

Novell NetWare® 6

www.novell.com

NOVELL STORAGE SERVICES
ADMINISTRATION GUIDE



Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1993-2002 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

U.S. Patent No. 5,157,663; 5,349,642; 5,455,932; 5,553,139; 5,553,143; 5,572,528; 5,594,863; 5,608,903; 5,633,931; 5,652,859; 5,671,414; 5,677,851; 5,692,129; 5,701,459; 5,717,912; 5,758,069; 5,758,344; 5,781,724; 5,781,724; 5,781,733; 5,784,560; 5,787,439; 5,818,936; 5,828,882; 5,832,274; 5,832,275; 5,832,483; 5,832,487; 5,850,565; 5,859,978; 5,870,561; 5,870,739; 5,873,079; 5,878,415; 5,878,434; 5,884,304; 5,893,116; 5,893,118; 5,903,650; 5,903,720; 5,905,860; 5,910,803; 5,913,025; 5,913,209; 5,915,253; 5,925,108; 5,933,503; 5,933,826; 5,946,002; 5,946,467; 5,950,198; 5,956,718; 5,956,745; 5,964,872; 5,974,474; 5,983,223; 5,983,234; 5,987,471; 5,991,771; 5,991,810; 6,002,398; 6,014,667; 6,015,132; 6,016,499; 6,029,247; 6,047,289; 6,052,724; 6,061,743; 6,065,017; 6,094,672; 6,098,090; 6,105,062; 6,105,132; 6,115,039; 6,119,122; 6,144,959; 6,151,688; 6,157,925; 6,167,393; 6,173,289; 6,192,365; 6,216,123; 6,219,652; 6,229,809. Patent Pending.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.

www.novell.com

Novell Storage Services Administration Guide
February 2002
103-000141-001

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

ConsoleOne is a trademark of Novell, Inc.

Hot Fix is a trademark of Novell, Inc.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.

Storage Management Services and SMS are trademarks of Novell, Inc.

Transaction Tracking System and TTS are trademarks of Novell, Inc.

Third-Party Trademarks

All third-party trademarks are the property of their respective owners.

Contents

Contents	5
Introducing Novell Storage Services	9
Documentation Conventions	9
1 Overview	11
NSS Features and Benefits	12
NSS Management	12
Combine Logical Volumes in a Single Storage Pool	12
Create Logical Volumes	12
Exceed the Limit of Volumes in a Storage Pool (Overbooking)	13
Deactivate/Activate Logical Volumes and Storage Pools	13
Fast Error Correction and Data Recovery	13
Immediately Save Data to Disk	13
Retain Previously Saved Files (Snapshot)	14
Transaction Tracking System (TTS)	14
Review the Modified File List	14
Use Clustering Feature	14
Enable File Compression	14
Prevent Access to Deleted Files (Data Shredding)	15
Set Up User Space Restrictions	15
Set Up Directory Space Restrictions	15
Use CD-ROMs as Read-only Logical Volumes	15
No Additional Memory Required for Mounting Volumes	15
Redirect Data from Unreliable Blocks (Hot Fix)	16
Repair Storage Pools Instead of Individual Volumes	17
Mirror Partitions	17
Stripe Data Across Storage Devices (RAID)	18
NSS Enhancements	19
NSS Architecture	19
2 Setting Up and Configuring Novell Storage Services	21
System Requirements	21
Setting Up Novell Storage Services	21
Updating NSS Volumes	21
Using Multipath Support	23
Transaction Tracking System (TTS)	25

Setting the Cache Buffers	25
Using ConsoleOne to Configure Novell Storage Services	29
Setting Up ConsoleOne on a Workstation	30
Configuring Disk Space	30
Renaming a Storage Pool	38
Renaming a Logical Volume	39
Increasing the Size of a Storage Pool	39
Deleting NSS Pools or Logical Volumes	39
Reviewing and Restoring Deleted Logical Volumes	40
Changing the Purge Delay Setting	41
Creating a Traditional Volume	41
Increasing the Size of a Traditional Volume	42
Renaming a Traditional Volume	43
Deleting a Traditional Volume	43
Dismounting a Traditional Volume	43
Repairing a Traditional Volume	44
Using DFS to Create Junctions	46
DFS Limitations	47
Creating Junctions	48
Managing the VLDB Management Service	50
Setting Up File Snapshot	51
Converting Traditional Volumes to Logical Volumes	52
Volume Conversion	53
Modifying NDS Objects	54
Setting Up User Space Restrictions	54
Setting Up Directory Space Restrictions	55
Mounting a Logical Volume	55
Deactivating/Activating NSS Pools and Volumes	56
Salvaging and Purging Deleted Files	56
Using NetWare Backup Services	57
Rebuilding NSS Storage Pools and Volumes	57
NSS Server Console Commands	58
Online Help Commands	58
NSS Configuration Information Commands	58
Start Up NSS Commands	58
Other NSS Commands	59
DOS FAT Commands	63
Other Volumes That NSS Creates	63
CD-ROMs As Logical Volumes	64
DVDs as Logical Volumes	65
DOS Partitions As Logical Volumes	65
3 Viewing Storage Objects	67
Viewing Storage Devices	67

Viewing Your NSS Configuration.	67
View Partitioned Free Space	67

4 Troubleshooting **71**

Introducing Novell Storage Services

As your network structure grows, you need larger storage capacity, more efficient file management, and faster volume mounting speed. For example, you might need large volumes for numerous directories and files, or have numerous files open simultaneously. Novell® Storage Services™ (NSS) is a 64-bit file storage and management system that meets your file system needs. It is a seamless, scalable, and flexible storage and file system that lets you continually configure, mount, and maintain large volumes and numerous files without disrupting the work of end users.

This book is divided into the following sections:

For a description of NSS, see [Chapter 1, “Overview,”](#) on page 11.

For setup and configuration information, see [Chapter 2, “Setting Up and Configuring Novell Storage Services,”](#) on page 21.

For information about monitoring the status of your storage objects and NSS configuration, see [Chapter 3, “Viewing Storage Objects,”](#) on page 67.

For troubleshooting information, see [Chapter 4, “Troubleshooting,”](#) on page 71.

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

Also, a trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

1

Overview

Novell® Storage Services™ is a storage and file system that provides an efficient way to use all of the space on your storage devices. NSS is best used with systems that require the ability to store and maintain large volumes and numerous files or large databases.

Novell Storage Services is the default storage and file system for NetWare® 6. You use it to create, store, and maintain both traditional and NSS volumes. The NSS volumes are called logical volumes. When you install NSS, it creates a storage pool SYS: and an equal sized volume SYS:.

To mount a volume, the traditional file system scans every file and directory. Then to provide quick access to the files, NetWare stores all the meta data in memory. File system scanning slows the mounting speed. Storing the meta data for all files requires more memory as the number of files increase.

NSS speeds up the mounting process by not scanning the entire file system when you mount volumes. The memory requirements are lower because NSS does not load Meta data into the memory until you access the files.

You do not run VREPAIR on NSS volumes when the file system crashes because NSS keeps a journal of all file system transactions. After the crash, NSS scans the journal to ensure that all transactions are either completed or undone. This way volumes do not require any repair when you mount them again after the crash.

NSS uses free space from multiple storage devices. This lets you create unlimited volumes and store up to 8 trillion files in a single volume—up to 8 terabytes in size. NSS also lets you mount up to 255 volumes simultaneously.

IMPORTANT: In NetWare 5, NSS runs parallel with the traditional NetWare file system. In NetWare 6, NSS is the primary storage and management system. However, you can continue to maintain traditional partitions and volumes along

with NSS. Although this documentation focuses primarily on NSS storage pools and logical volumes, it includes instructions for setting up and maintaining traditional volumes.

NSS Features and Benefits

This section describes the features and benefits of NSS.

NSS Management

ConsoleOne™ is the primary utility used to configure and maintain both traditional and NSS logical volumes. See [“Using ConsoleOne to Configure Novell Storage Services” on page 29](#).

NOTE: NWCONFIG and NSS Menu are not compatible with some of the NetWare 6 NSS features. Use ConsoleOne for all NSS configuration and management procedures.

Combine Logical Volumes in a Single Storage Pool

NSS uses storage pools. A storage pool is a specified amount of space you obtain from all your storage devices. Then you place all NSS logical volumes into the storage pool. This way you do not have to limit the number of volumes you have in a partition. You can have only one storage pool on a partition, but you can place unlimited logical volumes in the storage pool. The size of logical volumes cannot exceed the size of the storage pool.

Create Logical Volumes

The volumes you create from NSS storage pools are called logical volumes. A logical volume is a volume that is either set to a specific size, or one that can grow dynamically according to the amount of physical space assigned to the pool. This lets you add and store any size or any number of files you need without having to create other partitions. You can add any number of volumes to a storage pool as long as you have available physical space in the pool.

Exceed the Limit of Volumes in a Storage Pool (Overbooking)

Individual logical volumes cannot exceed the size of a storage pool. However, you can create numerous volumes in a storage pool. Individual logical volumes must be smaller than the storage pool, but the sum of multiple logical volumes can exceed the pool size. This feature, called overbooking, can be an efficient way to manage your file system.

For example, you might have a number of users who are individually assigned to volumes with a limited amount of space for each volume. You can assign volumes that collectively exceed the pool size if not all users fill up their volumes. Essentially, the users can borrow space from other volumes in the pool as long as there are volumes that are not filled to the limit. This way it is not always necessary to add more disk space when some users reach or exceed their volume limit.

Deactivate/Activate Logical Volumes and Storage Pools

You might need to temporarily prevent user access to storage pools or volumes to do maintenance. Instead of bringing down the server, you can deactivate individual storage pools to run the REBUILD and VERIFY utilities. When you deactivate a storage pool, users will not have access to any of the volumes in that pool. The REBUILD and VERIFY utilities generate log files to the root of the DOS drive.

Fast Error Correction and Data Recovery

NSS quickly recovers data after a file system crash. Instead of scanning an entire volume for corruption, NSS quickly replays the latest set of changes to make sure they were written correctly. The file system either recovers the changed information, or it returns the data to its original settings before the transaction began. This is beneficial for applications such as mail services because there is less server down time.

Immediately Save Data to Disk

The Flush Files Immediately feature saves your file data to disk immediately after you close the file instead of waiting for the next flush cycle. This ensures that your data is not at risk of being lost between flush cycles.

Retain Previously Saved Files (Snapshot)

The File Snapshot feature allows your backup utility to keep a consistent copy of all files. This is done by having the file system keep an original copy of all open files. For example, when your system backs up or crashes at a time when you are modifying a file, the backup could contain both old and new data. Snapshot captures the most recent closed copy of the file. This way, if you lose data between the backup cycles, you still have a solid copy of the previously saved file.

Transaction Tracking System (TTS)

Transaction Tracking System™ protects database applications by backing out transactions that are incomplete due to a system failure. TTS™ can be available on either all traditional or all NSS logical volumes. It cannot be available on all volumes if you use both types of volumes. If your system uses both traditional and logical volumes, by default it can only work on traditional volumes. If you want to use it on logical volumes on a system that has both, you must enter `DISABLE TTS` on the `AUTOEXEC.NCF` then enable the feature on each logical volume you want it on. Unlike traditional volumes, TTS does not apply to logical volumes by default. To enable it on logical volumes, enter `nss/transaction=volname`.

