

# Novell Open Enterprise Server

[www.novell.com](http://www.novell.com)

April 12, 2006

NOVELL STORAGE SERVICES™ FILE  
SYSTEM 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.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. Please refer to [www.novell.com/info/exports/](http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2005 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.

Novell, Inc., has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents/> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.  
404 Wyman Street, Suite 500  
Waltham, MA 02451  
U.S.A.  
[www.novell.com](http://www.novell.com)

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

## **Novell Trademarks**

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

eDirectory is a trademark of Novell, Inc.

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

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

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

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

Novell Client is a trademark of Novell, Inc.

Novell Cluster Services is a trademark of Novell, Inc.

Novell Core Protocol and NCP are trademarks of Novell, Inc.

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

Novell Storage Services is a trademark of Novell, Inc.

Storage Management Services is a trademark of Novell, Inc.

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

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

## **Third-Party Materials**

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

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org>).



# Contents

<b>About This Guide</b>	<b>19</b>
<b>1 Overview of NSS</b>	<b>23</b>
1.1 Benefits of NSS	23
1.2 Understanding NSS	24
1.2.1 Storage Pools	24
1.2.2 NSS Volumes	25
1.3 NSS Features	25
1.3.1 Use Less Memory and Gain More Speed	26
1.3.2 Improve Storage Availability	26
1.3.3 Prevent Unauthorized Access	27
1.3.4 Protect Data from Corruption or Loss	27
1.3.5 Maximize Available Space	28
1.4 What's Next	28
<b>2 What's New</b>	<b>29</b>
2.1 What's New (OES SP2 and NetWare 6.5 SP5)	29
2.2 What's New (OES SP1 and NetWare 6.5 SP4)	29
2.2.1 Linux	29
2.2.2 NetWare	30
2.3 What's New (OES and NetWare 6.5 SP3)	30
2.3.1 Linux	30
2.3.2 NetWare	32
<b>3 Upgrading the Media Format (NetWare)</b>	<b>33</b>
3.1 Guidelines for Upgrading the Media Format of NSS Volumes on NetWare	33
3.1.1 Support for the NSS Metadata Structure Upgrade	33
3.1.2 Before Upgrading the Media Format	34
3.1.3 Which NSS Volumes to Upgrade	34
3.1.4 After Upgrading the Media Format	35
3.1.5 If You Do Not Upgrade the Media Format	35
3.2 Upgrading the Media Format for NSS Volumes on NetWare	35
3.2.1 Upgrading the Structure for Currently Active NSS Volumes	36
3.2.2 Upgrading the Structure for New NSS Pools and Volumes	36
<b>4 Management Tools for NSS</b>	<b>39</b>
4.1 Novell iManager Storage Plug-In	39
4.1.1 Prerequisites for Using the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager 2.5	40
4.1.2 Downloading and Installing Plug-in Updates	42
4.1.3 Accessing Novell iManager	45
4.1.4 Accessing NSS Storage Management Role and Tasks	46
4.1.5 Selecting a Server to Manage	46
4.1.6 Understanding iManager Storage Management Tasks	47
4.2 NSS Management Utility	52
4.2.1 NSSMU for Linux	53
4.2.2 NSSMU for NetWare	54

4.3	NSS Commands and Utilities . . . . .	56
4.3.1	NSS Commands . . . . .	57
4.3.2	NSS Utilities . . . . .	57
4.3.3	NSS Console for OES Linux . . . . .	57
4.3.4	NetWare Server Console . . . . .	58
4.4	Novell NetStorage . . . . .	58
4.4.1	Accessing NetStorage . . . . .	58
4.4.2	Configuring File System Trustees, Trustee Rights, and Attributes . . . . .	59
4.4.3	Purging and Salvaging Deleted Files . . . . .	59
4.4.4	Browsing Directories and Files . . . . .	59
4.4.5	Additional Information . . . . .	59
4.5	Novell Remote Manager . . . . .	60
4.5.1	Novell Remote Manager for Linux . . . . .	60
4.5.2	Novell Remote Manager for NetWare . . . . .	60
4.5.3	Accessing Novell Remote Manager . . . . .	60
4.6	Novell Client . . . . .	62
4.6.1	Novell Client for Linux . . . . .	62
4.6.2	Novell Client for Windows 2000/XP/2003 . . . . .	62
4.7	ConsoleOne . . . . .	62
4.8	Virtual File Services, APIs, and Scripts . . . . .	63
<b>5</b>	<b>Planning for NSS Storage Solutions</b>	<b>65</b>
5.1	Prerequisites . . . . .	65
5.2	Planning to Move NSS Pools and Volumes Cross-Platform . . . . .	66
5.3	Planning for Your Storage Needs . . . . .	66
5.3.1	Devices . . . . .	67
5.3.2	Software RAID Devices . . . . .	67
5.3.3	Multiple I/O Paths (NetWare) . . . . .	68
5.3.4	Free Space . . . . .	68
5.3.5	NSS Pools and Volumes . . . . .	68
5.3.6	NSS Encrypted Volumes . . . . .	69
5.4	Creating NSS Storage Objects in eDirectory . . . . .	69
5.5	Naming NSS Storage Objects . . . . .	69
5.6	What's Next . . . . .	71
<b>6</b>	<b>Coexistence and Migration Issues</b>	<b>73</b>
6.1	Compatibility with Operating Systems and Services . . . . .	73
6.2	NSS Kernel Modules in SLES 9 . . . . .	74
6.3	Comparison of NSS to Other File Systems . . . . .	74
6.4	Access Control Issues for NSS on OES Linux . . . . .	74
6.4.1	NSS and eDirectory . . . . .	75
6.4.2	NSS, the Admin User, and the Root User . . . . .	75
6.4.3	NSS and File System Users . . . . .	76
6.4.4	NSS and NCP Server . . . . .	77
6.4.5	NSS and Linux User Management . . . . .	77
6.4.6	Enforcing File Ownership and User Space Restrictions . . . . .	78
6.4.7	Accessing Files as the Root User . . . . .	78
6.4.8	Accessing Files via NCP Server Only . . . . .	78
6.4.9	Accessing Files with Other Protocols or from Linux . . . . .	79
6.4.10	Accessing Linux Traditional Volumes with NCP Server and Services . . . . .	79
6.5	Compatibility Issues for the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager . . . . .	79
6.5.1	Novell iManager 2.5 . . . . .	80
6.5.2	Web Browser Language Setting . . . . .	80

6.5.3	Protocols . . . . .	80
6.6	Compatibility Issues for Using NSS Cross-Platform . . . . .	82
6.6.1	Guidelines for File Systems on Linux Servers. . . . .	82
6.6.2	Guidelines for File Systems on NetWare Servers . . . . .	82
6.6.3	Guidelines for Cross-Platform Compatibility of NSS . . . . .	83
6.6.4	Guidelines for Moving Devices Cross-Platform . . . . .	85
6.7	Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux. . . . .	87
6.7.1	Prerequisites . . . . .	87
6.7.2	Setting Up File Access For Users on the OES Linux Server. . . . .	88
6.7.3	Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . . .	89
6.7.4	Recommissioning Each NSS Pool and Its Volumes on the Destination Server . . . . .	90
6.7.5	Using Scripts to Decommission and Recommission NSS Volumes . . . . .	90
6.8	Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later . . . . .	92
6.8.1	Prerequisites . . . . .	92
6.8.2	Setting Up File Access For Users on the OES Linux Server. . . . .	93
6.8.3	Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . . .	93
6.8.4	Recommissioning Each NSS Pool and Its Volumes on the Destination Server . . . . .	94
6.9	Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare . . . . .	95
6.9.1	Prerequisites . . . . .	96
6.9.2	Setting Up File Access For Users on the Destination Server . . . . .	96
6.9.3	Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . . .	96
6.9.4	Recommissioning Each NSS Pool and Its Volumes on the Destination Server . . . . .	97
6.10	Moving Clustered Devices with NSS Volumes Cross-Platform . . . . .	98
6.11	Upgrading NetWare 5.x NSS Volumes and NetWare Traditional Volumes to NSS Volumes . . . . .	98

## **7 Managing Devices 99**

7.1	Understanding Devices . . . . .	99
7.2	Making Devices Available to EVMS (Linux). . . . .	100
7.3	Viewing Devices on a Server. . . . .	101
7.4	Viewing Details for a Device . . . . .	102
7.5	Scanning for Devices. . . . .	103
7.6	Initializing a Disk . . . . .	104
7.7	Sharing Devices. . . . .	104
7.7.1	Planning for Device Sharing . . . . .	105
7.7.2	Configuring the Device's Share State . . . . .	105
7.8	Viewing Partitions on a Device . . . . .	106
7.9	Viewing Pools on a Device . . . . .	107
7.10	What's Next . . . . .	107

## **8 Managing Partitions 109**

8.1	Understanding Partitions . . . . .	109
8.1.1	NSS Partitions. . . . .	109
8.1.2	Viewing Partitions on a Device . . . . .	110
8.1.3	Understanding Types of Partitions . . . . .	110
8.1.4	Understanding Partition Details. . . . .	111
8.2	Creating NSS Partitions. . . . .	111
8.3	Viewing a List of Partitions. . . . .	112
8.3.1	Viewing Partitions on a Device . . . . .	112
8.3.2	Viewing Partitions in a Software RAID Device . . . . .	113
8.3.3	Viewing Partitions in an NSS Pool . . . . .	113
8.4	Viewing Details for a Partition . . . . .	114
8.5	Labeling a Partition . . . . .	114
8.6	Deleting a Partition . . . . .	114

8.6.1	Deleting NSS Partitions in an NSS Software RAID Device . . . . .	114
8.6.2	Deleting NSS Partitions in a Pool. . . . .	115
8.6.3	Deleting Cluster Services Partitions. . . . .	115
8.6.4	Deleting an iSCSI Partition. . . . .	116
8.7	Adding Other Types of File System Partitions (Linux) . . . . .	116
<b>9</b>	<b>Managing Software RAID Devices</b>	<b>117</b>
9.1	Understanding Software RAID Devices . . . . .	117
9.2	Planning for a Software RAID Device. . . . .	119
9.2.1	General Guidelines for Software RAID Devices. . . . .	119
9.2.2	Guidelines for Software RAID 1 Devices . . . . .	120
9.2.3	Drive Restrictions for NSS Software RAID 0 and 1 Devices on Linux . . . . .	121
9.2.4	Choosing a Software RAID Solution . . . . .	121
9.2.5	Determining the Partition Size . . . . .	122
9.2.6	Determining the Number of Partitions . . . . .	123
9.2.7	Determining the Stripe Size for RAID 0 and RAID 5 . . . . .	123
9.3	Viewing Software RAID Devices on a Server. . . . .	124
9.4	Viewing Details of a Software RAID Device . . . . .	125
9.5	Creating a Software RAID Device Using iManager . . . . .	126
9.6	Renaming a Software RAID Device . . . . .	128
9.7	Increasing the Size of a Software RAID Device . . . . .	128
9.8	Restriping a Software RAID 0 or 5 Device . . . . .	129
9.9	Replacing a Failed Segment in a Software RAID. . . . .	130
9.10	Deleting a Software RAID Device. . . . .	132
9.11	Viewing Pools on a Software RAID Device . . . . .	133
9.12	Viewing Partitions on a Software RAID Device . . . . .	133
9.13	Deleting Partitions on a Software RAID Device . . . . .	134
9.14	Managing Software RAID Devices with NSSMU . . . . .	134
9.15	Creating Software RAID Devices with NSSMU . . . . .	135
9.16	Mirroring an Existing Pool with NSSMU . . . . .	136
9.17	Creating a Software RAID 10 with NSSMU . . . . .	136
9.18	Creating a Software RAID 15 with NSSMU . . . . .	137
9.19	Creating a RAID 1 Mirror to Duplicate Data . . . . .	137
<b>10</b>	<b>Managing Multiple Connection Paths to Devices (NetWare)</b>	<b>139</b>
10.1	Planning Your Multiple Connection Path Solution . . . . .	139
10.2	Configuring the Primary Path and Priorities for Failover Connection Paths . . . . .	140
10.3	Viewing Connection Path Details . . . . .	141
10.4	Setting a Connection Path as Up or Down . . . . .	141
10.5	Setting the Primary Path for a Device to Its Default Path. . . . .	142
10.6	Resetting the Server Registry with Default Priority Settings for a Device. . . . .	143
10.7	Managing Multiple Paths with Server Console Commands . . . . .	143
<b>11</b>	<b>Managing NSS Pools</b>	<b>147</b>
11.1	Creating Pools with iManager. . . . .	147
11.2	Viewing Pools on a Server . . . . .	150
11.3	Viewing Pool Details. . . . .	151
11.4	Viewing Partition Information for a Pool . . . . .	152
11.5	Viewing Volume Information for a Pool . . . . .	153
11.6	Viewing Device Information for a Pool . . . . .	153



11.7	Renaming a Pool . . . . .	154
11.8	Increasing the Size of a Pool . . . . .	154
11.9	Activating and Deactivating Pools . . . . .	155
11.10	Deleting a Pool . . . . .	156
11.11	Verifying and Rebuilding NSS Pools and Volumes . . . . .	157
11.11.1	When to Use Verify and Rebuild . . . . .	157
11.11.2	Mounting the Volume to Repair Journaled Errors . . . . .	157
11.11.3	Verifying the Pool to Identify Inconsistent Data Blocks . . . . .	158
11.11.4	Reviewing Log Files for Errors . . . . .	159
11.11.5	Restoring the Data from Backup Media . . . . .	160
11.11.6	Rebuilding NSS Pools to Repair Data Consistency . . . . .	160
11.12	ReZIDing Volumes in an NSS Pool . . . . .	162
11.12.1	Understanding ReZID . . . . .	162
11.12.2	When to ReZID Volumes During a Pool Rebuild . . . . .	162
11.12.3	ReZIDing Volumes on Linux . . . . .	163
11.12.4	ReZIDing Volumes on NetWare . . . . .	163
11.13	Updating eDirectory Pool Objects . . . . .	164
11.14	Preventing Pools from Activating on Multiple Servers . . . . .	164
11.14.1	Managing MSAP Using Server Console Commands . . . . .	165
11.15	What's Next . . . . .	166
<b>12</b>	<b>Managing NSS Pool Snapshots (NetWare)</b>	<b>167</b>
12.1	Benefits of Pool Snapshots . . . . .	167
12.2	How the Pool Snapshot Works . . . . .	168
12.3	Understanding Key Concepts for Pool Snapshots . . . . .	169
12.3.1	Creating a Pool Snapshot . . . . .	169
12.3.2	Naming Pool Snapshots . . . . .	169
12.3.3	Choosing the Stored-On Pool . . . . .	170
12.3.4	Maintaining the Stored-on Pool . . . . .	170
12.3.5	Activating Pool Snapshots . . . . .	171
12.3.6	Deleting Pool Snapshots . . . . .	171
12.4	Managing Pool Snapshots Using iManager . . . . .	172
12.4.1	Creating a New Pool Snapshot . . . . .	172
12.4.2	Viewing Pool Snapshot Details . . . . .	174
12.4.3	Renaming a Pool Snapshot . . . . .	175
12.4.4	Activating and Deactivating a Pool Snapshot . . . . .	175
12.4.5	Moving an Active Pool Snapshot Offline from the Pools Page . . . . .	177
12.4.6	Deleting a Pool Snapshot . . . . .	178
12.4.7	Modifying the Stored-On Pool . . . . .	178
12.5	Managing Pool Snapshots Using Commands . . . . .	179
<b>13</b>	<b>Managing NSS Volumes</b>	<b>181</b>
13.1	Understanding NSS Volumes . . . . .	181
13.1.1	Moving NSS Volumes Across Platforms . . . . .	182
13.2	Understanding Volume Attributes . . . . .	183
13.2.1	Encryption . . . . .	183
13.2.2	Data Shredding . . . . .	183
13.2.3	Compression . . . . .	183
13.2.4	Data Migration . . . . .	184
13.2.5	Directory Quotas . . . . .	184
13.2.6	User Space Quotas . . . . .	184
13.2.7	Salvage Files . . . . .	184
13.2.8	Lookup Name Space . . . . .	184
13.2.9	Read Ahead Blocks in NSS Volumes . . . . .	184

13.2.10	Flush Files Immediately . . . . .	185
13.2.11	User-Level Transaction Model (NetWare) . . . . .	185
13.2.12	File-Level Snapshot (NetWare) . . . . .	185
13.2.13	Modified File List . . . . .	186
13.2.14	Backup . . . . .	186
13.3	Creating and Configuring Unencrypted NSS Volumes with iManager . . . . .	186
13.4	Configuring Encrypted NSS Volumes with NSSMU . . . . .	188
13.5	Updating eDirectory Volume Objects . . . . .	188
13.6	Viewing the Details of an NSS Volume . . . . .	189
13.7	Viewing Properties of an NSS Volume . . . . .	189
13.8	Modifying Attributes of an NSS Volume . . . . .	191
13.9	Renaming an NSS Volume . . . . .	192
13.10	Activating and Deactivating an NSS Volume . . . . .	192
13.11	Activating and Deactivating CDs and DVDs as NSS Volumes (NetWare) . . . . .	193
13.12	Activating and Deactivating CD and DVD Image Files as NSS Volumes (NetWare) . . . . .	194
13.13	Activating and Deactivating DOS Partitions as NSS Volumes (NetWare) . . . . .	195
13.14	Mounting and Dismounting an NSS Volume (Linux) . . . . .	195
13.15	Using NFS Mount for NSS Volumes (Linux) . . . . .	196
13.16	Mounting an NSS Volume on Linux to Support Case Insensitive File Names (Linux) . . . . .	197
13.17	Mounting a Disk in a USB Floppy Drive as a Local FAT Volume (NetWare) . . . . .	198
13.18	Deleting an NSS Volume . . . . .	199
13.19	Verifying and Rebuilding NSS Pools and Volumes . . . . .	200
13.20	What's Next . . . . .	200

## **14 Managing Encrypted NSS Volumes 201**

14.1	Understanding Encrypted Volume Support . . . . .	201
14.1.1	Encryption Method . . . . .	202
14.1.2	Encryption Password . . . . .	202
14.1.3	How Encrypted Volume Support Works . . . . .	202
14.1.4	Guidelines for Using Encrypted Volumes . . . . .	202
14.2	Security Considerations for Encrypted Volumes . . . . .	203
14.2.1	Choosing a Strong Encryption Password . . . . .	203
14.2.2	Backing Up Data from an Encrypted Volume . . . . .	203
14.2.3	Excluding the NSS Cache Memory from Core Dumps . . . . .	203
14.2.4	Disabling Logs (Linux) . . . . .	203
14.2.5	Archiving Files from an Encrypted Volume . . . . .	204
14.2.6	Using Direct I/O to an Encrypted Volume . . . . .	204
14.2.7	Sharing Encrypted NSS Volumes in a Cluster . . . . .	204
14.3	Creating an Encrypted Volume . . . . .	204
14.4	Mounting an Encrypted NSS Volume with NSSMU (Linux) . . . . .	205
14.5	Activating an Encrypted NSS Volume with NSSMU (NetWare) . . . . .	205
14.6	Mounting Encrypted NSS Volumes with NSS Commands (Linux) . . . . .	206
14.7	Activating Encrypted NSS Volumes with NSS Commands (NetWare) . . . . .	206
14.8	Dismounting an Encrypted NSS Volume with NSSMU (Linux) . . . . .	207
14.9	Deactivating an Encrypted NSS Volume with NSSMU (NetWare) . . . . .	208
14.10	Using Encrypted Volumes in a Server Cluster (Linux) . . . . .	208
14.11	Using Encrypted Volumes in a Server Cluster (NetWare) . . . . .	209
14.12	Removing Encrypted Volumes . . . . .	209
14.13	What's Next . . . . .	210

## **15 Securing Access to NSS Volumes, Directories, and Files 211**

15.1	Controlling Access to the Admin Volume and Files.cmd (NetWare) . . . . .	211
------	--	-----

15.1.1	Assigning Volume Rights with Trustees.xml (NetWare) . . . . .	211
15.1.2	Restricting Access to Files.cmd (NetWare). . . . .	214
15.2	Using Data Shredding to Prevent Access to Deleted Files . . . . .	215
15.2.1	Setting the Data Shredding Attribute When You Create a Volume . . . . .	215
15.2.2	Setting the Data Shredding Attribute for an Existing Volume . . . . .	215
15.2.3	Disabling Data Shredding for an Existing Volume. . . . .	216
15.3	Managing File System Access to Directories and Files . . . . .	216
<b>16</b>	<b>Using Volume Attributes that Help Ensure Data Integrity</b>	<b>217</b>
16.1	Using the Flush Files Immediately Attribute to Protect Against Possible Data Corruption on System Failure. . . . .	217
16.1.1	Setting the Flush Files Immediately Attribute When You Create a Volume . . . . .	217
16.1.2	Setting the Flush Files Immediately Attribute for an Existing Volume . . . . .	217
16.1.3	Enabling and Disabling the Flush Files Immediately Attribute for All Volumes on a Server . . . . .	218
16.2	Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare) . . . . .	219
16.3	Using the File-Level Snapshot Attribute to Enable the Backup of Open Files. . . . .	220
16.3.1	Understanding File-Level Snapshots . . . . .	220
16.3.2	How the File Snapshot Attribute Works. . . . .	221
16.3.3	Sizing the Volume to Accommodate File Snapshots. . . . .	221
16.3.4	Assessing I/O Overhead for File Snapshots . . . . .	222
16.3.5	Configuring Volumes with the File Snapshot Attribute . . . . .	223
16.3.6	Configuring Backup Software to Use File Snapshots . . . . .	224
16.3.7	Restoring Volumes from Backup Media . . . . .	225
<b>17</b>	<b>Managing Compression on NSS Volumes</b>	<b>227</b>
17.1	Understanding File Compression . . . . .	227
17.1.1	Compression Settings . . . . .	227
17.1.2	Compression and Decompression Processes. . . . .	228
17.1.3	Compression Criteria. . . . .	229
17.1.4	Decompression Criteria . . . . .	229
17.1.5	Compression and Storage Backup . . . . .	230
17.2	Planning for File Compression. . . . .	230
17.2.1	Inactivity Determines Which Files Are Eligible for Background Compression. . . . .	230
17.2.2	Some Files Do Not Compress Well. . . . .	231
17.2.3	Some Volumes Are Not Good Candidates for Compression. . . . .	231
17.2.4	Decompression Activity Depends on Available Space . . . . .	231
17.2.5	Immediate Compression Impacts CPU Performance . . . . .	231
17.2.6	Files Remain Compressed during Backup and Restore . . . . .	231
17.2.7	If Set, the Compression Attribute Persists for the Life of the Volume . . . . .	232
17.2.8	How to Monitor Compression Activity . . . . .	232
17.3	Configuring Compression Parameters for a Server. . . . .	232
17.3.1	Understanding Compression Parameters . . . . .	232
17.3.2	Using Monitor (NetWare). . . . .	234
17.3.3	Using Novell Remote Manager (NetWare) . . . . .	235
17.3.4	Using the Set Command (NetWare) . . . . .	235
17.4	Configuring a Volume for Compression. . . . .	235
17.4.1	Enabling Compression for a New Volume. . . . .	235
17.4.2	Enabling Compression for an Existing Volume . . . . .	236
17.5	Suspending Compression for Volumes or Files. . . . .	236
17.6	Disabling Compression for a Volume. . . . .	237
17.7	Restoring Data to a Uncompressed Volume . . . . .	237
17.8	Configuring Compression Preferences for Directories and Files . . . . .	237

17.8.1	Using the Novell Client . . . . .	237
17.8.2	Using ATTRIB (Linux) . . . . .	238
17.8.3	Using FLAG (NetWare) . . . . .	239
17.9	Using NSS Commands to Configure and Monitor Compression . . . . .	241
17.10	Copying Compressed Files between Volumes with the NCOPY Command (NetWare). . . . .	242
17.11	Viewing the File Attributes with the NDIR Command (NetWare) . . . . .	242
17.12	Repairing Compressed Volumes with the Compfix Utility. . . . .	242
<b>18</b>	<b>Conserving Disk Space with Volume, Directory, or User Space Quotas</b>	<b>243</b>
18.1	Understanding Space Quotas . . . . .	243
18.2	Managing NSS Volume Quotas . . . . .	245
18.3	Managing Directory Quotas . . . . .	246
18.3.1	Setting the Directory Quotas Attribute for an NSS Volume . . . . .	246
18.3.2	Setting a Directory Quota . . . . .	247
18.3.3	Deleting All Directory Quotas for an NSS Volume . . . . .	250
18.4	Managing User Space Quotas . . . . .	250
18.4.1	Setting the User Space Quotas Attribute for an NSS Volume . . . . .	251
18.4.2	Viewing User Space Quotas . . . . .	251
18.4.3	Configuring User Space Quotas . . . . .	252
18.4.4	Modifying a User Quota . . . . .	253
18.4.5	Deleting a User Quota . . . . .	253
18.4.6	Setting Quotas for Users Who Were Not Initially Enabled for Linux Access (Linux) . . . . .	254
<b>19</b>	<b>Salvaging and Purging Deleted Volumes, Directories, and Files</b>	<b>255</b>
19.1	Configuring the NSS Volume Salvage Area . . . . .	255
19.2	Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool. . . . .	255
19.3	Setting the Low and High Watermarks for Purging the File System Salvage. . . . .	257
19.4	Salvaging or Purging Deleted Directories from the File System Salvage Area (OES SP1 NetWare, NetWare 6.5 SP4, or Later) . . . . .	257
19.5	Salvaging or Purging Deleted Files from the File System Salvage Area . . . . .	259
19.5.1	Salvaging or Purging Deleted Files with NetStorage . . . . .	259
19.5.2	Salvaging or Purging Deleted Files with the Novell Client . . . . .	259
19.5.3	Salvaging or Purging Deleted Files with Novell Remote Manager for NetWare (NetWare). . . . .	260
<b>20</b>	<b>Using Distributed File Services to Move and Split Volumes (NetWare)</b>	<b>261</b>
20.1	Understanding DFS . . . . .	261
20.1.1	Volume Location Database . . . . .	262
20.1.2	DFS Junctions . . . . .	262
20.1.3	Move Volume Jobs. . . . .	262
20.1.4	Split Volume Jobs . . . . .	263
20.1.5	Original and New Volumes. . . . .	263
20.2	Prerequisites and Guidelines for DFS . . . . .	263
20.2.1	Prerequisites for Creating and Using DFS Junctions. . . . .	263
20.2.2	Prerequisites for Pointing DFS Junctions at an NSS Volume on OES Linux. . . . .	264
20.2.3	Guidelines for Creating DFS Junctions Manually. . . . .	264
20.2.4	Guidelines for Original Volumes and New Volumes . . . . .	265
20.2.5	Guidelines for Managing Move Volume or Split Volume Jobs . . . . .	265
20.2.6	Guidelines for Files in Nearline Storage. . . . .	265
20.3	Creating a DFS Management Context for NCP . . . . .	266
20.4	Enabling DFS Junction Support for CIFS Clients. . . . .	266
20.5	Creating a DFS Junction Manually with ConsoleOne. . . . .	267

20.6	Viewing a DFS Junction . . . . .	268
20.7	Deleting a DFS Junction Manually with ConsoleOne . . . . .	268
20.8	Managing the Volume Location Database Service (NetWare) . . . . .	269
20.8.1	VLDB Commands . . . . .	269
20.8.2	VLDB Management Service in ConsoleOne . . . . .	269
20.9	Moving Volumes with iManager . . . . .	270
20.10	Splitting Volumes with iManager . . . . .	272
20.11	Managing Move or Split Jobs . . . . .	274
20.11.1	Viewing the Status of Move or Split Jobs . . . . .	275
20.11.2	Pausing a Move or Split Job . . . . .	276
20.11.3	Resuming a Move or Split Job . . . . .	277
20.11.4	Rescheduling a Move or Split Job . . . . .	277
20.11.5	Viewing Files Skipped by a Move or Split Job . . . . .	278
20.11.6	Deleting a Move or Split Job . . . . .	278
20.11.7	Finishing a Move or Split Job . . . . .	279
<b>21</b>	<b>Managing Hard Links</b>	<b>281</b>
21.1	Understanding Hard Links . . . . .	281
21.1.1	Hard Links and the Primary Link . . . . .	282
21.1.2	Hard Links and Linux Users . . . . .	282
21.1.3	Hard Links and File Ownership . . . . .	282
21.1.4	Hard Links and File System Trustees . . . . .	283
21.1.5	Hard Links and Directory Space Quotas . . . . .	283
21.1.6	Hard Links and User Space Quotas . . . . .	283
21.1.7	Hard Links and the Hard Link Attribute (OES SP1 NetWare, NetWare 6.5 SP4) . . . . .	284
21.1.8	Hard Links and File Salvage and Purge . . . . .	284
21.1.9	New Metadata Structure Supports Up to 65,535 Hard Links for a File (OES SP1 NetWare, NetWare 6.5 SP4) . . . . .	284
21.1.10	Old Metadata Structure Supports Limited Hard Links for a File . . . . .	285
21.2	Enabling or Disabling the Hard Links Attribute (OES SP1 NetWare, NetWare 6.5 SP4) . . . . .	285
21.3	Creating a Hard Link Using Ln on a Linux Client . . . . .	286
21.4	Creating a Hard Link Using a zLink API . . . . .	287
21.5	Creating a Hard Link for Testing Purposes (OES SP1 NetWare, NetWare 6.5 SP4) . . . . .	287
21.6	Viewing Hard Links for a Volume (OES SP1 NetWare, NetWare 6.5 SP4) . . . . .	288
21.7	Deleting a Hard Link . . . . .	288
<b>22</b>	<b>Using NSS on Devices Managed by Non-EVMS Volume Managers</b>	<b>291</b>
22.1	Guidelines for Using NSS with Non-EVMS Volume Managers . . . . .	291
22.1.1	Guidelines for the Server Install . . . . .	291
22.1.2	NSS Management Considerations . . . . .	291
22.1.3	NSS Capabilities That Are Unavailable . . . . .	292
22.2	Creating and Mounting NSS Pools and Volumes . . . . .	292
22.2.1	Using the Linux mkfs Command to Create NSS Pools . . . . .	293
22.2.2	Creating a Partition . . . . .	293
22.2.3	Creating and Mounting an NSS Pool . . . . .	294
22.2.4	Creating and Mounting an NSS Volume . . . . .	295
22.3	Configuring Default Mount Settings for NSS Pools and Volumes . . . . .	295
22.3.1	Understanding Entries in the /etc/fstab Configuration File . . . . .	296
22.3.2	Adding NSS Pool and Volume Mount Information to /etc/fstab . . . . .	297
<b>23</b>	<b>Using NSS Server Console Commands and Utilities</b>	<b>299</b>
23.1	Online Help Commands . . . . .	299
23.2	NSS Startup Commands . . . . .	300

23.2.1	Issuing NSS Startup Commands while OES NetWare Boots . . . . .	300
23.2.2	Making NSS Startup Commands Persist through an OES NetWare Reboot. . . . .	300
23.2.3	Load Commands . . . . .	301
23.3	Status Commands . . . . .	301
23.4	Maintenance Commands . . . . .	302
23.5	Maintenance Commands (NetWare) . . . . .	303
23.6	System Volume Recovery Commands (NetWare) . . . . .	304
23.7	Compression Commands . . . . .	305
23.8	Encrypted Volume Support Commands . . . . .	305
23.9	Salvage and Purge Commands . . . . .	306
23.10	Read Ahead Commands . . . . .	307
23.11	Cache Balancing Commands (NetWare) . . . . .	308
23.12	Commands for Mounting CDs and DVDs as NSS Volumes (NetWare) . . . . .	309
23.13	Commands for CD and DVD Image Files as NSS Volumes (NetWare) . . . . .	310
23.14	DOS FAT Commands (NetWare) . . . . .	311
23.15	Multiple Server Activation Prevention Commands . . . . .	312
23.16	Multipath Connection Failover Commands (NetWare) . . . . .	312
23.17	Pool Snapshot Commands (NetWare) . . . . .	314
23.18	Other NSS Commands . . . . .	315
23.19	Security Equivalence Vector Update Commands (Linux) . . . . .	316
<b>24</b>	<b>Monitoring the NSS File System Status . . . . .</b>	<b>319</b>
24.1	Monitoring Status of NSS Devices, Pools, and Volumes with iManager . . . . .	319
24.2	Monitoring Compression and Salvage Statistics . . . . .	320
24.3	Monitoring Quotas and Space Usage for NSS Pools and Volumes . . . . .	322
24.4	Monitoring File System Parameters . . . . .	324
24.4.1	Monitoring NSS File System Parameters with Novell Remote Manager . . . . .	325
24.4.2	Monitoring NSS File System Parameters with Novell NetStorage . . . . .	325
<b>25</b>	<b>Configuring the System Cache to Fine-Tune NSS Performance (NetWare) . . . . .</b>	<b>327</b>
25.1	Do I Need to Tune NSS? . . . . .	327
25.2	NSS Tuning Guidelines . . . . .	327
25.2.1	Before You Adjust Cache Parameters . . . . .	328
25.2.2	Hints for Low Memory NetWare Systems . . . . .	328
25.2.3	The Tuning Process . . . . .	328
25.2.4	Restoring Settings After a Failed Tuning Attempt . . . . .	329
25.3	Setting the Cache Balance . . . . .	329
25.4	Setting the Cache Buffers . . . . .	333
25.5	Setting the Closed File Cache Size . . . . .	335
25.6	Setting the File/Buffer Flush Timer . . . . .	336
25.7	Setting the Name Cache . . . . .	336
<b>26</b>	<b>Configuring the System Cache to Fine-Tune NSS Performance (Linux) . . . . .</b>	<b>339</b>
26.1	Understanding How NSS Uses Cache on Linux . . . . .	339
26.2	Setting the Minimum Number of Cache Buffers . . . . .	339
<b>27</b>	<b>Troubleshooting the NSS File System . . . . .</b>	<b>341</b>
27.1	Cannot See NSS Devices, Pools, or Volumes . . . . .	341
27.2	Cannot Create an NSS Pool or Volume . . . . .	342

27.3	eDirectory Error 613 Returned When Creating an NSS Pool or Volume on Linux . . . . .	343
27.4	NSS Volumes on Linux with Large Directories Mount Slowly with the UNIX Name Space . .	343
27.5	Linux Fails to Execute Dismount, Umount, or Delete Commands for NSS Volumes . . . . .	343
27.6	Slow Performance on SCSI Devices and RAID Controllers (Linux) . . . . .	344
27.7	File Compression Is Not Working . . . . .	345
27.8	Software RAID 1 Fails to Recognize a Replacement Device . . . . .	345
27.9	Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers . . . . .	345
27.10	Tuning NSS Volumes for GroupWise Servers . . . . .	345
27.11	Problems Following DFS Junctions in Windows 2000/XP Releases (NetWare) . . . . .	346
<b>28</b>	<b>Security Considerations</b>	<b>349</b>
28.1	Security Features of NSS . . . . .	349
28.2	Preventing Unauthorized Access . . . . .	350
28.3	Protecting Data During Backup and on Backup Media . . . . .	351
28.4	Preventing Exposure of Sensitive Data in a Core Dump . . . . .	351
28.5	Preventing Exposure of the Encryption Password in a Log (Linux) . . . . .	351
28.5.1	NSS Logging . . . . .	352
28.5.2	NSS Logging and Security Implications . . . . .	352
28.5.3	Logging Communications between NSS and the _ADMIN Volume . . . . .	352
28.5.4	Logging Communications between NSS and eDirectory, NCI, or Linux User Management . . . . .	354
28.6	Using Data Shredding to Prevent Unauthorized Access to Purged Deleted Files . . . . .	355
28.7	Protecting the Admin Volume and Files.cmd (NetWare) . . . . .	355
28.8	Acquiring Security Equivalence Vectors for NSS Users . . . . .	356
28.9	Protecting Modules Responsible for Security Equivalence Vectors (Linux) . . . . .	356
28.10	Controlling File System Access and Attributes for NSS Volumes . . . . .	356
28.11	Displaying Directory and File Attributes for NSS Volumes on Linux . . . . .	357
28.12	Controlling Physical Access to Servers and Resources . . . . .	357
28.13	Securing Access to the Servers with a Firewall . . . . .	357
28.14	Creating Strong Passwords . . . . .	357
<b>A</b>	<b>NSS Utilities</b>	<b>359</b>
A.1	ABORT REMIRROR (NetWare) . . . . .	359
A.1.1	Syntax . . . . .	359
A.1.2	Option . . . . .	360
A.1.3	Example . . . . .	360
A.2	ATTRIB (Linux) . . . . .	360
A.2.1	Syntax . . . . .	360
A.2.2	Options . . . . .	360
A.2.3	Attributes . . . . .	360
A.2.4	Example . . . . .	362
A.3	CDDVD (NetWare) . . . . .	362
A.3.1	Syntax . . . . .	362
A.3.2	Notes . . . . .	362
A.4	COMPFIX . . . . .	362
A.4.1	Purpose . . . . .	363
A.4.2	COMPFIX for Linux . . . . .	363
A.4.3	COMPFIX for NetWare . . . . .	366
A.5	DISMOUNT . . . . .	368
A.5.1	Syntax . . . . .	368
A.5.2	Option . . . . .	368
A.5.3	Notes . . . . .	368

A.6	FILER (NetWare)	368
A.6.1	Options	369
A.7	FLAG (NetWare)	369
A.7.1	Syntax	369
A.7.2	Options	370
A.7.3	Attributes	371
A.7.4	Examples	373
A.8	METAMIG (Linux)	374
A.8.1	Syntax	374
A.8.2	Arguments	374
A.8.3	Options	374
A.8.4	Examples	375
A.9	MIRROR STATUS (Netware)	376
A.9.1	Syntax	376
A.9.2	Option	376
A.9.3	Status Messages	376
A.10	MOUNT	376
A.10.1	Syntax	377
A.10.2	Notes	377
A.10.3	Examples	377
A.11	NCOPY (NetWare)	377
A.11.1	Syntax	377
A.12	NDIR (NetWare)	378
A.12.1	Syntax	378
A.13	RAVSUI (Linux)	378
A.13.1	Syntax	378
A.13.2	Arguments	378
A.13.3	Options	378
A.13.4	Example	379
A.14	RAVVIEW (Linux)	380
A.14.1	Syntax	380
A.14.2	Arguments	380
A.14.3	Options	380
A.14.4	Example	381
A.15	REBUILD (NetWare)	381
A.15.1	Syntax	382
A.15.2	Starting and Using REBUILD	382
A.16	REMIRROR PARTITION (NetWare)	382
A.16.1	Syntax	383
A.16.2	Example	383
A.17	RIGHTS	383
A.17.1	RIGHTS (Linux)	383
A.17.2	RIGHTS (NetWare)	386
A.18	SET (NetWare)	388
A.18.1	Syntax	388
A.18.2	Using SET	388
A.18.3	Common File System Parameters	389
A.18.4	Disk Parameters	392
A.18.5	Novell Storage Services Parameters	395
A.19	VERIFY (NetWare)	396
A.19.1	Syntax	396
A.19.2	Starting and Using VERIFY	397
A.20	VOLUME	397
A.20.1	Syntax	397
A.20.2	Using VOLUME	397



<b>B</b>	<b>Upgrading Legacy NSS and NetWare Traditional Volumes</b>	<b>399</b>
B.1	Migrating Legacy NSS Volumes and NetWare Traditional Volumes to OES Linux . . . . .	399
B.2	Copying and Upgrading Legacy Volumes with Volume Copy Upgrade . . . . .	399
B.2.1	Understanding Volume Copy Upgrade . . . . .	400
B.2.2	Using VCU from the Command Line . . . . .	401
B.2.3	Using VCU from a Java Application . . . . .	402
B.2.4	After the Volume Copy or Volume Upgrade . . . . .	406
B.3	Upgrading Legacy NSS Volumes . . . . .	406
B.4	Upgrading Legacy NSS Volumes in Clusters . . . . .	407
<b>C</b>	<b>Using EVMS to Manage Devices with NSS Volumes (Linux)</b>	<b>409</b>
C.1	Installing Linux with EVMS as the Volume Manager on a System Hard Drive . . . . .	409
C.2	Viewing Free Space Available to EVMS-Managed Devices . . . . .	409
C.3	Making Devices Available to EVMS . . . . .	410
C.3.1	Mounting Your Root File System with EVMS . . . . .	410
C.3.2	Converting LVM Volumes to EVMS Volumes . . . . .	414
C.3.3	Patching the 2.6 Kernel (Unsupported) . . . . .	415
C.4	Additional EVMS Documentation . . . . .	415
<b>D</b>	<b>Comparison of NSS on NetWare and NSS on Linux</b>	<b>417</b>
<b>E</b>	<b>Comparison of NSS for Linux and Linux Traditional File Systems</b>	<b>425</b>
<b>F</b>	<b>Comparison of NSS on NetWare and the NetWare Traditional File System</b>	<b>429</b>
F.1	Features of NSS and Traditional File Systems on NetWare . . . . .	429
F.2	SET Parameters Shared by NSS and Traditional File Systems . . . . .	433
F.3	SET Parameters for the NSS File System . . . . .	434
F.4	SET Parameters for the Traditional File System . . . . .	435
F.5	Comparison of Directory and File Attributes . . . . .	436
<b>G</b>	<b>NSS Nomenclature</b>	<b>439</b>
<b>H</b>	<b>Documentation Updates</b>	<b>441</b>
H.1	April 12, 2006 (Updates) . . . . .	441
H.1.1	Management Tools for NSS . . . . .	441
H.1.2	Coexistence and Migration Issues . . . . .	442
H.1.3	Using EVMS to Manage Devices with NSS Volumes (Linux) . . . . .	442
H.2	April 3, 2006 (Updates) . . . . .	442
H.2.1	Overview of NSS . . . . .	442
H.2.2	Coexistence and Migration Issues . . . . .	443
H.2.3	Conserving Disk Space with Volume, Directory, or User Space Quotas . . . . .	443
H.2.4	Troubleshooting . . . . .	443
H.2.5	Using EVMS to Manage Devices with NSS Volumes (Linux) . . . . .	444
H.3	December 23, 2005 (OES SP2, NetWare 6.5 SP5) . . . . .	444
H.3.1	Management Tools for NSS . . . . .	444
H.3.2	Coexistence and Migration Issues . . . . .	444
H.3.3	Managing Compression on NSS Volumes . . . . .	445
H.3.4	Using Distributed File Services to Move and Split Volumes (NetWare) . . . . .	445
H.3.5	Using Non-EVMS Volume Managers for NSS on Linux . . . . .	445
H.3.6	NSS Utilities . . . . .	446

H.4	November 1, 2005	446
H.5	September 29, 2005 (Updates)	446
H.5.1	About This Guide	446
H.5.2	Management Tools for NSS	446
H.5.3	Coexistence and Migration Issues	447
H.5.4	Managing Encrypted NSS Volumes	447
H.5.5	Conserving Disk Space with Volume, Directory, or User Space Quotas	447
H.5.6	Using NSS Server Console Commands and Utilities	448
H.5.7	Configuring the System Cache to Fine-Tune NSS Performance (Linux)	448
H.5.8	Comparison of NSS on NetWare and NSS on Linux	448
H.6	August 19, 2005 (OES SP1 and NetWare 6.5 SP4)	449
H.6.1	Overview of NSS	449
H.6.2	What's New	450
H.6.3	Upgrading the Media Format for OES SP1 NetWare and NetWare 6.5 SP4	450
H.6.4	Management Tools for NSS	450
H.6.5	Planning for NSS Storage Solutions	450
H.6.6	Coexistence and Migration Issues	451
H.6.7	Managing Partitions	452
H.6.8	Managing Software RAID Devices	452
H.6.9	Managing NSS Pools	452
H.6.10	Managing NSS Volumes	452
H.6.11	Managing Encrypted NSS Volumes	452
H.6.12	Using Volume Attributes that Help Ensure Data Integrity	452
H.6.13	Conserving Disk Space with Volume, Directory, or User Space Quotas	453
H.6.14	Salvaging and Purging Deleted Volumes, Directories, and Files	453
H.6.15	Using Distributed File Services to Move and Split Volumes (NetWare)	453
H.6.16	Managing Hard Links	454
H.6.17	Configuring the System Cache to Fine-Tune NSS Performance (NetWare)	454
H.6.18	Configuring the System Cache to Fine-Tune NSS Performance (Linux)	455
H.6.19	Troubleshooting the NSS File System	455
H.6.20	Security Considerations	455
H.7	June 1, 2005 (Updates)	455
H.7.1	Management Tools for NSS	455
H.7.2	Securing Access to NSS Volumes and Files	455
H.8	March 28, 2005 (Updates)	456
H.8.1	Management Tools for NSS	456
H.8.2	Planning for NSS Storage Solutions	456
H.8.3	Managing Software RAID Devices	457
H.8.4	Managing NSS Volumes	457
H.8.5	Comparison of NSS on NetWare and NSS on Linux	457

# About This Guide

This documentation describes how to use Novell® Storage Services™ File System (NSS) to configure, activate, and maintain large volumes and numerous files without disrupting the work of users. Unless otherwise specified, the information in this guide applies to NSS on both NetWare® and Linux\*.

- Chapter 1, “Overview of NSS,” on page 23
- Chapter 2, “What’s New,” on page 29
- Chapter 3, “Upgrading the Media Format (NetWare),” on page 33
- Chapter 4, “Management Tools for NSS,” on page 39
- Chapter 5, “Planning for NSS Storage Solutions,” on page 65
- Chapter 6, “Coexistence and Migration Issues,” on page 73
- Chapter 7, “Managing Devices,” on page 99
- Chapter 8, “Managing Partitions,” on page 109
- Chapter 9, “Managing Software RAID Devices,” on page 117
- Chapter 10, “Managing Multiple Connection Paths to Devices (NetWare),” on page 139
- Chapter 11, “Managing NSS Pools,” on page 147
- Chapter 12, “Managing NSS Pool Snapshots (NetWare),” on page 167
- Chapter 13, “Managing NSS Volumes,” on page 181
- Chapter 14, “Managing Encrypted NSS Volumes,” on page 201
- Chapter 15, “Securing Access to NSS Volumes, Directories, and Files,” on page 211
- Chapter 16, “Using Volume Attributes that Help Ensure Data Integrity,” on page 217
- Chapter 17, “Managing Compression on NSS Volumes,” on page 227
- Chapter 18, “Conserving Disk Space with Volume, Directory, or User Space Quotas,” on page 243
- Chapter 19, “Salvaging and Purging Deleted Volumes, Directories, and Files,” on page 255
- Chapter 20, “Using Distributed File Services to Move and Split Volumes (NetWare),” on page 261
- Chapter 21, “Managing Hard Links,” on page 281
- Chapter 22, “Using NSS on Devices Managed by Non-EVMS Volume Managers,” on page 291
- Chapter 23, “Using NSS Server Console Commands and Utilities,” on page 299
- Chapter 24, “Monitoring the NSS File System Status,” on page 319
- Chapter 25, “Configuring the System Cache to Fine-Tune NSS Performance (NetWare),” on page 327
- Chapter 26, “Configuring the System Cache to Fine-Tune NSS Performance (Linux),” on page 339
- Chapter 27, “Troubleshooting the NSS File System,” on page 341

- Chapter 28, “Security Considerations,” on page 349
- Appendix A, “NSS Utilities,” on page 359
- Appendix B, “Upgrading Legacy NSS and NetWare Traditional Volumes,” on page 399
- Appendix C, “Using EVMS to Manage Devices with NSS Volumes (Linux),” on page 409
- Appendix D, “Comparison of NSS on NetWare and NSS on Linux,” on page 417
- Appendix E, “Comparison of NSS for Linux and Linux Traditional File Systems,” on page 425
- Appendix F, “Comparison of NSS on NetWare and the NetWare Traditional File System,” on page 429
- Appendix G, “NSS Nomenclature,” on page 439
- “Documentation Updates” on page 441

## Audience

This guide is intended for network administrators.

## Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation, or go to [www.novell.com/documentation/feedback.html](http://www.novell.com/documentation/feedback.html) and enter your comments there.

## Documentation Updates

For the most recent version of the *Novell Storage Services Administration Guide* for OES, see the latest [Novell Open Enterprise Server documentation \(http://www.novell.com/documentation/oes/index.html\)](http://www.novell.com/documentation/oes/index.html).

## Additional Documentation

For more information about documents referenced in this guide, see the following:

- *File Systems Management Guide for OES*
- *OES Server Disks and Storage Devices for NetWare*
- *Novell Archive and Version Services 2.0 for NetWare Administration Guide for OES*
- *Novell Archive and Version Services for NetWare User Guide for OES*
- *NetWare Traditional File System Administration Guide for OES*
- *Novell OES Planning and Implementation Guide* (<http://www.novell.com/documentation/oes/implgde/data/front.html>)
- *OES for NetWare Installation Guide* (<http://www.novell.com/documentation/oes/install-nw/data/hz8pck9v.html>)
- *OES Linux Installation Guide* ([http://www.novell.com/documentation/oes/install\\_linux/data/front.html](http://www.novell.com/documentation/oes/install_linux/data/front.html))
- *Novell Server Consolidation and Migration Toolkit Administration Guide* (<http://www.novell.com/documentation/scmt/index.html>)

- *Novell eDirectory 8.7.3 Administration Guide* (<http://www.novell.com/documentation/edir873/edir873/data/a2iii88.html>)
- *NCP Server for Linux Administration Guide* ([http://www.novell.com/documentation/oes/ncp\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/ncp_lx/data/h9izvdye.html))
- *Novell Linux User Management Technology Guide* (<http://www.novell.com/documentation/oes/lumadgd/data/front.html>)
- Enterprise Volume Management System for Linux
  - EVMS Project at SourceForge.net (<http://sourceforge.net/projects/evms/>)
  - *EVMS User Guide* ([http://evms.sourceforge.net/user\\_guide](http://evms.sourceforge.net/user_guide))
  - *EVMS GUI Screen Shots* ([http://evms.sourceforge.net/gui\\_screen/](http://evms.sourceforge.net/gui_screen/))
- *NDK: Virtual File Services for NetWare* (<http://developer.novell.com/ndk/vfs.htm>)
- *NDK: Novell Storage Architecture Component (Media Manager and NWP)* (<http://developer.novell.com/ndk/storarch.htm>)
- *Novell iManager 2.5 Installation Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html))
- *Novell iManager 2.5 Administration Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_admin\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html))
- *Novell Remote Manager Administration Guide for Linux for OES* ([http://www.novell.com/documentation/oes/remotemgr\\_lx/data/front.html](http://www.novell.com/documentation/oes/remotemgr_lx/data/front.html))
- *Novell Remote Manager for NetWare Administration Guide for OES* (<http://www.novell.com/documentation/oes/remotemgr/data/a7hjvxo.html>)
- *OES NetStorage Administration Guide for Linux* ([http://www.novell.com/documentation/oes/netstor\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/netstor_lx/data/h9izvdye.html))
- *OES NetStorage Administration Guide for NetWare* (<http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html>)
- *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html))
- *OES Novell Cluster Services Administration Guide for NetWare* ([http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html))
- *iSCSI 1.1.3 Administration Guide for NetWare 6.5* ([http://www.novell.com/documentation/iscsi1\\_nak/iscsi/data/h4hgu4hs.html](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html))
- *SBCON Administration Guide* ([http://www.novell.com/documentation/oes/sbcon\\_admin/data/front.html](http://www.novell.com/documentation/oes/sbcon_admin/data/front.html))

## Documentation Conventions

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

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

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux or UNIX\*, should use forward slashes as required by your software.



# Overview of NSS

# 1

Novell® Open Enterprise Server provides Novell Storage Services™ (NSS) File System for both Linux and NetWare® kernels. Its many features and capabilities include visibility, a trustee access control model, multiple simultaneous name space support, native Unicode\*, user and directory quotas, rich file attributes, multiple data stream support, event file lists, and a file salvage subsystem. These features can help you effectively manage your shared file storage for any size of organization, scaling management of the system for even the largest of organizations with hundreds of thousands of employees.

NSS volumes are cross compatible between kernels. You can mount an NSS data volume on either kernel—Linux or NetWare—and move it between them. In a clustered SAN with Novell Cluster Services, volumes can fail over between kernels, allowing for full data and file system feature preservation when migrating data to Linux.

You can manage all storage management functions in the Web-based Novell iManager utility. NSS also supports third-party tools on both kernels for advanced data protection and management, virus scanning, and traditional archive and backup solutions.

- [Section 1.1, “Benefits of NSS,” on page 23](#)
- [Section 1.2, “Understanding NSS,” on page 24](#)
- [Section 1.3, “NSS Features,” on page 25](#)
- [Section 1.4, “What’s Next,” on page 28](#)

## 1.1 Benefits 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.

Businesses today demand more storage space and faster and easier access to data. To meet the demands, you need a file system that can scale to a growing business, is easily maintained, and is better protected against corruption. NSS provides a variety of features that can be combined to provide a robust and reliable solution for your business.

NSS provides the following benefits on Linux and NetWare:

- A journaling file system that lets you create bigger volumes that activate (mount) quicker, store more data, and resist corruption better than non-journaling file systems
- Encrypted volume support to meet the legal standard of making data inaccessible to software that circumvents normal access control, such as if the media were stolen
- An unlimited number of NSS volumes
- Software RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping), RAID 10 (mirroring RAID 0 devices), and RAID 15 (mirroring RAID 5 devices)
- Multiple server activation prevention (MSAP) to help protect pools from being concurrently activated by multiple servers that do not share a cluster relationship
- An NSS volume in a pool of storage that spans multiple storage devices

- Up to 4 billion (10E9) files in a single directory
- Faster access to data, regardless of file size or volume size
- Lower memory requirements: 1 MB of RAM can activate an NSS volume
- Directory space restrictions
- User space restrictions
- Salvage support for deleted volumes and files

NSS provides the following additional benefits on NetWare:

- Multiple I/O path support
- Pool snapshots that capture point-in-time versions of files in the pool
- Transaction Tracking System™ (TTS™)
- Mountable media and image files activated as NSS volumes
- Novell Distributed File Services that allow you to move and split volumes to better manage storage growth
- Novell Archive and Version Services that allow you to archive interval-based versions of user files and make them available to users for individual retrieval. For information, see the *Novell Archive and Version Services 2.0 for NetWare Administration Guide for OES*.

## 1.2 Understanding NSS

NSS is a 64-bit file system that can manage a virtually unlimited number of file objects. On each physical storage device, NSS abstracts up to four physical NetWare partitions to make them appear as contiguous free space. NSS recognizes physical and logical devices up to 2 TB in size from which you can create any number of virtual storage resources, called pools. You can choose space from at least four logical devices of up to 2 TB each to create a pool with a maximum pool size of 8 TB. A pool can contain any number of volumes, depending on the pool size where the minimum volume size is 10 MB. If the pool spans devices by using space from them for the pool, the volumes automatically span the devices. A single volume can contain up to 8 trillion files and grow to 8 TB in size, depending on the size of the pool and space consumed by other volumes in the pool.

### 1.2.1 Storage Pools


During the NetWare installation, NSS automatically creates a system pool (`sys`) and an NSS volume (`sys:`) that can grow to the size of the pool. We recommend that you reserve the system pool and volume for operating system software and extensions. Create additional pools to effectively store applications, files, and databases.

You create additional storage pools by assigning areas of free space obtained from one or more of a server's storage devices. You can create one or more NSS volumes from the space in the storage



pool. The following figure shows how NSS uses free space on multiple devices to create a storage pool.

**Figure 1-1** NSS Pool Architecture



## 1.2.2 NSS Volumes

The logical volumes you create on NSS storage pools are called NSS volumes. You can specify a maximum storage quota for the volume, or allow the volume to grow dynamically to the size of its pool. 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 administrative size of all the volumes taken together is larger than the physical size of the storage pool itself. NSS refers to this process as “overbooking.” If you overbook space in the pool, the individual administrative size of a volume cannot exceed the size of the storage pool.

NSS allocates space from the pools to the volumes only as needed. Typically, user consumption of a volume’s available space ebbs and flows; it is unlikely that users concurrently consume volumes at 100% of their available capacity. Each volume consumes the space it needs as it needs it. By overbooking space, NSS provides a flexible and cost effective way to accommodate expanding storage needs.

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, 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 reach 150 GB, but the other volume stays under 100 GB, your overbooking plan worked.

Suppose you expect one of the volumes might exceed its share of the pool. You can overbook the pool by creating one NSS volume with a quota of 200 GB and a second NSS volume that can grow to the size of the pool. As the combined size nears the size of the pool, you can extend the size of the pool to allow more space for the larger, expanding volume. Your overbooking plan works because you built in the opportunity to expand the pool and volume, according to your business needs.

## 1.3 NSS Features

NSS helps improve the scalability, flexibility, and availability of your storage devices. This section identifies specific NSS features that will help you do the following:

- **Use Less Memory and Gain More Speed**

- **Improve Storage Availability**
- **Prevent Unauthorized Access**
- **Protect Data from Corruption or Loss**
- **Maximize Available Space**

### 1.3.1 Use Less Memory and Gain More Speed

NSS requires only about 1 MB of server memory to activate a volume, independent of the number of files it contains. With NSS, you can activate up to 256 NSS volumes concurrently per server, up to the available server memory.

Whenever you activate an NSS volume, it takes only seconds to mount a volume instead of minutes. NSS uses a journaling file system and does not need to scan the entire file system to create a directory entry table (DET) to load the volume. NSS loads a file's metadata into the memory *only* when you access the file.

NSS reads the file system journal only if server goes down abnormally. Instead of slowly searching the volume for errors, NSS reads the journal to identify any incomplete transactions. It either completes the transaction or backs it out. This results in less server down time and is beneficial for applications such as mail services.

### 1.3.2 Improve Storage Availability

NSS provides three key features to improve I/O performance and provide fault tolerant access to your data:

- Multiple I/O paths support for adapters (NetWare only)

Use multiple channels to connect the server to local media or a storage area network. For information, see **“Managing Multiple Connection Paths to Devices (NetWare)” on page 139**.

- Software RAID support for RAIDs 0, 1, 5, 10, and 15

Use software RAID devices to improve performance and availability. For information, see **“Managing Software RAID Devices” on page 117**.

- Shared-disk storage

Make devices shareable for use in a cluster. For information, see **Section 7.7, “Sharing Devices,” on page 104**.

- Multiple name space support

NSS provides full support for filenames in the Long, UNIX, DOS, and Macintosh\* name spaces. You can specify the default name space to use when mounting the volume. Long name space is recommended.

- Rich file metadata support

NSS provides full support for all file attributes and multiple data streams for DOS, Windows, UNIX, and Macintosh simultaneously.

### 1.3.3 Prevent Unauthorized Access

NSS includes the following features to help prevent access to data that circumvents normal access control:

- Encrypted Volume Support

Encrypt data volumes, meeting U.S. Government security standards. For information, see [“Managing Encrypted NSS Volumes” on page 201](#).

- Data shredding (up to 7 times) for deleted files

Erase files completely, meeting U.S. Government security standards. For information, see [Section 15.2, “Using Data Shredding to Prevent Access to Deleted Files,” on page 215](#).

- Multiple server access prevention for pools

Multiple Server Access Prevention ensures data integrity by preventing unauthorized access to shared media in a storage area network. For information, see [Section 11.14, “Preventing Pools from Activating on Multiple Servers,” on page 164](#).

- Trustee model for access control on NSS volumes

NSS uses the Novell Trustee Model to greatly simplify access control management in the file system. It restricts visibility of data structures so that users only see subdirectories they have rights to see, not the whole tree like all other file systems.

For information about the Novell Trustee Model and NSS file system rights, see the [File Systems Management Guide for OES](#).

Some additional steps are necessary to configure access control for NSS on Linux. For information, see [Section 6.4, “Access Control Issues for NSS on OES Linux,” on page 74](#).

### 1.3.4 Protect Data from Corruption or Loss

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

- Pool snapshots to provide point-in-time views of data (NetWare only)

Back up files from snapshots of data so that all files, including open ones, can be backed up. For information, see [“Managing NSS Pool Snapshots \(NetWare\)” on page 167](#).

- Immediate data saves

On file closure, write data from memory to the file immediately instead of waiting for the next scheduled write to disk. For information, see [Section 16.1, “Using the Flush Files Immediately Attribute to Protect Against Possible Data Corruption on System Failure,” on page 217](#).

- Transaction tracking support for user applications (NetWare only)

Monitor application transactions and roll back uncompleted application transactions after a system failure. For information, see [Section 16.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback \(NetWare\),” on page 219](#).

- File-level snapshots to allow backup of files that are not open exclusively

Back up files from snapshots of data so that all files, including open ones, can be backed up. For information, see [Section 16.3, “Using the File-Level Snapshot Attribute to Enable the Backup of Open Files,” on page 220](#).

- Salvage file subsystem

Recover files, directories, and volumes that were deleted by applications or from the terminal/console commands. For information, see [“Salvaging and Purging Deleted Volumes, Directories, and Files” on page 255](#).

## 1.3.5 Maximize Available Space

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

- File compression

Compress inactive files, according to preset parameters, to conserve space in a volume. For information, see [“Managing Compression on NSS Volumes” on page 227](#).

- Volume space restrictions

Limit the amount of space a volume can consume in its pool. For information, see [Section 18.2, “Managing NSS Volume Quotas,” on page 245](#).

- Directory space restrictions

Limit the amount of space a subdirectory can consume, regardless of broader volume and user constraints. For information, see [Section 18.3, “Managing Directory Quotas,” on page 246](#).

- User space restrictions

Limit the amount of space a user’s data can consume, regardless of broader directory or volume constraints. For information, see [Section 18.4, “Managing User Space Quotas,” on page 250](#).

## 1.4 What’s Next

See [Chapter 2, “What’s New,” on page 29](#) to learn about new and modified features in this release of NSS.

Review the following sections before designing and implementing your storage solution:

- [Chapter 3, “Upgrading the Media Format \(NetWare\),” on page 33](#)
- [Chapter 4, “Management Tools for NSS,” on page 39](#)
- [Chapter 5, “Planning for NSS Storage Solutions,” on page 65](#)
- [Chapter 6, “Coexistence and Migration Issues,” on page 73](#)

This section discusses additions and modifications to the Novell® Storage Services™ File System for Novell Open Enterprise Server.

- [Section 2.1, “What’s New \(OES SP2 and NetWare 6.5 SP5\),” on page 29](#)
- [Section 2.2, “What’s New \(OES SP1 and NetWare 6.5 SP4\),” on page 29](#)
- [Section 2.3, “What’s New \(OES and NetWare 6.5 SP3\),” on page 30](#)

## 2.1 What’s New (OES SP2 and NetWare 6.5 SP5)

There are no new features for NSS in this release.

## 2.2 What’s New (OES SP1 and NetWare 6.5 SP4)

This section discusses additions and modifications to NSS for OES SP1 and to NSS for NetWare 6.5 SP4.

- [Section 2.2.1, “Linux,” on page 29](#)
- [Section 2.2.2, “NetWare,” on page 30](#)

### 2.2.1 Linux

The following NSS features and capabilities have been added or modified for OES Linux:

- Encrypted volume support. For information, see [“Managing Encrypted NSS Volumes” on page 201](#).
- User space quota support. For information, see [Section 18.4, “Managing User Space Quotas,” on page 250](#).
- RAID 5 support. For information, see [Section 9.5, “Creating a Software RAID Device Using iManager,” on page 126](#).
- Hard links support. For information, see [“Managing Hard Links” on page 281](#).
- Lookup Name Space. Mount an NSS volume in any of the names paces, using NSSMU or the Storage plug-in to iManager 2.5. Long name space is highly recommended. For information, see [Section 13.2.8, “Lookup Name Space,” on page 184](#).
- Metadata Migration (METAMIG, `metamig`) utility. For information, see [Section A.8, “METAMIG \(Linux\),” on page 374](#).
- Direct partitions format upgrade for NetWare 6.0 NSS volumes moved to an OES SP1 Linux server. For information, see [Section 6.8, “Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later,” on page 92](#).
- Performance improvements such as cache integration and NCP Server integration.
- Storage plug-in requires Novell iManager 2.5. For OES SP1, the Storage plug-in has been redesigned. To take advantage of the changes, you must upgrade from iManager 2.0.2 to iManager 2.5. For information, see [“Upgrading to iManager 2.5” \(http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/ah8bnjv.html\)](http://www.novell.com/documentation/imanager25/imanager_install_25/data/ah8bnjv.html) in the *Novell iManager*

[2.5 Installation Guide \(http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html\)](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html).

- Opportunistic locking (oplocks) support
- RAVSUI and RAVVIEW man pages

## 2.2.2 NetWare

The following NSS features and capabilities have been added or modified for OES NetWare:

- New Media Format. A new metadata structure (Beast Version 3) supports enhanced hard link capability. You must upgrade the media format after upgrading the server to OES SP1 NetWare or NetWare 6.5 SP4. Some restrictions apply. For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33](#).
- DFS junctions can point to NSS volumes on Linux.
- Purge and salvage deleted directories in Novell Remote Manager. For information, see [Section 19.4, “Salvaging or Purging Deleted Directories from the File System Salvage Area \(OES SP1 NetWare, NetWare 6.5 SP4, or Later\),” on page 257](#).
- Hard links support. Enhanced hard links support provides the following:
  - A Hard Links attribute for a volume allows the administrator to enable or disable the ability to create hard links for it. For information, see [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 285](#).
  - A new metadata structure that allows up to 65,535 hard links per volume. For information, see [Section 21.1.9, “New Metadata Structure Supports Up to 65,535 Hard Links for a File \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 284](#).
- Storage plug-in requires Novell iManager 2.5. For OES SP1, the Storage plug-in has been redesigned. To take advantage of the changes, you must upgrade from iManager 2.0.2 to iManager 2.5. For information, see [“Upgrading to iManager 2.5” \(http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/ah8bnjv.html\)](#) in the *Novell iManager 2.5 Installation Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html)).

## 2.3 What’s New (OES and NetWare 6.5 SP3)

This section discusses additions and modifications to NSS for OES.

- [Section 2.3.1, “Linux,” on page 30](#)
- [Section 2.3.2, “NetWare,” on page 32](#)

### 2.3.1 Linux

NSS now supports Linux in addition to NetWare®. Review the following:

- [“Coexistence and Migration Issues” on page 31](#)
- [“Feature Comparison” on page 31](#)
- [“Install of NSS on Linux” on page 31](#)
- [“NSS Configuration and Management for Linux” on page 31](#)
- [“NSS Management Utility for Linux” on page 31](#)

- “NSS Console for Linux” on page 31
- “RIGHTS Utility for Linux” on page 31
- “ATTRIB Utility for Linux” on page 32
- “RAVSUI Utility for Linux” on page 32
- “RAVVIEW Utility for Linux” on page 32

## Coexistence and Migration Issues

Most of the NSS features and capabilities on Linux parallel those that are available on NetWare, but there are some differences. For information, see [Chapter 6, “Coexistence and Migration Issues,”](#) on [page 73](#).

## Feature Comparison

For a comparison of features and capabilities, see [“Comparison of NSS on NetWare and NSS on Linux”](#) on [page 417](#).

## Install of NSS on Linux

For information on installing NSS on Linux, see [“Installing Linux with EVMS as the Volume Manager of the System Device”](#) in the *OES Linux Installation Guide*.

To configure the system disk for EVMS by patching the kernel after installation, see [Section 7.2, “Making Devices Available to EVMS \(Linux\),”](#) on [page 100](#).

You can also install NSS after the OES Linux install by selecting NSS from the software options in YaST.

## NSS Configuration and Management for Linux

The configuration and management of NSS on Linux are the same as for NetWare, except as noted in specific tasks and procedures described herein. If a feature or procedure is specific to only one platform, the heading identifies the platform.

## NSS Management Utility for Linux

The NSS Management Utility (NSSMU, `nssmu`) for Linux is available for managing NSS data volumes after you have installed OES, but it is not available during the install. It is similar to NSSMU for NetWare. You can load NSSMU from the command line of the server console. For information, see [Section 4.2.1, “NSSMU for Linux,”](#) on [page 53](#).

## NSS Console for Linux

The NSS Console (NSSCON, `nsscon`) for Linux is a utility that provides an interface for issuing NSS commands in a terminal window. You can load `nsscon` from the command line of the server console. For information, see [Section 4.3.3, “NSS Console for OES Linux,”](#) on [page 57](#).

## RIGHTS Utility for Linux

The RIGHTS (`rights`) utility for Linux is available for managing file system trustees for directories and files in NSS volumes on Linux. It is similar to the RIGHTS utility for NetWare. You

can load RIGHTS from the command line of the server console. For information, see [in the Section A.17.1, “RIGHTS \(Linux\),” on page 383.](#)

### **ATTRIB Utility for Linux**

The ATTRIB (`attrib`) utility for Linux is available for managing file system attributes directories and files in NSS volumes on Linux. It is similar to the FLAG utility for NetWare. You can load ATTRIB from the command line of the server console. For information, see [Section A.2, “ATTRIB \(Linux\),” on page 360.](#)

### **RAVSUI Utility for Linux**

The Rebuild and Verify Simple User Interface (RAVSUI, `ravsui`) utility allows you to rebuild and verify NSS pools on Linux. For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378.](#)

### **RAVVIEW Utility for Linux**

The Rebuild and Verify View (RAVVIEW, `ravview`) utility allows you to reformat the RTF, RTFN, VBF, and VBFN files of the rebuild and verify process for NSS pools on Linux to human-readable format. For information, see [Section A.14, “RAVVIEW \(Linux\),” on page 380.](#)

## **2.3.2 NetWare**

The following NSS features and capabilities have been added or modified for OES NetWare:

- [“Encrypted Volume Support on OES NetWare” on page 32](#)
- [“Multiple Device Selection in NSSMU for OES NetWare” on page 32](#)
- [“CD and DVD Image Files as NSS Volumes” on page 32](#)

### **Encrypted Volume Support on OES NetWare**

NSS Encrypted Volume Support on OES NetWare meets the legal standard of making data inaccessible to software that circumvents normal access control, such as if the media were stolen. On OES NetWare, you can encrypt newly created NSS data volumes, but you cannot encrypt the `sys :` volume. EVS is not available on Linux in this release.

For information, see [“Managing Encrypted NSS Volumes” on page 201.](#)

### **Multiple Device Selection in NSSMU for OES NetWare**

NSSMU for NetWare allows you to select one device or multiple devices for Initialize or Share actions. For information, see [Section 4.2, “NSS Management Utility,” on page 52.](#)

### **CD and DVD Image Files as NSS Volumes**

The `cddvd.nss` module allows you to activate CD and DVD image files as NSS volumes. For information, see [Section 13.12, “Activating and Deactivating CD and DVD Image Files as NSS Volumes \(NetWare\),” on page 194.](#)



# Upgrading the Media Format (NetWare)

# 3

For OES SP1 NetWare and later and for NetWare 6.5 SP4 and later, a new metadata structure (Beast Version 3) provides improved support for hard links. After you upgrade your operating system, you must separately upgrade the media format for your NSS volumes to use the new metadata structure; some restrictions apply.

- [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33](#)
- [Section 3.2, “Upgrading the Media Format for NSS Volumes on NetWare,” on page 35](#)

## 3.1 Guidelines for Upgrading the Media Format of NSS Volumes on NetWare

Before upgrading the media format of your NSS volumes on NetWare, make sure you understand the following guidelines:

- [Section 3.1.1, “Support for the NSS Metadata Structure Upgrade,” on page 33](#)
- [Section 3.1.2, “Before Upgrading the Media Format,” on page 34](#)
- [Section 3.1.3, “Which NSS Volumes to Upgrade,” on page 34](#)
- [Section 3.1.4, “After Upgrading the Media Format,” on page 35](#)
- [Section 3.1.5, “If You Do Not Upgrade the Media Format,” on page 35](#)

### 3.1.1 Support for the NSS Metadata Structure Upgrade

The new metadata structure (Beast Version 3) is available only for OES SP1 NetWare and later and NetWare 6.5 SP4 and later, as shown in the following table:

**Table 3-1** *Support for the NSS Metadata Structure Upgrade*

Operating System	NSS Metadata Upgrade
OES NetWare SP1 and later	Supported
NetWare 6.5 SP4 and later	
OES Linux and later	Not supported
OES NetWare (initial release)	Not supported
NetWare 6.5 SP3 and earlier	

### 3.1.2 Before Upgrading the Media Format

When you upgrade the operating system, NSS does not automatically upgrade the media format to use the new metadata structure because it might create system conflicts for system roll backs, shared volumes, migrating volumes cross-platform to OES SP2 Linux and earlier, or migrating devices to earlier releases of NetWare.

Before you upgrade the media format of NSS volumes to the new data structure, make sure the server is performing as expected with the upgrade to OES SP1 NetWare or later or to NetWare 6.5 SP4 or later. This allows you to roll back to the previous NetWare release if necessary.

---

**WARNING:** After the media format is upgraded to the new metadata structure, you cannot roll back to a previous release.

---

If you attempt to upgrade the media format for a shared NSS volume on a cluster with Novell Cluster Services™, the upgrade is refused until all servers on the system are configured with a supported operating system. Make sure to upgrade all cluster nodes to NetWare 6.5 SP4 or OES SP1 NetWare or later before attempting to upgrade the shared volume.

Until you upgrade the structure for an NSS volume, the volume's format uses the same structure (Beast Version 2) as is used on earlier NetWare releases and on OES SP2 Linux and earlier. Any existing hard links on your system's NSS volumes are visible, and they can be opened, closed, read, and written. However, until you upgrade the NSS volume to the new structure, you cannot create new hard links, and you cannot rename or delete existing hard links. Attempts to do so are rejected with an appropriate error.

### 3.1.3 Which NSS Volumes to Upgrade

With a few exceptions as noted below, it is highly recommended that you upgrade the NSS volume to the new metadata structure after you upgrade the operating system to OES SP1 NetWare or later or to NetWare 6.5 SP4 or later. The next major (non support pack) release of NSS will force the upgrade anyway.

Do not upgrade the media format of the NSS volume to the new metadata structure if any of the following conditions exist:

- You have not yet verified that the performance of your system after upgrading the operating system is as expected, and you might need to roll back to an earlier NetWare release.
- You plan to migrate one or more devices containing the NSS volume to an OES Linux server after upgrading to OES SP1 NetWare. The new metadata structure is not supported on OES SP2 Linux and earlier.
- You need to share this volume with an OES SP2 Linux or earlier server in a mixed cluster with Novell Cluster Services. The cluster software prevents the media upgrade unless all operating systems in the cluster support the new media format.
- You have not yet upgraded all nodes in a NetWare cluster to OES SP1 NetWare or later or to NetWare 6.5 SP4 or later. The cluster software prevents the media upgrade until all operating systems in the cluster support the new media format.

### 3.1.4 After Upgrading the Media Format

After you upgrade the volume to use the new metadata structure, the volume supports only the new media format. The following constraints apply for its use:

- The upgraded NSS volume cannot be rolled back to use the old media format.
- You cannot roll back the operating system to a previous NetWare version.
- You cannot migrate a device containing the upgraded volume to a system with OES SP2 Linux and earlier or to an earlier unsupported NetWare release.
- Upgraded shared volumes cannot be used in clusters with nodes that use unsupported operating systems, such as OES NetWare, NetWare 6.5 SP3 and earlier, and OES SP2 Linux and earlier. Upgraded shared NSS volumes cannot be used in a mixed cluster with Novell Cluster Services.
- You must enable the Hard Links attribute for your NSS volume on NetWare before you can create hard links. When you upgrade the NSS volume to use the new media format, if any old-style hard links are detected, the Hard Links attribute is automatically enabled. Otherwise, the volume is upgraded, but the attribute is disabled and must be enabled before you can create hard links. For information, see [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 285.](#)

### 3.1.5 If You Do Not Upgrade the Media Format

If you do not upgrade the media format for an NSS volume on an OES SP1 NetWare or later server or on a NetWare 6.5 SP4 or later server, the volume’s format uses the same metadata structure (Beast Version 2) as is used on earlier NetWare releases and on OES SP1 Linux and earlier. Any existing hard links on your system’s NSS volumes are visible, and they can be opened, closed, read, and written. However, until you upgrade the NSS volume to the new structure, you cannot create new hard links, and you cannot rename or delete existing hard links. Attempts to do so are rejected with an appropriate error.

If the non-upgraded NSS volume is shared in a mixed cluster with Novell Cluster Services, hard links can be created, renamed, or deleted by first mounting the volume on a node in the cluster with an operating system that is compatible with the old media format.

If desired, devices containing non-upgraded NSS volumes can be moved cross-platform to servers with operating systems compatible with the old media format. For information about moving media cross-platform, see [Section 6.7, “Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux,” on page 87.](#)

## 3.2 Upgrading the Media Format for NSS Volumes on NetWare

Before you upgrade an NSS volume to use the new metadata structure (Beast Version 3), make sure you understand the [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33.](#) This section describes how to upgrade to your NSS volume to use the new metadata structure available for NSS volumes on OES SP1 NetWare and later or on NetWare 6.5 SP4 and later.

- [Section 3.2.1, “Upgrading the Structure for Currently Active NSS Volumes,” on page 36](#)
- [Section 3.2.2, “Upgrading the Structure for New NSS Pools and Volumes,” on page 36](#)

### 3.2.1 Upgrading the Structure for Currently Active NSS Volumes

The following commands perform the one-time upgrade for currently active volumes and are permissible only when used at the server console. The process runs in the background. The elapsed time for the upgrade process varies according to the number of storage objects in the volume; the more objects a volume contains, the longer it takes. Volumes with a large number of objects can take a substantial time to upgrade. You cannot enable the Hard Links attribute for the volume until the media upgrade is complete. A message is sent to the logger screen when the process is successfully completed.

#### Syntax

---

Command and Description
-------------------------

---

<pre>nss /ZLSSUpgradeCurrentVolumeMediaFormat=<i>volumentname</i></pre>
---

Upgrades the specified active volume. If it is a shared volume, the clustering system vetoes the upgrade until all nodes in the cluster are running a supported operating system.

<pre>nss /ZLSSUpgradeCurrentVolumeMediaFormat=all</pre>
---

Upgrades all currently active volumes. If a shared volume is encountered, the clustering system vetoes the upgrade until all nodes in the cluster are running a supported operating system.

<pre>nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=shared</pre>
---

Upgrades all currently active volumes that are part of a shared pool, but only if the clustering system permits it.

<pre>nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=local</pre>
--

Upgrades all currently active volumes that are not part of a shared pool.

---

### 3.2.2 Upgrading the Structure for New NSS Pools and Volumes

The following commands determine the media version type for new pools and volumes. Make sure to use the commands on the NetWare servers where you plan to use hard links. Issue the command from the server console or from the `nssstart.cfg` file. If the command is issued from the command line, it persists until a server reboot. If the command is placed in the `nssstart.cfg` file, it persists across server reboots.

#### Syntax

---

Command and Description
-------------------------

---

<pre>nss /ZLSSUpgradeNewVolumeMediaFormat</pre>
---

Creates all new NSS volumes with the upgraded structure. This applies to volumes on both local and shared pools.

---

---

**Command and Description**

---

`nss /ZLSSUpgradeNewVolumeMediaFormat /include=shared`

Creates all new NSS volumes with the upgraded structure if the new volume is on a shared pool. However, the clustering system vetoes the upgrade until all nodes in the cluster are running a supported operating system.

`nss /ZLSSUpgradeNewVolumeMediaFormat /include=local`

Creates all new NSS volumes with the upgraded structure if the new volume is not on a shared pool.

---



# Management Tools for NSS

# 4

This section identifies the various tools for managing your Novell® Storage Services™ file system:

- Section 4.1, “Novell iManager Storage Plug-In,” on page 39
- Section 4.2, “NSS Management Utility,” on page 52
- Section 4.3, “NSS Commands and Utilities,” on page 56
- Section 4.4, “Novell NetStorage,” on page 58
- Section 4.5, “Novell Remote Manager,” on page 60
- Section 4.6, “Novell Client,” on page 62
- Section 4.7, “ConsoleOne,” on page 62
- Section 4.8, “Virtual File Services, APIs, and Scripts,” on page 63

---

**NOTE:** Legacy NetWare® tools, such as NWADMIN and NWCONFIG, might not be compatible with some of the NSS features.

---

NSS also supports third-party tools on both kernels for advanced data protection and management, virus scanning, and traditional archive and backup solutions.

## 4.1 Novell iManager Storage Plug-In

Novell iManager 2.5 is a Web browser-based tool used for configuring, managing, and administering Novell eDirectory™ objects on your network. The Storage plug-in is the primary tool used to manage NSS devices, software RAIDs, pools, and volumes.

Novell iManager gives you the ability to assign specific tasks or responsibilities to user accounts 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 2.5 Administration Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_admin\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html)).

---

**NOTE:** The Storage plug-in does not support Mobile iManager.

---

This section discusses the following:

- Section 4.1.1, “Prerequisites for Using the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager 2.5,” on page 40
- Section 4.1.2, “Downloading and Installing Plug-in Updates,” on page 42
- Section 4.1.3, “Accessing Novell iManager,” on page 45
- Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46
- Section 4.1.5, “Selecting a Server to Manage,” on page 46
- Section 4.1.6, “Understanding iManager Storage Management Tasks,” on page 47

### 4.1.1 Prerequisites for Using the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager 2.5

The requirements in this section apply to the following storage-related plug-ins for iManager 2.5:

- Storage plug-in
- Archive Versioning plug-in
- File Protocols plug-in (NetWare)
- Clustering plug-in

---

**IMPORTANT:** The Storage, Archive Versioning, File Protocols (NetWare), and Cluster Services plug-ins share common code. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins. If you remove only one of the plug-ins, it removes the common code and breaks the remaining installed plug-ins.

---

#### Upgrading to Novell iManager 2.5

For OES SP1, the storage-related plug-ins have been redesigned. To take advantage of the changes, you must upgrade from iManager 2.0.2 to iManager 2.5. For information, see “Upgrading to iManager 2.5” ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/ah8bnjv.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/ah8bnjv.html)) in the *Novell iManager 2.5 Installation Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html)).

#### EVMS

For OES Linux, the Storage plug-in requires that EVMS be the volume manager of devices that contribute space to NSS pools and volumes. It does not recognize devices managed by non-EVMS volume managers.

#### Web Browser Language Setting

The iManager plug-in might not operate properly if the highest priority Language setting for your Web browser is set to a language other than one of the supported languages. To avoid problems, in your Web browser, click *Tools > Options > Languages*, and then set the first language preference in the list to a supported language.

#### Protocols

The following table provides information about the protocols needed to use iManager to manage storage in a heterogeneous environment. A protocol annotated with an asterisk is the default and is configured automatically on the servers. The protocols you use must be loaded and running on both the iManager server and the target server you want to manage.



**Table 4-1** Interoperability of Protocols for iManager Server and Target Servers

iManager 2.5 or later Server	Target Server			
	OES Linux and later	OES NetWare and later; NetWare 6.5 SP4 and later	NetWare 6.5 SP3	NetWare 6.5 SP2
OES Linux and later	* WBEM	* WBEM	WBEM (Start WBEM.)	
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
OES NetWare and later; NetWare 6.5 SP4 and later	* WBEM	* WBEM	WBEM (Start WBEM.)	* NCP
	CIFS	CIFS	* NCP CIFS	CIFS (Field Patch 2B)
NetWare 6.5 SP3	* WBEM	* WBEM	WBEM (Start WBEM.)	* NCP
	CIFS	CIFS	* NCP CIFS	CIFS (Field Patch 2B)
NetWare 6.5 SP2	Not available	* NCP	* NCP	* NCP

## WBEM

Where WBEM is the default protocol, WBEM is loaded and runs automatically when you start the server. Otherwise, you must start WBEM to use the protocol.

To start WBEM:

- 1 At the server console, enter

```
openwbem
```

If you receive file protocol errors, it might be because WBEM is not running.

For information about installing WBEM, see “Setting Up OpenWBEM” (<http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html>) in the *OpenWBEM Services Administration Guide for OES* (<http://www.novell.com/documentation/oes/cimom/data/front.html>).

## CIFS

Where it is available, CIFS must be configured before you can use it. An additional CIFS setup requirement for Field Patch 2B is noted where it is required. For information, see the *OES Native File Access Protocols Guide* (<http://www.novell.com/documentation/oes/native/data/h9izvdye.html>).

## NCP

NCP™ is the default protocol when the iManager server and target server is NetWare 6.5 SP3 or SP2.

## Stop EVMSGUI When Working in iManager or NSSMU

You cannot run the Enterprise Volume Management System Graphical User Interface (EVMSGUI) at the same time that iManager or NSSMU is running. If you use EVMSGUI (`evmsgui` command) to manage NSS devices, close any instance of iManager or NSSMU, and vice versa.

### 4.1.2 Downloading and Installing Plug-in Updates

This section contains the instructions for downloading and installing updates to the Storage plug-in for iManager 2.x. However, it can be generalized and applied to updates for the following:

- **NSS (`nssmgmt.npm`):** Contains the Storage plug-in for managing NSS on NetWare 6.5 Support Pack 3, OES NetWare, and OES Linux.

This file also contains the File Protocols plug-in, which is used to manage CIFS services for end users of data stored on your NetWare 6.0 and later servers. For information about CIFS, see the *OES Native File Access Protocols Guide* (<http://www.novell.com/documentation/oes/native/data/h9izvdye.html>).

- **Archive Versioning (`arkmgmt.npm`):** Contains the Archive Versioning plug-in for managing Novell Archive and Version Services for NetWare.

For information, see the *Novell Archive and Version Services 2.0 for NetWare Administration Guide for OES*.

- **Cluster Services (`ncsmgmt.npm`):** Contains the Clustering plug-in for managing Novell Cluster Services™.

For information, see the following:

- *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html))
- *OES Novell Cluster Services Administration Guide for NetWare* ([http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html))

#### Downloading the Plug-In

- 1 Click the *Download* link, log in, and review the download agreement.

If you do not have a Novell account, you can register for free.

- 2 On the system where Novell iManager is installed, download the NPM file.

---

**NOTE:** When downloading this iManager plug-in, the browser might rename the downloaded file by adding a `.zip` extension (i.e. `pluginname.npm.zip`). If this occurs, you can either install the `.zip` file in iManager or you can rename the file and install it. Any plug-in with the additional `.zip` extension can still be installed with iManager without any additional steps.

---

- 3 You can leave the NPM in the downloaded location, or move it to the iManager packages location.

## Installing the Plug-In for iManager 2.5 When RBS Is Not Configured

If you do not have Role-based Services (RBS) configured for Novell eDirectory, install the NPM file as follows:

- 1 Download the NPM update file.

For information, see [“Downloading the Plug-In” on page 42](#).

- 2 In a Web browser, log in to iManager on the system where you downloaded the file.
- 3 Click the *Configure* icon (person seated behind a desk).
- 4 Under *Module Installation*, click *Available Novell Plug-In Modules*.
- 5 If the Storage and File Protocols plug-ins (nssmgt . npm file) do not appear in the list, click *New*, browse to select the nssmgt . npm file, then click *OK*.

You are returned to the Available Novell Plug-in Modules page.

- 6 Select the Storage and File Protocols plug-ins (nssmgt . npm file) from the list of available modules, then click *Install*.

This install takes a few minutes. You should receive a message confirming a successful install.

- 7 Click *OK* to dismiss the message, then close iManager.
- 8 Restart Tomcat 4, using one of the following methods:

Platform	Restart Command
NetWare 6.5 or later	Enter
	TC4STOP
	Wait at least one minute, then enter
	TOMCAT4
OES Linux	Enter
	/etc/init.d/novell-tomcat4 stop
	Wait at least one minute, then enter
	/etc/init.d/novell-tomcat4 start
Windows*	Stop and start the Tomcat service.
Solaris*	Enter
	/etc/init.d/imgr stop
	Wait at least one minute, then enter
	/etc/init.d/imgr start
HP*-UX*	Enter
	/opt/hpws/tomcat/bin/shutdown.sh
	Wait at least one minute, then enter
	/opt/hpws/tomcat/bin/startup.sh

- 9 Tomcat sometimes requires several minutes to fully initialize. Wait at least 5 minutes before trying to log into iManager.

For more information, see “Downloading and Installing an NPM” ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/bnpta1r.html#bnpta1r](http://www.novell.com/documentation/imanager25/imanager_install_25/data/bnpta1r.html#bnpta1r)) in the *Novell iManager 2.5 Installation Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html)).

## Installing a Plug-In When RBS Is Configured

If you are running iManager in Assigned Mode and have RBS configured for eDirectory, complete the following steps to install the NPM file.

---

**IMPORTANT:** To re-install an existing plug-in, you must first delete the rbsModule object for that plug-in from eDirectory, using the *Module Configuration > Delete RBS Module* task.

---

- 1 Download the NPM update file.

For information, see “Downloading the Plug-In” on page 42.

- 2 In a Web browser, log in to iManager as an RBS Collection Owner on the system where you downloaded the NPM file.
- 3 Click the *Configure* icon (person seated behind a desk).
- 4 Under *Module Installation*, click *Available Novell Plug-In Modules*.
- 5 If the Storage and File Protocols plug-ins (nssmgt . npm file) do not appear in the list, click *New*, browse to select the nssmgt . npm file, then click *OK*.  
You are returned to the Available Novell Plug-in Modules page.

- 6 Select the Storage and File Protocols plug-ins (nssmgt . npm file) from the list of available modules, then click *Install*.

This install takes a few minutes. You should receive a message confirming a successful install.

- 7 Click *OK* to dismiss the message, then close iManager.
- 8 Restart Tomcat 4, using one of the following methods:

Platform	Restart Command
NetWare 6.5 or later	Enter
	TC4STOP
	Wait at least one minute, then enter TOMCAT4
OES Linux	Enter
	/etc/init.d/novell-tomcat4 stop
	Wait at least one minute, then enter /etc/init.d/novell-tomcat4 start
Windows*	Stop and start the Tomcat service.

Platform	Restart Command
Solaris*	<p>Enter</p> <pre>/etc/init.d/imgr stop</pre> <p>Wait at least one minute, then enter</p> <pre>/etc/init.d/imgr start</pre>
HP*-UX*	<p>Enter</p> <pre>/opt/hpws/tomcat/bin/shutdown.sh</pre> <p>Wait at least one minute, then enter</p> <pre>/opt/hpws/tomcat/bin/startup.sh</pre>

- 9 Tomcat sometimes requires several minutes to fully initialize. Wait at least 5 minutes before trying to log into iManager.
- 10 After Tomcat initializes, in a Web browser, log in to iManager as a Collection Owner again.
- 11 Click the *Configure* icon.
- 12 Under *Role-Based Services*, select *RBS Configuration*.  
The table on the *2.x Collections* tabbed page displays modules ready to update.
- 13 Locate the collection where you want to install the plug-in, then click its *Out-of-Date* number.  
The Storage and File Protocols plug-ins (*nssmgmt . npm*) should be displayed under *Modules Not Yet Installed* column.
- 14 Select the plug-ins.
- 15 Click *Update*.
- 16 Wait for the *Completed* message, then click *OK* to continue.  
The *Storage* role and *File Protocols* role should appear on the *Roles and Tasks* pane.

### 4.1.3 Accessing Novell iManager

- 1 Launch a Web browser.
- 2 Click *File > Open*, then enter  

```
https://server-IP-address/nps/iManager.html
```

Replace *server-IP-address* with the actual server DNS name or IP address. For example:  

```
https://192.168.1.1/nps/iManager.html
```

The iManager Login page opens.
- 3 Log in, using your administrator username and password, to the eDirectory tree that contains the server you want to manage.  


In Novell iManager, you can access only the roles and tasks you are authorized to manage. For full access to all available Novell iManager features, you must log in as Supervisor of the tree.

## 4.1.4 Accessing NSS Storage Management Role and Tasks

- 1 Access iManager, then log in to the eDirectory tree where the server you want to manage resides.

For information, see [Section 4.1.3, “Accessing Novell iManager,”](#) on page 45.

- 2 In *Roles and Tasks*, click *Storage*.



The Storage role expands to reveal its main tasks, and opens to the *Devices* task.

As you work in the Storage plug-in, use the navigation links at the top of the page, referred to as “breadcrumbs,” to return to pages you recently visited, or use the links in *Roles and Tasks*. If you use the *Refresh* and *Back* features of your Web browser to navigate, iManager returns you to the initial page you encountered after login.

- 3 To activate the options on the selected page, select a server to manage.

For information, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.

- 4 Continue with Device management, or click one of the other tasks to access its associated management options:



- *Pools*
- *Volumes*
- *Volume Job Control* (NetWare, requires Distributed File Services)
- *Partitions*
- *Software RAIDS*
- *Devices*
- *Scan for Devices*

For information about Storage tasks, see [Section 4.1.6, “Understanding iManager Storage Management Tasks,”](#) on page 47.

## 4.1.5 Selecting a Server to Manage

Before you can access the management options on a selected Storage task page, you must select a server to manage that is in the same Novell eDirectory tree where you are currently logged in.

- 1 Use one of the following methods to select a server in the tree where you are logged in:

Server:   

- Type the Novell eDirectory distinguished server name for the server you want to manage, then click *Apply*. For example: `svr1.novell`.

- Click the *Search* icon to open the eDirectory Object Selector. Browse or search the list to locate the server you want to manage, then click the server name.
  - Click the *Object History* icon to select a server you have recently managed.
- 2 Wait for iManager to retrieve information about that server and display the appropriate information to the task page you are in.

It might take several seconds to retrieve the information, depending on the size and complexity of your storage solution.

## 4.1.6 Understanding iManager Storage Management Tasks

The NSS 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 users, you can create virtual devices, storage pools, and volumes that abstract the physical storage devices in a mixed-media environment.

Use the Novell iManager Storage plug-in to manage the NSS file system for your server. In iManager, the Storage Management role comprises six key tasks:

- “Pools” on page 47
- “Volumes” on page 48
- “Volume Job Control (NetWare)” on page 49
- “Partitions” on page 50
- “Software RAIDs” on page 50
- “Devices” on page 51
- “Scan for Devices” on page 52

### Pools

You can create and manage storage pools to efficiently use all free space. For NetWare, you can also enable the Pool Snapshot feature to preserve point-in-time views of data pools and to support data recovery and backup.

Use the following table to find information about subtasks available from the Pools task.

**Table 4-2** *Explanation of Pool Management Tasks*

Subtask	Description	Reference
<i>Pools</i>	A list of all pools on the selected server.	<a href="#">Viewing Pools on a Server</a>
<i>Details</i>	View information about a selected pool.	<a href="#">Viewing Pool Details</a>
<i>New</i>	Create a new pool on the selected server.	<a href="#">Creating Pools with iManager</a>
<i>Delete</i>	Delete a selected pool and all of its volumes and their data.	<a href="#">Deleting a Pool</a>
<i>Rename</i>	Rename a selected pool.	<a href="#">Renaming a Pool</a>

Subtask	Description	Reference
<i>Activate</i>	Mount and activate a selected deactive or unmounted pool.	<a href="#">Activating and Deactivating Pools</a>
<i>Deactivate</i>	Deactivate a selected active pool.	<a href="#">Activating and Deactivating Pools</a>
<i>Increase Size</i>	Add more partitions from available devices to expand the size of a pool.	<a href="#">Increasing the Size of a Pool</a>
<i>Snapshot</i>	Opens the Pool Snapshot Management page where you can create and manage pool snapshots.	<a href="#">Managing NSS Pool Snapshots (NetWare)</a>
<i>Update eDirectory</i>	If a storage object is not recognized automatically, update eDirectory objects to force recognition.	<a href="#">Updating eDirectory Pool Objects</a>
<i>Deleted Volume</i>	View and salvage deleted volumes in a pool.	<a href="#">Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool</a>
<i>Offline</i>	For a pool snapshot that is online as an active pool, take it offline. This does not delete the pool snapshot.	<a href="#">Moving an Active Pool Snapshot Offline from the Pools Page</a>
<i>Partitions</i>	View a list of the partitions comprising the pool's storage space.	<a href="#">Viewing Partition Information for a Pool</a>
<i>Volumes</i>	View a list of all volumes on a selected pool.	<a href="#">Viewing Volume Information for a Pool</a>
<i>Devices</i>	View a list of the devices that contribute space to a selected pool.	<a href="#">Viewing Device Information for a Pool</a>

## Volumes

You can create and manage NSS volumes, including their key attributes.

Use the following table to find information about subtasks available from the Volumes task.

**Table 4-3** *Explanation of Volume Management Tasks*

Subtask	Description	Reference
<i>Volumes</i>	View a list of volumes on the selected server.	<a href="#">Managing NSS Volumes</a>
<i>Details</i>	View information about a selected volume.	<a href="#">Viewing the Details of an NSS Volume</a>
<i>New</i>	Create a new unencrypted volume.	<a href="#">Creating and Configuring Unencrypted NSS Volumes with iManager</a>
<i>Delete</i>	Delete a selected volume and all of its contents.	<a href="#">Deleting an NSS Volume</a>
<i>Rename</i>	Rename a selected volume.	<a href="#">Renaming an NSS Volume</a>
<i>Activate</i>	Mount and activate a deactive or unmounted volume.	<a href="#">Activating and Deactivating an NSS Volume</a>



Subtask	Description	Reference
<i>Deactivate</i>	Deactivate an active volume.	<a href="#">Activating and Deactivating an NSS Volume</a>
<i>Mount</i>	Mount an unmounted volume. (A volume must be mounted to view its details.)	<a href="#">Mounting and Dismounting an NSS Volume (Linux)</a>
<i>Dismount</i>	Dismount a mounted volume.	<a href="#">Mounting and Dismounting an NSS Volume (Linux)</a>
<i>Move</i>	Move a selected volume to reorganize and redistribute storage on the same server (or to other servers) in response to changing business needs.	<a href="#">Moving Volumes with iManager</a>
<i>Split</i>	Split a selected volume to reorganize and redistribute storage on the same server (or to other servers) in response to changing business needs.	<a href="#">Splitting Volumes with iManager</a>
<i>Properties</i>	View or set attributes for a volume, including the Lookup Name Space. View usage statistics for a volume.	<a href="#">Viewing Properties of an NSS Volume</a>
<i>User Quotas</i>	View and manage user quotas.	<a href="#">Configuring User Space Quotas</a>
<i>Offline</i>	Take a selected snapshot volume offline, where it remains active in the background.	<a href="#">Moving an Active Pool Snapshot Offline from the Pools Page</a>
<i>Update eDirectory</i>	If a storage object is not recognized automatically, update eDirectory objects to force recognition.	<a href="#">Updating eDirectory Volume Objects</a>

## Volume Job Control (NetWare)

You can use NSS Distributed File Services to move an NSS volume to a different pool. You can also use DFS to split part of a volume, including all data beneath a specified directory in a volume's file structure, to a new volume in a different pool.

