



SnapManager[®] 4.0 for Microsoft[®] Exchange

Best Practice Guide

Shannon Flynn, NetApp
January, 2007 | TR-3541



Table of Contents

1. Executive Summary.....	4
1.1 Purpose and Scope.....	4
1.2 Intended Audience	4
2. SME 4.0 New Features.....	5
2.1 Exchange Server 2007 Support	5
2.2 Storage Groups and Databases.....	5
2.3 Exchange Server 2007 High-Availability Features	5
2.4 Integrity Verification on a SnapMirror® Destination.....	9
2.5 Restore Backup Sets from a Different Exchange Server.....	9
2.6 Recovery Storage Group Integration.....	9
3. Exchange and NetApp Storage Design	10
3.1 Aggregates and Flexible Volumes.....	11
3.2 LUN Sizing	13
3.3 Volume Sizing	14
3.4 Effects of Transaction Log Archiving Using NTFS Hard Links on Volume Size	15
3.5 Fractional Space Reservation	15
4. SnapManager 4.0 for Exchange Installation.....	16
4.1 SME Interactive Mode	16
4.2 SME Unattended Mode.....	17
5.0 SnapManager 4.0 User Interface.....	17
5.1 SME 4.0 Dashboard View	18
5.2 SME 4.0 PowerShell Interface	19
6.0 Configuration	19
6.1 Configuration Wizard Rules.....	19
6.2 Volume Mount Point Support	20
6.3 Fractional Space Reservation Policies	21
6.4 Copyless Transaction Log Archiving	23
7. Backup and Restore	24

7.1 Backup	24
7.2 Verification	26
7.3 Restore.....	28
8. SnapMirror.....	29
9. Summary	30

1. Executive Summary

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.

Because of the increasing importance of Microsoft Exchange Server for any business, Exchange data protection, disaster recovery, and high availability are of increasing concern. As the importance of Exchange within the organization increases, companies expect quick recovery times with little or no data loss. With Exchange databases growing rapidly in size every day, it is increasingly difficult to complete time-consuming backup operations in a reasonable amount of time. When an outage occurs, it can take days to restore service from slower media such as tape, even assuming that all of the backup tapes are available and error free.

Network Appliance offers a comprehensive suite of hardware and software that enables an organization 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 4.0 software is available for Microsoft Exchange Server 2003 and 2007. SnapManager 4.0 for Exchange (SME) has achieved a [Certified for Windows](#) logo for Windows® Server 2003 Standard and Enterprise editions. SME is also a Microsoft [SimpleSAN](#) designated and a Windows Server 2003 certified backup and recovery solution for Exchange server. SME is tightly integrated with Microsoft Exchange, which allows for consistent online backups of your Exchange environment while leveraging NetApp Snapshot™ copy technology. SnapManager 4.0 for Exchange supports Microsoft Volume Shadow Copy Services (VSS), Virtual Disk Services (VDS), and other Microsoft storage technologies. (For details on VSS, see [Microsoft KB article 822896](#).) SME is a VSS (Snapshot copy) requestor, which means that it uses the Microsoft-supported VSS subsystem to initiate backups. SME provides a complementary feature set for the new Microsoft Exchange 2007 data replication features. SME works with Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR) replica databases and provides a rich feature set to leverage these new technologies.

1.1 Purpose and Scope

The success or failure of any software or infrastructure deployment hinges on making the proper design and architecture decisions in the planning phase. This guide provides best practices for deploying and using SnapManager 4.0 for Microsoft Exchange with a NetApp storage appliance and any supporting software. Organizations that want to get the most out of their NetApp storage investment for Exchange will benefit from putting into practice the recommendations in this report.

1.2 Intended Audience

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

- *SnapManager for Exchange Installation and Administration Guide*
- *SnapDrive Installation and Administration Guide*
- *Data ONTAP System Administrators Guide*

Readers of this best practice guide should have a solid understanding of the Exchange storage architecture and Exchange administration, as well as Exchange backup and restore concepts. The recommendations in this document are best practices to assist with the design, implementation, and configuration of SnapManager for Exchange in Windows Server 2003 environments with Microsoft Exchange Server 2003 and Microsoft Exchange Server 2007.

Note:

- The SnapDrive® and SnapManager for Exchange installation and administration guides can be found on the NOW™ (NetApp on the Web) site: <http://now.netapp.com>.

2. SME 4.0 New Features

2.1 Exchange Server 2007 Support

SME 4.0 provides support for the newly released Exchange Server 2007 from Microsoft. Exchange Server 2007 introduced many new features that help customers achieve flexibility in their deployment scenarios and provide built-in high-availability features. SME 4.0 is tightly integrated with these new features, providing a complementary feature set for Exchange 2007.

2.2 Storage Groups and Databases

Exchange 2007 extends the number of storage groups and databases from 4 storage groups and 5 databases per storage group (20 maximum databases per server) to 50 storage groups and 50 databases. This provides the Exchange administrator more options for where to store groups of mailboxes when laying out an Exchange deployment. SME 4.0 supports the migration, backup, and restore of up to 50 storage groups and 50 databases. For Exchange 2003, the current limits of 4 storage groups and 5 databases per storage group remain.

Note:

- Larger numbers of storage groups and Exchange databases may increase backup and restore times. This also depends on the size of the data set, how powerful the Exchange Server is, and how busy the storage controller may be.

2.3 Exchange Server 2007 High-Availability Features

Local Continuous Replication (LCR)

LCR, a new feature for Exchange 2007, is a single-server-based replication feature using asynchronous log shipping. A passive copy, or replica, resides on a secondary storage destination connected to the same Exchange Server as the active copy of the Exchange data. Once transaction logs are written into the active database, they are shipped to the passive database and replayed into that database.

In the case of a primary LCR storage failure where the active Exchange database becomes unavailable, the Exchange administrator can quickly perform a manual switchover to the passive Exchange database and restore mail flow in the organization.

To enable LCR, both the active and passive databases must be seeded. Seeding the databases synchronizes the passive database to the active database. When that process is complete, LCR shows a healthy state and replication proceeds. For more information on seeding in an LCR configuration, see the [Microsoft TechNet article on how to seed a LCR copy](#).

Migration

Migrating an LCR-enabled Exchange database involves the same processes as migrating a normal Exchange database. However, to distinguish an LCR-enabled Exchange database, SME appends \LCR to the storage group name.

Note:

- The active and passive copies of your Exchange data should never reside on the same storage appliance volume.

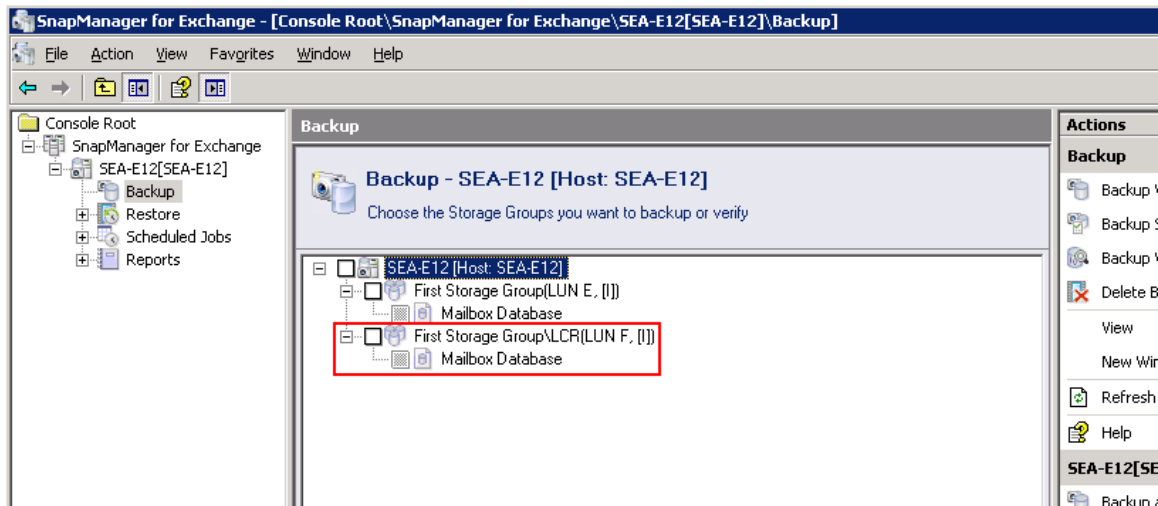


Figure 1) SME 4.0 appends \LCR to an LCR-enabled storage group.

Backup

Using SME, you can perform regular backups on the active Exchange database, the passive Exchange database, or both simultaneously. When a backup is performed on both the active and passive Exchange databases at the same time, two backup sets are created. The passive backup set is created first. When it is completed, the active database backup set is created and given the `_recent` backup name.

Note:

- Backing up both the active and passive Exchange databases at the same time creates two SME backup sets. This consumes additional storage space on the LUNs that contain the Exchange data files.
- SME does not perform any special operation to delete transaction log files on a passive database after performing a backup operation on the active database. When a backup set is created on an LCR-enabled Exchange database, the Exchange writer handles truncating the transaction logs on both the active and passive databases.

LCR is not a replacement for your backup and recovery solution. If a virus or another type of corruption passes the checksum verification for CCR, you can potentially have a corrupt active and passive database. The ability to back up one or both of the active and passive Exchange databases provides an additional layer of data protection for Exchange. SME backup sets can be scheduled at regular times, providing a timestamp-based history of your backup sets. If the need arises, you can choose a backup set from a previous point in time to recover your Exchange data. You can also restore your SME backup sets to another Exchange Server in case of disaster.

SME can verify backup sets from both the active and passive Exchange databases and mark those backup sets as verified. If you want to verify a backup set from the active Exchange database, you must schedule it at an appropriate time in order not to affect the performance of your Exchange storage. Performing the verification on the passive node, however, offloads the verification process onto the spindles that contain the passive Exchange data, thus eliminating the load on the production spindles.

Note:

- Although the verification load is now offloaded onto a second set of spindles, the load still remains on the production Exchange server. Take this into consideration when using this new feature. Implementing the remote verification functionality in SME can help to eliminate the load on the production Exchange server as well.

Restore

When restoring a backup of an LCR-enabled Exchange database, you can restore the active copy or the passive copy of the Exchange data. The restore process is identical to a normal SME restore process. However, when you restore the passive copy of the Exchange data, LCR is disabled and must be reseeded before replication can occur again.

Cluster Continuous Replication (CCR)

CCR is a high-availability feature of Microsoft Exchange 2007 that combines asynchronous log shipping with the failover and management features of the Microsoft Cluster service. A passive copy, or replica, resides on a secondary storage destination connected to a passive server. This passive server must be a member of the same cluster as the active Exchange server. Similarly to LCR, as a transaction log is written into the active database, it is then shipped to the passive node and replayed into the passive database.

