AUTHORIZED DOCUMENTATION

Novell Storage Services™ File System Administration Guide

Novell®
Open Enterprise Server

2 SP1

December 2008

www.novell.com

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. See the Novell International Trade Services Web page (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 © 2005–2008 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 on the Novell Legal Patents Web page (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

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For a list of Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

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

About This Guide 25

1 Overview of NSS 29

1.1 Introduction to NSS . . . 29
1.2 Benefits of NSS . . . 29
1.3 Understanding NSS . . . 30
1.3.1 Storage Pools . . . 31
1.3.2 NSS Volumes . . . 31
1.4 NSS Features and Capabilities . . . 32
1.4.1 Use Less Memory and Gain More Speed . . . 32
1.4.2 Improve Storage Availability . . . 32
1.4.3 Prevent Unauthorized Access . . . 33
1.4.4 Protect Data from Corruption or Loss . . . 33
1.4.5 Maximize Available Space . . . 34
1.5 Comparison of NSS to Other File Systems . . . 34
1.6 What's Next . . . 35

2 What's New 37

2.1 What's New for NSS (OES 2 SP1 and NetWare 6.5 SP8) . . . 37
2.1.1 Linux and NetWare . . . 37
2.1.2 Linux . . . 37
2.1.3 NetWare . . . 39
2.2 What's New for NSS (OES 2 and NetWare 6.5 SP7) . . . 39
2.2.1 Linux and NetWare . . . 39
2.2.2 Linux . . . 41
2.2.3 NetWare . . . 43

3 Installing and Configuring Novell Storage Services 45

3.1 Requirements for Installing NSS . . . 45
3.1.1 Device Requirements for NetWare and Linux . . . 45
3.1.2 Requirements for NSS on NetWare . . . 46
3.1.3 Requirements for NSS on Linux . . . 46
3.1.4 Requirements for Storage-Related iManager Plug-Ins . . . 47
3.2 Installing and Configuring NSS on NetWare . . . 47
3.3 Installing and Configuring NSS on Linux . . . 48
3.3.1 Installing Linux with EVMS as the Volume Manager of the System Device . . . 48
3.3.2 Selecting the NSS Pattern During the OES 2 Linux Install . . . 48
3.3.3 Installing NSS on an Existing OES 2 Linux Server . . . 50
3.3.4 Enabling or Disabling NSS on Linux . . . 51
3.4 Upgrading the Media Format for Hard Link Support . . . 51
3.5 Enabling Users for Linux Utilities and Services . . . 51
3.6 Updating NSS on OES 2 Linux . . . 52
3.6.1 Parameter Settings . . . 52
3.6.2 Reboot Server or Restart jstcpd, adminusd, and volmnd . . . 52
3.6.3 Storage-Related Plug-Ins . . . 52
3.7 Upgrading from OES 1 Linux to OES 2 Linux . . . 53
3.7.1 Parameter Settings . . . 53

3.7.2 Read Ahead Blocks Setting . . . 53
3.8 Updating the NSS Kernel Module (km_nss) for SUSE Linux Enterprise Server . . . 54

4 Upgrading the NSS Media Format 55

4.1 Guidelines for Upgrading the Media Format of NSS Volumes . . . 55
4.1.1 Cross-Platform Support for the NSS Media Upgrade . . . 55
4.1.2 Which NSS Volumes to Upgrade . . . 55
4.1.3 Before Upgrading the Media Format . . . 56
4.1.4 After Upgrading the Media Format . . . 57
4.1.5 If You Do Not Upgrade the Media Format . . . 57
4.2 Upgrading the Media Format for Existing NSS Volumes . . . 57
4.2.1 Prerequisite . . . 58
4.2.2 Upgrading the Media Format . . . 58
4.2.3 Enabling Hard Links After the Media Upgrade . . . 59
4.2.4 Post Upgrade Task . . . 60
4.3 Upgrading the Media Format Automatically for New NSS Volumes . . . 60

5 Planning NSS Storage Solutions 63

5.1 Guidelines for NSS Storage . . . 63
5.1.1 Devices . . . 63
5.1.2 Software RAID Devices . . . 64
5.1.3 Device Partitions . . . 64
5.1.4 Multiple I/O Paths (NetWare) . . . 65
5.1.5 NSS Pools and Volumes . . . 65
5.1.6 NSS Encrypted Volumes . . . 66
5.1.7 Storage Features . . . 66
5.2 Compatibility and Interoperability Issues for NSS . . . 66
5.3 Creating NSS Storage Objects in eDirectory . . . 67
5.4 Naming NSS Storage Objects . . . 67
5.4.1 Case Insensitive Names . . . 67
5.4.2 Number of Characters Allowed . . . 68
5.4.3 Conventions for Valid Names of NSS Storage Objects . . . 69
5.4.4 Other Naming Guidelines . . . 69
5.4.5 UTF-8 Naming Considerations in Mixed-Language Environments (NetWare) . . . 70
5.5 Access Control for NSS on Linux . . . 71
5.5.1 Administrator User and Root User Roles for NSS . . . 71
5.5.2 NSS File System Users . . . 72
5.5.3 Novell Trustee Model . . . 73
5.5.4 POSIX Permissions . . . 74
5.5.5 How NSS Uses Novell Linux User Management . . . 74
5.6 File Access for Users . . . 74
5.6.1 NCP . . . 75
5.6.2 Novell AFP . . . 76
5.6.3 Novell CIFS . . . 76
5.6.4 Novell Domain Services for Windows . . . 76
5.6.5 Samba . . . 76
5.6.6 SSH (Secure Shell) . . . 77
5.6.7 Accessing Files with Linux Services, Utilities, and Protocols . . . 77
5.7 Antivirus Support for NSS . . . 77
5.8 Backup Support for NSS . . . 78
5.9 NSS Support for Memory Mapped Files . . . 78

6 Using NSS in a Virtualization Environment 79

6.1 Guidelines for Using NSS in a Xen Virtualization Environment . . . 79
6.1.1 Host Server Issues . . . 79
6.1.2 Virtual Machine Issues . . . 81
6.1.3 Guest Server Issues . . . 81
6.2 Installing OES 2 Linux on a Virtual Machine . . . 82
6.3 Initializing New Virtual Disks on the Guest Server . . . 83
6.4 What's Next . . . 84

7 Cross-Platform Issues for NSS 87

7.1 Cross-Platform Issues for NSS Pool Snapshots . . . 87
7.2 Cross-Platform Issues for NSS Volumes . . . 87
7.2.1 Guidelines for File Systems on Linux Servers . . . 88
7.2.2 Guidelines for File Systems on NetWare Servers . . . 88
7.3 Cross-Platform Issues for NSS Features . . . 88
7.3.1 Multipath I/O to Devices . . . 89
7.3.2 NSS Support for Removable Media . . . 89
7.3.3 NSS Support for Transaction Tracking . . . 89
7.3.4 Modified File List . . . 89
7.4 Cross-Platform Issues for File Access . . . 89
7.5 Cross-Platform Issues for Management Tools . . . 90
7.5.1 Storage-Related Plug-Ins for Novell iManager 2.7 . . . 90
7.5.2 Interoperability of Protocols for the iManager Server and Target Server . . . 90
7.5.3 Do Not Run EVMS Tools When Working in iManager or NSSMU . . . 92
7.5.4 Management Tools for Managing Pool Snapshots . . . 92
7.5.5 Management Capabilities for Software RAIDs . . . 92

8 Cluster-Enabling Shared NSS Devices and Pools with Novell Cluster Services 93

8.1 Cluster-Enabling NSS Pools and Volumes . . . 93
8.2 Guidelines for Cluster-Enabling NSS . . . 93

9 Management Tools for NSS 97

9.1 Novell iManager and Storage-Related Plug-Ins . . . 97
9.1.1 Understanding Storage-Related Plug-Ins . . . 98
9.1.2 Accessing Novell iManager . . . 103
9.1.3 Accessing Roles and Tasks in iManager . . . 104
9.1.4 Selecting a Server to Manage . . . 104
9.1.5 Storage Plug-In Quick Reference . . . 105
9.1.6 Files and Folders Plug-In Quick Reference . . . 110
9.2 NSS Management Utility (NSSMU) Quick Reference . . . 112
9.2.1 NSSMU for Linux Quick Reference . . . 113
9.2.2 NSSMU for NetWare Quick Reference . . . 115
9.3 NSS Commands and Utilities . . . 117
9.3.1 Command Consoles . . . 117
9.3.2 NSS Commands . . . 118
9.3.3 NSS Utilities . . . 118
9.4 Novell NetStorage . . . 118
9.4.1 Prerequisites . . . 119
9.4.2 Accessing NetStorage . . . 119
9.4.3 Configuring File System Trustees, Trustee Rights, and Attributes . . . 119
9.4.4 Purging and Salvaging Deleted Files . . . 120
9.4.5 Browsing Directories and Files . . . 120

9.4.6 Additional Information . . . 120
9.5 Novell Remote Manager . . . 120
9.5.1 Prerequisites for Using Novell Remote Manager . . . 120
9.5.2 Novell Remote Manager for Linux . . . 122
9.5.3 Novell Remote Manager for NetWare . . . 122
9.5.4 Accessing Novell Remote Manager . . . 123
9.5.5 Starting, Stopping, or Restarting Novell Remote Manager on Linux . . . 123
9.6 Novell Client . . . 124
9.6.1 Novell Client for Linux . . . 124
9.6.2 Novell Client for Windows XP/2003 and Vista . . . 124
9.7 ConsoleOne . . . 124
9.8 Virtual File Services, APIs, and Scripts . . . 125

10 Managing Devices 127

10.1 Understanding Devices . . . 127
10.1.1 Device Size Limit . . . 127
10.1.2 Device Types . . . 128
10.2 Making Devices Available to EVMS (Linux) . . . 131
10.3 Viewing Devices on a Server . . . 131
10.4 Viewing Details for a Device . . . 132
10.5 Scanning for Devices (NetWare) . . . 133
10.6 Scanning for Devices (Linux) . . . 134
10.7 Initializing a Disk . . . 134
10.8 Sharing Devices for NSS Pools . . . 135
10.8.1 Understanding Sharing . . . 135
10.8.2 Planning for Device Sharing . . . 136
10.8.3 Configuring the Device's Share State . . . 136
10.9 Viewing Partitions on a Device . . . 136
10.10 Viewing Pools on a Device . . . 138
10.11 Enabling Write-Through Cache Management on SCSI Devices and RAID Controllers (Linux) . . . 138
10.12 What's Next . . . 139

11 Migrating NSS Devices from NetWare to OES 2 Linux 141

11.1 Guidelines for Moving Devices Cross-Platform . . . 141
11.1.1 Media Format . . . 141
11.1.2 Pool Snapshots . . . 142
11.1.3 Cross-Platform Issues . . . 142
11.2 Moving Non-Clustered Devices From NetWare Servers to OES 2 Linux Servers . . . 142
11.2.1 Prerequisites . . . 143
11.2.2 Setting Up File Access For Users on the OES 2 Linux Server . . . 143
11.2.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . 144
11.2.4 Recommissioning Each NSS Pool and Its Volumes on the Destination Server . . . 145
11.2.5 Using Scripts to Decommission and Recommission NSS Volumes . . . 146
11.3 Moving Non-Clustered Devices From NetWare 6.0 to OES 2 Linux . . . 147
11.3.1 Prerequisites . . . 147
11.3.2 Setting Up File Access For Users on the OES 2 Linux Server . . . 148
11.3.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . 148
11.3.4 Recommissioning Each NSS Pool and Its Volumes on the OES Linux Server . . . 149
11.4 Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare . . . 150
11.4.1 Prerequisites . . . 151
11.4.2 Setting Up File Access For Users on the Destination Server . . . 151
11.4.3 Decommissioning Each NSS Pool and Its Volumes on the Original Server . . . 151

11.4.4 Recommissioning Each NSS Pool and Its Volumes on the Destination Server . . . 152
11.5 Moving Clustered Devices with NSS Volumes Cross-Platform . . . 153
11.6 Upgrading NetWare 5.1 NSS Volumes and NetWare Traditional Volumes to NSS Volumes 153

12 Managing Partitions 155

12.1 Understanding Partitions . . . 155
12.1.1 NSS Partitions . . . 155
12.1.2 Viewing Partitions on a Device . . . 156
12.1.3 Understanding Types of Partitions . . . 156
12.1.4 Understanding Partition Details . . . 157
12.2 Creating Partitions . . . 158
12.2.1 Creating Virtual Device Partitions for Software RAIDs . . . 158
12.2.2 Creating NSS Partitions for Pools . . . 158
12.2.3 Creating Cluster Services Partitions . . . 158
12.2.4 Creating iSCSI Partitions (NetWare) . . . 158
12.3 Viewing a List of Partitions . . . 159
12.3.1 Viewing Partitions on a Device . . . 159
12.3.2 Viewing Partitions in a Software RAID Device . . . 159
12.3.3 Viewing Partitions in an NSS Pool . . . 160
12.4 Viewing Details for a Partition . . . 160
12.5 Labeling a Partition . . . 160
12.6 Deleting a Partition . . . 161
12.6.1 Deleting NSS Partitions in an NSS Software RAID Device . . . 161
12.6.2 Deleting NSS Partitions in a Pool . . . 162
12.6.3 Deleting Cluster Services Partitions (NetWare) . . . 162
12.6.4 Deleting an iSCSI Partition (NetWare) . . . 162
12.7 Adding Other Types of File System Partitions (Linux) . . . 162

13 Managing NSS Software RAID Devices 165

13.1 Understanding Software RAID Devices . . . 165
13.2 Planning for a Software RAID Device . . . 167
13.2.1 General Guidelines for Software RAID Devices . . . 167
13.2.2 Guidelines for Software RAID 1 Devices . . . 168
13.2.3 Drive Restrictions for NSS Software RAID 0 and 1 Devices on Linux . . . 169
13.2.4 Choosing a Software RAID Solution . . . 169
13.2.5 Determining the Partition Size . . . 170
13.2.6 Determining the Number of Partitions . . . 170
13.2.7 Determining the Stripe Size for RAID 0 and RAID 5 . . . 171
13.3 Viewing a List of Software RAID Devices on a Server . . . 171
13.4 Viewing Details of a Software RAID Device . . . 172
13.5 Creating Software RAID Devices with iManager . . . 174
13.6 Creating Software RAID Devices with NSSMU . . . 176
13.7 Mirroring an Existing Pool with NSSMU . . . 177
13.8 Creating a RAID 1 Mirror to Duplicate Data . . . 178
13.9 Creating a Software RAID 10 with NSSMU (NetWare) . . . 178
13.10 Creating a Software RAID 15 with NSSMU (NetWare) . . . 179
13.11 Renaming a Software RAID Device . . . 180
13.12 Increasing the Size of a Software RAID Device . . . 180
13.13 Restriping a Software RAID 0 or 5 Device . . . 182
13.14 Replacing a Failed Segment in a Software RAID . . . 183
13.15 Deleting a Software RAID Device . . . 185
13.16 Viewing Pools on a Software RAID Device . . . 185
13.17 Viewing Partitions on a Software RAID Device . . . 186

13.18 Deleting Partitions on a Software RAID Device 186
13.19 Managing Software RAID Devices with NSSMU 187

14 Managing Multipath I/O to Devices (NetWare) 189

14.1 Planning Your Multipath I/O Solution for NSS on NetWare 189
14.2 Enabling Multipath 190
14.3 Configuring the Primary Path and Priorities for Failover Connection Paths. 191
14.4 Viewing Connection Path Details 192
14.5 Setting a Connection Path as Up or Down. 193
14.6 Setting the Primary Path for a Device to Its Default Path 194
14.7 Resetting the Server Registry with Default Priority Settings for a Device 194
14.8 Managing Multiple Paths with Server Console Commands 194

15 Managing Multipath I/O to Devices (Linux) 195

15.1 Understanding Multipath I/O on Linux 195
15.2 NSS Errors When Linux Multipath Is Not Configured. 196
15.3 Configuring Multipath on Linux. 196

16 Managing NSS Pools 197

16.1 Guidelines for Creating a Pool 197
16.2 Creating a Pool 198
16.3 Activating and Deactivating Pools 201
16.4 Increasing the Size of a Pool 202
16.5 Renaming a Pool 203
16.6 Deleting a Pool. 204
16.6.1 Prerequisites for Deleting a Pool 204
16.6.2 Procedure 204
16.7 Viewing Pools on a Server 204
16.8 Viewing Pool Details 205
16.9 Viewing Partition Information for a Pool 206
16.10 Viewing Volume Information for a Pool 207
16.11 Viewing Device Information for a Pool 207
16.12 Preventing Pools from Activating on Multiple Servers 208
16.12.1 Understanding MSAP 208
16.12.2 Enabling or Disabling MSAP for All NSS Pools 209
16.12.3 Enabling or Disabling MSAP for a Given NSS Pool 209
16.12.4 Rebuilding the MSAP Block for a Given NSS Pool 211
16.12.5 Determining If MSAP Is Enabled or Disabled on NSS Pools. 212
16.12.6 Managing MSAP with XML or APIs 212
16.12.7 Additional Information 212
16.13 Updating eDirectory Pool Objects 213
16.14 What's Next 213

17 Verifying and Rebuilding NSS Pools and Volumes 215

17.1 When to Use Verify and Rebuild 215
17.1.1 What Verifying Does Not Do 215
17.1.2 What Rebuilding Does Not Do. 215
17.1.3 Before You Verify a Pool 216
17.1.4 Before You Rebuild a Pool 216
17.2 Verifying and Rebuilding an NSS Pool and Its Volumes 217

17.2.1 Mounting the Volume to Repair Journaled Errors . . . 217
17.2.2 Ruling Out Hardware Causes . . . 217
17.2.3 Verifying the Pool to Identify Metadata Inconsistencies . . . 218
17.2.4 Reviewing Log Files for Metadata Consistency Errors . . . 219
17.2.5 Rebuilding NSS Pools to Repair Metadata Consistency . . . 221
17.3 ReZIDing Volumes in an NSS Pool . . . 223
17.3.1 What Is a ZID? . . . 223
17.3.2 Understanding ReZID . . . 224
17.3.3 When to ReZID . . . 224
17.3.4 Viewing the Highest ZID for a Volume . . . 225
17.3.5 ReZIDing Volumes on Linux . . . 225
17.3.6 ReZIDing Volumes on NetWare . . . 226

18 Managing NSS Pool Snapshots 229

18.1 Understanding Pool Snapshots . . . 229
18.1.1 How Snapshots Work . . . 229
18.1.2 Benefits of Using Snapshots . . . 230
18.2 Guidelines for Using and Managing Pool Snapshots . . . 231
18.2.1 Differences Between Snapshots on Linux and NetWare . . . 232
18.2.2 Guidelines for Creating a Pool Snapshot . . . 233
18.2.3 Guidelines for Creating Pool Snapshots of Clustered Pools . . . 233
18.2.4 Guidelines for Naming Pool Snapshots . . . 233
18.2.5 Guidelines for Choosing the Stored-On Location . . . 234
18.2.6 Guidelines for Maintaining the Stored-On Location . . . 235
18.2.7 Guidelines for Onlining Pool Snapshots . . . 236
18.2.8 Guidelines for Deleting Pool Snapshots . . . 237
18.2.9 Guidelines for Viewing a List of Snapshots that are Stored in a Pool (NetWare) . . . 237
18.2.10 Guidelines for Shredding Deleted Snapshots (NetWare) . . . 237
18.3 Creating a New Pool Snapshot . . . 238
18.3.1 Prerequisites for Creating a Pool Snapshot . . . 238
18.3.2 Creating a Pool Snapshot on Linux . . . 238
18.3.3 Creating a Pool Snapshot on NetWare . . . 240
18.4 Viewing a List of Snapshots for a Given Pool . . . 242
18.5 Viewing Pool Snapshot Details . . . 244
18.6 Renaming a Pool Snapshot (NetWare) . . . 245
18.7 Modifying the Stored-On Location for Snapshots . . . 246
18.7.1 Modifying the Stored-On Partition (Linux) . . . 246
18.7.2 Modifying the Stored-On Pool (NetWare) . . . 246
18.8 Onlining or Offlining a Pool Snapshot . . . 246
18.8.1 Using iManager to Online a Pool Snapshot . . . 246
18.8.2 Using iManager to Offline a Pool Snapshot . . . 247
18.8.3 Using NSSMU for Linux to Online or Offline a Pool Snapshot . . . 247
18.9 Viewing and Managing an Online Pool Snapshot . . . 248
18.10 Restoring Data from an Online Pool Snapshot . . . 250
18.11 Deleting a Pool Snapshot . . . 250
18.11.1 Using iManager to Delete a Pool Snapshot . . . 250
18.11.2 Using NSSMU to Delete a Pool Snapshot (Linux) . . . 251
18.12 Shredding a Deleted Pool Snapshot (NetWare) . . . 251
18.13 Using Commands to Manage Pool Snapshots (NetWare) . . . 252

19 Managing NSS Volumes 253

19.1 Understanding Volume Properties . . . 253
19.1.1 Volume Attributes . . . 254
19.1.2 Encryption Support . . . 257

19.1.3 Enhanced Hard Link Support . . . 258
19.2 Guidelines for NSS Volumes . . . 258
19.2.1 Guidelines for Sizing Volumes . . . 258
19.2.2 Guidelines for Name Spaces . . . 259
19.2.3 Guidelines for NSS Volumes in a Mixed-Node Cluster . . . 259
19.3 Creating Unencrypted NSS Volumes . . . 259
19.4 Configuring Encrypted NSS Volumes with NSSMU . . . 262
19.5 Updating eDirectory Volume Objects . . . 262
19.6 Viewing the Details of an NSS Volume . . . 263
19.7 Viewing Properties of an NSS Volume . . . 263
19.8 Modifying Attributes of an NSS Volume . . . 265
19.9 Modifying the NSS Volume Size . . . 266
19.10 Configuring the Name Space for an NSS Volume . . . 267
19.11 Mounting NSS Volumes with Linux Commands (Linux) . . . 268
19.12 Renaming an NSS Volume . . . 269
19.13 Renaming (Modifying) the Mount Point for an NSS Volume (Linux) . . . 270
19.13.1 Renaming the Mount Point for a New Volume . . . 270
19.13.2 Enabling the Mount Point for the NSS Volume to Be Renamed . . . 270
19.13.3 Renaming the Mount Point for an Existing NSS Volume . . . 270
19.14 Activating and Deactivating an NSS Volume . . . 271
19.15 Mounting and Dismounting an NSS Volume (Linux) . . . 271
19.16 Exporting NSS Volumes for NFS Access (NetWare) . . . 272
19.17 Exporting and Importing NSS Volumes for NFS Access (Linux) . . . 273
19.17.1 Understanding NFS Export and Mount Options . . . 273
19.17.2 Exporting NSS Volumes on Linux for NFSv3 . . . 275
19.17.3 Importing NSS Volumes . . . 277
19.18 Deleting an NSS Volume . . . 278
19.19 Finding the Filename for a Given ZID . . . 279
19.20 Verifying or Rebuilding NSS Volumes . . . 279
19.21 Moving Volumes with DFS . . . 279
19.22 Splitting Volumes with DFS . . . 279
19.23 What's Next . . . 279

20 Managing Encrypted NSS Volumes 281

20.1 Understanding Encrypted Volume Support . . . 281
20.1.1 Encryption Method . . . 282
20.1.2 Encryption Password . . . 282
20.1.3 How Encrypted Volume Support Works . . . 282
20.1.4 Guidelines for Using Encrypted Volumes . . . 283
20.2 Security Considerations for Encrypted Volumes . . . 283
20.2.1 Choosing a Strong Encryption Password . . . 283
20.2.2 Backing Up Data from an Encrypted Volume . . . 283
20.2.3 Excluding the NSS Cache Memory from Core Dumps . . . 283
20.2.4 Disabling Logs (Linux) . . . 284
20.2.5 Archiving Files from an Encrypted Volume . . . 284
20.2.6 Using Direct I/O to an Encrypted Volume . . . 284
20.2.7 Sharing Encrypted NSS Volumes in a Cluster . . . 284
20.3 Creating an Encrypted Volume . . . 284
20.4 Mounting an Encrypted NSS Volume with NSSMU (Linux) . . . 285
20.5 Activating an Encrypted NSS Volume with NSSMU (NetWare) . . . 286
20.6 Mounting Encrypted NSS Volumes with NSS Commands (Linux) . . . 286
20.7 Activating Encrypted NSS Volumes with NSS Commands (NetWare) . . . 287
20.8 Dismounting an Encrypted NSS Volume with NSSMU (Linux) . . . 288

20.9 Deactivating an Encrypted NSS Volume with NSSMU (NetWare) . . . 288
20.10 Using Encrypted Volumes in a Server Cluster (Linux) . . . 288
20.11 Using Encrypted Volumes in a Server Cluster (NetWare) . . . 289
20.12 Removing Encrypted Volumes . . . 290
20.13 What's Next . . . 290

21 Managing Removable Media (NetWare) 291

21.1 Managing CDs and DVDs as NSS Volumes (NetWare) . . . 291
21.1.1 Understanding CDs and DVDs on NetWare . . . 291
21.1.2 Mounting a CD or DVD as an NSS Volume . . . 292
21.1.3 Creating a Volume Object for a CD or DVD . . . 292
21.1.4 Specifying the Load Order of Formats for Dual-Formatted . . . 292
21.1.5 Loading and Unloading the cddvd.nss Module . . . 292
21.2 Activating and Deactivating ISO Image Files as NSS Volumes (NetWare) . . . 293
21.2.1 Guidelines for Using Image Files as NSS Volumes . . . 293
21.2.2 Mounting an Image Volume . . . 293
21.2.3 Removing an Image Volume . . . 294
21.3 Activating and Deactivating DOS Partitions as NSS Volumes (NetWare) . . . 294
21.4 Mounting a USB Disk as a Local FAT Volume or NSS Volume (NetWare) . . . 295

22 Securing Access to NSS Volumes, Directories, and Files 297

22.1 Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes . . . 297
22.1.1 Prerequisites for Configuring Trustees . . . 297
22.1.2 Viewing Properties of a File or Folder . . . 298
22.1.3 Configuring File or Folder Attributes . . . 298
22.1.4 Configuring Rights Properties (File System Trustees, Trustee Rights, and Inherited Rights Filter) . . . 300
22.1.5 Viewing Effective Rights for a Trustee . . . 303
22.2 Configuring the Security Equivalence Vector Update Frequency (Linux) . . . 303
22.2.1 Understanding the SEV . . . 304
22.2.2 Enabling or Disabling the Background SEV Update for NSS Volumes on Linux . . . 304
22.2.3 Configuring the Background SEV Update Interval . . . 305
22.2.4 Forcing a Background SEV Update . . . 305
22.3 Using Data Shredding to Prevent Access to Purged Files . . . 306
22.3.1 Setting the Data Shredding Attribute When You Create a Volume . . . 306
22.3.2 Setting the Data Shredding Attribute for an Existing Volume . . . 306
22.3.3 Disabling Data Shredding for an Existing Volume . . . 307
22.4 Enabling or Disabling LAF Audit Log Messages for Trustee Events (Linux) . . . 307
22.4.1 Understanding NSS Audit Log Messages . . . 307
22.4.2 Enabling or Disabling LAF Audit Messages for Trustee Events (Linux) . . . 310
22.4.3 Viewing LAF Audit Messages . . . 310
22.4.4 Additional Information . . . 311
22.5 Controlling Access to the _admin Volume and Files.cmd (NetWare) . . . 311
22.5.1 Assigning Volume Rights with Trustees.xml (NetWare) . . . 311
22.5.2 Restricting Access to Files.cmd (NetWare) . . . 313

23 Managing Data Integrity on NetWare Servers 315

23.1 Enabling Flush Files Immediately to Write Data to the Disk on Close . . . 315
23.1.1 Setting the Flush Files Immediately Attribute When You Create a Volume . . . 316
23.1.2 Setting the Flush Files Immediately Attribute for an Existing Volume . . . 316
23.1.3 Enabling and Disabling the Flush Files Immediately Attribute for All Volumes on a Server . . . 317

23.2 Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare) . . . 317
23.2.1 Enabling TTS . . . 317
23.2.2 Disabling TTS . . . 318
23.2.3 TTS Log File and Transaction Recovery . . . 318
23.3 Using the File-Level Snapshot Attribute to Enable the Backup of Open Files . . . 318
23.3.1 Understanding File-Level Snapshots . . . 319
23.3.2 Guidelines for Using File-Level Snapshots . . . 320
23.3.3 Enabling the File Snapshot Attribute . . . 321
23.3.4 Configuring Backup Software to Use File Snapshots . . . 322
23.3.5 Restoring Volumes from Backup Media . . . 323

24 Managing Compression on NSS Volumes 325

24.1 Understanding Compression . . . 325
24.1.1 Compression and Decompression Processes . . . 325
24.1.2 Compression Settings . . . 326
24.1.3 Guidelines for Compression . . . 327
24.1.4 Factors Affecting Compression . . . 329
24.1.5 Factors Affecting Decompression . . . 330
24.1.6 Monitoring Compression Activity . . . 330
24.2 Configuring Compression for a Server . . . 330
24.2.1 Understanding Server-Level Compression Parameters . . . 331
24.2.2 Configuring Server-Level Compression Parameters with Commands . . . 334
24.2.3 Configuring Server-Level Compression Parameters with Novell Remote Manager (NetWare) . . . 335
24.2.4 Configuring Server-Level Compression Parameters with the Set Command (NetWare) . . . 335
24.2.5 Configuring Server-Level Compression Parameters with Monitor (NetWare) . . . 336
24.3 Configuring a Volume for Compression . . . 336
24.3.1 Enabling Compression for a New Volume . . . 336
24.3.2 Enabling Compression for an Existing Volume . . . 336
24.4 Suspending Compression for Volumes or Files . . . 337
24.5 Disabling Compression for a Volume . . . 338
24.6 Restoring Data to a Uncompressed Volume . . . 338
24.7 Configuring Compression Preferences for Directories and Files . . . 338
24.7.1 Using the Novell Client . . . 338
24.7.2 Using ATTRIB (Linux) . . . 339
24.7.3 Using FLAG (NetWare) . . . 340
24.8 Using NSS Commands to Configure and Monitor Compression . . . 342
24.9 Copying Compressed Files between Volumes with the NCOPY Command . . . 343
24.10 Viewing Compression Statistics with the NDIR Command . . . 343
24.11 Repairing Compressed Volumes with the Compfix Utility . . . 344
24.12 Backing Up Compressed Files . . . 344

25 Managing Space Quotas for Volumes, Directories, and Users 345

25.1 Understanding Space Quotas . . . 345
25.2 Managing NSS Volume Quotas . . . 346
25.3 Managing Directory Quotas . . . 348
25.3.1 Enabling or Disabling the Directory Quotas Attribute for an NSS Volume . . . 349
25.3.2 Configuring Directory Quotas . . . 349
25.3.3 Removing a Directory Quota . . . 352
25.3.4 Removing All Directory Quotas for an NSS Volume . . . 353
25.4 Managing User Space Quotas . . . 353
25.4.1 Setting the User Space Quotas Attribute for an NSS Volume . . . 353

25.4.2 Viewing User Space Quotas . . . 354
25.4.3 Configuring a User Space Quota . . . 356
25.4.4 Modifying a User Space Quota . . . 357
25.4.5 Deleting a User Space Quota . . . 358
25.4.6 Configuring User Space Quotas on Volumes After Upgrading or Migrating from OES 1 Linux (Linux) . . . 358

26 Salvaging and Purging Deleted Volumes, Directories, and Files 361

26.1 Understanding the NSS Salvage System . . . 361
26.1.1 Volume Salvage versus File Salvage . . . 361
26.1.2 Trustees and Rights Handling for Salvaged Data . . . 363
26.1.3 Understanding Purging Triggers . . . 363
26.1.4 Comparison of Salvage for NSS Volumes versus NetWare Traditional Volumes . . . 364
26.2 Configuring the Purge Behavior for NSS . . . 364
26.2.1 Setting the Purge Delay for All Deleted Volumes . . . 364
26.2.2 Setting the Immediate Purge of Deleted Files for All NSS Volumes . . . 365
26.2.3 Setting the Low and High Salvage Watermarks for Automatically Purging Deleted Directories and Files . . . 365
26.2.4 Setting the Purge Immediate Attribute for a Directory or File . . . 367
26.3 Enabling or Disabling the Salvage Attribute for an NSS Volume . . . 367
26.4 Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool . . . 368
26.5 Salvaging or Purging Deleted Files with iManager . . . 369
26.5.1 Prerequisites . . . 369
26.5.2 Salvaging a Deleted File . . . 370
26.5.3 Purging Deleted Files . . . 371
26.6 Salvaging or Purging Deleted Files with Other Tools . . . 371
26.6.1 Using NetStorage . . . 371
26.6.2 Using the Novell Client . . . 372
26.6.3 Using Novell Remote Manager for NetWare . . . 372
26.7 Purging Deleted Files or Directories (NetWare) . . . 373
26.7.1 Purging All Deleted Files . . . 373
26.7.2 Purging Specific Directories or Files . . . 373

27 Managing Hard Links 375

27.1 Understanding Hard Links . . . 375
27.1.1 Hard Links and the Primary Link . . . 376
27.1.2 Hard Links and File Ownership . . . 376
27.1.3 Hard Links and File System Trustees . . . 377
27.1.4 Hard Links and Directory Space Quotas . . . 377
27.1.5 Hard Links and User Space Quotas . . . 377
27.1.6 Hard Links and the Hard Links Attribute . . . 378
27.1.7 Hard Links and File Salvage and Purge . . . 378
27.1.8 Hard Links and DFS Move Volume . . . 378
27.2 Upgrading the Media Format for Hard Links Support . . . 379
27.2.1 New Metadata Structure Supports Up to 65,535 Hard Links for a File . . . 379
27.2.2 Old Metadata Structure Supports Limited Hard Links for a File . . . 379
27.3 Enabling or Disabling the Hard Links Attribute . . . 380
27.3.1 Prerequisite . . . 380
27.3.2 Hard Links Attribute Commands . . . 380
27.3.3 Viewing the Hard Link Attribute Setting . . . 380
27.4 Creating a Hard Link Using Ln on a Linux Client . . . 381
27.5 Creating a Hard Link Using a zLink API . . . 382
27.6 Creating a Hard Link for Testing Purposes . . . 382
27.6.1 Prerequisites . . . 382

27.6.2 Procedure . . . 383
27.7 Viewing Hard Links for a File . . . 383
27.7.1 Linux . . . 384
27.7.2 NetWare . . . 384
27.8 Deleting a Hard Link . . . 384
27.9 Deleting a Primary Link . . . 384

28 Managing Files and Folders on NSS Volumes 387

28.1 Creating a Folder on an NSS Volume . . . 387
28.1.1 Prerequisites . . . 387
28.1.2 Procedure . . . 387
28.2 Deleting a File or Folder on an NSS Volume . . . 388
28.2.1 Prerequisites . . . 388
28.2.2 Procedure . . . 388
28.3 Uploading Files to an NSS Volume . . . 389
28.3.1 Prerequisites . . . 389
28.3.2 Procedure . . . 389
28.4 Downloading Files from an NSS Volume . . . 390
28.4.1 Prerequisites . . . 390
28.4.2 Procedure . . . 390
28.5 Viewing or Modifying File or Folder Properties . . . 391
28.6 Viewing or Modifying File Ownership . . . 394
28.7 Viewing, Adding, Modifying, or Removing a Directory Quota . . . 396

29 Managing Backup and Restore for Data and Trustee Information 399

29.1 Using Novell Storage Management Services . . . 399
29.2 Using the Event File List to Refine the Backup . . . 400
29.3 Using METAMIG to Save and Restore Trustee Information on NSS and Linux POSIX File Systems (Linux) . . . 400
29.4 Using Extended Attributes (xAttr) Commands (Linux) . . . 400
29.4.1 Enabling NSS Support for Linux xAttr . . . 401
29.4.2 Disabling NSS Support for Linux xAttr . . . 401
29.4.3 Additional Information . . . 402
29.5 Backing Up Files Without Altering the Access Time . . . 402
29.5.1 Linux . . . 402
29.5.2 NetWare . . . 402
29.6 Additional Information . . . 403

30 Tuning NSS Performance on NetWare 405

30.1 Do I Need to Tune NSS? . . . 405
30.2 NSS Tuning Guidelines . . . 405
30.2.1 Before You Adjust Cache Parameters . . . 405
30.2.2 Hints for Low Memory NetWare Systems . . . 406
30.2.3 The Tuning Process . . . 406
30.2.4 Restoring Settings After a Failed Tuning Attempt . . . 407
30.3 Setting the Cache Balance . . . 407
30.4 Setting the Closed File Cache Size . . . 410
30.5 Setting the File and Buffer Flush Timers . . . 410
30.6 Setting the Name Cache Size . . . 411
30.7 Configuring NSS Extended System Memory . . . 411
30.8 Tuning NSS Write Performance on NetWare . . . 412

31 Tuning NSS Performance on Linux 413

31.1 Tuning Cache Buffers for NSS on Linux . . . 413
31.1.1 Understanding How NSS Uses Cache on Linux . . . 413
31.1.2 Setting the Minimum Number of Cache Buffers to Use for Kernel Memory . . . 413
31.2 Configuring High Memory Cache . . . 414
31.2.1 Configuring High Memory Cache Type and Size . . . 414
31.2.2 Monitoring the High Memory Cache Use . . . 415
31.3 Configuring or Tuning Group I/O . . . 415
31.3.1 Viewing the Metadata Area Size . . . 416
31.3.2 Configuring the Journal Group Write Timer . . . 417
31.3.3 Configuring the Metadata Group Write Timer and Limit . . . 417
31.3.4 Configuring the User Data Group Write Timer . . . 419
31.3.5 Viewing Group Write Policies . . . 419

32 Monitoring the Status of the NSS File System and Services 421

32.1 Monitoring Status of NSS Devices, Pools, and Volumes with iManager . . . 421
32.2 Monitoring Compression and Salvage Statistics . . . 422
32.3 Monitoring Quotas and Space Usage for NSS Pools and Volumes . . . 424
32.4 Monitoring File System Parameters . . . 426
32.4.1 Using iManager to Monitor NSS File System Parameters . . . 426
32.4.2 Using Novell Remote Manager to Monitor NSS File System Parameters . . . 426
32.4.3 Using Novell NetStorage to Monitor NSS File System Parameters . . . 427

33 Troubleshooting the NSS File System 429

33.1 Cannot Connect to Target Servers from iManager . . . 430
33.2 Cannot Create an NSS Pool or Volume . . . 430
33.2.1 General Issues . . . 430
33.2.2 Linux Issues . . . 430
33.2.3 NetWare Issues . . . 431
33.3 Cannot See NSS Devices, Pools, or Volumes . . . 431
33.3.1 Connection Failure . . . 431
33.3.2 NSS Tools Recognize Only EVMS-Managed Devices (Linux) . . . 431
33.3.3 EVMS Not Running (Linux) . . . 432
33.3.4 Pools Not Mounted by EVMS (Linux) . . . 432
33.4 Data Corruption Is Possible with Xen-Based File-Backed Block Devices on NetWare Virtual Machines . . . 432
33.5 eDirectory Errors When Creating NSS Pools or Volumes on Linux . . . 433
33.5.1 eDirectory Error 613 When Creating an NSS Pool or Volume on Linux . . . 433
33.5.2 eDirectory Error 672 When Creating an NSS Pool on Linux . . . 433
33.6 File Compression Is Not Working . . . 434
33.7 Linux Fails to Execute Dismount, Umount, or Delete Commands for NSS Volumes . . . 434
33.8 Multipath Devices Are Not Resolved . . . 434
33.9 NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux) . . . 434
33.10 NSS Volume Disappears from the NCP Console (ncpcon) (Linux) . . . 435
33.11 Pathname Exceeds 255-Character Limit . . . 436
33.12 Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers . . . 436
33.13 Server Hangs When Using an NSS Volume as a Netatalk Share . . . 436
33.14 Slow Mount Performance for NSS Volumes Using the UNIX Name Space (Linux) . . . 437
33.15 Software RAID 1 Fails to Recognize a Replacement Device . . . 437
33.16 Tuning NSS Volumes for GroupWise Servers . . . 437
33.17 Unknown Users as File Owners . . . 437
33.18 Using Scripts to Create or Manage NSS Pools and Volumes on Linux . . . 438

34 Security Considerations 439

34.1 Security Features of NSS . . . 439
34.2 Preventing Unauthorized Access . . . 441
34.3 Securing Sensitive Information During Remote Management Sessions . . . 441
34.4 Protecting Data During Backup and on Backup Media . . . 442
34.5 Preventing Exposure of Sensitive Data in a Core Dump . . . 442
34.6 Preventing Exposure of the Encryption Password in a Log (Linux) . . . 442
34.6.1 NSS Logging . . . 442
34.6.2 NSS Logging and Security Implications . . . 443
34.6.3 Logging Communications between NSS and the _ADMIN Volume . . . 443
34.6.4 Logging Communications between NSS and eDirectory, NICI, or Linux User Management . . . 445
34.7 Using Data Shredding to Prevent Unauthorized Access to Purged Files . . . 446
34.8 Protecting the Admin Volume and Files.cmd (NetWare) . . . 447
34.9 Acquiring Security Equivalence Vectors for NSS Users . . . 447
34.10 Protecting Modules Responsible for Security Equivalence Vectors (Linux) . . . 447
34.11 Controlling File System Access and Attributes for NSS Volumes . . . 447
34.12 Displaying Directory and File Attributes for NSS Volumes on Linux . . . 448
34.13 Security Best Practices for zAPIs . . . 448
34.14 Controlling Physical Access to Servers and Resources . . . 449
34.15 Securing Access to the Servers with a Firewall . . . 449
34.16 Creating Strong Passwords . . . 449

A NSS Commands 451

A.1 Using NSS Commands . . . 452
A.1.1 Issuing NSS Commands at Command Consoles . . . 452
A.1.2 Issuing NSS Commands at Server Startup (NetWare) . . . 452
A.1.3 Making NSS Commands Persist Through a Reboot . . . 453
A.1.4 Permissions . . . 453
A.1.5 Descriptions . . . 453
A.2 Help and Find Commands . . . 454
A.3 Access Time Command . . . 454
A.4 Background File System Checker Commands . . . 455
A.5 Cache Management Commands (Linux) . . . 455
A.5.1 Cache Command . . . 455
A.5.2 High Memory Cache Commands . . . 455
A.5.3 ID Cache Commands . . . 456
A.5.4 Cache Monitoring Commands . . . 457
A.5.5 UnplugAlways Command for the Read Queue . . . 457
A.6 Cache Management Commands (NetWare) . . . 457
A.6.1 Buffer Cache Commands . . . 458
A.6.2 File Cache Commands . . . 458
A.6.3 Extended System Memory (ESM) Commands . . . 459
A.6.4 ID Cache Commands . . . 459
A.6.5 Name Cache Command . . . 460
A.6.6 Time Cache Command . . . 460
A.6.7 Cache Monitoring Commands . . . 460
A.7 Compression Commands . . . 460
A.7.1 Server-Level Compression Parameters . . . 460
A.7.2 Volume-Level Compression Parameters . . . 462
A.8 Data Shredding Commands . . . 463
A.9 Daylight Savings Time Commands . . . 463
A.10 eDirectory Storage Object ID Commands . . . 464

A.11 Extended Attributes (XAttr) Commands (Linux) . . . 464
A.11.1 CtimeIsMetadataModTime Option . . . 465
A.11.2 ListXattrNWmetadata Option . . . 465
A.11.3 Additional Information . . . 467
A.12 Event File List (EFL) Command . . . 467
A.13 File-Level Snapshot Commands . . . 467
A.14 Hard Links Commands . . . 468
A.14.1 Hard Links Media Upgrade Commands for Existing Volumes . . . 468
A.14.2 Hard Links Media Upgrade Commands for New Volumes . . . 468
A.14.3 Hard Links Attribute Commands . . . 469
A.14.4 Hard Links Management Commands . . . 469
A.15 I/O Monitoring Commands . . . 470
A.16 I/O Write Commands (Linux) . . . 470
A.17 I/O Write Commands (NetWare) . . . 471
A.18 LAF Audit Log Messages Commands (Linux) . . . 472
A.19 Load Commands for the nssstart.cfg File . . . 473
A.20 Low Storage Alert Messages Commands . . . 474
A.21 Migration Support Commands for Near-Line Storage . . . 474
A.22 Modified File List (MFL) Commands (NetWare) . . . 474
A.23 Multipath I/O Failover Commands (NetWare) . . . 475
A.24 Multiple Server Activation Prevention (MSAP) Commands . . . 477
A.25 noatime and atime Commands . . . 478
A.25.1 Using noatime or atime at the Command Line . . . 478
A.25.2 Using noatime in a Cluster Load Script . . . 478
A.25.3 Viewing the atime or noatime Setting . . . 479
A.26 noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux) . . . 479
A.26.1 Linux open(2) Command . . . 479
A.26.2 Linux mount Command . . . 480
A.26.3 Linux nfsmount Command . . . 480
A.26.4 Linux /etc/fstab File . . . 480
A.26.5 Patches for km_nss . . . 481
A.27 Opportunistic Locking (Oplock) Commands . . . 481
A.28 Pool Freeze and Thaw Commands . . . 481
A.29 Pool Management Commands . . . 482
A.29.1 Pool Status . . . 482
A.29.2 PoolAuto Commands for Load Time . . . 482
A.30 Pool Snapshot Commands (NetWare) . . . 483
A.31 Pool Verify and Rebuild Commands (NetWare) . . . 484
A.32 POSIX Permission Mask Command (Linux) . . . 485
A.33 Quotas Commands . . . 485
A.33.1 Sys: Volume Quota Command . . . 485
A.33.2 Directory Quotas Commands . . . 485
A.33.3 User Quotas Commands . . . 486
A.34 Read Ahead Blocks and Allocate Ahead Blocks Commands . . . 486
A.34.1 Read Ahead Blocks . . . 486
A.34.2 Allocate Ahead Blocks . . . 487
A.35 Removable Media Commands (NetWare) . . . 487
A.35.1 CDs and DVDs Commands (NetWare) . . . 487
A.35.2 CD and DVD Image Files Commands (NetWare) . . . 488
A.35.3 DOS FAT Volumes Commands (NetWare) . . . 489
A.36 Salvage and Purge Commands . . . 489
A.37 Security Equivalence Vector Update Commands (Linux) . . . 490
A.38 Sendfile API Support Command . . . 491
A.39 Status Commands . . . 492
A.40 Transaction Tracking System (TTS) Commands (NetWare) . . . 492

A.41 Visibility Rebuild Command . . . 493
A.41.1 Description . . . 493
A.41.2 Syntax . . . 494
A.41.3 Additional Information . . . 494
A.42 Volume Management Commands . . . 494
A.42.1 System Volume Commands (NetWare) . . . 494
A.42.2 Volumes Command . . . 495
A.42.3 Volume Activity Commands . . . 496
A.42.4 Encrypted Volume Activity Commands . . . 496
A.42.5 VolumeAuto Commands for Load Time . . . 497
A.43 Upgrade Commands for NetWare 5.*x* or 6.*x* to NetWare 6.5 . . . 498
A.44 ZID Commands . . . 498

B NSS Utilities 499

B.1 ABORT REMIRROR (NetWare) . . . 499
B.1.1 Syntax . . . 499
B.1.2 Option . . . 500
B.1.3 Example . . . 500
B.2 ATTRIB (Linux) . . . 500
B.2.1 Syntax . . . 500
B.2.2 Options . . . 500
B.2.3 Attributes . . . 500
B.2.4 Example . . . 502
B.3 CDDVD (NetWare) . . . 502
B.3.1 Syntax . . . 502
B.3.2 Notes . . . 502
B.4 COMPFIX . . . 502
B.4.1 Purpose . . . 503
B.4.2 COMPFIX for Linux . . . 503
B.4.3 COMPFIX for NetWare . . . 506
B.5 DISMOUNT (NetWare) . . . 508
B.5.1 Syntax . . . 508
B.5.2 Option . . . 508
B.5.3 Notes . . . 508
B.6 FILER (NetWare) . . . 508
B.6.1 Options . . . 508
B.7 FLAG (NetWare) . . . 509
B.7.1 Syntax . . . 509
B.7.2 Options . . . 510
B.7.3 Attributes . . . 511
B.7.4 Examples . . . 513
B.8 METAMIG (Linux) . . . 514
B.8.1 Syntax . . . 514
B.8.2 Arguments . . . 514
B.8.3 Options . . . 514
B.8.4 Examples . . . 516
B.9 MIRROR STATUS (Netware) . . . 516
B.9.1 Syntax . . . 516
B.9.2 Option . . . 516
B.9.3 Status Messages . . . 516
B.10 MOUNT (NetWare) . . . 517
B.10.1 Syntax . . . 517
B.10.2 Notes . . . 517
B.10.3 Examples . . . 517
B.11 NCOPY (NetWare) . . . 518

B.11.1 Syntax . . . 518
B.12 NDIR . . . 518
B.12.1 Syntax . . . 518
B.13 NSSCON (Linux) . . . 519
B.13.1 Adding /opt/novell/nss/sbin to the PATH Environment Variable . . . 520
B.13.2 Starting nsscon . . . 520
B.13.3 Using nsscon in a Script . . . 520
B.14 NSSMU . . . 521
B.15 NSSUPDATE (Linux) . . . 521
B.15.1 Syntax . . . 521
B.15.2 Availability . . . 521
B.15.3 Options . . . 521
B.15.4 Example . . . 522
B.16 RAVSUI (Linux) . . . 522
B.16.1 Syntax . . . 522
B.16.2 Arguments . . . 523
B.16.3 Options . . . 523
B.16.4 Files . . . 524
B.16.5 Note . . . 525
B.16.6 Example . . . 525
B.17 RAVVIEW (Linux) . . . 525
B.17.1 Syntax . . . 526
B.17.2 Arguments . . . 526
B.17.3 Options . . . 527
B.17.4 Files . . . 527
B.17.5 Note . . . 528
B.17.6 Example . . . 528
B.18 REBUILD (NetWare) . . . 528
B.18.1 Syntax . . . 528
B.18.2 Starting and Using REBUILD . . . 529
B.19 REFRESHIDS (Linux) . . . 529
B.19.1 Syntax . . . 529
B.20 REMIRROR PARTITION (NetWare) . . . 529
B.20.1 Syntax . . . 530
B.20.2 Example . . . 530
B.21 RIGHTS . . . 530
B.21.1 RIGHTS (Linux) . . . 530
B.21.2 RIGHTS (NetWare) . . . 534
B.22 SET (NetWare) . . . 536
B.22.1 Syntax . . . 536
B.22.2 Using SET . . . 536
B.22.3 Common File System Parameters . . . 537
B.22.4 Disk Parameters . . . 541
B.22.5 Novell Storage Services Parameters . . . 543
B.23 VERIFY (NetWare) . . . 544
B.23.1 Syntax . . . 544
B.23.2 Starting and Using VERIFY . . . 545
B.24 VOLUMES (Linux, NCP Utility) . . . 545
B.24.1 Syntax . . . 545
B.24.2 Using VOLUMES . . . 545
B.24.3 Using VOLUME Name . . . 546

C Upgrading Legacy NSS and NetWare Traditional Volumes 547

C.1 Migrating Legacy NSS Volumes and NetWare Traditional Volumes to OES Linux . . . 547
C.2 Copying and Upgrading Legacy Volumes with Volume Copy Upgrade . . . 547

C.2.1 Understanding Volume Copy Upgrade . . . 548
C.2.2 Using VCU from the Command Line . . . 549
C.2.3 Using VCU from a Java Application . . . 551
C.2.4 After the Volume Copy or Volume Upgrade . . . 554
C.3 Upgrading Legacy NSS Volumes . . . 554
C.4 Upgrading Legacy NSS Volumes in Clusters . . . 555

D Using EVMS to Manage Devices with NSS Volumes (Linux) 557

D.1 Installing Linux with EVMS as the Volume Manager on the System Device . . . 557
D.2 Using Linux POSIX File Systems and NSS Pools and Volumes on the Same Device . . . 558
D.3 Viewing Free Space Available to EVMS-Managed Devices . . . 560
D.4 Making Devices Available to EVMS . . . 561
D.4.1 Mounting Your System Device with EVMS . . . 561
D.4.2 Converting LVM2 Volumes to EVMS Volumes . . . 566
D.4.3 Patching the 2.6 Kernel (Unsupported) . . . 567
D.5 Additional EVMS Documentation . . . 567

E Using NSS on Devices Managed by Non-EVMS Volume Managers (Linux) 569

E.1 FAQs About Using EVMS with NSS . . . 569
E.1.1 Why Should I Use EVMS with NSS? . . . 569
E.1.2 NSS Limitations for Non-EVMS Volume Managers . . . 570
E.2 Guidelines for Using NSS with Non-EVMS Volume Managers . . . 570
E.2.1 Guidelines for the Server Install . . . 571
E.2.2 NSS Management Considerations . . . 571
E.3 Creating and Mounting NSS Pools and Volumes Using Linux Commands . . . 571
E.3.1 Using the Linux mkfs Command to Create NSS Pools . . . 572
E.3.2 Creating a Partition . . . 572
E.3.3 Creating and Mounting an NSS Pool . . . 573
E.3.4 Creating and Mounting an NSS Volume . . . 574
E.4 Configuring Default Mount Settings for NSS Pools and Volumes . . . 575
E.4.1 Understanding Entries in the /etc/fstab Configuration File . . . 575
E.4.2 Adding NSS Pool and Volume Mount Information to /etc/fstab . . . 576
E.5 Expanding NSS Pools on Non-EVMS-Managed Devices . . . 577

F Comparison of NSS on NetWare and NSS on Linux 579

G Comparison of NSS on NetWare and the NetWare Traditional File System 585

G.1 Features of NSS and Traditional File Systems on NetWare . . . 585
G.2 SET Parameters Shared by NSS and Traditional File Systems . . . 589
G.3 SET Parameters for the NSS File System . . . 590
G.4 SET Parameters for the Traditional File System . . . 591
G.5 Comparison of Directory and File Attributes . . . 593

H Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems 595

I NSS Nomenclature 601

J Documentation Updates 605

J.1 December 2008 (OES 2 SP1, NetWare 6.5 SP8) . . . 605

J.1.1 Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems . . 606
J.1.2 Comparison of NSS on NetWare and NSS on Linux 606
J.1.3 Cross-Platform Issues for NSS 607
J.1.4 Installing and Configuring Novell Storage Services 607
J.1.5 Management Tools for NSS 607
J.1.6 Managing Backup and Restore for Data and Trustee Information 607
J.1.7 Managing Files and Folders on NSS Volumes 608
J.1.8 Managing NSS Pool Snapshots 608
J.1.9 Managing NSS Pools 608
J.1.10 Managing NSS Software RAID Devices 609
J.1.11 Managing NSS Volumes 610
J.1.12 Managing Space Quotas for Volumes, Directories, and Users 610
J.1.13 Migrating NSS Devices from NetWare to OES 2 Linux 610
J.1.14 NSS Commands 611
J.1.15 NSS Utilities 611
J.1.16 Planning for NSS Storage Solutions 612
J.1.17 Salvaging and Purging Deleted Volumes, Directories, and Files 612
J.1.18 Securing Access to NSS Volumes, Directories, and Files 612
J.1.19 Troubleshooting the NSS File System 613
J.1.20 Tuning NSS Performance on NetWare 613
J.1.21 Using EVMS to Manage Devices with NSS Volumes (Linux) 613
J.1.22 Using NSS in a Virtualization Environment 614
J.1.23 What's New 614
J.2 June 27, 2008 614
J.2.1 Managing Devices 614
J.2.2 Managing Multipath I/O to Devices (NetWare) 614
J.2.3 Managing Space Quotas for Volumes, Directories, and Users 615
J.2.4 NSS Commands 615
J.2.5 Salvaging and Purging Deleted Volumes, Directories, and Files 615
J.2.6 Using NSS in a Virtualization Guest Server Environment 615
J.3 May 14, 2008 616
J.3.1 Installing and Configuring Novell Storage Services 616
J.3.2 Managing Devices 616
J.3.3 NSS Commands 616
J.3.4 NSS Utilities 617
J.4 February 22, 2008 617
J.4.1 Installing and Configuring Novell Storage Services 617
J.4.2 Management Tools for NSS 618
J.4.3 NSS Utilities 618
J.4.4 Troubleshooting the NSS File System 618
J.5 January 14, 2008 618
J.5.1 Managing Devices 619
J.5.2 Managing Files and Folders 619
J.5.3 Management Tools for NSS 619
J.5.4 Preface 619
J.5.5 Securing Access to NSS Volumes, Directories, and Files 620
J.5.6 What's New for NSS 620
J.6 December 14, 2007 620
J.6.1 Installing and Configuring Novell Storage Services 620
J.6.2 Managing Backup and Restore for Data and Trustee Information 621
J.6.3 Management Tools for NSS 621
J.6.4 NSS Commands 621
J.6.5 Planning Your Multipath I/O Solution for NSS on NetWare 621
J.7 November 14, 2007 621
J.7.1 Comparison of NSS on NetWare and NSS on Linux 622
J.7.2 Extended Attributes (XAttr) Commands (Linux) 622

About This Guide

This documentation describes how to use Novell® Storage Services™ File System (NSS) to manage software RAIDs, pools, and volumes on a Novell Open Enterprise Server 2 server or a NetWare® 6.5 SP7 server. Unless otherwise specified, the information in this guide applies to NSS on both NetWare and Linux* platforms.

IMPORTANT: OES 2 NetWare and NetWare 6.5 SP 7 share the same code base and are the same in every way. Installing the OES 2 NetWare product or associated support pack is the same as installing the simultaneously released NetWare 6.5 product or associated support pack.

- Chapter 1, “Overview of NSS,” on page 29
- Chapter 2, “What’s New,” on page 37
- Chapter 3, “Installing and Configuring Novell Storage Services,” on page 45
- Chapter 4, “Upgrading the NSS Media Format,” on page 55
- Chapter 5, “Planning NSS Storage Solutions,” on page 63
- Chapter 6, “Using NSS in a Virtualization Environment,” on page 79
- Chapter 7, “Cross-Platform Issues for NSS,” on page 87
- Chapter 8, “Cluster-Enabling Shared NSS Devices and Pools with Novell Cluster Services,” on page 93
- Chapter 9, “Management Tools for NSS,” on page 97
- Chapter 10, “Managing Devices,” on page 127
- Chapter 11, “Migrating NSS Devices from NetWare to OES 2 Linux,” on page 141
- Chapter 12, “Managing Partitions,” on page 155
- Chapter 13, “Managing NSS Software RAID Devices,” on page 165
- Chapter 14, “Managing Multipath I/O to Devices (NetWare),” on page 189
- Chapter 15, “Managing Multipath I/O to Devices (Linux),” on page 195
- Chapter 16, “Managing NSS Pools,” on page 197
- Chapter 17, “Verifying and Rebuilding NSS Pools and Volumes,” on page 215
- Chapter 18, “Managing NSS Pool Snapshots,” on page 229
- Chapter 19, “Managing NSS Volumes,” on page 253
- Chapter 20, “Managing Encrypted NSS Volumes,” on page 281
- Chapter 21, “Managing Removable Media (NetWare),” on page 291
- Chapter 22, “Securing Access to NSS Volumes, Directories, and Files,” on page 297
- Chapter 23, “Managing Data Integrity on NetWare Servers,” on page 315
- Chapter 24, “Managing Compression on NSS Volumes,” on page 325
- Chapter 25, “Managing Space Quotas for Volumes, Directories, and Users,” on page 345
- Chapter 26, “Salvaging and Purging Deleted Volumes, Directories, and Files,” on page 361
- Chapter 27, “Managing Hard Links,” on page 375

- Chapter 28, “Managing Files and Folders on NSS Volumes,” on page 387
- Chapter 29, “Managing Backup and Restore for Data and Trustee Information,” on page 399
- Chapter 30, “Tuning NSS Performance on NetWare,” on page 405
- Chapter 31, “Tuning NSS Performance on Linux,” on page 413
- Chapter 32, “Monitoring the Status of the NSS File System and Services,” on page 421
- Chapter 33, “Troubleshooting the NSS File System,” on page 429
- Chapter 34, “Security Considerations,” on page 439
- Appendix A, “NSS Commands,” on page 451
- Appendix B, “NSS Utilities,” on page 499
- Appendix C, “Upgrading Legacy NSS and NetWare Traditional Volumes,” on page 547
- Appendix D, “Using EVMS to Manage Devices with NSS Volumes (Linux),” on page 557
- Appendix E, “Using NSS on Devices Managed by Non-EVMS Volume Managers (Linux),” on page 569
- Appendix F, “Comparison of NSS on NetWare and NSS on Linux,” on page 579
- Appendix G, “Comparison of NSS on NetWare and the NetWare Traditional File System,” on page 585
- Appendix H, “Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems,” on page 595
- Appendix I, “NSS Nomenclature,” on page 601

Audience

This guide is intended for network administrators. Chapter 34, “Security Considerations,” on page 439 is intended for security administrators or anyone who is using NSS storage objects and is responsible for the security of the system.

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 and enter your comments there.

Documentation Updates

For the most recent version of the *OES 2: Novell Storage Services File System Administration Guide*, see the latest Novell Open Enterprise Server 2 documentation (http://www.novell.com/documentation/oes2/index.html).

Additional Documentation

For information about planning and implementing storage solutions in Novell Open Enterprise Server 2, see the following:

- The “Storage and File Systems” section in the *OES 2 SP1: Planning and Implementation Guide* describes considerations for choosing a storage solution and system-wide caveats for implementing the different storage solutions.

- *OES 2 SP1: Storage and File Services Overview* describes typical requirements for system storage, and identifies the various storage products and services in Novell Open Enterprise Server 2 that address those requirements.
- *File System Primer* (http://wiki.novell.com/index.php/File_System_Primer) describes the variety of file systems available on Linux and which ones are the best to use for which workloads and data.

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

- The *OES 2: Novell Distributed File Services Administration Guide* describes how to configure and manage DFS services for NSS volumes on NetWare or Linux.
- The *OES 2 SP1: Dynamic Storage Technology Administration Guide* describes how to configure NSS volumes as shadow volumes by using Dynamic Storage Technology.
- Novell Archive and Version Services provides interval-based archiving for user data. See the following guides:
  - *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*
  - *OES 2: Novell Archive and Version Services 2.1 for NetWare Administration Guide*
- The *OES 2 SP1: File Systems Management Guide* describes the Novell trustee model and how to configure file system trustees, trustee rights, and attributes for NSS volumes on Linux and NetWare, NetWare Core Protocol™ (NCP™) volumes on Linux, and NetWare Traditional volumes on NetWare.
- The *OES 2 SP1: NCP Server for Linux Administration Guide* describes how to manage NCP™ connections for NSS volumes on Linux.
- The *OES 2 SP1: NetWare Server Disks and Storage Devices* describes how disks are handled by the NetWare operating system.
- The *OES 2 SP1: Novell Linux User Management Technology Guide* describes how to Linux-enable users for an OES 2 Linux server.
- The *NDK: Virtual File Services* (http://developer.novell.com/wiki/index.php/Virtual_File_Services_for_NetWare) describes the software APIs for creating software applications and scripts to manage NSS volumes and services on Linux and NetWare.
- The *NDK: Novell Storage Architecture Component (Media Manager and NWPA)* (http://developer.novell.com/wiki/index.php/Storage_Architecture_Components_%28Media_Manager_and_NWPA%29) describes software APIs for creating storage-related applications.
- Novell Storage Services Error Codes (http://www.novell.com/documentation/nwec/nwec/data/al3s3ui.html)
- The *SLES 10 SP2: Storage Administration Guide* describes storage services such as the EVMS volume manager, UUIDs, Linux multipath I/O for devices; and Linux software RAIDs 0, 1, 5, 6, and 10.
- Enterprise Volume Management System for Linux documentation:
  - *EVMS User Guide* (http://evms.sourceforge.net/user_guide)
  - *EVMS GUI Screen Shots* (http://evms.sourceforge.net/gui_screen/)

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 2 provides the Novell Storage Services™ (NSS) file system for both Linux and NetWare® operating systems. This section describes benefits and key features

- Section 1.1, “Introduction to NSS,” on page 29
- Section 1.2, “Benefits of NSS,” on page 29
- Section 1.3, “Understanding NSS,” on page 30
- Section 1.4, “NSS Features and Capabilities,” on page 32
- Section 1.5, “Comparison of NSS to Other File Systems,” on page 34
- Section 1.6, “What’s Next,” on page 35

1.1 Introduction to NSS

The NSS file system and services provide 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 capabilities can help you effectively manage your shared file storage for any size of organization, scaling from small businesses to even the largest of organizations with hundreds of thousands of employees.

NSS volumes are cross-compatible between Linux and NetWare. You can mount an NSS data volume on either platform. In a mixed-platform cluster with Novell Cluster Services™, NSS volumes can fail over between Linux and NetWare, allowing for full data, trustee, and file system feature preservation when migrating data to Linux.

NSS devices and storage can be managed in the Web-based Novell iManager utility. NSS also supports third-party tools on both platforms for advanced data protection and management, virus scanning, and traditional archive and backup solutions.

1.2 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.
- Access control and visibility management using the Novell trustee model.

- An unlimited number of NSS volumes, with up to 255 mounted concurrently.
- Lower memory requirements: 1 MB of RAM can activate an NSS volume.
- Pools of storage that span multiple devices and support dynamic resizing to grow the pool and volumes.
- Pool snapshots that capture point-in-time versions of files in the pool.
- Software RAID support, including RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping). NetWare tools also support creating 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.
- Up to 4 billion (10E9) files in a single directory, so how you organize files is limited only by the application or file browser, not the file system.
- Faster access to data, regardless of file size or volume size.
- Directory space restrictions.
- User space restrictions.
- Salvage support for deleted volumes and files.
- Data compression.
- Novell Distributed File Services allows you to better manage storage growth by defining virtual file structures with junctions, moving volumes, and splitting volumes. For information, see the *OES 2: Novell Distributed File Services Administration Guide*.
- Novell Archive and Version Services allows you to archive interval-based versions of user files on NSS volumes, and to make them available to users for individual retrieval. For information, see the following:
  - *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*
  - *OES 2: Novell Archive and Version Services 2.1 for NetWare Administration Guide*

NSS provides the following additional benefits for NetWare:

- Multiple I/O path support for devices.
- Removable media (CD, DVD, USB drives, and so on) and image files activated as NSS volumes.
- Transaction Tracking System™ (TTS™) for tracking file content changes to user data and ensuring that the transactions are completed.

1.3 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 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. 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.

1.3.1 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 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*

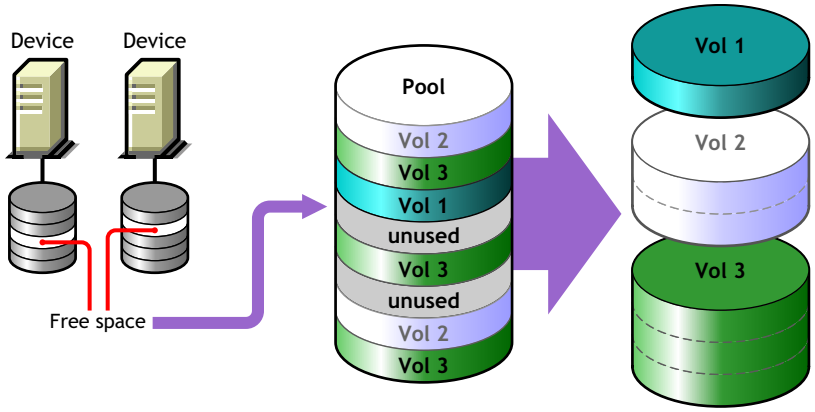

1.3.2 NSS Volumes

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 by adding another segment to it, allowing 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.

1.4 NSS Features and Capabilities

NSS helps improve the scalability, flexibility, and availability of your storage devices. This section identifies specific NSS features that 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

1.4.1 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 (the system volume plus 255 other 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 a 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.

1.4.2 Improve Storage Availability

NSS provides the following features to improve I/O performance and provide fault-tolerant access to your data:

- Multipath I/O support (NetWare only)

  Uses multiple channels to connect the server to local media or a storage area network. For information, see "Managing Multipath I/O to Devices (NetWare)" on page 189.

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

  Uses software RAID devices to improve performance and availability. For information, see Chapter 13, "Managing NSS Software RAID Devices," on page 165.

- Shared-disk storage

  Makes devices shareable for use in a cluster. For information, see Section 10.8, "Sharing Devices for NSS Pools," on page 135.

- Multiple name space support

NSS provides full support for filenames in the Long, UNIX, DOS, and Macintosh* name spaces. Long name space is the default. For information, see Section , “Lookup Namespace,” on page 257.

- Rich file metadata support

  NSS provides full support for all file attributes and multiple simultaneous data streams for DOS, Windows*, UNIX, and Macintosh. For information, see Section 22.1, “Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes,” on page 297.

1.4.3 Prevent Unauthorized Access

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

- Encrypted Volume Support

  Encrypts data volumes, meeting U.S. Government security standards. For information, see “Managing Encrypted NSS Volumes” on page 281.

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

  Erases files completely, meeting U.S. Government security standards. For information, see Section 22.3, “Using Data Shredding to Prevent Access to Purged Files,” on page 306.

- Multiple Server Access Prevention for pools

  Ensures data integrity by preventing unauthorized access to shared media in a storage area network. For information, see Section 16.12, “Preventing Pools from Activating on Multiple Servers,” on page 208.

- 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 *OES 2 SP1: File Systems Management Guide*.

  Some additional steps are necessary to configure access control for NSS on Linux. For information, see Section 5.5, “Access Control for NSS on Linux,” on page 71.

1.4.4 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

  Backs up files from snapshots of data so that all files, including open ones, are backed up. For information, see “Managing NSS Pool Snapshots” on page 229.

- Immediate data saves

  On file closure, immediately writes data from memory to the file instead of waiting for the next scheduled write to disk. For information, see Section 23.1, “Enabling Flush Files Immediately to Write Data to the Disk on Close,” on page 315.

- Transaction tracking support for user applications (NetWare only)

Monitors application transactions and rolls back uncompleted application transactions after a system failure. For information, see Section 23.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare),” on page 317.

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

  For some file types, you can back up files from snapshots of data so that all files, including open ones, can be backed up. For information, see Section 23.3, “Using the File-Level Snapshot Attribute to Enable the Backup of Open Files,” on page 318.

- Salvage file subsystem

  Recovers 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 361.

1.4.5 Maximize Available Space

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

- File compression

  Compresses inactive files, according to preset parameters, to conserve space in a volume. For information, see “Managing Compression on NSS Volumes” on page 325.

- Volume space restrictions

  Limits the amount of space a volume can consume in its pool. For information, see Section 25.2, “Managing NSS Volume Quotas,” on page 346.

- Directory space restrictions

  Limits the amount of space a subdirectory can consume, regardless of broader volume and user constraints. For information, see Section 25.3, “Managing Directory Quotas,” on page 348.

- User space restrictions

  Limits the amount of space a user’s data can consume, regardless of broader directory or volume constraints. For information, see Section 25.4, “Managing User Space Quotas,” on page 353.

1.5 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 POSIX File Systems plus NCP Server	NetWare Traditional File System
Comparison of NSS on NetWare and NSS on Linux (page 579)	✓	✓		
Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems (page 595)		✓	✓	
Comparison of NSS on NetWare and the NetWare Traditional File System (page 585)	✓			✓

1.6 What's Next

See Chapter 2, "What's New," on page 37 to learn about new and modified features in this release of NSS.

Review the following sections to help you plan your storage solution:

- For information about installing and configuring NSS on your server, see Chapter 3, "Installing and Configuring Novell Storage Services," on page 45.
- For guidelines and instructions about upgrading the media format of NSS volumes to use hard links, see Chapter 4, "Upgrading the NSS Media Format," on page 55.
- For guidelines and instructions about migrating NSS volumes from NetWare servers to OES 2 Linux servers, see Chapter 11, "Migrating NSS Devices from NetWare to OES 2 Linux," on page 141.
- For guidelines about setting up NSS volumes and services on a virtual server, see Chapter 6, "Using NSS in a Virtualization Environment," on page 79.
- For management tools overviews and quick references, see Chapter 9, "Management Tools for NSS," on page 97.
- For information to help you plan storage services to use the NSS file system and services, see Chapter 5, "Planning NSS Storage Solutions," on page 63.

What's New

2

This section describes enhancements and changes to the Novell® Storage Services™ file system and services for Novell Open Enterprise Server 2 (OES 2).

- Section 2.1, "What's New for NSS (OES 2 SP1 and NetWare 6.5 SP8)," on page 37
- Section 2.2, "What's New for NSS (OES 2 and NetWare 6.5 SP7)," on page 39

2.1 What's New for NSS (OES 2 SP1 and NetWare 6.5 SP8)

The features described in this section were added or modified for NSS in OES 2 SP 1 Linux and OES 2 SP1 NetWare® (same as NetWare 6.5 SP8) since the NSS release for OES 2 Linux and OES 2 NetWare (same as NetWare 6.5 SP 7).

- Section 2.1.1, "Linux and NetWare," on page 37
- Section 2.1.2, "Linux," on page 37
- Section 2.1.3, "NetWare," on page 39

2.1.1 Linux and NetWare

The enhancements and changes described in this section are available on both OES 2 SP1 Linux and OES 2 SP1 NetWare (same as NetWare 6.5 SP8).

- "noatime and atime Options" on page 37
- "eDirectory 8.8.4" on page 37

noatime and atime Options

This NSS command option allows the administrator to control whether access times are updated when files and directories are read. Using `/noatime` is useful for backup, news servers, and mail servers where the extra disk activity associated with updating the access time is not desired. Avoiding the writes associated with updating the access time can result in measurable performance gains. For information, see Section A.25, "noatime and atime Commands," on page 478.

eDirectory 8.8.4

NSS supports eDirectory™ 8.8.4 in the OES 2 SP1 release.

2.1.2 Linux

In addition to bug fixes, the following enhancements and changes were made for NSS on OES 2 SP1 Linux.

- "Long Is the Default Name Space for Mounting NSS Volumes on Linux" on page 38
- "PosixPermissionMask Option" on page 38
- "UnplugAlways Option" on page 38

- "Support for 64-Bit eDirectory, Novell CIFS, and Novell AFP" on page 38
- "Support for Novell AFP for Linux" on page 38
- "Support for Novell CIFS for Linux" on page 38
- "Support for Novell Domain Services for Windows" on page 39

Long Is the Default Name Space for Mounting NSS Volumes on Linux

Beginning with OES 2 SP1 Linux, the default name space that NSS uses when mounting NSS volumes on Linux has changed from UNIX to Long. This matches the name space handling of NSS volumes on NetWare. Using the Long name space as primary improves performance over using the UNIX name space, especially if you expect to store millions of files on the volume. To mount a volume with a different name space (that is, DOS, Macintosh, or UNIX), you must explicitly specify the name space to use.

PosixPermissionMask Option

The `/PosixPermissionMask=`*`mask`* command option allows the administrator to control which POSIX* permission bits can be set. This is useful if you need to modify the default setting to support a different value for the POSIX Other field that is required by Linux utitilites such as SSH. For information, see Section A.32, "POSIX Permission Mask Command (Linux)," on page 485.

UnplugAlways Option

The `/(No)UnplugAlways` command option allows NSS to unplug the device queue after queuing each read. It is enabled by default for OES 2 SP1. This improves performance significantly on certain workloads. For information, see Section A.5.5, "UnplugAlways Command for the Read Queue," on page 457.

Support for 64-Bit eDirectory, Novell CIFS, and Novell AFP

Support was added for 64-bit eDirectory 8.8.4, Novell CIFS, and Novell AFP.

Support for Novell AFP for Linux

Novell AFP for Linux is available for accessing files on NSS volumes on OES 2 SP1 Linux. It works similarly to Novell AFP for NetWare. Linux-enabling users with Linux User Management is not required.

For information about installing and managing Novell AFP for Linux, see the *OES 2 SP1: Novell AFP For Linux Administration Guide*.

Support for Novell CIFS for Linux

Novell CIFS for Linux is available for accessing files on NSS volumes on OES 2 SP1 Linux. It works similarly to Novell CIFS for NetWare. Linux-enabling users with Linux User Management is not required.

For information about installing and managing Novell CIFS for Linux, see the *OES 2 SP1: Novell CIFS for Linux Administration Guide*.

Support for Novell Domain Services for Windows

Novell Domain Services for Windows is available for accessing files on NSS volumes with the Samba/CIFS protocol. For guidelines, see Section 5.6.4, "Novell Domain Services for Windows," on page 76.

For information about installing and managing Domain Services for Windows, see the *OES 2 SP1: Domain Services for Windows Administration Guide*.

2.1.3 NetWare

There were bug fixes for NSS on OES 2 SP1 NetWare (same as NetWare 6.5 SP8).

2.2 What's New for NSS (OES 2 and NetWare 6.5 SP7)

The features described in this section were added or modified for NSS in OES 2 and NetWare 6.5 SP8 since the NSS release for Novell Open Enterprise Server 1 (OES 1) Linux and OES 1 NetWare (same as NetWare 6.5 SP 3).

- Section 2.2.1, "Linux and NetWare," on page 39
- Section 2.2.2, "Linux," on page 41
- Section 2.2.3, "NetWare," on page 43

2.2.1 Linux and NetWare

The enhancements and changes described in this section are available on both OES 2 Linux and OES 2 NetWare (same as NetWare 6.5 SP7).

- "Media Upgrade for Enhanced Hard Links Support" on page 39
- "Novell Distributed File Services Plug-In" on page 40
- "Files and Folders Plug-in" on page 40
- "Storage Plug-In" on page 40
- "User Quota Task in the Storage Plug-In" on page 41

Media Upgrade for Enhanced Hard Links Support

Enhanced hard links support was added for the following releases (and later versions):

- OES 2 Linux and NetWare
- OES 1 SP 1 NetWare
- NetWare 6.5 SP4

A new metadata structure allows up to 65,535 hard links per file. A media upgrade is required to take advantage of this feature. For information about the media upgrade, see Chapter 4, "Upgrading the NSS Media Format," on page 55.

For information about managing hard links, see Chapter 27, "Managing Hard Links," on page 375.

Novell Distributed File Services Plug-In

Novell Distributed File Services (DFS) can now be managed fully within the Distributed File Services plug-in for Novell iManager 2.7. The plug-in is included in the `nssmgmt.nlm` file. You can do the following with the new DFS plug-in:

- Create a DFS management context.
- Manage, monitor, and repair the Volume Location Database (VLDB) services, replica sites, and the VLDB.
- Create, modify, or delete DFS junctions for NSS volumes.
- Configure trustees, trustee rights, and inherited rights filters for junctions and junction target locations.
- Manage move volume or split volume jobs.

For information about DFS for NSS volumes, see the *OES 2: Novell Distributed File Services Administration Guide*.

Files and Folders Plug-in

The Files and Folders plug-in (`filemanager.npm`) for Novell iManager 2.7 provides the following capabilities for NSS volumes and NCP™ (NetWare Core Protocol™) volumes:

- Manage trustees, trustee rights, and inherited rights filters for files and directories
- View the inherited rights at every level in the file system tree for a given trustee of a selected file or directory
- View the effective rights for a given trustee of a selected file or directory
- Manage file attributes
- Configure directory quotas (only for NSS volumes where the volume's Directory Quotas attribute is enabled)
- Browse directories and files
- Create directories
- Upload and download files
- Delete files
- Salvage and purge deleted files (only for NSS volumes where the volume's Salvage attribute is enabled)

For an overview of these features for the Files and Folders plug-in, see Section 9.1.6, "Files and Folders Plug-In Quick Reference," on page 110.

Storage Plug-In

The Storage plug-in to Novell iManager 2.7 requires the `nssmgmt.npm` file and the `storagement.npm` file. For information, see Section 9.1.1, "Understanding Storage-Related Plug-Ins," on page 98.

User Quota Task in the Storage Plug-In

The User Quota option in the Storage plug-in for iManager 2.7 now allows users to access information via iManager about their personal user quotas on NSS volumes. For information, see Section 25.4.2, “Viewing User Space Quotas,” on page 354.

2.2.2 Linux

The enhancements and changes described in this section were made for NSS on OES 2 Linux.

- “64-Bit Support” on page 41
- “New Location for the NSS Configuration Files” on page 41
- “Encrypted Volumes” on page 41
- “Extended Attributes (XAttr)” on page 42
- “Group I/O Management” on page 42
- “Enhanced Hard Links” on page 42
- “High Memory Management” on page 42
- “LAF Audit Log Messages” on page 42
- “Metadata Migration (METAMIG) Utility” on page 42
- “NOATIME and NODIRATIME Support” on page 42
- “Novell Archive and Version Services” on page 42
- “Novell Distributed File Services” on page 43
- “Pool Snapshots” on page 43
- “Software RAIDs” on page 43
- “User Quotas” on page 43
- “User Quotas and Linux User Management” on page 43

64-Bit Support

Selecting NSS as part of a 64-bit installation on OES 2 Linux installs 64-bit NSS.

New Location for the NSS Configuration Files

In this release, NSS configuration files have moved from the `/opt/novell/nss/conf` directory to the `/etc/opt/novell/nss` directory.

Encrypted Volumes

Encrypted NSS volume support was added for OES SP1 Linux and later. Previously, encrypted volumes were supported only on NetWare. For information, see Chapter 20, “Managing Encrypted NSS Volumes,” on page 281.

Extended Attributes (XAttr)

The Extended Attributes (XAttr) extension for NSS provides accessibility into many extended attributes for NSS on Linux. It allows you to read, back up, and restore extended attributes of files on NSS. For information, see Section A.11, “Extended Attributes (XAttr) Commands (Linux),” on page 464.

Group I/O Management

In OES 2 Linux, NSS performs group writes in three categories: journal, metadata, and user data. By setting policies for group writes, you can improve the performance of the file system for your particular environment. For information, see Section 31.3, “Configuring or Tuning Group I/O,” on page 415.

Enhanced Hard Links

Enhanced support for hard links is now available for NSS volumes on OES 2 Linux. For information about configuring and managing hard links for NSS volumes on Linux or NetWare, see Chapter 27, “Managing Hard Links,” on page 375.

High Memory Management

NSS for Linux supports metadata caching in high memory for servers that are low-memory constrained. This can greatly improve metadata reads on 32-bit systems with large numbers of files. For information, see Section 31.2, “Configuring High Memory Cache,” on page 414.

LAF Audit Log Messages

NSS for OES 2 Linux supports Lightweight Auditing Format (LAF) for NSS trustee changes for NSS volumes on OES 2 Linux. When it is enabled, NSS reports events for adding or removing trustees and setting the inherited rights filter. For information, see Section 22.4, “Enabling or Disabling LAF Audit Log Messages for Trustee Events (Linux),” on page 307.

Metadata Migration (METAMIG) Utility

The Metadata Migration (METAMIG, `metamig`) utility is available for OES 1 Linux and later. For information, see Section B.8, “METAMIG (Linux),” on page 514.

NOATIME and NODIRATIME Support

NSS on OES 2 Linux supports the `O_NOATIME` option for the Linux `open(2)` command, and the `noatime` and `nodiratime` options for the `mount` and `nfsmount` command and the `/etc/fstab` file. All of these options have the same objective—that is, to prevent the access time from being updated unless the access involves a modification of a file’s or directory’s metadata or content. For information, see Section A.26, “noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux),” on page 479.

Novell Archive and Version Services

Novell Archive and Version Services is now supported for NSS volumes on OES 2 Linux. For information, see the *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*.

Novell Distributed File Services

Novell Distributed File Services (DFS) is now supported for NSS volumes on OES 2 Linux. For information about using DFS on Linux or NetWare, see the *OES 2: Novell Distributed File Services Administration Guide*.

Pool Snapshots

NSS now supports pool snapshots for NSS pools on OES 2 Linux. Pool snapshots were previously available only on NetWare. You can use NSSMU or the Storage plug-in for Novell iManager 2.7 to manage pool snapshots for Linux. Use iManager to manage pool snapshots for NSS on NetWare.

For information about creating, managing, or deleting snapshots on Linux or NetWare, see Chapter 18, “Managing NSS Pool Snapshots,” on page 229.

Software RAIDs

Support for NSS software RAIDs 5 and 15 was added for OES SP1 Linux and later. Previously, software RAIDs 5 and 15 were supported only on NetWare. For information, see Chapter 13, “Managing NSS Software RAID Devices,” on page 165.

User Quotas

User space quotas support was added for OES SP1 Linux and later. Previously, user quotas were supported only on NetWare. For information, see Section 25.4, “Managing User Space Quotas,” on page 353.

User Quotas and Linux User Management

Changes in NCP Server for Linux now make it possible for the file ownership to be tracked for users of NSS volumes on Linux without requiring the user to be Linux-enabled with Linux User Management.

2.2.3 NetWare

The enhancements and changes described in this section were made for NSS on OES 2 NetWare (same as NetWare 6.5 SP7).

- “Extended System Memory Management” on page 43

Extended System Memory Management

In OES 2 NetWare and NetWare 6.5 SP7, commands have been added to allow you to manage extended system memory, the memory above 4 GB. For information, see Section 30.7, “Configuring NSS Extended System Memory,” on page 411.

Installing and Configuring Novell Storage Services

3

This section describes how to install and configure Novell® Storage Services™ on Novell Open Enterprise Server 2.

- Section 3.1, “Requirements for Installing NSS,” on page 45
- Section 3.2, “Installing and Configuring NSS on NetWare,” on page 47
- Section 3.3, “Installing and Configuring NSS on Linux,” on page 48
- Section 3.4, “Upgrading the Media Format for Hard Link Support,” on page 51
- Section 3.5, “Enabling Users for Linux Utilities and Services,” on page 51
- Section 3.6, “Updating NSS on OES 2 Linux,” on page 52
- Section 3.7, “Upgrading from OES 1 Linux to OES 2 Linux,” on page 53
- Section 3.8, “Updating the NSS Kernel Module (km_nss) for SUSE Linux Enterprise Server,” on page 54

3.1 Requirements for Installing NSS

Make sure your system and storage solution meets the requirements in this section.

- Section 3.1.1, “Device Requirements for NetWare and Linux,” on page 45
- Section 3.1.2, “Requirements for NSS on NetWare,” on page 46
- Section 3.1.3, “Requirements for NSS on Linux,” on page 46
- Section 3.1.4, “Requirements for Storage-Related iManager Plug-Ins,” on page 47

3.1.1 Device Requirements for NetWare and Linux

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

- ❑ Devices can be up to 2 TB in size. For information, see Section 10.1.1, “Device Size Limit,” on page 127.
- ❑ At least 10 MB of free space is needed on the storage media for each NSS pool you plan to create.
- ❑ At least 12 MB of free space is needed on the storage media for each software RAID segment you plan to create.
- ❑ For information about devices to use in a virtual environment, see Chapter 6, “Using NSS in a Virtualization Environment,” on page 79.

3.1.2 Requirements for NSS on NetWare

The following additional requirements apply to NSS on NetWare:

- ❑ A physical server or virtual server running OES 2 NetWare, or a physical server running NetWare 6.5 SP7.
- ❑ The NSS file system is the default file system. The NetWare operating system must be installed on an NSS volume.

3.1.3 Requirements for NSS on Linux

- "General Requirements" on page 46
- "EVMS Requirements" on page 46
- "Non-EVMS Volume Managers" on page 47

General Requirements

The following general requirements apply to NSS on Linux:

- ❑ A physical server or virtual server running OES 2 Linux.
- ❑ NSS is not installed by default. You can select it during the YaST install, or install it at any time from the *YaST > Open Enterprise Server > OES Install and Configuration*.

  For information about install options, see Section 3.3, "Installing and Configuring NSS on Linux," on page 48.
- ❑ The NSS file system is used only for data volumes on OES 2 Linux. The Linux operating system requires a Linux POSIX file system for its system volume, such as Ext3.
- ❑ After installing OES 2 Linux, install only approved updates. Refer to the *OES2 SP1: Linux Installation Guide* to install the approved updates.

  For information about updating the NSS Kernel Module, see Section 3.8, "Updating the NSS Kernel Module (km_nss) for SUSE Linux Enterprise Server," on page 54.
- ❑ 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.

EVMS Requirements

The device where you want to create NSS volumes must be managed by EVMS in order to use the Storage plug-in to iManager or NSSMU to create and manage NSS partitions, pools, and volumes. Novell Cluster Services™ for Linux also requires EVMS.

IMPORTANT: If you move devices that contain NSS pools cross-platform (such as reassigning SAN-based devices from a NetWare server to an OES 2 Linux server), NSS recognizes the pools and automatically uses EVMS to manage the devices.

Updates to EVMS for OES 2 SP1 Linux, are received through the update channel for SUSE Linux Enterprise Server 10 SP 2. Make sure that you install the latest patches for EVMS before you create pools and volumes for this server.

WARNING: NSS requires EVMS version 2.5.5-24.54.5 or later. This mandatory update is available in the SUSE Linux Enterprise Server 10 Support Pack 1 patch channel as of March 21, 2008. OES 2 SP1 Linux automatically installs or updates EVMS to EVMS version 2.5.5-24.62.

