NETAPP TECHNICAL REPORT

SnapManager 5.0 for Microsoft Exchange Best Practices Guide

Shannon Flynn, NetApp
November 2008 | TR-3730

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® 5.0 for Microsoft Exchange software is available for Microsoft Exchange Server 2003 and 2007. SnapManager 5.0 for Microsoft Exchange (SME) has earned Certified for Windows® Server 2008 accreditation with the Hyper-V™ designation. SME 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 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. SME also supports standby continuous replication (SCR), providing a rich feature set for the active node only.

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 5.0 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 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 MIGRATING EXCHANGE DATA ONTO NETAPP

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.

2.1 LAYOUT RECOMMENDATIONS

Demand for higher availability, increased performance, and improved data protection is becoming the norm, requiring careful planning and consideration in preparation for deploying an Exchange environment. Recovery point objective (RPO) and recovery time objective (RTO) times need to be factored into this planning.

Best Practice
It is recommended, but not required, to place database and transaction log volumes on separate aggregates.

RAID-DP® safeguards data residing on an aggregate by providing double disk failure protection, while still maintaining the performance requirements required by Exchange. Placing database volumes and transaction log volumes on separate aggregates provides high data availability for Exchange in the extremely rare event that more than two disks in a RAID-DP aggregate fail. For more information on RAID-DP, please refer to the RAID-DP product page on the NetApp website.

Number of Aggregates	Benefits	Tradeoffs
Single aggregate for all Exchange data	**Minimize the number of parity disks required** **A single aggregate contains all of the disks that can be used to handle the I/O requirements for Exchange**	**Requires restoring from an archived backup to restore a lost aggregate (very unlikely)**
Database volumes and transaction log volumes on separate aggregates	**Dedicated aggregate handling the I/O for database or transaction logs** **If a single aggregate is lost (extremely rare), Exchange data can be recovered from other aggregate**	**More parity disks required because of more aggregates**

Best Practice
Place database LUN and transaction log/SnapInfo LUN in separate volumes.

Best Practice
When separate LUNs are used for the Exchange transaction log files and the SnapInfo directory, place those LUNs in the same volume. Both of these LUNs will have a similar I/O profile, allowing them to share the same volume. And for disaster recovery scenarios, having the entire log set for Exchange on the same volume will help achieve SLAs.

TRANSACTION LOG ARCHIVING

During the backup process, SME archives the transaction log files generated by Exchange and by default stores those files in the SnapInfo directory. The size of the SnapInfo directory depends on the number of transaction logs generated for a storage group. The space required to archive the transaction logs must be taken into consideration when creating volumes and LUNs for an Exchange deployment.

When the transaction log directory resides on a different NTFS volume than the SnapInfo directory, SME will copy each transaction log into the SnapInfo directory. Once the transaction logs have been copied, Exchange will delete the original copy of the log files that were committed to the database.

When the transaction log directory resides on the same NTFS volume as the SnapInfo directory, SME will utilize NTFS hard links to avoid the file copy operation and maximize storage utilization. During the backup process, SME will modify the file pointers for the transaction logs to add a pointer to the SnapInfo directory, thus avoiding a physical file copy. Once the backup of the storage group is complete, Exchange will truncate its transaction logs by removing the file pointers for the transaction log files that were committed to the database. As a result, it will appear as if the transaction logs were moved from one directory to the other. Once this copyless transaction log operation is complete, SME will continue its backup process.

Best Practice

It is recommended to use NTFS hard links by placing the SnapInfo directory on the same LUN as the transaction log directory whenever possible. This will increase storage utilization, eliminate the physical copy overhead incurred on the Exchange Server, and increase backup performance.

