Exploring the Distributed File System

in Microsoft Windows Server 2003 R2

The Distributed File System (DFS) in Microsoft® Windows Server™ 2003 Release 2 (R2) introduces DFS Namespaces and DFS Replication. These two enhancements are designed to help administrators manage distributed file server resources efficiently while enabling fast, fault-tolerant access with low-bandwidth replication.

BY MIN-JOHN LEE AND MAHESH VELLORE

Related Categories:

Microsoft Windows Server 2003

System deployment

Visit www.dell.com/powersolutions for the complete category index.

Microsoft Windows Server 2003 Release 2 (R2) is the first release update for Windows Server 2003, and its features are designed to integrate seamlessly into the Windows Server 2003 environment. At the core of Windows Server 2003 R2 is the Distributed File System (DFS), which has been enhanced with the goal of providing increased performance and improved availability while reducing issues commonly linked to branch-office server deployments.

For example, users may have trouble locating the files they need on distributed file servers. In addition, administrators often contend with low-bandwidth connections for providing file replication and file availability over wide area networks when servers at the central or branch offices fail. To help resolve such concerns, Windows Server 2003 R2 provides DFS Namespaces and DFS Replication, which together enable fault-tolerant access to distributed files with low-bandwidth replication:

- **DFS Namespaces:** Formerly known as the Distributed File System, the DFS Namespaces service allows administrators to group shared folders that are located on different servers. Folders are presented as a virtual tree called a namespace, so users no longer have to remember physical file locations.
- **DFS Replication:** A follow-up to the File Replication Service that was introduced in the Microsoft Windows® 2000 OS, the DFS Replication service is designed to address low network bandwidth using the remote differential compression (RDC) algorithm. RDC enables DFS Replication to

transfer only the changes that have been made since the last file update. Another feature, cross-file RDC, identifies files that are similar to the one being replicated. DFS Replication can then use portions of those similar files to replicate the file, helping reduce the amount of data transferred over the network.

Planning a DFS deployment

Administrators must consider a wide variety of factors when planning a DFS deployment. Major concerns include deployment of the Microsoft Active Directory® service, connection availability between branch offices, bandwidth availability between branch offices, file-share access in heterogeneous network environments, cost-effective backup procedures, capacity planning and scalability requirements, regional failover, and file types. By weighing alternatives carefully, administrators can determine a suitable deployment strategy for their particular enterprise requirements.

Active Directory considerations. Some DFS features exist only in the Active Directory environment. DFS Namespaces can be stand-alone or domain based. Stand-alone DFS Namespaces should be selected if an organization does not use the Active Directory service or if it needs to host more than 5,000 folders within a single namespace. Because of performance and directory service synchronization requirements, Microsoft best practices recommend using domain-based DFS Namespaces to host up to 5,000 folders. If using a server cluster for high availability, an organization must choose stand-alone DFS Namespaces.

By weighing alternatives carefully, administrators can determine a suitable deployment strategy for their particular enterprise requirements.

The DFS Replication service must be deployed in an Active Directory environment because the RDC technology can function only in an Active Directory domain. Active Directory enables DFS Replication to store configuration objects and to delegate user rights precisely.

Connection availability between branch offices. In an environment where the network connectivity could be either unavailable due to maintenance or unreliable due to the quality of the Internet service provider, DFS allows administrators to configure multiple target file servers for a given namespace folder. In this way, administrators can configure target priority so that data access requests fail over to a desired next-available target when the primary target cannot be reached. This feature enables zero-downtime server maintenance without interrupting data availability. In addition, once the primary target comes back online, data access requests can automatically route back to the primary DFS server by enabling client failback.[1]

Bandwidth availability between branch offices. In many circumstances, branch offices are connected to the enterprise's hub file server using on-demand or permanent virtual private network (VPN) connections over the Internet. In such network environments, data synchronization incurs a certain minimum latency due to network bandwidth and availability. In Windows Server 2003 R2, the DFS Replication service's RDC algorithm detects changes made by users in the replicated files and sends only changes that have been made since the last update. This approach consumes considerably less network bandwidth than the full file replication used in previous versions of DFS, and the reduced network overhead in turn helps improve the reliability of file replication in a low-bandwidth environment.

File-share access in heterogeneous network environments. For networks supporting both Windows and Linux® or UNIX® operating systems, Windows Server 2003 R2 provides Network File System (NFS) services and a user-mapping service that allows non-Windows platforms to access file shares on systems running Windows Server 2003. The combination of the NFS Client and NFS Server components in Windows Server 2003 R2 with DFS Namespaces and DFS Replication allows Linux- and UNIX-based systems to take advantage of the RDC algorithm and benefit from the high availability these services are designed to provide.