In some cases, reverting to EVMS versions earlier than 54.5 might result in data corruption. In the .46 and .49 versions, data corruption can occur if NSS pools and volumes are mounted when you install EVMS updates. For details and workarounds, see the following issues in the *OES 2 Readme* (http://www.novell.com/documentation/oes2/pdfdoc/readme-20/readme-20.pdf):

- Possible Data Corruption Can Occur If NSS Pools Are Mounted During an EVMS Update
- EVMS Updates Are Mandatory to Prevent Data Corruption for NSS Pools and Volumes

Linux uses Linux Volume Manager (LVM) as its default volume manager. Beginning with the Linux 2.6 kernel, any disk managed by the LVM2 cannot be managed by Enterprise Volume Management System (EVMS). Thus, disks where the boot partition (such as `/boot` for Grub) and system partition (such as for the swap and / (`root`) 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" in the *OES2 SP1: Linux Installation Guide*.

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.

Non-EVMS Volume Managers

NSS is not supported for devices managed by non-EVMS volume managers. For more information, see Chapter E, "Using NSS on Devices Managed by Non-EVMS Volume Managers (Linux)," on page 569.

3.1.4 Requirements for Storage-Related iManager Plug-Ins

For information about installing and using the storage-related iManager plug-ins, see Section 9.1, "Novell iManager and Storage-Related Plug-Ins," on page 97.

3.2 Installing and Configuring NSS on NetWare

NSS is the default file system for NetWare. It is installed and enabled by default. No action is required.

For information about installing or upgrading the NetWare operating system, see the *OES 2 SP1: NetWare Installation Guide*.

3.3 Installing and Configuring NSS on Linux

This section describes only those steps in the install that are directly related to installing Novell Storage Services and its dependencies. For information about installing OES 2 services on Linux, see the *OES2 SP1: Linux Installation Guide*.

- Section 3.3.1, "Installing Linux with EVMS as the Volume Manager of the System Device," on page 48
- Section 3.3.2, "Selecting the NSS Pattern During the OES 2 Linux Install," on page 48
- Section 3.3.3, "Installing NSS on an Existing OES 2 Linux Server," on page 50
- Section 3.3.4, "Enabling or Disabling NSS on Linux," on page 51

3.3.1 Installing Linux with EVMS as the Volume Manager of the System Device

If you are installing OES 2 Linux on a system with a single device where you also plan to create NSS volumes, that device must be managed by the Enterprise Volume Management System (EVMS) volume manager instead of the default Linux Volume Manager 2 (LVM2) volume manager. The volume manager for the system device is configured in the Partitioning section of the Installation Settings page for the YaST install. For more information and instructions, see "Installing Linux with EVMS as the Volume Manager of the System Device" in the *OES2 SP1: Linux Installation Guide*.

3.3.2 Selecting the NSS Pattern During the OES 2 Linux Install

1. In the YaST install, on the *Installations Settings* page, click *Software* to go to the *Software Selections and System Tasks* page.

   For information about the entire OES 2 Linux install process, see the *OES2 SP1: Linux Installation Guide*.

2. From the OES Services options, select *Novell Storage Services*. Selecting NSS as part of a 64-bit installation automatically installs NSS 64-bit support.

   The following additional OES 2 services are automatically selected:

   - Novell Backup / Storage Management Services™

     SMS makes it possible to back up trustee and other extended attributes for data on NSS volumes. It is also used by Novell Distributed File Services for moving or splitting NSS volumes.

   - Novell eDirectory

     eDirectory supports authentication of users who connect to NSS volumes.

   - Novell Linux User Management

     LUM allows eDirectory users to be enabled for Linux services, such as access via CIFS/Samba, FTP, and so on. The administrator user for the server is automatically Linux-enabled with LUM. Non-NCP users must be Linux-enabled with LUM in order to access data on NSS volumes with Linux services, utilities, or protocols, such as CIFS/Samba, NFS, or AFP (requires a third-party solution).

IMPORTANT: LUM is required even if the administrator user is the only LUM user on the server.

- NCP Server / Dynamic Storage Technology

  NCP Server provides support to NSS for access control, shadow volumes, commands, and file access. It is required even if you are not using NCP clients to connect to the NSS volume.

- Novell Remote Manager (NRM)

  NRM provides the NCP Server management plug-in that allows you to create shadow volumes using NSS volumes on Linux. You can also use it to manage NCP connections to the NSS volumes.

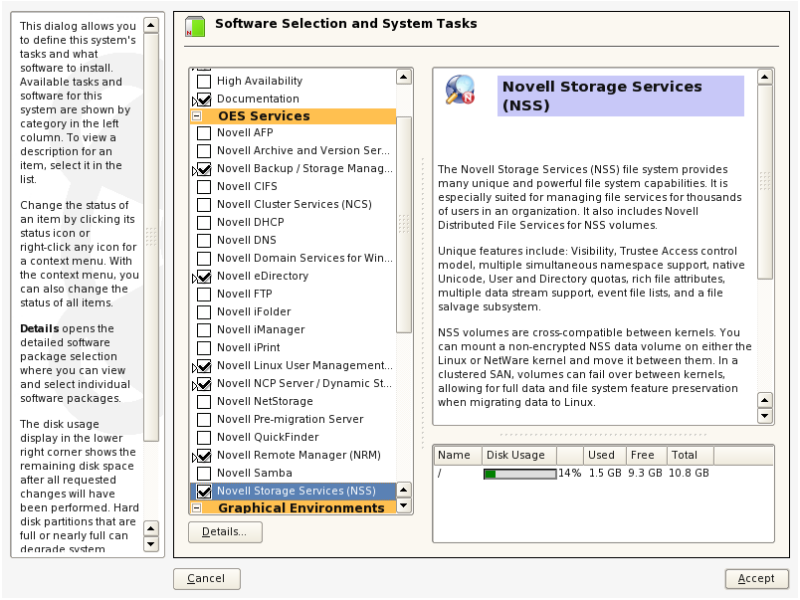

3 Optionally select *Novell iManager* to be installed on the server.

You must install iManager somewhere in the same tree as the server. If you install iManager and NSS on the same server, the storage-related plug-ins are automatically installed.

If you install iManager on a different server, make sure you install the storage-related plug-ins that you need to manage NSS file system and services. For information about installing storage-related plug-ins on an existing server, see Section 9.1, “Novell iManager and Storage-Related Plug-Ins,” on page 97.

4 Optionally select non-NCP file access services to be installed on the server.

NSS requires NCP Server to be installed and running on the server even if you select one or more of these alternate methods for user access.

- **Novell AFP:** Allows Macintosh users to connect to NSS volumes with the AFP (Apple Filing Protocol). Novell AFP is available beginning in OES 2 SP1 Linux. For information about configuring and managing AFP, see the *OES 2 SP1: Novell AFP For Linux Administration Guide*.

- **Novell CIFS:** Allows CIFS/Samba users to connect to NSS volumes with the CIFS/Samba protocol. Novell CIFS is available beginning in OES 2 SP1 Linux. For information about configuring and managing Novell CIFS, see the *OES 2 SP1: Novell CIFS for Linux Administration Guide*.
- **Novell Samba:** Allows CIFS/Samba users to connect to NSS volumes with the CIFS/Samba protocol. This service is based on Linux Samba and requires users to be Linux-enabled with Linux User Management. For information about configuring Samba during the install and configuring users for CIFS/Samba access after the install, see the *OES2 SP1: Samba Administration Guide*.

  IMPORTANT: Novell Samba and Novell CIFS are different file access services that allow CIFS/Samba users to connect to NSS volumes. You can select only one of the two on a given server because of port contention issues.

  ---

5. Optionally select Novell Cluster Services to be installed on the server.

   Install NCS if you plan to share NSS pools in a cluster. For information about installing NCS and configuring shared devices and pools using NCS, see the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.
6. Click *Accept* to return to the *Installation Settings* page.

   Licensing dialog boxes might open where you are prompted to accept proprietary modules being installed.
7. Continue with the OES 2 Linux installation.
8. After the install, use the Software Updater (or other update methods) to install any NSS patches from the OES 2 Linux patch channel and any EVMS patches from the SUSE® Linux Enterprise Server 10 SP2 patch channel.

3.3.3 Installing NSS on an Existing OES 2 Linux Server

If you did not install Novell Storage Services during the OES 2 Linux installation, you can install it later by using *YaST > Open Enterprise Server > OES Install and Configuration*.

1. Log in to the server as the `root` user.
2. In YaST, select *Open Enterprise Server > OES Install and Configuration*.
3. In the Selection window under *OES Services*, click *Novell Storage Services* and any other OES components that you want to install.

   Follow the instruction for selecting NSS and its dependencies described in Section 3.3.2, “Selecting the NSS Pattern During the OES 2 Linux Install,” on page 48.
4. Click *Accept* to begin the install, then click *Continue* to accept changed packages.
5. Follow the on-screen instructions to complete the install.
6. After the install, enter `rcnovell-smdrd restart` at the command prompt, or reboot the server before performing any backups, restores, or server consolidations on the NSS file system.
7. Use the Software Updater (or other update methods) to install patches from the OES 2 Linux patch channel and the SUSE Linux Enterprise Server 10 SP2 patch channel (such as for EVMS).

3.3.4 Enabling or Disabling NSS on Linux

When you install NSS during the initial install, NSS and its dependencies are automatically enabled in the Linux System Services (Runlevel) with Runlevels 2, 3, and 5. NSS is not automatically enabled if you post-install NSS on the server.

Although you can uninstall OES 2 Linux service RPMs through YaST, we do not recommend it because so many modules have interdependencies. Uninstalling services can leave the server in an undesirable state. If you no longer plan to use a server, we recommend disabling the service.

WARNING: NSS must be enabled to use any components or tools for the NSS file system.

To enable or disable NSS on Linux:

1. Log in to the server as the `root` user, then start YaST.
2. Click *System* > *System Services (Runlevel)*, then click *Expert Mode*.
3. Select NSS, then click *Set/Reset*.
4. Select one of the following options from the *Set/Reset* menu:
   - *Disable the service*
   - *Enable the service*
5. Click *Finish* to save and apply your changes, then exit the YaST Runlevel tool.

3.4 Upgrading the Media Format for Hard Link Support

The enhanced hard links support for NSS volumes requires a media format upgrade after you have installed or upgraded the operating system. For guidelines and media upgrade instructions, see Chapter 4, “Upgrading the NSS Media Format,” on page 55.

3.5 Enabling Users for Linux Utilities and Services

On Linux, eDirectory users must also have a Linux identity in order to access NSS volumes via Linux services and utilities such as Samba, SSH, and FTP. OES 2 Linux provides the Linux User Management (LUM) technology that creates the local Linux user identity and stores the UID for the user in eDirectory. The Administrator user for the server is automatically Linux-enabled with LUM and added to a LUM administrator group for the server as part of the installation process. Before users of the NSS volumes can access NSS volumes on Linux with Linux services and utilities, you must enable both the service and the users with LUM.

For information, about how to enable users and Linux services with LUM, see the *OES 2 SP1: Novell Linux User Management Technology Guide*. For more information about why LUM is necessary for Linux services and utilities, see Section 5.5, “Access Control for NSS on Linux,” on page 71.

3.6 Updating NSS on OES 2 Linux

You can get NSS patches in the OES 2 Linux update channel or from the Novell Download Web site (http://download.novell.com). For information about updating OES 2 Linux services, see "Updating an OES 2 SP1 Linux Server" in the *OES2 SP1: Linux Installation Guide*.

Consider the following issues when updating NSS on OES 2 Linux:

- Section 3.6.1, "Parameter Settings," on page 52
- Section 3.6.2, "Reboot Server or Restart jstcpd, adminusd, and volmnd," on page 52
- Section 3.6.3, "Storage-Related Plug-Ins," on page 52

3.6.1 Parameter Settings

When you update an OES 2 Linux server with a Support Pack or apply NSS patches, all NSS-related parameter settings remain the same as they were before the update or patch. For example, server-level, pool, and volume settings are not modified.

3.6.2 Reboot Server or Restart jstcpd, adminusd, and volmnd

If you do not reboot the server as part of the update or patch process, some NSS functions and tools might not work properly until you restart the `jstcpd`, `adminusd`, and `volmnd` daemons.

For example, NSSMU or the Novell Distributed File Services (DFS) volume location database might hang when you create a volume. DFS is delivered and updated as a part of the NSS package on OES 2 Linux (and later versions) servers. If the server is a VLDB replica site, the vldb might not work properly or cause hangs when creating new NSS volumes. For information, see "DFS may not function properly after upgrading NSS on OES 2 Linux and later." in the *OES 2: Novell Distributed File Services Administration Guide*.

To ensure that NSS and DFS is functioning properly after updating OES 2 Linux to a support pack or applying NSS patches:

1. Log in as the `root` user, then open a terminal console.
2. At the terminal console prompt, restart the following daemons in the order indicated:

   ```
   /opt/novell/nss/sbin/jstcpd
   /opt/novell/nss/sbin/adminusd
   /opt/novell/nss/sbin/volmnd
   ```

3. If the server is a VLDB replica site for a Novell Distributed File Services management context, restart the VLDB by entering

   ```
   vldb stop service

   vldb start service
   ```

3.6.3 Storage-Related Plug-Ins

The File Protocols plug-in for the Native File Access Protocols for NetWare service has been replaced in OES 2 SP1 by three plug-ins: Novell AFP (`afpmgmt.npm`), Novell CIFS (`cifsmgmt.npm`), and NFS for NetWare (`nfsmgmt.npm`). The Novell AFP and Novell CIFS plug-ins support AFP and CIFS services for NSS volumes on both Linux and NetWare.

The three protocols' plug-ins also require the NSS Management (`nssmgmt.npm`) and Storage Management (`storagemgmt.npm`) plug-ins. Other storage-related plug-ins include Archive Versioning (`arkmgmt.npm`), Cluster Services (`ncsmgmt.npm`), and Distributed File Services (`dfsmgmt.npm`). All storage-related plug-ins share code in common with the Storage Management plug-in.

You must uninstall the existing storage-related plug-ins, then install the new plug-ins at the same time to make sure that the common code works for all plug-ins.

3.7 Upgrading from OES 1 Linux to OES 2 Linux

Consider the issues in this section when upgrading from OES 1 Linux (with the latest support packs and patches) to OES 2 Linux:

- Section 3.7.1, "Parameter Settings," on page 53
- Section 3.7.2, "Read Ahead Blocks Setting," on page 53

3.7.1 Parameter Settings

When you upgrade the server from OES 1 Linux to OES 2 Linux, all NSS-related parameter settings remain the same as they were before the upgrade. For example, server-level, pool, and volume settings are not modified. All future pools and volumes are created with the default settings for OES 2 Linux.

3.7.2 Read Ahead Blocks Setting

The Read Ahead Blocks parameter specifies the number of data blocks that NSS reads ahead for any open file on which read operations are ongoing in the specified server. Its default setting was changed from 2 blocks in OES 1 Linux to 16 blocks on OES 2 Linux and in a patch on OES 1 SP2 Linux. Volumes created after the patch was installed on OES 1 SP2 Linux have a default setting of 16; however, older volumes might be set to the original default, or to manual settings under 16. After upgrading to OES 2 Linux, if you experience read performance problems with an NSS volume, check the volume's Read Ahead Blocks parameter setting and change it to at least 16 blocks. A read-ahead value of 2 is too low in most OES 2 Linux environments.

On Linux, you can view or modify the Read Ahead Blocks value by using NSSMU as follows:

1. In NSSMU, select *Volumes* to view a list of volumes.
2. Select the volume, then press *Enter* to view its *Volume Properties* list.
3. Press the arrow keys to go to the *Read Ahead Count in Blocks* parameter, then press *Enter* to access the setting.
4. Type the new count value, then press *Enter*.

   Valid values are 0 to 1024 blocks, where a block is 4 KB. However, a count of 128 is the practical maximum value. 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.
5. Press the arrow keys to go to *Apply*, then press *Enter* to save and apply the new setting.

You can also modify the value by using the Read Ahead Blocks switches in the NSS Console. For information, see Section A.34, “Read Ahead Blocks and Allocate Ahead Blocks Commands,” on page 486.

3.8 Updating the NSS Kernel Module (km_nss) for SUSE Linux Enterprise Server

The NSS kernel module (`km_nss`) file is delivered in SUSE Linux Enterprise Server. The NSS kernel module is installed by default because it is part of the Linux kernel.

The NSS kernel module is delivered in Linux because of the following:

- The NSS kernel module provides NSS support for Linux utilities and commands.
- The NSS kernel module is open source to meet the GPL (GNU Public License) requirements.
- If you add a kernel patch for a third-party module, application, or service, the NSS source code must be in the kernel code in order to be recompiled with the patch. Otherwise, an OES 2 Linux deployment using NSS would break.

Although the basic NSS code for Linux is delivered in the Linux kernel, it is not usable without OES 2 services on the Linux server. 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 Novell eDirectory to authenticate non-`root` users for connections to the volume. In addition, OES 2 Linux provides Linux User Management, NetWare Core Protocol™ (NCP) Server, and volume and user space management tools that make NSS volumes usable on a Linux server.

Updates for the `km_nss` module are delivered in major releases of SUSE Linux Enterprise Server and as patches for the Linux kernel in the Software Updater. You can also download patches for the `km_nss` module from the *Novell SUSE Patch Support Database* (PSDB) (http://support.novell.com/linux/psdb/bydate.html).

Upgrading the NSS Media Format

4

An enhanced Novell® Storage Services™ (NSS) media format is available that provides improved support for hard links. After you install or upgrade your operating system to Novell Open Enterprise Server 2, you can decide whether to upgrade the media format for your NSS volumes to use the new metadata structure; some restrictions apply.

- Section 4.1, “Guidelines for Upgrading the Media Format of NSS Volumes,” on page 55
- Section 4.2, “Upgrading the Media Format for Existing NSS Volumes,” on page 57
- Section 4.3, “Upgrading the Media Format Automatically for New NSS Volumes,” on page 60

4.1 Guidelines for Upgrading the Media Format of NSS Volumes

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

- Section 4.1.1, “Cross-Platform Support for the NSS Media Upgrade,” on page 55
- Section 4.1.2, “Which NSS Volumes to Upgrade,” on page 55
- Section 4.1.3, “Before Upgrading the Media Format,” on page 56
- Section 4.1.4, “After Upgrading the Media Format,” on page 57
- Section 4.1.5, “If You Do Not Upgrade the Media Format,” on page 57

4.1.1 Cross-Platform Support for the NSS Media Upgrade

The NSS media upgrade for enhanced hard links support is available for the following operating platforms (and later versions):

- Novell Open Enterprise Server 2 Linux and NetWare
- Novell Open Enterprise Server 1 SP1 NetWare
- NetWare® 6.5 SP4

If the NSS volume is used in a cluster with Novell Cluster Services™, all nodes in the cluster must be upgraded to a supported platform before you upgrade the media format for any shared volumes.

The media upgrade is not available on OES 1 Linux, NetWare 6.5 SP3, and earlier versions. In a mixed-version cluster, the version of NSS must be able to be loaded on any of the platforms. After you upgrade the media format on an NSS volume, it cannot be mounted on an unsupported platform.

4.1.2 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 a supported platform.

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 your system is performing as expected after upgrading the operating system, and you might need to roll back to an earlier release.
- You plan to migrate one or more devices containing the NSS volume to an unsupported platform.
- You need to share this volume with a mixed cluster with Novell Cluster Services where there are some unsupported platforms in the mix. The cluster software prevents the media upgrade unless all operating systems in the cluster support the new media format.

4.1.3 Before Upgrading the Media Format

- "Opportunity to Roll Back Before the Media Upgrade" on page 56
- "Hard Link Behavior without a Media Upgrade" on page 56
- "Clusters and the Media Upgrade" on page 56

Opportunity to Roll Back Before the Media Upgrade

When you upgrade the operating system, NSS does not automatically upgrade the media format to use the new metadata structure. This allows you to the opportunity to roll back to the previous release if necessary. Before you upgrade the media format of NSS volumes to the new data structure, make sure the server is performing as expected.

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

Hard Link Behavior without a Media Upgrade

Until you upgrade the media format for enhanced hard link support, any existing hard links on the NSS volume are visible, and they can be opened, closed, read, and written. However, 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.

Clusters and the Media Upgrade

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 a supported platform before attempting to upgrade the shared volume.

4.1.4 After Upgrading the Media Format

After you upgrade the volume to use the media format with enhanced hard links, the following constraints apply for its use:

- The Hard Links attribute must be enabled for the upgraded NSS volume 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 27.3, “Enabling or Disabling the Hard Links Attribute,” on page 380.
- The upgraded NSS volume cannot be rolled back to use the old media format.
- You cannot roll back the operating system to a previous version.
- You cannot migrate a device containing the upgraded volume to a system with an unsupported operating system.
- Only nodes that have a supported operating system can be added to a cluster where shared volumes use the upgraded media format.

4.1.5 If You Do Not Upgrade the Media Format

If you do not upgrade the media format for an NSS volume, the volume’s format uses the same metadata structure as is used on earlier releases. 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.

It is possible to move devices that contain non-upgraded NSS volumes cross-platform to servers with operating systems compatible with the old media format. For information about moving media cross-platform, see Section 11.2, “Moving Non-Clustered Devices From NetWare Servers to OES 2 Linux Servers,” on page 142.

4.2 Upgrading the Media Format for Existing NSS Volumes

- Section 4.2.1, “Prerequisite,” on page 58
- Section 4.2.2, “Upgrading the Media Format,” on page 58
- Section 4.2.3, “Enabling Hard Links After the Media Upgrade,” on page 59
- Section 4.2.4, “Post Upgrade Task,” on page 60

4.2.1 Prerequisite

The media format upgrade modifies the transaction record structure. For NSS volumes on NetWare, you must disable the Transaction Tracking System™ (TTS™) attribute for the volume and leave it disabled while the upgrade is in progress. After a successful media upgrade, you can enable TTS again.

Before you upgrade the media format for existing NSS volumes on NetWare where TTS is enabled, do the following to disable TTS:

1. Determine the names of all volumes where TTS is enabled.
   1a. At the server console, enter

   ```
   nss /volumes
   ```

   1b. Examine the Attributes column to see which volumes have User Transactions enabled.
2. For each affected volume, enter the following command to disable TTS on the specified volume:

   ```
   nss /NoTransaction=volumename
   ```

   A message is displayed noting that this request takes effect on the next volume deactivation. This prevents the `UXACTION.LOG` file from being opened on subsequent volume activates, and makes it subject to deletion. You delete the file in Step 5.
3. For each affected volume, enter the following command to deactivate the specified volume:

   ```
   nss /VolumeDeactivate=volumename
   ```
4. For each affected volume, enter the following command to activate the specified volume:

   ```
   nss /VolumeActivate=volumename
   ```
5. For each affected volume, use any convenient means to delete the existing user transaction log file (*volname*`:\UXACTION.LOG`), which is located at the root of each volume.

4.2.2 Upgrading the Media Format

The media format upgrade 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.

The commands in this section perform the one-time upgrade for currently active NSS volumes. A message is sent to the logger screen when the process is successfully completed.

Issue the commands at the server console on NetWare as the administrator user, or in the NSS Console (`nsscon`) in a Linux terminal console as the `root` user.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=*volumename*

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.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all

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.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=shared

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

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=local

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

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=local,shared

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

4.2.3 Enabling Hard Links After the Media Upgrade

After the media is upgraded successfully, you must set the Hard Links attribute on volumes where you want to create hard links. The Hard Links attribute is automatically enabled if there are existing hard links on the volume. For information about using hard links on NSS volumes, see Chapter 27, "Managing Hard Links," on page 375.

IMPORTANT: Do not attempt to enable the Hard Links attribute until the upgrade process is complete.

1 Issue on of the following commands at the server console on NetWare as the administrator user, or in the NSS Console (`nsscon`) in a Linux terminal console as the `root` user.

Command	Description
`nss /HardLinks=`*`volumename`*	Enables the Hard Links attribute for a specified volume. This enables hard links to be created on the volume.
`nss /HardLinks=all`	Enables 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.

2 You can verify that the hard links attribute is set for the volume by entering the following command at the server console on NetWare, or at the NSS Console (nsscon) on Linux:

```
volumes
```

The Hard Links attribute appears in the Attributes column for volumes where it is enabled.

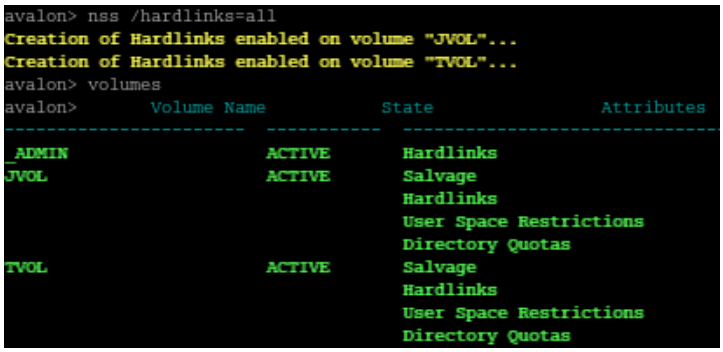

4.2.4 Post Upgrade Task

If you disabled the Transaction Tracking System on an NSS volume on NetWare prior to the media format upgrade, you can enable it again after the upgrade completes successfully.

1. For each affected volume, enter the following command to enable TTS on the specified volume:

   ```
   nss /Transaction=volumename
   ```

2. Verify that TTS is enabled on the volumes.

   2a. At the server console, enter

   ```
   nss /volumes
   ```

   2b. Examine the Attributes column to verify that the volumes have User Transactions enabled.

4.3 Upgrading the Media Format Automatically for New NSS Volumes

You can enable all future NSS volumes to be automatically created with the new media format. Make sure to use the commands on the servers where you plan to use hard links.

IMPORTANT: You must manually enable the Hard Links attribute on a volume before you are able to create hard links on it. For instructions, see Section 27.3, “Enabling or Disabling the Hard Links Attribute,” on page 380.

Issue the following commands from the server console on NetWare, the NSS Console (`nsscon`) on Linux, or 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.

nss /ZLSSUpgradeNewVolumeMediaFormat

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

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.

nss /ZLSSUpgradeNewVolumeMediaFormat /include=local,shared

Creates all new NSS volumes with the upgraded structure whether the volume is local or shared. In a cluster, the clustering system vetoes the upgrade until all nodes in the cluster are running a supported operating system.

Planning NSS Storage Solutions

5

Consider what your storage needs are and how you can effectively manage and divide your storage space to best meet your needs. Use the information in this section to help plan your storage deployment by using the Novell® Storage Services™ file system.

- Section 5.1, “Guidelines for NSS Storage,” on page 63
- Section 5.2, “Compatibility and Interoperability Issues for NSS,” on page 66
- Section 5.3, “Creating NSS Storage Objects in eDirectory,” on page 67
- Section 5.4, “Naming NSS Storage Objects,” on page 67
- Section 5.5, “Access Control for NSS on Linux,” on page 71
- Section 5.6, “File Access for Users,” on page 74
- Section 5.7, “Antivirus Support for NSS,” on page 77
- Section 5.8, “Backup Support for NSS,” on page 78
- Section 5.9, “NSS Support for Memory Mapped Files,” on page 78

5.1 Guidelines for NSS Storage

Use the guidelines in this section when planning your NSS storage solution:

- Section 5.1.1, “Devices,” on page 63
- Section 5.1.2, “Software RAID Devices,” on page 64
- Section 5.1.3, “Device Partitions,” on page 64
- Section 5.1.4, “Multiple I/O Paths (NetWare),” on page 65
- Section 5.1.5, “NSS Pools and Volumes,” on page 65
- Section 5.1.6, “NSS Encrypted Volumes,” on page 66
- Section 5.1.7, “Storage Features,” on page 66

5.1.1 Devices

NSS recognizes devices up to 2 TB in size (where 1 TB = 2E40 bytes = 1,099,511,627,776 bytes). The restriction for NSS is the size that the device reports to the operating system. If you have physical, logical, or virtual devices larger than 2 TB in size, you must carve them into logical devices of less than 2 TB each by using the vendor-provided or other third-party disk-carving software. For more information, see Section 10.1.1, “Device Size Limit,” on page 127.

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). For information about common device types, see Section 10.1.2, “Device Types,” on page 128.

A local hard drive typically contains the operating system software and can optionally be used for applications and user data.

If your system does not have sufficient power loss protection, we require using write-through cache management for SCSI devices to minimize the risk of losing data in the event of power failure. Write-Through cache management assures the file system that writes are being committed to disk as required. For information, see Section 10.11, “Enabling Write-Through Cache Management on SCSI Devices and RAID Controllers (Linux),” on page 138.

Table 5-1 *System Devices on Linux and NetWare*

Platform	Guidelines
Linux	The Linux operating system is stored on the system device by using a Linux POSIX file system (such as Ext3). By default, the system device is managed by the Linux Volume Manager 2 (LVM2), but you can configure it to be managed by EVMS by modifying the default partitioning scheme at install time. For information, see “ Installing Linux with EVMS as the Volume Manager of the System Device ” in the *OES 2: Linux Installation Guide*. NSS data pools and volumes are supported only on devices managed by EVMS.
NetWare	The operating system is stored on the system hard drive in the NSS `sys` pool and `sys:` volume.

Understanding how much free space you will need from each device helps you during the disk carving phase of the NSS configuration. For information about space availability, see Section 10.3, “Viewing Devices on a Server,” on page 131.

For Linux, only EVMS-managed devices show up in the NSSMU and the Storage plug-in to iManager. If you use LVM2 and EVMS to manage different devices on the server, EVMS might display free space on the LVM-managed devices that EVMS cannot actually manage. For information, see Section D.3, “Viewing Free Space Available to EVMS-Managed Devices,” on page 560.

For additional information, see Section 10.1, “Understanding Devices,” on page 127.

5.1.2 Software RAID Devices

NSS supports software RAIDs 0, 1, and 5 on Linux and NetWare®. NetWare also supports nested software RAIDs 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.

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 more information, see Section 13.1, “Understanding Software RAID Devices,” on page 165 and Section 13.2, “Planning for a Software RAID Device,” on page 167.

5.1.3 Device Partitions

NSS management tools automatically create and partitions for you on devices when you create and delete pools. For information, see Section 12.1, “Understanding Partitions,” on page 155.

5.1.4 Multiple I/O Paths (NetWare)

For NetWare, Media Manager provides management of multiple I/O paths for fault-tolerant connections between the server and its storage devices. It is disabled by default. For guidelines, see Section 14.1, “Planning Your Multipath I/O Solution for NSS on NetWare,” on page 189.

To manage multipath I/O on Linux, use Linux multipath utilities. For information, see Chapter 15, “Managing Multipath I/O to Devices (Linux),” on page 195.

5.1.5 NSS Pools and Volumes

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

Table 5-2 *NSS Pools and Volume*

Platform	Guidelines for Use
Linux	You can create NSS pools and volumes to store data on devices managed by EVMS. You cannot use NSS for the operating system and applications; they are stored on Linux POSIX volumes.
NetWare	The `sys` pool and `sys:` volume should be reserved for your system software. You should use additional pools to store applications and user data.

For prerequisites for creating a pool, see Section 16.1, “Guidelines for Creating a Pool,” on page 197.

When creating a pool, you can assign free space from multiple devices to create the maximum-sized pool of 8 TB. You can grow a pool dynamically by adding free space from the same device or different devices.

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 given volume belongs to only one pool.

Pools can be overbooked. If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool size by 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.

When creating a volume, assign it a fixed volume 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.

For guidelines for using volume attributes, see Section 19.1, “Understanding Volume Properties,” on page 253.

For more guidelines for creating and managing NSS volumes, see Section 19.2, “Guidelines for NSS Volumes,” on page 258.

5.1.6 NSS Encrypted Volumes

Encrypted Volume Support is available for data volumes. The `sys:` volume on NetWare cannot be encrypted. You cannot create an encrypted NSS volume during the install (or upgrade) for the system volume or for data volumes. Create encrypted volumes only after you verify a successful system install or upgrade. For information, see "Understanding Encrypted Volume Support" on page 281.

5.1.7 Storage Features

Descriptions of the NSS storage features and guidelines for their use are located in sections that discuss the how to manage them. Table 5-3 identifies the features and provides links to the guidelines.

Table 5-3 *Guidelines for Using NSS Storage Features*

Storage Feature	Refer to
Pool snapshots	Section 18.1, "Understanding Pool Snapshots," on page 229 Section 18.2, "Guidelines for Using and Managing Pool Snapshots," on page 231 Section 7.1, "Cross-Platform Issues for NSS Pool Snapshots," on page 87
Compression	Section 24.1, "Understanding Compression," on page 325
Quotas	Section 25.1, "Understanding Space Quotas," on page 345
Salvage and purge	Section 26.1, "Understanding the NSS Salvage System," on page 361
Hard links	Section 27.1, "Understanding Hard Links," on page 375
Security	Chapter 22, "Securing Access to NSS Volumes, Directories, and Files," on page 297 Chapter 23, "Managing Data Integrity on NetWare Servers," on page 315 Chapter 34, "Security Considerations," on page 439
Performance tuning	Chapter 30, "Tuning NSS Performance on NetWare," on page 405 Chapter 31, "Tuning NSS Performance on Linux," on page 413

5.2 Compatibility and Interoperability Issues for NSS

Table 5-4 lists references for compatibility and interoperability issues for NSS.

Table 5-4 *Compatibility and Interoperability Issues for NSS*

Known Issues	Refer to
Virtualization environments	Chapter 6, “Using NSS in a Virtualization Environment,” on page 79
Cross-platform issues	Chapter 7, “Cross-Platform Issues for NSS,” on page 87
Clustering NSS pools and volumes	Chapter 8, “Cluster-Enabling Shared NSS Devices and Pools with Novell Cluster Services,” on page 93

5.3 Creating NSS Storage Objects in eDirectory

When you use NSSMU or iManager to create an NSS 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.

IMPORTANT: An NSS volume must have a Storage object in eDirectory to be able to participate in Novell Distributed File Services.

For more information about Novell eDirectory, see “Object Classes and Properties” in the *Novell eDirectory 8.8 Administration Guide*.

5.4 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.

- Section 5.4.1, “Case Insensitive Names,” on page 67
- Section 5.4.2, “Number of Characters Allowed,” on page 68
- Section 5.4.3, “Conventions for Valid Names of NSS Storage Objects,” on page 69
- Section 5.4.4, “Other Naming Guidelines,” on page 69
- Section 5.4.5, “UTF-8 Naming Considerations in Mixed-Language Environments (NetWare),” on page 70

5.4.1 Case Insensitive Names

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 also saves 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 19.11, "Mounting NSS Volumes with Linux Commands (Linux)," on page 268.

5.4.2 Number of Characters Allowed

Use the guidelines in Table 5-5 to determine the length requirements for names of NSS Storage objects.

Table 5-5 *Storage Object Name and Password Length*

NSS Storage Object	Minimum Number of Characters	Maximum Number of Characters
Device name for a physical or logical device	2	15
Device name for a software RAID 0 and RAID 5	1	128 Longer names are truncated.
Device name for a software RAID 1, RAID 10, and RAID 15 **NOTE:** RAIDs 10 and 15 can be created only from NSSMU on NetWare.	1	80 Longer names are truncated.
Partition label	2	128
Pool name	2	15
Volume name	2	15
Encryption password for encrypted NSS volumes. Use standard ASCII characters	2 (a minimum of 6 is recommended)	16
Pathnames for files, including the server name, volume name, path delineators (such as colons, slashes, and dots), directory names, filename, and file extension	1	255

5.4.3 Conventions for Valid Names of NSS Storage Objects

Valid device, pool, and volume object names conform to the following naming conventions. We recommend that you also consider the character conventions for the software RAID names in order to have consistent naming policies on your system.

- Use only valid characters:

  ```
  ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789_!@#$%&()
  ```

  IMPORTANT: Special characters (non-alphanumeric) can create confusion or problems for some configuration files, command line utilities, applications, and scripts. For example, the ampersand (&), at (@), dollar sign ($), exclamation point (!), percent (%), and number sign (#) characters should be avoided for this reason. For more information, see Section 5.4.4, "Other Naming Guidelines," on page 69

- Do not use spaces in the object name.
- Do not begin or end the object name with an underscore (_).
- Do not use multiple contiguous underscores (__) anywhere in the object name.

5.4.4 Other Naming Guidelines

- In general, we recommend that you avoid using reserved names or words as names of Storage objects in order to avoid confusion.

  For example, the following case-insensitive names are reserved names on NetWare:

  ALL
  AUX
  CLOCK
  COM1
  COM2
  COM3
  COM4
  CON
  LPT1
  LPT2
  LPT3
  NETQ
  NSS_ADMIN
  NUL
  PIPE
  PRN
  SYS (This is the default name of the system pool and volume on NetWare.)

- Some characters on Linux, such as the ampersand (&), dollar sign ($), exclamation point (!), and number sign (#) characters, can cause problems in some configuration files, command line utilities, applications, and scripts. You might need to use different techniques in each case to

make the name be accepted in the manner intended. Refer to the documentation for the specific consumer application or utility to find how to treat names that contain special characters in that environment.

To avoid this extra effort, we recommend that you avoid using special characters in names of Storage objects.

- Because the “at sign” (`@`) character (also called “the at symbol”) 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 `@` 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.

- If spaces are used in User or Group object names, you must enclose the object name in double quotation marks (`"`) in order for it to be recognized in command line utilities, scripts, and applications.

- If special characters are used in User or Group object names and passwords, you might need to use different escape techniques in command line utilities (such as Bash on Linux) to make the name be accepted in the manner intended. Refer to the documentation for the specific command line utility to find how to escape special characters in that environment.

  For example, enclosing the name in double quotation marks and preceding the character with a backslash are common techniques for escaping special characters when parsing command lines.

  To avoid this extra effort, we recommend that you avoid using special characters in names of User and Group objects and in passwords.

5.4.5 UTF-8 Naming Considerations in Mixed-Language Environments (NetWare)

NSS supports UTF-8 (8-bit Unicode Transformation Format) compatible filenames. Unicode allows the characters for multiple languages to be represented using a single Unicode representation. We recommend that you use the same language environments on the server and on the client workstations that access the server. UTF-8 support is enabled by default in the Novell Client™ 4.91 SP2 for Windows and later.

In NetWare 6.5 SP1 or earlier versions of NetWare, NSS used code-page translation to access the file system information on the server. The best practice was to use the same code pages on the servers and clients. However, workstation clients could use any local code page settings to read and write file system data to the server. Because each workstation might use a different code page, it was possible for file system data to be stored at the server based on different code pages. Unfortunately, this practice resulted in behavior that appeared to work, but in reality, it stored the wrong Unicode values in the NSS filenames. Users had to change their local code pages to see the names of the files as stored by another user with a different code page. The solution to this issue was for Novell to move to UTF-8 in NetWare 6.5 SP2.

If you upgrade data from NetWare 6.5 SP1 or earlier versions where you allowed users to use different local language codes than the language code of the server, you might encounter filename problems. For information, see *Supporting Mixed Language Environments with Novell NetWare (TID 10097059)* (http://support.novell.com/docs/Tids/Solutions/10097059.html) in the Novell Support Knowledgebase.

5.5 Access Control for NSS on Linux

This section describes how Novell eDirectory, NCP™ (NetWare Core Protocol™) Server, and Linux User Management (LUM) work with Novell Storage Services to provide access to NSS volumes on OES 2 Linux servers.

- Section 5.5.1, "Administrator User and Root User Roles for NSS," on page 71
- Section 5.5.2, "NSS File System Users," on page 72
- Section 5.5.3, "Novell Trustee Model," on page 73
- Section 5.5.4, "POSIX Permissions," on page 74
- Section 5.5.5, "How NSS Uses Novell Linux User Management," on page 74

5.5.1 Administrator User and Root User Roles for NSS

The Administrator user and the Linux `root` user are two very different concepts. It is important to understand the role of each in managing your NSS volume on OES 2 Linux servers.

- "Administrator User" on page 71
- "Root User" on page 72

Administrator User

The Administrator user is an eDirectory user who is given all file system trustee rights for the server, including the Supervisor right. The Administrator user account, or the Administrator equivalent user account, is given the following privileges:

- The user identity and credentials are defined in eDirectory.
- The user is assigned as a trustee of the NSS volume and given all file system trustee rights for that volume. You can also create a group for administrators with equivalent rights, and assign the user to that group.
- The username must be Linux-enabled with Linux User Management (LUM), which gives the user both an eDirectory GUID and a POSIX UID on the server.
- The user belongs to the Administrator group for the server that is Linux-enabled with LUM.

The Administrator user who installs NSS on OES 2 Linux is automatically given these privileges. Any other administrator, including the Tree Administrator user, who you want to be able to manage the NSS storage must be manually configured with the same privileges.

IMPORTANT: The Tree Administrator user is not automatically granted permissions to OES Linux servers installed in the tree, as is the case for NetWare servers.

For more information about Linux-enabled eDirectory users, see Section 5.5.2, "NSS File System Users," on page 72.

For a Linux server, the administrator logs in to iManager as the Administrator user (or Administrator equivalent user) to manage the NSS volume on Linux. The Administrator user can also use the iManager Files and Folders plug-in, NetStorage, and the Novell Client to manage file system trustee assignments, trustee rights, inherited rights masks, and file and directory attributes. These tools can also be used to purge and salvage files for volumes where the Salvage attribute is enabled.

Root User

The `root` user is a local Linux user who 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 with LUM. This allows you to log in to the server as the `root` user when eDirectory services are not available. The `root` user is the only local Linux user who is allowed to access NSS via the VFS layer without having an eDirectory GUID.

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

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 that 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.

IMPORTANT: 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.

5.5.2 NSS File System Users

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

- "eDirectory Users" on page 72
- "Linux-Enabled eDirectory Users" on page 73
- "Local Linux Users" on page 73

eDirectory Users

NSS uses the eDirectory GUID of a user to control access by using the Novell Trustee model. Users of the NSS volume and the Administrator user (or Administrator equivalent user) who manages the volume must be defined as users in Novell eDirectory. For information about managing users with eDirectory, see the *Novell eDirectory 8.8 Administration Guide*.

To grant access to eDirectory users, you must assign them to be trustees in the file system, grant them file system trustee rights, and set inherited rights filters. For more information about configuring trustees for NSS on Linux, see Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297.

Linux-Enabled eDirectory Users

Linux-enabled eDirectory users are users who are defined in eDirectory, granted file system trustee rights to the NSS volume, and Linux-enabled with Linux User Management. Linux-enabled eDirectory users have both a POSIX UID and an eDirectory GUID. You must Linux-enable users who need to access NSS volumes via Samba, NFS, third-party AFP solutions, or who need to use Linux utilities, commands, or services. NetStorage for Linux requires that users be Linux-enabled if NetStorage is configured to use OpenSSH for user access.

IMPORTANT: A Linux service or utility must also be enabled for LUM in order for users to access the file system with it.

For OES 2 Linux and later, it is no longer necessary to Linux-enable the users with LUM in order for user quotas (space restrictions) to be enforced. NCP Server for Linux has been modified to provide the GUID information that NSS needs for file ownership. NSS uses file ownership information to enforce user space restrictions based on a user's eDirectory username.

Users who create hard links must be Linux-enabled in order to use the `ln` command on the server. It is not necessary to Linux-enable users if they are only consumers of the hard link.

For OES 2 Linux and later, if users are Linux-enabled with LUM, the file creator, modifier, and deleter fields are recorded with the username of the user who performs the action. For users who are not Linux-enabled with LUM, the file create and modifier are recorded as the actual user; however, the `root` user or the Nobody user (if it exists) is listed as the deleter in the file's metadata instead of the actual user.

NOTE: In OES 1 Linux, the modifier field and deleter field are reported as the `root` user or Nobody user for non-LUM-enabled users.

For information about installing and configuring Linux User Management and enabling users and groups for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

Local Linux Users

Local Linux users are 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.

5.5.3 Novell Trustee Model

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 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 Using Trustees " in the *OES 2 SP1: File Systems Management Guide*.

5.5.4 POSIX Permissions

For NSS volumes on Linux, the POSIX directory and file permissions are not used to determine access permission. Access control is based on the Novell trustee model and not on the POSIX permissions or access control lists (ACLs). NSS uses the POSIX permission fields to display Read Only, Read/Write, Execute, and Hidden attributes for directories and files. NSS does not use the Group ID field. Group ids associated with POSIX have no effect on files stored on NSS.

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.

By default, NSS sets the POSIX permissions fields for directories to 0777 (`drwxrwxrwx`). Some Linux services specify permissions needed to use the service. NSS provides the `nss /PosixPermissionMask=`*`mask`* option that allows you to change the default POSIX permissions, such as for the Group or Other fields.

For example, SSH requires that the POSIX permissions on home directories be set so that the Other field has no permissions. When you use NSS volumes as home directories, you must change the permission to 0770 on the home directories. You can use the `nss /PosixPermissionMask=0770` command in the NSS Console (`nsscon`) to modify the permissions.

For information and examples of how to interpret POSIX settings on your NSS volume on Linux, see "Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions" in the *OES 2 SP1: File Systems Management Guide*.

5.5.5 How NSS Uses Novell 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.

Linux-enabled eDirectory users have both UIDs as local Linux users and GUIDs as eDirectory users. NSS needs the UID to execute protocols and services that communicate to NSS through the VFS layer only. 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 finds the GUID for the user who issues the command, then executes the command. Without LUM, NSS cannot identify a GUID for the UID it receives, and rejects the command with an error.

For information about installing and configuring LUM, enabling Linux services and utilities, and enabling users and groups for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

5.6 File Access for Users

NSS supports access via NCP and other protocols to eDirectory users and Linux-enabled eDirectory users.

IMPORTANT: NSS uses the Novell trustee model for file access. Users must be made file system trustees and granted trustee rights to data on the NSS volume that you want them to be able to access. Rights management can be done in multiple management tools, including iManager, Novell Remote Manager, the Novell Client and other NCP services, and command line commands. For information, see Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297.

- Section 5.6.1, "NCP," on page 75
- Section 5.6.2, "Novell AFP," on page 76
- Section 5.6.3, "Novell CIFS," on page 76
- Section 5.6.4, "Novell Domain Services for Windows," on page 76
- Section 5.6.5, "Samba," on page 76
- Section 5.6.6, "SSH (Secure Shell)," on page 77
- Section 5.6.7, "Accessing Files with Linux Services, Utilities, and Protocols," on page 77

5.6.1 NCP

NCP (NetWare Core Protocol) is the default protocol for accessing data on NSS volumes. NCP Server is required for NSS even if users access the volume via other protocols. Users access data on NSS volumes by using the Novell Client software on their Windows, Vista*, or Linux workstations. This document refers collectively to those workstations as "Novell clients".

IMPORTANT: NSS uses NCP Server by default and requires that NCP Server be running even if your users are accessing the volume via other protocols.

NCP Server for Linux is installed by selecting *NCP Server and Dynamic Storage Technology* from the *OES Services* menu in the YaST installation interface. For information about NCP Server for Linux, see the *OES 2 SP1: NCP Server for Linux Administration Guide*.

NCP Server for NetWare is automatically installed during the NetWare install. For information about NCP Server for NetWare, see the *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*.

NCP Server works with Novell eDirectory, the Novell Client, and other NCP-based services such as NetStorage to authenticate and manage user sessions. 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 services. NSS cooperates with NCP Server to track file ownership and file system trustee assignments, trustee rights, and inherited rights based on the Novell trustee model.

The Linux file system interface uses UTF-8 encoding for all filenames. When accessing files with NCP, make sure to use the UTF-8 enabled NCP software that is available in the latest Novell Client. For more information, see Section 5.4.5, "UTF-8 Naming Considerations in Mixed-Language Environments (NetWare)," on page 70.

For information about configuring and managing NCP Server, see the *OES 2 SP1: NCP Server for Linux Administration Guide*.

5.6.2 Novell AFP

NSS supports access to NSS volumes using the Novell AFP (Apple Filing Protocol). Novell AFP is installed automatically during the NetWare install. For OES 2 SP1 Linux and later, Novell AFP for Linux is installed by selecting *Novell AFP* from the OES Services menu in the YaST install interface.

For information about Novell AFP for Linux, see the *OES 2 SP1: Novell AFP For Linux Administration Guide*. For information about Novell AFP for NetWare, see "Working with Macintosh Computers" in the *OES 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide*.

5.6.3 Novell CIFS

NSS supports access to NSS volumes using Novell CIFS. Novell CIFS for NetWare is installed automatically during the NetWare install. For OES 2 SP1 Linux and later, Novell CIFS for Linux is installed by selecting *Novell CIFS* from the OES Services menu in the YaST install interface.

For information about Novell CIFS for Linux, see the *OES 2 SP1: Novell CIFS for Linux Administration Guide*. For information about Novell CIFS for NetWare, see "Working with UNIX Machines" in the *OES 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide*.

5.6.4 Novell Domain Services for Windows

NSS supports access to NSS volumes using Novell Domain Services for Windows (DSfW). DSfW configures Samba access for Samba/CIFS users. Administrators must export NSS volumes over Samba so that domain users (eDirectory users in the DSfW domain partition) can access NSS volumes over Samba/CIFS.

Samba/CIFS users under the domain are Linux-enabled with Linux User Management. The Domain Users group must be associated with the UNIX Workstation objects of the server (or servers if the volume is used in a cluster) where the volume is mounted in order to give the users access to the NSS volume via Samba/CIFS.

5.6.5 Samba

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 ACLs and standard POSIX permissions for Samba access to an NSS volume. Instead, the Administrator user or Administrator 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 Samba service must also be enabled in LUM.

For information about configuring and managing Samba services for your OES 2 Linux server, see the *OES2 SP1: Samba Administration Guide*.

5.6.6 SSH (Secure Shell)

You can give users SSH (Secure Shell) access to NSS volumes by Linux-enabling users and the SSH utility in Linux User Management. For information, see see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

In addition, SSH requires that the POSIX permissions on home directories be set so that the Other field has no permissions. By default, NSS sets the POSIX permissions to 0777 and SSH is disabled in Linux User Managerment. If you use NSS volumes for home directories and you want users to have SSH access to them, you must modify the POSIX permissions on NSS volumes to 0770. You must also enable SSH with Linux User Management.

Add the following command in the `/etc/opt/novell/nss/nssstart.cfg` file to turn off all of the bits corresponding to the Other field:

```
/PosixPermissionMask=0770
```

The setting applies to all NSS volumes on the server. If the volume is shared in a cluster, make sure to add the command to the `nssstart.cfg` file on all nodes and to Linux-enable SSH on all nodes.

5.6.7 Accessing Files with Linux Services, Utilities, and Protocols

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. Users must be Linux-enabled with Linux User Management in order to use any of the standard Linux protocols, utilities, commands, services, or APIs for the NSS volume.

IMPORTANT: Any Linux service or utility that you want users to have access to must also be enabled in Linux User Management.

For information about installing and configuring Linux User Management, enabling users and groups for Linux, and enabling Linux services and utilities, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

5.7 Antivirus Support for NSS

For information about antivirus issues for NSS on Linux, see *Providing Anti-Virus Support for NSS on Linux* (http://developer.novell.com/wiki/index.php/Providing_Anti-Virus_Support_for_NSS_on_Linux).

For a current list of antivirus software vendors that support Novell Open Enterprise Server, see *Novell Open Enterprise Server Partner Support: Backup and Antivirus Support* (http://www.novell.com/products/openenterpriseserver/partners_communities.html). This list is updated quarterly.

The Apple Filing Protocol (AFP) support for NSS files on OES 2 SP1 Linux is implemented via a technology that bypasses the real-time scanning employed by most OES 2 antivirus solutions. NSS files shared through an AFP connection might be protected by on-demand scanning on the OES 2 server or by real-time and on-demand scanning on the Apple client.

5.8 Backup Support for NSS

For information about OES 2 backup support for NSS, see Chapter 29, "Managing Backup and Restore for Data and Trustee Information," on page 399.

For a current list of backup software vendors that support Novell Open Enterprise Server, see *Novell Open Enterprise Server Partner Support: Backup and Antivirus Support* (http://www.novell.com/products/openenterpriseserver/partners_communities.html). This list is updated quarterly.

5.9 NSS Support for Memory Mapped Files

NSS has limited support for memory mapped files, primarily to support loading programs. NSS does not fully support memory mapped files especially if the application uses sparse files.

For example, the CopyCat application used by Netatalk uses sparse files for its database. Netatalk tries to create a CopyCat database as a sparse file called `.AppleDB` in the root of the volume by using memory mapped IO. This can cause the server to hang if you are using an NSS volume as the Netatalk share because of the limited support in NSS for this combination.

Using NSS in a Virtualization Environment

6

Use the information in this section to help you deploy Novell® Storage Services™ file system and services in a virtualization environment.

For information about Xen* virtualization, see *SUSE Linux Enterprise Server 10 SP2: Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/data/bookinfo.html)*.

- Section 6.1, “Guidelines for Using NSS in a Xen Virtualization Environment,” on page 79
- Section 6.2, “Installing OES 2 Linux on a Virtual Machine,” on page 82
- Section 6.3, “Initializing New Virtual Disks on the Guest Server,” on page 83
- Section 6.4, “What’s Next,” on page 84

6.1 Guidelines for Using NSS in a Xen Virtualization Environment

Consider the following guidelines when planning to use NSS in a virtualization environment:

- Section 6.1.1, “Host Server Issues,” on page 79
- Section 6.1.2, “Virtual Machine Issues,” on page 81
- Section 6.1.3, “Guest Server Issues,” on page 81

6.1.1 Host Server Issues

- “Running NSS on the Host Server Is Not Supported” on page 79
- “Using RAIDs” on page 80
- “Using Multipath Devices” on page 80

Running NSS on the Host Server Is Not Supported

NSS pools and volumes are not supported on the Xen host server in a Xen virtualization environment. You can install NSS on the guest servers from inside the guest server environment, just as you would if the guest servers were physical servers.

When you create a virtual machine, you must assign devices to it. If you plan to use the virtualization guest server as a node in a cluster and you need to be able to fail over cluster resources to different physical servers, you must assign SAN-based physical devices to the virtual machine. You create the NSS pools and volumes from within the guest server.

If you install Novell Cluster Services™ in the host server environment, the cluster resources use shared Linux POSIX volumes, and do not use shared NSS pools.

If you install Novell Cluster Services in the guest server environment, the guest server is a node in the cluster. The disk sharing is managed by Novell Cluster Services from within the guest server environment. You can use shared NSS pools as cluster resources that run on the guest server and on other nodes in that cluster.

For information about deployment scenarios using shared NSS pools in clusters in a virtualization environment, see "Configuring Novell Cluster Services in a Xen Virtualization Environment" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.

Using RAIDs

In a Xen virtualization environment, if you need to use RAIDs for device fault tolerance in a high-availability solution, we recommend that you use standard hardware RAID controllers. Hardware RAIDs provide better performance over using software RAIDs on the virtualization host server or guest server.

To get the best performance from a software RAID, create a RAID device on the Xen host and present that device to the guest VM. Each of the RAID's segments must be on different physical devices. It is best to present the entire physical RAID device or a physical partition of the RAID device to the guest VM, and to not present just a file-backed virtual device.

NSS is not supported to run in the virtualization host server environment, so NSS software RAIDs cannot be used there. Xen supports using Linux `mdadm` for software RAIDs on the host server.

If you attempt to create and manage a software RAID on the guest server in a production environment, make sure to present different physical devices to the guest VM that you want to use for the software RAID. Using segments from virtual devices that actually reside on the same physical device on the host server slows performance and provides no protection against failed hardware devices. The maximum number of disks that can be presented to the VM is 16 (`xvda` to `xvdp`). Xen provides a mechanism to dynamically add and remove drives from a VM, but that capability is currently not supported in paravirtualized NetWare.

Using NSS software RAIDs in a virtualization guest server environment has not been tested.

Using Multipath Devices

If it is available, use your storage vendor's multipath I/O management solution for the storage subsystem. In this case, the multiple paths are resolved as a single device that you can assign to a virtual machine.

Do not use multipath management tools in the guest environment.

If a storage device has multiple connection paths between the device and the host server that are not otherwise managed by third-party software, use Linux multipathing to resolve the paths into a single multipath device. When assigning the device to a VM, select the device by its multipath device node name (`/dev/mapper/mpathN`). The guest server operating system is not aware of the underlying multipath management being done on the host. The device appears to the guest server as any other physical block storage device. For information, see "Managing Multipath I/O for Devices" in the *SLES 10 SP2: Storage Administration Guide*.

6.1.2 Virtual Machine Issues

- "Assigning Physical Disks or Disk Partitions to the Virtual Machine" on page 81
- "Assigning File-Backed Disk Images for Virtual Devices" on page 81

Assigning Physical Disks or Disk Partitions to the Virtual Machine

For the best performance on a Xen guest server, NSS pools and volumes on NetWare should be created on block storage devices that are local SCSI devices, Fibre Channel devices, iSCSI devices, or partitions on those types of devices.

SATA or IDE disks have slower performance because special handling is required when working through the Xen driver to ensure that data writes are committed to the disk in the order intended before it reports back.

Assigning File-Backed Disk Images for Virtual Devices

Novell supports file-backed disk images on virtual machines, but does not recommend using them for important data because the volume can become corrupt after a power failure or other catastrophic failure. For example, file-backed volumes might be useful for training and sales demonstrations.

WARNING: Data corruption can occur if you use Xen file-backed disk images for NSS volumes on the guest server in the event of a power failure or other catastrophic failure.

6.1.3 Guest Server Issues

Unless otherwise indicated, the issues in this section apply to both OES 2 Linux and NetWare, and to NetWare 6.5 SP7.

- "Initializing Virtual Disks" on page 81
- "Configuring Write Barrier Behavior for NetWare in a Guest Environment" on page 81
- "NSS Features that Are Not Supported in a Virtualization Environment" on page 82

Initializing Virtual Disks

The primary virtual disk (the first disk you assign to the virtual machine) is automatically recognized when you install the guest operating system. The other virtual devices must be initialized before any space is shown as available for creating a pool. Without initializing the devices, no space is shown as available for pool creation. For information, see Section 6.3, "Initializing New Virtual Disks on the Guest Server," on page 83.

Configuring Write Barrier Behavior for NetWare in a Guest Environment

Write barriers are needed for controlling I/O behavior when writing to SATA and ATA/IDE devices and disk images via the Xen I/O drivers from a guest NetWare server. This is not an issue when NetWare is handling the I/O directly on a physical server.

The XenBlk Barriers parameter for the `SET` command controls the behavior of XenBlk Disk I/O when NetWare is running in a virtualization environment. The setting appears in the Disk category when you issue the `SET` command in the NetWare server console.

Valid settings for the XenBlk Barriers parameter are integer values from 0 to 255, with a default value of 16. A non-zero value specifies the depth of the driver queue, and also controls how often a write barrier is inserted into the I/O stream. A value of 0 turns off XenBlk Barriers.

A value of 0 (no barriers) is the best setting to use when the virtual disks assigned to the guest server's virtual machine are based on physical SCSI, Fibre Channel, or iSCSI disks (or partitions on those physical disk types) on the host server. In this configuration, disk I/O is handled so that data is not exposed to corruption in the event of power failure or host crash, so the XenBlk Barriers are not needed. If the write barriers are set to zero, disk I/O performance is noticeably improved.

Other disk types such as SATA and ATA/IDE can leave disk I/O exposed to corruption in the event of power failure or a host crash, and should use a non-zero setting for the XenBlk Barriers parameter. Non-zero settings should also be used for XenBlk Barriers when writing to Xen LVM-backed disk images and Xen file-backed disk images, regardless of the physical disk type used to store the disk images.

To configure XenBlkBarriers:

1 In the server console on the guest NetWare server, enter

```
SET XenBlkBarriers=value
```

For example, to turn off XenBlk Barriers for virtual disks based on physical SCSI, Fibre Channel, and iSCSI disks, enter

```
SET XenBlkBarriers=0
```

NSS Features that Are Not Supported in a Virtualization Environment

Some NSS features are not supported in a Xen guest server environment.

Table 6-1 *NSS Feature Support in a Guest Server Environment*

NSS Feature	NSS on Linux	NSS on NetWare
Data shredding	Not supported	Not supported
Multipath I/O	Not applicable; not supported on Linux	Not supported
Software RAIDs	Not tested	Not tested

6.2 Installing OES 2 Linux on a Virtual Machine

When you install OES 2 Linux on a virtual machine, we recommend that you configure a virtual machine with multiple devices. Use the primary disk on the guest server as the system device with LVM2 (the YaST install default) as the volume manager. After the install, assign additional storage resources from the host to the virtual machine. The additional disks can use LVM2 or EVMS as needed. In this scenario, NSS volumes are created only on the data disks for the guest server, not on the system disk that you are using for the guest server's system device.

IMPORTANT: When you create the virtual machine, make sure to configure the size of the primary virtual disk according to the amount of space you need for the boot (`/boot`), `swap`, and root (`/`) volumes.

If you decide to use EVMS for the system device on the virtual machine, follow the install instructions in "Installing Linux with EVMS as the Volume Manager of the System Device" in the *OES2 SP1: Linux Installation Guide*.

For information about creating a Xen virtual machine, see *SUSE Linux Enterprise Server 10 SP2: Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/data/bookinfo.html)*.

6.3 Initializing New Virtual Disks on the Guest Server

A new virtual disk can appear as an unformatted disk to the guest server if it does not have a partition table associated with it. You must initialize the device on the guest server just as you do for a blanked-out device on a physical server.

The primary virtual disk (the first disk you assign to the virtual machine) is automatically recognized when you install the guest operating system. After the install, use NSS tools to initialize additional blanked-out virtual devices where you plan to create NSS pools and volumes. On Linux, the NSS initialize function makes EVMS the volume manager of the device, and adds a NetWare Segment Manager.

You can initialize the disk by using the *Initialize Disk* function in NSSMU or in the Storage plug-in to iManager. For general instructions for initializing disks, see Section 10.7, "Initializing a Disk," on page 134.

To initialize devices for the guest server:

1. On the host, use the virtualization management tool to create and allocate virtual devices for the virtual machine.

   For information, see *SUSE Linux Enterprise Server 10 SP2: Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/data/bookinfo.html)*.
2. If the guest server is not running, boot the guest server now.
3. In iManager, click *Storage* > *Devices*.
4. Browse to locate and select the guest server to view a list of its devices.

   The virtual server has a Server object in the Novell eDirectory™ database, just like a physical server.
5. In the *Devices* list, select the newly added virtual device to view its details.
6. Verify that the device you selected is the new unformatted device, and not your system device or a formatted device.

   WARNING: Do not initialize the system disk.

   For example, for an unformatted device, the *Free Space* size is reported as 0.00 Bytes.

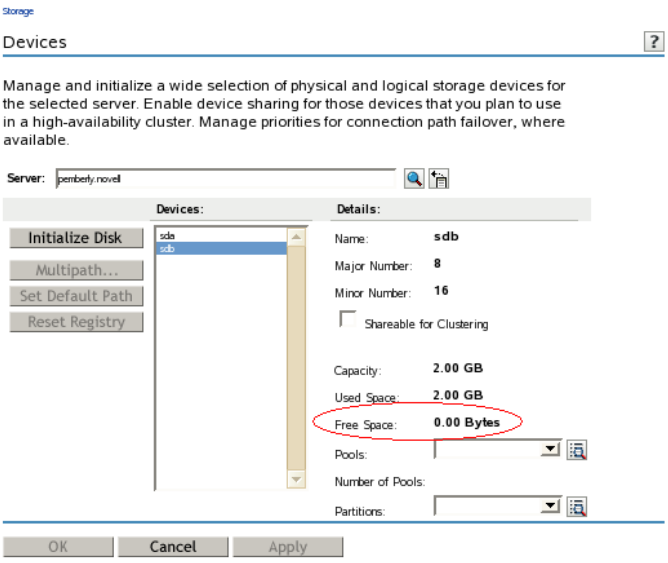

7 Click *Initialize*.

When the page refreshes, the device is initialized and available for further configuration with NSS pools and volumes.

8 Verify that the *Free Space* is now reported properly.

For example, after the device is initialized, the *Free Space* is reported to be the same as *Capacity*.

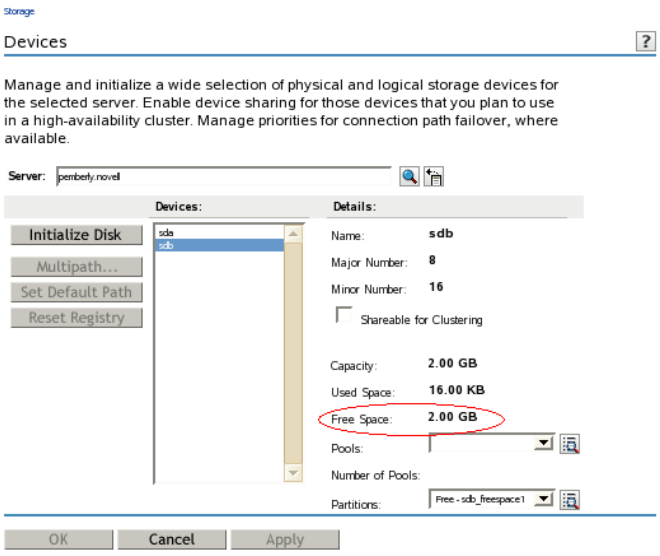

6.4 What's Next

To get started with virtualization, see *SUSE Linux Enterprise Server 10 SP2: Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/data/bookinfo.html)*

For information on setting up NetWare on a Xen-based virtual guest server, see "Installing and Managing OES 2 SP1 NetWare on a Xen-based VM Host Server" in the *OES 2 SP1: NetWare Installation Guide*.

For information on setting up OES 2 Linux on a Xen-based virtual guest server, see "Installing, Upgrading, or Updating OES 2 SP1 Linux on a Xen-based Virtual Machine" in the *OES2 SP1: Linux Installation Guide*.

Cross-Platform Issues for NSS 7

This section describes cross-platform compatibility issues for the Novell® Storage Services™ file system and services between NetWare® and Novell Open Enterprise Server 2 Linux servers.

- Section 7.1, “Cross-Platform Issues for NSS Pool Snapshots,” on page 87
- Section 7.2, “Cross-Platform Issues for NSS Volumes,” on page 87
- Section 7.3, “Cross-Platform Issues for NSS Features,” on page 88
- Section 7.4, “Cross-Platform Issues for File Access,” on page 89
- Section 7.5, “Cross-Platform Issues for Management Tools,” on page 90

7.1 Cross-Platform Issues for NSS Pool Snapshots

Different pool snapshot technologies are used for NSS pools on NetWare and NSS pools on Linux. You can create pool snapshots on either platform, but you should not move them to another platform. Pool snapshots taken on Linux do not work on NetWare, and vice versa.

Consider these guidelines when working with NSS pool snapshots:

- The snapshots taken on a given platform are unusable if you move the pool’s devices cross-platform. Before you move a pool with existing snapshots taken of it to a different platform, delete all existing snapshots for the pool.

  WARNING: You might not be able to open the original pool on the other platform if you do not delete the snapshots.

- On NetWare, do not use the Pool Snapshot feature for a clustered pool in a mixed-platform cluster.
- On Linux, NSS does not support using pool snapshots for clustered pools.

  You must remove any existing pool snapshots for a clustered pool on NetWare before you cluster migrate the pool cluster resource from a NetWare server to a Linux server during a rolling cluster conversion.

7.2 Cross-Platform Issues for NSS Volumes

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

- Section 7.2.1, “Guidelines for File Systems on Linux Servers,” on page 88
- Section 7.2.2, “Guidelines for File Systems on NetWare Servers,” on page 88

7.2.1 Guidelines for File Systems on Linux Servers

OES 2 Linux requires a Linux POSIX file system volume for the operating system, such as Ext3. 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 Linux cross-platform.
- You cannot move Linux POSIX file systems cross-platform to a NetWare server.
- You can move NSS volumes cross-platform to a Linux server.
- NSS volumes that were originally created on NetWare can be moved cross-platform from NetWare to Linux. If both platforms support the same media format, Linux volumes can be moved cross-platform from Linux to NetWare.

7.2.2 Guidelines for File Systems on NetWare Servers

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 POSIX file system volume or on an NSS volume on Linux.
- You cannot create or mount Linux POSIX file systems on a NetWare server.
- You cannot SAN boot NetWare cross-platform.
- You should not move an NSS system volume from NetWare to Linux unless you intend to use it as a data volume (or not at all) while it is mounted on the Linux server.
- If you move an NSS system pool cross-platform, any volumes it contains function as data volumes on Linux, including the `sys:` volume.
- You can create an NSS data volume on NetWare, and then move the volume cross-platform to Linux.
- If you use shared pools in a cluster, only pools that are originally created on NetWare can be migrated or failed back from Linux to NetWare. Mixed-platform clusters are supported only for rolling cluster conversions from NetWare to Linux.

7.3 Cross-Platform Issues for NSS Features

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

- Section 7.3.1, “Multipath I/O to Devices,” on page 89
- Section 7.3.2, “NSS Support for Removable Media,” on page 89
- Section 7.3.3, “NSS Support for Transaction Tracking,” on page 89
- Section 7.3.4, “Modified File List,” on page 89

7.3.1 Multipath I/O to Devices

If a device has multiple connection paths between the hardware controller and the HBA on the server, each path appears as a separate device to the operating system. You must use multipath management tools to resolve the apparent devices into a single multipath device.

Media Manager provides multipath I/O for devices on NetWare. Media Manager is not available on Linux. For information, see Chapter 14, "Managing Multipath I/O to Devices (NetWare)," on page 189.

On Linux, use Linux multipath I/O management tools. You should configure multipath I/O before using NSS management tools to create NSS software RAIDs, pools, or volumes on the devices. For information, see Chapter 15, "Managing Multipath I/O to Devices (Linux)," on page 195.

7.3.2 NSS Support for Removable Media

Removable media such as CDs, DVDs, CD and DVD image files, and DOS partitions are typically mounted as file systems native to the Linux or NetWare platform on the server.

On NetWare, CDs, DVDs, CD and DVD image files, and DOS partitions are mounted as NSS volumes. Disks in USB drives are mounted as local DOS FAT partitions. For information, see Chapter 21, "Managing Removable Media (NetWare)," on page 291.

On Linux, removable media and partitions are mounted by using Linux POSIX file systems options. For information, see "Some Other Supported File Systems" in the *SLES 10 SP2 Installation and Administration Guide*.

7.3.3 NSS Support for Transaction Tracking

NSS provides file transaction tracking with the Transaction Tracking System™ (TTS™) for NSS volumes on NetWare. TTS is not available for NSS on Linux. For information about using TTS, see Section 23.2, "Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare)," on page 317.

7.3.4 Modified File List

The Modified File List attribute for NSS volumes is available only for NSS volumes on NetWare.

This feature is seldom used since the introduction of the Event File List support in OES 1 Support Pack 1. Consider using the Event File List instead, which is available for NSS on Linux and NetWare. For information, see "FileEvents.xml Definitions" (http://developer.novell.com/documentation/vfs/vfs__enu/data/ak7gh2x.html) in *NDK: Virtual File Services* (http://developer.novell.com/documentation/vfs/vfs__enu/data/bktitle.html).

7.4 Cross-Platform Issues for File Access

On Linux, users of the NSS volume must be Linux-enabled with Linux User Management if you want to give users access via any non-NCP Linux protocol or any Linux service or utility, such as Samba, FTP, or SSH. You must also LUM-enable the Linux service or utility. For information, see Section 5.5, "Access Control for NSS on Linux," on page 71.

7.5 Cross-Platform Issues for Management Tools

- Section 7.5.1, “Storage-Related Plug-Ins for Novell iManager 2.7,” on page 90
- Section 7.5.2, “Interoperability of Protocols for the iManager Server and Target Server,” on page 90
- Section 7.5.3, “Do Not Run EVMS Tools When Working in iManager or NSSMU,” on page 92
- Section 7.5.4, “Management Tools for Managing Pool Snapshots,” on page 92
- Section 7.5.5, “Management Capabilities for Software RAIDs,” on page 92

7.5.1 Storage-Related Plug-Ins for Novell iManager 2.7

The following storage-related plug-ins for OES 2 require Novell iManager 2.7:

Storage-Related Plug-In	NSS on Linux	NSS on NetWare
Archive and Versioning (Archive and Version Services)	Yes	Yes
Clustering (Novell Cluster Services™)	Yes	Yes
Distributed File Services (Novell Distributed File Services)	Yes	Yes
Novell AFP	Yes (OES 2 SP1 Linux and later)	Yes
Novell CIFS	Yes (OES 2 SP1 Linux and later)	Yes
Novell NFS	No Use Linux NFS.	Yes
Files and Folders	Yes	Yes
Storage (NSS file system)	Yes	Yes

The storage-related plug-ins share common code in the `storagemgmt.npm` file. 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.

7.5.2 Interoperability of Protocols for the iManager Server and Target Server

The following table provides information about the protocols needed to use iManager when managing 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 7-1 *Interoperability of Protocols for the iManager Server and Target Servers*

iManager Server	Target Server			
	OES 1 Linux or Later	**OES 1 NetWare or Later, or NetWare 6.5 SP4 or Later**	**NetWare 6.5 SP3**	**NetWare 6.5 SP2**
OES 1 Linux or Later	* WBEM CIFS	* WBEM CIFS	WBEM (Start OpenWBEM.) CIFS	 CIFS (Field Patch 2B)
OES 1 NetWare and Later, or NetWare 6.5 SP4 and Later	* WBEM CIFS	* WBEM NCP™ CIFS	WBEM (Start OpenWBEM.) * NCP CIFS	 * NCP CIFS (Field Patch 2B)
NetWare 6.5 SP3	* WBEM CIFS	* WBEM NCP CIFS	WBEM (Start WBEM.) * NCP CIFS	 * NCP 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.

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

To check the status of WBEM:

1 As `root` in a console shell, enter

```
rcowcimomd status
```

To start WBEM:

1 As `root` in a console shell, enter

```
rcowcimomd start
```

For information about installing OpenWBEM, see "Setting Up OpenWBEM" in the *OES 2 SP1: OpenWBEM Services Administration Guide*.

CIFS (NetWare)

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 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide*.

NCP

NetWare Core Protocol™ (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.

7.5.3 Do Not Run EVMS Tools When Working in iManager or NSSMU

Do not run the Enterprise Volume Management System (EVMS) management tools at the same time that iManager or NSSMU is running. Before you use `evmsgui`, close any instance of iManager or NSSMU, and vice versa.

7.5.4 Management Tools for Managing Pool Snapshots

Pool snapshots for NSS on Linux can be managed in NSSMU or iManager.

Pool snapshots for NSS on NetWare can be managed in iManager. NSSMU for NetWare does not provide a Snapshot option.

7.5.5 Management Capabilities for Software RAIDs

NSSMU for Linux does not support creating nested RAID 10 and 15 devices.

Cluster-Enabling Shared NSS Devices and Pools with Novell Cluster Services

8

Shared Novell® Storage Services™ devices and pools can be used in a cluster environment by using Novell Cluster Services™ on your Novell Open Enterprise Server 2 (OES2) server. The NSS software is not clustered and must be installed and running on every server in the cluster.

- Section 8.1, "Cluster-Enabling NSS Pools and Volumes," on page 93
- Section 8.2, "Guidelines for Cluster-Enabling NSS," on page 93

8.1 Cluster-Enabling NSS Pools and Volumes

For information about installing Novell Cluster Services and cluster-enabling shared NSS devices and pools in the Novell Cluster Services clusters, see the following:

- **Linux:** "Configuring Cluster Resources for Shared NSS Pools and Volumes" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.
- **NetWare:** "Setting Up Cluster Resources for Novell Cluster Services" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*.

8.2 Guidelines for Cluster-Enabling NSS

Novell Cluster Services must already be installed and configured on the server. Table 8-1 provides references for cluster-related tasks for NSS.

Table 8-1 *Clustering Guidelines for NSS*

NSS Feature	Description	Reference
Shared device	Enable the Shareable for Clustering parameter to support high-availability server clusters with Novell Cluster Services.	Section 10.8, "Sharing Devices for NSS Pools," on page 135
Shared pools	Enable the pool for clustering when you create the pool. Devices contributing space to the pool must already be marked as shareable in order to be able to create a shared pool. Unshared pools can be created on shared devices. Pools created on NetWare® can fail over to a Linux node in a mixed-node cluster, but only pools that were originally created on NetWare can fail back from Linux to NetWare.	Section 16.2, "Creating a Pool," on page 198
Multiple Server Activation Prevention (MSAP) for pools	MSAP prevents some accidental activations of a pool on more than one server at a time. MSAP is enabled by default.	Section 16.12, "Preventing Pools from Activating on Multiple Servers," on page 208

NSS Feature	Description	Reference
Pool snapshot	On NetWare, when you take a snapshot of a shared pool, the stored-on location for the snapshot must be the same shared pool. Snapshots of unshared pools can be stored only on unshared pools. The technologies used for pool snapshots are different for NetWare and Linux. Pool snapshots taken on Linux do not work on NetWare, and vice versa. On Linux, NSS does not support using pool snapshots for clustered pools. You must remove any existing pool snapshots for a clustered pool on NetWare before you cluster migrate the pool cluster resource from a NetWare server to a Linux server during a rolling cluster conversion. **WARNING:** You might not be able to open the original pool on Linux if you do not delete the snapshots before you attempt to cluster migrate the pool cluster resource from NetWare to Linux.	“Stored-On Pool (NetWare)” on page 235
Shared volumes	You must create at least one shared volume in a cluster-enabled pool. Typically, all volumes are created when you initially set up the cluster resource and before you need to cluster migrate or fail over the resource to other servers in the cluster. You can add volumes to the pool later by cluster migrating the pool cluster resource back to the original server node in the cluster where the pool was created. Otherwise, you get an eDirectory error because the tools only look for the Pool object under its current server node, and not under the original node where it it was created. To create or modify home directories, Distributed File Services junctions, or any other elements that are managed using eDirectory objects, you must cluster migrate the pool resource back to the node where it was created before you perform those management tasks. This restriction also applies to management tasks like renaming a pool or volume that changes information in the eDirectory objects for the shared pool or volume.	
Shared encrypted volume	When shared pools contain encrypted volumes, you must provide the encryption password the first time that a volume is mounted after a reboot. Thereafter, the nodes in the cluster share the key.	“Sharing Encrypted NSS Volumes in a Cluster” on page 284 “Using Encrypted Volumes in a Server Cluster (Linux)” on page 288 “Using Encrypted Volumes in a Server Cluster (NetWare)” on page 289

NSS Feature	Description	Reference
Shadow volume pair using Dynamic Storage Technology	When a shared pool contains a volume that is part of a shadow volume pair, the other volume in the shadow pair can reside on the same pool or in a different pool on the same server. If it is on a different pool, both pools must be managed by the same cluster resource. In a cluster, shadow volumes can reside on and fail over only to nodes running OES 2 Linux.	*OES 2 SP1: Dynamic Storage Technology Administration Guide*

Management Tools for NSS

9

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

- Section 9.1, “Novell iManager and Storage-Related Plug-Ins,” on page 97
- Section 9.2, “NSS Management Utility (NSSMU) Quick Reference,” on page 112
- Section 9.3, “NSS Commands and Utilities,” on page 117
- Section 9.4, “Novell NetStorage,” on page 118
- Section 9.5, “Novell Remote Manager,” on page 120
- Section 9.6, “Novell Client,” on page 124
- Section 9.7, “ConsoleOne,” on page 124
- Section 9.8, “Virtual File Services, APIs, and Scripts,” on page 125

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

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

9.1 Novell iManager and Storage-Related Plug-Ins

Novell iManager 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.

NOTE: The storage-related plug-ins do not support Mobile iManager.

This section describes the following:

- Section 9.1.1, “Understanding Storage-Related Plug-Ins,” on page 98
- Section 9.1.2, “Accessing Novell iManager,” on page 103
- Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104
- Section 9.1.4, “Selecting a Server to Manage,” on page 104
- Section 9.1.5, “Storage Plug-In Quick Reference,” on page 105
- Section 9.1.6, “Files and Folders Plug-In Quick Reference,” on page 110

9.1.1 Understanding Storage-Related Plug-Ins

Storage-related plug-ins share some management code in common. It is necessary to manage these plug-ins together when installing or updating any of the plug-ins.

- "Storage-Related Plug-Ins for iManager" on page 98
- "Downloading and Installing Plug-In Updates" on page 100
- "Prerequisites for Using the Storage-Related Plug-Ins" on page 100
- "Files for Storage-Related Plug-Ins" on page 102

Storage-Related Plug-Ins for iManager

Table 9-1 identifies the storage-related plug-ins for Novell iManager 2.7 in OES 2 SP1 and later.

Table 9-1 *Storage-Related Plug-Ins for iManager*

Storage-Related Plug-In	NPM File	Role in iManager	Use to Manage
Archive Versioning Management	`arkmgmt.npm`	Archive Versioning	Novell Archive and Version Services
Cluster Services Management	`ncsmgmt.npm`	Clusters	Novell Cluster Services™
Distributed File Services Management	`dfsmgmt.npm`	Distributed File Services	Novell Distributed File Services
AFP Management	`afpmgmt.npm`	File Protocols > AFP	Novell AFP for Linux Novell AFP for NetWare
CIFS Management	`cifsmgmt.npm`	File Protocols > CIFS	Novell CIFS for Linux Novell CIFS for NetWare
NFS Management	`nfsmgmt.npm`	File Protocols > NFS	Novell NFS for NetWare
NSS Management	`nssmgmt.npm`	Storage	Novell Storage Services
Storage Management	`storagemgmt.npm`	No role. Required when using any combination of storage-related plug-ins	Contains common code for all storage-related plug-ins

IMPORTANT: The storage-related plug-ins share code in common in the `storagemgmt.npm` file. 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.

Archive Versioning Management

The Archive Versioning Management (`arkmgmt.npm`) file contains the Archive Versioning plug-in for Novell Archive and Version Services. This plug-in requires the NSS Storage Management plug-in and the Storage Management plug-in.

For information about using this plug-in, see the following:

- *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*
- *OES 2: Novell Archive and Version Services 2.1 for NetWare Administration Guide*

Cluster Services Management

The Cluster Services Management (`ncsmgmt.npm`) file contains the Clustering plug-in for managing Novell Cluster Services. This plug-in requires the Storage Management plug-in. The NSS Storage Management plug-in is required for cluster-enabling NSS pools and volumes.

For information about using this plug-in, see the following:

- *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*
- *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*

Distributed File Services Management

Use the Distributed File Services plug-in to manage Novell DFS for NSS volumes. For information about using the DFS plug-in, see *OES 2: Novell Distributed File Services Administration Guide*.

The DFS plug-in also requires the NSS Storage Management and Storage Management plug-ins.

File Protocols

The File Protocols plug-in for the Native File Access Protocols for NetWare service has been replaced in OES 2 SP1 by three plug-ins: Novell AFP (`afpmgmt.npm`), Novell CIFS (`cifsmgmt.npm`), and NFS for NetWare (`nfsmgmt.npm`). The Novell AFP and Novell CIFS plug-ins support AFP and CIFS services for NSS volumes on both Linux and NetWare.

The three protocols' plug-ins also require the NSS Storage Management (`nssmgmt.npm`) and Storage Management (`storagemgmt.npm`) plug-ins.

IMPORTANT: Make sure to uninstall the existing File Protocols plug-in, NSS plug-in, and Storage Management plug-in before you upgrade to these new plug-ins.

For information about managing these services, see the following guides:

- *OES 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide* (AFP, CIFS, and NFS on NetWare)
- *OES 2 SP1: Novell AFP For Linux Administration Guide*
- *OES 2 SP1: Novell CIFS for Linux Administration Guide*

NSS Management

The Novell Storage Services Management (`nssmgmt.npm`) plug-in allows you to manage NSS services (devices, software RAIDs, pools, and volumes) on NetWare 6.5, OES 1, and OES 2. Information about using this plug-in is the focus of this guide.

Storage Management

The Storage Management (`storagemgmt.npm`) file contains common code that is shared by the storage-related plug-ins. If you use more than one of these storage-related plug-ins, you should install, update, or remove the `storagemgmt.npm` file and all installed storage-related `.npm` files at the same time.

IMPORTANT: If you remove any one of the installed storage-related plug-ins, it removes the common code (`storagemgmt.npm`) file, which breaks the remaining installed plug-ins.

Downloading and Installing Plug-In Updates

For information, see "Downloading and Installing Plug-Ins During Installation" in the *Novell iManager 2.7 Installation Guide*.

Prerequisites for Using the Storage-Related Plug-Ins

The requirements in this section apply to the storage-related plug-ins for iManager 2.7 that are described in Section , "Storage-Related Plug-Ins for iManager," on page 98.

- "EVMS (Linux)" on page 100
- "Web Browser Language Setting" on page 100
- "Protocols for iManager Communications" on page 101

EVMS (Linux)

NSS for Linux uses the Enterprise Volume Management Service (EVMS) version 2.5.5-24.54.5 or later to manage devices that contain NSS pools and volumes and Novell Cluster Services cluster resources. For information about upgrading EVMS, see "EVMS Requirements" on page 46.

NSS tools work transparently with EVMS to manage your devices, pools, and volumes.

IMPORTANT: You cannot run the EVMS graphical user interface (`evmsgui`), ncurses, or console commands at the same time that iManager or NSSMU is running. If you use EVMSGUI to manage NSS devices, close any instance of iManager or NSSMU, and vice versa.

NSS does not recognize devices that are being managed by non-EVMS volume managers. Only unconfigured devices or EVMS-managed devices with free unconfigured space are available to you when you create or expand NSS pools and software RAID devices.

For information about how to make EVMS the volume manager for your existing non-EVMS devices, see Section D.4, "Making Devices Available to EVMS," on page 561.

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 iManager's 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.

Protocols for iManager Communications

Table 9-2 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 that you use must be loaded and running on both the iManager server and the target server that you want to manage.

Table 9-2 *Interoperability of Protocols Used to Connect the iManager Server and Target Servers*

iManager Server Operating Platform	**Protocol Used to Connect to the Target Server Based on Its Operating Platform (* indicates the default)**			
	OES 1 Linux and Later	**OES 1 SP1 NetWare, NetWare 6.5 SP4, and Later**	**OES 1 NetWare and NetWare 6.5 SP3**	**NetWare 6.5 SP2**
OES 1 Linux and Later	* WBEM	* WBEM	WBEM (Start WBEM)	
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
OES 1 SP1 NetWare, NetWare 6.5 SP4, and Later	* WBEM	* WBEM	WBEM (Start WBEM)	
		NCP™	* NCP	* NCP
	CIFS	CIFS	CIFS	CIFS (Field Patch 2B)
OES 1 NetWare and NetWare 6.5 SP3	*WBEM	*WBEM	WBEM (Start WBEM)	
		NCP	* NCP	* NCP
	CIFS	CIFS	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.

For OES 2 Linux, the storage-related plug-ins for iManager require CIMOM connections for tasks that transmit sensitive information (such as a username and password) between iManager and the `_admin` volume on the OES 2 Linux that server you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

If CIMOM is not currently running when you click *OK* or *Finish* for the task that sends the sensitive information, you get an error message explaining that the connection is not secure and that CIMOM must be running before you can perform the task.

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

To check the status of WBEM on Linux:

1 As `root` in a console shell, enter

```
rcowcimomd status
```

To start WBEM on Linux:

1 As `root` in a console shell, enter

```
rcowcimomd start
```

To check the status of WBEM on NetWare:

1 As the Administrator user or equivalent, enter the following at the server console:

```
modules owcimomd
```

To start WBEM on NetWare:

1 As the Administrator user or equivalent, enter the following at the server console:

```
openwbem
```

For information about installing WBEM, see "Setting Up OpenWBEM" in the *OES 2 SP1: OpenWBEM Services Administration Guide*.

CIFS (NetWare)

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 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide*.

NCP

NetWare Core Protocol™ (NCP) is the default protocol when the iManager server and target server are NetWare 6.5 SP3 or SP2.

Files for Storage-Related Plug-Ins

- "File Locations on Linux" on page 102
- "File Locations on NetWare" on page 103
- "Java and Tomcat Files" on page 103

File Locations on Linux

The module files (see Table 9-1 on page 98) are located in the `/var/opt/novell/iManager/nps/portal/modules/` directory.

The Java* JAR files are located in the `/var/opt/novell/iManager/nps/WEB-INF/lib/` directory.

The Tomcat TLD files are located in the `/var/opt/novell/iManager/nps/WEB-INF/` directory.

File Locations on NetWare

The module files (see Table 9-1 on page 98) are located in the `\sys\tomcat\4\webapps\portal\modules\` directory.

The Java JAR files are located in the `\sys\tomcat\4\webapps\WEB-INF\lib` directory.

The Tomcat TLD files are located in the `\sys\tomcat\4\webapps\WEB-INF\` directory.

Java and Tomcat Files

Table 9-3 *Java and Tomcat Files for Storage-Related Plug-Ins*

Storage-Related Plug-In	Java Files	Tomcat Files

Novell AFP	`afpGadgets.jar`
	`afpManageLib.jar`
	`afpTags.jar`
		`afp.tld`

Archive and Version Services	`arkGadgets.jar` `arkManageLib.jar`	`ark.tld`
Novell CIFS	`cifsGadgets.jar` `cifsManageLib.jar`	`cifs.tld`

Novell Cluster Services	`ncsGadgets.jar`
	`ncsManageLib.jar`
	`ncsTags.jar`
		`ncs.tld`

Novell Distributed File Services	`dfsGadgets.jar`
	`dfsManageLib.jar`
	`nasGadgets.jar`
		`dfs.tld`

Novell NFS for NetWare	`nfsGadgets.jar` `nfsManageLib.jar`	`nfs.tld`
Novell Storage Services	`nssGadgets.jar` `nssManageLib.jar`	`nss.tld`

Storage Management (common to all)	`nssAdminClient.jar`
	`nssGadgetLib.jar`
	`nssTags.jar`	

9.1.2 Accessing Novell iManager

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

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

   The URL is case sensitive. 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 Use your administrator username and password to log in 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.

