

# Configuring and Using DFS Replication

Mitch Tullock

In my previous article Implementing DFS Replication, we looked at the new DFS Replication component of Windows Server 2003 R2 and how it can be combined with the DFS Namespaces component. We learned about the improvements to file replication in R2 that make it fast and reliable even over slow WAN connections, and we walked through an example of how to use DFS Replication to provide fault-tolerance for folder targets in a namespace. This article continues the previous one by first examining some of the configuration options for DFS Replication and then looking at two scenarios where you might use DFS Replication in enterprise environments.

## Configuring DFS Replication

In the previous article we used DFS Replication to provide fault-tolerance for \\r2.local\Accounting\Billing\Invoices, a folder within the \\r2.local\Accounting namespace in the r2.local domain. This folder originally had only one folder target, the shared folder \\BOX163\Invoices, and to make this folder redundant we had to do two things:

- Add a second folder target, namely the shared folder \\BOX162\Invoices, so that if the first folder target was unavailable, client machines could obtain a referral from the namespace server so they could connect to the second target instead.
- Replicate the contents of the first folder target (\\BOX163\Invoices) to the second folder target (\\BOX162\Invoices) and keep the contents of these two shared folders in sync so that if one of them becomes unavailable, clients can still access the files stored in this namespace folder.

Using the Replicate Folder Wizard, we saw how easy this task is to perform. Let's spend a moment looking more closely at the results of the previous walkthrough and the configuration options that are available to us. The first figure shows the two folder targets created for the \\r2.local\Accounting\Billing\Invoices of our namespace:

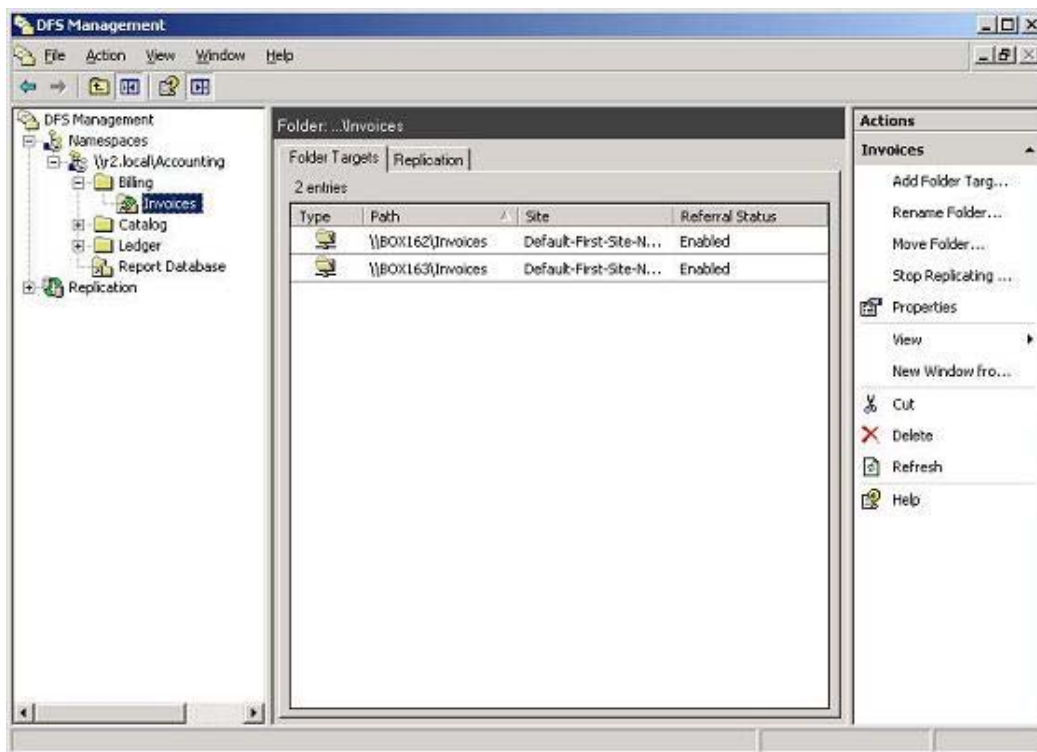
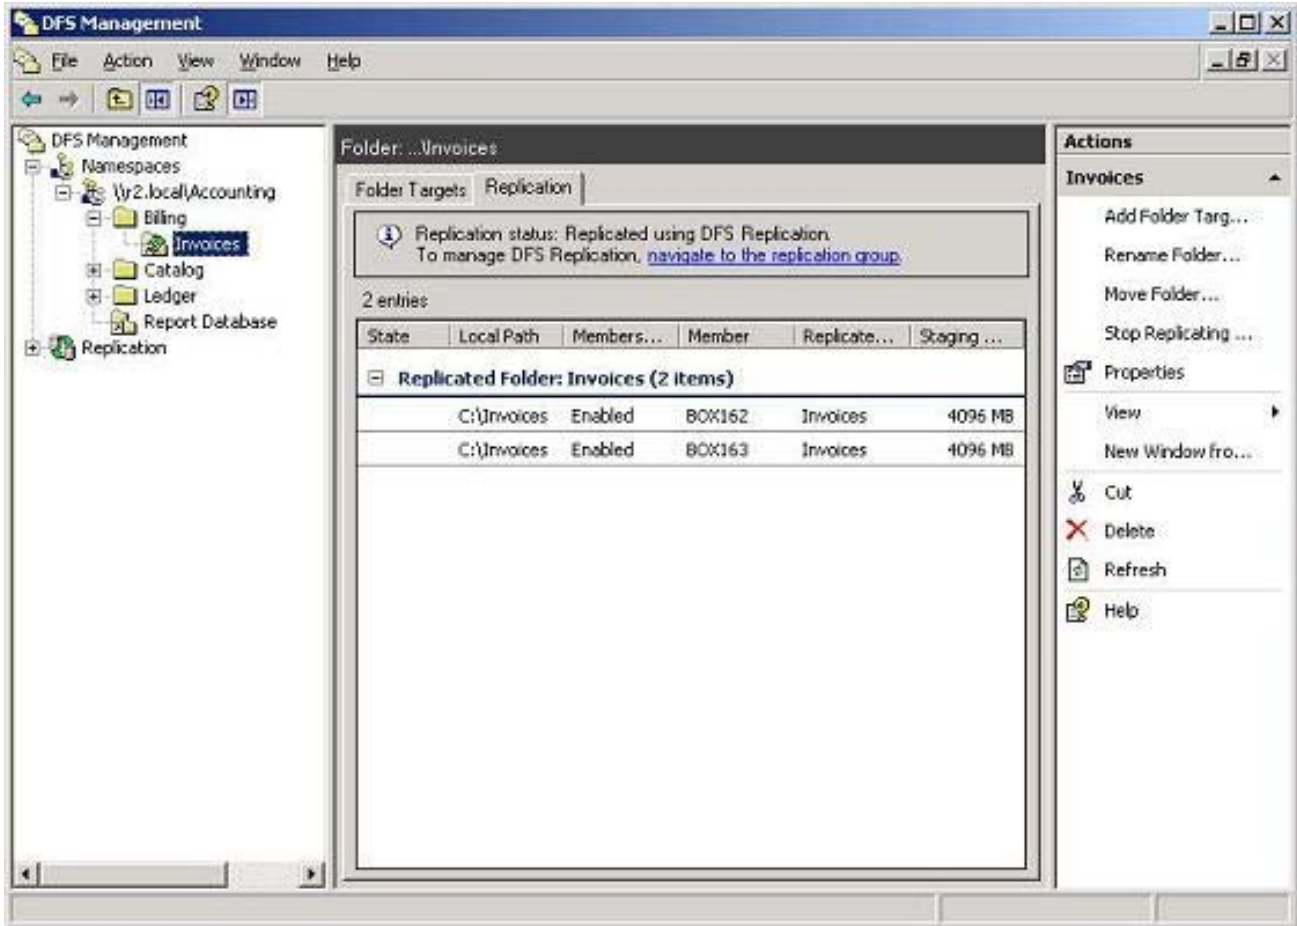


Figure 1: The Invoices Folder With Its Two Folder Targets

# Configuring and Using DFS Replication

## Mitch Tullock

Clicking the Add Folder Target link in the Actions pane at the right (or right-clicking on the Invoices folder in the console tree and selecting Add Folder Target) lets us easily add additional folder targets for this folder if greater redundancy is needed. Let's switch to the Replication tab in the middle pane and see what options are available:

