

**WHY THE NEW
DFS IS SO
ATTRACTIVE
TO WINDOWS
SYSTEMS
ADMINS**

not your father's DFS



- 02 Not your father's DFS
- 10 New DFS terminology
- 13 Storage price drops; where's the downside?
- 14 Storage management tasks that'll waste your time

not your father's DFS

By Laura E. Hunter

While the Distributed File System (DFS) has been around since the days of Windows NT 4.0, improvements made to DFS in Windows Server 2003 R2 have made it a more attractive storage feature to Windows systems administrators.

In particular, major advances in the area of replication and file availability now make DFS an ideal solution for low-bandwidth branch office environments. Let's discuss the enhancements made to DFS in R2, and how DFS can meet the ever-increasing storage needs of your organization.

The original intent of DFS was to make file access across multiple file servers more transparent to the users in your organization. Consider a typical small network environment that has three file servers named FS1, FS2 and FS3, containing the following file shares:

\\FS1

\\FS1\accounting

\\FS1\marketing

\\FS1\training

\\FS2

\\FS2\hr

\\FS2\payroll

\\FS3

\\FS3\systems

\\FS3\production



DFS terminology has changed in R2. You'll find the new DFS lexicon on p. 10.

**IMPROVEMENTS
MADE TO DFS IN
WINDOWS SERVER
2003 HAVE MADE
IT A MORE
ATTRACTIVE
FEATURE
TO WINDOWS
SYSTEMS
ADMINISTRATORS.**

In this environment, users who require access to both the Accounting and the Payroll shares (or even two shares on the same server) would need to maintain and remember two separate connections, either by manually specifying the Universal Naming Convention (UNC) path of each share or by mapping two separate drive letters within a logon script.

This can become clumsy for users who need access to many different file shares, particularly if the locations of those shares need to change over time. For example, if the FS3 file share is running out of space and you need to move the Systems share to the new FS4 server, you'd have to communicate this

DFS greatly simplifies the view of shared folders on your network.

change to your users or modify the necessary logon scripts.

You can improve this situation by deploying the Distributed File Service, which creates a unified logical namespace across multiple physical file servers. In our example, by deploying DFS, you can create a single DFS root that can then reference multiple file shares underneath it. A DFS root takes the format of `\\<Domain Name>\<Root Sharename>`. For example, within the company.com domain, we can create a domain DFS root called `\\company.com\shared`, then create DFS links to the shares stored on the three physical servers as follows:

```
\\company.com\shared
\\company.com\shared\Accounting
\\company.com\shared\Marketing
\\company.com\shared\Training
\\company.com\shared\HR
\\company.com\shared\Payroll
\\company.com\shared\System
\\company.com\shared\Production
```

As you can see, this greatly simplifies the view of shared folders on your network. Your users can now specify a single UNC name to access all the shares configured beneath it or have a single drive letter mapped within a logon script.

If you have to move the Systems share from one physical server to another, its DFS link will remain the same regardless of its new physical location. This gives you flexibility in serving up shared files to your users, since you're no longer tied to the physical location of a file or folder when providing access to it.

IMPROVEMENTS TO DFS IN R2

In R2, Microsoft split the DFS service into two components: DFS-Namespaces (DFS-N) and DFS-Replication (DFS-R). This division allows you to make a more granular decision about

the services you deploy on your file servers. If you only use a unified namespace, you can skip installing the DFS-R component unless (and until) your environment expands to the point that it requires replication capabilities.

The DFS Namespaces feature in R2 offers the following updated capabilities:

TARGET PRIORITY. If DFS detects that a particular link target or folder target is inaccessible, it will automatically route cli-


ents to another target server. In previous versions of the server OS, if you had multiple link targets specified for a particular link, you could not specify the order in which referrals should take place. In R2, you can specify a priority list of targets that the client will be referred to.

CLIENT FAILBACK. In previous versions of DFS, if a client was routed past an unavailable link target to another in the list of link targets, the client would continue to use that server until the client was rebooted or until its *DFS referral cache* was cleared. In R2, clients can fail back to a preferred local server once its availability has been restored. However, client failback is available only for clients running Windows XP SP2 or Windows Server 2003 SP1 (including R2). Both OSes require a hotfix that's available from Microsoft.

NOTE: You can configure client failback for an entire namespace, and this setting will be inherited by every folder within the namespace. You can also configure client failback for only specific folders and their folder targets.

DELEGATION OF AUTHORITY.

