

High Availability for Exchange 2007 Mailbox Servers

Solutions in this chapter:

- Managing the Local Continuous Replication Feature
- Managing a Cluster Continuous Replication-Based Setup
- Managing a Single Copy Cluster-Based Setup

- Summary
- Solutions Fast Track
- Frequently Asked Questions

Introduction

The availability requirements for messaging and collaboration servers have increased drastically over the years, with the result that these servers are now among the most mission-critical servers in the datacenter. Several recent reports have concluded that e-mail is more important to end users than their phones. So it's not rocket science; it's in the interests of you as the Exchange Administrator to achieve as high an uptime as possible. Each of these facts played an important role when the Exchange Product Group developed Exchange Server 2007, so it's no surprise that when speaking of high availability as well as disaster recovery, we can find many improvements as well as new functionality in the Exchange Server 2007 product.

Exchange Server 2007 includes three primary high-availability solutions relating to the Mailbox Server role, although one of these features isn't really new at all but has instead changed name and been further improved since Exchange Server 2003. We're referring to the Single Copy Cluster (SCC) functionality, which is a clustered solution that uses a single copy of a storage group on storage that is shared between the nodes in a cluster. Those of you with just a little bit of Exchange cluster experience would say that the SCC solution is similar to a traditional Exchange 2000/2003 active/passive cluster setup, and you're right.

SOME INDEPENDENT ADVICE

With Exchange Server 2007, active/active clusters are no longer supported; only active/passive clusters are supported. If you have experience deploying Exchange 2000/2003 in an active/active cluster, most likely you understand why this was dropped in Exchange 2007. An Exchange cluster configured with two active nodes has never performed as well as one would have expected, since the failover causes the remaining node to take on additional processing operations. Constraints such as number of concurrent user connections per node and average CPU load per server limits also play an important role in the reason that active/active Exchange cluster setups have never been successful.

Then we have Local Continuous Replication (LCR), which is a solution that uses the new continuous replication technology introduced in Exchange 2007. LCR is a new functionality that uses built-in asynchronous log shipping and log replay technology to create and maintain a replica of a storage group on a second set of disks that are connected to the same server as the production storage group. As mentioned, the LCR solution uses log shipping and log replay and gives you the option of switching to the passive copy of the storage group in a matter of minutes, should the database in the active storage group become corrupted and shut down for one reason or another. The interesting thing about LCR is that this solution doesn't require more than a single Exchange 2007 server with the Mailbox Server role installed.

Finally, we have the Clustered Continuous Replication (CCR) solution, which, like LCR, uses the new Exchange 2007 continuous replication technology, but as the name implies, CCR is a clustered solution that eliminates the single point of failure that exists in traditional Exchange cluster setups today. This is done by maintaining a copy of the database on the active node; in the event of a database corruption, this allows both services and databases to fail over to the passive node. CCR can only be deployed in a two-node active/passive cluster.

Managing the Local Continuous Replication Feature

In this first section of the chapter we'll take a closer look at the architecture behind the new Local Continuous Replication (LCR) feature. We'll then go through the steps necessary to enable this feature; finally, we'll look at how we can take advantage of LCR should the database in the active copy of the storage group fail.

Local Continuous Replication under the Hood

The Exchange Product group developed the Local Continuous Replication (LCR) technology to provide a native data availability solution that can be used to recover an Exchange database on an Exchange 2007 standalone server in a matter of a few minutes. In Exchange 2003 as well as previous versions, you needed to recover the lost database by restoring it from backup, which, depending on the database size, could take up to many hours. With LCR, you will be able to switch over to an exact replica (that is, a fully updated copy) of the crashed database by running a simple Exchange 2007 task.

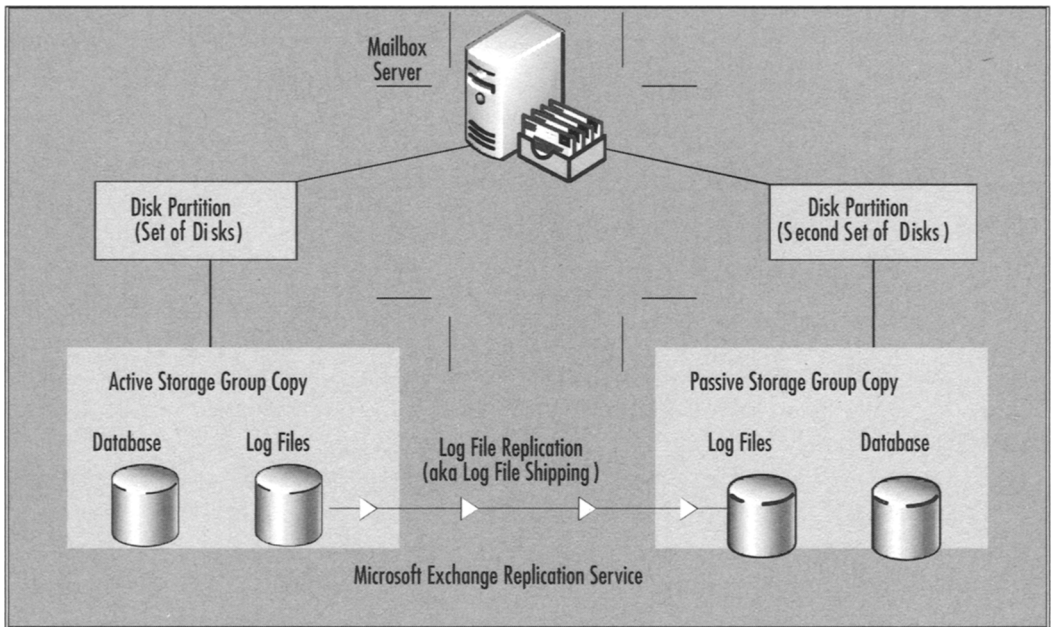
So how does this LCR magic work? As most of us know, the database type Exchange uses is Extensible Storage Engine (ESE). ESE employs transaction log files, which means that every time a modification is made, a transaction log file is generated (instead of the change being committed directly to the database). The reason is that when the ESE database is modified, the modification won't be made directly in the physical database but instead in memory of the respective Exchange 2007 Mailbox Server. This means that should the database for some reason become corrupted or shut down, Exchange always will be able to recover the lost data (which is held in memory, remember) by using the log files.

Each log file that is generated because of a modification in the database belonging to the active copy of the storage group is replicated (copied) from the source log folder (the log folder defined for the Storage Group containing the respective database) to a target log folder associated with the passive copy of the storage group. This isn't the entire truth, because each log file is first copied to an inspector log folder located beneath the target log folder, where it is inspected to make sure it is correct. (If it isn't correct, the log file will be recopied). Finally the file is copied to the target log folder and from there replayed into the database belonging to the passive copy of the storage group.

The target log folder also contains an IgnoredLogs folder that holds any valid log files that for some reason cannot be replayed. A typical reason is that the particular log is too old. In addition, the subfolder can contain an InspectionFailed and an E00OutOfDate folder. The first is a folder that holds any log files that failed inspection. When this happens, an event 2013 will be logged in the application log. The E00OutOfDate folder will hold any E00.log files that are present in the target log folder when a failover occurs.

A new Exchange 2007 service called the Microsoft Exchange Replication Service will be installed on any Exchange 2007 servers with the Mailbox Server role installed. These are responsible for replicating the log files to the target log folder. As you can see, we've tried to illustrate the basic architecture of LCR in Figure 8.1.

Figure 8.1 The Basic Local Continuous Replication Architecture



The log files that are replicated from the active copy to the passive copy of the storage group will be replayed in batches in order to provide the best performance possible.

Since LCR keeps an exact replica of the active copy of the storage group, the number of Exchange backups needed is also reduced drastically. But it's important to understand that LCR in no way eliminates traditional backups of the databases on your Exchange 2007 Mailbox servers; instead, it provides you with the option of taking weekly instead of daily backups, for example.

SOME INDEPENDENT ADVICE

Bear in mind that if you want to enable LCR for a storage group, the storage group may not contain more than one mailbox or public folder database. This is because LCR doesn't support multiple databases in the same storage group. Actually, you won't even be able to enable LCR on a storage group containing multiple databases. In addition, you cannot enable LCR for a storage group containing a Public Folder database if more than one Public Folder database exists in the organization. The reason is that LCR and Public Folder replication cannot run at the same time.

When you're partitioning the disks that should be storing the passive copies your storage groups, it is best practice to take advantage of mount points, because they will let you surpass the 26-drive-letter limitation that exists on a Windows 2003 server. If you end up in a situation where you need to switch to a passive copy of a storage group, using mount points will make the recovery process much more painless because you can quickly change drive letters and paths.

As has been the case with mailbox stores and log files in previous versions of Exchange, it's also recommended that you place the databases and log files for a passive copy of a storage group on separate disks, just as you do with active copies of storage groups.

You should, of course, also make sure that you partition the disks that are to be used for the passive copies of the storage groups, so they are at least the same size as the disks holding the active storage group copies. Finally, keep in mind that a Mailbox Server with LCR enabled will use approximately 30–40 percent more CPU and memory than a Mailbox Server on which LCR hasn't been enabled. These extra resources are primarily used by log file verification as well as log file replay.



TIP

LCR enables you to offload Volume ShadowCopy Service (VSS) backups from the active storage group to the passive storage group, which will preserve disk I/O on the disks on which the active storage group is located. This also means that you can perform restores from a passive copy of a storage group.

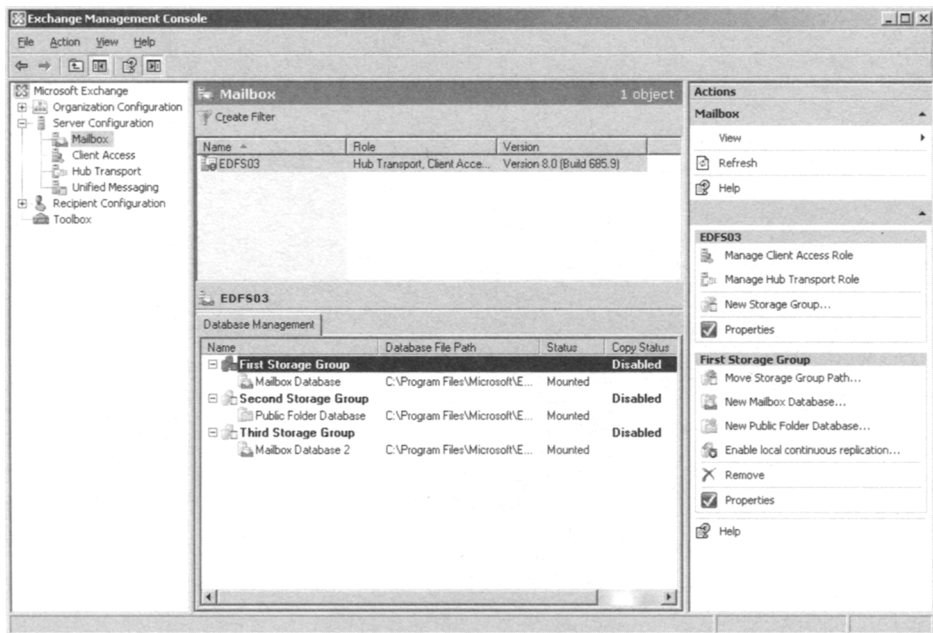
As you can understand, LCR is an ideal solution for small or medium-sized organizations because the functionality allows rapid recovery from database issues and requires only an extra set of disks for the database copies. LCR increases the availability of databases on an Exchange 2007 standalone server in an affordable way. For small shops that don't have a big fancy server with multiple sets of disks, it is possible to keep the LCR copy on an external USB disk.

Enabling Local Continuous Replication on a Storage Group

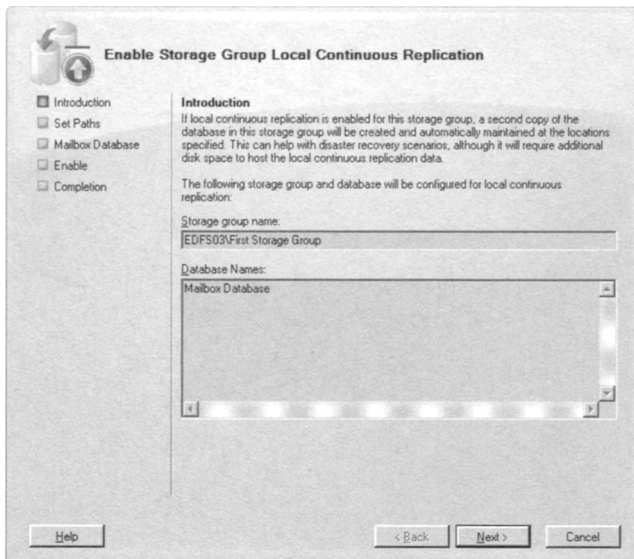
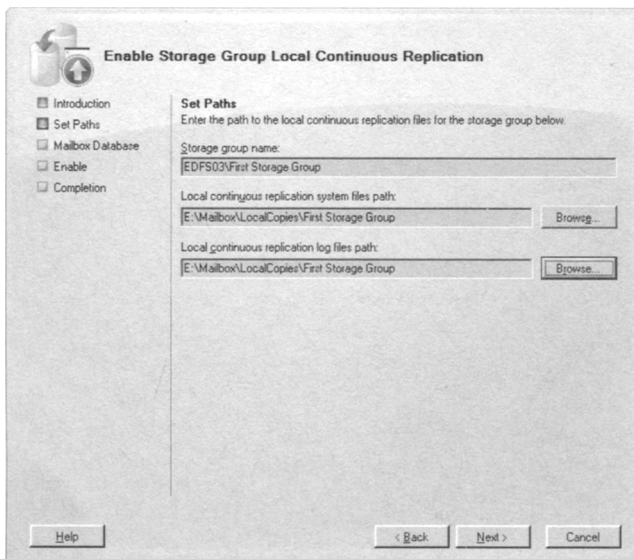
The LCR feature is enabled on a Storage Group level under the Mailbox subnode, located beneath the Server Configuration work center node in the left pane of the Exchange System Management Console, as shown in Figure 8.2.

1. To enable LCR for the First Storage Group, select it in the work pane, and click **Enable local continuous replication** in the Action pane.

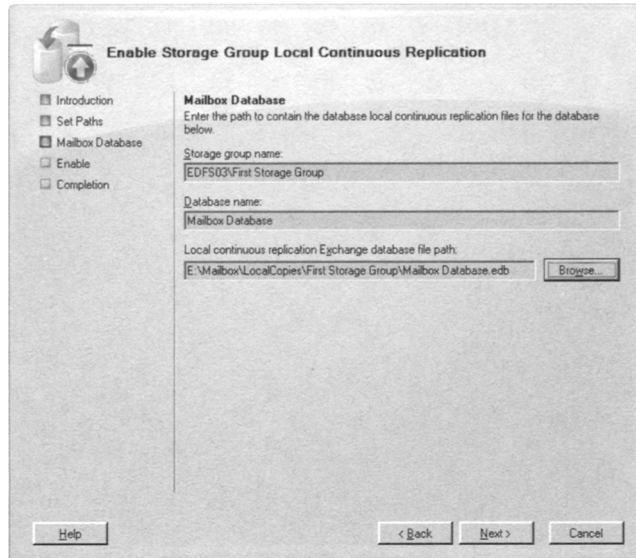
Figure 8.2 The Local Continuous Replication Link in the Action Pane



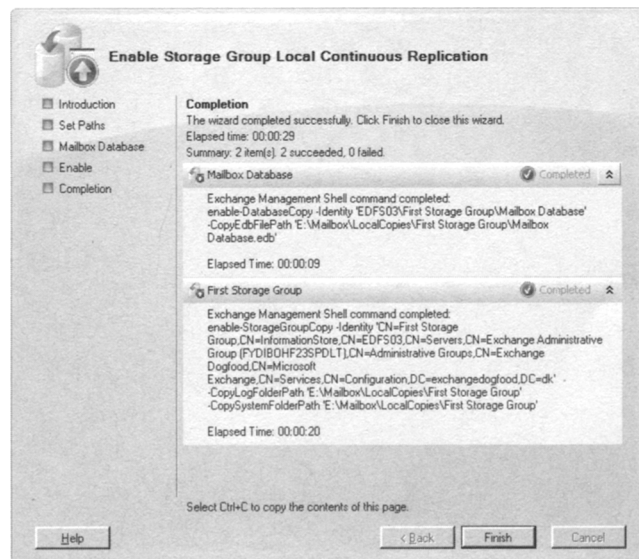
2. This will bring up the Local Continuous Replication Wizard's Introduction page, shown in Figure 8.3. As you can see, this page shows us the storage group as well as mailbox database name. Because there aren't many interactions on this page, simply click **Next**.
3. Now let's specify the path to the LCR files for the respective storage group (see Figure 8.4). For the purpose of this example, we're simply specifying the E: drive, which is a second set of disks on the server. When the location has been specified, we can click **Next**.

Figure 8.3 Enable Storage Group Local Continuous Replication**Figure 8.4** Specifying the Paths for the Replicated Log and System Files

4. On the Mailbox Database page, we have to specify the path to the location of the second copy of the database, as shown in Figure 8.5. When you have done so, click **Next**.

Figure 8.5 Specifying the Path for the Database Copy

5. We have now reached the step where we enable LCR for the storage group, so let's do so by clicking **Enable** and see what happens. As shown in Figure 8.6, the Local Continuous Replication Wizard completed successfully. Click **Finish**.

Figure 8.6 The Local Continuous Replication Feature Was Enabled with Success

If you would rather enable LCR for a storage group via the EMS, you will have to do so using the *Enable-DatabaseCopy* and *Enable-StorageGroupCopy* cmdlets. To enable LCR for the First Storage Group, you would need to first run the following command:

```
Enable-DatabaseCopy -Identity "EDFS03\First Storage Group\Mailbox Database" -
CopyEDBFilePath:"E:\Mailbox\LocalCopies\First Storage Group\Mailbox Database.edb"
```

Then type:

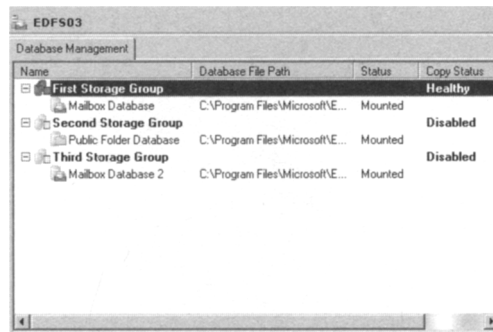
```
Enable-StorageGroupCopy -Identity "EDFS03\First Storage Group"
-CopyLogFolderPath:"E:\Mailbox\LocalCopies\First Storage Group"
-CopySystemFolderPath:"E:\Mailbox\LocalCopies\First Storage Group"
```

NOTE

Even though we're dealing with a secondary copy of a production database, it's still a best practice to keep the log files and database separated on their own set of disks.

Now notice that the copy status for the First Storage Group has change to Healthy (see Figure 8.7).

Figure 8.7 The Copy Status for the Storage Group Is Healthy



Name	Database File Path	Status	Copy Status
First Storage Group			Healthy
Mailbox Database	C:\Program Files\MicrosoftE...	Mounted	
Second Storage Group			Disabled
Public Folder Database	C:\Program Files\MicrosoftE...	Mounted	
Third Storage Group			Disabled
Mailbox Database 2	C:\Program Files\MicrosoftE...	Mounted	

Viewing the Status for a Local Continuous Replication Copy

To view basic health and status information for an LCR copy, you can bring up the Properties page for the storage group on which LCR has been enabled. To do this, select the respective storage group and click the **Properties** link in the Action pane. On the Properties page, select the **Local continuous replication** tab, as shown in Figure 8.8. Here you can see the basic health for an LCR copy.