CCR is differentiated from LCR by providing not only a passive copy of the Exchange data for recovery, but by providing server-level protection as well. In the case of a disaster where you lose the Exchange data, the Exchange server, or both, you can quickly fail over to the passive server and Exchange data and be fully operational.

As with LCR, to enable CCR, the passive database must be seeded. Seeding the databases synchronizes the passive database to the active database. Once that is complete, CCR shows a healthy state and replication proceeds. For more information on seeding in a CCR configuration, refer to the [Microsoft TechNet article on how to seed a CCR copy](#).

CCR is restricted to a Majority Node Set (MNS) cluster only. Shared Copy Cluster (SCC), which uses traditional shared resources and disks, does not support CCR. To create an MNS cluster, you need two servers, one active and one passive, and an additional server that serves as a File Share Witness. The File Share Witness serves as the tie-breaker and is used to decide which node is the owner of the resources and disks for that cluster. For more information about MNS, see the [Microsoft TechNet's Majority Node Set article](#).

Migration

SME 4.0 recognizes both the CCR active node and the passive node when launched. It will migrate both the active node and the passive node onto NetApp storage at the same time. The Migration Wizard walks you step by step through the migration process.

Note:

- The CCR databases must be in a healthy state before you can migrate them.
- LUNs on different MSCS nodes should not share the same storage volume.
- The same LUN on each node must have identical storage system layouts.
 - For example, if LUNs E, F, and G on NodeA are sharing a storage system volume. LUNs E, F, and G on NodeB must also share a storage system volume.

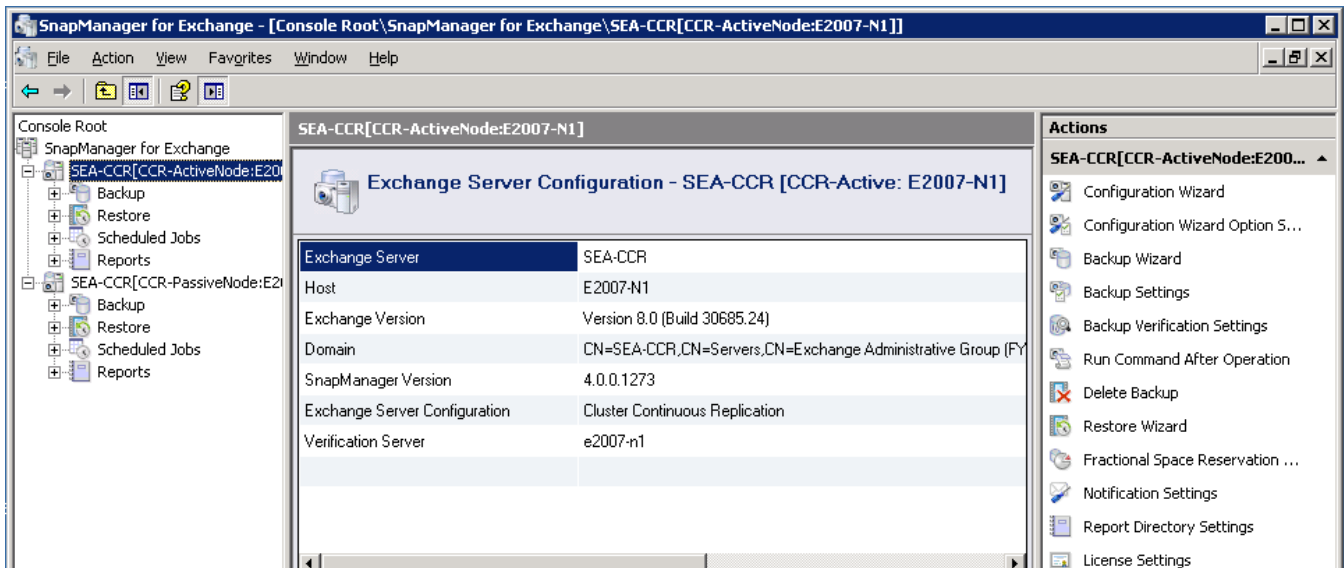


Figure 2) SME 4.0 showing a CCR-enabled Exchange server configuration.

Backup

Similar to LCR, CCR is not a replacement for your backup and recovery solution. If a virus or another type of corruption passes the checksum verification for CCR, you may potentially have a corrupt active and passive database. SME backup sets mitigate this problem by providing a timestamp-based history of your Exchange server backups. This allows you to go back in time if necessary and recover your Exchange data from that point.

During a backup or verification operation, you can choose either the active node or the passive node to back up, but not both. Performing a backup operation from the passive node eliminates a potential load incurred during the backup and verification process. If you choose to do a verification using the passive node backup, those processes are performed on the passive server and the spindles, thus eliminating the load on both the production Exchange server and the spindles.

Restore

When restoring a backup set in a CCR environment, take into consideration the following points:

- A backup set always resides on the node where it was created.
- When you select either the active or the passive node, SME restores backup sets back to the node where the backup was originally taken.
- SME moves the Exchange virtual server resources to the node where the restore is taking place before the restore operation occurs.

Rules for the CCR restore process

The following rules explain the seeding requirements after a type of restore has been performed on either the active database or the passive database.

Restore Type	Seeding required?
Up-to-the-minute restore	No
Point-in-time restore	Yes

Seeding is required with a point-in-time restore because that type of restore does not have an up-to-the-minute record of your transaction logs. Therefore a seeding operation is required to resync both the active and passive databases.

2.4 Integrity Verification on a SnapMirror® Destination

SME 4.0 utilizes SnapDrive and FlexClone™ to allow you to perform an integrity verification on a SnapMirror destination volume. Using FlexClone, SME creates a clone of the SnapMirror destination volume. Creating a clone avoids having to break the SnapMirror relationship between the source and destination volumes. The writable cloned volume can be accessed by SME, and the LUN containing the database to be verified can be mounted at the destination. This provides the necessary access for SME to perform an integrity verification on the backup set. Once verified, the clone is destroyed, and based on the results of the verification job, SME reports back a success or failure to the production SME instance. SME then marks the backup set on the primary accordingly.

Integrity verifications on a SnapMirror destination volume can be managed in the SME interface by:

- The Backup and Verification tab or wizard
- Deferred verification
- The Test Restore button or Restore wizard
- The Restore tab

The requirements to take advantage of integrity verifications on a SnapMirror destination volume are:

- A SnapMirror license and an existing SnapMirror relationship in place
- A FlexClone license on the SnapMirror destination storage system
- The destination volumes must be accessible via SnapDrive

If these requirements are not met, SME displays an error message and the verification operation fails for that backup set.

Note:

- When you are using the Restore wizard, SME prompts you to perform an integrity verification on either the source volume or the destination volume if it detects that a SnapMirror relationship is configured.

2.5 Restore Backup Sets from a Different Exchange Server

SME 4.0 has the ability to restore backup sets created from a different Exchange server onto another Exchange server located in the same Active Directory domain that is running SME. During the restore process, you select the Exchange server that originally created the backup sets. This allows SME to properly restore the backup set and take the appropriate restore options.

During the restore process, you are given the following mount options:

- Recover and mount database after restore: This option simply mounts the databases after restore and brings them online. There are no additional processes on those mailboxes.
- Update user accounts associated with mailboxes in restored databases to point to mailbox server with new name: This option runs the Exchange 2007 `Move-Mailbox-ConfigureOnly` cmdlet to update the properties of the newly restored mailboxes so that they now reference the new restore target Exchange server.

Note:

- To use this new feature, the database, transaction log, and SnapInfo LUNs must be mounted with the same drive letters and/or volume mount point paths as the Exchange server where the backup set was created. In a SnapMirror environment, those LUNs must be replicated to the destination site.

2.6 Recovery Storage Group Integration

The Recovery Storage Group (RSG) functionality in Exchange server allows you to mount a second copy of a storage group and perform various manual recovery functions such as mailbox recovery. Mounting an RSG is manual process that can be time consuming. For Exchange Server 2007, SME 4.0 automates and accelerates these manual steps. SME takes the following steps:

- Creates an RSG
- Mounts a database into the RSG

- Adds databases to an existing RSG using SME Snapshot copies
- When all manual recovery operations are completed in the RSG, SME destroys the RSG

By default, SME destroys an existing RSG before creating a new one and mounting a database into that RSG.

Note:

- SME cannot restore a database into an existing RSG. If SME did not create the RSG, it cannot restore any database into that RSG.
- You cannot restore a public folder database into an RSG using SME. SME does not support restoring public folder databases.

In an MSCS environment, the LUNs that contain the database, transaction log, and storage group system files must be mounted to the node that owns the Exchange virtual server. These LUNs do not have the Affect the Group checkbox selected because they are connected to just one node to enable the RSG operation. If the Exchange virtual server is moved to another node, these resources go to a failed state until the Exchange virtual server is moved back to the node where they are connected.

In Exchange Server 2003, there is no integration into SME 4.0. The steps to create, mount, and destroy RSGs must be done manually using the Exchange System Manager.

NetApp offers a complementary solution called Single Mailbox Recovery (SMBR). This product is capable of restoring mailboxes and messages directly from SME Snapshot copies, without the need to use RSGs. For more information about SMBR, see <http://www.netapp.com/ftp/SMBR.pdf>.

3. Exchange and NetApp Storage Design

Microsoft Exchange servers require responsiveness from the underlying storage to support user activity, backup and restore operations, database maintenance, etc. The requirements for all of these actions can be distilled into two basic categories:

- Capacity
- I/O performance (measured in data flowing into and out of the file system per second, or IOPS)

Common examples that require capacity:

- The number of users multiplied by the mailbox quota size (in megabytes) provides a starting measurement for database size (for mailbox stores)
- Policy changes that result in an increase in mailbox quota size (it can be safely assumed that mailboxes will become full)
- Growth in the number of users, resulting in larger or additional mailbox stores
- Public folders
- Third-party Exchange add-on applications such as unified messaging servers, fax servers, and workflow applications that require additional storage to operate

Common examples that generate I/O load:

- User-related loads such as users sending and retrieving e-mails, attachments, and calendaring events, working on collaborative data in public folders, etc.
- Mobile devices that access e-mail, such as Blackberry devices and smart phones
- Server-specific events such as online database maintenance, indexing, etc.

