9

Backup and Recovery Operations

Systems do not always run as smoothly as you would like. Hardware failures, software failures, human error, hacker attacks, and sometimes even natural disasters can disrupt your electronic mail (e-mail) environment. Routine hardware maintenance, disciplined system management, and educated users can reduce risk, but the potential for failures can never be eliminated completely. Disasters will happen, and you must be prepared to respond quickly. Regular backups are a key part of disciplined system management, and they protect data from accidental loss, hardware failures, and other disasters. You should regularly back up your Exchange databases and other critical files so that you can quickly restore them if data are accidentally deleted. If you accidentally delete data, you can recover a single database from the backup media. If a server fails, you can rebuild key components or the entire server.

The Exchange 2003 Information Store can be partitioned, allowing Exchange to overcome one of the more serious limitations of Exchange 5.5. Each Exchange 5.5 server was limited to a single large private database—stored as a single Windows NT file—containing the mailboxes for all users. The size of the Information Store grew in direct relation with the number of messages retained by the users, resulting in databases that exceeded many gigabytes in size. The large database size increased the time required to restore the database from backup tapes, which constrained the number of users you could safely put on a single server. Exchange 2003 solves this problem by allowing you to partition the Information Store (Figure 9.1).

In Exchange 2003, the Information Store for each Exchange server or cluster can be partitioned into up to four Storage Groups. Each of the Storage Groups can have up to five private or public database sets. This provides

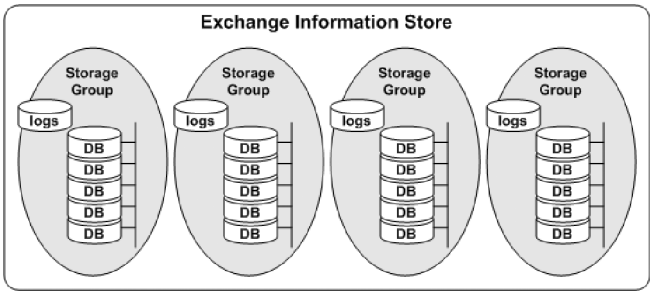

Figure 9.1
Exchange Information Store partitioning

a theoretical limit of 20 databases per server—four Storage Groups, each with five databases.

The databases are transaction based and fault tolerant. Each database change is recorded in a transaction log before the change is applied to the database itself. If a system failure occurs, the database recovery process uses the transaction logs to restore the changes that have occurred since the last successful backup. To reduce the overhead of multiple sets of log files, all database sets in the Storage Group share a common set of transaction log files. Partitioning results in smaller databases, reducing recovery time, and thus user impact.

From a backup and recovery perspective, each of the individual databases is independent. A database can be mounted or dismounted at any time, allowing a failed database to be restored while other databases remain operational. In other words, users with mailboxes in other databases can continue to send and receive e-mail during the recovery.

The changes to the Information Store since Exchange 5.5 was released affect the backup and recovery strategy for Exchange. When you install Exchange, the installation process extends the standard Windows backup utility to support the Exchange Information Store. This Windows "Exchange-aware" backup utility understands the relations among the Information Store, storage groups, databases, and transaction logs. The backup utility knows to delete the transaction logs after a successful normal (full) backup. A normal backup is the proper way to recover the disk space used by log files.

> **Note**: *If you elect to use third-party backup software, make sure that it will work with the Exchange Information Store, storage groups, databases, and transaction logs. Not all third-party products are Exchange-aware. Some third-party providers sell their Exchange-aware versions as add-on agents at additional cost.*

Exchange backups are designed to be done while Exchange is running. You do not—and should not—stop any Exchange services or dismount any Exchange databases when you do a backup. Because Exchange is still running, your users can continue to send and receive e-mail while the backup is in progress. Although you cannot have multiple instances of Windows Backup running simultaneously, Backup will allow you to select multiple databases to back up.

The Windows Backup utility supports several types of backup. Each type has advantages and disadvantages in terms of the time required to perform the backup, the amount of storage space needed on the backup media, and the time required to restore a database. Types of backups include the following:

- **Normal**. A normal backup is a full backup that copies all selected Storage Groups and databases, along with the associated transaction log files. After backing up the log files, the backup procedure merges pending transactions (messages) from the transaction logs into the Information Store and then deletes the log files from the disk. A normal backup is the proper way to recover the disk space used by log files. Because normal backups copy more data than other backup types, they take longer to complete. However, normal backups are strongly recommended because they minimize the number of tapes required to recover data. If you create daily normal (full) backups, you need only one tape to restore an Exchange database. Both incremental and differential backups require multiple tapes to recover the same amount of data. Because users cannot send or receive e-mail while a recovery is in progress, reducing the length of the recovery process is highly desirable.
- **Differential**. Differential backups are used in conjunction with normal backups. To use differential backups, you also must periodically make a normal backup. The differential backup is then used to copy the transaction log files that have changed since your normal backup. The database itself is not copied, and transaction logs are not deleted from the disk after being copied to tape. If you use differential backups, the recovery process requires your most recent normal backup tape and your most recent differential backup tape. Because this recovery process only requires two tapes, it is the second fastest recovery process.
- **Incremental**. Incremental backups also are used in conjunction with normal backups. The incremental backup copies the log files that have changed since your most recent normal backup or incremental backup. The database itself is not copied. If you use incremental backups, the

recovery process requires your most recent normal backup tape and each subsequent incremental backup tape. Because of the number of tapes involved, this is the slowest and most error-prone recovery method.

- **Copy**. A copy backup is the same as a normal backup, except that the transaction logs are not deleted from the disk at the end of the backup process. Copy backups are not the best method for restoring a database. They are most useful for taking a snapshot of the database.
- **Daily**. A daily backup backs up only files that have been changed that day, but it does not mark them as being backed up.

Using the Windows Backup utility, you can back up files and databases to a tape drive, a file on another hard drive, a removable disk, a CD-RW, or any other Windows storage device. The backup device must be directly connected to the system where you are running the Backup utility. If you are backing up Active Directory configuration data and the Windows registry, the backup device must be directly connected to the server you are backing up.

If you are backing up Exchange databases, the Exchange server can be anywhere on the network. This allows you to use a single backup server to back up the databases from several different Exchange servers. If you elect to do backups over the network, you may want to install a second network card in each of your Exchange servers and implement an isolated, high-bandwidth network just for backup traffic. This keeps the normal network traffic from slowing the backup and keeps the backup from affecting other network activity.

9.1 Minimizing risk

Hardware failures, software failures, human error, and sometimes even natural disasters can disrupt your e-mail environment. Disasters happen, and you must be prepared to respond quickly. By using the following practices, you can reduce the risk and impact of potential disasters.

- Ensure that circular logging is turned off for all Storage Groups. With circular logging enabled, transaction logs are overwritten to save disk space. However, overwriting transaction logs prevents the overwritten logs from being used during recovery operations.
- Perform daily full (normal) backups of the Exchange Information Store.
- Perform periodic full backups of Windows and Exchange configuration data.

- Select server-class hardware for your servers rather than high-end desktop systems. Redundant power supplies, multiple processors, and hardware RAID are worth the extra cost to ensure server availability in the event of a component failure.
- Install all Exchange servers in a controlled environment consistent with the manufacturer's recommendations. Protect the servers with Uninterruptible Power Supplies. Physically secure the environment that houses the servers.
- Protect databases using hardware RAID-5 (disk striping plus parity) or RAID 0+1 (disk striping and mirroring) technology.
- Keep transaction log files on separate hard drives from the databases. Protect the log files using RAID-1 (disk mirroring) technology.
- Keep the Windows operating system files on separate hard drives and protect them using RAID-1 (disk mirroring).
- Ensure that your Exchange servers have adequate disk space, including sufficient space to support recovery operations.
- Have multiple Windows domain controllers (DCs) for each domain to provide redundancy in the event of a single failure. Three DCs are recommended. If you only have two DCs, then you are at risk whenever you take one DC offline for maintenance. With three DCs, you are still protected if one of the DCs fails while you have one temporarily offline for maintenance.
- Maintain up-to-date documentation for your server configurations.
- Have a dedicated recovery server with the same configuration as your production servers. A dedicated recovery system is one that is only used when a disaster occurs. The server is not connected to the network.
- Fully document your recovery procedures and regularly practice disaster recoveries.

By following these practices, you can reduce the risk and impact of disasters, but you cannot completely avoid disruptions.

9.2 Preparing for disaster

If you have regularly backed up your Exchange databases and other critical files, you can quickly restore them if data become corrupted. If one of your servers fails or is physically damaged, the recovery process is more complex and requires

more preparation. You should prepare for this worst-case situation by creating and maintaining a disaster recovery toolkit containing the following items:

- A replacement server with the same configuration as the failed production server
- Windows installation CD-ROM
- Exchange Server installation CD-ROM
- All Service Packs and hot fixes that you have applied to the system
- An up-to-date full backup of your system drive (i.e., the drive where Windows is installed)
- An up-to-date full backup of the Windows System State; a System State backup includes the registry, Internet Information Server metabase, and COM + registrations
- An up-to-date full backup of the Windows and Exchange configuration data. Configuration data include settings for administrative groups, servers, security, and virtual servers. Configuration data are stored in the Active Directory and the registry
- An up-to-date full backup of the Exchange Information Store databases
- Written procedures for recovering a mailbox, restoring a database, and rebuilding an Exchange server after a disaster

Collecting and maintaining this list of CD-ROMs, backup media, and procedures is only the first step toward being prepared for disaster. The second, and equally important, step is to periodically practice recovering mailboxes, restoring databases, and rebuilding servers. In the midst of a disaster is not the time to be testing your procedures for the first time. When a disaster strikes, you should already be comfortable with the recovery process. Remember that your users cannot send or receive e-mail during the recovery process. Unless you enjoy responding to impatient users, you should do everything possible to ensure that the recovery process will be quick and painless. Having practiced the recovery process also will allow you to make confident predictions on how soon the server will be available.

9.3 Backing up the Exchange Information Store

Performing regular backups of the Exchange Information Store is an important part of creating a fault-tolerant messaging environment. You should schedule a daily normal (full) backup of the Information Store. Scheduling

the backup reduces the amount of human interaction and reduces the possibility that someone may forget to perform the backup. As with any type of backup, it is important that you always verify the success of the backup operation.

Exchange is designed to be backed up while it is running. You do not—and should not—stop any Exchange services or dismount any Exchange databases when you do a backup. Because Exchange is still running, your users can continue to send and receive e-mail while the backup is in progress. You can use the following procedure to schedule backups for an Exchange Information Store.

1. Start the Backup process from the Windows Start menu by selecting All Programs → Accessories → System Tools → Backup (Figure 9.2).
2. On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*
3. In the Backup Utility window, select the Backup tab (Figure 9.4).
4. Expand the Microsoft Exchange Server section to display the Exchange servers in your organization.

Figure 9.2 *Welcome to the Backup or Restore Wizard*

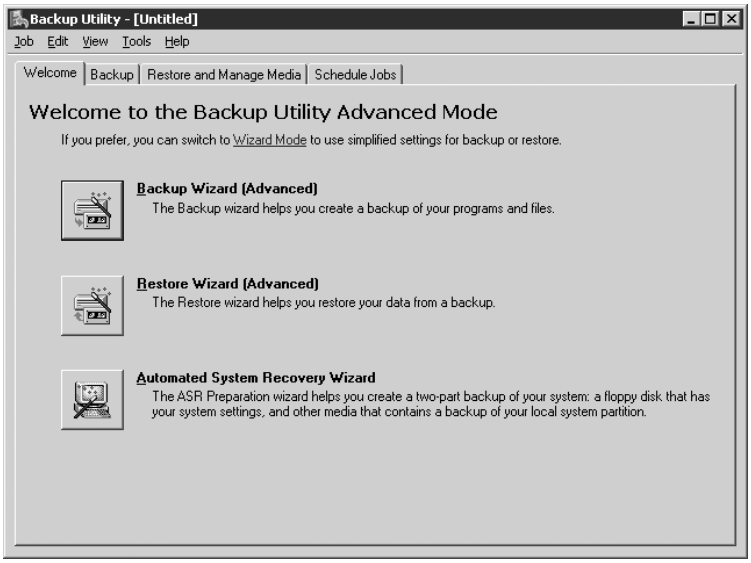

Figure 9.3
Backup Utility – Welcome tab

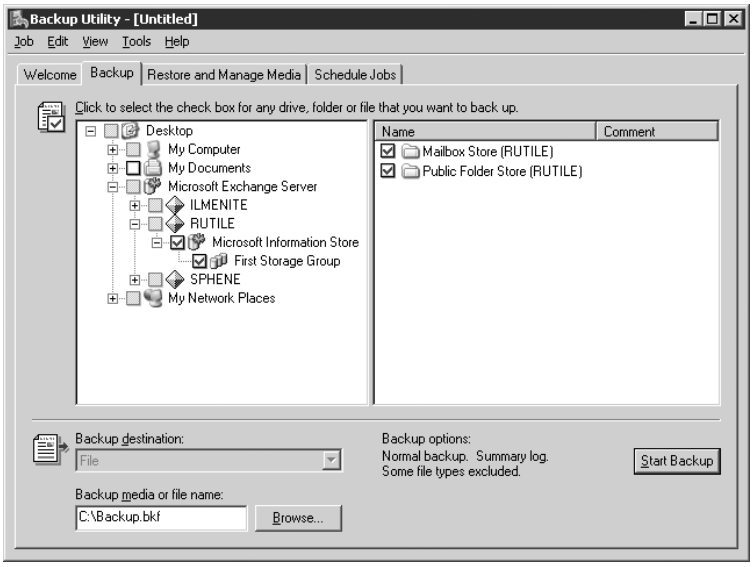

Figure 9.4
Backup Utility – Backup tab

5. Expand the server containing the Information Store you want to back up.
6. Expand the Microsoft Information Store section to display the Storage Groups contained within the Information Store.