Figure 8.8 The LCR Status Properties Page

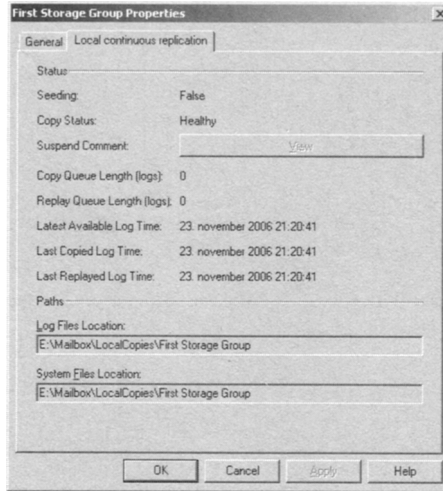


Table 8.1 lists the health and status information, with a short description of each.

Table 8.1 Local Continuous Replication Health and Status Information

Health/Status Information	Description
Seeding	Used to indicate whether seeding of the passive database occurs or not. Can have a status of True or False.
Copy Status	Used to indicate whether log file copying has started. Can have a status of Healthy, Suspended, or Broken.
Suspend Comment	Can be used to view <i>suspend</i> comment if LCR has been suspended.
Copy Queue Length (logs)	Used to display the number of log files that are waiting to be copied to the passive storage group's log file folder. Note that a copy is not considered complete until it has been inspected for corruption.
Replay Queue Length (logs)	Used to display the number of log files waiting to be replayed into the passive storage group's database.
Latest Available Log Time	Used to display the time stamp on the active storage group of the most recently detected new log file.
Last Copied Log Time	Used to display the time stamp on the active storage group of the last successful copy of a transaction log file.
Last Replayed Log Time	Used to display the time stamp on the passive storage group of the last successful replay of a log file.

In addition, you can see the path to the log file and system file location for the passive storage group copy.

If you want even more information about the health and status of an LCR copy, you can open the EMS and type **Get-StorageGroupCopyStatus -Identity <Storage Group> | FL**, as shown in Figure 8.9.

Figure 8.9 Retrieving LCR Status Information via the Exchange Management Shell

```

Machine: EDF503 | Scope: exchangedogfood.dk
[PS] C:\>Get-StorageGroupCopyStatus -Identity "First Storage Group" | FL

Identity                : EDF503\First Storage Group
StorageGroupName        : First Storage Group
SummaryCopyStatus       : Healthy
CCRTargetNode           :
Failed                  : False
FailedMessage           :
Seeding                 : False
Suspend                 : False
SuspendComment          :
CopyQueueLength         : 0
ReplayQueueLength       : 0
LatestAvailableLogTime  : 25-11-2006 20:04:56
LastCopyNotifiedLogTime : 25-11-2006 20:04:56
LastCopiedLogTime       : 25-11-2006 20:04:56
LastInspectedLogTime    : 25-11-2006 20:04:56
LastReplayedLogTime     : 25-11-2006 20:04:56
LastLogGenerated        : 50
LastLogCopyNotified     : 50
LastLogCopied           : 50
LastLogInspected        : 50
LastLogReplayed         : 50
LatestFullBackupTime    : 24-11-2006 14:29:20
LatestIncrementalBackupTime :
SnapshotBackup          : False
IsValid                 : True
ObjectState              : Unchanged

[PS] C:\>_
  
```

Going through each information field returned by the *Get-StorageGroupCopyStatus* CMDlet is outside the scope of this book, so if you want to dig deeper into these topics, we recommend that you refer to the Exchange 2007 Help file.

Switching to the Passive Storage Group Copy When Disaster Strikes


When disaster strikes and the database or log files in the active copy of the storage group have become corrupted and have shut down, you have the option to recover database availability by switching to the LCR copy (the passive copy of the storage group).

You can recover from corruption of either one or more log files or the database using a variety of methods, depending on whether you use mount points or not. One method is to run the *Restore-StorageGroupCopy* CMDlet with the *ReplaceLocations* parameter, which will activate the LCR copy as the active storage group copy in one step. To activate the LCR copy as the active storage group, you first need to make sure that the active database is dismounted, which should already be the case if it's corrupted. If this is not the case, you should dismount it now. When you have done so, we're ready to run the *Restore-StorageGroupCopy* CMDlet, which in the case of this example is done for the First Storage Group. So the command to run in the EMS is:

```
Restore-StorageGroupCopy -Identity "First Storage Group" -ReplaceLocations:$true
```

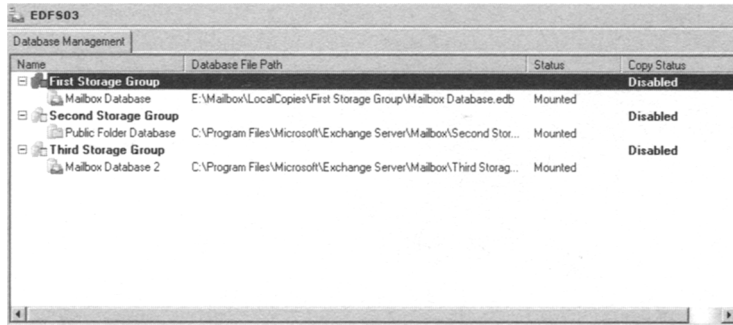
An integrity check will now be passed for the log files, and if it's completed without errors, the storage group copy switch will be completed and the production paths will be updated, as shown in Figure 8.10.

Figure 8.10 Switching to the LCR Copy Using the *Restore-StorageGroupCopy* CMDlet



All there is to do now is to mount the database using either the EMC or the EMS. Now notice that the Database File Path will have changed, as shown in Figure 8.11.

Figure 8.11 Database File Path Change



NOTE

When you have run the *Restore-StorageGroupCopy* CMDlet against a storage group, LCR for the respective storage group will be disabled. So remember to re-enable LCR for the particular storage group after you perform a switch to the LCR copy.

Although this method is straightforward and fully supported, Microsoft actually recommends that instead you use a method whereby you run the *Restore-StorageGroupCopy*

CMDlet without the *ReplaceLocations* parameter, to activate the copy in its current location, and then either move the files manually, change drive letters, or use mount point assignments to have the copy files reflected under the respective production paths so that the production database is maintained in the expected location. Following this method means that the active storage group copy will continue to have meaningful filenames that represent that they indeed are active production copies. Why is this the preferred method? Because Microsoft believes that using the *Restore-StorageGroupCopy* CMDlet with the *ReplaceLocations* parameter could lead to future confusion in distinguishing the active copy of the data from the passive copy of the data, and to be honest, we agree. That said, we cannot see why you shouldn't use the *ReplaceLocations* parameter if you know what you're doing; just make sure that you switch back to the original disk set again.

Let's examine an example of how you would use the recommend method. First, make sure that the production database is dismounted. Then open the EMS and type **Restore-StorageGroupCopy -Identity "First Storage Group"**.

This command will activate the copy and leave the path for the production storage group intact. Now you can choose between either moving the LCR copy files to the location of the original production database manually using Windows Explorer or using Xcopy or a similar tool. Just be sure to move or delete the files in the folder you move the files to first. When the files have been moved, you simply need to mount the database again, and that's it.

The second option available when using the *Restore-StorageGroupCopy* CMDlet without the *ReplaceLocations* parameter is to change the drive letter for the partition holding the LCR copy to the drive letter used by the production storage group. This can be done using either the Disk Management MMC snap-in or the Diskpart tool.

1. To do so using the MMC snap-in, click **Start | Run** and type **Diskmgmt.msc**. This will bring up the MMC snap-in shown in Figure 8.12. Now right-click the partition holding the production storage group and its database, then select **Change drive letter and paths** in the context menu.
2. In the Change Drive Letter and Paths For window, click **Change**, then specify an unallocated drive letter and click **OK**, as shown in Figure 8.13.

Figure 8.12 The Disk Management MMC Snap-in

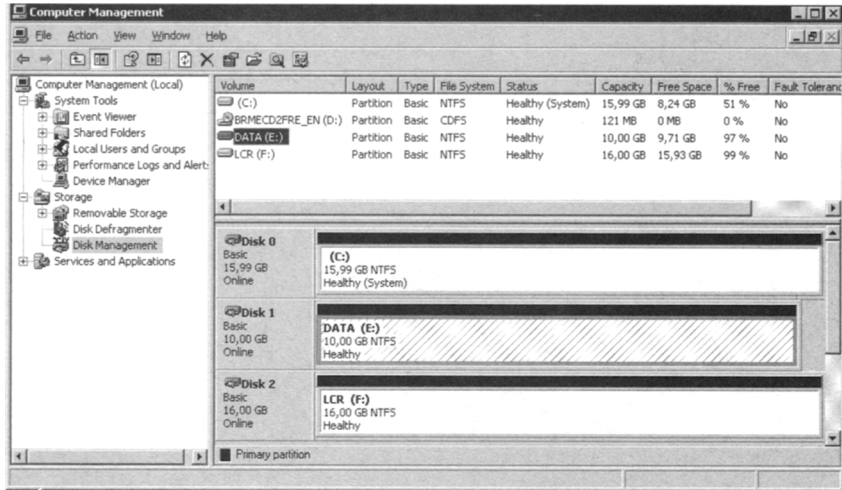
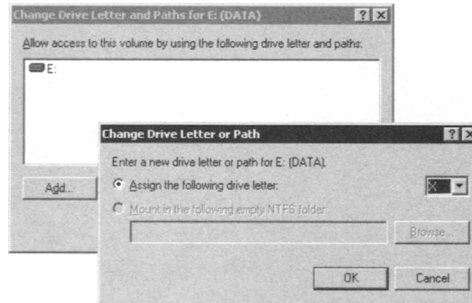


Figure 8.13 Specifying the Drive Letter for the Partition



3. Click **OK** to the confirmation message and click **OK** to close the Change Drive Letter and Paths window.
4. Now change the drive letter for the partition holding your LCR copy to the drive letter that originally was assigned the partition that holds the production storage group, which in this example is **E:**.

It's important that the partition for which you change the drive letter for doesn't contain any other data used by other applications. If it does, you will most likely destroy functionality for the respective applications!

When you have changed the drive letter, all there is to do is to mount the database again, but remember, the paths for the active and passive storage groups must be the same on each partition.

NOTE

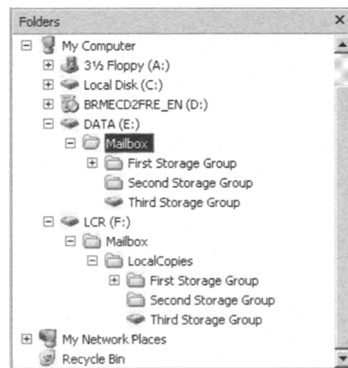
A restart of the server might be required for you to be able to assign the E: drive to the partition holding the LCR copy.

The last option available involves the use of mount points. A *mount point* is a feature with which you can surpass the 26-drive-letter limitation that exists in Windows 2003 Server. Using volume mount points, you can graft, or mount, a target partition into a folder on another physical disk. Since volume mount points are transparent to Exchange 2007 as well as most other programs, they are pretty popular, especially in deploying Exchange 2000/2003 cluster environments.

To use mount points to switch LCR storage group copies, you must already have configured the partitions holding the storage group copies to use them. If you haven't done so, the mount point option cannot be used. In this example, the Third Storage Group's folder as well as the LCR copy for this storage group, which is called Third Storage Group, point to an NTFS volume mount point.

You can see whether a particular folder in Windows Explorer is a mount point because the icon is represented as a disk and not the normal yellow folder icon (see Figure 8.14).

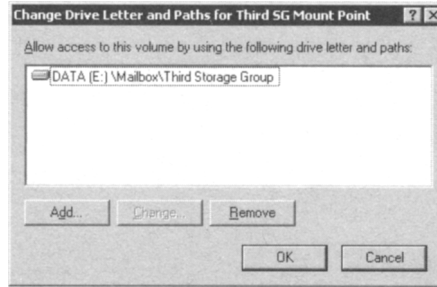
Figure 8.14 The Mount Point Icon in Windows Explorer



1. As is the case with the options we have covered, the first thing you should do before switching the storage group copies using NTFS volume mount points is to make sure that the database is in a dismounted state. If this is not the case, you should dismount it manually now. The next step is to open the EMS and type **Restore-StorageGroupCopy -Identity "Third Storage Group"** (which is the storage group used in this example).
2. Next open the Disk Management MMC snap-in, right-click the partition that is used as the NTFS volume mount point by the production storage group, then

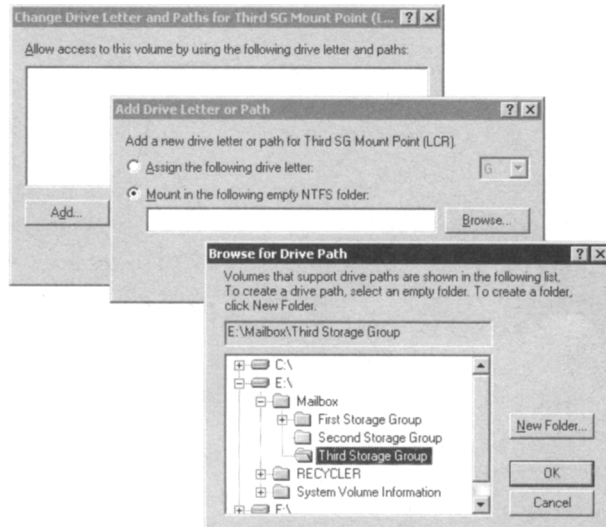
select **Change Drive Letter or Paths** in the context menu. In the Change Drive Letter and Paths window, remove the existing path by highlighting it, then click the **Remove** button (see Figure 8.15).

Figure 8.15 Changing the NTFS Volume Mount Point Path



3. You now need to confirm that you want to remove the path. Click **Yes**.
4. Now remove the mount point for the partition used for the LCR copy, using the same steps. This is required to be able to use the LCR copy path as a mount point for the Production Storage Group copy.
5. We're now ready to mount the LCR copy to the Production Storage Group. We do so by right-clicking the partition that was used for the LCR copy, then choosing **Change Drive Letter or Paths** in the context menu. Now click **Add** and select **Mount** in the following empty NTFS folder. Click **Browse** and specify the path to the production storage group (see Figure 8.16). Finally, click **OK** twice and close the Disk Management MMC snap-in.

Figure 8.16 Specifying the New Path for the NTFS Volume Mount Point



6. Now verify that the folder within Windows Explorer contains the expected data, and then mount the database again.

Is that cool or what?

Suspending Local Continuous Replication

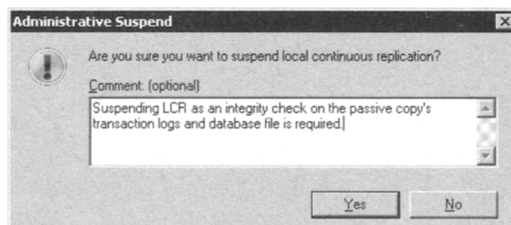
On occasion, you might need to suspend LCR for a storage group. You need to suspend LCR should either the active or passive storage group copy for some reason become unavailable. Suspending LCR is also necessary if you need to seed the LCR copy (seeding is covered next in this chapter). Finally, you need to suspend LCR when you're performing an integrity check on the passive copy's transaction logs and database file, which is a recommended practice now and then.

NOTE

Suspending LCR means that all log file shipping as well as log file replaying is halted.

Suspending LCR is a straightforward process; it's done by selecting the respective storage group in the EMC, then clicking **Suspend Local continuous replication** in the Action pane. When you click this link, you'll need to confirm that you really want to suspend LCR. In addition, you'll have the option of specifying why LCR was suspended. This comment can be viewed by clicking the **View Comment** button on the Properties page of the storage group (shown in Figure 8.17).

Figure 8.17 Suspending Local Continuous Replication



If you'd rather to suspend LCR for a storage group via the EMS, you'll need to do so using the `Suspend-StorageGroupCopy` CMDlet. To suspend LCR for the First Storage Group, where the comment shown in Figure 8.17 is specified, you should run the following command:


```
Suspend-StorageGroupCopy -Identity "First Storage Group" -SuspendComment
"Suspending LCR as an integrity check on the passive copy's transaction logs and
database file is required."
```

Again, you need to confirm that you really want to suspend LCR for the storage group. To do so, type **Y** for Yes and press **Enter**.

Resuming Local Continuous Replication

When the active or passive storage group is available again or when you have performed the integrity check or whatever type of maintenance you have completed, you need to resume LCR for the storage group. Again, this can be done via either the EMC or the EMS. To perform this task using the EMC, select the respective storage group and click **Resume local continuous replication** in the Action pane. When you do, the warning message shown in Figure 8.18 will appear. Click **Yes** and watch the Copy Status change to **Healthy** once again. Both log file shipping and log file replay have now been resumed.

Figure 8.18 Resuming Local Continuous Replication



To resume LCR for a storage group via the EMS, type **Resume-StorageGroupCopy -Identity "First Storage Group"**.

Manually Seeding a Database Copy

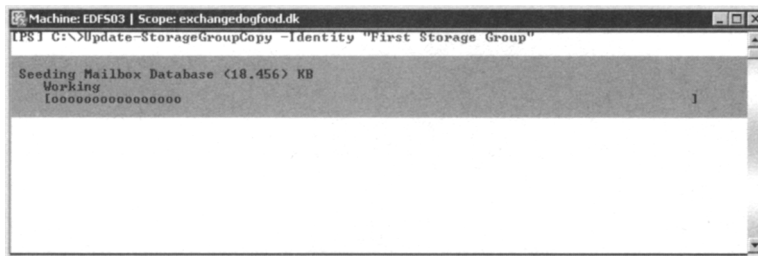
Before we start talking about how to perform a manual seeding of a database copy, it would be a good idea to define the term *seeding* in terms of LCR. Seeding is the process whereby a database is added to a storage group copy. This can be a blank database or a copy of the database the storage group uses as the production database. When you enable LCR on a storage group using the EMC or via the EMS using the *Enable-DatabaseCopy* and *Enable-StorageGroupCopy* CMDlets, seeding normally takes place automatically. If it happens automatically, why should we even care about it, then? The answer is that there are a few situations in which manually seeding is required. The first is after you have performed an offline defragmentation of the production database belonging to the storage group for which you have enabled LCR. The second is if or when Exchange detects a corrupt log file, which the Microsoft Exchange Replication Service cannot replay into the database copy. The third is after a page scrubbing of a database on the active node in a Cluster Continuous Replication (CCR) setup occurs, and you then want to propagate these

changes to the passive node in the CCR setup. Yes, you're right, the last one isn't really related to LCR but only continuous replication in clustered environments, where CCR is used. We'll talk much more about CCR later in this chapter.

Seeding a database copy manually can be done using the *Update-StorageGroupCopy* CMDlet in the EMS. Before doing so, you must suspend LCR for the respective storage group and then remove any .log, .chk, .jrs, and .edb files from the passive storage group's database copy, log files, and system files paths. To seed the database copy for the First Storage Group, you use the *Update-StorageGroupCopy* CMDlet and type **Update-StorageGroupCopy -Identity: "First Storage Group"**.

Running this command will create a temporary temp-seeding folder, and after a little while the seeding will take place, as shown in Figure 8.19.

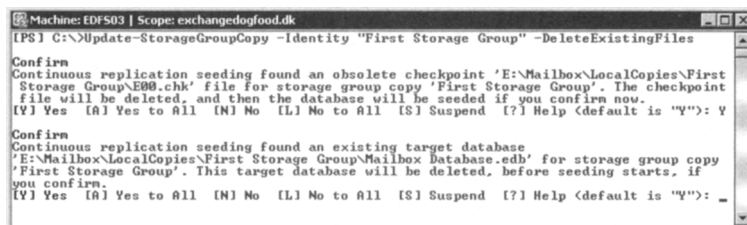
Figure 8.19 Seeding a Mailbox Database Copy



When seeding has taken place, the Microsoft Exchange Replication Service will start to replicate any .log, .chk, and .jrs files to the folder paths. When it's finished, you can resume LCR for the storage group, and you're back in business.

