

---



Microsoft  
**Windows Server 2003**

## Simplifying Infrastructure Complexity with the Windows Distributed File System

*Microsoft Corporation*  
*Published: March 2003*

---

### **Abstract**

The Microsoft® Distributed File System (DFS) is a strategic storage management architecture that creates administrative flexibility and reduces the long-term cost of storage ownership for enterprises. Just as a file system provides uniform access to collections of sectors on disks, DFS provides uniform, simplified access to servers, shares, and files. This document reviews DFS and describes its benefits for IT professionals, business decision makers, and anyone who wants to learn how to manage file servers and shares more efficiently.

*The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.*

*This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.*

*Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.*

*Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.*

*© 2003 Microsoft Corporation. All rights reserved.*

*Microsoft, Active Directory, MSN, Win32, Win64, Windows, Windows logo, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.*

*The names of actual companies and products mentioned herein may be the trademarks of their respective owners.*

---

## Contents

<b>Introduction</b> .....	<b>1</b>
About This Document .....	2
<b>Is Your Network Storage Maximizing Resources?</b> .....	<b>3</b>
A Loss of Time and Productivity .....	3
Unused Storage Space.....	3
Added Costs.....	3
<b>How DFS Simplifies Complexity and Streamlines Network Storage</b> .....	<b>4</b>
How DFS Works.....	4
A Secure Architecture .....	7
DFS Setup.....	8
<b>Summary</b> .....	<b>11</b>
<b>Glossary of DFS Terminology</b> .....	<b>12</b>
<b>Related Links</b> .....	<b>13</b>

---

## Introduction

If there were an easy way to view, access, and manage files scattered across the multiple file servers and shares that make up most enterprise networks, wouldn't you want to take advantage of it? There is. It's called the Distributed File System (DFS), an efficient and cost effective storage architecture for Microsoft® Windows environments—and you probably already have it. DFS does more than help you manage files across shares and servers. It also provides fault-tolerant access, an advantage for users and administrators alike, as well as other cost-saving enterprise benefits.

To appreciate why DFS can be a file management breakthrough for your IT organization, it helps to understand the problems that it solves. The goal of most IT groups is to manage file and server resources efficiently while keeping them available and secure for users. As networks expand to include more users as well as multiple physical devices—located on site or remotely—IT administrators find it increasingly difficult to keep users connected to the files they need. On one hand, distributing resources across a network makes them more available to more people and promotes cross-organizational efforts—that is, everyone can have access to the information they need. On the other hand, file servers seem to crop up overnight in today's wide area networks (WAN), making it hard for users to know where to look for information but also easy for them to create and duplicate files. Vast quantities of network storage space and bandwidth are quickly consumed in a manner that's difficult for IT groups to support. Simply swapping out an old server becomes a major communication chore when users across an enterprise must be notified to update necessary links and file paths.

From this perspective, the success of the distributed paradigm has become its liability as organizations struggle to manage their decentralized resources. But DFS changes all that. DFS is a strategic storage management solution that gives administrators a more flexible way to centrally manage their distributed resources. With DFS, administrators can create simplified views of folders and files—that is, a virtual organization called a **namespace**—regardless of where those files physically reside in a network. A namespace makes file paths transparent from a user's perspective. For example, when the name of a file server changes, the user's intuitive, virtual DFS paths do not break. DFS stores the path names logically as a single namespace, a significant change from the traditional data management approaches.

With globally available DFS namespaces, file servers become highly available for access across the enterprise. Organizations of any size can benefit from the following enhanced benefits of DFS:

- **Simplified access to files and folders.** DFS provides an intuitive way for users to access multiple file servers throughout your organization. File and folder replicas can even be made available across multiple sites to provide inexpensive access to users within each site.
- **Easier maintenance.** Because DFS eliminates machine name dependencies and provides a uniform, virtual folder space for end users, you can perform maintenance tasks, such as server consolidation or data migration, without affecting how users and applications locate data. Because the shares are virtual, administrators have more flexibility to move shares. Servers aren't tied directly to share names. As a result, servers can be replaced, changed, or moved without impacting the share name or availability.
- **Increased cost savings.** Managing storage in a distributed environment can cost as much as 75 percent of an organization's IT budget according to a Gartner Group study. DFS can reduce this expense through more efficient management of existing resources.

