Microsoft® Exchange 2007 Disaster Recovery Model Using NetApp Solutions

Amarnath Rampratap, Network Appliance, Inc.

July 2007 | TR-3584

Table of Contents

Introduction 3

Purpose and Scope 3

Intended Audience 3

Business Continuity and High-Availability Planning 4

Disaster Recovery Model for Microsoft Exchange Server 2007 4

Technology Components 4

Disaster Recovery Model: Objective 5

Architecture 5

Disaster Recovery Scenario: Exchange 2007 Case Study 6

Business Case One: Five-Minute RPO/30-Minute RTO Target 6

Implementation details 7

Setting Up SnapMirror Relationship 9

Setting Up SME Scheduled Backups and Rolling Snapshot Copies 10

Rolling Snapshot Copies 10

Replication of Hub Transport Boot LUN 10

Rate of Change 11

Recovery Methodology in the Event of a Disaster 12

Restore Time 13

Business Case Two: Zero-Minute RPO/30-Minute RTO Target 14

Implementation details 14

Replication of Hub Transport Server Using ReplicatorX 15

Recovery Methodology in the Event of a Disaster 17

Restore Time 18

Summary 18

Appendix A: Best Practices 19

Appendix B: Reference 20

Introduction

Today's business requirements, such as high availability, business continuity, and disaster recovery, are more extensive than ever before. As Microsoft® Exchange has become a mission-critical application in recent years, Microsoft Exchange Server downtime costs companies millions of dollars per year. IT organizations work hard to eliminate or lessen the impact of both planned and unplanned downtime through implementation of high-availability strategies and disaster recovery solutions.

Microsoft Exchange Server 2007 has matured to a new level, providing operational efficiency, delivering a messaging system that addresses performance, security, availability and cost. Microsoft Exchange Server 2007, with its new server roles and the new command-line interface, provides improved manageability and simplify maintenance. The built-in features such as local continuous replication (LCR) and cluster continuous replication (CCR) provide high availability, quick recovery, and resiliency for Exchange 2007 mailbox servers. NetApp along with Microsoft has extended its reach to increase productivity and keep information close to hand, flexible enough to meet your organization's administrative model.

This technical report delivers an overview of a disaster recovery model for Microsoft Exchange 2007 CCR using NetApp solutions.

Purpose and Scope

The purpose of this technical report is to provide a disaster recovery scenario for Microsoft Exchange Server 2007 CCR setup using a NetApp solution that is designed to achieve multiple levels of RPO and RTO objectives, discussed later in this document. The scope of the solution provided is limited to the following:

1. An operational disaster recovery solution for migrating Clustered Mailbox Server (CMS) onto a standby CCR cluster in a secondary location. In this case the new CCR cluster will be in our recovery site in the DR site.
2. A fully implemented disaster recovery solution for Microsoft Exchange Server 2007 based on a NetApp solution using SnapMirror®, ReplicatorX™ for replication of the production data to the DR site, and SnapManager® for Exchange and SnapDrive® for backup and recovery.

SnapManager for Exchange restore options covered in this document are limited to an up-to-the-minute restore using volume SnapRestore® to recover database and transaction log volumes in the secondary location.

This technical report will not cover the following:

- Detailed setup of Exchange Server recovery environment, including Windows® Server 2003 cluster. For detailed information on recovering Exchange 2007 CMS on a standby cluster in a disaster recovery location, refer to the Microsoft TechNet article on how to recover Exchange CMS on a standby cluster.

- Synchronous SnapMirror and Semi-Synchronous SnapMirror.

- Failback to the primary site.

Intended Audience

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

- Exchange 2007 architecture
- Exchange storage architecture and administration
- Service-level expertise of Microsoft Exchange recovery operations

Working knowledge on NetApp solutions, including the following:

- Data ONTAP®
- SnapDrive for Windows
- SnapManager for Exchange backup and restore procedures
- SnapMirror
- ReplicatorX

Business Continuity and High-Availability Planning

Business continuance (referred to as business continuity) describes the process and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business continuity planning seeks to prevent disruption of mission-critical services and to reinstate full functioning, quickly and efficiently. High availability is a system design protocol and associated implementation that ensures a certain degree of operational continuance during a given measurement period.

Business continuity and high availability are not a specific technology and should integrate a variety of strategies and technologies to address all potential causes of outage, balancing cost vs. acceptable risk resulting into a resilient infrastructure. As a first step in business continuity, high-availability planning is deciding which of the organization's functions are essential to be available and operational during a crisis. Once the crucial/mission-critical components are identified, it is essential to identify your RPO and RTO objectives for the identified crucial/mission-critical apportioning in terms of cost and acceptable risk. To appropriately architect a disaster recovery solution, one must be familiar with the following terms.

Availability
Generally, a degree to which a system, subsystem, service, or equipment is in an operable state for a proportion of time in a functional condition. It refers to the ability of the user community to access the system.

Disaster Recovery (DR)
A process of regaining access to the data, hardware, and software necessary to resume critical business operations after a disaster. A disaster recovery plan should also include methods or plans of copying necessary mission-critical data to a recovery site to regain access to such mission-critical data after a disaster.

High Availability (HA)
A system design protocol and associated implementation that ensure a certain absolute degree of operational continuity of a system, service, or equipment during a given measurement period. High-availability planning should include strategies to prevent single points of failure that could potentially disrupt the availability of mission-critical business operations.

RPO (Recovery Point Objective)
The recovery point objective (RPO) describes a point in time to which data must be restored/recovered in order to be acceptable to the organization's process supported by the data.

RTO (Recovery Time Objective)
The recovery time objective (RTO) is the frontier of time and service level within which service availability must be accomplished to avoid undesirable consequences associated with a break in continuity of a service/process.

Service Level Agreement (SLA)
A formal negotiated agreement between a service provider and a user (typically customers), specifying the levels of availability, serviceability, performance, and operation of a system, service, or application.

Disaster Recovery Model for Microsoft Exchange Server 2007

When architecting a disaster recovery solution for Microsoft Exchange Server 2007, it is important to review your current SLA to derive RTO/RPO objectives. We'll discuss multiple NetApp components that were used to achieve two different levels of RTO/RPO targets below.

Technology Components

SnapManager for Exchange
NetApp SnapManager for Exchange (SME) has achieved the certified Windows logo for Windows Server 2003. SME is a Windows simple SAN-designated Windows Server 2003 certified for backup and recovery solution for Microsoft Exchange Server. SME tightly Integrates with Microsoft Exchange, which allows consistent online backup for Microsoft Exchange environments while leveraging NetApp Snapshot™ copy and SnapMirror technologies. These technologies are critical in protecting Microsoft Exchange Server data, allowing Exchange server administrators to quickly back up and mirror Exchange data faster and efficiently.

