



Microsoft Windows Storage Server 2003 R2

IT Pro White Paper

*Microsoft Corporation
Published: March 2006*

Abstract

Microsoft® Windows® Storage Server 2003 R2 provides organizations of all sizes a more cost-effective solution to storage management. Windows Storage Server 2003 R2 provides an industry-leading storage management platform by implementing a set of enterprise-ready storage management features and improvements that allow storage managers to reduce their total lifecycle cost of storage through easier management, faster performance, greater reliability, and stronger security. Windows Storage Server 2003 R2 takes advantage of the stability and security enhancements of the proven code base of the Windows Server™ family of operating systems while extending functionality and control in new areas and is available only to OEMs. Windows Storage Server 2003 R2 offers all of the benefits of Windows Storage Server 2003 with SP1 while greatly improving the search functionality, reducing space consumed by duplicate files, offering new and improved Network File System (NFS), directory quotas, file screening, reporting, new DFS replication and namespace capabilities, and collaboration. Windows Storage Server 2003 R2 is easy to integrate into an existing Windows Server 2003 environment and provides IT service providers increased return on investment in Windows-attached storage over the long term, while offering similar application compatibility, manageability, and serviceability as existing servers that run Windows Server 2003 with SP1. NAS devices running Windows Storage Server 2003 R2 are available from a wide range of Original Equipment Manufacturers (OEMs).

This white paper introduces the reader to the new and enhanced features of Windows Storage Server 2003 R2., illustrating some use cases of the product as well as what the reader can do with it.

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 White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, 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.

© 2006 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, MSN, SharePoint, Windows, the Windows logo, Windows NT, Windows Server, and Windows Server System 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

| | |
|--|-----------|
| FIGURES | V |
| INTRODUCTION TO NETWORK ATTACHED STORAGE | 1 |
| OVERVIEW OF WINDOWS STORAGE SERVER 2003 R2 | 4 |
| Key Features of Windows Storage Server 2003 R2 | 4 |
| WINDOWS STORAGE SERVER 2003 R2 FEATURES | 7 |
| Single Instance Storage | 7 |
| Performance Optimization | 7 |
| Search Enhancements | 7 |
| Windows SharePoint Services | 8 |
| SharePoint sites – file storage including collaboration | 8 |
| SharePoint Central Administration – Web browser interface for managing servers | 8 |
| Efficient Storage Management | 8 |
| Storage Manager for SANs | 9 |
| File Server Resource Management | 9 |
| Storage Resource Manager quotas vs. NTFS disk quotas | 10 |
| Distributed File System | 10 |
| DFS Namespaces | 10 |
| DFS Replication with Remote Differential Compression | 11 |
| Simplified Branch Server Management | 12 |
| Network Security | 12 |
| File Serving across Different Operating Systems | 12 |
| Microsoft Services for Network File System | 12 |
| File Sharing Protocols | 13 |
| Remote Hardware Management | 13 |
| Print Management Console | 13 |
| WINDOWS STORAGE SERVER 2003 R2: PREDICTABLE AND EASY TO ADOPT | 14 |
| SUMMARY | 15 |
| RELATED LINKS | 16 |
| APPENDIX A: SINGLE INSTANCE STORAGE IMPLEMENTATION | 17 |

| | |
|--|----|
| SIS Design | 17 |
| Creating SIS Links | 17 |
| Accessing SIS files | 18 |
| Breaking SIS links | 18 |
| The SIS Common Store and Common-Store Files | 18 |
| The SIS Groveler | 18 |
| Scanning Process | 19 |
| Backup/Restore Support | 19 |
| Using a backup application that does not support SIS | 19 |
| Interoperability with other features | 19 |
| SIS Limitations | 21 |

APPENDIX B: WINDOWS STORAGE SERVER 2003 R2 LICENSING RESTRICTIONS22

Figures

| | |
|--|----|
| Figure 1: Network Attached Storage | 1 |
| Figure 2: Storage Area Network | 2 |
| Figure 3: NAS Gateway on a SAN | 2 |
| Figure 4: Distributed File System | 11 |
| Figure 5: Creating SIS Links with Reparse Points | 17 |
| Figure 6: Breaking a SIS Link..... | 18 |

Introduction to Network Attached Storage

Developments in technologies delivering storage over the network have revolutionized the availability, distribution and accessibility of storage resources. Attaching storage directly to the company network helps to eliminate or reduce the drawbacks business experience when using directly attached storage. These drawbacks include:

- Limited storage capacity and scaling constraints as data assets grow.
- Proliferation of computer hardware and software to meet increased storage capacity needs.
- Proliferation of hardware and software necessary to protect the data allocated to new storage.
- Increased management complexity for the system administrator(s) configuring and maintaining IT client and server systems.
- Poor or inconsistent protection of data on client desktops and laptops.
- Limited or laborious access to data, especially between company workgroups, departments and branches.
- Temporary or permanent loss of data resulting from hardware failures, data corruption or user error.
- Spiraling costs associated with capital expenditure, increased staffing, and production downtime resulting from data loss.

Network attached storage (NAS) is one solution to these challenges. NAS storage and file serving devices can be attached directly to the organization's IP network. Storage directly attached to the network becomes accessible to all computers that can access the network. NAS servers are designed for ease of deployment: they can be plugged directly into the network without disruption of services, management is minimal and simplified, and they are largely maintenance free. NAS devices are an ideal means by which to consolidate file servers and backup equipment and to expand storage capacity.

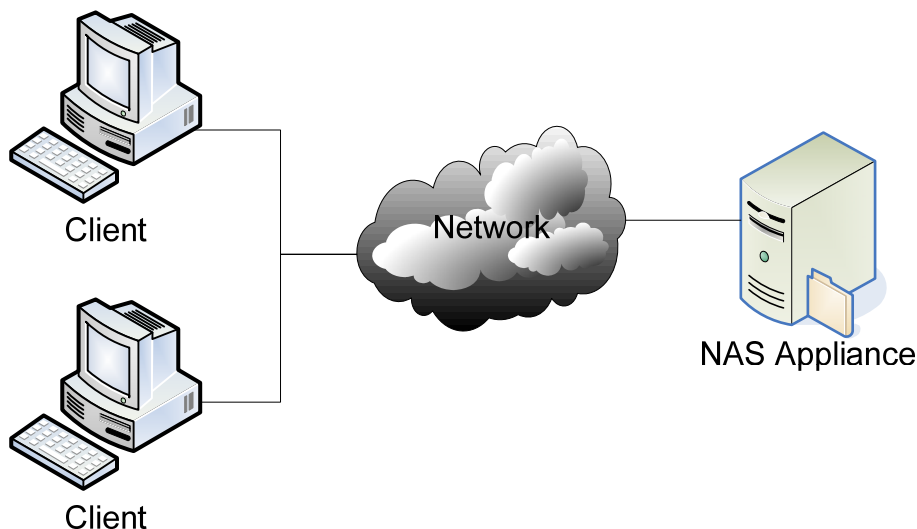


Figure 1: Network Attached Storage

A second networked-storage solution is the installation of a storage area network (SAN). Unlike NAS solutions, these dedicated storage networks require considerably more planning to deploy, and their management is more complex. SAN solutions are ideal for database and on-line processing

applications requiring rapid data access. SAN storage devices serve block-level data rather than files, and thus require either a Fibre Channel network to transport the SCSI (Small Computer System Interface) commands or a Gigabit Ethernet infrastructure to do so via iSCSI (Internet SCSI).

