DIGILIANT

Windows Storage Server 2003 R2

User's Guide

1 Overview of Windows Storage Server 2003 R2

The Digiliant Storage Servers are Network Attached Storage (NAS) appliances that utilize the Microsoft Windows Storage Server 2003 R2 server platform. A NAS appliance built with Windows Storage Server 2003 R2 is designed to perform without requiring a monitor, keyboard and mouse; this type of device is typically referred to as a 'headless' appliance. A headless appliance can be fully managed remotely. In the case of Windows Storage Server 2003 this is accomplished via Remote Desktop; once the minimal configuration tasks required for setup are complete, the monitor, keyboard and mouse can be unplugged. Unlike application servers which require a great deal of planning and coordination to properly implement, NAS devices are designed to be deployed quickly, typically in 15 minutes or less, and can be attached directly to a company's Local Area Network (LAN) with no interruption to services. Digiliant Storage Servers provide performance gains over general purpose servers by using only the highest quality hardware components optimized for the demands of the production environment; this serves as a solid foundation for the Windows Storage Server 2003 R2 platform. The integration of Digiliant Storage Servers into a network typically results in improvements in the performance of existing servers in the network. This is due to the fact that Digiliant Servers are optimized specifically for file serving tasks; this frees other servers in your network of this resource intensive task, thereby increasing the efficiency of the information system on a large scale. In addition, Digiliant Storage Servers are capable of operating in diverse computing environments using a variety of communication protocols from basic Microsoft Windows Workgroups to complicated multiprotocol Domains using DFS, NFS, FTP, HTTP, and Microsoft SMB. The varieties of clients that can be serviced by Digiliant Servers are equally diverse, to include Microsoft Windows, UNIX, Linux, Novell, and Macintosh.

This chapter provides an overview of these environments and typical deployments strategies.

Networking

Network Data Transmission Protocols

In order for clients to access resources on the Storage Server, they must be connected to the network and have the necessary network protocols installed and configured. The most common method used to connect to the Storage Server utilizes TCP/IP over Ethernet. Depending on the system platforms of other computers on your network as well, as other variables, it may become necessary to utilize other communication protocols. The Windows Storage Server 2003 R2 platform allows your Digiliant Storage Server to communicate with non-Microsoft client and server machines, even those that may require Unix or Apple network protocols (Table 1-1), thus allowing nearly transparent cross-platform network communication.

Table 1-1 Networking Protocols Supported by Windows Storage Server 2003

Network Protocols	Additional Information
TCP/IP	Used to connect hosts to the Intranet and Internet.
AppleTalk	Networking protocol for Apple computers.
IPX	Internet packet exchange, most common in Netware network.
NetBEUI	NetBIOS Extended User Interface used for Windows environments.
SNMP	Simple Network Management Protocol, internet standard for network management.
Telnet	Provides remote terminal access to host.
Fiber Channel	Transmits block-level data; most common in SAN configurations.
Ethernet	Physical network, transmits files; most common transport for NAS.

Network Security

Digiliant Storage Servers use the most up-to-date and effective authentication services provided by the latest Microsoft Server technology; this helps ensure that only those with permission may access your data. An additional layer of

protection, beyond authentication, is available through the data encryption capabilities of the Windows Storage Server 2003 R2 operating system (Table 1-2). Windows Storage Server 2003 has a robust set of security features, including tools for working with UNIX and Macintosh platforms as well as earlier versions of Windows.

Table 1-2 Security and Authentication Support in Windows Storage Server 2003.

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 OS, acts as central authority for network security.
Windows NT Domain	Administration (including security and authentication) of users, groups, servers etc in NT environments.
NTLM	Windows NT LAN Manager provides security for connections between NT clients and servers.
NIS	Network Information System (for Unix servers).
Apple UAM	User Authentication Module.

File Serving

File Sharing Protocols

File serving could be described as the process of making storage media available via network connection to a variety of clients for the purpose of storing and retrieving data in the form of discreet packages, or 'files'. In order to make this media available for storage and retrieval, common file sharing protocols must be used by the client and the server; protocols, like languages, allow the various computers to speak among themselves in a meaningful way. Such protocols allow servers and clients to establish an appropriate level of trust, thereby allowing or denying certain services or levels of access based on the client system's credentials.

Many NAS devices are limited in that they can only support the most basic file systems, generally the Common Internet File System (CIFS) and Network File System (NFS). In contrast, Windows Storage Server 2003 R2, with its support for multiple file sharing protocols (Table 1-3), allows for complex multi-platform file processing. In addition to Windows clients, Windows Storage Server 2003 supports file processing for UNIX, Macintosh, and Web HTTP clients, among others.

Table 1-3 File Sharing Protocols Supported in Windows Storage Server 2003

Protocols	Additional Information
SMB/CIFS	Enables Windows-based file sharing.
FTP	Enables file transfer.
DFS	Enables simplified, fault-tolerant file access and replication.
NFS	Enables Unix/Linux-based file sharing. NFS 3.0 supported in Windows Storage Server 2003.
AppleTalk	Enables Apple file sharing.
HTTP	Enables web file sharing.
WebDAV	Enables desktop users to manage web based files using HTTP.
NetWare	Enables Novell-based file sharing. Administration through Remote Desktop.

Supported Utilities and Applications

The specialized nature of NAS offers key advantages as a dedicated file server. Among these advantages are system capabilities that are dedicated to protecting data on the server from corruption and the possibility of hardware failure.

Based on the latest Microsoft Server technology, Windows Storage Server 2003 R2 is able to support a variety of anti-virus and backup utilities (Table 1-4).

Table 1-4 Software Supported in Windows Storage Server 2003

Software Support	Additional Information
Anti-virus	Via third party software.
Backup	Via third party software. Includes the backup utility NTBACKUP which uses VSS and SCSF to backup both system and user data with shadow copies.
Active Directory	Integration with the Active Directory service provides role-based authentication and authorization with single sign-on to system resources.
Virtual Disk Service	Manage heterogeneous hardware vendors with a single set of management tools, even when your environment includes devices from multiple vendors.
Exchange Feature Pack	The Feature Pack allows Microsoft Exchange Server 2003 databases and transaction logs to be stored on a Storage Server running Microsoft Windows Storage Server 2003. A single Storage Server computer running the Feature Pack can host the databases and transaction logs of up to two Exchange servers and up to 1,500 Exchange mailboxes. The Feature Pack installs new components on both the Storage Server computer and Exchange Server 2003. These components provide tools and services that allow Exchange databases and transaction logs to be moved to a Storage Server computer, and they perform the necessary configuration updates to give Exchange Server 2003 access to the remotely stored files.
iSCSI Target Feature Pack	The iSCSI Feature Pack allows Windows Storage Server 2003 powered Digiliant Storage Server to be functioned as NAS device and iSCSI Target at same time.

Storage

Managing Storage Devices

The Windows Storage Server 2003 R2 platform helps simplify your file server migration and consolidation efforts by streamlining the management of your storage infrastructure and providing a single platform for DAS, NAS, and SAN architectures.

Table 1-5 Disk and Deployment Management Capabilities in Windows Storage Server 2003

Disk Management	Additional Information
Windows Storage Management Console	Provides one place to manage files or print serving components. The console is accessible using Remote Desktop. The Storage Management page provides a portal to: • File Server Resource Manager • DFS Management • Disk and Volume Management • Index Service • MSNFS (under “Share Folder”) • Cluster Management (under “Utility”) • Shares • Sessions • Open Files
File Server Management Console (FSM)	Provides a centralized tool for managing your file server. Using FSM, you can perform many file server management tasks, such as formatting volumes, creating and managing shares, defragmenting volumes, setting quota limits, creating storage utilization reports, replicating data to and from the file server, managing Storage Area Networks (SANs), and sharing files with UNIX and Macintosh systems.

File Server Resource Manager	File Server Resource Manager 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 Storage Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports. By using Storage Resource Manager, you can perform the following tasks: Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached and exceeded. Create quotas that limit the space allowed for a volume or folder and that generate notifications when the quota limits are approached and exceeded. Schedule periodic storage reports to allow users to identify trends in disk usage and to monitor attempts to save unauthorized files, or generate the reports on demand.
Search enhancements	The Indexing service is tuned for additional indexing and query performance. Prior to the R2 release, if the Indexing service on a Windows Storage Server was not entirely up-to-date, the client-side search engine needed to "walk through" all the files within the scope of the search on the server. With the performance tuning in R2, the Indexing service no longer needs to be entirely up-to-date.
Storage Area Network Support (SAN)	Control growing volumes with a SAN-friendly configuration, a benefit that protects volumes from unintentional access. Improved handling of fiber channel SANs and improved SAN Host Bus Adapter (HBA) interoperability also help ease administration.
Single Instance Storage	Single Instance Storage (SIS) provides a copy-on-write link between multiple files. Disk space is recovered by reducing the amount of redundant data stored on a server. If a user has two files sharing disk storage by using SIS, and someone modifies one of the files, users of the other files do not see the changes. The underlying shared disk storage that backs SIS links is maintained by the system and is only deleted if all the SIS links pointing to it are deleted. SIS automatically determines that two or more files have the same content and links them together.
Windows Share Point Services	Windows Share Point Services is an integrated set of collaboration and communication services designed to connect people, information, processes, and systems, both within and beyond the organization firewall.
Storage Area Network Support (SAN)	Control growing volumes with a SAN-friendly configuration, a benefit that protects volumes from unintentional access (San) support improves handling of fiber channel SANs and improved SAN Host Bus Adapter (HBA) interoperability also help ease administration.

Ensuring Availability of Stored Data

A number of new features have been added to Windows Storage Server 2003 to help keep your business's data highly available, while many existing technologies have been enhanced. These technologies include tools to ensure hardware components are operating at peak efficiency such as MPIO (Multipath I/O), technologies enabling hardware and software redundancy such as RAID (Redundant Array of Inexpensive Disks), as well as data availability techniques for redundancy such as clustering (Table 1-6).

Table 1-6 Features Enhancing System Reliability and Data Availability

Reliability/Availability	Additional Information
Distributed File System	Distributed File System (DFS) helps users find data using simplified file paths regardless of where the data is stored—decreasing business losses from network outages and individual server failures.
Distributed File System Replication	Distributed File System Replication (DFSR) helps IT staff to configure file servers to automatically replicate documents between servers and provide transparent access anytime and anywhere to information stored on multiple file servers.
Volume Shadow Copy Service	Volume Shadow Copy Service (VSS) is an infrastructure feature that enables IT staff to create "point-in-time" volume snapshots for easier, more reliable backups and smaller backup windows when used with compatible applications.
Clustering	Provides application failover. Windows Storage Server 2003 Enterprise Edition supports up to 8 nodes

Software RAID 0, 1, 5	RAID types provide differing levels of data protection and redundancy. VDS also enables hardware based RAID.
Multipath I/O (MPIO)	Enables high performance and high availability through multiple paths to storage.
System Monitoring	Monitors performance of the operating system. Allows system administrator to assess I/O performance with different devices.
Watchdog Timer	Detects system hangs; can be programmed to reboot system after a given time.

Deployment Scenarios

The Storage Server is deployed in one of two security modes:

- Workgroup
- Domain (Windows NT® Domain or Active Directory Domain)

The Storage Server uses standard Windows user and group administration methods in each of these environments.

Workgroup

In a Workgroup environment, users and groups are stored and managed separately, on each member server of the Workgroup. Workgroups are typical for very small deployments where little or no computing environment planning is required.

Domain

When operating in an Active Directory Domain environment, the Storage Server is a member of the Domain and the Domain controller is the repository of all account information. Client machines are also members of the Domain and users log on to the Domain through their Windows based client machines. The Domain controller also administers user accounts and appropriate access levels to resources that are a part of the Domain, to include files located on the NAS.

The Storage Server obtains user account information from the Domain controller when deployed in a Domain environment. *The Storage Server itself cannot act as a Domain controller, backup Domain controller, or the root of an Active Directory tree as these functions are disabled in the operating system.*

Windows Storage Server 2003 Editions

The following table lists, by edition, the Windows components that are either pre-installed (P), available (A), or unavailable (U) with Windows Storage Server 2003 R2.

Table 1-7 Component Comparison between Editions

Component	Express	Workgroup	Standard	Enterprise
Software				
File server role	P	P	P	P
File Server Management (FSM)	P	P	P	P
Print Management Console	P	P	P	P
Microsoft Services for Network File System (NFS)	P	P	P	P
Microsoft .NET Framework 2.0	P	P	P	P
Indexing Service	P	P	P	P
File Server Resource Manager (FSRM)	U	P	P	P
Distributed File Systems (DFS) Management	U	P	P	P
DFS Replication	U	P	P	P
DFS Replication Diagnostics and Configuration Tools	U	P	P	P
Storage Manager for Storage Area Networks (SAN)	U	U	P	P
Single Instance Storage (SIS)	U	U	P	P

Windows SharePoint Services	U	U	A	A
Microsoft Exchange Feature Pack	P	P	P	U
Transportable snapshots	U	U	U	P
Hardware				
RAM	4 GB	4 GB	4 GB	8 GB
Processors	1	1	1 to 4	1–8

2 Configuration

Start up

- Start the Storage Server with a monitor, a keyboard and a mouse connected to the back of the server.
- The start up wizard appears upon startup to guide you through the setup process. You will need to approve the Software License Agreement and enter the activation key code, as in other Windows software installations. The activation key code can be found on the case of the server.
- After entering the activation key code, when asked to set the IP address, note that 'Normal Installation' assumes that you have a DHCP server on your network to assign the Storage Server an IP address. If this is not the case, you can use the "Custom Settings" process to assign a static IP address to the Storage Server. When the wizard is finished, the Storage Server will reboot automatically.
- If your Storage Server acquires the IP address from a DHCP server, you can find the Storage Servers' IP address by running the "`ipconfig /all`" command under command prompt.
- The default computer name for the Storage Server is '***nasserver'***. The default administrator's password is '***digiliant***". After login to the server with the default password, you can change the password and server name. If you change the password, you should record it, or you will not be able to access the Storage Server later.
- Once the setup Wizard configuration is complete and you have the IP address for the server, you can disconnect the monitor, keyboard and mouse. You should now be able to manage the Storage Server through Remote Desktop using any computer on the network.

User interfaces

The Storage Server desktop can be accessed by either of two methods:

- Directly connecting a keyboard, mouse, and monitor;
- Over the network, via Remote Desktop.

Remote Desktop

The Remote Desktop client is automatically installed on Windows XP. For older version of Windows, you can download the Remote Desktop client following the link list bellow:

http://www.microsoft.com/windowsxp/downloads/tools/rdclientdl.mspx

One way to make a Remote Desktop connection to the Digiliant Storage Server is to following these steps:

- Choose 'Start' → 'All Programs' → 'Accessories' → 'Communications' and click 'Remote Desktop Connection'. This displays the 'Remote Desktop Connection' dialog box.
- In the 'Computer' text box, type "**nasserver**" or the IP address of

Figure 2-1 Windows Storage Server 2003 R2 Desktop

the Storage Server.

- The 'Login User Name' is "**Administrator**," the 'Password' is "**digiliant**".

Windows Storage Server Management Console

The Windows Storage Server Management Console is a suite of tools that allows administrators to understand, control, and manage the data stored on their servers. To access Windows Storage Server manage console, you can either click the 'Storage Server Management' icon on the desktop, or choose 'Start' → 'All Programs' → 'Administrative Tools' and then click 'Windows Storage Server Management'. This opens the console shown in Figure 2-2.

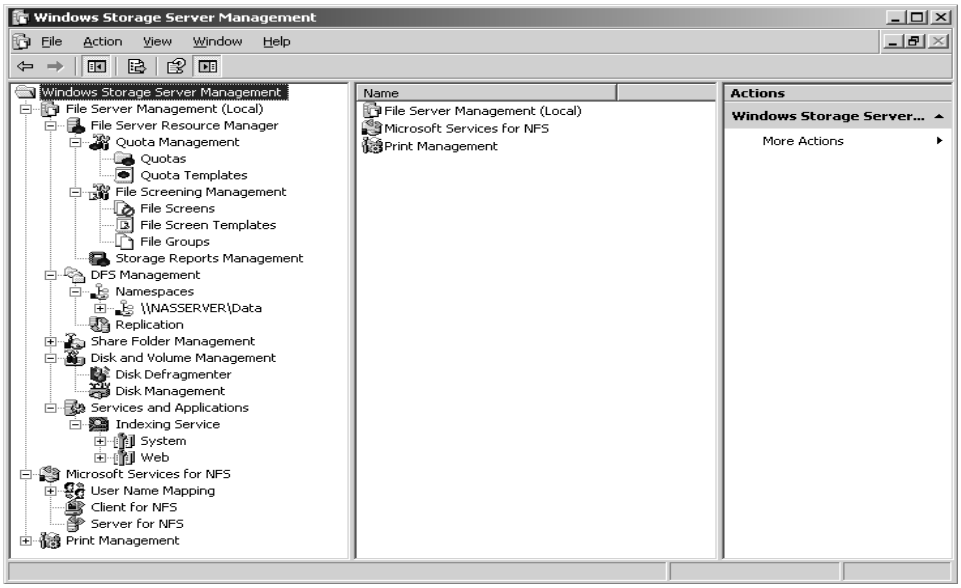

Figure 2-2 Windows Storage Server Management Console

Server Name and Domain

You can use the 'My Computer' on the Windows Storage Server 2003 Desktop to set the server name and the name of the Domain you want the server to joint:

1. In the Windows Storage Server 2003 Desktop, right-click the 'My Computer' icon, and select 'Properties'. Click the 'Computer Name' tab, and then click the Change button. This opens the dialog box show in Figure 2-3.

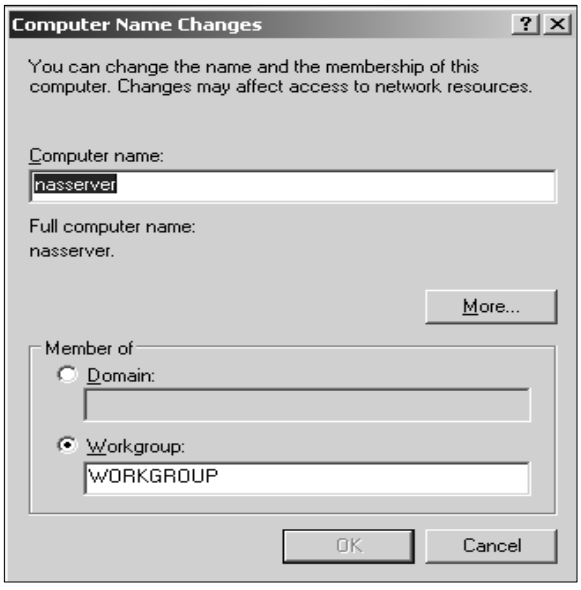

2. Enter a name for the Storage Server in the 'Computer Name' field.
3. Indicate whether to join the Storage Server to a Workgroup or a Domain. The Storage Server can join one of the following types of Domains or Workgroups:
   - For Workgroup (used also for Linux, UNIX, or if no Domain or Workgroups exist on the network), enter any name.
   - Microsoft Windows NT 4 Domain, Microsoft Active Directory Domain.

   If you are joining the Storage Server to a Domain, you either must be, or be able to furnish the user name and password of an authorized Domain Administrator.
4. Click '**OK'**, when prompted to reboot the Storage Server, you may either accept or cancel the reboot. If you cancel, the changes will not take effect until the next reboot.

Administrator Account

The administrator's password should be changed periodically where system security is a concern. To change the administrator's password, press **Control** + **Alt** + **Del** if you are logged on locally or press **Control** + **Alt** + **End** if you are logged on through Remote Desktop. This opens the dialog show in Figure 2-4.

Figure 2-4 Administrator Password

1. For User Name, enter the name you want the administrator to use for the logon name.
2. For Old Password, enter the password ("**digiliant**") you used to log onto the Storage Server.
3. For New Password and for Confirm Password, enter the password you want the administrator to use for logging onto the Storage Server. There is no explicit confirmation, but an error message will appear if you are not successful.

If you are logged on as a Domain user, you may receive an error message saying that the account name or password "*cannot be changed for this Domain account*." You must be logged on as the Storage Server administrator to change the administrator password. Only one administrator account has administrative privileges on the Storage Server; this dialog can be used for only one administrator. However, the Storage Server software does have a set of default accounts, and if you add a Domain Account to the 'Local Administrators Group', other Domain Users can be set up to administer the Storage Server.

Network

Setting the IP Address

Computers use Internet Protocol (IP) address to communicate over TCP/IP. If your network has a Dynamic Host Configuration Protocol (DHCP) server you can have the this server assign IP and Default Gateway addresses for the Storage Server; alternatively can just enter a static IP address manually. Note, however, that changing the Storage Server IP address will break your Remote Desktop connection and the connections of all clients using the Storage Server at that time. If you change the IP address, you will need to use the new IP address to re-establish the connection.

To set or change the IP settings:

1. Access 'Network Connections' in the 'Control Panel'.
2. Select or double-click the connection with which you want to work.
3. Click the 'Properties' button. This opens the Properties dialog for the desired network connection. Double-click 'Internet Protocol (TCP/IP)' in the list box of the 'Properties' dialog, or select 'Internet Protocol (TCP/IP)' from the list box and then click the 'Properties' button. This opens the 'Internet Protocol (TCP/IP) Properties' dialog box, shown in Figure 2-5.
4. Select 'Use the following IP Address' and type the IP

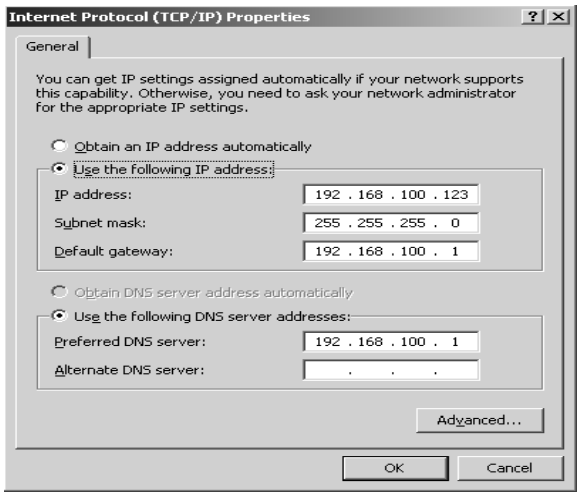

Table 2-5 IP Address Dialogue

address for the Storage Server in the 'IP Address' field. The IP address you assign to the Storage Server must not be used any-where else on the network.

5. The 'Subnet Mask' field ensures that the Storage Server communicates over the network properly. Windows Storage Server 2003 should insert a default value for the subnet mask into the 'Subnet Mask' field. If the network doesn't use subnets, the default value should suffice; however, if it *does* use subnets you'll need to change this value as appropriate for your network.
6. If the computer needs to access other TCP/IP networks, the Internet, or other subnets, you must specify a 'Default Gateway'. Type the IP address of the network's default router in the 'Default Gateway' field.
7. Domain name services are needed for Domain name resolution. Type a preferred and alternate Domain Name System (DNS) server address in the fields provided.
8. When you're finished, click '**OK**'. Repeat this process for other network adapters you want to configure. Keep in mind that each network adapter must have a unique IP address.
9. Configure Windows Internet Name Service (WINS) as necessary. You might also need to set advanced options for DNS.

DNS Resolution

DNS is a host name resolution service. You use DNS to determine a computer's IP address from its host name. This allows users to work with host names, such as http://www.Digiliant.com, rather than an IP address, such as 192.168.15.3. DNS is the primary name service for Windows Storage Server 2003 R2 and Internet. The basic DNS configuration is covered under Setting IP Address section. You configure advanced DNS settings by using the DNS tab in the Advanced TCP/IP Setting dialog box shown in Figure 2-6. You use the field of the DNS tab as follows:

"DNS server addresses, in order of use:"

Use this area to specify the IP address of the DNS servers that are used for Domain name resolution. Use the 'Add' button to add a server IP address to the list. Use the 'Remove' button to remove a server from the list. Use the 'Edit' button to edit the selected entry. You can specify multiple servers to use for DNS resolution. These servers are used in the order displayed; if the first server isn't available to respond to a host name resolution request the next DNS server on the list is accessed, and so on. It's important to note that *TCP/IP doesn't go to the next server if the first server can't resolve the name, only if the first server doesn't respond*. To change the position of a server in the list box, click to highlight it and use the 'Up' or 'Down' arrow button.

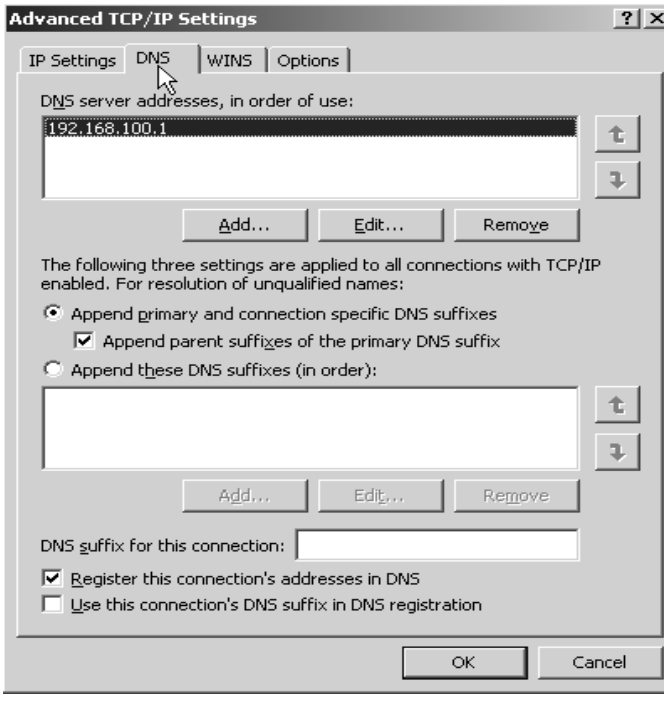

Figure 2-6. Advanced DNS settings

"Append primary and connection specific DNS suffixes"

Select this option to resolve unqualified computer names in the primary Domain. For example, if the computer name "**nas**" were used and the parent Domain was 'digiliant.com.', the computer name would resolve to 'nas.digiliant.com.' If the fully qualified computer name doesn't exist in the parent Domain, the query fails. The parent Domain used is the one set in the 'Network Identification' tab of the 'System Properties' dialog box. Normally, this option is selected by default.

"Append parent suffixes of the primary DNS suffix"

Select this option to resolve unqualified computer names using the parent-child Domain hierarchy. If a query fails in the immediate parent Domain, the suffix for the parent of the first parent Domain is used to again try to resolve the query.

This process continues until the top of the DNS Domain hierarchy is reached. For example, if the computer name "**nas**" ware used in 'tech.digiliant.com.', DNS would attempt to resolve the computer name to 'nas.tech.digiliant.com.' If this didn't work, DNS would attempt to resolve the computer name to 'nas.digiliant.com'. This option is typically selected by default.

"Append these DNS suffixes (in order)"
Select this option to set specific DNS suffixed to use rather than resolving through the parent Domain. Use the 'Add' button to add a Domain suffix to the list. Use the 'Remove' button to remove a Domain suffix from the list. Use the 'Edit' button to edit the selected entry. You can specify multiple Domain suffixes. These suffixes are used in priority order. If the first suffix doesn't resolve properly, DNS attempts to use the next suffix in the list. If this fails, the next suffix is used, and so on. To change the order of the Domain suffixes, select the suffix, and then use the 'Up' and 'Down' arrow to change its position.

"DNS suffix for this connection"
Sets a specific DNS suffix for the connection that overrides DNS names already configured for use on this connection. You'll usually want to set the DNS Domain name through the Network Identification tab in the System Properties dialog box instead.

"Register this connection's address in DNS"
Select this option if you want all IP address for this connection to be registered in DNS under the computer's fully qualified Domain name. This option is selected by default.

"Use this connection's DNS suffix in DNS Registration"
Select this option if you want all IP address for this connection to be registered in DNS under the parent Domain.

WINS Resolution

You can use WINS to resolve a computer's NetBIOS name to an IP address, or to help computers on a network determine the addresses of other computers if a WINS server is installed on the network. Although WINS is supported on all versions of Windows, Windows Storage Server 2003 uses WINS primarily for backward compatibility. You can also configure computers running Windows Storage Server 2003 to use the local file 'LMHOSTS' to resolve NetBIOS computer names. However, LMHOSTS is consulted only if normal name resolution methods fail. In a properly configured network, these files are rarely used. Thus, the preferred method of NetBIOS computer name resolution is WINS in conjunction with a WINS server.

You can configure WINS by completing the following steps:

1. Access the 'Advanced TCP/IP Settings' dialog box, and then click the 'WINS' tab. This displays the dialog box shown in Figure 2-7.
2. The box named 'WINS Addresses, in order of use' allows you to specify the IP address of the WINS servers that are used for NetBIOS name resolution. Use the 'Add' button to add a server IP address to the list. Use the 'Remove' button to remove a server from the list. Use the 'Edit' button to edit the selected entry.
3. You can specify multiple servers to use for WINS resolution. These servers are used in priority order. If the first server isn't available to respond to a NetBIOS name resolution request, the computer accesses the next WINS server on the list, and so on. It's important to note that *TCP/IP doesn't go to the next server if the first server can't resolve the name,*

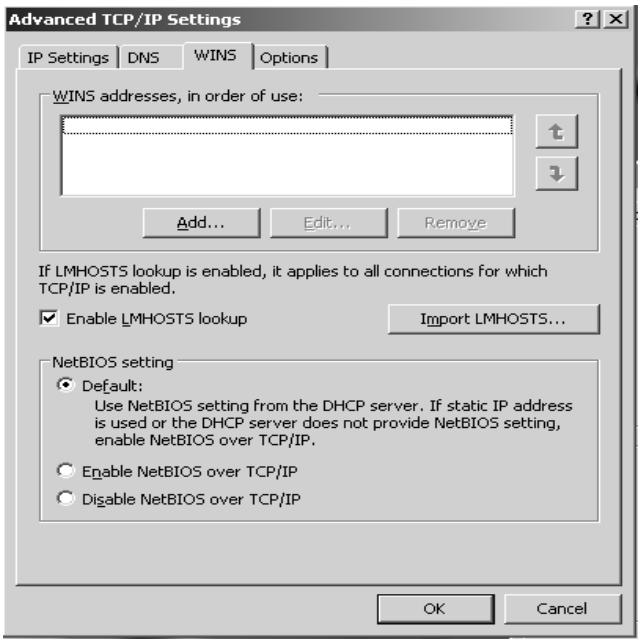

Figure 2-7 TCP/IP Hosts File

only if the first server doesn't respond. To change the position of a server in the list box, select it and then use the Up or Down arrow button to move the server in this list.