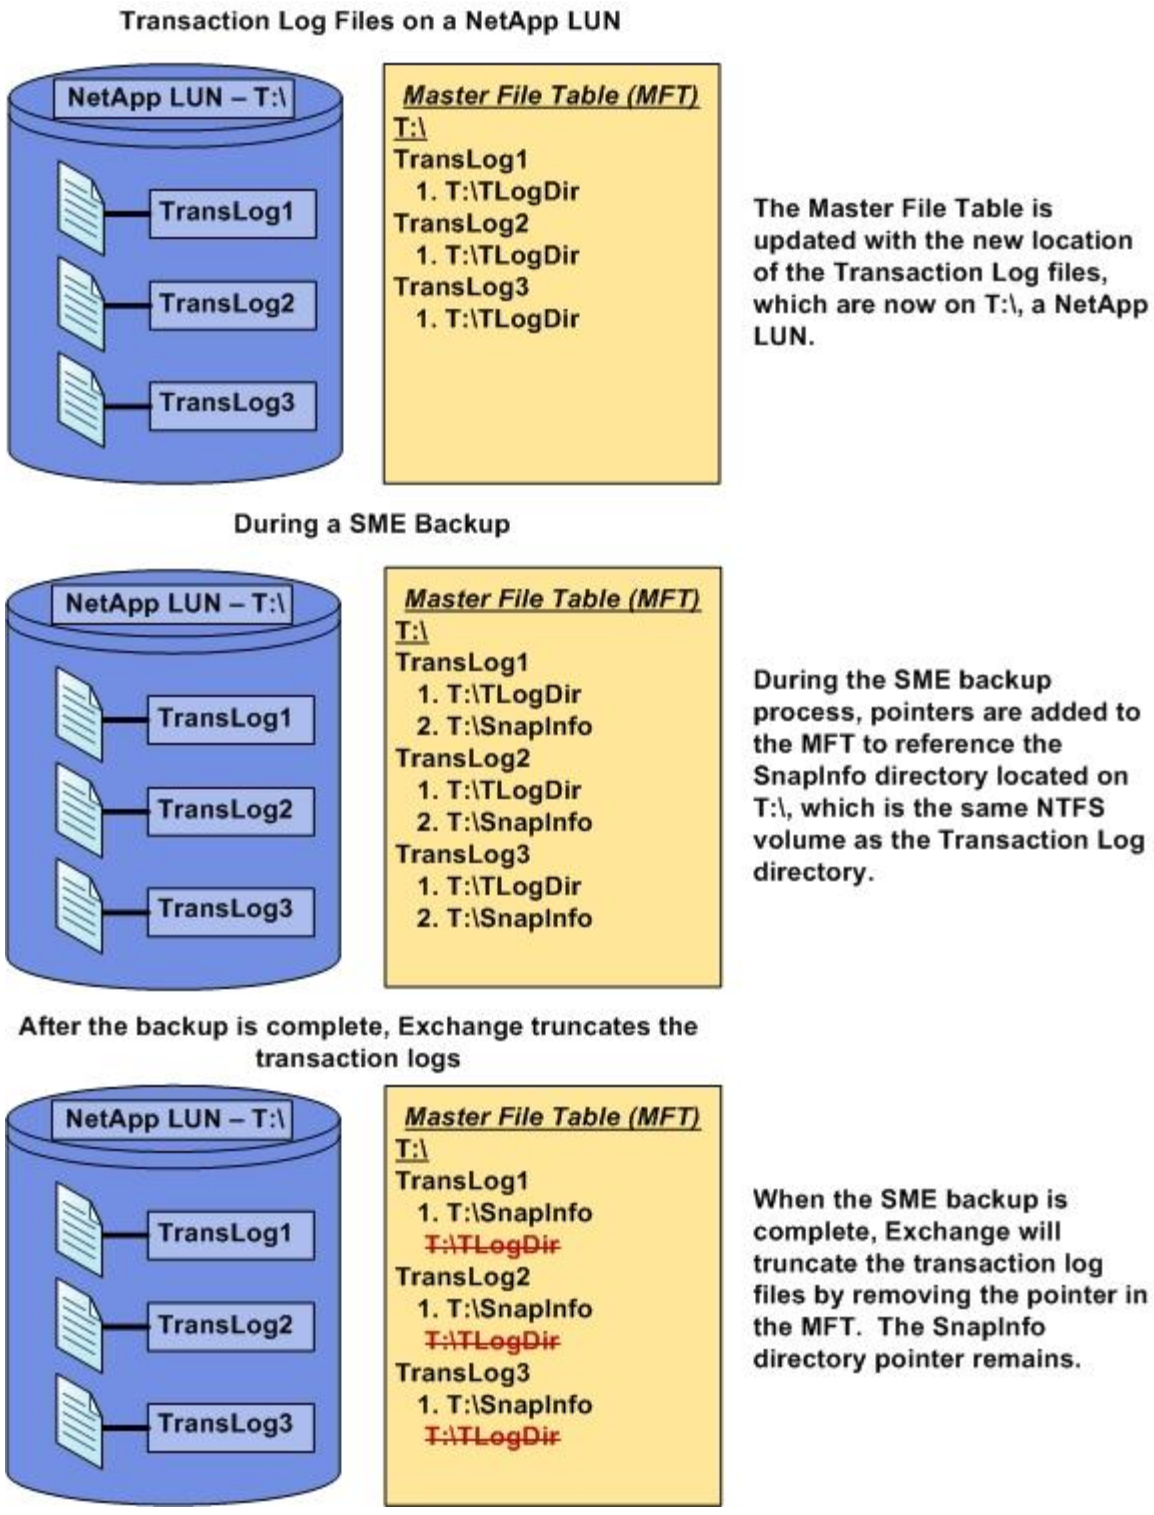

Figure 1) NTFS hard links.

To avoid transaction log archiving using NTFS hard links, the Exchange transaction log file directory and the SnapInfo directory must reside on different NTFS volumes. This means an additional LUN will be needed for the SnapInfo directory.

FRACTIONAL SPACE RESERVATIONS

NetApp recommends proper space management policies be set for volumes hosting LUNs containing Exchange data. This guarantees sufficient space for all write operations on Exchange data volumes and provides zero downtime for Exchange users.

A fractional space reservation policy can be implemented in an Exchange environment. The benefits of

doing this are greater space utilization and a potentially lower initial cost of storage. Determining the exact space management policies will depend on your specific Exchange business requirements.

SME provides a method to monitor the overwrite reserve utilization on fractionally space-reserved volumes. SME will take appropriate actions to prevent a LUN from becoming inaccessible due to a no free space condition on the hosting volume. SME can be configured to take the following actions.

- Automatic deletion of Exchange backup sets
  - This option allows SME to automatically delete the oldest Exchange backup sets it created to free up space on the volume.
  - SME will retain the most recent backup of any database on the fractionally reserved volume.
  - SME will also keep the last backup of any database that no longer exists.
- Automatic dismounting of Exchange databases
  - This option allows SME to automatically dismount a database residing on a fractionally reserved volume that is nearing a no free space condition.
  - This prevents Exchange from forcibly dismounting a storage group due to an out of space error.

SME can also be configured to execute both of the options above. When using both options, the threshold limits for the automatic deletion policy must be set lower than those for the automatic dismount policy. This will make sure that SME will attempt to delete older backup sets to free up space on the volume before it dismounts Exchange databases.

Best Practice

Choose an SME fractional space reservation policy that works best for the Exchange environment. If the backup set deletion policy is triggered, make sure that a minimum of one verified backup set remains on the volume.

For additional fractional space reservation information, please refer to TR-3578: Microsoft Exchange Server 2007 Best Practices Guide. Although this technical report focuses on Exchange Server 2007, the fractional space reservation calculations also apply for Exchange Server 2003 configurations.

Work with your local NetApp professional to make sure your storage is sized correctly, all factors are taken into consideration, and a successful fractional space reservation policy is implemented.

DRIVE LETTERS AND VOLUME MOUNT POINTS

Exchange environments that require multiple storage groups or clustered configurations can quickly use the available drive letters for a given server. When there are no available drive letters to assign to a LUN, you must use volume mount points (VMPs). VMPs are directories on a Windows volume that map to a mounted LUN. VMPs can be used at any time, whether drive letters are available or not.

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.

Best Practice

Make the transaction log LUNs the volume mount point root, because they will be backed up on a regular basis. This helps make sure that your volume mount points are preserved in a backup set and can be restored if necessary.

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.

2.2 LARGE-SCALE OR SIMILAR DEPLOYMENTS

When deploying large-scale Exchange environments, the task of configuring each Exchange Server can be very repetitive, especially when dealing with Exchange Servers with similar configurations. SME 5.0 helps alleviate that repetitiveness with a new control file–based configuration option. The XML-based control file contains all of the configuration information that is set through the configuration wizard. The file is exported from an installed and configured Exchange Server and then can be used in multiple ways, such as:

- Unattended installations
- Replicating an existing configuration to another new Exchange Server
- Rebuilding an existing Exchange Server
- Moving an Exchange Server onto new hardware

Certain parameters in the exported control file need to be modified prior to importing the control file on a different server. Examples of parameters that might need to be modified are server names, storage system names, volume names, and so on.

Best Practice

It is recommended to review the settings in the configuration file before running it for the first time. This will help make sure that the file is free of errors and will configure your Exchange data without issues.