Review the Modified File List

The Modified File List is a list of files that changed since the previous backup. The backup utility accesses this list instead of searching the entire file system.

Use Clustering Feature

Some storage devices exist in a storage area network (SAN). These devices can be shared by multiple servers in a cluster. For more information about clustering, see [NetWare 6](#).

Enable File Compression

NSS supports file compression. This lets you decide whether to compress the files in your volumes to create additional space. Once you enable file compression, you cannot turn it off without recreating the volumes.

Prevent Access to Deleted Files (Data Shredding)

The Data Shredding feature overwrites purged disk blocks with random patterns of hexadecimal characters. This prevents unauthorized users from using a disk editor to access purged files. You can place up to seven data shred patterns over deleted data.

Set Up User Space Restrictions

User space restrictions let you limit the space users have on volumes. This is beneficial for systems that accommodate a large number of users, such as students. When you create volumes, you can select the user space restriction option in ConsoleOne.

Set Up Directory Space Restrictions

Directory space restrictions let you limit the space users can have in a directory or subdirectory.

Use CD-ROMs as Read-only Logical Volumes

NSS has full CD-ROM support for ISO9660 and Macintosh* HFS formats.

No Additional Memory Required for Mounting Volumes

When you mount volumes using the traditional NetWare file system, it scans every file and directory during the mounting process. Then it loads all the files into memory for quick access. If you increase the number or size of files, you need additional memory to maintain the performance level you want.

NSS does not require large amounts of memory to mount volumes because it does not scan the entire file system during the mounting process. After the mounting is complete, NSS does not load files into memory until you access them. Therefore, no additional memory is required when you add files and mount volumes.

Redirect Data from Unreliable Blocks (Hot Fix)

As your file system constantly reads and writes data to disk, some disk storage blocks lose their ability to reliably store data. NSS uses Hot Fix to prevent data from being written to unreliable blocks.

Hot Fix™ redirects the original block of data (still in memory) to the Hot Fix Redirection Area of the partition where the data can be stored correctly. To redirect a block of data, the operating system records the address of the defective block. Then the server no longer attempts to store data in that block.

If you do not enable Hot Fix when you create a partition, then no Hot Fix object is created and no data redirection can occur. You need to set up Hot Fix when you create the partition. To add Hot Fix later, you must delete the volumes from the partition, add Hot Fix, then restore the volumes from a backup.

The Hot Fix feature is combined with partition mirroring. When you select Hot Fix, you also enable mirroring for that partition.

Hot Fix redirection can occur during a write request, a read request, or a read-after-write verification. The following are conditions of data redirection:

- ◆ Write redirection occurs when the disk reports an error during a write request. The system marks the block as bad and redirects the data to a different block.
- ◆ Read redirection occurs when a disk error happens during a read request. If the disk is mirrored, the system retrieves the data from the mirrored disk and redirects the data on the primary disk. If the disk is not mirrored, the data is lost, and the block is marked as unusable.
- ◆ Read-after-write-verify redirection occurs after data is written to disk. The data on disk is then read and compared to that in memory. If the two don't match, the system marks the block as bad and redirects the data to a different block.

By default, 2 percent of a disk's space is set aside as the Hot Fix Redirection Area. You can increase or decrease this amount.

If Hot Fix is enabled, then it is always active unless the disk fails or the redirection area is full. You can view Hot Fix activity in MONITOR by selecting Storage Devices and then a Hot Fix partition. For information on using MONITOR, see **MONITOR** in the *Utilities Reference*.

You can specify no Hot Fix Redirection Area when you create disk partitions. Having no Hot Fix Redirection Area saves partition space. This might be useful on large RAID systems that provide a method of fault tolerance.

IMPORTANT: If Hot Fix is not enabled, then the partition will not have the fault tolerance provided by NetWare. Data on corrupted sectors of the disk won't be redirected to the Hot Fix Redirection Area. In addition, you cannot mirror partitions that do not use Hot Fix. It is easiest to enable Hot Fix when you create the partition. To add Hot Fix later, you must delete all the volumes on the partition.

Repair Storage Pools Instead of Individual Volumes

You use VREPAIR to fix traditional volumes. For instructions on using VREPAIR, see [VREPAIR Utilities Reference](#). For logical volumes, you use repair utilities called VERIFY and REBUILD on the storage pools that contain the logical volumes. For instructions on rebuilding logical volumes, see [“Rebuilding NSS Storage Pools and Volumes” on page 57](#).

VERIFY checks the file system integrity for an NSS pool by searching for inconsistent data blocks or other errors. This utility indicates if there are problems with the file system.

REBUILD verifies and uses the existing leaves of an object tree to rebuild all the other trees in the system. You need to deactivate pools (and all the volumes in the pools) before you run Rebuild so users cannot access the volumes you are rebuilding. When you deactivate a storage pool, all the volumes in the pool automatically deactivate.

These utilities are not the same as VREPAIR. You should use REBUILD only as a last resort to recover the file system after data corruption. When you use REBUILD and VERIFY, a log file is generated to the root of the DOS drive.

Mirror Partitions

You can protect your data by mirroring your partitions. Mirroring stores the same data on separate disks using the same controller channel. Using ConsoleOne, you can mirror both traditional and NSS partitions.

The following is a list of requirements for mirroring partitions:

- ◆ Mirrored partitions must have the same partition type as you mirror the partition to. In other words, you can only mirror NSS partitions to other NSS partitions—the same applies to traditional partitions.

- ◆ You can only mirror partitions. However, you might want to mirror an entire storage pool. Because you can have only one storage pool on a partition, the only way to mirror that storage pool is to mirror all of the partitions the pool resides on.
- ◆ To mirror partitions, you need to select an option that makes the partitions compatible for mirroring when you create them—you cannot change that mirroring option after you create a partition. You can mirror to an existing group or create a new mirror group for the partition. You cannot combine mirror groups (existing groups with multiple mirrored partitions).
- ◆ Mirrored partitions must have compatible data area size. This means the new partition must be at least the same size or slightly larger than the other partitions in the group. For example, the physical size (combined data and hot fix size) must be at least 100 KB bigger, but no more than 120 MB bigger than the data size of the existing partitions in the mirror group.

The file system adjusts the Hot Fix size to the legal ranges in order to make the data area identical to the other partitions in the mirror group.

- ◆ Mirrored partitions must have the same sharable for clustering status.
- ◆ The partitions you add to a group cannot be part of an existing group—they must be individual mirrored objects.

Stripe Data Across Storage Devices (RAID)

NSS lets you create an additional storage device by striping data across multiple disk drives on your system. This option, called RAID, is a software configuration that emulates an actual hardware RAID 0 system. This RAID configuration occurs at the software level instead of using a RAID machine.

You set up a RAID device by obtaining space from all your physical storage devices then putting segments on the combined space. A RAID segment is the amount of space you obtain from each storage device. You can have up to eight segments on a RAID device. The file system places data on the RAID disks sequentially. This is called disk striping. The RAID stripe size is the amount of data the file system places on a disk before moving to the next disk.

You can use the software RAID 0 on both logical and traditional volumes.

Important Information about NSS RAID

- ◆ Each element in the RAID configuration should come from a different device. NSS will let you obtain RAID elements from the same device, but this will severely impede the performance of your file system.
- ◆ Currently, the NSS RAID supports RAID 0.
- ◆ The stripe size ranges between 4 KB and 256 KB in increments of 2 KB.

NSS Enhancements

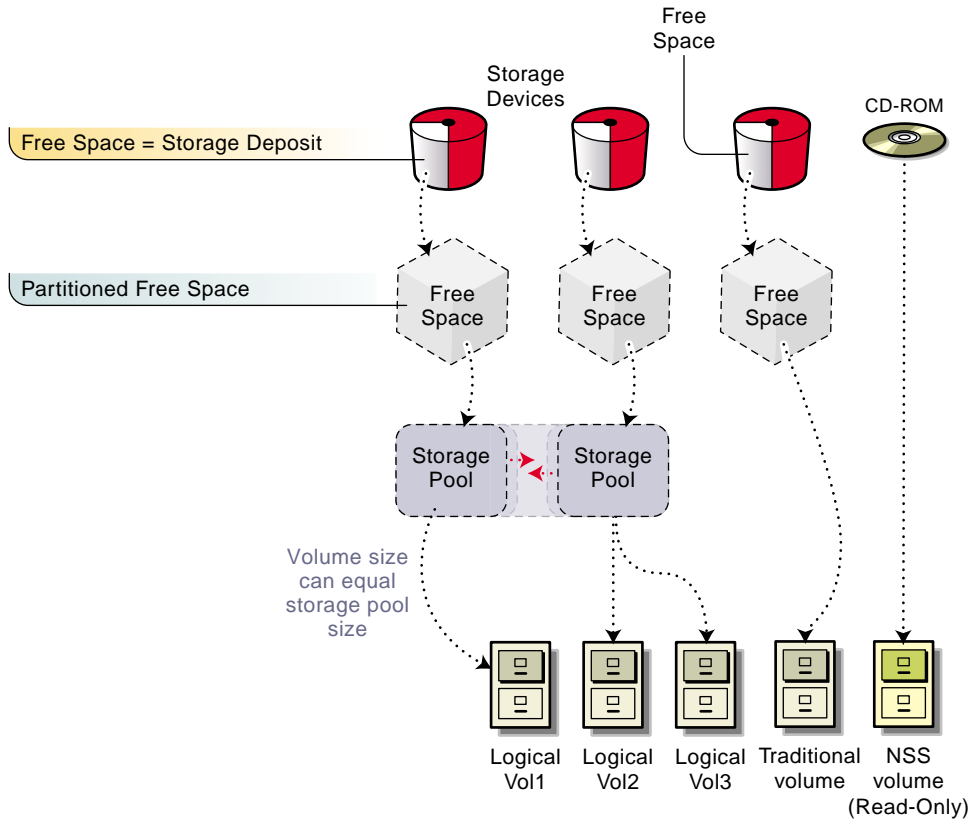
The following features that were available in the traditional NetWare file system behave differently in NetWare 6:

- ◆ Block suballocation
NSS uses 4K blocks, unlike the traditional file system which uses 64K blocks, thus, suballocation is no longer needed.
- ◆ Auditing
Novell relies on the newer Novell Advanced Auditing Services (NAAS) NetWare 6 does not work with the older audit system of past NetWare operating systems.
- ◆ File name locks
Because applications now open and lock files with deny-read and deny-write flags, support for file name locks has been removed in NSS. NSS supports the current file locking architecture.

NSS Architecture

With NSS, you can use storage pools, logical volumes, traditional volumes, and remote devices to maintain your file system. You store logical volumes in storage pools on NSS partitioned space. You store traditional volumes directly in space partitioned for traditional volumes. The following diagram illustrates the NSS configuration with the available file system options:

Figure 1 Novell Storage Services Configuration



2

Setting Up and Configuring Novell Storage Services

Novell® Storage Services™ installs automatically when you install NetWare® 6. This section provides instructions for setting up and configuring Novell Storage Services (NSS) after the installation. It also contains information about configuring and maintaining traditional partitions and volumes.