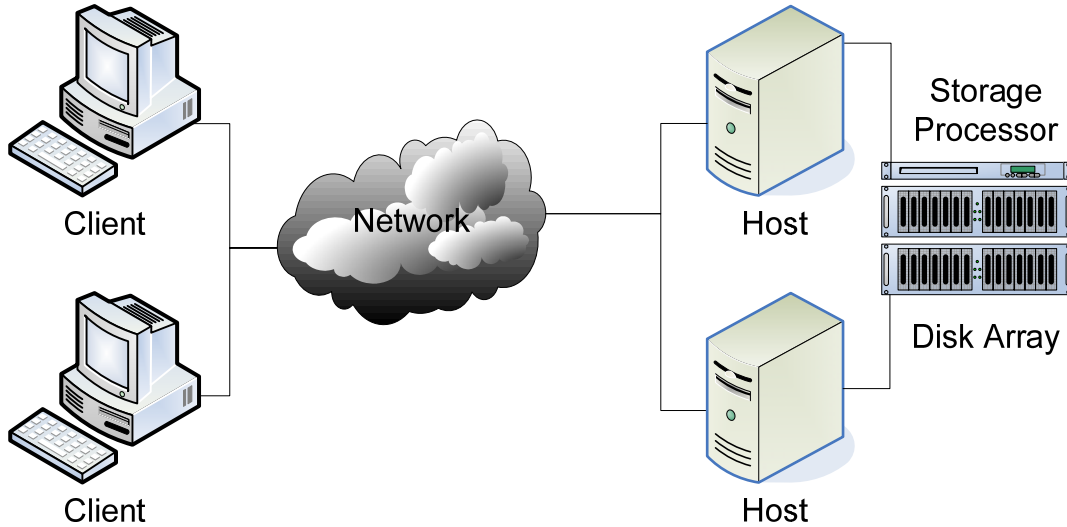


Figure 2: Storage Area Network

A third traditional storage network solution is a hybrid one: NAS devices can be attached to SANs. In this configuration, the NAS components necessary for filing are physically separated from the components necessary for storing: a NAS “head” or “gateway” (containing the filing functionality) attaches to the LAN network, and behind that lies the storage component—a “backend” SAN consisting of either the Fibre Channel network (wiring and switches) or the Gigabit Ethernet infrastructure for iSCSI and the storage disks

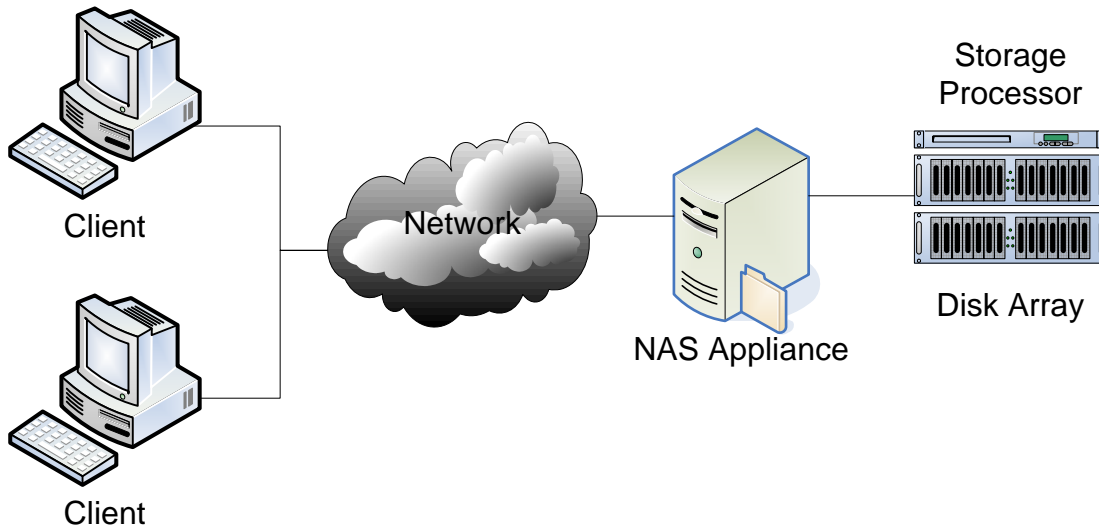


Figure 3: NAS Gateway on a SAN

The advantages of a traditional NAS design include:

- Increased scalability: NAS is designed to provide high disk capacity up to tens of terabytes, which means that a silo of data can be made available centrally.

- Designed for file serving performance: NAS is designed for data throughput. Its components are usually manufactured with multiple high-speed and multiple data paths, including internal components and external network adapters.
- Increased availability through redundancy: NAS is usually designed to offer high levels of availability through redundancy in internal components such as disks, fans, and even operating systems.
- Provision of data integrity: Typical NAS devices have a range of data backup, replication, and mirroring options, including LAN-free and server-less backup to offline tape or other storage devices.

Overview of Windows Storage Server 2003 R2

Microsoft® Windows® Storage Server 2003 R2 provides original equipment manufacturers (OEMs) with an operating system optimized for NAS functionality. Windows Storage Server 2003 R2 is the latest release in this operating system series and provides significant enhancements in storage management scenarios, as well as ease of installation and improved support processes for OEMs. Windows Storage Server 2003 R2 is part of the Windows Server 2003 family, which ensures that NAS devices built upon the Windows Storage Server 2003 R2 operating system have all the performance and scalability benefits associated with Windows Server 2003 R2.

A NAS appliance built with Windows Storage Server 2003 R2 is designed to perform without requiring a monitor, keyboard and mouse. Instead, this “headless” appliance is remotely managed through a new and intuitive Microsoft Management Console (MMC)-based Windows Storage Server Management user interface (UI) via Remote Desktop for Administration.

Windows Storage Server 2003 R2 takes advantage of the stability and security enhancements of a proven code base while extending connectivity and control into new areas of functionality. Windows Storage Server 2003 R2 offers the benefits of the Windows Server 2003 family, including Microsoft Active Directory® integration, while greatly improving access management, collaboration, branch server solutions, and storage setup and management both inside and outside your organization’s traditional network, storage, and management boundaries.

Windows Storage Server 2003 R2 also contains new Windows Storage Management features that build upon the existing features in Windows Server 2003 R2 and provide more optimization for the minimal configuration tasks required for setup. Unlike application servers which require proper planning to implement well, NAS devices built with Windows Storage Server 2003 R2 are designed to be deployed in under 15 minutes, and can be attached directly to the company local area network (LAN) with no interruption to services. Once plugged in, these NAS devices require minimal maintenance.