SnapDrive for Windows
NetApp SnapDrive for Windows is an enterprise-class storage and data management solution for Microsoft Windows Server environments. SnapDrive enables storage and system administrators to quickly and easily manage, map, and migrate data.

NetApp SnapMirror
NetApp SnapMirror delivers the disaster recovery and data replication solution that today's global enterprises need. By replicating data at high speeds over LAN and WAN, SnapMirror provides the highest possible data availability and fastest recovery for mission-critical applications.

NetApp ReplicatorX
ReplicatorX is an enterprise-class replication solution that can near synchronously replicate block data over any distance, across heterogeneous infrastructures, without operational disruption. ReplicatorX can be used for disaster recovery (DR), data migration, and development/test or other cloning purposes.

Cluster Continuous Replication (CCR)
Cluster continuous replication is a high-availability feature of Microsoft Exchange Server 2007 that combines the asynchronous log shipping and replay technology built into Exchange Server 2007 with failover management features provided by Microsoft Windows Cluster Service. CCR is limited to a two-node, active-passive cluster using the majority node set quorum removing the shared storage technology barrier, providing no single point of failure. transaction logs on the active node are copied to the passive node upon closing the logs and replayed on the passive node to maintain a readily available copy of the database.

Transport Dumpster
The transport dumpster is an essential component for CCR and is a feature built into the hub transport server role. The transport dumpster helps recover the lost data that occurs during an automatic recovery. The transport dumpster takes advantage of the redundancy in the environment to reclaim the lost data from the dumpster queue maintained by the hub transport server.

Disaster Recovery Model: Objective

The primary objective of this disaster recovery model is to achieve highest degree of operational continuance at the primary site with no single points of failure and to have a recovery site and replicate the production Exchange Server data for recovery in case of a disaster. Two scenarios with multiple NetApp components were tested to achieve two different levels of RTO/RPO objectives outlined below.

Business Case 1 (Overview)

To meet a five-minute RPO and a 30-minute RTO, SME backups were scheduled and replicated to the DR site every four hours, and SnapDrive rolling Snapshot copies of the transaction log volume were replicated to the DR site every 30 minutes using SnapMirror. The hub transport server was SAN booted and the boot LUN of the hub transport server was replicated every five minutes.

Business Case 2 (Overview)

To meet a near zero-minute RPO and a 35-minute RTO, SME backups were scheduled and replicated to the DR site every four hours, and SnapDrive rolling Snapshot copies of the transaction log volume were replicated to the DR site every 30 minutes using SnapMirror. The hub transport server system partition was replicated near synchronously to the DR site using ReplicatorX.

Architecture

The architecture used in this model to provide high availability and disaster recovery for Microsoft Exchange Server 2007 contained a primary site for production and a DR site for recovery of Exchange Server 2007. A CCR cluster in the primary site hosting 1,500 users deployed on NetApp storage cluster to provide resiliency on the server hardware, application, and storage levels. A standby Windows MNS cluster deployed with Exchange passive mailbox server roles were deployed in the DR site with the NetApp storage cluster.

The following sections will discuss the two business case scenarios outlined earlier with implementation details, recovery methodologies, and timelines.

Disaster Recovery Scenario: Exchange 2007 Case Study

In this section we'll demonstrate two business cases for Microsoft Exchange Server 2007 taking advantage of the high-availability features like CCR and transport dumpster along with NetApp hardware and software solutions to build a resilient infrastructure in the primary site and a DR site for recovery of Microsoft Exchange Server 2007 in case of a disaster with different RPO/RTO objectives.

Business Case One: Five-Minute RPO/30-Minute RTO Target

The following section will demonstrate a disaster recovery model for Microsoft Exchange Server 2007 using NetApp solutions with the architecture discussed earlier to provide the highest degree of operational continuance in the primary site and replicate the Exchange Server data to the DR site for recovery in case of a disaster with a five-minute RPO and a thirty-minute RTO objective.

The following diagram shows the basic architecture use in this scenario:

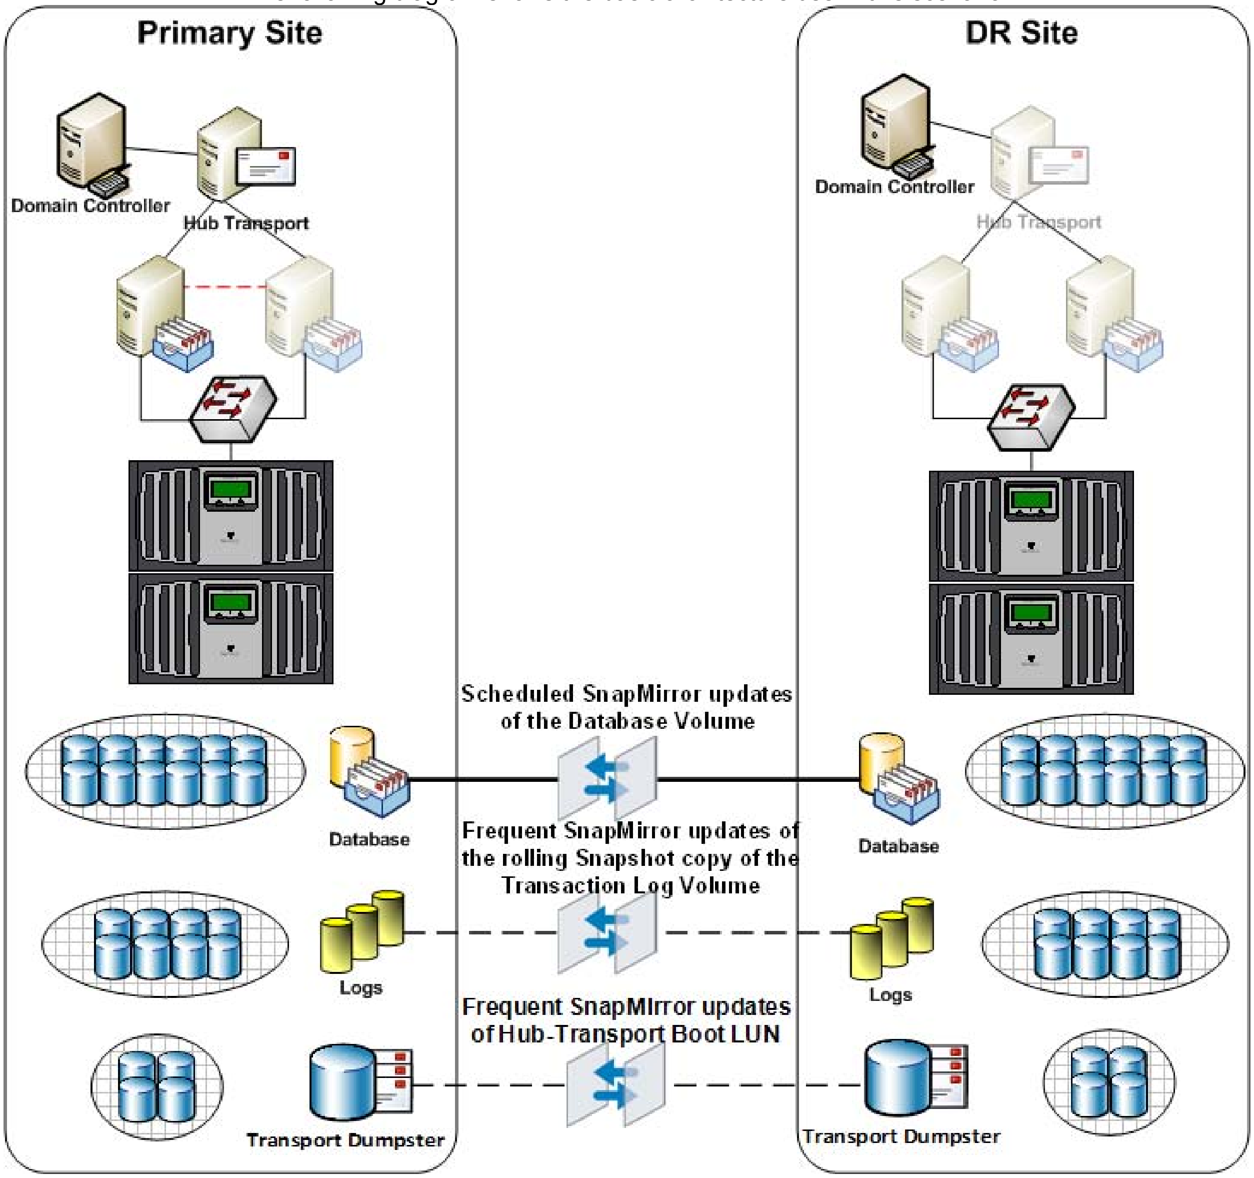
Primary Site
DR Site
Domain Controller
Hub Transport
Domain Controller
Hub Transport
Scheduled SnapMirror updates of the Database Volume
Database
Database
Frequent SnapMirror updates of the rolling Snapshot copy of the Transaction Log Volume
Logs
Logs
Frequent SnapMirror updates of Hub-Transport Boot LUN
Transport Dumpster
Transport Dumpster

Implementation details

As per the architecture discussed earlier, the following sections will provide a detailed overview of the implementation of this scenario in the primary site and the DR site.

Primary Site

The primary objective of this setup is having the primary production site operational with no single-points of failure. Microsoft Exchange Server is deployed in a resilient fashion to provide resiliency in server hardware, application, and storage level.

- A Windows 2003 Active Directory forest was built using one domain controller (raised to 2000 or 2003 forest functional level).
- Two Windows Server 2003 SP1 enterprise edition with Exchange Server 2007 installed with one active mailbox role and one passive mailbox role.
- One Windows Server 2003 with Exchange Server 2007 hub transport server role also acting as a file share witness for the CCR cluster.
- Exchange Server databases hosted on NetApp FAS 6070 active-active cluster.

The following table provides a quick snapshot of the business problems and how they are addressed on the primary site providing a resilient architecture.

Business Problem	***Addressed?***	***How?***	***Description***
Single Point of Failure	✓	CCR + NetApp storage	CCR addressing server resiliency and NetApp storage cluster addressing resiliency on the storage level providing no single point of failure on the application, server hardware, and storage
Disaster Recovery	✓	SDW and NetApp SnapMirror	Replicating the database and log files to the DR site using Asynchronous SnapMirror replication, SME to provide faster backups and rapid restores
Business Continuity	✓	Exchange CCR + NetApp storage	Exchange CCR hosted on NetApp storage cluster provides no single point of failure in case of a server hardware or application or a storage failure
Storage Resiliency	✓	NetApp storage cluster	NetApp storage cluster provides no single point of failure on the storage level
Fast Backup/Recovery	✓	SME 4.0	SME as a simple SAN application integrates well with Exchange HA features providing faster backups and rapid restores
Five-Minute RPO	✓	SnapMirror	Schedule SME backups every four hours and SnapDrive rolling Snapshot updates using SnapMirror every 30 minutes and SnapMirror update of the hub transport server not LUN every five minutes
Less RTO	✓	SME/SDW and SnapMirror	Volume SnapRestore providing an instantaneous restore

Hub Transport Server and Transport Dumpster

When using CCR in your environment an important step is to enable the transport dumpster on the hub transport server. Since the transport dumpster cannot be clustered or replicated at this time, we chose to SAN boot the hub transport server so that the hub transport server boot LUN can be replicated to the remote site for transport dumpster recovery in case of a disaster for the passive node to reclaim lost e-mails.

Transport Database Configuration

When deploying a CCR cluster, the transport dumpster configurations should be taken into consideration. The storage capacity of the hub transport server should contain enough capacity to store long enough for all storage groups in its site, so that messages can be recovered at the passive CCR node in the event of an unscheduled outage.

There are two important settings on the transport server that control how long a message remains in the transport dumpster:

```
MaxDumpsterSizerPerStorageGroup
MaxDumpsterTime
```

Note: By default MaxDumpsterSizerPerStorageGroup is set to 18. To size the transport dumpster properly, Microsoft recommends that you configure the MaxDumpsterSizerPerStorageGroup to 1.5 times of the max message size.

In this scenario we configured the MaxDumpsterSizerPerStorageGroup to 25MB and the max dumpster time to 7 days (default).

Data Replication

To achieve aggressive RPO targets for Microsoft Exchange, scheduled SnapMirror updates of the Exchange database volume and the hub transport boot LUN with the dumpster database and consistent SME backups replicated using SnapMirror were performed along with frequent rolling Snapshot updates of the transaction log volumes.

Replication Schedule

The Exchange database volumes were replicated every four hours, and SnapDrive rolling Snapshot copies were updated using SnapMirror every 30 minutes and the HT boot LUN was replicated every five minutes to the DR site to make sure that at any given point of time the maximum data loss would be only five minutes..