System Requirements

- ❑ A server running NetWare 6
- ❑ A least 10 MB free space to create an NSS storage pool and logical volume

Setting Up Novell Storage Services

Updating NSS Volumes

After you upgrade to NetWare® 6 and reboot the server, you need to update your NSS volumes. NSS volumes from previous releases will not function until you update them to NetWare 6 volumes—they do not upgrade automatically.

Checking the NDS Status

Before you update NSS volumes, you must ensure that NDS® is running properly; otherwise, the upgrade process might remove your trustees. To check ensure NDS is running, you need to check the time stamp for the backlink process.

Do the following to ensure NDS is running:

- 1 Use iMonitor to determine that the upgrade process is complete.
 - 1a Open iMonitor in a network browser by entering the IP address of the server followed by /nds. For example: `http://155.555.1/nds`.
 - 1b Click Agent Process Status > External Reference Status.
 - 1c Make sure the time stamp (under the Time column) shows a time after the time you began the upgrade—specifically when NDS unloaded during the upgrade. The following graphic provides an example of the time stamp:

Version	Flags	Time	Error	Activity	Remote Entry ID	Target Server ID	Target Server	Entry ID	Entry
0	Marker	08/22/01 08:55:06	Unknown	(0)					

If the time stamp shows a time earlier than when you began the upgrade, the backlink process is not complete. Repeat **Step 1a** through **Step 1c** until the backlink is complete.

Upgrading NSS Volumes

To upgrade NSS volumes, enter one of the following commands at the NetWare 6 server console:

- ◆ **NSS /ZLSSVOLUMEUPGRADE=ALL**
- ◆ **NSS /ZLSSVOLUMEUPGRADE=*volumename***

Using Multipath Support

When you use the server console to assign multiple paths, you can also assign a priority level to each path, according to your preference of which paths you want used first and last. The lower the number, the higher the priority. This number determines which path to take; if the current path fails, the path with the next highest priority is used. Multiple paths can have the same priority level. These values are persistent and they are stored in the registry.

The following commands are available from the server console.

- ◆ List Failover Devices

This command lists all of the failover devices. It also lists all the paths associated with each device, their assigned priorities, and the path that is selected.

If you want to see the paths associated with one particular device, enter the above command followed by the device ID. For example, if you wanted to see the paths associated with device Ox12, you would enter **List Failover Devices Ox12** at the server console and see the following results:

```
Ox12 [V132-A1-D2:20] ABC DEVICE rev:5678
```

```
Up Ox11 [V132-A1-D2:20] ABC DEVICE rev:5678 Priority = 1  
selected
```

```
Up Ox26 [V132-A1-D2:20] ABC DEVICE rev:5678 Priority = 2
```

This information tells you that there are two paths assigned to device Ox12. Out of those two devices, number 11 is *Up*, and it has also been assigned 1st priority. The word *Up*, simply means that this path is available for use. The word *selected*, means that any disk requests that are trying to reach the Ox12 device automatically goes through the number 11 path, unless another path is *selected* by using the MM Set Failover Path command.

- ◆ MM Set Failover Priority *pathid = number*

This command sets the priority level for each path. The *pathid* must be a valid path ID and the *number* is a decimal number between 1 and 4 million.

- ◆ MM Set Failover Priority *pathid = number /insert*

The */insert* option inserts a new path with the specified priority. For example, if you have four existing paths that are assigned the priorities of 1, 2, 3, and 4, and you enter a new path followed by a number 2 and the */insert* option, then the paths that were previously assigned to 2, 3, and 4 priority are bumped to priority 3, 4, and 5 in order to accommodate the new path that has been assigned priority 2.

- ◆ MM Set Failover State *pathid = Up /setpath*

This command sets the state of the path to Up. The *pathid* must be a valid path ID. If the path is up, it can be taken down, meaning, it will not be used and another path will be selected. If the Up option is followed by the */setpath* option, the highest priority path is automatically selected from the paths that are currently up.

- ◆ MM Set Failover State *pathid = Down /setpath*

This command option is useful if you want to do preventative maintenance or reroute cables on good, functional path. If the path is bad, NSS automatically moves the path to a down status.

If the Down option is followed by the */setpath* option, the highest priority path is automatically selected from the paths that are currently up.

- ◆ MM Set Failover Path *pathid*

This command will mark the path you entered as *selected*. The *pathid* must be a valid path that is up. The priority of a path remains static regardless of whether a path is considered up or down.

For more information about selected paths, see the List Failover Devices command of this bulleted list.

- ◆ MM Restore Failover *Path deviceid*

This command causes the device to reselect the highest priority path that is online. The *deviceid* must be a valid device ID.

- ◆ MM Set Multi-path Support = *Off*

This command will disable multipath support. Using this command, you can turn the multipath support on or off. In order for this command to

work, you must put this set parameter in the STARTUP.NCF file before any drivers are loaded.

Transaction Tracking System (TTS)

Transaction Tracking protects database applications by backing out transactions that are incomplete due to a system failure. TTS can be available on either all traditional or all NSS logical volumes. It cannot be available on all volumes if you use both types of volumes. If your system uses both traditional and logical volumes, by default it can only work on traditional volumes. If you want to use it on logical volumes on a system that has both, you must enter DISABLE TTS on the AUTOEXEC.NCF then enable the feature on each logical volume you want it on. Unlike traditional volumes, TTS does not apply to logical volumes by default. To enable it on logical volumes, enter the following command at the server console:

```
nss/transaction=<volname>.
```

Setting the Cache Buffers

A cache buffer is a block (4 Kb) of NetWare server memory (RAM) that temporarily stores data. When you install NSS, it consumes 60 percent of the the cache buffers by default.

The following are options for changing the cache buffers:

- ◆ You can set a maximum of 1048576 cache buffers.
- ◆ You can set the buffers in percentages rather than integers.

Use the following commands at the server console to set your cache buffers or tune your cache parameters. You can set or tune cache buffers at anytime.

Task	Purpose	Command
Set number of cache buffers	This sets the minimum number of cache buffers allocated for NSS file system access. The default and minimum number of cache buffers is 512. The range is 256 to 1048576.	<code>nss / MinBufferCacheSize= value</code>

Task	Purpose	Command
Look up cache buffers	To quickly look up cache buffers, the cache uses a hashing algorithm that maps names to buckets. Sometimes multiple names map to the same bucket. When this occurs, all names must be searched to determine if the entry is in the cache. The number of entries in the hash table is proportional to the amount of free memory available when you load NSS.	nss /NumHashBuckets
Set cache buffers for traditional file system	You can set the minimum number of cache buffers that NSS allocates for traditional file system access. To accomplish this, NSS returns pages to the operating system. This occurs until the operating system has enough pages, or until NSS reaches the minimum number of pages. The default is 1024. The range is 1024 to 1048576.	nss /MinOSBufferCacheSize
Set the cache for file and directory entries	This sets the number of Name cache entries. NSS keeps a cache of the file and directory names it recently looked up. This speeds up opening files and path searches. The default is 2111. The range is 3 to 65521.	nss /NameCacheSize
Set name cache to on or off	This sets Name Cache to ON or OFF. The default is ON.	nss / (No)NameCache
Use percentages instead of integers for cache balancing	This sets the buffers to percentages rather than integers for dynamic balancing of free memory for the buffer cache. The default is ON. Use this switch with /MinBufferCacheSize.	nss / (No)CacheBalance

Task	Purpose	Command
Set the amount of cache NSS takes from the traditional file system	This is the percent of cache between NSS and the server. The CacheBalance percentage determines how many cache blocks NSS will take from the traditional file system for its own cache. A high cache balance percentage will impede the performance of the traditional file system. A low cache balance will impede the performance of NSS. We recommend that you set the cache balance parameter to equal the percentage of the total disk space you allocate for NSS. However, do not exceed 80 percent. You might want to initially set the parameter at 50 percent and adjust this number as necessary. The default is 60 percent. The range is from 1 to 99 percent.	nss / CacheBalance=value
Limit the number of cache buffers	This is the maximum number of cache buffers you want. Every time you do a cache balance, limit the number of cache buffers to this amount.	nss / cacheBalanceMaxBuffer

Task	Purpose	Command
Show the cache hit ratio at the server console	<p>This is the number of cache hits divided by the number of attempts to get a file cache buffer from the cache. Any amount above 90 percent is good. This means the system is spending less time going to the disk to retrieve data. Lower numbers indicate NSS needs more memory. If this occurs, you might want to set the cache balance percentage higher.</p> <p>Cache hit is the number of times that NSS finds the desired file block in the cache buffers.</p> <p>Cache miss is the number of times that NSS does not find the desired block in the cache buffers.</p> <p>If you allocate more memory to NSS, the cache hit percentage increases. You can add more memory by adding more physical memory, increasing the cache balance, or increasing the minimum number of cache buffers for NSS. We do not recommend lowering the minimum number of OS cache buffers.</p>	nss /CacheHit
Set the number of Authorization cache entries	<p>This parameter specifies the number of cache buffers that will be used to speed up authorization requests. If many trustees have been set on different files and directories, we recommend that you increase this number. The default is 1024. The range is 16 to 50000.</p>	nss /AuthCacheSize

Task	Purpose	Command
Determine how often NSS rebalances the amount of cache it uses	This timer lets you determine how often (in seconds) NSS will rebalance the amount of cache it uses according to the total amount of memory on the server. If you have difficulty loading NLMs because of low memory, you might want to lower this timer setting. A lower setting forces NSS to rebalance more often. If your server is stable, you might want to increase this timer so NSS does not work as hard; this will ensure the cache remains balanced. NSS checks the total number of cache buffers to determine if the CacheBalance percentage is met. It then adjusts the timer to provide the appropriate number of cache buffers. The default is 30. The range is 1 to 3600.	<code>nss / CacheBalanceTimer</code>
Set the flush time for modified cache buffers	This sets the flush time for modified cache buffers in seconds. The default is 1 second. The range is 1 to 3600 seconds.	<code>nss /BufferFlushTimer</code>
Show cache buffer statistics	This shows the cache buffer statistics.	<code>nss /CacheStats</code>
Reset cache statistics	This command lets you clear the current statistics and start over. This is useful for checking the results of changes to the system.	<code>nss /reset</code>

Using ConsoleOne to Configure Novell Storage Services

ConsoleOne™ is the primary utility you use to configure and maintain partitions, storage pools, and volumes in NSS. To open ConsoleOne and access the options you need, complete the following:

- 1** Open ConsoleOne.
- 2** Open the tree you want to work with.
- 3** Right-click on the server object you want.
- 4** Click Properties.

This opens the ConsoleOne interface where you can access the snap-ins for managing NSS. The NSS configuration features are in the Media snap-in.

Setting Up ConsoleOne on a Workstation

To use ConsoleOne on a client workstation, you need to load ConsoleOne 1.3 or later. Then you must copy the NSS JAR files from your server to your client machine.