7. You can back up the entire Information Store, selected Storage Groups, or selected databases. If you select multiple databases, Backup will write them to the backup media one after another.
   - Select Microsoft Information Store to back up all Storage Groups and databases within the Information Store.
   - Select a Storage Group to back up all databases within the Storage Group.
   - In the details pane, select a database to back up the single database.

   Because all databases in a storage group share the same set of transaction log files, you can improve the speed of the backup process by backing up an entire storage group at the same time.
8. Use the *Backup destination* drop-down list to select a backup device or file where the backup data will be written. If you do not have a backup device, the Backup Utility automatically selects File by default.
9. Use the Browse button to enter the file name for the new backup file into the *Backup media or file name* field.
10. Select the Start Backup button to display the Backup Job Information dialog box (Figure 9.5).
11. The *Backup description* field contains the default description of the backup set. You can change this description if necessary.
12. If the backup media or file contains a previous backup, select one of the following options:
    - Select *Append this backup to the media* to keep the previous backup and append the new backup.
    - Select *Replace the data on the media with this backup* to replace the previous backup.

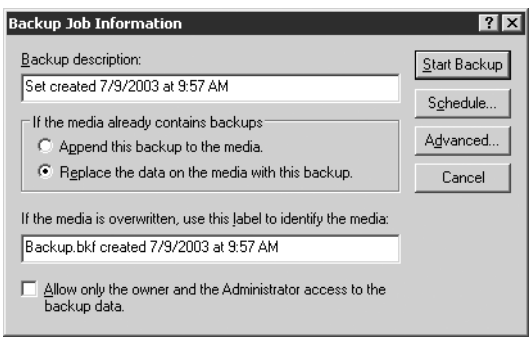

Figure 9.5 *Backup Job Information dialog box*

13. If you are creating a new backup (i.e., you are not appending this backup to a previous one), you can select the *Allow only the owner and the Administrator access to the backup data* check box to limit access to the backup.
14. The field near the bottom of the dialog box contains the default label that will be used to identify the media. You can change this description if necessary.
15. Select the Advanced button to display the Advanced Backup Options dialog box (Figure 9.6).
16. Use the *Backup Type* drop-down list to select the type of backup you want to perform. The choices are Normal, Copy, Incremental, Differential, or Daily. Normal backups (also known as full backups) are strongly recommended for two primary reasons:
    - Normal backups minimize the number of tapes required to recover data, thus minimizing the time required to recover the data. Both incremental and differential backups require multiple tapes to recover the same amount of data.
    - After backing up the transaction log files, the normal backup deletes the log files from the disk, thus recovering the disk space.

    Exchange continues to run and database changes can occur during the backup process. To capture these changes, the database engine maintains a patch file that logs these last minute changes. The backup utility writes the patch file to the backup media after copying the transaction log files.

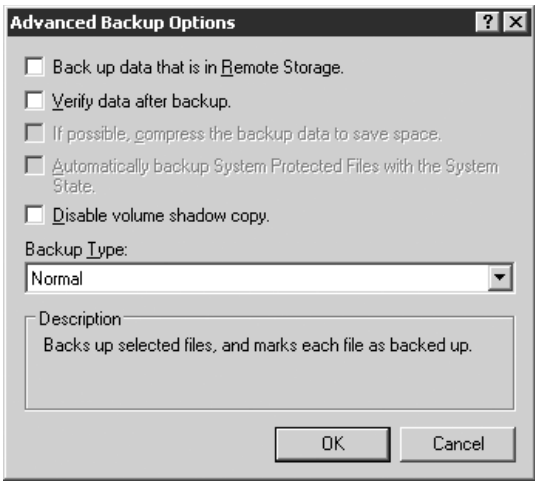

Figure 9.6
Advanced Backup Options dialog box

17. Select the *Verify data after backup* check box. Verification reads the backed up data to verify its integrity. This takes extra time, but it helps to ensure that you will be able to recover data from this backup media.
18. If you are using a tape drive capable of compressing the data, you can select the *If possible, compress the backup data to save space* check box. Compression allows you to store more data on the backup tape. However, you can only restore compressed backup tapes using drives that support the same type of compression.
19. Select OK to return to the Backup Job Information dialog box.
20. Select the Schedule button.
21. The Backup Utility will ask you to save the current backup job information and will display a Set Account Information dialog box asking for an account and password. The account will be used to run the backup job. Enter the account and password for the backup job security context. Once you have entered the account information, the Backup Utility will display the Scheduled Job Options dialog box (Figure 9.7).

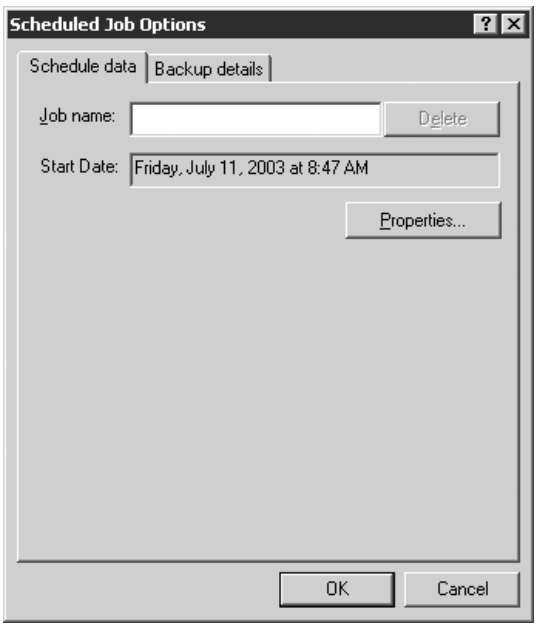

Figure 9.7 *Schedule Job Options dialog box*

22. In the *Job name* field, enter a name for this backup job.
23. Select Properties to display the Schedule Job dialog box (Figure 9.8).
24. You can use the *Schedule Task* drop-down list to elect to perform the backup just once, daily, weekly on selected days, or monthly on selected days of the month. For each of these options, you can specify the time when the backup should start. Optionally, you can select to perform the backup each time the system starts up, each time you log on to the system, or whenever the system is idle for a specified number of minutes. Because backups can affect server performance, you should schedule the backup for a time when there is moderate to low load. Select OK when you have selected the appropriate schedule.
25. The Backup Utility will display a Set Account Information dialog box asking for an account and password. The account will be used to run the backup job. Enter the account and password for the backup job security context. Once you have entered the account information, the Backup Utility will return to the Scheduled Job Options dialog box (see Figure 9.7).
26. Select OK on the Schedule Job Options dialog box. The Backup Utility will schedule the backup for the time you selected.

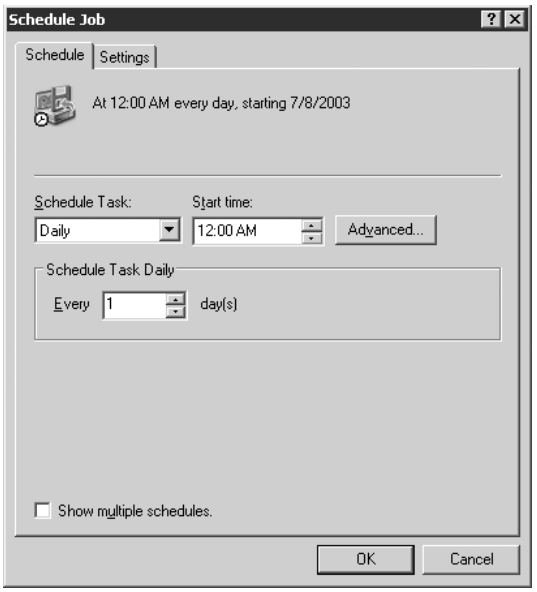

Figure 9.8
Schedule Job dialog box

9.4 Backing up configuration data

Exchange configuration data are stored in the Active Directory, the Windows registry, and the System State. Periodically, you should back up the system disk, any application disks, the Active Directory, the Windows registry, and the System State. In addition to periodic backups, it is especially important that you back up these data any time you make hardware or software configuration changes. These backups provide the capability for restoring the disks or individual files and are essential for rebuilding a complete server. The following procedure can be used to schedule backups for configuration data.

1. Start the Backup process from the Windows Start menu by selecting All Programs → Accessories → System Tools → Backup (see Figure 9.2).
2. On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (see Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*
3. In the Backup Utility window, select the Backup tab (Figure 9.9).

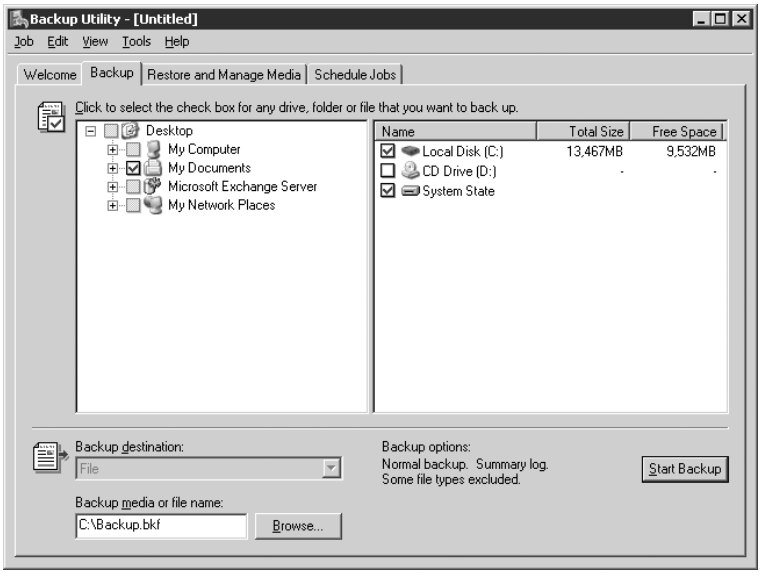

Figure 9.9
Backup Utility – Backup tab

4. Select the system you want to back up.
5. In the details pane, select the system devices you want to back up. Select the System State item to back up the Windows registry and the Active Directory settings.
6. Use the *Backup destination* drop-down list to select a backup device or file where the backup data will be written. If you do not have a backup device, the Backup Utility automatically selects File by default.
7. Use the Browse button to enter the file name for the new backup file into the *Backup media or file name* field.
8. Select the Start Backup button to display the Backup Job Information dialog box (Figure 9.10).
9. The *Backup description* field contains the default description of the backup set. You can change this description if necessary.
10. If the backup media or file contains a previous backup, select one of the following options:
    - Select *Append this backup to the media* to keep the previous backup and append the new backup.
    - Select *Replace the data on the media with this backups* to replace the previous backup.
11. If you are creating a new backup (i.e., you are not appending this backup to a previous one), you can select the *Allow only the owner and the Administrator access to the backup data* check box to limit access to the backup.
12. The field near the bottom of the dialog box contains the default label that will be used to identify the media. You can change this description if necessary.

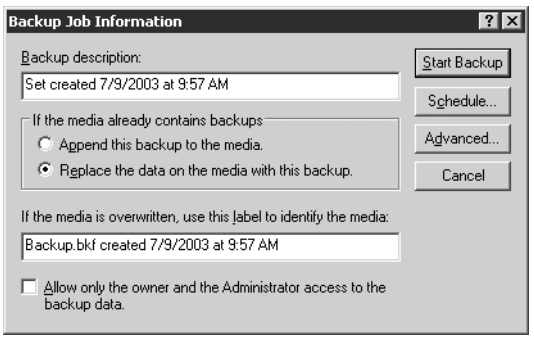

Figure 9.10 *Backup Job Information dialog box*

13. Select the Advanced button to display the Advanced Backup Options dialog box (Figure 9.11).
14. Use the *Backup Type* drop-down list to select the type of backup you want to perform. The choices are Normal, Copy, Incremental, Differential, or Daily. Normal backups are recommended.
15. Select the *Verify data after backup* check box. Verification reads the backed up data to verify its integrity. This takes extra time, but it helps to ensure that you will be able to recover data from this backup media.
16. If you are using a tape drive capable of compressing the data, you can select the *If possible, compress the backup data to save space* check box. Compression allows you to store more data on the backup tape. However, you can only restore compressed backup tapes using drives that support the same type of compression.
17. Select OK to return to the Backup Job Information dialog box.
18. Select the Schedule button.
19. The Backup Utility will ask you to save the current backup job information and will display a Set Account Information dialog box asking for an account and password. The account will be used to run the backup job. Enter the account and password for the backup job security context. Once you have entered the account information, the Backup Utility will display the Scheduled Job Options dialog box (Figure 9.12).
20. In the *Job name* field, enter a name for this backup job.

