

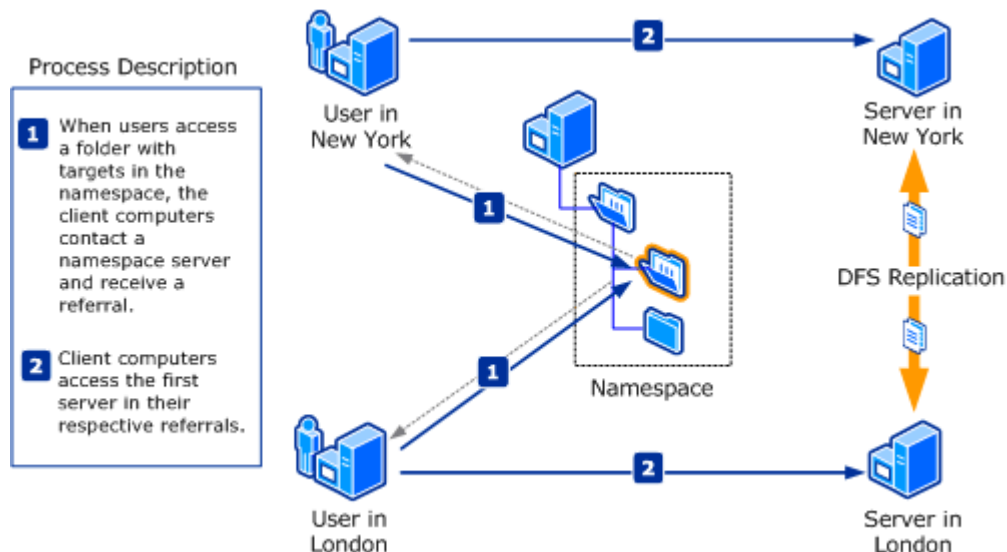
# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

One of the goals of information technology (IT) groups in medium and large organizations is to manage file servers and their resources efficiently while keeping them available and secure for users. As organizations expand to include more users and servers—whether they are located in one site or in geographically distributed sites—administrators find it increasingly difficult to keep users connected to the files they need. On one hand, storing files on distributed servers makes files available to more users and decreases latency and bandwidth use when the servers are located near users. On the other hand, as the number of distributed servers increases, users have difficulty locating files they need, and operational costs increase.

Administrators who manage these distributed, remote servers need a solution that helps them limit network traffic over slow WAN connections, ensure the availability of files during WAN outages or server failures, and ensure that branch servers are backed up correctly. The Distributed File System solution in the Microsoft® Windows Server™ 2003 R2 operating system helps administrators address these challenges by providing two technologies, DFS Namespaces and DFS Replication, which, when used together, offer simplified, fault-tolerant access to files, load sharing, and WAN-friendly replication.

- DFS Replication is a new state-based, multimaster replication engine that supports replication scheduling and bandwidth throttling. DFS Replication uses a new compression protocol called Remote Differential Compression (RDC), which can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and re-arrangements of data in files, thereby enabling DFS Replication to replicate only the changes when files are updated. Additionally, a function of RDC called *cross-file RDC* can help reduce the amount of bandwidth required to replicate new files.
- DFS Namespaces, formerly known as Distributed File System, allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a *namespace*. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration.

The following figure illustrates how DFS Namespaces and DFS Replication work together. The processes marked 1 and 2 are described in more detail following the figure.



As the figure shows, when a user attempts to access a folder in the namespace (1), the client computer contacts a namespace server. The namespace server sends the client computer a referral that contains a list of servers that host the shared folders (called folder targets) associated with the folder. The client computer caches the referral and then contacts

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

the first server in the referral (2), typically a server in the client's own site unless no same-site servers exist or the administrator configures target priority.

The highlighted folder in the figure shows that it is hosted by shared folders on two servers, one in New York and one in London, to provide users in those sites with fast, reliable access to files. The shared folders are kept synchronized by DFS Replication. The fact that multiple servers host the folder is transparent to users, who see only a single folder in the namespace. If one of the servers becomes unavailable, the client computer fails over to the remaining server.

## Key Scenarios

Although DFS Namespaces and DFS Replication are two separate technologies, when used together they provide solutions for a number of key scenarios in medium and large organizations.