Unlike proprietary solutions, Windows Storage Server 2003 R2 works with standard hardware from multiple original equipment manufacturers (OEMs) and standard Windows Server utilities and tools. This gives businesses maximum flexibility in choosing among vendors with the hardware solution that best meets their needs.

Note: The Windows Storage Server 2003 R2 is meant primarily as a file server. It will not run most line-of-business applications. For more information about Storage Server functionality limitations, see Appendix B.

Key Features of Windows Storage Server 2003 R2

Windows Storage Server 2003 R2 is designed for simplicity, reliability, and performance. NAS appliances built on Windows Storage Server operating system integrate seamlessly into the IT network to provide one of the most economical file serving and network attached storage solutions available to departmental and enterprise-sized businesses.

- **Single Instance Storage:** Single Instance Storage (SIS) recovers disk space by reducing the amount of redundant data stored on a server, by identifying identical files, storing a single copy of the file in a central repository (the “common store”) and replacing the files with pointers to the file in the SIS common store.
- **Searching and Indexing:** Indexing Service extracts the information from a set of documents and organizes it in a way that makes it quick and easy to access that information through the Search function for computers running Microsoft Windows 2000 or Microsoft Windows XP.
- **Ease of Deployment:** Depending on the expertise of the system administrator and the complexity of the computing environment, installation of an application or general purpose server can take anywhere from several hours to a day or more of work. Because Windows Storage Server 2003 R2 comes preconfigured, other than using the new MMC

interface to set up users and shares, the only installation work necessary is plugging the device into the company LAN. In less than 15 minutes, gigabytes to terabytes of storage can be made available to users across multiple OS platforms.

- **Improved, Simplified Management:** Windows Storage Server 2003 R2 is managed remotely through Terminal Services sessions using a new, intuitive and user friendly MMC interface from any desktop on the network. Administrators already familiar with Windows do not have to learn a new NAS operating system as they will already be familiar with the features and functionality included in the Windows Storage Server 2003 R2 operating system. Also included is a new Distributed File System (DFS) Management snap-in that provides an intuitive graphical interface allowing administrators to easily configure DFS Namespaces (DFS-N) and DFS Replication (DFS-R) as well as offering other enhancements to configure target priority, delegation, client failback, and the management of existing features that were only configurable through the command line in Windows Server 2003. The new and improved File Server Resource Manager MMC provides storage administrators with the ability to configure file screens, reporting, and disk quotas without the need for third-party utilities.
- **Dependability:** Windows Storage Server 2003 R2 is designed to fully support redundant hardware components—disks, power supplies, and fans—to provide continuous and uninterrupted availability should a hardware failure occur. Because NAS devices are designed this way, the potential points of overall system failure are fewer than with general purpose servers.
- **Enhanced Data Protection:** Through built-in, point-in-time shadow copy technology, Windows Storage Server 2003 R2 helps businesses keep their data online 24x7 year around. Using the infrastructure provided by Windows Server 2003 R2 Volume Shadow Copy Service (VSS), the system administrator can use point-in-time shadow copy technologies to make up to 512 snapshots per volume (of which 64 are reserved for Shadow Copies for Shared Folders) using applications that have a VSS requestor. These shadow copy backups are available for rapid restores should the need arise. Unlike tape backups which can take hours to restore, these shadow copy backups can be restored in minutes. The end user can also benefit from point in time imaging technologies, using the Shadow Copies for Shared Folders (SCSF) feature. SCSF enables users to restore accidentally deleted or overwritten files or entire folders without the need for IT intervention. A maximum of 64 SCSF per volume can be created.
- **ISV Utility Support:** NAS devices developed with Windows Storage Server 2003 R2 include all of the benefits of application support available in Windows Server 2003 R2. Critical ISV utilities, such as antivirus, and backup are immediately available and supported in Windows Storage Server 2003 R2. NAS devices *not* based on Windows Storage Server 2003 are not able to support these kinds of products without special versions, or those operational procedures must be performed from application servers on the NAS files and file.
- **Robust Security:** Because it can be integrated seamlessly with Active Directory services, Windows Storage Server 2003 R2 can take advantage of the Windows security features such as data and file encryption, network authentication, secure network transport, and network wide group policies.
- **Load Balancing and Robust File Replication.** Windows Storage Server 2003 R2 includes a completely rewritten replication engine for the DFS: DFS Replication (DFS-R) provides a robust multi-master file replication service that is significantly more scalable and efficient in synchronizing file servers than its predecessor, FRS. DFS-R schedules and throttles replication schemes, supports multiple replication topologies, and utilizes Remote Differential Compression (RDC) to increase WAN efficiency. Moreover, if WAN connections fail, data can be stored and forwarded when WAN connections become available.

- **Advanced Compression Technologies:** RDC is a WAN-friendly compression technology that replicates only the changes needed to ensure global file consistency. RDC thus provides significant WAN efficiencies, including enhanced performance with replicated file size, an area of primary importance for branch office servers based on customer feedback.
- **Server Fail Over:** Windows Storage Server 2003 R2, when configured to use file replication, allows failover to a central server in case a local server fails. With Storage Server 2003 R2 Enterprise Edition, Windows failover clustering is also available.
- **Improved Performance:** Performance optimizations are automatically applied during clean installs of Windows Storage Server 2003 R2 and the existing Indexing Service is tuned for additional indexing and query performance.
- **Windows SharePoint Services:** Microsoft Windows SharePoint® Services, is a powerful web-based team collaboration environment that is included in Windows Storage Server 2003 R2. Windows SharePoint Services provides an integrated portfolio of collaboration and communication services designed to connect people, information, processes, and system both within and beyond the organizational firewalls.
- **WebDAV:** Web-based Distributed Authoring Versioning (WebDAV) allows customers to access and manage web based files using HTTP including data stored in Microsoft Windows SharePoint Services and MSN® Communities, as if they were standard file servers. The update to WebDAV also helps prevent customers' credentials (user name, password) from being transmitted over unencrypted channels during such exchanges.
- **Backup:** Backup programs that are compatible with the Windows Server 2003 family will run on Window Storage Server 2003 R2.
- **Antivirus:** As Windows Storage Server 2003 R2 is built upon the Windows Server family of operating systems, customers are also able to benefit from the increased security provided by their existing antivirus software from a variety of supported Microsoft partners. For more information on Microsoft's antivirus partners, see <http://www.microsoft.com/security/partners/antivirus.asp>.

Windows Storage Server 2003 R2 Features

This section details the Windows Storage Server 2003 features included in the R2 release of Windows Storage Server 2003.