Use the following table to find information about subtasks available from the Move/Split Job Control task.

**Table 4-4** *Explanation of Volume Job Control Tasks*

Subtask	Description	Reference
<i>Create a Move/Split Job</i>	Move or split an NSS volume.	<a href="#">Prerequisites and Guidelines for DFS</a>
<i>Move/Split Job Report</i>	After creating a move volume or split volume job, view the job's status report.	<a href="#">Managing Move or Split Jobs</a>
<i>Pause</i>	Pause active move and split jobs.	<a href="#">Pausing a Move or Split Job</a>

Subtask	Description	Reference
<i>Resume</i>	Resume previously paused move and split jobs.	<a href="#">Resuming a Move or Split Job</a>
<i>Reschedule</i>	Reschedule pending move and split jobs.	<a href="#">Rescheduling a Move or Split Job</a>
<i>Delete</i>	Cancel and delete pending move and split jobs.	<a href="#">Deleting a Move or Split Job</a>
<i>Finish</i>	Manually force a job completion to occur for a selected move or split job that cannot complete without manual intervention. Typically, this occurs if files are open during the move or split process and cannot be copied to the new location.	<a href="#">Finishing a Move or Split Job</a>

## Partitions

NSS creates and deletes partitions for you when you work with software RAIDS and pools in iManager and NSSMU. To delete partitions manually, use the NSSMU Partitions page.

Use the following table to find information about subtasks available from the Partitions task.

**Table 4-5** *Explanation of Partition Management Tasks*

Subtask	Description	Reference
<i>Partitions</i>	View a list of partitions on the selected server.	<a href="#">Viewing a List of Partitions</a>
<i>Details</i>	View information about a selected partition.	<a href="#">Viewing Details for a Partition</a>
<i>Label</i>	Add a label to a selected partition.	<a href="#">Labeling a Partition</a>

## Software RAIDs

You can create and manage software RAID 0, 1, and 5 devices to improve storage performance and reliability. You must use NSSMU to create a software RAID 10 or 15 pool.

Use the following table to find information about subtasks available from the Software RAIDs task.

**Table 4-6** *Explanation of Software RAID Management Tasks*

Subtask	Description	Reference
<i>Software RAIDs</i>	View a list of software RAID devices on the selected server.	<a href="#">Viewing Software RAID Devices on a Server</a>
<i>Details</i>	View the details of a selected software RAID device.	<a href="#">Viewing Details of a Software RAID Device</a>

Subtask	Description	Reference
<i>New</i>	Create a new software RAID 0, 1, or 5 device for the selected server.	<a href="#">Creating a Software RAID Device Using iManager</a>
<i>Rename</i>	Rename a selected software RAID device.	<a href="#">Renaming a Software RAID Device</a>
<i>Increase Size</i>	<p>Expand an existing software RAID device by adding a partition to the RAID (up to the limit for that type of RAID). If there are no devices available, the button is disabled.</p> <p>Each partition you add must reside on a different device. You can add partitions that match the shared state of current member devices. They must be all local or all shared; you cannot mix them.</p>	<a href="#">Increasing the Size of a Software RAID Device</a>
<i>Restripe</i>	Complete a restriping process for a RAID 0 or RAID 5 device that has been paused.	<a href="#">Restriping a Software RAID 0 or 5 Device</a>
<i>Delete</i>	Delete the selected software RAID device and remove the RAID relationship between member partitions and the underlying storage structures. All data on the member partitions is lost.	<a href="#">Deleting a Software RAID Device</a>
<i>Pools</i>	Create and manage software RAID devices.	<a href="#">Viewing Pools on a Software RAID Device</a>
<i>Partitions</i>	View details about partitions in the RAID. In some cases, you can also delete a partition to repair a RAID.	<a href="#">Viewing Partitions on a Software RAID Device</a>

## Devices

You can configure, mount, and maintain a wide selection of storage devices, including direct-attached-storage devices, network-attached storage devices, networked storage devices in a Fibre Channel or iSCSI storage area network (SAN), and hardware device arrays.

Use the following table to find information about subtasks available from the Devices task.

**Table 4-7** *Explanation of Device Management Tasks*

Subtask	Description	Reference
<i>Devices</i>	A list of all local and external devices available on the selected server.	<a href="#">Viewing Devices on a Server</a>
<i>Details</i>	View information about a selected device.	<a href="#">Viewing Details for a Device</a>

Subtask	Description	Reference
<i>Initialize Disk</i>	Initialize the selected device by erasing its partition table and effectively destroying all its data.	<a href="#">Initializing a Disk</a>
<i>Multipath</i>	For network configurations with multiple paths between network devices, opens the Multipath page where you can set the primary path and path failover priorities for fault tolerance of connections between host bus adapters and storage devices. You can also bring paths up and down.	<a href="#">Managing Multiple Connection Paths to Devices (NetWare)</a>
<i>Set Default Path</i>	Sets the connection to selected device to its user-defined default primary path.	<a href="#">Setting the Primary Path for a Device to Its Default Path</a>
<i>Reset Registry</i>	Resets the multipath priority settings for a selected device in the server registry to its user-defined defaults.	<a href="#">Resetting the Server Registry with Default Priority Settings for a Device</a>
<i>Shareable for Clustering</i>	Enable device sharing to support high-availability server clusters.	<a href="#">Sharing Devices</a>
<i>Pools</i>	View a list of pools on a device.	<a href="#">Viewing Pools on a Device</a>
<i>Partitions</i>	View information about partitions configured on a device.	<a href="#">Viewing Partitions on a Device</a>

## Scan for Devices

You can scan for devices not found automatically by the Media Manager.

Use the following table to find information about subtasks available from the Scan for Devices task.

**Table 4-8** *Explanation of the Scan for Devices Task*

Subtask	Description	Reference
<i>Scan (Continue)</i>	After you select a server to manage, scan for devices added since you booted the server.	<a href="#">Scanning for Devices</a>

## 4.2 NSS Management Utility

The Novell Storage Services Management Utility (NSSMU) is a console-based utility for managing NSS storage media on a server. You can use NSSMU at any time as an alternative to the browser-based iManager Storage plug-in.

- [Section 4.2.1, “NSSMU for Linux,” on page 53](#)
- [Section 4.2.2, “NSSMU for NetWare,” on page 54](#)

## 4.2.1 NSSMU for Linux

For OES Linux, NSSMU is installed when you install NSS. The Linux install creates symlinks in the `/opt/novell/nss/sbin` folder for common nss utilities, including NSSMU. Symlinks allow the path to NSSMU to become part of the root user's path, which allows you to run it with the command `nssmu`.

The following table identifies key functions available in NSSMU for Linux.

**Table 4-9** *Summary of Management Options in NSSMU for Linux*

Management Options	Description
<i>Device Management</i>	Initialize and maintain physical storage devices and software RAID devices available to this server. See Software RAID Devices to create, repair, or delete RAID.
<i>F1</i> = Help	
<i>F3</i> = Initialize device (do not initialize your system device)	
<i>F5</i> = Refresh display	
<i>F6</i> = Share (shareable/not shareable for clustering)	
<i>Enter</i> = Show partitions	
<i>Pool Management</i>	Create, delete, rename, and expand NSS storage pools to efficiently use all free space in the available devices.
<i>Ins</i> = Create a pool	
<i>Del</i> = Delete a pool	
<i>F3</i> = Expand a pool (add partitions)	
<i>F4</i> = Update NDS™/eDirectory	
<i>F5</i> = Refresh details of the pool	
<i>F6</i> = Rename a pool	
<i>F7</i> = Activate/deactivate pools	
<i>F8</i> = More (list more options)	
<i>F9</i> = Show/Hide deleted volumes	
<i>F10</i> = List devices (that a pool resides on)	
<i>Enter</i> = Show volumes (on a server or pool)	

Management Options	Description
<i>Volume Management</i>	Use this menu option to create, delete, rename, activate/deactivate, and mount/dismount NSS volumes and to set their attributes.
<i>Ins</i> = Create an unencrypted volume	To store data in encrypted format, specify a password when you create the volume. This enables the Encryption attribute. The encryption setting persists for the life of the volume. The encryption password can be 2 to 16 standard ASCII characters, with a suggested minimum of 6. The password generates a 128-bit NCI key for encryption. You must supply the password on the first volume activate after a system boot or reboot.
<i>Del</i> = Delete a volume	
<i>F2</i> = Rename mount point for the volume (new path with volume name)	
<i>F3</i> = Rename volume	
<i>F4</i> = Update NDS/eDirectory	
<i>F5</i> = Refresh details of the volume	
<i>F6</i> = View compression statistics	
<i>F7</i> = Dismount/mount a volume. If it is encrypted, the volume prompts for a password on the first mount after a system boot or reboot.	
<i>F8</i> = More (list more options)	
<i>F9</i> = Name Space - choose Long, UNIX, DOS, Mac	
<i>Enter</i> = Set or view volume properties	On Linux, you can mount encrypted volumes only from NSSMU on the first time after a system reboot. Provide the password when needed. Until you provide a password for encrypted volumes, you cannot mount multiple encrypted volumes at a time.
<i>Software RAID Management</i>	A software RAID device emulates a hardware RAID device. RAID devices combine partitioned space on multiple physical devices into a single virtual device that you manage like any device. Each member device contributes an equal amount of space and only a single partition to the RAID.
<i>Ins</i> = Create a software RAID (0, 1, or 10)	
<i>Del</i> = Delete a software RAID device	
<i>F3</i> = Expand a RAID device (add partitions)	
<i>F4</i> = Rename a RAID device	
<i>F5</i> = Refresh details of the software RAID device	
<i>F6</i> = Restripe (resume restriping for paused RAID 0)	
<i>Enter</i> = Show segments (list member partitions for selected device)	
<i>Partition Management</i>	Use this option to view details about partitions. You can also delete a single partition at a time when repairing a failed software RAID partition.
<i>Ins</i> = Create an NSS partition	To delete all software RAID partitions, delete the RAID itself.
<i>Del</i> = Delete an NSS partition	
<i>F5</i> = Refresh details of the partition	
<i>F6</i> = Label	
<i>Enter</i> = Show volumes	

## 4.2.2 NSSMU for NetWare

NSSMU is the interface you use to set up your basic NSS storage solution during the NetWare installation. For example, you must use NSSMU to create the `sys` pool and `sys :` volume. NSSMU allows limited creation and management of the following NSS components: devices (including software RAID 0, 1, 5, 10, and 15), partitions (NSS and iSCSI), pools, and volumes.

To set up iSCSI partitions, use the Partitions option in NSSMU. You must have iSCSI target software loaded to make your server into an iSCSI disk server. You cannot create iSCSI devices

during the initial installation. For information, see the *iSCSI 1.1.3 Administration Guide for NetWare 6.5* ([http://www.novell.com/documentation/iscsi1\\_nak/iscsi/data/h4hgu4hs.html](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html)).

---

**IMPORTANT:** NSSMU is not meant to replace iManager or any other NetWare file system management utilities. However, it is the only management tool that can access your NetWare server if you accidentally delete or rename the `sys :` volume.

---

Load the NSSMU module (`nssmu.nlm`) from the command line of the server console by entering

`nssmu`

The following table identifies key functions available in NSSMU for NetWare.

**Table 4-10** *Summary of Management Options in NSSMU for NetWare*

Function Keys	Description
<i>Device Management</i>	Use this menu option to initialize and maintain a wide selection of physical storage devices and software RAID devices available to this server. See Software RAID Devices to create, repair, or delete RAID.
<i>F1</i> = Help	
<i>F2</i> = Scan for devices	
<i>F3</i> = Initialize device (not for your <code>sys</code> pool or <code>sys :</code> volume)	
<i>F4</i> = Multipath (available only if there are multiple paths for the device)	
<i>F5</i> = Select/deselect one or more devices for Initialize or Share actions	
<i>F6</i> = Share (shareable/not shareable for clustering)	
<i>F9</i> = Refresh details of the device	
<i>Enter</i> = Show partitions	
<i>Pool Management</i>	Use this menu option to create, delete, rename, and expand NSS storage pools to efficiently use all free space in the available devices.
<i>Ins</i> = Create a pool	
<i>Del</i> = Delete a pool	
<i>F3</i> = Expand a pool (add partitions)	
<i>F4</i> = Update NDS/eDirectory	
<i>F5</i> = Refresh details of the pool	
<i>F6</i> = Rename a pool	
<i>F7</i> = Activate/deactivate pools	
<i>F8</i> = More (list more options)	
<i>F9</i> = Show/hide deleted volumes	
<i>F10</i> = List devices (that a pool resides on)	
<i>Alt+F8</i> = Update NDS/eDirectory	
<i>Enter</i> = Show volumes (on a server or pool)	

Function Keys	Description
<i>Volume Management</i>	Use this menu option to create, delete, rename, activate/deactivate, and mount/dismount volumes and to set volume attributes.
<i>Ins</i> = Create an encrypted or unencrypted NSS volume	
<i>Del</i> = Delete a volume	
<i>F3</i> = Rename a volume	
<i>F5</i> = Refresh details of the volume	To store data in encrypted format, specify a password when you create the volume. This enables the Encryption attribute. The encryption setting persists for the life of the volume. The encryption password can be 2 to 16 standard ASCII characters, with a suggested minimum of 6. The password generates a 128-bit NCI key for encryption. Supply the password on the first volume activate after a system boot or reboot.
<i>F6</i> = Deactivate/activate the volume. If it is encrypted, the volume prompts for a password on the first activation after a system boot or reboot.	
<i>F7</i> = Dismount/mount a volume	
<i>F8</i> = More (list more options)	
<i>F9</i> = Name Space - choose Long, UNIX, DOS, Mac	
<i>Alt+F8</i> = Update NDS/eDirectory	
<i>Enter</i> = Properties (set volume properties)	
<i>Software RAID Management</i>	A software RAID device emulates a hardware RAID device. RAID devices combine partitioned space on multiple physical devices into a single virtual device that you manage like any device. Each member device contributes an equal amount of space and only a single partition to the RAID.
<i>Ins</i> = Create a software RAID (0, 1, or 5)	
<i>Del</i> = Delete a software RAID device	
<i>F3</i> = Expand a RAID device (add segments)	
<i>F5</i> = Refresh details of the software RAID device	
<i>F6</i> = Restripe (resume restriping for paused RAID 0 or 5 device)	
<i>F8</i> = Show pools (list pools on selected RAID 0 or 5; for RAID 1, its single member pool appears in device details)	Create a software RAID 10 device by mirroring software RAID 0 devices. Create a software RAID 15 device by mirroring software RAID 5 devices.
<i>Enter</i> = Show segments (list member partitions for selected device)	
<i>Partition Management</i>	Use this option to create and manage NSS partitions and iSCSI partitions. To create iSCSI partitions, the server must run iSCSI target software and be the server you plan to make your iSCSI disk server.
<i>Ins</i> = Create an NSS or an iSCSI partition	
<i>Del</i> = Delete an NSS or an iSCSI partition	
<i>F3</i> = Mirror (create a software RAID-1 device)	
<i>F5</i> = Refresh details of the partition	
<i>Enter</i> = Show volumes	

## 4.3 NSS Commands and Utilities

Command line instructions are available to control most NSS functions, using the following consoles:

- [Section 4.3.1, “NSS Commands,” on page 57](#)
- [Section 4.3.2, “NSS Utilities,” on page 57](#)
- [Section 4.3.3, “NSS Console for OES Linux,” on page 57](#)
- [Section 4.3.4, “NetWare Server Console,” on page 58](#)



### 4.3.1 NSS Commands

To view a list of NSS commands and options, enter the following command at the NetWare server console command prompt or the NSSCON prompt:

```
nss /help
```

For information about commands, see [Chapter 23, “Using NSS Server Console Commands and Utilities,” on page 299](#).

### 4.3.2 NSS Utilities

In addition to NSSMU and NSSCON, NSS provides several utilities to support command line management of NSS pools and volumes.

For information about NSS utilities for Linux, see the following:

- Attributes (ATTRIB) Utility ([Section A.2, “ATTRIB \(Linux\),” on page 360](#))
- Compression Fix (COMPFIX) Utility ([Section A.4.2, “COMPFIX for Linux,” on page 363](#))
- Metadata Migration (METAMIG) Utility ([Section A.8, “METAMIG \(Linux\),” on page 374](#))
- Rebuild and Verify Simple User Interface (RAVSUI) Utility ([Section A.13, “RAVSUI \(Linux\),” on page 378](#))
- Rebuild and Verify Viewer (RAVVIEW) Utility ([Section A.14, “RAVVIEW \(Linux\),” on page 380](#))
- NSS File System Trustee Rights (RIGHTS) Utility ([Section A.17.1, “RIGHTS \(Linux\),” on page 383](#))

For information about all utilities, see [Appendix A, “NSS Utilities,” on page 359](#).

### 4.3.3 NSS Console for OES Linux

The NSS Console (NSSCON) utility for OES Linux provides a command line interface in a console environment familiar to NetWare users. Use it to issue NSS commands and to monitor NSS activity through console messages.

The Linux install creates symlinks in the `/opt/novell/nss/sbin` folder for common NSS utilities, including NSSCON. Symlinks allow the path to the NSSCON to become part of the root user’s path, which allows you to run NSSCON by entering `nsscon` at the system prompt.

The NSSCON utility uses a device file (`/dev/nsscmd`), which is set up to allow access only for the root user. Thus, only root can run `nsscon`. If you want to give access to a group of local users to run `nsscon`, use the `chmod` command on `/dev/nsscmd` to change the POSIX permissions for that group.

You can optionally add `/opt/novell/nss/sbin` to the PATH environment variable:

**1** At a shell prompt, log in as the root user.

**2** From a Bash shell, set the path with the command

```
export PATH=$PATH:/opt/novell/nss/sbin
```

This allows you to run `nsscon` by entering `./nsscon` at the system prompt.

To run NSSCON:

- 1 At a shell prompt, log in as the root user.
- 2 Use one of the following methods to start NSSCON:
  - At a shell prompt, load NSSCON when its path is symlinked by entering  
`nsscon`
  - At a shell prompt, load NSSCON by its absolute pathname by entering  
`/opt/novell/nss/sbin/nsscon`
  - At a shell prompt, load NSSCON when its path is in the PATH environment variable by entering  
`./nsscon`

### 4.3.4 NetWare Server Console

All NSS commands and utilities can be issued from the NetWare server console. Use the server console to monitor NSS activity through console messages.

## 4.4 Novell NetStorage

Novell NetStorage provides a Web-based interface to access directories and files on your NSS volumes on NetWare. You can also manage file system trustees, file system trustee rights, and directory and file attributes for the NSS file system on NetWare.

### 4.4.1 Accessing NetStorage

#### In iManager (OES SP1)

Beginning in OES SP1, NetStorage is accessible from within iManager.

- 1 Log in to iManager in the eDirectory tree of the servers that you want to manage.  
For information, see [Section 4.1.3, “Accessing Novell iManager,” on page 45](#).
- 2 In *Roles and Tasks*, select *NetStorage*.

#### Direct URL

To avoid conflicts, the date and time on the workstation being used to access NetStorage should be reasonably close (within a few hours) to the date and time on the server running NetStorage.

- 1 Launch your Web browser and open it to the following location:

```
http://192.168.1.1/oneNet/NetStorage
```

Replace *192.168.1.1* with the actual DNS name or IP address of your NetStorage server or the IP address for Apache-based services. If Apache-based services use a port other than 80, you must also specify that port number with the URL.

For example, if the port number is 51080, the URL would be in the form

```
http://192.168.1.1:51080/oneNet/NetStorage
```

- 2 Log in with your administrator username and password to manage file system access for directories and files on NSS volumes.

NetStorage uses Novell eDirectory for authentication. You can also log in as any username with equivalent rights to the administrator. This limitation does not apply if you have created a Storage Location object using SSH (Secure Shell).

## 4.4.2 Configuring File System Trustees, Trustee Rights, and Attributes

Using NetStorage, you can set file system trustees, trustee rights, and attributes for directories and files on NSS volumes on your Linux or NetWare servers, using the NetWare Info tab and NetWare Rights tab in Properties. Although the option label refers to NetWare, use the option for both your Linux and NetWare NSS volumes.

For information about file system trustees, trustee rights, and attributes for directories and files on NSS Volumes, see the *File Systems Management Guide for OES*.

### Directory or File Attributes

- 1 In NetStorage, select the file or directory, then click the *NetWare Info* tab to view or modify NSS directory or file attributes.

### NSS File System Trustee Rights

- 1 In NetStorage, select the file or directory, then click the *NetWare Rights* tab to view or modify NSS file system trustee rights.

## 4.4.3 Purging and Salvaging Deleted Files

Using NetStorage, you can purge and possibly undelete, or *salvage*, NSS files that were previously deleted. For information, see [Section 19.5.1, “Salvaging or Purging Deleted Files with NetStorage,” on page 259](#).

## 4.4.4 Browsing Directories and Files

Administrators and users can use NetStorage to browse directories and files in an NSS volume.

## 4.4.5 Additional Information

For information, see:

- *File Systems Management Guide for OES*
- *OES NetStorage Administration Guide for Linux* ([http://www.novell.com/documentation/oes/netstor\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/netstor_lx/data/h9izvdye.html))
- *OES NetStorage Administration Guide for NetWare* (<http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html>)

## 4.5 Novell Remote Manager

Novell Remote Manager (NRM) is a browser-based management utility for server health monitoring.

- [Section 4.5.1, “Novell Remote Manager for Linux,” on page 60](#)
- [Section 4.5.2, “Novell Remote Manager for NetWare,” on page 60](#)
- [Section 4.5.3, “Accessing Novell Remote Manager,” on page 60](#)

### 4.5.1 Novell Remote Manager for Linux

With the NCP™ Server plug-in, Novell Remote Manager for Linux allows you to browse NSS volumes on your Linux servers. It does not support the following tasks for NSS on Linux:

- Management of partitions, pools, and volumes
- Configuration of file system trustees and attributes for directories and files
- Purging and salvaging deleted files

For information, see the *Novell Remote Manager Administration Guide for Linux for OES*.

### 4.5.2 Novell Remote Manager for NetWare

Novell Remote Manager for NetWare provides most of the functionality of the Monitor utility and other functionality from server-console-based utilities. It is the primary management tool for NetWare Traditional File System volumes.

NRM for NetWare also supports some common tasks for managing NSS pools and volumes. For example, you can configure devices, NSS pools, and NSS encrypted or unencrypted volumes. However, these tasks are best managed with iManager or NSSMU. It is also used for purging and salvaging deleted volumes, directories, or files. For information, see [“Salvaging and Purging Deleted Volumes, Directories, and Files” on page 255](#).

For NSS, NRM for NetWare is commonly used to manage file system trustees, file system trustee rights, and attributes for directories and files in your NSS file system. For information, see the *File Systems Management Guide for OES*.

You can also use Novell Remote Manager to monitor your server’s health, change the configuration of your server, or perform diagnostic and debugging tasks.

For information, see the *Novell Remote Manager for NetWare Administration Guide for OES*.

### 4.5.3 Accessing Novell Remote Manager

#### Prerequisites

Your configuration must satisfy the following prerequisites:

- Make sure SSL 3.0 (where available) or SSL 2.0 is enabled in your Web browser.

Novell Remote Manager requires an SSL connection between your Web browser and the target server where it is running. You must enable SSL services for your Web browser; otherwise, the browser displays an error when it tries to display the Novell Remote Manager Web pages later.

- For an OES Linux server, you must enable Linux User Management (LUM) in your eDirectory tree and install and configure LUM on the target server to be able to log in to NRM using your eDirectory credentials.

For information, see the *Novell Linux User Management 2.2 Technology Guide*.

- Ports 8008 (insecure) and 8009 (secure) are the default ports used for accessing Novell Remote Manager. If you change the port number assigned to it, make sure you specify the same value for the port number when you log in.

## Prerequisites for Admin User Access on Linux Servers

For users to log in to an OES Linux server as an Admin user or equivalent, one of the following conditions must be met:

- The Admin user must be associated to the group that has the Supervisor right for the Entry Rights property for the UNIX Workstation object.
- The Admin user must have the Supervisor right for the Entry Rights to the NCPTM object that represents the Linux server in the eDirectory tree.

eDirectory users who have the Supervisor right to UNIX Workstation object or the NCPTM object have rights equivalent to Root on the Linux server.

Users who are enabled for Linux User Management can be distinguished in iManager as follows:

- They have a Linux Profile tab on their Modify User page in iManager.
- Their eDirectory object is associated with the UNIX Workstation object that represents the Linux server.

Use iManager or the LUM command line utility `namuseradd` to enable users for Linux User Management.

For information, see the *Novell Linux User Management 2.2 Technology Guide*.

## Prerequisite for Admin User Access on NetWare Servers

To access all pages necessary to manage your server remotely, log in as a user with the Supervisor right to the Server object. Usually, this is the Admin or a user with rights equivalent to the Admin user.

## Procedure

- 1 From your Web browser, enter one of the following

`http://server-ip-address:8008`

`https://server-ip-address:8009`

Replace *server-ip-address* with the IP address of the server you want to manage. If you have Domain Name Services (DNS) installed on your network for server name-to-IP address resolution, you can optionally use the server's DNS name instead of the IP address.

- 2 Determine the authenticity of the SSL certificate, then accept if the certificate is valid.
- 3 When the Login page appears, type the username and password of the Admin user or equivalent.

Log in as user Root or as an eDirectory Admin user that is Linux User Management enabled.

- 4 Click *OK* to log in to the target server and initiate your SSL session.

The management interface opens in your Web browser. After logging in, your SSL session for Novell Remote Manager remains open until you close all your browser windows at that workstation.

## 4.6 Novell Client

- [Section 4.6.1, “Novell Client for Linux,” on page 62](#)
- [Section 4.6.2, “Novell Client for Windows 2000/XP/2003,” on page 62](#)

### 4.6.1 Novell Client for Linux

The Novell Client™ for Linux software allows users of Linux workstations to access and use all of the services available on servers running Novell eDirectory. The Novell Client brings the full power, ease of use, manageability, and security of eDirectory to Linux workstations. The Novell Client for Linux fully supports NetWare, OES, and eDirectory services and utilities on a Linux workstation, including security, file, and print services through Novell iPrint.

For information, see the *Novell Client for Linux Installation and Administration Guide* ([http://www.novell.com/documentation/linux\\_client/linuxclient/data/h4rudg93.html](http://www.novell.com/documentation/linux_client/linuxclient/data/h4rudg93.html))

### 4.6.2 Novell Client for Windows 2000/XP/2003

In combination with NCP Server, the Novell Client for Windows 2000/XP supports the following:

- Management of file system trustees, trustee rights, and inherited rights filters for directories and files on NSS volumes
- Management of attributes for directories and files on NSS volumes
- Purge and salvage of deleted files on NSS volumes, if the volume is configured to support it by enabling the Salvage Files attribute for a volume
- Drive mapping for NSS volumes
- Login scripts for automatic drive mapping on login

For information, see the *Novell Client for Windows Installation and Administration Guide* (<http://www.novell.com/documentation/noclienu/noclienu/data/h4rudg93.html>).

## 4.7 ConsoleOne

ConsoleOne® is a Java\* application that runs on the server or a workstation. You can use ConsoleOne to manage trustees and attributes for directories and files in your NSS file system. For information about trustees and file and directory attributes for NSS, see the *File Systems Management Guide for OES*.

In NetWare, you can also use ConsoleOne to manage the Distributed File Services Management Context and Junctioning functions, including the following options:

- Create one or more DFS management contexts where servers can host the Volume Location Database (VLDB).

- Create, maintain, and monitor active VLDB services to locate and transparently redirect traffic to volumes.
- Optionally create and manage DFS junctions. When you use iManager's Storage plug-in to move or split NSS volumes, it automatically creates the necessary DFS junctions.

## Loading ConsoleOne

To use ConsoleOne on a client workstation, you need to load ConsoleOne 1.3 or later, then 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\nsscllib.jar
```

```
public\mgmt\consoleone\1.3\lib\nssjavallib.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\nsscllibres.jar
```

- 3 Copy the following .jar file to your local snap-ins directory:

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

## Enabling the ConsoleOne Media Snap-In for Managing DFS

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

This opens the ConsoleOne interface, where you can access the Media snap-in for configuring DFS.

- 5 Select the *Media* snap-in.

## Additional Information

For more information, see the *ConsoleOne User Guide* ([http://www.novell.com/documentation/consol13/c1\\_enu/data/hk42s9ot.html](http://www.novell.com/documentation/consol13/c1_enu/data/hk42s9ot.html)).

## 4.8 Virtual File Services, APIs, 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 as `_admin`), that exists on each server. 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, you can potentially control all server management functions.

For more information and instructions, see *NDK: Virtual File Services for NetWare* (<http://developer.novell.com/ndk/vfs.htm>) in the Novell Developers Kit (NDK) documentation.

For NetWare, other APIs are available in the *NDK: Novell Storage Architecture Component (Media Manager and NWP4)* (<http://developer.novell.com/ndk/storarch.htm>).



# Planning for NSS Storage Solutions

# 5

Use the information in this section to plan your storage deployment using the NSS file system.

- [Section 5.1, “Prerequisites,” on page 65](#)
- [Section 5.2, “Planning to Move NSS Pools and Volumes Cross-Platform,” on page 66](#)
- [Section 5.3, “Planning for Your Storage Needs,” on page 66](#)
- [Section 5.4, “Creating NSS Storage Objects in eDirectory,” on page 69](#)
- [Section 5.5, “Naming NSS Storage Objects,” on page 69](#)
- [Section 5.6, “What’s Next,” on page 71](#)

## 5.1 Prerequisites

### NetWare and Linux

The following requirements apply to NSS on NetWare® and Linux:

- ☐ A server running OES NetWare or OES Linux
- ☐ At least 10 MB free space on the server storage for each NSS volume you plan to create, up to the combined space you want to allocate to NSS volumes
- ☐ A physical or logical device up to 2 TB in size and with free space available

### Linux

The following requirements apply to NSS on Linux, regardless of the volume manager used:

- ☐ You must select NSS from the Software options on the Installation Settings page in the YaST install; it is not installed by default.

If you did not install NSS during the OES installation, you can install it and its dependencies later from YaST.

- ☐ OES Linux uses a Linux traditional file system for its system volume, such as ReiserFS.
- ☐ After installing OES Linux, install only approved updates. If unapproved updates are installed, NSS cannot load. Refer to the *OES Linux Installation Guide* ([http://www.novell.com/documentation/oes/install\\_linux/data/front.html](http://www.novell.com/documentation/oes/install_linux/data/front.html)) to install the approved updates.
- ☐ The server must have a Novell® eDirectory™ Read/Write replica available when you create an NSS pool or volume on Linux; otherwise, NCP™ cannot map to the pool or volume.

### Linux with EVMS as the Volume Manager

The device where you want to create NSS volumes must be managed by EVMS if you plan to use the Storage plug-in to iManager or NSSMU to create and manage NSS partitions, pools, and volumes.

Linux uses Linux Volume Manager (LVM) as its default volume manager. With the Linux 2.6 kernel, any disk managed by the LVM cannot be managed by Enterprise Volume Management System (EVMS). Thus, the disk(s) where the boot partition (such as /boot for Grub) and system partition (such as /root for the swap and system volumes) reside are typically unavailable to NSS.

To be able to create an NSS data volume on the same device as your boot partition or system partition, make sure to configure the device for EVMS during the install. For information, see “[Installing Linux with EVMS as the Volume Manager of the System Device](http://www.novell.com/documentation/oes/install_linux/data/front.html)” in the *OES Linux Installation Guide* ([http://www.novell.com/documentation/oes/install\\_linux/data/front.html](http://www.novell.com/documentation/oes/install_linux/data/front.html)).

When your data volumes are on non-system devices, do not configure devices during the install. Instead, leave the devices as unconfigured free space and do not assign a volume manager for them. After the install, create the volumes with NSSMU or the Storage plug-in to iManager.

### Linux with Other Volume Managers

If you want to use other volume managers instead of EVMS, you must manage NSS pools and volumes with Linux and NSS commands and utilities. Some capabilities of NSS are unavailable. For more information, see [Chapter 22, “Using NSS on Devices Managed by Non-EVMS Volume Managers,”](#) on page 291.

## 5.2 Planning to Move NSS Pools and Volumes Cross-Platform

If you plan to move NSS pools and volumes cross-platform between OES NetWare and OES Linux servers on non-clustered devices, see the following:

- [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,”](#) on page 82
- [Section 6.7, “Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux,”](#) on page 87
- [Section 6.8, “Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later,”](#) on page 92
- [Section 6.9, “Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare,”](#) on page 95
- [Section 6.11, “Upgrading NetWare 5.x NSS Volumes and NetWare Traditional Volumes to NSS Volumes,”](#) on page 98

If you plan to move NSS pools and volumes cross-platform in a mixed-cluster configuration, see the *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html)).

## 5.3 Planning for Your Storage Needs

Before you begin to subdivide, or 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. Use the following guidelines when planning how to implement your storage solution:

- [Section 5.3.1, “Devices,”](#) on page 67
- [Section 5.3.2, “Software RAID Devices,”](#) on page 67
- [Section 5.3.3, “Multiple I/O Paths \(NetWare\),”](#) on page 68

- [Section 5.3.4, “Free Space,” on page 68](#)
- [Section 5.3.5, “NSS Pools and Volumes,” on page 68](#)
- [Section 5.3.6, “NSS Encrypted Volumes,” on page 69](#)

## 5.3.1 Devices

NSS recognizes devices up to 2 TB in size. If you have physical or logical devices larger than 2 TB in size, you must carve them into logical devices of less than 2 TB each, using the vendor-provided or other third-party disk-carving software.

Storage devices can be local to the server, such as a system hard drive, or external to the server, such as with direct-attached storage or in a Fibre Channel or iSCSI storage area network (SAN).

A local hard drive typically hosts the system software and can optionally host applications and user data.

---

Linux	<p>The operating system is stored on the system hard drive in a Linux traditional file system (such as ReiserFS or EXT3). The device is typically managed by LVM, but you can configure it to be managed by EVMS. For information, see “<a href="#">Installing Linux with EVMS as the Volume Manager of the System Device</a>” in the <i>OES Linux Installation Guide</i>.</p> <p>Use NSS data pools and volumes only on devices managed by EVMS.</p>
NetWare	<p>The operating system is stored on the system hard drive in the NSS <code>sys</code> pool and <code>sys : volume</code>.</p>

---

NSS supports server clusters built with Novell Cluster Services™. Both local and SAN storage devices can be shared between the servers, depending on network connections. However, you cannot SAN boot cross-platform. For information, see the following:

- [OES Novell Cluster Services Administration Guide for Linux](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html) ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html))
- [OES Novell Cluster Services Administration Guide for NetWare](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html) ([http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html))

For information about using NSS pools and volumes cross-platform, see [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#).

## 5.3.2 Software RAID Devices

NSS supports software RAID 0, 1, 5, 10, and 15

If you use hardware RAID devices, software RAID devices are unnecessary. You can use both hardware and software RAID devices on the same server.

For more information, see [Chapter 9, “Managing Software RAID Devices,” on page 117](#).

### 5.3.3 Multiple I/O Paths (NetWare)

For NetWare, NSS provides multiple I/O paths for fault-tolerant connections between the server and its storage devices. For information, see [Chapter 10, “Managing Multiple Connection Paths to Devices \(NetWare\),” on page 139](#).

### 5.3.4 Free Space

Understanding how much free space you will need from each device helps you during the disk carving phase of the NSS configuration.

To maximize the performance benefits of software RAID devices, partitions used for the RAID should come from different physical devices. For software RAID 1 devices, the mirrored partitions cannot share any disks in common.

For information about space availability, see [Section 7.3, “Viewing Devices on a Server,” on page 101](#).

For Linux, only EVMS-managed devices show up in the NSSMU and the Storage plug-in to iManager. For information, see [Section C.2, “Viewing Free Space Available to EVMS-Managed Devices,” on page 409](#).

### 5.3.5 NSS Pools and Volumes

NSS is the primary storage file system on NetWare. NSS is a data storage file system on Linux.

Linux	Linux requires a Native Linux volume, such as Reiser or EXT2, for the operating system. Use NSS pools and volumes as data volumes only on devices managed by EVMS.
NetWare	The <code>sys</code> pool and <code>sys:</code> volume should be reserved for your system software. Create separate pools for applications and user data.

When creating a pool, assign up to 2 TB from free space for each partition. Use at least 4 partitions of up to 2 TB each to create the maximum-sized pool of 8 TB. To grow a pool, add more partitions of space, up to the limits.

Pools can contain space from multiple devices, using various amounts of free space from each. Pools can use multiple partitions from the same device.

To mirror pools, each pool must use partitions from different devices; mirrored pools can have no devices in common.

Pools can contain multiple volumes, but a volume belongs to only one pool.

When creating a volume, assign it a fixed administrative quota or allow the volume to grow dynamically to the size of the pool. Any given volume’s quota cannot exceed the size of the pool.

Pools can be overbooked. If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool size, using the overbooking feature, although real total size is bound by physical limitations. Because space is allocated to volumes as needed, a volume might not reach its quota.

You can create an NSS data volume on an OES Linux server or an OES NetWare server and then move the volume cross-platform. For information, see [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82.](#)

### 5.3.6 NSS Encrypted Volumes

Encrypted Volume Support is available only in NetWare 6.5 Support Pack 2 or later, OES NetWare and later, and OES SP1 Linux and later. Create encrypted volumes only after you verify a successful system install or upgrade. For information, see [“Managing Encrypted NSS Volumes” on page 201.](#)

## 5.4 Creating NSS Storage Objects in eDirectory

When you create a logical pool or volume on a server, a Storage object is automatically created in Novell eDirectory. By default, the name of the Storage object is the server’s name with an underscore and the object’s name appended (for example, *myserver\_sys*). A Storage object represents a logical or physical object on a server, whether it is a writable disk, a CD, or other storage medium.

For more information about Novell eDirectory, see [“Object Classes and Properties” \(http://www.novell.com/documentation/edir873/edir873/data/fbabihe.html\)](http://www.novell.com/documentation/edir873/edir873/data/fbabihe.html) in the *Novell eDirectory 8.7.3 Administration Guide*.

## 5.5 Naming NSS Storage Objects

Storage object names must be unique on a server. If the server is part of a cluster, then all pools and volumes must have unique names across all servers in the cluster, whether they are in shared relationships or not.

### Case Insensitive Naming Convention

NSS storage object names are case insensitive. Names such as AURORA, Aurora, and aurora are the same. NSS saves pool, volume, and software RAID device names in uppercase. NSS saves names and labels in uppercase. For example, if you enter vol2 as the name, it is saved as VOL2.

---

**IMPORTANT:** Because Linux treats filenames as case sensitive, when using NSS volumes on Linux, make sure to mount the volume with the Long name space (ns=long) option so that file queries are case insensitive. For information, see [Section 13.16, “Mounting an NSS Volume on Linux to Support Case Insensitive File Names \(Linux\),” on page 197.](#)

---

### Number of Characters Allowed

- Device, pool, and volume names consist of at least 2 characters, but not more than 15 characters
- Software RAID 0 and 5 device names consist of 2 to 128 characters.
- RAID 1 device names consist of 2 to 80 characters.
- Partition labels consist of 2 to 128 characters.
- The encryption password for encrypted NSS volumes can be 2 to 16 standard ASCII characters, with a recommended minimum of 6.

## Conventions for Valid Names

Valid storage object names use the following conventions:

- Can contain only valid characters:

ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789\_!@#\$\$%&()

---

**IMPORTANT:** Some characters in NSS pool and volume names on OES Linux can cause problems in the `/etc/fstab` configuration file on Linux. Do not use ampersand (&) and pound (#) characters in pool and volume names.

---

- Cannot contain spaces
- Cannot begin or end with the underscore (\_)
- Cannot contain multiple contiguous underscores (\_\_\_)
- Cannot be one of these case-insensitive reserved names:

ALL

AUX

CLOCK

COM1

COM2

COM3

COM4

CON

LPT1

LPT2

LPT3

NETQ

NSS\_ADMIN

NUL

PIPE

PRN

SYS

## Other Guidelines

- Because the commercial “at” character (@) is an element of electronic mail addresses, such as `code@engineer.com`, it might cause confusion and possible problems in a storage object name. A Web browser or other application could mistake it for an e-mail address. We recommend that you do not use the commercial “at” character in storage object names.
- The percent character (%) might cause problems if it is passed in a format string to an application routine that uses it to delineate parameters. For example, if a volume name that contains the percent character, such as `store%sales`, is passed to an `(s)printf` routine, the `(s)printf` routine might look for parameters that are not there and crash. We recommend that you do not use the percent character in storage object names.

## 5.6 What's Next

Continue with the next section, “Coexistence and Migration Issues” on page 73.





# Coexistence and Migration Issues

# 6

This section discusses the issues involved in the coexistence and migration of Novell® Storage Services™ for Novell Open Enterprise Server.

- [Section 6.1, “Compatibility with Operating Systems and Services,” on page 73](#)
- [Section 6.2, “NSS Kernel Modules in SLES 9,” on page 74](#)
- [Section 6.3, “Comparison of NSS to Other File Systems,” on page 74](#)
- [Section 6.4, “Access Control Issues for NSS on OES Linux,” on page 74](#)
- [Section 6.5, “Compatibility Issues for the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager,” on page 79](#)
- [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#)
- [Section 6.7, “Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux,” on page 87](#)
- [Section 6.8, “Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later,” on page 92](#)
- [Section 6.9, “Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare,” on page 95](#)
- [Section 6.10, “Moving Clustered Devices with NSS Volumes Cross-Platform,” on page 98](#)
- [Section 6.11, “Upgrading NetWare 5.x NSS Volumes and NetWare Traditional Volumes to NSS Volumes,” on page 98](#)

For a general discussion of coexistence and migration issues in OES, see the *OES Coexistence and Migration Guide*.

## 6.1 Compatibility with Operating Systems and Services

The following table summarizes the compatibility of NSS with various network operating systems and directory services versions, along with any dependencies on other products or services.

Linux Operating System	OES SP2 Linux and earlier  Does not support the media format upgrade that occurred in OES SP1 NetWare
NetWare® Operating System	OES SP1 NetWare and later  NetWare 6.5 SP4 and later  Beginning in OES SP1, NSS volumes on NetWare have a new metadata format that is not compatible on earlier NetWare systems and on OES SP2 Linux and earlier.
Directory Services	Novell eDirectory™ 8.7.3 or later

Management Services	Novell iManager 2.5
	Linux User Management (LUM) for Linux-enabling eDirectory users
Protocols	<b>Linux:</b> NCP or other protocols such as NFS or Samba
	<b>NetWare:</b> NCP™ or Native File Access Protocols (AFP, CIFS, NFS)

## 6.2 NSS Kernel Modules in SLES 9

The NSS kernel modules (km\_nss) and their sources are delivered with SLES 9 for the following reasons:

- Kernel modules need to be open-source to meet the GPL requirements.
- If a customer adds a kernel patch for some third-party module, application, or service, the NSS sources must be in the kernel code in order to get recompiled with the patch. Otherwise, an OES deployment using NSS would break.

Although the NSS kernel modules are included in SLES 9, they are not usable without OES, because NSS is a unique file system that is tightly integrated with identity management. The root user is the only local user who can see NSS volumes on a Linux server. NSS needs Linux User Management and Novell eDirectory to establish non-root connections to the volume. It is OES that provides LUM, eDirectory, NCP Server, and volume and user space management tools that make NSS volumes usable on an OES Linux server.

## 6.3 Comparison of NSS to Other File Systems

Use the following table to find comparisons of NSS to other file systems in OES:

Comparison	NSS on NetWare	NSS on Linux	Linux Traditional File Systems	NetWare Traditional File System
Comparison of NSS on NetWare and NSS on Linux (page 417)	X	X		
Comparison of NSS for Linux and Linux Traditional File Systems (page 425)		X	X	
Comparison of NSS on NetWare and the NetWare Traditional File System (page 429)	X			X

## 6.4 Access Control Issues for NSS on OES Linux

This section describes how to use Novell eDirectory, NCP Server, and Linux User Management (LUM) with NSS to provide access to NSS volumes on OES Linux.

- [Section 6.4.1, “NSS and eDirectory,” on page 75](#)
- [Section 6.4.2, “NSS, the Admin User, and the Root User,” on page 75](#)
- [Section 6.4.3, “NSS and File System Users,” on page 76](#)
- [Section 6.4.4, “NSS and NCP Server,” on page 77](#)

- Section 6.4.5, “NSS and Linux User Management,” on page 77
- Section 6.4.6, “Enforcing File Ownership and User Space Restrictions,” on page 78
- Section 6.4.7, “Accessing Files as the Root User,” on page 78
- Section 6.4.8, “Accessing Files via NCP Server Only,” on page 78
- Section 6.4.9, “Accessing Files with Other Protocols or from Linux,” on page 79
- Section 6.4.10, “Accessing Linux Traditional Volumes with NCP Server and Services,” on page 79

## 6.4.1 NSS and eDirectory

Users of the NSS volume and the Admin user (or Admin equivalent user) who manages the volume must be defined as users in Novell eDirectory. The locally defined Linux users, except for the root user, cannot see or access an NSS volume.

NSS controls access to data based on the Novell Trustee Model, which uses file system trustee assignments, trustee rights, and inherited rights filters to control file access. The trustee model depends on the secure directory services provided by eDirectory to manage the file system users. For example, eDirectory users must be authenticated by eDirectory to connect to the server, and then NSS uses the effective file system rights of the user to control access to specific files or directories. For information about the Novell Trustee Model, see “[Understanding File System Access Control for NSS and NetWare Traditional File Systems](#)” in the *File Systems Management Guide for OES*.

For information about managing users with eDirectory, see the *Novell eDirectory 8.7.3 Administration Guide* (<http://www.novell.com/documentation/edir873/edir873/data/a2iii88.html>).

## 6.4.2 NSS, the Admin User, and the Root User

Admin and root are two very different concepts. It is important to understand the role of each in managing your NSS volume on OES Linux.

### Admin User

The Admin user is an eDirectory user who is given all file system trustee rights. The Admin User account, or the Admin equivalent user account, must be defined in eDirectory, assigned as a trustee of the NSS volume, and given all file system trustee rights for that volume. In addition, to administer an OES Linux server, the Admin User account must be Linux-enabled with Linux User Management to have both an eDirectory GUID and a POSIX UID on the server.

For a Linux server, the administrator logs in to iManager, NetStorage, or the Novell Client as the Admin user (or Admin equivalent user) to manage the NSS volume on Linux. NetStorage and the Novell Client are used to manage file system trustee assignments, trustee rights, inherited rights masks, and file and directory attributes. They can also be used to purge and salvage files.

### Root User

Root is a local Linux user. The root user is the all-powerful connection when running on the Linux server. The root user is hardcoded internally in NSS to have all access rights to all files. In this way, the root user on Linux is similar to the Link Connection “0” user on NetWare. The root user is not defined as a user in eDirectory, and the root user is not Linux-enabled.

The root user logs in locally to use NSS utilities (such as NSSCON, NSSMU, RIGHTS, ATTRIB, METAMIG, RAVSUI, and RAVIEW) and NSS commands from the terminal console or the NSS Console (NSSCON). The root user can also execute applicable Linux commands and utilities.

NCP Server uses the root user internally for all Linux VFS (virtual file system) layer communication to the NSS volume, but it provides NSS with UIDs for file ownership via POSIX system calls, and it provides NSS with trustee information via XML calls.

### 6.4.3 NSS and File System Users

In addition to the root user and Admin user, file system users fall into three categories:

#### Local Linux Users

Users who are defined locally for the Linux server. The root user is the only local Linux user who can see and access the NSS volume.

#### eDirectory Users

Users who are defined in eDirectory and granted file system trustee rights to the NSS volume. When eDirectory and NCP Server are running, these users can access NSS volumes via the Novell Client or other NCP-based services such as NetStorage.

NSS uses the eDirectory GUID of a user to control access using the Novell Trustee model. Because they are not Linux-enabled, these users do not have a POSIX UID. Therefore, NSS cannot work with Linux to specify the user as the actual owner of a file. Whenever the user creates a file, NSS assigns the file ownership to the root user, using the root UID. NSS also cannot support user space quotas for these users.

If you move an NSS volume from a NetWare server to a Linux server, the file ownership information is unchanged in the existing files. However, because eDirectory users do not have a UID, NSS cannot report that information correctly to Linux. Instead, NSS identifies the owner as the Nobody user (if it exists) or the root user.

#### Linux-Enabled eDirectory Users

Users who are defined in eDirectory, granted file system trustee rights to the NSS volume, and Linux-enabled with Linux User Management. These users can access an NSS volume with NCP, CIFS/Samba, NFS, AFP (requires a third-party solution), or Linux utilities, commands, or services.

Linux-enabled eDirectory users have both an POSIX UID and an eDirectory GUID, which enables NSS to assign and track ownership by user on Linux. NSS supports hardlinks and user space quotas for Linux-enabled eDirectory users.

---

**IMPORTANT:** Make sure to Linux-enable the users before setting hardlinks or user space quotas.

---

For more information, see “[Access Control for NSS on Linux](#)” in the *File Systems Management Guide for OES*.

## 6.4.4 NSS and NCP Server

NCP Server works with Novell eDirectory, the Novell Client, and other NCP-based services such as NetStorage to authenticate and manage user sessions with OES Linux servers. When NCP Server is running, eDirectory users who have been granted file system trustee access can access an NSS volume with the Novell Client or NCP-based services.

NCP Server enforces access to files and directories based on the Novell Trustee model. NCP Server uses the root user identity for all communication via the Linux VFS (virtual file system) layer to the NSS volume, and then it provides NSS with UIDs for file ownership via POSIX system calls, and it provides NSS with trustee information via XML calls.

NSS does not rely on NCP Server to track file ownership and file system trustee assignments, trustee rights, and inherited rights. NSS must track this information, regardless of what NCP is doing, because it must be capable of delivering the files with other protocols, such as CIFS/Samba or NFS. Also, NSS must be able to work with applications running on the server. File ownership is not tracked by NCP; the information is stored by the file system. NSS tracks file ownership and enforces user space restrictions based on a user's eDirectory GUID.

The Linux file system interface uses UTF-8 encoding for all file names. When accessing files with NCP, make sure to use the UTF-8 enabled NCPs that are available in the latest Novell NCP Client.

## 6.4.5 NSS and Linux User Management

Novell Linux User Management is a directory-enabled application that simplifies and unifies the management of user profiles on Linux-based platforms. Only eDirectory users who are Linux-enabled with LUM can access files with non-NCP protocols such as CIFS/Samba, NFS, or AFP (requires a third-party solution) or Linux utilities, commands, or services. In addition, users must be Linux-enabled before you set up any NSS hard links for the volume and before you set user quotas.

---

**IMPORTANT:** If you Linux-enable a user who has been logged into the system before being Linux-enabled, make sure execute the `resetidcache` command from the NSS Console (`nsscon`) utility afterwards. This allows the proper reporting of ownership because it resets the mapping of user identities in the ID cache and forces it to update with the Linux UID for the user.

---

Linux-enabled eDirectory users have both UIDs as local Linux users and GUIDs as eDirectory users. The UID is needed to execute protocols and services that communicate to NSS through the VFS layer only. The user's UID is also needed to map it to a GUID for tracking file ownership, managing hard links, and enforcing user space restrictions. In addition, NSS uses the GUID to enforce access to the files and directories based on the Novell Trustee Model, which uses file system trustee assignments, trustee rights, and inherited rights filters.

With non-NCP protocols and services, the UID is passed to NSS via the VFS layer. There is no back-end XML call to exchange GUID information as there is with the NCP interface. NSS uses a LUM API to translate the UID to a GUID, and then caches the result for fast mapping on subsequent access by the same UID. With the GUID-UID mapping, NSS can find the GUID for the user who issues the command, then executes the command. If the user is not Linux-enabled, NSS cannot identify a GUID for the UID it receives, and then rejects the command with an error. Only the root user is allowed to access NSS via the VFS layer without having an eDirectory GUID.

## 6.4.6 Enforcing File Ownership and User Space Restrictions

Only Linux-enabled users can have their user space restrictions enforced. You must Linux-enable the users first, then set the user space quota.

If a user is not Linux-enabled, any files created by that user are owned by the root user. When users are not Linux-enabled, NSS cannot record changes in ownership issued from the Linux interface. For example, if the root user issues a `chown` command and provides a UID of a user that is not Linux-enabled, even if that user is in eDirectory, the LUM API cannot find a matching UID, so NSS rejects the request.

## 6.4.7 Accessing Files as the Root User

When accessing an NSS volume from the Linux environment, the root user observes some information differently, depending on whether the eDirectory user is Linux-enabled, or not. Any native Linux commands run from a terminal console on the NSS volume, such as the `ls` command, are sent via the VFS layer. If the users are not Linux-enabled, instead of seeing the local UID of the eDirectory user who owns the file, the root user sees all files as belonging to either the Nobody user (if it exists) or the root user. NSS reports the Nobody UID or the root user UID for display purposes only; it does not change the true file ownership information stored as the user's eDirectory GUID in the metadata of the file system.

For NSS volumes on Linux, the POSIX directory and file permissions are not used to determine access permission. NSS uses the permission fields to store Read Only, Read/Write, Execute, and Hidden attributes for directories and files. NSS does not allow the Linux system to set typical access control permissions in the POSIX fields. It interprets Linux `chmod` commands to apply the values as NetWare directory and file attributes, according to the way NSS maps them to the User, Group, and Other permission fields. For information, see “[Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions](#)” in the *File Systems Management Guide for OES*.

## 6.4.8 Accessing Files via NCP Server Only

When the NCP Server is running and the NCP protocol is the only protocol or service used to access the NSS volume, the NCP Server controls access to the server for eDirectory users. If all of the following conditions exist, it is not necessary to Linux-enable the users with Linux User Management:

- NCP Server must be running.
- Only eDirectory users who have been granted file system trustee access to the volume and the root user have access to the NSS volume.
- eDirectory users access files only through a Novell Client, NetStorage with NCP, or any other NCP-based service.
- NCP Server enforces access to the volume based on the Novell Trustee Model for eDirectory users.
- eDirectory users never need to execute Linux utilities, commands, or services on the volume.
- eDirectory users never need to access the volume via other protocols such as CIFS/Samba, NFS, or AFP (requires a third-party solution).

## 6.4.9 Accessing Files with Other Protocols or from Linux

Only the root user and Linux-enabled eDirectory users who have been granted trustee access can see and access the NSS volume from a Linux interface. eDirectory users must be Linux-enabled with LUM if they access files on the NSS volume with any protocol other than NCP, regardless of whether NCP Server is running. Except where NCP Server is controlling a process, NSS enforces the Novell Trustee Model for user access based on the UID, as mapped to the GIUD. NSS needs users to be Linux-enabled to correctly interpret the UIDs coming in from the VFS layer and match the owner's UID with a GUID that ties back into eDirectory.

Users must also be Linux-enabled if the administrator wants users to be able to use any of the standard Linux utilities, commands, services, or APIs for the NSS volume. Users need a standard Linux UID to issue these commands, and NSS needs a GUID to execute it. Linux-enabled eDirectory users have both. NSS further controls access based on the effective file system trustee rights of the user. An example would be the `chown` command, which requires a username that is changed to a UID by the application, and then is sent via the VFS layer to the NSS volume. NSS matches the UID with a GUID, then is able to enforce the effective rights of the user and execute the command accordingly.

Because NSS controls access based on file system trustee rights, not by the POSIX permissions, Samba connections do not work until this trustee system has been configured for the Linux-enabled eDirectory users of the NSS file system. You cannot set up the standard POSIX permissions for Samba access to an NSS volume. Instead, the Admin user or Admin user equivalent must set up users in eDirectory and make file system trustee assignments, grant trustee rights, and configure inherited rights masks on directories. The eDirectory users must also be Linux-enabled.

## 6.4.10 Accessing Linux Traditional Volumes with NCP Server and Services

NCP Server also works with file systems other than NSS, such as Linux traditional volumes, which have no capability for storing trustees. If the eDirectory user is not Linux-enabled with LUM, when the user creates a file on a Linux traditional volume with the Novell Client (or NCP-based services), the NCP Server cannot identify the UID of the user, so it sets the owner UID to the UID of the root user. NCP Server cannot report back file owners as specific users on the Linux traditional volume if the user is in eDirectory but is not Linux-enabled for the server.

If the administrator wants NCP Server to record and report the true UID of a user who creates the file via the Novell Client, the user must be Linux-enabled. Otherwise, the UID that is saved is the UID of the root user.

## 6.5 Compatibility Issues for the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager

The requirements in this section apply to the following storage-related plug-ins for Novell iManager in OES:

- Storage plug-in
- Archive Versioning plug-in
- File Protocols (CIFS) plug-in (NetWare)

- Clustering plug-in

---

**IMPORTANT:** The Storage, Archive Versioning, CIFS, and Cluster Services plug-ins share common code. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins. If you remove only one of the plug-ins, it removes the common code and breaks the remaining installed plug-ins.

---

## 6.5.1 Novell iManager 2.5

The Storage plug-in for OES SP1 requires Novell iManager 2.5. You cannot use the OES SP1 version of the plug-in with earlier releases of iManager. For information on how to upgrade to iManager 2.5, see “Upgrading to iManager 2.5” ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/ah8bnjv.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/ah8bnjv.html)) in the *Novell iManager 2.5 Installation Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_install\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_install_25/data/hk42s9ot.html)).

## 6.5.2 Web Browser Language Setting

The iManager plug-in might not operate properly if the highest priority Language setting for your Web browser is set to a language other than one of the supported languages. To avoid problems, in your Web browser, click *Tools > Options > Languages*, then set the first language preference in the list to a supported language.

---

**NOTE:** This issue was resolved for the Archive Versioning plug-in for iManager 2.5 in OES SP1.

---

## 6.5.3 Protocols

- “Interoperability of Protocols for the iManager Server and Target Server” on page 80
- “WBEM” on page 81
- “CIFS” on page 81
- “NCP” on page 81
- “Stop EVMSGUI When Working in iManager or NSSMU” on page 81

### Interoperability of Protocols for the iManager Server and Target Server

The following table provides information about the protocols needed to use iManager to manage storage in a heterogeneous environment. A protocol annotated with an asterisk is the default and is configured automatically on the servers. The protocols you use must be loaded and running on both the iManager server and the target server you want to manage.

**Table 6-1** *Interoperability of Protocols for the iManager Server and Target Servers*

iManager 2.5 Server	Target Server			
	OES Linux or Later	OES NetWare or Later, or NetWare 6.5 SP4 and later	NetWare 6.5 SP3	NetWare 6.5 SP2



iManager 2.5 Server	Target Server			
<b>OES Linux or Later</b>	* WBEM	* WBEM	WBEM (Start OpenWBEM.)	
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
<b>OES NetWare and Later, or NetWare 6.5 SP4 and later</b>	* WBEM	* WBEM	WBEM (Start OpenWBEM.)	
		NCP	* NCP	* NCP
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
<b>NetWare 6.5 SP3</b>	* WBEM	* WBEM	WBEM (Start WBEM.)	
		NCP	* NCP	* NCP
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
<b>NetWare 6.5 SP2</b>	Not available	* NCP	* NCP	* NCP

## WBEM

Where WBEM is the default protocol, OpenWBEM is loaded and runs automatically when you start the server. Otherwise, you must start OpenWBEM to use the protocol.

To start WBEM: At the server console, enter

```
openwbem
```

If you receive file protocol errors, it might be because OpenWBEM is not running. For information about installing OpenWBEM, see “[Setting Up OpenWBEM](http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html)” (<http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html>) in the *OpenWBEM Services Administration Guide for OES* (<http://www.novell.com/documentation/oes/cimom/data/front.html>).

## CIFS

Where it is available, CIFS must be configured before you can use it. An additional CIFS setup requirement for Field Patch 2B is noted where it is required. For information, see the *OES Native File Access Protocols Guide* (<http://www.novell.com/documentation/oes/native/data/h9izvdye.html>).

## NCP

NCP is the default protocol when the iManager server or the target server you want to manage are NetWare 6.5 SP3 or NetWare 6.5 SP2.

## Stop EVMSGUI When Working in iManager or NSSMU

You cannot run the Enterprise Volume Management System Graphical User Interface (EVMSGUI) at the same time that iManager or NSSMU is running. If you use EVMSGUI (`evmsgui` command) to manage NSS devices, close any instance of iManager or NSSMU, and vice versa.

## 6.6 Compatibility Issues for Using NSS Cross-Platform

If you plan to move NSS pools and volumes cross-platform between OES NetWare and OES Linux servers, consider the following guidelines:

- [Section 6.6.1, “Guidelines for File Systems on Linux Servers,” on page 82](#)
- [Section 6.6.2, “Guidelines for File Systems on NetWare Servers,” on page 82](#)
- [Section 6.6.3, “Guidelines for Cross-Platform Compatibility of NSS,” on page 83](#)
- [Section 6.6.4, “Guidelines for Moving Devices Cross-Platform,” on page 85](#)

### 6.6.1 Guidelines for File Systems on Linux Servers

OES Linux requires a Linux traditional file system volume for the operating system. The following guidelines apply:

- You cannot install the Linux operating system on an NSS volume.
- Use NSS on Linux only as data pools and volumes.
- You cannot SAN boot cross-platform.
- You cannot move Linux traditional file systems cross-platform to a NetWare server.
- For OES SP1, you cannot create a mixed cluster if the shared pool resides on software RAID 0 or RAID 5 devices. This configuration is not supported by NSS on Linux.
- Do not use the DFS move/split functions for a shared NSS volume in a mixed cluster because DFS is not supported for NSS on Linux.

### 6.6.2 Guidelines for File Systems on NetWare Servers

OES NetWare uses an NSS pool (`sys`) and volume (`sys:`) for the operating system. Upgraded NetWare systems can optionally keep a NetWare Traditional volume as the system volume. The following guidelines apply:

- You cannot install the NetWare operating system on a Linux traditional file system volume or on an NSS volume on Linux.
- Use NSS on NetWare for system or data pools and volumes.
- You cannot use Linux traditional file systems on a NetWare server.
- You cannot SAN boot cross-platform.
- When using the move/split functions in a NetWare cluster, perform the move/split only from an active node.
- You should not move an NSS system volume cross-platform unless you intend to use it as a data volume, or not at all, while it is on a Linux server.
- If you move a non-clustered NSS system pool cross-platform, any volumes it contains function as data volumes on Linux within the [Section 6.6.3, “Guidelines for Cross-Platform Compatibility of NSS,” on page 83](#).
- You can move non-clustered NSS data pools and volumes cross-platform within the [Section 6.6.3, “Guidelines for Cross-Platform Compatibility of NSS,” on page 83](#).

### 6.6.3 Guidelines for Cross-Platform Compatibility of NSS

The following features of NSS are available for NSS on NetWare but are not supported for NSS on Linux.

- “New Media Format for OES SP1 NetWare, NetWare 6.5 SP4, and Later” on page 83
- “NSS Hard Links” on page 83
- “NSS Multiple Path I/O to Devices” on page 83
- “NSS Pool Snapshots” on page 84
- “CD and DVD Media and Image Files” on page 84
- “Novell Distributed File Services” on page 84
- “Novell Archive and Version Services” on page 85
- “Interoperability of NSS Volumes on OES Linux with Novell Audit 2.0 and Novell NSure Audit 1.0.x” on page 85

#### New Media Format for OES SP1 NetWare, NetWare 6.5 SP4, and Later

Beginning with OES SP1 NetWare and NetWare 6.5 SP4, a new metadata structure (Beast Version 3) for NSS volumes on NetWare requires a media format upgrade after you upgrade the operating system. For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33](#).

The media format upgrade is not supported for NSS volumes on OES SP1 Linux. You should not upgrade an NSS volume on NetWare to the new media format if either of the following conditions exist:

- You plan to move one or more devices containing an NSS volume from the NetWare server to a server with OES Linux or later, an earlier version of NetWare 6.5, or OES NetWare
- The volume is a shared in a cluster with nodes with OES Linux or later, an earlier version of NetWare 6.5, or OES NetWare.

#### NSS Hard Links

Beginning with OES SP1 NetWare and NetWare 6.5 SP4, a new metadata structure (Beast Version 3) for NSS volumes on NetWare supports a Hard Links attribute and up to 65,535 hard links for a file, regardless of the length of the filename and the number of name spaces. After you upgrade the operating system, you must upgrade the NSS volume on NetWare to use the new media format. Some restrictions apply. For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33](#).

The new NSS media format is not available for NSS on OES SP1 Linux. Hard link support for OES SP1 Linux currently uses the old metadata structure. For information, see [Section 21.1.10, “Old Metadata Structure Supports Limited Hard Links for a File,” on page 285](#).

#### NSS Multiple Path I/O to Devices

NSS for Linux does not support NSS multiple path I/O to devices. For NSS on Linux, use the same multipath tools you use for Linux traditional file systems.

The devices used by NSS pools can take advantage of multiple path I/O support only when they are assigned to a NetWare server. The NSS multipath priorities and failover preferences do not apply when the devices are on the Linux server.

For information about multiple path I/O for NSS on NetWare, see [Chapter 10, “Managing Multiple Connection Paths to Devices \(NetWare\),” on page 139](#).

## **NSS Pool Snapshots**

NSS on Linux does not support NSS pool snapshots. For NSS on NetWare, pool snapshots preserve a point-in-time copy of specified data in an NSS pool.

NSS pools that are a source pool or a destination pool for NSS pool snapshots on NetWare cannot move cross-platform if you want to keep the pool snapshots. A pool snapshot is no longer available if you move its source pool or destination pool to a Linux server. The snapshot no longer works even after you move the pools back to NetWare.

Before you move an NSS pool cross platform, make sure you delete all of its snapshots stored on other pools and all snapshots for other pools that it might contain.

For information about pool snapshots for NSS on NetWare, see [Chapter 12, “Managing NSS Pool Snapshots \(NetWare\),” on page 167](#).

## **CD and DVD Media and Image Files**

On NetWare, CDs, DVDs, CD and DVD image files, and DOS partitions are mounted as NSS volumes. Disks in USB floppy drives are mounted as local DOS FAT partitions. For information, see the following:

- [Section 13.11, “Activating and Deactivating CDs and DVDs as NSS Volumes \(NetWare\),” on page 193](#)
- [Section 13.12, “Activating and Deactivating CD and DVD Image Files as NSS Volumes \(NetWare\),” on page 194](#)
- [Section 13.13, “Activating and Deactivating DOS Partitions as NSS Volumes \(NetWare\),” on page 195](#)
- [Section 13.17, “Mounting a Disk in a USB Floppy Drive as a Local FAT Volume \(NetWare\),” on page 198](#)

On Linux, these alternative media and partitions are mounted using Linux traditional file systems options. For information, see [“Some Other Supported File Systems” \(http://www.novell.com/documentation/beta/oes/sles\\_admin/data/sec-filestystems-add.html\)](http://www.novell.com/documentation/beta/oes/sles_admin/data/sec-filestystems-add.html) in the *SUSE LINUX Enterprise Server 9 Administration Guide*.

## **Novell Distributed File Services**

NSS on Linux does not support Novell Distributed File Services (DFS). For NSS on NetWare, DFS allows you to move and split NSS volumes without modifying scripts and maps with the new physical location of data.

You cannot create a DFS junction on an NSS volume on Linux. NSS for Linux does not recognize DFS junctions. If the device you moved contains an NSS volume with a DFS junction in it, you cannot follow the existing junction while the volume is on a Linux server.

VLDB can include a GUID for an NSS volume on Linux. This allows you to create DFS junctions on NetWare servers that point to the volumes on the Linux machine. If you use DFS anywhere in the same tree as your NSS on Linux volume, run VLDB Repair on the primary VLDB server to add the volume's GUID to the VLDB.

If you plan dual use of an NSS volume, such as in a mixed cluster with Novell Cluster Services, do not use DFS to move or split NSS volumes while they are on the NetWare server. Perform the physical relocation without using DFS, and update all related scripts and maps with the new path.

For information about DFS for NSS on NetWare, see [Chapter 20, “Using Distributed File Services to Move and Split Volumes \(NetWare\),” on page 261](#).

## **Novell Archive and Version Services**

NSS on Linux does not support Novell Archive and Version Services. For NSS on NetWare, Archive and Version Services is used to save interval-based versions of specified files in an archive database and to make the file versions available to users without the assistance of an IT administrator.

Do not move an NSS pool cross-platform if it contains source volumes with data that is subject to versioning for archival with Novell Archive and Version Services. If you move the pool cross-platform, the versioning jobs for its volumes fail. If you move a source volume cross-platform, you should stop the versioning job for the volume, delete and purge the job to remove the volume's file versions from the archive database.

Do not move an NSS pool cross-platform if it contains the archive volume (archive database and data) for the archive server. Moving the archive volume, which contains the archive database and archived data, causes the ArkManager and ArkSQL to fail.

For information about Novell Archive and Version Services for NSS on NetWare, see the *Novell Archive and Version Services 2.0 for NetWare Administration Guide for OES*.

## **Interoperability of NSS Volumes on OES Linux with Novell Audit 2.0 and Novell NSure Audit 1.0.x**

Filesystem events for NSS volumes on OES Linux cannot be audited with Novell Audit 2.0 or Novell NSure Audit 1.0.x. At this time, the Linux versions of the auditing tools are instrumented to audit eDirectory events, but they are not instrumented to monitor filesystem events. This includes filesystem events for NSS volumes on OES Linux, NCP shares, and other Linux filesystems.

For the most recent information about this issue, see [TID 10099858 \(http://support.novell.com/techcenter/search/Docs/Tids/Solutions/10099858.html\)](http://support.novell.com/techcenter/search/Docs/Tids/Solutions/10099858.html) at Novell Technical Support.

## **6.6.4 Guidelines for Moving Devices Cross-Platform**

You can move devices containing NSS volumes between NetWare servers and OES Linux servers. When you move an unshared device to a different server, you must decommission its volumes in eDirectory for the current server, then recommission them for the new server. For shared NSS pools and volumes, Novell Cluster Services provides this service automatically.

NSS supports device moves between servers with the following platforms:

**Table 6-2** *Cross-Platform Compatibility of NSS Volumes by Operating Platforms*

Source Platforms	Target Platforms	For Information
From NetWare 6.0 servers	<p>To any of these servers:</p> <ul style="list-style-type: none"> <li>• OES SP1 Linux, with a post-move upgrade of the partition format, using the EVMSGUI</li> <li>• NetWare 6.5 SP3 and earlier, with a post-move partitions format upgrade</li> <li>• OES NetWare 6.5 SP4, with a post-move partitions format upgrade and a media format upgrade</li> <li>• OES NetWare, with a post-move partitions format upgrade</li> <li>• OES SP1 NetWare, with a post-move partitions format upgrade and a media format upgrade</li> </ul>	<p>Section 6.8, "Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later," on page 92</p> <p>Section 6.9, "Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare," on page 95</p> <p>Section 6.10, "Moving Clustered Devices with NSS Volumes Cross-Platform," on page 98</p>
<p>From any of these servers:</p> <ul style="list-style-type: none"> <li>• OES Linux and later</li> <li>• NetWare 6.5 SP3 and earlier</li> <li>• NetWare 6.5 SP4 without the media upgrade</li> <li>• OES NetWare (initial release)</li> <li>• OES SP1 NetWare without the media format upgrade</li> </ul> <p>For information about the media format upgrade, see <a href="#">Section 3.1, "Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,"</a> on page 33</p>	<p>To any of these servers:</p> <ul style="list-style-type: none"> <li>• OES Linux and later</li> <li>• NetWare 6.5 SP3 and earlier</li> <li>• NetWare 6.5 SP4 or later, without the media upgrade</li> <li>• OES NetWare (initial release)</li> <li>• OES SP1 NetWare or later, without the media format upgrade</li> </ul>	<p>Section 6.7, "Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux," on page 87</p> <p>Section 6.10, "Moving Clustered Devices with NSS Volumes Cross-Platform," on page 98</p> <p>You can generalize these instructions to fit the platforms of the source and destination servers.</p>
<p>From either of these NetWare servers:</p> <ul style="list-style-type: none"> <li>• NetWare 6.5 SP4 or later, after the media upgrade</li> <li>• OES SP1 NetWare or later, after the media format upgrade</li> </ul> <p>For information, see <a href="#">Section 3.1, "Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,"</a> on page 33.</p>	<p>To either of these NetWare servers:</p> <ul style="list-style-type: none"> <li>• NetWare 6.5 SP4 or later, after the media format upgrade</li> <li>• OES SP1 NetWare or later, after the media format upgrade</li> </ul>	<p>Section 6.7, "Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux," on page 87</p> <p>Section 6.10, "Moving Clustered Devices with NSS Volumes Cross-Platform," on page 98</p> <p>You can generalize these instructions to fit the platforms of the source and destination servers.</p>

## 6.7 Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux

NSS supports moves of devices containing NSS volumes between any servers that support a compatible media format. For information, see [Section 6.6.4, “Guidelines for Moving Devices Cross-Platform,” on page 85](#).

To preserve the NSS pool and volumes on the device when you move it, you must modify the volumes’ Storage objects in eDirectory. You “decommission” the volume by removing its related Storage object from eDirectory for the original server. You “recommission” the volume by creating a new Storage object in eDirectory for the destination server. When moving clustered devices cross-platform, such as in a mixed cluster configuration, Novell Cluster Services automatically manages the Storage object updates to eDirectory.

---

**NOTE:** The “decommission” and “recommission” terminology is used only to illustrate the process; it does not represent a particular technology or tool.

---

---

**IMPORTANT:** Similar handling is necessary when moving devices with NSS pools between any two OES servers (NetWare - Linux, NetWare - NetWare, or Linux - Linux).

---

This section describes how to move devices cross platform from a NetWare 6.5 or later server to an OES Linux or later server:

- [Section 6.7.1, “Prerequisites,” on page 87](#)
- [Section 6.7.2, “Setting Up File Access For Users on the OES Linux Server,” on page 88](#)
- [Section 6.7.3, “Decommissioning Each NSS Pool and Its Volumes on the Original Server,” on page 89](#)
- [Section 6.7.4, “Recommissioning Each NSS Pool and Its Volumes on the Destination Server,” on page 90](#)
- [Section 6.7.5, “Using Scripts to Decommission and Recommission NSS Volumes,” on page 90](#)

### 6.7.1 Prerequisites

The prerequisites in this section apply to moving multiple devices from a NetWare 6.5 or later server to an OES Linux or later server.

---

**IMPORTANT:** When moving a non-clustered device, you must also move any other devices that contribute segments to the NSS pools on the device you are moving.

---

#### Compatibility Issues for Using NSS Volumes Cross-Platform

Before you begin, make sure you understand the [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#).

#### Original NetWare 6.5 or OES NetWare Server

You can move NetWare 6.5 NSS media to an OES Linux and later server if they share the same NSS media format. NSS for OES Linux or later servers do not support the new media format of NetWare 6.5 SP4 and OES SP1 NetWare.



---

**IMPORTANT:** Do not upgrade the media format of any NSS volumes on the device you want to move to the new NSS media format after upgrading the server operating system to NetWare 6.5 SP4 or OES SP1 NetWare.

---

For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,”](#) on page 33.

## Destination OES Linux Server

NSS must be installed on an EVMS-managed device on the destination server. For information, see “[Installing Linux with EVMS as the Volume Manager of the System Device](#)” in the *OES Linux Installation Guide*.

## 6.7.2 Setting Up File Access For Users on the OES Linux Server

Before or after you move an NSS volume from NetWare to Linux, you need to set up file access for users on the OES Linux server.

### Set Up Users in eDirectory

The original server and the destination server can be in the same or different eDirectory trees.

If the destination server is in the same tree as the original server, the file system trustees and trustee rights continue to work after the move.

If the destination server is in a different tree, use eDirectory to enable or reassign affected users for access in the destination tree. For information, see the *Novell eDirectory 8.7.3 Administration Guide* (<http://www.novell.com/documentation/edir873/edir873/data/a2iii88.html>).

### Set Up Protocols and Services

To provide access for users on the OES Linux server, do one or more of the following, depending on your network environment:

- **NCP Server and Services:** Install and configure NCP Server to allow the users to access the volume with the Novell Client or other NCP services. For information, see the *NCP Server for Linux Administration Guide* ([http://www.novell.com/documentation/oes/ncp\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/ncp_lx/data/h9izvdye.html)).
- **Other Protocols and Services:** Install and configure other protocols, such as Samba (CIFS), NFS, or AFP, and enable users with Linux User Management to allow the users to access the volume with the non-NCP protocols and to execute Linux commands and services on the volume.

In this case, if you want to preserve file ownership, such as for configuring and enforcing user space quotas after the move, you must perform this configuration before you move the devices.

For information about enabling users for Linux, see the *Novell Linux User Management Technology Guide* (<http://www.novell.com/documentation/oes/lumadgd/data/front.html>).

For guidelines about users and access, see [Section 6.4, “Access Control Issues for NSS on OES Linux,”](#) on page 74.



## 6.7.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server

For each NSS pool, decommission the pool and its volumes from the original server.

- 1 If you use non-NCP protocols or Linux services for user access on the destination OES Linux server, you must Linux-enable the current users of the volumes before you move the devices if you want file ownership information to be usable after the move, such as for enforcing user space restrictions on a volume.

---



**IMPORTANT:** If you use only NCP Server and NCP services for user access, this step is not necessary.

---

Use one of the following methods to Linux-enable users of the volumes on the device where you want to preserve file ownership information:

- To enable multiple users at once, use the `nambulkadd` command.  
User IDs are automatically refreshed after the enabling process ends.
- To enable a single user at a time, use iManager.

For information, see the [Novell Linux User Management Technology Guide \(http://www.novell.com/documentation/oes/lumadgd/data/front.html\)](http://www.novell.com/documentation/oes/lumadgd/data/front.html).



- 2 Deactivate the pool on the device.
  - 2a In iManager, click *Roles and Tasks* .
  - 2b Click *Storage > Pools*.
  - 2c Browse to select the original server where the NSS pool resides.
  - 2d Select the pool you want to decommission, then click *Deactivate*.
- 3 Remove the eDirectory Storage objects for the NSS pool and each of its volumes.
  - 3a In iManager, click *Roles and Tasks* .
  - 3b Click *eDirectory Administration > Delete Object*.
  - 3c Specify the name and context of the object or objects you want to delete.
  - 3d Click *OK*.
- 4 Repeat **Step 2** and **Step 3** for each pool on the devices you plan to move.
- 5 If you are using DFS in the tree where the original server is located, run the `vldb repair` command.  
On the primary VLDB server, at the command prompt, enter  

```
vldb repair
```

  
This removes a GUID entry for each of the decommissioned volumes from the VLDB.
- 6 Remove or reallocate the devices from the original server. Depending on your storage configuration, this might require a server shutdown.

## 6.7.4 Recommissioning Each NSS Pool and Its Volumes on the Destination Server

For each NSS pool, recommission the pool and its volumes on the destination server.

- 1 Relocate or reassign the devices to the destination server.
- 2 Reboot the destination server to mount the devices.
- 3 If a pool on the devices you moved is not automatically activated, activate the pool.
  - 3a In iManager, click *Roles and Tasks* .
  - 3b Click *Storage > Pools*.
  - 3c Browse to select the destination server.
  - 3d Select the pool, then click *Activate*.
- 4 Create the eDirectory Storage objects for the NSS pool and each of its volumes.
  - 4a In iManager, click *Roles and Tasks* .
  - 4b Click *Storage > Pools*.
  - 4c Browse to select the destination server.
  - 4d Select the pool, then click *Update eDirectory*.
  - 4e In the lower right, select *View Volume Details* to view all volumes on the selected pool.  
iManager opens to the Volumes page with the server and pool preselected.
  - 4f For each volume in the selected pool, select the volume, then click *Update eDirectory*.
  - 4g Repeat **Step 4d** through **Step 4f** for each NSS pool and its volumes.
- 5 Allow the eDirectory tree to stabilize.  
This can take several minutes.
- 6 Run the `vldb repair` command.  
On the primary VLDB server, at the command prompt, enter  

```
vldb repair
```

  
This adds a GUID entry for each of the recommissioned volumes to the VLDB. Having the GUIDs in the VLDB allows you to create DFS junctions on NetWare servers that point to the volumes on the Linux machine. You cannot create a DFS junction on an NSS volume on Linux. If the device you moved contains an NSS volume with a DFS junction in it, you cannot follow the existing junction while the volume is on a Linux server.

## 6.7.5 Using Scripts to Decommission and Recommission NSS Volumes

Scripts are available to automate the process of decommissioning and recommissioning NSS volumes that are not cluster-enabled, see [Decommissioning Script and Recommissioning Script for Moving NSS Volumes Cross-Platform](http://www.novell.com/documentation/oes/script/decom_recom.zip) ([http://www.novell.com/documentation/oes/script/decom\\_recom.zip](http://www.novell.com/documentation/oes/script/decom_recom.zip)). This `decom_recom.zip` file contains two Perl scripts:

- **decom.pl:** The decommissioning script deactivates the specified pool, removes eDirectory Storage objects for a specified NSS pool and each of its volumes on the original server, then it repairs the VLDB, if it exists, to remove the volumes' information from the VLDB. You provide the pool name, and the script automatically gets the list of volumes on the pool.

- **recom.pl:** The recommissioning script activates the specified pool, creates eDirectory Storage objects for a specified NSS pool and each of its volumes on the destination server, then it repairs the VLDB, if it exists, to add the volumes' information to the VLDB. You provide the pool name, and the script automatically gets the list of volumes on the pool.

The scripts support moving NSS volumes on OES NetWare or NetWare 6.5 SP3 to OES Linux. You can modify the scripts to move volumes between any two non-clustered OES servers:

- NetWare to Linux
- Linux to NetWare
- NetWare to NetWare
- Linux to Linux

### Decommissioning NSS Pools on the Original Server with `decom.pl`

For each NSS pool, decommission the pool and its volumes from the original server.

- 1 If you use non-NCP protocols or Linux services for user access on the destination OES Linux server, you must Linux-enable the current users of the volumes before you move the devices if you want file ownership information to be usable after the move, such as for enforcing user space restrictions on a volume.

---

**IMPORTANT:** If you use only NCP Server and NCP services for user access, this step is not necessary.

---

Use one of the following methods to Linux-enable users of the volumes on the device where you want to preserve file ownership information:

- To enable multiple users at once, use the `nambulkadd` command.  
User IDs are automatically refreshed after the enabling process ends.
- To enable a single user at a time, use iManager.

For information, see the *Novell Linux User Management Technology Guide* (<http://www.novell.com/documentation/oes/lumadgd/data/front.html>).

- 2 For each NSS pool on the device you are moving, run the `decom.pl` script and specify the name of the pool to decommission.
- 3 Remove or reallocate the devices from the original server. Depending on your storage configuration, this might require a server shutdown.

### Recommissioning NSS Pools on the Destination Server with `recom.pl`

For each NSS pool, recommission the pool and its volumes on the destination server.

- 1 Relocate or reassign the devices to the destination server.
- 2 Reboot the destination server to mount the devices.
- 3 For each NSS pool on the device you moved, run the `recom.pl` script and specify the name of the pool to recommission.

## 6.8 Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later

NSS supports moves of devices containing NSS volumes from a NetWare 6.0 server to an OES SP1 Linux or later server. For information, see [Section 6.6.4, “Guidelines for Moving Devices Cross-Platform,” on page 85](#).

To preserve the NSS pool and volumes on the device when you move it, you must modify the volumes’ Storage objects in eDirectory. You “decommission” the volume by removing its related Storage object from eDirectory for the original server. You “recommission” the volume by creating a new Storage object in eDirectory for the destination server. When moving clustered devices cross-platform, such as in a mixed cluster configuration, Novell Cluster Services automatically manages the Storage object updates to eDirectory.

---

**NOTE:** The “decommission” and “recommission” terminology is used only to illustrate the process; it does not represent a particular technology or tool.

---

---

**IMPORTANT:** Similar handling is necessary when moving devices with NSS pools between any two OES servers (NetWare - Linux, NetWare - NetWare, or Linux - Linux).

---

This section describes how to move devices cross platform from a NetWare 6.0 server to an OES SP1 Linux or later server:

- [Section 6.8.1, “Prerequisites,” on page 92](#)
- [Section 6.8.2, “Setting Up File Access For Users on the OES Linux Server,” on page 93](#)
- [Section 6.8.3, “Decommissioning Each NSS Pool and Its Volumes on the Original Server,” on page 93](#)
- [Section 6.8.4, “Recommissioning Each NSS Pool and Its Volumes on the Destination Server,” on page 94](#)

### 6.8.1 Prerequisites

The prerequisites in this section apply to moving multiple devices from a NetWare 6.0 server to an OES SP1 Linux server.

---

**IMPORTANT:** When moving a non-clustered device, you must also move any other devices that contribute segments to the NSS pools on the device you are moving.

---

#### Compatibility Issues for Using NSS Volumes Cross-Platform

Before you begin, make sure you understand the [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#).

#### Destination OES Linux Server

NSS must be installed on an EVMS-managed device on the destination server. For information, see [“Installing Linux with EVMS as the Volume Manager of the System Device”](#) in the *OES Linux Installation Guide*.

## 6.8.2 Setting Up File Access For Users on the OES Linux Server

Before or after you move an NSS volume from NetWare 6.0 to OES SP1 Linux, you need to set up file access for users on the OES Linux server.

### Set Up Users in eDirectory

The original server and the destination server can be in the same or different eDirectory trees.

If the destination server is in the same tree as the original server, the file system trustees and trustee rights continue to work after the move.

If the destination server is in a different tree, use eDirectory to enable or reassign affected users for access in the destination tree. For information, see the *Novell eDirectory 8.7.3 Administration Guide* (<http://www.novell.com/documentation/edir873/edir873/data/a2iii88.html>).

### Set Up Protocols and Services

To provide access for users on the OES Linux server, do one or more of the following, depending on your network environment:

- **NCP Server and Services:** Install and configure NCP Server to allow the users to access the volume with the Novell Client or other NCP services. For information, see the *NCP Server for Linux Administration Guide* ([http://www.novell.com/documentation/oes/ncp\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/ncp_lx/data/h9izvdye.html)).
- **Other Protocols and Services:** Install and configure other protocols, such as Samba (CIFS), NFS, or AFP, and enable users with Linux User Management to allow the users to access the volume with the non-NCP protocols and to execute Linux commands and services on the volume.

In this case, if you want to preserve file ownership, such as for configuring and enforcing user space quotas after the move, you must perform this configuration before you move the devices. For information about enabling users for Linux, see the *Novell Linux User Management Technology Guide* (<http://www.novell.com/documentation/oes/lumadgd/data/front.html>).

For guidelines about users and access, see [Section 6.4, “Access Control Issues for NSS on OES Linux,” on page 74](#).

## 6.8.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server

For each NSS pool, decommission the pool and its volumes from the original server.

- 1 On the NetWare 6.0 server, if the device you want to move contains any mirrored partitions, you must first break the mirror (remove all but 1 segment) before upgrading.
- 2 If you use non-NCP protocols or Linux services for user access on the destination OES Linux server, you must Linux-enable the current users of the volumes before you move the devices if you want file ownership information to be usable after the move, such as for enforcing user space restrictions on a volume.

---

**IMPORTANT:** If you use only NCP Server and NCP services for user access, this step is not necessary.

---

Use one of the following methods to Linux-enable users of the volumes on the device where you want to preserve file ownership information:

- To enable multiple users at once, use the `nambulkadd` command.  
User IDs are automatically refreshed after the enabling process ends.
- To enable a single user at a time, use iManager.

For information, see the *Novell Linux User Management Technology Guide* (<http://www.novell.com/documentation/oes/lumadgd/data/front.html>).

- 3** Deactivate the pool: At a server console, enter

```
nss /pooldeactivate=pool
```

- 4** Remove the eDirectory Storage objects for the NSS pool and each of its volumes.

**4a** In iManager, click *Roles and Tasks* .

**4b** Click *eDirectory Administration > Delete Object*.

**4c** Specify the name and context of the object or objects you want to delete.

**4d** Click *OK*.

- 5** Repeat **Step 2** and **Step 3** for each pool on the devices you plan to move.

- 6** If you are using DFS in the tree where the original server is located, run the `vldb repair` command.

On the primary VLDB server, at the command prompt, enter

```
vldb repair
```

The VLDB Repair removes a GUID entry for each of the decommissioned volumes from the VLDB.

- 7** Remove or reallocate the devices from the original server. Depending on your storage configuration, this might require a server shutdown.

## 6.8.4 Recommissioning Each NSS Pool and Its Volumes on the Destination Server

For each NSS pool, recommission the pool and its volumes on the destination server.

- 1** Relocate or reassign the devices to the destination server.
- 2** Reboot the destination server to mount the devices.
- 3** If a pool on the devices you moved is not automatically activated, activate the pool. At a server console, enter

```
nsscon
```

```
nss /poolactivate=pool
```

- 4** Create the eDirectory Storage objects for the NSS pool and each of its volumes.

**4a** In iManager, click *Roles and Tasks* .

**4b** Click *Storage > Pools*.

**4c** Browse to select the destination server.

- 4d** Select the pool, then click *Update eDirectory*.
- 4e** In the lower right, select *View Volume Details* to view all volumes on the selected pool.  
iManager opens to the Volumes page with the server and pool preselected.
- 4f** For each volume in the selected pool, select the volume, then click *Update eDirectory*.
- 4g** Repeat **Step 4d** through **Step 4f** for each NSS pool and its volumes.
- 5** Allow the eDirectory tree to stabilize.  
This can take several minutes.
- 6** Run the `vldb repair` command.  
On the primary VLDB server, at the command prompt, enter  
`vldb repair`  
The VLDB Repair adds a GUID entry for each of the recommissioned volumes to the VLDB. Having the GUIDs in the VLDB allows you to create DFS junctions on NetWare servers that point to the volumes on the Linux machine. You cannot create a DFS junction on an NSS volume on Linux. If the device you moved contains an NSS volume with a DFS junction in it, you cannot follow the existing junction while the volume is on a Linux server.
- 7** Make sure the devices are up and working as expected on the Linux System.
- 8** For each device, you can optionally upgrade the format for its partitions (or segments, as specified in EVMSGUI) on the device, or leave the partitions in the old format.  
Upgrading the partitions on the device to the NetWare 6.5 and later partition format is strongly recommended, but it is not necessary unless you want to mirror a partition on the device. Before you can create the mirror, you must upgrade the partition format of the partitions on the device you moved.
  - 8a** At a server console, open the EVMSGUI by entering  
`evmsgui`
  - 8b** Right-click the segment, then select *Upgrade*.
  - 8c** Repeat **Step 8b** for every partition on the device that was moved to the OES Linux server.
  - 8d** Click *Save* to save the changes.
- 9** If you want to mirror the upgraded partition from **Step 8**, see **Section 9.16, “Mirroring an Existing Pool with NSSMU,”** on page 136.

## 6.9 Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare

NSS supports moves of devices containing NSS volumes from a NetWare 6.0 server to a NetWare 6.5 or OES NetWare server. For information, see **Section 6.6.4, “Guidelines for Moving Devices Cross-Platform,”** on page 85.

To preserve the NSS pool and volumes on the device when you move it, you must modify the volumes’ Storage objects in eDirectory. You “decommission” the volume by removing its related Storage object from eDirectory for the original server. You “recommission” the volume by creating a new Storage object in eDirectory for the destination server. When moving clustered devices cross-platform, such as in a mixed cluster configuration, Novell Cluster Services automatically manages the Storage object updates to eDirectory.

---

**NOTE:** The “decommission” and “recommission” terminology is used only to illustrate the process; it does not represent a particular technology or tool.

---

---

**IMPORTANT:** Similar handling is necessary when moving devices with NSS pools between any two OES servers (NetWare - Linux, NetWare - NetWare, or Linux - Linux).

---

This section describes how to move devices cross platform from a NetWare 6.0 server to a NetWare 6.5 or OES NetWare server:

- [Section 6.9.1, “Prerequisites,” on page 96](#)
- [Section 6.9.2, “Setting Up File Access For Users on the Destination Server,” on page 96](#)
- [Section 6.9.3, “Decommissioning Each NSS Pool and Its Volumes on the Original Server,” on page 96](#)
- [Section 6.9.4, “Recommissioning Each NSS Pool and Its Volumes on the Destination Server,” on page 97](#)

## 6.9.1 Prerequisites

When moving a non-clustered device, you must also move any other devices that contribute segments to the NSS pools on the device you are moving.

## 6.9.2 Setting Up File Access For Users on the Destination Server

Before or after you move a device, you need to set up file access for users on the destination server. The original server and the destination server can be in the same or different eDirectory trees.

If the destination server is in the same tree as the original server, the file system trustees and trustee rights continue to work after the move.

If the destination server is in a different tree, use eDirectory to enable or reassign affected users for access in the destination tree. For information, see the *Novell eDirectory 8.7.3 Administration Guide* (<http://www.novell.com/documentation/edir873/edir873/data/a2iii88.html>).

## 6.9.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server

For each NSS pool, decommission the pool and its volumes from the original server.

- 1** Deactivate the pool: At a server console, enter

```
nss /pooldeactivate=pool
```

- 2** Remove the eDirectory Storage objects for the NSS pool and each of its volumes.

- 2a** In iManager, click *Roles and Tasks* .

- 2b** Click eDirectory Administration > Delete Object.

- 2c** Specify the name and context of the object or objects you want to delete.

- 2d** Click *OK*.



- 3 Repeat **Step 2** and **Step 3** for each pool on the devices you plan to move.
- 4 If you are using DFS in the tree where the original server is located, run VLDB Repair.  
On the primary VLDB server, at the command prompt, enter  


```
vldb repair
```

  
The VLDB Repair removes a GUID entry for each of the decommissioned volumes from the VLDB.
- 5 Remove or reallocate the devices from the original server. Depending on your storage configuration, this might require a server shutdown.

## 6.9.4 Recommissioning Each NSS Pool and Its Volumes on the Destination Server

For each NSS pool, recommission the pool and its volumes on the destination server.

- 1 Relocate or reassign the devices to the destination server.
- 2 Reboot the destination server to mount the devices.
- 3 If a pool on the devices you moved is not automatically activated, activate the pool. At a server console, enter  

```
nss /poolactivate=pool
```
- 4 Create the eDirectory Storage objects for the NSS pool and each of its volumes.
  - 4a In iManager, click *Roles and Tasks* .
  - 4b Click *Storage > Pools*.
  - 4c Browse to select the destination server.
  - 4d Select the pool, then click *Update eDirectory*.
  - 4e In the lower right, select *View Volume Details* to view all volumes on the selected pool.  
iManager opens to the Volumes page with the server and pool preselected.
  - 4f For each volume in the selected pool, select the volume, then click *Update eDirectory*.
  - 4g Repeat **Step 4d** through **Step 4f** for each NSS pool and its volumes.
- 5 Allow the eDirectory tree to stabilize.  
This can take several minutes.
- 6 Run VLDB Repair.  
On the primary VLDB server, at the command prompt, enter  

```
vldb repair
```

  
The VLDB Repair adds a GUID entry for each of the recommissioned volumes to the VLDB. Having the GUIDs in the VLDB allows you to create DFS junctions on NetWare servers that point to the volumes on the Linux machine. You cannot create a DFS junction on an NSS volume on Linux. If the device you moved contains an NSS volume with a DFS junction in it, you cannot follow the existing junction while the volume is on a Linux server.
- 7 Make sure the devices are up and working as expected on the destination server.
- 8 For each device, you can optionally upgrade the format for its partitions on the device, or leave the partitions in the old format.

Upgrading the partitions on the device to the NetWare 6.5 and later partition format is strongly recommended, but it is not necessary unless you want to mirror a partition on the device. Before you can create the mirror, you must upgrade the partition format of the partitions on the device you moved.

At a server console, enter

