



NetworkAppliance®

The evolution of storage.™

SnapManager® for Microsoft® Exchange 2000 Technical Overview

Network Appliance's Data Management
Solution for Microsoft Exchange 2000

Andrew McCumiskey | September 2002 | TR3198

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Abstract

Network Appliance provides a complete data management solution for Microsoft Exchange 2000 that addresses the entire spectrum of Exchange 2000 data management operations including data availability, disaster recovery, storage capacity, performance, ease of administration, backup, restoration, and configuration. This paper provides a technical overview of the technology, products, and features of Network Appliance's solution for Microsoft Exchange 2000.



Table of Contents



1) List of Figures and Tables	3
2) Introduction	3
3) Network Appliance's SnapManager for Microsoft Exchange 2000 Solution	4
Network Appliance Filers and Data ONTAP Software	4
SnapDrive 2.0	5
SnapManager for Microsoft Exchange 2000	5
4) SnapManager for Microsoft Exchange 2000 V1.1 Components	6
Exchange 2000 Data Storage Concepts	6
SnapManager for Microsoft Exchange 2000 Configurator Wizard	7
SnapManager for Microsoft Exchange 2000 Backup	9
SnapManager for Microsoft Exchange 2000 Restore	10
4) Conclusion	11

1) List of Figures and Tables

Figure 1 - SnapManager for Microsoft Exchange 2000 Architecture	4
Figure 2 - SnapDrive 2.0 Management Interface/MMC Plug-In	5
Figure 3 - SnapManager for Microsoft Exchange 2000 Main Menu	6
Figure 4 - Exchange 2000 Data Migration Steps with the Configurator Wizard	8
Figure 5 - SnapManager Backup Component	10
Figure 6 - SnapManager Restore Component	11

2) Introduction

The management of Microsoft Exchange 2000 information store data requires considerable diligence on behalf of today's Exchange administrators to ensure the proper operation, performance, and recoverability of Exchange 2000 data. Common concerns of Exchange administrators include a number of daunting tasks. Storage group and database volume layouts must be planned to ensure optimal performance. A plan must be defined and executed to provide backups of ever-growing Exchange 2000 information stores while at the same time considering shrinking backup windows. Perhaps most important, a strategy must be devised to provide rapid recovery of Exchange 2000 information stores in the event of a disaster or system failure.

Network Appliance provides a complete Microsoft Exchange 2000 data management solution that addresses all core data management functions. The solution leverages the Snapshot™, performance, and availability features of Network Appliance™ filers, the storage virtualization features of SnapDrive™ 2.0, and the integration of Microsoft Exchange backup and restore APIs in SnapManager for Microsoft Exchange 2000. Support of Microsoft Cluster Services and Network Appliance filer clusters are included. The solution can also leverage SnapMirror® capabilities of NetApp® filers to provide on-site or off-site SnapManager backup-set mirroring for disaster recovery.

This paper provides a technical overview of the SnapManager for Microsoft Exchange 2000 V1.1 solution and how it leverages NetApp filer capabilities and SnapDrive 2.0 to provide a complete Exchange 2000 data management solution.

3) Network Appliance's SnapManager for Microsoft Exchange 2000 Solution

Network Appliance's solution addresses all core Microsoft Exchange 2000 data management tasks: storage configuration and migration, backup, and restore. The physical storage and instantaneous Snapshots of Exchange data are delivered by Network Appliance file serving appliances and Data ONTAP™ software. The storage virtualization, logical disk management, and Snapshot operations are provided by Network Appliance's SnapDrive 2.0 storage virtualization software. Finally, the configuration, migration, and online backup and restore of Microsoft Exchange 2000 data are delivered by SnapManager for Microsoft Exchange 2000. Each product is tightly integrated with the others to provide Exchange 2000 with an easy-to-manage storage platform with the highest availability.