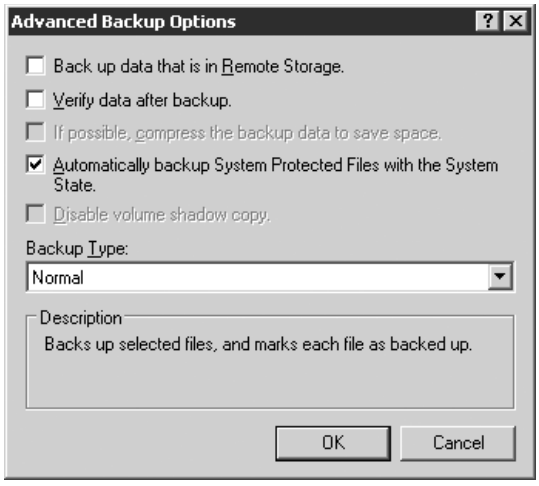

Figure 9.11
Advanced Backup Options dialog box

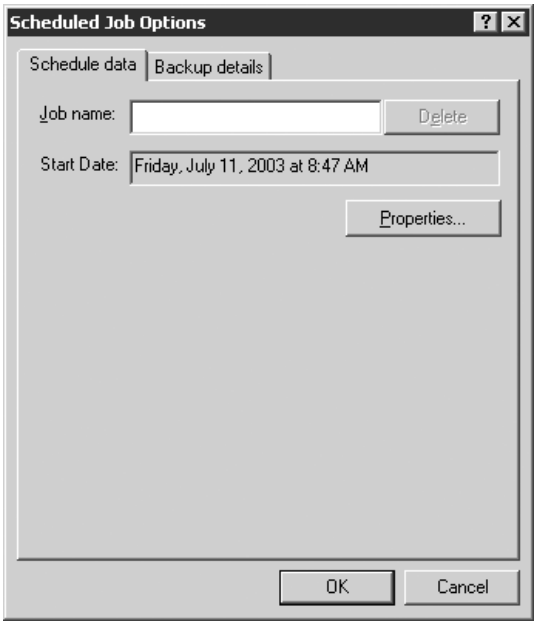

Figure 9.12
Schedule Job Options dialog box

21. Select Properties to display the Schedule Job dialog box (see Figure 9.8).
22. You can use the *Schedule Task* drop-down list to elect to perform the backup just once, daily, weekly on selected days, or monthly on selected days of the month. For each of these options, you can specify the time when the backup should start. Optionally, you can select to perform the backup each time the system starts up, each time you log on to the system, or whenever the system is idle for a specified number of minutes. Because backups can affect server performance, you should schedule the backup for a time when there is moderate to low load. Select OK when you have selected the appropriate schedule.
23. The Backup Utility will display a Set Account Information dialog box asking for an account and password. The account will be used to run the backup job. Enter the account and password for the backup job security context. Once you have entered the account information, the Backup Utility will return to the Scheduled Job Options dialog box (see Figure 9.12).
24. Select OK on the Schedule Job Options dialog box. The Backup Utility will schedule the backup for the time you selected.

9.5 Verifying backup success

Your ability to recover servers and restore data depends on the quality of your backups. The problem with backups is that they may sometimes fail, and this failure may go undetected. A series of unsuccessful backups leaves you unprotected against disasters and allows the log files to consume an ever-increasing amount of disk space. Therefore, it is important that you always verify the successful completion of the backup operation and that you verify that the backup media contain usable data.

You should always examine the Backup log to verify that all scheduled backups actually completed. You can view the Backup log using the following procedure.

1. Start the Backup process from the Windows Start menu by selecting All Programs → Accessories → System Tools → Backup (see Figure 9.2).
2. On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (see Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*
3. In the Backup Utility window, select Report from the Tools menu to display a list of the backup logs.
4. Double-click on the backup log you want to view.
5. If any errors are listed in the log, or if the backup did not complete successfully, the problem should be investigated immediately. Because the backup process accesses every page of the database, it is often the first process to discover a database corruption.
6. If the backup completes successfully, you should label the backup media and store it in a safe and secure location—preferably an off-site location.

Just because your backup processes regularly complete without error, do not assume that your backup media actually contain usable data. Tapes do not last forever. The usable lifetime of a tape should be available from the tape manufacturer. On rare occasions, bad tapes and malfunctioning hardware can produce unusable tapes, leaving you with an unwarranted sense of protection. At least once a month, you should verify the data integrity by restoring the data to your recovery server. Recovery testing also provides your support personnel with an opportunity to become familiar with the recovery procedure.

9.6 Modifying scheduled backups

You can use the following procedure to modify a scheduled backup job.

1. Start the Scheduled Tasks utility from the Windows Start menu by selecting All Programs → Accessories → System Tools → Schedule Tasks (Figure 9.13).
2. Double-click a scheduled job to view details about the job.

Task tab

3. The *Run* field contains the command that will be executed to perform the backup (Figure 9.14). The *Run as* field contains the security context under which the backup job will be run.

Schedule tab

4. Select the Schedule tab to view details about the backup job schedule (Figure 9.15).
5. You can use the Schedule tab to change the backup schedule. You can perform the backup just once, daily, weekly on selected days, or monthly on selected days of the month. For each of these options, you can specify the time when the backup should start. Optionally, you can select to perform the backup each time the system starts up, each time you logon to the system, or whenever the system is idle for a specified number of minutes.

9.7 Deleting scheduled backups

You can use the following procedure to delete a scheduled backup job:

1. Start the Scheduled Tasks utility from the Windows Start menu by selecting All Programs → Accessories → System Tools → Schedule Tasks (see Figure 9.13).
2. Right-click the job you want to delete and select Delete.

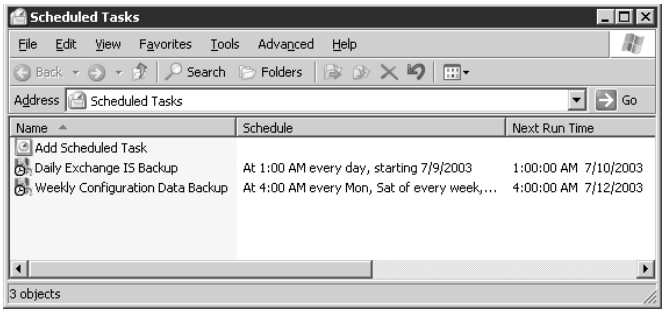

Figure 9.13 *Scheduled Tasks dialog box*

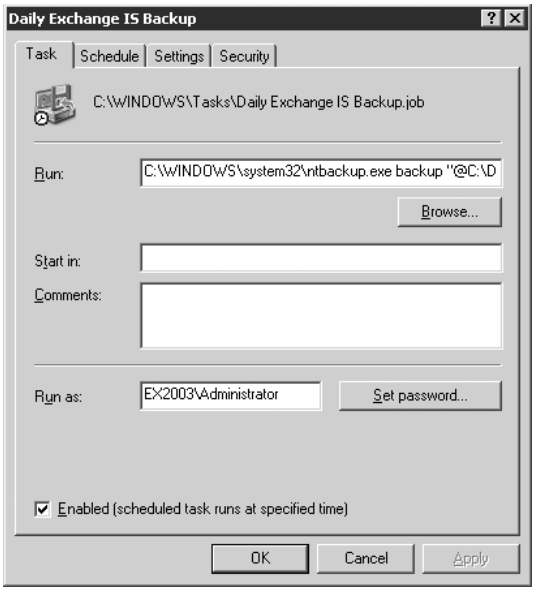

Figure 9.14
Scheduled Task – Task tab

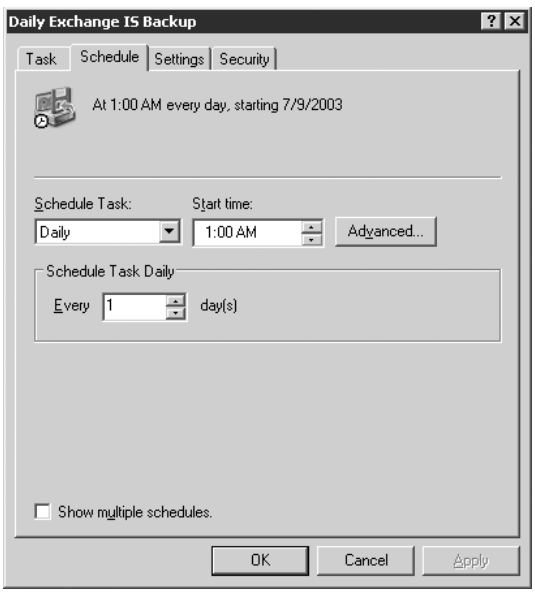

Figure 9.15
Scheduled Task – Schedule tab

9.8 Exchange 2003 and Volume ShadowCopy Services

The time required to back up an Information Store or to recover a lost or damaged Information Store is directly related to the amount of data to be copied and the speed of the backup device. Administrators are always seeking better ways to improve this process, and Storage Area Network hardware vendors have provided technologies—specifically snapshots and clones—that enable more rapid recovery. However, Exchange backups (regardless of whether they are normal backups, snapshots, or clones) require coordination with Exchange software. Exchange has always provided Application Programming Interfaces so that backup utilities could perform Exchange-aware backups. However, until Windows 2003 and Exchange 2003, Microsoft did not provide support for snapshots or clones. The Windows 2003 Volume ShadowCopy Service provides Storage Area Network and software vendors with the Application Programming Interfaces they need to develop complete Exchange-aware snapshot and clone solutions. Windows 2003 does not provide the snapshot and clone solution; it only supplies the Application Programming Interfaces needed to create the solution.

A snapshot is not a complete redundant copy of your Exchange Information Store and therefore does not provide protection for a lost or damaged Information Store. A snapshot is a metadata mapping and is designed to maintain a point-in-time view of the data in the snapshot.

On a disk volume for which you have created a snapshot, when a block of data is changed, the changed block is written to another location that is allocated from free volume pool space. The original (unchanged) block is maintained as part of the snapshot. In this manner, the original blocks that represent the point-in-time snapshot are preserved. As shown in Figure 9.16, the production data set consists of the original unchanged blocks remaining in the disk volume, plus the changed blocks. The point-in-time snapshot consists of the original blocks.

Clones are based on RAID 0+1 concepts. RAID 0+1 is a combination of striping (RAID 0), which interleaves data across multiple disks for better performance, and mirroring (RAID 1), which provides complete duplication of data. For example, Figure 9.17 shows a two-member RAID 0+1 set that consists of four disks mirrored to four disks.

You can create a clone by adding an additional member (or members) to this RAID 0+1 set and then separating one of the members from the set (Figure 9.18). The separated member (the clone) is a complete standalone copy of the data that you can use to restore an Exchange Information Store.

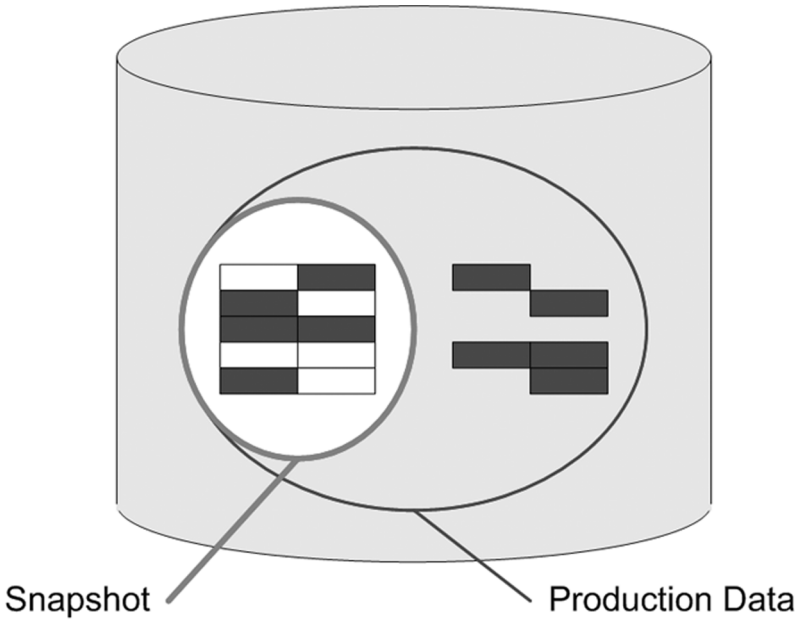

Figure 9.16 *Snapshot technology*

A complete backup/recovery solution using clone technology requires the Windows 2003 Volume ShadowCopy Service support, Storage Area Network hardware technology, and Volume ShadowCopy Service–aware (and Exchange-aware) backup software. The backup procedures will be specific to the hardware/software vendor.

9.9 Recovering a storage group or database

Restoring a storage group or database is sometimes necessary when a database becomes corrupt. You can restore your Exchange databases using your most recent full backup tape or backup tapes if you are using incremental or differential backups.