9.1.3 Accessing Roles and Tasks in iManager

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

For information, see Section 9.1.2, "Accessing Novell iManager," on page 103.

2 In *Roles and Tasks,* click the *Storage* role to expand its main tasks

As you work in the storage-related plug-ins, 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 9.1.4, "Selecting a Server to Manage," on page 104.

9.1.4 Selecting a Server to Manage

Before you can access the management options on a selected 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:

- Type the Novell eDirectory distinguished server name for the server you want to manage, then press *Tab* or click somewhere on the page outside of the *Server* field to enter your selection. For example:

  ```
  svr1.company
  ```

- 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.

9.1.5 Storage Plug-In Quick Reference

The Storage role comprises seven key tasks:

- “Pools” on page 105
- “Volumes” on page 106
- “User Quotas” on page 107
- “Partitions” on page 108
- “Software RAIDs” on page 108
- “Devices” on page 109
- “Scan for Devices (NetWare)” on page 110

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 backup and recovery.

Table 9-4 *Pool Management Tasks*

Task	Description	Reference
Pools	Displays a list of all pools on the selected server.	“Viewing Pools on a Server” on page 204
Details	Displays information about a selected pool.	“Viewing Pool Details” on page 205
New	Creates a new pool on the selected server.	“Creating a Pool” on page 198
Delete	Deletes a selected pool and all of its volumes and their data.	“Deleting a Pool” on page 204
Rename	Renames a selected pool.	“Renaming a Pool” on page 203
Activate	Mounts and activates a selected deactive or unmounted pool.	“Activating and Deactivating Pools” on page 201
Deactivate	Deactivates a selected active pool.	“Activating and Deactivating Pools” on page 201
Increase Size	Allows you to select one or more partitions from available devices in order to expand the size of a pool.	“Increasing the Size of a Pool” on page 202

Task	Description	Reference
Snapshot	Opens the Pool Snapshots page where you can create and manage pool snapshots. This option is disabled for online snapshot pools, because you cannot create a snapshot of a snapshot.	“Managing NSS Pool Snapshots” on page 229
Update eDirectory	Updates (replaces) the eDirectory object for a selected pool. Use only if a storage object is not recognized or has been lost.	“Updating eDirectory Pool Objects” on page 213
Deleted Volume	Displays a list of deleted volumes in a pool, and allows you to salvage or purge them. You can also pause and resume the autopurging of deleted volumes.	“Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool” on page 368
Offline	For a pool snapshot that is online as an active pool, takes it offline. This does not delete the pool snapshot.	“Viewing and Managing an Online Pool Snapshot” on page 248
Partitions	Displays a list of the partitions comprising the pool’s storage space.	“Viewing Partition Information for a Pool” on page 206
Volumes	Lists all volumes on a selected pool.	“Viewing Volume Information for a Pool” on page 207
Devices	Displays a list of the devices that contribute space to a selected pool.	“Viewing Device Information for a Pool” on page 207

Volumes

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

Table 9-5 *Volume Management Tasks*

Task	Description	Reference
Volumes	Displays a list of all volumes on the selected server.	“Managing NSS Volumes” on page 253
Details	Displays information about a selected volume.	“Viewing the Details of an NSS Volume” on page 263
New	Creates a new unencrypted volume. To create an encrypted NSS volume, use NSSMU.	“Creating Unencrypted NSS Volumes” on page 259
Delete	Deletes a selected volume and all of its contents.	“Deleting an NSS Volume” on page 278
Rename	Renames a selected volume.	“Renaming an NSS Volume” on page 269

Task	Description	Reference
Activate	Mounts and activates a deactive or unmounted volume.	"Activating and Deactivating an NSS Volume" on page 271
Deactivate	Deactivates an active volume.	"Activating and Deactivating an NSS Volume" on page 271
Mount	Mounts an unmounted volume. A volume must be mounted to view its details.	"Mounting and Dismounting an NSS Volume (Linux)" on page 271
Dismount	Dismounts a mounted volume.	"Mounting and Dismounting an NSS Volume (Linux)" on page 271
Move	Moves a selected NSS volume for the purpose of reorganizing and redistributing storage on the same server (or to other servers) in response to changing business needs.	For requirements, guidelines, and procedures for splitting volumes, see "Using DFS to Move NSS Volumes" in the *OES 2: Novell Distributed File Services Administration Guide*.
Split	Splits a selected NSS volume for the purpose of reorganizing and redistributing storage on the same server (or to other servers) in response to changing business needs.	For requirements, guidelines, and procedures for splitting volumes, see "Using DFS to Split NSS Volumes" in the *OES 2: Novell Distributed File Services Administration Guide*.
Properties	Displays a list of volume attributes, and allows you to set the attributes and name space to use for a volume. It also displays usage statistics for a volume.	"Viewing Properties of an NSS Volume" on page 263
User Quotas	Displays user quotas and space consumed for users of the volume. Administrators can view and manage user quotas. Users can view their own user space quotas.	"Configuring a User Space Quota" on page 356
Offline	Takes a selected snapshot volume offline, where it remains active in the background.	"Viewing and Managing an Online Pool Snapshot" on page 248
Update eDirectory	Updates (replaces) the eDirectory object. Use only if a storage object is not recognized or has been lost.	"Updating eDirectory Volume Objects" on page 262

User Quotas

The User Quotas task (see Table 9-6) allows managers to view and manage user space restrictions. It can be specified as an iManager role-based task for administrators. An individual user can log in to iManager under his or her own username to view the user's user space restrictions.

Table 9-6 *User Quotas Task*

Task	Description	Reference
User Quotas	For the Administrator user, displays quotas and allows the administrator user to manage user quotas for all users. For the user, displays the user's own space restrictions.	"Configuring a User Space Quota" on page 356

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.

Table 9-7 *Partitions Tasks*

Task	Description	Reference
Partitions	Displays a list of all partitions on a selected server.	"Viewing a List of Partitions" on page 159
Details	Displays information about a selected partition.	"Viewing Details for a Partition" on page 160
Label	Adds a label for a selected partition.	"Labeling a Partition" on page 160

Software RAIDs

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

Table 9-8 *Software RAID Management Tasks*

Task	Description	Reference
Software RAIDs	Displays a list of software RAID devices on the selected server.	"Viewing a List of Software RAID Devices on a Server" on page 171
Details	Displays the details of a selected software RAID device.	"Viewing Details of a Software RAID Device" on page 172
New	Creates a new software RAID 0, 1, or 5 device for the selected server.	"Creating Software RAID Devices with iManager" on page 174
Rename	Renames a selected software RAID device.	"Renaming a Software RAID Device" on page 180

Task	Description	Reference
Increase Size	Expands 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. 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.	“Increasing the Size of a Software RAID Device” on page 180
Restripe	Completes a restriping process for a RAID 0 or RAID 5 device that has been paused.	“Restriping a Software RAID 0 or 5 Device” on page 182
Delete	Deletes the selected software RAID device and removes the RAID relationship between member partitions and the underlying storage structures. All data on the member partitions is lost.	“Deleting a Software RAID Device” on page 185
Pools	Lists pools on a selected software RAID devices.	“Viewing Pools on a Software RAID Device” on page 185
Partitions	Lists details about partitions (member segments) in the RAID. In some cases, you can also delete a partition to repair a RAID.	“Viewing Partitions on a Software RAID Device” on page 186

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.

Table 9-9 *Device Management Tasks*

Task	Description	Reference
Devices	Displays a list of all local and external devices available on the selected server.	“Viewing Devices on a Server” on page 131
Details	Displays information about a selected device.	“Viewing Details for a Device” on page 132
Initialize Disk	Initializes a selected device by erasing its partition table, effectively destroying all of its data. If devices are present but not showing up for creating pools and volumes, you should initialize the disk.	“Initializing a Disk” on page 134

Task	Description	Reference
Multipath (NetWare)	For network configurations with multiple paths between network devices and your NetWare server, 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.	"Managing Multipath I/O to Devices (NetWare)" on page 189
Set Default Path	Sets the connection to selected device to its user-defined default primary path.	"Setting the Primary Path for a Device to Its Default Path" on page 194
Reset Registry	Resets the multipath priority settings for a selected device in the server registry to its user-defined defaults.	"Resetting the Server Registry with Default Priority Settings for a Device" on page 194
Shareable for Clustering	Enables device sharing to support high-availability server clusters.	"Sharing Devices for NSS Pools" on page 135
Pools	Displays a list of the pools on a device.	"Viewing Pools on a Device" on page 138
Partitions	Displays information about partitions that are configured on a device.	"Viewing Partitions on a Device" on page 136

Scan for Devices (NetWare)

You can scan for devices on a NetWare server that were not found automatically by the Media Manager.

Table 9-10 *Scan for Devices Task*

Task	Description	Reference
Scan	After you select a server to manage, scans for devices added since you booted the server.	"Scanning for Devices (NetWare)" on page 133

9.1.6 Files and Folders Plug-In Quick Reference

The Files and Folders plug-in for iManager 2.7 provides the Files and Folders role for Linux and NetWare. It is also integrated in iManager as the *View Objects* option in the iManager toolbar. File browsing in iManager is available for file systems that have a Volume object defined in eDirectory, such as for NSS volumes on Linux and NetWare and for NCP volumes on Linux.

The Files and Folders Manager NPM file (`filemanager.npm`) is automatically installed in iManager. For information about manually installing NPM files for iManager, see the *Novell iManager 2.7 Installation Guide*.

Click the *Files and Folders* role to select tasks first, then search for the file or folder you want to manage. Click the *View Objects* icon to view the Tree, Browse, and Search view of a server's eDirectory objects in the left pane. In the Tree view, click a Volume object to see the hierarchical file system tree view of the volume's folders and files. Click the plus (+) or minus (-) icon next to a directory name to expand or collapse the view of its subdirectories. Locate the file or folder you want to manage, then specify the action you want to perform for it.

The Files and Folders plug-in for Novell iManager 2.7 provides the tasks described in this section. All of the tasks and actions that are available under the *Files and Folders* role are also available from the *View Objects* tree view.

- "Delete" on page 111
- "Deleted Files" on page 111
- "Download" on page 111
- "New Folder" on page 111
- "Properties" on page 111
- "Upload" on page 112

Delete

Deletes a file or folder on an NSS volume or an NCP volume (NCP share on a Linux POSIX file system). For information, see Section 28.2, "Deleting a File or Folder on an NSS Volume," on page 388.

Deleted Files

Salvages or purges deleted files on an NSS volume. Salvage and purge of deleted files and directories is available only for NSS volumes where the volume's Salvage attribute is enabled. Other NSS settings determine how long deleted files and directories are available.

For information about configuring salvage and purge behavior for NSS volumes, see Chapter 26, "Salvaging and Purging Deleted Volumes, Directories, and Files," on page 361.

For non-NSS volumes, the Deleted Files report is empty (no deleted files).

Download

Downloads a selected file from an NSS volume or NCP volume to a specified location on your local drive or mapped network drive. For information, see Section 28.4, "Downloading Files from an NSS Volume," on page 390.

New Folder

Creates a folder on an NSS volume or NCP volume. For information, see Section 28.1, "Creating a Folder on an NSS Volume," on page 387.

Properties

Adds, removes, or modifies file system trustees, trustee rights, inherited rights filters, and file system attributes for files and folders on NSS volumes and NCP volumes. See Table 9-11 for a complete list of tasks you can perform from the Properties page.

Table 9-11 Properties Tasks

Properties Tab	Task Description
Information	Displays or modifies a directory quota for the selected folder. Directory quotas management is available only for NSS volumes where the volume's Directory Quotas attribute is enabled. For information, see Section 25.3, "Managing Directory Quotas," on page 348. Displays information about a selected file or folder, such as: ◆ Current size ◆ Time stamps for when the file was created, modified, accessed, and archived Displays or modifies the owner of a selected file or folder. Displays or modifies the file system attributes for a file or folder. For information, see Section 22.1.3, "Configuring File or Folder Attributes," on page 298.
Rights	Displays, adds, or removes file system trustees for a selected file or directory. Displays, grants, or revokes file system trustee rights for trustees of the selected file or directory. Displays or modifies the inherited rights filter for a selected file or directory. For information, see Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297.
Inherited Rights	Displays or modifies the inherited rights filters at every level of the path for a selected file or directory. For information, see "Configuring the Inherited Rights Filter for a File or Directory" on page 302. Displays the effective rights for the selected file or directory. For information, see Section 22.1.5, "Viewing Effective Rights for a Trustee," on page 303.

Upload

Uploads a specified file from your local drive or a mapped network drive to a specified location on an NSS volume or NCP volume. For information, see Section 28.3, "Uploading Files to an NSS Volume," on page 389.

9.2 NSS Management Utility (NSSMU) Quick Reference

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 9.2.1, "NSSMU for Linux Quick Reference," on page 113
- Section 9.2.2, "NSSMU for NetWare Quick Reference," on page 115

9.2.1 NSSMU for Linux Quick Reference

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 `nssmu` from a terminal console as the `root` user.

Table 9-12 identifies key functions available in NSSMU for Linux. This quick reference is also available in the `nssmu(8)` man page. To access the man page, enter the following at a terminal console prompt:

```
man 8 nssmu
```

Table 9-12 *Summary of Management Options in NSSMU for Linux*

Management Options	Description
Devices *F3* = Initialize EVMS-managed device (Do not initialize your system device.) *F5* = Refresh display *F6* = Share (shareable/not shareable for clustering) *Enter* = Show partitions *Esc* = Previous menu	Use this option to initialize and maintain physical storage devices and software RAID devices available to this server. Use the Software RAID Devices option to create, repair, or delete RAIDs. Initialize the selected device by erasing its partition table, effectively destroying all of its data. If devices are present but not showing up for creating pools and volumes, you should initialize the disk.
Partitions *Ins* = Create an NSS partition (disabled) *Del* = Delete an NSS partition *F3* = Mirror partition (that contains an existing pool) *F5* = Refresh details of the partition *F6* = Label *Enter* =Show volumes *Esc* = Previous menu	Use this option to display details about partitions. All types of partitions are displayed, including those for Linux file systems. The Create option is disabled. NSS partitions are automatically created for you as you define NSS pools or software RAIDs. You can delete a single partition at a time when repairing a failed software RAID partition. To delete all partitions for a software RAID, you should delete the RAID itself from the Software RAIDs page; otherwise, the RAID is not cleanly deleted. The Mirror option lets you specify 1 to 3 partitions to mirror an existing partition that contains an NSS pool. Effectively, you are creating a RAID1 mirror device for the pool. Each segment of the defined RAID is a complete mirror of the original pool and is the same size as the original partition. After you mirror the partition, manage the RAID from the Software RAIDs page.

Management Options	Description
Pools *Ins* = Create a pool *Del* = Delete a pool *F3* = Expand a pool (by adding space) *F4* = Update NDS™/eDirectory *F5* = Refresh details of a pool *F6* = Rename a pool *F7* = Activate/deactivate a pool *F8* = More (list more options) *F9* = Show/Hide deleted volumes (then salvage, purge, or pause/resume autopurging) *Enter* = Show volumes for a pool *Esc* = Previous menu	Use this option to create, delete, rename, and expand NSS storage pools to efficiently use all free space in the available devices. After you create a pool, you can expand it by adding free space from the same or different device to increase its size. Select from the available free space to allocate it to the pool. Each device can contribute a different amount of space to the pool. Devices that contribute space must be in the same share state as the pool, that is, *Shared* or *Not Shared*. You can increase the size of a pool, but you cannot reduce it.
Volumes *Ins* = Create a new volume *Del* = Delete a volume *F2* = Rename mount point for the volume (new path with volume name) *F3* = Rename volume *F4* = Update NDS/eDirectory *F5* = Refresh details of the volume *F6* = View compression statistics *F7* = Dismount/mount a volume. If it is encrypted, the volume prompts for a password on the first mount after a system boot or reboot. *F8* = More (list more options) *F9* = Name Space - choose Long (default), UNIX, DOS, or Macintosh *Enter* = Set or view volume properties *Esc* = Previous menu	Use this option to create, delete, rename, activate/deactivate, and mount/dismount NSS volumes and to set their attributes. 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 NICI key for encryption. You must supply the password on the first volume activate after a system reboot. 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.
RAID Devices *Ins* = Create a software RAID *Del* = Delete a software RAID device *F3* = Expand a RAID device (add partitions) *F4* = Rename a RAID device *F5* = Refresh details of the software RAID device *F6* = Restripe (resume restriping for paused RAID 0) *Enter* = Show segments (list member partitions for selected device) *Esc* = Previous menu	Use this option to create and manage NSS software RAID devices. 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.

Management Options	Description
Snapshot *Ins* = Create a pool snapshot *Del* = Delete a pool snapshot *F5* = Refresh display *F7* = Mount or dismount the pool snapshot as an active pool. The snapshot functions continue whether the snapshot is mounted or dismounted. *Esc* = Previous menu	Use this option to create, delete, mount, and dismount pool snapshots for NSS pools. Up to 15 snapshots of a pool are allowed to concurrently coexist. On Linux, snapshots are stored on a separate partition that you specify, not another pool. After it is created, the partition for the snapshot pool cannot be expanded.

9.2.2 NSSMU for NetWare Quick Reference

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 RAIDs 0, 1, 5, 10, and 15), partitions (NSS and iSCSI), pools, and volumes.

To set up NetWare iSCSI partitions, use the *Partitions* option in NSSMU. You must have NetWare 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 *OES 2 SP 1: iSCSI 1.1.3 for NetWare Administration Guide*.

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
```

IMPORTANT: Use iManager to manage pool snapshots on NetWare. With a server selected, go to *Storage* > *Pools*, select the pool, then click *Snapshots*.

The following table identifies key functions available in NSSMU for NetWare. This quick reference is available as help by pressing *F1* while working in NSSMU. Help is localized according to the languages supported by the OES 2 release.

Table 9-13 *Summary of Management Options in NSSMU for NetWare*

Function Keys	Description
Devices *F1* = Help *F2* = Scan for devices *F3* = Initialize device (do not use for the `sys` pool or `sys:` volume) *F4* = Multipath (available only if there are multiple paths for the device) *F5* = Select/deselect one or more devices for Initialize or Share actions *F6* = Share (shareable/not shareable for clustering) *F9* = Refresh details of the device *Enter* = Show partitions	Use this option to initialize and maintain a wide selection of physical storage devices and software RAID devices available to this server. Use the Software RAID Devices option to create, repair, or delete RAIDs. Initialize the selected device by erasing its partition table, effectively destroying all of its data. If devices are present but not showing up for creating pools and volumes, you should initialize the disk.
Pools *Ins* = Create a pool *Del* = Delete a pool *F3* = Expand a pool (add partitions) *F4* = Update NDS/eDirectory *F5* = Refresh details of the pool *F6* = Rename a pool *F7* = Activate/deactivate pools *F8* = More (list more options) *F9* = Show/hide deleted volumes (then salvage, purge, or pause/resume autopurging) *F10* = List devices (that a pool resides on) *Alt+F8* = Update NDS/eDirectory *Enter* = Show volumes (on a server or pool)	Use this option to create, delete, rename, and expand NSS storage pools to efficiently use all free space in the available devices.
Volumes *Ins* = Create an encrypted or unencrypted NSS volume *Del* = Delete a volume *F3* = Rename a volume *F5* = Refresh details of the volume *F6* = Deactivate/activate the volume. If it is encrypted, the volume prompts for a password on the first activation after a system boot or reboot. *F7* = Dismount/mount a volume *F8* = More (list more options) *F9* = Name Space - choose Long (default), UNIX, DOS, or Macintosh *Alt+F8* = Update NDS/eDirectory *Enter* = Properties (set volume properties)	Use this menu option to create, delete, rename, activate/deactivate, and mount/dismount volumes and to set volume attributes. 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 NICI key for encryption. Supply the password on the first volume activate after a system boot or reboot.

Function Keys	Description
Software RAIDs *Ins* = Create a software RAID (0, 1, or 5) *Del* = Delete a software RAID device *F3* = Expand a RAID device (add segments) *F5* = Refresh details of the software RAID device *F6* = Restripe (resume restriping for paused RAID 0 or 5 device) *F8* = Show pools (list pools on selected RAID 0 or 5; for RAID 1, its single member pool appears in device details) *Enter* = Show segments (list member partitions for selected device)	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. Create a software RAID 10 device by mirroring software RAID 0 devices. Create a software RAID 15 device by mirroring software RAID 5 devices.
Partitions *Ins* = Create an NSS or an iSCSI partition *Del* = Delete an NSS or an iSCSI partition *F3* = Mirror (create a software RAID-1 device) *F5* = Refresh details of the partition *Enter* =Show volumes	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.

9.3 NSS Commands and Utilities

Command line instructions and utilities are available to control most NSS functions.

- Section 9.3.1, “Command Consoles,” on page 117
- Section 9.3.2, “NSS Commands,” on page 118
- Section 9.3.3, “NSS Utilities,” on page 118

9.3.1 Command Consoles

NSS commands and utilities are issued from command line interfaces referred to as consoles in this guide.

- “Linux” on page 117
- “NetWare” on page 118

Linux

On Linux, all NSS commands and utilities are issued while logged in as the `root` user or a user with equivalent privileges.

Linux Terminal Console

NSS utilities for Linux are issued at the command prompt of a Linux terminal console.

If you are not running a graphical interface on the server, the terminal console is simply the command line prompt displayed when you log in to the server.

If you are using a graphical interface, you can open a terminal console by using one of these methods:

- Right-click on the Linux desktop, then select *Open Terminal* from the menu.
- From the Linux desktop, click the *Computer* menu, then select *Terminal (Command Line Terminal)* from the *Applications* menu.

NSS Console

The NSS Console (NSSCON, `nsscon(8)`) utility for 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.

To start NSSCON, enter the following at a terminal console prompt:

```
nsscon
```

For more information, see Section B.13, “NSSCON (Linux),” on page 519.

NetWare

All NSS commands and utilities can be issued from the NetWare server console, except those specified as Linux only. Use the server console to monitor NSS activity through console messages.

9.3.2 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 on Linux:

```
nss /help
```

For information about NSS commands, see Appendix A, “NSS Commands,” on page 451.

9.3.3 NSS Utilities

The most well-known NSS utilities are the NSS Management Utility (`nssmu`) for managing storage and the NSS Console (`nsscon`) utility for issuing NSS commands on Linux. NSS provides other utilities to support more complex command line management tasks for NSS pools and volumes.

For information about NSS utilities, see Appendix B, “NSS Utilities,” on page 499.

9.4 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.

- Section 9.4.1, “Prerequisites,” on page 119
- Section 9.4.2, “Accessing NetStorage,” on page 119
- Section 9.4.3, “Configuring File System Trustees, Trustee Rights, and Attributes,” on page 119
- Section 9.4.4, “Purging and Salvaging Deleted Files,” on page 120

- Section 9.4.5, “Browsing Directories and Files,” on page 120
- Section 9.4.6, “Additional Information,” on page 120

9.4.1 Prerequisites

For NSS on Linux, users must be Linux-enabled with Linux User Management in order to use NetStorage. For information about installing and configuring Linux User Management and enabling users and groups for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

9.4.2 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 9.1.2, “Accessing Novell iManager,” on page 103.
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).

9.4.3 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 by using the *NetWare Info* tab and *NetWare Rights* tab in the *Properties* dialog box.

IMPORTANT: The label of Netware refers to the NetWare Core Protocol (NCP) that is used for trustee management. Use the option for both NSS volumes on both Linux and NetWare.

For information about file system trustees, trustee rights, and attributes for directories and files on NSS Volumes, see the *OES 2 SP1: File Systems Management Guide*.

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.

9.4.4 Purging and Salvaging Deleted Files

Using NetStorage, you can purge and possibly undelete (salvage) NSS files that were previously deleted. For information, see Section 26.6.1, “Using NetStorage,” on page 371.

9.4.5 Browsing Directories and Files

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

9.4.6 Additional Information

For information, see:

- *OES 2 SP1: File Systems Management Guide*
- *OES 2 SP1: NetStorage for Linux Administration Guide*
- *OES 2: NetStorage for NetWare Administration Guide*

9.5 Novell Remote Manager

Novell Remote Manager (NRM) is a browser-based management utility for monitoring server health, changing the configuration of your server, or performing diagnostic and debugging tasks. NRM also allows you to create NSS pools and volumes and to manage some capabilities for NSS volumes on Linux and NetWare.

- Section 9.5.1, “Prerequisites for Using Novell Remote Manager,” on page 120
- Section 9.5.2, “Novell Remote Manager for Linux,” on page 122
- Section 9.5.3, “Novell Remote Manager for NetWare,” on page 122
- Section 9.5.4, “Accessing Novell Remote Manager,” on page 123
- Section 9.5.5, “Starting, Stopping, or Restarting Novell Remote Manager on Linux,” on page 123

9.5.1 Prerequisites for Using Novell Remote Manager

- “Prerequisites for Remote Administration” on page 121

- “Prerequisites for Administrator User Access on Linux Servers” on page 121
- “Prerequisite for Administrator User Access on NetWare Servers” on page 121

Prerequisites for Remote Administration

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

Prerequisites for Administrator User Access on Linux Servers

You can log into Novell Remote Manager for Linux as the `root` user or equivalent for the OES Linux server you are managing.

You can alternately log in to Novell Remote Manager with your eDirectory credentials if you first enable Linux User Management (LUM) in your eDirectory tree and install and configure LUM on the target server. The Administrator user or equivalent must be Linux-enabled and at least one of the following conditions must be met:

- The Administrator user (or equivalent user) must be associated to the eDirectory group that has the Supervisor right for the Entry Rights property for the UNIX Workstation object in eDirectory.
- The Administrator user (or equivalent user) must have the Supervisor right for the Entry Rights property to the NCP object that represents the Linux server in the eDirectory tree.

To see if a user is Linux-enabled, go to iManager, select the User role, then select the user to see if the following is true:

- The user has a *Linux Profile* tab on the Modify User page in iManager.
- The user’s eDirectory object is associated with the UNIX Workstation object that represents the Linux server.

For information about configuring Linux User Management and enabling users for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

Prerequisite for Administrator 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 Administrator user or a user with rights equivalent to the Administrator user.

9.5.2 Novell Remote Manager for Linux

Novell Remote Manager for Linux allows you to browse NSS volumes on your Linux servers. It requires that the NCP Server and NCP Server plug-in for Novell Remote Manager be installed and running.

Tasks

The NCP Server plug-in supports the following tasks:

- Managing connections to NSS volumes and viewing open files for a connection.

  For information, see "Managing Connections for NCP Volumes and NSS Volumes" in the *OES 2 SP1: NCP Server for Linux Administration Guide*.

- Creating or managing shadow volumes with NSS volumes as the primary and secondary storage areas.

  For information, see the *OES 2 SP1: Dynamic Storage Technology Administration Guide*.

Novell Remote Manager for Linux does not support the following tasks for NSS on Linux:

- Configuring directory quotas
- Salvaging and purging deleted files and directories
- Configuring file system trustees and attributes for directories and files
- Creating and managing partitions, pools, and volumes

Additional Information

For detailed information about Novell Remote Manager on Linux, see the *OES 2 SP1: Novell Remote Manager for Linux Administration Guide*.

9.5.3 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.

Tasks

Novell Remote Manager for NetWare supports the following tasks for managing NSS pools and volumes on NetWare servers:

- Configuring directory quotas.

  This requires that the Directory Quotas attribute be enabled on the NSS volume. For information, see Section 25.3, "Managing Directory Quotas," on page 348.

- Salvaging and purging deleted files and directories.

  This requires that the Salvage attribute be enabled on the NSS volume. For information, see "Salvaging and Purging Deleted Volumes, Directories, and Files" on page 361.

- Configuring file system trustees, trustee rights, inherited rights filter, and file and folder attributes.

For information, see the *OES 2 SP1: File Systems Management Guide*.

- Creating and managing some features for devices, partitions, pools, and volumes.

  For full management and feature support, use NSSMU and iManager to manage storage.

- Managing connections to NSS volumes and viewing open files for a connection.

  For information, see "Managing Connections to the Server" in the *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*.

Additional Information

For detailed information, see the *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*.

9.5.4 Accessing Novell Remote Manager

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 it if the certificate is valid.

3. When the Login page appears, do one of the following:

   - **Linux:** Type the username and password of the `root` user for that server, or type the username and password of the Administrator user (or equivalent user) who is an eDirectory user and who has been Linux-enabled.
   - **NetWare:** Type the username and password of the Administrator user or equivalent.

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.

9.5.5 Starting, Stopping, or Restarting Novell Remote Manager on Linux

Novell Remote Manager on Linux is installed and runs by default. If it hangs, you can use the `/etc/init.d/novell-httpstkd` script to get status or to stop, start, or restart `httpstkd`. For the latest information about `httpstkd`, see "Starting or Stopping HTTPSTKD" in the *OES 2 SP1: Novell Remote Manager for Linux Administration Guide*.

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, enter the command for the task you need to perform:

Task	Command
Status	`rcnovell-httpstkd status`
Start	`rcnovell-httpstkd start`

Task	Command
Stop	`rcnovell-httpstkd stop`
Restart	`rcnovell-httpstkd restart`

9.6 Novell Client

- Section 9.6.1, “Novell Client for Linux,” on page 124
- Section 9.6.2, “Novell Client for Windows XP/2003 and Vista,” on page 124

9.6.1 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 2.0 SP1 for Linux Administration Guide*.

9.6.2 Novell Client for Windows XP/2003 and Vista

In combination with NCP Server, the Novell Client for Windows XP/2003 and the Novell Client for Vista support 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 4.91 SP5 for Windows XP/2003 Installation and Administration Guide* and the *Novell Client SP1 for Windows Vista Administration Guide*.

9.7 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 *OES 2 SP1: File Systems Management Guide*.

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

2. Copy these two `.jar` files to your local `resources` directory:

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

3. Copy the following `.jar` file to your local `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 1.3.x User Guide*.

9.8 Virtual File Services, APIs, and Scripts

Virtual File Services (VFS) provides methods that allow you to manage services such as NSS by 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, 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* (http://developer.novell.com/wiki/index.php/Virtual_File_Services_for_NetWare) in the Novell Developers Kit (NDK) documentation.

For NetWare, other APIs are available in the *NDK: Novell Storage Architecture Component (Media Manager and NWPA)* (http://developer.novell.com/wiki/index.php/Storage_Architecture_Components_%28Media_Manager_and_NWPA%29).

Managing Devices

10

Novell® Storage Services™ is the default file system for NetWare® 6 and later. This section describes how to manage devices where you want to create or manage NSS storage objects.

- Section 10.1, “Understanding Devices,” on page 127
- Section 10.2, “Making Devices Available to EVMS (Linux),” on page 131
- Section 10.3, “Viewing Devices on a Server,” on page 131
- Section 10.4, “Viewing Details for a Device,” on page 132
- Section 10.5, “Scanning for Devices (NetWare),” on page 133
- Section 10.6, “Scanning for Devices (Linux),” on page 134
- Section 10.7, “Initializing a Disk,” on page 134
- Section 10.8, “Sharing Devices for NSS Pools,” on page 135
- Section 10.9, “Viewing Partitions on a Device,” on page 136
- Section 10.10, “Viewing Pools on a Device,” on page 138
- Section 10.11, “Enabling Write-Through Cache Management on SCSI Devices and RAID Controllers (Linux),” on page 138
- Section 10.12, “What’s Next,” on page 139

10.1 Understanding Devices

A block storage device is the physical, logical, or virtual storage media available to a server. A device can be directly attached to the server or connected via storage networking protocols such as Fibre Channel and iSCSI.

To manage iSCSI devices on your NetWare server, the server must run iSCSI target software and be the server that you plan to make your iSCSI disk server. For information about managing and using iSCSI devices, see the *OES 2 SP 1: iSCSI 1.1.3 for NetWare Administration Guide*.

- Section 10.1.1, “Device Size Limit,” on page 127
- Section 10.1.2, “Device Types,” on page 128

10.1.1 Device Size Limit

NSS can recognize physical, logical, or virtual devices up to 2 TB in size (where 1 TB = 2E40 bytes = 1,099,511,627,776 bytes). The restriction for NSS is the size that the device reports to the operating system. If the a device’s size is larger than 2 TB, NSS cannot see the device. You must use the storage vendor’s or third-party disk carving tools to carve the device into logical devices (such as LUNs) that are each up to 2 TB in size.

Different manufacturers report device sizes differently. The actual device size varies with the hardware design and the applications and software drivers that manage the device. Many vendors report sizes using a definition where 1 TB = 10E12 bytes = 1,000,000,000,000 bytes. Space can also be consumed by metadata that is added to manage the device. The location on the device where the metadata is stored can also vary by hardware manufacturer and software vendor. After you format

the drive, yet another size might be reported. Third-party product documentation might state the maximum size limits of devices it supports before or after making accommodations for any management data or space lost to formatting. The size of devices you ultimately carve out for use with NSS depends on all these factors.

IMPORTANT: Make sure to refer to the documentation of the device manufacturer, application vendor, and software driver vendor for other limitations on the device size.

If you combine physical partitions or disks to represent a single device to the server, such as for RAIDs, the resulting device must be less than 2 TB in size in order for NSS to see it. For example, if you create a RAID device that is larger than 2 TB in size, NSS cannot recognize the device even if its member devices are each smaller than 2 TB. If the RAID size is larger than 2 TB, you must carve it into multiple LUNs or logical devices of up to 2 TB each in order for NSS to recognize them.

Do not expand devices beyond the 2 TB limit if they contain NSS pools and volumes. On NetWare, the devices are not recognized and you cannot access the data on them. On Linux, devices larger than 2 TB are recognized by the Linux volume manager, but NSS cannot see them.

WARNING: Attempting to expand any of the devices that contribute space to an NSS pool beyond 2 TB in size can result in data loss on the associated NSS pool.

10.1.2 Device Types

The following are examples of common types of devices that are subject to the 2 TB maximum device size:

- "Server Disks" on page 128
- "Direct-Attached Storage Devices" on page 128
- "LUN Devices" on page 129
- "iSCSI Devices" on page 129
- "RAID Devices" on page 129
- "Multipath Devices" on page 130
- "Removable Media" on page 130
- "Virtual Disks" on page 130

Server Disks

Server disks include physical disks on the server or logical disks carved from the server disk.

Direct-Attached Storage Devices

Physical or logical disks can be directly attached to the server as individual devices or in a storage array.

LUN Devices

A LUN (logical unit number) can be either a physical or a logic disk drive. NetWare does not differentiate between the two, and all LUNs are treated as physical disk drives. Refer to the iSCSI SAN or Fibre Channel SAN documentation for information about creating and managing LUNs for your SAN implementation.

A metaLUN is a controller-managed group of multiple LUNs or of multiple hardware RAIDs that are striped or concatenated together to be presented as a single LUN device to the server. Refer to the hardware manufacturer's documentation for information about creating metaLUNs.

iSCSI Devices

An iSCSI device is a remote target disk or tape drive on an iSCSI disk server that is made available across an IP network by iSCSI initiator software running on the server. After connecting to the disk server, you can view the devices in the *Devices* list and add NSS pools and volumes as you would with any device.

For information about managing and using iSCSI devices on NetWare, see the *OES 2 SP 1: iSCSI 1.1.3 for NetWare Administration Guide*.

For information about managing and using iSCSI devices on OES 2 Linux, see "Setting Up an iSCSI Target" and "Configuring iSCSI Initiator" in the *SLES 10 SP2 Installation and Administration Guide*. See also the Linux iSCSI Project (http://linux-iscsi.sourceforge.net/).

RAID Devices

A RAID (redundant array of independent disks) is a logical device that combines space from multiple devices by using special hardware, software, or both. Data is striped or replicated across all member devices to improve data reliability, increase I/O performance, or provide device fault tolerance. All RAID types require configuration using a RAID management tool made for the specific hardware or software used in the RAID.

Hardware RAID Devices

In a hardware RAID, the RAID functionality and management are in firmware within the storage cabinet. Refer to the hardware manufacturer's documentation for information about creating hardware RAIDs.

Controller RAID Devices

Controller RAID devices are also known as *BIOS RAIDs*, *fakeRAIDs*, *hostRAID*, and *quasi-hardware RAIDs*.

In controller RAIDs, the functionality and management are in the HBA or controller BIOS/firmware. If the controller does not contain an on-board CPU resource to use for RAID management, the controller RAID consumes server CPU resources to manage the RAID.

Refer to the hardware manufacturer's documentation for information about configuring Controller RAID devices. For information about using Controller RAIDs with OES 2 Linux, see *TID 3626577: BIOS RAID Support* (http://www.novell.com/support/search.do?cmd=displayKC&docType=kc&externalId=3626577&sliceId=SAL_Public&dialogID=44924286&stateId=0%200%2044926638) in the Novell Support Knowledgebase.

Software RAID Devices

Software RAIDs are controlled by special software in the server's OS such as in the HBA driver or in upper level module such as NSS. Software RAIDs consume CPU resources to manage the RAID.

NSS software RAIDs are supported on Linux and NetWare. For information about creating and managing NSS software RAIDs, see Chapter 13, “Managing NSS Software RAID Devices,” on page 165.

On Linux, you can optionally use Linux tools to create and manage Linux software RAIDs. Linux software RAIDs must be managed by the EVMS volume manager and have a Cluster Segment Manager or a NetWare Segment Manager on it in order to be visible to NSS. For information about using Linux tools to create Linux software RAIDs on OES 2 Linux, see the *SLES 10 SP2: Storage Administration Guide*.

Multipath Devices

If there are multiple connection paths between a device’s hardware controller and the server, each path presents a given device to the server as a separate device. You must use a multipath management tool to resolve the multiple apparent devices to a single multipath device. Use the multipath device UUID or alias when you are creating NSS pools and volumes. Multipath tools also provide automatic path management for path failover, failback, and reconfiguration.

On NetWare, the Media Manager multipath I/O service automatically provides multipath resolution and management for devices with multiple connection paths. You can specify priorities for the primary and failover paths to use. For information, see Chapter 14, “Managing Multipath I/O to Devices (NetWare),” on page 189.

On Linux, use Linux multipath I/O tools to create the multipath device. Afterwards, you must configure the multipath device to be managed by EVMS and add a Cluster Segment Manager or NetWare Segment Manager on it in order for it to be recognized by NSS. For information, see Chapter 15, “Managing Multipath I/O to Devices (Linux),” on page 195 and the *SLES 10 SP2: Storage Administration Guide*.

Removable Media

Removable media devices include CDs, DVDs, or CD/DVD image files.

On NetWare, removable media are mounted as NSS devices. For information about managing removable media on NetWare, see Chapter 21, “Managing Removable Media (NetWare),” on page 291.

On Linux, removable media are mounted as Linux POSIX file systems. Use Linux native tools to manage removable media on Linux.

Virtual Disks

In a Xen virtual environment, you use the Virtual Machine Manager in YaST to allocate storage devices from the host to the virtual machine. The devices that you want to use for the NSS file system on the guest operating system cannot exceed the 2 TB limit, even if the host operating system and guest operating system can handle larger devices. For information about storage considerations in virtual environment, see Chapter 6, “Using NSS in a Virtualization Environment,” on page 79.

10.2 Making Devices Available to EVMS (Linux)

You can view and manage devices for NSS on Linux with the Storage plug-in to iManager if they are managed by the Enterprise Volume Management System (EVMS). For information about making the devices where you want to use NSS available to EVMS, see Appendix D, "Using EVMS to Manage Devices with NSS Volumes (Linux)," on page 557.

For information about management capabilities that are not available in NSS when the device is managed with a non-EVMS volume manager, see Section E.1, "FAQs About Using EVMS with NSS," on page 569.

10.3 Viewing Devices on a Server

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

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   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 109.

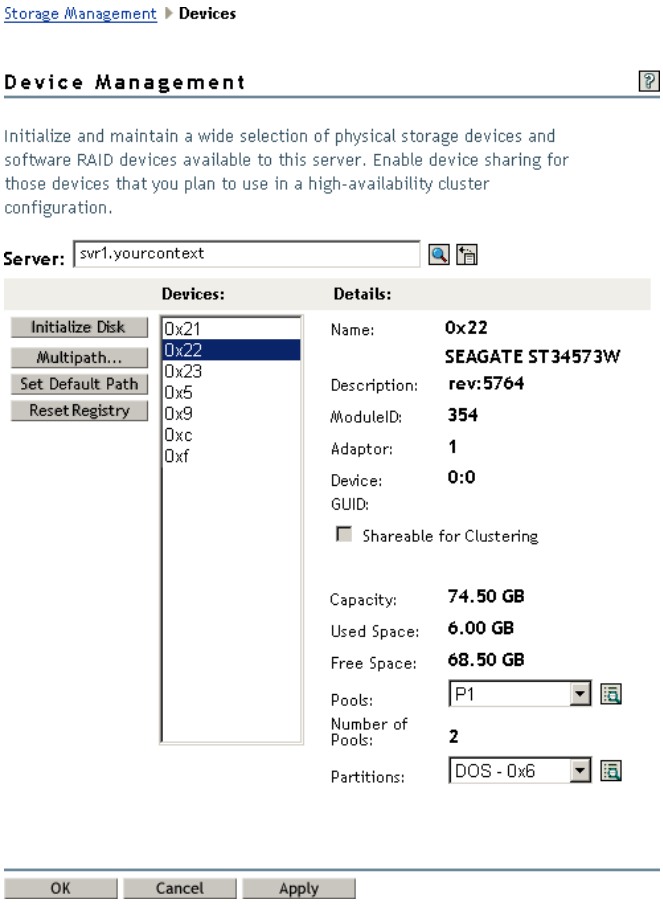

10.4 Viewing Details for a Device

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

Table 10-1 *Explanation of Device Details*

Device Detail	Description
Name	The device name assigned by NetWare.
Description	The physical description of the storage object. This is the physical device identification. For software RAIDs, the description might be RAID0, RAID1, or RAID5.
Module ID	The Novell driver identification number for the module, which indicates what type of driver you are using.
Adapter	The sequential number registered for the adapter. The number indicates which drive is on which adapter.
Device	The server address of the selected device. 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. 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).
GUID	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.
Shareable for Clustering	The attribute of a device that indicates whether the selected device can be shared by multiple computers in a cluster solution.
Capacity	The total available storage space of the selected device.
Used Space	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.
Free Space	The total amount of space on the device that is currently not in use.
Pools	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 *View Details* icon to go to the *Pools* page for that pool.
Number of Pools	The total number of pools that use this device.
Partitions	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 *View Details* icon to go to the Partition Information page for that partition.
Mirror Status	For a RAID1 device, this field shows its status: ◆ **In Sync:** The mirror group is fully synchronized. ◆ **Partial Sync:** The mirror group is only partially synchronized. ◆ **Not Mirrored:** The device is not mirrored (only one partition).

To view a device's details:

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

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   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.

10.5 Scanning for Devices (NetWare)

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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2. Select a NetWare server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   The browser opens to the Scan Devices Page.

   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.

   You can only view the list. 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.

10.6 Scanning for Devices (Linux)

If you add more disks to the SAN, use the following procedure to scan the devices and make them available without rebooting the system:

1. On the storage subsystem, use the vendor's tools to allocate the devices and update its access control settings to allow the Linux system access to the new storage. Refer to the vendor's documentation for details.
2. On the Linux system, use the HBA driver commands to scan the SAN to discover the new devices. The exact commands depend on the driver.

   For example, for a QLogic 2300 HBA, the command is

   ```
   echo scsi-qlascan >/proc/scsi/qla2xxx/<host number>
   ```

   At this point, the newly added device is not known to the higher layers of the Linux kernel's SCSI subsystem and is not yet usable.
3. Scan all targets for a host to make its new device known to the middle layer of the Linux kernel's SCSI subsystem. At a terminal console prompt, enter

   ```
   echo "- - -" >/sys/class/scsi_host/host<hostnumber>/scan
   ```
4. If the devices have multiple paths, run the Multipath tool to recognize the devices for Device Mapper Multipath I/O (DM-MPIO) configuration. At a terminal console prompt, enter

   ```
   multipath
   ```

   For information about configuring multipathing for devices on your Linux server, see "Managing Multipath I/O for Devices" in the *SLES 10 SP2: Storage Administration Guide*.

10.7 Initializing a Disk

If you can see a device listed in the *Devices* list, but the device is not available for creating pools and volumes, you probably need to initialize the disk.

On the *Devices* 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* process also deletes data on all of the partitions of the entire pool, volume, or device, even if they reside on separate devices.

WARNING: For NetWare, do not initialize the device that contains your `sys:` volume. Initializing the `sys:` volume destroys the operating system and all the data in it. For Linux, do not initialize the device that contains a system volume (such as `/boot`, `swap`, and `/` (root)).

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 from iManager:

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

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage to view the *Devices* list.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

   A list of devices appears in the *Devices* list.
3. In the *Devices* list, select the device that you want to initialize.
4. Click *Initialize Disk*.

To initialize a disk from NSSMU:

1. In NSSMU, click *Devices*.

   A list of devices appears in the *Devices* list.
2. In the *Devices* list, select a device.
3. Press F3, *Initialize Disk*.

10.8 Sharing Devices for NSS Pools

- Section 10.8.1, “Understanding Sharing,” on page 135
- Section 10.8.2, “Planning for Device Sharing,” on page 136
- Section 10.8.3, “Configuring the Device’s Share State,” on page 136

10.8.1 Understanding Sharing

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

- *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*
- *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*

IMPORTANT: The system hardware does not specify that disk drives 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.

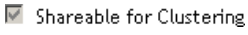

WARNING: Marking a device as shareable for clustering sets all of the pools on this device to shareable. If any of these pools span multiple devices, you must make sure that each device is set to the same share state as this one, or the device can become unusable.

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

10.8.2 Planning for Device Sharing

By default, devices are not shared. Use the following guidelines when planning whether to share devices:

- The system pool (`sys`) and all of its member devices cannot be marked as *Shareable for Clustering*.
- Devices that contain NetWare Traditional partitions and volumes can be marked as *Shareable for Clustering*, but Novell Cluster Services supports clustering only for NSS volumes on those devices.
- You cannot mix space from shared and unshared devices to create a pool. If a pool spans multiple storage devices, all of the member devices in that pool must be marked as *Shareable for Clustering,* or all marked as *Not Shareable for Clustering*.
- You cannot mix space from shared and unshared devices to create a software RAID. All devices that contribute space to the RAID must be marked as *Shareable for Clustering*, or all marked as *Not Shareable for Clustering*.
- Do not mark a device as *Shareable for Clustering* if it is not capable of being shared, such as when the device contributes space to the system pool (`sys`), to an unshared software RAID, or to an unshared pool.

10.8.3 Configuring the Device's Share State

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

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select the server that you want to manage to view a list of its devices.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Devices* list, select a device to view information about it.
4. Depending on the current state of the device, do one of the following:
   - 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 change is not implemented.

10.9 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 155.

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

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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.

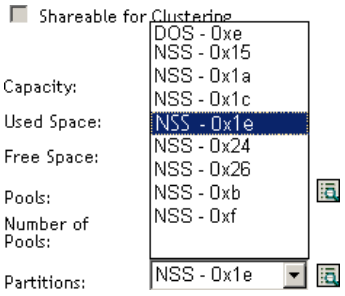

5 Select a partition, then click *View Details*.

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

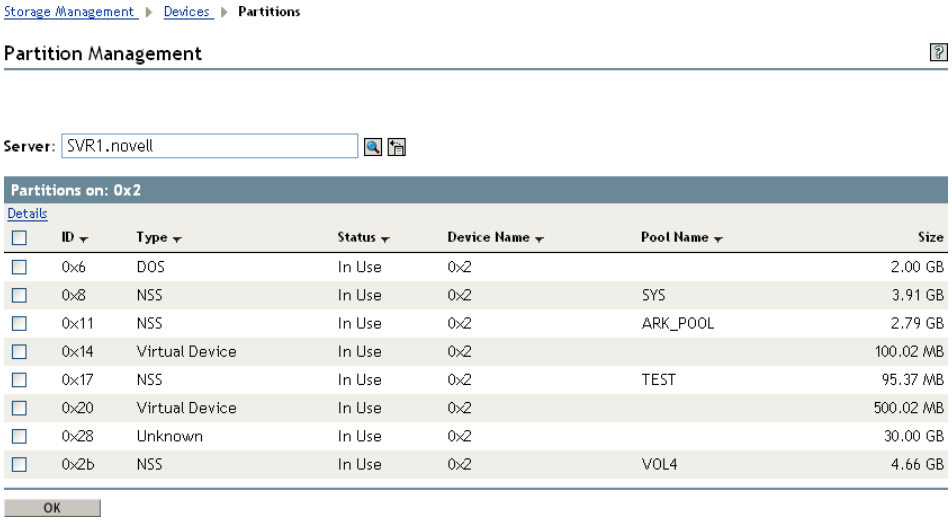

ID	Type	Status	Device Name	Pool Name	Size
0x6	DOS	In Use	0x2		2.00 GB
0x8	NSS	In Use	0x2	SYS	3.91 GB
0x11	NSS	In Use	0x2	ARK_POOL	2.79 GB
0x14	Virtual Device	In Use	0x2		100.02 MB
0x17	NSS	In Use	0x2	TEST	95.37 MB
0x20	Virtual Device	In Use	0x2		500.02 MB
0x28	Unknown	In Use	0x2		30.00 GB
0x2b	NSS	In Use	0x2	VOL4	4.66 GB

To view details about a partition:

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

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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 *Partitions* 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.

10.10 Viewing Pools on a Device

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

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select the server that you want to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
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 *Pools* page where you can view the details of the pool and manage it.

   For information about pool management, see “Managing NSS Pools” on page 197.

10.11 Enabling Write-Through Cache Management on SCSI Devices and RAID Controllers (Linux)

Any journal based file system (including NSS) requires that when writes occur, they must be committed to disk in order to prevent corruption in event of a power failure.

Using write-back cache management can improve performance by allowing data to be held in cache rather than being written to disk. However, write-back cache management introduces the risk of losing data if the power fails. Many array controllers have an on-board battery backup, which can reduce the risk of data loss when using write-back, but it might not eliminate the risk. It is up to you to determine if the power backup is sufficient for power loss scenarios in your production environment.

If your system does not have sufficient power loss protection, we require using write-through cache management for SCSI devices to minimize the risk of losing data in the event of power failure. Write-Through cache management assures the file system that writes are being committed to disk as required.

If the server uses a RAID controller, enable Write-Through (disable Write-Back) cache management when configuring the RAID device by 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. If the `scsi-config` utility is not already installed, install it using the `xscsi` RPM.

   The `xscsi` RPM, which contains the `scsi-config` command, is not installed by default.

   2a In YaST, open the *Various Linux Tools* section.

   2b Install the `xscsi` RPM package, then close YaST.

The `xscsi` package installs the `scsi-config` utility in `/user/bin/scsi-config`.

3 Enable Write-Through (disable Write-Back) cache management for each SCSI device by performing the following for each device where you plan to use NSS volumes:

3a At a terminal console prompt, enter

```
scsi-config
```

3b In the window that opens, browse to select drive you want to manage, then click *Continue*.

3c Click *Cache Control Page*.

3d Enable Write-Through cache management mode by deselecting the *Write cache enabled* check box.

Write-Through cache management is enabled by default, so the *Write cache enabled* check box should be deselected. If the *Write cache enabled* check box is selected, Write-Back cache management mode is enabled and you deselect the box to disable Write-Back cache management.

3e Click *Quit* > *Save Changes*.

3f When prompted to confirm the change, click *Go Ahead and Save* > *Quit*.

3g To verify the setting, at a terminal console prompt, enter

```
scsiinfo -c /dev/sdx
```

Replace `/dev/sdx` with the device you are checking.

A value of 0 for *Write Cache* means that the drive is in Write-Through cache management mode.

10.12 What's Next

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 14, "Managing Multipath I/O to Devices (NetWare)," on page 189. When you are done, continue with creating software RAIDs and pools.

To configure software RAID devices, see Chapter 13, "Managing NSS Software RAID Devices," on page 165.

To create pools of storage on the devices, see Chapter 16, "Managing NSS Pools," on page 197.

Migrating NSS Devices from NetWare to OES 2 Linux

11

This section describes the issues involved in migrating Novell® Storage Services™ devices from NetWare® to Novell Open Enterprise Server 2 Linux servers.

IMPORTANT: For a general discussion of migration issues in OES 2, see "Migrating and Consolidating Existing Servers and Data" in the *OES 2 SP1: Planning and Implementation Guide*.

- Section 11.1, "Guidelines for Moving Devices Cross-Platform," on page 141
- Section 11.2, "Moving Non-Clustered Devices From NetWare Servers to OES 2 Linux Servers," on page 142
- Section 11.3, "Moving Non-Clustered Devices From NetWare 6.0 to OES 2 Linux," on page 147
- Section 11.4, "Moving Non-Clustered Devices From NetWare 6.0 to NetWare 6.5 or OES NetWare," on page 150
- Section 11.5, "Moving Clustered Devices with NSS Volumes Cross-Platform," on page 153
- Section 11.6, "Upgrading NetWare 5.1 NSS Volumes and NetWare Traditional Volumes to NSS Volumes," on page 153

11.1 Guidelines for Moving Devices Cross-Platform

You can move devices containing NSS volumes between NetWare servers and OES 2 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.

- Section 11.1.1, "Media Format," on page 141
- Section 11.1.2, "Pool Snapshots," on page 142
- Section 11.1.3, "Cross-Platform Issues," on page 142

11.1.1 Media Format

The NSS media upgrade for enhanced hard links support is available for the following operating platforms (and later versions):

- Novell Open Enterprise Server 2 Linux and NetWare
- Novell Open Enterprise Server 1 SP1 NetWare
- NetWare® 6.5 SP4

The media upgrade is not available on OES 1 Linux, NetWare 6.5 SP3, and earlier versions. You cannot move a device that has been upgraded to the new media format to a platform that does not support it. For information about the media format upgrade, see Chapter 4, "Upgrading the NSS Media Format," on page 55.

11.1.2 Pool Snapshots

Different pool snapshot technologies are used for NSS pools on NetWare and NSS pools on Linux. You can create pool snapshots on either platform, but the snapshots are unusable if you move the devices cross-platform.and are invalid if you move the volume back.

WARNING: Before moving a device cross-platform make sure to delete any existing pool snapshots for all pools on it. You might not be able to see the original pools on the Linux platform if you do not delete the snapshots before moving the device.

11.1.3 Cross-Platform Issues

For information about differences to expect when using NSS cross-platform, see Chapter 7, "Cross-Platform Issues for NSS," on page 87.

11.2 Moving Non-Clustered Devices From NetWare Servers to OES 2 Linux Servers

This section describes how to move devices cross-platform from a NetWare 6.5 SP7 or OES 2 NetWare server to an OES 2 Linux server. NSS supports moves of devices containing NSS volumes between any servers that support a compatible media format. For information, see Section 11.1, "Guidelines for Moving Devices Cross-Platform," on page 141.

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

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.

- Section 11.2.1, "Prerequisites," on page 143
- Section 11.2.2, "Setting Up File Access For Users on the OES 2 Linux Server," on page 143
- Section 11.2.3, "Decommissioning Each NSS Pool and Its Volumes on the Original Server," on page 144

- Section 11.2.4, “Recommissioning Each NSS Pool and Its Volumes on the Destination Server,” on page 145
- Section 11.2.5, “Using Scripts to Decommission and Recommission NSS Volumes,” on page 146

11.2.1 Prerequisites

The prerequisites in this section apply to moving multiple devices from a NetWare 6.5 SP7 or OES 2 NetWare server to an OES 2 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 7.2, “Cross-Platform Issues for NSS Volumes,” on page 87.

Original NetWare Server

You can move NetWare 6.5 SP 4 or later NSS media to an OES 2 Linux server if the operating platform can support the NSS media format. NetWare 6.5 SP3, OES 1 SP2 Linux, and earlier servers do not support the new media format.

For information, see Section 4.1, “Guidelines for Upgrading the Media Format of NSS Volumes,” on page 55.

Destination OES 2 Linux Server

- NSS and other OES 2 services it needs must be installed on the OES 2 Linux server where you want to move the NSS volume. For information, see Section 3.3, “Installing and Configuring NSS on Linux,” on page 48.
- EVMS must be installed and running on the OES 2 Linux server.

11.2.2 Setting Up File Access For Users on the OES 2 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 2 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 8.8 Administration Guide*.

Set Up Protocols and Services

To provide access for users on the OES 2 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 *OES 2 SP1: NCP Server for Linux Administration Guide*.
- **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.

  For information about enabling users for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

For guidelines about users and access, see Section 5.5, "Access Control for NSS on Linux," on page 71.

11.2.3 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 2 Linux server, you must Linux-enable the current users of the volumes before you move the devices.

   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:

   - 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 *OES 2 SP1: Novell Linux User Management Technology Guide*.
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 from the VLDB for each of the decommissioned volumes.

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

11.2.4 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.

At the server command prompt on the primary VLDB server, enter

```
vldb repair
```

This adds a GUID entry to the VLDB for each of the recommissioned volumes.

11.2.5 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 devices cross-platform (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 1 NetWare or NetWare 6.5 to OES 2 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 2 Linux server, you must Linux-enable the current users of the volumes before you move the devices.

   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:

   - 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 *OES 2 SP1: Novell Linux User Management Technology Guide*.
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.

11.3 Moving Non-Clustered Devices From NetWare 6.0 to OES 2 Linux

NSS supports moves of devices containing NSS volumes from a NetWare 6.0 server to an OES 2 Linux server. For information, see Section 11.1, "Guidelines for Moving Devices Cross-Platform," on page 141.

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 to Linux, NetWare to NetWare, or Linux to Linux).

This section describes how to move devices cross platform from a NetWare 6.0 server to an OES 2 Linux or later server:

- Section 11.3.1, "Prerequisites," on page 147
- Section 11.3.2, "Setting Up File Access For Users on the OES 2 Linux Server," on page 148
- Section 11.3.3, "Decommissioning Each NSS Pool and Its Volumes on the Original Server," on page 148
- Section 11.3.4, "Recommissioning Each NSS Pool and Its Volumes on the OES Linux Server," on page 149

11.3.1 Prerequisites

The prerequisites in this section apply to moving multiple devices from a NetWare 6.0 server to an OES 2 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 Section 7.2, "Cross-Platform Issues for NSS Volumes," on page 87.

Destination OES 2 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 *OES2 SP1: Linux Installation Guide*.

11.3.2 Setting Up File Access For Users on the OES 2 Linux Server

Before or after you move an NSS volume from NetWare 6.0 to OES 2 Linux, you need to set up file access for users on the OES 2 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 8.8 Administration Guide*.

Set Up Protocols and Services

To provide access for users on the OES 2 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 *OES 2 SP1: NCP Server for Linux Administration Guide*.
- **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.

  For information about enabling users for Linux, see the *OES 2 SP1: Novell Linux User Management Technology Guide*.

For guidelines about users and access, see Section 5.5, "Access Control for NSS on Linux," on page 71.

11.3.3 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 2 Linux server, you must Linux-enable the current users of the volumes before you move the devices.

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:

- 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 *OES 2 SP1: Novell Linux User Management Technology Guide*.

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.

At the command prompt on the primary VLDB server, enter

```
vldb repair
```

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

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

11.3.4 Recommissioning Each NSS Pool and Its Volumes on the OES Linux 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.

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 2 Linux server.

8d Click *Save* to save the changes.

9 If you want to mirror the upgraded partition from Step 8, see Section 13.7, “Mirroring an Existing Pool with NSSMU,” on page 177.

11.4 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 11.1, “Guidelines for Moving Devices Cross-Platform,” on page 141.

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 and Linux, NetWare and NetWare, or Linux and 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 11.4.1, "Prerequisites," on page 151
- Section 11.4.2, "Setting Up File Access For Users on the Destination Server," on page 151
- Section 11.4.3, "Decommissioning Each NSS Pool and Its Volumes on the Original Server," on page 151
- Section 11.4.4, "Recommissioning Each NSS Pool and Its Volumes on the Destination Server," on page 152

11.4.1 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.

11.4.2 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 8.8 Administration Guide*.

11.4.3 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.

11.4.4 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.

At the command prompt on the primary VLDB server, enter

```
vldb repair
```

The VLDB Repair adds a GUID entry to the VLDB for each of the recommissioned volumes.

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 recommend that you consider upgrading the media format to take advantage of the enhanced hard link support.

For information, see Chapter 4, “Upgrading the NSS Media Format,” on page 55.

10 If you want to mirror the upgraded partition from Step 9, see Section 13.7, “Mirroring an Existing Pool with NSSMU,” on page 177.

11.5 Moving Clustered Devices with NSS Volumes Cross-Platform

You can use a shared NSS data pool and volume in a mixed cluster configuration by using Novell Cluster Services for Linux if all nodes in the cluster support the same NSS media format.

For information about upgrading to the latest media format, see Section 4.1, “Guidelines for Upgrading the Media Format of NSS Volumes,” on page 55.

For information, see the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.

11.6 Upgrading NetWare 5.1 NSS Volumes and NetWare Traditional Volumes to NSS Volumes

Upgrading NSS Volumes from NetWare 5.1 to NetWare 6.x Format

To upgrade NSS volumes from NetWare 5.1 to NetWare 6.*x*, simply mount the NSS volume on your NetWare 6.*x* server. The mounting process starts the upgrade. The upgrade takes some time. You must leave the upgraded NSS volume mounted on NetWare 6.x afterwards; the media format is not backwards compatible with your NetWare 5.1 server.

Upgrading NetWare Traditional Volumes to NetWare 6.x Format

Use Volume Copy Upgrade to upgrade legacy NetWare 5.1 Traditional volumes to NSS volumes on OES NetWare. For information, see “Upgrading Legacy NSS and NetWare Traditional Volumes” on page 547.

Migrating Upgraded NSS Volumes to OES 2 Linux

You can mount NSS volumes that have not been upgraded to the new format directly on OES 2 Linux and OES 1 Linux servers. You can mount NSS volumes from NetWare 6.0 servers directly onto OES 2 Linux and later. For information, see the following:

- Section 7.2, “Cross-Platform Issues for NSS Volumes,” on page 87
- Section 11.5, “Moving Clustered Devices with NSS Volumes Cross-Platform,” on page 153

NetWare 6.5 SP3, OES 1 SP2 Linux, and earlier servers do not support the new media format for NSS. For information, see Section 4.1, “Guidelines for Upgrading the Media Format of NSS Volumes,” on page 55.

Managing Partitions

12

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 12.1, “Understanding Partitions,” on page 155
- Section 12.2, “Creating Partitions,” on page 158
- Section 12.3, “Viewing a List of Partitions,” on page 159
- Section 12.4, “Viewing Details for a Partition,” on page 160
- Section 12.5, “Labeling a Partition,” on page 160
- Section 12.6, “Deleting a Partition,” on page 161
- Section 12.7, “Adding Other Types of File System Partitions (Linux),” on page 162

12.1 Understanding Partitions

- Section 12.1.1, “NSS Partitions,” on page 155
- Section 12.1.2, “Viewing Partitions on a Device,” on page 156
- Section 12.1.3, “Understanding Types of Partitions,” on page 156
- Section 12.1.4, “Understanding Partition Details,” on page 157

12.1.1 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 the *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 automatically managed by NSS whenever you create pools. You do not manage NSS partitions directly. Table 12-1 identifies the tools used by NSS for managing devices and partitions:

Table 12-1 *Device and Partition Managers*

Platform	Partition Management Tool
Linux	Enterprise Volume Management System (EVMS). EVMS discovers and mounts the partition in the `/dev/evms/` directory. **IMPORTANT:** If you create and mount pools from the command line, you must modify your mount point to the partition’s location under `/dev/evms/`.
NetWare	Media Manager

12.1.2 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 12-1 *Example Partitions Page*

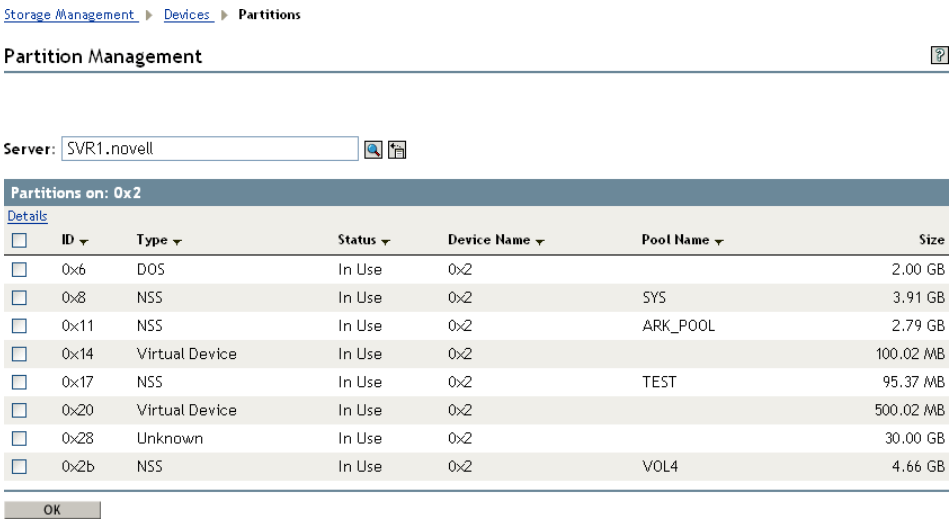

12.1.3 Understanding Types of Partitions

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 automatically create a Virtual Device 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 12-2 *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
iSCSI	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.
NSS	The primary partition type for NSS file systems.
Reiser	The partition type for Linux Reiser file systems.
System Configuration	A vendor-specific partition for maintaining metadata about the server configuration.
Traditional	A partition type found in NetWare systems with legacy NetWare traditional file systems. These might exist in a mixed NetWare environment.
Unknown	An partition type that is unknown to the current operating system.
Upgrade Traditional	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.

12.1.4 Understanding Partition Details

You can view the following information about partitions:

Table 12-3 *Explanation of Partition Details*

Partition Detail	Description
Partition ID	The partition name assigned by NetWare.
Partition Name	The physical descriptive name of the partition that corresponds to the device's physical descriptive name, followed by the type of partition it is.
Type	The abbreviated name of the partition type.
Status	Specifies if a partition is In Use or Available.
Label	The partition name assigned by the administrator.
Starting Offset	Amount of space on the disk that precedes the beginning of the selected partition.
Size	The storage capacity of this partition.
Device Name	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.
Device ID	The device name assigned by NetWare.
Pool Name	For NSS partitions, specifies the name of the pool that uses the partition.

12.2 Creating Partitions

Generally, you cannot create or modify partitions with iManager or NSSMU tools. However, you must use NSSMU to create iSCSI partitions on the iSCSI disk server that is running iSCSI target software.

- Section 12.2.1, "Creating Virtual Device Partitions for Software RAIDs," on page 158
- Section 12.2.2, "Creating NSS Partitions for Pools," on page 158
- Section 12.2.3, "Creating Cluster Services Partitions," on page 158
- Section 12.2.4, "Creating iSCSI Partitions (NetWare)," on page 158

12.2.1 Creating Virtual Device Partitions for Software RAIDs

The NSS tools automatically create a Virtual Device partition when you create a software RAID device. NSS automatically creates, manages, and deletes the partitions as you manage the RAID itself.

For information about creating NSS software RAIDs, see Section 13.5, "Creating Software RAID Devices with iManager," on page 174.

12.2.2 Creating NSS Partitions for Pools

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, and manages them as you manage the pool. There might be multiple NSS partitions that are aggregated and managed underneath the single pool of space.

For information about creating pools, see Section 16.2, "Creating a Pool," on page 198.

12.2.3 Creating Cluster Services Partitions

You can view Novell Cluster Services™ partitions with NSS tools, but you cannot create them. Use the Novell Cluster Services management interface to create and manage Cluster Services partitions. For information, see the following:

- *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*
- *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*

12.2.4 Creating iSCSI Partitions (NetWare)

Use NSSMU to create an iSCSI device and to manage its partitions. To create iSCSI partitions, the server must run iSCSI target software and be the server you plan to make your iSCSI disk server.

1. Start the NSSMU utility by entering `nssmu` at the target server console.
2. Select *Partitions* from the *Main* menu.
3. Press *Insert* to open a list of available devices.
4. Select the device where you want to create the partition.

5 Select *iSCSI* as the partition type.

6 Specify the partition size in MB, then select *Create* to create the partition.

When you are done creating iSCSI partitions, you must load iSCSI target software, configure access control to the target, and then create pools and volumes on the target from an iSCSI initiator. For information about installing iSCSI target software and managing iSCSI devices, see the *OES 2 SP 1: iSCSI 1.1.3 for NetWare Administration Guide*.

12.3 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.

In iManager, you must access partitions through a related task, such as the following:

- Section 12.3.1, "Viewing Partitions on a Device," on page 159
- Section 12.3.2, "Viewing Partitions in a Software RAID Device," on page 159
- Section 12.3.3, "Viewing Partitions in an NSS Pool," on page 160

12.3.1 Viewing Partitions on a Device

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

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select the server that you want to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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.

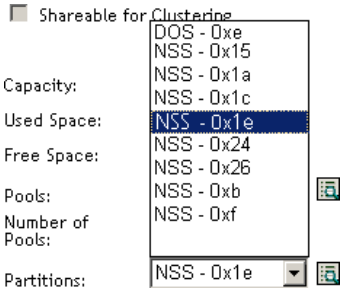

5 Select a partition, then click the *View Partition Details* icon.

The *Partitions* page displays a list of all the partitions that currently exist on the selected device.

12.3.2 Viewing Partitions in a Software RAID Device

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select the server that you want to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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 *Partitions* 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.

12.3.3 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select the server that you want to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Pools* list, 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 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 *Partitions* 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.

12.4 Viewing Details for a Partition

1 From the *Partitions* page, select a partition.

2 Click *Details* to view the *Partition Information* page.

For information about the details, see Section 12.1, “Understanding Partitions,” on page 155.

12.5 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* page, select the partition, then click *Edit*.

2 Type the new label.

3 Click *OK*.

12.6 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 POSIX file systems partitions on your Linux server or non-NetWare partitions on your NetWare server.

This section describes the following:

- Section 12.6.1, "Deleting NSS Partitions in an NSS Software RAID Device," on page 161
- Section 12.6.2, "Deleting NSS Partitions in a Pool," on page 162
- Section 12.6.3, "Deleting Cluster Services Partitions (NetWare)," on page 162
- Section 12.6.4, "Deleting an iSCSI Partition (NetWare)," on page 162

12.6.1 Deleting NSS Partitions in an NSS Software RAID Device

For NSS software RAID devices, use the Software RAIDs 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 13.14, "Replacing a Failed Segment in a Software RAID," on page 183

To delete partitions for a RAID 0, use the Software RAIDs 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 13.15, "Deleting a Software RAID Device," on page 185.

Deleting Partitions with iManager

To delete a single partition in a software RAID 1 or RAID 5 device:

1. In iManager, click *Storage* > *Software RAIDs*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select the server that you want to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Software RAIDs* list, 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* page for the selected software RAID.
5. Select a partition, then 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.

12.6.2 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 16.6, "Deleting a Pool," on page 204.

12.6.3 Deleting Cluster Services Partitions (NetWare)

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 "Console Commands for Novell Cluster Services" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*.

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.

12.6.4 Deleting an iSCSI Partition (NetWare)

Deleting an iSCSI partition on a NetWare server 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: If the iSCSI device is shared in a Novell Cluster Services cluster, make sure to offline the shared NSS pool before deleting the partition.

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.

12.7 Adding Other Types of File System Partitions (Linux)

NSS and Linux POSIX file systems can co-exist on a device that is managed by EVMS. Use NSSMU or the iManager Storage plug-in to manage NSS pools and volumes on OES Linux. Use the EVMS GUI (evmsgui) management tool to create Linux POSIX file systems on the device. For more information, see Section D.2, "Using Linux POSIX File Systems and NSS Pools and Volumes on the Same Device," on page 558.

Do not use YaST or LVM2 management tools to create or manage NSS pools and volumes on EVMS-managed devices. 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. This can make your system unbootable.

For example, a command like this is added to the `/etc/fstab` file for each NSS pool:

```
/dev/evms/poolname /nss/.pools/poolname ext2 defaults 1 2
```

where *poolname* is the name of the pool.

Instead of `ext2`, the partition type should be `nss`:

```
/dev/evms/poolname /nss/.pools/poolname nss defaults 1 2
```

To work around this problem when using the YaST partitioner, after you create a Linux POSIX 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 2 Linux server in Single User mode.
2. Mount the root (`/`) file system with the `remount` and read/write (`rw`) options by entering the following at a terminal console prompt:

   ```
   mount -n -o remount,rw /
   ```
3. Edit the `/etc/fstab` file to do the following, then save the file:
   - If a line wrongly identifies an NSS partition as an Ext2 partition, correct the entry by changing `ext2` to `nss`.
   - If a line wrongly identifies an NSS partition as an Ext2 partition, and if the entry duplicates a correct entry for the NSS partition, then remove or comment out the line that wrongly identifies the NSS partition as an Ext2 partition.
4. Reboot the server to apply the changes.

Managing NSS Software RAID Devices

13

RAID devices help provide data fault tolerance for storage devices. In some RAID configurations, the read/write performance is also improved. The Novell® Storage Services™ (NSS) File System supports software RAIDs 0, 1, 5, 10, and 15 on NetWare and RAIDs 0, 1, and 5 on Linux.

This section describes the following:

- Section 13.1, “Understanding Software RAID Devices,” on page 165
- Section 13.2, “Planning for a Software RAID Device,” on page 167
- Section 13.3, “Viewing a List of Software RAID Devices on a Server,” on page 171
- Section 13.4, “Viewing Details of a Software RAID Device,” on page 172
- Section 13.5, “Creating Software RAID Devices with iManager,” on page 174
- Section 13.6, “Creating Software RAID Devices with NSSMU,” on page 176
- Section 13.7, “Mirroring an Existing Pool with NSSMU,” on page 177
- Section 13.8, “Creating a RAID 1 Mirror to Duplicate Data,” on page 178
- Section 13.9, “Creating a Software RAID 10 with NSSMU (NetWare),” on page 178
- Section 13.10, “Creating a Software RAID 15 with NSSMU (NetWare),” on page 179
- Section 13.11, “Renaming a Software RAID Device,” on page 180
- Section 13.12, “Increasing the Size of a Software RAID Device,” on page 180
- Section 13.13, “Restriping a Software RAID 0 or 5 Device,” on page 182
- Section 13.14, “Replacing a Failed Segment in a Software RAID,” on page 183
- Section 13.15, “Deleting a Software RAID Device,” on page 185
- Section 13.16, “Viewing Pools on a Software RAID Device,” on page 185
- Section 13.17, “Viewing Partitions on a Software RAID Device,” on page 186
- Section 13.18, “Deleting Partitions on a Software RAID Device,” on page 186
- Section 13.19, “Managing Software RAID Devices with NSSMU,” on page 187

13.1 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 RAIDs use standard host adapters and do not require any special RAID hardware.

The following table describes the software RAID devices supported by NSS:

Table 13-1 *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 (NetWare only)	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 re-create 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 (NetWare only)	2 to 4 RAID 5 devices	Mirroring RAID 5 devices	Provides full data redundancy for failover and instant recovery. If a single data disk fails, the RAID 5 device remains up and mirrored. Its performance is degraded until you replace the failed disk. It can handle multiple disk failures, depending on the number of failures and which disks fail.	Requires separate channels for each member disk to achieve best I/O performance. RAID 5 devices that you mirror can share no disks in common. If multiple data disks fail concurrently on the same segment, 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.

13.2 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 describes the following:

- General Guidelines for Software RAID Devices (page 167)
- Guidelines for Software RAID 1 Devices (page 168)
- Choosing a Software RAID Solution (page 169)
- Determining the Number of Partitions (page 170)
- Determining the Partition Size (page 170)
- Determining the Stripe Size for RAID 0 and RAID 5 (page 171)

13.2.1 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 13-1 *Striping Data on a Software RAID 0 Device*

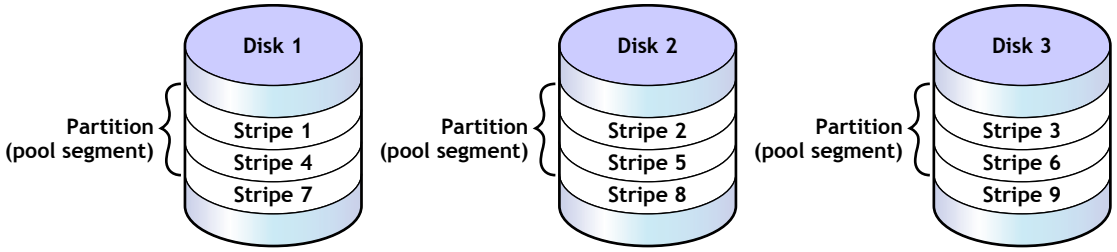

Consider the following general guidelines when creating a software RAID device:

- 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 (`/`) or `/boot` partitions on Linux or the `SYS` pool on NetWare).
- 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 RAIDs 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.

13.2.2 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.

13.2.3 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.

13.2.4 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 table highlights the key data fault tolerance, performance, and configuration issues associated with each RAID type.

Table 13-2 *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	Not supported	Not supported
NSS on NetWare	Yes, for data and system volumes	Yes, for data and system volumes	Yes, for data and system volumes	Yes, for data volumes	Yes, for data 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
Valid names	2 to 128 characters	2 to 80 characters	2 to 128 characters	2 to 80 characters	2 to 80 characters

Requirement	RAID 0	RAID 1	RAID 5	RAID 10	RAID 15
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	0.66 TB, if 3 segments	2 TB for each mirror	0.66 TB, if 3 segments
Maximum RAID size (total for combined segments)	2 TB	2 TB	2 TB	2 TB	2 TB
Minimum segment size	12 MB	12 MB	12 MB	12 MB	12 MB

13.2.5 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 TB (1 TB = 2E40 bytes). For example, when using two segments to create a RAID 0, the maximum partition size of each member segment is 1 TB. For a RAID 1, each segment is a mirror, so the maximum size of each segment is 2 TB.

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.

13.2.6 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) that can be seen by 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 RAIDs, 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 13.14, “Replacing a Failed Segment in a Software RAID,” on page 183.

13.2.7 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.

13.3 Viewing a List of Software RAID Devices on a Server

Use the Software RAIDs task in the iManager Storage plug-in to create and manage your software RAID devices.

1. In iManager, click *Storage* > *Software RAIDs*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. Select a server in the Novell eDirectory™ tree where you are logged in.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

   A list of software RAID devices appears in the *Software RAIDs* 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 RAIDs* list.

   For an overview of the subtasks available from this page, see “Software RAIDs” on page 108.

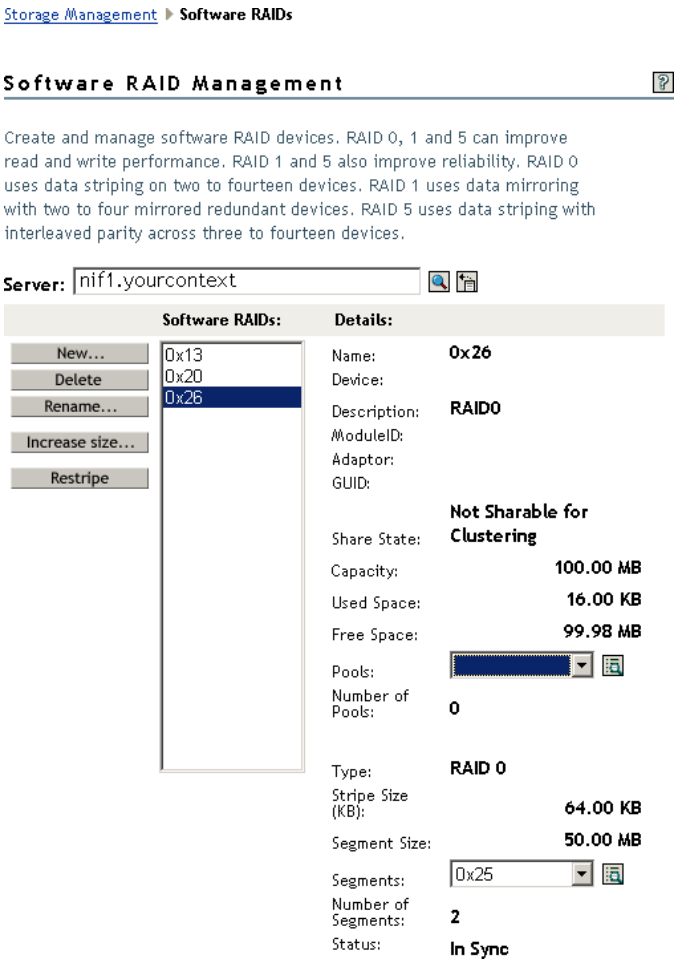

13.4 Viewing Details of a Software RAID Device

You can view the following information about a selected software RAID device:

Table 13-3 *Explanation of Details for a Software RAID Device*

Software RAID Device Detail	Description
Name	On Linux, this is the administrator-specified descriptive name for the RAID. In iManager, if you do not specify a name for the device at create time, a name is autogenerated in the format of RAID <type> Device <number>. On NetWare®, this is the device ID assigned by the operating system, such as 0x2. This name appears in the *Devices* list.
Device	On NetWare, this field displays the RAID type and the number of member devices.
Description	On NetWare, this is the administrator-specified descriptive name for the RAID. In iManager, if you do not specify a name for the device at create time, a name is autogenerated in the format of RAID <type> Device <number>.

Software RAID Device Detail	Description
Module ID	The Novell driver identification number for the module; it indicates what type of driver you are using.
Adapter	The sequential number registered for the adapter; it indicates which drive is on which adapter.
GUID	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.
Share State	*Shareable for Clustering* or *Not Shareable for Clustering*. The share state can be modified on the Devices 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 all marked as *Not Shareable for Clustering*.
Capacity	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.
Used Space	The amount of space on the device that is currently in use by NSS partitions.
Free Space	The total amount of space on the device that is currently not in use.
Partitions	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 *View Details* icon.
Partition Size	The size (in MB) of partitioned space per drive.
Stripe Size	The maximum size (in KB) of a data write, as configured for a RAID 0 or RAID 5 device.
Number of Partitions	The total number of partitions in the selected software RAID device.
Status	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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2. Select a server in the eDirectory tree where you are logged in.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   This opens the *Software RAIDs Management* page.

   The *Software RAIDs* list displays the virtual RAID devices on the selected server. The list might include any RAID O, 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 RAIDs* list to view information about that device, then wait for the page to refresh.

Viewing Details in NSSMU

1 In NSSMU, select *Software RAIDs*.

2 Select the RAID device you want to manage and wait for the information to be displayed.

13.5 Creating Software RAID Devices with iManager

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server in the eDirectory tree where you are logged in.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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, type a name for the RAID device, then click *Next*.

If you leave the Name field blank, NSS creates a unique name for the device in the form of RAID <type> Device <sequential_number>. For information about choosing names, see Section 5.4, “Naming NSS Storage Objects,” on page 67.

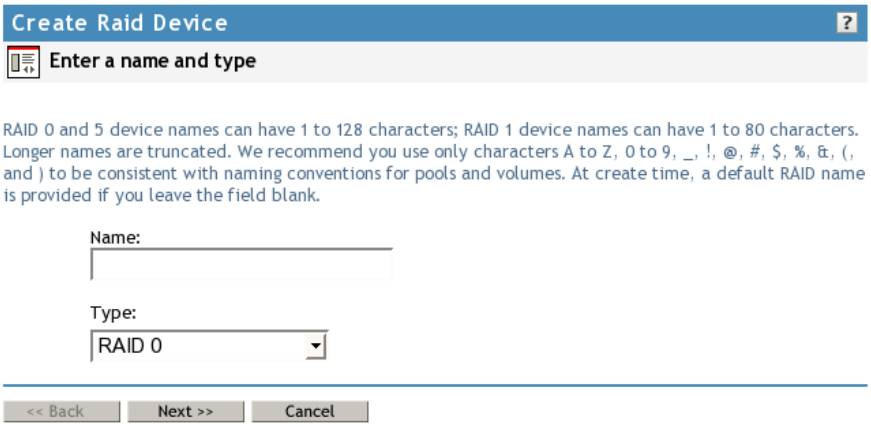

This opens the *Devices and Space* page.

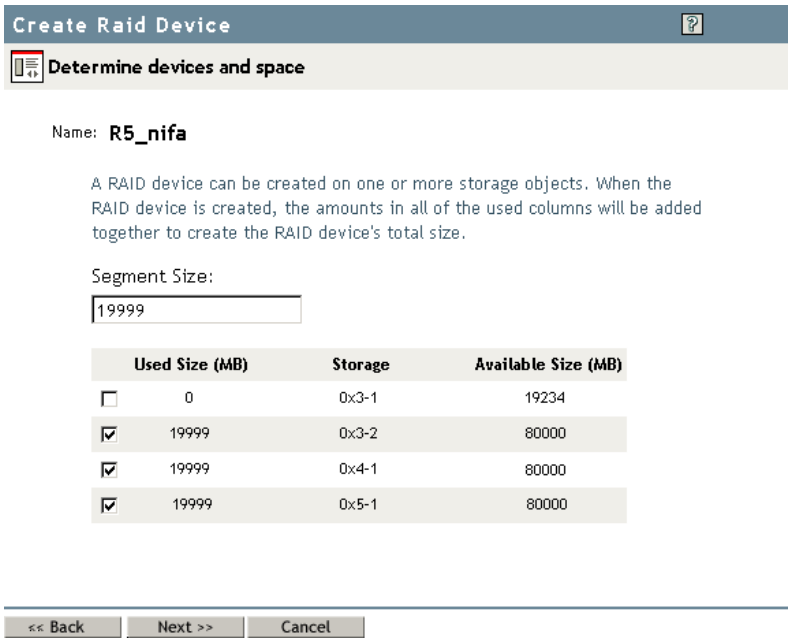

5 Devices and Space: Select devices and the amount of space to use from each, then click *Next*.

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 170.

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.

5b 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. Each segment must be more than 12 MB. The combined space from all segments must be less than 2 TB.

If a device's available space is smaller than the specified partition size, it is disabled (dimmed) so that you cannot select it.

IMPORTANT: Unallocated partitions (that is, partitions that are not mirrored and do not contain pools or other file systems), are deleted in order to present the unused space as free space for use by the RAID. No data loss occurs by this action.

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.

5c Click *Next*

6 **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 171.

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. The name you provided for the RAID appears in the Description field. This is the device name displayed when the RAID is listed in the NSSMU Software RAIDs page.

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 16, "Managing NSS Pools," on page 197.

13.6 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. Each segment must be more than 12 MB. The combined space from all segments must be less than 2 TB.

IMPORTANT: Unallocated partitions (that is, partitions that are not mirrored and do not contain pools or other file systems), are deleted in order to present the unused space as free space for use by the RAID. No data loss occurs by this action.

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 the RAID window.

13.7 Mirroring an Existing Pool with NSSMU

1 Before you begin, review the Section 13.2.2, “Guidelines for Software RAID 1 Devices,” on page 168.

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.

7 Select the RAID device, then view its details to make sure that synchronization has begun.

The remirroring status shows a percentage is greater than 0. It is fully synchronized at 100%

If you are mirroring a shared partition (that is, if it contains a cluster-enabled pool), the synchronization to the mirror does not begin automatically. Continue with Step 8a.

8 If you are mirroring a shared partition (that is, if it contains a cluster-enabled pool), start the remirroring manually by doing the following:

8a Use one of the following methods to initiate mirroring for the newly created mirror:

- At the server console of a cluster node, enter the following to migrate the cluster resource to another node:

  ```
  cluster migrate cluster_resource destination_node_name
  ```

  Migrating the pool causes load scripts to be executed and causes the mirroring to start on the new node.

- At the server console of the cluster node where the pool is currently active, enter