Figure 1 - SnapManager for Microsoft Exchange 2000 Architecture

To understand the benefits of deploying this solution, a quick overview of each Network Appliance product and the value it brings to Exchange 2000 deployments is beneficial.

Network Appliance Filers and Data ONTAP Software

Network Appliance filers provide highly scalable, high-performance, and reliable data storage for a wide variety of applications. Filer systems are disk storage arrays capable of scaling to multiterabytes of RAID protected data. NetApp filers run the Data ONTAP microkernel operating system, which implements the WAFL® (Write Anywhere File Layout) file system. Data ONTAP allows data to be

accessed and shared by a variety of file access protocols including CIFS, FCP (Fibre Channel Protocol), VLD (Virtualized Local Disk), NFS, HTTP, and FTP. Perhaps one of the most important features of Data ONTAP and the WAFL file system is the ability to take Snapshots. Snapshots are taken in seconds regardless of whether the data volume is 3GB or 3TB in size. Each volume can store up to 31 Snapshots and each Snapshot's data can be restored by a simple file copy, a volume-wide reset (SnapRestore[®]) or single-file reset (Single-File SnapRestore). Filers can also be configured with a number of options including data mirroring (SnapMirror) and clustering (Clustered Failover).

SnapDrive 2.0

Network Appliance's SnapDrive 2.0 software provides storage virtualization of filer volumes via the VLD or Fibre Channel file access protocols. The software allows administrators to define virtual disks that are presented to the Windows[®] 2000 operating system as basic logical disks. Once created, the virtual disks can be completely managed by the SnapDrive 2.0 MMC plug-in and the Windows 2000 Disk Administrator MMC plug-in.

SnapDrive 2.0 also provides integration between the Windows 2000 operating system and the filer to ensure consistent data Snapshots. Snapshots can be created from within the SnapDrive 2.0 MMC plug-in or from applications such as SnapManager for Microsoft Exchange 2000 that utilize the SnapDrive 2.0 Snapshot APIs.

In the event of a low disk space situation, SnapDrive 2.0 provides the ability to expand virtual disks on demand and instantly present the new disk space to Windows 2000. This capability assists in storage planning by allowing the administrators to add storage as needed.

Finally, SnapDrive 2.0 integrates with Microsoft Cluster Services to provide highly available storage solutions in clustered Microsoft Windows environments.

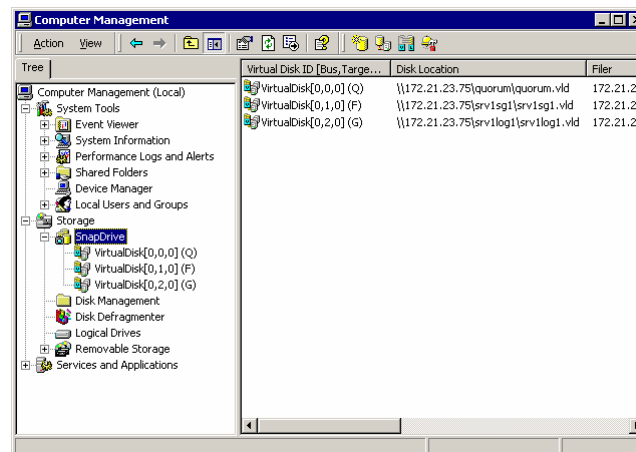


Figure 2 - SnapDrive 2.0 Management Interface/MMC Plug-In

SnapManager for Microsoft Exchange 2000

SnapManager for Microsoft Exchange 2000 integrates Microsoft Exchange 2000 backup and restore APIs with many of the features of Network Appliance filers and SnapDrive 2.0. The rapid Snapshot capabilities of the filer and the virtual disk management of SnapDrive 2.0 are leveraged to provide easy Exchange 2000 storage group configuration and migration, near-instantaneous backup of hundreds of gigabytes of data, and rapid restoration and recovery of that data. SnapManager for Microsoft

Exchange 2000 also includes an enhanced application-level verification function, which ensures that SnapManager backup-sets are capable of being restored successfully. Remote administration of SnapManager for Microsoft Exchange 2000 is accomplished by installing the product on a remote workstation. The remote administration feature allows manipulation of SnapManager backup and restore operations.