With DFS-Namespaces in R2, you can delegate the ability to create namespaces, as well as the ability to administer existing namespaces, by setting the necessary permissions within Active Directory (for a domain-based namespace) or in the server Registry (for a stand-alone namespace). By default, you need Domain Admin rights to manage domain-based namespaces or to be a member of the local administrators group of the server that's hosting a standalone namespace.



The DFS Namespaces feature offers updated capabilities in the areas of target priority, client failback and delegation of authority.

DFS-R: THIS REALLY ISN'T YOUR FATHER'S DFS

In R2, it's in the realm of DFS-Replication (DFS-R) that the new DFS really begins to shine. A new replication algorithm provides incredible performance gains for bandwidth-challenged environments such as a branch office separated from corporate headquarters

Because File Replication Service traffic is not compressed even when traversing site boundaries, replication of large files can be a tricky process, one that often creates replication errors or inconsistencies.

by a low-speed or heavily used WAN link.

Prior to R2, DFS used the File Replication Service (FRS) to replicate files between multiple link targets. FRS is the service used to replicate the information stored in the AD SYSVOL share: logon/logoff scripts and Group Policy Objects.

FRS uses RPC over TCP/IP to replicate files within a single site as well as between sites; FRS creates its own replication topology with its own schedule and connection objects that are controlled separately from AD replication.

FRS will trigger replication whenever a file is closed, with changes held in a 3-second aging cache to allow for files that are changed frequently. Once this 3-second "waiting period" is up, the FRS service on the server hosting the changed file notifies its FRS replication partners, and the

file is replicated across the FRS replication topology.

For small files stored on lightly used servers, this process works quite well. But because FRS traffic is not compressed even when traversing site boundaries, replication of large files can be a tricky process, one that often creates replication errors or inconsistencies.

DFS in R2 changes all of this by introducing a replication algorithm called Remote Differential Compression (RDC). RDC breaks files up into small chunks, then replicates only the individual chunks of a file that have changed from one replication cycle to the next.

Consider a Microsoft Word file that contains the line of text: The quick brown fox jumps over the lazy white dog. If someone changed the sentence to read "The slow black fox," then RDC would replicate only that particular chunk, rather

than sending the entire .doc file across the wire. RDC does this by computing MD4 hashes of these small chunks of files, then comparing those hashes between servers that are attempting to replicate.

If a particular chunk of a file has changed, the MD4 hash of that chunk will change while the hashes for the remaining chunks in the file remain the same. This allows an RDC-enabled server's replication partner to request only those chunks whose hashes have changed since the last time

replication took place. For larger files that only need to replicate small changes, this will reduce replication time drastically, while improving performance for your users.

Imagine a 4 MB Visio document in which you need to change the title of one or two sections. In the FRS world, that would prompt the entire 4 MB file to replicate. However, RDC needs only a few seconds to replicate the changed sections of the file. For environments with branch offices to support, particularly where bandwidth is at a premium, DFS-R in R2 can more than justify making the move to the new OS.

NOTE: Because DFS Replication is triggered on file close, it's not efficient for replicating files that are always locked and in use, as in the case of a database or another file used by an "always-on" service. Nor can you

use DFS-R to replicate the AD SYSVOL share; Logon scripts and Group Policy Objects still need to be replicated via FRS. However, FRS and DFS-R can co-exist comfortably on the same server.

Now let's examine some of the other improvements in DFS-Replication in Windows Server 2003 R2:

BANDWIDTH THROTTLING AND REPLICATION SCHEDULING. To gain more control over the use of your bandwidth, you can specify replication schedules similar to those you'd set up between sites in AD. You can specify these schedules for an entire replication group or create a custom schedule for an individual replication connection. You can also cap on the amount of bandwidth that DFS-R replication can take up.



Improvements to DFS will give you more control over the use of your bandwidth as well as help you collect data from branch sites to perform centralized backup.

SUPPORT FOR REPLICATION

GROUPS. You can configure one or more sets of data and servers as a replication group with a common configuration for replicated folders, replication schedules and bandwidth throttling. Each DFS server can support a maximum of 256 replication groups, and each of these groups can contain up to 256 replicated folders.

COLLECTING DATA FOR BACKUP

PURPOSES. You can use replication groups to collect data from branch sites to perform centralized backups. Rather than relying on remote sites to maintain their own backup hardware and perform their own backups, you can create a separate replication group to replicate their data to a central location. By disabling replication from the hub site back to the branch server, you'll create a "one-way" replication agreement that pre-

vents any inadvertent changes made at the backup site from replicating back to the remote server.