  ```
  dmsetup message raid_device_name 0 remirror=on
  ```

WARNING: Issue this command only on the node where the pool is currently active. Issuing the command on multiple nodes can corrupt the mirror.

8b Verify that the remirroring has begun by opening NSSMU on the node where the pool is currently active, open the RAID page, then select the RAID device.

The remirroring status shows a percentage is greater than 0. It is fully synchronized at 100%.

13.8 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 13.5, “Creating Software RAID Devices with iManager,” on page 174.
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 and uses 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.

13.9 Creating a Software RAID 10 with NSSMU (NetWare)

For NSS on NetWare, you can mirror your software RAID 0 devices to create a nested RAID 10 device. 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.

IMPORTANT: NSS management tools on Linux do not support creating nested software RAID 10 device.

1. In NSSMU, create a software RAID 0 device with 2 to 14 segments.

For information, see Section 13.6, “Creating Software RAID Devices with NSSMU,” on page 176.

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 RAIDs 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.

13.10 Creating a Software RAID 15 with NSSMU (NetWare)

For NSS on NetWare, you can mirror your software RAID 5 devices for your server to create a nested RAID 15 device. 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.

IMPORTANT: NSS management tools on Linux do not support creating nested software RAID 15 device.

---

1 In NSSMU, create a software RAID 5 device with 3 to 14 segments.

For information, see Section 13.6, “Creating Software RAID Devices with NSSMU,” on page 176.

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 RAIDs 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.

13.11 Renaming a Software RAID Device

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server in the eDirectory tree where you are logged in.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 Click *Rename*.

This opens a dialog box where you can enter the new name.

4 Type the new name.

Do not leave the field blank when you are renaming because a default RAID name is not generated for a rename procedure. For information about choosing names, see Section 5.4, “Naming NSS Storage Objects,” on page 67.

5 Click *Finish*. (Or click *Cancel* at any time to back out of the process.)

NSS renames the software RAID device, then opens to the Software RAIDs Page. The details for the renamed software RAID device are displayed on the page, with the new name in the Description field.

13.12 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.

IMPORTANT: If the software RAID device is shared in a cluster, connect to the node where the RAID is currently active to manage the RAID and increase the size of the RAID.

To add partitions to an existing software RAID:

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server in the eDirectory tree where you are logged in.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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 Device*s 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. Click *Cancel*.

  A device must be the same size or larger than the segment size being used in the RAID. You might need to add or initialize a new device, then try again.

- 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 170 and “Determining the Partition Size” on page 170.

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. While restriping, the new device is considered a failed device until it is completely resynchronized. After the restriping is complete, the RAID’s capacity includes the added partition.

While expanding a RAID 5 device, if one of the drives goes down (either one of the existing segments or the newly added segment), the pool deactivates. If you remove any device from a RAID 5 other than the one that was just added for restripe, it considers that as a two-disk error, and deactivates the RAID and the pool.

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.

13.13 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.

When expanding a RAID 5 on Linux, if the newly added drive goes down during the restripe, the restriping continues without the new partition and puts the RAID in a degraded state with one partition missing. If the same partition comes back online, it finishes the restripe. If the partition has completely failed, after the degraded restriping is complete, you can add a new replacement partition, and the RAID restripes to fix it.

When expanding a RAID 5 on NetWare, if the newly added drive fails during the restripe, this is considered a fault and the device and its pools are automatically deactivated. If the same partition comes back online, it finishes the restripe. If the partition fails to come back online, you must re-create the RAID to the desired size and recover from backup. (Handling this type of failure with a degraded restripe recovery (same as for RAID 5 on Linux) is planned for a future release.)

To manually resume or pause the Restripe process:

1. In iManager, click *Storage* > *Software RAIDs*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server in the eDirectory tree where you are logged in.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
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 RAIDs* 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.

13.14 Replacing a Failed Segment in a Software RAID

For some RAIDs, 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 13-4 *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	The last segment of a RAID 1 device is the primary device that contains the data. You can use the 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. 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.
RAID 5	One data segment at a time	You can temporarily remove one segment at a time in a RAID 5 device to replace a failed data disk. 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. If multiple data segments fail concurrently, 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. For example, if a second segment fails before the restriping is completed for the first drive replacement, this is considered a two-drive failure. You must recover data from a backup copy.
RAID 10	All but one mirrored segment; cannot remove disks from the underlying RAID 0	The last segment of a RAID 10 is a RAID 0 device that contains the original data. 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.

Software RAID	Remove Segments	Recommended Action
RAID 15	All but one mirrored segment; can safely remove one segment at a time per segment	The last segment of a RAID 15 is a RAID 5 device that contains the original data. 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 RAID 5. If multiple disks in a mirrored segment fail concurrentlys, 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.

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 13.15, “Deleting a Software RAID Device,” on page 185.

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 13.15, "Deleting a Software RAID Device," on page 185.

13.15 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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server in the eDirectory tree where you are logged in.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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 RAIDs* list.

5 Click *Delete*.

13.16 Viewing Pools on a Software RAID Device

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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 *Pools* page where you can view the details of the pool and manage it. See Section 16.2, "Creating a Pool," on page 198 for a sample *Pools* page.

For information about pool management, see "Managing NSS Pools" on page 197.

13.17 Viewing Partitions on a Software RAID Device

1 In iManager, click *Storage* > *Software RAIDs*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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 *Partitions* 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.

13.18 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 RAIDs*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

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* 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*.

13.19 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 13-5 *Software RAID Management Tasks in NSSMU*

Keystroke	Description
Enter	Shows the RAID partitions associated with the selected RAID device.
F3	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 severely impedes the performance of your file system.
F4 (Linux)	Rename a selected RAID device.
F6	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.
F8	Show pools (list pools on selected RAID 0 or 5; for RAID 1, its single member pool appears in device details)
Ins	Creates a software RAID (0, 1, or 5)
Del	Deletes an entire RAID device. When you delete a RAID device, all partitions, volumes, and pools associated with that device are also deleted.
F5	Refreshes the display.

Managing Multipath I/O to Devices (NetWare) 14

Novell® Storage Services™ for NetWare® supports multipath I/O (MPIO), a fault-tolerant and high-availability storage solution. This section describes how to configure and manage multipath I/O for devices on NetWare.

- Section 14.1, “Planning Your Multipath I/O Solution for NSS on NetWare,” on page 189
- Section 14.2, “Enabling Multipath,” on page 190
- Section 14.3, “Configuring the Primary Path and Priorities for Failover Connection Paths,” on page 191
- Section 14.4, “Viewing Connection Path Details,” on page 192
- Section 14.5, “Setting a Connection Path as Up or Down,” on page 193
- Section 14.6, “Setting the Primary Path for a Device to Its Default Path,” on page 194
- Section 14.7, “Resetting the Server Registry with Default Priority Settings for a Device,” on page 194
- Section 14.8, “Managing Multiple Paths with Server Console Commands,” on page 194

14.1 Planning Your Multipath I/O Solution for NSS on NetWare

Media Manager for NetWare supports multipath I/O 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.

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 by 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. Media Manager multipathing is interoperable with third-party storage products such as SANs and hardware RAIDs.

IMPORTANT: The NSS Media Manager supports fault tolerance with its multipath solution; it does not support load balancing across the multiple paths. Some third-party multipathing systems support load balancing across the multiple connection paths. When you use third-party multipath solutions, make sure you do not turn on Media Manager multi MPIO.

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. Media Manager multipath support enables you to dynamically manage multiple paths according to the priorities you set. If multiple paths have the same highest-priority setting, the path is chosen randomly from among them.

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. If multiple paths have the same highest-priority setting, the path is chosen randomly from among them.

Set a path's priority to determine its position in the failover sequence of the alternate available paths. The default priority is zero (0), or Off. Assign a value between 1 (highest priority) to 4 million (lowest priority). You should set a non-zero priority on all of the paths for the device. If you set a non-zero priority on any of the paths, zero is considered the highest priority until all paths are set to a non-zero value.

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.

14.2 Enabling Multipath

Multipath I/O is disabled by default. If your devices do not have multiple paths, then multipath support should be disabled; you do not need it.

Before enabling multipath support, you need to determine if the a device's HAM driver handles multiple paths to devices. Check to see if the HAM driver presents only one path to a given device even if there are multiple adapters and multiple ports connected to the same storage device.

To check the HAM driver, you can scan devices (`SCAN ALL`), then list them (`LIST DEVICES`). If you can see a given LUN multiple times, then you need to enable multipath support. If a device should have multiple paths and does not, either you have a hardware configuration problem or a driver problem, and you need to resolve the issue before you enable multipath support.

In order to enable multipath support, use the SET parameter under *Disk* > *Multi-Path Support*, or issue the following command at the NetWare server console:

```
set multi-path support = on
```

After enabling the parameter, you need to scan for new devices to allow Media Manager to recognize the new paths to the multipathed devices. For some storage arrays, a reboot might be required.

For persistence across reboots, you can add the following line near the top of the `startup.ncf` file and before any lines that load drivers:

```
set multi-path support = on
```

The scan for new devices occurs automatically on start-up.

If the HAM adapter is handling multipath for the devices, but you want NetWare MPIO to handle failover between paths, you need to disable the HAM driver's multipath handling. Make sure the SAN manager is not managing the HBA devices, then add the /ALLPATHS and /PORTNAMES HBA load options (if appropriate for your configuration) to the HAM driver load lines in order to expose all device paths and ports to NetWare:

The /ALLPATHS option disables the HAM adapter's path failover and reports to NetWare all devices on all adapter paths. All adapters report the same devices, which allows upper layer modules to fail over across a NetWare server's multiple adapters.

The /PORTNAMES option tracks devices internally by port name rather than node name. It disables the HAM adapter's storage port failover and reports all storage ports for each device on the reporting adapters. This is required when storage LUNs do not have a one-to-one correspondence across port names. This allows upper layer modules to fail over across a subsystem's multiple ports.

For example:

```
LOAD QL2x00.HAM SLOT=3 /LUNS /ALLPATHS /PORTNAMES
```

Depending on your hardware configuration, specify one or both options if you want MPIO to handle paths to multiple adapters, multiple ports, or both multiple adapters and multiple ports.

The AEN (Asynchronous Event Notification) parameter for `scsihd.cdm` is also advisable if you are using Fibre Channel. When AEN is enabled, the HAMs that can, such as Fibre Channel, report the loss of a hard disk drive or the discovery of a new one even without active IO. The default value for AEN is off.

If `scsihd.cdm` is already loaded without the /AEN option, you can enter the following at the NetWare server console:

```
load scsihd.cdm /aen
```

or

```
load scsihd.cdm /aen
```

AEN functionality is seamlessly added to the driver without the driver re-loading or re-initializing.

14.3 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 direct attached storage devices.

To set the primary path and set priorities for failover:

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

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3. In the *Devices* list, 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.

14.4 Viewing Connection Path Details

You can view the following information in the Failover Device Path details:

Table 14-1 *Explanation of Failover Device Path Details*

Failover Device Path Detail	Description
Device Name	The device name assigned by NetWare.
Path Name	The pathname assigned by NetWare.
Adapter ID	The ID assigned to the host bus adapter in the server that is a component of the selected interconnect.
Port	The port address for the storage device that is a component of the selected interconnect.
Priority	The integer value you set for the failover path.
Status	The state of the path. - **Available:** A path is marked as Up if it is available, or Down if it is not. - **Selected:** Only one path serves as the primary path at any given time. The selected path is the primary path.

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

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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.

14.5 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Devices* list, 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.

14.6 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
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.

14.7 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
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.

14.8 Managing Multiple Paths with Server Console Commands

You can optionally use server console commands to manage multipath failover for your NetWare server. For a list of commands, see Section A.23, “Multipath I/O Failover Commands (NetWare),” on page 475.

Managing Multipath I/O to Devices (Linux) 15

Novell® Storage Services™ for Linux does not provide multipath I/O (MPIO) support. This section describes how to use Linux tools to configure and manage multipathing for devices, and how to configure Linux multipathed devices for use with EVMS and NSS.

- Section 15.1, “Understanding Multipath I/O on Linux,” on page 195
- Section 15.2, “NSS Errors When Linux Multipath Is Not Configured,” on page 196
- Section 15.3, “Configuring Multipath on Linux,” on page 196

15.1 Understanding Multipath I/O on Linux

Multipath I/O software resolves multiple paths to a device into a single device and manages the traffic flow across the paths transparently for file systems on the devices. NSS on Linux does not provide an EVMS-based software solution for managing multiple paths like the Media Manager multipath solution on NetWare®. Instead, you can use Linux multipath I/O tools to configure and manage multiple paths for devices where you want to create NSS software RAIDs, pools, and volumes. You can also use solutions from the storage array vendor or third-party vendor.

Devices have multiple connection paths when you implement hardware configurations such as the following:

- The server has multiple host bus adapters for connection to external devices.
- The external storage device has multiple interconnects for connection to one or more host bus adapters.
- A server with multiple host bus adapters is connected to a storage device through intermediate devices, such as a Fibre Channel SAN switch.

In a Linux host, when there are multiple paths to a storage controller, each path appears as a separate block device, which results in multiple block devices for single LUN. The Device Mapper Multipath utility detects multiple paths with the same LUN WWID, and creates a new multipath device with that WWID.

For example, a host with two HBAs attached to a storage controller with two ports via a single unzoned Fibre Channel switch sees four block devices:

```
/dev/sdb
/dev/sdc
/dev/sdd
/dev/sde
```

Device Mapper Multipath creates a single block device, `/dev/mpath/mpath1` that reroutes I/O through those four underlying block devices.

15.2 NSS Errors When Linux Multipath Is Not Configured

If you have not started multipathing before you attempt to configure NSS pools or volumes, NSS cannot resolve the multiple paths and attempts the command on all the paths. You get the following error:

```
Error 21621: zERR_MSAP_POOL_ALREADY_IN_USE.
```

After you have configured Linux multipathing for a device, the multipath device appears in NSSMU or iManager.

15.3 Configuring Multipath on Linux

For detailed information about configuring and managing multipath I/O for devices on Linux, see "Managing Multipath I/O for Devices" in the *SLES 10 SP2: Storage Administration Guide*.

Managing NSS Pools

16

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 Section 9.1, "Novell iManager and Storage-Related Plug-Ins," on page 97.

You can also use the console-based NSS Management Utility to configure and manage NSS pools. For information, see Section 9.2, "NSS Management Utility (NSSMU) Quick Reference," on page 112.

This section describes how to configure and manage NSS pools by completing the following tasks:

- Section 16.1, "Guidelines for Creating a Pool," on page 197
- Section 16.2, "Creating a Pool," on page 198
- Section 16.3, "Activating and Deactivating Pools," on page 201
- Section 16.4, "Increasing the Size of a Pool," on page 202
- Section 16.5, "Renaming a Pool," on page 203
- Section 16.6, "Deleting a Pool," on page 204
- Section 16.7, "Viewing Pools on a Server," on page 204
- Section 16.8, "Viewing Pool Details," on page 205
- Section 16.9, "Viewing Partition Information for a Pool," on page 206
- Section 16.10, "Viewing Volume Information for a Pool," on page 207
- Section 16.11, "Viewing Device Information for a Pool," on page 207
- Section 16.12, "Preventing Pools from Activating on Multiple Servers," on page 208
- Section 16.13, "Updating eDirectory Pool Objects," on page 213
- Section 16.14, "What's Next," on page 213

For information about pool snapshots for NSS on NetWare®, see "Managing NSS Pool Snapshots" on page 229.

16.1 Guidelines for Creating a Pool

Devices must be initialized before any space is shown as available for creating a pool. Without initializing the devices, no space will be shown available for pool creation. For instructions, see Section 10.7, "Initializing a Disk," on page 134.

Novell NCP Server must be installed, configured, and running. For information, see Section 5.6.1, "NCP," on page 75.

Novell CIFS must be installed, configured, and running before you can use the CIFS option when cluster-enabling an NSS pool. For information, see Section 5.6.3, "Novell CIFS," on page 76.

Novell AFP must be installed, configured, and running before you can use the AFP option when cluster-enabling an NSS pool. For information, see Section 5.6.2, “Novell AFP,” on page 76.

16.2 Creating a Pool

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

A list of pools appears in the *Pools* list.

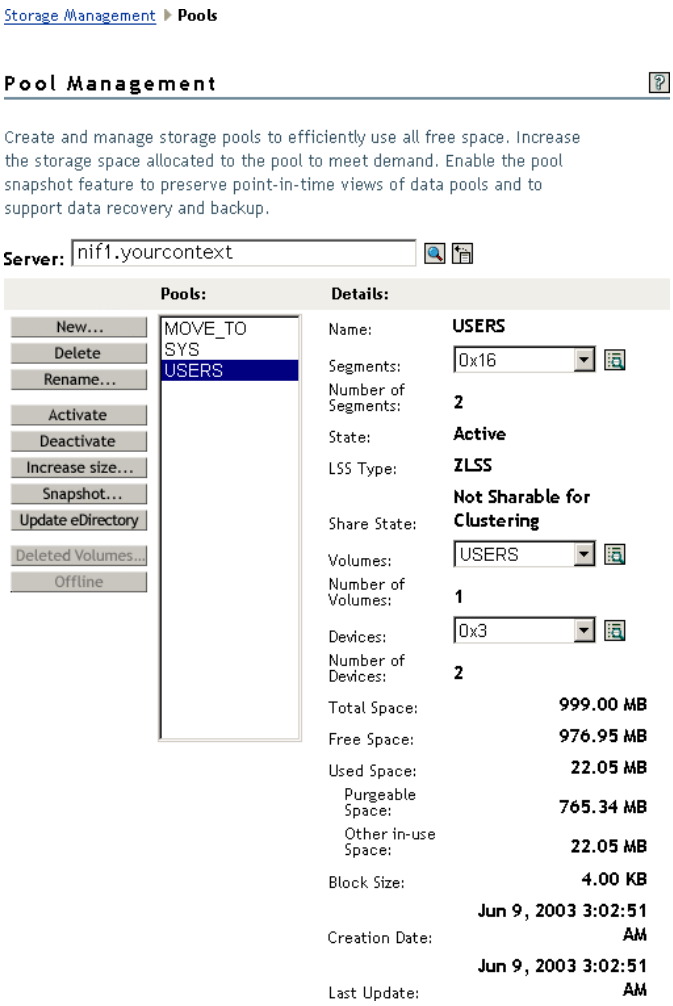

3 To create a new pool, click *New*.

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 5.4, “Naming NSS Storage Objects,” on page 67.

5 Specify device parameters and the space to use, then click *Next*.

5a Select the check box next to one or more of the available devices you want to use in the pool.

5b In *Used Size*, specify the amount of space in megabytes (MB) to add to the pool from each device you selected, up to the amount of free space available for that device.

To update the *Total Pool Size* as you enter the device's *Used Size*, click anywhere within the Wizard dialog box. If any entry exceeds a device's available space, the pool expansion fails and returns an error message.

The pool itself can be up to 8 TB. NSS recognizes devices up to 2 terabytes (TB) in size, so the free space available on any given device is 2 TB or less. To create an 8 TB pool, you would use least four segments of free space of up to 2 TB each.

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, it might be because you need to initialize a recently added device, or it might be that there is no space available on any devices. Cancel the Wizard, add more devices to the server or free up space on existing devices, then return to the *Pools* page to increase the size of this pool.

5c If the selected device is shareable, the *Cluster Enable on Creation* check box is automatically selected so the pool can be shared in a cluster configuration. Deselect the check box if you do not want to cluster-enable this pool for sharing.

5d Select *Activate on Creation* (*Mount on Creation* for Linux) to activate (mount) the device automatically after it is created.

This parameter is automatically enabled. Deselect the check box to turn it off.

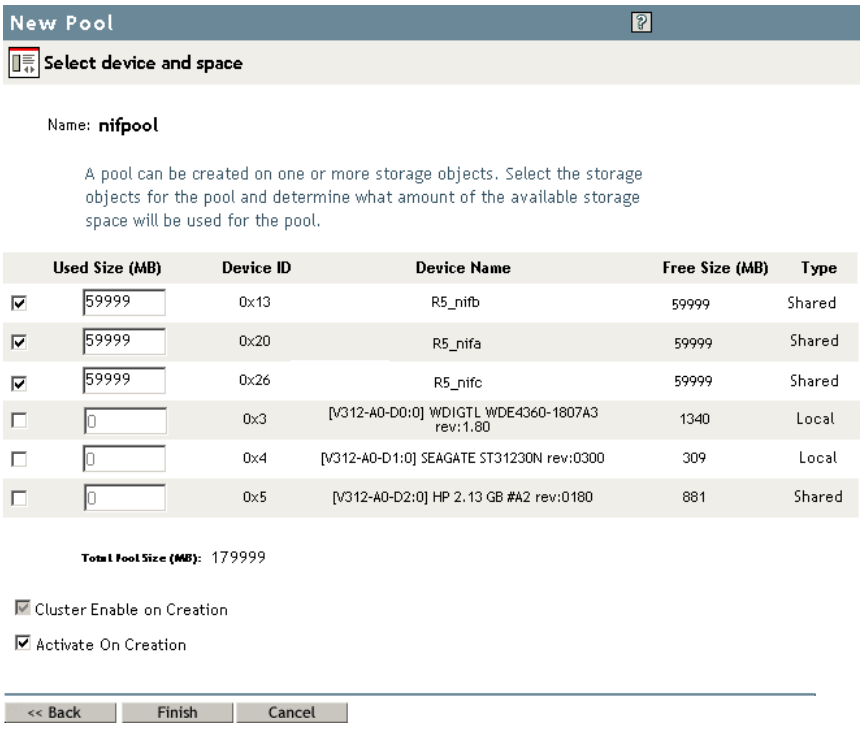

5e If the pool is cluster-enabled, click *Next* to specify its cluster parameters; otherwise, click *Finish*.

6 If the *Cluster Enable on Creation* check box is selected, an additional page appears that allows you to specify the cluster information.

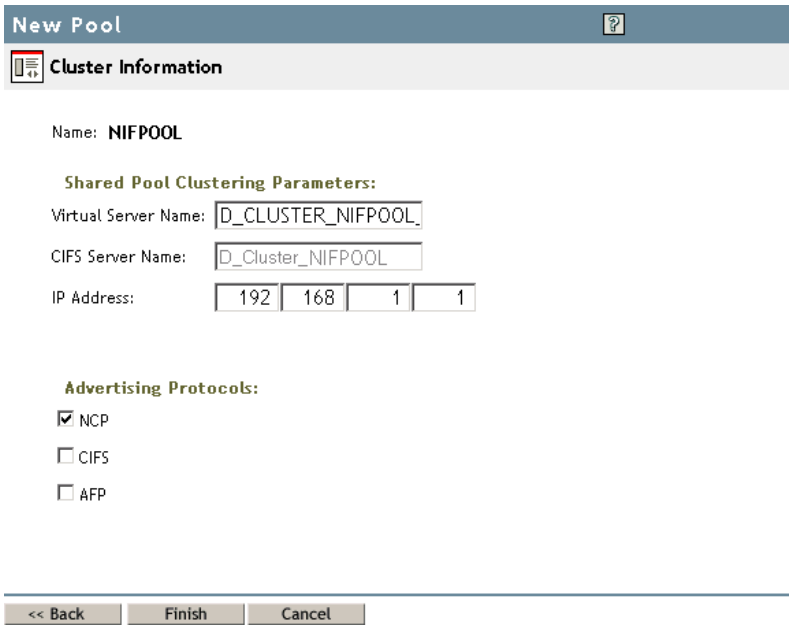

Specify the following shared pool clustering parameters:

- **Virtual Server Name:** The name assigned to the virtual server that represents the shared pool in the cluster.

  When you cluster-enable a pool, a virtual Server object is automatically created in Novell eDirectory™ and given the name of the Cluster object plus the name of the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

- **CIFS Virtual Server Name:** The name assigned 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.

- **IP Address:** The IP address that you want to assign the virtual server.

  Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

  IMPORTANT: The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

  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 the following:

  ```
  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.

NOTE: For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available. CIFS and AFP check boxes can be selected, but CIFS and AFP functionality does not apply to Linux. Selecting the check boxes has no effect.

- NetWare Core Protocol™ (NCP™) is the Novell networking protocol used by the Novell Client™. It is selected by default. Selecting NCP causes commands to be added to the pool-resource load and unload scripts to activate the NCP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to Novell clients.
- CIFS is the Windows networking protocol. Selecting CIFS causes commands to be added to the pool-resource load and unload scripts to activate the CIFS protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to CIFS/Samba clients.
- Apple* Filing Protocol (AFP) is the Macintosh networking protocol. Selecting AFP causes commands to be added to the pool-resource load and unload scripts to activate the AFP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to AFP clients.

7 Click *Finish*.

For NSS on Linux, the create time might take longer than expected. Typically, the pool creation takes less than a minute, and the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

8 Create a volume on the pool. For information, see Section 19.3, “Creating Unencrypted NSS Volumes,” on page 259 or Section 20.3, “Creating an Encrypted Volume,” on page 284.

16.3 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 *Pools* page makes the selected pools and all the volumes in them available for user access. The *Deactivate* option on the *Pools* 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Pools* list, 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 *Actions* > *Activate*.
- Click *Deactivate* > *Actions*.

16.4 Increasing the Size of a Pool

Using the *Increase Size* option on the *Pools* 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3. In the *Pools* list, 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.

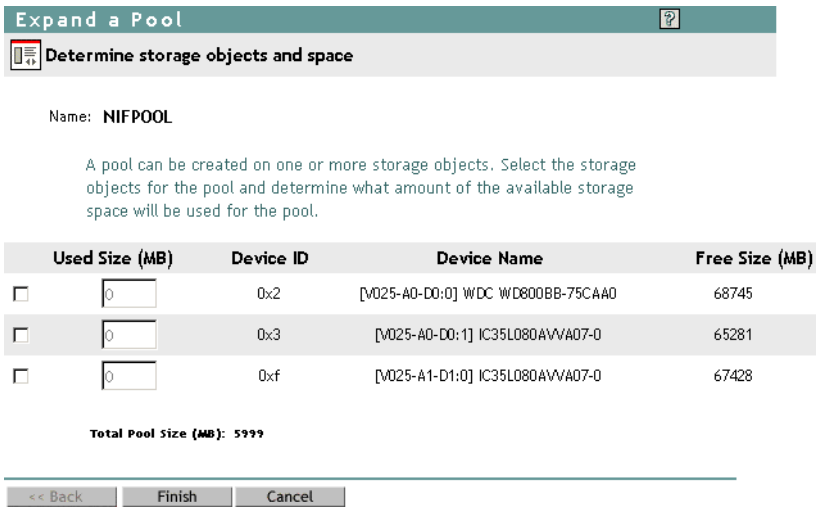

5. Select the devices you want to use and the amount of space to use from each device.

   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* 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 *Pools* page to increase the size of this pool.

6. Click *Finish*, or click *Cancel* at any time to back out of the process.

16.5 Renaming a Pool

The *Rename* option on the *Pools* 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. The pool must be in the active state when you rename the pool so that eDirectory can be updated.

For an NSS pool on Linux, EVMS must unload and reload the pool in order to rename it. Depending on the pool's load-time behavior and share state, the pool might be in a deactive state after the rename and require administrator action to reload the pool and its volumes. Because the volumes are temporarily unavailable, it is best to perform a pool rename on Linux during a period of little or no user activity. See Table 16-1 to determine what actions to take after renaming a pool on Linux:

Table 16-1 *Actions Required after Renaming an NSS Pool on Linux*

Pool Share State	Pool Load-Time State	Pool State After a Rename	Action Required
Unshared	Autoloaded	Active with volumes dismounted	Mount the pool's volumes
Unshared	Not autoloaded	Deactive	Activate the pool, then mount its volumes
Shared	Load and unload controlled by Novell Cluster Services™ for Linux	Deactive	Activate the pool, then mount its volumes

Because renaming involves changing information in the Pool object in eDirectory, you must also make sure that a shared pool resource is loaded on the original server in the cluster where it was created.

For an NSS pool on NetWare, Media Manager updates the pool name without unloading and reloading the pool, so both unshared and shared pools remain in the active state during and after the rename. User activity is not disrupted. No further action is required after the rename.

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Pools* list, select the pool that you want to rename.

   Wait for the page to refresh and display the details.
4. If the pool is deactive, activate it by clicking Activate.

   Wait for the page to refresh and display the details.
5. Click *Rename*.

   This opens the *Rename a Pool* Wizard that guides you through the process.
6. 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 5.4, "Naming NSS Storage Objects," on page 67.

7 If the *Pools* page does not automatically update to show the new name for the pool, in *Roles and Tasks*, click *Pools* to refresh the current page.

8 For an NSS pool on Linux, activate the pool if it is deactive, then mount the pools's volumes.

16.6 Deleting a Pool

You might need to delete an NSS pool to create more free space for other pools. The *Delete* option on the *Pools* page removes one or more selected pools from the server, including all member partitions and the data on them. 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.

WARNING: Deleting a pool destroys all volumes in the pool and destroys all the data in them. These volumes cannot be restored.

You cannot use iManager to delete a system (`sys`) pool on NetWare. 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.

16.6.1 Prerequisites for Deleting a Pool

If the pool is shared in a Novell Cluster Services cluster, you must offline the cluster resource before you attempt to delete the clustered pool or its cluster resource.

If the pool has pool snapshots, you must delete the pool snapshots before you delete the pool. For information, see Section 18.11, "Deleting a Pool Snapshot," on page 250.

16.6.2 Procedure

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 In the *Pools* list, 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.

16.7 Viewing Pools on a Server

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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.

16.8 Viewing Pool Details

The *Pool Details* area of the *Pools* page displays information about a selected pool in the *Pools* list, as described in the following table:

Table 16-2 *Explanation of Pool Details*

Device Detail	Description
Name	The descriptive pool name assigned by the administrator.
Mount Point (Linux)	On your Linux system, this is the mount location for the NSS pool. The default mount location is `/opt/novell/nss/mnt/.pools/poolname`, where *poolname* is the name of the selected pool.
Partitions	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 *View Details* icon.
Number of Partitions	The total number of partitions currently assigned to the selected pool.
State	The current state of the selected pool, as Active or Deactive. Active pools are available to the users; deactive pools are not available to users.
LSS Type	The type of Loadable Storage System, such as ZLSS, CDDVD, or DOSFAT.
Share State	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 Devices page to set this device attribute.
Volumes	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 *View Details* icon. You can also select a pool, then click *Volumes* in *Roles and Tasks*.
Number of Volumes	The total number of volumes residing in the selected pool.

Device Detail	Description
Devices	A list of the descriptive device names of all logical devices contributing space in the selected pool.
Number of Devices	The total number of devices currently assigned to the selected pool.
Total Space	The total amount of space assigned to the selected pool.
Free Space	The total amount of space that is currently not in use on the selected pool.
Used Space	The total amount of space that is in use on the selected pool. ◆ **Purgeable Space:** 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. ◆ **Other In-Use Space:** The total amount of space in the selected pool that is currently in use and cannot be easily deleted without destroying data.
Block Size	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.
Creation Date	The time stamp (date and time) that the pool was created.
Last Update	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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. In the *Pools* list, select a pool to view its details.

   Wait for the page to refresh and display the pool’s details.
4. The pool must be active to display its details. If the *Details* area is empty, select the pool, then click *Activate*.

   When the page refreshes, you can view the pool’s details.

16.9 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. In the *Pools* list, 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 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.

16.10 Viewing Volume Information for a Pool

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 In the *Pools* list, 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 *Volumes* page opens with the server and volume preselected.

- Click *Storage* > *Volumes*.

  The *Volumes* 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 253.

16.11 Viewing Device Information for a Pool

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

When the page refreshes, a list of pools appears in the *Pools* list.

3 In the *Pools* list, 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) Select the device, then click *View Details* to view its details.

The *Devices* 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 127.

16.12 Preventing Pools from Activating on Multiple Servers

Multiple Server Activation Prevention (MSAP) prevents some accidental activations of a pool on more than one server at a time. You should never purposely attempt to activate a pool on two servers at the same time.

- Section 16.12.1, “Understanding MSAP,” on page 208
- Section 16.12.2, “Enabling or Disabling MSAP for All NSS Pools,” on page 209
- Section 16.12.3, “Enabling or Disabling MSAP for a Given NSS Pool,” on page 209
- Section 16.12.4, “Rebuilding the MSAP Block for a Given NSS Pool,” on page 211
- Section 16.12.5, “Determining If MSAP Is Enabled or Disabled on NSS Pools,” on page 212
- Section 16.12.6, “Managing MSAP with XML or APIs,” on page 212
- Section 16.12.7, “Additional Information,” on page 212

16.12.1 Understanding MSAP

MSAP is enabled by default for all pools on the server. When enabled, it helps prevent some accidental activations of a pool on more than one server at a time. It does not catch all multiple activations. 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.

16.12.2 Enabling or Disabling MSAP for All NSS Pools

By default, MSAP is enabled for all pools on the server when the server is booted.

To manually enable or disable MSAP for all pools on the server, issue the following MSAP console commands at the server console on NetWare, or in the NSS Console (`nsscon`) on Linux.

nss /msapserver

Enables MSAP for all the pools on the server. By default, MSAP is enabled for every pool on the server.

nss /nomsapserver

Disables MSAP for all the pools on the server. This command remains in effect only until the server is next rebooted.

IMPORTANT: We recommend that you never disable MSAP.

16.12.3 Enabling or Disabling MSAP for a Given NSS Pool

Use the procedures in this section to enable or disable MSAP for a given pool.

The `/PoolMSAP` option enables MSAP for a given pool on the server. Use the command when the pool is activated. MSAP is enabled the next time the pool is activated.

The `/NoPoolMSAP` option disables MSAP for a given pool. Use the command when the pool is activated. MSAP is disabled the next time the pool is activated.

IMPORTANT: We recommend that you never disable MSAP.

Linux

1. On Linux, open a terminal console, then log in as the `root` user.
2. If the pool is not active, activate it now.
   - **2a** Start NSSMU by entering the following at the terminal console prompt:

     ```
     nssmu
     ```
   - **2b** Go to the *Pools* page.
   - **2c** Select the pool, then activate it by pressing *F7*.
   - **2d** Exit NSSMU.
3. Enable or disable MSAP for a given pool.
   - **3a** Start the NSS Console by entering the following at the terminal console prompt:

     ```
     nsscon
     ```
   - **3b** At the `nsscon` prompt, do one of the following:
     - **Enable MSAP:** Enter

       ```
       nss /poolmsap=poolname
       ```
     - **Disable MSAP:** Enter

```
nss /nopoolmsap=poolname
```

3c Close the NSS Console by entering

```
exit
```

4 Deactivate the pool, then activate it again.

4a Start NSSMU by entering the following at the terminal console prompt:

```
nssmu
```

4b Go to the *Pools* page.

4c Select the pool, then deactivate it by pressing *F7*.

4d Select the pool, then activate it by pressing *F7* again.

4e Exit NSSMU.

MSAP is now enabled or disabled, depending on your action in Step 3.

5 Verify that MSAP is enabled or disabled for the given pool.

5a Start the NSS Console by entering the following at the terminal console prompt:

```
nsscon
```

5b At the `nsscon` prompt, enter

```
nss /pools
```

5c Review the messages to determine if the pool was successfully enabled or disabled as follows:

- **MSAP Enabled:** The *Multi-Use Detect* message is displayed for the pool.
- **MSAP Disabled:** The *Multi-Use Detect* message is not displayed for the pool.

5d Close the NSS Console by entering

```
exit
```

NetWare

1 If the pool is not active, activate it now by entering the following at the server console prompt:

```
nss /poolactivate=poolname
```

2 Enable or disable MSAP by entering one of the following commands at the server console prompt:

- **Enable MSAP:** Enter

```
nss /poolmsap=poolname
```

- **Disable MSAP:** Enter

```
nss /nopoolmsap=poolname
```

```
nss /poolmsap=poolname
```

3 Deactivate the pool by entering

```
nss /pooldeactivate=poolname
```

4 Activate the pool by entering

```
nss /poolactivate=poolname
```

MSAP is now enabled or disabled, depending on your action in Step 2.

5 Verify that MSAP is enabled or disabled for the given pool.

5a At the server console prompt, enter

```
nss /pools
```

5b Review the messages for the pool to determine if MSAP was successfully enabled or disabled as follows:

- **MSAP Enabled:** The *Multi-Use Detect* message is displayed for the pool.
- **MSAP Disabled:** The *Multi-Use Detect* message is not displayed for the pool.

16.12.4 Rebuilding the MSAP Block for a Given NSS Pool

If the MSAP block for a pool becomes corrupt, it prevents a pool from going into the Maintenance state. Use the `/MSAPRebuild` option to rebuild a pool's corrupt MSAP block. Before issuing the command to rebuild, you must deactivate the pool. Rebuilding an MSAP block does not give the rebuilder ownership of the pool.

Linux

1 On Linux, open a terminal console, then log in as the `root` user.

2 Deactivate the pool.

2a Start NSSMU by entering the following at the terminal console prompt:

```
nssmu
```

2b Go to the *Pools* page.

2c Select the pool, then deactivate it by pressing *F7*.

2d Exit NSSMU.

3 Rebuild the MSAP block for the pool.

3a Start the NSS Console by entering the following at the terminal console prompt:

```
nsscon
```

3b At the `nsscon` prompt, enter

```
nss /msaprebuild=poolname
```

3c Close the NSS Console by entering

```
exit
```

4 Activate the pool.

4a Start NSSMU by entering the following at the terminal console prompt:

```
nssmu
```

4b Go to the *Pools* page.

4c Select the pool, then activate it by pressing *F7*.

4d Exit NSSMU.

The pool should now be able to be placed in Maintenance mode.

NetWare

1. Deactivate the pool by entering the following at the server console prompt:

   ```
   nss /pooldeactivate=poolname
   ```

2. At the NetWare server console prompt, enter

   ```
   nss /msaprebuild=poolname
   ```

3. Activate the pool by entering the following at the server console prompt:

   ```
   nss /poolactivate=poolname
   ```

   The pool should now be able to be placed in Maintenance mode.

16.12.5 Determining If MSAP Is Enabled or Disabled on NSS Pools

The `nss /pools` command displays the message *Multi-Use Detect* for NSS pools that have MSAP enabled.

1. At the NetWare server console prompt, or at the NSS Console prompt on Linux, enter

   ```
   nss /pools
   ```

2. For each pool, review the messages to determine whether MSAP is enabled or disabled as follows:
   - **MSAP Enabled:** The *Multi-Use Detect* message is displayed for the pool.
   - **MSAP Disabled:** The *Multi-Use Detect* message is not displayed for the pool.

16.12.6 Managing MSAP with XML or APIs

The `_admin\manage_nss\pool\`*`poolname`*`\zlss\msap.xml` file contains MSAP statistics for the pool. One file exists for each 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.

16.12.7 Additional Information

For more information about the MSAP commands used in this section, see Section A.24, “Multiple Server Activation Prevention (MSAP) Commands,” on page 477.

16.13 Updating eDirectory Pool Objects

On the *Pools* 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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Pools* list, 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*.

16.14 What's Next

You can now create volumes in the NSS pools you created. For information, see Chapter 19, "Managing NSS Volumes," on page 253.

Verifying and Rebuilding NSS Pools and Volumes

17

This section describes how to verify and rebuild Novell® Storage Services™ pools to restore the consistency of a pool's metadata structure, and thus, the metadata structure of its volumes.

- Section 17.1, "When to Use Verify and Rebuild," on page 215
- Section 17.2, "Verifying and Rebuilding an NSS Pool and Its Volumes," on page 217
- Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223

17.1 When to Use Verify and Rebuild

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. All of the volumes on the pool are part of the verify or rebuild process.

The purpose of the Pool Verify and Pool Rebuild utilities is to make sure you have a valid metadata structure for a pool. Use the utilities only when you have problems with the pool's metadata structure.

- Section 17.1.1, "What Verifying Does Not Do," on page 215
- Section 17.1.2, "What Rebuilding Does Not Do," on page 215
- Section 17.1.3, "Before You Verify a Pool," on page 216
- Section 17.1.4, "Before You Rebuild a Pool," on page 216

17.1.1 What Verifying Does Not Do

Verifying a pool does not fix any problems. It is a read-only assessment of the pool's metadata structure to identify the types of errors, the severity of errors, and in what volumes the errors occur.

17.1.2 What Rebuilding Does Not Do

Rebuilding a pool restores the consistency of the pool's metadata structure. Rebuilding a pool does not restore lost data and does not repair the data integrity of corrupted data.

Rebuilding a pool does not fix problems for the following:

- Journaling errors
- Hardware and media errors
- File system trustee assignments, trustee rights, and inherited rights filters
- File system attributes for files and directories
- Opportunistic locking
- Content of files

17.1.3 Before You Verify a Pool

Volume errors are typically transactions left unfinished during a system crash of some kind. Most volume errors are fixed automatically during volume mount when NSS resolves the journaled errors. If the pool can be mounted, mount its volume to allow the NSS journaling feature to repair any transactional errors that occurred during a system failure.

Afterwards, there if there are still problems, use diagnostic tools to rule out hardware problems as the cause.

If non-hardware errors persist, and if you have a viable backup to restore the pool to the last known good state, restore the backup to recover the pool and restore the data. It is probably not necessary to verify or rebuild the pool.

If non-hardware errors persist, and if you do not have a viable backup, use the following Pool Verify utilities to identify any errors in the pool's metadata:

Operating System	Utilities
Linux	RAVSUI (verify option)
	RAVVIEW (reformats log files to human-readable format)
NetWare®	PoolVerify

17.1.4 Before You Rebuild a Pool

Review the verification log to determine the type and severity of problems with the pool's metadata.

If all of the following conditions exist, then you should rebuild the pool to restore its metadata integrity.

- Errors were not corrected by mounting the volume, or you could not mount the volume.
- Errors were not caused by media or hardware problems, or they persisted after correcting any media or hardware issues.
- You have no viable backup of the pool's volumes to restore the pool to the last known good state.
- The Pool Verify process reports errors in the physical integrity of any of the volumes' metadata that would definitely cause data corruption if no action is taken.
- More data will be lost from continued data corruption than will be lost from rebuilding the pools now.

WARNING: You should rebuild a pool only as a last resort to restore the consistency of the pool's metadata. The rebuild repairs the metadata; it does not recover lost data or repair the integrity of the data itself. Data loss occurs during a rebuild if the utility must prune leaves in the data structure to restore metadata consistency.

If the Pool Verify process did not report errors, but you cannot create files or directories, you should run rebuild with the ReZID option. For information, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.

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. Then if the rebuild is acceptable, you can replace the pool with the rebuilt snapshot. For information about pool snapshots, see Chapter 18, “Managing NSS Pool Snapshots,” on page 229.

If necessary, rebuild the pool’s metadata by using the following utilities:

Operating System	Utilities
Linux	RAVSUI (build option)
	RAVVIEW (reformats log files to human-readable format)
NetWare	PoolRebuild

17.2 Verifying and Rebuilding an NSS Pool and Its Volumes

- Section 17.2.1, “Mounting the Volume to Repair Journaled Errors,” on page 217
- Section 17.2.2, “Ruling Out Hardware Causes,” on page 217
- Section 17.2.3, “Verifying the Pool to Identify Metadata Inconsistencies,” on page 218
- Section 17.2.4, “Reviewing Log Files for Metadata Consistency Errors,” on page 219
- Section 17.2.5, “Rebuilding NSS Pools to Repair Metadata Consistency,” on page 221

17.2.1 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 NSS journaling feature. Journaling in NSS handles the same level of problems as `Vrepair` does on NetWare Traditional volumes.

If errors persist after you mount the volume, or if you cannot mount the volume, first rule out hardware causes for the problems. For information, see Section 17.2.2, “Ruling Out Hardware Causes,” on page 217.

17.2.2 Ruling Out Hardware Causes

If a volume cannot be mounted or problems persist after journaling errors are resolved, check the hardware for faulty media or controller problems.

1. Make sure you have a good backup of the data.
2. Use the latest diagnostic software and utilities from the manufacturer of your hard drives and controllers to troubleshoot the hard drives without destroying the data.

   For example, verify the media integrity and that devices are operating correctly.
3. If necessary, repair the media or controllers.

If errors persist after you have ruled out hardware causes, and you do not have a viable backup to restore to the last known good state, you should check the pool for metadata inconsistencies. For information, see Section 17.2.3, “Verifying the Pool to Identify Metadata Inconsistencies,” on page 218.

17.2.3 Verifying the Pool to Identify Metadata Inconsistencies

The verify process is a read-only assessment of the pool. The Pool Verify option searches the pool for inconsistent data blocks or other errors in the file system's metadata and reports data in the verification log. For information on where to find the verification log and how to interpret any reported errors, see Section 17.2.4, "Reviewing Log Files for Metadata Consistency Errors," on page 219.

Following one of these procedures to verify the pool:

- "Linux" on page 218
- "NetWare" on page 219

Linux

1 For a 32-bit machine, make sure you have enough space available in the Linux kernel cache memory to run a pool verify.

When running `ravsui(8)` for a pool verify or a pool rebuild on Linux, the utility needs contiguous space in kernel memory separate from the space allocated to the core NSS process. The larger the pool, the larger the space that is needed. To make space available, you might need to reduce the space used by other processes. You can optionally reduce the minimum number of buffers reserved for the core NSS process to as little as 10,000 4-KB buffers.

1a Open a terminal console as the `root` user.

1b At the console prompt, enter

```
nsscon
```

1c In nsscon, enter

```
nss MinBufferCacheSize=10000
```

2 Place the pool in maintenance mode.

2a At a terminal prompt, enter

```
nsscon
```

2b In nsscon, enter

```
nss /PoolMaintenance=poolname
```

3 Start the pool verify by entering the following at the terminal console prompt:

```
ravsui verify poolname
```

4 Use RAVVIEW to read the logs.

For information about using RAVVIEW, see Section B.17, "RAVVIEW (Linux)," on page 525.

5 Do one of the following:

- If the log reports no errors with the pool's metadata, it is safe to activate the pool and mount the volumes.

- If the log reports no errors with the pool's metadata, but you still cannot create files or directories, run a Pool Rebuild with the ReZID option. For information, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.
- If the log reports errors with the pool's metadata, the affected volumes remain in Maintenance mode. Decide whether to rebuild the pool based on the type of error and potential outcomes. For information about rebuilding the pool, see Section 17.2.5, "Rebuilding NSS Pools to Repair Metadata Consistency," on page 221.

6 For a 32-bit machine, if you modified the MinBufferCacheSize setting in Step 1, you can change it back to its original setting now, unless you are continuing with a pool rebuild.

6a Open a terminal console as the `root` user.

6b At the console prompt, enter

```
nsscon
```

6c In nsscon, enter

```
nss MinBufferCacheSize=value
```

Replace *value* with the desired minimum number of 4-KB buffers. The default value is 30000.

NetWare

1 Place the pool in maintenance mode by entering the following at the server console prompt:

```
nss /PoolMaintenance=poolname
```

2 Verify the pool by entering the following at the server console prompt:

```
nss /poolverify=poolname
```

3 Review any errors on-screen or in the *volume_name*`.rlf` file, located at the root of the DOS drive.

4 Do one of the following:

- If the log reports no errors with the pool's metadata, the pools and volumes are automatically activated. It is safe to mount the volumes.
- If the log reports no errors with the pool's metadata, but you still cannot create files or directories, run a Pool Rebuild with the ReZID option. For information, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.
- If the log reports errors with the pool's metadata, the volumes affected remain in Maintenance mode. Decide whether to rebuild the pool based on the type of error and potential outcomes. For information about rebuilding the pool, see Section 17.2.5, "Rebuilding NSS Pools to Repair Metadata Consistency," on page 221.

17.2.4 Reviewing Log Files for Metadata Consistency 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 On-Screen Display" on page 220
- "Warnings Reported" on page 220

- “Errors Reported” on page 221
- “No Errors Reported, but Cannot Create Files or Directories” on page 221

Log Files and On-Screen Display

Messages are written to the following logs:

Table 17-1 *Location of Log Files for the NSS Pool Verify and Pool Rebuild Utilities*

Platform	Log	Purpose
Linux	`/var/opt/novell/log/nss/rav/`*`filename`*`.vbf` 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 Section B.17, “RAVVIEW (Linux),” on page 525.
	`/var/opt/novell/log/nss/rav/`*`filename`*`.rtf`	Log of the pool rebuild process using `ravsui rebuild`. This log contains information about data that has been lost during a rebuild by the pruning of leaves in the data structure.
NetWare	`filename.vlf`, located at the root of the server’s DOS drive.	Log of the pool verify process using `poolverify`. If a volume in the pool has errors, the errors are displayed on the screen and written to this log file of errors and transactions.
	`filename.rlf`, located at the root of the server’s DOS drive.	Log of the pool rebuild process using `poolrebuild`. This log contains information about data that has been lost during a rebuild by the pruning of leaves in the data structure.

Whenever you verify or rebuild a pool, the new information is appended at the end of the log file. If you want to keep old log files intact, rename the log file or move it to another location before you start the verify or rebuild process.

Warnings Reported

Warnings indicate that there are problems with the metadata, but that there is no threat of data corruption. Performing a data restore from a backup tape or rebuilding the pool’s metadata are optional. However, rebuilding a pools’s metadata typically results in some data loss.

Errors Reported

Errors indicate that there are physical integrity problems with the pool's metadata, and data corruption will definitely occur, or it will continue to occur, if you continue to use the pool as it is.

If you decide to rebuild the pool, use the Pool Rebuild utility. For information, see Section 17.2.5, "Rebuilding NSS Pools to Repair Metadata Consistency," on page 221.

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 volumes in the pool, it might be because the files' ID numbers have exceeded the maximum size of file numbering field. You might need to rebuild the pool with the ReZID option. For information about how to decide if a ReZID is needed, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.

17.2.5 Rebuilding NSS Pools to Repair Metadata Consistency

The purpose of a pool rebuild is to repair the metadata consistency of the file system. 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.

A pool rebuild depends on many variables, so it is difficult to estimate how long it might take. The number of storage objects in a pool, such as volumes, directories, and files, is the primary consideration in determining the rebuild time, not the size of the pool. This is because a pool rebuild is reconstructing the metadata for the pool, not its data. For example, it would take longer to rebuild the metadata for a 200 GB pool with many files than for a 1 TB pool with only a few files. Other key variables are the number of processors, the speed of the processors, and the size of the memory available in the server.

You do not need to bring down the server to rebuild a pool. NSS allows you to temporarily place an individual storage pool in maintenance mode while you verify or rebuild it. While the pool is deactivated, users do not have access to any of the volumes in that pool.

If you do not place the pool in maintenance mode before issuing the rebuild or verify commands, you receive NSS Error 21726:

```
NSS error: PoolVerify results

   Status: 21726

     Name: zERR_RAV_STATE_MAINTENANCE_REQUIRED

   Source: nXML.cpp[1289]
```

Following one of these procedures to rebuild the pool:

- "Linux" on page 222
- "NetWare" on page 223

Linux

1. Depending on the nature of the reported errors, you might want to open a call with Novell Support before you begin the rebuild process.
2. For a 32-bit machine, make sure you have enough space available in the Linux kernel cache memory to run a pool rebuild.

   When running `ravsui(8)` for a pool verify or a pool rebuild on Linux, the utility needs contiguous space in kernel memory separate from the space allocated to the core NSS process. The larger the pool, the larger the space that is needed. To make space available, you might need to reduce the space used by other processes. You can optionally reduce the minimum number of buffers reserved for the core NSS process to as little as 10,000 4-KB buffers.

   2a Open a terminal console as the `root` user.

   2b At the console prompt, enter

   ```
   nsscon
   ```

   2c In nsscon, enter

   ```
   nss MinBufferCacheSize=10000
   ```
3. Place the pool in maintenance mode.

   3a At a terminal prompt, enter

   ```
   nsscon
   ```

   3b In nsscon, enter

   ```
   nss /PoolMaintenance=poolname
   ```
4. Start the pool rebuild. At the terminal console prompt, enter

   ```
   ravsui rebuild poolname
   ```

   For information, see Section B.16, “RAVSUI (Linux),” on page 522 for options to set the pruning parameters for the rebuild.

   Rebuilding can take several minutes to several hours, depending on the number of storage 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 17.2.4, “Reviewing Log Files for Metadata Consistency Errors,” on page 219.
6. Do one of the following:

   - **No Errors:** If errors do not exist at the end of the rebuild, the pool’s volumes are available for mounting.
   - **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 Novell Support for assistance.
7. For a 32-bit machine, if you modified the MinBufferCacheSize setting in Step 2, you can change it back to its original setting.

   7a Open a terminal console as the `root` user.

   7b At the console prompt, enter

```
nsscon
```

7c In nsscon, enter

```
nss MinBufferCacheSize=value
```

Replace *value* with the desired minimum number of 4-KB buffers. The default value is 30000.

NetWare

1. Depending on the nature of the reported errors, you might want to open a call with Novell Support before you begin the rebuild process.
2. Place the pool in maintenance mode. At a terminal prompt, enter

   ```
   nss /PoolMaintenance=poolname
   ```

3. 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.
4. 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 17.2.4, “Reviewing Log Files for Metadata Consistency Errors,” on page 219.
5. Do one of the following:
   - **No Errors:** If errors do not exist at the end of the rebuild, the pool is activated automatically. It is safe to mount the volumes.
   - **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 Novell Support for assistance.

17.3 ReZIDing Volumes in an NSS Pool

- Section 17.3.1, “What Is a ZID?,” on page 223
- Section 17.3.2, “Understanding ReZID,” on page 224
- Section 17.3.3, “When to ReZID,” on page 224
- Section 17.3.4, “Viewing the Highest ZID for a Volume,” on page 225
- Section 17.3.5, “ReZIDing Volumes on Linux,” on page 225
- Section 17.3.6, “ReZIDing Volumes on NetWare,” on page 226

17.3.1 What Is a ZID?

When a file is created, it is assigned a unique file number, called a ZID. In NSS, the maximum number of file ZIDs 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 support up to 4 billion (4E9) ZIDs. NSS restricts ZIDs, and 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.

The ZID of a file is an internal file system bit of information and is not viewable. However, the highest ZID number in use for each volume is reported when you verify a pool.

17.3.2 Understanding ReZID

The ReZID option for a pool rebuild changes the ZIDs for all the files on the volume, thus freeing ZIDs so they are available for creating new files and directories. The rezid does not modify any other metadata on the volumes, nor does it modify any file's content. The reZID is unrelated to any other rebuild activities that might occur.

IMPORTANT: The reZID step in a rebuild adds a third review of the pool and can increase the time of a rebuild by 50%.

17.3.3 When to ReZID

After verifying a pool, the log reports the highest ZID (highestZID parameter) for each of the pool's volumes. If the highest ZID number is close to the 4 billion ZID limit (4E9), you should reZID the volume. For NSS volumes on Linux, if the highest ZID is 2 billion (2E9) or greater, the reZID option occurs automatically when you rebuild a pool. You can optionally specify a different ZID limit to trigger the rezid, or use the /noReZID option to stop the reZID from occurring with that rebuild.

There are no errors reported if ZIDs are nearing the upper limit of 4 billion for a volume. You might get errors creating a file or directory that suggest a reZID needs to be done. For example:

- NDS database is locked.
- Server hangs at the end of load stage 1.
- Cannot copy to a volume.
- Cannot copy to `sys:` volume.

NSS API calls return `Error 20108 zERR_ZID_GREATER_THAN_32_BITS`, which means that the ZID numbering has exceeded the 4 billion (4E9) limit. NSS also sends a volume alert to the server console that reZID needs to be done on a specified volume. The calling application gets only a generic error when it attempts and fails to create 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 again, then check the highest ZID number reported for the pool's volumes to know that each is well under the 4 billion ZIDs limit.

If you do not place the pool in maintenance mode before rebuilding the pool with the ReZID option, you receive NSS Error 21726:

```
NSS error: PoolVerify results

   Status: 21726

     Name: zERR_RAV_STATE_MAINTENANCE_REQUIRED
```

```
Source: nXML.cpp[1289]
```

17.3.4 Viewing the Highest ZID for a Volume

When you verify a pool, look in the log to find the highest ZID value that has been assigned for each of the pool's volume. Look at each value to see whether you should rezid the pool as part of the rebuild process.

You should be aware of the rate at which you are consuming ZIDs by creating and deleting files. If the highest ZID for a given volume reaches the limit of 4 billion (4E9), you cannot create new files on the volume until you rezid the pool.

17.3.5 ReZIDing Volumes on Linux

1. For a 32-bit machine, make sure you have enough space available in the Linux kernel cache memory to run a pool rebuild.

   When running `ravsui(8)` for a pool verify or a pool rebuild on Linux, the utility needs contiguous space in kernel memory separate from the space allocated to the core NSS process. The larger the pool, the larger the space that is needed. To make space available, you might need to reduce the space used by other processes. You can optionally reduce the minimum number of buffers reserved for the core NSS process to as little as 10,000 4-KB buffers.

   1a Open a terminal console as the `root` user.

   1b At the console prompt, enter

   ```
   nsscon
   ```

   1c In nsscon, enter

   ```
   nss MinBufferCacheSize=10000
   ```

2. Place the pool in maintenance mode.

   2a At a terminal prompt, enter

   ```
   nsscon
   ```

   2b In nsscon, enter

   ```
   nss /PoolMaintenance=poolname
   ```

3. If you have not already verified the volume, enter the following at a command prompt:

   ```
   ravsui verify poolname
   ```

   For information, see Section B.16, "RAVSUI (Linux)," on page 522.

4. Review any errors on-screen or in the *`filename`*`.vbf` file, located where you specified.

   For information, see Section 17.2.4, "Reviewing Log Files for Metadata Consistency Errors," on page 219.

5. 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 `0xeffffffe`. For information, see Section B.16, "RAVSUI (Linux)," on page 522 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 the rebuild, which increases the time to rebuild a volume by about 50%.

6 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 17.2.4, "Reviewing Log Files for Metadata Consistency Errors," on page 219.

7 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 Novell Support for assistance.
- **No Errors:** If errors do not exist, the pool's volumes are mounted automatically.

8 For a 32-bit machine, if you modified the MinBufferCacheSize setting in Step 1, you can change it back to its original setting.

8a Open a terminal console as the root user.

8b At the console prompt, enter

```
nsscon
```

8c In nsscon, enter

```
nss MinBufferCacheSize=value
```

Replace *value* with the desired minimum number of 4-KB buffers. The default value is 30000.

17.3.6 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 17.2.4, "Reviewing Log Files for Metadata Consistency Errors," on page 219.

4 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 the 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 17.2.4, “Reviewing Log Files for Metadata Consistency Errors,” on page 219.

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 Novell Support for assistance.
- **No Errors:** If errors do not exist, the pool’s volumes are mounted automatically.

Managing NSS Pool Snapshots 18

Novell® Storage Services™ supports pool snapshots to improve backup and restore services. This section describes the following:

- Section 18.1, “Understanding Pool Snapshots,” on page 229
- Section 18.2, “Guidelines for Using and Managing Pool Snapshots,” on page 231
- Section 18.3, “Creating a New Pool Snapshot,” on page 238
- Section 18.4, “Viewing a List of Snapshots for a Given Pool,” on page 242
- Section 18.5, “Viewing Pool Snapshot Details,” on page 244
- Section 18.6, “Renaming a Pool Snapshot (NetWare),” on page 245
- Section 18.7, “Modifying the Stored-On Location for Snapshots,” on page 246
- Section 18.8, “Onlining or Offlining a Pool Snapshot,” on page 246
- Section 18.9, “Viewing and Managing an Online Pool Snapshot,” on page 248
- Section 18.10, “Restoring Data from an Online Pool Snapshot,” on page 250
- Section 18.11, “Deleting a Pool Snapshot,” on page 250
- Section 18.12, “Shredding a Deleted Pool Snapshot (NetWare),” on page 251
- Section 18.13, “Using Commands to Manage Pool Snapshots (NetWare),” on page 252

18.1 Understanding Pool Snapshots

- Section 18.1.1, “How Snapshots Work,” on page 229
- Section 18.1.2, “Benefits of Using Snapshots,” on page 230

18.1.1 How Snapshots Work

A pool snapshot is a metadata copy of a storage data pool that preserves a point-in-time view of a data pool. 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 snapshot. 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 snapshot.

Before the snapshot can occur, the snapshot function must quiesce the original pool 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 if you activate the snapshot for access.

After the snapshot, 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 temporarily suspends the write activity while it copies the original block data to the designated pool where it stores the pool snapshot. The snapshot storage area is referred to as the *stored-on pool* on NetWare® and as the *stored-on partition* on Linux. After the data is copied, the snapshot function allows the write to that block in the original pool. This copy-on-write process keeps the snapshot metadata consistent in time with the exact instant the snapshot was taken.

As data on the original pool changes, the snapshot can theoretically 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. A combination of up to 500 snapshots on NetWare can exist on any given stored-on pool and 15 snapshots on OES 2 Linux on any given stored-on partition. The amount of space required depends on the number of snapshots, the snapshot retention policy, and the turnover rate for data in the original pool.

While the snapshot exists, the performance for volumes on the pool can decrease slightly because the number of disk writes increases for the copy-on-write activity. The decrease depends on the volatility of your data and the number of pool snapshots that exist for the original pool.

18.1.2 Benefits of Using 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" on page 230
- "Archiving Data" on page 230
- "Restoring Data" on page 231
- "Re-Creating Operational and Development Environments" on page 231
- "Testing and Training" on page 231

Supporting Backup Operations

A pool snapshot facilitates non-disruptive backups because the snapshot becomes the source of the backup. When you back up volumes in a pool from a 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.

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 maintain a point-in-time history of the changes made to the original data pool.

Restoring Data

Pool snapshots can serve as a source for recovering a point-in-time version of a file. 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. To restore data, manually copy the old version of the file from the online snapshot volume to the original volume. For information, see Section 18.8, "Onlining or Offlining a Pool Snapshot," on page 246.

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 restore 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 snapshot.

Testing and Training

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

18.2 Guidelines for Using and Managing Pool Snapshots

Use the guidelines in this section when planning your snapshot solution:

- Section 18.2.1, "Differences Between Snapshots on Linux and NetWare," on page 232
- Section 18.2.2, "Guidelines for Creating a Pool Snapshot," on page 233
- Section 18.2.3, "Guidelines for Creating Pool Snapshots of Clustered Pools," on page 233
- Section 18.2.4, "Guidelines for Naming Pool Snapshots," on page 233
- Section 18.2.5, "Guidelines for Choosing the Stored-On Location," on page 234
- Section 18.2.6, "Guidelines for Maintaining the Stored-On Location," on page 235
- Section 18.2.7, "Guidelines for Onlining Pool Snapshots," on page 236
- Section 18.2.8, "Guidelines for Deleting Pool Snapshots," on page 237
- Section 18.2.9, "Guidelines for Viewing a List of Snapshots that are Stored in a Pool (NetWare)," on page 237
- Section 18.2.10, "Guidelines for Shredding Deleted Snapshots (NetWare)," on page 237

18.2.1 Differences Between Snapshots on Linux and NetWare

For NSS on Linux, snapshot volumes are not automatically mounted on reboot as they are in NetWare. The snapshots are active and performing their snapshot functions, but they are not mounted. If a snapshot was mounted when the server went down and you want the snapshot mounted after reboot, you must mount it manually. Mounting a snapshot is necessary only if users require access to the point-in-time version of the data.

In NSSMU, devices that contain NSS pool snapshots cannot be re-initialized. To initialize the device, you must first delete all NSS pool snapshots on the device. For information about deleting snapshots, see Section 18.11, "Deleting a Pool Snapshot," on page 250.

Table 18-1 identifies the differences for using snapshots for NSS pools on Linux and NetWare:

Table 18-1 *Comparison of NSS Pool Snapshots on Linux and NetWare*

Capability	NetWare	Linux
Rename a snapshot	Supported	Not supported. If you attempt to rename a snapshot, you get an eDirectory error because eDirectory objects are not automatically created for pool snapshots on Linux.
Snapshot stored-on location	An NSS pool is designated as the stored-on pool for a given original pool.	An EVMS-managed Linux partition is designated as the stored-on partition for a given original pool.
Snapshots of cluster-enabled pools	Supported The stored-on pool must be the same as the original pool.	Not supported
eDirectory object for the snapshot	eDirectory objects are automatically created for the snapshot pools and volumes.	No eDirectory objects are created. In order to allow users to access data on an NSS pool snapshot (snapshots can only be accessed as read only, of course), you must first activate and mount the snapshot as an NSS pool, then use the *Update eDirectory* option on the *Storage* > *Pools* page to create an eDirectory object for the snapshot pool or volume.
Taking new snapshots	The stored-on pool must be activated and mounted.	The stored-on Linux partition must be mounted on Linux.
Deleting snapshots	Delete pool snapshots in a first-created, first-deleted order.	Snapshots can be deleted in any sequence.
Shredding deleted snapshots	Supported	Not supported

18.2.2 Guidelines for Creating a Pool Snapshot

Create a pool snapshot when you want to capture a point-in-time view of a active data pool. The original pool must be active when you create the snapshot. On NetWare, the stored-on pool must also be active when you make the snapshot. For instructions on creating a pool snapshot, see "Creating a New Pool Snapshot" on page 238.

18.2.3 Guidelines for Creating Pool Snapshots of Clustered Pools

Pool snapshots are supported for clustered NSS pools on NetWare.

Pool snapshots are not supported for clustered NSS pools on Linux.

18.2.4 Guidelines for Naming Pool Snapshots

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

- "Default Naming Scheme When Using iManager" on page 233
- "Alternate Naming Scheme" on page 234
- "Considerations for Naming" on page 234

Default Naming Scheme When Using iManager

When you create a snapshot in iManager, the snapshot name is by default a modified version of the original pool's name, which allows a simple identification of all snapshots for any given pool. NSS adds a letter and number designator (_*Sn*) to the original poolname. The *S* indicates that it is a snapshot. The *n* represents an incremental number (1 to 500) of snapshots taken for this pool.

IMPORTANT: When you create a snapshot in NSSMU, no default name is suggested. You can optionally adopt the default naming scheme when you provide a name for the pool snapshot.

For example, pool snapshot names for `POOLA` might be `POOLA_S1`, `POOLA_S2`, and so on.

If you delete snapshots out of sequence, it is possible that the numbers could be reused. A simple sort by snapshot name could be confusing. Make sure to verify the create stamp on the pool when you work with pool snapshots that use the default naming scheme.

The snapshot name can be 2 to 16 characters in length and must adhere to the same character conventions as for poolnames. If the poolname is too long to allow the snapshot identifier to be appended, the poolname is truncated so that the length of the pool snapshot name is 16 characters. For example, if the poolname is `POOL_EUR_MANUF` (14 characters), its name would be truncated then the snapshot identifier appended. The number of characters to be truncated would depend upon the pool snapshot number, such as POOL_EUR_MANU_S1, POOL_EUR_MAN_S12, or POOL_EUR_MA_S102

If you bring a pool snapshot online, its volume names are automatically renamed to indicate that they are snapshot volumes. 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 pool snapshot name is `users_s1`. If its volumes are named `users_aj` and `users_ kz,` the volumes in the snapshot pool are `users_aj_sv`, `users_aj_sv001` and `users_kz_sv`, `users_kz_sv001` and so on.

Alternate Naming Scheme

You can optionally adopt your own naming convention for pools. If you create multiple snapshots of a pool each day, consider using a logical naming convention that identifies the poolname and numbers that allow sequential listing based on the order the snaps were taken. Of course, the time stamp shows the exact time that the snapshot was taken, and you can always refer to it to be sure you have the right snapshot.

Considerations for Naming

It is also important to consider the names of existing pools and pool snapshots and your naming conventions when you name new data pools and volumes. You should get errors if you attempt to create pools with names that are in use by pool snapshots, and vice versa.

If a volume with the same name as a snapshot volume exists on the server when you mount a snapshot pool, NSS automatically appends the snapshot volume name with an additional sequential number (such as `VOL1_SV001`) or characters (such as `VOL1_SV_SV`) as it onlines the volume. This makes the snapshot name unique with respect to the active volumes while the pool snapshot is online.

If name conflicts occur, you might need to rename a pool or pool snapshot to a unique name in order to bring the pool snapshot online.

18.2.5 Guidelines for Choosing the Stored-On Location

- "Stored-On Partition (Linux)" on page 234
- "Stored-On Pool (NetWare)" on page 235

Stored-On Partition (Linux)

The stored-on partition is the partition on an EVMS-managed device where you want to store a given pool snapshot for a given original pool.

IMPORTANT: On Linux, creating snapshots of clustered NSS pools is not supported.

When you create a pool's snapshot, you select the device and specify the size to use. You cannot increase the size of the partition later, so make sure you allow sufficient space.

If the destination partition that stores the snapshot data runs out of space, it causes write errors to happen on the original pool. 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 original pool.

Stored-On Pool (NetWare)

The stored-on pool is the active data pool where you want to store the pool snapshots for a given original 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 snapshot.

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 original pool you are snapshotting. The exception to this is clustering, where the destination pool for the snapshot data must be the same as the original pool.

If the destination pool that stores the snapshot data runs out of space, it causes write errors to happen on the pools being snapshot. 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 original pool.

18.2.6 Guidelines for Maintaining the Stored-On Location

- "Stored-on Partition (Linux)" on page 235
- "Stored-on Pool (NetWare)" on page 236

Stored-on Partition (Linux)

The status of any given pool snapshot partition is closely tied to the operational status of the original pool. You can deactivate the original pool, as needed, without adversely impacting the pool snapshot or the status of the stored-on partition. If the original pool is deactive, there are no active transactions for the pool snapshot function to process. For Linux, the stored-on partition hosts only the a single snapshot, so it can be safely deactivated after you deactivate its original pool. Make sure that you re-activate the stored-on partition first when bringing the original pool back into service.

In contrast, deactivating the stored-on partition first can cause the ungraceful deactivation of the corresponding original pool.

IMPORTANT: The stored-on partition should remain active as long as it hosts any pool snapshots. You can deactivate it safely after its original pool is deactivated, and for the duration of the pool's deactivation. Activate the stored-on partition before re-activating the original pool.

Stored-on Pool (NetWare)

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.

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.

18.2.7 Guidelines for Onlining Pool Snapshots

For pool snapshots, Online and Offline are conditions related to the visibility of the pool snapshot to users. The pool snapshot is offline by default. The snapshot functions are working in the background to capture any changes being made to the original pool whether the pool is offline or online.

You activate a pool snapshot as a pool by bringing it *online* whenever you want to access the data on it, such as for data retrieval and data backup. After the pool snapshot is online, it appears by its snapshot name in the pool list. Treat it as you would any pool to activate and mount its volumes. Because you are working with a snapshot and not the original pool and its volumes, other management tasks are limited.

The names of volumes on the pool snapshot are a modified version of the volumes on the original pool. By default, the characters *_SV* (snapshot volume) are appended to the volumes' names. When you deactivate the pool snapshot, the corresponding snapshot volumes are automatically deactivated, and the pool snapshot is no longer listed in the pool list.

For NSS on NetWare, when you bring the pool snapshot online, the Pool object and Volume objects are automatically created. They are automatically deleted when you bring the pool snapshot offline again.

For NSS on Linux, the Pool object and Volume object for a snapshot are not automatically created in Novell eDirectory™ when you bring an NSS volume snapshot online. These objects are needed only if you want to verify the NSS metadata information on a snapshot volume. For this case, you must create the objects manually by using the *Update eDirectory* option to create the storage objects for the online snapshot NSS pool and each of its volumes. For information, see Section 16.13, "Updating eDirectory Pool Objects," on page 213 and Section 19.5, "Updating eDirectory Volume Objects," on page 262.

If you reboot the server while a pool snapshot is online, the snapshot might be online or offline after the restart, depending on the platform. On NetWare, if the pool snapshot is online when the server goes down, the pool snapshot is automatically set in the Online state on restart. On Linux, if the pool snapshot is online when the server goes down, the pool snapshot is automatically set in the Offline state after a reboot.

For instructions, see Section 18.8, “Onlining or Offlining a Pool Snapshot,” on page 246.

18.2.8 Guidelines for Deleting Pool Snapshots

Delete pool snapshots as follows:

- Delete all pool snapshots before you move devices that contain the original pool cross-platform. Different technologies are used for pool snapshots on Linux and NetWare, so existing pool snapshots cannot be used on a different platform.

  WARNING: Without first deleting the pool’s snapshots, you might not be able to access or manage the original pool after you move the pool’s device cross-platform.

- Delete a pool snapshot when you no longer need it.
- On NetWare, delete a pool snapshot when you need to free space in the stored-on pool.
- On Linux, delete a pool snapshot when you need free space on the device where the stored-on partition exists.

Make sure that the pool snapshot is not mounted as an online pool before you delete it.

For NetWare, delete pool-level snapshots in a first-created, first-deleted order, deleting the oldest snapshot of the pool first. For Linux, snapshots can be deleted in any sequence.

For instructions, see Section 18.11, “Deleting a Pool Snapshot,” on page 250.

18.2.9 Guidelines for Viewing a List of Snapshots that are Stored in a Pool (NetWare)

If the original pool and the stored-on pool are the same pool, you can view a list of snapshots stored on the pool by viewing the pool's Pool page. However, a stored-on pool does not show snapshots that it contains from other pools. If you want to be able to know which snapshots are stored on a stored-on pool, we recommend the following best practices:

- Specify the original pool as the stored-on pool. This is mandatory for clustered pools.
- If you want to use a different pool to store snapshots, use a given pool as the stored-on pool for only a single original pool at a time. Don't mix snapshots from multiple original pools in a single stored-on pool.
- When you create a stored-on pool to use for snapshots, name it similarly to the original pool so that you know the association by the naming convention.

18.2.10 Guidelines for Shredding Deleted Snapshots (NetWare)

In NetWare 6.5 SP3 and later, when a snapshot is deleted, Media Manager automatically zeroes out all the blocks that it used. 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 snapshot shredding. If shredding is enabled, 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. For instructions, see Section 18.12, “Shredding a Deleted Pool Snapshot (NetWare),” on page 251.

IMPORTANT: Data shredding is not supported for pool snapshots on Linux.

18.3 Creating a New Pool Snapshot

- Section 18.3.1, “Prerequisites for Creating a Pool Snapshot,” on page 238
- Section 18.3.2, “Creating a Pool Snapshot on Linux,” on page 238
- Section 18.3.3, “Creating a Pool Snapshot on NetWare,” on page 240

18.3.1 Prerequisites for Creating a Pool Snapshot

- The pool you want to snapshot must already exist and be active.
- On NetWare, the pool you want to use as the stored-on pool must already exist. It can be the same pool as the pool you want to snapshot. In a cluster, you must use the same pool.
- On Linux, free space must be available on an EVMS-managed device that you want to use as the stored-on partition.
- On Linux, you cannot create snapshots of shared NSS pools.

18.3.2 Creating a Pool Snapshot on Linux

- “Using iManager” on page 238
- “Using NSSMU for Linux” on page 240

Using iManager

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. In the server field on the *Pools* page, select a server to manage to view a list of pools.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. If the pool you want to snapshot is not active, select the pool from the *Pools* list, then click *Activate*.
4. In the *Pools* list, select the active pool that you want to snapshot, then click *Snapshot* to go to the *Snapshots for: poolname* page.

   IMPORTANT: If the selected pool is a online pool snapshot, the *Snapshot* option is not available. Select the original pool instead.

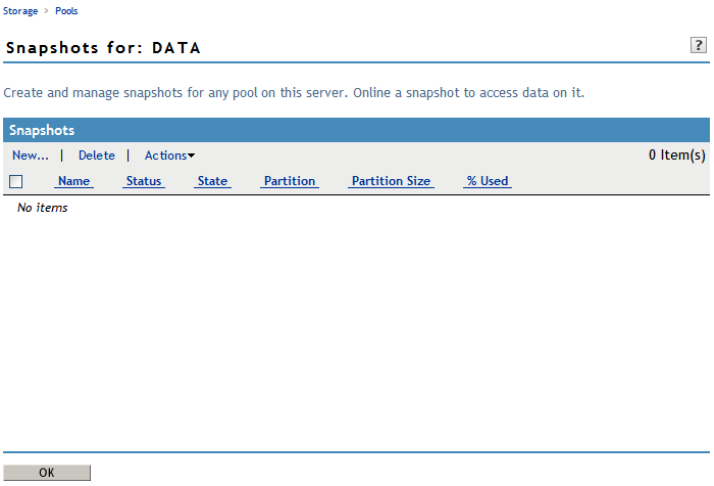

5 In the *Snapshots* menu, select *New* to open the *New Snapshot for: <poolname>* page.

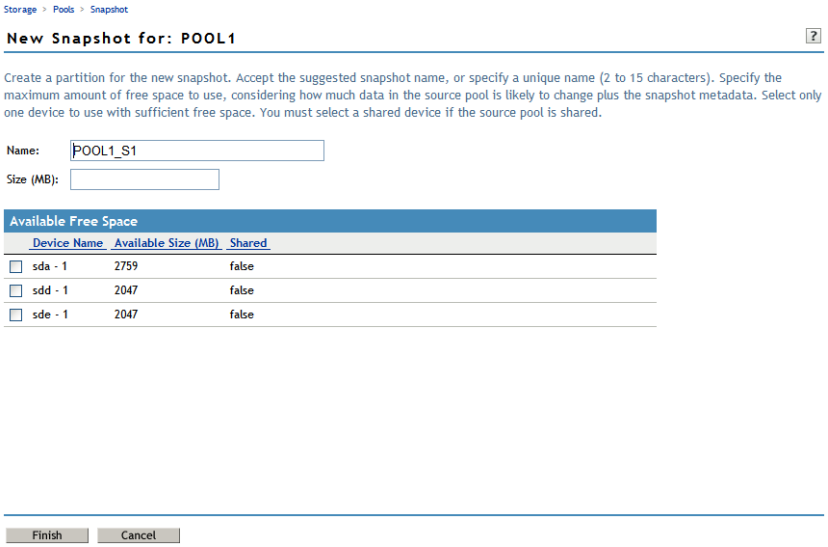

6 On the *New Snapshot* page, specify the following:

- **Name:** Optionally modify the default snapshot name.

  For information about pool snapshot names, see "Guidelines for Naming Pool Snapshots" on page 233.

- **Size:** Type the amount of free space (in MB) to use for the stored-on partition.

- **Stored-on Partition:** From the list of active devices, select the device where you want to create the stored-on partition.

  On Linux, each snapshot is stored on a separate partition. The partition for the snapshot cannot be expanded after it is created. If the pool is shared in a cluster, the snapshot feature is not supported. For information, see Section , "Stored-On Partition (Linux)," on page 234.

7 Click *Finish* to create the snapshot, or click *Cancel* to back out of the process.

8 After NSS creates the pool snapshot, NSS automatically opens to the *Snapshots* page so that you can further manage the snapshot. The *Snapshots* list contains the newly created snapshot.

IMPORTANT: You might see an error message if the iManager connection to the server you are managing times out before the snapshot is created. The pool snapshot creation should continue on the managed server. If a timeout error occurs, navigate to the *Snapshots* page to view and manage the snapshot.

By default, the snapshot on Linux is always *Offline* and *Active*. This means that the snapshot is functioning, but that the pool snapshot is not mounted as an online pool.

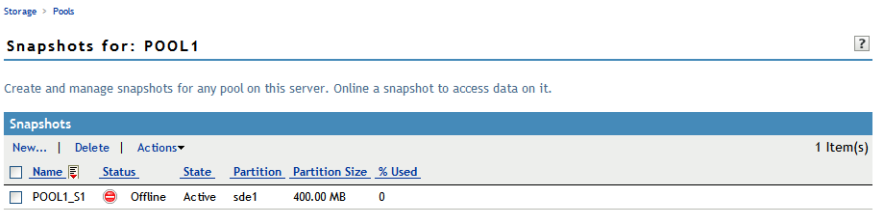

Using NSSMU for Linux

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, enter

   ```
   nssmu
   ```
3. If the pool that you want to snapshot is not currently mounted, mount it now by using your normal mount methods.
4. In NSSMU, select *Snapshot*.
5. Press *Insert* to begin the create process.
6. Specify a name for the pool snapshot.
7. From the *Pools* list, select the pool you want to snapshot, then press *Enter*.
8. From the *Devices* list, select the device where you want to create a partition for the snapshot stored-on partition, then press *Enter*.

   On Linux, each snapshot is stored on a separate partition, not on a pool.
9. Specify how much space in MB to allocate to the partition.

   The partition for the snapshot cannot be expanded after it is created.
10. Press *Enter* to create the snapshot.

    The newly created snapshot appears in the *Snapshots* list.

18.3.3 Creating a Pool Snapshot on NetWare

You must use iManager to create and manage pool snapshots for NSS volumes on NetWare. NSSMU for NetWare does not support snapshots.

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. In the server field on the *Pools* page, select a server to manage to view a list of pools.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 If the pool you want to snapshot is not active, select the pool from the *Pools* list, then click *Activate*.

4 For NetWare, if the pool you want to make the stored-on pool is not active, select the pool from the *Pools* list, then click *Activate*.

5 In the *Pools* list, select the active pool that you want to snapshot, then click *Snapshot* to go to the *Snapshots for: <poolname>* page.

IMPORTANT: If the selected pool is a online pool snapshot, the *Snapshot* option is not available. Select the original pool instead.

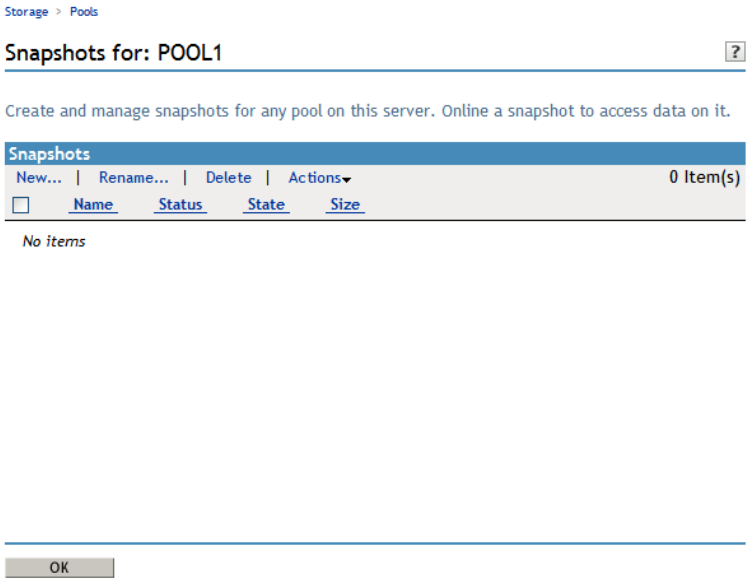

6 In the *Snapshots* menu, select *New* to open the *New Snapshot for: poolname* page.

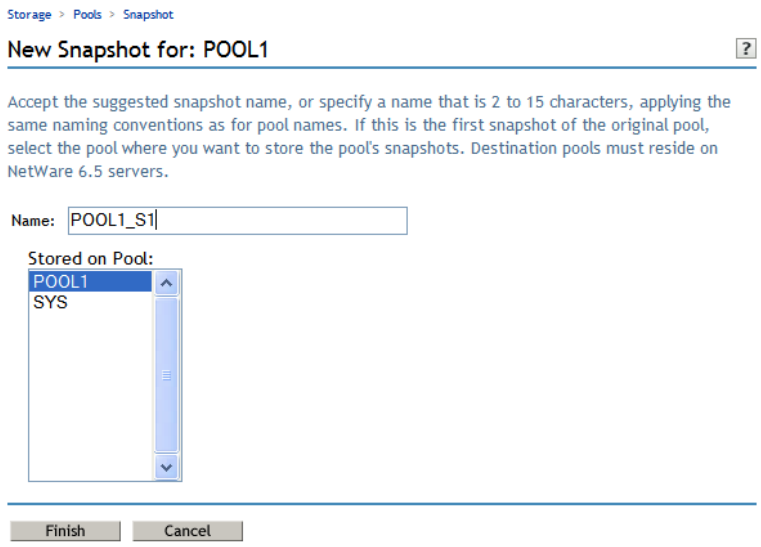

7 On the *New Snapshot* page, specify the following:

- **Name:** Optionally modify the default snapshot name.

For information about pool snapshot names, see "Guidelines for Naming Pool Snapshots" on page 233.

- **Stored-on Pool:** From the list of active pools, select the pool you want to use as the stored-on pool. If the pool is shared in a cluster, the stored-on pool must be the same pool as the pool you are snapshotting.

  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 series of pool snapshots. For information, see Section , "Stored-On Pool (NetWare)," on page 235.

8 Click *Finish* to create the snapshot, or click *Cancel* to back out of the process.

9 After NSS creates the pool snapshot, NSS opens to the *Snapshots* page so that you can further manage the snapshot. The *Snapshots* list contains the newly created snapshot.

IMPORTANT: You might see an error message if the iManager connection to the server you are managing times out before the snapshot is created. The pool snapshot creation should continue on the managed server. If a timeout error occurs, navigate to the *Snapshots* page to view and manage the snapshot.

By default, the snapshot on NetWare is always *Offline* and *Deactive*. This means that the snapshot is functioning, but that the pool snapshot is not activated as an online pool.

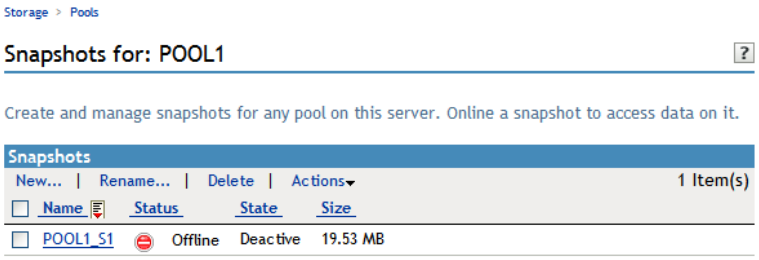

18.4 Viewing a List of Snapshots for a Given Pool

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 In the server field on the *Pools* page, select a server to manage to view a list of pools.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 If the original pool for the snapshots you want to view is not active, select the pool from the *Pools* list, then click *Activate*.

4 For NetWare, if the pool that contains the stored-on pool is not active, select the pool from the *Pools* list, then click *Activate*.

5 In the *Pools* list, select the active pool that has snapshots you want to manage, then click *Snapshot* to go to the *Snapshots for: poolname* page.

IMPORTANT: If the selected pool is an online pool snapshot, the *Snapshot* option is not available. Select the original pool instead.

On Linux, the *Snapshots* report includes the snapshot name, status (*Offline* (default) or *Online*), state (*Active* (default) or *Deactive*), the name and size of the stored-on partition, and the percent of space used on the partition.

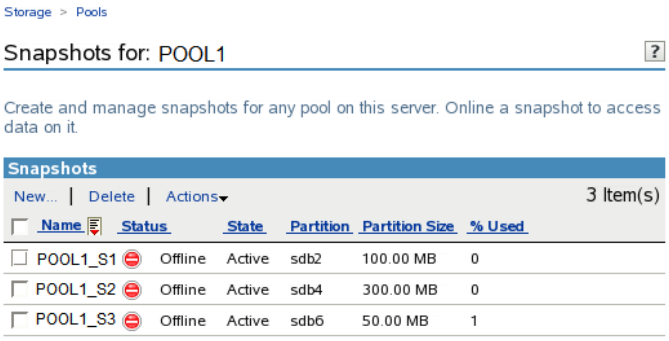

On NetWare, the *Snapshots* report includes the snapshot name, status (*Offline* (default) or *Online*), state (*Deactive* (default) or *Active*), and the size of the snapshot. For details, access the *Details* report by clicking *Actions* > *Details*. For more information, see Section 18.5, "Viewing Pool Snapshot Details," on page 244.

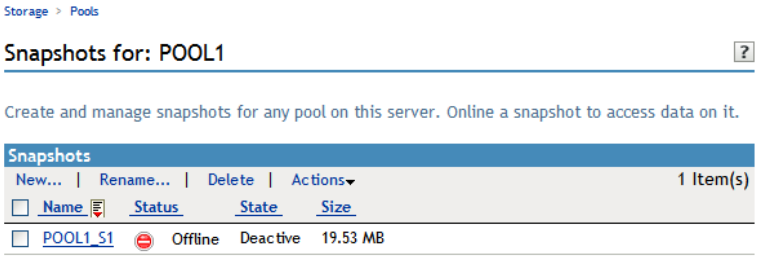

18.5 Viewing Pool Snapshot Details

In iManager, you can view the following details about a pool snapshot:

Table 18-2 *Explanation of Pool Snapshot Details*

Snapshot Details	Parameter	Description
Snapshot	Status	Offline (default) or Online. Online pool snapshots appear in the *Pools* page and its snapshot volumes appear in the *Volumes* page.
	State	Deactive (default) or Active. Active indicates that the pool snapshot operation is active. On Linux, Disabled/Full means the snapshot is invalid because the stored-on partition is full. You should delete the invalid snapshot. Other snapshots continue working. Disabled/Full snapshots are ignored by the pool and do not impact the I/O performance for the pool.
	Size	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.
	Creation Date	The time stamp that shows the date and time the original pool was snapped.
Snapshot of	Name	The name of the original pool that was snapped.
	Total space	The total space allocated for the original pool.
	Used space	The amount of space currently in use for the original pool.
Stored-On Location (pool or partition)	Name	The pool where the snapshot metadata resides. This can be the same as the original pool.
	State	The current state of the stored-on pool, either Active (default) or Deactive.
	Total space	The total space allocated for the pool where the snapshot is stored.
	Used space	The amount of space currently in use for the pool where the snapshot is located. The used space represents all uses of the pool, not just the snapshot.

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 *Pools* page.
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*.

5 On the *Snapshots* page, select the snapshot, then click *Actions* > *Details* to open the Details dialog box.

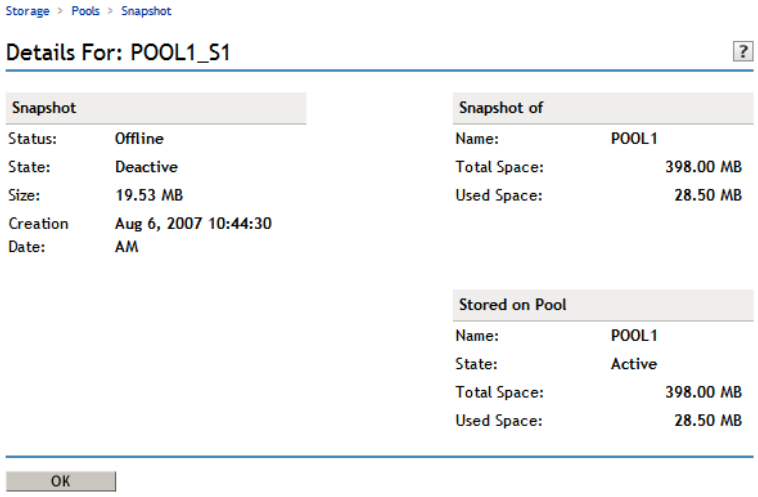

18.6 Renaming a Pool Snapshot (NetWare)

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. Renaming a pool snapshot is not supported on Linux.

1 In iManager, log on to the tree for the server you want to manage.

2 In *Roles and Tasks*, click *Storage* > *Pool*s to open the *Pools* page.

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*.