- **Fault-tolerant file access.** DFS transparently routes clients to the next available server in the event of failure so users can always access shares while remaining unaware of any problem.
- **High availability and performance.** DFS provides load distribution, making data highly available. In environments where multiple servers have thousands of users, file server response is enhanced.
- **Improved file path performance.** DFS provides the added benefit of reduced delays that occur when user's access heavily used shared folders.
- **Smart access to resources.** Through DFS, clients can access the closest resources and file servers, reducing stress on network bandwidth.

## What's New in Windows Server 2003 for DFS

There are number of enhanced features in the DFS in Windows Server 2003. These include:

- **Multiple namespaces per DFS Server.** A single Windows Server 2003 DFS server can now host hundreds of DFS namespaces (supported by Windows Server 2003 Enterprise and Datacenter editions only), which increases namespace flexibility without requiring additional servers.
- **Enhanced multiple-root DFS.** DFS now allows for more than 16 root targets and the number of DFS targets can be scaled across widely separated geographic locations. The number of root targets can range into the hundreds and will depend on several factors including the size of the root target DFS namespace metadata.
- **Integrated with Microsoft Active Directory® link costing.** When DFS is implemented in conjunction with an organization's existing Active Directory service, DFS ranks all available client-server connections by the site link cost defined in Active Directory. Users can then transparently access data from the nearest available file replica.
- **Dynamic site selection.** Windows Server 2003 DFS root servers dynamically detect when DFS root servers or link targets change sites. This makes it easier to relocate servers to other parts of your network and ensures that site selection remains predictable and efficient.
- **Enhanced script-based management.** Microsoft Windows Server 2003 offers significantly enhanced command line tools for scripting creation, updates, and deletions of DFS namespaces. For even easier management, scripts can be used to monitor and back up DFS.
- **Improved replication management.** The DFS snap-in in Windows Server 2003 supports configurable replication topologies including ring, full mesh, hub-and-spoke, and custom. In addition Windows Server 2003 includes a range of new tools for monitoring file replication.

## About This Document

If your organization already has file servers running Microsoft Windows NT® 4 or Windows 2000, but you have not yet implemented DFS, this technical article can help you understand the benefits of adopting DFS. This document also introduces you to the many improvements to DFS in the Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition, operating systems (collectively referred to as Windows Server 2003 in this article).

---

## Is Your Network Storage Maximizing Resources?

Consider how a network grows. From the initial plan, an enterprise WAN often grows haphazardly as organizations within the enterprise acquire their own file servers and users and applications quickly take advantage of and use up the new resources. Many organizations have hundreds of file servers throughout the enterprise. Shared folders are often distributed on many of those file servers across a WAN. Managing a file system like this becomes a familiar IT headache. But the pain is shared by users and business decision makers alike when it becomes apparent that this key enterprise resource is wasting space, money, and time. The good news is that DFS can help organizations manage their file systems much more efficiently and easily.

### A Loss of Time and Productivity

An inefficient storage architecture is a waste of users' and administrators' time. When files are spread across a WAN, not only do they become difficult for network administrators to support, users waste time trying to find what they need. One solution has been for users to create shortcuts to the physical servers containing file shares they require. For example, a user might map a drive letter to a server and share. The full Universal Naming Convention (UNC) path might be something like \\Server\Share\Path\Filename, but users can map drive x: to \\Server\Share. By doing so they avoid remembering long or meaningless names, and they can navigate faster to the information they need. Some additional navigation is usually required to locate a particular file on the share.

For small groups of users, this solution may suffice, but as networks continue to grow in size and as organizations begin to use existing internal and external storage for purposes such as intranets, mapping a single drive letter to individual shares scales poorly. Even if users are comfortable using UNC paths directly, they can still become overwhelmed by the number of places where files can be stored. And each time a network administrator needs to move, rename, or upgrade a file server, they quickly discover how inflexible this solution is, while the entire user community feels the change as a productivity loss.