Single Instance Storage

Single Instance Storage (SIS) recovers disk space by reducing the amount of redundant data stored on a volume by identifying identical files, storing only a single copy of the file in the SIS common store and replacing the files with pointers to the file in the SIS common store. Consider the following scenario:

1. Two users receive the same e-mail with an attachment. They both save the attachment to their home folder. SIS runs in the background and detects the two identical files on the volume and moves one of the copies into the SIS common store and replaces both files with a pointer to the file in the SIS common store.
2. One of the users makes a change to the file. SIS immediately detects that an update is pending for the file removes the pointer in the users home folder and replaces it with a copy of the file in the SIS common store. The updates are then applied to a fresh copy of the original file. This is completely transparent to the application.
3. The other user's file remains in the SIS common store with a pointer in the user's home folder, even if there is only one pointer to the file. When the second user updates the file (assuming there are no other pointers) the pointer is deleted and replaced with a copy of the original file and the file in the SIS common store is deleted.

SIS works automatically without any end-user involvement. The administrator can enable SIS on a per volume basis, and for best results the administrator should use a backup application that supports SIS.

Performance Optimization

As of the publication date of this paper, the operating system under girding Windows Storage Server 2003 R2 can achieve a NetBench benchmark of 4.11 gigabits per second for Common Internet File System (CIFS). Moreover, Windows Storage Server 2003 R2 has achieved a SpecSFS benchmark of 22,416 operations per second. For full disclosure on this benchmark test, visit the [fourth-quarter 2005 Standard Performance Evaluation Corporation SFS97_R1 results](http://www.spec.org/sfs97r1/results/res2005q4/) Web site at <http://www.spec.org/sfs97r1/results/res2005q4/>

The following optimizations are applied during the setup by the OEM:

- 8.3 filenames are disabled – disabling 8.3 filenames avoids the short-name attribute creation for the file and changes the way NTFS displays and manages the file.
- Paged pool size is limited – file cache space and paged pool space share a common area in the system virtual address. Limiting the paged pool allows for a larger system cache causes more content to be cached, and thus allows faster serving of files.
- Disk alignment is configured.
- Interrupt affinity driver is applied.

Search Enhancements

Indexing Service extracts the information from a set of documents and organizes it in a way that makes it quick and easy to access that information through the Search function for computers running Windows 2000 or Windows XP. This information can include text from within a document (its

contents), and the characteristics and parameters of the document (its properties), such as the author's name. Once the index is created you can query the index for documents that contain key words, phrases, or properties. For example, you can query all documents containing the word "product," or you can query for all Microsoft Office documents written by a specific author. Indexing Service returns a list of all documents that meet your search criteria.

Note: The amount of disk space required by the Indexing Service depends on how much text is in the data files that are being indexed. Pure text files (.txt) typically require an additional 30% of disk space for the index, but these are the most demanding file type. Other file types which also include non-text data (graphics, audio, binary data etc) require far less index space.

Windows SharePoint Services

Many organizations with large document management and team collaboration needs rely on Windows SharePoint Services, which is an ideal candidate to extend functionality with Windows Storage Server 2003 R2. This powerful, out-of-the-box implementation has the potential to open new chapters for organizations and their partners, vendors, and customers in terms of efficiency and collaboration.

Microsoft Windows SharePoint Services is a powerful web-based team collaboration environment that is included in Windows Storage Server 2003 R2. Windows SharePoint Services provides an integrated portfolio of collaboration and communication services designed to connect people, information, processes, and system both within and beyond the organizational firewalls.

SharePoint sites – file storage including collaboration

Web sites based on Windows SharePoint Services provide a place where teams can communicate, share documents, and work together on a project. SharePoint functionality includes:

- Team collaboration features including event calendars, contacts, Web links, discussions, issues lists, and announcements.
- Document libraries – places where users can store and retrieve documents while taking advantage of rich features such as check-in and check-out, version history, custom metadata, and flexible, customizable views.
- Web Parts which can provide data access, Web services, and many other applications and content to SharePoint sites.

Site users can contribute to the site by using nothing more than a Web browser. However, if users have Windows SharePoint Services-compatible client programs, such as Microsoft Office 2003, installed on their computers they can work seamlessly with the site, saving files to libraries, editing documents in the client program, and moving or linking that information to the site.

SharePoint Central Administration – Web browser interface for managing servers

Administrators can manage a single server or an entire server farm running Windows SharePoint Services from a Web browser interface called SharePoint Central Administration. Use SharePoint Central Administration to extend a virtual server, create sites (or turn on Self-Service Site Creation so users can create their own sites), manage security settings, manage the list of servers in a server farm, and so on. If administrators prefer, they can also use the Stsadm.exe command-line utility to manage their servers running Windows SharePoint Services.

Efficient Storage Management

To help accommodate the widely varying storage management capabilities found throughout the IT industry, Windows Storage Server 2003 R2 includes new tools designed to provide 1) a centralized view of storage; 2) simplified storage planning, provisioning and maintenance; and 3) improved monitoring and reporting. These benefits are captured in convenient new tools that enable administrators to more efficiently manage storage across IT resources and optimize storage space on

those resources. Storage management extends to two key new features contained in Windows Server 2003 R2:

- **Storage Manager for SANs** – Storage Manager for SANs enables customers to provision storage on one or more storage subsystems on a storage area network (SAN). Based on Microsoft Virtual Disk Service (VDS) technology, Storage Manager for SANs allows provisioning on Fibre Channel and Internet SCSI (iSCSI) storage subsystems.
- **File Server Resource Manager** – File Server Resource Manager (FSRM) enables system administrators to understand how storage is being used and to manage the use of their storage by generating storage reports, applying quotas to volumes and folders, and screening files on the server. Using FSRM you can better plan and optimize storage by creating quotas, creating file screens, and scheduling storage reports.

Centralized data stores reduce the management costs of geographically disbursed mini data centers. The replication of branch office data is automated to a central location at specified intervals when there is available bandwidth. In this way round trips are minimized.

Increased end-user productivity is provided by branch servers providing reliable and consistent access to the latest data that end users and applications rely on. The server employs local data to handle local requests or central servers in the event that a local server becomes unresponsive.

Storage Manager for SANs

Administered through a conventional snap-in, Storage Manager for SANs can be used on SAN-based storage arrays that support VDS using a hardware VDS provider. Because of hardware, protocol, transport layer and security differences, configuration and LUN management differ for the two types of supported environments. Storage management of SANs is done with built-in:

- LUN management for Fibre Channel subsystems
- Support for multiple I/O paths
- LUN management for iSCSI subsystems
- Securing data on iSCSI SANs

File Server Resource Management

With the increasing demand on storage resources as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time tracking the kind of information it contains. Managing storage resources has come to include not only data size and availability but also the enforcement of company policies and a very good understanding of how existing storage is utilized. This allows for sound strategic planning and proper response to organizational changes.

File Server Resource Manager (FSRM) is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports. This set of advanced utilities not only helps the administrator efficiently monitor existing storage resources, but also aids in the planning and implementation of future policy changes.