NOTE: DFS-R can replicate data across multiple forests within the same forest; you're not restricted to replicating within a single domain.

CROSS-FILE RDC. This takes the performance improvement of RDC to the next logical level. Say you have a file stored in a DFS namespace called **2006 Board of Directors.doc** detailing the names and biographical information of your company's board for that year. You need to create a similar file for the 2007 board, so you save the 2006 file as **2007 Board of Directors.doc** and make a few changes to reflect two new board members.

Now there's a new file that needs to be replicated within

the DFS namespace. But is it really brand new? By using cross-file RDC, DFS can use the contents of the 2006 Board of Directors file to seed replication for the new file, using the "chunking and hashing" method already described to send over the wire only the information that's different between the two files. (This feature is possible because comparing the MD4 hashes created by two files is far more efficient than comparing the actual contents of the files.)

FILE AND SUBFOLDER FILTERS.

You can specify individual subfolders or filenames that should not be included in DFS Replication, either by explicitly listing the name of the file or folder or by using the * wildcard symbol. By default, DFS-R will not replicate any folder that begins with the tilde (~) character, as well as any files with a .TMP file

extension.

Other files and file types that will always be excluded from DFS Replication include:

- ✦ any EFS-encrypted files;
- ✦ any file that has had the temporary attribute set;
- ✦ any reparse points used by Single Instance Storage or Hierarchical Storage Management (The reparse points used by DFS itself are not affected by this.); and
- ✦ any NTFS-mounted drive paths where you've added a new drive to a system and assigned its space as a folder within an existing drive letter, rather than assigning it a drive letter of its own.

LIMITATIONS STILL EXIST

Despite the many improvements to DFS in R2, the service still has its limitations. You should be aware of these in determining

whether DFS is appropriate for your environment.

Perhaps the most important limitation is this: DFS Replication is suitable for environments that can tolerate a certain loose consistency between different copies of a particular document. Even given the performance enhancements of the Remote Differential Compression replication, there will still be a slight amount of inconsistency between servers in a replication group as a change is replicated between them.

Consider this situation: You've configured a replication schedule that only allows for DFS Replication between 11 p.m. and 5 a.m. for a site whose WAN link is fairly saturated during the day. If a user makes a change to a file stored on ServerA at 1 p.m. on a Tuesday, and another user accesses the copy of that file that's stored on ServerB at 2 p.m. on the same day, the second user

DFS still has limitations, which may determine whether it's appropriate for your environment.

will not see the changes that were made on ServerA. The changes have not replicated yet.

Even if you allow DFS Replication to take place 24x7, limitations in network transmission speed mean that ServerA's copy of a document might not be precisely in sync with ServerB's copy at any given second. If you're working with documents that have zero tolerance for this type of loose consistency, you might want to consider a document collaboration solution (one that allows for document check-in and check-out), such as Microsoft SharePoint. But in most cases, the replication improvements offered by R2

DFS will provide an easy way of synchronizing files across multiple locations.

But what happens if two people create a replication conflict by managing to modify the same file at the same time while working from two separate servers? Similar to AD replication, the RDC algorithm resolves conflicts by taking the “last writer wins” approach: Whichever file was modified most recently is the one that will win the conflict and be replicated throughout the namespace. The “losing” file will be renamed and stored in a **Conflict and Deleted** folder on the server that processed the conflict. Details of the file will be stored in a ConflictandDeletedManifest.xml file. This folder has a default quota of 660 MB and will be automatically purged when its size reaches 90% of that limit.

ABOUT THE AUTHOR

Laura E. Hunter (CISSP, MCSE: Security, MCDBA, Microsoft MVP) is employed as an Active Directory architect. She is a two-time recipient of the prestigious Microsoft “Most Valuable Professional” award in the area of Windows Server-Networking. She is the author of the Active Directory Field Guide, published by Apress.com.



WINSTORAGE

is brought
to you by

SearchWinComputing.com. The stories Not Your Father's DFS, Storage Management Tasks that are a Waste of Time, and Good News/Bad News to Dropping Price of Storage all originally appeared on SearchWinComputing.com.

EDITOR

Peter Bochner

COPY EDITOR

Marty Moore

DESIGN DIRECTOR

Ronn Campisi
www.ronncampisi.com

DFS terminology has changed in R2. To see what the new DFS features can do, you need to know its new vocabulary.

