Exchange Server 2003 Site Recovery Using Volume SnapRestore®

Brad Garvey, Network Appliance, Inc.
Shannon Flynn, Network Appliance, Inc.

October 2006 | TR-3526

Table of Contents

1. Introduction

Many organizations have come to rely on Microsoft® Exchange Server to facilitate critical business e-mail communication processes, group scheduling, and calendaring on a 24x7 basis. System failures may result in unacceptable operational and financial losses.

There are several approaches to protect data and maintain data availability in the event of hardware, software, or even site failures. Backups provide a way to recover lost data from an archival medium such as disk or tape. Redundant hardware technologies also mitigate the damage caused by hardware failures. Data replication or mirroring provides a third way to ensure data availability and minimize downtime.

Because of the increasing importance of Microsoft Exchange Server in the enterprise, Exchange data protection, disaster recovery, and availability are of increasing concern. As the importance of Exchange within the organization increases, companies expect quick recovery times with little to no data loss.

Network Appliance offers the following comprehensive suite of hardware and software that enable organizations to keep pace with the increasing data availability demands of an ever-expanding Exchange environment, as well as scale to accommodate future needs while reducing cost and complexity.

Network Appliance™ SnapManager® for Microsoft Exchange (SME) software is available for Microsoft Exchange Server 2003. SME 3.2 has achieved a Certified for Windows® logo for Windows Server 2003 Standard and Enterprise editions. SME is also a Microsoft SimpleSAN designated and a Windows Server 2003 certified backup and recovery solution for Exchange Server. SME is tightly integrated with Microsoft Exchange, which allows for consistent online backups of your Exchange environment while leveraging NetApp Snapshot™ copy and SnapMirror® technologies. These technologies are critical in protecting Exchange data, allowing Exchange Server administrators to quickly back up and mirror Exchange data in a fast, efficient manner.

NetApp SnapMirror provides a fast and flexible enterprise solution for data replication over local area and wide area IP and Fibre Channel (FC) networks. If a disaster strikes at a primary site, replicated mission-critical data can be made quickly available at a remote site, ensuring uninterrupted operation and data availability.

When recovering data to a secondary site, lower RTO times can be achieved by using volume SnapRestore during the restore process. By using this method, lower recovery times can be achieved, and overhead that may be placed on the storage controller during the process is minimized.

1.1 Purpose and Scope

The purpose of this document is to provide a disaster recovery scenario for Exchange Server 2003 that is designed to achieve stringent RPO/RTO objectives. The scope of the solution provided is limited to the following.

1. An operational disaster recovery solution for migrating Exchange Virtual Server (EVS) onto a standby cluster in a secondary location. In this case, the new Windows Server 2003 Cluster will be in our standby recovery site. NetApp KB10542 and

Microsoft KB555603 were consulted to ensure that the recovery environment was correctly implemented and will recover the Exchange Servers according to best practices and Microsoft standards.

2. A fully implemented disaster recovery solution for Exchange Server 2003 based on NetApp solutions. This scenario using SnapMirror to replicate SME VSS initiated backup sets supplemented with periodic replication of the Exchange Server 2003 transaction logs that can be applied during a recovery operation.

SME restore options covered in this document are limited to an up-to-the-minute restore using volume SnapRestore to recovery database and transaction log volumes in the secondary location.

This version will not cover the following items:

- Detailed setup of Exchange Server recovery environment including the Windows Server 2003 Cluster. For detailed information on moving Exchange Virtual Server to a secondary cluster in a disaster recovery location, refer to NetApp KB10542 and Microsoft TechNet article on how to move Exchange Virtual Servers from a Production Exchange 2003 Cluster to a Standby Exchange 2003 Cluster.
- Server side recovery options such as usage of Recovery Storage Groups.
- Sync SnapMirror and Semi-Synchronous SnapMirror.
- Failback to the primary site.

1.2 Intended Audience

This document is intended for information technology professionals and storage professionals responsible for corporate Exchange messaging environments. For methods and procedures mentioned in this document it is assumed that the reader has a working knowledge of the following:

- Exchange storage architecture and administration
- Service level expertise of Microsoft Exchange recovery operations
- Working knowledge of NetApp solutions including the following:
    - DATA ONTAP®
    - SnapMirror
    - SnapManager for Exchange backup and restore procedures
    - SnapDrive® for Windows

2. Business Requirements

The first requirement in planning for High Availability or Business Continuity is identifying metrics to measure what needs to be achieved. Generally, an effective Business Continuity blue print starts with the description of metrics chosen, followed by strategy planning and implementation, and concludes with tests to ensure that what is being proposed actually works in practice.

To appropriately architect a disaster recovery solution for Microsoft Exchange Server, one should be familiar with the following terms.

Availability: It is important to remember that levels of availability are discrete in nature and fairly arbitrary. But, they generally apply to a wide range of real world scenarios.

High Availability (HA): High availability systems protect against application outages caused by loss of hardware, loss of software, or human error. HA solutions do not focus on events that cause the loss of an entire site. For example, a High Availability solution for Microsoft Exchange in a NetApp environment would include Microsoft Cluster Services to protect against physical server loss and SME to provide backup and recovery services in the event of data loss or corruption.

Disaster Recovery (DR): Protection against the loss of an application due to the loss of an entire site, be it a room, building, city, or region of the country. Commonly, Disaster Recovery combines high availability system design techniques with geographic separation between equipment at different sites. Protection requires both recovery of lost data and restoration of application service. Disaster Recovery is the highest level of availability and provides the most expensive level of protection.

The primary difference between HA and DR is: Two complete copies of data at a single site provide high availability, while two complete copies of data at two different sites provide disaster recovery protection. Geographic separation is required for DR because a disaster assumes the loss of an entire site, not just a piece of equipment. SnapMirror provides data replication and synchronization services to replicate SME Snapshot copies to an alternate site to provide recovery in the event of a site failure.