You can use FSRM to perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and to generate notifications when the quota limits are approached and exceeded.
- Create file screens to filter the files that users can save on volumes and in folders and to send notifications when users attempt to save blocked files.
- Create periodic or on-demand storage reports that allow you to identify trends in disk usage and to monitor attempts to save unauthorized files.

Storage Resource Manager quotas vs. NTFS disk quotas

The Windows 2000 and Windows Server 2003 operating systems support disk quotas, which are used to track and control disk usage per user on NTFS volumes. The following table outlines the advantages of using the quota management tools in Storage Resource Manager.

| Quota features | Storage Resource Manager | NTFS disk quotas |
|-------------------------|---|----------------------|
| Quota tracking | By folder or by volume | Per user on a volume |
| Disk usage calculation | Actual disk space | Logical file size |
| Notification mechanisms | E-mail, custom reports, command execution, event logs | Event logs only |

Distributed File System

One of the goals of information technology groups in medium and large organizations is to manage file server resources efficiently while keeping them available and secure for users. As organizations expand to include more users and servers—whether they are located in one site or in geographically distributed sites—administrators find it increasingly difficult to keep users connected to the files they need. On one hand, storing files on distributed servers makes files available to more users and decreases latency and bandwidth use when the servers are located near users. On the other hand, as the number of distributed servers increases, users have difficulty locating the files they need. In organizations with servers located in remote or branch offices, administrators face additional challenges, such as limiting network traffic over slow wide area network (WAN) connections, ensuring the availability of files during WAN outages or server failures, and ensuring that branch servers are backed up correctly.

The Distributed File System (DFS) solution in Windows Storage Server 2003 R2 helps administrators address these challenges by providing two technologies, DFS Namespaces and DFS Replication, which, when used together, offer simplified, fault-tolerant access to files, load sharing, and WAN-friendly replication.

The Distributed File System (DFS) solution in Windows Storage Server 2003 R2 provides simplified, fault-tolerant access to files and WAN-friendly replication. Distributed File System consists of two technologies:

- DFS Namespaces allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a namespace.
- DFS Replication, the successor to File Replication service (FRS), is a new state-based, multi master replication engine that supports scheduling and bandwidth throttling.

DFS Namespaces

DFS Namespaces allow administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a namespace. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration. Users can navigate these virtual namespaces without having to keep track of the names of the physical servers or shared folders hosting the data.

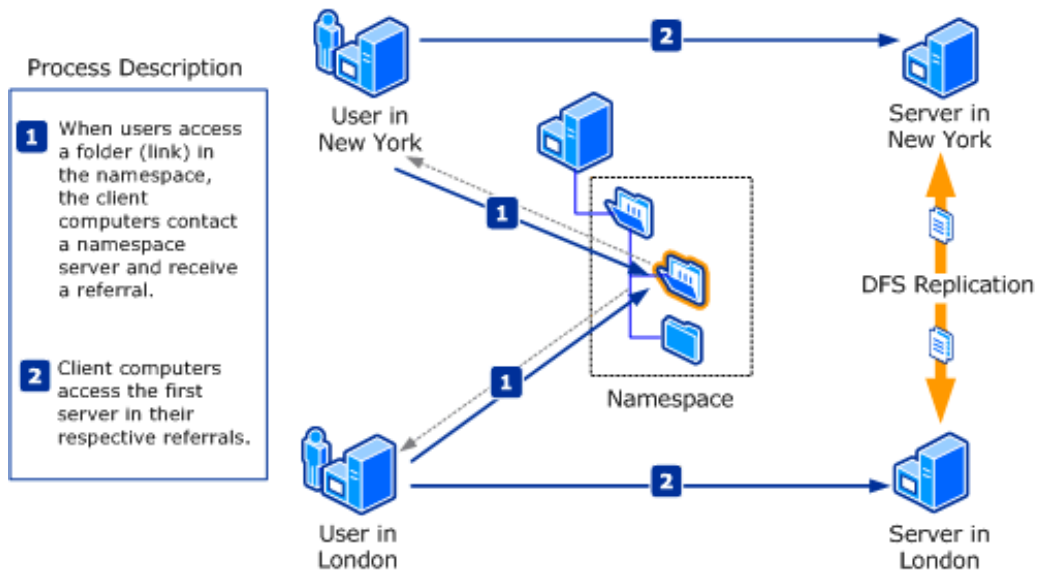


Figure 4: Distributed File System

If local servers become unavailable, DFS Namespace configurations provide for client failover by closest site selection and failback to a preferred server. For example, if failback is enabled on a DFS link that has targets in both the branch and the hub, branch clients will automatically failover to the hub when the service is unavailable.

For Windows Server 2003 R2, DFS failback functionality allows administrators to set server priorities for root and link target referrals, including setting high and low priority servers. In this case, servers are first ordered by site cost and then by priority within each site. Clients failback to the branch server when availability is restored.

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

Note: The functionality available through the DFS Management snap-in is also available from the command line.

DFS Replication with Remote Differential Compression

DFS Replication (DFS-R), the successor to File Replication Services (FRS), is a new state-based, multi-master replication engine that supports replication scheduling and bandwidth throttling. DFS-R schedules and throttles replication schemes and supports multiple replication topologies. If WAN connections fail, data can be stored and forwarded when WAN connections become available.

DFS Replication uses a new compression algorithm known as Remote Differential Compression (RDC). RDC is a “diff over the wire” protocol that can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, re-arrangements of data in files, enabling DFS Replication to replicate only the deltas (changes) when files are updated.

RDC can also copy any roughly similar file from any client or server to another using data that is common to both computers. This effectively reduces the size of the data sent and the overall bandwidth requirements for the transfer. Local differencing techniques — sometimes called “patching” — compute the differences between two local files, detecting insertions, removals, and

rearrangements of data. The differences can then be used to transform the old version to a new version. The differences between two known versions of a file are calculated on a server, and then sent to the client.

DFS includes a highly scalable, multi-master state-based file replication service that synchronizes file servers. It supports:

- Automatic recovery from database loss or corruption.
- On-demand scheduling and bandwidth throttling for replication schemes.
- Multiple replication topologies.

Note: Only Windows-based devices can be roots in DFS; NAS devices based on third-party software cannot serve as DFS roots.

Simplified Branch Server Management

Windows Storage Server 2003 R2 offers a new distributed file system implementation (DFS-N & DFS-R), with a new UI, WAN traffic compression based on RDC. This makes it easy to access your storage resources from anywhere by virtualizing the namespace and efficient data replication from branch to datacenter

Additionally, print management console is integrated into the core functionality of Windows Storage Server 2003 R2. It allows you to manage multiple printers from a single console. PMC works efficiently over WAN for branch office scenario, it also implements a push printer capability.

