NetApp
Go further, faster

Technical Report

SnapManager 6.0 for Microsoft Exchange Best Practices Guide

Sourav Chakraborty, NetApp
May 2010 | TR-3845

TABLE OF CONTENTS

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 might 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. Companies require 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. NetApp offers a comprehensive suite of hardware and software solutions that enable 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.

NetApp® SnapManager® 6.0 for Microsoft Exchange software is available for Microsoft Exchange Server 2007 and 2010. SME 6.0 is tightly integrated with Microsoft Exchange, which allows for consistent online backups of your Exchange environment while leveraging NetApp Snapshot™ copy technology. SME is a Volume Shadow Service (VSS) (Snapshot copy) requestor, which means that it uses the Microsoft VSS subsystem to initiate backups. For details on VSS, see Microsoft KB article 822896. SME provides a complementary feature set for the new Microsoft Exchange 2010 database availability group (DAG). Apart from that, SME 6.0 also supports local continuous replication (LCR) and cluster continuous replication (CCR) environments in Exchange Server 2007.

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 recommended best practices for deploying and using SnapManager 6.0 for Microsoft Exchange with a NetApp storage system 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 2003/2008 environments with Microsoft Exchange Server 2007 and Microsoft Exchange Server 2010.

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 MIGRATING EXCHANGE DATA TO NETAPP STORAGE

The process of migrating Exchange data files from location to location can be a time-consuming and lengthy process. There are many manual steps that need to be taken to make sure the Exchange data files are in the proper state to be moved, and more manual steps need to be performed to bring those files back online and handling Exchange traffic. SME automates the migration process, eliminating any potential user errors. Once the data is migrated, SME automatically mounts the Exchange data files and allows Exchange to continue to serve e-mail.

Storage layout for Exchange data should be designed for optimal performance before the actual migration of Exchange data can be carried out. The other key drivers of storage layout design are high availability and improved data protection. In this section we discuss the best practices for designing storage layout for Exchange environments.

2.1 LAYOUT RECOMMENDTATION

The first phase in planning for migration of Exchange data to NetApp is to plan the storage layout. In this section we describe the best practices for Exchange data layout on NetApp.

Before we discuss each of the best practices, it is advised that you read the following sections of the Installation and Administration Guide:

- Rules for Exchange Server storage groups and databases enforced by the Configuration wizard.
- NTFS volume mounts points.
- SnapManager support for volume mount points.

These sections present a comprehensive list of guidelines that must be followed while designing the Exchange data layout on NetApp storage for SnapManager for Exchange (SME) to work successfully.

Best Practice

For Exchange Server 2007, mailbox databases belonging to the same storage group should be kept on either:

The same LUN

or

Individual LUNs on the same volume

Preferably the first of the two options should be used. This is because an SME backup job iterates through the volumes over which the data is located. The second option needs to be considered in those cases where databases in the same storage group (for Exchange Server 2007) need to be restored individually. For Exchange Server 2010, multiple mailbox databases cannot be put on the same LUN; hence, the second of the two preceding options should be used.

The reason for the preceding best practice is that the larger the number of LUNs, the more the amount of time it takes for backups to complete in SnapManager for Exchange due to various disk-related operations that are carried out as a part of the backup operation.

To further minimize the time taken to back up Exchange data, the following best practice is suggested.

Best Practice

Minimize the number of volumes used to store log files for Exchange databases or storage groups. For instance:

- In Exchange 2007, put the entire log LUNs for different storage groups on a common log volume.
- In Exchange 2010, use the preceding principle for databases since there are no storage groups.

SME will determine the number of backup groups based on the log volume count. For example, let's assume the following scenarios for Exchange 2007 with five storage groups:

- Scenario 1: The storage groups are spread across five volumes where each volume has one data LUN and one log LUN.
- Scenario 2: The storage groups are spread across five volumes where each volume has one data LUN. The log LUNs for these databases is placed on one common volume.