- Public folder usage

3.1 Aggregates and Flexible Volumes

Data ONTAP® 7G provides new functionality that allows the creation of flexible volumes for managing data without the need to assign physical disks to the volumes. Instead, the FlexVol® volumes gain performance benefits from a larger pool of physical disks, called an *aggregate*.

This results in the following additional benefits for Microsoft Exchange environments:

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

Examples Using Traditional Versus Flexible Volumes

The volume structure shown in Figure 3 illustrates a scenario using traditional volumes that yields the greatest operational flexibility. Using traditional volumes, each disk is associated with a specific volume and is for the exclusive use of that volume only.

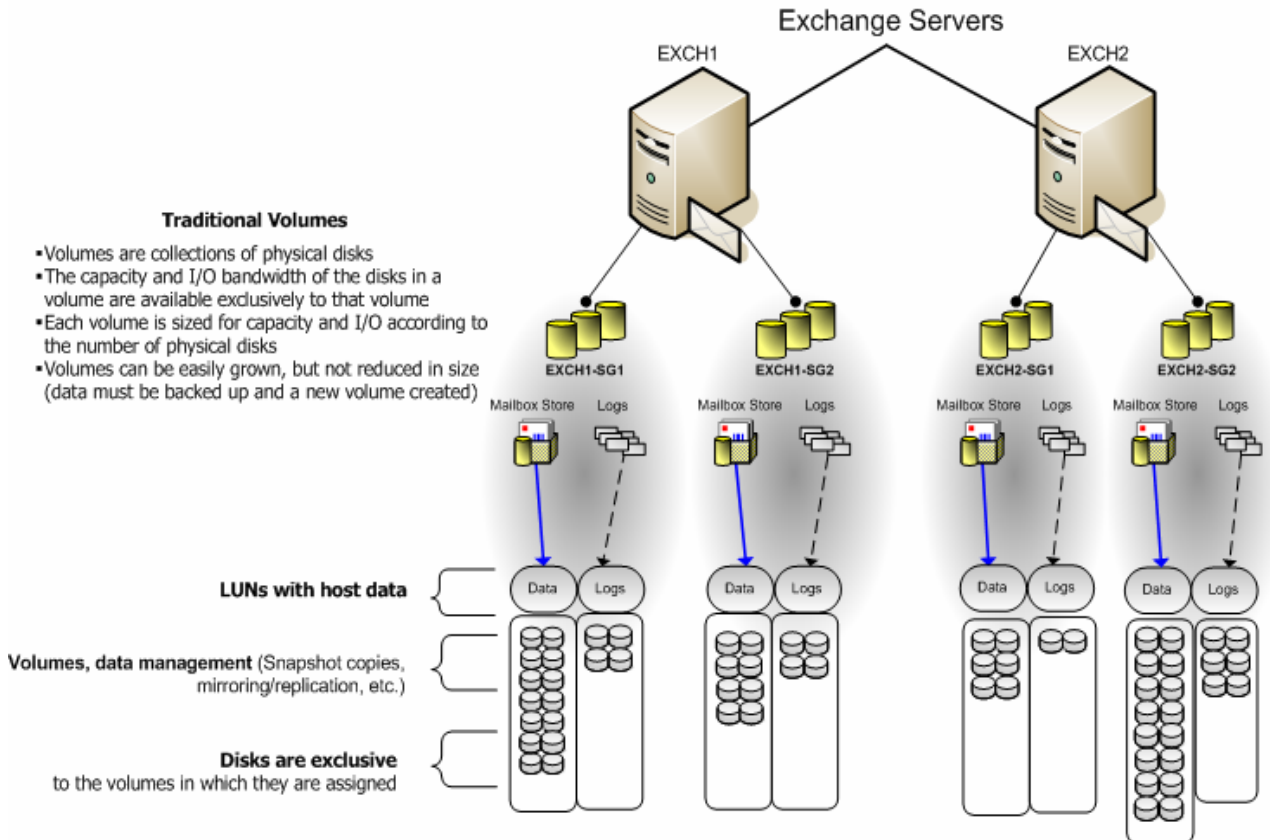


Figure 3) Using Data ONTAP 7G with traditional volumes.

Figure 4 shows a simplified example of how the same volume structure that existed previously could be used in a configuration with flexible volumes. This configuration yields the same operational flexibility for SME, but there are other significant differences and advantages. When using flexible volumes, the capacity and performance bandwidth of a large collection of fast disks can be made available to all volumes, no matter what size they are. Even very small volumes have the benefit of a very large numbers of disks. Volumes can be better leveraged for managing data and still realize the performance benefit of the aggregate's total spindle count.

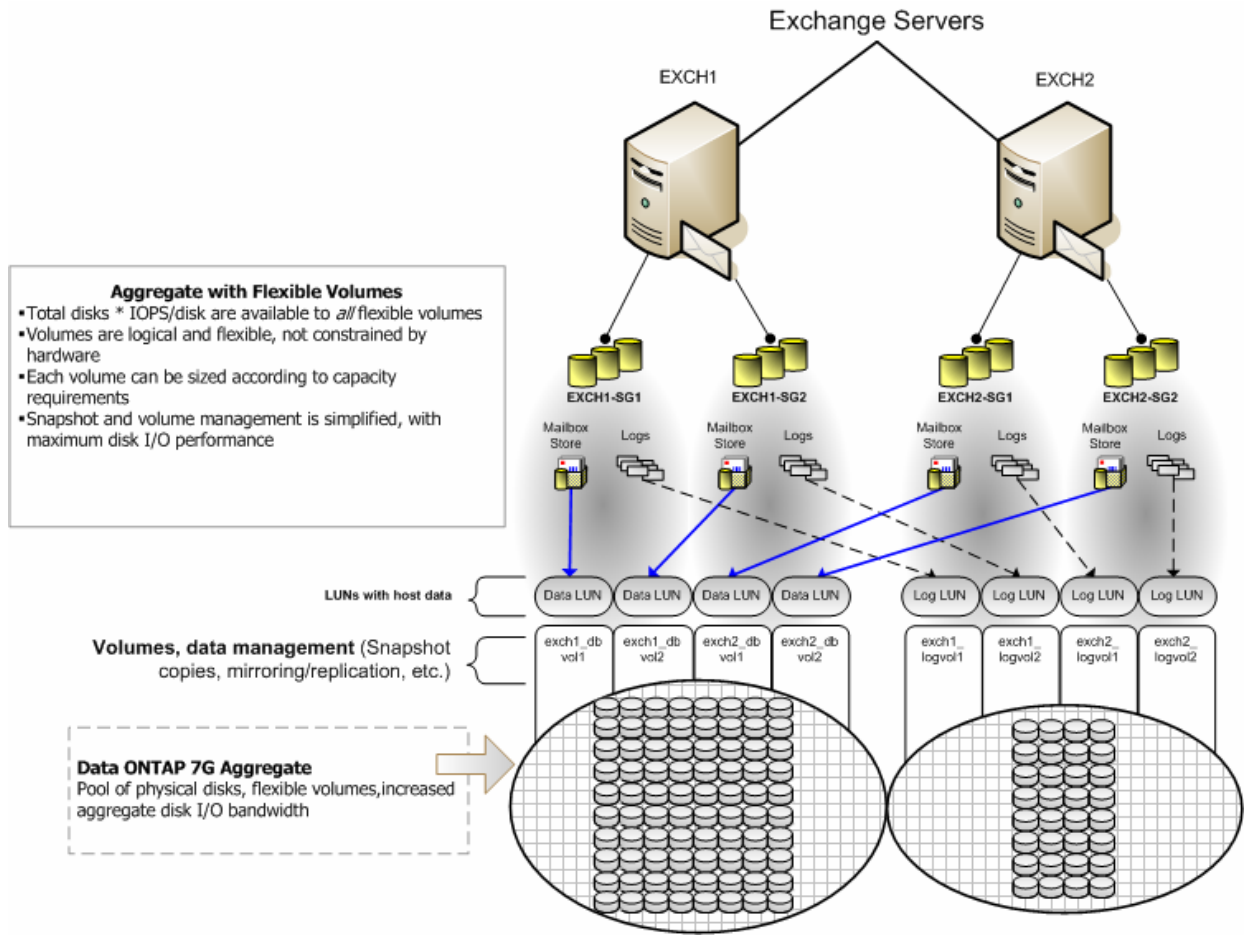


Figure 4) Using Data ONTAP 7G with flexible volumes.

Best Practice

When using aggregates, it is recommended that you separate the database flexible volumes and LUNs from the transaction log flexible volumes and LUNs. In the unlikely event that an aggregate is lost, Exchange data is still available for disaster recovery.

Best Practice

It is recommended that you have different database and transaction log volumes for different Exchange servers, to aid in avoiding a potential busy Snapshot copy problem. Because there are separate volumes for each server, you don't have to be as concerned with Snapshot schedules overlapping each other for different servers. For more information about busy Snapshot copy issues, see section 7.1 of this report, or see Chapter 12 of the *SME Installation and Administrator Guide*.

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

For accurate sizing and deployment of your Exchange storage, consult your NetApp Global Services representative (1-88-4NETAPP). They are the best resource to ensure a successful deployment that is sized to suit your storage needs.

3.2 LUN Sizing

Proper sizing of LUNs for Exchange databases and transaction logs is key to a successful Exchange deployment. If you size too small, you will obviously run out of space and Exchange will dismount that storage group. If you size too large, you are not utilizing your storage efficiently. Optimizing LUN size helps reduce wasted storage, increase backup and restore times, and, in a SnapMirror environment, decrease data replication times. The following calculations are for properly sized LUNs.

In the following calculations, the *logsize* variable differs between Exchange Server 2003 and Exchange Server 2007.

- Exchange Server 2003 = 5MB logsize
- Exchange Server 2007 = 1MB logsize

Database LUN Sizing

Formula:

Database size = ((number of mailboxes * mailbox quota) / single instance ratio) + deleted items cache

Example:

- 2,500 users
- 250MB mailbox size
- Single Instance Ratio of 1.2
- Deleted Items cache is 20%

Database Size = ((2,500 * 250) / 1.2) * 1.2 = 62,500MB or ~611GB

Note:

- In the preceding calculation, the addition of 20% of the database size (for the deleted items cache) is expressed by multiplying the estimated storage group size by a factor of 1.2 (120%).