Network Security

Windows Storage Server 2003 R2 uses the authentication services of Windows Server 2003 R2 to ensure that only those users with permission to access data can do so. An additional layer of protection is available through the data encryption capabilities of the Windows operating system (see Table 7). The NAS Windows Storage Server is able to support security features for Unix environments, as well as earlier Windows platforms.

| Security and Authentication | Additional Information |
|-----------------------------|--|
| Kerberos | Network authentication protocol for client server configurations. |
| SSL | Secure Socket Layer protocol, provides connection security for Web servers. |
| IPsec | Internet Protocol Security, provides encryption for network transmission. |
| Active Directory | Directory service for Windows operating system; acts as central authority for network security. |
| Windows NT® Domain | Administration (including security and authentication) of users, groups, servers etc in Windows NT environments. |
| NTLM | Windows NT LAN Manager provides security for connections between NT clients and servers. |
| NIS | Network Information System (for Unix servers). |

Table 1: Security and Authentication Support in Windows Storage Server 2003 R2

File Serving across Different Operating Systems

Microsoft Services for Network File System

Microsoft Services for Network File System (MSNFS) provides Windows-based implementations of both the client and server aspects of Network File System (NFS), as well as related services and

utilities. The primary purpose of MSNFS is to provide an interoperability solution for enterprise businesses that have both Windows-based and UNIX-based clients. MSNFS supports the Network File System (NFS) protocol and provides file sharing interoperability between Windows and UNIX machines. It also provides a strategy for migrating from mixed or UNIX-based client environments to Windows. MSNFS is now an integrated part of Windows Storage Server 2003 R2.

MSNFS can provide:

- 64-bit support
- Better interoperation between Server Message Block (SMB) and NFS systems
- Improved reliability
- Support for NFS Devices (MKNOD)

File Sharing Protocols

File serving, as discussed earlier, is the process of opening, closing, reading and writing files. In order for file serving to be enabled over the network, file sharing protocols must be supported. These protocols enable clients to make requests and receive responses from remote devices, such as a NAS storage box.

Many NAS devices enable only simple file processing because they include support only for the standard file systems: the Network File System (NFS) and Common Internet File System (CIFS). In contrast, Windows Storage Server 2003 R2, with its support for multiple file sharing protocols (Table 8), allows for complex multi-platform file processing. In addition to Windows clients, Windows Storage Server 2003 supports file processing for Unix, and Web HTTP clients, among others.

| Protocols | Additional Information |
|-----------|---|
| CIFS/SMB | Enables Windows-based file sharing. |
| NFS | Enables Unix/Linux-based file sharing. NFS 3.0 supported in Windows Storage Server 2003. |
| HTTP | Enables web file sharing. |
| WebDAV | Enables desktop users to manage web based files using HTTP. |

Table 2: File Sharing Protocols Supported in Windows Storage Server 2003 R2

Note: To move from a fully or partially Novell® NetWare®-based file-serving environment, explore migration to Windows Storage Server 2003 R2 with your OEM.

Remote Hardware Management

Hardware Management is a new technology that enables system administrators to manage remote computers using SOAP-based protocols and standard Web services. In addition, Hardware Management provides a command-line interface for common management tasks, providing a quick and easy alternative to writing Windows Management Instrumentation (WMI) scripts. From the command prompt, an administrator can connect to a remote NAS appliance (using the HTTPS secure protocol) and do everything from retrieving a list of services installed on that computer to modifying the size of any page files found on that computer.

Print Management Console

Through the Printer Management Console (PMC), administrators have a central interface for managing all printers connected to all print servers within an organization. With PMC, administrators can monitor printer errors, deploy printer connections to clients, automatically find and install printers on a local branch office subnet, and run configuration scripts. PMC allows branch servers to perform as print servers because they are remotely manageable on a one-to-many basis.

Windows Storage Server 2003 R2: Predictable and Easy to Adopt

Windows Storage Server 2003 R2 was built on the efficient, reliable code base of Windows Server 2003 with SP1. Since it shares the same code base, there is no need for separate application compatibility testing, as all applications tested for Windows Server 2003 SP1 will function with Windows Storage Server 2003 R2.

Note: While all applications that function on Windows Server 2003 SP1 will function on Windows Storage Server 2003 R2, some will not be useable. In order to streamline Windows Storage Server 2003 R2 to the role of file and print server, some applications and services that would otherwise function on Windows Storage Server 2003 R2 are blocked from running.

This all translates into rapid realization of significant efficiencies for the teams and networks in organizations that have deployed or are planning to deploy Windows Storage Server 2003 R2.

Note: For information about how to upgrade a previous version of Windows Storage Server, contact your OEM.

Summary

Microsoft Windows Storage Server 2003 R2 provides OEMs with an operating system optimized for NAS devices. Windows Storage Server 2003 R2 is the latest release in this operating system series and provides significant enhancements in storage management scenarios, as well as ease of installation and improved support processes for OEMs. Customers will immediately benefit from the improved Server Resource Manager, significant DFS improvements including more efficient WAN replication and the integration of Windows SharePoint Services which provides power Web-based team collaboration environment.

Microsoft Windows Storage Server 2003 R2 lowers engineering and support costs by making setup and servicing easier and more tightly integrated with core Windows Server 2003 operating system management features and components like DFS.

Built on the Windows Server 2003 family, Windows Storage Server 2003 R2 takes advantage of the stability and security enhancements of a proven code base while extending connectivity and control into new areas. Windows Storage Server 2003 R2 offers all the benefits of the Windows Server 2003 family, while greatly improving branch server solutions, and storage setup and management.

Related Links

For more information on Windows Storage Server 2003 R2, Windows Server 2003, and the Windows Server System™, see the following:

- “[Windows Server 2003](http://www.microsoft.com/windowsserver2003/)” on the Microsoft Windows Server 2003 Web site at <http://www.microsoft.com/windowsserver2003/>
- “[Windows Server System](http://www.microsoft.com/windowsserversystem/)” on the Microsoft Windows Server System Web site at <http://www.microsoft.com/windowsserversystem/>
- “[Windows Server 2003 R2](http://www.microsoft.com/windowsserver2003/R2/)” on the Microsoft Windows Server 2003 Web site at <http://www.microsoft.com/windowsserver2003/R2/>
- “Windows Storage Server 2003” on the Microsoft Windows Server System site at <http://www.microsoft.com/windowsserversystem/wss2003/default.aspx>

Appendix A: Single Instance Storage Implementation

SIS Design

Single Instance Storage uses a file system filter driver (SIS.SYS) that can replace duplicate files with copy-on-write links (referred to as SIS links) that point to a single backing file, reducing the disk and cache overhead of those files. This backing file is contained in a common store. The implementation of the SIS architecture makes use of reparse points that contain information about the SIS links.

SIS links are implemented as sparse files, usually with most regions of the file unallocated, and a reparse point. The structure and contents of a reparse point is opaque to applications, for most usages the SIS.sys filter driver transparently provides the intended SIS behavior

One exception is that backup applications can send and retrieve data within these reparse points to and from SIS API functions that process the information in the reparse points. The information in a reparse point refers to a single backing file that contains the actual file data. This backing file is called a "common-store file", and it exists in the SIS common store.

