange (SPM) to size all Exchange server deployments utilizing NetApp storage.

5.9.1 Backup Design Considerations

In this section we focus on the backup of databases and transaction logs on NetApp storage using SnapManager for Exchange. It is important to consider the following factors for planning a backup strategy:

- Service-level agreement (SLA)
- High availability and disaster recovery planning
- Backup verification policy

The recovery time objective (RTO) to return a database to service is affected by the number of transaction logs that need to be replayed. A more frequent backup window will reduce the number of transaction logs that need to be replaced, shrinking the RTO.

Database Availability Group

In Exchange Server 2010, mailbox servers can be grouped to form a DAG. A DAG is a high-availability feature of Microsoft Exchange Server 2010 that provides database-level recovery from failures and data corruption. A DAG can contain up to 16 mailbox servers in which each server can have a copy of a database. The DAG is created for mailbox databases and not for public folder databases.

A DAG through the active manager provides automatic recovery from a database, server, or network failure. The current active database and its copies use the same path on each server.

Database Verification

Database verification is not a support requirement for databases with at least two copies in a database availability group. By default, when performing a DAG backup with SnapManager for Exchange, verification will be off. You can monitor both locally running and remote verification jobs through the main SnapManager for Exchange management console. A single Exchange mailbox server can only run one verification process at a time on a particular verification server. A verification server can simultaneously run one verification job from each Exchange mailbox server. More than one verification server can be used in order to simultaneously verify more than one backup job on a single Exchange mailbox server. Many customers utilize virtual machines to offload the verification.

SnapManager for Exchange also supports the verification of backups that have been archived to the SnapMirror and SnapVault secondary location.

5.9.2 Preparing Exchange Server Databases for Migration to NetApp Storage

SnapManager for Exchange makes it easy to move databases from local storage to NetApp storage utilizing the Configuration Wizard. All the databases are automatically mounted after the migration is completed, and NetApp recommends backing up the databases soon after the migration.

Best Practices
• The Exchange Server 2010 database and transaction log path must be unique for each database. • SnapManager for Exchange uses a host-based licensing mechanism, which means SnapManager for Exchange licenses should be purchased for each member server in the DAG even if SnapManager for Exchange is not intended to be installed on each member server.

5.9.3 Snapshot Retention Guidelines

Primary Storage

The recovery point objective (RPO) will guide how frequently a backup is taken. NetApp flexible volumes running Data ONTAP operating in 7-Mode can store a maximum of 255 Snapshot copies per flexible volume. The amount of storage needed for Snapshot copies will depend on the rate of change.

Consult a local NetApp Exchange expert or your NetApp partner to provide accurate volume sizing and layout for Exchange environments.

Secondary Storage

SnapManager for Exchange backups can be archived with SnapVault using the SnapManager for Exchange integration with NetApp OnCommand Protection Manager.

Long-Term Archival to Tape

Long-term archival of SnapManager for Exchange backups to tape can be done by either using an NDMP-based backup to copy the LUN in the Snapshot copy created by SnapManager for Exchange to tape or by mounting the LUNs in the Snapshot copy created by the SnapManager for Exchange backup and then streaming to tape.

Best Practices
• Use the business requirements established by the Exchange stakeholders to help determine the number of Snapshot copies to keep online. • Use NetApp OnCommand Protection Manager to archive SnapManager for Exchange backup sets from primary to secondary storage. • If mounting LUNs in the Snapshot copy created by the SnapManager for Exchange backup to archive the SnapManager for Exchange backup to tape, license the controller with FlexClone so that there are no busy Snapshot copies. • When performing an up-to-the-minute restore and database verification is utilized, restore from your most recently verified backup.

5.9.4 Restore Guidelines

SnapManager for Exchange’s restore functionality allows you to recover your Exchange databases and transaction logs from backups that it created or from SnapVault archives. There are two types of restore operations in SnapManager for Exchange:

- **Up-to-the-minute.** Selected by default, an up-to-the-minute restore mounts the database, and Exchange replays the transaction logs from the backup set and from the transaction log directory and applies them to the database. A contiguous set of transaction logs is required for an up-to-the-minute restore to succeed.

- **Point-in-time.** This option allows you to restore your Exchange data to a chosen point in time. Any Exchange data past that point is not restored. This option is particularly useful when trying to restore to a point before something such as data corruption occurred. A point-in-time restore only replays and applies to the database those transaction logs that existed in the active file system when the backup was created up to the specified point in time. All transaction logs beyond that point in time are discarded.