Transaction Log LUN Sizing

Formula:

Transaction log LUN size = (number of transaction logs generated * logsize)

Example:

- 2,500 transaction logs between backups

Transaction log LUN size = (2,500 * 5MB) = 125,000MB or ~13GB

or

Transaction log LUN size = (2,500 * 1MB) = 2,500MB or ~2.4GB

SnapInfo Directory LUN Sizing

Formula:

SnapInfo LUN size = ((number of transaction logs generated * logsize) * (number of backups to keep online))

Example:

- 2,500 transaction logs between backups
- 7 backups kept online

SnapInfo LUN size = (2,500 * 5MB) * 7 = 87,500MB or ~86GB

or

SnapInfo LUN size = (2,500 * 1MB) * 7 = 17,500MB or ~17GB

Transaction Log and SnapInfo Directory LUN Sizing

Formula:

Transaction log and SnapInfo LUN size = ((number of transaction logs generated * logsize) * (number of backups to keep online + 1 for the active file system))

Example:

- 2,500 transactions logs between backups
- 7 backups kept online

Transaction log and SnapInfo LUN size = ((2,500 * 5MB) * (7 + 1)) = 100,000MB or ~98GB

or

Transaction log and SnapInfo LUN size = ((2,500 * 1MB) * (7 + 1)) = 20,000MB or ~20GB

All of these LUN sizings show the exact LUN size to support a particular Exchange profile. However, it is always best practice to add free space to a LUN for possible growth. This free space amount can be agreed upon by all parties involved in the sizing, then added to the overall LUN size. An appropriate number to increase the LUN size by would be approximately 15%, which should provide enough free space to handle file growth in a LUN.

3.3 Volume Sizing

There are many factors to consider when designing a Microsoft Exchange environment on a NetApp storage appliance. Important factors include volume sizing for Exchange data, transaction logs, and SME SnapInfo data. As stated earlier, best practice is to place the Exchange data volumes and transaction log/SnapInfo data on separate aggregates. This ensures that if an aggregate is lost, which is rare, a portion of your Exchange data is still available for disaster recovery.

Database Volume Sizing

Formula:

Minimum database volume size = ((2 * LUN size) + (number of online backups * data change percentage * max database size))

Example:

- Using the database LUN size from above + 15% for growth: 611GB + 15% = 703GB
- 10% data change between backups
- 7 backups kept online

Minimum database volume size = ((2 * 703) + (7 * 10% * 611GB)) = 1,834GB

Transaction Log Volume Sizing

Formula:

Transaction log volume size = ((transaction log LUN size in MB * 2) + (2 * number of transaction logs generated * logsize * number of backups to keep online))

Examples:

For Exchange 2003:

- Using the transaction log LUN size from above + 15% growth: 12,500MB + 15% = 14,375MB
- 2500 transaction logs between backups
- 7 backups kept online

Transaction log volume size = ((14,375MB * 2) + (2 * 2,500 * 5MB * 7)) = 203,750MB or ~199GB

For Exchange 2007:

- Using the transaction log LUN size from above + 15% growth: 2,500MB + 15% = 2,875MB
- 2500 transaction logs between backups
- 7 backups kept online

Transaction log volume size = ((2875MB * 2) + (2 * 2,500 * 1MB * 7)) = 40,750MB or ~40GB

3.4 Effects of Transaction Log Archiving Using NTFS Hard Links on Volume Size

NTFS hard links can be utilized in supported configurations where the transaction logs reside on the same NTFS volume and LUN as the SnapInfo directory. During a normal backup, SME copies the transaction log files to the appropriate SnapInfo directory for archiving. With NTFS hard links, and in supported configurations, when a backup is initiated, links are created to the transaction logs rather than copying the file itself.

Since fewer blocks are changed on the storage system volume, the formula to calculate the transaction log volume size has to be modified slightly. The difference now is that the number of transaction logs generated does not need to be multiplied by 2, as with the original formula.

Formula:

Transaction log volume size = (transaction log LUN size in MB * 2) + (number of transaction logs generated * logsize * number of backups to keep online)

Examples:

For Exchange Server 2003:

- Using the transaction log LUN size from above + 15% growth: 12,500MB + 15% = 14,375MB
- 2500 transaction logs between backups
- 7 backups kept online

Transaction log volume size = ((14,375MB * 2) + (2,500 * 5MB * 7)) = 116,250MB or ~114GB

For Exchange Server 2007:

- Using the transaction log LUN size from above + 15% growth: 2,500MB + 15% = 2,875MB
- 2500 transaction logs between backups
- 7 backups kept online

Transaction log volume size = ((2,875MB * 2) + (2,500 * 1MB * 7)) = 23,250MB or ~23GB

3.5 Fractional Space Reservation

Previous versions of SME enforced a 100% space reservation on volumes hosting LUNs. Fractional space reservation enables you to reserve less than 100% on volumes hosting LUNs, allowing for greater space utilization for those volumes. However, you do run a higher risk of unexpectedly running out of disk space on a volume, which in turn will halt any writes to a LUN. This means downtime for Exchange due to the inability to write any data to the affected LUNs.

Best practices have always been, and continue to be, to keep your volumes at 100% space reservations. This minimizes the risk of running out of free space on the volume and taking your Exchange storage groups offline.

Fractional space reservation affects two volume sizing formulas:

- Database volume size
- Transaction logs and the SnapInfo directory volume size

Effects of Fractional Space Reservation on Database Volume Sizing

Formula:

Database volume size = ((1 + fractional space reservation percentage) * LUN size) + (number of online backups * data change percentage * max database size)

Example:

- Using the database LUN size from above + 15% for growth: 611GB + 15% = 703GB
- 10% data change between backups
- 7 backups kept online
- Fractional space reservation set to 65%

Database volume size = $((1 + 0.65) * 703) + (7 * 10\% * 611\text{GB}) = 1,588\text{GB}$

Effects of Fractional Space Reservation on Transaction Log and SnapInfo Directory Volume Sizing

Formula:

Transaction log volume size = (transaction log LUN size in MB * (1 + fractional space reservation percentage)) + (2 * number of transaction logs generated * logsize * number of backups to keep online)

Examples:

For Exchange Server 2003:

- Using the transaction log LUN size from above + 15% growth: 12,500MB + 15% = 14,375MB
- 2,500 transaction logs between backups
- 7 backups kept online
- Fractional space reservation set to 65%

Transaction log LUN size = $((14,375\text{MB} * (1 + 0.65)) + (2 * 2,500 * 5\text{MB} * 7)) = 198,719\text{MB}$ or ~ 194GB

For Exchange 2007:

- Using the transaction log LUN size from above + 15% growth: 12,500MB + 15% = 2,875MB
- 2,500 transaction logs between backups
- 7 backups kept online
- Fractional space reservation set to 65%

Transaction log LUN size = $((2,875\text{MB} * (1 + 0.65)) + (2 * 2,500 * 1\text{MB} * 7)) = 39,744\text{MB}$ or ~39GB

If you plan to implement a fractional space reservation policy in conjunction with an Exchange solution, contact your NetApp Global Service representative (1-888-4NETAPP). They are experts in handling configurations like these and will help make this process as safe as possible.

4. SnapManager 4.0 for Exchange Installation

The steps for installing SME are detailed in the *SnapManager 4.0 for Exchange Installation and Administration Guide*. Those steps and the best practices listed in this guide are the best resources to deploy and manage your SME implementation.

Note:

- To install SME 4.0, MMC v3.0 and PowerShell 1.0 must be installed on the server.

4.1 SME Interactive Mode

Although the procedures to install SME in interactive mode are described in the *Installation and Administration Guide*, it is important to emphasize the following key points:

- Do not use terminal services when installing or administering SME. Important or critical pop-up boxes may appear only on the local system console and could be missed by a terminal services session. It is best practice to be at the local system console when installing or administering SME.
- When installing SME, there is no reason to stop any Exchange services. Exchange continues to run during an SME installation and after it is installed.
- You do not need a license when you first install SME. You can bypass this step during the initial installation; however, you will be required to provide either a server-side license or a storage-side license when you run SME.

- In the SnapManager Server Identity window, it is best practice to use the prepopulated account. This account is the same as the SnapDrive service account. All necessary privileges have been given to this account, and therefore it is the best to use for SME as well.

4.2 SME Unattended Mode

You can deploy SME in unattended mode, or CLI (command line interface) mode. This mode allows you to script the installation of SME on any server that requires an installation. By setting appropriate parameters, you can specify fields like the following:

- Username and Password
- End User License
- Serial Number
- Installation Directory
- Reboot options

Here is an example script for the unattended installation:

```
C:\installation\SME4.0_x64.exe /v'AGREETOLICENSE=Yes SVCUSERNAME=corp\SMEAdmin  
SVCUSERPASSWORD=SMEpass SVCCONFIRMUSERPASSWORD=SMEpass  
ISX_SERIALNUM=ABCDEFGHIJKLMN /qb
```

This example installs SME4.0. It automatically agrees to the End User License prompt, sets the user account to `corp\SMEAdmin` with the password `SMEpass`, provides a server-side license of `ABCDEFGHIJKLMN`, and completes the installation. If a reboot is required, the unattended installation automatically reboots the server. To prevent the reboot of the Exchange server, use the `REBOOT=0` parameter.

5.0 SnapManager 4.0 User Interface

SME 4.0 is now a snap-in-based application that leverages the Microsoft Management Console 3.0 framework. This new framework allows SME 4.0 to be added seamlessly into the same management console as the new Microsoft Exchange Management Console for Exchange 2007. SME 4.0 also integrates with the new Exchange 2007 PowerShell 1.0 interface. This integration allows easy scripting of SME functionality.