### Unused Storage Space

The power of distributed networks is in how they connect multiple users and devices across great distances. The downside may be in how they have inadvertently encouraged data proliferation. While network storage requirements keep increasing, according to the Gartner Group, only 30 to 40 percent of storage in a distributed environment is used compared to 80 percent in mainframe environments. That means there is a lot of wasted space that's virtually impossible to consolidate or manage.

### Added Costs

Storage inefficiencies waste money. As network storage capacity gets filled, many organizations choose to purchase more and bigger hardware, which seems like a reasonable response, particularly since hardware prices have gone down in recent years. However, another Gartner Group study shows that for every \$1 US spent on storage hardware, another \$3.50 US is spent on administration. Most organizations cannot afford the added expense. Disk space aside, network bandwidth is wasted as well. When accessing shares during peak usage times, clients may find that response time suffers when there is no intelligence to distribute loads or direct clients to other available file servers.

## How DFS Simplifies Complexity and Streamlines Network Storage

The solution is a centrally managed distributed scheme, which may seem to be a paradox. DFS offers a storage management platform that meets the above challenges. Simply put, DFS is a platform for distributed file management. DFS does for servers and shares what file systems do for hard disks. Just as file systems provide uniform named access to collections of sectors on disks, DFS provides a uniform naming convention and mapping for collections of servers, shares, and files.

DFS works by allowing administrators to create uniform, meaningful names for the shared folders on a network. Because the names are independent of the underlying physical server descriptions, the physical location of data becomes transparent to applications and users, who are no longer required to know the physical name assigned to each server they need to access.

### How DFS Works

The DFS service consists of a client component and a server component. The client component is included with all Windows clients and allows the client to make requests to the DFS server. The server component is included with Windows NT, Windows 2000, and Windows Server products. The DFS server component receives a client request and redirects or refers it to a physical target, similar to the way a browser receives a DNS call and refers the client to a Web site.

The view of shared folders on different servers is called the DFS namespace. Another way to think of a DFS namespace is as an intuitive view of shared folders on different servers. Or think of it as a virtual UNC path. A namespace is much easier to use as Figure 1 below shows. For example, an administrator can create a single namespace for commonly accessed corporate documents called `\\myCompany.com\corpdata\Sales` that maps to physical resources residing on multiple servers that could be located just about anywhere.

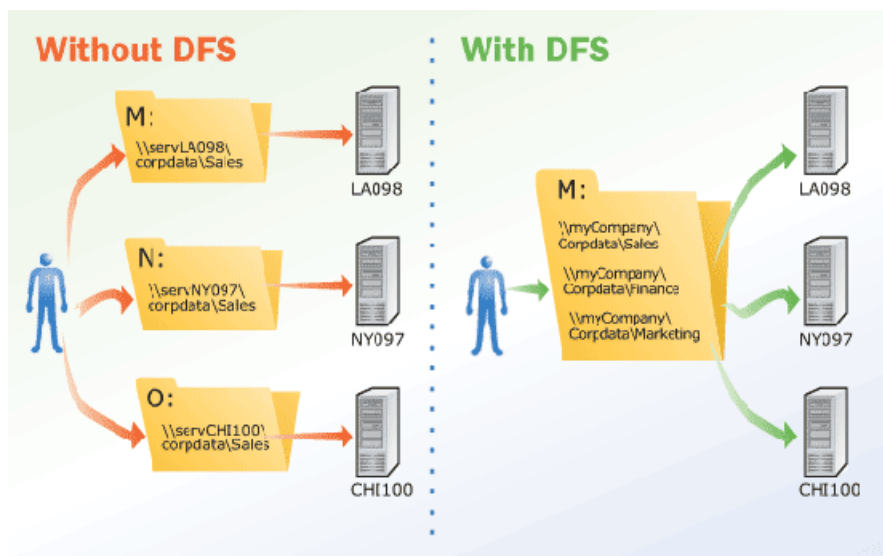


Figure 1. DFS simplifies access to frequently accessed file shares with a unified namespace

When DFS is not used, users are able to select file servers without regard to geography or user load, much less logical namespace shown below in Figure 2.