If you don't want to delete any .log, .chk, .jrs, and .edb files manually before running the *Update-StorageGroupCopy* CMDlet, you can tell the CMDlet to do it for you using the *DeleteExistingFiles* parameter. This method requires that you confirm the deletion of these files, as shown in Figure 8.20. The method you use is up to you, since they do the same thing.

Figure 8.20 Specifying That the *StorageGroupCopy* CMDlet Delete Any Existing Files



In addition, you can use the *ManualResume* parameter if you don't want replication to occur automatically on the storage group copy.

Another method available for seeding a database copy is to dismount the database in the EMC, suspend LCR for the storage group containing the database, and then copy the .edb file to the LCR copy folder using Windows Explorer. When the file has been copied, you then mount the database again using the EMC and resume LCR. Bear in mind that if you choose this method, your end users will be disconnected until the database is mounted. So unless there's a specific reason that you would use this method, we recommend that you use the *StorageGroupCopy* CMDlet.

Performing an Integrity Check of the Passive Copy Using Eseutil

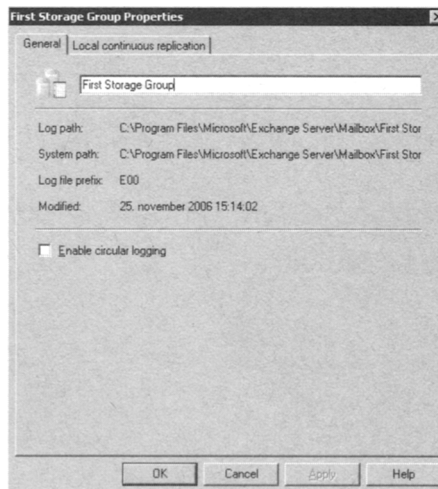
It's a recommended best practice to periodically verify the integrity of the passive storage group copy to make sure neither the database copy nor any of the log files are corrupted. This is done by running a physical consistency check against both the database copy as well as the log files using Exchange Server Database Utilities (Eseutil.exe).

As mentioned earlier in this chapter, you need to suspend LCR on the storage group for which you want to verify the integrity of the passive database and log files.

To verify the physical integrity of the log files that have been replicated to the passive copy of the storage group, you'll need to open either a Command Prompt window or the EMS. In either the Command Prompt window or the EMS you should run Eseutil with the */k* switch followed by the log file prefix of the storage group.

The log file prefix for a storage group can be found under the General tab of the respective storage group, as shown in Figure 8.21.

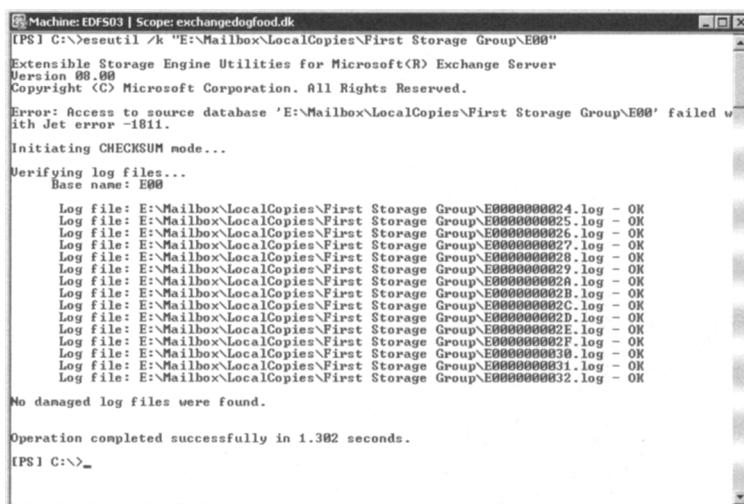
Figure 8.21 Log File Prefix



As you can see, the log file prefix for the First Storage Group typically is E00. To see the path for the log files, refer back to Figure 8.8. For the purpose of this example, the path is E:\Mailbox\LocalCopies\First Storage Group, so we'll need to type **Eseutil /k "E:\Mailbox\LocalCopies\First Storage Group\E00"**.

This will initiate checksum mode and start verifying each log file located under the specified path, as shown in Figure 8.22. If no corrupted log files are detected, the operation will complete successfully after a few seconds or minutes, depending on how many log files are contained in the respective folder.

Figure 8.22 Integrity Check of the LCR Log Files



```
[PS] C:\>eseutil /k "E:\Mailbox\LocalCopies\First Storage Group\E00"
Extensible Storage Engine Utilities for Microsoft(R) Exchange Server
Version 08.00
Copyright (C) Microsoft Corporation. All Rights Reserved.

Error: Access to source database 'E:\Mailbox\LocalCopies\First Storage Group\E00' failed with Jet error -1811.

Initiating CHECKSUM mode...

Verifying log files...
Base name: E00
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000024.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000025.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000026.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000027.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000028.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000029.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002A.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002B.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002C.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002D.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002E.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E000000002F.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000030.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000031.log - OK
Log file: E:\Mailbox\LocalCopies\First Storage Group\E0000000032.log - OK

No damaged log files were found.

Operation completed successfully in 1.302 seconds.

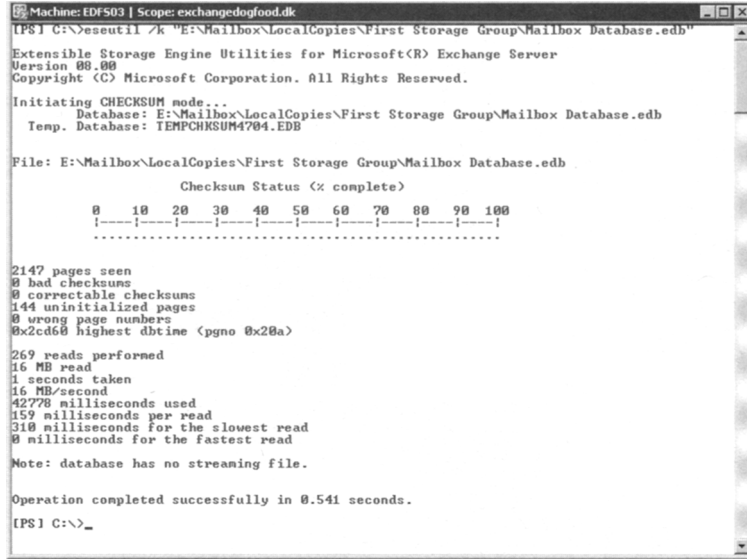
[PS] C:\>_
```

When the log files have been verified, we can move on to checking the integrity of the database copy. This is also done by running Eseutil with the /k switch but instead followed by the full path the database copy. In this example, we need to run the following command: **Eseutil /k "E:\Mailbox\LocalCopies\First Storage Group\Mailbox Database.edb"**.

Eseutil will once again initiate checksum mode and then create a temporary database so that the database copy can be checked for any errors (see Figure 8.23). Again, the time required for the integrity check depends on the size of the database.

When you have performed an integrity check of both the log files and the database copy (and hopefully Eseutil.exe hasn't found too many corrupted log files or issues with the database copy), you should make sure that LCR for the respective storage group is resumed again. Should you be so unlucky that Eseutil.exe finds one or more corrupted log files or corruption in the database copy, you need to disable LCR on the storage group, then remove the corrupted log files and/or database copy file. When the files have been removed, you can re-enable LCR, which will create a database copy and seed it as well as replicate any existing log files from the active copy of the storage group to the specified path.

Figure 8.23 Integrity Check of the LCR Database Copy




We’ll bet that most of you understand the importance of during periodically integrity checks of both the log files as well as the database copy—right?

Disabling Local Continuous Replication on a Storage Group

There might come a time when you no longer want to have the LCR feature enabled for a particular storage group. Luckily, it’s a painless process to disable this feature once it’s enabled.

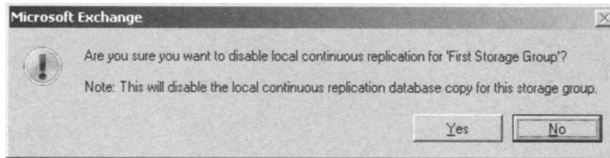
You can disable LCR for a storage group via either the EMC or the EMS. To disable LCR using the EMC, you need to select the Storage Group level under the Mailbox subnode, located beneath the Server Configuration work center node; you then click **Disable local continuous replication** in the Action pane, as shown in Figure 8.24.

Figure 8.24 Disable LCR Action Link



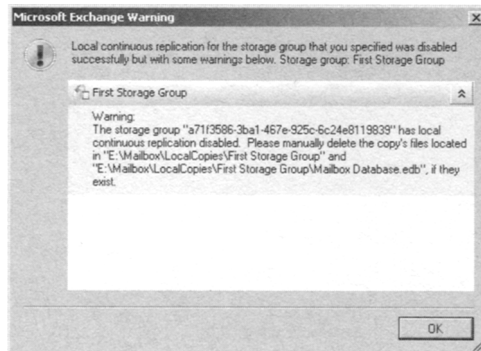
When we disable LCR for a storage group, we'll get the warning message shown in Figure 8.25, which tells us that LCR will be disabled for the replication database copy for the respective Storage Group. Since this is exactly what we want to do, click **Yes**.

Figure 8.25 Disabling Local Continuous Replication Confirmation



After we click **Yes**, believe it or not we'll get an additional warning message. This one informs us that we must delete the files (that is, the log files, EDB database, and so on) manually from the path (which in this example is `E:\Mailbox\LocalCopies\First Storage Group`) we specified when we originally enabled LCR (see Figure 8.26). Once you have clicked **OK** and deleted these files, LCR has been properly disabled.

Figure 8.26 Disabling Local Continuous Replication



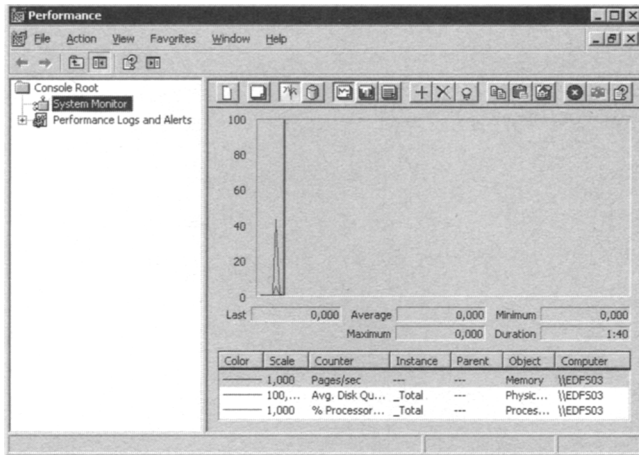
If you want to disable LCR for a Storage Group via the EMS, you need to do so using the `Disable-StorageGroupCopy` cmdlet. To disable the LCR for the First Storage Group, type `Disable-StorageGroupCopy -Identity "First Storage Group"`. When you do, you'll get the same warning message as the one shown in Figure 8.25.

Local Continuous Replication Performance Objects and Counters

When the Exchange 2007 Mailbox Server role is installed, setup adds two LCR-related performance objects to the Windows 2003 Performance Monitor. To open the Performance Monitor, either click **Start** | **Run** and type `Perfmon` or click **Start** | **Administrative**

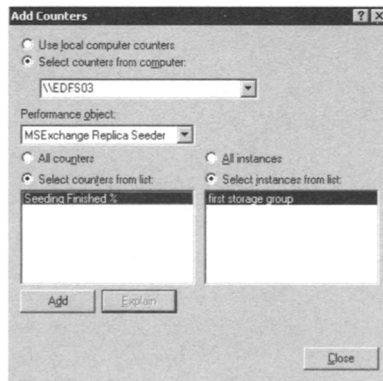
Tools and select **Performance**. This will bring up the Performance Monitor, shown in Figure 8.27.

Figure 8.27 The Performance Monitor



The first object is the MExchange Replica Seeder performance object, which, as you can see in Figure 8.28, contains only one counter, called Seeding Finished %. This counter is used to show the progress of database seeding in percent. When you add this counter, you can choose which instance (in this case, the particular storage group) you want to view the database seeding for.

Figure 8.28 Continuous Replication Performance Objects



The MExchange Replication performance object contains at least 14 different counters (see Table 8.2).

Table 8.2 Continuous Replication Performance Counters

Performance Counter	Description
Copy Queue Exceeds Mount Threshold (CCR only)	Copy Queue Exceeds Mount Threshold (CCR Only) is 1 if the copy queue length is larger than the Mount Threshold specified by the Auto Database Mount Dial. This counter is used only with CCR. It will always be 0 with LCR.
<i>CopyGenerationNumber</i>	Copy Generation Number is the generation of the last log file that has been copied.
<i>CopyNotificationGenerationNumber</i>	Copy Notification Generation Number is the generation of the last log file the copier knows about.
<i>CopyQueueLength</i>	Copy Queue Length is the number of log generations waiting to be both copied and inspected successfully.
<i>Failed</i>	Failed is 1 if the replica instance is set to failed, otherwise 0.
<i>InspectorGenerationNumber</i>	Inspector Generation Number is the generation of the last log file that has been inspected.
<i>ReplayBatchSize</i>	Replay Batch Size is the number of log generations replayed together.
<i>ReplayGenerationNumber</i>	Replay Generation Number is the generation of the last log file that has been replayed successfully.
<i>ReplayGenerationsComplete</i>	Replay Generations Complete is the number of log generations already played in the current replay batch.
<i>ReplayGenerationsPerMinute</i>	Replay Generations Per Minute is the rate of replay in log generations per minute in the current replay batch.
<i>ReplayGenerationsRemaining</i>	Replay Generations Remaining is the number of log generations remaining to be played in the current replay batch.
<i>ReplayNotificationGenerationNumber</i>	The generation of the last log file that replay knows about.
<i>ReplayQueueLength</i>	Replay Queue Length is the number of log generations waiting to be replayed.
<i>Suspended</i>	Suspended is 1 if the continuous replication is suspended. When the continuous replication is suspended, logs are not copied and replayed into the passive copy.

As you can see, all these counters can be used to determine how replication for an LCR-enabled storage group have progressed, but a high-availability feature such as LCR should really be monitored using a proactive and automated monitoring solution such as Microsoft Operation Manager (MOM) with the Exchange Server 2007 Management Pack installed.

Managing a Cluster Continuous Replication-Based Setup

Exchange Server 2007 introduces another new high-availability feature called Cluster Continuous Replication (CCR). This feature takes the new Exchange Server 2007 log file shipping and replay mechanisms (known as continuous replication) and combines them with the features that are available in a more traditional two-node Windows 2003 server active/passive cluster setup. A traditional two-node active/passive cluster has its benefits but has also always had one major drawback: You still have a single point of failure when it comes to the information stores. CCR provides redundancy for both Exchange Services and the information stores.

As is the case with traditional Exchange clusters, CCR uses Windows Clustering Services to provide virtual servers (which, in Exchange 2007, are called clustered mailbox servers) and failover capabilities. CCR has one big difference from traditional clusters, though, and that is that functionality doesn't require any kind of shared storage subsystem, because each node contains a local copy of the information stores. This eliminates the dependency on SAN technology in the cluster design, which makes CCR a more cost-efficient solution because you can use a storage option such as Direct Attached Storage (DAS) or Serial Attached SCSI.

With CCR, the transaction logs generated on the active node are replicated to the information store on the passive node using log file shipping. These replicated log files are then posted into the database(s) on the passive node using the log file replay technology. This means that should the active node or a database on this node fail or for some other reason go offline, an automatic failover to the passive node will occur. When the passive node becomes the active node, the replication of log files will happen from the new active node to the passive node.

Another thing worth mentioning about CCR is that the feature supports stretched clustering (called *geoclustering*), but bear in mind that the nodes must belong to the same subnet. This means that as the cluster is stretched between the locations, the subnet must be stretched, too.

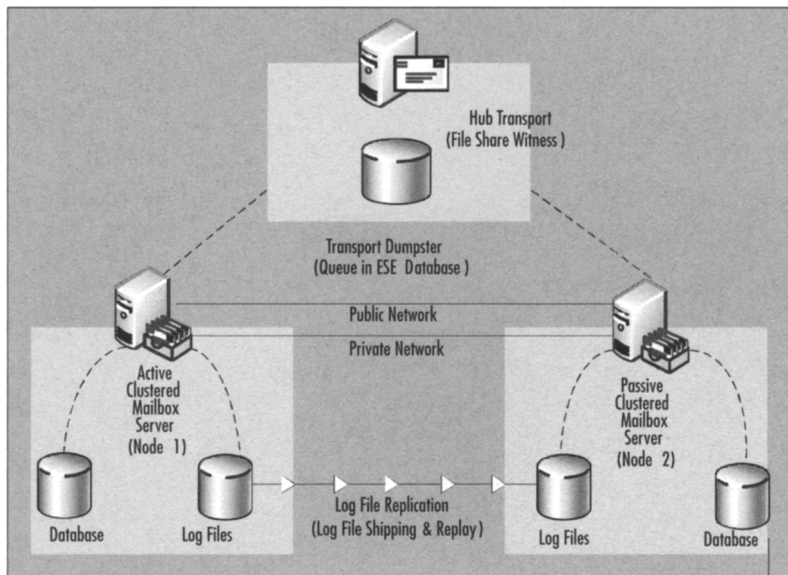


When Exchange 2007 supports Longhorn server (which will be provided via a service pack when the Longhorn product has been released), we will be able

to take advantage of stretched clustering spanning multiple subnets, both on the public as well as the private network (also called the heartbeat network).

Last but not least, you can reduce the frequency of backups and restores as well as perform backups of the databases on the passive node, and thereby not impact the performance of the active node. In Figure 8.29 you can see a basic CCR scenario.

Figure 8.29 A Basic Cluster Continuous Replication Scenario



Prerequisites

To set up a CCR-based cluster, the following are required:

- A Windows 2003 Active Directory forest with at least one domain controller (raised to 2000 or 2003 forest functional level)
- Two Windows 2003 Server R2 Enterprise Editions or Windows 2003 Server SP1 Enterprise Editions
- One Windows File Share Witness, which is recommended to be an Exchange 2007 Hub Transport Server in the existing Exchange 2007 organization; note that CCR-based clusters don't use a shared quorum as traditional clusters do
- A Cluster Service Account in the Active Directory forest (we'll create this one later in this section)

You also need to apply the update mentioned in MS KB article 921181 to both servers that will act as nodes in the Exchange Server 2007 Clustered Mailbox setup. The update adds a new file share witness feature to the current Majority Node Set (MNS) quorum model. The file share witness feature lets you use a file share that is external to the cluster as an additional “vote” to determine the status of the cluster in a two-node MNS quorum cluster deployment, which is a requirement to use the CCR functionality in Exchange Server 2007.

To deploy CCR, the following hardware requirements must be met:

- Two network interface cards (NICs) installed in each node—one for the public and one for the private cluster network (the heartbeat network)
- Extra sets of disks or a DAS, SAN, or Serial SCSI solution to hold the database and transaction log files

In addition to the software and hardware requirements, you also should be aware of the following general requirements:

- When dealing with CCR environments, you must and can only use one database per storage group.
- You cannot create a public folder database in a CCR environment if you already have more than one public folder database in your organization.
- In a CCR environment, Microsoft recommends that you create no more than 30 storage groups and databases (one database per storage group) on the clustered mailbox server.
- The cluster on which Exchange 2007 is installed cannot contain Exchange Server 2000/2003 or any version of Microsoft SQL Server. Running Exchange 2007 in a cluster with any of these other applications is simply not supported.