4. To enable LMHOSTS lookups, select the 'Enable LMHOSTS Lookup' check box. If you want the computer to use an existing LMHOSTS file defined somewhere on the network, retrieve this with the 'Import LMHOSTS' button. You generally use LMHOSTS only when other name resolution methods fail.
5. 'NetBIOS over TCP/IP' services are required for WINS name resolution. Choose one of the following options to configure WINS name resolution using NetBIOS:
    - If you use DHCP, you can get the NetBIOS setting from the DHCP server. Select 'Default', then 'Use NetBIOS Setting from the DHCP Server'.
    - If you use static IP addressing or the DHCP server doesn't provide NetBIOS settings, select 'Enable NetBIOS over TCP/IP'.
    - If WINS and NetBIOS aren't used on the network, select Disable NetBIOS Over TCP/IP. This eliminates the NetBIOS broadcasts that the computer would otherwise send.

Network Interface Controllers Properties

You can access 'Network Connections' in the 'Control Panel' to change and view the specific network interface controller's properties. By right clicking on the Network Interface Controller (NIC) you are interested in, you can bring up a dialog that will allow you change the name, or check the status of, the controller. If you select 'Status' you can check the link speed, the working status and the working duration of the network card.

Network Cards Teaming

Some Storage Servers are equipped with two built-in network interface controllers. There is a network card functionality know as "Teaming" that allows administrators to configure and monitor Ethernet interfaces in Windows-based operating system to provide additional functionality and options for increasing fault tolerance and load balancing. Presently, there are two companies that manufacture NICs that support Teaming, Intel and Broadcom.

Teaming for Intel based NICs

To access Intel's NIC utility, choose 'Start' → 'Control Panel' → 'Network Connections' and click the 'Local Area Connection' icon you want to manage. This opens the 'Local Area Connection Properties' dialog; click the 'Configuration' button. This opens the 'Intel Network Connection Properties' dialog, as shown in Figure 2-8.

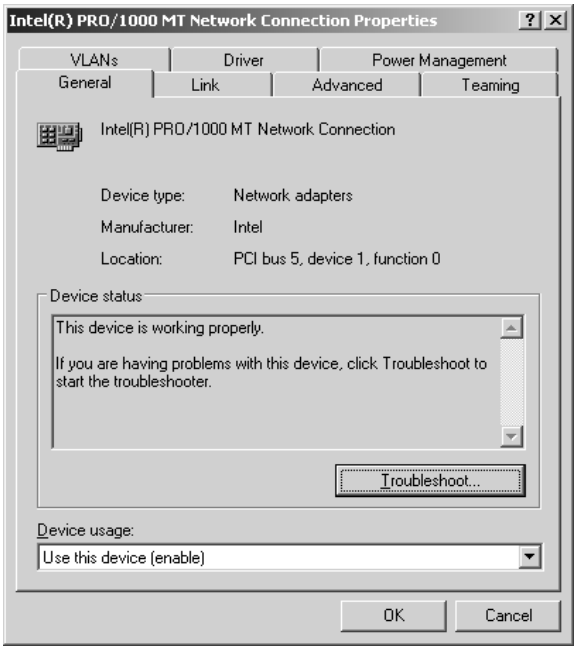

Figure 2-8 Intel NIC Utility

You can configure Teaming by completing the following steps:

1. Click the 'Teaming' tab on 'Intel Network Connection Properties' dialog.
2. Select 'Team with other adapters' option, and then click 'New Team'. This opens the 'New Team Wizard'.
3. Type a name for the new team, and then click 'Next'.
4. In 'Select adapters to include in this team' window, check all adapters which you want them to joint the team, and then click 'Next'.
5. In 'Select a team mode' window, select a NIC teaming mode. You can scroll through the 'Advanced Networking Service Team Types' window to understand which team type would work better for you. Click 'Next' when you finish.
6. Click 'Finish'.

You can remove a NIC team by completing the following steps:

1. Choose 'Start' → 'Control Panel' → 'Network Connections', right-click the network team you want to manage, and select 'Properties'. This opens the 'Local Area Connection Properties' dialog.
2. Click the 'Configuration' button; this opens the 'Team Properties' dialog.
3. Select the Setting tab; this displays the team members, team type and following action buttons:
   - **Remove Team:** Click this button to remove the team.
   - **Details:** Click this button to display team type and Ethernet address of this team.
   - **Modify Team:** Click this button to modify the team properties, for example change the team type.
   - **Test switch:** Click to button to test the status of this team.
   - **Adapter Properties:** Displays the network connection properties of this team.
4. To remove the team, click the 'Remove Team' button. Click 'Yes' to confirm the action.

Teaming for Broadcom based NICs

Broadcom Advanced Control Suite (BACS) provides powerful tool for trouble shooting network problem and configuring NIC teaming. To access BACS, choose Start, All Programs, Broadcom, and then click Broadcom Advanced Control Suite 2. This opens BACS program windows, as shown in Figure 2-9.

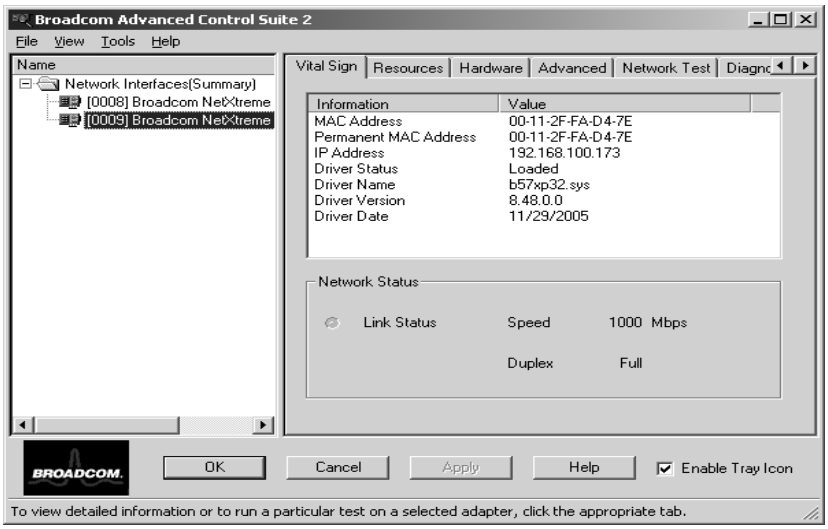

Figure 2-9 Broadcom Advanced Control Suite

You can configure NIC teaming by completing the following steps:

1. Select 'Tools' menu, and click 'Create a Team'. This opens the 'New Team Configuration' dialog, as shown in Figure 2-10.
2. Type a name for the new team, as well as team type, and click 'Apply'.
3. In the 'Member Assignment' dialog, highlight the network adapter you want to joint the team, and click the add button ('>') to add the desired network adapter into team member list. For a Storage Server with an Asus server motherboard, you may see a warning message about adding a network adapter into the new team. Ignore the warning message. This message is due an inconsistency between Asus and Broadcom design specifications and will not impact performance or functionality.

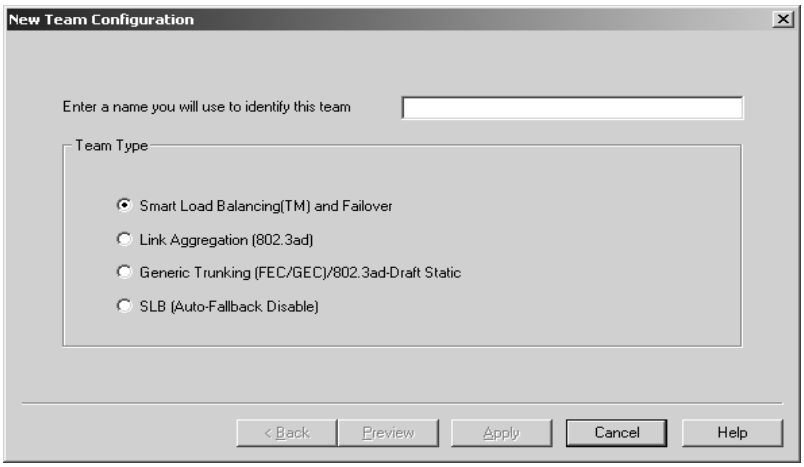

Figure 2-10 New Team Configuration

4. When finished click the 'Apply' button; click 'Yes' to confirm the action.

You can remove a NIC team by completing the following steps:

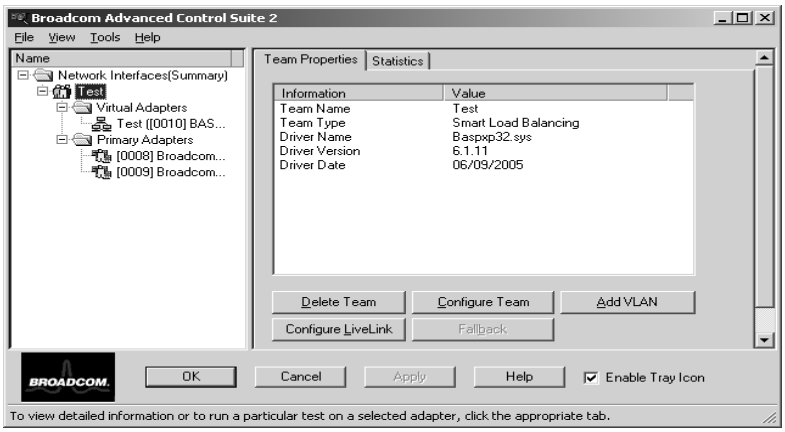

Figure 2-11 Change Team Properties

1. Open 'Broadcom Advanced Control Suite' and then click on the team you want to delete.
2. On the 'Team Properties' tab pane, you have options to: 'Delete Team', 'Configure Team', 'Add VLAN' and 'Configure LiveLink'. Click the 'Delete Team' button to delete the team. This action removes the team from the Network Interface list. The previous member NICs would show on the Network Interface list.
3. Click 'Apply' to make the change permanent.
4. Click 'Yes' to confirm the action.

Maintenance

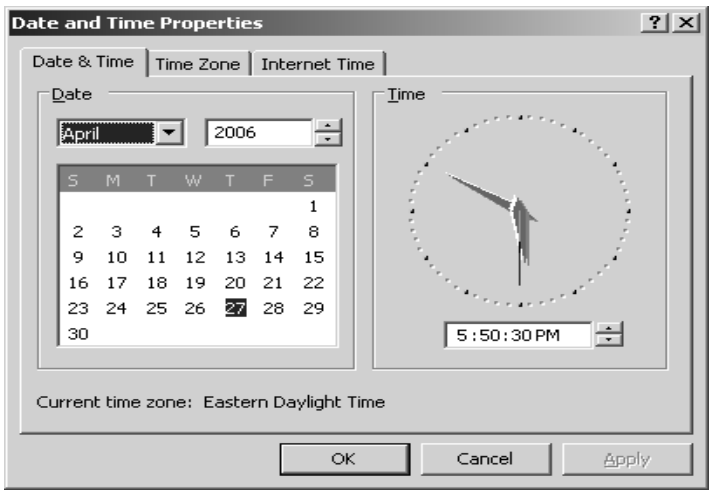

Figure 2-10 Date and Time Setting

Setting the system date and time

Double click the time display at the right lower corner of the Windows Storage Server 2003's desktop. This will bring up the Date and Time Properties dialog. You can set the date, time and time zone of the Storage Server. The Internet Time tab gives you the ability to synchronize the time on Storage Server to one of the public time server.

Viewing and maintaining event logs

A variety of event logs are provided on the Storage Server through Event Viewer. Event logs provide historical information that can help you track down system and security problems. Each log has viewing, clearing, printing, and saving options. To access the Event Viewer Dialog, Choose Start, All Programs, Administrative Tools, and then Event View. This displays the Event Viewer dialog.

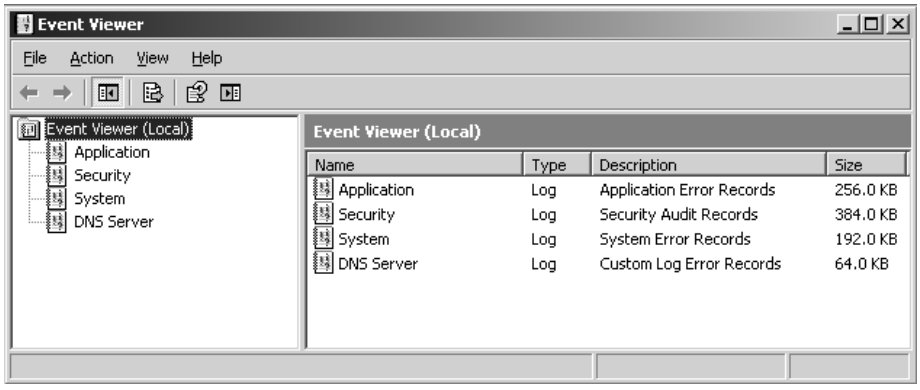

Figure 2-11. Event Logs

3 User and Group Management

With the Windows Server 2003 R2, there are two basic network environments for users: the Workgroup and the Domain. Because users and groups in a Domain environment are managed through standard Windows or Active Directory Domain administration methods, this document will discuss only local users and groups, which are stored and managed on the Storage Server. For information on managing users and groups in an Active Directory Domain, refer to Microsoft's website for resources.

Domain compared to Workgroup environments

When a Storage Server is deployed into a Workgroup environment, all user and group account access permissions to file resources are stored locally on the server.

By contrast, when a Storage Server is deployed into a Domain environment it uses the account database from the Domain controller, with user and group accounts stored outside the server; the Storage Server integrates with the Domain controller infrastructure.

For Domain environments, the Storage Server can't act as a Domain controller for other servers on the network. If user and group account information is stored locally, those accounts may be used only to authenticate logons to the Storage Server itself, resulting in a *de facto* Workgroup configuration.

User and group name planning

Effective user and group management depends upon how well the user and group names are organized. Administrators typically create a small number of groups on the network and then assign users to the appropriate group or groups. File system and share permissions can then be applied at the group level, rather than at the user level. If the number of groups is small, assigning the appropriate permissions to the selected group, or groups, is more efficient than assigning permissions to every user individually.

Although each organization has specific conventions, following general guidelines makes administration simpler and more efficient. Because CIFS/SMB is dependent on users and groups to grant appropriate access levels to file shares, CIFS/SMB administration benefits from a consistent user and group administration strategy.

Managing user names

Usernames should reflect a logical relationship between the username and the person who uses the account. It is important that rules are established to ensure that usernames are:

- Systematic.
- Easy to follow and implement.
- Easy to remember.

Using a combination of the user's first name, middle initial, and last name results in systematic usernames for every member of a particular organization. For example, first initial followed by last name (jdoe for John Doe).

Guidelines must be in place for instances when two users have the same initials or name. For example, a number can be added to the end of the username (jdoe1, jdoe2, etc.).

Naming conventions can be applied in a number of ways depending on the situation; just ensure that, however it is done, it is done systematically and consistently.

Managing group names

Group management follows many of the same principles as user management.

Group naming conventions should be systematic and easy to understand. Make the group name convey some logical information about the function or purpose of the group. Table 3-1 provides examples of group names.

Table 3-1 Group name examples

Group Name	Description
Administrators	All designated administrators on the server
Users	All standard server users
Power users	All standard server users requiring advanced access levels

Using tags is a helpful convention that indicates the specific access that a particular user has to a network resource. For example, if there is a data share on a device, the network administrator can create a 'Data Users Only' group for those who are authorized to view the stored information, as well as a 'Data Users Write' group to contain users that can both read from and write to the storage media.

Workgroup user and group management

In a Workgroup environment, local users and groups are managed through the Computer Management console. To open the Computer Management console, choose 'Start' → 'All Programs' → 'Administrative Tools' and then select 'Computer Management'.

Managing local users

To manage local users, open the Computer Management Console expand 'System Tools', 'Local Users and Groups', and then select 'Users' to display all the local users on the Storage Server. Figure 3-1 shows the local users page. All Workgroup user administration tasks are performed here.

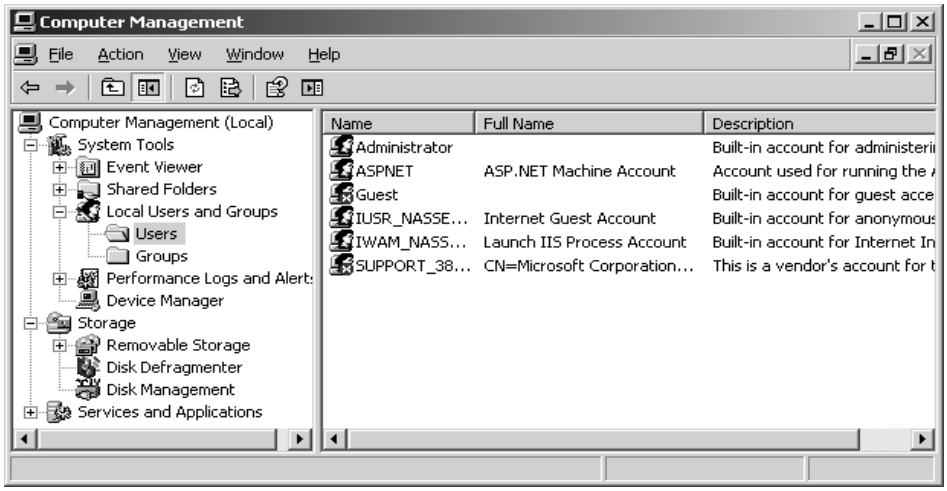

Figure 3-1 Local Users page

Adding a new user

To add a user:

1. Right-click 'Users' and then select 'New User'.
2. Click 'Create' when you're finished configuring the new account.

Deleting a user

To delete a user:

1. On the 'Local Users' page, right-click the user you want to delete, and then select 'Delete'. This opens the 'Delete User' dialog box, including a warning note about deleting users.
2. To delete the user, click 'Yes'.

Modifying a user password

To modify a user password:

1. On the 'Local Users' page, right-click the user whose password needs to be changed, and then select 'Set Password'. This opens the 'Set Password' dialog box, including a warning note about change password for users.

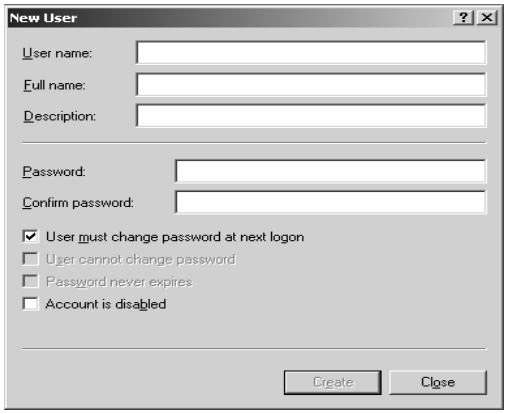

Figure 3-2 Create New User page

2. Enter the password, and then click ‘OK’.

Modifying user properties

To modify other user properties:

1. On the Local Users page, right-click the user whose record needs to be modified, and then select Properties. This opens the User Properties dialog box, shown in Figure 3-3.

2. Complete the changes, and then click **OK.**

Figure 3-3. User Properties page

Managing local groups

To manage local groups, open the Computer Management console; expand System Tools, Local Users and Groups, and then select Groups to display all the local groups on the Storage Server. Figure 3-4 shows the local groups page. All Workgroup user group administration tasks are performed here.

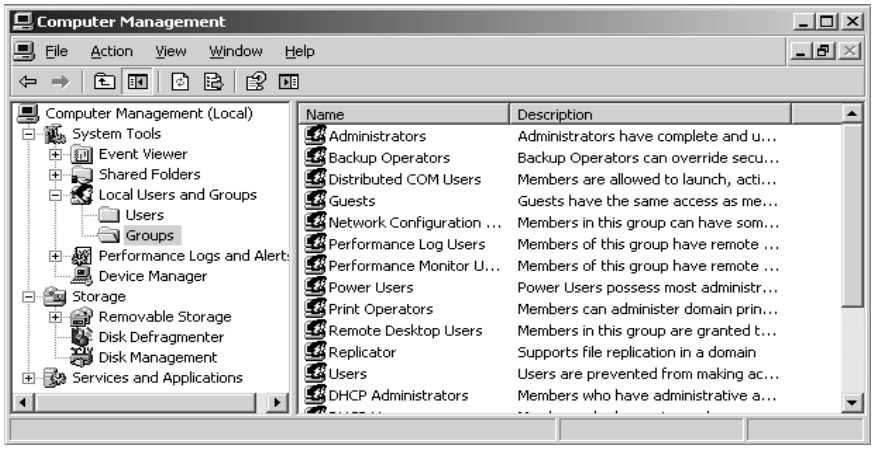

Figure 3-4 Local Groups page

Adding a new group

To add a group:

1. In Computer Management console tree, right-click Groups and select New Group. This opens the New Group dialog, shown in Figure 3-5.
2. Enter the group name and description.
3. Click Add to add users to be the members.
4. After all group information is entered, click Create.

Figure 3-5 Create New Group page

Deleting a group

To delete a group:

1. On the Local Groups page, right-click the group you want to delete, and then select Delete. This opens the Delete Group dialog box, including a warning note about deleting group.
2. Verify that this is the intended group, and then click Yes.

Modifying group properties

To modify group properties:

1. On the 'Local Groups' page, right-click the group whose record needs to be modified, and then select 'Properties'. This opens the 'Group Properties' dialog box, shown in Figure 3-6.
2. You can change the description and add or remove group members. When finished, click 'OK'.

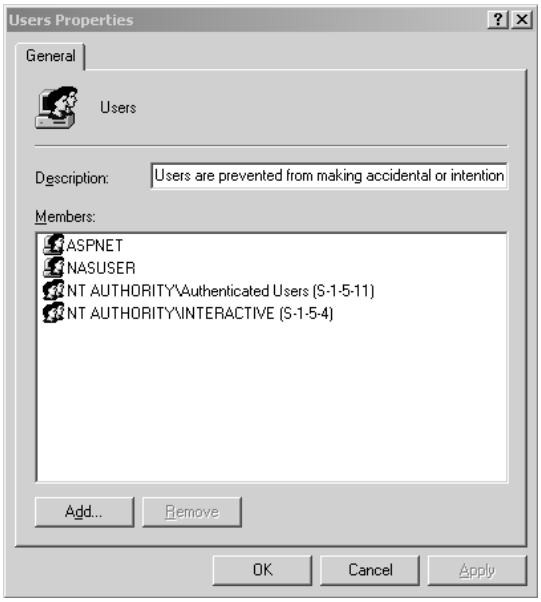

Figure 3-6. Group Properties page, 'General' tab

4 Disks and Volumes

RAID

On Storage Servers, physical disks can be arranged as RAID arrays, and then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the Storage Server. Because some array types enhance performance while others improve reliability, it is important to consider your needs when planning your array configuration.

Here is a list of common array types:

- **RAID0:** RAID0, or striping, provides the highest performance but no data redundancy. Data in the array is striped (distributed) across several physical drives. RAID0 arrays are useful for holding information such as the operating system paging file where performance is extremely important but redundancy is not.
- **RAID1:** RAID1, or mirroring, mirrors data on a partition of one disk to another. RAID1 is useful when there are only two disks available and data integrity is more important than storage capacity.
- **RAID10:** RAID10 is also known as RAID (0+1) or striped mirror sets. This array type combines mirrors and stripe sets. RAID10 allows multiple drive failures, up to 1 failure in each mirror that has been striped. This array type offers better performance than a simple mirror because of the extra drives. RAID10 requires twice the disk space of RAID0 to offer redundancy.
- **RAID5:** RAID5, also known as a stripe with parity, stripes data as well as parity across all drives in the array. Parity information is interspersed across the drive array. In the event of a failure, the controller can restore the lost data of the failed drive from the other surviving drives. This array type offers exceptional read performance as well as redundancy. In general, write performance is not an issue due to the tendency of operating systems to perform many more reads than writes. This array type requires only one extra disk to offer redundancy. For most systems with four or more disks, this is the correct choice as array type.

Digiliant storage servers use several different types of RAID controllers based on the server model and customer requirements. Consult the RAID controller manual that shipped with your server for instructions on how to properly manage the disk array. In most cases the RAID array management software can be accessed locally or via Remote Desktop by simply clicking the management software icon.

Disks and Volumes

The RAID controller presents the array to Windows Storage Server 2003 R2 as a raw logical disk. In order for the server to be able to utilize the array, this raw disk space must be configured; this is accomplished by partitioning the logical disk and formatting the allocated space. A disk *partition* is a section of a disk that functions as an independent unit, while *formatting* refers to the process of adding structural information to the partition that will support the necessary file system. This process of partitioning and formatting make the storage space provided by the RAID volume accessible to the operating system.

Windows Storage Server 2003 R2 supports *two types of partitions*: Master Boot Record (**MBR**) and Globally Unique Identifier Partition Table (**GPT**). The key difference between the GPT partition style and MBR partition style has to do with how partition data is stored. With GPT, critical partition data is stored in the individual partitions while redundant primary and backup partition tables are created for improved structural integrity. MBR disks support up to 4 partitions with maximum size of 4 TB (terabytes). GPT disks support volumes of up to 18 EB (exabytes) and up to 128 partitions. Any Digiliant Storage Server exceeding 2Tb in size will be converted to GPT disk.

After you set the partition type for a specific logical disk it must be configured. Windows Storage Server 2003 R2 supports two types of disk configurations.

- **Basic Disk:** The standard disk type used in previous version of Windows. Basic disks are divided into partitions and can be used with previous versions of Windows.
- **Dynamic Disk:** An enhanced disk type for Windows Storage Server 2003 R2 that can be updated without having to restart the system. Dynamic disks are divided into volumes and can be used only with Windows 2000 and Windows Storage Server 2003 R2.

Disk configuration tasks that you can perform with basic and dynamic disks are different.

With basic disks, you can:

- Format partitions and mark them active.
- Create and delete primary and extended partitions.
- Create and delete logical drives within extended partitions.
- Convert from a basic disk to a dynamic disk.

With dynamic disks, you can:

- Extend simple or spanned volumes.
- Split a volume into two volumes.
- Repair mirrored or RAID-5 volumes.
- Reactivate a missing or offline disk.
- Revert to basic disk from a dynamic disks (requires deleting volumes and reload).

With either disk type, you can:

- View properties of disks, partitions, and volumes.
- Make drive letter assignments.
- Configure security and drive sharing.

You'll use Disk Management to configure disks. One way to access the Disk Management console is through the Windows Storage Server Management console. In the Windows Storage Server Management console tree, expand 'File Server Management'→ 'Disk and Volume Management' and select 'Disk Management'. This displays Disk Management console shown in Figure 4-1.

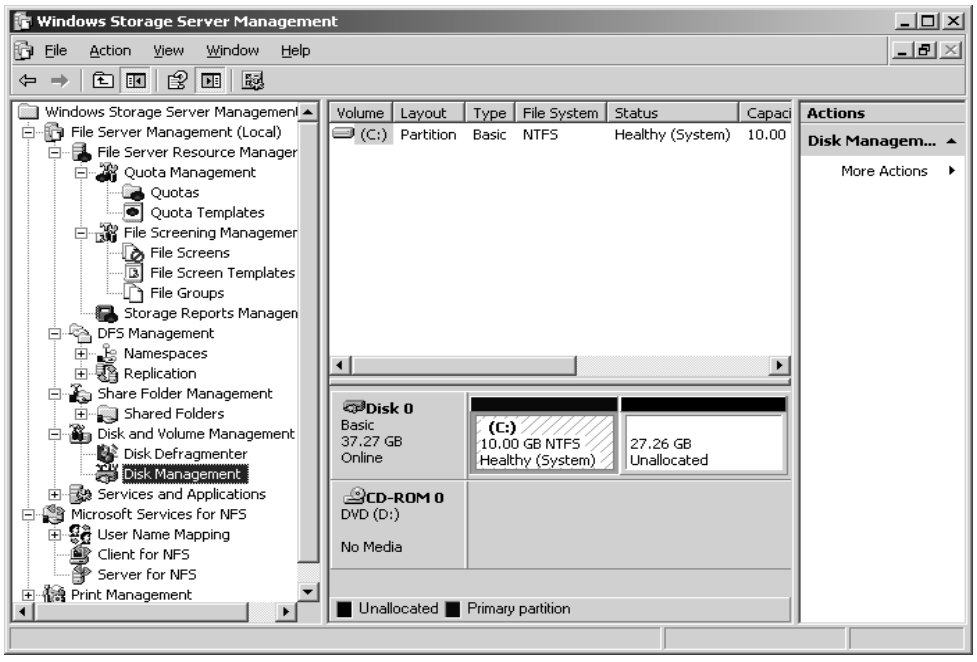

Figure 4-1 Disks and Volumes

Changing Disk Types

With Disk Manager, you can convert a basic disk to a dynamic disk and change a dynamic disk back to a basic disk. When you convert to a dynamic disk, partitions are changed to volumes of the appropriate type automatically. You can't change these volumes back to partitions. Instead, you must delete the volumes on the dynamic disk and then change the disk back to a basic disk. Deleting the volumes destroys all the information on the disk.

Converting a Basic Disk to a Dynamic Disk

In general, the system disk (OS drive) should be kept as basic disk. To convert a basic disk to a dynamic disk, complete the following steps:

1. In Disk Manager, right-click a basic disk that you want to convert, either in the 'Disk List' view or in the left pane of the 'Graphic View'. Then select 'Convert to Dynamic Disk'.
2. In the 'Convert to Dynamic Disk' dialog box, select the check boxes for the disks you want to convert. If you're converting a spanned, striped, mirrored, or RAID-5 volume, be sure to select all the basic disks in this set. *You must convert the set together*. Click 'OK' when you're ready to continue.
3. The 'Disk to Convert' dialog box shows the disks you're converting. The buttons and columns on this dialog box contain the following information:
    - **Name:** Shows the disk number.
    - **Disk Contents:** Shows the type and status of partitions, such as boot, active, or in use.
    - **Will Convert:** Specifies whether the drive will be converted. If the drive doesn't meet the criteria, it won't be converted, and you might need to take corrective action.
    - **Details:** Shows the volumes on the selected drive.
    - **Convert:** Starts the conversion.
4. If you're ready to begin the conversion, click 'Convert'. Disk Management warns you that after you finish the conversion you won't be able to boot previous versions of Windows from volumes on the selected disks. Click 'Yes' to continue.
5. Disk Management will restart the computer if a selected drive contains the boot partition, system partition, or a partition in use.

Changing a Dynamic Disk Back to a Basic Disk

Before you can change a dynamic disk back to a basic disk, you must delete all dynamic volumes on the disk. After you do this, right-click the disk and select the 'Convert to Basic Disk' command. This changes the dynamic disk to a basic disk and you can then create new partitions and logical drives on the disk.