- 1 Copy these three NSS JAR files to your local LIB directory:

```
public\mgmt\consoleone\1.3\lib\nssadmin.jar
```

```
public\mgmt\consoleone\1.3\lib\nssc1lib.jar
```

```
public\mgmt\consoleone\1.3\lib\nssjavali.jar
```

- 2 Copy these two JAR files to your local RESOURCES directory:

```
public\mgmt\consoleone\1.3\resources\nssadminres.jar
```

```
public\mgmt\consoleone\1.3\resources\nssc1libres.jar
```

- 3 Copy the following JAR file to your local SNAPINS directory:

```
public\mgmt\consoleone\1.3\snapins\nssadminreg.jar
```

Configuring Disk Space

To set up the NSS storage and file system, complete the following tasks in the given order:

1. Create a RAID device (optional).
See [“Creating a RAID Device” on page 31](#).
2. Create partitions.
See [“Creating a Partition” on page 33](#).
3. Mirror partitions.
See [“Creating a Partition” on page 33](#).
4. Create storage pools.
See [“Creating a Storage Pool” on page 35](#).
5. Create logical and traditional volumes.
See [“Creating a Logical Volume” on page 36](#) and [“Creating a Traditional Volume” on page 41](#).

Creating a RAID Device

NSS lets you create a RAID 0 device by striping data across multiple drives on your system. This RAID configuration occurs at the software level. You can use this RAID feature for both logical and traditional volumes.

Important information about NSS RAID

- ◆ Each element in the RAID configuration should come from a different device. NSS will let you obtain RAID elements from the same device, but this will severely impede the performance of your file system.
- ◆ Currently, the NSS RAID supports RAID 0. This means there is no disk redundancy. If one element in the RAID device malfunctions, the entire configuration is unusable.
- ◆ NSS RAID can accommodate a maximum of eight segments (amount of space used from each device) on a device per RAID 0 logical device.
- ◆ The stripe size (amount of data the file system places on one device before moving to the next device) ranges between 4 KB and 256 KB in increments of 2 KB.

1 From ConsoleOne, open the tree you want.

2 Right-click on the server object and select Properties.

3 Click Media > Raid Devices > New.

4 Enter the amount of space you want to obtain from each storage device in megabytes.

5 Select a segment size.

6 Select a stripe size.

7 Select a RAID type.

The RAID type indicates the method of how data is striped across the different disks.

8 Click Next.

9 Select the devices you want to use for the RAID configuration.

You can select one or more devices, but select only the ones you want to use for the RAID configuration.

10 Click Finish.

The file system first creates a virtual device unit, then creates a RAID device.

Deleting a RAID Device

When you delete a RAID device, the file system deletes all the partitions and volumes from all physical storage devices used in the RAID configuration.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Raid Devices.
- 4** Select the Raid device you want to delete.
- 5** Click Delete.

The file system displays a warning message to ensure you want to delete the RAID device.

- 6** Click Yes or No.

Increasing the Size of the RAID Device

You can increase the size of a RAID device by adding more segments from another storage device. You should not place more than one RAID segment on a disk; this will severely impede the performance of your file system.

- 1** From ConsoleOne, click Media > Raid Devices.
- 2** Select the RAID device you want to increase.
- 3** Click Increase Size.
- 4** Select the device you want to add to the RAID configuration.
- 5** Click Finish.

After you add another physical storage device to your RAID device, you must restripe the configuration. This stripes the data to the other device you added.

NOTE: The file system performance might slow down during the restriping process. Make sure you add all the physical storage devices you want to the RAID device before you click Restripe.

- 6** Click Restripe.

The Restripe option recognizes new devices in your RAID configuration and distributes the data across all the devices.

Marking a Device as Sharable for Clustering

Some storage devices exist in a storage area network (SAN). These devices can be shared by multiple servers in a cluster.

IMPORTANT: The system hardware does not specify disk drives to come up automatically as sharable for clustering or not sharable for clustering. You must manually set different devices according to the configuration of your storage system.

The following is important information about making drives Sharable for Clustering:

- ◆ All portions of an NSS pool must be on a storage device using the Sharable for Clustering feature. This means if you have a storage pool that spans multiple devices, all these devices must be either sharable for clustering or not.
- ◆ The Sharable for Clustering option is not available for devices that contain the storage pool SYS: or Volume SYS:.
- ◆ ConsoleOne does not provide the Sharable for Clustering option for devices that contain traditional partitions and volumes.
- ◆ When you add partitions to a mirror group, all partitions in that group must be either sharable for clustering or not.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Devices.
- 4** Select a device.
- 5** Select Sharable for Clustering.
- 6** Click OK.

Creating a Partition

The first task for setting up NSS is to create partitions on your storage devices for the NSS storage pools and logical volumes.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click the server object and select Properties.
- 3** Click Media > Partitions > New.
- 4** Select a device to create a partition on.

- 5** Select the type of partition you want to create (NSS, Traditional, Remote Storage Device).
- 6** Enter the size of the partition in bytes (B), kilobytes (KB), megabytes (MB), or gigabytes (GB).
- 7** To reserve space for the Hot Fix error correction feature, select Hot Fix and enter the amount of space or percentage you want to reserve.

Mirrored partitions must be compatible in data area size. This means the new partition must be at least the same size or slightly larger than the other partitions in the group. The physical size (combined data and Hot Fix size) of the partition must be at least 100 KB, but no more than 120 MB larger than the data size of the existing partitions in the mirror group.

- 8** To mirror the partition, select Mirror and select one of the following options:
 - ◆ **Create New Mirror.** This option means you are making the partition capable of being part of a mirror group. You do not actually create the group until you add another mirrored partition to the partition you are creating.
 - ◆ **Existing Mirror Group.** (If you select this option, also select the ID of the mirrored partition.) This shows a list of existing mirror groups that are compatible in data area size. This option lets you add this new partition to one of the mirror groups in the list.
- 9** Select a label for the partition (optional).
- 10** Click OK.

Deleting a Partition

When you delete a partition you remove all volumes and data on that partition. If the partition is mirrored, the other partitions in the mirror group will retain the data from the deleted partition. If you want to delete a mirrored partition, you must unmirror that partition before you delete it.

Unmirror a partition

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click the server object and select Properties.
- 3** Click Media > Partitions > Mirror.
- 4** Select the partition you plan to delete, click Remove.

This removes the partition from the mirror group. You can now delete the partition.

Delete a Partition

- 1 From ConsoleOne, open the tree you want.
- 2 Right-click the server object and select Properties.
- 3 Click Media > Partitions.
- 4 Select the partition you want to delete, click Delete.

Creating a Storage Pool

After creating and setting up your partitions, you can create storage pools within the partitions. When you create a storage pool, NSS lets you determine how much space to use from each storage device on your server.

When you create storage pools, you use NSS partitioned or unpartitioned disk space. However, if you choose unpartitioned disk space, a partition and storage pool is created for you at the same time. The storage pool will be the same size as the partition because you can have only one partition per storage pool.

- 1 From ConsoleOne, open the tree you want.
- 2 Right-click the Server object and then click Properties.
- 3 Click Media > NSS Pools > New.
- 4 Enter a name for the new storage pool and then click Next.
ZLSS is the default pool (LSS) type.
- 5 Select the existing NSS partitions or unpartitioned disk space that you want to use for the pool and then click Next.

If you select an unpartitioned device, you can double-click the number in this column to edit the current amount of space you want to obtain from the storage devices. The storage pool size will be the total of all devices in the pool.

- 6 Do one of the following:
 - ♦ If the partition or free space that you selected is *not* sharable for clustering, the Create Pool dialog box appears with the Activate on Creation check box automatically checked. Click Finish if you want to activate this pool on creation; otherwise, uncheck the box and click Finish.
 - ♦ If the partition or free space that you selected *is* sharable, the Create Pool dialog box appears with the Activate on Creation and the

Cluster Enable on Creation check boxes checked. You must fill in additional clustering parameters and then click Finish.

For more information on clustering parameters, see Setting Up Novell Cluster Services > Create NSS Pools in the Novell Cluster Services Overview and Installation Guide.

Creating a Logical Volume

After creating storage pools, you can create any number of logical volumes for them according to the physical space available. When you create a logical volume, you can either assign it a space or allow it to expand to the pool size.

The size of a single volume cannot exceed the size of a storage pool. However, the combined sizes of multiple volumes can exceed the size of the storage pool. This is called *overbooking*. For example, you might have an 800 MB storage pool with eight volumes set at 100 MB each. That is the maximum amount of space you have in the pool. To overbook the pool, you can add volumes or increase the size of existing volumes knowing that not all of the volumes in the storage pool will fill up. If the volumes fill up, you need to add more disk space.

NSS recognizes DOS, Macintosh*, UNIX*, and long name spaces.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click the Server object and then click Properties.
- 3** Click Media > NSS Logical Volumes > New.
- 4** Enter a name for the volume and click Next.
- 5** Complete the following steps:
 - 5a** Select the existing storage pool, unpartitioned space, or NSS partitioned space where you want to create your volume and then click Next.

If you select unpartitioned space, NSS creates an NSS partition, then the storage pool, and then the volume.
 - 5b** Enter a quota size for the volume if you want to limit the size of the volume *or* check “Allow Volume Quota to Grow to the Pool Size” if you want the volume to expand to the size of the pool.
 - 5c** Click Next.

- 6** (Conditional) If you chose in [Step 5a on page 36](#) to create your volume on NSS partitioned space or free space, enter in the name and size of your NSS pool and then click OK.

Because all volumes must be contained in a storage pool, if you selected free space that did not already contain a volume, NSS will create a pool at the same time that it creates the volume.

- 7** Do one of the following:

- ◆ If the device that you selected is *not* sharable for clustering, skip to [Step 8](#).
- ◆ If the device you selected *is* sharable, determine whether you want to cluster enable the volume on creation. Cluster enabling a volume means that you will assign the volume an IP address so it can be accessed through the Novell client.

If the Cluster Enable on Creation check box is checked, you must fill in additional clustering parameters. When you have filled in the additional clustering parameters, click Finish and continue with [Step 8](#).

For more information on clustering parameters, see [Setting Up Novell Cluster Services > Create NSS Pools in the Novell Cluster Services Administration Guide](#).

- 8** Select the attribute settings you want for the new logical volume from the following list:

- ◆ **Backup.** Indicates if the volume should be backed up. Set this flag if the volume contains data you want to back up. Clear this flag if the volume is empty or if backing up the data is unnecessary. This backup option is separate from the third-party backup vendor you use. This means your backup system might not recognize this option whether you select it or not.
- ◆ **Compression.** Activates file compression in the logical volumes.
- ◆ **Data Shredding.** Activates the Data Shredding security feature. This feature scrambles any data that you delete to prevent anyone from accessing the information with a disk reader. Enter the number of times you want the data shredder to scramble your deleted files. You can select between one and seven times.
- ◆ **Directory Quotas.** Activates a feature that lets you assign a maximum quota of space a directory can have

- ◆ **Flush Files Immediately.** Activates the feature that flushes all data in a file to disk immediately when you close the file.
- ◆ **Modified File List (MFL).** This provides a list of all files modified since the previous backup. A third-party vendor must implement the Modified File List feature in order for you to use it.
- ◆ **Salvage Files.** Activates the feature that lets you salvage deleted files. The file system keeps all deleted files in an allocated space until that space is needed for other data. The Salvage feature tracks the files and lets you retrieve the data for a time until the space is needed for other data. You must select this option to salvage the files in the volume.
- ◆ **Snapshot-File Level (File Snapshot).** Activates the Snapshot feature at the file level. The Snapshot feature allows the backup utility to capture a snapshot of the last closed version of a file. For example, if your system backs up or crashes while you have a file open, this feature will save a copy of the file before you opened it. You might lose some new information, but you will retain all the previous information.
- ◆ **User Space Restrictions.** Activates the user space restrictions feature on the volumes you create.
- ◆ **Activate.** Activates logical volumes as soon as you create them.
- ◆ **Mount.** Mounts logical volumes as soon as you create them.