NAMESPACE ROOT (formerly DFS Root) provides the “starting point” of a DFS environment. (The term is also used to refer to the DFS environment as a whole.) A namespace root can have one or more *namespace targets* (formerly *root targets*), which are physical servers that host the configuration information for a particular namespace. In our example, `\\company.com\shared` is a namespace root. On each namespace target, DFS creates a physical folder that corresponds to the name of the namespace root.

DFS NAMESPACE refers to the overall “view” of a DFS environment as it appears to the user. You can create either a domain-based namespace or a standalone namespace.

A domain-based namespace is published to Active Directory (AD) and supports the file replication and fault tolerance features. A standalone namespace

stores its configuration information in the Registry of the namespace target that hosts it. Standalone namespaces do not integrate with AD. A domain namespace can have multiple namespace targets configured for the same namespace, with the configuration information for the namespace replicated to each domain root as a part of AD replication. A standalone namespace can only have a single namespace target configured.

DFS FOLDER (formerly DFS link) is the portion of the DFS namespace that maps to a link target. In the current example, `\\company.com\shared\training` is a DFS folder. On each namespace target, a physical folder is created that corresponds to the name of the DFS folder.

This is a reparse point used by Windows to redirect the client to the appropriate location. If users try to access this folder

Return to page 2

directly, they'll receive an error.

NOTE: When planning for DFS scalability, Microsoft recommends that you configure no more than 5,000 DFS folders within a domain-based namespace or up to 50,000 within a standalone namespace. The limitation for domain-based namespaces is designed to limit the space taken up by DFS configuration information within the AD database.

DFS FOLDER TARGET (formerly DFS link target) refers to the physical location that a DFS folder points to. The `\\fs1\training share` is the folder

When planning for DFS scalability, Microsoft recommends that you configure no more than 5,000 DFS folders in a domain-based namespace or up to 50,000 in a standalone namespace.

target of the `\\company.com\shared\Training` DFS folder. In a domain-based namespace, you can configure a single DFS folder to point to multiple folder targets, and the DFS service will replicate changes between the different physical locations.

DFS REFERRALS is the process that DFS uses to take a DFS link submitted by a client and to point (refer) them to the correct physical server. A little-known feature of DFS is that referrals can be enabled or disabled for a particular link without actually deleting the link. This is particularly useful if you have multiple folder targets defined for a particular folder and you need to bring one of the targets offline for maintenance. Rather than allowing DFS to continue to refer clients to the offline server, creating an error for the user, you can temporarily configure DFS to refer clients to only those link targets that are still operational.

DFS REFERRALS are also site-aware, meaning that if a DFS folder has folder targets in multiple cities, DFS will refer any connecting clients to the closest folder target based on the sites and subnets configured in AD.

GET YOUR LIFE BACK

AUTOMATE DFS-RELATED TASKS WITH BROCADE STORAGEX



If you use or plan to use Microsoft Distributed File System (DFS), you'll love Brocade StorageX.

Brocade® StorageX® takes a good thing and makes it even better—building on DFS to give you the industry's most comprehensive file management solution. It can help you automate key tasks such as:

- Easily migrating and consolidating data
- Optimizing backup/replication at remote offices
- Providing high availability for critical data

Brocade StorageX provides a wide range of additional capabilities, such as:

- Building and managing enterprise-class DFS roots
- Integrating with both CIFS and NFS
- Integrating with and managing NetApp filers
- Auditing actions, operations, and administrative policies
- Integrating logical-physical file views
- Supporting namespace backup/restore
- Automating creation of the DFS consolidation root
- Replicating to shared cluster storage
- Providing local group processing for files and directories
- Automating storage tiering

TRY BROCADE STORAGEX YOURSELF!

E-mail bcrgroup@brocade.com today to get your free evaluation download!

“Brocade StorageX is a simple, elegant solution to the complex problem of file system management. It merely rides on top of what is already there. The price is such that a quick, back-of-the-envelope ROI calculation should yield a solid multiplier.”

— *Michael Fisch, Director of Storage and Networking, The Clipper Group*



BROCADE

good news/bad news to dropping price of storage

By Peter Bochner

The continually dropping price of disk storage has made storage managers and network administrators very happy. But can storage get too cheap?

The continuing price drop may possibly have a downside, according to Jacob Farmer, chief technology officer for Cambridge Computer Services, Inc., an integrator specializing in storage solutions, headquartered in Waltham, Mass.

Farmer says that as disk drives get bigger and prices get cheaper, some IT folks out there may be getting a bit sloppy.

"The good news about inexpensive disk," says Farmer, "is that first of all, you don't have to worry about managing it nearly as much, and secondly, it means you're redundant - you have more copies of your data every which way from Sunday."