Service Level Agreement (SLA): An SLA is a contract between IT service providers and users for application level performance. For Disaster Recovery SLAs, both RPO (protection) and RTO (Recovery time) are typically defined. The SLA for DR specifies “how long do I have to wait to start working and how much work do I have to re-do because it’s been lost.”

Recovery Time Objective (RTO): RTO is the maximum elapsed time to recover an application from outage and ensure that recovered technology components are functional enough to resume business transactions. In other words, RTO is an indicator of how much downtime your business can endure. Equipment failover and restart, application failover or restart, and recovery of data by replaying transaction logs all contribute to the time required before the application is restored. For DR planning, a RTO is defined considering recovery times for equipment, software and additional recovery steps. Frequent testing and tuning of the plans and recovery procedures ensure the proposed RTO can be met at the time of disaster.

Recovery Point Objective (RPO): RPO defines the acceptable amount of data loss in the event of a disaster. Because DR planning requires replication of data across two or more sites, some amount of data could very well be in transition when an event occurs. The RPO defines how much data may possibly be lost. For example, if daily tape backup is the only recovery mechanism, the RPO will be the amount of data changed since the last backup and could be as much as 24 hours. If synchronous mirroring is the data-protection mechanism, the RPO could be zero because no completed transactions are ever lost. For transaction-oriented systems like Exchange, database files may have an RPO of several hours, while the transaction log for the database has an RPO close to zero. With such a design, it is possible to recover the database records by replaying the transaction log from the last checkpoint. This approach can lengthen the RTO but helps reduce the cost of the entire solution.

The following diagram outlines RPO, RTO and timelines associated with various stages of Business continuity.

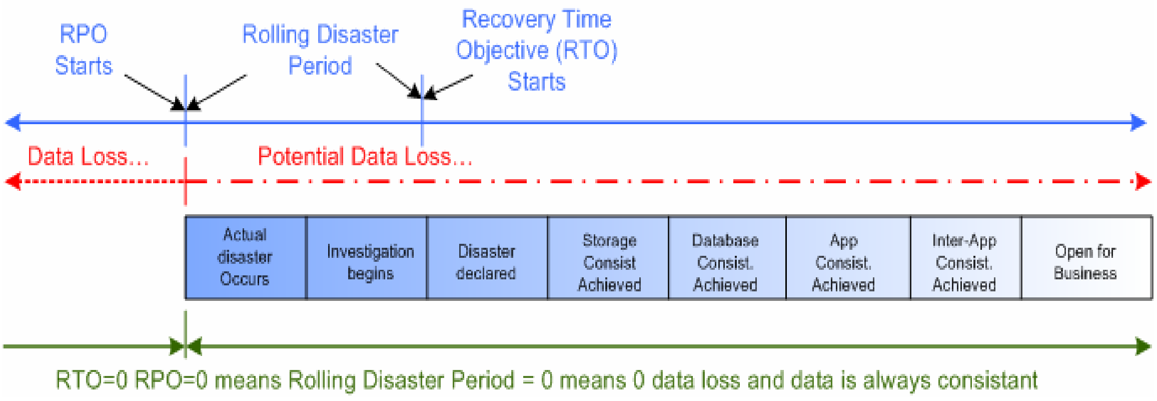

Figure 1) RPO/RTO Timeline

3. Meeting Stringent RPO/RTO Objectives

When architecting an Exchange Server disaster recovery solution based on NetApp products, there are three main components which aid in meeting desired RPO/RTO times.

- **SnapManager for Exchange (SME):** SME ensures that consistent backups are created of the Exchange data. These backups can be used to easily and quickly restore your Exchange environment, and can be quickly replicated to a DR site using SnapMirror.
- **SnapMirror Replication:** SnapMirror provides a means by which consistent Exchange Snapshot copies can be mirrored to a disaster recovery site. To help achieve an aggressive RTO, Rolling Snapshot copies of the Exchange transaction logs can be configured to bridge the gap between full SME backups.
- **Restore options:** There are two core methods that can be used in an Exchange DR situation: LUN-Clone-Split and volume SnapRestore.

3.1 Replication of Exchange Data to the DR Site

To meet stringent RPO objectives, it is critical to have your Exchange data required for a restore operation replicated to the recovery site. Increasing the frequency of your SME backups can help you achieve a lower RTO but, for a variety of reasons, is not the optimal way to achieve this objective. This may cause an increase in latency times and a higher risk of overlapping SME operations.

As an alternative to increasing the frequency of SnapManager backups, the rolling Snapshot copy capability of SnapDrive will provide frequent replication of the volumes that contain the Exchange transaction log LUNs. These additional rolling Snapshot copies complement the regular full SME backups that are replicated to the recovery site. With the most current transaction logs replicated, a restore from the destination volume replays those transaction logs and results in an up-to-the-minute restore from the time of the last transaction log update. This method enables you to keep your destination volume more current without running the risk of overlapping your SnapManager operations. For more information on rolling Snapshot copies, please refer to the SnapDrive Installation and Administration Guide, Chapter 11, Page 205.

There are three major advantages to using SnapDrive rolling Snapshot copies to augment your SnapMirror replication schedule for SnapManager:

- **Fewer changes are made to the Exchange data between replications**
  When you use SnapDrive to initiate frequent updates of the transaction logs, a restore from the destination volume replays those transaction logs and results in an up-to-the-minute restore for the time of the last transaction log update.
- **Fewer SnapManager backups are required**
  It is not recommended to create SnapManager backups every few minutes. This can result in overlapping SnapManager operations and more difficulty managing the large number of resulting backups.
- **Fewer Snapshot copies are retained**
  A maximum of two rolling Snapshot copies are retained at any time. This provides easier management of rolling Snapshot copies and also saves disk space on the storage appliances.

