TANDBERG VIKING SERIES NAS APPLIANCE

Storage Server Administration Manual

TANDBERG DATA ASIA
7 Tai Seng Drive #02-00
Singapore 535218
Phone +65 6593 4700
Telefax +65 6281 7358

Part No. 65 82 X2 - 02
April 2009

Related publications available from Tandberg Data Asia:

Part No.	Title
6582B7	Tandberg Viking FS-1600 NAS Hardware User Manual
6582A7	Tandberg Viking FS-1500 NAS Hardware User Manual
6582C7	Tandberg Viking FS-420U NAS Hardware User Manual

Table of Contents

1. System Overview

1.1. Product Information

The Tandberg Viking Series NAS appliance can be used in many types of computing environments, from basic Microsoft Windows workgroups to complicated multi-protocol domains using DFS, NFS, FTP, HTTP, and Microsoft SMB. The corresponding varieties of clients that can be serviced include any Windows, UNIX, Linux, Novell, or Macintosh variant.

This chapter provides an overview of these environments and deployments and includes brief descriptions of the available user interfaces. The Viking Series NAS appliance is the disk-based storage for remote office or small to medium business class NAS solutions that provide reliable performance, manageability, and fault tolerance.

The Viking Series NAS appliance provides performance gains over general purpose servers by integrating optimized hardware components and specialized operating software. Integrating NAS appliance into the network improves the performance of existing servers because NAS appliances are optimized for file serving tasks.

Notes

The Viking Series NAS appliance has been specifically designed to function as a Network Attached Storage server. Except as specifically authorized by Tandberg Data, you may not use the server software to support additional applications or significant functionality other than system utilities or server resource management or similar software that you may install and use solely for system administration, system performance enhancement, and/or preventative maintenance of the appliance.

Your Viking Series NAS appliance comes preinstalled with either the Windows® Storage Server™ 2003 R2 operating system (32-bit version) or the Microsoft® Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system.

Microsoft Windows Storage Server 2003 R2 extends the Windows Storage Server 2003 operating system, providing a more efficient way to manage and control access to local and remote resources. In addition, Windows Storage Server 2003 R2 provides a scalable, security-enhanced Web platform for simplified branch server management, improved identity and access management, and more efficient
storage management.

Notes

The Microsoft® Windows® Storage Server 2003 x64 Edition operating system is designed to support 32–bit applications without modification; however, any 32–bit applications that are run on this operating system should be thoroughly tested before releasing the storage server to a production environment.

Microsoft® Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system provides unified storage server management capabilities, simplified setup and management of storage and shared folders, and support for Microsoft iSCSI Software Target.

Notes

For more information about Microsoft® Windows® Unified Data Storage Server 2003 operating system, see “Using iSCSI Software Target” on chapter 9.

1.1.1. Product Manageability

The Viking Series NAS appliance ships with the following utilities and features that ease the administration tasks associated with managing the system:

- The Recovery Disc (factory image) contains the preconfigured default settings of either the Windows Storage Server 2003 R2 operating system or the Microsoft® Windows® Unified Data Storage Server 2003 operating system. This is a quick and easy way to setup or bring the Viking Series NAS appliances back to the factory default configuration.
- Using the Windows Remote Desktop client to establish administrative session with the appliance without physically connecting to it.
- Ability to connect directly to the NAS appliance’s console.

1.2. Redundancy

The Viking Series NAS appliance is specifically designed to perform file serving tasks for networks, using industry standard components to ensure reliability. Other industry standard features, such as redundant array of independent drives (RAID) and remote manageability, further enhance the overall dependability of the NAS appliance.

To ensure redundancy and reliability, it is recommended that the hard drives installed in the Viking Series NAS appliance are configured so that a single drive failure will not cause data loss or system failure.

Depending on the model of Viking Series NAS used, it is capable of support from 4 HDD to 16 HDD in a single chassis:

- Viking FS-412 supports up to 4 HDD (1U)
- Viking FS-1600 / FS-1610 supports up to 16 HDD (3U)

1.2.1. Configuring RAID using 4 HDD (1U)

For the 1U appliance (FS-412) Tandberg Data recommends two logical Volume sets consisting of System Volume (SV) and Data Volume (DV) to be created within the RAID set. The RAID set is first created which comprises of member HDD (number of HDD in a set). The two logical volumes (SV and DV) are then allocated from aggregated capacity of the RAID set with the desired RAID levels. The System Volume, installed with the host operating system, is organized as RAID 0+1 and the Data Volume is organized as RAID 5. This is to allow OS redundancy if more than one HDD fails. Please see the Viking Series NAS Hardware User Manual for details.

The volume layout for the appliance is illustrated below.

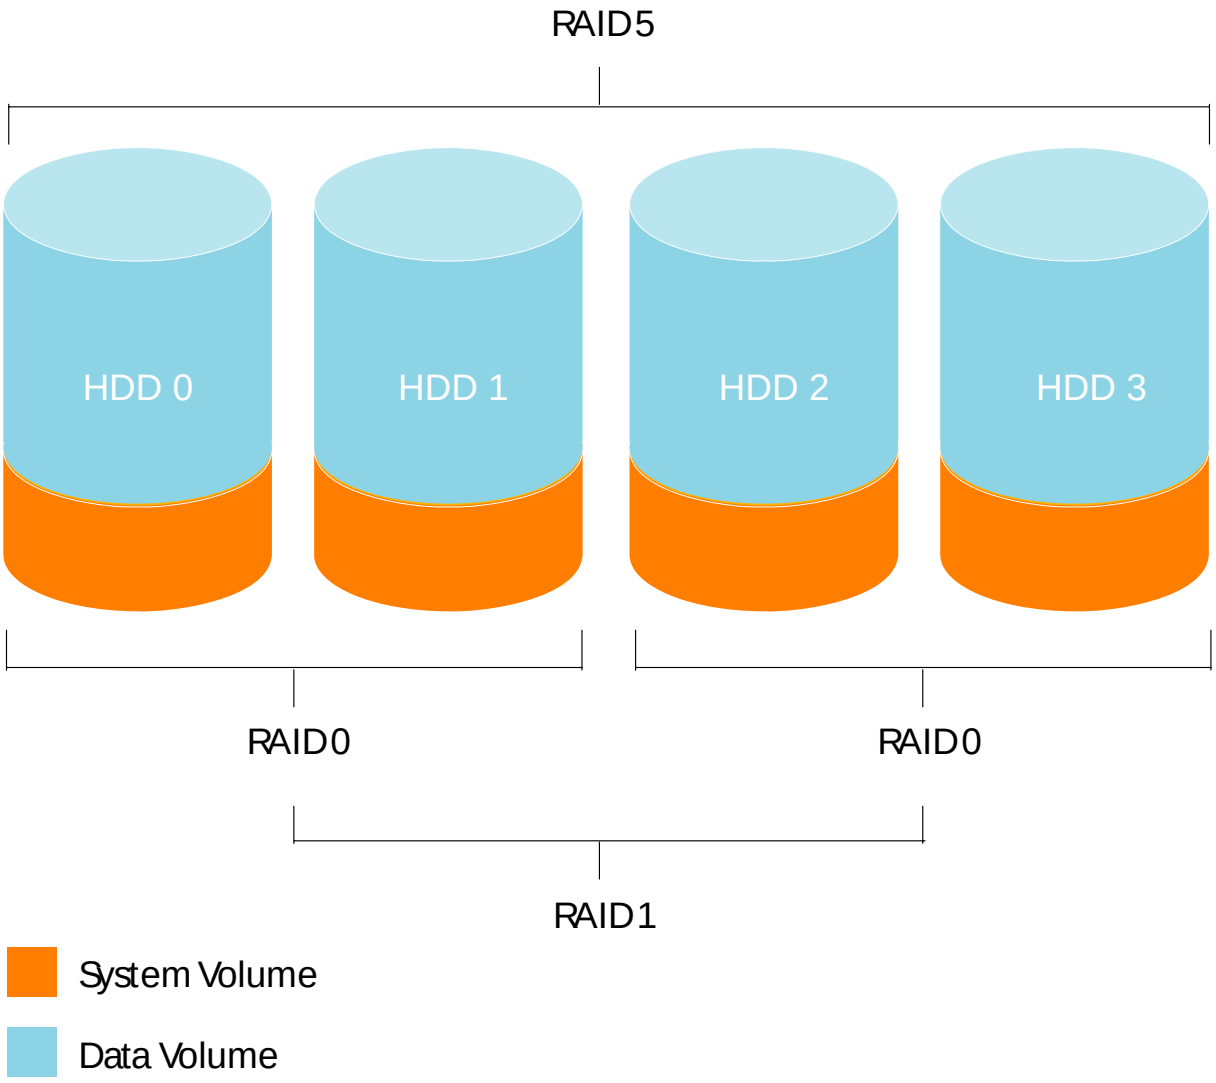

1.2.2. Configuring RAID using 16 HDD (3U)

The 3U appliance uses (FS-1600 / FS-1610) supports up to a maximum of 16 HDD. This provides greater degree of flexibility in RAID configurations:

- Allow HDD to be grouped into multiple RAID sets with respect to their usage role: System Volume, Data Volume(s) or RAW Volumes (for iSCSI target volume provisioning). HDD of differing capacity can also be grouped into similar RAID sets.

- Assigning unused HDD as hot-spares.

- Multiple RAID level support on Volume sets.

Depending on the number of HDD used and their capacity, the RAID Set grouping and the RAID level of each Volume Set is easily customizable.

Below is a working example:

Using 16 HDD in which 4 is 750GB and 12 is 1,000GB in capacity. The HDD can be grouped into their respective RAID Set:

RAID Set#00 is created to house the OS as well as a File-system for sharing. Volume Set#00 is created with RAID level 0+1 which contains a usable capacity of 200GB. The remaining RAID Set#00 capacity is sliced into Volume Set#01 of RAID level 5 that yield a resultant capacity of 1,950GB. Volume Set#00 is the System Volume (SV) and Volume Set#01 is the Data Volume (DV).

- RAID Set#00 = 4 x 750GB
- Volume Set#00 = 200GB (RAID 0+1) SV
- Volume Set#01 = 1,950GB (RAID 5) DV

RAID Set#01 is created with iSCSI target LUN provisioning in mind. Ten (10) of the 1,000GB HDD are grouped into RAID Set#01 and Volume Set#00 of RAID level 6 is created out of it. This yields a usable capacity of 8,000GB after the RAID 6 overhead.

- RAID Set#01 = 10 x 1000GB
- Volume Set#00 = 8,000GB (RAID 6)

The remaining two (2) of the 1,000GB HDD are set as hot-spares.

- Hot-spare = 2 x 1,000GB

The hot-spare can be assigned to any degraded RAID Set for rebuilding in event that one (or more) of its member disk have failed.

Notes

While assigning hot-spare to a degraded RAID Set, ensure that the capacity of the hot-spare must be equal or greater than the lowest capacity of the member disk of that RAID Set.

1.2.3. System Volume (SV)

This volume is organized as a RAID 0+1 volume. This volume includes the operating system image. The minimum recommended size for this volume is 10 GB (default). The RAID level can be customized using the RAID controller BIOS or web-based RAID manager. The web-based RAID manager is only available after OS installation or recovery.

1.2.4. Data Volume (DV)

The remaining hard drives space can be organized as the Data Volume using RAID 5 or other RAID level definable by user. For a FS-412 NAS appliance installed with four 250 GB hard drives, the data volume size can be set to 735 GB using the recommended RAID and Volume settings.

1.4. Deployment

Various deployment scenarios are possible. Typical application of NAS appliances include:

1.4.1. File Server Consolidation

As businesses continue to expand their information technology (IT) infrastructures, they must find ways to manage larger environments without a corresponding increase in IT staff. Consolidating many servers into a single NAS appliance reduces the number of points of administration and increases the availability and flexibility of storage space.

1.4.2. Multi-protocol Environments

Some businesses require several types of computing systems to accomplish various tasks. The multi-protocol support of the NAS appliance allows it to support many types of client computers concurrently.

1.4.3. Protocol and platform transitions

When a transition between platforms is being planned, the ability of the NAS appliance to support most file sharing protocols allows companies to continue to invest in file storage space without concerns about obsolescence. For example, an administrator planning a future transition from Windows to Linux can deploy the NAS appliance with confidence that it can support both CIFS and NFS simultaneously, assuring not only a smooth transition, but also a firm protection of their investment.

1.4.4. Remote office deployment

Frequently, branch offices and other remote locations lack dedicated IT staff members. An administrator located in a central location can use the Microsoft Terminal Services, and other remote administration methods to configure and administer all aspects of the NAS server.

1.5. Environment scenarios

The NAS appliance is deployed in one of two security modes:

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

The NAS appliance uses standard Windows user and group administration methods in each of these environments. For procedural instructions on managing users and groups, see Chapter 5 of this Manual.

Regardless of the deployment, the NAS appliance integrates easily into multi-protocol environments, supporting a wide variety of clients. The following protocols are supported:

- Distributed File System (DFS)
- Network File System (NFS)
- Hypertext Transfer Protocol (HTTP)
- File Transfer Protocol (FTP)
- Microsoft Server Message Block (SMB)

1.5.1. 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.

1.5.2. Domain

When operating in a Windows NT or Active Directory domain environment, the NAS appliance 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. Additional information about planning for domain environments can be found at Microsoft web site.

The NAS appliance obtains user account information from the domain controller when deployed in a domain environment. The NAS 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.

1.6. User Interfaces

There are several user interfaces that administrators can use to access and manage the Viking Series NAS appliance. Two of these interfaces are:

- Using direct attached method
- Using Windows Remote Desktop

Each interface contains the same or similar capabilities, but presents them in a different manner. Each of these interfaces are illustrated in the following sections.

1.6.1 Direct Attached Method

The Viking Series NAS appliance can be accessed directly be connecting a keyboard, mouse, and monitor.

The default user name is **"Administrator"**. The default password is **"1234"**.

1.6.2. Using Windows Remote Desktop Console

The NAS appliance desktop console can be accessed remotely using Windows Remote Desktop. This requires the use of the Windows Remote Desktop client on the machine it is accessing from.

Remote Desktop provides the ability for you to log onto and remotely administer your server, giving you a method of managing it from any client. Installed for remote administration, Remote Desktop allows only two concurrent sessions. Leaving a session running takes up one license and can affect other users. If two sessions are running, additional users will be denied access.

To connect the storage server to a network using the Remote Desktop method:

- On the PC client, select **Start** > **Run**. At Open, type `mstsc`, then click **OK**.
- Type the IP address of the NAS appliance you will be connecting to in the **Computer** box and click **Connect**.
- Use the default user name **"Administrator"** and default password **"1234"** to gain system access.

Notes
When using Remote Desktop to connect to the NAS appliance desktop do not use the window close feature (X). Click on Start/Log Off Administrator to exit Remote Desktop. See "Improper Closure of Remote Desktop" in Chapter 2.

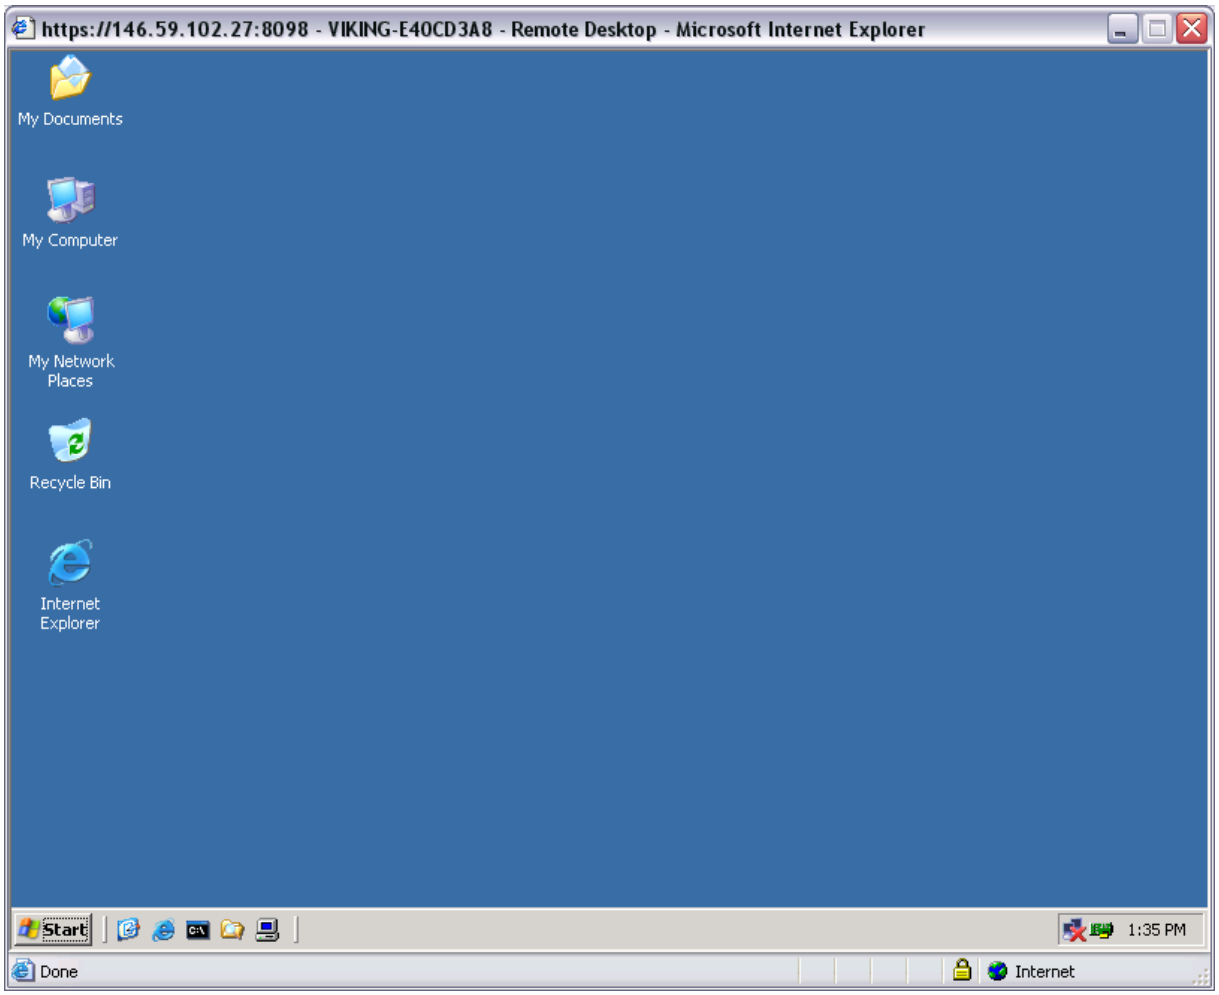

Remote Desktop provides two options when closing a client: you can either disconnect or log off the system.

Disconnecting leaves the session running on the server. You can reconnect to the server and resume the session. If you are performing a task on the server, you can start the task and disconnect from the session. Later, you can log back on the server, re-enter the session and either resume the task or check results. This is especially helpful when operating over a remote connection on a long-distance toll line.

Ending the session is known as logging off. Logging off ends the session running on the server. Any applications running within the session are closed, and unsaved changes made to open files will be lost. The next time you log onto the server, a new session is created.

Remote Desktop requires that all connecting users be authenticated, which is why users must log on each time they start a session.

2. Setting Up & Administration

Basic system administration functions are discussed in this chapter.

This chapter also continues the process of setting up the system that was started using the Viking Series NAS Hardware User Manual by discussing additional setup procedures and options.

Unless otherwise instructed, all procedures are performed using the Windows Remote Desktop Interface.

Notes
The NAS appliance desktop can be accessed via a directly connected keyboard, mouse, and monitor.

Unlike the Windows Storage Server 2003, Windows Storage Server 2003 R2 uses as new Windows Storage Server Management Console that configures and manages shares, storage, quotas, various file sharing protocol as well as print service. The Storage Server Management Console can only be accessed using direct attached method or Windows Remote Desktop.

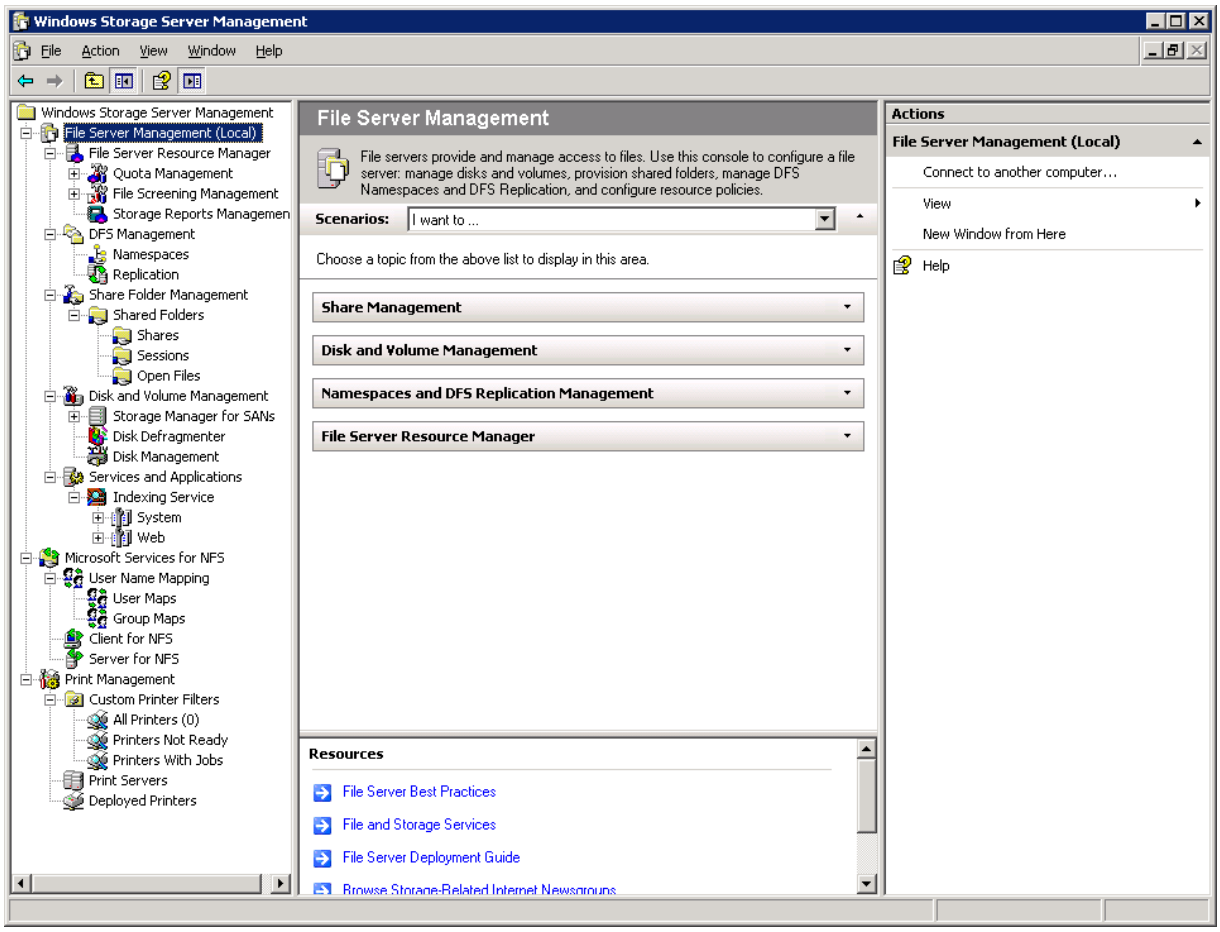

2.1. Using Windows Remote Desktop

Remote Desktop is provided to allow for additional remote system administration and installation of approved third-party applications. Backup software and antivirus programs are examples of approved applications.

To open a Remote Desktop session from a connecting machine using Windows XP:

- Select **Start** > **Run**. At Open, type `mstsc`, then click **OK**.
- Type the IP address of the NAS appliance you will be connecting to in the **Computer** box and click **Connect**.
- Use the default user name **"Administrator"** and default password **"1234"** to gain system access.

Notes

Two open sessions of Remote Desktop are allowed to operate at the same time. After completing an application do not use the window close feature (X) to close that session of Remote Desktop. Click **Start** then **Log Off** Administrator to exit Remote Desktop.

2.1.1. Improper Closure of Remote Desktop

Certain operations can leave the utilities running if the browser is closed versus exiting from the program via the application menu or logging off the Remote Desktop session. A maximum of two Remote Desktop sessions may be used at any given time. Improper exit from a session can result in the sessions becoming consumed. Sessions and processes can be terminated using the Terminal Services Manager via **Start** > **Programs** > **Administrator Tools**.

Note

The Terminal Services Manager must be accessed via the direct attached method.

2.2. Telnet Server

Telnet Server is a utility that lets users connect to machines, log on, and obtain a command prompt remotely. Telnet Server is preinstalled on the storage server, but must be activated before use.

Notes

For security reasons, the Telnet Server is disabled by default. The service needs to be modified to enable access to the storage server with Telnet.

2.2.1. Enabling Telnet Server

The Telnet Server service needs to be enabled prior to its access. The service can be enabled by opening the services MMC:

1. Select **Start** > **Run**, and then enter `services.msc`.

2. Locate and right-click the Telnet service and then select Properties.

3. Choose one of the following:

   - For the Telnet service to start up automatically on every reboot, in the Startup Type drop-down box, click **Automatic**, and then click **OK**.

   - For the Telnet service to be started manually on every reboot, in the Startup Type drop-down box, click **Manual**, and then click **OK**.

On the storage server, access the command line interface, either by Remote Desktop or a direct connection, and then enter the following command:

```
net start tlntsvr
```

The sessions screen provides the ability to view or terminate active sessions.

2.3. Setup Completion

After the NAS appliance is physically set up and the basic configuration is established, additional setup steps must be completed. Depending on the deployment scenario of the NAS appliance, these steps may vary.

Additional setup steps may include:

- Managing system storage
- Creating and managing users and groups
- Creating and managing file shares
- Running Microsoft Windows Update
- Using Ethernet NIC teaming
- Installing third-party software applications

Each of these setup steps is discussed in the following sections.

2.4. Managing System Storage

The NAS administrator uses Disk Management to manage volumes, and Shadow Copies to manage snapshots. See the following chapters for more detailed information on managing system storage:

- Chapter 3 discusses storage and file server management procedures
- Chapter 4 discusses snapshot (shadow copy) management procedures
- Chapter 6 discusses folder and share management procedures

2.5. Creating and Managing Users and Groups

User and group information and permissions determine whether a user can access files. If the NAS appliance is deployed into a workgroup environment, this user and group information is stored locally on the device. By contrast, if the NAS device is deployed into a domain environment, user and group information is stored on the domain usually the Primary Domain Controller.

To enter local user and group information, see Chapter 5.

2.2.1. Joining Workgroup & Domain

These are the two system environments for users and groups. Because users and groups in a domain environment are managed through standard Windows or Active Directory domain administration methods, this document discusses only local users and groups, which are stored and managed on the storage server. For information on managing users and groups on a domain, see the domain documentation available on the Microsoft web site.

2.6. Creating and Managing File Shares

Files shares must be set up, granting and controlling file access to users and groups. See Chapter 6 for complete information on managing file shares.
UNIX specific information is discussed in the “Microsoft Services for NFS” chapter.

2.7. Running Microsoft Windows Update

Tandberg Data highly recommends that you run Microsoft Windows updates to identify, review, and install the latest, applicable, critical security updates on the storage server. For recommendations, instructions, and documentation to help manage the software update, hotfix, and security patches process on the storage server, see documentation of Microsoft Software Updates available on Microsoft web site.

2.8. Using Ethernet NIC Teaming

All Viking Series NAS appliances are equipped with an Intel or Broadcom NIC Teaming utility. The utility allows administrators to configure and monitor Ethernet network interface controller (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput.

2.9. Installing 3rd Party Software Applications

For example, these might include an antivirus or backup application that you install.

This Page Intentionally Left Blank

3. Storage & File Server Management

This chapter provides an overview of some of the component that make up the storage structure of the Viking Series NAS appliance.

3.1. Storage Management Element

Storage is divided into four major divisions:

- Physical storage elements
- Logical storage elements
- File system elements
- File sharing elements

Each of these elements is composed of the previous level's elements.

3.1.1. Physical Storage Element

The lowest level of storage management occurs at the physical drive level. Minimally, choosing the best disk carving strategy includes the following policies:

- Analyze current corporate and departmental structure.
- Analyze the current file server structure and environment.
- Plan properly to ensure the best configuration and use of storage.
    - Determine the desired priority of fault tolerance, performance, and storage capacity.
    - Use the determined priority of system characteristics to determine the optimal striping policy and RAID level.
- Include the appropriate number of physical drives in the arrays to create logical storage elements of desired sizes.

3.1.1.1. Disk Array

With an array controller installed in the system, the capacity of several physical drives can be logically combined into one or more logical units called arrays. When this is done, the read/write heads of all the constituent physical drives are active simultaneously, dramatically reducing the overall time required for data transfer.

> Notes
>
> Depending on the storage server model, array configuration may not be possible or necessary.

Because the read/write heads are simultaneously active, the same amount of data is written to each drive during any given time interval. Each unit of data is termed a block. The blocks form a set of data stripes over all the hard drives in an array.

For data in the array to be readable, the data block sequence within each stripe must be the same. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each hard drive in a given array contains the same number of data blocks.

> Notes
>
> If one hard drive has a larger capacity than other hard drives in the same array, the extra capacity is wasted because it cannot be used by the array.

3.1.1.2. Fault Tolerance

Drive failure, although rare, is potentially catastrophic. For example, using simple striping with several HDD, failure of any hard drive leads to failure of all logical drives in the same array, and hence to data loss.

To protect against data loss from hard drive failure, storage servers should be configured with fault tolerance.

3.1.1.3. Online Spares (Hot-Spares)

Further protection against data loss can be achieved by assigning an online spare (or hot-spare) to any configuration except RAID 0. This hard drive contains no data and is contained within the same storage subsystem as the other drives in the array. When a hard drive in the array fails, the controller can then automatically rebuild information that was originally on the failed drive onto the online spare. This quickly restores the system to full RAID level fault tolerance protection. However, unless RAID 6 is being used, which can support two drive failures in an array, in the unlikely event that a third drive in the array should fail while data is being rewritten to the spare, the logical drive still fails.

3.1.2. Logical Storage Element

Logical storage elements consist of those components that translate the physical storage elements to file system elements. The storage server uses the Window Disk Management utility to manage the various types of disks presented to the file system. There are two types of LUN presentation: basic disk and dynamic disk. Each of these types of disk has special features that enable different types of management.

3.1.2.1. Logical Drives (LUNs)

While an array is a physical grouping of hard drives, a logical drive consists of components that translate physical storage elements into file system elements. It is important to note that a LUN may extend over (span) all physical drives within a storage controller subsystem, but cannot span multiple storage controller subsystems.

Through the use of basic disks, primary partitions or extended partitions can be created. Partitions can only encompass one LUN. Through the use of dynamic disks, volumes can be created that span multiple LUNs. The Windows Disk Management utility can be used to convert disks to dynamic and back to basic, and manage the volumes residing on dynamic disks. Other options include the ability to delete, extend, mirror, and repair these elements.

3.1.2.2. Partitions

Partitions exist as either primary partitions or extended partitions and can be composed of only one basic disk no larger than 2 TB. Basic disks can also only contain up to four primary partitions, or three primary partitions and one extended partition. In addition, the partitions on them cannot be extended beyond the limits of a single LUN. Extended partitions allow the user to create multiple logical drives. These partitions or logical disks can be assigned drive letters or be used as mount points on existing disks. If mount points are used, it should be noted that Services for UNIX (SFU) does not support mount points at this time. The use of mount points in conjunction with NFS shares is not supported.

3.1.2.3. Volumes

When planning dynamic disks and volumes, there is a limit to the amount of growth a single volume can undergo. Volumes are limited in size and can have no more than 32 separate LUNs, with each LUN not exceeding 2 terabytes (TB), and volumes totaling no more than 64 TB of disk space.

