

Novell NetWare 6.5

www.novell.com

103-000141-001

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.

You may not export or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside.

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

Storage Services Administration Guide
[April 2003 \(Public Beta\)](#)

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.

About This Guide

This documentation describes how to use Novell® Storage Services™ (NSS) to configure, activate, and maintain large volumes and numerous files without disrupting the work of end users.

This guide is intended for network administrators and is divided into the following sections:

- ◆ Chapter 1, “Overview of NSS,” on page 7
- ◆ Chapter 2, “Configuring NSS,” on page 15
- ◆ Chapter 3, “Additional NSS Configuration,” on page 19
- ◆ Chapter 4, “Viewing Storage Objects,” on page 31
- ◆ Chapter 5, “Choosing a Management Utility,” on page 35
- ◆ Chapter 6, “Choosing a Backup Service,” on page 39
- ◆ Chapter 7, “Managing RAID Devices,” on page 47
- ◆ Chapter 8, “Managing NSS Pools and Volumes,” on page 51
- ◆ Chapter 9, “Using NSS Server Console Commands,” on page 59
- ◆ Chapter 10, “Fine Tuning NSS Performance,” on page 67
- ◆ Chapter 11, “Upgrading and Copying Volumes,” on page 77
- ◆ Chapter 12, “Troubleshooting,” on page 85

Additional Documentation

NSS is the primary storage solution for NetWare 6 and later. For documentation on configuring the Traditional File Services, see the [Traditional File Services Administration Guide](#).

For documentation on configuring storage media, see the [Server Disks and Storage Devices Administration Guide](#)

Documentation Updates

For the most recent version of the *Novell Storage Services Administration Guide*, see the latest [NetWare \(http://www.novell.com/documentation/beta/nw65\)](http://www.novell.com/documentation/beta/nw65) documentation.

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 of NSS

Files are at the heart of every company, large or small. Whether your network spans continents or a few cubicles, your files become the foundation of your business. No one can afford unreliable file service—especially when the files you manage are continually growing and requiring more and more storage space.

Consequentially, these demands for more storage space have forced administrators to demand more from their file systems—they want to access their files faster and easier. More importantly, they want a file system that can scale to a growing business, be easily maintained, and better protected against corruption.

In response to this demand, Novell® Storage Services™ (NSS) promises a set of features that will help you effectively manage your storage usage and growth.

NSS Benefits

- ◆ A new journaling file system that lets you can create bigger volumes that activate (mount) quicker, store more, and resist corruption better
- ◆ Allows up to four NetWare partitions per disk
- ◆ Unlimited volumes per NetWare® partition
- ◆ Volumes can hold trillions of files in a single directory
- ◆ Faster access to data, regardless of file size
- ◆ Lower memory size: 32 MB of RAM can activate an NSS volume
- ◆ Treats free space on multiple storage devices as a single volume
- ◆ Provides enhanced CD-ROM and DVD support; CDs and DVDs are recognized and activated (mounted) automatically as volumes

For more information, see the [Novell Storage Services Product home page \(http://www.novell.com/products/netware/nw6_w_storage.html\)](http://www.novell.com/products/netware/nw6_w_storage.html).

NSS Features

NSS is concerned primarily with improving the scalability, flexibility, and reliability of your storage devices. This section describes specific NSS features that will help you do the following:

- ◆ “Use Less Memory and Gain More Speed” on page 8
- ◆ “Improve Your File System Reliability” on page 8
- ◆ “Protect Data from Corruption” on page 8
- ◆ “Maximize Free Space” on page 8

Use Less Memory and Gain More Speed

NSS lets you activate an unlimited number of volumes—a marked improvement over the traditional file system. (However, if you are using traditional APIs to access your NSS file system, you can mount only up to 255 volumes simultaneously.)

Unlike traditional volumes, NSS does not need to scan the entire file system when it activates a volume. Thus, the speed with which you can activate a volume is dramatically increased. In addition, NSS memory requirements are lower because NSS loads files into the memory *only* when you access them. Therefore, no additional memory is required when you add files and activate volumes.

NOTE: Traditional volumes use File Allocation Tables (FATs) to organize storage. FATs are filled with Directory Entry Tables (DETs) that list where files are located. When the traditional file system mounts a volume, 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.

In fact, the file system journal is read only if there is a crash. When NSS reads the journal, it finds any incomplete transactions and either completes the transaction or backs it out. This means the file system comes up after the journal fixes are made, avoiding a slow volume search for errors. This results in less server down time and is beneficial for applications such as mail services, etc.

Improve Your File System Reliability

The one thing that both you and your users agree on is how important access—quick access—to network files is. NSS has implemented the following features to help you quickly reach your data:

- ◆ Multiple paths support for adapters
- ◆ Clustering storage devices

Protect Data from Corruption

NSS includes the following features to ensure that the most current copy of your data is recoverable after a crash:

- ◆ Save data immediately
- ◆ Snapshot
- ◆ Transaction Tracking System (TTS)

Maximize Free Space

NSS includes the following features to help you maximize your free space:

- ◆ File compression
- ◆ Space restrictions (user and directory)
- ◆ Ability to activate CDs and DVDs as NSS volumes

What's New

The following features have been added for this release of NSS:

NOTE: Block suballocation, auditing, and filename locks are features that were available in the traditional file system but are not currently supported by NSS.

- ◆ “Software RAID 1 (Mirroring)” on page 9
- ◆ “Software Raid 5” on page 9
- ◆ “Moving and Splitting Volumes: DFS Enhancements” on page 10
- ◆ “Command Line Management Utility” on page 10
- ◆ “NSS Sys: Recovery Commands” on page 10
- ◆ “Pool Level Snapshot” on page 10

Software RAID 1 (Mirroring)

RAID 1 was enhanced in order to decrease the large amounts of memory that was required to mirror large partitions. Because disk sizes continue to increase to meet our storage demands, RAID 1 now uses the same amount of memory to mirror partitions, regardless of the size.

In the event of a server failure, you no longer have to remirror all the partitions because RAID 1 automatically tracks and remirrors any blocks that were in the process of being written to when the server crashed. Also, if the server crashes during a mirroring process, when the server comes back up, the mirroring process resumes where it left off.

Additional RAID 1 benefits:

- ◆ Multi-processor enabled
- ◆ Mirrors only those partitions that are out of sync
- ◆ Mirrors partitions that had no previous mirror

Create software RAID devices using the Web-based iManager utility or the console-based NSS Management Utility (NSSMU). For more information, see [“Creating a RAID Device” on page 48](#).

Software Raid 5

RAID 5 improves input/output speeds and data protection striping data to multiple segments and interleaving parity. The parity segment is used to store the parity of all the other data segments. If a single segment is lost, the parity segment is combined with the existing segments to reconstruct the lost data. The parity segment can reconstruct only one lost segment at a time. The parity features adds data protection to RAID 5 that RAID 0 does not have in the event of a disk failure. RAID 0, however, provides faster performance than RAID 5 because it does not have the overhead of the parity maintenance. The choice of whether to use RAID 0 or RAID 5 will depend on whether you want faster performance or data protection.

You need three segments to create a RAID 5 device. The storage capacity of the RAID 5 device will be the size of an individual segment times the amount of segments you have minus one (you lose one segment worth of data to store the parity). You can have up to eight segments in our implementation of a software RAID device. Remember that RAID requires that all segments be the same size.

Create software RAID devices using the Web-based iManager utility or the console-based NSS Management Utility (NSSMU). For more information, see [Chapter 7, “Managing RAID Devices,” on page 47](#).

Moving and Splitting Volumes: DFS Enhancements

The ability to move and split volumes makes it easier to get rid of old hardware or free up disk space.

If you split a volume using Distributed File Services (DFS), NSS creates a junction that points to the location in the same or different server where you moved the data. The existence of junctions is typically transparent to the user. Whenever a user logs in to the server to access a volume, the junction redirects the user's request to the appropriate volume and server without requiring the user to know the new location.

For example, suppose you have a server that holds all of your users' home directories. You know that you have enough free space for the data; however, the files are accessed so frequently that it increases the server load and decreases performance. Using DFS, you can split the volume, moving the home directories of the users with last names starting with L-Z to another volume in the same server or a different server. After the volume is split, NSS uses the junction to redirect traffic from the L-Z users. They can log in to the original server even though their data is now located on a different volume.

To learn about the specific server commands required to move and split volumes, see [“Moving and Splitting Volumes” on page 23](#).

To learn more about DFS, see [“Using Distributed File Services \(DFS\)” on page 21](#).

Command Line Management Utility

The NSS Management Utility (nssmu.nlm) is loaded from the command line of the server console and gives you limited creation and management of the following NSS components: devices (including RAID), partitions, volumes, and pools.

This utility is not meant to replace ConsoleOne, iManager, or any other Web-based management utility. However, this utility is the only management tool that can access your server if you accidentally delete or rename the SYS: volume.

NSS Sys: Recovery Commands

New NSS sys: recovery commands have been added to give you additional control over your sys: volume.

For more information, see [“System Volume Recovery Commands” on page 64](#).

Pool Level Snapshot

Pool level snapshots is another way to keep a consistent copy of your data for back up purposes.

For more information, see [“Pool Snapshots” on page 40](#).

Don't Overlook These Changes

The following NSS features now behave differently than they did in previous releases:

- ◆ [“Cddvd” on page 11](#)
- ◆ [“Vcu.nlm” on page 11](#)
- ◆ [“Name Caching” on page 11](#)

- ◆ “CIFS and AFP Protocols” on page 12
- ◆ “Admin Volume” on page 12

Cddvd

You no longer need to use `udf.nss`, `cd9660.nss`, or `cdhfs.nss` CD, DVD, and Macintosh* support. The functionality of these three NSS commands has been combined into one command: `cddvd.nss`

When you load `cddvd.nss`, it uses the same set of switches (`UpdateCDToNDS`, `LoginUpdateToNDS`, `LeaveCDInNDS`, and `LockCDMediaIn`) that were used in the previous commands. In addition, the `LoadISOVolumeFirst` switch has been added.

Dual formatted CDs are activated as one volume and its format is decided by `cddvd.nss`. By default, `cddvd.nss` loads the formats in descending order, with the UDF format first, ISO9660 next, and CDHFS last. If you have a UDF bridge disk, which supports both UDF and ISO9660 formats, you can instruct NSS to load ISO9660 first by entering **LoadISOVolumeFirst** at the server console.

Vcu.nlm

The following changes have been made to `vcu.nlm`:

- ◆ A new file, called `irf.out` (Inherited Rights Filter), is now created by `vcu.nlm`. This file contains information about whether the Inherited Rights Filters were moved correctly during a volume copy or restore. Mismatches will be preceded by three asterisks (***)

The `irf.out` is stored in the same location as `errors.out` (at the root directory of the destination volume).

- ◆ You can now use `vcu.nlm` to copy and restore NSS volumes.

With previous versions, you could use `vcu.nlm` to copy only traditional volumes to NSS volumes.

- ◆ You can now run `vcu.nlm` from a Windows* client.
- ◆ The optional input of `dsContainer`, `dsPoolName`, and `dsVolName` has been deleted.
- ◆ The `/i` switch has been deleted.

NOTE: For a listing of available switches and their functionality, at the server console, enter `vcu /h`

- ◆ The `/m` switch has been added.

This switch lets users rename a volume’s default eDirectory™ object.

For more information on how to use `vcu.nlm` to copy volume data from traditional or NSS volumes to NSS pools, see [Chapter 11, “Upgrading and Copying Volumes,” on page 77](#).

Name Caching

The `/NSS NameCachSize=` command is no longer a start-up command. You can now specify the maximum number of entries the system can use in the cache. Additionally, the size of the cache has now been automated; the cache automatically resizes itself during set intervals based on the file system usage.

Mirroring Partitions and Hot Fix

Hot Fix is not supported in NetWare 6.5 and later because the function is provided in modern hardware at the device level. This means that you will no longer need to specify a Hot Fix space in order to mirror a partition.

CIFS and AFP Protocols

With this release, the CIFS and AFP protocols will now be installed, by default, when you install your NetWare 6.5 server. Additionally, these protocols can access only NSS pools and volumes.

Admin Volume

A new file called `files.cmd` was added to the Admin volume (`Admin:manage_nss\files\files.cmd`). Applications like NetStorage and other end user browser-based tools use this file to communicate with the file server. `files.cmd` is a virtual file and generates dynamic information that lets users see and modify their files through a browser.

Every user on the file server will have access to this file and will thus inherit the following default rights: Read, Write, and Scan. This also means that the Admin volume will now be visible to end users—in previously versions of NetWare, this was not the case. End users cannot delete this file nor modify its properties.

If you want to restrict access to this file or the Admin volume, you must do so by modifying a configuration file called `trustees.xml`. For more information, see “[Assigning Rights with Trustees.xml](#)” on page 55.

For more information on virtual files, see the [Virtual File Services for NetWare \(http://developer.novell.com/ndk/doc/vfs/index.html?page=/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html\)](http://developer.novell.com/ndk/doc/vfs/index.html?page=/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html) located in the NDK documentation.

Understanding NSS

NSS differs from the traditional file system in many ways; however, the most important differences are in the way NSS utilizes space—free or otherwise.