Reactivating Dynamic Disks

If the status of a dynamic disk displays as 'Online (Errors)' or 'Offline', you can often reactivate the disk to correct the problem. You reactivate a disk by completing the following steps:

1. In Disk Management, right-click the dynamic disk you want to reactivate, and then select 'Reactivate Disk'. Confirm the action when prompted.
2. If the drive status doesn't change you might have to reboot the computer. If this still doesn't resolve the problem, check for problems with the drive, its controller, and/or the cables.

Rescanning Disks

Rescanning all drives on a system updates the drive configuration information on the computer. It can sometimes resolve a problem with drives that show a status of 'Unreadable'. You can rescan disks on a computer by selecting 'Rescan Disk' from the 'Disk Management' 'Action' menu.

Using Basic Disks and Partition

Windows Storage Server 2003 R2, a drive using the MBR partition style can have up to four primary partitions and up to one extended partition. Drives with GPT partition style can have up to *128* partitions.

Creating Partitions

In Disk management, you create partitions and logical drives by completing the following steps;

1. In the Disk Management Graphical view, right-click an area marked 'Unallocated' and then choose 'New Partition', or right-click a free space in an extended partition and select 'New Logical Drive'. This starts the 'New Partition Wizard'.
2. Click 'Next'. As shown in Figure 4-2, you can now select a partition type as follow:
    - **Primary Partition:** A primary partition can fill an entire disk, or you can size it as appropriate.
    - **Extended partition:** Each drive can have one extended partition. This extended partition can contain one or more logical drives, which are simply sections of the partition with their own file system.
    - **Logical Drive:** This will create a logical drive within an extended partition.

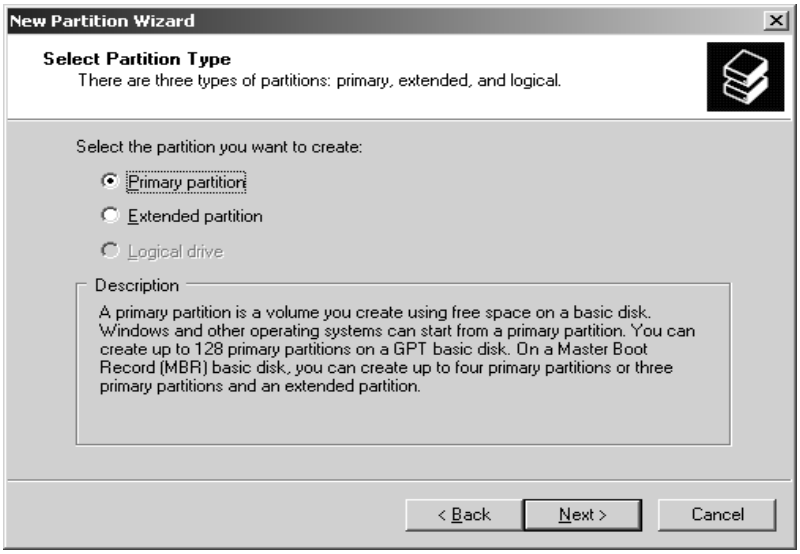

Figure 4-2 New Partition Wizard, Select Partition Type

3. You should see the 'Specify Partition Size' page. This page specifies the minimum and maximum size for the partition in megabytes and lets you size the partition within these limits. Size the partition using the 'Amount of Disk Space to Use' field.
4. Specify whether you want to assign a drive letter or path, as shown in Figure 4-3. You use these options as follows:
    - **Assign The Following Drive Letter:** To assign a drive letter, choose this option and then select an available drive letter in the selection list provided.
    - **Mount in the Following Empty NTFS Folder:** To assign a drive path, choose this option and then type the path to an existing folder or click Browse to search for or create a folder.
    - **Do Not Assign A Drive Letter Or Drive Path:** To create the partition without assigning a drive letter or path, choose this option. You can assign a drive letter or path later, if necessary.

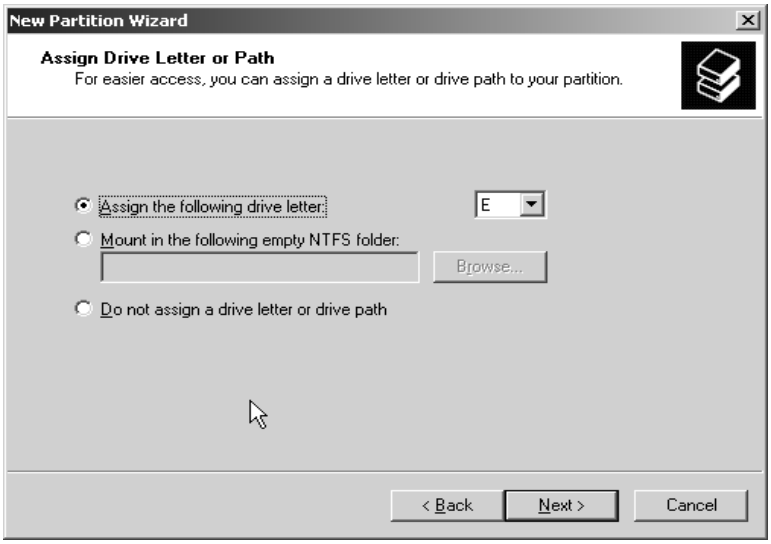

Figure 4-3 Assign Drive Letter or Path for Partition

5. Determine whether the partition should be formatted in the Format Partition page, shown in Figure 4-7. You use the formatting fields as follows:
    - **Volume Label**: This is the partitions' volume name.
    - **File System:** You should select NTFS as the file system. NTFS is the native file system for Windows Storage Server 2003 R2.
    - **Allocation Unit Size:** Specifies the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the size of the volume and is set dynamically prior to formatting. To override this feature, you can set the allocation unit size to a specific value. If you use lots of small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space.
    - **Perform A Quick Format:** Tells Windows Storage Server 2003 R2 to format without checking the partition for errors. With large partitions this option can save you some time.
    - **Enable File And Folder Compression:** Turns on compression for the disk. Built-in compression is available only for NTFS.
6. Click 'Next', then 'Finish'. If you add partitions to a drive that contains the Windows Storage Server 2003 operating system, you might inadvertently change the number of the boot partition. Windows Storage Server 2003 will display a prompt warning you that the number of the boot partition will change. Click 'Yes'.

Managing Existing Partitions and drives

Disk Management provides many ways to manage existing partitions and drives. Use these features to assign drive letters, delete partitions, set the active partition, defragmenting a drive, clean-up unused disk space and more.

Assigning Drive Letters and Paths

You can assign drives one drive letter and one or more drive paths, provided the drive paths are mounted on NTFS drives. Drives don't have to be assigned a drive letter or path. A drive with no designators is considered to be 'unmounted', and you can mount it by assigning a drive letter or path at later date.

To manage drive letters and paths, right-click the drive you want to configure in Disk Management, and then choose 'Change Drive Letter and Paths'. This opens the dialog box shown in Figure 4-4. You can now:

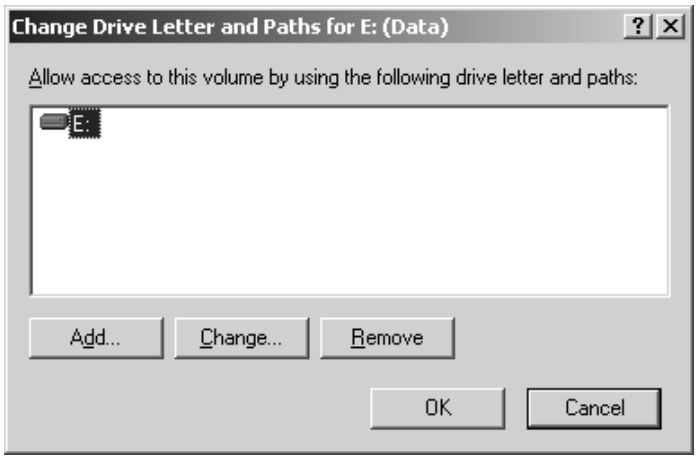

Figure 4-4 Change Drive Letter and Paths

- **Add a drive path:** Click 'Add', select 'Mount in the Following Empty NTFS Folder', and then type the path to an existing folder or click 'Browse' to search for or create a folder.
- **Remove a drive path:** Select the drive path to remove, click 'Remove', and then click 'Yes'.
- **Assign a drive letter:** Click 'Add', select 'Assign a Drive Letter', and then choose an available letter to assign to the drive.
- **Change the drive letter:** Select the current drive letter then click 'Change'. Select 'Assign a Drive Letter' then choose a different letter to assign to the drive.
- **Remove a drive letter:** Select the current drive letter, click 'Remove', and then click 'Yes'.

Changing or Deleting the Volume Label

The volume label is a text descriptor for a drive. Using Disk Management, you can change or delete a label by following these steps:

1. Right-click the partition, and then choose Properties.
2. In the general tab of the Properties dialog box, use the Label text box to type a new label for the volume or delete the existing label. Click OK.

Deleting Partitions and Drives

To change the configuration of an existing drive that's fully allocated, you might need to delete existing partitions and logical drives. Deleting a partition or a drive removes the associated file system, and all data in the file system is lost. So before you delete a partition or a drive, you should back up any files and directories that the partition or drive contains.

You can delete a primary partition or logical drive by following these steps:

1. In Disk Management, right-click the partition or drive you want to delete, and then choose Delete Partition or Delete Logical Drive, as appropriate.
2. Confirm that you want to delete the partition by clicking 'Yes'.

To delete an extended partition, follow these steps:

1. Delete all the logical drives on the partition following the steps listed above.
2. You should now be able to select the extended partition area itself and delete it.

Defragmenting Disks

Anytime you add files to or remove files from a drive, the data on the drive can become fragmented. When a drive is fragmented, large files can't be written to a single continuous area on the disk. As a result, the operating system must write the file to several smaller areas on the disk, which means more time is spent reading the file from the disk. To reduce fragmentation, you should periodically analyze and defragment disks using 'Disk Defragmenter'.

Using Volumes and Volume sets

When you work with Windows Storage Server 2003, you often need to create a volume set or setting up a software RAID array. You can create volumes sets and RAID arrays on dynamic drives.

- With a volume set, you can create a single volume that spans multiple drives. Users can access this volume as if it were a single drive, regardless of how many drives the actual volume is spread over. A volume that's on a single drive is referred to as a simple volume. A volume that spans multiple drives is referred to as spanned volume.
- With software RAID array, you can add another level of protection on top of Digiliant's hardware RAID. Combine two hardware RAID supported drive into a software RAID 0, you can improve the performance. Windows Storage Server 2003 supports three levels of RAID 0, 1, and 5. RAID arrays are implemented as mirrored, striped, and striped with parity volumes.

You create and manage volumes in much the same way as partitions. With volumes, you can:

- Assign drive letters as the same ways as partitions.
- Assign drive path as the same ways as partitions.
- Create any number of volumes on a disk as long as you have free space.
- Create volumes that span two or more disks and, if necessary, configure fault tolerance.
- Extend volumes to increase the volume's capacity.

Creating Volumes and Volume sets

You create volumes and volume sets by completing the following steps:

1. In the Disk Management Graphic View, right-click an area marked 'Unallocated' on a dynamic disk and then choose 'New Volume'. This starts the 'New Volume Wizard'. Read the 'Welcome to the New Volume Wizard' page and then click 'Next'.

2. As shown in Figure 4-5, select 'Simple' to create a volume on a single disk or 'Spanned' to create a volume set on multiple disks. You can format simple volumes as FAT, FAT32, or NTFS. To make management easier, you should format volumes that span multiple disks as NTFS. NTFS formatting allows you to expand the volume set if necessary. Click 'Next'.

3. You should see the 'Select Disks' page shown in Figure 4-6. Use this page to select dynamic disks that are a part of the volume and to size the volume segments on those disks.

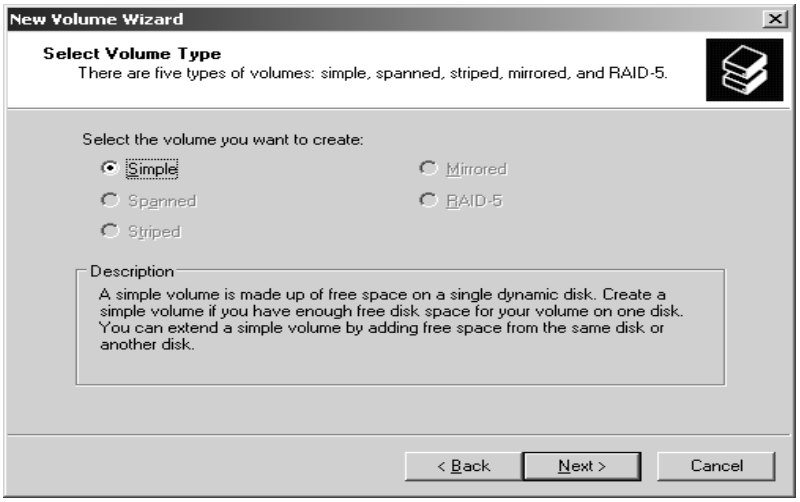

Figure 4-5 Select Volume Type

4. Available dynamic disks are shown in the 'Available' list box. If necessary, select a disk in this list box and click 'Add' to add the disk to the 'Selected' list box. If you make a mistake, you can remove disks from the 'Selected' list box by selecting the disk and then clicking 'Remove'.

5. Select a disk in the selected list box and then use the 'Select the amount of space in MB' combo box to specify the size of the volume on the selected disk. The 'Maximum' field shows you the largest area of free space available on the selected disk. The 'Total Volume Size' field shows you the total disk space selected for use with the volume. Click 'Next'.

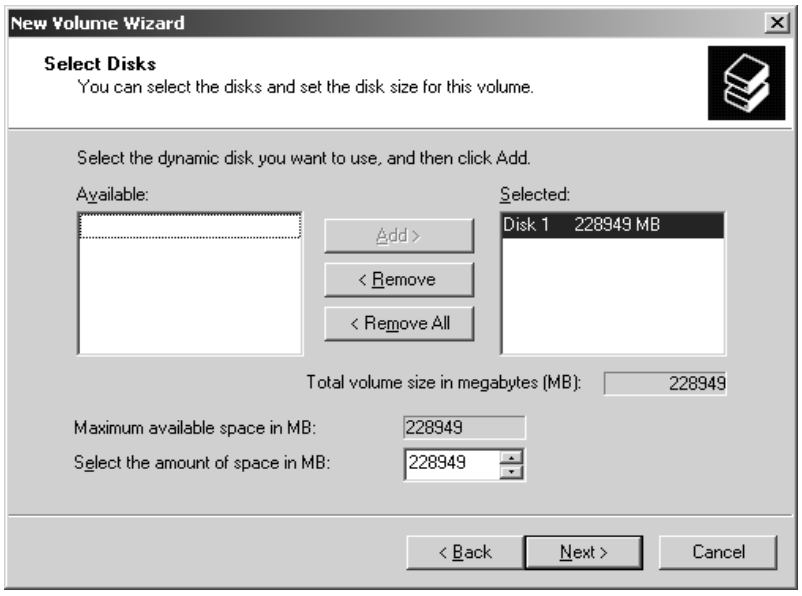

Figure 4-6 Add Disks to Volume

6. Specify whether you want to assign a drive letter or path to the volume, and then click 'Next'. You use these options as follows:

   - **Assign The Following Drive Letter:** To assign a drive letter, choose this option, and then select an available drive letter in the selection list provided.
   - **Mount in the Following Empty NTFS Folder:** To assign a drive path, choose this option, and then type the path to an existing *empty* folder or click Browse to search for or create a folder.
   - **Do Not Assign A Drive Letter Or Drive Path:** To create the volume without assigning a drive letter or path, choose this option. You can assign a drive letter or path later, if necessary.

7. As shown in Figure 4-7, determine whether the volume should be formatted. If you elect to format the volume, use the following fields to set the formatting options:

- **Volume Label:** This is the partition's volume name.
- **File System:** specifies the file system type. NTFS is the only option within Disk Management.
- **Allocation Unit Size:** Specifies the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the size of the volume and is set dynamically prior to formatting. To override this feature, you can set the allocation unit size to a specific value. If you use lots of small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space.
- **Perform A Quick Format:** Tells Windows Storage Server 2003 R2 to format without checking the partition for errors. With large partitions this option can save you some time.
- **Enable File And Folder Compression:** Turns on compression for the disk. Built-in compression is available only for NTFS.

8. Click 'Next' and then click 'Finish'. If you add volumes to a drive that contains the Windows Storage Server 2003 operating system, you might inadvertently change the boot volume's number. Windows Storage Server 2003 R2 will display a prompt warning you that the boot volume's number will change. Click 'Yes'.

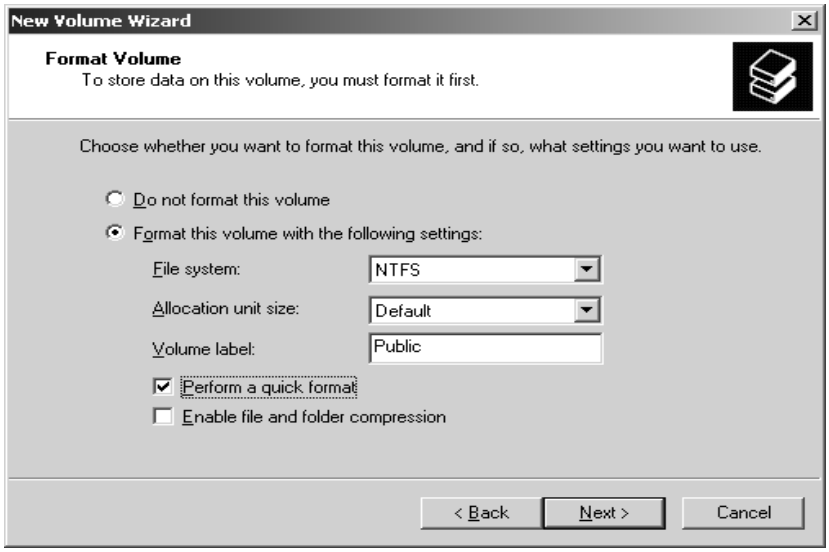

Figure 4-7 Format Volume

Deleting Volumes and Volume Sets

You delete all volumes using the same technique, whether they're simple, spanned, mirrored, striped, or RAID 5. Deleting a volume set removes the associated file system and all associated data is lost. So before you delete a volume set you should back up any files and directories that the volume set contains.

To delete volumes, follow these steps:

1. In Disk Management, right-click any volume in the set and then choose 'Delete Volume'. You can't delete a portion of a spanned volume without deleting the entire volume.
2. Confirm that you want to delete the volume by clicking 'Yes'.

Extending a Simple or Spanned Volume

Windows Storage Server 2003 R2 provides several ways to extend NTFS volumes that aren't part of a mirror set or striped set. You can extend both an individual simple volume and existing volume sets. When you extend volumes, you add free space to them. There are certain limitations regarding expanding volumes. You can't extend boot or system volumes. You can't extend volumes that use mirroring or striping. You can't extend a volume onto more than 32 disks.

To extend an NTFS volume, complete the following steps:

1. In Disk Management, right-click the simple or spanned volume that you want to extend, and then select 'Extend Volume'. This starts the 'Extend Volume Wizard'. Read the 'Welcome to the Extended Volume Wizard' page, and then click 'Next'.

2. You can now select dynamic disks that are a part of the volume, and size the volume segments on those disks. The size you set for ‘Select the amount of space in MB’ is the amount of space you want to add to the volume, which is reflected in the ‘Total Volume Size in Megabytes’ field. For example, if you created a 500MB volume and then set the ‘Select the amount of space in MB’ field to 300, the total volume size would be 800MB.
3. Click ‘Next’ and then click ‘Finish’.

5 Folder and Share Management

The Digiliant Storage Server supports several file sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This chapter covers general information as well as procedural instructions for the setup and management of the file shares for supported protocols. Security at the file level and at the share level is also touched upon.

Folder management

Volumes and folders on any system are used to organize data. Regardless of system size, organization is key; the systematic structuring and naming of volumes and files will ease the burden of administrating your Digiliant Storage Server. Moving from volumes to folders to shares increases the level of granularity of data types stored on the server as well as the level of security access allowed.

Although a variety of methods can be used to create and manage file folders on the Storage Server, this document will use Windows Explorer.

Managing system volumes and file folders includes the following tasks:

Navigating to a specific volume or folder

When you work with volumes and folders, the first task is to gain access to the desired volume or folder.

The steps are the same, whether navigating to a volume or a folder:

1. Start the Windows Explore by click 'My Computer' from the Windows Storage Server 2003 R2 desktop.
2. Click on the appropriate volume. This displays a list of all of the folders within that volume.
3. To navigate to a subfolder, click the folder in which the subfolder resides. Repeat this searching and opening process until the desired folder is opened.

Creating a new folder

To create a new folder:

1. Access Windows Explorer; navigate to the volume or folder on which the new folder will reside.
2. Right-click on a blank spot in the Windows Explorer list window, select 'New', and then select 'Folder'. This action adds a new folder named 'New Folder'. Change the folder name from 'New Folder' to a desired name.

Deleting a folder

To delete a folder:

1. Access Windows Explorer; navigate to the volume or folder on which the desired folder is resided.
2. Right-click on the folder you want to delete and choose 'Delete'. Click 'Yes' to confirm the action.

Modifying folder properties

To modify folder properties:

1. Access Windows Explorer; navigate to the folder whose properties need to be edited. Right-click the folder you want to manage, and then select 'Properties'. This opens the folder properties dialog as shown in Figure 5-1.

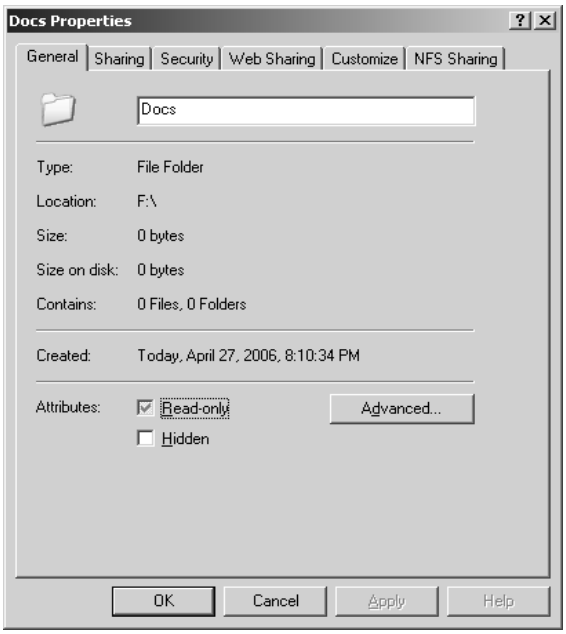

Figure 5-1 Folder Properties

2. On the ‘General’ tab, you can change the name of the folder and Attributes. For changing other advanced attributes (archiving, indexing, compression and encryption, etc.) you can click the ‘Advanced’ button. This opens the ‘Advanced Attributes’ dialog as shown in Figure 5-2.
3. Other properties such as ‘Sharing’ and ‘Security’ will be discussed later in this chapter.

Figure 5-2 Advanced Folder Properties, Attributes

Share management

There are several ways to set up and manage shares. Windows Storage Server Management Console provides a centralized place for setting up and managing shares. Additional methods include using a command line interface, Windows Explorer, or Computer Management Console. This guide demonstrates using the Windows Storage Server Management Console to set up and manage shares.

Share considerations

Planning the content, size, and distribution of shares on the Storage Server can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: (a) having too many shares of a very specific nature, or (b) having very few shares of a generic nature. For example, shares for general usage are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. For example, if it is sufficient to create a single share for user home directories, create a ‘`\\homes`’ share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the Storage Server is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as what type of access is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating local file system security into Windows Domain environments

ACLs include properties specific to users and groups from a particular Workgroup server or Domain environment. In a multi-Domain environment, user and group permissions from several Domains can apply to files stored on the same device. Users and groups local to the Storage Server can be given access permissions to shares managed by the device. The Domain name of the Storage Server supplies the context in which the user or group is understood. Permission configuration depends on the network and Domain infrastructure where the server resides. File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL, and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

List all share folders

To list all the share folders, access Windows Storage Server Management console, select File Server Management, Share Folder Management, Share Folders, and then finally Shares. This displays list of Share Folders in the Storage Server, shown in Figure 5-3.

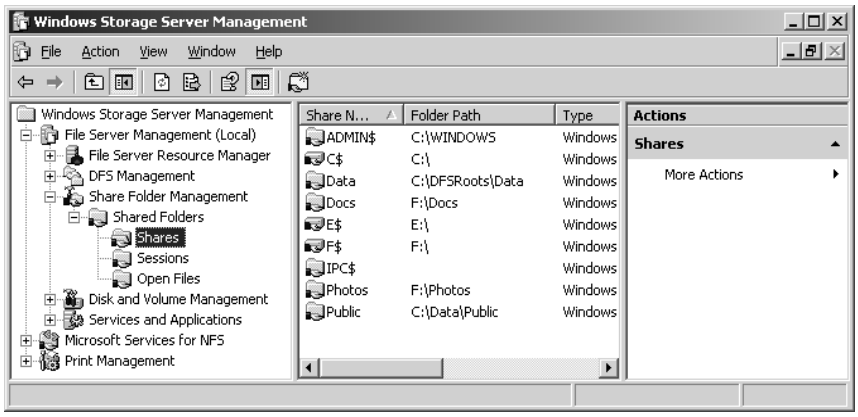

Figure 5-3 Share Folder List

Creating a new share folder

To create a new share folder:

1. Access Shares under Windows Storage Server Management, right-click on Shares and select New Share. This starts the Share a Folder Wizard. Click Next. Enter the folder path when prompted. Click Next.
2. In the Share Name text box, type a name for the share, as shown in Figure 5-4. This is the name of the folder to which users will connect. If you like you can type the description for the share.
3. By default, the share is configured so that only files and program that users specify are available for offline use. If you want to prohibit the offline use of files or programs in the share or specify that all files and programs in the share are available for offline use, click Change, and then select the appropriate options in the Offline Settings dialog box.
4. Click Next and then set basic permissions for the share. As shown in Figure 5-5, the available options are as follows:

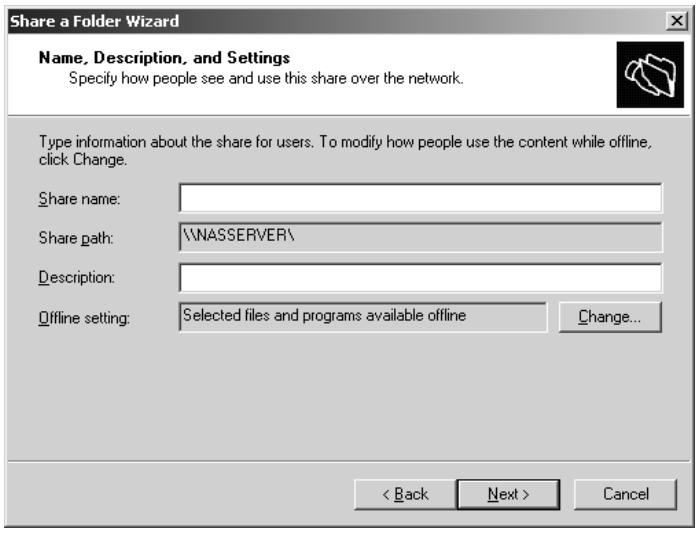

Figure 5-4 Create a New Folder page

- **All users have read-only access:** Gives users access to view files and read data. They can't create, modify, or delete files and folders.
- **Administrators have full access; other users have read-only access:** Gives administrators full access to create, modify, and delete files and folders. It also gives administrators the right to change permissions and to take ownership of files and folders. Other users can only view files and read data. They can't create, modify, or delete files and folders.
- **Administrators have full access; other users have read and write access**: Gives administrators complete control over the share and allows other users to create, modify, or delete files and folders.
- **Use custom share and folder permissions:** Allows you to configure access for specific users and groups, which is usually the best technique to use.

5. When you click 'Finish', the wizard displays a status report, which should state 'Sharing Was Successful'. Click 'Close'.

Managing Share Permissions

Share permissions set the maximum allowable actions available within a shared folder. By default, when you create a share, every one with access to the network has read access to the share's contents. You can assign share permissions to users and groups.

Figure 5-5 Folder Share Permissions

The Different Share Permissions

Share permissions available, from the most restrictive to the least restrictive, are:

- **No Access:** No permissions are granted for the share.
- **Read:** With this permission, users can:
  1. View file and subfolder names.
  2. Access the subfolders of the share.
  3. Read file data and attributes.
  4. Run program files.
- **Change:** Users have Read permissions and the additional ability to:
  1. Create files and subfolders.
  2. Modify files.
  3. Change attributes on files and subfolders.
  4. Delete files and subfolders.
- **Full Control:** Users have Read and Change permissions, as well as the following additional capabilities on NTFS volumes:
  1. Change file and folder permissions.
  2. Take ownership of files and folders.

View Share Permissions

To view share permissions, follow these steps:

1. Access Shares under Windows Storage Server Management Console.
2. Right-click the share you want to view, and then select 'Properties'.
3. In the 'Properties' dialog box, click the 'Share Permission' tab, shown in Figure 5-6.

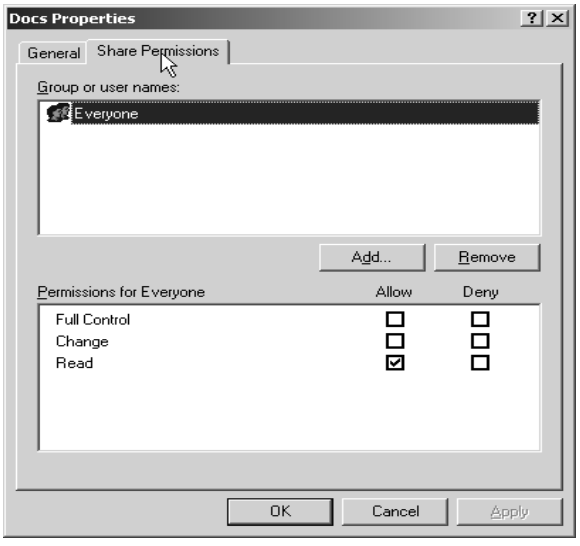

Figure 5-6 Share Permissions properties

Configuring Share Permissions

In Windows Storage Server Management, you can add user, computer, and group permissions to shares by completing the following steps:

1. Right-click the share you want to manage and then select 'Properties'.
2. In the 'Share Properties' dialog box, click the 'Share Permissions' tab.
3. Click 'Add'. This opens the 'Select Users or Groups' dialog box shown in Figure 5-7.
4. Type the name of user, computer, or group in the current Domain and then click 'Check Names'.
    - If a single match is found, the dialog box is automatically updated as appropriate and the entry is underlined.
    - If no matches are found, you've either entered an incorrect name part or you're working with an incorrect location. Modify the name and try again, or click 'Locations' to select a new location.
    - If multiple matches are found, select the name(s) you want to use and then click 'OK'. To assign permissions to other users, computers, or groups, type a semicolon (;) and then repeat this step.
5. Click 'OK'. The users and groups are added to the Name list for the share.
6. Configure access permissions for each user, computer, and group by selecting an account name and then allowing or denying access permissions. Keep in mind that you're setting the maximum allowable permissions for a particular account.

