

Chapter 9

Disaster Recovery with Exchange Server 2007

Solutions in this chapter:

- **Backing Up Exchange 2007 Using Windows 2003 Backup**
- **Restoring Exchange 2007 Storage Groups and Databases Using Windows 2003 Backup**
- **Repairing a Corrupt or Damaged Exchange 2007 Database Using Eseutil**
- **Restoring Mailbox Data Using the Recovery Storage Group Feature**
- **Recovering an Exchange 2007 Server Using the RecoverServer Switch**
- **Recovering an Exchange 2007 Cluster Using the RecoverCMS Switch**
- **Restoring Mailbox Databases Using the Improved Database Portability Feature**

Introduction

As mentioned in the previous chapter, the messaging and collaboration servers are mission critical, being perhaps the most vital servers in our datacenters today. It's therefore of the utmost importance that these servers be up and running all the time. Most service level agreements today require more than 99.99 percent uptime when it comes to the messaging and collaboration servers in the organization. In the previous chapter we showed you some of the options available to provide high availability of the Exchange 2007 mailbox Servers. But even if you have HA solutions such as CCR-based mailbox servers available, a disaster can still strike in your environment, and if this happens, you better be prepared since downtime typically means lost productivity and revenue. In this chapter, we'll go through the steps necessary to back up the different Exchange 2007 Server roles in your organization, and, just as important, look at how you restore Exchange 2007 servers and data should it be required.

Backing Up Exchange 2007 Using Windows 2003 Backup

Frequent backups of the Exchange 2007 servers in an organization are important operational tasks that, though a bit trivial, should be taken very seriously. We can only imagine one thing worse than a complete failure of an Exchange 2007 server, and that's a complete failure of an Exchange 2007 server without any backups to restore from. In the first section of this chapter, we'll take a look at what you must back up, depending on which Exchange 2007 Server roles were deployed in your organization.

Backing Up an Exchange 2007 Mailbox Server

One of the most important things to back up regarding Exchange 2007 Mailbox Servers are the databases, which hold user mailboxes and public data. As you saw in the previous chapter, Exchange 2007 provides a new continuous replication functionality that keeps a second copy of one or more databases in a storage group in sync with the active versions of the databases using log file shipping and replay. This provides an extra level of protection for Mailbox and Public Folder databases. However, although the new functionality allows you to make less frequent backups of your databases, it doesn't eliminate the *need* for database backups. In this section, we'll show you how to perform a backup of the databases on an Exchange 2007 server.

SOME INDEPENDENT ADVICE

Another reason why it's crucial to conduct frequent full backups of your Exchange databases with an Exchange-aware backup application is to commit and delete any transaction log files generated since the last full backup. If

these log files aren't committed, they will take up more and more space on your disks, and when there's no more disk space for the log files, the database will be dismounted.

Since Exchange 2007 databases still use ESE, you can (just as with previous versions of Exchange), back them up using the Exchange-aware native Windows 2003 backup tool. Exchange 2007 supports two different backup methods. The first is a legacy streaming backup method based on the ESE application programming interface (API), which allows you to back up one or more storage groups at the same time. However, only one backup job can run against a specific storage group. Most of us are familiar with this type of backup since it's the one we have used for ages when referring to Exchange databases. The ESE API backup method is supported by the Windows 2003 backup tool, as well as most third-party backup products.

Then we have the Volume Shadow Copy Service (VSS) backup method, which some of you may know from Exchange 2003 where it was first introduced. The interesting thing about VSS is that this method, in addition to what the legacy streaming backup method offers, can also make an online backup of the copy database when using either Local Continuous Replication or Cluster Continuous Replication in your setup. This means you can schedule the backup windows anytime you want since taking a backup of the database copy has no performance-related impact on the active database. Unfortunately, this method isn't supported by the Windows 2003 backup tool when speaking Exchange databases (only file level backups), and Microsoft doesn't offer any products capable of using VSS, at least not at the time of this writing.

NOTE

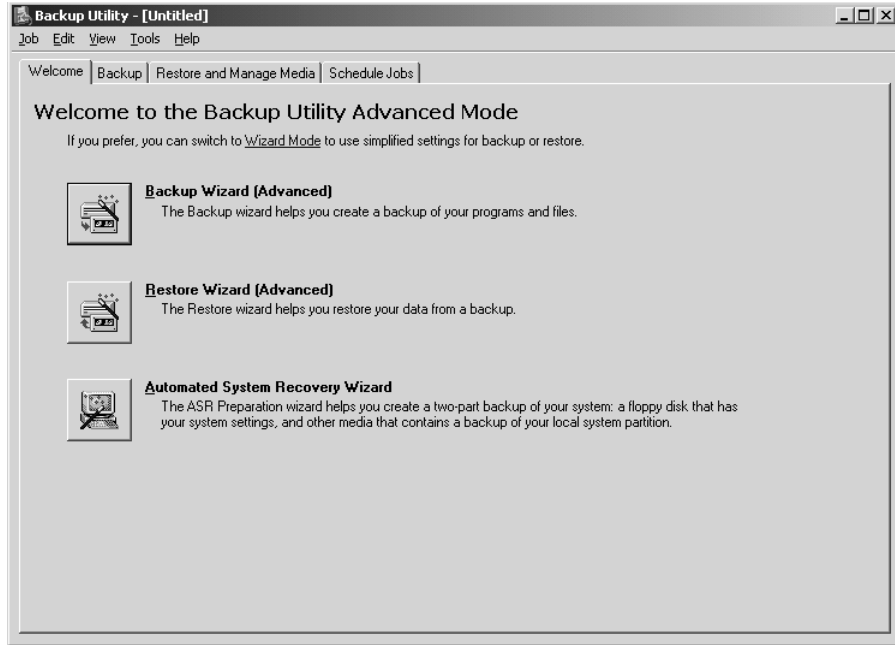
The Data Protection Manager (DPM) v2 product will support VSS backups, however. DPM is a server software application that enables disk- and tape-based data protection and recovery for file servers, servers running Microsoft Exchange, and servers running Microsoft SQL Server in an Active Directory Domain Services (AD DS) domain. DPM performs replication, synchronization, and recovery point creation to provide reliable protection and rapid recovery of data for both system administrators and end users.

Let's go through the steps necessary to back up an Exchange 2007 Mailbox and Public Folder database on an Exchange 2007 Mailbox Server. The first thing you need to do is launch the Windows 2003 backup tool, which can be done by clicking **Start** | **Run** and

484 Chapter 9 • Disaster Recovery with Exchange Server 2007

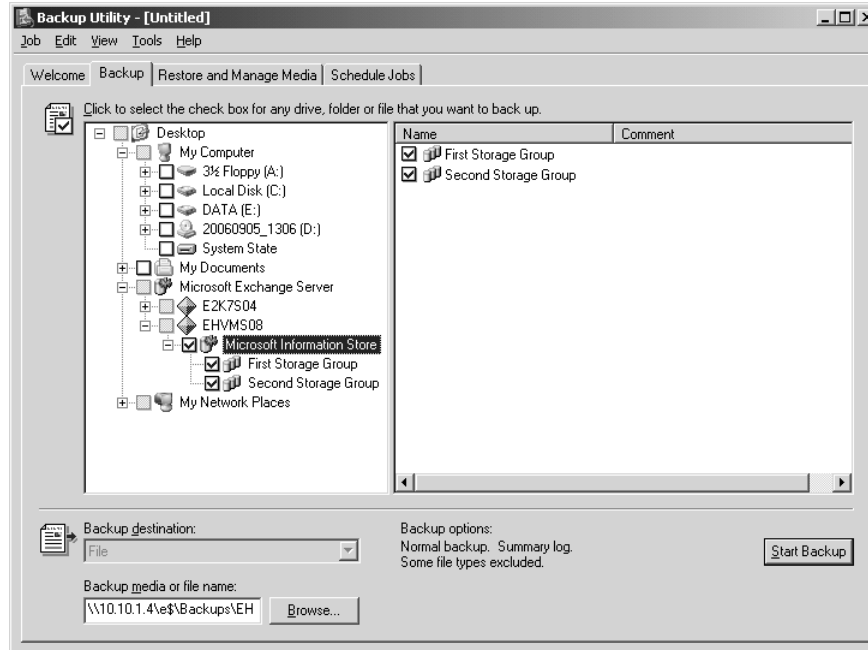
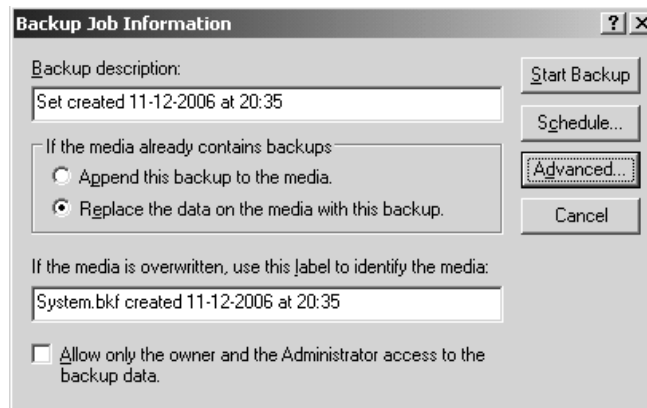
typing **NTBackup**. Now click **Switch to Advanced Mode** and then click the **Backup** tab shown in Figure 9.1.

Figure 9.1 Windows 2003 Backup Tool



Under the **Backup** tab expand **Microsoft Exchange Server | Mailbox Server | Microsoft Information Store** and check the storage group(s) containing the **Mailbox** and **Public Folder** database (Figure 9.2). Now specify the backup media or filename you want to perform the backup to, and then click **Start Backup**.

As you can see in Figure 9.3, you now have the option of entering a description for the respective backup job, as well as specify whether the backed-up data should be appended to an existing backup. In addition, you can create a scheduled backup job so it runs, let's say, every day at midnight. By clicking the **Advanced** button, you also have the option of having the backed-up data verified when the job completes.

Figure 9.2 Selecting the Storage Groups to Be Backed Up**Figure 9.3** Backup Job Information

Typically, you should set up an automated backup job schedule, but for the purpose of this example we'll just choose to back up the databases once. When ready, click **Start Backup**.

When the backup job has completed, you can view a report, which will contain any warnings or errors that might occur during the backup.

486 Chapter 9 • Disaster Recovery with Exchange Server 2007

That's how you back up the Mailbox and Public Folder databases, as well as commit and delete any existing transaction log files using the Windows 2003 Backup tool. Sounds simple, right?

Some of you might wonder whether there isn't anything else you need to back up on an Exchange 2007 Mailbox Server? The answer is no critical files at least since you can always recover an Exchange 2007 Mailbox Server using the *Setup /Mode:RecoverServer* command (shown later in the chapter), but it's always a good idea to back up the System State of the respective server as well.

Backing Up an Exchange 2007 Hub Transport Server

Since an Exchange 2007 Server with the Hub Transport Server role installed was designed to store all configuration data in the Active Directory configuration container, not much needs to be backed up on a server with this role installed either. But just as with the Mailbox server role, you should back up the System State.

Some of you may be wondering why I haven't mentioned anything about backing up the message queues stored in an ESE database on an Exchange 2007 Hub Transport Server... Well, there shouldn't be any need to do so since you can mount the message queues on another existing, or newly installed, Hub Transport server if required. You just need to retrieve the mail.queue (which, by default, is located under `C:\Program Files\Microsoft\Exchange Server\TransportRoles\data\Queue`) from the failed Hub Transport server.

NOTE

Step-by-step instructions on how to move a message queue from a failed Hub Transport server to another Hub Transport server in the organization is outside the scope of this book, but you can find information on the topic by searching under "Working with the Queue Database on Transport Servers" in the Exchange 2007 Documentation Help file.

One thing you might want to back up regarding an Exchange 2007 Hub Transport Server is the Message Tracking and Protocol logs which, by default, are located under `C:\Program Files\Microsoft\Exchange Server\TransportRoles\Logs`. These files can be backed up using a file level backup.

As is the case with a Mailbox Server, you can recover a Hub Transport server using the *Setup /Mode:RecoverServer* command.