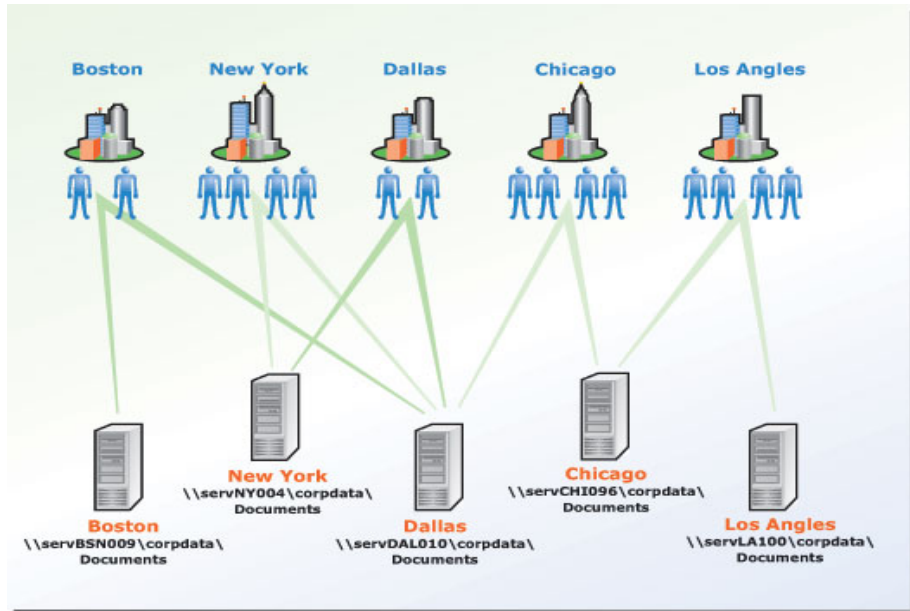


Figure 2. Before implementing a DFS namespace

After DFS is implemented, however, users are automatically routed to the closest server, as Figure 3 shows. Moreover, if a server becomes unavailable, DFS ensures that users are routed to the next closest server by using site-costing. This capability is available when using Windows Server 2003 in either a stand-alone or domain-based DFS configuration. Windows Server 2003 will also use the site and costing information in Active Directory to determine whether sites are linked by inexpensive, high-speed links or by expensive WAN links.

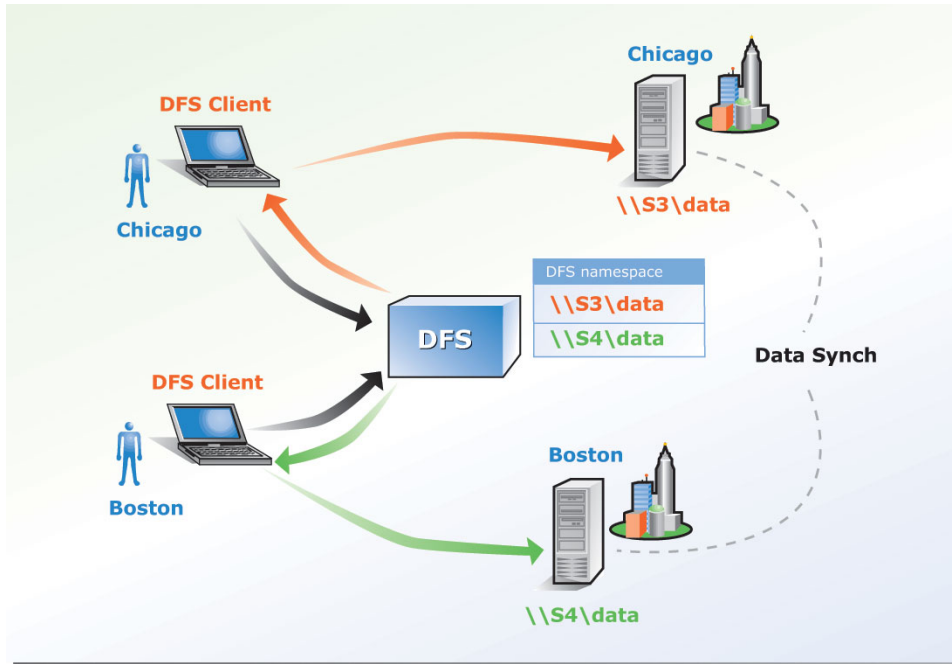


Figure 3. With DFS, users are automatically routed to the closest server

## A Reliable Architecture

In today's business environment, it is critical to ensure data is always available, even if a server, disk drive, or file occasionally fails or otherwise becomes unavailable. DFS delivers several benefits which make your data more reliable and your users more productive.