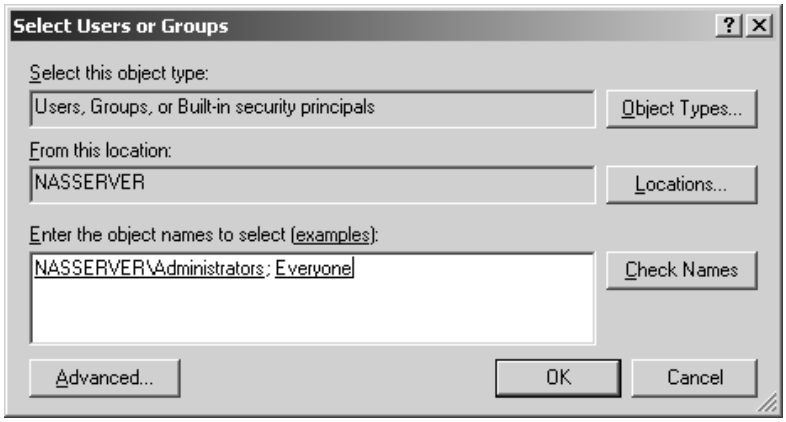

Figure 5-7 Add Objects to Share Permissions

Modifying Existing Share Permissions

You can change the share permissions you assign to users, computers, and groups by using the 'Share Properties' dialog box. In Windows Storage Server Management Console, follow these steps:

1. Right-click the share you want to manage, and then select 'Properties'.
2. In the 'Share Properties' dialog box, click the 'Share Permissions' tab.
3. In the Name list box, select the user, computer, or group you want to modify.
4. Use the check boxes in the 'Permissions' area to allow or deny permissions.
5. Repeat for other users, computers, or groups, and then click 'OK' when you're finished.

Remove Share Permissions for Users and Groups

1. Right-click the share you want to manage, and then select 'Properties'.
2. In the 'Share Properties' dialog box, click the 'Share Permissions' tab.
3. In the 'Name' list box, select the user, computer, or group you want to remove, and then click 'Remove'.
4. Repeat for other users or groups, and then click 'OK' when you're finished.

Managing Existing Shares

As an administrator, you'll often have to manage shared folders. The common administrative tasks of managing shares are covered in this section.

Understanding Special Shares

The Windows Storage Server 2003 R2 comes with some special shares which are created by the operating system automatically. These shares are also known as administrative shares and hidden shares. These shares are designed to help makes system administration easier. You can't set access permissions on automatically created special shares,

Windows Storage Server 2003 R2 assigns access permissions (You can create your own hidden shares by typing '$' as the last character of resource name.).

Connecting to Special Shares

Special shares end with the '$' symbol. Although these shares aren't displayed in Windows Explorer, administrators and certain operators can connect to them. To connect to a special share, follow these steps:

1. In Windows Explorer, from the 'Tools' menu, select 'Map Network Drive'. This opens the page shown in Figure 5-8.
2. From the Drive drop-down list, select a free drive letter. This drive letter is used to access the special share.
3. In the 'Folder' text box, type the Universal Naming Convention (UNC) path to the desired share. For example, to access the 'C$' share on the Storage Server called 'NASSERVER', you'd use the path \\NASSERVER\C$. Click 'Finish'.

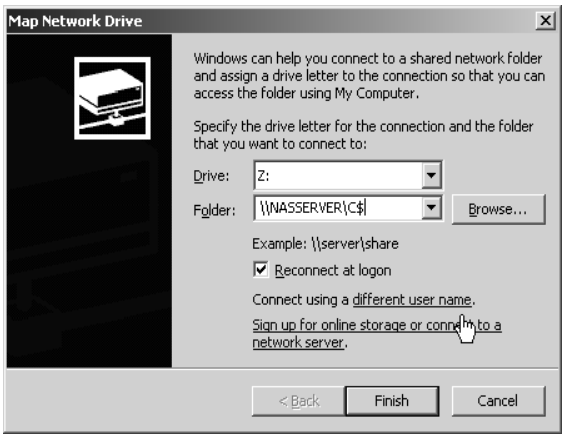

Figure 5-8 Map Network Drive

After you connect to a special share, you can access it as you would any other drive. Because special shares are protected, you don't have to worry about ordinary users accessing these shares. The first time you connect to the shares, you might be prompted for a user name and password. If you are, provide that information.

Viewing User and Computer Sessions

You can use Windows Storage Server Management Console to track all connections to shared resources on a Windows Storage Server 2003 R2 system. Whenever a user or computer connects to a shared resource, Windows Storage Server 2003 R2 lists a connection in the 'Sessions' node.

To view connections to shared resources, follow these steps:

1. Access Windows Storage Server Management Console. In the console tree: 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then select 'Sessions'.
2. As shown in Figure 5-9, you can view connections to shares for users and computers.

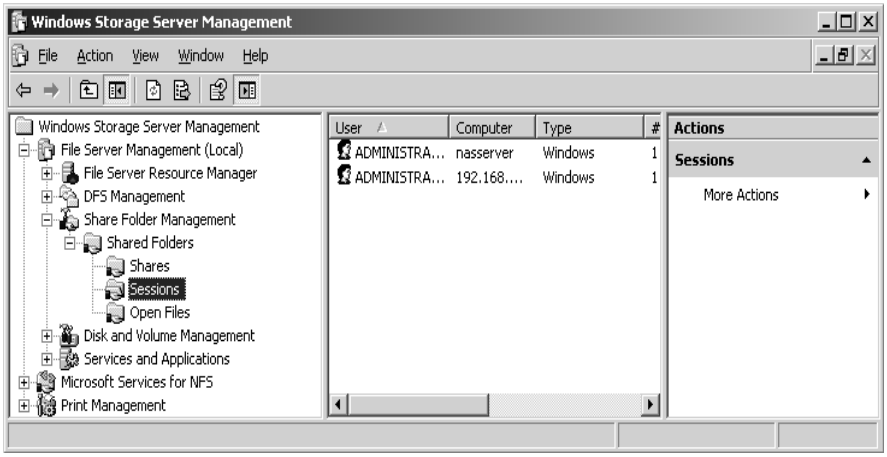

Figure 5-9 Connections to Shares

The 'Sessions' node provides important information about user and computer connections. The columns of this node provide the following information:

- **User:** The names of users or computers connected to shared resources. Computer names are shown with '$' suffix to differentiate them from users.
- **Computer:** The name of the computer being used.
- **Type:** The type of network connection being used.
- **#Open Files:** The number of files with which the user is actively working.

- **Connected Time:** The time that has elapsed since the connection was established.
- **Idle Time:** the time that has elapsed since the connection was last used.
- **Guest:** Whether the user is logged on as guest.

Managing Sessions and Shares

Managing sessions and shares is a common administrative task. Before you shut down a server, or an application running on a server, you might want to disconnect users from shared resources. You may also need to disconnect users when you plan to change access permissions or delete a share entirely. Another reason to disconnect users is to break locks on files. You disconnect users from shared resources by ending the related user sessions.

Ending Individual Sessions

To disconnect individual users from shared resources, follow these steps:

1. Access Windows Storage Server Management Console. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then select 'Sessions'.
2. Right-click the user sessions you want to end and then choose 'Close Session'.
3. Click 'Yes' to confirm the action.

Ending All Sessions

To disconnect all users from shared resources, follow these steps:

1. Access Windows Storage Server Management Console. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then right-click 'Sessions'.
2. Choose 'Disconnect All Sessions' and then click 'Yes' to confirm the action.

Managing Open Resources

Any time users connect to shares, the individual file and object resources with which they're actively working are displayed in the 'Open Files' node. The 'Open Files' node might show the files the user has open but isn't currently editing.

You can access the Open Files node by completing the following steps:

1. Access Windows Storage Server Management Console.
2. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then select 'Open Files'. This displays the 'Open Files' node, shown in Figure 5-10.

The 'Open Files' node provides the following information about resource usage:

- **Open File:** The file or folder path to the open file on local system. It might also be a named pipe, such as \PIPE\spools, which is used for printer spooling.

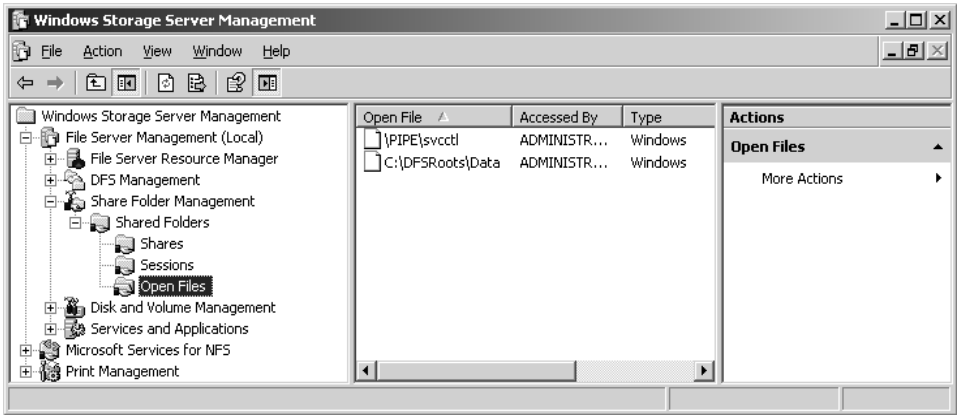

Figure 5-10. Properties dialog box, Security tab

- **Accessed By:** The name of user accessing the file.
- **Type:** The type of network connection being used.

- **# Locks:** the number of locks on the resource.
- **Open Mode:** The access mode used when the resource was opened, such as read, write or write+read mode.

Close an Open File

To close an open file on a computer's shares, follow these steps;

1. Access Windows Storage Server Management Console.
2. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then select 'Open Files'.
3. Right-click the open file you want to close, and then choose 'Close Open File'.
4. Click 'Yes' to confirm the action.

Close All Open File

To close all open files on a computer's shares, follow these steps;

1. Access the Windows Storage Server Management Console.
2. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then right-click 'Open Files'.
3. Choose 'Disconnect All Open File's' and then click 'Yes' to confirm the action.

Stopping File and Folder Sharing

To stop sharing a folder, follow these steps:

1. Access the Windows Storage Server Management Console.
2. In the console tree, expand 'File Server Management' → 'Share Folder Management' → 'Shared Folders', and then select 'Shares'.
3. Right-click the share you want to remove and then choose 'Stop Sharing'. Click 'Yes' to confirm the action.

Disk Quotas

Disk quotas track and control disk space use in volumes. You can set up the volumes to:

- Limit a user's disk space usage
- Log events when the user exceeds a specified early warning limit.
- Log an event when a user exceeds a specified warning level.

When you enable disk quotas, you can set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user may use, and the warning level specifies the point at which a warning occurs. For example, if you set a user's disk quota limit to 50 MB and the disk quota warning level to 45 MB, the user can store no more than 50 MB on the volume; when he reaches 45 MB the disk quota system logs a system event. You also can also allow a user to exceed the quota if you prefer logging usage without actually blocking the user from access.

Disk quotas are not retroactive. When you enable disk quotas on a volume, every user's disk volume usage is monitored and treated differently, depending on the quota management settings for the specific user.

You can set NTFS disk quotas on a per-volume basis. Only NTFS volumes can have disk quotas.

Enabling NTFS Disk Quotas on NTFS Volumes

To enable quota management on a volume, follow these steps:

1. In Disk Management, right-click the volume you want to work with and then select ‘Properties’.
2. Click the ‘Quota’ tab and then select ‘Enable Quota Management’ check box, as shown in Figure 5-11.
3. To set a default disk quota limit for all users, select ‘Limit Disk Space to’ and then use the text boxes provided to set a limit in KB, MB, GB, TB, PB, or EB. Afterward, use the ‘Set Warning Level to’ text boxes to set the default warning limit. Again, you’ll usually want the disk quota warning limit to be 90-95 percent of the disk quota limit.
4. To enforce the disk quota limit and prevent users from going over the limit, select the ‘Deny Disk Space to Users Exceeding Quota Limit’ check box.
5. To configure logging when users exceed a warning limit or the quota limit, select the ‘Log Event’ check box. Click ‘OK’ to save your changes.

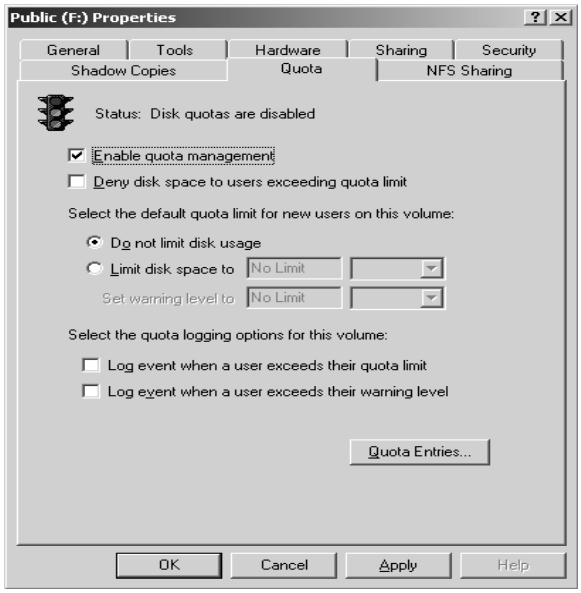

Figure 5-11 Enable Quota Management on Volume

6. If the quota system isn’t currently enabled, you’ll see a prompt asking you to enable the quota system. Click ‘OK’ to allow Windows to rescan the volume and update disk usage statistics. Actions might be taken against users who exceed the current limit or warning levels. Theses actions can include preventing additional writing to the volume, notifying them the next time they access the volume, and logging applicable events in the Application log.

Viewing Disk Quota Entries

Disk space usage is tracked on a per user basis. When disk quotas are enabled, each user storing data on a volume has an entry in the disk quota file. This entry is updated periodically to show the current disk space used, the applicable quota limit, the applicable warning level, and the percentage of allowable space being used. As an administrator, you can modify disk quota entries to set different limits and warning levels for particular users. You can also create disk quota entries for users who haven’t yet saved data on a volume. The key reason for creating entries is to ensure that when a user does make use of a volume the user has an appropriate limit and warning level.

To view the current disk quota entries for a volume, follow these steps:

1. In Disk Management, right-click on the volume with which you want to work and then select ‘Properties’.
2. In the ‘Quota’ tab, click ‘Quota Entries’. This displays the ‘Quota Entries’ dialog box. Each quota entry is listed according to a status. The status is meant to quickly depict whether a user has gone over a limit. A status of ‘OK’ means the user is working within the quota boundaries. Any other status usually means the user has reached the warning level or the quota limit.

Creating Disk Quota Entries

You can create disk quota entries for users who haven’t yet saved data on a volume. This allows you to set custom limits and warning levels for a particular user. You’ll usually use this feature when a user frequently stores more information than other users and you want to allow the user to go over the normal limit or when you want to set a specific limit for administrators.

To create a quota entry on a volume, follow these steps:

1. Access the ‘Quota Entries’ dialog box as discussed in the section of this chapter entitled “Viewing Disk Quota Entries.” Current quota entries for all users are listed. To refresh the listing, press ‘F5’ or select ‘Refresh’ from the ‘View’ menu.

2. If the user doesn't have an existing entry on the volume, you can create it by selecting 'New Quota Entry' from the 'Quota' menu. This opens the 'Select Users' dialog box.
3. In the 'Select Users' dialog box, type the name of the user you want to use in the 'Name' text box and then click 'Check Names'. If matches are found, select the account you want to use and then click 'OK'. If no matches are found, update the name you entered and try searching again. Repeat this step as necessary and then click 'OK' when you're finished.
4. After you've selected a user, the 'Add New Quota Entry' dialog box is displayed as shown in Figure 5-12. You can remove all quota restrictions for this user by selecting 'Do Not Limit Disk Usage'. Or you can set a specific limit and warning level by selecting 'Limit disk space to' and then entering the appropriate values in the fields provided. Click 'OK'.

Deleting Disk Quota Entries

When you've created disk quota entries on a volume and a user no longer needs to use the volume, you can delete the associated disk quota entry. When you delete a disk quota entry, all files owned by the user are collected and displayed in a dialog box so that you can permanently delete the files, take ownership of the files, or move the files to a folder on a different volume.

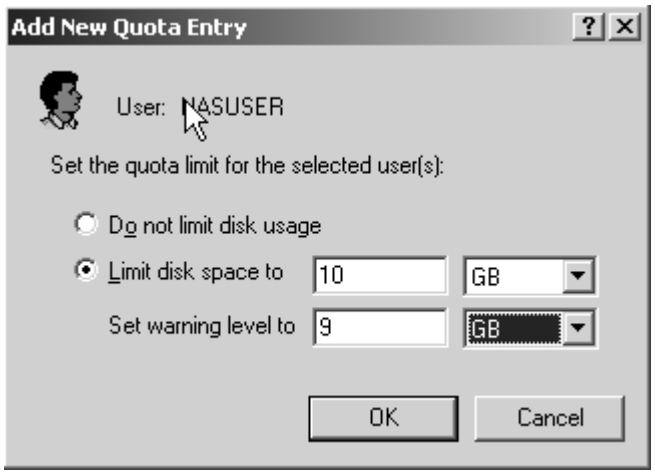

Figure 5-12 Add New Quota Entry

To delete a disk quota entry for a user and manage the user's remaining files on the volume, follow these steps:

1. Access 'Quota Entries' dialog box. Current quota entries for all users are listed. To refresh the listing, press 'F5' or select 'Refresh' from the 'View' menu.
2. Select the disk quota entry that you want to delete and then press the 'Delete' key or select 'Delete Quota Entry' from the 'Quota' menu. You can select multiple entries using 'Shift' and 'Ctrl' keys.
3. When prompt to confirm the action, click 'Yes'. This displays the 'Disk Quota' dialog box with a list of current files owned by the selected user or users.
4. Use the 'List files owned by' selection list to display files for a user whose quota entry you're deleting. You must now specify how the files for the user are to be handled. You can handle each file separately by selecting individual files and then choosing an appropriate option.
    - **Permanently Delete Files:** Select the files to delete and then press 'Delete'. When prompted to confirm the action, click 'Yes'.
    - **Take Ownership Of Files:** Select the files of which you want to take ownership and then click 'Take Ownership'.
    - **Move Files To:** Select the files that you want to move and then enter the path to a folder on a different volume in the filed provided. If you don't have the path that you want to use, click 'Browse' to display the 'Browse for Folder' dialog box, which you can use to find the folder. Once you find the folder, click 'Move'.
5. Click 'Close' when you're finished managing the files. If you've appropriately handle all user files, the disk quota entries will be deleted.

Exporting and Importing NTFS Disk Quota Setting

Rather than recreating custom disk quota entries on individual volumes, you can export the settings from a source volume and then import the settings on another volume. You must format both volumes using NTFS. The steps you follow to export and then import disk quota entries are the following:

1. Access the Quota Entries dialog box as discussed in the section of this chapter entitled "Viewing Disk Quota Entries." Current quota entries for all users are listed. To refresh the listing, press 'F5' or select 'Refresh' from the 'View' menu.
2. Select 'Export' from the 'Quota' menu. This displays the 'Export Quota Settings' dialog box. Use the 'Save In' drop-down list to choose the save location for the file containing the quota settings and then set a name for the file using the 'File Name' text box. Afterward, click 'Save'.
3. On the 'Quota' menu, select 'Close' to exit the 'Quota Entries' dialog box.
4. Right-click 'Computer Management' in the console tree. On the shortcut menu, select 'Connect to Another Computer'. In the 'Select Computer' dialog box, choose the computer containing the target volume. The target volumes the one that you want to use the exported settings.
5. As explained preciously, access the 'Properties' dialog box for the target volume. Then click 'Quota Entries' in the 'Quota' tab. This displays the 'Quota Entries' dialog box for the target volume.
6. Select 'Import' on the 'Quota' menu. Then, in the 'Import Quota Settings' dialog box, select the quota settings file that you saved previously. Click 'Open'.
7. If the volume had previous entries, you'll have the opportunity to replace existing entries or keep existing entries. When prompted about a conflict, click 'Yes' to replace an existing entry or click 'No' to keep the existing entry. You can apply the option to replace or keep existing entries to all entries on the volume by selecting the 'Do This For All Quota Entries' check box prior to clicking 'Yes' or 'No'.

Disabling NTFS Disk Quotas

You can disable quota for individual users or all users on a volume. When you disable quotas for a particular user, the user is no longer subject to the quota restrictions but disk quotas are still tracked for other users. When you disable quotas on a volume, quota tracking and management are completely removed. To disable quotas for a particular user, follow the technique outlined in the section of this chapter entitled "Creating Disk Quota Entries." To disable quota tracking and management on a volume, follow these steps:

1. Start Computer Management. If necessary, connect to a remote computer.
2. Display the 'Properties' dialog box for the volume on which you want to disable NTFS quotas.
3. In the 'Quota' tab, clear the 'Enable Quota Management' checkbox. Click 'OK'. When prompted to confirm, click 'OK' again.

Managing Disk Quota Templates

You use disk quota templates to define quota properties, including the limit, quota type, and notification thresholds. In 'File Server Management', you can view the currently defined disk quota templates by expanding the 'File Server Resource Manager' and 'Quota Management' nodes and then selecting 'Quota Templates'. Table 14-7 provided a summary of the default disk quota templates.

You can modify existing disk quota templates by completing the following steps:

1. In 'File Server Management', expand the 'File Server Resource Manager' and 'Quota Management' nodes and then select 'Quota Templates'.
2. Currently defined disk quota templates are listed by name, limit, and quota type.
3. To modify disk quota template properties, double-click the disk quota template name. This displays a related properties dialog box, as shown in Figure 5-13.

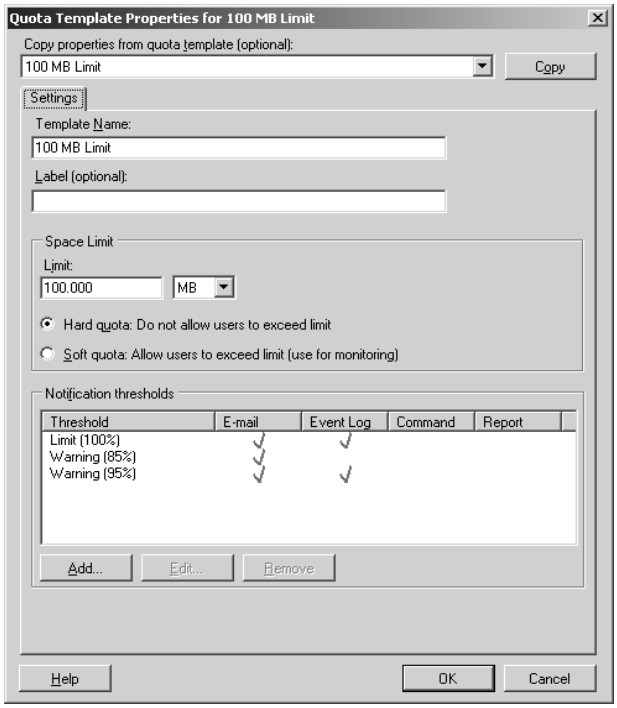

Figure 5-13 Properties dialog box, Security tab

4. In the 'Settings' tab, you can set the template name, limit, and quota type. Current notification thresholds are listed. To modify an existing threshold, select it and then click 'Edit'. To define a new threshold, click 'Add'.
5. When you're finished modifying the quota template, click 'OK' to save the changes.

You can create a new disk quota template by completing the following steps:

1. In 'File Server Management', expand the 'File Server Resource Manager' and 'Disk Management' nodes and then select 'Quota Templates'.
2. On the 'Action' menu or in the 'Actions' pane, select 'Create Quota Template'. This displays the 'Create Quota Template' dialog box.
3. In the 'Settings' tab, set the template name, limit, and quota type.
4. A limit threshold is already created. You should edit this threshold first and then create additional warning thresholds as necessary. Select 'Limit ND', then click 'Edit' to define the limit threshold.
5. Click 'Add' to add warning thresholds. In the 'Add Threshold' dialog box enter a percentage value under 'Generate Notifications When Usage Reaches (%)'. Warning thresholds are considered to be any percentage of the limit that is less than 100 percent. Limit thresholds occur when the limit reached is 100 percent.
6. In the 'E-mail Message' tab, you can configure notification as follows:
   - To notify an administrator when the disk quota is triggered, select the 'Send E-mail to the Following Administrators' check box and then type the e-mail address or addresses to use. Be sure to separate multiple e-mail addresses with a semicolon. Use the value [Admin E-mail] to specify the default administrator as configured previously under the global options.
   - To notify users, select the 'Send E-mail to the User Who Attempted to Save an Unauthorized File' check box.
   - To specify the contents of the notification message, use the 'Subject' and 'Message Body' text boxes. Table 13-6, "File Screen Variables and Their Meaning," in Chapter 13, "Managing Files and Folders," lists available variables and their meaning.
7. In the Event Log tab, you can configure event logging. Select the 'Send Warning to Event Log' check box to enable logging and then use the 'Log Entry' text box to specify the text of the log entry.
8. In the 'Report' tab, select the 'Generate Reports' check box to enable incident reporting and then select the types of reports to generate. Incident reports are stored under:

   "%SystemDreve%\StorageReports\Incident "

   by default, and they can also be sent to designated administrators. Use the value 'Admin E-mail' to specify the default administrator as configured previously under the global options.

9. Repeat Steps 5-8 to define additional notification thresholds. Click 'OK' when you're finished creating the template.

Creating Disk Quotas

You use disk quotas to designate file paths that have specific usage limits. In 'File Server Management', you can view current disk quotas by expanding the 'File Server Resource Manager' and 'Quota Management' nodes and then selecting 'Quotas'. Before you define disk quotas, you should specify screening file groups and disk quota templates that you will use, as discussed in Chapter 13 under "Managing the File Groups to Which Screens Are Applied" and in this chapter under "Managing Disk Quota Templates" respectively.

After you've defined the necessary file groups and disk quota templates, you can create a disk quota by completing the following steps:

1. In 'File Server Management', expand the 'File Server Resource Manager' and 'Quota Management' nodes and then select 'Quotas'.
2. Select 'Create Quota' on the 'Action' menu or in the 'Actions' pane.
3. In the 'Create Quota' dialog box, set the local computer path for the quota by clicking 'Browse' and then using the 'Browse For Folder' dialog box to select the desired path, such as `C:\Data`. Click 'OK'.
4. Use the 'Derive Properties from this Quota Template' drop-down list to choose the disk quota template that defines the quota properties you want to use Click 'Create'.

File Screening and Storage Report

File screening is a tool you can use in the effort to keep networks safe from malicious programs and to block unauthorized types of content. Using file screening, you can monitor and block the usage of certain types of files. You can configure file screening in one of two models:

1. **Active Screening:** Does not allow users to save unauthorized files.
2. **Passive Screening:** Allows users to save unauthorized files but monitors or warns about usage (or both).

You actively or passively screen files by defining a file screen. All file screens have a file screen path, which is a folder that defines the base file path to which the screening is applied. Screening applies to the designated folder and all subfolders of the designated folders. The particulars of how screening works and what is screened are derived from a source template that defined the file screen.

Table 5-1 File Screen Templates

File Screen Template name	Screening Type	File Group Action
Block Audio And video Files	Active	Block: Audio and Video files
Block E-Mail Files	Active	Block: E-Mail files
Block Executable Files	Active	Block: Executable files
Block Image Files	Active	Block: Image Files
Monitor Executable And System Files	Passive	Warn: Executable files, System files

File screen templates or custom properties define:

- Screening type: active or passive
- File groups to which screening is applied
- Notifications: e-mail, event log, or both

Each file group has a predefined set of files to which it applies. You can modify the included file types and create additional file groups as necessary using 'File Server Resource Manager' under Windows Storage Server Management Console.

You can configure exception paths as well to designate specially allowed save locations for blocked file type. You can use this feature to allow specific users or all users to save blocked file types to designated locations. As an example, you might want to deter illegal downloading of music and movies within the organization. To do this, you might want to prevent users from saving audio and video files and thereby prevent them from downloading music and movies. However, if your organization has a multimedia department that needs to be able to save audio and video files, you could configure an exception to allow files to be saved on a folder accessible only to members of this group.

You can generate storage reports as part of quota and file screening management. There are some standard storage reports available. The three general types of storage reports that can be generated based on one of the standard storage reports are:

1. **Incident Reports**: Generated automatically when a user tries to save an unauthorized file or when a user exceeds a quota.
2. **Scheduled reports**: Generated periodically based on a scheduled report task.
3. **On-Demand reports**: Generated manually upon request.

Managing File Screening and Storage Reporting

You manage file screening and storage reporting using the 'File Server Resource Manager' node in the Windows Storage Server Management Console.

File screening and storage reporting management can be divided into these key areas:

- **Global options**: Control global settings for file server resources, including e-mail notification, storage report default parameters, report locations, and file screen auditing
- **File groups**: Control the types of file to which screens are applied
- **File screens templates**: Control screening properties (screening type: active or passive, file groups to which screening is applied; notifications: e-mail, event log, or both)
- **File screens**: Control file paths that are screened
- **File screen exceptions**: Control file paths that are screening exceptions
- **Report generation**: Controls whether and how storage reports are generated

Managing Global File Resource Setting

You use global file resource options to configure e-mail notification, storage report default parameters, report locations, and file screen auditing. You should configure these global setting prior to configure quotas, file screens, and storage reporting.