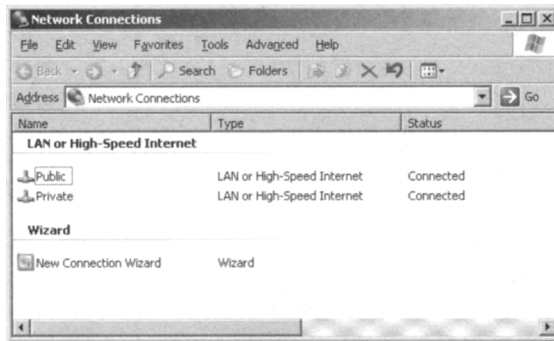
SOME INDEPENDENT ADVICE

Some of you might wonder whether the licensing rules have changed regarding Exchange 2007 cluster setups. Unfortunately, this isn't the case; you still have to purchase an Exchange 2007 Enterprise Edition CAL for each node in your cluster (also any passive nodes). The reason is that the passive node still runs Exchange code although the node is the passive one.

Configuring the Network Interface for Each Node

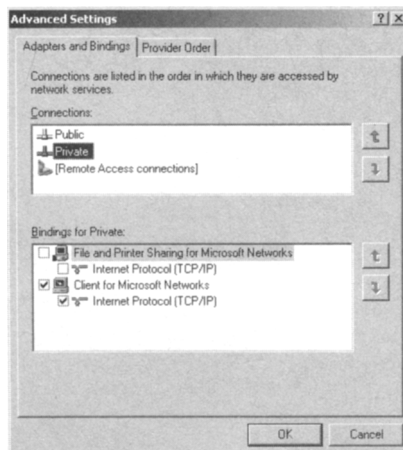
When you start the servers that are to be the nodes in the cluster, begin by naming the machines EDFS07 and EDFS08 or whatever naming scheme you want to use. (These names have nothing to do with the Exchange server name that your clients will be configured to connect to later.) Now name the two network connections Public and Private (see Figure 8.30) for the external and the internal networks, respectively. Remember to do this on both nodes.

Figure 8.30 Network Connections



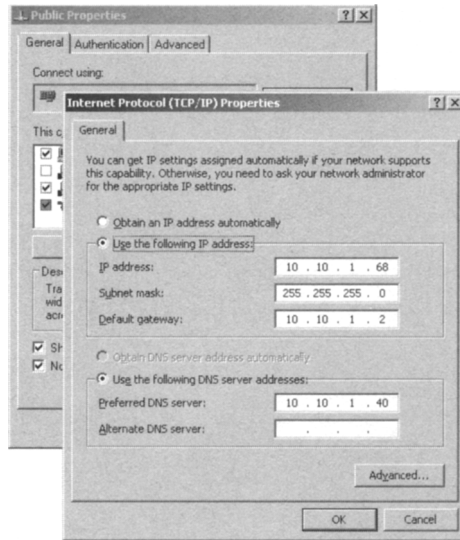
1. Click **Advanced | Advanced Settings**. If it's not already the case, make sure Public is listed first on the binding order list, then Private, and Remote Access Connections last. Also make sure that you clear the **File and Printer Sharing** check box for Microsoft Networks for the Private network connection, as shown in Figure 8.31.

Figure 8.31 Binding Order



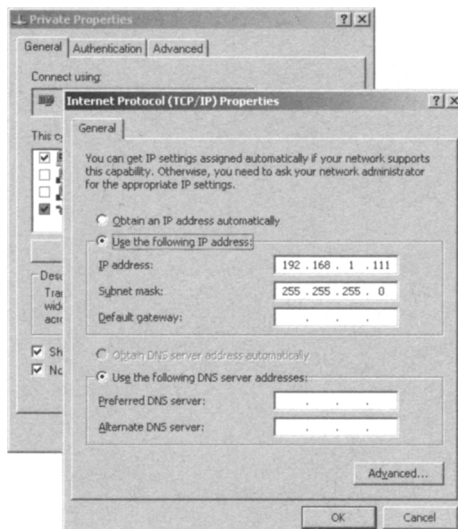
2. Now configure the **Public** network with the TCP/IP settings that should be used in your environment (see Figure 8.32).

Figure 8.32 Configuring the Public Network Interface



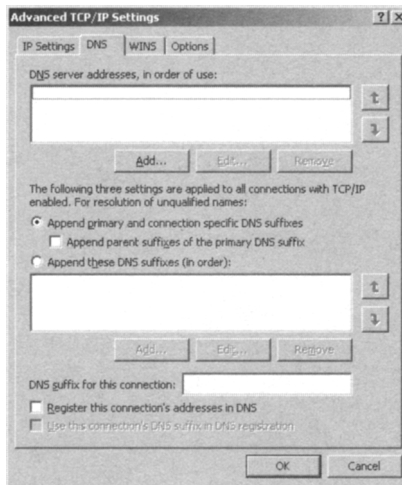
We also need to configure the Private network with an IP address and a subnet mask, as shown in Figure 8.33. Nothing else is required, since this network is used only for communication (heartbeats) between the nodes in the cluster.

Figure 8.33 Configuring the Private Network Interface



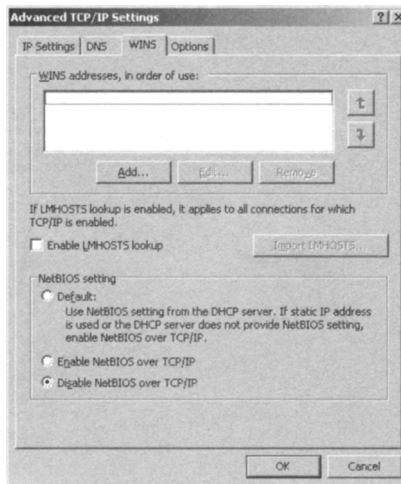
1. Click **Advanced**, then select the **DNS** tab. Here you should clear both the **Register this connection's addresses in DNS** and **Use this connection's DNS suffix** check boxes, as shown in Figure 8.34.

Figure 8.34 Configuring DNS Settings for the Private Network Interface



2. Click the **WINS** tab. Clear the **Enable LMHOSTS lookup** option and select **Disable NetBIOS over TCP/IP**, as shown in Figure 8.35.

Figure 8.35 Configuring WINS Settings for the Private Network Interface

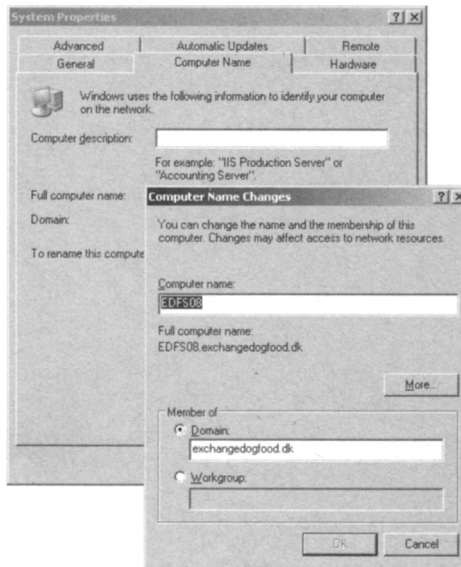


3. Click **OK** three times and close the Network Connections window.

Adding the Servers to the Active Directory Domain

Since a CCR setup requires both nodes to be part of the same Active Directory domain, now would be a good time to make this the case. You can add the nodes to the domain by right-clicking **My Computer** and selecting **Properties** in the context menu. Now click the **Computer Name** tab (see Figure 8.36), then the **Change** button, and specify the domain.

Figure 8.36 Adding the Nodes to the Domain



When you have added both servers to the domain as well as rebooted each, we can move on to creating the necessary cluster service account.

Creating a Cluster Service Account

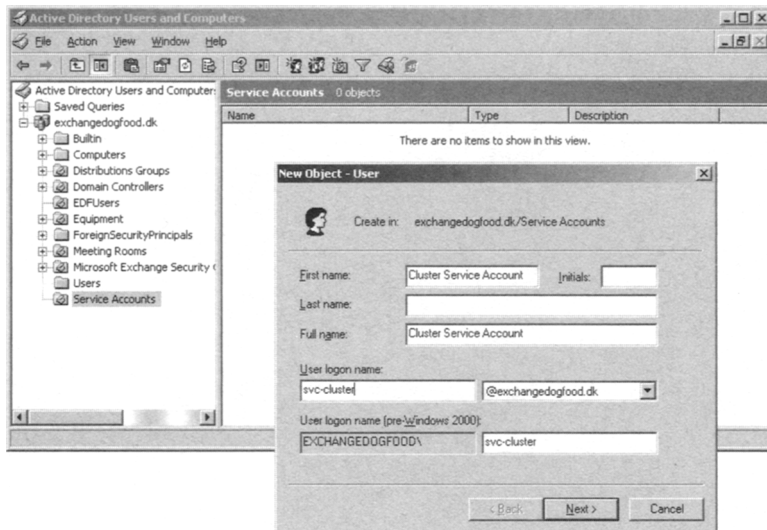
Because each node belonging to the cluster needs to use the same account, we need to create a cluster service account.

The cluster service account must be a member of either the Exchange Server Administrators (ServerName) group or the Exchange Organization Administrators group. In addition, it must be a member of the local administrators group on each node in the cluster. For our purposes, we'll add it to the Exchange Organization Administrators group.

To create the cluster service account:

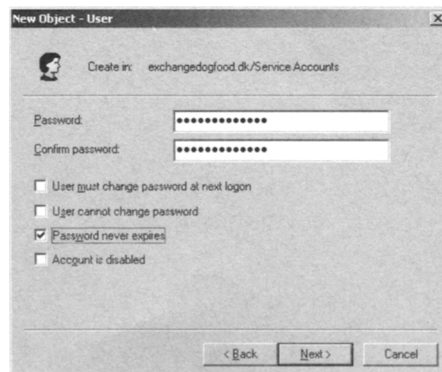
1. Log onto a domain controller in the respective Active Directory domain, then click **Start | Run** and type **DSA.msc** to open the Active Directory Users and Computers MMC snap-in. Now right-click the Organizational Unit (OU) in which you want the service account to be created, then choose **New | User** in the context menu. Give the account a meaningful name and user logon name (such as **Cluster Service Account** and **svc-cluster**), as shown in Figure 8.37. Now click **Next**.

Figure 8.37 Creating the Cluster Service Account



2. Give the service account a complex password and uncheck **User Must change password at next logon**, then check **Password never expires**, as shown in Figure 8.38. Click **Next**.

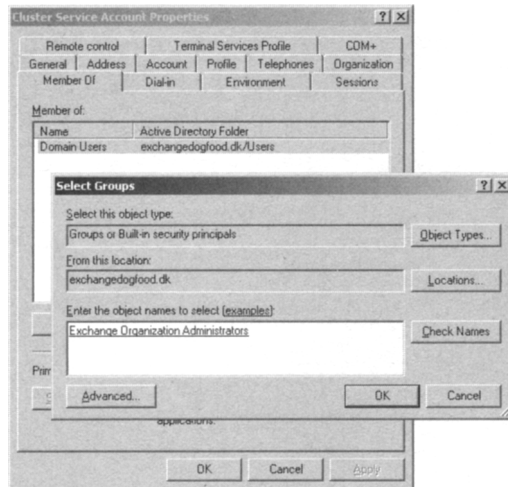
Figure 8.38 Specifying the Password for the Cluster Service Account



On the New User object completion page click Finish.

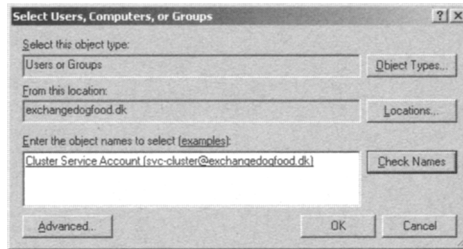
- Now we need to give the new cluster service account the appropriate permissions. To do so, open the **Properties** page for the user object and select the **Member Of** tab. Make sure it's the respective Active Directory domain that's shown in the **From this location** field, then click the **Add** button and type **Exchange Organization Administrators**, as shown in Figure 8.39. Click **OK**.

Figure 8.39 Adding the Cluster Service Account to the Exchange Organization Administrators Group



- Now switch over to the server that will be the first node in the cluster and click **Start | Run**. Type **compmgmt.msc**. Expand **Local Users and Groups** and select the **Groups** container. Open the **Properties** page for the Administrators group object in the right pane, then click the **Add** button. Make sure that the Active Directory domain is shown in the **From this location** field, as shown in Figure 8.40, and type **Cluster Service Account** (or whatever name you gave the account in your setup). click **Check Names** to verify that it resolves successfully. Click **OK** and close the Computer Management MMC snap-in.

Figure 8.40 Adding the Cluster Service Account to the Local Administrators Group



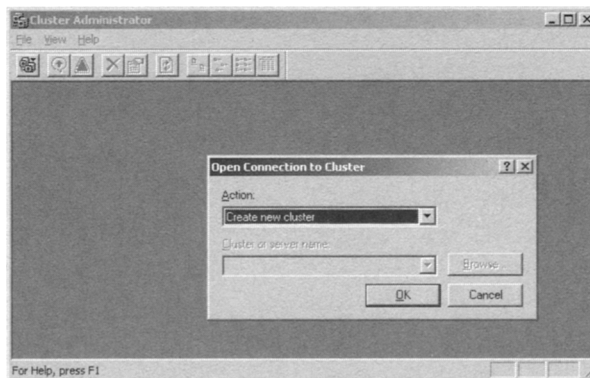
5. Repeat Steps 1–4 for the server that will be the second node in the cluster.

Creating and Configuring the Windows 2003 Server Cluster

Now that the two servers are ready to act as nodes in a Windows 2003 cluster, it's time to create the actual Windows 2003 Server Cluster. To do so:

1. Log onto EDFS07 with a domain admin account, then click **Start | Administrative Tools | Cluster Administrator**, and select **Create new cluster** in the drop-down box. Click **OK** and then click **Next**, as shown in Figure 8.41.

Figure 8.41 Creating a New Cluster

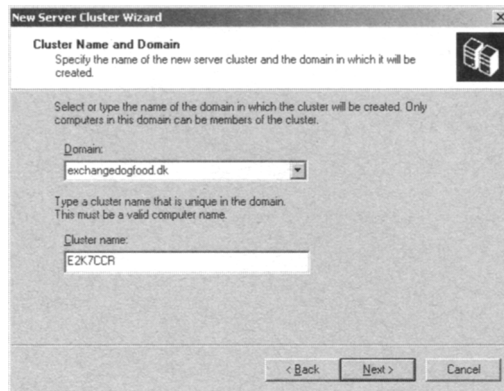


NOTE

You can also open a command prompt and type **Cluster.exe /create /wizard** to start the Cluster Wizard.

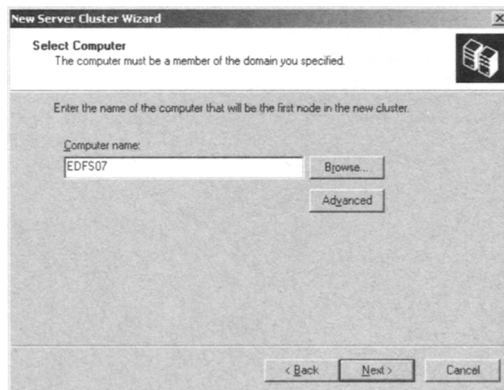
2. Now specify the domain name as well as the cluster name (the name for the Windows 2003 cluster, *not* the Exchange cluster name to which the clients will connect) as shown in Figure 8.42, then click **Next**.

Figure 8.42 Specifying the Cluster Name and Domain



3. If it's not already entered, type the name of the Windows 2003 server that is to be the first node in the cluster (in this case, **EDFS07**), then click **Next** (see Figure 8.43).

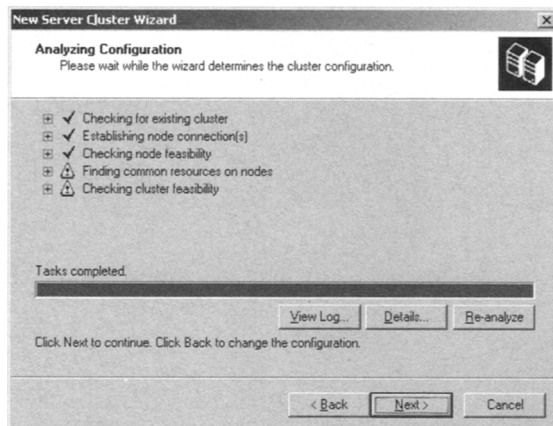
Figure 8.43 Adding the First Cluster Node to the New Cluster



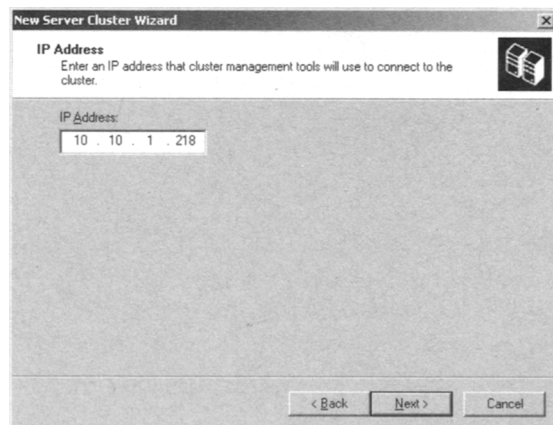
4. Let the Cluster Wizard determine the cluster configuration and click **Next**.

NOTE

You can ignore the two warnings shown in Figure 8.44, since the nodes in a cluster continuous replication-based mailbox server setup aren't going to share the same disk subsystem.

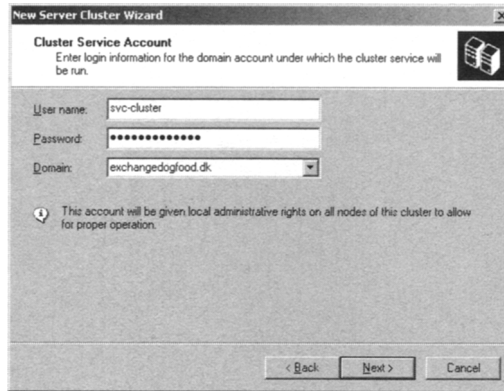
Figure 8.44 Analyzing Cluster Configuration

5. Now enter the IP address that the cluster management tools should use to connect to the cluster (in this case, **10.10.1.218**) and click **Next** (see Figure 8.45).

Figure 8.45 Specifying the IP Address to Which the Cluster Management Tools Should Connect

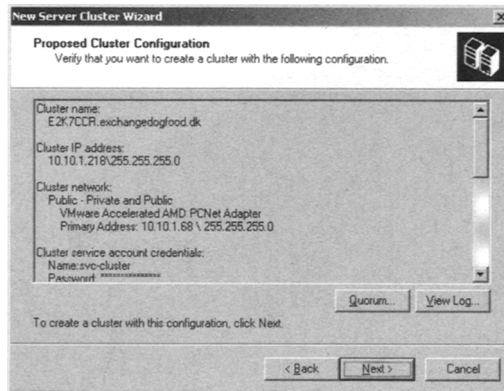
6. Enter the credentials of the cluster service account and click **Next** (see Figure 8.46).

Figure 8.46 Entering the Credentials of the Cluster Service Account



7. Now click **Quorum** and select **Majority Node Set** as the resource type, then click **OK** and **Next** (see Figures 8.47 and 8.48).

Figure 8.47 Proposed Cluster Configuration




NOTE

The Majority Node Set resource type manages cluster configuration data that might or might not be on a cluster storage device. For example, the Majority Node Set resource type can manage cluster configuration data that is actually stored on multiple nodes in a cluster at the same physical location or in a

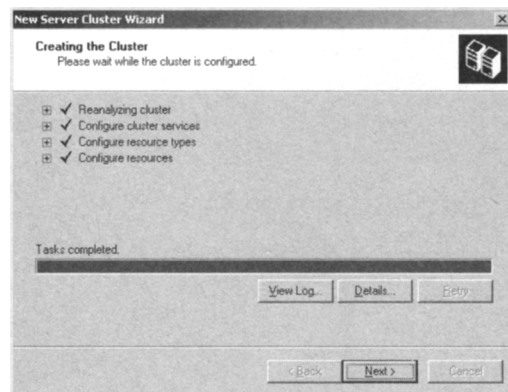
geographically dispersed cluster. The Majority Node Set resource ensures that the cluster configuration data is kept consistent across the various nodes.