Best Practice
When performing an up-to-the-minute restore, restore from your most recently verified backup to minimize the number of transaction logs that must be replayed.

5.9.5 Single Mailbox and Item-Level Recovery

The NetApp Single Mailbox Recovery (SMBR) tool allows the customer to extract e-mails and other items from an Exchange database and place them in a PST file or a live mailbox. In order to extract e-mail from an Exchange database, you need the LUN to be mounted using SnapDrive for Windows.

The NetApp Single Mailbox Recovery Administrative Server (SMAS) is a framework that can host centralized services to multiple clients and provide both client and server support for NetApp Single Mailbox Recovery and NetApp Single Mailbox Recovery Extract Wizard users. Single Mailbox Recovery Administrative Server activates mailbox permissions, centralizes administration of certain application settings, and provides auditing services for Single Mailbox Recovery and Single Mailbox Recovery Extract Wizard clients.

Best Practice
Always use SnapDrive for Windows to mount the LUN to access the Exchange databases.

The archive mailbox can be placed in a different database than a user's primary mailbox starting with Exchange Server 2010 SP1. The following steps are necessary to import items to an archive mailbox.

SnapManager for Exchange 6.0 and later handles all the back-end work of presenting the Snapshot copies and then cleaning up after the queries, and mail is moved to the appropriate mailbox or PST file.

Figure 5) SMBR recovery.

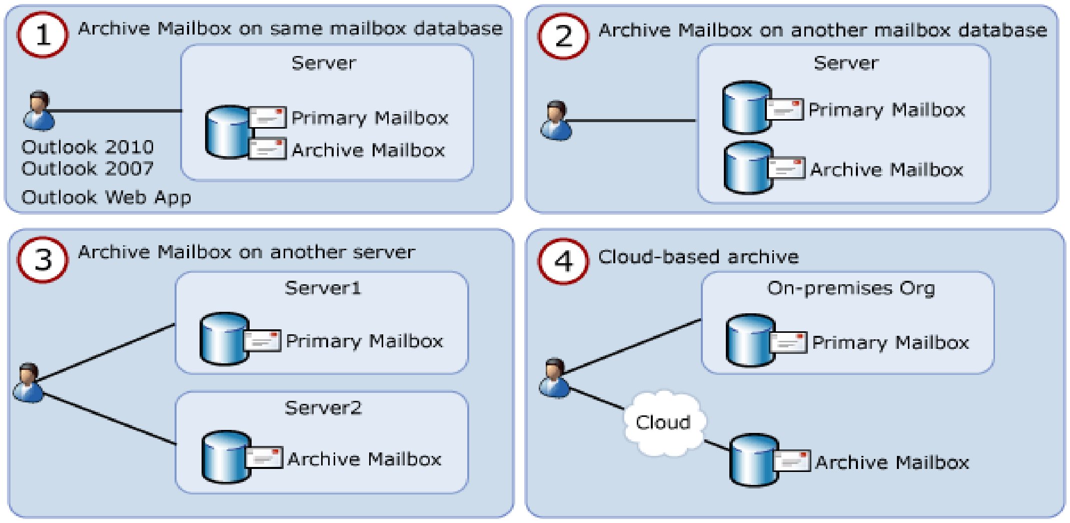

This is a two-step process:

- Export the mail from EDB to a PST.
- Run the following Exchange Windows PowerShell cmdlet to import the PST directly to the archive mailbox using the following command:

```
New-MailboxImportRequest -Mailbox “JohnDoe” –IsArchive -FilePath
\\Server_name\PSTFileShare\John_Doe.pst
```

The IsArchive switch specifies that you're importing the PST file into the user's archive mailbox.

5.9.6 Troubleshooting

SnapManager for Exchange reports list the step-by-step details of every SnapManager for Exchange operation that is performed, their final status, and any error messages that are encountered during the operation. The SnapManager for Exchange report directory provides subfolders that group the reports for each operation type.

The following troubleshooting steps can be followed to gather additional information:

1. Enable debug logging on all nodes.
2. Restart the SME service on all nodes.
3. Identify which operation on which node failed, based on the SME operation sequence.
4. Go to the node with the failure and find the backup report and debug log under \Backup\Server_name\, and \Debug\Server_name\.
5. Use the server-level backup report and debug log to find the root cause of the problem.

For more information on troubleshooting, refer to the NetApp Knowledge Base.