3.2 Restore Options

There are two methods that can be used to do the restore of Exchange databases in a recovery site. SnapManager for Exchange by default will perform a restore using the LUN Clone Split restore method. Alternatively a Volume Snap Restore can be done manually to minimize the affect of additional overhead placed on the storage array during the process and help meet more stringent RPO and RTO times.

LUN Clone Split Restore (LCSR)

Once your data is replicated to your Disaster Recovery site, both NetApp SnapDrive and SME will perform a recovery of the replicated LUNs. The default process for this recovery is called a LUN Clone Split operation. Starting with Data ONTAP 7.1, when a LUN restore operation is initiated, Data ONTAP will create a LUN clone in place of the original LUN. The LUN clone shares space with the LUN in the backing Snapshot copy. The clone does not require additional disk space until new writes occur. Data ONTAP then splits the clone from the backing Snapshot copy and data is copied from the Snapshot copy to the clone. After the split operation is complete, both the backing Snapshot copy and the clone occupy their own space. This improvement means the LUN restore operation is now nearly instantaneous and the requesting host and applications can immediately start using the restored LUN to service reads and writes from the Exchange Server.

After the LUN Clone Split is initiated, it will continue to run in the background to complete the restore operation. During this time, the storage array will experience an increased CPU load which may increase latency times. Also, while the process is running in the background, you will not be able to perform an SME backup of your newly restored Exchange environment.

The LUN Clone Split method of doing a restore should be used when recovering a single LUN or mailbox database. Restore operations such as this usually occur in a single site and fit under the umbrella of traditional backup and recovery, not disaster recovery.

At the time this paper was written, deployments containing more than six LUNs and average LUN sizes of 100GB or greater, the SME recovery and resulting LUN Clone Split operation may cause high utilization on the storage controller which in some cases has been known to result in unacceptable response times. This can adversely affect other production workloads that may be on the same storage system during the LUN Clone split background operation.

When failing over to a secondary site and restoring the whole Exchange cluster, an alternative method called volume SnapRestore method can be used as a means of decreasing the time required to complete the restore while minimizing overhead on the storage controller.

Volume SnapRestore (VSR)

In contrast to a LUN Clone split, which restores a single LUN from a Snapshot copy, VSR works by restoring the entire volume to a given point in time. This process is very efficient and quick, with virtually no performance impact on the storage device. Since this type of restore occurs at the volume level, careful consideration for LUN layout for Exchange databases and transaction logs is critical. Improper LUN layout may result in loss of data.

To take advantage of VSR restores for Exchange DR, the following guidelines should be followed:

- LUNs containing Exchange database files should never share the same volume as LUNs containing transaction logs or the SnapInfo directories.
- Storage groups that share the same volume must be backed up in the same backup operation.
- Exchange Virtual Server data files must reside on separate volumes from other Exchange Virtual Server data files.

Currently, volume SnapRestore is not integrated into SME. Manual steps using a storage console session will need to be performed in order to take advantage of the benefits of a volume-based SnapRestore. The details of this process will be covered in the test scenario described below.

4. Requirements and Solutions Overview

The goal of the test environment was to architect a solution that would meet a 30 minute RPO and 1 hour RTO time. The following test case was designed using two different sites, a primary site for the production and a secondary site for recovery. An Exchange Server Cluster hosting 1500 users connected to a FAS3050 was located in the primary location. The secondary DR site was set up with a standby cluster connected to a FAS980. Scheduled SME backups were configured to replicate to the disaster recovery site every four hours. To provide up-to-the-minute restore capabilities, rolling Snapshot copies were also used to replicate incremental changes for the transaction log volumes only every thirty minutes.

4.1 Deployed Architecture

The following diagram shows the basic architecture used for the Exchange Server disaster recovery scenario used in this document. This environment was used to test and validate the Volume SnapRestore process and the use of rolling Snapshot copies to replicate incremental changes of the transaction log volume(s). For additional information on setting up the recovery environment or moving the EVS to a recovery cluster, please reference both NetApp KB10542 and the Microsoft TechNet article How to Move Exchange Virtual Servers from a Production Exchange 2003 Cluster to a Standby Exchange 2003 Cluster for exact details and steps required.