Figure 8.48 Setting Majority Node Set as the Resource Type



8. Now wait for the cluster to be configured, then click **Next** (see Figure 8.49).

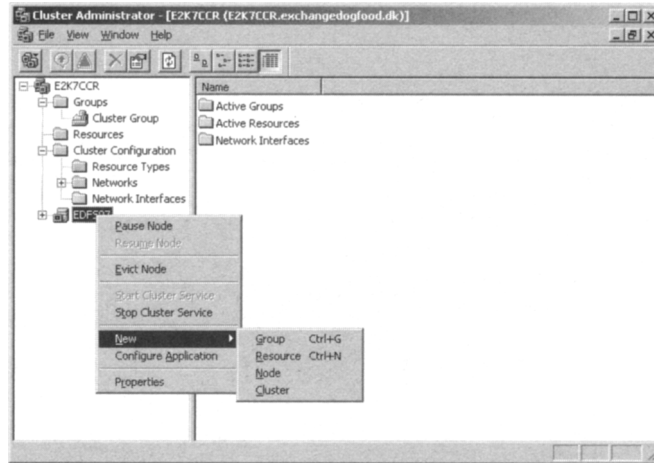
Figure 8.49 Creating the Cluster



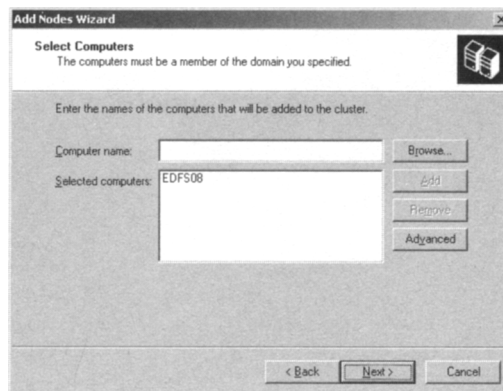
9. When the cluster has been completed successfully, you can click **Finish**.

We now have a full working Windows 2003 cluster running, but since there's only one node, it's not very fault tolerant. So let's add the second Windows 2003 server too. Do the following:

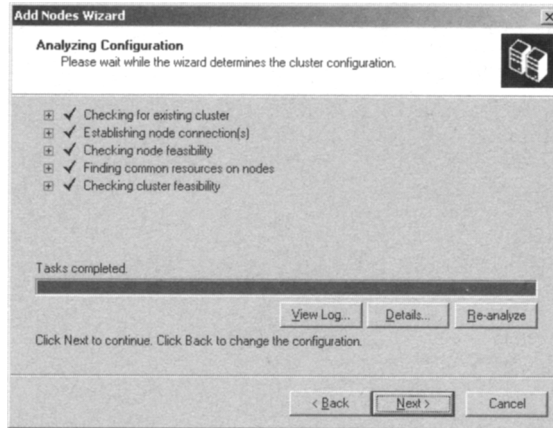
1. Right-click **EDFS07** in the left pane of the Cluster Administrator, then selecting **New | Node**, as shown in Figure 8.50.

Figure 8.50 Adding a Second Node to the Cluster

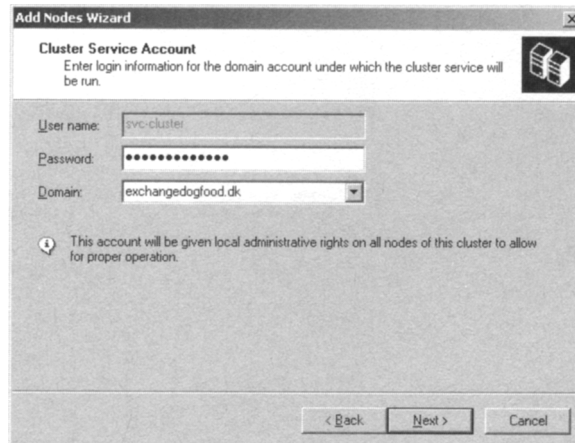
2. The Add Nodes Wizard will launch and you can click **Next**. Enter the name of the server that is going to be the second node (for the purpose of this example, **EDFS08**), then click **Next** (see Figure 8.51).

Figure 8.51 Entering the Name of the Second Node

3. Again, let the Add Notes Wizard determine the cluster configuration, then click **Next** (see Figure 8.52).

Figure 8.52 Analyzing Cluster Configuration

- 4 Enter the password for the cluster service account (in this case, **svc-cluster**, which we created earlier in the chapter), then click **Next** (see Figure 8.53).

Figure 8.53 Entering the Password for the Cluster Service Account

5. When you are verified, you'll want to add the second node to the cluster with the configuration shown in Figure 8.54. Click **Next**.
6. When the cluster has been configured properly without any errors or warnings (see Figure 8.55), click **Next**.

Figure 8.54 Proposed Cluster Configuration for Node Two

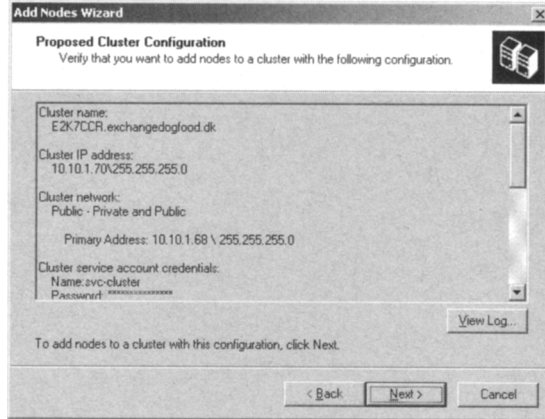
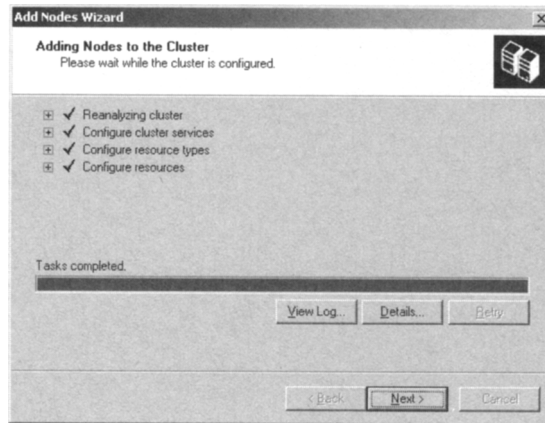
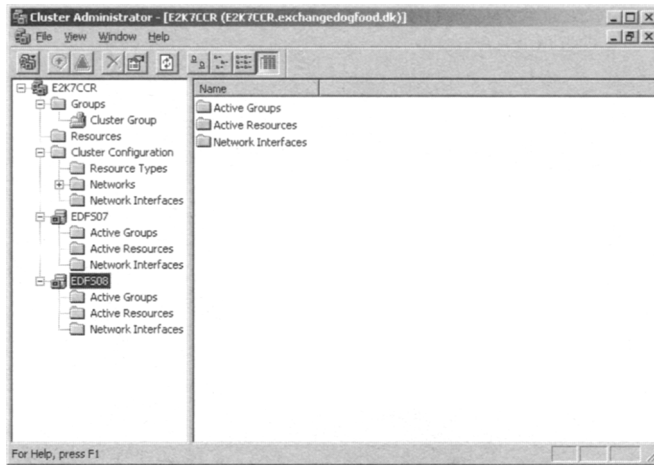


Figure 8.55 The Cluster Is Configured for the Second Node



7. When the Add Notes Wizard has completed successfully, click **Finish**.

The second Windows server is now part of the cluster, as shown in Figure 8.56.

Figure 8.56 Cluster Administrator with Two Nodes

Installing the Necessary Windows Components

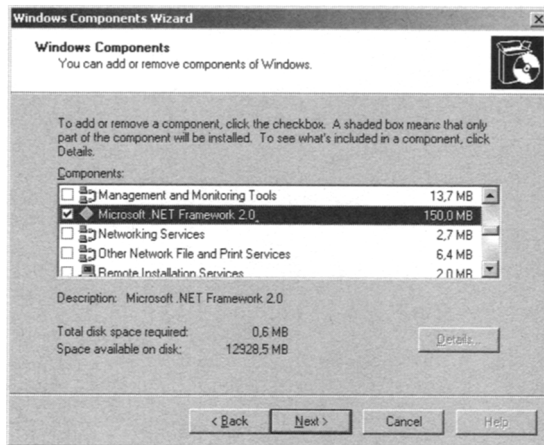
Before we move on to install the Exchange Server 2007 binaries, we need to make sure that the required Windows components have been installed. All types of Exchange Server 2007 installations (no matter what server role we're talking about) need the Microsoft .NET Framework 2.0 component installed.

If you have installed Windows Server 2003 Enterprise Edition with Service Pack 1 on the nodes, you need to download the Microsoft .NET Framework Version 2.0 Redistributable Package (x86) from Microsoft.com, since it's only a standard Windows component for Windows Server 2003 R2. If you're using Windows Server 2003 R2-based servers, you can install the component by clicking **Start | Control Panel | Add or Remove Programs | Add/Remove Windows Components**, checking the **Microsoft .NET Framework 2.0** check box as shown in Figure 8.57, then clicking **Next**.

Since we're deploying a clustered mailbox server, we also need to install the following IIS 6.0 components on each node:

- Enable network COM+ access
- Internet Information Services
- World Wide Web Service

Figure 8.57 Installing the Microsoft .NET Framework 2.0 Windows Component



When you have done so, you can move on to configure the File Share Witness.

Configuring the Majority Node Set Quorum with File Share Witness

No doubt some of you are thinking: What the heck is a Majority Node Set quorum with File Share Witness? We can understand why; this is a completely new type of quorum model that is made available by installing the update (MS KB article 921181) mentioned in the beginning of this chapter section. The update makes it possible to use a file share witness that is external to the cluster as an additional “vote” to determine the status of the cluster in a two-node MNS quorum cluster deployment, which is a requirement for using the CCR functionality in Exchange Server 2007.

The file share for this file share witness can be located on any type of Windows server in your environment, but best practice is to use an Exchange 2007 Hub Transport server in the Active Directory server site containing the nodes in the respective cluster. We'll also use a Hub Transport server in this example.

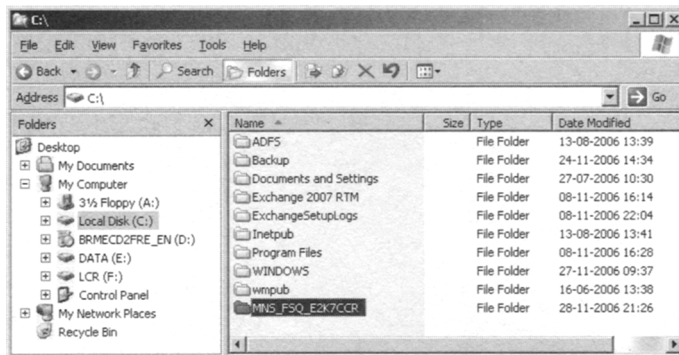
The first thing you need to do is to create the file share on the Hub Transport server. You can do this either via the CLI or by using the GUI. In this example we'll use the GUI:

1. Log on to the Hub Transport server with a domain admin account, then open Windows Explorer and create a new folder called **MNS_FSQ_E2K7CCR** on the C: drive or wherever you want it to be created, as shown in Figure 8.58.

NOTE

It's recommended that you use the *MNS_FSQ_clustername* naming convention when you create this folder.

Figure 8.58 The Majority Node Set File Share Quorum Folder



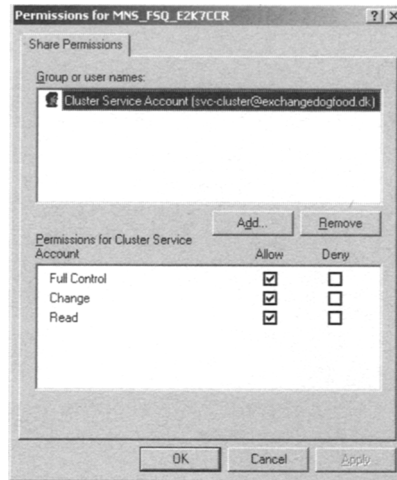
2. Now open the **Properties** page for the newly created folder and click the **Sharing** tab (see Figure 8.59).

Figure 8.59 The Majority Node Set File Share Quorum Folder Share



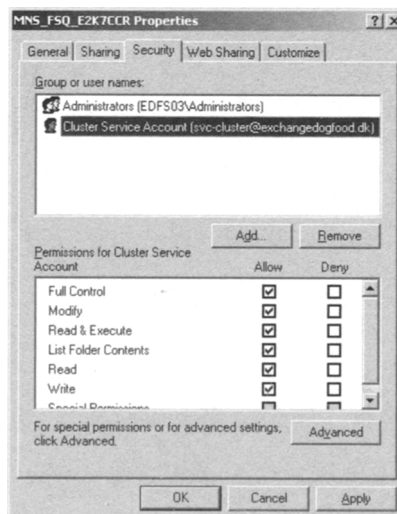
3. Click **Permissions** and configure the share permissions so that only the Cluster Service Account is allowed access to this share (see Figure 8.60).

Figure 8.60 Share Permissions for the Majority Node Set File Share Quorum Folder



4. Click **OK**, then select the **Security** tab. Here you should give Full Control to the local administrator and the cluster service account, as shown in Figure 8.61. Make sure you clear **Allow inheritable permissions from the parent to propagate to this object and all child objects** when doing so, then click **OK** twice and log off the server.

Figure 8.61 Security Permissions to the Majority Node Set File Share Quorum Folder



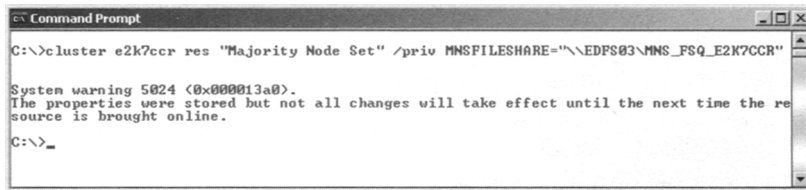
- Back on EDFS07 or EDFS08, we now need to set the Majority Node Set Private Property attribute to point to the file share we just created. We do so by opening a command prompt, then issuing the command **Cluster res "Majority Node Set" /priv MNSFileShare=\\EDFS03\MNS_FSQ_E2K7CCR**.

NOTE

Make sure to replace the server name so that it matches the name of the Hub Transport server in your environment.

You will get a warning that all properties were stored but not all changes will take effect until the next time the resource is brought online, just as is shown in Figure 8.62.

Figure 8.62 Configuring the Majority Node Set on EDFS07



```

C:\>cluster e2k7ccr res "Majority Node Set" /priv MNSFILESHARE="\\EDFS03\MNS_FSQ_E2K7CCR"

System warning 5024 (0x000013a0).
The properties were stored but not all changes will take effect until the next time the resource is brought online.

C:\>_
  
```

SOME INDEPENDENT ADVICE

In a couple of the CCR-based cluster deployments I've done, I have gotten an error message similar to the following when running the command *Cluster res "Majority Node Set" /priv MNSFileShare=\\EDFS03\MNS_FSQ_E2K7CCR*:

Too many command line parameters have been specified for this option.

See **"CLUSTER RESOURCE /?"** for correct syntax

Should you experience this error, too, you should be able to get going using the following command syntax instead:

```
Cluster <ClusterName> res "Majority Node Set" /priv
MNSFileShare=UNCPath
```

- To force all changes to take effect, we will move the cluster group from one node to the other (taking the cluster group offline and online again). Do this using the command **Cluster Group "Cluster Group" /Move**. When you have done so,

you will see that the cluster group is now online on E2K7Node2, as shown in Figure 8.63.

Figure 8.63 Moving the Cluster Group from One Node to the Other

```

C:\>Cluster group "Cluster Group" /move
Moving resource group 'Cluster Group' ...
Group           Node           Status
-----
Cluster Group   EDFS08        Online
C:\>_
  
```

- Now let's verify that the `/Priv` property is set correctly. This can be done by issuing the command **Cluster Res "Majority Node Set" /Priv**.

As you can see in Figure 8.64, this property has been set correctly for the purposes of our example.

Figure 8.64 Verifying That the Property of `/Priv` Is Set Correctly

```

C:\>Cluster res "Majority Node Set" /priv
Listing private properties for 'Majority Node Set':
T Resource           Name                               Value
-----
S Majority Node Set  MNSFileShare                      \\EDFS03\MNS_FSQ_E2K7CCR
D Majority Node Set  MNSFileShareCheckInterval         240 (0xf0)
D Majority Node Set  MNSFileShareDelay                  4 (0x4)
C:\>
  
```

Configuring the Transport Dumpster

When deploying a CCR-based cluster in your environment, an important step is to enable the Transport Dumpster on the Hub Transport server.

The Transport Dumpster is a new feature of the Exchange 2007 Hub Transport server that can submit recently delivered mail after an unscheduled outage. For an e-mail message to be able to be retained in the Transport Dumpster, at least one of the message recipients must have his or her mailbox located on a CCR-based mailbox cluster server, because the Transport Dumpster works only with mailboxes located on a CCR-based mailbox server cluster. As mentioned earlier in this chapter, with CCR the replication of mailbox data from the active node to the passive node is *asynchronous*, which means that the passive node will always lag behind the active node (although not by much). This means that should a failure

of the active node occur, there's a chance that not all transaction log files will have been replicated to the passive node before this happens. This is where the Transport Dumpster comes into the picture. It can resubmit recently delivered mail and thereby constitute for the majority of the changes in the database(s). When a failure of the active node results in a lossy failover to the passive node, the cluster mailbox server will ask all the Hub Transport servers in the site to redeliver any lost mail.

NOTE

Should any of the messages that are being resubmitted to the cluster mailbox server be duplicates, the store is intelligent enough to discard any duplicates it finds.

The Transport Dumpster is enabled by default; you can see the default configured settings by running the *Get-TransportConfig* CMDlet.

Microsoft recommends that you configure the *MaxDumpsterSizePerStorageGroup* parameter, which specifies the maximum size of the Transport Dumpster queue for each storage group to a size that is 1.25 times the size of the maximum message that can be sent. For example, if the maximum size for messages is 10 megabytes (MB), you should configure the *MaxDumpsterSizePerStorageGroup* parameter with a value of 12.5 MB. In addition, Microsoft recommends that you configure the *MaxDumpsterTime* parameter, which specifies how long an e-mail message should remain in the Transport Dumpster queue, to a value of 07.00:00:00, which is seven days. This amount of time is sufficient to allow for an extended outage to occur without loss of e-mail. When you use the Transport Dumpster feature, additional disk space is needed on the Hub Transport server to host the Transport Dumpster queues. The amount of storage space required is roughly equal to the value of *MaxDumpsterSizePerStorageGroup* multiplied by the number of storage groups.

You use the *Set-TransportConfig* CMDlet to enable and configure the Transport Dumpster. So, for example, to configure the maximum size of the dumpster per storage group to 25 MB with a dumpster life of 10 days, you would need to run the command *Set-TransportConfig -MaxDumpsterSizePerStorageGroup 25MB -MaxDumpsterTime 10.00:00:00*.

To see the *MaxDumpsterSizePerStorageGroup* and *MaxDumpsterTime* configuration settings, you can type **Get-TransportConfig**, as shown in Figure 8.65.