Restoring a database involves taking the corrupt database offline, replacing it with the good database from the backup media, replaying the transaction logs since the backup was taken, and then bringing the database back online. In Exchange 5.5, you had to stop all Exchange services before restoring the Information Store, meaning that all users with mailboxes on the failed server could not send or receive e-mail until the recovery process was

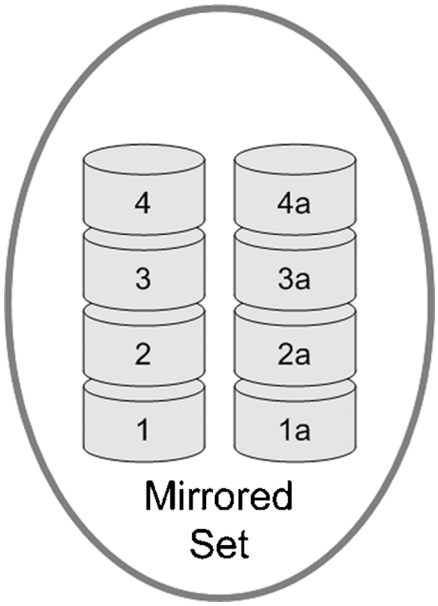

Figure 9.17
RAID 0+1 technology

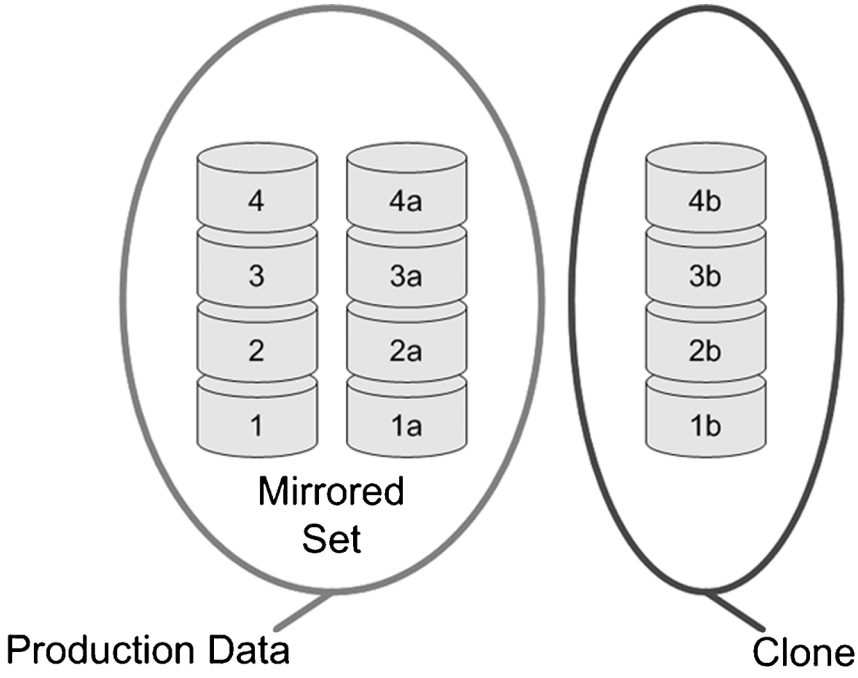

Figure 9.18
Clone technology

completed. With Exchange 2003, you do not—and should not—stop any Exchange services, and the only users affected are those with mailboxes in the corrupted database. Your other users can continue to send and receive e-mail while the recovery is in progress.

The following procedure can be used to recover an Exchange database to its original location. If you restore databases or log files to their original locations, any existing databases or log files are overwritten.

1. Verify that the Exchange services are running.
2. Use the following procedure to dismount the mailbox store you want to recover:
   - Start the Exchange System Manager (ESM) console from the Windows Start menu by selecting All Programs → Microsoft Exchange → System Manager.
   - Expand the Administrative Groups section. Expand the administrative group (e.g., First Administrative Group) that contains the server where the database is located. Expand the Servers section. Expand the server where the database is located. Expand the Storage Group where the database is located.
   - Right-click on the database and select All Tasks → Dismount Store. Select Yes when asked whether you want to continue. The dismount process may take a few minutes.
   - Find and mount the correct backup media.
3. From the Windows Start menu, select All Programs → Accessories → System Tools → Backup (see Figure 9.2).
4. On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (see Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*
5. In the Backup Utility window, select the Restore and Manage Media tab (Figure 9.19).
6. On the Restore and Manage Media tab, double-click the backup file containing the files you want to restore. Use the check boxes to select the data that you want to restore.
7. Select Start Restore to display the Restoring Database Store dialog box (Figure 9.20).

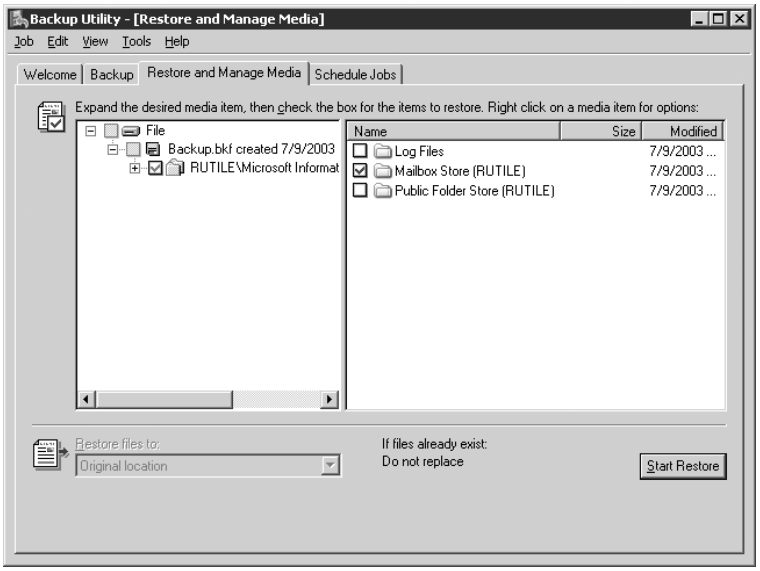

Figure 9.19
Backup Utility – Restore and Manage Media tab

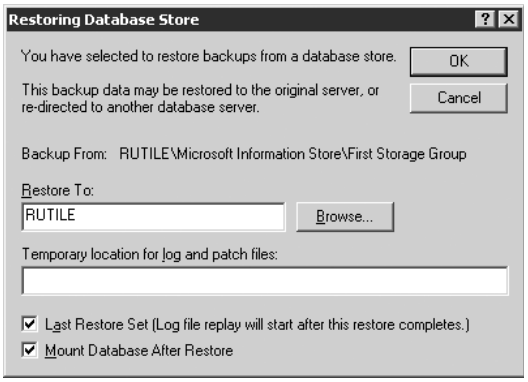

Figure 9.20
Restoring Database Store dialog box

8. Enter a directory name in the *Temporary location for log and patch files* field. This directory should be different from the one where the original log files are stored and should have sufficient disk space to store the files. During the restore process, Exchange will first apply the older transaction logs from the temporary directory and then apply the more recent logs from the original location.
9. The recovery procedure replays the transaction logs once all files have been written back to the disk. Select the *Last Restore Set* check box and

the *Mount Database After Restore* check box if any of the following conditions apply:

- You are restoring from a normal (full) backup without any incremental or differential backups
- You are restoring from the final incremental backup
- You are restoring from the final differential backup

10. Select OK to begin restoring the database. The recovery process will copy the database from the backup media. If the transaction logs recorded since the backup was taken are still intact (i.e., they were not affected by the database corruption), then the recovery process will replay these recent transaction logs to bring the database back to the state it was in when the corruption occurred. No data will be lost. However, if the recent transaction logs are also corrupted or unavailable, then the recovery procedure restores the database back to its current state.

9.10 Recovering a deleted mailbox

Sometimes system administrators make mistakes by erroneously deleting a mailbox. To recover a deleted mailbox on an Exchange 5.5 server, the administrator had to restore the entire mailbox store on a recovery server, export the mailbox contents to a .PST file, and import the mailbox contents back into a newly created mailbox. The process is easier for Exchange 2003 (and Exchange 2000) because the administrator can set a deleted mailbox retention period (the default period is 30 days) as a mailbox store property. When you delete a mailbox, Exchange hides the deleted mailbox and keeps its contents in their original mailbox store until the deletion period expires. As long as the deleted mailbox retention period has not expired and you have not deleted the associated user object, you can use the following procedure to easily reestablish the connection between the mailbox and the user.

1. Start ESM from the Windows Start menu by selecting All Programs → Microsoft Exchange → System Manager.
2. Expand the Tools section.
3. Right-click on Mailbox Recovery Center and select Add Mailbox Store.
4. Enter the name of the mailbox store that contains the deleted mailboxes and select OK. The Mailbox Recovery Center will display a list of all deleted mailboxes from the selected mailbox store (Figure 9.21). The deleted mailboxes are all marked with a red circle and an "x."

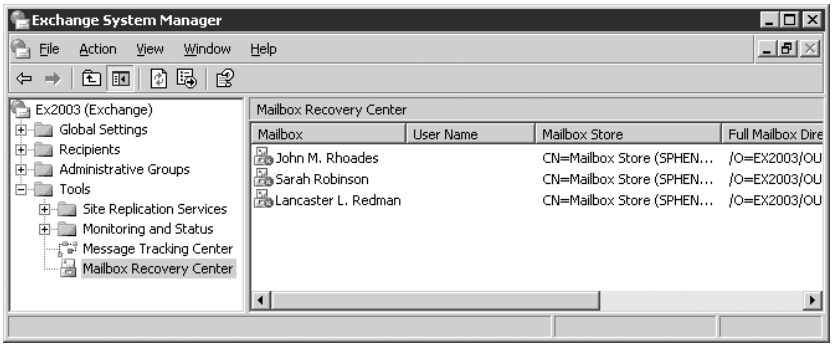

Figure 9.21
Mailbox Recovery Center

(The Mailbox Cleanup Agent, which runs nightly as part of normal background maintenance, marks the deleted mailboxes.)

5. In the details pane, right-click on one of the deleted mailboxes and select Find Match to start the Exchange Mailbox Matching Wizard.
6. Select Next on the Exchange Mailbox Matching Wizard welcome window. The wizard immediately begins searching for the associated user account and displays the completion window (Figure 9.22) when it has found a match.

Figure 9.22
Exchange Mailbox Matching Wizard

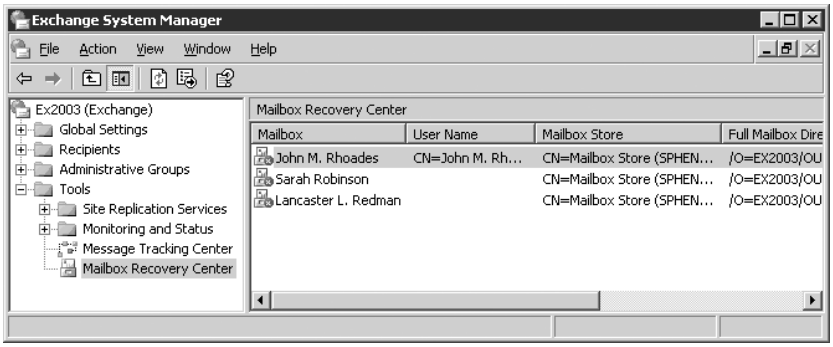

Figure 9.23 *Mailbox Recovery Center – User account added*

7. Select Finish to accept the user account found by the wizard. The Mailbox Recovery Center adds the user account as shown in Figure 9.23.
8. In the Mailbox Recovery Center, right-click on the matched mailbox and select Reconnect to start the Exchange Mailbox Reconnect Wizard to relink the deleted mailbox with its user account.
9. Select Next on the Exchange Mailbox Reconnect Wizard welcome window.
10. In the *Ready to proceed* window (Figure 9.24), select Next to reconnect the deleted mailbox with its user account.

9.11 Recovering deleted messages and mailboxes from backup media

Sometimes users delete important messages (or their entire mailbox). Subsequently, they may come to you to recover the deleted items from the backup media. Unfortunately, retrieving a single mailbox or a single message is not a simple process. Thankfully, there are often ways to avoid having to recover the data from the backup media.

When an Outlook user deletes a message, the message is moved into the user's Deleted Items folder. Often, the user can recover the missing items by looking in this folder. Even after the Deleted Items folder has been emptied, it may be possible to recover the deleted items without using the backup media. You can configure Exchange 2003 so that it does not immediately delete mail or mailboxes for a specified number of days. This period is known

Figure 9.24
Exchange Mailbox Reconnect Wizard

as the deleted item retention period, and you can tailor the duration to meet the needs of your organization. If the missing item was deleted within the deleted item retention period, the user can recover the item without your assistance. Most requests to restore user mailboxes or individual messages can be avoided if users are familiar with the Deleted Items folder and the deleted item retention period.

Before Exchange 2003, restoring a mailbox or a single message from backup media required that you restore the Information Store to a different Exchange server to avoid affecting other users on the production Exchange server. Exchange 2003 makes this process much easier. The procedure for Exchange 2003 uses the following major steps:

- Create a Recovery Storage Group (RSG).
- Enable the RSG.
- Restore the backed up database into the RSG.
- Extract selected messages from the recovered database.

- Merge the recovered messages into the production database.

The details for these major steps are described in the following sections.

Create Recovery Storage Group

1. Start ESM from the Windows Start menu by selecting All Programs → Microsoft Exchange → System Manager.

   Note: *By default, administrative groups and routing groups are not displayed. If you have not already enabled these, right-click on the Exchange organization and select Properties to display the organizational properties. Select the* Display administrative groups *check box to allow the administrative groups to be displayed, and select the Display routing groups check box to display the routing groups. You must restart ESM after enabling display of administrative groups and routing groups.*

2. Expand the Administrative Groups section.
3. Expand the administrative group (e.g., First Administrative Group) that contains the server where the deleted mailbox was located.
4. Expand the Servers section.
5. Right-click on the server where the deleted mailbox was located and select New → Recovery Storage Group. The RSG is a new feature with Exchange 2003. The RSG is simply a storage group that provides a context within the production Exchange organization for recovering individual items (e.g., mailboxes, folders, messages) or entire databases from backup media. You can create one RSG on each server, even for those servers that already have the maximum four storage groups. Because of the RSG, you no longer need a separate recovery server.
6. In the RSG Properties dialog box (Figure 9.25), select OK to create the RSG.

Enable Recovery Storage Group

7. To use the RSG, you must enable an Exchange database that will be associated with the RSG. To enable the RSG, right-click on the newly created RSG in ESM and select *Add Database to Recover*.
8. You can only use the RSG to recover the database you associate with the RSG. In the *Select database to recover* dialog box (Figure 9.26), select the database to be recovered and then select OK. ESM will display the Mailbox Store Properties dialog box (Figure 9.27).