9 Click Finish.

You can now store and manage files in the NSS system.

Renaming a Storage Pool

You can rename existing storage pools. For example, you might want to assign a storage pool name that relates to a department name change.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Pools.
- 4** Select a pool to rename.
- 5** Click Rename.
- 6** Enter a new name for the pool.
- 7** Click Finish.

Renaming a Logical Volume

You can rename logical volumes. For example, you might want to change the name of a volume to reflect the department or organization that uses it.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Logical Volumes.
- 4** Click Rename.
- 5** Enter a new name for the logical volume.
- 6** Click Finish.

Increasing the Size of a Storage Pool

You can increase the size of your storage pools, but you cannot reduce their size.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Pools.
- 4** Select the storage pool you want to increase, then click Increase Size.
- 5** Select the storage objects you want to obtain the space from.
- 6** Enter the amount of space you want from each storage object in the Used column.
- 7** Click Finish.

Deleting NSS Pools or Logical Volumes

You might need to delete a storage pool or a logical volume to create more free space for other storage pools. When you delete a logical volume or a traditional volume from a partition, that partition still owns the space, but you can reassign the space to another storage pool or traditional file system volume. When you delete a partition, ownership of the space in that partition is removed.

After you delete a logical or traditional volume, free space becomes available. You can assign this space to other storage pools to increase their size. You can also use the free space (if you removed NSS ownership) for a traditional file

system volume. If necessary, you can restore a volume, see [“Reviewing and Restoring Deleted Logical Volumes” on page 40](#). If you delete a storage pool, you cannot restore any of the volumes in that pool.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Pools or Logical Volumes.
- 4** Select the pool or volume you want and click Delete.

Reviewing and Restoring Deleted Logical Volumes

If you delete a volume, the file system removes it from the storage pool. However, for a specified amount of time, called the Purge Delay time, you can review and even restore the contents of the volume you deleted. For example, you might want to restore information in a volume that was deleted in error. You must retrieve the volume before the delay time elapses; otherwise, the volume is removed from the system and you can no longer access it.

If you delete a storage pool, you delete all the volumes in that pool. You cannot restore volumes.

The default setting for the Purge Delay time is two days. After this time expires, NSS automatically purges the volume. You can change the Purge Delay time to extend or reduce the time for the automatic purging cycle. See [“Changing the Purge Delay Setting” on page 41](#). You can also manually purge deleted volumes. You purge or restore volumes in the “Deleted Logical Volumes on Pool” screen.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Logical Volumes.
- 4** Click Deleted Volumes.
- 5** Select one of the following options:
 - ◆ **Purge.** Click Purge to immediately purge all the deleted volumes.
 - ◆ **Prevent Purge/Allow Purge.** Click Prevent Purge to stop the volume purging process. For example, you might want to access the deleted volume again, or wait until there is less activity on the server. After you click Pause Purge, the button changes to Allow Purge. Click this button to purge the volume.

- ◆ **Salvage.** Click Salvage to restore the deleted volume.
 - ◆ **Refresh.** Click Refresh to rescan the volumes that have been deleted and update the list in the panel.
- 6** Click Close when you are finished.

Changing the Purge Delay Setting

The default for the Purge Delay setting is two days. This is the amount of time that you can still access the volume before it is removed from the system. To change the Purge Delay time, enter the following at the server console:

```
NSS/logicalVolumePurgeDelay=ddd
```

In this command, you replace ddd with the actual number of seconds to delay the auto purge. The default is 172800 (2 days). For example, if you want to change the Purge Delay time for one day, enter 86400.

The Purge Delay change command is not permanent. You must enter the command each time you restart the server. If you want to make the new setting permanent, enter the command in the AUTOEXEC.NCF file.

Creating a Traditional Volume

NSS is the default file storage and management system in NetWare 6. However, you might want to continue maintaining traditional partitions and volumes in NSS. The following is a list of guidelines for creating traditional volumes in NetWare 6:

- ◆ NWCONFIG and NSS Menu are not compatible with NetWare 6. Use ConsoleOne to create, mount, modify, and rename traditional volumes.
- ◆ You can use VREPAIR to fix traditional volumes, but you need to use VERIFY and REBUILD for all logical volumes. For instructions on using VREPAIR, see **VREPAIR** in the *Utilities Reference*.
- ◆ You cannot put traditional volumes in an NSS storage pool.
- ◆ If you create a traditional volume in NetWare 6, you cannot access that volume from previous releases of NetWare.

You can use ConsoleOne to create, mount, and manage traditional volumes.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.

- 3** Click Media > Traditional Volumes.
- 4** Click New.
- 5** Enter a name for the volume and click Next.
- 6** Select either partitioned or unpartitioned.
- 7** Enter a size for the volume in the Used column and click Next.
- 8** Select a block size from the pull-down menu.
- 9** Select any of the following options:
 - ♦ **Compression.** Enables the file system to compress the files in the volume. You set up file compression when you create volumes. Once you enable file compression, you cannot turn it off without creating the volume again.
 - ♦ **Migration.** Enables the operating system to move infrequently accessed data to remote areas on your server. This creates space for new and more commonly accessed data. Selecting this option only enables the attribute. The data migration feature uses a third-party software package that does the migration function.
 - ♦ **Suballocation.** Enables the file system to divide partially used disk blocks into suballocation blocks of 512 bytes. These blocks can be used by other data files.
 - ♦ **Mount Volume on Creation.** Instructs the operating system to mount the volume when you create it. Otherwise, you can mount it later.
- 10** Click Finish.

Increasing the Size of a Traditional Volume

To increase the size of a traditional volume, you need to add another segment to that volume.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Traditional Volumes.
- 4** Select a volume and click Increase size.
- 5** Select the storage device you want to obtain space from.
- 6** Double-click on the Used column next to the selected storage device.

- 7** Enter the amount of space you want.
- 8** Click Finish.

Renaming a Traditional Volume

You can rename a traditional volume.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Traditional Volumes.
- 4** Select the volume you want to rename and click Rename.
- 5** Enter a name.
- 6** Click Finish.

Deleting a Traditional Volume

You can delete traditional volumes.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Traditional Volumes.
- 4** Select the volume you want to delete and click Delete.

Dismounting a Traditional Volume

To repair a traditional volume, you need to dismount that volume.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Traditional Volumes.
- 4** Select the volume you want to dismount and click Dismount.
The button label changes to mount.

Repairing a Traditional Volume

Typically, you can't mount a volume if it has even minor damage. Occasionally, a damaged volume mounts and causes errors in the process.

Use VREPAIR to correct volume problems or to remove name space entries from File Allocation Tables (FATs) and Directory Entry Tables (DETs). For instructions on using VREPAIR, see **VREPAIR** in the *Utilities Reference*.

NOTE: If you need to repair an NSS volume, see “**Rebuilding NSS Storage Pools and Volumes**” on page 57.

You can run VREPAIR on a damaged volume while other volumes are mounted. Following are typical instances when VREPAIR can help:

- ◆ A hardware failure either prevented a volume from mounting or caused a disk read error.
NOTE: Although VREPAIR can't fix hardware problems, VREPAIR can sometimes fix related volume damage.
- ◆ A power failure caused a corrupted volume.
- ◆ The server console displays a mirroring error when the server boots. This mirroring refers to the two copies of FATs and DETs that the operating system keeps (if disks are mirrored, NetWare keeps four copies).

If a volume fails to mount as the server is booting, VREPAIR loads automatically and attempts to repair the volume.

When VREPAIR autoloads, it uses the default options. If you want to use an alternate option, load VREPAIR manually and set the alternate option before running VREPAIR.

NOTE: If you don't want VREPAIR to automatically repair a volume that won't mount, use the SET parameter named Automatically Repair Bad Volumes to change the default.

Prerequisites

- The volume you want to repair must be dismounted.
- If the volume to be repaired has name space support, the corresponding VREPAIR name space module (`V_namespace.NLM`) must be located in either the SYS:SYSTEM directory or in a search path directory.

Example modules include `V_MAC.NLM` and `V_LONG.NLM`.

Procedure

- 1 At the server console prompt, enter

```
VREPAIR [volume_name] [logfile_name]
```

(Optional) Replace *volume_name* with the name of the volume to repair. If there is only one volume that is dismounted, you don't need to specify this parameter, since VREPAIR will attempt to repair that volume.

(Optional) If you want to save the error log, replace the *logfile_name* with the name of the file you want VREPAIR to create. VREPAIR creates a log of errors it finds. VREPAIR displays the errors on screen and will write them to a file if you specify a filename.

When you load VREPAIR, an Options menu is displayed.

- 2 Accept the default options, or select alternate options, as appropriate.

The first time you try to repair a volume, accept the default options. If the default options fail to repair the volume, select alternate options.

2a To accept the default options, continue with Step 3.

2b To set alternate options at the Options menu, choose Set VRepair Options by typing

```
2 Enter
```

- 3 To begin the repair process, choose Repair A Volume from the Options menu.

- ♦ If more than one volume is dismounted, select the volume to repair from those listed.
- ♦ If only one volume is dismounted, VREPAIR assumes it is the volume that needs repairing and begins the repair.

As the volume is being repaired, the server console screen displays a message indicating VREPAIR activity.

- 4 (Optional) Modify error log settings after the repair has started.

If VREPAIR finds many errors during the repair process, you might want to change some of the run-time error settings. To modify these settings after the repair has started, press F1 to display the Current Error Settings menu.

- ♦ Select Option 1 if you do not want VREPAIR to pause after each error.
- ♦ Select Option 2 if you want VREPAIR to log errors in a text file.

- ◆ Select Option 3 to stop the repair of the volume.
 - ◆ Select Option 4 to continue with a volume repair after you have stopped it.
- 5** When the repair is complete, answer Y when prompted to write repairs to the disk.
 - 6** If VREPAIR has found errors, run VREPAIR again by repeating **Step 2** through **Step 6**. Repeat until VREPAIR finds no errors.

If you are unable to mount the volume after running VREPAIR several times, you must delete the volume, re-create the volume using ConsoleOne.

Using DFS to Create Junctions

Distributed File Services (DFS) is a service, installed by default with NetWare 6, that provides access to remote volumes. This process, called *junctioning*, allows volumes to be joined together so that they appear as subdirectories in one file system regardless of their physical location.

A *junction* is simply a virtual directory or folder that is assigned a pseudo filename within a file system. The junction contains information that redirects file requests to a file system located at the root of another NetWare volume.

Junctions can point to the root of any NetWare 5.1 or 6 traditional or NSS volumes that reside on another server within the *same* NDS tree. However, a system administrator can *create* junctions *only* on NetWare 6 NSS volumes. For a complete list of DFS limitations, see "**DFS Limitations**" below.

When a junction is created, the Volume Location Database (VLDB) assigns the junction a unique ID, which is automatically stored in the VLDB and also in the corresponding NDS Volume object. When a user views the contents of a junction, the client looks at the unique ID, then queries the database to find the path to the physical volume associated with the unique ID, and then displays the contents of that physical volume.