However, the downside to the dropping cost of storage is a certain wastefulness. "I've seen people putting in centralized storage systems that are really ill-conceived," Farmer said.

"The bigger the hard disk,

the worse it performs," he said. As an example, he pointed to the storage manager who gets excited he "can buy a 750-gig hard drive for a few hundred dollars, and can build a RAID array with that and get a couple of terabytes, divide that up and share that up among all his servers."

Sure, said Farmer, that manager is handling all his capacity needs. "But from a performance and reliability standpoint, he's running on much less robust gear than when he had direct attached storage. And he has far fewer disk spindles."

In short, Farmer said, some IT folks out there can buy capacity for so little money they forget what it really costs to build a robust high-performance storage system.

• Peter Bochner is the site editor of SearchWinComputing.com

storage management tasks that are a waste of time

By Serdar Yegulalp

For systems administrators, some tasks related to storage and storage management aren't worth the time and effort to perform them. Here's a list of **five tasks** that you get away with cutting out of your schedule.

- Serdar Yegulalp is editor of the Windows Power Users Newsletter, which is devoted to hints, tips and tricks for Windows 2000, NT, XP and now Vista users and administrators. He contributes regularly to SearchWinComputing.com and other TechTarget Web sites.

1 FULL FORMATTING A DRIVE. Once upon a time, putting a drive into service that hadn't had every single sector physically tested was madness. Go back far enough and you can find drives that came from the factory with a defect list pasted on the top of the drive, so you could pass appropriate instructions on which sectors to exclude during the format.

But each new generation of hard drives have become more thoroughly spec'd out and reliable, to the point where it's no longer necessary to physically test a new one. Most drives—especially those in high-end RAID arrays—have self-checking mechanisms that can automatically detect physical problems and relocate data if it's at risk. If

you're determined to waste several hours performing a disk test that might simply be redundant, it won't do any harm, but it'll sure slow you down. (Check your drive array's documentation for the straight dope on how new drives are provisioned.)

2 DEFRAGMENTING WORKSTATIONS MORE THAN ONCE A WEEK. On a system less than three years old, defragmenting more than once a week in a workstation-type environment doesn't provide any justifiable benefits. Defragmenting a workstation more than once a week is probably not going to help.

The real performance killer is fragmentation *plus* low free space; workstations with drives that are more than 75% full need to either be cleaned off or upgraded. Servers, on the other hand, *can* benefit from being defragmented more aggressively, but only when it's not at the expense of performance. Defragment servers during off-peak hours (e.g., 4 a.m.) to keep the process from slowing other things down.

3 PERFORMING SPURIOUS OFFLINE DEFRAGMENTATION OF EXCHANGE DATABASES. Exchange 2000 and 2003 defragment themselves internally once a day, at 2 a.m. Some Exchange administrators get real twitchy about the amount of space used up by the Exchange database, since the only way to compact the database files is to run ESEUTIL (which essentially recreates the database in an entirely new file). This takes 1 GB to 7 GB of database space per hour (estimates vary widely), and you're going to have to take the whole database offline to do it.

Unless there's an overriding reason to run ESEUTIL—i.e., as part of a larger error-checking or crash-recovery operation, or when the database can't be started—and unless you have a current backup of the database, it's a waste of time to run it just for the sake of reclaiming free space. One way to determine if an offline defrag will be worth it is to inspect the application log on the Exchange server and look for event 1221, which contains an estimate of how much free space might be recovered.

Unless there's an overriding reason to run ESEUTIL, it's a waste of time to run it just for the sake of reclaiming free space.

4

USING CHKDSK /R TO PERFORM SURFACE TESTS ON SOME RAID ARRAYS. The /R option in CHKDSK performs a surface test in an attempt to locate bad sectors and recover any data in them. It sounds like a good idea, but it's really slow. On some RAID arrays, it's both redundant and slow. For instance, the HP StorageWorks 1000 performs background surface tests on connected disks, and moves out data if it finds a bad sector to prevent future problems. Running surface tests through CHKDSK on such a drive is like polishing a no-wax floor. If the manufacturer has surface-test tools of its own, use those instead.

5

USING MEMORY "OPTIMIZER" UTILITIES. I'm amazed how often I hear people talking about this one. A bunch of utilities—some freeware, some shareware, some commercial software—claim to "defragment your system memory," mostly by allocating and then deallocating large blocks of physical memory. *Do not use them!* Why? Because they try to second-guess the way the memory manager works in Windows, which simply trades one set of problems for another. If you need more physical memory, buy it. Memory is cheap.