Backing Up an Exchange 2007 Client Access Server

When using Exchange 2007 Server with the Client Access Server role installed, there are several files you should back up. The first, and perhaps most important, to back up is the IIS Metabase, which among other things is used to store OWA Virtual Directory configuration data. You can back up the IIS configuration on a CAS using the following command:

```
get-owavirtualdirectory "owa (default web site)" | export-xml owa.xml -depth 1
```

In order to restore the IIS configuration from the owa.xml file, you need to use a Windows PowerShell script similar to the following (save it as Restore-OWA.PS1 or use some other meaningful name):

```
$ErrorActionPreference = 'stop'
$savedprops = @(
'DirectFileAccessOnPublicComputersEnabled',
'DirectFileAccessOnPrivateComputersEnabled',
'WebReadyDocumentViewingOnPublicComputersEnabled',
'WebReadyDocumentViewingOnPrivateComputersEnabled',
'ForceWebReadyDocumentViewingFirstOnPublicComputers',
'ForceWebReadyDocumentViewingFirstOnPrivateComputers',
'RemoteDocumentsActionForUnknownServers',
'ActionForUnknownFileAndMIMEtypes',
'WebReadyFileTypes',
'WebReadyMimeTypes',
'WebReadyDocumentViewingForAllSupportedTypes',
'AllowedFileTypes',
'AllowedMimeTypes',
'ForceSaveFileTypes',
'ForceSaveMimeTypes',
'BlockedFileTypes',
'BlockedMimeTypes',
'RemoteDocumentsAllowedServers',
'RemoteDocumentsBlockedServers',
'RemoteDocumentsInternalDomainSuffixList',
'LogonFormat',
'ClientAuthCleanupLevel',
'DefaultDomain',
'FormsAuthentication',
'BasicAuthentication',
'DigestAuthentication',
```

488 Chapter 9 • Disaster Recovery with Exchange Server 2007

```
'WindowsAuthentication',
'GzipLevel',
'FilterWebBeaconsAndHtmlForms',
'NotificationInterval',
'DefaultTheme',
'UserContextTimeout',
'ExchwebProxyDestination',
'VirtualDirectoryType',
'RedirectToOptimalOWAServer',
'DefaultClientLanguage',
'LogonAndErrorLanguage',
'UseGB18030',
'UseISO885915',
'OutboundCharset',
'CalendarEnabled',
'ContactsEnabled',
'TasksEnabled',
'JournalEnabled',
'NotesEnabled',
'RemindersAndNotificationsEnabled',
'PremiumClientEnabled',
'SpellCheckerEnabled',
'SearchFoldersEnabled',
'SignaturesEnabled',
'ThemeSelectionEnabled',
'JunkEmailEnabled',
'UMIntegrationEnabled',
'WSSAccessOnPublicComputersEnabled',
'WSSAccessOnPrivateComputersEnabled',
'ChangePasswordEnabled',
'UNCAccessOnPublicComputersEnabled',
'UNCAccessOnPrivateComputersEnabled',
'ActiveSyncIntegrationEnabled',
'AllAddressListsEnabled',
'InternalUrl',
'ExternalUrl'
)

$vdirex = import-clixml $args[0]

'Recreating "' + $vdirex.name + '"' + ' owa version: ' + $vdirex.owaversion
if ($vdirex.owaversion -eq 'Exchange2007') {
```

```

new-owavirtualdirectory -website $vdir.website -internalurl
$vdir.internalurl -externalurl $vdir.externalurl
}
else {
new-owavirtualdirectory -website $vdir.website -owaversion $vdir.
owaversion -name $vdir.displayname -virtualdirectorytype $vdir.
virtualdirectorytype
}
$new = get-owavirtualdirectory $vdir.name
'Restoring properties'
foreach ($prop in $savedprops) {
if ($prop -eq 'ExchwebProxyDestination' -or $prop -eq
'VirtualDirectoryType') {
continue
}
$new.$prop = $vdir.$prop
}
$new | set-owavirtualdirectory

```

To restore the IIS configuration data that were saved in the owa.xml file, type **Restore-OWA.PS1 owa.xml**.

In addition to the IIS metabase, you should back up the System State and the files listed in Table 9.1.

Table 9.1 Files Needed to Restore the IIS Configuration

Data	Location
Microsoft Office Outlook Web Access Web site, and Web.config file	C:\ProgramFiles\Microsoft\Exchange Server\ClientAccess\Owa
IMAP4 and POP3 protocol settings	C:\Program Files\Microsoft\Exchange Server\ClientAccess\
Availability service	Active Directory configuration container and file system, including the Web.config file C:\Program Files\Microsoft\Exchange Server\ClientAccess\exchweb\ews
Autodiscover	IIS metabase

Continued

Table 9.1 continued Files Needed to Restore the IIS Configuration

Data	Location
Exchange ActiveSync	Active Directory configuration container File system, including the Web.config file in the \ClientAccess\Sync folder IIS metabase
Outlook Web Access virtual directories	Active Directory configuration container and file system C:\Program Files\Microsoft\Exchange Server\ClientAccess\
Web services configuration	IIS metabase

Like a Mailbox or Hub Transport Server, a Client Access Server can be restored using the *Setup /Mode:RecoverServer* command.

Backing Up an Exchange 2007 Unified Messaging Server

Exchange 2007 servers with the Unified Messaging (UM) role installed store most of the configuration data in the Active Directory, which means it's very limited what you need to back up on the UM server itself.

Table 9.2 lists the files you need to back up.

Table 9.2 Files to Back Up on Unified Messaging Server

Data	Location
Custom audio prompts: Custom audio files (.wav) for UM Dial Plans and UM Auto Attendants Custom audio files (.wav) for telephone user interface (TUI) or Voice Access	C:\Program Files\Microsoft\Exchange Server\UnifiedMessaging\Prompts
Incoming calls: .eml and .wma files for each voicemail	C:\Program Files\Microsoft\Exchange Server \UnifiedMessaging\temp

In addition, you should back up the System State.

The rest of the configuration data is, as mentioned previously, stored in Active Directory, which makes it possible to restore using the *Setup /Mode:RecoverServer* command.

Backing Up an Exchange 2007 Edge Transport Server

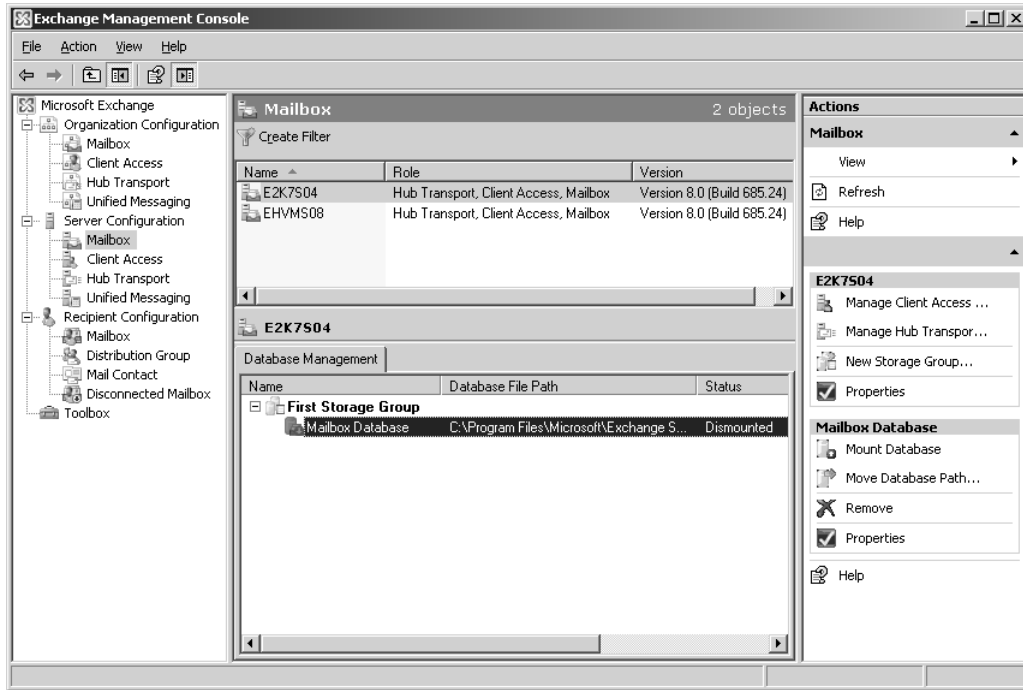
An Exchange 2007 Server with the Edge Transport Server role installed can be restored by using a Cloned Configuration (employing the ImportEdgeConfig.ps1 script). For step-by-step instructions on how you deal with clone configuration, see Chapter 7. In addition to cloned configuration, you should back up System State as well as the Message Tracking and protocol logs, which are located in C:\Program Files\Microsoft\Exchange Server\TransportRoles\Logs. The message queues that are stored in an ESE database just like message queues on a Hub Transport server can be mounted on another Edge Transport server.

Restoring Exchange 2007 Storage Groups and Databases Using Windows 2003 Backup

So now that you have seen how to back up Mailbox and Public Folder databases, you should of course also be aware of how you restore these databases properly should you experience a database corruption or find them unusable in some other way. In this section, I'll show you how to perform a restore of a Mailbox database from the backup set we created earlier in this chapter. When you restore a Mailbox or Public Folder database from a backup set, any associated transaction log files are restored as well. It's important you understand that a restore of a Mailbox database will copy the database file (.EDB) into its original location on the disk, and thereby overwrite any existing .EDB file. In addition, any transaction log files will be copied to a temporary location, which can be specified when doing the actual restore. Upon the restore's completion (hopefully without any serious warnings or errors!), the log files will be replayed into the restored version of the database. In addition to the log files, a file called Restore.env will also be copied to the specified temporary folder. This file keeps control of which storage group the log files belong to, as well as the database paths and range of log files that have been restored.

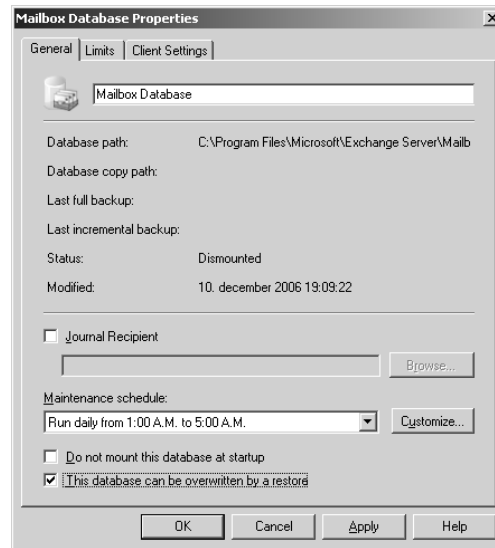
In order to restore the aforementioned Mailbox database, we need to perform the following steps. First, open the **Exchange Management Console**, expand **Server Configuration**, and then select the **Mailbox** subnode. Now choose the respective Mailbox server in the **Result** pane, and then dismount the Mailbox database, as shown in Figure 9.4.

Figure 9.4 Dismounting the Mailbox Database



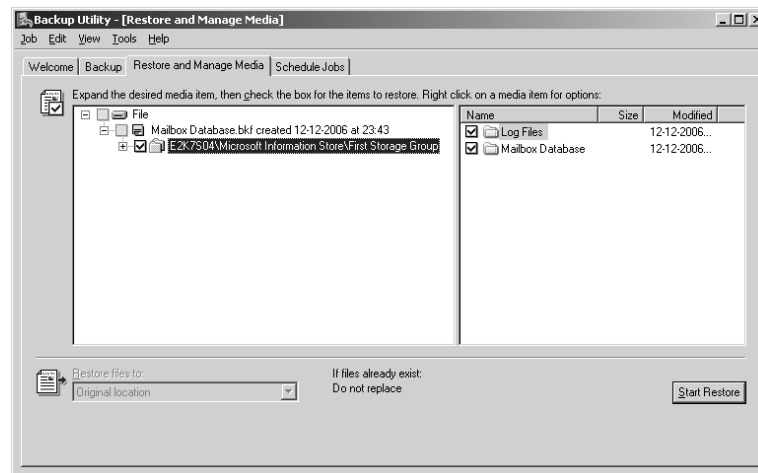
Now open the properties page for the Mailbox database. Check **This database can be overwritten by a restore** (Figure 9.5) and click **OK**.

Figure 9.5 Allowing the Mailbox Database to Be Overwritten by a Restore



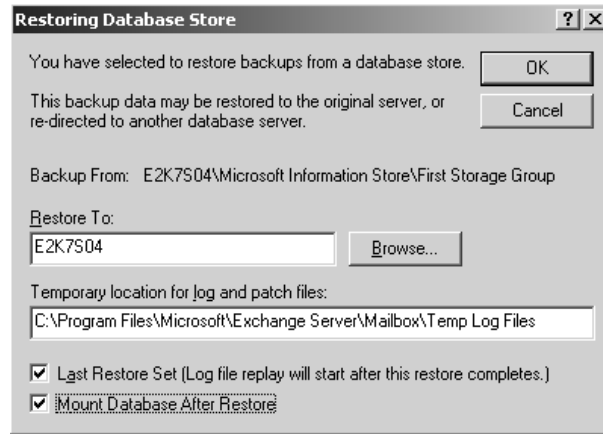
We're now ready to restore the databases using the Windows 2003 Backup tool, so let's launch this tool by clicking **Start | Run** and typing **NTBackup**, and then selecting the **Restore and Manage Media** tab. Expand the desired media item and backup set, then check the **log files** and **mailbox database**, as shown in Figure 9.6. We can then click **Start Restore**.