In the end, the junction appears to the user as a new subdirectory and its contents include the entire volume that the junction points to. Users are unable to create or manipulate these junctions. The access rights to data within the new volume is still controlled by normal rights assigned to that user on that volume; the junction serves merely as a shortcut to access the volume.

Junctions eliminate the end user's need to know the path to the physical location of the data. For example, if John's data is located on server's X, Y, and Z, you can create junctions on server X that point to all of his data on servers Y and Z. That way, John only has to remember the path to server X, because with junctions, it appears as if the data is all located in one place.

The benefit of this feature is two-fold. Not only does it decrease administration costs by allowing the administrator to move a volume to a different server without making any announcements or having to reeducate users, but it also simplifies the number of paths a user has to remember if his data is located in multiple places on the network.

DFS Limitations

Before you begin creating junctions, make sure you are familiar with the following DFS limitations:

- ❑ At this time, you can create junctions only on NetWare 6 NSS volumes; however, you can create junctions on NetWare 6 NSS volumes that point to either NetWare 5.1 NSS or traditional volumes.
- ❑ You can create junctions only within the same NDS tree.
- ❑ DFS does not support the Novell or Macintosh client.

However, it does support Web-based services, like WebDAV or NetStorage, via the XTier framework or the Microsoft Client via the CIFS protocol.

In order to use DFS to create junctions on your NetWare 6 servers, you must either

- ◆ Be accessing the DFS files through WebDAV, NetStorage, or any other Web-based service that supports the XTier protocol
 - ◆ Have the Microsoft Client installed on the workstation that is accessing the files
- ❑ Use ConsoleOne to create or delete junctions. Network Administrator and Remote Manager do not yet work with DFS and junctions.

Creating Junctions

Before you can begin to create an actual junction, you must create at least one DFS Management Context at an O or OU level in the NDS tree. You can create more than one Management Context if you have a geographically diverse company. That way, each geographic area can manage and control the junctions and the VLDB Management Service within its own domain.

IMPORTANT: If you are creating more than one Management Context, be aware that a server cannot be a part of two different Management Contexts. If you are junctioning to volumes that are physically located on servers outside your Management Context, and volume repairs need to be made, you will have to contact the network administrator responsible for the Management Context that contains that specific server.

When you create a Management Context, you specify which servers will run the VLDB Management Service and hold the actual database. It is the activated VLDB Management Service, which includes the database, that allows you to create junctions.

Creating a DFS Management Context with ConsoleOne

- 1** Right-click an O or an OU in the NDS tree and then click New > DFS Management Context.
- 2** Select one or more servers from the Available list, click the Right-arrow to move them into the Selected list, and then click Next.

The servers that you selected will run the VLDB Management Service and house the VLDB database. Select more than one if you want to keep a copy of the database on another server to prevent data loss or corruption.

- 3** Specify where you want to store the database on the server, and then click Finish.

The default location is SYS:\ETC.

Once you create the Management Context, attributes are created in the container object and the service is started on the servers that you selected.

Once the Management Context is created and the VLDB Management Service is up and running, you can create junctions.

Creating Junctions with ConsoleOne

- 1** In ConsoleOne, create a subdirectory or use an existing subdirectory of any NetWare 6 NSS volume.
IMPORTANT: If you installed a NetWare 6 server within the DFS Management Context after the DFS Management Context has already been activated, you might get an error when you try to create a junction because those NetWare 6 volumes do not have a DFS attribute. To fix this, you must rebuild the VLDB. Follow the instructions in ["Repairing the Volume Location Database" on page 51](#).
- 2** Right-click the subdirectory and then click > New > Shortcut > Junction.
You can also create a junction at the root of a volume as well. However, to do this, instead of right-clicking the subdirectory, select the volume and then click the Up-arrow in the ConsoleOne toolbar.
- 3** Verify that the NDS tree name is correct, browse to the volume that you are pointing to from the junction, and then click Next.
- 4** Name the junction and then click Finish.
- 5** To view the newly created junction, from the Console One toolbar, click View > Refresh.

When DFS is activated, it creates the following files in SYS:ETC:

- ◆ VLDB.DAT
- ◆ VLDBCFG.DAT
- ◆ VLRPR.LOG

In addition, each NetWare volume will contain a 24 byte, hidden, read-only file called ~DFSINFO.8-P.

IMPORTANT: Do not modify or delete any of these four files. If you accidentally delete any of these files, you must rebuild the VLDB from the VLDB Management Service on the volume that you have deleted the file from to re-create the file. To rebuild the VLDB, follow the instructions in ["Repairing the Volume Location Database" on page 51](#).

Deleting a Junction

- 1** Right-click the junction in ConsoleOne.
- 2** Click Delete Shortcut > Junction.

Continue with the next section, "[Managing the VLDB Management Service.](#)"

Managing the VLDB Management Service

To access the VLDB management service, in ConsoleOne, right-click any server that has the VLDB management service running and then click Properties > Supported Services > NSS VLDB.

The management service shows you basic information, such as the version of the service, whether the service is running or not, and the database location.

The management service also allows you to perform the following administrative tasks:

- ◆ Adjust processing threads
- ◆ Start and stop the VLDB Management Service
- ◆ Add and remove replicas
- ◆ Repair volumes

HINT: You can also manage the VLDB Management Service from the server console. To view a listing of the available commands, enter `VLDB HELP` at the server console.

Adding and Removing Replicas

At this time, to add or remove a replica from a server, you must access the VLDB Management Service from the server that you want to modify.

- 1** In ConsoleOne, right-click the server you want to modify, and then click Properties > Supported Services > NSS VLDB.
- 2** In the VLDB Management Service, click either Add Replica or Remove Replica.

Starting and Stopping the VLDB Management Service

To start the VLDB Management Service, enter `vldb` at the server console.

To stop the VLDB Management Service, enter `vldb exit` at the server console.

Repairing the Volume Location Database

You can repair the VLDB in the following three ways:

- ◆ Low-level Repair

The repair option restores the last saved copy of your backup files to the database.

- ◆ Get Database from Another Replica

You can use this feature only if you have the VLDB Management Service running on more than one server. The VLDB Management Service gets a copy of the database from another server that is currently running the service.

- ◆ Rebuild

When you rebuild the VLDB, the VLDB Management Service walks the NDS tree, looks at volume and server objects, and then completely rebuilds the database from scratch.

Setting Up File Snapshot

The File Snapshot feature enables your backup utility to keep an original copy of all data files. For example, when your system backs up or crashes at a time when you are modifying a file, the file backup could contain both old and new data. Snapshot captures the most recent closed copy of the file. This way, if you lose data between the backup cycles, you still have a solid copy of the previously saved file.

If you select File Snapshot on a volume, the backup utility copies all the meta-data (owner, created date/time, modifier, etc.). During write requests, only the data that is being written to is copied. NSS recognizes how to return the appropriate data when read requests come in for the real file or the Snapshot file.

Use the following NSS command at the server console to enable the File Snapshot for a volume:

```
nss /FileCopyOnWrite=<volName>
```

or

```
nss /FileCopyOnWrite=all
```

To disable File Snapshot, enter the following:

```
nss /NoFileCopyOnWrite=<volName>
```

or

```
nss /NoFileCopyOnWrite=all
```

After you set up File Snapshot on your volumes, you need to deactivate the volumes. Then you need to activate and mount the volumes again to ensure there are no open files without a snapshot.

To open a snapshot file for backup, you need to set a bit in the requested rights parameter in addition to `READ_ACCESS_BIT`. This bit is `0x04000000`. It is defined in the NSS header files as `zPR_READ_ACCESS_TO_SNAPSHOT`. The default setting for File Snapshot is when your file system receives an open for write request, it will wait for the backup open to close the file. If you want to return an error, the backup open should also set the deny the write bit.

To get information on a snapshot file, use the `getInfo` command. Any `getInfo` call that returns packed information with the `RNewStyle` bit set in the info mask is supported. You need to set an additional bit `0x20000000` in the information mask. This bit is defined at `RNSSInfoOnSnapshot` in NSS header files. Both `open` and `getinfo` calls require the bit for requested rights and the bit for additional info mask.

The following are some additional calls that return information about File Snapshot:

- ◆ `GetFileSize`
- ◆ `GetFileHoles`

You can also use `ConsoleOne` to select the File Snapshot feature. See [“Creating a Logical Volume” on page 36](#).

Converting Traditional Volumes to Logical Volumes

You can convert traditional volumes to logical volumes. This is not an automatic process or an in-place upgrade. You need to ensure you have enough space to transition from a traditional volume to a logical volume and maintain your attributes. The conversion utility is a server command that copies the data from a traditional volume to an NSS logical volume in an NSS storage pool. (For example, if you want to move a 2 GB volume, you need to have at least 2 GB of free space available. The original volume will be renamed `volumename_old`. The new logical volume will keep the original volume name.

Volume Conversion

After you copy the traditional volume to a logical volume, you can restart the server to ensure the volume converted properly. If the volume converted, you can remove the traditional volume. After you copy traditional volume data to a logical volume in NetWare 6, you cannot access the new logical volume in NetWare 5.1. You would need to either retain the data in the traditional volume, or copy the data back to the traditional volume.

The volume conversion utility can convert volumes with long names.

The following is a sample command line with all the available parameters for converting traditional volumes:

```
vcu /p /l /i /d /r originalVolume nssPool [ds_container [ds_volName]]
```

- ◆ /p: Do not print directory file names.
- ◆ /l: Do not write errors to log file (Dst_Vol:error.out).
- ◆ /i: Keep file COMPRESS_FILE_IMMEDIATELY_BIT.
- ◆ /d: Delete original volume if copy process is successful. If you choose to delete the traditional volume, the new volume will retain the name of the that volume.
- ◆ /r: Remove xxx_new NSS volume and restore xxx traditional volume (use this option if you want to keep the original volume name for the new logical volume name).
- ◆ ds_container: Original volume's DS container
- ◆ ds_volumeName: If specified, the vcu conversion utility uses this name to rename or delete the original volume's DS object. Otherwise, vcu.nlm uses serverame_originalVolName as the default DS name

The parameters in this list are optional; you can include any of the parameters in the list.

To convert a volume, you specify the volume to convert and then specify the NSS pool where you want to store the data from the volume. Enter the following at the server console:

```
vcu volume_name pool_name
```

For example,

```
vcu apps poolone
```

In this example, the volume conversion utility copies the data in the volume apps and places it in Poolone.

If you want to keep the name of the traditional volume and use it for the new logical volume, enter the following at the server console:

```
vcu /r xxxx yyyy_old[ds_container[ds_volName]]
```

In this example, *xxxx* is the original traditional volume name, *yyyy_old* is the modified traditional volume name. If *xxxx* is longer than 11 characters, *yyyy_old* might be VCU0000_old, otherwise *yyyy_old* will be *xxxx_old*.

Modifying NDS Objects

The volume conversion utility lets you modify the NDS objects and physical names for both logical and traditional volumes except the VOL SYS. Use the following commands VCU commands if you want to change the NDS objects with the volume conversion process.