5.1 SME 4.0 Dashboard View

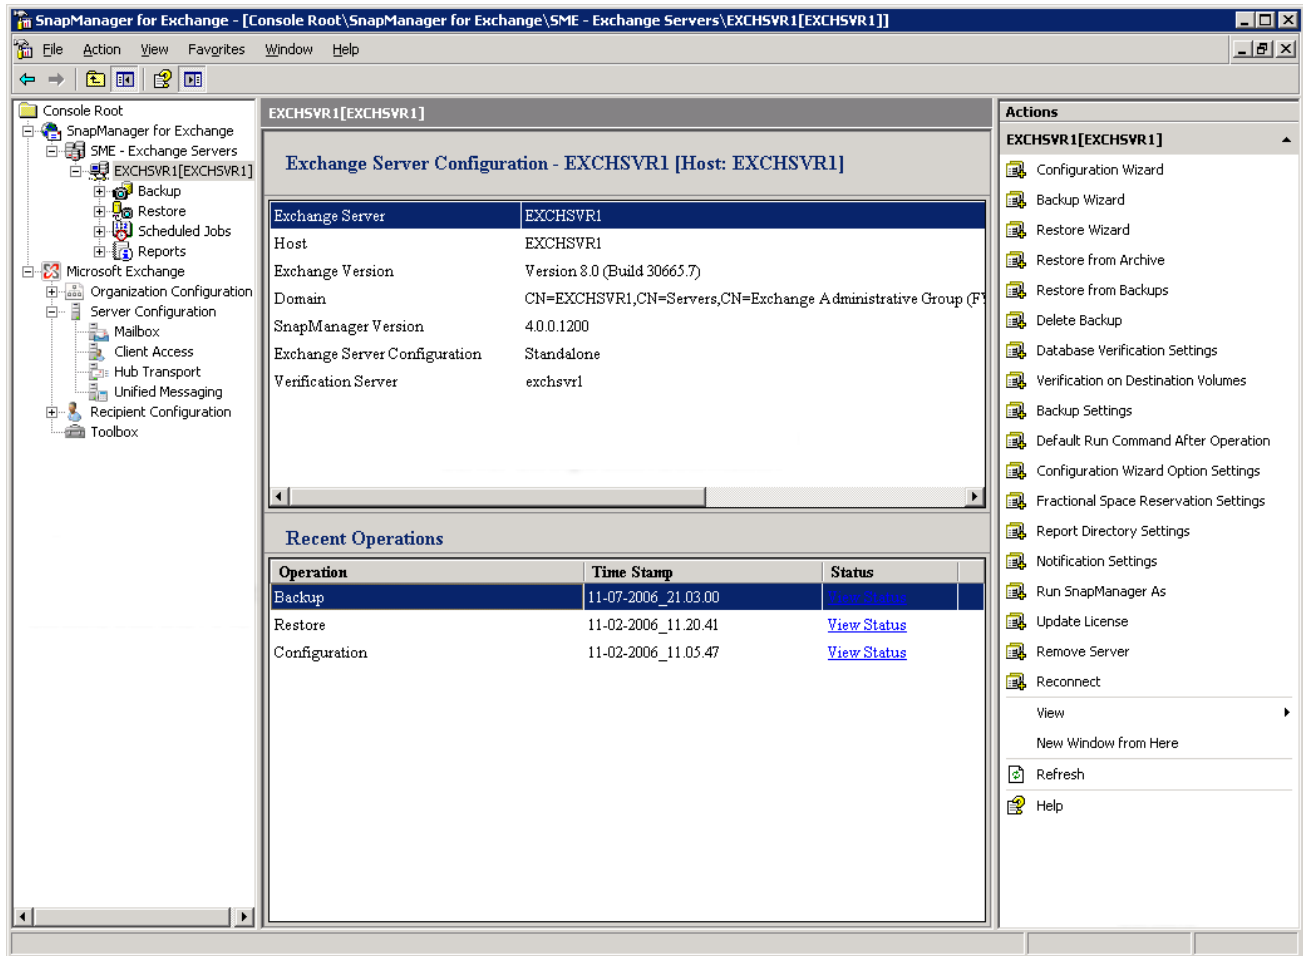
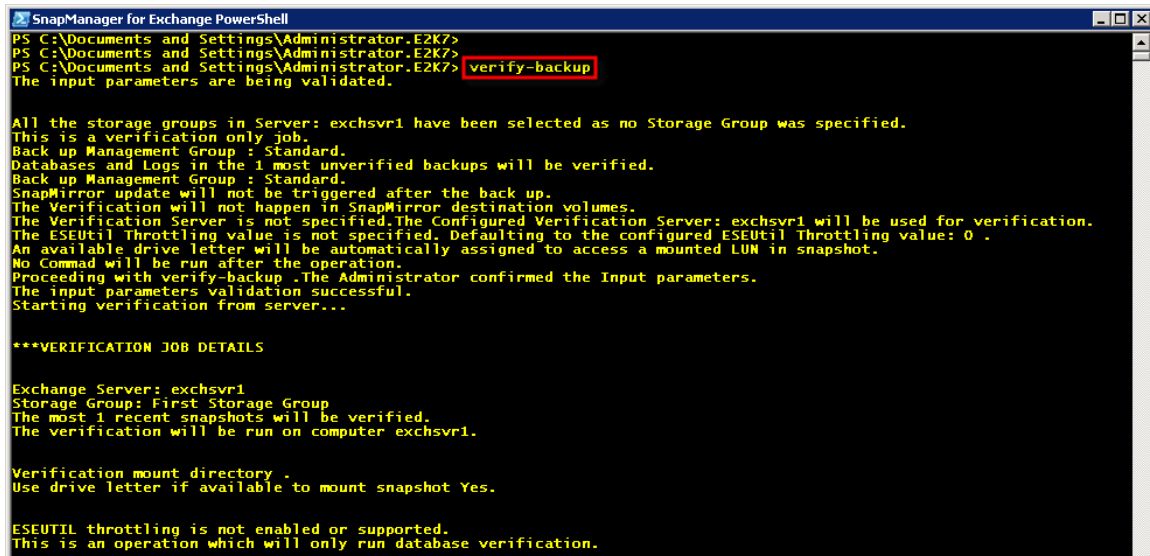


Figure 5) The new SME 4.0 UI with Dashboard view.

The Dashboard view is composed of three main sections:

- **Scope pane:** On the left side, the Scope pane lists the SME 4.0 snap-in. By expanding the snap-in, you can see any servers that SME is currently managing. You can also add the Exchange Management Console snap-in and have it show up in the snap-in list. This enables you to have one management console that manages your entire Exchange environment. Expanding the SnapManager for Exchange snap-in reveals a tree list of common operation tasks.
- **Overview/Results pane:** The center of the SME snap-in gives an overview of the current Exchange server and a list of recent operations. This view gives you quick access to pertinent information about the Exchange server and recently completed SME tasks.
- **Actions pane:** The right side shows the Actions pane, listing all the tasks that SME 4.0 can execute. This organized list gives you fast and easy access to SME tasks that you want to perform. It also lists the settings and options that you can use to customize SME to your environment.

5.2 SME 4.0 PowerShell Interface



```
SnapManager for Exchange PowerShell
PS C:\Documents and Settings\Administrator.E2K7>
PS C:\Documents and Settings\Administrator.E2K7>
PS C:\Documents and Settings\Administrator.E2K7> verify-backup
The input parameters are being validated.

All the storage groups in Server: exchsvr1 have been selected as no Storage Group was specified.
This is a verification only job.
Back up Management Group : Standard.
Databases and Logs in the 1 most unverified backups will be verified.
Back up Management Group : Standard.
SnapMirror update will not be triggered after the back up.
The Verification will not happen in SnapMirror destination volumes.
The Verification Server is not specified.The Configured Verification Server: exchsvr1 will be used for verification.
The ESEutil Throttling value is not specified. Defaulting to the configured ESEutil Throttling value: 0 .
An available drive letter will be automatically assigned to access a mounted LUN in snapshot.
No Command will be run after the operation.
Proceeding with verify-backup .The Administrator confirmed the Input parameters.
The input parameters validation successful.
Starting verification from server...

***VERIFICATION JOB DETAILS

Exchange Server: exchsvr1
Storage Group: First Storage Group
The most 1 recent snapshots will be verified.
The verification will be run on computer exchsvr1.

Verification mount directory .
Use drive letter if available to mount snapshot Yes.

ESEUTIL throttling is not enabled or supported.
This is an operation which will only run database verification.
```

Figure 6) SME 4.0 PowerShell interface.

The new SME 4.0 PowerShell interface makes scripting SME actions easier and more integrated into Exchange Server 2007. For a complete list of SME cmdlets for PowerShell, see the *SME Installation and Administrator Guide*.

6.0 Configuration

Careful planning of your Exchange data layout is an integral part of a successful Exchange deployment.

6.1 Configuration Wizard Rules

The following rules are enforced by the Configuration Wizard when migrating your Exchange data:

- Exchange databases:
 - Cannot share a LUN with databases from different storage groups
 - Cannot share the same LUN containing transaction log files

Note: All databases on a LUN are restored together. In order to restore an individual database, it must reside on its own LUN.
- Exchange transaction log files:
 - Transaction log files from different storage groups can share the same LUN
- SnapInfo directory:
 - The SnapInfo directory can reside on its own LUN or can share a LUN with transaction log files
 - When using NTFS hard links, the SnapInfo directory cannot share a LUN with Exchange database files
- SAN boot LUN:
 - You cannot place database files, transaction log files, or a SnapInfo directory on a SAN boot LUN
- Exchange storage group system files:
 - The Exchange storage group system files must be on a LUN containing the transaction log files for that storage group

Best Practice

Place your SMTP and MTA queues on any LUN or disk that is not hosting any Exchange database files. For standalone systems, a dedicated LUN or the default installation location is typically fine. In an MSCS

environment, the queues must be located on a shared LUN that does not contain any Exchange database files. You can also place the SMTP and MTA queues on a transaction log or on the SnapInfo LUN.

Best Practice

When installing Exchange and SME in an MSCS environment, it is best practice to have at least one passive MSCS node available to the cluster. This Microsoft best practice ensures that if one node fails, there is another node ready to take over for the failing node. This minimizes any performance degradation that may occur in active/active configurations. Without the passive node available in the cluster, an active node is selected to take over duties for the failing node. In most cases, the increased load on that active node causes Exchange performance to decrease.

6.2 Volume Mount Point Support

SME 4.0 is tightly integrated with volume mount point support in NetApp SnapDrive. Volume mount point support allows SnapManager for Exchange to mount up to 120 NetApp LUNs per server, even if there are a limited number of drive letters available. This functionality is especially beneficial to Microsoft Clustering Services (MSCS), where shared drive letter availability can be a problem.

With mounted volume support, you don't have to worry about having enough available drive letters for things like:

- Maximum layout configuration where the database layout is configured for individual database restore
- Transaction log and SnapInfo information
- Remote verification of more than one backup set

Volume mount point support gives you more flexibility when dealing with restore and verification of backup sets.

Migrating from Drive Letters to Volume Mount Points

Existing SME configurations may have Exchange data LUNs mapped to drive letters. You can easily migrate those mapped LUNs to volume mount points without having to copy data from one LUN to another. The affected storage group is taken offline briefly to remap the paths. This process saves time by not having to copy physical data and minimizing the time that the storage groups are offline.

Note that not all drive letters are eliminated. A minimum of one drive letter remains mapped to a LUN that serves as the volume mount point root. As a best practice, this LUN should serve as a transaction log LUN. For environments with more than one storage group, consider using one LUN per storage group mapped to a drive letter that serves as the transaction log LUN and the volume mount point root.

In MSCS environments, the entire Exchange virtual server is taken offline during the volume mount point migration. The System Attendant cluster resource depends on any Exchange physical disk resources. When those resources are taken offline, the System Attendant is taken offline as well. This also takes the remaining Exchange virtual server resources offline.