Figure 8.65 Transport Configuration Settings

```

Machine: EDF503 | Scope: exchangedogfood.dk
[PS] C:\>Get-TransportConfig

ClearCategories                : True
GenerateCopyOfDSNFor          : <>
InternalSMTPServers           : <>
JournalingReportNdrTo         : <>
MaxDumpsterSizePerStorageGroup : 25MB
MaxReceiveTime                 : 10:00:00:00
MaxReceiveSize                 : unlimited
MaxRecipientEnvelopeLimit     : unlimited
MaxSendSize                    : unlimited
TLSReceiveDomainSecureList    : <>
TLSSendDomainSecureList      : <>
VerifySecureSubmitEnabled     : False
VoicemailJournalingEnabled    : True
Xexch50Enabled                 : True

[PS] C:\>_

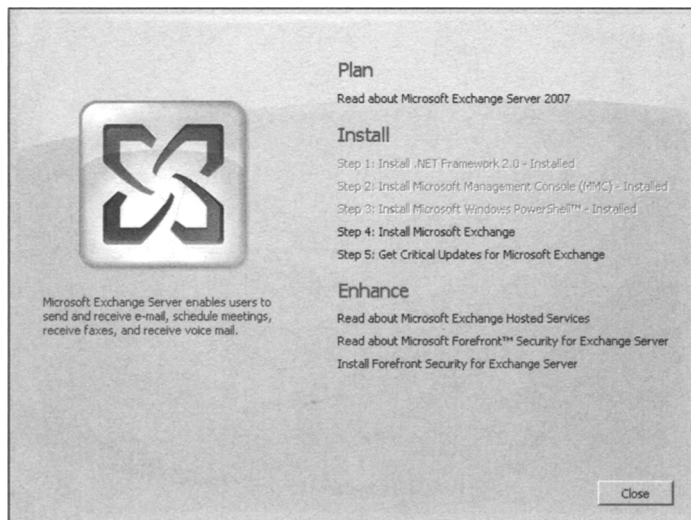
```

Installing Exchange 2007 on the Active Node

It's time to install the Exchange Server 2007 binaries on each node. We'll start with EDF507, which is the active node. To do so:

1. Double-click **Setup.exe** on the network share or the DVD media containing the Exchange 2007 setup files.
2. The Exchange Server 2007 Installation Wizard splash screen will launch, and as you can see in Figure 8.66, Step 1: Install .NET Framework 2.0, Step 2: Install Microsoft Management Console (MMC), and Step 3: Install Windows PowerShell have already been completed.

Figure 8.66 The Exchange Server 2007 Splash Screen



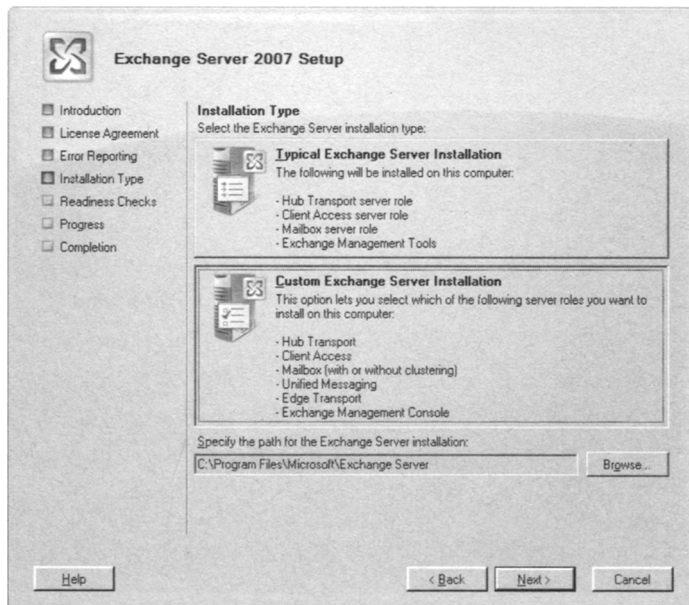
NOTE

If you have installed Windows Server 2003 with Service Pack 1 on each node, you need to download Microsoft Management Console (MMC) 3.0 and install it manually (by following the link in Step 2). But since I'm using Windows 2003 R2 Servers in my environment, the MMC 3.0 is installed by default.

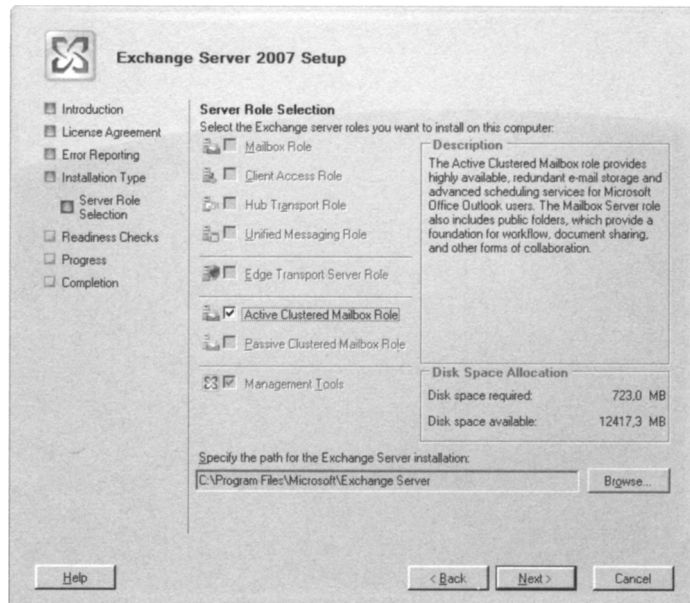
Click **Step 4: Install Microsoft Exchange**. Then click **Next** and accept the **License Agreement**. Click **Next** once again. Decide whether you want to enable error reporting or not (it's a good idea to enable this function, since the Exchange Product Group will receive any obscure errors you should experience in your CCR setup), then click **Next**.

3. Now select **Custom Exchange Server Installation** (see Figure 8.67) and click **Next**.

Figure 8.67 Selecting a Custom Exchange Server Installation



4. Check **Active Clustered Mailbox Role** as shown in Figure 8.68 and click **Next**.

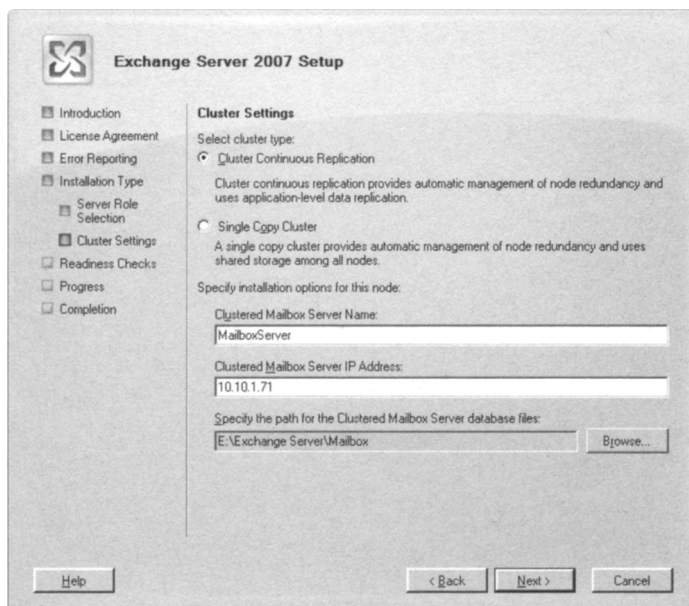
Figure 8.68 Selecting to Install an Active Clustered Mailbox Role

- Now select **Cluster Continuous Replication**, then specify a name for the mailbox server (the name you want your Outlook clients to connect to) and a unique IP address on your public network. Finally, specify the path for the clustered mailbox server database files (which in the example is **E:**) or use the default path (see Figure 8.69).

If you're installing CCR in a production environment, you should keep the transaction log files and database on separate disks, but if you're deploying CCR in a test environment, you simply use the default path.

- Let the readiness check complete, and if no issues are found, click **Next** to begin the installation.

Figure 8.69 Selecting to Install a Cluster Continuous Replication Cluster and Specifying the Name and IP Address of the Clustered Mailbox Server

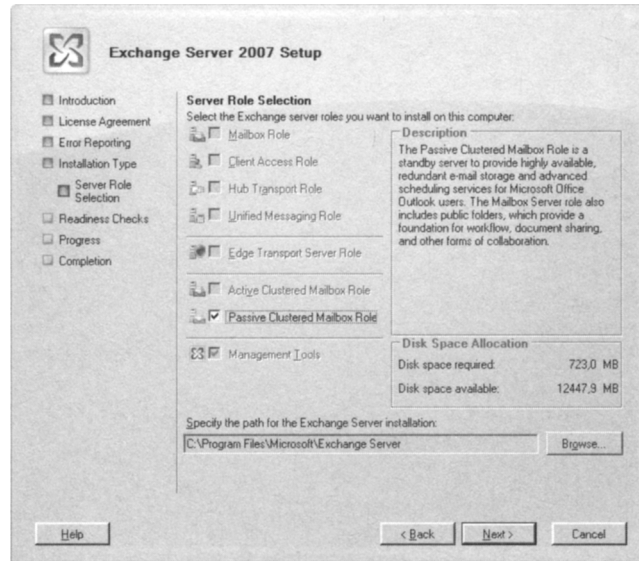


The Exchange Server 2007 Installation Wizard will now copy the needed Exchange files, install and configure the Mailbox Role, and finally create and configure the clustered mailbox server resources locally and create the object in Active Directory. After all steps have been completed, untick **Exit Setup** and open Exchange System Manager (yes, this will be corrected in a later build), then click **Finish**. We don't want to open the EMC just yet; we'll install Exchange on the second node first.

Installing Exchange 2007 on the Passive Node

Log on to EDFS08 with a domain admin account and do the same steps as we did when installing Exchange Server 2007 on EDFS07. The only difference is that you should select **Passive Clustered Mailbox Role** instead of Active Clustered Mailbox Role, as shown in Figure 8.70.

Figure 8.70 Installing the Passive Clustered Mailbox Role on the Second Node



Testing the Functionality of the Clustered Mailbox Server

It's time to verify that our Exchange 2007 clustered mailbox server is working as expected. Let's first open the Cluster Administrator and check whether the respective Exchange resources have been created. If you take a look at Figure 8.71, it looks good; we have both nodes listed in the left pane and all Exchange resources have been created and are currently owned by EDFS07.

Try to open the EMS by clicking **Start | All Programs | Microsoft Exchange Server 2007 | Exchange Management Shell** on one of the nodes, then type **Get-ClusteredMailboxServerStatus -Identity MailboxServer**. As you can see in Figure 8.72, the status of the clustered mailbox server is Online, and EDFS7 is currently the active node.

Figure 8.71 Listing All Exchange Cluster Resources in the Cluster Administrator

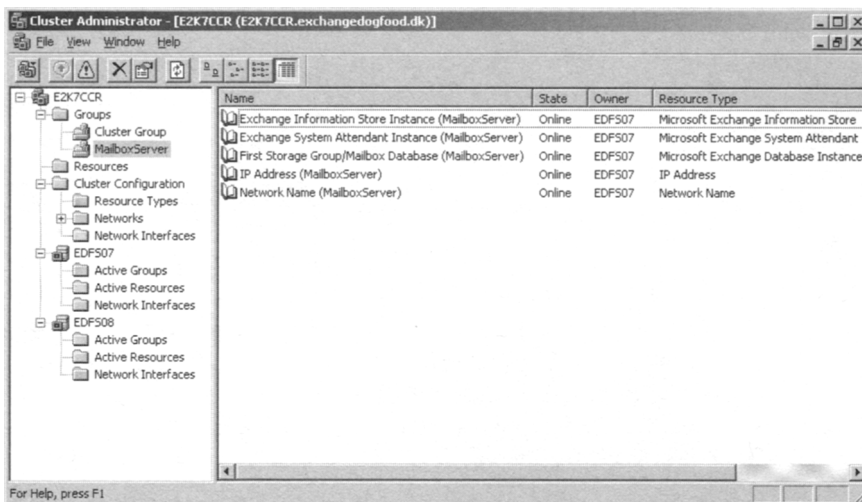


Figure 8.72 Requesting the Online Status of the Clustered Mailbox Server

```
Machine: EDF507 | Scope: exchangedogfood.dk
[PS] C:\>Get-ClusteredMailboxServerStatus -Identity MailboxServer

Identity                : MailboxServer
ClusteredMailboxServerName : MAILBOXSERVER.exchangedogfood.dk
State                   : Online
OperationalMachines     : <EDFS07 <Active, Quorum Owner>, EDF508>
FailedResources         : <>
IsValid                  : True
ObjectState              : Unchanged

[PS] C:\>
```

Now that we have verified that the clustered mailbox server is online, let's try to move the Exchange resources from node one to node two using the *Move-ClusteredMailboxServer* cmdlet. In the environment used in this chapter, we do so by issuing the command *Move-ClusteredMailboxServer -Identity:MailboxServer -TargetMachine:EDFS08 -MoveComment:"Verifying the Move Clustered Mailbox Server Functionality!"*

You're then asked to confirm this action. Type **Yes**, then press **Enter**. After a while the clustered mailbox resources will have been moved to the second node (EDFS08), as shown in Figure 8.73.

Figure 8.73 Moving the Clustered Mailbox Resources to the Second Node

```

Machine: EDF507 | Scope: exchangedogfood.dk
[PS] C:\>Move-ClusteredMailboxServer -Identity: MailboxServer -TargetMachine: EDF508 -Move
Comment: "Verifying the Move Clustered Mailbox Server Functionality!"
Confirm
Are you sure you want to perform this action?
Moving clustered mailbox server "MailboxServer" to target node "EDF508" with move comment
"Verifying the Move Clustered Mailbox Server Functionality!".
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): Y
[PS] C:\>_

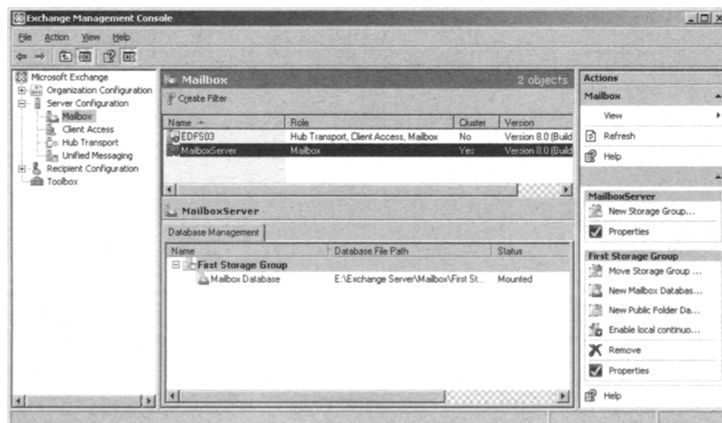
```

WARNING

Even though it's possible to move the cluster resource groups between nodes using the Cluster Administrator console, you should *always* do so using the `Move-ClusteredMailboxServer` CMDlet, because the Move Group task in the Cluster Administrator console isn't Exchange 2007 aware.

Viewing the Clustered Mailbox Server From Within the Exchange Management Console

Let's take a look at the clustered mailbox server in the EMC. To do so, click **Start | All Programs | Microsoft Exchange Server 2007 | Exchange Management Console**, then drill down to **Server Configuration | Mailbox**. Notice that the clustered mailbox server we named MailboxServer is listed in the Results pane and that it's recognized as a cluster server, as shown in Figure 8.74.

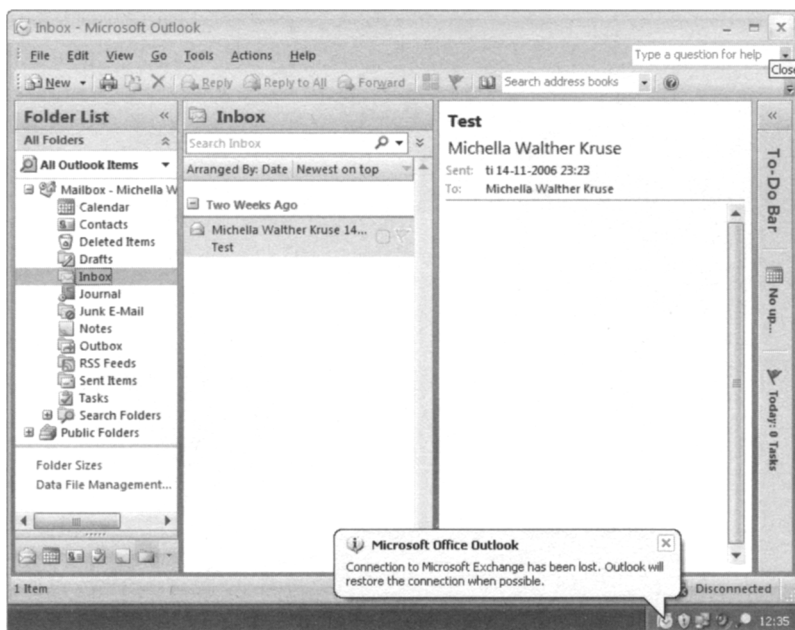
Figure 8.74 Viewing the Clustered Mailbox Server in the Exchange Management Console

Simulating a Failover from One Node to the Other

Now let's try to simulate a failover from EDFS08 (currently the active node) to EDFS07 so that we can see what will happen from the Outlook client perspective. To switch from one node to the other, we'll issue the CMDlet we used earlier in the chapter: *Move-ClusteredMailboxServer -Identity:MailboxServer -TargetMachine:EDFS07 -MoveComment:"Simulating a failover from one node to the other, seen from the end-user perspective"*.

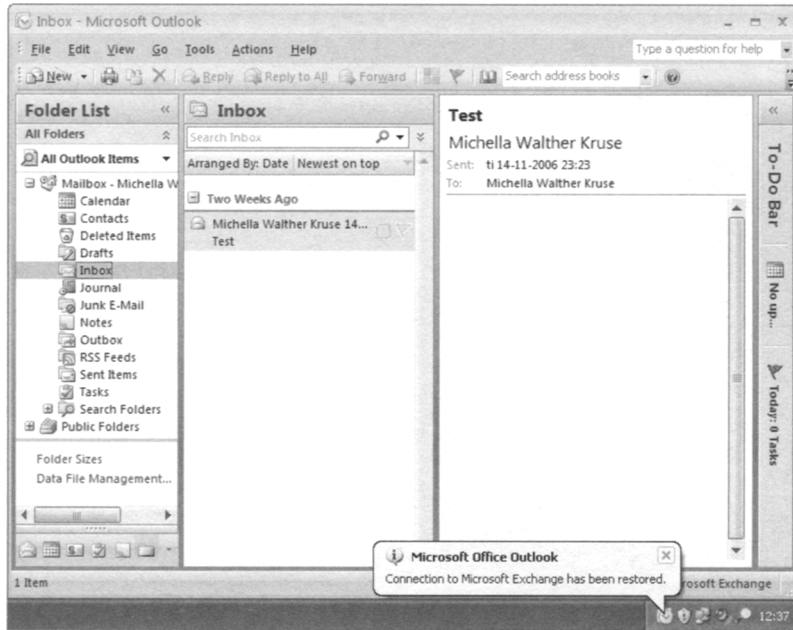
When a manual move or a failover occurs, the balloon shown in Figure 8.75 will appear because all services need to be stopped on EDFS07 before they can be moved and brought online on EDFS08.

Figure 8.75 Connection to the Exchange Server Has Been Lost



Depending on the number as well as the size of the databases in your Cluster Continuous Replication setup, this will take somewhere between 10 seconds to a couple of minutes, which shouldn't cause panic for the end users in the organization.

When EDFS08 has taken over, the end users will be notified that the connection to the Exchange Server has been restored (see Figure 8.76).

Figure 8.76 Connection to the Exchange Server Has Been Restored

As you have seen you throughout this chapter, you benefit from several advantages when you choose to install the Exchange 2007 Mailbox Server role in a Cluster Continuous Replication setup in your organization. The primary benefit is that you no longer have a single point of failure in regard to the Mailbox/Public Folder databases. Should the database on one node crash, an automatic failover to the other node containing the secondary database will be completed. This also means that you no longer need to use a shared storage system in the CCR setup, as is the case with Exchange 2007 Single Copy Clusters as well as cluster setups in previous versions of Exchange. In addition, the two nodes in the CCR setup can even be placed in two different locations, as long as they belong to the same subnet. Not only that, the installation of the Exchange 2007 cluster has also been further simplified over previous versions. Since the CCR setup uses log file shipping and replay to a secondary database, you also don't have to do full online backups as often as was the case in Exchange 200x and earlier versions. Last but certainly not least, the failover process has been improved in several areas now that the new file share witness model has been introduced.