### Data collection

The data collection scenario helps address the need to eliminate the use of tape backup in branch offices and significantly reduce operational costs for branch offices. To accomplish this, data is replicated from a server in a branch office to a server in a hub office or data center. Administrators at the hub office can use backup software to back up the branch server's data from a hub server, eliminating the often error-prone process of having non-IT administrators performing the backups at branch offices that are not staffed by trained IT personnel. Centralizing backups at the hub office also allows organizations to consolidate backup hardware and related management tasks.

Thanks to RDC, DFS Replication replicates only the differences (or changes) between the two servers; as a result, bandwidth used during replication is minimized, an important consideration for branch offices that use low-bandwidth WAN connections to the hub office. In addition, replication schedules and bandwidth throttling can be used to set up replication windows during off-hours and regulate the amount of bandwidth used during replication, providing administrators more control over WAN traffic.

When DFS Replication is used in conjunction with DFS Namespaces, administrators can configure a namespace so that branch clients always connect to the branch server. If the branch server becomes unavailable, branch clients fail over to the hub server. And, using a DFS Namespaces enhancement known as *client fallback*, branch clients running a client fallback hotfix will fail back to the branch server after it is restored.

### Data distribution

DFS Namespaces and DFS Replication can be used to publish documents, software, and line-of-business data to users throughout an organization. Although DFS Replication alone is sufficient to distribute data, using DFS Namespaces allows administrators to configure the namespace so that a folder in the namespace is hosted by multiple servers, which increases data availability and distributes the client load across servers. When browsing the namespace, users see a single folder and are not aware the folder is hosted by multiple servers. When a user opens the folder, the client computer is automatically referred to a server in its own site; if no same-site servers are available, administrators can configure the namespace to refer the client to a server that has the lowest connection cost as defined in the Active Directory® directory service. And, using a DFS Namespaces enhancement known as *target priority*, administrators can specify the priority of servers so that a specific server is always placed first or last in the list of servers (known as a referral) the client receives when accessing a folder in the namespace.

### Sharing files across branch offices

In large organizations that have branch offices located around the world, users in one branch office often need access to shared folders and files stored at another branch or at the hub office or data center. Although in some cases the need is collaborative, in other cases users simply need access to those files. If low latency and data availability during WAN

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

outages is worth the tradeoff of network traffic and disk space used in the branch office, administrators can use DFS Replication to replicate files between branch servers, providing users fast access to files in their respective branch offices. Having files in multiple branch offices also benefits users who travel from one branch office to another. The users can make changes to their files in one branch office, and those changes are replicated back to the user's branch office.

When the shared folders in a branch office are defined as folders in a namespace, users access them by browsing what appears to be a single server, giving the users a consistent place to begin browsing for files. For example, a user who travels from one branch office to another will access what appears to be the same folder in the namespace, but the user is actually referred to the closest server that contains the data. If that data is replicated using DFS Replication, the user's changes are replicated back to the user's branch office.

Collaboration scenarios using DFS Replication are recommended only if users can tolerate some inconsistency of the files as the changes are replicated to other branch servers. In addition, because DFS Replication uses a heuristic of last writer wins for changes made to the same file on multiple servers, this scenario works best when a file is updated from a single server or by a single user. Setting shared folder permissions can help ensure that conflicting changes don't occur. For collaboration scenarios that require file locking, we recommend using Microsoft® Windows® SharePoint® Services.