Creating SIS Links

The SIS filter driver will set the size of link files (underneath the reparse point) to be equal to the size of the file in the common store, and set the files to be sparse. To SIS-link a file, it is copied into the SIS common store, renamed with a Globally Unique Identifier (GUID), and given the .sis file name extension. The identical files on the volume are then changed to reparse points. For more information, see "The SIS Common Store and Common-store Files" later in this guide.

The following figure shows three files that were combined with a reparse pointing to the SIS common store.

In the example below, each client computer contains a copy of a company handbook. Since the copies of the company handbook are all identical, the file has been placed in the SIS common store and the original files are changed to reparse points with referrals to the <guid>.sis file.

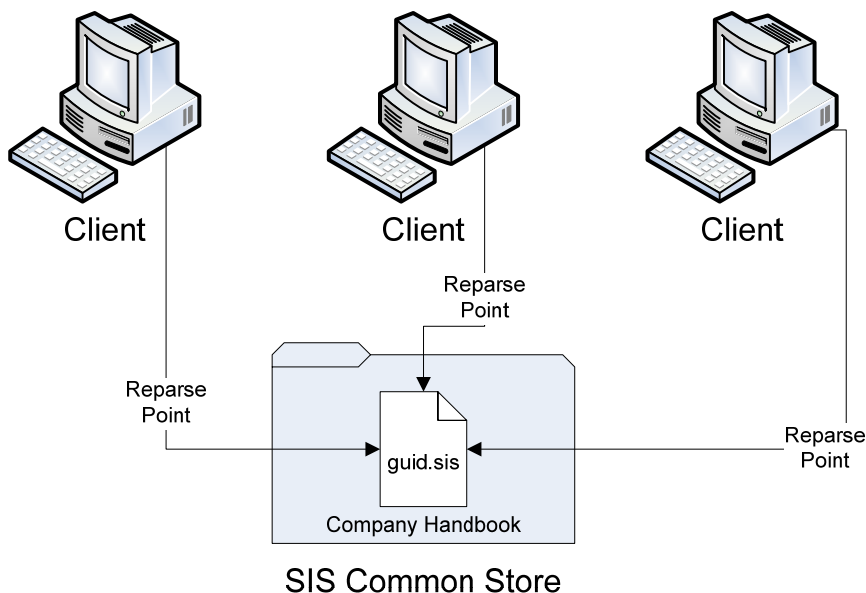


Figure 5: Creating SIS Links with Reparse Points

Accessing SIS files

When an application tries to open the original file, the sis.sys filter driver redirects any file input or output to the <guid>.sis file in \SIS Common Store. The filter driver is in all IO paths on volumes it is attached to.

Breaking SIS links

When a file consolidated by SIS-links is modified or its contents replaced, the SIS driver replaces the modified SIS-linked file with a copy of the <guid>.sis file from the SIS common store and removes its reparse point. The changes are then applied to the new copy of the original file.

The following figure shows an example of a file with no reparse point after modification.

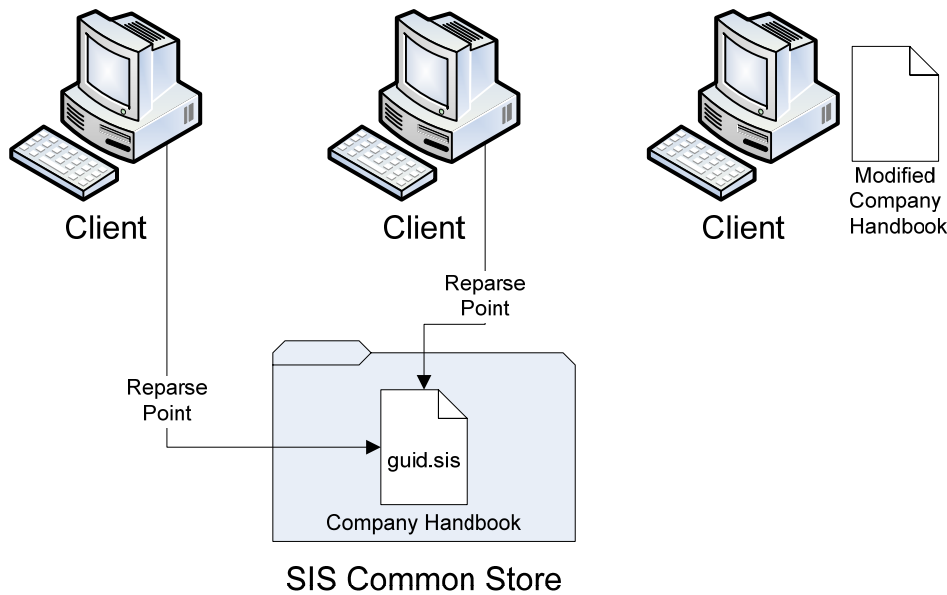


Figure 6: Breaking a SIS Link

The other reparse point or points for the original file are not changed, even if only one reparse point remains pointing to the <guid>.sis file. After the final instance of the original file is modified or deleted, the <guid>.sis file in the SIS Common Store folder is deleted.

Note: A change of timestamp or ACLs will not decouple a file from SIS: SIS only copies the data of the primary stream. SIS does not touch any metadata so any metadata changes (other than file size) will not affect SIS functionality on the file. This is done so that each file instance tied to SIS will still contain the correct time stamps, ACLs, attributes, among other data.

The SIS Common Store and Common-Store Files

All backing files maintained by SIS are called common-store files. One common store exists on each SIS-maintained volume and contains all of the common-store files that exist on that volume. This directory is located in the root directory of the volume and is called "\SIS Common Store". The common store is implemented as a directory that is restricted to allow access only to SYSTEM. Only the SIS filter and backup applications need access to this directory.

Note: The common-store file is removed only when the last SIS link pointing to it is deleted.

The SIS Groveler

The SIS Groveler searches for and combines files that are identical to files on the NTFS volume. It then reports those files to the SIS filter driver. It does most of its work when the system is not busy.

This is referred to as "background mode." It is possible to run the groveler at maximum capacity by using the sisadmin.exe tool. This is referred to as "foreground mode." After the groveler completes its work in foreground mode, it resumes normal operation in background mode.

Note: The groveler does not dynamically recognize newly configured SIS volumes. The sisadmin.exe tool will restart the groveler appropriately as configuration changes are made. The joint engine technology (JET) database that the groveler uses limits the groveler to managing six volumes at the same time. The groveler will only attach to NTFS volumes that are mounted as a drive letter. There is no support for SIS on volumes mounted using mount points.

Scanning Process

The SIS Groveler uses an Update Sequence Number (USN) Journal to know when files on a volume have been created or changed. If the journal wraps (e.g. the groveler misses some journal entries), the groveler will automatically re-scan the entire volume and recreate the signature database.

Backup/Restore Support

Because SIS creates reparse points from identical files, backup programs must understand this capability of the NTFS file system in order to efficiently backup and restore these files. SIS includes the SISBKUP.DLL modules to support backup/restore applications. For information about processing requirements for backup/restore applications, see the platform SDK.