SnapManager for Microsoft Exchange 2000 does not completely eliminate the need for other archive technologies. SnapManager for Microsoft Exchange 2000 stores its Snapshots on the same volumes that contain the production Exchange data files. In the event of a disaster that disabled the use of a volume, restoration from another archive technology would be required. Complementary archive technologies such as the use of NetApp's SnapMirror data mirroring technology, NetApp's NearStore™ near-online data storage product line, and tape backup technology are valid options for disaster recovery of SnapManager for Microsoft Exchange 2000 backup-sets.

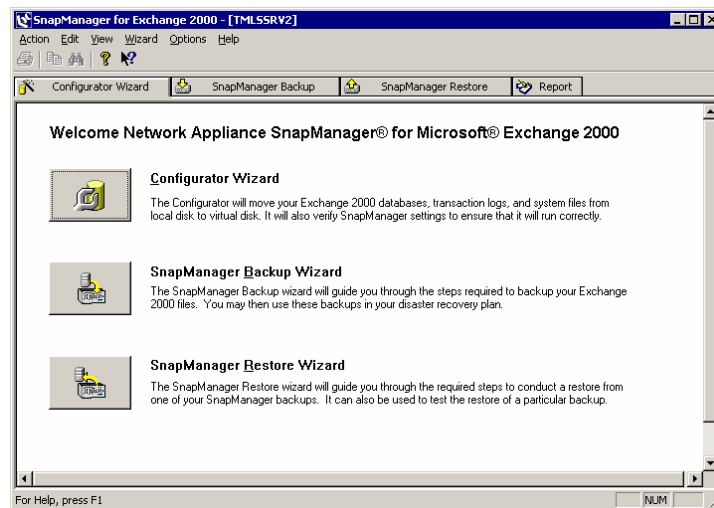


Figure 3 - SnapManager for Microsoft Exchange 2000 Main Menu

4) SnapManager for Microsoft Exchange 2000 V1.1 Components

A complete description of the technical aspects of SnapManager for Exchange 2000 would not be complete without a brief discussion of Exchange 2000 data storage concepts. The following section will provide a fundamental understanding of Exchange 2000 storage concepts to help provide a context for the functions and features within SnapManager for Microsoft Exchange 2000.

Exchange 2000 Data Storage Concepts

The primary goals for Exchange 2000 administrators that design Exchange storage system layouts are manageability, scalability, and recoverability while maintaining high performance. In order to accomplish these goals, Microsoft's Exchange development team has included the ability to partition data into smaller, more manageable data units called databases. Each Exchange 2000 database essentially contains a subset of a given Exchange server's mailboxes. Multiple databases may then be grouped together in a "storage group." All databases within a storage group share a common transaction log file set. Exchange 2000 Enterprise server allows for the deployment of a maximum of four storage groups, yielding the potential of having to manage up to 20 databases on a single Exchange 2000 server.

Behind this logical partitioning of Exchange 2000 information store data are a number of physical files. Each database consists of two physical files: the property store file and the streaming store file. The property store file has an .EDB file extension. This file contains the data for each user's mailbox, such as headers, message content, and attachments that are sent from other internal Exchange servers. The streaming store file contains raw file attachments that are sent to Exchange 2000 via Internet protocols such as SMTP. This file carries a .STM file extension.

The property and streaming store data files constitute the largest files found in an Exchange 2000 server. These files may grow to be tens or hundreds of gigabytes in size as Exchange users collect and send e-mail. Therefore, large disk volumes capable of supporting highly random access patterns must be available to ensure capacity, scalability, and performance.