Configuring E-mail Notifications

Notifications and storage reports are e-mailed through a Simple Mail Transfer Protocol (SMTP) server. For this process to work you must designate which organizational SMTP server to use, as well as default administrative recipients and the 'From' address to be used in mailing notifications and report. To configure these settings, follow these steps:

1. In the Windows Storage Server Management Console and select 'File Server Resource Manager node. On the 'Action' menu or in the 'Actions' pane, click 'Configure Options'. This displays the 'File Server Resource Manager Options' dialog box with the 'Email Notifications' tab selected by default, as shown in Figure 5-14.
2. In the SMTP Server Name or IP Address text box, type the fully qualified Domain name (FQDN) of the organization's mail server, such as smtp.digiliant.com, or the IP address of this SMTP server, such as 192.168.100.100.
3. In the 'Default Administrator Recipients' field, type the e-mail address of the default administrator for notification, such as john@digiliant.com. Typically, you'll want this to be a separate mailbox that this monitored by an administrator or a distribution list that goes to the specific administrators responsible for file

4. server resource management. You can also enter multiple e-mail address. Be sure to separate each e-mail address with a semicolon.
5. In the Default "From" e-mail address field, type the e-mail address you want the server to use. Remember, users as well as administrators may receive notifications.
6. To test the settings, click 'Send Test E-Mail'. The test e-mail should be delivered to the default administrator recipients almost immediately. If it isn't, check to ensure that the e-mail address used is valid and that the 'From' e-mail address is acceptable to the SMTP server as a valid sender. Click 'OK'.

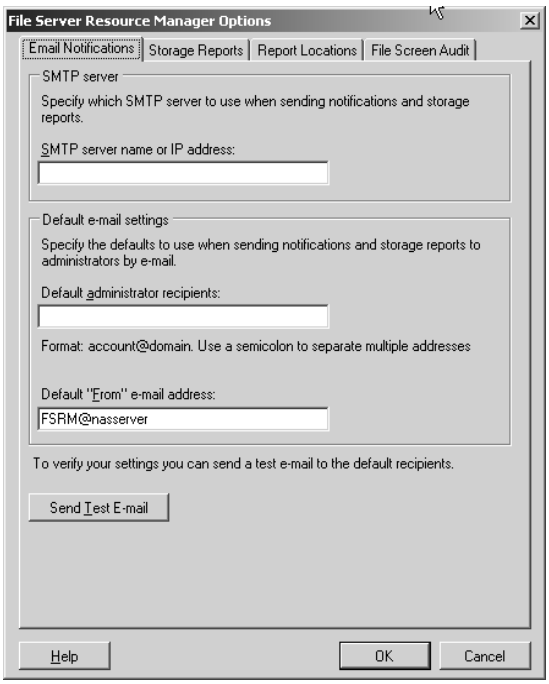

Figure 5-14 File Server Resource Manage Options

Reviewing Reports and Configuring Storage Report Parameters

Each storage report has a default configuration that you can review and modify using 'File Server Resource Manager Options'. Default parameter changes apply to all future incident reports and any existing report tasks that use the default configuration. You're able to override the default settings as necessary if you subsequently schedule a report task or generate a report on demand.

You can access the standard storage reports and change their default parameters by completing the following steps:

1. In Windows Storage Server Management Console, select the 'File Server' resource management node. On the 'Action' menu or in the 'Actions' pane, click 'Configure Options'. This displays the 'File Server Resource Manager Options' dialog box.
2. Click the 'Storage Reports' tab.
3. To review a report's current setting, select the report name in the 'Reports' list and then click 'Review Reports'.
4. To modify a report's default parameters, select the report name in the reports list and then click 'Edit Parameters'.
5. When you're finished, click 'Close' or 'OK' as appropriate.

Configuring Report Locations

By default, incident, scheduled, and on-demand reports are stored on the server on which notification is triggered in separate subfolders under `%SystemDrive%\StorageReports`. You can review or modify this configuration by completing the following steps:

1. In the Windows Storage Server Management Console, select 'File server Resource Manager' node. On the 'Action' menu or in the 'Actions' pane, click 'Configure Options'. This displays the 'File Server Resource Manager Options' dialog box.
2. Click the report 'Locations' tab.
3. The report folders currently in use are listed under 'Report Locations'. To specify a different local folder for a particular report type, type a new folder path or click 'Browse' to search for the folder path you want to use. Click 'OK'.

Configuring File Screen Auditing

All file screening activity can be recorded in an auditing database for later review by running a 'File Screen Auditing' report. This auditing data is tracked on a per server basis, so that the server on which the activity occurs is the one where the activity is audited. To enable or disable file screen auditing, follow these steps:

1. In the Windows Storage Server Management Console, select the 'File Server Resource Manager' node. On the 'Action' menu or in the 'Actions' pane, click 'Configure Options'. This displays the 'File Server Resource Manager Options' dialog box.
2. Click the 'File Screen Audit' tab.
3. To enable auditing, select the 'Record File Screening Activity in Auditing Database' check box.
4. To disable auditing, clear the 'Record File Screening Activity in Auditing Database' check box. Click 'Ok'.

Managing the File Groups to Which Screens Are Applied

You use file groups to designate sets of similar file types to which screening can be applied. In 'File Server Management', you can view the currently defined screening file groups by expanding the 'File Server resource Manager and File Screening Management' nodes and then selecting 'File Groups'.

You can modify existing file groups by completing the following steps:

1. In the Windows Storage Server Management Console → 'File Server Resource Manager' → 'File Screening Management' and then select 'File groups'.
2. Currently defined file groups are listed along with included and excluded files.
3. To modify file group properties, double-click the file group name. This displays a related properties dialog box similar to the one shown in Figure 5-15.
4. In the Files To Include text box, type the file extension of an additional file type to screen such as ***.pdf**, or the file name pattern, such as **Archive*.***. Click 'Add'. Repeat this step to specify other file types to screen.
5. In the 'Files To Exclude' text box, type the file extension of a file type to exclude from screening such as **.doc**, or the file name pattern, such as **Report*.***. Click 'Add'. Repeat this step to specify other file types to exclude from screening.
6. Click 'OK' to save the changes.

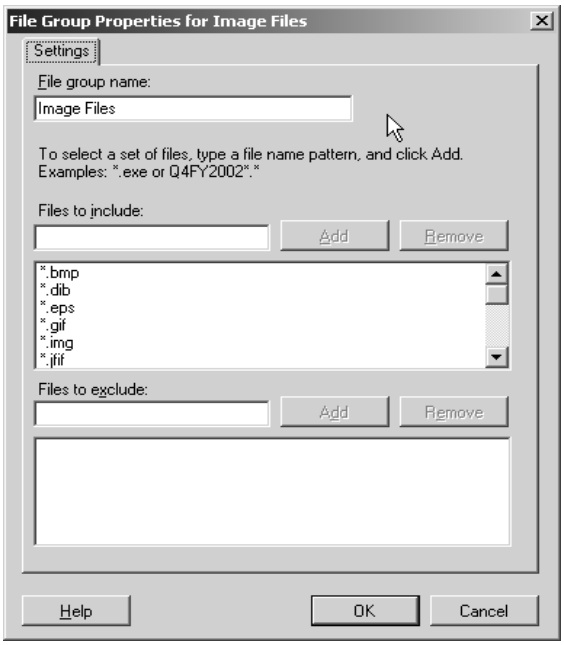

Figure 5-15 File Group Properties

You can specify additional file groups to screen by completing the following steps:

1. In Windows Storage Server Management Console, expand 'File Server Resource Manager' → 'File Screening Management', then select File groups.
2. On the 'Action' menu or in the 'Actions' pane, click 'Create File Group'. This displays the 'Create File Group' properties dialog box.
3. In the 'File Group' name text box, type the name of the file group you're creating.
4. In the 'Files to Include' text box, type the file extension of an additional file type to screen' such as .**pdf,** or the file name pattern, such as **Archive***.*. Click 'Add'. Repeat this step to specify other file types to screen.
5. In the 'Files To Exclude' text box, type the file extension of a file type to exclude from screening such as **.doc**, or the file name pattern, such as **Report***.*. Click 'Add'. Repeat this step to specify other file types to exclude from screening.
6. Click 'OK' to save the changes.

Managing File Screen templates

You use file screen templates to define screening properties, including the screening type, the file groups to which screen is applied, and notification. In Windows Storage Server Management Console, you can view the currently defined file screen templates by expanding 'File Server Resource Manager' → 'File Screening Management', then selecting 'File Screen Templates'.

You can modify existing file screen templates by completing the following steps:

1. In Windows Storage Server Management Console, expand 'File Server Resource Manager' → 'File Screening Management', and then select 'File Screen Templates'.
2. Currently defined file screen templates are listed by name, screening type, and file groups affected.
3. To modify file screen template properties, double-click the file screen template name. This displays a related properties dialog box (shown in Figure 5-16).
4. In the setting tab, you can set the template name, screen type, and file groups affected using the text boxes provided.
5. In the 'E-Mail Message' tab, you can configure notification:
    - To notify an administrator when the file screen is triggered, select 'Send E-Mail to the Following Administrators' check box and then type the e-mail address or addresses to use. Be sure to separate multiple e-mail addresses with a semicolon. Use the value to specify the default administrator as configured previously under the global options.
    - To notify users, select the 'Send E-Mail to the User Who Attempted to Save an Unauthorized File' check box.
    - To specify the contents of the notification message, use the 'Subject and Message Body' text boxes.

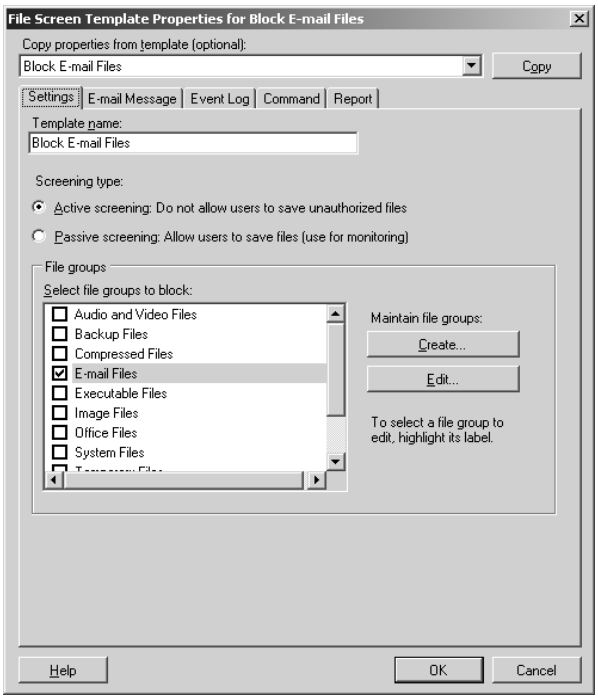

Figure 5-16 File Screen Template Properties

6. In the 'Event Log' tab, you can configure event logging. Select 'Send Warning to Event Log to enable logging and then use the 'Log Entry' field to specify the text of log entry.
7. In the 'Report' tab, select the 'Generate Report' check box to enable incident reporting and select the check boxes for the types of reports you want to generate. Incident reports are stored under `%SystemDrive%\Storage\Reports\Incident` by default and can also be sent to designated administrators. Use the value to specify the default administrator as configured previously under the global options.
8. Click 'OK' when you're finished modifying the template.

Creating File Screens

After you've defined the necessary file groups and file screen templates, you can create a file screen by completing the following steps:

1. In the Windows Storage Server Management Console, expand the 'File Server Resource Manager' → 'File Screening Management', and then select 'File Screens'.
2. Click 'Create File Screen' on the action menu or in the 'Actions' pane.

3. In the ‘Create File Screen’ dialog box, set the local computer path to screen by clicking browse and then using the ‘Browse for Folder’ dialog to select the path to screen, such as `C:\Data`.
4. Use the ‘Derive Properties’ selection list to choose the file screen template that defines the screening properties you want to use.
5. Click ‘Create’.

Defining File Screening Exceptions

You use exception paths to specifically designate folder locations where it’s permitted to save blocked file types. Based on the NTFS permissions on the excepted file path, you can use this feature to allow specific users to save blocked file types to designated locations or to allow all users to save blocked file types to designated locations.

You can create a file screen exception by completing the following steps:

1. In Windows Storage Server Management Console, expand the ‘File Server Resource Manager’ → ‘File Screening Management’, then select ‘File Screens’.
2. Click ‘Create File Screen Exception’ on the ‘Action’ menu or in the ‘Actions’ pane.
3. In the ‘Create File Screen Exception’ dialog box, set the local path to exclude from screening by clicking ‘Browse’ and the using ‘Browse For Folder’ dialog box to select the path to exclude from screening, such as `C:\Data\Media`.
4. Select the file groups to exclude from screening on the designated path. Click ‘OK’.

Scheduling and generating Storage Reports

Incident reports are generated automatically when triggered, as defined in the Reports tab properties of a file screen template. Scheduled and on-demand reports are configured separately.

You can schedule reports on per volume or folder basis by completing the following steps:

1. In Windows Storage Server management Console, expand the ‘File Server Resource Manager’ node and select ‘Storage Reports Management’.
2. On the ‘Action’ menu or in the ‘Actions’ pane, click ‘Schedule a New Report Task’. This displays the ‘Storage Reports Task Properties’ dialog box shown in Figure 5-17.

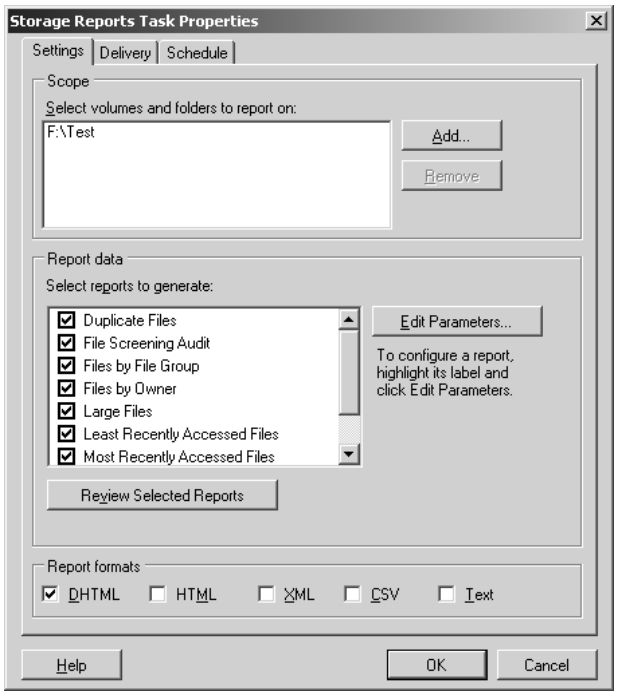

Figure 5-17 Storage Report Task Properties

3. In the ‘Setting’ tab under ‘Scope’, click ‘Add’. Use the ‘Browse for Folder’ dialog box to select the volume or folder on which you want to generate scheduled storage reports. Repeat to add other volumes or folders.
4. Under ‘Report Data’, select the types of reports to generate.
5. Under ‘Report Formats’, select the format for the report.
6. By default, Windows Storage Server’s scheduled storage reports are generated in the `%SystemDrive%\StorageReports\Scheduled` folder. If you’d also like to deliver reports by e-mail to administrators, click the ‘Delivery’ tab and then select the ‘Send Reports To The Following Administrators’ check box. Enter the e-mail address or addresses to which reports should be delivered.
7. In the schedule tab, click ‘Create Schedule’. In the ‘Schedule’ dialog box, click ‘New’ and then define the

run schedule for reporting.

8. Click 'OK' twice to schedule the report task.

You can generate an on-demand report by following these steps:

1. In Windows Storage Server management Console, expand 'File Server Resource Manager' and select 'Storage Reports Management'.
2. On the 'Action' menu or in the 'Actions' pane, click 'Generate Reports Now'. This displays the 'Storage Reports Task Properties' dialog box.
3. In the 'Setting' tab under 'Scope', click 'Add'. Use the 'Browse For folder' dialog box to select the volume or folder on which you want to generate on-demand storage reports. Repeat to add other volumes or folders.
4. Under 'Report Data', select the types of reports to generate.
5. Under 'Report Formats', select the format for the report.
6. By default, Windows Storage Server stores on-demand storage reports in the `%SystemDrive%\StorageReports\Interactive` folder. If you'd also like to deliver reports by e-mail to administrators, click the 'Delivery' tab and then select the 'Send Reports to The Following Administrators' check box. Enter the e-mail address or addresses to which reports should be delivered.
7. Click 'OK'. When prompted, specify whether to wait for the reports to be generated and then display them or to generate the reports in the background for later access. Click 'OK'.

Managing the Distributed File System

The Distributed File System (DFS) solution in Windows Storage Server 2003 R2 is a combination of two technologies, DFS Namespaces and DFS Replication; together these technologies provide a fault tolerant, virtual file system composed of file shares on multiple file servers that are kept in sync by a WAN friendly replication algorithm that is vastly more efficient and robust than the file replication services.

Each DFS namespace is a shared group of network shares residing under a DFS root, which serves as a container for the namespace and performs much the same function for the distributed file system that a root folder serves for a physical volume. DFS root contains links to the shares (local and remote) that form the distributed file system. Each link appears as a subfolder of the root share. For example, if documents are scattered across multiple servers in a Domain, DFS can make it appear as though the documents all resides on a single server. This eliminates the need for users to go to multiple locations on the network to find the information.

A server that hosts a DFS root is called a *DFS server*. You can create root targets on other servers to replicate a DFS namespace and provide redundancy in the event that the host server becomes unavailable. While DFS by default does not provide replication of a DFS root or any replicas associated with a given DFS link, you can configure DFS to replicate entire DFS roots or individual shared folders.

A user can access the DFS root using a UNC pathname in the form *server**root*, where '*server'* is the network name for the server hosting the DFS root and '*root'* is the root folder's name. For example, if you created a root named 'Shares' on a server name 'Nasserver', users would access the DFS namespace from their computers using the UNC pathname `\\Nasserver\Shares`.

Accessing the DFS namespace from other computers

In addition to the server-based DFS component of the Windows Storage Server 2003 family, there is a client-based DFS component. The DFS client caches a referral to a DFS root or a DFS link for a specific length of time, defined by the administrator.

The DFS client component runs on a number of different Windows platforms. In the case of older versions of Windows, the client software must be downloaded to run on that version of Windows. Newer versions of Windows have client software built-in.

Unfortunately, however, clients running non-Windows operating systems (such as Linux/UNIX) cannot access the DFS namespace; this is because DFS is dependent on a native Windows component to function.

Deploying DFS

A distributed file system can be implemented as a standalone root distributed file system or as a Domain root distributed file system. A standalone namespace stores all namespace information on the registry of the namespace server instead of in Active Directory. Standalone namespace can host more folders (up to 50,000 folders with target), but the only way to provide redundancy for a standalone namespace root is to use a server cluster. You cannot use multiple namespace to host a standalone namespace as you can with a Domain-based namespace. Domain-based namespace roots differ from standalone namespace roots in a couple of ways. First, you must host Domain-based namespace roots on a member server or Domain controller of an Active Directory Domain. Second, Domain-based namespace roots automatically publish the DFS topology in Active Directory. This arrangement provides fault tolerance and network performance optimization by directing clients to the nearest target, as discussed in the next section.

Choose a standalone namespace if the network does not use Active Directory, if the namespace contains more than 50,000 folders with targets, and you want to host the namespace on a server cluster. Otherwise, choose a Domain-based namespace to use multiple servers for redundancy and to take advantage of Active Directory for site aware client referrals. You can also combine the two, for example, you can create a Domain-based namespace that includes a standalone roots as folder.

DFS Management is the management tool of DFS namespace and replication. It can be found under 'Administrative Tools', or in the Windows Storage Server Management Console.

Creating or Opening a Namespace Root

To create a namespace or open an existing namespace, following these steps:

1. In the Windows Storage Server Management Console, expand → 'File Server Management' → 'DFS Management' then 'Namespaces'.
2. To open an existing namespace root, right-click 'Namespace' and choose 'Add Namespace to Display'.
3. To create a new existing namespace root, right-click 'Namespaces' and choose 'New Namespace'. The 'New Namespace Wizard' appears.
4. On the 'Namespace Server' page, type the name of the server that you want to host the namespace root then click 'Next'. If the DFS service is disabled, click 'Yes' in the 'Warning' dialog box to start the DFS service and set its startup mode to 'Automatic'.

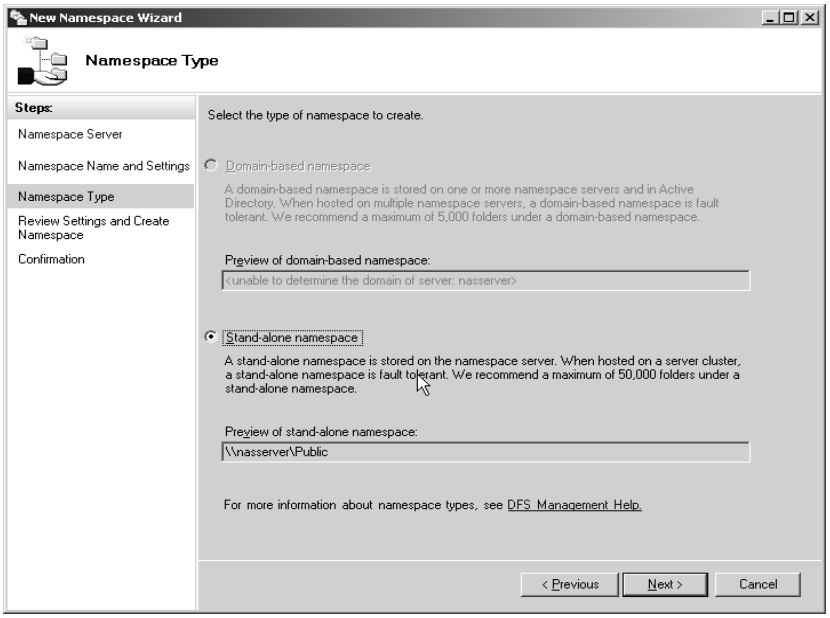

Figure 5-18 Namespace Type, New Namespace Wizard

5. On the 'Namespace Name and Settings' page, type the name to use for the namespace root. This name appears as the share name to users-for example, `\\Digiliant.local\shares`.

The 'New Namespace Wizard' creates the namespace root in the `%SystemDrive%:\DFSRoots\name` folder and gives all users read-only permissions. To change these setting, click 'Edit Setting'.

6. On the 'Namespace Type' page (shown in Figure 5-18), choose whether to create a Domain-based namespace or a standalone namespace:
    - Select 'Domain-Based Namespace' to store the namespace on multiple servers in Active Directory. An example of a Domain-based namespace is `\\Digiliant.local\shares`.
    - Select the 'Stand-Alone Namespace' option to create the namespace on a single server or server cluster. An example of a standalone namespace is `\\nasserver\shares`.
7. On the 'Review Settings and Create Namespace' page, click 'Create'. The 'New Namespace Wizard' creates the namespace root location. Review any errors and then click 'Close'.

Adding Namespace Servers

The namespace root is the most important part of the namespace. Without it, clients cannot access any DFS folders. Because of this, the first step in creating a more fault-tolerant namespace is to add namespace servers to the namespace root. If possible, add at least one namespace server on each site where users need access to the DFS namespace.

1. In the 'DFS Management' console, navigate to 'Namespace', right-click the domain-based namespace root you want to replicate, and then choose 'Add Namespace Server'.
2. In the 'Add Namespace Server' dialog box, type the path to the namespace server, and then click 'OK'. Windows creates the namespace root on the target server in the `%SystemDrive%:\DFSRoots\name` folder and gives all users read-only permissions. To change these settings, click 'Edit Settings'.
3. If the DFS service is disabled, click 'Yes' in the 'Warning' dialog box to start the DFS service and set its startup mode to 'Automatic'.

Adding DFS Folders

DFS folders allow users to navigate from the namespace root to other file shares on the network without leaving the DFS namespace structure. To create a DFS folder, follow these steps:

1. Right-click the namespace root to which you want to add a folder, and then choose 'New Folder'. This displays the 'New Folder' dialog box, shown in Figure 5-19.
2. Type a name for the folder in the Name box. To create a folder that contains other DFS folders, click 'OK' without adding any target folders. This creates a layer of structure to the namespaces.
3. To add target folders, click 'Add', and then type the shared folders' UNC or DNS path in the second text box, or click 'Browse' to browse to the shared folder.
4. Add any additional folder targets and then click 'OK'.
5. If you added multiple folder targets, click 'Yes' in the 'Replication' dialog box to create a replication group for the folder targets, or click 'No' to set up a replication group later (or not at all). If you click 'Yes', the 'Replicate Folder Wizard' appears with some settings already entered.

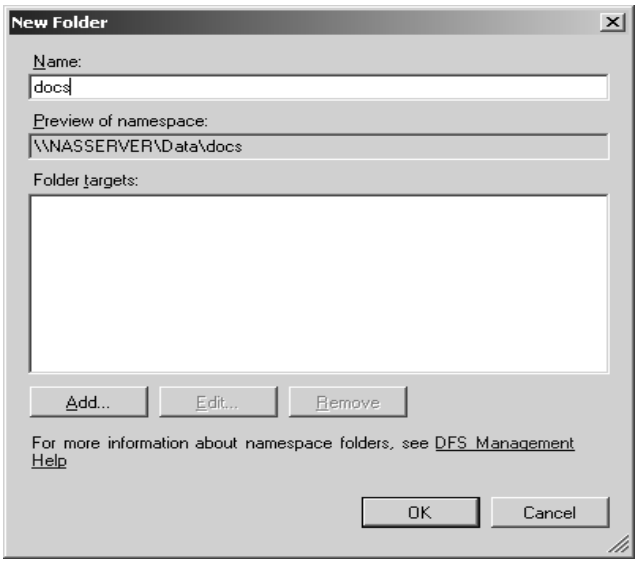

Figure 5-19 New Folder Under Namespace Root

DFS Replication

An easy-to-use, fault-tolerant, and high-performance file system is not worth much if the data you want to access is unavailable or out of date. To ensure that files are available to users even if a server goes down, create additional folder targets and use DFS Replication to keep the folder targets in sync. You can also use DFS Replication to synchronize folders that are not part of a DFS namespace-for example, to replicate data from a branch office to a server in the main office that you back up regularly and reliably.

To use DFS Replication, select the 'DFS Replication' node in the Windows Storage Server management Console, right-click replication, and choose 'Add Replication Group to Display' to open any existing replication groups, and then use the following sections.

There are two ways to create a replication group using DFS Management. You can add a DFS folder to a new replication group, or you can use the 'New Replication Group Wizard' to create a branch office replication group or a multipurpose replication group.

Replicating a DFS Folder

To create a replicated folder in a new replication group that replicates a DFS folder, use the following steps:

1. Right-click the appropriate folder under the 'Namespace' node of 'DFS Management', and choose 'Replicate Folder'. The 'Replicate Folder Wizard' appears.
2. On the 'Replication Group and Replicated Folder Name' page, confirm the name for the replication group and for the replicated folder.
3. On the 'Primary Folder Target' page, select the server that holds the data that you want to use as the seed for initial replication.

   If other members of the replication group have data in the replicated folders, Windows takes the following actions during the initial replication:

   - If an identical file already exists on the target server (any server other than the primary member), the primary member does not replicate the file.
   - If a file already exists on a target server but the file is not identical to the version on the primary member, Windows moves the file on the target server to the local conflict folder and then replicates the primary member's version of the file, even if this file is older than the version on the target server.
   - If a file exists on the target server that is not present on the primary member, Windows does not replicate it during the initial replication but does replicate it during subsequent replications to other members, including the primary member.

   After the initial replication, the primary member role goes away and replication is multiple-master-based. Do not delete, rename, or move files on the primary member or any member that has already replicated until the first replication is complete. *Deleting, renaming, or moving files before the first replication is complete can cause the files to reappear if they existed on a target that had not yet replicated*.
4. On the 'Topology Selection' page, select one of the following replication topologies:

   - **Hub-and-Spoke:** Spoke servers replicate with one or two central hub servers. Hub servers replicate with all other hub servers by using the 'Full Mesh' topology, as well as with designated spoke servers. Choose this topology in large network environments and environments with multiple brand office. This topology requires a minimum of three members.
   - **Full Mesh:** All servers replicate with all other servers. Choose this topology when there *are less than 10 servers in the replication group* and all links have low enough costs to allow each server to replicate with every other server instead of a central hub server.

- **No Topology:** This option does not specify a topology; in addition, it postpones replication until you specify a topology manually. To specify a replication topology after creating the replication group, right-click the replication group in the 'DFS Management' snap-in and then choose 'New Topology'.

5. On the 'Hub Members' page that appears if you chose the 'Hub-and-spoke topology, specify the hub servers.
6. On the 'Hub-and-Spoke Connections' page that appears if you chose the 'Hub-and-Spoke' topology, verify that the wizard lists the proper spoke servers. To change the required hub server with which a spoke member replicates preferentially, or the optional hub member with whom a spoke member replicates if the required hub member is unavailable, select the spoke server, click 'Edit', and then specify the required hub and the optional hub.
7. On the 'Replication Group Schedule and Bandwidth' page, choose when to replicate and the maximum amount of bandwidth you want DFS Replication to use.

   To create a custom schedule, choose 'Replicate During the Specified Days and Times' and then click 'Edit Schedule'. You can create a custom schedule that uses Coordinated Universal Time (UTC) or the local time of the receiving server.
8. On the 'Review Setting and Create Replication Group' page, review the settings, and then click 'Create'. Review any errors and then click 'Close'.

   Windows then replicates topology and replication settings to all Domain controllers. A replication group member polls its nearest Domain controller regularly. By default, replication group members perform a lightweight poll every 5 minutes for Subscription objects under the local computer container and a full poll every hour. It receives the setting after Windows updates the Domain controller. To change the replication polling interval, use the `dfsdiag` command.

Creating a branch Office replication Group

To create a replication group that replicates a single branch server with a single hub server, use the following steps;

1. In the 'DFS Management' snap-in, right-click 'Replication' and choose 'New Replication Group'. The 'New Replication Group Wizard' appears.
2. On the 'Replication Group Type' page, choose 'Replication Group for Data Collection'.
3. On the 'Name and Domain' page, type a name for the replication group that is unique on the Domain, specify in which Domain to host the replication group, and optionally type a description of the replication group.
4. On the 'Branch Server' page, type the name of the branch server that holds the data you want to replicate with the hub server.
5. On the 'Replicated Folders' page, click 'Add', and then use the 'Add Folder to Replicate' dialog box to specify the local folder on the branch server to replicate with the hub server, as shown in Figure 5-20. Click 'OK' when you are finished.