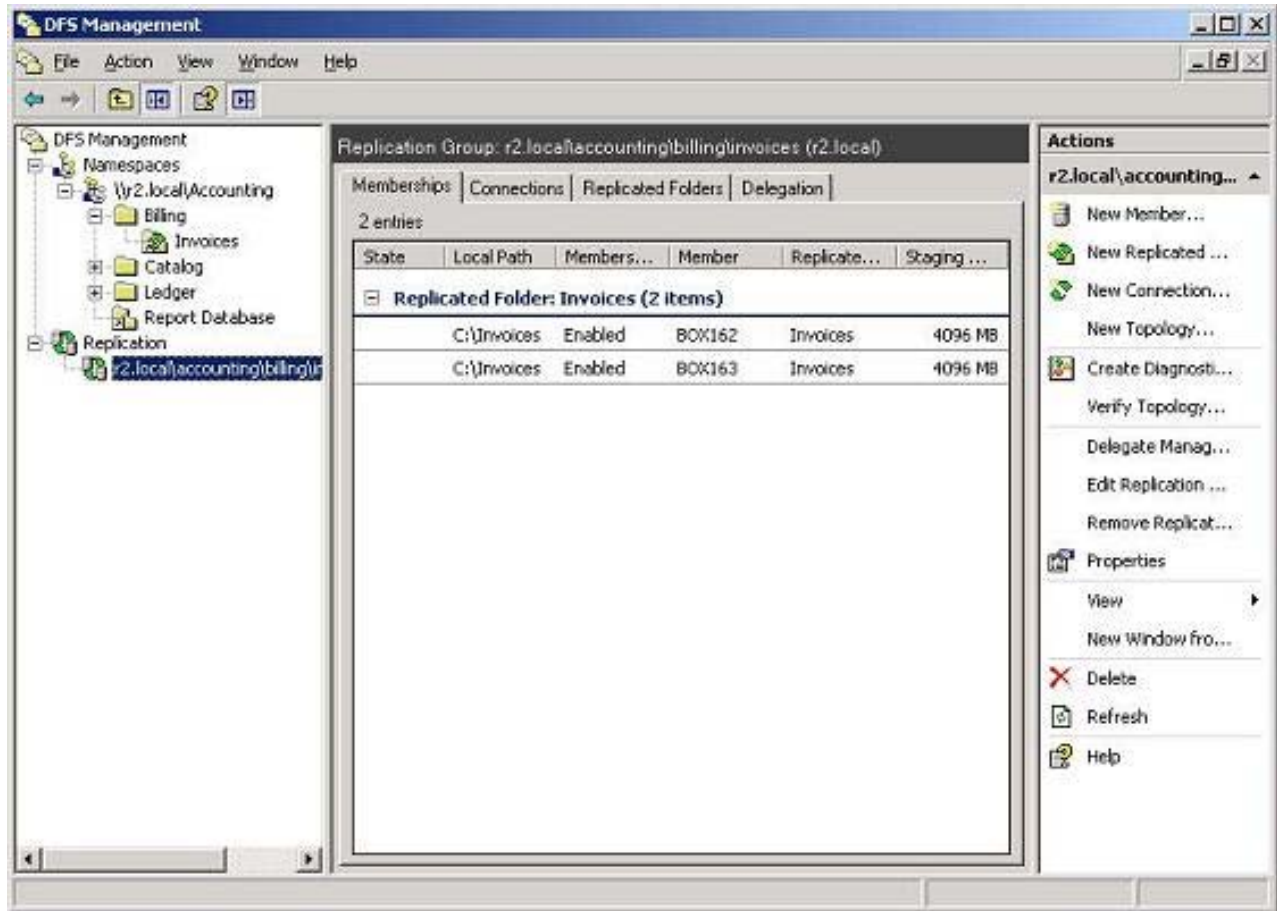


**Figure 2: Replication Status Of Invoices Folder**

Here we see the two folder targets that replicate with each other to provide fault-tolerance for the Invoices folder. Right-clicking on either of these folder targets and selecting Properties will bring up read-only properties for these targets. Clicking the Stop Replicating link in the Action pane will stop replication between these folder targets and delete the replication group that the Replicate Folder Wizard created earlier for these targets. To further configure replication settings, we first need to select the previously created replication group in the console tree, which can be found under the Replication node as shown here:

# Configuring and Using DFS Replication

## Mitch Tullock



**Figure 3: Configuring The Replication Group Created Previously**

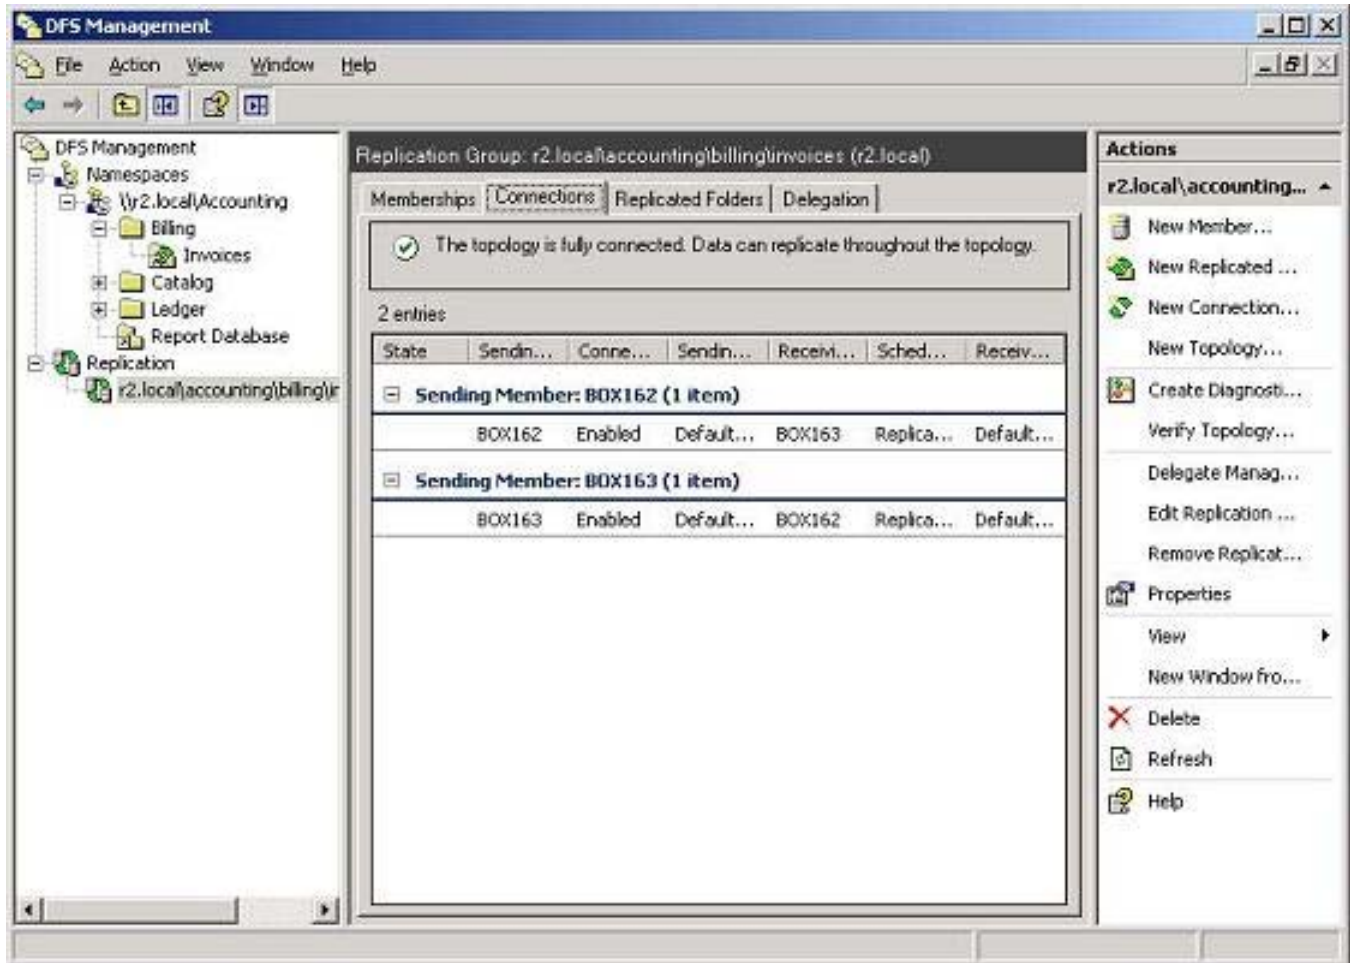
The Memberships tab lets us add additional members to this replication group. A member is a server that participates in replication of folders, so for example if we want to create a third folder target for Invoices and point this folder target to a share on a server named BOX164, we first need to make BOX164 a member of the replication group for Invoices and then add the folder target to the share on that box. DFS Replication makes this easy—click the New Member to start a wizard that guides you through the process of adding a new server to the replication group.