## Note

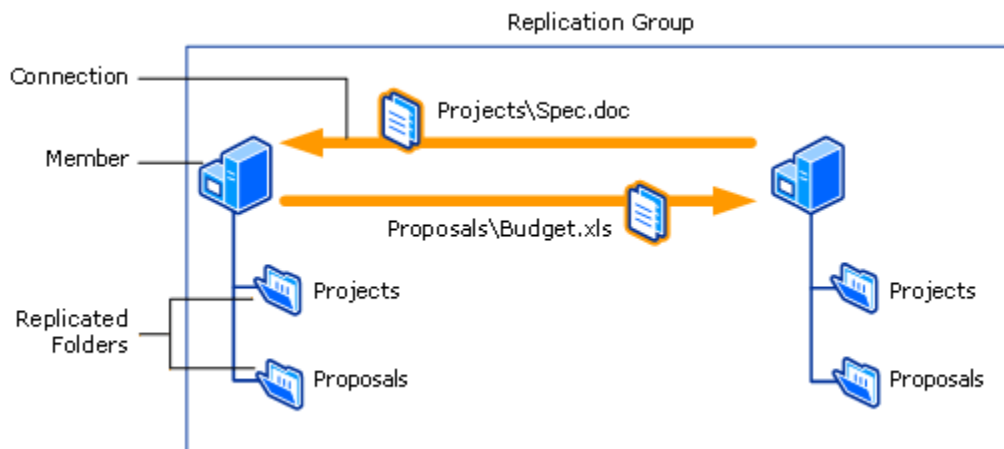
Because DFS Replication only replicates a file after it is closed, DFS Replication is not recommended for replicating database files or any files that are held open for long periods of time.

## DFS Replication Benefits in Windows Server 2003 R2

DFS Replication is a new replication engine that provides substantial improvements over the File Replication service (FRS). The benefits of using DFS Replication are described below.

### Simplified process for replicating discrete folders to the same set of servers

The process of setting up replicated folders is simplified in Windows Server 2003 R2 by the introduction of *replication groups* and *replicated folders*, which are illustrated in the following figure.



As this figure shows, a *replication group* is a set of servers, known as *members*, that participates in the replication of one or more replicated folders. A *replicated folder* is a folder that is kept synchronized on each member. In the previous figure,

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

there are two replicated folders, Projects and Proposals. As data changes in each replicated folder, the changes are replicated across *connections* between the members. The connections between all members form the replication *topology*.

Creating multiple replicated folders in a single replication group simplifies the process of deploying replicated folders because the topology, schedule, and bandwidth throttling for the replication group are applied to each replicated folder. Each replicated folder also has its own settings, such as file and subfolder filters, so that administrators can filter out different files and subfolders for each replicated folder. To deploy additional replicated folders, administrators use a brief wizard to define the local path and permissions for each new replicated folder.

The replicated folders stored on each member can be located on different volumes in the member, and the replicated folders do not need to be shared folders or part of a namespace, though the DFS Management snap-in makes it easy to share replicated folders and optionally publish them in an existing namespace.

## Differential replication of changes to files

DFS Replication uses RDC to replicate only the differences (or changes) between the members. This allows branch offices with slow WAN connections to participate in replication using minimal bandwidth.

RDC is especially efficient when small changes are made to large files. For example, a change that is made to a 2-megabyte (MB) PowerPoint® presentation can result in only 60 kilobytes (KB) being sent across the network, a 97 percent savings in bytes transferred. A test was run on a mix of 780 Office files (.doc, .ppt, and .xls) replicating from a source server to a target server using DFS Replication with RDC. The goal was to determine how effective RDC is for Office files on which typical edits are made. The target server had version *x* of the files and the source server had version *x+*, and the two versions differed with significant edits. The percent savings in bytes transferred was on average 50 percent and significantly better for large files.

Also, by using RDC and compression, DFS Replication provides bandwidth savings for distributing software files. For example, the IT department at Microsoft uses DFS Replication to replicate software files to Microsoft offices in several countries/regions. Although bandwidth used varies over time, in one typical case, a 63.14 percent reduction in bandwidth was measured in network traffic, specifically, 16.22 gigabytes (GB) replicated instead of 44.0 GB.

### Note

RDC is not used on files smaller than 64 KB and might not be beneficial on high-speed LANs where network bandwidth is not contended. RDC can be disabled on a per-connection basis using the DFS Management snap-in.

## Reduction of bandwidth used for replicating new files

Another function of RDC, known as cross-file RDC, uses a heuristic to identify files that are similar to the file that needs to be replicated. This is useful when a new file is created on one server and needs to be replicated to another server. Instead of replicating the entire file, DFS Replication can use portions from files that are similar to the replicating file to minimize the amount of data transferred over the WAN.

Cross-file RDC is available only if one of the servers (in a replicating pair of servers) is running Windows Server 2003 R2, Enterprise Edition; Windows Server 2003 R2, Datacenter Edition; or Windows® Storage Server R2, Enterprise Edition. For example, if a branch office server is running Windows Server 2003 R2, Standard Edition, and the hub server is running Windows Server 2003 R2, Enterprise Edition, cross-file RDC will be used.