NSS uses free space from multiple storage devices, which allows you to create unlimited volumes and store up to 8 trillion files in a single volume—up to 8 terabytes in size. Similarly, this 64-bit file system handles a virtually unlimited number of file objects.

NSS volumes, which are sometimes referred to as Logical volumes, are stored in storage pools that are located on NSS partitioned space.

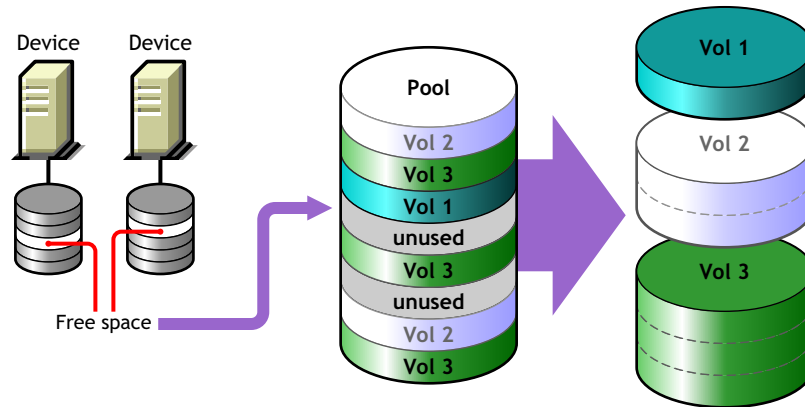
IMPORTANT: In NetWare 5, NSS runs parallel with the traditional NetWare file system. Beginning with NetWare 6, NSS is the primary storage and management system. However, you can continue to maintain traditional partitions and volumes along with NSS. This documentation focuses primarily on NSS storage pools and NSS volumes. If you need instructions on setting up and maintaining traditional volumes, see the [Traditional File Services Administration Guide](#).

Storage Pools

A storage pool is an amount of space that is obtained from one, many, or all of your storage devices. After you have created a storage pool from combining your storage space, you can begin to create NSS volumes from the space in the storage pool.

The following graphic is a visual representation of how NSS uses free space on multiple devices to create a storage pool:

Figure 1 NSS Architecture



When you create a storage pool, it is automatically created on an NSS partition. You can only have one storage pool per partition.

When you install NSS, a storage pool named `sys:` and an equal-sized volume named `sys:` is automatically created. Supporting the `sys:` volume makes management simpler, because now all volumes on the NetWare server can be NSS volumes.

NSS (Logical) Volumes

The volumes you create on NSS storage pools are called NSS volumes (these are sometimes referred to as Logical volumes). An NSS volume can be either set to a specific size or set to grow dynamically to the size of a pool it belongs to. You can add any number of volumes to a storage pool.

Because there is no limit to the number of volumes you can create, it is possible that the combined size of all the volumes taken together is larger than the storage pool itself. This is allowed; however, the size of each *individual* NSS volume *cannot* exceed the overall size of the storage pool.

NSS refers to this process as "overbooking." NSS overbooks volumes by utilizing space from another volume in the storage pool whose space is not being fully used. This process is based on the assumption that there is a constant ebb and flow in user consumption. It is quite unlikely that all volumes will be accessed all at once and at 100% of their capacity. In this sense, NSS employs a type of borrowing system amongst the NSS volumes of a storage pool. This way, all of your disk space is used by those who need it most by balancing user needs. Furthermore, it is not always necessary to add more disk space when some users reach or exceed their volume limit.

For example, suppose you have a 300 GB storage pool. From this storage pool, you create two NSS volumes of 200 GB. You can define two 200 GB NSS volumes out of a storage pool of only 300 GB only if you feel comfortable that the NSS volumes will not both exceed 75 percent capacity (150 GB) and therefore exceed the overall size of the storage pool. If one NSS volume does hit 150 GB, but the other volume stays under 100 GB, your overbooking plan worked.

2

Configuring NSS

Novell® Storage Services™ (NSS) is the default file system that ships with NetWare® 6. This section describes how to configure NSS, by completing the following steps:

1. [Meet the prerequisites.](#)
2. [Plan for your storage needs.](#)
3. [Create a storage pool.](#)
4. [Create an NSS volume.](#)

Meet the Prerequisites

- A server running NetWare 6
- A least 10 MB free space on the server (to create an NSS storage pool and a logical volume)

Plan For Your Storage Needs

Before you begin to carve your physical disks, you should consider what your storage needs are and how you can effectively manage and divide your storage space to best meet your needs.

The following table lists some decisions that you should consider when planning for your storage needs:

Storage Components	Things to think about
RAID Devices	NetWare supports a software version of RAID 0, 1, and 5.
NOTE: You might have hardware that contains the RAID functionality, therefore no RAID software would be necessary.	RAID 0 is simply disk striping. Disk striping improves the performance of your servers by striping your data across all of your storage devices; however, RAID 0 does not provide any protection against data loss. If one element in the RAID device malfunctions, the entire configuration is unusable. RAID 1 is our mirroring strategy. It is helpful to determine your mirroring strategy before you begin disk carving. If your hardware has mirroring capabilities, you do not need to set up an additional software RAID configuration. RAID 5 is disk striping and mirroring combined. For more information, see Chapter 7, “Managing RAID Devices,” on page 47.

Storage Components	Things to think about
Free Space	<p>It is good to have an idea of how much free space you will need from each device. This will help you during the disk carving phase of the NSS configuration.</p> <p>For more information on viewing your storage objects, see "Viewing Storage Devices" on page 31.</p>
Pools	<p>How do you want to organize your pools?</p> <p>Do not allow your volumes to span more than one pool.</p>

Creating a Storage Pool

After creating and setting up your partitions (sometimes referred to as pool segments), 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 partitioned disk space. However, you can select unpartitioned space. This will create a partition and make the storage pool 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 on the server object and select Properties.
- 3** Click Media > NSS Pools > New.
- 4** Enter a name for the new storage pool and click Next.
- 5** Select the existing NSS partitions or unpartitioned free space that you want to use for the pool and then click Next.
- 6** (Conditional) You will fill in different information depending on whether the partition or free space that you selected in the previous step is sharable for clustering.

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 selected. Click Next if you want to activate this pool on creation; otherwise, uncheck the box and click Next.

If the partition or free space that you selected *is* sharable, the Create Pool dialog box appears with the Activate on Creation and Cluster Enable on Creation check boxes, checked. When the Cluster Enable on Creation check box is checked, you must fill in the following shared pool clustering parameters:

- ◆ Virtual Server Name
- ◆ CIFS Server Name
- ◆ IP Address
- ◆ Advertising Protocol

- 7** Click Finish.

Continue with the next section, "[Creating a Logical \(NSS\) Volume](#)."

Creating a Logical (NSS) Volume

After creating storage pools, you can create any number of logical volumes for the storage pools 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 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 on the server object and select Properties.
- 3** Click Media > NSS Logical Volumes > New.
- 4** Enter a name for the volume and click Next.
- 5** Select the storage pool where you want to store the unpartitioned space or NSS partitioned space, and then do *one* of the following:
 - ♦ Enter a quota size for the volume if you want to limit the size of the volume.
 - ♦ Mark “Allow volume quota to grow to the pool size” if you want the volume to expand to the size of the pool.
- 6** Click Next.
- 7** (Conditional) If you selected a device that does not already contain a pool, a new dialog box appears, asking you to specify a pool name and size. Enter the information and then click OK.

Then, if the space that you selected is on a shared device, you will see a dialog box with Cluster Enable on Creation checked by default. If you uncheck this option, the clustering parameters disappears. If you keep this option checked, you must fill in the following shared pool cluster parameters:
 - ♦ Virtual Server Name
 - ♦ CIFS Server Name
 - ♦ IP Address
 - ♦ Advertising ProtocolsClick OK when finished.
- 8** Click Next.
- 9** Select the attribute settings you want for the new logical volume, then click Finish.

Volume Attributes

- ♦ **Backup.** 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 in order 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.

10 Click Finish.

You can now store and manage files in the NSS system. If desired, continue with the next section, [Chapter 3, “Additional NSS Configuration,” on page 19.](#)

3

Additional NSS Configuration

This section describes the following NSS configuration options:

- ◆ “Using Multiple Paths to Access Devices” on page 19
- ◆ “Marking a Device as Sharable for Clustering” on page 21
- ◆ “Using Distributed File Services (DFS)” on page 21
- ◆ “Changing the Purge Delay Setting” on page 26
- ◆ “Using the Flush Files Immediately Feature” on page 26
- ◆ “Setting Up Space Restrictions” on page 26
- ◆ “Using TTS” on page 27
- ◆ “Preventing Unauthorized Access to Deleted Files (Data Shredding)” on page 28
- ◆ “Enabling File Compression” on page 28
- ◆ “Enabling Pool Multiple Server Activation Prevention” on page 28
- ◆ “Activating CD-ROM, DVD, and DOS Partitions as NSS Volumes” on page 29

Using Multiple Paths to Access Devices

Before NetWare 6, if you wanted to have multi-path support to a device, you had to manually configure each driver to include multiple paths. However, this process was time-consuming and did not have a failover plan.

With the multi-path support in NSS, you can assign priority levels to each path from the server console. That way, if you have two adapters pointing to the same device and one of the adapters fails or the cable goes bad, the second adapter can still communicate with the device, thereby increasing the fault tolerance of your network. This functionality is useful if you have more than one adapter connected to the same SCSI bus, or if you have a SAN with multiple adapters pointing to the same switch in your network.

The NLMs needed to take advantage of multi-path support (MM.NLM and NWPA.NLM) are installed and loaded when you install NetWare 6 and later.

Assigning Multiple Paths

When you use the server console to assign multiple paths, you can also assign a priority level to each path, according to 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
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 go 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 a 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 in 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.

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. These devices can be shared by multiple servers in a cluster. For more information about clustering, see the [Novell Cluster Services Overview and Installation Guide \(http://www.novell.com/documentation/lg/ncs6p/index.html\)](http://www.novell.com/documentation/lg/ncs6p/index.html).

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.

To cluster a device:

- 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.

Using Distributed File Services (DFS)

- ◆ [“Overview” on page 21](#)
- ◆ [“DFS Requirements” on page 22](#)
- ◆ [“Creating Junctions” on page 22](#)
- ◆ [“Deleting a Junction” on page 23](#)
- ◆ [“Moving and Splitting Volumes” on page 23](#)
- ◆ [“Managing the VLDB Service” on page 25](#)

Overview

Distributed File Services (DFS) is a service, installed by default with NetWare 6 and later, that allows you to create junctions. A junction functions like an alias in Novell® eDirectory™ that

appears to be a normal eDirectory object. A DFS junction appears to be any normal folder or directory in your file system. However, the junction is simply a virtual directory that points to the root of a volume where you relocated data. The new location can be in a different pool on the same server or on a different server in the same eDirectory context.

Whenever you create a junction, the Volume Location Database (VLDB) service assigns the junction a unique ID, which is automatically stored in a database. Whenever a user views the contents of a junction, the client looks at the unique ID, 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.

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 servers 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 you 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 spread among different volumes or servers.

DFS Requirements

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

- You can create junctions only on NetWare 6 and later NSS volumes; however, you can create junctions on NetWare 6 and later NSS volumes that point to either NetWare 5.1 NSS volumes or NetWare 5.1 traditional volumes.
- You can create junctions only within the same eDirectory tree.
- You can access DFS junctions with the latest Novell Client™ via the NCP protocol, Web-based services via the XTier protocol, and the Microsoft Client via the CIFS protocol.

If you are using an unsupported protocol to access a DFS junction, you will see something that looks like a small file, but you will be unable to read or open it.

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 eDirectory 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 service within its own domain.

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

To create a Management Context:

- 1** Open ConsoleOne®.
- 2** Right-click an O or an OU in the eDirectory tree and then click New > DFS Management Context.
- 3** 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 service and house the database. Select more than one server if you want to keep a copy of the database on another server to prevent data loss or corruption.

- 4 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 service is up and running, you can create junctions.

To create junctions:

- 1 In ConsoleOne, create a subdirectory or use an existing subdirectory of any NetWare 6 and later NSS volume.
- 2 Right-click the subdirectory and then click New > Shortcut > Junction.
You can also create a junction at the root of a volume. 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 eDirectory 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.

Deleting a Junction

To delete a junction, right-click the junction and then click Delete Shortcut > Junction.

Moving and Splitting Volumes

The following information applies to moving and splitting volumes:

- ♦ You can split only NSS volumes.
- ♦ When you move or split volumes to a destination server, the destination volume must be an NSS volume.
You can move Traditional volumes as long as the destination volume is an NSS volume.
- ♦ The destination volume must physically exist before you move data to it—it does not get created during the volume move or split process.
- ♦ When you enter your source path and username, use the *full* context, including leading and trailing dots. For example, if your destination server was named music-master, you would specify the full context of where this server resides.
For example: **.music-master.novell.music-tree.**
- ♦ You must create a DFS Management Context before you can move or split volumes. For more information, see [“Creating Junctions” on page 22](#).
- ♦ If the server crashes during a volume move or split operation, when the server comes back up, the operation will resume where it left off.