```
mm upgrade partitions
```

- 9 If you upgrade the partitions, we strongly recommend that you consider upgrading the media format to the format for OES SP1 NetWare or NetWare 6.5 SP4.

For information, see [Chapter 3, “Upgrading the Media Format \(NetWare\),” on page 33](#).

- 10 If you want to mirror the upgraded partition from [Step 9](#), see [Section 9.16, “Mirroring an Existing Pool with NSSMU,” on page 136](#).

## 6.10 Moving Clustered Devices with NSS Volumes Cross-Platform

You can use a shared NSS data pool and volume in a mixed cluster configuration, using Novell Cluster Services for Linux as long as all nodes in the cluster support the media format of the device. OES Linux and later servers do not support the new media format of NetWare 6.5 SP4 and OES SP1 NetWare. For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,” on page 33](#).

For information, see the *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html)).

## 6.11 Upgrading NetWare 5.x NSS Volumes and NetWare Traditional Volumes to NSS Volumes

Use Volume Copy Upgrade to upgrade legacy NetWare 5.x NSS volumes and NetWare Traditional volumes to OES NetWare NSS volumes. For information, see [“Upgrading Legacy NSS and NetWare Traditional Volumes” on page 399](#).

After you upgrade legacy NSS and Traditional volumes to OES NetWare, you can move the volumes cross-platform to Linux. For information, see the following:

- [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#)
- [Section 6.10, “Moving Clustered Devices with NSS Volumes Cross-Platform,” on page 98](#)

# Managing Devices

# 7

Novell® Storage Services™ is the default file system for NetWare® 6 and later. The Web-browser-based iManager Storage plug-in is the primary storage management interface for devices on your Novell OES server on both Linux and NetWare kernels. For information about the plug-in, see [Section 4.1, “Novell iManager Storage Plug-In,” on page 39](#)

You can also use the console-based NSS Management Utility (NSSMU) to manage your NSS file system. You can create pools, volumes, iSCSI partitions for iSCSI disk servers, software RAID devices, and other NSS management functions. For information, see [Section 4.2, “NSS Management Utility,” on page 52](#).

You must use the NSSMU interface to create iSCSI disk partitions for an iSCSI disk server. An iSCSI disk server requires iSCSI target software. For information about managing and using iSCSI devices, see the *iSCSI 1.1.3 Administration Guide for NetWare 6.5* ([http://www.novell.com/documentation/iscsi1\\_nak/iscsi/data/h4hgu4hs.html](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html)).

- [Section 7.1, “Understanding Devices,” on page 99](#)
- [Section 7.2, “Making Devices Available to EVMS \(Linux\),” on page 100](#)
- [Section 7.3, “Viewing Devices on a Server,” on page 101](#)
- [Section 7.4, “Viewing Details for a Device,” on page 102](#)
- [Section 7.5, “Scanning for Devices,” on page 103](#)
- [Section 7.6, “Initializing a Disk,” on page 104](#)
- [Section 7.7, “Sharing Devices,” on page 104](#)
- [Section 7.8, “Viewing Partitions on a Device,” on page 106](#)
- [Section 7.9, “Viewing Pools on a Device,” on page 107](#)
- [Section 7.10, “What’s Next,” on page 107](#)

## 7.1 Understanding Devices

A device is the physical or virtual storage media available to a server. Physical devices include direct attached storage devices (such as server disks, disk arrays, hardware RAID (redundant arrays of independent disks) devices, and SAN (storage area network) devices. Virtual devices include software RAID devices and iSCSI devices. The following table discusses the management methods to use with NSS for different types of devices:

**Table 7-1** *Management Methods for Device Types*

Type of Device	Management Method
Physical device	To manage physical devices: In iManager, click <i>Storage &gt; Devices</i> .

Type of Device	Management Method
Software RAID device	<p>To create and manage software RAID devices: In iManager, click <i>Storage &gt; Software RAID</i>s. For information, see <a href="#">“Managing Software RAID Devices” on page 117</a>.</p> <p>After creating the device, you can add pools and volumes as you would for a physical device.</p> <p>To create a software RAID 10 pool, use NSSMU. For information, see <a href="#">Section 9.17, “Creating a Software RAID 10 with NSSMU,” on page 136</a>.</p>
iSCSI device	<p>An iSCSI disk server requires iSCSI target software. To create and manage iSCSI partitions for an iSCSI disk server: At the server console, enter</p> <pre>nssmu</pre> <p>In NSSMU, click <i>Partition Management</i>, press <i>Insert</i> to open the Create dialog, then allocate space to the partition</p> <p>For information about using NSSMU, see <a href="#">Section 4.2, “NSS Management Utility,” on page 52</a>.</p> <p>To access iSCSI storage devices for a file and print server, you first must connect to an iSCSI disk server. This requires iSCSI initiator software. After connecting, you can view the devices on the disk server in the Devices list and add NSS pools and volumes as you would with any device.</p> <p>For information about managing and using iSCSI devices, see the <a href="#">iSCSI 1.1.3 Administration Guide for NetWare 6.5</a> (<a href="http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html">http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html</a>).</p>

## 7.2 Making Devices Available to EVMS (Linux)

OES Linux uses Linux Volume Management (LVM) for volume management, but NSS requires the Enterprise Volume Management System (EVMS). The 2.6 kernel prevents multiple volume managers from managing the same device. As a result, EVMS might display free space that it cannot actually manage.

To prevent display of space that is unavailable to EVMS, edit the `/etc/evms.conf` file to add the device to the exclude list of the `sysfs_devices` section.

For example, if your device is your boot device and contains the root volume, it cannot be used to create pools on. Edit the `/etc/evms.conf` file as follows:

```
sysfs_devices {
...
exclude = [hda]
}
```

To make a device that is being managed by another device manager available to EVMS, you can use one of the following options:

- Install OES Linux with EVMS as the default volume manager

See [Mounting Your Root Filesystem through EVMS \(http://evms.sourceforge.net/install/root.html\)](http://evms.sourceforge.net/install/root.html) for more information.

- Patch the 2.6 kernel

See [BD-Claim Patch \(http://evms.sourceforge.net/install/kernel.html#bdclaim\)](http://evms.sourceforge.net/install/kernel.html#bdclaim) for more information.

- Convert LVM volumes to EVMS volumes after OES Linux installation

See [Converting to EVMS \(http://evms.sourceforge.net/convert.html\)](http://evms.sourceforge.net/convert.html) for more information.

## 7.3 Viewing Devices on a Server

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

Depending on the number of devices, it can take a few seconds to display the list of devices. Avoid clicking again in the page until it refreshes and displays the Devices list.

For an overview of the subtasks available from this page, see [“Devices” on page 51](#).

[Storage Management](#) ► **Devices**

### Device Management



Initialize and maintain a wide selection of physical storage devices and software RAID devices available to this server. Enable device sharing for those devices that you plan to use in a high-availability cluster configuration.

Server:

Devices:	Details:
<div><div>Initialize Disk</div><div>Multipath...</div><div>Set Default Path</div><div>Reset Registry</div></div> <div><div>0x21</div><div><b>0x22</b></div><div>0x23</div><div>0x5</div><div>0x9</div><div>0xc</div><div>0xf</div></div>	<div><div>Name: <b>0x22</b></div><div>Description: <b>SEAGATE ST34573W</b></div><div>ModuleID: <b>rev:5764</b></div><div>Adaptor: <b>354</b></div><div>Device: <b>1</b></div><div>GUID: <b>0:0</b></div><div><input type="checkbox"/> Shareable for Clustering</div></div> <div><div>Capacity: <b>74.50 GB</b></div><div>Used Space: <b>6.00 GB</b></div><div>Free Space: <b>68.50 GB</b></div><div>Pools: <input type="text" value="P1"/> </div><div>Number of Pools: <b>2</b></div><div>Partitions: <input type="text" value="DOS - 0x6"/> </div></div>

## 7.4 Viewing Details for a Device

The Details field in the Devices Management page displays information about each device in the Devices list, as described in the following table.

**Table 7-2** *Explanation of Device Details*

Device Detail	Description
<i>Name</i>	The device name assigned by NetWare.
<i>Description</i>	The physical description of the storage object. This is the physical device identification. For software RAIDs, the description might be RAID0, RAID1, or RAID5.
<i>Module ID</i>	The Novell driver identification number for the module, which indicates what type of driver you are using.
<i>Adapter</i>	The sequential number registered for the adapter. The number indicates which drive is on which adapter.
<i>Device</i>	<p>The server address of the selected device.</p> <p>For a SCSI device, the first number is the SCSI ID number and the second number is the Logical Unit Number (LUN), which is commonly used for backup tape drives.</p> <p>For an IDE device, the first number is the channel number and the second number indicates if the device is a master (0) or a slave (1).</p>
<i>GUID</i>	The Global Unique Identifier (GUID) number that NSS assigns to the storage object. This number is necessary so your file system can locate the specific device.
<i>Shareable for Clustering</i>	The attribute of a device that indicates whether the selected device can be shared by multiple computers in a cluster solution.
<i>Capacity</i>	The total available storage space of the selected device.
<i>Used Space</i>	The amount of space on the device that is currently in use by partitions, including NSS partitions, Traditional NetWare partitions, and non-NetWare partitions such as DOS partitions.
<i>Free Space</i>	The total amount of space on the device that is currently not in use.
<i>Pools</i>	The drop-down list shows all pools that exist on this device. To view a pool's details or to manage a pool, select the pool from the list, then click the <i>View Details</i> icon to go to the Pool Management page for that pool.
<i>Number of Pools</i>	The total number of pools that use this device.
<i>Partitions</i>	The drop-down list shows all partitions that exist on this device. To view a partition's details, select the partition from the list, then click the <i>View Details</i> icon to go to the Partition Information page for that partition.
<i>Mirror Status</i>	<p>For a RAID1 device, this field shows its status:</p> <ul style="list-style-type: none"><li>• <b>In Sync:</b> The mirror group is fully synchronized.</li><li>• <b>Partial Sync:</b> The mirror group is only partially synchronized.</li><li>• <b>Not Mirrored:</b> The device is not mirrored (only one partition).</li></ul>

To view a device's details:

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes the devices are listed in the Devices list. Depending on the number of devices, this can take several seconds. Wait for the page to load before moving to the next step.

- 3 Select a device to view its details.

The page must refresh to display the details, which might take several seconds.

## 7.5 Scanning for Devices

NetWare typically recognizes all devices on reboot or after you create them. If you add devices to a server and NetWare does not automatically detect them, you might need to scan for devices.

- 1 In iManager, click *Storage > Scan for Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

The browser opens to the Scan Devices Page.

[Storage Management](#) ► **Scan for Devices**

### Scan for Devices



Scan for devices attached to this server, if not found automatically by the media manager. Scanning can take several minutes to complete. Allow an elapsed time of about 30 seconds for each adapter, depending on the number of devices on the adapter.

Server:

Continue

Cancel

After iManager connects to the server, it scans for devices and displays them in the Devices list. The scan can take several seconds, depending on the number of adapters and disks on your systems. Click *Cancel* at any time to back out of the process.

- 3 View the list of devices.

This list is only viewable; there are no actions to make from this page.

- 4 When you are done:

- Click *Continue* to go to the Storage role main page.
- Click *Storage > Devices* in *Roles and Tasks* to manage the devices.

## 7.6 Initializing a Disk

On the Devices Management page, the Initialize Disk option initializes the selected device and completely removes all the partitions it contains. All the data stored on the device is lost. If the device contains a partition of an NSS pool, a Traditional volume, or a software RAID device, the Initialize command also deletes data on all of the partitions of the entire pool, volume, or device, even if they reside on separate devices.

---

**WARNING:** Do not initialize the device that contains your `sys :` volume. Initializing the `sys :` volume destroys the operating system and all the data in it.

---

This option is disabled (dimmed) if the selected device contains any of the following:

- System pool or `sys :` volume
- DOS partition
- Software RAID 1 (mirrored) device

It can also be disabled if there is no space available based on each partition's size, or if you already have the maximum number of partitions allocated in a software RAID device.

To initialize a disk:

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 In the Devices list, select a device.

- 4 Click *Initialize Disk*.

## 7.7 Sharing Devices

Storage devices that exist in a storage area network (SAN) can be shared by multiple servers in a cluster. For more information about clustering, see the following:

- *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html))
- *OES Novell Cluster Services Administration Guide for NetWare* ([http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html))

---

**IMPORTANT:** The system hardware does not specify disk drives to come up automatically as *Shareable for Clustering* or *Not Shareable for Clustering*. You must manually set this value for each device, according to the configuration of your storage system.

---



Making a device shareable enables device sharing for those devices in high-availability clusters that you want to be part of a shared-disk storage solution. If the *Shareable for Clustering* option is enabled (selected), the selected storage device can be shared by multiple computers in a cluster.

☒ Shareable for Clustering

Marking a device as shareable for clustering sets all of the pools on this device to shareable. If any of these pools span multiple devices, you must make sure that each device is set to the same share state as this one, or the device can become unusable.

If this device is a member of a software RAID device, marking a device as shareable for clustering also sets all the other member devices as shareable for clustering.

## 7.7.1 Planning for Device Sharing

- Do not assign a device as *Shareable for Clustering* if it is not capable of being shareable.
- If you create a pool that spans multiple storage devices, each of the member devices in that pool must either be marked as *Shareable for Clustering* or as *Not Shareable for Clustering*.
- The system pool and all of its member devices cannot be marked as *Shareable for Clustering*.
- Devices that contain Traditional partitions and volumes cannot be marked as *Shareable for Clustering*.
- If you assign partitions to a software RAID device, all the devices for those member partitions must either be marked as *Shareable for Clustering* or as *Not Shareable for Clustering*.

## 7.7.2 Configuring the Device's Share State

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of devices appears in the Devices list.

- 3 In the Devices list, select a device.

After the page refreshes, information about the selected device appears in the Devices list.

- 4 Depending on the current state of the device:

- To set a device's share state to On, select the *Shareable for Clustering* check box, then click *Apply* or click *OK*.
- To set a device's share state to Off, deselect the *Shareable for Clustering* check box, then click *Apply* or click *OK*.

If you click *Apply*, iManager saves the change and remains on the device page. If you click *OK*, iManager saves the change and takes you to the main Storage page. If you do not click *Apply* or *OK*, the setting is not implemented; the device remains not shareable.

## 7.8 Viewing Partitions on a Device

In NetWare 6.5 and later, NSS abstracts all partition creation and deletion in iManager; there are no actions to perform on partitions. For information about partitions, see [“Managing Partitions” on page 109](#).

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

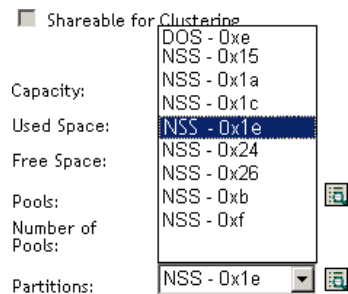
- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 In the Devices list, select a device.

- 4 In the *Details* area, click the arrow on the *Partitions* drop-down list to expand it.



- 5 Select a partition, then click *View Details*.

This opens the Partition Management page. It displays a list of all the partitions that currently exist on the selected device.

[Storage Management](#) > [Devices](#) > [Partitions](#)

### Partition Management

Server:

Partitions on: 0x2						
<a href="#">Details</a>						
<input type="checkbox"/>	ID ▾	Type ▾	Status ▾	Device Name ▾	Pool Name ▾	Size
<input type="checkbox"/>	0x6	DOS	In Use	0x2		2.00 GB
<input type="checkbox"/>	0x8	NSS	In Use	0x2	SYS	3.91 GB
<input type="checkbox"/>	0x11	NSS	In Use	0x2	ARK_POOL	2.79 GB
<input type="checkbox"/>	0x14	Virtual Device	In Use	0x2		100.02 MB
<input type="checkbox"/>	0x17	NSS	In Use	0x2	TEST	95.37 MB
<input type="checkbox"/>	0x20	Virtual Device	In Use	0x2		500.02 MB
<input type="checkbox"/>	0x28	Unknown	In Use	0x2		30.00 GB
<input type="checkbox"/>	0x2b	NSS	In Use	0x2	VOL4	4.66 GB
OK						

To view details about a partition:

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46.](#)

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46.](#)

A list of devices appears in the Devices list.

- 3 In the Devices list, select a device.
- 4 In the *Details* area, click the arrow on the *Partitions* drop-down list to expand it.
- 5 Select a partition, then click *View Details*.

This opens the Partition Management page and displays a list of all the partitions that currently exist on the selected device.

- 6 Select a partition from the Partitions list, then click *Details* to view its details.

## 7.9 Viewing Pools on a Device

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46.](#)

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46.](#)

A list of devices appears in the Devices list.

- 3 In the Devices list, select a device.
- 4 In the *Details* area, click the arrow on the *Pools* drop-down list to expand it.



- 5 Select a pool, then click *View Details*.

This opens the Pool Management page where you can view the details of the pool and manage it.

For information about pool management, see [“Managing NSS Pools” on page 147.](#)

## 7.10 What’s Next

You can create software RAID devices to use for pools. For information, see [Chapter 9, “Managing Software RAID Devices,” on page 117.](#)

If your server provides multiple I/O paths between the server and its storage devices, configure the primary path and path priorities for I/O failover for each device. For information, see [Chapter 10, “Managing Multiple Connection Paths to Devices \(NetWare\),” on page 139.](#) When you are done, continue with creating pools on your devices.

You can now create pools of storage on the devices you created. For information, see [Chapter 11, “Managing NSS Pools,” on page 147.](#)



# Managing Partitions

# 8

Novell® Storage Services™ automatically manages partitions you use to build software RAID devices and NSS pools, whether you create RAIDs and pools in the Storage plug-in for iManager or the NSS Management Utility.

- [Section 8.1, “Understanding Partitions,” on page 109](#)
- [Section 8.2, “Creating NSS Partitions,” on page 111](#)
- [Section 8.3, “Viewing a List of Partitions,” on page 112](#)
- [Section 8.4, “Viewing Details for a Partition,” on page 114](#)
- [Section 8.5, “Labeling a Partition,” on page 114](#)
- [Section 8.6, “Deleting a Partition,” on page 114](#)
- [Section 8.7, “Adding Other Types of File System Partitions \(Linux\),” on page 116](#)

## 8.1 Understanding Partitions

- [Section 8.1.1, “NSS Partitions,” on page 109](#)
- [Section 8.1.2, “Viewing Partitions on a Device,” on page 110](#)
- [Section 8.1.3, “Understanding Types of Partitions,” on page 110](#)
- [Section 8.1.4, “Understanding Partition Details,” on page 111](#)

### 8.1.1 NSS Partitions

A partition is a logical division of a physical hard drive. In NetWare® 6.5 and later, NSS abstracts all partition creation and deletion in iManager and NSSMU through the Pools page and Software RAIDs page. When you create NSS pools or NSS software RAID devices, NSS automatically creates the NSS partitions on the devices you specify. In NetWare 6.5 SP2 and later, you can view and label these NSS partitions from the Partitions page.

Partitions are managed below iManager and NSSMU by the following partition management tools:


Platform	Partition Management Tool
Linux	<p>On Linux, NSS manages a pool's partitions via the Enterprise Volume Management System (EVMS). EVMS discovers and mounts the partition in the <code>/dev/evms/</code> directory.</p> <hr/> <p><b>IMPORTANT:</b> If you create and mount pools from the command line, you must modify your mount point to the partition's new apparent location under <code>/dev/evms/</code>.</p> <hr/>
NetWare	<p>On NetWare, NSS manages a pool's partitions via the Media Manager.</p>



## 8.1.2 Viewing Partitions on a Device

The Partitions page in iManager reports the partitions that it finds on your system and information about them. Select a device to view information about the partition in the Details area.

**Figure 8-1** Example Partitions Management Page

Storage Management ► Devices ► Partitions

Partition Management 

Server:   

Partitions on: 0x2

[Details](#)

<input type="checkbox"/>	ID ▼	Type ▼	Status ▼	Device Name ▼	Pool Name ▼	Size
<input type="checkbox"/>	0x6	DOS	In Use	0x2		2.00 GB
<input type="checkbox"/>	0x8	NSS	In Use	0x2	SYS	3.91 GB
<input type="checkbox"/>	0x11	NSS	In Use	0x2	ARK_POOL	2.79 GB
<input type="checkbox"/>	0x14	Virtual Device	In Use	0x2		100.02 MB
<input type="checkbox"/>	0x17	NSS	In Use	0x2	TEST	95.37 MB
<input type="checkbox"/>	0x20	Virtual Device	In Use	0x2		500.02 MB
<input type="checkbox"/>	0x28	Unknown	In Use	0x2		30.00 GB
<input type="checkbox"/>	0x2b	NSS	In Use	0x2	VOL4	4.66 GB

## 8.1.3 Understanding Types of Partitions

The following table describes the variety of partition types that iManager reports as being on your system. The Linux partition types are found only on OES Linux servers, not on OES NetWare servers.

The Partitions function in iManager and NSSMU is intended simply as a reporting tool so that you can see the types of partitions that are being virtualized by higher-level storage entities such as pools or software RAID devices. Generally, you cannot create or modify partitions with iManager or NSSMU tools. On NetWare, NSSMU allows you to create an NSS partition, but this capability is seldom used. For NSS on either platform, NSS partitions are created for you automatically when you create a pool. There might be multiple NSS partitions that are aggregated and managed underneath the single pool of space. Similarly, the tools creates a Virtual Device automatically when you create a RAID device.

---

**IMPORTANT:** You cannot create any of the reported partition types in iManager; you can only view and label NSS partitions from the Partitions page.

---

**Table 8-1** Explanation of Partition Types

Partition Type	Description
Cluster Service	A partition that NetWare uses to monitor cluster connectivity and services; it appears only in shared devices in the cluster.
DOS	A conventional DOS partition that NetWare uses for installing and upgrading NetWare on the server.
EXT3	The partition type for Linux Extended File System 3.

Partition Type	Description
<i>iSCSI</i>	A partition in a target disk server in an iSCSI storage area network; it appears as an iSCSI device to file servers with iSCSI initiator software.
<i>NSS</i>	The primary partition type for NSS file systems.
<i>Reiser</i>	The partition type for Linux Reiser file systems.
<i>System Configuration</i>	A vendor-specific partition for maintaining metadata about the server configuration.
<i>Traditional</i>	A partition type found in NetWare systems with legacy NetWare traditional file systems. These might exist in a mixed NetWare environment.
<i>Unknown</i>	An partition type that is unknown to the current operating system.
<i>Upgrade Traditional</i>	An NSS partition in mid-upgrade state between different versions of NetWare.
Virtual Device	A partition that serves as a partition in a software RAID 0 or RAID 5 device.

## 8.1.4 Understanding Partition Details

You can view the following information about partitions:

**Table 8-2** *Explanation of Partition Details*

Partition Detail	Description
<i>Partition ID</i>	The partition name assigned by NetWare.
<i>Partition Name</i>	The physical descriptive name of the partition that corresponds to the device's physical descriptive name, followed by the type of partition it is.
<i>Type</i>	The abbreviated name of partition type.
<i>Status</i>	Specifies if a partition is In Use or Available.
<i>Label</i>	The partition name assigned by the administrator.
<i>Starting Offset</i>	Amount of space on the disk that precedes the beginning of the selected partition.
<i>Size</i>	The storage capacity of this partition.
<i>Device Name</i>	The physical descriptive name of the device where the partition exists. For software RAIDs, the description might include RAID 0, RAID 1, or RAID 5.
<i>Device ID</i>	The device name assigned by NetWare.
<i>Pool Name</i>	For NSS partitions, specifies the name of the pool that uses the partition.

## 8.2 Creating NSS Partitions

You can view the details of an NSS partition, but NSS abstracts the creation of partitions for pools and software RAID devices.

Use the Software RAID options to create a new software RAID device and to manage its partitions. For information, see [Section 9.5, “Creating a Software RAID Device Using iManager,” on page 126](#).

Use the Pools option to create a new pool and to manage its partitions. For information, see [Section 11.1, “Creating Pools with iManager,” on page 147](#).

Use the Novell Cluster Services management interface to create a new Cluster Services partition. For information, see the following:

- *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html))
- *OES Novell Cluster Services Administration Guide for NetWare* ([http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html))

Use NSSMU to create an iSCSI device and to manage its partitions. For information, see the *iSCSI 1.1.3 Administration Guide for NetWare 6.5* ([http://www.novell.com/documentation/iscsi1\\_nak/iscsi/data/h4hgu4hs.html](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/h4hgu4hs.html)).

## 8.3 Viewing a List of Partitions

The Partitions page in the Storage plug-in to iManager and in NSSMU is a reporting tool that allows you to view a list of partitions. You can view the types of partitions that are being virtualized by higher-level storage entities such as NSS pools or software RAID devices.

Generally, you cannot create or modify partitions with iManager or NSSMU tools. On NetWare, NSSMU allows you to create an NSS partition, but this capability is seldom used. If you attempt to create a partition from NSSMU or iManager, NSS prompts you to create a pool. NSS partitions are created for you automatically when you create a pool. There might be multiple NSS partitions that are aggregated and managed underneath the single pool of space. Similarly, the tools create a Virtual Device partition automatically when you create a software RAID device.

In iManager, you must access partitions through a related task, such as the following:

- [Section 8.3.1, “Viewing Partitions on a Device,” on page 112](#)
- [Section 8.3.2, “Viewing Partitions in a Software RAID Device,” on page 113](#)
- [Section 8.3.3, “Viewing Partitions in an NSS Pool,” on page 113](#)

### 8.3.1 Viewing Partitions on a Device

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

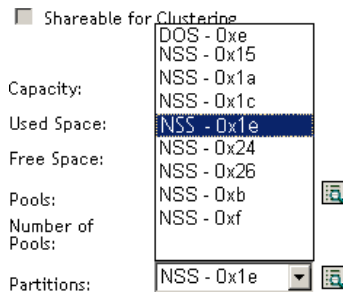
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 In the Devices list, select a device.



- 4 In the Details area, click the arrow on the Partitions drop-down list to expand it.



- 5 Select a partition, then click the *View Partition Details* icon.

This opens the Partition Management page. It displays a list of all the partitions that currently exist on the selected device.

### 8.3.2 Viewing Partitions in a Software RAID Device

- 1 In iManager, click *Storage > Software RAID*s.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Software RAIDs list.

- 3 In the Software RAIDs list, select a device.

Wait for the page to refresh before continuing.

- 4 In the Details area, click the arrow on the Partitions drop-down list to expand it.

- 5 To view details about partitions, click the View Partition Details icon.

This opens the Partition Management page. It displays a list of all the partitions that currently exist on the selected device.

- 6 Select a partition from the Partitions list, then click Details to view more information.

### 8.3.3 Viewing Partitions in an NSS Pool

Although NSS abstracts the partitions underlying the pool structure, you can view information about those partitions.

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select the pool you want to manage.

Wait for the page to refresh and display the details. The pool must be active to see partition details.

- 4 (Conditional) If the pool is deactive, make sure the pool is selected, then click *Activate*.

After the page refreshes, the Partitions drop-down list is available.

- 5 Click on the arrow next to the Partitions drop-down list to expand the list.

- 6 To view details about the partitions, click the View Partition Details icon.

A Partition Management page opens where you can view details about the pool's partitions.

- 7 Select a partition from the Partitions list, then click Details to view more information.

## 8.4 Viewing Details for a Partition

- 1 From the Partitions Management page, select a partition.

- 2 Click Details to view the Partition Information page.

For information about the details, see [Section 8.1, “Understanding Partitions,” on page 109](#).

## 8.5 Labeling a Partition

Partition labels are optional, but if you use them, the label must be unique on a server. If the server is part of a cluster, then partition labels must have unique names across all servers in the cluster, whether they are in shared relationships or not. Partition labels can have 2 to 128 characters.

- 1 On the Partitions Management page, select the partition, then click Edit.

- 2 Type the new label.

- 3 Click *OK*.

## 8.6 Deleting a Partition

Deleting a partition results in data loss. The Storage plug-in to iManager does not allow you to delete partitions from the Partitions page because it abstracts all partition management. Instead, go to the management page of the storage structure that is using the partition to access its partitions and delete them. NSSMU allows you to delete any NSS or NetWare-related partition directly from the Partitions page. Use a third-party management tool to delete Linux traditional file systems partitions on your Linux server or non-NetWare partitions on your NetWare server.

This section discusses the following:

- [Section 8.6.1, “Deleting NSS Partitions in an NSS Software RAID Device,” on page 114](#)
- [Section 8.6.2, “Deleting NSS Partitions in a Pool,” on page 115](#)
- [Section 8.6.3, “Deleting Cluster Services Partitions,” on page 115](#)
- [Section 8.6.4, “Deleting an iSCSI Partition,” on page 116](#)

### 8.6.1 Deleting NSS Partitions in an NSS Software RAID Device

For NSS software RAID devices, use the Software RAID page in iManager or NSSMU to access and delete its partitions.

You can delete all but one partition of a RAID 1 (mirror). You can also delete one failed partition at a time for a RAID 5 to repair and restore the RAID. For information, see [Section 9.9, “Replacing a Failed Segment in a Software RAID,” on page 130](#)

To delete partitions for a RAID 0, use the Software RAID page to delete its RAID.

To delete all of the partitions in a software RAID device, you must delete the software RAID itself. For information, see [Section 9.10, “Deleting a Software RAID Device,” on page 132](#).

### Deleting Partitions with iManager

To delete a single partition in a software RAID 1 or RAID 5 device:

- 1 In iManager, click *Storage > Software RAID*s.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to view a list of its software RAID devices.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3 Select the software RAID 1 or RAID 5 device you want to manage.

Wait for the page to refresh and display the details.

- 4 Click the View Partition Details icon to open the Partitions Management page for the selected software RAID.

- 5 Select a partition to delete.

- 6 Click Delete to delete the partition and its contents.

### Deleting Partitions with NSSMU

- 1 In NSSMU, select Partitions.

- 2 Select the partition you want to delete.

- 3 Click Delete to delete the partition and its contents, then click *OK* to confirm deletion.

## 8.6.2 Deleting NSS Partitions in a Pool

You cannot shrink the size of a pool by deleting its partitions. Use the Pools page in iManager or NSSMU to delete the pool. For information, see [Section 11.10, “Deleting a Pool,” on page 156](#).

## 8.6.3 Deleting Cluster Services Partitions

The Cluster Services partition is a small partition used by Novell Cluster Services™ to pass status information between cluster nodes. Deleting a Cluster Services partition disrupts services for the cluster. If you have a cluster on your server, you should not delete this partition unless it is corrupted and you plan to rebuild it. If you have deleted your cluster and no longer use Novell Cluster Services, the partition can be deleted safely.

If you delete this partition in error, you can rebuild it with Cluster Services console commands. For information, see [“Novell Cluster Services Console Commands”](#) in the *OES Novell Cluster Services 1.8.1 Administration Guide for NetWare*.

- 1 In NSSMU, select Partitions.

- 2 Select the Cluster Services partition you want to delete.
- 3 Click Delete to delete the partition and its contents, then click *OK* to confirm deletion.

### 8.6.4 Deleting an iSCSI Partition

Deleting an iSCSI partition deletes the NSS partitions on the iSCSI partition, destroys all data in the NSS partitions, and makes the iSCSI target inaccessible to the iSCSI initiators. After deletion, servers (cluster or otherwise) no longer have access to the shared storage in your storage area network that was previously made available by the iSCSI target.

---

**WARNING:** Do not delete an iSCSI partition unless it is your intention to remove your shared storage, data and all.

---

If you are not sure about which partitions to delete, you should delete the iSCSI target directly from its management interface.

- 1 In NSSMU, select Partitions.
- 2 Select the iSCSI partition you want to delete.
- 3 Click Delete to delete the partition and its contents, then click *OK* to confirm deletion.

## 8.7 Adding Other Types of File System Partitions (Linux)

After you install NSS and create NSS pools on a device, make sure to use EVMS to manage any other types of file system partitions on that device, such as EXT3, ReiserFS, or other Linux traditional file systems. Continue to use NSSMU or the iManager Storage plug-in as the preferred methods of managing NSS pools and volumes on OES Linux.

If you use another volume manager, your system can become unbootable. During the partition creation process, the YaST partitioner modifies the `/etc/fstab` configuration file to incorrectly identify NSS pools as EXT2 partitions instead of NSS partitions. For example, a command like this is added to the `/etc/fstab` file for each NSS pool:

```
/dev/evms/MYPOOL /nss/.pools/MYPOOL ext2 defaults 1 2
```

To work around this problem when using the YaST partitioner, after you create a Linux traditional file system and before you reboot your system, edit the `/etc/fstab` file to remove or comment out the lines that identify NSS partitions as ext2 partitions.

To recover your system if you reboot your server before editing the `/etc/fstab` file:

- 1 Boot your OES Linux server in Single User mode.
- 2 Mount the root file system with remount and read/write options. At the command console, enter  

```
mount -n -o remount,rw /
```
- 3 Edit the `/etc/fstab` file to remove or comment out the lines that identify NSS partitions as ext2 partitions.
- 4 Reboot the server.

# Managing Software RAID Devices

# 9

RAID devices can help improve your storage availability and performance by enhancing data fault tolerance and I/O performance. The Novell® Storage Services™ (NSS) File System supports software RAID 0, 1, 5, 10, and 15.

This section discusses the following:

- [Section 9.1, “Understanding Software RAID Devices,” on page 117](#)
- [Section 9.2, “Planning for a Software RAID Device,” on page 119](#)
- [Section 9.3, “Viewing Software RAID Devices on a Server,” on page 124](#)
- [Section 9.4, “Viewing Details of a Software RAID Device,” on page 125](#)
- [Section 9.5, “Creating a Software RAID Device Using iManager,” on page 126](#)
- [Section 9.6, “Renaming a Software RAID Device,” on page 128](#)
- [Section 9.7, “Increasing the Size of a Software RAID Device,” on page 128](#)
- [Section 9.8, “Restriping a Software RAID 0 or 5 Device,” on page 129](#)
- [Section 9.9, “Replacing a Failed Segment in a Software RAID,” on page 130](#)
- [Section 9.10, “Deleting a Software RAID Device,” on page 132](#)
- [Section 9.11, “Viewing Pools on a Software RAID Device,” on page 133](#)
- [Section 9.12, “Viewing Partitions on a Software RAID Device,” on page 133](#)
- [Section 9.13, “Deleting Partitions on a Software RAID Device,” on page 134](#)
- [Section 9.14, “Managing Software RAID Devices with NSSMU,” on page 134](#)
- [Section 9.15, “Creating Software RAID Devices with NSSMU,” on page 135](#)
- [Section 9.16, “Mirroring an Existing Pool with NSSMU,” on page 136](#)
- [Section 9.17, “Creating a Software RAID 10 with NSSMU,” on page 136](#)
- [Section 9.18, “Creating a Software RAID 15 with NSSMU,” on page 137](#)
- [Section 9.19, “Creating a RAID 1 Mirror to Duplicate Data,” on page 137](#)

## 9.1 Understanding Software RAID Devices

A software RAID is a configuration for storage devices that emulates a hardware RAID device. A software RAID combines partitioned space from multiple physical devices into a single virtual device that you manage like any device. Each member device contributes an equal amount of space to the RAID. You can create partitions, pools, and volumes on a RAID device, just as you would with any physical storage device. Unlike hardware RAID devices, software RAID 0s use standard host adapters and do not require any special RAID hardware.

The following table describes the software RAID devices supported by NSS:

**Table 9-1** *RAID Characteristics*

Type of RAID	Number of Segments	Purpose	Advantages	Disadvantages
RAID 0	2 to 14	Data striping	Improves I/O performance for both reads and writes, which occur concurrently in parallel to its member devices.	Does not provide data redundancy for data fault tolerance.  If a single disk fails, the data cannot be recovered. You must re-create the RAID 0 and restore its volumes from a backup copy before you can use it again.
RAID 1	2 to 4	Data mirroring	Provides full data redundancy for failover and instant recovery.  Improves read performance.  Equivalent write performance is possible with a duplex connection, which provides a separate channel for each member disk.	To achieve the best I/O performance, it requires separate channels for each member disk; otherwise, write performance decreases slightly.  Each mirror must be on a separate device; it can share no disks in common.  Can be a member of only one pool.
RAID 5	3 to 14	Data striping with parity	Provides limited data recovery for one member disk at a time. If a single drive in the RAID fails, its volumes and pools remain active, but with degraded performance because the RAID must use parity to reconstruct the missing data. You must remove the failed segment, replace the disk, add the new segment, and restripe the data to reconstruct the data on the replacement drive.  Improves read performance if all drives are present and working properly. If a drive fails, read performance is reduced because of parity reads and data reconstruction.	Read responses are the same only if data happens to be in cache when called; otherwise it is slightly reduced for parity checking.  I/O performance for writes is reduced because it takes time to calculate and write parity to disk. The more writes to the drive, the greater is the burden to CPU.  If multiple disks fail, the data cannot be recovered. You must re-create the RAID 5 and restore its volumes from a backup copy before you can use it again.
RAID 10	2 to 4 RAID 0 devices	Mirroring RAID 0 devices	Provides full data redundancy for failover and instant recovery.  Improves I/O performance for both reads and writes, but is slower than an unmirrored RAID 0 device.	Requires separate channels for each member disk to achieve best I/O performance.  RAID 0 devices that you mirror can share no disks in common.  If a single disk fails, you must recreate the RAID 0 and remirror the entire device. The data is restored through mirroring.

Type of RAID	Number of Segments	Purpose	Advantages	Disadvantages
RAID 15	2 to 4 RAID 5 devices	Mirroring RAID 5 devices	<p>Provides full data redundancy for failover and instant recovery.</p> <p>If a single data disk fails, the RAID 5 device remains up and mirrored. Its performance is degraded until you replace the failed disk.</p> <p>It can handle multiple disk failures, depending on the number of failures and which disks fail.</p>	<p>Requires separate channels for each member disk to achieve best I/O performance.</p> <p>RAID 5 devices that you mirror can share no disks in common.</p> <p>If multiple data disks fail concurrently on the same segment or if the parity disk fails, you must remove the damaged segment from the mirror, re-create the RAID 5, and then mirror the RAID 5. The data synchronizes with the mirrors.</p>

## 9.2 Planning for a Software RAID Device

Before you create your software RAID device, you must evaluate your storage requirements and determine which RAID solution best fits your performance and fault tolerance needs.

This section discusses the following:


- [General Guidelines for Software RAID Devices \(page 119\)](#)
- [Guidelines for Software RAID 1 Devices \(page 120\)](#)
- [Choosing a Software RAID Solution \(page 121\)](#)
- [Determining the Number of Partitions \(page 123\)](#)
- [Determining the Partition Size \(page 122\)](#)
- [Determining the Stripe Size for RAID 0 and RAID 5 \(page 123\)](#)

### 9.2.1 General Guidelines for Software RAID Devices

To set up a RAID device, you allocate free space from any of your physical storage devices. NSS transparently presents the allocated free space as virtual partitions that represent NSS-managed physical partition areas on the participating drives. These partitions are the basic elements of a software RAID device. How you allocate each of the partitions to pools depends on the nature of the pools (shared or not shared for clustering) and the type of RAID device it is.

As an example, the RAID 0 device, illustrated in the following figure, consists of three RAID partitions. It stripes data across three physical drives. The stripes are written and read in the order designated.

**Figure 9-1** *Striping Data on a Software RAID 0 Device*



Consider the following general guidelines when creating a software RAID device:

- 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.
- Do not use space from a drive that contains your system partition (such as the / root partition on Linux or the SYS pool on NetWare) or boot partition (Linux).
- You can use any combination of IDE or SCSI devices in a software RAID device. Make sure these devices have similar performance characteristics; otherwise, your performance might decrease.
- In a clustered solution using Novell Cluster Services™, for software RAID 0 or 5 on shared disks:
  - You can have only one pool associated with that RAID device.
  - You must create an NSS pool and volume on that RAID device from the same server node before the pool can be migrated to other nodes in the cluster.

## 9.2.2 Guidelines for Software RAID 1 Devices

The following is a list of requirements for mirroring partitions with software RAID 1 devices:

- 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, each from its own NetWare partition. If a storage pool spans multiple devices, each of the individual partitions that make up that pool can be mirrored independently. All of the pool's partitions 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 partition from your free space but not by adding an existing mirror group to the current group.
- All of the devices that participate in a mirror must be marked a shared or not shared for clustering for each mirror group.
- Avoid setting up multiple mirror groups on a single device. Such configuration heavily degrades the performance of the file system.



- To mirror software RAID 0 devices, the member devices must have no drives in common.

### 9.2.3 Drive Restrictions for NSS Software RAID 0 and 1 Devices on Linux

When you create or expand an NSS software RAID device, do not use space from the drive that contains your boot partition or system partition. In a worst-case scenario, you might need to reinitialize a drive if the partition in the RAID failed.

---

**WARNING:** Reinitializing a drive destroys its contents.

---

If your boot and system partitions are on drives managed by LVM, the drives are not eligible to contribute space to the software RAID.

If the boot and system partitions are on drives managed by EVMS, the drives might appear in the list of available devices if they contain enough space. Do not specify space from those drives as segments for the NSS software RAID device you are creating or expanding.

### 9.2.4 Choosing a Software RAID Solution

When choosing a software RAID solution, determine whether you need to address file system performance, data fault tolerance, or both. The following tables highlights the key data fault tolerance, performance, and configuration issues associated with each RAID type.

**Table 9-2** RAID Performance Characteristics

Requirement	RAID 0	RAID 1	RAID 5	RAID 10	RAID 15
NSS on Linux	Yes, for data volumes only	Yes, for data volumes only	Yes, for data volumes only	Yes, for data volumes only	Yes, for data volumes only
NSS on NetWare	Yes, for data and system volumes	Yes, for data and system volumes	Yes, for data and system volumes	Yes, for data and system volumes	Yes, for data and system volumes
Data fault-tolerance	No	Redundancy	Parity	Redundancy	Redundancy and parity
Read I/O performance	Best improved (parallel reads)	Improved if parallel channels to each mirror	Improved, if all segments are present and working properly	Improved, with RAID 0 read advantage	Improved if parallel channels to each mirror, with RAID 5 read advantage
Write I/O performance	Best improved (parallel writes)	Same if parallel channels; otherwise, slightly decreased	Somewhat decreased by parity calculation	Slightly improved, depending on channel configuration	Somewhat decreased, depending on channel configuration and parity calculations

Requirement	RAID 0	RAID 1	RAID 5	RAID 10	RAID 15
Valid names	2 to 128 characters	2 to 80 characters	2 to 128 characters	2 to 80 characters	2 to 80 characters
Number of segments	2 to 14	2 to 4	3 to 14	Mirror 2 to 4 software RAID 0 devices	Mirror 2 to 4 software RAID 5 devices
Maximum segment size	1 TB, if 2 segments	2 TB for each mirror	1 TB, if 3 segments (2 for data and 1 for parity)	2 TB for each mirror	1 TB, if 3 segments (2 for data and 1 for parity)
Maximum device size	2 TB	2 TB	2 TB	2 TB	2 TB

## 9.2.5 Determining the Partition Size

The space that a member device contributes to a software RAID is called a partition or segment. Each physical device should contribute only one partition to the RAID; otherwise, it negates the benefits of the RAID. A software RAID device can contain only one partition per device. All member partitions in a software RAID device must be the same size.

The size of physical and logical devices used for NSS and Traditional file systems cannot exceed 2 terabytes (TB). Therefore, the maximum partition size is 1 TB for software RAID 0 and 5 devices and 2 TB for software RAID 1. For information, see [Section 9.2.4, “Choosing a Software RAID Solution,” on page 121](#).

The capacity of the RAID device depends on the RAID type and the number of member partitions:

- **RAID 0:** Capacity equals the number of partitions times the partition size.
- **RAID 1:** Capacity equals one partition size.
- **RAID 5:** Capacity equals the number of partitions minus one, times the partition size.
- **RAID 10:** Capacity equals one partition size of space taken from the RAID 0; it is not limited to the partition size of partitions in the RAID 0 itself.
- **RAID 15:** Capacity equals one partition size of space taken from the RAID 5; it is not limited to the partition size of partitions in the RAID 5 itself.

For example, the following table shows the maximum size of partitions needed to create a software RAID device of up to 2 TB.

**Table 9-3** Maximum Segment Sizes Based on Number of Segments and RAID Type

Maximum Segment Size					
Number of Segments	RAID 0	RAID 1	RAID 5	RAID 10	RAID 15
2	1 TB (1000000 MB)	2 TB (where each mirrored element is a 2 TB device)	Not applicable	2 TB (where each mirrored element is a 2 TB RAID 0)	Not applicable

Maximum Segment Size					
3	666666 MB	2 TB (2000000 MB)	1 TB (1000000 MB)	2 TB	2 TB (where each mirrored element is a 2 TB RAID 5)
4	500000 MB	2 TB	666666 MB	2 TB	2 TB
5	400000 MB	Not applicable	500000 MB	Not applicable	Not applicable
6	333333 MB	Not applicable	400000 MB	Not applicable	Not applicable
7	285714 MB	Not applicable	333333 MB	Not applicable	Not applicable
8	250000 MB	Not applicable	285714 MB	Not applicable	Not applicable
9	222222 MB	Not applicable	250000 MB	Not applicable	Not applicable
10	200000 MB	Not applicable	222222 MB	Not applicable	Not applicable
11	181818 MB	Not applicable	200000 MB	Not applicable	Not applicable
12	166666 MB	Not applicable	181818 MB	Not applicable	Not applicable
13	153846 MB	Not applicable	166666 MB	Not applicable	Not applicable
14	142857 MB	Not applicable	153846 MB	Not applicable	Not applicable
Maximum device size	2 TB	2 TB	2TB	2 TB	2 TB

## 9.2.6 Determining the Number of Partitions

Each software RAID device comprises multiple partitions. You must specify at least the minimum number of partitions to create the type of RAID you choose. The maximum number of partitions is limited by the maximum number supported by that RAID type and the maximum device size (2 TB) for NSS and Traditional file systems.

After you set up the software RAID device, you can increase its size by adding segments. In iManager, click *Storage > Software RAIDs > Increase Size*, and then add segments up to the maximum number of segments for each type of RAID, or until you reach the maximum device size of 2 TB.

You cannot remove segments in a RAID device to decrease its size. In general, to reduce the size of a RAID device: Back up its data, delete the RAID, re-create the RAID with a smaller segment size or fewer segments, and then restore its data from the backup copy.

For some RAID types, you can replace a failed partition by removing the segment from the RAID, replacing the failed disk, and then adding a segment to the RAID to replace the failed one. For information, see [Section 9.9, “Replacing a Failed Segment in a Software RAID,” on page 130](#).

## 9.2.7 Determining the Stripe Size for RAID 0 and RAID 5

In RAID 0 and RAID 5 configurations, NSS writes data to each member device in turn. The maximum amount of data (in KB) committed to each write to a partition is called a stripe. Striping is unrelated to file block sizes that you might set on your storage device.

Set the stripe size in increments of powers of two, between 4 KB and 256 KB (4, 16, 32, 64, 128, 256). The default stripe size is 64 KB.

To maximize performance of the RAID, set the stripe size to correspond with your typical data write requirements. In general, use smaller stripe sizes for data servers and medium-to-large sizes for file servers. For most implementations, 64 KB provides the best performance.

### 9.3 Viewing Software RAID Devices on a Server

Use the Software RAID task in the iManager Storage plug-in to create and manage your software RAID devices.

- 1 In iManager, click *Storage > Software RAID*s.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server in the Novell eDirectory™ tree where you are logged in.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).



A list of software RAID devices appears in the Software RAID list. Depending on the number of devices, this can take a few seconds. Avoid clicking again in the page until it refreshes and displays the Software RAID list.



For an overview of the subtasks available from this page, see [“Software RAID” on page 50](#).

[Storage Management](#) ▶ **Software RAID**s

#### Software RAID Management

Create and manage software RAID devices. RAID 0, 1 and 5 can improve read and write performance. RAID 1 and 5 also improve reliability. RAID 0 uses data striping on two to fourteen devices. RAID 1 uses data mirroring with two to four mirrored redundant devices. RAID 5 uses data striping with interleaved parity across three to fourteen devices.

Server:   

Software RAID:	Details:
<div><div>New...</div><div>Delete</div><div>Rename...</div><div>Increase size...</div><div>Restripe</div></div> <div><div>0x13</div><div>0x20</div><div>0x26</div></div>	<div><div>Name: 0x26</div><div>Device:</div><div>Description: RAID0</div><div>ModuleID:</div><div>Adaptor:</div><div>GUID:</div><div>Share State: Not Sharable for Clustering</div><div>Capacity: 100.00 MB</div><div>Used Space: 16.00 KB</div><div>Free Space: 99.98 MB</div><div>Pools: <div><div></div><div></div></div> </div><div>Number of Pools: 0</div><div>Type: RAID 0</div><div>Stripe Size (KB): 64.00 KB</div><div>Segment Size: 50.00 MB</div><div>Segments: <div><div>0x25</div><div></div></div> </div><div>Number of Segments: 2</div><div>Status: In Sync</div></div>

## 9.4 Viewing Details of a Software RAID Device

You can view the following information about a selected software RAID device:

**Table 9-4** *Explanation of Details for a Software RAID Device*

Software RAID Device Detail	Description
<i>Name</i>	The device name assigned by NetWare. This is the Device ID from the list box.
<i>Device</i>	The RAID type and the number of member devices.
<i>Description</i>	The physical description of the storage object. This is the administrator-provided name on creation or on rename. If no name is provided, NetWare provides a default name.
<i>Module ID</i>	The Novell driver identification number for the module; it indicates what type of driver you are using.
<i>Adapter</i>	The sequential number registered for the adapter; it indicates which drive is on which adapter.
<i>GUID</i>	The Global Unique Identifier number that NSS assigns to the storage object. This number is necessary so your file system can locate the specific device.
<i>Share State</i>	Shareable for Clustering or Not Shareable for Clustering. The share state can be modified on the Devices Management page.  If you assign partitions to a software RAID device, all the devices for those member partitions must either be marked as Shareable for Clustering or as Not Shareable for Clustering.
<i>Capacity</i>	The total storage capacity of the device that is reserved for data storage. For a RAID 0, the storage capacity is equal to the sum of its partitions. For RAID 1, the storage capacity is equal to a single partition size; the duplicate partitions are mirrors. For RAID 5, the storage capacity is equal to the sum of its partitions minus one partition for parity.
<i>Used Space</i>	The amount of space on the device that is currently in use by NSS partitions.
<i>Free Space</i>	The total amount of space on the device that is currently not in use.
<i>Partitions</i>	Lists the member partitions of the selected software RAID device in a drop-down list. To view a partition's details, select the partition in the drop-down list, then click the <i>View Details</i> icon.
<i>Partition Size</i>	The size (in MB) of partitioned space per drive.
<i>Stripe Size</i>	The maximum size (in KB) of a data write, as configured for a RAID 0 or RAID 5 device.
<i>Number of Partitions</i>	The total number of partitions in the selected software RAID device.
<i>Status</i>	Shows the status of a RAID 1 (mirrored) device's partitions: In Sync or Out of Sync.

## Viewing Details in iManager

- 1 In iManager, click *Storage > Software RAIDs*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server in the eDirectory tree where you are logged in.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

This opens the Software RAID Management page.

The Software RAID list displays the virtual RAID devices on the selected server. The list might include any RAID 0, RAID 1, or RAID 5 devices that you created. It does not list any hardware RAID devices in this list.

- 3 Select a virtual storage device in the Software RAID list to view information about that device, then wait for the page to refresh.

## Viewing Details in NSSMU

- 1 In NSSMU, select Software RAID.

- 2 Select the RAID device you want to manage and wait for the information to be displayed.

# 9.5 Creating a Software RAID Device Using iManager

- 1 In iManager, click *Storage > Software RAIDs*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server in the eDirectory tree where you are logged in.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3 Click New.


This opens a wizard that steps you through the process.

- 4 **Device Type and Name:** Specify the type of RAID device you want to create and name the RAID, then click *Next*.

If you do not provide a name, NSS creates a unique name for the device. For information about choosing names, see [Section 5.5, “Naming NSS Storage Objects,” on page 69](#).

This opens the Devices and Space page.

Create Raid Device

 Determine devices and space

Name: **R5\_nifa**

A RAID device can be created on one or more storage objects. When the RAID device is created, the amounts in all of the used columns will be added together to create the RAID device's total size.

Segment Size:

19999

	Used Size (MB)	Storage	Available Size (MB)
<input type="checkbox"/>	0	0x3-1	19234
<input checked="" type="checkbox"/>	19999	0x3-2	80000
<input checked="" type="checkbox"/>	19999	0x4-1	80000
<input checked="" type="checkbox"/>	19999	0x5-1	80000

<< Back

Next >>

Cancel

**5 Devices and Space:** Select devices and the amount of space to use from each, then click *Next*.

**5a** In the *Partition Size* field, type the amount of space in MB to use from each physical device.

NSS identifies devices that have enough free space to meet the partition-size requirements and are eligible for inclusion in the RAID. For information, see “**Determining the Partition Size**” on page 122.

If the amount you specify exceeds the amount of free space available on a minimum number of physical devices, the RAID creation fails and returns an error message.

**5a** Select the check box next to each of the storage devices you want to obtain space from.

You can obtain space from multiple devices. Select only devices that have enough space available to meet your needs.

If a device’s available space is smaller than the specified partition size, it is disabled (dimmed) so that you cannot select it.

A single physical device can offer multiple free space areas in the list. After you select a device, all other free space on that device is disabled for that RAID. Each device should contribute only one partition to the RAID; otherwise, it defeats the purpose of improved performance and data protection that a RAID affords.

If the server has both local and shared devices, the partitions of a RAID can reside only on shared devices or only on local devices. If you select the check box next to a local storage device, the shared devices are dimmed so that you cannot select them. If you select the box next to a shared storage device, the local devices are dimmed.

Only devices that have free space appear in the list. If no devices are listed, there is no space available to create a software RAID device. *Cancel* the wizard, add more devices to the server or free up space on existing devices, then return to the Software RAIDs page to create a RAID device.

**5a** Click *Next*

- 6 (Conditional) Stripe Size:** If this is a RAID 0 or 5 device, set the stripe size. For information, see [“Determining the Stripe Size for RAID 0 and RAID 5” on page 123.](#)
- 7** Click *Finish*. (Or click *Cancel* at any time to back out of the process.)  
NSS creates the software RAID device, then opens to the Software RAIDs task. Your newly created RAID device appears in the Software RAID Devices list.

After you create the RAID, manage it as you would a physical device in terms of using it for pools and volumes. For information on configuring an NSS file system on your RAID, see [Chapter 11, “Managing NSS Pools,” on page 147.](#)

## 9.6 Renaming a Software RAID Device

- 1** In iManager, click *Storage > Software RAIDs*.  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46.](#)
- 2** Select a server in the eDirectory tree where you are logged in.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46.](#)
- 3** Click *Rename*.  
This opens a dialog box where you can enter the new name.
- 4** Type the new name.  
For information about choosing names, see [Section 5.5, “Naming NSS Storage Objects,” on page 69.](#)
- 5** Click *OK*. (Or click *Cancel* at any time to back out of the process.)  
NSS renames the software RAID device, then opens to the Software RAID Management task. Your newly renamed software RAID device appears in the Software RAIDs list.

## 9.7 Increasing the Size of a Software RAID Device

You can increase the capacity of an existing software RAID 0, 1 or 5 device by adding partitions, up to the maximum number for the type of RAID. You cannot modify the size of an individual partition after the device is created.

To add partitions to an existing software RAID:

- 1** In iManager, click *Storage > Software RAIDs*.  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46.](#)
- 2** Select a server in the eDirectory tree where you are logged in.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46.](#)
- 3** Make sure that there is no I/O for the volumes on the RAID device by deactivating the volume.
  - 3a** Click *Volumes*.
  - 3b** Select the volumes in the RAID device you want to expand.
  - 3c** Click *Deactivate* (NetWare) or *Dismount* (Linux).



**4** Select a device in the Software RAID Devices list.

If the device contains the maximum number of partitions, the Increase Size option is dimmed. You cannot expand the RAID. Do not proceed.

**5** Click *Increase Size*.

This opens the Expand a RAID wizard to let you choose from available free space on devices that are not already members in the RAID and that contain enough free space to meet the RAID's current partition size.

**6** Do one of the following:

- If there are no devices available, you cannot expand the RAID. Proceed to [Step 7](#).
- If the RAID already contains the maximum number of partitions, the Increase Size option is dimmed. You cannot expand the RAID. Proceed to [Step 7](#).
- Select the check box next to each of the storage devices you want to obtain space from.

The partition size is predetermined by the existing RAID. The partition you choose must be the same size as other partitions comprising the device.

Stripe size is fixed at its current value for the duration of the expansion. If you want to change the stripe size, restripe after the expansion.

You can choose multiple partitions up to the maximum for the type of RAID it is. For information, see [“Determining the Number of Partitions” on page 123](#) and [“Determining the Partition Size” on page 122](#).

**7** Click *Finish*.

After you add a partition, the RAID's data is restriped across both existing and new partitions. During the restriping, the RAID's capacity does not include the added partition. After the restriping is complete, the RAID's capacity includes the added partition.

## Remirroring and Restriping Temporarily Impacts System Performance

For software RAID 1 devices, the additional mirror begins to collect data immediately. Performance is slightly impacted, but data is available.

For software RAID 0 or RAID 5 devices, the system automatically begins restriping the disks to accommodate the partition you just added. This severely impacts performance until the striping is complete. The capacity of the RAID is not updated until the restriping is done. If the restriping process is interrupted before it completes, it begins automatically on system reboot.

## 9.8 Restriping a Software RAID 0 or 5 Device

In general, there are three reasons for restriping of software RAID 0 and 5 devices:

- **Partition Replacement:** If a partition fails, you must replace it. Restriping can recover the data in a single lost partition in a RAID 5 by using parity. However, the data must be restriped from a backup tape if a partition fails in a RAID 0.
- **RAID Expansion:** If you expand a RAID 1 or 5 device, the RAID restripes the data across all members.
- **RAID Stripe Size:** If you increase or decrease the stripe size of a RAID 1 or 5 device, the RAID restripes the data across all members. This happens infrequently, unless you are measuring performance with different striping sizes to determine which best fits your operational needs.

If the restriping process is interrupted, the RAID recognizes that when the system reboots, and automatically continues the restriping process. You can also use iManager to pause and resume a restriping process.

To manually resume or pause the Restripe process:

- 1 In iManager, click *Storage > Software RAID*s.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server in the eDirectory tree where you are logged in.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3 Make sure that there is no I/O for the volumes on the RAID device by deactivating the volume.

**3a** Click *Volumes*.

**3b** Select the volumes in the RAID device you want to expand.

**3c** Click *Deactivate* (NetWare) or *Dismount* (Linux).

- 4 Select a device in the Software RAID Devices list.

Wait for the page to refresh before continuing.

- 5 Click *Restripe*.

The restriping process begins or pauses immediately. Restriping severely degrades I/O performance until the restriping is complete.

## 9.9 Replacing a Failed Segment in a Software RAID

For some RAID types, you can replace a failed segment by removing the segment from the RAID, replacing the failed disk, and then adding a segment to the RAID to replace the failed one. The following table provides recommended actions for each RAID type.

**Table 9-5** Recommended Actions on Segment Failure in a Software RAID

Software RAID	Remove Segments	Recommended Action
RAID 0	Not allowed	If one or more disks in a RAID 0 fails, all data is lost on that device. You must re-create the RAID and restore its data from a backup copy.
RAID 1	All but one mirrored segment	<p>The last segment of a RAID 1 device is the primary device that contains the data.</p> <p>You can use NSSMU Partitions option to remove mirrored segments to free up the related space on those segments, but the data on the last remaining segment is no longer mirrored.</p> <p>If a disk fails in one of the segments, all of its data is lost on that segment. The remaining mirrors continue to operate normally. Remove the segment, replace the failed disk, and then add the segment as an element in the mirror. The data is synchronized over time until it is fully mirrored on the new segment.</p>

Software RAID	Remove Segments	Recommended Action
RAID 5	One data segment at a time	<p>You can temporarily remove one segment at a time in a RAID 5 device to replace a failed data disk.</p> <p>For example, you can replace a single failed data disk while the system is still operational. Use the NSSMU Partitions option to remove the failed segment, replace the failed disk, add a segment to the RAID 5, and then restripe the RAID. The parity is used during restriping to restore the missing data. Read and write performance is degraded on the failed segment until the data is recovered because of parity-related reads and calculations.</p> <p>If multiple data segments fail concurrently or if the parity segment fails, all data is lost on that device. You must delete the RAID 5 and re-create it with good disks. Recover its data from a backup copy.</p>
RAID 10	All but one mirrored segment; cannot remove disks from the underlying RAID 0	<p>The last segment of a RAID 10 is a RAID 0 device that contains the original data.</p> <p>If one or more disk fails in one of the mirrored segments, all data is lost on that segment. The remaining mirrors continue to operate normally. Remove the failed RAID 0 from the RAID 10. Delete the RAID 0, replace the failed disks, re-create the RAID 0, and then add the RAID 0 segment as an element in the mirror. The data is synchronized over time until it is fully mirrored on the repaired RAID 0 segment.</p>
RAID 15	All but one mirrored segment; can safely remove one segment at a time per segment	<p>The last segment of a RAID 15 is a RAID 5 device that contains the original data.</p> <p>If a single data segment fails in a RAID 5 that is an element in a RAID 15, repair the RAID 5 while it is operational, as you would with any <b>RAID 5</b>.</p> <p>If multiple disks in a mirrored segment fail concurrently or if the parity segment fails, all data is lost on that segment. The remaining mirrors continue to operate normally. Remove the failed RAID 5 from the RAID 15. Delete the RAID 5, replace the failed disks, re-create the RAID 5, and then add the repaired RAID 5 as a segment in the RAID 15. The data is synchronized over time until it is fully mirrored on the repaired RAID 5 segment.</p>

## A Segment Fails in a RAID 0

If a segment fails in a RAID 0, you must delete the software RAID 0 device, create a new RAID 0 device, then copy your data to the RAID from backup media. For information, see [Section 9.10, “Deleting a Software RAID Device,” on page 132](#).

## A Segment Fails in a RAID 1

- 1 From the command console, enter `nssmu`.
- 2 From the NSSMU main menu, select Software RAIDs.
- 3 Remove the bad segment.
  - 3a Select the software RAID 1 device that you want to manage.
  - 3b Press Enter to show its member segments. The bad segment should show a status of Bad - Unavailable Partition.

- 3c** Select the bad segment, then press Delete.
- 4** Expand the RAID with a replacement segment.
  - 4a** Select the software RAID 1 device that you want to manage.
  - 4b** Press *F3* to increase the size of the RAID.
  - 4c** From the list of available devices, select the device you want to use for the new segment. The segment size defaults to the size of existing partitions in the RAID 1.
  - 4d** Select *OK* twice.
- 5** The data begins mirroring automatically and continues until the segment is 100% mirrored.

### A Single Data Segment Fails in a RAID 5

To replace a single failed data segment in a software RAID 5:

- 1** From the command console, enter `nssmu`.
- 2** From the NSSMU main menu, select Software RAIDs.
- 3** Remove the bad segment.
  - 3a** Select the software RAID 5 device that you want to manage.
  - 3b** Press Enter to show its member segments. The bad segment should show a status of Bad - Unavailable Partition.
  - 3c** Select the bad segment, then press Delete.
- 4** Expand the RAID with a replacement segment.
  - 4a** Select the software RAID 5 device that you want to manage.
  - 4b** Press *F3* to increase the size of the RAID.
  - 4c** From the list of available devices, select the device you want to use for the new segment. The partition size defaults to the size of existing partitions in the RAID 5.
  - 4d** Select *OK* twice.
- 5** The restriping should begin automatically. If it does not, from the Software RAIDs page, select the RAID 5 device, then press *F6* to Restripe.

### Multiple Segments Fail in a RAID 5

If two or more segments fail concurrently in a RAID 5 or if the parity partition fails, you must delete the software RAID 5 device, create a new RAID 5 device, then copy your data to the RAID from backup media. For information, see [Section 9.10, “Deleting a Software RAID Device,” on page 132](#).

## 9.10 Deleting a Software RAID Device

If you delete a software RAID device, it ends the RAID relationship, and it destroys the NSS file structure on member partitions. All data is lost. Make sure to back up your data or move it to another location before deleting the software RAID device.

- 1** In iManager, click *Storage > Software RAIDs*.
 

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).
- 2** Select a server in the eDirectory tree where you are logged in.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.

- 3 Make sure that there is no I/O for the volumes on the RAID device by deactivating the volume.
  - 3a Click *Volumes*.
  - 3b Select the volumes in the RAID device you want to expand.
  - 3c Click *Deactivate* (NetWare) or *Dismount* (Linux).
- 4 Select a device in the Software RAID Devices list.
- 5 Click Delete.

## 9.11 Viewing Pools on a Software RAID Device

- 1 In iManager, click *Storage > Software RAIDs*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on page 46.

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.

A list of devices appears in the Software RAIDs list.

- 3 In the Software RAIDs list, select a RAID device.
- 4 In the Details area, click the arrow on the Pools drop-down list to expand it.



- 5 Select a pool, then click *View Details*.

This opens the Pool Management page where you can view the details of the pool and manage it. See [Section 11.1, “Creating Pools with iManager,”](#) on page 147 for a sample Pool Management page.

For information about pool management, see [“Managing NSS Pools”](#) on page 147.

## 9.12 Viewing Partitions on a Software RAID Device

- 1 In iManager, click *Storage > Software RAIDs*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on page 46.

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.

A list of devices appears in the Software RAIDs list.

- 3 In the Software RAIDs list, select a device.

Wait for the page to refresh before continuing.
- 4 In the Details area, click the arrow on the *Partitions* drop-down list to expand it.

- 5 To view information about partitions, click the *Partitions View Details* icon.  
This opens the Partition Management page. It displays a list of all the partitions that currently exist on the selected device.
- 6 Select a partition from the Partitions list, then click *Details* to view its details.

## 9.13 Deleting Partitions on a Software RAID Device

You can delete all but one partition of a RAID 1 (mirror) and only one partition at a time for a RAID 5. To delete a RAID 1 partition, you must delete its RAID.

- 1 In iManager, click *Storage > Software RAID*s.  
  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).
- 2 Select a server to manage.  
  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).  
A list of devices appears in the Software RAIDs list.
- 3 Make sure that there is no I/O for the volumes on the RAID device by deactivating the volume.
  - 3a Click *Volumes*.
  - 3b Select the volumes in the RAID device you want to expand.
  - 3c Click *Deactivate* (NetWare) or *Dismount* (Linux).
- 4 In the Software RAIDs list, select a device.  
Wait for the page to refresh before continuing.
- 5 Click the *Partitions View Details* icon.  
This opens the Partition Management page. It displays a list of all the partitions that currently exist on the selected device.
- 6 Select a partition from the Partitions list, click *Delete*, then click *OK*.

## 9.14 Managing Software RAID Devices with NSSMU

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

**Table 9-6** *Software RAID Management Tasks in NSSMU*

Keystroke	Description
<i>Enter</i>	Shows the RAID partitions associated with the selected RAID device.
<i>F3</i>	Expands a RAID device by adding more partitions from another storage device.  You should not place more than one RAID partition on a disk; this will severely impede the performance of your file system.

Keystroke	Description
<i>F4</i> (Linux)	Rename a selected RAID device.
<i>F6</i>	Restripes or remirrors a RAID device.  Make sure all partitions have been added to your RAID device before restriping a device.  The file system performance might slow down during the restriping/remirroring process.
<i>F8</i>	Show pools (list pools on selected RAID 0 or 5; for RAID 1, its single member pool appears in device details)
<i>Ins</i>	Creates a software RAID (0, 1, or 5)
<i>Del</i>	Deletes an entire RAID device.  When you delete a RAID device, all partitions, volumes, and pools associated with that device are also deleted.
<i>F5</i>	Refreshes the display.

## 9.15 Creating Software RAID Devices with NSSMU

- 1 In NSSMU, select RAID Devices from the NSSMU main menu.
- 2 Press Insert (Ins) to create a new device.
- 3 Select the RAID type (0, 1, or 5), then press Enter.
- 4 (Conditional) If this is a RAID 0 or 5 device, specify the stripe size, then press Enter.  
The default stripe size is 64 KB, which typically provides the best performance for devices with NSS volumes.
- 5 Use the arrow keys to select the partitions that you want to contribute space to the RAID.  
If no partitions appear, it is an indication that either there are no partitions large enough or there are no free partitions.

After space is selected from a device, other free space associated with that device might not appear. This prevents you from adding more than one segment from a single physical device, and consequently, helps ensure the optimum performance of your file system.

The following table lists the number of segments that you can include in a RAID device:

RAID Type	Number of Segments
RAID 0	Minimum of 2 and a maximum of 14
RAID 1	Maximum of 4
RAID 5	Minimum of 3 and a maximum of 14; one segment is used for parity

- 6 Specify the amount of space to use, then press Enter.  
The segment is created and added to the Segments Included in RAID window.

## 9.16 Mirroring an Existing Pool with NSSMU

**1** Before you begin, review the [Section 9.2.2, “Guidelines for Software RAID 1 Devices,”](#) on [page 120](#).

**2** In NSSMU, select Partitions from the NSSMU main menu.

**3** From the list of existing partitions, select the NSS partition for the pool you want to mirror.

**4** Press F3 to create the RAID 1 device and mirror the partition where the pool resides.

**5** From the available devices, select up to three additional devices that you want to use as a segment, then press Enter.

The space assigned is the same size as the existing partition. The segments must reside on different devices. If no partitions appear, it is an indication that either there are no partitions large enough or no free space exists on other devices.

**6** To confirm the RAID 1 device, select RAID Devices from the NSSMU main menu. The RAID 1 device ID appears in the RAID Devices window.

## 9.17 Creating a Software RAID 10 with NSSMU

If you want to mirror your software RAID 0 devices, you must use NSSMU to mirror the partition used by the pool on a RAID 0 device. The following procedure describes how to create the RAID 10 in NSSMU. You can also create the RAID 0 devices first in iManager, but you must use NSSMU to be able to select them for a mirror.

**1** In NSSMU, create a software RAID 0 device with 2 to 14 segments.

For information, see [Section 9.15, “Creating Software RAID Devices with NSSMU,”](#) on [page 135](#).

**2** Repeat [Step 1](#) one to three times to create 2 to 4 RAID 0 devices.

The RAID 0 devices you use to create the mirror must have no drives in common. Each drive you use to create the RAID 0 devices can belong to only one of the RAID 0 devices.

**3** In NSSMU, create a pool on one of the RAID 0 devices.

**3a** In NSSMU, select Pools from the NSSMU main menu.

**3b** Press *Insert* (Ins) to create a pool.

**3c** From the list of available devices, select one of the RAID 0 devices.

**3d** Assign all of the available space to the pool, then press Enter.

**4** Create a RAID 1 device to mirror the pool.

**4a** In NSSMU, select Partitions from the NSSMU main menu.

**4b** Select the NSS partition for the pool you want to mirror.

**4c** Press F3 to create the RAID 1 device and mirror the partition.

**4d** From the available devices, select one to three of the remaining RAID 0 devices you created above, then press Enter.

**4e** Press F3 to initialize and create the RAID 1 (mirror) device.

After the RAID device is created, the device ID appears in the RAID Devices window. This window is viewed from the RAID Devices NSSMU main menu. The RAID is a RAID 10.



## 9.18 Creating a Software RAID 15 with NSSMU

If you want to mirror your software RAID 5 devices for your server, you must use NSSMU to mirror the partition used by the pool on a RAID 5 device. The following procedure describes how to create the RAID 15 in NSSMU. You can also create the RAID 5 devices and NSS pool for the RAID 5 devices in the Storage plug-in for iManager, and then switch to NSSMU to mirror the pool's partition.

- 1 In NSSMU, create a software RAID 5 device with 3 to 14 segments.

For information, see [Section 9.15, “Creating Software RAID Devices with NSSMU,” on page 135](#).

- 2 Repeat [Step 1](#) one to three times to create 2 to 4 RAID 5 devices.

The RAID 5 devices you use to create the mirror must have no drives in common. Each drive you use to create the RAID5s can belong to only one of the RAID 5 devices.

- 3 In NSSMU, create a pool on one of the RAID 5 devices.

**3a** In NSSMU, select Pools from the NSSMU main menu.

**3b** Press Insert (Ins) to create a pool.

**3c** From the list of available devices, select one of the RAID 5 devices.

**3d** Assign all of the available space to the pool, then press Enter.

- 4 Create a RAID 1 device to mirror the pool.

**4a** In NSSMU, select Partitions from the NSSMU main menu.

**4b** Select the NSS partition for the pool you want to mirror.

**4c** Press F3 to create the RAID 1 device and mirror the partition.

**4d** From the available devices, select one to three of the remaining RAID 5 devices you created above, then press Enter.

**4e** Press F3 to initialize and create the RAID 1 (mirror) device.

After the RAID device is created, the device ID appears in the RAID Devices window. This window is viewed from the RAID Devices NSSMU main menu. The RAID is a RAID 15.

## 9.19 Creating a RAID 1 Mirror to Duplicate Data

You can create a RAID 1 mirror to duplicate data on a new device, such as to duplicate data on a new storage array.

- 1 Set up the RAID 1 mirror between the initial storage element and new storage element.

For information, see [Section 9.5, “Creating a Software RAID Device Using iManager,” on page 126](#).

- 2 Let the RAID create a duplicate of the data on the mirror.

- 3 (Optional) If you want to retain the data on the initial storage element and use the element elsewhere, remove the disk it is on from the server.

---

**WARNING:** If you leave the disk attached to the server while deleting the element from the RAID, its data is destroyed.

---

- 4 Use NSSMU or iManager to delete the initial storage element from the mirror, leaving only the new storage element active on the server as a single-element mirror.

The RAID 1 group remains in Media Manager, using only 1 KB of memory. The new array performs normally, without performance degradation and without consuming additional resources.

---

**WARNING:** Leave the RAID 1 group active because deleting the RAID 1 group deletes all of its member partitions and destroys the data on them.

---

In iManager and from the command line, the new array reports that it is Not Mirrored. In NSSMU, the new array reports that it is In-Sync and 100% remirrored, even though there is only a single element.

# Managing Multiple Connection Paths to Devices (NetWare)

# 10

Novell® Storage Services™ supports multipathing, a fault-tolerant and high-availability storage solution. NSS enables you to dynamically manage multiple, redundant connection paths between a server and its external storage devices, according to priorities you set. A single path serves as the primary, active interconnect; the other paths serve as failover paths. In the event of a failure in the active interconnect, NSS automatically and dynamically recovers to another available failover path. This feature is only available if you have multiple interconnect paths in your server-to-storage configuration.

This section discusses the following:

- [Section 10.1, “Planning Your Multiple Connection Path Solution,” on page 139](#)
- [Section 10.2, “Configuring the Primary Path and Priorities for Failover Connection Paths,” on page 140](#)
- [Section 10.3, “Viewing Connection Path Details,” on page 141](#)
- [Section 10.4, “Setting a Connection Path as Up or Down,” on page 141](#)
- [Section 10.5, “Setting the Primary Path for a Device to Its Default Path,” on page 142](#)
- [Section 10.6, “Resetting the Server Registry with Default Priority Settings for a Device,” on page 143](#)
- [Section 10.7, “Managing Multiple Paths with Server Console Commands,” on page 143](#)

## 10.1 Planning Your Multiple Connection Path Solution

Connection failures can occur if an adapter, cable, or switch in the path between the server and the storage device fails for any reason. For example, if you have two adapters in a server pointing to the same device and the adapter in the primary path fails, the connection between the server and the device fails. The multipath connection feature automatically fails over to the secondary path, using the second adapter and its related path to continue communications. The failover is transparent to users and applications.

To achieve the desired availability, your storage solution can implement one or more host bus adapters in the server and multiple adapters in a SAN. These adapters can interconnect with cables from the server to the device, with cables from each set of adapters to the same switch, or with cables from each set of adapters homed to different switches in your SAN. NSS multipathing is interoperable with third-party storage products such as SANs and hardware RAIDs.

NSS automatically identifies multiple connection paths and randomly selects one path to serve as the primary, active interconnect. You can specify path priorities to control which path serves as the primary path and to set up the failover sequence of the alternate available paths. NSS multipath support enables you to dynamically manage multiple paths according to the priorities you set.

In the event of a failure in the active interconnect, NSS automatically and dynamically recovers to another available failover path. The path with the highest priority of the paths that are available (up)

becomes the primary path. The paths not selected remain as failover device paths for the new primary path.

Set a path's priority to determine its position in the failover sequence of the alternate available paths. The default priority is zero (0), which is no priority. Assign a value between 1 (highest priority) to 4 million (lowest priority).

The NLM™ programs needed to take advantage of multipath support (mm.nlm and nwpa.nlm) are installed and loaded when you install NetWare® 6 and later. You can manage multipathing using the browser-based iManager Storage plug-in, NSS Management Utility (NSSMU), or server console commands.

## 10.2 Configuring the Primary Path and Priorities for Failover Connection Paths

You can configure multiple connection paths for a selected device by implementing designs such as the following:

- Configure the server with multiple host bus adapters for connection to external devices.
- Use a storage device with multiple interconnects for connection to one or more host bus adapters.
- Interconnect a server's multiple host bus adapters to storage devices through one or more intermediate network devices, such as a Fibre Channel switch.

Typically, multipath solutions involve storage area network interconnects, but other solutions are possible with external SCSI devices.

To set the primary path and set priorities for failover:

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 Select a device to manage.

Wait for the page to refresh before continuing.

- 4 Click *Multipath*.

This opens the Multipath page. All of the paths available between the selected server and the selected device appear in the Failover Device Paths list.

- 5 Select a path to manage.

Wait for the page to refresh before continuing.

- 6 Do one or both of the following:

- **Select Path:** Click *Select Path* to make the selected path the primary path between the server and the external storage device. If the selected path is already specified as the primary path, this option is dimmed.

- **Priority:** Click *Priority* to set the priority for the selected path. In the Multipath Priority dialog box, enter a valid priority range of 1 to 4 million, where 1 is the highest priority and 4 million is the lowest. The default value of 0 (zero) is no priority.

## 10.3 Viewing Connection Path Details

You can view the following information in the Failover Device Path details:

**Table 10-1** *Explanation of Failover Device Path Details*

Failover Device Path Detail	Description
<i>Device Name</i>	The device name assigned by NetWare.
<i>Path Name</i>	The path name assigned by NetWare.
<i>Adapter ID</i>	The ID assigned to the host bus adapter in the server that is a component of the selected interconnect.
<i>Port</i>	The port address for the storage device that is a component of the selected interconnect.
<i>Priority</i>	The integer value you set for the failover path.
<i>Status</i>	<p>The state of the path.</p> <ul style="list-style-type: none"> <li>• <b>Available:</b> A path is marked as <i>Up</i> if it is available, or <i>Down</i> if it is not.</li> <li>• <b>Selected:</b> Only one path serves as the primary path at any given time. The selected path is the primary path.</li> </ul>

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 Select a device to manage.

Wait for the page to refresh before continuing.

- 4 Click *Multipath*.

This opens the Multipath page. All of the paths available between the selected server and the selected device appear in the Failover Device Paths list.

- 5 Select a path to manage.

When the page refreshes, the path’s details appear in the Details area.

## 10.4 Setting a Connection Path as Up or Down

When a path’s status is *Up*, it is available to serve as the primary path or as a failover path between the server and the storage device. A connection path’s status might be *Down* because a component

along the path has failed. You might specify a functioning path as Down to perform maintenance on components in a path.

You can also designate a path as Down if you want to remove the path from consideration as a failover path. For example, if multiple adapters in the server are connected to a target hardware device consisting of multiple devices, NSS would see multiple paths to the devices through each adapter. You might want to allocate which adapter serves which device. To accomplish this, you could down the paths you want the server to ignore.

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 Select a device to manage.

Wait for the page to refresh before continuing.

- 4 Click *Multipath*.

This opens the Multipath page. All of the paths available between the selected server and the selected device appear in the Failover Device Paths list.

- 5 Select a path to manage.

Wait for the page to refresh before continuing.

- 6 Select one of the following:

- **Up:** Make the path available to be selected as the primary path or as one of the failover device paths. If the selected path is already up, this option is dimmed.
- **Down:** Make the path unavailable to be selected as the primary path or as one of the failover device paths. If the selected path is already down, this option is dimmed.

## 10.5 Setting the Primary Path for a Device to Its Default Path

The Set Default Path option resets the connection paths to selected devices to their user-defined default primary paths. For example, after an automatic failover, you can reset the device to use its default primary connection path for current connectivity.

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 In the Devices list, select one or more devices that you want to manage.

- 4 Click *Set Default Path*.

The action affects only those devices that actually have multiple paths available.

## 10.6 Resetting the Server Registry with Default Priority Settings for a Device

The Reset Registry option resets the values for path priorities in the registry to their user-defined default values for selected devices. For example, after an automatic failover, you can reset the device to use its default connection priorities for future failovers.

- 1 In iManager, click *Storage > Devices*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of devices appears in the Devices list.

- 3 In the Devices list, select one or more devices that you want to manage.

- 4 Click *Reset Registry*.

The action affects only those devices that actually have multiple paths available.

## 10.7 Managing Multiple Paths with Server Console Commands

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:

Media Manager Multipath I/O Console Command	Description
<pre>list failover devices device_id</pre>	<p>This command lists all of the failover devices, all the paths associated with each device, their assigned priorities, and the selected primary path.</p> <p>To list the paths associated with all the devices on the server, enter the command without specifying the <code>device_id</code>.</p> <p>To list the paths associated with a given device, enter the command followed by the <code>device_id</code>.</p> <p>For example, to see the paths associated with a device with <code>device_id</code> 0x12, enter the following at the server console</p> <pre>list failover devices 0x12</pre> <p>The list provides output similar to the following results:</p> <pre>Ox12 [V132-A1-D2:20] ABC DEVICE rev:5678 Up 11 [V132-A1-D2:20] ABC DEVICE rev:5678 Priority = 1 selected Up 26 [V132-A1-D2:20] ABC DEVICE rev:5678 Priority = 2</pre> <p>In this example, two paths are assigned to device 0x12. Number 11 is <b>Up</b>, and it has been assigned the highest priority of 1. The word <b>Up</b> simply means that this path is available for use. The word <b>selected</b> means that any disk requests that are trying to reach the 0x12 device automatically go through the 11 path, because it is set as the primary path. In this example, you could use the <code>mm set failover path</code> command to override the priorities and make the number 26 path to be the primary path used.</p>
<pre>mm set failover priority pathid=number</pre>	<p>This command sets the priority level for the selected path. The <i>pathid</i> must be a valid path ID. The <i>number</i> is the priority to set for the selected path. The number must be a decimal integer value between 1 and 4 million, with 1 being the highest priority and 4 million being the lowest priority.</p>
<pre>mm set failover priority pathid=number /insert</pre>	<p>The <code>/insert</code> option inserts the specified path with its 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 <code>/insert</code> 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.</p>
<pre>mm set failover state pathid=up /setpath</pre>	<p>This command sets the state of the specified path to Up. The <i>pathid</i> must be a valid path ID.</p> <p>If the Up option is followed by the <code>/setpath</code> option, NSS automatically selects the highest priority path from among all the paths that are in the Up state, including the new path.</p>



Media Manager Multipath I/O Console Command	Description
mm set failover state <i>pathid</i> =down /setpath	<p>This command enables you to designate a path to be in the Down state so you can perform preventative maintenance or reroute cables on a good, functional path. If the path is bad, NSS automatically moves the path to a Down state.</p> <p>If the Down option is followed by the /setpath option, NSS automatically selects the highest priority path from among the other paths that are currently in the Up state.</p>
mm set failover path <i>pathid</i>	<p>This command activates the selected path as <b>selected</b>, which means it serves as the primary path. The <i>pathid</i> must be a valid path that is currently in the Up state. The priority of a path remains static regardless of whether a path is considered up or down or whether a path is selected or not selected as the primary path.</p>
mm restore failover path <i>device_id</i>	<p>This command forces the device to reselect the highest priority path that is online. The <i>deviceid</i> must be a valid device ID.</p>
mm reset failover registry	<p>This command resets the multipath settings in the operating system registry.</p>



# Managing NSS Pools

# 11

Novell® Storage Services™ uses storage pools to efficiently acquire and use all free space available on devices. A pool is an area of storage that consists of space, called a *partition*, obtained from one or more of the storage devices available on a server. The amount of space that each storage device contributes can differ for each member device.

Use the iManager Storage plug-in to configure and manage NSS pools. For information about iManager, see the *Novell iManager 2.5 Administration Guide* ([http://www.novell.com/documentation/imanager25/imanager\\_admin\\_25/data/hk42s9ot.html](http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html)).

You can also use the console-based NSS Management Utility to configure and manage NSS pools. For information, see **Section 4.2, “NSS Management Utility,” on page 52.**

This section describes how to configure and manage NSS pools by completing the following tasks:

- **Section 11.1, “Creating Pools with iManager,” on page 147**
- **Section 11.2, “Viewing Pools on a Server,” on page 150**
- **Section 11.3, “Viewing Pool Details,” on page 151**
- **Section 11.4, “Viewing Partition Information for a Pool,” on page 152**
- **Section 11.5, “Viewing Volume Information for a Pool,” on page 153**
- **Section 11.6, “Viewing Device Information for a Pool,” on page 153**
- **Section 11.7, “Renaming a Pool,” on page 154**
- **Section 11.8, “Increasing the Size of a Pool,” on page 154**
- **Section 11.9, “Activating and Deactivating Pools,” on page 155**
- **Section 11.10, “Deleting a Pool,” on page 156**
- **Section 11.11, “Verifying and Rebuilding NSS Pools and Volumes,” on page 157**
- **Section 11.12, “ReZIDing Volumes in an NSS Pool,” on page 162**
- **Section 11.13, “Updating eDirectory Pool Objects,” on page 164**
- **Section 11.14, “Preventing Pools from Activating on Multiple Servers,” on page 164**
- **Section 11.15, “What’s Next,” on page 166**

For information about pool snapshots for NSS on NetWare, see **“Managing NSS Pool Snapshots (NetWare)” on page 167.**

## 11.1 Creating Pools with iManager

- 1 In iManager, click *Storage > Pools*.

For instructions, see **Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46.**

- 2 Select a server to manage.



For instructions, see **Section 4.1.5, “Selecting a Server to Manage,” on page 46.**




A list of pools appears in the Pools list.

[Storage Management](#) ► **Pools**

## Pool Management

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

Server:   

Pools:	Details:
<div><div>New...</div><div>Delete</div><div>Rename...</div><div>Activate</div><div>Deactivate</div><div>Increase size...</div><div>Snapshot...</div><div>Update eDirectory</div><div>Deleted Volumes...</div><div>Offline</div></div> <div><div>MOVE_TO</div><div>SYS</div><div>USERS</div></div>	<div>Name: <b>USERS</b></div> <div>Segments: <input type="text" value="0x16"/> </div> <div>Number of Segments: <b>2</b></div> <div>State: <b>Active</b></div> <div>LSS Type: <b>ZLSS</b></div> <div>Share State: <b>Not Sharable for Clustering</b></div> <div>Volumes: <input type="text" value="USERS"/> </div> <div>Number of Volumes: <b>1</b></div> <div>Devices: <input type="text" value="0x3"/> </div> <div>Number of Devices: <b>2</b></div> <div>Total Space: <b>999.00 MB</b></div> <div>Free Space: <b>976.95 MB</b></div> <div>Used Space: <b>22.05 MB</b></div> <div>Purgeable Space: <b>765.34 MB</b></div> <div>Other in-use Space: <b>22.05 MB</b></div> <div>Block Size: <b>4.00 KB</b></div> <div>Creation Date: <b>Jun 9, 2003 3:02:51 AM</b></div> <div>Last Update: <b>Jun 9, 2003 3:02:51 AM</b></div>

- 3 To create a new pool, click *New*.

The New Pool wizard opens to guide you through the process.

- 4 Specify a name for the new storage pool, then click *Next*.

For guidelines about naming pools, see [Section 5.5, “Naming NSS Storage Objects,” on page 69](#).

- 5 Select the devices you want to use and the amount of space to use from each device.

In the Used Size field, type the amount of space in megabytes (MB) to add, up to the amount of free space available for that device. If any entry exceeds a device’s available space, the pool expansion fails and returns an error message.

The size of a given storage partition cannot exceed 2 terabytes (TB). If you want to create an 8 TB pool or volume, you must use at least four partitions when you create the pool.

The Total Pool Size is the sum of the partitions you define plus the current pool size. Initially, the Total Pool Size field displays the current size of the pool. To update the Total Pool Size as you enter values in the Used Size field, click anywhere within the wizard dialog box.

You can obtain space from one or more of the devices listed. Only devices that have free space appear in the list. If no devices are listed, there is no space available to increase the size of the

pool. *Cancel* the wizard, add more devices to the server or free up space on existing devices, then return to the Pool Management page to increase the size of this pool.

**6** Specify if you want to activate the device and if you want to cluster-enable shared devices by doing the following:

- **Local:** The Create Pool dialog box appears with the *Activate on Creation* check box automatically selected. Deselect the box if you do not want to activate this pool on creation.

If the partition that you selected is *not* shareable for clustering (local), click *Finish*.

- **Shared:** If the partition that you selected is shareable, the Create Pool dialog box appears with the *Activate on Creation* and *Cluster Enable on Creation* check boxes selected.

Deselect the *Activate on Creation* check box if you do not want to activate this pool on creation.

If you want to cluster-enable the pool, click *Next*; otherwise, deselect the box, then click *Next*.

New Pool ?

Select device and space

Name: **nifpool**

A pool can be created on one or more storage objects. Select the storage objects for the pool and determine what amount of the available storage space will be used for the pool.

	Used Size (MB)	Device ID	Device Name	Free Size (MB)	Type
<input checked="" type="checkbox"/>	59999	0x13	R5_nifb	59999	Shared
<input checked="" type="checkbox"/>	59999	0x20	R5_nifa	59999	Shared
<input checked="" type="checkbox"/>	59999	0x26	R5_nifc	59999	Shared
<input type="checkbox"/>	0	0x3	[V312-A0-D0:0] WDIGTL WDE4360-1807A3 rev:1.80	1340	Local
<input type="checkbox"/>	0	0x4	[V312-A0-D1:0] SEAGATE ST31230N rev:0300	309	Local
<input type="checkbox"/>	0	0x5	[V312-A0-D2:0] HP 2.13 GB #A2 rev:0180	881	Shared

**Total Pool Size (MB):** 179999

☒ Cluster Enable on Creation

☒ Activate On Creation

<< Back

Finish

Cancel

- 7 (Conditional) If the Cluster Enable on Creation check box is selected, an additional page appears to enable you to specify the cluster information.

**New Pool** ?

**Cluster Information**

Name: **NIFPOOL**

**Shared Pool Clustering Parameters:**

Virtual Server Name:

CIFS Server Name:

IP Address:

**Advertising Protocols:**

☒ NCP

☐ CIFS

☐ AFP

<< Back Finish Cancel

Fill in the following shared pool clustering parameters:

- **Virtual Server Name:** The name assigned by NetWare® to the virtual server that represents the shared pool in the cluster.
- **CIFS Virtual Server Name:** The name assigned by NetWare to the virtual server for handling CIFS (Common Internet File System) requests. This is the name of the server as it appears in a Windows system. (For Linux, CIFS is via Samba.)
- **IP Address:** The IP address that you want to assign the virtual server.

To specify an IP address, tab between the different entries; no dot is required in the fields. For example, if the IP address is 192.168.1.1, type

192 168 1 1

- **Advertising Protocols:** Protocols that give users native file access to data.

Specify one or more advertising protocols by selecting the check boxes of the protocols you want to enable for data requests to this shared pool.

- NetWare Core Protocol™ (NCP™) is the Novell networking protocol.
- CIFS is the Windows networking protocol
- Apple\* File Protocol (AFP) is the Macintosh networking protocol. (AFP is not supported on OES Linux.)

- 8 Click *Finish*.

## 11.2 Viewing Pools on a Server

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

## 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of pools appears in the Pools list. Depending on the number of pools, this can take a few seconds. Avoid clicking again in the page until it refreshes and displays the Pools list.

## 3 If the pool is not in the list, you might need supply the pool name.

For Linux, EVMS discovers and mounts pools at system startup. If you later create and mount a pool from the command line, EVMS does not find the pool and that pool does not show up in the list. You must repeat this process if you dismount the manually created pool at any time. Alternately, you can continue to manage the pool from the command line.

To help EVMS discover a pool not in the list:

**3a** Click *Mount*.

**3b** Type the pool name, then click *OK*.

There is no search function to find the pool; you must supply the pool name.

# 11.3 Viewing Pool Details

The Pool Details area of the Pool Management page displays information about a selected pool in the pools list, as described in the following table:

**Table 11-1** *Explanation of Pool Details*

Device Detail	Description
<i>Name</i>	The descriptive pool name assigned by the administrator.
<i>Mount Point (Linux)</i>	On your Linux system, this is the mount location for the NSS pool. The default mount location is <code>/opt/novell/nss/mnt/.pools/<i>poolname</i></code> , where <i>poolname</i> is the name of the selected pool.
<i>Partitions</i>	A list of all of the partitions that are part of the selected pool. To view information about any of the partitions, select the partition in the drop-down list, then click the <i>View Details</i> icon.
<i>Number of Partitions</i>	The total number of partitions currently assigned to the selected pool.
<i>State</i>	The current state of the selected pool: Active or Deactive. Active pools are available to the users; deactive pools are not available to users.
<i>LSS Type</i>	The type of Loadable Storage System, such as ZLSS, CDDVD, or DOSFAT.
<i>Share State</i>	Shows whether the selected pool is on a device that is marked as Shareable for Clustering or as Not Shareable for Clustering. The system pool cannot reside on a device that is shareable for clustering. Use the Device Management page to set this device attribute.
<i>Volumes</i>	A list of all existing volumes residing in the selected pool. To view information about any of the volumes or to manage any of the volumes, select the volume in the drop-down list, then click the <i>View Details</i> icon. You can also select a pool, then click <i>Volumes</i> in <i>Roles and Tasks</i> .
<i>Number of Volumes</i>	The total number of volumes residing in the selected pool.

Device Detail	Description
<i>Devices</i>	A list of the descriptive device names of all logical devices contributing space in the selected pool.
<i>Number of Devices</i>	The total number of devices currently assigned to the selected pool.
<i>Total Space</i>	The total amount of space assigned to the selected pool.
<i>Free Space</i>	The total amount of space that is currently not in use on the selected pool.
<i>Used Space</i>	<p>The total amount of space that is in use on the selected pool.</p> <ul style="list-style-type: none"> <li>• <b>Purgeable Space:</b> The total amount of space in the selected pool that is currently in use as a salvage area or partitioned space that is not yet otherwise assigned.</li> <li>• <b>Other In-Use Space:</b> The total amount of space in the selected pool that is currently in use and cannot be easily deleted without destroying data.</li> </ul>
<i>Block Size</i>	The maximum amount of data committed to a single write. Possible sizes include 4, 8, 16, 32, or 64 KB. The default setting for NSS is 4 KB.
<i>Creation Date</i>	The time stamp (date and time) that the pool was created.
<i>Last Update</i>	The time stamp (date and time) that the pool was last modified by a management action.

## Procedure

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select a pool to view its details.

Wait for the page to refresh and display the pool’s details.

- 4 (Conditional) The pool must be active to display its details. If the Details field is empty, select the pool, then click *Activate*.

When the page refreshes, you can view the pool’s details.

## 11.4 Viewing Partition Information for a Pool

Although NSS abstracts the partitions underlying the pool structure, you can view information about those partitions.

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.



For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select the pool you want to manage.

Wait for the page to refresh and display the details. The pool must be active to see partition details.

- 4 (Conditional) If the pool is deactive, make sure the pool is selected, then click *Activate*.

After the page refreshes, the Partitions drop-down list is available.

- 5 Click on the arrow next to the Partitions drop-down list to expand the list.

- 6 To view details about a partition, select the partition, then click *View Details*.

A Partition Information page opens where you can view details about the partition.

## 11.5 Viewing Volume Information for a Pool

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of pools appears in the Pools list.

- 3 From the Pool Management page, select a pool.

Wait for the page to refresh and display the details in the Details area.

- 4 To view the volumes in the pool, use one of these methods:

- In the *Details* area, click the arrow on the Volumes drop-down list.

To view details for a volume in the list: Select the volume, then click *View Details*. The Volume Management page opens with the server and volume preselected.

- Expand Storage in *Roles and Tasks*, then click *Volumes*.

The Volume Management page opens with the server preselected. To view details for a volume in the Volumes list: Select the volume, then wait for the page to refresh.

For information about Volume management, see [“Managing NSS Volumes” on page 181](#).

## 11.6 Viewing Device Information for a Pool

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of pools appears in the Pools list.

- 3 From the Pool Management page, select a pool.

Wait for the page to refresh and display the pool’s details in the Details area.

- 4 To view the list, in the Details area, click the arrow on the Devices drop-down list.
- 5 (Optional) To view details for a device in the list: Select the device, then click *View Details*.  
The Device Management page opens with the server and device preselected. Wait for the page to refresh to view the device's details.

For information about device management, see [“Managing Devices” on page 99](#).

## 11.7 Renaming a Pool

The *Rename* option on the Pool Management page lets you to modify the name of the selected pool. For example, you might want to assign a pool name that relates to a department name change.

- 1 In iManager, click *Storage > Pools*.  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).
- 2 Select a server to manage.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).  
A list of pools appears in the Pools list.
- 3 Select the pool that you want to rename.  
Wait for the page to refresh and display the details.
- 4 Click *Rename*.  
This opens the Rename a Pool wizard that guides you through the process.
- 5 Specify a name, then click *Finish*.  
If the name is valid and unique, the pool is successfully renamed.  
If not, you receive an error, and you must repeat the process. For information about valid pool names, see [Section 5.5, “Naming NSS Storage Objects,” on page 69](#).
- 6 (Conditional) If the Pool Management page does not automatically update to show the new name for the pool, in *Roles and Tasks*, click *Pools* to refresh the current page.

## 11.8 Increasing the Size of a Pool

Using the Increase Size option on the Pool Management page expands the storage capacity of a selected pool by adding new partitions. You can increase the size of your storage pools, but you cannot reduce their size.

- 1 In iManager, click *Storage > Pools*.  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).
- 2 Select a server to manage.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).  
A list of pools appears in the Pools list.
- 3 Select the pool that you want to expand.  
Wait for the page to refresh. It displays the pools's details and enables management options.
- 4 Click *Increase Size*.

This opens an Expand a Pool wizard that guides you through the process of adding partitions from available devices.

**Expand a Pool**

**Determine storage objects and space**

Name: **NIFPOOL**

A pool can be created on one or more storage objects. Select the storage objects for the pool and determine what amount of the available storage space will be used for the pool.

	Used Size (MB)	Device ID	Device Name	Free Size (MB)
<input type="checkbox"/>	0	0x2	[V025-A0-D0:0] WDC WD800BB-75CAA0	68745
<input type="checkbox"/>	0	0x3	[V025-A0-D0:1] IC35L080AVVA07-0	65281
<input type="checkbox"/>	0	0xf	[V025-A1-D1:0] IC35L080AVVA07-0	67428

Total Pool Size (MB): **5999**

<< Back Finish Cancel

- 5 Select the devices you want to use and the amount of space to use from each device.

In the Used Space field, type the amount of space in megabytes (MB) to add, up to the amount of Free Space available for that device. If any entry exceeds a device's available space, the pool expansion fails and returns an error message.

Software RAID 1 (mirrored) devices can contain only one pool per device. If you select a RAID-1 device to add a partition to your pool, NSS automatically allocates all of the available space to the pool.

The Total Pool Size is the sum of the partitions you define plus the current pool size. Initially, the Total Pool Size field displays the current size of the pool. To update the Total Pool Size as you enter values in the Used Size, click anywhere within the wizard dialog box.

You can obtain space from one or more of the devices listed. Only devices that have free space appear in the list. If no devices are listed, there is no space available to increase the size of the pool. Cancel the wizard, add more devices to the server or free up space on existing devices, then return to the Pool Management page to increase the size of this pool.

- 6 Click *Finish*, or click *Cancel* at any time to back out of the process.

## 11.9 Activating and Deactivating Pools

You might need to temporarily restrict user access to an pool. Instead of bringing down the server, you only need to deactivate the specific pool.

The *Activate* option on the Pool Management page makes the selected pools and all the volumes in them available for user access. The *Deactivate* option on the Pool Management page takes the selected pools and all the volumes in them temporarily unavailable to users. It does not destroy volumes in the pools, nor does it destroy the data contained in the volumes.

To change the state of a pool:

- 1 In iManager, *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of pools appears in the Pools list.

- 3 Select the pool that you want to activate or deactivate.

Wait for the page to refresh. It displays the pool’s details and enables management options. The State field shows whether the device is Active or Deactive.

- 4 Depending on the pool’s current state, to change the state of the pool:

- Click *Activate*.
- Click *Deactivate*.

## 11.10 Deleting a Pool

You might need to delete an NSS pool to create more free space for other pools. After you delete a pool, ownership of the space for the pool is removed and the free space becomes available. You can assign this space to other NSS pools, as desired, to increase their size.

---

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

---

The Delete option on the Pool Management page removes one or more selected pools from the server, including all member partitions and the data on them.

---

**WARNING:** Deleting a pool destroys all volumes in the pool and destroys all the data in them. These volumes cannot be restored.

---

You cannot delete a system pool using iManager. Deleting a system pool removes the operating system. You need to reinstall NetWare to delete a system pool. However, if the system pool you want to delete does not contain the active operating system (for example, a remnant on a previously used disk), you can rename the pool, then delete it.

Deleting a pool removes the ownership of the space it occupied, freeing the space for reassignment. If the pools you want to delete are active, deactivate them before you delete them.

- 1 In iManager, *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select one or more pools that you want to delete.

Wait for the page to refresh. It displays the pools’s details and enables its management options.

- 4 Click Delete.

- 5 Click Yes to confirm the deletion, or click No to cancel the deletion.

---

**WARNING:** If you click Yes, the pool and all the volumes and data on it are immediately destroyed.

---

## 11.11 Verifying and Rebuilding NSS Pools and Volumes

This section discusses the following:

- [Section 11.11.1, “When to Use Verify and Rebuild,” on page 157](#)
- [Section 11.11.2, “Mounting the Volume to Repair Journaled Errors,” on page 157](#)
- [Section 11.11.3, “Verifying the Pool to Identify Inconsistent Data Blocks,” on page 158](#)
- [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#)
- [Section 11.11.5, “Restoring the Data from Backup Media,” on page 160](#)
- [Section 11.11.6, “Rebuilding NSS Pools to Repair Data Consistency,” on page 160](#)

### 11.11.1 When to Use Verify and Rebuild

Volume errors are typically transactions left unfinished during a system crash of some kind. Most volume errors are fixed automatically during volume mount. If errors persist, use the following utilities to resolve any problems:

Linux	RAVSUI (provides both verify and rebuild options)
	RAVVIEW (reformats log files to human-readable format)
NetWare	PoolVerify
	PoolRebuild

---

**WARNING:** You should rebuild a pool only as a last resort to recover the file system after data corruption. Data loss might occur during a rebuild.

---

You do not need to bring down the server to rebuild pools. NSS allows you to temporarily deactivate individual storage pools while you verify or rebuild the data consistency of the volume. However, when you deactivate a storage pool, users do not have access to any of the volumes in that pool.

### 11.11.2 Mounting the Volume to Repair Journaled Errors

Volume errors are typically transactions left unfinished during a system crash of some kind. This type of error is fixed automatically during volume mount by the journaling feature. Journaling in NSS handles the same level of problems as Vrepair does on NetWare Traditional volumes.

If the errors persist, NSS requires you to use the Verify and Rebuild utilities to resolve any problems.

### 11.11.3 Verifying the Pool to Identify Inconsistent Data Blocks

The Pool Verify option searches the pool for inconsistent data blocks or other errors on the file system. If there are no problems with the pool, it is safe to activate the pool. The verify process is a read-only assessment of the pool.

#### Linux

- 1 Before you verify or rebuild, you must place the pool in maintenance mode.

- 1a At a terminal prompt, enter