Figure 9.6 Restoring the First Storage Group

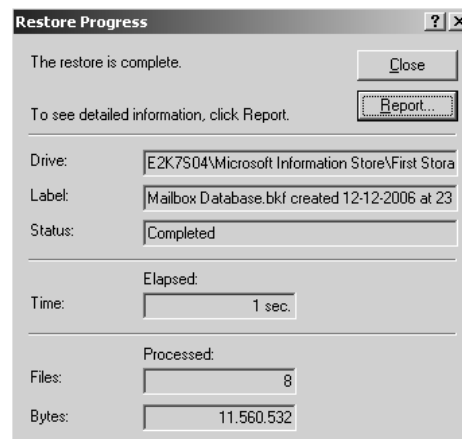


You'll be faced with a screen similar to the one shown in Figure 9.7. Here, you need to specify the server you want to restore the database to (the local server on which the Windows 2003 Backup tool is run is typically pre-entered here), and the temporary location for log and patch files. In addition, you need to specify whether the restore you're performing is the last restore set. If you select this option, all the restored log files will be replayed automatically into the database after the restore has completed. You typically want to do this if you don't have any incremental or differential backups of the database's log files you need to restore after this restore. Finally, you have the option of specifying that the database should be mounted automatically after the restore has occurred. When you have made your selections, click **OK**.

The restore will now begin. Depending on the size of the database, it will take some time to complete. Since the database in this example is under 11MB, the restore took less than a second, as you can see in Figure 9.8. When the restore has completed, you can click the **Report** button to see a detailed log of the restore process. When ready, click **Close**.

Figure 9.7 Restoring Database Store Options

If your restore completed successfully, you can now switch back to the Exchange Management Console, where the restored Mailbox database should have been mounted automatically, and we can call the restore a success.

Figure 9.8 Restore Completed Successfully**NOTE**

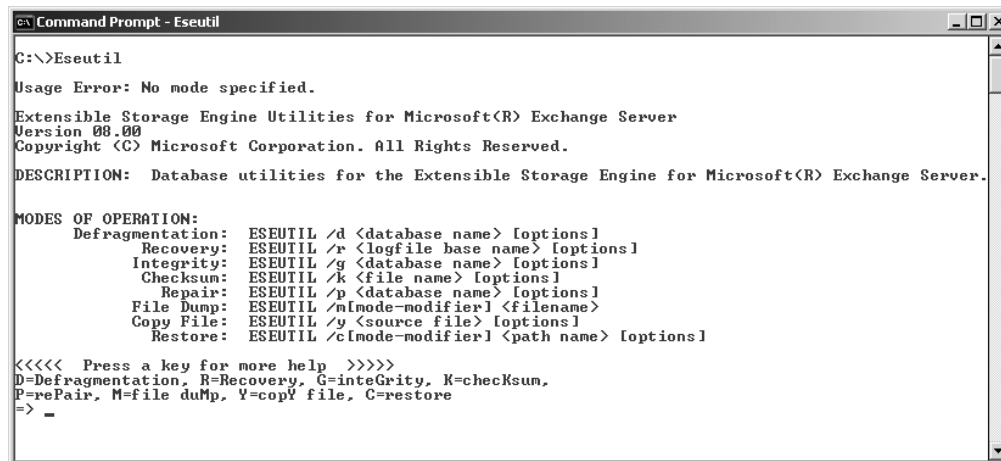
It's beyond the scope of this book to show the steps necessary to restore a database to its last good known state using a combination of a full backup set and incremental or differential backups.

Repairing a Corrupt or Damaged Exchange 2007 Database Using Eseutil

There may be situations where you either don't have a proper backup set to restore a particular database from, or perhaps you found out that the database you just restored to replace a corrupt or damaged database is also corrupt or damaged. This is where Extensible Storage Engine Utilities for Microsoft Exchange Server (Eseutil) comes in. Eseutil is a command-line utility that can be used to perform a range of database tasks including repair, offline defragmentation, and integrity checks. Eseutil hasn't changed much from Exchange 2003 since Exchange still uses ESE databases when speaking Exchange 2007. This means that pretty much all of the switches and parameters available in Eseutil are the same as in previous versions. Since there are plenty of books and online documentation describing how you should approach fixing a corrupt database using Eseutil, I won't include comprehensive information on how to use this utility in this book. Instead, I'll provide you with the most common Eseutil switches, as well as a few examples.

Eseutil, as in previous versions, is located in the Bin folder under your Exchange installation path, which in Exchange 2007, by default, is C:\Program Files\Microsoft\Exchange Server. However, you no longer need to run the tool from that path; you can just open a Command Prompt window and type **Eseutil**, as shown in Figure 9.9.

Figure 9.9 Eseutil Modes



```

C:\>Eseutil

Usage Error: No mode specified.

Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.00
Copyright (C) Microsoft Corporation. All Rights Reserved.

DESCRIPTION: Database utilities for the Extensible Storage Engine for Microsoft(R) Exchange Server.

MODES OF OPERATION:
  Defragmentation: ESEUTIL /d <database name> [options]
  Recovery: ESEUTIL /r <logfile base name> [options]
  Integrity: ESEUTIL /g <database name> [options]
  Checksum: ESEUTIL /k <file name> [options]
  Repair: ESEUTIL /p <database name> [options]
  File Dump: ESEUTIL /m[mode-modifier] <filename>
  Copy File: ESEUTIL /y <source file> [options]
  Restore: ESEUTIL /c[mode-modifier] <path name> [options]

<<<< Press a key for more help >>>>
D=Defragmentation, R=Recovery, G=integrity, K=checksum,
P=rePair, M=file duMp, Y=copY file, C=restore
=> -
  
```

NOTE

You can also run Eseutil directly from the Exchange Management Shell.

496 Chapter 9 • Disaster Recovery with Exchange Server 2007

Before we move on, we want to stress that it's very important you always try to restore your databases from a backup if possible, since there's a good chance you will lose some data when performing a repair of a database. The reason for this is that Eseutil often needs to discard rows from tables or even entire tables. In addition, you should have a repaired database running in your production environment only for a temporary period, which means that after you have repaired a database, you should move all mailboxes from the database to a new one. Needless to say, you should also be sure to make a copy of the database before performing a repair using Eseutil.

NOTE

Did you know that when a database corruption occurs, 99.9 percent of the time it's caused by the underlying hard disk drive subsystem? Yes, it's true! This means there's a pretty good chance the database corruption experienced is caused by an I/O issue on the disk set in your Exchange 2007 server. You should therefore always examine the Application and System logs, searching for any events that might indicate this to be the problem.


Eseutil /P can, in addition to the Mailbox and Public Folder databases, also be run against the ESE database-based message queues on either a Hub Transport or Edge Transport server in your Exchange 2007 organization.

To repair a corrupted or otherwise damaged database, run Eseutil with the /P switch. So, to repair a database called Mailbox Database.edb located in E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group, you would need to type:

Eseutil /P "E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\Mailbox Database.edb"

After pressing **Enter**, you would receive the warning message shown in Figure 9.10.

Figure 9.10 An Eseutil Repair Warning

**NOTE**

You must have the necessary amount of free space (equal to 110 percent of the database file size) on the disk containing the database before you can run Eseutil /P and Eseutil /D.

Click **OK** to proceed, and then wait until Eseutil has repaired the database. If the database is completed successfully, it's highly recommended you perform a full backup of the database, since restoring a backup made before the repair would roll the database back to the state it was in at the time of the backup, which wouldn't be very smart.

After you have run Eseutil /P against a database, also run Eseutil /D in order to fully rebuild indexes and defragment the database. In order to run Eseutil /D against the database, type:

Eseutil /D "E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\Mailbox Database.edb"

When an offline defragmentation has been completed, there's one additional thing to do: repair the database at the application level (repair information and relationships between mailboxes, folders, items, and attachments) by running the Information Store Integrity Checker (Isinteg) utility with the *-fix* parameter. Figure 9.11 shows the parameters and syntaxes available for the Isinteg utility.

Figure 9.11 Isinteg Switches

```

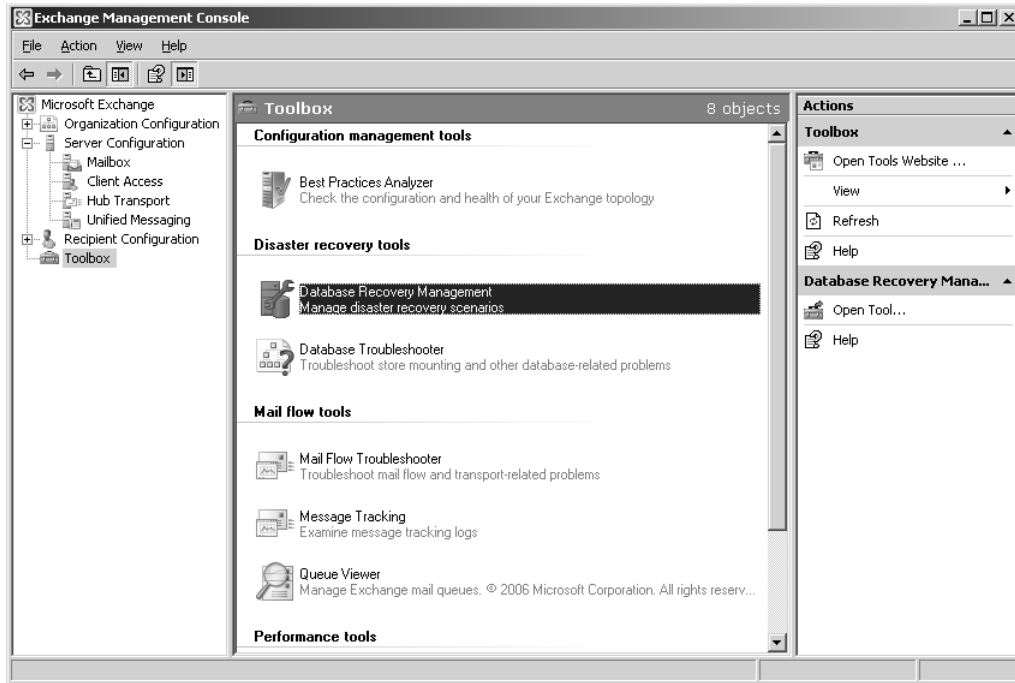
C:\>isinteg
Microsoft Exchange Information Store Integrity Checker v08.00.0685.024
Copyright (c) 1986-2000 Microsoft Corp. All rights reserved.
Usage:
isinteg -s ServerName [-fix] [-verbose] [-l logfilename] -test testname[[, testname]...]
-s ServerName
-fix check and fix (default - check only)
-verbose report verbosely
-l filename log file name (default - .\isinteg.pri/pub)
-t refdblocation (default - the location of the store)
-test testname,...
  folder message aclitem mailbox<pri only> delfld acllist
  rcvfld<pri only> timedev rowcounts attach morefld ooflist<pri only>
  global searchq dlvrto replstate<pub only>
  peruser artidx<pub only> search newsfeed<pub only> dumpsterprops
  Ref count tests: msgref msgsoftref attachref acllistref aclitenref
  newsfeedref<pub only> fldrcv<pri only> fldsub dumpsterref
  Groups tests: allfoldertests allacltests
isinteg -dump [-l logfilename] (verbose dump of store data)

C:\>_

```

If you aren't comfortable running the Eseutil and Isinteg utilities manually on your databases, you also have the option of performing a repair using a wizard-driven interface. This is where the new Disaster Recovery Management tool, a sibling of tools such as the Exchange Best Practices Analyzer Tool (ExBPA), comes into play. To invoke this tool, click the **Toolbox** work center node in the navigation tree in the **Exchange Management Console**, then open the tool by selecting it in the Result pane and clicking **Open Tool** in the **Actions** pane (Figure 9.12).

Figure 9.12 Disaster Recovery Management Tool



The tool will now check if there is any tool or configuration file updates available on Microsoft.com, and if so, apply them without requiring a restart. Once any updates have been applied, click the **Go to Welcome Screen** link, then enter an identifying label for the activity, and click **Next**. When the tool has connected to the Active Directory, you will be presented with the task list shown in Figure 9.13. Here, you should select the **Repair Database** task.

Now select the storage group that contains the database you wish to repair, click **Next**, and on the **Select Databases to Repair** page, check the respective database, as I did in Figure 9.14. Then, click **Next**.