Backup Choices in a CCR Setup

When you deployed a cluster with Exchange 2003, the only option available when the stores were going to be backed up was to take a backup of the stores running on the production servers. With CCR (and LCR), you have the option of taking a backup of the database

copies on the passive node, thereby eliminating any heavy load on the active node, both in terms of I/O to the disk spindles as well as CPU usage.

Keep in mind, though, that you can only perform a backup on the passive node using VSS, which means that Windows Backup cannot be used for this purpose. Instead you need to use Microsoft Data Protection Manager version 2 (DPM v2) or a third-party backup application that supports VSS backups.

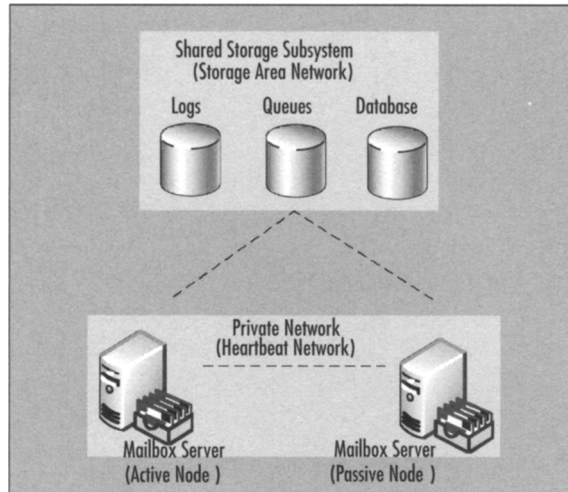
It's also worth mentioning that any backups performed via the passive node will be backups of the database copies, not the databases on the active node. So, you might wonder, what will happen to the transaction log files on the active node? When the backups have been performed on the passive node, all log files associated with the respective storage group on the active node will be truncated. In addition, the database header on the active node will be modified, and this will generate a log file that will be replicated to the passive node and then modify the database header on the passive node afterward.

To read more about how you back up the databases in Exchange 2007, see Chapter 10.

Managing a Single Copy Cluster-Based Setup

In addition to the CCR type of setup, Exchange 2007 supports the Single Copy Clusters (SCC) type, which, as mentioned in the beginning of the chapter, is more or less identical to the traditional active/passive clusters we know from previous versions of Exchange. This means that a SCC-based cluster only provides service failover and still has a single point of failure when it comes to the databases, unless a shared storage solution that provides redundancy in other means is used in the environment. An SCC-based cluster using a fault-tolerant SAN is just as good as a CCR-based cluster in terms of data availability, but such a solution is much more expensive than a CCR solution.

An SCC is basically a clustered mailbox server that consists of two or more servers (known as *nodes*) that share the same storage (for databases and log files). The shared storage subsystem is typically a SAN. In Figure 8.77 you can see what the architecture behind a typical SCC scenario looks like.

Figure 8.77 A Basic Single Copy Cluster Scenario**NOTE**

We know we mentioned it in the beginning of this chapter, but because it's important that you understand this concept, we repeat: Exchange Server 2007 no longer supports active/active clusters. Only active/passive clusters are supported in Exchange 2007.

The primary benefit of an SCC is that it provides high availability of server resources because one node takes over should the active node be taken offline or fail for some reason. In addition, you can apply hotfixes, service packs, and the like to the nodes without having any downtime of your mission-critical mailbox servers. However, bear in mind that an SCC is susceptible to failure of the shared storage subsystem. This means that no matter how many nodes are part of your cluster, you'll always have a single point of failure when you're using SCC opposite a CCR-based cluster, which, as we demonstrated, provides storage group failover via the new log file shipping and replay functionality.

Since most of you don't have the necessary hardware for a cluster, before you actually decide to deploy one in your environment, we thought it would be a clever idea to show you how to install an SCC in a Virtual Server 2005 R2 environment. Pretty much all the steps in this section can be used to install the SCC on real hardware, too.

SOME INDEPENDENT ADVICE

Some of you might wonder whether standby blusters are supported in Exchange 2007, just as they were in Exchange 2003. A *standby cluster* is a Windows cluster that matches a production Exchange cluster in terms of hardware and software configuration, including Windows and Exchange versions and any updates or hotfixes that have been applied. In addition, a standby cluster has the Exchange program files installed but has not yet been configured with any Exchange Virtual Servers (EVS). Lastly, a standby cluster can only be used when all Exchange Virtual Servers on the production cluster are offline.

So, is a standby cluster supported in Exchange 2007? The answer is no, but then it's really not that useful anymore, since Exchange 2007 gives us the ability to recover an Exchange 2007 cluster using the new *Exsetup /RecoverCMS* switch (which is similar to the */DisasterRecovery* switch we know from previous versions of Exchange). Even better, the */RecoverCMS* switch can be used to recover both Exchange 2007 CCR and SCC-based cluster setups. We'll take a closer look at the */RecoverCMS* switch in Chapter 10.

Prerequisites

To follow the steps throughout this section, you need the following:

- One physical machine running Virtual Server 2005 R2. Since this product is free to download from the Microsoft Web site, getting it shouldn't be a problem. You can download Virtual Server 2005 R2 from the following link: www.microsoft.com/windowsserversystem/virtualsever/software/default.msp.
- A Windows 2003 Active Directory forest with at least one domain controller (raised to 2000 or 2003 forest functional level).
- At least one existing Exchange 2007 Hub Transport/Client Access server already installed in the aforementioned forest.
- Two virtual guests running Windows 2003 R2 or Windows 2003 SP1 Enterprise Edition with at least 512MB RAM and two virtual NICs each—one for the Public network and one for the Private network (the heartbeat network). This means that you need to create an additional virtual network on the virtual host server; None (Guest Only) is sufficient for this network.

NOTE

To install a Exchange 2007 Single Copy Cluster, you also need to install the cluster hotfix mentioned in MS KB article 898790, which at the time of this writing can be requested by contacting Microsoft Product Support Services. Microsoft is working on making it public.

Configuring the Network Settings for each Network Interface

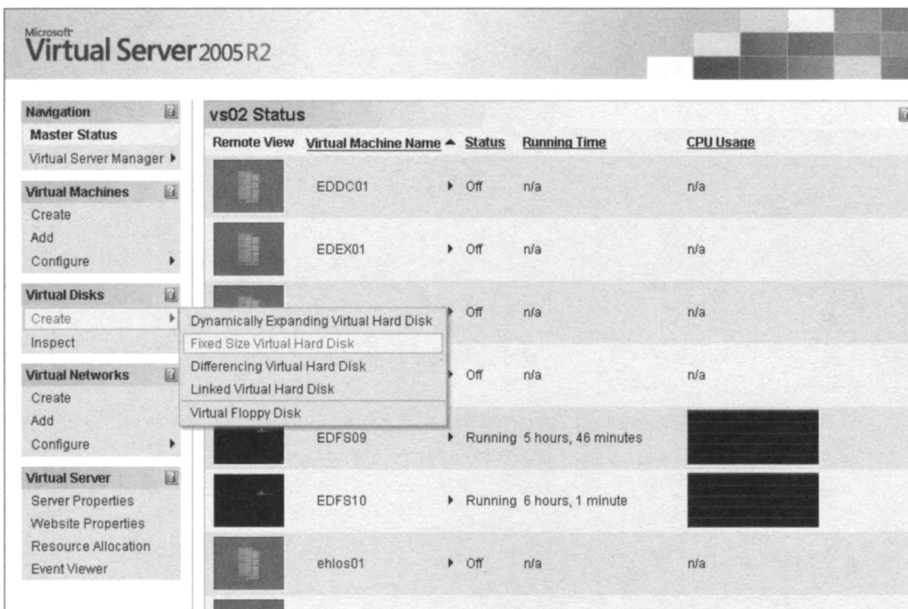
In this example, we'll create an SCC consisting of two active/passive clusters that will be part of the same Exchange organization as the CCR-based cluster we discussed previously in this chapter. This means that you will need to install two NICs in each node (which we recommend you call *public* and *private* so that you can see what belongs to which network) and then configure the private and public interfaces for each of the two nodes identically to the network interfaces we configured on the two nodes in the CCR-based cluster setup. The only difference would be the IP addresses, since using the same ones would result in IP conflicts, but everything from the binding order, WINS, DNS, and so on should be the same for each interface. So instead of going through all the steps again, refer back to the “Configuring the Network Interfaces for Each Node” subsection of the “Managing a Cluster Continuous Replication-Based Cluster Setup” section of this chapter.

Creating the Shared Cluster Disks

As those of you with cluster experience are aware, a Windows cluster requires a quorum cluster disk. This quorum disk is used to store cluster configuration database checkpoints and log files that help manage the cluster as well as maintain consistency. Since we're dealing with a virtual environment, we need to create this disk in the Virtual Server 2005 R2 Web console. This is done by following these steps:


1. Open the **Virtual Server Manager**, then click **Create | Fixed Size Virtual Hard Disk** under **Virtual Disks**, as shown in Figure 8.78.

Figure 8.78 Creating a Fixed-Size Virtual Hard Disk



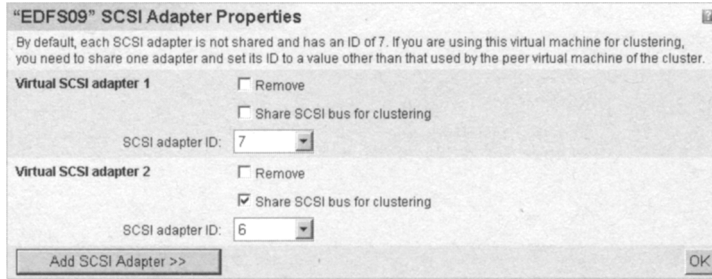
- Place the virtual hard disk file (.VHD) in the folder containing your two virtual Windows 2003 Servers, then set the size to **500MB** (or less if you're low on disk space). Then click **Create** (see Figure 8.79).

Figure 8.79 Specifying the Virtual Hard Disk Filename and Size



- We now need to add the virtual quorum disk to each of the two virtual Windows 2003 Servers. Let's add it to EDFS09 first. We do this by clicking **Master Status | EDFS09 | Edit Configuration**. Since this disk needs to be shared between the nodes, we need to click **SCSI Adapters**, then **Add SCSI Adapter** (see Figure 8.80). Under the new SCSI adapter, check **Share SCSI Bus for Clustering**, then set the SCSI adapter ID to **6** (or whatever SCSI adapter ID is unused in your environment). Click **OK**.

Figure 8.80 Adding an Additional Shared SCSI Adapter



4. We now need to make the new disk visible on each node, so click **Hard disks | Add disk**, then select **SCSI 1 ID 0** in the **Attachment** drop-down menu. Finally, specify the path to the virtual Quorum disk, which in this example is **E:\E2K7SCC\Shared Disks\Quorum.vhd**, as shown in Figure 8.81. Click **OK**.

Figure 8.81 Adding a Virtual Hard Disk



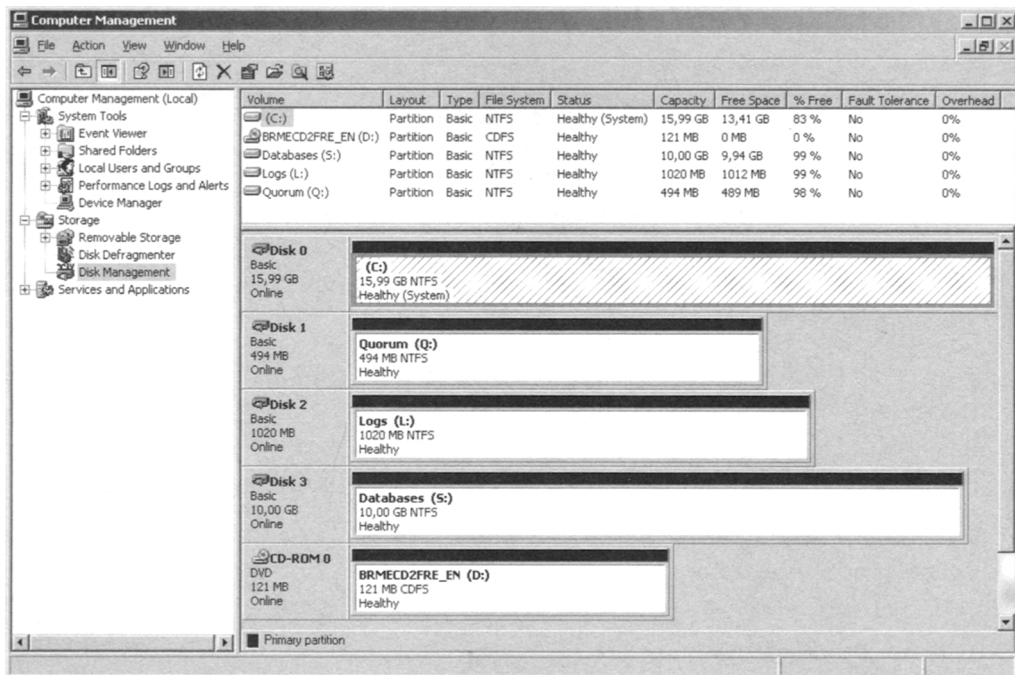
SOME INDEPENDENT ADVICE

If you're installing the SCC in a Virtual Server 2005 R2 environment like I do in this example, you need to create a virtual SCSI adapter for each disks you want to share between the nodes. Since you should place the databases and log files on share disks as well, I recommend you create two additional virtual fixed sized disks more, one called Logs.vhd and one called Databases.vhd. When these have been created you need to add two additional virtual SCSI adapters on each virtual guest, and since the two disks should be shared between the nodes this should have Share SCSI bus for clustering enabled and configured with SCSI adapter ID 6 like the adapter for the quorum disk

we already created. When you have done so, you will be able to add the two disks under Virtual Hard Disk Properties on each node respectively.

5. We now need to partition the Quorum disk in the Disk Management console on EDFS09, so start the virtual machine, log on using a domain admin account, click **Start | Run**, and type **Compmgmt.msc**. Under **Storage**, click **Disk Management** (see Figure 8.82). Click **Next** three times in the Initialize and Convert Disk Wizard that appears, then click **Finish**.
6. The detected disk now needs to be partitioned. To do so, right-click the unallocated space then select **New partition**.
7. Click **Next** three times and select the drive letter **Q** (for quorum), then click **Next** again. Use **NTFS** as the file system type and type **Quorum** in the Volume label field. To speed up the formatting process, it's a good idea to tick **Perform a quick format**.

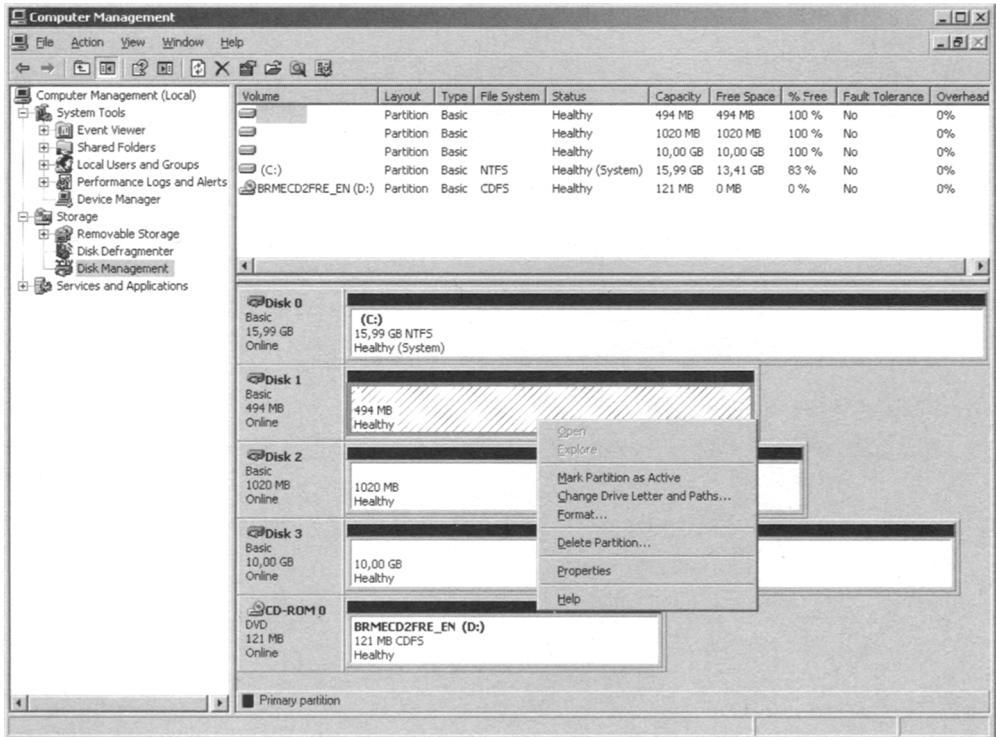
Figure 8.82 Partitioning the Shared Disks and Assigning Drive Letters



8. Now turn off EDFS09, then turn on EDFS10 and log on to the server with a domain admin account. Again, click **Start | Run** and type **Compmgmt.msc**.

Under **Storage**, click **Disk Management** **Mark the Quorum disk (disk 1) active** and assign it the drive letter **Q** (see Figure 8.83).

Figure 8.83 Allocating Drive Letters to the New Partitions on the Second Node



Now verify that you can access the Q: drive from Windows Explorer. Also try to create a test file on each server and make sure you can see it both ways.


Creating the Windows Server 2003 Cluster

We have reached the point where we can create the actual Windows 2003 cluster. To do so:

1. Turn off EDFS10, then log on to EDFS09 with a domain admin account. Now click **Start | Administrative Tools | Cluster Administrator**, then select **Create new cluster** in the drop-down box and click **OK**, then click **Next**.
2. If it's not already the case, specify the domain in which the two Windows 2003 Servers are members, then type the name of the cluster (in this case, **E2K7SCC**), then click **Next**.

3. If it's not already entered, type the name of the Windows 2003 Server, which will be the first node in the cluster (in this case, **EDFS09**), then click **Next**.
4. The Cluster Wizard will now determine the cluster configuration, and after a while you should get a check mark in each checked configuration step. We can now click **Next**.
5. Now enter an IP address that cluster management tools will use to connect to the cluster and click **Next**.
6. Enter the cluster service account and password, then click **Next**.
7. You now see a screen with the proposed cluster configuration. Click the **Quorum** button and make sure that the cluster configuration quorum is set to **Disk Q**, as shown in Figure 8.84. Then click **Next**.

Figure 8.84 Selecting the Resource Type Used for the Quorum Resource



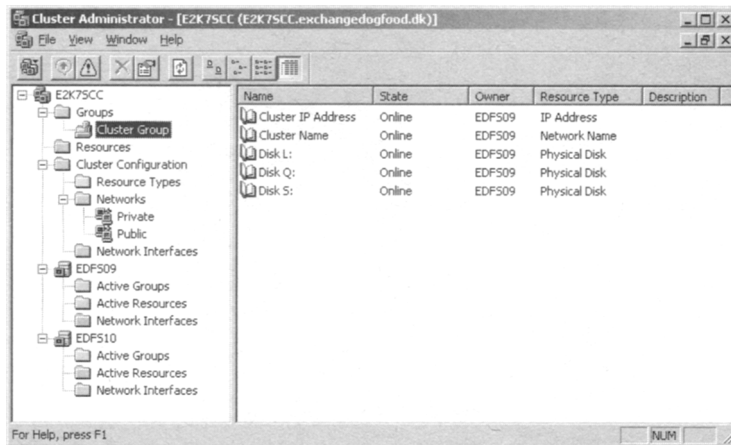
8. The cluster will now be created. Again, you need to wait for each step to complete, then click **Next** | **Finish**.