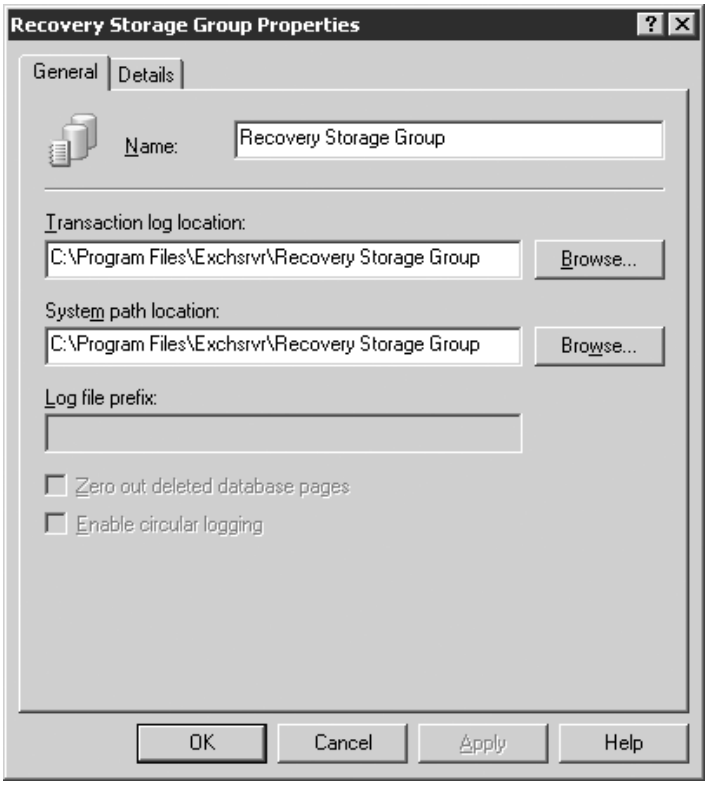

Figure 9.25
Recovery Storage Group Properties

9. Select OK in the database properties dialog box. ESM will list the unmounted database under the RSG (Figure 9.28).
10. Right-click on the newly created unmounted database and select *Properties*. Select the Database tab (Figure 9.29).
11. On the Database tab, select the *This database can be overwritten by a restore* check box and then select OK.

Restore backed up database into Recovery Storage Group

12. From the Windows Start menu, select All Programs → Accessories → System Tools → Backup. On the Backup or Restore Wizard welcome window, select the *Advanced Mode* hyperlink to start the Backup Utility.

    Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*

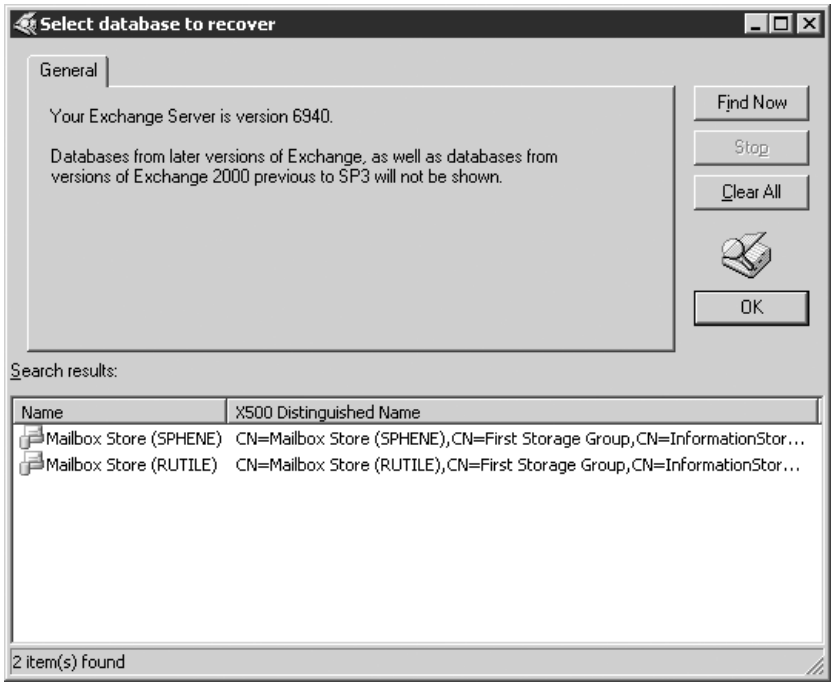

Figure 9.26
Select Database to Recover dialog box

13. In the Backup Utility window, select the Restore and Manage Media tab (Figure 9.30).
14. On the Restore and Manage Media tab, double-click the backup file containing the files you want to restore. Use the check boxes to select the data that you want to restore.
15. Select Start Restore to display the Restoring Database Store dialog box (Figure 9.31).
16. Enter a directory name in the *Temporary location for log and patch files* field. This directory should be different from the one where the original log files are stored and should have sufficient disk space to store the files. During the restore process, Exchange will first apply the older transactions logs from the temporary directory and then apply the more recent logs from the original location.

    You will notice that the recovery procedure does not ask where it should restore the backup files. When you have enabled the RSG, the recovery process always restores backups to the databases located in the

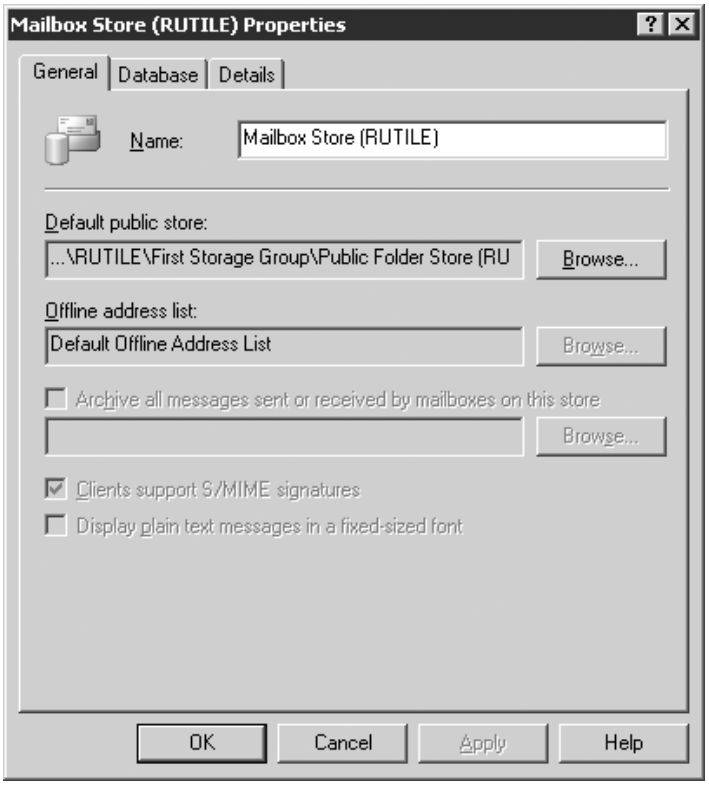

Figure 9.27
Mailbox Store Properties dialog box

RSG. To restore directly to a production database, you have to delete the RSG or set an undocumented registry key to bypass the RSG.

17. The recovery procedure replays the transaction logs once all files have been written back to disk. Select the *Last Restore Set* check box and the *Mount Database After Restore* check box if any of the following conditions apply:
    - You are restoring from a normal (full) backup without any incremental or differential backups.
    - You are restoring from the final incremental backup.
    - You are restoring from the final differential backup.
18. Select OK to begin restoring the database. The recovery process will copy the database from the backup media.
19. When you restart ESM, you can expand the RSG and the recovered database to view the mailboxes (Figure 9.32). Verify the contents of the restored database.

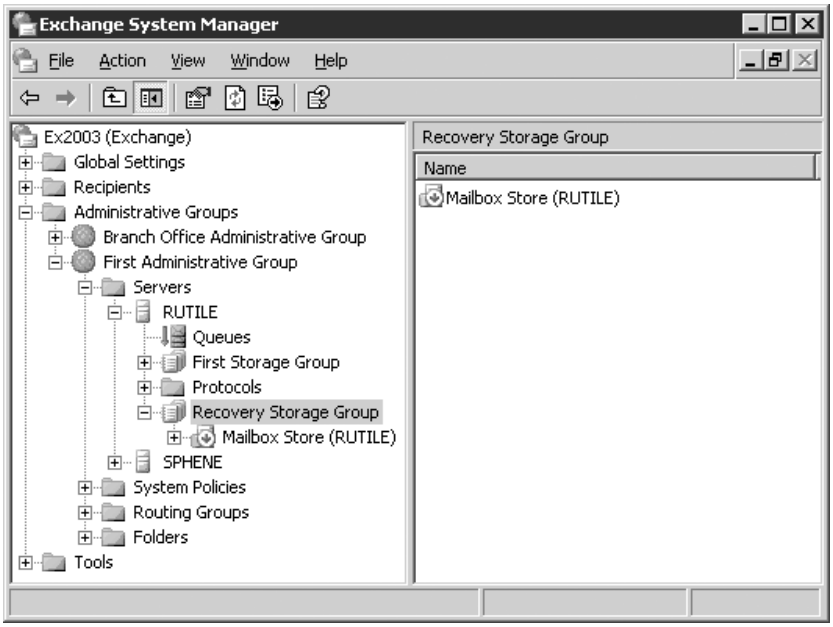

Figure 9.28
Unmounted recovery database

Extract selected messages from recovered database

20. Start the Exchange Mailbox Merge Wizard (ExMerge.exe). ExMerge allows you to extract data from mailboxes in one Exchange database and then merge the extracted data into mailboxes in another Exchange database. When combined with the Exchange 2003 RSG and backup tapes, ExMerge allows you to recover data, such as deleted messages or deleted mailboxes, that are no longer in your deleted item recovery area. ExMerge can use a significant percentage of the processor; therefore, whenever possible, you should avoid running ExMerge on production Exchange servers.

    Note: *ExMerge.exe is located in the \support\utils\i386\ directory on the Exchange 2003 CD. You must copy Exmerge.exe and Exmerge.ini to the C:\Program Files\exchsrvr\BIN folder.*

21. In the ExMerge welcome window, select Next to continue.

22. In the Procedure Selection window (Figure 9.33), select *Extract or Import (Two Step Procedure)* and then select Next. The two-step procedure creates intermediate personal stores (.PST files) in the first step and then merges the .PST file data into the destination production store during the second step.

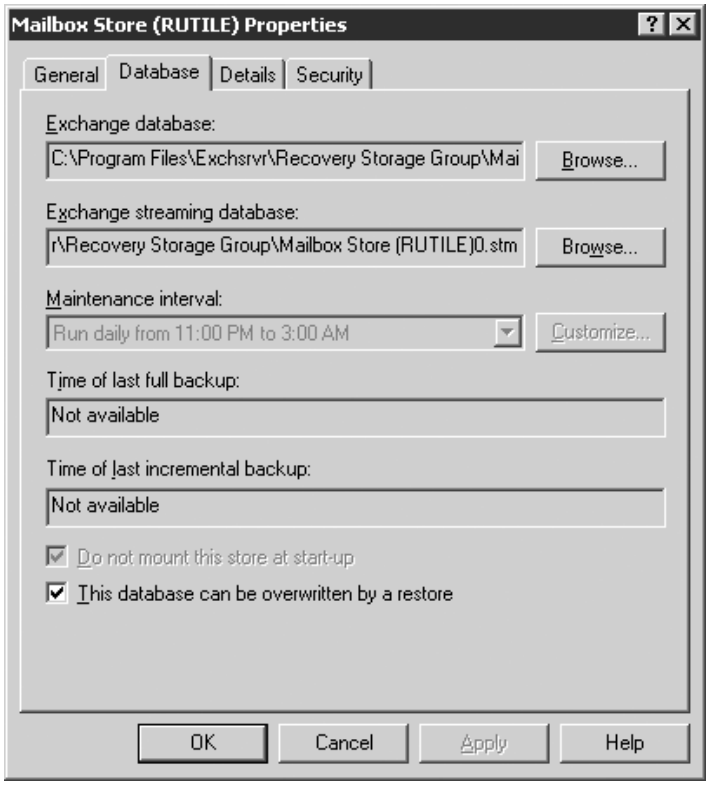

Figure 9.29
Mailbox Store Properties – Database tab

23. In the Two Step Procedure window (Figure 9.34), you can select to extract data to Personal Folders or to import data from Personal Folders. Select *Step 1: Extract data from an Exchange Server Mailbox* and then select Next.
24. In the Source Server window (Figure 9.35), enter the name of the Exchange server from which you want to extract data. You also can specify an Active Directory DC and a port number to use for LDAP queries. If you do not specify a DC, ExMerge will use the first available DC. If you do not specify a port number, ExMerge will use port 389. ExMerge will use the DC to extract a list of storage groups and databases available on the Exchange server.
25. Select Options to display the Data Selection Criteria window. You can use the five tabs on the Data Selection Criteria window to specify the criteria ExMerge will use to select the data that should be extracted from the source store.