5 Select the snapshot in the *Snapshots* list.

6 Click *Rename* to open the *Rename Snapshot* page.

For information about naming snapshots, see “Guidelines for Naming Pool Snapshots” on page 233.

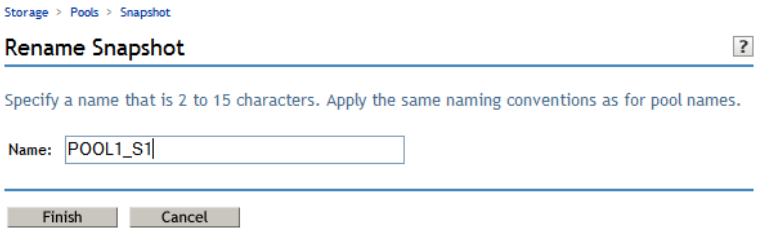

7 Specify the new name you want for the pool snapshot.

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.

18.7 Modifying the Stored-On Location for Snapshots

- Section 18.7.1, “Modifying the Stored-On Partition (Linux),” on page 246
- Section 18.7.2, “Modifying the Stored-On Pool (NetWare),” on page 246

18.7.1 Modifying the Stored-On Partition (Linux)

After you specify a partition to use for a given snapshot, that partition cannot be changed or resized. When you delete the snapshot, the stored-on partition is also deleted.

18.7.2 Modifying the Stored-On Pool (NetWare)

After you specify a stored-on pool to use for a given original pool, that location is the default stored-on pool for all subsequent snapshots made of the original pool. To modify the stored-on pool that is associated with an original pool, you must delete all existing snapshots for the original pool, then begin again to create a new family of pool snapshots. Deleting all snapshots on the stored-on pool does not delete the pool itself.

To resize the stored-on pool, you can add segments to the pool to expand it just as you would for any pool. For information, see Section 16.4, “Increasing the Size of a Pool,” on page 202.

18.8 Onlining or Offlining a Pool Snapshot

You can mount a pool snapshot as a pool in order to make the point-in-time versions of the volumes available.When it is mounted, the pool snapshot appears on the *Pools* page by its snapshot name. The volumes on the pool appear on the *Volumes* page with an `_SV` appended to the name, such as `VOL1_SV`.

For example, you might want to mount a snapshot as a pool to back up data from the snapshot version of the volumes. Snapshot functions occur while the snapshot pool is active, whether the snapshot is mounted or dismounted. Snapshots are typically not mounted for general user access purposes.

- Section 18.8.1, “Using iManager to Online a Pool Snapshot,” on page 246
- Section 18.8.2, “Using iManager to Offline a Pool Snapshot,” on page 247
- Section 18.8.3, “Using NSSMU for Linux to Online or Offline a Pool Snapshot,” on page 247

18.8.1 Using iManager to Online a Pool Snapshot

You can use iManager to online or offline pool snapshots for NSS volumes on Linux or NetWare.

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 *Pools* page.
3. Select the server that contains the original pool of the snapshot you want to manage.
4. In the *Pools* list, select the original pool of the snapshot you want to manage, then click *Snapshot* to open the *Snapshots* page.
5. Select one or more pool snapshots that you want to manage, then click *Actions* > *Online*.

This activates the selected pool snapshots and their volumes. The volumes are not mounted automatically.

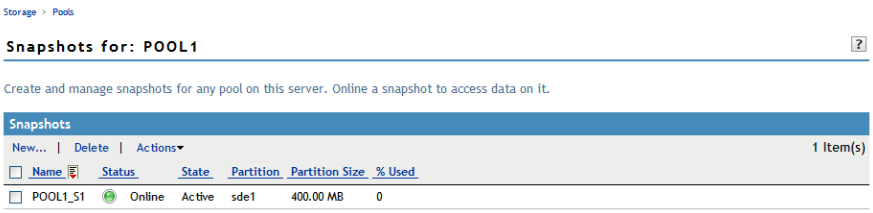

6 If you need to access a pool snapshot volume, mount the volume.

6a In *Roles and Tasks*, click *Storage* > *Volumes* to open the *Volumes* page.

6b On Linux, if you need to verify NSS metadata information for the snapshot volume while it is online, select the snapshot volume, then click *Update eDirectory* to create a Volume object for the volume.

On NetWare, the Volume object is created automatically when you online the pool snapshot.

6c On the *Volumes* page, select the snapshot volume you want to manage, then click *Mount*.

18.8.2 Using iManager to Offline a Pool Snapshot

You can use iManager to offline pool snapshots for NSS volumes on Linux or NetWare.

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 *Pools* page.

3 Select the server that contains the original pool of the snapshot you want to manage.

4 If the snapshot volumes for the pool snapshot are currently mounted, go to the *Volumes* page, select the mounted snapshot volumes, then click *Dismount*.

5 In the *Pools* list, select the original pool of the snapshot you want to manage, then click *Snapshot* to open the *Snapshots* page.

All snapshot volumes are automatically offlined at this time.

6 Select one or more pool snapshots that you want to manage, then click *Actions* > *Offline*.

This makes the selected pool snapshots and all the volumes in them unavailable to users. It does not destroy volumes in the snapshots, nor does it destroy the data contained in the volumes.

18.8.3 Using NSSMU for Linux to Online or Offline a Pool Snapshot

On Linux, you can use NSSMU for Linux to online or offline the pool snapshot.

1 Open a terminal console, then log in as the `root` user.

2 At the terminal console prompt, enter

```
nssmu
```

3 In NSSMU, select *Snapshot*.

4 Select the pool snapshot that you want to manage, the press *F7* to mount or dismount the pool snapshot.

The pool snapshot functions continue even though the snapshot is not mounted for general user access.

18.9 Viewing and Managing an Online Pool Snapshot

When online, pool snapshots appear and function as a pool in the *Pools* list on the *Pools* page.

1 In iManager, click *Storage* > *Pools*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage to view a list of its pools.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Pools* list, select the online pool snapshot.

Wait for the page to refresh. It displays the pool snapshot’s details and enables its management options. The Name is followed by *snapshot* in parens to indicate that the selected pool is a snapshot.

The following is an example pool snapshot on Linux.

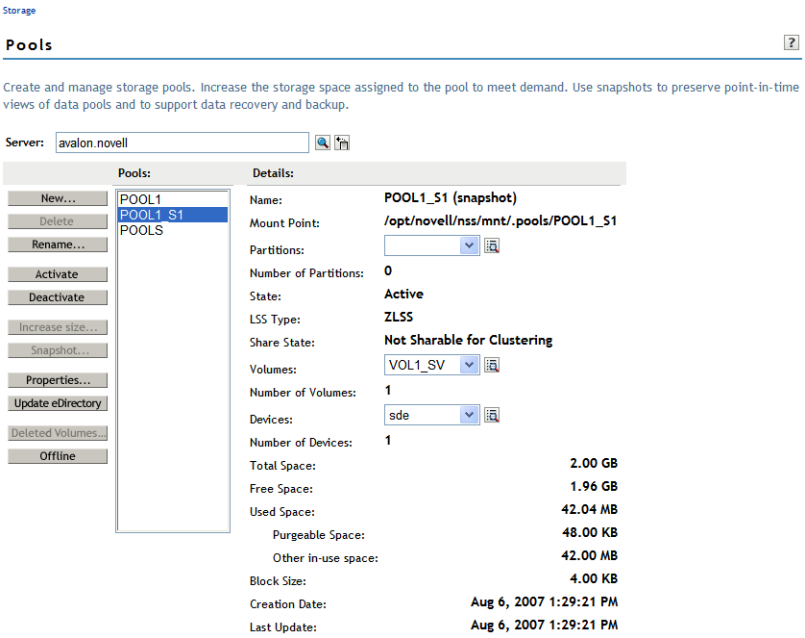

The following is an example pool snapshot on NetWare.

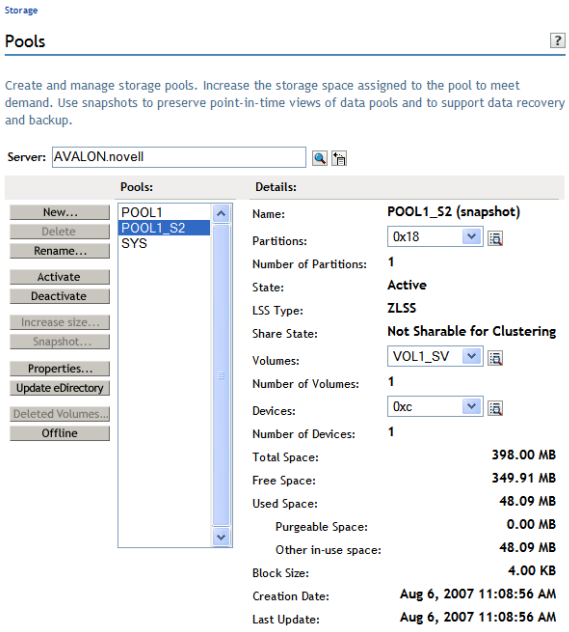

4 Click *Storage* > *Volumes* to go to the volumes page.

The snapshot volumes are listed in the *Volumes* list. They are deactive and unmounted by default.

5 Optionally select the snapshot volume, then click *Mount* to mount the snapshot volume so that you are able to access its data.

The following is an example snapshot volume mounted on Linux.

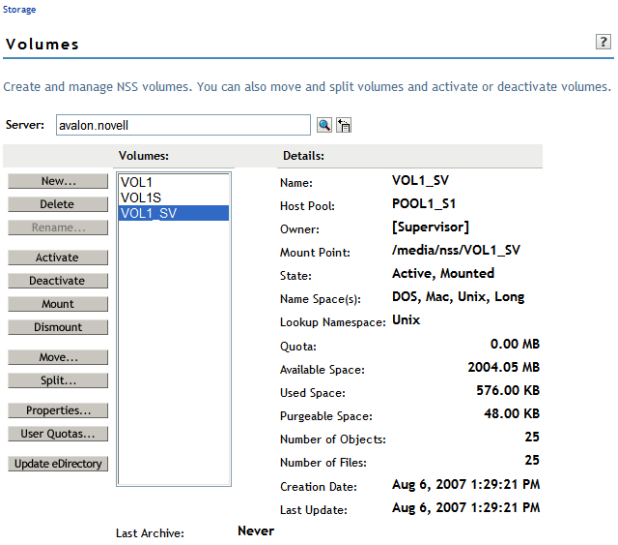

6 When you are done, go to the *Pools* page, select the snapshot pool, then click *Offline* to take the pool snapshot offline and dismount its snapshot volumes.

18.10 Restoring Data from an Online Pool Snapshot

You can restore a point-in-time version of data from a pool snapshot by manually copying the data from an online snapshot volume to another location.

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. Select a server to manage to view a list of its pools.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3. Online the pool snapshot that contains the version of the file you want to restore.

   For information, see Section 18.8.1, “Using iManager to Online a Pool Snapshot,” on page 246.

4. Select *Storage* > *Volumes* to go to the volumes page.

5. Select the snapshot volume (such as `VOL1_SV`) from the *Volumes* list, then click *Mount*.

   Snapshot volumes are mounted Read Only. You cannot modify the content of files on the snapshot.

6. Use any normal method to copy the file of interest from the mounted snapshot volume to a new location.

18.11 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. For guidelines, see “Guidelines for Deleting Pool Snapshots” on page 237.

IMPORTANT: Delete the oldest snapshot first in a first-created, first-deleted manner.

- Section 18.11.1, “Using iManager to Delete a Pool Snapshot,” on page 250
- Section 18.11.2, “Using NSSMU to Delete a Pool Snapshot (Linux),” on page 251

18.11.1 Using iManager to Delete a Pool Snapshot

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. In the server field on the *Pools* page, select a server to manage to view a list of its pools.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3. If the pool that has snapshots you want to manage is not active, select the pool from the *Pools* list, then click *Activate*.

4. If a pool snapshot for the selected pool is currently online, do the following to offline the pool snapshot:

   4a If the snapshot volumes for the pool snapshot are currently mounted, go to the *Volumes* page, select the mounted snapshot volumes, then click *Dismount*.

   4b On the *Pools* page, select the pool snapshot from the *Pools* list, then click *Offline*.

5 On the *Pools* page, select the active pool from the *Pools* list that has snapshots you want to manage, then click *Snapshot* to go to the *Snapshots for: <poolname>* page.

IMPORTANT: If the selected pool is a online pool snapshot, the *Snapshot* option is not available. Select the original pool instead.

6 In the *Snapshots* list, select one or more snapshots that you want to delete.

7 Click *Delete*, then click *Yes* to confirm the delete.

18.11.2 Using NSSMU to Delete a Pool Snapshot (Linux)

For NSS pool snapshots on OES 2 Linux, you can use NSSMU for Linux to delete snapshots.

1 Open a terminal console, then log in as the `root` user.

2 At the terminal console prompt, enter

```
nssmu
```

3 If the pool that has snapshots you want to manage is not active, select *Pools* from the main menu, select the pool, then press F7 to activate the pool.

4 In NSSMU, click *Snapshots*.

5 If the pool snapshot you want to delete is currently online, select the pool snapshot from the *Snapshots* list, the press F7 to dismount the pool.

6 Select the pool snapshot that you want to delete from the *Snapshots* list, then press *Delete* to delete the snapshot.

7 Click *Y (Yes)* to confirm the delete.

18.12 Shredding a Deleted Pool Snapshot (NetWare)

In NetWare 6.5 SP2 and earlier, Media Manager automatically turned on shredding for deleted snapshots. In NetWare 6.5 SP3 and later, NSS 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 the following at the NetWare server console:

```
mm snap shred snapname=on
```

Replace *snapname* 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 a given 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 snapshot, allow more time for larger amounts of data to shred. The typical delay time is a few minutes.

If any pool snapshot was created for a pool prior to the NetWare 6.5 SP3 release, shredding is already turned on. You can turn shredding off by entering the following at the NetWare server console:

```
mm snap shred snapname=off
```

Replace *snapname* 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.

18.13 Using Commands to Manage Pool Snapshots (NetWare)

Command line options are available to manage pool snapshots for your NetWare server. For information about pool snapshot commands for NetWare, see Section A.30, "Pool Snapshot Commands (NetWare)," on page 483.

Command line options are not available for NSS pool snapshots on Linux. You must manage the snapshots with the *Snapshot* page in NSSMU for Linux. For information, see "Snapshots" in Table 9-12 on page 113.

Managing NSS Volumes

19

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 19.1, “Understanding Volume Properties,” on page 253
- Section 19.2, “Guidelines for NSS Volumes,” on page 258
- Section 19.3, “Creating Unencrypted NSS Volumes,” on page 259
- Section 19.4, “Configuring Encrypted NSS Volumes with NSSMU,” on page 262
- Section 19.5, “Updating eDirectory Volume Objects,” on page 262
- Section 19.6, “Viewing the Details of an NSS Volume,” on page 263
- Section 19.7, “Viewing Properties of an NSS Volume,” on page 263
- Section 19.8, “Modifying Attributes of an NSS Volume,” on page 265
- Section 19.9, “Modifying the NSS Volume Size,” on page 266
- Section 19.10, “Configuring the Name Space for an NSS Volume,” on page 267
- Section 19.11, “Mounting NSS Volumes with Linux Commands (Linux),” on page 268
- Section 19.12, “Renaming an NSS Volume,” on page 269
- Section 19.13, “Renaming (Modifying) the Mount Point for an NSS Volume (Linux),” on page 270
- Section 19.14, “Activating and Deactivating an NSS Volume,” on page 271
- Section 19.15, “Mounting and Dismounting an NSS Volume (Linux),” on page 271
- Section 19.16, “Exporting NSS Volumes for NFS Access (NetWare),” on page 272
- Section 19.17, “Exporting and Importing NSS Volumes for NFS Access (Linux),” on page 273
- Section 19.18, “Deleting an NSS Volume,” on page 278
- Section 19.19, “Finding the Filename for a Given ZID,” on page 279
- Section 19.20, “Verifying or Rebuilding NSS Volumes,” on page 279
- Section 19.21, “Moving Volumes with DFS,” on page 279
- Section 19.22, “Splitting Volumes with DFS,” on page 279
- Section 19.23, “What’s Next,” on page 279

19.1 Understanding Volume Properties

- Section 19.1.1, “Volume Attributes,” on page 254
- Section 19.1.2, “Encryption Support,” on page 257
- Section 19.1.3, “Enhanced Hard Link Support,” on page 258

19.1.1 Volume Attributes

Figure 19-1 shows the volume attributes for an NSS volume that can be set in iManager when you create volumes. An explanation of each attribute is provided below.

Figure 19-1 *Volume Attributes*

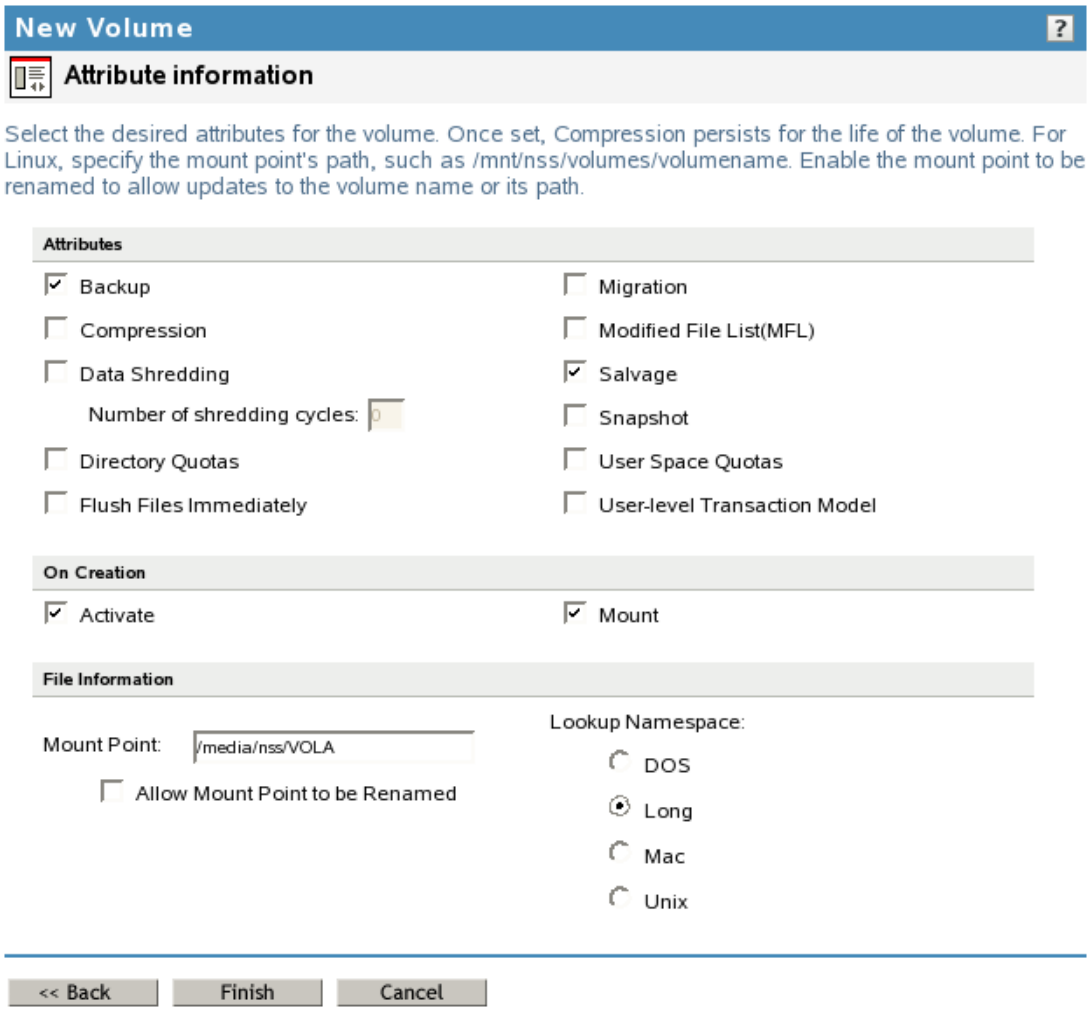

- "Backup" on page 255
- "Compression" on page 255
- "Data Shredding" on page 255
- "Directory Quotas" on page 255
- "Flush Files Immediately" on page 255
- "Migration (to Third-Party Near-Line or Offline Storage)" on page 255
- "Modified File List (NetWare)" on page 255
- "Salvage Files" on page 256
- "Snapshot (File-Level)" on page 256
- "User Space Quotas" on page 256

- "User-Level Transaction Mode (NetWare)" on page 257
- "Lookup Namespace" on page 257

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. The Backup attribute is enabled by default.

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 325.

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 22.3, "Using Data Shredding to Prevent Access to Purged Files," on page 306.

Directory Quotas

The Directory Quotas attribute enables you to assign a maximum quota of space that a directory can consume. For information, see "Managing Space Quotas for Volumes, Directories, and Users" on page 345.

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 23.1, "Enabling Flush Files Immediately to Write Data to the Disk on Close," on page 315.

IMPORTANT: On Linux, a group write function controls how writes to disk occur. For information, see Section 31.3, "Configuring or Tuning Group I/O," on page 415.

Migration (to Third-Party Near-Line or Offline Storage)

The 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.

Modified File List (NetWare)

The Modified File List (MFL) attribute enables NSS on NetWare to create a list of all files modified since the previous backup. The log is available only through third-party software.

NOTE: This feature is seldom-used since the introduction of the Event File List support in OES 1 Support Pack 1 for Linux and NetWare (same as NetWare 6.5 SP4). Consider using the Event File List instead, which is available for NSS on Linux and NetWare. For information, see "FileEvents.xml Definitions" (http://developer.novell.com/documentation/vfs/vfs__enu/data/ak7gh2x.html) in *NDK: Virtual File Services* (http://developer.novell.com/documentation/vfs/vfs__enu/data/bktitle.html).

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. Salvage is enabled by default.

If the Salvage Files attribute is disabled, deleted files are purged immediately on deletion.

IMPORTANT: The Salvage Files attribute does not affect whether deleted volumes can be salvaged or purged. Salvage for deleted volumes is determined at the server level with the `nss /ImmediatePurgeOfDeletedFiles=<on | off>` setting. For more information, see Section 26.2.2, "Setting the Immediate Purge of Deleted Files for All NSS Volumes," on page 365.

For information, see Chapter 26, "Salvaging and Purging Deleted Volumes, Directories, and Files," on page 361.

Snapshot (File-Level)

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.

IMPORTANT: Not all third-party backup software can take advantage of the file snapshot attribute, even if you set it.

For more information, see Section 23.3, "Using the File-Level Snapshot Attribute to Enable the Backup of Open Files," on page 318.

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 "Managing Space Quotas for Volumes, Directories, and Users" on page 345.

User-Level Transaction Mode (NetWare)

The User-Level Transaction mode enables the Transaction Tracking System™ (TTS™) function for NSS volumes on Netware. TTS logs changes made to a file contents, and protects database applications by backing out transactions that are incomplete because of a system failure. For information, see Section 23.2, "Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare)," on page 317.

IMPORTANT: This capability is available only on NetWare. If you enable this attribute for an NSS volume on Linux, it is ignored.

When working on OES 2 Linux, if your data requires content-level transaction tracking and you want to continue to control access to data with trustees, you can use NCP volumes on Linux Reiser or Ext3 file systems. Configure the journaling mode for the file system to the Journaling level, which provides content-level transaction tracking. For information about creating and managing NCP volumes on Linux, see "Managing NCP Volumes" in the *OES 2 SP1: NCP Server for Linux Administration Guide*.

Lookup Namespace

NSS provides multiple name spaces for the volume: Long, UNIX, DOS, and Macintosh. The Lookup Namespace attribute sets the primary name space to use when you mount the volume, but all name spaces are available for use by various applications.

For NSS volumes, the Long name space is highly recommended because names on NSS are case insensitive by default. The UNIX name space supports case-sensitive naming.

For NetWare, Long is the default setting for the primary name space when mounting an NSS volume. For OES 2 SP1 and later, Long is also the default. Using the Long name space as primary improves performance over using the UNIX name space, especially if you expect to store millions of files on the volume.

NOTE: In OES 2 Linux and earlier, UNIX was the default name space for mounting NSS volumes on Linux.

For Linux, NCP tools can be used only when Long or UNIX is set as the primary name space. With DOS or Mac set as the primary name space, you cannot view or manage the volume from Novell Remote Manager, and users are unable to map to the volume using NCP clients. If you use the Long or UNIX name space, the DOS and Mac name spaces are still available, but they are not the primary.

The UNIX name space supports some special characters that are not allowed in the Long name space, such as characters 0x01 through 0x07 and 0x10 through 0x1f. If you need to use these special characters in filenames, choose UNIX as the default name space.

19.1.2 Encryption Support

Encryption provides 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. Encrypted volume support is provided for NSS volumes in NetWare 6.5 Support Pack 2 and later. It is available on OES 1 SP1 Linux and later.

Encrypted volumes can be created only from NSSMU. Encrypted volumes require special handling on the first activation after startup, but all attributes are available for encrypted volumes and are managed the same as for unencrypted volumes. For information about creating and activating encrypted volumes, see “Managing Encrypted NSS Volumes” on page 281.

19.1.3 Enhanced Hard Link Support

Enhanced 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. NSS supports zero to 65,535 hard links per file on NSS volumes on NetWare or Linux.

You must upgrade the media format before you can set the Hard Links attribute and create new hard links on a volume. For information about upgrading the media format, see Chapter 4, “Upgrading the NSS Media Format,” on page 55.

After the media upgrade for enhanced hard links support, the Hard Links attribute must be enabled or disabled by using commands. The attribute cannot be enabled or disabled in NSSMU or in iManager. For information about enabling Hard Link support for a volume, see Section 27.3, “Enabling or Disabling the Hard Links Attribute,” on page 380.

After the volume has been enabled for enhanced hard links, you can create hard links. For information about creating and managing hard links, see Chapter 27, “Managing Hard Links,” on page 375.

Beginning in OES 2 SP1, Novell Storage Management Services™ supports the backup and restore of hard links on NSS volumes.

Hard links are lost when you use the Move Volume or Split Volume features of Distributed File Services.

19.2 Guidelines for NSS Volumes

- Section 19.2.1, “Guidelines for Sizing Volumes,” on page 258
- Section 19.2.2, “Guidelines for Name Spaces,” on page 259
- Section 19.2.3, “Guidelines for NSS Volumes in a Mixed-Node Cluster,” on page 259

19.2.1 Guidelines for Sizing 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 quota as the maximum size, 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.

If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool’s maximum size by overbooking, 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.

Because volume sizes can be overbooked in a pool, NSS automatically considers what space is remaining in a pool in order to report the maximum size that is currently possible for the volume. In addition to other volumes that can consume space in a pool, NSS snapshots and third-party snapshots can consume space that might not be reported in all of the management tools that report space. NSS reports the total space possible for the volume and the amount of space used by the volume so that tools can properly calculate the maximum free space available.

19.2.2 Guidelines for Name Spaces

NSS recognizes DOS, Macintosh, UNIX, and Long name spaces. Volume names, directory names, and filenames in NSS are case insensitive. This differs from Linux POSIX file systems, which are case sensitive by default. For information, see Section , “Lookup Namespace,” on page 257.

19.2.3 Guidelines for NSS Volumes in a Mixed-Node Cluster

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 about using NSS volumes cross-platform, see the following:

- Section 7.2, “Cross-Platform Issues for NSS Volumes,” on page 87
- Section 7.3, “Cross-Platform Issues for NSS Features,” on page 88
- Section 7.4, “Cross-Platform Issues for File Access,” on page 89

For information about clustering, see the following:

- *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*
- *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*

19.3 Creating Unencrypted NSS Volumes

This section describes how to create an unencrypted volume with iManager. Encrypted volumes can be created only in NSSMU. For information on creating encrypted volumes, see “Managing Encrypted NSS Volumes” on page 281.

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

   A list of volumes appears in the *Volumes* list, as illustrated in the following figure.

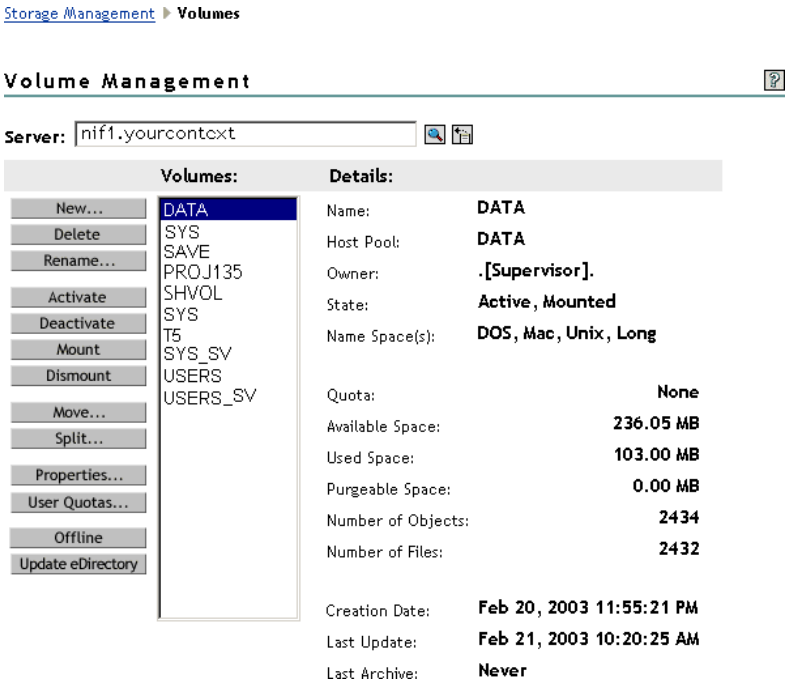

3 To create a new volume, click *New*.

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 5.4, “Naming NSS Storage Objects,” on page 67.

5 Do one of the following:

- Select an existing pool from the list where you want the new volume to reside.
- If no pools exist, click *New Pool*, create a pool to use, select the pool.
- If existing pools do not have sufficient space for the volume you want to create, click *Cancel* to close the Wizard. You must add more segments of free space to the pool, then return to the *Volumes* page to create the new volume.
- If no pools exist and no space is available to create one, click *Cancel* to close the Wizard. You must add more devices to the server or free up space on existing pools, then return to the Volumes page to create the new volume.

6 Specify the size of the volume:

- **No Volume Quota:** Select *Allow Volume Quota to Grow to the Pool Size* if you want the volume to expand to the size of the pool. This is the default.

  Pools can be overbooked; each volume can potentially grow to the size of the pool. NSS allocates space as it is needed.
- **Volume Quota:** Deselect *Allow Volume Quota to Grow to the Pool Size*, then type a *Volume Quota* size in MB for the volume if you want to limit the size of the volume.

7 On the *Attribute Information* page under the *Attributes* section, set the attributes for the new volume you are creating. The Backup and Salvage attributes are selected by default.

For information about volume attributes, see Section 19.1, “Understanding Volume Properties,” on page 253.

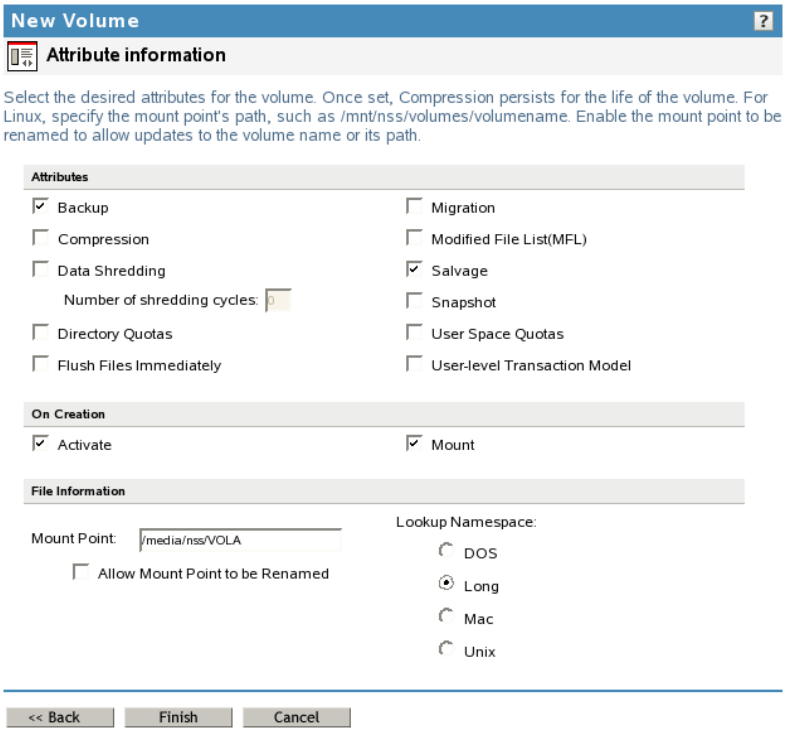

8 On the *Attribute Information* page under the *On Creation* section, set the following preferences:

- *Activate*. Activates logical volumes as soon as you create them.
- *Mount*. Mounts logical volumes as soon as you create them.

9 On the *Attribute Information* page under *File Information*, specify the following parameters:

- **Mount Point (Linux):** For a Linux server, specify the mount point for the NSS volume, such as `/media/nss/VOLA`.

  The default mount path for NSS volumes on Linux is `/media/nss/volumename`, where `volumename` is the name of the volume. You can optionally specify another path as the mount point.

- **Allow Mount Point to Be Renamed (Linux):** For a Linux server, select this option if you want to allow the mount point to be renamed after it has been created.

  Renaming a mount point means that you can specify another path as the mount point, such as `/media/uservols/volumename`, but the volume would continue to be associated with the same Volume object in eDirectory.

- **Lookup Name Space:** Select the name space to use when you mount the volume. The name spaces are UNIX, Long, DOS, or Macintosh. The default name space is Long.

  The recommended setting is Long. This setting ensures that filenames are case insensitive whether the volume is mounted on a Linux server or NetWare server. It also improves performance over using UNIX, especially if you expect to store millions of files on the volume.

10 Click *Finish*.

11 (Linux) If you enabled the *Directory Quotas* attribute, restart NCP2NSS by entering at a terminal prompt:

```
/etc/init.d/ncp2nss restart
```

19.4 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 281.

19.5 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 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 Volumes page allows you to add or replace a Volume object for a selected volume at the same context level as the server.

IMPORTANT: When you delete (or delete and replace) a Volume object in eDirectory, the home directory attribute is removed in the User objects for any users that reference that Volume. The home directory attribute points to a particular Volume object. When that Volume object is deleted, eDirectory needs to clean up all references to the object being deleted.

When you select Update eDirectory, NSS searches for the object.

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Volumes* list, 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.

19.6 Viewing the Details of an NSS Volume

1 In iManager, click *Storage* > *Volumes*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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) Activate the volume, select the volume, then click *Activate*.

19.7 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Volumes* list, select a volume that you want manage.

4 Click *Properties*.

The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. It opens by default to the *Attributes* tab.

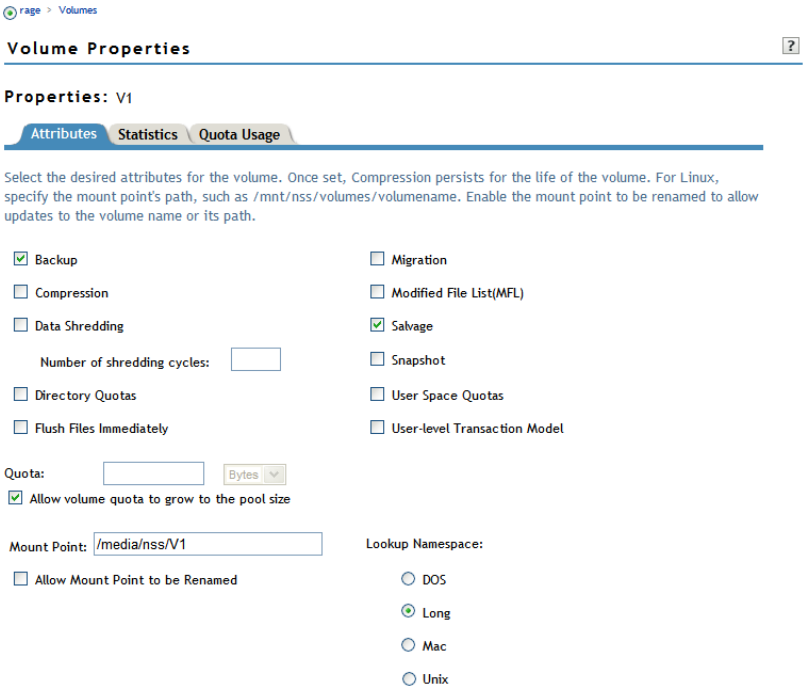

Use the *Attributes* page to view the volume's attribute configuration, the volume quota, and the volume mount point (Linux). For information about modifying attributes, see Section 19.8, "Modifying Attributes of an NSS Volume," on page 265.

5 Click the *Statistics* tab to view the current space usage statistics for the selected volume.

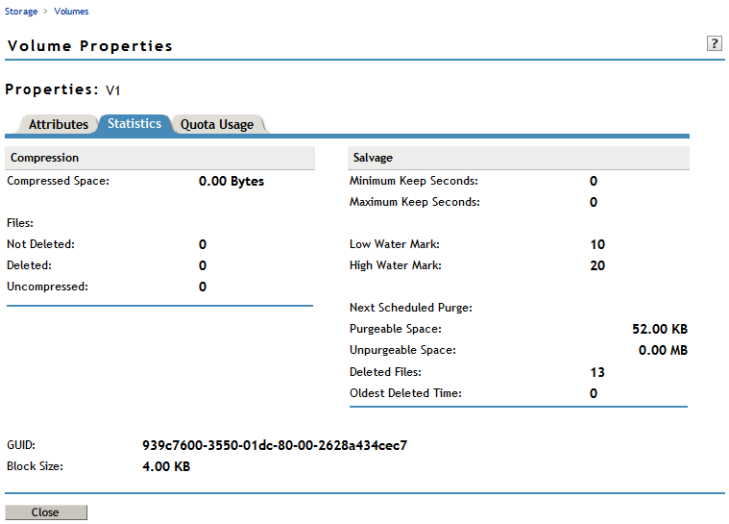

If the Salvage attribute is enabled, values are displayed for the salvage parameters. The low and high watermark displays the default settings for the pool-level watermarks for the pool where the volume resides. For information about managing salvage parameters, see Chapter 26, "Salvaging and Purging Deleted Volumes, Directories, and Files," on page 361.

If the Compression attribute is enabled, statistics are displayed for the compression data. For information about configuring compression parameters, see Chapter 24, “Managing Compression on NSS Volumes,” on page 325.

6 Click the *Quota Usage* tab to view the volume and pool space usage for the selected volume.

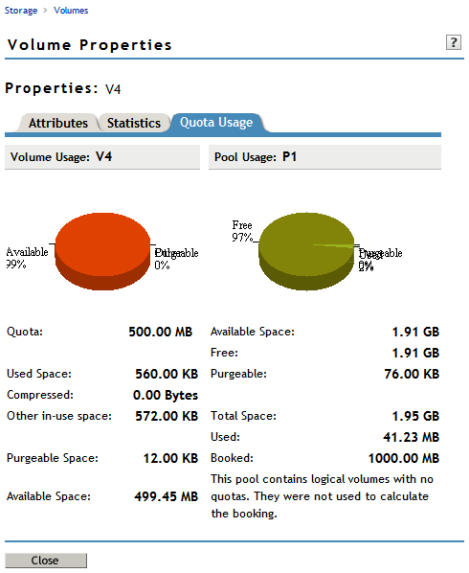

19.8 Modifying Attributes of an NSS Volume

After you set up and configure NSS volumes, you can modify most of the attribute settings. The Encrypted Volume Support attribute and the 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.

1 In iManager, click *Storage* > *Volumes*.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Volumes* list, select a volume that you want manage.

4 Click *Properties*.

The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. Use the *Attributes* page to view or modify the attributes for the selected volume.

5 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 19.1, “Understanding Volume Properties,” on page 253.

- Specify a volume quota, then click *Apply*.

- Specify the default *Lookup Namespace* to use when mounting the volume, then click *Apply*.

  The next time the volume is mounted, this will be the name space used. The default name space is Long.
- (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 the Directory Quotas attribute or User Space Quotas attribute, see "Managing Space Quotas for Volumes, Directories, and Users" on page 345.

19.9 Modifying the NSS Volume Size

1 In iManager, click *Storage* > *Volumes*.

  For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

  For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 In the *Volumes* list, select a volume that you want manage.

4 Click *Properties*.

  The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. It opens to the *Attributes* tab.

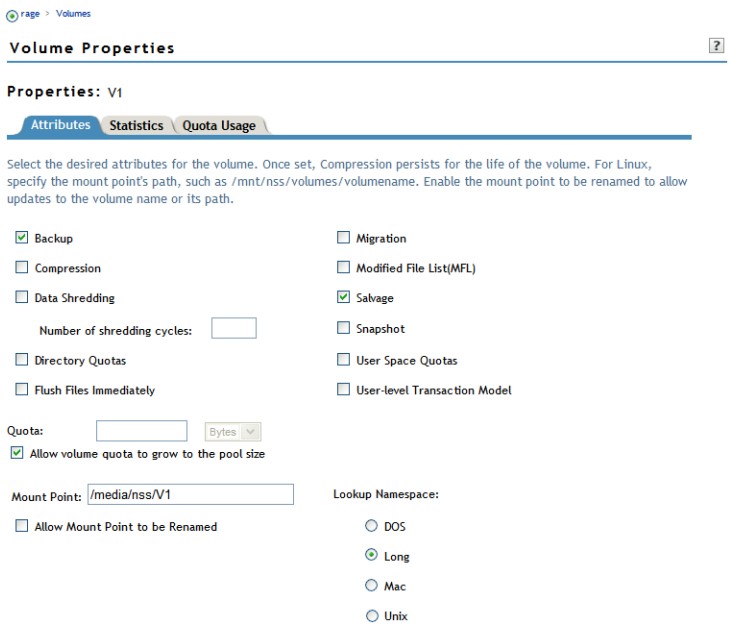

5 In the *Quota* field, do one of the following:

- **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.

  If you set a quota that is less than the volume's current size, no files can be saved to the file until you purge files to make room on the volume.

6 Click *Apply*.

7 Click the *Quota Usage* tab to view the volume and pool space usage for the selected volume and to verify the new setting.

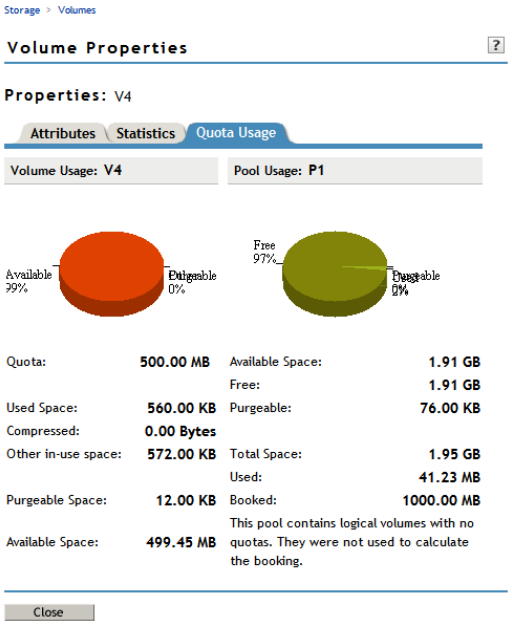

19.10 Configuring the Name Space for an NSS Volume

NSS supports the Long, DOS, UNIX, and Macintosh name spaces. By default, names on the NSS file system are case insensitive, which is supported by the Long name space. The Long name space is the default setting used when mounting NSS volumes on Linux or NetWare. In order to mount a volume with a different name space, you must specify the name space explicitly in the mount command, or you can specify the name space to use as a property of the NSS volume.

On Linux, the UNIX name space is typically used for Linux POSIX file systems. If your volume contains large directories with millions of files, using the default UNIX name space on NSS volumes can cause volumes to mount very slowly. Using the Long name space allows the NSS volume to mount normally. Unless you need to support case sensitive filenames, we strongly recommend using the Long name space.

The preferred name space can be set when you create the volume and set its attributes, or at any time by modifying the Lookup Namespace attribute in the volume's properties.

For OES Linux, you can also specify mount the volume by specifying the name space to use as an option of the mount command. For instructions, see Section 19.11, “Mounting NSS Volumes with Linux Commands (Linux),” on page 268.

To view or modify the Lookup Namespace attribute for the NSS volume:

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage to view a list of its volumes.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. In the *Volumes* list, select a volume that you want manage.
4. Click *Properties* to view the volume’s properties.

   The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*.
5. On the *Attributes* page, view the current setting of the *Lookup Namespace* for the selected volume.
6. On the *Attributes* page, optionally modify the *Lookup Namespace* to use by selecting the radio button next to it.
   - *Long* (recommended, default)
   - *DOS*
   - *UNIX*
   - *Mac*

   This is the new value that is applied automatically whenever you mount the volume.
7. Click *Apply* to save your changes.
8. On the *Volumes* page, click *Dismount* to unmount the volume.

   Wait until the volume unmounts gracefully before continuing.
9. On the Volumes page, click *Mount* to mount the volume to mount it the new name space.

19.11 Mounting NSS Volumes with Linux Commands (Linux)

When mounting an NSS volume on Linux, specify the Long name space to make its directory names and filenames case insensitive. Long is the default name space, and improves performance over using the UNIX name space. If your volume contains large directories with millions of files, using the UNIX name space can cause volumes to mount very slowly and can degrade performance.

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 primary name space type. Valid name space options 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.

19.12 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
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.

19.13 Renaming (Modifying) the Mount Point for an NSS Volume (Linux)

The default mount point for NSS volumes is `/media/nss/volumename`. You must enable the *Allow the Mount Point to Be Renamed* option for the volume to allow the mount point to be renamed

- Section 19.13.1, "Renaming the Mount Point for a New Volume," on page 270
- Section 19.13.2, "Enabling the Mount Point for the NSS Volume to Be Renamed," on page 270
- Section 19.13.3, "Renaming the Mount Point for an Existing NSS Volume," on page 270

19.13.1 Renaming the Mount Point for a New Volume

The default mount point for NSS volumes is `/media/nss/volumename`. You can specify a different mount point (that is, modify the volume's directory path) as you create the volume if you create the volume in iManager. Creating the volume in the NSSMU (`nssmu`) does not allow for the mount point to be changed during the volume setup, but you can change it afterwards.

19.13.2 Enabling the Mount Point for the NSS Volume to Be Renamed

The *Allow the Mount Point to Be Renamed* option enables the NSS volume's mount point to be renamed. This option is disabled by default. Enable the option as you create the volume, or enable it at any time for an existing volume by modifying the setting on the *Attributes* page (*Storage* > *Volumes* > *Properties* > *Attributes*) in iManager.

1. In iManager, click *Storage*, then click *Volumes*.
2. Select the server you want to manage to view a list of its volumes.
3. From the *Volumes* list, select the volume, then click *Properties* to view the volume attributes.
4. On the *Attributes* page, select *Allow the Mount Point to Be Renamed*.
5. Click *Apply* to save the change.

19.13.3 Renaming the Mount Point for an Existing NSS Volume

Whenever you change the mount point for an existing NSS volume on Linux, you must also restart Novell eDirectory to update the NetWare Core Protocol™ (NCP™) Server cache. When an NSS volume is created, the NCP Server gets the path to the volume and caches it, assuming that it never changes. When you later run `ncpcon` and enter the `volume` command, it reports which volumes are still found at their respective mount points. Only the volumes that are still valid as compared to the list in cache are reported. Restarting eDirectory forces the NCP Server volume cache to update, so that the correct path is stored for reporting volume status.

1. Use either iManager or `nssmu` to change the volume's mount point.

   The following instructions are for iManager.

   1a. In iManager, click *Storage* > *Volumes*.

   1b. Select the server you want to manage to view a list of its volumes.

   1c. From the *Volumes* list, select the volume, then click *Properties* to view the volume attributes.

1d If the *Allow the Mount Point to Be Renamed* option is not selected, select it now, then click *Apply*.

1e In *Mount Point*, type the new mount point.

The default mount point for NSS volumes is `/media/nss/volumename`. The new path should also include the *volumename*.

```
/mnt/nss/volumes/volumename
```

1f Click *Apply* to save the change.

2 Open a terminal console on the server, then log in as the `root` user or equivalent.

3 Restart eDirectory by entering

```
/etc/init.d/ndsd restart
```

Restarting eDirectory causes the NCP Server's volume cache to be updated.

19.14 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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 In the *Volumes* list, select one or more volumes that you want to make active or deactive.

4 Click *Activate* or *Deactivate*.

After the page refreshes, you can see that each volume's state matches the state you specified. If a selected volume is already in the specified state, no change occurs. The details of deactive volumes are not available.

19.15 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. In the *Volumes* list, select one or more volumes that you want to mount or dismount.
4. Click *Mount* or *Dismount*.

   After the page refreshes, you can see that the volume’s state changed. If a selected volume is already in the specified state, no change occurs. The details of dismounted volumes are not available.

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. 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. Press *F7* to dismount the mounted volume.

19.16 Exporting NSS Volumes for NFS Access (NetWare)

NSS on NetWare supports NFSv2 and NFSv3 as specified in “Working with UNIX Machines” in the *OES 2 SP1: AFP, CIFS, and NFS for NetWare (NFAP) Administration Guide*. See the referenced section for information on how to configure NIS and NFS for your NetWare server, the NFS mode options, and how to export NSS volumes on NetWare for access by NFS users.

19.17 Exporting and Importing NSS Volumes for NFS Access (Linux)

NSS volumes on OES 2 Linux and their directories are NFSv3 exportable and accessible from remote systems. NFSv4 is not supported for NSS, but exports for other file systems using NFSv4 can coexist with NSS exports using NFSv3.

- Section 19.17.1, "Understanding NFS Export and Mount Options," on page 273
- Section 19.17.2, "Exporting NSS Volumes on Linux for NFSv3," on page 275
- Section 19.17.3, "Importing NSS Volumes," on page 277

19.17.1 Understanding NFS Export and Mount Options

- "Host Options" on page 273
- "Mount Options for Export via NFSv3" on page 273
- "Mount Options for Import via NFSv3" on page 274
- "Additional Information" on page 275

Host Options

The following table describes options for specifying which servers on the network can import the NFS volume. For more information, see the `exports(5)` man page.

Table 19-1 *Host Options for NFSv3 Export of NSS Volumes on Linux*

Mount Option	Description
Single host	Specify a single host by its fully qualified domain name or its IP address.
Netgroups	Specify NIS netgroups as `@groupname`, such as `@trusted`.
Wildcards	Specify a asterisk (*) to specify all hosts. Use the wildcard characters of asterisk (*) and question marks (?) in server names to match multiple servers. For example, `proj*.example.com` matches all hosts in the domain `example.com` that begin with `proj`.
IP networks	Specify all hosts on a network or subnetwork by specifying the IP address and netmask pair as *address/netmask*. For example: 10.10.10.1/255.255.252.0.

Mount Options for Export via NFSv3

Table 19-2 describes mount options available for mounting NSS volumes for export via NFSv3. For more information, see the `exports(5)` man page and the `mount(8)` man page.

Table 19-2 *Mount Options for NFSv3 Export of NSS Volumes on Linux*

Mount Option	Description
rw	Mount the NSS file system with Read/Write (`rw`) access.

Mount Option	Description
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=value	Importing with the `fsid` option works around the fact that there is no device associated with a logical volume. You must import the NSS volume or directory with the FSID option set on it for export: `fsid=`*n* Replace *n* with an integer value greater than 0. The numbers do not need to be sequential. For example, `fsid=1` and `fsid=10`. Make sure to use a unique `fsid` number for each NSS volume or directory you are exporting. **IMPORTANT:** FSID=0 is reserved for NFSv4 as the pseudo root of the exported file system for exported volumes on the Linux server.

Mount Options for Import via NFSv3

Table 19-3 describes mount options available for mounting NSS volumes for import via NFSv3. For more information, see the `mount(8)` man page.

Table 19-3 *Mount Options for NFSv3 Import of NSS Volumes on Linux*

Mount Option	Description
rw	Mount the NSS file system with Read/Write (`rw`) access.
sync	Specify the Sync (`sync`) option, which requires all file system writes to be committed to disk before the request can be completed.
noatime	NSS also supports the optional use of the `noatime` for importing and mounting NSS volumes on Linux using NFS. The noatime option disables the updating of the access time for files so that reading a file does not update its inode access time (`atime`). For more information, see Section A.26, “noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux),” on page 479.

Additional Information

When you use NFS to export or import NSS volumes on Linux, other supporting services are needed, including DNS, NIS, and NFS. For information about configuring and managing these services on Linux, see the following sections in the *SLES 10 SP2 Installation and Administration Guide*:

- "The Domain Name System"
- "Using NIS"
- "Sharing File Systems with NFS"

19.17.2 Exporting NSS Volumes on Linux for NFSv3

1. In a terminal console, log in as the `root` user.
2. In YaST, select *Network Services*, then select *NFS Server* to open the *NFS Server Configuration* page.
3. If NFS Server is not started and enabled, you must configure the NFS Server.

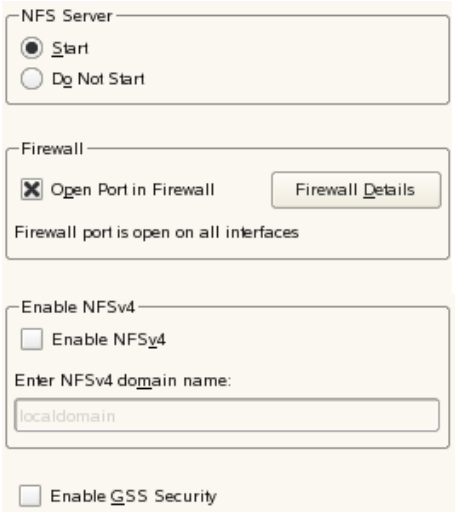

   3a. **NFS Server:** Select *Start*.

   3b. **Firewall:** Select *Open Port in Firewall* to allow access to the NFS service from remote computers, then click *Firewall Details* to specify the network interfaces where you want to open the port.

   3c. **Enable NFSv4:** Make sure that *Enable NFSv4* is not selected if you are exporting only via NFSv3.

   IMPORTANT: NFSv4 is not supported for NSS, but exports for other file systems using NFSv4 can coexist with NSS exports using NFSv3. If you enable NFSv4, make sure that you enter the NSS directories for export with options that use non-zero settings for their FSIDs, and do not bind the NSS directories to paths in the pseudo-root file system that you set up for NFSv4 exports.

3d **Enable GSS Security:** To enable Kerberos secure access to the server, click *Enable GSS Security*. A prerequisite for this is to have Kerberos installed in your domain and both the server and the clients are kerberized.

3e Click *Next* to continue to the *Directories to Export* page.

4 On the *Directories to Export* page, do the following for each NSS volume on the server that you want to export via NFSv3.

4a Under *Directories*, click *Add Directory*, to open a dialog box where you can configure the settings for a volume.

4b Specify the NSS volume that you want to export by typing the Linux path for the volume, or browse the Linux file system to locate and select the NSS volume, then click *OK*.

The default location of NSS volumes on Linux is `/media/nss/volumename`, such as `/media/nss/VOL1`.

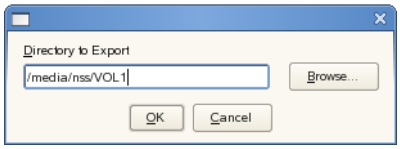

4c In the Host Wildcard field, specify the servers where you want to be able to mount the NSS volume via NFSv3.

A default asterisk (*) wildcard indicates all servers. You can specify a single host, netgroups, wildcards, or IP networks. For information, see Section , “Host Options,” on page 273.

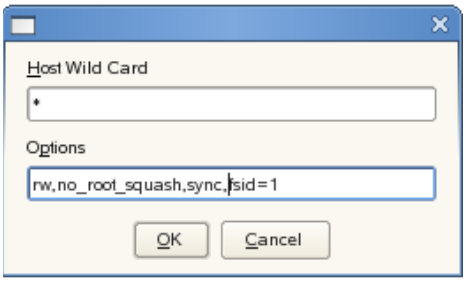

4d Enter the following required mount options:

```
rw,no_root_squash,sync,fsid=value
```

For NFSv3, make sure you do not include the fsid=0 or bind=`/pseudo_rootdir/volumename` options. Not using these options allows the export to be processed as an NSFv3 export.

For example:

```
rw,no_root_squash,sync,fsid=1
```

Separate the options with commands and no spaces. For information, see Section , “Mount Options for Export via NFSv3,” on page 273.

4e Click *OK* to save your settings and return to the *Directories to Export* page.

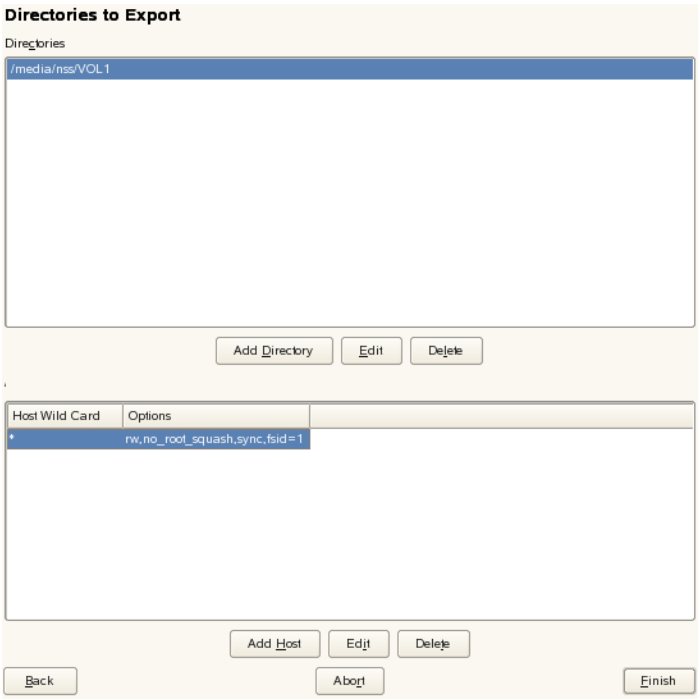

5 On the *Directories to Export* page, click *Finish* to apply the settings.

19.17.3 Importing NSS Volumes

1 On the OES 2 Linux server where you want to import the NSS volume via NFS, open YaST.

2 In YaST, select *Network Services*, then select *NFS Client* to open the *NFS Client Configuration* page.

3 Near the bottom of the page, select *Open Port in Firewall* to allow access to the NFS service from remote computers, then click *Firewall Details* to specify the network interfaces where you want to open the port

4 Do the following for each NSS volume on remote servers that you want to import via NFS.

4a Click *Add* to open a dialog box where you can specify the information for the volume you want to import.

4b In *NFS Server Hostname*, specify the remote server where the volume resides. Type the fully distinguished name (such as *servername.ou_context.o_context.com*), or click *Choose*, select the NFS server from a list of servers, then click *OK*.

4c In *Remote File System*, specify the path on the remote server where the volume resides. Type the full path such as `/media/nss/VOL1`, or click *Select* to open the Exported Directories dialog box, then select the path from the list of NFS exported directories on the selected server, then click *OK*.

4d In *Mount Point (local)*, specify the path on the server (the NFS Client location) where you want to mount the remote volume, such as `/mnt/nfs/volumename`, or click *Browse* to locate and select the location.

The *Browse* option allows you to create a new folder on the server for the target path.

4e Enter the following required mount options:

`rw,sync`

You can optionally specify the `noatime` option. For information, see Table 19-3 on page 274 and the `mount(8)` man page.

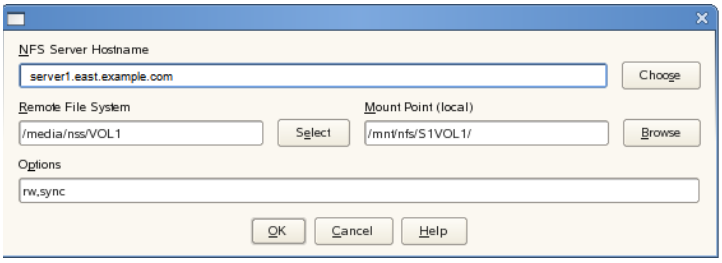

4f Click *OK* to save your settings and return to the *NFS Client Configuration* page.

The entry you just made should appear in the list.

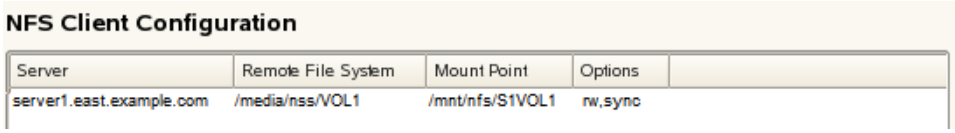

4g When you are done adding volumes to be imported, continue with the next step.

5 On the *NFS Client Configuration* page, click *Finish* to apply the settings.

19.18 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 is salvageable until one of the following events occurs:

- Volume Purge Delay times out. The deleted volume is purged automatically. For information, see Section 26.2.1, “Setting the Purge Delay for All Deleted Volumes,” on page 364.
- You manually purge the deleted volume. For information, see Section 26.4, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 368.

During the purge delay time, the deleted volume is salvageable, but the space belonging to the deleted volume is not available to other volumes and. When 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 26.4, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 368.

You cannot use iManager to delete a `sys:` volume. 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

3 In the *Volumes* list, 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.

19.19 Finding the Filename for a Given ZID

You might get an error report that identifies the ZID for a file, but not the filename and path on the volume. To find the associated filename and full path for the file on a given volume and name space, use the /ZIDtoFilename option.

1 At the NetWare server console, or in the NSS Console (`nsscon`) on Linux, enter the following command (all on the same line, of course):

   ```
   nss /ZIDtoFilename=ZIDnumber /ZIDNameSpace=namespace /ZIDVolumeName=volumename
   ```

   Replace *ZIDnumber* with the ZID of the file. Replace *namespace* with the Long, UNIX, Macintosh, or DOS name space to use for the search. Replace *volumename* with the name of the volume for the search.

19.20 Verifying or Rebuilding NSS Volumes

You cannot rebuild or verify an NSS volume independently of other volumes in the same pool. For guidelines and procedures for verifying and rebuilding NSS pools and volumes, see Chapter 17, “Verifying and Rebuilding NSS Pools and Volumes,” on page 215.

19.21 Moving Volumes with DFS

The Move Volume function uses Novell Distributed File Services to move a volume’s file structure, data, and the file system trustee rights information from the original location to a new volume in the network. For guidelines and procedures for moving volumes, see “Using DFS to Move NSS Volumes” in the *OES 2: Novell Distributed File Services Administration Guide*.

19.22 Splitting Volumes with DFS

The Split Volume function uses Novell Distributed File Services to move a specified part of a volume’s file structure, data, and the file system trustee rights information from the original location to a new volume in the network. A DFS junction replaces the selected directory and its contents in the source volume. The data and metadata in the directory are moved to the target location, which can be the root directory or other directory in the destination volume. For guidelines and procedures for splitting volumes, see “Using DFS to Split NSS Volumes” in the *OES 2: Novell Distributed File Services Administration Guide*.

19.23 What’s Next

For information about advanced volume features, see the following:

- “Managing Encrypted NSS Volumes” on page 281
- “Securing Access to NSS Volumes, Directories, and Files” on page 297
- “Managing Data Integrity on NetWare Servers” on page 315

- “Managing Compression on NSS Volumes” on page 325
- “Managing Space Quotas for Volumes, Directories, and Users” on page 345
- “Salvaging and Purging Deleted Volumes, Directories, and Files” on page 361
- “Managing Hard Links” on page 375
- *OES 2: Novell Distributed File Services Administration Guide*
- *OES 2 SP1: Dynamic Storage Technology Administration Guide*

Managing Encrypted NSS Volumes

20

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 describes the following:

- Section 20.1, "Understanding Encrypted Volume Support," on page 281
- Section 20.2, "Security Considerations for Encrypted Volumes," on page 283
- Section 20.3, "Creating an Encrypted Volume," on page 284
- Section 20.4, "Mounting an Encrypted NSS Volume with NSSMU (Linux)," on page 285
- Section 20.5, "Activating an Encrypted NSS Volume with NSSMU (NetWare)," on page 286
- Section 20.6, "Mounting Encrypted NSS Volumes with NSS Commands (Linux)," on page 286
- Section 20.7, "Activating Encrypted NSS Volumes with NSS Commands (NetWare)," on page 287
- Section 20.8, "Dismounting an Encrypted NSS Volume with NSSMU (Linux)," on page 288
- Section 20.9, "Deactivating an Encrypted NSS Volume with NSSMU (NetWare)," on page 288
- Section 20.10, "Using Encrypted Volumes in a Server Cluster (Linux)," on page 288
- Section 20.11, "Using Encrypted Volumes in a Server Cluster (NetWare)," on page 289
- Section 20.12, "Removing Encrypted Volumes," on page 290
- Section 20.13, "What's Next," on page 290

20.1 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.

Dynamic Storage Technology (NSS) does not support using encrypted volumes in a DST shadow volume pair.

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.

- Section 20.1.1, “Encryption Method,” on page 282
- Section 20.1.2, “Encryption Password,” on page 282
- Section 20.1.3, “How Encrypted Volume Support Works,” on page 282
- Section 20.1.4, “Guidelines for Using Encrypted Volumes,” on page 283

20.1.1 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 storage object. After the cryptographic data is wrapped for the activated volume, EVS eliminates the password from memory.

20.1.2 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.

20.1.3 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 the time 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.

20.1.4 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.

20.2 Security Considerations for Encrypted Volumes

- Section 20.2.1, “Choosing a Strong Encryption Password,” on page 283
- Section 20.2.2, “Backing Up Data from an Encrypted Volume,” on page 283
- Section 20.2.3, “Excluding the NSS Cache Memory from Core Dumps,” on page 283
- Section 20.2.4, “Disabling Logs (Linux),” on page 284
- Section 20.2.5, “Archiving Files from an Encrypted Volume,” on page 284
- Section 20.2.6, “Using Direct I/O to an Encrypted Volume,” on page 284
- Section 20.2.7, “Sharing Encrypted NSS Volumes in a Cluster,” on page 284

20.2.1 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 34.16, “Creating Strong Passwords,” on page 449.

20.2.2 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 34.4, “Protecting Data During Backup and on Backup Media,” on page 442.

20.2.3 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 34.5, “Preventing Exposure of Sensitive Data in a Core Dump,” on page 442.

20.2.4 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 the 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 34.6, “Preventing Exposure of the Encryption Password in a Log (Linux),” on page 442.

20.2.5 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.

For information about managing Archive and Version Services, see the following:

- *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*
- *OES 2: Novell Archive and Version Services 2.1 for NetWare Administration Guide*

20.2.6 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.

20.2.7 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 Novell Cluster Services™ 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.

20.3 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. You cannot encrypt the `sys:` volume.

EVS is supported by NSS volumes on NetWare 6.5 SP2 and later, OES NetWare and later, and OES SP1 Linux and later operating systems. If you choose to encrypt a volume, you cannot roll back the system to earlier versions of NetWare or OES Linux without taking steps to preserve your data before the rollback. For information, see Section 20.12, "Removing Encrypted Volumes," on page 290.

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 *Volumes*, 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 20.4, "Mounting an Encrypted NSS Volume with NSSMU (Linux)," on page 285.

On NetWare, you can also enter the password as needed when you activate a volume from the command line. For information, see Section 20.6, "Mounting Encrypted NSS Volumes with NSS Commands (Linux)," on page 286.

For information about entering the password for a volume in a cluster, see Section 20.10, "Using Encrypted Volumes in a Server Cluster (Linux)," on page 288.

20.4 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. 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.

20.5 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. 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.

20.6 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.

IMPORTANT: 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 20.4, “Mounting an Encrypted NSS Volume with NSSMU (Linux),” on page 285.

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.

Table 20-1 *Volume Mount Commands*

Command	Description
`mount` *`volname`*	Mounts 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.

Command	Description
`mount all`	Activates and mounts 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.
`nss /volumes`	Displays a list of encrypted and unencrypted NSS volumes, showing their attributes. The encrypted volume returns a status of Encrypted.

20.7 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`.

Table 20-2 *Volume Activation Commands*

Command	Description
`nss /activate=`*`volname`*	Activates 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.
`nss /volumeactivate=`*`volname`*	Activates 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.
`nss /activate=all`	Activate all encrypted NSS volumes that have been previously activated with their passwords.
`nss /forceactivate=`*`volname`*	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.

Command	Description
`nss /volumes`	Displays a list of encrypted and unencrypted NSS volumes, showing their attributes. The encrypted volume returns a status of Encrypted.

20.8 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, start NSSMU, then select *Volumes*.
   2b. In the *Volumes* list, select the encrypted volume that you want to dismount.
   2c. Press *F7* to dismount the mounted volume.

20.9 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. Press *F6* to deactivate the active volume.

20.10 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 Novell Cluster Services 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 that 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 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 Press *F7* to dismount the shared volume.

3 Follow the normal procedures to activate the cluster resources.

For information, see the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*. 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.

20.11 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 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*. 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.

20.12 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 20.3, “Creating an Encrypted Volume,” on page 284.

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.

20.13 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 253.

Managing Removable Media (NetWare)

21

On NetWare®, CDs, DVDs, CD and DVD image files, and DOS partitions are mounted as Novell® Storage Services™ volumes. Disks in USB floppy drives are mounted as local DOS FAT partitions.

IMPORTANT: On Linux, removable media are mounted as Linux POSIX file systems and partitions. For a list of supported formats, see "Some Other Supported File Systems" in the *SLES 10 SP2 Installation and Administration Guide*.

This section describes how to mount removable media on NetWare:

- Section 21.1, "Managing CDs and DVDs as NSS Volumes (NetWare)," on page 291
- Section 21.2, "Activating and Deactivating ISO Image Files as NSS Volumes (NetWare)," on page 293
- Section 21.3, "Activating and Deactivating DOS Partitions as NSS Volumes (NetWare)," on page 294
- Section 21.4, "Mounting a USB Disk as a Local FAT Volume or NSS Volume (NetWare)," on page 295

21.1 Managing CDs and DVDs as NSS Volumes (NetWare)

When you insert CDs and DVDs into the drive of a NetWare server, the removable media are loaded as read-only NSS volumes.

- Section 21.1.1, "Understanding CDs and DVDs on NetWare," on page 291
- Section 21.1.2, "Mounting a CD or DVD as an NSS Volume," on page 292
- Section 21.1.3, "Creating a Volume Object for a CD or DVD," on page 292
- Section 21.1.4, "Specifying the Load Order of Formats for Dual-Formatted," on page 292
- Section 21.1.5, "Loading and Unloading the cddvd.nss Module," on page 292

21.1.1 Understanding CDs and DVDs on NetWare

The `cddvd.nss` module must be loaded to support the handling of removable media on NetWare. 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.

The `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.

21.1.2 Mounting a CD or DVD as an NSS Volume

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.

21.1.3 Creating a Volume Object for a CD or DVD

Although the NSS volume is activated, it is not automatically added to Novell eDirectory™. You can add it to eDirectory by enabling the `UpdateCDtoNDS` switch.

1 At the server console prompt, enter

```
nss /UpdateCDtoNDS=cdname
```

21.1.4 Specifying the Load Order of Formats for Dual-Formatted

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 with the `LoadISOVolumeFirst` switch.

To Load ISO 9660 format first:

1 At the server console, enter

```
nss /LoadISOVolumeFirst
```

21.1.5 Loading and Unloading the cddvd.nss Module

The `cddvd.nss` module loads automatically when NSS starts.

To load the `cddvd.nss` module manually:

1 At the server console prompt, enter

```
load cddvd.nss
```

To unload `cddvd.nss`:

1 At the server console prompt, enter

```
unload cddvd.nss
```

21.2 Activating and Deactivating ISO 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.

- Section 21.2.1, "Guidelines for Using Image Files as NSS Volumes," on page 293
- Section 21.2.2, "Mounting an Image Volume," on page 293
- Section 21.2.3, "Removing an Image Volume," on page 294

21.2.1 Guidelines for Using Image Files as NSS Volumes

The image volume functions the same as any other NSS volume, except that it is read-only. While the image is mounted as a volume, you can mount and dismount the volume like a normal volume. You cannot modify the ISO file when it is in use as a mounted disk.

IMPORTANT: While the image file is mounted as a volume, make sure you access it and its contents only as a mounted volume. Accessing the image file itself as a file causes it to be immediately dismounted and removed as a volume.

While the CD or DVD image file is mounted on the server as an NSS volume, the image file is immediately dismounted and removed as an NSS volume if the following occurs:

- The NSS volume that contains the image source file is dismounted, deactivated, or deleted.

  If the image file is removed as a volume, mount the parent volume, then mount the image file as a volume.

- Accessing the image file itself by any application. This includes selecting the image file in a file browser as a file on the parent volume, or backing up the parent volume.

  If the image file is removed as a volume, close the application or browser that is accessing the file, then mount the image file as a volume.

21.2.2 Mounting an Image Volume

- "Prerequisite" on page 293
- "Procedure" on page 293

Prerequisite

The image files must reside on an existing NSS volume on NetWare, and that volume must be mounted.

Procedure

To activate a CD or DVD image file as an NSS volume on NetWare, use the `MountImageVolume` switch.

NOTE: It is not possible to mount ISO files that reside on another server this way. Use NFS, CIFS, or similar protocols to do that.

1 At the server console prompt, enter

```
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:\suse10_1.iso
```

21.2.3 Removing an Image Volume

To deactivate a CD or DVD image file as an NSS volume on NetWare, use the `RemoveImageVolume` switch. This removes the image volume, but it does not delete the image file itself.

1 At the server console prompt, enter

```
nss /RemoveImageVolume=path_imagefilename
```

Replace *path_imagefilename* with the path on the NSS volume where the image file resides, including its filename. For example:

```
nss /RemoveImageVolume=nssvol:\suse10_1.iso
```

21.3 Activating and Deactivating DOS Partitions as NSS Volumes (NetWare)

The `dosfat.nss` module allows any DOS FAT partitions to be 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_f`.

The dosfat.nss module provides the following benefits:

- **Long Filenames:** DOS filenames for file systems such as OS/2*. This uses the Windows 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.

The `dosfat.nss` module loads automatically when NSS starts.

To load the `dosfat.nss` module manually:

1 At the server console prompt, enter

```
load dosfat.nss
```

To unload `dosfat.nss`:

1 At the server console prompt, enter

```
unload dosfat.nss
```

21.4 Mounting a USB Disk as a Local FAT Volume or NSS Volume (NetWare)

NetWare 6.5 SP5 and later supports storage media attached via USB connection, such as USB flash drives, floppy drives, or hard drives. The USB drives are automatically recognized and mounted if the drive contains unpartitioned free space or is formatted with the FAT16, FAT32, or NSS file systems. Other file systems such as NTFS are not supported. Mount USB devices to drive letters above C:.

Any unpartitioned free space larger than 10 MB can be configured with NSS pools and volumes, just as for any local drive on NetWare. Use the standard NSS management tools, commands, and utilities to control NSS volumes on USB-attached media.

The SET parameter and command called Local FAT Volume Mount (`lfvmount`) controls whether automatic detection and mount is enabled (On, default) or disabled (Off) for local FAT volumes. Local FAT volumes are available to the server but are not seen by users. Use the commands in Table 21-1 to control the detection and mounting behavior for local FAT volumes on your NetWare server:

Table 21-1 *Local FAT Volume Commands*

Command	Description
`set auto lfvmount=value`	Determines whether automatic detection and mounting is enabled (default) or disabled for local FAT volumes. If `lfvmount` is set to On (default), NetWare automatically detects and mounts local FAT volumes as you insert the media, and assigns them a drive letter. If `lfvmount` is set to Off, you must enter the `lfvmount` command at a system console prompt to manually mount FAT volumes as needed. This setting persists until the server is rebooted. Put the command in the `startup.ncf` file if you want it to persist through a server reboot. You can also set this parameter using the MONITOR utility. Value: On (default), Off
`lfvmount`	Mounts local FAT volumes. For example, you would re-enter this command each time you change a floppy disk in the USB floppy drive.
`lfvolumes`	Displays a list of all local FAT volumes that are currently mounted and provides details about them.
`help lfvmount`	Displays syntax information for using the `lfvmount` command.

Securing Access to NSS Volumes, Directories, and Files

22

This section describes measures you can use to help secure access to your Novell® Storage Services™ (NSS) volumes and user data.

- Section 22.1, “Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes,” on page 297
- Section 22.2, “Configuring the Security Equivalence Vector Update Frequency (Linux),” on page 303
- Section 22.3, “Using Data Shredding to Prevent Access to Purged Files,” on page 306
- Section 22.4, “Enabling or Disabling LAF Audit Log Messages for Trustee Events (Linux),” on page 307
- Section 22.5, “Controlling Access to the _admin Volume and Files.cmd (NetWare),” on page 311

22.1 Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes

NSS uses the Novell Trustee model for controlling access to user data. As an administrator or a user with the Supervisor right or Access Control right, you can use the Files and Folders plug-in to iManager to manage file system trustees, trustee rights, inherited rights filters, and attributes for a file or folder on an NSS volume. A user who has only the Access Control right cannot modify the rights of another user who has the Supervisor right.

IMPORTANT: For more information and alternate methods for configuring file system trustees and attributes for directories and files on NSS volumes, see the *OES 2 SP1: File Systems Management Guide*.

- Section 22.1.1, “Prerequisites for Configuring Trustees,” on page 297
- Section 22.1.2, “Viewing Properties of a File or Folder,” on page 298
- Section 22.1.3, “Configuring File or Folder Attributes,” on page 298
- Section 22.1.4, “Configuring Rights Properties (File System Trustees, Trustee Rights, and Inherited Rights Filter),” on page 300
- Section 22.1.5, “Viewing Effective Rights for a Trustee,” on page 303

22.1.1 Prerequisites for Configuring Trustees

- The volume that you want to manage must be in the same tree where you are currently logged in to iManager.

- You must have trustee rights for the volume, folder, and file that you want to manage.
- The volume must be a file system that uses the Novell trustee model for file access, such as an NSS volume on OES 2 NetWare or Linux, an NSS or NetWare traditional volume on NetWare® 6.5, or an NCP™ (NetWare Core Protocol™) volume (an NCP share on a Linux POSIX file system) on OES 2 Linux.

22.1.2 Viewing Properties of a File or Folder

1. In iManager, click *Files and Folders* > *Properties* to open the *Properties* page.
2. Click the *Search* icon to browse and locate volume, folder or file from the Storage objects, then click the name link of the object to select it.

   The pathname of the object appears in the *Name* field.
3. View the following properties in three *Properties* tabs:

Properties Tabs	Description	For Information
Information	◆ View details about the selected volume, folder, or file. ◆ Configure directory quotas for folders on NSS volumes where the Directory Quotas attribute is enabled. ◆ Modify the file owner. ◆ Configure file or directory attributes.	See Section 28.5, "Viewing or Modifying File or Folder Properties," on page 391. See Section 28.7, "Viewing, Adding, Modifying, or Removing a Directory Quota," on page 396. See Section 28.6, "Viewing or Modifying File Ownership," on page 394. See Section 22.1.3, "Configuring File or Folder Attributes," on page 298.
Rights	◆ View details about trustees, trustee rights, and inherited rights filter for the selected volume, folder, or file. ◆ Add or remove trustees. ◆ Grant or revoke trustee rights for one or more trustees. ◆ Configure the inherited rights filter.	See Section 22.1.4, "Configuring Rights Properties (File System Trustees, Trustee Rights, and Inherited Rights Filter)," on page 300.
Inherited Rights	◆ View details about explicitly assigned trustee rights and inherited rights at all levels along the path from the selected file or folder to the `root` of the volume. ◆ View the effective rights for a given trustee for the selected volume, folder, or file.	See Section 22.1.5, "Viewing Effective Rights for a Trustee," on page 303.

22.1.3 Configuring File or Folder Attributes

File attributes determine how the file or folder behaves when accessed by any user. File attributes apply universally to all users. For example, a file that has a read-only attribute is read-only for all users.

Attributes can be set by any trustee with the Modify right to the directory or file, and attributes stay set until they are changed. Attributes do not change when you log out or when you down a file server.

For example, if a trustee with the Modify right enables the Delete Inhibit attribute for a file, no one, including the owner of the file or the network administrator, can delete the file. However, any trustee with the Modify right can disable the Delete Inhibit attribute to allow the file's deletion.

1. In iManager, click *Files and Folders* > *Properties* to open the *Properties* page.
2. Click the *Search* icon to browse and locate volume, folder or file from the Storage objects, then click the name link of the object to select it.

   The pathname of the object appears in the *Name* field. For example:

   ```
   VOL1:dir1\dirB\filename.ext
   ```
3. Click the *Information* tab to view or modify the file or folder attributes. Enable or disable an attribute by selecting or deselecting the check box next to it.

   IMPORTANT: Changes do not take effect until you click *OK* or *Apply*. If you click a different tab before you save, changes you make on this page are lost.

The following table defines file system attributes and whether they apply to files, folders, or both files and folders.

Attribute	Description	Files	Folders
Read Only	Prevents a file from being modified. This attribute is typically used in combination with Delete Inhibit and Rename Inhibit.	Yes	No
Archive	Identifies files and folders that have been modified since the last backup. This attribute is assigned automatically.	Yes	Yes
Hidden	Hides directories and files so they do not appear in a file manager or directory listing.	Yes	Yes
Shareable	Allows more than one user to access the file at the same time. This attribute is usually used with Read Only.	Yes	No
Transactional	Allows a file on an NSS volume or a NetWare Traditional volume to be tracked and protected by the Transaction Tracking System™ (TTS™) for NetWare. For NSS, the TTS attribute for the volume must be enabled in order for this setting to be enforced. TTS is not available for NSS on Linux.	Yes	No
Purge Immediate	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
Rename Inhibit	Prevents the directory or filename from being modified.	Yes	Yes
Delete Inhibit	Prevents users from deleting a directory or file. This attribute overrides the file system trustee Erase right. When Delete Inhibit 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 attribute to allow the directory or file to be deleted.	Yes	Yes
Copy Inhibit	Prevents users from copying a file. This attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare. This attribute overrides the trustee Read right and File Scan right. A trustee with the Modify right must disable this attribute to allow the file to be copied.	Yes	No

4 If you modified any settings, click *Apply* or *OK* to save your changes.

22.1.4 Configuring Rights Properties (File System Trustees, Trustee Rights, and Inherited Rights Filter)

File system trustees, trustee rights, and inherited rights filters are used to determine access and usage for directories and files on NSS volumes on OES 2 NetWare and Linux, NCP volumes on OES 2 Linux, and NSS and NetWare Traditional volumes on NetWare 6.5. If you modify any settings, you must click *Apply* or *OK* to save the changes.

Viewing, Adding, or Removing File System Trustees

A trustee is any Novell eDirectory object (such as a User object, Group object, Organizational Role object, or other container object) that you grant one or more rights for a directory or file. Trustee assignments allow you to set permissions for and monitor user access to data.

1. In iManager, click *Files and Folders*, then click *Properties* to open the *Properties* page.
2. On the *Properties* page, select a volume, folder, or file to manage.

   For instructions, see Section 22.1.2, "Viewing Properties of a File or Folder," on page 298.
3. Click the *Rights* tab to view the trustees, trustee rights, and inherited rights filter for the selected volume, folder, or file.
4. Add trustees.
   - **4a** Scroll down to the *Add Trustees* field.
   - **4b** Use one of the following methods to add usernames as trustees:
     - Click the *Search* icon, browse to locate the usernames of the users, groups, or roles that you want to add as trustees, click the name link of the objects to add them to the *Selected Objects* list, then click *OK*.
     - Click the *History* icon to select usernames from a list of users, groups, or roles that you recently accessed.
     - Type the typeless distinguished username (such as username.context) in the *Add Trustees* field, then click the *Add* (+) icon.

     The usernames appear in the Trustees list, but they are not actually added until you click *Apply* or *OK*. Each of the usernames has the default Read and File Scan trustee rights assigned.
   - **4c** On the *Properties* page, click *Apply* to save the changes.
5. Remove trustees.
   - **5a** Scroll down to locate and select the username of the user, group, or role that you want to remove as a trustee.
   - **5b** Click the *Remove* (red X) icon next to the username to remove it as a trustee.

     The username disappears from the list, but it is not actually removed until you click *Apply* or *OK*.
   - **5c** On the *Properties* page, click *Apply* to save changes.

Viewing, Granting, or Revoking File System Trustee Rights

Administrator users and users with the Supervisor right or the Access Control right can grant or revoke file system trustee rights for a volume, folder, or file. Only the administrator user or user with the Supervisor right can grant or revoke the Access Control right.

1. In iManager, click *Files and Folders*, then click *Properties* to open the *Properties* page.
2. On the *Properties* page, select a volume, folder, or file to manage.

   For instructions, see Section 22.1.2, "Viewing Properties of a File or Folder," on page 298.
3. Click the *Rights* tab to view the trustees, trustee rights, and inherited rights filter for the selected volume, folder, or file.
4. Scroll to locate the username of the trustee you want to manage.

5 In the check boxes next to the trustee name, select or deselect the rights you want to grant or revoke for the trustee.

IMPORTANT: Changes do not take effect until you click *OK* or *Apply*. If you click a different tab before you save, any changes you have made on this page are lost.

Trustee Right	Description
Supervisor (S)	Grants the trustee all rights to the directory or file and any subordinate items. The Supervisor right cannot be blocked with an inherited rights filter (IRF) and cannot be revoked. Users who have this right can also grant other users any rights to the directory or file and can change its inherited rights filter. Default=Off
Read (R)	Grants the trustee the ability to open and read files, and open, read, and execute applications. Default=On
Write (W)	Grants the trustee the ability to open and modify (write to) an existing file. Default=Off
Erase (E)	Grants the trustee the ability to delete directories and files. Default=Off
Create (C)	Grants the trustee the ability to create directories and files and salvage deleted files. Default=Off
Modify (M)	Grants the trustee the ability to rename directories and files, and change file attributes. Does not allow the user to modify the contents of the file. Default=Off
File Scan (F)	Grants the trustee the ability to view directory and filenames in the file system structure, including the directory structure from that file to the root directory. Default=On
Access Control (A)	Grants the trustee the ability to add and remove trustees for directories and files and modify their trustee assignments and inherited rights filters. Default=Off

6 Click *Apply* or *OK* to save changes.

Configuring the Inherited Rights Filter for a File or Directory

File system trustee rights assignments made at a given directory level flow down to lower levels until they are either changed or masked out. This is referred to as inheritance. The mechanism provided for preventing inheritance is called the inherited rights filter. Only those rights allowed by

the filter are inherited by the child object. The effective rights that are granted to a trustee are a combination of explicit rights set on the file or folder and the inherited rights. Inherited rights are overridden by rights that are assigned explicitly for the trustee on a given file or folder.

1. In iManager, click *Files and Folders*, then click *Properties* to open the *Properties* page.
2. On the *Properties* page, select a volume, folder, or file to manage.

   For instructions, see Section 22.1.2, "Viewing Properties of a File or Folder," on page 298.
3. Click *Information*, then scroll down to view the inherited rights filter.

   The selected rights are allowed to be inherited from parent directories. The deselected rights are disallowed to be inherited.
4. In the *Inherited Rights Filter*, enable or disable a right to be inherited from its parent directory by selecting or deselecting the check box next to it.
5. Click *Apply* or *OK* to save the changes.

22.1.5 Viewing Effective Rights for a Trustee

Effective rights are the explicit rights defined for the trustee plus the rights that are inherited from the parent directory. The *Inherited Rights* page shows the inheritance path for a trustee for the selected file or folder and the effective rights at each level from the current file or directory to the root of the volume. You can use this information to help identify at which directory in the path a particular right was filtered, granted, or revoked.

1. In iManager, click *Files and Folders*, then click *Properties* to open the *Properties* page.
2. On the *Properties* page, select a volume, folder, or file to manage.

   For instructions, see Section 22.1.2, "Viewing Properties of a File or Folder," on page 298.
3. On the *Properties* page, click the *Inherited Rights* tab to view the effective rights for a given trustee.

   By default, the page initially displays the effective rights for the username you used to log in to iManager.
4. On the *Inherited Rights* page, click the *Search* icon next to the *Trustee* field to browse for and locate the username of the trustee you want to manage, then select the username by clicking the name link.

   The path for the selected file or folder is traced backwards to the root of the volume. At each level, you can see the rights that have been granted and inherited to create the effective rights for the trustee.
5. If you make any changes, click *Apply* or *OK* to save them.

22.2 Configuring the Security Equivalence Vector Update Frequency (Linux)

The Security Equivalence Vector (SEV) is used to validate the user against the trustee rights of the directory and file the user is attempting to access. You can use commands in the NSS Console utility (`nsscon`) on Linux 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.

- Section 22.2.1, "Understanding the SEV," on page 304

- Section 22.2.2, “Enabling or Disabling the Background SEV Update for NSS Volumes on Linux,” on page 304
- Section 22.2.3, “Configuring the Background SEV Update Interval,” on page 305
- Section 22.2.4, “Forcing a Background SEV Update,” on page 305

22.2.1 Understanding the SEV

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

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.

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 automatically retrieves the user’s SEV from the NetWare connection structure, then deletes it when the session ends.

On Linux, the SEV behavior differs because NSS does not have the same integrated relationship to the connection infrastructure as it does on NetWare. NSS caches the SEV locally in the server memory, where it remains until the server is rebooted or the user is deleted from eDirectory. NSS polls eDirectory at a specified interval for updates to the SEVs that are in cache.