Figure 9.13 Exchange Troubleshooting Assistant Tasks

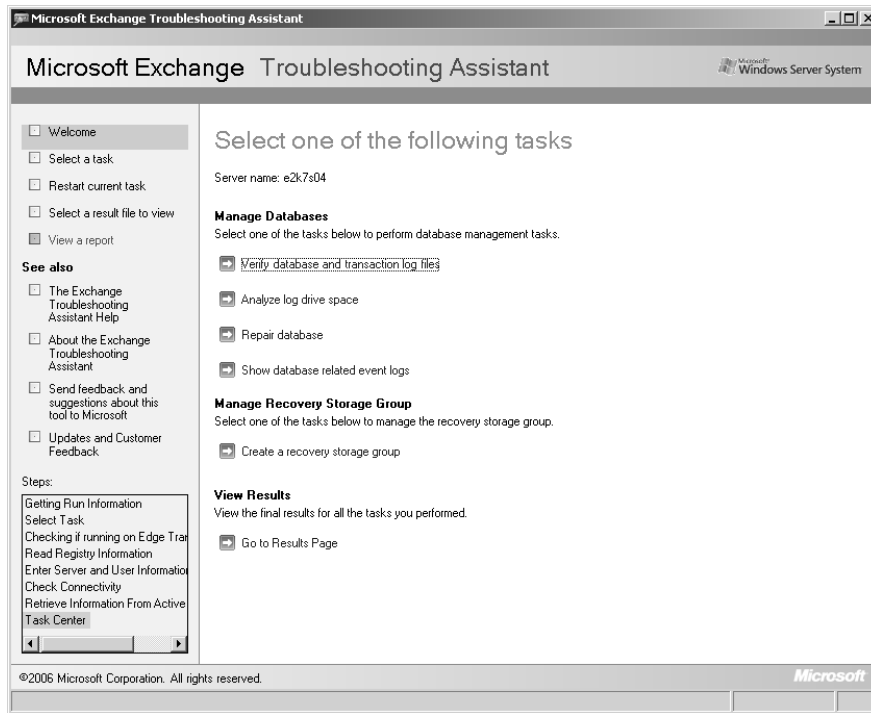
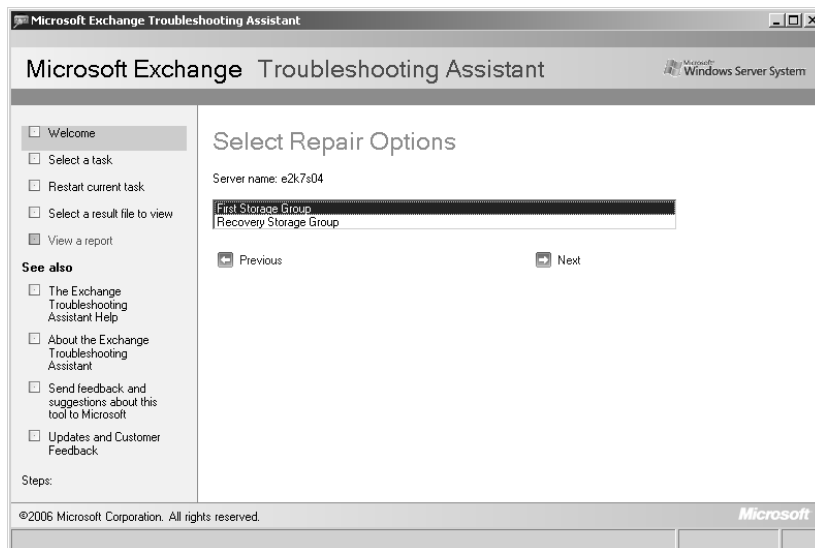



Figure 9.14 Selecting the Database to Repair



500 Chapter 9 • Disaster Recovery with Exchange Server 2007

You will now need to read a repair task warning. I suggest you read it carefully. When you have done so, choose **Continue to Perform Repair Task**, and then click **OK** in the confirmation dialog box shown in Figure 9.15.

Figure 9.15 ExTRA Confirmation



The tool will now run Eseutil /P and then Eseutil /D, followed by Isinteg -fix -test all-tests against the respective database, just like we did manually earlier in this section. After a while, depending on the size of the database, you will be taken to a Report Repair Results page where you can see if the actions completed without any issues, and if not, it will show an explanation why it didn't.

Restoring Mailbox Data Using the Recovery Storage Group Feature

The Recovery Storage Group (RSG) feature, which was originally introduced back in Exchange 2003, gives you, the Exchange administrator, the option of mounting a second copy of a mailbox database (typically a mailbox database restored from backup) so you can extract data from one or more mailboxes in the respective database during working hours without affecting the production databases.

Depending on how much you have used the new Exchange 2007 Management Console (EMC), you may have noticed you can no longer create an RSG from within the EMC. With Exchange 2007, this is instead done using the new Database Recovery Management tool, which as you saw in the previous section, is found under the Exchange Toolbox work center, or by using the Exchange Management Shell (EMS).

When mounting a copy of a Mailbox database to an RSG, you can extract the data from a mailbox and then merge the data with another mailbox located in a mailbox database in a production storage group. You can also extract the data and copy it to a specific folder in another mailbox. With Exchange 2003 RTM, the data was extracted, copied, and merged with another mailbox or mailbox folder using the Microsoft Exchange Server Mailbox Merge Wizard (ExMerge) tool, but in Exchange 2003 SP1 the process was integrated into the Exchange 2003 System Manager GUI.

There are a few things you should be aware of when dealing with RSGs. First, they cannot be accessed by any protocols other than MAPI, and although they can be accessed using MAPI, this doesn't mean you can connect to a mailbox stored in a recovery database using an Outlook MAPI client. MAPI is strictly used to access mailboxes using the Exchange Troubleshooting Assistant and the respective Exchange Management Shell cmdlets. In addition,

you should be aware that you still cannot use RSGs to restore Public Folder data, only mailbox data. It's also worth mentioning that even though you can create up to 50 storage groups on an Exchange 2007 Enterprise edition server, you're limited to one RSG per server. However, it's supported to add multiple mailbox databases to an RSG as long as all databases belong to the same storage group. Finally, you should note that although it's possible to add a restored mailbox database to an RSG on another Exchange 2007 server, it's important you understand that the Exchange 2007 server must belong to the same Active Directory forest.

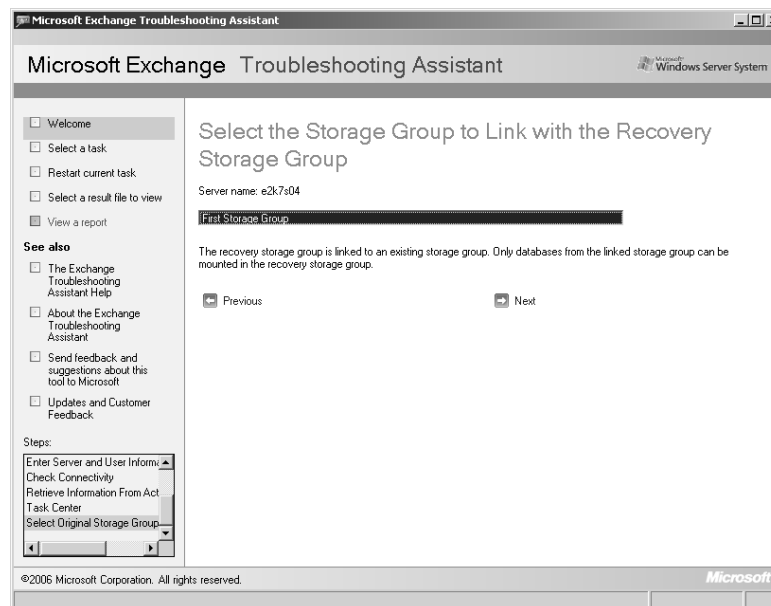
With the preceding in mind, let's move on and see how you manage RSGs.

Managing Recovery Storage Groups Using the Exchange Troubleshooting Assistant

You can create a Recovery Storage Group (RSG) either by using the Disaster Recovery Management tool, which is based on the Microsoft Exchange Troubleshooting Assistant (ExTRA), or by running the *New-StorageGroup* cmdlet with the *-Recovery* parameter in the Exchange Management Shell.

To create the RSG using the Disaster Recovery Management tool, you should first launch it from beneath the Toolbox work center in the navigation tree of the Exchange Management Console (EMC). Let the tool check for any tool or configuration file updates available, and then click the **Go to Welcome** screen link. Enter an identifying label for this activity (such as Create RSG), and then click **Next**. In the **Tasks** list that appears, click **Create a Recovery Storage Group**, and then select the storage group you want to link with the recovery storage group, as shown in Figure 9.16. Then, click **Next** once again.

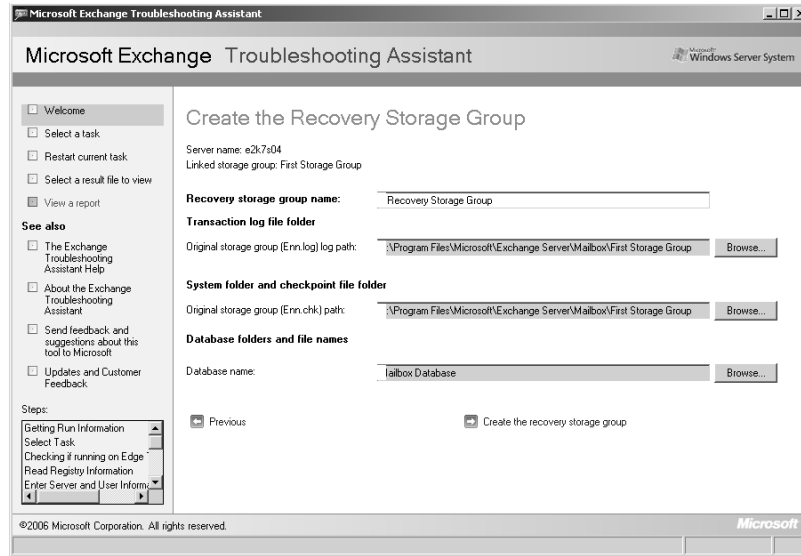
Figure 9.16 Selecting the Storage Group to Link with the RSG



502 Chapter 9 • Disaster Recovery with Exchange Server 2007

Now it's time to create the RSG, but before doing so you need to give it a name (the default name is Recovery Storage Group, which should be okay in most situations). When you have entered an appropriate name, click **Create the recovery storage group** (Figure 9.17).

Figure 9.17 Creating the RSG



After a little while, you will be presented with a screen similar to the one in Figure 9.18, and the RSG for the respective Mailbox database has now been created.

With the RSG created, we can move, copy, or restore database and transaction log files to the recovery storage group paths. To see the path for the recovery storage group log and database files, click **Show Create Recovery Storage Group Information**. By default, the path is C:\Program Files\Microsoft\Exchange Server\Mailbox\<Storage Group>\RSGxxxxxxx, as you can see in Figure 9.19. The RSGxxxxxxx folder will appear empty in Windows Explorer until you have moved, copied, or restored the database and transaction log files to it.

Figure 9.18 RSG Result

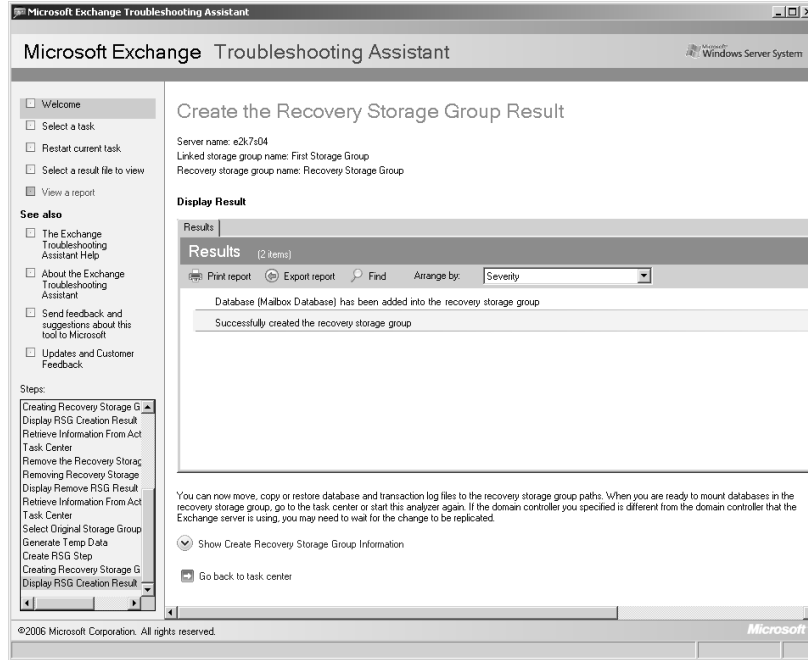
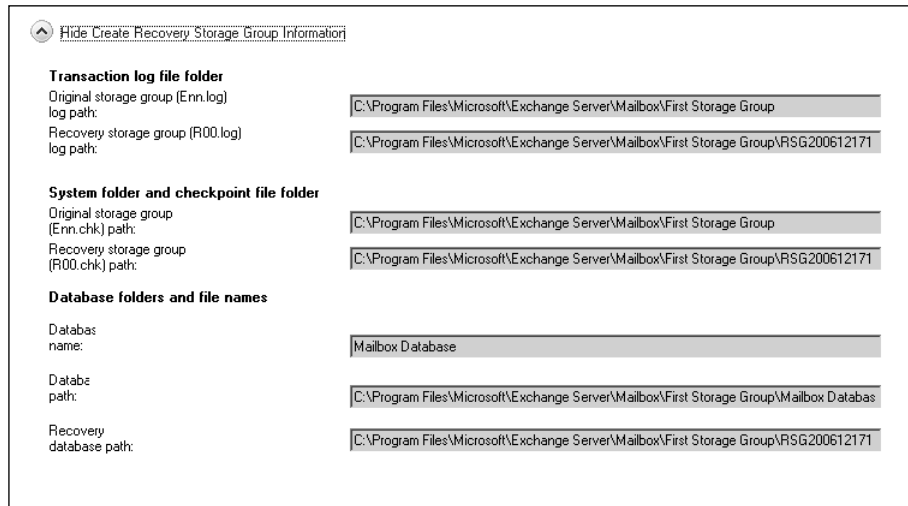


Figure 9.19 Storage Group and Recovery Storage Group Paths



For the purpose of this example, we will restore a Mailbox database from a backup using the Windows 2003 Backup tool. So let's launch the Windows 2003 Backup tool in advanced mode, and then click the **Restore and Manage Media** tab. Here we need to select the

504 Chapter 9 • Disaster Recovery with Exchange Server 2007

Mailbox database and log files we want to restore. When you have done so, click the **Start Restore** button.