- ◆ You can enter as many requests as you want; however, DFS can run only four active operations at one time. After four requests, any additional requests will have to wait for one of the running operations to complete or you will need to pause an operation in order to let another operation run. For more information on pausing an operation, see “[Server Console Commands](#)” on page 24.

Open Files and Moving or Splitting Volumes with DFS

If files are in use during a volume move or split operation, the operation will complete; however, DFS does not copy those files to the new location. Instead, the open files are left behind in the old location. DFS creates a list of the files that it does not move.

Whenever DFS completes copying all the files it can, it automatically returns one time to the remaining files and tries again to copy the files to the new location. If any of these files are in use and still cannot be copied, then the move or split operation goes into the Files Skipped state where the administrator must take action.

The administrator can Resume the process to try again to copy the files. If files still cannot be copied, the process returns to the Files Skipped state. The retry can be repeated, as necessary, until it copies all the files successfully.

The administrator can Cancel the process from the Files Skipped state to force the process to complete the move or split operation, but the files in the skipped list will not be copied. After the DFS junction is in place, the user can no longer access the files that remain behind, but the administrator can. The administrator must decide what to do with the remaining files, such as to move them manually to the new location, to leave them where they are, or to delete them.

NSS does not track files might have been newly created during the move or split process. These files will remain behind in the old location. Because they did not exist at the beginning of the move or split process, NSS does not include these files in the list of files it could not move.

NSS does not track modifications the user makes to files while the move or split process is in progress. It is possible for these changes to be lost if a user accesses and modifies a file in the original location after NSS copies the original file to the new location, but before NSS activates the DFS junction and deletes the original files. Because NSS does not compare the time stamp of the files before deleting them, the changed file would be deleted as if it were the original file.

Server Console Commands

The following command line parameters have been added to the DFS functionality.

WARNING: DFS uses junctions to move or split volumes. If you don't have DFS setup, you could lose data when trying to move or split volumes using DFS commands. For instructions on setting up DFS, see “[Creating Junctions](#)” on page 22.

To view these commands, enter **volmn help** at the server console.

- ◆ **volmn cancel *opID*** terminates a running operation. The *opID* is a unique number that is automatically assigned every time you begin an operation.

To view the unique ID of a particular operation, enter **volmn status**.

- ◆ **volmn move *srcVol destVol destServer*** moves a volume to another volume on the same server or a different server.

For example, if you wanted to move your Elvis volume to the Music volume on the music-master server, you would enter the following at the server console:


```
volmn move elvis music .music-master.novell.music-tree.
```

- ◆ **volmn pause *opID comment*** pauses a running operation. The *opID* is a unique number that is automatically assigned every time you begin an operation. You might want to pause an operation if it's creating too much network traffic or load on the server.

To view the unique ID of an particular operation, enter **volmn status**. You also have the option of adding a comment at the end of the command to identify the reason why the operation was paused.

- ◆ **volmn resume *opID*** continues a paused operation. The *opID* is a unique number that is automatically assigned every time you begin an operation.

To view the unique ID of an particular operation, enter **volmn status**.

- ◆ **volmn split *srcPath destVol destServer*** splits and moves all of the contexts below the specified source path to the designated destination volume and server. It then replaces the original subdirectory with a DFS junction.

For example, if you wanted to move the Beatles subdirectory of your MP3 volume to the music volume on your music-master server, you would enter the following at the server console:

```
volmn split mp3:beatles music .music-master.novell.music-tree.
```

- ◆ **volmn status** displays the status of previous operations.

Managing the VLDB Service

To access the VLDB management service:

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

- 1 In ConsoleOne, right-click any server that has the VLDB service running.
- 2 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 and the database location.

The management service also allows you to perform the following:

- ◆ Adjust processing threads
- ◆ Start and stop the VLDB service
- ◆ Add and remove replicas

IMPORTANT: At this time, you must access the VLDB service from the server that you want to remove or add a replica to. In ConsoleOne, right-click the server and then click Properties > Supported Services > NSS VLDB. Once you are in the management service, click either Add Replica or Remove Replica.

- ◆ Repair Database

You can repair a database in the following 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 service running on more than one server. The VLDB service gets a copy of the database from another server that is currently running the service.

- ◆ Rebuild

When you rebuild a database, the VLDB service walks the eDirectory tree, looks at volume and server objects, and then completely rebuilds the database from scratch.

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=xx
```

In this command, you replace xx with the actual number of seconds to delay the auto purge. The default is 345600 (4 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.

Using the Flush Files Immediately Feature

The Flush Files Immediately feature helps protect against data corruption by saving your file data to a disk immediately after you close the file, instead of waiting for the next flush cycle. It also protects you against the risk of data being lost between flush cycles.

Normally, when data is written to disk by a user or process, the data isn't actually written to disk immediately. Server operating systems juggle hundreds or thousands of demands at once and use memory to speed hardware-dependent functions, like disk writes, whenever possible. So a file that gets written to disk may stay in the server memory for a short time (a matter of seconds or less) before actually changing bit patterns on a hard drive.

Optionally, with Flush Files Immediately, this server-juggling process can be overridden, forcing file changes out to the disk controller hardware immediately after the file is closed. The performance penalty comes on the rest of the server operations, as they wait for a disk write, or dozens of disk writes, every few seconds on a busy server.

To enable the Flush Files Immediately feature, enter

```
nss /FlushFilesImmediately
```

To disable the feature, enter

```
nss /NoFlushFilesImmediately
```

Setting Up Space Restrictions

NSS lets you put a hard limit on the amount of disk space a user can own. 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.

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

User Space Restrictions

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.

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 > Facts.
- 4** Select Restrict Size.
- 5** Enter the amount you want to limit the size to and click Apply.
- 6** Click Close.

Using TTS

The Transaction Tracking System™ protects simple database applications by backing out transactions that are incomplete due to a system failure. By default, TTS™ works with the traditional file system. TTS will not run on NSS and traditional volumes at the same time. If you have both traditional and NSS volumes on your system and you want to run it on NSS volumes instead of traditional volumes, you must first disable TTS on all traditional volumes.

To do this, add **DISABLE TTS** to the AUTOEXEC.NCF and then enable TTS on *each* NSS volume you want it on by entering **nss /transaction=volname** at the server console.

Preventing Unauthorized Access to Deleted Files (Data Shredding)

Data Shredding hides deleted and purged files by overwriting them 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. Data Shredding truly erases files. Be warned that only files that have been purged are shredded, so there remains a grace period between when the file is deleted and then purged for users who might delete the wrong file.

Also, there is a performance penalty for using the disk and system resources needed to overwrite the shredded file. Unless you must use this feature for security reasons, it should be disabled, as data shredding consumes a great deal of disk I/O bandwidth.

Data shredding can be disabled as a volume attribute from ConsoleOne, or via this command:

```
nss /nodatashredding=volume_name
```

Enabling 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. To enable file compression, enter

```
nss /Compression
```

To disable file compression on certain volumes, enter

```
nss /NoCompression=volume_name
```

To disable file compression on all volumes, enter

```
nss /NoCompression=all
```

To view the NSS compression statistics screen, enter

```
nss /CompScreen
```

Enabling Pool Multiple Server Activation Prevention

If enabled, the pool Multiple Server Activation Prevention (MSAP) prevents some accidental activations of a pool on more than one server at a time. It will not catch all multiple activations. You should never purposely attempt to activate a pool on two servers at the same time. MSAP is not meant as a replacement of clustering software that controls shared pools.

MSAP protects pools on systems that do not have clustering installed but are attached to a shared disk by accident. For example, a pool might not be marked with the Sharable for Clustering attribute, but it exists on shared disks seen by multiple servers.

Pool MSAP also protects against dangerous conflicts that can occur if you disable the Sharable for Clustering flag in order to force an activation of a shared pool, or use `override=shared` when activating a pool. If it detects a conflict, it deactivates the pool before massive corruption occurs.

If you unload Novell Cluster Services NLM software or are not running it, pool MSAP provides an extra level of protection. The clustering software watches pools that are marked with the

Sharable for Clustering attribute; MSAP detects conflicting I/O from multiple servers and prevents corruption of pools even on devices that are marked as Not Sharable for Clustering.

In some cases, the MSAP software causes pools to take up to 30 seconds to activate. This delay might occur on the next pool activation after the Server ID or the Cluster ID changes for a given server pool. The Server ID changes if the registry is corrupted. The Cluster ID changes if the Cluster NDS object ID is lost.

If MSAP is enabled, all active ZLSS pools will be read every 14 seconds. If your storage media are not shared between multiple servers such as in a SAN, you can clear the `zpool_feature_msap` bit. You should not clear this bit if your pools are on physically shared storage media.

If a pool can be accessed by older servers not running the Support Pack with the MSAP software, then multiple pool activations can still occur.

MSAP does not protect against partition conflicts for pools. It does not prevent multiple servers from creating a pool in the same partition.

Activating CD-ROM, DVD, and DOS Partitions as NSS Volumes

NSS now allows you to activate CD-ROM and DVD drives as NSS volumes through CDDVD.NLM, which is loaded automatically when NSS starts.

CD-ROMs as NSS Volumes

The new CDDVD.NSS loads the ISO966, CDHFS, CD-ROM, and Macintosh Hierarchical File System (HFS) file formats. When NSS recognizes a CD-ROM, it creates a Read-only NSS volume for that CD-ROM. This is done automatically by CDDVD.NSS. However, this volume is not automatically added to eDirectory, but it can be added later by enabling a CDDVD.NSS switch. For more information on the CDDVD.NLM switches and formatting, see [“Cddvd” on page 11](#).

If a CD-ROM is in the drive when CDDVD.NSS loads, it is immediately activated as a NSS volume. The volume name is the CD-ROM volume label if one exists and is acceptable to NetWare. Otherwise, NSS generates a volume name in this form: `CD_XXXX` (where *X* is a number). This volume functions the same as any other NetWare volume, except that it is Read-only.

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

```
UNLOAD CDDVD.NSS
```

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

```
LOAD CDDVD.NSS
```

DVDs as NSS Volumes

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

Requirements

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

- A device that reads and supports DVDs

- ❑ The CDDVD.NSS module loaded

Activating a DVD

At the server console, enter

```
LOAD CDDVD.NSS
```

To unload CDDVD.NSS, at the server console enter

```
UNLOAD CDDVD.NSS
```

DOS Partitions as NSS Volumes

After you load the DOSFAT.NSS module, any DOS FAT partitions are dynamically made available as NSS 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.

4

Viewing Storage Objects

Occasionally, you might need to review the status of your NSS 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.

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 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.

Viewing 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.

Viewing 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.

Viewing 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.

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

Viewing 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.

Viewing 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.

5

Choosing a Management Utility

ConsoleOne[®] or iManager is the primary utility used to configure and maintain NSS logical volumes.

NOTE: NWADMIN, NWCONFIG, and NSS Menu are not compatible with some of the NSS features.

You can manage NSS through the following utilities:

- ◆ “iManager” on page 35
- ◆ “ConsoleOne” on page 36
- ◆ “NetWare Remote Manager” on page 37
- ◆ “Command Line Management Utility” on page 37
- ◆ “Virtual File Services (VFS) and Scripts” on page 37
- ◆ NetStorage

iManager

Novell[®] iManager is a browser-based tool used for administering, managing, and configuring Novell eDirectory[™] objects. Novell iManager gives you the ability to assign specific tasks or responsibilities to users and to present the user with only the tools (with the accompanying rights) necessary to perform those sets of tasks.

For more information, see the *Novell iManager Administration Guide* (<http://www.novell.com/documentation/lg/imanager/index.html?imanager/data/a6160f7.html>).

Using iManager for Storage Management

Use the Novell iManager Storage Management interface to manage the physical, virtual, and logical storage infrastructures for each server.

The NetWare file storage and management system provides an easy, reliable, and secure way to organize, consolidate, and manage data for your enterprise. To organize storage in ways that make sense to end users, you can create virtual devices, storage pools, and logical volumes that virtualize the storage devices in a mixed-media environment.

Manage Devices

Configure, mount, and maintain a wide selection of storage devices, including direct attached storage devices, networked storage devices in a Fibre Channel or iSCSI storage area network (SAN), and hardware device arrays. Other device management options include the following:

- ◆ Initialize devices
- ◆ Scan for devices, if not found automatically by the media manager

- ◆ Enable device sharing to support high-availability server clusters

Manage Software Raid Devices

Create and manage software RAID devices (RAID-0, RAID-1, and RAID-5) to improve storage performance and reliability.

Manage Pools

Create and manage storage pools to efficiently use all free space. Enable the pool snapshot feature to preserve point-in-time views of data pools and to support data recovery and backup.

Manage Volumes

Create and manage logical volumes, including these and other key attributes:

- ◆ Compressing volumes to extend storage capacity
- ◆ Moving and splitting volumes to reorganize and redistribute storage on the same server (or to other servers) in response to changing business needs
- ◆ Data shredding to scramble any data that end users delete to prevent subsequent access to the information via a disk

Manage Distributed File Services

DFS junctions use a Volume Location Database (VLDB) service to locate and transparently redirect traffic to volumes. DFS management includes the following options:

- ◆ Create one or more DFS management context(s) where servers can host the VLDB
- ◆ Create, maintain, and monitor active VLDB services
- ◆ Create and manage DFS junctions

ConsoleOne

ConsoleOne is the primary utility used to configure and maintain partitions, storage pools, and volumes in NSS. It is a GUI-based application that runs on the server or 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 workstation.

- 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\nssjavalib.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
```

To manage NSS from ConsoleOne:

- 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.