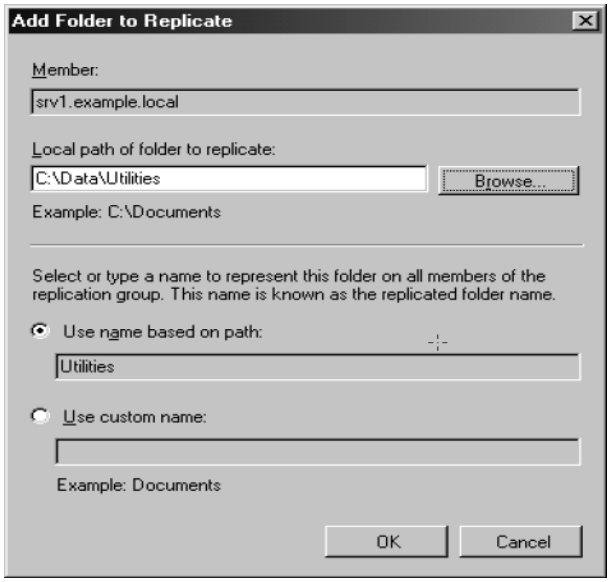

Figure 5-20 Add Folder To Replicate

6. On the Hub Server page that appears if you chose 'Replication Group for Data Collection' on the 'Replication Group Type' page, type the name of the hub server that serves as a replication target for the replicated folders.
7. On the 'Target Folder on Hub Server' page, specify the local folder on the hub server in which you want to place replicated data from the branch server. This folder is usually located in a folder or volume that you back up regularly.
8. On the 'Replication Group Schedule and Bandwidth' page, choose when to replicate and the maximum amount of bandwidth you want to allow DFS Replication to use.

   To create a custom schedule, choose 'Replicate During the Specified Days and Times', and then click 'Edit Schedule'. You can create a custom schedule that uses Coordinated Universal Time (UTC) or the local time of the receiving server.
9. On the 'Review Settings and Create Replication Group' page, review the settings, and then click 'Create'. Review any errors and then click 'Close'.

   Windows then replicates topology and replication settings to all Domain controllers. A replication group member polls its nearest Domain controller regularly; by default, replication group members perform a lightweight poll every 5 minutes for Subscription objects under the local computer container and a full poll every hour. It receives the settings after Windows updates the Domain controller. To change the replication polling interval, use the `dfsdiag` command.

Creating a Multipurpose replication Group

To create a replication group that replicates any number of servers with any number of other servers, use the following steps:

1. In the 'DFS Management' snap-in, right-click 'Replication' and choose 'New replication Group'. The 'New Replication Group Wizard' appears.
2. On the 'Replication Group Type' page, choose 'Multipurpose Replication Group'.
3. On the 'Name and Domain' page, type a name for the replication group that is unique on the Domain, specify in which Domain to host the replication group, and optionally type a description of the replication group.
4. On the 'Replication Group Members' page, add the servers on which you want to replicate content.
5. On the 'Topology Selection' page (shown in Figure 5-21), select one of the following replication topologies:
   - **Hub-and-Spoke:** Spoke servers replicate with one or two central hub servers. Hub servers replicate with all other hub servers by using the Full Mesh topology, as well as with designated spoke servers. Choose this topology to reduce network usage when there are more than 10 members of the replication group, or when members of the replication group are in a site connected via a WAN connection. This topology requires a minimum of three members.
   - **Full Mesh:** All servers replicate with all other servers. Choose this topology when there are less than 10 servers in the replication group and all links have low enough costs to allow each server to replicate with every other server. The Full Mesh topology minimizes the time it takes to propagate changes to all members of the replication group by increasing network usage.

- **No Topology:** This option does not specify a topology and postpones replication until you specify a replication topology manually. Use this setting if you want to create a custom topology from scratch instead of modifying an existing topology.

6. On the 'Hub Members' page that appears if you chose the Hub And Spoke topology, specify the hub servers.
7. On the 'Hub And Spoke Connections' page that appears if you chose the Hub And Spoke topology, verify that the wizard lists the proper spoke servers. To change the required hub server with which a spoke member replicates preferentially, or the optional hub member with which a spoke member replicates if the required hub member is unavailable, select the spoke server, click 'Edit', and then specify the required hub and the optional hub.

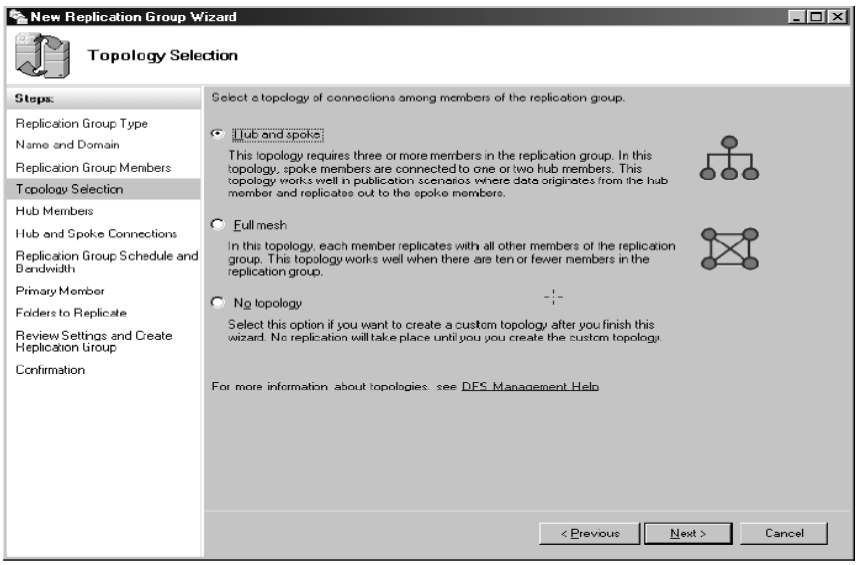

Figure 5-21 Properties dialog box, Security tab

8. On the 'Replication Group Schedule and Bandwidth' page, choose when to replicate and the maximum amount of bandwidth you want to allow DFS Replication to use.

   To create a custom schedule, choose 'Replicate During the Specified Days and Times' and then click 'Edit Schedule'. You can create a custom schedule that uses Coordinated Universal Time (UTC) or the local time of the receiving server.

9. On the 'Primary Member' page, select the server that holds the data that you want to use as the seed for the initial replication.
10. On the 'Folder to Replicate' page, click 'Add', and then use the 'Add Folder to Replicate' dialog box to specify the folder to replicate. Click 'OK' when you are finished.
11. On the 'Local Path of Folder On Other Members' page, select a replication member that you want to participate in replication of the specified folder, click 'Edit', and then use the 'Edit Local Path' dialog box to enable replication and specify the local folder on the target server in which to place replicated data from the hub server.

    Repeat this step for every replicated folder you specify in the 'Replicated Folders' page.

12. On the 'Review Settings and Create Replication Group' page, review the settings, and then click 'Create'. Review any errors and then click 'Close'.

    Windows then replicates topology and replication settings to all Domain controllers. A replication group member polls its nearest Domain controller regularly; by default, replication group members perform a lightweight poll every 5 minutes for Subscription objects under the local computer container and a full poll every hour. It receives the settings after Windows updates the Domain controller. To change the replication polling interval, use the `dfsdiag` command.

Managing Replication Groups

Select a replication group, and then use the 'Memberships', 'Connections', 'Replicated Folders' and 'Delegation' tabs of the DFS Management console to manage the replication group, as discussed in the following list.

- Use the following options on the Memberships tab to view and manage the member servers for each replicated folder:
    - To disable a member of the replication group, right-click the member and then choose 'Disable'. Disabled members are those who do not need to replicate a specific replicated folder. Do not disable members temporarily and then enable them; doing so causes roughly one kilobyte of replication traffic per file in the replicated folder and *will overwrite all changes on the disabled member*.
    - To delete a member of the replication group, right-click it and then choose 'Delete'.
    - To add a member server that participates in replication, right-click the replication group in the DFS Management console, choose 'New Member', and then use the 'New Member Wizard' to specify the local path of the replicated folders, connections, and schedule.
    - To change the size of the conflict or staging folders or to disable the retention of deleted files, right-click the member, choose 'Properties', click the 'Advanced' tab, and then use the 'Quota' boxes. The conflict folder stores the "losing" files that Windows deletes when it encounters two versions of the same file during replication as well as the most recently deleted files in the replicated folder, and the staging folder queues replication data.
    - To create a report showing the replication heath as well as RDC efficiency, right-click the replication group, choose 'Create Diagnostic Report', and then use the 'Diagnostic Report Wizard' to create the report.
    - To verify the replication topology, right-click the replication group and then choose 'Verify Topology'.
- Use the 'Connections' tab to view and manage all replication connections. To add a new replication connection between two members of a replication group, right-click the replication group and choose 'New Connection'. Then use the 'New Connection' dialog box to specify the sending member, the receiving member, the schedule, and whether or not to create a one-way replication connection or two-way connection.
- Use the following options on the 'Replicated Folders' tab to view and manage all replicated folders:
    - To add a new replicated folder to the replication group, right-click the replication group in the DFS Management console, choose 'New Replicated' Folder, and then use the 'New Replicated Folder Wizard' to specify the primary member and the local folders to replicate.
    - To omit certain file types or subfolders from replication, click the 'Replicated Folders' tab, right-click the replicated folder, choose 'Properties', and then use the 'File Filter' and 'Subfolder Filter' boxes on the 'General tab'.
    - To share a replicated folder on the network and optionally add the folder to a DFS namespace, right-click the replicated folder, choose 'Share and Publish in Namespace', and then use the 'Share or Publish Replicated Folder Wizard'.
- Use the 'Delegation' tab to view and manage administrative permissions. See the DFS Namespaces section of this chapter for information about the delegation tab.

Volume Shadow Copies

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots of volumes. A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the Shadow Copy mechanism is managed at the server, previous versions of files and folders are only available over the network from clients and as seen on a per folder or file level and an entire volume.

Shadow Copy only works on NTFS Volumes, *not on Fat or Fat32 volumes*.

The Shadow Copy feature uses data blocks. As changes are made to the file system, the Shadow Copy Service copies the original blocks to a special cache file to maintain a consistent view of the file at a particular point in time. Because the

snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space because blocks are not moved until an update to the disk occurs.

By using shadow copies, a Storage Server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted; previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file; a previous version of that file can be accessed.
- Compare several versions of a file while working; use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup archive or business recovery system, but they can help to simplify restore procedures. Because a snapshot only contains a portion of the original data blocks, shadow copies can not protect against data loss due to media failures. However the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow copy planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?
- How frequently will shadow copies be made?

Identifying the volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.

Shadow copies are designed for volumes that store user data such as home directories and ‘My Documents’ folders that are redirected by using ‘Group Policy’ or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the `\\servername\sharename` path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating disk space

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily. If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, no shadow copy is created.

Administrators should also consider user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

The minimum amount of storage space that can be specified is 350 megabytes (MB). The default storage size is 10 percent of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the *storage* volume instead of the *source* volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.

Converting basic storage disks to dynamic disks

When using a basic disk as a storage area for shadow copies and converting the disk into a dynamic disk, it is important to take the following precaution to avoid data loss:

- If the disk is a non-boot volume and is a different volume from where the original files reside, first dismount and take offline the volume containing the original files before converting the disk containing shadow copies to a dynamic disk.
- The volume containing the original files must be brought back online within 20 minutes; otherwise, the data stored in the existing shadow copies is lost.
- If the shadow copies are located on a boot volume, the disk can be converted to dynamic without losing shadow copies.

Identifying the storage area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on `H:`, another volume such as `S:` can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used Storage Server.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to 'No Limit' to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

By keeping the shadow copy on the same volume there is a potential gain in ease of setup and maintenance; however there may be a reduction in performance and reliability.

Determining creation frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a *maximum of 64 shadow copies per volume*, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the Storage Server creates shadow copies at 0700 and 1200, Monday through Friday. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs. To modify these schedules see 'Scheduling Shadow Copies' later in this chapter.

Managing shadow copies

To manage shadow copies on a volume:

1. Access 'Disk Management'.
2. Select the volume, then right-click on it.
3. Select 'Properties' tab.
4. Select 'Shadow Copies' tab under the 'Local Disk Properties' dialog.

Figure 5-22 shows Shadow Copies properties for a select volume. Through this dialog, you can performance the following functions.

- Enable or disable Shadow Copies on the selected volume.
- View, Create, delete and revert a Shadow Copy.
- Set schedules on Shadow Copies for the selected volume.

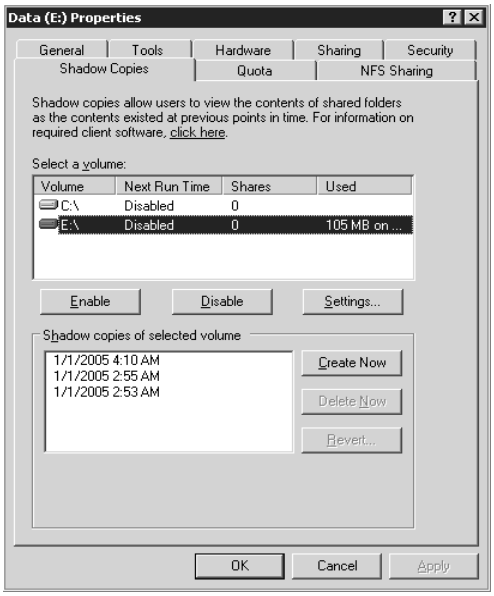

Figure 5-22 Shadow Copies page

Shadow Copies for Shared Folders

Shadow Copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this would include HTTP, FTP, AppleTalk, and NetWare Shares. To have SMB support a client side application denoted as 'Shadow Copies for Shared Folders' is required. The client side application is currently only available for Windows XP and Windows 2000 SP3+.

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.

SMB shadow copies

Windows users can independently access previous versions of files stored on SMB shares via the 'Shadow Copies for Shared Folders' client. After the 'Shadow Copies for Shared Folders' client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its 'Properties' window, clicking the 'Previous Versions' tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

'Shadow Copies for Shared Folders' preserves the permissions set in the Access Control List (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

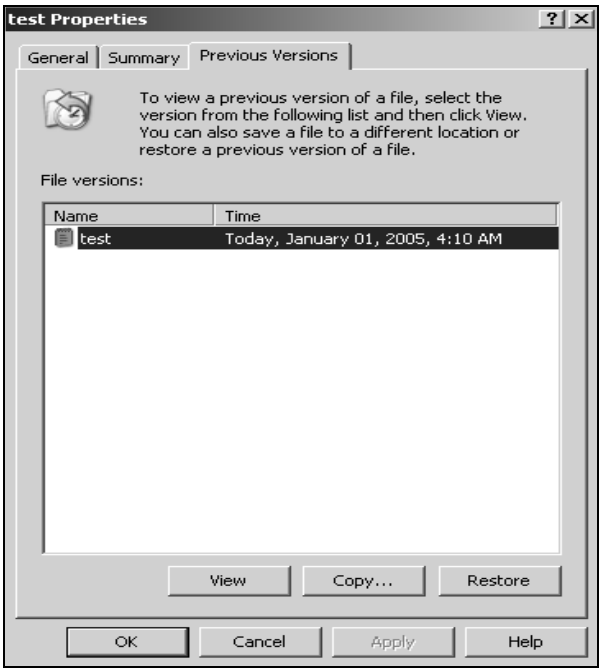

Figure 5-23 Client G U I

The 'Shadow Copies for Shared Folders' client pack installs a 'Previous Versions' tab in the 'Properties' window of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting 'View', 'Copy', or 'Restore', from the 'Previous Versions' tab. Both individual files and folders can be restored.

When users view a network folder hosted on the Storage Server for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history—a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS shadow copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format `.@GMT-YYYY.MM.DD-HH:MM:SS`. To prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named `'NFSShare'` with three shadow copies, taken on January 27, 28, and 29 of 2005 at 4 a.m.

```
NFSShare
.@GMT-2005.01.27-04:00:00
.@GMT-2005.01.28-04:00:00
.@GMT-2005.01.29-04:00:00
```

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of Files or Folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation.
- Accidental file replacement, which may occur if a user selects Save instead of Save As.
- File corruption.

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

Recovering a Deleted File or Folder

To recover a deleted file or folder within a folder:

1. Navigate to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file is selected.
3. Right-click and select ‘Properties’ from the bottom of the menu, and then click the ‘Previous Versions’ tab.
4. Select the version of the folder that contains the file before it was deleted, and then click ‘View’.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Click ‘Restore’ to restore the file or folder to its original location. Click ‘Copy’ to allow the placement of the file or folder to a new location.

Recovering an Overwritten or Corrupted File

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file:

1. Right-click the overwritten or corrupted file and then click ‘Properties’.
2. Click ‘Previous Versions’.
3. To view the old version, click ‘View’. Click ‘Copy’ to replace the current version with the older version, click ‘Restore’.

Recovering a folder

To recover folders use the following procedure:

1. Position the cursor so that it is over a blank space in the folder to be recovered. If the cursor hovers over a file, that file is selected.
2. Right-click, select ‘Properties’ from the bottom of the menu, and then click the ‘Previous Versions’ tab.
3. Click either ‘Copy’ or ‘Restore’.

Clicking ‘Restore’ enables the user to recover everything in that folder as well as all subfolders. Clicking ‘Restore’ does not delete any files.

Backup and Shadow Copies

Shadow copies are only available on the network via the client application and only at a file or folder level as opposed to the entire volume. Hence the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, shadow copies are available for back up in two situations. If the backup software in question supports the use of shadow copies and can communicate with the underlying block device, it is supported and the previous version of the file system will be listed in the backup application as a complete file system snapshot. Lastly, if the built-in backup application `NTbackup` is utilized, the backup software forces a snapshot and then uses the snapshot as the means for back up. The user is unaware of this activity and it is not self evident although it does address the issue of open files.

6 Data Backup and Recovery

Introduction

The loss of data can bring an enterprise down – be it for a few hours, days, or even weeks. Backup is, without a doubt, one of the most significant aspects of a business continuity plan. Essentially, the term 'backup' is used to describe a complete duplicate copy of all key information, both physical (paper) and computer records. An organization must create a decent set of backup procedures to ensure that data is protected.

Planning for Backup and Recovery

When planning for backup and recovery, you must ask yourself a number of questions to help you make a decision on how, when and what data will be backed up. These include the following:

When is the most convenient time to schedule backups?

Backing up data should ideally be done at off-peak hours when system usage is low. However, due to the type of data being backed up this is not always possible. Considerations must be made as to when key system data is to be backed up.

Will you store backups off-site?

It is highly recommended that backup media (containing data) is stored off-site in case of a natural disaster, fire, leak, etc. It would also be wise to keep a copy of the software required to install and restore operating systems, database servers, backup recovery, and so on.

How important is the data your systems contain?

Classifying the importance of your data will help you decide if the data needs to be backed up, how it should be backed up and when it should be backed up. Critical data (such as financial data, databases, etc) will take priority and should have a long-term redundant backup set, whereas data of less importance should be backed up daily and easily be recoverable.

How fast will restoring data from backup need to be?

Bringing a critical system back online would normally be done as soon as possible. Your backup plan depends a lot on the time it takes to recover a system. Data should be classified by priority and restored in sequence.

How regularly does data change?

Data that changes daily should be backed up daily; the rate at which your data changes will reflect your decision on how often the data should be backed up.

What type of information does the data residing on the systems contain?

It is important to consult others who may use data stored on the system when planning or reviewing a backup regiment. New projects and changes in business may create new priorities from time to time…ensure your information regarding the file system is current. This will help you determine when and how certain data should be backed up as well.

Do you have what's necessary to perform backups?

Make sure that you have the right hardware and enough media needed to perform a backup. Choosing backup media is an important factor in the backup and recovery process. Backup tapes are a common form of media since they can store large amounts of data and are cheap. They are however much slower than alternative options.

Basic Types of Backup

Each file or folder on your system consists of what is called an 'Archive' attribute. If this attribute is enabled then the file or folder may require backing up at the next backup time.

To view the archive attribute in Windows Server 2003 R2, right click a file or folder and select 'Properties'. Press the 'Advanced' button to bring up the 'Advanced Attributes' dialog box. This will allow you to select whether you want the object to be ready for archiving.

There are five backup types which you can use, depending on the importance of the data you are backing up and how convenient you want the restoration process to be.

Daily - Backs up files that have changed since the last daily backup. If a file is modified on the same day as the backup, it will be backed up. The archive attribute of the files is not changed.

Incremental - Backs up files that have changed since the most recent full (normal) or incremental backup. If the archive attribute is present then it means the file has been modified – only files with this attribute are backed up. Once the file has been backed up, the archive attribute is cleared and only set once the data has been modified again.

Full (Normal) - Backs up all files that have been selected, despite the archive attribute setting. Once the file has been backed up, the archive attribute is cleared until the file is modified. When the archive attribute is set again, it indicates that the file needs to be backed up.

Differential - Backs up files that have changed since the last Full backup. If the archive attribute is present, it means that the data has been modified and files having this attribute set will be backed up. However, in this case the attribute is not cleared so as to allow other types of backups to take place on this data at a later stage.

Copy - Backs up all files that have been selected, despite the archive attribute setting. The archive attribute is not changed, so that other types of backup can be performed on the same data.

Keep in mind, a backup procedure is never considered complete until it has been fully tested. What good is it if you backup data but can't restore it?

Types of Backup Media

There are numerous tools and media available for backing up data. When making your selection, there are five fundamental factors that you should base your decision on: *Speed, Reliability, Capacity, Extensibility and Cost*.

The most common types of backup media available on the market today include: Tape drives, Disk drives, Removable Disks, DAT (Digital Audio Tape) drives and Autoloader tape systems.

Backup Tips

1. Draw up a simple (easy to understand) plan of who will do what in the case of an emergency.
2. Be organized! Keep a record of what was backed up, when it was backed up and which backup media contains what data. You can also make a calendar of which type of backup is due on a certain date.
3. Utilize the 'Volume Shadow Copy' service in Windows Storage Server 2003 R2. This feature allows you to create point-in-time copies of data so that they can be restored and reverted to at any given time. For instance, if I created a Word document yesterday and decide I want to revert to it today, I can do so using VSS.
4. Select the option to verify backup, the process will take a little longer but it's definitely worth the wait.
5. Create a reference point where you know everything is working properly. It will be quicker to restore the changes from tape.
6. Select the option to restrict restoring data to owner or administrator and also set the Domain Group Policy to restrict the Restore privilege to system administrators only. This will help to reduce the risk of someone being able to restore data should the media be stolen.
7. Create a step-by-step guideline (a flowchart for example) clearly outlining the sequence for the retrieval and restoration of data depending on the state of the system.

About the Windows 2003 Backup Utility

There are many third-party backup software packages out there – HP, VERITAS, and CA being some of the big name players on the market. Depending on the size and budget of your enterprise you may wish to choose any of these. If however, you are after a simple solution to backup individual systems and data on shared folders then why not use the

backup utility that comes free with the Windows Server 2003 R2 operating system? Alternatively, why not use 'Backup' in conjunction with another software backup to provide a complete backup and recovery solution?

The backup utility in Windows Server 2003 R2 will allow you to, among other things, archive files and folders on the current system or any remote shared folders to a hard disk and then to restore these files to any accessible disk sometime in the future; to create a copy of the system state, the system/boot partition (and any files needed to start up your system in the event of a system failure), to schedule automated backups, create a log file of what was backed up and when, and also create an ASR (Automated System Recovery) disk that will save system files and configuration settings. Backup provides extensions for working with special types of data including system state data, Exchange Server data, removable storage data and remote storage data.

Backing up Data

To open the backup utility, go to the Start menu, navigate to 'Programs' → 'Accessories' → 'System Tools' and click 'Backup'. This will start the 'Backup and Restore Wizard' or go straight to the 'Backup and Restore Utility' (depending on your previously chosen settings). As an administrator, you'll want to use 'Advanced Mode', which gives you more options. If you start in 'Wizard' mode, click 'Advanced Mode' to switch. To perform backup and recovery operations, you must have certain permissions and user rights. Members of the Administrators and the Backup Operators groups have full authority to back up and restore any type of file. Other users can only back up files they own or those for which they have Read, Read & Execute, Modify, or Full Control permissions.

Setting backup options

Select 'Tools' → 'Options' to open the 'Options' dialog box and select your backup preferences. The 'General' tab will allow you to choose whether you want to verify backup after the backup process has completed – this is a good idea as it will compare the data on the source with that of the destination to make sure an identical copy has been made. The 'Restore' tab gives you the option to replace files, not to replace files or to replace files on disk if they are older than those on the backup media. The 'Backup Type' tab allows you to select which default backup type you want to use – choose from 'Normal', 'Copy', 'Differential', 'Incremental' and 'Daily' (as discussed in Part 1 of this series). From the 'Backup Log' tab you can set the level of logging you want for a backup – choose from 'Detailed' to log all information, 'Summary' to log the most important information and 'None' to log nothing at all. Finally, the Exclude Files tab will let you set which files to exclude from being backed up.

The image in Figure 6-1 shows the 'General' tab of 'Backup Options'.

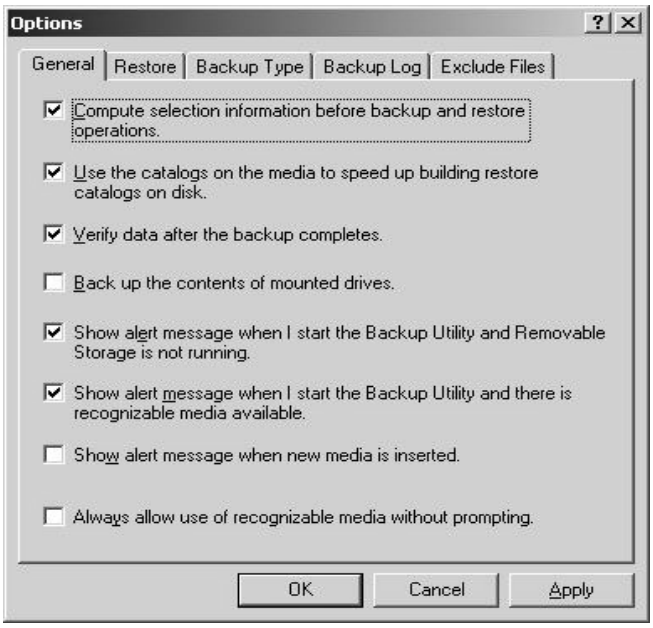

Figure 6-1 Backup Options

Performing an interactive backup

From the 'Backup' tab you can choose which drive, file or folder you want to back up and to which destination. In the left hand pane, click the checkboxes for which drive, file or folder you want to be backed up. The details for the selected folder appear in the right hand pane, as seen in the image below:

In this example the contents of 'My Documents' on the local machine and a share located on another computer in my Workgroup are chosen to be backed up. If you want to back up system state data, select 'System State' below the 'My Computer' node. If you're backing up an Exchange Server, select the 'Microsoft Exchange' icon below the 'My

Computer' node. When you do this, you'll be prompted to type the Universal Naming Convention (UNC) name of the Exchange server, such as `\\CorpMail`.

Use the 'Backup Destination' selection list to choose the media type (File or a storage device such as Tape) for backup.

In the 'Backup Media or File Name' text box, select the backup file or media. For file, type a path and file name for the backup file or click Browse to find a file. Backup files usually have an extension of **.bkf**, but you can change it to whatever extension you like when assigning the file name.

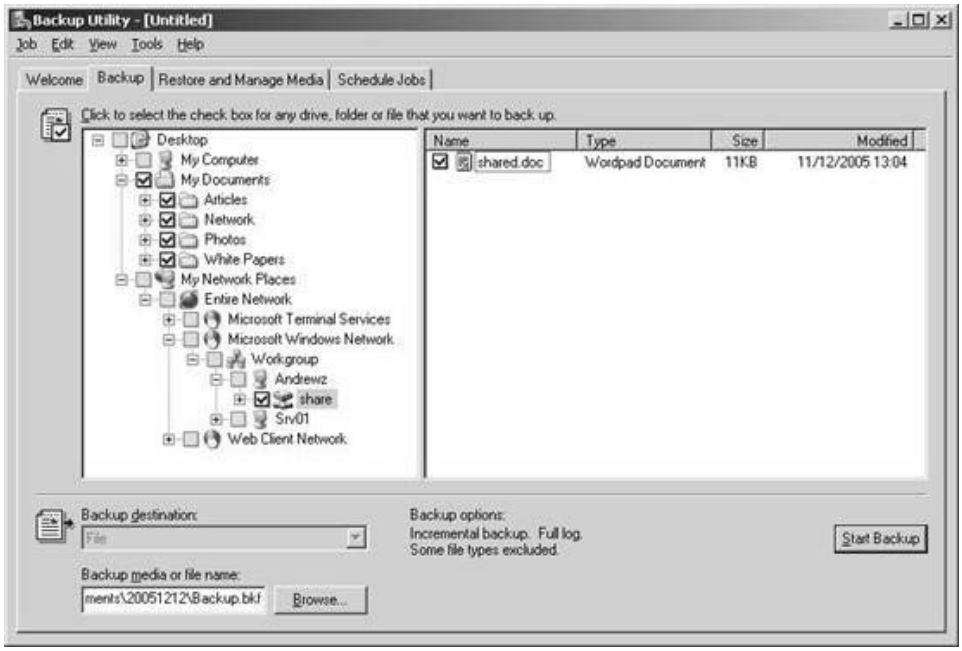

Figure 6-2 Backup

Scheduling a Backup job

To save you from having to manually backup files, you can schedule a backup job and let the backup utility do everything for you automatically. At a certain point in time the backup utility will start, and initialize the backup job. This is great if you want to perform routine backups – like a weekly full backup of all drives, for example.

Once you have selected which files to backup and pressed 'Start Backup', click the 'Schedule' button in the 'Backup Job Information' screen. After you have saved the backup job you will be asked to enter the username and password of the account you want the job to run under. From the 'Scheduled Job Options' dialog, type a job name and click the 'Properties' button to set the date, time, and frequency of this job.

NOTE:

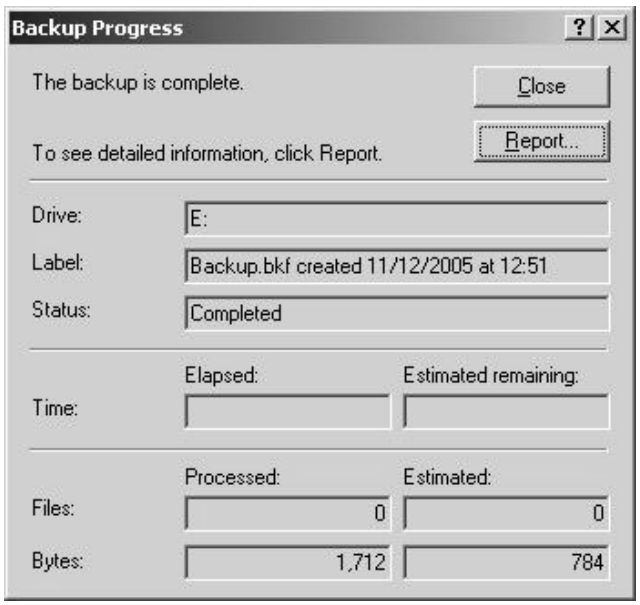

Figure 6-3 Backup Progress

You will have to be a member of the local Administrators or Backup Operators group to perform this task.