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

## Efficient and scalable replication

When two members of a replication group begin to synchronize with each other, they use an efficient algorithm for determining which files need to be replicated. The amount of metadata exchanged is minimal, and the possibility of sending changes unnecessarily (due to the order the changes occur) is eliminated because the synchronization is state-based instead of event-based, as in File Replication service (FRS).

The introduction of state-based synchronization, along with RDC, allows DFS Replication to support replicating more files to more members than FRS. The tested scalability figures are as follows:

- Each server can be a member of up to 256 replication groups.
- Each replication group can contain up to 256 replicated folders.
- Each server can have up to 256 connections (for example, 128 incoming connections and 128 outgoing connections).
- On each server, the number of replication groups multiplied by the number of replicated folders multiplied by the number of simultaneously replicating connections must be kept to 1024 or fewer. If the replication schedule is staggered, you do not need to count the connections that are not replicating due to a closed schedule.
- A replication group can contain up to 256 members.
- A volume can contain up to 8 million replicated files, and a server can contain up to 1 terabyte of replicated files. These are tested numbers and recommended guidelines for performance and scalability reasons.

## Flexible scheduling and bandwidth throttling

DFS Replication supports replication scheduling and bandwidth throttling in 15-minute increments during a 7-day period. When specifying a replication interval, administrators choose the start and stop times as well as the bandwidth to use during that interval. The settings for bandwidth usage range from 16 kilobits per second (Kbps) to 256 megabits per second (Mbps) as well as full (unlimited) bandwidth. Administrators can configure a default schedule and bandwidth that applies to all connections between members and, optionally, create a custom schedule and bandwidth for individual connections.

To help administrators configure replication windows for servers in different time zones, administrators can set the schedule so that the server that initiates replication interprets the schedule as being in either Coordinated Universal Time (UTC) or in the server's local time.

## Supported in stand-alone and domain-based namespaces and on individual folders

DFS Replication can be used in both stand-alone and domain-based namespaces as well as for folders that are not part of any namespace. The folders to be replicated can be shared or unshared.

## Self-healing after USN journal wraps and database corruption

DFS Replication provides "self-healing" for update sequence number (USN) journal wraps and Jet database corruption. Although replication temporarily stops during this healing process, the service recovers without any administrator intervention. To repair itself, DFS Replication scans the file system and re-creates the DFS Replication database that contains metadata associated with the files in the replicated folder. The database must then be synchronized with a database on another member. During the synchronization process, the amount of metadata sent across the network is dictated by the number of files under the replicated folder's local path (that is, the number of ID records in the database) and the size of the metadata to be sent per file. The metadata size for a file is two times the length of the file name plus approximately 144 bytes. Additional RPC and TCP overhead results in a roughly 5 percent overhead; therefore, in the worst case, for 1 million files in the database with an average file name size of 50 bytes, approximately 194 megabytes (MB) of metadata is sent across the network.

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

## Ease of member recovery

DFS Replication stores its global configuration settings, such as the topology and replication schedule, in Active Directory. These settings are also cached in a local .xml file on each member; DFS Replication can rebuild this file (using the settings stored in Active Directory) if the file becomes corrupted or if the member is restored after a failure. This type of self-healing allows for greater server uptime and reliability and makes it easier to rebuild a member of a replication group during disaster recovery. DFS Replication also uses the .xml file to store member-specific settings, such as debug log settings or RPC port settings that are configured using Windows Management Instrumentation (WMI).

## Simple and flexible prestaging of new servers

Before adding a new server to a replication group, administrators can prestage the replicated folders on the target servers by copying the data to the servers, restoring a backup, or copying files from a tape, DVD, or removable hard disk. As described earlier, the synchronization process is very efficient in terms of bandwidth usage and metadata exchanged, resulting in minimal WAN traffic during the initial synchronization for files that are the same on the source server (known as the *primary member*) and target servers. If the files on the target server are out-of-date, DFS Replication will use RDC to replicate only the changes that occurred since the data was prestaged. Any prestaged files that existed on the target servers, but not on the source server, are moved to the PreExisting folder under the target server's replicated folder path.