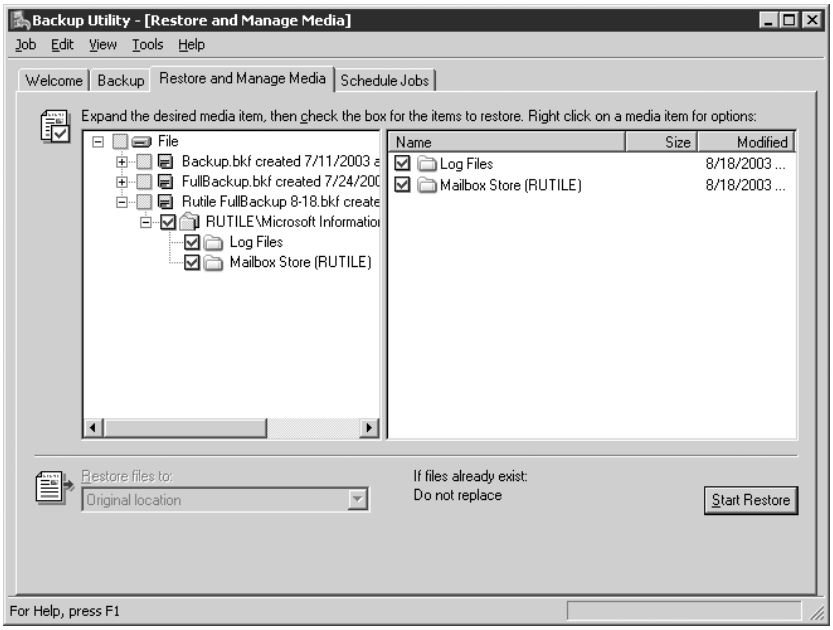

Figure 9.30 *Backup Utility – Restore and Manage Media tab*

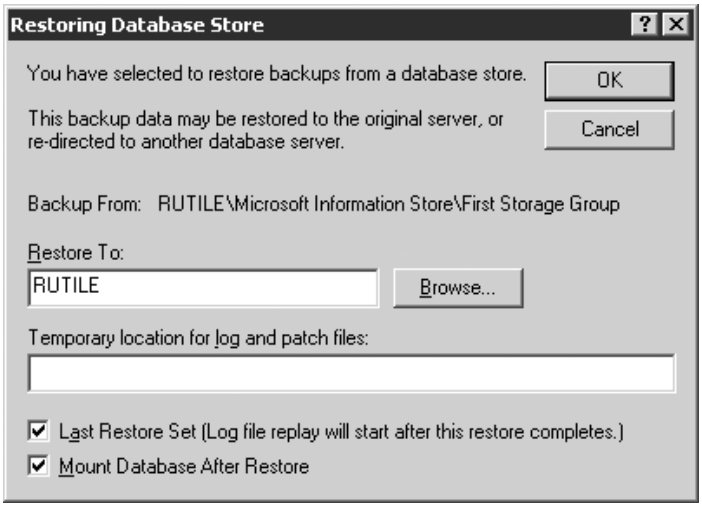

Figure 9.31 *Restoring Database Store dialog box*

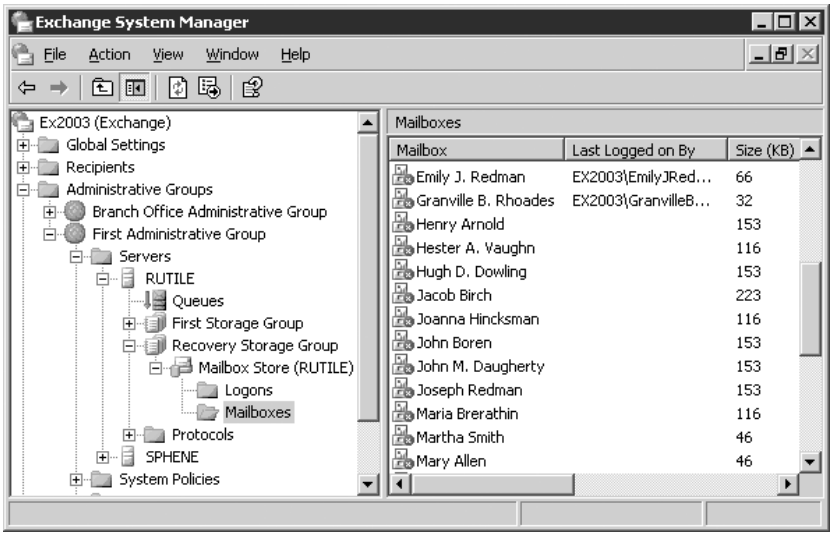

Figure 9.32
Recovered database

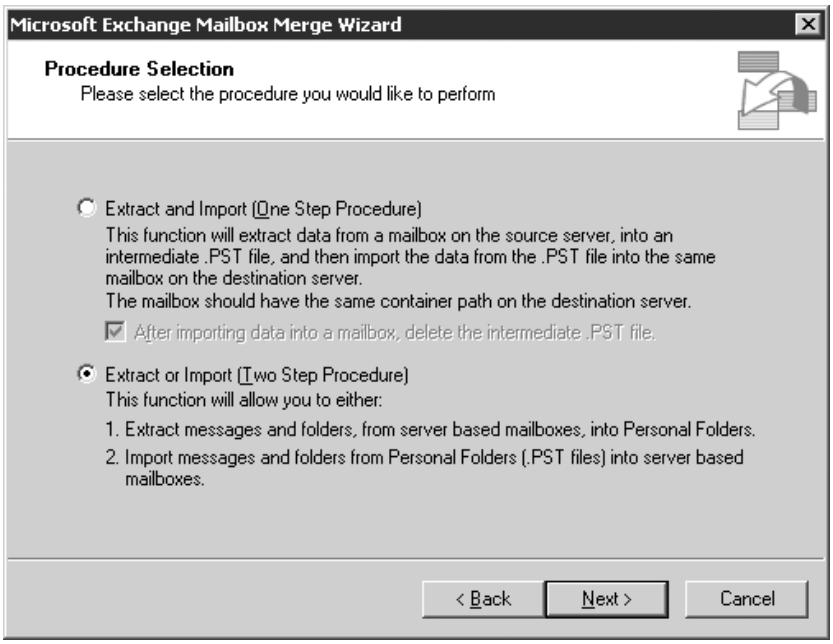

Figure 9.33
Exchange Mailbox Merge Wizard – Procedure Selection

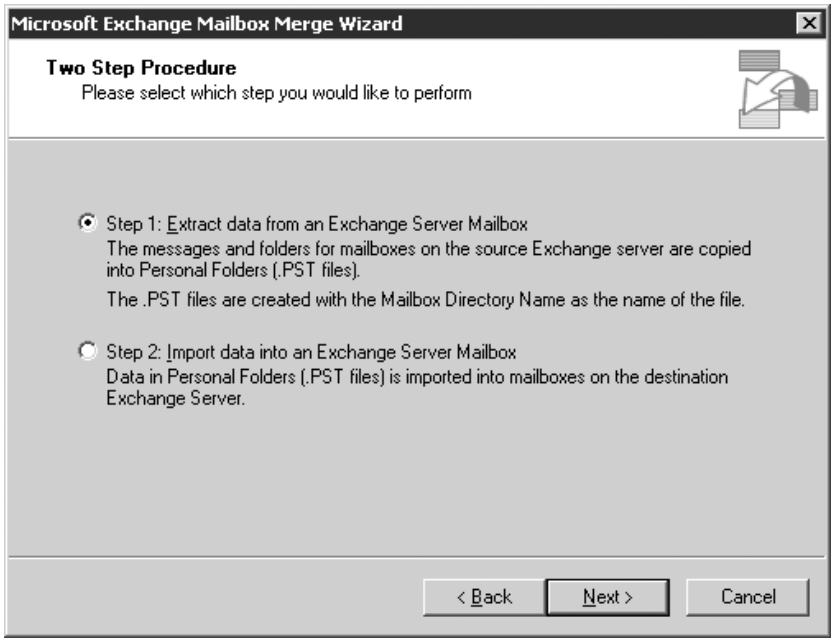

Figure 9.34 *Exchange Mailbox Merge Wizard – Two Step Procedure*

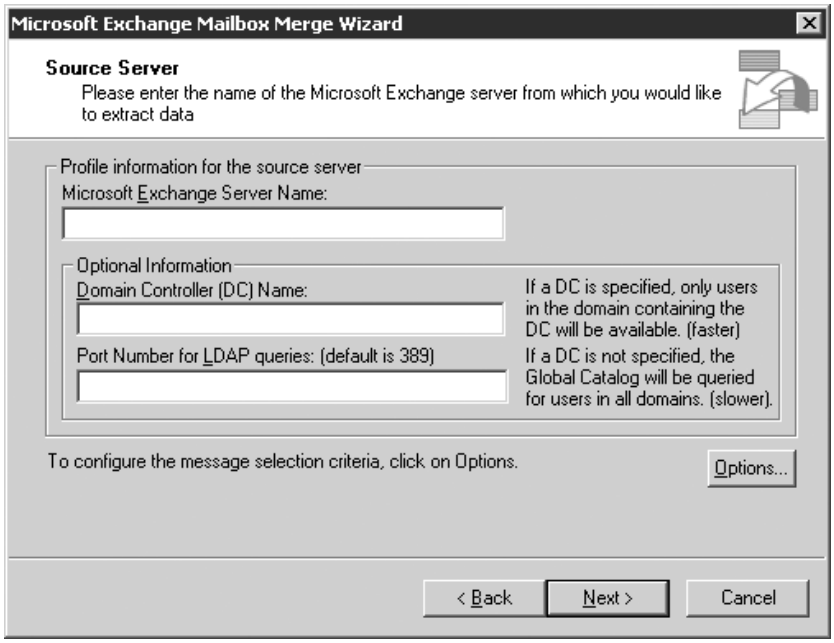

Figure 9.35 *Exchange Mailbox Merge Wizard – Source Server*

Data tab

26. On the Data tab (Figure 9.36), you can select the types of data to be extracted from the source store. You can select any combination of the following options:
    - **User messages and folders**. Select this check box to extract all types of messages, including e-mail messages, contacts, appointments, tasks, notes, and journal items. This is the default option.
    - **Associated folder messages**. Select this check box to extract associated messages in user folders. Associated messages are special, hidden messages that are used to store user settings, such as folder rules and folder views.
    - **Folder permissions**. Select this check box to extract folder permissions. If you select this check box, all existing permissions on the target store will be replaced by the permissions from the source store. This option is most useful when extracting data from a backup and importing it into a server in the original site.

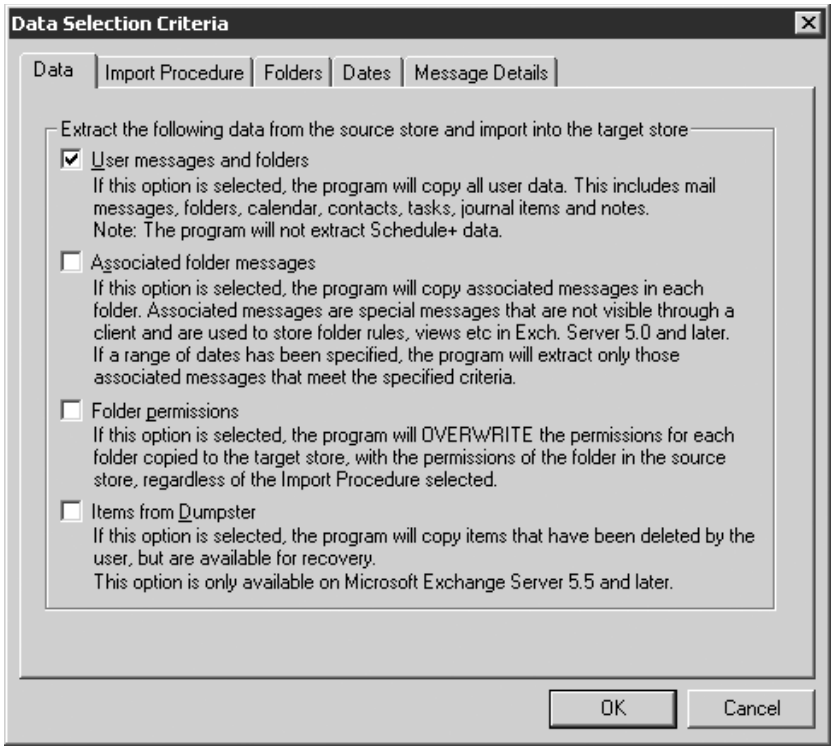

Figure 9.36 *Data Selection Criteria – Data tab*

- **Items from Dumpster**. Select this check box to extract data from the deleted item recovery area—commonly referred to as the *dumpster*. Even after a user empties the Deleted Items folder, it may be possible to recover the deleted items without using the backup media. You can configure Exchange so it does not immediately delete mail or mailboxes for a specified number of days. This period is known as the deleted item retention period, and you can tailor the duration to meet the needs of your organization. If the missing item was deleted within the deleted item retention period, the user can recover the item from the deleted item recovery area without your assistance. The *Items from Dumpster* check box only extracts messages that were deleted from the Deleted Items folder. If a user permanently deleted messages from other folders, ExMerge cannot extract them from the deleted item recovery area.