6 Performance

Accurately sizing NetApp storage controllers for Exchange workloads is essential for good Exchange performance and so that Exchange service levels are met. Consult a local NetApp Exchange expert to provide accurate performance sizing and layout for Exchange environments.

6.1 SATA Performance Considerations

SATA-based deployments of Exchange must take into account that SATA drives have a lower I/O profile than SAS and FC disk. The I/O profile of a 7,200-RPM SATA drive is around 45–55 IOPS at a 20ms response time.

Exchange 2010 utilizes background database maintenance (BDM) to maintain the consistency of the databases. BDM applies a per-database performance tax on the storage system that must be taken into account when sizing the storage for Exchange. Having fewer, larger databases in the Exchange database design helps reduce the amount of background database maintenance I/O, which in some cases can exceed the transactional I/O generated by users. This is typically seen in designs in which there are a large number of small databases.

To help improve the storage efficiency and read I/O performance and latency of SATA-based deployments, Flash Cache™ should be used. Flash Cache is a read cache that can be installed on certain models of NetApp storage controllers. Flash Cache enables fewer SATA disks to be used in SATA-based deployments, because a percentage of the Exchange database working set is cached in the Flash Cache, thus greatly reducing the amount of read I/O on the SATA disk. NetApp recommends Flash Cache and SATA for deployments exceeding 1,000 mailboxes or when SATA-based designs are bounded by performance instead of capacity.

Best Practice
NetApp recommends using Flash Cache when placing Exchange Server 2010 database files on SATA physical disk drives.

6.2 Database Sizing Considerations

Using a smaller number of larger databases can help reduce the amount of background database maintenance I/O as well as reduce the complexity of the storage design. NetApp recommends using a database size of at least 2TB with at least two copies in a DAG, or 200GB for non-DAG databases. A database size of 2TB is a practical size that can be restored in minutes with SnapManager for Exchange. Microsoft recommends 2TB as a maximum size, but supports up to 16TB for databases on both standard and enterprise server editions.

Best Practice
NetApp recommends using a database size of at least 2TB. With the protection of clustered storage controllers, RAID-DP, and Snapshot backups, many customers have deployed databases larger than 2TB, which significantly reduces the database maintenance I/O.

6.3 Aggregate Sizing and Configuration Considerations

Aggregates are sized for performance and storage capacity in storage designs that support Exchange workloads as well as maintain data protection for the Exchange data.

As mentioned earlier in this document, Exchange databases and transaction logs can be placed on the same aggregate. There is marginal benefit in locating transaction logs and databases on separate aggregates. However, putting DAG database copies on separate aggregates and/or controllers enables at least one copy of the Exchange data to exist if an aggregate is lost due to a catastrophic failure. Placing database copies on separate aggregates also helps isolate background database maintenance I/O to the aggregates where the database copies are located.

If Exchange is virtualized, place the Exchange VMs on a separate aggregate from the Exchange data. Certain mailbox roles, such as the hub transport server, might affect Exchange performance if the virtual machine is located on the same aggregate as the Exchange data.
So that Exchange performance is not affected by disk drive reconstruction, determine that there are at least two spare disks per controller and that the Data ONTAP option disk.maint_center.enable is enabled. (It is on by default but requires two hot spares.) Maintenance Center is a feature in Data ONTAP that can prefail a disk if the disk does not pass a certain number of diagnostic tests. For more details on hot spares, see the System Management Guide for the version of Data ONTAP that is installed.

The aggregates for Exchange should be configured for the RAID-DP RAID level. This enables maximum data protection for the Exchange data so that an aggregate can survive a double disk failure in any RAID group of that aggregate.

The RAID group size of the aggregate affects the level of data protection, speed of recovery, and available data storage space. Configuring an optimum RAID group size for an aggregate requires a trade-off of factors. Adding more data disks to a RAID group increases the striping of data across those disks, which typically improves I/O performance. Additionally, a smaller percentage of disks is used for parity rather than data. However, with more disks in a RAID group, there is a greater risk that one of the disks might fail. The recommendation is to use the default RAID group size when the Exchange aggregate is created, because this balances storage efficiency and performance.

Exchange workloads can run effectively on both 32-bit and 64-bit aggregates if the aggregates and controller heads are properly sized. NetApp recommends that 64-bit aggregates supporting an Exchange workload be used only in configurations that are supported in the NetApp Exchange sizing tool. This is so that the storage is properly sized for the anticipated Exchange workload. NetApp recommends consulting a local NetApp Exchange expert to provide accurate performance sizing when considering the use of 64-bit aggregates for Exchange environments.