```
nsscon
```

- 1b In nsscon, enter

```
nss /poolmaintenance=poolname
```

- 2 To verify a pool, enter the following at a command prompt

```
ravsui verify poolname
```

- 3 Use RAVVIEW to read the logs to determine whether a restore from backup media or a pool rebuild is necessary.

For information, see [Section A.14, “RAVVIEW \(Linux\),” on page 380](#).

- 4 To rebuild a pool, enter the following at a command prompt

```
ravsui rebuild poolname
```

For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378](#).

#### NetWare

- 1 Place the pool in maintenance mode. At a terminal prompt, enter

```
nss /poolmaintenance=poolname
```

- 2 Verify the pool.

- 2a At the server console, enter

```
nss /poolverify=poolname
```

- 2b Select the pool name from the displayed list.

- 3 Review any errors on screen or in the *volume\_name.rlf* file, located at the root of the DOS drive.

- 4 For each volume, do one of the following:

- **Volume with Errors:** If a volume has errors, its error messages are reported on screen and in the Verify Log. The volume remains in maintenance mode. Read the *filename.vlf* file at the root of the DOS drive on your server for information about the errors and transactions. Read the *filename.rlf* file for information about data that has been lost by the pruning of leaves in the data structure.
- **Volume without Errors:** If a volume has no errors, it automatically reverts to the active state.

## 11.11.4 Reviewing Log Files for Errors

Make sure to check the error log whenever an NSS volume does not come up in active mode after a verify or rebuild.

### Log Files and Onscreen Display

Messages are written to the following logs:

**Table 11-2** Location of Log Files for the NSS Pool Rebuild and Verify Commands

Platform	Log	Purpose
Linux	<code>/var/opt/novell/log/nss/ rav/filename.vbf</code>  This is the default location, but you can specify the location and the filename.	Log of the pool verify process using ravsui verify. If a volume has errors, the errors are displayed on the screen and written to this log file of errors and transactions.  On Linux, use the RAVVIEW utility to read logs. For information, see <a href="#">Section A.14, “RAVVIEW (Linux),” on page 380</a> .
	<code>/var/opt/novell/log/nss/ rav/filename.rtf</code>	Log of the pool rebuild process using ravsui rebuild. It contains information about data that has been lost during a rebuild by the pruning of leaves in the data structure.
NetWare	<code>filename.vlf</code> , located at the root of the server's DOS drive	Log of the pool verify process using poolverify. If a volume has errors, the errors are displayed on the screen and written to this log file of errors and transactions.
	<code>filename.rlf</code> , located at the root of the server's DOS drive	Log of the pool rebuild process using poolrebuild. It contains information about data that has been lost during a rebuild by the pruning of leaves in the data structure.

Whenever you rebuild a volume, the new information is appended at the end of the log. If you want to keep old log files intact, rename it or move it to another location.

### Warnings Reported

Warnings indicate that there are problems but that there is no threat of data corruption. Performing a data restore or a pool rebuild is optional.

### Errors Reported

Errors indicate that there are physical integrity problems and data corruption will definitely occur, or it will continue to occur, if you continue to use the volume as it is. Do one of the following:

- Restore the pool or volume from an existing copy on backup media. For information, see [Section 11.11.5, “Restoring the Data from Backup Media,” on page 160](#).
- If restoring from backup is not an option, use the Rebuild utility. For information, see [Section 11.11.6, “Rebuilding NSS Pools to Repair Data Consistency,” on page 160](#).

## No Errors Reported, but Cannot Create Files or Directories

If the verify log does not report errors, but you continue to be unable to create files or directories on a volume, it might be because the file numbers have exceeded the maximum file numbering field. You might need to rebuild with the ReZID option. For information about how to decide if a ReZID is needed, see [Section 11.12, “ReZIDing Volumes in an NSS Pool,” on page 162](#).

## 11.11.5 Restoring the Data from Backup Media

After a pool verify reports problems with a pool’s physical integrity, consider whether there is a satisfactory backup of the pool available to restore your system. Use your standard methods of restoring pools from backups. If restoring from backup is not an option, use the Rebuild utility.

## 11.11.6 Rebuilding NSS Pools to Repair Data Consistency

Rebuild uses the existing leaves of an object tree to rebuild all the other trees in the system to restore visibility of files and directories. It checks all blocks in the system. Afterwards, the NSS volume remains in maintenance mode if there are still errors in the data structure; otherwise, it reverts to the active state.

---

**WARNING:** Data will be lost during the rebuild.

---

### What Rebuilding Does Not Do

Rebuilding a pool does not fix problems for the following:

- File system trustee assignments, trustee rights, and inherited rights filters
- File system attributes for files and directories
- Opportunistic locking
- Content of files

### Before You Rebuild a Volume

If all of the following conditions exist, then you should rebuild the pool to restore its data integrity.

- If the pool can be mounted, mount its volume to allow the NSS journaling feature repair any transactional errors that occurred during a system failure. Afterwards, there if there are still problems, verify the pool.
- The Pool Verify process reports errors in the physical integrity of the volume that would definitely cause data corruption if no action is taken.
- You have no viable backup media to restore the pool to the last known good state.
- More data will be lost from continued data corruption than will be lost from rebuilding the pools now.
- (NetWare) If you are not sure whether you can tolerate a system rebuild, take a pool snapshot and run the rebuild against the pool snapshot instead. (Pool snapshot is not available on Linux.)

If the Pool Verify process did not report errors, but you cannot create files or directories, you should run rebuild with the ReZID option instead. For information, see [Section 11.12, “ReZIDing Volumes in an NSS Pool,” on page 162](#).



## Linux

- 1 Before you verify or rebuild, you must place the pool in maintenance mode.
  - 1a At a terminal prompt, enter  

```
nsscon
```
  - 1b In nsscon, enter  

```
nss /poolmaintenance=poolname
```
- 2 If you have not already run a pool verify, verify the pool. At a terminal console, enter  

```
ravsui verify poolname
```

For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378](#).
- 3 Review any errors on screen or in the *filename.vbf* file, located where you specified.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).
- 4 To rebuild a pool, enter the following at a command prompt  

```
ravsui rebuild poolname
```

For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378](#) for options to set the pruning parameters for the rebuild.

Rebuilding can take several minutes to several hours, depending on the number of objects in the pool.
- 5 Review the log on screen or in the *filename.rtf* file to learn what data has been lost during the rebuild.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).
- 6 Do one of the following:
  - **Errors:** If errors still exist, the pool remains in the maintenance state. Repeat the pool verify to determine the nature of the errors, then call Tech Support for assistance.
  - **No Errors:** If errors do not exist, the pool’s volumes are available for mounting.

## NetWare

- 1 Place the pool in maintenance mode. At a terminal prompt, enter  

```
nss /poolmaintenance=poolname
```
- 2 If you have not already run a pool verify, verify the pool to determine if a rebuild is needed.

At the server console, enter  

```
nss /poolverify=poolname
```
- 3 Review any errors on screen or in the *filename.vlf* file, located at the root of the DOS drive.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).
- 4 To run Rebuild, enter the following command at the server console:  

```
nss /poolrebuild=poolname
```

Replace *poolname* with the name of the pool you want to rebuild.

Rebuilding can take several minutes to several hours, depending on the number of objects in the pool.

- 5 Read the `filename.rlf` file at the root of the DOS drive on your server for information about data that has been lost.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).

- 6 Do one of the following:
  - **Errors:** If errors still exist, the pool remains in the maintenance state. Repeat the pool verify to determine the nature of the errors, then call Tech Support for assistance.
  - **No Errors:** If errors do not exist, the pool is activated automatically. Mount the volumes.

## 11.12 ReZIDing Volumes in an NSS Pool

When a file is created it is assigned a unique file number, called a ZID. In NSS, the maximum number of file IDs available is a 64-bit number, which provides for up to 8 trillion (8E12) ZIDs, so NSS was designed to not re-use ZIDs. However, NCP clients and other traditional applications can only work with 32-bit IDs, which supports up to 4 billion (4E9) ZIDs. NSS restricts ZIDs, thus the number of files, to the lower value.

It is possible for a file system to reach the 32-bit limit on the ZID number. For example, if a lot of files and directories are regularly created, exist for a short time, and are then deleted, the ZIDs are used up at that rate. Otherwise, reaching the upper limit is something that happens rarely.

### 11.12.1 Understanding ReZID

The ReZID option for a pool rebuild changes the ZIDs for all the files on the volume, freeing ZIDs in the less than 32-bit range so they are available for creating new files and directories. The ZID of a file is an internal file system bit of information and is not viewable. However, you can view the highest ZID number for each volume is reported when you verify a pool.

---

**IMPORTANT:** Adding the ReZID step adds a third review of the pool and can increase the time of a rebuild by 50%.

---

### 11.12.2 When to ReZID Volumes During a Pool Rebuild

There are no errors reported if ZIDs are nearing the upper limit of 4 billion for a volume. However, you can view the `highestZid` field for the volume, and if the number is close to the 32-bit limit (4 billion), you should reZID the volume.

For OES on Linux, if the highest ZID is 2 billion or greater, the reZID option occurs automatically when you rebuild a pool.

You might get errors creating a file or directory that will suggest a reZID needs to be done. For example:

- NDS database is locked.
- Server hangs at the end of loadstage 1.
- Cannot copy to a volume.
- Cannot copy to `sys:` volume.

The NSS error that NSS (internally) returns is 20108 `zERR_ZID_GREATED_THAN_32_BITS`. It also sends a volume alert that reZID needs to be done to the server console. The application only gets a generic error when creating the file.

After rebuilding a pool with the ReZid option, the errors you were getting when creating files and directories no longer occur. You can also verify the pool, then check the `highestZID` number to know that it is well under the 4 billion ZIDs limit.

### 11.12.3 ReZIDing Volumes on Linux

- 1 Before you verify or rebuild, you must place the pool in maintenance mode.

- 1a At a terminal prompt, enter

```
nsscon
```

- 1b In `nsscon`, enter

```
nss /poolmaintenance=poolname
```

- 2 If you have not already verified the volume, enter the following at a command prompt

```
ravsui verify poolname
```

For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378](#).

- 3 Review any errors on screen or in the `filename.vbf` file, located where you specified.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).

- 4 Rebuild a pool by entering the following at a command prompt

```
ravsui --rezid=zid rebuild poolname
```

Replace `zid` with the value of a threshold to cause a rezid of a volume. The default value is `0xffffffff`. For information, see [Section A.13, “RAVSUI \(Linux\),” on page 378](#) for options to set the pruning parameters for the rebuild.

For NSS on OES Linux, rebuild automatically causes a reZID of a volume if rebuild finds a ZID over 2 billion.

This checks all blocks in the system. Rebuilding can take several minutes to several hours, depending on the number of objects in the pool. For all systems, rezid adds a third pass to rebuild which increases the time to rebuild a volume by about 50%.

- 5 Review the log on screen or in the `filename.rtf` file to learn what data has been lost during the rebuild.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).

- 6 Do one of the following:

- **Errors:** If errors still exist, the pool remains in the maintenance state. Repeat the pool verify to determine the nature of the errors, then contact Technical Support for assistance.
- **No Errors:** If errors do not exist, the pool’s volumes are mounted automatically.

### 11.12.4 ReZIDing Volumes on NetWare

- 1 Place the pool in maintenance mode. At a terminal prompt, enter

```
nss /poolmaintenance=poolname
```

- 2 If you have not already run a pool verify, verify the pool to determine if a rebuild is needed.

At the server console, enter

```
nss /poolverify=poolname
```

- 3 Review any errors on screen or in the *filename.vlf* file, located at the root of the DOS drive.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).

- 4 To run Rebuild with the ReZID option, enter the following command at the server console:

```
nss /poolrebuild=poolname /rezid
```

Replace *poolname* with the name of the pool you want to rebuild.

This checks all blocks in the system. Rebuilding can take several minutes to several hours, depending on the number of objects in the pool. For all systems, rezid adds a third pass to rebuild which increases the time to rebuild a volume by about 50%.

- 5 Review the log on screen or in the *filename.rtf* file to learn what data has been lost during the rebuild.

For information, see [Section 11.11.4, “Reviewing Log Files for Errors,” on page 159](#).

- 6 Do one of the following:

- **Errors:** If errors still exist, the pool remains in the maintenance state. Repeat the pool verify to determine the nature of the errors, then call technical support for assistance.
- **No Errors:** If errors do not exist, the pool’s volumes are mounted automatically.

## 11.13 Updating eDirectory Pool Objects

On the Pool Management page, the Update eDirectory option adds or updates the Novell eDirectory™ pool object at the same context level as the server. NSS searches for the object. If the pool object exists, NSS prompts you with two options: delete and replace the existing object, or retain the existing object. If the pool object does not exist, NSS adds the object to the context level.

Use this option to update eDirectory after you have modified a pool’s parameters or renamed it.

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select the pool you want to update.

Wait for the page to refresh. It displays the pools’s details and enables its management options.

- 4 Click *Update eDirectory*.

## 11.14 Preventing Pools from Activating on Multiple Servers

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 does 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 Shareable 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 Shareable for Clustering flag in order to force an activation of a shared pool, or if you use `override=shared` when activating a pool. If MSAP detects a conflict, it deactivates the pool before massive corruption occurs.

If you unload Novell Cluster Services™ NLM™ software, or if you are not running it, pool MSAP provides an extra level of protection. The clustering software watches pools that are marked with the Shareable for Clustering attribute; MSAP detects conflicting connection from multiple servers and prevents corruption of pools even on devices that are marked as Not Shareable 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 eDirectory object ID is lost.

If MSAP is enabled, all active NSS pools are 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.

---

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

---

## 11.14.1 Managing MSAP Using Server Console Commands

The management file for pool MSAP is `_admin\manage_nss\pool\poolname\zlss\msap.xml`. One file exists for each pool. The `msap.xml` file contains MSAP statistics for the pool. The MSAP attribute is displayed in the Enabled Attributes (`<enabledAttributes>`) tag of the `poolinfo.xml` management file.

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

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

### Syntax

Use the following table understand MSAP console commands.

MSAP Console Command	Description
<code>nss /msapserver</code>	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.
<code>nss /nomsapserver</code>	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.
<code>nss /msaprebuild</code>	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 rebuilder ownership of the pool.
<code>nss /poolmsap=<i>poolname</i></code>	Enables MSAP for a given pool on the server. MSAP is enabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now enabled.)
<code>nss /nopoolmsap=<i>poolname</i></code>	Disables MSAP for a given pool. Use the command when the pool is activated. MSAP is disabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now disabled.)
<code>nss /pools</code>	Displays the message “Multi-use detect” for pools that have MSAP enabled.

## 11.15 What's Next

You can now create volumes in the NSS pools you created. For information, see [Chapter 13, “Managing NSS Volumes,”](#) on page 181.

# Managing NSS Pool Snapshots (NetWare)

# 12

In NetWare® 6.5, Novell Storage Services™ supports pool snapshots to improve backup and restore services. A pool snapshot is a metadata copy of a storage data pool that preserves a point-in-time view of a data pool. When you back up a pool from its pool snapshot, your backup can capture every file in the pool, even those that are in use at the time. You can create, manage, and delete a pool snapshot for any pool on your server.

A pool snapshot also supports file recovery. After you take a snapshot, you can activate it at a later time to access the original pool's data as it existed at the time of the snapshot. Both the pool and its snapshots 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.

When pool snapshots are active, server performance can decrease slightly because the number of disk writes increases, depending on the volatility of your data and the number of pool snapshots that are active.

This section discusses the following:

- [Section 12.1, “Benefits of Pool Snapshots,” on page 167](#)
- [Section 12.2, “How the Pool Snapshot Works,” on page 168](#)
- [Section 12.3, “Understanding Key Concepts for Pool Snapshots,” on page 169](#)
- [Section 12.4, “Managing Pool Snapshots Using iManager,” on page 172](#)
- [Section 12.5, “Managing Pool Snapshots Using Commands,” on page 179](#)

## 12.1 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 back up, 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 snapshots 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 snapshot 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 user. With traditional backups, applications might 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 might 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 might 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 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.

## Re-Creating 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 snapshot version of the data. For example, in a software development environment, engineers might 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.

# 12.2 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 snapshot is a metadata-based copy that stores only the blocks of data that change subsequent to the instant of the snap. The snapshot 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 snapshot 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. Any open files are seen by the snapshot feature as being closed after these outstanding writes occur. Then, it snapshots the now-stable pool, and allows data transaction activity to resume.

The quiescence process provides a transactionally consistent image at the instant the snapshot is made. Because the snapshot is consistent, it is not necessary to check the consistency of the file system or database when you activate the snapshot 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 snapshot exists before it is archived or deleted. Currently, a combination of up to 500 snapshots can exist on any given stored-on pool. The more snapshots you must traverse to find data, the longer it takes to retrieve the data, and the slower the perceived response time to the user.

## 12.3 Understanding Key Concepts for Pool Snapshots

This section discusses the following topics:

- [Section 12.3.1, “Creating a Pool Snapshot,” on page 169](#)
- [Section 12.3.2, “Naming Pool Snapshots,” on page 169](#)
- [Section 12.3.3, “Choosing the Stored-On Pool,” on page 170](#)
- [Section 12.3.4, “Maintaining the Stored-on Pool,” on page 170](#)
- [Section 12.3.5, “Activating Pool Snapshots,” on page 171](#)
- [Section 12.3.6, “Deleting Pool Snapshots,” on page 171](#)

### 12.3.1 Creating a Pool Snapshot

Create a pool snapshot when you want to capture a point-in-time view of a active data pool Both the original pool and the stored-on pool must be active when you make the snapshot. For instructions on creating a pool snapshot, see [“Creating a New Pool Snapshot” on page 172](#).

### 12.3.2 Naming Pool Snapshots

You name a pool snapshot at the time you order the snap. Specify a unique name for each snapshot. Because the name also serves as the snapshot’s pool name when active, the name you give it should be unique among snapshots and among pools. The combination of snapshot name and time stamp when the snapshot was taken can 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 for snaps. This is particularly important if you plan to create a series of snapshots for multiple pools. The snapshot name is generally a modified version of the original pool’s name to support easy identification of all snapshots for any given pool.

---

**IMPORTANT:** The snapshot name can be 2 to 16 characters.

---

#### Default Naming Scheme

For example, by default, NSS adds 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 snapshots taken for this pool. In this scenario, snapshot names might be PoolA\_S1, PoolA\_S2, and so on for snapshots of PoolA. If you delete older snapshots and reuse the lower sequence numbers, a simple sort by snapshot 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.

## Alternate Naming Scheme

For some scenarios, it might be desirable to adopt more meaningful and sortable naming conventions. Currently, you can sort by snapshot name only, not by time stamp. If you create multiple snapshots 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 “\_Sddtttt”. Snapshot names for PoolA might be PoolA\_S300600, PoolA\_S301800, PA\_S310600, PA\_S311800, and so on. Of course, the time stamp shows the exact time that the snapshot was taken, because there is a slight delay between ordering a snapshot and the instant of the snapshot itself.

## Considerations for Naming

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 occurs when you try to activate a pool snapshot or another regular pool with the same name. To resolve the conflict, the Media Manager dynamically modifies the snapshot name with an \_*n* (sequential number). This makes the snapshot name unique with respect to the active pool list during the time that the pool snapshot is active and functioning as a pool.

### 12.3.3 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 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.

When creating your first NSS pool-level snapshot for a pool, we highly recommend that the destination pool specified for the snapshot data is different than the source pool you are snapshotting. The exception to this is clustering, where the destination pool for the snapshot data must be the same as the source pool you are snapshotting.

If the destination pool that stores the snapshot data runs out of space, it causes write errors to happen on the pool(s) being snapshot. Please allow ample space for snapshots to grow over time when sizing the snapshot’s destination pool. The amount of space required depends on the number of snapshots, the snapshot retention policy, and the turnover rate for data in the source pool.

### 12.3.4 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 are 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 has 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 causes 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.

---

### 12.3.5 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. After the pool snapshot is active, it appears by its user-assigned snapshot 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 snapshot name is no longer listed in the pool list.

### 12.3.6 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. If you create a pool snapshot and specify its stored-on pool as the same pool where its original data resides, and more than one pool snapshot exists when you are ready to delete a snapshot, delete its oldest snapshot in a first created, first deleted sequence.

---

**IMPORTANT:** Delete pool-level snapshots in a first-created, first deleted manner, deleting the oldest snapshot of the pool first.

---

Users can access the snapshot for reads and writes when the snapshot is activated as a pool, just as they would any pool. Because any given pool snapshot can rotate into an archived status or be deleted, you need to be aware if the snapshot is in use before you delete it.

In NetWare 6.5 SP3 and later, when a snapshot is deleted, Media Manager automatically zeroes out all the blocks that it used. It no longer turns on shredding for the internal volume when creating a snapshot. You can apply the system shredding feature to a snapshot pool by entering

```
mm snap shred snapshotname = on
```

Replace *snapshotname* with the unique name of the pool snapshot. Shredding occurs only in the internal volume of the pool where you are storing the snapshot.

If you turn on snapshot shredding for a snapshot, you can have only one snapshot at a time for its original pool. When you delete the snapshot, do not create another snapshot until the system has had enough time to complete the shredding process. Depending on the size of the snap, allow more time for larger amounts of data to shred. The typical delay time is a few minutes.

If any pool snapshot was created on this pool prior to the NetWare 6.5 SP3 release, shredding is already turned on. You can turn shredding off by entering

```
mm snap shred snapshotname = off
```

Replace *snapshotname* with the unique name of the pool snapshot. Shredding is turned off only in the internal volume of the pool where you are storing the snapshot. It does not affect other snapshots.

If the system goes down during the delete, the zeroing process does not continue after the system restarts. Some nonzeroed data blocks might remain. If it is important that all snapped data be removed from the system, you can turn on shredding, with the limitations as noted above.

If shredding is turned on, the delete process zeroes out the data, and then starts shredding the blocks where data was stored. If the system goes down during the delete, the zeroing process does not continue after the system restarts, but the shredding automatically continues after the system restarts.

## 12.4 Managing Pool Snapshots Using iManager

This section discusses the following management tasks:

- [Section 12.4.1, “Creating a New Pool Snapshot,” on page 172](#)
- [Section 12.4.2, “Viewing Pool Snapshot Details,” on page 174](#)
- [Section 12.4.3, “Renaming a Pool Snapshot,” on page 175](#)
- [Section 12.4.4, “Activating and Deactivating a Pool Snapshot,” on page 175](#)
- [Section 12.4.5, “Moving an Active Pool Snapshot Offline from the Pools Page,” on page 177](#)
- [Section 12.4.6, “Deleting a Pool Snapshot,” on page 178](#)
- [Section 12.4.7, “Modifying the Stored-On Pool,” on page 178](#)

### 12.4.1 Creating a New Pool Snapshot

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 In the server field on the Pool Management page, select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3** In the Pools list, select the active pool that you want to be the original pool to be snapped, then click Snapshot.


Both the original pool and the stored-on pool must be active when you make the snap. The original pool and stored-on pool can be the same. To activate a pool, select the pool, then click *Activate*.

---

**IMPORTANT:** If the selected pool is a pool snapshot, this option is not available. You cannot take a pool snapshot of a snapshot.

---

- 4** In the Define Snapshot dialog box, name the snapshot and specify its stored-on pool.



- 4a** Name the snapshot.

By default, the name is the name of the original pool, appended with `_Sn`, where `n` represents a number. You can change this name; it should be something that you can easily relate to the original pool and the sequence in which the snapshot was taken. By default, `_SV` is added to volume names to indicate the storage object is a volume in a pool snapshot.

For example, if your original pool is named `users`, its default snapshot name is `users_s1`.

You can change the default name, if desired. For information about naming considerations, see [“Naming Pool Snapshots” on page 169](#).

- 4b** Select the stored-on pool from the list of pools available on the selected server.

By default, NSS specifies the stored-on pool to be the same as the original pool. You can choose any of the pools on the server.

---

**IMPORTANT:** After you specify a stored-on pool, that pool is the default stored-on pool for all subsequent snapshots made of the original pool. To modify the stored-on pool associated with an original pool, you must delete all existing snapshots for the original pool, then begin again to create a new pool.

---

- 4c** Click *OK* to create the snapshot, or click *Cancel* to back out of the process.

- 5 After NSS creates the pool snapshot, NSS opens to the Pool Snapshot Management page so that you can further manage the snapshot. The Snapshots list contains the newly created snapshot.

[Storage Management](#) ► [Pool](#) ► [Snapshot](#)

## Pool Snapshot



Create and manage snapshots for any pool on this server. All snapshots for a given pool reside on the pool you select as the Stored on Pool when you create its first snapshot. Activate a snapshot to access data on it.

Pools:

Snapshots:		Details:
New...	USERS_S1	Snapshot name: <b>USERS_S1</b>
Delete		Snapshot of Pool: <b>USERS</b>
Rename...		Pool Size: <b>999.97 MB</b>
Online		Stored on Pool: <b>USERS</b>
Offline		State: <b>Deactive</b>
		Creation Date: <b>Jun 11, 2003 12:44:13 PM</b>
		Size: <b>19.53 MB</b>

OK

## 12.4.2 Viewing Pool Snapshot Details

In iManager, you can view the following details about a pool snapshot:

**Table 12-1** Explanation of Pool Snapshot Details

Snapshot Details	Description
<i>Snap Name</i>	The descriptive pool snapshot name assigned by the administrator.
<i>Snapshot of Pool</i>	The name of the original pool that was snapped.
<i>Pool Size</i>	The total size allocated for the pool snapshot.
<i>Stored-On Pool</i>	<p>The pool where the snapshot metadata resides. This can be the same as the original pool.</p> <p>The current default in iManager is that the stored-on pool and the original pool are the same. If you want to store snapshots on a different pool, you must use the command line interface. For instructions, see <a href="#">Section 12.5, "Managing Pool Snapshots Using Commands,"</a> on page 179.</p>

Snapshot Details	Description
<i>State</i>	The current state of the selected pool: Show or Hide. If you show pool snapshots (Online), they appear in the pool list. If you hide pool snapshots (Offline), they do not appear in the pools list.
<i>Creation Date</i>	The time stamp that shows the date and time the original pool was snapped.
<i>Size</i>	The current size of the pool snapshot. The pool's snapshot can continue to add metadata for the snapshot until it reaches the size allocated.

To view details of a pool snapshot:

- 1 In iManager, log on to the tree for the server you want to manage.
- 2 In *Roles and Tasks*, click *Storage > Pools* to open the Pool Management task
- 3 Select the server that contains the original pool of the snapshot you want to view.
- 4 In the *Pools* list, select the original pool of the snapshot you want to view, then click *Snapshot*.  
This opens the Pool Snapshot page.
- 5 Select the snapshot, then view its details in the Details field.  
It can take a few seconds for iManager to load the details about the snapshot.

### 12.4.3 Renaming a Pool Snapshot

You can rename a pool snapshot. For example, you might want to assign a new name for the pool snapshot to correspond to a similar name change to its original pool.



- 1 In iManager, log on to the tree for the server you want to manage.
- 2 In *Roles and Tasks*, click *Storage > Pools* to open the Pool Management task
- 3 Select the server that contains the original pool of the snapshot you want to view.
- 4 In the *Pools* list, select the original pool of the snapshot you want to view, then click *Snapshot*.  
This opens the Pool Snapshot page.
- 5 Select the snapshot in the Snapshots list.  
It can take a few seconds for iManager to load the details about the snapshot.
- 6 Click *Rename*.  
This opens the Pool Snapshot Rename wizard. For information about naming conventions in NetWare 6.5, see [Section 5.5, “Naming NSS Storage Objects,” on page 69](#).
- 7 Specify the new name you want for the pool snapshot.
- 8 Click *Finish* to accept the new name, or click *Cancel* to back out of the process.  
When the page refreshes, you can see the renamed pool snapshot in the Snapshots list.

### 12.4.4 Activating and Deactivating a Pool Snapshot

Although the pool snapshot is deactive by default, the snapshot functions are working in the background to capture any changes being made to the original pool. For pool snapshots, Deactive and Active are conditions related to the visibility of the pool snapshot to users as a pool. Active pool snapshots are online; deactive pool snapshots are offline.

To deactivate a pool snapshot and take it offline from the Pool Management page, see “[Moving an Active Pool Snapshot Offline from the Pools Page](#)” on page 177.

To activate or deactivate a pool snapshot for user access as a pool:

- 1** In iManager, log on to the tree for the server you want to manage.
- 2** In *Roles and Tasks*, click *Storage > Pools* to open the Pool Management task
- 3** Select the server that contains the original pool of the snapshot you want to view.
- 4** In the Pools list, select the original pool of the snapshot you want to view, then click *Snapshot*.  
This opens the Pool Snapshot page.
- 5** Select one or more pool snapshots that you want to manage, then modify the state:
  -  To change the state of the pool from Deactive to Active, click *Online*.  
This makes the selected pool snapshots and all the volumes in them available for user access.
  -  To change the state of the pool from Active to Deactive, click *Offline*.  
This makes the selected pool snapshots and all the volumes in them temporarily unavailable to users. It does not destroy volumes in the snapshots, nor does it destroy the data contained in the volumes.



## 12.4.5 Moving an Active Pool Snapshot Offline from the Pools Page



When online, pool snapshots appear and function as a regular pool in the Pools list on the Pools Management page.




**Figure 12-1** Sample Pool Snapshot When Online

[Storage Management](#) ► **Pools**

### Pool Management

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

Server:   

Pools:	Details:
<div><div>New... Delete Rename... Activate Deactivate Increase size... Snapshot... Update eDirectory Deleted Volumes... Offline</div><div>MOVE_TO SYS USERS <b>USERS_S1</b></div></div>	<div><div>Name: <b>USERS_S1 (snapshot)</b></div><div>Segments: <input type="text" value=""/></div><div>Number of Segments: <b>0</b></div><div>State: <b>Active</b></div><div>LSS Type: <b>ZLSS</b></div><div>Share State: <b>Not Sharable for Clustering</b></div><div>Volumes: <input type="text" value="USERS_SV"/></div><div>Number of Volumes: <b>1</b></div><div>Devices: <input type="text" value="0x4"/></div><div>Number of Devices: <b>1</b></div><div>Total Space: <b>999.00 MB</b></div><div>Free Space: <b>957.42 MB</b></div><div>Used Space: <b>41.58 MB</b></div><div>Purgeable Space: <b>765.34 MB</b></div><div>Other in-use Space: <b>41.58 MB</b></div><div>Block Size: <b>4.00 KB</b></div><div>Creation Date: <b>Jun 11, 2003 12:45:20 PM</b></div><div>Last Update: <b>Jun 11, 2003 12:45:20 PM</b></div></div>

To move an active pool snapshot offline:

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of pools appears in the Pools list. Online pools also appear in the list.

- 3 Select the pool snapshot that you want to move offline.

Wait for the page to refresh. It displays the pools's details and enables its management options.

- 4 Verify that the pool you selected is the pool snapshot and not the original pool.

If it is a snapshot, in the Details area, the name of the pool should be followed “(snapshot)” and the Offline button is activated.

- 5 Click *Offline*.

## 12.4.6 Deleting a Pool Snapshot

Use the Delete option to permanently remove one or more selected pool snapshots from the server. Deleting a pool snapshot removes the ownership of the space it occupied, freeing the space for reassignment.

---

**IMPORTANT:** If you create a pool snapshot and specify its stored-on pool as the same pool where its original data resides, and more than one pool snapshot exists when you are ready to delete a snapshot, delete its oldest snapshot in a first created, first deleted sequence.

---

To delete a pool snapshot:

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 In the Pools list, select the pool that has a snapshot you want to delete.

Wait for the page to refresh. It displays the pools's details and enables its management options.

If a pool snapshot is online, it might appear in the Pools list. If the selected pool is a snapshot pool, in the Details area, the name of the pool should be followed “(snapshot)” and the Offline button is activated. If you select a pool snapshot in the Pools list, the page refreshes and enables the Snapshot option.

- 4 Click *Snapshot*.

If the selected pool has snapshots available or is an online pool snapshot, the Pool Snapshot page opens. If the selected pool has no pool snapshots associated with it, the Snapshot

- 5 On the Pool Snapshot page, select one or more snapshots in the Snapshots list that you want to delete.

Verify that the pool snapshots you selected are the versions you want to delete.

- 6 Click *Delete*.

- 7 A message requests you to confirm (*Yes*) or back out of (*Cancel*) this action.

## 12.4.7 Modifying the Stored-On Pool

After you specify a stored-on pool, that pool is the default stored-on pool for all subsequent snapshots made of the original pool. To modify the stored-on pool associated with an original pool, you must delete all existing snapshots for the original pool, then begin again to create a new pool.

## 12.5 Managing Pool Snapshots Using Commands

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

### Syntax

Snapshot Management 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.
<code>mm snap rename <i>snapname newsnapname</i></code>	Renames the pool as specified
<code>mm snap shred <i>snapname</i> = on</code>	Turns on shredding for the internal volume of the pool where the snapshot is stored.  For information, see <a href="#">“Deleting Pool Snapshots” on page 171</a> .
<code>mm snap shred <i>snapname</i> = off</code>	Turns off shredding for the internal volume of the pool where the snapshot is stored.  For information, see <a href="#">“Deleting Pool Snapshots” on page 171</a> .
<code>mm snap list</code>	Lists the snapshots on the server.

### Attributes

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

Attributes for Pool Snapshot Commands	Description
<code>snappool</code>	The unique name of the original data pool that you are snapping; must be 2 to 16 characters.
<code>datapool</code>	The unique name of the stored-on pool for the pool snapshot; must be 2 to 16 characters.
<code>snapname</code>	The unique snapshot 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.



Novell® Storage Services™ uses storage volumes to logically organize your data. After creating NSS pools, you can create any number of NSS volumes for each pool, depending on the physical space available.

This section describes how to configure and manage NSS volumes by completing the following tasks:

- [Section 13.1, “Understanding NSS Volumes,” on page 181](#)
- [Section 13.2, “Understanding Volume Attributes,” on page 183](#)
- [Section 13.3, “Creating and Configuring Unencrypted NSS Volumes with iManager,” on page 186](#)
- [Section 13.4, “Configuring Encrypted NSS Volumes with NSSMU,” on page 188](#)
- [Section 13.5, “Updating eDirectory Volume Objects,” on page 188](#)
- [Section 13.6, “Viewing the Details of an NSS Volume,” on page 189](#)
- [Section 13.7, “Viewing Properties of an NSS Volume,” on page 189](#)
- [Section 13.8, “Modifying Attributes of an NSS Volume,” on page 191](#)
- [Section 13.9, “Renaming an NSS Volume,” on page 192](#)
- [Section 13.10, “Activating and Deactivating an NSS Volume,” on page 192](#)
- [Section 13.11, “Activating and Deactivating CDs and DVDs as NSS Volumes \(NetWare\),” on page 193](#)
- [Section 13.12, “Activating and Deactivating CD and DVD Image Files as NSS Volumes \(NetWare\),” on page 194](#)
- [Section 13.13, “Activating and Deactivating DOS Partitions as NSS Volumes \(NetWare\),” on page 195](#)
- [Section 13.14, “Mounting and Dismounting an NSS Volume \(Linux\),” on page 195](#)
- [Section 13.15, “Using NFS Mount for NSS Volumes \(Linux\),” on page 196](#)
- [Section 13.16, “Mounting an NSS Volume on Linux to Support Case Insensitive File Names \(Linux\),” on page 197](#)
- [Section 13.17, “Mounting a Disk in a USB Floppy Drive as a Local FAT Volume \(NetWare\),” on page 198](#)
- [Section 13.18, “Deleting an NSS Volume,” on page 199](#)
- [Section 13.19, “Verifying and Rebuilding NSS Pools and Volumes,” on page 200](#)
- [Section 13.20, “What’s Next,” on page 200](#)

## 13.1 Understanding NSS Volumes

NSS volumes are logical storage media that acquire space from pools of storage. When you create a logical volume, you can either assign it a fixed amount of space or allow it to expand to the pool size. To grow a volume, you might need to add new segments to grow the pool first, up to the maximum pool size of 8 TB. On NetWare®, you can also use Novell Distributed File Services Move

and Split functions to relocate the volume (or part of a volume) to a new pool on the same or different server in the same tree.

If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool size, using the overbooking feature, although real total size is bound by physical limitations. Because space is allocated to volumes as needed, a volume might not reach its quota. As the overbooked volumes consume the available physical space, you need to add more disk space to the pool to accommodate the growth, or consider moving or splitting volumes to move data to other pools.

For example, suppose you have an 800 MB storage pool with eight volumes set at 100 MB each. The administrative size equals the physical limits. To overbook the pool, you can add volumes, set one or more of the volumes to expand to the pool size, or increase the size of existing volumes, with the understanding that these are administrative maximum sizes, not physical sizes.

NSS recognizes DOS, Macintosh, UNIX, and Long name spaces. Volume names, directory names, and file names in NSS are case insensitive. This differs from Linux traditional file systems, which are case sensitive. For information, see [Section 13.16, “Mounting an NSS Volume on Linux to Support Case Insensitive File Names \(Linux\),” on page 197.](#)

After you create NSS pools, use the iManager Storage plug-in to configure NSS storage volumes. For information about iManager, see the [Novell iManager 2.5 Administration Guide \(http://www.novell.com/documentation/imanager25/imanager\\_admin\\_25/data/hk42s9ot.html\)](http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html).

You can also use the console-based NSS Management Utility to configure and manage NSS volumes. For information, see [Section 4.2, “NSS Management Utility,” on page 52.](#)

### 13.1.1 Moving NSS Volumes Across Platforms

For information, see the following in [“Coexistence and Migration Issues” on page 73:](#)

- [Section 6.6, “Compatibility Issues for Using NSS Cross-Platform,” on page 82](#)
- [Section 6.7, “Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux,” on page 87](#)
- [Section 6.8, “Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later,” on page 92](#)
- [Section 6.9, “Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare,” on page 95](#)
- [Section 6.10, “Moving Clustered Devices with NSS Volumes Cross-Platform,” on page 98](#)
- [Section 6.11, “Upgrading NetWare 5.x NSS Volumes and NetWare Traditional Volumes to NSS Volumes,” on page 98](#)

In a clustered storage area network with Novell Cluster Services™, NSS volumes can fail over between kernels, allowing for full data and file system feature preservation when migrating data to Linux. However, you cannot SAN boot cross-platform. For information, see the following:

- [OES Novell Cluster Services Administration Guide for Linux \(http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html\)](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html)
- [OES Novell Cluster Services Administration Guide for NetWare \(http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html\)](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html)

## 13.2 Understanding Volume Attributes

The following volume attributes determine the services and characteristics of an NSS volume:

- [Section 13.2.1, “Encryption,” on page 183](#)
- [Section 13.2.2, “Data Shredding,” on page 183](#)
- [Section 13.2.3, “Compression,” on page 183](#)
- [Section 13.2.4, “Data Migration,” on page 184](#)
- [Section 13.2.5, “Directory Quotas,” on page 184](#)
- [Section 13.2.6, “User Space Quotas,” on page 184](#)
- [Section 13.2.7, “Salvage Files,” on page 184](#)
- [Section 13.2.8, “Lookup Name Space,” on page 184](#)
- [Section 13.2.9, “Read Ahead Blocks in NSS Volumes,” on page 184](#)
- [Section 13.2.10, “Flush Files Immediately,” on page 185](#)
- [Section 13.2.11, “User-Level Transaction Model \(NetWare\),” on page 185](#)
- [Section 13.2.12, “File-Level Snapshot \(NetWare\),” on page 185](#)
- [Section 13.2.13, “Modified File List,” on page 186](#)
- [Section 13.2.14, “Backup,” on page 186](#)

### 13.2.1 Encryption

Encrypted volume support is provided for NSS volumes in NetWare 6.5 Support Pack 2 and later. It is available on OES SP1 Linux and later.

The Encryption attribute enables password-protected activation of encrypted NSS volumes. Encryption can be activated at creation time only, and this choice persists for the life of the volume. The `sys :` volume cannot be encrypted.

This attribute is currently not available if you create an NSS volume with iManager; use NSSMU to create an encrypted volume. For information, see [“Managing Encrypted NSS Volumes” on page 201](#).

### 13.2.2 Data Shredding

The Data Shredding attribute allows you to electronically overwrite deleted and purged data areas to prevent unauthorized users from using a disk editor to access purged files. You can specify the number of times (1 to 7) to shred data. For information, see [Section 15.2, “Using Data Shredding to Prevent Access to Deleted Files,” on page 215](#).

### 13.2.3 Compression

The Compression attribute activates file compression in NSS volumes. Compression can be activated at creation time only and this choice persists for the life of the volume. Data in the volume might be stored normally or in compressed form, depending on how frequently it is used. Compression parameters can be set at the server level to control compression behavior. For information, see [“Managing Compression on NSS Volumes” on page 227](#).

## 13.2.4 Data Migration

The Data Migration attribute sets a flag that indicates to third-party software that this volume's data can be migrated to near-line or offline storage media after it is inactive for specified lengths of time. This attribute requires third-party software to take advantage of the capability.

## 13.2.5 Directory Quotas

The Directory Quotas attribute enables you to assign a maximum quota of space that a directory can consume. For information, see [“Conserving Disk Space with Volume, Directory, or User Space Quotas” on page 243](#).

## 13.2.6 User Space Quotas

The User Space Quotas (user space restrictions) attribute enables you assign a maximum quota of space that a user's data can consume across all directories in the volume.

For information, see [“Conserving Disk Space with Volume, Directory, or User Space Quotas” on page 243](#).

## 13.2.7 Salvage Files

The Salvage Files attribute enables deleted files to remain on the volume until the Purge Delay time expires or until space is needed on the volume for other data. Until the Purge Delay time expires, the Salvage feature tracks the deleted files and allows the deleted files to be salvaged and restored. If space is needed, the oldest deleted files are purged to clear space.

If the Salvage Files attribute is disabled, deleted files are purged immediately on deletion.

For information, see [Section 19.1, “Configuring the NSS Volume Salvage Area,” on page 255](#) and [Section 19.5, “Salvaging or Purging Deleted Files from the File System Salvage Area,” on page 259](#).

## 13.2.8 Lookup Name Space

The Lookup Name Space attribute sets the name space to use when you mount the volume. The name spaces are UNIX, Long, DOS, or Mac. For NSS volumes, the Long name space is highly recommended.

If you create an NSS volume on Linux, the name space defaults to UNIX, but you should consider choosing the Long names pace to improve performance, especially if you expect to store millions of files on the volume.

## 13.2.9 Read Ahead Blocks in NSS Volumes

The Read Ahead Blocks attribute specifies the number of data blocks that NSS reads ahead for any open file on which read operations are ongoing in the specified server. The Read Ahead Blocks in Volumes attribute is not yet available in the Storage plug-in for iManager. It is enabled by default and set at 2 blocks. To modify the value, you must set it from the command line using the ReadAheadBlocks option.



The default block count of 2 blocks (8 kilobytes) satisfies the performance requirements for a typical volume. The valid range for block count is 0 blocks to 1024 blocks, where a block count of zero (0) implies no read ahead. However, the practical maximum is 128 blocks.

The most efficient value for block count depends on your hardware. In general, we recommend a block count of 8 to 16 blocks for large data reads; 2 blocks for CDs, 8 blocks for DVDs, and 2 blocks for ZLSS.

The block count should not exceed 128 blocks for most implementations. Aggressive read ahead is optimal for sequential access but it degrades random access. Block counts greater than 128 can starve other system components for memory or buffers, which can result in poor performance or can cause system hangs. As the number of concurrent connections to unique files increases, you should reduce the read-ahead block count.

To modify this parameter: At the server console, enter

```
nss /ReadAheadBlks=VolName:Count
```

Replace *VolName* with the volume name, such as *data*. Replace *Count* with the number of blocks to read ahead. The valid range is 0 (no read ahead) to 1024 blocks. Your setting is persistent; the value you set for a volume persists after reboot

### 13.2.10 Flush Files Immediately

The Flush Files Immediately attribute enables NSS to immediately write to disk all data in cache that is pending writes to the file when you close the file. Otherwise, the data in cache must wait until the next write cycle to be written to the disk, putting the information at risk for loss during the interim, for example, if the server failed. For information, see [Section 16.1, “Using the Flush Files Immediately Attribute to Protect Against Possible Data Corruption on System Failure,” on page 217](#).

### 13.2.11 User-Level Transaction Model (NetWare)

The User-Level Transaction Model enables the Transaction Tracking System™ (TTS™) function for this volume. TTS protects database applications by backing out transactions that are incomplete because of a system failure. For information, see [Section 16.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback \(NetWare\),” on page 219](#).

### 13.2.12 File-Level Snapshot (NetWare)

The File-level Snapshot attribute enables a backup utility to capture the last closed version of a file that is open at the time a backup is in progress. You must manually deactivate the volume, then activate the volume after setting this attribute to let the volume set up the virtual volume for the metadata about file snapshots.

If the File Snapshot attribute is enabled, Novell Storage Management Services™ (SMS) saves the snapshot version of the file to backup media if a file is in use when the backup occurs. Some third-party backup applications can also take advantage of this attribute.

For more information, see [Section 16.3, “Using the File-Level Snapshot Attribute to Enable the Backup of Open Files,” on page 220](#).

### 13.2.13 Modified File List

The Modified File List (MFL) attribute enables NSS to create a list of all files modified since the previous backup. The log is available only through third-party software.

### 13.2.14 Backup

The Backup attribute sets a flag to indicate to the backup software that the volume contains data you want to back up. Disable this flag if the volume is empty or if backing up the data is unnecessary. This backup flag is independent of the third-party backup system you use; your backup system might not recognize this option, even if you select it.

## 13.3 Creating and Configuring Unencrypted NSS Volumes with iManager

This section describes how to create an unencrypted volume. For information on creating encrypted volumes, see [“Managing Encrypted NSS Volumes” on page 201](#).

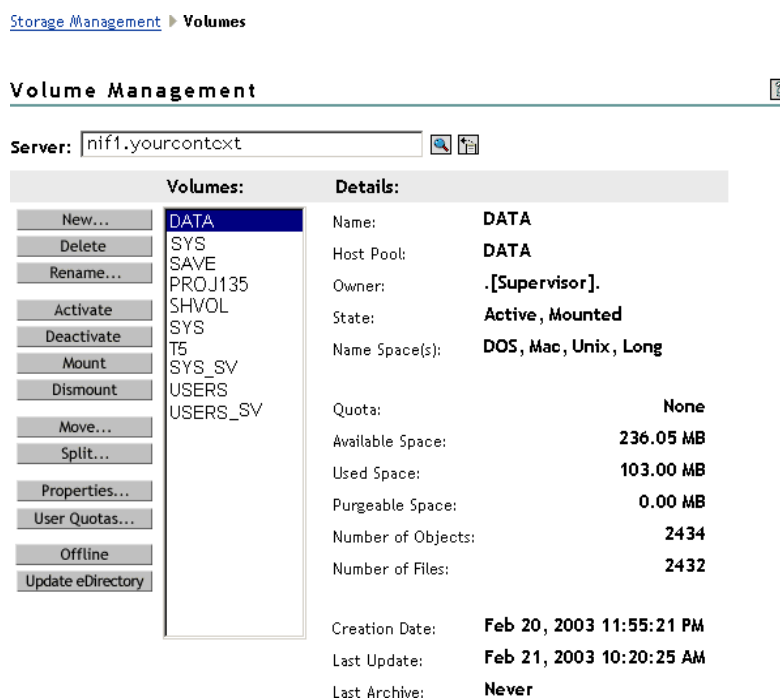
- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list, as illustrated in the following figure.



- 3 To create a new volume, click *New*.

This opens the New Volume wizard to guide you through the process.

4 Specify a name for the new volume.

If the name you provide is not unique, you receive an error message. For guidelines about naming volumes, see [Section 5.5, “Naming NSS Storage Objects,”](#) on page 69.

5 Specify the pool on the new location where you want the new volume to reside, then click *Next*.

Only pools that have free space appear in the list. If no pools are listed, there is no space available to create a volume in the New Location. Cancel the wizard, add more devices to the server or free up space on existing pools, then return to the Volume Management page to create the new volume.

6 Specify the size of the volume:

- Specify a quota size for the volume if you want to limit the size of the volume.
- Select *Allow Volume Quota to Grow to the Pool Size* if you want the volume to expand to the size of the pool.

Pools can be overbooked; each volume can potentially grow to the size of the pool. NSS allocates space as it is needed.

7 On the Volume Attributes page, set the attributes for the new volume you are creating.

For information, see [Section 13.2, “Understanding Volume Attributes,”](#) on page 183.

**New Volume** [?]

**Attribute information**

New Volume: **NIFVOL**

<input checked="" type="checkbox"/> Backup	<input type="checkbox"/> Migration
<input type="checkbox"/> Compression	<input type="checkbox"/> Modified File List(MFL)
<input type="checkbox"/> Data Shredding	<input checked="" type="checkbox"/> Salvage Files
Number of shredding cycles <input type="text"/>	<input type="checkbox"/> Snapshot - File Level
<input type="checkbox"/> Directory Quotas	<input type="checkbox"/> User Space Quotas
<input type="checkbox"/> Flush Files Immediately	<input type="checkbox"/> User-level Transaction Model

**On Creation**

☒ Activate Volume

☒ Mount Volume

<< Back   Finish   Cancel

8 In the *On Creation* section, set the following preferences:

- *Activate*. Activates logical volumes as soon as you create them.
- *Mount*. Mounts logical volumes as soon as you create them.

9 If this is an NSS volume on Linux, specify the mount point, such as `/media/nss/NIFVOL`.

10 Click *Finish*.

11 (Linux) If you enabled or disabled the Directory Quotas attribute, restart NCP2NSS by entering at a terminal prompt:

```
/etc/init.d/ncp2nss restart
```

## 13.4 Configuring Encrypted NSS Volumes with NSSMU

NSS Encrypted Volume Support allows you to create encrypted NSS volumes using NSSMU. You cannot encrypt the `sys :` volume. For information, see “[Managing Encrypted NSS Volumes](#)” on [page 201](#).

## 13.5 Updating eDirectory Volume Objects

In Novell eDirectory™, each NSS volume is represented by a Volume object. Volume objects are leaf objects that represent a physical or logical volume on the network.

The Volume object’s properties contains the following information:

- The NetWare server where the volume resides
- The volume name recorded when the volume was initialized on the server (for example, `sys :`)
- The volume’s owner (login username of the administrator who created it)
- Space use restrictions for users
- A description of the volume’s use
- Statistical information on disk space availability, block size, directory entries, name space support, and so on.

Usually, NSS creates the Novell eDirectory Volume object when you create the volume, and it updates the properties of the volume as needed. The *Update eDirectory* option on the Volume Management page allows you to add, keep, or replace a Volume object for a selected volume at the same context level as the server.

When you select Update eDirectory, NSS searches for the object.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 Select the volume you want to update.

Wait for the page to refresh and display the details.

- 4 Click *Update eDirectory*.

Wait while NSS searches for the Volume object in the server context.

- 5 Do one of the following:

- If the Volume object does not exist, NSS adds the Volume object to the context level. Confirm the addition.
- If the Volume object exists, NSS prompts you with two options: Delete and Replace the existing object or Retain the existing object. Select one option and confirm your choice.

## 13.6 Viewing the Details of an NSS Volume

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 In the Volumes list, select a volume that you want manage.

When the page refreshes, the details for the volume appear in the *Details* area. The volume must be mounted and active for the details to be available.

- 4 (Conditional) To activate the volume, select the volume, then click *Activate*.

## 13.7 Viewing Properties of an NSS Volume

After you set up and configure NSS volumes, you can view the properties, such as attribute settings, volume statistics, and volume usage.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 In the Volumes list, select a volume that you want manage.

- 4 Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. It opens by default to the Attributes tab.

[Storage Management](#) > [Volumes](#) > [Properties](#)

---

Volume Properties ?

---

Properties: USER\_VOL

**Attributes**   Statistics   Quota Usage

<input checked="" type="checkbox"/> Backup	<input type="checkbox"/> Migration
<input type="checkbox"/> Compression	<input type="checkbox"/> Modified File List(MFL)
<input type="checkbox"/> Data Shredding	<input checked="" type="checkbox"/> Salvage
Number of shredding cycles: <input type="text"/>	<input type="checkbox"/> Snapshot - File Level
<input checked="" type="checkbox"/> Directory Quotas	<input checked="" type="checkbox"/> User Space Quotas
<input type="checkbox"/> Flush Files Immediately	<input type="checkbox"/> User-level Transaction Model

Quota:  Bytes ↓

☒ Allow volume quota to grow to the pool size

Mount Point:

---

Use the Attributes page to view the volume's attribute configuration, the volume quota, and the volume mount point (Linux). For information about modifying attributes, see [Section 13.8, "Modifying Attributes of an NSS Volume,"](#) on page 191.

- 5 Click the *Statistics* tab to view the current space usage statistics for the selected volume.

[Storage Management](#) ▶ [Volumes](#) ▶ **Volume Properties**

## Volume Properties ?

Properties: NIFVOL

Attributes		Statistics		Quotas Usage	
Compression		Salvage			
Compressed Space:		0.00 Bytes		Minimum Keep Seconds:	
				Maximum Keep Seconds: 0	
Files:					
Not Deleted:		0		Low Water Mark: 10	
Deleted:		0		High Water Mark: 20	
Uncompressed:		0		Next Scheduled Purge	
				Purgeable Space: 0.00 MB	
				Unpurgeable Space: 0.00 MB	
				Deleted Files: 0	
				Oldest Deleted Time: 0	
GUID:		A7942770-9C3E-01D7-80-00-1A1A3BF816F4			
Block Size:		4.00 KB			
Close					

- Click the *Quotas Usage* tab to view the volume and pool space usage for the selected volume.

[Storage Management](#) ► [Volumes](#) ► **Volume Properties**

## Volume Properties




Properties: NIFVOL

Attributes Statistics **Quotas Usage**

Volume Usage: NIFVOL

Pool Usage: NIFPOOL



Quota:	none	Available Space:	89.93 MB
Used Space:	576.00 KB	Free:	89.93 MB
Compressed:	0.00 Bytes	Purgeable:	0.00 MB
Other in-use space:	576.00 KB	Total Space:	99.00 MB
Purgeable Space:	0.00 MB	Used:	9.07 MB
Available Space:	89.93 MB	Booked:	0.00 MB

This pool contains logical volumes with no quotas. They were not used to calculate the booking.

Close

## 13.8 Modifying Attributes of an NSS Volume

After you set up and configure NSS volumes, you can modify most of the attribute settings. Encrypted Volume Support attribute and Compression attribute can be set only at the time the volume is created. If you try to modify those settings, iManager or NSSMU returns an error message.

You can specify a Volume Quota. For NSS on Linux, you can modify the mount point.

- In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- In the Volumes list, select a volume that you want manage.

- Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. Use the Attributes page to view or modify the attributes for the selected volume.

- Do one or more of the following:

- Select or deselect a modifiable attribute, then click *Apply*.

The Encryption and Compression attributes can be set only at the time the volume is created. If you try to modify those settings, iManager returns an error message.

For information about attributes, see [Section 13.2, “Understanding Volume Attributes,” on page 183](#).

- Specify a volume quota, then click *Apply*.
- (Linux) Specify a new *Mount Point* for your volume, then click *Apply*. For example:

```
/media/nss/VOL1
```

- 6** (Linux) If you enabled or disabled the *Directory Quotas* attribute, restart NCP2NSS by entering at a terminal prompt:

```
/etc/init.d/ncp2nss restart
```

For information about setting quotas after you have enabled Directory Quotas or User Space Quotas, see [“Conserving Disk Space with Volume, Directory, or User Space Quotas” on page 243](#).

## 13.9 Renaming an NSS 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** In iManager, *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2** Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3** In the Volumes list, select a volume.

- 4** Click *Rename*.

The Rename Volume wizard opens.

- 5** Specify the new name of the volume.

- 6** Click *Finish*.

After the page refreshes, the volume appears in the Volumes list with its new name.

## 13.10 Activating and Deactivating an NSS Volume

After you set up and configure NSS volumes, you can activate and deactivate volumes to make them available to users and applications. To view details of a volume, it must be active.

- 1** In iManager, *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2** Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.



- 3 In the Volumes list, select one or more volumes that you want to make active or deactivate.
- 4 Depending on the state of the volume, click *Activate* or *Deactivate*.

After the page refreshes, you can see that each volume's state matches the state you specified.

## 13.11 Activating and Deactivating CDs and DVDs as NSS Volumes (NetWare)

You can activate CDs and DVDs as NSS volumes with the `cddvd.nss` module. The `cddvd.nss` module replaces `cdrom.nlm` and `cd9660.nss` for CD support, `udf.nss` for DVD support, and `cdhfs.nss` for Macintosh CD support. When you load `cddvd.nss`, it uses the same set of switches (`UpdateCDToNDS`, `LoginUpdateToNDS`, `LeaveCDInNDS`, and `LockCDMediaIn`) that were used in the previous commands.

`Cddvd.nss` supports the following disk formats:

- Universal Disc Format (UDF) for DVDs
- ISO 9660 for CDs
- Hierarchical File System (HFS) for Macintosh CDs

By default, `cddvd.nss` loads the formats in sequential order, with the UDF format for DVDs first, ISO 9660 for CDs next, and CDHFS for Macintosh CDs last.

If a CD or DVD is in a drive when `cddvd.nss` loads, the disk is immediately activated as an NSS volume. Whenever the `cddvd.nss` module recognizes a newly inserted CD or DVD, it immediately activates the disk as a read-only NSS volume.

The volume name is the CD or DVD volume label, if one exists, and is acceptable to NetWare. Otherwise, `cddvd.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.

Although the NSS volume is activated, it is not automatically added to Novell eDirectory. You can add it to eDirectory later by enabling a `cddvd.nss` switch, such as

```
nss /UpdateCDtoNDS=cname
```

### Loading and Unloading the `Cddvd.nss` Module

The `cddvd.nss` module loads automatically when NSS starts. To load `cddvd.nss` manually, enter the following at the server console:

```
load cddvd.nss
```

To unload `cddvd.nss`, enter the following at the server console:

```
unload cddvd.nss
```

### Load Order of Formats for Dual-Formatted Drives

Disks in dual-formatted CD/DVD drives are activated as one volume and the format is decided by `cddvd.nss`. If you have an ISO/UDF bridge disk, which supports both ISO 9660 and UDF formats, you can instruct NSS to load ISO 9660 first by entering the following command at the server console:

```
nss /LoadISOVolumeFirst
```

## 13.12 Activating and Deactivating CD and DVD Image Files as NSS Volumes (NetWare)

In NetWare 6.5 Support Pack 2 and later, the `cddvd.nss` module allows you to activate a CD or DVD image file as a read-only NSS volume, if the source file resides on an existing NSS volume. You add the image file as a volume with the Mount Image Volume switch. This volume functions the same as any other NSS volume, except that it is read-only. You remove it as a volume with the Remove Image Volume switch.

While the CD or DVD image file exists on the server as an NSS volume, the following actions cause the image file to be removed as an NSS volume:

- Any file operations the image source file, such as open, read, or write
- Deactivation of the NSS volume where the image source file resides

To activate a CD or DVD image file as an NSS volume, enter at the system prompt:

```
nss /MountImageVolume=path_imagefilename
```

Replace *path\_imagefilename* with the path on the NSS volume where the image file resides, including its filename. For example:

```
nss /MountImageVolume=nssvol:\suse9_1.iso
```

You can mount and dismount the volume as with normal volumes, until you remove it by entering the following:

```
nss /RemoveImageVolume
```

NSS Command Options	Description
<code>/MountImageVolume=<i>path_imagefilename</i></code>	<p>Add and mount the specified CD or DVD image file as a read-only NSS volume.</p> <p>Replace <i>path_imagefilename</i> with the path on the NSS volume where the image file resides, including its filename. For example:</p> <pre>nss /MountImageVolume=   <i>nssvol:\suse9_1.iso</i></pre>
<code>/RemoveImageVolume=<i>path_imagefilename</i></code>	<p>Dismount and remove the specified CD or DVD image file as a read-only NSS volume.</p> <p>Replace <i>path_imagefilename</i> with the path on the NSS volume where the image file resides, including its filename. For example:</p> <pre>nss /RemoveImageVolume=   <i>nssvol:\suse9_1.iso</i></pre>

## 13.13 Activating and Deactivating DOS Partitions as NSS Volumes (NetWare)

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:** DOS filenames for file systems such as OS/2\*. This uses the Windows 95/98/ME 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.

## 13.14 Mounting and Dismounting an NSS Volume (Linux)

After you set up and configure NSS volumes, you can mount and dismount volumes to make them available to users and APIs. After you mount a volume, it is only available to APIs until you activate it. Dismounting a volume makes it unavailable to users and to APIs.

### Dismounting an NSS Volume on Linux from the NCP Server

Before you can dismount an NSS volume on Linux, you must dismount the volume from NCP™ Server; otherwise, the dismount function fails.

- 1 At the server prompt, open the NCP Console by entering

```
ncpcon
```

- 2 Dismount the volume from NCP.

The volume is no longer accessible or visible to NCP clients.

### Mounting or Dismounting an NSS Volume with iManager (NetWare)

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 In the Volumes list, select one or more volumes that you want to mount or dismount.
- 4 Depending on the state of the volume, click *Mount* or *Dismount*.

After the page refreshes, you can see that the volume's state changed. The details of dismounted volumes are not available on the iManager Volume Management page.

### Mounting an Encrypted NSS Volume with NSSMU (Linux)

For NSS volumes on Linux, you can mount the volume only from NSSMU the first time it is started after a reboot. NSSMU allows you to mount only one volume at a time so that you can enter its password.

- 1 In NSSMU, select *Volumes*.
- 2 In the Volumes list, select the encrypted volume that you want to mount.
- 3 Select or press *F7* to mount the volume.
- 4 If you are prompted to enter the password, enter the password, then click *OK*.  
You are prompted for the password on the first time the volume is mounted after a system reboot.

### Dismounting an Encrypted NSS Volume with NSSMU (Linux)

For NSS volumes on Linux, you can mount the volume only from NSSMU the first time it is started after a reboot.

- 1 In NSSMU, select *Volumes*.
- 2 In the Volumes list, select the encrypted volume that you want to mount.
- 3 Select or press *F7* to dismount the mounted volume.

## 13.15 Using NFS Mount for NSS Volumes (Linux)

NSS volumes and their directories are NFS exportable and accessible from remote systems.

In YaST, click *Network Services > NFS Server > NFS Mount*, specify the volume, and then enter the following options:

```
rw,no_root_squash,sync,fsid=value
```

### RW Attribute

Mount the NSS file system with Read/Write (*rw*) access.

### No Root Squash

Disable root squashing for the superuser with the No Root Squash (*no\_root\_squash*) option. This allows root users on client computers to have root access on the server. With the No Root Squash option, mount requests for root are not mounted to the anonymous user (*nobody*). This option is needed for diskless clients.

NSS volumes are logical volumes. They are not directly mounted on devices, but are associated with pools, which are mounted on devices. Because NSS volumes do not have a device directly associated with them, NFS treats the volume like a diskless client, which makes the *no\_root\_squash* option necessary when you mount NSS volumes.

## Sync

Specify the Sync (`sync`) option, which requires all file system writes to be committed to disk before the request can be completed.

## FSID

You must export the NSS volume or directory with the FSID option:

```
fsid=n
```

Replace *n* with an integer value greater than or equal to 0. The numbers do not need to be sequential. For example, `fsid=1` and `fsid=10`.

---

**IMPORTANT:** Make sure to use a unique `fsid` number for each NSS volume or directory you are exporting.

---

Exporting with the `fsid` option works around the fact that there is no device associated with a logical volume.

## 13.16 Mounting an NSS Volume on Linux to Support Case Insensitive File Names (Linux)

When mounting an NSS volume on Linux, specify the Long name spaces to make its directory names and filenames case insensitive. If your volume contains large directories with millions of files, using the default UNIX name space can cause volumes to mount very slowly. Use the Long name space to allow the volume to mount normally.

### Mounting an NSS Volume

To mount an NSS volume from a terminal command line, enter

```
mount -t nssvol volname mount_point -o name=volname,ns=long
```

For the `-t` option, *nssvol* is the file system type for NSS volumes.

*Volname* is the name of the NSS volume.

The *mount\_point* is the full path with the volume name where you want to mount the volume, such as `/media/nss/VOL1`. The default mount location for NSS volumes is in the `/media/nss/` directory.

For the `-o` option, specify the volume name and the name space type. Valid name spaces are `dos`, `mac`, `long`, or `unix`.

For example, to mount an NSS volume named `VOL1` as case insensitive, enter the following at a terminal prompt:

```
mount -t nssvol VOL1 /media/nss/VOL1 -o name=VOL1,ns=long
```

### Mounting an NSS Volume Automatically on System Reboot

You can automatically mount the NSS volume on system reboot by adding a line to the `/etc/fstab` file in the following general format:

```
label mount_point fstype mount_options dump_frequency fsck_order
```

For example:

```
VOL1 /media/nss/VOL1 nssvol noauto,rw,name=VOL1,ns=long 0 0
```

## Using Samba with NSS Volumes

When using Samba, make sure to do the following:

- Mount the NSS volume as case insensitive by using the Long name space.
- Specify `Case Sensitive=No` when exporting Samba shares for NSS volumes with case insensitive name spaces.

Edit the `/etc/samba/smb.conf` file to set the `Case Sensitive` parameter to `No`.

This improves performance for your NSS volumes on Linux, especially those with larger directories.

## 13.17 Mounting a Disk in a USB Floppy Drive as a Local FAT Volume (NetWare)

If your server uses a USB floppy drive, NetWare 6.5 does not automatically mount floppies during the server installation. You might need to manually mount a floppy disk to install the NetWare 6.5 license.

Local File Allocation Table (FAT) volumes are available to the server but are not seen by users. By default, NetWare 6.5 does not automatically mount local FAT volumes such as the floppy disk in a USB floppy drive. Use the Local Fat Volume Mount (LFVMount) command to detect and mount any local FAT volumes not recognized by DOS and assign them a local drive letter.

To mount local FAT volumes, enter the following at the system console prompt:

```
lfvmount
```

The LFVMount command affects only the current local FAT volumes. For example, you must re-enter this command each time you change a floppy disk in the USB floppy drive.

Use the associated SET parameter (`set auto lfvmount=<on | off (default)>`) to configure automatic detection. If it is set to `On`, NetWare automatically detects and mounts local FAT volumes as you insert the media into the drive and then assigns drive letters to them. If it is set to `Off` (default), NetWare does not automatically detect and mount local FAT volumes; you must use the LFVMount command to manually mount them as needed.

To automatically mount local FAT volumes for a server, enter the following at the system console prompt:

```
set auto lfvmount=on
```

To prevent the automatic mounting of local FAT volumes for a server, enter the following at the system console prompt:

```
set auto lfvmount=off
```

If entered at the system console, the associated SET parameter remains in effect until you change it or until you reboot the server.

To make the `SET` parameter persistent through a server reboot, add one of the following commands in the `startup.ncf` file:

- `set auto lfvmount=on`
- `set auto lfvmount=off`

If no value is set in the `startup.ncf` file, NetWare does not automatically mount the local FAT volumes. You must use the `LFVMount` command to manually mount local FAT volumes as needed.

The Local Fat Volumes (`lfvolumes`) command displays a list of local FAT volumes that are currently mounted and provides detailed information about them.

To list information about local FAT volumes mounted on the server, enter the following at the system console prompt:

```
lfvolumes
```

The `LFVMount` command appears in the console-based help for console commands.

For help with the `LFVMount` command, enter one of the following at the system console prompt:

- `help mount`
- `help lfvmount`

## 13.18 Deleting an NSS Volume

Deleting a volume removes the data in the volume and frees the space to be used by other volumes in the same pool. When you delete a volume, it remains for salvageable until one of the following events occurs:

- Volume Purge Delay times out. The deleted volume is purged automatically. For information, see [Section 19.1, “Configuring the NSS Volume Salvage Area,” on page 255](#).
- The available space on the server falls below the Low Watermark. Deleted volumes are purged until enough space is cleared so that the available space increases above the High Watermark. For information, see [Section 19.3, “Setting the Low and High Watermarks for Purging the File System Salvage,” on page 257](#).
- You manually purge the deleted volume. For information, see [Section 19.2, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 255](#).

During that time, the space belonging to the deleted volume is not available to other volumes and the volume is salvageable. Once the purging process begins, the volume is no longer salvageable.

If it is necessary, you can restore a deleted volume before it is purged. See [Section 19.2, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 255](#).

You cannot delete a `sys :` volume using iManager. Deleting a `sys :` volume removes the operating system. You need to reinstall NetWare to delete a `sys :` volume. However, if the `sys :` volume you want to delete does not contain the active operating system (for example, a remnant on a previously used disk), you can rename the volume, then delete it.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 Select one or more volumes that you want to delete.

Wait for the page to refresh and make the *Delete* option available for the selected volume.

- 4 Click *Delete*.

- 5 Click *Yes* to confirm the deletion, or click *No* to cancel the deletion.

## 13.19 Verifying and Rebuilding NSS 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 do not have access to any of the volumes in that pool. You cannot rebuild or verify an NSS volume independently of other volumes in the same pool.

---

**IMPORTANT:** A pool rebuild does not restore the data itself. Instead, it restores consistency to the data structure itself after data corruption. Data is likely to be lost during a rebuild so rebuild pools only as a last resort.

---

For more information, see [Section 11.11, “Verifying and Rebuilding NSS Pools and Volumes,” on page 157](#)

## 13.20 What’s Next

For information about advanced volume features, see the following:

- [“Managing Encrypted NSS Volumes” on page 201](#)
- [“Securing Access to NSS Volumes, Directories, and Files” on page 211](#)
- [“Using Volume Attributes that Help Ensure Data Integrity” on page 217](#)
- [“Managing Compression on NSS Volumes” on page 227](#)
- [“Conserving Disk Space with Volume, Directory, or User Space Quotas” on page 243](#)
- [“Salvaging and Purging Deleted Volumes, Directories, and Files” on page 255](#)
- [“Using Distributed File Services to Move and Split Volumes \(NetWare\)” on page 261](#)



# Managing Encrypted NSS Volumes

# 14

Novell® Storage Services™ provides optional Encrypted Volume Support (EVS) for NSS volumes on NetWare 6.5 SP2 and later, OES NetWare and later, and OES SP1 Linux and later operating systems.

This section discusses the following:

- [Section 14.1, “Understanding Encrypted Volume Support,” on page 201](#)
- [Section 14.2, “Security Considerations for Encrypted Volumes,” on page 203](#)
- [Section 14.3, “Creating an Encrypted Volume,” on page 204](#)
- [Section 14.4, “Mounting an Encrypted NSS Volume with NSSMU \(Linux\),” on page 205](#)
- [Section 14.5, “Activating an Encrypted NSS Volume with NSSMU \(NetWare\),” on page 205](#)
- [Section 14.6, “Mounting Encrypted NSS Volumes with NSS Commands \(Linux\),” on page 206](#)
- [Section 14.7, “Activating Encrypted NSS Volumes with NSS Commands \(NetWare\),” on page 206](#)
- [Section 14.8, “Dismounting an Encrypted NSS Volume with NSSMU \(Linux\),” on page 207](#)
- [Section 14.9, “Deactivating an Encrypted NSS Volume with NSSMU \(NetWare\),” on page 208](#)
- [Section 14.10, “Using Encrypted Volumes in a Server Cluster \(Linux\),” on page 208](#)
- [Section 14.11, “Using Encrypted Volumes in a Server Cluster \(NetWare\),” on page 209](#)
- [Section 14.12, “Removing Encrypted Volumes,” on page 209](#)
- [Section 14.13, “What’s Next,” on page 210](#)

## 14.1 Understanding Encrypted Volume Support

NSS Encrypted Volume Support meets the legal standard of making data inaccessible to software that circumvents normal access control, such as if the media were stolen. EVS is available only for newly created NSS volumes. EVS stores user data in encrypted format on the NSS volume, yet works transparently with most applications, NLM™ programs, and backup utilities that currently work with NSS.

Any NSS volume, except the `sys:` volume, can be designated at volume creation time to be an encrypted volume. The Encrypted attribute stays with the volume throughout its life. An encrypted volume cannot later be converted to be unencrypted, nor can an unencrypted volume later be converted to be encrypted. This is a creation-time-only decision.

Encryption is transparent above the physical read/write layer of an NSS volume. It requires no changes for applications. All the rules of file system trustee assignments, trustee rights, ownership, sharing, visibility, locking, transactions, and space restrictions remain the same. Performance for an encrypted volume is slightly degraded compared to an unencrypted volume under the same conditions.

### 14.1.1 Encryption Method

Encrypted volume support uses the NICI libraries for all cryptographic support. NICI generates a 128-bit AES key for encryption that persists for the life of the volume. You cannot change the password because it is the key used to encrypt data. NICI uses the password to wrap the key and other volume-specific cryptographic information into a 128-bit package that is persistently stored in two locations on the NSS media: the Volume Data Block and the Volume Locator Block. After the cryptographic data is wrapped for the activated volume, EVS eliminates the password from memory.

### 14.1.2 Encryption Password

The encryption password can be 2 to 16 standard ASCII characters, with a suggested minimum of 6. The password generates a 128-bit NICI key for encryption. The password is set when you create the volume. It persists for the life of the volume; it cannot be changed later.

### 14.1.3 How Encrypted Volume Support Works

On the first activation after a system reboot, you must enter a valid password. When the volume is activated, NSS loads the volume's persistent data from the Volume Data Block. If the Encrypted attribute is enabled for a volume, NSS searches in memory for a known key in the list of volume names and keys. If the key is present, it is used. If no key is present, NSS consults the list of volumes and passwords. If a password is available, it is used to unwrap the key from the persistent data and the new key is placed in the list of volumes and keys. The password is eliminated from memory.

After the encrypted volume is activated, all encryption operations on user data are transparent to file system applications that use normal file I/O functions. Data written to files is held in cache until such time as it would be normally written. At physical write time, the data is encrypted to a temporary write buffer and written to the volume in encrypted format.

During reads, the cache is consulted, as it would normally be, to determine if a requested block is already in memory. If the requested data block is in cache, the clear-text data is transferred. If it is not, a physical read request is made, with the read directed to a temporary buffer. After read completion, but before control is returned to the calling program, the encrypted data in the temporary buffer is decrypted into a cache buffer. The read proceeds normally, with clear-text data being made available to all future requestors.

### 14.1.4 Guidelines for Using Encrypted Volumes

- For NSS volumes on NetWare, the `sys:` volume cannot be encrypted. If your IT environment requires that you encrypt data, make sure to store data on its own encrypted volume.
- We recommend that you avoid mixing encryption and compression features in a volume. Use one or the other, but not both.
- You can enable the Encryption attribute only at volume creation time.
- If it is enabled, the Encrypted volume attribute persists for the life of the volume.
- To encrypt an existing volume, you must create a new encrypted volume, then migrate existing data from the unencrypted volume to the encrypted volume.
- The encryption password is 6 to 16 standard ASCII characters.

## 14.2 Security Considerations for Encrypted Volumes

- [Section 14.2.1, “Choosing a Strong Encryption Password,” on page 203](#)
- [Section 14.2.2, “Backing Up Data from an Encrypted Volume,” on page 203](#)
- [Section 14.2.3, “Excluding the NSS Cache Memory from Core Dumps,” on page 203](#)
- [Section 14.2.4, “Disabling Logs \(Linux\),” on page 203](#)
- [Section 14.2.5, “Archiving Files from an Encrypted Volume,” on page 204](#)
- [Section 14.2.6, “Using Direct I/O to an Encrypted Volume,” on page 204](#)
- [Section 14.2.7, “Sharing Encrypted NSS Volumes in a Cluster,” on page 204](#)

### 14.2.1 Choosing a Strong Encryption Password

The encryption password is 6 to 16 standard ASCII characters. Make sure to employ security best practices for passwords. For information, see [Section 28.14, “Creating Strong Passwords,” on page 357](#).

### 14.2.2 Backing Up Data from an Encrypted Volume

Make sure to encrypt the data from an encrypted volume on backup media. Backups of an encrypted volume are not encrypted, unless it is a feature of the backup software you use. For information, see [Section 28.3, “Protecting Data During Backup and on Backup Media,” on page 351](#).

### 14.2.3 Excluding the NSS Cache Memory from Core Dumps

Make sure that you exclude the NSS cache memory from core dumps; otherwise, encrypted NSS volume data might be displayed in the clear. For information, see [Section 28.4, “Preventing Exposure of Sensitive Data in a Core Dump,” on page 351](#).

### 14.2.4 Disabling Logs (Linux)

When working with encrypted volumes on Linux, it is important to realize that the volume password and key information is exchanged between user and kernel space as encrypted volumes are created and/or mounted. If you have logging enabled on the Linux server when you enter the encryption password, your password and volume key information might show up in log file.

Even though the logging mechanisms are root user protected, we strongly recommend that you make sure logging is disabled when creating an encrypted volume or mounting the encrypted volume after a system reboot in order to protect the secrecy of your password credentials at these critical times when you are entering the encryption password.

For information, see [Section 28.5, “Preventing Exposure of the Encryption Password in a Log \(Linux\),” on page 351](#).

## 14.2.5 Archiving Files from an Encrypted Volume

If you use Novell Archive and Version Services to archive files from an encrypted volume, the destination path for Archive Manager must be on an encrypted volume. If the destination path is an unencrypted volume, the versioned data is stored in an unencrypted state.

Make sure to encrypt the archive volume when you create it on the archive server. Otherwise, the file versions are not secure.

## 14.2.6 Using Direct I/O to an Encrypted Volume

Direct I/O to an encrypted volume bypasses the EVS encryption engine and allows data to be stored in unencrypted format on the encrypted volume. This capability is useful for diagnostic, repair, or special-purpose applications, but should be avoided otherwise.

You should avoid using direct-I/O applications on encrypted volumes, especially for user data that you intend to be stored in encrypted format.

## 14.2.7 Sharing Encrypted NSS Volumes in a Cluster

When you mount the shared volume and enter the password, NSS uses the password to create a key, which it stores in the server memory. The NCS software passes the key to other nodes. After all servers hold the key, the volume is available while any one of the servers is still participating actively in the cluster. If all servers in the cluster fail, you must repeat this procedure when you recover the cluster and restart services.

## 14.3 Creating an Encrypted Volume

NSS Encrypted Volume Support allows you to create encrypted NSS volumes using NSSMU version 3.20 build 940 or later. You can create encrypted user data volumes only after the installation or upgrade process. On NetWare, you cannot encrypt the `sys :` volume.

EVS is not supported by NSS volumes on NetWare 6.5 SP1 and earlier servers or OES Linux servers. If you choose to encrypt a volume, you cannot roll back the system to these earlier versions of NetWare or OES Linux without taking steps before the rollback to preserve data. For information, see [Section 14.12, “Removing Encrypted Volumes,” on page 209](#).

---

**WARNING:** We strongly recommend that you verify that your system is working as desired before creating encrypted volumes on the system.

---

- 1 In NSSMU, select *Volume Management*, then press *Enter*.
- 2 To create a new volume, press the *Insert* key.  
A query asks if you want to encrypt the volume.
- 3 To encrypt the new volume, select *Yes*, then press *Enter*.  
NSS enables the Encrypted attribute for the volume, then prompts you to enter a password for the volume.
- 4 Enter an encryption password, then enter it again to verify it.

The encryption password can be 2 to 16 standard ASCII characters, with a suggested minimum of 6. The password generates a 128-bit NICI key for encryption. The password persists for the life of the volume; it cannot be changed later.

- 5 Set the volume size and other attributes, as desired.

When you are done, the encrypted volume is active and mounted.

You must supply the encryption password for the volume on the first volume activation (or mount on Linux) after a system boot or reboot. For information, see [Section 14.4, “Mounting an Encrypted NSS Volume with NSSMU \(Linux\),” on page 205](#).

On NetWare, you can also enter the password as needed when you activate a volume from the command line. For information, see [Section 14.6, “Mounting Encrypted NSS Volumes with NSS Commands \(Linux\),” on page 206](#).

For information about entering the password for a volume in a cluster, see [Section 14.10, “Using Encrypted Volumes in a Server Cluster \(Linux\),” on page 208](#).

## 14.4 Mounting an Encrypted NSS Volume with NSSMU (Linux)

Mount only one volume at a time so that you can enter its password.

---

**IMPORTANT:** For encrypted NSS volumes on Linux, you can mount the volume only from NSSMU the first time it is mounted after a reboot.

---

- 1 In NSSMU, select *Volumes*.
- 2 In the Volumes list, select the encrypted volume that you want to mount.
- 3 Select or press F7 to mount the volume.
- 4 If you are prompted to enter the password, enter the password, then click *OK*.

You are prompted for the password on the first time the volume is mounted after a system reboot. The password is stored on the system until the next system reboot. You can mount the volume without the password until a system reboot occurs.

## 14.5 Activating an Encrypted NSS Volume with NSSMU (NetWare)

Activate only one volume at a time so that you can enter its password.

- 1 In NSSMU, select *Volumes*.
- 2 In the Volumes list, select the encrypted volume that you want to activate.
- 3 Select or press F6 to activate the volume.
- 4 If you are prompted to enter the password, enter the password, then click *OK*.

You are prompted for the password on the first time the volume is mounted after a system reboot. The password is stored on the system until the next system reboot. You can activate the volume without the password until a system reboot occurs.

## 14.6 Mounting Encrypted NSS Volumes with NSS Commands (Linux)

You must enter a password only on the first activation following a system reboot. Thereafter, other environmental security and authentication measures control access to user data.

On Linux, the NSS Console (`nsscon`) does not support entering the password from the command line. You must mount the encrypted volume from NSSMU on the first time after a system reboot. Thereafter, you can use the commands in this section without supplying the password. For information, see [Section 14.4, “Mounting an Encrypted NSS Volume with NSSMU \(Linux\),” on page 205](#).

### Syntax

The following table provides the syntax for NSS commands to use with encrypted volumes on subsequent mounts of the volume until the system reboots. Enter the commands from `nsscon`. In each case, replace *volname* with the name of the encrypted NSS volume.

You cannot use the wildcard option of “all” as the *volname* before an encrypted volume is mounted with its password following each system reboot. The All option does not find the volume and does not execute the command.

Command	Description
<code>mount volname</code>	Mount an encrypted or unencrypted NSS volume.  The mount command is usable for encrypted volumes only after a previous activation with password. Otherwise, it returns an error message, requesting more information.
<code>mount all</code>	Activate and mount all encrypted NSS volumes that have been previously activated with their passwords.  Encrypted NSS volumes that were not previously activated are not mounted. Mount them from NSSMU where you can provide the encryption password.
<code>nss /volumes</code>	Display a list of encrypted and unencrypted NSS volumes, showing their attributes.  The encrypted volume returns a status of Encrypted.

## 14.7 Activating Encrypted NSS Volumes with NSS Commands (NetWare)

You must enter a password only on the first activation following a system reboot. Thereafter, other environmental security and authentication measures control access to user data.

## Syntax

The following table provides the syntax for NSS commands to use when activating encrypted volumes. In each case, replace *volname* with the name of the encrypted NSS volume.

You cannot use the wildcard option of “all” as the *volname* if you are activating volumes the first time after a system reboot. Until an encrypted volume is activated with its password following each system reboot, the All option does not find the volume and does not execute the command. The system returns an error message requesting more information.

You cannot use wildcard options and commands to activate encrypted volumes, such as `nss / VolumeAutoActivate`.

Command	Description
<code>nss /activate=volname</code>	Activate the encrypted NSS volume.
If prompted for a password, enter the encryption password for the encrypted volume.	The password is required only on the first activation following a system reboot.
<code>nss /volumeactivate=volname</code>	Activate the encrypted NSS volume.
If prompted for a password, enter the encryption password for the encrypted volume.	The password is required only on the first activation following a system reboot.
<code>nss /activate=all</code>	Activate all encrypted NSS volumes that have been previously activated with their passwords.
<code>nss /forceactivate=volname</code>	Forces activation for an encrypted NSS volume that has been previously activated with its password.  Encrypted NSS volumes that were not previously activated return an error message, requesting more information.
<code>nss /volumes</code>	Display a list of encrypted and unencrypted NSS volumes, showing their attributes.  The encrypted volume returns a status of Encrypted.

## 14.8 Dismounting an Encrypted NSS Volume with NSSMU (Linux)

Before you can dismount an NSS volume on Linux, you must dismount the volume from NCP™ Server; otherwise, the dismount function fails.

**1** If NCP Server is running, dismount the volume from NCP Server.

**1a** At the server prompt, open the NCP Console by entering

```
ncpcon
```

**1b** Dismount the volume from NCP.

The volume is no longer accessible or visible to NCP clients.

**2** Dismount the volume.

**2a** From a terminal console: In NSSMU, select *Volumes*.

- 2b** In the Volumes list, select the encrypted volume that you want to dismount.
- 2c** Select or press *F7* to dismount the mounted volume.

## 14.9 Deactivating an Encrypted NSS Volume with NSSMU (NetWare)

- 1** In NSSMU, select *Volumes*.
- 2** In the Volumes list, select the encrypted volume that you want to deactivate.
- 3** Select or press *F6* to deactivate the active volume.

## 14.10 Using Encrypted Volumes in a Server Cluster (Linux)

If you use an encrypted NSS volume in a Novell Cluster Services™ cluster, you must manually enter the password for the volume on one of the servers only when you first start or restart the cluster. You use NSSMU to mount the encrypted volume on one of the OES Linux servers and enter the volume password, then dismount volume before you can bring the cluster resource online for the first time.

NSS uses the password to create a key, which it stores in the server memory. The NCS software passes the key to other nodes. After all servers hold the key, the volume is available while any one of the servers is still participating actively in the cluster. If all servers in the cluster fail, you must repeat this procedure when you recover the cluster and restart services.

- 1** Boot or restart the servers in the cluster.

If you automated the loading of cluster resources, the cluster reports each resource is comatose because it cannot bring the corresponding encrypted volume online.

If you opt to manually start cluster resources, the cluster resources are not yet active.

- 2** From one of the nodes in the cluster, repeat the following steps for each of the encrypted volumes in the cluster.

- 2a** In NSSMU, select *Volumes*.

- 2b** In the Volumes list, select the shared volume you want to mount.

- 2c** Select or press *F7* to mount the shared volume.

- 2d** When prompted, enter the password, then click *OK*.

- If the server already knows the key for the volume, you are not prompted for the password.

- 2e** In the Volumes list, select the shared volume that you want to dismount.

- 2f** Select or press *F7* to dismount the shared volume.

- 3** Follow the normal procedures to activate the cluster resources.

For information, see the *OES Novell Cluster Services Administration Guide for Linux* ([http://www.novell.com/documentation/oes/cluster\\_admin\\_lx/data/h4hgu4hs.html](http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html)). The node passes the key information to the other nodes. While at least one of the servers is actively participating in the cluster, you do not need to reenter the encryption password again.



## 14.11 Using Encrypted Volumes in a Server Cluster (NetWare)

If you use an encrypted NSS volume in a Novell Cluster Services™ cluster, you must manually enter the password for the volume on one of the servers only when you first start or restart the cluster. You activate the NSS pool and volume on one of the servers, enter the volume password, then deactivate the pool and volume before you can bring the cluster resource online for the first time.

NSS uses the password to create a key, which it stores in the server memory. The NCS software passes the key to other nodes. After all servers hold the key, the volume is available while any one of the servers is still participating actively in the cluster. If all servers in the cluster fail, you must repeat this procedure when you recover the cluster and restart services.

- 1** Boot or restart the servers in the cluster.

If you automated the loading of cluster resources, the cluster reports each resource is comatose because it cannot bring the corresponding encrypted volume online.

If you opt to manually start cluster resources, the cluster resources are not yet active.

- 2** From one of the nodes in the cluster, activate the encrypted volumes.

- 2a** Activate the cluster pool and its encrypted volumes by entering the following commands at the server console:

```
nss /poolactivate=poolname
nss /volumeactivate=volumename
```

At the prompt, enter the password for the encrypted volume. NSS creates the 128-bit encryption key and stores it in the server's memory.

If the server already knows the key for the volume, you are not prompted for the password.

- 2b** To deactivate the cluster pool, enter

```
nss /pooldeactivate=poolname
```

- 3** Follow the normal procedures to activate the cluster resources.

For information, see the [OES Novell Cluster Services Administration Guide for NetWare \(http://www.novell.com/documentation/oes/cluster\\_admin/data/h4hgu4hs.html\)](http://www.novell.com/documentation/oes/cluster_admin/data/h4hgu4hs.html). The node passes the key information to the other nodes. While at least one of the servers is actively participating in the cluster, you do not need to reenter the encryption password again.

## 14.12 Removing Encrypted Volumes

If a rollback becomes necessary, you must remove the encrypted volume from the server before you perform the rollback.

Encrypted volumes require NetWare 6.5 Support Pack 2 and later or OES SP1 Linux or later. Because earlier releases of NSS cannot activate an encrypted volume, you cannot roll back the system to the earlier release. If you do, the encrypted volume fails to activate or mount, and its pool cannot be repaired.

To prevent this potential data loss, make sure that the system upgrade to a supported platform is active and performing as desired before creating encrypted volumes.

- 1 Create an unencrypted volume where you want to copy the data.

For information, see [Section 14.3, “Creating an Encrypted Volume,” on page 204](#).

- 2 Use one of these methods to save the encrypted volume’s data on the unencrypted volume:
  - Back up the volume’s data in unencrypted format on backup media, then restore the data to the unencrypted volume.
  - Make a volume-to-volume copy of the data from the encrypted volume to the unencrypted volume.
- 3 Delete the encrypted volume.
- 4 Perform the system rollback.

## 14.13 What’s Next

Manage other NSS features of your encrypted volume as you would for an unencrypted volume. For information, see [“Managing NSS Volumes” on page 181](#).

# Securing Access to NSS Volumes, Directories, and Files

# 15

This section discusses NSS volume attributes and other measures you can use to help secure access to your NSS volume and the files in it.

- [Section 15.1, “Controlling Access to the Admin Volume and Files.cmd \(NetWare\),” on page 211](#)
- [Section 15.2, “Using Data Shredding to Prevent Access to Deleted Files,” on page 215](#)
- [Section 15.3, “Managing File System Access to Directories and Files,” on page 216](#)

## 15.1 Controlling Access to the Admin Volume and Files.cmd (NetWare)

On your NetWare server, the `_admin:\manage_nss\files.cmd` file enables applications, such as NetStorage and other end-user browser-based tools, to communicate with the file server. The `files.cmd` file is a virtual file; it generates dynamic information that lets users see and modify their files through a browser.

Every user on the file server has access to the `_admin:\manage_nss\files.cmd` and thus inherits the following default file system trustee rights: Read, Write, and Scan. These rights make the Admin volume (`_admin:`) visible to end users. This behavior differs from NetWare 6.0 and earlier servers, where the `files.cmd` file is not visible to end users. Although end users can see the `files.cmd` file, they cannot delete this file or modify its properties.

To restrict access to the `_admin:\manage_nss\files.cmd` file or the Admin volume, you must modify a configuration file called `trustees.xml`. For more information, see [Section 15.1.1, “Assigning Volume Rights with Trustees.xml \(NetWare\),” on page 211](#).

For information about virtual files, see the *Virtual File Services for NetWare* ([http://developer.novell.com/ndk/doc/vfs/vfs\\_\\_enu/data/h5ubo7tt.html](http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html)) in the Novell Developer Kit Documentation Web site (<http://developer.novell.com/ndk/doc.htm>).

### 15.1.1 Assigning Volume Rights with Trustees.xml (NetWare)

Use the `sys:\etc\trustees.xml` file to assign file system trustee rights to the root of an NSS volume on NetWare. `Trustees.xml` is executed every time the server boots.

The `sys:\etc\trustees.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 an NSS volume.

Refer to the following table for information about `sys:\etc\trustees.xml` tags:

XML Tags	Explanation
<name>	<p>Enter the full context and username of the user. You can enter the information in slash-separated format or Novell dot format.</p> <hr/> <p><b>NOTE:</b> This tag cannot contain multiple usernames.</p> <hr/> <p>For example, if John Smith existed in the novell context in the <code>aurora_tree</code>, you would enter one of the following:</p> <p><code>\aurora_tree\novell\jsmith</code></p> <p>or</p> <p><code>jsmith.novell.aurora_tree</code></p>
<filename>	<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 <code>etc</code> directory located on the <code>sys:</code> volume, you would enter the following into the brackets</p> <p><code>&lt;filename&gt;sys:\etc\hosts&lt;/filename&gt;</code></p>
<rights>	<p>The child elements to the Rights tag include the following tags for file system rights:</p> <ul style="list-style-type: none"> <li>• <code>&lt;accesscontrol/&gt;</code> for the Access Control right</li> <li>• <code>&lt;create/&gt;</code> for the Create right</li> <li>• <code>&lt;edit/&gt;</code> for the Edit right</li> <li>• <code>&lt;filescan/&gt;</code> for the File Scan right</li> <li>• <code>&lt;modify/&gt;</code> for the Modify Metadata right</li> <li>• <code>&lt;read/&gt;</code> for the Read right</li> <li>• <code>&lt;supervisor/&gt;</code> for the Supervisor right</li> <li>• <code>&lt;write/&gt;</code> for the Write right</li> </ul>

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