Cost-effective backup procedures. A comprehensive DFS implementation should include backup considerations. In a distributed environment, file updates are scattered across multiple office locations. Careful planning of DFS Replication procedures can help reduce staff requirements at branch locations. When data is replicated to a central DFS server, IT staff is needed only where the enterprise's hub file server is located to perform administrative data backup tasks. Windows Server 2003 R2 DFS Replication allows organizations to establish replication groups and replication topology and to set up filters to help prevent synchronizing unnecessary content such as MP3 files.

Capacity planning and scalability requirements. When planning for storage services, an organization's two major tasks are meeting immediate storage requirements and planning for future growth. Because Windows Server 2003 R2 DFS allows organizations to configure multiple target servers for a given folder, administrators can upgrade the storage capacity of a major DFS namespace server while data access requests fail over to the secondary namespace

[1] To enable support for client failback on a Microsoft Windows XP system, administrators must install Windows XP Service Pack 2 or later as well as the hot fix described in Microsoft Knowledge Base article KB898900. For more information, visit www.microsoft.com/downloads/details.aspx?FamilyID=7d3f51e3-2d33-48c4-8b5f-fe2345b0a35e&DisplayLang=en.

server. As a result, storage capacity can be resized quickly and flexibly in response to immediate business requirements, and administrators can scale enterprise storage capacity in cost-effective increments as future business needs materialize.

Regional failover. DFS in Windows Server 2003 R2 can help organizations provide data redundancy in different geographic areas—for example, replicating to several remote locations as well as to the enterprise's hub file server to maintain high availability and safeguard business continuity despite regional disasters. The RDC algorithm enables data on DFS namespace servers to be synchronized reliably between remote locations, even in low-bandwidth network environments.

File types. Windows Server 2003 R2 DFS is designed to work effectively with file types that are not affected by latency due to network bandwidth and availability issues. For example, file types that require real-time data synchronization, such as transaction-oriented databases, are not particularly well suited for DFS deployments.

Magazine.com DFS topology

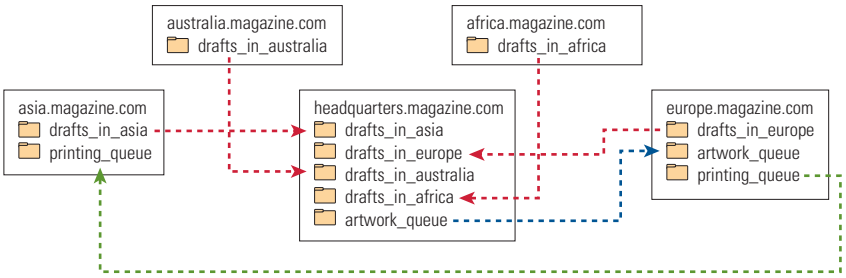

Magazine.com DFS namespace definition

Figure 1. Global document workflow process for example scenarios

Customizing DFS deployments

The three DFS deployment scenarios described in this section demonstrate common branch-office requirements. Each scenario involves a global magazine company that has writers located around the world, article reviewers at the company's headquarters in North America, an artwork team in Europe, and a printing team in East Asia. DFS servers are deployed on each continent to transfer data around the world.

Scenario 1: Implementing business processes

The magazine's remote DFS servers can collect article drafts from writers around the world. A writer in Asia uses a notebook computer to establish a connection to the company's network and maps the DFS namespace—\\magazine.com\drafts_in_asia—in Microsoft Windows Explorer. The writer saves draft articles in the namespace as soon as they are ready for review. On the back end, the draft article is actually saved in a share folder hosted by a regional DFS server called \\asia.magazine.com\drafts_in_asia. This server is defined as a target folder in the \\magazine.com\drafts_in_asia namespace. By doing this, the writer does not have to remember all the DFS server names—only a company domain name. This setup can be duplicated on each continent.

At the North American headquarters, a reviewer uses Windows Explorer to map a drive to the same \\magazine.com\drafts_in_asia namespace. When this connection is established, the DFS Namespaces service on the back end actually establishes a connection to a local target folder hosted in a DFS server located at headquarters. The local target folder resides at \\headquarters.magazine.com\drafts_in_asia. The reviewer can view articles submitted by writers from Asia because a DFS Replication group is defined such that all files in \\asia.magazine.com\drafts_in_asia are replicated to \\headquarters.magazine.com\drafts_in_asia. Organizations can duplicate such a setup to implement a sophisticated document workflow process—for example, a process for legal review.

