

Implementing DFS Replication

Mitch Tullock

In two previous articles we examined how to set up and configure DFS Namespaces in Windows Server 2003 R2. DFS Namespaces is the new name in R2 for one component of the Distributed File System (DFS) in Windows Server 2003 and lets you create a single virtual tree of shared folders that maps to actual (physical) shared folders located on multiple different file servers. Using DFS Namespaces lets you consolidate shared resources on your network into a single virtual resource that can be browsed and accessed easily by clients regardless of where the actual shared resources are located.

There's a second component of DFS in R2 however, and this is called DFS Replication. (In pre-R2 Windows Server 2003 both components were simply called DFS). DFS Replication is basically just a service that can be used to replicate files from one server to another so you can maintain multiple copies in different locations of a single file. Of course, in pre-R2 Windows Server 2003 there was already such a service named (quite obviously) the File Replication Service or FRS, and this was used for two different purposes:

- Replicating the contents of shared folders from one DFS tree to another.
- Replicating the contents of the SYSVOL folder between one domain controller and another.

The trouble with FRS is that it did a good job of the second task but a lousy job of the first. For example, say you had two DFS trees set up to replicate using FRS, and the shared folders in these trees contained very large files like video files. If you made a small change to one such file (maybe by editing out a single video frame) then FRS had to replicate the entire file from one tree to the other, which meant lots of bandwidth being used for replication. This could be a particular problem when the two trees were located in different sites connected by a slow WAN link—replication of files could easily eat up the entire bandwidth of the link! On top of that, FRS in pre-R2 was also a bit shaky and prone to failure, which meant sometimes DFS didn't work properly at all, which could be a pain to say the least.

Enter Remote Differential Compression

File replication has changed however as of Windows Server 2003 R2. In particular, there are now two different replication services built into the platform, namely:

- DFS Replication. This service is new and can be used not just for replicating DFS trees to provide fault tolerance and better performance, but even simply just replicating files between servers for any purpose. More about this new service in a moment.
- File Replication Service (FRS). This is the original replication service, first released with Windows Server 2000, though somewhat enhanced. In R2, the FRS is only used for replicating SYSVOL content on domain controllers and nothing else.

What's great is that the new DFS Replication service has a totally revamped replication engine that uses a new replication algorithm called Remote Differential Compression (RDC). This new algorithm replicates only the changes to files and not the files themselves, which means that DFS now works much better over slow WAN links than before. In addition, the new replication engine supports bandwidth throttling and replication scheduling, plus it operates on a multimaster replication model. The overall result is that DFS in Windows Server 2003 is much more solid in terms of its reliability and performance than before.

Implementing DFS Replication

Mitch Tullock

Setting Up DFS Replication

To illustrate how to configure and use DFS Replication, let's use the namespace we created in my previous article titled Implementing DFS Namespaces. In that scenario, the Accounting department used two file servers, BOX 162 and 163, with BOX162 having the shares Payables and Receivables while BOX163 had the shares Invoices, Inventory, and Reports. These shares were consolidated using DFS Namespaces into a virtual folder tree rooted in a shared folder named Accounting, which was created on BOX162 when we created the DFS namespace itself on that server. The following table describes how the folders in the DFS tree map to the shared folders on the two servers:

Shared Folder	Server	DFS Path
Accounting	BOX162	\\r2.local\Accounting (root)
Payables	BOX162	\\r2.local\Accounting\Ledger\Accounts Payable
Receivables	BOX162	\\r2.local\Accounting\Ledger\Accounts Receivable
Invoices	BOX163	\\r2.local\Accounting\Billing\Invoices
Inventory	BOX163	\\r2.local\Accounting\Catalog\Inventory
Reports	BOX163	\\r2.local\Accounting\Report Database

Note that the Invoices share is located on BOX163 and can be accessed using the DFS path \\r2.local\Accounting\Billing\Invoices where r2.local is the name of the Active Directory domain we are working with (this example uses a domain-based DFS namespace and the domain controller is BOX161). In other words, the Invoices share (C:\Invoices on BOX163) is the folder target associated with the \\r2.local\Accounting\Billing\Invoices folder in the \\r2.local\Accounting namespace.

Note: The terminology for DFS differs in R2 than in Windows 2000 Server and pre-R2 Windows Server 2003. In the old version of DFS, the Invoices share was called the link target and \\r2.local\Accounting\Billing\Invoices the link associated with that target, while \\r2.local\Accounting was called the root. So be sure you understand the new DFS terminology R2 uses before you start working with it. In other words, roots are now called namespaces, links are called folders, and link targets are folder targets!