Volume copy

```
vcu VCUTargetVolName NSSPoolName [DS_container  
DS_NSSPool_name DS_volname]
```

Volume restore

```
vcu /r VCUTargetVolName VCUTargetVolName_old [DS_container  
DS_NSSPool_name DS_Volname]
```

Setting Up User Space Restrictions

You can limit the amount of space users can have on a volume. This is beneficial for systems that accommodate a large number of users, such as students.

When you create volumes, you can select the user space restrictions option.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click the server object and select Properties.
- 3** Right-click on the volume that contains the user object you want.
- 4** Click Properties > Attributes > NSS Attributes.
- 5** Select User Space Restrictions.

6 Click Users with Space Restrictions.

7 Select the user you want to apply restrictions to.

If the user you want is not in the list, do the following:

7a Click Add.

7b Browse for the user you want.

7c Double-click on the user.

The User Space Restriction dialog box opens. You can enter or modify the User Space Restriction parameters in this box.

8 Click Modify.

9 Enter or edit the user space amount in the Limit field and click OK.

Setting Up Directory Space Restrictions

You can limit the amount of space occupied by a directory and the files it contains.

1 Open ConsoleOne.

2 Right-click on the directory you want.

3 Click Properties and then Facts.

4 Select Restrict Size.

5 Enter the amount you want to limit the size to and click Apply.

6 Click Close.

Mounting a Logical Volume

After you set up and configure logical volumes, you can mount them.

1 From ConsoleOne, open the tree you want.

2 Right-click the server object and select Properties.

3 Click Media > NSS Logical Volumes.

4 Choose the logical volume you want to mount and click Mount.

Deactivating/Activating NSS Pools and Volumes

You might need to temporarily restrict user access to a storage pool or volume. Instead of bringing down the server, you only need to deactivate the specific volume or pool. You can deactivate both storage pools and volumes; however, you can only do maintenance (rebuild and verify) on storage pools. This means that users will be unable to access any of the volumes in the pool while you run REBUILD and VERIFY. When you use REBUILD and VERIFY a log file is generated to the root of the DOS drive. For instructions on running REBUILD and VERIFY, see [“Rebuilding NSS Storage Pools and Volumes” on page 57](#).

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Pools or Logical Volumes.
- 4** Select the storage pool or volume you want to activate or deactivate.
- 5** Click Activate or Deactivate.
- 6** Click Activate when you are ready to restore access to the volume or pool.

Salvaging and Purging Deleted Files

You can salvage files that were deleted by users, or purge deleted files to improve your file system performance. Deleted files are in a DELETED.SAV directory in the volume's root directory.

- 1** In ConsoleOne, open the tree you want.
- 2** Select the volume that contains the deleted file.
- 3** Click View > Deleted File View.

The Deleted File section lists the deleted file, the file size, the deletion date, and the user's name.
- 4** Select the file you want to Salvage or Purge.
- 5** Click Salvage or Purge on the ConsoleOne Toolbar.

Using NetWare Backup Services

You can use the backup services provided with NetWare (Backup and Restore Services) to back up your NSS data. For instructions for using Backup and Restore services, see the Storage Management Services Administration section in the *Storage Management Services Administration Guide*.

IMPORTANT: Be sure to use the version of the Target Service Agent (TSA600). Earlier versions do not work with NSS.

Rebuilding NSS Storage Pools and Volumes

To repair storage pools and logical volumes, NSS uses the VERIFY and REBUILD utilities.

VERIFY checks the file system integrity for an NSS pool by searching for inconsistent data blocks or other errors. This utility indicates if there are problems with the file system.

REBUILD verifies and uses the existing leaves of an object tree to rebuild all the other trees in the system. You need to deactivate pools (and all the volumes in the pools) before you run REBUILD so users cannot access the volumes you are rebuilding. When you deactivate a storage pool, all the volumes in the pool automatically deactivate.

REBUILD also copies errors and transactions into an error file called *volume_name.rlf* at the root of the DOS drive on your server. Every time you rebuild a particular NSS volume, the previous error file is overwritten. If you want to keep old error files, move them to another location. You can check the error file whenever an NSS volume does not come up in active mode after a rebuild.

REBUILD is not equivalent to VREPAIR. You should use REBUILD only as a last resort to recover the file system. If you use it to recover from data corruption, you will likely lose some data in the process.

When you use REBUILD and VERIFY a log file is generated to the root of the DOS drive.

- 1 To run rebuild, enter the following command at the server console:

```
nss/poolrebuild=<poolname>
```

This verifies and accounts for all blocks in the system. If the volume has errors, the errors appear on the screen. The NSS volume remains in

maintenance mode; otherwise, it reverts to the active state. You then need to mount the volume again.

You can also verify a pool by entering **nss /verify** for a list of pool names, and then selecting the pool. Verify is a read-only assessment of the pool.

NSS Server Console Commands

Use the following commands at the server console to view NSS information and perform tasks.

Online Help Commands

Enter **nss /help** or **nss /?** to access Help.

NSS Configuration Information Commands

- ♦ **nss /modules** Lists the providers, loadable storage subsystems, and semantic agents.
- ♦ **nss /status** Lists the current NSS status.
- ♦ **volumes** Lists all NetWare and NSS volumes that are mounted, including Admin volume.
- ♦ **nss volumes** Lists all NSS volumes, including the Admin volume.

Start Up NSS Commands

The following commands are only valid when NSS is loading:

- ♦ **nss /NumWorkToDos=value** Sets the number of WorkToDo entries which may be concurrently executing. NSS uses WorkToDo entries for tasks such as flushing file metadata to disk in the background. Increasing the number of WorkToDo entries might be useful on a system that is heavily used. NSS always reserves 20 WorkToDo entries. The default is 40. The range is 5 to 100.
- ♦ **nss /MailBoxSize=value** Sets the size of your mailbox. The default is 228. The range is 64 to 256.
- ♦ **nss /NumAsyncios=value** Sets the number of asynchios entries to allocate. The default is 2048. The range is 4 - 65536.

- ◆ **nss /Numbonds=value** Sets the number of bond entries to allocate. The default is 5000. The Range is 512 - 2097152.
- ◆ **nss / (No)skipLoadModules** Skips the auto loading of the NSS support modules. StartupOnly Value=Off.
- ◆ **noLSS** If specified, does not load any LSS modules.
- ◆ **allLSS** Loads all the LSS modules.
- ◆ **defaultNLMS** Loads only default NSS NLMS and any additional LSS modules specified on the command line.
- ◆ **CD9660** Loads only those modules that are essential for CD9660 support.
- ◆ **CDHFS** Loads only those modules that are essential for CDHFS support.
- ◆ **CDROM** Loads only those modules that are essential for CD9660 and CDHFS support.
- ◆ **DOSFAT** Loads only those modules that are essential for DOSFAT support.
- ◆ **UDF** Loads only those modules that are essential for UDF support.
- ◆ **zLSS** Loads only those modules that are essential for zLSS support.

Modifying NSS Startup Commands After NetWare 6 Has Booted

- 1** Create a file called NSSSTART.CFG.
- 2** Enter any NSS startup commands that you want to modify after NetWare 6 has booted.

Each NSS command should be preceded by a forward slash (/) and ended with a space. For example:

```
/xxxxx /xxxxx /xxxxx
```

- 3** Place the file in the same directory as SERVER.EXE.

Other NSS Commands

IMPORTANT: This is not a complete list of NSS commands. To see a complete listing, type `nss /help` at the server console.

- ◆ **nss /Activate=volume_name** Activates an NSS volume.
- ◆ **nss /Deactivate=volume_name** Deactivates an NSS volume.

- ◆ **nss /Maintenance=volume_name** Switches the specified NSS volume to maintenance mode.
- ◆ **nss /ForceActivate=volume_name** Forces an NSS volume to become active.
- ◆ **nss /StorageAlarmThreshold=value** Lets you set the threshold for a low storage space warning. The default is 10. The range is 0 to 1000000.
- ◆ **nss /StorageResetThreshold=value** Lets you reset the threshold for a low storage space warning. The default is 10. The range is 1 to 1000000.
- ◆ **nss / (No)StorageAlertMessages** Turns ON or OFF the low storage message to users. The default is ON.
- ◆ **nss /FileFlushTimer=value** Sets the flush time for modified open files in seconds. Increasing this number might reduce the number of writes to disk; however, it increases the amount of data that will be lost if the system crashes. The default is 10 seconds. The range is 1 to 3600 seconds.
- ◆ **nss /OpenFileHashShift=value** Sets the size of the Open File hash table (in powers of 2). If many files are used concurrently on the server, we recommend that you increase this number. The default is 11. The range is 8 to 20.
- ◆ **nss /ClosedFileHashShift=value** Sets the number of closed files that can be cached in memory. The default is 512. The range is 1 to 100000.
- ◆ **nss /Salvage** Enable salvage of deleted files on volumes.
- ◆ **nss /NoSalvage** Disable salvage of deleted files on volumes.
- ◆ **nss /UserSpaceRestrictions** Enable user space restrictions on the volume.
- ◆ **nss /NoUserSpaceRestrictions** Disable user space restrictions on the volume.
- ◆ **nss /DirectoryQuotas** Enable directory quotas on the volume.
- ◆ **nss /NoDirectoryQuotas** Disable directory quotas on the volume.

- ◆ **nss /DataShredding** Specify <vol name>[:<count>]. Enables data shredding for the volume. This is the number of times you want to shred data. The default value is 1. The maximum value is 7.
- ◆ **nss /NoDataShredding** Disable data shredding on the volume.
- ◆ **nss /FlushFilesImmediately** Synchronously flush files in the volume when they are closed.
- ◆ **nss /NoFlushFilesImmediately** Disable synchronously flushing files in the volume.
- ◆ **nss /FileCopyOnWrite** Create a copy of files in the volume when they are open for write.
- ◆ **nss /NoFileCopyOnWrite** Do not create a copy of files on the volume when they are open for write.
- ◆ **nss /UpgradeObjectsOnVolume** Upgrade objects on volumes to the current version.
- ◆ **nss /Transaction=volume name/all** Enable transactional writes in files on the volumes.
- ◆ **nss /NoTransaction=volume name/all** Disable transactional writes in files on the volumes.
- ◆ **nss /Migration=volume name/all** Enable migration of files on the volumes.
- ◆ **nss /NoMigration=volume name/all** Disable migration of files on the volumes.
- ◆ **nss /AllocAheadBlks=value** Set the number of blocks to allocate ahead on writes. The default is 15. The range is 0 - 63.
- ◆ **nss /Compression** Enable file compression on the volume.
- ◆ **nss /NoCompression=volume name/all** Disable file compression on the volume.
- ◆ **nss /NoBGCompeession** Start or stop background compression. Stop BGCompression will stop all the compression requests in the queue.
- ◆ **nss /nobackgroundchecking** Turns the background checker off.
- ◆ **nss /backgroundchecking** Turns the background checker on.
- ◆ **nss /CompScreen** Start the NSS compression statistics screen.
- ◆ **nss /Pools** Shows all of the currently available NSS pools.