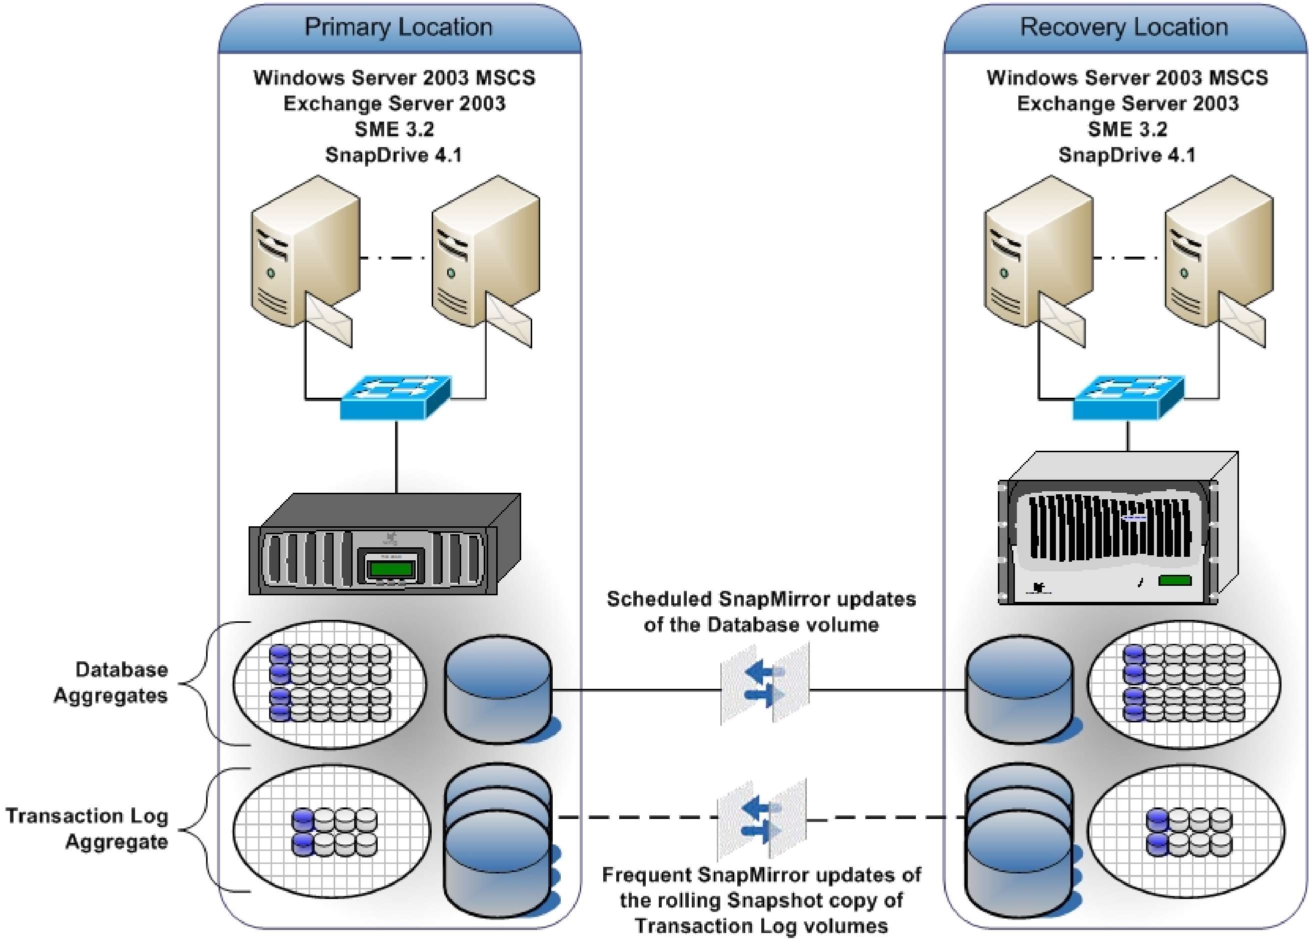

Figure 2) Disaster Recovery layout

4.2 Hardware

The primary site consisted of the following storage configuration.

Storage infrastructure:

- 2 FAS3050 Storage Controllers acting as an Active/Active Cluster
- Storage OS: Data ONTAP 7.2
- 6 DS14 MK2 Shelves (3 per controller) with 144GB 15K FC Drives

Data was recovered to a secondary site with the following storage configuration:

- 2 FAS980 Storage Controllers acting as an Active/Active Cluster
- Storage OS: Data ONTAP 7.2
- 6 DS14 MK2 Shelves (3 per controller) with 144GB 15K FC Drives

4.3 Data Layout

The following volume layout was used in the test environment.

	Used Capacity	Total Capacity	LUNs
DBVOL	379 GB	1920GB	D: - SG1 Mailbox store
			E: - SG2 Mailbox Store
			F: - SG3 Mailbox Store
LOGVOL1	20GB	32GB	G: - SG1 Logs
LOGVOL2	20GB	32GB	H: - SG2 Logs
LOGVOL3	20GB	32GB	I: - SG3 Logs

Table 1) Data file layout

Best Practice

When laying out your Exchange data onto the storage appliance, take into careful consideration the factors that were outlined in the above sections. Things like rate of change, bandwidth available for replication, and RPO/RTO for different storage groups all affect storage layout. For example, if you have executive users that require a higher RPO/RTO than normal users, it's best practice to place those users into their own storage group, place that storage group onto its own set of LUNs, and place those LUNs into their own dedicated volumes that can be replicated to the DR site more frequently.

4.4 Exchange Load Simulation

Microsoft LoadSim was used to create mailboxes for 1500 users and simulate Exchange MAPI traffic during the test procedure. For more information on Microsoft LoadSim, please see the Exchange Server 2003 Performance Tools page on the Microsoft TechNet website. The following configuration was used in the LoadSim configuration:

Exchange Server 2003 SP2

- 1500 mailboxes
- 3 Storage Groups
- 1 Mailbox database per storage group

Mailbox Characteristics

- 500 Users/Mailbox database
- Average mailbox size of 100MB

After the initialization process completed, LoadSim was run for 12 hours to simulate traffic. During the initial run period SME was configured to create Snapshot copies every 4 hours in order to provide data needed to calculate the rate of change.

4.5 SME Backup Schedules and Rate of Change

When determining RPO/RTO times for site, the following items must be taken into consideration:

- Rate of change of the database volumes
- Rate of change of the transaction log volumes
- SnapMirror performance
- Total data sizes

4.6 Backup Schedule

The following table shows the schedule that was set up to achieve a 30 minute RPO time. Full SME backups were scheduled every 4 hours with rolling Snapshot copies of the transaction log volume scheduled every 30 minutes to achieve the desired RPO time of 30 minutes.

NOTE: When setting up a schedule, rolling Snapshot copy schedules must not conflict with SME full backups. A rolling Snapshot copy is not necessary when a full SME backup occurs. When a full SME backup is created, both the Exchange database files and transaction log files are backed up and replicated to the DR site. In this schedule, the rolling Snapshot copies started a half an hour after the full backup.