The RAID level of the LUNs included in a volume must be considered. All of the units that make up a volume should have the same high-availability characteristics. In other words, the units should all be of the same RAID level. For example, it would not be a good practice to include both a RAID 0+1 and a RAID 5 array in the same volume set. By keeping all the units the same, the entire volume retains the same performance and high-availability characteristics, making managing and maintaining the volume much easier. If a dynamic disk goes offline, the entire volume dependent on the one or more dynamic disks is unavailable. There could be a potential for data loss depending on the nature of the failed LUN.

Volumes are created out of the dynamic disks, and can be expanded on the fly to extend over multiple dynamic disks if they are spanned volumes. However, after a type of volume is selected, it cannot be altered. For example, a spanning volume cannot be altered to a mirrored volume without deleting and recreating the volume, unless it is a simple volume. Simple volumes can be mirrored or converted to spanned volumes. Fault-tolerant disks cannot be extended either. Therefore, selection of the volume type is important. The same performance characteristics on numbers of reads and writes apply when using fault-tolerant configurations, as is the case with controller-based RAID. These volumes can also be assigned drive letters or be mounted as mount points off existing drive letters.

The administrator should carefully consider how the volumes will be carved up and what groups or applications will be using them. For example, putting several storage-intensive applications or groups into the same dynamic disk set would not be efficient. These applications or groups would be better served by being divided up into separate dynamic disks, which could then grow as their space requirements increased, within the allowable growth limits.

3.1.3. File System Element

File system elements are composed of the folders and subfolders that are created under each logical storage element (partitions, logical disks, and volumes). Folders are used to further subdivide the available file system, providing another level of granularity for management of the information space. Each of these folders can contain separate permissions and share names that can be used for network access. Folders can be created for individual users, groups, projects, and so on.

3.1.4. File Sharing Element

The storage server supports several file sharing protocols, including Distributed File System (DFS), Network File System (NFS), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), and Microsoft Server Message Block (SMB). On each folder or logical storage element, different file sharing protocols can be enabled using specific network names for access across a network to a variety of clients. Permissions can then be granted to those shares based on users or groups of users in each of the file sharing protocols.

3.2. Volume Shadow Copy Service Overview

The Volume Shadow Copy Service (VSS) provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. VSS supports 64 shadow copies per volume.

Shadow Copies of Shared Folders resides within this infrastructure, and helps alleviate data loss by creating shadow copies of files or folders that are stored on network file shares at pre-determined time intervals. In essence, a shadow copy is a previous version of the file or folder at a specific point in time.

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.

Shadow copies should not replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. For example, shadow copies cannot protect against data loss due to media failures; however, recovering data from shadow copies can reduce the number of times needed to restore data from tape.

3.3. Using Storage Element

The last step in creating the element is determining its drive letter or mount point and formatting the element. Each element created can exist as a drive letter, assuming one is available and/or as mount points off of an existing folder or drive letter. Either method is supported. However, mount points cannot be used for shares that will be shared using Microsoft Services for Unix. They can be set up with both but the use of the mount point in conjunction with NFS shares causes instability with the NFS shares.

Formats consist of NTFS, FAT32, and FAT. All three types can be used on the storage server. However, VSS can only use volumes that are NTFS formatted. Also, quota management is possible only on NTFS.

3.4. Network Adapter Teaming

Network adapter teaming is software-based technology used to increase a server's network availability and performance. Teaming enables the logical grouping of physical adapters in the same server (regardless of whether they are embedded devices or Peripheral Component Interconnect (PCI) adapters) into a virtual adapter. This virtual adapter is seen by the network and server-resident network-aware applications as a single network connection.

3.5. Disk Management

Disk Management can be accessed after administrator login using Remote Desktop. Disk Management is accessed by right-clicking on the **My Computer** icon on the Desktop and then selects **Manage** option or using the Windows Storage Server Management Console.

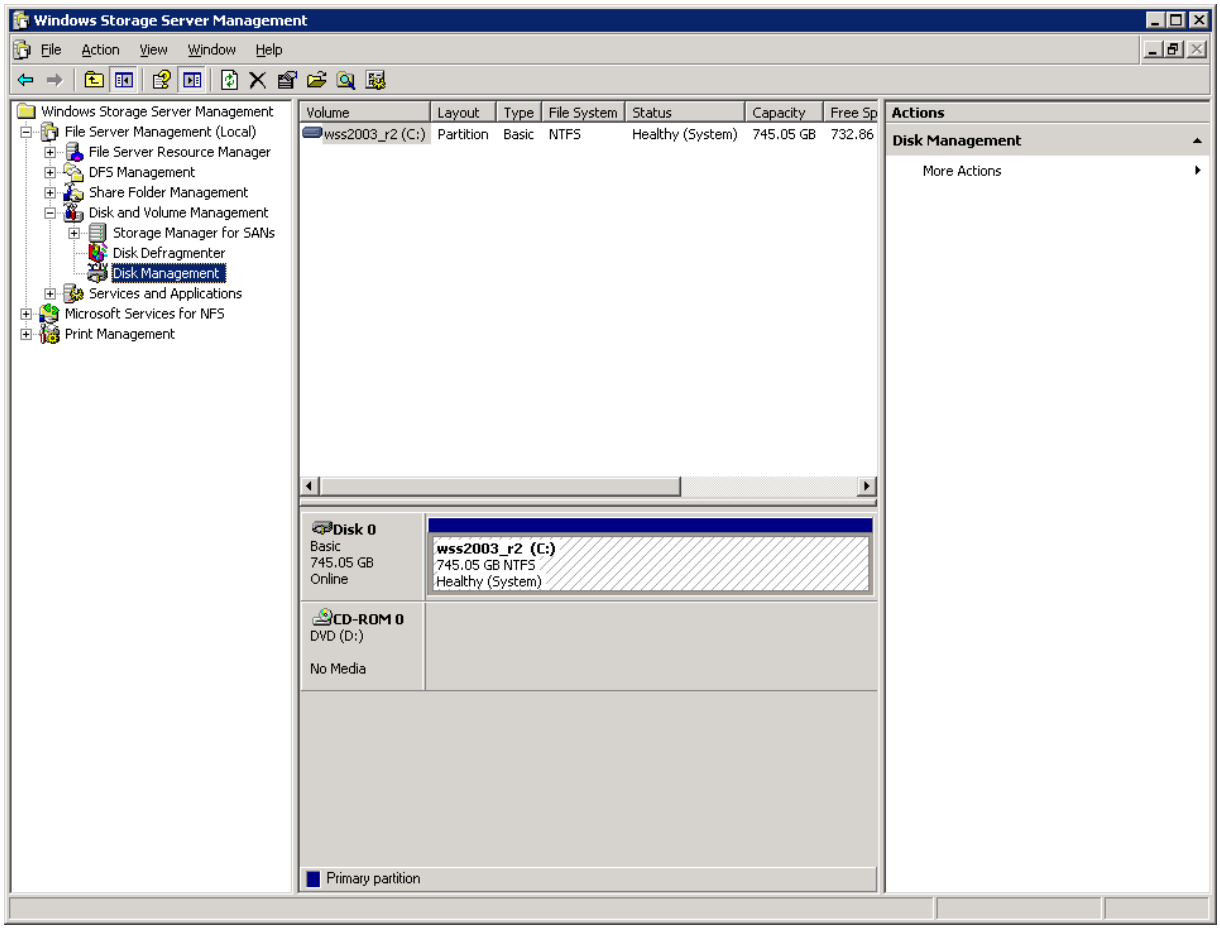

The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions that they contain. Disk Management is used to initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Most disk-related tasks can be preformed in Disk Management without restarting the system or interrupting users; most configuration changes take effect immediately. A complete online help facility is provided with the Disk Management Utility for assistance in using the product

Note

When the Disk Management utility is accessed, the Remote Desktop connection assumes a dedicated mode and can only be used to manage disks and volumes on the server. Navigating to another page during an open session closes the session.

Note

It may take a few moments for the Remote Desktop Connection session to log off when closing Disk Management.

3.6. Guidelines for Managing Disks & Volumes

When managing disks and volumes:

- Read the online Disk Management Help found in the utility.
- Do not alter the Operating System Disk labeled C:. This logical drive is configured for the storage server operating system and should not be altered in any manner, unless return to factory image is desired.
- Tandberg Data does not recommend spanning arrays (volume set) with dynamic volumes. The use of software RAID-based dynamic volumes is not recommended. Use the array controller instead; it is more efficient.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. For example, volume F: might be named “Disk F:.” Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in cases when OS recovery is necessary.
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic without bringing the system offline or loss of data, but the volume will be unavailable during the conversion. However, it cannot be converted back to basic without deleting all data on the disk.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of snapshots, performance, and defragmentation.
- NTFS formatted drives are recommended since they provide the greatest level of support for snapshots, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32.

3.7. RAID & Volume Management

The RAID and Volume set can be managed using web-based RAID manager. It can be accessed using the following on your web browser:

http://<your NAS machine name or IP Address>:81/

The default user name is **"admin"**. The default password is **"0000"**.

The web-based RAID manager is use to administer further configurations i.e. changing the RAID level on the System Volume or the Data Volume, and monitor controllers as well as the RAID re-building progress.

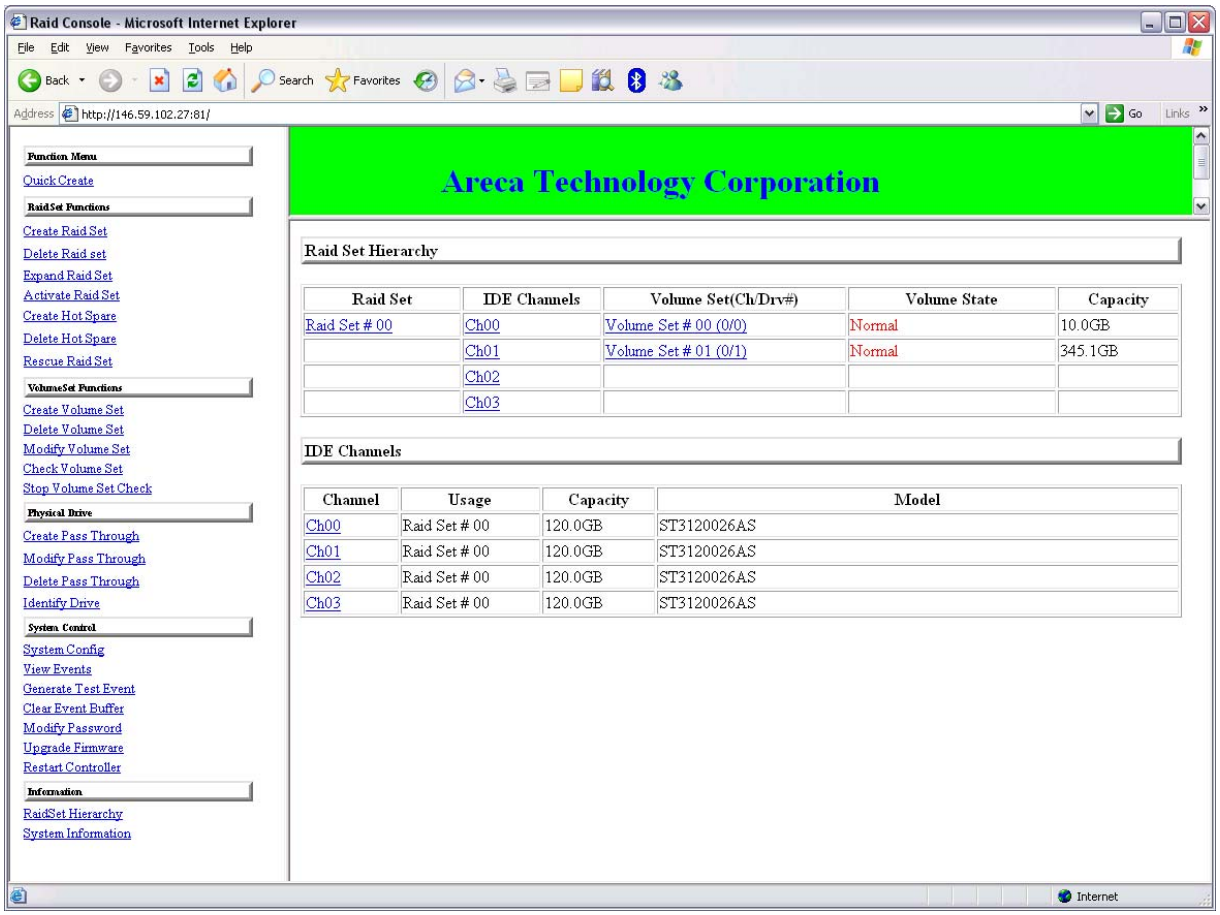

The default password can be changed using the **Modify Password** field located on the left column strip of the web-based RAID manager.

Prior to using the web-based RAID manager, ensure that the web service of RAID manager is started by accessing the Remote Desktop using **Start** > **Programs** > **Areca Technology Corp** > **archttppci** > **archttppci**.

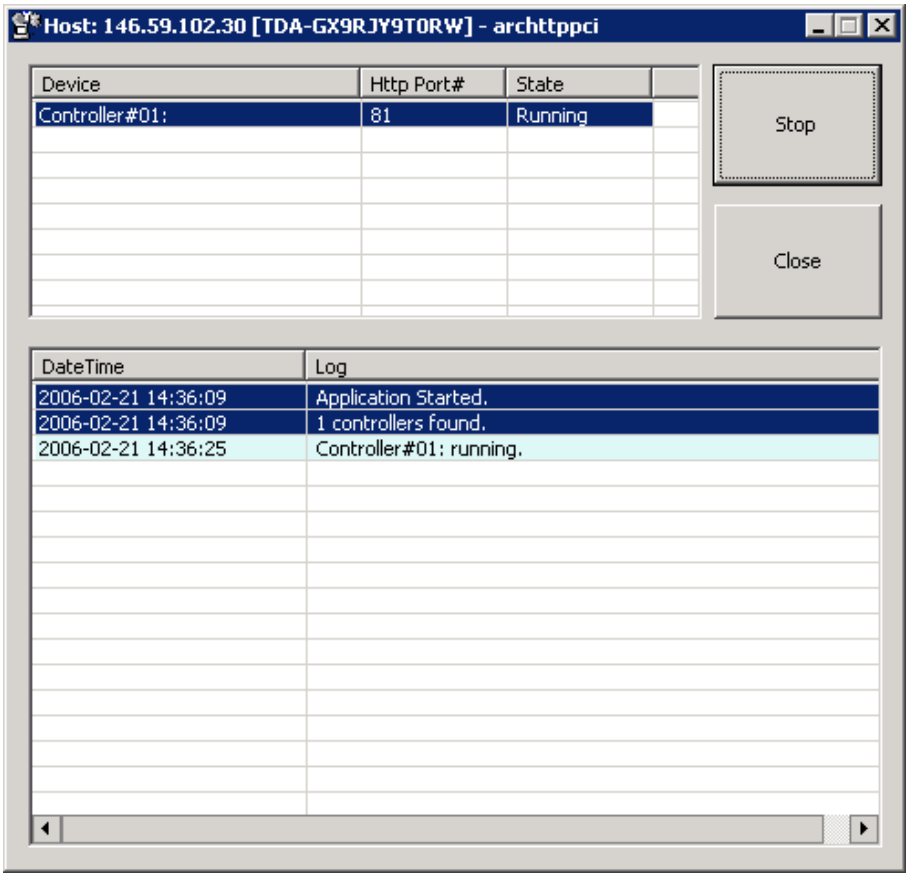

Closing the **archttppci** application during a web session with the array controller will result in lost of unsaved information.

3.8. Scheduling Defragmentation

Defragmentation is the process of analyzing local volumes and consolidating fragmented files and folders so that each occupies a single, contiguous space on the volume. This improves file system performance. Because defragmentation consolidates files and folders, it also consolidates the free space on a volume. This reduces the likelihood that new files will be fragmented.

Defragmentation for a volume can be scheduled to occur automatically at convenient times. Defragmentation can also be done once, or on a recurring basis.

Notes

Scheduling defragmentation to run no later than a specific time prevents the defragmentation process from running later than that time. If the defragmentation process is running when the time is reached, the process is stopped. This setting is useful to ensure that the defragmentation process ends before the demand for server access is likely to increase.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger during the format. Otherwise defragmentation registers as a change by the Shadow Copy process. This increase in the number of changes forces Shadow Copy to delete snapshots as the limit for the cache file is reached.

Warning

Allocation unit size cannot be altered without reformatting the drive. Data on a reformatted drive cannot be recovered.

Note

NTFS compression is supported only if the cluster size is 4 KB or smaller.

3.9. Disk Quotas

Disk quotas track and control disk space use in volumes.

Note

To limit the size of a folder or share, see “Directory Quotas” in Chapter 6.

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

Note

When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

3.10. Using Diskpart

Diskpart.exe is a text-mode command interpreter that enables the administrator to manage disks, partitions, or volumes.

When using the list commands, an asterisk (*) appears next to the object with focus. Select an object by its number or drive letter, such as disk 0, partition 1, volume 3, or volume C.

When selecting an object, the focus remains on that object until a different object is selected. For example, if the focus is set on disk 0 and volume 8 on disk 2 is selected, the focus shifts from disk 0 to disk 2, volume 8. Some commands automatically change the focus. For example, when creating a new partition, the focus automatically switches to the new partition.

Focus can only be given to a partition on the selected disk. When a partition has focus, the related volume (if any) also has focus. When a volume has focus, the related disk and partition also have focus if the volume maps to a single specific partition. If this is not the case, focus on the disk and partition is lost.

Some of the common Diskpart commands are:

- `add disk`
  Mirrors the simple volume with focus to the specified disk.

- `assign`
  Assigns a drive letter or mount point to the volume with focus.

- `convert basic`
  Converts an empty dynamic disk to a basic disk.

- `convert dynamic`
  Converts a basic disk into a dynamic disk. Any existing partitions on the disk become simple volumes.

- `create volume simple`
  Creates a simple volume. After creating the volume, the focus automatically shifts to the new volume.

- `exit`
  Exits the DiskPart command interpreter.

- `help`
  Displays a list of the available commands.

- `list disk`
  Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table. The disk marked with an asterisk (*) has focus.

- `list partition`
  Displays the partitions listed in the partition table of the current disk. On dynamic disks these partitions may not correspond to the dynamic volumes on the disk. This discrepancy occurs because dynamic disks contain entries in the partition table for the system volume or boot volume (if present on the disk). They also contain a partition that occupies the remainder of the disk in order to reserve the space for use by dynamic volumes.

- `list volume`
  Displays a list of basic and dynamic volumes on all disks.

- `rem`
  Provides a way to add comments to a script.

- `retain`
  Prepares an existing dynamic simple volume to be used as a boot or system volume.

- `select disk`
  Selects the specified disk and shifts the focus to it.

For a complete list of Diskpart commands, go to the Windows Storage Server 2003 R2 Desktop on the NAS appliance via Remote Desktop and select Start > Help and Support, search on Diskpart.

3.10.1. Example of using Diskpart

The following example shows how to configure a volume on the NAS server. In the cmd window, type:

```
c:\>diskpart
DISKPART>rescan
DISKPART>select disk 1
DISKPART>convert dynamic
DISKPART>rem create a simple volume
DISKPART>create volume simple size=4000
DISKPART>rem assign drive letter F: to the volume
DISKPART>assign letter=F
DISKPART>list vol
DISKPART>exit
```

3.11. Adding Storage

Expansion is the process of adding physical disks to an array that has already been configured. Extension is the process of adding new storage space to an existing logical drive on the same array, usually after the array has been expanded.

Storage growth may occur in three forms:

- Extend unallocated space from the original logical disks or LUNs.
- Alter LUNs to contain additional storage.
- Add new LUNs to the system.

The additional space is then extended through a variety of means, depending on which type of disk structure is in use.

NOTE:

This section addresses only single storage node configuration. If your server has Windows Storage Server 2003 R2 Enterprise Edition, see the Cluster Administration chapter for expanding and extending storage in a cluster environment.

3.11.1. Expanding Storage

Expansion is the process of adding physical disks to an array that has already been configured. The logical drives (or volumes) that exist in the array before the expansion takes place are unchanged, because only the amount of free space in the array changes. The expansion process is entirely
independent of the operating system.

Notes

See your storage array hardware user documentation for further details about expanding storage on the array.

3.11.2. Extending Storage using Disk Management

The Disk Management snap-in provides management of hard disks, volumes or partitions. It can be used to extend a dynamic volume only.

Notes
Disk Management cannot be used to extend basic disk partitions.

Guidelines for extending a dynamic volume:

- Use the Disk Management utility.
- You can extend a volume only if it does not have a file system or if it is formatted NTFS.
- You cannot extend volumes formatted using FAT or FAT32.
- You cannot extend striped volumes, mirrored volumes, or RAID 5 volumes.

For more information, see the Disk Management online help.

3.12. File Services Features in Windows Storage Server 2003 R2

This section begins by identifying file services in Windows Storage Server 2003 R2. The remainder of it describes the many tasks and utilities that play a role in file server management.

3.12.1. Storage Manager for SANs

The Storage Manager for SANs (also called Simple SAN) snap-in enables you to create and manage the LUNs that are used to allocate space on storage arrays. Storage Manager for SANs can be used on SANs that support Virtual Disk Server (VDS). It can be used in both Fibre Channel and iSCSI environments.

For more information on Storage Manager for SANs, see the online help. A Microsoft document titled Storage Management in Windows Storage Server 2003 R2: File Server Resource Manager and Storage Manager for Storage Area Networks is available at:

http://download.microsoft.com/download/7/4/7/7472bf9b-3023-48b7-87be-d2cedc38f15a/WS03R2_Storage_Management.doc .

Notes
Storage Manager for SANs is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

3.12.2. 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.

Notes

Single Instance Storage is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

3.12.3. File Server Resource Manage

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 to generate notifications when the quota limits are approached and exceeded.

- Create file screens to screen the files that users can save on volumes and in folders and to send notifications when users attempt to save blocked files

- Schedule periodic storage reports that allow users to identify trends in disk usage and to monitor attempts to save unauthorized files, or generate the reports on demand.

3.12.4. Windows SharePoint Services

Windows SharePoint 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.

NOTE:
Windows SharePoint Services is only available on Standard and Enterprise editions of Windows Storage Server 2003 R2.

3.12.5. Windows Storage Server Management Console

The Windows Storage Server Management Console is a user interface in Windows Storage Server 2003 R2 and Windows Unified Data Storage Server 2003 that provides one place to manage files or print serving components. The console is accessible using Remote Desktop or direct attached method.

The Storage Management page provides a portal to:

- File Server Resource Manager
- DFS Management
- Disk and Volume Management
- Single Instance Storage
- Indexing Service
- MSNFS (under Share folder)

The Share Folder Management page provides a portal to Shared Folders, consisting of:

- Shares
- Sessions
- Open Files

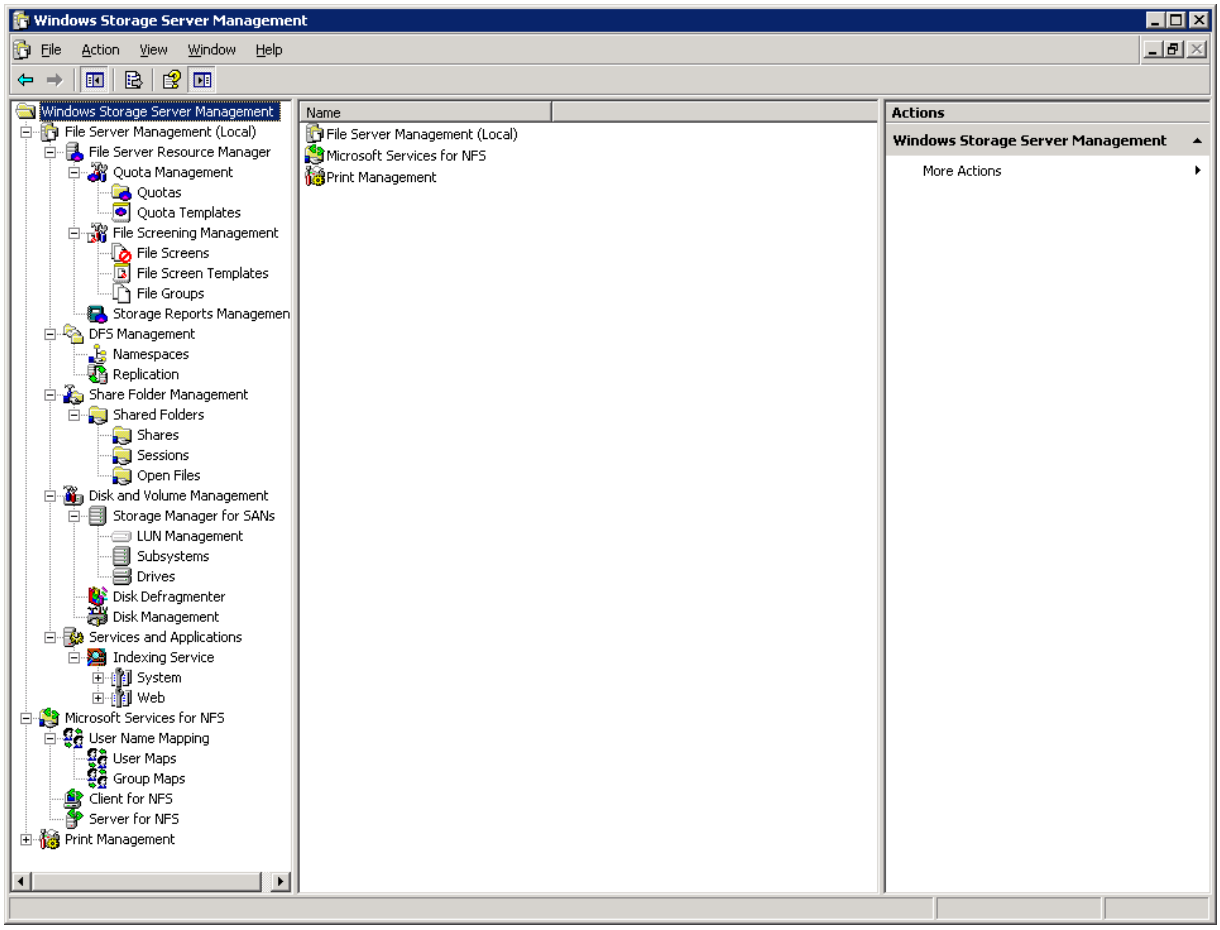
Windows Storage Server Management
File Action View Window Help
Windows Storage Server Management
File Server Management (Local)
File Server Resource Manager
Quota Management
Quotas
Quota Templates
File Screening Management
File Screens
File Screen Templates
File Groups
Storage Reports Managemen
DFS Management
Namespaces
Replication
Share Folder Management
Shared Folders
Shares
Sessions
Open Files
Disk and Volume Management
Storage Manager for SANs
LUN Management
Subsystems
Drives
Disk Defragmenter
Disk Management
Services and Applications
Indexing Service
System
Web
Microsoft Services for NFS
User Name Mapping
User Maps
Group Maps
Client for NFS
Server for NFS
Print Management
Name
File Server Management (Local)
Microsoft Services for NFS
Print Management
Actions
Windows Storage Server Management
More Actions

This Page Intentionally Left Blank

4. Volume Shadow Copy Service(VSS)

4.1. Volume Shadow Copy Service Overview

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume under Windows® Storage Server 2003 environment.

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 Viking Series NAS appliance (see the “Managing Shadow Copy” section in this chapter), previous versions of files and folders are only available over the network from clients and are seen on a per folder or file level and not as an entire volume.

The Shadow Copy feature works at the block level. As changes are made to the file system, the Shadow Copy Service copies out the original blocks to a special cache file, to maintain a consistent view of the file at a particular point in time. Since 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 since blocks are not moved until an update to the disk occurs.

By using shadow copies, the NAS appliance 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. Since 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.

4.2. Planning for use of Shadow Copy

Before Shadow Copy Service is initiated on the NAS appliance and the client interface is made available to end users, consider the followings:

- 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?

4.2.1. 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.

> Note
>
> Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

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 (the NAS appliance) must also specify the \\nasname\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.

> Note
>
> Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

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

4.2.2. Allocating Disk Space

When shadow copies are enabled on a volume, the maximum amount of volume space to be used for the shadow copies can be specified. The default limit is 10 percent of the source volume (the volume being copied). The limit for volumes in which users frequently change files should be increased. Also, note that setting the limit too low causes the oldest shadow copies to be deleted frequently, which defeats the purpose of shadow copies and frustrates users.

If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, then no shadow copy is created. Therefore, administrators should carefully consider the amount of disk space they want to set aside for shadow copies, and keep in mind 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 and other backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

> Note
>
> Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 100 megabytes (MB). The default storage size is 10% 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.

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.

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

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 to can be converted to dynamic without losing shadow copies.

Note

Use the `mountvol` command with the `/p` option to dismount the volume and take it offline. Mount the volume and bring it online using the `mountvol` command or the Disk Management snap-in.

4.2.3. Identifying the Storage Area

To store the shadow copies of another volume on the same NAS appliance, 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 NAS appliance.

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, however, a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

For example, by keeping the shadow copy on the same volume, although there is a potential gain in ease of setup and maintenance, there may be a reduction in performance and reliability.

Caution

If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

4.2.4. 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 NAS appliance will create shadow copies at 0700 and 1200, Monday through Friday when the feature is enabled for a volume. 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 the section on “Shadow Copy Schedules” documented later in this chapter.

Note

The more shadow copies are created, the more disk space the shadow copies can consume, especially if files change frequently.

4.2.4. Shadow Copies & Disk Defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Utilizing this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

Note

To check the cluster size of a volume, use the following text-mode command:

```
C:\>fsutil fsinfo ntfsinfo <volume pathname or drive letter>
```

To change the cluster size on a volume that contains data, backup the data on the volume, reformat it using the new cluster size, and then restore the data.

4.2.5. Mounted Drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder F:\data\users, and the Users folder is a mount point for G:\. If shadow copies are enabled on both F:\ and G:\, F:\data is shared as \\nas1\data, and G:\data\users is shared as \\nas1\users. In this example, users can access previous versions of \\nas1\data and \\nas1\users but not \\nas1\data\users.

4.3. Managing Shadow Copies

The vssadmin tool provides a command line capability to create, list, resize, and delete volume shadow copies.

The system administrator can make shadow copies available to end users through a feature called “Shadow Copies for Shared Folders.” The administrator uses the Properties menu (see below) to turn on the Shadow Copies feature, select the volumes to be copied, and determine the frequency with which shadow copies are made.