22.2.2 Enabling or Disabling the Background SEV Update for NSS Volumes on Linux

By default, the SEV is updated in the background and whenever the server is rebooted. You can optionally disable the background updating. If it is disabled, the user access can become unsynchronized over time, so that users might have less or more access than you have configured. We recommend that you leave the SEV updating feature enabled, then modify the polling frequency to best meet the security needs of your production environment.

To enable or disable the setting:

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, start the NSS Console by entering

   ```
   nsscon
   ```
3. At the `nsscon` prompt, do one of the following:
   - **Enable:** This is the default. To enable the background updating of the SEV in addition to the default update at server reboot, enter

```
nss /SecurityEquivalenceUpdating
```

- **Disable:** To disable the background updating, enter

```
nss /NoSecurityEquivalenceUpdating
```

The SEV Update is enabled when you first reboot the server. If you disable SEV updates and want the setting to persist across server reboots, include the `/SecurityEquivalenceUpdating` option in the `/etc/opt/novell/nss/nssstart.cfg` file.

22.2.3 Configuring the Background SEV Update Interval

You might want to modify the background SEV update interval to make the polling for eDirectory updates to be more or less frequent. Polling too frequently can impact performance. Polling too infrequently can cause delays in granting or restricting access for certain users. To avoid possible security violations, you can also force an update at any time by using the /ForceSecurityEquivalenceUpdate command. For information, see Section 22.2.4, "Forcing a Background SEV Update," on page 305.

The interval for the background updating of the SEV is the elapsed time between the last update and the next one. At the end of the elapsed time, NSS requires updated SEVs from eDirectory. The default interval is 7200 seconds (2 hours). The valid range is 300 (5 minutes) to 7776000 (90 days).

To set the interval to use until the next server reboot:

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, start the NSS Console by entering

   ```
   nsscon
   ```

3. At the `nsscon` prompt, enter

   ```
   nss /UpdateSecurityEquivalenceInterval=value
   ```

   Replace *value* with the desired interval.

To make the interval setting persistent across server reboots, include the `/UpdateSecurityEquivalenceInterval=value` option in the `/etc/opt/novell/nss/nssstart.cfg` file.

22.2.4 Forcing a Background SEV Update

If you modify user's access control settings or remove a user from eDirectory in between SEV update intervals, you can to force the SEV to be updated immediately after that to avoid possible security violations. Use the `/ForceSecurityEquivalenceUpdate` option to force an immediate update for all users in the NSS file system so that your changes can be reflected immediately in the user's active SEV for this server.

To force an immediate update:

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, start the NSS Console by entering

   ```
   nsscon
   ```

3. At the `nsscon` prompt, enter

```
nss /ForceSecurityEquivalenceUpdate
```

22.3 Using Data Shredding to Prevent Access to Purged Files

Data shredding hides 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. If Salvage is enabled, there remains a purge delay between when the file is deleted and purged during which users can still salvage deleted files.

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 describes the following:

- Section 22.3.1, “Setting the Data Shredding Attribute When You Create a Volume,” on page 306
- Section 22.3.2, “Setting the Data Shredding Attribute for an Existing Volume,” on page 306
- Section 22.3.3, “Disabling Data Shredding for an Existing Volume,” on page 307

22.3.1 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 19.3, “Creating Unencrypted NSS Volumes,” on page 259.

22.3.2 Setting the Data Shredding Attribute for an Existing Volume

1. In iManager, click *Storage* > *Volumes* to open the Volumes page.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

   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.

22.3.3 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.

1 At the server console, enter

```
nss /nodatashredding=volumename
```

where *volumename* is the name of the volume where you want to prevent the shredding capability.

22.4 Enabling or Disabling LAF Audit Log Messages for Trustee Events (Linux)

Use the NSS audit log messages commands to enable or disable messages via Lightweight Auditing Format (LAF) for NSS trustee changes for NSS volumes on OES 2 Linux.

- Section 22.4.1, “Understanding NSS Audit Log Messages,” on page 307
- Section 22.4.2, “Enabling or Disabling LAF Audit Messages for Trustee Events (Linux),” on page 310
- Section 22.4.3, “Viewing LAF Audit Messages,” on page 310
- Section 22.4.4, “Additional Information,” on page 311

22.4.1 Understanding NSS Audit Log Messages

When the LAFAuditTrustee parameter is enabled, NSS reports changes for the following subset of NSS events:

- Adding trustees (AddTrustee)
- Removing trustees (RemoveTrustee)
- Setting the inherited rights mask (SetInheritedRightsMask)

Comma separated name value pairs are used for the NSS audit log messages. The messages are written to the `/var/log/audit/audit.log` file.

The types of information reported are described below:

- “Message Type and ID” on page 308
- “Add Trustee Event Messages” on page 308
- “Remove Trustee Event Messages” on page 308
- “Set Inherited Rights Mask Event Messages” on page 308

- "Trustee Rights" on page 309
- "Inherited Rights Mask for Trustee Rights" on page 309
- "Special Rights" on page 309
- "Inheritance Attributes" on page 310

Message Type and ID

All NSS Audit Log messages are of the type AUDIT_KERNEL_OTHER (1316) for LAF. For example, the log messages begin

```
type=UNKNOWN[1316] msg=audit(message_id):
```

Add Trustee Event Messages

The general format of NSS audit log messages for a single AddTrustee event is:

```
NSS: AddTrustee: fsuid=<user requesting the
operation>,vol=<VOLNAME>,path=<FULL_PATH (relative to the
volume)>,trustee=<typeful Fully Distinguished eDirectory username of the trustee
being added>,rights=<RIGHTS>,attributes=<ATTRIBUTES>
```

For example, the following message is for a single event for adding a trustee:

```
type=UNKNOWN[1316] msg=audit(1164926678.066:7): NSS: AddTrustee:
fsuid=0,vol=NSS1,path=/abc/
a,trustee=.CN=user5.O=company.T=COMPANY_TREE.,rights=0x1fb,attributes=0xc000
```

In this example, the trustee user5.company.company_tree is assigned the SRWCEMFA rights, totaling 0x1fb. For a map of rights to values, see "Trustee Rights" on page 309.

Remove Trustee Event Messages

The general format of NSS audit log messages for a single RemoveTrustee event is:

```
NSS: RemoveTrustee: fsuid=<user requesting the
operation>,vol=<VOLNAME>,path=<FULL_PATH (relative to the
volume)>,trustee=<typeful Fully Distinguished eDirectory username of the trustee
being removed>
```

For example, the following message is for a single event for removing a trustee:

```
type=UNKNOWN[1316] msg=audit(1164926734.422:8): NSS: RemoveTrustee:
fsuid=0,vol=NSS1,path=/abc/a,trustee=.CN=user5.O=company.T=COMPANY_TREE.
```

Set Inherited Rights Mask Event Messages

The general format of NSS audit log messages for a single SetInheritedRightsMask event is:

```
NSS: SetInheritedRightsMask: fsuid=<user>,vol=<VOLNAME>,path=<FULL_PATH(relative
to the volume)>,inheritedRightsMask=<RIGHTS>
```

For example, the following message is for a single event for changes to the inherited rights mask:

```
type=UNKNOWN[1316] msg=audit(1164926882.005:10): NSS: SetInheritedRightsMask:
fsuid=0,vol=NSS1,path=/abc/a,inheritedRightsMask=0x149
```

In this example, the trustee rights settings can be inherited from the parent directory for the Supervisor (0x0100), Read (0x0001), Create (0x0008), and File Scan (0x0040) rights, totaling 0x0149.

Trustee Rights

The file system trustee rights setting in the message is a hexadecimal value that represents the combination of rights assigned.

The following table maps the trustee rights to hexadecimal values. The values for enabled rights are added to get the reported value for the Rights and Inherited Rights Mask.

Trustee Right	Hexadecimal Value
Supervisor (S)	0x0100
Read (R)	0x0001
Write (W)	0x0002
Create (C)	0x0008
Erase (E, Delete)	0x0010
Modify (M)	0x0080
File Scan (F, See Files)	0x0040
Access Control (A)	0x0020

For example, if the trustee has SRWCEMFA rights, the value is the sum of these or 0x1fb in hexadecimal.

Inherited Rights Mask for Trustee Rights

An inherited rights mask (IRM) specifies which trustee rights are allowed to be inherited downward through a directory. If a trustee bit is set in the IRM of a directory, that bit can be inherited downward in the tree. If a trustee bit is not set in the IRM of a directory, then that right cannot be inherited by the directory's contents, even if a higher level in the directory tree had that right.

The bit definitions for inherited rights masks are the same bits as the trustee rights themselves as described in “Trustee Rights” on page 309. For example, if the Read and File Scan rights can be inherited, the inherited rights mask value is 0x0041 in hexadecimal.

Special Rights

In addition to trustee rights, the following are special rights that might be reported in the Rights field. They cannot be inherited.

Special Right	Hexadecimal Value
Salvage	0x0200
Secure	0x8000

Inheritance Attributes

The attributes reported in the log are flags that tell the trustee how it gets inherited. (They are not file system attributes.) By default, the NetWare trustee model inherits downward and upward (visibility inherits upward; actual rights inherit downward).

The following table maps the inheritance attributes to hexadecimal values. The values for enabled inheritance attributes are added to get the reported value for the Attributes parameter.

Inheritance Attribute	Hexadecimal Value
Inherit Down Make rights inherit downward.	0x8000
Inherit Up Make directories above this file visible.	0x4000
Negative Rights (Not currently used) All other bits are ignored if this parameter is set.	0x2000

For example, an Attribute value of 0xc000 in the audit message indicates that both the Inherit Down and Inherit Up parameters are enabled. This is the typical setting for NSS file systems.

22.4.2 Enabling or Disabling LAF Audit Messages for Trustee Events (Linux)

Enable or disable the generation of audit messages via LAF for NSS trustee changes. After you enable the audit log messages, the setting persists until the server reboot. After a server reboot, the audit log is disabled again by default. To make the command persist across reboots, add it to the `/etc/opt/novell/nss/nssstart.cfg` file. The messages are written to the`/var/log/audit/audit.log` file.

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, start the NSS Console by entering

   ```
   nsscon
   ```
3. At the `nsscon` prompt, do one of the following:
   - **Enable:** To enable audit messages for an NSS volume, enter

     ```
     nss /LAFAuditTrustee=volumename
     ```
   - **Disable:** To disable audit messages for an NSS volume, enter

     ```
     nss /NoLAFAuditTrustee=volumename
     ```

22.4.3 Viewing LAF Audit Messages

View audit messages in the `/var/log/audit/audit.log` file.

For information about the format of the messages, see Section 22.4.1, “Understanding NSS Audit Log Messages,” on page 307.

22.4.4 Additional Information

For information about configuring Linux Audit, see the *Linux Audit Quick Start* (http://www.novell.com/documentation/sled10/auditqs_sp1/data/art_audit_qs.html).

For information about the Linux Audit Framework, see *The Linux Audit Framework* (http://www.novell.com/documentation/sled10/audit_sp1/data/book_sle_audit.html).

22.5 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 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 22.5.1, "Assigning Volume Rights with Trustees.xml (NetWare)," on page 311.

For information about creating software applications and scripts to manage NSS volumes and services on Linux and NetWare, see the *NDK: Virtual File Services* (http://developer.novell.com/ndk/doc/vfs/vfs__enu/data/h5ubo7tt.html) in the Novell Developer Kit Documentation Web site (http://developer.novell.com/wiki/index.php/Category:Novell_Developer_Kit).

22.5.1 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>	Enter the full context and username of the user. You can enter the information in slash-separated format or Novell dot format. **NOTE:** This tag cannot contain multiple usernames. For example, if John Smith existed in the novell context in the aurora_tree, you would enter one of the following: \aurora_tree\novell\jsmith or jsmith.novell.aurora_tree
<filename>	Enter the name of volume, directory, and file that you are assigning rights to. For example, if you are assigning rights to the hosts file in the etc directory located on the sys: volume, you would enter the following into the brackets <filename>sys:\etc\hosts</filename>
<rights>	The child elements to the Rights tag include the following tags for file system rights: ◆ <accesscontrol/> for the Access Control right ◆ <create/> for the Create right ◆ <erase/> for the Erase right ◆ <filescan/> for the File Scan right ◆ <modify/> for the Modify Metadata right ◆ <read/> for the Read right ◆ <supervisor/> for the Supervisor right ◆ <write/> for the Write right

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/>
         <erase/>
      </rights>
   </addTrustee>
```

```
   <addTrustee>
      <name>xxxx.novell.kona_tree</name>
      <fileName>sys:\etc\group</fileName>
      <rights>
         <read/>
         <write/>
         <filescan/>
      </rights>
   </addTrustee>


   <addTrustee>
      <name>fred.novell.kona_tree</name>
      <fileName>sys:\etc\group</fileName>
      <rights>
         <read/>
         <write/>
         <filescan/>
      </rights>
   </addTrustee>

   <addTrustee>
      <name>[public]</name>
      <fileName>_admin:manage_nss\files.cmd</fileName>
      <rights>
         <read/>
         <write/>
         <filescan/>
      </rights>
      <background/>
   </addTrustee>

   <addTrustee>
      <name>fred.novell.kona_tree</name>
      <fileName>sys:\etc\hosts</fileName>
      <rights>
         <read/>
         <write/>
         <filescan/>
      </rights>
   </addTrustee>

</specialTrustees>
```

22.5.2 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 who 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/>
        <filescan/>
      </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>
        <filescan/>
     </rights>
     <background/>
  </addTrustee>
```

For further explanation of the XML tags, see Section 22.5.1, “Assigning Volume Rights with Trustees.xml (NetWare),” on page 311.

Managing Data Integrity on NetWare Servers

23

This section describes Novell® Storage Services™ volume attributes that help ensure data integrity.

- Section 23.1, “Enabling Flush Files Immediately to Write Data to the Disk on Close,” on page 315
- Section 23.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare),” on page 317
- Section 23.3, “Using the File-Level Snapshot Attribute to Enable the Backup of Open Files,” on page 318

Use Novell Storage Management Services™ (SMS) for backup, restore, and data migration services. For information, see the *OES 2 SP1: SBCON Administration Guide*.

23.1 Enabling Flush Files Immediately to Write Data to the Disk on Close

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.

NOTE: On Linux, a group write function controls how writes to disk occur. For information, see Section 31.3, “Configuring or Tuning Group I/O,” on page 415.

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.

- Section 23.1.1, “Setting the Flush Files Immediately Attribute When You Create a Volume,” on page 316
- Section 23.1.2, “Setting the Flush Files Immediately Attribute for an Existing Volume,” on page 316
- Section 23.1.3, “Enabling and Disabling the Flush Files Immediately Attribute for All Volumes on a Server,” on page 317

23.1.1 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 19.3, "Creating Unencrypted NSS Volumes," on page 259.

23.1.2 Setting the Flush Files Immediately Attribute for an Existing Volume

1. In iManager, click *Storage* > *Volumes* to open the Volumes page.

   For information, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   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.

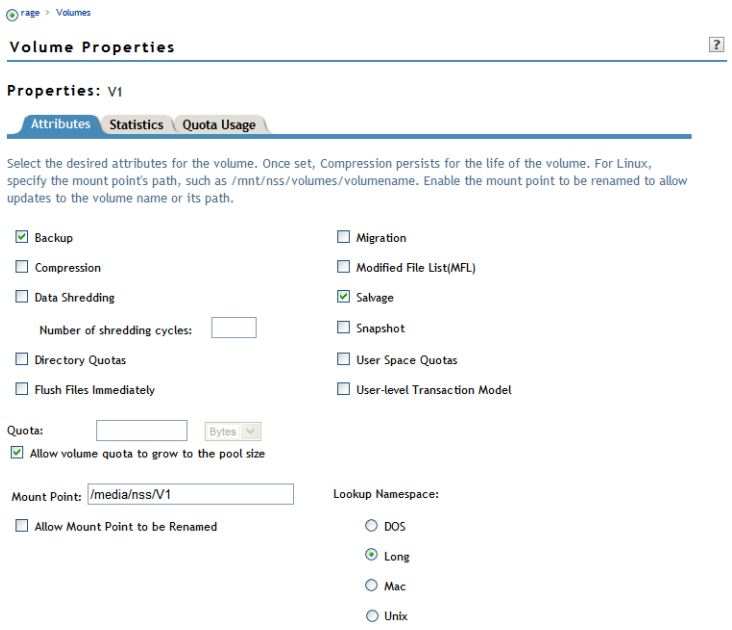

5. Select the *Flush Files Immediately* check box.
6. Click *Apply* or *OK* to save the change, or click *Cancel* to back out of the process.

   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.

23.1.3 Enabling and Disabling the Flush Files Immediately Attribute for All Volumes on a Server

Use the NSS *Flush Files Immediately* attribute to enable and disable the capability for all volumes on the server.

To enable the *Flush Files Immediately* attribute, enter the following at the server console:

```
nss /FlushFilesImmediately
```

To disable the *Flush Files Immediately* attribute, enter the following at the server console:

```
nss /NoFlushFilesImmediately
```

23.2 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" in the *OES2 SP1: Server Operating System for NetWare Administration Guide*.

NOTE: When working on OES 2 Linux, if your data requires content-level transaction tracking and you want to continue to control access to data with trustees, you can use NCP volumes on Linux Reiser or Ext3 file systems. Configure the journaling mode for the file system to the Journaling level, which provides content-level transaction tracking. For information about creating and managing NCP volumes on Linux, see "Managing NCP Volumes" in the *OES 2 SP1: NCP Server for Linux Administration Guide*.

- Section 23.2.1, "Enabling TTS," on page 317
- Section 23.2.2, "Disabling TTS," on page 318
- Section 23.2.3, "TTS Log File and Transaction Recovery," on page 318

23.2.1 Enabling TTS

To enable TTS on an NSS volume, enter the following at the server console prompt:

```
nss /Transaction=volname
```

23.2.2 Disabling TTS

To disable TTS on an NSS volume, enter the following at the server console prompt:

```
nss /NoTransaction=volname
```

To disable TTS on all Traditional volumes, add `disable tts` to the `autoexec.ncf` file.

23.2.3 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:\XLog yyyymmddhhmmss.log
```

Corrupt TTS log files are relatively large, about 10 MB. TTS stores the corrupt TTS log file in the root directory of its transactioned 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.

23.3 Using the File-Level Snapshot Attribute to Enable the Backup of Open Files

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.

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” on page 229.

- Section 23.3.1, “Understanding File-Level Snapshots,” on page 319
- Section 23.3.2, “Guidelines for Using File-Level Snapshots,” on page 320
- Section 23.3.3, “Enabling the File Snapshot Attribute,” on page 321

- Section 23.3.4, “Configuring Backup Software to Use File Snapshots,” on page 322
- Section 23.3.5, “Restoring Volumes from Backup Media,” on page 323

23.3.1 Understanding File-Level Snapshots

- “Advantages File Level Snapshot” on page 319
- “How the File Snapshot Attribute Works” on page 319

Advantages File Level Snapshot

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. The copy of the file is temporary. When the user or API closes the file, NSS deletes the file-snapshot copy.

IMPORTANT: Not all third-party backup software can take advantage of the file snapshot attribute, even if you set it.

How the File Snapshot Attribute Works

The file snapshot attribute uses block-level copy-on-write technology, which is a technique of capturing images of data before 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, other information about the file, and an empty Blocks Changed list.

Whenever the file opens for a Write call, NSS makes a copy of the block’s contents before writing. NSS identifies writes at the block level. The NSS file snapshot attribute updates the metadata for the file, then monitors for writes to the file. 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 it is 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 is 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.

23.3.2 Guidelines for Using File-Level Snapshots

- "Using File-Level Snapshots for Backup of Unstructured Data" on page 320
- "Sizing the Volume to Accommodate File Snapshots" on page 320
- "Assessing I/O Overhead for File Snapshots" on page 321

Using File-Level Snapshots for Backup of Unstructured Data

The copy-on-write technology used in file-level snapshots is designed for most files that spend their lifetimes in a closed state or read-only state. However, copy-on-write is neither desirable nor useful for database backups, which are perpetually held in an open state for writing. Because the copy occurs when the write call is issued for a block, the copy is not updated if subsequent writes are made to the same block before the file is closed. The longer the database is open, the more out of date the copied block becomes.

For example, in a Novell GroupWise® system, the longer the database is open, the more chronologically disjointed the contents of the copied blocks are from the actual contents in them. If you make a backup of the database without first closing the file, you get a backup image, but it is not useful.

Sizing the Volume to Accommodate File Snapshots

Initially, the virtual volume occupies very little storage space. It contains only an empty blocks changed 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 blocks changed 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)

You should 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.

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.

When assessing overhead, it is important to remember that every time a file is closed, NSS erases the file's 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, consider the following questions:

- 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?

23.3.3 Enabling the File Snapshot Attribute

File snapshots apply only to NSS volumes on Netware. After the attribute is set, the file-snapshot process is automatic.

- "Enabling the File Snapshot Attribute for a New Volume" on page 322

- "Enabling the File Snapshot Attribute for an Existing Volume" on page 322
- "Enabling the File Snapshot Attribute Using Commands" on page 322

Enabling the File Snapshot Attribute for a New Volume

You can enable the File Snapshot attribute when you create the volume by selecting the *Snapshot - File Level* check box on the *Attributes* page. Whenever the volume is first activated, NSS creates the virtual volume and the file snapshot process occurs automatically.

Enabling 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. You must deactivate and activate the volume to enable NSS to set up the virtual volume area where file snapshots are stored. 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. In iManager, select the server for the volume you want to manage.
2. In *Roles and Tasks*, select *Storage* > *Volumes*.
3. Select a volume in the *Volumes* list.
4. If the volume is not active and mounted, click *Mount*.

   The *Mount* option 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.
5. Click *Properties* > *Attributes*.
6. Select the *Snapshot - File Level* check box, click *Apply*, then click *OK* to confirm the change.
7. On the *Volumes* page, select *Storage* > *Volumes*.
8. Select the volume, then click *Deactivate*.
9. 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.

Enabling the File Snapshot Attribute Using Commands

For information, see Section A.13, "File-Level Snapshot Commands," on page 467.

23.3.4 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 waits 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 information 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 information mask.

The following calls also return information about File Snapshot:

- `GetFileSize`
- `GetFileHoles`

23.3.5 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

24

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 describes the following:

- Section 24.1, “Understanding Compression,” on page 325
- Section 24.2, “Configuring Compression for a Server,” on page 330
- Section 24.3, “Configuring a Volume for Compression,” on page 336
- Section 24.4, “Suspending Compression for Volumes or Files,” on page 337
- Section 24.5, “Disabling Compression for a Volume,” on page 338
- Section 24.6, “Restoring Data to a Uncompressed Volume,” on page 338
- Section 24.7, “Configuring Compression Preferences for Directories and Files,” on page 338
- Section 24.8, “Using NSS Commands to Configure and Monitor Compression,” on page 342
- Section 24.9, “Copying Compressed Files between Volumes with the NCOPY Command,” on page 343
- Section 24.10, “Viewing Compression Statistics with the NDIR Command,” on page 343
- Section 24.11, “Repairing Compressed Volumes with the Compfix Utility,” on page 344
- Section 24.12, “Backing Up Compressed Files,” on page 344

24.1 Understanding Compression

This section describes the following:

- Section 24.1.1, “Compression and Decompression Processes,” on page 325
- Section 24.1.2, “Compression Settings,” on page 326
- Section 24.1.3, “Guidelines for Compression,” on page 327
- Section 24.1.4, “Factors Affecting Compression,” on page 329
- Section 24.1.5, “Factors Affecting Decompression,” on page 330
- Section 24.1.6, “Monitoring Compression Activity,” on page 330

24.1.1 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 non-peak hours, but it can occur at any time you prefer.

A file must be idle for the period specified in the Days Untouched before Compression parameter before it is considered eligible for a scheduled compression. 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. After the compression completes successfully, NSS deletes the uncompressed file, and keeps 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 24.11, “Repairing Compressed Volumes with the Compfix Utility,” on page 344.

NSS retains the compressed file when it is opened. Before NSS opens a compressed file, it verifies that the file system has enough space available for both the compressed file and the uncompressed file to temporarily coexist. If there is not enough space available, the file is not opened. You must make free space available for the volume before the file can be opened. Use any measures you normally take to make space available for the volume to use.

24.1.2 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 for the server level, for the volume, and optionally for individual directories and files.

Common Service Compression Parameters

At the server level, the settings for compression parameters in Common Services govern when and how compression works for the NSS volumes where the compression attribute is enabled. For information about these parameters, see Section 24.2, “Configuring Compression for a Server,” on page 330.

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 *OES 2 SP1: NetWare Traditional File System Administration Guide*.

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. For information about setting attributes for existing volumes, see Section 24.3, “Configuring a Volume for Compression,” on page 336.

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. For information about how to set compression preferences for individual directories and files, see Section 24.7, "Configuring Compression Preferences for Directories and Files," on page 338.

24.1.3 Guidelines for Compression

To effectively use compression for your NSS volumes, you must understand the following key concepts:

- "Some Volumes Are Not Good Candidates for Compression" on page 327
- "After It Is Set, the Compression Attribute Persists for the Life of the Volume" on page 327
- "Inactivity Determines Which Files Are Eligible for Background Compression" on page 327
- "Some Files Do Not Compress Well" on page 328
- "Decompression Activity Depends on Available Space" on page 328
- "Immediate Compression Impacts CPU Performance" on page 329
- "Files Remain Compressed during Backup and Restore" on page 329

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.

After It Is 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, by using the Enable File Compression parameter. This parameter suspends compression for all volumes on the server. For information, see Section 24.4, "Suspending Compression for Volumes or Files," on page 337 and Section 24.5, "Disabling Compression for a Volume," on page 338.

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.

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 and qualify for compression, then are accessed and uncompressed. 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.

If the volume is on a shared pool, its files might be queued for compression on one node, when its pool is cluster migrated or failed over to another node. The Compression Queue is non-persistent, so on the new node, the volume does not have information about the compression queue from the old node. Therefore, the files must re-qualify for compression when the Background Compression starts on the new node.

On NetWare, use the `flag` command to enable immediate compression for a file or directory. For information, see Section 24.7, “Configuring Compression Preferences for Directories and Files,” on page 338.

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.

Some database files become unavailable when they are compressed, such as Sybase* database files.

IMPORTANT: If you use Sybase database files in a volume, such as for ZENworks® databases, do not enable compression on the volume, or mark each database file with the Don’t Compress (Dc) attribute so that it is never compressed even if compression is enabled for the volume. For details, see *Technical Information Document 10075966* (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10075966.htm).

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.

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 24.2, "Configuring Compression for a Server," on page 330.

Files Remain Compressed during Backup and Restore

Novell Storage Management Services™ (SMS) backs up and restores compressed files in their compressed state. It does not compress uncompressed files for backup; they are stored and restored in their uncompressed state. For more information, see Section 24.12, "Backing Up Compressed Files," on page 344.

24.1.4 Factors Affecting Compression

Typically, files that have been inactive for a specified period of time set in the Days Untouched before Compression parameter are eligible for compression.

The Immediate Compression attribute for a file or directory can also 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 has not been inactive for a period longer than the value in Days Untouched before Compression parameter when the compression daily check begins and compares the file's time stamp to the starting time. The daily check is controlled by the Compression Daily Check Starting Hour parameter.
- The file is flagged with a Don't Compress (Dc) attribute. For information, see Section 24.7, "Configuring Compression Preferences for Directories and Files," on page 338.
- 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. For information, see Section 24.4, "Suspending Compression for Volumes or Files," on page 337.
- The queue of files marked for compression is long and cannot be completed during the specified hours set aside for compression activities. For information on how to modify the Compression Daily Check Stop Hour, see Section 24.2, "Configuring Compression for a Server," on page 330.

- 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. For information on how to modify this setting, see Section 24.2, "Configuring Compression for a Server," on page 330.

24.1.5 Factors Affecting Decompression

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.

24.1.6 Monitoring Compression Activity

Monitor compression activity using the `nss /compscreen` command. For instructions, see Section 24.8, "Using NSS Commands to Configure and Monitor Compression," on page 342.

24.2 Configuring Compression 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.

Before you set parameters, make sure you understand how compression works for NSS. For information, see Section 24.1, "Understanding Compression," on page 325.

- Section 24.2.1, "Understanding Server-Level Compression Parameters," on page 331
- Section 24.2.2, "Configuring Server-Level Compression Parameters with Commands," on page 334
- Section 24.2.3, "Configuring Server-Level Compression Parameters with Novell Remote Manager (NetWare)," on page 335

- Section 24.2.4, “Configuring Server-Level Compression Parameters with the Set Command (NetWare),” on page 335
- Section 24.2.5, “Configuring Server-Level Compression Parameters with Monitor (NetWare),” on page 336

24.2.1 Understanding Server-Level Compression Parameters

The following table describes each compression parameter, its purpose, supported values, and default value.

Table 24-1 *Explanation of Compression Parameters*

Parameter	Description
Days Untouched Before Compression	Specifies the number of days the system waits after a file was last accessed before it is compressed. 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. When background compression starts, it first evaluates which files meet this inactivity requirement to determine which files are to be compressed during the compression period. 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. **Range:** 0 to 100000 **Default:** 14
Compression Daily Check Starting Hour	Specifies the hour when you want the file compressor to start scanning enabled volumes for files that need to be compressed and to compress them. 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. **Range:** 0 to 23 Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.). **Default:** 0
Compression Daily Check Stop Hour	Specifies 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. **Range:** 0 to 23 Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.). **Default:** 6

Parameter	Description
Enable File Compression	Specifies whether file compression is enabled or suspended for all volumes where the Compression attribute is enabled. After an NSS volume's Compression attribute is enabled, it cannot be turned off because the volume contains compressed files and metadata about compression. The server-level Enable File Compression parameter allows you to turn off the compression of more files on the server's compressed volumes. While file compression is suspended, files that would have been compressed are queued for compression, then are compressed only when (or if) the Enable File Compression parameter is reset to On. Files that are already compressed remain compressed, unless they are decompressed when they are opened and used. **Range:** On (default) or Off The On setting allows file compression activity to occur on volumes where the Compression attribute is enabled. It does not enable the Compression attribute on the server's volumes. The Off setting suspends compression on volumes where the Compression attribute is enabled. Immediate compression requests are queued until the value is reset to On, when the files meeting criteria are compressed. The Off setting does not disable the Compression attribute on individual volumes, and it does not prevent you from enabling the Compression attribute for a volume. **Default:** On
Minimum Compression Percentage Gain	Sets the minimum percentage a file must compress to remain in a compressed state. **Range:** 0 to 50 **Default:** 20
Maximum Concurrent Compressions	Specifies the maximum concurrent or simultaneous compressions allowed. **Range:** 1 to 8 **Default:** 2

Parameter	Description
Convert Compressed to Uncompressed Option	Specifies what the file system does with an uncompressed version of a file after the server has decompressed it. **IMPORTANT:** Before a compressed file can be opened, there must be sufficient space available on the volume for the uncompressed and compressed copies of the file to coexist while the file is open. **Range:** 0, 1 (default), or 2 0 = Always leave the file compressed. While the file is open, both the uncompressed and compressed copies of the file coexist on the volume. If the file is closed without changes, the uncompressed copy of the file is discarded. If changes are saved, the compressed copy of the file is discarded. After the modified file is closed, it is queued for immediate compression. Sufficient space must be available for both the compressed and uncompressed copies of the file to temporarily coexist on the volume in order for the compression to occur. After successful compression, the uncompressed copy of the modified file is discarded. 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. This is the default behavior for compression. While the file is open, both the uncompressed and compressed copies of the file coexist on the volume. The first time that the file is closed without changes in the specified period, the uncompressed copy of the file is discarded. The second time that the file is closed without changes in the specified period, the compressed copy of the file is discarded. If changes are saved, the compressed copy of the file is discarded. The uncompressed file remains uncompressed until it meets requirements for being compressed. 2 = Always leave the file uncompressed. While the compressed file is open, both the uncompressed and compressed copies of the file coexist on the volume. When the file is closed or when changes are saved, the compressed copy of the file is discarded. The uncompressed file remains uncompressed until it meets requirements for being compressed.
Decompress Percent Disk Space Free to Allow Commit	Specifies 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. **Range:** 0 to 75 **Default:** 10
Decompress Free Space Warning Interval	Specifies the time between alerts when the file system is not changing compressed files to decompressed because of insufficient disk space. **Range:** 0 seconds to 29 days 15 hours 50 minutes 3.8 seconds Setting the interval to 0 turns off the alert. **Default:** 31 minutes 18.5 seconds

Parameter	Description
Deleted Files Compression Option	Specifies whether and when deleted files are compressed. **Range:** 0, 1, or 2 0 = Do not compress deleted files 1 = Compress deleted files the next day 2 = Compress deleted files immediately **Default:** 1

24.2.2 Configuring Server-Level Compression Parameters with Commands

Use the following commands to modify server-level compression parameters from the command line in NetWare server console for NSS on NetWare, or in `nsscon` for NSS on Linux. For details about each parameter, see Section 24.2.1, "Understanding Server-Level Compression Parameters," on page 331.

Command	Values
nss /DaysUntouchedBeforeCompression=*value*	**Range:** 0 to 100000 (in days) **Default:** 14
nss /CompressionDailyCheckStartingHour=*value*	**Range:** 0 to 23 Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.). **Default Value:** 0
nss /CompressionDailyCheckStopHour=*value*	**Supported Values:** 0 to 23 Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.). **Default Value:** 6
nss /EnableFileCompression=<On | Off>	**Supported Values:** On (default) or Off **Default Value:** On
nss /MinimumCompressionPercentageGain=*value*	**Supported Values:** 0 to 50 **Default Value:** 20
nss /MaximumConcurrentCompressions=*value*	**Supported Values:** 1 to 8 **Default Value:** 2

Command	Values
nss /ConvertCompressedToUncompressedOption=*value*	**Supported Values:** 0, 1, or 2 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 **Default Value:** 1
nss /DecompressPercentDiskSpaceFreeToAllowCommit=*value*	**Supported Values:** 0 to 75 **Default Value:** 10
nss /DecompressFreeSpaceWarningInterval=*value*	**Supported Values:** 0 seconds to 29 days 15 hours 50 minutes 3.8 seconds Setting the interval to 0 turns off the alert. **Default Value:** 31 minutes 18.5 seconds
nss /DeletedFilesCompressionOption=*value*	**Supported Values:** 0, 1, or 2 0 = Do not compress deleted files 1 = Compress deleted files the next day 2 = Compress deleted files immediately **Default Value:** 1

24.2.3 Configuring Server-Level Compression Parameters with 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 *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*.

24.2.4 Configuring Server-Level Compression Parameters with 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. For a list of Set parameters for server-level compression settings on NetWare, see Section B.22, "SET (NetWare)," on page 536.

Enter the commands using the following syntax at the NetWare 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.

On NetWare, you can use the NWCONFIG utility to edit both the `autoexec.ncf` and the `startup.ncf` file. For information, see " NWCONFIG " in the *OES 2 SP1: Utilities Reference for NetWare*.

24.2.5 Configuring Server-Level Compression Parameters with 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 Section 24.2.2, "Configuring Server-Level Compression Parameters with Commands," on page 334.

For information about using Monitor, see "MONITOR" in the *OES 2 SP1: Utilities Reference for NetWare*.

24.3 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. For more information, see Section 24.4, "Suspending Compression for Volumes or Files," on page 337.

IMPORTANT: You cannot use file compression on a volume on a CD or DVD drive.

- Section 24.3.1, "Enabling Compression for a New Volume," on page 336
- Section 24.3.2, "Enabling Compression for an Existing Volume," on page 336

24.3.1 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 19.3, "Creating Unencrypted NSS Volumes," on page 259.

24.3.2 Enabling Compression for an Existing Volume

1 In iManager, click *Storage* > *Volumes* to open the Volumes page.

For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2 Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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.

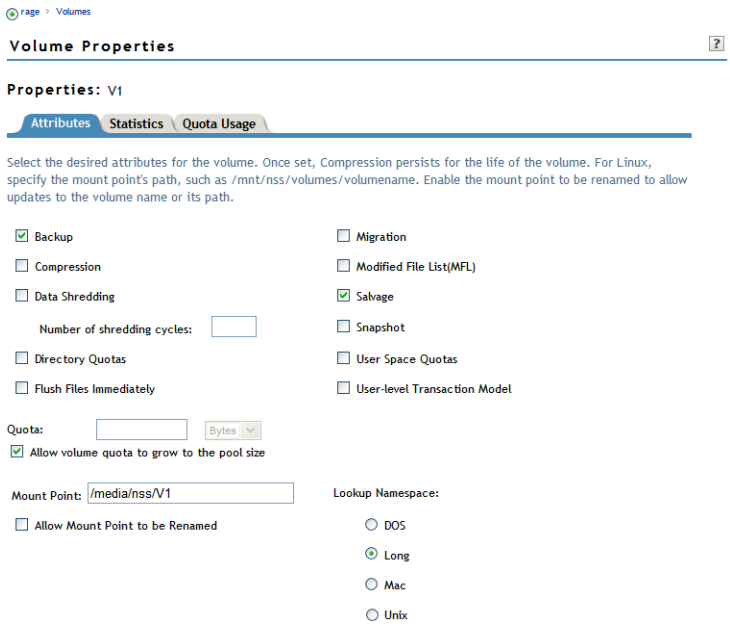

5 Select the *Compression* check box to enable file compression for the selected volume.

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 24.4, “Suspending Compression for Volumes or Files,” on page 337 or Section 24.6, “Restoring Data to a Uncompressed Volume,” on page 338.

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.

24.4 Suspending Compression for Volumes or Files

You can temporarily suspend file compression for all volumes by 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 24.7, “Configuring Compression Preferences for Directories and Files,” on page 338.

You can suspend compression for a file by setting its Don't Compress flag. For information, see Section 24.7, "Configuring Compression Preferences for Directories and Files," on page 338.

24.5 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 24.6, "Restoring Data to a Uncompressed Volume," on page 338.
- Set Enable File Compression to Off. Eventually, files are decompressed as they are used and remain uncompressed.

24.6 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.

24.7 Configuring Compression Preferences for Directories and Files

File and directory compression preferences override the compression settings for the volume and server.

- Section 24.7.1, "Using the Novell Client," on page 338
- Section 24.7.2, "Using ATTRIB (Linux)," on page 339
- Section 24.7.3, "Using FLAG (NetWare)," on page 340

24.7.1 Using the Novell Client

1. From a workstation, click the Novell Client™ icon (the red N in the notification area), select *Novell Map Network Drive*, then map a drive to the NSS volume by using the login and password of the Administrator user.
2. Use either of the following methods to open the *NetWare Info* dialog box 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 *Information* 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.

24.7.2 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.

Table 24-2 *Operations Options for the ATTRIB Utility*

Option	Description
-s, --set=ATTRIBUTES	Sets the attributes on the specified file or directory.
-c, --clear=[ATTRIBUTES | all]	Clears the attributes on the specified file or directory.

Compression Attributes Options

Table 24-3 *Compression Attributes Options for the ATTRIB Utility*

Attribute	Description	Applies to Files	Applies to Directories
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
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

For example, to set the Don't Compress attribute for all files in the current directory, enter

```
attrib --set=dc
```

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.

Table 24-4 *Compression Status for the ATTRIB Utility*

Status	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
```

24.7.3 Using FLAG (NetWare)

Use the FLAG utility (`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 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 `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.

Table 24-5 *Compression Attributes Options for the FLAG Utility*

Parameter	Description
(no parameter)	If you type `flag` with no parameters, the attribute status of all the files in the current directory is displayed.
path	Specifies the path to the file or directory whose attributes or search modes you want to modify.
+	Sets 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.
–	Clears 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.
attribute	Specifies the compression-related attribute, such as Immediate Compression (Ic) or Don't Compress (Dc).
/option	Replaces 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 *OES 2 SP1: Utilities Reference for NetWare*.

Configuring Compression Attributes for a Directory

Use the Flag command to set or clear the Immediate Compression attribute for a specified directory.

Table 24-6 *Directory Compression Option for the FLAG Utility*

Attribute	Description
Ic (Immediate compress)	Compresses 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.

To set the Immediate Compression attribute for the *sys:\course\winter* directory, enter

```
flag sys:\course\winter +Ic
```

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.

Table 24-7 *File Compression Options for the FLAG Utility*

Attribute	Description
Dc (Don't Compress)	Prevents a specified file from being compressed.
Ic (Immediate compress)	Compresses a file as soon as the operating system can.

To set the Don't Compress attribute to all files in the current directory, enter

```
flag *.* +Dc
```

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.

Table 24-8 *Compression Status for the FLAG Utility*

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
```

24.8 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`.

Table 24-9 *Compression Management Commands*

Command	Description
`nss /Compression=volume_name`	Enables the Compression attribute for the specified volume.
`nss /Compression=all`	Enables the Compression attribute for all volumes on the server.
`nss /StopNormalCompression`	Stops all queued compression for files, based on the compression triggered by a file open or close.
`nss /BGCompression`	Allows compression to occur in the background at any time, instead of only within specified hours.
`nss /NoBGCompression`	Stops background compression. Allow compression to occur only within the specified hours.
`nss /CompScreen`	Displays the NSS volume compression statistics on the compression screen.

24.9 Copying Compressed Files between Volumes with the NCOPY Command

The NCOPY (`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 B.11, “NCOPY (NetWare),” on page 518.

24.10 Viewing Compression Statistics with the NDIR Command

The NDIR (`ndir`) command allows you to view file attributes and compression information for all files in a volume or directory. You can use NDIR for NSS volumes on NetWare and for NSS volumes on Linux when NCP Server is running. NDIR is an MS-DOS* command that you run from a Windows desktop where you have mapped a drive letter to the network volume (or directory in a volume) that you want to manage.

Two useful statistics that NDIR provides are *Space Used By Compressed Files* and *Space Used If Files Were Not Compressed*. Subtracting one value from the other indicates how much space you need if all files are decompressed.

The following is an example report:

```
M:\>ndir /comp /vol
Statistics for fixed volume SRV1:
 Space statistics are in KB (1024 bytes).
 Total volume space: 979,628,112 100.00%
 Space used by 131,072,000 entries: 223,924,804 22.86%
 Deleted space not yet purgeable: 0 0.00%
 ------------ --------
 Space remaining on volume: 755,703,308 77.14%
```

```
Space available to <username>: 755,703,308 77.14%
Maximum directory entries: 2,147,483,647
Available directory entries: 2,016,411,647 93.90%
Space used if files were not compressed: 0
Space used by compressed files: 0
-------------
Space saved by compressing files: 0 0.00%
Uncompressed space used: 223,924,804
Name spaces loaded: Long
```

To generate a report on compression statistics for the volume:

1. For NSS on Linux volumes, make sure that NCP™ Server is configured and running.

   The NDIR command requires the NCP protocol.
2. On your Windows desktop, open Explorer, then click *Tools* > *Map Network Drive* to map a drive letter to the NSS volume (or directory) you want to manage.
3. On your Windows desktop, click *Start* > *Run* to open an MS-DOS command window.
4. Navigate to the mapped drive using the Change Directory (`cd`) command. For example, if the mapped drive letter is `N:`, enter

   ```
   cd N:
   ```
5. Get a report on compression statistics by entering

   ```
   ndir /comp /vol
   ```

   For details about NDIR, see Section B.12, “NDIR,” on page 518.

24.11 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 B.4, “COMPFIX,” on page 502.

24.12 Backing Up Compressed Files

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 “Backing Up Compressed files” in the *OES 2 SP1: Storage Management Services Administration Guide*, or see your third-party backup application’s documentation.

Managing Space Quotas for Volumes, Directories, and Users

25

This section describes how to manage space quotas for volumes, directories, and users of Novell® Storage Services™ (NSS) volumes.

- Section 25.1, “Understanding Space Quotas,” on page 345
- Section 25.2, “Managing NSS Volume Quotas,” on page 346
- Section 25.3, “Managing Directory Quotas,” on page 348
- Section 25.4, “Managing User Space Quotas,” on page 353

25.1 Understanding Space Quotas

You can control how space is allocated in an NSS pool or 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 quota does not reserve the space.

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.

Quotas restrict the actual physical space that the volume, directory, or user is allowed to consume. When enforcing quotas, NSS considers only the actual physical blocks consumed (in 4 KB blocks), not the file’s logical size. If you have sparse files or compressed files, only the actual physical space they consume is counted against the quota. In order for a compressed file to be uncompressed, there must be enough space available in the most restrictive of the quotas set (whether volume, directory, or user) to accommodate the decompression process and the uncompressed file size. Otherwise, the user is not able to open the file. For more information about the space requirements for compressing and decompressing files, see Section 24.1.5, “Factors Affecting Decompression,” on page 330.

As the amount of space consumed by a user’s files approaches the user’s space quota, the user 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.

WARNING: If storing a file would cause a quota to be exceeded, only part of the file is actually saved, resulting in data corruption.

If the Salvage attribute is enabled for a volume, deleted files are not immediately purged from the volume. Deleted files on the volume are not counted against quotas.

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 25.2, "Managing NSS Volume Quotas," on page 346.

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 32.3, "Monitoring Quotas and Space Usage for NSS Pools and Volumes," on page 424.

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.

For information about configuring attributes when you create a volume, see "Understanding Volume Properties" on page 253.

For information about setting directory quotas, see Section 25.3, "Managing Directory Quotas," on page 348

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. For information about setting user space quotas, see Section 25.4.3, "Configuring a User Space Quota," on page 356.

You must first enable the User Space Quotas attribute on the NSS volume where you want to configure user space restrictions. You can set the attribute at any time. For information about configuring the User Space Quotas attribute for an existing volume, see "Modifying Attributes of an NSS Volume" on page 265.

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.

25.2 Managing NSS Volume Quotas

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

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 *Quota Usage*. It opens to the *Attributes* tab.

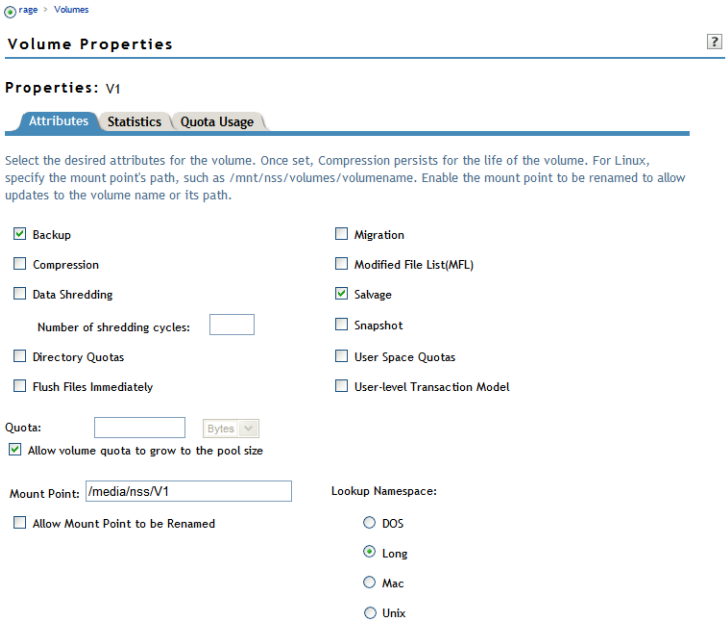

5 In the *Quota* field, do one of the following:

- **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.

  If you set the quota to a value less than the current volume size, you cannot save files to the volume until you purge some files to make room.

6 Click *Apply*.

7 Click the *Quota Usage* tab to view the volume and pool space usage for the selected volume and to verify the new setting.

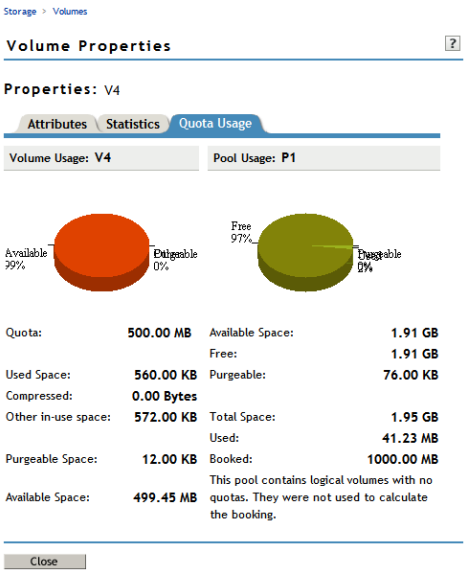

25.3 Managing Directory Quotas

A directory quota limits the amount of space on a volume that can be consumed by all of the files and folders in that directory. You can specify any postive value for the quota. If the current size of the directory exceeds the specified limit, users cannot save data to the directory until space is cleared by removing files from the directory. If the specified directory quota exceeds the volume quota, the volume quota overrides the directory quota in determining whether data can be saved to the directory as data is written. Because overbooking is allowed for directory and volume quotas, the physical space limits for the pool might also prevent the directory from growing to its specified maximum. Physical space limits for the pool override the volume and directory quotas.

For example, lets assume that volume `VOL1:` has a volume quota of 50 GB in pool `POOL1`. You set the directory quota to 100 GB for directory `VOL1:\finance\sales`. Users cannot save 100 GB of files in the `sales` directory because NSS reaches the volume quota long before it reaches the directory quota. Because of overbooking, other directories are competing for the 50 GB of space in the volume, and `VOL1:` is competing for space with other volumes in `POOL1`. You might not be able to put as much as 50 GB of data in the `sales` directory.

Before you can set directory quotas, you must enable the volume's Directory Quotas attribute. As the administrator user, you can view and configure directory quotas with the Files and Folders plug-in for iManager, NetStorage, and the Novell Client™. For NSS on NetWare, you can also use Novell Remote Manager for NetWare.

This section describes the following:

- Section 25.3.1, "Enabling or Disabling the Directory Quotas Attribute for an NSS Volume," on page 349
- Section 25.3.2, "Configuring Directory Quotas," on page 349
- Section 25.3.3, "Removing a Directory Quota," on page 352
- Section 25.3.4, "Removing All Directory Quotas for an NSS Volume," on page 353

25.3.1 Enabling or Disabling the Directory Quotas Attribute for an NSS Volume

Before setting directory quotas on a volume, you must enable the Directory Quotas attribute for the volume. You can set the attribute at create time or at any time for an existing volume.

To set the Directory Quotas attribute for an existing volume:

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage to view a list of NSS volumes on the server.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Volumes* list, select a volume that you want manage.

   Wait for the volume details to be displayed before you continue.
4. Click *Properties*.

   The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. 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 the following at a terminal console prompt:

   ```
   /etc/init.d/ncp2nss restart
   ```

25.3.2 Configuring Directory Quotas

- "Adding or Modifying a Directory Quota with iManager" on page 349
- "Adding or Modifying a Directory Quota with Novell NetStorage" on page 350
- "Adding or Modifying Directory Quotas with the Novell Client" on page 350
- "Adding or Modifying a Directory Quota with Novell Remote Manager for NetWare" on page 352

Adding or Modifying a Directory Quota with iManager

1. In iManager, select *Files and Folders* > *Properties*.
2. Click the *Search* icon, then browse to locate and select the folder you want to manage on an NSS volume.
3. View the current status of the Directory Quota.

   If a Directory Quota is set, the *Restrict Size* field is selected and the *Limit* field shows the quota size in KB.

   If the Directory Quota is not set, the *Restrict Size* field is deselected and the *Limit* field is dimmed (grayed out).
4. Do one of the following:
   - **Add a Quota:** On the *Information* tab, select *Restrict Size* to enable space restrictions for the selected directory. In the *Limit* field, type the directory quota in KB. The value must be an increment of 4 KB; that is, it must be divisible by 4 with no remainder.

- **Modify an Existing Quota:** In the *Limit* field, type the new directory quota in KB. The value must be an increment of 4 KB; that is, it must be divisible by 4 with no remainder.
- **Remove a Quota:** On the *Information* tab, deselect *Restrict Size* to disable space restrictions for the selected directory. The *Limit* field is automatically dimmed (grayed out).

5 On the *Information* page, click *Apply* or *OK* to apply the changes.

Adding or Modifying a Directory Quota with 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 2 SP1: NetStorage for Linux Administration Guide*
- *OES 2: NetStorage for NetWare Administration Guide*

To create or modify NSS directory quotas with NetStorage:

1 In a Web browser, connect to NetStorage.

For information, see Section 9.4, “Novell NetStorage,” on page 118.

2 Log in to NetStorage with the username and password of the Administrator user or equivalent 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.

Adding or Modifying Directory Quotas with the Novell Client

The Novell Client for Windows 2000/XP allows the Administrator 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 red N icon in the notification area), then select *Novell Map Network Drive*.

1b Specify the network path to the directory. For example: `192.168.1.1/users`.

1c Specify the username of the Administrator user or equivalent user, then click *Map*.

1d When prompted, enter the user’s 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.

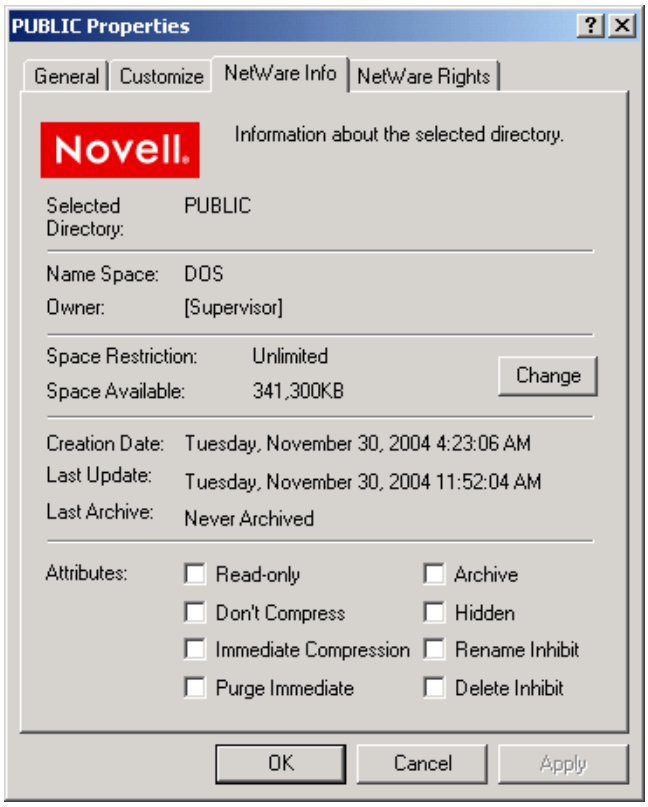

3 In the *Space Restriction* field, click *Change* to open the *Space Restriction* dialog box.

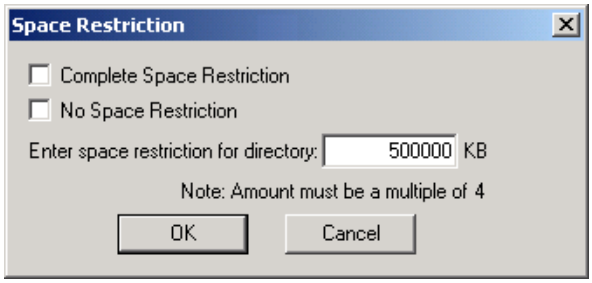

4 Do one of the following to configure the directory quota:

- **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.

Adding or Modifying a Directory Quota with Novell Remote Manager for NetWare

1. Access Novell Remote Manager for your NetWare server.

   For information, see Section 9.5, “Novell Remote Manager,” on page 120.
2. Click *Volumes* to open the *Volumes* 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 Entry Information* page.
5. In the *Directory Entry Information* field, click the *File Space Limit* link.

Figure 25-1 *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	SRWCEMFA
File space limit	None
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.

25.3.3 Removing a Directory Quota

1. In iManager, select *Files and Folders* > *Properties*.
2. Click the *Search* icon, then browse to locate and select the folder you want to manage on an NSS volume.
3. On the *Information* tab, deselect *Restrict Size* to disable space restrictions for the selected folder.
4. Click *Apply* or *OK* to apply the changes.

25.3.4 Removing 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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

   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 *Quota Usage*. 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 the following at a terminal prompt:

   ```
   /etc/init.d/ncp2nss restart
   ```

25.4 Managing User Space Quotas

User space quotas are the space restrictions you optionally set for users of an NSS volume where the User Space Quotas attribute is enabled. The Users with Quotas page of the Storage plug-in reports the quota setting and space usage for each user who has space restrictions in place for a specified volume, whether the user has data stored on the volume or not.

This section describes the following:

- Section 25.4.1, "Setting the User Space Quotas Attribute for an NSS Volume," on page 353
- Section 25.4.2, "Viewing User Space Quotas," on page 354
- Section 25.4.3, "Configuring a User Space Quota," on page 356
- Section 25.4.4, "Modifying a User Space Quota," on page 357
- Section 25.4.5, "Deleting a User Space Quota," on page 358
- Section 25.4.6, "Configuring User Space Quotas on Volumes After Upgrading or Migrating from OES 1 Linux (Linux)," on page 358

25.4.1 Setting the User Space Quotas Attribute for an NSS Volume

The administrator user can view and modify the User Space Quotas attribute at any time.

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage to view a list of its volumes.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.

3 In the *Volumes* list, select a volume that you want manage, then wait for the page to refresh to view its details.

4 Click *Properties*.

The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. 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*.

25.4.2 Viewing User Space Quotas

As the Administrator user or a user with the Supervisor right for the volume can view and configure user space quotas with the Storage plug-in for iManager. Users can view their own user quotas and space usage statistics, but they cannot modify the quota settings.

The *Users with Quotas* page reports the quota setting and space usage for each user who has space restrictions in place for a specified volume, whether the user has data stored on the volume or not. The *All Users* page reports the current usage of all users on the volume who have data stored on the volume, whether the users have a quota on the volume or not.

The tracking of user space usage and quotas is an expensive operation in terms of performance. For this reason, NSS does not begin tracking user space usage until you create the first user quota on the volume. If you have never assigned a user quota to a volume, the All Users page has no information to report. After you create the first user quota on the volume, NSS begins tracking all of the user space used on the volume. From then on, the All Users page reports usage for all users with data on the volume.

The *Users with Quotas* report and *All Users* report include the following information:

Table 25-1 *Report of User Space Quotas*

Parameter	Description
User Name	The distinguished user name. For example: userid.context, jsmith.geo.example, asantiago.example
Quota	Indicates the amount of space in MB that the user can use on the selected volume. The user's data can grow only to the size of the quota or to the amount of available physical space on the volume, whichever is less. NSS allocates space only as needed. You can set user quotas that total more than the total available space on the volume. This is called overbooking. However, you cannot set any single quota to be greater than the volume size.
Used	Reports the current amount of space used by the user's data in all directories combined on the specified volume. Reports the current amount of space used by the user's data in all directories combined on the specified volume.
Available	Reports the free, unused space within the user's quota for the specified volume. Because of overbooking, other users might be competing for a portion of this space.

Viewing User Quotas as the Administrator User

1 Log in to iManager as the Administrator user or equivalent user.

2 In iManager, click *Storage* > *User Quotas*.

3 Click the *Volume* browser, then select the volume that you want to manage.

Wait for the page to refresh to see the user space restrictions for all users with quotas for the selected volume.

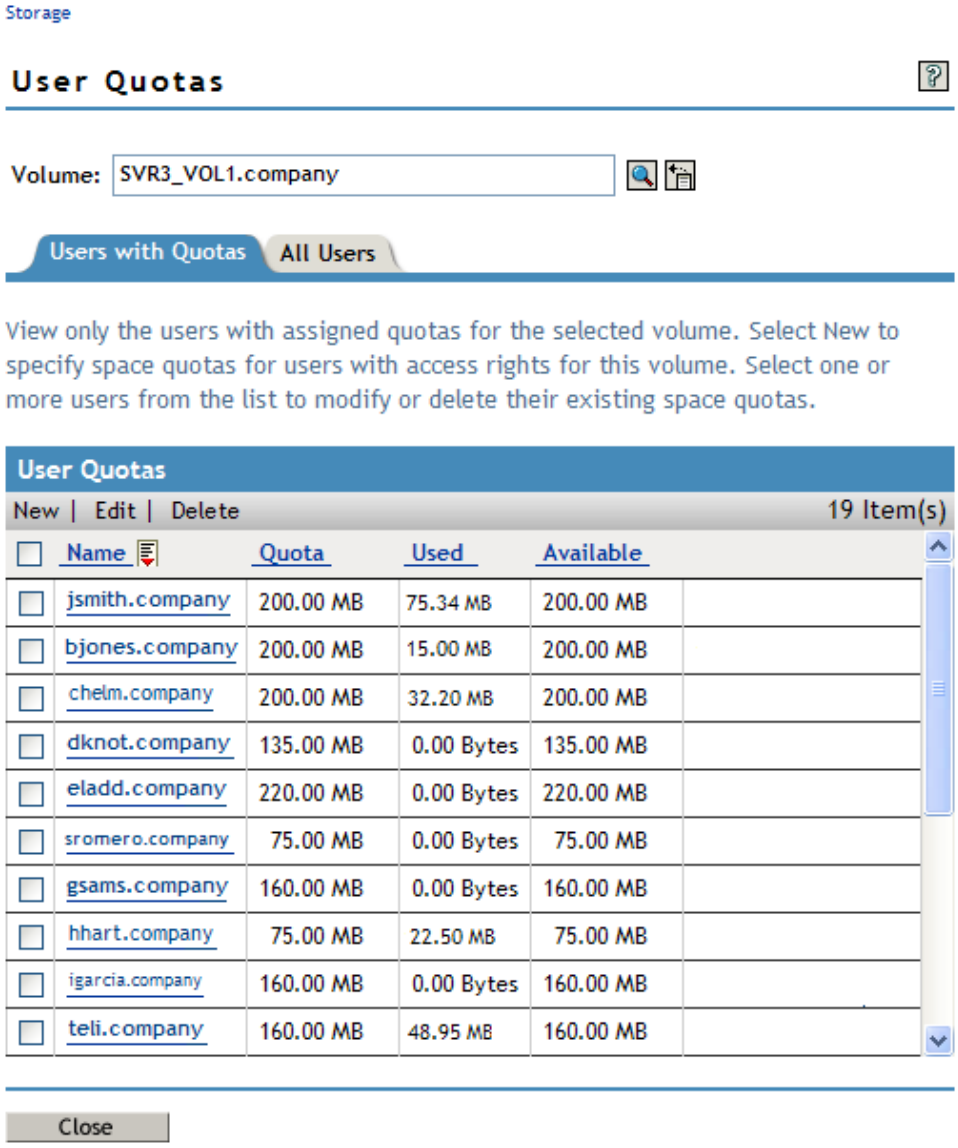

4 To view all users (with or without quotas) who have data stored on that volume, click *All Users*.

Viewing User Quotas as an Individual User

1 Log in to iManager as an individual user with your eDirectory username and password.

2 In iManager, click *Storage* > *User Quotas*.

3 Click the Volume browser, then select the volume that you want to manage.

Wait for the page to refresh to see the individual user's space restrictions for the selected volume on the *Users with Quotas* page.

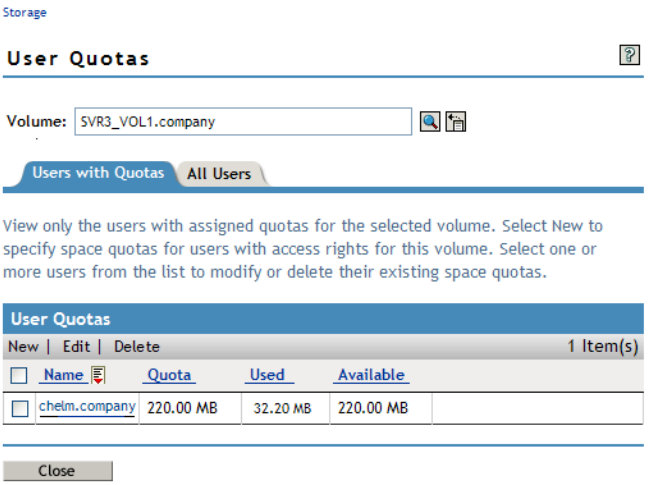

25.4.3 Configuring a User Space Quota

The Administrator user or equivalent user can configure the user space quota for one user or multiple users at a time.

1. In iManager, click *Storage* > *User Quotas*, then select the volume you want to manage.
2. Click *New* to open the *Add User Quota* dialog box.

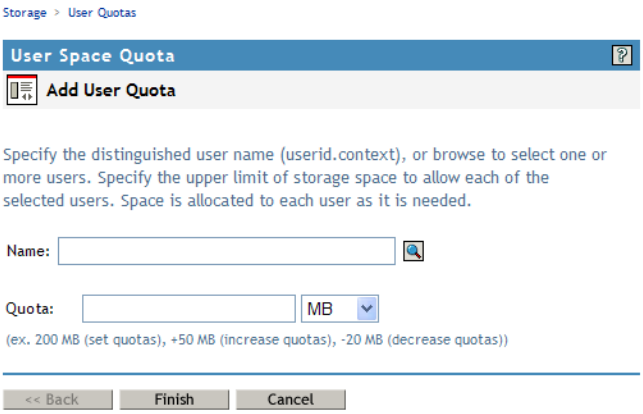

3. Click the *Browse* icon to open the User object list, then browse to select one or more users who you want to share the same quota.
4. In *User Quotas*, specify the upper limit of storage space that you want to assign to each of the selected users.
5. Use the *Units* drop-down list to select the units to use for the quota you specified in Step 4. Options are bytes, kilobytes, megabytes, and gigabytes.
6. Click *Finish* to apply the quota, or click *Cancel* to back out of the transaction.

   Each of the usernames you selected now has user space quotas assigned to it.

25.4.4 Modifying a User Space Quota

The Administrator user or equivalent user can modify the quota for one user or multiple users at a time. Modify the current quotas for the selected users by setting a new quota for all users, increasing each quota by the same amount, or decreasing each quota by the same amount.

1. In iManager, click *Storage* > *Volumes*, select a server, then select the volume you want to manage.
2. Click *Quotas* > *Users with Quotas*.
3. Open the *Edit User Quotas* dialog box by doing one of the following:
   - **Single User:** Select the name link for the user.
   - **Multiple Users:** Select the *User* check box next to one or more user names whose user space quota you want to delete, then click *Edit*.

4. Specify the change that you want to apply to each of the selected users by doing one of the following:
   - **Modify Quota:** Specify a value to set each selected user's quota to the specified value, such as 200 MB.

- **Increase Quota:** Use a plus (+) with the value to increase each selected user's quota by that amount, such as +50 MB.

  If any individual quota or the total of all quotas exceeds the volume size, the increase is processed, because space is allocated to the users as needed, on a first-come-first-served basis.

  If you attempt to increase the quota for a user with an Unlimited quota, no changes are made for that user's quota. The setting remains Unlimited.

- **Decrease Quota:** Use a minus (-) with the value to decrease each selected user's quota by that amount, such as -20 MB.

  If a reduction takes a user's quota to 0 MB, then the user is fully restricted. The user cannot add any new files no matter how many existing files are deleted. To remove the restriction, set a non-zero quota or delete the quota for the user.

  If a reduction takes a user's quota to a negative value, then an error message reports the quota as an Invalid Quota. You must repeat the process to set a valid quota for the user.

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.

25.4.5 Deleting a User Space Quota

The Administrator user or equivalent user can delete the user space quota for one or more users. Deleting the user space quota for a user removes the space restriction for the user on the selected volume. Of course, any directory quotas or volume quotas still apply.

1 In iManager, click *Storage* > *User Quotas*, then select the volume you want to manage.

2 Select the *User* check box next to one or more user names whose user space quota you want to delete, then click *Delete*.

3 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*.

To remove all user space quotas on the volume at the same time, you can simply disable the User Space Quota attribute for the volume. For information about disabling the attribute, see Section 25.4.1, “Setting the User Space Quotas Attribute for an NSS Volume,” on page 353.

25.4.6 Configuring User Space Quotas on Volumes After Upgrading or Migrating from OES 1 Linux (Linux)

Beginning with OES 2 Linux, file ownership is properly assigned for access via the NCP protocol whether the user is Linux-enabled or not. User quotas for NCP users no longer require that a user be Linux-enabled.

For OES 1 Linux, using user space quotas required usernames to be Linux-enabled with Linux User Management (LUM). If an NCP user was not Linux-enabled on OES 1 Linux, any files the user created were assigned as the `root` user identity (after verifying the username's trustee rights to do so, of course) instead of the actual username.

If you migrate an NSS volume where NCP users were not Linux-enabled from OES 1 Linux to OES 2 Linux, you will observe that the file ownership is now recorded as the actual user. However, existing files might still belong to the `root` user and are not counted against user space quotas.

IMPORTANT: Directories and files that were created while users were not Linux-enabled on the OES 1 Linux server are owned by the `root` user and are not counted against the user space quotas you set unless you reassign file ownership to the individual users.

---

1. Use one of the following tools to change ownership of any files that the user normally uses that were previously assigned to the `root` user as owner:
   - Novell Client
   - Files and Folders plug-in to iManager

You cannot use the Linux `chown` command to change the creator field for the NSS file system. It changes the root user's view of who is reported as the owner user in the Linux path, but the change has no effect on the NSS metadata. The Linux `chown` command also does not modify the owner group.

Salvaging and Purging Deleted Volumes, Directories, and Files

26

This section describes how to configure the salvage system for Novell® Storage Services™ file systems on Novell Open Enterprise Server 2.

- Section 26.1, “Understanding the NSS Salvage System,” on page 361
- Section 26.2, “Configuring the Purge Behavior for NSS,” on page 364
- Section 26.3, “Enabling or Disabling the Salvage Attribute for an NSS Volume,” on page 367
- Section 26.4, “Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool,” on page 368
- Section 26.5, “Salvaging or Purging Deleted Files with iManager,” on page 369
- Section 26.6, “Salvaging or Purging Deleted Files with Other Tools,” on page 371
- Section 26.7, “Purging Deleted Files or Directories (NetWare),” on page 373

26.1 Understanding the NSS Salvage System

- Section 26.1.1, “Volume Salvage versus File Salvage,” on page 361
- Section 26.1.2, “Trustees and Rights Handling for Salvaged Data,” on page 363
- Section 26.1.3, “Understanding Purging Triggers,” on page 363
- Section 26.1.4, “Comparison of Salvage for NSS Volumes versus NetWare Traditional Volumes,” on page 364

26.1.1 Volume Salvage versus File Salvage

The NSS salvage system makes it possible to retain deleted files for a specified period of time or until space is needed. The volume salvage and file salvage subsystems function separately.

For volume salvage, the NSS volumes are automatically retained on deletion. The deleted volume can be salvaged for a period of time that is determined by the server-level Logical Volume Purge Delay setting. Administrators with the Supervisor right can salvage or purge deleted volumes at any time before the purge delay elapses.

Table 26-1 describes parameters that control the volume salvage behavior for NSS volumes. On Linux, server-level settings apply to directories and files on all NSS volumes. On NetWare, server-level settings apply to directories and files on all NSS and NetWare traditional volumes.

Table 26-1 *Volume Salvage Parameters*

Salvage Policy	Range of Influence	Refer to
Logical Volume Purge Delay	Server-level Default: 345600 seconds (4 days) The automatic purging delay applies to deleted NSS volumes.	Section 26.2.1, "Setting the Purge Delay for All Deleted Volumes," on page 364
Logical Volume Purge Delay After Continue	Server-level Default: 900 seconds Seconds to delay purging a deleted volume after a continue.	Section 26.2.1, "Setting the Purge Delay for All Deleted Volumes," on page 364
Logical Volume Purge Delay After Load	Server-level Default: 7200 seconds Seconds to delay purging a deleted volume after an NSS load if the purge delay elapses while NSS is disabled.	Section 26.2.1, "Setting the Purge Delay for All Deleted Volumes," on page 364

Salvage for directories and files is controlled by each volume's Salvage attribute. You can enable the Salvage attribute when you create the volume, or modify the setting later in the volume's properties. Deleted directories and files are retained and can be salvaged until space is needed in the pool where the volume resides, as determined by the pool's available-space watermark settings. The administrator user or any user who is a trustee with the Create right can salvage deleted directories and files.

Table 26-2 describes parameters that control the file-system salvage behavior for NSS volumes.

Table 26-2 *Files Salvage Parameters*

Salvage Policy	Range of Influence	Refer to
Salvage attribute	Volume-level Default: Enabled	Section 26.3, "Enabling or Disabling the Salvage Attribute for an NSS Volume," on page 367
Immediate Purge of Deleted Files	Server-level Default: Disabled	Section 26.2.2, "Setting the Immediate Purge of Deleted Files for All NSS Volumes," on page 365

Salvage Policy	Range of Influence	Refer to
Low and High Watermarks	Pool-level Default: low 10%; high 20% Volume-level watermarks are not available.	Section 26.2.3, "Setting the Low and High Salvage Watermarks for Automatically Purging Deleted Directories and Files," on page 365
Purge Immediate file system attribute	Individual directory or file where the attribute is enabled.	Section 26.2.4, "Setting the Purge Immediate Attribute for a Directory or File," on page 367

26.1.2 Trustees and Rights Handling for Salvaged Data

When you salvage a volume, the data and metadata is exactly the same as it was at delete time, with no changes.When salvaging deleted directories or files, the content, trustees, trustee rights, and the inherited rights filters are just as they were before the file was deleted. If the rights in the tree above the salvaged file have changed, then the inherited rights for the salvaged deleted file are calculated based on the current rights above it in the directory tree.

26.1.3 Understanding Purging Triggers

Purging is triggered to begin by the following events. After the deleted data enters a Purge state by manually starting a purge or by autopurging, deleted files can no longer be salvaged (do not return to a Salvageable state), even if you pause the autopurging process.

- The Logical Volume Purge Delay setting times out for a deleted volume. Autopurging begins automatically and can take some time, depending on the size of the volume.

  The elapsed time between the delete and the purge is called the purge delay. The server-level LogicalVolumePurgeDelay parameter applies to all NSS volumes. For information, see Section 26.2.1, "Setting the Purge Delay for All Deleted Volumes," on page 364.

- The pool's Low Salvage Watermark setting is reached, indicating that the amount of free space is below the administrator-specified minimum. NSS automatically purges the deleted files and directories for all volumes in the pool with a first deleted, first purged policy until the free space reaches the high watermark, or until all of the existing deleted directories and files are purged, whichever occurs first.

  You configure thresholds for space-based purging on each pool. Low and high watermarks determine when to begin and stop automatic purging of deleted files to free up space on the pool. For information, see Section 26.2.3, "Setting the Low and High Salvage Watermarks for Automatically Purging Deleted Directories and Files," on page 365.

- A user or administrator purges the deleted file or directory.

  Deleted directories and files can be purged by the administrator user or by any user who is a trustee with the Erase right at any time before the automatic purge begins. For information, see Section 26.5.3, "Purging Deleted Files," on page 371.

- An administrator purges the deleted volume.

  Delete volumes can be purged by the administrator with Supervisor right to the volume. For information, see Section 26.5.3, "Purging Deleted Files," on page 371.

- Salvage is disabled at the server level for directories and files.

Directories and files are purged immediately on deletion if the server-level ImmediatePurgeOfDeletedFiles parameter is enabled. If the Salvage attribute is enabled for a volume, this setting overrides it. For information, see Section 26.2.2, “Setting the Immediate Purge of Deleted Files for All NSS Volumes,” on page 365.

- Salvage is disabled for the individual directory or file.

  If the Salvage attribute is enabled for an NSS volume, you can set the PurgeImmediate file-system attribute on individual directories and files so that they are purged immediately on deletion. For information, see Section 26.2.4, “Setting the Purge Immediate Attribute for a Directory or File,” on page 367.

26.1.4 Comparison of Salvage for NSS Volumes versus NetWare Traditional Volumes

It is important to understand the differences between the salvage systems for NSS volumes as compared to NetWare® traditional volumes if you are upgrading your media to the NSS file system.

For NSS volumes, when you delete a volume, a directory (and its contents), or a file, the metadata for the deleted item is marked as deleted, but its metadata and data remain in the same physical location until it is purged or salvaged. This allows the deleted file or directory to be salvaged to the directory where it existed when it was deleted. You can also salvage a deleted directory to make its deleted files and subdirectories available for salvage.

For NetWare Traditional volumes, deleted files and directories are copied to a separate storage area, unlike the method used for NSS volumes. When you delete a file on a Traditional NetWare volume, the deleted file is moved into the `deleted.sav` directory, which is known as the *file salvage area*. No record is kept of where the file originated, so salvaging files by their original directory is not possible.

When you migrate files from a Traditional volume to an NSS volume, the `deleted.sav` directory and its contents are not migrated; thus, they are essentially purged.

NOTE: If you upgrade a Traditional NetWare volume to an NSS volume, make sure to salvage any deleted files you want to keep before the migration.

26.2 Configuring the Purge Behavior for NSS

- Section 26.2.1, “Setting the Purge Delay for All Deleted Volumes,” on page 364
- Section 26.2.2, “Setting the Immediate Purge of Deleted Files for All NSS Volumes,” on page 365
- Section 26.2.3, “Setting the Low and High Salvage Watermarks for Automatically Purging Deleted Directories and Files,” on page 365
- Section 26.2.4, “Setting the Purge Immediate Attribute for a Directory or File,” on page 367

26.2.1 Setting the Purge Delay for All Deleted Volumes

The Purge Delay setting for the NSS volume determines the amount of time (in seconds) that you can still access the deleted volume before it is removed from the system. The default value for the Purge Delay setting is 345600 (4 days). The volume name is changed during delete so that a new volume with the same name can be immediately created. The management tool used to delete the

volume should clean up any Novell eDirectory™ Storage objects at delete time. Use NSSMU or the Storage plug-in to iManager to purge or salvage the deleted volume before the Purge Delay time elapses.

To configure the Purge Delay time, enter the following at the server console:

```
nss /LogicalVolumePurgeDelay=value
```

In this command, replace *value* with the actual number of seconds to delay the purge. For example, if you want to change the Purge Delay time from the default of 4 days to 1 day, set the value to 86400 by entering:

```
nss /LogicalVolumePurgeDelay=86400
```

The Purge Delay change command is not permanent when entered from the command line. You must enter the command each time you restart the server. To make the new setting permanent, add the command to the `autoexec.ncf` file.

26.2.2 Setting the Immediate Purge of Deleted Files for All NSS Volumes

The salvage capability for directories and files can be turned on and off for NSS at the server level by using the /(No)ImmediatePurgeOfDeletedFiles flag. By default, the setting is disabled (set to NoImmediatePurgeOfDeletedFiles). You might want to enable this setting if you have Salvage enabled for multiple volumes, but want to disable salvage across all of them without separately changing the volumes' Salvage attribute settings.

This server-level salvage setting overrides the settings for the volume-level Salvage attribute. It does not affect deleted NSS volumes.

Table 26-3 *Server-Level Salvage Parameter*

Parameter	Description
`nss /NoImmediatePurgeOfDeletedFiles`	**Supported Values:** Enabled or disabled (default)
`nss /ImmediatePurgeOfDeletedFiles`	If this parameter is enabled, it overrides the Salvage attribute setting for individual NSS volumes so that the directories and files are immediately purged on deletion.

26.2.3 Setting the Low and High Salvage Watermarks for Automatically Purging Deleted Directories and Files

Salvage watermarks are parameters associated with the salvage feature of NSS pools. Automatic purging of deleted files begins when the pool's low watermark is reached and continues until its high watermark is reached, or until all existing deleted files and directories have been purged. Files and directories are purged on a first-deleted, first-purged bases. If you have deleted volumes in the pool, any deleted files and directories that existed when the volume was deleted are considered for this space-based purge, but the deleted volume is not.

When free space in the pool falls below a low watermark, NSS begins autopurging the deleted files. When enough files are purged so that the free space increases to a percentage equal to or greater than the high watermark, the autopurge stops. The autopurge also stops before the high watermark is reached if there are no more deleted files or directories to be purged. Autopurging does not start again until free space again drops below the low watermark.

Figure 26-1 *How Autopurging Works*

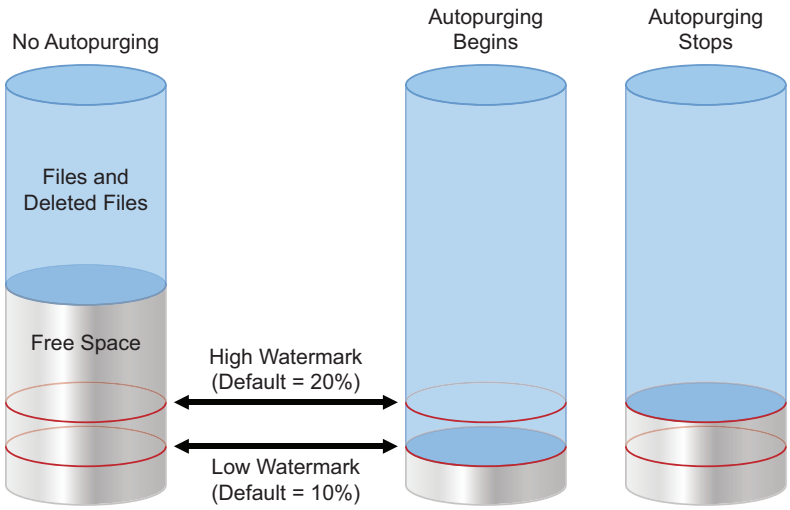

The high and low watermarks can be configured only at a pool level. The default low watermark is 10% of the maximum pool size. The default high watermark is 20% of the maximum pool size. The low watermark's percentage can range from a minimum of 0% to a maximum of 98%. The high watermark's percentage 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.

IMPORTANT: The pool's watermark settings are enforced only at the pool level. Volume-level watermarks are not supported. The Storage plug-in to iManager displays default values for volume watermarks, but they have no effect. In order for a volume to benefit from the watermark protection, set the volume's size to grow to the size of the pool.

At the minimum setting of 0%, the low watermark activates the autopurge only when the pool is totally out of free space. If the watermark is set this low, users are likely to get out-of-space errors when they try to save files. Setting the low watermark to a percentage a little higher than 0% guarantees that autopurging begins before free space is completely used, and users are less likely to get out-of-space errors.

The low and high salvage watermarks for a pool are set on boot to the default levels. Commands to modify the watermarks can be issued from the command line or placed in a startup file. Settings in the startup file persist across server reboots. Commands issued at the command line persist until the next reboot, or until the command is issued again, whichever occurs first.

Use the following commands to configure the high and low watermarks for pools. For Linux, issue the commands in the NSS Console (`nsscon`) as the `root` user from within a terminal console. For NetWare, enter the commands at the server console.

```
nss /PoolHighWaterMark=poolname:Percent
```

```
nss /PoolLowWaterMark=poolname:Percent
```

Replace the *poolname* with the name of an individual pool or with All to set the value for all pools. 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
```

26.2.4 Setting the Purge Immediate Attribute for a Directory or File

The Purge Immediate file-system attribute flags a directory or file to be erased from the system as soon as it is deleted. Purged directories and files cannot be recovered. When this attribute is enabled, it overrides the salvage settings at the volume and server level. When it is disabled, the server and volume salvage settings apply. In order to modify this setting, you must be the administrator user or a user who is a trustee with the Erase right.

1. In iManager, select *Files and Folders* > *Properties*.
2. Click the *Search* icon to locate and select the directory or file you want to manage.
3. Click the *Information* tab to view the properties for the selected directory or file.
4. Scroll down to view the *Attributes* section, then select the *Purge Immediate* attribute to enable the selected directory or file to be purged immediately on deletion, or deselect it to allow the salvage settings to control the fate of the deleted directory or file.
5. Click *Apply* or *OK* to save your changes.

26.3 Enabling or Disabling the Salvage Attribute for an NSS Volume

You can enable the Salvage attribute when you create the volume, or modify the setting later in the volume's properties.

1. In iManager, click *Storage* > *Volumes*.

   For instructions, see Section 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Volumes* list, select a volume that you want manage.
4. Click *Properties*.

   The Properties page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. Use the Attributes page to view or modify the attributes for the selected volume.
5. Select (enable) or deselect (disable) the Salvage attribute, then click *Apply*.

26.4 Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool

If you delete a volume, NSS removes it from the pool. During the Purge Delay (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. When you salvage a volume, the data and metadata are exactly the same as they was at delete time, with no changes. After the Purge Delay time elapses, NSS automatically purges the deleted volume from the system and you can no longer access it. You can also manually purge any volumes you have deleted.

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.

You can change the Purge Delay time to extend or reduce the time for the automatic purging cycle. For information, see Section 26.2.2, “Setting the Immediate Purge of Deleted Files for All NSS Volumes,” on page 365.

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.

To manage the deleted volumes in a pool:

1. In iManager, click *Storage* > *Pools*.

   For instructions, see Section 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.

2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.

   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.

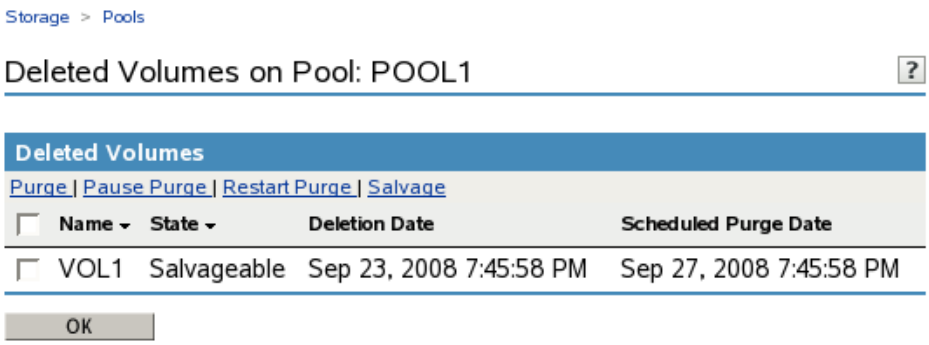

5. Select a volume from the *Deleted Volumes* list.

   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 deletion states are 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. After this option is selected, the deleted volume can no longer be salvaged (does not return to a Salvageable state). You cannot stop the purging by clicking *Pause Purge* for the autopurging process; the two options are unrelated.
- **Pause Purge/Restart Purge:** *Pause Purge* disables autopurging for the selected volumes so that purging does not begin automatically for a deleted volume when its Purge Delay time elapses. During the purge delay or while autopurging is disabled, the deleted volume is salvageable.

  Restart Purge enables autopurging for the selected volumes. The deleted volume is purged when the purge delay time elapses.

  This option does not make it possible to salvage a deleted volume that has already entered a Purge state.
- **Salvage/Pause Salvage:** Begins the restore process for deleted volumes you selected.

  If you salvage a volume, you can assign a new name to that volume, or reuse the old one if no current volume is using that name. 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.

  After you start the salvage process, the *Salvage* button changes to *Pause Salvage.* Click *Pause Salvage* to pause the salvage process temporarily; click *Salvage* again to resume.
- **Close:** Closes the Deleted Volumes page.

26.5 Salvaging or Purging Deleted Files with iManager

As an administrator, you can use the Files and Folders plug-in to iManager to salvage or purge deleted files from an NSS volume where the Salvage attribute is enabled. When salvaging deleted files, the file content, trustees, trustee rights, and inherited rights filter are just as they were before the file was deleted. If the rights in the tree above the salvaged file have changed, then the inherited rights for the salvaged deleted file are calculated based on the current rights above it in the directory tree.

- Section 26.5.1, “Prerequisites,” on page 369
- Section 26.5.2, “Salvaging a Deleted File,” on page 370
- Section 26.5.3, “Purging Deleted Files,” on page 371

26.5.1 Prerequisites

- The NSS volume that you want to manage must be in the same tree where you are currently logged in.

- You must have trustee rights for the file that you want to manage.
- The NSS volume must be configured for salvage at the time the files, directories, or volumes are deleted in order for deleted files to be available. Enable the Salvage attribute by going to the volume's Attributes page (*Storage* > *Volumes* > *Properties* > *Attributes*), select *Salvage*, then click *OK*.
- Deleted files are typically purged according to the Purge Delay settings on the server. When the delay time elapses, the deleted data is no longer available for salvage.
- Deleted files can be salvaged by any trustee for the file with the Create right. If another user has purged the deleted file, it is no longer available for salvage.
- Deleted files can be purged by any trustee for the file with the Erase right. If another user has purged the deleted file, it is no longer available for salvage.
- If the Purge Immediate attribute is set for a file or folder, the file or folder is immediately and permanently removed from the volume upon deletion.

26.5.2 Salvaging a Deleted File

You can salvage a deleted file and restore it to the directory from which it was deleted if you are a trustee of the file with the Create write. You can choose to overwrite any existing copies of the file in that location, or to rename the deleted file before it is salvaged. Review the guidelines in Section 26.5.1, "Prerequisites," on page 369 to understand when deleted files are available for salvage.

1. In iManager, click *Files and Folders*, then click *Deleted File* to open the Deleted File page.
2. On the Deleted File page, use one of the following methods to locate the folder on an NSS volume where the deleted file existed when it was deleted:
   - Click the *Search* icon to browse and locate the folder, then click the name link of the folder to select it.
   - Click the *History* icon to select a folder from the list of folders that you recently accessed.

   The *Deleted Files* report lists the deleted files in the folder and shows who deleted each file and when it was deleted.
3. Browse the list of deleted files to locate the version of the file you want to salvage.
4. Select the deleted file that you want to salvage, then click *Salvage*.
5. If a current file in the folder is named the same as the salvaged file, you are prompted to do one of the following:
   - Type a new name for the salvaged file, then click *OK*.
   - Click *OK* to overwrite the current file with the salvaged file.