```
<specialTrustees>
  <addTrustee>
    <name>admin.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>
      <read/>
      <write/>
      <filescan/>
      <edit/>
    </rights>
  </addTrustee>
</specialTrustees>
```

```

    </rights>
</addTrustee>
<addTrustee>
    <name>xxxx.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>
        <read/>
        <write/>
        <filesan/>
    </rights>
</addTrustee>
<addTrustee>
    <name>fred.novell.kona_tree</name>
    <fileName>sys:\etc\group</fileName>
    <rights>
        <read/>
        <write/>
        <filesan/>
    </rights>
</addTrustee>
<addTrustee>
    <name>[public]</name>
    <fileName>_admin:manage_nss\files.cmd</fileName>
    <rights>
        <read/>
        <write/>
        <filesan/>
    </rights>
    <background/>
</addTrustee>
<addTrustee>
    <name>fred.novell.kona_tree</name>

```

```

    <fileName>sys:\etc\hosts</fileName>

    <rights>
        <read/>
        <write/>
        <filesan/>
    </rights>
</addTrustee>
</specialTrustees>

```

### 15.1.2 Restricting Access to Files.cmd (NetWare)

If you want to restrict access to the `_admin:\manage_nss\files.cmd`, you must do so by modifying the `<addTrustee>` tag in the `sys:\etc\trustees.xml` file.

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

```

<addTrustee>
    <name>[public]</name>
    <fileName>_admin:manage_nss\files.cmd</fileName>
    <rights>
        <read/>
        <write/>
        <filesan/>
    </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>
        <filesan/>
    </rights>
    <background/>

```

</addTrustee>

For further explanation of the XML tags, see [Section 15.1.1, “Assigning Volume Rights with Trustees.xml \(NetWare\),” on page 211](#).

## 15.2 Using Data Shredding to Prevent Access to Deleted Files

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.

If the Data Shredding attribute for an NSS volume is disabled, unauthorized access to data is possible. An individual can extend a file, `LSEEK` to the end of the existing file data, and then read the data. This returns the decrypted leftover data that is in the block.

You can place up to seven data shred patterns over deleted data. Data shredding truly erases files. 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.

Data shredding consumes a great deal of disk connection bandwidth, resulting in a performance penalty for using the disk and system resources needed to overwrite the shredded file. Unless you must use data shredding for security reasons, the Data Shredding attribute for your NSS volume can be disabled or set to a lower number of shredding passes.

This section discusses the following:

- [Section 15.2.1, “Setting the Data Shredding Attribute When You Create a Volume,” on page 215](#)
- [Section 15.2.2, “Setting the Data Shredding Attribute for an Existing Volume,” on page 215](#)
- [Section 15.2.3, “Disabling Data Shredding for an Existing Volume,” on page 216](#)

### 15.2.1 Setting the Data Shredding Attribute When You Create a Volume

When you create a volume, simply select the *Data Shredding* check box and specify the number of shredding cycles with an integer number between 1 and 7 times (or specify 0 to indicate no shredding capability) when you set the volume’s attributes. For more information, see [Section 13.3, “Creating and Configuring Unencrypted NSS Volumes with iManager,” on page 186](#).

### 15.2.2 Setting the Data Shredding Attribute for an Existing Volume

- 1 In iManager, click *Storage > Volumes* to open the Volume Management page.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

Wait until the page refreshes with a list of volumes in the Volumes list.

- 3 From the *Volumes* list, select the volume that you want to manage.
- 4 Click *Properties > Attributes*.  
This opens the Volume Properties page to the Attributes tab.
- 5 Select the *Data Shredding* check box.
- 6 Specify the number of shredding cycles, where 0 is no shredding and 1 to 7 are the valid number of cycles to shred data.
- 7 Click *Apply* or *OK* to save the change, or click *Cancel* to back out of the process.  
If you click *Apply*, iManager saves the change and remains on the device page. If you click *OK*, iManager saves the change and takes you to the main Storage page. If you do not click *Apply* or *OK*, the setting is not implemented.

### 15.2.3 Disabling Data Shredding for an Existing Volume

---

**WARNING:** If you disable data shredding, an individual can recover leftover data on the drive and secure data might be exposed.

---

At the server console, enter

```
nss /nodatashredding=volumename
```

where *volumename* is the name of the volume where you want to prevent the shredding capability.

## 15.3 Managing File System Access to Directories and Files

For information about file system trustees and attributes for directories and files on NSS volumes, see the *File Systems Management Guide for OES*.

For information about access control on Linux, see [Section 6.4, “Access Control Issues for NSS on OES Linux,”](#) on page 74.



# Using Volume Attributes that Help Ensure Data Integrity

# 16

This section discusses NSS volume attributes that help ensure data integrity.

- [Section 16.1, “Using the Flush Files Immediately Attribute to Protect Against Possible Data Corruption on System Failure,” on page 217](#)
- [Section 16.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback \(NetWare\),” on page 219](#)
- [Section 16.3, “Using the File-Level Snapshot Attribute to Enable the Backup of Open Files,” on page 220](#)

Use Novell Storage Management Services™ (SMS) for backup, restore, and data migration services. For information, see the [SBCON Administration Guide](http://www.novell.com/documentation/oes/sbcon_admin/data/front.html) ([http://www.novell.com/documentation/oes/sbcon\\_admin/data/front.html](http://www.novell.com/documentation/oes/sbcon_admin/data/front.html)).

## 16.1 Using the Flush Files Immediately Attribute to Protect Against Possible Data Corruption on System Failure

The Flush Files Immediately attribute for volumes helps protect against data corruption by saving your file data from memory 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 is not actually written immediately to disk. Server operating systems juggle hundreds or thousands of demands that use memory to speed hardware-dependent functions, like disk writes, whenever possible. A file that gets written to disk might 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.

### 16.1.1 Setting the Flush Files Immediately Attribute When You Create a Volume

When you create the volume, simply select the Flush Files Immediately check box when you set the volume’s attributes. For information, see [Section 13.3, “Creating and Configuring Unencrypted NSS Volumes with iManager,” on page 186](#).

### 16.1.2 Setting the Flush Files Immediately Attribute for an Existing Volume

- 1 In iManager, click *Storage > Volumes* to open the Volume Management page.

For information, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on [page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on [page 46](#).

A list of volumes appears in the Volumes list.

- 3 From the Volumes list, select the volume that you want to manage.

- 4 Click *Properties > Attributes*.

This opens the Volume Properties page to the Attributes tab.

[Storage Management](#) ▶ [Volumes](#) ▶ **Volume Properties**

---

**Volume Properties** ?

**Properties:** NIFVOL

**Attributes** **Statistics** **Quotas Usage**

---

<input checked="" type="checkbox"/> Backup	<input type="checkbox"/> Migration
<input type="checkbox"/> Compression	<input type="checkbox"/> Modified File List(MFL)
<input type="checkbox"/> Data Shredding	<input checked="" type="checkbox"/> Salvage Files
Number of shredding cycles <input type="text"/>	<input type="checkbox"/> Snapshot - File Level
<input type="checkbox"/> Directory Quotas	<input type="checkbox"/> User Space Quotas
<input type="checkbox"/> Flush Files Immediately	<input type="checkbox"/> User-level Transaction Model

Quota:  Bytes ▾

☒ Allow volume quota to grow to the pool size

---

- 5 Select the *Flush Files Immediately* check box.

- 6 Click *Apply* or *OK* to save the change, or click *Cancel* to back out of the process.

If you click *Apply*, iManager saves the change and remains on the device page. If you click *OK*, iManager saves the change and takes you to the main Storage page. If you do not click *Apply* or *OK*, the setting is not implemented.

### 16.1.3 Enabling and Disabling the Flush Files Immediately Attribute for All Volumes on a Server

Use the NSS Flush Files Immediately option to enable and disable the capability for all volumes on the server.

To enable the Flush Files Immediately attribute, enter at the server console:

```
nss /FlushFilesImmediately
```

To disable the Flush Files Immediately attribute, enter at the server console:

```
nss /NoFlushFilesImmediately
```

## 16.2 Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare)

The Novell Transaction Tracking System™ (TTS™) protects simple database applications by backing out application-based transactions that are incomplete because of a system failure. For example, a banking application might require that all steps in a user's session be complete before updating the database and ending the transaction. If a system failover occurred and the transaction did not complete all steps successfully, all of the steps would be rolled back to their condition before the transaction began. If all steps of the transaction are complete, the TTS replays the steps and completes the transaction successfully.

By default, TTS works with the Traditional file system. TTS cannot run on NSS and Traditional volumes at the same time on the same server. If you have both Traditional and NSS volumes on your system and you want to run TTS on NSS volumes instead of Traditional volumes, you must first disable TTS on all Traditional volumes, and then enable TTS on your NSS volumes.

For more information, see [Using the Transaction Tracking System \(http://www.novell.com/documentation/beta/oes/sos\\_\\_enu/data/h0wbzsqh.html\)](http://www.novell.com/documentation/beta/oes/sos__enu/data/h0wbzsqh.html) in the *Server Operating System for NetWare Administration Guide for OES* ([http://www.novell.com/documentation/beta/oes/sos\\_\\_enu/data/h1hge52k.html](http://www.novell.com/documentation/beta/oes/sos__enu/data/h1hge52k.html)).

### Enabling TTS

To enable TTS on each NSS volume where you want it, enter `nss /transaction=volname` at the server console.

### Disabling TTS

To disable TTS on all Traditional volumes, add `disable tts` to the `autoexec.ncf` file.

### TTS Log File and Transaction Recovery

If a system failure occurs, TTS replays the TTS log file to recover lost application-based transactions as part of the failover recovery.

Before replaying the TTS log to recover lost transactions, TTS checks the integrity of the transaction file blocks in the log. If TTS detects no problems, the replay continues.

If TTS detects a bad block, it does not use the log file. TTS creates a new TTS log file to use as the system continues to operate.

---

**IMPORTANT:** Although the system continues to run, data might have been lost.

---

TTS renames the corrupted TTS log file to indicate the date and time that it detected the problem. For example, it renames

`volume:\UXaction.log`

to

`volume:\XLogyyyymmddhhmmss.log`

Corrupt TTS log files are relatively large, about 10 MB. TTS stores the corrupt TTS log file in the root directory of its transacted volumes. If users report finding TTS log files in the root directory of their volumes, it indicates that a problem occurred and that TTS detected it.

---

**NOTE:** If a system failure causes corruption in the transaction log file, please report to Novell Support the network circumstances that lead to the corruption of the TTS log file. We want to examine the event to see if we can establish why the corruption might have occurred.

---

## 16.3 Using the File-Level Snapshot Attribute to Enable the Backup of Open Files

This section discusses the following key concepts and tasks:

- [Section 16.3.1, “Understanding File-Level Snapshots,” on page 220](#)
- [Section 16.3.2, “How the File Snapshot Attribute Works,” on page 221](#)
- [Section 16.3.3, “Sizing the Volume to Accommodate File Snapshots,” on page 221](#)
- [Section 16.3.4, “Assessing I/O Overhead for File Snapshots,” on page 222](#)
- [Section 16.3.5, “Configuring Volumes with the File Snapshot Attribute,” on page 223](#)
- [Section 16.3.6, “Configuring Backup Software to Use File Snapshots,” on page 224](#)
- [Section 16.3.7, “Restoring Volumes from Backup Media,” on page 225](#)

The file snapshot attribute is not the same as the pool snapshot feature. Pool snapshot is an NSS feature that makes block-level snapshots of a pool that persist until intentionally deleted. For information, see [“Managing NSS Pool Snapshots \(NetWare\)” on page 167](#).

### 16.3.1 Understanding File-Level Snapshots

File-level snapshot is an attribute of an NSS volume that allows a temporary copy of an open file to be made at the time of a volume backup so that the last-known-good copy of the file is available for backup while the file is in use.

Normally, if a file is open when backup services begin, that open file cannot be backed up. With the file snapshot attribute set for a volume, NSS creates and maintains a virtual volume where it writes a partial or full copy of a file’s original data whenever a file on the source volume is open for write. This temporary copy of the file’s original data is made available to the backup software instead of the current contents of the open file. Consequently, open files can be copied to the backup media, even though changes in progress for the files are not backed up.

All file snapshot activity happens transparently to the user and to non-backup applications. The user and applications always read the current content of the file.

Not all backup software can take advantage of the file snapshot attribute, even if you set it.

The copy of the file is temporary. When the user or API closes the file, NSS deletes the file-snapshot copy.

## 16.3.2 How the File Snapshot Attribute Works

The file snapshot attribute uses block-level copy-on-write technology, a technique of capturing before images of data as it is modified by applications. The primary purpose of these images is to enable point-in-time backups of file systems while they are in use.

Before the file is open, the physical volume contains a file with data that fills multiple storage blocks on the device. The virtual volume contains only metadata about the file. Metadata includes owner, created date/time, and other information about the file and an empty Blocks Changed list.

Whenever the file opens, the NSS file snapshot attribute updates the metadata for the file, then monitors for writes to the file. NSS identifies writes at the block level. When it intercepts a write command, NSS does not immediately write to the physical location. Instead, NSS reads the data in the target block and copies the existing data in that block to the virtual volume. It adds the addresses of the changed blocks to the Blocks Changed list. At this point in the process, the old data exists in two places: the target block and a block in the virtual volume.

Next, NSS allows the new data to overwrite the old data in the target block. Now, the target block contains new data. The file on the physical volume has some blocks that contain old data, and some blocks that contain new data, just as it would after any normal write process. All that is stored in the virtual volume is the metadata about the file and the blocks of data that represent what the file looked like at the moment it was opened (the blocks of data that NSS copied to the virtual volume).

With the file snapshot attribute disabled, your backup software normally cannot capture a file to backup media, if the file is open and changes are in progress. The backup passes over the file and does not copy it to the backup media. The state of the file is not captured to the backup media.

When the file snapshot attribute is enabled and your backup software takes advantage of the feature, the backup knows to read blocks from the virtual volume and from the physical volume for the open file. From the physical volume, the backup software reads the blocks that are as yet unchanged by any block writes. From the virtual volume, the backup software reads the blocks that contain the file's data as it was when the file opened. The total information saved to the backup media represents the state of the file when the file was opened.

The virtual volume and its contents are not directly accessible to users or non-backup APIs. It is only accessible to Novell Storage Management Services or third-party backup software that takes advantage of the file-snapshot flag. All other applications ignore the virtual volume. If accessing an open file, the backup software reads original data from the virtual volume, but any other application reads the file's current contents.

When the open file closes, NSS updates the file's metadata and erases the blocks of original data stored in the virtual volume.

This is how the last-known-good version of the file gets saved to the backup media when it would have been passed over previously. The backup media contains real data; it does not back up the virtual volume per se. It only copies data from the virtual volume to re-create the complete last-known-good copy of the open file.

## 16.3.3 Sizing the Volume to Accommodate File Snapshots

Initially, the virtual volume occupies very little storage space. It contains only an empty changed-block list. As an application opens a file to update its data, NSS dynamically allocates space in the virtual volume where it can copy blocks of the original data. As blocks of data change, NSS copies

the original data to the virtual volume and updates the changed-block list to indicate which blocks have changed.

The size of the virtual volume depends on several factors:

- The size of the original file when it is opened
- How the files are used (such as how often they change, when they change, and who changes them)
- How many blocks of data that a single change affects (up to the number of blocks in the original file)
- How the application saves changes (such as by blocks or by the entire file)

It is a good idea to set the physical volume's size to grow dynamically to the size of the pool. Otherwise, you must periodically monitor the volume size to make sure it can accommodate the peak size of the virtual volume as the volume contents fluctuate at peak productivity times. Make sure your volume's size can also accommodate any seasonal productivity peaks.

### 16.3.4 Assessing I/O Overhead for File Snapshots

The amount of I/O overhead related to the dynamic file snapshots depends on the usage of the volume. Every time a file opens for write (not read), NSS makes a copy of the metadata about the file, but not the data itself. NSS makes a copy of only that part of the data in the file that a write changes. The algorithm is optimized in such a way that NSS does not make copies of data too often, given the behavior of most applications such as word processors.

Applications write data in different ways. Some close the file after each write, while others close the file intermittently and on demand. As soon as the application closes the file, even if the file is still open for read or for another write, NSS erases any related blocks for that file that are saved in the virtual volume. On executing the next open for write, the process begins again for the file, working with the last closed copy of the file.

Other applications write only differential data at the (logical) end of the file, leaving the existing data blocks intact. In this case, no existing data blocks are overwritten by the application, and there is no writing of data (only metadata) to the virtual volume.

However, in all cases, the backup software can still back up the last closed version of the file.

The thing to remember in assessing overhead is that every time a file is closed, NSS erases its metadata and the original data blocks that were copied to the virtual volume. The data in the virtual volume is not a perpetual copy that is constantly updated.

To assess overhead related to the File Snapshot attribute of an NSS volume:

- What types of applications are in use? How do those applications write to files? How are closes executed?
- How many files are there in the volume, categorized by application?
- How frequently are files changed? Does the frequency change by user or user group? Time of year? Type of application?
- How much data in the file changes typically with each use? How often are files opened for Read versus for Write?

## 16.3.5 Configuring Volumes with the File Snapshot Attribute

File snapshots apply only to volumes. You must configure a volume with this attribute in order to take advantage of the feature. After the attribute is set, the file-snapshot process is automatic.

You can set the file snapshot attribute for a volume as you create the volume, or you can opt to set the attribute later.

### Setting the File Snapshot Attribute for a New Volume

To enable the file snapshot attribute when you first create the volume:

- 1 Open the New Volume wizard in iManager, following the normal setup. See “[Managing NSS Volumes](#)” on page 181.
- 2 On the Attributes page in the New Volumes wizard, also select the *Snapshot - File Level* check box.
- 3 Complete the New Volume wizard.

Whenever the volume is first activated, NSS creates the virtual volume and the file snapshot process occurs automatically.

### Setting the File Snapshot Attribute for an Existing Volume

To set up the file snapshot attribute after a volume is in use requires some pre-planning. Because you must deactivate and activate the volume to allow NSS to set up the virtual volume, you should initiate the file snapshot attribute at a time when no I/O operations are in progress and all cached data is also reflected in the storage media. For example, to make sure no transactions are lost:

- Allow in-progress transactions to complete gracefully.
- Temporarily block any new transactions.
- Flush any file changes that are cached in memory to the volume.

To enable the file snapshot attribute on a given volume:

- 1 Launch a Web browser.
- 2 Click *File > Open*, then enter

`http://svrname.your-company-name.com/iManager.html`

Replace *svrname* with the actual server host name. Replace *your-company-name.com* with your company’s actual domain name. You can use the server’s IP address (such as 192.168.1.1) instead.

- 3 In iManager, log in to the server context that the server belongs to.
- 4 In *Roles and Tasks*, select *Storage > Volumes*.
- 5 Use one of these methods to select a server in the tree where you are logged in:
  - Type the Novell eDirectory server object name for the server you want to manage, then click *Apply*. For example: `svr1.novell`.
  - Click the *Search* icon to open the eDirectory Object Selector. Browse or search the list to locate the server object of the server you want to manage, then click the server name.
  - Click the *Object History* icon to select a server you have recently managed.

After iManager connects to the server, the Volumes field lists all of the volumes for the selected server. Let the page load completely before proceeding to the next step.

**6** Select a volume in the Volumes list.

**7** (Conditional) Review the *Details* area to make sure the volume is active and mounted. If it is not, click *Mount*.

The Mount command activates and mounts the volume in one step. The volume should still be selected when the page reloads. If it is not, select the volume again.

**8** Click *Properties > Attributes*.

**9** Select the *Snapshot - File Level* check box.

**10** Click *Apply*.

**11** Click *OK*.

This should return you to Volumes page. If not, select *Storage > Volumes*.

**12** Make sure the volume is selected, then click *Deactivate*.

You must deactivate and activate the volume to enable NSS to set up the virtual volume area where file snapshots are stored.

When the page reloads, the volume is in a deactive state.

**13** If not already selected, select the volume, then click *Activate*.

NSS creates the metadata about files in the virtual volume during the activation process. When the page reloads, the volume is in the activated state and the virtual volume is ready for file-snapshot process.

## Using Console Commands to Configure the File Snapshot Attribute

To enable the File Snapshot Attribute for one or multiple volumes, enter at the server console:

```
nss/FileCopyOnWrite=VolName
```

or

```
nss/FileCopyOnWrite=all
```

To disable File Snapshot, enter at the server console:

```
nss/NoFileCopyOnWrite
```

After you set up File Snapshot on your volume, deactivate the volume, then activate the volumes again. This ensures that there are no open files without a snapshot.

## 16.3.6 Configuring Backup Software to Use File Snapshots

NSS sets a flag that Novell Storage Management Services™ (SMS) and some third-party backup software can read to determine if there are copies of open files available to be backed up.

---

**IMPORTANT:** Even if you set this flag, not all third-party backup software takes advantage of the service.

---

Novell SMS automatically checks for the File Snapshot attribute when it backs up a volume. If the attribute is set, it takes advantage of the feature.



To open a snapshot file for backup, set a bit in the requested rights parameter and set the `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 whenever your file system receives an open for write request, it will wait for the backup open to close the file. 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. NSS supports any `getInfo` call that returns packed information with the `RNewStyle` bit set in the info mask. 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`

### **16.3.7 Restoring Volumes from Backup Media**

Whenever backup software tries to read an open file, it reads data from the virtual volume and the volume. The virtual volume contains a list of changed blocks for the file. It reads the original data for those blocks from the virtual volume. Then, it reads the other blocks of data from the volume itself. The data saved to backup media is the file's original content before it was open.

When you restore the volume from its backup media, it restores just as it does normally. It requires no special action.



# Managing Compression on NSS Volumes

# 17

Novell® Storage Services™ file compression uses algorithms to reduce the amount of space a file consumes in your storage system. Compression can optionally be used to conserve disk space and increase the amount of data a volume can store. No data in the file is permanently eliminated to compress the file; all original data is recovered when the file is decompressed.

This section discusses the following:

- [Section 17.1, “Understanding File Compression,” on page 227](#)
- [Section 17.2, “Planning for File Compression,” on page 230](#)
- [Section 17.3, “Configuring Compression Parameters for a Server,” on page 232](#)
- [Section 17.4, “Configuring a Volume for Compression,” on page 235](#)
- [Section 17.5, “Suspending Compression for Volumes or Files,” on page 236](#)
- [Section 17.6, “Disabling Compression for a Volume,” on page 237](#)
- [Section 17.7, “Restoring Data to a Uncompressed Volume,” on page 237](#)
- [Section 17.8, “Configuring Compression Preferences for Directories and Files,” on page 237](#)
- [Section 17.9, “Using NSS Commands to Configure and Monitor Compression,” on page 241](#)
- [Section 17.10, “Copying Compressed Files between Volumes with the NCOPY Command \(NetWare\),” on page 242](#)
- [Section 17.11, “Viewing the File Attributes with the NDIR Command \(NetWare\),” on page 242](#)
- [Section 17.12, “Repairing Compressed Volumes with the Compfix Utility,” on page 242](#)

## 17.1 Understanding File Compression

This section discusses the following:

- [Section 17.1.1, “Compression Settings,” on page 227](#)
- [Section 17.1.2, “Compression and Decompression Processes,” on page 228](#)
- [Section 17.1.3, “Compression Criteria,” on page 229](#)
- [Section 17.1.4, “Decompression Criteria,” on page 229](#)
- [Section 17.1.5, “Compression and Storage Backup,” on page 230](#)

### 17.1.1 Compression Settings

Although the cost of storage media is decreasing, you might consider compression to store more information on media where available space is limited. File compression requires configuration at for the server level, for the volume, and optionally for individual directories and files.

## Common Service Compression Parameters

At the server level, the compression parameters in Common Services settings govern when and how compression works for the NSS volumes where the compression attribute is enabled.

On a NetWare® server, the server-level compression settings also apply to NetWare Traditional volumes where the compression attribute is enabled. For information on NetWare Traditional volumes and compression, see *NetWare Traditional File System Administration Guide for OES*.

## The Volume's Compression Attribute

The volume's Compression attribute determines if its files can be compressed. You can enable the attribute when you create a new volume or add it at any time for an existing volume. After it is set, the Compression attribute persists for the life of the volume.

## Compression Preferences for Directories and Files

For individual directories and files, you can optionally set compression preferences that allow file compression to occur immediately for specified files, regardless of the server's compression parameters. You can also specify restrictions for individual files that make them ineligible for compression.

### 17.1.2 Compression and Decompression Processes

File compression and decompression processes occur in the background, as needed for each compressed volume, to support normal file access and immediate file compression settings. Scheduled compression occurs during a specified time each day. Typically, the scheduled period is set to nonpeak hours, but it can occur at any time you prefer.

During scheduled compression, NSS evaluates file time stamps and all compression settings to determine which files qualify for compression. When it queues eligible files for compression, the compression process begins and handles as many compression tasks as it can in the available time. Any remaining queued files wait for the next compression opportunity.

To minimize the impact of compression and decompression on system performance, you can limit the maximum number of concurrent process threads the system devotes to it. The system queues compression and decompression requests, then processes them as process threads become available.

NSS retains the uncompressed file during the compression process. Before NSS compresses a file, it verifies that the file system has enough space available for both the uncompressed file and the compressed file to temporarily coexist. If there is not enough space available, the file is not compressed. If compression completes successfully, NSS replaces the uncompressed file with the compressed file. If errors occur during compression, NSS discards the compressed file and marks the uncompressed file with a Cannot Compress (Cc) attribute.

NSS does not attempt to compress a file while its Cc attribute is set to On. If the file is opened and saved, its Cc attribute is reset to Off. You can also run the Compfix utility to clear the Cc attribute. For information, see [Section 17.12, "Repairing Compressed Volumes with the Compfix Utility," on page 242](#).

### 17.1.3 Compression Criteria

Typically, files that have been inactive for a specified period of time are eligible for compression. Other criteria can be used to identify files for compression:

- Files residing in a directory marked for immediate compression
- Files residing in subdirectories of a directory marked recursively for immediate compression
- Individual files that are marked for immediate compression

Several factors prevent an uncompressed file from being compressed, even if it meets inactivity criteria:

- The file is flagged with a Don't Compress (Dc) attribute.
- The file is an NSS sparse file, so its physical storage size is already minimized. A sparse file contains numerous contiguous zeros that NSS stores in a special way to conserve space.
- The amount of space freed by compressing a file does not meet the specified minimum reduction criteria. For example, if the Minimum Compression Percentage Gained parameter is set to 20%, a file would not be compressed if compression reduced its file size by only 10%.
- The file compression service is suspended.
- The queue of files marked for compression is long and cannot be completed during the specified hours set aside for compression activities.
- The volume does not contain enough space to hold both the original version and the compressed version of the file while compression occurs.
- The file has been deleted and the Deleted File Compression Option parameter does not allow compression of deleted files.

### 17.1.4 Decompression Criteria

Decompression occurs as needed to support file access, but other factors affect whether the uncompressed version or compressed version of the file remains in the volume after the access. The files remain compressed in the following cases:

- The percentage of free disk space available on the volume is insufficient to allow a decompressed file to remain in its uncompressed state.
- The Convert Compressed to Uncompressed Option parameter requires that compressed files always remain compressed.
- The file was opened for the first time for viewing only and the Convert Compressed to Uncompressed Option parameter requires that the file must be opened at least twice for viewing or opened once for modification before the file remains uncompressed after access.

Whenever you open a compressed file, NSS decompresses the file, but it keeps the compressed copy of the file while the file is open. The first time you open a compressed file for viewing only, NSS discards the decompressed copy of the file when you close the file. The compressed file remains on the system. If you open the file a second time for viewing only, the file is considered active. When you close the file, NSS keeps the uncompressed file and discards the compressed copy of the file. If you modify the file and save it, NSS saves the uncompressed file, then discards the compressed file.

## 17.1.5 Compression and Storage Backup

When you back up a compressed volume, files are written to the backup media in compressed or uncompressed format, according to how they are currently stored. To back up compressed files in uncompressed format, you must decompress the files first, then back up the files.

When you recover a volume, the files are restored in their saved format to the destination volume. If you try to restore a compressed file to a volume without compression, the file is not readable. For information, see “[Guidelines to Back Up Compressed Files](http://www.novell.com/documentation/nw65/smsadmin/data/hwcq1ix9.html)” (<http://www.novell.com/documentation/nw65/smsadmin/data/hwcq1ix9.html>) in the *NetWare 6.5 Storage Management Services Administration Guide* (<http://www.novell.com/documentation/nw65/smsadmin/data/hjc2z4tu.html>), or see your third-party backup application’s documentation.

## 17.2 Planning for File Compression

To effectively use compression for your NSS volumes, you must understand the following key concepts:

- [Section 17.2.1, “Inactivity Determines Which Files Are Eligible for Background Compression,” on page 230](#)
- [Section 17.2.2, “Some Files Do Not Compress Well,” on page 231](#)
- [Section 17.2.3, “Some Volumes Are Not Good Candidates for Compression,” on page 231](#)
- [Section 17.2.4, “Decompression Activity Depends on Available Space,” on page 231](#)
- [Section 17.2.5, “Immediate Compression Impacts CPU Performance,” on page 231](#)
- [Section 17.2.6, “Files Remain Compressed during Backup and Restore,” on page 231](#)
- [Section 17.2.7, “If Set, the Compression Attribute Persists for the Life of the Volume,” on page 232](#)
- [Section 17.2.8, “How to Monitor Compression Activity,” on page 232](#)

### 17.2.1 Inactivity Determines Which Files Are Eligible for Background Compression

NSS compresses files based on the interval of time that a file remains inactive. With background compression, files automatically pass in and out of their compressed state as they are unused, then used. It is not necessary to separate application files from data files for file compression. Most application files are used regularly and are not inactive long enough to qualify for compression.

Use the compression parameter named **Days Untouched before Compression** to set the length of the interval of inactivity. The parameter uses the date the file was last accessed for reading or writing to determine if a file is inactive, and therefore, eligible for compression.

To determine the optimal period of inactivity to use, consider the frequency of use of different types of files and your compression goals. Application files tend to be used more frequently, while user data is used less frequently. For example, the shorter the period of inactivity is, the higher the frequency of compression. The longer the period of inactivity is, the lower the frequency of compression, and the less likely it is that files are ever compressed.

On NetWare, use the `flag` command to enable immediate compression for a file or directory. For information, see [Section 17.8, “Configuring Compression Preferences for Directories and Files,” on page 237](#).

## 17.2.2 Some Files Do Not Compress Well

A file must be larger than 8 KB and smaller than 256 MB to be eligible for compression. The compression algorithm determines these limits.

To avoid the overhead of decompressing files that do not compress well, the system calculates the compressed size of a file before actually compressing it. If no disk space is saved by compression, or if the size difference does not meet the value specified by the set command's **Minimum Percentage Compression Gain parameter**, the file is not compressed.

NSS does not compress NSS sparse files. A sparse file contains numerous contiguous zeros that NSS stores in a special way to conserve space. A sparse file's logical size is larger than its physical disk usage. If a sparse file were compressed, it would actually consume more storage space than it normally does.

## 17.2.3 Some Volumes Are Not Good Candidates for Compression

Compression is not recommended for the `sys :` volume. Reserve compression for user data volumes. You cannot use compression on an NSS volume on a CD or DVD drive.

## 17.2.4 Decompression Activity Depends on Available Space

Compressed files are uncompressed as they are needed, then remain uncompressed until they are inactive for a designated period. For a file to be uncompressed, there must be enough free space on the volume to accommodate the decompression process and the uncompressed file size.

## 17.2.5 Immediate Compression Impacts CPU Performance

Compression requires processor resources, memory resources, and hard disk space during the compression and decompression processes. Compression is usually a low priority process thread because of compression's impact on performance. If you flag many large files for immediate compression during peak system usage, CPU performance might deteriorate.

You can configure the server's compression parameters to control how compression services use resources. For example, you can schedule compression to occur only during non-peak hours to free CPU resources during peak and normal hours. For information, see **Section 17.3, "Configuring Compression Parameters for a Server,"** on page 232.

## 17.2.6 Files Remain Compressed during Backup and Restore

Backup applications that use Novell Storage Management Services™ (SMS) back up and restore files in their compressed state. Uncompressed files on the volume are not compressed at backup time; they are stored and restored in their uncompressed state.

## 17.2.7 If Set, the Compression Attribute Persists for the Life of the Volume

The Compression attribute for a volume can be set when you create the NSS volume, or it can be set at any time thereafter. After you set the Compression attribute for a volume, you cannot turn it off; the parameter is in effect for the life of the volume.

You can suspend the compression activity, as needed, using the [Enable File Compression](#) parameter. This parameter suspends compression for all volumes on the server. For information, see [Section 17.5, “Suspending Compression for Volumes or Files,” on page 236](#) and [Section 17.6, “Disabling Compression for a Volume,” on page 237](#).

If you want to turn off file compression permanently, you must uncompress the data, back up the volume in its uncompressed state, then restore the uncompressed data to a new volume on which the Compression attribute is not set.

## 17.2.8 How to Monitor Compression Activity

Monitor compression activity using the `nss /compscreen` command. For instructions, see [Section 17.9, “Using NSS Commands to Configure and Monitor Compression,” on page 241](#).

# 17.3 Configuring Compression Parameters for a Server

The server’s compression parameters govern compression behavior for all NSS volumes on your server. The server-level settings apply to all files and directories in compression-enabled NSS volumes, but some settings can be overridden by individual file or directory attributes.

To configure the server’s compression parameters, use one of the following management utilities:

- Monitor (NetWare)
- Novell Remote Manager (NetWare)
- Set Command (NetWare)

This section discusses the following:

- [Section 17.3.1, “Understanding Compression Parameters,” on page 232](#)
- [Section 17.3.2, “Using Monitor \(NetWare\),” on page 234](#)
- [Section 17.3.3, “Using Novell Remote Manager \(NetWare\),” on page 235](#)
- [Section 17.3.4, “Using the Set Command \(NetWare\),” on page 235](#)

## 17.3.1 Understanding Compression Parameters

The following table describes each compression parameter, its purpose, supported values, and default value.



**Table 17-1** *Explanation of Compression Parameters*

Parameter	Description
Compression Daily Check Stop Hour	<p>Specify the hour when you want the file compressor to stop scanning enabled volumes for files that need to be compressed and to stop compressing them.</p> <p><b>Supported Values:</b> 0 to 23</p> <p>Hours are specified by a 24-hour clock: (0 = midnight; 23 = 11 p.m.).</p> <p><b>Default Value:</b> 6</p>
Compression Daily Check Starting Hour	<p>Specify the hour when you want the file compressor to start scanning enabled volumes for files that need to be compressed and to compress them.</p> <p>If the Compression Daily Check Stop Hour parameter is the same as the Compression Daily Check Starting Hour, then the file compressor starts checking every day at the Compression Daily Starting Hour time and runs as long as necessary to finish all files that meet the compressible criteria.</p> <p><b>Supported Values:</b> 0 to 23</p> <p>Hours are specified by a 24-hour clock: (0 = midnight; 23 = 11 p.m.).</p> <p><b>Default Value:</b> 0</p>
Minimum Compression Percentage Gain	<p>Set the minimum percentage a file must compress to remain in a compressed state.</p> <p><b>Supported Values:</b> 0 to 50</p> <p><b>Default Value:</b> 20</p>
Enable File Compression	<p>Specify whether file compression for the volume is suspended. While file compression is suspended, files that would have been compressed are queued, then compressed when compression is re-enabled.</p> <p><b>Supported Values:</b> On or Off</p> <p>On allows file compression on compression-enabled volumes.</p> <p>Off suspends compression; immediate compression requests are queued until the value is reset to On, when the files meeting criteria are compressed.</p> <p><b>Default Value:</b> On</p>
Maximum Concurrent Compressions	<p>Specify the maximum concurrent or simultaneous compressions allowed.</p> <p><b>Supported Values:</b> 1 to 8</p> <p><b>Default Value:</b> 2</p>

Parameter	Description
Convert Compressed to Uncompressed Option	<p>Specify what the file system does with an uncompressed version of a file after the server has decompressed it.</p> <p><b>Supported Values:</b> 0, 1, or 2</p> <p>0 = Always leave the file compressed</p> <p>1 = Leave the file compressed until second access if it is read only once during the time specified by the Days Untouched Before Compression parameter</p> <p>2 = Always leave the file decompressed</p> <p><b>Default Value:</b> 1</p>
Decompress Percent Disk Space Free to Allow Commit	<p>Specify the percentage of free disk space required on a volume for file decompression to permanently change compressed files to decompressed. This parameter prevents newly decompressed files from filling up the volume.</p> <p><b>Supported Values:</b> 0 to 75</p> <p><b>Default Value:</b> 10</p>
Decompress Free Space Warning Interval	<p>Specify the time between alerts when the file system is not changing compressed files to decompressed because of insufficient disk space.</p> <p><b>Supported Values:</b> 0 seconds to 29 days 15 hours 50 minutes 3.8 seconds</p> <p>Setting the interval to 0 turns off the alert.</p> <p><b>Default Value:</b> 31 minutes 18.5 seconds</p>
Deleted Files Compression Option	<p>Specify whether and when deleted files are compressed.</p> <p><b>Supported Values:</b> 0, 1, or 2</p> <p>0 = Do not compress deleted files</p> <p>1 = Compress deleted files the next day</p> <p>2 = Compress deleted files immediately</p> <p><b>Default Value:</b> 1</p>
Days Untouched Before Compression	<p>Specify the number of days the system waits after a file was last accessed before it is compressed. To effectively stop compression for a volume, set the elapsed time very high. Eventually, files are decompressed and remain uncompressed because they never cross the inactivity threshold.</p> <p><b>Supported Values:</b> 0 to 100000</p> <p><b>Default Value:</b> 14</p>

## 17.3.2 Using Monitor (NetWare)

Monitor is a console-based management tool that allows you to configure and monitor server parameters, such as compression. In Monitor, select Common File Systems from the Available

Options menu. The menu options include the ten compression parameters, as defined in the preceding section.

For information about using Monitor, see “**MONITOR**” in the *Utilities Reference for OES*.

### 17.3.3 Using Novell Remote Manager (NetWare)

You can modify compression parameter values using the Set Parameters option in Novell Remote Manager. For information on using Set Parameters, see “**Viewing or Changing Set Parameters**” in the *Novell Remote Manager for NetWare Administration Guide for OES*.

### 17.3.4 Using the Set Command (NetWare)

You can use the Set command at the server console to configure and view Common File System parameters for compression. Compression parameters apply to both the NSS file system and the Traditional file system.

Enter the following at the server console

```
set parameter=value
```

Replace *parameter* with the compression parameter’s name. Replace *value* with its setting.

The system is immediately configured to the specified setting. Commands issued at the command prompt override any setting in the autoexec.ncf file. Parameter settings are persistent in NSS. If the server goes down, your setting modifications are not lost.

You can place Set commands that you execute at the console prompt in the autoexec.ncf or the startup.ncf file. Whenever a parameter value is set in the autoexec.ncf file, the server configures itself to that setting each time the server is booted, unless the value is modified. Use the NWCONFIG utility to edit both the autoexec.ncf and the startup.ncf file. For information, see “**NWCONFIG**” in the *Utilities Reference for OES*.

## 17.4 Configuring a Volume for Compression

To use compression on a volume, set the volume’s Compression attribute to On. You can set the Compression attribute when you create a new volume or enable the attribute for an existing non-compressed volume. After you enable compression for a volume, you can suspend compression, but you cannot turn the Compression attribute off.

---

**IMPORTANT:** You cannot use file compression on a volume on a CD or DVD drive.

---

For information, see [Section 17.5, “Suspending Compression for Volumes or Files,” on page 236](#).

### 17.4.1 Enabling Compression for a New Volume

When you create a new volume, simply select the Compression check box when you set the volume’s attributes. For information, see [Section 13.3, “Creating and Configuring Unencrypted NSS Volumes with iManager,” on page 186](#).

## 17.4.2 Enabling Compression for an Existing Volume

- 1 In iManager, click *Storage > Volumes* to open the Volume Management page.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on page 46.

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.

A list of volumes appears in the Volumes list.

- 3 From the Volumes list, select the volume that you want to manage.

- 4 Click *Properties > Attributes*.

This opens the Volume Properties to the Attributes page.

[Storage Management](#) > [Volumes](#) > **Volume Properties**

---

**Volume Properties** ?

**Properties:** NIFVOL

**Attributes** **Statistics** **Quotas Usage**

<input checked="" type="checkbox"/> Backup	<input type="checkbox"/> Migration
<input type="checkbox"/> Compression	<input type="checkbox"/> Modified File List(MFL)
<input type="checkbox"/> Data Shredding	<input checked="" type="checkbox"/> Salvage Files
Number of shredding cycles <input type="text"/>	<input type="checkbox"/> Snapshot - File Level
<input type="checkbox"/> Directory Quotas	<input type="checkbox"/> User Space Quotas
<input type="checkbox"/> Flush Files Immediately	<input type="checkbox"/> User-level Transaction Model

Quota:  Bytes

☒ Allow volume quota to grow to the pool size

---

- 5 Select the *Compression* check box to enable file compression for the selected volume.

---

**IMPORTANT:** After compression is enabled for a volume, you cannot deselect the Compression check box to turn off compression. You can suspend compression or restore uncompressed data to an uncompressed volume. For information, see [Section 17.5, “Suspending Compression for Volumes or Files,”](#) on page 236 or [Section 17.7, “Restoring Data to a Uncompressed Volume,”](#) on page 237.

---

- 6 Click *Apply* or *OK* to save the change, or click *Cancel* to back out of the process.

If you click *Apply*, iManager saves the change and remains on the Volumes page. If you click *OK*, iManager saves the change and takes you to the main Storage page. If you do not click *Apply* or *OK*, the setting is not implemented; the volume remains uncompressed.

## 17.5 Suspending Compression for Volumes or Files

You can temporarily suspend file compression for all volumes, using the `Set` command's **Enable File Compression** parameter. While file compression is suspended, files that would have been compressed are queued for compression, then compressed when compression is re-enabled. For

information, see [Section 17.8, “Configuring Compression Preferences for Directories and Files,” on page 237](#).

You can suspend compression for a file by setting its **Don't Compress** flag. For information, see [Section 17.8, “Configuring Compression Preferences for Directories and Files,” on page 237](#).

## 17.6 Disabling Compression for a Volume

After compression is enabled for an NSS volume, you cannot deselect the Compression attribute check box to turn off compression. Use one of the following methods to disable compression for a volume.

- To effectively stop compression for a volume, set the **Days Untouched Before Compression** parameter's elapsed time to the maximum value. Eventually, files are decompressed as they are used and remain uncompressed because they never cross the threshold for inactivity.
- Restore the data to an uncompressed volume. For information, see [Section 17.7, “Restoring Data to a Uncompressed Volume,” on page 237](#).
- Set **Enable File Compression** to Off. Eventually, files are decompressed as they are used and remain uncompressed.

## 17.7 Restoring Data to a Uncompressed Volume

- 1 Decompress the volume data.
- 2 Back up the uncompressed data.
- 3 Create a new volume with the Compression attribute disabled.
- 4 Restore the old volume contents as uncompressed data from your backup media.

## 17.8 Configuring Compression Preferences for Directories and Files

File and directory compression preferences override the compression settings for the volume and server.

- [Section 17.8.1, “Using the Novell Client,” on page 237](#)
- [Section 17.8.2, “Using ATTRIB \(Linux\),” on page 238](#)
- [Section 17.8.3, “Using FLAG \(NetWare\),” on page 239](#)

### 17.8.1 Using the Novell Client

- 1 From a workstation, click the Novell Client icon (the big red N in the status bar), select Novell Map Network Drive, then map a drive to the NSS volume, using the login and password of the admin user.
- 2 Use either of the following methods to open the NetWare Info dialog for the file or directory you want to manage:
  - In a file manager, navigate to the directory or file you want to manage. Right-click the directory or file, select Properties, then select the NetWare Info tab in the Properties window.

- Click the Novell Client icon, select NetWare Utilities > Object Properties, navigate to and select the directory or file you want to manage, click *OK*, then select the NetWare Info tab in the NetWare Services window.

**3** Do any of the following:

- Select its check box to enable (set) the Don't Compress or Immediate Compression attribute.
- Deselect its check box to disable (clear) the Don't Compress or Immediate Compression attribute.

**4** Click *OK* or *Apply* to accept the changes, or click *Cancel* to back out of the process.

## 17.8.2 Using ATTRIB (Linux)

Use the `attrib` utility at the Linux terminal console prompt to view or modify file and directory attributes for compression on NSS volumes on Linux where compression is enabled.

### Syntax

```
attrib [options] [filename]
```

### Options

If both the `set` and `clear` options are selected, the `clear` option is completed before the `set` option. If the `filename` is not specified, the operation is completed on the current directory.

Option	Description
<code>-s, --set=ATTRIBUTES</code>	Sets the attributes on the specified file or directory.
<code>-c, --clear=[ATTRIBUTES   all]</code>	Clears the attributes on the specified file or directory.

### Configuring Compression Attributes

Attribute	Description	Applies to Files	Applies to Directories
<code>dc</code>	Don't Compress keeps data from being compressed. This attribute overrides settings for automatic compression of files not accessed within a specified number of days.	Yes	No
<code>ic</code>	Immediate Compression sets data to be compressed as soon as a file is closed. If applied to a directory, every file in the directory is compressed as each file is closed.  The files in the specified directory are compressed as soon as the operating system can perform the operation after the file is closed. This does not apply to the directory's subdirectories and the files in them.	Yes	Yes

For example, to set the Don't Compress attribute for all files in the current directory, enter

```
attrib --set=dc
```

For example, to clear the Immediate Compression attribute from the `/usr/course/winter/students.sxi` file, enter

```
attrib --clear=ic /usr/course/winter/students.sxi
```

## Viewing Compression Status for Files and Directories

Enter `attrib` without options to show the compression attribute information for a specified file or for all files in the directory. You cannot modify Status attributes.

Status Attribute	Description
cc	Cannot Compress (status display only) displays if the file cannot be compressed because of limited space savings.
cm	Compressed (status display only) indicates whether the file is currently stored in compressed format.

To view the compression status of a file, enter the following at the server console:

```
attrib filename
```

Replace *filename* with the path to the file. For example, to view the attributes of the `/usr/course/winter/students.sxi` file, enter

```
attrib /usr/course/winter/students.sxi
```

To view the compression attributes of all files in the current directory, enter

```
attrib
```

## 17.8.3 Using FLAG (NetWare)

Use the `flag` command (`sys:\public\flag.exe`) at a DOS command prompt to view or modify file and directory attributes for compression on NSS volumes on NetWare where compression is enabled.

FLAG is client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the Novell Map Network Drive option in the client menu to map a local drive letter to the `sys:\public` directory. Open a DOS Command Prompt window, `cd` to the mapped drive, then issue the `flag` command from the command line.

### Syntax

```
flag path [[+ | -] attribute...] [/option...]
```

### Options

The following table provides information about compression-related options for the `flag` command.

Parameter	Description
(no parameter)	If you type FLAG with no parameters, the attribute status of all the files in the current directory is displayed.
<i>path</i>	Specify the path to the file or directory whose attributes or search modes you want to modify.
+	Set specified attributes for a specified file or directory with the + (plus), or specify neither + nor - (minus).  If you set and clear attributes in the same command, group all + (plus) attributes together and all - (minus) attributes together.
-	Clear attributes from a specified file or directory with the - (minus).  If you set and clear attributes in the same command, group all + (plus) attributes together and all - (minus) attributes together.
<i>attribute</i>	Specify the compression-related attribute, such as Immediate Compression (Ic) or Don't Compress (Dc).
<i>/option</i>	Replace option with any available option.  Use the /S option to search the subdirectories in the specified directory recursively to apply the command to files in the directory and to its subdirectories and the files beneath them.  For other options, see "FLAG" in the <i>Utilities Reference for OES</i> .

## Configuring Compression Attributes for a Directory

Use the Flag command to set or clear the Immediate Compression attribute for a specified directory.

Attribute	Description
Ic (Immediate compress)	Compress the files in the specified directory as soon as the operating system can.  This does not apply to the directory's subdirectories and the files in them. Use Ic in combination with the /S option to apply immediate compression to the directory's subdirectories and their files recursively.

For example, to set the Immediate Compression attribute for the *sys:\course\winter* directory, enter

```
flag sys:\course\winter +Ic
```

For example, to clear the Immediate Compression attribute from the *sys:\course\winter* directory, enter

```
flag sys:\course\winter -Ic
```

## Configuring Compression Attributes for a File

Use the flag command to set or clear the following compression attributes for a specified file.



Attribute	Description
Dc (Don't Compress)	Prevent a specified file from being compressed.
Ic (Immediate compress)	Compress a file as soon as the operating system can.

For example, to set the Don't Compress attribute to all files in the current directory, enter

```
flag *.* +Dc
```

For example, to clear the Immediate Compression attribute from the `sys:\course\winter\students.sxi` file, enter

```
flag sys:\course\winter\students.sxi -Ic
```

### Viewing Compression Status for Files and Directories

Status flags show the compression attribute information for a file or directory. You cannot modify Status flags.

Status Attribute	Indicates
Cc (Cannot compress)	The file or directory cannot be compressed because of limited space savings.
Co (Compressed)	The file is currently stored in compressed format.

To view the compression status of a file, enter the following at the server console:

```
flag path
```

Replace *path* with the path to the directory or file. For example, to view the attributes of the `students.sxi` file in `sys:\course\winter` directory, enter

```
flag sys:\course\winter\students.sxi
```

To view the attributes of all files in the current directory, enter

```
flag
```

## 17.9 Using NSS Commands to Configure and Monitor Compression

NSS offers the following commands for configuring and monitoring compression. Enter the commands at the server console. On Linux, issue the commands from `nsscon`.

Command	Description
<code>nss /Compression=volume_name</code>	Enable the Compression attribute for the specified volume.
<code>nss /Compression=all</code>	Enable the Compression attribute for all volumes on the server.

Command	Description
<code>nss /StopNormalCompression</code>	Stop all queued compression for files, based on the compression triggered by a file open or close.
<code>nss /BGCompression</code>	Allow compression to occur in the background at any time, instead of only within specified hours.
<code>nss /NoBGCompression</code>	Stop background compression. Allow compression to occur only within the specified hours.
<code>nss /CompScreen</code>	Display the NSS volume compression statistics on the compression screen.

## 17.10 Copying Compressed Files between Volumes with the NCOPY Command (NetWare)

The `ncopy` command provides two attributes to specify your intention for copying a compressed file to a compressed volume or a non-compressed volume. For information, see [Section A.11, “NCOPY \(NetWare\),” on page 377](#).

## 17.11 Viewing the File Attributes with the NDIR Command (NetWare)

The `ndir` command allows you to view file attributes and compression information for all files in a directory. For information, see [Section A.12, “NDIR \(NetWare\),” on page 378](#).

## 17.12 Repairing Compressed Volumes with the Compfix Utility

Use the `Compfix` utility to repair compression information for compressed volumes or to clear the Cannot Compress attribute for files in the compressed volume. For information, see [Section A.4, “COMPFIX,” on page 362](#).

# Conserving Disk Space with Volume, Directory, or User Space Quotas

# 18

This section discusses the following:

- [Section 18.1, “Understanding Space Quotas,” on page 243](#)
- [Section 18.2, “Managing NSS Volume Quotas,” on page 245](#)
- [Section 18.3, “Managing Directory Quotas,” on page 246](#)
- [Section 18.4, “Managing User Space Quotas,” on page 250](#)

## 18.1 Understanding Space Quotas

You can conserve space in an NSS volume by restricting the amount of space available to a particular volume, directory, or user. These space restrictions, or quotas, work independently, with the lower value being the most restrictive if all constraints apply. NSS allocates the space as it is needed.

The following table indicates quotas supported by NSS on NetWare® and Linux:

Space Restriction	NSS on NetWare	NSS on Linux
Volume quotas	Yes	Yes
Directory quotas	Yes	Yes
User space quotas	Yes	Yes, for OES SP1 Linux and later  Before setting quotas, make sure to Linux-enable the users with Linux User Management so that file ownership can be tracked and space consumption can be enforced.  If you Linux-enable a user who has been logged into the system before being Linux-enabled, make sure execute the <code>resetidcache</code> command from the NSS Console (nsscon) utility afterwards. This allows the proper reporting of ownership because it resets the mapping of user identities in the ID cache and forces it to update with the Linux UID for the user.

If you set a quota to a value equal to or less than the current size of space in use for the specified volume, directory, or user, users cannot add files until enough files are deleted to free up space in the volume, directory, or user files. Users can continue to access existing files for which they are authorized users, but they cannot save them.

**WARNING:** As the amount of space consumed by a user’s files approaches the user’s space quota, users should use caution when saving files. Data loss can occur if the user attempts to save a file that

is too large for the remaining unused space. If storing the file would cause the quota to be exceeded, only part of the file is actually saved, resulting in data corruption.

---

## Volume Quotas

When you create an NSS volume, you have the option of setting a space quota for the volume or letting it grow to the size of the pool. At any time thereafter, you can view and configure the volume quota from the *Storage > Volumes > Properties > Attributes* tab in iManager. For information, see [Section 18.2, “Managing NSS Volume Quotas,” on page 245](#).

If you set a volume quota to grow to the pool size, you can also add segments to the volume’s pool to expand its size, and, therefore, expand the volume quota.

As a volume nears its quota, automatic controls can be configured to manage space. For information, see [Section 24.3, “Monitoring Quotas and Space Usage for NSS Pools and Volumes,” on page 322](#).

## Directory Quotas

Directory quotas limit the space available in an individual NSS directory. To use directory quotas on an NSS volume, you must first enable the Directory Quotas attribute.

---

**IMPORTANT:** On Linux, after you enable or disable the Directory Quotas attribute, you must restart eDirectory to enforce the change.

---

For information about configuring attributes when you create a volume, see [“Managing NSS Volumes” on page 181](#).

For information about setting directory quotas, see [Section 18.3, “Managing Directory Quotas,” on page 246](#)

## User Space Quotas

User space restrictions limit the space available to a user of the NSS volumes across all directories and files owned by the user. To apply user space restrictions, you must first enable the User Space Quotas attribute.

User space quotas are supported on NetWare and on OES SP1 Linux and later. On Linux, users must be additionally enabled for Linux User Management so that NSS can track file ownership and enforce space consumption.

For information about configuring attributes when you create a volume, see [“Managing NSS Volumes” on page 181](#).

For information, see [Section 18.4.3, “Configuring User Space Quotas,” on page 252](#).

## Example of Directory and User Space Quotas

Quotas are beneficial for systems where you want to control how your storage resources are used. In environments such as a university, where you set up a common work area for a large number of students and you want to limit the space that directory can consume, set a Directory Quota. You might also limit the amount of space an individual user’s work can consume by setting the User Quota.

For example, if a directory's quota is 500 MB and the user's quota is 1 GB, the user is limited to up to 500 MB in the specific directory. If the user can access multiple directories, each with a 500 MB quota, the maximum space the user's work can consume for all directories combined is limited to the user's 1 GB administrative limit.

## 18.2 Managing NSS Volume Quotas

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 In the Volumes list, select a volume that you want manage.

- 4 Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. It opens to the Attributes tab.

[Storage Management](#) ► [Volumes](#) ► **Volume Properties**

**Volume Properties** ?

**Properties:** NIFVOL

**Attributes** **Statistics** **Quotas Usage**

<input checked="" type="checkbox"/> Backup	<input type="checkbox"/> Migration
<input type="checkbox"/> Compression	<input type="checkbox"/> Modified File List(MFL)
<input type="checkbox"/> Data Shredding	<input checked="" type="checkbox"/> Salvage Files
Number of shredding cycles <input type="text"/>	<input type="checkbox"/> Snapshot - File Level
<input type="checkbox"/> Directory Quotas	<input type="checkbox"/> User Space Quotas
<input type="checkbox"/> Flush Files Immediately	<input type="checkbox"/> User-level Transaction Model

Quota:  Bytes ▾

☒ Allow volume quota to grow to the pool size

- 5 In the *Quota* field, do one of the following:

- **No Quota:** Select *Allow Volume Quota to Grow to the Pool Size*. NSS pools allow overbooking so the administrative sum of all volumes' quotas in a pool can exceed the physical pool quota.
- **Quota:** Deselect *Allow Volume Quota to Grow to the Pool Size*, then specify the maximum size you want to allow volume to grow. The quota cannot exceed the pool size.

- 6 Click *Apply*.

- Click the *Quotas Usage* tab to view the volume and pool space usage for the selected volume and to verify the new setting.

[Storage Management](#) ▶ [Volumes](#) ▶ [Volume Properties](#)

### Volume Properties

Properties: NIFVOL

Attributes

Statistics

Quotas Usage

Volume Usage: NIFVOL

Pool Usage: NIFPOOL



Quota:	none	Available Space:	89.93 MB
		Free:	89.93 MB
Used Space:	576.00 KB	Purgeable:	0.00 MB
Compressed:	0.00 Bytes		
Other in-use space:	576.00 KB	Total Space:	99.00 MB
		Used:	9.07 MB
Purgeable Space:	0.00 MB	Booked:	0.00 MB
Available Space:	89.93 MB	This pool contains logical volumes with no quotas. They were not used to calculate the booking.	

Close

## 18.3 Managing Directory Quotas

Before you can set directory quotas, you must enable the volume's Directory Quotas attribute. As the Admin user, you can view and configure directory quotas with NetStorage and the Novell Client™. For NSS on NetWare, you can also use Novell Remote Manager for NetWare.

This section discusses the following:

- [Section 18.3.1, “Setting the Directory Quotas Attribute for an NSS Volume,” on page 246](#)
- [Section 18.3.2, “Setting a Directory Quota,” on page 247](#)
- [Section 18.3.3, “Deleting All Directory Quotas for an NSS Volume,” on page 250](#)

### 18.3.1 Setting the Directory Quotas Attribute for an NSS Volume

After you set up and configure NSS volumes, you can view and modify attribute settings such as the Directory Quotas attribute.

- In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- In the Volumes list, select a volume that you want manage.

4 Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. It opens to the Attributes tab.

5 On the *Attributes* tab, select or deselect the *Directory Quotas* check box, then click *Apply*.

6 (Linux) If you enabled or disabled the Directory Quotas attribute, restart NCP2NSS by entering at a terminal prompt:

```
/etc/init.d/ncp2nss restart
```

## 18.3.2 Setting a Directory Quota

- “Novell NetStorage” on page 247
- “Novell Client” on page 248
- “Novell Remote Manager for NetWare” on page 249

### Novell NetStorage

Using Novell NetStorage, you can manage directory quotas for directories in an NSS volume from any computer with a supported Web browser. This requires you to first configure a NetStorage server in the same context. For information, see the following:

- *OES NetStorage Administration Guide for Linux* ([http://www.novell.com/documentation/oes/netstor\\_lx/data/h9izvdye.html](http://www.novell.com/documentation/oes/netstor_lx/data/h9izvdye.html))
- *OES NetStorage Administration Guide for NetWare* (<http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html>)

To create or change NSS directory quotas with NetStorage:

1 In a Web browser, connect to NetStorage.

For information, see [Section 4.4, “Novell NetStorage,” on page 58](#).

2 Log in to NetStorage with the admin username and password.

You can also log in as any user with equivalent rights to the admin user.

3 Navigate to the directory you want to manage.

4 Right-click the directory, then select *Properties*.

5 Click the *NetWare Info* tab.

Use the NetWare Info tab whether your server is OES Linux or NetWare.

6 Do one of the following to configure the directory quota:

- **Space Restriction:** Select *Restrict Size*, then specify the directory quota in KB. The value must be a multiple of 4.
- **No Space Restriction:** Deselect *Restrict Size* to set the directory quota to Unlimited.
- **Complete Space Restriction:** Select *Restrict Size*, then specify the directory quota as 0 KB. If the directory already contains files and subdirectories, the directory cannot grow beyond the current space consumed.

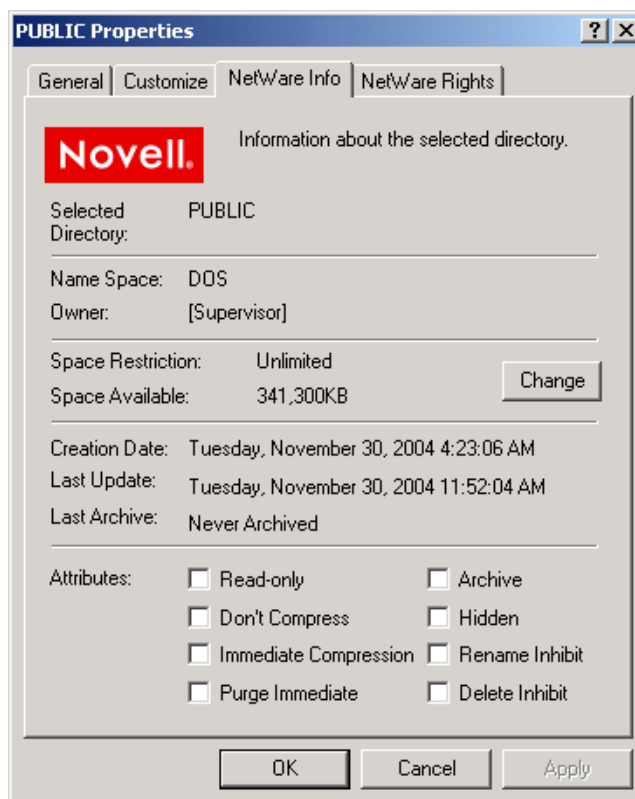
7 Click *Apply* to accept the directory quota configuration.

## Novell Client

The Novell Client for Windows 2000/XP allows the admin user to manage directory quotas for directories in an NSS volume from a Windows computer.


- 1** In the Novell Client, map a drive to the NSS directory you want to manage, or map to its parent directory.
  - 1a** Right-click the *Novell Client* icon (the N icon on the status bar), then select Novell Map Network Drive.
  - 1b** Specify the network path to the directory. For example: 192.168.1.1/users.
  - 1c** Specify admin as the username, then click Map.
  - 1d** When prompted, enter the admin user password.
- 2** In a file browser, locate and right-click the directory you want to manage, then click *Properties* > *NetWare Info*.

Use the NetWare Info tab whether your server is OES Linux or NetWare.






- 3 In the *Space Restriction* field, click *Change* to open the Space Restriction dialog box.



- 4 Do one of the following to configure the directory quota:
  - **Space Restriction:** Specify the directory quota in KB. The value must be a multiple of 4.
  - **No Space Restriction:** Select *No Space Restriction* to set the directory quota to Unlimited.
  - **Complete Space Restriction:** Select *Complete Space Restriction* to set the directory quota to 0 KB. If the directory already contains files and subdirectories, the directory cannot grow beyond the current space consumed.
- 5 Click *OK* to accept the directory quota.

## Novell Remote Manager for NetWare

- 1 Access Novell Remote Manager for your NetWare server.

For information, see [Section 4.5, “Novell Remote Manager,” on page 60](#).
- 2 Click *Volumes* to open the Volume Management page.
- 3 Locate the volume where the directory resides, then click the volume name link to view a directory listing for the volume.
- 4 Locate the directory you want to manage, then click the *Directory Information* icon  to the left of the directory to open the Directory management page.
- 5 In the *Directory Entry Information* field, click the *File Space Limit* link.

**Figure 18-1** NetStorage Directory Quota's File Space Limit Link

### Directory entry information

Owner	ADMIN
Creation date and time	Feb 1, 2005 5:11 pm
Effective rights	SRWCEMFA
Inherited rights filter	<a href="#">SRWCEMFA</a>
File space limit	<a href="#">None</a>
File space in use	0 KB

This link is not available if you did not enable the Directory Quotas attribute when you created the parent volume.

- 6 Do one of the following:
  - **Space Restriction:** Specify the maximum size (in kilobytes) for the selected directory, then click *Set Space Restriction* to apply the setting.
  - **No Space Restriction:** Click *No Quota* to remove space restrictions for the selected directory.
  - **Complete Space Restriction:** Specify the maximum size as 0 KB. If the directory already contains files and subdirectories, the directory cannot grow beyond its current space consumed.

### 18.3.3 Deleting All Directory Quotas for an NSS Volume

To delete the directory quotas for all directories on an NSS volume without dealing individually with each directory, you can simply disable the Directory Quotas attribute for the NSS volume.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).
- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.
- 3 In the Volumes list, select a volume that you want manage.
- 4 Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. It opens to the Attributes tab.
- 5 On the *Attributes* tab, deselect the *Directory Quotas* check box, then click *Apply*.
- 6 (Linux) Restart NCP2NSS by entering at a terminal prompt:

```
/etc/init.d/ncp2nss restart
```

## 18.4 Managing User Space Quotas

Before you can set user space quotas, you must enable the volume’s User Space Quotas attribute. As the Admin user, you can view and configure user space quotas in Storage plug-in for iManager.

User space quotas are available on NetWare and on OES SP1 Linux and later. For Linux, users must be Linux-enabled eDirectory users to allow NSS to track file ownership and enforce user space restrictions. For information, see [Section 6.4.6, “Enforcing File Ownership and User Space Restrictions,” on page 78](#).

This section discusses the following:

- [Section 18.4.1, “Setting the User Space Quotas Attribute for an NSS Volume,” on page 251](#)
- [Section 18.4.2, “Viewing User Space Quotas,” on page 251](#)
- [Section 18.4.3, “Configuring User Space Quotas,” on page 252](#)
- [Section 18.4.4, “Modifying a User Quota,” on page 253](#)
- [Section 18.4.5, “Deleting a User Quota,” on page 253](#)

- [Section 18.4.6, “Setting Quotas for Users Who Were Not Initially Enabled for Linux Access \(Linux\),” on page 254](#)

## 18.4.1 Setting the User Space Quotas Attribute for an NSS Volume

After you set up and configure NSS volumes, you can view and modify attribute settings such as the User Space Quotas attribute.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of volumes appears in the Volumes list.

- 3 In the Volumes list, select a volume that you want manage.

- 4 Click *Properties*.

The Properties page has three tabs: Attributes, Statistics, and Volume Quota. It opens to the Attributes tab.

- 5 On the Attributes tab, select (enable) or deselect (disable) the *User Space Quotas* check box, then click *Apply*.

## 18.4.2 Viewing User Space Quotas

- 1 In iManager, click *Storage > Volumes* to open the Volume Management page.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to view and manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

When the page refreshes, a list of volumes appears in the Volumes list.

- 3 In the Volumes list, select the volume that you want to manage.


Wait for the page to refresh to see details and options available for that volume.

- 4 View users and their quotas:

- To view only users with quotas, click *User Quotas > Users with Quotas*.

This opens the User Quotas page to the Users with Quotas tab.

[Storage Management](#) ▸ [Volumes](#) ▸ [User Restriction](#)

**User Quotas** 

**Users with Quotas** **All Users**


A storage quota restricts the amount of space on the volume available to a user. Because NSS allocates space as needed, each quota can be as big as you want to make it. View all users or only the users with quotas. Then, add, modify, and delete quotas for users.

	User Name	Quota	Space in Use	Space Available
<input type="checkbox"/>	acattt.novell	1000.00 MB	0.00 Bytes	1000.00 MB
<input type="checkbox"/>	bpeep.novell	1000.00 MB	0.00 Bytes	1000.00 MB
<input type="checkbox"/>	pcharming.novell	1000.00 MB	0.00 Bytes	1000.00 MB

[Add...](#) [Modify...](#) [Delete](#)

- To view all users (with or without quotas) for that volume, click *User Quotas > All Users*. This opens the User Quotas page to the All Users tab.

[Storage Management](#) ▸ [Volumes](#) ▸ [User Restriction](#)

**User Quotas** 

**Users with Quotas** **All Users**

A storage quota restricts the amount of space on the volume available to a user. Because NSS allocates space as needed, each quota can be as big as you want to make it. View all users or only the users with quotas. Then, add, modify, and delete quotas for users.

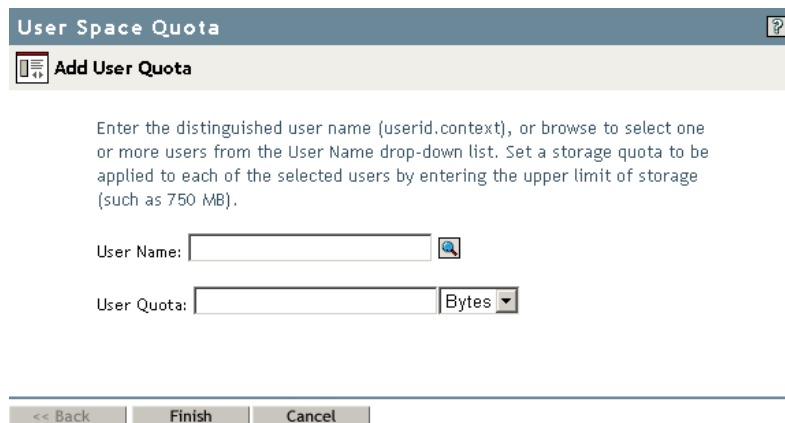
	User Name	Quota	Space in Use	Space Available
<input type="checkbox"/>	. [Supervisor].	No User Quotas	943.51 MB	
<input type="checkbox"/>	admin.novell	1000.00 MB	6.98 MB	
<input type="checkbox"/>	bpeep.novell	1000.00 MB	0.00 Bytes	1000.00 MB
<input type="checkbox"/>	pcharming.novell	1000.00 MB	0.00 Bytes	1000.00 MB
<input type="checkbox"/>	NFAUser.novell	1000.00 MB	0.00 Bytes	

[Add...](#) [Modify...](#) [Delete](#)

### 18.4.3 Configuring User Space Quotas

- 1 In iManager, click *Storage > Volumes*, select a server, then select the volume you want to manage.
- 2 Click *Quotas > All Users*.

- 3 Click *Add* to open the User Space Quota wizard.



The screenshot shows a window titled "User Space Quota" with a sub-header "Add User Quota". Below the header, there is instructional text: "Enter the distinguished user name (userid.context), or browse to select one or more users from the User Name drop-down list. Set a storage quota to be applied to each of the selected users by entering the upper limit of storage (such as 750 MB)." Below this text are two input fields: "User Name:" followed by a text box and a magnifying glass icon, and "User Quota:" followed by a text box and a dropdown menu currently set to "Bytes". At the bottom of the window are three buttons: "<< Back", "Finish", and "Cancel".

- 4 Click the *Browse* icon to open the User object list, then browse to select one or more users who you want to share the same quota.
- 5 In *User Quotas*, specify the maximum quota you want to assign to each of the selected users.
- 6 Use the *Units* drop-down list to select the units to use for the quota you specified in [Step 5](#). Options are bytes, kilobytes, megabytes, and gigabytes.
- 7 Click *Finish*.  
Each of the usernames you selected now has user quotas assigned to it.

## 18.4.4 Modifying a User Quota

- 1 In iManager, click *Storage > Volumes*, select a server, then select the volume you want to manage.
- 2 Click *Quotas > Users with Quotas*.
- 3 Select the check box next to one or more user names whose user space quota you want to delete, then click *Modify* to open the Modify User Quota wizard.  
The Modify User Quota wizard is similar to the Add User Quota wizard.
- 4 Specify the user space quota you want to apply individually to each of the selected users.
- 5 Use the *Units* drop-down list to select the units to use for the quota you specified in [Step 4](#). Options are bytes, kilobytes, megabytes, and gigabytes.
- 6 Click *Finish*.  
Each of the users you selected now has the new user quota assigned to his or her individual account on this volume.

## 18.4.5 Deleting a User Quota

- 1 In iManager, click *Storage > Volumes*, select a server, then select the volume you want to manage.
- 2 Click *Quotas > Users with Quotas*.
- 3 Select the check box next to one or more user names whose user space quota you want to delete, and then click *Delete*.

4 Do one of following:

- To confirm, click *Yes* to remove the quotas. The user accounts no longer have quotas assigned to them for the selected volume.
- To back out of the process, click *No*.

## 18.4.6 Setting Quotas for Users Who Were Not Initially Enabled for Linux Access (Linux)

Users who are not enabled for Linux access with Linux User Management can work on NSS volumes through NCP clients and Web-based file services, such as NetStorage. Files created by these users are not counted against User Quotas because the system cannot map file ownership between POSIX and eDirectory, and so all newly created files are owned by the root user. For information, see [Section 6.4, “Access Control Issues for NSS on OES Linux,” on page 74](#).

If you have created an NSS volume on an OES Linux server or moved a volume from NetWare to OES Linux, and you did not enable volume users for Linux prior to allowing users to access the volume, be aware of the following points:

- File ownership for files initially created on NSS volumes while on NetWare servers remains intact.
- Files created on OES Linux servers by users who are not enabled for Linux access are owned by the root user.

This applies both to files on newly created NSS volumes and to files on NSS volumes moved from NetWare servers.

- User quotas in force while an NSS volume was on a NetWare server are no longer in force after the volume is moved to OES Linux.

If you decide to assign User Quotas to users who created files on an NSS volume while they were not enabled for Linux access, you must do the following:

- 1 Disconnect any connected users you are enabling for Linux access.
- 2 In iManager, use the Linux User Management plug-in to Linux-enable the user.
- 3 If you want to change the ownership of any files the user normally uses that were previously assigned to the root user as owner, use the Novell Client to modify ownership of those files. Alternately, at the server's shell prompt, log in as the root user, then use the `chown` command to modify file ownership.

---

**IMPORTANT:** Directories and files created while users were not Linux enabled are owned by the root user and are not counted against the quotas you set unless you reassign file ownership.

---

- 4 Establish quotas for each user by following the instructions in [Section 18.4.3, “Configuring User Space Quotas,” on page 252](#).
- 5 At the server's shell prompt, start the `nsscon` utility and run the `resetidcache` command. This is required for displaying the correct owner name. Otherwise, the owner is listed as Nobody.
- 6 If you want to verify the file ownership from the Linux interface, at a server's shell prompt, log in as the root user, then use the `ls` command to list the files in the user's directory.
- 7 Activate the volume for user access by following the instructions in [Section 13.14, “Mounting and Dismounting an NSS Volume \(Linux\),” on page 195](#).

# Salvaging and Purging Deleted Volumes, Directories, and Files

# 19

This section discusses the following:

- [Section 19.1, “Configuring the NSS Volume Salvage Area,” on page 255](#)
- [Section 19.2, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 255](#)
- [Section 19.3, “Setting the Low and High Watermarks for Purging the File System Salvage,” on page 257](#)
- [Section 19.4, “Salvaging or Purging Deleted Directories from the File System Salvage Area \(OES SP1 NetWare, NetWare 6.5 SP4, or Later\),” on page 257](#)
- [Section 19.5, “Salvaging or Purging Deleted Files from the File System Salvage Area,” on page 259](#)

## 19.1 Configuring the NSS Volume Salvage Area

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

```
nss /logicalVolumePurgeDelay=value
```

In this command, replace *value* 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, you would 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, add the command to the `autoexec.ncf` file.

The volume salvage and file system salvage function separately. For information about configuring automatic purging for directories and files, see [Section 19.3, “Setting the Low and High Watermarks for Purging the File System Salvage,” on page 257](#).

## 19.2 Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool

If you delete a volume, NSS removes it from the pool. However, during a specified amount of time, called the Volume Purge Delay time, you can review and restore the volumes you have deleted. After the purge delay time elapses, NSS automatically removes volumes from the system and you can no longer access them. You can also manually purge any volumes you have deleted.

During the Purge Delay time (by default, 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 desired.

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. The default value for the Purge Delay is 345600 seconds, which is four days. Specify the time in seconds. See [Section 19.1, “Configuring the NSS Volume Salvage Area,” on page 255](#).

The Deleted Volumes option on the Pools page opens a separate Deleted Volumes page where you can purge or salvage the deleted volumes for the pool. This option is only available if the selected pool has deleted volumes on it.

---

**WARNING:** If you delete an entire pool, all the volumes are deleted with it. You cannot restore a deleted pool or any deleted volumes in it.

---

To manage the deleted volumes in a pool:

- 1 In iManager, click *Storage > Pools*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of pools appears in the Pools list.

- 3 Select the pool that you want to manage.

Wait for the page to refresh and display the details. The Deleted Volumes button is active only when there are deleted volumes in that list.

- 4 If the button is available, click *Deleted Volumes*.

This opens the Deleted Volumes page.

[Storage Management](#) ► [Pools](#) ► Deleted Volumes

Deleted Volumes

Purge

Pause Purge

Restart Purge

Salvage...

Close

PRO J46

PRO J135

SHVOL

TEMP30

TEMP60

TEMP90

Details:

Pool:

PRO J46

Volume was deleted:

Jun 9, 2003 4:19:07 PM

Scheduled Purge Time:

Jun 13, 2003 4:19:07 PM

Current deletion state:

Salvageable

- 5 Select a volume from the Deleted Volumes list.

The browser refreshes to display details in the Details area.

- **Pool:** The pool name.
- **Volume Was Deleted:** The time stamp when the volume was deleted.
- **Scheduled Purge Time:** The time that the Purge Delay expires for this deleted volume and the purging process is scheduled to begin.
- **Current Deletion State:** The deleted volume can be salvageable, purging, or paused.



6 Select one or more volumes, then perform one or more of these actions:

- **Purge:** Immediately begins the process of purging the selected volumes and their contents.
- **Pause Purge/Restart Purge:** *Pause Purge* stops the volume purging process for selected volumes. *Restart Purge* purges the selected volumes.
- **Salvage:** Begins the restore process for deleted volumes you selected.

If you salvage a volume, you need to assign a new name to that volume. A wizard opens to allow you to name the salvaged volume.

The volume salvage process can slow the server response time, so you might want to do this when there is less server activity.

- **Close:** Closes the Deleted Volumes page.

## 19.3 Setting the Low and High Watermarks for Purging the File System Salvage

Purging begins for the file system salvage area when the volume's low watermark is reached and continues until its high watermark is reached. The default low watermark is when free space is 10% or less of the total disk space. The default high watermark is when free space is at least 20% of the total disk space.

When free disk space falls below a low watermark, NSS begins autopurging the salvage area. When enough files are purged such that the free disk space increases to a percentage equal to or greater than the high watermark, the autopurge stops. Autopurging does not start again until free space again drops below the low watermark.

The high and low watermarks can be configured at a pool level. The low watermark's percentage can range from a minimum of 0% to a maximum of 98%. The high watermark's percentage can range from a minimum 2% to a maximum of 100%. The high and low watermarks must be at least 2% apart from each other.

To configure the high and low watermarks, enter the following at the server console:

```
nss /PoolHighWaterMark=<poolname>:Percent
```

```
nss /PoolLowWaterMark=<poolname>:Percent
```

For example, to set the low watermark to 5% and the high watermark to 10% for pool p\_users, enter

```
nss /PoolHighWaterMark=p_users:10
```

```
nss /PoolLowWaterMark=p_users:5
```

## 19.4 Salvaging or Purging Deleted Directories from the File System Salvage Area (OES SP1 NetWare, NetWare 6.5 SP4, or Later)

In Novell Remote Manager for OES SP1 NetWare, NetWare 6.5 SP4, or later, you can salvage or purge deleted directories for NSS volumes.

When your users delete files from the server, the files are retained and can be recovered (salvaged) until they have been purged. The server starts purging files if a volume starts to fill up; the oldest files that have been deleted are automatically purged. If you do not want files to be recovered after deletion, you can set the *Immediate Purge of Deleted Files* parameter value to *On*.

If your server is running out of disk space or you want to immediately purge files that have been deleted, you can easily do so on each volume from the *Volume Information* page.

### Purging All Deleted Files

To immediately purge all files that have been deleted from a volume:

---

**IMPORTANT:** Files that have been purged can no longer be recovered.

---

- 1 On the *Volume Management* page, click the *Information* icon to left of the volume name that you want to manage.
- 2 On the *Volume Information* page, click the *Purge Deleted Files* link.

### Purging Specific Directories or Files

To immediately purge specific files that have been deleted from a directory or all files from a specific directory:

---

**IMPORTANT:** Files that have been purged can no longer be recovered.

---

- 1 On the *Volume Management* page, click the volume name link where the file was deleted.
- 2 In the *Volume* directory list, click the *Information* icon to left of the directory name where the file was deleted.
- 3 On the *Directory Information* page, find the *Salvageable Files* field, then click the *Select for List* link.
- 4 Do one of the following:
  - Locate the file that you want to purge, and click the *Purge* button next the file name.
  - Click the *Purge All Files* button to purge all the files in the directory.
  - Click the *Purge* link for each directory that you want to purge.

### Salvaging a Deleted File

When your users delete files from the server, they are retained and can be recovered (salvaged) until they have been purged. If a user accidentally deletes a file and wants to recover it, use the following procedure before the file is purged.

- 1 On the *Volume Management* page, click the volume name link where the file was deleted.
- 2 From the directory list, click the *Information* icon to left of the directory name where the file was deleted.
- 3 On the *Directory Information* page, find the *Salvageable Files* field, then click the *Select for List* link.
- 4 Locate the file that you want to salvage, then click the *Salvage* button next to the file name.

## 19.5 Salvaging or Purging Deleted Files from the File System Salvage Area

Whenever users delete files from the server, the files are retained and can be recovered (salvaged) until they have been purged.

---

**IMPORTANT:** Files that have been purged can no longer be restored from salvage.

---

Deleted files are purged from the volume when one of the following events occurs:

- The Purge Delay setting times out. Deleted files are purged automatically.
- The server starts purging files if a volume starts to fill up; the oldest files that have been deleted are automatically purged.
- If you do not want a file to be restored after deletion, you can enable the Immediate Purge of Deleted Files parameter. The file is deleted and purged immediately.
- You manually purge the deleted files.

Use any of the following methods to salvage or purge deleted files.

- [Section 19.5.1, “Salvaging or Purging Deleted Files with NetStorage,” on page 259](#)
- [Section 19.5.2, “Salvaging or Purging Deleted Files with the Novell Client,” on page 259](#)
- [Section 19.5.3, “Salvaging or Purging Deleted Files with Novell Remote Manager for NetWare \(NetWare\),” on page 260](#)

### 19.5.1 Salvaging or Purging Deleted Files with NetStorage

Using NetStorage, the Admin user, equivalent Admin user, and users can purge and possibly undelete NSS files that were previously deleted on your Linux or NetWare® server.

**1** Access NetStorage.

For information, see [Section 4.4, “Novell NetStorage,” on page 58](#).

**2** In the left column, select the directory where the deleted files were.

**3** Click *View*, then click *Show Deleted Files*.

**4** Select the check box next to one or more files you want to undelete or purge.

**5** Click *File*, then click *Undelete* or click *Purge*.

### 19.5.2 Salvaging or Purging Deleted Files with the Novell Client

Using the Novell Client™ for Windows 2000/XP/2003, Admin users, equivalent Admin users, and users can purge and possibly undelete NSS files that were previously deleted on your Linux or NetWare server.

**1** Right-click the Novell Client icon (the red *N*) in the notification area to display the menu.

Use the NetWare utility even if the NSS volume resided on a Linux server.

- 2 If you want to salvage a deleted file, click *NetWare Utilities* > *Salvage*, browse to locate the directory where the deleted file resided, then do one of the following:
  - To restore one or multiple deleted files, select the deleted files, then click *Salvage File*.
  - To restore all deleted files in the directory, click *Salvage All*.

When you are done, click *OK*.

- 3 If you want to purge a deleted file, click *NetWare Utilities* > *Purge*, browse to locate the directory where the deleted file resided, then do one of the following:
  - To purge one or multiple deleted files, select the deleted files, then click *Purge File*.
  - To purge all deleted files in the directory, click *Purge All*.
  - To purge the directory's subdirectories and all deleted files in them, click *Purge Subdirectories*.
- 4 When you are done, click *OK*.

### 19.5.3 Salvaging or Purging Deleted Files with Novell Remote Manager for NetWare (NetWare)

Using Novell Remote Manager for NetWare, the Admin user or equivalent Admin user can purge and possibly undelete NSS files that were previously deleted on your NetWare server.

- 1 In a Web browser, connect to Novell Remote Manager on the NetWare server where the deleted file resides.

For information, see [Section 4.5.3, “Accessing Novell Remote Manager,” on page 60](#).

- 2 On the Volume Management page, click the volume name link where the file was deleted
- 3 From the directory list, click the *Information* icon to left of the directory name where the file was deleted.
- 4 On the *Directory Information* page, find the *Salvageable Files* field, then click the *Select for List* link.
- 5 Do one of the following:
  - **Salvage:** Locate the deleted file you want to restore, then click the **Salvage** button next to the filename.
  - **Purge:** Locate the deleted file you want to purge, then click the *Purge* button next to the filename.
  - **Purge All:** Click the *Purge All Files* button to purge all the files in the directory.

# Using Distributed File Services to Move and Split Volumes (NetWare)

# 20

For Novell® Storage Services™ on NetWare®, Novell Distributed File Services (DFS) allows you to develop a logical view of your storage solution that abstracts the underlying physical complexities from users.

DFS is not supported for NSS volumes on Linux, but DFS junctions can point to them. For information, see [Section 20.2.2, “Prerequisites for Pointing DFS Junctions at an NSS Volume on OES Linux,”](#) on page 264.

This section discusses the following topics:

- [Section 20.1, “Understanding DFS,”](#) on page 261
- [Section 20.2, “Prerequisites and Guidelines for DFS,”](#) on page 263
- [Section 20.3, “Creating a DFS Management Context for NCP,”](#) on page 266
- [Section 20.4, “Enabling DFS Junction Support for CIFS Clients,”](#) on page 266
- [Section 20.5, “Creating a DFS Junction Manually with ConsoleOne,”](#) on page 267
- [Section 20.6, “Viewing a DFS Junction,”](#) on page 268
- [Section 20.7, “Deleting a DFS Junction Manually with ConsoleOne,”](#) on page 268
- [Section 20.8, “Managing the Volume Location Database Service \(NetWare\),”](#) on page 269
- [Section 20.9, “Moving Volumes with iManager,”](#) on page 270
- [Section 20.10, “Splitting Volumes with iManager,”](#) on page 272
- [Section 20.11, “Managing Move or Split Jobs,”](#) on page 274

## 20.1 Understanding DFS

Novell Distributed File Services provides tools to help you modify the underlying physical organization to maximize the use and performance of available storage resources. With DFS, you can split a volume at any directory and move its contents to a new volume. It places a junction file in the directory’s location and uses a Volume Location Database (VLDB) to redirect queries automatically to the data’s new location, which preserves the logical file organization from the user perspective. DFS tools include the VLDB, DFS Junctions, and the Move Volume and Split Volume functions.

In OES NetWare, DFS comprises the following key components.

- [Section 20.1.1, “Volume Location Database,”](#) on page 262
- [Section 20.1.2, “DFS Junctions,”](#) on page 262
- [Section 20.1.3, “Move Volume Jobs,”](#) on page 262
- [Section 20.1.4, “Split Volume Jobs,”](#) on page 263
- [Section 20.1.5, “Original and New Volumes,”](#) on page 263

## 20.1.1 Volume Location Database

With DFS, you can create a Volume Location Database (VLDB) for NetWare 6.0 and later NSS volumes in the same Novell eDirectory™ tree. The VLDB maps the physical location of NetWare 6.0 and later NSS volumes on all servers in the tree. It is the VLDB service and database that allows you to create junctions. DFS junctions retrieve the current location of a volume from the VLDB in order to dynamically redirect queries to the physical location of the volume it represents.

## 20.1.2 DFS Junctions

A DFS junction is a logical placeholder for data stored on a NetWare 6.0 or later NSS volume. It is a virtual directory that points to the root of a destination volume. The junction itself can be located anywhere in the source volume. A junction functions like an alias in Novell eDirectory that appears to be a normal eDirectory object. To the user, a DFS junction appears to be any normal directory in your NetWare file system.

The servers of the source volume and destination volume must reside in a management context. That is, all volumes that participate in DFS, whether they contain junctions or are the target of junctions, must exist in some (though not necessarily the same) management context. Every destination volume must be located beneath a management context because the junction file just stores the DFS GUID of the volume that it points to, and the client looks up the physical location by presenting that GUID to the VLDB of the management context where the destination volume resides. The junction file also contains a hint about where to find the target volume's management context in the eDirectory tree, and the client walks up the tree from the specified hint until it hits the top (failure) or finds a management context.

For DFS junctions created manually with ConsoleOne®, the destination volume can be any NetWare 6.x NSS volume in the same eDirectory tree; it can be in any management context. For junctions created automatically by a Split Volume process, the destination volume can be any NetWare 6.x NSS volume in the same management context as the source volume.

Whenever you create a junction, the VLDB service assigns the junction a unique ID, which is automatically stored in its 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 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 needs 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 twofold: Not only does it decrease administration costs by allowing you to move a volume to a different server without making any announcements or needing to reeducate users, but it also simplifies the number of paths a user needs to remember if the data is spread among different volumes or servers.

## 20.1.3 Move Volume Jobs

With DFS, you can move a NetWare 6.x NSS volume and its contents to a new NSS volume faster than with a normal copy. You can move a crowded volume to a different pool that is has space available or is expandable. You can move volumes to different servers in the same context to balance associated traffic and workload across multiple servers. After the move, you must modify

scripts and mappings for the original volume to point to the new volume. Any DFS junctions that pointed to, or are in, the original volume are automatically updated in the VLDB with the new volume information.

## 20.1.4 Split Volume Jobs

With DFS, you can split a NetWare 6.x NSS volume at a specified directory and relocate the directory contents to a new volume on the same or different server anywhere in the same eDirectory management context. The new volume typically resides in a different pool. After a successful relocation of directory contents, DFS automatically creates a DFS junction at the split point, which replaces the directory. The DFS junction redirects queries to the new location. Users can continue to access their data on the new volume, without modifying the familiar logical paths.

## 20.1.5 Original and New Volumes

When you move or split a volume, the data is relocated from the original location to a new location and new volume. The new volume typically has a different name than the original volume; the names must be different if the new volume resides on the same server as the original volume.

Element	Description
Original location	The NetWare 6.0 or later server where the original NSS volume resides when the Move Volume or Split Volume job is initiated
Original volume	The source of the content for the new volume. For moves, it is the NetWare 6.0 or later NSS or Traditional volume that you want to relocate. For splits, it is the NetWare 6.0 or later NSS volume that you want to split.
New location	The destination NetWare 6.0 or later server where the new NSS volume resides
New volume	The NetWare 6.0 or later NSS volume that is created to contain the data from the original volume

## 20.2 Prerequisites and Guidelines for DFS

This section discusses the following prerequisites and guidelines for using DFS:

- [Section 20.2.1, “Prerequisites for Creating and Using DFS Junctions,” on page 263](#)
- [Section 20.2.2, “Prerequisites for Pointing DFS Junctions at an NSS Volume on OES Linux,” on page 264](#)
- [Section 20.2.3, “Guidelines for Creating DFS Junctions Manually,” on page 264](#)
- [Section 20.2.4, “Guidelines for Original Volumes and New Volumes,” on page 265](#)
- [Section 20.2.5, “Guidelines for Managing Move Volume or Split Volume Jobs,” on page 265](#)
- [Section 20.2.6, “Guidelines for Files in Nearline Storage,” on page 265](#)

### 20.2.1 Prerequisites for Creating and Using DFS Junctions

- In OES, DFS is available only on NSS volumes on NetWare.

- You can create DFS junctions only in NSS volumes that reside on your NetWare 6.0 or later server that is part of a DFS management context.
- If you move the NSS volume that contains junctions to a Linux server, the junctions do not work.
- Volumes must have an eDirectory Volume object to be usable with DFS. For example, if you create a DFS junction manually in ConsoleOne, you specify the destination volume for the data by selecting its eDirectory Volume object. Typically, whenever you create an NSS volume, NSS creates a corresponding eDirectory object. It would be unusual not to have an eDirectory Volume Object for your volume.

To update the eDirectory Volume object, see [Section 13.5, “Updating eDirectory Volume Objects,” on page 188](#).

- You must create a DFS Management context before you can use DFS. All volumes that participate in the junctions, whether they contain junctions or are targets of junctions, must be located on servers in a management context.

For information, see [Section 20.3, “Creating a DFS Management Context for NCP,” on page 266](#).

- If you are using CIFS, you must enable DFS support for CIFS clients.

For information, see [Section 20.4, “Enabling DFS Junction Support for CIFS Clients,” on page 266](#).

- If you are using CIFS, you might need to make sure CIFS users use a Windows logon username that is the same as their eDirectory username. On some earlier builds of Windows 2000/XP, Windows passes the local logon username to the DFS junctions rather than the eDirectory login username.

## 20.2.2 Prerequisites for Pointing DFS Junctions at an NSS Volume on OES Linux

You can point a DFS junction in a source NSS volume on NetWare to a target NSS volume on Linux, if all of the following conditions exist:

- The volumes are in the same Volume Location Database (VLDB) context.

After you create an NSS volume on Linux, make sure to run VLDB Repair on your VLDB server to update the VLDB database to include the volume. For information, see [Section 20.8, “Managing the Volume Location Database Service \(NetWare\),” on page 269](#).

- The target NSS volume on Linux has a volume object in eDirectory.
- The source and target servers share a network protocol in common, such as CIFS/Samba or NCP.

## 20.2.3 Guidelines for Creating DFS Junctions Manually

For DFS junctions you create manually, the following additional criteria must be met:

- The original and destination locations must be in the same Novell eDirectory tree.
- The destination volume must have an object defined in eDirectory.
- The destination volume’s server and the DFS junction’s server must share a networking protocol common to both, such as NCP™ or CIFS.



## 20.2.4 Guidelines for Original Volumes and New Volumes

Consider the following guidelines for original volumes and new volumes:

- You can split only NetWare 6.0 or later NSS volumes.
- You can move NetWare 6.0 or later NSS volumes or Traditional volumes.
- For both moves and splits, the destination volume must be an NSS volume on a NetWare 6.0 or later server.
- The destination volume can reside in the same or different pool than the original NSS volume.
- The destination volume can reside on the same or different server than the original NSS volume.
- The original NSS volume and the destination NSS volume must be in the same Novell eDirectory DFS Management Context.
- The access rights for the original and destination locations must be the same. The access rights for the original location are transferred automatically to the new volume. Make sure you have the appropriate rights on the destination server.
- The destination volume should be configured with the same attributes as the original volume. Review the attribute settings for the original volume before you begin a move or split job to be sure you know which attributes to set.
- The original volume and new volume must reside on servers configured to share the same networking protocol, such as NCP or CIFS.

## 20.2.5 Guidelines for Managing Move Volume or Split Volume Jobs

Consider the following guidelines for managing your move or split jobs:

- If the server crashes during a volume move or split operation, the operation resumes where it left off when the server comes back up.
- You can set up as many move or split jobs as you want; however, DFS can run only four active move or split operations at one time. After four requests, any additional requests must wait for one of the running operations to complete, or you must pause a job and start the operation you want to run.
- After you move a volume, you must modify mappings and scripts to point to the new physical location of the moved volume. If the volume contained DFS junctions, the junctions continue to work normally.
- When you split a volume, DFS automatically creates a DFS junction at the directory point where you split the volume. If the directory or its subdirectories contained DFS junctions, the junctions continue to work normally.

## 20.2.6 Guidelines for Files in Nearline Storage

If the Data Migration attribute is set for the NSS volume you want to move or split, some of its data might reside on tertiary storage media as nearline storage. The Move and Split processes demigrate the files from the tertiary storage and copies them to the new location. Make sure the original volume and new volume have enough space to incorporate both inline and nearline files during the demigration.

## 20.3 Creating a DFS Management Context for NCP

Before you create a DFS junction, you must create at least one DFS Management Context at an O or OU level in the Novell eDirectory tree where you want to create the junction. You can create more than one Management Context if you have a geographically diverse company. In this way, each geographic area can manage and control the junctions and the VLDB service within its own domain.

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

- 1 On the server that you want to make a VLDB server, open ConsoleOne and log in with the administrator username.
- 2 Navigate to the eDirectory container where you want to enable DFS support.
- 3 Right-click an O-level or an OU-level context in the eDirectory tree, then click **New > DFS Management Context**.
- 4 Select the server from the Available list, click the right-arrow to move it into the Selected list, then click *Next*.

The server you select runs the VLDB service and houses the database.

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

The default location is the `sys:\etc` directory.

After you create the DFS Management Context, DFS creates attributes for volume objects in the eDirectory container object for the specified server, and it starts DFS on the server you selected.

- 6 (Conditional) For a larger eDirectory tree, with servers located in multiple eDirectory containers, manually run `VLDB Repair` from the server console because a username and password are required in this situation.

---

**IMPORTANT:** You must start the VLDB Repair command from the server console immediately after creating the DFS Management Context; do not start it from ConsoleOne.

---

Logging in from the server console is necessary to enter the The VLDB Repair adds a GUID to every volume object in the eDirectory containers for other servers in the tree and enables the DFS Junction function for all the NetWare 6.x servers in the tree.

- 7 (Optional) You can specify a second server in the same LAN as a VLDB server, which allows you to maintain a copy of the database on another server to prevent data loss or corruption.
  - 7a On the second server, open ConsoleOne.
  - 7b Open the server's Property page.
  - 7c Click *Add Replica*.

## 20.4 Enabling DFS Junction Support for CIFS Clients

If you have clients that use CIFS networking protocol instead of a NetWare client, you must enable DFS Junction support for CIFS clients on all servers where you created DFS junctions. You must

enable CIFS protocols for the destination servers, too. DFS Junction support for clients is available on NetWare 6.5 or later servers.

If new volumes are added within the DFS Management Context container after you enable DFS Junction support for CIFS clients, you might need to run VLDB Repair to update the VLDB database, depending on the utility that was used to create the new volume.

To enable DFS Junction support for CIFS clients on NetWare 6.5 or later servers:

- 1** On your NetWare 6.5 or later server, edit the `sys:etc\cifsctxs.cfg` file.  
Add all CIFS/DFS users contexts to the file. CIFS uses a flat tree structure.
- 2** Enable DFS Junction support for CIFS for each server:
  - 2a** In iManager, click File Protocols.  
For instructions, see [Section 4.1.3, “Accessing Novell iManager,” on page 45](#).
  - 2b** Select a server to manage.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).
  - 2c** Select the *Enable DFS* check box, then click *Apply*.
  - 2d** Repeat [Step 2b](#) and [Step 2c](#) for each of the servers.
- 3** If you have not created a DFS Management Context, do it now.  
For information, see [Section 20.3, “Creating a DFS Management Context for NCP,” on page 266](#).
- 4** In ConsoleOne, set the user Simple Password.  
If Universal Password is enabled for the user container, it automatically synchronizes eDirectory and Simple Password.

---

**IMPORTANT:** Make sure the Windows logon user and password match the Simple Password’s username and password. DFS only passes the default Windows logon username and password.

---
- 5** In ConsoleOne, create junctions, if desired.  
For information, see [Section 20.5, “Creating a DFS Junction Manually with ConsoleOne,” on page 267](#).

## 20.5 Creating a DFS Junction Manually with ConsoleOne

You can create DFS junctions manually with ConsoleOne to impose a logical overlay for your storage structure. When you use Novell iManager to split volumes, the process creates DFS junctions automatically to replace the directory where you split the volume.

- 1** (Conditional) If it does not already exist, you must create a DFS Management context and make sure the VLDB service is up and running. For information, see [Section 20.3, “Creating a DFS Management Context for NCP,” on page 266](#).
- 2** In ConsoleOne, create a subdirectory or use an existing subdirectory of any NetWare 6.0 or later NSS volume.
- 3** Right-click the subdirectory, 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, then click the up-arrow in the ConsoleOne toolbar.

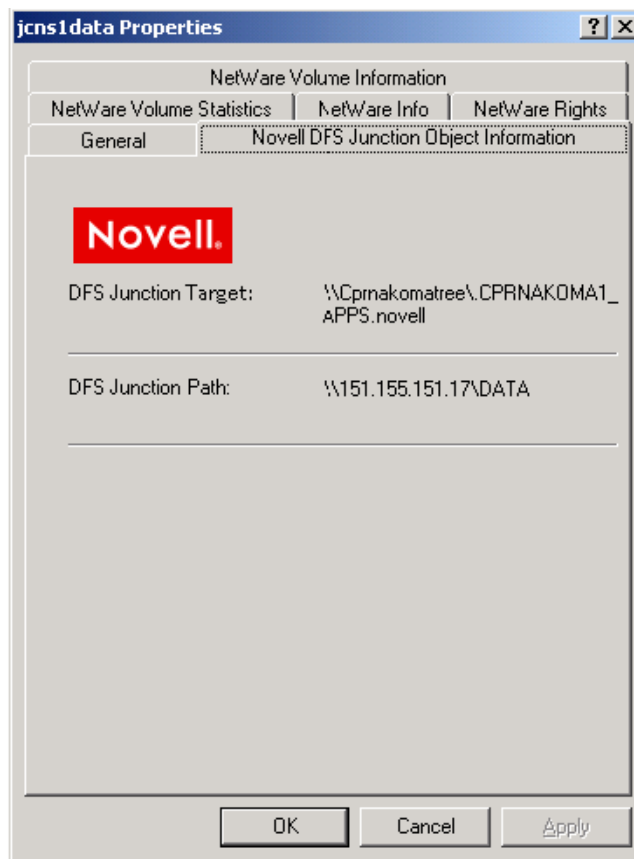
- 4 Verify that the eDirectory tree name is correct, browse to the volume that you are pointing to from the junction, then click *Next*.
- 5 Name the junction, then click *Finish*.
- 6 To view the newly created junction, from the Console One toolbar, click *View > Refresh*.

## 20.6 Viewing a DFS Junction

You can access DFS junctions with the latest Novell Client™ via the NCP protocol, Web-based services via Novell NetStorage, and the Microsoft\* Client via the CIFS protocol.

- 1 In a file browser, right-click the directory, then select Properties.
- 2 Click the Junction Properties tab to view the junction's property page.

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



## 20.7 Deleting a DFS Junction Manually with ConsoleOne

To delete a junction:

- 1 In ConsoleOne, locate and right-click the junction.

2 Select *Delete Shortcut > Junction*.

## 20.8 Managing the Volume Location Database Service (NetWare)

VLDB services run only on NetWare servers. You can manage the VLDB service from the server console or using the VLDB management service.

### 20.8.1 VLDB Commands

To view a listing of the available commands, enter at the server console:

```
vldb help
```

To repair the VLDB, enter the following at the server console of the primary VLDB server:

```
vldb repair
```

The VLDB Repair adds a GUID entry for all volumes in the eDirectory tree to the VLDB.

The VLDB contains GUID for NSS volumes on both Linux and NetWare. You can create DFS junctions on NetWare servers that point to the volumes on the Linux machine. You cannot create a DFS junction on an NSS volume on Linux. If you move a device from NetWare to Linux that contains an NSS volume with a DFS junction in it, you cannot follow the existing junction while the volume is on a Linux server.

### 20.8.2 VLDB Management Service in ConsoleOne

To access the VLDB management service:

- 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 tasks:

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

To add or remove replicas, 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, then click *Properties > Supported Services > NSS VLDB*. After you are in the management service, click either *Add Replica* or *Remove Replica*.

- Repair a database in one or more of 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.

## 20.9 Moving Volumes with iManager

When you move a volume, DFS moves the volume's file structure, data, and user rights information from the original location to a new volume in the network.

---

**IMPORTANT:** You cannot move the `sys :` volume.

---

- 1 Make sure you meet the [Section 20.2, “Prerequisites and Guidelines for DFS,”](#) on page 263.
- 2 You must create at least one DFS Management Context at an O or OU class level in the Novell eDirectory tree. Both the original volume location and the new volume location must be in the same DFS Management Context. You must create at least one Volume Location Database (VLDB) server in this management context.

For information, see [Section 20.3, “Creating a DFS Management Context for NCP,”](#) on page 266.

- 3 (Optional) If you plan to temporarily retain the deleted original volume in the salvage area after a successful volume move, verify that your server's *Purge Delay* setting meets your needs. The default *Purge Delay* setting is 4 days (345600 seconds), but the actual setting might differ. If desired, modify the *Purge Delay* setting.

At the server console, enter

```
nss /logicalVolumePurgeDelay=value
```

where *value* is the number of seconds until a deleted volume is purged from the salvage area. For information, see [Section 19.1, “Configuring the NSS Volume Salvage Area,”](#) on page 255.

- 4 Make sure users and applications are not accessing files on the volume throughout the move or split process.
- 5 In iManager, click *Storage > Volumes*.  
For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on page 46.
- 6 Select a server to manage.  
For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on page 46.
- 7 From the Volumes list, select the volume that you want to move.
- 8 Click *Properties > Attributes* to view the current attribute settings for the volume.  
Remember this information for [Step 13](#), where you specify the attributes for the destination volume.

- 9 From the Volumes list, select the volume, then click *Move* to open the Move Volume wizard.

The screenshot shows the 'Move Volume' wizard dialog box. The title bar says 'Move Volume'. Below the title bar is a tabbed interface with the first tab selected, labeled 'Enter location, start date and time'. The dialog contains the following fields and controls:

- 'Original Volume:' with the text 'PROJ135'.
- 'New Location:' with a text box containing 'PROJSVR12.yourcontext' and a browse button (folder icon).
- Two radio buttons for scheduling:
  - 'Start Now' (selected).
  - 'Select Start Date and Time'.
- If 'Select Start Date and Time' is chosen, there are fields for 'Day:' (containing 'March 24, 2003') and 'Time:' (containing '20:00 (8:00 PM)').
- At the bottom are three buttons: '<< Back', 'Next >>', and 'Cancel'.

- 10 Specify the following parameters, then click *Next*.

- **New Location:** Specify the eDirectory common name of the server where you want to move the selected volume.
- **Schedule:** Select *Start Now* to begin the move immediately, or specify the date and time you want to schedule the move.

Make sure that the volume is active at the time that the move is to begin.

- **Purge Immediately:** Select *Purge Immediately* to purge the original volume's deleted data from the server's salvage area immediately following the successful completion of a volume move job.

If Purge Immediately is disabled (deselected), DFS delays the purge of the volume's deleted data according to the server's *Purge Delay* setting and the normal purge process.

You can manually purge or restore the original volume's deleted data at any time during the purge delay period, using the normal method of purging deleted volumes. For information, see [Section 19.2, "Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,"](#) on page 255.

- 11 Specify a unique name for the new volume.

Typically, the new volume's name is different from the original volume's name. The new volume's name must meet the uniqueness requirements in the new location and conform to volume naming conventions. For guidelines about naming volumes, see [Section 5.5, "Naming NSS Storage Objects,"](#) on page 69. It is possible, though not necessarily advisable, to use the same name under some conditions, such as when the original location and the new location are different servers.

---

**IMPORTANT:** If the name you provide is not unique, you receive an error message. You must click *Cancel* to back out of the wizard, then begin the move process again.

---

- 12 Specify the pool on the new location where you want the new volume to reside, specify the volume quota, then click *Next*.

Only pools that have free space appear in the list. You can select an existing pool from the list or create a new pool. If you create a new pool, the wizard guides you through steps similar to the process for creating a new pool. For instructions, see [Section 11.1, “Creating Pools with iManager,” on page 147](#).

If no pools are listed, there is no space available to create a volume in the new location. Cancel the wizard, add more devices to the server and expand the desired pool, or free up space by deleting existing pools, then return to the Volume Management page to begin the move process from the beginning.

- 13** Specify the attributes for the new volume you are creating, based on the volume attributes of the original volume that you viewed in [Step 8](#).

The wizard displays only the default settings for volume attributes. The attribute settings might not be the same as those set for the volume you are moving. Settings, such as Compression, that cannot be changed after they are set for a volume must be the same on the original volume and the new volume.

- 14** Click *Finish*.

The move can take a few minutes to several hours, depending on how much data needs to be moved.

- 15** To view the job’s status or to pause and resume the job, click *Volume Job Control*.

For information, see [Section 20.11, “Managing Move or Split Jobs,” on page 274](#).

- 16** (Optional) After the job completes successfully, if you disabled the *Purge Volume Immediately* option, you can manually purge or restore the original volume’s deleted data at any time during the purge delay period, using the normal method of purging deleted volumes.

For information, see [Section 19.2, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 255](#).

- 17** Manually update script files, configuration files, or mappings by modifying the location of the original volume to the location of the new volume.

## 20.10 Splitting Volumes with iManager

When you split a volume, NSS creates a DFS junction to serve as a placeholder for that volume and its contents. This might involve any directory level of the file system structure, from the volume’s root directory to its most distant subdirectory. Everything below the specified directory, including its subdirectories and the data therein, moves from the original location to a new volume in the network.

- 1** Make sure you meet the [Section 20.2, “Prerequisites and Guidelines for DFS,” on page 263](#).
- 2** You must create at least one DFS Management Context at an O or OU class level in the Novell eDirectory tree. Both the original volume location and the new volume location must be in the same DFS Management Context. You must create at least one Volume Location Database (VLDB) server in this management context.

For information, see [Section 20.3, “Creating a DFS Management Context for NCP,” on page 266](#).

- 3** Make sure users and applications are not accessing files on the volume throughout the move or split process.
- 4** In iManager, *Storage > Volumes*.



For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,”](#) on [page 46](#).

**5** Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,”](#) on [page 46](#).

**6** From the Volumes list, select a volume that you want to split.

**7** Click *Properties > Attributes* to view the current attribute settings for the volume.

Remember this information for [Step 12](#), where you specify the attributes for the destination volume.

**8** Click *Split* to open the Split Volume wizard.

The screenshot shows the 'Split Volume' wizard dialog box. The title bar says 'Split Volume'. Below the title bar is a tabbed interface with the first tab selected, labeled 'Enter location, split path, start date and time'. The dialog contains the following fields and options:

- Original Volume: **PROJ135**
- New Location:  (with browse and help icons)
- Split volume at:
- Start Now (selected radio button)
- Select Start Date and Time (unselected radio button)
  - Day:  (with calendar icon)
  - Time:  (with time selection icons)

At the bottom are three buttons: '<< Back', 'Next >>', and 'Cancel'.

**9** Specify the following parameters, then click *Next*.

- **New Location:** Specify the eDirectory common name of the server where you want to move the selected volume.
- **Split Volume At:** Specify the directory in the selected volume where you want the DFS junction to occur. All data below that point moves to the new volume created at the new location.

---

**IMPORTANT:** Do not include a leading slash in the path.

---

- **Schedule:** Select *Start Now* to begin the move immediately, or specify the date and time you want to schedule the move.

---

**WARNING:** Make sure that the volume is active but not in use by users or APIs at the time that the move is to begin.

---

**10** Create and name the new volume.

Typically, the new volume’s name is different from the original volume’s name. The new volume’s name must meet the uniqueness requirements in the new location and conform to volume naming conventions. For guidelines about naming volumes, see [Section 5.5, “Naming NSS Storage Objects,”](#) on [page 69](#). It is possible, though not necessarily advisable, to use the

same name under some conditions, such as when the Original Location and the New Location are different servers.

---

**IMPORTANT:** If the name you provide is not unique, you receive an error message. You must click *Cancel* to back out of the wizard, then begin the split process again.

---

- 11** Specify the pool on the new location where you want the new volume to reside, specify the volume quota, then click *Next*.

Only pools that have free space appear in the list. You can select an existing pool from the list or create a new pool. If you create a new pool, the wizard guides you through steps similar to the process for creating a new pool. For instructions, see [Section 11.1, “Creating Pools with iManager,” on page 147](#).

If no pools are listed, there is no space available to create a volume in the New Location. Cancel the wizard, add more devices to the server and expand the desired pool, or free up space by deleting existing pools, then return to the Volume Management page to begin the move process from the beginning.

- 12** Specify the attributes for the new volume, based on the volume attributes of the original volume that you viewed in [Step 7](#).

The wizard displays only the default settings for volume attributes. The attribute settings might not be the same as those set for the volume you are moving. Settings, such as Compression, that cannot be changed after they are set for a volume must be the same on the original volume and the new volume.

- 13** Click *Finish*.

The split can take a few minutes to several hours, depending on how much data needs to be relocated.

- 14** To view the job’s status or to pause and resume the job, click *Volume Job Control*.

For information, see [Section 20.11, “Managing Move or Split Jobs,” on page 274](#).

## 20.11 Managing Move or Split Jobs

After you initiate a Move job or a Split job using the Volume Management page in iManager, use the Move/Split Job Control page to view its status. The report displays the status of all the active jobs and recently completed jobs that were initiated from the selected server. You can pause, resume, reschedule, finish, or delete a process, depending on the state it is in.

From the Move/Split Job Control Management page, you can perform the following tasks:

- [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#)
- [Section 20.11.2, “Pausing a Move or Split Job,” on page 276](#)
- [Section 20.11.3, “Resuming a Move or Split Job,” on page 277](#)
- [Section 20.11.4, “Rescheduling a Move or Split Job,” on page 277](#)
- [Section 20.11.5, “Viewing Files Skipped by a Move or Split Job,” on page 278](#)
- [Section 20.11.6, “Deleting a Move or Split Job,” on page 278](#)
- [Section 20.11.7, “Finishing a Move or Split Job,” on page 279](#)

## 20.11.1 Viewing the Status of Move or Split Jobs

The Move/Split Job Control Status Report shows the following information:

- **Type:** Distinguishes the job as a Move or Split job.
- **Name:** If it is a Move, this is the name of the original volume. If it is a Split, it is the name of the original volume plus the path to the original directory location where you are splitting the volume.
- **Percent Complete:** The estimated percentage of data to be copied from the original volume that has been copied to the new volume as of the instant the status report was created.

To update the status, click *Move/Split Job Control* in *Roles and Tasks*.

- **State:** States are defined in logical tasks so that the Move or Split job can start and stop at several points in the process and go back to repeat any subprocesses, as needed.

Possible states include the following:

- Scheduled
  - Running
  - Files Skipped (administrator action required)
  - Failed
  - Deleted
  - Completed
- **Comment:** A comment you typed when you issued a *Pause* command.

To access a report of move and split jobs:

- 1 In iManager, click *Storage > Move/Split Job Control*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

A list of move and split jobs appears.

[Storage Management](#) ► **Move / Split Job Control**

### Move / Split Job Control Management

To move and split volumes, use the Move and Split functions on the Volume Management page. The Move/Split Job Control page reports the status of all the active move and split jobs initiated from this server. You can pause/resume, reschedule, finish, or delete a job, depending on the state it is in.

Server:



<input type="checkbox"/>	Type ▾	Name ▾	Percent Complete ▾	State ▾	Time ▾	Comment
<input checked="" type="checkbox"/>	Move	MOVE_FROM:	0	Scheduled	7/30/03 8:00 PM	
<input type="checkbox"/>	Split	SPLIT_FROM:install\javabin	0	Paused	7/30/03 6:30 PM	pause it again

- 3 (Optional) Click the Sort arrow next to the column heading that you want to sort jobs by.
- 4 When you are done, click *OK*.

## 20.11.2 Pausing a Move or Split Job

Pause suspends one or more Move/Split jobs until you manually resume or delete them. Only four combined Move or Split jobs can be running on a given server at a time. The Volume Manager performs the jobs in the order they are scheduled. If four operations are in progress and you want to activate others that you decide are a higher priority, you can pause one or more of the active jobs, thus allowing the new job to run immediately.

You might want to pause a Move or Split job when users need access to the content in the original volume at times when Move/Split jobs are normally running.

- 1 In iManager, open the Move/Split Job Control Management page to view a job status report.  
For instructions, see [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#).
- 2 Select the *Job* check box next to one or more active jobs that you want to pause.
- 3 Click *Pause*.
- 4 Type a Comment to be displayed in the status report, such as the reason you are pausing the job.
- 5 Click *OK*.

After the page refreshes, the jobs report their status as Pausing or Paused. If you inadvertently selected some jobs that are not eligible to be rescheduled, iManager ignores the Pause command for those jobs. It applies only to those jobs that the command is valid for.

### 20.11.3 Resuming a Move or Split Job

Resume continues one or more paused Move/Split jobs so that they can continue from wherever they were in the Move or Split process when you paused them.

- 1 In iManager, open the Move/Split Job Control Management page to view a job status report.

For instructions, see [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#).

- 2 Select the Job check box next to one or more paused jobs that you want to resume.
- 3 Click *Resume*.
- 4 Click *OK*.

After the page refreshes, the jobs report their status as Scheduled or Running, depending on when the job was originally scheduled to run. If you inadvertently selected some jobs that are not eligible to be rescheduled, iManager ignores the Resume command for those jobs. It applies only to those jobs that the command is valid for.

You cannot resume a completed, failed, or deleted job.

### 20.11.4 Rescheduling a Move or Split Job

Reschedule changes the date and time that the selected jobs should run. It applies the same date and time to all of the selected jobs.

---

**IMPORTANT:** You cannot reschedule a completed, failed, or deleted job.


---

- 1 In iManager, open the Move/Split Job Control Management page to view a job status report.

For instructions, see [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#).

- 2 Select the *Job* check box next to one or more uncompleted jobs that you want to reschedule.
- 3 Click Reschedule.

This opens a Reschedule Jobs dialog box.



- 4 To set the new schedule:
  - To start the job immediately, select *Start Now*.
  - To start the job at a future date or time, specify the start date and time when you next want the move or split job to begin.
- 5 Click *OK*.

After the page refreshes, the jobs report their status with the new schedule. If you inadvertently selected some jobs that are not eligible to be rescheduled, iManager ignores the Reschedule option for those jobs. It applies only to those jobs for which the command is valid.

## 20.11.5 Viewing Files Skipped by a Move or Split Job

If a Move or Split Volume job reports a Files Skipped status, some files could not be moved because they were open at the time that DFS attempted to copy them.

To view a list of files:

- 1 To retrieve the Operation ID number of the Move or Split Volume job, enter the following command at the server command prompt:

```
volmn status
```

NSS displays a list of the Move and Split jobs that are in progress or were initiated or run in the past seven days.

- 2 Find the Move or Split job of interest and make a note of its Operation ID (OpID) number.  
For example, Move(Vol1) might have an Operation ID number of 104211375.
- 3 To display a list of the files skipped, enter the following at the server command prompt:

```
volmn list OpID
```

For example, enter

```
volmn list 104211375
```

## 20.11.6 Deleting a Move or Split Job

Delete cancels a selected job before it begins or up to a certain point in the Move or Split process. If the process is beyond a certain state, it returns an error message to prevent you from deleting the process.

If the deletion continues, it stops the job at the next convenient step in the process, but performs no cleanup. The target volume exists and contains all files copied to it before the you deleted the job. Deleted jobs appear in the status report for a week by default.

---

**IMPORTANT:** If deletion continues, the original data is still intact. There is no loss of service or need to copy data from the destination volume back to the original volume.

---

- 1 In iManager, open the Move/Split Job Control Management page to view a job status report.  
For instructions, see [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#).
- 2 Select the *Job* check box next to one or more scheduled or paused jobs that you want to cancel.
- 3 Click *Delete*.

After the page refreshes, the selected jobs report their status as Deleted. Some jobs might not allow themselves to be deleted, depending on how far into the move or split process they are when you click delete.

If you inadvertently selected some jobs that are not eligible to be deleted, iManager ignores the Delete command for those jobs. It applies only to those jobs that the command is valid for. You cannot delete a completed, failed, or previously deleted job.

## 20.11.7 Finishing a Move or Split Job

Finish continues a Move or Split job that has reached a state that requires manual intervention to finish even though some files might remain behind.

- 1 In iManager, open the Move/Split Job Control Management page to view a job status report.

For instructions, see [Section 20.11.1, “Viewing the Status of Move or Split Jobs,” on page 275](#).

- 2 Select the *Job* check box next to one or more jobs waiting for administrator intervention that you want to complete, acknowledging and understanding the exceptions noted.
- 3 Click *Finish*.

After the page refreshes, the selected jobs report their status as Scheduled, Running, or Complete, depending on how much work was left to do in the job.

If you inadvertently selected some jobs that are not eligible to be finished, iManager ignores the *Finish* command for those jobs. It applies only to those jobs for which the command is valid.





Novell® Storage Services™ supports zero or more hard links for files. This section discusses the following:

- [Section 21.1, “Understanding Hard Links,” on page 281](#)
- [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 285](#)
- [Section 21.3, “Creating a Hard Link Using Ln on a Linux Client,” on page 286](#)
- [Section 21.4, “Creating a Hard Link Using a zLink API,” on page 287](#)
- [Section 21.5, “Creating a Hard Link for Testing Purposes \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 287](#)
- [Section 21.6, “Viewing Hard Links for a Volume \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 288](#)
- [Section 21.7, “Deleting a Hard Link,” on page 288](#)

## 21.1 Understanding Hard Links

Hard link support for an NSS volume allows users to create multiple names for a single, existing file object in the same directory or in multiple directories in the same NSS volume.

---

**IMPORTANT:** Hard links to directories, data streams, and extended attributes are not allowed.

---

The alternate names for the file object link to the primary file inode in the NSS file system. One file can have different filenames in the same directory or multiple directories as long as all of the directories reside on the same volume. It is not possible to create hard links from different volumes and have them point to the same file.

This section discusses the following:

- [Section 21.1.1, “Hard Links and the Primary Link,” on page 282](#)
- [Section 21.1.2, “Hard Links and Linux Users,” on page 282](#)
- [Section 21.1.3, “Hard Links and File Ownership,” on page 282](#)
- [Section 21.1.4, “Hard Links and File System Trustees,” on page 283](#)
- [Section 21.1.5, “Hard Links and Directory Space Quotas,” on page 283](#)
- [Section 21.1.6, “Hard Links and User Space Quotas,” on page 283](#)
- [Section 21.1.7, “Hard Links and the Hard Link Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 284](#)
- [Section 21.1.8, “Hard Links and File Salvage and Purge,” on page 284](#)
- [Section 21.1.9, “New Metadata Structure Supports Up to 65,535 Hard Links for a File \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 284](#)
- [Section 21.1.10, “Old Metadata Structure Supports Limited Hard Links for a File,” on page 285](#)

### 21.1.1 Hard Links and the Primary Link

Every file has what is considered to be the primary link, which is originally the path and filename assigned at file creation time. The primary parent directory is originally the path and directory name where the file is created. When you create hard links for the file, the file's primary link and all of its hard links share the file content by pointing to the file's inode. All links share all of the file's metadata except for name information; each link accesses only its own filename information.

After you create hard links for a file, you can delete the file's primary link or its hard links, but the file's content and metadata are not deleted until the last remaining link is deleted. If you delete a hard link, the link name is deleted, not the file. If hard links exist, whenever you delete a primary link, the most recently created hard link automatically becomes the new primary link.

Whenever a primary link is reassigned, the following occurs:

- The file owner does not change.
- The explicit file system trustee assignments, trustee rights, and file attributes for the file do not change.
- The file system inherited rights filter for the new primary parent directory is used to calculate effective rights for the file. The inherited rights filter for the old primary parent directory is no longer used. For information, see [Section 21.1.4, "Hard Links and File System Trustees," on page 283](#).
- For the directory space quota, the total disk space used by the file is reassigned from the old primary parent directory to the new primary parent directory. For information, see [Section 21.1.5, "Hard Links and Directory Space Quotas," on page 283](#).
- For the user space quota, the total disk space used by the file continues to be associated with the file's assigned owner.

### 21.1.2 Hard Links and Linux Users

When hard links are used on NSS volumes on Linux, users must be Linux-enabled eDirectory users. For information, see [Section 6.4.5, "NSS and Linux User Management," on page 77](#).

### 21.1.3 Hard Links and File Ownership

In the NSS file system, the user who creates a file is assigned as its owner at file creation time. File ownership is not affected by transferring the primary link, creating hard links, or deleting hard links. A file can have only one owner, even if it has multiple hard links.

File ownership rarely, if ever, changes, but an administrator or administrator equivalent user can assign a new owner when necessary. Changing file ownership requires the Supervisor right for the primary parent directory and the file. Use whatever tools you normally use to modify the file's ownership.

NSS uses the Novell trustee model to secure access to the file, not file ownership. For information about how access control is affected by hard links, see [Section 21.1.4, "Hard Links and File System Trustees," on page 283](#).

File ownership allows accounting of disk block usage per user and enables user space quotas to operate effectively. For information, see [Section 21.1.6, "Hard Links and User Space Quotas," on page 283](#).

File ownership has no affect on the directory space quotas. For information, see [Section 21.1.5, “Hard Links and Directory Space Quotas,” on page 283](#).

## 21.1.4 Hard Links and File System Trustees

Explicit file system trustee assignments, trustee rights, and attributes for a file are stored in the file’s metadata where the information is shared equally by the primary link and all hard links. For users who match the trustee, the file is visible along any path of the file’s primary link and hard links. Explicit access security is enforced equally for any of these paths.

Inherited trustee assignments and trustee rights are inherited only through the primary parent directory. When calculating effective rights to a file, rights are granted based on explicit rights and rights inherited from the primary parent directory path.

If the primary link is deleted and reassigned, the inheritance changes to the settings of the newly assigned primary parent directory.

---

**IMPORTANT:** Make sure you verify the inherited rights on the new primary link whenever you delete a primary link.

---

## 21.1.5 Hard Links and Directory Space Quotas

If you set a directory space quota for a file’s primary parent directory, the file’s size is charged against the quota. All metadata increases associated with hard links are included in the file’s total disk space usage. For hard links in the primary parent directory, the file has an alternate name, but the charge for space usage is not duplicated. For hard links in other directories, quotas for their parent directories are not charged for space usage.

Whenever you delete a primary link, the most recently created hard link automatically becomes the new primary link. At that time, the file’s size is subtracted from the space usage reported on old primary parent directory and added to the space usage reported on the newly assigned primary parent directory. The file’s space usage is no longer charged against the directory space quota on the old primary parent directory. If a quota is set on the new primary parent directory, the file’s space usage is charged against it.

When you delete a primary link, the transfer of responsibility for space usage to the new primary parent directory is allowed even if the added file size overbooks its directory space quota. NSS allows the quota to be overbooked, and then enforces the directory quota restrictions.

---

**IMPORTANT:** Make sure you verify the directory quota on the new primary parent directory whenever you delete a primary link.

---

## 21.1.6 Hard Links and User Space Quotas

User space restrictions are charged based on ownership of the file. Ownership is tracked with the user’s GUID, not the username. If the username ever becomes invalid, the file continues to be charged to the GUID. In a space usage report, the value of the GUID appears in place of where the a valid username would normally be.

If other users create hard links to the file, they are not charged a quota on that file. The file size is charged to the file owner’s user space quota, even if the owner no longer has the necessary rights to

access the file. If hard links exist, deleting the primary link does not delete the file, and the owner continues to be responsible for the space used.

The file's users can delete the primary link and hard links to the file only in directories where they have the necessary access rights. Hard links from directories the owner cannot access causes the file to be retained, and the owner continues to be charged for its quota. The file continues to be charged against the owner's user space quota until the file itself is deleted or until the system administrator reassigns another user as the file owner.

When file ownership is changed, the new owner's user space quota is checked before the change of ownership is allowed. If the file size will overbook the user space quota, the change is not allowed.

---

**IMPORTANT:** Make sure you verify the user space quota on the new owner whenever you change ownership of a file.

---

### **21.1.7 Hard Links and the Hard Link Attribute (OES SP1 NetWare, NetWare 6.5 SP4)**

Beginning with OES SP1 NetWare and NetWare 6.5 SP4, NSS volumes on NetWare can be configured to enable or disable hard link creation. For information, see [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),”](#) on page 285.

### **21.1.8 Hard Links and File Salvage and Purge**

The interaction between hard links and the salvage system depends on how many links exist for a specified file. If a file has multiple links pointing to it, and one of the links is deleted, the hard link name is not added to the salvage system, and cannot be recovered later via a salvage operation. When the last and only remaining link to a file is deleted, that last name is a candidate for the salvage system. If the Salvage Files attribute is enabled on the volume, and an immediate purge is not in force, then the last link to be deleted is added into the salvage system. Under this final primary link and filename, the file can be salvaged or purged until the file is autopurged from the system.

### **21.1.9 New Metadata Structure Supports Up to 65,535 Hard Links for a File (OES SP1 NetWare, NetWare 6.5 SP4)**

Beginning with OES SP1 NetWare® and NetWare 6.5 SP4, NSS volumes on NetWare support up to 65,535 hard links for a file, regardless of the length of the filename and the number of name spaces used. To use hard links on these operating systems, you must upgrade the NSS volume to use a new metadata structure (Beast Version 3). For information, see [Section 3.1, “Guidelines for Upgrading the Media Format of NSS Volumes on NetWare,”](#) on page 33.

The new metadata structure overcomes a prior limitation on the number of hard links for a file that exists in the old metadata structure (Beast Version 2). The old data structure is used for OES SP1 Linux and earlier, OES NetWare, and NetWare 6.5 SP3 and earlier. For information, see [Section 21.1.10, “Old Metadata Structure Supports Limited Hard Links for a File,”](#) on page 285.

## 21.1.10 Old Metadata Structure Supports Limited Hard Links for a File

In OES SP1 Linux and earlier, OES NetWare, and NetWare 6.5 SP3 and earlier, NSS volumes support only a limited number of hard links for a file. These volumes use a metadata structure (Beast Version 2) that provides limited space for the file's filenames across all namespaces.

With the old metadata structure, the number of hard links you can use for a file depends on the length of the names you choose for a file, and how those names can be shared across the UNIX, Long, DOS, and Macintosh namespaces. It is also limited by the amount of space available in the metadata structure.

---

**IMPORTANT:** The sum total of all metadata for a file cannot exceed 2 KB. In the old metadata structure, this includes all hard link names. The longer the names are, the fewer hard links you can create.

---

If you reach the 2 KB metadata limit, you cannot create new hard links until you make room for them. Use any of the following methods:

- Delete one or more hard links for the file.
- Create hard links for the file with shorter names that can be stored optimally in each namespace.
- Rename the original file with a shorter name that can be stored optimally in each name space.

## 21.2 Enabling or Disabling the Hard Links Attribute (OES SP1 NetWare, NetWare 6.5 SP4)

Beginning with OES SP1 Netware and NetWare 6.5 SP4, you must enable the Hard Links attribute for your NSS volume on NetWare before you can create hard links. When you upgrade the NSS volume to use the new media format, if any old-style hard links are detected, the Hard Links attribute is automatically enabled. Otherwise, the volume is upgraded, but the attribute is disabled and must be enabled before you can create hard links.

Use console commands to enable and disable the Hard Links attribute for NSS volumes that have been upgraded to the new media format on OES SP1 Netware and NetWare 6.5 SP4. The Hard Links attribute cannot be set or viewed in NSSMU or in the Storage plug-in to iManager 2.5.

### Configuring the Hard Link Attribute

To enable or disable the Hard Links attribute, use the following commands at a server console:

Command	Description
<code>nss /HardLinks=<i>volumename</i></code>	Enables the Hard Links attribute for the specified volume. This enables hard links to be created on the volume.
<code>nss /HardLinks=all</code>	Sets the Hard Links attribute for all NSS volumes on the server. This enables hard links to be created on any volume on the server. Any given hard link can point only to a file on the same volume.

Command	Description
<code>nss /NoHardLinks=<i>volume</i></code>	Disables the Hard Links attribute for the specified volume. Existing hard links continue to function, but no new hard links can be created on the specified volume.
<code>nss /NoHardLinks=all</code>	Disables the Hard Links attribute for all NSS volumes on the server. Existing hard links continue to function, but no new hard links can be created on any NSS volume on the server.

### Verifying the Hard Link Attribute Setting

To verify that the Hard Links attribute is set for the NSS volume: At the server console, enter

```
nss /volumes
```

In the Attributes column, the Hard Links attribute is listed if the attribute is set.

## 21.3 Creating a Hard Link Using Ln on a Linux Client

Typically, you create hard links, using clients (like NFS) that have existing commands to manipulate links. This section explains how to create a hard link in a volume by accessing the volume from a Linux workstation, then using the Link command (ln) to create a hard link.

Creating a hard link to a file does not copy the contents of the file to the new location; it simply makes a new name by which the file can be known, in addition to the file's existing name or names. All of the names are equally real, and no one of them is preferred to the others except regarding access control and directory space quotas. For information, see [Section 21.1, “Understanding Hard Links,” on page 281](#).

The actual content of the file is maintained in only one place: the original file. Users in different directories can use hard links to access and share the same file. A single user can use hard links to access a file from different directories.

### Prerequisite

For NSS volumes on NetWare servers that use the new metadata structure, you must enable the Hard Links attribute on the volume to be able to create hard links. For information, see [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 285](#).

For NSS volumes on Linux servers, for hard links to work properly, users must be Linux-enabled eDirectory users. For information, see [Section 6.4.5, “NSS and Linux User Management,” on page 77](#).

### Procedure

To create a hard link, use the Link (ln) command from a Linux client. At a terminal console, enter

```
ln /path/filename /linkpath/linkfilename
```

Replace */path/filename* with the path name for the original file. Replace */linkpath/linkfilename* with the path name for the hard linked filename. To maximize the number of hard links you can use, choose short names that can be stored optimally in each name space. For information, see [Section 21.1, “Understanding Hard Links,” on page 281](#).

For help with syntax and options, refer to the Link Manual Page from a Linux client. At a terminal console, enter

```
man 2 ln
```

For example, from a Linux client, you can use the Link command to link the file */tmp/timetest* to */usr/tmp/t1*, using the following command:

```
ln /tmp/timetest /usr/tmp/t1
```

The file can now be referred to by either name. You can delete the original name or any link name, but the contents of the file are not removed until the final name is deleted.

## 21.4 Creating a Hard Link Using a zLink API

To create a hard link, using a zLink API, see the *NDK: File System Services (64-bit)* ([http://developer.novell.com/ndk/doc/fs64/fs64\\_enu/data/a5p4x94.html#a5p4x94](http://developer.novell.com/ndk/doc/fs64/fs64_enu/data/a5p4x94.html#a5p4x94)).

## 21.5 Creating a Hard Link for Testing Purposes (OES SP1 NetWare, NetWare 6.5 SP4)

For testing purposes only, you can use a server console command to create hard links on OES SP1 NetWare or NetWare 6.5 SP4.

### Prerequisite

- The NSS volume on NetWare has been upgraded to use the new metadata structure. For information, see [“Upgrading the Media Format \(NetWare\)” on page 33](#).
- The Hard Links attribute is enabled for the volume. For information, see [Section 21.2, “Enabling or Disabling the Hard Links Attribute \(OES SP1 NetWare, NetWare 6.5 SP4\),” on page 285](#).

### Procedure for Creating a Hard Link for Testing Purposes

To create a hard link from the server console: At a command prompt, enter

```
nss /CreateHardLink=PrimaryFilePath|HardLinkFilePath
```

Replace *PrimaryFilePath* with the complete volume name, pathname, and filename of the primary file. Replace *HardLinkFilePath* with the complete volume name, pathname, and filename of the new hard link. Hard links must be created on the same volume as the original file.

---

**IMPORTANT:** The file paths are specified using the Long name space.

---

For example, to create a hard link on VOL1:

```
nss /CreateHardLink=VOL1:\path\file.ext|VOL1:\newpath\newfile.ext
```

## 21.6 Viewing Hard Links for a Volume (OES SP1 NetWare, NetWare 6.5 SP4)

To view a list of all hard links that exist for a volume: At the server console, enter

```
/nss /ListHardLinks=vol:path\filename.ext
```

Replace *path* with the file's primary link path or one of its hard link paths where you want to start the search. Replace *filename.ext* with the actual filename of the file, including the extension.

The command returns a list of hard links, where the first link in the list is the primary link. The second link is the most recently created hard link. This is the hard link that becomes the new primary link if the primary link is deleted. All links for the file are listed in the order they appear in the metadata, which is also the order in which a new primary name would be assigned.

For example, if you wanted the hard link listed as name 4 in the list to become the primary link path, names 1 (the primary link), name 2 (the most recently created hard link), and name 3 (the next most recently created link), would each need to be deleted without creating any new links in the meantime.

---

**IMPORTANT:** Make sure that you are aware of the impact of the link deletion on the inherited rights for the file as you are deleting links and the primary link is reassigned to the next hard link in the list.

---

The list reports the following information for the file:

Line	Information
Line 1	The ZID (file number) of the inode that contains the metadata for the hard link file set and the number of links associated with that inode.
Line 2	The complete name of the primary link, including the path and filename
Line 3	The complete name of the most recently created hard link name, including the path and filename. This is the next candidate in line to become the primary link if the primary link is deleted.
Line 4 - 65,536 (up to 65,535 links)	The complete name of each of the hard link names, including the path and filename, in the reverse chronological order of their creation time.

## 21.7 Deleting a Hard Link

If multiple links exist for a file, deleting a hard link to a file deletes only the link name, not the file's contents. NSS keeps count of how many hard links a file has and does not delete the file contents until all hard link names for the file and its primary link have been deleted (the link count goes to zero).

There are no special commands required to delete a hard link. In a file manager, locate the link name to be deleted, then delete the file, using the delete commands and procedures native to the client you use to access the file.



Whenever you delete a primary link for a file with hard links, the most recently created hard link automatically becomes the new primary link. This change potentially impacts the effective file system access rights for the file and the directory space quota for the new parent directory. Before deleting the primary link for a file with hard links, make sure you understand the following:

- [Section 21.1.1, “Hard Links and the Primary Link,” on page 282](#)
- [Section 21.1.4, “Hard Links and File System Trustees,” on page 283](#)
- [Section 21.1.5, “Hard Links and Directory Space Quotas,” on page 283](#)



# Using NSS on Devices Managed by Non-EVMS Volume Managers

# 22

The Novell® Storage Services™ file system can exist on devices managed by any volume manager. However, when using non-EVMS volume managers, a few of the NSS features are not available because the NSS management tools require EVMS to carry out related commands.

This section discusses the following:

- [Section 22.1, “Guidelines for Using NSS with Non-EVMS Volume Managers,” on page 291](#)
- [Section 22.2, “Creating and Mounting NSS Pools and Volumes,” on page 292](#)
- [Section 22.3, “Configuring Default Mount Settings for NSS Pools and Volumes,” on page 295](#)

## 22.1 Guidelines for Using NSS with Non-EVMS Volume Managers

If you use a volume manager other than EVMS to manage devices, some capabilities of NSS are not available. However, you can create and manage NSS pools and volumes on those devices, using the guidelines in this section.

- [Section 22.1.1, “Guidelines for the Server Install,” on page 291](#)
- [Section 22.1.2, “NSS Management Considerations,” on page 291](#)
- [Section 22.1.3, “NSS Capabilities That Are Unavailable,” on page 292](#)

### 22.1.1 Guidelines for the Server Install

#### Partition Setup

Make sure to leave unformatted free disk space so that you can create partitions for NSS after the install, using the Linux `mkfs` command.

#### Software Setup

NSS is not installed by default. To add NSS to the Software list at install time, open the Software setup option, search for the NSS package, select the NSS check box, then click Accept.

You can optionally install NSS after the initial install, using the *Software Updates* feature in YaST.

### 22.1.2 NSS Management Considerations

If you want to create NSS pools and volumes on devices managed by non-EVMS volume managers, consider the following management differences:

- Use the volume manager’s tools or Linux commands and utilities to manage devices, software RAID devices, and partitions. The NSS tools, utilities, or commands that manage devices, software RAID devices, and partitions are not available without EVMS.

- Use the Linux `mkfs` command to create an NSS pool, then use NSSMU create its storage objects in Novell eDirectory™. You cannot create pools with NSS tools, utilities, or commands without EVMS. For information, see [Section 22.2.1, “Using the Linux mkfs Command to Create NSS Pools,” on page 293](#).
- Manage existing NSS pools and create and manage NSS volumes with NSSMU and NSS commands and utilities.

For information, see the following:

- [Section 4.2.1, “NSSMU for Linux,” on page 53](#)
- [Chapter 23, “Using NSS Server Console Commands and Utilities,” on page 299](#)
- File system management tools are unaffected; they work with any volume manager.

For information, see the following:

- *File Systems Management Guide for OES*
- [Section 4.4, “Novell NetStorage,” on page 58](#)
- [Section 4.6.2, “Novell Client for Windows 2000/XP/2003,” on page 62](#)

### 22.1.3 NSS Capabilities That Are Unavailable

The following NSS capabilities are unavailable when the device is managed by a non-EVMS volume manager:

- **Management Tools:** The following NSS management tools are affected:
  - The Storage plug-in to iManager requires that EVMS; it does not work with other volume managers. It recognizes only NSS pools and volumes created on devices managed by EVMS.
  - NSSMU works only for the Volume management options and on existing pools that you created from the server’s command line interface. It requires EVMS to manage partitions, to manage devices, to create and manage software RAIDs, and to create pools.
- **NSS Software RAIDs:** NSS requires EVMS to create and manage software RAID devices. If you want to create software RAIDs, you must use the capability in whatever non-EVMS volume manager you use.
- **Cross-Platform Compatibility:** Pools created on devices managed by non-EVMS volume managers cannot be moved to a NetWare server.
- **Clusters:** Novell Cluster Services™ (NCS) for Linux requires EVMS for clustering. Pools created on devices managed by non-EVMS volume managers cannot be used in NCS clusters. Third-party clustering systems that integrate with third-party cluster volume managers should be able to work fine with NSS volumes.

## 22.2 Creating and Mounting NSS Pools and Volumes

Use the following procedure to create and mount an NSS pool and volume.

- [Section 22.2.1, “Using the Linux mkfs Command to Create NSS Pools,” on page 293](#)
- [Section 22.2.2, “Creating a Partition,” on page 293](#)
- [Section 22.2.3, “Creating and Mounting an NSS Pool,” on page 294](#)

- [Section 22.2.4, “Creating and Mounting an NSS Volume,” on page 295](#)

## 22.2.1 Using the Linux mkfs Command to Create NSS Pools

The Linux `mkfs` command is used to build a file system on Linux. This section describes how to use the Linux `mkfs` command to create an NSS file system pool.

---

**IMPORTANT:** Be very careful with this command. The `mkfs` command destroys any existing data on the specified device or partition.

---

### Syntax

```
mkfs [ -t filesystem ] [ fs-options ] filesystem
```

Option or Parameter	Description
<code>-t <i>fstype</i></code>	Specifies the type of file system to be built, such as <code>nsspool</code> or <code>nssvol</code> . For example:  <code>-t nsspool</code> <code>-t nssvol</code>
<code><i>fs-options</i></code>	File-system-specific options to be passed to the real file system builder.  When creating an NSS pool, use this fs-option: <code>-n <i>poolname</i></code> Replace <i>poolname</i> with the actual name of the pool you want to create.  When creating an NSS volume, use this fs-option: <code>-n <i>volname</i></code> Replace <i>volname</i> with the actual name of the volume you want to create.
<code><i>filesystem</i></code>	For <code>nsspool</code> , this is device name ( <i>devname</i> such as <code>/dev/hda1</code> or <code>/dev/sdb2</code> ).  For <code>nssvol</code> , this is the pool name ( <i>poolname</i> such as <code>POOL2</code> ).

### Examples

```
mkfs -t nsspool -n poolname devname
```

```
mkfs -t nssvol -n volname poolname
```

## 22.2.2 Creating a Partition

Use the following procedure to create a `/dev/partition` device, such as `/dev/hda2` or `/dev/sda5`.

- 1 Log in to the server as the root user, or use the `su` command to become root.
- 2 Go to the *YaST > System Partitioner*, then select *Create a Partition*.
- 3 Select a device with free space available.

- 4 Create a partition. The name is automatically specified, such as *hda2* or *sda5*.
- 5 If you do not want to use all of the available free space, specify the amount of space to use.  
Make sure the partition size is sufficient for the NSS pool you want to create later; the partition size determines the pool size.
- 6 Select *Unformatted*.
- 7 Do not specify the Mount Point; leave the field blank.
- 8 Click *OK* to create the partition.
- 9 Continue with [Section 22.2.3, “Creating and Mounting an NSS Pool,” on page 294](#).

### 22.2.3 Creating and Mounting an NSS Pool

- 1 Log in to the server as the root user, or use the `su` command to become root.
- 2 At a system command prompt, enter

```
mkfs -t nsspool -n poolname devname
```

Replace *poolname* with the name you want to use, such as `POOL2`. Replace *devname* with the device you created in [Section 22.2.2, “Creating a Partition,” on page 293](#).

---

**IMPORTANT:** Do not use ampersand (&) and pound (#) characters in pool and volume names; it creates problems in the `/etc/fstab` file. For information about other naming conventions, see “Naming NSS Storage Objects” on page 61.

---

For example, enter

```
mkfs -t nsspool -n POOL2 /dev/hda2
```

where `POOL2` is the pool name and `/dev/hda2` is the device name.

- 3 If the location where you want to mount the NSS pool does not already exist, create the mount point. At the server command prompt, enter

```
mkdir /mnt/pooldir
```

The `/mnt` directory is the default location for mounting devices. If you are using a different location, replace `/mnt` with that path. Replace *pooldir* with the name (path) you want to use, such as `POOL2`. If you want to make it a hidden directory, begin the directory name with a period, such as `.POOL2`.

For example, enter

```
mkdir /mnt/.POOL2
```

where `/mnt/.pool2` is the mount point for your pool.

- 4 Mount the NSS pool. At a system command prompt, enter

```
mount -t nsspool devname mountpoint -o name=poolname
```

For example, enter

```
mount -t nsspool /dev/hda2 /mnt/.POOL2 -o name=POOL2
```

- 5 Use NSSMU to create a Storage object in eDirectory for the newly created pool.

- 5a At a terminal prompt, enter

```
nssmu
```

- 5b From the NSSMU menu, select `Pools`.

- 5c Press `F4` (NDS Update)

- 6 Continue with [Section 22.2.4, “Creating and Mounting an NSS Volume,”](#) on page 295.

## 22.2.4 Creating and Mounting an NSS Volume

- 1 Log in to the server as the root user, or use the su command to become root.
- 2 Make sure the pool where you want to create the volume is mounted.

For information, see [Section 22.2.3, “Creating and Mounting an NSS Pool,”](#) on page 294.

- 3 Create an NSS volume. At the system console, enter

```
mkfs -t nssvol -n volname poolname
```

Replace *volname* with the name you want to use, such as NSSV1. Replace *poolname* with the NSS pool where the volume will reside, such as POOL2. This is the pool you created in [Section 22.2.3, “Creating and Mounting an NSS Pool,”](#) on page 294.

---

**IMPORTANT:** Do not use ampersand (&) and pound (#) characters in pool and volume names; it creates problems in the /etc/fstab file. For information about other naming conventions, see “Naming NSS Storage Objects” on page 61.

---

For example, enter

```
mkfs -t nssvol -n NSSV1 POOL2
```

- 4 If the location where you want to mount the NSS volume does not already exist, create the directory path. At the server command prompt, enter

```
mkdir /media/nss/volname
```

The /media/nss path is the default location for NSS volumes. If you are using a different location, replace /media/nss with that path. Replace *volname* with the name of the volume you created in [Step 3](#).

For example, enter

```
mkdir /media/nss/NSSV1
```

- 5 Mount the NSS volume. At a system command prompt, enter

```
mount -t nssvol VOL volmountpoint -o name=volname
```

For example, enter

```
mount -t nssvol VOL /media/nss/NSSV1 -o name=NSSV1
```

## 22.3 Configuring Default Mount Settings for NSS Pools and Volumes

Whenever you create NSS pools and volumes from the command line, their mount information is not added by default to the /etc/fstab configuration file. After creating the pool and volume, make sure to edit the /etc/fstab configuration file to add entries for them.

- [Section 22.3.1, “Understanding Entries in the /etc/fstab Configuration File,”](#) on page 296
- [Section 22.3.2, “Adding NSS Pool and Volume Mount Information to /etc/fstab,”](#) on page 297

## 22.3.1 Understanding Entries in the /etc/fstab Configuration File

The `/etc/fstab` file is a configuration file that contains information about all the devices and partitions on your Linux computer. Each line is an entry that describes where and how to mount one device or one partition. The following table describes the field information needed for NSS pools and volumes.

Column in <code>/etc/fstab</code>	Description	Examples
Device name	The location of the device or partition you want to mount	<code>/dev/hda3</code> <code>/dev/sdb1</code>
Mount point	The default location where the device or partition is to be mounted if the mount point is not otherwise specified in a mount command	<code>/mnt/.pool2</code> <code>/media/nss/NSSV1</code>
File system type	The file system type of the device or partition	<code>nsspool</code> <code>nssvol</code>
Mount options	<p>Lists the comma-delimited mount options for the device or partition. Use a comma without spaces between options.</p> <p><b>auto or noauto:</b> Use <code>auto</code> if you want the volume to mount automatically after a system reboot. Use <code>noauto</code> if you want the device to be mounted only when you explicitly issue the mount command. Do not use the <code>auto</code> command for encrypted NSS volumes; they require a password to be entered on the first mount after a system reboot.</p> <p><b>rw:</b> Use <code>rw</code> to mount the device as read-write.</p> <p><b>name=&lt;poolname   volname&gt;:</b> Specify the <i>poolname</i> or <i>volname</i> of the partition to be mounted.</p>	<code>noauto,rw,name=POOL2</code> <code>noauto,rw,name=NSSV1</code>
Dump option number	<p>Designates if the file system should be dumped with the Linux Dump utility. A value of 0 indicates that Dump should ignore this file system.</p> <p>Set this value to 0 for NSS pools and volumes.</p>	0
File system check number	<p>Designates if the file system should be checked with the Linux <code>fsck</code> utility. A value of 0 indicates that <code>fsck</code> should ignore this file system.</p> <p>Set this value to 0 for NSS pools and volumes.</p>	0



## 22.3.2 Adding NSS Pool and Volume Mount Information to /etc/fstab

For information about completing the fields for an entry in the /etc/fstab file, see [Section 22.3.1, “Understanding Entries in the /etc/fstab Configuration File,”](#) on page 296.

- 1 Log in as the root user, or use the su command to become root.
- 2 In a text editor, open the /etc/fstab configuration file.

---

**IMPORTANT:** When working in fstab, make sure not to leave any stray characters or spaces in the file. This is a configuration file, and it is highly sensitive to such mistakes.

---

- 3 Locate the area in /etc/fstab where partitions are defined, then add a line defining the mount information for the NSS pool.

*devname mountpoint fstype mountoptions dump# fsck#*

For example, suppose you created POOL2 on device /dev/hda2 and mounted it at /mnt/.POOL2. In the /etc/fstab file, the line to add would be

```
/dev/hda2 /mnt/.POOL2 nsspool noauto,rw,name=POOL2 0 0
```

If you want to mount an NSS pool automatically after a system reboots, make sure to use the auto command.

- 4 Locate the area in /etc/fstab below where you entered pool information, then add a line defining the mount information for the NSS volume.

*devname mountpoint fstype mountoptions dump# fsck#*

Make sure to place the NSS volume entry after its pool entry to ensure that the pool is mounted before the volumes in it.

For example, suppose you created volume NSSV1 and mounted it at /media/nss/NSSV1. In the /etc/fstab, the line to add would be

```
NSSV1 /media/nss/NSSV1 nssvol noauto,rw,name=NSSV1 0 0
```

If you want to mount a non-encrypted NSS volume automatically after a system reboots, make sure to use the auto command on both the pool it is in and the volume.

- 5 Save the /etc/fstab file to accept your changes.



# Using NSS Server Console Commands and Utilities

# 23

This section discusses the Novell® Storage Services™ console commands for the NetWare® server console or the NSS Console (`nsscon`) on Linux.

---

**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 `CacheStats` command.

---

- [Section 23.1, “Online Help Commands,” on page 299](#)
- [Section 23.2, “NSS Startup Commands,” on page 300](#)
- [Section 23.3, “Status Commands,” on page 301](#)
- [Section 23.4, “Maintenance Commands,” on page 302](#)
- [Section 23.5, “Maintenance Commands \(NetWare\),” on page 303](#)
- [Section 23.6, “System Volume Recovery Commands \(NetWare\),” on page 304](#)
- [Section 23.7, “Compression Commands,” on page 305](#)
- [Section 23.8, “Encrypted Volume Support Commands,” on page 305](#)
- [Section 23.9, “Salvage and Purge Commands,” on page 306](#)
- [Section 23.10, “Read Ahead Commands,” on page 307](#)
- [Section 23.11, “Cache Balancing Commands \(NetWare\),” on page 308](#)
- [Section 23.12, “Commands for Mounting CDs and DVDs as NSS Volumes \(NetWare\),” on page 309](#)
- [Section 23.13, “Commands for CD and DVD Image Files as NSS Volumes \(NetWare\),” on page 310](#)
- [Section 23.14, “DOS FAT Commands \(NetWare\),” on page 311](#)
- [Section 23.15, “Multiple Server Activation Prevention Commands,” on page 312](#)
- [Section 23.16, “Multipath Connection Failover Commands \(NetWare\),” on page 312](#)
- [Section 23.17, “Pool Snapshot Commands \(NetWare\),” on page 314](#)
- [Section 23.18, “Other NSS Commands,” on page 315](#)
- [Section 23.19, “Security Equivalence Vector Update Commands \(Linux\),” on page 316](#)

This section does not include a complete list of NSS commands. To see a complete list, enter `nss /help` at the server console on NetWare or the NSS Console (`nsscon`) on Linux.

## 23.1 Online Help Commands

To access online Help for NSS console commands: enter one of the following commands:

```
nss /help
```

```
nss /?
```

The interface returns a list of NSS commands, each with a brief definition and any constraints such as the default value, a range of valid values, and persistence information.

### **Default Value**

The default value is the setting used for a given server configuration parameter unless you modify it for your particular implementation. Default values are the best choice for a majority of server configurations, but you should review the settings as you fine-tune the server.

### **Range of Valid Values**

The range of valid values establishes the constraints for any particular variable setting.

### **Persistence**

If a command's setting is persistent, the values you set remain in effect for the server through any subsequent server reboots until you next modify the settings. If a command is not persistent, the setting remains in effect only until the next server reboot.

## **23.2 NSS Startup Commands**

NSS provides the following startup commands. Startup commands are only valid when NSS is loading.

### **23.2.1 Issuing NSS Startup Commands while OES NetWare Boots**

NSS Startup commands can run as part of the OES NetWare `c:\nwserver\server.exe` command by using the `-z` option. However, these commands are not persistent through a server reboot; you must enter them each time you boot your server.

At the server command line, enter

```
server -z"/parameter=value"
```

In the above example, the `-z` tells the server to look for specific NSS startup commands. The information in quotes is the actual command with command values, if any. For example:

```
server -z"CDDVD"
```

### **23.2.2 Making NSS Startup Commands Persist through an OES NetWare Reboot**

NSS Startup commands issued at the command line during a server boot do not persist through a server reboot. To make non-persistent command settings persist automatically through a server reboot, place the commands in the `nssstart.cfg` file, which NSS reads on startup.

- 1 In a text editor, create a file called `nssstart.cfg` in the `c:\nwserver` directory.

- 2 Enter any NSS startup commands that you want to persist through server reboots.

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

```
/numworktodos=40 /cddvd /dosfat
```

3 Save and close your `c:\nwserver\nssstart.cfg` file.

## 23.2.3 Load Commands

Commands	Description
CDDVD	<p>If it is specified in the <code>nssstart.cfg</code> file, loads the <code>cddvd.nss</code> module that is essential for CD, DVD, and Macintosh support.</p> <p>For more information, see <a href="#">Section 13.11, “Activating and Deactivating CDs and DVDs as NSS Volumes (NetWare),” on page 193</a>.</p>
defaultNLMs	<p>If it is specified in the <code>nssstart.cfg</code> file, loads only default NSS NLM™ programs and any additional LSS modules specified on the command line.</p>
DOSFAT	<p>If it is specified in the <code>nssstart.cfg</code> file, loads only those modules that are essential for accessing FAT 16 and FAT 12 drives. NetWare does not support FAT 32 drives.</p> <p>DOSFAT supports only hard drives. DOSFAT does not support USB drives.</p>
/allLSS	<p>If it is specified in the <code>nssstart.cfg</code> file, it loads all LSS modules.</p>
/noLSS	<p>If it is specified in the <code>nssstart.cfg</code> file, it does not load the LSS modules.</p>
/NumAsyncIOs= <i>value</i>	<p>Sets the number of asynchronous I/O entries to allocate.</p> <p>[Default=2048; Range=4 to 65536]</p>
/NumBonds= <i>value</i>	<p>Sets the number of bond entries to allocate.</p> <p>[Default=5000; Range=512 to 2097152]</p>
/NumWorkToDos= <i>value</i>	<p>If used in the <code>nssstart.cfg</code> file, sets the number of WorkToDo entries. Entries can execute concurrently.</p> <p>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.</p> <p>[Default=50; Range=5 to 100]</p>
/zLSS	<p>If it is specified in the <code>nssstart.cfg</code> file, loads only those modules that are essential for zLSS support.</p>

## 23.3 Status Commands

NSS provides the following commands to show the status of various NSS parameters. To view a complete list, use the `nss /help` command.

Commands	Description
Status	Lists the current NSS status information.
Pools	Lists all of the NSS pools that are currently available on the server.
PoolAutoDisplay	Displays load-time policies for pools on the server.
Volumes	Lists all of the NSS volumes that are currently mounted and active, including the Admin volume.
SpaceInformation	Lists the amount of space on active pools and their associated volumes.
ListFreeSpace	Lists the amount of available space that has not been assigned to a pool.
nss /StorageAlertMessages	Enables NSS to send Low Storage messages to all users. [Default=On]
nss /NoStorageAlertMessages	Disables NSS from sending Low Storage messages to all users.
CompScreen	Display the NSS Compression Statistics GUI. For more information, see <a href="#">Section 23.7, "Compression Commands," on page 305</a> .
CacheStats	Shows the caching statistics for buffers.
FileCacheStats	Shows the caching statistics for file system objects.
NameCacheStats	Shows the caching statistics for names.
ResetStats	Resets caching and file statistics.
nss /ErrorCode= <i>code</i>	Translates and describes the specified error code.
nss /GetMFLStatus= <i>volume</i>	Shows the modified file list status for the specified volume.
Modules	Lists the providers, loadable storage subsystems, and semantic agents.
Version	Displays the version information for NSS.

## 23.4 Maintenance Commands

NSS provides the following maintenance for pools and volumes. To view a complete list, use the nss /help command.

Commands	Description
nss /BackgroundChecking	Enables the background file system checker. [Default=On]
nss /NoBackgroundChecking	Disables the background file system checker.

Commands	Description
ForceBackgroundCheck	Force background file system checking to start.
nss /FixDirectoryQuotas= <i>volume</i>	Recomputes used space for directory quotas for the specified volume.
nss /FixMFL	Repairs the modified file list to be consistent with the file system.
nss /MFL= <i>volume</i>	Enables modified file list maintenance for the specified volume.
nss /NoMFL= <i>volume</i>	Disables modified file list maintenance for the specified volume.
nss /MFLVerify= <i>volume</i>	Compares the modified file list with the specified volume and reports any inconsistencies.
nss /IncludeType= <i>type</i>	Include pools of the specified type for the given command. Possible types are SHARED, CORRUPT, and QUESTIONS.
nss /OverrideType= <i>type</i>	Override pools of the specified type for the given command. Possible types are SHARED, CORRUPT, and QUESTIONS.
pools	Shows all of the currently available NSS pools.
nss /VolumeMaintenance= <i>volume</i>	Places a specified volume into maintenance mode. Volumes can be put in maintenance mode, but maintenance occurs only at the storage pool level.
nss /UpgradeObjectsOnVolume	Upgrades objects on volumes to the current version.

## 23.5 Maintenance Commands (NetWare)

NSS provides the following maintenance commands for pools and volumes on NetWare. To view a complete list, use the nss /help command.

Commands	Description
nss /PoolAutoDisplay	Displays current pool load time policies.
nss /PoolAutoActivate= <i>pool</i>	Activates specified pools at load time. (pools with only unencrypted volumes)
nss /PoolAutoDeactivate= <i>pool</i>	Leaves specified pools deactivated at pool load time. (pools with only unencrypted volumes)
nss /PoolAutoMaintenance= <i>pool</i>	Places specified pools in maintenance mode at pool load time. (pools with only unencrypted volumes)
nss /PoolAutoVerify	Verifies the specified pool's physical integrity at startup time. (pools with only unencrypted volumes)
nss /PoolRebuild= <i>pool</i>	Rebuilds the data structure of the specified pool.

Commands	Description
PoolRebuild	Select the pools to rebuild from a menu, then it rebuilds them.
nss /PoolVerify= <i>pool</i>	Verifies the physical integrity of the specified pool.
poolverify	Allows you to select pools from a menu and verifies their physical integrity.
nss /VolumeActivate= <i>volume</i>	Activates the specified NSS volume.
nss /VolumeDeactivate= <i>volume</i>	Deactivates the specified NSS volume.
nss /VolumeAutoActivate= <i>volume</i>	Activates the specified volume at load time.
nss /VolumeAutoDeactivate= <i>volume</i>	Deactivates the specified volume at load time.
nss /ForceActivate= <i>volume</i>	Forces an NSS volume to become active.
VolumeAutoDisplay	Displays the current volume's load time policies.

## 23.6 System Volume Recovery Commands (NetWare)

NSS provides the following commands for recovering the system volume on NetWare. To view a complete list, use the `nss /help` command.

Commands	Description
nss /ChangeSysQuota= <i>size</i>	Lets you change the quota (in MB) for the <code>sys:</code> volume. Setting this value to zero sets the Quota to none and allows the <code>sys:</code> volume to grow to the size of the pool.
nss /ExpandSys	Lets you expand the pool that contains your <code>sys:</code> 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 <code>/ListFreeSpace</code> command.
nss /ListFreeSpace	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 <code>/ExpandSYS</code> command.
nss /RenameToSys	Lets you rename the <code>sys:</code> volume back to <code>sys:</code> if you have previously renamed it. At the prompt, enter the name of the volume you want to rename to <code>sys:</code> .
nss /SalvageSys	Lets you restore the <code>sys:</code> volume if you have deleted it (only if it has not yet been purged).



## 23.7 Compression Commands

NSS offers the following commands for configuring and monitoring compression. To view a complete list, use the `nss /help` command.

Command	Description
<code>CompScreen</code>	Displays the NSS volume compression statistics on the compression screen.
<code>nss /Compression=volume</code>	Enables the Compression attribute for the specified volume. After you enable the Compression attribute, the setting persists for the life of the volume. You cannot disable compression, but you can set parameters to effectively turn it off. For information, see <a href="#">Chapter 17, “Managing Compression on NSS Volumes,”</a> on page 227.
<code>nss /Compression=all</code>	Enables the Compression attribute for all volumes on the server.
<code>StopNormalCompression</code>	Stops all queued compression for files, based on the compression triggered by a file open or close.
<code>nss /BGCompression</code>	Allows compression to occur in the background at any time, instead of only within specified hours.
<code>nss /NoBGCompression</code>	Stops background compression and clears any queued background compression requests. Allow compression to occur only within the specified hours.

## 23.8 Encrypted Volume Support Commands

You can use NSS console commands to display volume status and to activate, mount, deactivate, or dismount encrypted volumes. You must enter a password only on the first activation following a system reboot. Thereafter, other environmental security and authentication measures control access to user data.

The following table provides the syntax for NSS commands to use with encrypted volumes. In each case, replace *volname* with the name of the encrypted NSS volume. In general, you cannot use the wildcard option of “all” as the *volname*. Until an encrypted volume is activated with its password following each system reboot, the All option does not find the volume and does not execute the command. The system returns an error message, requesting more information.

Command	Description
<code>nss /activate=volname</code>	On NetWare, activate the encrypted NSS volume in two steps.
If prompted, enter the encryption password for the encrypted volume.	The password is required only on the first activation following a system reboot.

Command	Description
<code>nss /volumeactivate=volname</code>	On NetWare, activate the encrypted NSS volume in two steps.  The password is required only on the first activation following a system reboot.
<code>nss /activate=all</code>	Activate all encrypted NSS volumes that have been previously activated with their passwords.
<code>nss /volumes</code>	View the status of an encrypted and unencrypted NSS volumes.  The encrypted volume returns a status of Encrypted.
<code>mount volname</code>	On Linux, mount an encrypted or unencrypted NSS volume.  The mount command is usable for encrypted volumes only after a previous mount with password in NSSMU. Otherwise, it returns an error message, requesting more information.
<code>mount all</code>	Activate and mount all encrypted NSS volumes that have been previously activated with their passwords.  Encrypted NSS volumes that were not previously activated return an error message, requesting more information.

You cannot use wildcard options and commands to activate encrypted volumes, such as `nss /VolumeAutoActivate`.

## 23.9 Salvage and Purge Commands

NSS offers the following commands for tuning the purging processes for a volume. To view a complete list, use the `nss /help` command.

Commands	Description
<code>nss /Salvage=volume/all</code>	Enables salvage of deleted files on volumes. <i>Volume</i> enables the Salvage command on the specified NSS volume on the specified server. All enables the Salvage command on all NSS volumes on the specified server.
<code>nss /NoSalvage=volume/all</code>	Disables salvage of deleted files on volumes. <i>Volume</i> disables the Salvage command on the specified NSS volume on the specified server. All disables the Salvage command on all NSS volumes on the specified server.

Commands	Description
<code>nss /LogicalVolumePurgeDelay=value</code>	<p>The number of seconds before deleted logical volumes are purged. This allows time to reverse the deletion.</p> <p>[Default=345600]</p>
<code>nss /LogicalVolumePurgeDelayAfterLoad</code>	<p>The number of seconds after NSS loads before deleted logical volumes are purged. This allows time to reverse the deletion.</p> <p>[Default=7200]</p>
<code>nss / LogicalVolumePurgeDelayAfterContinue=value</code>	<p>The number of seconds to delay purging a logical volume after clicking Continue. After a volume starts to purge, it cannot be salvaged.</p> <p>[Default=900]</p>
<code>nss /PoolHighWaterMark=poolname:Percent</code>	<p>Purging begins in the salvage area when the volume's low watermark is reached and continues until its high watermark is reached. When enough files are purged such that the free disk space increases to a percentage equal to or greater than the high watermark, the autopurge stops. Autopurging does not start again until free space again drops below the low watermark.</p> <p>The high and low watermarks must be at least 2% apart from each other.</p> <p>[Default=20; Range=2 to 100]</p>
<code>nss /PoolLowWaterMark=poolname:Percent</code>	<p>Purging begins in the salvage area when the volume's low watermark is reached and continues until its high watermark is reached. When free disk space falls below a low watermark, NSS begins autopurging the salvage area.</p> <p>The high and low watermarks must be at least 2% apart from each other.</p> <p>[Default=10; Range=0 to 98]</p>

## 23.10 Read Ahead Commands

NSS offers the following Read Ahead commands for tuning your storage solution. To view a complete list, use the `nss /help` command.

Commands	Description
<code>nss /AllocAheadBlks=<i>value</i></code>	Sets the number of blocks to allocate ahead on writes.  [Default=15; Range=0 to 63]
<code>nss /ReadAheadBlks=<i>VolName:Count</i></code>	Specify <i>VolName</i> as the name of the volume that you are setting the attribute for. Specify the <i>Count</i> to be the number of 4 KB blocks that you want to NSS to read ahead. The default block count is 2 blocks, which satisfies the performance requirements for a typical volume. The valid range for a block count is 0 blocks to 1024 blocks, where a block count of zero (0) implies no read ahead. However, 128 blocks is the practical maximum.  Read-ahead block counts higher than 128 can starve other system components for memory or buffers, which can impair performance or cause the system to hang. As the number of concurrent connections to unique files increase, you should reduce the number of read-ahead blocks.  [Default=2; Maximum=1024; Practical maximum=128]

## 23.11 Cache Balancing Commands (NetWare)

NSS provides commands to balance the cache for NSS on NetWare.

**WARNING:** If you give NSS more memory than you allow for the server, NSS fails to load. Your NSS-based `sys:` volume becomes inoperable until you load the `c:\nwserver\server.exe` with an NSS override to correct the situation. For information, see [Section 23.2, “NSS Startup Commands,” on page 300](#).

Commands	Description
<code>nss /CacheBalance=<i>value</i></code>	Enables the dynamic balancing of free memory for the NSS buffer cache. Sets the specified percentage of free kernel memory that NSS uses for its buffer cache. [Default=85; Range=1 to 99]
<code>nss /NoCacheBalance</code>	Disables the dynamic balancing of free memory for the NSS buffer cache.
<code>nss /MinBufferCacheSize=<i>value</i></code>	Sets the specified minimum number of NSS buffer cache entries. [Default=512; Range=384 to memory size in KB/4 KB]
<code>nss /MinOSBufferCacheSize=<i>value</i></code>	Sets the specified minimum number of NetWare buffer cache entries. [Default=256; Range=256 to memory size in KB/4 KB]

Commands	Description
<code>nss /NameCacheSize=value</code>	Sets the specified maximum number of Name cache entries. [Default=100000; Range=17 to 1000000]
<code>nss /CacheUserMaxPercent=value</code>	Sets the specified maximum percentage of buffer cache that can be used for user data. [Default=80; Range=10 to 100]
<code>nss /ClosedFileCacheSize=value</code>	Sets the number of closed files that can have their beast objects cached in memory. On average, each Closed File Cache entry consumes 0.4 - 1KB of RAM.  [Default=50000; Range=16 to 1000000; Recommended=100000]
<code>CacheStats</code>	Shows the caching statistics for buffers.
<code>FileCacheStats</code>	Shows the caching statistics for file system objects.
<code>NameCacheStats</code>	Shows the caching statistics for names.
<code>ResetStats</code>	Resets caching and file statistics.
<code>nss /FlushFilesImmediately=volume</code>	Synchronously flush data to the specified volume when its files are closed instead of waiting for the next scheduled write of data to the volume.
<code>nss /NoFlushFilesImmediately=volume</code>	Disables synchronously flushing files to the specified volume when its files are closed. Instead, data is written to the volume on the next scheduled write of data to the volume.
<code>nss /FileFlushTimer=value</code>	Sets the flush time in seconds for writing data to modified open files. Increasing this number might reduce the number of writes to disk; however, it increases the amount of data that is lost if the system crashes.  [Default=10 seconds; Range=1 to 3600 seconds]
<code>nss /IDCacheResetInterval</code>	Set the number of seconds between invalidation of the ID cache. [Default=90000 Range=0-200000000]

## 23.12 Commands for Mounting CDs and DVDs as NSS Volumes (NetWare)

Cddvd.nss is loaded automatically on startup by default. Use the following commands to manage CDs, DVDs, and Macintosh CDs as NSS volumes on Netware.

**NOTE:** The cddvd.nss module replaces cdrom.nlm and cd9660.nss for CD support, udf.nss for DVD support, and cdhfs.nss for Macintosh CD support. The following commands for cddvd.nss are the same ones used for previous CD and DVD management modules.

Commands	Description
CDDVD	If CDDVD is specified in the nssstart.cfg file, loads only those modules essential for CD and DVD support.
nss / (No) LeaveCDinNDS	<p>If the switches UpdateCDToNDS or LoginUpdateCDToNDS is enabled, enable this switch to retain the CD or DVD volume objects when the actual disks are removed.</p> <p>[Default=Off]</p>
nss / (No) LoadISOVolumeFirst	<p>If this switch is enabled, tries to load UDF-bridge CD as and ISO9660 volume first. By default, a UDF-bridge CD is loaded as DVD format first and CD format second.</p> <p>[Default=Off]</p>
nss / (No) LockCDMediaIn	<p>If this switch is enabled, whenever CDs or DVDs are activated or deactivated as NSS volumes, locks or unlocks the media in the device.</p> <p>[Default=Off]</p>
nss / (No) UpdateCDToNDS	<p>If this switch is enabled, whenever CDs or DVDs are activated or deactivated as NSS volumes, updates their Novell eDirectory objects.</p> <p>[Default=Off]</p>
nss / (No) LoginUpdateCDtoNDS	<p>If this switch is enabled, whenever CDs or DVDs are activated or deactivated as NSS volumes, updates their Novell eDirectory objects. You must login as admin.</p> <p>[Default=Off]</p>
nss / (No) MountISODiscSetInOneVolume	<p>If this switch is enabled, cddvd.nss attempts to load an ISO9660 multiple-disk set as a single NSS volume. This volume is not mounted until all disks among the set are available to cddvd.nss.</p> <p>If this switch is disabled, cddvd.nss mounts each disk in the set as one individual NSS volume.</p> <p>[Default=Off]</p>

## 23.13 Commands for CD and DVD Image Files as NSS Volumes (NetWare)

Cddvd.nss is loaded automatically on startup by default. Cddvd.nss provides the following commands for adding and removing CD and DVD image files as NSS Volumes.

While the CD or DVD image source file is loaded as an NSS volume, the following actions cause the image file to be removed as an NSS volume:

- Any file operations on the image source file, such as open, read, or write
- Deactivation of the NSS volume where the image source file resides

Commands	Description
<code>nss /MountImageVolume=path_imagefilename</code>	<p>Adds and mounts the specified CD or DVD image file as a read-only NSS volume.</p> <p>Replace <i>path_imagefilename</i> with the path on the NSS volume where the image file resides, including its filename. For example:</p> <pre>nss /MountImageVolume= nssvol:\suse9_1.iso</pre>
<code>nss /RemoveImageVolume=path_imagefilename</code>	<p>Dismounts and removes the specified CD or DVD image file as a read-only NSS volume.</p> <p>Replace <i>path_imagefilename</i> with the path on the NSS volume where the image file resides, including its filename. For example:</p> <pre>nss /RemoveImageVolume= nssvol:\suse9_1.iso</pre>

## 23.14 DOS FAT Commands (NetWare)

NSS provides the following commands for managing DOS FAT partitions.

Commands	Description
DOSFAT	<p>If DOSFAT is specified in the <code>nssstart.cfg</code> file, loads only those modules essential for accessing FAT 16 and FAT 12 drives. OES NetWare does not support FAT 32 drives.</p> <p>DOSFAT supports only hard drives. DOSFAT does not support USB drives.</p>
<code>nss /FATPartition=partition_type_number</code>	Supports up to three additional partition types containing 16-bit FATs, such as /FATPartition=12,13.
<code>nss / (No) FATLongNames</code>	<p>Enables long filenames on FAT volumes.</p> <p>[Default=On]</p>

## 23.15 Multiple Server Activation Prevention Commands

The management file for pool MSAP is

`_admin\manage_nss\pool\poolname\zlss\msap.xml`. One file exists for each pool. This file contains MSAP statistics for the pool. The MSAP attribute is displayed in the `enabledAttributes` tag 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.

MSAP Console Command	Description
<code>nss /msapserver</code>	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.
<code>nss /nomsapserver</code>	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.
<code>nss /msaprebuild</code>	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 rebuilder ownership of the pool.
<code>nss /poolmsap=pool</code>	Enables MSAP for a given pool on the server. MSAP is enabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now enabled.)
<code>nss /nopoolmsap=pool</code>	Disables MSAP for a given pool. Use the command when the pool is activated. MSAP is disabled the next time the pool is activated. (Enter the command, deactivate the pool, then activate the pool. MSAP is now disabled.)
<code>nss /pools</code>	Displays the message "Multi-use detect" for pools that have MSAP enabled.

## 23.16 Multipath Connection Failover Commands (NetWare)

NSS provides the following Media Manager commands for managing multipath connection failover.



Media Manager Multipath I/O Console Command	Description
<code>list failover devices <i>device_id</i></code>	<p>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.</p> <p>To list the paths associated with all the devices on the server, enter the command without specifying the <i>device_id</i>.</p> <p>To list the paths associated with a given device, enter the command followed by the <i>device_id</i>.</p>
<code>mm set failover priority <i>pathid=number</i></code>	<p>This command sets the priority level for the selected path. The <i>pathid</i> must be a valid path ID. The <i>number</i> is the priority to set for the selected path. The number must be a decimal integer value between 1 and 4 million, with 1 being the highest priority and 4 million being the lowest priority.</p>
<code>mm set failover priority <i>pathid=number</i> /insert</code>	<p>The /insert option inserts the specified path with its 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.</p>
<code>mm set failover state <i>pathid=up</i> / setpath</code>	<p>This command sets the state of the specified path to Up. The <i>pathid</i> must be a valid path ID.</p> <p>If the Up option is followed by the /setpath option, NSS automatically selects the highest priority path from among all the paths that are in the Up state, including the new path.</p>
<code>mm set failover state <i>pathid=down</i> /setpath</code>	<p>This command enables you to designate a path to be in the Down state so you can perform preventative maintenance or reroute cables on a good, functional path. If the path is bad, NSS automatically moves the path to a Down state.</p> <p>If the Down option is followed by the /setpath option, NSS automatically selects the highest priority path from among the other paths that are currently in the Up state.</p>
<code>mm set failover path <i>pathid</i></code>	<p>This command activates the selected path as <b>selected</b>, which means it serves as the primary path. The <i>pathid</i> must be a valid path that is currently in the Up state. The priority of a path remains static regardless of whether a path is considered up or down or whether a path is selected or not selected as the primary path.</p>

Media Manager Multipath I/O Console Command	Description
<code>mm restore failover path <i>device_id</i></code>	This command forces the device to reselect the highest priority path that is online. The <i>deviceid</i> must be a valid device ID.
<code>mm reset failover registry</code>	This command resets the multipath settings in the operating system registry.

## 23.17 Pool Snapshot Commands (NetWare)

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

Snapshot Management 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.
<code>mm snap rename <i>snapname newsnapname</i></code>	Renames the pool as specified.
<code>mm snap list</code>	Lists the snapshots on the server.
<code>mm snap shred <i>snapname</i> = on</code>	<p>Turns on shredding for the internal volume of the pool where you are storing the specified snapshot.</p> <p>If this option is enabled:</p> <ul style="list-style-type: none"> <li>• You can have only one snapshot at a time of its original pool.</li> <li>• After you delete the snapshot, allow enough time for the shredding process before you create a new one.</li> </ul>
<code>mm snap shred <i>snapname</i> = off</code>	<p>Turns off shredding for the internal volume of the pool where you are storing the specified snapshot.</p> <p>If any pool snapshot was created on this pool prior to the NetWare 6.5 SP3 release, shredding is already turned on. Use this command to turn off shredding.</p>

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

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

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

## 23.18 Other NSS Commands

NSS provides these additional commands for managing pools and volumes. To view a complete list, use the `nss /help` command.

Commands	Description
<code>nss /NoClassicDirectoryQuotas</code>	Turns emulation of Traditional NetWare volumes on or off.  [Default=Off]
<code>nss /DataShredding=volume:count</code>	Enables the Data Shredding attribute for the specified volume. Specify the number of times you want to shred data.  Data shredding overwrites purged files with bit patterns up to seven times. 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.  [Default=1; Range=0 to 7, where 0 indicates no shredding]
<code>nss /NoDataShredding=volume</code>	Disables data shredding on the specified 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>	Does not create a copy of files on the volume when they are open for write.
<code>nss /LVDeleteStatusBasic</code>	Displays information about deleted logical volumes.
<code>nss /LVDeleteStatusSalvagable</code>	Displays information about salvagable logical volumes.
<code>nss /Migration=volume/all</code>	Enables migration of files on the volumes.
<code>nss /NoMigration=volume/all</code>	Disables migration of files on the volumes.
<code>nss /StorageAlarmThreshold=value</code>	Lets you set the threshold for a low storage space warning.  [Default=10; Range=0 to 1000000]
<code>nss / (No) StorageAlertMessages</code>	Turns On or Off the low storage message to users.  [Default=On]

Commands	Description
<code>nss /StorageResetThreshold=<i>value</i></code>	Lets you reset the threshold for a low storage space warning.  [Default=10; Range=1 to 1000000]
<code>nss /NoTransaction=<i>volume/all</i></code>	Disables transactional writes in files on the volumes.
<code>nss /UserSpaceRestrictions</code>	Enables user quotas on the volume.
<code>nss /NoUserSpaceRestrictions</code>	Disables user quotas on the volume.
<code>nss /Transaction=<i>volume/all</i></code>	Enables transactional writes in files on the volumes.
<code>nss /ZLSSIOStatus</code>	Displays current NSS IO status information.
<code>nss /ZLSSPendingWriteIOs</code>	Number of write inputs dropped to the storage subsystem. (0 = all available)  [Default=1000; Range=0 to 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.

## 23.19 Security Equivalence Vector Update Commands (Linux)

The Security Equivalence Vector (SEV) is calculated for each NSS user based on information in the user's profile in Novell eDirectory. It is a list of eDirectory GUIDs, for example:

- the user's own GUIDs
- GUIDs of groups that include the user
- GUIDs of parent containers for the user and his or her groups
- security equivalent GUIDs

The SEV is used to validate the user against the trustee rights of the directory and file the user is attempting to access.

After you boot the Linux server, when a user first attempts to connect to the NSS file system, NSS contacts Novell eDirectory to retrieve the user's Security Equivalence Vector (SEV). eDirectory calculates the user's effective rights for the NSS volume, creates the SEV, and passes it to NSS. NSS compares the user's SEV with file system trustees and trustee rights for the specified file or directory to determine if the user can access the resource.

NSS caches the SEV locally in the server memory, where it remains until the server is rebooted or unless the user is deleted from eDirectory. NSS polls eDirectory at a specified interval for updates to the SEVs that are in cache.

In contrast, for NetWare, whenever a user connects to the NSS file system, NetWare retrieves the user's SEV from eDirectory and maintains it as part of the connection structure for the user's session. NSS retrieves the user's SEV from the connection structure.

Command line switches are available in the NSS Console utility (nsscon) to enable or disable the update, to set the update interval from 5 minutes to 90 days (specified in seconds), and to force an immediate update of security equivalence vectors. Polling too frequently can impact performance. Polling too infrequently can cause delays in granting or restricting access to certain users.

NSSCON Command	Description
<code>nss /SecurityEquivalenceUpdating</code>	<p>Enables SEV updates to occur in the background in addition to updates that occur when the system reboots. It is enabled by default.</p> <p>To make it persistent, include the command in the <code>/opt/novell/nss/conf/nssstart.cfg</code> file.</p>
<code>nss /NoSecurityEquivalenceUpdating</code>	<p>Disables SEV updates to occur in the background. Instead, SEV updates occur only on system reboot.</p> <p>To make it persistent, include the command in the <code>/opt/novell/nss/conf/nssstart.cfg</code> file.</p>
<code>nss /UpdateSecurityEquivalenceInterval=value</code>	<p>Sets the SEV update interval to the specified value in seconds. At the end of the elapsed time, NSS requires updated SEVs from eDirectory.</p> <p>To make it persistent, include the command in the <code>/opt/novell/nss/conf/nssstart.cfg</code> file.</p> <p>The default value is 7200 (2 hours). The valid range is 300 (5 minutes) to 7257600 (90 days).</p>

NSSCON Command	Description
<code>nss /ForceSecurityEquivalenceUpdate</code>	<p>Forces the SEV update to occur immediately for all users in the NSS file system. Use this command if you modify a user's access control settings in eDirectory and want those changes to be reflected immediately in the user's active SEV for this server.</p> <p>This command is invalid if used in the <code>/opt/novell/nss/conf/nssstart.cfg</code> file.</p> <p>A unique abbreviation such as</p> <p><code>nss /ForceS</code></p> <p>also works.</p>

# Monitoring the NSS File System Status

# 24

This section discusses the following methods for monitoring the status of Novell® Storage Services™:

- [Section 24.1, “Monitoring Status of NSS Devices, Pools, and Volumes with iManager,” on page 319](#)
- [Section 24.2, “Monitoring Compression and Salvage Statistics,” on page 320](#)
- [Section 24.3, “Monitoring Quotas and Space Usage for NSS Pools and Volumes,” on page 322](#)
- [Section 24.4, “Monitoring File System Parameters,” on page 324](#)

## 24.1 Monitoring Status of NSS Devices, Pools, and Volumes with iManager

- 1 Use the following table to determine where to go to view the status of your NSS storage devices, pools, and volumes.

To monitor the status of:	Refer to:
Devices	<a href="#">Section 7.4, “Viewing Details for a Device,” on page 102</a>
Partitioned free space	<a href="#">Section 7.8, “Viewing Partitions on a Device,” on page 106</a>
Partition	<a href="#">Section 8.4, “Viewing Details for a Partition,” on page 114</a>
Software RAIDs	<a href="#">Section 9.3, “Viewing Software RAID Devices on a Server,” on page 124</a>
Software RAID details	<a href="#">Section 9.4, “Viewing Details of a Software RAID Device,” on page 125</a>
NSS pools	<a href="#">Section 11.2, “Viewing Pools on a Server,” on page 150</a>
NSS pool details	<a href="#">Section 11.3, “Viewing Pool Details,” on page 151</a>
NSS volumes in a pool	<a href="#">Section 11.5, “Viewing Volume Information for a Pool,” on page 153</a>
NSS volume details	<a href="#">Section 13.6, “Viewing the Details of an NSS Volume,” on page 189</a>
NSS volume quota and space usage	<a href="#">Section 13.7, “Viewing Properties of an NSS Volume,” on page 189</a> , then click the Quotas tab

To monitor the status of:	Refer to:
NSS volume attributes	<a href="#">Section 13.7, “Viewing Properties of an NSS Volume,” on page 189</a> , then click the Attributes tab

## 24.2 Monitoring Compression and Salvage Statistics

For each volume, the Volume Properties Statistics page in iManager reports statistics about the compressed and salvageable files in the volume, the GUID of the volume, and the block size being used.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3 In the Volumes list, select a volume that you want manage.

- 4 Click Properties.

The Properties page has three tabs: Attributes, Statistics, and Quotas Usage. It opens by default to the Attributes tab.



- 5 Click the Statistics tab to view the compression and salvage statistics for the selected volume.

[Storage Management](#) ► [Volumes](#) ► **Volume Properties**

## Volume Properties



Properties: NIFVOL

Attributes		Statistics		Quotas Usage	
Compression		Salvage			
Compressed Space:		0.00 Bytes		Minimum Keep Seconds:	
Files:				Maximum Keep Seconds:	
Not Deleted:		0		Low Water Mark:	
Deleted:		0		High Water Mark:	
Uncompressed:		0			
				Next Scheduled Purge	
				Purgeable Space:	
				Unpurgeable Space:	
				Deleted Files:	
				Oldest Deleted Time:	

GUID: **A7942770-9C3E-01D7-80-00-1A1A3BF816F4**

Block Size: **4.00 KB**

Close

If the compression attribute is set, the Compression report shows statistics of all the compressed files for the selected volume.

Statistic	Description
<i>Compressed Space</i>	The amount of space in the volume in use by compressed files
<i>Files</i>	<p>The total number of files in the volume and information by the following categories:</p> <ul style="list-style-type: none"><li>• <b>Not Deleted:</b> The total number of files in the volume that are currently available to users</li><li>• <b>Deleted:</b> The total number of files in the volume that are deleted but not yet purged from the system</li><li>• <b>Uncompressed:</b> The total number of files in the volume that are not stored in compressed form</li></ul>

If the Salvage Files attribute is enabled, the Salvage report shows statistics about deleted files that have not yet been purged.

Statistic	Description
<i>Minimum Keep Seconds</i>	Minimum time (in seconds) to keep deleted files

Statistic	Description
<i>Maximum Keep Seconds</i>	Maximum time (in seconds) to keep deleted files
<i>Low Water Mark</i>	If the amount of free space drops below this percentage, the file system begins purging deleted files.
<i>High Water Mark</i>	If there are files to delete, the autopurging process stops when the amount of free space reaches this percentage.
<i>Next Scheduled Purge</i>	Date and time of the next purge
<i>Purgeable Space</i>	Amount of space in the volume that is occupied by deleted files that are queued for purging
<i>Unpurgeable Space</i>	Amount of space in the volume that is occupied by files
<i>Deleted Files</i>	The number of deleted files in salvage.
<i>Oldest Deleted Time</i>	Time line for deleted files. The file system purges the files in the same order they were deleted.
<i>GUID</i>	The Global Unique Identifier (GUID) number that NSS assigns to the volume. This number is necessary so your file system can locate the specific volume.
<i>Block Size</i>	The maximum amount of data committed to a single write process. Possible sizes include 4, 8, 16, 32, or 64 KB.

Other information reported includes the volume's GUID and block size.

Statistic	Description
<i>GUID</i>	The Global Unique Identifier (GUID) number that NSS assigns to the volume. This number is necessary so your file system can locate the specific volume.
<i>Block Size</i>	The maximum amount of data committed to a single write process. Possible sizes include 4, 8, 16, 32, or 64 KB.

## 24.3 Monitoring Quotas and Space Usage for NSS Pools and Volumes

For each volume, the Volume Properties Quotas Usage page in iManager reports the space usage for the selected volume and the pool that contains the volume.

- 1 In iManager, click *Storage > Volumes*.

For instructions, see [Section 4.1.4, “Accessing NSS Storage Management Role and Tasks,” on page 46](#).

- 2 Select a server to manage.

For instructions, see [Section 4.1.5, “Selecting a Server to Manage,” on page 46](#).

- 3 In the Volumes list, select a volume that you want manage.

- 4 Click Properties.

The Properties page has three tabs: Attributes, Statistics, and Quotas Usage. It opens by default to the Attributes tab.

- 5 Click the Quotas Usage tab to view the current space usage statistics for the selected volume.

[Storage Management](#) ► [Volumes](#) ► **Volume Properties**

## Volume Properties ?

**Properties:** NIFVOL

Attributes
Statistics
Quotas Usage

Volume Usage: NIFVOL

Pool Usage: NIFPOOL



Quota:	<b>none</b>	Available Space:	<b>89.93 MB</b>
		Free:	<b>89.93 MB</b>
Used Space:	<b>576.00 KB</b>	Purgeable:	<b>0.00 MB</b>
Compressed:	<b>0.00 Bytes</b>		
Other in-use space:	<b>576.00 KB</b>	Total Space:	<b>99.00 MB</b>
		Used:	<b>9.07 MB</b>
Purgeable Space:	<b>0.00 MB</b>	Booked:	<b>0.00 MB</b>
Available Space:	<b>89.93 MB</b>	This pool contains logical volumes with no quotas. They were not used to calculate the booking.	

Close

Volume Usage reports the amount of space on the volume, categorized by usage: Compressed, Other In-use, Purgeable, and Available.

Statistic	Description
<i>Quota</i>	Indicates whether the volume has a quota. If there is a quota, the volume can grow only to the size of the quota. If there is no quota, the volume can grow to the amount of available physical space in the pool.

Statistic	Description
<i>Used Space</i>	<p>The amount of space currently in use and information by the following categories:</p> <ul style="list-style-type: none"> <li>• <b>Compressed:</b> If the Compressed attribute is enabled, reports the amount of space in the volume containing data that is compressed.</li> <li>• <b>Other In-Use Space:</b> The amount of space in the volume containing data that is not compressed.</li> <li>• <b>Purgeable Space:</b> The amount of space in the Salvage system that you can use as free space. You can manually purge deleted files to free space.</li> <li>• <b>Available:</b> Available free space that is not in the Salvage system.</li> </ul>

Pool Usage reports the amount of space on the pool, categorized by usage: Free, Purgeable, Used, and Booked.

Statistic	Description
<i>Available Space</i>	<p>The amount of space in the pool that is not currently in use and information by the following categories:</p> <ul style="list-style-type: none"> <li>• <b>Free:</b> The total amount of free space that is available on the pool.</li> <li>• <b>Purgeable:</b> The amount of space in the Salvage system that you can use as free space. You can manually purge deleted files to free space.</li> </ul>
<i>Total Space</i>	<p>The total amount of space allocated to the pool and information by the following categories:</p> <ul style="list-style-type: none"> <li>• <b>Used:</b> The total amount of space currently in use by all volumes on the pool.</li> <li>• <b>Overbooked/Booked:</b> If the amount of space assigned to the pool's volumes exceeds the amount of physical space available in the pool, the Overbooked field shows the amount of exceeded space. Otherwise, the Booked field shows the total amount of space in all volumes in the pool. If any of the volumes do not have a quota, these volumes are not calculated in the total combined quota.</li> </ul>

## 24.4 Monitoring File System Parameters

- [Section 24.4.1, “Monitoring NSS File System Parameters with Novell Remote Manager,” on page 325](#)
- [Section 24.4.2, “Monitoring NSS File System Parameters with Novell NetStorage,” on page 325](#)

## 24.4.1 Monitoring NSS File System Parameters with Novell Remote Manager

Use Novell Remote Manager for NetWare to monitor the status of the NSS file system parameters listed in the following table. For information about using Novell Remote Manager, see the *Novell Remote Manager for NetWare Administration Guide for OES*.

Use this link	To do these tasks
Volumes	<ul style="list-style-type: none"><li>• Monitor disk space usage by directory</li><li>• Purge or salvage deleted volumes or files</li><li>• Configure file system trustee rights and attributes for directories and files</li><li>• View or change a list of Set parameters</li></ul>

You can browse directories and files in NSS Volumes on Linux, using Novell Remote Manager for Linux. For information, see “[Viewing File Systems](#)” in the *Novell Remote Manager Administration Guide for Linux for OES*.

## 24.4.2 Monitoring NSS File System Parameters with Novell NetStorage

Use the Novell NetStorage to monitor the status of the NSS file system parameters listed in the following table.

Use this link	To do these tasks
Volumes	<ul style="list-style-type: none"><li>• Monitor disk space usage by directory (<a href="http://www.novell.com/documentation/oes/netstor/data/bv3twro.html">http://www.novell.com/documentation/oes/netstor/data/bv3twro.html</a>)</li><li>• Purge or salvage deleted volumes or files</li><li>• Configure file and directory Trustee rights and Inherited Rights filters</li><li>• View or change a list of Set parameters</li></ul>



# Configuring the System Cache to Fine-Tune NSS Performance (NetWare)

# 25

Under certain conditions, the Novell® Storage Services™ performance can be tuned to great benefit without sacrificing NetWare's 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 Performance Monitoring and Tuning* (<http://support.novell.com/techcenter/articles/ana20020701.html>).

- [Section 25.1, “Do I Need to Tune NSS?,” on page 327](#)
- [Section 25.2, “NSS Tuning Guidelines,” on page 327](#)
- [Section 25.3, “Setting the Cache Balance,” on page 329](#)
- [Section 25.4, “Setting the Cache Buffers,” on page 333](#)
- [Section 25.5, “Setting the Closed File Cache Size,” on page 335](#)
- [Section 25.6, “Setting the File/Buffer Flush Timer,” on page 336](#)
- [Section 25.7, “Setting the Name Cache,” on page 336](#)

## 25.1 Do I Need to Tune NSS?

There are many factors that contribute to decreasing server performance; however, if 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 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://support.novell.com/techcenter/articles/ana20020701.html>).

After you have determined that your storage subsystem needs tuning, determine what NSS parameters can be adjusted to improve your current NSS performance. Before adjusting parameter settings, review the [Section 25.2, “NSS Tuning Guidelines,” on page 327](#).

## 25.2 NSS Tuning Guidelines

Consider the following before you adjust the cache management parameters:

- [Section 25.2.1, “Before You Adjust Cache Parameters,” on page 328](#)
- [Section 25.2.2, “Hints for Low Memory NetWare Systems,” on page 328](#)
- [Section 25.2.3, “The Tuning Process,” on page 328](#)
- [Section 25.2.4, “Restoring Settings After a Failed Tuning Attempt,” on page 329](#)

## 25.2.1 Before You Adjust Cache Parameters

- Make sure that you are running the latest version of NSS before adjusting parameters.
- Do not split your resources between two file systems. Use NSS volumes only on your server, not a combination of NSS volumes and NetWare Traditional volumes.

When the file system module (filesys.nlm) loads on NetWare, RAM that is not allocated elsewhere is given to the legacy cache. When NSS loads, it also takes about 60% of the RAM that is not allocated elsewhere for NSS. Because 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 might enable you to balance RAM between them more effectively.

- Purchase more memory, if needed to improve performance. If this is not an option, consider the tips in [Section 25.2.2, “Hints for Low Memory NetWare Systems,”](#) on page 328.

## 25.2.2 Hints for Low Memory NetWare Systems

If your server has limited RAM resources, you can set some parameters to minimize the amount of RAM used for NSS.

NetWare ship with NSS configured for good performance on servers with at least 1 GB of memory. However, some servers might have less than 1 GB of memory, or memory might be dedicated to other services.

Consider the following when you need to reduce the memory footprint of NSS:

**ClosedFileCacheSize:** The default value for ClosedFileCacheSize is 100000. Each entry uses about 1 KB of memory. On systems that only access a few files, this number can be reduced to 10000, which saves 90 MB of memory. This lower setting is not recommended for mail servers.

**MinBufferCacheSize:** Some administrators turn off cache balancing but do not change the default value for MinBufferCacheSize. The default setting gives NSS too few cache buffers to use and results in poor performance. Set the MinBufferCacheSize to at least 10000 when your system uses only NSS volumes. Set it below 10000 only if you have NetWare Traditional file systems on your server. Because NSS freely gives cache buffers to the system, playing with cache balancing only helps systems running out of virtual address space.

## 25.2.3 The Tuning Process

Tuning NSS is a matter of modifying only one or two parameters at a time with a gradual, minor adjustment. Reset the cache performance statistics, then monitor the server for performance gains or degradation.

---

**IMPORTANT:** If you assign more memory to NSS than the server has NSS fails to load. To recover, see [Section 25.2.4, “Restoring Settings After a Failed Tuning Attempt,”](#) on page 329.

---

Adjusting the following parameters can improve your NSS performance:

- [Section 25.3, “Setting the Cache Balance,”](#) on page 329
- [Section 25.4, “Setting the Cache Buffers,”](#) on page 333



- [Section 25.5, “Setting the Closed File Cache Size,” on page 335](#)
- [Section 25.6, “Setting the File/Buffer Flush Timer,” on page 336](#)
- [Section 25.7, “Setting the Name Cache,” on page 336](#)

## 25.2.4 Restoring Settings After a Failed Tuning Attempt

If you give NSS more memory than the server has, NSS fails to load. Because your `sys :` volume is NSS-based, your server is inoperable until you load the server with an NSS override to correct the situation.

If you need to restore a setting after a failed tuning attempt: At the server console, enter