The backup tool that is included with Windows 2000 and Windows Server 2003, Ntbackup.exe, handles these files properly.

Using a backup application that does not support SIS

If a backup program is not aware of SIS files, it would make duplicate copies of the SIS files in the actual backup data file. When restored, they would create normal files instead of the reparse points. The SIS Groveler would eventually combine these files, but possibly not in time to prevent the restore operation from running out of disk space.

Interoperability with other features

The following table provides information regarding interoperability issues SIS may have with other Windows Storage Server 2003 R2 features.

| Feature | Interop with SIS |
|-------------------------------|---|
| Windows Server 2003 R2 Quotas | Quota will only count the size of the reparse point towards the quota. If a directory quota is preventing SIS from bringing a changed file back from the SIS Common Store, part of the file will be brought back. The remaining part will stay in the SIS Common Store. SIS will retry this operation every time the file is changed. |
| Duplicate Files Report | The duplicate files report will report SIS'd files as being duplicated. In actuality, they are not because they have been SIS'd. Administrators can use Sisadmin to identify which duplicates have been SIS'd. |

| | |
|---|--|
| SIS Filter (Cluster) | Since SIS only links files within a single NTFS volume, the NTFS volume can failover even with SIS links. If the receiving node is running the SIS filter driver, then the files can be accessed normally. If the receiving node does not have the SIS driver, the files will not be accessible. |
| Groveler (Cluster) | In order to automatically SIS-link new files, the Groveler would need to be configured on the receiving node to scan the appropriate volume. For more information, see "The SIS Groveler" earlier in this guide. |
| Backup (Cluster) | In order to correctly backup or restore SIS files, the receiving node needs to have the SISBKUP.DLL installed. For more information, see "Backup/Restore Support" earlier in this guide. |
| Backup/Restore | See "Backup/Restore Support" earlier in this guide. |
| System Center Data Protection Manager (DPM) | DPM will operate with SIS if the appropriate hot fix is installed. Be sure to only configure SIS on the computer running the DPM agent and not on the DPM server. |
| Security | <p>When the SIS driver attempts to open the link file as part of the normal file create process, NTFS will do the standard permission checks at that time. SIS prevents users from creating bogus SIS links by verifying the checksum in the link against the common store file to which the link refers before allowing an open to complete. Any reparse points with bogus checksums will be fixed.</p> <p>Common store files (in fact the entire \SIS Common Store directory) are set to be accessible only to SYSTEM. As a result, users can't circumvent SIS and read the common store files directly unless they have full access to the local file system, in which case they could read the raw disk or override security on the user's links anyway. Admins can only access the SIS common directory if they first set the ACLs on the common store directory.</p> |
| Sparse Files | Sparse files will not be SIS'd. |
| Files that already have reparse points | Files that already have reparse points, such as HSM files, will not be SIS'd. |
| EFS File Encryption | EFS files will not be SIS'd. |

| | |
|---------------------------|--|
| Network File System (NFS) | Files with extended attributes will not be SIS'd. |
| Extended Attributes | Files with extended attributes will not be SIS'd. When components attempt to set extended attributes on SIS'd files, the files remain SIS'd and the attempt to set extended attributes fails. For more information, see "Network File System (NFS)" earlier in this table. |

Table 3: SIS Interoperability Issues**SIS Limitations**

SIS has the following limitations:

- SIS will only work with NTFS disk volumes.
- SIS will not work on a remote disk drive, such as a USB drive.
- SIS will not work on the boot/system drive.
- The SIS Groveler is limited to six disk drives.

Appendix B: Windows Storage Server 2003 R2 Licensing Restrictions

Windows Storage Server 2003 R2 has a different licensing structure from other Windows Server operating systems; some of the licensing restrictions are locked down in the operating system.

| Precluded Functionality | Examples/Specific Sevices Blocked | Notes/Exceptions |
|---------------------------------|--|---|
| Authentication Services | <ul style="list-style-type: none"> Using a NAS appliance running Windows Storage Server 2003 R2 as a domain controller Making use of dcpromo.exe | |
| Directory Services | <ul style="list-style-type: none"> Microsoft Active Directory | <p>Note: The following Windows Server 2003 R2 components are not available on all Windows Storage Server 2003 R2 editions:</p> <ul style="list-style-type: none"> Active Directory Application Mode (ADAM) Active Directory Federation Service (ADFS) |
| Network Infrastructure Services | <ul style="list-style-type: none"> Routing and Remote Access Service (RRAS) Windows Internet Name Service (WINS) | <p>Exception: Dynamic Host Configuration Protocol (DHCP) functionality is permitted.</p> |
| Terminal Services | <ul style="list-style-type: none"> Windows Server 2003 Terminal Services | <p>Exception: The limited extent necessary to administer Windows Storage Server 2003 R2 in “remote administration” mode is permitted.</p> |
| Network Load Balancing | <ul style="list-style-type: none"> Windows Server 2003 Network Load Balancing network driver | <p>Notes:</p> <ul style="list-style-type: none"> DFS-based, WAN-node failover is executable with NAS devices running Windows Storage Server 2003 R2. Failover between clustered NAS devices running Windows Storage Server 2003 R2 is permitted. |
| Enterprise Database Engines | <ul style="list-style-type: none"> Microsoft SQL Server | <p>Exceptions:</p> <ul style="list-style-type: none"> Non-enterprise database engines that are licensed to support not more than twenty-five (25) concurrent users—such as Microsoft SQL Server Desktop Engine (MSDE)—are permitted. |

| | | |
|---|--|--|
| | | <ul style="list-style-type: none"> Enterprise database engines that are integrated in and operate only in support of the system utilities, server resource management, storage management software or similar software that is useable solely for administration, performance enhancement, and/or preventive maintenance of the software or server are permitted. <p>Note: Windows Storage Server 2003 R2 is permitted to store enterprise database software data files (for example, Microsoft SQL Server data files).</p> |
| Enterprise Mail, Messaging, and Team Collaboration Software | <ul style="list-style-type: none"> Exchange Server 2003: <ul style="list-style-type: none"> Store.exe Mad.exe Setup.exe Lotus® Notes: <ul style="list-style-type: none"> Nserver.exe SharePoint Portal Server: <ul style="list-style-type: none"> Setup.exe Owstimer.exe | <p>Exception: Windows SharePoint Services is permitted.</p> <p>Note: Windows Storage Server 2003 R2 is permitted to support storage and management of data files in connection with team collaboration software (for example, Microsoft Exchange data files).</p> |
| Line-of-business Applications | <ul style="list-style-type: none"> Custom, third-party business software Microsoft enterprise resource planning (ERP) solutions | |
| Serving Web-based time management applications that address appointments, meetings, and other calendaring items | <ul style="list-style-type: none"> Microsoft Outlook® Web Access Server | |

Table 4: Windows Storage Server 2003 R2 licensing restrictions