The configuration file can contain information for all sections of the configuration or a subset of sections. The use of an XML editor is highly recommended when editing the control file. Edit the appropriate sections, like the server name, and then run the configuration wizard with the control file option to configure that Exchange Server.

```
<?xml version="1.0" ?>
- <SMECONFIG xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xmlns:xsd="http://www.w3.org/2001/XMLSchema">
    <SERVER_NAME>excms</SERVER_NAME>
    <HOST_NAME>W2K8-EXSVR1</HOST_NAME>
  + <STORAGE_LAYOUT>
  - <COMMON_SETTINGS>
    - <NOTIFICATION>
      - <SEND_EMAIL_NOTIFICATION>
          <NOTIFY_AUTO>false</NOTIFY_AUTO>
          <LONG_MSG>false</LONG_MSG>
          <AS_ATTACHMENT>false</AS_ATTACHMENT>
          <SEND_ON_FAILURE>false</SEND_ON_FAILURE>
        </SEND_EMAIL_NOTIFICATION>
        <EMS_ENABLED>true</EMS_ENABLED>
        <ASUP_ENABLED>true</ASUP_ENABLED>
        <ASUP_ON_FAIL>true</ASUP_ON_FAIL>
      </NOTIFICATION>
    - <VERIFICATION>
      - <VERIFICATION_CLIENT_SETTING>
          <VERIFICATION_SERVER>w2k8-exsvr2</VERIFICATION_SERVER>
          <VER_SERVER_NTAUTH>false</VER_SERVER_NTAUTH>
        </VERIFICATION_CLIENT_SETTING>
      - <VERIFICATION_SERVER_SETTING>
          <AUTO_DRIVELETTER>false</AUTO_DRIVELETTER>
          <MP_DIR>C:\Program Files\NetApp\SnapManager for Exchange\SnapMgrMountPoint</MP_DIR>
          <ESEUTIL_PATH>C:\Program Files\Microsoft\Exchange Server\bin\eseutil.exe</ESEUTIL_PATH>
          <THROTTLE>false</THROTTLE>
          <IO_PAUSE>0</IO_PAUSE>
        </VERIFICATION_SERVER_SETTING>
      </VERIFICATION>
      <REPORT_DRIECTORY>C:\Program Files\NetApp\SnapManager for Exchange\Report</REPORT_DRIECTORY>
    - <BACKUP>
      - <BACKUP_CLIENT_SETTING>
          <NAMING_CONVENTION>0</NAMING_CONVENTION>
          <BACKUP_SET_TO_KEEP>8</BACKUP_SET_TO_KEEP>
          <BACKUP_SET_TO_KEEP_IN_DAYS>0</BACKUP_SET_TO_KEEP_IN_DAYS>
          <DELETE_BACKUPS_OPTION>0</DELETE_BACKUPS_OPTION>
          <BACKUP_SET_TO_VERIFY>5</BACKUP_SET_TO_VERIFY>
        </BACKUP_CLIENT_SETTING>
      </BACKUP>
    - <VERIFICATION_ON_DESTINATION>
      - <SELECTED_DESTINATIONS>
        - <SELECTED_DESTINATION>
            <SOURCE_FILER>fas980-svl21</SOURCE_FILER>
            <SOURCE_VOLUME>db_vol</SOURCE_VOLUME>
            <DESTINATION_FILER>fas980-svl22</DESTINATION_FILER>
```

Figure 2) Control file example.

3 BACKUP AND RECOVERY

3.1 FREQUENT RECOVERY POINT BACKUPS

Recovery point objectives (RPOs) 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 (SLAs) and RPO times, SME 5.0 now has frequent recovery points (FRPs). These 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. The higher frequency of FRP backups reduces RPO times.

Frequent recovery points are also beneficial in multisite business continuance (BC) configurations, especially when using NetApp SnapMirror®. SnapMirror replication times are based on a few factors such as the speed of the connection between the source and destination storage systems and the amount of data that needs to be replicated across that connection. By minimizing the amount of data that needs to be replicated, you can dramatically reduce the time to complete the replication. FRPs help reduce the amount of data for a SnapMirror replication by only replicating the transaction log files that have been created since the last full backup or the last FRP backup. The net result of this is a shorter replication time for SnapMirror.

Best Practice

It is recommended to use frequent recovery points whenever possible in an Exchange environment where RPOs are very stringent and/or SnapMirror replication times and/or data set sizes need to be reduced.

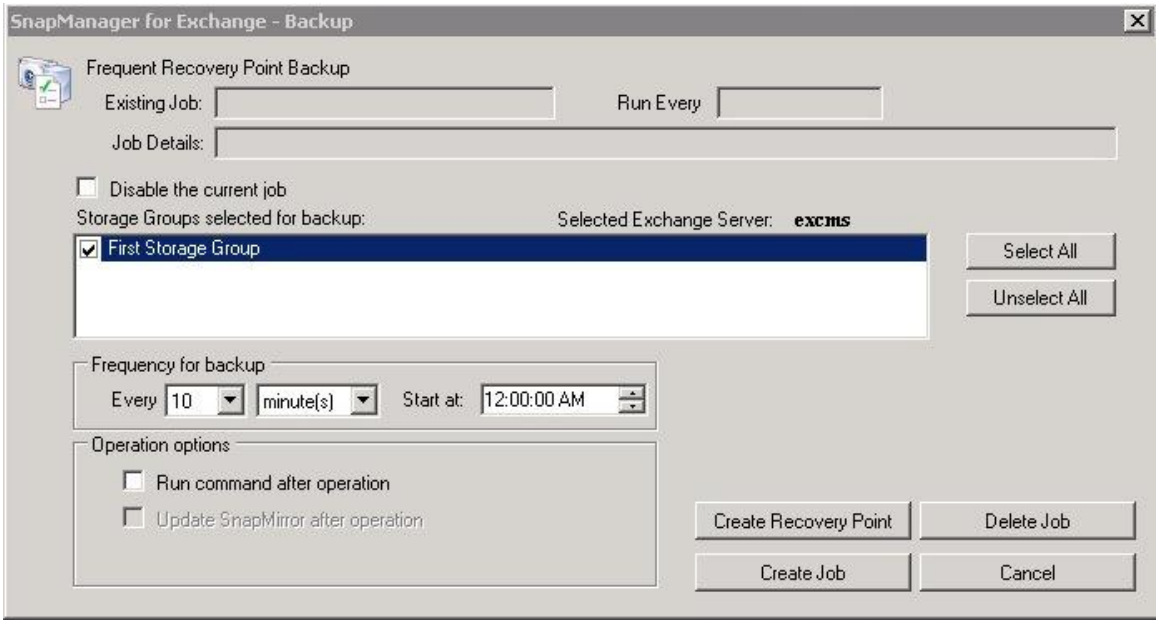

Figure 3) Frequent recovery point scheduler.