6.4 Volume Configuration Considerations

NetApp recommends setting the volume option read_realloc on each database volume. This is particularly helpful in environments with many databases and the corresponding sequential read due to the background database maintenance.

7 Virtualization

7.1 Microsoft Support for Exchange Server 2010 in Virtualized Environments

The documentation for support for Exchange 2010 in virtualized environments can be found in the Microsoft TechNet article Exchange 2010 System Requirements.

Here is a high-level list of some important considerations:

- Exchange Server 2010 SP2 virtual machines (including Exchange mailbox virtual machines that are part of a DAG) may be combined with host-based failover clustering and migration technology as long as the virtual machines are configured such that they will not save and restore state on disk when moved or taken offline.
- All storage used by an Exchange guest machine for storage of Exchange data must be block-level storage because Exchange 2010 doesn't support the use of network-attached storage (NAS) volumes. Also, NAS storage that's presented to the guest as block-level storage using the hypervisor isn't supported.
- Microsoft does not support the use of dynamic virtual disks to store Exchange data.
- Microsoft does not support the use of differencing disks or Snapshot copies of virtual disks storing Exchange data.
- Microsoft does not support the use of virtual machine Snapshot copies of Exchange virtual machines.
- Microsoft recommends that both shared memory and hypervisor-based autotuning be disabled.

8 High Availability

In Exchange 2010, the database availability group (DAG) feature was implemented to support mailbox database resiliency, mailbox server resiliency, and site resiliency. The DAG consists of two or more servers, and each server can store up to one copy of each mailbox database.

Transaction log replication is used by the DAG so that each database copy is identical. The DAG also leverages a feature on the Exchange hub transport servers called shadow redundancy. Shadow redundancy is enabled by default and is used to store copies of messages until the message is delivered and replicated to each DAG member.

The DAG Activation Manager manages the database and mailbox failover and switchover processes. A failover is an unplanned failure, and a switchover is a planned administrative activity to support maintenance activities.

The database and server failover process is an automatic process when a database or mailbox server incurs a failure. The order in which a database copy is activated is set by the administrator.

For more information on Exchange 2010 DAGs, refer to the Microsoft TechNet article Understanding Database Availability Groups.

8.1 Exchange 2010 Database Availability Group 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 utilizing 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 enables 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 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 will provide site resiliency but also provide high availability in each site.

For additional information on DAG layout planning, refer to the Microsoft TechNet article Database Availability Group Design Examples.

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

Deployment	
Best practice	In a multisite scenario it is a best practice to 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 will provide site resiliency but also provide high availability in each site.
Storage Design	
Best practice	Design identical storage for active and passive copies of the mailboxes in terms of capacity and performance.
Best practice	Provision the active and passive LUNs identically regarding path, capacity, and performance.
Best practice	Place flexible volumes for active and passive databases onto separate aggregates. If a single aggregate is lost, only the database copies on that aggregate are affected.
Volume Separation	
Best practice	Place active and passive copies of the database into separate volumes.
Backup	
Best practice	Perform a SnapManager for Exchange full backup on one copy of the database and a copy-only backup on the rest of the database copies.
Best practice	Verification of database backups is not required if Exchange 2010 is in a DAG configuration with at least two copies of the databases, with Exchange background database maintenance enabled. Verification of transaction log backups is still required.
Best practice	Verification of database backups and transaction log backups is required if Exchange 2010 is in a standalone (non-DAG) configuration.
Best practice	In Exchange 2010 standalone environments using SnapMirror, configure database backup and transaction log backup verification to occur on the SnapMirror destination storage.

9 Exchange 2010 Disaster Recovery

Extending an Exchange 2010 DAG across multiple sites provides site resiliency of Exchange services. The DAG functionality relies on transaction log shipping as the data replication mechanism for high availability and site resiliency. NetApp SnapMirror does not integrate with the Exchange 2010 third-party replication API, so SnapMirror is not used for data replication between DAG nodes.

For environments in which Exchange 2010 is deployed in standalone (non-DAG) configurations, SnapMirror replication can be used with SnapManager for Exchange to provide site resiliency for Exchange services. The SnapManager for Exchange business continuance module is not supported with Exchange 2010.

NetApp SnapMirror

NetApp SnapMirror is a storage-based replication mechanism that allows data replication to occur between two NetApp storage controllers. SnapManager for Exchange uses SnapMirror in asynchronous mode only.