The following picture provides an overview of the replication schedules used in this architecture for a four-hour wind

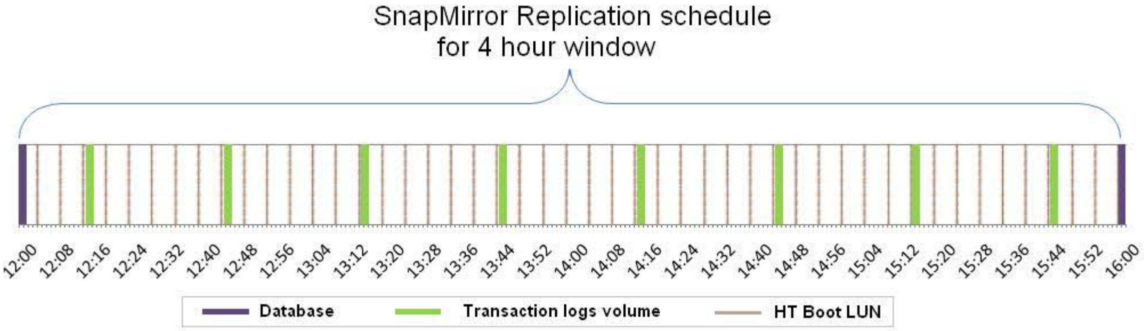

When setting up SME backups and rolling Snapshot copy schedules, it is important to take the following into consideration.

The first Snapshot copy replication must not begin until the initial backup operation is complete.

Subsequent Snapshot copy replication must not begin before the previous replication is complete.

If the replication of the previous Snapshot copy did not complete is still running, then SnapMirror will wait for the replication to complete for it to start the replication of the latest.

Note: The “Snapmirror status” command can be used to monitor the status of the replication of the Snapshot copies and depending on the bandwidth and data sizes, it may be necessary to change the replication schedules for the replication to complete within a reasonable amount of time.

Load Generation

Load Generator was used to generate a 1,500-mailbox load on the Exchange CMS in the test environment with 100MB storage quota. After the initialization process LoadGen was run for eight hours to simulate messaging traffic. SME was configured to create Snapshot copies every four hours in order to provide data need to calculate the rate of change.

DR Site

As the objective of this setup is to have an operational DR site for recovery in case of a disaster, the following are required:

- A Windows 2003 additional domain controller
- Two-node Windows Server 2003 Enterprise Edition with SP1 MNS cluster with Exchange Server 2007 installed with passive mailbox server roles
- One server hardware identical to the hub transport server on the primary server to be SAN booted from the hub transport server Boot LUN replicated from the primary site.

Storage Layout

The following storage layout was used in the test environment.

	Total Capacity	Used Capacity	LUNs
Hub Transport Boot LUN	80GB	56%	Boot LUN
CCR Active Node			
DB VOL	819GB	54%	E:\SG1 Mailbox Store
			F:\SG2 Public Store
Logs Vol1	200GB	49%	G:\SG1 Logs
Logs Vol2	56GB	20%	H:\SG2 Logs
CCR Passive Node			
DB VOL	819GB	54%	E:\SG1 Mailbox Store
			F:\SG2 Public Store
Logs Vol1	200GB	49%	G:\SG1 Logs
Logs Vol2	56GB	20%	H:\SG2 Logs

Note: The storage layout in the DR site is an exact replica of the primary site.

Setting Up SnapMirror Relationship

The following section describes how to set up a SnapMirror relationship for the test environment.
Create the SnapMirror destination volumes to be the DR site.
Note: These volumes have to be equal or greater than the size of the source volumes.

On the source storage controller console, use the options snapmirror.access command to specify the hostnames of the storage systems that are allowed to copy data directly from the source storage system. For example:

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

Restrict the volumes to allow SnapMirror to access them using vol restrict command:

```
Vol restrict <volume_name>
```

Initialize the SnapMirror Process

From the destination storage controller console, use the SnapMirror initialize command to create an initial seed of the source on the destination and start the mirroring process.

```
Snapmirror initialize –s <src_storage_name>:<src_vol> <dest_storage_name>:<dest_vol>
```

Note: You can use the SnapMirror Status command to check the status of the SnapMirror initialization as shown below:

```
NB-6070-1> snapmirror status

Snapmirror is on

Source               Destination          State        Lag          Status
```

```
NB-6070-1:HTPSLUN   NB-6070-2:HTRSLUN   snapmirrored 00:05:06      Idle
NB-6070-1:CCRNBDB   NB-6070-2:CCRNCDB   source       transferring (276 MB done)
NB-6070-1:CCRNBLOGS NB-6070-2:CCRNCLOGS source       transferring (56 MB done)
```

Setting Up SME Scheduled Backups and Rolling Snapshot Copies

SME backups can be scheduled to run from the SME console using the Windows Task Scheduler.
The following steps were taken to create an SME full backup every four hours to satisfy our DR requirements:

1. Open the SME console and click the Backup and Verification settings.
2. Select the storage group that you want to have backed up and replicated.
3. Ensure to check the Update SnapMirror after operation checkbox and then click the Schedule button.
4. SME will then create a new Windows task for this backup operation.

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`*
2. *DriveLetterList* is a list of space-separated drive letters that contain the transaction logs.
3. **Example:** `sdcli snap update_mirror -D g h i`

4. Select Start Menu > Settings > Control Panel > Scheduled Tasks.
5. Double-click Add Scheduled Task.
6. 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.
7. After the following panel appears, select from the list of frequencies, taking into consideration your RPO objectives, then click Next.
   For an RPO of 33 minutes, the schedule is set to every 30 minutes.
8. 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.
9. Then, type the user name and password for the scheduled task, then click Next.
10. 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.