In the preceding example, scenario 1 will internally lead to the generation of five backup groups, compared to one in scenario 2. This means that scenario 2 will be five times as fast as scenario 1 in terms of time taken to complete a backup job.

Typically in DAG environments, backups are run on one of the replica nodes. Based on this assumption, the transaction log LUN on the passive node is correctly sized by taking the backup factor into account, but this factor is ignored for other nodes. This can lead to situations where the need to create backups from other nodes in the DAG can lead to issues with insufficient space on the transaction log LUNs on those nodes.

Best Practice

For Exchange Server 2010, it is strongly recommended that the same amount of space should be provisioned for transaction log LUNs on all the nodes of a DAG. The same holds true to nodes in an Exchange Server 2007 CCR cluster.

For additional information on Exchange data layout and sizing, refer to the Exchange Server 2010 Best Practice Guide.

The Snap Info directory in SME is the central repository. It contains two main pieces of data:

- Backup metadata
- Transaction log backups

In SnapManager for Exchange, if the Snap Info directory is placed in the same path as the log folder for a storage group (in Exchange 2007) or database (in Exchange 2010), then SME stores NTFS hard links to log files in the Snap Info directory while backing up that particular storage group or database. This not only saves space but also makes the log backups run faster.

Best Practice

Place the Exchange transaction log files and the Snap Info directory on the same LUN.

As a summary of the best practices discussed till now, we present the following composite view of those best practices:

Best Practice

In order to reduce the time taken to back up Exchange data using SME:

- Place Exchange data (database and transaction log files) into as few flexible volumes as possible.
- Place the SnapInfo directory in the LUN that contains the transaction logs.

The disk-to-disk backup archival feature in SME requires that the Exchange database files, log files, and SnapInfo directory be placed on QTree LUNs. This requirement arises from the design of the SnapVault® technology. Since many organizations do not plan backup archival as a part of the initial deployment, it might save time and effort later on, especially in the provisioning cycle, if this suggestion is followed.

Best Practice

Place database and transaction log files on QTree LUNs. In addition, place the Exchange log folder and SnapInfo folder together on a QTree LUN.

In order for SME to work, it needs to have a service logon account with appropriate permissions. As of now the least set of permissions is as follows:

- The SME service logon account must be an Exchange administrator in Exchange 2007/2010 and a member of the "organization management" role in Exchange 2010.
- In SMBR, access control to mailboxes is done using SMBR-Administrative Server (SMBR-AS). Hence, it is strongly suggested that SMBR-AS be used along with SMBR.

However, if the "create mailbox" functionality of SMBR needs to be used, the user's logon account must have "organization management" permissions in the Exchange organization. More details about the features and steps for installation of SMBR-AS are in the user guide.

SMBR-AS is very important to install in a hosted application environment where Exchange mailbox services are being provided to a large number of customers. Note that in this environment, one instance of SMBR-AS has to be installed in each individual Exchange partition that belongs to a particular customer.

Best Practice
In hosted Exchange environments, provision a virtual machine in each customer's environment and install SMBR-AS on the VM to centrally audit and control mailbox data access using SMBR. Note that one instance of SMBR-AS should exist in each individual domain.

2.2 MIGRATION PREPARATION

Once the Exchange data layout is planned and the appropriate storage provisioning tasks are completed, migration of Exchange data to the appropriate LUNs can be started. In this section we discuss some of the best practices around the same.

Best Practice
On Exchange Server 2010, use the same drive letters or mount points for the Exchange data LUNs on all the nodes of a DAG. Also, use the same drive letters or mount points for the Exchange data LUNs on all the nodes of a CCR cluster in Exchange Server 2007.

The preceding recommendation is crucial for seamless migration of Exchange data across all the nodes in a DAG environment. The preceding is also a prerequisite for adding a mailbox database copy to a DAG, as stated in http://technet.microsoft.com/en-us/library/dd298105.aspx.