### Higher Data Availability and Easier Maintenance

Having critical data located in a single file share on a single file server or disk volume can mean a significant loss of productivity if that data were suddenly unavailable to the information worker. DFS is designed to deliver enhanced data availability by pointing to multiple volumes that can be alternates for each other. In this way, should a volume on which your important data is stored become unavailable, DFS hands the request for that data to an alternate volume.

DFS allows you to mount multiple copies of read-only shares under the same logical DFS name to provide alternate locations for accessing data. Not only does this enhanced data availability have value for the information work, but the IT administrator benefits as well because individual shares participating in DFS can be taken offline without affecting the remaining portion of the volume namespace. So, DFS delivers higher data availability and a flexible volume administration model as can be seen in Figure 4 below.

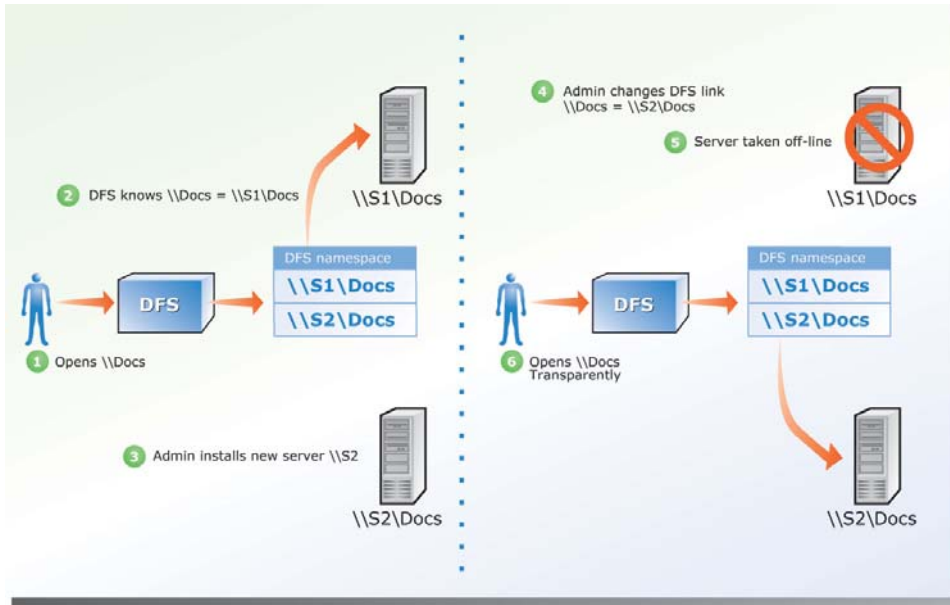


Figure 4. DFS simplifies maintenance of servers without impacting your users

### Optimizing Bandwidth Usage in Wide Area Networks

Having large servers and terabytes of data does not necessarily ensure your users of high performance access to important data. Network bandwidth and throughput play an important role in assuring that your file systems deliver peak performance for the information worker.

DFS provides load distribution performance gains because it can evenly distribute client access to DFS volumes across multiple alternate network shares. Administrators can mount multiple read-only copies of shares under the same logical DFS name which permits limited load balancing between drives or servers. As users request files from the DFS volume, they are transparently referred to one of the network shares comprising the DFS volume. So for example, if three hundred users require access to one volume DFS can split the users among copies of two or more servers to balance the load.

### A Secure Architecture

Security is built into DFS, no additional overhead is necessary to implement the security features. DFS is secure in part because it only redirects names. In addition, a user connecting to a DFS namespace is permitted to access only the files or folders for which he or she has appropriate permissions.

Technically speaking, DFS is a multi-protocol architecture that uses the default, built-in authentication protocols server message block (SMB) and LAN Manager (LM) to communicate between a DFS client and a DFS server as Figure 5 illustrates. In this figure, Step 1 shows the client making a request of the DFS namespace, in Step 2, DFS returns the appropriate path to the data (including AD site-costing information when AD is in use), and Step 3 illustrates the client making a connection to the server and share.