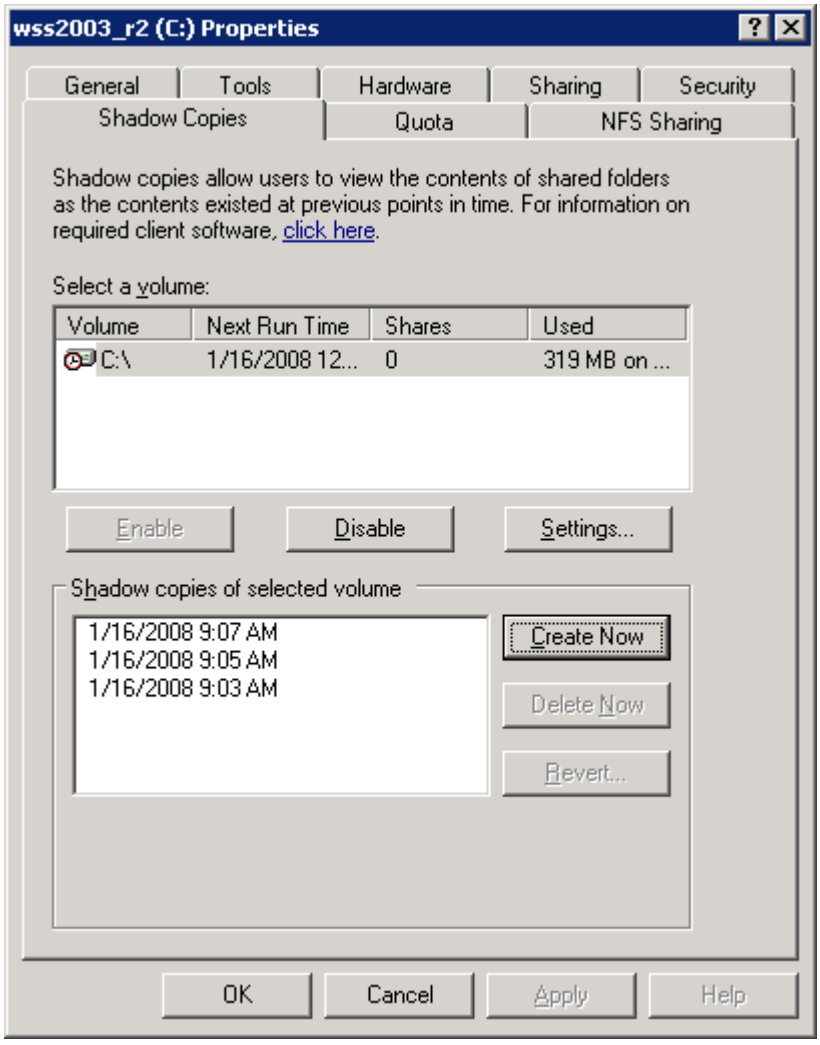
wss2003_r2 (C:) Properties
General
Tools
Hardware
Sharing
Security
Shadow Copies
Quota
NFS Sharing
Shadow copies allow users to view the contents of shared folders as the contents existed at previous points in time. For information on required client software, click here.
Select a volume:
Volume
Next Run Time
Shares
Used
C:\
1/16/2008 12...
0
319 MB on ...
Enable
Disable
Settings...
Shadow copies of selected volume
1/16/2008 9:07 AM
1/16/2008 9:05 AM
1/16/2008 9:03 AM
Create Now
Delete Now
Revert...
OK
Cancel
Apply
Help

4.3.1. Shadow Copy Cache File

The default shadow copy settings allocate 10% of the source volume being copied (with a minimum of 350 MB), and store the shadow copies on the same volume as the original volume. See figure below. The cache file is located in a hidden protected directory entitled “System Volume Information” off of the root of each volume for which Shadow Copy is enabled.

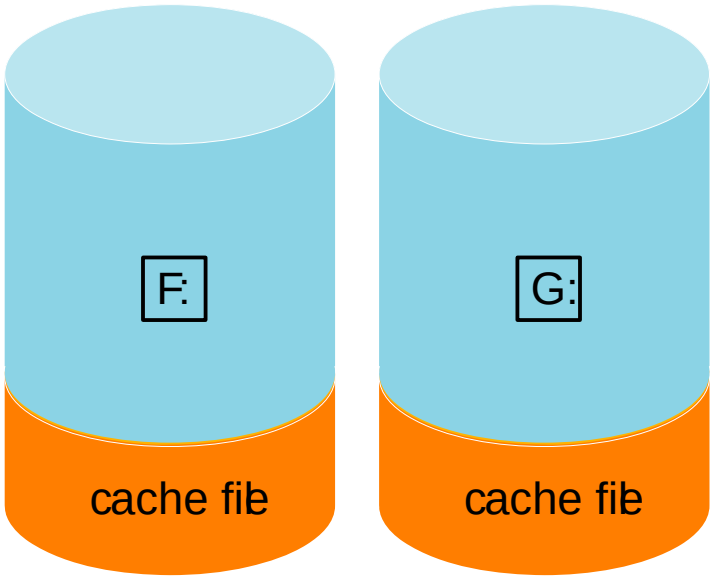

As mentioned previously, the cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. See figure below.

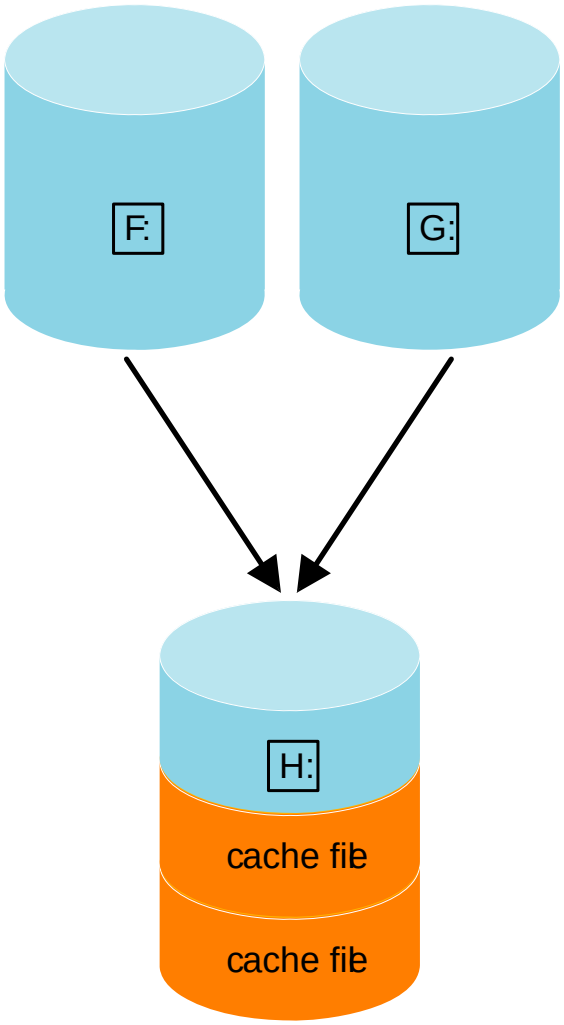

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space for Shadow Copies may be manage separately, limits can generally be set higher, or set to No Limit. See the properties tab of the shadow copy page for a volume to alter the cache file location, covered later in this chapter.

Caution

If the data on the separate volume H: is lost, the shadow copies cannot be recovered.

4.4. Enabling and Creating Shadow Copies

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume
- Sets the maximum storage space for the shadow copies
- Schedules shadow copies to be made at 7 A.M. and 12 noon on weekdays.

> Note
>
> Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

To enable shadow copies on a volume:

1. Access Disk Management.
2. Select the volume of the logical drive which you want to enable shadow copy service and right-click on it.
3. Select **Properties**.
4. Select **Shadow Copies** tab.

> Note
>
> After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See “Viewing Shadow Copy Properties” in this chapter.

5. Click **Enable**.

To create a shadow copy on a volume:

1. On **Shadow Copies** tab, click on the **Create Now** button.

4.5. Viewing List of Shadow Copies

The list of Shadow Copies can be view on a volume using the following method:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.
2. Click the **Shadow Copies** tab.

All shadow copies are listed, sorted by the date and time they were created.

> Note
>
> It is also possible to create new shadow copies or delete shadow copies from this page.

4.6. Scheduling Shadow Copies

Shadow Copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow-copy schedule to allow for these differences.

It is recommended that shadow copies be scheduled not more frequently than once per hour.

> Notes
>
> When deleting a shadow copy schedule, that action has no effect on existing shadow copies (which have taken).

4.6.1. Setting Shadow Copy Schedule

When the Shadow Copies service is enabled on a volume, it automatically schedules shadow copies to be made each weekday at 7 A.M. and 12 noon.

To add or change a shadow copy schedule for a volume:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.
2. Click the **Shadow Copies** tab.
3. Click on the **Setting…** button.
4. Click on the **Schedule…** button.
5. On the Shadow Copy Schedules page, click **New**.
6. Select a frequency: **Once**, **Daily**, **Weekly**, or **Monthly**.

7. Use the **Advance…** controls to specify the recurrence pattern and the starting date and time. The available controls change according to the frequency selected.

8. Click **OK**.

4.6.2. Deleting Shadow Copy Schedule

To delete a shadow copy schedule on a volume:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.

2. Click the **Shadow Copies** tab.

3. Click on the **Setting…** button.

4. Click on the **Schedule…** button.

5. From the Schedule drop down box, select the schedule to be deleted, and click **Delete**.

Note
When deleting a shadow copy schedule, that action has no effect on existing shadow copies.

4.6.3. Viewing Properties of Shadow Copies

To view shadow copy properties on a volume:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.

2. Click the **Shadow Copies** tab.

3. The Shadow Copy Properties screen lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

4. Click on the **Setting…** button.

   The maximum size limit for all shadow copies cab be changed by defining a new cache size in the box, or choose **No limit**.

   For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. See “The Shadow Copy Cache File” earlier in this chapter. The list of available disks and the space available on each is presented at the bottom of the page. Managing the cache files on a separate disk is recommended.

Note

If shadow copies have already been enabled, the cache file location is grayed out. To change this location after shadow copies have been enabled, all shadow copies must be deleted and cannot be recovered. Remember enabling Shadow Copies creates a Shadow Copy by default.

5. Click **OK** to save changes, or click **Cancel** to discard changes.

Caution

Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted

4.6.4. Redirecting Shadow Copies to Alternate Volume

IMPORTANT

Shadow copies must be initially disabled on the volume before redirecting to an alternate volume. If shadow copies are enabled and you disable them, a message appears informing you that all existing shadow copies on the volume will be permanently deleted.

To redirect shadow copies to an alternate volume:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.
2. Click the **Shadow Copies** tab.
3. Select the volume that you want to redirect shadow copies from and ensure that shadow copies are disabled on that volume; if enabled, click **Disable**.
4. Click on the **Setting…** button.
5. In the **Located on this volume** field, select an available alternate volume from the list.

Notes

To change the default shadow copy schedule settings, click Schedule.

6. Click **OK**.
7. On the **Shadow Copies** tab, ensure that the volume is selected, and then click **Enable**.

Shadow copies are now scheduled to be made on the alternate volume.

4.6.4. Disabling Shadow Copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

To disable shadow copies on a volume:

1. Right-clicked on the volume or logical drive with shadow copy service enabled and select **Properties**.
2. Click the **Shadow Copies** tab.
3. Click **Disable**.
4. Click **Yes** to confirm.

Caution
When the Shadow Copies service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

4.7. 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. For 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.

Note

Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.

Note

Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files on behalf of these users.

4.7.1. SMB Access to 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 dialog, selecting 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.

The Shadow Copies of Shared Folders client pack installs a **Previous Versions** tab in the **Properties** dialog box 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. See figure below. Both individual files and folders may be restored.

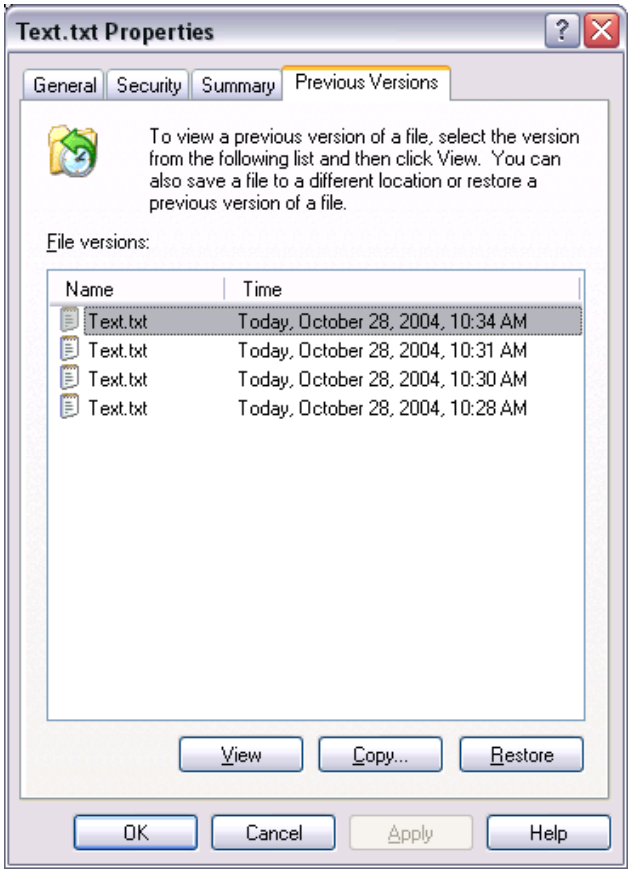

When users view a network folder hosted on the NAS appliance for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history-a list of read-only, point-in-time copies 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.

4.7.2. NFS Access to 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. Note that, 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 April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare
@ GMT-2003.04.27-04:00:00
@ GMT-2003.04.28-04:00:00
@ GMT-2003.04.29-04:00:00

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.

4.8. Recovery of Files and 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.

As documented previously, the use of the snapshots are from the network and are based on shares created on the NAS appliance.

4.8.1. Recovering Deleted Files or Folders

To recover a deleted file or folder within a folder:

1. Navigate to the folder where the deleted file was stored (on the NAS appliance).
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file will be selected.
3. Right-click the mouse and select **Properties** from the bottom of the menu. Select 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. Select **Restore** to restore the file or folder to its original location. Selecting **Copy…** will allow the placement of the file or folder to a new location.

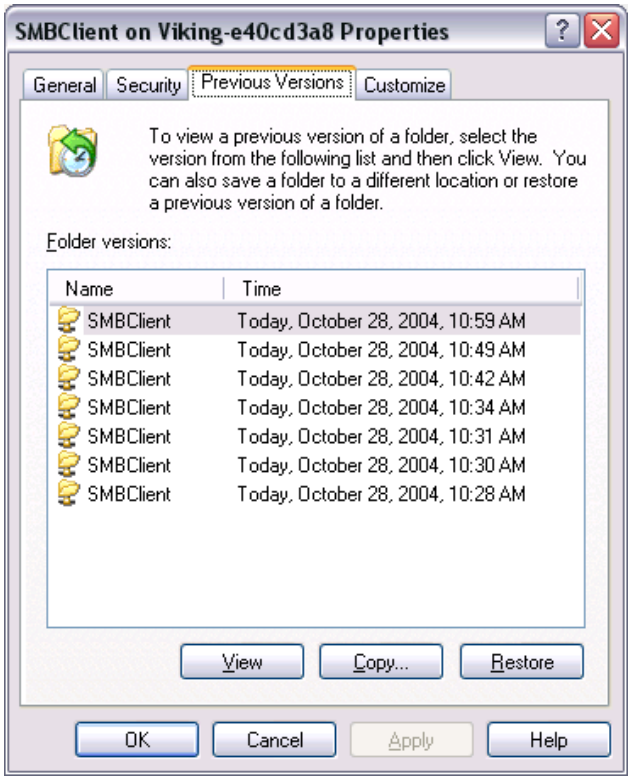

4.8.2. Recovering Overwritten or Corrupted Files

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 use the following procedure:

1. Right-click the overwritten or corrupted file and click **Properties**.
2. Select **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

4.8.3. Recovering Folders

To recover a folder, use the following procedure:

1. Position the cursor so that it is over a blank space in the folder that will be recovered. If the cursor hovers over a file, that file will be selected.

2. Right-click the mouse, select **Properties** from the bottom of the menu, then click the **Previous Versions** tab.

3. Choose either **Copy** or **Restore**.

4. Choosing **Restore** enables the user to recover everything in that folder as well as all subfolders. Selecting **Restore** will not delete any files.

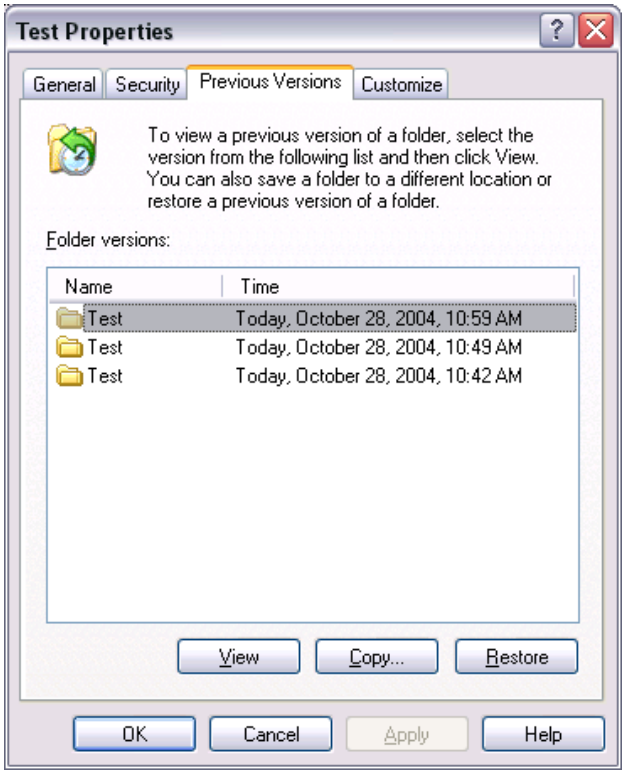

4.9. Backup & Shadow Copies

As mentioned previously, 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 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 backup. The user is unaware of this activity and it is not self evident although it does address the issue of open files.

4.10. Shadow Copy Transport

Shadow Copy Transport provides the ability to transport data on a Storage Area Network (SAN). With a storage array and a VSS-aware hardware provider, it is possible to create a shadow copy on one server and import it on another server. This process, essentially “virtual” transport, is accomplished in a matter of minutes, regardless of the size of the data.

Notes

Shadow copy transport is supported only on Windows Server 2003 Enterprise Edition, Windows Storage Server 2003 Enterprise Edition, and Windows Server 2003 Datacenter Edition. It is an advanced solution that works only if it has a hardware provider on the storage array.

A shadow copy transport can be used for a number of purposes, including:

- Tape backups
  An alternative to traditional backup to tape processes is transport of shadow copies from the production server onto a backup server, where they can then be backed up to tape. Like the other two alternatives, this option removes backup traffic from the production server. While some backup applications might be designed with the hardware provider software that enables transport, others are not. The administrator should determine whether or not this functionality is included in the backup application.

- Data mining
  The data in use by a particular production server is often useful to different groups or departments within an organization. Rather than add additional traffic to the production server, a shadow copy of the data can be made available through transport to another server. The shadow copy can then be processed for different purposes, without any performance impact on the original server.

The transport process is accomplished through a series of DISKRAID command steps:

1. Create a shadow copy of the source data on the source server (read-only).
2. Mask off (hide) the shadow copy from the source server.
3. Unmask the shadow copy to a target server.
4. Optionally, clear the read-only flags on the shadow copy.

The data is now ready to use.

5. User & Group Management

5.1. Overview

There are two system environments for users and groups: workgroup and domain. Because users and groups in a domain environment are managed through standard Windows or Active Directory domain administration methods, this document discusses only local users and groups, which are stored and managed on the NAS appliance. For information on managing users and groups on a domain, refer to the domain documentation available on the Microsoft website.

5.2. Domain versus Workgroup Environments

There are two system environments for users and groups: workgroup and domain. Because users and groups

NAS appliances can be deployed in workgroup or domain environments. When in a domain environment, the appliance is a member of the domain. The domain controller is a repository of accounts and account access for the NAS appliance. Client machines are also members of the domain, and users log on to the domain through their Windows clients. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain.

In a CIFS/SMB environment, when mapping a network drive or a client machine, a user sends a logon credential to the NAS appliance. This credential includes the username, password, and if appropriate, domain information. Using the credential, the NAS appliance authenticates and provides the corresponding access to the user.

When a NAS appliance is deployed into a workgroup environment, all user and group account access permissions to file resources are stored locally on the appliance.

In contrast, when a NAS appliance is deployed into a domain environment it uses the account database from the domain controller, with user and group accounts stored outside the appliance. The NAS appliance integrates with the domain controller infrastructure.

> Note
>
> The NAS appliance cannot 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 NAS appliance, resulting in a workgroup configuration.

Administering users and groups in a domain environment is similar in a mechanical sense to administering them in a workgroup environment. If using an Active Directory domain controller, the Computer Management tool allows for adding, modifying, and removing users in the same context as in a workgroup environment. The concepts, however, are very different.

Additional information about planning for domain environments can be found at Microsoft® using the following URL:

http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx

The configuration of the domain controller is reflected on the NAS appliance because it obtains user account information from the domain controller when deployed in a domain environment. As mentioned previously, the server cannot act as a domain controller itself.

5.3. User & Group Name Planning

Effective user and group management is dependent 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 selected group, or groups, is more efficient than assigning permissions to each user.

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.

5.3.1. 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. Common examples include:

- First initial followed by last name (jdoe for John Doe)
- First initial followed by middle initial and last name (jqpublic for John Q. Public)
- First name followed by last name, separated by a period (john.smith for John Smith)

- Last name followed by first initial (doej for Jane 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 and jdoe2).
Other conventions can be applied. Just ensure that conventions are both systematic and consistent.

5.3.2. Managing Group Names

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

It is recommended that group naming conventions be systematic and easy to understand. Make the group name convey some logical information about the function or purpose of the group. The following table provides examples of group names.

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 the device, the network administrator can create a “Data Users ROnly” group and a “Data Users RWrite” group to contain users that have read only or read write access on the share, respectively.

5.4. Workgroup User & Group Management

In a workgroup environment, users and groups are managed through the Microsoft Management Console of the NAS server. Within the Users option, there are two choices:

- Managing local users
- Managing local groups

User and group administrative tasks include adding, deleting, and modifying user and group information. Managing local users and managing local groups.

This Page Intentionally Left Blank

6. Folder & Share Management

The Tandberg Viking Series NAS appliance supports several file sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This chapter discusses overview information as well as procedural instructions for the setup and management of the file shares for the supported protocols. In addition, discussions on security at the file level and at the share level are included in this chapter.

Abbreviated information on creating NFS file shares is included in this chapter; for detailed information on setting up and managing NFS file shares, see the “Microsoft Services for NFS” chapter.

More information about Windows file system security is available on the Microsoft® website:

http://www.microsoft.com

All procedures in this chapter are documented using the WebUI. In addition to this guide, you may use the WebUI online help.

6.1. Folder Management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Although a variety of methods can be used to create and manage file folders on the NAS appliance, this document discusses using the Windows Remote Desktop or direct attached method interface.

Managing system volumes and file folders includes the following tasks:

- Navigating to a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder

6.2. Managing File Level Permissions

Security at the file level is managed using Windows Explorer available from the Desktop of the NAS appliance. To access the NAS server Desktop from use either the direct attached method or Windows Remote Desktop.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, navigate to the folder or file that needs to be changed and then right-click the folder.

2. Select **Properties**, and then select the **Security** tab.

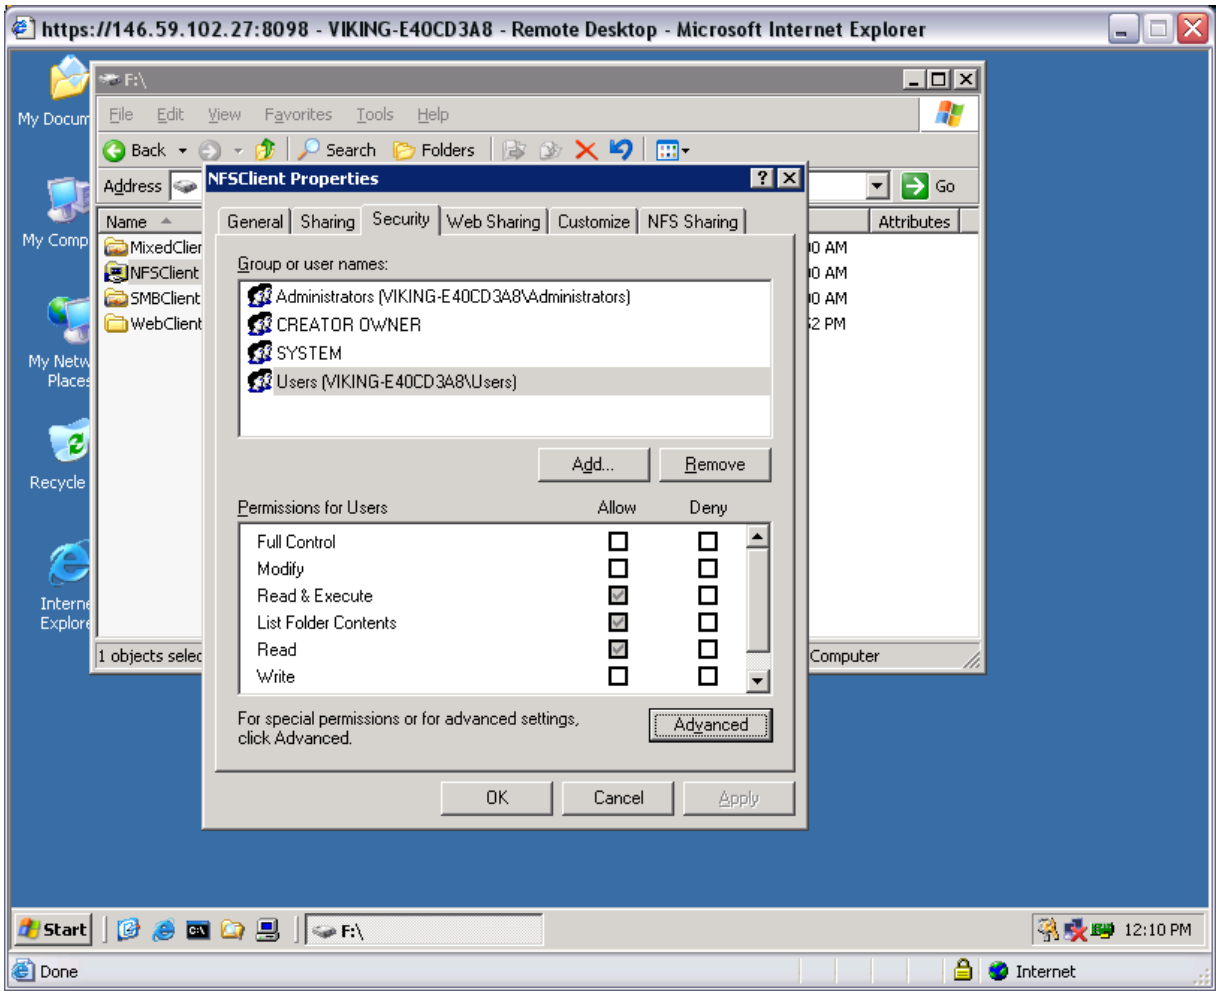

3. Several options are available in the **Security** tab dialog box:

    - To add users and groups to the permissions list, click **Add**. Then follow the dialog box instructions.
    - To remove users and groups from the permissions list, highlight the desired user or group and then click **Remove**.
    - The center section of the Security tab provides a listing of permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file access levels.
    - To modify ownership of files or to modify individual file access level permissions, click **Advanced**.

4. Click **Advanced**. Figure below illustrates the properties available on the **Advanced Security Settings** page.

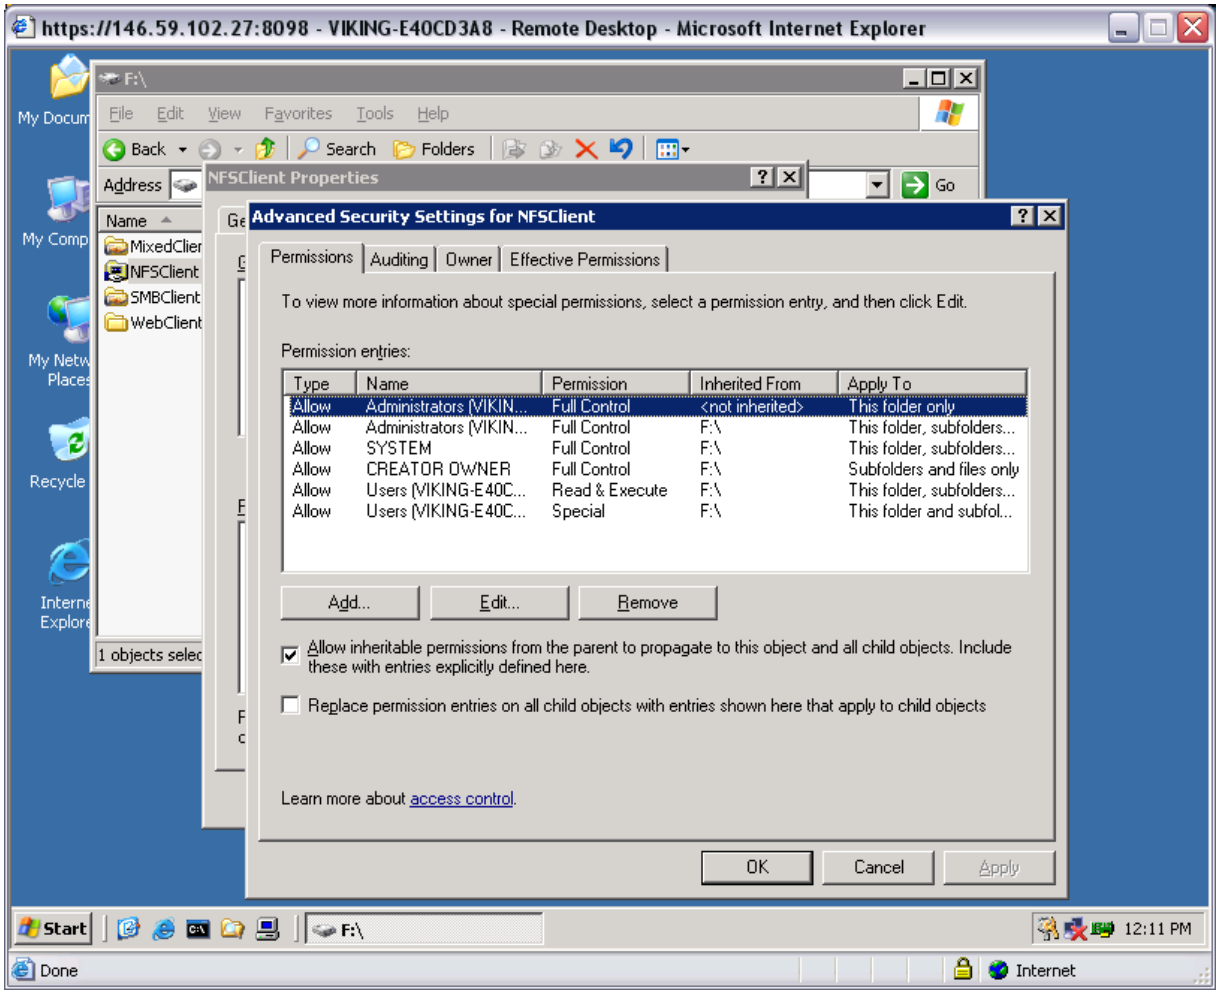

To modify specific permissions assigned to a particular user or group for a selected file or folder in the **Advanced** screen:

1. Select the desired user or group.

2. Click **Edit**.

3. Check all the permissions that you want to enable, and clear the permissions that you want to disable. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. Figure below illustrates the **Edit** screen and some of the permissions.