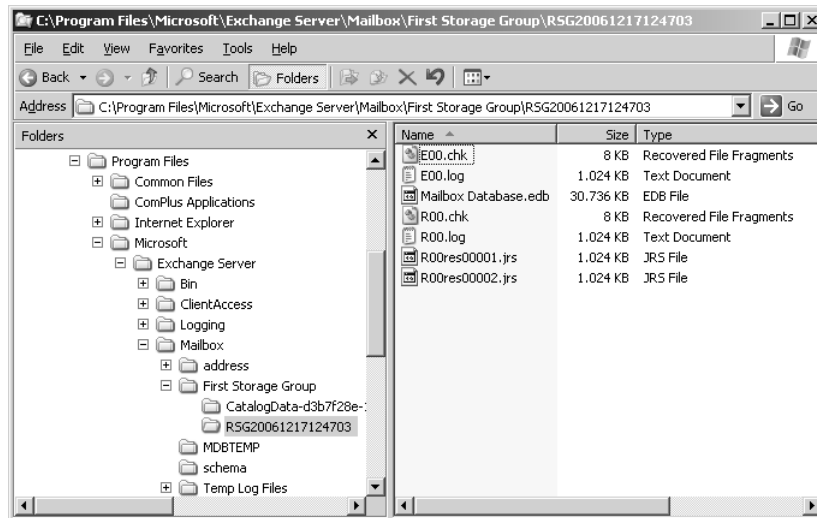
NOTE

Note that the Restore Files To: Drop-Down box is set to Original Location. Also notice we cannot change this selection. But does that mean the Mailbox database currently in production will be replaced by the one we restore from backup? No, this is not the case. First, we haven't dismantled the production Mailbox database, and second, we haven't enabled the *This Database Can Be Overwritten By A Restore* option on the Mailbox database property page. Because of this, the Mailbox database will be restored to the recovery storage group we just created.

Now specify the Exchange Server to which you want to restore the respective Mailbox database, and then enter a temporary location for the log and patch files. Lastly, check **Last Restore Set** (Log File Replay will start after this restore completes) since this is the last restore set. When you are done, click **OK** and wait for the restore job to complete. Then, click the **Close** button.

The respective files have now been restored to the RSGxxxxxxx folder, as you can see in Figure 9.20.

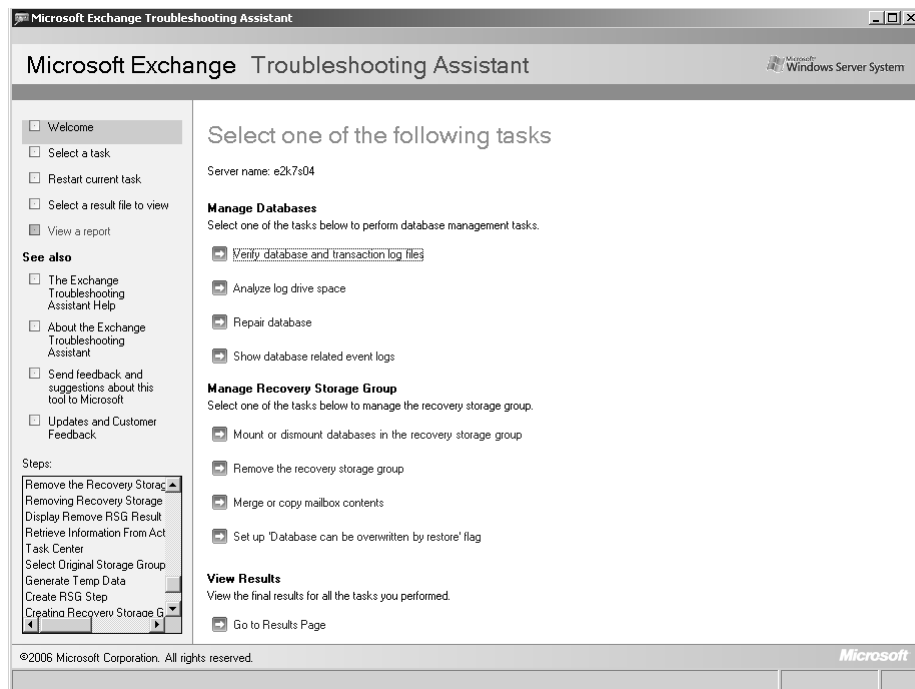
Figure 9.20 The Restored Mailbox Database in Windows Explorer



Since we didn't check the *Mount Database After Restore* option, the Mailbox database will now be in a dismounted state. With this in mind, let's switch back to the ExTRA Task Center. As shown in Figure 9.21, we now have several new recovery storage group–related tasks available. Since the Mailbox database needs to be mounted before we can extract data from it, we have to click **Mount or dismount databases in the recovery storage group**.

On the **Mount or Dismount Database** page, check the respective Mailbox database and click **Mount selected database** (Figure 9.22).

Figure 9.21 Selecting Mount or Dismount Databases in the Recovery Storage Group



Once the Mailbox database has been mounted, click **Go back to task center**, and then select **Merge or copy mailbox content**. This will bring us to a screen similar to the one shown in Figure 9.23, here you should just make sure the Mailbox database you wish to extract data from is selected, and then click **Gather merge information**.

Figure 9.22 Mounting the Mailbox Database Using the ExTRA Tool

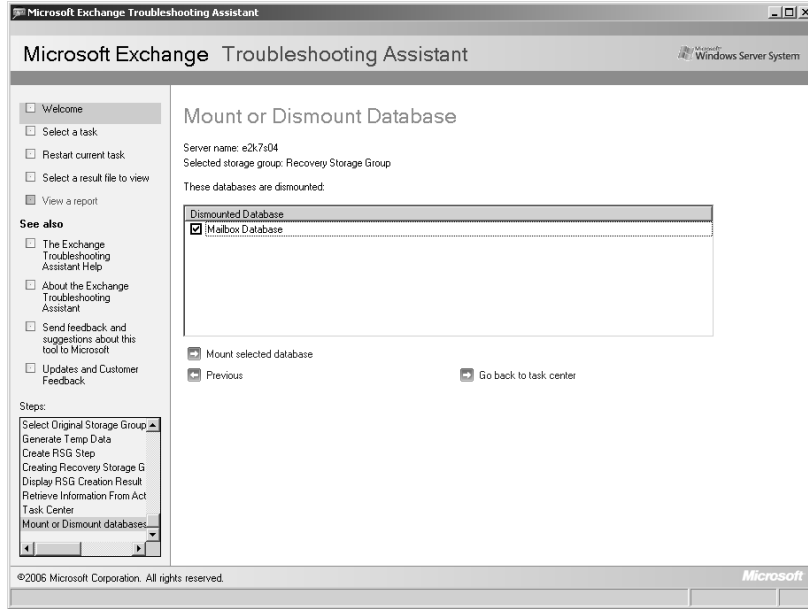
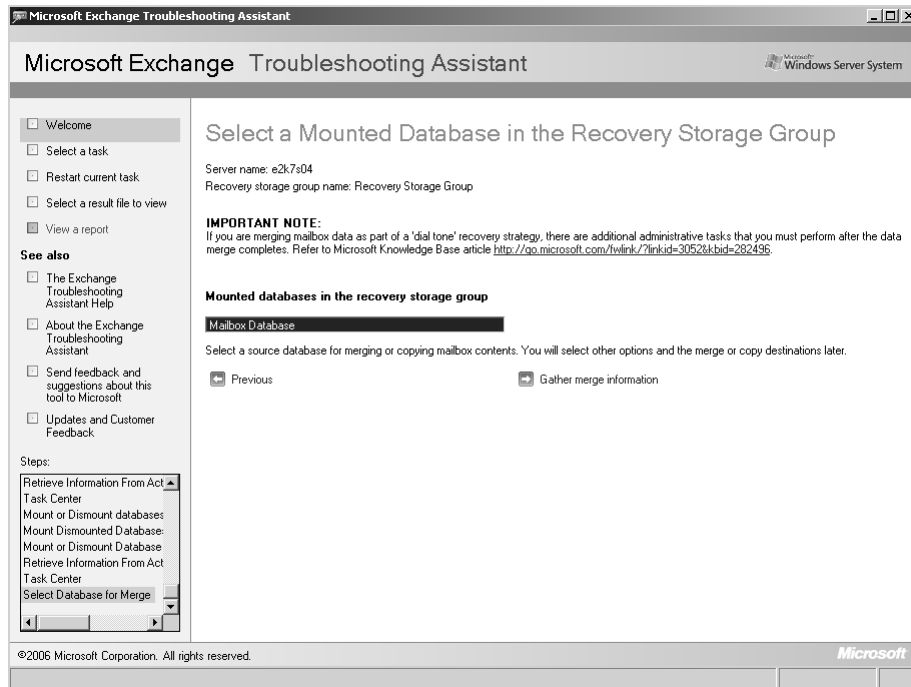
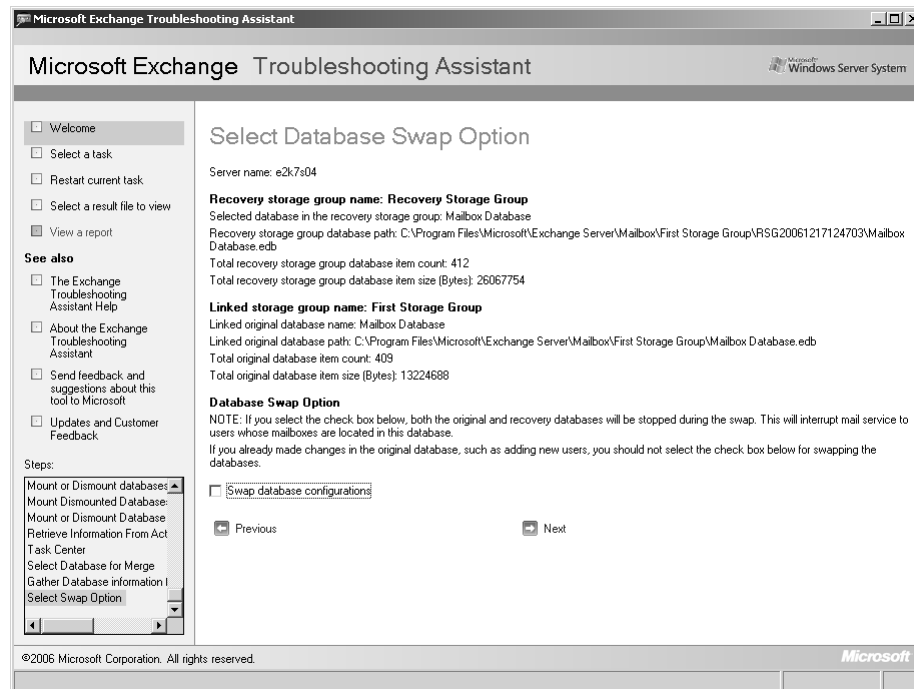


Figure 9.23 Selecting a Mounted Database in the Recovery Storage Group



We now have the option of swapping the Mailbox database mounted to the RSG and the linked production Mailbox database (a recommended step if you're performing a dial-tone database restore) by checking Swap Database Configurations, as can be seen in Figure 9.24. Since this option will swap the two databases, both of them need to be dismounted, which will affect mail service to the end users whose mailboxes are stored in the respective database.