When migrating from drive letters to volume mount points, you lose the ability to restore those LUNs from previously taken backups. It is best practice to take a full backup of all Exchange databases both before and after a successful migration. When the migration is successful, you can discard the earlier backups. You cannot restore your Exchange databases from these backups. Consult your NetApp Global Services representative (1-888-4NETAPP) to assist in planning for a migration.

Detailed steps for volume mount point migrations are found in the *SME 4.0 Installation and Administration Guide*, Chapter 8: "Migrating an existing SnapManager for Exchange configuration from using drive letters to using mount points".

Restrictions

You cannot store Exchange database files on a mount point root volume, only in transaction log files.

Best Practice

Make the transaction log LUN (or LUNs) your volume mount point root because it will be backed up on a regular basis. This ensures that your volume mount points are preserved in a backup set and can be restored if necessary.

Also, if Exchange databases reside on a LUN, do not add mount points to that LUN. If you have to complete a restore of a database residing on a LUN with volume mount points, the restore operation removes any mount points that were created after the backup, disrupting access to the data on the mounted volumes referenced by these volume mount points. This is true only of mount points that exist on volumes that hold database files.

The following figure shows how volume mount points are displayed in SnapDrive and in Windows Explorer.

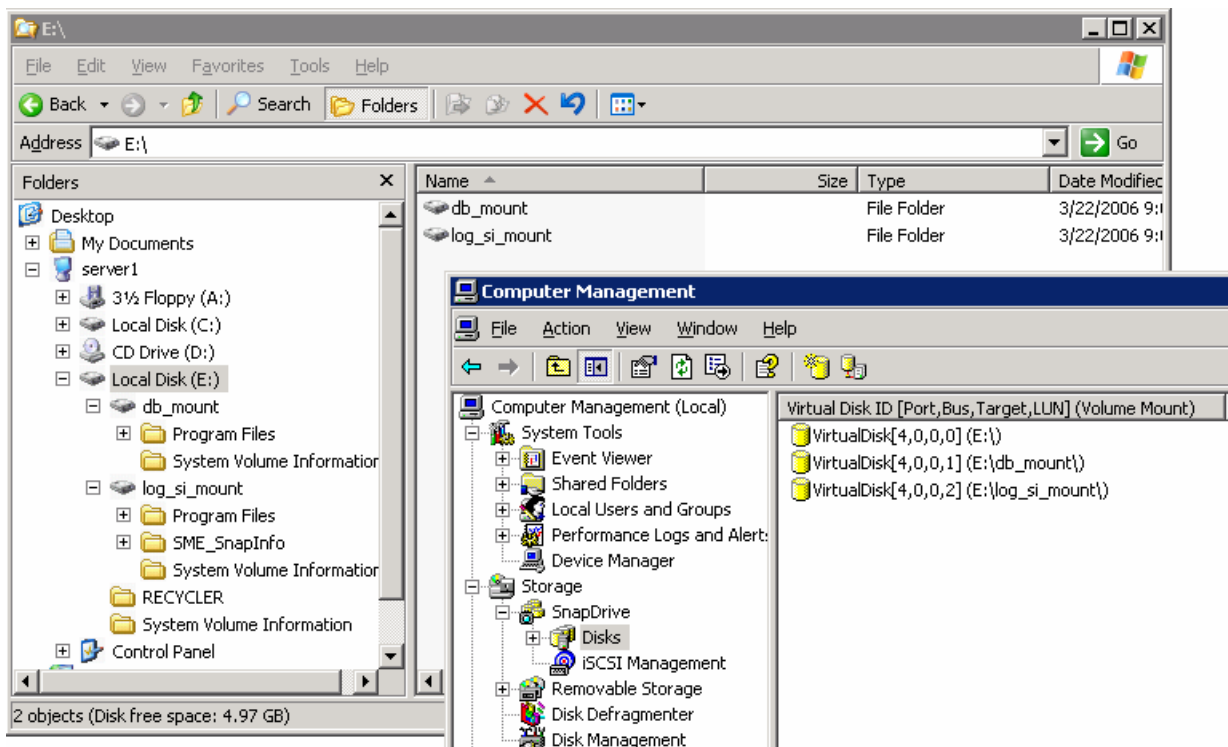


Figure 7) Volume mount points in SnapDrive and Windows Explorer.

6.3 Fractional Space Reservation Policies

SME 4.0 takes advantage of fractional space reservation capabilities in Data ONTAP 7.1 and later. These policies allow SME to take appropriate action when available space on a volume becomes limited. These configurable options give greater flexibility to administrators who are monitoring storage utilization very closely.

Best Practice

Set your fractional space reservation to 100%. When a LUN is fully space reserved (fractional space reservation set to 100%), write operations to that LUN are guaranteed against failure caused by an out-of-space condition due to Snapshot copy disk space consumption. If you do require a fractional space reservation policy, work with your NetApp support representative. The representative can help implement a fractional space reservation policy that fits your environment.

Fractional space reservation is set to 100% (disabled) by default. In order to leverage fractional space reservation policies, you need to set the fractional reserve parameter to less than 100%.

SME provides its own space management tool for monitoring overwrite reserve utilization on fractionally space-reserved volumes. SME takes appropriate action to prevent a LUN from becoming inaccessible as a result of no free space being available on the hosting volume. SME automatically either deletes Exchange backup sets or dismounts Exchange databases when the appropriate, user configurable, trigger is hit.

Automatic Deletion of Exchange Backup Sets

SME can be configured to automatically delete Exchange backup sets when the fractional space reservation meets a threshold specified in the Options menu. SME deletes the Exchange backup sets as follows:

- Deletes the oldest backup Snapshot copies first
- Retains the specified number of total backup Snapshot copies on the volume
- Retains the most recent backup of any database (if it resides on that volume)
- Retains the last backup of any database that no longer exists

Best Practice

Choose a backup retention level based on your SME backup creation and verification schedule. If a Snapshot copy deletion occurs, you should ensure that a minimum of one verified backup remains on the volume. Otherwise, you run a higher risk of not having a usable backup set to restore from in case of a disaster.

SME executes defined deletion policies on a per-volume basis. If a backup set is spanned across multiple volumes, you must have identical deletion policies across those volumes. If this practice is not followed, you may end up with mismatching backup sets on different volumes. This may cause a backup set to be rendered useless and unable to be restored.

Automatic Dismounting of Exchange Databases

SME can be configured to automatically dismount Exchange databases when the corresponding threshold is reached. This option is configurable in the fractional space reservation policy. This action stops Exchange from writing more data to that LUN in a volume where the overwrite reserve space approaches critical levels. This policy is used as a last-resort action to prevent further consumption of the overwrite reserve space.

When both Automatic Deletion of Exchange Backup Set” and Automatic Dismounting of Exchange Databases are enabled, the automatic dismounting policy must be set at a higher threshold than the automatic deletion policy. This ensures that SME first uses automatic backup deletion to regain space in the overwrite reserve. If this does not relieve the space issues, SME resorts to dismounting the affected databases to prevent further consumption of the overwrite reserve space. You should set the threshold level low enough to give Exchange enough time to write outstanding transactions to the database and dismount the database before running out of disk space. This is important because if Exchange encounters a disk full error, the affected databases are taken offline. Depending on the state of the shutdown, you may be required to run Eseutil or perform a restore operation to recover the affected databases.

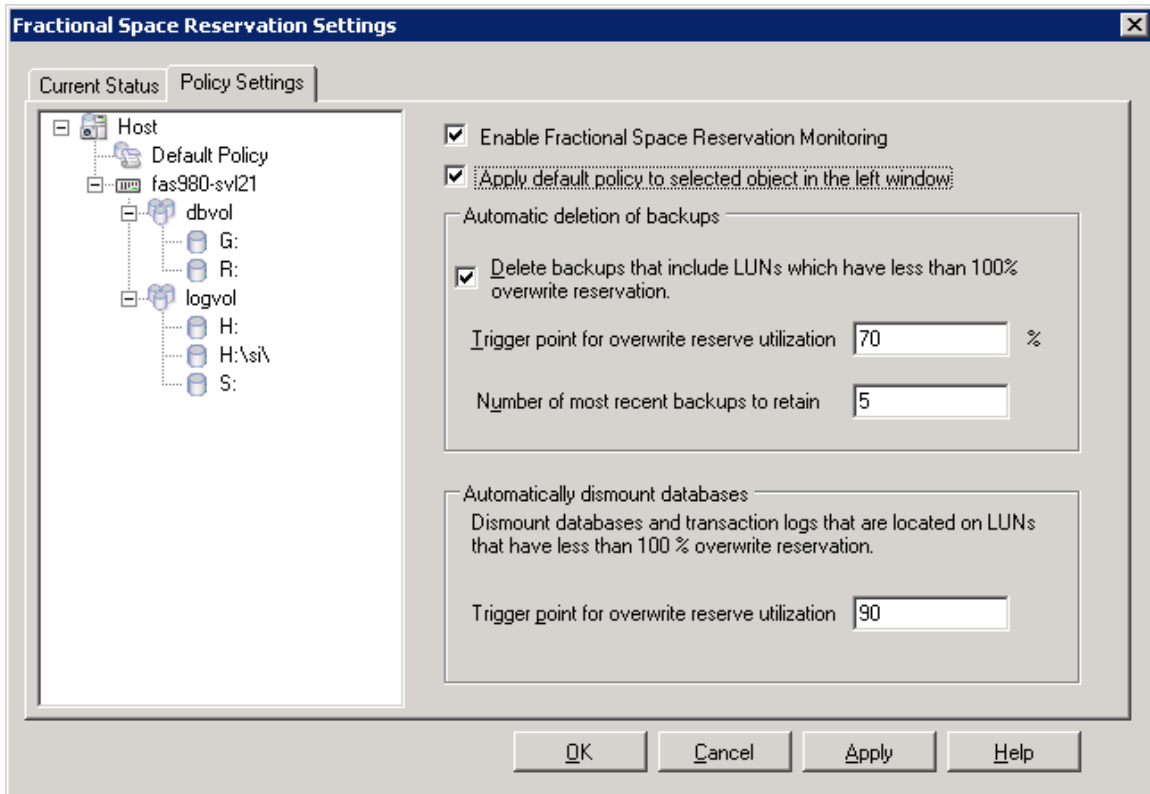


Figure 8) Fractional space reservation settings.

Figure 8 shows the trigger point for automatic deletion set to 70%. When the fractional overwrite reserve is 70% consumed, SME automatically deletes backup sets, keeping five of the most recent backup sets. The automatic dismount database trigger point is set to 90%. If automatically deleting backups cannot free enough space in the fractional overwrite reserve, SME dismounts the databases when the reserve is 90% consumed. This prevents an out-of-space condition on LUNs hosting Exchange databases and transaction log files.