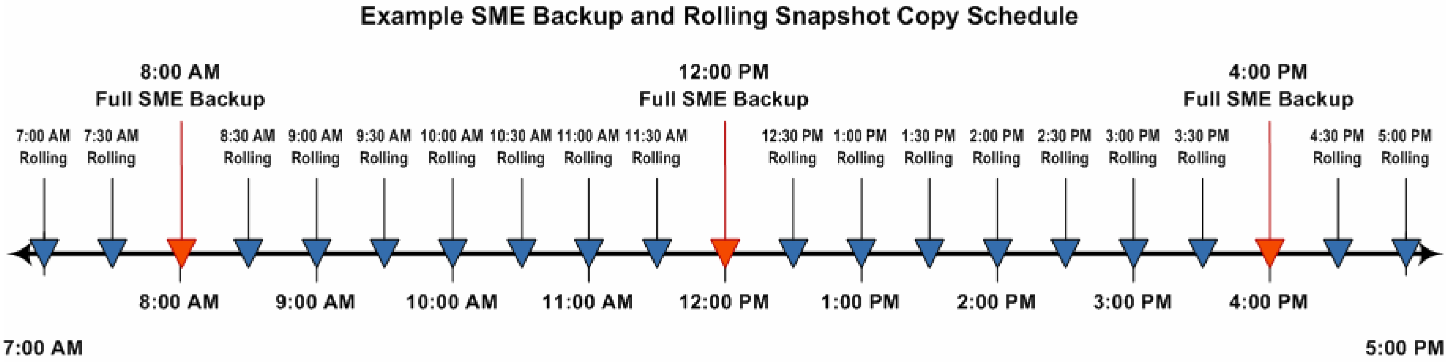

Figure 3) Example schedule for full database backups and rolling Snapshot copies

When setting up SME backup and rolling Snapshot copy schedules, it is important to take the following into consideration:

- The first Snapshot copy replication must not begin until the initial backup operation is complete.
- Subsequent Snapshot copy replications must not begin before the previous replication is complete.

NOTE: Depending on available bandwidth and data sizes it may be necessary to adjust the replication schedules to complete within a reasonable amount of time.

To accurately calculate the rate of change of the database volume, SME backups should be scheduled and monitored for at least a 12 hour period during peak usage period.

During this period, the `snap delta` command can be used to get the size of each Snapshot copy. The sizes of each Snapshot copy will aid in estimating how much data is changing throughout the twelve hour period and how much data needs to be replicated during the SnapMirror updates.

The `snap delta` command shown in the figure below was used to display the rate of change between all Snapshot copies on the database volume over the course of the 12 hour LoadSim.

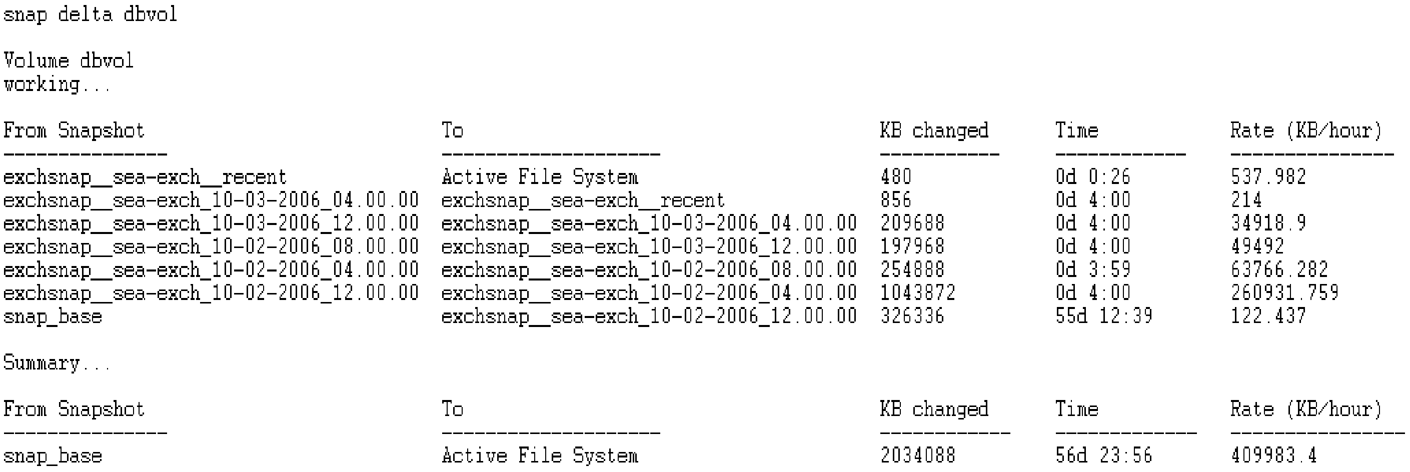

```
snap delta dbvol

Volume dbvol
working...

From Snapshot                         To                                    KB changed     Time           Rate (KB/hour)
---------------                       --------------------                  -----------    -----------    ---------------
exchsnap__sea-exch__recent            Active File System                    480            0d 0:26        537.982
exchsnap__sea-exch_10-03-2006_04.00.00 exchsnap__sea-exch__recent           856            0d 4:00        214
exchsnap__sea-exch_10-03-2006_12.00.00 exchsnap__sea-exch_10-03-2006_04.00.00 209688       0d 4:00        34918.9
exchsnap__sea-exch_10-02-2006_08.00.00 exchsnap__sea-exch_10-03-2006_12.00.00 197968       0d 4:00        49492
exchsnap__sea-exch_10-02-2006_04.00.00 exchsnap__sea-exch_10-02-2006_08.00.00 254888       0d 3:59        63766.282
exchsnap__sea-exch_10-02-2006_12.00.00 exchsnap__sea-exch_10-02-2006_04.00.00 1043872      0d 4:00        260931.759
snap_base                             exchsnap__sea-exch_10-02-2006_12.00.00 326336        55d 12:39      122.437

Summary...

From Snapshot                         To                                    KB changed     Time           Rate (KB/hour)
---------------                       --------------------                  ------------   ------------   ----------------
snap_base                             Active File System                    2034088        56d 23:56      409983.4
```

Figure 4) snap `delta` console output

The first row of the snap delta output displays the rate of change between the most recent Snapshot copy and the active file system. The following rows provide the rate of change between successive Snapshot copies. Each row displays the names of the two Snapshot copies that are compared; the amount of data that has changed between them, the time elapsed between the two Snapshot copies, and how fast the data changed between the two Snapshot copies.