A full backup needs to be created before you can create frequent recovery point backups for that backup set. FRPs are scheduled through SME using the Windows task scheduler to execute the task. You can create a FRP backup schedule to run as frequently as every 10 minutes,

Best Practice

Schedule FRP backups at the lowest possible time interval required for the backup to complete. A 10-minute FRP backup will achieve a 10-minute RPO. FRP backups reduce how often full backups need to be created; one or two full daily backups are sufficient to achieve a low RTO as well. As always, check with your local NetApp professional to determine an FRP backup schedule that suits your Exchange environment.

When restoring Exchange databases that have FRP backups associated with them, SME will list the available FRP point-in-time backups for a given database. The backups will be displayed in 24-hour increments for each backup set. Selecting an FRP backup set will restore the last full backup set, then replay the transaction log files up to the selected recovery point.

3.2 BACKUP VERIFICATIONS

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. Common practice is to perform verification operations during off-peak hours to minimize the impact to the active Exchange Server and not adversely affect Exchange latencies.

SME has many ways to assist an Exchange administrator in mitigating the I/O load for verifications.

- Deferred verification: This option allows you to schedule your verifications to run at a later, more convenient time. By scheduling the verification process, the Exchange administrator does not have to worry about starting the verification process manually through the SME backup interface.
- Verification throttling: This option allows the Exchange administrator to throttle the amount of I/O load the verification process places on an Exchange Server.

- Backup verification on a SnapMirror destination volume: When replicating Snapshot copies using SnapMirror, SME can use the destination storage system to verify the backup.
- Remote verification server: This allows an Exchange Administrator to run verifications on a separate server, thus removing the I/O load from the active Exchange Server altogether.

MULTIPLE BACKUP SET VERIFICATIONS

SME 5.0 can now perform multiple backup set verifications simultaneously. This functionality allows multiple Exchange Servers to offload the verification process to a single verification server. Please 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.

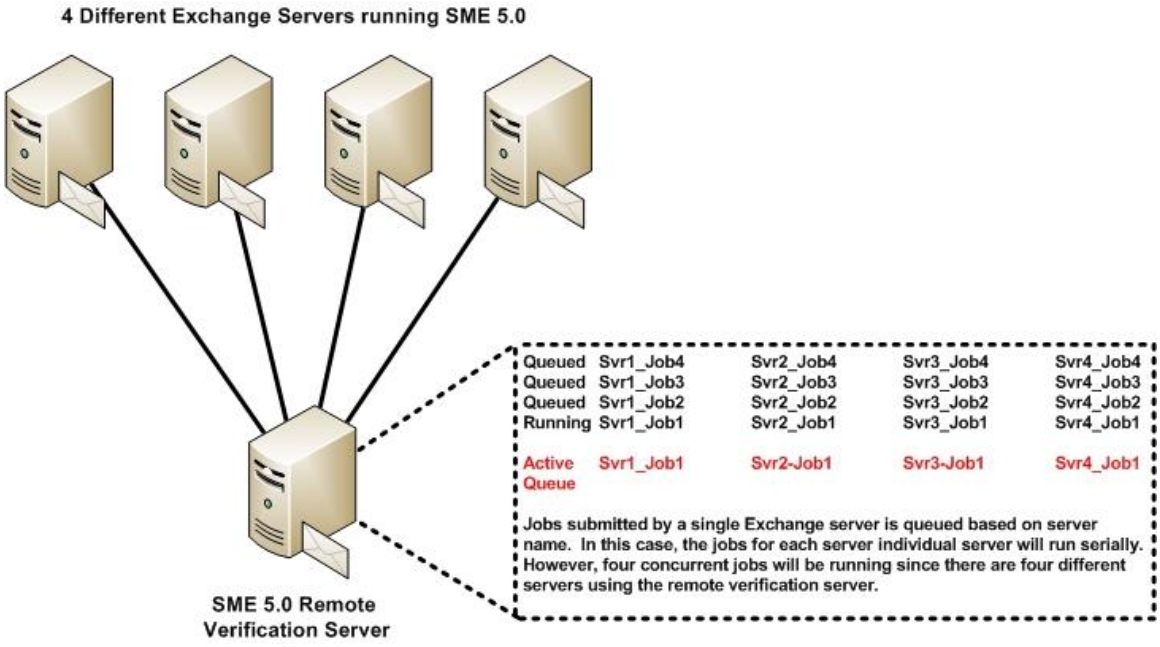

Figure 4) Concurrent verification behavior in SME 5.0.

The verification server can handle up to four concurrent verification operations. Any verification jobs that are submitted while four verification operations are running will be queued and handled on a first come, first served basis.

You can monitor concurrent verification jobs through the main SME management console. You have the option to monitor both locally running jobs and remote jobs for the remote verification server.

3.3 RECOVERING EXCHANGE DATA

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 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.

Best Practice

When performing an up–to-the-minute restore, restore from your most recently verified backup. This is the fastest way to restore your Exchange Server. Using an older verified backup slows the restore time because it requires more transaction logs to be replayed and applied to the database. Using the most recently verified backup helps make sure of the quickest recovery of your Exchange database.

3.4 RESTORE BACKUPS TO ANOTHER SERVER

In the event that an Exchange Server 2007 fails, complex manual steps are required in order to recover the data from the failed Exchange Server. SME automates this manual process by allowing an Exchange administrator to restore a backup set to another Exchange Server. The only prerequisite for this process is for the LUNs that contain the Exchange data to be mapped to the new target Exchange Server. Once the database, transaction log, and SnapInfo LUNs are mapped to and available on the target server, an Exchange administrator can simply run the restore wizard and chose the `Restore backup created on a different server` option in the restore wizard to recover the Exchange data to the new server.