```
server -z parameter
```

## 25.3 Setting the Cache Balance

The cache performance statistics are divided into system cache and user cache. System data is everything that is not user data. It includes the roll-back ZLOG (journal), 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
Set the amount of cache that NSS takes from the file system	<p>Every 30 seconds (by default) or whenever the server loads an NLM™, NSS rebalances the NSS file system cache to the percentage specified by the CacheBalance parameter. This command enables the dynamic balancing of free memory for the NSS buffer cache, and sets the specified percentage of free kernel memory that NSS uses for its buffer cache.</p> <hr/> <p><b>NOTE:</b> You can change the default time interval with the <a href="#">CacheBalanceTimer</a> command.</p> <hr/> <p>A low cache balance impedes the performance of NSS.</p> <p>[Default=85; Range=1 to 99]</p>	<code>nss /CacheBalance=value</code>
Control the division of the file system cache between user and system data	<p>Sets the specified maximum percentage of buffer cache that can be used for user data.</p> <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>[Default=80; Range=10 to 100]</p>	<code>nss /CacheUserMaxPercent=value</code>

Task	Purpose	Command
Set the number of cache buffers allocated for NSS	<p>This sets the minimum number of 4 KB cache buffers allocated for NSS file system access. Make sure this value is at least 10000 unless you also have NetWare traditional volumes on your system.</p> <p>[Default=512; Range=256 to memory size in KB/ 4 KB]</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 <b>MinOSBufferCacheSize</b> command to set a hard limit for the amount of RAM that can be used exclusively by the operating system.</p> <hr/> <p><b>TIP:</b> LRU Sitting Time is found in the Monitor utility under Disk Cache Utilization. Under normal circumstances, the server should be up for over 24 hours, so this statistic should never drop below 24 hours.</p> <hr/>	nss /MinBufferCacheSize=value
Set a minimum amount of RAM exclusively for the operating system	<p>If you allocate too much memory to your NSS cache, the server can run out of memory. To prevent this, set aside a minimum number of 4 KB cache buffers for the exclusive use of the server operating system. Recommended is 20000.</p> <p>[Default=256; Range=256 to memory size in KB/4 KB]</p>	nss /MinOSBufferCacheSize=value

Task	Purpose	Command
Set a maximum amount of memory to transfer to NSS during a cache balancing session	<p>Usually when the CacheBalance is recalculated, a small amount of RAM is either given or taken away from the NSS Cache. This parameter controls the maximum cache balance transfer (in the number of 4 KB cache buffers) given to the NSS file system cache per cache balance.</p> <p>If your server has applications running on it that can consume significant amounts of RAM for a short period of time, you might want to increase this parameter. Otherwise your server may take several minutes to get back to the correct amount of cache dictated by the CacheBalance parameter.</p> <p>This parameter does not affect memory requests made of NSS by the OS. For example, if NSS is asked to give back 5 MB of RAM, it will comply at the time of the request in one transfer.</p> <p>[Default=1024; Range=16 to memory size in KB/4 KB]</p>	<pre>nss / CacheBalanceMaxBuffersPerSession =value</pre>
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>	<pre>nss /CacheHit</pre>

Task	Purpose	Command
Determine how often NSS rebalances the amount of cache it uses	<p>This lets you determine how often (in seconds) NSS rebalances the amount of cache it uses according to the total amount of memory on the server.</p> <p>If you have difficulty loading NLM programs 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 ensures that the cache remains balanced. NSS monitors 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.</p> <p>[Default=30; Range=1 to 3600]</p>	<code>nss /CacheBalanceTimer=value</code>
Reset cache statistics	<p>This 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>

## 25.4 Setting the Cache Buffers

A cache buffer is 4 KB block of RAM that temporarily stores data. When you install NSS, it consumes 85% of the cache buffers by default.

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 memory size in KB/4 KB. Recommended is 20000.</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> <hr/> <p><b>TIP:</b> LRU Sitting Time is found in the Monitor utility under Disk Cache Utilization. Under normal circumstances, the server should be up for over 24 hours, so this statistic should never drop below 24 hours.</p> <hr/>	nss /MinBufferCacheSize=value
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 4 KB blocks) that is given to the NSS file system cache per cache balance. The default is 1024 (4 MB) per session. The range is 16 to memory size in KB/4 KB. Recommended is 20000.</p> <p>If your server has applications running on it that can consume significant amounts of RAM for a short period of time, you might want to increase this parameter.</p>	nss/ CacheBalanceMaxBufferPerSession= value

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.	<code>nss/NumHashBuckets</code>
Set the cache for file and directory entries	This sets the number of name cache entries.  [Default=100000; Range=17 to 1000000]	<code>nss/NameCacheSize=value</code>
Use percentages instead of integers for cache balancing	Enables the dynamic balancing of free memory for the NSS buffer cache. Sets the specified percentage of free kernel memory that NSS uses for its buffer cache. [Default=85; Range=1 to 99]  The default is On. Use this command in combination with the <code>/MinBufferCacheSize</code> switch.	<code>nss / (No)CacheBalance</code>
Set the number of authorization cache entries	This specifies the number of cache buffers to be used to speed up authorization requests.  If many trustees are set on different files and directories, we recommend that you increase this number to at least 20000.  [Default=1024; Range=16 to 50000]	<code>nss /AuthCacheSize=value</code>
Show cache buffer statistics	This shows the cache buffer statistics.	<code>nss /CacheStats</code>
Reset cache statistics	This lets you clear the current statistics and start over.  This is useful for checking the results of changes to the system.	<code>nss /reset</code>

## 25.5 Setting the Closed File Cache Size

The `ClosedFileCacheSize` parameter keeps the beast objects in cache so that NSS does not need to go to disk and unpack beast information again when it wants the same files. At the server console, enter

```
nss /ClosedFileCacheSize=value
```

The default is 100000. The range is 16 - 1000000. On average, each Closed File Cache entry consumes 0.4 - 1 KB of RAM.

Set this parameter to 100000 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.

---

**TIP:** The Closed File Cache grows 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.

---

## 25.6 Setting the File/Buffer Flush Timer

Increasing the flush time settings can improve performance, but does it at the risk of data loss during a system failure. The timers directly relate to the amount of uncommitted data in NSS Cache.

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. Increasing this number might reduce the number of writes to disk; however, it increases the amount of data that is lost if the system crashes.  Increasing the /FileFlushTimer setting beyond a minute or two can lead to overflowing the ZLOG file and throttling system operations.  [Default=10; Range=1 to 3600]	<code>nss /FileFlushTimer=value</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.  [Default=1; Range=1 to 3600]	<code>nss /BufferFlushTimer=value</code>

## 25.7 Setting the Name Cache

The NameCache parameter specifies the amount of recently used Name Tree entries that NSS caches. Each entry uses about 150 bytes of memory.

At the server console, enter

```
nss /NameCacheSize=value
```



[Default=100000]

---

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

---



# Configuring the System Cache to Fine-Tune NSS Performance (Linux)

# 26

This section discusses the following:

- [Section 26.1, “Understanding How NSS Uses Cache on Linux,” on page 339](#)
- [Section 26.2, “Setting the Minimum Number of Cache Buffers,” on page 339](#)

## 26.1 Understanding How NSS Uses Cache on Linux

The following table describes how Novell® Storage Services™ uses the cache on Linux.

OES Release	How NSS Uses the Cache Buffer
OES SP1 Linux	<p>Beginning with OES SP1 Linux, for file data, NSS uses the Linux cache page manager to gain access to all available memory in the system. There are some limits in place so that when copying large files, NSS does not starve other user applications for memory, which is also how it works in NetWare. For metadata, NSS is limited to a percentage of the kernel memory limit, which is 1 GB cache.</p> <p>You can adjust the percentage to be used by NSS with the <code>MinBufferCacheSize</code> parameter. The percentage allotted for OES SP1 can be lower because only metadata is stored in this portion of the cache, not both metadata and file data as it was in the initial release of OES.</p> <p>This is the same method used in other Linux file systems such as ReiserFS, Polyserve, XFS, with the exception of EXT.</p>
OES Linux	<p>NSS for OES Linux does not support cache balancing for NSS cache buffers. NSS does not use Linux HighMem. Cache management for NSS cache buffers is limited to a percentage of the cache that is assigned for kernel memory on Linux. The kernel takes a maximum of 1 GB.</p> <p>You can adjust the percentage to be used by NSS with the <code>MinBufferCacheSize</code> parameter.</p>

## 26.2 Setting the Minimum Number of Cache Buffers

To set the minimum number of cache buffers used by NSS on OES Linux, enter at the server console:

```
nss /MinBufferCacheSize=value
```

where *vaue* is the number of 4 KB buffers to assign for NSS. The maximum setting is the amount of memory in KB divided by 4 KB.



# Troubleshooting the NSS File System

# 27

This section describes some issues you might experience with Novell® Storage Services™ (NSS) and provides suggestions for resolving or avoiding them.

- Section 27.1, “Cannot See NSS Devices, Pools, or Volumes,” on page 341
- Section 27.2, “Cannot Create an NSS Pool or Volume,” on page 342
- Section 27.3, “eDirectory Error 613 Returned When Creating an NSS Pool or Volume on Linux,” on page 343
- Section 27.4, “NSS Volumes on Linux with Large Directories Mount Slowly with the UNIX Name Space,” on page 343
- Section 27.5, “Linux Fails to Execute Dismount, Umount, or Delete Commands for NSS Volumes,” on page 343
- Section 27.6, “Slow Performance on SCSI Devices and RAID Controllers (Linux),” on page 344
- Section 27.7, “File Compression Is Not Working,” on page 345
- Section 27.8, “Software RAID 1 Fails to Recognize a Replacement Device,” on page 345
- Section 27.9, “Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers,” on page 345
- Section 27.10, “Tuning NSS Volumes for GroupWise Servers,” on page 345
- Section 27.11, “Problems Following DFS Junctions in Windows 2000/XP Releases (NetWare),” on page 346

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

## 27.1 Cannot See NSS Devices, Pools, or Volumes

If you cannot see your volumes or the devices associated with those volumes, you might have a connection failure. Connection failures can occur if an adapter, cable, or switch in the path between the server and the storage device fails for any reason. If there is a connection failure, repair or reconfigure the equipment.

### Linux

For OES Linux, if the Storage plug-in for iManager or NSSMU for Linux do not recognize a device, it might be because NSS recognizes only what the Enterprise Volume Management System (EVMS) recognizes. For this release, NSS recognizes hard drives, CD drives, DVD drives, USB drives, and Zip\* drives that are managed by EVMS. If your hard drive is managed by Linux Volume Management (LVM), it cannot be managed by EVMS, and NSS does not recognize it.

For OES Linux, whenever you reboot your system, EVMS discovers and mounts partitions under `/dev/evms/`. The iManager Storage plug-in and NSSMU work through EVMS to manage partitions, and this behavior is expected. However, if you create and mount partitions from the command line, your partitions might appear to be lost after a reboot and EVMS discovery. For example, if you use commands to create and mount a partition at `/dev/sdb2/`, after the reboot, EVMS discovers and mounts the partition at `/dev/evms/sdb2`.

---

**IMPORTANT:** If you mount pools from the command line, you must modify your mount point to the partition's new apparent location under `/dev/evms/`.

---

## NetWare

If you have multiple interconnect paths in your server-to-storage configuration, turn on Multiple Connection Path Support. For instructions, see [“Managing Multiple Connection Paths to Devices \(NetWare\)” on page 139](#). Then repair the equipment.

In OES, NSS does not provide Multiple Connection Paths Support for OES Linux. Your system does not have the appropriate multipath drivers.

## 27.2 Cannot Create an NSS Pool or Volume

If you cannot create an NSS pool or volume, consider these possible causes:

- Your server might not have enough free space to create more storage pools or NSS volumes. Ensure you have enough free space to create another NSS storage pool or NSS volume. See [Section 7.8, “Viewing Partitions on a Device,” on page 106](#).
- Each NSS volume must be part of an NSS pool. If you cannot create the pool at the same time as the volume, try creating a pool to use first.

## Linux

If you cannot create a pool on a device you can see in NSSMU, NSS might not own the free space you want to use for a storage pool. If the device you want to use is managed by LVM instead of EVMS, NSS does not recognize the device as being available and does not create the pool.

SLES 9 uses LVM, but NSS requires EVMS. The Linux 2.6 kernel prevents multiple volume managers from managing the same device. As a result, EVMS might display free space that it cannot actually manage.

To prevent display of space that is unavailable to EVMS, edit the `/etc/evms.conf` file to add the device to the exclude list of the `sysfs_devices` section.

For example, if your `hda` device is your boot device and contains the root volume, it cannot be used to create pools on. Edit the `/etc/evms.conf` file as follows:

```
sysfs_devices {  
    ...  
    exclude = [hda]  
}
```

To make a device that is being managed by another device manager available to EVMS, you can use one of the following options:

- Install SLES 9 with EVMS as the default volume manager

See [Mounting Your Root File System through EVMS \(http://evms.sourceforge.net/install/root.html\)](http://evms.sourceforge.net/install/root.html) for more information.

- Patch the 2.6 kernel

See [BD-Claim Patch \(http://evms.sourceforge.net/install/kernel.html#bdclaim\)](http://evms.sourceforge.net/install/kernel.html#bdclaim) for more information.

- Convert LVM volumes to EVMS volumes after SLES installation

See [Converting to EVMS \(http://evms.sourceforge.net/convert.html\)](http://evms.sourceforge.net/convert.html) for more information.

## NetWare

If you cannot create an NSS volume on your NetWare<sup>®</sup> server, you might not be able to create and mount an NSS volume if you exceed the maximum number of 256 volumes mounted at a time on the server.

## 27.3 eDirectory Error 613 Returned When Creating an NSS Pool or Volume on Linux

When creating an NSS pool or volume with NSSMU on your Linux server, an Error 613 is returned if the server has no eDirectory<sup>™</sup> Read/Write replica. The error occurs because NCP<sup>™</sup> cannot map to the pool or volume.

To avoid this problem, make sure the server has an eDirectory Read/Write replica. You can also add the NSS volume path to the ncpserv.conf file.

## 27.4 NSS Volumes on Linux with Large Directories Mount Slowly with the UNIX Name Space

Normally, NSS volumes mount in seconds, even if the volume contains large directories with millions of files. You might observe that an NSS volume on Linux mounts slowly if it contains large directories and if you are mounting it with a UNIX name space.

To avoid this problem, mount NSS volumes on Linux with a Long name space. For information, see [Section 13.16, “Mounting an NSS Volume on Linux to Support Case Insensitive File Names \(Linux\),” on page 197.](#)

## 27.5 Linux Fails to Execute Dismount, Umount, or Delete Commands for NSS Volumes

When NCP Server is active, it always keeps a file open on a volume. As a result, issuing Dismount, Umount, or Delete commands for an NSS volume fails, whether the command is issued from the command line, in NSSMU, or in iManager.

This occurs because Linux does not allow you to dismount a volume if a file is open on that volume. Linux does not provide a method of identifying which files are open on volumes.

To dismount or delete an NSS volume on your OES Linux server:

- 1 At the server prompt, open the NCP Console by entering

```
ncpcon
```

- 2 Dismount the volume from NCP.

The volume is no longer accessible or visible to NCP clients.

- 3 In iManager, dismount or delete the volume, as desired.

## 27.6 Slow Performance on SCSI Devices and RAID Controllers (Linux)

If you plan to use space from SCSI devices—local or in a RAID controller—for your NSS pools and volumes on Linux, make sure the devices and controller use Write-Through cache management, not Write-Back. Using Write-Back cache management on SCSI drives with a journaled file system such as NSS can actually degrade file system performance instead of improving it.

Write-Through and Write-Back (also known as Write-Behind) are cache management schemes that can affect the manner in which data is written to the hard drives. Although Write-Back can sometimes improve performance, it does so by allowing data to be held in cache rather than being written to disk. Data loss can occur if the server fails before the cached data is written to disk, such as with a power failure.

To minimize the potential for data loss when Write-Back is being used, a journaled file system might alter its manner of committing writes to disk. Instead of improving performance, Write-Back cache management actually creates work for the journaled file system, and the additional workload can degrade file system performance. It seems counter-intuitive, but this additional effort is necessary to help ensure data integrity. Write-Through cache management assures the file system that writes are being committed to disk as intended. The additional effort to protect against potential data loss associated with Write-Back is no longer required, and the file system performance is not degraded.

If the server uses a RAID controller, when configuring the RAID device, enable Write-Through (disable Write-Back) cache management, using the controller's BIOS setup routine or configuration utility.

To enable Write-Through cache management for local devices:

- 1 Log in to the server as root.
- 2 In YaST, open the *Various Linux Tools* section.
- 3 Install the `xscsi` RPM package, then close YaST.

The `xscsi` package installs the `scsi-config` utility in `/usr/bin/scsi-config`.

- 4 In a terminal console, enter

```
scsi-config
```

- 5 Enable Write-Through (disable Write-Back) cache management for each SCSI device.



## 27.7 File Compression Is Not Working

If you cannot compress a file, you might not have set the Compression attribute when you created your NSS volume. To apply the file compression option to an existing NSS volume: From iManager, click *Storage > Volumes > Properties > Attributes*, then select *Compression*.

For information, see [“Managing Compression on NSS Volumes” on page 227](#).

## 27.8 Software RAID 1 Fails to Recognize a Replacement Device

If a drive fails that is part of a software RAID 1 device, your system might not recognize the replacement drive and does not begin remirroring automatically.

To recover:

- 1 Reboot the system to rescan devices and recognize the replacement drive.
- 2 Unmount the NSS volumes on the software RAID device to stop the I/O.
- 3 Expand the software RAID 1 device and specify space from the replacement drive as a new RAID segment. Remirroring begins automatically.
- 4 Allow the remirror to complete before mounting the NSS volumes.

## 27.9 Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers

NetWare 6.5 data pools are normally compatible for clustering purposes with NetWare 6.0, both backwards and forwards. The Media Manager changed between NetWare 6.0 and 6.5, but the media format is the same.

If you move a device from the NetWare 6.5 server to a NetWare 6.0 server, and if you can see the NetWare 6.5 pool on the disk, but cannot activate it on your NetWare 6.0 server, the most likely causes might be a name conflict or that the pool was mirrored.

Pool and volume names must be unique on the server. Renaming the NetWare 6.5 pool should allow it to be activated on the NetWare 6 server.

If the pool was mirrored, the old media manager on NetWare 6.0 probably would have problems with it because the information is different. You should remove the mirror before trying to move the disk to NetWare 6.0.

## 27.10 Tuning NSS Volumes for GroupWise Servers

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

```
nss /NoSalvage
```

## 27.11 Problems Following DFS Junctions in Windows 2000/XP Releases (NetWare)

Windows 2000 Service Pack 4 and Windows XP Service Pack 2 might have problems following DFS junctions over CIFS because of a defect in Windows. (This problem exhibits itself in a pure Windows environment.)

When using DFS with NetWare CIFS, the CIFS server and Windows clients are on different IP subnets. In this case, it is required that the client have a way to resolve the CIFS server name in order for DFS to work. This is a Microsoft/CIFS requirement not a NetWare CIFS requirement. There are multiple ways the client can do this, including 1) both client and NW server are configured for the same WINS server, 2) both client/server using same DNS server, or 3) the lmhosts workaround described in this topic.

To avoid this problem for your users with these releases of Windows, install the Novell Client on their desktops.

Another option is to modify the lmhosts file for all client computers with appropriate entries for any volumes on OES NetWare servers that use DFS junctions.

**1** In a text editor, open the lmhosts file.

- **Windows 2000:** `c:\WINNT\system32\drivers\etc\lmhosts`
- **Windows XP:** `c:\windows\system32\drivers\etc\lmhosts`

If you do not have an lmhosts file, use the `lmhosts.sam` file as an example.

**2** For each volume the user connects to that has a DFS junction, add a line at the end of the file that identifies the IP address and NetBIOS name of the data server where the volume resides.

```
192.168.1.1      servername-W
```

Replace `192.168.1.1` with the actual IP address and `servername` with the host name of your server.

For example, suppose you have the following server with VOL1 that contains one or more DFS junctions:

- Server IP address: `10.10.1.1`
- Server name: `USERSVR`
- NetBIOS server name: `USERSVR-W`
- Volume name: `VOL1`

The line you add to the lmhosts file would be:

```
10.10.1.1 USERSVR-W
```

**3** Save and close the lmhosts file.

**4** If necessary, repeat **Step 1** to **Step 3** on each client computer, or create an lmhosts file and distribute it to the client machines.

**5** On each client, map a network drive to the user's data volume.

Continuing the example above, the user could map to `\\10.10.1.1\VOL1` or to `\\USERSVR-W\VOL1`.

**5a** In the Windows Explorer file manager, click *Tools > Map Network Drive*.

**5b** In Folder, type one of the following:

`\\192.168.1.1\volumename`

`\\servername-W\volumename`

Replace `192.168.1.1` with the actual IP address or `servername` with the host name of your server.

**5c** Select *Reconnect at Logon*.

**5d** Click *Finish*.



This section discusses the following security issues for NSS and is intended for security administrators:

- [Section 28.1, “Security Features of NSS,” on page 349](#)
- [Section 28.2, “Preventing Unauthorized Access,” on page 350](#)
- [Section 28.3, “Protecting Data During Backup and on Backup Media,” on page 351](#)
- [Section 28.4, “Preventing Exposure of Sensitive Data in a Core Dump,” on page 351](#)
- [Section 28.5, “Preventing Exposure of the Encryption Password in a Log \(Linux\),” on page 351](#)
- [Section 28.6, “Using Data Shredding to Prevent Unauthorized Access to Purged Deleted Files,” on page 355](#)
- [Section 28.7, “Protecting the Admin Volume and Files.cmd \(NetWare\),” on page 355](#)
- [Section 28.8, “Acquiring Security Equivalence Vectors for NSS Users,” on page 356](#)
- [Section 28.9, “Protecting Modules Responsible for Security Equivalence Vectors \(Linux\),” on page 356](#)
- [Section 28.10, “Controlling File System Access and Attributes for NSS Volumes,” on page 356](#)
- [Section 28.11, “Displaying Directory and File Attributes for NSS Volumes on Linux,” on page 357](#)
- [Section 28.12, “Controlling Physical Access to Servers and Resources,” on page 357](#)
- [Section 28.13, “Securing Access to the Servers with a Firewall,” on page 357](#)
- [Section 28.14, “Creating Strong Passwords,” on page 357](#)

## 28.1 Security Features of NSS

Issue/Feature	Description	Recommendation	For More Information
Encrypted volume support	Encrypted NSS volumes meet the legal standard of making data inaccessible to software that circumvents normal access control, such as if the media were stolen.	Encrypt data volumes that contain mission critical data or sensitive data.  Use a strong encryption password and protect the password.	<a href="#">“Managing Encrypted NSS Volumes” on page 201</a>  <a href="#">Section 28.14, “Creating Strong Passwords,” on page 357</a>

Issue/Feature	Description	Recommendation	For More Information
Storage plug-in for iManager 2.5	iManager requires eDirectory authentication and SSL connections between your Web browser and the iManager server and between the iManager server and the target server being managed.	Use an Admin user identity and a strong password.	<p><a href="http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html">Novell iManager 2.5 Administration Guide (http://www.novell.com/documentation/imanager25/imanager_admin_25/data/hk42s9ot.html)</a></p> <p>Section 28.14, “Creating Strong Passwords,” on page 357</p>
NSSMU	NSSMU is executed locally on the server.	Secure physical access to servers.	<p>Section 28.12, “Controlling Physical Access to Servers and Resources,” on page 357</p>
Novell Remote Manager	Novell Remote Manager requires eDirectory authentication and SSL connections between your Web browser and the Novell Remote Manager running on the server being managed.	Use an Admin user identity or equivalent, and use a strong password.	<p>Section 4.5.3, “Accessing Novell Remote Manager,” on page 60</p> <p>Also, see the following:</p> <ul style="list-style-type: none"> <li>• Section 28.14, “Creating Strong Passwords,” on page 357</li> <li>• <i>Novell Remote Manager Administration Guide for Linux for OES</i></li> <li>• <i>Novell Remote Manager for NetWare Administration Guide for OES</i></li> </ul>

## 28.2 Preventing Unauthorized Access

NSS includes the following features to help prevent access to data that circumvents normal access control:

- Encrypted Volume Support

Encrypt data volumes, meeting U.S. Government security standards. For information, see [“Managing Encrypted NSS Volumes” on page 201](#).

- Data shredding (up to 7 times) for deleted files

Erase files completely, meeting U.S. Government security standards. For information, see [Section 15.2, “Using Data Shredding to Prevent Access to Deleted Files,” on page 215](#).

- Multiple server access prevention for pools

Multiple Server Access Prevention ensures data integrity by preventing unauthorized access to shared media in a storage area network. For information, see [Section 11.14, “Preventing Pools from Activating on Multiple Servers,” on page 164](#).

- Trustee model for access control on NSS volumes

NSS uses the Novell Trustee Model to greatly simplify access control management in the file system. It restricts visibility of data structures so that users only see subdirectories they have rights to see, not the whole tree like all other file systems.

For information about the Novell Trustee Model and NSS file system rights, see the *File Systems Management Guide for OES*.

## 28.3 Protecting Data During Backup and on Backup Media

Backups of NSS volumes are not encrypted, unless it is a feature of the backup software or hardware you use. Although data is stored encrypted on an encrypted NSS volume, its data is transmitted and backed up in unencrypted format.

Use backup methods that protect data transmitted between the server and the backup media, according to your security needs.

Use one of the following methods to encrypt the data for backup:

- Use backup software that is able to encrypt data when you back it up. This method has performance and manageability challenges, especially for managing encryption keys.
- Use an encryption appliance that encrypts sensitive backup media as data is backed up.

If you transport and store media offsite, use a company that specializes in media shipment and storage. This way, your tapes are tracked via barcodes, stored in environmentally friendly conditions, and are handled by a company whose reputation rests on its ability to handle your media properly.

## 28.4 Preventing Exposure of Sensitive Data in a Core Dump

When a core dump occurs for an encrypted NSS volume, data from the encrypted volume might be included in the core dump as unencrypted data dumped from cache memory. To prevent encrypted data exposure in the dump, when prompted to select writing all of memory in the core dump or to exclude NSS cache, select to exclude cache during a core dump.

This applies also for volumes that are not encrypted but contain confidential data. While the data is not normally encrypted, you might not want to allow unauthenticated access to the information.

## 28.5 Preventing Exposure of the Encryption Password in a Log (Linux)

This section discusses NSS debugger logging features so you can identify when these logs are turned on and turn them off in your operational environment.

- [Section 28.5.1, “NSS Logging,” on page 352](#)

- [Section 28.5.2, “NSS Logging and Security Implications,” on page 352](#)
- [Section 28.5.3, “Logging Communications between NSS and the \\_ADMIN Volume,” on page 352](#)
- [Section 28.5.4, “Logging Communications between NSS and eDirectory, NICI, or Linux User Management,” on page 354](#)

## 28.5.1 NSS Logging

On OES Linux, most of the NSS code runs in kernel space, but some portions are required to run in user space. To communicate across the boundary between user and kernel space, some internal mechanisms were implemented. For debugging purposes, some logging features were added to track these communications between user and kernel space. These logging features are slow and cumbersome, and are intended for use by Novell support engineers to help diagnose any problems that arise. They are not intended for everyday use, and seriously impact performance when they are turned on.

There are two main areas where logging is built into the system. The first is the capacity to log all XML communication to/from the \_ADMIN volume. The second is the capacity to log NSS kernel requests to communicate with eDirectory, NICI, and LUM, all of which run in user space.

## 28.5.2 NSS Logging and Security Implications

When working with encrypted volumes on Linux, it is important to realize that the volume password and key information is exchanged between user and kernel space as encrypted volumes are created and/or mounted. If you have logging enabled on the Linux server when you enter the encryption password, your password and volume key information might show up in log file.

You must be the root user or an equivalent user with root user privileges to perform the steps required to enable logging, disable logging, or read /var/log/messages. This prevents ordinary users from manipulating the logging environment. We strongly recommend that you protect the physical access to the server and the root user passwords to prevent unauthorized access to your servers.

Even though the logging mechanisms are root user protected, we strongly recommend that you make sure logging is disabled whenever you plan to enter the encryption password for an encrypted NSS volume on your system. You enter an encryption password when you create the volume and when you mount the volume for the first time after any system start or reboot.

## 28.5.3 Logging Communications between NSS and the \_ADMIN Volume

Applications such as NSSMU and Perl\* scripts communicate with NSS via the \_ADMIN volume. In these communications, the volume's encryption password is passed in the clear. There are two utilities that can log these exchanges, the adminusd daemon and the nss /vfs commands in NSSCON. Logs are written to /var/log/messages.

### Prerequisite

You must be the root user or an equivalent user with root user privileges to perform the steps required to enable logging, disable logging, or read /var/log/messages. This prevents ordinary users from manipulating the logging environment.



## Enabling or Disabling adminusd Logging

On your OES Linux server, an NSS daemon called `adminusd` is installed into `/opt/novell/nss/sbin` directory. It is run from the `startnss.bsh` script. Output data is written to the `/var/log/messages` directory.

### Enable adminusd Logging

At a Linux terminal console, do the following to enable `adminusd` logging:

- 1 Log in as the root user.
- 2 Kill the `adminusd` daemon.
- 3 Run the daemon with logging turned on by entering

```
adminusd -l
```

Using the `-l` option enables logging of all communication to and from the `_ADMIN` volume in the `/var/log/messages` directory.

### Disable adminusd Logging

At a Linux terminal console, do the following to disable `adminusd` logging:

- 1 Log in as the root user.
- 2 Kill the `adminusd` daemon.
- 3 Run the daemon with logging turned off by entering  

```
adminusd
```

Not using the `-l` option turns logging off.
- 4 Delete and purge the `adminusd` log files in `/var/log/messages`.

## Enabling or Disabling VFS Logging

In the NSS Console (NSSCON), the VFS option for NSS can log communications between NSS and the `_ADMIN` volume. The logged data is displayed on the NSSCON screen and is also written to the `var/log/messages` directory.

### Enabling VFS Logging

At a Linux terminal console, do the following to enable VFS logging:

- 1 Log in as the root user, then enter

```
nsscon
```

- 2 In NSSCON, enter

```
nss /vfs
```

Logging is turned on.

### Disabling VFS Logging

At a Linux terminal console, do the following to disable VFS logging:

- 1 Log in as the root user, then enter

```
nsscon
```

- 2 In NSSCON, enter

```
nss /novfs
```

Logging is turned off.

- 3 Exit NSSCON.

- 4 If the terminal console logging feature was on, turn it off, then delete and purge the logged session.

- 5 Delete and purge the VFS log files in `/var/log/messages`.

## 28.5.4 Logging Communications between NSS and eDirectory, NICI, or Linux User Management

All internal NSS kernel space requests for Novell eDirectory™, NICI, and Linux User Management are routed through an interface called the NDP (Novell Data Portal). NDP has a user space daemon (`ndpapp`) and a kernel module (`ndpmod`). In communications between `ndpapp` and `ndpmod`, the volume's encryption password is obscured, but it can be easily broken. Both `ndpapp` and `ndpmod` have a logging capacity, and both of them write their log data to `/var/log/messages`.

### Prerequisite

You must be the root user or an equivalent user with root user privileges to perform the steps required to enable logging, disable logging, or read `/var/log/messages`. This prevents ordinary users from manipulating the logging environment.

### Enabling or Disabling `ndpapp` Logging

On your OES Linux server, an NSS daemon called `ndpapp` is installed into `/opt/novell/nss/sbin` directory. It is run from the `startnss.bsh` script.

### Enable `ndpapp` Logging

At a Linux terminal console, do the following to enable `ndpapp` logging:

- 1 Log in as the root user.
- 2 Kill the `ndpapp` daemon.
- 3 Run the daemon with logging turned on by entering

```
ndpapp --debug=nn
```

Replace `nn` with the log level desired. Set the log level to 1 and above to turn logging on. The higher the number, the greater and more detailed is the logged output.

### Disable `ndpapp` Logging

At a Linux terminal console, do the following to disable `ndpapp` logging:

- 1 Log in as the root user.
- 2 Kill the `ndpapp` daemon.
- 3 Run the daemon with logging turned off by entering

```
ndpapp
```

Running `ndpapp` without the `--debug` option turns logging off.

- 4 Delete and purge the log files in `/var/log/messages`.

## Enabling or Disabling `ndpmod` Logging

### Enabling `ndpmod` Logging

At a Linux terminal console, do the following to enable `ndpmod` logging:

- 1 Log in as the root user, then enter

```
echo nn >/proc/driver/ndp/debug
```

Replace `nn` with the log level desired. Set the log level to 1 and above to turn logging on. The higher the number, the greater and more detailed is the logged output.

### Disabling `ndpmod` Logging

At a Linux terminal console, do the following to disable `ndpmod` logging:

- 1 Log in as the root user, then enter

```
echo 0 >/proc/driver/ndp/debug
```

Setting the Log Level field to 0 turns logging off.

- 2 Delete and purge the log files in `/var/log/messages`.

## 28.6 Using Data Shredding to Prevent Unauthorized Access to Purged Deleted Files

If the Data Shredding attribute for an NSS volume is disabled, unauthorized access to purged deleted files is possible. An individual can extend a file, `LSEEK` to the end of the existing file data, and then read the data. This returns the decrypted leftover data that is in the block.

To secure this vulnerability, make sure to enable Data Shredding for your NSS volumes by specifying an integer value of 1 to 7 times for the Data Shredding attribute. A value of 0 disables Data Shredding.

For information, see [Section 15.2, “Using Data Shredding to Prevent Access to Deleted Files,” on page 215](#).

## 28.7 Protecting the Admin Volume and `Files.cmd` (NetWare)

The `_admin:\manage_nss\files.cmd` file is visible to users, but they cannot delete this file nor modify its properties. You should set trustee rights to prevent users from seeing this volume. For information, see [Section 15.1, “Controlling Access to the Admin Volume and `Files.cmd` \(NetWare\),” on page 211](#).

## 28.8 Acquiring Security Equivalence Vectors for NSS Users

When a user authenticates to the network, the system calculates the user's Security Equivalence Vector (SEV) based on information in the user's profile in Novell eDirectory. NSS validates the user's SEV against the file system trustee rights of the directory and file the user is attempting to access. In OES, SEVs are acquired differently for NSS on NetWare and NSS on Linux.

For NSS on NetWare, whenever a user connects to the NSS file system, NetWare retrieves the user's SEV from eDirectory and maintains it as part of the connection structure for the user's session. NSS automatically retrieves the user's SEV from the connection structure.

For NSS on Linux, NSS caches the SEV locally in the server memory, where it remains until the server is rebooted or unless the user is deleted from eDirectory. NSS polls eDirectory at a specified interval for updates to the SEVs that are in cache. Command line switches are available in the NSS Console utility (nsscon) to enable or disable the update, to set the update interval from 5 minutes to 90 days (specified in seconds), and to force an immediate update of security equivalence vectors. For information, see [Section 23.19, "Security Equivalence Vector Update Commands \(Linux\)," on page 316](#).

## 28.9 Protecting Modules Responsible for Security Equivalence Vectors (Linux)

The Linux modules in user space that are responsible for providing Security Equivalence Vectors for NSS users can be replaced without the kernel module being aware of it. Make sure that the directory /opt/novell/nss/sbin/ and the files involved (ndpapp and idbrokerd) can only be modified by root. For example, make root the owner and set permissions to restrict access for Group and Other users.

## 28.10 Controlling File System Access and Attributes for NSS Volumes

To ensure that users have the appropriate effective file system rights to data on NSS volumes, make explicit file system trustee assignments, grant security equivalences for users, and filter inherited rights. To simplify the assignment of rights, you can create Group and Organizational Role objects in Novell eDirectory, then assign users to the groups and roles.

Set file system attributes for directories and files on an NSS volume to specify how a file or directory is used.

For information about controlling file system access and attributes for NSS volumes, see ["Understanding File System Access Control for NSS and NetWare Traditional File Systems"](#) in the *File Systems Management Guide for OES*.

For information about access control issues for NSS on Linux, see [Section 6.4, "Access Control Issues for NSS on OES Linux," on page 74](#).

## 28.11 Displaying Directory and File Attributes for NSS Volumes on Linux

With NCP Server, NSS on Linux supports the Novell Trustee Model, which is the same file system trustee rights and attributes for its directories and files as does NSS on NetWare. Management tools provide similar methods on each platform for configuring rights and attributes. For information, see the *File Systems Management Guide for OES*.

NSS on Linux displays some of the NSS file system directory and file attributes in the Linux POSIX directory and file permissions, including the Hidden, Read Only, Read/Write, and Execute attributes. These are not intended as a direct mapping of POSIX rights and behave differently. NSS does not support the POSIX set-user-ID mode bit and set-group-ID mode bit. For information, see “[Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions](#)” in the *File Systems Management Guide for OES*.

## 28.12 Controlling Physical Access to Servers and Resources

- Servers must be kept in a physically secure location with access by authorized personnel only.
- The corporate network must be physically secured against eavesdropping or packet sniffing.

## 28.13 Securing Access to the Servers with a Firewall

Use firewalls between public access points and servers to prevent direct access to data by a would-be third-party intruders.

## 28.14 Creating Strong Passwords

Make sure to employ security best practices for passwords, such as the following:

- **Length:** The minimum recommended length is 6 characters. A secure password is at least 8 characters; longer passwords are better.
- **Complexity:** A secure password contains a mix of letters and numbers. It should contain both uppercase and lowercase letters and at least one numeric character. Adding numbers to passwords, especially when added to the middle and not just at the beginning or the end, can enhance password strength. Special characters such as &, \$, and > can greatly improve the strength of a password.

Do not use recognizable words, such as proper names or words from a dictionary, even if they are bookended with numbers. Do not use personal information, such as phone numbers, birth dates, anniversary dates, addresses, or zip codes. Do not invert recognizable information; inverting bad passwords does not make them more secure.

- **Uniqueness:** Do not use the same passwords for all servers. Make sure to use separate passwords for each server so that if one server is compromised, all of your servers are not immediately at risk.



This section details the syntax and options for the following Novell® Storage Services™ utilities:

- [Section A.1, “ABORT REMIRROR \(NetWare\),” on page 359](#)
- [Section A.2, “ATTRIB \(Linux\),” on page 360](#)
- [Section A.3, “CDDVD \(NetWare\),” on page 362](#)
- [Section A.4, “COMPFIX,” on page 362](#)
- [Section A.5, “DISMOUNT,” on page 368](#)
- [Section A.6, “FILER \(NetWare\),” on page 368](#)
- [Section A.7, “FLAG \(NetWare\),” on page 369](#)
- [Section A.8, “METAMIG \(Linux\),” on page 374](#)
- [Section A.9, “MIRROR STATUS \(Netware\),” on page 376](#)
- [Section A.10, “MOUNT,” on page 376](#)
- [Section A.11, “NCOPY \(NetWare\),” on page 377](#)
- [Section A.12, “NDIR \(NetWare\),” on page 378](#)
- [Section A.13, “RAVSUI \(Linux\),” on page 378](#)
- [Section A.14, “RAVVIEW \(Linux\),” on page 380](#)
- [Section A.15, “REBUILD \(NetWare\),” on page 381](#)
- [Section A.16, “REMIRROR PARTITION \(NetWare\),” on page 382](#)
- [Section A.17, “RIGHTS,” on page 383](#)
- [Section A.18, “SET \(NetWare\),” on page 388](#)
- [Section A.19, “VERIFY \(NetWare\),” on page 396](#)
- [Section A.20, “VOLUME,” on page 397](#)

## A.1 ABORT REMIRROR (NetWare)

Use at the server console to stop remirroring a logical partition. Abort Remirror can be used on any NetWare mirror, regardless of type as long as all the partition types in the mirror match. This includes NSS, TRAD, SBD, and iSCSI devices and any other NetWare partition types that are mirrored.

### A.1.1 Syntax

```
ABORT REMIRROR logical_partition_number
```

## A.1.2 Option

Option	Description
<i>logical_partition_number</i>	Specify the number of the logical partition you want to stop remirroring.

## A.1.3 Example

To stop remirroring logical partition 0x2, enter

```
ABORT REMIRROR 2
```

## A.2 ATTRIB (Linux)

Use the Attribute (ATTRIB) utility to set NSS file system directory and file attributes on Linux.

**IMPORTANT:** This utility works only with directories and files in the NSS file system on Linux.

### A.2.1 Syntax

```
attrib [options] [filename]
```

If both the set and clear options are selected, the clear option is completed before the set option. If the filename is not specified, the operation is completed on the current directory.

### A.2.2 Options

Option	Description
-s, --set=ATTRIBUTES	Sets the attributes on the file.
-c, --clear=[ATTRIBUTES   all]	Clears the attributes on the file.
-l, --long	Displays a long version of the file attributes.
-q, --quiet	Does not display any normal output.
-v, --version	Displays the program version information.
-h, --help	Displays the ATTRIB help screen.

### A.2.3 Attributes

Multiple attributes are separated with commas.



Attribute	Description	Applies to Files	Applies to Directories
aa	Attribute Archive identifies that a file's metadata has been modified since the last backup. This attribute is assigned automatically.	Yes	No
all	All (used only for the Clear option) represents all attributes that can be modified.	Yes	Yes
ar	Archive identifies files that have modified content since the last backup. This attribute is assigned automatically.	Yes	No
cc	Cannot Compress (status display only) displays if the file cannot be compressed because of limited space savings.	Yes	No
ci	Copy Inhibit prevents users from copying a file. This attribute overrides the Read and File Scan trustee rights.	Yes	No
cm	Compressed (status display only) indicates whether the file is currently stored in compressed format.	Yes	No
dc	Don't Compress keeps data from being compressed. This attribute overrides settings for automatic compression of files not accessed within a specified number of days.	Yes	No
di	Delete Inhibit prevents users from deleting a directory or file.  This attribute overrides the Erase trustee right. When it is enabled, no one, including the owner and network administrator, can delete the directory or file. A trustee with the Modify right must disable this right to allow the directory or file to be deleted.	Yes	Yes
ex	Execute indicates program files, such as .exe or .com files.	Yes	No
hi	Hidden hides directories and files so they do not appear in a file manager or directory listing.	Yes	Yes
ic	Immediate Compression sets data to be compressed as soon as a file is closed. If applied to a directory, every file in the directory is compressed as each file is closed.  The files in the specified directory are compressed as soon as the operating system can perform the operation after the file is closed. This does not apply to the directory's subdirectories and the files in them.	Yes	Yes
ip	Immediate Purge flags a directory or file to be erased from the system as soon as it is deleted. Purged directories and files cannot be recovered.	Yes	Yes
ln	Link (status display only) indicates a symbolic link (soft link).	Yes	No
mg	Migrated (status display only) displays if the file or directory is migrated to near-line media.	Yes	Yes
mi	Migrate Inhibit prevents directories and files from being migrated from the server's disk to a near-line storage medium.	Yes	Yes
ri	Rename Inhibit prevents the file or directory name from being modified.	Yes	Yes

Attribute	Description	Applies to Files	Applies to Directories
ro	Read Only prevents a file from being modified. This attribute automatically sets Delete Inhibit and Rename Inhibit.	Yes	No
sd	Subdirectory (status display only) indicates that the entry is a directory, not a file.	No	Yes
sh	Sharable allows more than one user to access the file at the same time. This attribute is usually used with Read Only.	Yes	No
sy	System hides the directory or file so it does not appear in a file manager or directory listing. This attribute is normally used with operating system files, such as Linux or NetWare system files.	Yes	Yes
tr	Transactional allows a file to be tracked and protected by the Transaction Tracking System™ (TTS™).	Yes	No
vo	Volatile indicates that a file can change without being written to so that opportunistic locks cannot be set on it.	Yes	No

## A.2.4 Example

```
attrib /designs/topsecret -c=all -s=ro,di
```

This command clears all attributes, then sets Read Only and Delete Inhibit on the /designs/topsecret file.

## A.3 CDDVD (NetWare)

Use at the server console to allow a CD or DVD disk to function as a read-only Novell® Storage Services (NSS) volume. Immediately mounts the NSS volume.

This version of cddvd.nss replaces other CD modules used in previous NetWare releases. The NetWare 6 version of this module is cdrom.nlm and should be used only for NetWare 6 installation because cdrom loads the nss module, cd9660.nss, and cdhfs.nss, and these modules are not needed in NetWare 6.5. The NetWare 5 version was named cdinst.nlm and should be used only for NetWare 5 installation.

### A.3.1 Syntax

```
[LOAD] CDDVD
```

### A.3.2 Notes

- Supports CD and DVD volumes mounted with the Macintosh\* and ISO 9660 name spaces
- Supports High Sierra, ISO 9660, and HFS (Apple\*) file system formats

## A.4 COMPFIX

- [Section A.4.1, “Purpose,” on page 363](#)
- [Section A.4.2, “COMPFIX for Linux,” on page 363](#)

- [Section A.4.3, “COMPFIX for NetWare,” on page 366](#)

## A.4.1 Purpose

Use the COMPFIX utility to repair compression information for compressed NSS volumes or to clear the Cannot Compress attribute for files in the compressed NSS volume. This tool can help identify which compressed files are corrupted and can be fixed; however, not all corrupted compressed files are fixable.

## A.4.2 COMPFIX for Linux

### Prerequisite for Computing Compression Statistics

Before using the COMPFIX utility to compute compression statistics, make sure to put the volume in maintenance mode in order for statistics to be computed accurately.

- 1 On the server, open a terminal console, then log in as the root user by entering the `su` command and providing the root password when prompted.

- 2 At the terminal console prompt, start the NSS Console (NSSCON) utility by entering

```
nsscon
```

- 3 In the NSS console, enter

```
nss /poolmaintenance=poolname
```

### Syntax

On Linux, run the COMPFIX utility (`/opt/novell/nss/sbin/compfix`) from the terminal console prompt.

```
compfix [HOPTION]
```

```
compfix [GOPTION]... [VOPTION] VOLUMENAME
```

```
compfix [GOPTION]... [FOPTION] FILENAME
```

Mandatory arguments to long options are also mandatory for short options.

### Parameters

Parameter	Description
<i>VOLUMENAME</i>	Specify the volume name of the compressed volume you want to repair, such as <code>VOL1:</code> .
<i>FILENAME</i>	Specify the full path to the individual compressed file that has compression errors you want to repair, including its filename and extension. The filename must be in the Long name space.  For example:  <code><i>VOL1:/dirpath/filename.xxx</i></code>

## Help Options (HOPTION)

Option	Description
-h, --help	Displays help information and exits.
-v, --version	Displays version information and exits.

## General Options (GOPTION)

Multiple general options can be selected.

Option	Description
-H, --no-header	If this option is specified, Compfix does not validate compression headers. Use this option on volumes restored from scan files that do not have user data blocks.
-p, --logpath= <i>path</i>	Specifies the location of the log file.  The default location is at the root of the compressed volume you are analyzing or fixing.  Default: <code>./compfix.log</code>

## Volume-Level Options (VOPTION)

Only one volume-level option can be selected.

Option	Description
-D, --delete-all	Delete all non-fixable compressed files on the specified volume.
-F, --fix-all	Fix all repairable compressed files on the specified volume.
-L, --list-all	List all problematic compressed files on the specified volume.
-S, --fix-stats	Fix volume compression-related statistics. For accurate results, make sure the volume is in maintenance mode before issuing this command.
-C, --clear-all	Clear the CC (Cannot_Compress_File) attribute for all files on the specified volume.

## File-Level Options (FOPTION)

Only one file-level option can be selected.

Option	Description
-b, --background	Check if the specified file is eligible for the next background compression process.
-d, --delete	Delete the specified compressed file if it is non-fixable.
-f, --fix	Try to fix the specified file's compression-related problem.

Option	Description
-l, --list	List the specified file's compression-related information.
-c, --clear	Clear the specified file's Cc (Cannot_Compress_File) attribute.

## Examples

The following table illustrates typical uses of the Compfix utility. The commands in the left column should be written all on the same line, of course.

Command	Description
compfix --fix-stats VOL1:	Fix compression statistics for the specified volume, VOL1:. Make sure the volume is in maintenance mode to get accurate results.  Log the results in the default location of VOL1:/compfix.log.
compfix --list-all VOL1: --logpath=LOGS:/compfix.log	List all corrupted compressed files on the specified volume, VOL1:.  Log the results in the specified location of LOGS:/compfix.log.
compfix -f VOL1:\dir1\dir2\myfile.xxx	Fix an individual compressed file, myfile.xxx.  Log the results in the default location of VOL1:/compfix.log.
compfix -F VOL1: --logpath=LOGS:/compfix.log	Fix all fixable corrupted compressed files on the specified volume, VOL1:.  Log the results in the specified location of LOGS:/compfix.log.
compfix -b VOL1:/dir1/dir2/myfile.xxx	Check whether an individual compressed file, myfile.xxx, is eligible for the next background compression process.  Log the results in the default location of VOL1:/compfix.log.
compfix -C VOL1: --logpath=LOGS:/compfix.log	Clears the Cannot Compress (Cc) attribute for files in the specified volume, VOL1:.  Log the results in the specified location of LOGS:/compfix.log.

## A.4.3 COMPFIX for NetWare

### Prerequisite for Computing Compression Statistics

Before using the Compfix utility to compute compression statistics, make sure to put the volume in maintenance mode in order for statistics to be computed accurately.

- 1 At the server console, enter

```
nss /poolmaintenance=poolname
```

### Syntax

Load `compfix.nlm` on your NetWare® server, then enter the Compfix command at the server console. Compfix.nlm runs on processor 0 only.

```
compfix [-attribute] [-noheader] path [log_volume]
```

### Parameters

Parameter	Description
<i>path</i>	<p>Specify the volume name of the compressed volume you want to repair, or specify the full path to the individual file that has compression errors you want to repair. The filename must be in the Long name space.</p> <p>For example:</p> <pre>VOL1:</pre> <pre>VOL1:\dirpath\filename.xxx</pre>
<i>log_volume</i>	<p>If the <i>log_volume</i> is specified, the compfix log data is written to the <i>log_volume</i>:\compfix.txt file. For example, if LOGS: is the <i>log_volume</i>, then the file is written to LOGS:\compfix.txt.</p> <p>If the <i>log_volume</i> is not specified, the Compfix log data is written to the root of the compressed volume such as <i>volumename</i>:\compfix.txt file. For example, VOL1:\compfix.txt.</p>

### Option

Option	Description
-noheader	<p>If -noheader is specified, Compfix does not validate compression headers. Use this on volumes restored from scan files that do not have user data blocks.</p> <p>If -noheader is not specified, Compfix validates compression headers.</p>

### Attributes

Multiple attributes can optionally be used for a single command.

Attribute	Description
-fs	Fix compression statistics. Make sure the volume is in maintenance mode before you begin in order for statistics to be properly computed.
-lc	List all corrupted compressed files on the volume.
-df	Diagnose an individual compressed file.
-ff	Fix an individual compressed file.
-fc	Fix all fixable corrupted compressed files on the volume.
-bg	Diagnose the background compression effect for an individual file to determine if the specified file is eligible for the next background compression process.
-cc	Clears the <code>Cannot Compress (Cc)</code> attribute for files in the specified volume or for a specified file.

## Examples

The following table illustrates typical uses of the Compfix utility. The commands in the left column should be written all on the same line, of course.

Command	Description
<code>compfix -fs VOL1:</code>	Fix compression statistics for the specified volume, <code>VOL1:</code> . The volume should be in maintenance mode to get accurate results.  Log the results in the default location of <code>VOL1:\compfix.txt</code> .
<code>compfix -lc VOL1: LOGS:</code>	List all corrupted compressed files on the specified volume, <code>VOL1:</code> .  Log the results in the specified location of <code>LOGS:\compfix.txt</code> .
<code>compfix -ff VOL1:\dir1\dir2\myfile.xxx</code>	Fix an individual compressed file, <code>myfile.xxx</code> .  Log the results in the default location of <code>VOL1:\compfix.txt</code> .
<code>compfix -fc VOL1: LOGS:</code>	Fix all fixable corrupted compressed files on the specified volume, <code>VOL1:</code> .  Log the results in the specified location of <code>LOGS:\compfix.txt</code> .
<code>compfix -bg VOL1:\dir1\dir2\myfile.xxx</code>	Diagnose the background compression effect for an individual file, <code>myfile.xxx</code> .  Log the results in the default location of <code>VOL1:\compfix.txt</code> .

Command	Description
<code>compfix -cc VOL1: LOGS:</code>	Clears the Cannot Compress (Cc) attribute for files in the specified volume, VOL1:.  Log the results in the specified location of LOGS:\compfix.txt.

## A.5 DISMOUNT

Use at the server console to make a volume unavailable to users. For NSS on Linux, issue this command from the `nsscon` utility.

### A.5.1 Syntax

```
dismount volume_name
```

### A.5.2 Option

Parameter	Description
<code>volume_name</code>	Specify the name of the volume you want to dismount.

### A.5.3 Notes

- `Dismount` allows you to maintain or repair a volume and to upgrade disk drivers while the NetWare server is up.
- Use this command to dismount all volumes residing on a removable drive before you remove the drive.
- If you are running out of memory and cannot add more RAM, consider dismounting a volume that is used infrequently.

Mounted volumes take up memory allocated for file caching and thus affect server performance. NSS volumes use much less memory than Traditional volumes.

## A.6 FILER (NetWare)

This legacy DOS utility can be used at a workstation to manage files and directories on NSS for NetWare volumes. `FILER` works with the file system; it doesn't affect eDirectory objects.

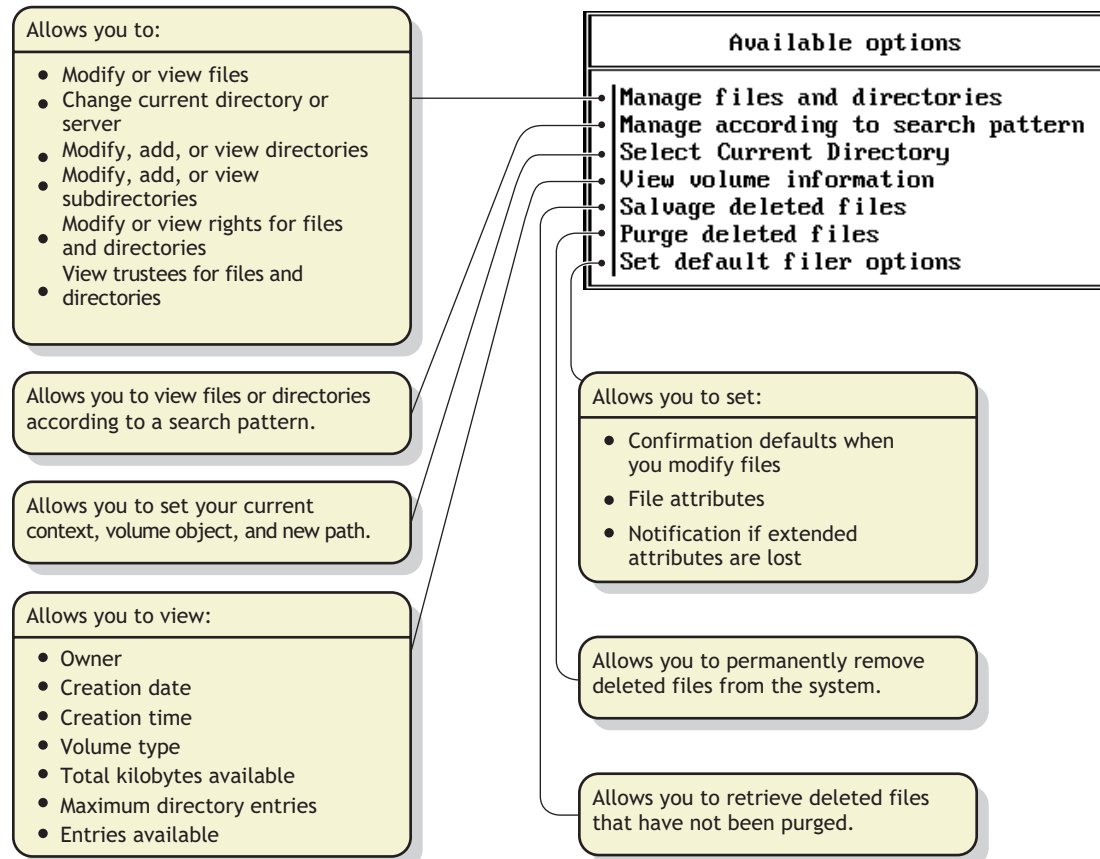
Filer is a client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the Novell Map Network Drive option in the client menu to map a local drive letter to the `sys:\public` directory. Open a DOS Command Prompt window, `cd` to the mapped drive, then issue the `filer` command from the command line.



## A.6.1 Options

The following figure shows the functions of FILER.

**Figure A-1** *FILER Functions*



## A.7 FLAG (NetWare)

This legacy DOS utility can be used at a workstation to do the following:

- View or modify file and directory attributes to determine which operations can be performed with the file or directory
- Modify the owner of a file or directory
- View or modify the search mode of executable files to determine how the program uses search drives when looking for a file

### A.7.1 Syntax

The `flag` command is a DOS command-line command utility, `sys:\public\flag.exe`.

FLAG is client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the Novell Map Network Drive option in the client menu to map a local drive letter to the `sys:\public` directory. Open a DOS Command Prompt window, `cd` to the mapped drive, then issue the `flag` command from the command line.

`flag path [[+ | -] attribute...] [/option...] [/? | /VER]`

## A.7.2 Options

- Use FLAG options to set or clear directory and file attributes.
- You can use wildcard characters.

---

**NOTE:** For help with tasks, press F1 in the utility.

---

### Parameters

Parameters	Description
(no option)	If you type FLAG with no parameters, the attribute status of all the files in the current directory appears.
<i>path</i>	Specify the path to the file or directory whose attributes or search modes you want to modify.
+, -	Set attributes for a file or directory with the + (plus).
	Clear attributes from a file or directory with the - (minus).
	If neither + nor - is used, the attributes will be assigned to the file/directory as specified.
	If you set and clear attributes in the same command, group all + (plus) attributes together and all - (minus) attributes together.
<i>attribute</i>	Specify one or more file or directory attributes.  See <a href="#">Directory Attributes</a> and <a href="#">File Attributes</a> .
<i>/option</i>	Replace option with any available option. See <a href="#">General Options</a> .
<i>/?</i>	View online help. All other parameters are ignored when <i>/?</i> is used.
<i>/VER</i>	View the version number of the utility and the list of files it uses to execute. All other parameters are ignored when <i>/VER</i> is used.

### General Options

Option	Description
<i>/NAME=name</i>	Change the owner of a file or directory.
<i>/GROUP=name</i>	
<i>/D</i>	View details about a file or directory.
<i>/DO</i>	View or modify only directories in the specified path.
<i>/FO</i>	View or modify only files in the specified path.
<i>/OWNER=name</i>	View all files or directories owned by a user.

Option	Description
<code>/M=mode</code>	Modify search modes of executable files. See <a href="#">Search Modes for Executable Files</a> .
<code>/S</code>	Search the subdirectory in the specified path and any subdirectories below that level.
<code>/C</code>	Scroll continuously through output.

## Search Modes for Executable Files

**IMPORTANT:** The syntax is `/M=mode`. Replace *mode* with a mode number from the following table.

Mode	Description
0	Search for instructions in the <code>net.cfg</code> file (the default mode).
1	Search the path specified in the file. If there is no path, the file searches the default directory, and then all search drives.
2	Search the path specified in the file. If there is no path, the file searches only the default directory.
3	Search the path specified in the file. If there is no path, the file searches the default directory; then if the open request is read only, the file searches the search drives.
4	Reserved; don't use.
5	First search the path specified and then all search drives. If there is no path, the file searches the default directory and then all search drives.
6	Reserved; don't use.
7	First search the path specified. If the open request is read only, the file searches the search drives. If there is no path, the file searches the default directory and then all search drives.

## A.7.3 Attributes

When you set or clear multiple attributes, include a space between attribute abbreviations.

### Directory Attributes

Option	Description
ALL	Specify the Di, H, Ic, P, Ri, and Sy attributes as a group. Primarily used to assign directories these specific attributes.

Option	Description
Di (Delete Inhibit)	Prevent the directory from being deleted.  This attribute overrides the trustee Erase right. When it is enabled, no one, including the owner and network administrator, can delete the directory or file. A trustee with the Modify right must disable this right to allow the directory or file to be deleted.
Dm (Don't Migrate)	Prevent the directory from being migrated to a secondary backup system (regardless of what the volume is set to).
H (Hidden)	Prevent the directory from being seen in the directory listing.
Ic (Immediate Compression)	Compress the files in the specified directory as soon as the OS can.  This does not apply to the directory's subdirectories and the files in them. Use Ic in combination with the /S option to apply immediate compression to the directory's subdirectories and their files recursively.
N (Normal)	Specify no attributes.
P (Purge)	Purge the directory or files in a directory immediately when deleted.
Ri (Rename Inhibit)	Prevent the directory from being renamed.
Sy (System)	System hides the directory so it does not appear in a file manager or directory listing. System is normally used with operating system files, such as Linux or NetWare system files.

## File Attributes

Option	Description
ALL	Specify the A, Ci, Di, H, Ic, P, Ri, Ro, Sh, Sy and T attributes as a group. Primarily used to assign files these specific attributes.
A (Archive needed)	Indicate that the file has been modified since the last backup.
Ci (Copy Inhibit)	Prevent files from being copied. This attribute overrides the trustee Read right and File Scan right.
Dc (Don't Compress)	Prevent a file from being compressed.
Di (Delete Inhibit)	Prevent a file from being deleted or copied over.  This attribute overrides the trustee Erase right. When it is enabled, no one, including the owner and network administrator, can delete the directory or file. A trustee with the Modify right must disable this right to allow the directory or file to be deleted.
Dm (Don't Migrate)	Prevent a file from being migrated to a secondary backup system (regardless of what the volume or directory is set to).
Ds (Don't Suballocate)	Prevent an individual file from being suballocated, even if suballocation is enabled for the system. Use on files which are enlarged or appended to frequently, such as certain database files.  This flag applies only to files in a NetWare Traditional File System volume.

Option	Description
H (Hidden)	Prevent a filename from being displayed with the DOS DIR command. The file can't be copied or deleted.
Ic (Immediate Compression)	Compress a file as soon as the OS can.
N (Normal)	Specify the Rw attribute.
P (Purge)	Purge a file immediately if the file is deleted.
Ri (Rename Inhibit)	Prevent a file from being renamed.
Ro (Read Only)	Allow a file to only be read; it can't be written to or deleted (in other words, Ro includes Ri and Di).
Rw (Read Write)	Allow a file to be read and written to.
Sh (Shareable)	Allow a file to be used by several users simultaneously. This attribute is usually used with Read Only.
Sy (System)	System hides the directory or file so it does not appear in a file manager or directory listing. System is normally used with operating system files, such as Linux or NetWare system files.
T (Transactional)	Protect a file by using the Transaction Tracking System™.
X (Execute)	Identifies executable files such as .exe and .com files.

## Status Attributes

Status attributes show information for a file or directory. You cannot directly modify these flags.

Status	Description
Cc (Cannot compress)	The file cannot be compressed because of limited space savings or failures that might have occurred during an unsuccessful compression attempt.  If this flag is set for a file, the system will not attempt compression again until the next time the file is saved. To allow the file to be reevaluated for compression, use the COMPFIX utility to reset (turn off) this flag. For information, see <a href="#">Section A.4, "COMPFIX," on page 362</a> .
Co (Compressed)	The file has been compressed.
M (Migrated)	The file has been migrated.

## A.7.4 Examples

- To add Rw and Ic and remove Di and A from all files in the current directory, enter  

```
FLAG *.* +R W Ic -Di A
```
- To view attributes for file STUDENTS in directory SYS:\COURSE\WINTER, enter  

```
FLAG SYS:\COURSE\WINTER\STUDENTS
```
- To give files in drive G: the Hidden attribute, enter  

```
FLAG G:*.* H /FO
```

- To set the search mode to 7 for all executable files in the current directory, enter  

```
FLAG *.* /M=7
```
- To set the search mode to 1 and assign all files in the current directory the Read Write attribute, enter  

```
FLAG *.* RW /M=1
```

## A.8 METAMIG (Linux)

The Novell Storage Services (NSS) File System Metadata Migration Utility (METAMIG) for OES Linux allows you to save and restore NSS file system trustee, user quota, and directory quota metadata. This utility works only with files in the NSS file system on OES Linux. It is available in OES SP1 and later.

### A.8.1 Syntax

METAMIG is located in the `/opt/novell/nss/sbin/metamig` directory.

```
metamig [OPTIONS]
```

```
metamig save volume [SOPTIONS]
```

```
metamig restore volume [ROPTIONS]
```

### A.8.2 Arguments

The first argument indicates the action to be taken on the specified NSS volume. Possible actions are Save and Restore.

Argument	Description
save	Save the indicated metadata to stdout.
restore	Restore the indicated metadata to stdin.

The second argument specifies the NSS volume name to be saved or restored.

The third argument specifies the path to a file. This is the file to be created on save, or the file to be restored from on restore.

### A.8.3 Options

Several option types are available.

#### General Options

The following options are general options available to actions related to save or restore.

Option	Description
<code>-v, --version</code>	Display the program version information.

Option	Description
-h, --help	Display the help screen.

## Options

The following options are options available to actions related to save.

Option	Description
-m, --meta=mask	The types of metadata to be saved. If a value is not specified, it defaults to all. For information, see <a href="#">Mask</a> .

## ROptions

The following options are options available to actions related to restore.

Option	Description
-m, --meta=mask	The types of metadata to be saved. If a value is not specified, it defaults to all. For information, see <a href="#">Mask</a> .
-t, --tree=treename	If this value is specified, this setting overrides the saved directory tree name.
-d, --details	Display all actions taken.

## Mask

The mask is a string of characters with each character representing a type of NSS file system metadata. Use the mask to specify values for the meta parameter in [Options](#) and [ROptions](#).

Option	Description
t	Trustees
u	User quotas
d	Directory quotas
a	All metadata

## A.8.4 Examples

To save the user quota data and directory quota data for a volume called VOL1 in the /backup/volquotas file, enter

```
metamig save VOL1 -m ud >/backup/volquotas
```

To restore the user quota data from the file /backup/volquotas to VOL1, enter

```
metamig restore VOL1 -m u </backup/volquotas
```

## A.9 MIRROR STATUS (Netware)

Use at the server console to do the following:

- View the status of mirrored disk partitions
- View the percentage of mirrored data on each partition

Mirror Status can be used on any NetWare mirror, regardless of type as long as all the partition types in the mirror match. This includes NSS, TRAD, SBD, and iSCSI devices and any other NetWare partition types that are mirrored.

### A.9.1 Syntax

```
MIRROR STATUS [logical_partition_number]
```

### A.9.2 Option

Option	Description
(no option)	View a list of all logical disk partitions with their mirrored status.
<i>logical_partition_number</i>	Display the partition's mirrored status and the devices that make up the members of the mirror group.

### A.9.3 Status Messages

The possible status messages are described in the following table.

Status	Description
Being remirrored	Remirroring is in progress; the percent completed is displayed.
Fully synchronized	The mirrored partitions have the same data. Remirroring is complete.
Not mirrored	Disk mirroring was not set up for this partition. It has no mirrored partner.
Orphaned state	A partition has been removed from a mirrored group, and the volumes on the partition have not been renamed. These volumes cannot be mounted unless they are renamed or remirrored.
Out of synchronization	The partition is out of synchronization with its mirrored partners and for some reason cannot be resynchronized.

## A.10 MOUNT

Use at the server console to make a volume available to users. For NSS on Linux, issue this command from the `nsscon` utility.



## A.10.1 Syntax

```
MOUNT volume_name | ALL
```

Parameter	Description
<code>volume_name</code>	Specify the name of the volume you want to mount.
<code>ALL</code>	Mount all volumes not mounted.

## A.10.2 Notes

- If you set up the `autoexec.ncf` file to mount all volumes automatically each time the server comes up, then you would use the `MOUNT` command only if you have dismounted a volume and want to remount it.
- After you replace a removable drive, use this command to mount all volumes residing on the removable drive.
- You can mount and dismount volumes while the server is running.

## A.10.3 Examples

- To mount the `STATISTICS` volume, enter

```
MOUNT STATISTICS
```

- To mount all unmounted volumes at once, enter

```
MOUNT ALL
```

# A.11 NCOPY (NetWare)

The `ncopy` command is a legacy DOS utility that provides two attributes to specify your intention for copying a compressed file to a compressed volume or a non-compressed volume.

`NCOPY` is client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the Novell Map Network Drive option in the client menu to map a local drive letter to the `sys:\public` directory. Open a DOS Command Prompt window, `cd` to the mapped drive, then issue the `ncopy` command from the command line. You can also run a GUI version of `NCOPY` from the Novell Client by selecting *NetWare Utilities > NetWare Copy*, then following the instructions.

## A.11.1 Syntax

Command	Description
<code>ncopy /R</code>	Retain compression on supported media.
<code>ncopy /R/</code>	Retain compression on unsupported media.

## A.12 NDIR (NetWare)

The `ndir` command is a legacy DOS utility that allows you to view file attributes and compression information for all files in a directory.

NDIR is a client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the Novell Map Network Drive option in the client menu to map a local drive letter to the `sys:\public` directory. Open a DOS Command Prompt window, `cd` to the mapped drive, then issue the `ndir` command from the command line.

### A.12.1 Syntax

Command	Description
<code>ndir /R</code>	View the file attributes for all files in a directory.
<code>ndir /COMP</code>	View the compression information for all files in a directory. It shows the amount of disk space saved by compressing data.
<code>ndir /COMP /VOL</code>	View the compression information for the volume. It also shows the disk space saved on the volume.

## A.13 RAVSUI (Linux)

Use the Rebuild and Verify Simple User Interface (`ravsui`) utility to rebuild or verify an NSS pool on Linux that is in a maintenance state.

### A.13.1 Syntax

```
/opt/novell/nss/sbin/ravsui  
  
ravsui [OPTION]... [ROPTION]... rebuild poolname  
ravsui [OPTION]... [VOPTION]... verify poolname
```

### A.13.2 Arguments

The first argument is the ACTION to perform. This can be either `rebuild` or `verify`.

`Rebuild` checks the integrity of the data in the pool and rebuilds the pool. A rebuild operation can take up to several hours, depending on the size of the pool.

`Verify` checks the integrity of the data in the pool; it does not repair. A verify operation can take several minutes, depending on the size of the pool.

The second argument is the TARGET of the action. Replace *poolname* with the name of an NSS pool to be rebuilt or verified.

### A.13.3 Options

Mandatory arguments for long options are also mandatory for short options.

## OPTION

The following are general options available to either rebuild or verify actions.

Option	Description
-a, --attach	Attaches to a rebuild or verifies that is running.
-D, --log-kernel=MASK	Controls the amount of log output from the kernel. Default : 0x70
-d, --log-application=LEVEL	Controls the amount of log output from the application. Default: 0x60
-h, --help	Displays this help information and exits.
-P, --path=PATH	Specifies the path to store logs and the results. Default: /var/opt/novell/log/nss/rav/
-v, --version	Displays version information and exits.

## ROPTION

The following options are available for the rebuild action.

Option	Description
-i, --iv-prune	Prunes internal volume.
-l, --loss-file-limit=LIMIT	Maximum number of files per volume to quietly prune. Default : 100
-p, --purge-deleted-files	Purges deleted files.
-r, --rezid=ZID	Threshold to cause a rezid of a volume. Default: 0xefffffff
-u, --unknown-loss-prune	Prunes if losses unknown.

## VOPTION

The following options are available for the verify action.

Option	Description
-q, --quick	Skips cross-tree validations.

## A.13.4 Example

```
ravsui -q verify POOL1
```

This command checks the integrity of the data in in POOL1 in the current tree; it skips cross-tree validations.

## A.14 RAVVIEW (Linux)

The Rebuild and Verify View (ravview) utility dumps specified rebuild or verify files in human-readable format.

### A.14.1 Syntax

```
/opt/novell/nss/sbin/ravview
```

```
ravview [OPTION]... rtf filename
```

```
ravview [OPTION]... [NOPTION]... rtfn poolname
```

```
ravview [OPTION]... [VOPTION]... vbf filename
```

```
ravview [OPTION]... [NOPTION]... [VOPTION]... vbfn poolname
```

### A.14.2 Arguments

The first program argument is the FILETYPE. This can be *rtf*, *rtfn*, *vbf*, or *vbfn*.

For *rtf*, the *filename* is a rebuild log file. RAVVIEW converts ZIDs in the *filename* into the file's name. It also displays any messages that occurred during rebuild.

For *rtfn*, RAVVIEW uses the newest rebuild log in *path* for the specified *poolname*.

For *vbn*, the *filename* is a verify report file. RAVVIEW converts this into a human-readable form.

For *vbfn*, RAVVIEW uses the newest verify report in *path* for the specified *poolname*.

The second argument is the TARGET of the action. Replace *poolname* with the name of an NSS pool to be rebuilt or verified.

### A.14.3 Options

Mandatory arguments for long options are mandatory for short options too.

#### OPTION

The following are general options available to actions related to rebuild or verify.

Option	Description
-h, --help	Displays help information and exits.
-V, --verbose=LEVEL	Controls the amount of output.
-v, --version	Displays version information and exits.

## VOPTION

The following options are available for the verify action.

Option	Description
-a, --actions	Displays action information.
-H, --histograms	Display histograms.
-o, --object-details	Displays object details.

## NOPTION

The following options are available for the newest action.

Option	Description
-P, --path=PATH	Specifies the path to RAV files.
	Default: /var/opt/novell/log/nss/rav/

### A.14.4 Example

```
ravview -V /var/opt/novell/log/nss/rav/ rtfn POOL1
```

In this command, RAVVIEW uses the newest rebuild log in the /var/opt/novell/log/nss/rav/ directory for POOL1.

## A.15 REBUILD (NetWare)

Use REBUILD to recover corrupted Novell Storage Services (NSS) pools. The REBUILD utility salvages the data it finds on your corrupted NSS pool and recovers it.

NetWare users may be used to using VREPAIR to repair and rebuild NetWare volumes. VREPAIR still works for NetWare Traditional volumes, but it does not work on NSS pools. For NSS, you must use REBUILD.

Rebuild verifies and uses the existing leaves of an object tree to rebuild all the other trees in the system. The NSS pools that are verified and rebuilt are placed in maintenance mode. This means the NSS pool and the volumes on it are unusable until this process is finished, and the pool is remounted.

After running REBUILD, you must run the VERIFY utility. VERIFY accounts for all blocks in the system. If errors are found, they are reported to the screen, and the NSS pool is left in maintenance mode. Run REBUILD again until no errors are found. If errors are not found, the pool is placed back in the active state. You may have to mount the pool and its volumes.

---

**NOTE:** This utility protects against system failures only, not hardware failures.

---

## A.15.1 Syntax

To rebuild a single pool, specifying the name, enter

```
nss /poolrebuild=pool_name
```

## A.15.2 Starting and Using REBUILD

Use REBUILD at the command line.

---

**NOTE:** Always back up your data. If an NSS pool uses several hard disks and one of the hard disks becomes corrupted, you must create a new NSS pool, then restore its NSS volumes from backup.

---

**1** Load NSS.

**2** Rebuild your NSS pool.

Select one of the following options:

- To rebuild a single pool, specifying the name, enter

```
nss /poolrebuild=pool_name
```

- To rebuild a single pool by selecting from a list of pool names, enter

```
nss /poolrebuild
```

A list of pools appears. Select the preferred pool. You must enter this command for each pool you want to rebuild.

- To rebuild more than one NSS volume at a time, enter

```
nss /poolrebuild=pool_name, pool_name
```

A process runs for each volume up to five NSS volumes.

Rebuilding a pool can take up to several hours, depending on the size of the pool.

**3** Check the output screen.

This screen indicates the time elapsed, the time remaining, the total elapsed time, the number of objects processed, etc.

**4** Verify your rebuilt NSS pool, by entering

```
nss /poolverify=[pool_name]
```

If you do not enter the NSS pool name parameter, you will be prompted to select it.

## A.16 REMIRROR PARTITION (NetWare)

Use at the server console to start the remirroring of a logical partition.

Because the server remirrors partitions automatically, you should use REMIRROR PARTITION only if you have stopped remirroring by using the ABORT REMIRROR command or something has caused your server to cease remirroring.

Remirror Partition can be used on any NetWare mirror, regardless of type as long as all the partition types in the mirror match. This includes NSS, TRAD, SBD, and ISCSI devices and any other NetWare partition types that are mirrored.

## A.16.1 Syntax

```
REMIRROR PARTITION logical_partition_number
```

Parameter	Description
<i>logical_partition_number</i>	Specify the number of the logical partition you want to start remirroring.

## A.16.2 Example

To start remirroring for logical partition 0x4, type

```
REMIRROR PARTITION 4
```

## A.17 RIGHTS

- [Section A.17.1, “RIGHTS \(Linux\),” on page 383](#)
- [Section A.17.2, “RIGHTS \(NetWare\),” on page 386](#)

### A.17.1 RIGHTS (Linux)

The Trustee Rights Utility (`rights`) for OES Linux allows you to specify trustee rights for directories and files in the NSS file system. This utility does not provide support for trustees on Linux file systems. It is also not meant to be used to set trustees for NSS volumes on OES NetWare. The trustee information is saved in the file and directory metadata in the NSS volume and works seamlessly with OES NetWare if the volume is moved to OES NetWare.

#### Syntax

```
rights [OPTIONS]
rights [TOPTIONS] trustee username
rights [DOPTIONS] delete username
rights [IOPTIONS] irf
rights [EROPTIONS] effective username
rights [SOPTIONS] show
```

#### Options

#### ACTIONS

The first argument indicates the action to be taken.

Option	Description
<code>trustee</code>	Adds or modifies a trustee on a file or directory.

Option	Description
delete	Removes a trustee from a file or directory.
irf	Sets the inherited rights filter on a directory.
effective	Displays a user's effective rights.
show	Displays the trustees and inherited rights filter.

## OPTIONS

Option	Description
-v, --version	Displays the program version information.
-h, --help	Displays the help screen.

## TOPTIONS

Option	Description
-r, --rights=MASK	<p>Specifies the rights to be given to this trustee. For more information, see <a href="#">"MASK" on page 385</a>.</p> <p>If the No Rights (n) option is assigned, the trustee is removed.</p> <p>If rights are not specified, the default assignment is Read and File Scan rights.</p>
-f, --file=filename	<p>Specifies the name of file or directory to assign trustees to. <i>Filename</i> is the path for the file or directory. For example:</p> <pre>-f /users/username/userfile.sxi</pre> <pre>--file=/designs/topsecret</pre> <p>If a file or directory is not specified, the current directory is used.</p>

## DOPTIONS

Option	Description
-f, --file=filename	<p>Specifies the name of file or directory to delete trustees from. <i>Filename</i> is the path for the file or directory.</p> <p>If a file or directory is not specified, the current directory is used.</p>



## IOPTIONS

Option	Description
<code>-r, --rights=MASK</code>	Specifies the rights to be passed through the filter. For more information, see <a href="#">"MASK" on page 385</a> .  If rights are not specified, the default assignment is All Rights.
<code>-f, --file=filename</code>	Specifies the name of the directory where the filter is to be applied. <i>Filename</i> is the path for the directory.  If a directory is not specified, the current directory is used.

## EROPTIONS

Option	Description
<code>-f, --file=filename</code>	Specifies the name of file or directory where effective right are to be calculated. <i>Filename</i> is the path for the file or directory.  If a file or directory is not specified, the current directory is used.

## SOPTIONS

Option	Description
<code>-f, --file=filename</code>	Specifies the name of the file or directory to display a list of trustees for that file or directory.  If a file or directory is not specified, the current directory is used.

## USERNAME

The username is the Full Distinguished Name of a Novell eDirectory™ object, including the tree name. For example: `username.context.treename` or `joe.engineer.acme_tree`.

## MASK

The mask is a string of characters, with each character representing a type of rights. The following table lists the rights, the letter to use for each right, and what the right is used for.

Right	Letter	Description
Supervisor	s	Grant all rights to the file or directory.
Read	r	Open and read files in the directory.
Write	w	Open and write to files in the directory.
Create	c	Create files and subdirectories.
Erase	e	Erase files and directories.

Right	Letter	Description
Modify	m	Rename files and directories, and change file attributes.
File Scan	f	View and search on file and directory names in the file system structure.
Access Control	a	Add and remove trustees, and change trustee rights to files and directories.
No Rights	n	Remove all rights.
All Rights	all	Add All rights except Supervisor.

### Example

```
rights -f /designs/topsecret -r rwfc trustee joe.engineer.acme_tree
```

This command assigns Read, Write, File Scan, and Create rights to the /designs/topsecret directory for user Joe in the engineer context of the acme\_tree eDirectory tree.

### See Also

For information about setting file system directory and file attributes, see [“ATTRIB \(Linux\)” on page 360](#).

## A.17.2 RIGHTS (NetWare)

Use at a workstation to do the following:

- View or modify user or group rights for files
- View or modify user or group rights for directories and volumes

RIGHTS is a client-side application, not a server-side application. Install the Novell Client for Windows on your workstation, then use the *Novell Map Network Drive* option in the client menu to map a local drive letter to the sys:/public directory. Open a DOS Command Prompt window, cd to the mapped drive, then issue the rights command from the command line.

### Syntax

```
RIGHTS path [[ + | - ] rights] [/option...] [/? | /VER]
```

Parameter	Description
<i>path</i>	Specify the path to the file, directory, or volume you want to modify or view rights to (you must always specify a path).
+   -	Add or delete the specified rights. See <a href="#">Using RIGHTS</a> .
<i>rights</i>	Specify one or more file or directory rights. See <a href="#">File and Directory Rights</a> .
<i>/option</i>	Replace <i>option</i> with any available option. See <a href="#">RIGHTS Options</a> .
<i>/?</i>	View online help. All other parameters are ignored when <i>/?</i> is used.

Parameter	Description
/VER	View the version number of the utility and the list of files it uses to execute. All other parameters are ignored when /VER is used.

## RIGHTS Options

Option	Description
/C	Scroll continuously through output.
/F	View the Inherited Rights Filter (IRF).
/I	View the trustee and group rights that created the inherited rights, and view where the inherited rights came from.
/NAME= <i>username</i>	View or modify rights for the user or group listed. Replace <i>username</i> with the name of the user or group whose rights you want to view or modify.
/S	View or modify subdirectories below the current level.
/T	View trustee assignments in a directory.

## File and Directory Rights

The following table lists the rights, the letter to use for each right, and what the right is used for.

Right	Letter	Description
Supervisor	S	Grant all rights to the file or directory.
Read	R	Open and read files in the directory.
Write	W	Open and write to files in the directory.
Create	C	Create files and subdirectories.
Erase	E	Erase files and directories.
Modify	M	Rename files and directories, and change file attributes.
File Scan	F	View and search on file and directory names in the file system structure.
Access Control	A	Add and remove trustees and change trustee rights to files and directories.
No Rights	N	Remove all rights.
Remove	REM	Remove the user or group as a trustee of the specified file or directory.
All	ALL	Add All rights except Supervisor.

## Using RIGHTS

- If you use + (plus) to add rights, the rights you list are added to the existing rights.

- If you use – (minus) to remove rights, the rights you list are deleted from the existing rights.
- If you add and delete rights in the same command, group all added rights together and all deleted rights together.
- If you list rights without using + or –, the rights you list replace the existing rights.
- You must always specify a path. You can use a period (.) to represent your current directory.
- You can use wildcard characters.

## Examples

- To set the trustee rights in the current directory for user JANICE to Read, Write, and File Scan, type

```
RIGHTS . R W F /NAME=JANICE
```

- To remove user ERNESTO from SYS:USERS, type

```
RIGHTS SYS:USERS REM /NAME=ERNESTO
```

- To see where user PATRICK's inherited rights came from for SYS:USERS\HOME, type

```
RIGHTS SYS:USERS\HOME /NAME=PATRICK /I
```

## A.18 SET (NetWare)

Use at the server console to view and configure operating system parameters.

---

**TIP:** You can also modify server parameter values with the *Set Parameter* link in Novell Remote Manager.

---

The default SET parameter values give maximum performance for most systems. Server parameter values should seldom need to be modified.

### A.18.1 Syntax

```
set [parameter] = [value]
```

#### Parameter Categories

This section discusses only those parameters related to managing NSS file systems on NetWare. For information about other SET options, see “SET” in the *Utilities Reference for OES*.

[Common File System Parameters \(page 389\)](#)

[Disk Parameters \(page 392\)](#)

[Novell Storage Services Parameters \(page 395\)](#)

### A.18.2 Using SET

Default values of the server parameters provide the desired performance for a majority of deployment scenarios.

Server parameter settings are persistent in NetWare. If the server goes down, any settings that you have made to tune your server for performance will not be lost.

## Displaying and Changing Current Settings

- If you type `SET` without a parameter, a list of numbered categories appears. When you select a category, the current settings for the server parameters in that category appear along with a brief description of each parameter, the range of valid values, and the default value.
- If you type `SET` with a parameter but no value, the current setting, range of valid values, and a brief description of the parameter are displayed.
- If you type `SET` with a parameter and a value, the operating system is reconfigured according to the specified value.

## Entering Parameters

- You can modify the values of most `SET` parameters at the console prompt. The system is immediately configured to that setting. Any setting in the `autoexec.ncf` file is overridden.
- You can enter `SET` commands that you execute at the console prompt in the `autoexec.ncf` file. When a parameter value is set in this file, the server configures itself to that setting each time the server is booted—unless the value is modified.
- Some commands can be also saved in the `startup.ncf` file. Use `NWCONFIG` to edit both the `autoexec.ncf` and the `startup.ncf` file.

## A.18.3 Common File System Parameters

The following Common File System parameters apply to both the NSS file system and the NetWare Traditional file system.

Parameter	Description
Maximum Transactions = <i>number</i>	<p>Specify how many transactions can occur at the same time.</p> <p><b>Supported values:</b> 100 to 10000 <b>Default:</b> 10000</p> <p><b>Owning NLM:</b> <code>filesys.nlm</code></p>
Maximum Concurrent Directory Cache Writes = <i>number</i>	<p>Specify how many write requests from directory cache buffers are put in the elevator before the disk head begins a sweep across the disk.</p> <p><b>Supported values:</b> 5 to 500 <b>Default:</b> 75</p> <p>A high number creates more efficient write requests. A low number creates more efficient read requests.</p> <p><b>Owning NLM:</b> <code>filesys.nlm</code></p>

Parameter	Description
Minimum File Delete Wait Time = <i>time</i>	<p>Specify how long a deleted file remains salvageable on the volume.</p> <p><b>Supported values:</b> 0 seconds to 7 days  <b>Default:</b> 1 minute 5.9 seconds</p> <p>Files deleted for less than this minimum aren't automatically purged even if the volume is full and users can't create new files.</p> <p><b>Owning NLM:</b> <code>filesys.nlm</code></p>
Immediate Purge of Deleted Files = <i>value</i>	<p><b>Supported values:</b> On, Off  <b>Default:</b> Off</p> <p>If this parameter is set to On, all files are purged immediately when they are deleted.</p> <p><b>Owning NLM:</b> <code>filesys.nlm</code></p>
Compression Daily Check Stop Hour = <i>number</i>	<p>Specify the hour when you want the file compressor to stop scanning enabled volumes for files that need to be compressed.</p> <p><b>Supported values:</b> 0 to 23  <b>Default:</b> 6</p> <p>Hours are specified by a 24-hour clock: (0 = midnight; 23 = 11 p.m.).</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> <code>filesys.nlm</code></p>
Compression Daily Check Starting Hour = <i>number</i>	<p>Specify the hour when you want the file compressor to start scanning enabled volumes for files that need to be compressed.</p> <p><b>Supported values:</b> 0 to 23  <b>Default:</b> 0</p> <p>Hours are specified by a 24-hour clock: (0 = midnight; 23 = 11 p.m.).</p> <p><b>Note:</b> If the Compression Daily Check Stop Hour parameter is the same as the Compression Daily Check Starting Hour, then the file compressor starts checking every day at the Compression Daily Starting Hour time and runs as long as necessary to finish all files that meet the compressible criteria.</p> <p>This parameter can be set in the startup.ncf file.</p>

Parameter	Description
Minimum Compression Percentage Gain = <i>number</i>	<p>Set the minimum percentage a file must compress to remain in a compressed state.</p> <p><b>Supported values:</b> 0 to 50 <b>Default:</b> 20</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Enable File Compression = <i>value</i>	<p>Specify whether file compression is suspended.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> On</p> <p>On allows file compression on compression-enabled volumes.</p> <p>Off suspends compression. Immediate Compression requests are queued until the value is reset to On, when the files meeting criteria are compressed.</p> <p>This parameter can be set in the startup.ncf.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Maximum Concurrent Compressions = <i>number</i>	<p>Specify the maximum concurrent or simultaneous compressions allowed.</p> <p><b>Supported values:</b> 1 to 8 <b>Default:</b> 2</p> <p>Concurrent compressions can occur only if there are multiple volumes.</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Convert Compressed to Uncompressed Option = <i>value</i>	<p>Specify what the file system does with an decompressed version of a file after the server has decompressed it.</p> <p><b>Supported values:</b></p> <p>0 = Always leave the file compressed 1 = Leave the file compressed until second access if it is read only once during the time specified by the Days Untouched Before Compression parameter 2 = Always leave the file decompressed</p> <p><b>Default:</b> 1</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>

Parameter	Description
Decompress Percent Disk Space Free to Allow Commit = <i>number</i>	<p>Specify the percentage of free disk space required on a volume for file decompression to permanently change compressed files to decompressed.</p> <p><b>Supported values:</b> 0 to 75 <b>Default:</b> 10</p> <p>This parameter prevents newly decompressed files from filling up the volume.</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Decompress Free Space Warning Interval = <i>number</i>	<p>Specify the time between alerts when the file system is not changing compressed files to decompressed because of insufficient disk space. Setting the interval to 0 turns off the alert.</p> <p><b>Supported values:</b> 0 seconds to 29 days 15 hours 50 minutes 3.8 seconds <b>Default:</b> 31 minutes 18.5 seconds</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Deleted Files Compression Option = <i>number</i>	<p>Specify whether and when deleted files are compressed.</p> <p><b>Supported values:</b></p> <ul style="list-style-type: none"> <li>0 = Don't Compress deleted files</li> <li>1 = Compress deleted files the next day</li> <li>2 = Compress deleted files immediately</li> </ul> <p><b>Default:</b> 1</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>
Days Untouched Before Compression = <i>number</i>	<p>Specify the number of days the system waits after a file was last accessed before it is compressed.</p> <p><b>Supported values:</b> 0 to 100000 <b>Default:</b> 14</p> <p>This parameter can be set in the startup.ncf file.</p> <p><b>Owning NLM:</b> filesys.nlm</p>

## A.18.4 Disk Parameters

The following Disk parameters control aspects of disk reads and writes.



Parameter	Description
Auto scan for devices = <i>value</i>	<p>Automatically scan for devices when a HAM is loaded. If turned off, operator must issue a Scan for New Devices on the command line before the devices will be seen.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> On</p> <p><b>Owning NLM:</b> <i>nwpa.nlm</i></p>
Auto load of CDM modules = <i>value</i>	<p>Automatically scan and load defined generic Custom Device Modules (CDMs). If turned off, the desired CDMs must be loaded manually or added to the startup.ncf file.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> On</p> <p><b>Owning NLM:</b> <i>nwpa.nlm</i></p>
Sequential Elevator Depth = <i>number</i>	<p>Set the maximum elevator depth for sequential requests.</p> <p>Media Manager sends the number of sequential requests up to this value to the same device. When the device contains this number of requests and another device in the mirror group is empty, Media Manager begins sending requests to the idle device.</p> <p><b>Supported values:</b> 0 to 4294967295 <b>Default:</b> 8</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM:</b> <i>mm.nlm</i></p>
Enable IO Handicap Attribute = <i>value</i>	<p>Enable drivers and applications to inhibit read requests from one or more devices.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> Off.</p> <p>Setting this parameter to On enables the Inhibit attribute to function. Setting this parameter to Off prevents the attribute from functioning.</p> <p>Do not set this attribute to On unless instructed to do so by a device manufacturer.</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM:</b> <i>mm.nlm</i></p>
Mirrored Devices Are Out of Sync Message Frequency = <i>time</i>	<p>Set the frequency (in minutes) for checking out-of-sync devices.</p> <p><b>Supported values:</b> 5 to 9999 <b>Default:</b> 28</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM:</b> <i>mm.nlm</i></p>

Parameter	Description
Remirror Block Size = <i>number</i>	<p>Set the remirror block size in 4 KB increments. (1=4 KB, 2=8 KB, 8=32 KB, etc.)</p> <p><b>Supported values:</b> 1 to 8  <b>Default:</b> 1</p> <p><b>Owning NLM :</b> mm.nlm</p>
Concurrent Remirror Requests = <i>number</i>	<p>Set the number of remirror requests per Mirror object.</p> <p><b>Supported values:</b> 2 to 32  <b>Default:</b> 32</p> <p><b>Owning NLM :</b> mm.nlm</p>
Ignore Partition Ownership = <i>value</i>	<p>Specify whether partitions owned by other servers in a cluster can be activated on this server.</p> <p>Setting this parameter to On allows partitions owned by other servers in the cluster to be activated on this server.</p> <p><b>Supported values:</b> On, Off  <b>Default:</b> Off</p> <p><b>Owning NLM :</b> mm.nlm</p>
Ignore Disk Geometry = <i>value</i>	<p>Create nonstandard and unsupported partitions.</p> <p><b>Supported values:</b> On, Off  <b>Default:</b> Off</p> <p>If you set this parameter to On before modifying or creating a partition, the software ignores disk geometry when creating the partition. This allows you to create nonstandard partitions.</p> <p><b>Caution:</b> Setting this parameter to On might harm other file systems contained on the disk.</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM :</b> mm.nlm</p>
Multi-path Support = <i>value</i>	<p>Enable or disable support for multipath devices or adapter failover.</p> <p><b>Supported values:</b> On, Off  <b>Default:</b> Off</p> <p><b>Owning NLM :</b> mm.nlm</p>

Parameter	Description
Enable Hardware Write Back = <i>value</i>	<p>Enable hardware write back, if supported.</p> <p>Hardware write back means that I/O write requests might be cached at the device and succeeded before data is committed to the media. Hardware write back usually improves write performance.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> Off</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM :</b> mm.nlm</p>
Enable Disk Read After Write Verify = <i>value</i>	<p>Control whether information written to disk is read back and compared with the original data.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> Off</p> <p>Setting this parameter to On might decrease performance significantly.</p> <p>You can set this parameter in the startup.ncf file.</p> <p><b>Owning NLM :</b> mm.nlm</p>
Auto LFMount	<p>Control whether any local FAT volumes not recognized by DOS are automatically mounted and assigned a local drive letter when they are detected.</p> <p><b>Supported values:</b> On, Off <b>Default:</b> Off</p> <p>On automatically mounts and assigns a local drive letter when the drive is detected.</p> <p>When set to Off, you must use the LFVMOUNT comand to manually mount local FAT volumes.</p> <p><b>Owning NLM :</b> fatfs.nlm</p>

## A.18.5 Novell Storage Services Parameters

The following are the default settings in NetWare 6.5 and OES NetWare for the NSS file system SET parameters.

Parameters	Default Value
NSS Auto Locks HFS CD-ROM Disc in device	Off
NSS Auto Update CDHFS Volume Objects To NDS	Off
Mount ISO 9660 multi disc set as one CD9660 volume	Off
NSS Auto Locks CD-ROM Disc in device	Off
NSS Auto Update CD9660 Volume Objects To NDS	Off

Parameters	Default Value
Emulate Classic NetWare Directory Quotas	Off
NSS Work To Do Count	50
NSS Maximum Cache Balance Buffers Per Session	1024
NSS Cache Balance Timer	30
NSS Cache Balance Percent	60
NSS Cache Balance Enable	On
NSS Buffer Flush Timer	1
NSS Minimum OS Cache Buffers	256
NSS Minimum Cache Buffers	256
NSS File Flush Timer	10
NSS Closed File Cache Size	50000
NSS Open File Hash Shift	16
NSS Auth Cache Size	1024
NSS ASCII Name Cache Enable	On
NSS Name Cache Enable	On
NSS Name Cache Size	2111
NSS Low Volume Space Alerts	On
NSS Low Volume Space Warning Reset Threshold	15
NSS Low Volume Space Warning Threshold	10

## A.19 VERIFY (NetWare)

Use VERIFY to validate all blocks in an NSS pool. The Verify option fixes an NSS pool and its volumes by searching for inconsistent data blocks or other errors. This utility indicates if there are problems with the file system. If there are no problems with the pool, it is safe to activate the pool.

---

**NOTE:** This utility protects against system failures only, not hardware failures.

---

### A.19.1 Syntax

The PoolVerify command is a read-only assessment of the pool. At the server console, enter

```
nss /poolverify=[pool_name]
```

If you do not enter the NSS pool name parameter, you will be prompted to select it.

During the verification process, the NSS pool and its volumes are placed in maintenance mode. This means the NSS pool and the volumes on it are unusable until this process is finished, and the pool is activated and mounted.

If errors are found, they are reported to the screen, and the NSS pool is left in maintenance mode. Run REBUILD until no errors are found.

## A.19.2 Starting and Using VERIFY

Use VERIFY at the command line.

---

**NOTE:** Always back up your data. If an NSS pool uses several hard disks and one of the hard disks becomes corrupted, you must create a new NSS pool, then restore its NSS volumes from backup.

---

**1** Load NSS.

**2** Verify your rebuilt NSS pool, by entering

```
nss /poolverify=[pool_name]
```

If you do not enter the NSS pool name parameter, you will be prompted to select it.

Verifying a pool can take a up to a few hours, depending on the size of the pool.

**3** Check the output screen.

This screen indicates the time elapsed, the time remaining, the total elapsed time, the number of objects processed, etc.

## A.20 VOLUME

Use this command at the System Console prompt (NetWare) or the NSSCON prompt (Linux) to list mounted volumes.

### A.20.1 Syntax

VOLUME [*name*]

VOLUMES

Command	Description
volumes	List general information about all mounted volumes.
volume <i>name</i>	Specify the volume you want to display information about.

### A.20.2 Using VOLUME

When you execute VOLUME without the *name* parameter, a message similar to the following appears:

Mounted Volumes	Name Spaces	Flags
SYS	DOS	Cp Sa Mg

The output lists each mounted volume, its name spaces, and the flags that are set for that volume. The three possible flags are:

Cp—indicates that file compression is enabled on the volume

Sa—indicates that block suballocation is enabled on the volume

Mg—indicates that migration is enabled on the volume

When you execute `volume name`, the screen displays more detailed information about the specific volume, such as the following:

**Figure A-2** *Volume Information*

**Volume SYS information:**

```
Block Usage: 4 FAT blocks, 34 Directories blocks, 0 EDS blocks
Blocks: 1569 used of 4084      Directory Entries: 1314 used of 2176
EDS: 0 used of 0 extants      Extended Attributes: 0
Data Streams: 590, 668 (Cp), 9 (Cp Limbo), 242 (Cant Cp)
Sub Alloc Blocks: 321
NameSpaces: DOS               Flags: Cp Sa
```

# Upgrading Legacy NSS and NetWare Traditional Volumes

# B

This section discusses how to convert legacy NetWare® 5 NSS volumes and NetWare Traditional volumes to NSS volumes for use with Novell® Open Enterprise Server NetWare.

- [Section B.1, “Migrating Legacy NSS Volumes and NetWare Traditional Volumes to OES Linux,” on page 399](#)
- [Section B.2, “Copying and Upgrading Legacy Volumes with Volume Copy Upgrade,” on page 399](#)
- [Section B.3, “Upgrading Legacy NSS Volumes,” on page 406](#)
- [Section B.4, “Upgrading Legacy NSS Volumes in Clusters,” on page 407](#)

## B.1 Migrating Legacy NSS Volumes and NetWare Traditional Volumes to OES Linux

After upgrading a Netware server to OES NetWare, it is possible for a Traditional volume to still reside on that server. This creates a problem when migrating from OES NetWare to OES Linux because the Traditional volume data cannot be read on OES Linux. You should move the data from the Traditional volume to an NSS volume before upgrading, either by directly copying it, using the Volume Copy Utility (VCU), or using the Server Consolidation and Migration Utility (SCMU).

For information about VCU, see [Section B.2, “Copying and Upgrading Legacy Volumes with Volume Copy Upgrade,” on page 399](#).

For information about SCMU, see [Novell Server Consolidation and Migration Toolkit Administration Guide](#) (<http://www.novell.com/documentation/scmt/index.html>).

## B.2 Copying and Upgrading Legacy Volumes with Volume Copy Upgrade

Volume Copy Upgrade (VCU) is a volume conversion utility that lets you copy metadata and user data from legacy NetWare Traditional volumes and NSS volumes to OES NetWare NSS volumes. You can also use VCU to copy NSS volumes. Volume copy upgrade moves the volume to the NSS pool and all of the volume’s file and directory attributes remain the same. Copying the data to the new location can take several hours, depending on the size of the volume. You can run VCU from the server command line or from a GUI application that runs on the server or on a workstation.

- [Section B.2.1, “Understanding Volume Copy Upgrade,” on page 400](#)
- [Section B.2.2, “Using VCU from the Command Line,” on page 401](#)
- [Section B.2.3, “Using VCU from a Java Application,” on page 402](#)
- [Section B.2.4, “After the Volume Copy or Volume Upgrade,” on page 406](#)

## B.2.1 Understanding Volume Copy Upgrade

Volume Copy Upgrade (`vcu.nlm`) provides the following capabilities:

- VCU creates an inherited rights filter file, called `irf.out`. This file contains information about whether the Inherited Rights Filters were moved correctly during a volume copy or restore. Mismatches are preceded by three asterisks (\*\*).

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

---

**IMPORTANT:** When you use VCU to upgrade a Traditional NetWare `sys :` volume to the NSS `sys :` volume, the implied trustee rights do not transfer with this method. After you upgrade volume `sys :` to NSS, you must manually add a public trustee with Read and File Scan rights to the `sys :login` directory.

---

- You can use VCU to copy and restore NSS volumes and to copy NetWare Traditional volumes to NSS volumes.
- You can run `vcu.nlm` from a Windows client.

Before you begin to upgrade NetWare Traditional volumes, make sure that the NSS pool that you are copying data to has enough space to accommodate the source volume's uncompressed data. Use the `ndir` command to determine how much space is saved by using compression on the volume. For information, see [Section 17.11, "Viewing the File Attributes with the NDIR Command \(NetWare\)," on page 242](#). Add the space saved to the space used to find the total uncompressed volume size of the Traditional volume.

An NSS volume is about 110% of the size of NetWare Traditional volumes because files and directories stored on NSS have more metadata associated with them. For example, if you want to move 2 GB of viable data (disregard free blocks or salvage blocks), the destination NSS volume needs to have at least 2.2 GB of free space available.

---

**IMPORTANT:** Upgrade only one volume at a time.

---

The syntax for `vcu.nlm` is:

---

Copy Upgrade	<code>vcu [/p /l /d /m] origVolName nssPool</code>
Restore the Original Volume	<code>vcu [/r /m] origVolName renamedVolName</code>
Help	<code>vcu /h</code>

---

The VCU parameters are listed in the following table:

VCU Parameters and Help	Description
<code>origVolName</code>	The name of the original NetWare Traditional volume or NSS volume that you want to copy to a new NSS volume of the same name.
<code>nssPool</code>	The pool where the new NSS volume resides.
<code>renamedVol</code>	The name assigned to the original NetWare Traditional volume or NSS volume that you successfully copied to the NSS volume.



VCU Parameters and Help	Description
/h	Lists the syntax and options for the command.

The following table lists the possible VCU switches when copying or upgrading volumes. The switches are optional; you can include them in any order in your VCU command line.

VCU Option	Description
/d	Deletes the original volume if the copy process is successful. If you choose to delete the Traditional volume, the new NSS volume retains the name of the original volume.
/l	Does not write errors to log file ( <code>Dst_Vol:error.out</code> ).
/m	Renames the physical volume and its default Novell eDirectory™ objects.
/p	Does not display the entire copy status on the VCU screen. You get only a brief message after the volume copy is completed.

The following table lists the possible VCU options when restoring volumes:

VCU Option	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 its original name.)
/m	Renames the physical volume and its default Novell eDirectory objects.

For example, to convert a NetWare Traditional volume named `users` to an OES NetWare NSS volume named `users`, enter:

```
vcu /m users p_users
```

To continue the example, suppose that you decide to undo the copy and return to the legacy volume. To remove the NSS volume `users` and restore the original name to the NetWare Traditional volume, enter:

```
vcu /r /m users users_old
```

## B.2.2 Using VCU from the Command Line

- [“Copying Volumes” on page 401](#)
- [“Restoring Volumes” on page 402](#)

### 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 At the command line, specify the volume that you want to copy, then specify the existing NSS pool (on the same server) 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
```

For customization options, see [“Understanding Volume Copy Upgrade” on page 400](#).