When using SnapMirror with SnapManager for Exchange, make sure that the flexible volumes on the SnapMirror destination are configured with the same options as the flexible volumes on the primary storage controllers. It is also important to size the flexible volumes on the SnapMirror destination to be the same size as or greater than the flexible volumes on the primary storage controllers.

SnapManager for Exchange only supports asynchronous SnapMirror replication. Determine that SnapMirror schedules are set for manual update, so that SnapManager for Exchange will trigger replication updates after a successful backup.

Refer to the Data ONTAP documentation for information on how to configure and initialize SnapMirror replication. The NetApp Communities site has many Windows PowerShell scripts that leverage the Data ONTAP PowerShell Toolkit. One such script (Exchange 2010 Rapid Database Seeding) takes a healthy Exchange 2010 database, copies it with SnapMirror to a destination controller, and mounts it on the destination Exchange server. This is useful in environments with poor latency and when stretching the DAG across a WAN is not viable.

9.1 Exchange 2010 Disaster Recovery Process for DAG Configurations

When a single server or database is lost, the high-availability features of Exchange 2010 DAGs automatically perform switchovers to activate new database copies on the same server or on a different server to keep Exchange services online.

In the case of a primary data center loss, the disaster recovery process is a controlled event and is initiated manually. The process is called a data center switchover. Enabling data center activation coordination (DAC) mode on the DAG helps prevent split-brain DAG scenarios.

The process for data center switchovers can be found in the Microsoft TechNet article Datacenter Switchovers.

For more information on data center activation coordination mode and how to configure a DAG for DAC mode, see the Microsoft TechNet article Understanding Datacenter Activation Coordination Mode.

9.2 Exchange 2010 Disaster Recovery Process for Standalone Physical Mailbox Server Configurations

Exchange server recovery should be used with SnapManager for Exchange to support the disaster recovery process for standalone (non-DAG) Exchange 2010 mailbox physical servers.

Exchange server recovery prerequisites and procedures can be found in the Microsoft TechNet article Recover an Exchange Server.

Prerequisites

- The server on which recovery is being performed must be running the same operating system as the lost server. For example, you can't recover a server that was running Exchange Server 2010 and Windows Server 2008 on a server running Windows Server 2008 R2, or vice versa.
- The same disk drive letters and/or volume mount points on the failed server for mounted databases must exist on the server on which you're running recovery.
- The server on which recovery is being performed should have the same performance characteristics and hardware configuration as the lost server.
- The following procedure can be run on a server running Exchange Server 2010 that has the client access, hub transport, mailbox, or unified messaging server roles installed. You can't use **Setup /m:RecoverServer** to recover an edge transport server. For information about preserving edge transport server settings and applying saved settings to an edge transport server, see Understanding Edge Transport Server Cloned Configuration.

Recovery Procedure

The recovery procedure with SnapManager for Exchange and SnapMirror is as follows.

1. Reset the computer account for the lost server.
2. Install the proper operating system and name the new server with the same name as the lost server. Recovery won't succeed if the server on which recovery is being performed doesn't have the same name as the lost server.
3. Join the server to the same domain as the lost server.
4. Install the necessary prerequisites and operating system components.
5. Install NetApp Windows host utilities, SnapDrive for Windows, and SnapManager for Exchange.
6. Determine that the new server is connected properly using iSCSI or FCP to the SnapMirror destination storage.
7. Use SnapDrive for Windows to connect to the LUNs in the SnapMirror destination. Use the same drive letters or mount points as the original server. SnapDrive for Windows will automatically break the SnapMirror relationship.
8. Log on to the server being recovered and open a command prompt.
9. Navigate to the Exchange 2010 installation files and run the following command:
   `setup /m:RecoverServer.`
10. Use SnapManager for Exchange to recover from the most recent backups.

9.3 Exchange 2010 Disaster Recovery Process for Virtualized Standalone Mailbox Servers

Prerequisites

- The Microsoft requirements for virtualized Exchange mailbox servers can be found in section 7 of this document. For more information about support for Exchange 2010 in virtualized environments, see the Microsoft TechNet article Exchange 2010 System Requirements.
- After the virtualized Exchange servers are created and configured, turn off the virtual machine and create a NetApp Snapshot copy of the volume or LUN where the virtual machines are located.

- The flexible volume containing the virtual machines is in a SnapMirror configuration.
- A virtualization host (Hyper-V™, ESX®) is connected to the NetApp SnapMirror secondary storage.