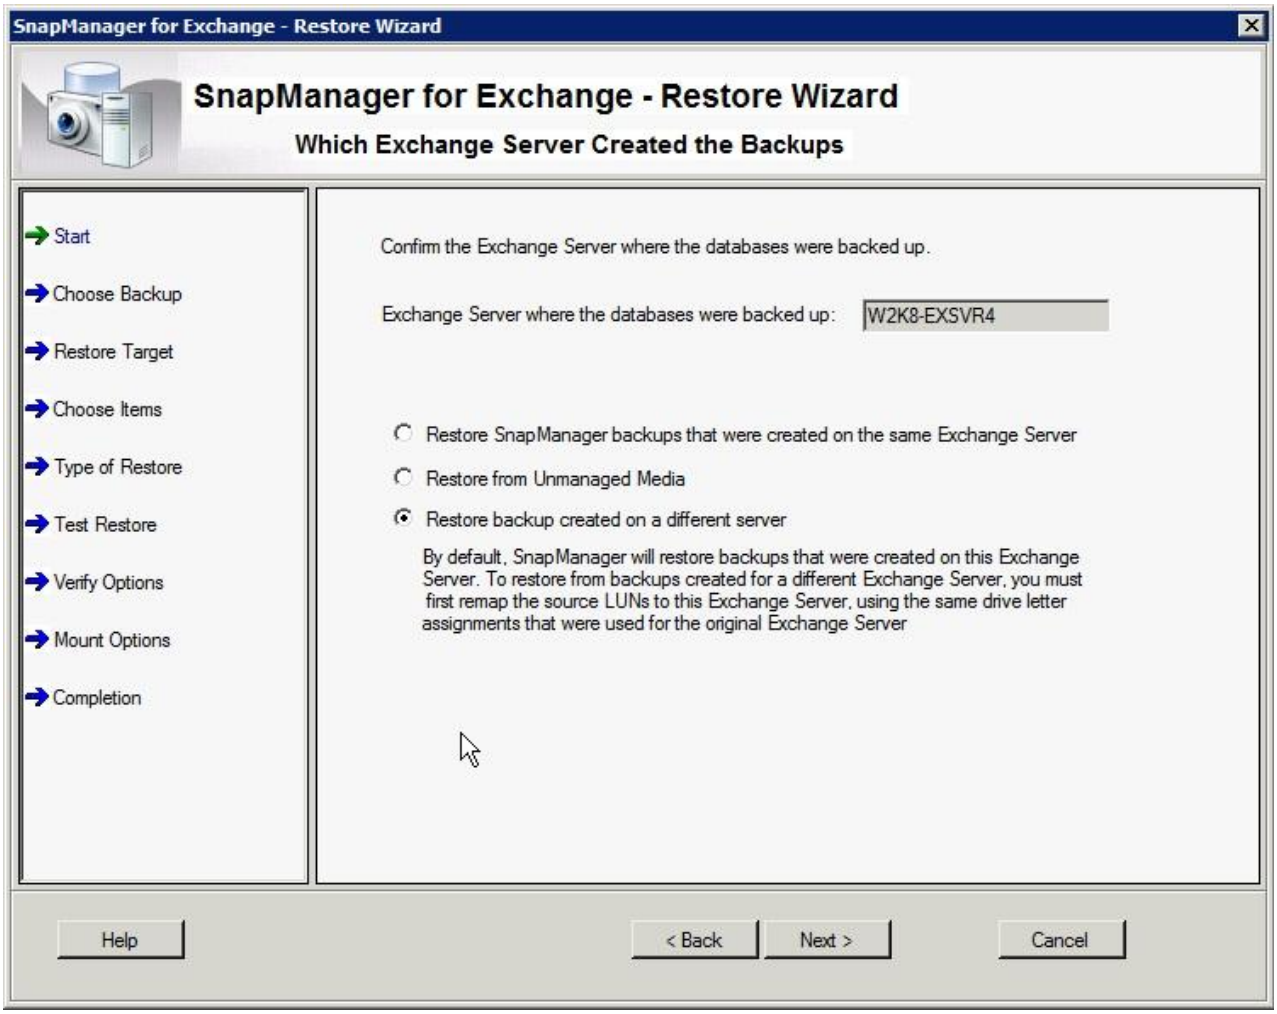

Figure 5) Restore wizard options.

REHOME MAILBOXES FOR EXCHANGE 2007

After the databases for a failed Exchange Server have been restored onto a new Exchange Server, the administrator must rehome affected mailboxes: update the user accounts in Active Directory for mailboxes to map to the new Exchange Server. Until the mailboxes are rehomed, users logging into Microsoft Outlook will be unable to retrieve their e-mail since Outlook will attempt to connect to the failed Exchange Server. SME 5.0 automates remapping of user accounts to mailboxes on the new Exchange Server for the administrator. Choosing the `Update user account associations` option in the SME restore wizard causes SME to automatically remap the user mailboxes to the new Exchange Server. This will update the Active Directory records so that the next time a user logs into Outlook, it will connect to the correct Exchange Server and retrieve the user's e-mail without error.