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 in 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.

- 3 Select Yes to rename your volumes.

After the volume copy is finished, VCU asks you if you want to rename your volumes. If you select *Yes*, VCU renames your original volume to *originalvolumename\_old* and the new NSS volume from *originalvolumename\_new* back to the name of the original volume. If you select *No*, you must manually rename the new volume later.

---

**NOTE:** 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*).

---

- 4 Review the information in [“After the Volume Copy or Volume Upgrade” on page 406](#) and perform any required follow-up actions.

## Restoring Volumes

At the server console, enter

```
vcu /r VCUTargetVolName VCUTargetVolName_old
```

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

```
vcu /r Vol2_new Vol2_old
```

VCU would delete Vol2\_new, then rename Vol2\_old back to its original name of Vol2.

---

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

---

## B.2.3 Using VCU from a Java Application

You can also upgrade volumes using the VCU 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, enter

vcux

- At the GUI console, click *Novell > Utilities > Volume Copy Upgrade*.

## Accessing VCU from a Workstation

- 1 Make sure that 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.

---

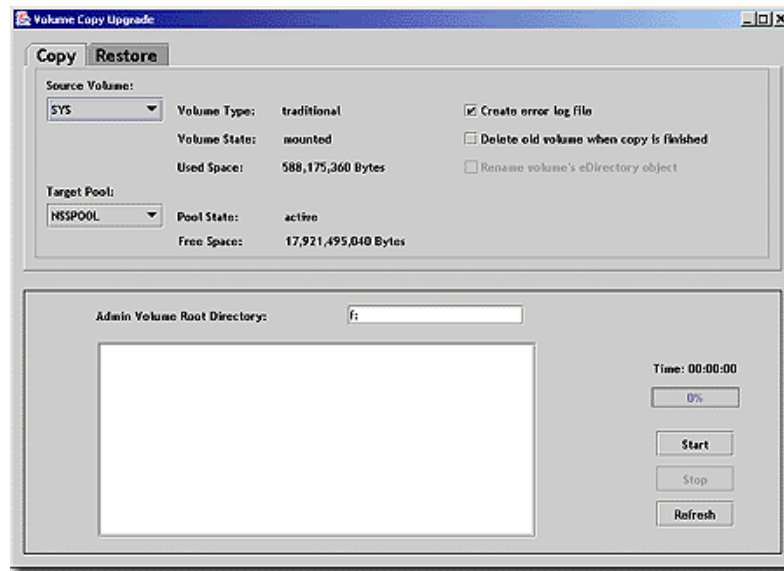
**WARNING:** Before using this utility to copy Traditional or NSS volumes, stop all applications that access the volume you are copying. This is especially important for the `sys:` volume. If any applications are open on the `sys:` volume when you begin the volume copy, your `sys:` volume might become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

---

- 2 Launch VCU.

For instructions, see “**Using VCU from a Java Application**” on page 402.

The following graphic is an example of the VCU Copy screen.



- 3 Click the *Copy* tab.
- 4 (Conditional) If you are running VCU from a workstation, specify the drive letter of the drive that you mapped to the `_Admin` volume, followed by a colon (:), in the *Admin Volume Root Directory* text box, then click *Refresh*.

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.

If you have recently added a pool or volume via 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 and saves it at the root directory of the destination volume.  For more information, see <a href="#">Section B.2, "Copying and Upgrading Legacy Volumes with Volume Copy Upgrade,"</a> on page 399.
Delete Old Volume When Copy Is Finished	Deletes the old (original) physical volume; however, the volume's eDirectory object remains.
Rename Volume's eDirectory Object	Renames the volume's default eDirectory object.  For more information, see <a href="#">"Copying Volumes"</a> on page 401.

- 7 Click Start to begin the volume copy.

If you stop the volume copy for any reason, then 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 asks you if you want to rename your volumes. If you click *Cancel*, VCU is finished. If you click *OK*, VCU renames 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 or Volume Upgrade](#)” on page 406 and perform any required follow-up actions.

## Restoring Volumes

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

---

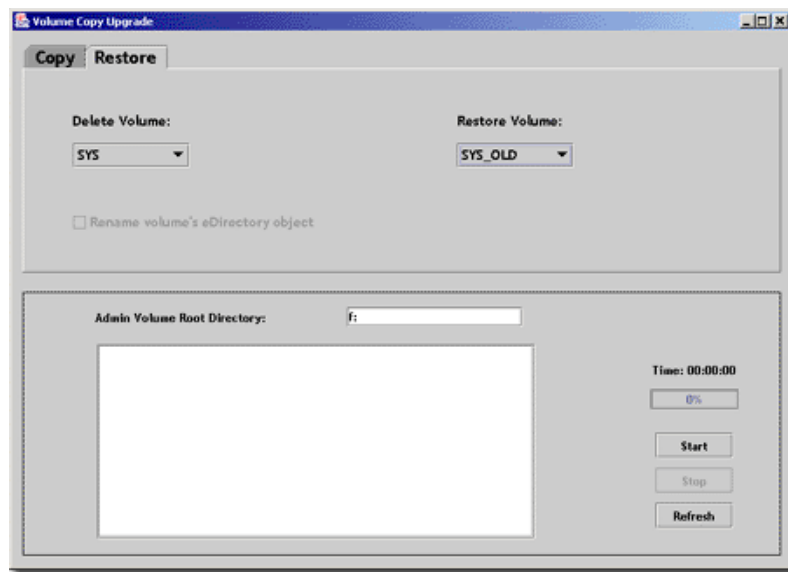
**WARNING:** Before using this utility to copy Traditional volumes or NSS volumes, stop all applications that access the volume you are copying. This is especially important for the `sys:` volume. If any applications are open on the `sys:` volume when you begin the volume copy, your `sys:` volume might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

---

- 2 Launch VCU.

For instructions, see “[Using VCU from a Java Application](#)” on page 402.

The following graphic is an example of the VCU Restore dialog box.



- 3 Click the *Restore* tab.
- 4 (Conditional) If you are running VCU from a workstation, specify the drive letter of the drive that you mapped to the `_Admin` volume, followed by a colon (:), in the *Admin Volume Root Directory* text box, then click *Refresh*.

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

- 5 From the drop-down list, select the volume that you want to delete and restore.

If you have recently added a pool or volume via 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 is automatically renamed to the name of the volume that you are deleting.

## B.2.4 After the Volume Copy or Volume Upgrade

- Restart the server if you are upgrading the `sys :` volume to ensure that 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 need to either retain the data in the Traditional volume, or copy the data back to the Traditional volume.

---

## B.3 Upgrading Legacy NSS Volumes

Follow this procedure to copy and upgrade your legacy NSS volumes to OES NetWare NSS volumes after you have upgraded the operating system to OES NetWare. When you upgrade the operating system of a NetWare 5.x server to an OES NetWare server, the NSS volumes are not automatically upgraded and do not function until you upgrade them.

If you upgrade your server from NetWare 6.0 to NetWare 6.5 or OES NetWare, the NSS media format upgrade occurs automatically in the background as you use the volume. Typically, this can take up to 21 days, depending on the amount of data you have on the volume and whether you are using compression. There are no negative performance implications for this gradual conversion.

---

**IMPORTANT:** You do not need to manually upgrade legacy NSS volumes in clusters after an upgrade. The cluster automatically calls to perform the upgrade.

---

- 1 View and verify 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 that eDirectory is running, you need to verify 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`.

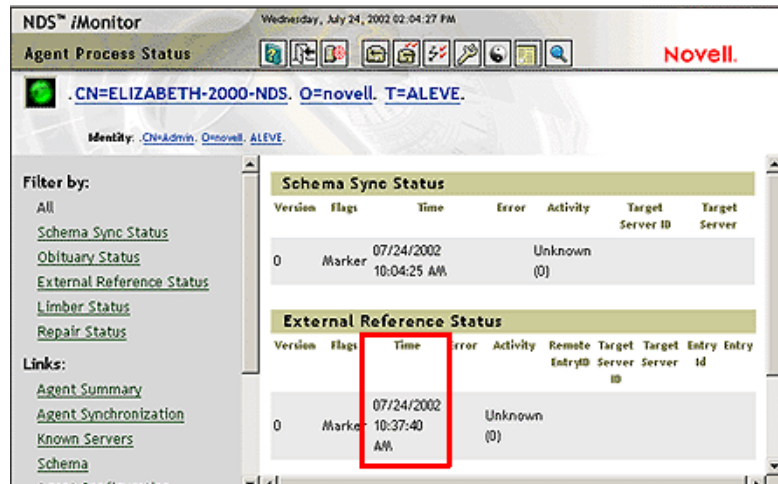
For example: `http://192.168.1.1/nds`

- 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 to OES NetWare. 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 to OES NetWare, 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.



2 At the server console, enter one of the following:

- `nss /zlssvolumeupgrade=all`

Specify “all” to upgrade all the NetWare 5 volumes on the selected server.

- `nss /zlssvolumeupgrade=volume`

Specify the volume to upgrade one of the NetWare 5 volumes on the selected server.

After you upgrade the volumes, use the Mount All command to mount all of the volumes. You can also mount the volumes individually.

## B.4 Upgrading Legacy NSS Volumes in Clusters

If your legacy NSS volumes are in clusters, the upgrade of your cluster services also automatically upgrades your legacy NSS volumes. Novell Cluster Services™ for OES NetWare includes a pre-upgrade step that backs up the Trustees for your files and directories. The cluster automatically calls `nss /zlssvolumeupgrade` to perform the upgrade from the NetWare 5 NSS volume media format to OES NetWare NSS volume media format. For information, see “[Upgrading Novell Cluster Services](#)” in the *OES Novell Cluster Services 1.8.1 Administration Guide for NetWare*.





# Using EVMS to Manage Devices with NSS Volumes (Linux)

# C

To use Novell® Storage Services™ data volumes on your Open Enterprise Server Linux system, the devices you use for those volumes must be managed by the Enterprise Volume Management System (EVMS). This section discusses how to configure EVMS services on Linux to allow you to use NSS.

- [Section C.1, “Installing Linux with EVMS as the Volume Manager on a System Hard Drive,” on page 409](#)
- [Section C.2, “Viewing Free Space Available to EVMS-Managed Devices,” on page 409](#)
- [Section C.3, “Making Devices Available to EVMS,” on page 410](#)
- [Section C.4, “Additional EVMS Documentation,” on page 415](#)

## C.1 Installing Linux with EVMS as the Volume Manager on a System Hard Drive

To use NSS on the same drive that contains your boot volume and system volume, that drive must be managed by EVMS. The default volume manager in OES Linux is the Linux Volume Manager (LVM), not EVMS. Unless you make your Linux swap and system volumes manageable by EVMS, you cannot create NSS data volumes on that drive.

The simplest solution is to put your swap and system volumes on one device, then use different devices for your NSS volumes.

You must install Linux with EVMS as the volume manager of the system hard drive in either of the following scenarios:

- Your system has a single hard drive and you want to use its free space for NSS data volumes.
- Your system has multiple hard drives and you want to use free space for NSS data volumes from any drive that contains the Linux boot, system, or swap volumes.

For information, see [“Installing Linux with EVMS as the Volume Manager of the System Device”](#) in the *OES Linux Installation Guide*.

## C.2 Viewing Free Space Available to EVMS-Managed Devices

When creating an NSS pool on Linux, you can use free space only from devices managed by EVMS. If you use multiple volume managers, make sure that the devices you select for the pool are managed by EVMS.

The OES Linux install uses LVM, but NSS requires EVMS. The Linux 2.6 kernel prevents multiple volume managers from managing the same device. If you use LVM and EVMS to manage different devices on the server, EVMS might display free space on the LVM-managed devices that EVMS cannot actually manage.

To avoid confusion, you can prevent display of space that is unavailable to EVMS. Edit the `/etc/evms.conf` file to add the device to the exclude list of the `sysfs_devices` section.

For example, in a typical OES Linux install, the `hda` device is your boot device and contains the root volume. Unless you have previously made EVMS the volume manager for that device, it is managed by LVM and cannot be used to create NSS pools. To prevent the display of space from `hda`, edit the `/etc/evms.conf` file as follows:

```
sysfs_devices {  
  
    ...  
  
    exclude = [hda]  
  
}
```

## C.3 Making Devices Available to EVMS

On your Linux system, if you've already installed Linux with a non-EVMS volume manager for one or more of the devices, use one of the following was to make the devices available to EVMS:

- [Section C.3.1, “Mounting Your Root File System with EVMS,” on page 410](#)
- [Section C.3.2, “Converting LVM Volumes to EVMS Volumes,” on page 414](#)
- [Section C.3.3, “Patching the 2.6 Kernel \(Unsupported\),” on page 415](#)

### C.3.1 Mounting Your Root File System with EVMS

If you did not configure EVMS as the root partition manager during installation, the following post-install workaround configures the device to be managed by EVMS. This enables you to create NSS data volumes on single disk systems.

The following procedures assume that you installed OES Linux with three partitions on a single SCSI device named `sda`:

```
/dev/sda1  reiserfs  /boot  
  
/dev/sda2  swap      swap  
  
/dev/sda3  reiserfs  /
```

---

**IMPORTANT:** Make sure to modify the following procedures as necessary for your specific setup.

---

- [“Disable boot.lvm and boot.md” on page 411](#)
- [“Enable the boot.evms Service” on page 411](#)
- [“Edit the /etc/init.d/boot.evms Script” on page 411](#)
- [“Edit the /etc/fstab File” on page 412](#)
- [“Reboot the Server” on page 412](#)
- [“Force the RAM Disk to Recognize the Root Partition” on page 412](#)
- [“Reboot the Server” on page 413](#)
- [“Verify that EVMS Manages the Boot, Swap, and Root Partitions” on page 413](#)

## Disable boot.lvm and boot.md

Disable `boot.lvm` (handles devices for Linux Volume Manager) and `boot.md` (handles multiple devices in software RAIDs) so they do not run at boot time. In the future, you want `boot.evms` to run at boot time instead.

**1** In YaST, click *System > Runlevel Editor > Expert Mode*.

**2** Select *boot.lvm*.

**3** Click *Set/Reset > Disable the Service*.

**4** Select *boot.md*.

**5** Click *Set/Reset > Disable the Service*.

**6** Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

**7** Continue with [“Enable the boot.evms Service” on page 411](#).

## Enable the boot.evms Service

Enable the `boot.evms` service so that it boots devices when you restart the server.

**1** In YaST, click *System > Runlevel Editor > Expert Mode*.

**2** Select *boot.evms*.

**3** Click *Set/Reset > Enable the Service*.

The *B runlevel* option is automatically selected.

**4** Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

**5** Continue with [“Edit the /etc/init.d/boot.evms Script” on page 411](#).

## Edit the /etc/init.d/boot.evms Script

**1** Open the `/etc/init.d/boot.evms` script in a text editor.

**2** Add the following lines to the Stop section:

```
mount -n -o remount,rw /
echo -en "\nDeleting devices nodes"
rm -rf /dev/evms
mount -n -o remount,ro /
```

For example, the Stop section looks like this after the edit:

```
stop)
    echo -n "Stopping EVMS"
    mount -n -o remount,rw /
    echo -en "\nDeleting devices nodes"
    rm -rf /dev/evms
    mount -n -o remount,ro /
    rc_status -v
;;
```

- 3 Save the file.

The changes do not take effect until the server is restarted. Do not reboot at this time.

- 4 Continue with [“Edit the /etc/fstab File” on page 412.](#)

### Edit the /etc/fstab File

When you boot the system, the kernel reads the `/etc/fstab` file to identify which file systems should be mounted and then mounts them. This file contains a table of file system information about the `/boot`, `swap`, and `root (/)` partitions plus other partitions and file systems you want to mount.

You must edit the `/etc/fstab` file to modify the mount location of these three partitions so they are mounted under the `/dev/evms` directory. For example, change `/dev/sda1` to `/dev/evms/sda1`.

Make sure to replace `sda1`, `sda2`, and `sda3` with the device names you used for your partitions.

---

**IMPORTANT:** When working in the `/etc/fstab` file, do not leave any stray characters or spaces in the file. This is a configuration file, and it is highly sensitive to such mistakes.

---

- 1 Open the `/etc/fstab` file in a text editor.

- 2 Locate the line that contains the `/boot` partition.

For example, if your `/boot` partition uses device `sda1` and the *Reiser* file system, look for a line similar to this:

```
/dev/sda1 /boot reiser defaults 1 1
```

- 3 In the *Device Name* column, modify the mount location of the `/boot` partition from `/dev` to `/dev/evms` so it can be managed by EVMS. Modify only the device name by adding `/evms` to the path:

```
/dev/evms/sda1 /boot reiser defaults 1 1
```

- 4 Repeat [Step 2](#) and [Step 3](#) to edit the Device Name entry in the lines for the `swap` and `root (/)` partitions.

For example, change `/dev/sda2` to `/dev/evms/sda2`, and change `/dev/sda3` to `/dev/evms/sda3`.

- 5 Save the file.

- 6 Continue with [“Reboot the Server” on page 413.](#)

### Reboot the Server

- 1 Reboot the server to apply the post-install configuration settings.

- 2 Continue with [“Force the RAM Disk to Recognize the Root Partition” on page 412.](#)

### Force the RAM Disk to Recognize the Root Partition

Although you modified the `/etc/fstab` file so that everything now mounts from the `/dev/evms` path, the RAM disk still tries to load the root (`/`) partition from the `/dev` path. You must force the RAM disk to update its information.

- 1 At a terminal prompt, enter the EVMS NCurses command:

```
evmsn
```

- 2 Review the output to verify that EVMS shows only the `/boot` and `swap` partitions as active in EVMS.

You should see the following devices mounted (with your own partition names, of course) for these two partitions:

```
/dev/evms/sda1
```

```
/dev/evms/sda2
```

- 3 Edit `/sbin/mkinitrd` to change line 2513 from

```
root_evms=
```

to

```
root_evms=1
```

then save the file.

- 4 At a terminal prompt, enter

```
/sbin/mkinitrd
```

This forces `initrd` to get the updated `/dev/evms` path information for the root (`/`) partition.

- 5 Edit `/sbin/mkinitrd` to change line 2513 from

```
root_evms=1
```

to

```
root_evms=
```

then save the file.

- 6 Continue with [“Reboot the Server” on page 413](#).

## Reboot the Server

- 1 Reboot the server to apply the post-install configuration settings.
- 2 Continue with [“Verify that EVMS Manages the Boot, Swap, and Root Partitions” on page 413](#).

## Verify that EVMS Manages the Boot, Swap, and Root Partitions

- 1 At a terminal prompt, enter the EVMS `NCurses` command:

```
evmsn
```

- 2 Review the output to verify that EVMS shows the `/boot`, `swap`, and `root (/)` partitions as active in EVMS.

You should see the following devices mounted (with your own partition names, of course) for these three partitions:

```
/dev/evms/sda1
```

```
/dev/evms/sda2
```

```
/dev/evms/sda3
```

Everything should now mount from the `/dev/evms` structure. You can now use the Storage plug-in to iManager to create and manage NSS pools and volumes, or use NSSMU and other NSS commands and utilities as desired.

## C.3.2 Converting LVM Volumes to EVMS Volumes

The following instructions were modified from [Converting to EVMS \(http://evms.sourceforge.net/convert.html\)](http://evms.sourceforge.net/convert.html) at the EVMS Project Web site.

- 1 In a terminal console, run the EVMSGUI by entering

```
evmsgui
```

- 2 In the Volume panel, review the names that EVMS reports for your existing volumes to find the volume that represents your system, swap, and boot volumes, then write down the names for future reference.

All volume names begin with the path `/dev/evms`. For example, an LVM volume called `/dev/system` appears in EVMS as `/dev/evms/lvm/system`.

- 3 In a text editor, edit the `/etc/fstab` file to use the EVMS volume names.

For example, change the following entry for an LVM volume from this

```
/dev/system / reiserfs defaults 1 2
```

to this

```
/dev/evms/lvm/system / reiserfs defaults 1 2
```

---

**IMPORTANT:** Make sure not to leave any stray characters or spaces in the line.

---

With these changes, each time your system boots, your filesystems will be mounted using the EVMS volume manager.

- 4 Update the boot scripts.

The command `evms_activate` must be run from your boot-scripts in order to activate your volumes so they can be mounted.

If you run software-RAID or LVM tools during your boot-scripts, remove or disable those commands.

- 5 Update the boot-loader.

To mount your root filesystem using EVMS, you must install the EVMS `init-ramdisk` image and update your boot-loader configuration. For information about updating the boot-loader, see [Configuring Your Boot-Loader for EVMS \(http://evms.sourceforge.net/install/boot.html\)](http://evms.sourceforge.net/install/boot.html) and [Mounting Your Root Filesystem Through EVMS \(http://evms.sourceforge.net/install/root.html\)](http://evms.sourceforge.net/install/root.html) (the root-filesystem sections of the installation instructions for EVMS).

- 6 Reboot your system.

Now that all of these items have been configured, you just need to reboot your machine one more time. When it boots, the kernel loads the `init-ramdisk`, which runs the EVMS tools to activate your volumes and mount your root filesystem. Then your boot-scripts run the EVMS tools once more to make sure your `/dev/evms/` directory correctly reflects the current state of your volumes. Finally, the remaining EVMS volumes are mounted as specified in your `/etc/fstab` file. Everything else on your system should start up as you would normally expect.

### C.3.3 Patching the 2.6 Kernel (Unsupported)

Apply the BD-Claim Patch, which is a reversal of the patch that prevents Device-Mapper and the kernel's built-in partitions from using the same disk at the same time. This patch is not supported by the kernel community, and in fact removes functionality that they specifically added. However, it will allow you to share your disks between EVMS and the kernel's built-in partitioning code, if that's the choice you wish to make for your system.

---

**WARNING:** Adding this patch makes your server an unsupported configuration with respect to Novell Technical Support.

---

See [BD-Claim Patch \(http://evms.sourceforge.net/install/kernel.html#bdclaim\)](http://evms.sourceforge.net/install/kernel.html#bdclaim) for more information.

## C.4 Additional EVMS Documentation

For more information about the Enterprise Volume Management System, see the following:

- [EVMS Project at SourceForge.net \(http://sourceforge.net/projects/evms/\)](http://sourceforge.net/projects/evms/)
- [EVMS GUI User Guide \(http://evms.sourceforge.net/user\\_guide/#assignseggui\)](http://evms.sourceforge.net/user_guide/#assignseggui)
- [EVMS GUI Screen Shots \(http://evms.sourceforge.net/gui\\_screen/\)](http://evms.sourceforge.net/gui_screen/)





# Comparison of NSS on NetWare and NSS on Linux

# D

This section compares features and capabilities of Novell® Storage Services™ on the NetWare® and Linux platforms for OES.

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
Management interfaces	Novell iManager	Novell iManager
	NSSMU for NetWare	NSSMU for Linux
		EVMS (evmsgui)
	Utilities in the server console (nssmu, rights, flag)	Utilities in the server console (nssmu, rights, nsscon, attrib, ravsui, ravview)
	Commands in the server console	Commands in the NSS Console (nsscon)
	Novell Remote Manager for NetWare (primarily for trustee management, file and directory purge and salvage)	Novell Remote Manager for Linux (browse only)
	Novell NetStorage	Novell NetStorage
	ConsoleOne® (DFS junctions, VLDB, and file system trustees and trustee rights)	Not available (DFS junction support is not available for NSS on Linux.)
File system trustees and trustee rights to control access to directories and files	Novell Remote Manager for NetWare	
	Novell NetStorage (via Web browser only, not WebDAV)	Novell NetStorage (via Web browser only, not WebDAV)
	Novell Client™. See “Using the Novell Client to Manage Trustees and Trustee Rights” in the <i>File Systems Management Guide for OES</i> .	Novell Client. “Using the Novell Client to Manage Trustees and Trustee Rights” in the <i>File Systems Management Guide for OES</i> .
	Rights utility for NetWare. See “Trustee Rights Utility for NetWare” in the <i>File Systems Management Guide for OES</i> .	Rights utility for Linux. See “Trustee Rights Utility for Linux” in the <i>File Systems Management Guide for OES</i> .
	ConsoleOne	

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
File system directory and file attributes to control functions available for directories and files	Novell NetStorage	Novell NetStorage
	Novell Client (NCP™)	Novell Client (NCP)
	Rights Utility for Linux	Rights Utility for Linux
	Novell Remote Manager for NetWare. See “ <a href="#">Managing Attributes for Directories and Files</a> ” in the <i>File Systems Management Guide for OES</i> .	Novell Remote Manager for Linux. See “ <a href="#">Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions</a> ” in the <i>File Systems Management Guide for OES</i>
	ConsoleOne	
Directory quotas management (requires the Directory Quotas attribute for the volume)	Novell NetStorage	Novell NetStorage
	Novell Client	Novell Client
	Novell Remote Manager for NetWare	
User space quota management (requires the User Space Quotas attribute for the volume )	Novell iManager	Novell iManager
		Make sure to Linux-enable users with Linux User Management before you set the quota.  User space quotas are available in OES SP1 and later.
Default mount location for NSS pools	Not applicable	/opt/novell/nss/mnt/.pools/
Default mount location for NSS volumes	Server root	/media/nss/
Default mount location for devices managed by EVMS	Not applicable	/dev/evms/
File system type (as recognized and reported by the operating system)	nss	nssvol
File protocols	NCP	NCP
	AFP, NFS, and CIFS	NFS and CIFS (Samba) require users to be Linux-enabled with Linux User Management.  AFP is available only through third-party software.
Interface	64-bit	64-bit
Character format	Unicode	Unicode
Maximum device size recognized (physical or logical)	2 TB	2 TB
Maximum partition size	2 TB	2 TB

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
Maximum number of partitions (logical or physical devices) per pool	No practical limit  A partition must have more than 10 MB and less than 2 TB of available space to be recognized by NSS.	No practical limit  A partition must have more than 10 MB and less than 2 TB of available space to be recognized by NSS.
Maximum pool size (using at least 4 partitions of up to 2 TB each)	8 TB	8 TB
Minimum pool size	10 MB	10 MB
Maximum size of a volume	Up to 8 TB, depending on the pool size and available space in the pool  Volume quotas can be overbooked. For information about overbooking, see <a href="#">Section 13.1, "Understanding NSS Volumes," on page 181</a> . For information, see <a href="#">Section 18.2, "Managing NSS Volume Quotas," on page 245</a> .	Up to 8 TB, depending on the pool size and available space in the pool  Volume quotas can be overbooked. For information about overbooking, see <a href="#">Section 13.1, "Understanding NSS Volumes," on page 181</a> . For information, see <a href="#">Section 18.2, "Managing NSS Volume Quotas," on page 245</a> .
Maximum file size	Up to 8 TB, depending on the volume size and available space in the volume	Up to 8 TB, depending on the volume size and available space in the volume
Maximum number of files per volume	Up to 8 trillion, regardless of how many name spaces are loaded	Up to 8 trillion, regardless of how many name spaces are loaded
Maximum number of files open concurrently	1 million	1 million
Maximum number of volumes per server	255 plus the <code>sys:</code> volume.  You can mount NSS volumes beyond 256, but they are not visible or accessible through the normal NetWare APIs.	No practical limit on the number of NSS data volumes
Time to mount a volume	Requires only a few seconds  NSS uses a journaling file system and does not need to scan the entire file system to create a directory entry table (DET) and to load a File Allocation Table (FAT).	Requires only a few seconds  NSS uses a journaling file system and does not need to scan the entire file system to create a directory entry table (DET) and to load a File Allocation Table (FAT).

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
Volume name space	<p>Accommodates all name spaces (DOS, Macintosh, Long, and UNIX).</p> <p>In OES SP1 and later, you can use the Lookup Name Space attribute for the volume to specify the default name space to use when activating the volume.</p> <p>Directory names and file names are case insensitive.</p>	<p>Accommodates all name spaces (DOS, Macintosh, Long, and UNIX).</p> <p>In OES SP1 and later, you can use the Lookup Name Space attribute for the volume to specify the default name space to use when mounting the volume.</p> <p>UNIX is the default name space, which is case sensitive. However, you can specify other name spaces on mounting the NSS volume to make its directory names and filenames case insensitive. For information, see <a href="#">Section 13.16, "Mounting an NSS Volume on Linux to Support Case Insensitive File Names (Linux)," on page 197.</a></p>
Minimum server memory required to activate a volume	Requires only 4 MB available RAM to activate a single volume of any size and any number of files. Loads a file's metadata into memory only as you access the file.	Requires only 4 MB available RAM to activate a single volume of any size and any number of files. Loads a file's metadata into memory only as you access the file.
File access time	Same for each file, regardless of its location on the volume.	Same for each file, regardless of its location on the volume.
Error correction and data recovery time on system failure	<p>Journaling file system logs changes.</p> <p>On system failure, replays the most recent transactions to confirm validity, then repairs errors or rolls back to the original condition, typically in 15 to 60 seconds, unless the volume is corrupted.</p>	<p>Journaling file system logs changes.</p> <p>On system failure, replays the most recent transactions to confirm validity, then repairs errors or rolls back to the original condition, typically in 15 to 60 seconds, unless the volume is corrupted.</p>
Repair of corrupted volume	Ongoing journaling of the file system; if corrupted, use verify and rebuild functions.	Ongoing journaling of the file system; if corrupted, use the ravsui utility to verify and rebuild the volume.
Time to repair corrupted volume	From a few seconds to several hours, depending on the volume size.	From a few seconds to several hours, depending on the volume size.
Multiple connection paths to storage media	Yes, NSS multipath I/O	NSS multipath I/O is not supported; use a Linux-based multiple path I/O solution.

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
Software RAID support	RAID 0 (striping) RAID 1 (mirroring) RAID 5 (striping with parity) RAID 10 (mirroring RAID 0 devices) RAID 15 (mirroring RAID 5 devices)	RAID 0 (striping) RAID 1 (mirroring) RAID 10 (mirroring RAID 0 devices) For OES SP1 and later, the following additional software RAIDs are supported: <ul style="list-style-type: none"> <li>RAID 10 (mirroring RAID 0 devices)</li> <li>RAID 15 (mirroring RAID 5 devices)</li> </ul>
Volume encryption	Yes	Yes, only for OES SP1 and later  You must mount encrypted volumes only from NSSMU on the first mount after a system reboot so that you can enter the password. The nsscon utility does not support entering a password from the command line.
Data shredding	Yes, up to 7 times	Yes, up to 7 times
File compression	Yes	Yes
Data migration	Yes	Yes (no known users at this time)
Directory quotas	Yes	Yes
User space quotas (user space restrictions)	Yes	Yes, only for OES SP1 and later  Make sure to Linux-enable users with Linux User Management before you set the quota.
Salvage or purge deleted files, directories, or volumes	Yes	Yes
Transaction Tracking System™ (TTS™)	Yes	Not supported
Read ahead blocks	Yes	Yes
File save time	Offers the Flush Files Immediately attribute for NSS volumes to write files to disk on save instead of waiting for the next disk write cycle. This helps prevent possible data loss between disk write cycles.	Offers the Flush Files Immediately attribute for NSS volumes to write files to disk on save instead of waiting for the next disk write cycle. This helps prevent possible data loss between disk write cycles.

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
File-level snapshot (make a temporary snapshot copy of an open file for backup)	Yes; allows back up of a copy of a file that is open nonexclusively so that it is not necessary to deactivate the volume for backup	Not supported
Pool snapshot (retain point-in-time version of a pool using block-level copy on write)	Yes; allows backup of block-level changes only, without deactivating the volume.  Uses a brief freeze-release process to capture information for last remaining open files.	Not supported
Backup support	Reviews only a list of files modified since the previous backup	No APIs are available.
Backup systems support	TSA, TSA600 for files over 2 GB, SMS, and Enhanced SBackup	Yes
Distributed File Services for moving and splitting NSS volumes	Yes	Not supported
Novell Archive and Version Services	Yes	Not supported
Device maintenance support	Activate and deactivate devices by pool.	Activate and deactivate devices by pool.
CD and DVD device recognition	Automatic process with full support for UDF, ISO 9660, and Macintosh HFS formats.  Use CDs and DVDs as read-only NSS volumes.	No; use Linux traditional file system options instead.
CD and DVD image files	Activate as read-only NSS volumes.	No; use Linux traditional file system options instead.
Operating system version detection	Default process	Default process

Feature Description	NSS for OES NetWare and NetWare 6.5	NSS for OES Linux
Cache balancing for NSS cache buffers	<p>Cache is limited to about 85% (by default) of the memory that is not allocated elsewhere. NSS on NetWare uses HighMem and recognizes up to 64 GB of RAM.</p> <p>You can set several tuning parameters for NSS on NetWare. For information, see <a href="#">“Configuring the System Cache to Fine-Tune NSS Performance (NetWare)” on page 327</a>.</p>	<p>For OES SP1 and later, for file data, NSS uses the Linux cache page manager to gain access to all available memory in the system. There are some limits in place so that when copying large files, NSS does not starve other user applications for memory, which is also how it works in NetWare. However, for metadata, NSS is still limited to what is assigned from the kernel below 1 GB cache, typically 60% of the kernel memory limit. (This is the same method used in other Linux file systems such as ReiserFS, Polyserve, XFS, with the exception of EXT.)</p> <p>For OES, cache for NSS is limited to about 60% of the cache assigned to the Linux kernel memory, which has its own maximum of 1 GB. NSS does not use Linux HighMem. It does not provide autobalancing.</p> <p>You can specify a minimum cache buffer size. For information, see <a href="#">“Configuring the System Cache to Fine-Tune NSS Performance (Linux)” on page 339</a>.</p>
Ability to access DOS partitions on the NetWare server	Load dosfat.nss to treat the partition as a standard NSS volume	No; use Native Linux options instead





# Comparison of NSS for Linux and Linux Traditional File Systems



This section compares features and capabilities of Novell® Storage Services™ on Linux to Linux traditional file systems, such as EXT3 and ReiserFS.

Feature Description	NSS on OES Linux with NCP Server	Linux Traditional File Systems with NCP Server
Management interfaces	Novell iManager Storage plug-in NSSMU for Linux EVMS (evmsgui) Utilities in the server console (nssmu, rights, nsscon, ravsui, ravview) Commands in the NSS Console (nsscon) Novell Remote Manager for Linux (browse only)	Novell iManager NCP™ Server plug-in YaST EVMS (evmsgui) Terminal commands Commands in the NCP Console (ncpcon)
File system trustees and trustee rights to control access to directories and files	Novell NetStorage Novell Client™. See <a href="#">“Using the Novell Client to Manage Trustees and Trustee Rights”</a> in the <i>File Systems Management Guide for OES</i> . Rights utility for Linux. See <a href="#">“Trustee Rights Utility for Linux”</a> in the <i>File Systems Management Guide for OES</i> .	NCP Server
File system directory and file attributes to control functions available for directories and files	Novell NetStorage Novell Client (NCP) Novell Remote Manager for Linux. See <a href="#">“Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions”</a> in the <i>File Systems Management Guide for OES</i> .	NCP Server
Directory quotas	Requires the Directory Quotas attribute. Novell NetStorage Novell Client	No

Feature Description	NSS on OES Linux with NCP Server	Linux Traditional File Systems with NCP Server
User space quotas (user space restrictions)	Not supported	Yes; differs from the NSS solution and is managed by Linux
Default mount location for NSS pools	/opt/novell/nss/mnt/.pools/	No
Default mount location for NSS volumes	/media/nss/	No
Default mount location for devices managed by EVMS	/dev/evms/	/dev/evms/
Interface	64-bit	64-bit
Character format	Unicode	UTF-8
Maximum device size recognized (physical or logical)	2 TB	2 TB
Maximum partition size	2 TB	2 TB
Maximum number of partitions (logical or physical devices) per pool	No practical limit  A partition must have more than 10 MB and less than 2 TB of available space to be recognized by NSS.	Not applicable
Maximum pool size (using at least 4 partitions of up to 2 TB each)	8 TB	16 TB
Minimum pool size	10 MB	4 KB
Maximum size of a volume	Up to 8 TB, depending on the pool size and available space in the pool.  Volume quotas can be overbooked. For information, see <a href="#">Section 13.1, "Understanding NSS Volumes," on page 181.</a>	Up to 16 TB, depending on the file system.
Maximum file size	Up to 8 TB, depending on the volume size and available space in the volume.	Up to 16 TB, depending on the volume size and available space in the volume.
Maximum number of files per volume	Up to 8 trillion, regardless of how many name spaces are loaded.	Up to 8 trillion, regardless of how many name spaces are loaded.
Maximum number of files open concurrently	1 million	Millions, depending on the file system
Maximum number of volumes per server	Unlimited NSS data volumes	Unlimited

Feature Description	NSS on OES Linux with NCP Server	Linux Traditional File Systems with NCP Server
Time to mount a volume	Requires only a few seconds  NSS uses a journaling file system and does not need to scan the entire file system to create a directory entry table (DET) and to load a File Allocation Table (FAT).	Depends on the file system; from a few seconds to a few minutes. Journalled file systems take only a few seconds.
Volume name space	UNIX is the default name space, which is case sensitive. However, you can specify other name spaces on mounting the NSS volume to make its directory names and filenames case insensitive. For example:  <code>mount ns=namespace</code>  Valid name space values are dos, mac, long, or unix.	UNIX; no support for case insensitive names.
File compression	Yes	No
Volume encryption	No	Yes; ReiserFS
Time to repair corrupted volume	Up to several hours, depending on the volume size.	Up to several hours, depending on the volume size
Multiple connection paths to storage media	No	No
Software RAID support	RAID 0, 1, and 10	RAID 0, 1, 3, and 5
Distributed File Services	No	No
Data shredding	Yes, up to 7 times	
Transaction Tracking System™ (TTS™)	No	No
File snapshot attribute (make a temporary snapshot copy of an open file for backup)	Yes; allows backup to continue without deactivating volume and back up open file	No
Pool snapshot (retain point-in-time version of a pool using block-level copy on write)	No	Depends on the file system
Backup support	Yes	No
Backup systems support		
Device maintenance support	Activate and deactivate devices by pool.	
CD and DVD device recognition	No; use Linux options instead	Yes

Feature Description	NSS on OES Linux with NCP Server	Linux Traditional File Systems with NCP Server
Operating system version detection	Default process	Default process
Cache balancing for NSS cache buffers	<p>Cache for NSS is limited to about 60% of the cache assigned to the Linux kernel memory, which has its own maximum of 1 GB. NSS does not use Linux HighMem.</p> <p>You can specify a minimum cache buffer size. For information, see <a href="#">“Configuring the System Cache to Fine-Tune NSS Performance (Linux)” on page 339</a>.</p>	Integrated with the Linux file system cache.
Ability to access DOS partitions on the NetWare server	No; use Linux options instead	Yes
Data migration	Yes	Yes
Novell Archive and Version Services	No	No

# Comparison of NSS on NetWare and the NetWare Traditional File System

# F

This section compares features and capabilities of the Novell® Storage Services™ (NSS) File System on NetWare® and the NetWare Traditional File System.

- Section F.1, “Features of NSS and Traditional File Systems on NetWare,” on page 429
- Section F.2, “SET Parameters Shared by NSS and Traditional File Systems,” on page 433
- Section F.3, “SET Parameters for the NSS File System,” on page 434
- Section F.4, “SET Parameters for the Traditional File System,” on page 435
- Section F.5, “Comparison of Directory and File Attributes,” on page 436

## F.1 Features of NSS and Traditional File Systems on NetWare

The following table describes key features of the NSS File System and the NetWare Traditional File System for OES NetWare and NetWare 6.5:

Feature Description	NSS File System	NetWare Traditional File System
Management interfaces	Novell iManager NSSMU for NetWare Novell Remote Manager for NetWare (limited feature set)	Novell Remote Manager for NetWare
Device and partition management	Novell iManager NSSMU for NetWare	Novell Remote Manager for NetWare
File system management for directories and files	Novell Client™ <i>OES NetStorage Administration Guide for NetWare</i> ( <a href="http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html">http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html</a> ) (via Web browser only, not WebDAV) Novell Remote Manager for NetWare Rights utility for NetWare	Novell Client <i>OES NetStorage Administration Guide for NetWare</i> ( <a href="http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html">http://www.novell.com/documentation/oes/netstor/data/h9izvdye.html</a> )
Interface	64-bit	32-bit

Feature Description	NSS File System	NetWare Traditional File System
File protocols	NCP™ NFAP (AFP, NFS, CIFS)	NCP
Character format	Unicode	ASCII double-byte
Maximum device size recognized (physical or logical)	2 TB	2 TB
Minimum device size recognized (physical or logical)	10 MB	10 MB
Maximum partition size	2 TB	1 TB
Minimum partition size	10 MB	10 MB
Maximum number of partitions (logical or physical devices) per NSS pool or Traditional volume	No practical limit, but the combined size cannot exceed the maximum NSS pool size of 8 TB.  A single partition must have more than 10 MB and up to 2 TB of available space to be recognized by NSS.	Up to 32 partitions, but the combined size cannot exceed the maximum Traditional volume size of 1 TB.
Maximum NSS pool or Traditional volume size	8 TB  An NSS pool of 8 TB requires at least 4 segments of up to 2 TB each. You can combine any number of partitions of 10 MB to 2 TB in size, but the combined size cannot exceed 8 TB.	1 TB  A Traditional volume of 1 TB requires 1 to 32 partitions of 10 MB to 1 TB in size, but the combined size cannot exceed 1 TB.
Minimum NSS pool or Traditional volume size	10 MB	10 MB
Maximum file size	Up to 8 TB per file, depending on the volume size and available space in the volume.	Up to 4 GB per file, depending on the volume size and available space in the volume.
Maximum number of files per volume	Up to 8 trillion, regardless of how many name spaces are loaded.	Up to 16 million files with 1 name space. The practical limit is up to 4 million files with 3 name spaces.
Maximum number of files open concurrently	1 million	100 thousand
Maximum number of data volumes mounted concurrently per server	255 plus the <code>sys:</code> volume.  You can mount NSS volumes beyond 256, but they are not visible or accessible through the normal Netware APIs.	255 plus the <code>sys:</code> volume.

Feature Description	NSS File System	NetWare Traditional File System
Maximum size of a volume	<p>Up to 8 TB, depending on the pool size and available space in the pool.</p> <p>If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool size, using the overbooking feature, although real total size is bound by physical limitations. Because space is allocated to volumes as needed, a volume might not reach its quota.</p>	<p>Up to 1 TB.</p> <p>The total size allocated for all volumes on the server cannot exceed the physical size of all devices combined, so overbooking is not possible. Unused space goes to waste if volumes grow unevenly.</p>
Time to mount a volume	<p>Only a few seconds.</p> <p>NSS uses a journaling file system and does not need to scan the entire file system to create a directory entry table (DET) and to load a File Allocation Table (FAT).</p>	<p>Up to several minutes, depending on volume size.</p> <p>The mount process scans the entire file system to create a DET before loading the volume. The time to mount increases based on the number of files, the number of name spaces used, and the size of the files.</p> <p>The mount process also loads the volume's FAT. Volumes with a smaller block size require more server memory to mount and manage, and it takes longer to mount the volume.</p>
Volume name space	<p>Accommodates all name spaces and requires less memory to mount than Traditional volumes.</p>	<p>Each name space you add to a Traditional volume requires additional server memory.</p> <p>If you add name space support to a volume and do not have enough memory, that volume cannot be mounted.</p>
Minimum server memory required to activate a volume	<p>Requires only 1 MB of available RAM to activate a single volume of any size and any number of files. Loads a file's metadata into memory only as you access the file.</p>	<p>Loads the DET into memory at the time of the mount. The amount of memory required increases with volume size and number of files. For a maximum volume size and number of objects, the memory required can be up to 1 GB just to activate the volume.</p>
File access time	<p>Same for each file, regardless of its location on the volume.</p>	<p>Depends on the file's location on the volume.</p>

Feature Description	NSS File System	NetWare Traditional File System
File save time	Offers the Flush Files Immediately attribute for NSS volumes to write files to disk on save instead of waiting for the next disk write cycle. This helps prevent possible data loss between disk write cycles and helps ensure data integrity.	Caches file in memory until the next disk write cycle.
File compression	Yes	Yes
Volume encryption	Yes	No
Error correction and data recovery time on system failure	<p>Journaling file system logs changes.</p> <p>On system failure, replays the most recent transactions to confirm validity, then repairs errors or rolls back to the original condition, typically in 15 to 60 seconds, unless the volume is corrupted.</p> <p>If the volume is corrupted, repair can take up to several hours, depending on the volume size.</p>	Must verify the entire volume contents with vrepair.nlm, which can take several hours, depending on the volume size.
Software RAID support	<p>RAID 0 (striping)</p> <p>RAID 1 (mirroring)</p> <p>RAID 5 (striping with parity)</p> <p>RAID 10 (mirroring RAID 0 devices)</p> <p>RAID 15 (mirroring RAID 5 devices)</p>	<p>RAID 0 (striping)</p> <p>RAID 1 (mirroring)</p> <p>RAID 10 (mirroring RAID 0 devices)</p>
Data shredding	Yes, up to 7 times	No
Transaction Tracking System™ (TTS™)	<p>Yes</p> <p>Requires exclusive use per server. To use TTS, first disable TTS for all Traditional volumes, and then enable TTS for your NSS volumes.</p>	<p>Yes</p> <p>Requires exclusive use per server. Use for Traditional or for NSS, but not for both on the same server. TTS is enabled by default.</p>
Directory quotas and User space quotas	Yes	Yes, enabled by default
File snapshot attribute (make a temporary snapshot copy of an open file for backup)	Yes; allows backup of open files without deactivating the volume.	No



Feature Description	NSS File System	NetWare Traditional File System
Pool snapshot (retain point-in-time version of a pool using block-level copy on write)	Yes; allows backup of block-level changes, without deactivating the volume.  Uses a brief freeze-release process to capture information for last remaining open files.	No
Backup support	Reviews a journaled list of files modified since the previous backup.	Does not have a modified file list. It must scan the entire file system to identify files modified since the previous backup.
Backup systems support	TSA, TSA600 for files over 2 GB, SMS, and Enhanced SBackup	TSA, SMS, and Enhanced SBackup
Device maintenance support	Activate and deactivate devices by pool.	Activate and deactivate devices by Traditional volume.
CD and DVD device recognition	Default process with full support for ISO 9660 and Macintosh HFS formats.  Use CDs and DVDs as read-only NSS volumes.	CDs and DVDs cannot be Traditional volumes in NetWare 6.5.
Operating system version detection	Default process	Manual process
DOS FAT compatible	Yes	No
Ability to access DOS partitions on the NetWare server	Load dosfat.nss to treat as a standard NSS volume.	No
Block suballocation	Unnecessary; superseded by technology.	Yes
Auditing	Unnecessary; superseded by technology.	Yes
File name locks	Unnecessary; superseded by technology.	Yes
Data migration	Yes	Yes

## F.2 SET Parameters Shared by NSS and Traditional File Systems

The SET parameters for the Common File System are shared by NSS and Traditional file systems. The following are the default settings in NetWare 6.5 for the Common File System SET parameters.

---

**IMPORTANT:** When modifying Common File System SET parameters, ensure that your planned settings satisfy the requirements for both your NSS and Traditional volumes on NetWare.

---

<b>Common File System SET Parameters</b>	<b>Default Value</b>
Maximum Transactions	10000
Maximum Concurrent Directory Cache Writes	75
Minimum File Delete Wait Time	1 minute 5.9 seconds
Immediate Purge Of Deleted Files	Off
Compression Daily Check Stop Hour	6
Compression Daily Check Starting Hour	0
Minimum Compression Percentage Gain	20
Enable File Compression	On
Maximum Concurrent Compressions	2
Convert Compressed To Uncompressed Option	1
Decompress Percent Disk Space Free To Allow Commit	10
Decompress Free Space Warning Interval	31 minutes 18.5 seconds
Deleted Files Compression Option	1
Days Untouched Before Compression	14

## F.3 SET Parameters for the NSS File System

The following are the default settings in NetWare 6.5 for the NSS file system SET parameters.

<b>Common File System SET Parameters for NSS File System</b>	<b>Default Value</b>
NSS Auto Locks HFS CD-ROM Disc in device	Off
NSS Auto Update CDHFS Volume Objects To NDS	Off
Mount ISO 9660 multi disc set as one CD9660 volume	Off
NSS Auto Locks CD-ROM Disc in device	Off
NSS Auto Update CD9660 Volume Objects To NDS	Off
Emulate Classic NetWare Directory Quotas	Off
NSS Work To Do Count	50
NSS Maximum Cache Balance Buffers Per Session	1024
NSS Cache Balance Timer	30
NSS Cache Balance Percent	60
NSS Cache Balance Enable	On
NSS Buffer Flush Timer	1
NSS Minimum OS Cache Buffers	256

<b>Common File System SET Parameters for NSS File System</b>	<b>Default Value</b>
NSS Minimum Cache Buffers	256
NSS File Flush Timer	10
NSS Closed File Cache Size	50000
NSS Open File Hash Shift	16
NSS Auth Cache Size	1024
NSS ASCII Name Cache Enable	On
NSS Name Cache Enable	On
NSS Name Cache Size	2111
NSS Low Volume Space Alerts	On
NSS Low Volume Space Warning Reset Threshold	15
NSS Low Volume Space Warning Threshold	10

## F.4 SET Parameters for the Traditional File System

The following are the default settings in NetWare 6.5 for the Traditional file system SET parameters.

<b>SET Parameters for the Traditional File System</b>	<b>Default Value</b>
Volume Log File State	1
Volume TTS Log File State	1
Volume Log File Overflow Size	4194304
Volume TTS Log File Overflow Size	4194304
Auto TTS Backout Flag	On
TTS Abort Dump Flag	Off
TTS Unwritten Cache Wait Time	1 minute 5.9 seconds
TTS Backout File Truncation Wait Time	59 minutes 19.2 sec
Dirty Directory Cache Delay Time	0.5 seconds
Directory Cache Allocation Wait Time	2.2 seconds
Directory Cache Buffer Nonreferenced Delay	5.5 seconds
Maximum Directory Cache Buffers	2000
Minimum Directory Cache Buffers	500
Maximum Number Of Internal Directory Handles	100
Maximum Number Of Directory Handles	20

<b>SET Parameters for the Traditional File System</b>	<b>Default Value</b>
Maximum Record Locks Per Connection	500
Maximum File Locks Per Connection	2500
Maximum Record Locks	20000
Maximum File Locks	200000
Read Ahead Enabled	On
Read Ahead LRU Sitting Time Threshold	10 seconds
Minimum File Cache Buffers	20
Maximum Concurrent Disk Cache Writes	750
Dirty Disk Cache Delay Time	3.3 seconds
Minimum File Cache Report Threshold	20
Automatically Repair Bad Volumes	On
File Delete Wait Time	5 minutes 29.6 seconds
Allow Deletion Of Active Directories	On
Maximum Percent of Volume Space allowed for Extended Attributes	10
Maximum Extended Attributes per File or Path	16
Purge Files On Dismount	Off
Fast Volume Mounts	On
Maximum Percent Of Volume Used By Directory	13
Maximum Subdirectory Tree Depth	25
Volume Low Warn All Users	On
Volume Low Warning Reset Threshold	256
Volume Low Warning Threshold	256
Turbo FAT Re-Use Wait Time	5 minutes 29.6 seconds
Allow Unowned Files To Be Extended	On
Auto Mount Mirrored Volume Containing Inactive Device	Off

## F.5 Comparison of Directory and File Attributes

<b>Attribute Code</b>	<b>Description</b>	<b>Applies to Files</b>	<b>Applies to Directories</b>	<b>NSS</b>	<b>Traditional</b>
A	Archive Needed identifies files that have been modified since the last backup. This attribute is assigned automatically.	Yes	No	Yes	Yes

Attribute Code	Description	Applies to Files	Applies to Directories	NSS	Traditional
Ci	Copy Inhibit prevents Macintosh users from copying a file. This attribute overrides Read and File Scan trustee rights.	Yes	No	Yes	Yes
Dc	Do Not Compress keeps data from being compressed. This attribute overrides settings for automatic compression of files not accessed within a specified number of days.	Yes	No	Yes	Yes
Di	Delete Inhibit means that the directory or file cannot be deleted. This attribute overrides the Erase trustee right.	Yes	Yes	Yes	Yes
Dm	Do Not Migrate prevents directories and files from being migrated from the server's server disk to another storage medium.	Yes	Yes	Yes	No migration support
Ds	Do Not Suballocate prevents data from being suballocated.	Yes	No	Not needed	Yes
H	The Hidden attribute hides directories and files so they cannot be listed using the DIR command.	Yes	Yes	Yes	Yes
I	Index allows large files to be accessed quickly by indexing files with more than 64 File Allocation Table (FAT) entries. This attribute is set automatically.	Yes	No	Not needed	yes
Ic	Immediate Compression sets data to be compressed as soon as a file is closed. If applied to a directory, every file in the directory is compressed as each file is closed.	Yes	Yes	Yes	Yes
N	Normal indicates the Read/Write attribute is assigned and the Shareable attribute is not. This is the default attribute assignment for all new files.	Yes	Yes	Yes	Yes
P	Purge flags a directory or file to be erased from the system as soon as it is deleted. Purged directories and files cannot be recovered.	Yes	Yes	Yes	Yes
Ri	Rename Inhibit prevents the directory or file name from being modified.	Yes	Yes	Yes	Yes
Ro	Read Only prevents a file from being modified. This attribute automatically sets Delete Inhibit and Rename Inhibit.	Yes	No	Yes	Yes
Rw	Read/Write allows you to write to a file. All files are created with this attribute.	Yes	No	Yes	Yes
Sh	Shareable allows more than one user to access the file at the same time. This attribute is usually used with Read Only.	Yes	No	Yes	Yes

Attribute Code	Description	Applies to Files	Applies to Directories	NSS	Traditional
Sy	System hides the directory or file so it does not appear in a file manager or directory listing. System is normally used with operating system files, such as Linux or NetWare system files.	Yes	Yes	Yes	Yes
T	Transactional allows a file to be tracked and protected by the Transaction Tracking System (TTS).	Yes	No	Explicit	Implicit and explicit
X	The Execute attribute indicates program files such as .exe or .com.	Yes	No	Yes	Yes

# NSS Nomenclature



This section discusses the nomenclature used for key Novell® Storage Services™ media objects in OES. This information can help you better understand the nature of error messages you might receive when using NSS. The table identifies the media object, defines it, and indicates the version of NetWare where the media object first appeared. All ZLSS (NSS Journaled Storage System) file blocks are 4 KB in size.

Media Object	Definition	NetWare Version Where First Used
Beast B-Tree	The Balanced Tree (B-Tree) that tracks all the file's metadata. This includes when the file was created, who created the file, the size of the file, and the location of the file's data.	5.0
Checkpoint	Four blocks (one in each <b>Superblock</b> ) that track where to start playing the journal if the server crashes. The checkpoint contains the metadata of the <b>Journal</b> .	5.0
Directory B-Tree	Used to implement an NSS volume's Directory Quota feature.	5.0 Support Packs
Entry ID (EID)	<p>An eDirectory™ entry within the FLAIM database (the underlying database for Directory Services 8.35 and later, including all eDirectory versions). EIDs are 32 bits and are unique for each server.</p> <p>EIDs are used within eDirectory to reference objects on a specific server. They are unique for each object for each server, meaning that object CN=Admin.O=EMA_CORP potentially has a different EID on server FS1 than on server FS2.</p> <p>This can be demonstrated with dsbrowse.nlm, a tool you can use to examine the local eDirectory database on the NetWare server. If you examine an object on one server and compare it with the same object on another server, the EID between the two instances of the object are probably dissimilar.</p>	6.0
Epoch File Log (EFL) B-Tree	Tracks files that change during an administrator-specified interval of time called an epoch. This feature is used by Novell Archive and Version Services.	6.5
Globally Unique ID (GUID)	<p>A globally unique identifier within eDirectory. The scope of this uniqueness is within one tree, although no actual checking is done to ensure this.</p> <p>GUIDs are 128 bits and are unique for each object. GUIDs allow an object to be referenced no matter which server you are accessing.</p>	5.0
Journal	The file used to quickly make the file system consistent after a server crash. The journal is sometimes referred to as a zlog.	5.0
Logged Pool Data Block	The block that tracks information about the pool, including the number of used blocks and salvageable blocks. Holds some of the items found in the zPoolInfo_s portion of the zInfo_s structure.	6.0

Media Object	Definition	NetWare Version Where First Used
Logged Volume Data Block	The block that tracks the number of files, used blocks, and compressed files of a volume. Holds some of the items found in the zVolumeInfo_s portion of the zInfo_s structure.	5.0
Modified File List (MFL) B-Tree	Tracks files and folders that have the Archive attribute set by the user. The Archive flag indicates that the file or folder needs to be backed up. NSS uses this list to quickly find files that need to be backed up during scheduled backups. This Archive file-and-folder attribute is unrelated to Novell Archive and Version Services.	6.0
Multiple Server Access Protection (MSAP)	The block used to reduce accidental use of a pool by more than one server at a time. A single copy of this block is stored in the second <b>Superblock</b> .	6.0 Support Packs
Name Tree B-Tree	The B-Tree that tracks the directory structure of a volume.	5.0
Pool Data Block	The block that tracks persistent pool configurable items. For example, a pool's features are stored here. Holds some of the items found in zPoolInfo_s portion of the zInfo_s structure.	6.0
Purge Log	Tracks transactions over an extended period so they can be completed after a crash. For example, the log records file deletes and truncates that need to be completed after a crash.	5.0
Purge Tree B-Tree	Used to store information about all salvageable files. This tree is used when the file system needs to auto purge files to create free blocks.	5.0
Snapshot	Used by the Media Manager to track pool snapshots. The object tracks which snapshots exist on a pool and where all blocks of a snapshot are stored. The root of this object resides in the first <b>Superblock</b> . All other blocks are allocated from a file on the internal volume.	6.5
Superblock	There are four Superblocks of 16 blocks each. The four Superblocks are replicas that reside in four fixed locations within the pool. The Superblock is used by the <b>Checkpoints</b> , <b>Superblock Header</b> , <b>Snapshot</b> , and <b>Multiple Server Access Prevention</b> .	5.0
Superblock Header	Four blocks (copies of each other, one in each Superblock) used to locate all other ZLSS media objects. These are the first blocks that the file system reads when a volume is loaded. Starting with NetWare 6.0, these blocks are read when a pool is loaded.	5.0
User B-Tree	Used to implement an NSS volume's User Quota feature. Starting with NetWare 6.0, the B-Tree also stores Novell eDirectory information related to User Quota.	5.0 Support Packs
Volume Data Block	The block that tracks configurable items in NSS volumes such as the volume's attributes and the high and low watermarks for salvage. It holds some of the items found in zVolumeInfo_s portion of the zInfo_s structure.	5.0
ZID	A numeric ID within the NSS file system, used to reference an object with the object store (also known as the <b>beast tree</b> ). ZIDs are 64 bits and are unique for each volume.	6.0



# Documentation Updates



This section contains information about documentation content changes made to the *Novell Storage Services Administration Guide for OES* since the initial release of Novell® Open Enterprise Server. If you are an existing user, review the change entries to readily identify modified content. If you are a new user, simply read the guide in its current state.

Refer to the publication date, which appears on the front cover and the Legal Notices page, to determine the release date of this guide. For the most recent version of the *Novell Storage Services Administration Guide for OES*, see the [Novell documentation Web site \(http://www.novell.com/documentation/oes/nss\\_enu/data/front.html\)](http://www.novell.com/documentation/oes/nss_enu/data/front.html).

In this section, content changes appear in reverse chronological order, according to the publication date. Within a dated entry, changes are grouped and sequenced, according to where they appear in the document itself. Each change entry provides a link to the related topic and a brief description of the change.

This document was updated on the following dates:

- [Section H.1, “April 12, 2006 \(Updates\),” on page 441](#)
- [Section H.2, “April 3, 2006 \(Updates\),” on page 442](#)
- [Section H.3, “December 23, 2005 \(OES SP2, NetWare 6.5 SP5\),” on page 444](#)
- [Section H.4, “November 1, 2005,” on page 446](#)
- [Section H.5, “September 29, 2005 \(Updates\),” on page 446](#)
- [Section H.6, “August 19, 2005 \(OES SP1 and NetWare 6.5 SP4\),” on page 449](#)
- [Section H.7, “June 1, 2005 \(Updates\),” on page 455](#)
- [Section H.8, “March 28, 2005 \(Updates\),” on page 456](#)

## H.1 April 12, 2006 (Updates)

Updates were made to the following section. The change is explained below.

- [Section H.1.1, “Management Tools for NSS,” on page 441](#)
- [Section H.1.2, “Coexistence and Migration Issues,” on page 442](#)
- [Section H.1.3, “Using EVMS to Manage Devices with NSS Volumes \(Linux\),” on page 442](#)

### H.1.1 Management Tools for NSS

The following change was made to this section.

Location	Change
<a href="#">“Prerequisites for Using the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager 2.5” on page 40.</a>	<b>IMPORTANT:</b> The Storage, Archive Versioning, File Protocols (NetWare), and Cluster Services plug-ins share common code. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins. If you remove only one of the plug-ins, it removes the common code and breaks the remaining installed plug-ins.

## H.1.2 Coexistence and Migration Issues

The following change was made to this section.

Location	Change
<a href="#">“Compatibility Issues for the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager” on page 79.</a>	<b>IMPORTANT:</b> The Storage, Archive Versioning, File Protocols (NetWare), and Cluster Services plug-ins share common code. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins. If you remove only one of the plug-ins, it removes the common code and breaks the remaining installed plug-ins.

## H.1.3 Using EVMS to Manage Devices with NSS Volumes (Linux)

The following change was made to this section.

Location	Change
<a href="#">“Mounting Your Root File System with EVMS” on page 410.</a>	This section is new.

## H.2 April 3, 2006 (Updates)

Updates were made to the following sections. The changes are explained below.

- [Section H.2.1, “Overview of NSS,” on page 442](#)
- [Section H.2.2, “Coexistence and Migration Issues,” on page 443](#)
- [Section H.2.3, “Conserving Disk Space with Volume, Directory, or User Space Quotas,” on page 443](#)
- [Section H.2.4, “Troubleshooting,” on page 443](#)
- [Section H.2.5, “Using EVMS to Manage Devices with NSS Volumes \(Linux\),” on page 444](#)

### H.2.1 Overview of NSS

The following change was made to this section.

Location	Change
<a href="#">“Benefits of NSS” on page 23</a>	<p>This statement was corrected from “trillions” to “4 billion (4E9)”:</p> <ul style="list-style-type: none"> <li>• Up to 4 billion (4E9) files in volume</li> </ul> <p>In NSS, the maximum number of file IDs available is a 64-bit number, which provides for up to 8 trillion (8E12) ZIDs, so NSS was designed to not re-use ZIDs. However, NCP clients and other traditional applications can only work with 32-bit IDs, which supports up to 4 billion (4E9) ZIDs. NSS restricts ZIDs, and thus the number of files, to the lower value.</p>

## H.2.2 Coexistence and Migration Issues

The following changes were made to this section.

Location	Change
<a href="#">“Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later” on page 92</a>	<p>Step 9 was added:</p> <p>If you want to mirror the upgraded partition from <a href="#">Step 8</a>, see <a href="#">Section 9.16, “Mirroring an Existing Pool with NSSMU,” on page 136</a>.</p>
<a href="#">“Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare” on page 95</a>	<p>Step 10 was added:</p> <p>If you want to mirror the upgraded partition from <a href="#">Step 9</a>, see <a href="#">Section 9.16, “Mirroring an Existing Pool with NSSMU,” on page 136</a>.</p>

## H.2.3 Conserving Disk Space with Volume, Directory, or User Space Quotas

The following change was made to this section.

Location	Change
<a href="#">“Deleting All Directory Quotas for an NSS Volume” on page 250</a>	This section is new.

## H.2.4 Troubleshooting

The following change was made to this section.

Location	Change
<a href="#">"Problems Following DFS Junctions in Windows 2000/XP Releases (NetWare)" on page 346</a>	When using DFS with NetWare CIFS, the CIFS server and Windows clients are on different IP subnets. In this case, it is required that the client have a way to resolve the CIFS server name in order for DFS to work. This is a Microsoft/CIFS requirement not a NetWare CIFS requirement. There are multiple ways the client can do this, including 1) both client and NW server are configured for the same WINS server, 2) both client/server using same DNS server, or 3) the lmhosts workaround described in this topic.

## H.2.5 Using EVMS to Manage Devices with NSS Volumes (Linux)

The following change was made to this section.

Location	Change
<a href="#">"Making Devices Available to EVMS" on page 410</a>	This section was modified for clarity.

## H.3 December 23, 2005 (OES SP2, NetWare 6.5 SP5)

Updates were made to the following sections. The changes are explained below.

- [Section H.3.1, "Management Tools for NSS," on page 444](#)
- [Section H.3.2, "Coexistence and Migration Issues," on page 444](#)
- [Section H.3.3, "Managing Compression on NSS Volumes," on page 445](#)
- [Section H.3.4, "Using Distributed File Services to Move and Split Volumes \(NetWare\)," on page 445](#)
- [Section H.3.5, "Using Non-EVMS Volume Managers for NSS on Linux," on page 445](#)
- [Section H.3.6, "NSS Utilities," on page 446](#)

### H.3.1 Management Tools for NSS

The following change was made to this section.

Location	Change
<a href="#">"EVMS" on page 40</a>	For OES Linux, the plug-in requires that EVMS be the volume manager of devices that contribute space to NSS pools and volumes.

### H.3.2 Coexistence and Migration Issues

The following change was made to this section.

Location	Change
Section 6.2, “NSS Kernel Modules in SLES 9,” on page 74	This section is new.

### H.3.3 Managing Compression on NSS Volumes

The following changes were made to this section.

Location	Change
“Using ATTRIB (Linux)” on page 238	This section is new.
“Copying Compressed Files between Volumes with the NCOPY Command (NetWare)” on page 242	This section’s explanation of the <code>NCOPY</code> utility was moved to the NSS Utilities section. For information, see <a href="#">Section A.11, “NCOPY (NetWare),” on page 377</a> .
“Viewing the File Attributes with the NDIR Command (NetWare)” on page 242	This section’s explanation of the <code>NDIR</code> utility was moved to the NSS Utilities section. For information, see <a href="#">Section A.12, “NDIR (NetWare),” on page 378</a> .
“Repairing Compressed Volumes with the Compfix Utility” on page 242	This section’s explanation of the <code>COMPFIX</code> utility was moved to the NSS Utilities section. For information, see <a href="#">Section A.4, “COMPFIX,” on page 362</a> . There is a <code>COMPFIX</code> utility for Linux and for NetWare.

### H.3.4 Using Distributed File Services to Move and Split Volumes (NetWare)

The following changes were made to this section.

Location	Change
“DFS Junctions” on page 262	The servers of the source volume and destination volume must reside in a management context.
“Prerequisites for Creating and Using DFS Junctions” on page 263	<p>You can create DFS junctions only in NSS volumes that reside on your NetWare 6.0 or later server that is part of a DFS management context.</p> <p>All volumes that participate in the junctions, whether they contain junctions or are targets of junctions, must be located on servers in a management context.</p>

### H.3.5 Using Non-EVMS Volume Managers for NSS on Linux

This section is new. See [Chapter 22, “Using NSS on Devices Managed by Non-EVMS Volume Managers,” on page 291](#).

## H.3.6 NSS Utilities

This section is new. See [Appendix A, “NSS Utilities,” on page 359](#).

## H.4 November 1, 2005

The entire guide was reformatted to comply with revised Novell documentation standards. The content is unchanged.

## H.5 September 29, 2005 (Updates)

Updates were made to the following sections. The changes are explained below.

- [Section H.5.1, “About This Guide,” on page 446](#)
- [Section H.5.2, “Management Tools for NSS,” on page 446](#)
- [Section H.5.3, “Coexistence and Migration Issues,” on page 447](#)
- [Section H.5.4, “Managing Encrypted NSS Volumes,” on page 447](#)
- [Section H.5.5, “Conserving Disk Space with Volume, Directory, or User Space Quotas,” on page 447](#)
- [Section H.5.6, “Using NSS Server Console Commands and Utilities,” on page 448](#)
- [Section H.5.7, “Configuring the System Cache to Fine-Tune NSS Performance \(Linux\),” on page 448](#)
- [Section H.5.8, “Comparison of NSS on NetWare and NSS on Linux,” on page 448](#)

### H.5.1 About This Guide

The following change was made to this section.

Location	Change
Additional Documentation	<a href="http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html">NDK: Virtual File Services for NetWare (http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html)</a>

### H.5.2 Management Tools for NSS

The following change was made to this section.

Location	Change
<a href="#">Section 4.8, “Virtual File Services, APIs, and Scripts,” on page 63</a>	<p>For NetWare, other APIs are available in the <a href="#">NDK: Novell Storage Architecture Component (Media Manager and NWPA) (http://developer.novell.com/ndk/storarch.htm)</a>.</p> <p>For OES Linux, the plug-in requires that EVMS be the volume manager of devices that contribute space to NSS pools and volumes.</p>

## H.5.3 Coexistence and Migration Issues

The following changes were made to this section.

Location	Change
Section 6.4, "Access Control Issues for NSS on OES Linux," on page 74	This section was updated for clarity. The following information was added:  If you Linux-enable a user who has been logged into the system before being Linux-enabled, make sure execute the <code>resetidcache</code> command from the NSS Console ( <code>nsscon</code> ) utility afterwards. This allows the proper reporting of ownership because it resets the mapping of user identities in the ID cache and forces it to update with the Linux UID for the user.
Section 6.6, "Compatibility Issues for Using NSS Cross-Platform," on page 82	The following issues were added: <ul style="list-style-type: none"><li>• For OES SP1, you cannot create a mixed cluster if the shared pool resides on software RAID 0 or RAID 5 devices. This configuration is not supported by NSS on Linux.</li><li>• Do not use the DFS move/split functions for a shared NSS volume in a mixed cluster because DFS is not supported for NSS on Linux.</li><li>• When using the move/split functions in a NetWare cluster, perform the move/split only from an active node.</li></ul>

## H.5.4 Managing Encrypted NSS Volumes

The following changes were made to this section.

Location	Change
Section 14.6, "Mounting Encrypted NSS Volumes with NSS Commands (Linux)," on page 206	Removed references to the <code>:volpassword</code> syntax for command lines for mounting encrypted volumes. This capability was removed before shipping as a security risk.
Section 14.7, "Activating Encrypted NSS Volumes with NSS Commands (NetWare)," on page 206	Removed references to the <code>:volpassword</code> syntax for command lines for mounting encrypted volumes. This capability was removed before shipping as a security risk.

## H.5.5 Conserving Disk Space with Volume, Directory, or User Space Quotas

The following changes were made to this section.

Location	Change
Section 18.1, "Understanding Space Quotas," on page 243	<p>The following information was added:</p> <p>Before setting quotas, make sure to Linux-enable the users with Linux User Management so that file ownership can be tracked and space consumption can be enforced.</p> <hr/> <p><b>WARNING:</b> As the amount of space consumed by a user's files approaches the user's space quota, users should use caution when saving files. Data loss can occur if the user attempts to save a file that is too large for the remaining unused space. If storing the file would cause the quota to be exceeded, only part of the file is actually saved, resulting in data corruption.</p> <hr/> <p>This section is new.</p>
Section 18.4.6, "Setting Quotas for Users Who Were Not Initially Enabled for Linux Access (Linux)," on page 254	

## H.5.6 Using NSS Server Console Commands and Utilities

The following change was made to this section.

Location	Change
Section 23.8, "Encrypted Volume Support Commands," on page 305	Removed references to the <code>:volpassword</code> syntax for command lines for mounting encrypted volumes. This capability was removed before shipping as a security risk.

## H.5.7 Configuring the System Cache to Fine-Tune NSS Performance (Linux)

The following change was made to this section.

Location	Change
Section 26.1, "Understanding How NSS Uses Cache on Linux," on page 339	This section was updated for clarity.

## H.5.8 Comparison of NSS on NetWare and NSS on Linux

This section was updated to reflect the availability of features as of OES SP1.



## H.6 August 19, 2005 (OES SP1 and NetWare 6.5 SP4)

Updates were made to the following sections. The changes are explained below.

- Section H.6.1, “Overview of NSS,” on page 449
- Section H.6.2, “What’s New,” on page 450
- Section H.6.3, “Upgrading the Media Format for OES SP1 NetWare and NetWare 6.5 SP4,” on page 450
- Section H.6.4, “Management Tools for NSS,” on page 450
- Section H.6.5, “Planning for NSS Storage Solutions,” on page 450
- Section H.6.6, “Coexistence and Migration Issues,” on page 451
- Section H.6.7, “Managing Partitions,” on page 452
- Section H.6.8, “Managing Software RAID Devices,” on page 452
- Section H.6.9, “Managing NSS Pools,” on page 452
- Section H.6.10, “Managing NSS Volumes,” on page 452
- Section H.6.11, “Managing Encrypted NSS Volumes,” on page 452
- Section H.6.12, “Using Volume Attributes that Help Ensure Data Integrity,” on page 452
- Section H.6.13, “Conserving Disk Space with Volume, Directory, or User Space Quotas,” on page 453
- Section H.6.14, “Salvaging and Purging Deleted Volumes, Directories, and Files,” on page 453
- Section H.6.15, “Using Distributed File Services to Move and Split Volumes (NetWare),” on page 453
- Section H.6.16, “Managing Hard Links,” on page 454
- Section H.6.17, “Configuring the System Cache to Fine-Tune NSS Performance (NetWare),” on page 454
- Section H.6.18, “Configuring the System Cache to Fine-Tune NSS Performance (Linux),” on page 455
- Section H.6.19, “Troubleshooting the NSS File System,” on page 455
- Section H.6.20, “Security Considerations,” on page 455

### H.6.1 Overview of NSS

The following changes were made to this section.

Location	Change
Improve Storage Availability	Software RAIDs 5 and 15 are now also supported on OES SP1 Linux and later.  You can now specify the default name space to use when mounting a volume.
Prevent Unauthorized Access	Encrypted volumes are now also supported on OES SP1 Linux and later.

Location	Change
Maximize Available Space	User space restrictions are now also supported on OES SP1 Linux and later.

## H.6.2 What's New

The following changes were made to this section.

Location	Change
What's New (OES SP1 and NetWare 6.5 SP4)	This section is new.
What's New (OES and NetWare 6.5 SP3)	This section was renamed.

## H.6.3 Upgrading the Media Format for OES SP1 NetWare and NetWare 6.5 SP4

This section is new. After you upgrade to NetWare 6.5 SP4 or OES SP1 NetWare, you must convert volumes that contain any existing hard links to the new hard link style before you use the volume; some restrictions apply. For information, see [Chapter 3, “Upgrading the Media Format \(NetWare\),” on page 33](#).

## H.6.4 Management Tools for NSS

The following changes were made to this section.

Location	Change
Novell iManager Storage Plug-In	<p>In OES SP1 and later, iManager 2.5 is required for the storage-related plug-ins.</p> <p>The Storage plug-in does not support Mobile iManager.</p> <p>This section was updated with instructions for installing updated plug-ins.</p> <p>This section was updated to document the <i>Lookup Name Space</i> option and features that are now available on Linux, including encrypted volumes, software RAID 5, and user space restrictions.</p>
NSS Management Utility	This section was updated to document the <i>Name Space</i> option (F9) and features that are now available on Linux, including encrypted volumes, software RAIDS 5 and 15, and user space restrictions.
Novell NetStorage	Beginning in OES SP1, NetStorage is accessible from within iManager. In iManager, select <i>NetStorage</i> in <i>Roles and Tasks</i> .

## H.6.5 Planning for NSS Storage Solutions

The following changes were made to this section.

Location	Change
Planning to Move NSS Pools and Volumes Cross-Platform	This section was updated to link to new sections in the Coexistence and Migration chapter.
Software RAID Devices	Software RAIDs 5 and 15 are now also supported on OES SP1 Linux and later.
NSS Encrypted Volumes	Encrypted volumes are now also supported on OES SP1 Linux and later.

## H.6.6 Coexistence and Migration Issues

The following changes were made to this section:

Location	Change
Compatibility with Operating Systems and Services	The operating systems for this release include OES SP1 NetWare, OES Linux SP1, and NetWare 6.5 SP4.
Access Control Issues for NSS on OES Linux	This section is new.
Compatibility Issues for the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager	In OES SP1 and later, iManager 2.5 is required for the storage-related plug-ins.
Web Browser Language Setting	This issue has been resolved for the Archive Versioning plug-in in the OES SP1 release.
Protocols	Table headings were modified to include the OSES for the OES SP1 release.
Guidelines for Cross-Platform Compatibility of NSS	This section is updated.
Guidelines for Moving Devices Cross-Platform	This section is new.
Moving Non-Clustered Devices From NetWare 6.5 or OES NetWare to OES Linux	This section was updated.
Moving Non-Clustered Devices From NetWare 6.0 to OES SP1 Linux or Later	This section is new.
Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare	This section is new.

## H.6.7 Managing Partitions

This section was updated for changes in the management interface in the Storage plug-in to iManager.

## H.6.8 Managing Software RAID Devices

This section was updated to document added support of RAID5 and 15 for OES SP1 Linux and later.

## H.6.9 Managing NSS Pools

The following changes were made to this section:

Location	Change
Linux	This section was updated for clarity.  Before you verify or rebuild, you must place the pool in maintenance mode. At a terminal prompt, start NSSCON by entering <code>nsscon</code> , then enter  <code>nss /poolmaintenance=poolname</code>
ReZIDing Volumes in an NSS Pool	This section is new.
Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool	This section is was moved to <a href="#">Section 19.2, "Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool," on page 255.</a>

## H.6.10 Managing NSS Volumes

The following changes were made to this section:

Location	Change
Encryption	Encrypted volumes are now also supported on OES SP1 Linux.
User Space Quotas	User space quotas are now also supported on OES SP1 Linux.
Lookup Name Space	The Lookup Name Space attribute is new for this release.

## H.6.11 Managing Encrypted NSS Volumes

Changes were made to this section to document additional procedures for managing encrypted NSS volumes on Linux.

## H.6.12 Using Volume Attributes that Help Ensure Data Integrity

The following change was made to this section:

Location	Change
Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare)	For more information, see <a href="http://www.novell.com/documentation/beta/oes/sos__enu/data/h0wbzsqh.html">Using the Transaction Tracking System</a> ( <a href="http://www.novell.com/documentation/beta/oes/sos__enu/data/h0wbzsqh.html">http://www.novell.com/documentation/beta/oes/sos__enu/data/h0wbzsqh.html</a> ) in the <i>Server Operating System for NetWare Administration Guide for OES</i> ( <a href="http://www.novell.com/documentation/beta/oes/sos__enu/data/h1hge52k.html">http://www.novell.com/documentation/beta/oes/sos__enu/data/h1hge52k.html</a> ).

## H.6.13 Conserving Disk Space with Volume, Directory, or User Space Quotas

The following change was made to this section:

Location	Change
Understanding Space Quotas	User space quotas are now also supported for OES SP1 Linux and later, if users are Linux-enabled with Linux User Management.
Managing User Space Quotas	The User Space Quotas function is fully supported for OES SP1 on Linux and NetWare.

## H.6.14 Salvaging and Purging Deleted Volumes, Directories, and Files

The following changes were made to this section:

Location	Change
Configuring the NSS Volume Salvage Area	The volume salvage and file system salvage function separately .
Setting the Low and High Watermarks for Purging the File System Salvage	The volume salvage and file system salvage function separately.
Salvaging or Purging Deleted Directories from the File System Salvage Area (OES SP1 NetWare, NetWare 6.5 SP4, or Later)	Novell Remote Manager for NetWare supports salvaging and purging deleted directories for NSS volumes on NetWare.

## H.6.15 Using Distributed File Services to Move and Split Volumes (NetWare)

The following change was made to this section:

Location	Change
<a href="#">Prerequisites for Pointing DFS Junctions at an NSS Volume on OES Linux</a>	<p>This section is new.</p> <p>DFS is not supported for NSS volumes on Linux, but DFS junctions can point to them if certain conditions exist.</p> <p>After you create an NSS volume on Linux, make sure to run VLDB Repair on your VLDB server to update the VLDB database to include the NSS volume on Linux. For information, see <a href="#">Section 20.8, “Managing the Volume Location Database Service (NetWare),”</a> on page 269.</p>

## H.6.16 Managing Hard Links

The following changes were made to this section:

Location	Change
<a href="#">Understanding Hard Links</a>	<p>In OES and NetWare 6.5 SP3 and earlier, all name information (names in each name space) combined for a file's hard link could not exceed 512 Unicode characters. However, updates made to the hard links function in this release remove that length limitation for OES SP1 NetWare and NetWare 6.5 SP4 and later, which can now support up to 65,535 hard links per volume. This requires a media format upgrade.</p> <p>This section has been updated to include important information about the behavior of hard links.</p>
<a href="#">Enabling or Disabling the Hard Links Attribute (OES SP1 NetWare, NetWare 6.5 SP4)</a>	This section is new.
<a href="#">Creating a Hard Link Using a zLink API</a>	This section is new.
<a href="#">Creating a Hard Link for Testing Purposes (OES SP1 NetWare, NetWare 6.5 SP4)</a>	This section is new.
<a href="#">Viewing Hard Links for a Volume (OES SP1 NetWare, NetWare 6.5 SP4)</a>	This section is new.

## H.6.17 Configuring the System Cache to Fine-Tune NSS Performance (NetWare)

The following change was made to this section:

Location	Change
<a href="#">Hints for Low Memory NetWare Systems</a>	This section is new.

## H.6.18 Configuring the System Cache to Fine-Tune NSS Performance (Linux)

The following change was made to this section:

Location	Change
<a href="#">Understanding How NSS Uses Cache on Linux</a>	Beginning with OES SP1 Linux, NSS uses all available assigned cache. The 1 GB cache buffer limit has been removed.

## H.6.19 Troubleshooting the NSS File System

The following change was made to this section:

Location	Change
Slow Performance with User Space Quotas (Linux)	This performance issue was resolved for OES SP1 Linux. The User Space Quotas function is fully supported for OES SP1 on Linux and NetWare.

## H.6.20 Security Considerations

This section is new. It is an overview of NSS security issues and is intended for security administrators.

## H.7 June 1, 2005 (Updates)

Updates were made to the following sections. The changes are explained below.

- [Section H.7.1, “Management Tools for NSS,” on page 455](#)
- [Section H.7.2, “Securing Access to NSS Volumes and Files,” on page 455](#)

### H.7.1 Management Tools for NSS

The following change was made to this section:

Location	Change
<a href="#">Prerequisites for Using the Storage, Archive Versioning, File Protocols, and Clustering Plug-Ins for iManager 2.5</a>	For information about installing WBEM, see “ <a href="http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html">Setting Up OpenWBEM</a> ” ( <a href="http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html">http://www.novell.com/documentation/oes/cimom/data/bv3wjre.html</a> ) in the <i>OpenWBEM Services Administration Guide for OES</i> ( <a href="http://www.novell.com/documentation/oes/cimom/data/front.html">http://www.novell.com/documentation/oes/cimom/data/front.html</a> ).

### H.7.2 Securing Access to NSS Volumes and Files

The following change was made to this section:

Location	Change
Controlling Access to the Admin Volume and Files.cmd (NetWare)	<p>The child elements to the Rights tag (&lt;rights&gt;) include the following tags for file system rights:</p> <ul style="list-style-type: none"> <li>• &lt;accesscontrol/&gt; for the Access Control right</li> <li>• &lt;create/&gt; for the Create right</li> <li>• &lt;edit/&gt; for the Edit right</li> <li>• &lt;filesCAN/&gt; for the File Scan right</li> <li>• &lt;modify/&gt; for the Modify Metadata right</li> <li>• &lt;read/&gt; for the Read right</li> <li>• &lt;supervisor/&gt; for the Supervisor right</li> <li>• &lt;write/&gt; for the Write right</li> </ul>

## H.8 March 28, 2005 (Updates)

Updates were made to the following sections. The changes are explained below.

- [Section H.8.1, “Management Tools for NSS,” on page 456](#)
- [Section H.8.2, “Planning for NSS Storage Solutions,” on page 456](#)
- [Section H.8.3, “Managing Software RAID Devices,” on page 457](#)
- [Section H.8.4, “Managing NSS Volumes,” on page 457](#)
- [Section H.8.5, “Comparison of NSS on NetWare and NSS on Linux,” on page 457](#)

### H.8.1 Management Tools for NSS

The following change was made to this section:

Location	Change
Downloading and Installing Plug-in Updates	This information is new.

### H.8.2 Planning for NSS Storage Solutions

The following change was made to this section:

Location	Change
Case Insensitive Naming Convention	<p><b>IMPORTANT:</b> Because Linux treats filenames as case sensitive, when using NSS volumes on Linux, make sure to mount the volume with the Long name space (ns=long) option so that file queries are case insensitive. For information, see <a href="#">Section 13.16, “Mounting an NSS Volume on Linux to Support Case Insensitive File Names (Linux),” on page 197.</a></p>



## H.8.3 Managing Software RAID Devices

The following change was made to this section:

Location	Change
Creating a Software RAID 15 with NSSMU	This information is new.

## H.8.4 Managing NSS Volumes

The following changes were made to this section:

Location	Change
Mounting an NSS Volume on Linux to Support Case Insensitive File Names (Linux)	<p>For the Mount command (<code>mount</code>) for NSS on Linux, key fields and options are now explained.</p> <p>A paragraph was added to explain how to place the mount command in the <code>/etc/fstab</code> file to automatically mount the volume as case insensitive whenever you restart the system.</p> <p>When using Samba, you should mount NSS volumes with the Long name space to improve performance.</p>
Verifying and Rebuilding NSS Pools and Volumes	Added information and links about using the <code>RAVSUI</code> utility to verify or rebuild an NSS Volume on Linux.

## H.8.5 Comparison of NSS on NetWare and NSS on Linux

The following changes were made to this section:

Location	Change
Maximum Size of a Voume	For information, see <a href="#">Section 18.2, "Managing NSS Volume Quotas," on page 245.</a>
File system type (as recognized and reported by the operating system)	For NetWare, the file system type is <code>nss</code> . For Linux, the file system type is <code>nssvol</code> .
Maximum number of volumes per server	<p>For NetWare: 255 plus the <code>sys:</code> volume.</p> <p>You can mount NSS volumes beyond 256, but they are not visable or accessible through the normal Netware APIs.</p>