We have created the cluster itself, but since it consists of only one node, we'll need to add the other Windows server as well. To do so:

1. Turn on EDFS10 and log in with a domain admin account. Now click **Start** | **Administrative Tools** | **Cluster Administrator**. Select **Add nodes to cluster** in the drop-down menu, then specify the cluster name in the **Cluster or server name** box and click **OK**.
2. Click **Next** in the Add Nodes Wizard.
3. Type **EDFS10** (or whatever you named the second server), then click **Add** and click **Next**.
4. When the configuration has been analyzed, click **Next**.
5. Enter the password for the cluster service account (in this case, the administrator account), then click **Next**.
6. Verify that you want to add the node to the cluster with the configuration shown on the proposed cluster configuration page, then click **Next**.
7. After a short period, the node will be added to the cluster. If it's not, you might want to expand the respective task as well as view the log. If each task has com-

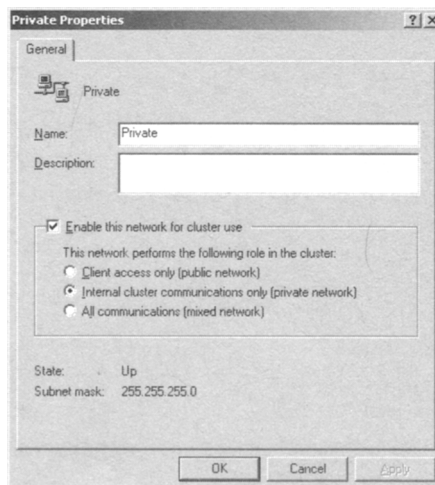
pleted successfully, click **Next | Finish** and verify that none of the nodes contains an error icon in the Cluster Administrator (see Figure 8.85).

Figure 8.85 The Cluster Administrator Will Cluster Resources Listed and Online



- There's one last thing you want to do before moving on, and that is to right-click and select **Properties** for the Private network in the left pane. Since the sole purpose of the Private network is to be used for communication between the internal cluster nodes, you should select **Internal cluster communications only (private network)**, as shown in Figure 8.86, then click **OK**. Do the same for the Public network, but set it to **Client access only (public network)**.

Figure 8.86 Changing the Cluster Role for the Private Network



We now have a fully operational two-node active/passive Windows cluster up and running.

Installing the Necessary Windows Components

Before we move on and try to install the Exchange Server 2007 Beta 2 bits, we need to make sure that the required Windows components have been installed. All types of Exchange Server 2007 installations (no matter what server role we're talking about) need the Microsoft .NET Framework 2.0 component installed.

NOTE

If you have installed Windows Server 2003 Enterprise Edition with Service Pack 1 on the nodes, you need to download the Microsoft .NET Framework Version 2.0 Redistributable Package (x86), since it's only a standard Windows component for Windows Server 2003 R2.

Since we're installing the Mailbox Server role in the cluster, we also need to install the following IIS 6.0 components:

- Enable network COM+ access
- Internet Information Services
- World Wide Web Service

NOTE

Remember to install these components on both cluster nodes.

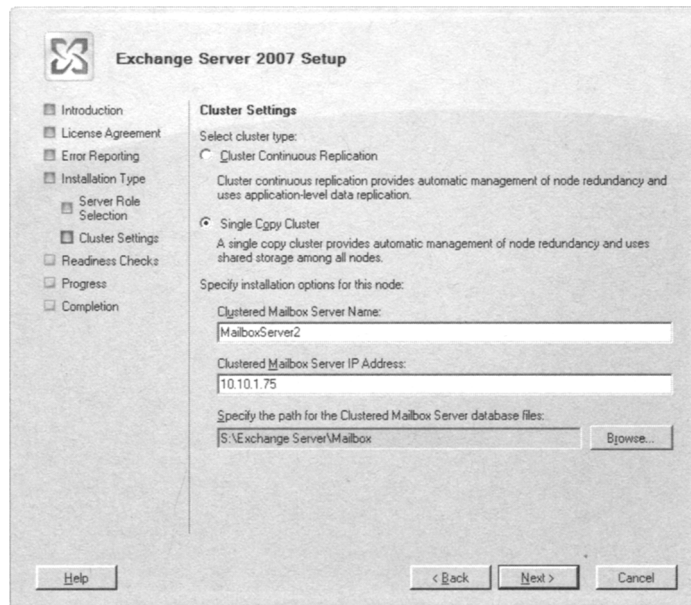
Installing Exchange Server 2007 on the Active Node

It's time to install the Exchange Server 2007 binaries on each node. Let's start with EDFS09. We'll do this using the GUI, so do the following:

1. Navigate to the network share or DVD media that contains the Exchange 2007 binaries, and double-click **Setup.exe**.

2. On the Exchange 2007 Setup splash screen, click **Step 4: Install Microsoft Exchange**. Then click **Next**. Accept the **License Agreement** and then click **Next** once again. Decide whether you want to enable error reporting or not (it's a good idea to enable this functionality since the Exchange Product Group will receive any obscure errors you should experience in your cluster setup), then click **Next**.
3. Now select **Custom Exchange Server Installation**, then click **Next**.
4. Check **Active Clustered Mailbox Role** and click **Next**.
5. Now select **Single Copy Cluster**, then specify a name for the mailbox server (the name you want your Outlook clients to connect to) and a unique IP address on your public network. Finally, specify the path for the clustered mailbox server database files (the virtual shared database disk you created earlier), then click **Next** (see Figure 8.87).

Figure 8.87 Specifying the Name and IP Address of the Clustered Mailbox Server



NOTE

To set the path for the clustered mailbox server database files, it's important that the cluster group containing the shared disks is owned by EDF509. The reason for this is that you aren't allowed to use the shared disks if the cluster group is currently owned by EDF510.

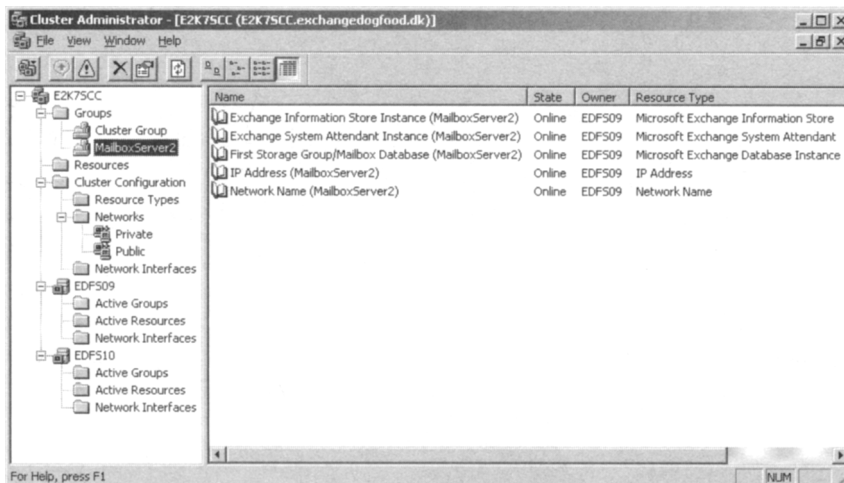
6. Let the readiness check complete, and if no issues are found, click **Next** to begin the installation.
7. The Exchange Server 2007 Installation Wizard will now copy the needed Exchange files, install and configure the Mailbox role, then create and configure the clustered mailbox server resources locally and create the object in Active Directory. When each step has been completed, clear the **Exit Setup and open Exchange System Manager** check box, then click **Finish**. We don't want to open the EMC just yet; we'll install Exchange on the second node first.
8. Log on to EDFS10 with a domain admin account and perform the same steps we did in installing Exchange Server 2007 on EDFS09. The only difference is that you should check **Passive Clustered Mailbox Role** instead of **Active Clustered Mailbox Role**.

When you have installed the Exchange Clustered Mailbox Role on the second node, we can move on to the next section, where we verify that the functionality of the clustered mailbox server works as expected.

Testing the Functionality of the Single Copy Cluster

It's time to verify that our Exchange 2007 clustered mailbox server is working as expected. Let's first open the Cluster Administrator and check whether the respective Exchange Resources have been created. If you take a look at Figure 8.88, it looks good; we have both nodes listed in the left pane and all Exchange resources have been created and are currently owned by EDFS09.

Figure 8.88 Listing All Exchange Cluster Resources in the Cluster Administrator



If you look closer at Figure 8.88, though, you can see that two cluster groups exist: one containing the cluster IP, name, and the shared disks, and one created by Exchange 2007 setup containing the Exchange Information Store, System Attendant, Storage Groups, and Database instances as well as the Exchange virtual server IP address and network name. WE recommend that you move all shared resources from the cluster group to the MailboxServer2 group (or whatever you called it); otherwise, you will have problems mounting the database when moving the clustered mailbox server from one node to the other (which we'll do in just a moment).

In addition, if you have assigned a shared disk specifically for the transaction log files, remember to change the path for these files. You can do so by selecting the respective storage group under **Server Configuration | Mailbox node** in the EMC, then click the **Move Storage Group** link in the Action pane. In the Move Storage Group Wizard, change the path for the log files to the **L:** drive or whatever drive you assigned them.

Now try to open the EMS by clicking **Start | All Programs | Microsoft Exchange Server 2007 | Exchange Management Shell** on one of the nodes, then type **Get-ClusteredMailboxServerStatus**. As you can see in Figure 8.89, the status of the clustered mailbox server is Online, and EDFS09 is currently the active node. This just keeps getting better and better, doesn't it?

Figure 8.89 Verifying That the Cluster Is Online



```
Machine: EDFS09 | Scope: exchangedogfood.dk
[PS] C:\>Get-ClusteredMailboxServerStatus

Identity                : MailboxServer2
ClusteredMailboxServerName : MAILBOXSERVER2.exchangedogfood.dk
State                   : Online
OperationalMachines     : <EDFS09 <Active, Quorum Owner>, EDFS10>
FailedResources         : <>
IsValid                 : True
ObjectState              : Unchanged

[PS] C:\>_
```

Now that we have verified that the clustered mailbox server is online, let's try to move the Exchange resources from node one to node two using the *Move-ClusteredMailboxServer* CMDlet. In the test environment we're using, we do so by issuing the command *Move-ClusteredMailboxServer -Identity:MailboxServer2 -TargetMachine:EDFS10 -MoveComment:"Testing the Move Clustered Mailbox functionality!"*.

You're then asked to confirm this action. Type **Yes**, then press **Enter** (see Figure 8.90). After a while the clustered mailbox resources will be moved to the second node.

Figure 8.90 Moving the Clustered Mailbox Resources to the Second Node

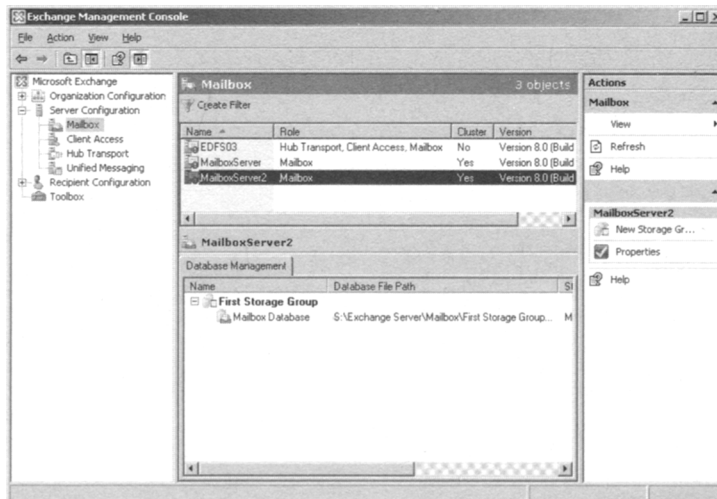

```
Machine: EDPS09 | Scope: exchangedogfood.dk
[PS] C:\>Move-ClusteredMailboxServer -Identity MailboxServer2 -TargetMachine: EDPS09 -MoveComment:"Testing the Move Clustered Mailbox funtionaliti!"

Confirm
Are you sure you want to perform this action?
Moving clustered mailbox server "MailboxServer2" to target node "EDPS09" with
move comment "Testing the Move Clustered Mailbox funtionaliti!"
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
(default is "Y")>Y
[PS] C:\>_
```

WARNING

Although it's possible to move the cluster resource group between the nodes using the Cluster Administrator console, you should always do so (just as is the case with CCR-based clusters) using the *Move-ClusteredMailboxServer* CMDlet because the Move Group task in the Cluster Administrator console isn't Exchange 2007 aware.

Let's also take a look at the clustered mailbox server in the EMC. To do so, click **Start | All Programs | Microsoft Exchange Server 2007 | Exchange Management Console**, then drill down to **Server Configuration | Mailbox**. Notice that the clustered mailbox server we named MailboxServer is listed in the Results pane and that it's recognized as a cluster server (see Figure 8.91). Also notice that the Mailbox Database for this server points to the S: drive, exactly as we specified during the installation of the Active Clustered Mailbox role.

Figure 8.91 Viewing the Clustered Mailbox Server in the Exchange Management Console

Summary

In this chapter we focused on the Mailbox server-related high-availability features included in Exchange Server 2007. First we took a look at how the Local Continuous Replication (LCR) feature works, and then we covered how it's implemented as well as managed. We then moved on to the new Cluster Continuous Replication (CCR) functionality, which makes it possible to deploy a mailbox server cluster, providing not only service availability but also database availability, which means that no single point of failure exists when using this type of cluster. We covered how to deploy a CCR-based cluster step by step as well as showed you how to manage it once deployed. Finally, we took a close look at the Single Copy Cluster (SCC) feature, which is similar to the traditional active/passive clusters we know from Exchange 2000 and 2003. We showed you the steps involved in deploying this type of cluster in a virtual server environment so that you can decide whether this is the type of cluster you want to use in your production environment.

Solutions Fast Track

Managing the Local Continuous Replication Feature

- ☑ The Exchange Product Group developed the Local Continuous Replication (LCR) technology to provide a native data availability solution that can be used to recover an Exchange database on an Exchange 2007 standalone server in a matter of a few minutes.
- ☑ Since LCR keeps an exact replica of the active copy of the storage group, the number of Exchange backups needed is also reduced drastically. But it's important to understand that LCR in no way eliminates traditional backups of the databases on your Exchange 2007 Mailbox servers; instead, it provides you with the option of taking weekly instead of daily backups, for example.
- ☑ As you can understand, LCR is an ideal solution for small or medium-sized organizations because the functionality allows rapid recovery from database issues and only requires an extra set of disks for the databases copies. LCR increases the availability of databases on an Exchange 2007 standalone server in an affordable way. For small shops that don't have a big fancy server with multiple sets of disks, it is possible to keep the LCR copy on an external USB disk.
- ☑ When disaster strikes and the database or log files in the active copy of the storage group become corrupted and shut down, you have the option of recovering database availability by switching to the LCR copy (the passive copy of the storage group).

- ☑ It's a recommended best practice to periodically verify the integrity of the passive storage group copy to make sure that neither the database copy nor any of the log files are corrupted. This is done by running a physical consistency check against both the database copy as well as the log files using Exchange Server Database Utilities (Eseutil.exe).
- ☑ When the Exchange 2007 Mailbox Server role is installed, setup adds two LCR-related performance objects to the Windows 2003 Performance Monitor.

Managing a Cluster Continuous Replication-Based Setup

- ☑ Exchange Server 2007 introduces a new high-availability feature called Cluster Continuous Replication (CCR). This feature combines the new Exchange Server 2007 log file shipping and replay mechanisms (known as continuous replication) with the features that are available in a more traditional two-node Windows 2003 server active/passive cluster setup.
- ☑ With CCR, the transaction logs generated on the active node are replicated to the information store on the passive node using log file shipping. These replicated log files are then posted into the database(s) on the passive node using the log file replay technology. This means that should the active node or a database on this node fail or for some other reason go offline, an automatic failover to the passive node will occur.
- ☑ A Majority Node Set (MNS) quorum with File Share Witness is a completely new type of quorum model that is made available by installing the update (MS KB article 921181) mentioned in this chapter. The update makes it possible to use a file share witness that is external to the cluster as an additional “vote” to determine the status of the cluster in a two-node MNS quorum cluster deployment, which is a requirement to use the CCR functionality in Exchange Server 2007.
- ☑ The Transport Dumpster is a new feature of the Exchange 2007 Hub Transport server that can submit recently delivered mail after an unscheduled outage. For an e-mail message to be able to be retained in the Transport Dumpster, at least one of the message recipients must have his or her mailbox located on a CCR-based mailbox cluster server, because the Transport Dumpster works only with mailboxes located on a CCR-based mailbox server cluster.
- ☑ Moving the Exchange resources from node one to node two should be done using the *Move-ClusteredMailboxServer* CMDlet. In the environment used in this chapter, we did so by issuing the cmdlet *Move-ClusteredMailboxServer -Identity:MailboxServer -TargetMachine:EDFS08 -MoveComment:"Verifying the Move Clustered Mailbox Server Functionality!"*.

- ☑ When we deployed a cluster with Exchange 2003, the only option available when the stores were going to be backed up was to take a backup of the stores running on the production servers. With CCR (and LCR), you have the option of taking a backup of the database copies on the passive node, thereby eliminating any heavy load on the active node related to both I/O to the disk spindles as well as CPU usage.

Managing a Single Copy Cluster-Based Setup

- ☑ Exchange 2007 supports the Single Copy Clusters (SCC) type, which is more or less identical to the traditional active/passive clusters we know from previous versions of Exchange. This means that a SCC-based cluster only provides service failover and still has a single point of failure when it comes to the databases, unless a shared storage solution that provides redundancy via other means is used in the environment. An SCC-based cluster using a fault-tolerant SAN is just as good as a CCR-based cluster in terms of data availability, but such a solution is much more expensive than a CCR solution.
- ☑ Exchange Server 2007 doesn't support active/active clusters anymore; only active/passive clusters are supported in Exchange 2007.
- ☑ Although it's possible to move the cluster resource group between the nodes using the Cluster Administrator console, you should always do so (as is the case with CCR-based clusters) using the *Move-ClusteredMailboxServer* CMDlet because the Move Group task in the Cluster Administrator console isn't Exchange 2007 aware.

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: Why would I want to deploy CCR instead of SCC?

A: Deploying CCR instead of SCC has several advantages. First, you no longer have a single point of failure regarding databases. Second, unlike SCC, CCR doesn't require a shared storage subsystem such as a SAN, because the nodes in a CCR don't share the same disks. Third, you have the option of spanning the CCR between two locations (although they must be on the same subnet, which means the subnet has to be stretched).

Q: You mentioned that it was possible to back up the passive copy of the databases in a CCR using a backup application with VSS support for Exchange databases. Is this also possible when we use LCR on a single Exchange 2007 box?

A: Yes, this is also supported on a single box with LCR enabled for one or more storage groups.

Q: How should I proceed when implementing storage design for a CCR-based setup?

A: To achieve storage resiliency, it is recommended that the passive copy be placed on a storage array that is completely isolated from the active copy's storage array. Isolating the arrays from one another also provides the flexibility to use a variety of storage solutions. If the storage solutions used by the active copy and the passive copy are isolated from each other, your storage solutions don't even need to be the same type or brand.

Q: Should I use an identical set of disks for the database copies in a CCR or LCR setup?

A: It's a best practice to size the active and passive storage solutions equivalently. The storage solution used by the passive copy should be sized in terms of both performance and capacity to handle the production load in the event of a failure.

- Q:** How many databases can I have in each storage group when I'm using either LCR or CCR?
- A:** You can only have one database in each storage group when you use either LCR or CCR. In addition, you cannot have more than one Public Folder database in the organization if you want to replicate a Public Folder database using continuous replication technology.
- Q:** Why would I want to use continuous replication technology in my Exchange environment?
- A:** Continuous replication provides service availability and service continuity for an Exchange 2007 mailbox server, without the cost and complexity of a shared storage cluster.