Folders tab

27. On the Folders tab (Figure 9.37), you can specify the set of folders for ExMerge to process or to ignore. ExMerge matches folders on the basis

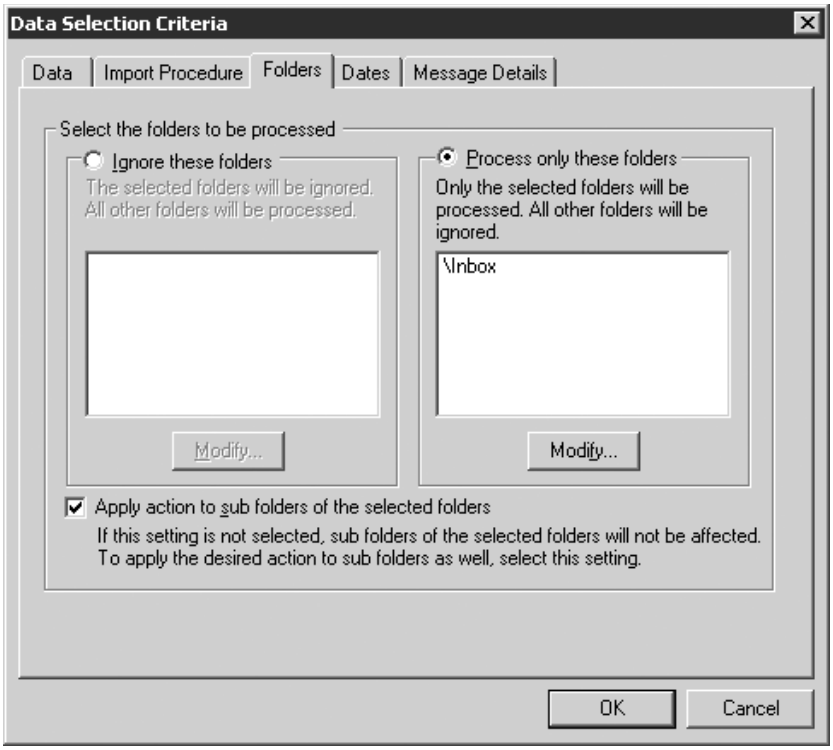

Figure 9.37 *Data Selection Criteria – Folders tab*

of an exact character match rather than folder type. For example, if you select the "\Inbox" folder, ExMerge will not match inbox folders in other languages or inbox folders that you have renamed. If you select the *Apply action to sub folders of the selected folders* check box, ExMerge will ignore subfolders whenever you have told ExMerge to ignore the parent folder and will process subfolders whenever you have told ExMerge to process the parent folder.

Dates tab

28. On the Dates tab (Figure 9.38), you can specify that you want ExMerge to select items from the source store on the basis of the delivery or modification date. ExMerge will ignore any messages that do not fall within the date range you specify. Specifying date selection criteria when importing .PST files into the target store will cause the import process to fail. To extract and import items that have a specific date range, you must enter the date range selection criteria when exporting items from the source store and then omit the date range selection criteria when importing the extracted .PST files into the target store. If you specify a range of dates

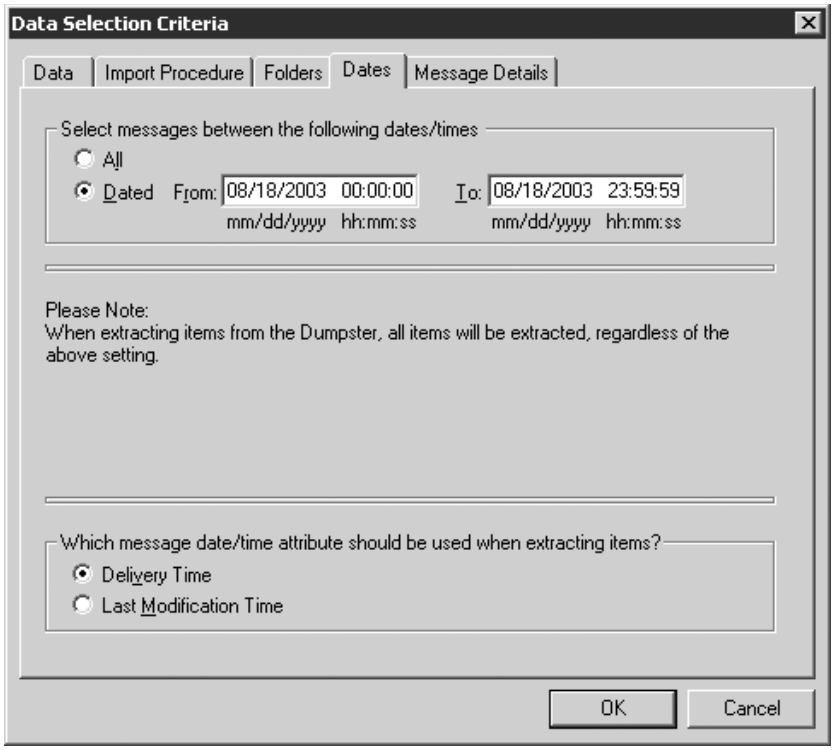

Figure 9.38 *Data Selection Criteria – Dates tab*

for selection criteria, you also can select *Delivery Time* or *Last Modification Time* to specify the date attribute ExMerge should compare.

Note: *If you selected the* Items from Dumpster *check box on the Data tab, ExMerge will ignore the range of dates when extracting items from the dumpster.*

Message Details tab

29. On the Message Details tab (Figure 9.39), you can specify that you only want ExMerge to select messages that include certain attachments, have a certain text string in the subject line, or both. Specifying message details criteria when importing .PST files into the target store will cause the import process to fail. To extract and import items that have a specific subject line or attachment, you must enter the message details criteria when exporting items from the source store and then omit the criteria when importing the extracted .PST files into the target store. You can use the *Subject string compare criteria* drop-down list and the *Attachment name string compare criteria* drop-down

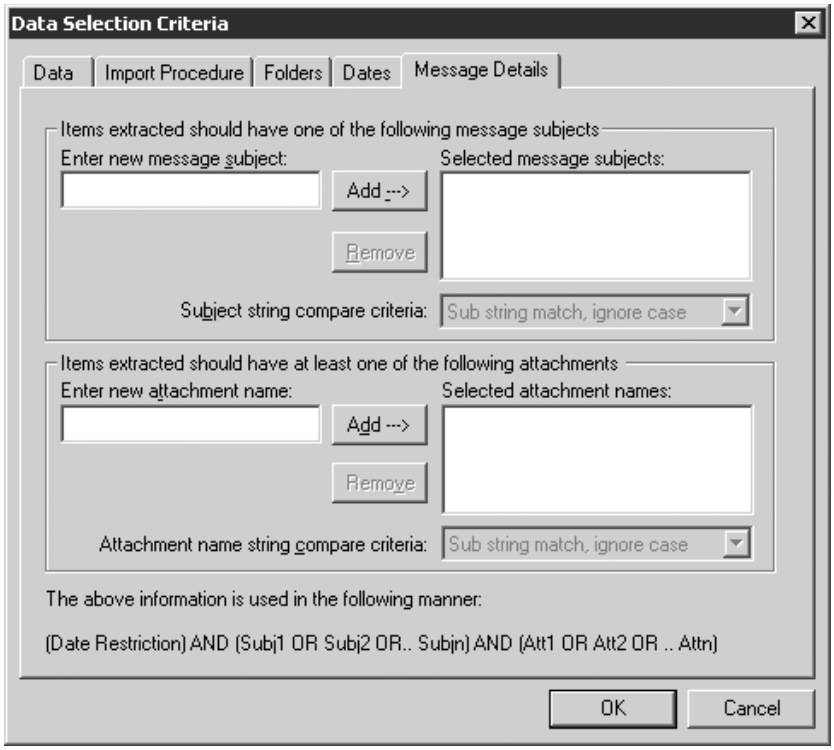

Figure 9.39 *Data Selection Criteria – Message Details tab*

list to select how ExMerge will perform string comparison. The options are as follows:

- **Substring match, ignore case**. ExMerge will look for items that have the specified text string within the subject line or attachment names. The comparison will be sensitive to case.
- **Full-string match, ignore case**. ExMerge will look for items that contain the entire specified text string. The comparison will not be sensitive to case.
- **Exact match**. ExMerge will look for items that contain the entire specified text string. The comparison will be sensitive to case.

30. When you have entered your data selection criteria, select OK to return to the Source Server window (see Figure 9.35). In the Source Server window, select Next to continue.
31. ExMerge will query the Active Directory DC to extract a list of storage groups and databases available on the Exchange server and will display the database selection options on the Database Selection window (Figure 9.40). To extract items recovered from backup tapes, select one or more databases from the RSG. Select Next to continue.

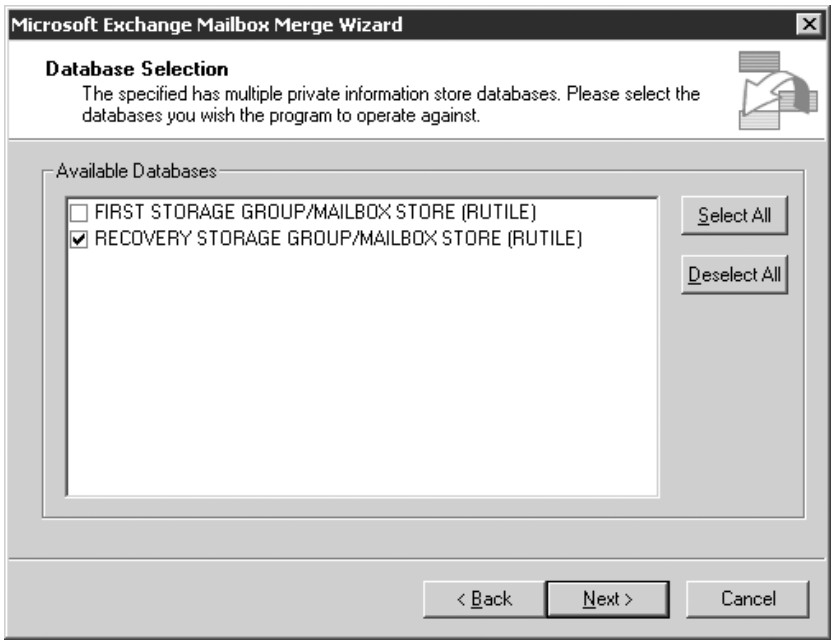

Figure 9.40
Exchange Mailbox Merge Wizard – Database Selection

32. ExMerge collects the list of mailboxes from the selected source databases and displays the list on the Mailbox Selection window (Figure 9.41). The mailbox size is a rough estimate; the actual size might be greater than the displayed size, especially if you are extracting data from the dumpster. Select the mailboxes from which ExMerge should extract data and then select Next to continue.
33. In the Locale Selection window, select the locale that ExMerge should use when connecting to a mailbox. Specifying the locale allows ExMerge to work with any supported language mailbox. ExMerge extracts data from the source mailbox using the locale with which the source mailbox was created. When importing data into the destination mailbox, the locale controls the language in which the mailbox folders will be created. Select Next to continue.
34. In the Target Directory window (Figure 9.42), select the folder where you would like ExMerge to store .PST files containing the extracted data. The *Required* field shows the estimated amount of free disk space that is required to hold the extracted .PST files. However, this is just a rough estimate, and if you have specified any selection criteria (e.g., date ranges, folders, subject line, or attachment name), ExMerge

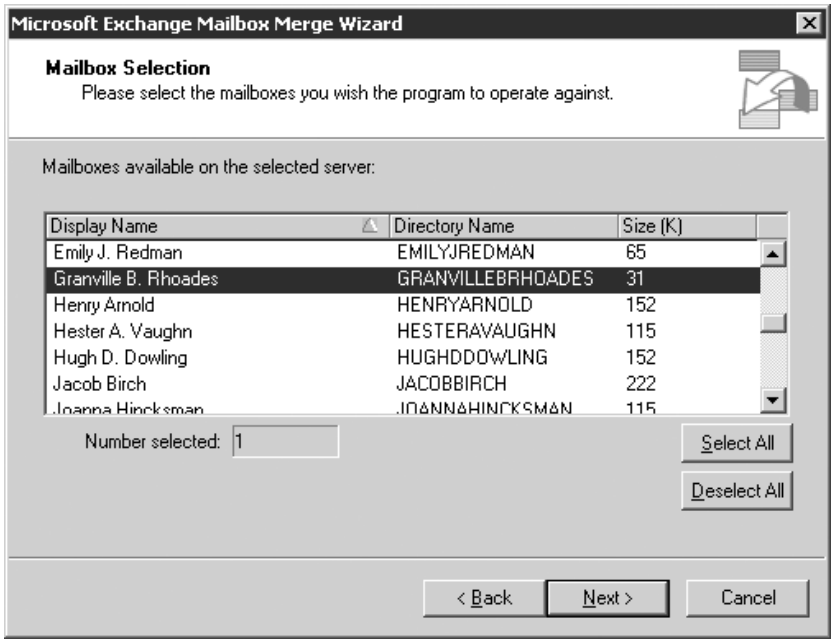

Figure 9.41
Exchange Mailbox Merge Wizard – Mailbox Selection

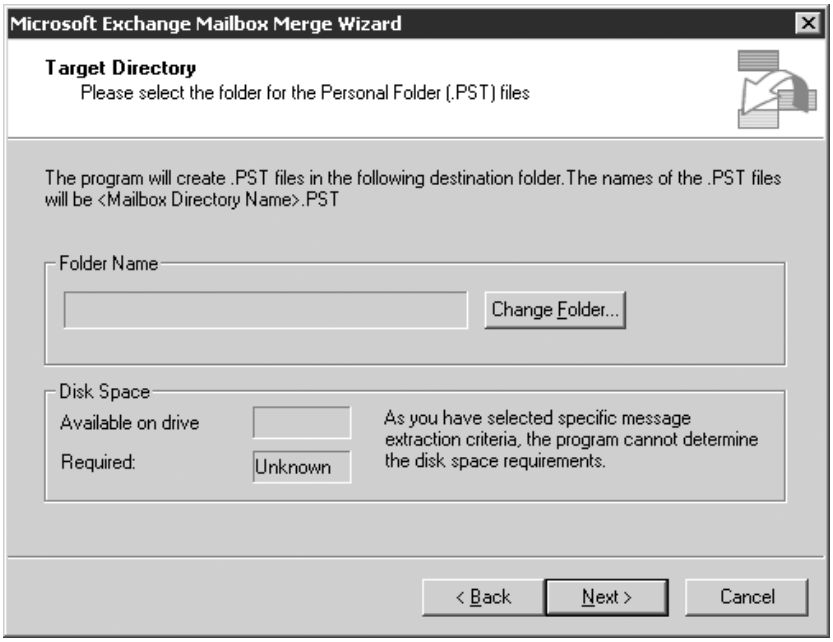

Figure 9.42
Exchange Mailbox Merge Wizard – Target Directory

will not display any estimate because it cannot calculate the required space before examining each mailbox. ExMerge will create a separate .PST file for each selected mailbox. The names of the .PST files will be *MailboxDirectoryName* .PST. Select Next to continue.