The database files by themselves do not provide enough protection from system failures or transaction rollbacks. To address this need, Microsoft Exchange 2000 implements a transaction logging mechanism. The transaction logging mechanism uses a write-ahead algorithm, ensuring every transaction is logged to a transaction log file set before being considered for addition into the property or streaming stores. If a system failure occurs before a transaction is completely written to the data store, the transaction log file set provides the information to complete the transaction once the system failure is corrected.

Each transaction log file set has a direct correlation to a single storage group. Since each storage group has the ability to support up to five databases, the transaction log files contain interleaved log entries from all databases within the storage group. These interleaved entries are stored in a set of 5MB files with a .LOG file extension. An Exchange administrator can also assign a file name designation code at the beginning of each transaction log file to ease the process of identifying a particular storage group's transaction log file set.

Exchange 2000 provides two options for logging: circular logging and noncircular logging. The use of circular logging in enterprise deployments is highly discouraged. Circular logging allows Exchange to overwrite previous log files once the data within the log files has been committed to the data files. Circular logging may significantly reduce the amount of disk space used by the transaction logs. However, the use of circular logging prevents full recovery in the event of a complete disaster.

The ability to completely recover Microsoft Exchange 2000 in the event of a system failure or virus attack is critical. Therefore, enterprise deployments of Microsoft Exchange 2000 always choose noncircular logging for mailbox data stores. The use of noncircular logging provides the highest recoverability but also consumes the most disk space.

Transaction log files are very rarely read by Exchange 2000 because they are only used in the event of a recovery. Because of their critical use in recovery, log files should be placed on RAID-protected or mirrored volumes capable of withstanding high sequential write activity.

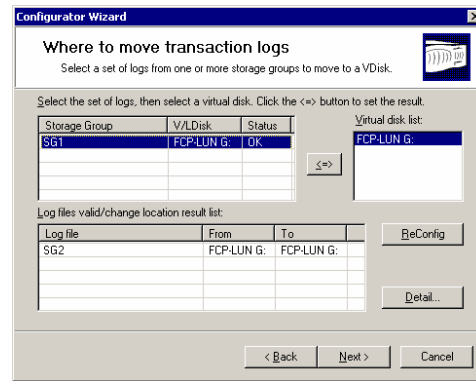
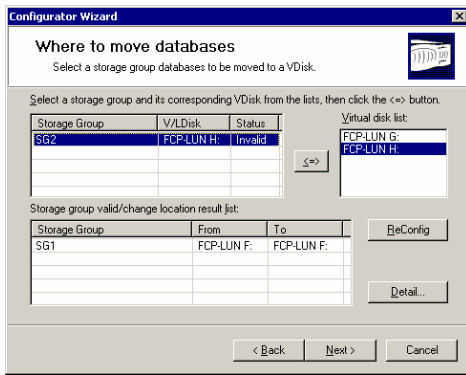
SnapManager for Microsoft Exchange 2000 Configurator Wizard

While the ability to partition Exchange 2000 yields a number of benefits, the Exchange 2000 administrator must realize that deploying a fully loaded Exchange 2000 server will result in 20 (4 storage groups multiplied by 5 databases each equals 20) different database entities that must be managed, backed up, and recovered. Of course, in order to take advantage of these features, an Exchange administrator must have a firm understanding of these data partitioning concepts and how to properly deploy them.

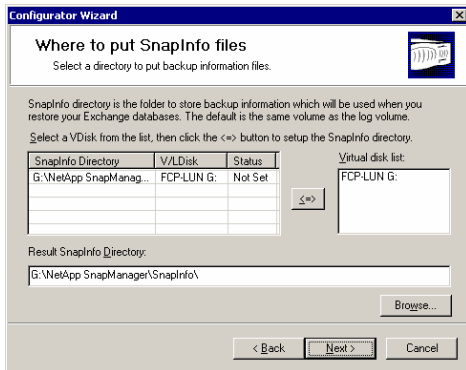
SnapManager for Microsoft Exchange 2000 includes the Configurator Wizard, which significantly eases the configuration and migration of Exchange 2000 information stores. This tool provides wizard-based navigation to assist an Exchange administrator through the steps of laying out and then migrating Exchange 2000 storage groups, databases, and transaction log file sets.