Procedure (Database and Transaction Log LUNs Managed by Virtual Machine)

1. Determine that the Hyper-V or ESX host in the primary site is offline.
2. Determine that the Hyper-V host or ESX host is connected to the NetApp SnapMirror secondary storage.
3. Use SnapDrive for Windows on the Hyper-V parent host or ESX host to connect to the LUN where the Exchange mailbox virtual machine is located on the NetApp SnapMirror secondary storage.
4. Import the Exchange mailbox virtual machine into the Hyper-V or ESX server and power on the virtual machine.
5. Configure the Exchange mailbox virtual machine to be on the network in the disaster recovery site and update DNS, if necessary.
6. Use SnapDrive for Windows within the Exchange mailbox virtual machine to set up the iSCSI network connections to the NetApp storage in the disaster recovery site.
7. Use SnapDrive for Windows within the Exchange mailbox virtual machine to connect to the database and transaction log LUNs.
8. Use SnapManager for Exchange to restore the mailbox databases in order to restore Exchange services.

Procedure (Database and Transaction Log LUNs Are Hyper-V Pass-Through Disks/ESX RDM LUNs)

1. Determine that the virtualization host in the primary site is offline.
2. Determine that the Hyper-V host or ESX host is connected to the NetApp storage.
3. Use SnapDrive for Windows on the Hyper-V host to connect to the LUN where the Exchange mailbox virtual machine is located on the NetApp SnapMirror secondary storage. For ESX, connect to the NFS datastore or VMFS LUN where the Exchange mailbox virtual machine is located.
4. Import the Exchange mailbox virtual machine into the Hyper-V or ESX server and power on the virtual machine.
5. Configure the Exchange mailbox virtual machine to be on the network in the disaster recovery site and update DNS, if necessary.
6. Configure SnapDrive for Windows to communicate with the new Hyper-V server or ESX server.
7. Use SnapDrive for Windows within the Exchange mailbox virtual machine to connect to the database and transaction log LUNs as Hyper-V pass-through disks or ESX RDM LUNs.
8. Use SnapManager for Exchange to restore the mailbox databases in order to restore Exchange services.

10 Summary

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. This guideline enables 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

SnapManager For Exchange

- SnapManager for Exchange uses a host-based licensing mechanism, which means that SnapManager for Exchange licenses should be purchased for each member server in the DAG even if SnapManager for Exchange is not intended to be installed on each member server.

- It is a best practice to install SnapManager for Exchange and SnapDrive for Windows on all member servers of the DAG. SnapDrive for Windows must be installed on all member servers of the DAG.

- If mounting LUNs in the Snapshot copy created by the SnapManager for Exchange backup to archive the SnapManager for Exchange backup to tape, license the controller with FlexClone so that there are no busy Snapshot copies.

- When performing an up-to-the-minute restore, restore from your most recently verified backup to minimize the number of transaction logs that must be replayed.

- Always use SnapDrive for Windows to mount the LUN to access the Exchange databases.

- Verification of database backups is not required if Exchange 2010 is in a DAG configuration with at least two copies of the databases with Exchange background database maintenance enabled. Verification of transaction log backups is still required.

- Verification of database backups and transaction log backups is required if Exchange 2010 is in a standalone (non-DAG) configuration.

- In Exchange 2010 standalone environments using SnapMirror, configure database backup and transaction log backup verification to occur on the SnapMirror destination storage.

- Perform a SnapManager for Exchange full backup of the active database copies and copy only backup of the passive database copies.

Storage Design and Layout

- NetApp recommends having at least 10% free space available in an aggregate hosting Exchange data for optimal storage performance.

- NetApp recommends having at least 10% free space available in a volume hosting Exchange data.

- NetApp recommends separating database and transaction logs from different servers into separate volumes to prevent a potential “busy” Snapshot copy problem. Utilizing separate volumes for each server reduces complexity, since there is no concern that Snapshot copy schedules of overlapping different servers might overlap.

- Place each database on a separate LUN in a separate volume.

- Place the transaction log files and the SnapInfo directory on the same LUN.
  - If a database’s transaction log files and the SnapInfo directory are placed on separate LUNs, place them both in the same volume.

- When creating LUNs, use volume mount points. There are a finite number of drive letters, and in a DAG each database path must be the same on every server that has a copy of that database.
  - The Exchange Server 2010 database and transaction log path must be unique for each database.

- Microsoft recommends approximately 20% free disk space in each LUN that has Exchange data.