Let's now use DFS Replication to replicate the contents of the Invoices share from BOX163 to BOX162. That way, should the share on BOX163 somehow become unavailable, users will still be able to access its content using BOX162. Of course, for true fault-tolerance you also need to replicate the namespace too—see my previous article titled Configuring DFS Namespaces for information on how to do this. What we'll do here though is simply create a second Invoices share on BOX162, replicate the contents of \\BOX163\Invoices to \\BOX162\Invoices, and add \\BOX162\Invoices to the list of folder targets for the \\r2.local\Accounting\Billing\Invoices folder in the namespace. That way if a client (for example XP-191) tries to access a file named Sample.doc found in \\r2.local\Accounting\Billing\Invoices on BOX163 but BOX163 is down, it can access the copy of the file on BOX162.

To accomplish this, the first thing you need to do is install the DFS Replication component if you haven't already done so. As mentioned in my previous articles, when you add or upgrade the File Server role you have the option of installing DFS Namespaces but leaving DFS Replication uninstalled, and that's what we did previously. So to add DFS Replication now, you can use Add or Remove Programs in Control Panel. Start this utility, select Add/Remove Windows Components, select Distributed File System, click Details and select DFS Replication Service (see Figure 1):

Implementing DFS Replication

Mitch Tullock

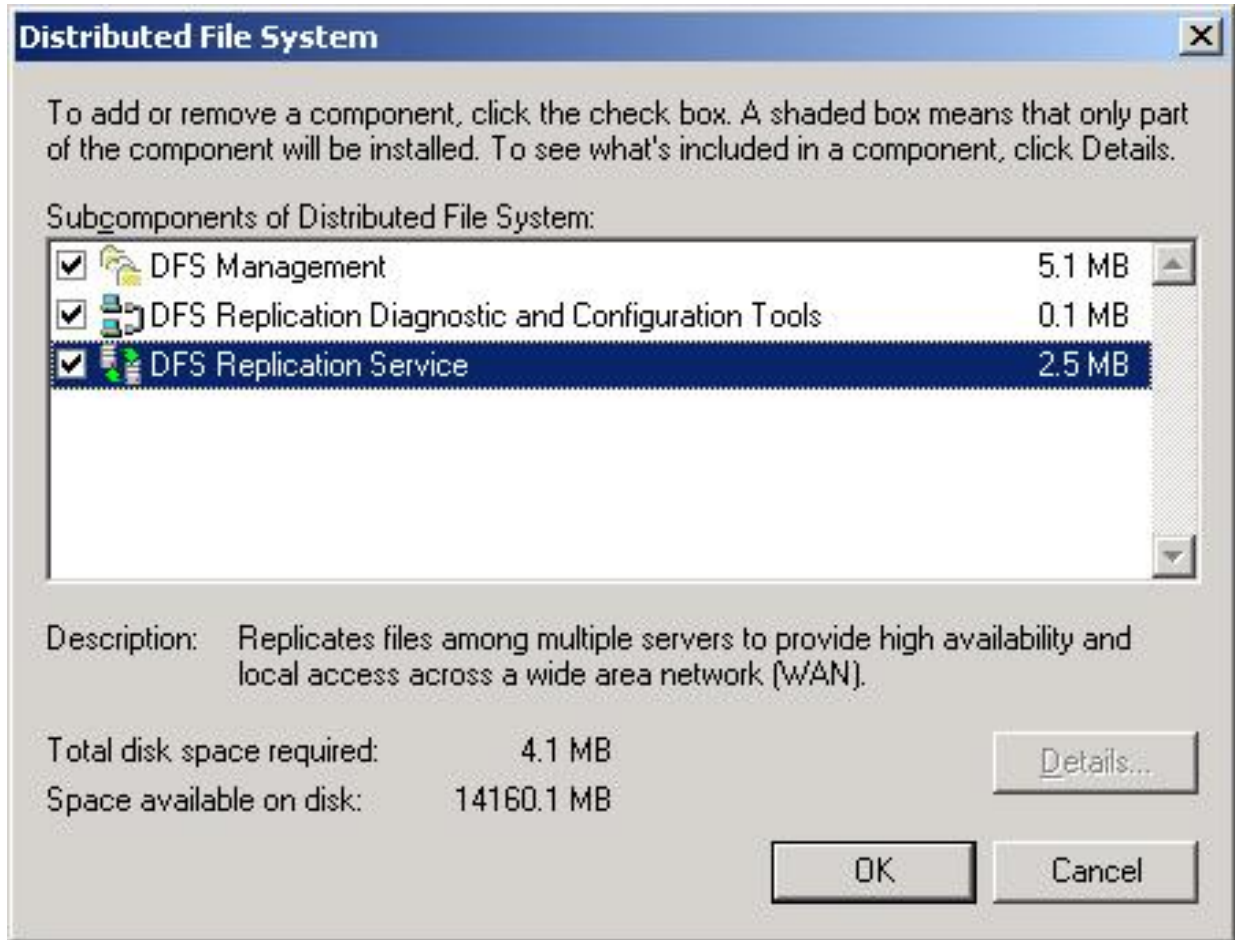


Figure 1
Installing DFS Replication