Time-saving tips

- Don't perform a disk test on today's drives.
- Don't defrag workstations more than once a week.
- Don't run ESEUTIL on Exchange databases just for the sake of reclaiming free space.
- Don't use the /R option in CHKDSK on RAID arrays.
- Beware utilities that claim to optimize your system memory.

Avoid utilities that try to second-guess the way the memory manager works in Windows.



**FILE AREA
NETWORK**

Optimizing and Managing File Storage in Windows Environments

A Powerful Solution Based on Microsoft DFS and Brocade StorageX.

The Microsoft Distributed File System (DFS) is a storage service that can help solve many complex problems in Windows environments. Organizations that deploy DFS can achieve immediate and significant file management benefits. In addition to making it easier to manage files, DFS provides a core set of capabilities on which powerful enterprise storage management solutions can be built.

Brocade® StorageX® is a solution for managing distributed file storage in Windows environments. Built on DFS, Brocade StorageX enables the integrated management of logical and physical storage elements, making it the most comprehensive Windows storage management solution available.

INTRODUCTION

Organizations in every industry are experiencing explosive growth in their requirements for file storage capacity. The proliferation of application data and user-generated documents—such as presentations, spreadsheets, graphics, and scanned documents—is driving continued expansion of network storage requirements. As a result, Windows storage configurations are growing in both size and complexity. This growth and complexity creates challenges for storage administrators as well as users. And it is a direct cause of increased storage management costs, as well as sub-optimal storage utilization and data availability.

ENTERPRISE STORAGE MANAGEMENT CHALLENGES

Storage requirements are outpacing most organizations' ability to efficiently manage them. There are four key challenges relative to the growing requirements for file storage, described in the following sections:

- **Explosive growth in enterprise storage:** Many factors are contributing to an enormous increase in network storage requirements, including the pervasiveness of wide area networking and users' ability to create and duplicate huge volumes of data.
- **Underutilization of distributed storage:** According to a Gartner Group study, only 30 to 40 percent of storage in distributed environments is utilized, which is low compared to 80 percent storage utilization in mainframes.
- **Complexity in distributed storage management:** Another Gartner Group study found that storage and storage-related management consumes as much as 75 percent of corporate IT budgets, and for every dollar spent on storage hardware, much more is spent on administrative costs. Given the fact that storage requirements are increasing as much as 100 percent annually, organizations cannot afford to continue “throwing hardware” at the problem or hiring more administrators to handle the additional workload.
- **Achieving cost-effective high availability for distributed storage:** Organizations face a daunting task trying to make data that is spread throughout the enterprise highly available. Doing so requires a well-planned, well-designed network storage architecture that can be centrally managed. Until now, it has not been possible to centrally manage distributed storage in Windows environments, so organizations have found it very difficult to create and manage distributed file storage environments.

DFS and Brocade StorageX provide a means for addressing and solving many of these storage management problems. Based on DFS, Brocade StorageX

is the most powerful automated storage management software for Windows environments. It can significantly enhance storage availability and scalability while reducing the cost and complexity of storage management.

LAYING THE FOUNDATION FOR FILE STORAGE MANAGEMENT WITH DFS

Historically, with the Universal Naming Convention (UNC), a user or application was required to specify a physical server and share in order to access file information (that is, the user or application had to specify `\\Server\Share\Path\Filename`). A UNC is typically mapped to a drive letter where `x:` might be mapped to `\\Server\Share`. Users must know the physical name/location of the machines and shares they want to access. And users must map to many drive letters (`d:`, `e:`, `f:`, `g:`, `h:`, etc.) to access information stored on different servers. As a result, users can be overwhelmed by complicated share names and the number of places where data can be stored.

This storage management paradigm creates major problems for both users and administrators. The administrator must forever maintain the server name once it is published to users. Changing a `\\servername\share` name requires the administrator to notify each user (perform “desktop touches”) so that their machines can be reconfigured to access the renamed server and shares. With this approach, administrators have very little flexibility in the way they manage users and storage. Machine name dependencies lead to inefficient administration and underutilized storage, both of which increase costs and reduce the effectiveness of storage management. However, DFS services in Windows Server 2003 and Windows 2000 provide a significant change to the traditional storage management paradigm.

A Platform for Distributed File Management

Windows Server 2003 and Windows 2000 include DFS, a storage service that can help solve many storage management challenges. DFS enables the creation of a logical file system that can unify multiple physical file systems. As a result, organizations can use it to improve storage-related functionality, reduce costs, and provide more comprehensive storage management.