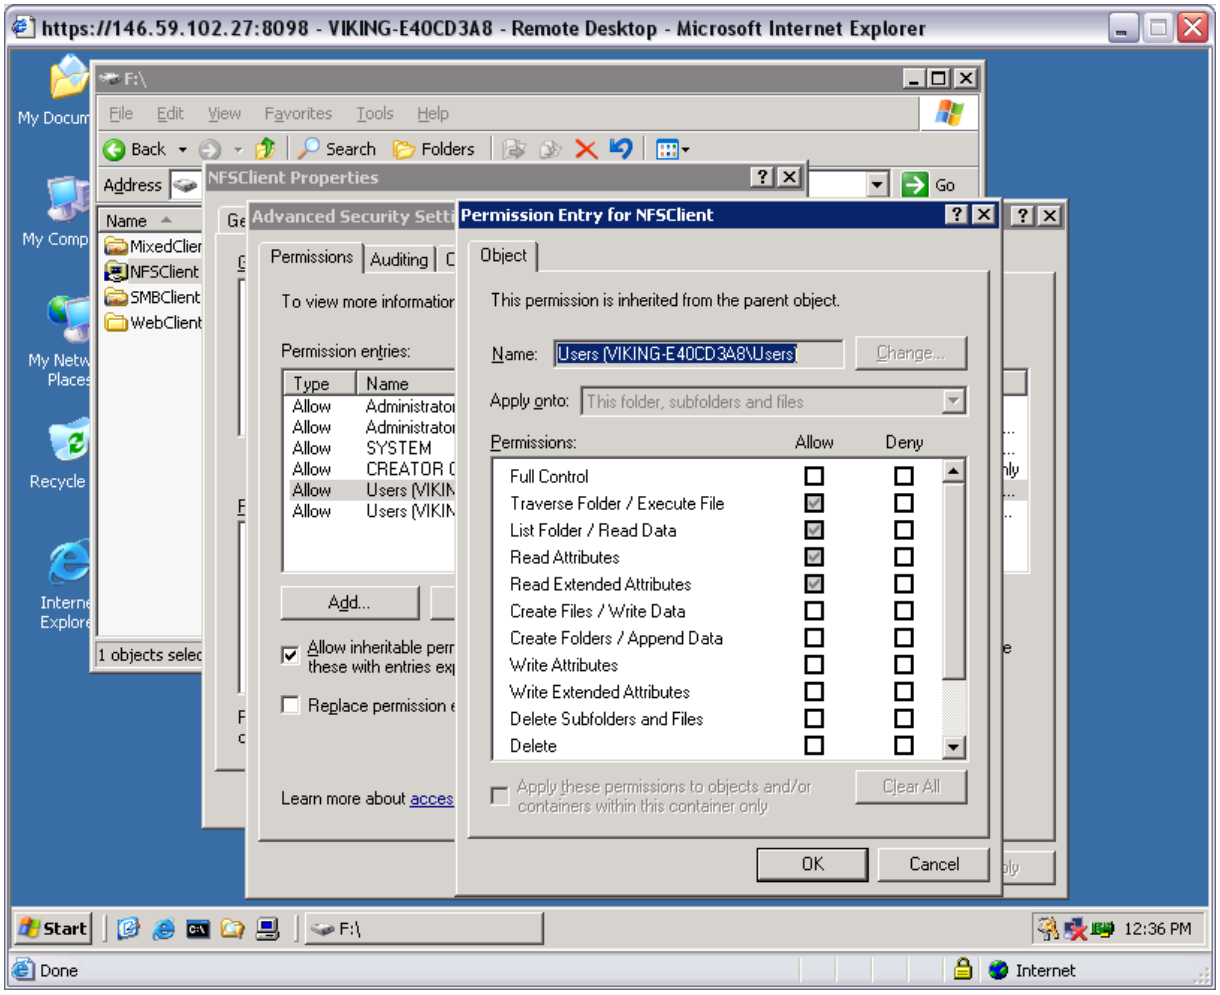

4. Other functionality available in the **Advanced Security Settings** tab includes:

   - **Add a new user or group**. Click **Add**, and then follow the dialog box instructions.

   - **Remove a user or group**. Click **Remove**.

- **Replace permission entries on all child objects with entries shown here that apply to child objects**. This allows all child folders and files to inherit the current folder permissions by default.

5. Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the advanced **Advanced Security Settings Auditing** tab. The **Auditing** tab dialog box is illustrated below.

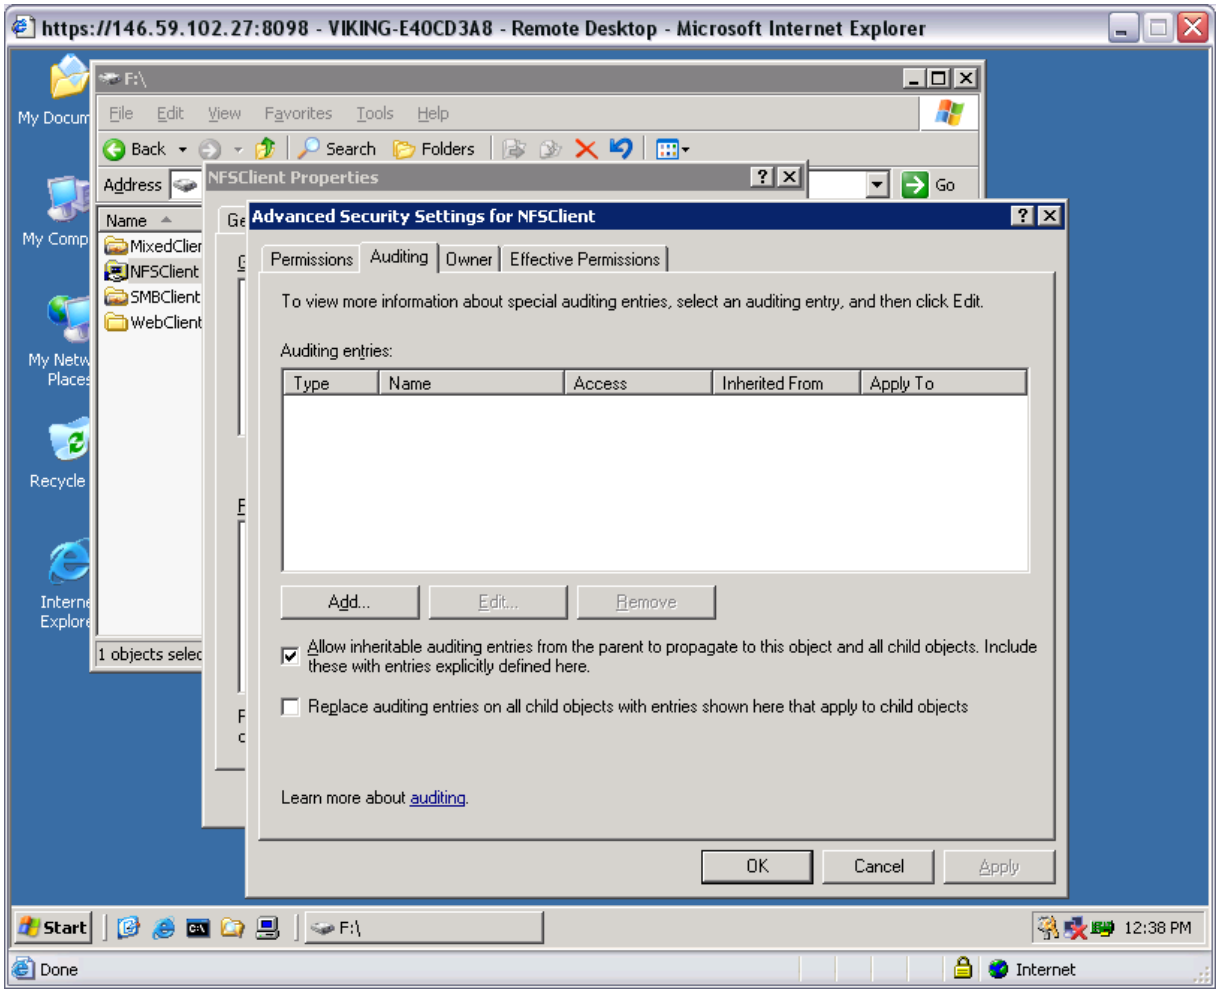

6. Click **Add** to display the Select User or Group dialog box.

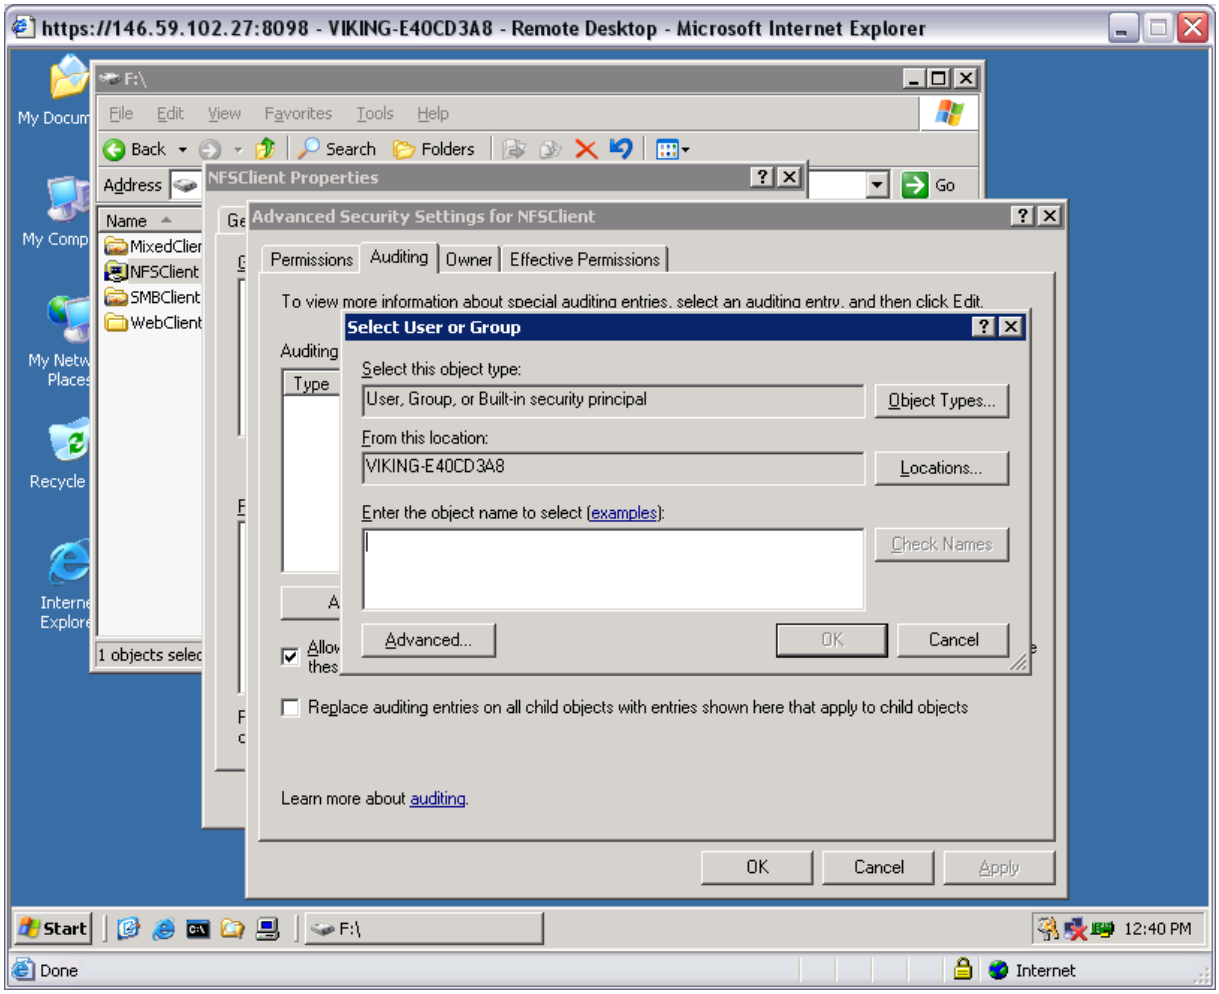

Note
Click Advanced to search for users or groups.

7. Select the user or group.

8. Click **OK**. Figure below illustrates the **Auditing Entry** screen that is displayed.

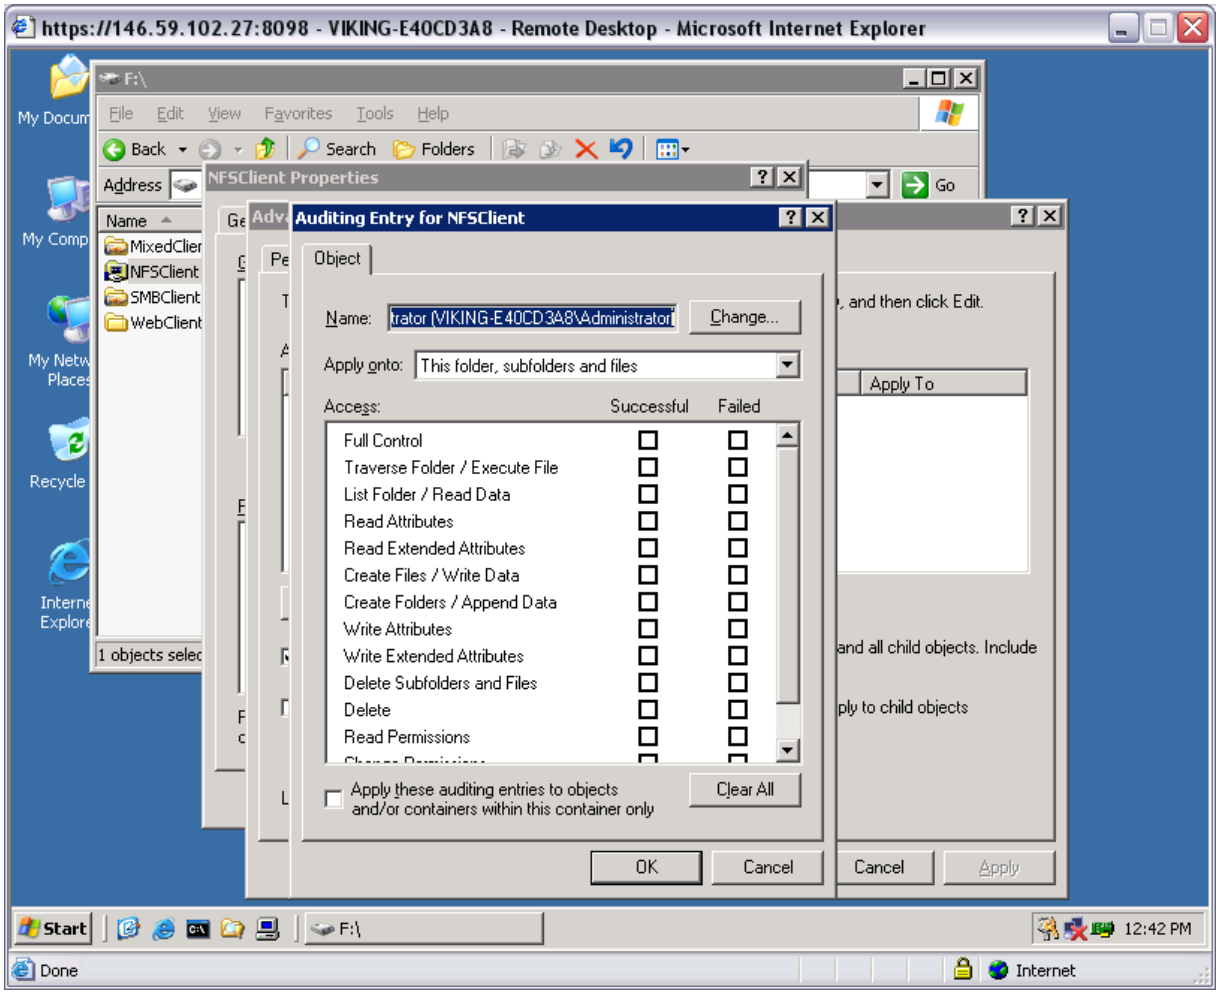

9. Select the desired **Successful** and **Failed** audits for the user or group as shown in Figure 40.

10. Click **OK**.

Note
Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the NAS appliance.

11. The **Owner** tab allows for taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files and then manually apply the appropriate security configurations. Figure below illustrates the **Owner** tab.

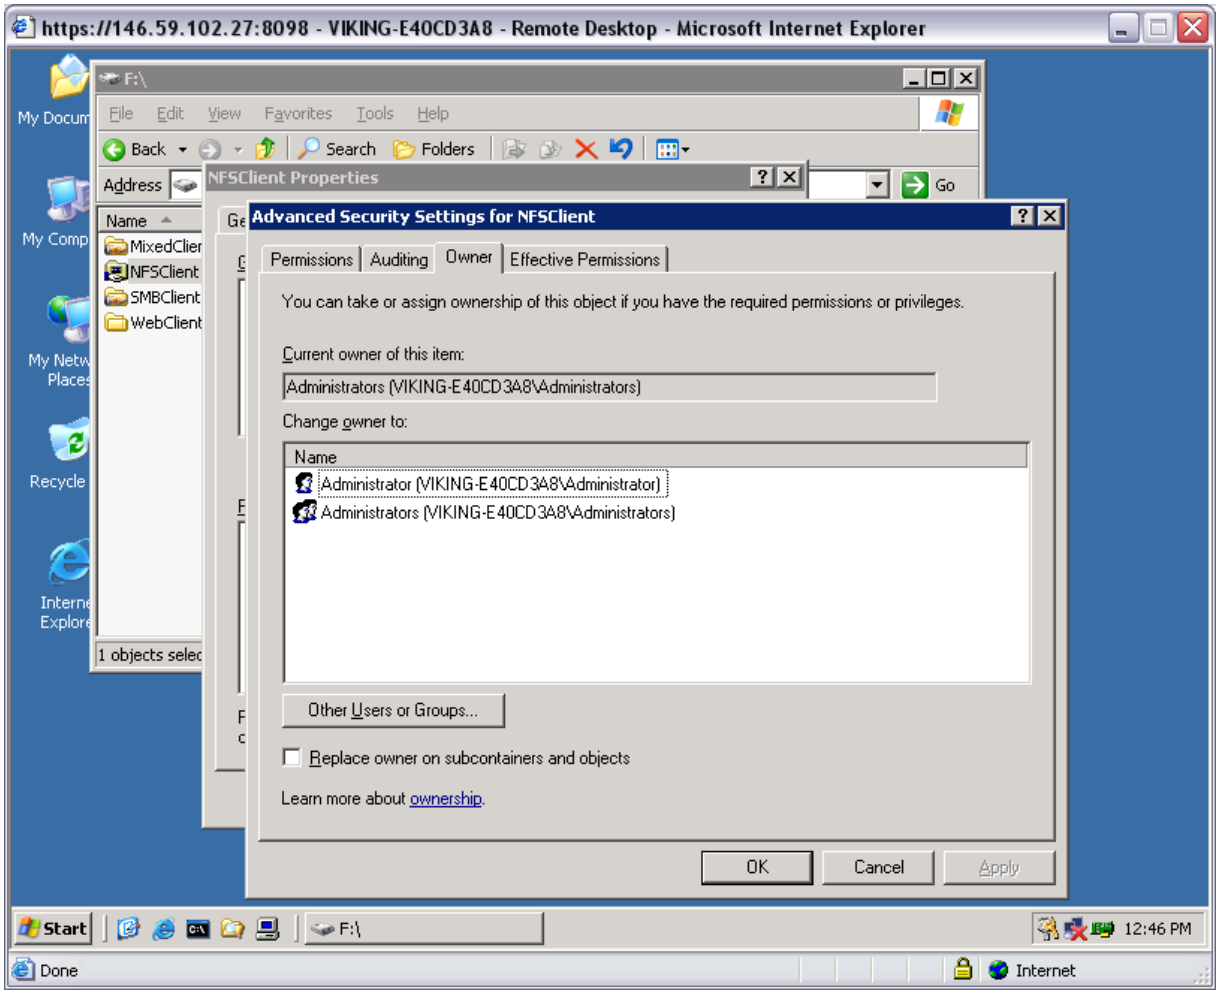

12. The current owner of the file or folder is listed at the top of the screen. To take ownership:

    a. Select the appropriate user or group from the **Change owner to** list.

    b. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.

    c. Click **OK** to execute the commands.

6.3. Share Management

There are several ways to set up and manage shares. Methods include using a command line interface, Windows Explorer, or NAS Management Console.

As previously mentioned, the file sharing security model of the NAS appliance is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security. See “Managing File Level Permissions” earlier in this chapter for information on file security.

Shares management topics include:

- Share Considerations
- Defining Access Control Lists
- Integrating Local File System Security into Windows Domain Environments
- Comparing Administrative and Standard Shares
- Planning for Compatibility between File-Sharing Protocols
- Managing Shares

6.3.1. Share Considerations

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

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature or of 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. Take care to avoid creating shares unnecessarily. 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 NAS appliance 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.

6.3.2. 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 the type of access that 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.

6.3.3. 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 multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the NAS appliance can be given access permissions to shares managed by the device. The domain name of the NAS appliance 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.

Note

Share permissions and file level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file level permissions override the share permissions.

6.3.4. Comparing Administrative & Standard Shares

CIFS supports both administrative shares and standard shares. Administrative shares are shares with a last character of $. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server. Standard shares are shares that do not end in a $ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The NAS appliance supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the $ character when setting up the share. Do not type a $ character at the end of the share name when creating a standard share.

6.3.5. Planning for Compatibility between File Sharing Protocols

When planning for cross-platform share management on the NAS appliance, it is important to understand the different protocols and their associated constraints. Each additional protocol that is supported adds another level of constraints and complexity.

6.3.5.1. NFS Compatibility Issues

When planning to manage CIFS and NFS shares, consider two specific requirements.

Note
Further information, including details about the NFS Service and the User Mapping service, is available in the “UNIX File System Management” chapter.

- **NFS service does not support spaces in the names for NFS file shares.**

  NFS translates any spaces in an export into an underscore character. Additional translations can be set up for files. See the “OEM Supplemental Help” chapter of the SFU help, found on the NAS appliance. This feature is designed to ensure the greatest level of compatibility with NFS clients, because some do not work with NFS exports that contain a space in the export name.

  If you plan to use the same name when sharing a folder through CIFS, and then exporting it through NFS, do not put spaces in the CIFS share name.

- **NFS service does not support exporting a child folder when its parent folder has already been exported.**

  An NFS client can access a child folder by selecting the parent folder and then navigating to the child folder. If strict cross-platform compatibility is an administration goal, CIFS must be managed in the same way. Do not share a folder through CIFS if the parent folder is already shared.

6.3.6. Managing Shares

Shares can be managed through the Windows Storage Server Management Console. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties
- Publish in DFS (See “Publishing a new share in DFS”)

Each of these tasks is discussed in this section.

6.3.6.1. Creating a New Share

To create a new share:

1. From Windows Storage Server Management console, drop down the Share Folder Management. Select the **Shares** folder and then right-click. From the pop-up menu, click on **New Share**.

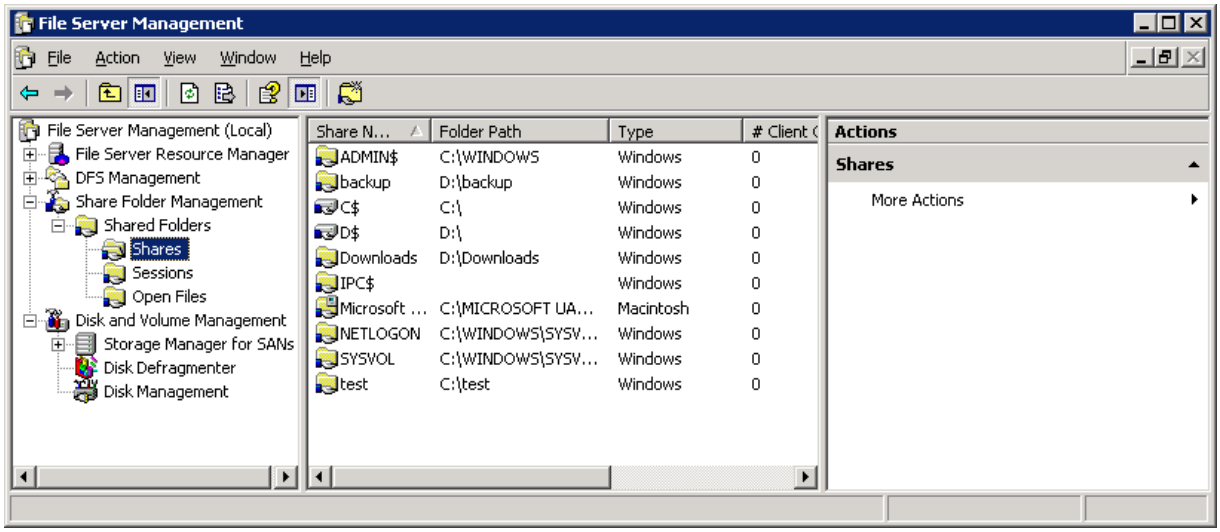

2. Follow the Share a Folder Wizard to create share folder by specifying the following information:
    - Share name
    - Share path
    - Client protocol types (SMB/CIFS and/or AppleTalk)
    - Access permission
3. At the end of the Wizard, click **Close**.

6.3.6.2. Stopping a Share

Caution

Before deleting a share, warn all users to exit that share and confirm that no one is using the share.

To delete a share:

1. On the Windows Storage Server Management console, right-click on the share that you want to stop sharing. Click **Stop Sharing** from the pop-up menu.

2. Click **Yes** to confirm.

Notes

When a share is stopped, the physical folder is NOT deleted. Use the Windows Explorer to manually delete the folder if desired.

6.3.6.3. Modifying Share Properties

To change share settings:

1. On the Windows Storage Server Management console, right-click on the share that you want to change settings. Click **Properties** from the pop-up menu.

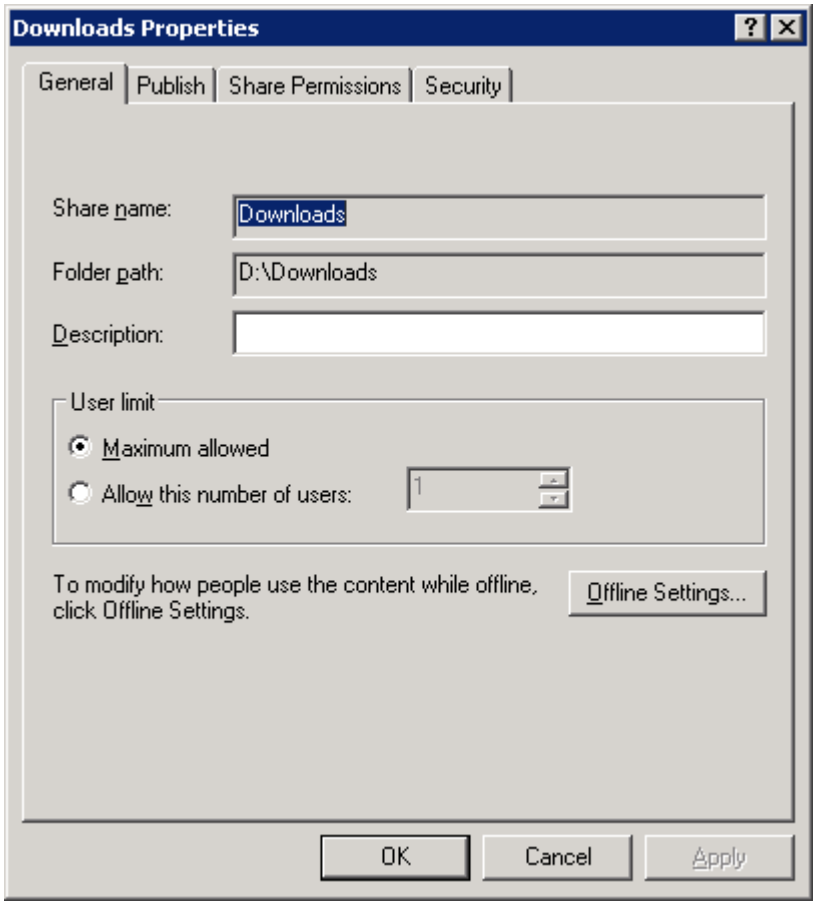

2. The name and path of the selected share is displayed.

3. After all share information has been entered, click **OK**. The setting is updated.

Windows Sharing

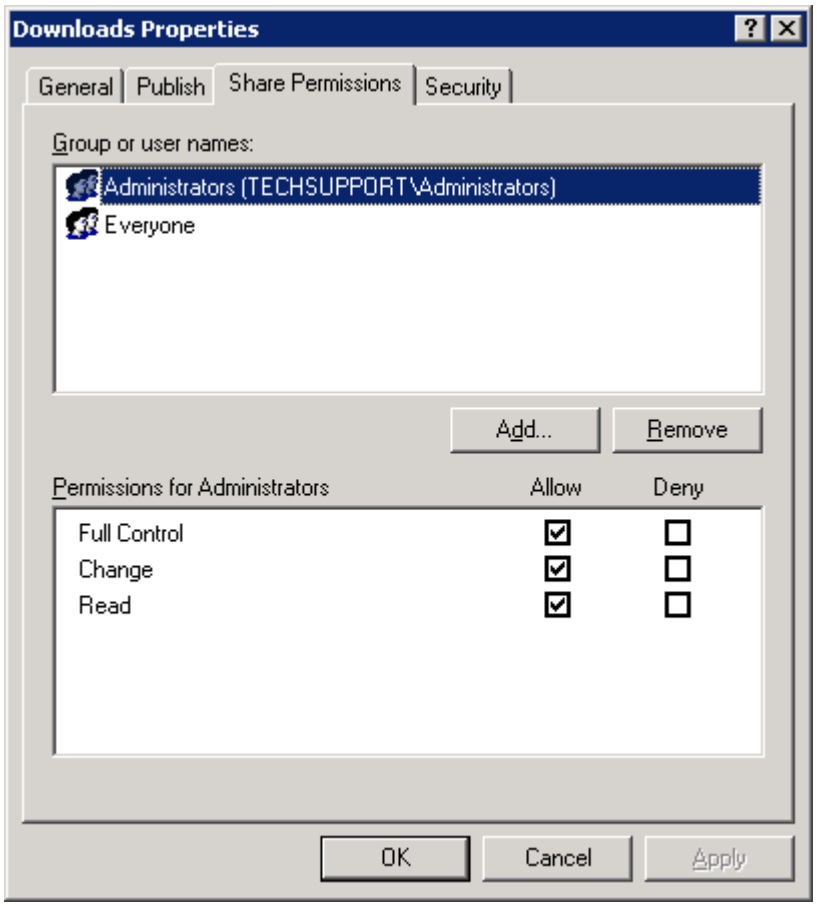

From the **Share Properties** menu:

1. Enter a descriptive Comment, and the User limit (optional). See figure below for an example of the screen display.

2. Select Offline settings.

3. Set the permissions. The **Share Permissions** tab lists the currently approved users for this share.

    - To add a new user or group, click **Add…**. Specify the group or user information and they will be added to the Permissions box.

    - To remove access to a currently approved user or group, select the user or group from the Permissions box and then click Remove.

    - To indicate the type of access allowed for each user or group, select them and specify the appropriate option.

4. After all Windows Sharing information is entered, click **OK**.

UNIX Sharing

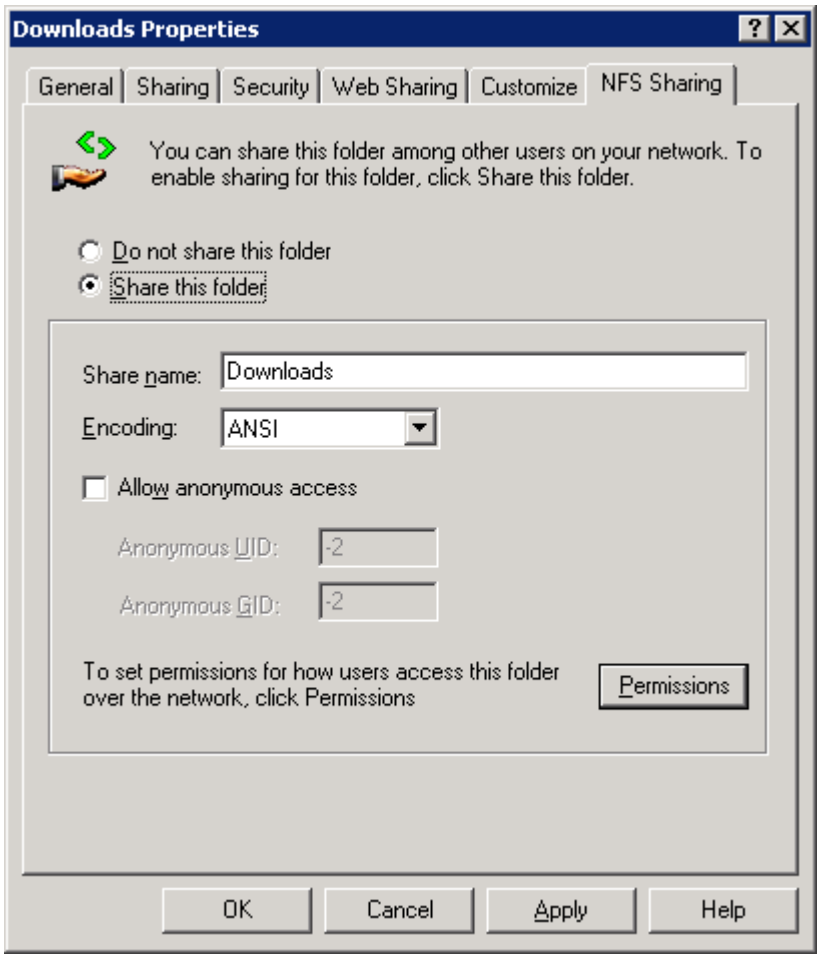

The UNIX (NFS) sharing need to be access using Windows Explorer:

1. Using the Windows Explorer, navigate to the share folder and right-click. Select **Properties**. Go to the **NFS Sharing** tab.