Viewing a Backup log

Once the backup is complete you can view the report by clicking on the Report button. This will show you details like what type of backup was performed and if the backup was a success.

To view previous backup reports, go to 'Tools' → 'Report' to open the 'Backup Reports' dialog window. Select a report and press 'View' to open the report in your default text editor or 'Print' to print to a file or print device.

The image below shows a simple report for an interactive backup job:

Restoring Data

Restoring data is a simple procedure using the Backup utility.

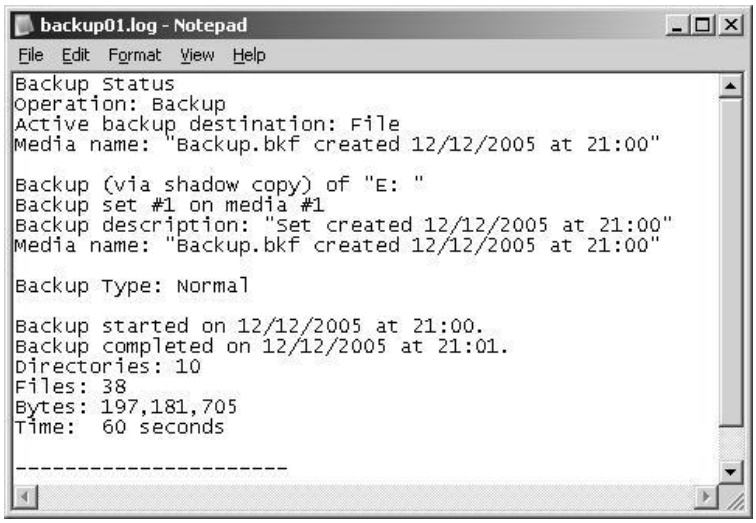

Figure 6-4 Backup Report

Go to the 'Restore and Manage Media' tab and select which media you wish to restore from – this will be displayed in the left hand pane. Once you have selected the backup media, the details will be displayed in the right hand pane, as seen below:

To restore system state data, select the check box for System State as well as other data you want to restore. If you're restoring to the original location, the system state data you're restoring will replace the current system state. If you restore to an alternate location, only the registry, `sysvol` and system boot file are restored. You can restore system state only on a local system.

If you're restoring Microsoft Exchange, select the 'Microsoft Exchange' data to restore. Before the restore starts, you'll see the 'Restoring Microsoft Exchange' dialog box. If you're restoring the 'Information Store', type the UNC name of the Exchange server such as `\\CorpMail`. If you're restoring to a different server, select 'Erase All Existing Data'. This destroys all existing data and creates a new 'Information Store'.

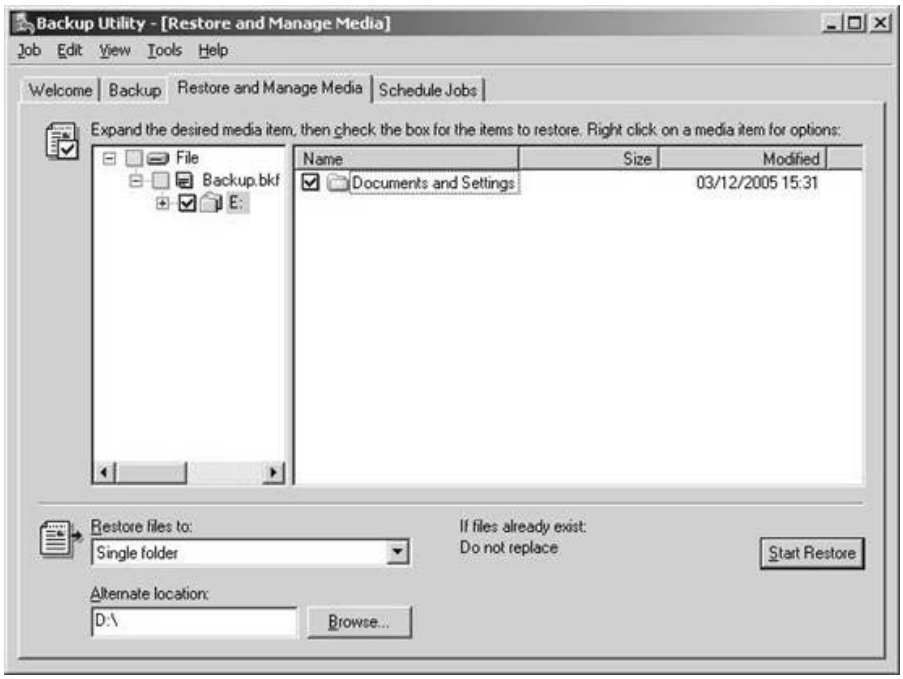

Figure 6-5 Restoring a Backup

Your next step is to choose where you want the data to be restored to. In the 'Restore Files' to drop-down list, choose 'Original location' for the files to be restored to the location which they were backed up from, 'Alternate Location' to restore the files to a different location (perhaps a different drive or folder) and keep the original folder structure, or 'Single Folder' if you want to restore the files to a folder and not keep the original folder structure (all files will be placed in the folder you choose).

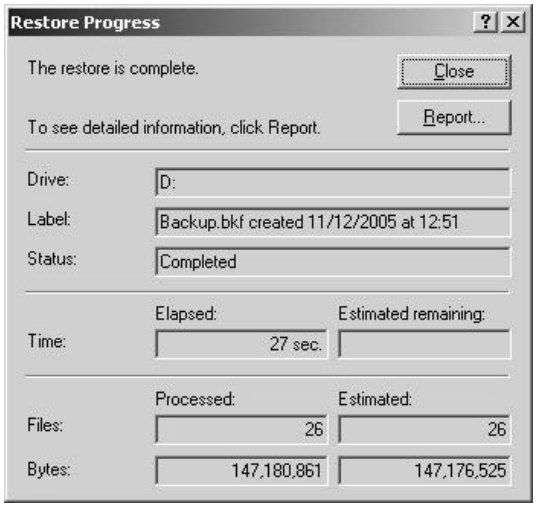

Figure 6-6

The final step would be to choose 'Start Restore'. After asking you to confirm whether you want to restore the data and giving you the option to change 'Advanced Settings', a dialog box will open and start the restoration process. You will be notified when it is complete by means of a "Restore is Complete" message, as shown in Figure 6-6.

Backing Up and Restoring Data on Remote Systems

You can use the Backup utility to backup data on remote systems. To do this, you must create network drives for the remote file systems before you begin the backup procedure. When backing up data on the network drives, be sure to select the 'General' option 'Back Up the Contents of Mounted Drives'. If you don't, only folder references are backed up and not the actual data.

You can also use Backup to restore data on remote systems. When you do this, you can select restore location in 'My Network Places'. If you're restoring to volume containing junction points, such as a volume containing the `sysvol`, be sure to select the advanced restore option 'Restore Junction Points', and 'Restore File and Folder Data Under Junction Points to the Original Location'. See Microsoft Knowledge Base article KB205524 for more information on junction points.

Disaster Recovery

Disaster Recovery is the ability to recover system operations after a disaster has occurred. One of the most significant aspects of disaster recovery is planning and designing a comprehensive backup plan that includes procedures, maintenance and backup storage methods.

During a recovery, it may not always be necessary to bring all systems and services back online at once. The most critical systems should be a top priority, with other systems (such as the public website) being of lesser priority. This will allow you to bring the core of the system back up and running again before turning your attention other services.

Usually companies in different locations have bi-lateral agreements that allow them to use each others site in case of a disaster. There are three types of sites available that will allow you to restore system operations in the event that a natural disaster destroys your main site. These are:

- Hot Site – A site that provides the ability to get back online and resume operations within a few hours of failure by having the equipment needed on stand by.
- Warm Site – A site that provides some capabilities in the event of a recovery. Everything will be in place for the organization to install and configure the systems to get operations up and running again.
- Cold Site – A physical location that has all the resources necessary to allow an organization to use it if the original site has been deemed impossible to use. The systems will have to be installed, setup and configured. A decision on moving to this alternate site is normally made within a few hours of the disaster.

When your system does not start properly, or if it does not start at all, then you can use the Windows Server 2003 R2 Recovery Console to help recover your system software and perform administrative tasks such as format drives, read and write data on a local drive, and enable/disable system services. Three of the most commonly used commands in the recovery console are:

- `fixmbr` - Repairs the master boot record (MBR) of a specific disk.
- `fixboot` - Writes a new boot sector onto a specific partition.
- `chkdsk` - Checks a local drive and displays a status report, and allows you to fix common disk errors.

To view all recovery console commands, enter the recovery console and type help at the command prompt.

Creating System Recovery Data

Automated System Recovery (ASR) data can often help you recover a system that won't boot. The recovery data includes essential system files, partition boot sector information, and the startup environment for a particular system. Normally, you'll want to update the recovery data when you install service packs, manipulate the boot drive, or modify the startup environment. Recovery data does not include user data files.

You can create ASR data using the Backup utility. ASR data is stored in two different forms: primary and secondary data. The primary data is store on a media you choose, such as a tape backup or disk drive. The secondary data is stored on a floppy disk and contains the files needed to boot the operating system and access the primary data.

You can create a system recovery data snapshot by completing the following steps:

1. Insert a blank 3.5", 1.44MB disk into the floppy drive.
2. Start the Backup utility. If 'Wizard' mode is enabled, click 'Next', select 'Prepare an Automated System Recovery Backup', and then click 'Next' again. If 'Wizard' mode is disabled, click 'Automated System Recovery Wizard' in the 'Welcome' tab and then click 'Next'.
3. On the 'Backup Destination' page, specify where the primary data should be stored. Select the backup media type and the location of the backup media.
4. Click 'Next' and then 'Finish'.

Using the Recovery Data to Restore a System

When you can't start or recover a system in safe mode, your next step is to try to recover the system using the last system recovery data snapshot you made. If the boot sector or essential system files are damaged, you might be able to use the recovery data to restore the system. If the startup environment is causing problems on a dual or multi-boot system, you may be able to recover the system as well. You can't recover a damaged registry, however; to do that, you must use the 'Recovery Console'.

You can repair a system using the recovery data by completing the following steps:

1. Insert the first boot disk into the appropriate drive, and then restart the computer.
2. When the Setup program begins, follow the prompts, and then choose the 'Repair or Recover' option by pressing 'R'.
3. Choose emergency repair by pressing 'R' and then do one of the following:
    - Press 'M' for 'Manual Repair' - Select this option to choose whether you want to repair system files, the partition boot sector, or the startup environment. Only advanced users or administrators should use this option.
    - Press 'F' for 'Fast Repair' – Select this option to have Windows Storage Server 2003 R2 attempt to repair problems related to system files, the partition boot section, and the startup environment.
4. Insert the 'System Recovery' floppy disk when prompted. Damaged or missing files are replaced with files from the Windows Server 2003 R2 CD or from the `%SystemRoot%\Repair` folder on the system partition. These replacement files won't reflect any configuration changes made after setup, and you might need to reinstall service packs and other updates.
5. If the repair is successful, the system is restarted and should boot normally. If you still have problems, you might need to use the 'Recovery Console'.

Media Pools

A 'media pool' is a collection of tapes or disks to which the same management properties apply. All media in a 'Removable Storage System' belong to a media pool and each media pool holds only one type of media. Data management programs use media pools to gain access to specific tapes or disks within a library

Using media pools, you can define properties that apply to a set of media. This is useful because 'Removable Storage' allows multiple programs to share the same media within a single library. A library can include media from different media pools, each with different properties. A single media pool can span multiple libraries. You can also create hierarchies of media pools, or media pools that contain other media pools. For example, you can create a media pool for each specific media type required by a program, and then create another media pool that contains this collection of media pools. Media pools can contain either media or other media pools, but not both.

A Removable Storage system provides two classes of media pools: system and application.

System media pools include free media pools, unrecognized media pools, and import media pools. Removable Storage creates one free, one unrecognized, and one import media pool for each media type in your Removable Storage system. The system media pools are used to hold media that are not currently being used by an application.

Application Media Pools are created by data management programs such as 'Backup' and 'Remote Storage'.

The different media pools are defined as follows:

- **Unrecognized:** Media pools containing blank (new) media and media that Removable Storage does not recognize. You should immediately move a new tape or disk from an unrecognized media pool to a free media pool so that the tape or disk can be used by applications, or remove it from the library. Unrecognized media are automatically deleted from the Removable Storage database when they are ejected from a library.
- **Import:** Media pools containing media that Removable Storage recognizes in the database but that have not been used before in a particular Removable Storage system. For example, media in an import media pool could be media from one office location that are introduced into a Removable Storage system at another office location. You can move media from import media pools to free media pools or application media pools so applications can use them.
- **Free:** Media pools containing media that are not currently in use by applications and do not contain useful data. Media in free media pools are available for use by applications. Application media pools can be configured to automatically draw media from free media pools when there are not sufficient media available in a particular application media pool. If this configuration is not implemented, you must manually move media from a free media pool when needed.
- **Application Media Pools**: created by data management applications (and by you), determine what media can be accessed by which applications.

  Media in an application media pool are controlled by that application or by an administrator. An application can use more than one media pool, and more than one application can share a single media pool. For example, Backup might use one media pool for full-backup and another media pool for incremental backup, each containing a different media type.

  There can be any number of application media pools in a Removable Storage system. Media that are currently reserved for use by an application, called *allocated media*, cannot be moved between media pools. Allocation controls how the media are used by applications.

If media have information that you don't need any more, you can initialize the media and prepare them for use in the 'Free' media pool. When you do this, you destroy the information on the media and move the media to the 'Free' media pool.

1. In' Computer Management', open 'Removable Storage' and then double-click 'Libraries'.
2. Expand the library and the library's media folder by double-clicking them.
3. Right-click the media you want to prepare and then click 'Prepare'.
4. Confirm the action by clicking 'Yes'.

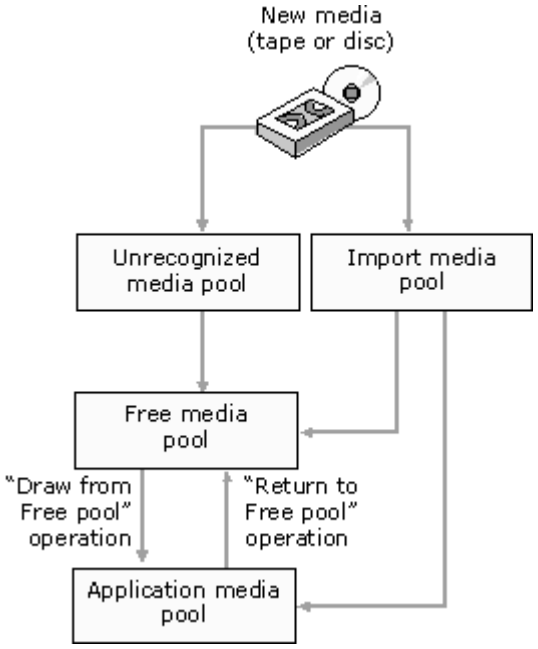

Preparing Media for Use in the Free Media Pool

Moving Media to a Different Media Pool

You can move media to a different media pool to make it available for use or to allocate it to an application. To do this, follow these steps:

1. In 'Computer Management', open 'Removable Storage' and then expand the 'Libraries and Media Pools' folder by double-clicking them.
2. Select the media pool that contains the media you want to move.
3. In the details pane, drag the media you want to the application media pool in the console tree.

   Caution: Moving media to the 'Free' media pool destroys the data on the media. Additionally, you can't move read-only media to the 'Free' media pool.

Creating Application Media Pools

1. In 'Computer Management', open 'Removable Storage'. In 'Full View', in the console tree, right-click 'Media Pools', and then click 'Create Media Pool'.
2. Or, to create a new media pool within another media pool, right-click the applicable media pool and then click 'Create Media Pool'.
3. On the 'General' tab, in 'Name', type a name for the new media pool; then, in 'Description', type a relevant description.
4. Under 'Media Information', click 'Contains Media of Type', and click the appropriate media type in the list.
5. Under 'Allocation/Reallocation Policy', do one or more of the following:
   - To automatically draw unused media from a free media pool when needed, select the 'Draw Media from Free Media Pool' check box.
   - To automatically return media to a free media pool when no longer needed, select the 'Return Media to Free Media Pool' check box.
   - To set an allocation limit for media in this media pool, select the 'Limit Reallocations' check box, and then change the default value as necessary.

You can create media pools within application media pools only. You cannot create new media pools within free, unrecognized, or imported media pools.

Deleting Application Media Pools

In ‘Removable Storage’, you delete Application Media Pools by right-clicking them and selecting ‘Delete’. Do this only if you no longer need the media pool.

7 Services for NFS/UNIX

Microsoft Services for NFS (MSNFS) and Windows Services for UNIX are comprehensive software packages designed to provide complete UNIX environment integration into a Windows NT, Windows 2000, Windows Storage Server 2003, or Active Directory Domain file server. Services for NFS manage tasks on both Windows and UNIX platforms. Tasks include creating NFS exports from Windows and administering user name mappings.

File services for MSNFS

The following use scenarios are supported by MSNFS file services:

- Allow UNIX clients to access resources on computers running Windows Server 2003 R2.

  Your company may have UNIX clients accessing resources, such as files, on UNIX file servers. To take advantage of new Windows Storage Server 2003 R2 features, such as Shadow Copies for Shared Folders, you can move resources from your UNIX servers to computers running Windows Server 2003 R2. You can then set up MSNFS to enable access by UNIX clients that are running NFS software. All of your UNIX clients will be able to access the resources using the NFS protocol with no changes required.

- Allow computers running Windows Server 2003 R2 to access resources on UNIX file servers.

  Your company may have a mixed Windows and UNIX environment with resources, such as files, stored on UNIX file servers. You can use MSNFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

MSNFS components

MSNFS comprises the following three main components:

- Username Mapping Server:

  Username Mapping Server maps user names between Windows and UNIX user accounts. In a heterogeneous network, users have separate Windows and UNIX security accounts. Users must provide a different set of credentials to access files and other resources, depending on whether they are stored on a Windows or UNIX file server. To address this issue, Username Mapping Server maps the Windows and UNIX user names so that users can log on with their Windows or UNIX credentials and then access resources regardless of whether they are stored on a Windows or UNIX file server.

- Server for NFS:

  Normally, a UNIX computer cannot access files on a Windows-based computer. A computer running Windows Server 2003 R2 and Server for NFS, however, can act as a file server for both Windows and UNIX computers.

- Client for NFS:

  Normally, a Windows-based computer cannot access files on a UNIX computer. A computer running Windows Server 2003 R2 and Client for NFS, however, can access files stored on a UNIX-based NFS server.

Managing User Name Mapping

The ‘User Name Mapping’ component provides centralized user mapping services for Server for NFS and Client for NFS. User Name Mapping lets you create maps between Windows and UNIX user and group accounts even though the user and group names in both environments may not be identical; it also allows you to maintain a single mapping database making it easier to configure account mapping for multiple computers running MSNFS.

In addition to one-to-one mapping between Windows and UNIX user and group accounts, User Name Mapping permits one-to-many mapping. This lets you associate multiple Windows accounts with a single UNIX account. This can be useful, for example, when you do not need to maintain separate UNIX accounts for individuals and would rather use a few accounts to provide different classes of access permission.

You can use simple maps, which map Windows and UNIX accounts with identical names. You can also create advanced maps to associate Windows and UNIX accounts with different names, which you can use in conjunction with simple maps.

User Name Mapping can obtain UNIX user, password, and group information from one or more Network Information Service (NIS) servers or from password and group files located on a local hard drive. The password and group files can be copied from a UNIX host or from a NIS server.

User Name Mapping periodically refreshes its mapping database from the source databases, ensuring that it is always kept up-to-date as changes occur in the Windows and UNIX name spaces. You can also refresh the database anytime you know the source databases have changed.

You can back up and restore User Name Mapping data at any time. Because the database is backed up to a file, you can use that file to copy the mapping database to another server. This provides redundancy for the sake of fault tolerance.

User Name Mapping associates Windows and UNIX user names for Client for NFS and Server for NFS. This allows users to connect to Network File System (NFS) resources without having to log on to UNIX and Windows systems separately.

To configure User Name Mapping, complete the following steps:

1. In Windows Storage Server Management Console, expand the 'Microsoft Service for NFS' to list all the components.
2. Right-click 'User Name Mapping' in the left pane, and select 'Properties' to open the dialog box shown in Figure 7-1.

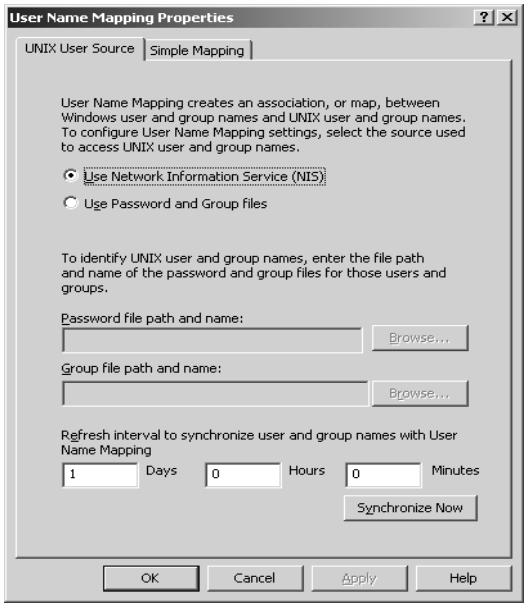

Figure 7-1 User Name Mapping Properties

3. Choose the UNIX authentication method that will be used, either User Network Information Service (NIS) or 'Use Password and Group Files'. If you select 'Use Password and Group Files', you'll need to set the path to the files. Click 'Apply' to apply any changes before changing screens.
4. Select the 'Simple Mapping' tab if your user names are identical in both UNIX and Windows; depending on the case, you can activate simple maps by selecting the check box and selecting the Domain from the drop-down list. If some user names are different, you'll need to use 'Advanced Maps' to handle those users.
5. To create advanced user maps, right-click 'User Maps' and select 'Create Map' from the 'Action' menu to open the 'Create Advanced User Mapping' dialog box shown in Figure 7-2.
6. Select the correct Windows Domain from the drop-down list, and click the 'List Windows Users' button to populate the 'Windows User' list box. Click the 'List UNIX Users' button to populate the 'UNIX User' list box.
7. Highlight a Windows user and a UNIX user, and click 'Add'. Repeat this to add additional advanced maps.
8. Click 'Close' when you have created all the maps you need.

To enable other computers to connect to the 'User Name Mapping Service' on a server, you must edit the `.maphosts` file. This file is in the `%SFUDIR%\Mapper` directory on SFU servers and in `%windir%\msnfs` in Windows Server 2003 R2 or Windows Storage Server 2003 R2 server. If the file is present, but empty, only the local computer can

connect to the User Name Mapping Server. Entries in `.maphosts` explicitly allow or disallow connection by remote computers. Table 7-1 gives the syntax and meaning of entries in the `.maphosts` file.

Table 7-1 .maphosts Entries

host	The host or hosts that resolve to host are permitted access to the User Name Mapping server. An explicit IP address can also be entered.
Host -	The host or hosts that resolve to host are prohibited from access to the User Name Mapping server. An explicit IP address can also be used.
+	All hosts can connect to the User Name Mapping server, unless explicitly excluded by a prior entry. All entries after this are ignored.
-	All hosts are prohibited from connecting to the User Name Mapping server unless explicitly granted access by a prior entry. All entries after this are ignored.
#	This begins a comment line.

You should initially enable all computers to access UNM until you have confirmed that everything is working as desired, and then restrict access to only those computers that need access, including the UNIX computers that will need to connect to UNM.

Connecting to an NFS Share

Microsoft Service for NFS includes and NFS client that enables Windows Storage Server 2003 R2 computers to connect to a shared (exported) NFS file system on a remote UNIX computer. You can use Windows Explore to locate and connect to a remote NFS resource. To use Windows Explorer to locate and connect to NFS shares, do the following.

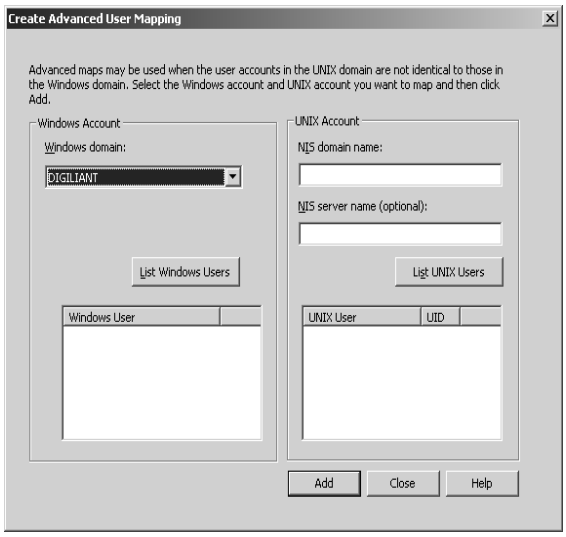

Figure 7-2 Advanced User Mapping

1. Browse to ‘My Network Places’, ‘Entire Network’, ‘NFS Network’.
2. Right-click ‘NFS Network’, and choose ‘Add/Remove NFS LANs’ from the ‘Action’ menu to open the dialog box shown in Figure 7-3.
3. Click ‘Add LAN’ to open the ‘Add Broadcast LAN’ dialog box shown in Figure 7-4.
4. Type in a name for the LAN, the IP address of an NFS server on the LAN, and the subnet mask for the LAN, as shown in Figure 7-4.
5. Browse the LAN to see available NFS exports.
6. Select an NFS export, and then, from the ‘Tools’ menu, select

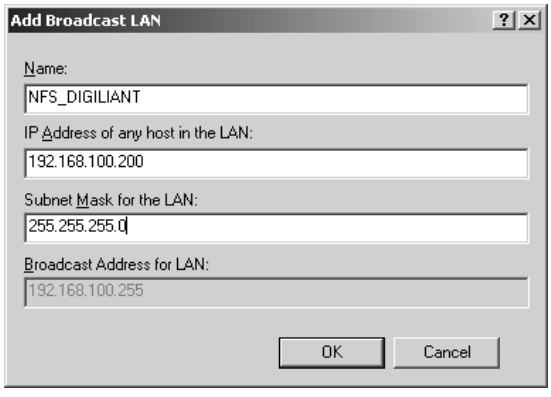

Figure 7-4 Add Broadcast LAN

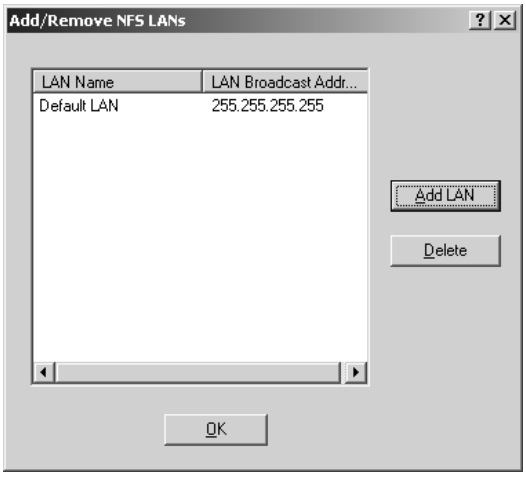

Figure 7-3 Add/Remove NFS LANs

'Map Network Drive' to map the NFS export to a drive letter.

Configuring Client for NFS

There are several options that you can configure in Client for NFS. These include the protocols to use, the default mount type, buffer size, and default file permissions. Any of these options can be manually set for a specific mount by using the mount command with the appropriate options. The options available and what they mean are described in Table 7-2. To set these options, open the 'Microsoft Service for NFS Management' console, right-click 'Client for NFS', and then select 'Properties' from the 'Action' menu.

Table 7-2 Client for NFS Setting

Setting	Page	Options
Transport Protocol	Client Settings	TCP, UDP, TCP+UDP
Use Soft Mounts	Client Settings	Number of times to retry a connection. Default is 1 attempt.
Use Hard Mounts	Client Settings	
Interval between retries	Client Settings	Default is 0.8 seconds, and applies to both hard and soft mounts.
Read buffer size	Client Settings	1-64KB. Default is 32KB.
Write buffer size	Client Settings	1-64KB. Default is 32KB.
Owner	File Permissions	Default file permissions for new files. Read(r), Write(w), and Execute(x) are true by default.
Group	File Permissions	Default file permissions for new files. Read(r), Write(w), and Execute(x) are true by default.
Others	File Permissions	Default file permissions for new files. Read(r), Write(w), and Execute(x) are true by default.

Creating an NFS Share

'Server for NFS' enables you to share (or export, to use the UNIX terminology) a folder or file system. Sharing is simple from either the command line or from Windows Explorer.

The NFS protocol doesn't support sharing the subdirectory of an already shared resource, so you need to ensure that you share from as far up the tree as necessary, because you won't be able to then share another folder within that directory structure. The one exception to this is that each drive letter is shared as the top of a file system.

To share a directory with NFS, complete the following steps;

1. Highlight the directory you want to share in Windows Explorer.
2. Right-click and select 'Sharing and Security' from the 'Action' menu.
3. Click the 'NFS Sharing' tab to open the 'NFS Sharing' dialog box shown in Figure 7-5.
4. Select 'Share Name', and enter the name for export. Set additional options as described in Table 7-3.

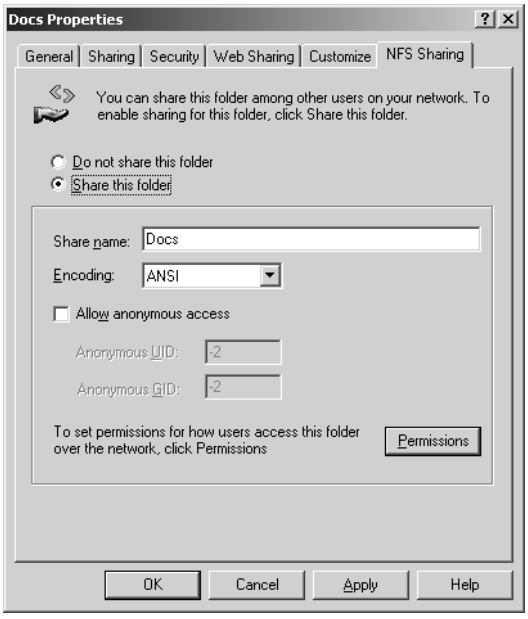

Figure 7-5 NFS Sharing Properties

Table 7-3 Properties of NFS Sharing

Property	Default	Options
Share name	The shortest name of the directory	
Encoding	ANSI	ANSI, Big 5, EUC-JP, EUC-KR, EUC-TW, GB2312-80, KSC5601, ShiftJIS
Allow anonymous	No	When true, you can also set the default UIS and GID for an anonymous user.
Permissions	All Machines	Read-only, ANSI, Root Access Denied. Click Add to add additional machines with specific permissions levels.
Type of access	Read-only	No access, Read-only, Read/Write.
Allow root access	No root access	Allow or deny root access.