## New management tools

Administrators can use the DFS Management snap-in to configure both DFS Namespaces and DFS Replication. The snap-in provides integration between the two Distributed File System components so that administrators can:

- Select an existing folder in a namespace and configure DFS Replication on the folder targets (shared folders) associated with the folder.
- Add a replicated folder to an existing namespace.

Administrators can also perform administrative tasks at the command line using Dfsradmin.exe or Dfsrdiag.exe. Both tools are part of Windows Server 2003 R2. Additional configurations can be made programmatically using WMI.

## Delegation of management tasks

Administrators who are not part of the Domain Admins group can be delegated the ability to create new replication groups in a domain, manage an existing replication group, or both. Members of the Domain Admins group can use either the DFS Management snap-in or the Dfsradmin.exe command-line tool to perform this delegation.

## Built-in health metrics and diagnostic events

DFS Replication provides built-in WMI providers for monitoring the health of DFS Replication. For example, the WMI providers can report USN journal wraps, database loss, insufficient disk space, network connectivity issues, sharing violations, excessive replication, and clock skew between members. These events are also reported in the DFS Replication event log, which is used exclusively for storing events related to replication.

There are two ways of monitoring DFS Replication: a built-in diagnostic report and the Windows DFS Replication Management Pack for Microsoft Operations Manager (MOM). The diagnostic report is an .html file generated using either the DFS Management snap-in or the Dfsradmin.exe command-line tool. The report contains a wealth of DFS Replication information, including error and warning events; service state and uptime; replication efficiency based on RDC and stream compression; backlogged sending and receiving transactions; free disk space, and more. An example health report is shown in the following figure:

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

## DFS Replication Health Report

(Show All)

<b>Data collected on:</b>	Tuesday, October 04, 2005 at 7:00:04 PM (GMT-8:00)
<b>Replication Group:</b>	SOFTWARE (na.corp.contoso.com)
<b>Reference member:</b>	DDS-1 (DDS-1.na.corp.contoso.com)
<b>Server scope:</b>	Selected 30 of 30 servers
<b>DFS Replication efficiency savings:</b>	<div style="display: flex; align-items: center;"> <div style="width: 66.98%; height: 15px; background-color: #0070C0; margin-right: 5px;"></div> <span>66.98% reduction (36.53 GB replicated instead of 111 GB)</span> </div>

<b>Server health:</b>	<span style="color: green;">✔</span> Servers with no errors or warnings (14) <span style="margin-left: 100px;"><span style="color: red;">⚠</span> Servers unavailable for reporting (2)</span>
	<span style="color: red;">✘</span> Servers with DFS Replication errors (5) <span style="margin-left: 100px;"><span style="color: orange;">⚠</span> Servers with DFS Replication warnings (14)</span>

✘ ERRORS (5 servers with errors)
(Show All)

☐ B01 (1 error) (View Server Details)

DFS Replication service stopped replication on replicated folder Software\Betas\Developers\Visual Studio.

☐ C01 (1 error) (View Server Details)

DFS Replication unable to replicate files for replicated folder Software\OS\Windows XP due to insufficient disk space.

☐ H01 (1 error) (View Server Details)

DFS Replication unable to replicate files for replicated folder Software\OS\Windows Server 2003 due to insufficient disk space.

☑ O01 (6 errors) (View Server Details)

☑ SRV-1 (7 errors) (View Server Details)

⚠ WARNINGS (14 servers with warnings)
(Show All)

⚠ SERVERS UNAVAILABLE FOR REPORTING (2 servers not reporting)
(Show All)

☰ SERVER DETAILS (30 servers)
(Show All)

☐ C-SRV-1
(Show All)

DNS name: c-srv-1.europe.corp.contoso.com  
 Domain name: europe.corp.contoso.com  
 Reference domain controller: dc-01.europe.corp.contoso.com  
 IP address: 151.67.11.58  
 Site: EU-CORP  
 Time zone: (GMT-8:00)

✘ **ERRORS (There are no errors to report)**

⚠ **WARNINGS (There are no warnings to report)**

☐ **INFORMATIONAL**

- ☐ Service state: Running
- ☐ DFS Replication service uptime: 4 days 9 hr.
- ☐ Summary of replicated folder status

The following table provides a high-level overview of replicated folder status on this server.