6.4 Copyless Transaction Log Archiving

When your transaction log files and SnapInfo directory reside on different NTFS volumes, SME physically copies the transaction log files for a backup set to the corresponding SnapInfo directory during a backup operation. Once the copy is complete, the transaction log files are truncated and the backup operation completes.

When archiving Exchange transaction logs that reside on the same NTFS volume as the SnapInfo directory, SME 4.0 uses NTFS hard links instead of copying the transaction logs. This technology maximizes storage space and increases backup performance during the transaction log archival phase.

SME archives transaction log files to the SnapInfo directory (on the same NTFS volume) by creating NTFS hard links to the live transaction log file, without physically copying the files. When an SME backup operation completes, Exchange truncates its committed transaction log files. However, the actual files are not deleted. NTFS detects that other hard links still exist and the physical files are retained. The original NTFS hard links to the live transaction files are removed, and the links to the SnapInfo directory remain intact and available for access for SME restore operations.

To use this optimized functionality, your transaction log files and your SnapInfo directory must reside on the same NTFS volume on the same LUN. SME automatically uses NTFS hard links in this configuration.

However, in most SME environments, there are multiple LUNs for transaction log files. To take advantage of the NTFS hard link functionality requires separate SnapInfo directories for each Exchange storage group, with each SnapInfo directory residing on the same NTFS volume and LUN as the transaction log files for the storage group.

Multiple SnapInfo directories are a configurable option with SME. There are two main benefits to having multiple SnapInfo directories:

- Facilitates the use of NTFS hard links within SME. Most SME environments have multiple LUNs for transaction log files. To use NTFS hard links, the SnapInfo directory must reside on the same LUN as the transaction log file. To achieve this, multiple SnapInfo directories are needed. One transaction log LUN per storage group containing the SnapInfo directory is the best way to take advantage of the performance benefits of NTFS hard links.
- The loss of backup sets is minimized when using multiple SnapInfo directories. If you do lose a SnapInfo directory, you have other directories to fall back on in case of emergency.

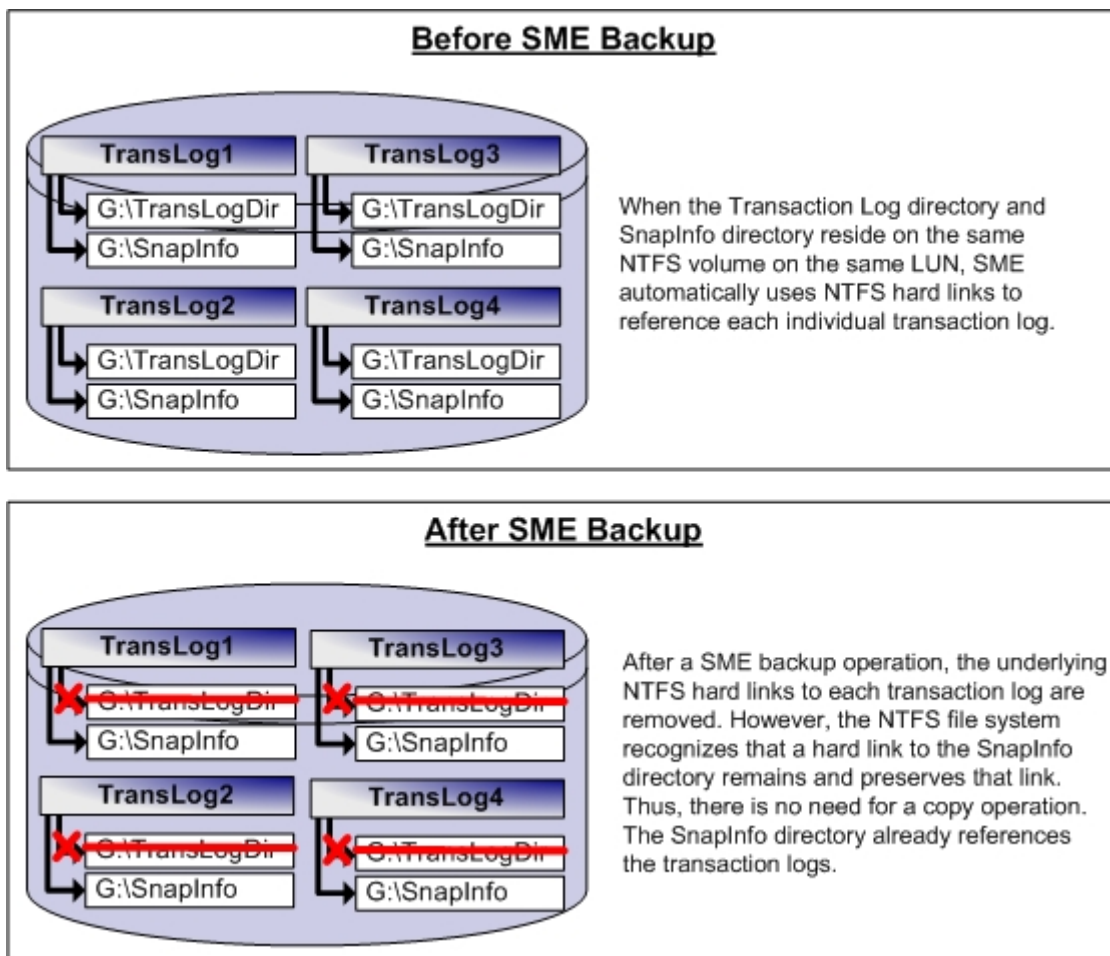


Figure 9) Copyless transaction log archiving.

7. Backup and Restore

7.1 Backup

SME performs backups of Exchange data in storage group sets. A storage group set is composed of all databases and transaction logs for that storage group. These databases can be distributed over multiple

LUNs and volumes, or they can reside on one LUN on a single volume. To implement a successful backup strategy, careful layout of storage volumes and LUNs is of paramount importance.

Volume-Wide Backup

When SME initiates a backup operation, a Snapshot copy is taken for the entire volume that hosts the LUN containing the Exchange data files. That backup is valid only for the Exchange server that performed the backup. Data from other Exchange servers that might reside on LUNs on that volume is not restorable from the Snapshot copy. Necessary operations to prepare Exchange for backup were not performed on the other Exchange servers and therefore the backups of those LUNs cannot be used for restore.

Multiple-Volume Backups

SME performs backups in parallel on all LUNs that belong to a storage group set. When a storage group set includes LUNs that span multiple volumes, the resulting backup set contains multiple Snapshot copies. All Snapshot copies are restorable as single entities.

Partial Backups

If the backup fails for some of the volumes in a storage group set, the other volumes that were successfully backed up are not affected and can be restored, even though all volumes were part of the same backup operation.

The number of LUNs per Exchange storage group is another key consideration that you must take into account. Depending on an Exchange environment, there may be individual LUNs per storage group database, or there may be a single LUN for all databases within a storage group.

Individual LUNs for each database in a storage group provide a high level of granularity when performing backup and restore operations. Restoring LUNs that are hosting individual databases can yield moderate performance increases because of fewer logs to replay after a restore operation. However, managing a configuration like this can become complicated. In Exchange Server 2007, you can have up to 50 storage groups and 50 databases. Placing each on individual LUNs results in a complicated environment with many LUNs and many Exchange objects to manage. Plan your Exchange layout carefully to meet your SLA requirements and for ease of administration. Consult your NetApp Global Services representative (1-888-4NETAPP) to assist in planning your Exchange layout.

Option to Back Up Transaction Logs That Exchange Truncates

As with previous versions of SME, by default all databases and transaction logs associated with the storage group are backed up. This process provides up-to-the-minute restore capability from any backup. If your environment does not require up-to-the-minute restore capability, you can now choose to skip the backup of the transaction logs that Exchange truncates. This option reduces the amount of disk space that is consumed by backing up those transaction logs. However, you lose the ability to restore up to the minute.

Note:

- If you deselect this option for a backup, the next time that you start SME, it is automatically reselected. If you want to perform another backup without backing up the transaction logs that Exchange truncates, you must deselect this option again before performing the backup.

Best Practice

It is best practice to leave the default option to back up transaction logs that Exchange truncates selected. This ensures that you have complete restore capability from backups without losing any Exchange data.

Scheduling

To minimize database latency spikes, it's best to perform an SME backup operation with verification of that backup during off peak hours—typically between 6:00 p.m. and 7:00 a.m. This minimizes the load on the Exchange server and therefore keeps latency times within the Microsoft suggested range. Other important factors that you must consider when performing verification operations are Exchange database maintenance, offline database defragmentation, and any other operations that can affect database latency times.

Also consider the following SME-specific recommendations for scheduling backups:

- Do not overlap any SME operations. SME does not support concurrent operations on the same server at the same time.
- Do not schedule a backup to occur while verifying a backup, even if the verification is performed on a remote server. This can cause problems deleting backups due to a busy Snapshot copy error.
- Do not schedule verifications during peak Exchange server usage hours. This causes performance degradation on the Exchange server.
- When possible, take advantage of performing backup verifications against a SnapMirror destination volume. For more information, see section 2.4, “Integrity Verification on a SnapMirror Destination.”

Scheduling Backup Operations in an MSCS

SME 4.0 scheduled backups on all nodes of an MSCS. Scheduling jobs in this manner ensures that if a node fails, your scheduled backups do not have to be re-scheduled on another node.

Expanding a LUN and Its Effect on Backups

If you expand a LUN through SnapDrive, you lose the ability to restore that LUN and any Exchange data on it through SME. After you expand your LUNs, take a full backup of the affected Exchange storage groups to ensure that you have a consistent backup from which you can restore if necessary. Consult your NetApp Global Services representative to aid in planning for a LUN expansion.

Problems Deleting Backups Due to Busy Snapshot Copy Error

This issue arises when you take a Snapshot copy of a LUN that is already backed by another Snapshot copy.

For example:

1. You take a Snapshot copy of drive E: (LUN1) and call it SNAP1.
2. Through SnapDrive, you mount LUN1 from SNAP1 to G:.
3. You take a Snapshot copy of the volume by any means (through SME, SnapDrive, the storage console, or FilerView®).
4. At this point, SNAP1 shows as (busy, LUNs) on the local NetApp storage console or in FilerView.

If any scheduled backups occur while the LUN backed by a Snapshot copy is still mounted, you are not able to delete that backup. To resolve this problem, you need to delete the more recently taken backup before the older backup can be deleted.