Alternatively, let's say you want to make another folder in the namespace redundant. For example, the folder `\\r2.local\Accounting\Catalog\Inventory` currently has a single folder target, namely `\\BOX163\Inventory`. If you wanted to add a second folder target for this folder, say `\\BOX162\Stuff`, simply click the New Replicated Folder link in the Action pane to start a wizard that guides you through the process. Once this is done, the new replicated folder will be displayed on the Replicated Folders tab.

The Connections tab displays the two one-way connections created by DFS Replication between `\\BOX163\Invoices` and `\\BOX162\Invoices`, the two folder targets for the Invoices folder:

# Configuring and Using DFS Replication

Mitch Tullock

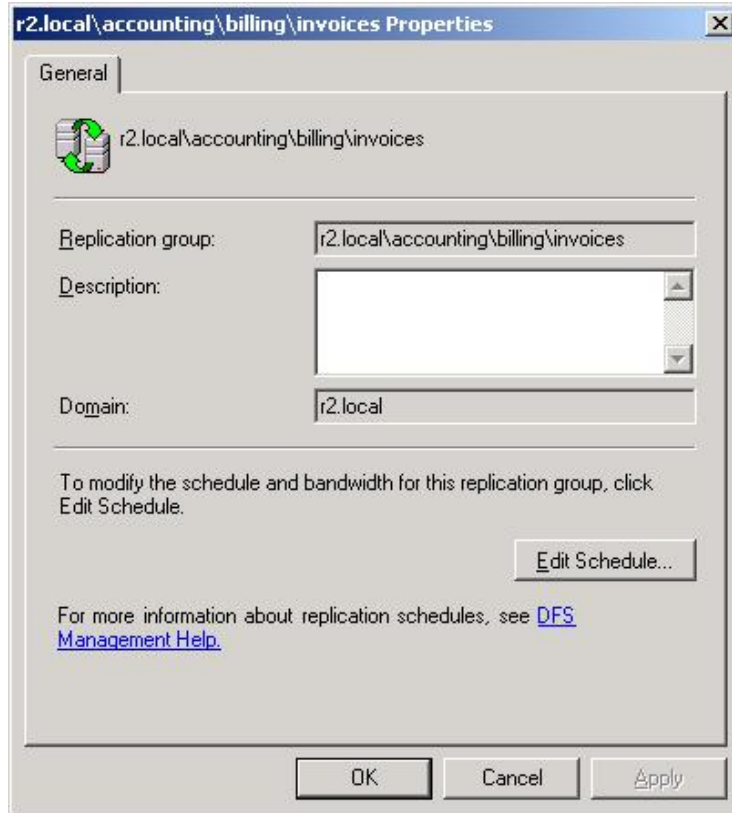


**Figure 4: The Two One-Way Connections For The Replication Group**

We'll see the advantage of having one-way connections in a moment. But first, recall that when we ran the Replicate Folder Wizard we chose to have files replicated continuously between the two targets for Invoices. What if we want to change this and schedule when replication occurs instead? This might be useful for example if the two members of the replication group are in different sites connected by a WAN link that is so highly utilized during the day that you would prefer to have replication occur only at night. You can configure your replication schedule in two ways. First, you can define a schedule for the entire replication group and all folders that it replicates by right-clicking the replication group under Replication in the console tree at the left and selecting Properties:

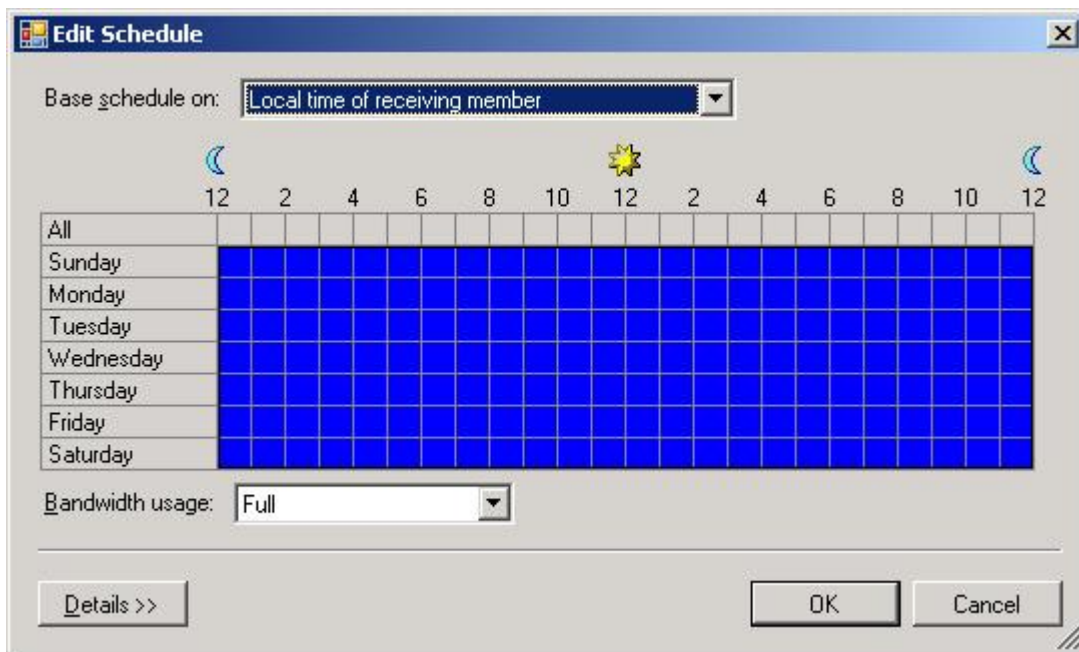
# Configuring and Using DFS Replication

Mitch Tullock



**Figure 5: The Properties Sheet For The Replication Group**

Then click the Edit Schedule button and edit the schedule for the replication group as desired.



**Figure 6: Editing The Schedule For A Replication Group**

# Configuring and Using DFS Replication

Mitch Tullock

Note that these schedules operate according to the schedule of the receiving member of a replication action over a replication connection i.e. DFS Replication is a pull process not a push process. You can also limit how much absolute bandwidth is used by the replication process, but be aware that this is only an average bandwidth and peak bandwidth may exceed this value occasionally. The second way of scheduling replication is to configure it on a per-connection basis, which can be done by opening the properties sheet of the appropriate connection instead. There are other DFS Replication settings we could look at, but let's move on now and talk about two enterprise scenarios where DFS Replication can be especially useful, namely branch office backups and publishing content.

## Using DFS Replication for Branch Office Backups

When we walked through the Replicate Folder Wizard in the previous article, we saw that two pre-defined topologies could be selected:

- Full Mesh. Every member of the replication group replicates with every other member of the group.
- Hub and Spoke. Every hub member replicates with the hub member, and if desired you can add a second hub member for fault tolerance (the two hub members replicate with each other).

The full mesh topology is useful mainly in large LAN environments where all subnets have high speed connectivity and you are using DFS Namespaces together with DFS Replication to provide fault-tolerant shared file resources to users. Note that Microsoft recommends that replication groups with full mesh topologies have no more than 10 members in them.

The hub and spoke topology is more interesting and can have a particular use for enterprises that have large headquarters where the company's permanent IT staff are located and multiple small branch offices with little or no on-site IT staff present. For such branch offices, one big concern is ensuring that reliable backups are done. Performing backups however costs money (tape drives and tapes) and effort (rotation and archiving of tapes) at such branch offices, and using DFS Replication you can work around this as follows:

- Create a replication group that has the hub member at headquarters and an additional spoke member at each branch office.
- Create a folder target on the hub member for each branch office member, where the branch office folder target points to the shared file resources that need to be backed up at that branch office.
- Configure DFS Replication to replicate between the hub and each branch office at 1 a.m. each night.
- Finally, disable the connections that go from the hub to each branch office so that replication occurs only from the branch offices to the hub and not vice versa.

Now all you have to do is configure your tape backup system at headquarters to back up the folder targets on the hub each night at 3 a.m. and you've automatically got each branch office's file resources backed up at headquarters without cost or effort to the branch offices themselves.

## Using DFS Replication to Publish Content

The hub and spoke topology can also be useful for publishing content from headquarters to branch offices in enterprise environments similar to the one described above. Just create a replication group with two hub members at headquarters (the second hub member is used to ensure fault-tolerance in case the first one fails) and with one spoke member at each branch office. DFS Replication will then automatically take files created at headquarters and replicate them out to branch offices. This is a

# Configuring and Using DFS Replication

Mitch Tullock

great way of disseminating information throughout an enterprise, and again you can disable inbound connections to the hubs if you only need to publish content to the branches and not the other way.

## Conclusion

DFS has undergone a significant overhaul in Windows Server 2003 R2. With the new DFS Namespaces and DFS Replication components, enterprises now have the tools they need to publish content, simplify remote backups, and create and manage virtual trees of shared folders simply and reliably.