The Configurator Wizard is flexible enough to allow an Exchange 2000 administrator to deploy almost any storage architecture that can be imagined. The Exchange 2000 administrator is still required to properly define an appropriate storage layout for a particular deployment. However, the Configurator Wizard will make deploying and migrating that layout extremely easy.

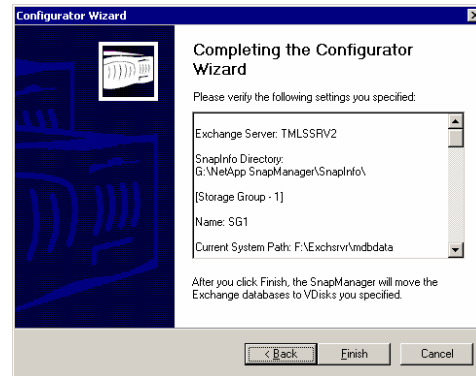
The screenshots below illustrates a complete Configurator Wizard session and shows the ease in which it configures and migrates Exchange 2000 data. The wizard breaks the process down into four steps: assign database files to volumes, assign transaction log file sets to volumes, define a location for SnapManager for Microsoft Exchange to store backup and restore operation files, and finally proceed with the migration operation.



Step 1: Database file location assignment



Step 2: Transaction log file location assignment



Step 3: SnapInfo directory location assignment

Step 4: Click Finish to migrate

Figure 4 - Exchange 2000 Data Migration Steps with the Configurator Wizard

SnapManager for Microsoft Exchange 2000 Backup

Exchange 2000 servers are deemed mission critical in many enterprises. Therefore, careful consideration must be taken when deciding on an appropriate backup strategy. Many deployments of Microsoft Exchange 2000 utilize backup technology that measures the completion of whole Exchange server backup jobs in terms of hours. While such deployments provide the benefits of being easy to deploy, they inevitably don't scale to information stores that grow to hundreds of gigabytes or aren't able of supporting shorter and shorter backup windows. Other deployments utilize copy-out mechanisms within the storage subsystem to provide quick backup jobs that provide higher scalability due to the shortened backup window. Unfortunately some of these solutions require that the Exchange server be taken offline while the backup is taken, causing an Exchange service outage.

SnapManager for Microsoft Exchange 2000 includes the SnapManager backup component, which provides nearly instantaneous online backups of Microsoft Exchange 2000 information stores. The component provides an easy-to-use interface for scheduled storage group backups, an enhanced backup-set verification feature that verifies backup sets at the application level, and a backup-set management feature that allows administrators to define and meet backup-set retention policies. Instead of backups completing in hours, whole Exchange 2000 server backups may now be performed in minutes and sometimes seconds thanks to NetApp's rapid Snapshot technology.

The operation of SnapManager Backup follows a particular sequence of events resulting in consistent Snapshots of Microsoft Exchange 2000 data files. The first step in the sequence utilizes native Microsoft Exchange 2000 online backup APIs to prepare the Exchange storage groups for backup. Once the storage groups are prepared, SnapManager Backup tells SnapDrive 2.0 to flush any outstanding NTFS buffers to the virtual disks that hold storage group data files. As soon as the NTFS buffers are flushed, SnapManager Backup places a request to SnapDrive 2.0 to conduct Snapshots of all of the data volumes containing Exchange 2000 data files. SnapDrive 2.0 receives the request from SnapManager Backup and coordinates and completes Snapshots of all filer volumes containing Exchange 2000 data files in a matter of seconds. After a successful return from the Snapshot request, SnapManager backs up the transaction log file sets to the SnapManager Snapinfo directory for future recovery. Upon completion of the log file copy, SnapManager Backup calls additional Microsoft Exchange 2000 online backup APIs to close the prepared storage groups and purge the committed transaction log files.