If you do not specify any Snapshot copies when you enter the `snap delta` command, the output also displays a table that summarizes the rate of change for the volume between the oldest Snapshot copy and the active file system.

Using the snap delta command to calculate the size of each Snapshot copy, the amount of data generated during a 12 hour period on the database volume was calculated to be 2.03 GB.

From Snapshot Copy	To Snapshot Copy	Size (KB)
exchsnap__sea-exch__recent	Active File System	480
exchsnap__sea-exch_10-03-2006_04.00.00	exchsnap__sea-exch__recent	856
exchsnap__sea-exch_10-03-2006_12.00.00	exchsnap__sea-exch_10-03-2006_04.00.00	209688
exchsnap__sea-exch_10-02-2006_08.00.00	exchsnap__sea-exch_10-03-2006_12.00.00	197968
exchsnap__sea-exch_10-02-2006_04.00.00	exchsnap__sea-exch_10-02-2006_08.00.00	254888
exchsnap__sea-exch_10-02-2006_12.00.00	exchsnap__sea-exch_10-02-2006_04.00.00	1043872
snap_base	exchsnap__sea-exch_10-02-2006_12.00.00	326336

Table 2) Snapshot copy sizes from database volume.

4.7 SnapMirror: Rolling Snapshot Copies and Rate of Change

To calculate the amount of data generated by transaction logs, the log volumes were analyzed for a 4 hour interval during the LoadSim run. The number of logs generated in each storage group per thirty minute interval was collected and the following calculations were made. During the 30 minute timeframe, the average total size of transaction logs generated was 741MB. To successfully meet the RPO time of 30 minutes, each rolling Snapshot copy must be able to replicate approximately 741MB of data within a 30 minute window. If this replication time is not met and rolling Snapshot replications overlap, the RPO time could increase.

Time	Number of logs generated for SG1	Number of logs generated for SG2	Number of logs generated for SG3	Total number of logs generated	Number of logs per user	Total Size of logs in MB
8:30 am	47	48	47	142	0.09	710
9:00 am	48	47	47	142	0.09	710
9:30 am	52	52	52	156	0.10	782
10:00 am	52	53	52	157	0.10	783
10:30 am	47	51	50	148	0.10	740
11:00 am	48	51	49	148	0.09	740
11:30 am	49	49	49	47	0.09	735
12:00 pm	50	49	48	47	0.10	735

Table 3) Transaction log metrics taken at thirty minute intervals.

4.8 Setting Up the SnapMirror Relationship

The following steps describe how to set up a SnapMirror relationship for the test environment.

1. Create the SnapMirror destination volumes at DR site.
   - These volumes have to be equal to or greater than the size of the source volumes that will be mirrored.

2. From the command console on the source storage controller, use the `options snapmirror.access` command to specify the host names of storage controllers that are allowed to copy data directly from the source. For example:

   ```
   options snapmirror.access host=<dest_storage>
   ```

3. Restrict those volumes on the destination storage controller to allow SnapMirror to access them by using the `vol restrict` command:

   ```
   vol restrict <vol_name>
   ```

   - Repeat this command for each volume that will serve as a SnapMirror destination.

4. Initialize SnapMirror process.
   - From the destination storage controller command console, use the SnapMirror initialize command to create an initial complete (baseline) copy of the source on the destination and start the mirroring process.

     ```
     snapmirror initialize –S
     <src_storage>:<src_vol><dest_storage>:<dest_vol>
     ```

To check that a destination volume has been initialized, you can use the `snapmirror status` command. If you specify no options or arguments, the `snapmirror status` command shows the status of the volumes on the storage controller, as shown in the example below.

```
SEA3050-2>snapmirror status

Snapmirror is on.

Source             Destination          State         Lag        Status
SEA3050-2:dbvol    fas980-svl21:dbvol   Snapmirrored  00:56:58   Idle
SEA3050-2:logvol1 fas980-svl21:logvol1 Source        Transferring (126
MB done)
```

4.9 Setting Up SME Scheduled Full Backups and Rolling Snapshot Copies

The following steps outline how to set up SME to create regularly scheduled full database backups and how to configure SnapDrive to create rolling Snapshot copies.

SME Scheduled Full Backups

Using the SME backup window, you can schedule backups using the Windows Task Scheduler. This is the most reliable way for regular backups to occur given a preset schedule. The following steps were taken for our environment. We have agreed that a full SME backup every 4 hours will satisfy our DR requirements.

1. Open SME and click on the Backup and Verification tab.
2. Select the storage groups that you want to have backed up and replicated.
3. Ensure that the Update SnapMirror after operation checkbox is selected, then click on the Schedule button.
4. SME will then create a new Windows Task for this backup operation.
    a. Set the appropriate schedule based on the data recovery model for your environment.
    b. In our environment, a full SME backup was scheduled every 4 hours.

Rolling Snapshot Copies

The following steps outline how to create a scheduled rolling Snapshot copy operation for SnapDrive using the Windows Task Scheduler.

1. Create a batch file (a file with a .bat extension) containing the following command on the Windows host on which you are scheduling the rolling Snapshot copies:

    ```
    sdcli snap update_mirror -D DriveLetterList
    ```

    - *DriveLetterList* is a list of space-separated drive letters that contain the transaction logs.