For more information, see the *ConsoleOne User Guide* (<http://www.novell.com/documentation/lg/consol13/index.html>).

NetWare Remote Manager

NetWare[®] Remote Manager is a utility that provides most of the functionality of MONITOR, along with some functionality from other utilities available at the server console, but is available from a Web browser. You can use NetWare Remote Manager to monitor your server's health, change the configuration of your server, or perform diagnostic and debugging tasks.

For information on accessing, understanding the layout, and configuring NetWare Remote Manager, see the Accessing NetWare Remote Manager section of the *NetWare Remote Manager Administration Guide* (http://www.novell.com/documentation/lg/nw6p/index.html?port_enu/data/a7mkeva.html).

Command Line Management Utility

The NSS Management Utility (nssmu.nlm) is loaded from the command line of the server console and gives you limited creation and management of the following NSS components: devices (including RAID), partitions, volumes, and pools.

This utility is not meant to replace ConsoleOne, iManager, or any other Web-based Management Utility. However, this utility is the only management tool that can access your server if you accidentally delete or rename the sys: volume.

Virtual File Services (VFS) and Scripts

Virtual File Services (VFS) provides methods that allow you to manage services such as NSS using standard file system functions. Using VFS and a scripting or GUI-based interface, you can view the status and statistics for your system and change the system parameters.

NSS provides a special administration volume—known as the admin volume (named `_admin`)—that exists on all servers. This volume uses no disk space and is created at startup time. Using VFS and the services provided by files that are created on the admin volume, your administrator can potentially control all server management functions.

For more information and instructions, see *Virtual File Services for NetWare* (http://developer.novell.com/ndk/doc/vfs/index.html?vfs__enu/data/h5ubo7tt.html) in the NDK documentation.

6

Choosing a Backup Service

You can use the backup services provided with NetWare (Backup and Restore Services) to back up your NSS data. For instructions on using Backup and Restore services, see the [\(link to eDirectory doc, Backup and Restore section\)](#).

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

Additionally, NSS has the following features to assist you in the back up process:

- ◆ “File Level Snapshot” on page 39
- ◆ “Pool Snapshots” on page 40

File Level Snapshot

The Snapshot feature lets you back up open files. When you use the Snapshot feature, a snapshot file is created of the open file that contains a copy of the metadata. When blocks are modified in the open file, the changes are saved in the snapshot file. That way, if the server crashes, you still have what existed before the crash and also any changes that you were in the process of making saved in the snapshot file.

Snapshot settings apply to volumes and requires some manual setup; however, after the parameters are defined, the process is automatic.