A DFS namespace can consolidate many types of share, including shares created by Network File System (NFS) servers, Services for Macintosh™ servers, and Netware™ Core Protocol (NCP) servers.

As long as the Windows client machine has a suitable client redirector driver installed supporting these other protocols, DFS can be used to provide a virtual namespace for these types of shares.

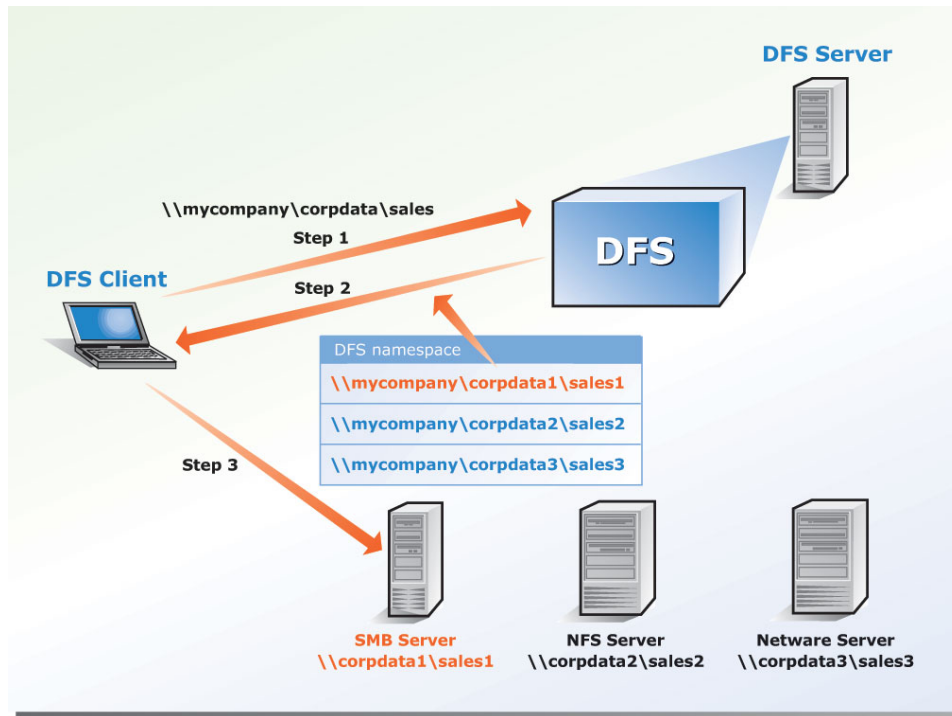


Figure 5. DFS uses SMB/LM protocol to communicate between a DFS client and the DFS server, but then uses the existing file system's protocol to connect to remote shares

## DFS Setup

Figure 4 shows the starting point of the DFS namespace. The namespace starts with a root that maps to one or more root targets. The root is often used to refer to the namespace as a whole. Below the root are links that map to their own targets, each of which corresponds to a shared folder on a separate server. The DFS root must reside on an NTFS volume. A DFS root has one of the following formats: `\\servername\rootname` or `\\domainname\rootname`.