Figure 9.24 The Database Swap Option



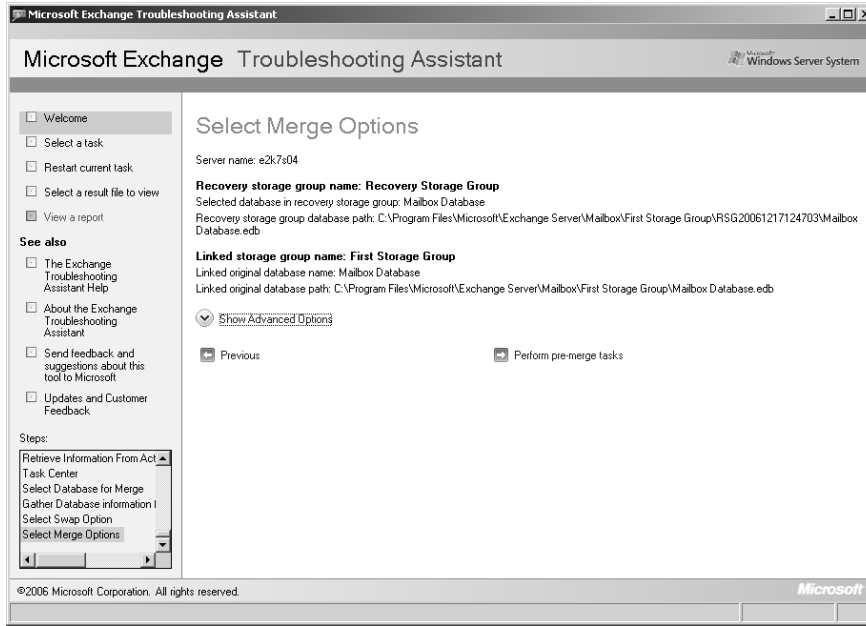
Since we aren't dealing with a dial-tone database restore in this example, just click **Next**. On the **Select Merge Options** page, click **Perform pre-merge tasks** (Figure 9.25).

NOTE

Note that you have the option of clicking Show Advanced Options. Under the Advanced options, we can specify different match and filtering options, as well as the bad item limit. This is also the place where you specified whether all merge mailbox data should be merged to the respective mailboxes in the production Mailbox database, or whether they should be copied to a single target mailbox.

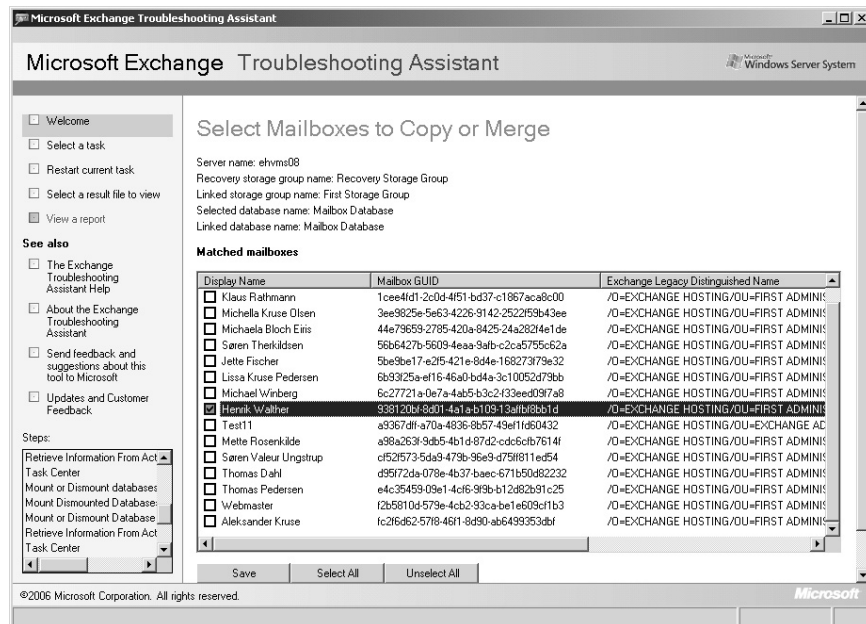
508 Chapter 9 • Disaster Recovery with Exchange Server 2007

Figure 9.25 Specifying Merge Options



The final step is to select the mailboxes you want to merge. You do this by checking the box to the left of each user name in the list, as shown in Figure 9.26.

Figure 9.26 Selecting the Mailboxes to Merge



Now wait for the tool to merge the mailbox data from the Mailbox database in the recovery storage group for the selected mailbox. When the mailbox data merge has completed, you should be able to see the content deleted from the production Mailbox database. You don't even need to restart the Outlook or OWA client for the restored data to appear!

When you have merged or copied the required Mailbox data, you can use ExTRA to dismount and then remove the recovery storage group. Be sure you delete the files in the RSG\xxxxxxx folder after you have removed it so the files don't take up valuable disk space.

Managing Recovery Storage Groups Using the Exchange Management Shell

As mentioned earlier in this chapter, you can also manage an RSG using the Exchange Management Shell (EMS). If you know your cmdlets, restoring mailbox data from a Mailbox database in a recovery storage group can be done a lot faster than when you're using ExTRA.

The first step is to create the RSG. In order to create an RSG via the EMS, you need to run the *New-StorageGroup* cmdlet with the *-Recovery* parameter. So, to create an RSG for the first storage group on a server named E2K7S04, type:

```
New-StorageGroup -Server E2K7S04 -LogFolderPath "E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG -Name "Recovery Storage Group" -SystemFolderPath "E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG" -Recovery
```

The *LogFolderPath* and *SystemFolderPath* parameters are used to specify where the RSG-related files should be located. As you can see, we specified they should them to be restored to a subfolder called RSG under E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG. If you intend to do the same, please make sure there's sufficient disk space available for the Mailbox database you're restoring from backup.

To see if a respective storage group is a recovery storage group (as well as many other types of information), you can use the *Get-StorageGroup* <storage group name> | *FL* command. If the storage group is a recovery storage group, it will say True under Recovery, as shown in Figure 9.27

The next step is to add a recovery database (either moved, copied, or restored from backup) to the RSG, this is done by running the *New-MailboxDatabase* cmdlet with the *MailboxDatabaseToRecover* parameter. So, to add a recovery database to the recovery storage group on a server named E2KS04 with the edb file path pointing to E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG, type:

```
New-MailboxDatabase -MailboxDatabaseToRecover "Mailbox Database" -StorageGroup "E2K7S04\Recovery Storage Group" -EDBFilePath "E:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG\Mailbox Database.edb"
```

Figure 9.27 Full List of Recovery Storage Group Information

```

Machine: e2k7s04 | Scope: exchangehosting.dk
C:\>Get-StorageGroup "Recovery Storage Group" | FL

LogFolderPath           : C:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG
                        : 20061217124703
SystemFolderPath        : C:\Program Files\Microsoft\Exchange Server\Mailbox\First Storage Group\RSG
                        : 20061217124703
CircularLoggingEnabled  : False
ZeroDatabasePages      : False
LogFilePrefix          : R00
LogFileSize             : 1024
RecoveryEnabled        : True
OnlineDefragEnabled     : True
IndexCheckingEnabled   : True
EventLogSourceID       : MSEXCHANGES
LogCheckpointDepth     : 20971520
CommitDefault          : False
DatabaseExtensionSize  : 256
PageFragment           : 8
PageTempDBMinimum     : 0
Server                 : E2K7S04
ServerName             : E2K7S04
CopyLogFolderPath      :
CopySystemFolderPath   :
Recovery               : True
Name                   : Recovery Storage Group
Replicated              : None
HasLocalCopy           : False
MinAdminVersion        : -2147453113
AdminDisplayName       :
ExchangeVersion        : 0.1 (8.0.535.0)
DistinguishedName     : CN=Recovery Storage Group,CN=InformationStore,CN=E2K7S04,CN=Servers,CN=Exc
                        hange Administrative Group (FYDIBOHF23SPDLT),CN=Administrative Groups,CN=E
                        xchange Hosting,CN=Microsoft Exchange,CN=Services,CN=Configuration,DC=exch
                        angehosting,DC=dk
Identity              : E2K7S04.Recovery Storage Group
Guid                  : b94ce131-faa0-411d-88d6-e80ad33b5159
ObjectCategory        : exchangehosting.dk/Configuration/Schema/ms-Exch-Storage-Group
ObjectClass            : <top, container, msExchStorageGroup>
WhenChanged           : 17-12-2006 12:47:12
WhenCreated            : 17-12-2006 12:47:12
OriginatingServer     : ehvms07.exchangehosting.dk
IsValid               : True
  
```

With the Mailbox Database created in the recovery storage group, we now need to configure it to allow overwrites by running the *Set-MailboxDatabase* cmdlet with the *-AllowRestore* parameter. To allow file restores for the recovery database just created, type:

```
Set-MailboxDatabase -Identity "E2K7S04\Recovery Storage Group\Mailbox Database" -AllowFileRestore $true
```

Now that we have created a recovery database in the recovery storage group and allowed it to be overwritten by a file restore, it's time to restore the mailbox database version from which you want to extract and copy or merge data to the mailbox database in production. To do so, launch the Windows 2003 Backup tool and restore the respective Mailbox database version using the same steps as we did when we used the ExTRA to recover Mailbox data.

We now need to mount the restore Mailbox database using the *Mount-Database* cmdlet. In order to do so, type:

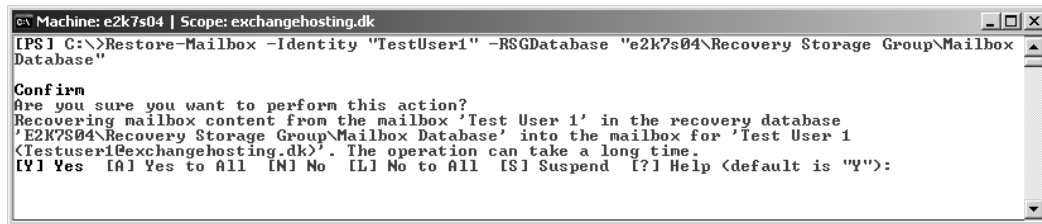
```
Mount-Database -Identity "E2K7S04\Recovery Storage Group\Mailbox Database"
```

With the Mailbox database mounted, we can now extract Mailbox data from it. For example, if you want to merge the mailbox data of an existing user in the recovery database to the production Mailbox database, you need to type:

Restore-Mailbox -Identity <username> -RSGDatabase “servername\RSG name\database name”

In Figure 9.28, we recovered mailbox data for a user called Test User 1 on a server named E2K7S04.

Figure 9.28 Restoring Mailbox Data from a Mailbox in a Recovery Storage Group



```

c:\ Machine: e2k7s04 | Scope: exchangehosting.dk
[PS] C:\>Restore-Mailbox -Identity "TestUser1" -RSGDatabase "e2k7s04\Recovery Storage Group\Mailbox Database"

Confirm
Are you sure you want to perform this action?
Recovering mailbox content from the mailbox 'Test User 1' in the recovery database 'E2K7S04\Recovery Storage Group\Mailbox Database' into the mailbox for 'Test User 1 (Testuser1@exchangehosting.dk)'. The operation can take a long time.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
```

NOTE

Depending on the size of the mailbox to be recovered, this merging process can take a long time.

If you need to recover mailbox data for all users in the RSG, you would need to use the following command:

Get-MailboxStatistics -Database “Recovery Storage Group\Mailbox Database” | Restore-Mailbox

Let’s suppose the mailbox in the recovery database that you want to recover data from has in the meantime been deleted from the production Mailbox database. In this case, you have the option of recovering the mailbox data to a target folder in another mailbox by using the following command:

Restore-Mailbox -RSGMailbox “Test User 1” -RSGDatabase “servername\RSG name\database name” -Identity “Test User 2” -TargetFolder “Test User 1 Recovered data”

Just as with recovering data using the ExTRA tool, when using the Exchange Management Shell you should remember to remove the RSG after the required data has been recovered. To do so, first run the command to remove the recovery database:

Remove-MailboxDatabase -Identity “E2K7S04\Recovery Storage Group\Mailbox Database”

Click **Yes** to the confirmation warning, and then type the following command in order to remove the RSG:

Remove-StorageGroup -Identity “E2K7S04\Recovery Storage Group”

Finally, delete the RSG folder manually using Windows Explorer.