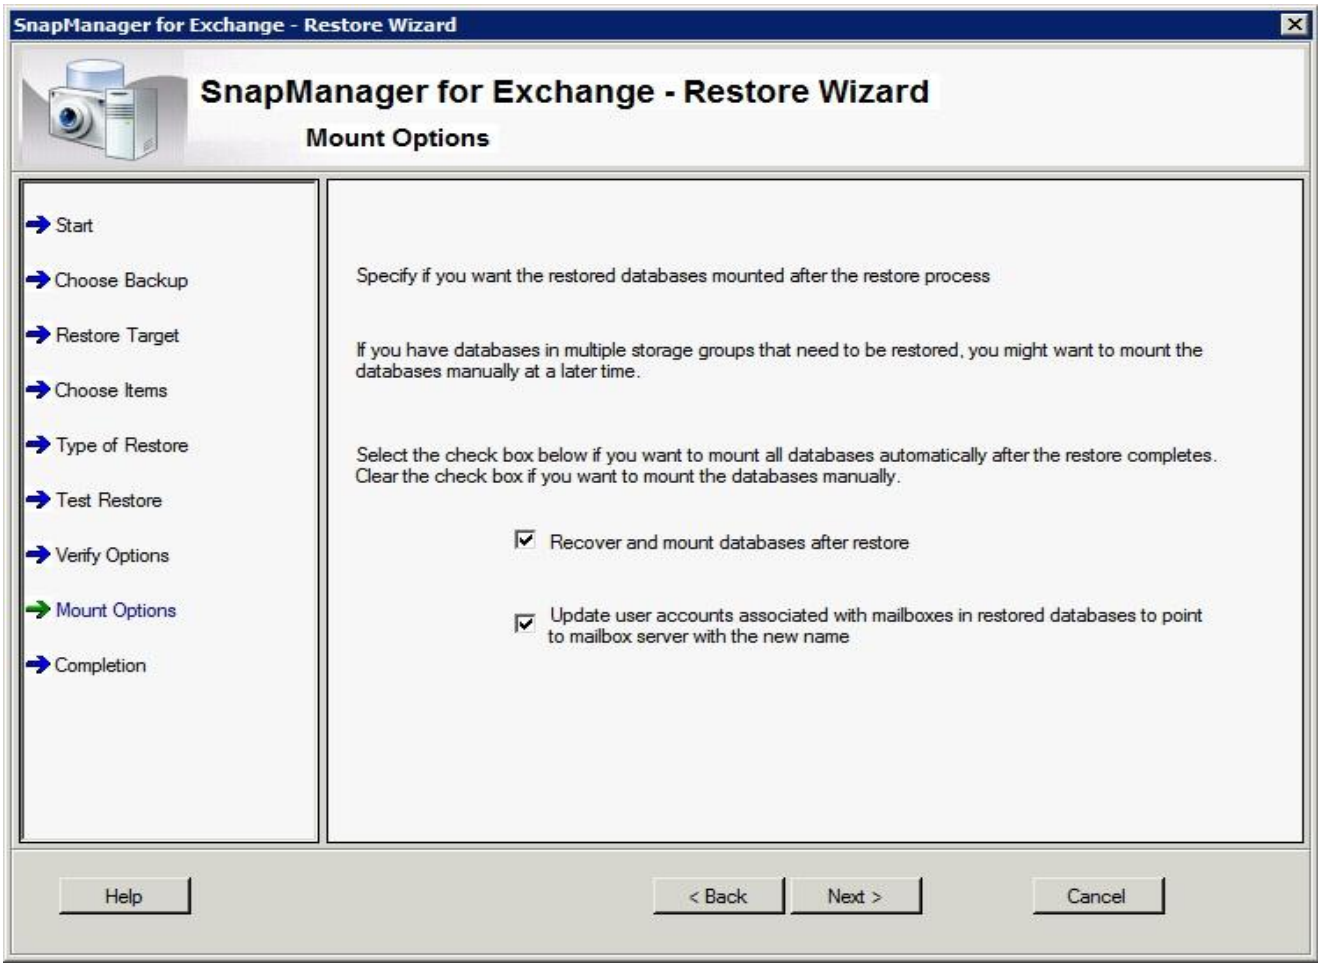

Figure 6) Rehome mailbox option.

4 BUSINESS CONTINUANCE AND HIGH AVAILABILITY

Microsoft Exchange is one of the most mission-critical applications for most companies. That means Exchange must be available at all times for users to access e-mail and schedules in order to do business in an efficient manner. High availability has become a primary concern for all Exchange administrators.

In the event of an Exchange failure, the ability to quickly recover from that failure and restore e-mail services is essential. Having a BC plan describing the steps to quickly recover and restore e-mail services in the event of a failure is a necessity.

NetApp storage systems and technologies, combined with SnapManager for Exchange, provide a very compelling high-availability/BC solution for Exchange. Core technologies such as Snapshot copies, SnapMirror, SnapVault®, RAID-DP, and more are built-in to the storage system. SME takes advantage of these features and automates their use in an easy-to-use interface.

4.1 REPLICATION

Replicating Exchange data to a secondary location is essential for protecting the business-critical e-mail data. Without that level of protection, a business risks extensive Exchange downtime, and this translates into lost revenues and productivity. NetApp SnapMirror is a key technology to protect critical data for applications such as Exchange. It has the ability to replicate data to multiple locations at high speeds in a simple, reliable, and cost-effective manner.

SME integrates with SnapMirror to provide an automated method to replicate successful backup sets to a destination storage system. When LUNs containing the database, transaction log, and SnapInfo files reside

on volumes that have a SnapMirror relationship, SME will provide an option to automatically update the SnapMirror relationship once the backup finishes successfully.

Best Practice

When creating the SnapMirror relationship, make sure the schedule parameters are set to ("- - - -"). This establishes the SnapMirror relationship and allows SME, communicating through SnapDrive, to update the mirror and manage scheduling of the updates.

An additional benefit for Exchange environments running SnapMirror is the ability to perform backup verifications on the SnapMirror destination storage system. This eliminates the I/O load on the production storage and utilizes the secondary storage system, which typically sits idle until it is needed. When used in conjunction with a remote verification server, you can completely remove the I/O load incurred by backup verifications from the production Exchange Server.

Best Practice

Whenever possible, perform database verifications on a SnapMirror destination. By offloading the database verification I/O to the secondary storage, the primary storage is free to handle regular Exchange traffic without the impact of additional verification I/O.

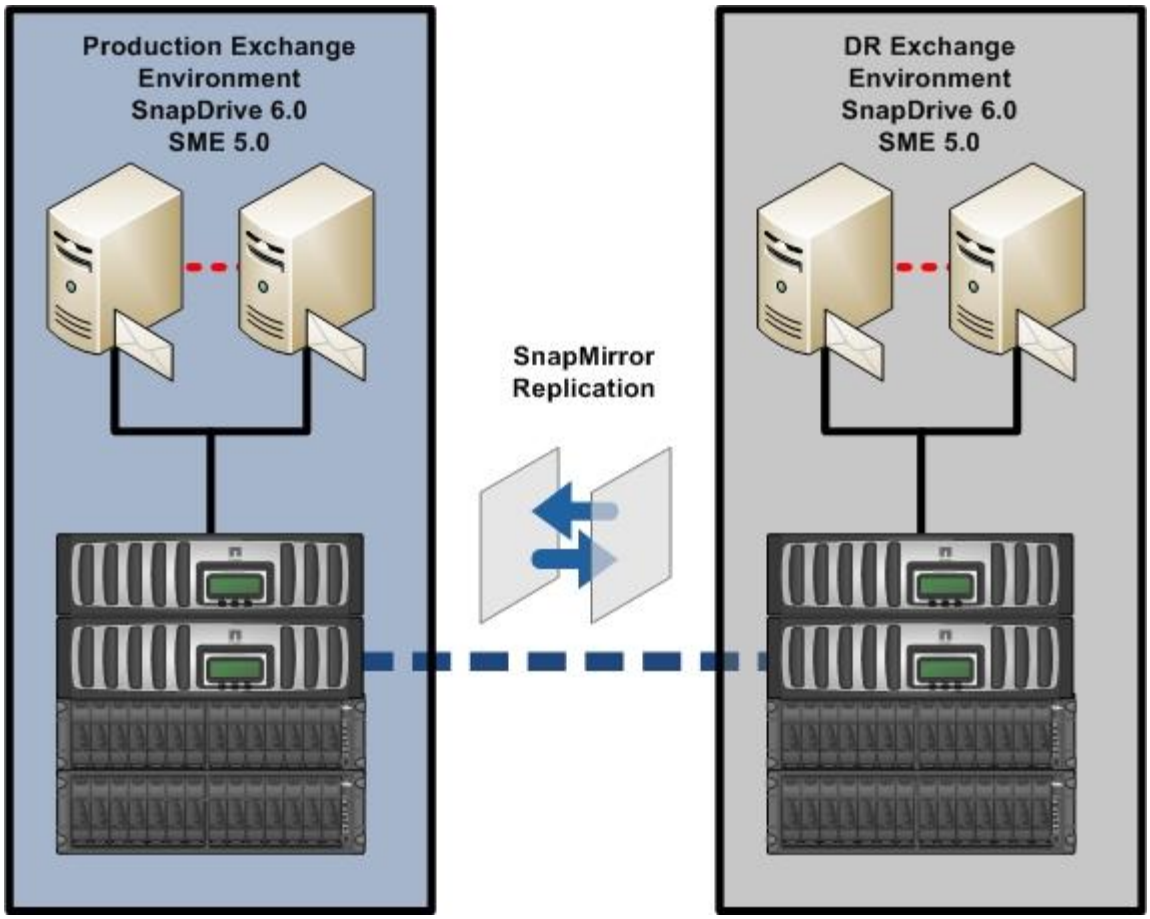

Figure 7) SnapMirror replication.