- Enough space must be available in the aggregate for the autosize option to succeed. NetApp recommends planning for additional buffer space when using thin provisioning for Microsoft Exchange Server 2010 environments.

- NetApp strongly recommends setting the autodelete trigger to volume.

Sizing and Capacity Planning

- Consult a local NetApp Exchange expert to assist in accurately sizing Exchange Server 2010. Use the NetApp Sizing Tool for Exchange to size all Exchange server deployments utilizing NetApp storage.

Database Maintenance

- Enable background database maintenance on each database.

- Use larger databases. Microsoft supports up to 16TB databases with a default size of 2TB. Many customers, including NetApp IT, run Exchange Server 2010 with larger than 2TB databases on NetApp storage.

Data Protection

- Use the business requirements that are established by the Exchange stakeholders to help determine the number of Snapshot copies to keep online.

- Use NetApp Protection Manager to archive SnapManager for Exchange backup sets from primary to secondary storage.

- 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.

- When using SnapMirror for replicating Microsoft Exchange Server 2010 databases, NetApp recommends not using the “disrupt” option for commitment, because SnapMirror baseline Snapshot copies can be destroyed by autodelete even though they will always be the last Snapshot copies deleted. In many configurations, deleting the last SnapMirror Snapshot copy is not desirable because a new full baseline copy will be required to resume mirroring operations. If, for example, the source and destination are at different sites, recreating this baseline can be a time-consuming and costly process.

High Availability (Deployment)

- In a multisite scenario it is a best practice to 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 but also provides high availability in each site.
- Replication of a deduplicated volume is supported by using SnapMirror. However, NetApp does not recommend using deduplication with synchronous SnapMirror because that can add substantial overhead on the storage subsystem and introduce performance overhead to Exchange Server 2010 databases.
- Use SnapDrive for Windows to create FlexClone volumes. This automates the creation of the FlexClone volume and connects the LUNs within the clone to the test and development host.

High Availability (Storage Design)

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

High Availability (Volume Separation)

- Place active and passive copies of the database into separate volumes.
- Place flexible volumes for active and passive databases onto separate aggregates. If a single aggregate is lost, only the database copies on that aggregate are affected.

High Availability (Backup)

- Perform VSS backups on one of the passive nodes.

SNMP and Windows PowerShell

```
-- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -

-- Snap Autodelete Notice

-- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -

         snapAutoDelete                               NOTIFICATION-TYPE

          OBJECTS                                     {productTrapData, productSerialNum}

          STATUS                                      current

          DESCRIPTION                                 "Snapshot Autodeleted"

          ::= { netapp 0 656 }


-- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -

-- Volume Autogrow Notice
-- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -
         volumeAutogrow                              NOTIFICATION-TYPE
          OBJECTS                                    {productTrapData, productSerialNum}
          STATUS                                     current
          DESCRIPTION                                "Volume is Autogrown"
          ::= { netapp 0 666 }
-- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -- -
```

Windows PowerShell Monitoring Scripts

This sample script displays a list of Snapshot copies in all volumes that are older than 10 days.

```
$Now = get-date
get-navol | get-naSnapshot | where-object {$_.AccessTimeDT -le $Now.AddDays(-10)}
```

This sample script displays a list of Snapshot copies that are greater than 10.

```
get-navol | get-naSnapshot | select Name -skip 10
```