Insert Disk 2 of R2 when prompted or browse to the \COMPONENTS\R2 folder on your network distribution share to complete the installation of the component. Then create a new folder named C:\Invoices on BOX162 and share it with Full Control permission for Everyone (this choice does not mean the folder is not secure as NTFS permission are really used to secure resources, not shared folder permissions). Then be sure to install the DFS Replication component on BOX163 also since every file server that needs to participate in replicating DFS content must have the DFS Replication Service installed and running. You can again use Add or Remove Programs to do this, or use Manage Your Server.

Now let's add \\BOX162\Invoices as a second folder target for \\r2.local\Accounting\Billing\Invoices. Open the DFS Management console and select the following node in the console tree: DFS Management, Namespaces, \\r2.local\Accounting, Billing, Invoices. Note that there is currently only one folder target for this folder (Figure 2):

Implementing DFS Replication

Mitch Tullock

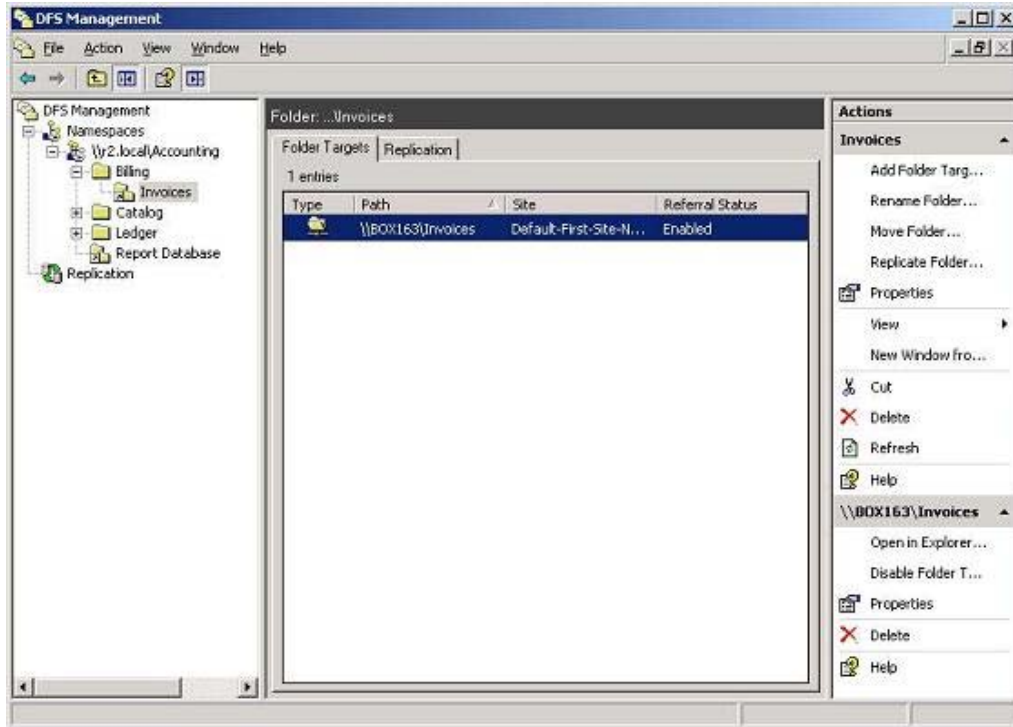


Figure 2
The Existing Folder Target

Verify from a client machine like XP-191 that a user can open the Sample.doc file by clicking Start, then Run, typing \\r2.local\Accounting\Billing\Invoices and double-clicking on the file. Now let's add a second folder target (\\BOX162\Invoices) for this folder as follows. Right-click the Invoices folder in the console tree and select Add Folder Target. Then specify the path to the new target as shown in Figure 3:



Figure 3
Adding a second folder target for the Invoices folder in the \\r2.local\Accounting namespace

Implementing DFS Replication

Mitch Tullock

Once the second target is added, you'll be prompted to create a replication group (see next Figure).