    Example:

    ```
    sdcli snap update_mirror -D g h i
    ```

2. Select Start Menu > Settings > Control Panel > Scheduled Tasks.
3. Double-click Add Scheduled Task.
4. In the Scheduled Task Wizard, click Browse, and navigate to the folder where the batch (.bat) file you created in Step 1 is located and select it.
5. After the following panel appears, select from the list of frequencies, taking into consideration your RPO objectives, then click Next.
    - For an RPO of 30 minutes, the schedule is set to every 30 minutes.
6. Next, enter a start time and complete the detailed frequency parameters. The option details displayed on this panel vary depending on the Snapshot copy frequency you picked in the previous panel.
7. Then, type the user name and password for the scheduled task, then click Next.
    - It is best to use the same SME user account and password for this task as it will have the correct permissions to execute the task.

Best Practice
When scheduling the SME backups and the rolling Snapshot copy backup jobs, stagger the run times so the two operations do not start at the same time. SME and SnapDrive do not support concurrent operations. If two operations do occur at the same time, the operation that is processed first will succeed; the other operation will fail.

5. Restoring the Exchange Environment in the Event of a Disaster

When using SME restores to recover multiple LUNs, the increased load placed on the storage appliance may at certain data size thresholds negatively affect RTO times and overall performance. Using Volume SnapRestore can help offset this load while decreasing RTO times.

Best Practice
During the recovery process the following things should be considered and accounted for when calculating RTO times.

- Added CPU overhead and increased read/write latency on the storage controller
- Added overhead on the server
- Size and Number of volumes and LUNs being restored
- Total size of the volumes being restored
- Amount of transaction logs to be replayed by the restore process
- Number of active nodes in the cluster group

5.1 Volume SnapRestore Method

The following steps were used to restore the Exchange Server at the secondary location using the Volume SnapRestore method. Prior to completing the following steps, ensure the Exchange Virtual Server is offline and any scheduled SME backup or rolling Snapshot copy jobs are disabled. Then follow the steps listed below:

1. Issue a SnapMirror break of the Exchange Virtual Server’s DB, Log, MTA/SMTP, and SnapInfo volumes:

   ```
   snapmirror break <vol_name>
   ```

   Examples:

   ```
   snapmirror break dbvol
   snapmirror break logvol1
   ```

   - Once this is complete, ensure the relationships have been removed by issuing the `snapmirror status command`. The status of the volumes should be `“broken-off”.`

2. Perform a Volume SnapRestore of the Exchange Virtual Server’s DB volume(s) to the most recently completed SME Snapshot copies. (i.e. __recent):

   ```
   snap restore –V  -s <most recent snap> <dest_vol_name>
   ```

   Examples:

   ```
   snap restore –V fas980-svl21(0101164971)_dbvol.3 dbvol
   snap restore –V exchsnap__sea-exch__recent db_vol
   ```

3. Perform a Volume SnapRestore of the Exchange Virtual Server's transaction log volume(s) to the most recently completed SnapDrive initiated Snapshot copy.
    - Typically, this will be the rolling Snapshot copy. These Snapshot copies are named with a "@snapmir@" prefix.
    - If a site failover is initiated shortly after an SME backup but before the next rolling Snapshot copy is initiated, there may be a case in which the SME eloginfo_timestamp and/or _recent Snapshot copy is more recent than the rolling Snapshot copy.
    - Ensure that the most current Snapshot copy initiated by SnapDrive is used to perform the volume SnapRestore operation.

    Examples:

    ```
    snap restore –V @snapmir@{ccc6c962-9c84-4875-a7fa-
    2ae66da4fe8b} logvol1
    snap restore –V eloginfo__sea-exch__recent logvol1
    ```

4. Ensure all recovered volumes are online.
5. Clear all LUN mappings that may have been carried over from the production site.
    - To do this, you can manually remove the mapping via the console or through FilerView®.
6. Turn SnapMirror off on the storage appliance. If this is not possible due to other SnapMirror relationships, please remove the Exchange destination volumes from the snapmirror.conf file either thru the command line or through FilerView.
7. Bring all of the Exchange Virtual Server's LUNs online.
8. From the DR Node that will host the Exchange Virtual Server, connect to all of the required LUNs using SnapDrive.
    - An SDCLI LUN connect script can be used to automate this process.

    ```
    SDCLI.exe disk connect -p <path_to_LUN> -d <drive_letter> -I
    <iqn_for_both_MSCS_nodes> -dtype shared -e <res_group_name>
    ```

    NOTE: LUNs must be mapped to the same drive letters used on the production site. If this is not done the SnapManager restore operation will fail.

    Example:

    ```
    sdcli.exe disk connect –p /vol/dbvol/dblun1.lun –d d: –I
    viaRPC.iqn.1991-05.com.microsoft:sea-
    exch1.exchdf.iop.eng.netapp.com viaRPC.iqn.1991-
    05.com.microsoft:sea-exch2.exchdf.iop.eng.netapp.com –d shared
    –e exch2k3
    ```

9. From the DR Node that will host the Exchange Virtual Server, use Cluster Administrator and verify that the LUNs appear in the correct Resource Groups.
10. Using Active Directory Users and Computers, reset the Exchange Virtual Server account that will be recovered.
11. Create the Exchange Virtual Server resources on the DR node.
12. Bring the newly created Exchange Virtual Server online.
    - At this point the EVS should come online and all mailbox databases should mount.
13. Start SME and rerun the Configuration Wizard.
    - Ensure that the SnapInfo location matches the production site.
    - This is especially important if a dedicated SnapInfo LUN is used.

14. Restart SME using the –restorearchive switch.
    - Select the first Storage Group and appropriate location for SnapInfo from the drop-down dialog box.
    - Recover with the “Up-to-the-Minute” checkbox selected and “Mount Databases automatically after the restore” checkbox cleared.
    - Using this switch and selecting the “Up-to-the-minute” checkbox, SME will restore the Exchange database for the storage from the last full backup and replay any transaction logs that were replicated since that backup was created.
15. Repeat previous step for all Storage Groups.
16. Using the Exchange System Manager, mount all of the Exchange storage groups.
17. Perform an SME backup of the newly recovered Exchange environment.