Make sure of the following:

Best Practice
Install SnapManager and SnapDrive for Windows on all member servers of the database availability group (DAG) in Exchange Server 2010 and cluster nodes of a CCR setup in Exchange Server 2007.

In SME 6.0, DAG is in itself a unit of management: all the nodes of a DAG can be managed by registering the DAG with SME 6.0. Optionally, each DAG node can also be managed individually, although this is not recommended.

There can be situations where all member servers of a DAG might not use NetApp storage to store Exchange data. In this case two key questions arise:

- What is the deployment strategy?
- What is the licensing policy?

The deployment strategy in these situations is as follows:

Best Practice

Install the complete SME stack (SME, SDW, and so on) on the node(s) on which NetApp storage is used.

Do not install any NetApp component, including SDW, on the DAG member servers that use third-party storage.

Note that in the preceding scenario, SME cannot be used to connect to the DAG anymore. The user must connect to individual mailbox servers on which SME is installed and perform tasks related to SME. This includes migration of mailbox databases as well. This operation must be performed individually on each and every member node that uses NetApp storage and has SME installed on it.

Furthermore, it must be noted that in the preceding scenario, SME can be licensed on only those nodes on which it is installed. However, there is a risk involved in those situations where multiple mailbox databases reside on the DAG and the backup policy dictates that backups should be created on the active server. In the event of database failover for one of the many databases that reside in the DAG, it is possible that a DAG member server that does not have SME can become the active server. In this case, even if the member server is connected to NetApp storage, SME cannot be used on it.

Best Practice

It is recommended that for the host-based licensing model, SME licenses should be purchased for each member server in the DAG even if SME is not intended to be installed immediately on each member server.

However, SME will work even if it is licensed on a subset of the member servers in the DAG.

3 BACKUP BEST PRACTICES

Backing up Exchange data needs careful planning such that backups can be created as frequently as possible with the aim of minimizing data loss in the event of an outage. In this section we discuss the concepts related to backing up Exchange data using SnapManager for Exchange.

3.1 BACKUP

It is important to consider the following factors for planning a backup strategy for the Exchange data in the organization:

- Organization SLA: This parameter will determine the frequency and the type of backups.
- Disaster recovery planning: This will determine whether backup sets need to be mirrored to a secondary location or not.
- Backup verification policy: This parameter will determine when to engage backup verification and when not to.

SnapManager for Exchange allows configuring backups based on all the preceding factors, and these factors will form the basis of further discussions in the section.

The time taken to restore Exchange data in the event of an outage is dependent on the number of transaction logs that need to be replayed. Hence, reducing the number of transaction logs that need to be replayed when restoring from a full backup is important. The only way to do this is to create more frequent backups.

Best Practice
For Exchange Server 2007, backups in a CCR cluster should be scheduled on the passive node. For Exchange Server 2010, backups should be scheduled on the passive nodes only if they are zero lag.

It is also important to note that public folders in Exchange 2010 cannot be backed up using SME when connected to a DAG. In order to back up public folders, connect to individual member nodes of a DAG.

An important aspect of making sure of continued data protection in an Exchange environment is to make sure that backups are getting created as planned. This calls for active monitoring of the backup jobs being run by SME. There are two ways of monitoring the health of backup jobs:

- SME notification system: This feature in SME allows the administrators to receive detailed e-mails for failures in executing operations in SME. The e-mail notification can be configured using the SME configuration wizard. Refer to the installation and administration guide for more details.