5. Click 'OK' to exit and make the share available to NFS clients on the network.

Configuring Server for NFS

There are several options that you can configure in Server for NFS. These include the protocols and NFS version to use, filename handling, locking options, and auditing. The options available and what they mean are described in Table 7-4. To set these options, open the 'Microsoft Services for NFS Management' console, right click 'Server for NFS', and select 'Properties' from the 'Action' menu.

Table 7-4 Server for NFS Properties

Setting	Page	Options
Protocols	Server Settings	TCP, UDP, TCP+UDP. Default is TCP+UDP
Enable NFS v3 Support	Server Settings	Enabled by default. When not selected, only NFSv2 support is available.
Authentication	Server Settings	Renewal period set in seconds or don't renew at all.
Translate File Names	Filename handling	Disabled by default. You can specify a character translation table.
Dot files as hidden	Filename handling	When checked, files that begin with a period (.) are marked hidden.
Enable case-sensitive lookups	Filename handling	Enabled by default.
Return filenames in:	File Permissions	Preserve case, all lowercase or all uppercase.
Waiting Period	Locking	45 seconds is the default period to wait for clients to reclaim locks after connectivity is interrupted.
Existing Client Locks	Locking	Any existing locks are shown and can be forced to release.
Log events to event log	Audit Logging	Disabled by default.
Log events to a text file	Audit Logging	Disabled by default. When enabled, you can specify a filename and maximum file size.
Event to log	Audit Logging	The options are mount, locking, read, write, create, delete, and All.

Enabling NFS event logging will have a serious impact on overall NFS performance and can rapidly fill the event log or create a very large text file. Logging should be enabled only for sufficient time to troubleshoot an issue and then disabled.

Windows Subsystem for UNIX-Based Applications

The Windows Subsystem for UNIX-Based Applications (SUA) is a full-featured, POSIX-compliant, UNIX application environment with more than 2000 UNIX APIs. The SUA shells and applications run as a full subsystem on the Windows kernel and support standard UNIX shell programs and applications.

SUA gives the UNIX programmer or user a familiar environment, with a single-rooted file system that supports typical file locations such as `/etc, /usr/bin`, and `/usr/local/bin`. SUA supports symbolic and hard links that are transparent to the UNIX user. SUA can also support full-case sensitivity and SUID behavior if required.

8 NetWare and Macintosh Interoperability

File and Print Services for Netware

File and Print Services for NetWare (FPNW) is one part of the Microsoft software package called 'Services for NetWare'. The most common use of the NetWare network operating system is as a file and print server. FPNW eases the addition of the Storage Server into a mixed infrastructure by providing a NetWare user interface (UI) to a Windows Storage Server 2003-based server; administrators and users see their same, familiar NetWare UI. Additionally, the same logon for clients is maintained without a need for any client configuration changes.

This service also provides the ability to create Novell volumes, which are actually NTFS shares, from which users can map drives and access resources. Novell Login scripts are supported on the Storage Server or through an existing NDS (Novell Directory Services) account.

Installing Services for NetWare

The installation of FPNW on the Storage Server allows for a smooth integration with existing Novell servers. FPNW allows a Windows Storage Server 2003 based server to emulate a NetWare file and print server to users, clients, and administrators. This emulation allows authentication from Novell clients, the use of Novel logon scripts, and the creation of Novell volumes (shares), the use of Novell file attributes, and many other Novell features.

To install Services for NetWare:

1. In Windows Storage Server 2003 desktop, choose 'Start' → 'Control Panel' → 'Network Connections' and right-click 'Local Area Connection', and then select 'Properties'.

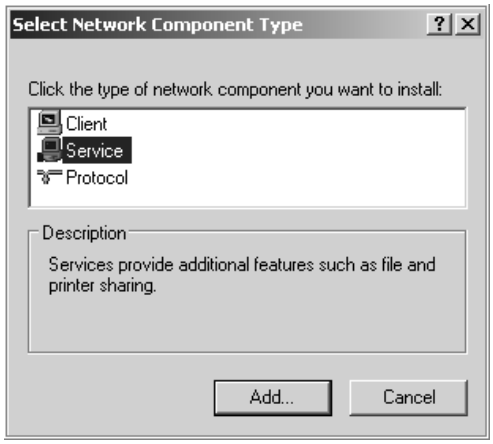

Figure 8-1. Local Area Connection Properties page, Install option

2. Click 'Install'. The 'Select Network Component Type' dialog box is displayed.
3. Click 'Service', and then click 'Add'.
4. Click the 'Have Disk' icon, and then navigate to the location of 'Services for NetWare'. 'Services for NetWare' is located on the Resource DVD.

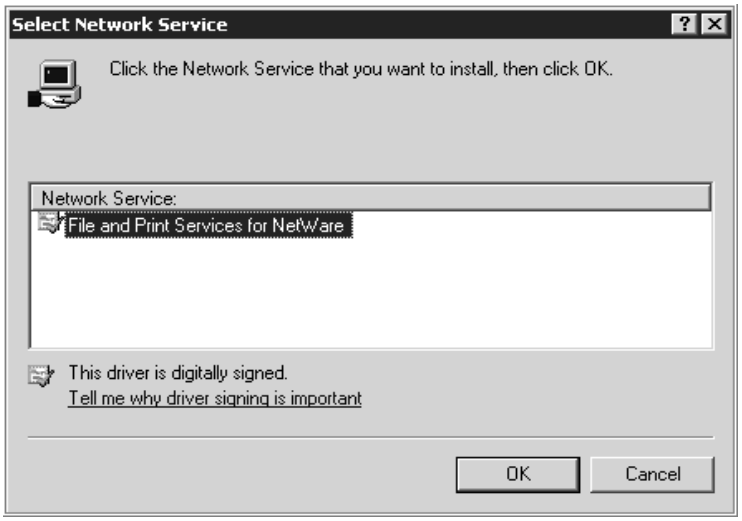

Figure 8-2. Installing File and Print Services for NetWare

5. Select the NETSFNTSRV file, and then click 'OK'.
6. 'File and Print Services for NetWare' should now be displayed as an option to install.
7. Select 'File and Print Services for NetWare', and then click 'OK'.

Managing File Services for NetWare

To access FPNW:

1. In Windows Storage Server 2003 desktop, choose 'Start' → 'All Programs' → 'Administrative Tools' and then select 'Server Manager'.

2. Select FPNW, and then click ‘Properties’.
3. Enter an ‘FPNW Server Name’ and ‘Description’. This server name must be different from the server name used by Windows or LAN Manager-based clients. If changing an existing name, the new name is not effective until stopping and restarting FPNW. For example, in Figure 8-3 the Windows server name is Storage Server and the FPNW server name is `NASSSERVER_FPNW`.
4. Indicate a ‘Home Directory Root Path’. This path is relative to where the `sysvol` volume is installed. This is the root location for the individual home directories. If the directory specified does not already exist, it must first be created.

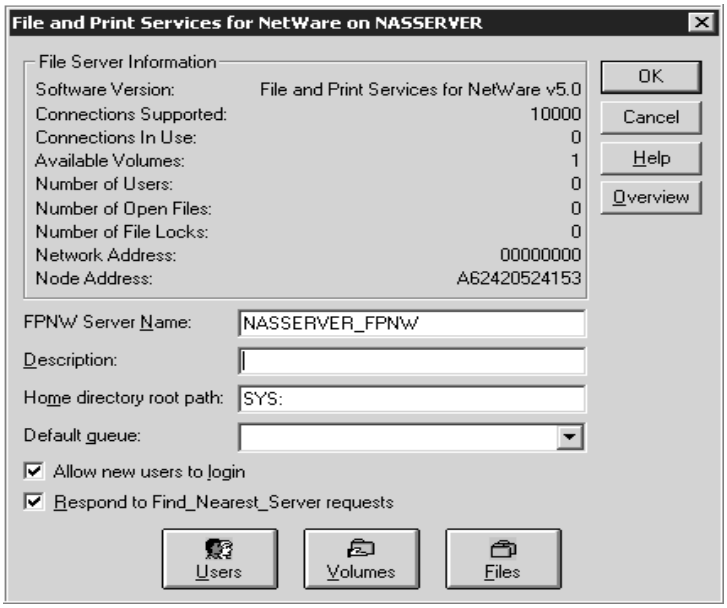

Figure 8-3 File and Print Services for NetWare dialog box

5. Click ‘Users’ to: see connected users, disconnect users, send broadcast messages to all users connected to the server, and to send a message to a specific user.
6. Click ‘Volumes’ to: See users connected to specific volume and to disconnect users from a specific volume.
7. Click ‘Files’ to: View open files and close open files.

Creating and managing NetWare users

To use ‘Services for NetWare’, the Novell clients must be entered as local users on the Storage Server.

Adding local NetWare users

1. Open ‘Computer Management’ console, expand the console tree to ‘Local Users and Groups’, right-click ‘Users’ and select ‘New User’. This opens the ‘New User’ dialog, shown in Figure 8-4.
2. Enter the user information, including the user's User name, Full name, Description, and Password.
3. Click ‘Create’.
4. Repeat these steps until all NetWare users have been entered.

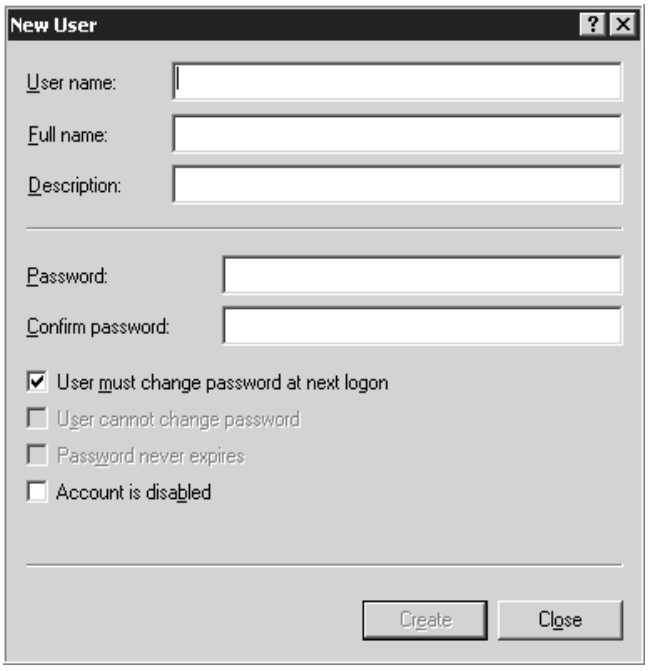

Figure 8-4 New User dialog box

Enabling local NetWare user accounts

1. Open ‘Computer Management’ console, expands the console tree to ‘Local Users and Groups’, and then click the ‘Users’. This lists all the local users.
2. Right-click an NCP client listed in the right pane of the screen, and then click ‘Properties’.
3. Click the ‘NetWare Services’ tab.

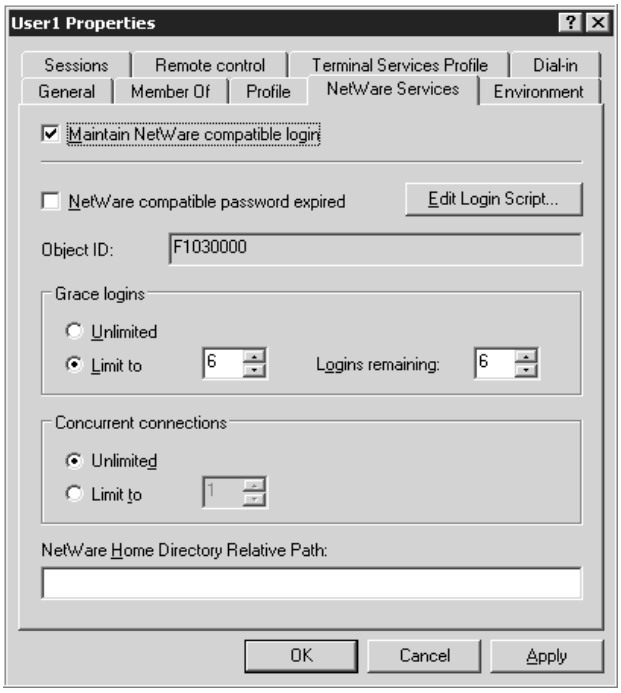

Figure 8-5 NetWare Services tab

3. Specify the volume name and path.
4. Click 'Permissions' to set permissions.
5. Click 'Add' to add additional users and groups, and to set their permissions.

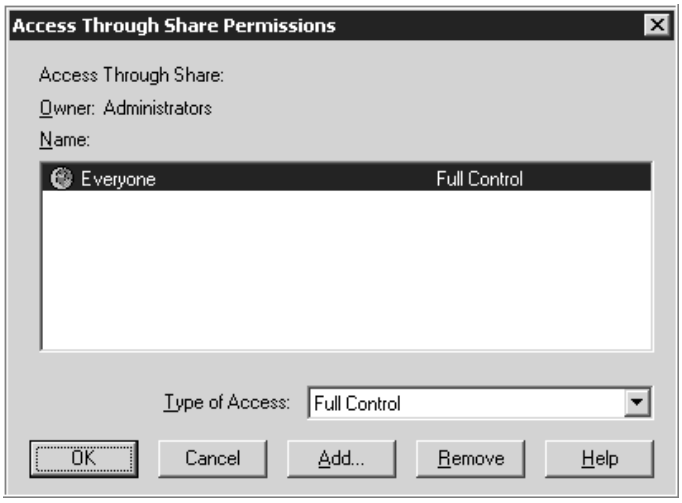

Figure 8-7 Access Through Share Permissions dialog box

4. Select Maintain NetWare compatible login.
5. Set other NetWare options for the user, and then click 'OK'.

Managing NCP volumes (shares)

NCP file shares are created the same way as other file shares; however, there are some unique settings. NCP shares can be created and managed using Server Manager.

Creating a new NCP share

To create a new file share:

1. In Windows Storage Server 2003 desktop, choose 'Start' → 'All Programs' → 'Administrative Tools' → 'Server Manager' and then select Shared Volumes from FPNW.
2. Click 'Create Volume'.

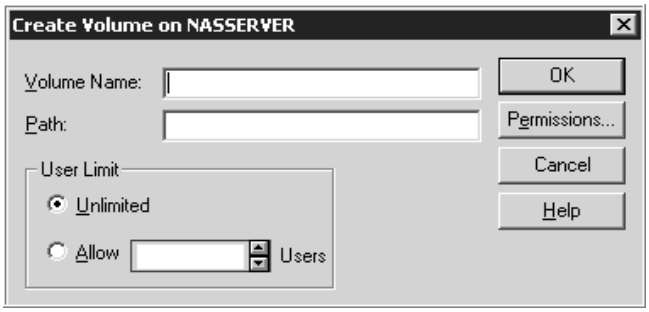

Figure 8-6 Create Volume dialog box

6. Highlight the desired user or group, and then click 'Add'.
7. Select the 'Type of Access' in the drop down list. The 'Type of Access' can also be set from the 'Access Through Share Permissions' dialog box.
8. Click 'OK' when all users and groups have been added.
9. Click 'OK' in the 'Create Volume' dialog box.
10. Click 'Close'.

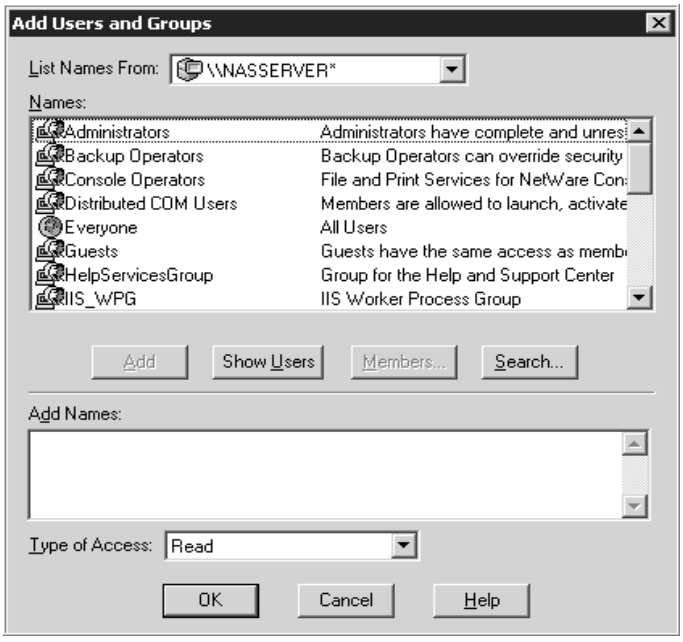

Figure 8-8 Add Users and Groups dialog box

Modifying NCP share properties

To modify a file share:

1. In Windows Storage Server 2003 desktop, choose 'Start' → 'All Programs' → 'Administrative Tools' and the select Server Manager.
2. Select 'File and Print Services for NetWare', and then click 'Shared Volumes'.
3. Highlight the volume to modify.
4. Click 'Properties'.

AppleTalk and File Services for Macintosh

The AppleTalk network integration allows the Storage Server to share files and printers between your server and any Apple Macintosh clients that are connected to your network. After installing Microsoft Windows Services for Macintosh, the administrator can use the AppleTalk protocol to configure the Storage Server to act as an AppleTalk server. The AppleTalk protocol is the communications protocol used by clients running a Macintosh operating system. The Macintosh computers need only the Macintosh OS software to function as clients; no additional software is required.

AppleTalk network integration simplifies administration by maintaining just one set of user accounts instead of separate user accounts, for example, one on the Macintosh server and another on the computer running Windows server software.

Installing the AppleTalk protocol

1. Access 'Network Connections' in 'Control Panel', right-click 'Local Area Connection' and then click 'Properties'.
2. Click 'Install'. Figure 8-9 is an example of the 'Select Network' component.
3. Select 'Protocol', and then click 'Add'.
4. Select 'AppleTalk Protocol', and then click 'OK'.

Installing File Services for Macintosh

To install File Services for Macintosh, perform the following steps:

1. Open 'Add or Remove Programs' from the 'Control Panel'.
2. Click 'Add or Remove Windows Components'.
3. Double-click 'Other Network File and Print Services'.
4. Select 'File Services for Macintosh', and then click 'OK'.

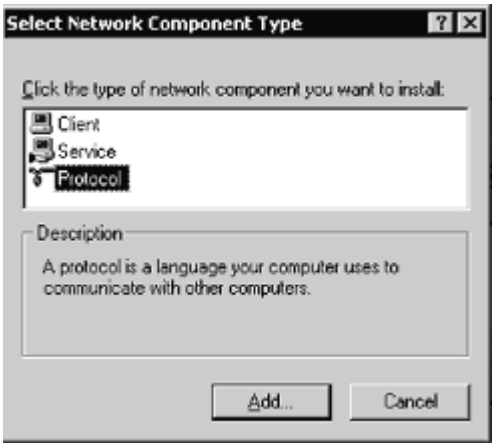

Figure 8-9 Local Area Connection Properties page, Install option

5. Click ‘Next’.
6. Click ‘Finish’.

Completing setup of AppleTalk protocol and shares

See the online Help to complete the following set up and configurations tasks:

- To set up AppleTalk protocol properties.
- To set up AppleTalk shares.
- To configure AppleTalk sharing properties.
- To allow client permission to an AppleTalk share.

9 Print Management

Print services

Printer services support network printers only and are not intended for use with locally attached printers (USB or Parallel port connected).

If the Storage Server is a part of an Active Directory Domain rather than Workgroup, the Storage Server enables the following management features:

- Restrict access to printer based Domain user accounts
- Publish shared printers to Active Directory to aid in search for the resource

Before adding a print server role, the following checklist of items should be followed:

- **Determine the operating system version of the clients that will send jobs to this printer.** This information is used to select the correct client printer drivers for the client and server computers utilizing the printer. Enabling this role on the print server allows the automatic distribution of these drivers to the clients. Additionally, the set of client operating systems determines which of these drivers need to be installed on the server during the print server role installation.
- **At the printer, print a configuration or test page that includes manufacturer, model, language, and installed options.** This information is needed to choose the correct printer driver. The manufacturer and model are usually enough to uniquely identify the printer and its language. However, some printers support multiple languages, and the configuration printout usually lists them. Also, the configuration printout often lists installed options, such as extra memory, paper trays, envelope feeders, and duplex units.
- **Choose a printer name.** Users running Windows-based client computers choose a printer by using the printer name. The wizard that you will use to configure your print server provides a default name, consisting of the printer manufacturer and model. The printer name is usually less than 31 characters in length.
- **Choose a share name.** A user can connect to a shared printer by entering this name, or by selecting it from a list of share names. The share name is usually less than 8 characters for compatibility with MS-DOS and Windows 3.x clients.
- (Optional) **Choose a location description and a comment.** These can help identify the location of the printer and provide additional information. For example, the location could be “Second floor, copy room” and the comment could be “Additional toner cartridges are available in the supply room on floor 1."

Configuring the print server

To set up a print server:

1. In Windows Storage Server Management console, expand the console tree to ‘Print Management’ → ‘Print Servers’ and then right-click ‘Print Servers’ and select ‘Add/Remove Servers’. This opens the ‘Add/Remove Servers’ dialog, show in Figure 9-1.
2. Type the print server’s name or IP address in the ‘Add Server’ box, or click ‘Browse’ to search the print servers on the network. To set this server as a print server, click ‘Add the Local Server’ button.
3. Click ‘OK’ when finished. The newly added print server should be listed on the ‘Print Servers’ list.
4. Right-click the newly added print server and select ‘Add Printer’. This opens the ‘Welcome to the Add Printer Wizard’. Click ‘Next’.
5. Select ‘Local Printer’, clear “automatically detect install my plug and play printers,” and then click ‘Next’.

6. Select 'Create a New Port', and then select 'Standard TCP/IP Port (recommended)'. The 'Add Standard TCP/IP Printer Port Wizard' starts.
7. Click 'Next'.
8. Enter the name or IP address of the printer. The IP address is usually listed on the printer configuration page. The wizard completes the 'Port Name' field. Click 'Next'.
9. The wizard attempts to connect to the printer. If the wizard is able to connect, the 'Completing the Add Standard TCP/IP Printer Port Wizard' page opens; click 'Finish'. If the wizard is not able to connect, the 'Additional Port Information Required' page opens.
    - Verify that the IP address or name is correct.
    - Select '**Standard'** to identify the printer network adapter. A list of manufacturers and models of the network adapters is displayed. Select the appropriate printer in the Standard list.

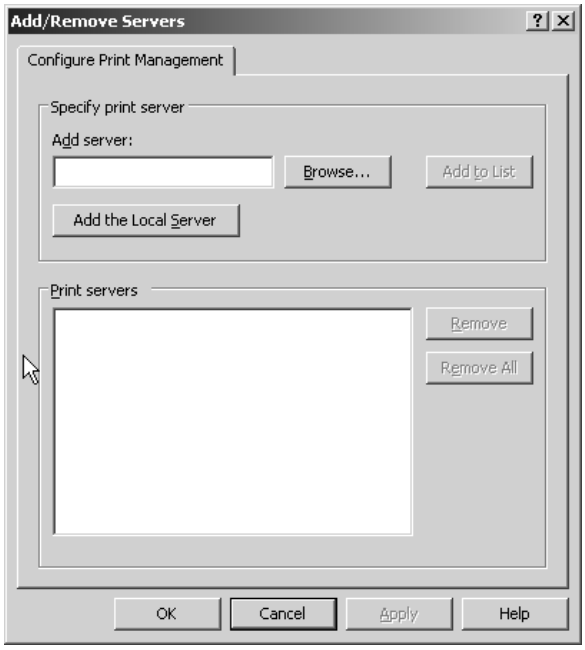

Figure 9-1 Add/Remove Print Server

    - If the printer network adapter uses nonstandard settings, click '**Custom',** and then click '**Settings'.** The '**Configure Standard TCP/IP Port Monitor'** page opens. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click '**OK'.**
    - Click '**Next'.**
10. Select the manufacturer and the type of printer in the presented list, and then click 'Next'. If the printer does not exist in the list, click 'Have Disk' and load the drivers, or select a compatible driver.
11. Enter the name of the desired printer to be presented on the Storage Server, and then click 'Next'.
12. Enter a Share Name for the printer to be used on the network, and then click 'Next'.
13. Enter a location description and a comment, and then click 'Next'.
14. Select 'Print a Test Page', and then click 'Next'.
15. Clear the 'Restart the Add Printer Wizard' if adding only one printer, and then click 'Finish'. A test page prints.
16. Click 'OK' if the page printed, otherwise click 'Troubleshoot'. If 'Restart the Add Printer Wizard' was selected, the wizard restarts to add an additional printer.
17. Repeat the steps above for adding an additional printer.

Removing the print server role

To remove the print server role:

1. In the Windows Storage Server Management console, expand the console tree to 'Print Management' → 'Print Servers' and then right-click 'Print Servers' and select 'Add/Remove Servers'. This opens the 'Add/Remove Servers' dialog, show in Figure 9-1.
2. Highlight the print server you want to remove from the 'Print Servers' list box, and then click 'Remove'.
3. Click 'OK'.

Adding an additional printer

To add additional printers to the Storage Server:

1. In the Windows Storage Server Management console, navigate to the 'Print Servers' list. Right-click the print server you want to manage, and select 'Add Printer' to start the 'Add Printer Wizard'.
2. Follow the steps on "Configuring a Print Server' to finish the wizard.

Adding additional operating system support

By default, support is added for Windows 2000 and Windows XP. If the client base is com posed of other Windows operating systems, additional printer drivers must be loaded. To load an additional driver for client download:

1. In the Windows Storage Server Management console, expand the console tree to 'Print Management' → 'Printer Servers' and then click on the printer server you want to manage.
2. Click 'Drivers' and select 'Add Driver'. This opens 'Welcome to the Add A Printer Driver Wizard'.
3. Follow the wizard to finish installation of additional drivers.

Print services for UNIX

Network clients with UNIX-based operating systems that use the client program line printer remote (LPR) can send printing jobs to the line printer daemon (LPD) on the Storage Server. LPR clients must comply with Request for Comments (RFC) 1 179. The combination of the LPR and LPD are included in print services for UNIX. 'Print Services for UNIX' is not pre-installed on the print server or the 'File Print Appliance'.

To install print services for UNIX:

1. Log on as administrator or as a member of the Administrators group.
2. Choose 'Start' → 'Control Panel' and then click 'Add or Remove Programs'.
3. Click 'Add/Remove Windows Components'.
4. In the 'Components' list, click 'Other Network File and Print Services' (but do not select or clear the check box), and then click 'Details'.
5. In the 'Subcomponents of Other Network File and Print Services' list, select 'Print Services for UNIX', if appropriate to the print services that you want to install.
6. 'Print Services for UNIX': This option permits UNIX clients to print to any printer that is available to the print server.
7. Click 'OK', and then click 'Next'.
8. Click 'Finish'.

Print Services for NetWare

With File and Print Services for NetWare installed, the print server or 'File Print Appliance' appears to a NetWare client as a NetWare 3.x–compatible print server. Print services presents the same dialog boxes to the client as a NetWare-based server uses to process a print job from a client. A user can display and search for printers on the print server or 'File Print Appliance' just like in a NetWare environment.

Installing Print Services for NetWare

Refer to the previous section "Installing Services for Netware" for information on installing Print Services for NetWare.

Print services for Macintosh

Macintosh clients can send print jobs to a print server or File Print Appliance (FPA) when Print Server for Macintosh is installed on the server. To the Macintosh-based client, the print server or FPA appears to be an AppleTalk printer on the network, and no reconfiguration of the client is necessary.

Installing Print Services for Macintosh

Consult the following resource for information about installing Print Services for Macintosh:

- How To: Install Print Services for Macintosh in Windows Storage Server 2003 R2 `http://support.microsoft.com/Default.aspx?scid=kb;en-us;323421`

Point and Print from UNIX, Netware and Macintosh to Windows Storage Server 2003 R2

Point-and-Print behavior from UNIX, Netware and Macintosh clients to Windows Storage Server 2003 R2 or Windows Storage Server 2003 is similar to the behavior for Windows 95, Windows 98, and Windows Millennium Edition clients, because all these clients create SMB connections. However, the non-Windows operating systems maintain their own driver model, so these clients do not automatically get the driver during Point and Print—they must install the driver locally. Like the Windows 95, Windows 98, and Windows Millennium clients, these non-Windows clients do not receive driver updates from the print server after a driver is initially downloaded. The same connection methods are available: drag and drop, the 'Add Printer Wizard', referencing a UNC path, or double-clicking the shared printer icon.

10 System Installation and Recovery

The Installation and Recovery DVD

The Digiliant Storage Server System Installation and Recovery DVD that is provided with your Storage Server will allow you to install an OS image or recover from a catastrophic failure. The DVD is used initially to install and configure the operating system and applications provided with your Storage Server.

At any later time, you may boot from the DVD and restore the server to the original factory configuration. This allows you to recover the system if all other means to boot the server fails.

While the recovery process makes every attempt to preserve the existing data volumes, you should have a backup of your data if at all possible before recovering the system.

To restore a factory image

1. Connect keyboard, monitor, and mouse directly to the Storage Server.
2. Insert the 'Restore DVD'. Reboot the computer.
3. Select 'Boot from CDROM'.
4. After the main window appears, choose the option to restore the factory image.
    - Restore OS to factory default: This option will wipe everything on the first logical disk (disk 0); and then restores the OS back to the factory default. After you chose this option, you will see a list of all logical disks. You should carefully review the list to make sure that disk 0 is selected. Confirm the action by typing "yes", followed by the Enter key, and the restore process will start. After the restore process finishes, you should see a message indicating the process was successful. If you experience any problem during this restoring process, please call Digiliant for support.
    - Go to command prompt: This option will take you to the "Command Prompt". You can connect this computer to other computers on the network through Microsoft's Net command. After you connect this computer to the network, you can backup to or restore from another computer on the network by using Imagex.exe command. At beginning of the command prompt, you will see a list of sample commands. By following those sample commands, you will be able to backup or restore the OS from the network. If you need to wipe the Disk 0 before restoring the image, use the command "diskpart /s x:\windows\system\disk.txt".
    - Restart the computer: Reboot the computer without modifying the system disk.

On systems with more than two physical drives, the system may be recovered without affecting data volumes. This is dependant on the systems hardware RAID configuration.