Recovering an Exchange 2007 Server Using the RecoverServer Switch

What could be worse than facing one or more seriously corrupted Exchange 2007 mailbox databases? Yes, you guessed right: facing a completely dead Exchange 2007 Server. In this section, I'll shine some light on the steps necessary to restore an Exchange 2007 Server that has experienced a major hardware failure causing a complete loss of data. As is the case with Exchange 2000 and 2003, you can recover an Exchange 2007 Server in a fairly straightforward way. As you probably know, we could use the DisasterRecovery switch to recover a dead Exchange 2000 or 2003 Server on new hardware, but with Exchange 2007 this switch no longer exists. Instead, it has been replaced by the new RecoverServer switch, which is similar to the DisasterRecovery switch. The interesting thing about the RecoverServer switch is that it can be used to recover all types of Exchange 2007 Server roles, except the Edge Transport Server role, which uses ADAM and not the Active Directory to store configuration data.

NOTE

To recover a server with the Edge Transport Server role installed, you must use the cloned configuration tasks to export and import configuration information. You can read more about the cloned configuration tasks in Chapter 7.

When you run Setup with the RecoverServer switch on a new Windows 2003 Server that is configured with the same name as the one that has crashed or is permanently down for some reason, Setup will read the configuration information for the respective Exchange 2007 server from the Active Directory. In addition to applying the roles and settings stored in Active Directory, Setup will, as is the case when installing an Exchange 2007 Server role without the RecoverServer switch, install the Exchange files and services required for the respective Exchange 2007 server role(s). This means that local customizations done on the server (such as Mailbox databases, Receive connectors, custom OWA settings, SSL certificates, and so on) need to be re-created or recovered manually afterward.

In this section, we'll go through the steps necessary to recover an Exchange 2007 server with the Hub Transport, Mailbox Server, and Client Access Server roles installed.

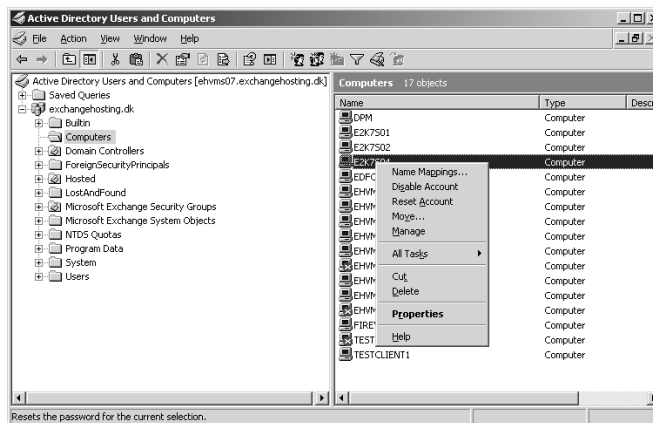
Restoring and Configuring the Operating System

When you have received a replacement server or replacements for the failed hardware components, it's important you configure and partition the disk sets in the new server so they are identical to the way they were configured in the failed server. When the hardware is configured according to the documentation you wrote for the failed Exchange 2007 (which you did write, right?), we can begin installing the operating system from the Windows 2003 Server 64-bit media. When Windows 2003 Server has been installed, it's important you install the Windows Components required by the Exchange Server 2007 Server roles, as well as any service packs and Windows updates that were applied on the failed server. For details about which Windows components are required for each server role, refer back to Chapter 2.

In addition to that already mentioned, you should also make sure you name the new server with the same server name. Before doing so, however, it's important the failed Exchange 2007 server be turned off. In addition, you should add the server to the respective Active Directory domain, first resetting the computer account for the respective Exchange 2007 server. In order to reset the computer account, you must follow these steps.

1. Log on to a domain controller or another server with the Adminpak installed in the Active Directory domain, and then open the **Active Directory Users and Computers (ADUC) MMC** snap-in.
2. In the ADUC MMC snap-in, navigate to the organizational unit (OU) containing your computer accounts (by default, the Computers OU), right-click the computer account that should be reset, and then select **Reset Account**, as shown in Figure 9.29.

Figure 9.29 Resetting the Computer Account in the ADUC MMC Snap-in



3. Click **Yes** to the warning in the dialog box that appears, and then click **OK**.

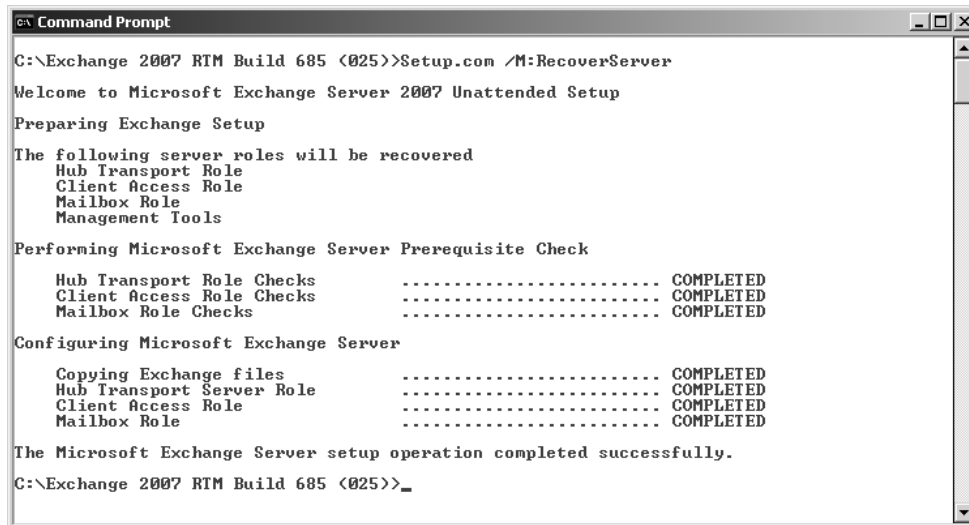
We can now join the new server to the domain without issues. Do so and perform the required reboot.

Installing Exchange 2007 Using the RecoverServer Switch

Now that Windows 2003 has been installed properly, we can move on and start installing Exchange 2007 by running Setup.exe with the RecoverServer switch. In order to do so, perform the following steps:

1. Click **Start | Run** and type **CMD**. Then, press **Enter**.
2. Change to the directory or media containing your Exchange 2007 Setup files, and then type **Setup.com /M:RecoverServer**. As can be seen in Figure 9.30, Exchange 2007 Setup will now prepare the Exchange 2007 setup, and then perform the mandatory prerequisite checks. Finally, it will begin to copy the Exchange files and then configure each Exchange 2007 Server role by reading the required configuration information from Active Directory.

Figure 9.30 Recovering an Exchange 2007 Server Using the RecoverServer Switch



```
Command Prompt
C:\Exchange 2007 RTM Build 685 <025>>Setup.com /M:RecoverServer
Welcome to Microsoft Exchange Server 2007 Unattended Setup
Preparing Exchange Setup
The following server roles will be recovered
Hub Transport Role
Client Access Role
Mailbox Role
Management Tools
Performing Microsoft Exchange Server Prerequisite Check
Hub Transport Role Checks ..... COMPLETED
Client Access Role Checks ..... COMPLETED
Mailbox Role Checks ..... COMPLETED
Configuring Microsoft Exchange Server
Copying Exchange files ..... COMPLETED
Hub Transport Server Role ..... COMPLETED
Client Access Role ..... COMPLETED
Mailbox Role ..... COMPLETED
The Microsoft Exchange Server setup operation completed successfully.
C:\Exchange 2007 RTM Build 685 <025>>_
```

NOTE

If you're recovering an Exchange 2007 server with the Hub Transport Server role installed, and this is the only Exchange 2007 server with this role installed, it's recommended you run `Setup.com /M:RecoverServer` with the `/DoNotStartTransport` syntax since there's a few post-recovery steps that should be completed before this role is made active.

When the Exchange setup has completed each phase successfully, we're close to calling the server recovery a success. However, there are a few post-recovery steps that need to be finished, depending on what Exchange 2007 Server roles are installed on the server. It's obvious a recovered server with the Mailbox Server role must have the respective Mailbox database and Public Folder database restored from backup, or copied back from the disks on the old server (if possible). If the Public Folders are replicated with other Exchange 2000/2003 or 2007 servers in the Exchange organization, you don't need to restore it since an empty Public Folder database will be backfilled from the other Public folder server(s).

NOTE

If you need to restore one or more Mailbox and/or Public Folder databases to the recovered server using the Windows 2003 Backup tool, note that you must catalog the respective backup (.BKF). This is done by selecting the **Restore and Manage Media** tab, and then clicking **Tools | Catalog** a backup file in the menu.

If the Hub Transport Server role is installed on the recovered Exchange 2007 server, you may also need to restore any saved message queue databases (which in Exchange 2007 are stored in an ESE database and not in the NTFS file system as was the case with Exchange 2000 and 2003) and place them in the right folder (should be done while the Microsoft Exchange Transport service is stopped, which is why it's a good idea to run the `RecoverServer` switch with the `/DoNotStartTransport` syntax if you're recovering an Exchange 2007 server with the Hub Transport Server role installed), as well as reconfigure any Receive connectors since these are stored locally on the Hub Transport Server and not in Active Directory, as is the case with Send Connectors. In addition, you may need to restore the Client Access Server settings (custom OWA files and/or virtual directories). Custom virtual folder settings can be restored by using the script method mentioned earlier in this chapter.

SOME INDEPENDENT ADVICE

Although it should be the most comprehensive, as well as fastest, way to recover a server using the RecoverServer switch, it's worth mentioning that it's fully supported to restore an Exchange 2007 Server by restoring the System State as well as all the Exchange installation files. Bear in mind, however, that this method requires you restore Exchange 2007 on the same hardware.

Recovering an Exchange 2007 Cluster Using the RecoverCMS Switch

To finish off this chapter, we wanted to talk a little about how you can recover an Exchange 2007 clustered mailbox server (both CCR and SCC) by using the *ExSetup.exe* command with the RecoverCMS switch. Since we're talking about restoring a cluster, many of you may think the tasks involved are terribly complex. As a matter of fact, it's a relatively simple task. The biggest challenge is rebuilding the Windows 2003 cluster itself, which as you learned in Chapter 8, is a rather harmless process. Once you have rebuilt the Windows 2003 cluster on new hardware, you need to install the Passive Clustered Mailbox Role on one of the Windows 2003 cluster nodes, navigate to the Exchange Bin folder (which, by default, is located under C:\Program Files\Microsoft\Exchange Server\), and then run the following command:

```
ExSetup.exe /RecoverCMS /CMSName:<name of the clustered mailbox server> /CMSIPAddress:<IP address of the clustered mailbox server>
```

When the clustered mailbox server has been recovered successfully (if the recovered clustered mailbox server is based on a CCR), you need to enable replication as replication, which, by default, will be in a suspended state after recovery using the RecoverCMS switch. In addition you must (both when recovering a CCR and SCC) start the Exchange System Attendant service manually since it will stop right after the clustered mailbox server has been recovered.

The next step is to restore the respective Mailbox and/or Public Folder databases that existed on the failed clustered mailbox server from backup, or move/copy them from their respective locations.

NOTE

If you're recovering a Single Copy Cluster (SCC) and stored the Mailbox and Public Folder databases on a storage area network (SAN), you won't need to restore the databases from backup as long as each node points to the same shared storage subsystem that the failed clustered mailbox server did.

When any required Mailbox and/or Public Folder databases have been restored, you should now install the Passive Clustered Mailbox Role on the second node (and if recovering an SCC, any additional nodes). If you recovered a clustered mailbox server that is based on SCC, we can now call the recovery of the clustered mailbox server a success, but if you use CCR, there's one final task to complete, and that is to reseed the replica and resume replication. To reseed the second copy of the database(s), you should run the following command in the Exchange Management Shell:

Update-StorageGroupCopy -Identity: <Servername\Name of StorageGroup>

When the storage group(s) have been reseeded, you can resume replication by running:

Resume-StorageGroupCopy -Identity:<Servername>\Name of Storage Group>

So, this was not as difficult as you had imagined it, right?

Restoring Mailbox Databases Using the Improved Database Portability Feature

As those of you with plenty of disaster recovery experience from Exchange 2003 might be aware, Mailbox database portability (that is mounting a Mailbox database to an alternative Exchange Server) was rather limited in this version of Exchange, actually the only options available were to mount the respective Mailbox database into a recovery storage group (RSG), into a storage group on a server with the same name as the failed server, or into the storage group on an Exchange Server in the same administrative group. Although mailbox databases were portable between Exchange 2003 servers (on the same service pack level) in the same administrative group, certain tasks were involved with this procedure. You had to rename the Mailbox databases appropriately, as well as re-link each mailbox in the database to an Active Directory user account before the mailbox could be accessible to an end user. In addition, several other issues might exist if the Mailbox database contained a System Attendant mailbox. Finally, depending on what type of third-party applications were running on the particular Exchange server, it was also best practice to reboot the server once the Mailbox database move was completed.

With Exchange 2007, the Mailbox database portability feature has been improved drastically. Now you can port and recover a Mailbox database to any server in the Exchange 2007 organization, and because of the new Autodiscover service (which we discussed in Chapter 5), all Outlook 2007 clients will be redirected to the new server automatically the first time they try to connect after the Mailbox database has been mounted on another Exchange 2007 server.