2. Select the radio button **Share this folder** to enable UNIX sharing (NFS). Define the share name, encoding.

3. At the Indicate the machines that will have access to this share. Select the machine to include in the **Select a group** box or manually enter the NFS client computer name or IP address. Then click **Add**.

4. Click the **Permissions** to define how users access the shared folder.

   From the **NFS Share Permissions** menu, define the type of access on the drop down box. The types of access are:

   - **Read-only**—Use this permission to restrict write access to the share.

   - **Read-write**—Use this permission to allow clients to read or write to the share.

- **No access**—Use this permission to restrict all access to the share.

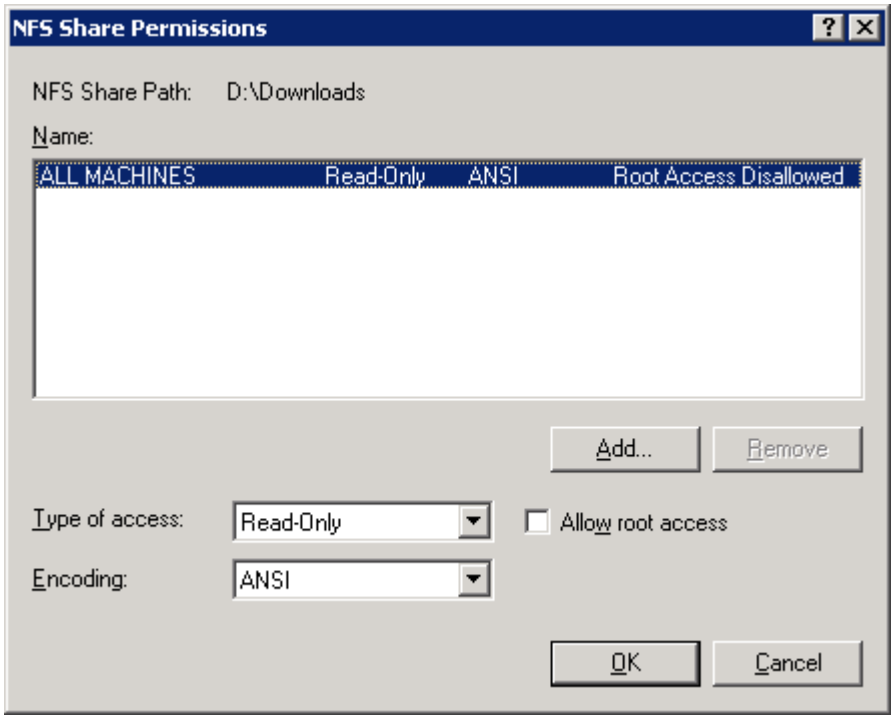

5. Select whether or not to allow root access by checking the box.

   - **Read-only + Root**—Use this permission to restrict write access to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.

   - **Read-write + Root**—Use this permission to allow clients to read or write to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.

6. After all UNIX sharing information is entered, click **OK**.

Web Sharing (HTTP)

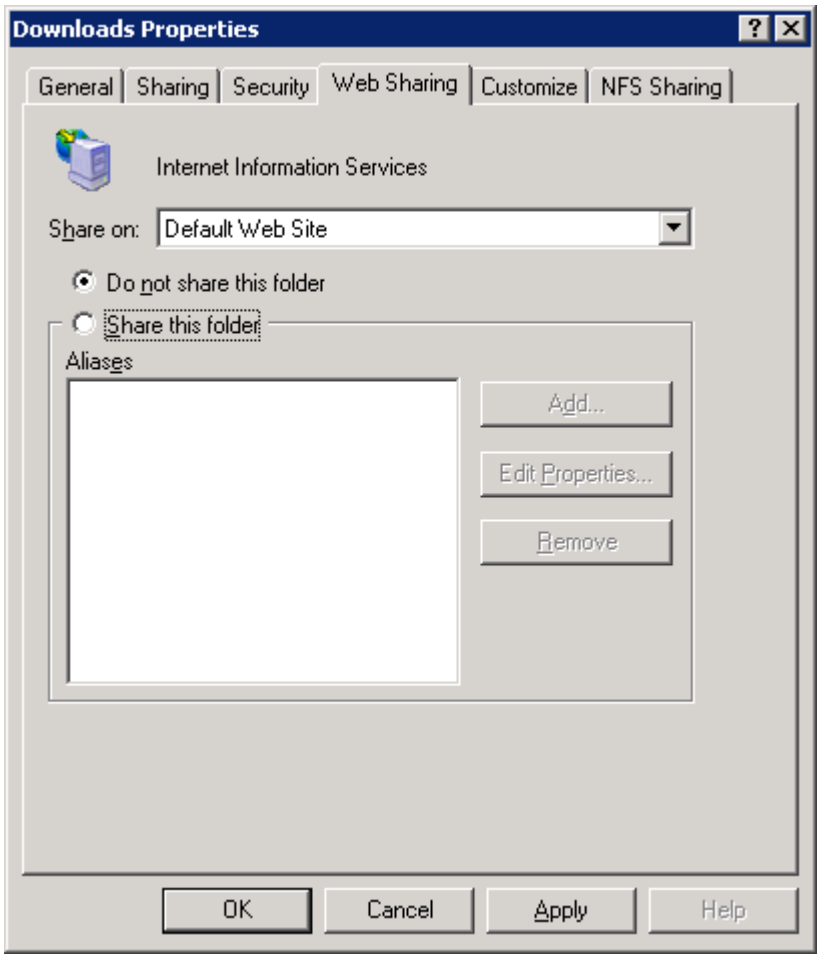

The HTTP sharing can be enabled using Windows Explorer:

1. Select Using the Windows Explorer, navigate to the share folder and right-click. Select **Properties**. Go to the **Web Sharing** tab.

2. Select the radio button **Share this folder** to enable web sharing (HTTP). Define the alias information.

3. Click **OK**.

AFP (Appletalk) Sharing

AppleTalk shares can be set up only after AppleTalk Protocol and File Services for Macintosh have been installed on the NAS appliance.

Note
AppleTalk shares should not be created on clustered resources as data loss can occur due to local memory use.

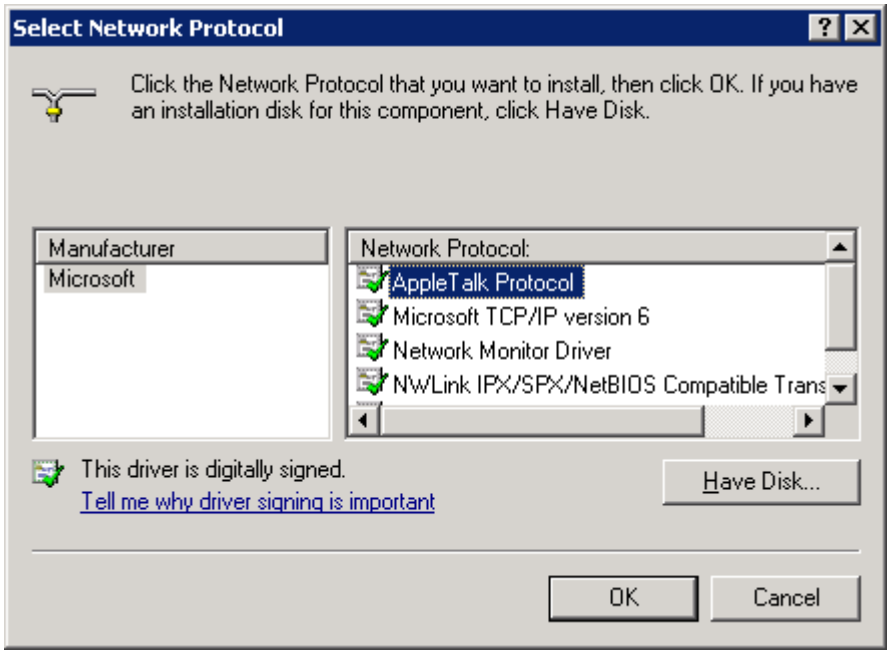

Instaling the AppleTalk Protocol

To install the AppleTalk Protocol:

1. From the desktop of the NAS server, click **Start**, navigate to **Control Panel > Network Connections > Local Area Connection**. Right-click and select **Properties** from the pop-up menu, and then click on **Install…**.

2. Select **Protocol** and click **Add…**.

3. Select **AppleTalk Protocol** and click **OK**.

4. Click **Close** to finish the installation.

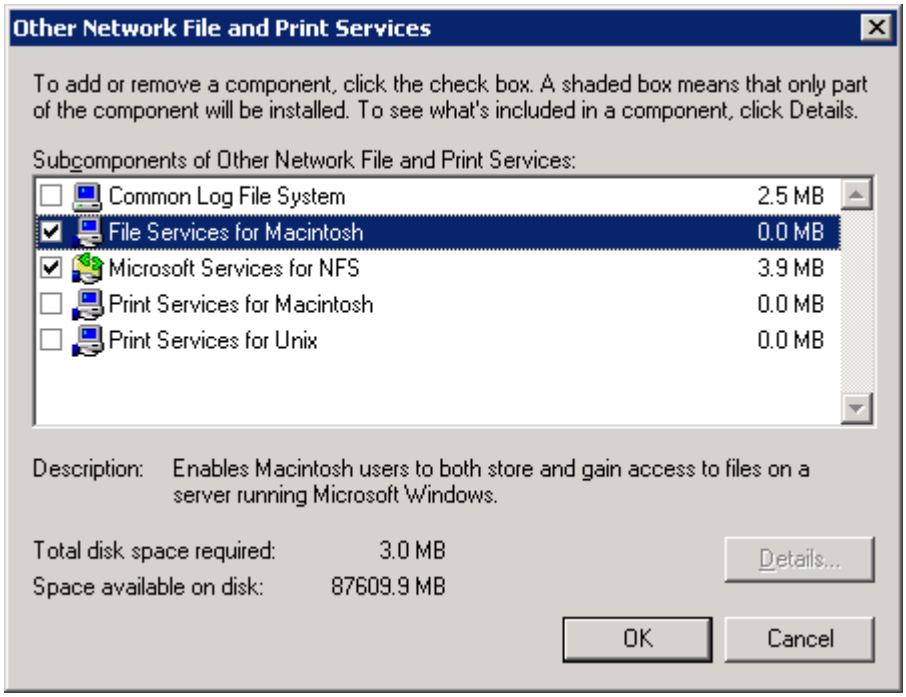

Instaling the File Services for Macintosh

To install the AppleTalk Protocol:

1. From the desktop of the NAS server, click **Start**, navigate to **Control Panel > Add or Remove Programs**. Click on **Add/Remove Windows Components** and navigate to the **Other Network and File and Print Services**, and then click on **Details…**.

2. Check the **File Services for Macintosh**.

3. Click **OK**.

4. Click **Next**.

5. Click **Finish** to exit **Windows Component Wizard** and commence installation.

Setting AppleTalk Protocol Properties

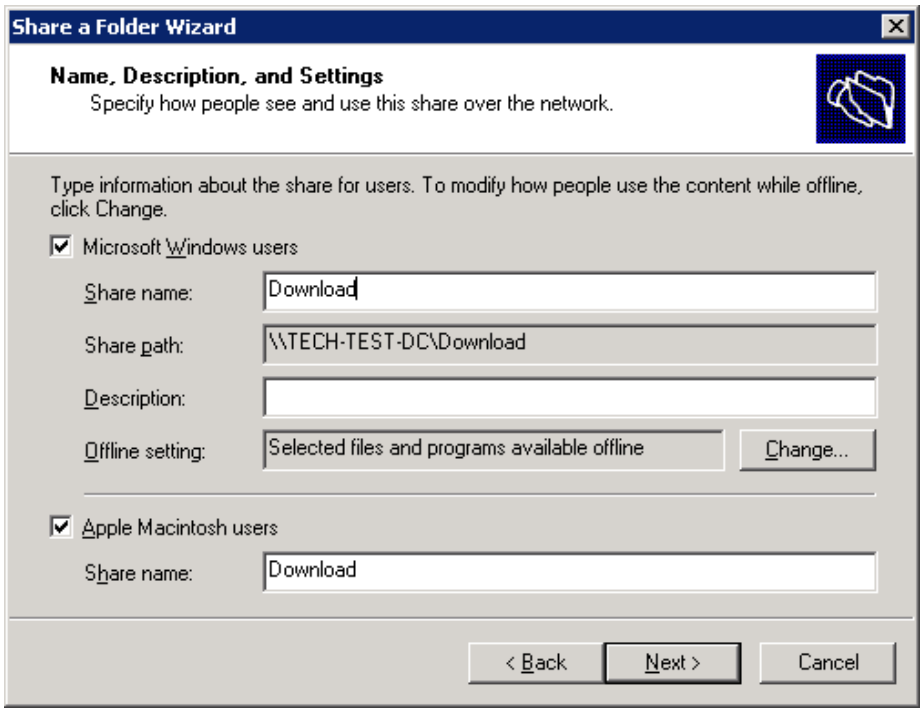

To set up AppleTalk shares:

1. From Windows Storage Server Management console, drop down the Share Folder Management. Select the **Shares** folder and then right-click. From the pop-up menu, click on **New Share**.

2. Step through the **Share a Folder Wizard** to create shares.

3. Type in the share name and share path.

4. Check **Apple MacIntosh users** checkbox and finish the share creation process.

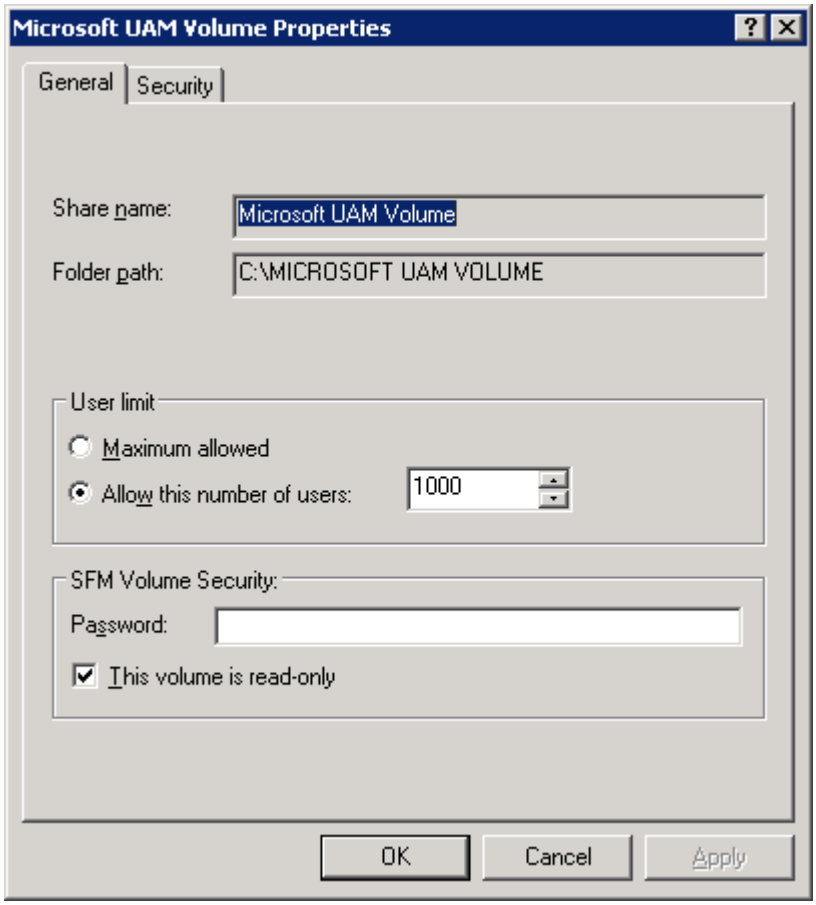

To change AppleTalk settings:

1. On the Windows Storage Server Management console, right-click on the AppleTalk share that you want to change settings. Click **Properties** from the pop-up menu.

2. Enter a user limit.

3. Enter password information.

4. Indicate whether the share has read only permission or read write permission by checking or un-checking the checkbox.

5. After all AppleTalk Sharing information is entered, click **OK**.

6.4. File Server Recourse Manager

With the increasing demand on storage resources, as organizations rely more heavily on data than ever before, IT administrators face the challenge of overseeing a larger and more complex storage infrastructure, while at the same time, tracking the kind of information available in it. Today, managing storage resources not only includes data size and availability but also the enforcement of company policies and a very good understanding of how existing storage is being used, allowing for sound strategic planning and proper response to organizational changes.

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 File Server Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports. This set of advanced instruments not only helps the administrator to efficiently monitor existing storage resources, but it also aids in the planning and implementation of future policy changes.

By using File Server 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 or exceeded.
- Automatically generate and apply quotas to all existing folders and any new subfolders in a volume or folder.
- Create file screens to control the type of files that users can save and send notifications when users attempt to save blocked files.
- Define quota and file screening templates that can be easily applied to new volumes or folders and reused across an organization.
- Schedule periodic storage reports that help identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or for a selected group of users.
- Generate storage reports instantly, on demand.

6.4.1. Using the File Server Resource Manager Component

This section tells you how to open and use File Server Resource Manager. The following tasks are included:

- Open File Server Resource Manager
- Perform storage management tasks
- Manage storage on a remote computer

6.4.2. Opening File Server Resource Manager

File Server Resource Manager is a Microsoft Management Console (MMC) snap-in and is located in Administrative Tools on the Start menu.

To open File Server Resource Manager:

- On the **Start** menu, click **Programs**, click **Administrative Tools**, and then click **File Server Resource Manager**.

File Server Resource Manager contains three main nodes:

- **Quota Management**. Use to create quotas that place size limits on volumes and folders.
- **File Screening Management**. Use to create file screens that block files from volumes and folders.
- **Storage Reports Management**. Use to schedule different types of storage reports and to create reports on demand.

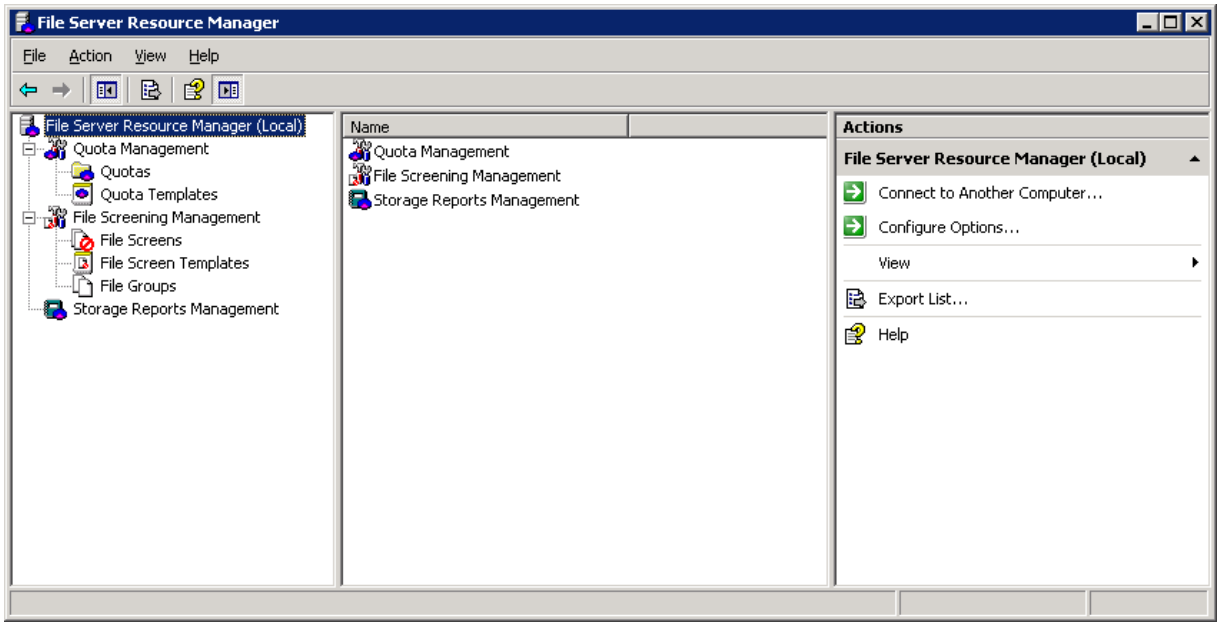

The following is a list of tasks that you can perform from each respective node in File Server Resource Manager:

6.4.3. Quota Management

- Create, manage, and obtain information about quotas, which are used to set a space limit on a volume or folder. By defining notification thresholds, you can send e-mail notifications, log an event, run a command or script, or generate reports when users approach or exceed a quota.
- Create and manage quota templates to simplify quota management.
- Create and manage auto quotas.

6.4.4. File Screening Management

- Create, manage, and obtain information about file screens, which are used to block selected file types from a volume or folder.
- Create file screening exceptions to override certain file screening rules.
- Create and manage file screen templates to simplify file screening management.
- Create and manage file groups. When used with file screens and file screening exceptions, the file groups determine which files will be blocked and which will be allowed. File groups also are used to select files to include and exclude from the Files by File Group Report and to sort file screens.

6.4.5. Storage Report Management

- Schedule and configure storage reports.
- Generate storage reports on demand.

6.4.6. Using Command-Line Tools for File Server Resource Manager

If you prefer to work from the command line, you can use the following tools:

- `Dirquota.exe`. Use to create and manage quotas and quota templates.
- `FileScrn.exe`. Use to create and manage file screens, file screening exceptions, and file groups.
- `StorRept.exe`. Use to configure report parameters and generate storage reports on demand. You can also create report tasks and then
- use `Schtasks.exe` to schedule the tasks.

The tools are added to the system path when you install File Server Resource Manager and can be run from the command prompt. They can be used to perform storage management tasks on remote computers that are running the same operating system.

To learn the parameters for a specific tool, include the `/?` switch.

7. Distributed File System (DFS)

7.1. Overview

One of the goals of information technology (IT) groups in medium and large organizations is to manage file servers and their resources efficiently while keeping them available and secure for users. As organizations expand to include more users and servers—whether they are located in one site or in geographically distributed sites—administrators find it increasingly difficult to keep users connected to the files they need. On one hand, storing files on distributed servers makes files available to more users and decreases latency and bandwidth use when the servers are located near users. On the other hand, as the number of distributed servers increases, users have difficulty locating files they need, and operational costs increase.

Administrators who manage these distributed, remote servers need a solution that helps them limit network traffic over slow WAN connections, ensure the availability of files during WAN outages or server failures, and ensure that branch servers are backed up correctly. The Distributed File System solution in the Microsoft® Windows Server™ 2003 R2 operating system helps administrators address these challenges by providing two technologies, DFS Namespaces and DFS Replication, which, when used together, offer simplified, fault-tolerant access to files, load sharing, and WAN-friendly replication.

- DFS Replication is a new state-based, multimaster replication engine that supports replication scheduling and bandwidth throttling. DFS Replication uses a new compression protocol called Remote Differential Compression (RDC), which can be used to efficiently update files over a limited-bandwidth network. RDC detects insertions, removals, and re-arrangements of data in files, thereby enabling DFS Replication to replicate only the changes when files are updated. Additionally, a function of RDC called cross-file RDC can help reduce the amount of bandwidth required to replicate new files.

- DFS Namespaces, formerly known as Distributed File System, allows administrators to group shared folders located on different servers and present them to users as a virtual tree of folders known as a namespace. A namespace provides numerous benefits, including increased availability of data, load sharing, and simplified data migration.

The following figure illustrates how DFS Namespaces and DFS Replication work together. The processes marked 1 and 2 are described in more detail following the figure.

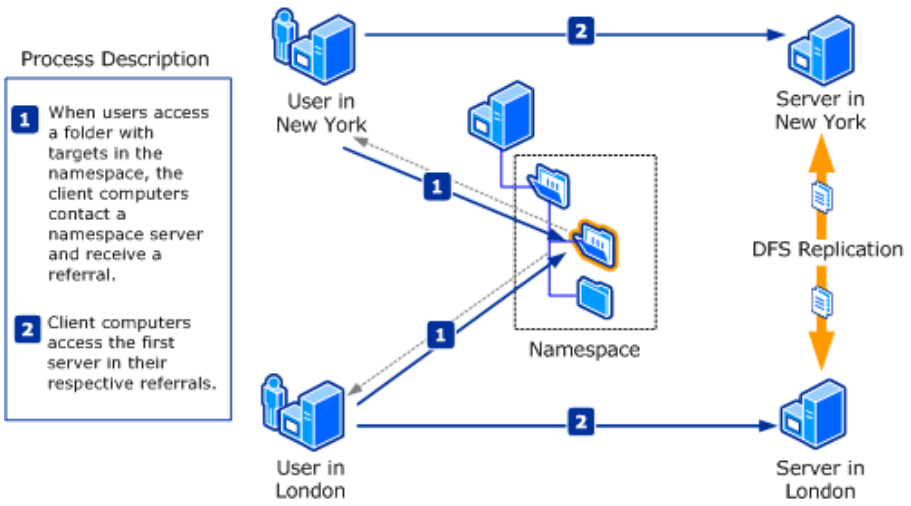

As the figure shows, when a user attempts to access a folder in the namespace (1), the client computer contacts a namespace server. The namespace server sends the client computer a referral that contains a list of servers that host the shared folders (called folder targets) associated with the folder. The client computer caches the referral and then contacts the first server in the referral (2), typically a server in the client's own site unless no same-site servers exist or the administrator configures target priority.

The highlighted folder in the figure shows that it is hosted by shared folders on two servers, one in New York and one in London, to provide users in those sites with fast, reliable access to files. The shared folders are kept synchronized by DFS Replication. The fact that multiple servers host the folder is transparent to users, who see only a single folder in the namespace. If one of the servers becomes unavailable, the client computer fails over to the remaining server.

7.2. DFS Namespaces

DFS Namespaces allows you to group shared folders located on different servers by transparently connecting them to one or more namespaces. A namespace is a virtual view of shared folders in an organization. When you create a namespace, you select which shared folders to add to the namespace, design the hierarchy in which those folders appear, and determine the names that the shared folders show in the namespace. When a user views the namespace, the folders appear to reside on a single, high-capacity hard disk. Users can navigate the namespace without needing to know the server names or shared folders hosting the data.

The path to a namespace is similar to a Universal Naming Convention (UNC) path of a shared folder, such as \\Server1\Public\Software\Tools. If you are familiar with UNC paths, you know that in this example the shared folder, Public, and its subfolders, Software and Tools, are all hosted on Server1. Now, assume you want to give users a single place to locate data, but you want to host data on different servers for availability and performance purposes. To do this, you can deploy a namespace similar to the one shown in the following figure. The elements of this namespace are described after the figure.

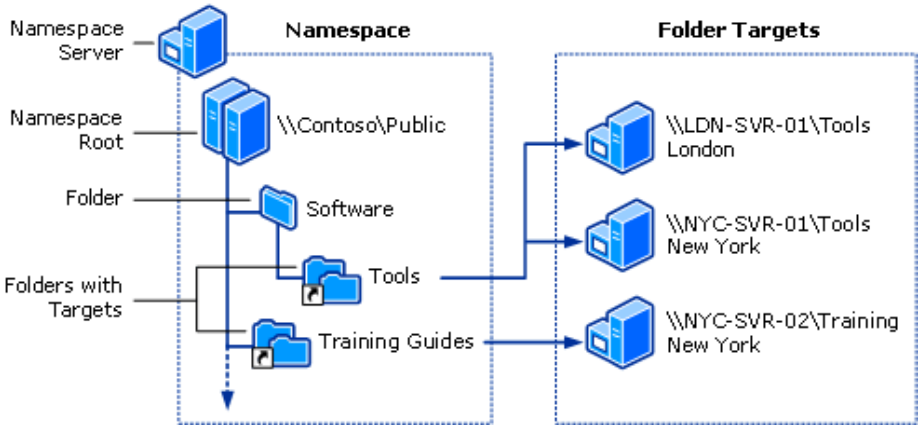

- **Namespace server**. A namespace server hosts a namespace. The namespace server can be a member server or a domain controller.

- **Namespace root**. The root is the starting point of the namespace. In the previous figure, the name of the root is Public, and the namespace path is \\Contoso\Public. This type of namespace is known as a domain-based namespace, because it begins with a domain name (for example, Contoso) and its metadata is stored in Active Directory. Although a single namespace server is shown in the previous figure, a domain-based namespace can be hosted on multiple namespace servers.

- **Folder**. Folders help build the namespace hierarchy. Folders can optionally have folder targets. When users browse a folder with targets in the namespace, the client computer receives a referral that directs the client computer to one of the folder targets.

- **Folder targets**. A folder target is a UNC path of a shared folder or another namespace that is associated with a folder in a namespace. In the previous figure, the folder named Tools has two folder targets, one in London and one in New York, and the folder named Training Guides has a single folder target in New York. A user who browses to \\Contoso\Public\Software\Tools is transparently redirected to the shared folder \\LDN-SVR-01\Tools or \\NYC-SVR-01\Tools, depending on which site the user is in.

7.3. DFS Replication

DFS Replication is the new state-based, multimaster replication engine in Windows Server 2003 R2. Although some DFS Replication concepts and processes are similar to the concepts and processes in File Replication service (FRS), there are several important differences that you should be aware of before you deploy DFS Replication.

7.3.1. DFS Replication Concept

First, let's review the basic concepts of DFS Replication. These concepts—replication groups, connections, members, and replicated folders—are illustrated in the following figure.

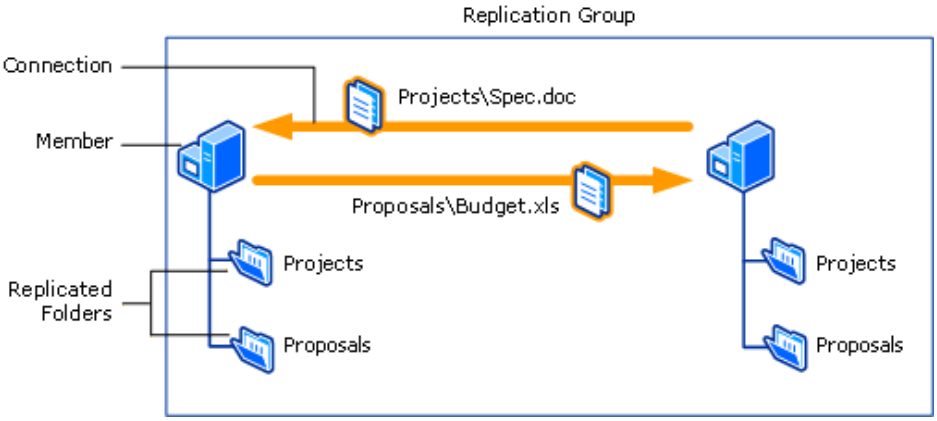

As this figure shows, a replication group is a set of servers, known as members, that participates in the replication of one or more replicated folders. A replicated folder is a folder that is kept synchronized on each member. In the previous figure, there are two replicated folders, Projects and Proposals. As data changes in each replicated folder, the changes are replicated across connections between the members. The connections between all members form the replication topology.

Creating multiple replicated folders in a single replication group simplifies the process of deploying replicated folders, because the topology, schedule, and bandwidth throttling for the replication group are applied to each replicated folder. To deploy additional replicated folders, you can use a short wizard to define the local path and permissions for the new replicated folder. Each replicated folder also has its own settings, such as file and subfolder filters, so that you can filter out different files and subfolders for each replicated folder.

The replicated folders stored on each member can be located on different volumes in the member, and the replicated folders do not need to be shared folders or part of a namespace, though the DFS Management snap-in makes it easy to share replicated folders and optionally publish them in an existing namespace. You will do both in one of the tasks later in this guide.