DFS also is a strategic file server virtualization infrastructure on which software vendors can build value-added storage management solutions. DFS consists of a client 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 2003. 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.

DFS does for servers and shares what file systems do for hard disks. File systems provide uniform named access to collections of sectors on disks. Likewise, DFS provides a uniform naming convention and mapping for collections of servers, shares, and files. Another way of thinking about it is that DFS is to file storage what DNS is to networking or Active Directory is to users and computers.

A Storage Solutions Development Platform

Because DFS is a component of Microsoft Server, vendors can easily build storage solutions based on open standards with DFS at the foundation. Vendors such as Brocade significantly leverage the DFS platform to create solutions that address critical storage issues such as disaster recovery, data migration, server consolidation, storage reconfiguration, and storage optimization.

The Compelling Case for Deploying DFS

DFS lays the foundation for enterprise file storage management by providing a core set of capabilities. In particular, DFS:

- **Eliminates machine name dependencies:** DFS removes the “once-published, forever-maintained” requirement for server names, which creates tremendous flexibility for administrators to add or move files and users without having to touch or reconfigure the desktop. Removing machine name dependencies enables the creation of a logical storage layer, which in turn enables the creation of strategic enterprise file storage solutions.
- **Enables a logical view of physical storage:** DFS separates the logical and physical aspects of storage to enable the creation of a logical layer. This means that administrators can create logical views of physical storage that match what users and applications want to see instead of how and where files are physically stored.
- **Protects investments in Windows software:** Because DFS is already included in all Windows clients and every Windows server product since Windows NT 4.0, organizations can take advantage of DFS without having to purchase, load, or deploy additional operating system software.
- **Protects investments in storage hardware:** DFS supports file storage across multiple storage types—NAS, DAS, server-attached SAN—from various vendors. This means that organizations can use DFS to aggregate and increase the utilization of their existing storage devices, which can significantly reduce hardware costs.
- **Enables the development of enterprise storage solutions that are easy to deploy:** With DFS as the foundation, vendors can develop storage solutions that reside above the operating system and that do not require kernel-level

software. This means that organizations can quickly and easily deploy these new storage solutions, as they will not have to retest applications or load new system-level software.

Brocade StorageX Adds Value to DFS

DFS is an immensely valuable service, but it is not a complete solution. Brocade StorageX adds value to DFS because it:

- Is a complete file virtualization solution that integrates management of the -logical and physical layers
- Provides comprehensive DFS management
- Makes DFS enterprise-ready by enabling its use in complex, large-scale - Windows environments
- Uses DFS to deliver a storage management solution

BROCADE STORAGEX IS A COMPLETE MANAGEMENT SOLUTION

Brocade StorageX is a complete solution for managing large-scale file storage in Windows environments. Built on DFS, Brocade StorageX enables organizations to solve some of their most pressing storage management problems, such as aggregating heterogeneous NAS devices so that they can be managed as a single storage pool, managing user storage throughout its lifecycle, migrating data, consolidating servers and storage, and providing near-continuous access to files in the event of a disaster.

Storage Transparency: What Does It Mean?

In the simplest terms, transparency is achieved by separating the logical and physical aspects of storage and enabling them to be scaled and managed separately. In other words, Brocade StorageX breaks the hard dependencies between file access and physical file locations.

The goals of this transparency are to shield users from the complexities of the storage architecture, and to enable administrators to manage the physical layer without affecting users’ access to data. In storage management terms, this is the act of “pooling” distributed storage so that it can be viewed and managed as a single unit. Brocade StorageX enables pooling of file storage across multiple, heterogeneous storage types (DAS, server-attached SAN, or NAS) from various vendors—increasing the flexibility to manage data. It also enables the administrator to manage both the logical and physical layers in an integrated fashion.

Integrated Management

Although the benefits of separating logical and physical storage components are well documented, the creation of a logical layer adds administrative burden. The key to solving storage management problems is to enable their integrated management. This means automatic synchronization between the logical and physical layers so that changes to physical files are automatically updated in all the logical links that reference those files. This integration reduces the administrator’s burden and increases the flexibility to manage users and files.

Brocade StorageX Uses DFS to Deliver an Enterprise Storage Solution Platform

Brocade StorageX adds several key services and applications to DFS to deliver an enterprise storage solution platform. Brocade StorageX includes the following:

- DFS management
- Windows storage management
- Enterprise namespace creation
- Data migration services
- Replication services
- Disaster recovery management
- Active Directory integration

By combining these services with DFS, Brocade StorageX provides a complete platform for storage solutions that can increase data availability, improve storage utilization, and reduce storage costs (see Figure 1).