- Monitoring host-side events: SME logs many events, each with its own event ID. These events get logged in the Windows application event log on the Exchange host. Monitoring some of the critical events using scripts, SCOM, MOM, and so on will help alert administrators of failures encountered by SME. The following table lists the most important events that must be monitored in an SME environment:

EVENT NAME	EVENT ID	DESCRIPTION
MSG_SME_VSSBACK_FAILED	204	Overall backup failed.
MSGERR_SME_SNAPMIRROR_UPDATE_FAIL	283	Update mirror failed when doing verification on SnapMirror® destination.
MSGERR_SME_SNAPMIRROR_UPDATE_TIMEOUT	286	Update mirror timeout when doing verification on SnapMirror destination.
MSG_SME_BACKUP_UPDATEMIRROR_FAIL	213	The request to update SnapMirror

		operation after backup failed.
MSG_SME_BACKUPFAIL	111	Overall backup failed.
MSG_SME_BACKUPFAIL2	139	Backup failed info specific to the SG, useful to narrow-down the error.
MSG_SME_ESEFILE_FAIL	131	
MSG_SME_ESEFILE_DEST_VOL_FAIL	132	Verification of SnapMirror destination volume failed.
MSG_SME_ESEFILE_REMOTE_FAIL	211	In case the verification server is a remote Exchange Server.
MSG_SME_ESELOG_REMOTE_FAILED	251	Same as event 132: applies to transaction log.
MSG_SME_ESELOG_FAILED	252	Same as event 211: applies to transaction log.
MSG_SME_RESTOREFAIL	112	Overall restore failed.

Recovery point objectives (Pros) have become a defining part of a data protection plan for Exchange. The ability to have a near zero RPO is highly desirable by Exchange administrators, as it minimizes the amount of data that is lost between the last full verified backup set and the point of failure. To help achieve desired service-level agreements (SLA) and RPO times, SME 6.0 now has frequent recovery points (FRPs). FRPs are optimized backup sets that are created through SME. The backup sets only contain the transaction log files that have been created since the last full backup or last FRP backup was created. Those transaction log files are copied into the SnapInfo directory, and then a Snapshot copy is created of the LUN containing the directory. And since the FRP backup sets contain a smaller amount of information, backups can be created very frequently, as often as every 10 minutes (highest frequency). The higher frequency of FRP backups reduces RPO times.

Best Practice

In order to reduce the time taken to exposure to data loss and create more frequent backups, it is recommended that frequent recovery points (FRPs) should be used. The recommended frequency of FRP backups is 15 minutes.

Best Practice

For Exchange Server 2007, use the deferred backup verification functionality to verify online backups. This saves resources as well as speeds up backups to a great extent.

In Exchange Server 2007, backup verification is a hard requirement when it comes to restore. However, in Exchange Server 2010 data can be restored from unverified backups.

It is neither required nor advised that backups be verified as a part of a backup job (although you can do so). Backup verification is a resource-intensive and time-consuming operation due to the way Exchange performs backup verification. Hence, deferring backups to later hours in a day can prove beneficial for optimizing backup jobs by reducing the actual duration of the backup job.

Observing verification throughputs in staging environments is a good option. This value is mentioned in the backup verification log in SnapManager for Exchange. This helps estimating completion times for backup verification jobs and hence planning a more aggressive backup verification strategy.

SME 6.0 supports both asynchronous and synchronous mirroring, with some caveats mentioned later.

If both the database and transaction log volumes are asynchronously mirrored:

Best Practice

Do not configure the SnapMirror replication schedule on the storage system. Create and initialize the SnapMirror relationships for all required SME volumes, but set the update schedule for each relationship in /etc/snapmirror.conf as - - - -. Refer to the Data ONTAP Installation and Administration Guide for more details on how to set up an asynchronous SnapMirror relationship. Additionally, Protection Manager can also be used to configure SnapMirror relationships.

If synchronous SnapMirror is to be used:

Best Practice

It is recommended that only the transaction log volume be synchronously copied with SnapMirror. The data volume can still be asynchronously copied with SnapMirror. In this case make sure that the "Update SnapMirror" option is selected for the backup plan.

Best Practice

If both data and transaction log volumes are synchronously copied with SnapMirror, then make sure that the "Update SnapMirror" option is not selected for the backup plan.

Best Practice

Synchronous SnapMirror is only supported with Data ONTAP® 7.3 or newer. Refer to the compatibility matrix for more information on corresponding SDW and host utility kit versions to use.