7.3.2. DFS Initial Replication

When you first set up replication, you must choose a primary member. Choose the member that has the most up-to-date files that you want replicated to all other members of the replication group, because the primary member's content is considered "authoritative." This means that during initial replication, the primary member's files will always win the conflict resolution that occurs when the receiving members have files that are older or newer than the same files on the primary member.
The following explanations will help you better understand the initial replication process:

- Initial replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

- Initial replication always occurs between the primary member and the receiving replication partners of the primary member. After a member has received all files from the primary member, then that member will replicate files to its receiving partners as well. In this way, replication for a new replicated folder starts from the primary member and then progresses out to the other members of the replication group.

- When receiving files from the primary member during initial replication, the receiving members that contain files that are not present on the primary member move those files to their respective DfsrPrivate\PreExisting folder. If a file is identical to a file on the primary member, the file is not replicated. If the version of a file on the receiving member is different from the primary member's version, the receiving member's version is moved to the Conflict and Deleted folder and remote differential compression (RDC) can be used to download only the changed blocks.

- To determine whether files are identical on the primary member and receiving member, DFS Replication compares the files using a hash algorithm. If the files are identical, only minimal metadata is transferred.

- After the initialization of the replicated folder, the "primary member" designation is removed. Initialization takes place after all files that exist before DFS Replication picks up the configuration are added to the DFS Replication database. The member that was previously the primary member is then treated like any other member and its files are no longer considered authoritative over those of other members that have completed initial replication. Any member that has completed initial replication is considered authoritative over members that have not completed initial replication.

7.4. DFS Management Snap-in

The DFS Management snap-in is the graphical user interface tool for managing DFS Namespaces and DFS Replication. This snap-in is new and differs from the Distributed File System snap-in in Windows Server 2003. Therefore, before you begin using DFS Namespaces and DFS Replication, you might want to review the components of this snap-in, which are shown in the following figure and described in the sections that follow.

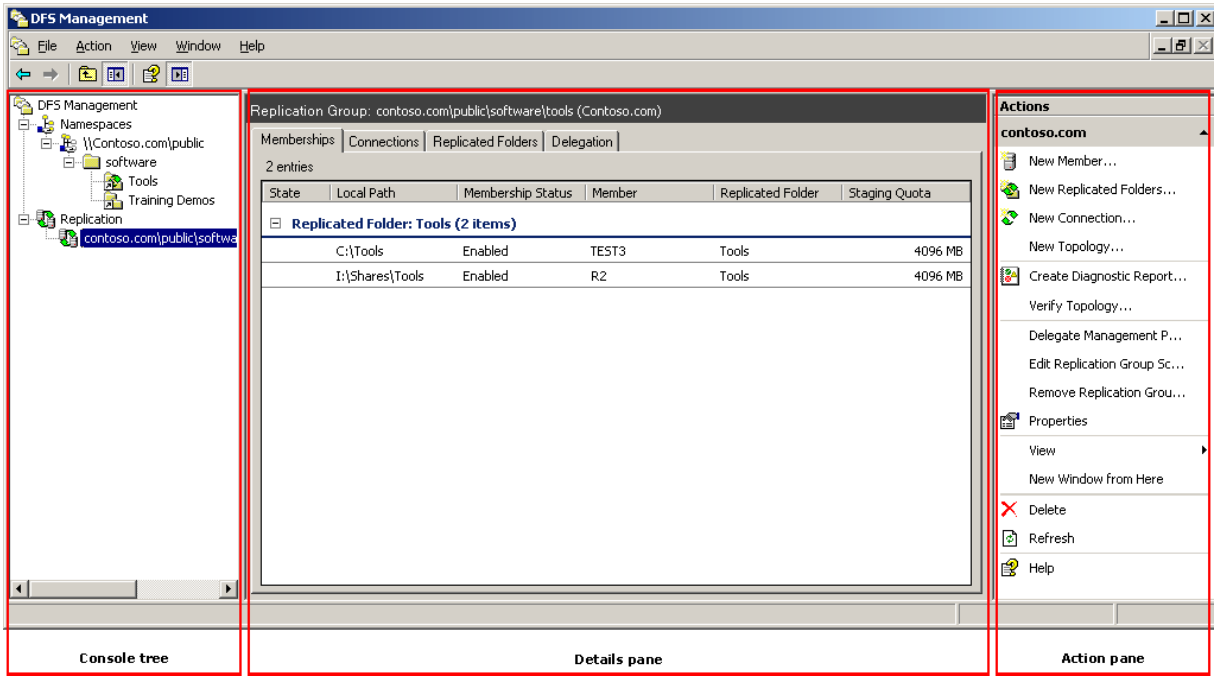

Console Tree

The console tree has two nodes, Namespaces and Replication, from which you can manage namespaces and DFS Replication.

Namespaces node

The following figure shows the elements under the Namespaces node in the console tree.

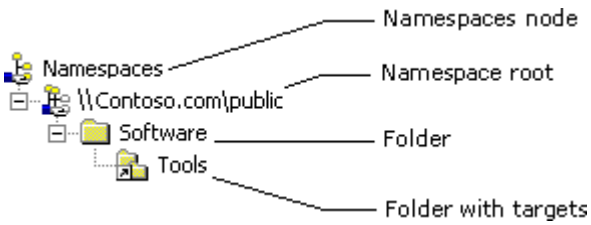

As the figure shows, the Namespaces node contains the namespaces you create as well as any existing namespaces you add to the console display. In the previous figure, one namespace is shown, \\Contoso.com\Public. Under each namespace is a hierarchical view of folders. Folders that have targets use a special icon to differentiate them from folders that do not have targets.

Notes

If you are not familiar with namespace terminology, see the section "DFS Namespaces" earlier section.

Replication node
The following figure shows the elements under the Replication node in the console tree.

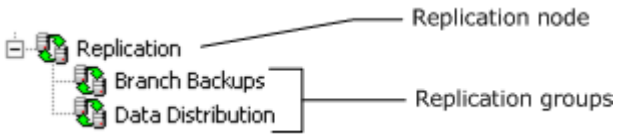

As the figure shows, the Replication node contains the replication groups you create as well as any existing replication groups that you add to the console display. A replication group represents a group of servers that participates in the replication of data. For more information about replication groups, see "DFS Replication" section later in this chapter.

Details Pane
The contents of the details pane change according to what you have selected in the console tree. For example, if you select a namespace in the console tree, you see tabs named Namespace, Namespace Servers, and Delegation in the details pane. If you select a replication group, you see tabs named Memberships, Connections, Replicated Folders, and Delegation. You can double-click objects in the details pane to view their properties.

Action Pane
The Action pane shows two types of tasks: common tasks and tasks that apply to the selected object. If the Action pane is not visible, you can open it using the following steps: click the View menu, click Customize, and then click the Action pane option in the Customize View dialog box.

Each DFS namespace requires a root. A DFS root is a starting point of the DFS namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a server. A root is implemented as a shared folder on the DFS server.

7.5. Deploying Namespace (Step-by-Step Guide)

The tasks in this section walk you through the process of deploying a namespace that looks similar to the namespace shown in the figure that appears in "DFS Namespaces" earlier in this chapter.

7.5.1. Create a Namespace

To create a namespace:

1. In the console tree of the DFS Management snap-in, right-click the **Namespaces** node, and then click **New Namespace**.

2. Follow the steps in the **New Namespace Wizard** and supply the information described in the following:

   - **Namespace Server** - Enter the name of the server to host the namespace. The server can be a domain controller or a member server.

   - **Namespace Name and Settings** - In **Name**, type **Public**.

   - **Namespace Type** - If Active Directory is deployed in your test lab and you are a member of the Domain Admins group or have been delegated permission to create domain-based namespaces, choose **Domain-based namespace**. Otherwise, choose **Stand-alone namespace**. To learn how a member of the Domain Admins group can delegate permission to create domain-based namespaces, see "Security requirements for creating and managing namespaces" in DFS Management Help.

   - **Review Settings and Create Namespace** - Click **Create** to create the namespace.

   - **Confirmation** - Click **Close** to close the wizard.

When the wizard finishes, your new namespace will be added to the console tree. Double-click the Namespaces node, if necessary, to view your namespace, which should be similar to the following figure.

To browse the new namespace, type the following command in the Run dialog box, substituting either the server name (if you created a stand-alone namespace) or the domain name (if you created a domain-based namespace) as appropriate:

\\server_or_domain**Public**

7.5.2. Add a Namespace Server (Domain)

If you created a domain-based namespace, perform this task to specify an additional server to host the namespace. Doing so increases the availability of the namespace and allows you to place namespace servers in the same sites as users. If you created a stand-alone namespace, you must skip this task because stand-alone namespaces only support a single namespace server.

To add a namespace server:

1. In the console tree of the DFS Management snap-in, right-click \\domain**Public**, and then click **Add Namespace Server**.

2. In **Namespace server**, type the name of another server to host the namespace, and then click **OK**.

After you finish this procedure, click the \\domain**Public** namespace in the console tree and review the contents of the **Namespace Servers** tab in the details pane, which should look similar to the following figure. Notice that two UNC paths are listed. The site of each namespace server is also displayed.

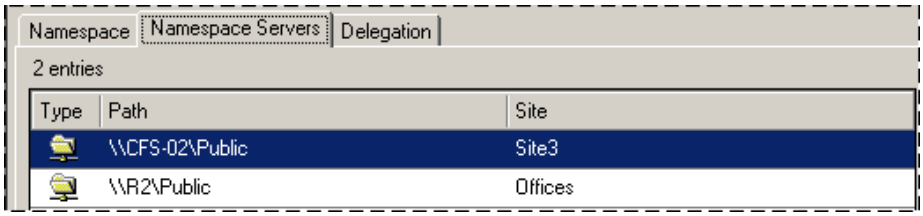

7.5.3. Delegate Management Permissions

You can delegate management permissions so that users who are not members of the Domain Admins group can create domain-based namespaces, and you can delegate management permissions so that users or groups can manage existing namespaces. In this section, you will delegate permissions to manage the namespace you created in the previous task.

To delegate permission to manage an existing namespace:

1. In the console tree of the DFS Management snap-in, right-click \\server_or_domain**Public**, and then click **Delegate Management Permissions**.

2. Type the name of a user or group that you want to manage the namespace, and then click **OK**.

After you finish this procedure, review the contents of the **Delegation** tab in the details pane. It should look similar to the following figure.

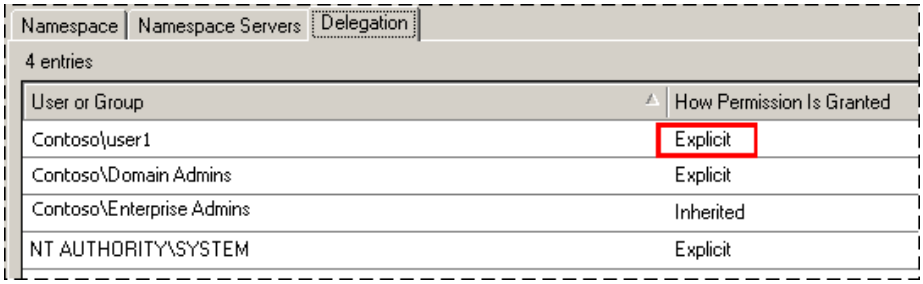

Notice that the user or group you added shows "Explicit" in the **How Permission Is Granted** column. "Explicit" means that you can remove the user or group from the delegation list by right-clicking the user or group, and then clicking **Remove**. Any users or groups that show "Inherited" have inherited management permissions from Active Directory, and you cannot remove them from the delegation list using the DFS Management snap-in.

Notes

To delegate the ability to create domain-based namespaces, see "Security requirements for creating and managing namespaces" in DFS Management Help.

7.5.4. Add Folders to Namespace

In this task, you add three folders to the namespace. Two of the folders will have folder targets. The hierarchy of the folders you will add is shown in the "Elements of a Namespace" figure earlier in this guide.

To create a folder named Software in the namespace:

1. In the console tree of the DFS Management snap-in, right-click \\server_or_domain**Public**, and then click **New Folder**.

2. In **Name**, type **Software**, and then click **OK**.

Note that the previous procedure creates a new folder in the namespace to build depth in the namespace hierarchy. You are not specifying the name of an existing folder, nor will you store data in this folder. This folder will not have folder targets that direct clients to other servers.

After you finish this procedure, the Software folder is added to the console tree as shown in the following figure. (You might need to double-click the \\server_or_domain**Public** root to display the Software folder.

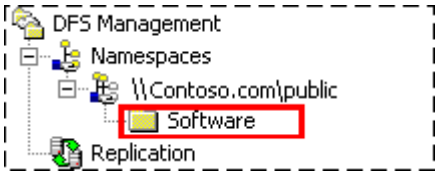

Next, you add two folders with targets to the namespace. You create one folder named Tools within the Software folder, and you create another folder named Training Guides directly under the root named Public.

To create a folder named Tools in the namespace:

1. In the console tree of the DFS Management snap-in, right-click the **Software** folder, and then click **New Folder**.
2. In **Name**, type **Tools**.
3. Click **Add** to add a folder target.
4. Click **Browse** to open the **Browse for Shared Folders** dialog box.
5. In **Server**, enter the name of the server that will host the Tools shared folder.
6. Click **New Shared Folder**.
7. In the **Create Share** dialog box, in the **Share name** box, type **Tools**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

After you finish this procedure, the Tools folder is added to the console tree as shown in the following figure. (You might need to double-click the Software folder to display the Tools folder.) Notice the icon next to the Tools folder and how it differs from the Software folder's icon. This icon appears next to all folders that have targets to differentiate them from folders that do not have targets.

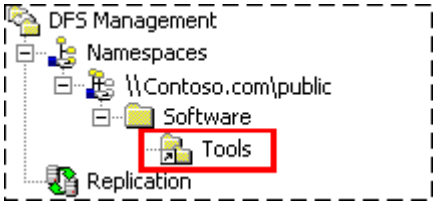

Now, select the Tools folder and review the contents of the **Folder Targets** tab in the details pane. Notice there is a single path shown. This means that only one server hosts the folder target that corresponds to the Tools folder. If that server becomes unavailable, the shared folder is also unavailable.

To increase the availability of the Tools folder, you can add a second folder target.

To add a second folder target to the Tools folder:

1. In the console tree of the DFS Management snap-in, right-click the **Tools** folder, and then click **Add Folder Target**.
2. Click **Browse** to open the Browse for Shared Folders dialog box.
3. In **Server**, enter the name of another server that will host the Tools shared folder. Be sure to enter a different server from the one you specified in the previous procedure.

4. Click **New Shared Folder**.

5. In the **Create Share** dialog box, in the **Share name** box, type **Tools**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

6. You are prompted to choose whether to create a replication group for these folder targets. For now, click **No**. You will enable DFS Replication on this folder in a later task.

To create a folder named Training Guides in the namespace:

1. In the console tree of the DFS Management snap-in, right-click \\server_or_domain**Public**, and then click **New Folder**.

2. In Name, type **Training Guides**.

3. Click **Add** to add a folder target.

4. Click **Browse** to open the Browse for Shared Folders dialog box.

5. In **Server**, enter the name of the server that will host the Training Guides shared folder.

6. Click **New Shared Folder**.

7. In the **Create Share** dialog box, in the **Share name** box, type **Training Guides**, and then enter the local path where you want the shared folder to be created. If the folder does not exist, you are prompted to create it. Click **OK** to close all dialog boxes.

When you finish these procedures, your namespace will look like the following figure.

7.5.5. Change How Target are Ordered in Referrals

In this task, you change settings that optimize how targets are ordered in referrals. If you are not familiar with referrals, a referral is a list of targets that a client computer receives from a domain controller or namespace server when the user accesses a namespace root or folder with targets in the namespace. The referral tells the client which servers host the associated root target or folder target. So, for example, when a client navigates to \\server_or_domain\Public, the client receives a root referral that contains a list of root targets on the namespace servers. When the client then navigates to the Tools folder, which has folder targets, the client receives a folder referral that contains a list of folder targets that correspond to the Tools folder.

When a client requests a referral, the Distributed File System service takes into account the site of the client and the site of the target and provides a referral with targets that are ordered according to the current referral ordering method. By default, targets in a client's site are listed first in a referral in random order, followed by a list of targets outside of the client's site, sorted by lowest cost.

To fine-tune how targets outside of a client's site are ordered, you can change the ordering method for an entire namespace or for individual folders with targets. Changing the ordering method is an important consideration in namespaces whose targets span sites. For example, there might be situations in which you want to prevent the client from accessing targets outside of its own site. If so, you can configure the namespace root or folder with targets so that clients receive referrals only for targets within their own site.

To further optimize how targets are listed in referrals, you can set target priority, which overrides the ordering method. For example, you can specify that a target is always first or last in a referral, regardless of the client's site, or you can specify that a target is always first or last among the targets that have the same connection cost. One common scenario for using target priority is when you have a "hot standby" server that is considered the server of last resort. In this scenario, you can specify that the standby server always appears last in referrals, and clients will fail over to this server only if all the other servers fail or become unavailable due to network outages.

In the following procedures, you verify the referral ordering method for the namespace and choose target priority of a folder target.

To verify the referral ordering method for a namespace:

1. In the console tree of the DFS Management snap-in, right-click \\server_or_domain**Public**, and then click **Properties**.

2. On the **Referrals** tab, in **Ordering method**, verify that **Lowest cost** is selected.

In the lowest cost ordering method, also called least expensive target selection or site costing in previous documentation, targets in a referral are ordered as follows:

1. Targets in the same site as the client are listed in random order at the top of the referral.

2. Targets outside of the client's site are listed in order of lowest cost to highest cost. Referrals with the same cost are grouped together and within each group the targets are listed in random order.

This method ensures that clients do not traverse expensive wide area network (WAN) links to access targets when lower-cost targets are available. This ordering method works in both stand-alone and domain-based namespaces, as long as all namespace servers and all domain controllers are running Windows Server 2003.

Notes

If you do not want clients to access folder targets outside of their site, you can override the ordering method for individual folders. To do this, right-click a folder with targets in the console tree, click **Properties**, click the **Referrals** tab, and then click **Exclude targets outside of the client's site**. Note that if no same-site targets are available, the client fails to access the folder because no folder targets are returned in the referral.

In the next procedure, you change the priority of one of the folder targets of the Tools folder.

To change the priority of a folder target:

1. In the console tree of the DFS Management snap-in, click the Tools folder.

2. In the details pane, on the Folder Targets tab, right-click one of the folder targets, and then click Properties.

3. On the Advanced tab, click Override referral ordering, and then click Last among all targets.

7.5.6. Rename and Move a Folder

You can use the DFS Management snap-in to rename folders or move folders to another location in the namespace. This is useful if you need to change a folder name or restructure the namespace.

In this task, you rename the Training Guides folder to Training Demos and move it to the Software folder. Currently, your namespace should look similar to the following figure.

To rename the Training Guides folder:

1. In the console tree of the DFS Management snap-in, right-click the **Training Guides** folder, and then click **Rename Folder**.

2. In the **Rename Folder** dialog box, in **New name**, type **Training Demos**.

To move the Training Demos folder:

- In the console tree of the DFS Management snap-in, click the **Training Demos** folder, and then drag it to the **Software** folder.

After you finish these procedures, your namespace should look like this:

7.5.7. Replicate a Folder in the Namespace Using DFS Replication

In this task, you enable DFS Replication on the Tools folder. If you recall from "7.5.4 Add Folders to the Namespace," you created two folder targets for the Tools folder. Because users can be directed to either one of the folder targets, you need to ensure that the contents of the folders are kept synchronized.

If you are familiar with File Replication service (FRS) in Windows Server 2003, you know that FRS is only supported in domain-based namespaces. In Windows Server 2003 R2, you can use DFS Replication in both stand-alone and domain-based namespaces. Therefore, you can complete this task regardless of the type of namespace you created in "7.5.1 Create a Namespace."

IMPORTANT

To perform this task, you need to have Active Directory deployed in your test lab environment, and you must be a member of the Domain Admins group or have been delegated the ability to create replication groups to perform this task.

IMPORTANT

After you complete this task, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and the long polling interval (60 minutes) on each member.

To enable DFS Replication on the Tools folder:

1. In the console tree of the DFS Management snap-in, right-click the **Tools** folder, and then click **Replicate Folder**.

2. Follow the steps in the Replicate Folder Wizard and supply the information in the following fields:

    - **Replication Group and Replicated Folder Name** - Accept the defaults.

    - **Replication Eligibility** - Accept the defaults.

    - **Primary Member** - If the folder targets are empty, choose either member. If both folder targets contain content, choose the member that has the most up-to-date content.

    - **Topology Selection** - Select **Full mesh**.

    - **Replication Group Schedule and Bandwidth** - Select **Replicate continuously using the specified bandwidth**.

    - **Review Settings and Create Replication Group** - Click **Create** to create the replication group.

    - **Confirmation** - Click **Close** to close the wizard.

    - **Replication Delay** - Click **OK** to close the dialog box that warns you about the delay in initial replication.

After you finish the previous procedure, navigate to the **Replication** node in the console tree. Notice that a new replication group has been created, as shown in the following figure.

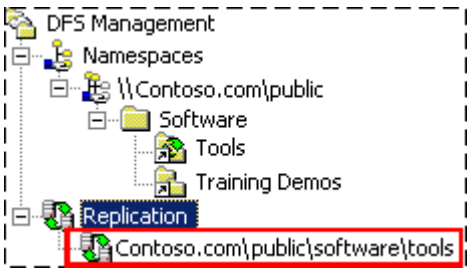

If you are not familiar with DFS Replication terminology, a replication group is a set of servers, known as members, that participates in the replication of one or more replicated folders. A replicated folder is a folder that is kept synchronized on each member. When you enable DFS Replication on a folder with targets, the servers that host the folder targets become members of the replication group, and the folder targets are associated with the replicated folder. The name of the replication group matches the namespace path (Contoso.com\Public\Software\Tools), and the name of the replicated folder matches the folder name (Tools).

From the **Replication** node, you can manage aspects of DFS Replication, such as the schedule and bandwidth usage, file and subfolder filters, and the topology (a framework of replication paths between members). On the **Replicated Folders** tab in the details pane, you can also view the namespace path that corresponds to the replicated folder, as shown in the following figure.

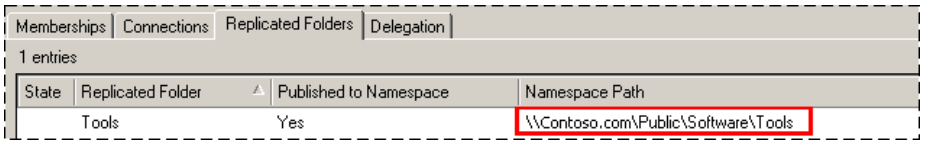

If you navigate back to the **Tools** folder in the **Namespaces** node, notice that the **Replication** tab in the details pane shows that the Tools folder is being replicated using DFS Replication.

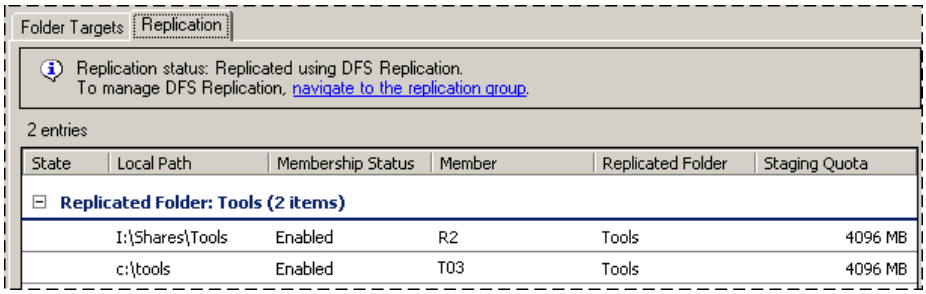

If one of the folders targets contained data when you enabled DFS Replication, you can verify that replication has completed by clicking the **Folder Targets** tab, right-clicking the folder target that initially held no data, and then clicking **Open in Explorer**. After the initial replication delay, the files in this folder target should match the files in the target that initially held the data.

Another way to view the status of replication is to create a diagnostic report. You will do this in the following task.

7.5.8. Create a Diagnostic Report

In this task, you create a diagnostic report to check the status of replication. The diagnostic report is an .html file that includes error and warning events, replication statistics, backlogged files, and other information for each member of the replication group.

To create a diagnostic report:

1. In the console tree of the DFS Management snap-in, under the **Replication** node, right-click the \\domain\Public\Software\Tools replication group, and then click **Create Diagnostic Report**.

2. Follow the steps in the Diagnostic Report Wizard and supply the information in the following fields:

    - **Path and Name** - Accept the defaults.

- **Members to Include** - Accept the defaults.
- **Options** - Ensure that **Yes, count backlogged files in this report** is selected, and also click the **Count the replicated files and their sizes on each member** check box.
- **Review Settings and Create Report** - Click **Create** to create the diagnostic report.
- **Confirmation** - The wizard closes automatically, and the diagnostic report appears.

Review the diagnostic report created for the Tools replication group. In particular, take a look at the following sections:

- Note the DFS Replication efficiency savings. This savings will change over time as files are added and changed.
- Review any errors or warnings, if any, for the members. These are typically event log errors that appear in the member's respective DFS Replication event log.
- In the informational section for each member, review the replicated folder status (the status will be "Normal" after initial replication is complete) and other information. Notice that the primary member will show different statistics from the non-primary member; this is because data originated from the primary member and replicated to the non-primary member during initial replication.

7.5.9. Browse the Namespace

In this task, you browse the namespace you created and view the referrals in the client's referral cache as you browse portions of the namespace. Viewing referrals cached on the client is useful in troubleshooting scenarios. The following procedures assume you are performing these tasks from a client computer running Windows XP or Windows Server 2003.

To browse to the namespace root and view the root referral:

1. Click **Start**, click **Run**, type \\server_or_domain**Public**, and then click **OK**. Windows Explorer opens and your view of the namespace looks similar to the following figure:

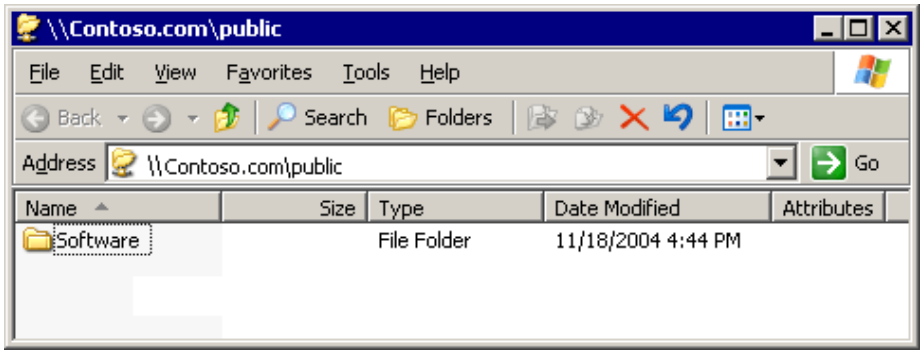

2. In Windows Explorer, click the **Folders** button to display the Public root in the folder tree.

3. In the folder tree, right-click **Public**, and then click **Properties** to open the **Properties** dialog box.

4. On the **DFS** tab, review the paths listed under **Referral list**. These are the root targets in the root referral that the client received when it accessed \\server_or_domain**Public**. These should match the root targets you created earlier in this guide. The target marked Active is the target currently connected to your client computer.

5. Click **OK** to close the dialog box.

To browse to the Tools and Training Demos folders and view their folder referrals:

1. In Windows Explorer, double-click the **Software** folder. You should see two folders, **Tools** and **Training Demos**.

2. Double-click the **Tools** folder to open it.

3. In the folder tree, right-click the **Tools** folder, and then click **Properties**.

4. On the **DFS** tab, review the paths listed under **Referral list**. These are the folder targets in the folder referral that the client received when it accessed \\server_or_domain**Public\Software\Tools**. These should match the folder targets you created earlier in this guide, and the target you marked as **Last among all targets**, when you set the target priority, should be listed last. The target marked Active is the target currently connected to your client computer.

5. Click **OK** to close the dialog box.

6. Click the **Training Demos** folder in the folder tree to open it.

7. Right-click the **Training Demos** folder in the folder tree, click **Properties**, and then click the **DFS** tab. Notice that only one folder target is listed in the referral list. Your client computer is currently connected to this folder target.

7.5.10. Test Failover

In this task, disable the network card or turn off the server that hosts one of the root targets for the \\domain**Public** namespace. Do the same for a server that hosts one of the folder targets for the Tools folder. After the network cards are disabled or the servers are turned off, repeat the procedures in "Task 9: Browse the Namespace." The procedures should work because another server continues to host the \\domain\Public namespace and the Tools folder.

7.6. Deploying DFS Replication (Step-by-Step Guide)

The tasks in this section walk you through the process of deploying DFS Replication, adding a member to a replication group, publishing a replicated folder in a namespace, and creating a diagnostic report.

7.6.1. Create a Multipurpose Replication Group and Two Replicated Folders

To enable DFS Replication, you use the New Replication Group Wizard to specify the members, topology, and default schedule and bandwidth for the replication group. In this task, you create a replication group named Data Distribution and two replicated folders named Antivirus Signatures and LOB Data.

IMPORTANT

When you create a new replication group, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and each member's long polling interval (60 minutes).

Before you enable replication, you will create two folders on one of the servers to be added to the replication group. You will then add files to the folders.

Create folders named Antivirus Signatures and LOB Data:

1. On one server, create two separate (non-overlapping) folders named Antivirus Signatures and LOB Data. Add some files to each folder but do not exceed the recommended limits described in "DFS Replication limits" in DFS Management Help.

2. Optionally, you can create the same folders on the second server. You can add the same files or different files from the primary member. If you add the same files, the files will be used for prestaging and will not be re-replicated. If you add files that don’t exist on the primary member, those files will be moved to the PreExisting folder on the second member. (This folder is in the DfsrPrivate folder under the replicated folder’s local path.) If you want to observe how the primary member's content becomes authoritative during initial replication, use updated versions of the files you added to the primary member. These updated files will be moved to the Conflict and Deleted folder on the non-primary members.

Next, create a replication group to replicate files between the two servers.

To create a replication group and two replicated folders:

1. In the console tree of the DFS Management snap-in, right-click the **Replication** node, and then click **New Replication Group**.

2. Follow the steps in the New Replication Group Wizard and supply the information in the following fields.

    - **Replication Group Type** - Select **Multipurpose replication group**.

    - **Name and Domain** - In **Name of replication group**, type **Data Distribution**.

    - **Replication Group Members** - Click **Add** to select at least two servers that will participate in replication. The servers must have the DFS Replication Service installed.

    - **Topology Selection** - Select **Full mesh**.

    - **Replication Group Schedule and Bandwidth** - Select **Replicate continuously using the specified bandwidth**.

    - **Primary Member** - Select the member that has the most up-to-date content that you want to replicate to the other member.

    - **Folders to Replicate** - Click **Add** to enter the local path of the LOB Data folder you created earlier on the first server. Use the name **LOB Data** for the replicated folder name. Repeat this procedure and enter the local path of the **Antivirus Signatures** folder.

    - **Local Path of LOB Data on Other Members** - On this page, you specify the location of the LOB Data folder on the other members of the replication group. To specify the path, click **Edit**, and then on the **Edit Local Path** dialog box, click **Enabled**, and then type the local path of the LOB Data folder.

    - **Local Path of Antivirus Signatures on Other Members** - On this page, you specify the location of the Antivirus Signatures folder on the other members of the replication group. To specify the path, click **Edit**, and then on the **Edit Local Path** dialog box, click **Enabled**, and then type the local path of the Antivirus Signatures folder.

    - **Review Settings and Create Replication Group** - Click **Create** to create the replication group.

    - **Confirmation** - Click **Close** to close the wizard.

    - **Replication Delay** - Click **OK** to close the dialog box that warns you about the delay in initial replication.