   A confirmation message confirms that the file was successfully saved.
6. Click *Repeat Task* to salvage or purge other deleted files, or click *OK* to dismiss the confirmation message.

26.5.3 Purging Deleted Files

You can purge a deleted file to remove it immediately from the volume if you are a trustee of the file with the Erase right. Purged files can no longer be salvaged. Review the guidelines in Section 26.5.1, “Prerequisites,” on page 369 to understand when deleted files are available for purging.

1. In iManager, click *Files and Folders*, then click *Deleted File* to open the Deleted File page.
2. On the Deleted File page, use one of the following methods to locate the folder on an NSS volume where the deleted file existed when it was deleted:
   - Click the *Search* icon to browse and locate the folder, then click the name link of the folder to select it.
   - Click the *History* icon to select a folder from the list of folders that you recently accessed.

   The *Deleted Files* report lists the deleted files in the folder and shows who deleted each file and when it was deleted.
3. Browse the list of deleted files to locate the version of the file you want to purge.
4. Select one or more deleted files that you want to purge, then click *Purge*.

   A confirmation message confirms that the file was successfully purged.
5. Click *Repeat Task* to salvage or purge other deleted files, or click *OK* to dismiss the confirmation message.

26.6 Salvaging or Purging Deleted Files with Other Tools

You can use any of the following methods to salvage or purge deleted files. To purge, the user must be a trustee of the file with the Erase right. To salvage, the user must be a trustee of the file with the Create right.

- Section 26.6.1, “Using NetStorage,” on page 371
- Section 26.6.2, “Using the Novell Client,” on page 372
- Section 26.6.3, “Using Novell Remote Manager for NetWare,” on page 372

26.6.1 Using NetStorage

Using NetStorage, the Administrator user, the Admin-equivalent user, and individual 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 9.4, “Novell NetStorage,” on page 118.
2. In the left column, select the directory where the deleted files were located when they were deleted.
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*.

26.6.2 Using the Novell Client

Using the Novell Client™ for Windows 2000/XP/2003, Administrator users, Admin-equivalent users, and individual 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 more deleted files, select the deleted files, then click *Salvage File*.
   - To restore all deleted files in the directory, click *Salvage All*.
3. When you are done, click *OK*.
4. 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 more 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*.
5. When you are done, click *OK*.

26.6.3 Using Novell Remote Manager for NetWare

Using Novell Remote Manager for NetWare, the Administrator user or equivalent 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 9.5.4, "Accessing Novell Remote Manager," on page 123.
2. On the Volumes page, click the volume name link of the volume 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 File*s field, then click the *Select for List* link.
5. Do one of the following:
   - **Salvage:** Select the deleted file you want to restore, then click the *Salvage* button next to the filename.
   - **Purge:** Select 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.

26.7 Purging Deleted Files or Directories (NetWare)

If your NetWare 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 in Novell Remote Manager for NetWare.

- Section 26.7.1, "Purging All Deleted Files," on page 373
- Section 26.7.2, "Purging Specific Directories or Files," on page 373

26.7.1 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. In Novell Remote Manager for NetWare, open the Volumes page.
2. Click the *Information* icon to left of the name of the volume that you want to manage.
3. On the Volume Information page, click the *Purge Deleted Files* link.

26.7.2 Purging Specific Directories or Files

To immediately purge specific files that have been deleted from a directory, or to immediately purge all files from a specific directory:

IMPORTANT: Files that have been purged can no longer be recovered.

---

1. In Novell Remote Manager for NetWare, open the Volumes page.
2. Click the volume name link of the volume where the file was deleted.
3. In the volume's *Directory* list, click the *Information* icon to left of the directory 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:
   - Select the file that you want to purge, and click the *Purge* button next to the filename.
   - 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.

Managing Hard Links

27

This section describes how to configure, create, and delete hard links on Novell® Storage Services™ volumes on Novell Open Enterprise Server 2 server.

- Section 27.1, “Understanding Hard Links,” on page 375
- Section 27.2, “Upgrading the Media Format for Hard Links Support,” on page 379
- Section 27.3, “Enabling or Disabling the Hard Links Attribute,” on page 380
- Section 27.4, “Creating a Hard Link Using Ln on a Linux Client,” on page 381
- Section 27.5, “Creating a Hard Link Using a zLink API,” on page 382
- Section 27.6, “Creating a Hard Link for Testing Purposes,” on page 382
- Section 27.7, “Viewing Hard Links for a File,” on page 383
- Section 27.8, “Deleting a Hard Link,” on page 384
- Section 27.9, “Deleting a Primary Link,” on page 384

27.1 Understanding Hard Links

Novell Storage Services™ supports zero to 65,535 hard links per file on NSS volumes on NetWare® or Linux. 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 describes the following:

- Section 27.1.1, “Hard Links and the Primary Link,” on page 376
- Section 27.1.2, “Hard Links and File Ownership,” on page 376
- Section 27.1.3, “Hard Links and File System Trustees,” on page 377
- Section 27.1.4, “Hard Links and Directory Space Quotas,” on page 377
- Section 27.1.5, “Hard Links and User Space Quotas,” on page 377
- Section 27.1.6, “Hard Links and the Hard Links Attribute,” on page 378
- Section 27.1.7, “Hard Links and File Salvage and Purge,” on page 378
- Section 27.1.8, “Hard Links and DFS Move Volume,” on page 378

27.1.1 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.

When multiple hard links exist 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.

When the primary link is deleted and the status of primary is given to the next hard link, the following occurs:

- The file owner does not change.
- The explicitly assigned file system trustee assignments, trustee rights, and file attributes for the file do not change.
- If the new primary link is in a different directory, the file's inherited rights filter applies to the new parent directory, so effective rights for the file's trustees can change. For information, see Section 27.1.3, "Hard Links and File System Trustees," on page 377.

  IMPORTANT: Make sure to check the consequences to effective rights before deleting the primary link.

- For the directory space quota, the total disk space used by the file is reassigned from the old primary link's parent directory to the new primary link's parent directory. For information, see Section 27.1.4, "Hard Links and Directory Space Quotas," on page 377.
- For the user space quota, the total disk space used by the file continues to be associated with the file's assigned owner.

27.1.2 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 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 27.1.3, "Hard Links and File System Trustees," on page 377.

File ownership allows accounting of disk block usage per user and enables user space quotas to operate effectively. For information, see Section 27.1.5, "Hard Links and User Space Quotas," on page 377.

File ownership has no affect on the directory space quotas. For information, see Section 27.1.4, "Hard Links and Directory Space Quotas," on page 377.

27.1.3 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 link's parent directory. When calculating effective rights to a file, rights are granted based on explicit rights and rights inherited from the primary link's parent directory path.

If the primary link is deleted and reassigned, the inheritance changes to the settings of the newly assigned primary link's parent directory.

IMPORTANT: Make sure you verify the effective rights on the new primary link whenever you delete a primary link.

27.1.4 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 an old primary link's parent directory and added to the space usage reported on the newly assigned primary link's parent directory. The file's space usage is no longer charged against the directory space quota on the old directory. If a quota is set on the new primary link's 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 link's parent directory whenever you delete a primary link.

27.1.5 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.

27.1.6 Hard Links and the Hard Links Attribute

Hard links support is available on the following platforms (or later versions):

- OES 2 Linux and OES 2 NetWare (same as NetWare 6.5 SP7)
- OES 1 SP1 NetWare (same as NetWare 6.5 SP4)

A media upgrade is required. For instructions, see Chapter 4, "Upgrading the NSS Media Format," on page 55.

The Hard Links attribute for the NSS volume must be enabled to allow hard links to be created. For information, see Section 27.3, "Enabling or Disabling the Hard Links Attribute," on page 380.

27.1.7 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.

27.1.8 Hard Links and DFS Move Volume

If you use the Novell Distributed File Services Volume Move operation to move a volume that has been upgraded to the new media format for hard links, consider the following guidelines:

- Before you create the DFS Move Volume job, make sure that NSS is set so that the new target volume is automatically created with the upgraded media format for enhanced hard links. For information and command options, see Section 4.3, "Upgrading the Media Format Automatically for New NSS Volumes," on page 60.

- If you moved the volume without enabling the new media format, you must upgrade the volume to the new media format after the move completes successfully. For information and command options, see Section 4.2, “Upgrading the Media Format for Existing NSS Volumes,” on page 57.
- In the initial release of OES 2, the Move Volume Wizard does not have an option to enable the Hard Links attribute for the new target volume. After the move is completed and the media format is upgraded for enhanced hard links support, you must manually enable the Hard Links attribute. For instructions, see Section 27.3, “Enabling or Disabling the Hard Links Attribute,” on page 380.

27.2 Upgrading the Media Format for Hard Links Support

An enhanced NSS media format is available that provides improved support for hard links. After you install or upgrade your operating system to Novell Open Enterprise Server 2, you can decide whether to upgrade the media format for your NSS volumes to use the new metadata structure; some restrictions apply. For information, see Section 4.1, “Guidelines for Upgrading the Media Format of NSS Volumes,” on page 55.

27.2.1 New Metadata Structure Supports Up to 65,535 Hard Links for a File

NSS volumes can 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.

27.2.2 Old Metadata Structure Supports Limited Hard Links for a File

Previously, NSS volumes support only a limited number of hard links for a file. These volumes use a metadata structure that provides limited space for the file’s filenames across all name spaces.

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 name spaces. 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 name space.
- Rename the original file with a shorter name that can be stored optimally in each name space.

27.3 Enabling or Disabling the Hard Links Attribute

The new media format for enhanced hard link support provides the Hard Links attribute for NSS volumes. The Hard Links attribute must be enabled in order to create and manage hard links on an NSS volume.

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 Hard Links attribute is disabled. The attribute must be enabled before you can create hard links.

- Section 27.3.1, "Prerequisite," on page 380
- Section 27.3.2, "Hard Links Attribute Commands," on page 380
- Section 27.3.3, "Viewing the Hard Link Attribute Setting," on page 380

27.3.1 Prerequisite

The NSS volume must be upgraded to the new media format for enhanced hard link support. For instructions, see Chapter 4, "Upgrading the NSS Media Format," on page 55.

27.3.2 Hard Links Attribute Commands

Use the commands in this section to enable or disable the Hard Links attribute for an NSS volume. The Hard Links attribute cannot be set or viewed in NSSMU or in the Storage plug-in to iManager.

On Linux, issue the commands as the `root` user from the NSS Console (`nsscon`) in a Linux terminal console. On NetWare, issue the commands at the server console.

`nss /HardLinks=`*`volumename`*

Enables the Hard Links attribute for the specified volume. This enables hard links to be created on the volume.

`nss /HardLinks=all`

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.

`nss /NoHardLinks=`*`volumename`*

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.

`nss /NoHardLinks=all`

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.

27.3.3 Viewing the Hard Link Attribute Setting

The Hard Links attribute cannot be viewed in NSSMU or in the Storage plug-in to iManager. Use the `nss /volumes` command to determine whether the Hard Links attribute is set for the NSS volume. In the *Attributes* column, the Hard Links attribute is listed if the attribute is enabled.

```
avalon> nss /hardlinks=all
Creation of Hardlinks enabled on volume "JVOL"...
Creation of Hardlinks enabled on volume "TVOL"...
avalon> volumes
avalon>       Volume Name           State                Attributes
----------------------  -----------  ------------------------------
_ADMIN                  ACTIVE       Hardlinks
JVOL                    ACTIVE       Salvage
                                     Hardlinks
                                     User Space Restrictions
                                     Directory Quotas
TVOL                    ACTIVE       Salvage
                                     Hardlinks
                                     User Space Restrictions
                                     Directory Quotas
```

Linux

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, enter the following to start the NSS Console:

   ```
   nsscon
   ```
3. At the `nsscon` prompt, enter either of the following commands:

   ```
   nss /volumes
   ```

   ```
   volumes
   ```

   In the *Attributes* column, the Hard Links attribute is listed if the attribute is enabled.

NetWare

1. At the server console, enter either of the following commands:

   ```
   nss /volumes
   ```

   ```
   volumes
   ```

   In the *Attributes* column, the Hard Links attribute is listed if the attribute is enabled.

27.4 Creating a Hard Link Using Ln on a Linux Client

Typically, you create hard links by 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 (`ln`) command to create a hard link. For help with syntax and options, refer to the Link Manual Page from a Linux client. At a terminal console, enter

```
man 2 ln
```

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 27.1, "Understanding Hard Links," on page 375.

The actual content of the file is maintained in only 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.

Prerequisites

- The NSS volume must be upgraded to use the new metadata structure. For information, see "Upgrading the NSS Media Format" on page 55.
- The Hard Links attribute must be enabled for the volume. For information, see Section 27.3, "Enabling or Disabling the Hard Links Attribute," on page 380.

Procedure

To create a hard link, use the Link (`ln`) command from a Linux client.

1 At a terminal console, enter

```
ln /path/filename /linkpath/linkfilename
```

Replace */path/filename* with the pathname for the original file. Replace */linkpath/linkfilename* with the pathname for the hard linked filename.

For example, to link the file `/tmp/timetest` to `/usr/tmp/t1`, enter

```
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.

27.5 Creating a Hard Link Using a zLink API

To create a hard link from an application or script, you can use a zLink API. For information, see the *NDK: File System Services (64-bit)* (http://developer.novell.com/ndk/doc/fs64/fs64_enu/data/a5p4x94.html).

You can also use the Linux `link(2)` command to create hard links from applications or scripts.

27.6 Creating a Hard Link for Testing Purposes

While working on a server, you can use the `nss /CreateHardLink` command to create hard links for testing purposes only. In a production environment, use the Link (`ln`) command from a client instead, as described in Section 27.4, "Creating a Hard Link Using Ln on a Linux Client," on page 381.

IMPORTANT: Do not use the `nss /CreateHardLink` command in a production environment.

- Section 27.6.1, "Prerequisites," on page 382
- Section 27.6.2, "Procedure," on page 383

27.6.1 Prerequisites

- The NSS volume must be upgraded to use the new metadata structure. For information, see "Upgrading the NSS Media Format" on page 55.

- The Hard Links attribute must be enabled for the volume. For information, see Section 27.3, "Enabling or Disabling the Hard Links Attribute," on page 380.
- Hard links must be created on the same volume as the original file.

27.6.2 Procedure

To create a hard link for testing purposes on your server:

1 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.

IMPORTANT: The file paths are specified using the Long name space.

For example, to create a hard link on a volume named VOL1 on NetWare:

```
nss /CreateHardLink=VOL1:\path\file.ext|VOL1:\newpath\newfile.ext
```

27.7 Viewing Hard Links for a File

You can view a report of hard links for a file to identify its primary link and the hard link that becomes the primary link if the primary link is deleted. The `nss /ListHardLinks` command returns a list of all hard links for a specified file, where the first link in the list is the primary link. The second link is the most recently created hard link. All other hard links follow in reverse chronological order of their create time. All links for the file are listed in the order that they appear in the metadata, which is also the order in which a new primary name would be assigned.

IMPORTANT: If the primary link is deleted, the most recently created hard link (the second link reported in the list) becomes the new primary link. Changing the primary link can impact the trustees and inheritance for the file. For more information, see Section 27.1.3, "Hard Links and File System Trustees," on page 377.

The list of hard links contains the following information for the file:

List of Hard Links	Information
Line 1 of the report	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 of the report	The complete name of the primary link, including the path and filename
Line 3 of the report	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.
Lines 4 through 65,536 (up to 65,535 links per file) of the report	The complete name of each of the hard link names, including the path and filename, are listed in the reverse chronological order of their creation time.

27.7.1 Linux

To view information about the primary link and hard links for a file on an NSS volume on Linux:

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, enter the following to open the NSS Console:

   ```
   nsscon
   ```

3. At the `nsscon` prompt, 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.

27.7.2 NetWare

To view information about the primary link and hard links for a file on an NSS volume on NetWare:

1. 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.

27.8 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 it with the delete commands and procedures native to the client you use to access the file.

27.9 Deleting a Primary Link

When multiple hard links exist for a file, deleting its primary link does not delete the content and metadata. These are deleted only when the last remaining link is deleted. While a file has multiple links, if you delete a primary link, the most recently created hard link automatically becomes the new primary link.

If the most recently created hard link is not the link you want to become the primary link, you can delete and re-create the preferred link, then delete the primary link. Make sure you do not create any other new hard links in the meantime.

Deleting a primary link has the following consequences:

- **Order of Ascendancy:** Whenever you delete a primary link for a file with hard links, the most recently created hard link automatically becomes the new primary link.

  To identify the pathname of the hard link that is next in line to become the new primary link, see Section 27.7, “Viewing Hard Links for a File,” on page 383.

  For information about primary links, see Section 27.1.1, “Hard Links and the Primary Link,” on page 376

- **Effective Rights for File System Trustees:** When the inherited rights filter is applied to the new primary link’s parent directory, the effective rights for trustees might change.

  For information, see Section 27.1.3, “Hard Links and File System Trustees,” on page 377

- **Directory Quotas:** When the new primary link is in a different directory, the directory quotas are affected for the old and new parent directories.

  For information, see Section 27.1.4, “Hard Links and Directory Space Quotas,” on page 377

There are no special commands required to delete a primary link. In a file manager, locate the link name to be deleted, then delete it with delete commands and procedures native to the client you use to access the file.

Managing Files and Folders on NSS Volumes

28

This section provides an overview of how to manage files and folders as an administrator of a Novell® Storage Services™ (NSS) volume using the Files and Folders plug-in for Novell iManager 2.7. For more information about using the Files and Folders plug-in to configure the file system trustees, trustee rights, inherited rights filters, and file and folder attributes, see Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297.

- Section 28.1, "Creating a Folder on an NSS Volume," on page 387
- Section 28.2, "Deleting a File or Folder on an NSS Volume," on page 388
- Section 28.3, "Uploading Files to an NSS Volume," on page 389
- Section 28.4, "Downloading Files from an NSS Volume," on page 390
- Section 28.5, "Viewing or Modifying File or Folder Properties," on page 391
- Section 28.6, "Viewing or Modifying File Ownership," on page 394
- Section 28.7, "Viewing, Adding, Modifying, or Removing a Directory Quota," on page 396

28.1 Creating a Folder on an NSS Volume

As an administrator, you can use the Files and Folders plug-in to iManager to create a folder on an NSS volume.

28.1.1 Prerequisites

- The destination NSS volume must be in the same tree where you are currently logged in to iManager.
- You must have trustee rights for the volume and destination location where you want to create the new folder. The Create right is required for creating files and folders.

28.1.2 Procedure

1 In iManager, click *Files and Folders*, then click *New Folder* to open the *New Folder* page.

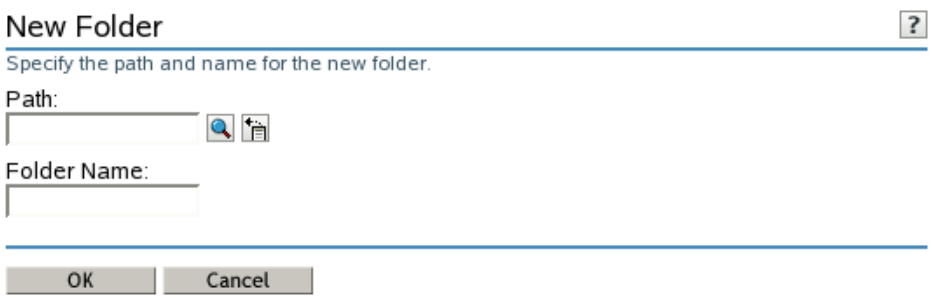

2 Use one of the following methods to specify the destination path on the NSS volume where you want to create the new folder:

- Click the *Search* icon to browse and locate the destination folder, then click the name link of the folder to select it.
- Click the *History* icon to select a folder from the list of folders that you recently accessed.

The pathname of the folder appears in the *Path* field.

3 In *Folder Name*, type the name the folder you want to create in the selected location.

4 Click *OK* to create the folder, or click *Cancel* to abandon it.

A message confirms when the folder has been successfully created.

5 Click *Repeat Task* to create another folder, or click *OK* to dismiss the confirmation message.

6 Click *Files and Folders*, then click *Properties* to set file system trustees, trustee rights, and attributes for the new folder or folders.

For instructions for configuring properties, see Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297.

28.2 Deleting a File or Folder on an NSS Volume

As an administrator, you can use the Files and Folders plug-in to iManager to delete a file or folder on an NSS volume.

28.2.1 Prerequisites

- The NSS volume must be in the same tree where you are currently logged in to iManager.
- You must have trustee rights for the file or folder that you want to delete. The Erase right is required to delete the file.
- A folder must be empty before it can be deleted.

28.2.2 Procedure

1 In iManager, click *Files and Folders*, then click *Delete* to open the *Delete File or Folder* page.

2 Use one of the following methods to specify the file or folder that you want to delete from the NSS volume:

- Click the *Search* icon to browse and locate the file or folder, then click the name link of the object to select it.
- Click the *History* icon to select a file or folder from the list of files and folders that you recently accessed.

The pathname of the folder appears in the *Name* field.

3 Click *OK* to delete the selected file or folder, or click *Cancel* to abandon the delete process.

   A message confirms when the file or folder has been successfully deleted.

4 Click *Repeat Task* to delete another folder, or click *OK* to dismiss the confirmation message.

28.3 Uploading Files to an NSS Volume

As an administrator, you can use the Files and Folders plug-in to iManager to upload files from your local computer to an existing folder on an NSS volume.

28.3.1 Prerequisites

- The destination NSS volume must be in the same tree where you are currently logged in to iManager.
- You must have trustee rights for the destination folder in order to be able to find the folder and upload the file. The Create right is required for file uploads.

28.3.2 Procedure

1 In iManager, click *Files and Folders*, then click *Upload* to open the *Upload File* page.

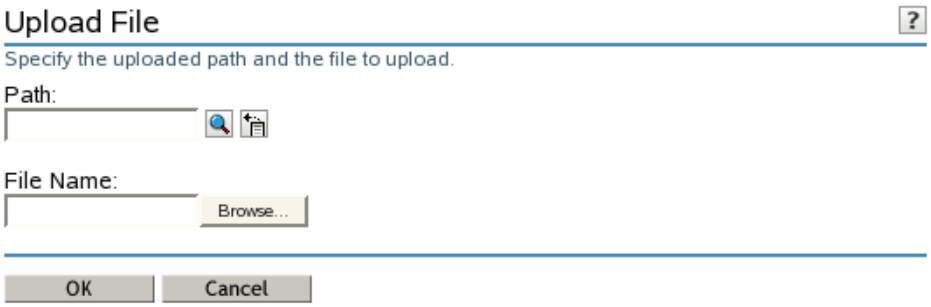

2 Use one of the following methods to specify the path to the folder on the NSS volume where you want to put the file:

   - Click the *Search* icon to browse and locate the folder, then click the name link of the folder to select it.
   - Click the *History* icon to select a folder from the list of folders that you recently accessed.

   The pathname appears in the Path field.

3 Select the file on your local computer that you want to upload:

   3a Click *Browse* to open a local file browser dialog box.

   3b Browse and locate the file.

   3c Select the file, then click *Open*.

   The local pathname for the selected file appears in the *File Name* field.

4 Click *OK* to begin the upload, or click *Cancel* to abandon the process.

   A message confirms when the file has been successfully uploaded. Wait until the upload completes before proceeding to other tasks.

5 Click *Repeat Task* to upload another file, or click *OK* to dismiss the confirmation message.

28.4 Downloading Files from an NSS Volume

As an administrator, you can use the Files and Folders plug-in to iManager to download a file from an NSS volume to your local computer.

28.4.1 Prerequisites

- The NSS volume must be in the same tree where you are currently logged in to iManager.
- You must have trustee rights for the file in order to be able to browse to and download the file.

28.4.2 Procedure

1 In iManager, click *Files and Folders*, then click *Download* to open the *Download File* page.

2 Use one of the following methods to select the file that you want to download from the NSS volume to your local drive:

- Click the *Search* icon to browse and locate the file, then click the name link of the file to select it.
- Click the *History* icon to select a file from the list of files that you recently accessed.

The pathname appears in the *File Name* field.

3 Click *OK* to open the *File Download* dialog box.

IMPORTANT: If the File Download dialog box does not open, make sure the security settings in your browser allow downloads from the server by adding the server as a trusted site, then try again.

---

4 Use one of the following methods to save the file to the local computer:

- Click *Open* to view the file in an appropriate application, then save the file by using the application's *File* > *Save* options.

  The application that opens the file must already be installed on your computer.
- Click *Save* to open the *Save As* dialog box, browse to an existing folder or create a new local folder where you want to save the file, then click *Save*.

  The browser's download manager manages the download and notifies you when the download is complete.

You can continue with other iManager tasks while the file is downloading.

28.5 Viewing or Modifying File or Folder Properties

1 In iManager, click *Files and Folders* > *Properties* to open the *Properties* page.

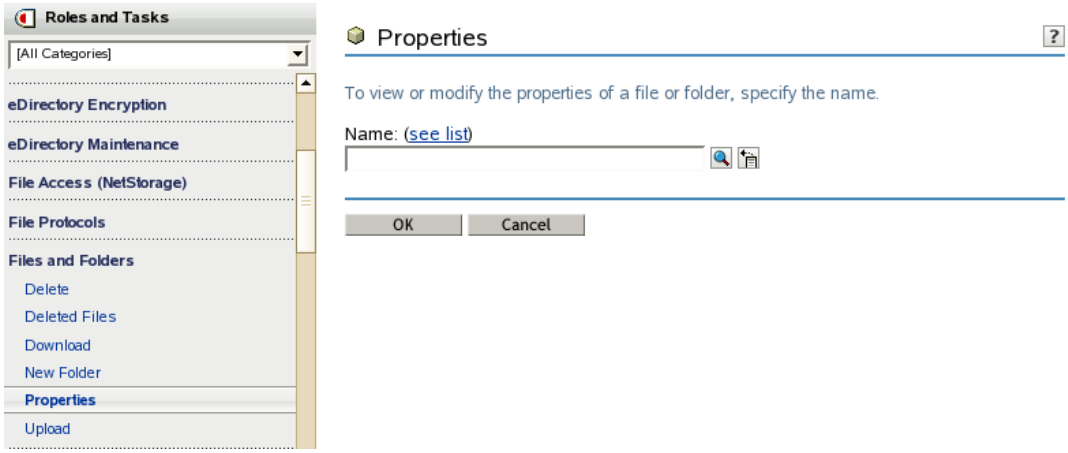

2 Click the *Search* icon to browse the Storage objects, locate and select the name link of the file or folder you want to manage, then click *OK* to view the Properties for the file.

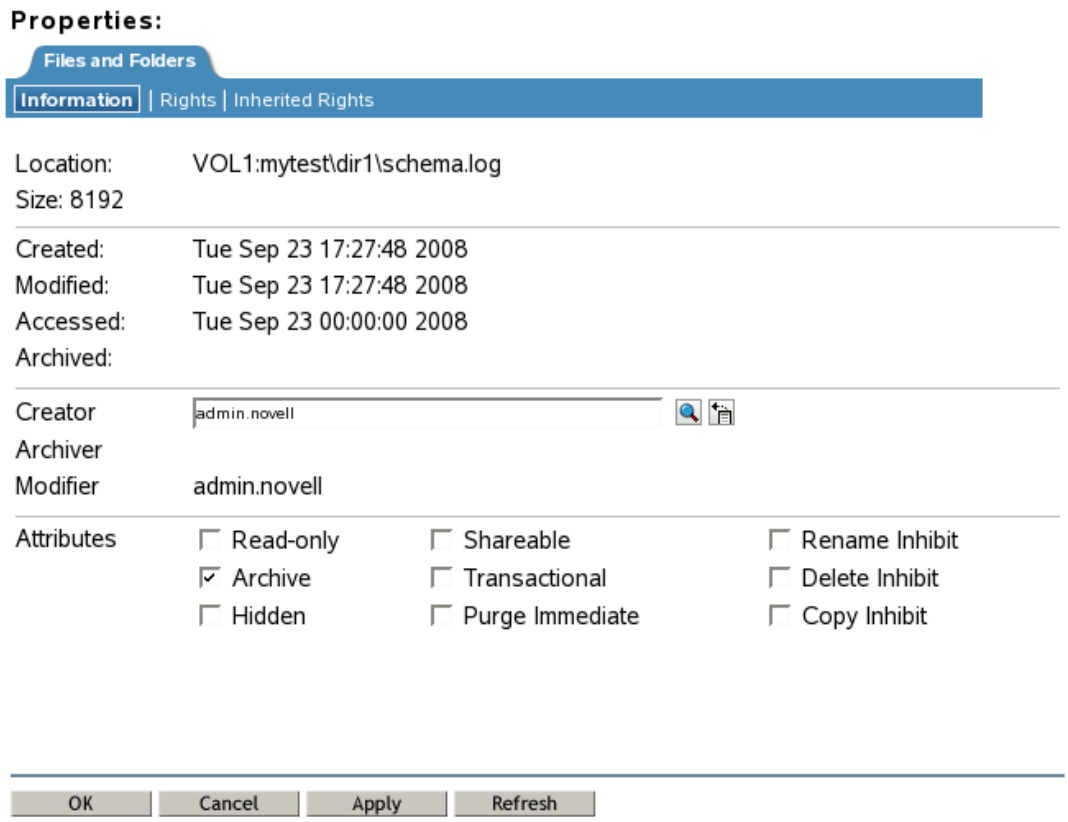

3 Click the *Information* tab to view or modify the following information for the selected folder or file:

Property	Description
Location	The pathname of the selected volume, folder, or file. For example: `VOL1:dir1\dirB\filename.ext`
Restrict Size (Enable or Disable a Directory Quota on a Folder)	Enable (select) or disable (deselect) a directory quota on the specified folder on an NSS volume where the Directory Quotas attribute is enabled. The default is Disabled. If this option is enabled, you must also specify a value for the quota in the Limit field. A directory quota limits the amount of space on a volume that can be consumed by all of the files and folders in that directory. The directory quota applies to files and folders created by any user of the directory. Select *Restrict Size* to enable a directory quota for the selected folder, specify the quota value in *Limit*, then click *Apply*. Deselect *Restrict Size* to disable a directory quota for the selected folder, then click *Apply*.
Limit (Set Limit for a Directory Quota on a Folder)	The maximum size allowed for the specified directory and its contents. Default: Disabled (not available unless *Restrict Size* is enabled). If you enable *Restrict Size* for the selected folder, you must specify a limit for the directory quota. Type a value in KB for the quota. The value must be an increment of 4 KB; that is, it must be divisible by 4 with no remainder. Click *Apply* to save the changes. If the directory quota exceeds the volume quota, the volume quota is enforced. If the current size of the selected folder exceeds the specified limit, users cannot save data to the folder until space is cleared by removing files from it. If a user quota is set for a user on the volume, the user space restriction overrides the directory quota. That is, the user cannot save data to the folder if doing so causes the user to exceed his or her user quota.
Created	The time stamp (MM/DD/YYYY hh:mm) for when the file or folder was created.
Modified	The time stamp (MM/DD/YYYY hh:mm) for when the file or folder was last modified.
Accessed	The time stamp (MM/DD/YYYY hh:mm) for when the file or folder was last accessed.
Archived	The time stamp (MM/DD/YYYY hh:mm) for when the file or folder was last archived.

Property	Description
Creator (View or Modify Ownership)	The typeless distinguished Novell eDirectory™ username (such as username.context) of the user who created the file or folder. If the username becomes invalid, such as if an employee leaves the company, the GUID of the username is reported. For NSS, any number of files or folders can be represented by GUIDs instead of valid usernames. User quotas for NSS volumes consider file ownership to enforce user space restrictions. You might need to change the ownership of a file or folder in order to make the space it consumes be charged against a different user. For NSS volumes (as for all volumes that use the Novell trustee model of access), all access to data is controlled by file system trustees and trustee rights instead of by ownership. When a user creates a file or folder, the trustees and trustee rights for accessing the file are automatically inherited from the directory where the file is created. If you intend different trustees and rights for the file, you must assign them explicitly by user, or assign the rights to a group and put the users into that group. For instructions, see Section 22.1.4, “Configuring Rights Properties (File System Trustees, Trustee Rights, and Inherited Rights Filter),” on page 300. Changing the ownership of the file or folder does not modify who can access it, but it does modify whose username is charged for the space it consumes. If you modify the ownership, you must click Apply or OK to save the changes.
Archiver	The distinguished username (such as username.context) of the user who modified the version of the file or folder that was last archived.
Modifier	The distinguished username (such as username.context) of the user who last modified the current version of the file or folder.
Attributes	File attributes determine how the file or folder behaves when accessed by any user. Enable or disable an attribute by selecting or deselecting the check box next to it. If you modify a setting, click *Apply* or *OK* to save the changes. File attributes apply universally to all users. For example, a file that has a read-only attribute is read-only for all users. Attributes can be set by any trustee with the Modify right to the directory or file, and attributes stay set until they are changed. Attributes do not change when you log out or when you down a file server. For example, if a trustee with the Modify right enables the Delete Inhibit attribute for a file, no one, including the owner of the file or the network administrator, can delete the file. However, any trustee with the Modify right can disable the Delete Inhibit attribute to allow the file’s deletion.

The following table defines file system attributes and whether they apply to files, folders, or both files and folders.

Attribute	Description	Files	Folders
Read Only	Prevents a file from being modified. This attribute is typically used in combination with Delete Inhibit and Rename Inhibit.	Yes	No
Archive	Identifies files and folders that have been modified since the last backup. This attribute is assigned automatically.	Yes	Yes
Hidden	Hides directories and files so they do not appear in a file manager or directory listing.	Yes	Yes
Shareable	Allows more than one user to access the file at the same time. This attribute is usually used with Read Only.	Yes	No
Transactional	Allows a file on an NSS volume or a NetWare Traditional volume to be tracked and protected by the Transaction Tracking System™ (TTS™) for NetWare. For NSS, the TTS attribute for the volume must be enabled in order for this setting to be enforced. TTS is not available for NSS on Linux.	Yes	No
Purge Immediate	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
Rename Inhibit	Prevents the directory or filename from being modified.	Yes	Yes
Delete Inhibit	Prevents users from deleting a directory or file. This attribute overrides the file system trustee Erase right. When Delete Inhibit 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 attribute to allow the directory or file to be deleted.	Yes	Yes
Copy Inhibit	Prevents users from copying a file. This attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare. This attribute overrides the trustee Read right and File Scan right. A trustee with the Modify right must disable this attribute to allow the file to be copied.	Yes	No

4 If you modified any settings, click *Apply* or *OK* to save your changes.

28.6 Viewing or Modifying File Ownership

The owner of a file is assigned by default to be the identity of the user who creates the file. Ownership does not determine who can access a file because the NSS file system uses the Novell trustee model to control access. However, user quotas for NSS volumes consider file ownership to enforce user space restrictions. You might need to change the ownership of a file or folder in order to make the space it consumes be charged against a different user. Changing the ownership of the file or folder does not modify who can access it, but it does modify whose username is charged for the space it consumes.

The Creator field shows the typeless distinguished Novell eDirectory™ username (such as username.context) of the user who owns the file or folder. If the username becomes invalid, such as if an employee leaves the company, the GUID of the username is reported. For NSS, any number of files or folders can be represented by GUIDs instead of valid usernames.

1. In iManager, click *Files and Folders*, then click *Properties* to open the *Properties* page.
2. Click the *Search* icon to browse and locate file from the Storage objects, click the name link of the file to select it.

   The pathname of the file or folder appears in the *Name* field.
3. Click *OK* to open the file's Properties page.

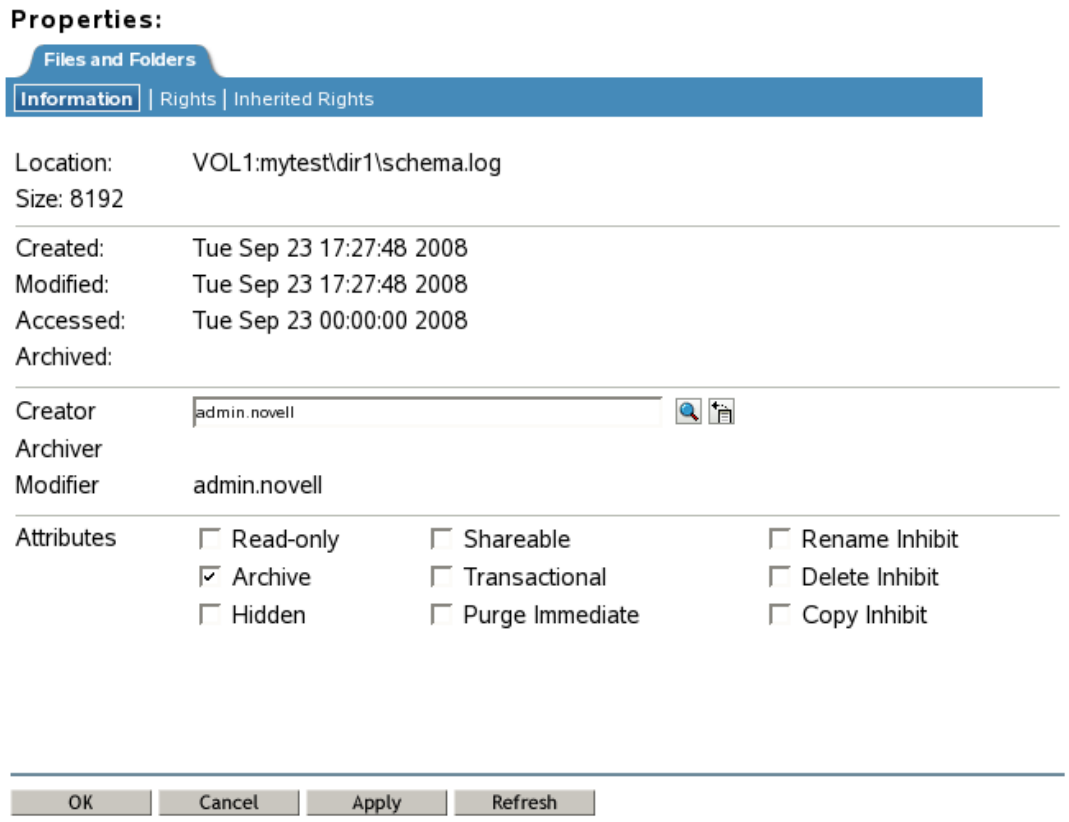

4. On the Information page, the *Creator* field shows the typeless distinguished username of the current owner, such as `username.context`.

5. If you want to modify the owner, click the *Search* icon to open the *Object Browser* dialog box, then locate and select the username of the new owner.
6. If you modified the owner, click *Apply* or *OK* on the Information page in order to save the change.

28.7 Viewing, Adding, Modifying, or Removing a Directory Quota

Directory quotas for NSS volumes require that the Directory Quotas attribute be set for the volume. For information, see Section 25.3.1, “Enabling or Disabling the Directory Quotas Attribute for an NSS Volume,” on page 349.

1. In iManager, select *Files and Folders* > *Properties*.
2. Click the *Search* icon, browse to locate and select the folder you want to manage on an NSS volume, then click *OK* to open the Properties page for the selected folder.

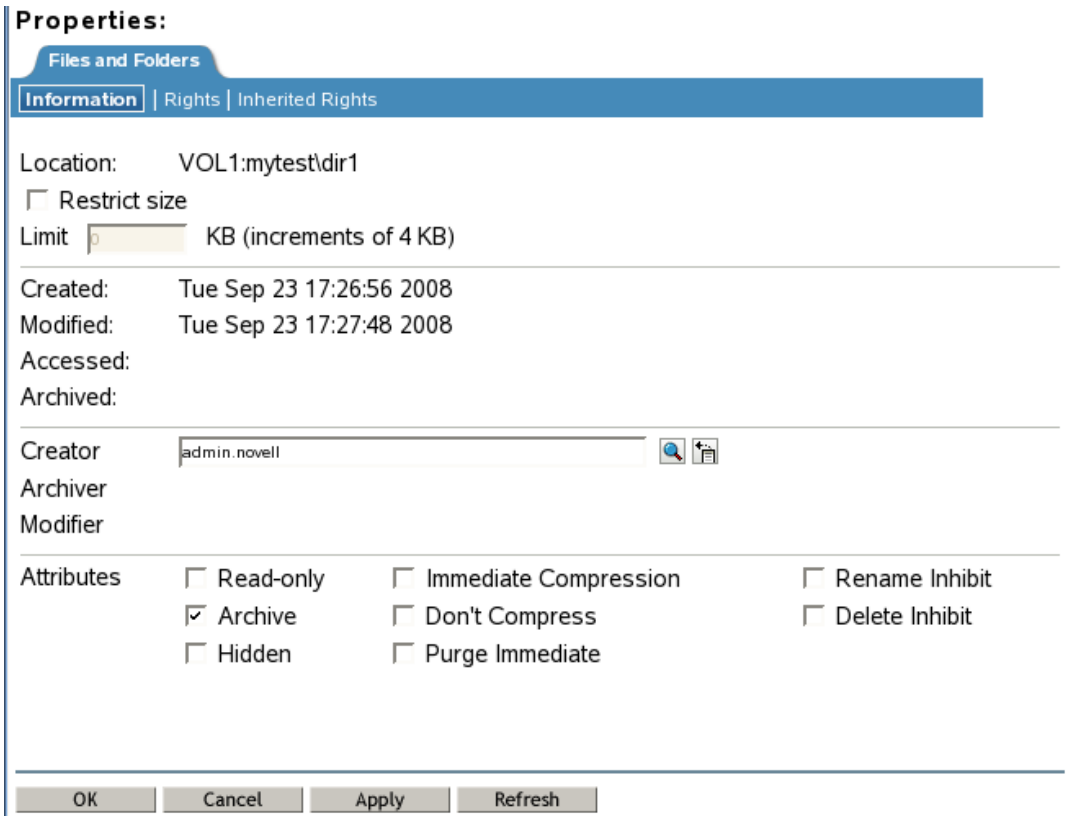

3. View the current status of the Directory Quota.

   If a Directory Quota is set, the *Restrict Size* field is selected and the *Limit* field shows the quota size in KB.

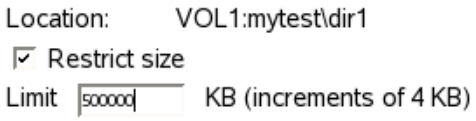

   If the Directory Quota is not set, the *Restrict Size* field is deselected and the *Limit* field is dimmed (grayed out).

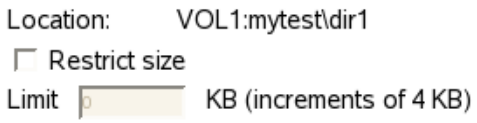

4 Do one of the following:

- **Add a Quota:** On the *Information* tab, select *Restrict Size* to enable space restrictions for the selected directory. In the *Limit* field, type the directory quota in KB. The value must be an increment of 4 KB; that is, it must be divisible by 4 with no remainder.
- **Modify an Existing Quota:** In the *Limit* field, type the new directory quota in KB. The value must be an increment of 4 KB; that is, it must be divisible by 4 with no remainder.
- **Remove a Quota:** On the *Information* tab, deselect *Restrict Size* to disable space restrictions for the selected directory. The *Limit* field is automatically dimmed (grayed out).

5 On the *Information* page, click *Apply* or *OK* to apply the changes.

Managing Backup and Restore for Data and Trustee Information

29

This section describes your backup and restore options for data and trustee information for Novell® Storage Services™ volumes on Novell Open Enterprise Server 2.

- Section 29.1, "Using Novell Storage Management Services," on page 399
- Section 29.2, "Using the Event File List to Refine the Backup," on page 400
- Section 29.3, "Using METAMIG to Save and Restore Trustee Information on NSS and Linux POSIX File Systems (Linux)," on page 400
- Section 29.4, "Using Extended Attributes (xAttr) Commands (Linux)," on page 400
- Section 29.5, "Backing Up Files Without Altering the Access Time," on page 402
- Section 29.6, "Additional Information," on page 403

29.1 Using Novell Storage Management Services

Novell® Storage Management Services™ (SMS) provide backup, restore, and data migration services for NSS volumes. For information, see the following:

- *OES 2 SP1: Storage Management Services Administration Guide*
- *OES 2 SP1: SBCON Administration Guide*
- *NBACKUP Utility for OES Linux* (`nbackup(1)`) (http://www.novell.com/documentation/oes/smsadmin/data/nbackup.1.html). The `nbackup` utility is included in the `novell-sms` RPM file under `/opt/novell/sms/bin`. It can be used to back up and restore NSS and non-NSS volumes by using the SMS framework. It backs up and restores NSS metadata, which includes file system trustees and trustee rights information. Trustees on NSS volumes are tied to Novell eDirectory™ users and objects, so you need to protect eDirectory in addition to backing up data. The man page for `nbackup` provides further details on its usage.

  Related utilities and configuration files include the following:

  - `smdrd(8)` (http://www.novell.com/documentation/oes/smsadmin/data/smdrd.8.html)
  - `smdrd.conf(5)` (http://www.novell.com/documentation/oes/smsadmin/data/smdrd.conf.5.html)
  - `sms(7)` (http://www.novell.com/documentation/oes/smsadmin/data/sms.7.html)
  - `smsconfig(1)` (http://www.novell.com/documentation/oes/smsadmin/data/smsconfig.1.html)
  - `tsafs(1)` (http://www.novell.com/documentation/oes/smsadmin/data/tsafs.1.html)
  - `tsafs.conf(5)` (http://www.novell.com/documentation/oes/smsadmin/data/tsafs.conf.5.html)

Encrypted NSS volumes store user data in encrypted format on the NSS volume, yet work transparently with most applications, NLM™ programs, and backup utilities that currently work with NSS.

SMS backs up and restores compressed files in their compressed state. It does not compress uncompressed files for backup; they are stored and restored in their uncompressed state.

29.2 Using the Event File List to Refine the Backup

NSS uses the Event File List (EFL) feature to track files that have changed on a volume during an interval called an *epoch*. It logs changes that are made to data and metadata for each active epoch on a specific NSS volume in the `_admin:manage_nss\volume\volumename\FileEvents.xml` file.

Your backup solution can take advantage of this file in order to get a list of modified files for NSS volumes on Linux and NetWare. You can use the API commands in scripts to start and stop an epoch, reset the event list for an epoch, and to affect how long epochs are retained.

For information about the Event File List (EFL) APIs for developers, see "FileEvent.xml Definitions" (http://developer.novell.com/documentation/vfs/vfs__enu/data/ak7gh2x.html) in *NDK: Virtual File Services (http://developer.novell.com/documentation/vfs/vfs__enu/data/bktitle.html)*.

See Cool Tools (search for EFL) (http://www.novell.com/communities/search/node/EFL) on the Novell Cool Solutions Web site for scripts that use the Event File List APIs.

29.3 Using METAMIG to Save and Restore Trustee Information on NSS and Linux POSIX File Systems (Linux)

The `metamig` utility allows you to save and restore trustee information for NSS volumes. You can also restore trustee information for any NCP™ volume that was backed up as raw data with a third-party backup application. For information, see Section B.8, "METAMIG (Linux)," on page 514.

For OES Linux, the NCP Server allows you to create NCP volumes for Linux POSIX file systems. NSS volumes are NCP volumes by default. You can assign trustees and trustee rights for NCP volumes based on Linux file systems just as you do for NSS volumes. The trustee information is located in a hidden file on the volume rather than being integrated in the volume. When you use a third-party backup application to backup files as raw data, the trustee file is also backed up as raw data. You can use the `metamig -ncp` option to restore trustee information for NCP volumes on OES Linux that use the NSS file system or Linux POSIX file systems.

For information about creating NCP volumes, see "Managing NCP Volumes" in the *OES 2 SP1: NCP Server for Linux Administration Guide*.

29.4 Using Extended Attributes (xAttr) Commands (Linux)

In OES SP2 and later, NSS supports the Linux extended attributes (XAttr) option that allows listing, saving, and restoring the trustee information that is stored in the `netware.metadata` extended attribute. Third-party backup software that supports the standard Linux Extended Attributes (xattr) can use this feature for NSS volumes to preserve trustees, trustee rights, file attributes, and quotas in backup and restore.

Support for Linux `xattr` is disabled by default. It is necessary to set the following NSS switches to enable this behavior:

```
nss /ListXattrNWMetadata
nss /CtimeIsMetadataModTime
```

If you issue the commands from the command line, the support is automatically disabled at the next server reboot. You can enable the support for Linux `xattr` across server reboots by adding the switches to the `/etc/opt/novell/nss/nssstart.cfg` file.

- Section 29.4.1, "Enabling NSS Support for Linux xAttr," on page 401
- Section 29.4.2, "Disabling NSS Support for Linux xAttr," on page 401
- Section 29.4.3, "Additional Information," on page 402

29.4.1 Enabling NSS Support for Linux xAttr

- "Using NSSCON" on page 401
- "Using the nssstart.cfg File" on page 401

Using NSSCON

To enable support for Linux `xattr` from NSSCON:

1. Open a Linux terminal console, then log in as the `root` user.
2. Start NSSCON by entering the following at the console prompt:

   ```
   nsscon
   ```

3. To enable the Linux `xattr` support for all NSS volumes on the server, enter

   ```
   nss /ListXattrNWMetadata
   nss /CtimeIsMetadataModTime
   ```

   The commands are enabled until the next server reboot. You can also issue commands that disable the support.

Using the nssstart.cfg File

You can enable the support for Linux `xattr` across server reboots by adding the following lines to the `/etc/opt/novell/nss/nssstart.cfg` file:

```
/ListXattrNWMetadata
/CtimeIsMetadataModTime
```

Make sure the switches are spelled correctly, and do not have spaces after the forward slash (`/`). If the switch names are entered incorrectly in the `nssstart.cfg` file, parsing errors can prevent the NSS pool from mounting.

29.4.2 Disabling NSS Support for Linux xAttr

- "Using NSSCON" on page 402
- "Using the nssstart.cfg File" on page 402

Using NSSCON

To disable support for Linux `xattr` from NSSCON:

1. Open a Linux terminal console, then log in as the `root` user.
2. Start NSSCON by entering the following at the console prompt:

   ```
   nsscon
   ```

3. To disable the Linux `xattr` support for all NSS volumes on the server, enter

   ```
   nss /noListXattrNWMetadata
   nss /noCtimeIsMetadataModTime
   ```

Using the nssstart.cfg File

If you added the switches to the `/etc/opt/novell/nss/nssstart.cfg` file, and you want the support for Linux xattr to be automatically disabled after a server reboot, remove the switches from the file. When the server reboots, the Linux xattr support is disabled, which is the default behavior.

29.4.3 Additional Information

For information, see Section A.11, “Extended Attributes (XAttr) Commands (Linux),” on page 464.

29.5 Backing Up Files Without Altering the Access Time

You can set the `noatime` option on NetWare or Linux to control whether the access time is updated when reading files and directories.

29.5.1 Linux

In OES 2 SP1 Linux, NSS provides the `/atime` and `/noatime` options for Linux. For information, see Section A.25, “noatime and atime Commands,” on page 478.

In OES SP2 and later, NSS provides `noatime` and `nodiratime` support for the Linux `open(2)` API command, `mount` command, and the `/etc/fstab` configuration file. Backup applications can take advantage of this option to back up a file without altering its access time. For information, see Section A.26, “noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux),” on page 479.

29.5.2 NetWare

In OES 2 SP1 NetWare (same as NetWare 6.5 SP8), NSS provides the `/atime` and `/noatime` options for NetWare. For information, see Section A.25, “noatime and atime Commands,” on page 478.

29.6 Additional Information

For a current list of backup software vendors that support Novell Open Enterprise Server, see *Novell Open Enterprise Server Partner Support: Backup and Antivirus Support* (http://www.novell.com/products/openenterpriseserver/partners/). This list is updated quarterly.

Tuning NSS Performance on NetWare

30

This section describes proactive strategies that can help improve Novell® Storage Services™ performance on a Novell Open Enterprise Server 2 NetWare® server.

- Section 30.1, “Do I Need to Tune NSS?,” on page 405
- Section 30.2, “NSS Tuning Guidelines,” on page 405
- Section 30.3, “Setting the Cache Balance,” on page 407
- Section 30.4, “Setting the Closed File Cache Size,” on page 410
- Section 30.5, “Setting the File and Buffer Flush Timers,” on page 410
- Section 30.6, “Setting the Name Cache Size,” on page 411
- Section 30.7, “Configuring NSS Extended System Memory,” on page 411
- Section 30.8, “Tuning NSS Write Performance on NetWare,” on page 412

30.1 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 30.2, “NSS Tuning Guidelines,” on page 405.

30.2 NSS Tuning Guidelines

Consider the following before you adjust the cache management parameters:

- Section 30.2.1, “Before You Adjust Cache Parameters,” on page 405
- Section 30.2.2, “Hints for Low Memory NetWare Systems,” on page 406
- Section 30.2.3, “The Tuning Process,” on page 406
- Section 30.2.4, “Restoring Settings After a Failed Tuning Attempt,” on page 407

30.2.1 Before You Adjust Cache Parameters

Use Novell Remote Manager for NetWare to check memory usage statistics for NSS volumes.

IMPORTANT: Do not use the Monitor utility to check memory configuration and usage when running NSS volumes. The Cache Buffers statistics in Monitor apply only to NetWare Traditional file systems. Instead, use Novell Remote Manager for NetWare.

Before you adjust cache parameters, consider the following guidelines:

- 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 cache for the NetWare Traditional file system. 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 halves 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 30.2.2, "Hints for Low Memory NetWare Systems," on page 406.

30.2.2 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 ships 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:

nss /ClosedFileCacheSize=*value*

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.

nss /MinBufferCacheSize=*value*

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.

30.2.3 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 available, NSS fails to load. To recover, see Section 30.2.4, “Restoring Settings After a Failed Tuning Attempt,” on page 407

30.2.4 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, enter the following command at the server console:

```
server -z parameter
```

Replace *parameter* with the parameter you want to reset to its previous value, such as

```
server -z CacheBalance
```

30.3 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 need to dedicate more RAM to NSS.

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. The following NSS parameters can be used to adjust the cache balance between user and system data. You can set or tune cache buffers at anytime.

nss /AuthCacheSize=*value*

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

nss /CacheBalance=*value*

Set the amount of cache that NSS takes from the file system. A low cache balance impedes the performance of NSS.

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.

NOTE: You can change the default time interval with the CacheBalanceTimer option.

Default: 85

Range: 1 to 99

The default is On. Use this command in combination with the /MinBufferCacheSize switch.

nss /NoCacheBalance

Disables the dynamic balancing of free memory for the NSS buffer cache.

nss /CacheBalanceMaxBuffersPerSession=*value*

Set a maximum amount of memory to transfer to NSS during a cache balancing session.

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.

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.

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 complies at the time of the request in one transfer.

Default: 1024

Range: 16 to memory size in KB/4 KB

nss /CacheBalanceTimer=*value*

Determine how often NSS rebalances the amount of cache it uses.

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.

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.

Default: 30

Range: 1 to 3600

nss /CacheHit

Show the cache hit ratio at the server console.

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.

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.

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.

CacheStats (or nss /CacheStats)

Show cache buffer statistics.

nss /CacheUserMaxPercent=*value*

Control the division of the file system cache between user and system data.

Sets the specified maximum percentage of buffer cache that can be used for user data.

If you have a server with a few files and volume attributes and a generous amount of cache RAM, you can increase this parameter.

Default: 80

Range: 10 to 100

nss /MinBufferCacheSize=*value*

Set the number of cache buffers allocated for NSS.

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.

Default: 512

Range: 256 to memory size in KB/ 4 KB

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.

To safeguard against this, use the MinOSBufferCacheSize option to set a hard limit for the amount of RAM that can be used exclusively by the operating system.

TIP: 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 only be below 24 hours right after a reboot.

---

nss /MinOSBufferCacheSize=*value*

Set a minimum amount of RAM exclusively for the operating system.

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.

Default: 256

Range: 256 to memory size in KB/4 KB

nss /reset

Reset cache statistics to clear the current statistics and start over. This is useful when checking the results of changes to the system.

30.4 Setting the Closed File Cache Size

The ClosedFileCacheSize parameter keeps the storage objects in cache so that NSS does not need to go to disk and unpack the information again when it wants the same files.

nss /ClosedFileCacheSize=*value*

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.

Default: 100000

Range: 16 - 1000000. On average, each Closed File Cache entry consumes 0.4 - 1 KB of RAM.

TIP: The Closed File Cache grows up to the specified amount. Unlike file system cache, the 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.

30.5 Setting the File and Buffer Flush Timers

Increasing the flush time settings can improve performance, but does it at the risk of data loss during a system failure. The timers relate directly to the amount of uncommitted data in the NSS cache. The higher the timer setting, the more dirty blocks are waiting to be written to disk. The lower the timer setting, the more frequently data is written, so there is less data queuing for the write and less uncommitted data that could potentially be lost if the server were to crash.

The default timer value of 10 seconds is best suited for most production environments. If you decrease the value (increase the number of writes), you should monitor the performance of the system to understand how more frequent writes are impacting the system. If you increase the value (decrease the number of writes), you should monitor the ZLOG file and cache statistics to understand how increasing the amount of uncommitted data might affect your users and data recovery procedures in the event of a crash. The maximum of 60 minutes (3600 seconds) is available for the discretionary flexibility of the system administrator.

It is up to the administrator to determine how frequently data is committed to the disk. Where you set the timer depends the criticality of the data, how much data there is to be written, and how fast dirty blocks are generated for writes (turnover of the data). It also depends on how sensitive your environment is to long reboot times should the server crash. The more unwritten data transactions that have to be resolved from the journal, the longer it takes to bring up the system after the crash.

The following NSS parameters can be used to adjust the file and buffer flush timers:

nss /FileFlushTimer=*value*

Set the maximum amount of time in seconds 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 (seconds)

Range: 1 to 3600 (seconds)

nss /BufferFlushTimer=*value*

Set the maximum amount of time in seconds that modified NSS file system cache buffers are kept in memory before they are written to disk.

```
nss /BufferFlushTimer=value
```

IMPORTANT: In OES 2 Linux, this timer is not used. It has been replaced by the group write timers for Journal, Metadata, and User Data. For information, see Section 31.3, “Configuring or Tuning Group I/O,” on page 415.

Default: 1 (second)

Range: 1 to 3600 (seconds)

30.6 Setting the Name Cache Size

The NSS Name Cache is responsible for caching the Name Tree information. This is the information that is read when you perform any kind of search by file or directory name. The Name Cache maps a name to a ZID (a unique file object ID). Directory listings do not do this as much as normal file opens that must resolve each name in the file path.

Use the NameCacheSize parameter to specify the amount of recently used Name Tree entries for files and directories that NSS caches. Each entry uses about 150 bytes of memory.

Increasing the maximum number of entries Name Cache entries does not necessarily improve the performance for getting directory listing information if NSS also needs to look up information about the file from a tree or structure outside of the name tree.

If you want to see how your name cache is performing, use the `nss /NameCacheStats` command.

nss /NameCacheSize=*value*

Specify the maximum number of recently used Name Tree entries for files and directories to cache. Name cache grows up to the specified limit. Unlike the file system cache, it does not take the maximum amount of memory allocated from the start.

Default: 100000

Range: 17 to 1000000

30.7 Configuring NSS Extended System Memory

Extended system memory (ESM) is the memory above 4 GB. By default, NSS on NetWare servers uses 80% of the ESM to cache user data for files greater than 128 KB in size. Only files with an End-of-File (EOF) files size greater than or equal to 128KB can be stored by ESM. To use ESM, the LSS must call it. ESM is called only by ZLSS and the `cddvd` LSSs. For example, large DOSFAT files are not stored in ESM memory.

nss /Above4GigMemoryPercent=*value*

Specifies the percentage of extended system memory used. This value can be set only at initialization time by adding the `Above4GigMemoryPercent` option in the `nssstart.cfg` file of your NetWare server.

For example:

```
nss /Above4GigMemoryPercent=90
```

Additional NSS command line options for managing NSS usage of ESM are:

nss /TrackESMCachePerf

Enables monitoring of the NSS ESM.

nss /NoTrackESMCachePerf

Disables monitoring of the NSS ESM.

nss /ESMCachePerf

Prints the statistics for NSS ESM.

30.8 Tuning NSS Write Performance on NetWare

For information about tuning OES 2 NetWare (NetWare 6.5 SP8) to improve write performance, see *Tuning OES NetWare for Write Performance* (Technical Information Document 3660811) (http://www.novell.com/support/search.do?searchString=3660811).

Tuning NSS Performance on Linux 31

This section describes how to tune the Novell® Storage Services™ cache buffers to improve performance on a Novell Open Enterprise Server 2 Linux server.

- Section 31.1, “Tuning Cache Buffers for NSS on Linux,” on page 413
- Section 31.2, “Configuring High Memory Cache,” on page 414
- Section 31.3, “Configuring or Tuning Group I/O,” on page 415

31.1 Tuning Cache Buffers for NSS on Linux

- Section 31.1.1, “Understanding How NSS Uses Cache on Linux,” on page 413
- Section 31.1.2, “Setting the Minimum Number of Cache Buffers to Use for Kernel Memory,” on page 413

31.1.1 Understanding How NSS Uses Cache on Linux

NSS manages cache buffers on Linux using methods similar to those used in other Linux file systems such as Reiser, Polyserve, and XFS, with the exception of EXT.

For file data, NSS uses the Linux cache page manager to gain access to 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. This is similar to the cache handling used in NetWare®.

For metadata, NSS uses kernel memory. NSS can use only a percentage of this space because other applications share this space. By default, NSS reserves a minimum buffer cache size of 30,000 4-KB buffers, which is about 120 MB of the kernel memory space. You can adjust the minimum number of buffers to be used by NSS with the MinBufferCacheSize parameter.

For a 32-bit machine, the kernel cache memory limit is 1 GB cache. Depending on what else is running, you might need to modify how much space you allocate for NSS.

For example, when running `ravsui(8)` for a pool verify or a pool rebuild, the utility needs contiguous space in kernel memory separate from the space allocated to the core NSS process. The larger the pool, the larger the space that is needed. On a 32-bit machine with a 1 GB limit, you might need to stop other processes temporarily to free up space so that the verify or rebuild process can run. You can optionally modify the amount of space used by the core NSS process by lowering the setting for MinBufferCacheSize to as little as 10000 4-KB buffers. When the verify or rebuild is done, you can change the setting back to its normal setting.

31.1.2 Setting the Minimum Number of Cache Buffers to Use for Kernel Memory

1. Open a terminal console as the `root` user.
2. Start `nsscon(8)`. At the console prompt, enter

   ```
   nsscon
   ```

3. Set the minimum number of cache buffers used by NSS on Linux. In `nsscon`, enter

`nss /MinBufferCacheSize=value`

where *value* is the number of 4-KB buffers to assign for NSS. The default value is 30000.

The maximum setting is the amount of memory in KB divided by 4 KB. For a 32-bit machine the maximum setting is 250000 buffers.

31.2 Configuring High Memory Cache

Linux limits the low memory to less than 1GB on 32-bit machines, but places all physical memory in the low-memory area for 64-bit machines. By default, NSS on Linux caches its metadata blocks in the server's low memory. When the working set of metadata does not fit into low memory, performance is affected. In particular, metadata caching on servers with large numbers of files can create performance issues for servers that are low memory constrained as for 32-bit machines.

NSS for Linux supports metadata caching in high memory for servers that are low-memory constrained. This can greatly improve metadata reads on 32-bit systems with large numbers of files.

31.2.1 Configuring High Memory Cache Type and Size

Enter the following command in the NSS Console (`nsscon`):

nss /HighMemoryCacheType=[private | linux | none]

Set the type of use of User memory in 32-bit OES Linux as `none`, `linux`, or `private`. On 64-bit OES Linux, NSS does not use User memory at all.

Cache Type Options

private

Set up a separate user-mode cache in memory outside of the Linux cache. This is the default for 32-bit machines. By default, NSS takes 20 percent of the total high memory. You can modify this setting with the `nss /PrivateCacheSizeInBlocks=value` command.

However, If the number of Private Cache blocks is less than twice the NSS Cache blocks, then the default high memory cache type is `linux`. If the high memory is less than twice the NSS Cache blocks, then the default high memory cache type is `none`.

Private Cache is best used for dedicated file-servers. The Private Cache increases the performance of NSS, but reduces the amount of memory available to other file systems and tasks.

linux

Integrate caching into the regular Linux caches. This is the default for 32-bit machines where the number of Private Cache blocks is less than twice the NSS Cache blocks.

This can be a problem on a dual purpose file-server and application-server system, because memory-hungry applications can cause the file-system cache to purge completely.

none

Use the same algorithm as the initial release of OES 1 Linux, which is to try and cache everything in Kernel-mode memory. This is the only option on 64-bit OES Linux. This is the default for 32-bit machines if the high memory is less than twice the NSS Cache blocks.

nss /PrivateCacheSizeInBlocks=*value*

Sets the number of metadata blocks to cache for a 32-bit OES Linux machine where you have set the HighMemoryCacheType=private.

By default, NSS takes 20% of all high memory to use for the Private Cache. We recommend against giving NSS 100 percent of the high memory. The actual percentage you should allot depends on your deployment scenario. For example:

- Which features of NSS are enabled--salvage, compression, and so on.
- Use of really long filenames.
- Use of extended attributes.

Default: 20 percent of all high memory blocks

Range: 0 blocks to the number of blocks available in high memory (total available high memory divided by 4KB)

31.2.2 Monitoring the High Memory Cache Use

To view the current settings for the high memory cash usage, enter

```
nss /status
```

View the following parameters in the Current NSS Status report:

```
Current NSS Status:
  ... (parameters omitted)
  Current Private Cache Size          = 387680k (96920)
  Private Cache Size Limit            = 596364k (149091)
  High Memory Cache Type              = private
 ... (parameters omitted)
```

The Current Private Cache Size indicates the amount of memory currently being used by the Private Cache. The Private Cache Size Limit shows the maximum amount of memory that the Private Cache can use. The High Memory Cache Type indicates the type of high memory cache that is being used as `none`, `linux`, or `private`.

31.3 Configuring or Tuning Group I/O

Group write is a technique of writing data to the volume at regular intervals in order to reduce the seek time on the drive. It also reduces the number of writes because more changes to the same block are made only to memory.

In NSS for OES 2 NetWare and earlier versions of NSS on Linux and NetWare, NSS writes are done on a block-based timer. A block is written one second after the block becomes *dirty* (modified by a user or process). This can cause lots of head movement because there is no control over the order of blocks being sent to disk.

In OES 2 Linux, NSS performs group writes in three categories: journal, metadata, and user data. By setting policies for group writes, you can improve the performance of the file system for your particular environment.

For information, see the following:

- Section 31.3.1, “Viewing the Metadata Area Size,” on page 416
- Section 31.3.2, “Configuring the Journal Group Write Timer,” on page 417
- Section 31.3.3, “Configuring the Metadata Group Write Timer and Limit,” on page 417
- Section 31.3.4, “Configuring the User Data Group Write Timer,” on page 419
- Section 31.3.5, “Viewing Group Write Policies,” on page 419

31.3.1 Viewing the Metadata Area Size

NSS for OES 2 Linux provides a logical read-ahead capability. NSS is designed to physically store logically related data near each other, such as files in the same directory. By reading ahead using the logical information, performance is increased. When a block is read, its logically related blocks are also read. The area read is determined by the default area size.

To improve performance for NSS on OES 2 Linux, metadata blocks use an area seed logic to make sure that related metadata blocks are physically stored near each other. The default area size for metadata blocks is 16 blocks that are 4 KB each, or 64 KB total.

For metadata blocks, the seed is set to the block number for the area. When metadata is written, the seed logic determines the closest free block in the area to use next. When the area is new, a new free area is found in a higher area in the pool, and a new seed marks this area. When the search for a free area reaches the end of the pool, it wraps back to start searching for free areas to use at the start of the pool. If no free space of sufficient size is found, the size is temporarily halved from 16 to 8, 4, 2,or 1 blocks progressively as needed until the temporary size is 1. A setting of 1 block indicates that the pool is essentially out of space. As space is freed or the pool increases in size, future space allocations use the default area size of 16 blocks

The maximum number of dirty data blocks that are allowed to accumulate is governed by the Metadata Group Write Limit parameter. By default, the limit is 20000 dirty blocks. For information, see Section 31.3.3, “Configuring the Metadata Group Write Timer and Limit,” on page 417.

You can view the metadata area size that is currently in use and the number of dirty blocks waiting to be written by viewing the Current Metadata Group Write Size parameter in the NSS status report. The information is reported in the following format:

```
Current Metadata Group Write Size  = areasize (number_dirty_blocks)
```

For example, with the default setting of 16 4-KB blocks, the metadata area is 64 KB. If 16000 dirty blocks are waiting to be written, the values are reported follows:

```
Current Metadata Group Write Size  = 64K (16000)
```

To view the Current Metadata Group Write Size information:

1. Open a terminal console, then log in as the `root` user.
2. At the console prompt, open the NSS Console by entering

   ```
   nsscon
   ```
3. At the `nsscon` prompt, enter

```
nss /status
```

4 In the NSS status report, look for the Current Metadata Group Write Size parameter to view the current values:

```
Current Metadata Group Write Size  = areasize (number_dirty_blocks)
```

31.3.2 Configuring the Journal Group Write Timer

For NSS, the journal keeps metadata consistent up to the time when its blocks are written to the device. The Journal Group Write Timer determines the elapsed time between writes of journal blocks. Thus, its timer policy determines how long ago that a consistent point is relative to a system crash.

Journal blocks are written by default as a group every second. Journal blocks might be written sooner than the one-second elapsed time if another timer policy triggers a write or if the journal gets full before the time elapses. Writing blocks as a group helps improve performance because it allows fewer writes, while ensuring that data is actually recorded to the device.

Use the following NSS command option to control the group write policy for journal blocks:

/JournalGroupWriteTime=*seconds*

Use the JournalGroupWriteTime parameter to specify the elapsed time to wait before group writes of journal blocks.

Journal Group Write Timer	Risk for Inconsistent Metadata (Time Elapsed Since Last Consistent Point)	File System Performance
1 second (default)	Minimized	Optimized for most scenarios
Greater than 1 second	Higher	Faster

To set the JournalGroupWriteTime parameter, issue the following command as the `root` user in the NSS Console (`nsscon`):

```
nss /JournalGroupWriteTime=seconds
```

Replace *seconds* with the maximum number of seconds to elapse before forcing journal blocks to be written to the volume. The default value of *seconds* is 1.

For example, to group write journal blocks every 2 seconds, enter

```
nss /JournalGroupWriteTime=2
```

31.3.3 Configuring the Metadata Group Write Timer and Limit

The metadata blocks are written by default as a group every 40 seconds, or when the MetadataGroupWriteLimit is reached, whichever occurs first. Metadata loss does not occur if the system crashes because all metadata changes are automatically recorded in the journal. However, increasing the timer setting increases the redo/undo time that is required to activate a pool (the mount time) after a crash because there is more unwritten metadata in the journal to be resolved.

IMPORTANT: Within a clustered environment, this means that the time to complete a failover is related to the setting of MetadataGroupWriteLimit parameter.

You can limit the amount of time it takes for a pool activation after a crash by decreasing the maximum number of metadata blocks that can be dirty in the MetadataGroupWriteLimit parameter. A group write is performed when the limit is reached.

You can increase performance of the file system by increasing the maximum number of metadata blocks that can be dirty.

Use the following NSS command options to control the group write behavior for metadata blocks:

/MetadataGroupWriteTime=*seconds*

Use the MetadataGroupWriteTime parameter to specify the elapsed time to wait before group writes of metadata blocks. Decreasing the metadata group write timer can help reduce the mount time for the volume after a crash.

Metadata Group Write Timer	Time to Mount After a System Crash	File System Performance
40 seconds (default)	Optimized for most scenarios	Optimized for most scenarios
Less than 40 seconds	Faster	Slower
More than 40 seconds	Slower	Faster

To set the MetadataGroupWriteTime parameter, issue the following command as the `root` user in the NSS Console (`nsscon`):

```
nss /MetadataGroupWriteTime=seconds
```

Replace *seconds* with the maximum number of seconds to elapse before forcing metadata blocks to be written to the volume. The default value of *seconds* is 40.

For example, to group write metadata blocks every 30 seconds, enter

```
nss /MetadataGroupWriteTime=30
```

/MetadataGroupWriteLimit=*blocks*

Use the MetadataGroupWriteLimit parameter to specify the maximum number of metadata blocks that can be dirty before a group write is performed. The following describes how the settings affect time to mount and file system performance:

Maximum Number of Dirty Metadata Blocks	Time to Mount After a System Crash	File System Performance
20000 blocks (default)	Optimized for most scenarios	Optimized for most scenarios
Less than 20000 blocks	Faster	Slower
More than 20000 blocks	Slower	Faster

To set the MetadataGroupWriteLimit parameter, issue the following command as the `root` user in the NSS Console (`nsscon`):

```
nss /MetadataGroupWriteLimit=blocks
```

Replace *blocks* with the maximum number of metadata blocks that can be dirty before forcing them to be written to the volume. The default value of *blocks* is 20000.

For example, to decrease the maximum number of dirty metadata blocks to 15,000 for the purpose of reducing the mount time, enter

```
nss /MetadataGroupWriteLimit=15000
```

For example, to increase the maximum number of dirty metadata blocks to 30,000 for the purpose of increasing the file system performance, enter

```
nss /MetadataGroupWriteLimit=30000
```

31.3.4 Configuring the User Data Group Write Timer

The user data blocks are written as a group every 3 seconds. This increases the risk of data loss on a crash compared to previous versions of NSS that write data blocks every 1 second. You can set the user data group write timer (UserDataGroupWriteTime) to 1 second to get the familiar NSS behavior for data writes.

Use the following NSS command option to control the group write behavior for user data blocks:

/UserDataGroupWriteTime=*seconds*

Use the UserDataGroupWriteTime parameter to specify the elapsed time to wait before group writes of user data blocks. Decreasing the user data group write timer can help reduce the risk of data loss for a volume after a crash.

User Data Group Write Timer	Risk of Data Loss After a Crash	File System Performance
3 seconds (default)	Optimized for most scenarios	Optimized for most scenarios
1 second	Lower, typical of NSS on NetWare and OES 1 Linux and NetWare	Slower
Greater than 3 seconds	Higher	Faster

To set the UserDataGroupWriteTimer parameter, issue the following command as the `root` user in the NSS Console (`nsscon`):

```
nss /UserDataGroupWriteTime=seconds
```

Replace *seconds* with the maximum number of seconds to elapse before forcing user data blocks to be written to the volume. The default value of *seconds* is 3.

For example, to group write user data blocks every 1 second, enter

```
nss /UserDataGroupWriteTime=1
```

31.3.5 Viewing Group Write Policies

1. Open a terminal console, then log in as the `root` user.
2. At the terminal console prompt, enter

```
nsscon
```

3 In nsscon, enter

```
nss /status
```

4 Look for the following settings in the *Current NSS Status* report:

```
Journal Flush Timer               = 1 second
Metadata Flush Timer              = 40 seconds
User Data Flush Timer             = 3 seconds
Current Metadata Group Write Size  = 64k (16)
Metadata Block Group Write Limit   = 80000k (20000)
```

Monitoring the Status of the NSS File System and Services

32

This section describes the following methods for monitoring the status of Novell® Storage Services™:

- Section 32.1, “Monitoring Status of NSS Devices, Pools, and Volumes with iManager,” on page 421
- Section 32.2, “Monitoring Compression and Salvage Statistics,” on page 422
- Section 32.3, “Monitoring Quotas and Space Usage for NSS Pools and Volumes,” on page 424
- Section 32.4, “Monitoring File System Parameters,” on page 426

32.1 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	Section 10.4, “Viewing Details for a Device,” on page 132
Partitioned free space	Section 10.9, “Viewing Partitions on a Device,” on page 136
Partition	Section 12.4, “Viewing Details for a Partition,” on page 160
Software RAIDs	Section 13.3, “Viewing a List of Software RAID Devices on a Server,” on page 171
Software RAID details	Section 13.4, “Viewing Details of a Software RAID Device,” on page 172
NSS pools	Section 16.7, “Viewing Pools on a Server,” on page 204
NSS pool details	Section 16.8, “Viewing Pool Details,” on page 205
NSS volumes in a pool	Section 16.10, “Viewing Volume Information for a Pool,” on page 207
NSS volume details	Section 19.6, “Viewing the Details of an NSS Volume,” on page 263
NSS volume quota and space usage	Section 19.7, “Viewing Properties of an NSS Volume,” on page 263, then click the Quotas tab

To monitor the status of:	Refer to:
NSS volume attributes	Section 19.7, “Viewing Properties of an NSS Volume,” on page 263, then click the *Attributes* tab
File and folder properties	Section 22.1, “Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes,” on page 297

32.2 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 9.1.3, “Accessing Roles and Tasks in iManager,” on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, “Selecting a Server to Manage,” on page 104.
3. In the *Volumes* list, select a volume that you want manage.
4. Click *Properties*.

   The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota 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.

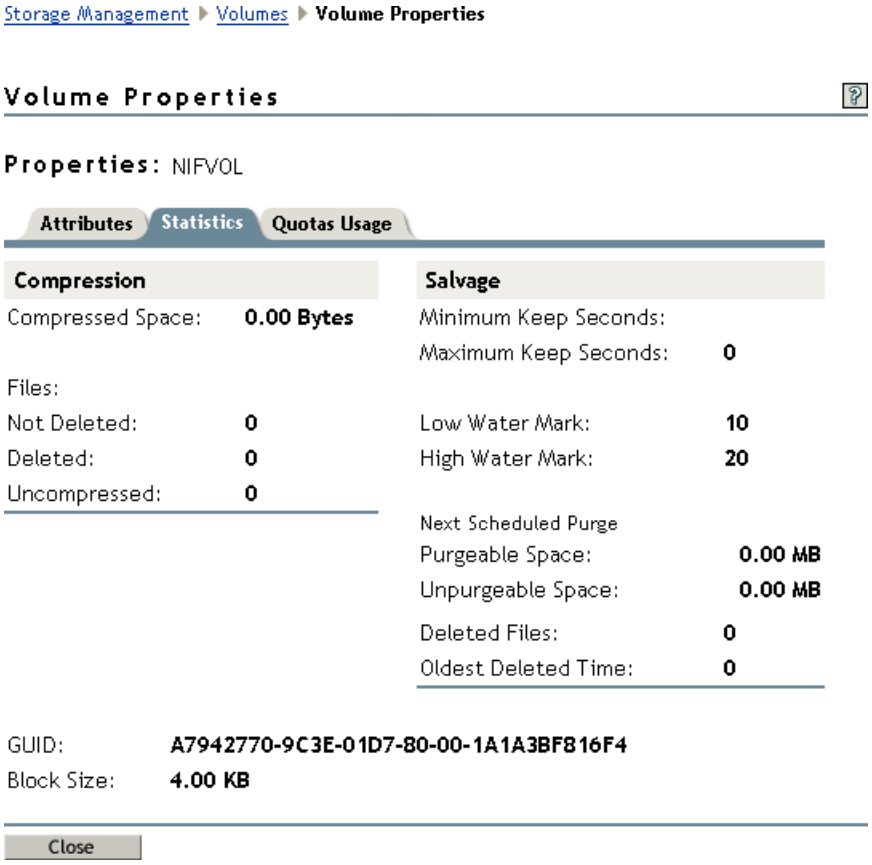

If the compression attribute is set, the Compression report shows statistics of all the compressed files for the selected volume.

Statistic	Description
Compressed Space	The amount of space in the volume in use by compressed files.
Files	The total number of files in the volume and information by the following categories: ◆ **Not Deleted:** The total number of files in the volume that are currently available to users. ◆ **Deleted:** The total number of files in the volume that are deleted but not yet purged from the system. ◆ **Uncompressed:** The total number of files in the volume that are not stored in compressed form.

If the Salvage Files attribute is enabled, the Salvage report shows statistics about deleted files that have not yet been purged.

Statistic	Description
Minimum Keep Seconds	Minimum time (in seconds) to keep deleted files.
Maximum Keep Seconds	Maximum time (in seconds) to keep deleted files.

Statistic	Description
Low Water Mark	If the amount of free space drops below this percentage, the file system begins purging deleted files.
High Water Mark	If there are files to delete, the autopurging process stops when the amount of free space reaches this percentage.
Next Scheduled Purge	Date and time of the next purge.
Purgeable Space	Amount of space in the volume that is occupied by deleted files that are queued for purging.
Unpurgeable Space	Amount of space in the volume that is occupied by files.
Deleted Files	The number of deleted files in salvage.
Oldest Deleted Time	Time line for deleted files. The file system purges the files in the same order they were deleted.
GUID	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.
Block Size	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
GUID	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.
Block Size	The maximum amount of data committed to a single write process. Possible sizes include 4, 8, 16, 32, or 64 KB.

32.3 Monitoring Quotas and Space Usage for NSS Pools and Volumes

For each volume, the *Volume Properties* > *Quota 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 9.1.3, "Accessing Roles and Tasks in iManager," on page 104.
2. Select a server to manage.

   For instructions, see Section 9.1.4, "Selecting a Server to Manage," on page 104.
3. In the *Volumes* list, select a volume that you want manage.
4. Click Properties.

   The *Properties* page has three tabs: *Attributes*, *Statistics*, and *Quota Usage*. It opens by default to the *Attributes* tab.
5. Click the *Quota Usage* tab to view the current space usage statistics for the selected volume.

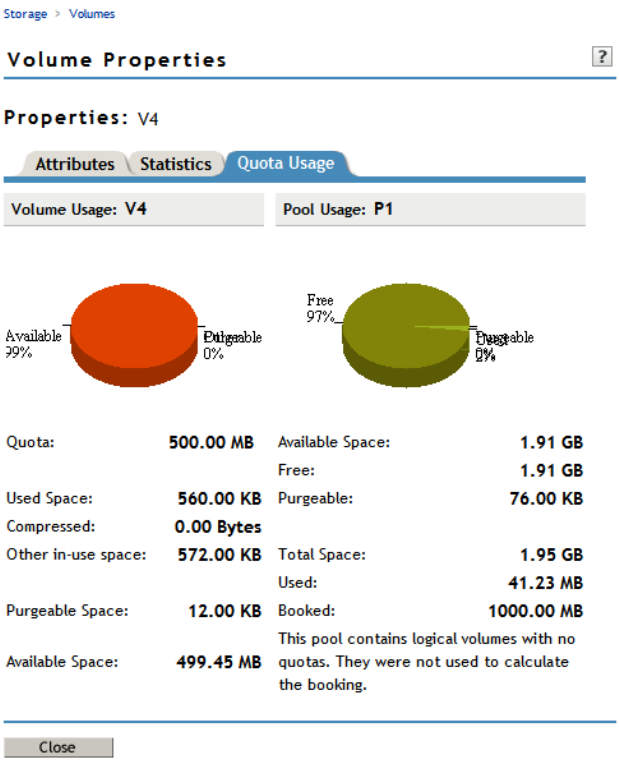

Volume Usage reports the amount of space on the volume, categorized by usage: Compressed, Other In-use, Purgeable, and Available.

Statistic	Description
Quota	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.
Used Space	The amount of space currently in use and information by the following categories: ◆ **Compressed:** If the Compressed attribute is enabled, this is the amount of space in the volume containing data that is compressed. ◆ **Other In-Use Space:** The amount of space in the volume containing data that is not compressed. ◆ **Purgeable Space:** The amount of space in the Salvage system that you can use as free space. You can manually purge deleted files to free space. ◆ **Available:** Available free space that is not in the Salvage system.

Pool Usage reports the amount of space on the pool, categorized by usage: Free, Purgeable, Used, and Booked.

Statistic	Description
Available Space	The amount of space in the pool that is not currently in use and information by the following categories: ◆ **Free:** The total amount of free space that is available on the pool. ◆ **Purgeable:** The amount of space in the Salvage system that you can use as free space. You can manually purge deleted files to free space.
Total Space	The total amount of space allocated to the pool and information by the following categories: ◆ **Used:** The total amount of space currently in use by all volumes on the pool. ◆ **Overbooked/Booked:** 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.

32.4 Monitoring File System Parameters

- Section 32.4.1, "Using iManager to Monitor NSS File System Parameters," on page 426
- Section 32.4.2, "Using Novell Remote Manager to Monitor NSS File System Parameters," on page 426
- Section 32.4.3, "Using Novell NetStorage to Monitor NSS File System Parameters," on page 427

32.4.1 Using iManager to Monitor NSS File System Parameters

Use the File Manager plug-in to iManager to browse files and directories, and to manage access control for them. For information, see the following:

- Section 22.1, "Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes," on page 297
- Chapter 28, "Managing Files and Folders on NSS Volumes," on page 387

32.4.2 Using Novell Remote Manager to Monitor NSS File System Parameters

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 *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*.

Use this link	To do these tasks
Volumes	◆ Monitor or restrict disk space usage by directory (directory quotas) ◆ Purge or salvage deleted volumes or files ◆ Configure file system trustee rights and attributes for directories and files ◆ View or change a list of Set parameters

You can browse directories and files on Linux by using Novell Remote Manager for Linux. For information, see " Viewing File Systems " in the *OES 2 SP1: Novell Remote Manager for Linux Administration Guide*.

32.4.3 Using Novell NetStorage to Monitor NSS File System Parameters

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	◆ Monitor or restrict disk space usage by directory (directory quotas) on Linux ◆ Monitor or restrict disk space usage by directory (directory quotas) on NetWare ◆ Purge or salvage deleted volumes or files ◆ Configure file and directory Trustee rights and Inherited Rights filters ◆ View or change a list of Set parameters

Troubleshooting the NSS File System

33

This section describes some issues you might experience with Novell® Storage Services™ (NSS) and provides suggestions for resolving or avoiding them.

For additional troubleshooting information, see the Novell Support Web site (http://support.novell.com).

- Section 33.1, “Cannot Connect to Target Servers from iManager,” on page 430
- Section 33.2, “Cannot Create an NSS Pool or Volume,” on page 430
- Section 33.3, “Cannot See NSS Devices, Pools, or Volumes,” on page 431
- Section 33.4, “Data Corruption Is Possible with Xen-Based File-Backed Block Devices on NetWare Virtual Machines,” on page 432
- Section 33.5, “eDirectory Errors When Creating NSS Pools or Volumes on Linux,” on page 433
- Section 33.6, “File Compression Is Not Working,” on page 434
- Section 33.7, “Linux Fails to Execute Dismount, Umount, or Delete Commands for NSS Volumes,” on page 434
- Section 33.8, “Multipath Devices Are Not Resolved,” on page 434
- Section 33.9, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux),” on page 434
- Section 33.10, “NSS Volume Disappears from the NCP Console (ncpcon) (Linux),” on page 435
- Section 33.11, “Pathname Exceeds 255-Character Limit,” on page 436
- Section 33.12, “Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers,” on page 436
- Section 33.13, “Server Hangs When Using an NSS Volume as a Netatalk Share,” on page 436
- Section 33.14, “Slow Mount Performance for NSS Volumes Using the UNIX Name Space (Linux),” on page 437
- Section 33.15, “Software RAID 1 Fails to Recognize a Replacement Device,” on page 437
- Section 33.16, “Tuning NSS Volumes for GroupWise Servers,” on page 437
- Section 33.17, “Unknown Users as File Owners,” on page 437
- Section 33.18, “Using Scripts to Create or Manage NSS Pools and Volumes on Linux,” on page 438

33.1 Cannot Connect to Target Servers from iManager

If you are having difficulty connecting to servers from iManager, it might be because you are using an unsupported protocol between the iManager server and the target server that you want to manage. For information, see "Interoperability of Protocols for the iManager Server and Target Server" on page 90 under the Section 7.5, "Cross-Platform Issues for Management Tools," on page 90.

33.2 Cannot Create an NSS Pool or Volume

If you cannot create an NSS pool or volume, consider these possible causes:

- Section 33.2.1, "General Issues," on page 430
- Section 33.2.2, "Linux Issues," on page 430
- Section 33.2.3, "NetWare Issues," on page 431

33.2.1 General Issues

- 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 10.9, "Viewing Partitions on a Device," on page 136.
- 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.

33.2.2 Linux Issues

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 LVM2 instead of EVMs, NSS does not recognize the device as being available and does not create the pool.

SUSE® Linux Enterprise Server uses LVM2 as the default device manager, 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 Linux with EVMS as the default volume manager

  See Mounting Your Root File System through EVMS (http://evms.sourceforge.net/install/root.html) for more information.

- Convert LVM2 volumes to EVMS volumes after Linux installation.

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

- Patch the 2.6 kernel on your Linux server to allow multiple device managers.

  WARNING: This solution is not recommended, tested, nor supported, but it is mentioned here for completeness.

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

33.2.3 NetWare Issues

If you cannot create an NSS volume on your NetWare® 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.

33.3 Cannot See NSS Devices, Pools, or Volumes

- Section 33.3.1, “Connection Failure,” on page 431
- Section 33.3.2, “NSS Tools Recognize Only EVMS-Managed Devices (Linux),” on page 431
- Section 33.3.3, “EVMS Not Running (Linux),” on page 432
- Section 33.3.4, “Pools Not Mounted by EVMS (Linux),” on page 432

33.3.1 Connection Failure

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.

33.3.2 NSS Tools Recognize Only EVMS-Managed Devices (Linux)

For NSS on 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. 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 NSS on Linux, whenever you reboot your system, EVMS discovers and mounts the device nodes under `/dev/evms/`. 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`. 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.

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/`.

33.3.3 EVMS Not Running (Linux)

NSS on Linux requires that EVMS be installed and running. If it is not running, you cannot see NSS devices, pools, or volumes on Linux.

To start EVMS:

1. Open a terminal console on the server, then log in as the `root` user.
2. At the console prompt, enter

   ```
   evms_activate
   ```

33.3.4 Pools Not Mounted by EVMS (Linux)

Randomly after a crash, you might find that EVMS fails to activate an NSS pool on a server restart and the pool and its volumes are not available. Make sure EVMS is running by entering `evms_activate` at a terminal console prompt, then use iManager or NSSMU to manually activate the pool and mount its volumes.

33.4 Data Corruption Is Possible with Xen-Based File-Backed Block Devices on NetWare Virtual Machines

NSS uses barrier flushes to ensure that metadata has actually been written on devices that cache writes and mis-inform the file system that the I/O has completed. Using barrier flushes helps prevent data corruption that can occur during a system crash because all confirmed writes are actually known to be on the device. This works as expected on physical devices.

Xen-based file-backed block devices do not support barrier flushes in the Xen version released in OES 2 Linux. If the system crashes, data corruption can occur because the flushes are not actually performed. We recommend against using file-backed devices for virtual machines. For information, see Section 6.1.2, "Virtual Machine Issues," on page 81.

You can determine whether NSS is using barrier flushes by looking for the following lines in the `/_admin/Manage_NSS/Pool/`*`poolname`*`/ZLSS/PhysicalIO.xml` file:

```
<barrierSetting value="1"/>
<barrierCommandLine value="1"/>
<barrierFlushes value="673"/>
```

If the barrierSetting value is true (1), then the pool is currently using barriers.

The barrierCommand value indicates whether the user turned off barriers when mounting the pool.

The barrierFlushes value is a persistent count of the number of times a barrier flush has been called, including calls that get returns of `Not Supported`.

33.5 eDirectory Errors When Creating NSS Pools or Volumes on Linux

- Section 33.5.1, "eDirectory Error 613 When Creating an NSS Pool or Volume on Linux," on page 433
- Section 33.5.2, "eDirectory Error 672 When Creating an NSS Pool on Linux," on page 433

33.5.1 eDirectory Error 613 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™ Read/Write replica. The error occurs because NCP™ (NetWare Control Protocol™) 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 `/etc/opt/novell/ncpserv.conf` file for NCP Server on Linux.

33.5.2 eDirectory Error 672 When Creating an NSS Pool on Linux

When creating an NSS pool on Linux by using NSSMU or iManager, an Error 672 is returned if there is no NSS Admin object in the Novell eDirectory database for the server (such as *`HOSTNAMEadmin.context`*). NSS requires that an NSS Admin object must exist for each and every server, or management does not work.

This situation occurs if you move a server across trees without also moving its NSS Admin object from one tree's eDirectory database to the other.

If you re-create the NSS Admin object, you are then able to successfully create pools.

To re-create the NSS Admin object, run `nssAdminInstall` at a Linux terminal console as the `root` user:

1. Open a terminal console, then log in as the `root` user.
2. At the console prompt, enter the following (all on the same line, of course):

   `nssAdminInstall -a adminname.context -p adminpassword -o HOSTNAMEadmin.context`

   For example, for a server with a hostname of *`SERVER1`* in the context of *`example`*, the NSS Admin object might be named *`SERVER1admin.example`*.

   `nssAdminInstall -a admin.example -p admin_pwd -o SERVER1admin.example`

After the NSS Admin object is created, update the eDirectory Pool object. For information, see Section 16.13, "Updating eDirectory Pool Objects," on page 213.

33.6 File Compression Is Not Working

If you cannot compress a file, check the following:

- Check the Compression attribute for the volume to make sure it is enabled. To apply the file compression option to an existing NSS volume: From iManager, click *Storage* > *Volumes* > *Properties* > *Attributes*, select *Compression*, then click *OK*.
- Check to see if the Do Not Compress (dc) attribute for the file or directory is set.
- Review other factors that affect compression, as detailed in "Factors Affecting Compression" on page 329.

33.7 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.

33.8 Multipath Devices Are Not Resolved

If you have multiple interconnect paths in your server-to-storage configuration, you must use multipath management software to resolve the multiple paths into a single mutlipath device.

On NetWare, see "Managing Multipath I/O to Devices (NetWare)" on page 189.

On Linux, see Chapter 15, "Managing Multipath I/O to Devices (Linux)," on page 195.

33.9 NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)

NSS requires the NDP user space module (`ndpapp`) to be loaded and running when NSS starts. If ndpapp is not running, modules in NSS that attempt eDirectory operations fail and prevent NSS from loading.

In some environments, when the NDP module (`ndpmod`) attempts to register the `/dev/ndp` device, the kernel routine misc_register() registers the device inside the kernel, but does not make it available in user space until about 17 seconds later. Because of the delay, the NDP user space

module kills itself for about 10 seconds. NSS cannot start until `ndpapp` reloads itself. When these events occur, you see a message in the `/var/log/messages` file similar to the following that indicates the `/dev/ndp` is not available in user space:

```
Sep  1 19:37:54 servername ndpapp[10052]: stat: /dev/ndp: 2: No such file or
directory
```

Increasing the UDEV event handling limits to 1024 has been tested and circumvents this problem. The `/dev/ndp` device gets created and is seen by ndpapp so that NSS can start.

1. Log in as the `root` user.
2. Open `/etc/sysconfig/udev` in a text editor, then locate the following lines:

   ```
   UDEVD_MAX_CHILDS=64
   UDEVD_MAX_CHILDS_RUNNING=16
   ```

3. Specify 1024 as the value for each parameter:

   ```
   UDEVD_MAX_CHILDS=1024
   UDEVD_MAX_CHILDS_RUNNING=1024
   ```

4. Save the file.
5. Reboot the server.

33.10 NSS Volume Disappears from the NCP Console (ncpcon) (Linux)

An NSS volume might not be found by the `ncpcon` utility if the volume's mount point is renamed (the directory path is changed). For example, this might occur if you change the mount point's directory path from `/media/nss/`*`volumename`* (the default path) to `/`*`targetpath/volumename`*.

When an NSS volume is created, the NCP server gets the path to the volume and caches it, assuming that it never changes. When you later run the NCP Console (`ncpcon`) utility and enter the `volume` command, it reports the volumes that are still found at their respective mount points. Only the volumes that are still valid as compared to the list in cache are reported.

Use the following methods to avoid or resolve this problem:

- If you know in advance that you want to modify default path of the volume's mount point, make sure to create the NSS volume with iManager and change the default mount point as you configure the volume. For example, if you are setting up an NSS volume to use for the home directories, change the mount point from `/media/nss/home` to `/home`.

  Creating the volume in NSSMU does not allow for the mount point to be changed during the volume setup, only afterwards.

- To modify the mount point for an existing volume, you can use either iManager or NSSMU to change it, then restart eDirectory by entering the following at a terminal console prompt:

  ```
  /etc/init.d/ndsd restart
  ```

  Restarting eDirectory causes the NCP volume cache to be updated.

33.11 Pathname Exceeds 255-Character Limit

Pathnames for files on the server can have up to 255 characters, including the server name, volume name, path delimiters, filename, and file extension. If a user maps a drive letter to a location deep down the directory path, and subsequently creates a pathname that exceeds the limit for the name on the server, the file cannot be saved. Even if the file's path does not exceed 255 characters with respect to the mapped drive letter, it exceeds the maximum length on the server.

When mapping a drive letter to a folder deep down the directory path, users should adopt conventions for naming files and directories and for creating directory structures so that they do not exceed the 255-character pathname limit on the server.

33.12 Problems Moving NetWare 6.5 NSS Pools to NetWare 6.0 Servers

NetWare 6.5 SP3 and earlier 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 the initial release 6.5, but the media format is the same through NetWare 6.5 SP3. In NetWare 6.5 SP4 and later, you can optionally upgrade the NSS media format for enhanced hard links support.

If you move a device from the NetWare 6.5 SP3 or earlier 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 has problems with it because the information is different. You should remove the mirror before trying to move the disk to NetWare 6.0.

If you have upgraded the media format for hard links on a NetWare 6.5 SP4 or later server, you cannot move its NSS pool to a NetWare 6.0 server, or to any platform where the enhanced media format is not supported.

33.13 Server Hangs When Using an NSS Volume as a Netatalk Share

The CopyCat application used by Netatalk uses sparse files for its database. Netatalk tries to create a CopyCat database as a sparse file called `.AppleDB` in the root of the volume by using memory mapped IO. This can cause the server to hang if you are using an NSS volume as the Netatalk share because of the limited support in NSS for this combination.

NSS has limited support for memory mapped files, primarily to support loading programs. NSS does not fully support memory mapped files especially if the application uses sparse files.

33.14 Slow Mount Performance for NSS Volumes Using the UNIX Name Space (Linux)

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 19.11, “Mounting NSS Volumes with Linux Commands (Linux),” on page 268.

33.15 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 re-scan 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.

33.16 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
```

33.17 Unknown Users as File Owners

In an NSS volume, a file’s owner is the user who created it. The Novell trustee model for file systems is used to control access to files, so ownership is a consideration only when enforcing user quotas.