Replicated Folder	Status	Backlogged Sending Transactions	Backlogged Receiving Transactions	# of Files Received	DFS Replication Efficiency Savings
Software\Applications\Server\ISA Server 2004	Normal	120	289	568	93.13%
Software\Applications\Server\SQL Server	Normal	84	416	632	65.54%

The Windows DFS Replication Management Pack is a real-time monitoring tool that monitors the health state and replication progress of DFS Replication on a per-member basis. Administrators can use this management pack to monitor the state of the DFS Replication service, replication groups, replicated folders, and the volumes where replicated folders are stored. Conditions that cause full or partial replication failures will temporarily change the state of objects; these objects will change back to a healthy state when the problem is resolved, either automatically by the service (as in the case of intermittent connection failures) or by the administrator.

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

The management pack is supported on Microsoft Operations Manager 2005 and will be available for download on the Microsoft Web site (<http://go.microsoft.com/fwlink/?LinkId=55220>).

## DFS Namespaces Enhancements in Windows Server 2003 R2

As described earlier, the Distributed File System technology in Windows Server™ 2003 has been renamed to DFS Namespaces. Although the underlying service and basic functionality are unchanged, there are a number of DFS Namespaces enhancements exposed in Windows Server 2003 R2. These enhancements, introduced as updated APIs in Windows Server 2003 Service Pack 1 (SP1), provide easier management and more flexibility for namespaces used in branch offices.

### New and updated management tools

The new DFS Management snap-in in Windows Server 2003 R2 provides an improved graphical user interface for managing namespaces and DFS Replication. The snap-in allows administrators to configure DFS Namespaces enhancements, such as target priority, delegation, and client failback, as well as existing features that in Windows Server 2003 were configurable only by using Dfsutil.exe. For example, administrators can use the DFS Management snap-in to configure how servers are ordered in a referral, such as by lowest cost or restricted to the same site as the client. Administrators can also enable root scalability mode, which reduces the load on the primary domain controller (PDC) emulator in large namespaces.

The command-line tool Dfsutil.exe, which is part of the Windows Support Tools in Windows Server 2003 SP1, is updated to include the DFS Namespaces enhancements. The operating system tool Dfscmd.exe is also updated to allow administrators to move or rename folders in the namespace.

### Client failback

Client failover in DFS Namespaces is the process in which clients attempt to access another server in a referral after one of the servers fails or is removed from the namespace. Unless client failback is configured, clients will continue using the server they failed over to unless the client is restarted or the client's referral cache is cleared. When client failback is configured, and clients have the appropriate client failback hotfix installed, clients will fail back to a preferred, local server when it is restored.

### Target priority

When a client accesses a namespace, the client receives a referral that contains a list of targets associated with the namespace root or folder. These targets are listed according to the current ordering method for the namespace or folder. To fine-tune how particular targets are ordered, administrators can specify whether a server appears first or last in a referral. Assigning target priority is useful in many scenarios, such as "hot-standby" scenarios where one server is considered the server of last resort. In this scenario, the administrator can specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages.

### Better delegation

Administrators can easily delegate the ability to create domain-based namespaces and manage individual stand-alone and domain-based namespaces. The DFS Management snap-in sets the appropriate permissions on either the DFS Namespace configuration objects in Active Directory or in the namespace server's registry (depending on the namespace type).

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

## Ability to restructure the namespace

Renaming or moving folders in the namespace is easy using the DFS Management snap-in. Administrators can restructure the namespace to correct mistakes or to adjust the hierarchy as business needs change or as new folders are added to the namespace. Administrators can also move namespace folders by using the updated version of the command-line tool, Dfs.cmd.exe.

## Configuration Requirements

The following sections describe how servers must be configured to support DFS Replication and DFS Namespaces in Windows Server 2003 R2.