4.2 BUSINESS CONTINUANCE MODULE

Maximizing Exchange availability is a core component of a well-planned and carefully deployed BC strategy. An Exchange administrator must keep Exchange available 24x7 in order to meet the business needs of

today's work environment. When an Exchange Server does fail, the administrator must be able to recover from the failure in a fast and efficient manner.

SME 5.0 features a new BC module that enables Exchange administrators to automate the process of failing over to a secondary Exchange Server in the event the source Exchange Server fails. Once a BC plan is defined and verified, an Exchange administrator can simply click the Execute link, and the BC module will fail over the entire Exchange environment and restore service on the secondary Exchange Server.

The BC module allows Exchange administrators to fail over to identical configurations. For example, it is possible for an Exchange Server 2007 single copy cluster to fail over to a standby Exchange Server 2007 single copy cluster. Likewise, a standalone Exchange Server can fail over to standby standalone Exchange Server. The BC module also allows for local storage failures, without having to fail over the entire Exchange Server. Currently, the BC module does not support CCR clusters.

Best Practice

In Exchange Server 2003 environments, place all Exchange components onto NetApp LUNs. This includes the Exchange installation directory, the MTA/SMTP queues, and the MSSEARCH directory. The LUN(s) containing these directories must be on the same volume as the transaction log/SnapInfo LUN(s). This helps make sure all required LUNs are regularly backed up and replicated to the secondary storage system.

Best Practice

Make sure the same installation paths are configured on both the production and standby Exchange environments.

The create plan wizard in the BC module is used to create a BC plan. The wizard steps the administrator through the process of specifying the Exchange Servers to protect, the SnapMirror relationships to use, the network configuration for the Exchange Servers, and setting a logical name for the BC plan. Once an administrator has created a BC plan, it is highly recommended that the plan be verified to make sure there are no errors or faulty configurations in the Exchange environment.

Best Practice

Verify the BC plan immediately after creating it. Then schedule regular verifications to make sure the plan is still valid for the Exchange environment. If the verification fails, address the issues in the environment or recreate the BC plan.

In Exchange Server 2007 environments, the BC module can rehome mailboxes as part of the recovery. Rehoming mailboxes should be used whenever user mailboxes are moved to another server and users access their mailboxes on the new Exchange Server. This option further automates the failover process, making the failover completely transparent to end users.

The BC module also has the capability to fail back to the original production Exchange Server utilizing the same BC plan. In order to perform the failback, the SnapMirror relationships must be resynchronized back to the original production storage. This can be accomplished in one of two ways:

- Resyncing the existing SnapMirror relationship: If the production storage system is still online, available, and has a common SnapMirror Snapshot copy, then the SnapMirror relationship can be resynced. During this resync process, only data that has changed since the SnapMirror Snapshot copy will be transferred back to the production storage system.
- Syncing a new SnapMirror relationship: If the production storage is lost and a new storage system is in place, a full sync of the data will be required. This will require a new SnapMirror relationship.

Best Practice

When failing back to the original production Exchange Server, make sure you select the Clean Up Destination task. This will help make sure all remnant Exchange configurations are cleaned up and prepared for failback.