5.2 Restore Time

During the restore tests, the following metrics were observed to track the total time to recover the Exchange Server storage on the DR site:

Initial Steps	Time to Completion
Break the SnapMirror relationship for all volumes	30 seconds
Perform SnapRestore of all volumes	30 seconds
Clear LUN mappings that may have been carried over from the production site	1 minute
Mount all six Exchange LUNs on the standby cluster	5 minutes
Total Time	**7 minutes**

Table 4) Initial DR steps for DR

Storage Group	Number of Transaction Logs To Replay	Size of Transaction Log Directory	Log Replay Time	Total SME Restore Time
SG1	825	4.01GB	10:30 minutes	14:20 minutes
SG2	809	3.93GB	11:25 minutes	15:30 minutes
SG3	811	3.94GB	11:10 minutes	15:10 minutes
Totals	**2445**	**11.88GB**	**33:05 minutes**	**45:00 minutes**

Table 5) Volume SnapRestore method timing.
Note: The Log Replay Time is included in the SME Restore Time.

Total recovery time for the scenario tested was 52 minutes. Note that the recovery time is largely affected by the number of logs to be replayed and LUN connect times. In our recovery scenario, there were a total of 11.88 GB of transaction logs to be replayed.

The number of active nodes in a cluster will impact LUN connect times. More active nodes in the cluster group will result in increased LUN connect times. When calculating RTO objectives careful consideration should be given to the amount of active nodes in a cluster group.

When planning RPO and RTO times, it is important to know the approximate number of logs that would need to be replayed and the rate at which logs replay. In our recovery environment, each storage group took approximately 11 minutes to replay its respective transaction logs into the database. The total restore time for all storage groups was 45 minutes. From this data we can calculate that logs for our particular test environment are replayed at the rate of approximately 80 transaction logs per minute, or 400 MB per minute.

To achieve a lower recovery time, more frequent SME backups would need to be created and the frequency of the rolling Snapshot copies increased, thus reducing the number of logs to be replayed. In our recovery environment, the bandwidth provided by the fractional T-1 link did not allow for more frequent replication of database backups and rolling Snapshot copies, which resulted in a larger number of logs to be replayed. Nonetheless, we were able to achieve an RPO of 30 minutes and an RTO of less than one hour.

6. Summary

Findings in this paper are based on one possible configuration. Other environment specific variables may affect recovery times and performance. However, the procedures and best practices still apply when planning an Exchange disaster recovery blueprint.

Planning for an Exchange Disaster Recovery process is a key part of protecting your valuable Exchange data. Exchange is a critical business application that can cripple operational productivity if it becomes unavailable. NetApp has proven data protection and disaster recovery tools for Exchange. SnapManager for Exchange backup and restore capabilities combined with SnapDrive and SnapMirror technologies provide a solid and robust solution for protecting and recovering your Exchange data while meeting stringent RPO and RTO objectives.

Appendix A – Best Practices

When planning and sizing a disaster recovery solution for Exchange Server environments the following best practices should be considered.

- Determine which data to be replicated: As the data needing to be replicated increases the time required for SnapMirror to successfully mirror that data also increases. This can impact available bandwidth and RPO/RTO objectives.

- Plan volume layout to help archive RPO/RTO objectives: Separate users or business units requiring faster recovery times and minimal data loss into separate storage groups and volumes. This adds flexibility to SnapMirror schedules and recovery objectives.

- Properly plan and size SnapMirror replication and schedules: Determine rate of change for each volume to ensure that the amount of data to be transferred by the SnapMirror process fits within desired incremental update times.

- When determining RTO objectives, ensure that the following processes are taken into account.
    - Time required for LUN recovery process.
    - Amount of time required to complete the LUN Clone Split process.
    - Number of transaction logs to be replayed during an up to the minute recovery. Increase the frequency of rolling Snapshot copies to decrease the amount of data to be replicated as well as the amount of transaction logs to be replayed during the restore process.

- When laying out your Exchange data onto the storage appliance, take into careful consideration the factors that were outlined in the above sections. Things like rate of change, bandwidth available for replication, and RPO/RTO for different storage groups all affect storage layout. For example, if you have executive users that require a higher RPO/RTO than normal users, it is best practice to put those executive users into their own storage group, place that storage group onto its own set of LUNs, and place those LUNs into their own dedicated volumes that can be replicated to the DR site more frequently.

- When scheduling the SME backups and the rolling Snapshot copy backup jobs, stagger the run time so the two operations do not start at the same time. SME and SnapDrive do not support concurrent operations. If two operations do occur at the same time, the operation that is processed first will succeed; the other operation will fail.

Appendix B – LoadSim Configuration

LoadSim was used to create, initialize, and simulate traffic for 1500 Mailboxes. To support this configuration, three storage groups each containing 500 users were created. During the testing period LoadSim was configured with an 8 hour daytime, 0 hour nighttime, and 8 hour duration.

Mailbox Initialization

Mailbox Setup	
Number of messages in Inbox	250.00
Number of messages in Deleted Items	1.00
Number of new folders	5
Number of messages per new folder	100.00
Number of Smart Folders	3
Number of rules in Inbox	3
Calendar Setup	
Number of Appointments	25
Contact Setup	
Number of Contacts	64

Tasks

Frequency	Task Name
10.00	Send Mail
12.00	Process Inbox
15.00	Browse Mail
1.00	Free/Busy
.10	Request Meetings
.20	Make Appointments
3.00	Browse Calendar
3.00	Smart Folders
3.00	Rules

Test Report

Approximate message traffic, per user, per day:

Total received:	141.34
Reply:	7.22
Reply all:	5.13
Forward:	7.22
Total submitted:	29.58

Average recipients per message (for all msgs): 4.78

Approximate receipts requested, per user, per day:

Read Receipts:	0.0
Delivery Receipts:	0.0

Attention: This report does not include the mail flow caused by rules, DDLs, and meeting requests/receipts.