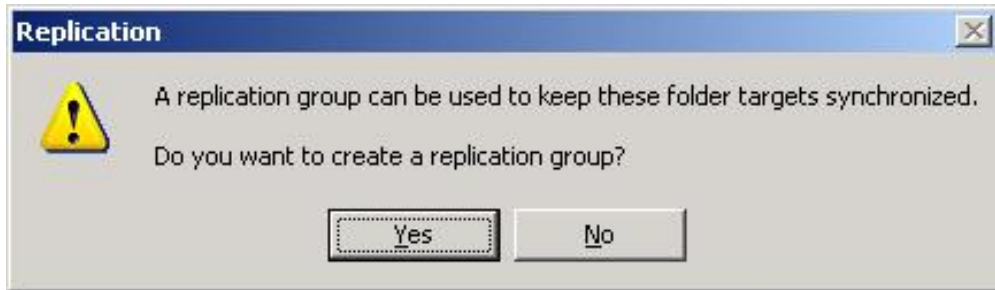


Figure 4

A replication group is a collection of file servers that participate in the replication of one or more folders in a namespace. In other words, if we want to replicate the contents of \\BOX163\\Invoices with \\BOX162\\Invoices, then BOX163 and BOX162 must first be added to a replication group. Replication groups can be created manually by right-clicking on the DFS Replication node in the DFS Management console, but it's easier here if we just create one on the fly by clicking Yes to this dialog box. This opens the Replicate Folder Wizard, an easy-to-use method for replicating DFS content on R2 file servers (see figure 5):

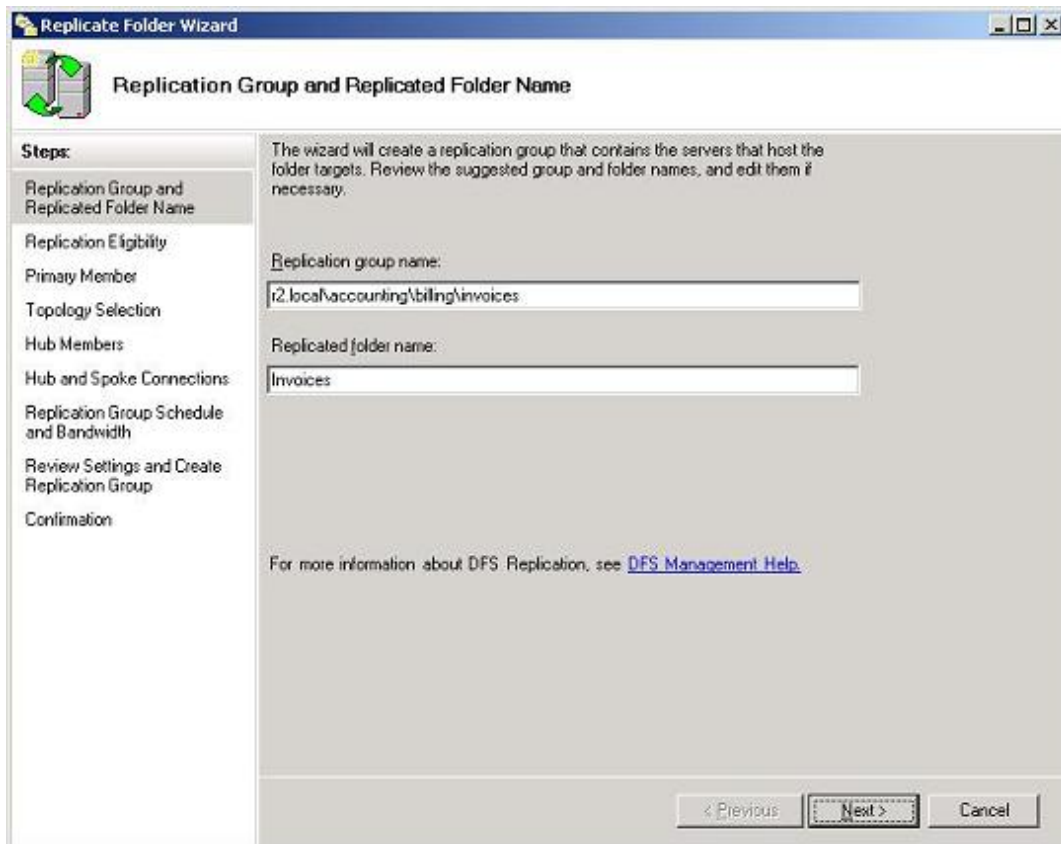


Figure 5
Wizard for Replicating Folders

Implementing DFS Replication

Mitch Tullock

We won't walk through all of the steps of this wizard here, instead we'll just summarize what the various steps of the wizard do:

- Replication Eligibility. Displays which folder targets can participate in replication for the selected folder (Invoices). Here the wizard displays \\BOX163\Invoices and \\BOX162\Invoices as expected.
- Primary Member. Makes sure the DFS Replication Service is started on the servers where the folder targets reside. One server is initially the primary member of the replication group, but once the group is established all succeeding replication is multimaster. We'll choose BOX163 as the primary member since the file Sample.doc resides in the Invoices share on that server (the Invoices share on BOX162 is initially empty).
- Topology Selection. Here you can choose full mesh, hub and spoke, or a custom topology you specify later. I'll talk more about replication topologies in my next article on WindowsNetworking.com. Selecting Full Mesh here jumps to the Replication Group Schedule and Bandwidth screen next.
- Replication Group Schedule and Bandwidth. Lets you replicate the content continuously up to a maximum specified bandwidth or define a schedule for replication (we'll choose the first option, continuous replication).

After reviewing your settings and confirming, two one-way replication connections are created between the two servers and initial replication is performed. Depending on how many other domain controllers you have and where they are located, initial replication may take awhile to complete.

Testing Replication

To view replication status for this folder, select the Replication tab as shown in this figure:

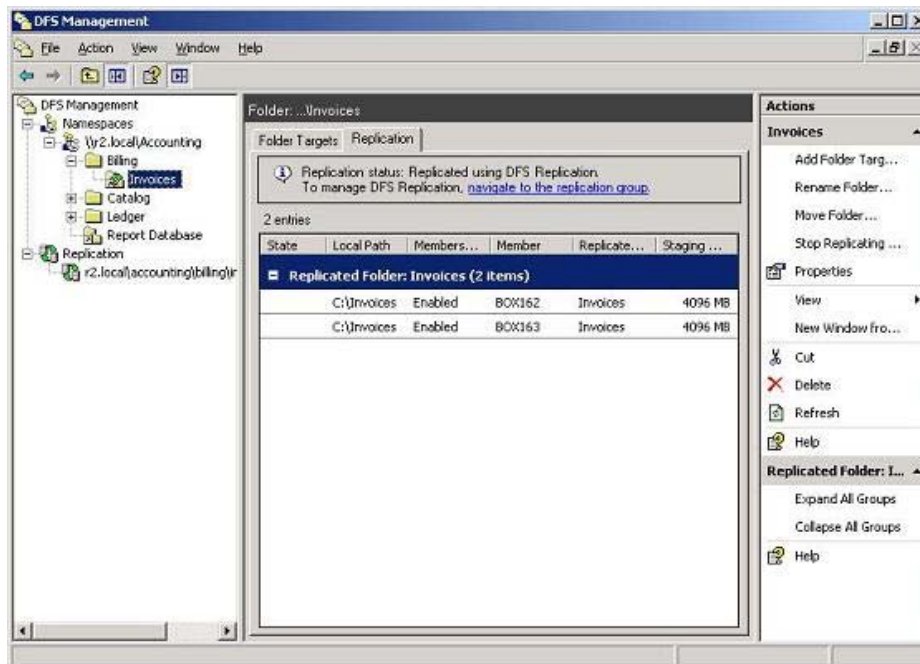


Figure 6
Replication Status of the Invoices Folder

Implementing DFS Replication

Mitch Tullock

A quick check of the C:\Invoices folder on BOX162 using Windows Explorer shows that the Sample.doc file has been replicated from BOX163 to BOX162 as expected. If you disable Local Area Connection on BOX163 and browse \\r2.local\Accounting\Billing\Invoices on client XP-191, the Sample.doc file is still available because the user is transparently referred to the next available folder target on BOX162. Finally, if a change is made to the Sample.doc file on BOX162 or BOX163, the change is replicated almost immediately to the folder target on the other box.