Once you have completed the above tasks, you should have two Windows Task Scheduler jobs created and ready to execute, as seen below:

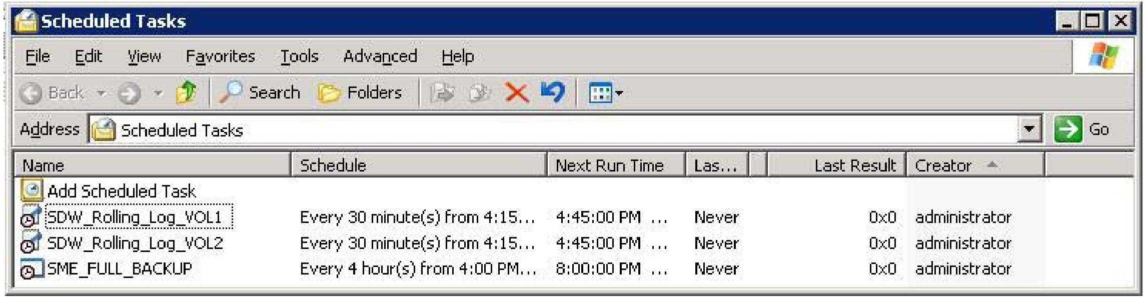

Address: Scheduled Tasks

Name	Schedule	Next Run Time	Las...		Last Result	Creator
Add Scheduled Task						
SDW_Rolling_Log_VOL1	Every 30 minute(s) from 4:15...	4:45:00 PM ...	Never		0x0	administrator
SDW_Rolling_Log_VOL2	Every 30 minute(s) from 4:15...	4:45:00 PM ...	Never		0x0	administrator
SME_FULL_BACKUP	Every 4 hour(s) from 4:00 PM...	8:00:00 PM ...	Never		0x0	administrator

Note: SME and SnapDrive do not operate concurrent operations, so it is important to make sure that the two operations do not start at the same time. If the two operations are scheduled to run at the same time, the first operation that this processed will start and the other operation will fail.

Replication of Hub Transport Boot LUN

The following steps outline how to create a SnapMirror update of the hub transport server boot LUN using the snamirror.conf file.

We used the `Snapmirror.conf` file to schedule the SnapMirror updates of the hub transport server boot LUN to the DR site.

- Edit the snapmirror.conf file on the destination storage using:
- `wrfile /etc/snapmirror.conf` and click enter
- Type `<src_storage>:<vol_name> <dest_storage>:<vol_name> - <Minute> <Hour> <week of the month> <day of the week>`

Example**: `<NB-6070-1>:<HTPSLUN> <NB-6070-3>:<HTRSLUN> - 0-59/5 * * *`**

Note: The above example will run SnapMirror updates to the destination every five minutes.

Rate of Change

Rate of change is the amount of changed data from the previous successful Snapshot copy to the current Snapshot copy. This information can be used to properly schedule SnapMirror times for replication of data to the DR site.

Rate of Change for Exchange Databases

In order to accurately calculate the rate of change for the databases volume, SME backups were scheduled and monitored for a 12-hour period. During this period the `snap delta` command can be used to estimate the rate of change between Snapshot copies.

The following command can be used to display the rate of change between all Snapshot copies on the database volume:

```
Snap delta CCRNCDB

Volume CCRNCDB
working...

From Snapshot                             To                    KB changed    Time         Rate
(KB/hour)
---------------                   --------------------         ----------- ------------ ---------------
exchsnap__ccr_05-18-2007_15.47.01         Active File System           9776    0d 17:55   545.544
exchsnap__ccr_05-04-2007_04.00.07 exchsnap__ccr_05-04-2007_08.00.07 2269236    0d 03:59   3725485.714
exchsnap__ccr_05-04-2007_00.00.03 exchsnap__ccr_05-04-2007_04.00.07 2636916    0d 03:59   659412.170
exchsnap__ccr_05-03-2007_20.00.03 exchsnap__ccr_05-04-2007_00.00.03 2451352    0d 03:58   616649.234
exchsnap__ccr_05-03-2007_16.00.03 exchsnap__ccr_05-03-2007_20.00.03 2173904    0d 04:01   541220.912

Summary...

From Snapshot   To                   KB changed  Time         Rate (KB/hour)
--------------- -------------------- ----------- ------------ ---------------
Snap_base        Active File System   49038424     17d 17:43   115190.059
```

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.

From Snapshot Copy	To Snapshot Copy	Size (KB)
exchsnap__ccr_05-18-2007_15.47.01	Active File System	9776
exchsnap__ccr_05-04-2007_04.00.07	exchsnap__ccr_05-04-2007_08.00.07	2269236
exchsnap__ccr_05-04-2007_00.00.03	exchsnap__ccr_05-04-2007_04.00.07	2636916
exchsnap__ccr_05-03-2007_20.00.03	exchsnap__ccr_05-04-2007_00.00.03	2451352
exchsnap__ccr_05-03-2007_16.00.03	exchsnap__ccr_05-03-2007_20.00.03	2173904

Note: 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.

Rate of Change for Transaction Logs

To calculate the amount of data generated by transaction logs, the logs volume were analyzed for four hours during the LoadGen run. Sample data was collected for four hours out of the eight-hour test run. The number of logs generated per storage group per hour was collected, and the following calculations were made.

Time	Total # of Logs Generated for First Storage Group	Total Size of Logs in MB
10:30 AM	1218	1218
11:00 AM	1198	1198
11:30 AM	1234	1234
12:00 PM	1247	1247
12:30 PM	1170	1170
1:00 PM	1227	1227
1:30 PM	1184	1184
2:00 PM	1187	1187

Note: When laying out Exchange data onto the storage appliance, take into careful consideration the factors outlined earlier. Things like rate of change, bandwidth available for replication, 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, its 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.

Recovery Methodology in the Event of a Disaster

This section will outline the recovery methodology on the DR site in case of a disaster. The following steps must be taken prior to running recovery on the DR site.

LUN drive letters must be the same as the primary site.
SME configuration must be completed.

Volume Snap-Restore

When failing over to a secondary site and restoring the whole Exchange cluster, the 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.

The following steps were used to recover the Exchange CMS at the DR site using Volume SnapRestore method:

Preparation

1. Issue a `SnapMirror break` of the HT boot LUN, Exchange Server DB, log, MTS/SMTP, and SnapInfo volumes.

Example: `snapmirror break <Vol_Name>`

Recover the Hub Transport Server

2. Remove any LUN mapping carried from the primary site and boot the hub transport server on the replicated boot LUN.

Recover the Exchange CMS

3. Perform a Volume `SnapRestore` of the Exchange Server log volume(s) with the most recently completed SnapDrive initiated Snapshot copy.
   a. Typically this will be the rolling Snapshot copy.
   b. If a site failover is initiated shortly after an SME backup but before the next rolling Snapshot copy is initiated, there may a case in which the `SME_eloginfo` and/or `_recent` Snapshot copy is more recent than the rolling Snapshot copy.
4. Ensure all recovered volumes are online.
5. Clear all LUN mappings that may have been carried over from the production site.
6. Turn `SnapMirror off` on the storage appliance.
   Note: If this is not possible due to other SnapMirror relationships, remove the Exchange destination volumes from the snapmirror.conf file.
7. Connect all the LUNs on the first recovery node, which will recover and host the Exchange CMS at the DR site using SnapDrive.
8. From the recovery node verify that all the LUNs are connected.
9. To recover the Clustered Mailbox Server (CMS), run **`setup.com /recoverCMS /CMSName:<CMS_Name> /CMSIPaddress:<IP_Address>`** from the bin directory under the Exchange program files.