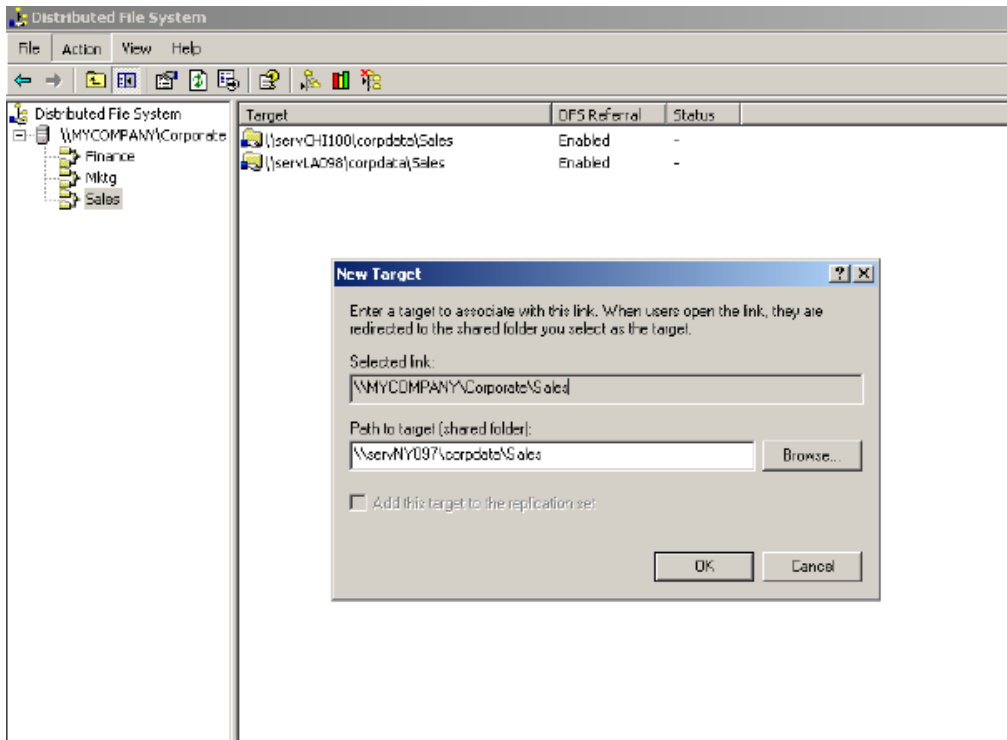


Figure 6. Easy-to-use tools make it simple for administrators to set up a DFS namespace

---

## How to Deploy DFS in Your Organization

DFS can be deployed flexibly to suit the degree of administrator access, namespace size, and fault tolerance that your organization requires. In addition, DFS can be implemented in conjunction with an organization's existing Active Directory directory service. The following considerations affect the way DFS is deployed:

- **For environments with Active Directory, use a domain-based deployment.** You can store the configuration information for a DFS namespace in Active Directory to create what is called a domain-based DFS deployment. This type of DFS namespace can include approximately 5,000 links (although the actual number of links depends on the number of characters in the root/link description—shorter names yield more links). Multiple DFS targets can be created in the same domain to ensure availability. Across the multiple targets, you can replicate files using scripts, the File Replication Service (FRS), or a third-party replication tool.
- **For a single namespace with more than 5,000 links, use a standalone deployment.** You can create a single namespace with more than 5,000 links by setting up DFS locally on a host server—the standalone DFS root. The local registry, not Active Directory, is used to store DFS configuration information. Broader administrator access is possible than with domain-based DFS, because members of the local Administrators group on the local server can implement a standalone DFS setup; membership in the Domain Administrators group is not required. A single, standalone DFS namespace can support as many as 50,000 links.
- **For maximum scalability, use a cascaded deployment.** By mixing standalone and domain-based DFS in a cascading configuration, you gain the highest scalability. Since DFS allows up to eight levels of cascading when used with Windows NT, Windows 2000, or Windows XP clients. DFS supports a virtually unlimited number of links.
- **For added fault tolerance in standalone DFS, use clusters.** To ensure the availability of the DFS namespace in a standalone deployment, you can create a standalone DFS root on a clustered file server. This is not required for domain-based DFS configurations – these use Active Directory to achieve high availability.

Implementation is made easier by the fact that many types of clients and servers are compatible with DFS. For example, any of the following servers works well for hosting targets: Windows Server 2003 family, Windows 2000, and Windows NT 4.0 with Service Pack 6a. In addition, any of the following clients can access targets in a DFS namespace: Windows Server 2003 family, Windows XP Professional, Windows 2000 (Professional and Server versions), Windows NT 4 with Service Pack 6a, and Windows 98/95 with the DFS client installed.

---

## Summary

By deploying DFS, enterprises can improve their file storage management, increase file storage availability, and lower storage costs. Because DFS creates a uniform, intuitive namespace independent of the underlying physical server descriptions, users and administrators alike benefit. Not only can users easily access and manage files that are physically distributed across a network, but administrators can maintain networked systems more efficiently. In the past, whenever a file server needed to be moved, renamed, or upgraded, the entire user community lost productivity. Today, the ability of DFS to point to multiple targets provides fault-tolerant access as well as a degree of load sharing for a new level of network flexibility.