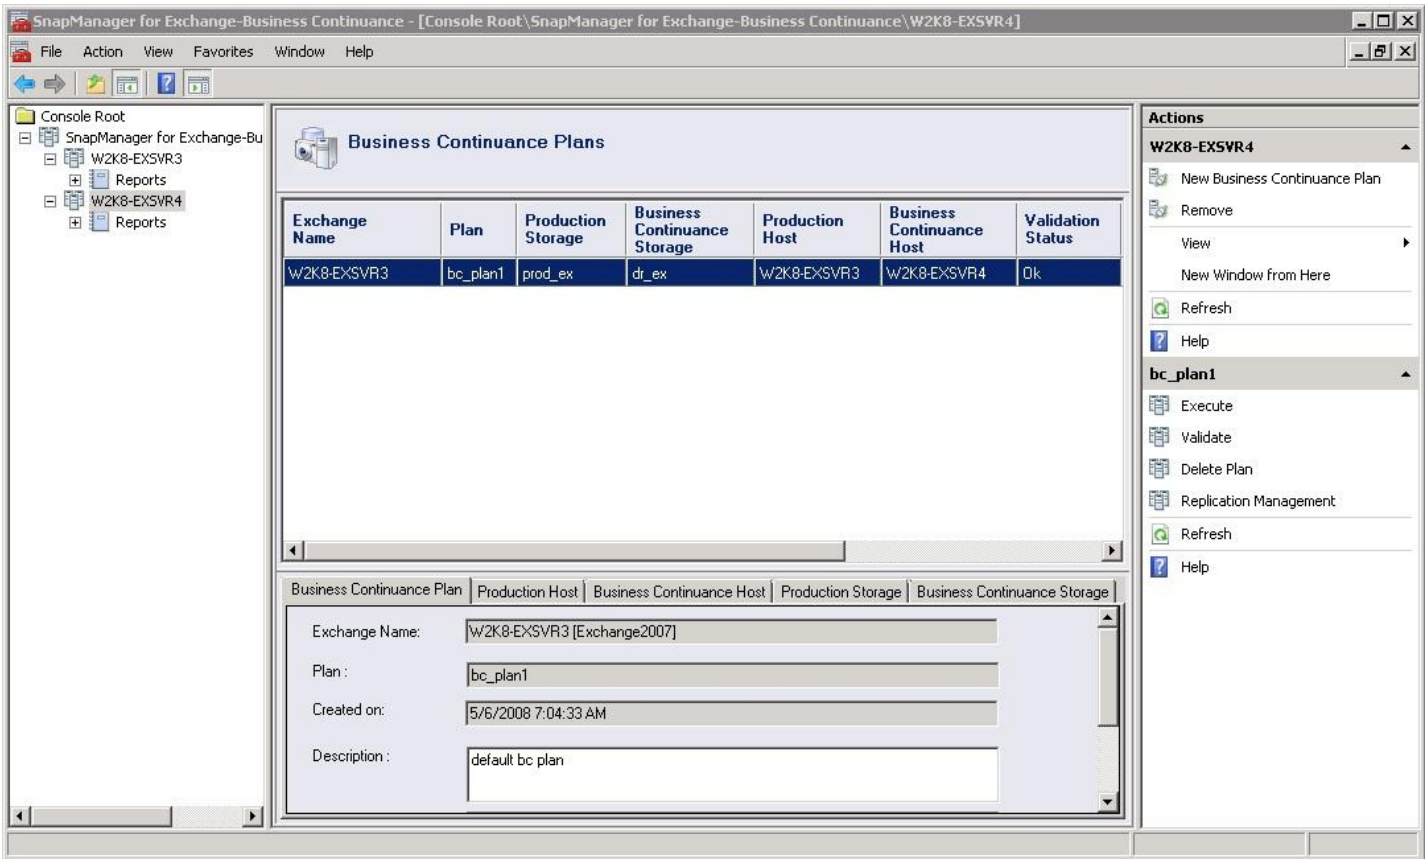

Figure 8) BC module interface.

5 ARCHIVING AND LONG-TERM DATA STORAGE

With e-mail containing business-critical data and with increasingly restrictive government data retention regulations, the need to archive and protect Exchange data is a must for most companies. Replication of that critical Exchange data onto a long-term storage solution must be fast, cost effective, and reliable. NetApp SnapVault meets these requirements by leveraging disk-based, block-level incremental backups to provide a low-overhead backup and recovery solution that is suitable for any environment.

5.1 DATA SET AND SNAPVAULT INTEGRATION

SME 5.0 integrates with SnapVault through NetApp Protection Manager. Providing global monitoring and alerting along with easy-to-use policies, Protection Manager simplifies common data protection tasks and helps automate SnapMirror and SnapVault management. Protection Manager manages operations using policies that define how data is protected. Policies are easily applied and updated across many systems and locations.

Best Practice

It is recommended that Protection Manager and the NetApp Management Console be installed on a dedicated server that is not an Exchange Server.

Through SME, a data set is created and assigned to a policy that has been defined in Protection Manager. The data set includes information such as storage system name, volume/qtree names, and LUN name for the storage groups which will be protected. Through Protection Manager, an administrator will assign a resource pool for the policy and choose the secondary storage that will host the archived data. Protection Manager will subsequently set up the relationships and apply the rules that are set forth in the policy that is assigned to the data set.

Once Protection Manager is aware of data sets created by SME, SME is able to protect new backup sets using the "`Archive local backup using SnapVault`" option in the backup window or backup wizard. Once the backup process has completed successfully, SME will communicate with Protection Manager, through SnapDrive, to archive the specified backup set. Protection Manager will identify the correct Snapshot copies on the storage system and update the SnapVault relationship accordingly.

Best Practice

If SnapVault relationships exist for volumes containing Exchange data that is protected with SME and Protection Manager, those relationships must be imported into Protection Manager. Failing this, new SnapVault relationships will be created for those volumes.

SME allows database verifications to be performed on the SnapVault destination volume. In order to take advantage of this feature and to utilize the secondary storage, a separate verification job must be scheduled after a successful backup operation completes with the archive option selected. The backup verification job will communicate with Protection Manager, through SnapDrive, all the necessary configuration information to perform the verification on the secondary storage. Once the verification job is complete, SME will mark the backup set on the primary storage as verified.

The data set and SnapVault integration with SME is supported on stand-alone Exchange Server and clustered Exchange Servers, both CCR-enabled clusters and single-copy clusters. In a clustered environment, an Exchange administrator has the option to create a data set on both nodes of the cluster. The data sets will be distinguished by appending the node name to the data set name that SME assigns.

The advantage of SnapVault integration with SME through Protection Manager is significant for an Exchange administrator. It simplifies archiving of backup sets created by SME on cheaper and larger secondary storage systems. You can archive more often and retain larger amounts of valuable Exchange data that needs to be retained and protected for extended periods of time. You can also manage that archived data using Protection Manager and monitor the environment using Operations Manager. The entire suite of products provides an administrator with a total data protection and archiving solution.

6 FLEXIBLE STORAGE OPTIONS

There are many factors that play into the decision-making process for buying storage systems. Because of that, various vendors might make up an entire Exchange solution. CCR can further complicate the storage solution since Microsoft recommends having separate storage for the active and passive nodes of the CCR cluster. Ideally, the storage for each node would reside on separate storage controllers. This could result in the need to have different vendor storage controllers on each node of a CCR cluster, which is now supported by SME 5.0.

An Exchange administrator can realize the benefits of SME 5.0, with its configuration, backup, and restore capabilities, while only having a single node of a CCR cluster on a NetApp storage controller. It does not matter which node (active or passive) is hosted by NetApp storage, but whichever node it is, that's the node where SME will run. So, if the active node resides on NetApp storage, SME will perform its actions on the active node. It will not be able to manage the passive node that is hosted by another vendor's storage system.

Best Practice

SnapDrive and SME must be installed on both nodes of the CCR cluster. SME will run from the node that is hosted by NetApp storage systems.

Best Practice

As required by CCR, both nodes of the cluster must have the same drive letters and file path locations. When deploying a heterogeneous configuration, make sure that both the third-party

storage and NetApp storage are capable of having identical paths on both nodes.

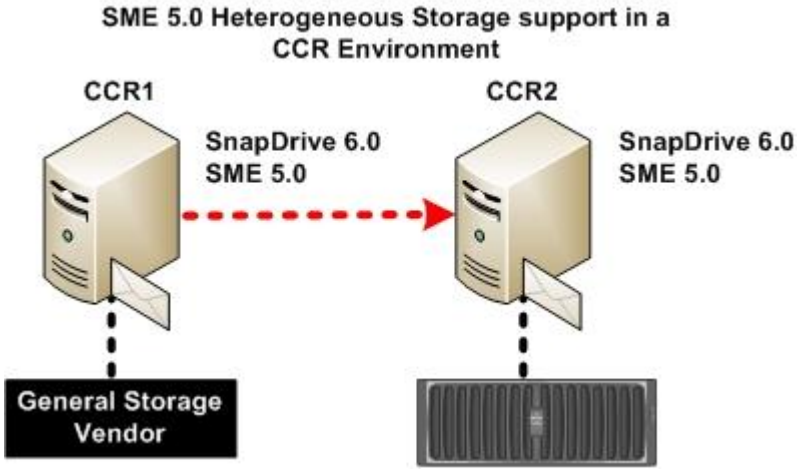

Figure 9) Heterogeneous storage.

7 SUMMARY

NetApp SnapManager 5.0 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.