Rapid Reseed Script

```
cls

#This script requires that the SnapManager for Exchange PS.snapin and Exchange
powershell cmdlets are loaded prior in the session prior to use

#Modify your powershell profile to add them with the following:

#Add-PSSnapin NetApp.SnapManager.Exchange.PS.Admin

#$Session = New-PSSession -ConfigurationName Microsoft.Exchange -ConnectionUri
http://mst-exch-02.mst.local/PowerShell/ -Authentication Kerberos

#import-pssession $session |out-null


#remove the broken one

Write-Host "Suspend-MailboxDatabaseCopy" -foreground "magenta"

Suspend-mailboxdatabasecopy –Identity “Exch_01\MST-EXCH-03” -Confirm:$False


Write-Host "Remove-MailboxDatabaseCopy" -foreground "magenta"

Remove-mailboxdatabasecopy –Identity "Exch_01\MST-EXCH-03" -Confirm:$False


Write-Host "The database replica on MST-EXCH-03 is now removed" -foreground "magenta"

Get-MailboxDatabaseCopyStatus -Identity "Exch_01"
```

```
#‘add the new one, but don’t start seeding yet
Write-Host "We now re-add the database on MST-EXCH-03 but with seeding postponed" -
foreground "magenta"
Add-mailboxdatabasecopy –Identity “Exch_01” –Mailboxserver MST-EXCH-03 –
seedingPostponed -Confirm:$False

Write-Host "Get-MailboxDatabaseCopyStatus" -foreground "magenta"
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"

Write-Host ""
Write-Host "It's now Failed and Suspended, as expected. We'll now leverage NetApp
snapshots to quickly restore the LUNs and then perform an incremental reseed"
Write-Host "Press a Key to Continue" -foreground "magenta"
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

$backupArray = get-backup -server mst-exch-03 -storagegroup Exch_01 | sort Backup -
descending
$targetSet = $backupArray[0].Backup.ToLower()
Write-Host "Determine lastest backup set with SnapManager PowerShell..."
write-host Most Recent Backup Set : $targetSet

#‘restore the last snapshot of the logs and database with snapdrive
# restore_volume = restore : restore = lun_clone_split
Write-Host "Restore the Database LUN" -foreground "magenta"
sdcli snap restore -m mst-exch-03 -d E -s $targetSet
Write-Host "Restore the Transaction Log LUN" -foreground "magenta"
sdcli snap restore -m mst-exch-03 -d F -s $targetSet

Write-Host "Press a Key to Continue" -foreground "magenta"
$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

#‘resume replication to play it up to current and get healthy
Write-Host "Resume-MailboxDatabaseCopy" -foreground "magenta"
Resume-mailboxdatabasecopy –Identity “Exch_01\MST-EXCH-03” -Confirm:$False

Write-Host "Get-MailboxDatabaseCopyStatus" -foreground "magenta"
Get-MailboxDatabaseCopyStatus -Identity "Exch_01"
```

```
Write-Host "Wait 5 seconds for incremental log reseed, then press any key" -foreground
"magenta"

$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

Get-MailboxDatabaseCopyStatus -Identity "Exch_01"


Write-Host "Wait 5 seconds replay queue to drain, then press any key" -foreground
"magenta"

$x = $host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")

Get-MailboxDatabaseCopyStatus -Identity "Exch_01"


Write-Host "Rapid Reseed complete" -foreground "magenta"


#‘now it’s fixed
```

References

Exchange 2010 System Requirements: http://technet.microsoft.com/en-us/library/aa996719.aspx

Use Deduplication with Synchronous SnapMirror: http://media.netapp.com/documents/tr-3326.pdf

Storage Management Guide: http://now.netapp.com/NOW/knowledge/docs/ontap/rel81rc2/pdfs/ontap/smg.pdf

Volume Shadow Copy Service Overview: http://msdn.microsoft.com/en-us/library/aa384649(v=VS.85).aspx

Understanding the Mailbox Database Cache: http://technet.microsoft.com/en-us/library/ee832793.aspx

Understanding Mailbox Database and Log Capacity Factors: http://technet.microsoft.com/en-us/library/ee832796.aspx

Exchange 2010 Mailbox Server Role Requirements Calculator: http://gallery.technet.microsoft.com/v144-of-the-Exchange-2010-1912958d

Exchange Supports Up to 16TB for Databases: http://www.microsoft.com/exchange/en-us/licensing-exchange-server-email.aspx

Exchange 2010 System Requirements: http://technet.microsoft.com/en-us/library/aa996719.aspx

Shadow Redundancy: http://technet.microsoft.com/en-us/library/dd351027.aspx

Data ONTAP documentation: http://now.netapp.com/NOW/knowledge/docs/docs.cgi

Datacenter Switchovers: http://technet.microsoft.com/en-us/library/dd351049.aspx

Understanding Datacenter Activation Coordination Mode: http://technet.microsoft.com/en-us/library/dd979790.aspx

Recover an Exchange Server: http://technet.microsoft.com/en-us/library/dd876880.aspx

Understanding Edge Transport Server Cloned Configuration: http://technet.microsoft.com/en-us/library/aa998622.aspx

VSS Frequently Asked Questions for Exchange Server 2010: http://msdn.microsoft.com/en-us/library/aa579091%28v=exchg.140%29.aspx

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

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.

Go further, faster®

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Cache, FlexClone, FlexVol, OnCommand, RAID-DP, SnapDrive, SnapManager, SnapMirror, Snapshot, SnapVault, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, Windows, and Windows Server are registered trademarks and Hyper-V and Windows PowerShell are trademarks of Microsoft Corporation. ESX is a registered trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4033-0313

NetApp®

www.netapp.com