You probably already have DFS because a version has been included in all Windows clients and every Windows server product since Windows NT 4. The Windows Server 2003 family makes DFS a strategic asset for enterprises through its ability to improve storage-related functionality, reduce costs, and make storage management much simpler and more comprehensive for IT organizations and end users alike.

---

## Glossary of DFS Terminology

**Active target:** The link or root target that is currently being used by the client to access the root or link.

**DFS Client:** A Windows-based computer system which has the client part of the DFS system software installed and active.

**DFS Metadata:** The metadata that describes the configuration of one DFS Namespace.

**DFS namespace:** A hierarchical namespace consisting of a DFS root, and many links and targets. The namespace starts with a DFS root that maps to one or more root targets. Below the root are links that map to their own targets.

**DFS Path:** A UNC path that starts with a DFS root.

**DFS Root:** The starting point of a DFS namespace (a DFS root is a part of a DFS Namespace). A DFS root maps to one or more root targets.

**DFS Server:** A Windows-based computer system which has the server part of the DFS system software installed and active and is hosting a DFS Root.

**Domain-based DFS:** A DFS Namespace hosted by an Active Directory domain.

**Link:** An element in a DFS namespace that lies below the DFS root and maps to one or more Link targets, each of which corresponds to a UNC name.

**Link Target:** The mapping destination of a DFS link, expressed as a UNC Path.

**Member Server:** A Server that belongs to an Active Directory domain.

**Mixed mode domain DFS:** A domain-based DFS Namespace with at least one root target hosted on Windows 2000 server and one root target hosted on Windows Server 2003.

**Parent Directory:** The parent directory for a DFS path is the parent directory of the entity on the shared folder that the DFS Client ultimately resolves the DFS path to. This means that in the case of a DFS path involving multiple link targets for a single link, there will be multiple candidate Parent Directories, and each DFS Client will have their own determination of the Parent directory for a given DFS path at any point in time.

**Root Target:** The mapping destination of a DFS root, expressed as a UNC Path.

**Site:** An Active Directory concept. A site is an object created in the Active Directory configuration container. From a conceptual point of view, it should map to a location. IP subnets are also objects in the configuration container. They can be used to define IP address ranges. One IP subnet can be assigned to exactly one site, multiple IP subnets can be assigned to the same site.

**Standalone DFS:** A DFS Namespace hosted by a single computer. Note that this may also include the clustered (MSCS) case.

**Target:** The mapping destination of a DFS root or link, expressed as a UNC Path.

**UNC path:** A pathname following the universal naming convention (UNC). This consists of \\servername\path. The server name must be a validly formed name of an entity on a network. Path must be a valid win32 pathname, with elements separated by “\” characters.

---

## Related Links

See the following resources for further information:

- For in-depth technical and troubleshooting information about DFS, see ““[Designing and Deploying File Servers](#),”” in the *Planning Server Deployments* book of the Windows Server 2003 Resource Kit <http://www.microsoft.com/windows/reskits/default.asp>
- For a list of the file and print benefits provided by the Windows Server 2003 family, see “What’s New in File and Print Services for Windows Server” at <http://www.microsoft.com/windowsserver2003/evaluation/overview/technologies/fileandprint.asp>
- For information about features in the Windows Server 2003 family, see “Introducing the Windows Server 2003 Family” at <http://www.microsoft.com/windowsserver2003/evaluation/overview/default.asp>
- For information about new storage management offerings in Windows Server, see “What’s New in Storage Management” at <http://www.microsoft.com/windowsserver2003/evaluation/overview/technologies/storage.asp>
- For details about Windows 2000 information storage technologies, see “Windows 2000 File and Print Services” at <http://www.microsoft.com/windows2000/technologies/fileandprint/default.asp>

For the latest information about Windows Server 2003, see the [Windows Server 2003 Web site](http://www.microsoft.com/windowsserver2003) at <http://www.microsoft.com/windowsserver2003>.