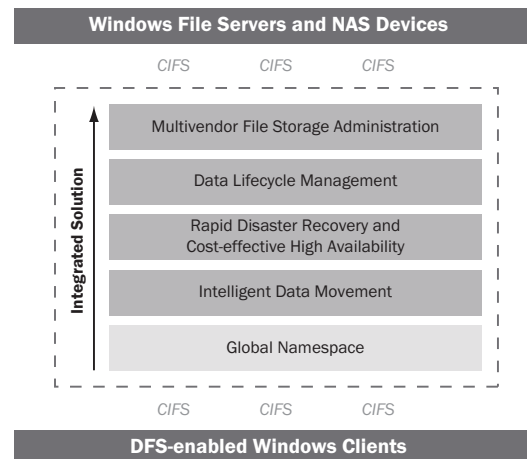


Figure 1. DFS and Brocade StorageX provide an integrated solution.

Brocade Storage X Provides Simple, Comprehensive DFS Management

Brocade StorageX increases the management and usability of DFS in significant ways. It provides a single point of management for distributed roots and enables logical visualization of an entire distributed storage environment. Brocade StorageX also includes powerful graphical tools for DFS server configuration and has a point-and-click, drag-and-drop interface to enhance DFS usability.

Brocade StorageX includes sophisticated management tools to monitor, diagnose, troubleshoot, and manage distributed DFS roots. And Brocade StorageX management reporting and Web-based administration make it easy to administer and monitor DFS networks. Organizations that want to deploy DFS often find Brocade StorageX invaluable because it makes DFS implementation, configuration, and management simple and reliable.

Brocade StorageX Makes DFS Enterprise-Ready

Brocade StorageX includes powerful features that enable DFS scalability and centralized management:

- **Enterprise view of a namespace:** Brocade StorageX provides an enterprise-wide view of file storage. It discovers and displays all DFS roots that exist in one or more namespaces and correlates logical views with physical machines for the entire enterprise.
- **Easy namespace creation:** Administrators can easily create a namespace with Brocade StorageX by using its graphical interface to define how logical links tie to physical storage. With the ease of dragging and dropping information from Windows Explorer and Brocade StorageX utilities such as Share Finder, an administrator can define an entire namespace in just minutes.
- **One-to-many management of DFS configurations:** Brocade StorageX provides a single point of management for multiple distributed DFS roots, which means that an entire namespace can be managed from a single, centralized location. Powerful features in Brocade StorageX enable an administrator to perform a single action simultaneously on groups of DFS roots, which provides scalability and reduces the amount of time required to set up and manage namespaces.
- **Cross-reference of logical views and physical storage:** Brocade StorageX provides a graphical view of DFS roots and links, and logical-to-physical as well as physical-to-logical views. The administrator can see which users and applications are dependent on specific machines. This ability to visualize logical views correlated with physical configurations gives the administrator the flexibility to change/manage physical storage without disrupting users' access to files.

- **Data migration and integrated management:** The Brocade StorageX Data Migration Wizard allows the administrator to migrate files from one storage device to another with all of the file attributes intact. In addition, it automatically updates the namespace and synchronizes the logical and physical layers.

CONCLUSION

Brocade StorageX is the most powerful automated storage management software available today for managing Windows storage. It can significantly enhance storage availability and scalability while reducing the cost and complexity of storage ownership—thereby benefiting both administrators and users.

As a DFS-based solution for managing large-scale storage in Windows environments, Brocade StorageX provides an easy way for administrators to create and manage both the logical and physical storage layers. It allows them to:

- Better utilize distributed storage capacity
- Gain flexibility by shielding users from changes in the storage implementation
- Easily support and administer heterogeneous storage devices
- Protect users from the complexity of the storage infrastructure
- Increase data availability
- Seamlessly accommodate growth and changes in storage requirements
- Decrease storage management costs

For more information, visit www.brocade.com.

Corporate Headquarters

San Jose, CA USA
T: (408) 333-8000
info@brocade.com

European Headquarters

Geneva, Switzerland
T: +41 22 799 56 40
emea-info@brocade.com

Asia Pacific Headquarters

Singapore
T: +65-6538-4700
apac-info@brocade.com

© 2007 Brocade Communications Systems, Inc. All Rights Reserved. 01/07 GA-WP-835-01

Brocade, the Brocade B weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and the Brocade B wing logo and Tapestry are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