To avoid this situation altogether, make sure that you do not take backups while you have any LUNs backed by Snapshot copies mounted. Common situations where a LUN backed by a Snapshot copy may be mounted include the following:

- During a verification on the primary copy
- Archiving from a LUN backed by a Snapshot copy
- Performing a Single Mailbox Restore using a LUN backed by a Snapshot copy

Best Practice

When you have recovered from a busy Snapshot copy error, take a full backup of your Exchange data. This ensures that you have full backup of your Exchange data from which to recover in case of disaster.

7.2 Verification

It is best practice to have at least one full database backup that is verified every 24 hours. In typical deployments, four Snapshot copies are taken per day and a total of seven copies are kept online and verified. Depending on your recovery point objective (RPO), you can adjust how many backups you need to take, and how many need to be verified.

When you perform more frequent backups, fewer Exchange transaction logs must be replayed into the database at restore time. This can result in faster restore times. However, you still must perform verifications

before that backup can be restored. More frequent database verifications increase the amount of load on an Exchange server, and may increase latency times.

For Exchange Server 2003 configurations, eseutil throttling is an option to help reduce the load that a database verification operation places on a server. This helps keep latency times from increasing beyond recommended ranges.

For Exchange Server 2007, SME supports the new ChkSgFiles functionality. Similar to eseutil, ChkSgFiles verifies the integrity of an Exchange database and can be throttled to help reduce load on the Exchange server.

Eseutil and ChkSgFiles verify each database checksum and transaction log checksum in 512K increments. The calculation for the maximum throughput rate when using the throttling option is:

Maximum throughput rate = 512KB per I/O * number I/Os between 1-second pauses

To decrease the maximum throughput rate, decrease the number of I/Os between 1-second pauses. The following table shows an example.

Maximum possible database verification speed		
Number of I/O operations between 1-second pauses	Calculation	Maximum Speed
100	512KB per I/O * 100 IO per sec = 51,200KB/sec	50MB/sec
150	512KB per I/O * 150 IO per sec = 76,800KB/sec	75MB/sec
200	512KB per I/O * 200 IO per sec = 102,400KB/sec	100MB/sec
250	51 KB per I/O * 250 IO per sec = 128,000KB/sec	125MB/sec

Note:

- Setting any throttling option increases the duration of database verifications. If SME is configured to perform database verification as part of the backup process, it takes longer to complete the backup process. However, if SME is configured to not perform database verification as part of the backup process, backup times are not affected.

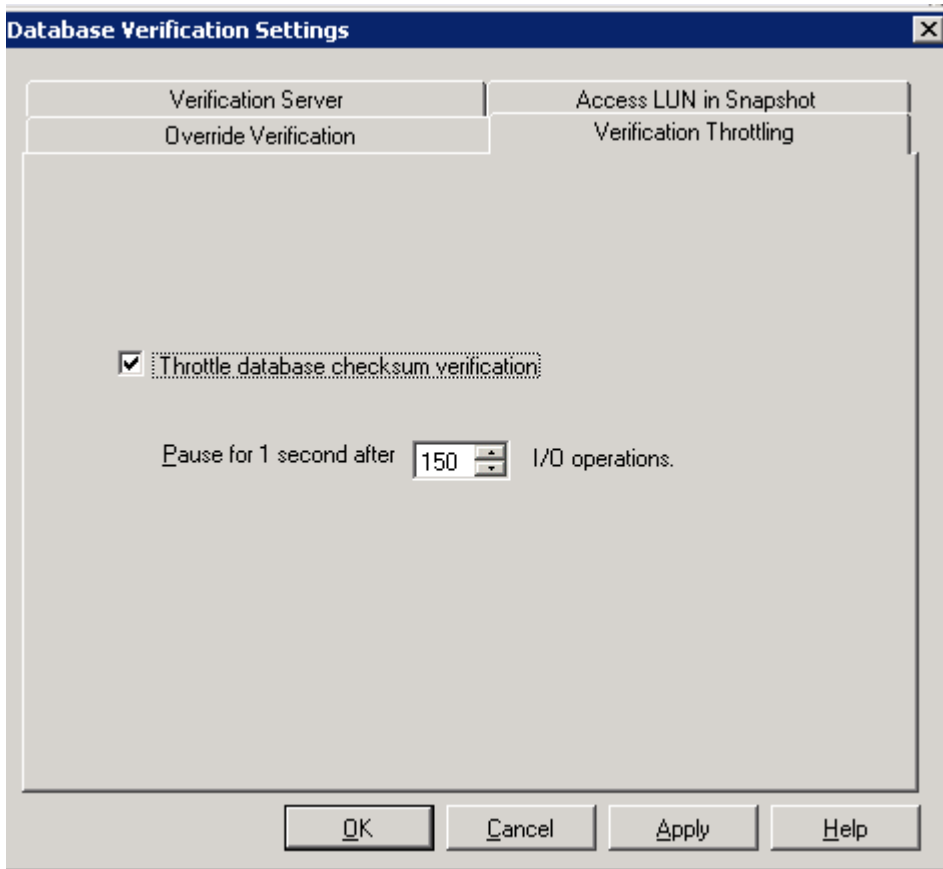


Figure 10) Database verification throttle settings.

The Exchange Server 2007 high-availability features, LCR and CCR, also help reduce the load of a backup verification in the production Exchange environment. By moving the load off of the production spindles and server (CCR only), you have greater flexibility to schedule a backup verification. For more information on LCR and CCR, see section 2.3, “Exchange Server 2007 High-Availability Features.”

If you have a SnapMirror implementation for your Exchange environment, SME 4.0 can now offload the verification process onto the SnapMirror destination volume. SME can connect to the SnapMirror destination and, using FlexClone, mount the LUN containing the backup set to be verified and perform the verification. After verification, SME marks the backup set as verified on the production Exchange server. For more information on database verification on a SnapMirror destination, see section 2.4, “Integrity Verification on a SnapMirror Destination”.

7.3 Restore

The ability to recover Exchange databases when necessary is a critical operation for an Exchange administrator. With SME Restore you can recover your Exchange databases and transaction logs from backups that it created, or from archive. SME supports restoring an entire storage group or an individual database.

There are two types of restores in SME:

- **Up-to-the-minute restore:** Selected by default, an up-to-the-minute restore replays any necessary and available transaction logs from the backup set and transaction log directory and applies them

to the database. A contiguous set of transaction logs is required for an up-to-the-minute restore to succeed.

- **Point-in-time restore:** With this option, only the transaction logs that existed in the active file system when the backup was taken are replayed and applied to the database. All transaction logs past the particular point in time chosen are discarded. You restore your Exchange data to a chosen point in time. Any Exchange data past that point is not restored. This option is particularly useful when you need to restore to a point before something like data corruption occurred.

Best Practice

When performing an up-to-the-minute restore, use the most recent verified backup, because this is the fastest way to restore the Exchange server. Using an older verified backup slows the restore time because there are often more transaction logs to be replayed and applied to the database. Using the most recent verified backup ensures a quick recovery of your Exchange database.

Renamed Storage Group

If the storage group name that you need to restore has changed since the backup was taken, you cannot recover it from that backup. If you do change the storage group name, be sure to perform a full backup immediately to protect your environment. If you do need to restore from a backup, and the storage group name is different, you must rename the storage group back to the name that matches the storage group name in the backup.

8. SnapMirror

SME takes advantage of asynchronous SnapMirror through SnapDrive. When SME performs a backup of Exchange data residing in a SnapMirror source volume, it can optionally request SnapDrive to initiate a SnapMirror update after the SME backup process has completed.

Note:

- The SME snapshot backup set is being mirrored, not the live Exchange data.

An update SnapMirror command is issued only after the SME backup operation completes, if that option is set. This may not be a feasible update schedule for protecting your valuable Exchange data. You can increase the number of backups that are taken during the day; however, you run a higher risk of overlapping SME operations, increasing load on the Exchange server, and potentially increasing latency times. To avoid these problems, you can take full advantage of the SnapDrive rolling Snapshot copy capability described later in this section.

When using SnapDrive to update SnapMirror, special Snapshot copies called *rolling Snapshot copies* are created. These Snapshot copies are used exclusively to facilitate frequent SnapMirror volume replication. Rolling Snapshot copies are replicated to the SnapMirror destination volume as soon as they are created. A maximum of two rolling Snapshot copies are maintained on the source volume.

Rolling Snapshot copies are best utilized on the transaction log volumes only. The transaction log volume is the most important volume to update because it contains the most recent journal of changes to your Exchange data. Replicating this volume more frequently maintains a very recent copy of the transaction logs on the SnapMirror destination volume. This allows SME to restore to a point that is more recent than the most recent backup set, as long as all data has been successfully replicated to the destination.

Best Practice

Implement rolling Snapshot copies on your transaction log volume only. You don't have to configure rolling Snapshot copies for the database volume or the SnapInfo volume (if separate from your transaction log volume), because the required data should exist on the destination. The database volume, transaction log volume, and SnapInfo volume should have been replicated after the last backup operation completed.

There are three major advantages to using the SnapDrive rolling Snapshot copy functionality:

- Fewer changes are made to the Exchange data between replications: When you use SnapDrive to initiate frequent updates to the transaction log volume, a restore from the destination volume replays the transaction logs, resulting in an up-to-the-minute restore from the time of the last transaction log volume update. Refer to [NetApp TR-3526](#) and [KB10542](#) for more information about replaying transaction logs into a recovered Exchange database using rolling Snapshot copies.
- Fewer SnapManager backups are required: It is not a best practice to perform SME backups every few minutes. You run a higher risk of overlapping SME operations, increased Exchange load, and potential higher latency times. Also, you will have a higher number of backup set to manage.
- Fewer Snapshot copies are retained: A maximum of two rolling Snapshot copies are retained at any time.

Determining the allowable amount of time between updates is important for your SnapMirror update schedule. You can recover your Exchange data up to the point of the last SnapMirror update. Any transaction logs or updates on the source SnapMirror volume that occurred before the last SnapMirror update will be lost in the event of a disaster. Understanding the importance of the data that is being mirrored also helps in determining your update schedule.

9. Summary

NetApp SnapManager 4.0 for Microsoft Exchange is an integral component of the NetApp data management solution for Microsoft Exchange Server environments. By reducing backup and restore times, minimizing Exchange outages, and consolidating Exchange storage, SME delivers a cost-effective solution for managing critical Exchange data.

The recommendations made in this paper are intended to be best practices for *most environments*. This paper should be used as a set of guidelines when designing, deploying, or administering SnapManager for Exchange. To ensure a supported and stable environment, review the concepts presented in this paper and involve an Exchange specialist if necessary.