Once a reviewer approves an article, the article is moved to another namespace called \\magazine.com\artwork_queue. This namespace has target folders at \\headquarters.magazine.com\artwork_queue and \\europe.magazine.com\artwork_queue. The artwork team in Europe maps a drive to the same \\magazine.com\artwork_queue namespace and can view the same content as reviewers by using a DFS Replication group between the two target folders. Once the artwork for an article is completed, the artwork team moves the final version of the article to the \\magazine.com\printing_queue namespace. With a replication group set up between the Europe branch server and the Asia branch server, the printing team in Asia can obtain the final version of the article. Figure 1 shows the DFS topology and namespace definition for this global document workflow process.

Scenario 2: Collecting data for central backup

DFS Replication can be configured to perform data collection from different branch-office locations so that the backup is centralized on the enterprise's hub file server. The IT department at the magazine's headquarters can set up DFS Replication groups so that the data in the shared folders—drafts, artwork_queue, and printing_queue—is replicated from the server at each branch office (data source) to the hub file server (data destination). This approach helps reduce operational costs at branch offices by consolidating hardware for backup and related management tasks. The data collection topology also allows data to be replicated from the hub server to each branch server if data recovery at a branch server is needed. Either setting up shared folder permissions or disabling the hub-to-branch connection in the replication group can prevent data from changing on the hub file server.

Replicating data to a branch-office server enables the branch office to provide fault-tolerant access by establishing failover to an office that is located in a different geographic region.

Scenario 3: Providing regional failover for high availability

The DFS Management snap-in within the Microsoft Management Console (MMC) can be used to configure DFS Replication–specific tasks such as creating replication groups, specifying replication group schedules, managing replication for specific connections, and managing replication filters. Implementing DFS Replication with DFS Namespaces can help achieve high data availability. If the branch server asia.magazine.com reaches its storage capacity, branch clients can fail over to the DFS server at headquarters while the Asia branch server is taken down to let administrators add storage space. With a DFS Namespaces enhancement known as client failback, branch clients can fail back to the branch server after the branch server is back online. Both failover and failback are transparent to end users, thus providing high data availability and business continuity. *Note:* DFS Replication for data collection is not recommended for database files or any other type of file that is held open for long periods of time because such files are replicated only after they are closed.

Enterprises with offices scattered across the globe must provide access to files that may be located at any remote office location. Replicating data to a branch-office server enables the branch office to provide fault-tolerant access by establishing failover to an office that is located in a different geographic region. This approach also helps provide an additional level of data protection if disaster strikes and distant offices are unaffected.

Extending DFS deployments with Storage Manager for SANs

A major requirement for branch offices is that storage be easily configured, provisioned, and managed. Windows Server 2003 R2 introduces the Storage Manager for Storage Area Networks (SANs) component, which helps simplify storage provisioning on external storage arrays. Using Storage Manager for SANs, administrators can host DFS replicated files on Fibre Channel and Internet SCSI (iSCSI) storage subsystems that support Virtual Disk Service (VDS).[2] This approach enables smooth integration into the Windows-based server, because Storage Manager for SANs is controlled through an MMC snap-in.

Enabling fast, fault-tolerant file access across the enterprise

Microsoft Windows Server 2003 R2 has enhanced the Distributed File System to address emerging needs for heightened performance, availability, and reliability of far-flung branch-office file servers while minimizing issues that are commonly linked to branch-office deployments, such as management overhead and limited connectivity. Two key capabilities introduced in Windows Server 2003 R2—DFS Namespaces and DFS Replication—are designed to simplify the management of distributed file server resources while providing easy and fault-tolerant access for remote branch-office locations. The three real-world scenarios described in this article show how DFS technologies enable enterprises to implement processes that help improve productivity while helping ensure high data availability and business continuity.

Min-John Lee is a software engineering consultant in the Server Operating Systems Engineering department in the Dell Product Group–Enterprise Software Development. Min-John has an M.S. in Electrical and Computer Engineering from Northwestern University.

Mahesh Vellore is a lead engineer in the Server Operating Systems Engineering Group at the Dell Bangalore Development Center. He has a diploma in Electronics Engineering and has been at Dell for more than two years.

FOR MORE INFORMATION

Microsoft Windows Server 2003 R2:
www.microsoft.com/windowsserver2003

[2] For more information about Storage Manager for SANs and VDS, visit www.microsoft.com/windowsserver2003/R2/storage/default.mspx.