3.2 BACKUP VERIFICATION

Microsoft requires backup sets to be verified. The process of verifying backup sets can be time consuming and cause significant I/O load on an Exchange Server and on the storage system. SnapManager for Exchange has many ways to assist an Exchange administrator in mitigating the I/O load for verifications, that is, deferred verification and Remote Verification Server.

SnapManager 5.0 for Exchange can now perform multiple backup set verifications simultaneously. This functionality allows multiple Exchange Servers to offload the verification process to a single verification server. Note that a single Exchange Server can only run a single verification process on the verification server. For example, ExchSvr1 and ExchSvr2 each submit a verification job to the remote verification server, which will both run concurrently. With those jobs still running, ExchSvr1 submits another verification job. That job will then be queued behind the currently running verification job previously submitted by ExchSvr1. Presently, the queue depth can be 4.

Best Practice

Use virtualized verification hosts to permit the use of dedicated virtualization hosts for each Exchange server or CMS to allow for maximum scheduling flexibility.

In order to speed up the backup verification server, the following strategy can be used:

1. Choose an appropriately powerful Hyper-V™ or VMware® server in the Exchange environment.

2. Create a golden image of a VM that has all the components of a remote verification server installed on it. Refer to remote verification prerequisites and remote verification server requirements in the Installation and Administration Guide for SME 6.0.

3. Using the golden image created earlier, create as many virtual machines as required and start these.

4. Now on each Exchange mailbox server in the environment, schedule a Windows scheduled task that is set to run the following SnapManager for Exchange cmdlet verify-backup. One of the arguments of this cmdlet is –VerificationServer. In this argument, put the name of one of the virtual machines from the preceding list.

Best Practice

Using the preceding solution you can concurrently schedule backup verification for four storage groups at a time per mailbox server if one remote verification host (VM) can be created for each mailbox server.

Note that it is possible to use iSCSI-based connectivity in the verification environment even if FCP is used in the production environment.

SnapManager for Exchange also supports verification of backups that have been archived to the SnapVault secondary location. Verification of archived backups depends on the amount of data held in the archived backups and the speed of the disks. Note that if the SnapVault secondary storage system has SATA disks, then the verification time could go up when compared to verification of online backups on the primary storage device.

4 SNAPVAULT INTEGRATION WITH SME

SnapManager 6.0 for Exchange Server is integrated with NetApp Protection Manager data sets. This integration allows Exchange administrators to archive database backups to a NetApp SnapVault secondary destination. Thus SME 6.0 offers a complete disk-to-disk backup archival solution integrated with the product. The benefits of this feature are:

- Backup archival is supported as a part of the backup workflow in SnapManager for Exchange, which enables administrators to easily archive online backups and optimize the usage of primary storage.
- Archived backups can also act as a DR strategy in addition to backups created with SnapMirror.
- Long-term backup retention can be applied to archived backups that allow requirements of legal compliance to be met with ease.

There are some prerequisites for planning disk-to-disk archiving using Protection Manager data sets:

- Target mailbox databases should reside on qtree LUNs with one LUN per qtree.
- NetApp Management Console should be installed in the organization and should be accessible to SnapManager.

At times, the base initialization of the SnapVault relationship created by SnapManager for Exchange might fail due to number of reasons such as:

- Inadequate space at the destination storage device. This might happen during baseline transfer to initialize the SnapVault relationship or during a later stage as Snapshot copies are being archived.
- Network issues such as other high-volume data transfers using the same network, general timeout issues, network slowdown, and so on.

Hence, planning for space on the SnapVault secondary is important. Before we begin there are three factors that we must be aware of for reasonably estimating the space requirements at the SnapVault secondary:

- Total baseline volume sizes (B): Sum of the LUN sizes of all the mailbox databases residing on the mailbox server. Note that this could be larger than the actual size of the mailbox database. In case more than one databases share the same LUN, then the LUN size should be counted only once for those databases.
- Maximum retention period (T): Maximum retention period set for the Protection Manager data set that corresponds to the SnapVault relationship.
- Total database daily change (R): Sum of the daily change rates for the mailbox databases on the server. The daily change rate value for each mailbox database can be obtained by using the formula mentioned in the Exchange 2010 Best Practices Guide.