If you select File Snapshot on a volume, the backup utility copies all the meta-data (owner, created date/time, modifier, etc.) to the snapshot file. 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
```

After you set up File Snapshot on your volumes, you need to deactivate the volumes. Then you need to activate 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 calls also return information about File Snapshot:

- ◆ `GetFileSize`
- ◆ `GetFileHoles`

Pool Snapshots

A pool snapshot is a metadata copy of a storage data pool. The snap preserves a point-in-time view of a data pool, and supports data backup and recovery for that pool. You can create, manage, and delete a pool snapshot for any pool on your server.

You can take a snapshot, and activate it later to access the original pool's data as it existed at the time of the snap. Both the pool and its snapshot(s) can be active and available concurrently. You access data on the active pool snapshot just as you would any other pool, even while data is changing on the original pool you snapped.

Be careful when using this features as its implementation will cause your server performance to decrease. When pool snapshot is active, expect your disk writes to double or even triple.

Benefits of Pool Snapshots

Pool snapshots save time and preserve data. They provide an instant copy of a pool that can help expedite routine maintenance procedures to backup, archive, and protect data on that pool. Because traditional methods of duplicating large amounts of data can be expensive and time-consuming, the efficiency of snapshots can be an important benefit for your enterprise. You can make snaps as frequently as needed to meet your data availability and resilience requirements.

You can use pool snapshots in a variety of ways to enhance your current storage infrastructure, including the following scenarios:

Supporting Backup Operations

A pool snapshot facilitates non-disruptive backups because the snap becomes the source of the backup. As contrasted to a traditional, full-data copy of the pool, the metadata copy only takes a moment to create and occurs transparently to the end user. With traditional backups, applications may be shut down throughout the backup routine. In comparison, the pool snapshot process makes the original pool available with almost imperceptible delay.

Archiving Data

You can archive pool snapshots to capture a history of the changes made to the original data pool.

Restoring Data

Pool snapshots can serve as a source for restoring information. Two common reasons to restore information are user error and application errors.

- ◆ A user may inadvertently make changes to a file that need to be reversed. Files can become corrupted or deleted. The pool snapshot provides a quick and easy way to locate and reinstate selected files.
- ◆ An application may be infected by a virus or be corrupted by other problems, causing the application to store erroneous data throughout the pool. With a pool snapshot, you easily can easily effect a rollback of all or part of the original pool to a point in time before the virus or problem was known to exist in the system.

Recreating Operational and Development Environments

You can also write to the pool snapshot, just as you would any pool. You can work with and modify the snap version of the data. For example, in a software development environment, engineers may want to repeat builds and tests of data within a given snap.

Testing and Training

Snapshots can provide a convenient source for testing and training environments and for data mining purposes.

How the Pool Snapshot Works

The pool snapshot function uses copy-on-write technology to enable the instantaneous block-level snapshot of a pool, while requiring only a fraction of the storage space of the original data pool. A pool snapshot does not save an exact copy of the original data pool. Instead, the snap is a metadata-based copy that stores only the blocks of data that change subsequent to the instant of the snap. The snap combines the metadata and stored block data with the unchanged data on the original pool to provide a virtual image of an exact copy of the data at the instant the snapshot was taken, plus any end-user modifications made to that snap.

Before the snap can occur, the snapshot function must render the original pool quiescent by briefly halting all data transaction activity when current transactions complete. It temporarily prevents new writes to the pool and flushes the file system cache to make the pool current with existing writes. Then, it snaps the now stable pool, and allows data transaction activity to resume.

The quiesce process provides a transactionally consistent image at the instant the snap is made. Because the snap is consistent, it is not necessary to check the consistency of the file system or database when you activate the snap for access.

After the snap, the snapshot function continues to track the transaction activity in the original pool. It determines which blocks in the original pool will change as data writes are made to the original pool. It suspends the write activity while it copies the original block data to the designated pool where it stores the pool snapshot, referred to as the "stored-on pool." Then, it allows a write to that block in the original pool. This keeps the snapshot metadata consistent in time with the exact instant the snapshot was taken.

As the original pool changes, the snapshot can grow to the size of the stored-on pool. The average disk space requirements for a pool snapshot are 10 percent to 20 percent of the original pool size. The actual space depends on how many writes are made to the original volume and the time period that the snap exists before it is archived or deleted. Currently, a combination of up to 500 snapshots can exist on any given stored-on pool.

Creating Pool Snapshots

Create a pool snapshot when you want to capture a point-in-time view of an active data pool. Both the original pool and the stored-on pool must be active when you make the snap.

IMPORTANT: If you reach the limit of 500 existing snapshots in the stored-on pool, you cannot create a new pool until you delete a pool from the stored-on pool. See the Archiving and Deleting Pool Snapshots section for more information about archiving and deleting snaps.

For pool snapshot commands to create and manage pool snapshots, refer to [Table 1, “Snapshot Management Commands,” on page 44](#).

Naming Pool Snapshots

You name a pool snapshot at the time you order the snap. Specify a unique snap name for each snap. Because the snap name also serves as the snap's pool name when active, the name you give it should be unique among snapshots and among pools. The combination of snap name and time stamp when the snap was taken will help you identify the version of a data pool you want to access when a pool snapshot is active.

We recommend that you apply a naming convention to name snaps. This is particularly important if you plan to create a series of snapshots for multiple pools. The snap name is generally a modified version of the original pool's name to support easy identification of all snaps for any given pool. The snap name can be 2 to 16 characters.

For example, add a letter and number designator to the original pool name, such as "_Sn". The "S" indicates that it is a snap. The "n" represents an incremental number of snaps taken for this pool. In this scenario, snap names might be PoolA_S1, PoolA_S2, and so on for snaps of PoolA. If you delete older snaps and reuse the lower sequence numbers, a simple sort by snap name can be confusing. Be sure to verify the time stamp on the pool when you work with pools that conform to this naming convention.

For some scenarios it might be desirable to adopt more meaningful and sortable naming conventions. Currently, you can sort by snap name only, not by time stamp. If you create multiple snaps of a pool each day, consider using a logical naming convention that identifies both the pool name and a combination of the date and an approximate 24-hour-clock time of the snap. For example, use a modified version of the pool name and add "_Syymmddtttt". Snap names for snaps of PoolA might be PA_S0210300600, PA_S0210301800, PA_S0210310600, PA_S0210311800, and so on. Of course, the time stamp will show the exact time that the snap was taken, because there is a slight delay between ordering a snap and the instant of the snap itself.

It is also important to consider the names of existing pool snapshots and your snap-naming conventions when you name new data pools. If an active data pool and a deactive pool snapshot share the same name, a conflict will occur when you try to activate the pool snapshot. To resolve the conflict, the snap name of the pool snapshot will be modified dynamically with an "_n" (sequential number). This will make the snap name unique with respect to the active pool list during the time that the pool snapshot it is active and functioning as a pool.

NOTE: Currently, the name modification of the active pool snapshot is temporary for the duration of its active state. This will change for FCS. The intent will be to make this renaming permanent to avoid future conflicts

Renaming Pool Snapshots

There is no interface at this time for renaming a pool snapshot with a different snap name. Please do not use other pool utilities to rename a pool snapshot at this time.

Choosing the Stored-on Pool

The stored-on pool is the active data pool where you want to store the pool snapshot. All pool snapshots for any given original pool must reside on the same stored-on pool. As noted previously, a combination of up to 500 snapshots can exist on any given stored-on pool. When you create a pool's first snapshot, you will select the stored-on pool for the first and any subsequent snapshots. Typically, you can select the pool itself or a different pool to be the stored-on pool, given that the pool you choose is active and has sufficient space available. There are two exceptions:

- ◆ If the original pool is a shared pool, you must select the original pool as the stored-on pool, and this pool must be active and have sufficient space available.
- ◆ If the original pool is a non-shared pool, you cannot select a shared pool as the stored-on pool.

In general, you can achieve better performance by selecting a data pool located on a different disk than the pool you want to snap.

Maintaining the Stored-on Pool

If the stored-on pool and the original pool are the same, the status of any given pool snapshot is closely tied to the operational status of the pool. However, if they are not the same, you need to consider how the status of each pool affects the other, and how they affect the status of the pool snapshot.

You can deactivate the original pool, as needed, without adversely impacting the pool snapshot or the status of the stored-on pool. If the original pool is deactive, there will be no active transactions for the pool snapshot function to process. If the stored-on pool hosts only the snapshots made for that deactivated original pool, it can be safely deactivated for the duration. Re-activate the stored-on pool first when bringing the pools back into service.

IMPORTANT: If you need to deactivate a stored-on pool, you must first deactivate each of the original pools that correspond to the families of pool snapshots stored on it.

In contrast, deactivating the stored-on pool first can cause the ungraceful deactivation of the corresponding original pool. Each pool snapshot must be able to dynamically expand its storage footprint in the stored-on pool, because its metadata grows in response to changes in the original pool. Dynamic allocation is not possible when its stored-on pool is deactive. For brief outages, it is possible that any given pool snapshot residing in that stored-on pool will have sufficient space available to accept new metadata, depending on the frequency of writes to its corresponding original pool. However, if you deactivate a stored-on pool, it increases the likelihood that one of the pool snapshots it hosts will run out of space. This will cause the ungraceful deactivation of the corresponding original pool.

IMPORTANT: The stored-on pool should remain active as long as it hosts any pool snapshots. You can deactivate it safely only after all original pools are deactivated, and for the duration of their deactivation. Activate the stored-on pool before re-activating any of the original pools.

Activating Pool Snapshots

You activate a pool snapshot whenever you want to access the data on it, such as for data retrieval, data modification, and data backup. Once the pool snapshot is active, it appears by its user-assigned snap name in the pool list. Treat it as you would any pool to manage the pool and to activate and mount its volumes.

The names of volumes on the pool snapshot are a modified version of the volumes on the original pool. Generally, an "_SV" (snapshot volume) is added to the volumes' names. When you deactivate

the pool snapshot, any snapshot volumes on it are automatically deactivated, and its snap name will no longer be listed in the pool list.

Archiving and Deleting Pool Snapshots

You delete a pool snapshot when you no longer need it, or if you want to make room to add other snapshots to its stored-on pool.

Remember that end users can access the snap for reads and writes, just as they would any pool. Because any given pool snapshot may rotate into an archived status or be deleted, you need to be aware of how end users are accessing and modifying data within that snap.

Using Pool Snapshot Commands

The following table shows the commands you can use to manage pool snapshots.

Table 1 Snapshot Management Commands

Commands	Description
<code>mm snap create <i>snappool datapool snapname</i></code>	Creates a pool snapshot.
<code>mm snap activate <i>snapname</i></code>	Activates a pool snapshot.
<code>mm snap deactivate <i>snapname</i></code>	Deactivates a pool snapshot.
<code>mm snap delete <i>snapname</i></code>	Deletes a pool snapshot.
Rename	Currently under development. This command is not functioning at this time. Please do not use other pool utilities to rename a pool snapshot at this time.
<code>mm snap list</code>	Lists the snapshots on the server.

The following table shows the attribute descriptions for pool snapshot commands listed in the table above.

Table 2 Attributes for Pool Snapshot Commands

Attributes	Description
snappool	The unique name of the original data pool that you are snapping; must be 2 to 16 characters.
datapool	The unique name of the stored-on pool for the pool snapshot; must be 2 to 16 characters.

Attributes	Description
snapname	The unique snap name of the pool snapshot, which will be used as its pool name when the pool snapshot is activated; must be 2 to 16 characters.

7


Managing RAID Devices

NSS lets you create a virtual storage device by striping data across multiple physical disk drives on your system. This option, called RAID, is a software configuration that emulates an actual hardware RAID 0 system; however, it uses standard host adapters and does not require any special RAID hardware. You can create partitions, pools, and volumes on a RAID device, just as you would with any physical storage device. For basic information on RAID types, see [“Understanding the Differences between RAID 0, 1 and 5” on page 48](#).

You can create a RAID 0, 1, or 5 device. When you set up a RAID device, you obtain partitioned free space from all your physical storage devices. Then, using this free space, you will create virtual partitions (also commonly referred to as pool segments). From these pool segments, you can create volumes and pools. These pool segments are the basic elements of a software RAID device. Once the segments are defined, data is then striped across the segments in blocks as large as the stripe size that you specify. Due to the striping, your input and output performance will increase.

The RAID 0 device illustrated below consists of three RAID partitions, or pool segments, across three physical drives. The stripes are written and read in the order designated.

Figure 2 Striping data on a software RAID 0 device



Important Information about RAID

- ◆ Each partition in the RAID configuration should come from a different device. NSS lets you obtain RAID partitions from the same device, but this severely impedes the performance of your file system. Similarly, creating more than one RAID device on a single physical device will severely decrease the performance of your file system.
- ◆ RAID can accommodate a maximum of eight segments (amount of space used from each device) on a virtual RAID device.
- ◆ The stripe sizes range from 4, 8, 16, 32, 64, 128, and 256 KB.
- ◆ You can use any combination of IDE or SCSI devices in a RAID device; however, make sure these devices that have similar performance drives. Otherwise, your performance might decrease.

Understanding the Differences between RAID 0, 1 and 5

You can create a RAID 0, 1, or 5 device. The type of device you choose depends on whether you value file system performance, fault tolerance, or both. The following tables lists the different fault tolerance and performance issues associate with each RAID type:

RAID Type	Description
RAID 0	<p>Provides the best performance but offers no fault tolerance.</p> <p>The performance increase in RAID 0 is due to the data striping which allows the data to be interleaved across various physical drives.</p>
RAID 1 (mirroring)	<p>Provides fault tolerance through data redundancy with the possibility of a slight performance increase when doing file system reads.</p>
RAID 5	<p>Provides both a significant performance increase as well as fault tolerance by using a parity segment (for more information, see “Software Raid 5” on page 9).</p> <p>The performance increase in RAID 5 is due to the data striping which allows the data to be interleaved across various physical drives.</p> <p>The fault tolerance is due to the fact that if you lose a device, the volumes and pools associated with that device remain active.</p>

RAID 1 (Mirroring) Requirements

The following is a list of requirements for mirroring partitions (sometimes referred to as pool segments) with RAID 1:

- ◆ Mirrored partitions must have the same partition type—NSS partitions to NSS partitions and traditional partitions to traditional partitions.
- ◆ Mirrored partitions should be set up on devices that have similar performance thresholds.
- ◆ You can mirror only partitions. If a storage pool spans multiple partitions, or pool segments, all the individual segments that make up that pool must be mirrored in order for the data in that pool to be fault tolerant.
- ◆ You cannot combine mirror groups (existing groups with multiple mirrored partitions). A mirror group is expanded by adding a segment from your free space but not by adding an existing mirror group to the current group.
- ◆ All mirrored partitions on a clustered system must be marked a shared or not shared for each mirror group.
- ◆ Avoid setting up multiple mirror groups on a single device. Such configuration heavily degrades the performance of the file system.

Creating a RAID Device

To create and manage RAID devices with the console-based NSS Management Utility (NSSMU):

- 1** (Conditional) If you are creating a RAID 1 device, review the “**RAID 1 (Mirroring) Requirements**” on page 48 before continuing.
- 2** Load `nssmu.nlm` by entering `nssmu` at the server console.
- 3** Do one of the following:
 - ♦ (RAID 1 only) To create a RAID 1 device on an existing pool or volume, select Partitions from the NSS main menu, select the existing partition you want to create the RAID 1 device on, press F3 to select the partitioned free space you want to include in your RAID 1 device, and then press Enter to create the RAID device (the partition is automatically mirrored when added to the RAID device).
 - ♦ To create a new RAID device from partitioned free space or to expand RAID devices, select RAID Devices from the `nssmu` main menu.
- 4** Press Insert (Ins) to create a new device.
- 5** Select the RAID type (RAID 0, 1, or 5) and then press Enter.
- 6** Enter the stripe size and then press Enter.

64 KB is the default size and will typically provide the best performance if you are using NSS volumes. The stripe size will vary according to the size of the requests that are being sent to the disk. If are sending 64 KB requests down to the disk, you would want to increase the strip size to 128 KB or larger to improve performance.
- 7** Select the partitions that you want to mirror or add by using the arrow keys and then pressing Enter.

IMPORTANT: If no partitions appear it is an indication that either there are no partitions large enough or there are no free partitions.

When you press Enter, the partition will appear in the Partition Segments Included in RAID window. The following table lists the number of segments that you can include in a RAID device:

RAID Type	Number of segments
RAID 0	Maximum of 8
RAID 1	Maximum of 4
RAID 5	Minimum of 3 (one of which is used as the parity) and a maximum of 8

IMPORTANT: Once a partition is selected from a device, other free partitions associated with that device might not appear. This prevents you from adding more than one partition from a single physical device, which will severely degrade the performance of your file system.

- 8** Press F3 to initialize and create the RAID device.

Once the RAID device is created, the device ID will appear in the RAID Devices window. This window is viewed from the RAID Devices NSSMU main menu.

Viewing, Expanding, Restriping, and Deleting a RAID Device

To manage RAID devices in `nssmu.nlm`, load the utility by entering `nssmu` at the server console and then selecting RAID Devices.

The following table lists the keyboard strokes that enable you to view, expand, restripe, and delete a RAID device:

Keyboard Stroke	Description
F7	<p>Shows the RAID segments associate with the highlighted RAID device.</p> <p>You can also delete a segment from this screen.</p>
F5	<p>Expands a RAID device by adding more segments from another storage device.</p> <p>You should not place more than one RAID segment on a disk; this will severely impede the performance of your file system.</p>
F6	<p>Restripes or remirrors a RAID device.</p> <p>Make sure all partitions have been added to your RAID device before restriping a device.</p> <p>NOTE: The file system performance might slow down during the restriping/remirroring process.</p>
Del	<p>Deletes an entire RAID device.</p> <p>When you delete a RAID device, all partitions, volumes, and pools associated with that device are also deleted.</p>
F5	<p>Refreshes the data with the changes that you made.</p>

8

Managing NSS Pools and Volumes

- ◆ “Renaming a Storage Pool” on page 51
- ◆ “Increasing the Size of a Storage Pool” on page 51
- ◆ “Deactivating/Activating Storage Pools and Volumes” on page 52
- ◆ “Deleting Storage Pools or Volumes” on page 52
- ◆ “Rebuilding Storage Pools and Volumes” on page 52
- ◆ “Renaming a Volume” on page 53
- ◆ “Activating a Volume” on page 53
- ◆ “Reviewing and Restoring Deleted Volumes” on page 54
- ◆ “Salvaging and Purging Deleted Files” on page 54
- ◆ “Assigning Rights with Trustees.xml” on page 55
- ◆ “Restricting Access to Files.cmd” on page 56

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.

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 and 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.

Deactivating/Activating Storage 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.

To deactivate/activate storage pools or volumes:

- 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.

Deleting Storage Pools or 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 Volumes” on page 54.](#))

IMPORTANT: If you delete a storage pool, you cannot restore any of the volumes in that pool.

To delete a storage pool or volume:

- 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.

Rebuilding Storage Pools and Volumes

NSS allows you to temporarily deactivate individual storage pools to fix volume problems instead of bringing down the server. However, when you deactivate a storage pool, users will not have access to any of the volumes in that pool.

Most volume errors, which are typically transactions left unfinished during a system crash of some kind, get fixed automatically during volume mount. If the errors persist, NSS uses the Rebuild and Verify utilities to resolve any problems.

Verify fixes NSS pools and volumes by searching for inconsistent data blocks or other errors. This utility indicates if there are problems with the file system. If you still have problems, use the Rebuild utility.

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 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).

IMPORTANT: The Verify and Rebuild utilities are not the same as VREPAIR (used to rebuild traditional volumes). You should use Rebuild only as a last resort to recover the file system after data corruption.

- 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.

- 2 Next, mount the volume.

Error and Log Files

When you use Rebuild and Verify a log file is generated at the root of the DOS drive. The error log, `sys:pool_name.vlf`, contains information about data that has been lost or recovered. Data loss, unfortunately, is a possibility during a rebuild.

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.

You can also verify a pool by entering `nss /verify` at the server console and then selecting the pool name from the displayed list. Verify is a read-only assessment of the pool.

Renaming a Volume

You can rename NSS 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.

Activating a Volume

After you set up and configure NSS 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 volume you want to mount and click Mount.

Reviewing and Restoring Deleted Volumes

If you delete a storage pool, you delete all the volumes in that pool; these volumes cannot be restored. However, if you delete just a volume, the file system removes it from the storage pool and for a specified amount of time, called the Purge Delay time, you can review and even restore the contents of that volume. After the Purge Delay time elapses, the volume is removed from the system and you can no longer access it.

During the Purge Delay time (four days after a volume is deleted), you can manually purge deleted volumes, view the volume contents, transfer files from the deleted volume to other volumes, or salvage the entire volume if you wish.

You can also change the Purge Delay time to extend or reduce the time for the automatic purging cycle. The Purge Delay time command can be changed by adding a line in the autoexec.ncf file. Default time, 345,600 seconds, is four days. The changed time must be listed in seconds. See [“Changing the Purge Delay Setting” on page 26](#).

- 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.

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 placed 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.

Assigning Rights with Trustees.xml

Trustees.xml is a file that lets you assign trustee rights to the root of a volume. Trustees.xml is located in sys:etc and is executed every time the server boots.

This XML file is useful for assigning rights to non-persistent files or directories, such as the Admin volume or a CD that you have activated as a volume.

The follow is a sample of the trustees.xml code:

```
<specialTrustees>
  <addTrustee>
    <name>admin.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>rwfe</rights>
  </addTrustee>
  <addTrustee>
    <name>xxxx.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>rwf</rights>
  </addTrustee>
  <addTrustee>
    <name>fred.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>rwf</rights>
  </addTrustee>
  <addTrustee>
    <name>[public]</name>
    <fileName>_admin:manage_nss\files.cmd</fileName>
    <rights>rwf</rights>
    <background/>
  </addTrustee>
  <addTrustee>
    <name>fred.novell.kona_tree</name>
    <fileName>sys:\etc\hosts</fileName>
    <rights>rwf</rights>
  </addTrustee>
</specialTrustees>
```

Refer to the following table for information about trustees.xml tags:

Table 3 Explanation of Trustees.xml Tags

XML Tags	Explanation
<code><name></code>	<p>Enter the full context and username of the user. You can enter the information in back slash or dot format.</p> <p>NOTE: This tag cannot contain multiple usernames.</p> <p>For example, if John Smith existed in the novell context in the kona_tree, you would enter one of the following:</p> <p><code>jsmith.novell.kona_tree</code></p> <p>or</p> <p><code>\kona_tree\novell\jsmith</code></p>
<code><filename></code>	<p>Enter the name of volume, directory, and file that you are assigning rights to.</p> <p>For example, if you are assigning rights to the hosts file in the etc directory located on the sys: volume, you would enter the following into the brackets</p> <p><code><filename>sys:\etc\hosts</filename></code></p>
<code><rights></code>	<p>The following rights can be assigned in this file:</p> <ul style="list-style-type: none"> ♦ A (access control) ♦ C (create) ♦ E (edit) ♦ F (file scan) ♦ M (modify metadata) ♦ R (read) ♦ S (supervisor) ♦ W (write)

Restricting Access to Files.cmd

If you want to restrict access to files.cmd, you must do so by modifying the `<addTrustee>` tag in a configuration file called trustees.xml.

Look for this code sample in the .xml file and modify [public] to reflect the user that needs restricted access to files.cmd:

```
<addTrustee>
  <name>[public]</name>
  <fileName>_admin:manage_nss\files.cmd</fileName>
  <rights>rwf</rights>
  <background/>
</addTrustee>
```


For example, if you want John Smith, a contract employee, to have only file scan rights to this file, you would enter the following:

```
<addTrustee>
  <name>jsmith.novell.kona_tree</name>
  <fileName>_admin:manage_nss\files.cmd</fileName>
  <rights>f</rights>
  <background/>
</addTrustee>
```

For further explanation of the xml tags, refer to [Table 3, “Explanation of Trustees.xml Tags,”](#) on [page 56](#).

9

Using NSS Server Console Commands

Use the following commands at the NetWare[®] server console to view NSS information and perform tasks.

- ◆ “Online Help Commands” on page 59
- ◆ “NSS Configuration Information Commands” on page 59
- ◆ “Start Up NSS Commands” on page 59
- ◆ “Other NSS Commands” on page 61
- ◆ “System Volume Recovery Commands” on page 64
- ◆ “DOS FAT Commands” on page 65
- ◆ “Pool Multiple Server Activation Prevention Commands” on page 65

NOTE: You can abbreviate NSS command options by providing just enough of the command to be unique. For example, NSS /CA would be unique enough to be an abbreviation for the NSS /CacheStats command.

Online Help Commands

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

NSS Configuration Information Commands

Commands	Description
<code>nss /modules</code>	Lists the providers, loadable storage subsystems, and semantic agents.
<code>nss /status</code>	Lists the current NSS status.
<code>volumes</code>	Lists all NetWare and NSS volumes that are mounted, including Admin volume.
<code>nss volumes</code>	Lists all NSS volumes, including the Admin volume.

Start Up NSS Commands

The following table lists commands that are only valid when NSS is loading:

Commands	Description
<code>nss /allLSS</code>	Loads all the LSS modules.
<code>nss /CDDVD</code>	Loads the <code>cddvd.nss</code> module that is essential for CD, DVD, and Macintosh support. For more information, see “Cddvd” on page 11 and “Activating CD-ROM, DVD, and DOS Partitions as NSS Volumes” on page 29 .
<code>nss /defaultNLMS</code>	Loads only default NSS NLMS and any additional LSS modules specified on the command line.
<code>nss /DOSFAT</code>	Loads only those modules that are essential for accessing FAT 16 and FAT 12 drives.
<code>nss /MailBoxSize=value</code>	Sets the size of your mailbox. The default is 228. The range is 64 to 256.
<code>nss /noLSS</code>	If specified, does not load any LSS modules.
<code>nss / (No)skipLoadModules=value</code>	Skips the auto loading of the NSS support modules. StartupOnly Value=Off.
<code>nss /NumAsyncios=value</code>	Sets the number of asynchronous I/O entries to allocate. The default is 2048. The range is 4 - 65536.
<code>nss /Numbonds=value</code>	Sets the number of bond entries to allocate. The default is 5000. The Range is 512 - 2097152.
<code>nss /NumWorkToDo=value</code>	Sets the number of WorkToDo entries (entries can 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.
<code>nss /zLSS</code>	Loads only those modules that are essential for zLSS support.

To Permanently Modify 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.

To Temporarily Modify NSS Startup Commands while NetWare 6 Is Booting

When loading the server.exe, you can modify NSS parameters as NetWare loads. However, this modification is only temporary and must be done each time you boot your server.

At the server command line, enter

```
Server -z"/XXXXX"
```

In the above example the -z tells the server to look for specific NSS startup commands and the information in the quotes is the actual command.

Other NSS Commands

NOTE: This is not a complete list of NSS commands. To see a complete listing, type `nss /help` at the server console. Also, you do not need to type the full command at the server console. You need only type enough of the name so that it is unique.

Commands	Description
<code>nss /Activate=volume_name</code>	Activates an NSS volume.
<code>nss /AllocAheadBlks=value</code>	Sets the number of blocks to allocate ahead on writes. The default is 15. The range is 0 - 63.
<code>nss /backgroundchecking</code>	Turns the background checker on.
<code>nss /nobackgroundchecking</code>	Turns the background checker off.
<code>nss /NoBGCompression</code>	Starts or stops background compression. Stop BGCompression will stop all the compression requests in the queue.
<code>nss / NoClassicDirectoryQuotas</code>	Turns emulation of traditional NetWare volumes on or off. The default is off.
<code>nss / ClosedFileHashShift=valu e</code>	Sets the number of closed files that can be cached in memory. The default is 512. The range is 1 to 100000.
<code>nss /Compression</code>	Enables file compression on the volume.
<code>nss / NoCompression=volume_nam e/all</code>	Disables file compression on the volume.
<code>nss /CompScreen</code>	Starts the NSS compression statistics screen.
<code>nss /DataShredding</code>	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.
<code>nss /NoDataShredding</code>	Disables data shredding on the volume.

Commands	Description
<code>nss / Deactivate=volume_name</code>	Deactivates an NSS volume.
<code>nss /DirectoryQuotas</code>	Enables directory quotas on the volume.
<code>nss /NoDirectoryQuotas</code>	Disables directory quotas on the volume.
<code>nss /FileCopyOnWrite</code>	Creates a copy of files in the volume when they are open for write.
<code>nss /NoFileCopyOnWrite</code>	Will not create a copy of files on the volume when they are open for write.
<code>nss /FileFlushTimer=value</code>	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.
<code>nss /FixMFL</code>	MFL maintenance for the volumes.
<code>nss / FlushFilesImmediately</code>	Synchronously flush files in the volume when they are closed.
<code>nss / NoFlushFilesImmediately</code>	Disables synchronously flushing files in the volume.
<code>nss / ForceActivate=volume_name</code>	Forces an NSS volume to become active.
<code>nss /GetMFLStatus</code>	Gets the status of MFL maintenance for the volumes.
<code>nss / LogicalVolumePurgeDelay=value</code>	The number of seconds before deleted logical volumes are purged. This allows time to reverse the deletion.
<code>nss / LogicalVolumePurgeDelayAfterLoad</code>	The number of seconds, after NSS loads, before deleted logical volumes are purged. This allows time to reverse the deletion.
<code>nss / LogicalVolumePurgeDelayAfterContinue=value</code>	The number of seconds to delay purging a logical volume after clicking continue. After a volume starts to purge, it cannot be salvaged.
<code>nss /LVDeleteStatusBasic</code>	Displays information about deleted logical volumes.
<code>nss / LVDeleteStatusSalvagable</code>	Displays information about salvagable logical volumes.
<code>nss / Maintenance=volume_name</code>	Switches the specified NSS volume to maintenance mode.
<code>nss /MFL=volume</code>	Enables MFL maintenance for the volumes.

Commands	Description
<code>nss /NoMFL=volume</code>	Disables MFL maintenance for the volumes.
<code>nss /MFLVerify=volume</code>	Gets the status of MFL maintenance for the volumes.
<code>nss /Migration=volume name/all</code>	Enables migration of files on the volumes.
<code>nss /NoMigration=volume name/all</code>	Disables migration of files on the volumes.
<code>nss / OpenFileHashShift=value</code>	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.
<code>nss /Pools</code>	Shows all of the currently available NSS pools.
<code>nss /PoolAutoActivate</code>	Activates specified pools at load time.
<code>nss / PoolAutoDeactivate=pool</code>	Leaves specified pools deactivated at pool load time.
<code>nss /PoolAutoDisplay</code>	Displays current pool load time policies.
<code>nss / PoolAutoMaintenance=pool</code>	Places specified pools in maintenance mode at pool load time.
<code>nss /PoolAutoVerify</code>	Verifies the specified pool's physical integrity at startup time.
<code>nss /PoolRebuild=pool</code>	Rebuilds specified pools.
<code>nss /PoolRebuild</code>	Selects pools from a menu and rebuild them.
<code>nss /PoolVerify=pool</code>	Verifies the specified pool's physical integrity.
<code>nss /PoolVerify</code>	Selects pools from a menu and verify their physical integrity.
<code>nss /Salvage</code>	Enables salvage of deleted files on volumes.
<code>nss /NoSalvage</code>	Disables salvage of deleted files on volumes.
<code>nss /space</code>	Shows the amount of space on pools and their associated volumes.
<code>nss / StorageAlarmThreshold=va lue</code>	Lets you set the threshold for a low storage space warning. The default is 10. The range is 0 to 1000000.
<code>nss / (No)StorageAlertMessages</code>	Turns ON or OFF the low storage message to users. The default is ON.

Commands	Description
<code>nss / StorageResetThreshold=volume</code>	Lets you reset the threshold for a low storage space warning. The default is 10. The range is 1 to 1000000.
<code>nss /NoTransaction=volume name/all</code>	Disables transactional writes in files on the volumes.
<code>nss / UpgradeObjectsOnVolume</code>	Upgrades objects on volumes to the current version.
<code>nss / UserSpaceRestrictions</code>	Enables user space restrictions on the volume.
<code>nss / NoUserSpaceRestrictions</code>	Disables user space restrictions on the volume.
<code>nss /Transaction=volume name/all</code>	Enables transactional writes in files on the volumes.
<code>nss /VolumeAutoActivate</code>	Activates volume at load time.
<code>nss /VolumeAutoDeactivate</code>	Deactivates volume at load time.
<code>nss /VolumeAutoDisplay</code>	Displays current volume load time policies.
<code>nss /VolumeMaintenance</code>	Places a volume into maintenance mode. Volumes can be put in maintenance mode, but maintenance only occurs at the storage pool level.
<code>nss /ZLSSIOStatus</code>	Displays current NSS IO status information.
<code>nss /ZLSSPendingWriteIOs</code>	Number of write IOs dropped to the storage subsystem. (0 = all available) [Value=1000 Range=0 - 3000].
<code>nss /ZLSSPoolScan</code>	Scans for and loads all ZLSS pools.
<code>nss /ZLSSVolumeUpgrade</code>	Allows the 5.x/ZLSS VolumeUpgrade to the NetWare 6 format to occur during the installation process.

System Volume Recovery Commands

Commands	Description
<code>nss /ChangeSYSQuota</code>	Lets you change the quota (in MB) for the sys: volume. Setting this value to zero will set the Quota to none and will allow the sys: volume to grow to the size of the pool.

Commands	Description
<code>nss /ExpandSYS</code>	Lets you expand the pool that contains your sys: volume. When prompted, enter the partition ID of the free space you plan to use to expand the pool. You can obtain the partition ID of all free partitions by using the /ListFreeSpace command.
<code>nss /ListFreeSpace</code>	Lists all free partitions and NSS partitions that have not been assigned to a pool. Use this command to obtain the partition ID required when using the /ExpandSYS command.
<code>nss /RenameToSYS</code>	Lets you rename the sys: volume back to sys: if you have previously renamed it. At the prompt, enter the name of the volume you want to rename to sys:.
<code>nss /SalvageSYS</code>	Lets you restore the sys: volume if you have deleted it (only if it has not yet been purged).
<code>nss /SYSHotFixSize</code>	Hot Fix is no longer supported because modern storage devices do the dynamic redirection in the hardware. Lets you change the hot fix size. This option is used in conjunction with /ExpandSYS. If you are expanding sys with a free partition, this option lets you change the default hot fix size. The default is 100KB.

DOS FAT Commands

Commands	Description
<code>nss / FATPartition=<i>partition_t</i> <i>type_number</i></code>	Supports up to three additional partition types containing 16-bit FATs, such as / FATPartition=12,13.
<code>nss /(No)FATLongNames</code>	Enables long filenames on FAT volumes. The default is ON.

Pool Multiple Server Activation Prevention Commands

The management file for pool MSAP is `_admin\manage_nss\pool\poolname\zls\msap.xml`. One file exists for each pool. This file contains MSAP statistics for the pool. The MSAP attribute is displayed in the `enabledAttributes` field of the `poolinfo.xml` management file.

For `manage.cmd`, the pool operation `getPoolInfo` returns the MSAP tag (`<msap>`) in the `supportedAttributes` tag (`<supportedAttributes>`) and the `enabledAttributes` tag (`<enabledAttributes>`).

For APIs, the pool feature `zpool_feature_msap` can be viewed and controlled using the `zGetInfo` and `zModifyInfo` commands.

Table 4 MSAP Console Commands

MSAP Console Command	Description
nss /msapserver	Enables MSAP for all the pools on the server. By default, MSAP is enabled for every pool on the server. We recommend that you never disable MSAP.
nss /nomsapserver	Disables MSAP for all the pools on the server. This command remains in effect only until the server is next rebooted. We recommend that you never disable MSAP.
nss /msaprebuild	Rebuilds a corrupt MSAP block. Before issuing the command to rebuild, you must deactivate the pool because a corrupt MSAP prevents a pool from going into Maintenance state. Rebuilding an MSAP block does not give the builder ownership of the pool.
nss /poolmsap=poolname	Enables MSAP for a given pool on the server. MSAP will be enabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now enabled.)
nss /nopoolmsap=poolname	Disables MSAP for a given pool. Use the command when the pool is activated. MSAP will be disabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now disabled.)
nss /pools	Displays the message “Multi-use detect” for pools that have MSAP enabled.

10 Fine Tuning NSS Performance

Novell Storage Services (NSS) is the default file system that ships with NetWare 6. Most customers who are running NetWare 6 have automatically chosen to run NSS; however, little information is available on how NSS works or how to tune it.

Under certain conditions, the performance of NSS can be tuned to great benefit without sacrificing NetWare's legendary performance. NSS tuning can be *proactive*, to get the highest possible speed from your storage system, or *reactive*, because there is a problem.

This section focuses on different proactive strategies that can improve NSS performance. In addition to this document, more in depth information on NSS tuning can be found in the July 2002 Appnote, *Novell Storage Services (NSS) Performance Monitoring and Tuning* (<http://developer.novell.com/research/appnotes/2002/july/01/a020701.htm>).

Do I Need to Tune NSS?

There are many factors that contribute to decreasing server performance; however, if you are noticing that your server is performing poorly and you suspect the storage subsystem (NSS), you can monitor the storage subsystem by using specific NSS command line options. These options will help you determine if any tuning is required.

For further details and explanation, see the Command Line Options for Monitoring NSS in the July 2002 Appnote, *Novell Storage Services (NSS) Performance Monitoring and Tuning* (<http://developer.novell.com/research/appnotes/2002/july/01/a020701.htm>).

Adjusting NSS Parameters

Once you have determined that your storage subsystem needs tuning, your next step is to determine what NSS parameters can be adjusted to improve your current NSS performance.

Before beginning, review the following section, [Important NSS Tuning Guidelines](#).

Important NSS Tuning Guidelines

- ♦ Make sure that you are running the latest version of NSS before adjusting parameters.
- ♦ Don't split your resources between two file systems.

When FILESYS loads, RAM that is not allocated elsewhere is given to the legacy cache. When NSS loads, it will also take (by default) 60% of the RAM that is not allocated elsewhere for NSS. As the two file systems are disparate, this will halve the performance of both caches. When applications load the amount of cache for both file systems, performance is reduced even further.

If you *must* run multiple file systems, try to get to know the usage patterns of each system. This may enable you to balance RAM between them more effectively.

- ◆ Purchase more RAM.

You might find that the performance tuning requires you to purchase more RAM for your server.

IMPORTANT: Do not add RAM over 4GB to improve cache performance. The Intel architecture has no DMA capabilities above 4GB; therefore, NSS deliberately does not use any memory above this amount.

Start with These Parameters

Tuning NSS is a matter of modifying a parameter or two, resetting the cache performance statistics, and then waiting and watching for performance gains or degradation.

WARNING: Before you begin to adjust settings, you need to be aware that if you give NSS more memory than the server has, NSS will fail to load. If your SYS volume is NSS-based, your server will be inoperable until you load the server with a NSS override to correct the situation. If you need to restore a setting after a failed tuning attempt, enter `server -z parameter` to restore the setting.

Adjusting the following parameters can improve your NSS performance:

- ◆ [“Setting the Cache Balance” on page 68](#)
- ◆ [“Setting the Cache Buffers” on page 71](#)
- ◆ [“Setting the Closed File Cache Size” on page 73](#)
- ◆ [“Setting the File/Buffer Flush Timer” on page 74](#)
- ◆ [“Setting the Name Cache” on page 74](#)

Setting the Cache Balance

The cache performance statistics are divided into SYS (system) and USER cache. System data is everything that isn't user data. It includes the roll-back ZLOG (journal) and the roll-forward Purge Log, and the various databases and structures where NSS keeps information about files—for example, the name tree, the beast tree, and the visibility list.

If user data from cached files leaves no space for the system data, the file system performance might suffer. Under normal circumstances, you should see a 98% hit rate for both user data cache and system data cache. In some cases a lower cache rate is normal—for example, in situations where backups have run, or when large file copies have just occurred, or in a business that deals with large files that consistently flush the cache. If you want to raise the cache statistics in these situations, you will need to dedicate more RAM to NSS.

The following table lists a few NSS parameters that can be used to adjust the cache balance between user and system data:

Task	Purpose	Command
Control the division of file system cache between User and System data	<p>If you have a server with a few files and volume attributes and a generous amount of cache RAM, you can increase this parameter.</p> <p>The default percentage for user data cache is 70%. And remember, that whatever is not system cache is user cache.</p>	<p><code>nss / CacheUserMaxPercent=xx</code></p>
Set the amount of cache NSS takes from the file system	<p>Every 30 seconds (by default) or when the server loads an NLM, NSS rebalances the NSS file system cache to the percentage specified by the <code>nss /CacheBalance</code> parameter.</p> <p>NOTE: You can change the default time interval with the <code>nss / CacheBalanceTimer</code> command. The range for this is 1 - 3,600 seconds.</p> <p>A low cache balance will impede the performance of NSS.</p> <p>We recommend that you set the cache balance parameter to equal the percentage of the total disk space you allocate for NSS. However, <i>do not</i> 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.</p>	<p><code>nss /CacheBalance=xx</code></p>

Task	Purpose	Command
Set the number of cache buffers allocated for NSS	<p>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.</p> <p>On a server with only NSS volumes, you can adjust the /CacheBalance parameter to 85% without any problem. However, you should monitor the Least Recently Used (LRU) Sitting Time on the NetWare server to ensure that you are not running the server too low on memory.</p> <p>To safeguard against this, use the nss /MinOSBufferCacheSize command to set a hard limit for the amount of RAM that can be used exclusively by the operating system.</p> <p>HINT: LRU Sitting Time can be found in the Monitor utility under Disk Cache Utilization. Because, under normal circumstances, the server should be up for over 24 hours, this statistic should never drop below 24 hours. It is important to understand that LRU Sitting Time is a measure of non-cache memory, not file system cache memory.</p>	<pre>nss /MinBufferCacheSize= xx</pre>
Set a minimum amount of RAM exclusively for the operating system	<p>If you allocate too much memory to your NSS cache, you might run the server out of memory. To prevent this, set the a minimum amount of RAM aside for the exclusive use of the server operation system.</p> <p>The default is 256. The range is 256 to 1048576.</p>	<pre>nss / MinOSBufferCacheSize=xx</pre>

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% is good. This means the system is spending less time going to the disk to retrieve data.</p> <p>Cache hit is the number of times that NSS finds the desired file block in the cache buffers. Lower numbers indicate that NSS needs more memory.</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>	<code>nss/CacheHit</code>
Determine how often NSS rebalances the amount of cache it uses	<p>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. I</p> <p>f 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.</p>	<code>nss/CacheBalanceTimer</code>
Reset cache statistics	<p>This command lets you clear the current statistics and start over.</p> <p>This is useful for checking the results of changes to the system.</p>	<code>nss /reset</code>

Setting the Cache Buffers

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

The following are guidelines 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 the number of cache buffers allocated for NSS	<p>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.</p> <p>On a server with only NSS volumes, you can adjust the /CacheBalance parameter to 85% without any problem. However, you should monitor the Least Recently Used (LRU) Sitting Time on the NetWare server to ensure that you are not running the server too low on memory.</p> <p>To safeguard against this, use the nss /MinOSBufferCacheSize command to set a hard limit for the amount of RAM that can be used exclusively by the operating system.</p> <p>HINT: LRU Sitting Time can be found in the Monitor utility under Disk Cache Utilization. Because, under normal circumstances, the server should be up for over 24 hours, this statistic should never drop below 24 hours. It is important to understand that LRU Sitting Time is a measure of non-cache memory, not file system cache memory.</p>	<code>nss /MinBufferCacheSize= xx</code>
Limit the number of cache buffers	<p>When the cache balance is recalculated, a small amount of RAM is either given or taken away from the NSS Cache.</p> <p>This parameter controls the maximum cache balance transfer (in 4KB blocks) that is given to the NSS file system cache per cache balance. The default is 1024 (4MB) per session. The range is 16 - 1048576.</p> <p>If your server has applications running on it that can consume significant amounts of RAM for a short period of time, you may want to increase this parameter.</p>	<code>nss/ CacheBalanceMaxBufferPe rSession=xx</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 the cache for file and directory entries	This sets the number of name cache entries. Because NSS keeps a cache of the file and directory names it has recently accessed, setting the number of name cache entries speeds up opening files and path searches. The default is 2111. The range is 3 to 65521.	nss/NameCacheSize
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
Set the number of authorization cache entries	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.	nss/AuthCacheSize
Show cache buffer statistics	This shows the cache buffer statistics.	nss/CacheStats
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.	nss/reset

Setting the Closed File Cache Size

This parameter keeps the beast objects in cache so that NSS doesn't have to go to disk and unpack beast information again when it wants the same files:

nss /ClosedFileCacheSize=xx

The default is 50,000. The range is 16 - 1,000,000. On average, each Closed File Cache entry consumes 0.4 - 1KB of RAM.

Set this parameter to 100,000 or more if you have applications on your server that consistently cycle through the same set of files and you suspect that this cache is being flushed when combined with normal server operations.

HINT: The Closed File Cache will grow up to the specified amount. Unlike file system cache, Closed File Cache does not take the maximum amount of memory it needs from the start. It also adapts—as memory is consumed by other processes, it dynamically reduces the number of entries so the system does not become starved for memory.

Setting the File/Buffer Flush Timer

Task	Purpose	Command
Set the amount of time before files are flushed.	This sets the maximum amount of time that modified files are kept in a cache buffer before flushing the files to disk. The default is 10 seconds. The range is 1 to 3600 seconds.	<code>nss/FileFlushTimer</code>
Set the amount of time before buffers are written.	This sets the maximum amount of time that modified NSS file system cache buffers are kept in memory before they are written to disk. The default is 1 second. The range is 1 to 3600.	<code>nss/BufferFlushTimer</code>

Setting the Name Cache

This parameter specifies the amount of recently used Name Tree entries that NSS will cache:

```
nss /NameCache=xx
```

The default is 2111. The range is 3 - 65,521. Each entry uses about 150 bytes of memory.

In medium to large servers with enough RAM, you might consider increasing this value to 40,000.

HINT: Name cache will grow up to the specified limit. Unlike the file system cache, it does not take the maximum amount of memory it needs from the start.

Tuning NSS for GroupWise Servers

NSS has been dramatically improved in NetWare 6.5. As a result, the only tuning that you need to do to enhance the performance of NSS and GroupWise is to disable the salvage feature by entering the following at the server command line:

```
nss /NoSalvage
```

NetWare 6 servers

If your GroupWise performance is suffering on a NetWare 6 server, you can adjust the following parameters to increase efficiency:

- ◆ Increase the CacheBalance size to 85
- ◆ Verify that the NameCache is set at 20,000
- ◆ Disable the salvage feature by entering **nss /NoSalvage** at the server console
- ◆ Increase the ClosedFileCacheSize to 100,000 and the OpenFileCacheSize to 17

11

Upgrading and Copying Volumes

This section discusses the following:

- ◆ How to upgrade existing NetWare® 5 NSS volumes after you have upgraded the operating system to NetWare 6.5.

IMPORTANT: When you upgrade the operating system of a NetWare 5 server to a NetWare 6.5 server, the NSS volumes are not automatically upgraded and will not function until you upgrade them. If you upgrade your server from NetWare 6 to NetWare 6.5, no NSS upgrade is necessary.

For instructions, see “[Upgrading NetWare 5 NSS Volumes](#)” on page 77.

- ◆ How to copy data from traditional or NSS volumes to NetWare 6.5 NSS volumes.

For instructions, see “[Copying Data from Existing Traditional or NSS Volumes to NetWare 6.5 NSS Volumes](#)” on page 78.

Upgrading NetWare 5 NSS Volumes

- 1 Check the time stamp.

Before you upgrade NSS volumes, you must ensure that Novell® eDirectory™ is running properly; otherwise, the upgrade process might remove your trustee assignments. To ensure eDirectory is running, you need to check the time stamp for the backlink process.

- 1a Open iMonitor in a network browser by entering the IP address of the server followed by /nds.

For example: `http://192.168.1.1/nds`

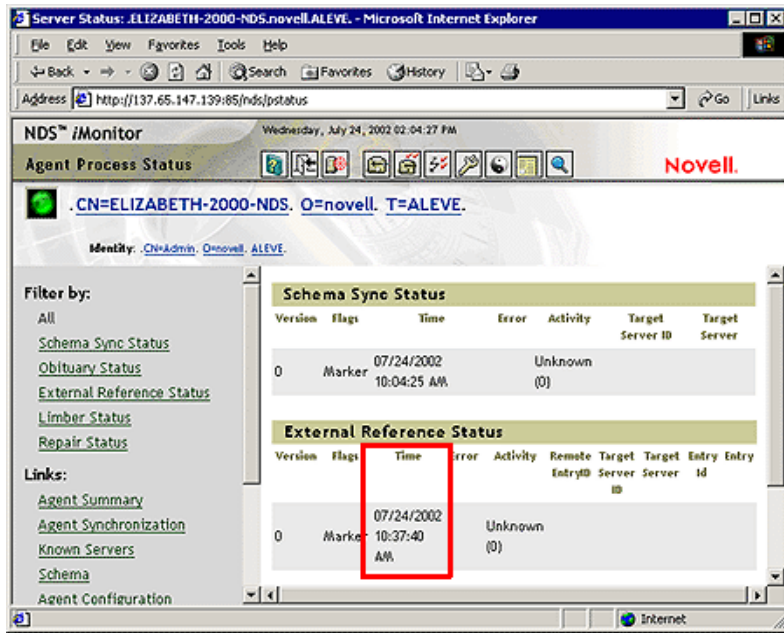
- 1b Click Agent Process Status > External Reference Status.

- 1c Make sure the time stamp (under the Time column) shows a time later than the time you began the upgrade—specifically, look for the time stamp when eDirectory unloaded during the upgrade.

IMPORTANT: If the time stamp shows a time earlier than when you began the upgrade, the backlink process is not complete. Do not upgrade your existing Netware 5 NSS volumes until the backlink is complete.

The following graphic provides an example of the time stamp:

Description: Example of time stamp



2 Enter one of the following at the server console:

- ◆ `nss /z1ssvolumeupgrade=all`
Specify "all" to upgrade all the NetWare 5 volumes on the selected server.
- ◆ `nss /z1ssvolumeupgrade=volumename`
Use volumename

Copying Data from Existing Traditional or NSS Volumes to NetWare 6.5 NSS Volumes

The Volume Copy Upgrade (VCU) is a volume conversion utility that lets you copy metadata and user data from preexisting NetWare 5 traditional volumes and NSS volumes to NetWare 6.5 NSS volumes. Before you begin, make sure that the NSS pool that you are copying data to has enough space to accommodate the source volume's data. For example, if you want to move 2 GB of viable data (disregard free blocks or salvage blocks), the NSS volume needs to have at least 2 GB of free space available.

VCU can be run from the server command line or from a GUI application that runs on the server or on a workstation. For instructions, see the following sections:

- ◆ "Using VCU from the Command Line" on page 79
- ◆ "Using VCU from a Java Application" on page 80

Using VCU from the Command Line

Copying Volumes

- 1 At the command line, specify the volume that you want to copy and then specify the NSS pool where you want to store the data from the volume (VCU can convert volumes with long names) by using this format:

```
vcu original_volume_name NSS_pool_name
```

HINT: For customization options, see [“Using Switches to Customize VCU” on page 79](#).

For example, if you wanted to copy the Apps volume and put the data on a volume in Pool2, you would enter the following at the server console:

```
vcu apps pool2
```

In this example, VCU copies the data in the volume Apps and places it on a volume on Pool2, named apps_new.

When volumes are copied using VCU, a new NSS volume named *originalvolumename_new* is automatically created for you in the specified pool. If the original volume name is longer than 10 characters, or if the newly created NSS volume name conflicts with another existing volumes, the new NSS volume will be named VCUXXXX_new, where the XXXX is a random number generated by NSS.

- 2 (Optional) Rename your volumes.

After the volume copy is finished, VCU will ask you if you want to rename your volumes. If you select No, VCU is finished. If you select Yes, VCU will rename your original volume to *originalvolumename_old* and the new NSS volume from *originalvolumename_new* back to the name of the original volume.

HINT: By using the /m switch, VCU also lets you rename the volume’s default eDirectory object for both NSS and traditional volumes (*except* for the SYS: volume).

- 3 Review the information in [“After the Volume Copy” on page 83](#).

Restoring Volumes

To restore a volume, enter the following at the server console:

```
vcu /r VCUTargetVolName VCUTargetVOLName_old
```

For example, if wanted to restore Vol2, you would enter the following:

```
vcu /r Vol2_new Vol2_old
```

VCU would then delete Vol2_new and then rename Vol2_old back to its original name: Vol2.

HINT: To rename the volume’s default eDirectory object, insert the /m switch into the command line.

Using Switches to Customize VCU

This section describes the different customization options (switches) that are available to you when you copy or restore volumes.

To view the available switches, enter **vcu /h** at the server console. These switches are optional; you can include any combination of them in your VCU command line:

The following tables list the possible VCU switches when copying and restoring volumes:

Table 5 VCU Switches for Copying Volumes

VCU Switch	Description
/d	Deletes 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.
/l	Does not write errors to log file (Dst_Vol:error.out).
/m	Renames the physical volume as well as it's default eDirectory objects.
/p	Does not display the entire copy status on the VCU screen. You will only get a brief message after the volume copy is completed.

Table 6 VCU Switches for Restoring Volumes

VCU Switch	Description
/r	Lets you restore the old volume to its original name. (Use this option if you want to delete the new NSS volume and rename the old one with it's original name.)
/m	Renames the physical volume as well as it's default eDirectory objects.

Using VCU from a Java Application

You can also upgrade volumes using VCU's Java* application called vcuApp.jar. You can access the VCU GUI application from a server or a workstation.

Accessing VCU from the Server

- 1 Use one of the following methods:
 - ◆ At the server console, type **vcux**, then press enter.
 - ◆ At the GUI console, click Novell > Utilities > Volume Copy Upgrade.

Accessing VCU from a Workstation

- 1 Make sure Java 1.4 or later is installed on your workstation.
- 2 Map a drive to the _Admin volume.
- 3 To run the VCU application from a server:
 - 3a Map a drive on the workstation to SYS:java\lib\ on the server.
 - 3b Open a directory browser to the mapped drive, then double-click the vcuApp.jar icon.
- 4 To run the VCU application locally on the workstation:
 - 4a Browse to SYS:\java\lib\vcuApp.jar on the server.
 - 4b Copy vcuApp.jar to a directory on your workstation.

4c To run the application, use one of these methods:

- ◆ Open the local directory, then double-click the vcuApp.jar icon.
- ◆ At the DOS prompt, enter

```
java -jar your_directory\vcuApp.jar
```

Continue with the next section, "[Copying Volumes](#)."

Copying Volumes

1 Stop all applications on the server that access the volumes you are copying.

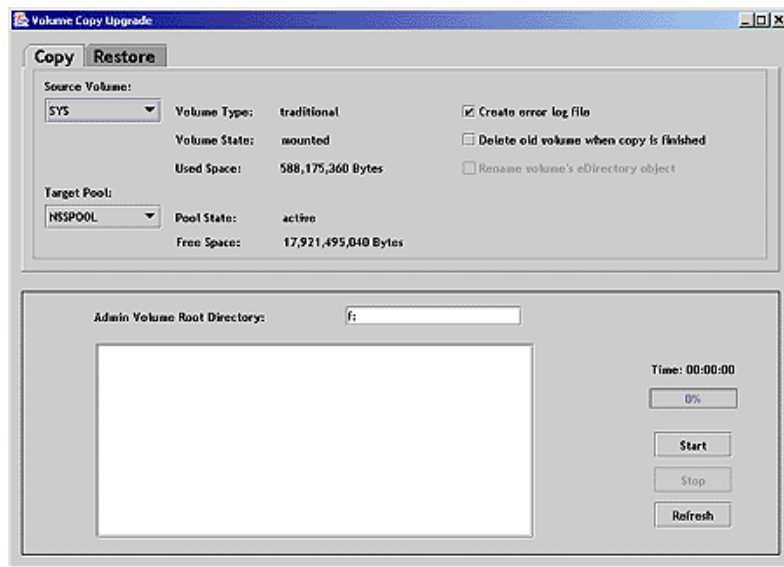
WARNING: Before using this utility to copy traditional or NSS volumes, *stop* all applications that access the volume you are copying. This is especially important for the sys: volume. If any applications are open on the sys: volume when you begin the volume copy, your sys: volume might become corrupted and render your server unusable.

2 Launch VCU.

For instructions, see "[Using VCU from a Java Application](#)" on page 80.

The following screen shot is an example of what the VCU Copy screen looks like.

Figure 3 VCU Copy screen



3 Click the Copy tab.

4 (Conditional) If you are running VCU from a workstation, enter the drive letter of the drive that you mapped to the _Admin volume, followed by a colon (:), into the Admin Volume Root Directory text box and then click Refresh.

HINT: If you are running VCU from the server, the Admin Volume Root Directory text box is automatically filled in for you.

5 Select the source volume and the target pool from the drop-down lists.

HINT: If you have recently added a pool or volume via ConsoleOne or iManager, click the Refresh button to see the new pool or volume appear in the drop-down list.

Make sure the target pool has enough space to accommodate the data that will be copied from the source volume.

- 6 Refer to the following table to determine whether you want to create an error log, rename the copied volume's default eDirectory object, or delete the old volume after the volume data is copied:

Options	Results
Create an Error Log	Creates an error log called ERROR.OUT is created and saved at the root directory of the destination volume. For more information, see "Vcu.nlm" on page 11 .
Delete Old Volume When Copy Is Finished	Deletes the old (original) physical volume is deleted; however, the volume's eDirectory object remains.
Rename Volume's eDirectory Object	Renames the volume's default eDirectory object is renamed. For more information, see "Copying Volumes" on page 79 .

- 7 Click Start to begin the volume copy.

HINT: If you stop the volume copy for any reason, when you click Start again, the volume copy does not resume where it left off, but starts again at the beginning.

- 8 (Optional) Rename your volumes.

After the volume copy is finished, VCU will ask you if you want to rename your volumes. If you click Cancel, VCU is finished. If you click Ok, VCU will rename your original volume to *originalvolumename_old* and the new NSS volume from *originalvolumename_new* back to the name of the original volume.

- 9 Review the information in ["After the Volume Copy" on page 83](#).

Restoring Volumes

- 1 Stop all applications on the server that access the volumes you are copying.

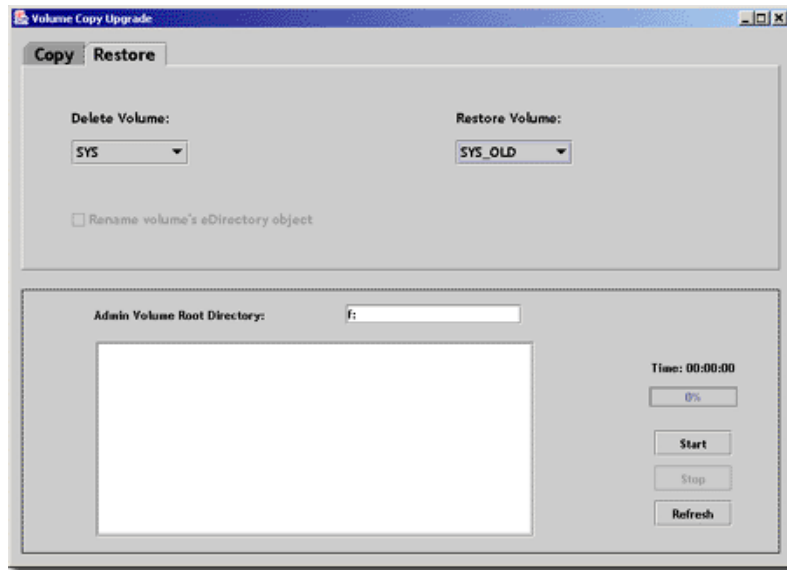
WARNING: Before using this utility to copy traditional or NSS volumes, *stop* all applications that access the volume you are copying. This is especially important for the SYS: volume. If any applications are open on the SYS: volume when you begin the volume copy, your SYS: volume might become corrupted and render your server unusable.

- 2 Launch VCU.

For instructions, see ["Using VCU from a Java Application" on page 80](#).

The following screen shot is an example of the VCU Restore screen.

Figure 4 VCU Restore screen



3 Click the Restore tab.

4 (Conditional) If you are running VCU from a workstation, enter the drive letter of the drive that you mapped to the `_Admin` volume, followed by a colon (:), into the Admin Volume Root Directory text box and then click Refresh.

HINT: If you are running VCU from the server, the Admin Volume Root Directory text box is automatically filled in for you.

5 Select the volume that you want to delete and restore from the drop-down lists.

HINT: If you have recently added a pool or volume via ConsoleOne or iManager, click the Refresh button to see the new pool or volume appear in the drop-down list.

6 Determine if you want to rename the volume's default eDirectory object.

7 Click Start.

When you click Start, the volume that you are restoring will automatically be renamed to the name of the volume that you are deleting.

After the Volume Copy

- Restart the server if upgrading the `SYS:` volume to ensure the volume copied properly.
- Verify that the volume data was copied successfully.

IMPORTANT: After you copy traditional volume data to an NSS volume, you cannot access the new NSS volume from NetWare 5.1. You would need to either retain the data in the traditional volume, or copy the data back to the traditional volume.

12 Troubleshooting

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

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 “Viewing Partitioned Free Space” on page 31. ◆ 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.
Cannot compress a file	You must choose the file compression option when you create a logical volume.	<p>Apply the file compression option to an existing logical volume:</p> <p>From ConsoleOne, click Media > NSS Logical Volumes > Properties > Attributes > NSS Attributes. Then select compression.</p>
Cannot see your volumes or devices associated with those volumes	You might not have the appropriate multipath drivers that support the multipath functionality.	<p>Turn off multipath support.</p> <p>For further instructions, see “Using Multiple Paths to Access Devices” on page 19.</p>

A NSS Change History

This section describes changes to the Novell Storage Service since the NetWare® 6 release.

January 2003

Expanded the functionality of the pool snapshot features. See [“Pool Snapshots” on page 40](#) for details.

Created a special-purpose server installation for Network Attached Storage server configurations.

Removed support for Hot Fix™. This feature is provided via hardware in modern solutions.