SnapManager Configuration and Restore

10. Start SME and rerun the configuration wizard.
    Note: Ensure that the snapinfo directory matches the production site. This is especially important if a dedicated SnapInfo LUN is used.

11. Restart SME and select Restore and select “First Storage Group” for the most recent Snapshot copy and click Restore.
    Note: Ensure that the up-to-the minute restore option is selected and the recover and mount database after restore option is unchecked and click Restore.
12. Repeat the previous step for all the storage groups.
13. Using Exchange Management Console mount all of the Exchange storage groups.
14. Perform an SME backup of the newly recovered Exchange environment.

Restore Time

A disaster simulation was perfumed two hours 15 minutes after the last SME backup and the following metrics were observed during the recovery at the DR site to track the total recovery time and the recovery point for the Exchange CMS:

Initial Steps	Time to Completion
Break the SnapMirror relationship for all volumes	30 Seconds
Map LUN and boot Hub Transport Server	4 Minutes
Perform a SnapRestore operation of all volumes	1 Min
Clear LUN mappings that may have carried from the Primary Site	1 Min
Connect and Mount all LUNs on Exchange Server	5 Min
Recover Exchange CMS	2 Min
Total Time	13:30 Minutes

Storage Group	# of Transaction Logs to Replay	Size of the Transaction Log directory	Log Replay Time	Total SME Restore Time
First Storage Group	4863	4.74GB	11:30 Minutes	15:20 Minutes
Totals	4863	4.74GB	11:30 Minutes	15:20 Minutes

15. Open Exchange Management Shell and run **`Update-StorageGroupCopy identity:<DomainName>\<CMSName>`** to seed the passive node.

The total recovery time for the scenario tested was approximately 27 minutes and the recovery point objective was 15 minutes without the transport dumpster and five minutes with the transport dumpster being replicated every five minutes from the production site.

When planning RTO and RPO, it is important to know the approximate number of logs that would need to replay and the rate at which the logs replay. In this recovery scenario, it approximately took 11 minutes to replay the logs and the total restore time was 15 minutes. From this data we can calculate the logs for this test scenario were replayed at a rate of approximately 405 logs per minute or 405MB per minute.

Business Case Two: Zero-Minute RPO/30-Minute RTO Target

The following section will demonstrate a disaster recovery model for Microsoft Exchange Server 2007 with the same architecture discussed earlier to achieve a zero-minute RPO and a 35-minute RTO targets with the use of ReplicatorX.

The following diagram shows the basic architecture used in this scenario:

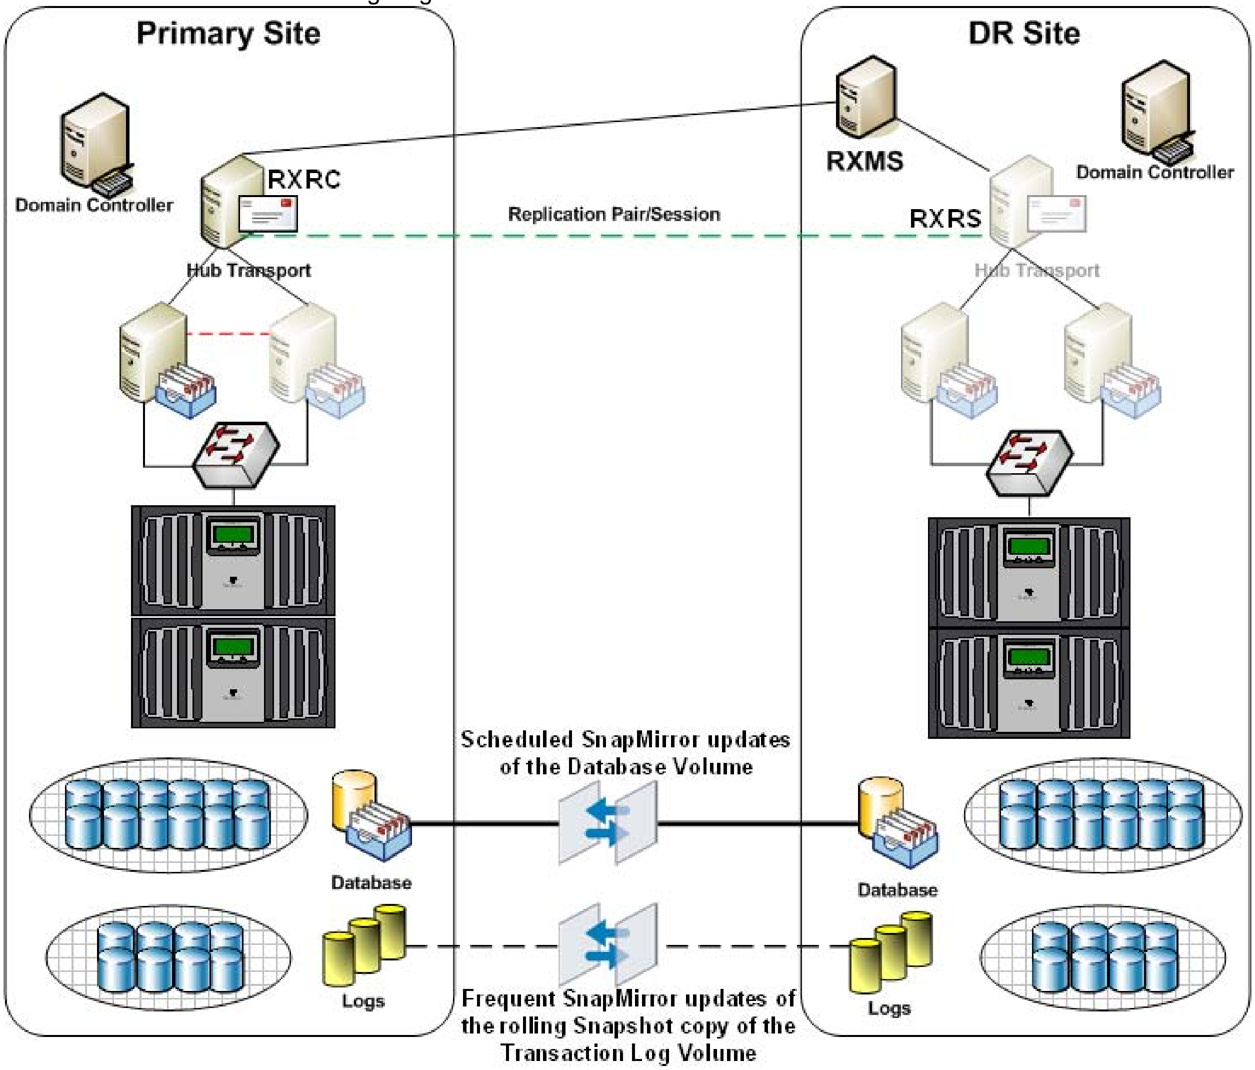

Implementation details

The implementation in this business scenario is an exact replica as the previous architecture except for the hub transport server replication. ReplicatorX was used to replicate the hub transport server to the DR site continuous asynchronous replication and provide the ability of replicating the system partition of the hub transport server to either NetApp or third-party storage. The following section will describe the implementation of ReplicatorX to replicate the hub transport server data to the DR site.

The following table provides a quick overview of the business problems and how they are addressed on the primary site providing a resilient architecture.

Business Problem	Addressed?	How?	Description
Single Point of Failure	✓	CCR + NetApp storage	CCR addressing server resiliency and NetApp storage cluster addressing resiliency on the storage level providing no single point of failure on the application, server hardware, and storage
Disaster Recovery	✓	SDW and NetApp SnapMirror	Replicating the database and log files to the DR site using Asynchronous SnapMirror replication, SME to provide Faster Backup's and rapid restores
Business Continuity	✓	Exchange CCR + NetApp Storage	Exchange CCR hosted on NetApp storage Cluster provides no single point of failure in case of a Server hardware or Application or a storage failure
Storage Resiliency	✓	NetApp Storage Cluster	NetApp Storage Cluster provides a No-Singe-Point of failure on the Storage level
Fast Backup/Recovery	✓	SME 4.0	SME as a simple SAN application integrates well with Exchange HA features providing faster backups and rapid restores
Zero RPO	✓	SnapMirror and ReplicatorX	Scheduled SnapMirror updates of the database LUN and rolling Snapshot updates of the transaction logs LUNs with continuous asynchronous replication of hub transport server using ReplicatorX
35-minute RTO	✓	SME/SDW and SnapMirror	Volume SnapRestore providing an instantaneous restore

Replication of Hub Transport Server Using ReplicatorX

ReplicatorX release 4.0 was used in this architecture to replicate the hub transport server data to the DR site.

ReplicatorX Components
RXRC: ReplicatorX replication client
Installed on the each primary site server which has direct access to the source volumes.

RXRS: ReplicatorX Replication Server
Installed on at least one server in the secondary site which has access to the target volumes.

RXMS: ReplicatorX Management Server
Each secondary site must have at least one server with RXMS installed on it.

Please refer to ReplicatorX documentation for more information on installation, configuration and administration.

The following steps outline the how to replicate the hub transport server to the DR site using ReplicatorX.

1. Install RXRC component on the hub transport server in the primary site.
2. Install RXRS component on the recovery server which has access to the target volume for the hub transport volume.
3. Install RXMS on an additional server to manage and administer ReplicatorX.
4. Open the ReplicatorX management GUI on the RXMS server.
    a. Open Internet Explorer.
    b. Type http://<Servername>/RepX_R4_0.
5. Add the primary site and add the RXRC client to the primary site.
6. Add the DR site and the RXRS client to the DR site.

Preparing RXRC

1. Right-click the RXRC server and click Add Volume.
2. Specify a name for the volume and specify the source volume. In this case it is the C drive of the hub transport server.

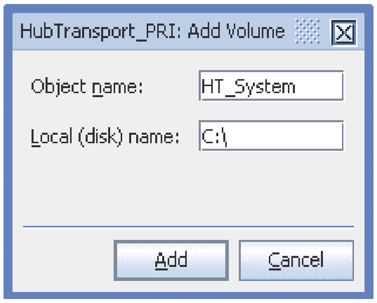

HubTransport_PRI: Add Volume

Object name: HT_System

Local (disk) name: C:\

3. Right-click Add Volume and click Add N: Volume.

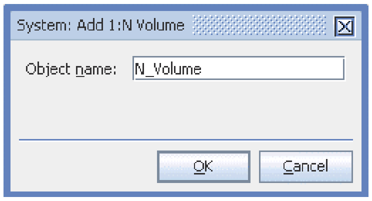

System: Add 1:N Volume

Object name: N_Volume

4. Right-click the additional volume on the RXRC server and click Set Role, and select Bitmap Role and click OK.
5. Right-click the RXRC server and click Add Bitmap File and specify the file name, size and location. (**Bitmap volumes should be RAW volumes and should not be formatted with a file system.**)

Note: In this scenario the bitmap volume size was 2.5GB (approximately 3% of the source volume). It is recommended to follow proper sizing and capacity planning guidelines when sizing bitmap volumes. Please refer to ReplicatorX Theory and operation Planning Guide for sizing guidelines.

Preparing RXRS

The RXRS server requires two volumes (UD-LOG and MD-LOG).

1. Right-click a volume in the RXRS server and click Set Role and select UD-Log and click OK.
2. Right-click another volume in the RXRS server and click Set Role and select MD-Log and click OK.

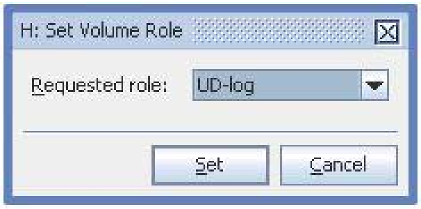

H: Set Volume Role

Requested role: UD-log

Note: In this scenario the MD-Log volume size was 2.5GB (approximately 3% of the source volume) and the UD Logs volume size was 14GB (approximately 20% of the source volume). It is recommended to follow proper sizing and capacity planning guidelines when sizing bitmap volumes. Please refer to ReplicatorX Theory and operation Planning Guide for sizing guidelines.

(UD-Log and MD-Log should be RAW volume and should not be formatted with a file system.)
Note: Target Volume should be a Primary basic partition and should not be formatted with a file system

Creating and Activating a Session/Pair

1. From the RXMS GUI, click Graphic-Configuration Mode, and select the volume which you want to replicate and drag It to the target volume.
   Note: The target volume should be of equal size or greater than the source volume.
2. Add Pair Dialog appears, specify the pair name and session name, check the Activate Pair check box and set the synchronization method to "**Online**" and click Next.