### DFS Replication requirements

Before DFS Replication can be deployed, administrators must configure servers and storage as follows:

- The Active Directory schema must be updated to include the new DFS Replication objects. These schema changes are provided on disc 2 of the Windows Server 2003 R2 operating system installation discs. This schema can be applied to domain controllers running Microsoft® Windows® 2000 Server, Windows Server 2003, and Windows Server 2003 R2.
- The servers that will participate in DFS Replication must run Windows Server 2003 R2. After you install Windows Server 2003 R2, you must install the DFS Replication Service on each server that will take part in replication, and you must install the DFS Management snap-in to manage replication. For information about installing the DFS Management snap-in on a client computer running Windows XP with Service Pack 2 (SP2), see "Administering Distributed File System from a computer running Windows XP" later in this document.
- Antivirus software must be compatible with DFS Replication; contact your antivirus software vendor to check for compatibility.
- Servers in a replication group must be in the same forest. You cannot enable replication across servers in different forests.
- Replicated folders must be stored on NTFS volumes.
- On server clusters, replicated folders should be located in the local storage of a node because the Distributed File System Replication service is not cluster aware and the service will not fail over to another node.

#### Important

DFS Replication is not supported for SYSVOL replication in Windows Server 2003 R2. Do not attempt to configure DFS Replication on SYSVOL by disabling FRS and setting up a replication group for SYSVOL. Continue to use FRS for SYSVOL replication on domain controllers running Windows Server 2003 R2. FRS and DFS Replication can co-exist on the same member server or domain controller.

### DFS Namespaces requirements

To enable all features in DFS Namespaces, you must configure lab servers and clients as follows:

# Overview of the Distributed File System Solution in Microsoft Windows Server 2003 R2 (Microsoft Corporation)

- Servers where the DFS Management snap-in will be used must run Windows Server 2003 R2 or Windows XP SP2. For information about installing the DFS Management snap-in on a client computer running Windows XP with SP2, see "Administering Distributed File System from a computer running Windows XP" later in this document.
- To support new namespace features, all servers that host namespaces must run Windows Server 2003 SP1 or Windows Server 2003 R2.
- To support new namespace features, all domain controllers must run Windows Server 2003 with SP1 or Windows Server 2003 R2.
- Namespaces must be created on NTFS volumes.
- Clients that access namespaces can run any of the supported client operating systems, but only clients running the following operating systems, service packs, and the appropriate client failback hotfix can be configured for client failback:
  - Windows XP SP2 and the Client Failback hotfix.
  - Windows Server 2003 SP1 and the Client Failback hotfix.

For more information about the Client Failback hotfix, see article 898900 in the Microsoft Knowledge Base on the Microsoft Web site at <http://go.microsoft.com/fwlink/?LinkId=53202>.

For more information about supported operating systems for DFS Namespaces clients, see "Evaluating Client and Server Compatibility" on the Microsoft Web site at <http://go.microsoft.com/fwlink/?LinkId=37455>.

It is possible to use DFS Namespaces when domain controllers and namespace servers run a mix of Windows 2000 Server, Windows Server 2003 with SP1, and Windows Server 2003 without SP1, but some functionality is disabled or available inconsistently, depending on the operating systems on the servers. Some examples of mixed-mode behavior are as follows:

- If domain controllers or namespace servers are running Windows Server 2003 without SP1, they cannot provide referrals that support target priority or client failback.
- If domain controllers or namespace servers are running Windows 2000 Server, they cannot provide referrals that support target server priority or client failback, nor can they order targets by lowest cost in referrals. Additional configuration is required to enable these namespace servers and domain controllers to detect the site of each target server in the namespace. For details, see the answer to the question "What are the issues to consider when I use multiple servers running Windows 2000 Server and Windows Server 2003 to host a domain-based DFS root?" in the DFS FAQ on the Microsoft Web site at <http://go.microsoft.com/fwlink/?LinkId=39465>.
- If the DFS Management snap-in connects to a namespace server that is not running Windows Server 2003 SP1 or Windows Server 2003 R2, none of the new configuration settings (such as client failback and target priority) can be enabled. Renaming or moving folders will not work, and delegation will not be effective.

## Administering Distributed File System from a computer running Windows XP

You can also manage DFS Namespaces and DFS Replication from a computer running Windows XP with SP2 by installing the Administration Tools Packs for Windows Server 2003 R2. When this pack is installed, the DFS Management snap-in is available as part of the File Server Management snap-in. For more information about installing this pack, see the Microsoft Web site at <http://go.microsoft.com/fwlink/?LinkId=55225>.