35. In the Save Settings window, you can save all of your program settings so that you can run ExMerge in batch mode at a later time. Select Next to begin the mailbox extraction process. ExMerge will display a Process Status window showing the status of the extraction.
36. When ExMerge is finished, select Finish in the Process Status window. ExMerge creates (or appends to) a log file containing any errors and messages indicating the progress of the current operation. By default, this log file is named ExMerge.log and is created in the same directory as the ExMerge.exe file.

Merge recovered messages into production database

Note: *The user running the following ExMerge procedure must have* Send As *and* Receive As *rights for the target mailboxes.*

37. Start ExMerge. In the ExMerge welcome window, select Next to continue.

38. In the Procedure Selection window (see Figure 9.33), select *Extract or Import (Two Step Procedure)* and then select Next.
39. In the Two Step Procedure window (see Figure 9.34), select *Step 2: Import data into an Exchange Server Mailbox* and then select Next.
40. In the Destination Server window (Figure 9.43), enter the name of the destination Exchange server. You also can specify an Active Directory DC and a port number to use for LDAP queries. If you do not specify a DC, ExMerge will use the first available DC. If you do not specify a port number, ExMerge will use port 389. ExMerge will use the DC to extract a list of storage groups and databases available on the destination Exchange server.
41. Select Options to display the Data Selection Criteria window. ExMerge lists the same five tabs on the Data Selection Criteria window that were available to specify the criteria that ExMerge used to select the extracted data from the source store. However, when importing data to a destination server, you should only use the Import Procedure tab. You use the Data, Folders, Dates, and Message Details tabs to specify the selection criteria (e.g., date ranges, folders, subject line, or attachment name) for

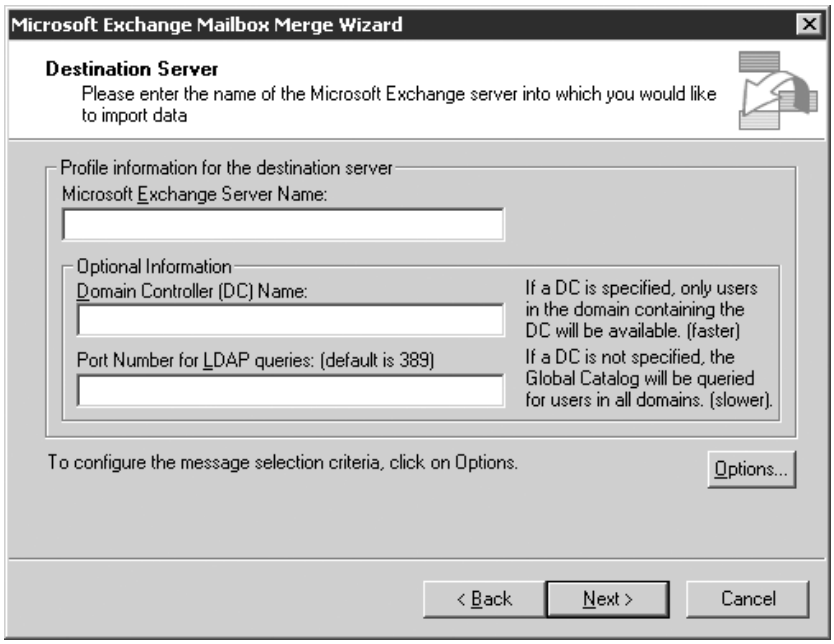

Figure 9.43
Exchange Mailbox Merge Wizard – Destination Server

extracting messages from the source store. To import items that have specific selection criteria, you must enter the selection criteria when exporting items from the source store and then omit the selection criteria when importing the extracted .PST files into the target store.

Import Procedure tab

42. On the Import Procedure tab (Figure 9.44), you can select how ExMerge should add items to the target store. You can select any one of the following options:
    - **Copy data into the target store**. Select this option to copy each item from the .PST file into the target store without checking whether the item already exists in the target store. Because this option can result in duplicate messages in the target store, you should use this option only if you are certain that the target store does not contain any of the messages in the .PST file.
    - **Merge data into the target store**. Select this option to merge .PST data into the target store. ExMerge will check whether an item already exists in the target database before copying the item from

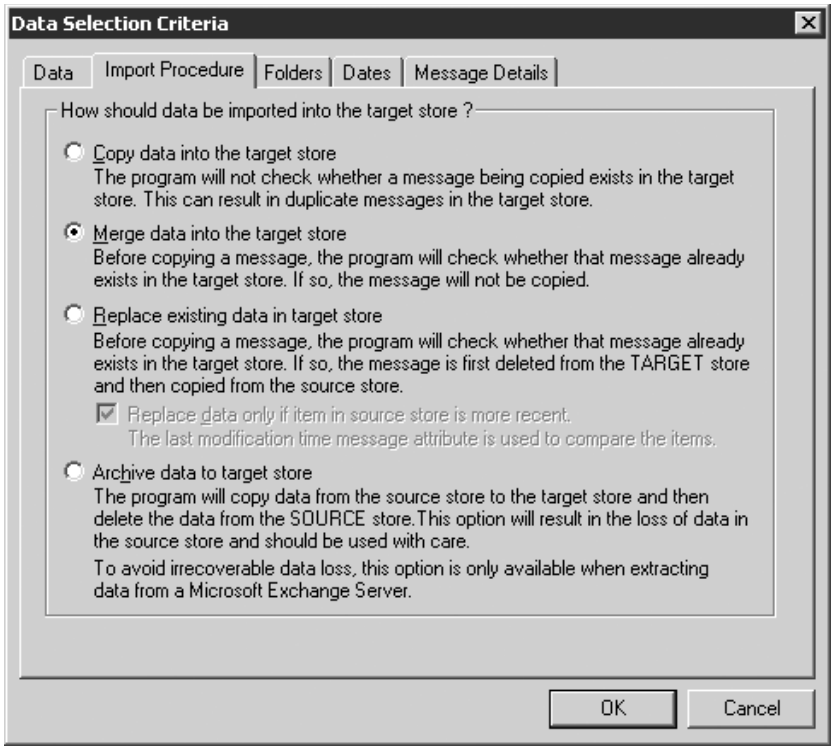

Figure 9.44 *Data Selection Criteria – Import Procedure tab*

the .PST file. This is the default (and preferred) option because it avoids creating duplicate messages.

- **Replace existing data in target store**. Select this option to overwrite any existing items in the target store. When ExMerge finds a duplicate item in the target store, it will delete the item and then copy the new item from the .PST file. This option can result in data loss if the items in the target store have been modified and are different from those in the .PST file. To avoid this potential problem, you can select the *Replace data only if item in source store is more recent* check box. If you select this check box, ExMerge will compare the last modified time for the duplicate items before replacing the item that already exists in the target store. ExMerge will overwrite the item in the target store only if the copy of the item in the .PST has been more recently modified.

43. When you have entered your import procedure options, select OK to return to the Destination Server window (see Figure 9.43). In the Destination Server window, select Next to continue.
44. ExMerge collects the list of mailboxes available on the destination Exchange server and displays the list on the Mailbox Selection window (Figure 9.45). Select the mailboxes into which ExMerge should merge data and then select Next to continue.
45. In the Locale Selection window, select the locale that ExMerge should use when connecting to a mailbox. Select Next to continue.
46. In the Target Directory window (see Figure 9.42), select the folder where you told ExMerge to store the .PST files containing the extracted messages. Select Next to continue.
47. In the Save Settings window, you can save all of your program settings so that you can run ExMerge in batch mode at a later time. Select Next to begin the mailbox merge process. ExMerge will display a Process Status window showing the status of the merge.
48. When ExMerge is finished, select Finish in the Process Status window.

9.12 Recovering from a disaster

On rare occasions, a server fails badly or is physically damaged to the point that you have no choice except to rebuild the system. The recovery process is complex, but if you have maintained the disaster recovery toolkit described in Section 9.2, the recovery should not be a problem.

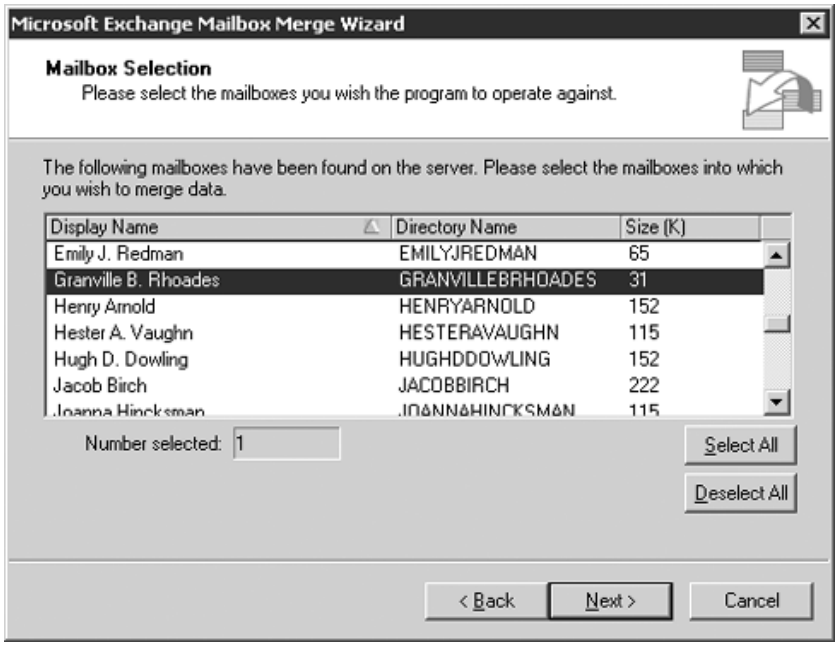

Figure 9.45
Exchange Mailbox Merge Wizard – Mailbox Selection

After fixing the system hardware (or acquiring replacement hardware), you must reinstall and reconfigure Windows, restore your disk drives using data from your backup media, recover the system state information from your backup media, reinstall Exchange in disaster recovery mode, and finally recover the Exchange Information Store databases from your backup media. The following procedure can be used to recover from a complete disaster. Many failures will not result in complete data loss, so you may need to adjust the following procedures to match your situation.

1. Reinstall Windows.
   - Reinstall the same version of Windows by running Windows Setup with the following options:
     - The hardware and software configuration should match the original Exchange server, including the same components, the same version of Windows, the same service packs and hot fixes, the same drive designations, and the same server name.
     - You should not join the Windows domain. Configure Windows as a standalone server in a workgroup. The server will automatically

rejoin the correct domain when you restore the System State from the backup media.

2. Restore the disk drives.
   - Find and mount the correct backup media.
   - From the Windows Start menu, select All Programs → Accessories → System Tools → Backup (see Figure 9.2).
   - On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (see Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*

   - In the Backup Utility window, select the Restore and Manage Media tab (see Figure 9.19).
   - On the Restore and Manage Media tab, double-click the backup file containing the files you want to restore. Use the check boxes to select each drive you want to restore. You should restore the system drive and any other drives containing data or applications.
   - Select Start Restore to begin recovering the drives.
3. Recover the system state.
   - Find and mount the correct backup media.
   - From the Windows Start menu, select All Programs → Accessories → System Tools → Backup (see Figure 9.2).
   - On the Backup or Restore Wizard Welcome window, select Advanced Mode to start the Backup Utility (see Figure 9.3).

   Note: *If you clear the* Always start in wizard mode *check box, you can avoid the Welcome to the Backup or Restore Wizard window in the future by going directly to the Backup Utility.*

   - In the Backup Utility window, select the Restore and Manage Media tab (see Figure 9.19).
   - On the Restore and Manage Media tab, double-click the backup file containing the files you want to restore. Select the *System State* check box. The System State backup includes Active Directory data, Windows registry data, and other data that are not usually backed up by file and drive backups.

- Select Start Restore to begin recovering the System State.
- Restart the server.

You will receive an error dialog box informing you that at least one service could not be started. The failing services are the ones that require Exchange. Windows incorrectly believes these services are configured on this server because they are listed in the System State backup. This problem will be corrected automatically when Exchange is reinstalled in disaster recovery mode.

4. Reinstall Exchange in Disaster Recovery Mode.
   - Insert the Exchange Server 2003 CD-ROM into your CD-ROM drive.
   - Select Run from the Windows Start menu. As the command to run, enter *x*:\setup\i386\setup.exe/DisasterRecovery, where *x* is your CD-ROM drive. Select OK to start the setup program.
   - You must install Exchange to the same drive and directory on which it was installed on the original server. Ensure that each component that was originally installed has an action of Disaster Recovery. If all of the originally installed components are not automatically set for Disaster Recovery, you must manually select them. Running Exchange Setup in Disaster Recovery mode restores the original Exchange system configuration and services. Once you have restored the Exchange configuration, you can recover the Exchange databases.
5. Recover Exchange Information Store databases.
   - Use the procedure outlined in Section 9.9 to recover the Exchange databases.