After you finish the New Replication Group Wizard, click the new replication group named **Data Distribution** located under the **Replication** node in the console tree as shown in the following figure:

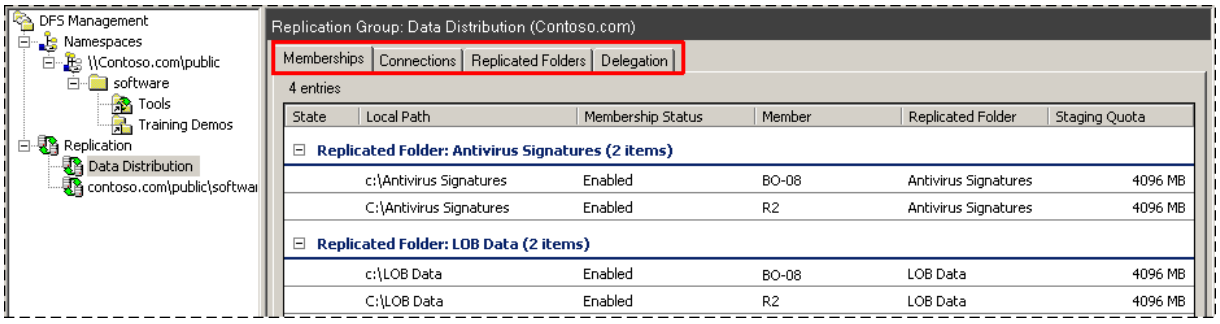

Notice the four tabs in the details pane: **Memberships**, **Connections**, **Replicated Folders**, and **Delegation**. Each of these tabs displays different details about the selected replication group, its members, and its replicated folders. Review the following details about each tab.

- On the **Memberships** tab, notice that entries on the tab are sorted by replicated folder and that there are two replicated folders listed. For example, the rows under the **Replicated Folder: Antivirus Signatures** heading are the members that host the Antivirus Signatures replicated folder. Double-click a member to view per-member, per-replicated folder properties on the **General**, **Replicated Folder**, and **Advanced** tabs. For example, on the **Advanced** tab, you can view the location and size of the staging folder and Conflict and Deleted folder on the selected member.

- On the **Connections** tab, two connections are listed. Each connection is a one-way replication path, so replication between two members requires two connections that replicate data in the opposite direction. Each connection has a schedule and other settings, such as a check box for enabling or disabling remote differential compression (RDC). Double-click a connection to view its settings.

- On the **Replicated Folders** tab, notice that two replicated folders are listed and that they are not published in a namespace. Double-click a replicated folder to view its properties, such as file and subfolder filters.

- On the **Delegation** tab, review the default users and groups granted permissions to manage the replication group. Any users or groups shown as "Inherited" have inherited management permissions from Active Directory, and you cannot remove them from the delegation list using the DFS Management snap-in.

7.6.2. Add a New Member to Replication Group

In this task, you add a third server to the Data Distribution replication group and specify where one of the two replicated folders, Antivirus Signatures, will be stored on the new member. You'll use a new feature in DFS Replication to specify that the LOB Data replicated folder is not replicated to the new member. You also will create a custom schedule that applies only to the connections to and from the new member.

IMPORTANT

Replication does not begin immediately on the new member. The DFS Replication settings for the new member must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency, the short polling interval (5 minutes) on the new member, and the long polling interval (60 minutes) on existing members.

To add a new member to the Data Distribution replication group:

1. In the console tree of the DFS Management snap-in, right-click the **Data Distribution** replication group, and then click **New Member**.

2. Follow the steps in the **New Member Wizard** and supply the information in the following fields.

    - **New Member** - Enter the name of the server to add to the replication group. The server must have the DFS Replication Service installed.

    - **Local Path of Replicated Folders** - Select the **Antivirus Signatures** replicated folder, click **Edit**, click **Enabled**, and then enter the local path of the replicated folder to be created on the new member. When you close the **Edit Local Path** dialog box, notice that the LOB Data replicated folder shows **<Disabled>**, which means that this replicated folder will not be replicated to the new member. Because you only want the Antivirus Signatures folder to be replicated to the new member, you can ignore the warning message that appears.

    - **Connections** - Under **Available members**, click a member, and then click **Add**. Repeat this step to add the second member. The new member will replicate directly with both existing members.

    - **Replication Schedule** - Select **Custom connection schedule**, and then click **Edit Schedule**. In the Edit Schedule dialog box, click **Details** to expand the schedule, and then select the entry that begins **Sunday 12:00 AM** and then click **Edit**. In the Edit Schedule dialog box, under **Bandwidth usage**, click **128 Mbps**.

    - **Review Settings and Create Member** - Click **Create** to add the new member to the Data Distribution replication group.

- **Confirmation** - Click **Close** to close the wizard.
- **Replication Delay** - Click **OK** to close the dialog box that warns you about the delay in initial replication.

After you finish the wizard, click **Data Distribution** in the console tree, and then review the contents of the **Connections** tab. It should look similar to the following figure:

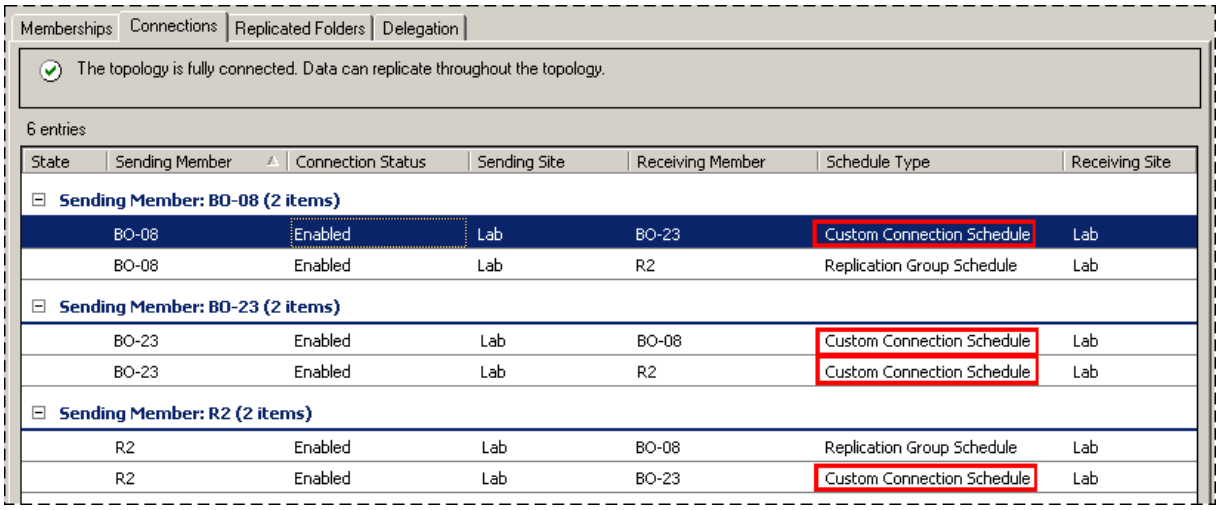

Notice that in the **Schedule Type** column, connections to and from the new member show **Custom Connection Schedule** instead of **Replication Group Schedule**. These show **Custom Connection Schedule** because you chose a custom schedule when you added the new member. Creating custom schedules for individual connections allows you to fine-tune the replication interval and bandwidth used when replicating to specific members. Although it isn't obvious in the user interface, each connection marked **Custom Connection Schedule** is a separate schedule. You can modify one schedule marked **Custom Connection Schedule**, but the other custom schedules are not affected.

Notes

To change how the items are grouped, click a column heading. For example, to group the items by schedule type, click the Schedule Type heading.

Entries marked **Replication Group Schedule** use the default replication schedule; this schedule is applied to all connections in the replication group that do not have a custom schedule. To modify the default replication schedule, right-click the Data Distribution replication group in the console tree, click **Properties**, and then click **Edit Schedule**. To change a connection schedule from a custom connection schedule to the replication group schedule or vice versa, on the **Connections** tab in the details pane, double-click the connection, click the **Schedule** tab, and then click **Replication group schedule** or **Custom connection schedule**.

7.6.3. Share and Publish Replicated Folders in a Namespace

When you created replicated folders in the previous tasks, you specified the local path of a folder on each member of the replication group. Unless the local path on each server was previously shared, users cannot access the replicated folders after they are created. To make replicated folders available to users, you must share them and, optionally, publish them in an existing namespace.

In this task, you publish the LOB Data replicated folder in the \\server_or_domain**Public** namespace that you created in the DFS Namespaces step-by-step section. If you did not complete the previous step-by-step section or do not have a namespace in your test lab, skip this procedure.

To share the LOB Data replicated folder and publish it in a namespace:

1. In the console tree of the DFS Management snap-in, under the **Replication** node, click the **Data Distribution** replication group.

2. In the details pane, click the **Replicated Folders** tab, right-click the **LOB Data** replicated folder, and then click **Share and Publish in Namespace**.

3. Follow the steps in the **Share and Publish Replicated Folder Wizard** and supply the information in the following fields.

    - **Publishing Method** - Select **Share and publish the replicated folder in a namespace**.

    - **Share Replicated Folders** - For each member that shows **[Shared Folder Needed]** in the **Action** column, select the member, and then click **Edit** to create the new shared folder and adjust shared folder permissions if necessary. If the **Action** column shows **Create shared folder: LOB Data or Existing Shared Folder**, you can click **Next**.

    - **Namespace Path** - In **Parent folder in namespace**, type \\server_or_domain**Public\Software**.

    - **Review Settings and Share Replicated Folder** - Click **Share** to share the replicated folders and publish the LOB Data replicated folder in the namespace.

    - **Confirmation** - Click **Close** to close the wizard.

After you finish the wizard, review the console tree and the **Replicated Folders** tab in the details pane. First, notice in the console tree that an LOB Data folder was added to the namespace and the folder icon indicates that the folder is replicated, as shown in the following figure.

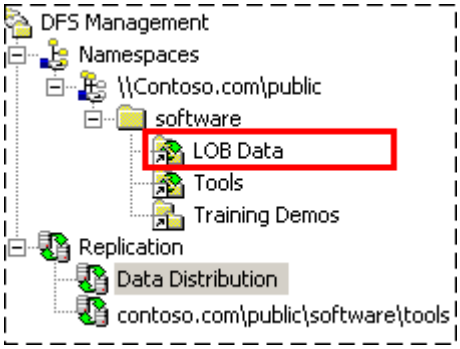

Next, review the namespace path listed in the **Replicated Folders** tab, which should look similar to the following figure.

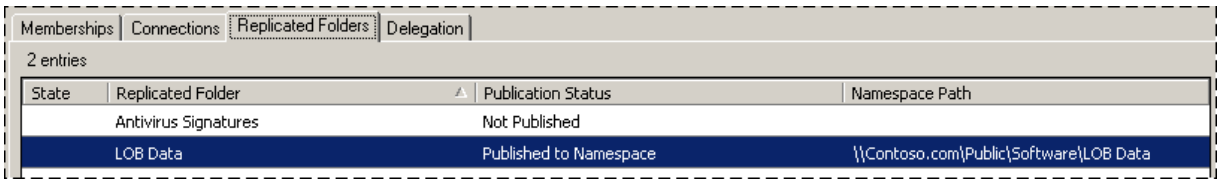

You can see that the LOB Data replicated folder is published in a namespace but the Antivirus Signatures replicated folder is not. To access the LOB Data folder in the namespace, in the Run dialog box, type \\server_or_domain**Public\Software\LOB Data**.

Notes

If you want to stop publishing the LOB Data replicated folder in the namespace, you can right-click the replicated folder and then click **Remove from Namespace**.

7.6.4. Create a Replication Group for Collection Purposes

In this task, assume that you have a hub server in a central hub or data center location and a branch server in a remote office. The branch server contains folders named Projects and Proposals that are very important to the branch office, but there is a concern that the backups performed at the branch office are performed incorrectly. You want to back up the Projects and Proposals folders from the data center to ensure that the backups are successful.

To accomplish this goal, you will set up a replication group for data collection purposes. This type of replication group consists of two members and one replicated folder for each folder that you want to back up from the hub server at the data center. The permissions that are set on the replicated folders on the branch server will be applied to the replicated folders on the hub server. You specify a single folder on the hub server under which subfolders for the replicated folders will be created. This allows you to back up multiple replicated folders from a single location on the hub server.

IMPORTANT

When you create a new replication group, replication does not begin immediately. The topology and DFS Replication settings must be replicated to all domain controllers, and each member in the replication group must poll its closest domain controller to obtain these settings. The amount of time this takes depends on Active Directory replication latency and each member's long polling interval (60 minutes).

To create a replication group to replicate the Projects and Proposals folders from a branch server to a hub server:

1. In Windows Explorer or from the command prompt, create a folder named **Projects** and a folder named **Proposals** on a server that will act as the branch server. The folders should be separate folders (that is, not nested in one another).

2. Add some data to the **Projects** and **Proposals** folders on the branch server.

3. In Windows Explorer or from the command prompt, create a folder named **Branch_Backups** on a server that will act as the hub server. Do not put data in this folder.

4. In the console tree of the DFS Management snap-in, right-click the **Replication** node, and then click **New Replication Group**.

5. Follow the steps in the New Replication Group Wizard and supply the information in following fields.

    - **Replication Group Type** - Select **Replication group for data collection**.

    - **Name and Domain** - In **Name of replication group**, type **Branch Backups**.

    - **Branch Server** - Type the name of a server that will act as the branch server.

    - **Replicated Folders** - Click **Add**. In the **Add Folder to Replicate** dialog box, type the local path of the Projects folder you created in Step 1. Repeat this step for the local path of the Proposals folder.

    - **Hub Server** - Type the name of a server that will act as the hub server. This is the server where you can back up the Projects and Proposals folders using backup software.

    - **Target Folder on Hub Server** - In **Target folder**, type the path of the folder you created in Step 3.

    - **Replication Group Schedule and Bandwidth** - Select **Replicate continuously using the specified bandwidth**.

    - **Review Settings and Create Replication Group** - Click **Create** to create the replication group.

    - **Confirmation** - Click **Close** to close the wizard.

- **Replication Delay** - Click **OK** to close the dialog box that warns you about the delay in initial replication.

After you finish the wizard, click the **Branch Backups** replication group in the console tree and view the **Memberships** tab in the details pane. Notice that two replicated folders were created, **Projects** and **Proposals**, as shown in the following figure:

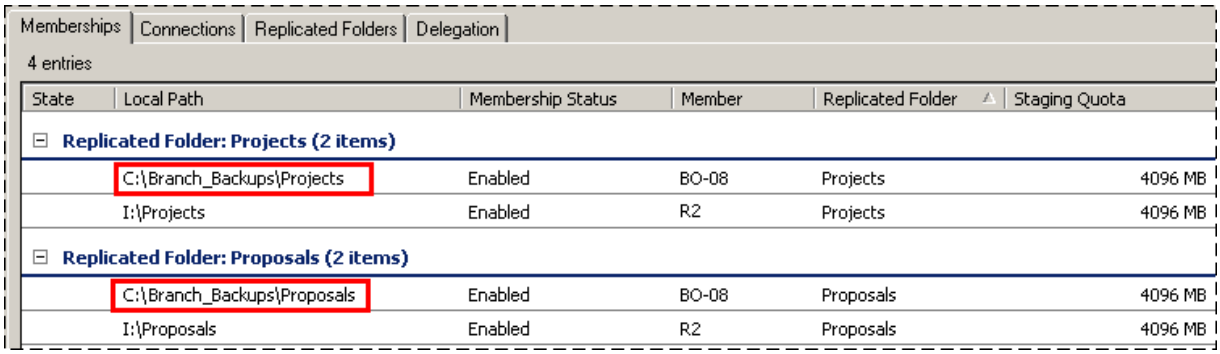

In the previous figure, notice that the path of the Projects and Proposals replicated folders on server BO-08 (the hub server) are both within the C:\Branch_Backups folder. This allows you to back up both replicated folders from a single location on the hub.

7.6.5. Create a Diagnostic Report

In this task, you create a diagnostic report to check the status of replication. The diagnostic report is an .html file that includes error and warning events, replication statistics, backlogged files, and so forth for each member of the replication group.

To create a diagnostic report

1. In the console tree of the DFS Management snap-in, under the **Replication** node, right-click the **Branch Backups** replication group, and then click **Create Diagnostic Report**.

2. Follow the steps in the Diagnostic Report Wizard and supply the information in the following fields.

   - **Path and Name** - Accept the defaults.

   - **Members to Include** - Accept the defaults.

   - **Options** - Ensure that **Yes, count backlogged files in this report** is selected, and also click the **Count the replicated files and their sizes on each member** check box.

   - **Review Settings and Create Report** - Click **Create** to create the diagnostic report.

   - **Confirmation** - The wizard closes automatically, and the diagnostic report appears.

Review the diagnostic report created for the Tools replication group. In particular, take a look at the following sections:

- Note the DFS Replication efficiency savings. This savings will change over time as files are added and changed.

- Review any errors or warnings, if any, for the members. These are typically event log errors that appear in the member's respective DFS Replication event log.

- In the informational section for each member, review the replicated folder status (the status will be "Normal" after initial replication is complete) and other information. Notice that the primary member will show different statistics from the non-primary member; this is because data originated from the primary member and replicated to the non-primary member during initial replication.

This Page Intentionally Left Blank

8. Microsoft Services for Network File System (MSNFS)

Microsoft Services for Network File System is a comprehensive software package 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 manages tasks on both Windows and UNIX platforms. Tasks include creating NFS exports from Windows and administering user name mappings.

This chapter discusses networking features in Microsoft Services for Network File System (MSNFS).

8.1. MSNFS Features

MSNFS is an update to the NFS components that were previously available in Services for UNIX 3.5.

MSNFS includes the following new features:

- Updated administration snap-in—MSNFS Administration
- Active Directory Lookup—The Identity Management for UNIX Active Directory schema extension, available in Microsoft Windows Server 2003 R2, includes UNIX user identifier (UID) and group identifier (GID) fields, which enables Server for NFS and Client for NFS to look up Windows-to-UNIX user account mappings directly from Active Directory. Identity Management for UNIX simplifies Windows-to-UNIX user account mapping management in Active Directory.
- Enhanced server performance—Microsoft Services for NFS includes a file filter driver, which significantly reduces common server file access latencies.
- UNIX special device support—Microsoft Services for NFS supports UNIX special devices (mknod).
- Enhanced UNIX support—Microsoft Services for NFS now supports the following versions of UNIX:
    - HewlettPackardHP-UXversion11i
    - IBMAIXversion 5L 5.2
    - RedHat Linux version9
    - Sun Microsystems Solaris version 9

The following features that were previously available in Services for UNIX 3.5 are not included in MSNFS:

- Gateway for NFS
- Server for PCNFS
- All PCNFS components of Client for NFS

8.1.1. UNIX Identity Management

Identity Management for UNIX makes it easy to integrate users of Windows operating systems into existing UNIX environments. It provides manageability components that simplify network administration and account management across both platforms.

With Identity Management for UNIX, the administrator can:

- Manage user accounts and passwords on Windows and UNIX systems using Network Information Service (NIS).
- Automatically synchronize passwords between Windows and UNIX operating systems.

UNIX Identity Management consists of the following components:

- Administration components
- Password synchronization
- Server for NIS

The UNIX Identity Management component is not enabled by default on the storage server. To install this component:

1. Access **Add/Remove Programs**.
2. Select **Add/Remove Windows Components > Active Directory Services > Details**.
3. Install **Identity Management for Windows**.

8.2. Microsoft Services for NFS usage scenarios

Microsoft Services for NFS enables you to support a mixed environment of Windows-based and UNIX-based operating systems. With Microsoft Services for NFS, you can also update your company's computers while supporting older technology during the transition phase. The following scenarios are examples of how enterprises can benefit from deploying Microsoft Services for NFS.

- **Enable UNIX-based client computers 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 features in Windows Server 2003 R2 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 Microsoft Services for NFS to enable UNIX clients that are running NFS software to access these computers. All of your UNIX clients will be able to access resources using the NFS protocol without additional configuration.

- **Enable 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 Microsoft Services for NFS to enable computers running Windows Server 2003 R2 to access these resources when the file servers are running NFS software.

8.3. Microsoft Services for NFS components

Microsoft Services for NFS includes the following three main components:

- **User Name Mapping**. User Name Mapping associates user accounts between Windows and UNIX domains. In a heterogeneous network, users have separate Windows and UNIX security accounts. Historically, users had to provide a different set of credentials to access files and other resources across system boundaries. To address this issue, User Name Mapping associates Windows and UNIX user names so users logged onto the UNIX domain can access NFS shared resources on Windows Server 2003 R2 without logging on separately to the Windows domain, and vice-versa.

- **Server for NFS**. Normally, a UNIX-based 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-based and UNIX-based computers.

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

8.4. Microsoft Services for NFS administrative tools

Microsoft Services for NFS provides a Microsoft Management Console (MMC) snap-in for administration, as well as several command-line tools.

8.4.1. Microsoft Services for NFS snap-in

With the Microsoft Services for NFS snap-in, you can administer each installed component of Microsoft Services for NFS. When you open the snap-in, the components installed on the local computer are available to administer.

To open Microsoft Services for Network File System:

- Click **Start**, point to **Administrative Tools**, and click **Microsoft Services for Network File System**.

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure. You can get help for an item in this snap-in by right-clicking the item and then clicking **Help**.

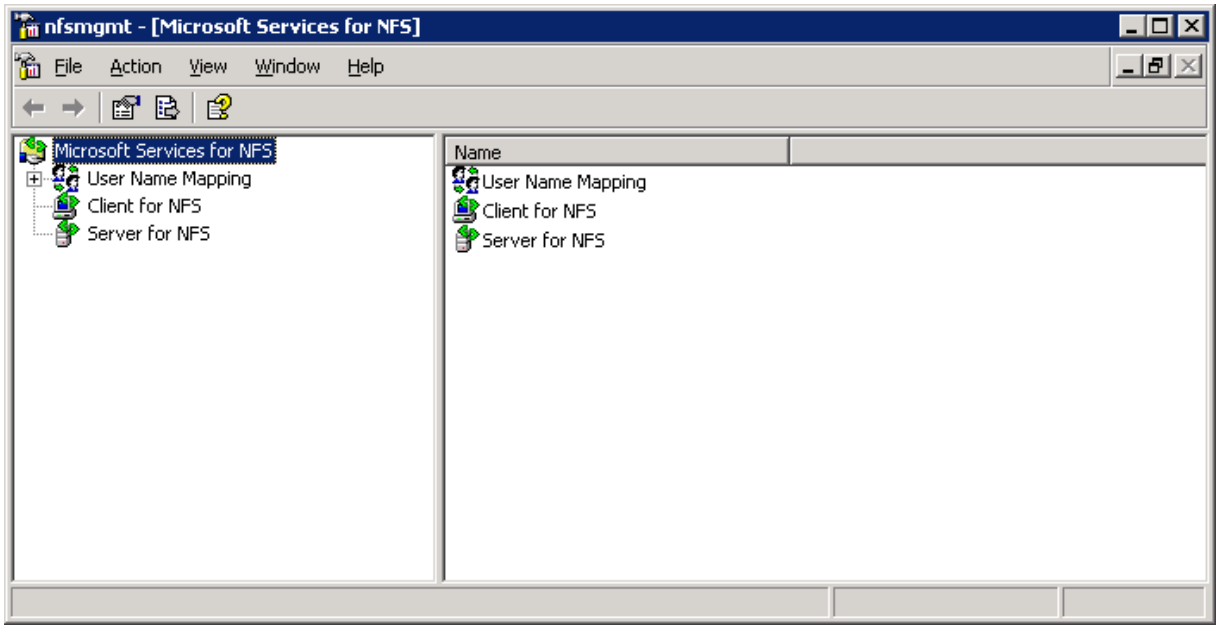

8.4.2. Microsoft Services for NFS command-line tools

Microsoft Services for NFS provides the following Windows command-line administration tools. To run a tool, type its name at the command prompt. For information about the available parameters, at the command prompt, type `toolname/?`.

- `mapadmin`. Administers User Name Mapping.
- `mount`. Mounts NFS network shares.
- `nfsadmin`. Manages Server for NFS and Client for NFS.
- `nfsshare`. Controls NFS shares.
- `nfsstat`. Displays or resets counts of calls made to Server for NFS.
- `showmount`. Displays mounted file systems exported by Server for NFS.

- `umount`. Removes NFS-mounted drives.

8.5. Test Scenario

This test scenario requires you to deploy Microsoft Services for NFS in a lab environment to assess how this technology would function if deployed in your production environment. The instructions provided in this document will help you:

- Set up User Name Mapping to map Windows and UNIX user accounts so that users can log on to either Windows or UNIX domains and access resources in both environments.

- Create an NFS shared resource on a computer running Windows Server 2003 R2 and Server for NFS that can be mounted and used by a UNIX computer.

- Create an NFS shared resource on a UNIX file server that can be mounted and used by a computer running Windows Server 2003 R2 and Client for NFS.

With the Microsoft Services for NFS snap-in, you can administer each installed component of Microsoft Services for NFS. When you open the snap-in, the components

8.6. Steps for Deploying and Testing Microsoft Services for NFS

This section describes how to set up a basic test environment for Microsoft Services for NFS. It discusses how to install and configure the Microsoft Services for NFS components and how to test the deployment.

8.6.1. Reviewing system requirements for Microsoft Services for NFS

Microsoft Services for NFS can be installed on computers running any edition of Windows Server 2003 R2. The three main components of Microsoft Services for NFS – User Name Mapping, Server for NFS, and Client for NFS – can be installed on the same computer or on separate computers.

IMPORTANT

Before installing Microsoft Services for NFS, you must remove any previously installed NFS components, such as NFS components that were included with Services for UNIX. We recommend that you back up or make a record of your configuration before removing NFS components, so that you can restore the configuration on Microsoft Services for NFS.

You can use Microsoft Services for NFS with UNIX computers which are running NFS client or server software which complies with version 2 or version 3 of the NFS protocol. NFS version 2 is defined in RFC 1094 and NFS version 3 is definied in RFC 1813.

Notes

By default, Server for NFS supports UNIX client computers using NFS version 2 or version 3. You can override this, however, and configure Server for NFS to allow access only to clients running NFS version 2. For instructions, see "Configuring Server for NFS" in the Microsoft Services for NFS Help. Client for NFS supports both versions, and this is not configurable.

8.6.2. Setting up the environment for Microsoft Services for NFS

The next step is to set up the environment for Microsoft Services for NFS by deploying computers and creating user accounts for testing.

8.6.2.1. Deploy computers

You need to deploy the following computers and connect them on a local area network (LAN):

- One or more computers running Windows Server 2003 R2 on which you will install the three main Microsoft Services for NFS components: User Name Mapping Server, Server for NFS, and Client for NFS. You can install the components on the same computer or on different computers. Installation instructions for installing all Microsoft Services for NFS components are provided later in this document.

- One or more UNIX computers running NFS client and NFS server software. The computer running the NFS client will access a Windows NFS shared resource provided by Server for NFS. The computer running NFS server will host a UNIX NFS shared resource, which will be accessed by a computer running Windows Server 2003 R2 and Client for NFS. You can install the NFS client and NFS server software on the same computer or on different computers.

- A Windows Server 2003 domain controller running at the Windows Server 2003 functional level. The domain controller will provide user authentication information for the Windows environment. Or, if you prefer, you can use local user accounts.

- A Network Information Service (NIS) server to provide user authentication information for the UNIX environment. Or, if you prefer, you can use Password and Group files that are stored on the computer running the User Name Mapping service.

The following diagram illustrates a simple test configuration for Microsoft Services for NFS. It includes a single Windows Server 2003 domain controller and a single NIS server. All Microsoft Services for NFS components are installed on a single computer running Windows Server 2003 R2, and NFS client and NFS server software is installed on a single UNIX computer.

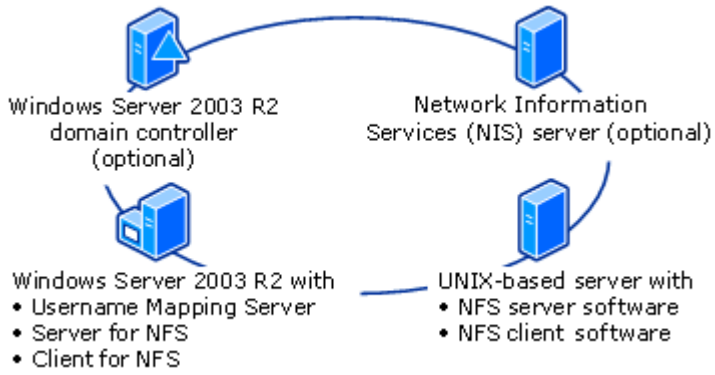

8.6.3. Create test user accounts

For the purposes of this test, you can create several fictitious users. For each user, you can create one Windows security account and one UNIX security account, giving the two accounts different user names. You can later use these accounts to test the advanced mapping feature of Microsoft Services for NFS. Advanced mapping allows you to map a given user's credentials between Windows and UNIX, even when the user name is different.

Notes

The alternative to advanced mapping is simple mapping. You can use simple mapping when Windows and UNIX user names for each user are the same. For more information about simple maps, see **User Name Mapping Administration in Network File System help**.

You can create the Windows user accounts on the Windows Server 2003 R2 domain controller. Or if you prefer, you can create local user accounts on each Windows-based computer in the deployment. For instructions on configuring user accounts, see your Windows Server 2003 R2 documentation.

You can create the UNIX user accounts either on the NIS server or in UNIX `/etc/passwd` and `/etc/group` files. For instructions on creating NIS user accounts, see the documentation for your NIS server. For instructions on creating `/etc/passwd` and `/etc/group` files, see the documentation for your UNIX operating system.

8.6.4. Installing Microsoft Services for NFS

You need to install Microsoft Services for NFS components on a computer running Windows Server 2003 R2. These instructions assume that you are installing all of the components on a single computer.

IMPORTANT

Before installing Microsoft Services for NFS, you must remove any previously installed NFS components, such as NFS components that were included with Services for UNIX. We recommend that you back up or make a record of your configuration before removing NFS components so that you can restore your settings on Microsoft Services for NFS.

To install Microsoft Services for NFS:

1. On the computer on which you want to install Microsoft Services for NFS, install Windows Server 2003 R2.

2. In Control Panel, double-click **Add or Remove Programs**.

3. Click **Add/Remove Windows Components**, click **Other Network File and Print Services**, and then click **Details**.

4. Click **Microsoft Services for NFS**, and then click **Details**.

5. Select **User Name Mapping**, **Server for NFS**, or **Client for NFS**, and then click **OK**.

Notes

When you select User Name Mapping, Server for NFS, or Client for NFS, the Windows Component Wizard will also select the appropriate combination of supporting subcomponents.

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority. If the computer is joined to a domain, members of the Domain Admins group might be able to perform this procedure. As a security best practice, consider using Run as to perform this procedure.

8.6.5. Configuring NFS authentication

The required configuration for this test uses a Windows Server 2003 domain controller or later running at the Windows Server 2003 functional level. For security reasons, we recommend installing Windows Server 2003 Service Pack 1 (SP1) and all the latest security updates.

8.6.6. Configuring User Name Mapping

The next step is to configure User Name Mapping and set up mapping between the Windows and UNIX user accounts that you created earlier. For the purposes of this test, we will configure advanced mapping. You must use advanced mapping when each user's account name in Windows is different than his or her account name in UNIX. You may use simple mapping when each user's account name in Windows is the same as his or her account name in UNIX.

8.6.6.1. Specify where UNIX user and group information is stored

You need to specify where UNIX user and group information is stored, either on a NIS server or in `/etc/passwd` and `/etc/group` files. If it is stored in password and group files, you must copy the files to the computer running User Name Mapping. For security reasons, you should only allow administrators access to the files. You must also specify the location of the files, so that User Name Mapping can access them.

To specify where UNIX user and group information is stored:

1. On the computer running User Name Mapping, open Microsoft Services for NFS. To open Microsoft Services for NFS, click **Start**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.

2. Right-click **User Name Mapping**, and then click **Properties**.

3. On the **UNIX User Source** tab, select the method used for storing UNIX user and group information: **Use Network Information Service (NIS)** or **Use Password and Group files**.

4. If you selected User Password and Group files, do the following:

    - In **Password file path and name**, type the full path of the `/etc/passwd` file.

    - In **Group file path and name**, type the full path of the `/etc/group` file.