At this stage of the process, Exchange 2000 has been completely backed up. Any overhead associated with having the Exchange 2000 storage groups open in online backup mode is no longer a concern. The final step in the sequence of a SnapManager backup is verification of the Exchange 2000 data files that were recently backed up. SnapManager Backup will conduct a complete page-by-page integrity check of the storage group data files using Microsoft's ESEFILE integrity check utility. This utility will tell Exchange 2000 administrators that the backup sets are capable of being recovered not only from the data file level, but from the Exchange 2000 application level.

When properly configured, up to 31 SnapManager backup-sets can be taken and available for recovery at any point in time. Because the entire SnapManager backup process is typically measured in minutes, it is common to have SnapManager Backup run multiple times throughout the day to enhance protection and availability of Exchange 2000 data.

SnapManager backup-sets reside on the filer hosting the Exchange 2000 data files. In order to properly address archive and disaster recovery concerns, data mirroring technologies such as NetApp's SnapMirror, NearStore appliances, or tape backups should be taken of the SnapManager backup-sets

for off-site recovery. SnapManager for Microsoft Exchange 2000 fully supports SnapMirror for archiving backup-sets to off-site filers or NearStore appliances.

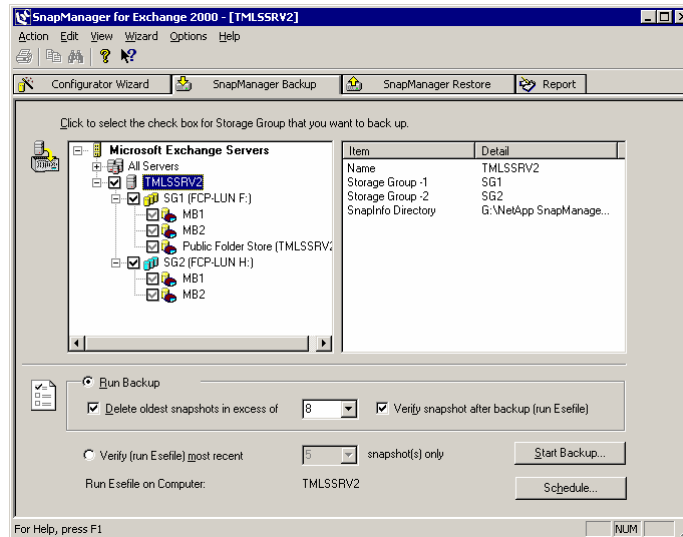


Figure 5 - SnapManager Backup Component

SnapManager for Microsoft Exchange 2000 Restore

The rapid restoration capabilities of the SnapManager Restore component of SnapManager for Microsoft Exchange 2000 provide Exchange administrators with the ability to deliver on today's stringent Exchange 2000 SLAs. By using all of the unique capabilities of NetApp filers, SnapDrive 2.0 and Microsoft Exchange 2000 restore APIs, SnapManager for Microsoft Exchange 2000 reduces the typical hours of waiting for a restore operation to complete to mere minutes. Recovery no longer requires finding all of the appropriate tapes in a differential or incremental backup archive. Gone are the days of waiting eight hours for a tape restore to complete, only to find that data corruption on the tape prevents recovery of the backup set.

The SnapManager Restore component provides an easy-to-use interface that allows selection and recovery of Exchange 2000 storage groups. The component is easy to navigate and requires minimal input from the Exchange administrator to ensure uptime is not limited because of product complexity.

SnapManager Restore provides the option of completing a point-in-time restore or an up-to-the-minute restore. Up-to-the-minute restores provide exactly that: The restore will start by restoring the data files to the time that the backup was completed and then roll forward every log file found up to the minute that Exchange 2000 experienced the system outage. A point-in-time restore will literally restore the data files to the point in time that the backup was taken and prevent further transaction logs from being played into the data files. The key consequence of the point-in-time restore method is any transactions that occurred after the backup job was completed will never be applied to the data files and are not recoverable. Therefore, up-to-the-minute restores are used at least 99% of the time.