NOTE


Since only Outlook 2007 clients can take advantage of the new Autodiscover service introduced in Exchange 2007, any legacy clients (Outlook 2003 and earlier) won't be redirected to the new server automatically.

Some of you might wonder if Exchange 2007 (unlike Exchange 2003) allows you to port or recover a Public Folder database to another server. The answer is no. Doing so is still not supported since it will break Public Folder replication. The proper method for moving a Public Folder database to another server is to add the respective server to the Public Folder replica list.

Okay, now that you have heard how cool the new Mailbox database portability improvements in Exchange 2007 are, let's take a look at the steps needed they entail:

First, it's important you make sure the Mailbox database you wish to port or recover to another server is in a clean shutdown state. If not, you must perform a soft recovery of the database, which is done by running `Eseutil /R <ENN>` against it. ENN is the prefix of the storage group to which you want to commit any existing transaction log files. One method you can use to find this prefix is to open the property page of the respective storage group containing the Mailbox database you wish to port or recover to another Exchange 2007 server (see Figure 9.31).

Figure 9.31 The Transaction Log Files Prefix

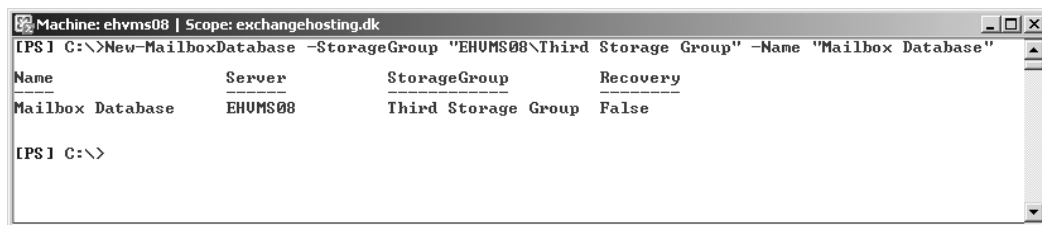


Once the Mailbox database is in a clean shutdown state, the next step is to move the Mailbox database (.EDB file, transaction log files, and Exchange Search catalog) to the system path folder of the respective storage group on the other server, and then create a new Mailbox database in the storage group using the following command:

New-MailboxDatabase -StorageGroup <Servername>\<Name of Storage Group> -Name <Name of Mailbox Database>

In this example, you will mount a database named Mailbox database to the Third Storage Group on an Exchange 2007 Server called EHVMS08. Therefore, the command we need to run is shown in Figure 9.32.

Figure 9.32 Creating a New Mailbox Database in the Third Storage Group



```
Machine: ehvms08 | Scope: exchangehosting.dk
[PS] C:\>New-MailboxDatabase -StorageGroup "EHVMS08\Third Storage Group" -Name "Mailbox Database"
Name                Server              StorageGroup        Recovery
-----                -
Mailbox Database    EHVMS08            Third Storage Group False

[PS] C:\>
```

Because Exchange 2007 won't create an .EDB file for a newly created Mailbox database before it's mounted for the first time, using the *New-MailboxDatabase* cmdlet to create a new Mailbox database, while the Mailbox Database.edb file is placed in the folder of the Third Storage Group will not conflict in any way. Actually, you can just move ahead and mount the ported Mailbox database.

NOTE

It's important that the name of the new Mailbox database you create using the *New-MailboxDatabase* cmdlet matches the name of the Mailbox database you ported or recovered from the old Exchange 2007 Server; otherwise, you won't be able to mount it.

To mount the Mailbox database, you can use the Mount-Database "Mailbox Database" or the Exchange Management Console. When the Mailbox database has been mounted appropriately, there's only one more task to complete, and that is to modify (re-link) the Active Directory user account objects associated with a mailbox in the Mailbox database that we ported to a new server, so they point to the correct server. This can be done by using the following command:

520 Chapter 9 • Disaster Recovery with Exchange Server 2007

Get-Mailbox -Database "E2K7S04\Mailbox Database" | Move-Mailbox -TargetDatabase "EHVMS08\Mailbox Database" -ConfigurationOnly: \$True

You then must confirm that you wish to perform this operation. Type **Y** for Yes, and press **Enter**.

NOTE

If you receive an error when trying to run this command, check to make sure the Mailbox database is mounted on the old Exchange 2007 server.

Now would be a good time to access a few mailboxes (using Outlook 2007 or OWA 2007) stored in the Mailbox database we ported so you can verify the end users still have mailbox connectivity.

Summary

In this chapter, we took a look at how to properly back up the different server roles in Exchange 2007. We then went through how you restore an Exchange 2007 Server with one or more server roles installed, as well as how you can restore a corrupt Mailbox or Public Folder database using the Windows 2003 Backup tool, and if this isn't an option, how you can repair a corrupt database using Eseutil. We also had walked through how you can recover mailbox data using the improved Recovery Storage Group (RSG) feature. In addition, I showed you how it's possible to recover a failed Exchange 2007 server using the RecoverServer and RecoverCMS switches. Lastly, we talked about the improvements that have been made regarding database portability in Exchange 2007.

Solutions Fast Track

Backing Up Exchange 2007 Using Windows 2003 Backup

- ☑ Frequent backups of the Exchange 2007 servers in an organization are important operational tasks, which perhaps can be a bit trivial, but should be taken very seriously. I can only imagine one thing that's worse than a complete failure of an Exchange 2007 server, and that's a complete failure of an Exchange 2007 server without having any backups to restore from.
- ☑ One of the most important things to back up regarding Exchange 2007 Mailbox Servers are the databases that hold user mailboxes and public data.
- ☑ Since Exchange 2007 databases still use ESE, you can (just as with previous versions of Exchange) back them up using the Exchange-aware native Windows 2003 backup tool.
- ☑ Exchange 2007 supports two different backup methods. The first is a legacy streaming backup, which is a backup method based on the ESE application programming interface (API) that allows you to back up one or more storage groups at the same time. However, only one backup job can run against a specific storage group. Then we have the Volume Shadow Copy Service (VSS) backup method, which some of you may know from Exchange 2003, where it was first introduced. The interesting thing about VSS is that this method, in addition to what the legacy streaming backup method offers, can also take an online backup of the copy database when using either Local Continuous Replication or Cluster Continuous Replication in your setup.

Restoring Exchange 2007 Storage Groups and Databases Using Windows 2003 Backup

- ☑ It's important you understand that a restore of a Mailbox database will copy the database file (.EDB) into its original location on the disk, and thereby overwrite any existing .EDB file.
- ☑ Once a restore has completed, the log files will be replayed into the restored version of the database. In addition to the log files, a file called Restore.env will also be copied to the specified temporary folder, and this file is the one that keeps control of which storage group the log files belong to, as well as the database paths and range of log files that have been restored.

Repairing a Corrupt or Damaged Exchange 2007 Database Using Eseutil

- ☑ There may be situations where you either don't have a proper backup set to restore a particular database from, or perhaps you have found out that the database you just restored, in order to replace a corrupt or damaged database, is corrupt or damaged itself. In such situations, you have the option of repairing the database using Extensible Storage Engine Utilities for Microsoft Exchange Server (Eseutil).
- ☑ Eseutil hasn't changed much from Exchange 2003 since Exchange still uses ESE databases when speaking Exchange 2007. This means that pretty much all of the switches and parameters available in Eseutil are the same as in previous versions.
- ☑ As in previous versions, Eseutil is located in the Bin folder under your Exchange installation path, which in Exchange 2007, by default, is C:\Program Files\Microsoft\Exchange Server.
- ☑ When a database corruption occurs, 99.9 percent of the time it's caused by the underlying hard disk drive subsystem.

Restoring Mailbox Data Using the Recovery Storage Group Feature

- ☑ The Recovery Storage Group (RSG) feature, which was originally introduced back in Exchange 2003, gives you (the Exchange administrator) the option of mounting a second copy of a mailbox database (typically a mailbox database restored from backup). This way, you can extract data during work hours from one

or more mailboxes in the respective database without affecting the production databases.

- ☑ With Exchange 2007, the RSG feature is accessed using the new Database Recovery Management tool, which is found under the Exchange Toolbox work center. You can also work with RSGs using the Exchange Management Shell (EMS).
- ☑ When you have merged or copied the required Mailbox data, you can use ExTRA to dismount and then remove the Recovery Storage Group. Be sure you delete the files in the RSGxxxxxxx folder again after you have removed it so the files don't take up valuable disk space.

Recovering an Exchange 2007 Server Using the RecoverServer Switch

- ☑ Just as with Exchange 2000 and 2003, you can recover an Exchange 2007 Server in a fairly straightforward way. As you perhaps know, we could use the DisasterRecovery switch to recover a dead Exchange 2000 or 2003 Server on new hardware, but with Exchange 2007 this switch no longer exists. Instead, it has been replaced by the new RecoverServer switch, which is similar to the DisasterRecovery switch.
- ☑ The RecoverServer switch can be used to recover all types of Exchange 2007 Server roles except for the Edge Transport Server role, which uses ADAM and not the Active Directory to store configuration data.
- ☑ If you're recovering an Exchange 2007 Server with the Hub Transport Server role installed, and this is the only Exchange 2007 Server with this role installed, it's recommended you run `Setup.com /M:RecoverServer` with the `/DoNotStartTransport` syntax since there's a few post-recovery steps that should be completed before this role is made active.
- ☑ When you run Setup with the RecoverServer switch on a new Windows 2003 Server that is configured with the same name as the one that has crashed or is permanently down for some reason, Setup will read the configuration information for the respective Exchange 2007 server from the Active Directory. In addition to applying the roles and settings stored in Active Directory, Setup will (just as when installing an Exchange 2007 Server role without the RecoverServer switch) install the Exchange files and services required for the respective Exchange 2007 server role(s).

Recovering an Exchange 2007 Cluster Using the RecoverCMS Switch

- ☑ You can recover an Exchange 2007 clustered mailbox server (both CCR and SCC) by using the *ExSetup.exe* command with the RecoverCMS switch.
- ☑ If you're recovering a Single Copy Cluster (SCC) and have stored the Mailbox and Public Folder databases on a storage area network (SAN), you won't need to restore the databases from backup as long as each node points to the same shared storage subsystem as the failed clustered mailbox server did.

Recovering Mailbox Databases Using the Improved Database Portability Feature

- ☑ With Exchange 2007, the Mailbox database portability feature has been improved drastically. Now you can port and recover a Mailbox database to any server in the Exchange 2007 organization, and because of the new Autodiscover service (which we discussed in Chapter 5), all Outlook 2007 clients will be redirected to the new server automatically the first time they try to connect after the Mailbox database has been mounted on another Exchange 2007 server.
- ☑ It's important that the name of the new Mailbox database you create using the *New-MailboxDatabase* cmdlet matches the name of the Mailbox database you ported or recovered from the old Exchange 2007 Server. Otherwise, you won't be able to mount it.

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to www.syngress.com/solutions and click on the "Ask the Author" form.

- Q:** Now that we have Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR), should you still take regular backups of the Exchange 2007 databases using a backup application?
- A:** It's important to understand that LCR and CCR aren't replacements for traditional regular backups. Instead, they are meant to serve as the primary fast recovery solution in case one or more of your production databases shuts down. But with LCR or CCR, you can change your backup schedule from daily to weekly backups.
- Q:** I heard you can take backups of the passive databases when using LCR or CCR, but I don't have the option of choosing the passive database in Windows 2003 Backup?
- A:** You're right in that LCR or CCR gives you the option of performing the backup of the passive database(s), but although the Windows 2003 Backup tool supports Volume Shadow Copy Service (VSS) backups, this is only the case when performing file-level-based backups of the databases. In order to perform a backup of the passive databases, you must use a third-party backup solution that supports VSS backups or Microsoft's Data Protection Manager version 2 (DPM v2), which at the time of this writing is still a beta product.
- Q:** How do you create and manage a Recovery Storage Group (RSG) in the Exchange 2007 Management Console?
- A:** You don't. With Exchange 2007, the RSG feature cannot be managed using the Exchange Management Console, as was the case in Exchange 2003. Instead, you must create and manage RSGs using the Database Recovery Management tool (which can be found beneath the Toolbox work center node) or the Exchange Management Shell.
- Q:** Is it possible to restore a Public Folder database to a Recovery Storage Group (RSG) in Exchange 2007?
- A:** No. Unfortunately, the RSG feature is still limited to Mailbox databases only.
- Q:** Can I recover all types of Exchange 2007 Server roles using the new RecoverServer switch?
- A:** Yes, almost. The only Exchange 2007 Server role that cannot be recovered using the RecoverServer switch is the Edge Transport server since this server doesn't belong to the Active Directory. To recover an Edge Transport server, you must instead use the cloned configuration method, which you can read more about in Chapter 7.