5. Specify the synchronization interval for User Name Mapping to synchronize UNIX user and group information and click **Apply**.

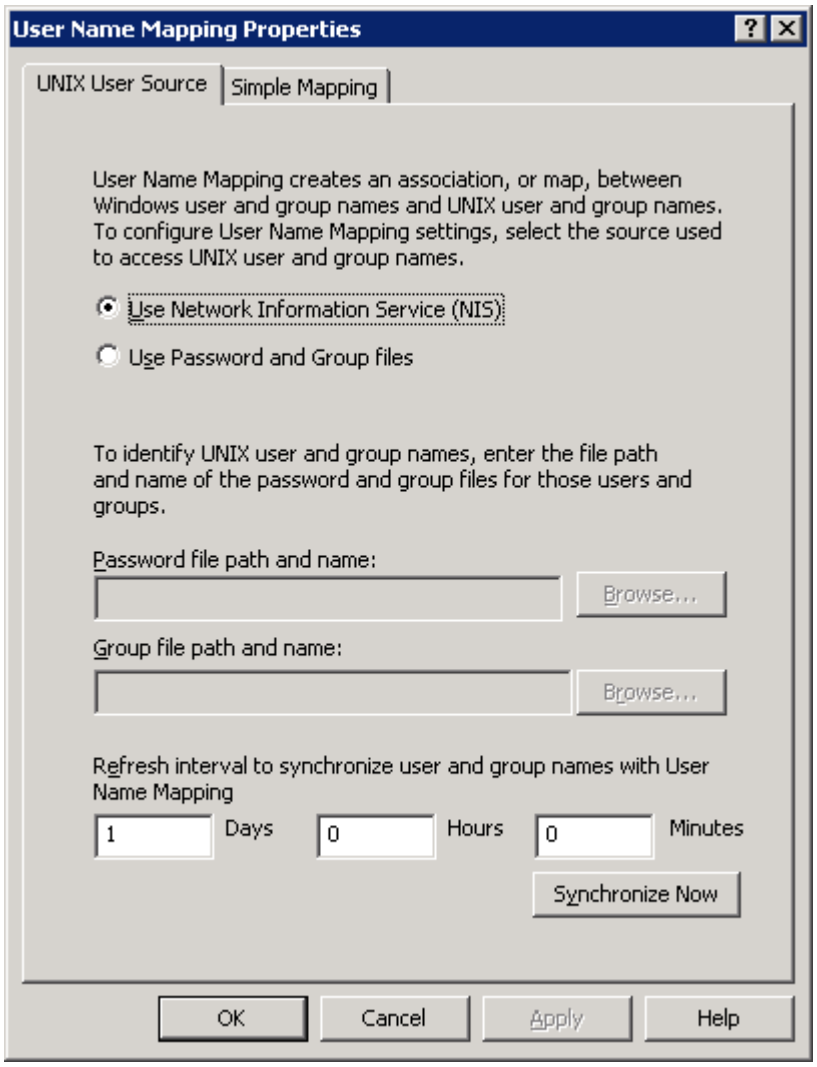

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

8.6.6.2. Edit the .maphosts file

Next, you need to add entries to the `.maphosts` file, which is installed with User Name Mapping. This file controls which computers on the network can access User Name Mapping, so you need to add an entry for each computer you are using in this test that is running either Server for NFS or Client for NFS.

You can edit the `.maphosts` file in a text editor. The `.maphosts` file is located in `%windir%\msnfs` on the computer running User Name Mapping.

The `.maphosts` file contains a list of one or more of the elements described in the following table, each on a separate line. The order of the elements is important because, when matching a computer making a request against the list, User Name Mapping searches from the top down until it finds a match.

Element	Description
`Host`	Specifies one or more computers that can access User Name Mapping. The `host` can be specified by an Internet Protocol (IP) address (IP version 4), or a host name that resolves to one or more IP addresses.
`host –`	Specifies one or more computers that are denied access to User Name Mapping. The `host` can be specified by an IP address (IP version 4), or a host name that resolves to one or more IP addresses. Note that there must be at least one blank space between `host` and the dash (–).
`+`	Specifies that all computers can access User Name Mapping unless disallowed by an earlier entry in the list. All entries in the list following this element are ignored.
`–`	Specifies that all computers are disallowed access to User Name Mapping unless allowed by an earlier entry in the list. All entries in the list following this element are ignored.

Examples
A `.maphosts` file with the following contents allows access only by computers named R2_host, UNIX_host, Windows_DC, and NIS_host:

```
R2_host
UNIX_host
Windows_DC
NIS_host
-
```

A `.maphosts` file with the following contents allows access by all computers except computers named rogue_host and malicious_wks:

```
rogue_host -
malicious_wks -
+
```

In the next example, rogue_host is denied access even though it appears in the list without a minus sign (**–**) because it follows a minus sign on its own line:

```
R2_host
UNIX_host
Windows_DC
NIS_host
-
rogue_host
```

8.6.6.3. Create a user map

The next step in setting up User Name Mapping is to create a user map that associates the Windows user name of each user to his or her UNIX user name.

To create a user map:

1. On the computer running User Name Mapping, open Microsoft Services for NFS. To open Microsoft Services for NFS, click **Start**, point to **Administrative Tools**, and then click **Microsoft Services for Network File System**.

2. In the console tree, expand **User Name Mapping**, right-click **User Maps**, and then click **Create Map**.

3. In the Windows domain list, click the domain for which you want to map user names, and then click **List Windows Users**.

4. If you are obtaining information about UNIX user names from an NIS server, type the NIS domain name in **NIS domain name**. If the NIS server is in a different subnet than the User Name Mapping server, type the DNS name or IP address of the NIS server in **NIS server name**. Otherwise, leave the **NIS Server name** box blank.

5. Click **List UNIX Users**.

6. In the **Windows Users** list, click the user name to map.

7. In the **UNIX Users** list, click the corresponding user name to map and click **Add**.

8. Repeat steps 6 and 7 for each user account to map and click **Apply**.

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

Notes

You can only map one user name to another user name at a time. If you want to map multiple user names with a single user name, you must repeat steps 6 and 7 in this procedure for each additional user name to map.

8.6.6.4. Create a group map

The last step in setting up a User Name Mapping server is to create a group map that maps Windows groups to UNIX groups.

To create a group map:

1. On the computer running the User Name Mapping service, open Microsoft Services for NFS. To open Microsoft Services for NFS, click **Start**, point to **Administrative Tools**, and then click Microsoft **Services for Network File System**.

2. In the console tree, expand **User Name Mapping**, right-click **Group Maps**, and then click **Create Map**.

3. In the Windows domain list, click the domain for which you want to map groups, and then click **List Windows Groups**.

4. Click **List UNIX Groups**.

5. In the **Windows Groups** list, click the Windows group to map.

6. In the **UNIX Groups** list, click the UNIX group to map to, and then click **Add**.

7. Repeat steps 5 and 6 for each group to map and click **Apply**.

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

8.6.6.5. Restart the User Name Mapping service

When you have finished all the steps to configure User Name Mapping that are described in this section, you need to restart the User Name Mapping service.

To restart the User Name Mapping service:

1. On the computer running the User Name Mapping service, click **Start**, point to **Administrative Tools**, and click **Services**.

2. In the list, right-click **User Name Mapping**, click **Stop**, wait for the service to stop, and then click **Start**.

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

8.6.7. Specifying the User Name Mapping server

On each computer running a component of Microsoft Services for NFS, you need to specify the name of the computer running the User Name Mapping service. This is necessary even in a simple installation where all Microsoft Services for NFS components are installed on the same computer.

To specify the User Name Mapping server:

1. On a computer running one or more components of Microsoft Services for NFS, open Microsoft Services for NFS. To open Microsoft Services for NFS, click **Start**, point to **Administrative Tools**, and then click Microsoft **Services for Network File System**.

2. In the console tree, right-click **Microsoft Services for Network File System** and click **Properties**.

3. In **User Name Mapping Server**, type the name of the computer running the User Name Mapping service, and then click **OK**.

4. Repeat these steps on each computer that is running one or more components of Microsoft Services for NFS.

Notes

To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

8.6.8. Creating an NFS shared folder

The next step is to use NFS sharing to create an NFS shared folder on the computer running Server for NFS. You can later mount this shared folder on a UNIX client and create a test file on it.

To create a shared folder using NFS sharing:

1. On the computer running Server for NFS, create a folder to use as the NFS shared folder.

2. Right-click the folder you created and click **NFS Sharing**.

3. Select **Share this folder**.

4. If you want to allow anonymous access, select **Allow anonymous access**.

5. Click **Permissions**, click **Add**, and then do either of the following:

    - In the **Names** list, click the clients and groups you want to add and click **Add**.

    - In the **Add Names** box, type the names of clients or groups you want to add, separating names in the list with a semicolon (;).

6. In the **Type of Access** list, click the type of access you want to allow the selected clients and groups.

7. Select **Allow Root Access** if you want a user identified as root to have access other than as an anonymous user. By default, the user identifier (UID) root user is coerced to the anonymous UID.

8. In the **Encoding list**, click the type of directory name and file name encoding to be used for the selected clients and groups.

9. Click **OK** twice, and then click **Apply**.

Notes
To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

Notes
To see a list of the members of a group, in the Names list, click a group, and then click Members.

8.6.9. Specifying default permissions for new files and folders

You can specify the default permissions that will be applied to any file or folder created on an NFS shared resource by the computer running Client for NFS. You can assign Read, Write, and Execute permissions to Owner, Group, and Others.

- **Owner**. The person creating the file. By default, Owner has Read, Write, and Execute permissions.
- **Group**. The primary group of the person creating the file. By default, Group has Read and Execute permissions.
- **Others**. Other file system users (equivalent to Everyone in Windows). By default, Others have Read and Execute permissions.

To specify default file permissions:

1. On the computer running Client for NFS, open Microsoft Services for NFS. To open Microsoft Services for NFS, click **Start**, point to **Administrative Tools**, and then click Microsoft **Services for Network File System**.
2. In the console tree, right-click **Client for NFS** and click **Properties**.
3. On the **File Permissions** tab, select the default file permissions to apply to each new file and folder created by this computer, and then click **OK**.

Notes
To perform this procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

8.6.10. Configuring Windows Firewall

After you install Microsoft Services for NFS, you must configure Windows Firewall to enable external computers to access the Microsoft Services for NFS services.

8.6.10.1. Open ports

On the computer(s) running User Name Mapping and Server for NFS, you must open ports in Windows Firewall. On a computer and running only User Name Mapping, you only need to open the portmapper port. On a computer running Server for NFS, you must open all of the ports listed in the following table.

To open ports in Windows Firewall:

1. On a computer running the User Name Mapping service or Server for NFS, click **Start**, click **Run**, type `firewall.cpl`, and then click **OK**.
2. Click the **Exceptions** tab, and then click **Add Port**.
3. In **Name**, type the name of a port to open, as listed in the following table.
4. In **Port number**, type the corresponding port number.
5. Select **TCP** or **UDP** and click **OK**.
6. Repeat steps 2 through 5 for each port to open, and then click **OK** when finished

Notes

Depending on your requirements, you may need to open Transmission Control Protocol (TCP) ports, User Datagram Protocol (UDP) ports, or both TCP ports and UDP ports. For testing purposes, we recommend that you open both TCP and UDP transports for all protocols.

Microsoft Services for NFS component	Port to open	Protocol	Port
User Name Mapping and Server for NFS	Portmapper	TCP, UDP	111
Server for NFS	Network Status Manager	TCP, UDP	1039
Server for NFS	Network Lock Manager	TCP, UDP	1047
Server for NFS	NFS Mount	TCP, UDP	1048
Server for NFS	Network File System	TCP, UDP	2049

8.6.10.2. Add mapsvc.exe to the exception list

On the computer running User Name Mapping, you must add the `Mapsvc.exe` program to the Windows Firewall exception list.

To add Mapsvc.exe to the exception list:

1. On the computer running the User Name Mapping service, click **Start**, click **Run**, type `firewall.cpl`, and then click **OK**.
2. Click the **Exceptions** tab, and then click **Add Program**.
3. Click **Browse**, click `mapsvc.exe`, and then click **Open**. By default, this file is located in `%windir%\System32`.
4. Fore testing purposes, click **Change scope**, select **Any computer**, and then click **OK**.

5. Click **OK** two times.

8.6.10.3. Enable file and printer sharing for administration tools

On the computer hosting the Microsoft Services for NFS snap-in and Microsoft Services for NFS command-line tools, you must enable file and printer sharing in Windows Firewall.

To enable file and printer sharing:

1. On a computer running Microsoft Services for NFS, click **Start**, click **Run**, type `firewall.cpl`, and then click **OK**.

2. Click the **Exceptions** tab, select the **File and Printer Sharing** check box, and then click **OK**.

3. Repeat these steps on each computer running Microsoft Services for NFS.

8.6.11. Testing your deployment

Now that everything is set up, you can test your deployment to verify its functionality. The following are some suggested basic tests.

8.6.11.1. Test 1: On the computer running Client for NFS, map a drive letter to a UNIX-based NFS shared resource.

The test is successful if you can map the drive and view the test file on the NFS shared resource from the computer running Client for NFS.

To map a drive letter to a UNIX-based NFS shared resource:

1. On a UNIX-based server running NFS software, create an NFS shared resource. Create a test file on the shared resource.

2. Log on to the computer running Windows Server™ 2003 R2 and Client for NFS with one of the Windows user accounts that you created for this test.

3. Open Windows Explorer (My Computer) and on the **Tools** menu, click **Map Network Drive**.

4. Type either the UNIX-style server and shared resource name (`hostname://sharedresourcename`) or the Universal Naming Convention (UNC) path of the NFS shared resource on the UNIX file server, and then click OK.

8.6.11.2. Test 2: On the computer running Client for NFS, create a test file and verify its permissions.

The test is successful if you can create a new document, and its ownership and permission match the default file permissions that you had specified.

To create a test file and verify its permissions:

1. Log on to the computer running Client for NFS with one of the Windows user accounts that you created for this test, and open the NFS shared resource that you used in Test 1.

2. Right-click in the file list, point to **New**, and then click **Text Document**.

3. Type a name for the file. Do not use spaces.

4. Right-click the file, click **Properties**, and then click **NFS Attributes**.

5. Verify that the NFS attributes match the default attributes that you specified earlier, as described in "Specifying default permissions for new files and folders." Also verify that the Owner UID and Group UID are correct.

8.6.11.3. Test 3: On a UNIX client computer, mount the Windows NFS shared resource.

The test is successful if you can mount the NFS shared resource.

To mount the Windows NFS share:

- In a command shell on a UNIX client running NFS client software, type:

- `mount hostname or IP:/sharename mountpoint`

Variable	Description
`Hostname or IP`	The name of the computer running Server for NFS, on which you previously created an NFS shared resourced, as described in "Creating an NFS shared folder."
`Sharename`	The name of the NFS shared resource.
`mountpoint`	The point in the file system where the command will mount the NFS shared resource, for example, /home/username/testshare.

8.6.11.4. Test 4: On a UNIX client, create a test file and verify the file permissions match, from both Windows and UNIX.

The test is successful if you can create the text file and the file permissions match from both Windows and UNIX

To create a test file and verify the file permissions match from both Windows and UNIX:

1. On the same UNIX client that you used in Test 3, create a text file by using a simple text editor. Save the file to the NFS shared resource that you mounted in Test 3.

2. On the computer running Server for NFS and hosting the NFS shared resource, open My Computer and browse to the NFS shared resource.

3. Right-click the file, click **Properties**, and then click **Security**.

4. Compare the file permissions reported through Windows with the file permissions reported through the same UNIX client you used in Test 3.

8.7. Using Remote Desktop for MSNFS

Windows Remote Desktop is available for remote administration of Services for UNIX. This service let users connect to machines, log on, and obtain command prompts remotely. See the following Table for a list of commonly used commands.

Caution

Two open sessions of Remote Desktop are allowed to operate at the same time. After completing an application do not use the window close feature () to close that session of Remote Desktop. Click Start/Log Off Administrator to exit Remote Desktop.

8.7.1. Using Remote Desktop

Microsoft Remote Desktop can be used to remotely access the NAS appliance desktop. This provides the administrator flexibility to automate setups and other tasks. Services for NFS file-exporting tasks and other Services for NFS administrative tasks can be accomplished using Remote Desktop to access the Services for NFS user interface from the NAS Desktop or from a command prompt.

The following Table describes some common Services for NFS commands.

Command	Function
`nfsstat /?`	Learn about viewing statistics by NFS operation type
`showmount /?`	View the format of the command to display NFS export settings on NFS servers
`showmount -a`	View users who are connected and what they currently have mounted
`showmount -e`	View exports from the server and their export permissions
`rpcinfo /?`	Learn how to display Remote Procedure Call (RPC) settings and statistics
`mapadmin /?`	View how to add, delete, or change user name mappings
`nfsshare /?`	Learn how to display, add, and remove exported shares

9. Using iSCSI Software Target

Some Tandberg Viking Series NAS appliances use the Microsoft® Windows® Unified Data Storage Server 2003 operating system. This operating system provides unified storage server management capabilities, simplified setup and management of storage and shared folders, and support for Microsoft iSCSI Software Target. It is specially tuned to provide optimal performance for network-attached storage and provides significant enhancements in share and storage management scenarios, as well as integration of storage server management components and functionality. This chapter describes features of the Microsoft® Windows® Unified Data Storage Server 2003 operating system.

Notes

Not all Viking Series NAS appliances use the Microsoft®Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system.

IMPORTANT

The Microsoft® Windows® Unified Data Storage Server 2003, Enterprise x64 Edition operating system is designed to support 32–bit applications without modification; however, any 32–bit applications that are run on this operating system should be thoroughly tested before releasing the storage server to a production environment.

9.1. Microsoft iSCSI Software Target

The Microsoft iSCSI Software Target snap-in is a standard feature of Windows Unified Data Storage Server 2003. This snap-in makes it possible not only for the storage server to connect to remote iSCSI targets, but also to serve as an iSCSI target. With Microsoft iSCSI Software Target, you can create and manage iSCSI targets, create and manage disks for storage, and implement backup and recovery support using snapshots.

9.1.1. Virtual Disk Storage

The disks you create using iSCSI Software Target are iSCSI virtual disks, which are files in the virtual hard disk (VHD) format. These virtual disks offer flexible and effective storage. They are dynamically extendable to provide extra capacity on demand, enable efficient storage utilization, and minimize the time required to create new disks and the down time typically required to install new disks.

9.1.2. Snapshots

To facilitate backup and recovery operations, you can schedule and create snapshots of iSCSI virtual disks. A snapshot is a point-in-time, read-only copy of an iSCSI virtual disk. Snapshots are typically used as interim copies of data that has been modified since the most recent backup. Snapshots offer the following advantages:

- Snapshots can be scheduled to be created automatically.
- Snapshots are space-efficient because they are differential copies.
- It is not necessary to close files or stop programs when creating snapshots, so application servers
- can continue servicing clients without disruption.
- Each snapshot is typically created in less than one minute—regardless of the amount of data.
- Snapshots are useful for fast system recovery of files and volumes, in case of accidental data
- deletion by a user, overwritten data, or data corruption resulting from a malicious program.
- Snapshots can be mounted locally or exported to facilitate backup and recovery operations.

9.1.3. Wizards

To support creation and management of iSCSI targets, virtual disks, and snapshots, the iSCSI Software Target snap-in provides several wizards.

9.1.4. Create iSCSI Target Wizard

This section describes how to create an iSCSI Target using the Create iSCSI Target Wizard.

1. Log on to the storage server using an account with administrative privileges.
2. Open the Microsoft iSCSI Software Target MMC snap-in by clicking **Start > Programs > Administrative Tools > Microsoft iSCSI Software Target**.
3. Click the **iSCSI Targets** node. On the details view (right pane), right-click and select **Create iSCSI Target**.
4. Click **Next** on the **Welcome** page of the wizard.
5. On the **iSCSI Target Identification** page, type a name and description for the iSCSI Target and then click **Next**.
6. On the **iSCSI Initiators Identifiers** page, type the iSCSI Qualified Name (IQN) of the iSCSI initiator requesting access to the iSCSI Target in the **IQN identifier** field. The IQN is found on the **General** tab of the Microsoft iSCSI Initiator interface.

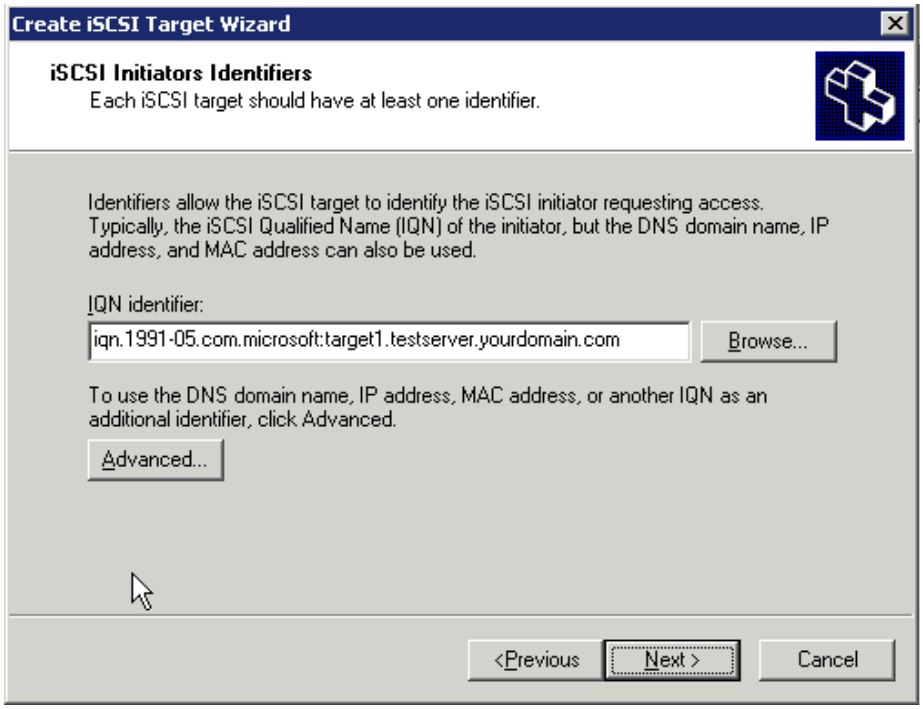

7. To enter additional identifiers, or if you are using an identifier other than an IQN (DNS domain name, IP address, or MAC address):

    - Click **Advanced**.
    - On the **Advanced Identifiers** page, click **Add**.

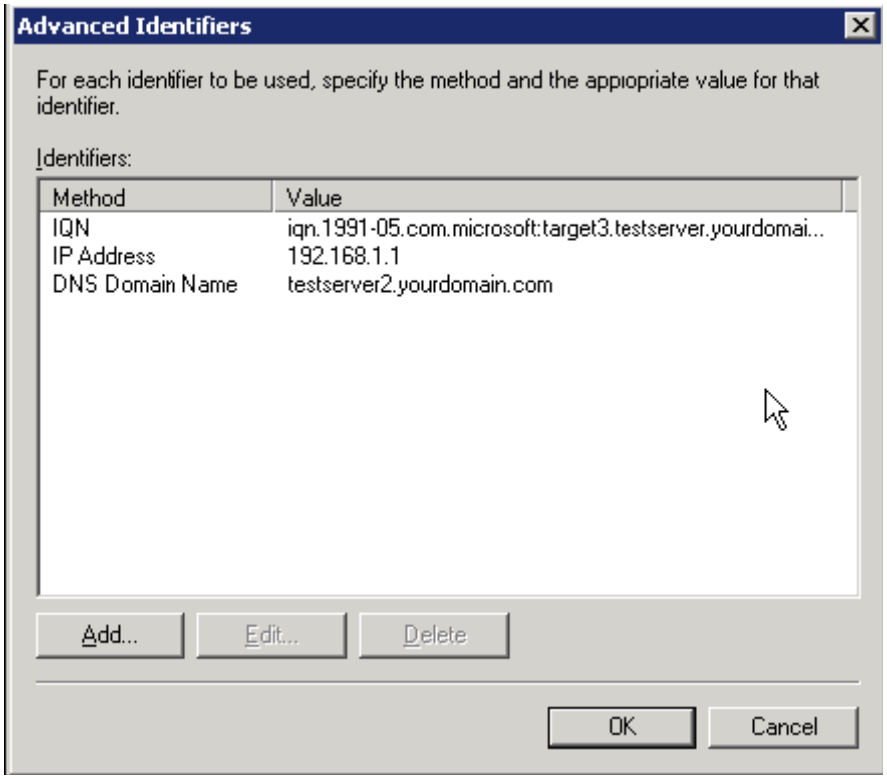

- Select the identifier type from the **Identifier Type** list and type the identifier in the **Value** field.

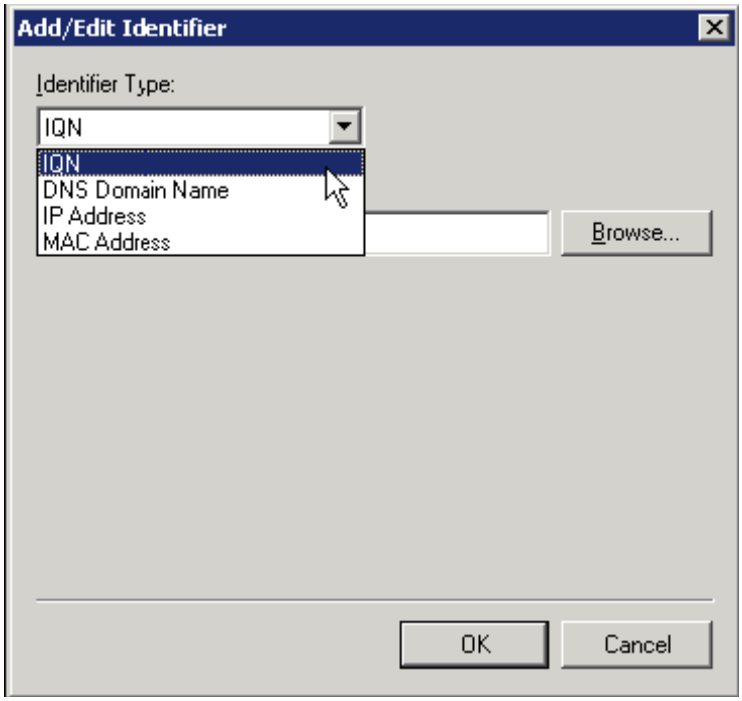

- Repeat steps b and c for each identifier you want to add.
- Click **OK**.
- Click **OK** again to close the **Advanced Identifiers** page.

8. Click **Next**.
9. Click **Finish** to complete the wizard and create the iSCSI Target.

9.1.5. Create Virtual Disk Wizard

This section describes how to create an iSCSI Virtual Disk using the Create Virtual Disk Wizard.

Notes
In order to create iSCSI Virtual Disks, it is required that physical disks are formatted as NTFS.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a volume and select **Create Virtual Disk**.

3. Click **Next** on the **Welcome** page of the wizard.
4. On the **File** page, specify the full path to use as the virtual disk and click **Next**.
5. On the **Size** page, specify the size to use for the virtual disk and click **Next**. If the file already exists, you cannot specify a new size.
6. Enter a description for the iSCSI virtual disk (optional) and click **Next**.
7. On the **Access** page, click **Add** to assign the iSCSI virtual disk to an iSCSI Target.
8. On the **Add Targets** dialog box, select a Target and click **OK**.
9. Click **Finish** to complete the wizard and create the iSCSI virtual disk.

Notes

If you delete a virtual disk, it is removed from the iSCSI Software Target MMC snap-in, but the virtual disk file (`.vhd`) is not removed from the physical disk. In order to permanently remove the virtual disk file, locate the file on the physical disk using Windows Explorer and manually delete it.

9.1.6. Import Virtual Disk Wizard

This section describes how to import a virtual disk using the Import Virtual Disk Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a volume and select **Import Virtual Disk**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Files** page, click **Browse**, navigate to the virtual disk file (`.vhd`) you want to import, select it, and then click **OK**.
5. Repeat step 4 for each virtual disk you want to import.
6. Click **Next** and then click **Finish** to complete the wizard and import the virtual disk(s).

9.1.7. Extend Virtual Disk Wizard

This section describes how to extend a virtual disk using the Extend Virtual Disk Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, click the **Devices** node.
2. On the details view (right pane) of the **Devices** node, right-click a virtual disk and select **Extend Virtual Disk**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Size** page, type the amount of space you want to add to the virtual disk in the **Additional virtual space capacity** field and then click **Next**.
5. Click **Finish** to complete the wizard and extend the virtual disk.

9.1.8. Schedule Snapshot Wizard

This section describes how to schedule a snapshot using the Schedule Snapshot Wizard.

1. In the **Microsoft iSCSI Software Target** MMC snap-in, expand the **Snapshots** node.
2. Right-click **Schedule** and select **Create Schedule**.
3. Click **Next** on the **Welcome** page of the wizard.
4. On the **Schedule Actions** page, specify whether the snapshots should be mounted locally or not.
5. On the **Name** page, type a name for the snapshot and then click **Next**.
6. On the **Virtual Disks** page, specify the virtual disks to include in the snapshot schedule.
7. On the **Frequency** page, select how often snapshots should be taken.
8. On the **Schedule** page, specify snapshot details according to the frequency selected on the previous page and then click **Next**.
9. Click **Finish** to complete the wizard and schedule snapshots.

9.2. Hardware Providers

To support advanced management of iSCSI virtual disks and snapshots, you can use the following hardware providers, which come preinstalled on the Viking Series NAS applaince:

- Microsoft iSCSI Software Target Virtual Disk Service Hardware Provider

Microsoft Windows Server 2003 introduced Virtual Disk Service (VDS), a set of application programming interfaces (APIs) that provides a single interface for managing disks. VDS provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks. The Microsoft iSCSI Software Target VDS Hardware Provider is required to manage virtual disks on a storage subsystem.

You install the Microsoft iSCSI Software Target VDS Hardware Provider on each iSCSI initiator computer running a storage management application (such as Storage Manager for SANs) that uses the hardware provider to manage storage.

- Microsoft iSCSI Software Target Volume Shadow Copy Service Hardware Provider

iSCSI snapshots are created using Volume Shadow Copy Service and a storage array with a hardware provider designed for use with Volume Shadow Copy Service. A Microsoft iSCSI Software Target VSS Hardware Provider is required to create transportable snapshots of iSCSI virtual disks and create application consistent snapshots from iSCSI initiators.

You install this hardware provider on the iSCSI initiator server and the server that is to perform backups. The backup software you use must support transporting snapshots.

This Page Intentionally Left Blank

10. Remote Access Methods & Monitoring

The Tandberg Data Viking Series NAS appliance comes from the factory with full remote manageability. Several methods of remote access are provided. These options let administrators use interfaces with which they are already familiar.

10.1. Remote Desktop

The NAS appliance supports Remote Desktop, with a license for two concurrently running open sessions. Remote Desktop provides the same capabilities as being physically present at the server console.

Use Remote Desktop to access:

- The NAS appliance desktop
- A command line interface
- Backup software
- Antivirus programs
- Telnet Server

Use the following logon credential to access the system.

1. Default Login: Administrator
   Default Password: 1234

2. RAID Storage Manager can be accessed using the following URL:

   http://<your NAS machine name or IP address>:81/

   Default Login: admin
   Default Password: 0000

10.2. Telnet Server

Telnet Server is a utility that lets users connect to machines, log on, and obtain a command prompt remotely. Telnet Server is preinstalled on the NAS server, but must be activated before use.

Caution

For security reasons, the Telnet Server service must be restarted each time the server is restarted.

Telnet Server can be enabled by using a Remote Desktop session or direct attached method to access a command line interface and enter the following command:

```
net start tlntsvr
```

The Telnet Server service needs to be enabled prior to running this command. The service can be enabled by opening the services MMC:

1. Select **Start**, **Run**, then type `services.msc`.
2. Locate the Telnet service, right-click on it, then select **Properties**.
3. In the startup type drop-down box, choose Manual, and click **OK**.

Note

The sessions screen provides the ability to view or terminate active sessions.

This Page Intentionally Left Blank