Therefore, the formula to estimate the space requirements at the SnapVault secondary is:

```
Space Needed = B + (R * T)
```

In addition, it is important to plan for unplanned data growth or retention policy. Hence, it is strongly advised that buffer space be planned at the SnapVault destination.

Best Practice

The SnapVault secondary storage device should have at least 10% more free space than the size of the source volume on which the Exchange data resides.

Best Practice

For low-bandwidth network links between source and destination storage devices, set the Throttle setting in Protection Manager to “Nightly Transfer Window.”

The most important aspects of backup archival are:

- Monitoring: SnapVault transfers initiated from SnapManager for Exchange (SME) can be monitored from the dashboard provided with SME. Note that SnapVault update (backup archival) is an asynchronous operation and SME queues the archival job in Protection Manager.
- Space consumption: If longer retention policies are required for the archived backups then monitoring the space usage becomes important. Protection Manager provisions a thinly provisioned volume at the SnapVault secondary that is of the same size as the aggregate on which the source volume resides.

Best Practice
Operations Manager should be used for active monitoring of space consumption at both the primary and secondary storage devices. Operations Manager provides predefined reports for monitoring SnapVault secondary devices. Refer to the Operations Manager guide for more information.

Note: Connect to individual DAG member servers in an Exchange 2010 setup to perform operations related to SnapVault. SnapVault integration is not present in the DAG view (when connected to a DAG) in SME 6.0.

5 RECOVERY

The ability to recover your Exchange databases when necessary is a critical operation for an Exchange administrator. SME restore functionality allows you to recover your Exchange databases and transaction logs from backups that it created or from archive.

There are two types of restore operations in SME:

- Up-to-the-minute: Selected by default, an up-to-the-minute restore replays any necessary and available transaction logs from the backup set and from the transaction log directory and applies them to the database. A contiguous set of transaction logs is required for an up-to-the-minute restore to succeed.
- Point-in-time: This option allows you to restore your Exchange data to a chosen point in time. Any Exchange data past that point is not restored. This option is particularly useful when trying to restore to a point before something such as data corruption occurred. A point-in-time restore only replays and applies to the database of those transaction logs that existed in the active file system when the backup was created up to the specified point in time. All transaction logs beyond that point in time chosen are discarded.

In this section we will discuss the best practices and solutions with respect to restoring Exchange data.

5.1 RESTORING PASSIVE COPIES IN A DAG SETUP

One of the key challenges in a DAG environment is to minimize reseeding of database copies in the event of a restore. This serves to greatly reduce the network bandwidth resources that are consumed by the reseeding operation. Consider a scenario where a passive copy of the database has been in the failed state for a while with no replication enabled between it and the primary active copy. Now in the event that we restore the passive copy there are two key challenges:

- The passive copy will become an active copy.
- A full database reseed will be required for the DAG setup to work again.

In order to minimize the network traffic generated by the reseed operation and recover the database copy at the passive-node, do the following:

1. Suspend the replication between the active database copy and the passive database copy.

2. Use SnapManager for Exchange to restore the passive database to a suitable point in time. Make sure that the “Recover and mount Database after restore” box is unchecked.

3. After the restore operation of the passive database copy completes, it shows up in the Exchange

Management Console in the “Dismounted” state. However, it becomes the current active copy of the database.

4. Activate the currently passive copy of the database that was the active copy of the database before the restore was initiated.

5. Mount the database copy that was just restored and resume replication. After some time the passive database copy becomes healthy. No reseed operation is required.

CONCLUSION

NetApp SnapManager 6.0 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 Microsoft Exchange. To make sure of a supported and stable environment, review the concepts presented in this paper and involve an Exchange specialist if necessary.