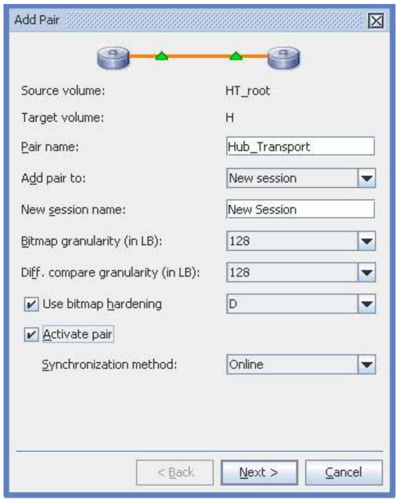

Add Pair

Source volume: HT_root

Target volume: H

Pair name: Hub_Transport

Add pair to: New session

New session name: New Session

Bitmap granularity (in LB): 128

Diff. compare granularity (in LB): 128

☑ Use bitmap hardening: D

☑ Activate pair

Synchronization method: Online

3. Select the compression method to be enabled and the DPR mode to slim only and click Finish.
4. Upon clicking finish the initial synchronization starts.

Recovery Methodology in the Event of a Disaster

The recovery methodology in this architecture is same as the previous business case. The following section will outline the recovery methodology of the hub transport server using ReplicatorX.

Recover the Hub Transport Server

1. Open the RXMS Management GUI.
2. Right-click the session in the Graphic Configuration mode and click the session name and click Freeze session.
3. Once the freeze is completed there will be an ❄ icon displayed on the target volume as shown in the figure below.

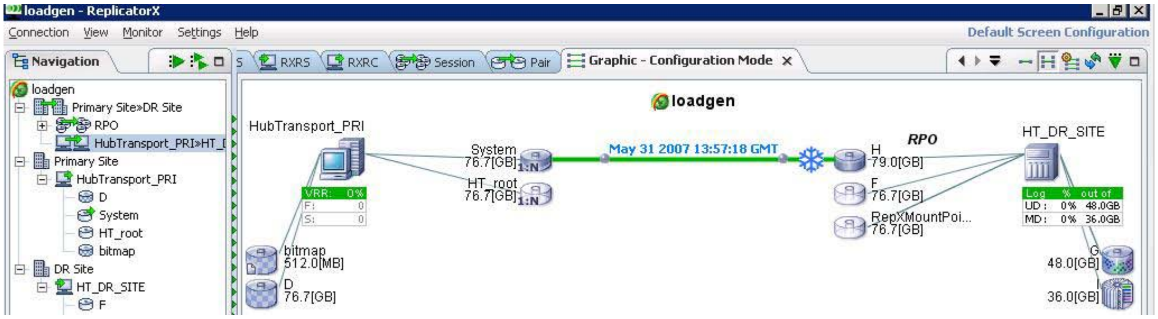

4. Login to the RXRS server and from the command line change directory to <X:\Program Files\RepX\R4_0\RepX Replication Server\bin> and run the disk signature command to change the 0th sector of the disk to make if usable.

```
disksignature tv=d:\ fixall=1
```

Note: Please refer to ReplicatorX Administration Guide for more information on the Disk Signature command like and options.

5. Reboot the server from the target volume (Which in this case is a Secondary hard drive on the hub-transport server.

Note: If the target volume resides on NetApp Storage, then the volume should be mapped to the DR server and booted up.

Exchange CMS (Clustered Mailbox Server) was recovered using the steps described in the previous business case.

Restore Time

A disaster simulation was perfumed two hours 15 minutes after the last SME backup and the following metrics were observed during the recovery at the DR site to track the total recovery time and the recovery point for the Exchange CMS:

Initial Steps	Time to Completion
Break the SnapMirror relationship for all volumes	30 Seconds
Freeze the ReplicatorX Session	10 Seconds
Run DiskSignature.exe to recover the Target Volume	10 Seconds
Reboot the Server from the target Volume	5 Minutes
Perform a SnapRestore of all volumes	1 Minute
Clear LUN mappings that may have carried from the Primary Site	1 Minute
Connect and Mount all LUNs on Exchange Server	5 Minutes
Recover Exchange CMS	2 Minutes
Total Time	**14:50 Minutes**

Storage Group	# of Transaction Logs to Replay	Size of the Transaction Log directory	Log Replay Time	Total SME Restore Time
First Storage Group	4863	4.74GB	11:30 Minutes	15:20 Minutes
Totals	**4863**	**4.74GB**	**11:30 Minutes**	**15:20 Minutes**

The total recovery time for the scenario tested was approximately 30 minutes 10 Seconds and the recovery point objective was 15 minutes without the transport dumpster and nearly zero with the transport dumpster being replicated using ReplicatorX asynchronously to the DR site.

When planning RTO and RPO, it is important to know the approximate number of logs that would need to replay and the rate at which the logs replay. In this recovery scenario, it approximately took 11 minutes to replay the logs and the total restore time was 15 minutes. From this data we can calculate the logs for this test scenario were replayed at a rate of approximately 405 logs per minute or 405MB per minute.

Summary

Microsoft Exchange is a mission-critical application and it can cripple the operational productivity if it becomes unavailable. NetApp has proven data protection and disaster recovery tools for Microsoft Exchange. SnapManager for Exchange backup and restore capabilities combined with SnapDrive, SnapMirror technologies and ReplicatorX provide a solid and robust solution for protecting and recovering your exchange data while meeting stringent RPO and RTO objectives based on your business requirements.

Appendix A: Best Practices

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

- Determine which data to be replicated and the replication methods (Synchronous or Asynchronous). This can impact the available bandwidth and RPO/RTO objectives.
- Plan volume layout to help archive RPO/RTO objectives: Separate users, business units requiring faster access 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: Reference

Microsoft Exchange 2007
Microsoft Exchange 2007 Architecture

Exchange 2007 CCR (Cluster Continuous Replication)

Exchange 2007 Hub Transport Server Role

How Transport Dumpster Works

Exchange 2007 Best Practices Guide

Data ONTAP
Data ONTAP Documentation

SnapDrive for Windows
SnapDrive for Windows 4.2.1 Installation and Administration Guide

SnapDrive for Windows Best Practices Guide

SnapManager for Exchange
SnapManager for Exchange Installation, Configuration, and Administration Guide

SnapManager for Exchange 4.0 Best Practices Guide

ReplicatorX
ReplicatorX Installation Guide

ReplicatorX Administration Guide

ReplicatorX Reference Guide

ReplicatorX Theory of Operation and Planning Guide