- ◆ **nss /PoolAutoActivate** Activates specified pools at load time.
- ◆ **nss /PoolAutoDeactivate=pool** Leaves specified pools deactivated at pool load time.
- ◆ **nss /PoolAutoDisplay** Displays current pool load time policies.
- ◆ **nss /PoolAutoMaintenance=pool** Places specified pools in maintenance mode at pool load time.
- ◆ **nss /PoolAutoVerify** Verifies the specified pool's physical integrity at startup time.
- ◆ **nss /PoolRebuild=pool** Rebuild specified pools.
- ◆ **nss /PoolRebuild** Select pools from a menu and rebuild them.
- ◆ **nss /PoolVerify=pool** Verify the specified pool's physical integrity.
- ◆ **nss /PoolVerify** Select pools from a menu and verify their physical integrity.
- ◆ **nss /LogicalVolumePurgeDelay=value** The number of seconds before deleted logical volumes are purged. This allows time to reverse the deletion.
- ◆ **nss /LogicalVolumePurgeDelayAfterLoad** The number of seconds, after NSS loads, before deleted logical volumes are purged. This allows time to reverse the deletion.
- ◆ **nss /LogicalVolumePurgeDelayAfterContinue=value** The number of seconds to delay purging a logical volume after clicking continue. After a volume starts to purge, it cannot be salvaged.
- ◆ **nss /VolumeMaintenance** Places a volume into a maintenance mode. Volumes can be put in maintenance mode, but maintenance only occurs at the storage pool level.
- ◆ **nss /MFL=volume** Enables MFL maintenance for the volumes.
- ◆ **nss /NoMFL=volume** Disable MFL maintenance for the volumes.
- ◆ **nss /MFLVerify=volume** Get the status of MFL maintenance for the volumes.
- ◆ **nss /FixMFL** MFL maintenance for the volumes.
- ◆ **nss /GetMFLStatus** Get the status of MFL maintenance for the volumes.

- ◆ **nss /VolumeAutoActivate** Activates volume at load time.
- ◆ **nss /VolumeAutoDeactivate** Deactivates volume at load time.
- ◆ **nss /VolumeAutoDisplay** Displays current volume load time policies.
- ◆ **nss /ZLSSIOStatus** Displays current NSS IO status information.
- ◆ **nss /ZLSSPendingWriteIOs** Number of write IOs dropped to the storage subsystem. (0 = all available) [Value=1000 Range=0 - 3000].
- ◆ **nss /ZLSSVolumeUpgrade** Allow the 5.x/ZLSS VolumeUpgrade to the NetWare 6 format to occur during the installation process.
- ◆ **nss /LVDeleteStatusBasic** Displays information about deleted logical volumes.
- ◆ **nss /LVDeleteStatusSalvageable** Displays information about salvageable logical volumes.
- ◆ **nss /ZLSSPoolScan** Scans for and loads all ZLSS pools.
- ◆ **nss /space** Shows the amount of space on pools and their associated volumes.
- ◆ **nss /NoClassicDirectoryQuotas** Turns emulation of traditional NetWare volumes on or off. The default is off.

DOS FAT Commands

- ◆ **nss / (No)FATLongNames** Enables long filenames on FAT volumes. The default is ON.
- ◆ **nss /FATPartition=*partition_type_number*** Supports up to three additional partition types containing 16-bit FATs, such as /FATPartition=12,13.

Other Volumes That NSS Creates

There are three Loadable Storage Subsystems (LSS) modules: ZLSS.NSS, CDHFS.NSS, CD9660.NSS, and DOSFAT.NSS. The ZLSS module provides access to logical volumes created by the user. The CD9660 module allows access to CD ROM drives as NSS volumes. The DOS FAT module allows access to FAT 16 and FAT 12 drives.

CD-ROMs As Logical Volumes

When NSS recognizes a CD ROM, it creates a Read-only logical volume for that CD ROM. This is done automatically when you load the CDRUM.NLM. However, this volume is not added to NDS.

The new CDRUM.NLM loads the CD9660.NSS and CDHFS.NSS modules. The CD9660.NSS module loads the ISO966, CDHFS, and Macintosh Hierarchical File System (HFS) file formats. Dual formatted CDs are mounted as two separate volumes. This means both Macintosh and Windows clients will see the same CD ROM volumes with two different layouts.

The CD9660.NSS module makes the contents of the standard ISO9660 CDRUM available as a Read-Only logical volume. If the CDRUM is in the drive when CD9660 loads, it will mount as a logical volume immediately.

The volume name displayed by the NetWare volumes command is the CD-ROM volume label if one exists. Otherwise, NSS generates a volume name in this form: CDXXXXXX (where X is a number). This volume functions the same as any other NetWare volume, except that it is Read-Only.

To load both the CD9660 and CDHFS modules, enter the following at the server console:

```
load cd-rom.nlm
```

To load CD9660.NSS alone, enter the following at the server console:

```
load cd9660.nss
```

To load CDHFS.NSS alone, enter the following at the server console:

```
load cdhfs.nss
```

To unload CD9660.NSS, enter the following at the server console:

```
unload cd9660.nss
```

To unload CDHFS.NSS, enter the following at the server console:

```
unload cdhfs.nss
```

The CD ROM unload process dismounts the volumes.

DVDs as Logical Volumes

Digital Versatile Disc (DVD), the newer optical storage technology, requires Universal Disc Format (UDF) support. DVD provides access time similar to CD-ROMs, has a bigger storage space, and records data in a universal format—meaning, you no longer need a plethora of different drivers to access the data.

Requirements

To mount a DVD as a logical volume, you need the following:

- ◆ A device that reads and supports DVDs
- ◆ The UDF.NSS module loaded

Mounting a DVD

At the server console, enter

```
LOAD UDF.NSS
```

To dismount a DVD, at the server console enter

```
UNLOAD UDF.NSS
```

DOS Partitions As Logical Volumes

After you load the DOSFAT.NSS module, any DOS FAT partitions are dynamically made available as logical volumes. The volume name displayed by the NetWare **volumes** command is DOSFAT_*X*, where *X* is the drive letter, for example DOSFAT_C.

The following are some benefits of the DOSFAT.NSS module:

- ◆ **Long filenames** (such as DOS filenames for file systems such as OS/2). This uses the Windows* 95/98 method for placing long names onto a DOS FAT partition.
- ◆ **NetWare trustee rights**. Only users with supervisor rights can access the volume.
- ◆ **NetWare utilities**. Any client or server utility that accesses or manages NetWare volumes can do the same with DOS FAT partitions.
- ◆ **MAP command**. NetWare clients can map to the volume and use it as any other NetWare mapping.
- ◆ **Local partition**. Allows access to the local partition.

3

Viewing Storage Objects

We recommend that you monitor the status of your storage objects and Novell® Storage Services™ configuration to ensure they work properly and have available disk space.

Viewing Storage Devices

You can use ConsoleOne™ to view your storage devices and see which ones have free space.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Show Devices.

The left panel lists all the storage devices on your server.

- 4** Select the device you want.

For information about the Devices screen, click Help.

Viewing Your NSS Configuration

Occasionally, you might need to review the status of your NSS system storage devices or storage space. This section explains how to use ConsoleOne to view the status and usage of your storage devices, storage pools, and volumes.

View Partitioned Free Space

You can view the free space on all NSS and traditional partitions. Partitioned free space is space within an NSS partition that is not occupied by a storage pool. On a traditional partition, space that is not occupied by a NetWare® volume is free space.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Free Space.

The left panel lists the storage devices that have free space.

- 4** Select the device you want.

For information about the Free Space screen, click Help.

Any time the Devices button appears, you can click it to open the Devices screen.

View the Status of Partitions

You can view the status of both NSS and traditional partitions. For example, you might want to see what space is available on either type of partition.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Partitions.
- 4** Select the partition you want.

The left panel lists both the NSS and traditional partitions.

For information about the Partitions screen, click Help.

View the Status of Your Traditional Volumes

You can view the status of traditional volumes. For example, you might want to look at the size of a particular volume, or see the creation and modification dates.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > Traditional Volumes.

The left panel lists all the traditional volumes set up on your server.

- 4** Select the volume you want.

For information about the Traditional volumes screen, click Help.

View the Status of Your NSS Storage Pools

You can view the status of NSS storage pools. For example, you might want to see what volumes are in a particular storage pool, or how much space is in all your storage pools.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Pools.

The left panel lists the NSS storage pools.

- 4** Select a storage pool.

For information about the NSS Pools screen, click Help.

When the Show Pool button appears, you can click it to open the NSS Pools screen.

View the Status of Logical Volumes

You can view the status of logical volumes. For example, you might want to see if the volume has a space quota. You might also want to see what space is available on a volume.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.
- 3** Click Media > NSS Logical Volumes.

The left panel lists both the storage pools and the logical volumes.

- 4** Select a volume.

For information about the NSS Logical Volumes screen, click Help.

Any time the Show Volume button appears, you can click it to see the status of the volume you selected.

View the Quota and Space Usage for Specific Volumes

You can open a screen that shows the status of a specific volume that has a space quota. You can also view the status of the storage pool that volume is in.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server object and select Properties.

- 3** Click Media > Logical Volumes.
- 4** Click Properties.
- 5** Click Statistics > NSS Quota Usage.

This screen shows the volume's quota, amount of space used, and the remaining space. The screen also shows the amount of space used and the available space on the storage pool.

For information about the NSS Quota Usage screen, click Help.

View and Modify Logical Volume Attributes

You can view and modify the attribute settings for specific volumes. These attributes include data compression, data shredding, directory and user space restrictions, and snapshot. For example, you might want to apply some of the features that were not included during the initial configuration of the volume.

- 1** From ConsoleOne, open the tree you want.
- 2** Right-click on the server properties and select Properties.
- 3** Click Media > Logical Volumes > Properties > Attributes > NSS Attributes.
- 4** Make any attribute changes you want.

For information about the NSS attributes screen, click Help.

4 Troubleshooting

The following table describes some problems you might experience with Novell® Storage Services™ (NSS) and provides suggestions for solving the problems.

For additional troubleshooting information, see the [Novell Knowledgebase \(http://support.novell.com\)](http://support.novell.com).

Problem	Explanation	Solution
NSS does not recognize a Device	NSS recognizes only what the Media Manager recognizes.	For this release, NSS recognizes only hard drives and CD ROMs.
NSS does not let you create a storage pool or logical volume	<ul style="list-style-type: none">◆ Your server might not have enough free space to create more storage pools or logical volumes.◆ All logical volumes must be part of a storage pool.◆ NSS might not own the free space you want to use for a storage pool.	<ul style="list-style-type: none">◆ Ensure you have enough free space to create another storage pool or logical volume. See “View Partitioned Free Space” on page 67.◆ Before you create a logical volume, create a storage pool.◆ Create an NSS partition for your storage pools and logical volumes.
Cannot configure a logical volume	NSS might not own enough free space for another logical volume.	<ul style="list-style-type: none">◆ Add another storage device.◆ Delete a logical or traditional volume to free up space for a storage pool.

Problem	Explanation	Solution
Cannot compress a file	You must choose the file compression option when you create a logical volume.	Apply the file compression option to an existing logical volume: From ConsoleOne, click Media > NSS Logical Volumes > Properties > Attributes > NSS Attributes. Then select compression.
Cannot see your volumes or devices associated with those volumes	You might not have the appropriate multipath drivers that support the multipath functionality.	Turn off multipath support. For further instructions, see “Using Multipath Support” on page 23.