There are a few instances where point-in-time restoration is necessary. The most common instance is when an Exchange 2000 data store is infected by a virus. In this case, doing an up-to-the minute restore will simply replay the virus infection back into the data files. If a SnapManager Backup was completed recently, then a point-in-time restore may be an option for an Exchange 2000 administrator.

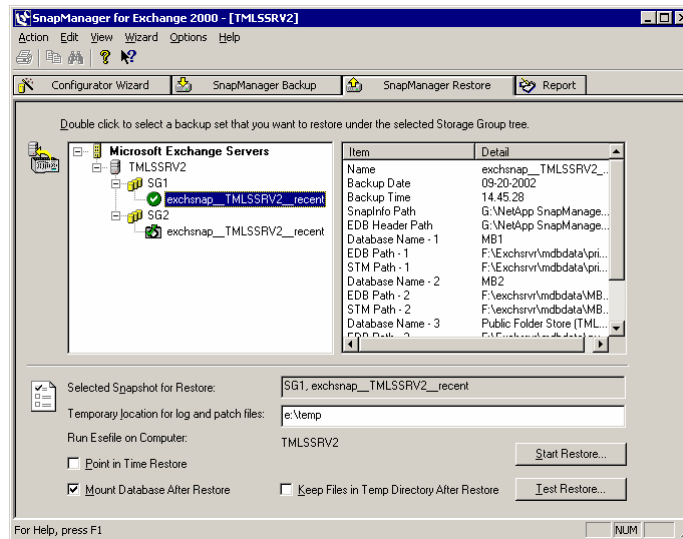


Figure 6 - SnapManager Restore Component

In order to restore a SnapManager backup-set, SnapManager Restore follows a sequence of steps to ensure successful recovery. The first operation SnapManager Restore completes is a full check of the storage group data and transaction log files to ensure the file signatures match and the transaction log files do not contain any gaps. The second step requires dismounting the storage group that will be restored if it has not already been dismounted. Next, Microsoft Exchange restore APIs are used to prepare the storage group for recovery. Once the storage groups are prepared, SnapManager Restore places a request to SnapDrive 2.0 to restore the Snapshots for the volume that contains the storage group data files (.EDB and .STM). SnapDrive 2.0 then coordinates and completes a SnapRestore of the virtual disk containing the .EDB and .STM files. The SnapRestore of a virtual disk usually completes in a few minutes. Upon notification of a successful restore of the virtual disk, SnapManager Restore locates the appropriate transaction log files necessary for restoration and copies them to the storage group's transaction log file directory. When the transaction log files are in the correct location, the Microsoft Exchange restore APIs are called to conduct a hard recovery of the Exchange 2000 storage group. The Exchange 2000 APIs verify the data files and then rolls the transaction log forward to the appropriate point in time based upon whether an up-to-the-minute or point-in-time restore was chosen.

SnapManager Restore also provides the ability to recover from tape or other archive media. In the event that the production filer is destroyed in a disaster, a command line option may be used to tell SnapManager Restore to use the files that have already been replaced from tape or other media. SnapManager restore fully supports the use of NetApp's SnapMirror data mirroring technology providing quick restoration of service to the original site over the network.

4) Conclusion

The SnapManager for Microsoft Exchange 2000 solution from Network Appliance provides the ability to easily manage Exchange 2000 servers and deliver unmatched availability. The solution is based upon tight integration between Network Appliance filers, SnapDrive 2.0 and SnapManager for Microsoft Exchange 2000. The benefits of deploying SnapManager for Microsoft Exchange 2000 include higher data availability, easier and more effective data management and rapid recovery in the event of a