Ownership is tracked with the user’s GUID, not the username. If the username ever becomes invalid (such as if the user is deleted from the system), the file continues to be charged to that user’s GUID. In a space usage report, the value of the GUID appears in place of where the a valid username would normally be. There is no limit on the space that can be associated with unknown users. Authorized users can continue to use the files without interruption or incident.

For the user space quota, the total disk space used by the file continues to be associated with the file’s assigned owner’s GUID. User quotas can be enforced only for valid users.You must change the file’s owner to a valid user if you want the files to be included in that user’s quota.

An administrator or administrator equivalent user can assign a new owner when necessary. Changing file ownership requires the Supervisor right for the file's parent directory and the file. Use whatever tools you normally use to modify the file's ownership.

33.18 Using Scripts to Create or Manage NSS Pools and Volumes on Linux

There is an XML interface that allows you to write scripts, such as Perl scripts, that will create and manage NSS pools and volumes. The API set can be downloaded from *NDK: Virtual File Services for NetWare* (http://developer.novell.com/wiki/index.php/Virtual_File_Services_for_NetWare). Although the title suggests that the Virtual File Services are only for NetWare, they are for NSS services on NetWare 6.0, NetWare 6.5, OES 1 Linux and NetWare, and OES 2 Linux and NetWare servers.

Look for sample scripts at the link above or in Cool Solutions. For example:

- *NSS Pool Lister for Linux* (http://www.novell.com/coolsolutions/tools/18074.html)
- *NSS Volume Lister for Linux* (http://www.novell.com/coolsolutions/tools/18082.html)

Security Considerations

34

This section describes security issues and recommendations for Novell® Storage Services™ for Novell Open Enterprise Server 2. It is intended for security administrators or anyone who is responsible for the security of the system. It requires a basic understanding of NSS. It also requires the organizational authorization and the administrative rights to effect the configuration recommendations.

- Section 34.1, “Security Features of NSS,” on page 439
- Section 34.2, “Preventing Unauthorized Access,” on page 441
- Section 34.3, “Securing Sensitive Information During Remote Management Sessions,” on page 441
- Section 34.4, “Protecting Data During Backup and on Backup Media,” on page 442
- Section 34.5, “Preventing Exposure of Sensitive Data in a Core Dump,” on page 442
- Section 34.6, “Preventing Exposure of the Encryption Password in a Log (Linux),” on page 442
- Section 34.7, “Using Data Shredding to Prevent Unauthorized Access to Purged Files,” on page 446
- Section 34.8, “Protecting the Admin Volume and Files.cmd (NetWare),” on page 447
- Section 34.9, “Acquiring Security Equivalence Vectors for NSS Users,” on page 447
- Section 34.10, “Protecting Modules Responsible for Security Equivalence Vectors (Linux),” on page 447
- Section 34.11, “Controlling File System Access and Attributes for NSS Volumes,” on page 447
- Section 34.12, “Displaying Directory and File Attributes for NSS Volumes on Linux,” on page 448
- Section 34.13, “Security Best Practices for zAPIs,” on page 448
- Section 34.14, “Controlling Physical Access to Servers and Resources,” on page 449
- Section 34.15, “Securing Access to the Servers with a Firewall,” on page 449
- Section 34.16, “Creating Strong Passwords,” on page 449

34.1 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.	“Managing Encrypted NSS Volumes” on page 281 Section 34.16, “Creating Strong Passwords,” on page 449

Issue/Feature	Description	Recommendation	For More Information
Storage plug-in for iManager	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 Administrator user identity and a strong password. Section 34.16, “Creating Strong Passwords,” on page 449	Section 9.1.5, “Storage Plug-In Quick Reference,” on page 105 *Novell iManager 2.7.1 Administration Guide*
Files and Folders plug-in for iManager	Manage trustees, trustee rights, and inherited rights filters. View effective rights. Manage file system attributes.	Use an Administrator user identity and a strong password. Section 34.16, “Creating Strong Passwords,” on page 449	Section 9.1.6, “Files and Folders Plug-In Quick Reference,” on page 110 Section 22.1, “Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes,” on page 297 *Novell iManager 2.7.1 Administration Guide*
NSSMU	NSSMU is executed locally on the server.	Secure physical access to servers.	Section 34.14, “Controlling Physical Access to Servers and Resources,” on page 449
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 Administrator user identity or equivalent, and use a strong password.	Section 9.5.4, “Accessing Novell Remote Manager,” on page 123 Section 34.16, “Creating Strong Passwords,” on page 449 *OES 2 SP1: Novell Remote Manager for Linux Administration Guide* *OES 2 SP1: Novell Remote Manager for NetWare Administration Guide*

34.2 Preventing Unauthorized Access

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

- **Encrypted volume support**

  Encrypted volume support encrypts the volume, which makes data inaccessible to software that circumvents normal access control, such as if the media were stolen. It meets U.S. Government security standards. For information, see "Managing Encrypted NSS Volumes" on page 281.

- **Data shredding**

  The Data Shredding attribute supports shredding of purged files (up to 7 times), which erases files completely. It meets the U.S. Government security standards. For information, see Section 22.3, "Using Data Shredding to Prevent Access to Purged Files," on page 306.

- **Multiple server access prevention for pools**

  Multiple server access prevention (MSAP) ensures data integrity by preventing unauthorized access to shared media in a storage area network. For information, see Section 16.12, "Preventing Pools from Activating on Multiple Servers," on page 208.

- **Novell Trustee model for access control**

  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 *OES 2 SP1: File Systems Management Guide*.

34.3 Securing Sensitive Information During Remote Management Sessions

When you are managing NetWare servers in iManager, any sensitive information (such as usernames and passwords) is encrypted with a shared-key technology before the sensitive data is passed between iManager and the `_admin` volume on the NetWare server you are managing. This ensures that sensitive data is not exposed during transmission. Other data is transmitted in clear text.

When you are managing OES-Linux and OES2-Linux servers in iManager, all information, including sensitive data, is typically sent via a Secure HTTP (HTTPS) connection between iManager and CIMOM on the Linux server you are managing. This ensures that sensitive data is not exposed during transmission. However, if CIMOM is not running on the Linux server you are managing, the plug-ins attempt to connect via NCP™ or CIFS. These connections are insecure and are a security concern only when transmitting sensitive information.

Effective in OES2-Linux, storage plug-ins have been modified to prevent this potential exposure of sensitive information. Where tasks involve the exchange of sensitive information between iManager and the Linux server you are managing, the plug-in now checks to see if CIMOM is running and available on the Linux server you are managing before it attempts to execute the command. If CIMOM is not running for some reason, it returns an error message and does not execute the task. The plug-ins do not allow sensitive data to be sent across insecure connections (such as NCP or CIFS/SAMBA) to the Linux server. You get an error message explaining that the connection is not secure and that CIMOM must be running before you can perform the task.

34.4 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.

34.5 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, select to exclude cache during a core dump when prompted to select writing all of memory in the core dump or to exclude NSS cache.

This applies also for volumes that are not encrypted but contain confidential data. Although the data is not normally encrypted, you might not want to allow unauthenticated access to the information.

34.6 Preventing Exposure of the Encryption Password in a Log (Linux)

This section describes NSS debugger logging features so you can identify when these logs are turned on and turn them off in your operational environment.

- Section 34.6.1, “NSS Logging,” on page 442
- Section 34.6.2, “NSS Logging and Security Implications,” on page 443
- Section 34.6.3, “Logging Communications between NSS and the _ADMIN Volume,” on page 443
- Section 34.6.4, “Logging Communications between NSS and eDirectory, NICI, or Linux User Management,” on page 445

34.6.1 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.

34.6.2 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 the 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.

34.6.3 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" on page 443
- "Enabling or Disabling adminusd Logging" on page 444
- "Enabling or Disabling VFS Logging" on page 444

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.

- "Enabling adminusd Logging" on page 444
- "Disabling adminusd Logging" on page 444

Enabling 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.

Disabling 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`.

34.6.4 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" on page 445
- "Enabling or Disabling ndpapp Logging" on page 445
- "Enabling or Disabling npdmod Logging" on page 446

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" on page 445
- "Disable ndpapp Logging" on page 446

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 npdmod Logging

- “Enabling ndpmod Logging” on page 446
- “Disabling ndpmod Logging” on page 446

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`.

34.7 Using Data Shredding to Prevent Unauthorized Access to Purged 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 22.3, “Using Data Shredding to Prevent Access to Purged Files,” on page 306.

34.8 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 or modify its properties. You should set trustee rights to prevent users from seeing this volume. For information, see Section 22.5, “Controlling Access to the _admin Volume and Files.cmd (NetWare),” on page 311.

34.9 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 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 A.37, “Security Equivalence Vector Update Commands (Linux),” on page 490.

34.10 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 the `root` user. For example, make `root` the owner and set permissions to restrict access for Group and Other users.

34.11 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 Using Trustees " in the *OES 2 SP1: File Systems Management Guide*.

For information about access control issues for NSS on Linux, see Section 5.5, "Access Control for NSS on Linux," on page 71.

34.12 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 *OES 2 SP1: File Systems Management Guide*.

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*.

34.13 Security Best Practices for zAPIs

zAPIs for NSS on Linux create the character special device `/dev/zapi`. Because zAPIs run at a level beneath where auditing tools track access and use, you should consider restricting access on the `/dev/zapi` directory to the `root` user and setting the device's POSIX permissions to mode=0400.

If you are using AppArmor, add the following line to the AppArmor profile of any program that uses zAPIs for NSS on Linux:

```
/dev/zapi r,
```

You should grant `root` access only to members of the administrative group called `wheel`. The `root` user is a member of the `wheel` group by default. Users in the `wheel` group can access the device by using the `su` or `sudo` commands to obtain `root` privileges for any necessary tasks.

To add a user to the `wheel` group:

1. Log in as the `root` user.
2. In a terminal console, enter

   ```
   usermod -G wheel username
   ```

   Replace *username* with the username of the user being added to the `wheel` group.

Regardless of the POSIX access rights set for the device, the Novell trustee model is enforced for the trustees and trustee access rights you define on `/dev/zapi` for individual users.

The key is specific to a user rather than a user-process pair. Therefore, two processes running as the same user can use the same key without requiring the second process to actually open the file. This behavior is the same as for zAPIs running for NSS on NetWare.

34.14 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.

34.15 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.

34.16 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 combination 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.

NSS Commands

A

This section describes commonly used Novell® Storage Services™ command line options for Novell Open Enterprise Server 2 servers and NetWare® 6.5 SP7 servers. The commands are grouped by management tasks.

- Section A.1, “Using NSS Commands,” on page 452
- Section A.2, “Help and Find Commands,” on page 454
- Section A.3, “Access Time Command,” on page 454
- Section A.4, “Background File System Checker Commands,” on page 455
- Section A.5, “Cache Management Commands (Linux),” on page 455
- Section A.6, “Cache Management Commands (NetWare),” on page 457
- Section A.7, “Compression Commands,” on page 460
- Section A.8, “Data Shredding Commands,” on page 463
- Section A.9, “Daylight Savings Time Commands,” on page 463
- Section A.10, “eDirectory Storage Object ID Commands,” on page 464
- Section A.11, “Extended Attributes (XAttr) Commands (Linux),” on page 464
- Section A.12, “Event File List (EFL) Command,” on page 467
- Section A.13, “File-Level Snapshot Commands,” on page 467
- Section A.14, “Hard Links Commands,” on page 468
- Section A.15, “I/O Monitoring Commands,” on page 470
- Section A.16, “I/O Write Commands (Linux),” on page 470
- Section A.17, “I/O Write Commands (NetWare),” on page 471
- Section A.18, “LAF Audit Log Messages Commands (Linux),” on page 472
- Section A.19, “Load Commands for the nssstart.cfg File,” on page 473
- Section A.20, “Low Storage Alert Messages Commands,” on page 474
- Section A.21, “Migration Support Commands for Near-Line Storage,” on page 474
- Section A.22, “Modified File List (MFL) Commands (NetWare),” on page 474
- Section A.23, “Multipath I/O Failover Commands (NetWare),” on page 475
- Section A.24, “Multiple Server Activation Prevention (MSAP) Commands,” on page 477
- Section A.25, “noatime and atime Commands,” on page 478
- Section A.26, “noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux),” on page 479
- Section A.27, “Opportunistic Locking (Oplock) Commands,” on page 481
- Section A.28, “Pool Freeze and Thaw Commands,” on page 481
- Section A.29, “Pool Management Commands,” on page 482
- Section A.30, “Pool Snapshot Commands (NetWare),” on page 483
- Section A.31, “Pool Verify and Rebuild Commands (NetWare),” on page 484

- Section A.32, “POSIX Permission Mask Command (Linux),” on page 485
- Section A.33, “Quotas Commands,” on page 485
- Section A.34, “Read Ahead Blocks and Allocate Ahead Blocks Commands,” on page 486
- Section A.35, “Removable Media Commands (NetWare),” on page 487
- Section A.36, “Salvage and Purge Commands,” on page 489
- Section A.37, “Security Equivalence Vector Update Commands (Linux),” on page 490
- Section A.38, “Sendfile API Support Command,” on page 491
- Section A.39, “Status Commands,” on page 492
- Section A.40, “Transaction Tracking System (TTS) Commands (NetWare),” on page 492
- Section A.41, “Visibility Rebuild Command,” on page 493
- Section A.42, “Volume Management Commands,” on page 494
- Section A.43, “Upgrade Commands for NetWare 5.x or 6.x to NetWare 6.5,” on page 498
- Section A.44, “ZID Commands,” on page 498

A.1 Using NSS Commands

- Section A.1.1, “Issuing NSS Commands at Command Consoles,” on page 452
- Section A.1.2, “Issuing NSS Commands at Server Startup (NetWare),” on page 452
- Section A.1.3, “Making NSS Commands Persist Through a Reboot,” on page 453
- Section A.1.4, “Permissions,” on page 453
- Section A.1.5, “Descriptions,” on page 453

A.1.1 Issuing NSS Commands at Command Consoles

Enter NSS commands at the NetWare server console or the NSS Console (NSSCON, `nsscon(8)`) on an OES 2 Linux server. For information about NSSCON, see Section B.13, “NSSCON (Linux),” on page 519.

A.1.2 Issuing NSS Commands at Server Startup (NetWare)

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"
```

A.1.3 Making NSS Commands Persist Through a Reboot

NSS commands issued at the command line 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. On NetWare, the file is in the `c:nwserver` directory. On Linux, it is in the `/etc/opt/novell/nss` directory. Some commands cannot be used in the `nssstart.cfg` file. Refer to the individual commands for information.

1. In a text editor, create a file called `nssstart.cfg` in the `c:nwserver` directory on NetWare, or in the `/etc/opt/novell/nss/` directory on Linux
2. Enter any NSS 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, on NetWare, the cddvd and dosfat options allow NetWare to automatically manage CDDVDs and DOS FAT partitions:

   ```
   /numworktodos=40 /cddvd /dosfat
   ```

   For example, on Linux, the `ListXattrNWmetadata` option enables the ability to return the `netware.metadata` extended attribute for a file or directory at `listxattr(2)` time. The `ZLSSUpgradeNewVolumeMediaFormat` option creates all new NSS volumes with the upgraded media structure that supports hard links. This applies to volumes on both local and shared pools.

   ```
   /numworktodos=40 /ListXattrNWmetadata /ZLSSUpgradeNewVolumeMediaFormat
   ```
3. Save and close your `nssstart.cfg` file.

A.1.4 Permissions

On Linux, you must be logged in as the `root` user, or as a Linux user with equivalent privileges.

On NetWare, you must be logged in as the administrator user for the server, or a user with equivalent Novell eDirectory™ permissions.

A.1.5 Descriptions

The descriptions of commands provide information about the default values, range of valid values, and persistence of the command.

Default Value

The default value is the setting used for a given server configuration parameter. Initially, the value reported is the default setting for the parameter. If you modify the value, it reports the actual value.

IMPORTANT: Default values are the best choice for a majority of server configurations, but you can modify the settings to meet your needs.

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 value or policy you set remains 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. Some commands can be issued in the `nssstart.cfg` file in order to make the settings persist across reboots.

On Linux, the file is `/etc/opt/novell/nss/nssstart.cfg`. On NetWare, the file is `c:nwserver\nssstart.cfg`.

A.2 Help and Find Commands

The `/help` and `/find` options provide information about various NSS switches and a brief description of the parameter. In the NSS Console, you can also simply enter `help` at the `nsscon` prompt.

/find *search_criteria*

To find a particular NSS switch, use the `/find` switch. Replace *search_criteria* with the characters to use in the search. Wild cards searches where you replace some characters with an asterisk (*) are allowed.

For example, the following command finds all of the NSS command options that contain the word "compress":

```
nss /find *compress*
```

/help or /?

To access online Help for NSS commands, enter one of the following:

```
nss /help
```

```
nss /?
```

A.3 Access Time Command

nss /(No)UpdateAccessTimeForReaddir=*volumename*

Enable or disable the ability to update the access time when enumerating directories. Enabled is the default POSIX behavior.

Default: On

Examples

To enable access time to be updated when files in a directory are accessed for listing, enter

```
nss /UpdateAccessTimeForReaddir=volumename
```

To disable access time to be updated when files in a directory are accessed for listing, enter

```
nss /NoUpdateAccessTimeForReaddir=volumename
```

A.4 Background File System Checker Commands

The background file system checker checks the integrity of user ID metadata, directory quota metadata, and files. On Linux, the output goes to the `nsscon` prompt. On NetWare, the output goes to the logger screen. There is no built-in report to follow the process of a background check.

nss /(No)BackgroundChecking

Enables or disables the background file system checker.

Default: On (enabled)

Examples

To enable background checking, enter

```
nss /BackgroundChecking
```

To disable background checking, enter

```
nss /NoBackgroundChecking
```

ForceBackgroundCheck

Forces the background file system checker to start.

A.5 Cache Management Commands (Linux)

Use the commands in this section to manage the cache for NSS volumes on Linux.

- Section A.5.1, “Cache Command,” on page 455
- Section A.5.2, “High Memory Cache Commands,” on page 455
- Section A.5.3, “ID Cache Commands,” on page 456
- Section A.5.4, “Cache Monitoring Commands,” on page 457
- Section A.5.5, “UnplugAlways Command for the Read Queue,” on page 457

A.5.1 Cache Command

nss /MinBufferCacheSize=*value*

Sets the specified minimum number of NSS buffer cache entries, where *value* is the number of 4-KB buffers to assign for NSS.

Default: 30000 for NSS on Linux

Range: 10000 for NSS on Linux to the amount of memory in KB divided by 4 KB (the block size). For a 32-bit machine, the maximum setting is 250000 buffers.

A.5.2 High Memory Cache Commands

nss /HighMemoryCacheType=[private | linux | none]

Set the type of use of User memory in 32-bit OES Linux as `none`, `linux`, or `private`. On 64-bit OES Linux, NSS does not use User memory at all.

Cache Type Options

private

Set up a separate user-mode cache in memory outside of the Linux cache. This is the default for 32-bit machines. By default, NSS takes 20 percent of the total high memory. You can modify this setting with the `nss /PrivateCacheSizeInBlocks=`*`value`* command.

However, If the number of Private Cache blocks is less than twice the NSS Cache blocks, then the default high memory cache type is `linux`. If the high memory is less than twice the NSS Cache blocks, then the default high memory cache type is `none`.

Private Cache is best used for dedicated file-servers. The Private Cache increases the performance of NSS, but reduces the amount of memory available to other file systems and tasks.

linux

Integrate caching into the regular Linux caches. This is the default for 32-bit machines where the number of Private Cache blocks is less than twice the NSS Cache blocks.

This can be a problem on a dual purpose file-server and application-server system, because memory-hungry applications can cause the file-system cache to purge completely.

none

Use the same algorithm as the initial release of OES 1 Linux, which is to try and cache everything in Kernel-mode memory. This is the only option on 64-bit OES Linux. This is the default for 32-bit machines if the high memory is less than twice the NSS Cache blocks.

nss /PrivateCacheSizeInBlocks=*value*

Sets the number of metadata blocks to cache for a 32-bit OES Linux machine where you have set the HighMemoryCacheType=private.

By default, NSS takes 20% of all high memory to use for the Private Cache. We recommend against giving NSS 100 percent of the high memory. The actual percentage you should allot depends on your deployment scenario. For example:

- Which features of NSS are enabled--salvage, compression, and so on.
- Use of really long filenames
- Use of extended attributes

Default: 20 percent of all high memory blocks

Range: 0 blocks to the number of blocks available in high memory (total available high memory divided by 4KB)

A.5.3 ID Cache Commands

Use the following command at the `nsscon` prompt in order to synchronize the cache of eDirectory IDs that is maintained for controlling access to NSS volumes.

nss /IDCacheResetInterval=*value*

Set the number of seconds between invalidation of the ID cache.

Default: 90000

Range: 0 to 200000000

ResetIDCache

Reset the various eDirectory ID caches.

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.

A.5.4 Cache Monitoring Commands

CacheStats

Shows the caching statistics for buffers.

ResetStats

Resets caching and file statistics.

A.5.5 UnplugAlways Command for the Read Queue

nss /(no)UnplugAlways

When enabled, this option allows NSS to unplug the device queue after queuing each read. This improves performance significantly on certain workloads, such as Linux copy (`cp`) command.

In OES 2 SP1 Linux and later, the UnplugAlways default setting is on (enabled).

For OES 2 Linux, the UnplugAlways default setting is off (disabled). The OES 2 Linux version of this option is available as patch *Novell Storage Services (NSS) and Novell Cluster Services (NCS) 20080806* (oes2-novell-nss-5503) for 32-bit and 64-bit architectures. The patch is available on the Novell Downloads Web site (http://download.novell.com).

Examples

Enable UnplugAlways by entering the following at the `nsscon` prompt as the `root` user:

```
nss /UnplugAlways
```

Disable UnplugAlways by entering the following at the `nsscon` prompt as the `root` user:

```
nss /NoUnplugAlways
```

A.6 Cache Management Commands (NetWare)

Use the commands in this section to manage the cache for NSS volumes on NetWare.

- Section A.6.1, “Buffer Cache Commands,” on page 458
- Section A.6.2, “File Cache Commands,” on page 458
- Section A.6.3, “Extended System Memory (ESM) Commands,” on page 459
- Section A.6.4, “ID Cache Commands,” on page 459
- Section A.6.5, “Name Cache Command,” on page 460
- Section A.6.6, “Time Cache Command,” on page 460
- Section A.6.7, “Cache Monitoring Commands,” on page 460

WARNING: If you give NSS more memory than you allow for the NetWare 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 A.1.2, "Issuing NSS Commands at Server Startup (NetWare)," on page 452.

A.6.1 Buffer Cache Commands

nss /(No)CacheBalance=*value*

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

nss /MinBufferCacheSize=*value*

Sets the specified minimum number of NSS buffer cache entries.

Default: 512 for NSS on NetWare

Range: 384 for NSS on NetWare to the amount of memory in KB divided by 4 KB (the block size). For a 32-bit machine, the maximum setting is 250000 buffers.

nss /MinOSBufferCacheSize=*value*

Sets the specified minimum number of NetWare buffer cache entries.

Default: 256

Range: 256 to memory size in KB/4KB

nss /CacheUserMaxPercent=*value*

Sets the specified maximum percentage of buffer cache that can be used for user data.

Default: 80

Range: 10 to 100

A.6.2 File Cache Commands

nss /ClosedFileCacheSize=*value*

Sets the number of closed files that can have their storage objects cached in memory. On average, each Closed File Cache entry consumes 0.4 to 1KB of RAM.

Default: 50000

Range: 16 to 1000000

Recommended: 100000

nss /(No)ClientFileCaching

Enable or disable client-side caching of opened files.

Default: On

A.6.3 Extended System Memory (ESM) Commands

Extended system memory (ESM) is the memory above 4 GB. By default, NSS on NetWare servers uses 80% of the ESM to cache user data for files greater than 128 KB in size. Only files with an End-of-File (EOF) files size greater than or equal to 128KB can be stored by ESM. To use ESM, the LSS must call it. ESM is called only by ZLSS and the `cddvd` LSSs. For example, large DOSFAT files are not stored in ESM memory.

/Above4GigMemoryPercent=*value*

Specifies the percentage of extended system memory used. This value can be set only at initialization time by adding the `Above4GigMemoryPercent` option in the `nssstart.cfg` file of the NetWare server.

```
nss /Above4GigMemoryPercent=value
```

For example:

```
nss /Above4GigMemoryPercent=90
```

Additional NSS command line options for managing NSS usage of ESM are:

/TrackESMCachePerf

Enables monitoring of the NSS ESM.

```
nss /TrackESMCachePerf
```

/NoTrackESMCachePerf

Disables monitoring of the NSS ESM.

```
nss /NoTrackESMCachePerf
```

/ESMCachePerf

Prints the statistics for NSS ESM.

```
nss /ESMCachePerf
```

A.6.4 ID Cache Commands

Use the following command to synchronize the cache of eDirectory IDs that is maintained for controlling access to NSS volumes.

nss /IDCacheResetInterval=*value*

Set the number of seconds between invalidation of the ID cache.

Default: 90000

Range: 0 to 200000000

ResetIDCache

Reset the various eDirectory ID caches.

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.

A.6.5 Name Cache Command

nss /NameCacheSize=*value*

Sets the specified maximum number of Name cache entries.

Default: 100000

Range: 17 to 1000000

A.6.6 Time Cache Command

ResetCachedTimes

Reset the various system caching for time conversions.

A.6.7 Cache Monitoring Commands

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.

A.7 Compression Commands

Use the commands in this section to manage compression parameters for NSS volumes where the Compression attribute has been enabled. After compression is enabled for a volume, it cannot be disabled.

- Section A.7.1, “Server-Level Compression Parameters,” on page 460
- Section A.7.2, “Volume-Level Compression Parameters,” on page 462

A.7.1 Server-Level Compression Parameters

Server-level compression parameters apply to all NSS volumes on the server. For details about each parameter, see Section 24.2.1, “Understanding Server-Level Compression Parameters,” on page 331.

nss /CompressionDailyCheckStartingHour=*value*

Default: 0

Range: 0 to 23

Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.).

nss /CompressionDailyCheckStopHour=*value*

Default: 6

Range: 0 to 23

Hours are specified by a 24-hour clock: (0=midnight; 23=11 p.m.).

nss /DaysUntouchedBeforeCompression=*value*

Default: 14

Range: 0 to 100000 (in days)

nss /EnableFileCompression

Enable file compression on volumes where the Compression attribute is enabled. Immediate Compress requests are queued until compression is allowed.

Default: On

Range: On or Off

nss /MinimumCompressionPercentageGain=*value*

The minimum percentage a file must compress in order to remain compressed.

Default: 20

Range: 0 to 50

nss /MaximumConcurrentCompressions=*value*

The number of simultaneous compressions allowed by the system (simultaneous compressions can only occur if there are multiple volumes).

Default: 2

Range: 1 to 8

nss /ConvertCompressedToUncompressedOption=*value*

Specify what the file system does with an uncompressed version of a file after the server has decompressed it.

IMPORTANT: Before a compressed file can be opened, there must be sufficient space available on the volume for the uncompressed and compressed copies of the file to coexist while the file is open.

Default: 1

Range: 0, 1, or 2

0 = Always leave the file compressed.

While the file is open, both the uncompressed and compressed copies of the file coexist on the volume. If the file is closed without changes, the uncompressed copy of the file is discarded. If changes are saved, the compressed copy of the file is discarded. After the modified file is closed, it is queued for immediate compression. Sufficient space must be available for both the compressed and uncompressed copies of the file to temporarily coexist on the volume in order for the compression to occur. After successful compression, the uncompressed copy of the modified file is discarded.

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. This is the default behavior for compression.

While the file is open, both the uncompressed and compressed copies of the file coexist on the volume. The first time that the file is closed without changes in the specified period, the uncompressed copy of the file is discarded. The second time that the file is closed without changes in the specified period, the compressed copy of the file is discarded. If changes are saved, the compressed copy of the file is discarded. The uncompressed file remains uncompressed until it meets requirements for being compressed.

2 = Always leave the file uncompressed.

While the compressed file is open, both the uncompressed and compressed copies of the file coexist on the volume. When the file is closed or when changes are saved, the compressed copy of the file is discarded. The uncompressed file remains uncompressed until it meets requirements for being compressed.

nss /DecompressPercentDiskSpaceFreeToAllowCommit=*value*

The percentage of disk space on a volume that is required to be free in order for file decompression to permanently change the compressed file version to uncompressed, which prevents newly uncompressed fils from entirely filling up the volume. Compressed files that are written to are always left uncompressed.

Default: 10

Range: 0 to 75

nss /DecompressFreeSpaceWarningInterval=*value*

The time interval (in minutes) between displaying warning alerts when the file system is not permanently changing compressed files to uncompressed files due to insufficient free disk space.

Setting the interval to 0 turns off the alert.

Default: 31 minutes

Range: 0 to 720 (0 seconds to 29 days 15 hours 50 minutes 3.8 seconds)

Setting the interval to 0 turns off the alert.

nss /DeletedFilesCompressionOption=*value*

Specifies whether and when to compress deleted files. This command presumes that you have enabled the Salvage attribute for NSS volumes. If Salvage is disabled, deleted files are purged immediately, so there are no deleted files to compress.

Default: 1

Range: 0, 1, or 2

0 = Do not compress deleted files.

1 = Compress deleted files the next day.

2 = Compress deleted files immediately.

A.7.2 Volume-Level Compression Parameters

NSS offers volume-level commands in this section for configuring and monitoring compression on a specified NSS volume.

CompScreen

Displays the NSS volume compression statistics on the compression screen.

nss /Compression=<*volumename* | all>

Enables the Compression attribute for the specified volume or for all volumes on the server. 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 Chapter 24, “Managing Compression on NSS Volumes,” on page 325.

nss /StopNormalCompression, or StopNormalCompression

Stops all queued compression for files, based on the compression triggered by a file open or close.

nss /(No)BGCompression

Allows compression to occur in the background at any time, instead of only within specified hours.

nss /NoBGCompression

Stops background compression and clears any queued background compression requests. Allow compression to occur only within the specified hours.

A.8 Data Shredding Commands

nss /(No)DataShredding=*volumename*:*count*

Enables or disables 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, because data shredding consumes a great deal of disk I/O bandwidth.

Default: 1

Range: 1 to 7, where 0 indicates no shredding

Examples

To enable data shredding on a volume VOL1 where the purged files are overwritten 7 times, enter

```
nss /DataShredding=VOL1:7
```

To disable data shredding for a volume VOL1, enter

```
nss /NoDataShredding=VOL1
```

A.9 Daylight Savings Time Commands

Use the commands in this section to manage daylight savings time for NSS volumes.

nss /DaylightSavingsTimeOffset=*value*

Issuing this command causes UTC time to be recalculated from local time. Specify the offset applied in time calculations when daylight savings time is in effect.

Default: +1 (one hour)

Range: 0 to 23

nss /StartOfDaylightSavingsTimeOffset=*value*

Local date and time when the switch on to daylight savings time should occur. Formats include a simple date and time enclosed in quotes, or rule enclosed in quotes and parenthesis. For example:

```
"April 1 2008 2:0:0 am"

"(April Sunday > 1 2:0:0 am)"

"(April Sunday First 2:0:0 am)"
```

Only rules cause rescheduling for the next year. You must set both the start and end dates before either is scheduled.

```
[Value=none]
```

nss /EndOfDaylightSavingsTimeOffset=*value*

Local date and time when the switch off of daylight savings time should occur. Formats include a simple date and time enclosed in quotes, or rules enclosed in quotes and parenthesis. For example:

```
"October 31 2008 2:0:0 am"

"(October Sunday <= 31  2:0:0 am)"

"(October Sunday Last 2:0:0 am)"
```

Only rules cause rescheduling for the next year. You must set both the start and end dates before either is scheduled.

[Value=none]

A.10 eDirectory Storage Object ID Commands

Use the commands in this section to remove or update the Storage objects in Novell eDirectory.

nss /RemoveObjectIDStore

Remove the object store.

nss /UpdateObjectIDStore

Scan and add all volume objects to an existing object store.

A.11 Extended Attributes (XAttr) Commands (Linux)

The Extended Attributes (XAttr) extension for NSS provides accessibility into many extended attributes for NSS on Linux. It allows you to read, back up, and restore extended attributes of files on NSS. This section describes options to determine how extended attributes are handled for NSS on Linux.

- Section A.11.1, "CtimeIsMetadataModTime Option," on page 465
- Section A.11.2, "ListXattrNWmetadata Option," on page 465
- Section A.11.3, "Additional Information," on page 467

A.11.1 CtimeIsMetadataModTime Option

By default, the Linux `ctime` is mapped to NSS create time (CreateTime). We prefer that `ctime` be based on the NSS metadata modified time (MetadataModifiedTime) instead of the NSS create time, but modifying the Linux `ctime` function might cause unknown complications. Thus, NSS provides the `CtimeIsMetadataModTime` option to allow an administrator to select to map the metadata modified time as the Linux `ctime` value, rather than the NSS create time when the different time stamp matters for your deployment.

The `CtimeIsMetadataModTime` option can be set persistently in the `/etc/opt/novell/nss/nssstart.cfg` file, or it can be set from `nsscon` by a user with `root` access.

nss /CtimeIsMetadataModTime

Maps the NSS metadata modified time to Linux `ctime`. This is the default behavior in OES 2 Linux and later, and in OES 1 SP3 Linux (not yet released).

nss /noCtimeIsMetadataModTime

Maps the NSS create time to Linux `ctime`. This is the default behavior in OES 1 SP2 Linux and earlier.

A.11.2 ListXattrNWmetadata Option

- "ListXattrNWmetadata Option" on page 465
- "Security Issues for ListXattrNWmetadata" on page 466
- "Using the Linux cp Command to Copy Files with Extended Attributes" on page 466

ListXattrNWmetadata Option

In OES 1 SP2 and later, the NetWare metadata (`netware.metadata`) extended attribute was added for files and directories. The `ListXattrNWmetadata` option for NSS allows a user or application with `root` access to select whether the `netware.metadata` extended attribute is returned for a file or directory at `listxattr(2)` time. The `ListXattrNWmetadata` option is disabled (OFF) by default. This option is intended for use by indexing or backup programs.

For users or applications without `root` access (without the `CAP_SYS_ADMIN` capability), the `listxattr(2)` command never lists the `netware.metadata` extended attribute, regardless of the `ListXattrNWmetadata` setting.

The `ListXattrNWmetadata` option can be set persistently in the `/etc/opt/novell/nss/nssstart.cfg` file, or it can be set from `nsscon` by a user with `root` access as follows:

nss /(No)ListXattrNWmetadata

Enables or disables the ability to return the `netware.metadata` extended attribute for a file or directory at `listxattr(2)` time.

The option is disabled by default in all OES versions. Enable the option if there is a need to use the Linux `xattr` functions to access or change NetWare metadata fields by name.

- **Off:** `listxattr()` does not return "netware.metadata" as an extended attribute for NSS files and directories. It is still possible to get extended attributes (`getxattr()`) and set extended attributes (`setxattr()`) by using the specific `xattr` name:

"netware.metadata"

- **On:** `listxattr()` returns “netware.metadata” as an extended attribute for NSS files and directories.

Examples

To enable the return of netware.metadata information, enter the following in the NSS Console:

```
nss /ListXattrNWmetadata
```

To disable the return of netware.metadata information, enter the following in the NSS Console:

```
nss /NoListXattrNWmetadata
```

Security Issues for ListXattrNWmetadata

The `ListXattrNWmetadata` option is available only to the user or application with `root` access (the `CAP_SYS_ADMIN` capability). It is disabled (off) by default.

When this feature is enabled (on) (such as by the backup user or by third-party backup software), and if the user or application has `root` user access, the following occurs:

- When copying NSS files or directories with the Linux `cp` utility from NSS volumes to NSS volumes, the `cp` utility copies the trustees assigned to a file or directory to the destination file or directory. This means that the old trustees of the file or directory now have visibility into the destination directory. In addition, the old trustees inherit trustee rights from the destination directory for other files in that directory.

  NOTE: For users or applications without `root` access (without the `CAP_SYS_ADMIN` capability), the trustee information is not copied to the destination directory.

- When copying NSS files with the `cp` utility from NSS volumes to non-NSS volumes, the `cp` utility issues a warning message advising that it could not apply the `netware.metadata` extended attribute.

  NOTE: For users or applications without `root` access (without the `CAP_SYS_ADMIN` capability), the `cp` utility does not attempt to apply the `netware.metadata` extended attribute.

There is no work-around for these two copy-related issues for the user or application with `root` access. This is how the Linux `cp` utility works.

Using the Linux cp Command to Copy Files with Extended Attributes

The Linux `cp` command has changed for OES 2. In OES 1, when listxattr is enabled, the extended attributes are also copied when you use the `cp` command as the `root` user to copy files. However, in OES 2, in order to copy the extended attributes, you must use the `--preserve` option.

The man page of `cp` on OES2 provides the following description of the `--preserve` option:

--preserve[=ATTR_LIST]
Preserve the specified attributes (default: mode,ownership,timestamps), if possible additional attributes: links, xattrs, all.

For example, after you enable listxattr, you can copy a file and its netware.metadata by logging in as the `root` user, then entering the following at a terminal console prompt:

```
cp --preserve=all /path/file1 /newpath/file1
```

A.11.3 Additional Information

For information about how to use the XAttr Extension for NSS, see the *NDK: XAttr Extension for NSS* (http://developer.novell.com/documentation/xattr/attr_enu/data/bktitle.html).

For the latest patches for the `km_nss` module, visit the Novell Support *SUSE Patch Support Database* (by date listing) (http://support.novell.com/linux/psdb/bydate.html).

For information about how to use the Linux `listxattr(2)` command, see the man page (enter `man 2 listxattr` at a terminal console prompt).

Novell Cool Solutions has a listxattrs tool you can use to check if you get the extended attributes after enabling `/ListXattrNWmetadata`. The listxattrs tool can be downloaded from the *Cool Solutions > Cool Tools > List Extended Attributes with xattr APIs* (http://www.novell.com/coolsolutions/tools/18206.html).

A.12 Event File List (EFL) Command

For information about the Event File List (EFL) feature for developers, see "FileEvent.xml Definitions" (http://developer.novell.com/documentation/vfs/vfs__enu/data/ak7gh2x.html) in *NDK: Virtual File Services (http://developer.novell.com/documentation/vfs/vfs__enu/data/bktitle.html)*.

nss /ResetEFLTree=*volumename*

Reset the Event File List (EFL) tree on the given volume.

A.13 File-Level Snapshot Commands

The File-level Snapshots attribute is supported for NSS volumes on Linux and NetWare. For more information about how this works, see Section 23.3, "Using the File-Level Snapshot Attribute to Enable the Backup of Open Files," on page 318.

nss /(No)FileCopyOnWrite=<*volumename* | all>

Enables or disables the ability to create a copy of files in the volume when they are open for write. The volume must be active and mounted before you issue the command. The All option does not alter settings for unmounted volumes.

After you enable the File Snapshot attribute on a volume, deactivate the volume, then activate it again. This ensures that there are no open files without a snapshot.

Examples

To enable the File Snapshot attribute for volume VOL1, enter the following at the server console:

```
nss /FileCopyOnWrite=VOL1
```

To enable file-level snapshots for all NSS volumes on the server, enter the following at the server console:

```
nss /FileCopyOnWrite=all
```

To disable file-level snapshots for volume VOL1, enter the following at the server console:

```
nss /NoFileCopyOnWrite=VOL1
```

A.14 Hard Links Commands

Use the commands in this section for managing hard links capability for NSS volumes on NetWare 6.5 SP4, OES 1 SP1 NetWare, OES 2 Linux and NetWare, and later versions.

- Section A.14.1, “Hard Links Media Upgrade Commands for Existing Volumes,” on page 468
- Section A.14.2, “Hard Links Media Upgrade Commands for New Volumes,” on page 468
- Section A.14.3, “Hard Links Attribute Commands,” on page 469
- Section A.14.4, “Hard Links Management Commands,” on page 469

A.14.1 Hard Links Media Upgrade Commands for Existing Volumes

Use the commands in this section to perform the one-time upgrade for currently active volumes.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=*volumename*

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.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all

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.

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=shared

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

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=local

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

nss /ZLSSUpgradeCurrentVolumeMediaFormat=all /include=local,shared

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

A.14.2 Hard Links Media Upgrade Commands for New Volumes

Use the commands in this section to enable all future NSS volumes to automatically be created with the new media format. Make sure to use the commands on the servers where you plan to use hard links. After a volume is created, you must manually enable the Hard Links attribute in order to create hard links on it. For command information, see Section A.14.3, “Hard Links Attribute Commands,” on page 469.

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.

nss /ZLSSUpgradeNewVolumeMediaFormat

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

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.

nss /ZLSSUpgradeNewVolumeMediaFormat /include=local,shared

Creates all new NSS volumes with the upgraded structure whether the volume is local or shared. In a cluster, the clustering system vetoes the upgrade until all nodes in the cluster are running a supported operating system.

A.14.3 Hard Links Attribute Commands

Use the commands in this section to enable or disable the Hard Links attribute for an NSS volume. The Hard Links attribute cannot be set or viewed in NSSMU or in the Storage plug-in to iManager.

nss /HardLinks=*volumename*

Enables the Hard Links attribute for the specified volume. This enables hard links to be created on the volume.

nss /HardLinks=all

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.

nss /NoHardLinks=*volumename*

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.

nss /NoHardLinks=all

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.

A.14.4 Hard Links Management Commands

You can view a report of hard links for a file to identify its primary link and the hard link that becomes the primary link if the primary link is deleted. For information, see Section 27.7, “Viewing Hard Links for a File,” on page 383.

/nss /ListHardLinks=vol:*path\filename.ext*

View information about the primary link and hard links for a file on an NSS volume.

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.

Use the `/volumes` option to verify that the Hard Links attribute is enabled.

nss /volumes

View a list of NSS volumes on the server and information about them. In the Attributes column, the *HardLinks* attribute is listed if it is enabled for the volume.

A.15 I/O Monitoring Commands

nss /ZLSSIOStatus

Displays current NSS I/O status information.

nss /ZLSSPendingWriteIOs=*value*

Number of write inputs dropped to the storage subsystem.

Currently, there is no way to see the current setting for this parameter. Initially, you can assume that the value is the default setting. If you set ZLSSPendingWriteIOs to a valid value within the valid range, you can assume that NSS correctly applied the new value if no error is reported when you issue the command. You must keep track manually of whatever valid value you set.

Default: 1000

Range: 0 to 3000 (0=all available)

A.16 I/O Write Commands (Linux)

Use the commands in this section to control the write behavior of dirty blocks from the cache to the disk. These commands are available only on OES 2 Linux and later. For more information about using these commands, see Section 31.3, "Configuring or Tuning Group I/O," on page 415.

nss JournalGroupWriteTime=*seconds*

Specify the elapsed time to wait before group writes of journal blocks.

Replace *seconds* with the maximum number of seconds to elapse before forcing journal blocks to be written to the volume. The default value of *seconds* is 1.

Example

To group write journal blocks every 2 seconds, enter

```
nss Journal GroupWriteTime=2
```

nss MetadataGroupWriteTime=*seconds*

Specify the elapsed time to wait before group writes of metadata blocks. Decreasing the metadata group write timer can help reduce the mount time for the volume after a crash.

Replace *seconds* with the maximum number of seconds to elapse before forcing metadata blocks to be written to the volume. The default value of *seconds* is 40.

Example

To group write metadata blocks every 30 seconds, enter

```
nss MetadataGroupWriteTime=30
```

nss UserDataGroupWriteTime=*seconds*

Specify the elapsed time to wait before group writes of user data blocks. Decreasing the user data group write timer can help reduce the risk of data loss for a volume after a crash.

Replace *seconds* with the maximum number of seconds to elapse before forcing user data blocks to be written to the volume. The default value of *seconds* is 3.

Example

To group write user data blocks every 1 second, enter

```
nss UserDataGroupWriteTime=1
```

nss MetadataGroupWriteLimit=*blocks*

Specify the maximum number of metadata blocks that can be dirty before a group write is performed.

Replace *blocks* with the maximum number of metadata blocks that can be dirty before forcing them to be written to the volume. The default value of *blocks* is 20000.

Examples

To decrease the maximum number of dirty metadata blocks to 15,000 for the purpose of reducing the mount time, enter

```
nss MetadataGroupWriteLimit=15000
```

To increase the maximum number of dirty metadata blocks to 30,000 for the purpose of increasing the file system performance, enter

```
nss MetadataGroupWriteLimit=30000
```

A.17 I/O Write Commands (NetWare)

Use the commands in this section to manage the write behavior for moving dirty blocks (blocks where information has been modified) from the cache buffer to the disk, referred to as flushing the buffer. For more information about how this works, see Section 23.1, “Enabling Flush Files Immediately to Write Data to the Disk on Close,” on page 315.

/(No)FlushFilesImmediately=*volume*

Enables or disables the ability to synchronously flush data to the specified volume when its files are closed. When it is disabled, NSS waits for the next scheduled write of data to the volume.

Examples

```
nss /FlushFilesImmediately=VOL1
```

```
nss /NoFlushFilesImmediately=VOL1
```

/BufferFlushTimer=*value*

Set the maximum amount of time in seconds that modified NSS file system cache buffers are kept in memory before they are written to disk.

IMPORTANT: In OES 2 Linux, this timer is not used. It has been replaced by the group write timers for Journal, Metadata, and User Data. For information, see Section 31.3, “Configuring or Tuning Group I/O,” on page 415.

Example

```
nss /BufferFlushTimer=10
```

Default: 1

Range: 1 to 3600

/FileFlushTimer=*value*

Sets the maximum amount of time in seconds that modified blocks are kept in a cache buffer before flushing them 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 seconds

Range: 1 to 3600 seconds

A.18 LAF Audit Log Messages Commands (Linux)

Use the NSS audit log messages commands to enable or disable messages via Lightweight Auditing Framework (LAF) for NSS trustee changes for NSS volumes on OES 2 Linux. When it is enabled, NSS reports changes for the following subset of NSS events:

- Adding trustees (AddTrustee)
- Removing trustees (RemoveTrustee)
- Setting the inherited rights mask (SetInheritedRightsMask)

The messages are stored in the `/var/log/audit/audit.log` file. For information about the content and format of messages in the log, see Section 22.4.1, “Understanding NSS Audit Log Messages,” on page 307.

nss /(No)LAFAuditTrustee=*volumename*

Enable or disable the generation of audit messages via Lightweight Auditing Framework for NSS trustee changes for NSS volumes on OES 2 Linux.

After you enable the audit log messages, the setting persists until the server reboot. After a server reboot, the audit log is disabled again by default. To make the command persist across reboots, add it to the `/etc/opt/novell/nss/nssstart.cfg` file.

To have the setting persist across reboots, add it to the `/etc/opt/novell/nss/nssstart.cfg` file.

Default: Off (disabled)

Values: On or Off

Examples

To enable NSS audit messages, enter the following at the `nsscon` prompt:

```
nss /LAFAuditTrustee=volumename
```

To disable NSS audit messages, enter the following at the `nsscon` prompt:

`nss /NoLAFAuditTrustee=volumename`

A.19 Load Commands for the nssstart.cfg File

Use the commands in this section to load the services automatically by placing them in the `nssstart.cfg` file. On NetWare, the `nssstart.cfg` file is in the `c:nwserver` directory, and on Linux, it is in the `/etc/opt/novell/nss` directory.

nss /(No)SkipLoadModules

If enabled, skips the auto-loading of the NSS support modules.

Default: Off

CDDVD

If it is specified in the `nssstart.cfg` file, loads the `cddvd.nss` module that is essential for CD, DVD, and Macintosh support on NetWare servers.

For more information, see Section 21.1, “Managing CDs and DVDs as NSS Volumes (NetWare),” on page 291.

defaultNLMs

If it is specified in the `nssstart.cfg` file, loads only default NSS NLM™ programs and any additional LSS modules specified on the command line.

DOSFAT

If it is specified in the `nssstart.cfg` file, loads only those modules that are essential for accessing FAT 16 and FAT 12 drives on NetWare servers. NetWare does not support FAT 32 drives.

DOSFAT supports only hard drives. DOSFAT does not support USB drives.

/allLSS

If it is specified in the `nssstart.cfg` file, it loads all LSS modules.

/noLSS

If it is specified in the `nssstart.cfg` file, it does not load the LSS modules.

/NumAsyncIOs=*value*

Sets the number of asynchronous I/O entries to allocate.

Default: 2048

Range: 4 to 65536

/NumBonds=*value*

Sets the number of bond entries to allocate.

Default: 5000

Range: 512 to 2097152

/NumWorkToDos=*value*

If used in the `nssstart.cfg` file, sets the number of WorkToDo entries. Entries can execute concurrently.

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.

Default: 50

Range: 5 to 100

/zLSS

If it is specified in the `nssstart.cfg` file, loads only those modules that are essential for zLSS support.

A.20 Low Storage Alert Messages Commands

nss /(No)StorageAlertMessages

Enables or disables NSS to send Low Storage messages to all users.

Default: On

nss /StorageAlarmThreshold=*value*

Sets the threshold (in MB) for a Low Storage space warning.

Default: 10

Range: 0 to 1000000

nss /StorageResetThreshold=*value*

Resets the threshold (in MB) for a Low Storage space warning.

Default: 10

Range: 0 to 1000000

A.21 Migration Support Commands for Near-Line Storage

nss /(No)Migration=<*volumename* | all>

Enables or disables migration of files on the specified volumes to a third-party near-line storage system.

This option is used only for third-party vendor applications that provide near-line storage. It is not a migration tool for migrating data from NetWare to Linux, and it is not associated with Dynamic Storage Technology.

A.22 Modified File List (MFL) Commands (NetWare)

Use the commands in this section to manage the Modified File List (MFL) feature for NSS volumes. The 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.

nss /(No)MFL=*volumename*

Enables or disables the Modified File List attribute for the specified volume.

Examples

To enable the MFL attribute for a given volume, enter

```
nss /MFL=volumename
```

To disable the MFL attribute for a given volume, enter

```
nss /NoMFL=volumename
```

nss /MFLVerify=*volumename*

Compares the modified file list with the specified volume and reports any inconsistencies.

nss /FixMFL=*volumename*

Repairs the modified file list to be consistent with the file system.

nss /GetMFLStatus=*volumename*

Shows the modified file list status for the specified volume.

A.23 Multipath I/O Failover Commands (NetWare)

Use the commands in this section for managing multipath I/O failover for NSS devices on NetWare. 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.

IMPORTANT: The Media Manager multipath I/O management is not available for Linux. For devices on Linux, use the Linux multipath I/O management tools. For information, Chapter 15, “Managing Multipath I/O to Devices (Linux),” on page 195.

---

list failover devices [*device_id*]

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.

To list the paths associated with all the devices on the server, enter the command without specifying the *device_id*.

To list the paths associated with a given device, enter the command followed by the *device_id*.

Example

To see the paths associated with a device with device_id 0x12, enter the following at the server console:

```
list failover devices 0x12
```

The list provides output similar to the following results:

```
0x12 [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
```

In this example, two paths are assigned to device `0x12`. Number 11 is up, and it has been assigned the highest priority of 1. The status of “Up” means that this path is available for use. The status of “selected” means the path is the primary path, and that any disk requests that are

trying to reach the 0x12 device automatically go through that path. In this example, you could use the `mm set failover path` command to override the priorities and make the number 26 path to be the primary path used.

mm set failover priority *pathid=number* [/insert]

This command sets the priority level for the selected path. The *pathid* must be a valid path ID. The *number* 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.

If the `/insert` option is not used, the specified path replaces the current path that is set at that priority level.

If the `/insert` option is used, the command inserts the specified path at the specified priority instead of replacing the current path set for that priority. For example, if you have four existing paths that are assigned the priorities of 1, 2, 3, and 4, and you enter a new path followed by a number 2 and the `/insert` option, then the paths that were previously assigned to 2, 3, and 4 priority are bumped to priority 3, 4, and 5 in order to accommodate the new path that has been assigned priority 2.

mm set failover state *pathid*=up [/setpath]

This command sets the state of the specified path to Up. The *pathid* must be a valid path ID.

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, to become the primary path.

mm set failover state *pathid*=down [/setpath]

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. The *pathid* must be a valid path ID.

If the Down option is followed by the `/setpath` option, NSS automatically selects the highest priority path from among all paths that are currently in the Up state to become the primary path.

Whenever a path is bad, NSS automatically moves the path to a Down state.

mm set failover path *pathid*

This command activates the specified path as “selected”, which means that it serves as the primary path. The *pathid* 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.

mm restore failover path *device_id*

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

mm reset failover registry

This command resets the multipath settings in the operating system registry (such as entries in `HKEY_LOCAL_MACHINE\Software\Novell\MediaManager\Failover\`), and re-creates them based on the current set of failover devices.

A.24 Multiple Server Activation Prevention (MSAP) Commands

Use the commands in this section to prevent NSS pools from being concurrently active on multiple nodes. 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.

nss /MSAPServer

Enables MSAP for all the pools on the server. By default, MSAP is enabled for every pool on the server. We recommend that you never disable MSAP.

nss /NoMSAPServer

Disables MSAP for all the pools on the server. This command remains in effect only until the server is rebooted. We recommend that you never disable MSAP.

nss /MSAPRebuild=*poolname*

Rebuilds a corrupt MSAP block. If the MSAP block for a pool becomes corrupt, it prevents a pool from going into Maintenance state. Use this command to rebuild the MSAP block for a given pool. Before issuing the command to rebuild, you must deactivate the pool. Rebuilding an MSAP block does not give the rebuilder ownership of the pool.

nss /PoolMSAP=*poolname*

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.)

nss /NoPoolMSAP=*poolname*

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.)

nss /MSAPServerID=*id*

Sets the ID for MSAP to use to uniquely identify the server.

nss /MSAPIDDisplay

Displays current IDs used by MSAP.

nss /pools

Displays the message *Multi-Use Detect* for NSS pools that have MSAP enabled.

A.25 noatime and atime Commands

The `/noatime` and `/atime` commands for a volume allow the administrator to control whether access times are updated when files and directories are read. They are available for NSS on Linux and NetWare beginning with OES 2 SP1 and NetWare 6.5 SP8. The setting persists across reboots.

Using `/noatime` is useful for backup, news servers, and mail servers where the extra disk activity associated with updating the access time is not desired. Avoiding the writes associated with updating the access time can result in measurable performance gains.

A.25.1 Using noatime or atime at the Command Line

Issue the commands at the NetWare server console or in the NSS Console on Linux.

nss /atime=*volumename*

This option enables the updating of access time for both files and directories in a specified volume so that reading a file updates its access time. This is enabled by default.

nss /noatime=*volumename*

This option disables the updating of access time for both files and directories in a specified volume so that reading a file does not update its access time.

A.25.2 Using noatime in a Cluster Load Script

You can also use the `noatime` option when mounting an NSS volume in a Novell Cluster Services cluster load script by using the `/opt` option. The `atime` setting is enabled by default, so it is not necessary to specify it explicitly.

For NetWare clusters, modify the `mount` command for the volume:

```
mount /opt="noatime" volume_name VOLID=number
```

For Linux clusters, modify the `ncpcon mount` command for the volume:

```
exit_on_error ncpcon mount /opt="noatime" volumename=volume_id
```

When you create a new volume on a cluster-enabled pool, Cluster Services automatically assigns it a volume ID that is unique in the entire cluster and writes the value to the cluster resource load script for the pool. When you add the `/opt` option to volume's mount line in the load script, the volume ID is already part of the command.

For example, for a volume named VOL1 and a volume ID of 254, the commands would be:

```
mount /opt="noatime" VOL1 VOLID=254
```

```
exit_on_error ncpcon mount /opt="noatime" VOL1=254
```

A.25.3 Viewing the atime or noatime Setting

You can view the current setting for the `atime` attribute by using the `nss /volumes` command at the NSS Console on Linux or the server console on NetWare.

A.26 noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux)

NSS on OES 2 Linux supports the `O_NOATIME` option for the Linux `open(2)` command, and the `noatime` and `nodiratime` options for the `mount` and `nfsmount` command and the `/etc/fstab` file. These options have the same objective—that is, to prevent the access time from being updated unless the access involves a modification of a file's or directory's metadata or content.

- **noatime:** Disables the updating of access time for both files and directories so that reading a file does not update their access time (atime).
- **nodiratime:** Disables updating of access time when opening directories so that the access time is not modified when enumerating directories. This routine also checks that the object is a directory, which slows down the routine.

These options are useful for backup, news servers, and mail servers where the extra disk activity associated with updating the access time is not desired. Avoiding the writes associated with updating the access time can result in measurable performance gains.

IMPORTANT: Typically, you need to use only the `noatime` option so that atime is not updated for the accessed file and its directory when the file is accessed. To determine if the `noatime` and `nodiratime` options can help the performance of a particular application, refer to the documentation and best practices for that application.

For information about applying these options, see the following:

- Section A.26.1, "Linux open(2) Command," on page 479
- Section A.26.2, "Linux mount Command," on page 480
- Section A.26.3, "Linux nfsmount Command," on page 480
- Section A.26.4, "Linux /etc/fstab File," on page 480
- Section A.26.5, "Patches for km_nss," on page 481

A.26.1 Linux open(2) Command

By default, the `open` command updates the access time whenever a file is opened or a directory is accessed. The `O_NOATIME` option disables the updating of access time, so that reading a file does not update its access time. Using this option allows you to back up a volume without modifying the access times (`ATIME`) of its files.

The man page for the `open(2)` command defines the `O_NOATIME` option as follows:

O_NOATIME

(Since Linux 2.6.8) Do not update the file last access time when the file is read. This flag is intended for use by indexing or backup programs, where its use can significantly reduce the amount of disk activity.

For information about how to use the `O_NOATIME` option, see the man page for the Linux `open(2)` command by entering `man 2 open` at a terminal console prompt.

A.26.2 Linux mount Command

The `noatime` and `nodiratime` options for the `mount` command are available for all Linux file systems, including NSS.

IMPORTANT: Typically, you need to use only the `noatime` option so that atime is not updated for the accessed file and its directory when the file is accessed. To determine if the `noatime` and `nodiratime` options can help the performance of a particular application, refer to the documentation and best practices for that application.

To enable `noatime` or `nodiratime` options when mounting an NSS volume:

1. Open a terminal console on the server, then log in as the `root` user.
2. At the terminal console prompt, enter (all on the same line, of course):

   ```
   mount -t nssvol VOL vol_mountpoint -o name=volname,noatime
   ```

   Replace *vol_mountpoint* with the mount point for the volume, such as `/media/nss/NSSV1`. Replace *volname* with the name of the volume, such as `NSSV1`.

   For example, enter:

   ```
   mount -t nssvol VOL /media/nss/NSSV1 -o name=NSSV1,noatime
   ```

   This command mounts the volume NSSV1 with the `noatime` option so that the file and directory access times are not updated when a file is accessed but not modified.

A.26.3 Linux nfsmount Command

1. In YaST, click *Network Services* > *NFS Server* > *NFS Mount*.
2. Specify the volume, then enter

   ```
   rw,no_root_squash,sync,fsid=value,noatime
   ```

   For information about the other `nfsmount` options used here, see Section 19.17, “Exporting and Importing NSS Volumes for NFS Access (Linux),” on page 273.
3. Click *OK*.

A.26.4 Linux /etc/fstab File

The `noatime` and `nodiratime` options are available as mount options in the `/etc/fstab` file for all Linux file systems, including NSS. To enable the `noatime` and `nodiratime` options as default mounting options for a volume so they are in effect at boot time, modify the entry for the NSS volume in the `/etc/fstab` file.

1. Open the `/etc/fstab` file in a text editor.
2. Modify the entry for the NSS volume by adding `noatime` or `nodiratime` as options.

   For example, type

   ```
   volname vol_mountpoint nssvol noauto,rw,name=volname,noatime 0 0
   ```

   Replace *volname* with the name of the volume, such as `NSSV1`. Replace *vol_mountpoint* with the mount location of the NSS volume, such as `/media/nss/NSSV1`.

For example (all on the same line, of course):

```
NSSV1 /media/nss/NSSV1 nssvol noauto,rw,name=NSSV1,noatime 0 0
```

3. Save the file.
4. Reboot the server to apply the changes.

A.26.5 Patches for km_nss

NSS support for access time options is delivered in the `km_nss` RPM in the Linux kernel. For the latest patches for the `km_nss` module in SUSE® Linux Enterprise Server 10, visit the Novell Support *SUSE Patch Support Database* (by date listing) (http://support.novell.com/linux/psdb/bydate.html).

A.27 Opportunistic Locking (Oplock) Commands

Use the commands in this section to enable or disable opportunistic locking (oplocks) for NSS volumes.

OpLocks are a type of client caching that manages concurrent read access to files. There is no real lock on the file. NSS implements two types of client caching: Level 1 and Level 2. For Level 1, only one client has the file open. When a second client opens the file, the oplock is broken, or if the client does not respond to the oplock break request, the oplock is broken for it. The client can then request a Level 2 oplock that can be shared by multiple clients as long as they send only read access requests to the file. When any of the clients send a write operation, the Level 2 oplock is automatically broken.

nss /(No)Level2OpLocksEnabled=*volumename*

Allow or disallow level 2 (shared) OpLocks.

Default: On

Examples

```
nss /Level2OpLocksEnabled=volumename
```

```
nss /NoLevel2OpLocksEnabled=volumename
```

A.28 Pool Freeze and Thaw Commands

Use the commands in this section to temporarily freeze (quiesce) and thaw activity on a pool.

nss /PoolFreeze=*poolname*

Temporarily quiesces activity on the specified pool.

nss /PoolThaw=*poolname*

Resume activity on the specified pool that has previously been frozen.

A.29 Pool Management Commands

Enter the following commands in a server console to perform maintenance for NSS pools and volumes on Linux and NetWare.

A.29.1 Pool Status

nss /LVScan=*poolname*

Scans for logical volumes within the specified active pool.

nss /pools, or pools

Shows all of the currently available NSS pools.

nss /ZLSSPoolScan

Scans for and loads all ZLSS pools.

nss /ZLSSPoolIOErrors

Shows the last 100 pool I/O errors reported from the block layer that have occurred since the server's last reboot. The block layer does not have many error codes (unlike the NetWare Media Manager), so the usefulness of the error may be limited. NSS stores some block layer errors in the `/var/log/messages` file. Look for `BIO_` in the log file to see if any pool I/O error messages are present..

A.29.2 PoolAuto Commands for Load Time

Use the pool commands in this section at load time. Any encrypted volumes on the pool cannot be automatically activated. Encrypted volumes remain deactive until you activate the volume manually and provide the password on the first activation after a reboot.

Place the PoolAuto commands in the `nssstart.cfg` file to make them persist across reboots. On NetWare, the `nssstart.cfg` file is in the `c:nwserver` directory, and on Linux, it is in the `/etc/opt/novell/nss` directory.

PoolAutoDisplay=<*poolname* | all>

Displays the pool's current load-time policies.

nss /PoolAutoActivate=<*poolname* | all>

Activates specified pools at load time.

nss /PoolAutoDeactivate=<*poolname* | all>

Leaves specified pools in the deactivated state at load time.

nss /PoolAutoMaintenance=<*poolname* | all>

Places specified pools in maintenance mode at pool load time.

nss /PoolAutoVerify=<*poolname* | all>

Verify the specified pool's physical integrity at startup time.

The following modifier commands can be used with the /PoolAuto commands to automatically handle pools of status type shared, corrupt, and questions.

nss <*/PoolAuto_command*> [/IncludeType=*type* | OverrideType=*type*]

Use the IncludeType or OverrideType modifier commands to include or override pools of the specified status type for a given /PoolAuto... command. Possible status types are SHARED, CORRUPT, and QUESTIONS.

A.30 Pool Snapshot Commands (NetWare)

The command line instructions in this section can be used to manage NSS pool snapshots on your NetWare server. We recommend that you use the Storage plug-in to Novell iManager 2.7 to manage pool snapshots.

IMPORTANT: NSS pool snapshots are supported on OES 2 Linux and later; however, there are no command line instructions available on Linux. You must use NSSMU to manage NSS pool snapshots on Linux. For general information about pool snapshots, see Chapter 18, "Managing NSS Pool Snapshots," on page 229. For an NSSMU for Linux quick reference on snapshots, see "Snapshots" in Table 9-12 on page 113.

Definitions of the variables used for commands in this section are as follows:

- **Datapool** The unique name of the stored-on pool for the pool snapshot; must be 2 to 16 characters.
- **Snapname:** The unique name of the pool snapshot, which is used as its pool name when the pool snapshot is activated; must be 2 to 16 characters. By default, the name is the pool name with an extension of `_Sn` where *n* is a sequential number (1 to 500) that is generated by NSS to indicate how many snapshots have been taken for a given pool. Numbers are reused as snapshots are deleted.
- **Snappool:** The unique name of the original data pool that you are snapping; must be 2 to 16 characters.

Issue the following Media Manager commands at the NetWare server console prompt:

mm snap create *snappool datapool snapname*

Creates a pool snapshot.

mm snap activate *snapname*

Activates a pool snapshot.

mm snap deactivate *snapname*

Deactivates a pool snapshot.

mm snap delete *snapname*

Deletes a pool snapshot.

mm snap rename *snapname newsnapname*

Renames the pool as specified.

mm snap list

Lists the snapshots on the server.

mm snap shred *snapname*=on

Turns on shredding for the internal volume of the pool where you are storing the specified snapshot.

If this option is enabled:

- You can have only one snapshot at a time of its original pool.
- After you delete the snapshot, allow enough time for the shredding process before you create a new one.

mm snap shred *snapname*=off

Turns off shredding for the internal volume of the pool where you are storing the specified snapshot.

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.

A.31 Pool Verify and Rebuild Commands (NetWare)

Use the commands in this section for verifying and rebuilding the metadata structure for NSS pools on NetWare. For NSS volumes on Linux, use the RAVSUI utility instead. For more information about rebuilding pools, see Chapter 17, "Verifying and Rebuilding NSS Pools and Volumes," on page 215.

PoolVerify

Allows you to select pools from a menu and verifies their physical integrity.

nss /PoolVerify=*poolname*

Verifies the physical integrity of the specified pool.

nss /PoolAutoVerify=*poolname*

Verifies the specified pool's physical integrity at startup time. Do not use for pools that contain any unencrypted volumes.

nss /PoolRebuild=*poolname* [/reZID]

Rebuilds the data structure of the specified pool. When the `/rezid` option is used with the pool rebuild command, the rebuild assigns new ZIDs to files on the pool's volumes while the pool is being rebuilt.

PoolRebuild [/reZID]

Allows you to select the pools to rebuild from a menu, then rebuild them. When the `/rezid` option is used with the pool rebuild command, the rebuild assigns new ZIDs to files on the pool's volumes while the pool is being rebuilt.

/(No)ReZID

Use this option only with the pool rebuild options (the `ravsui` utility on Linux or `nss /poolrebuild` on NetWare) to enable or disable the rebuild to reZID the volume(s) in the pool that is being rebuilt.

Default: Off

Range: On or Off

Examples

For guidelines and instructions for how to use the ReZID option with pool rebuild commands, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.

A.32 POSIX Permission Mask Command (Linux)

nss /PosixPermissionMask=*mask*

Specify the octal mask to control which bits in the POSIX permissions (`drwxrwxrwx`) are allowed to be set. The octal digits correspond to directory, user, group, and other fields. By default, NSS sets the POSIX permissions to 0777.

IMPORTANT: NSS uses the Novell trustee model to authenticate and give access to users, not the Linux ACLs and POSIX.

The command applies to all NSS volumes on the Linux server. In a cluster environment, make sure that the setting is the same on all nodes. This command should normally be added in the `/etc/opt/novell/nss/nssstart.cfg` file so that it persists across reboots.

Example

For example, SSH requires that the permissions in the Other field be disabled. If you use NSS volumes for home directories and you want users to have SSH access to them, you must modify the POSIX permissions to 0770. The following command in the `/etc/opt/novell/nss/nssstart.cfg` file turns off all of the bits corresponding to the Other field:

```
/PosixPermissionMask=0770
```

The setting applies to all NSS volumes on the server. You must also Linux-enable users and enable SSH with Linux User Management

If the volume is shared in a cluster, make sure to add the command to the `nssstart.cfg` file on all nodes and to Linux-enable SSH on all nodes.

A.33 Quotas Commands

Use the commands in this section to manage quotas (space restrictions) for NSS pools and volumes.

- Section A.33.1, “Sys: Volume Quota Command,” on page 485
- Section A.33.2, “Directory Quotas Commands,” on page 485
- Section A.33.3, “User Quotas Commands,” on page 486

A.33.1 Sys: Volume Quota Command

nss /ChangeSysQuota=*size*

Lets you change the quota (in MB) for the `sys:` volume. Setting this value to zero sets the Quota to none and allows the `sys:` volume to grow to the size of the `sys` pool.

A.33.2 Directory Quotas Commands

nss /(No)DirectoryQuotas=*volumename*

Enables or disables the Directory Quotas attribute on the volume.

nss /FixDirectoryQuotas=*volumename*

Recomputes used space for directory quotas for the specified volume.

This process can take time to recalculate, depending on the number of directory quotas set on the volume. There are no interim messages to report status.

nss /(No)ClassicDirectoryQuotas=*volumename*

For NSS on NetWare, enables or disables the emulation of directory quotas on NetWare Traditional volumes.

Default: Off (disabled)

A.33.3 User Quotas Commands

nss /(No)UserSpaceRestrictions=*volumename*

Enables or disables the User Quotas attribute on the volume. After you enable quotas, use go to the *Storage > User Quotas* page in iManager to set quotas for users.

A.34 Read Ahead Blocks and Allocate Ahead Blocks Commands

NSS offers the Read Ahead Blocks command for tuning read performance, and the Allocate Ahead Blocks command for tuning write performance. You can enter the commands at the server console on NetWare, or in the NSS Console (`nsscon`) on Linux.

- Section A.34.1, “Read Ahead Blocks,” on page 486
- Section A.34.2, “Allocate Ahead Blocks,” on page 487

A.34.1 Read Ahead Blocks

The Read Ahead Blocks parameter can be set differently on each NSS volume. It 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 parameter is enabled by default and set at 2 blocks on NetWare and 16 blocks on Linux. To modify the value, you must set it from the command line using the `ReadAheadBlks` switch. On Linux, you can also set the value from NSSMU in the volume properties.

The default block count 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.

nss /ReadAheadBlks=*volname:count*

Specify *VolName* as the name of the volume that you are setting the attribute for. Specify the *Count* to be the number of 4 KB blocks that you want to NSS to read ahead. 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 (NetWare); 16 (Linux)

Range: 0 to 1024 blocks, where 0=Off; Practical maximum: 128

IMPORTANT: The default ReadAheadBlks count for OES 1 Linux is 2 blocks. If you have read performance issues after upgrading from OES 1 Linux to OES 2 Linux, check the volume settings for ReadAheadBlks count. For information, see Section 3.7, "Upgrading from OES 1 Linux to OES 2 Linux," on page 53.

A.34.2 Allocate Ahead Blocks

Use the Allocate Ahead Blocks command for tuning write performance for all NSS volumes on the server.

nss /AllocAheadBlks=*value*

Sets the number of blocks to allocate ahead on writes.

Default: 15

Range: 0 to 63

A.35 Removable Media Commands (NetWare)

- Section A.35.1, "CDs and DVDs Commands (NetWare)," on page 487
- Section A.35.2, "CD and DVD Image Files Commands (NetWare)," on page 488
- Section A.35.3, "DOS FAT Volumes Commands (NetWare)," on page 489

A.35.1 CDs and DVDs Commands (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.

CDDVD

If CDDVD is specified in the `nssstart.cfg` file, loads only those modules essential for CD and DVD support.

nss /(No)LeaveCDinNDS

If the switches UpdateCDToNDS or LoginUpdateCDToNDS is enabled, enables this switch to retain the CD or DVD volume objects when the actual disks are removed.

Default: Off

nss /(No)LoadISOVolumeFirst

If this switch is enabled, it 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.

Default: Off

nss /(No)LockCDMediaIn

If this switch is enabled, it locks or unlocks the media in the device whenever CDs or DVDs are activated or deactivated as NSS volumes.

Default: Off

nss /(No)UpdateCDToNDS

If this switch is enabled, it updates their Novell eDirectory objects whenever CDs or DVDs are activated or deactivated as NSS volumes.

Default: Off

nss /(No)LoginUpdateCDtoNDS

If this switch is enabled, it updates their Novell eDirectory objects whenever CDs or DVDs are activated or deactivated as NSS volumes. You must login as admin.

Default: Off

nss /(No)MountISODiscSetInOneVolume

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`.

If this switch is disabled, `cddvd.nss` mounts each disk in the set as one individual NSS volume.

Default: Off

A.35.2 CD and DVD Image Files Commands (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

nss /MountImageVolume=*path_imagefilename*

Adds and mounts the specified CD or DVD image file as a read-only NSS volume.

Replace *`path_imagefilename`* with the path on the NSS volume where the image file resides, including its filename. For example:

```
nss /MountImageVolume=nssvol:\suse10_1.iso
```

nss /RemoveImageVolume=*path_imagefilename*

Dismounts and removes the specified CD or DVD image file as a read-only NSS volume.

Replace *path_imagefilename* with the path on the NSS volume where the image file resides, including its filename. For example:

```
nss /RemoveImageVolume= nssvol:\suse10_1.iso
```

A.35.3 DOS FAT Volumes Commands (NetWare)

NSS provides the following commands for managing DOS FAT partitions.

DOSFAT

If DOSFAT is specified in the `nssstart.cfg` file, loads only those modules essential for accessing FAT 16 and FAT 12 drives. OES NetWare does not support FAT 32 drives.

DOSFAT supports only hard drives. DOSFAT does not support USB drives.

nss /FATPartition=*partition_type_number*

Supports up to three additional partition types containing 16-bit FATs, such as /FATPartition=12,13.

nss /(No)FATLongNames

Enables long filenames on FAT volumes.

Default: On

A.36 Salvage and Purge Commands

Use the commands in this section for tuning the purging processes for NSS volumes. For more information about managing salvage for NSS volumes, see Chapter 26, “Salvaging and Purging Deleted Volumes, Directories, and Files,” on page 361.

nss /(No)ImmediatePurgeOfDeletedFiles

Enables or disables files to be purged immediately upon deletion.

NSS volumes on the server can optionally use the salvage feature of NSS to save deleted files until space is needed. When you enable the Immediate Purge of Deleted Files parameter, it affects all volumes on the server. If the Salvage attribute is enabled for a volume, enabling this flag overrides it so that deleted files and directories are purged immediately.

Default: Disabled

nss /(No)Salvage=<*volumename* | all>

Enables or disables the ability to salvage of deleted files on volumes. *Volume* 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.

nss /LogicalVolumePurgeDelay=*value*

The number of seconds before deleted logical volumes are purged. This allows time to reverse the deletion.

Default: 345600 (4 days)

nss /LogicalVolumePurgeDelayAfterLoad=*value*

The number of seconds after NSS loads before deleted logical volumes are purged. This allows time to pause autopurging.

Default: 7200

nss /LogicalVolumePurgeDelayAfterContinue=*value*

The number of seconds to delay purging a deleted logical volume after clicking Continue. After a volume starts to purge, it cannot be salvaged.

Default: 900

nss /PoolHighWaterMark=*poolname*:Percent

Purging begins in the salvage area when the pool's low watermark is reached, and continues until its high watermark is reached, or until there are no deleted files and volumes left to purge, whichever occurs first. Autopurging does not start again until free space again drops below the low watermark. Specify a given poolname or All to apply the setting to all pools.

The high and low watermarks must be at least 2% apart from each other.

The watermark is set to the default at boot unless you put the command in the startup file to set it to a different value.

Default: 20

Range: 2 to 100

nss /PoolLowWaterMark=*poolname*:Percent

Purging begins in the salvage area when the pool'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. Specify a given poolname or All to apply the setting to all pools.

The high and low watermarks must be at least 2% apart from each other.

The watermark is set to the default at boot unless you put the command in the startup file to set it to a different value.

Default: 10

Range: 0 to 98

nss /SalvageSys

Lets you restore the `sys:` volume if you have deleted it (if it has not yet been purged).

nss /LVDeleteStatusBasic, or LVDeleteStatusBasic

Displays information about deleted logical volumes.

nss /LVDeleteStatusSalvageable, or LVDeleteStatusSalvageable

Displays information about salvageable logical volumes.

A.37 Security Equivalence Vector Update Commands (Linux)

Use the Security Equivalence Vector (SEV) Update commands in the NSS Console utility (`nsscon`) on Linux 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. For more information about SEV, see Section 22.2, "Configuring the Security Equivalence Vector Update Frequency (Linux)," on page 303.

nss /(No)SecurityEquivalenceUpdating

Enables or disables SEV updates to occur in the background in addition to updates that occur when the system reboots. If it is disabled, SEV updates occur only at system reboots.

To make it persistent, include the command in the `/etc/opt/novell/nss/nssstart.cfg` file.

Default: On (enabled)

Examples

To enable background updating, enter

```
nss /SecurityEquivalenceUpdating
```

To disable background updating, enter

```
nss /NoSecurityEquivalenceUpdating
```

nss /UpdateSecurityEquivalenceInterval=*value*

Sets the SEV update interval to the specified value in seconds. At the end of the elapsed time, NSS requires updated SEVs from eDirectory.

To make it persistent, include the command in the `/etc/opt/novell/nss/nssstart.cfg` file.

Default: 7200 (2 hours)

Range: 300 (5 minutes) to 7776000 (90 days).

nss /ForceSecurityEquivalenceUpdate

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.

This command is invalid if used in the `/etc/opt/novell/nss/nssstart.cfg` file.

A unique abbreviation such as

```
nss /ForceS
```

also works.

ForceSecurityEquivalenceUpdate

Forces the user security equivalence background updating to start immediately. 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.

A.38 Sendfile API Support Command

nss /(No)SendfileSupport

Enable or disable support for the sendfile() API.

Default: On

A.39 Status Commands

Enter the following commands in a server console to show the status of various NSS parameters.

CompScreen

Displays the NSS volume compression statistics on the compression screen.

nss /pools, Pools

Lists all of the NSS pools that are currently available on the server.

PoolsAutoDisplay

Displays load-time policies for pools on the server.

nss /status, Status

Lists the current NSS status information.

nss /volumes, Volumes

Lists all of the NSS volumes that are currently mounted and active, including the `_admin` volume.

VolumesAutoDisplay

Displays load-time policies for volumes on the server.

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 /ErrorCode=*code*

Translates and describes the specified error code.

Modules

Lists the providers, loadable storage subsystems, and semantic agents.

Version

Displays the version information for NSS.

A.40 Transaction Tracking System (TTS) Commands (NetWare)

For information about using the Transaction Tracking System™ (TTS™), see Section 23.2, “Using the Transaction Tracking System for Application-Based Transaction Rollback (NetWare),” on page 317.

nss /(No)Transaction=<*volumename* | all>

Enables (or disables) transactional writes in files on the volumes. Use for the Transaction Tracking System only for NSS volumes on NetWare.

Examples

To enable the TTS attribute for volume VOL1, enter the following at the server console:

```
nss /Transaction=VOL1
```

To enable TTS for all NSS volumes on the server, enter the following at the server console:

```
nss /Transaction=all
```

To disable TTS for volume VOL1, enter the following at the server console:

```
nss /NoTransaction=VOL1
```

A.41 Visibility Rebuild Command

You might need to rebuild the visibility list in order to address problems where the computed value does not equal the stored value for the visibility lists. For example, if the `dmptrust volumename` command reports unknown trustees or many mismatched user GUIDs, you might need to run a visibility rebuild for the NSS volume.

You can use `checkvis.nlm` to check visibility for issues.

WARNING: The visibility rebuild process can be destructive. Users who could see directories before might not be able to afterwards. Run this operation only as required. Do not use visibility list rebuilds as a regular maintenance tool.

- Section A.41.1, “Description,” on page 493
- Section A.41.2, “Syntax,” on page 494
- Section A.41.3, “Additional Information,” on page 494

A.41.1 Description

When you perform a visibility rebuild, the system first deletes all current entries in the visibility list. Then, in a second pass, the system attempts to rebuild the visibility list for assigned trustees. The numbers you see at the conclusion of the rebuild give the following information:

Parameter	Description
Objects examined	Indicates the total number of file system objects looked at (a total of files and directories, plus special file system beasts).
Objects cleaned	Indicates the number of directories where visibility information was removed during the first pass.
Overflow objects removed	Indicates the number of visibility overflow objects removed during the first pass.
Trustees re-added	Indicates the number of trustees found on the volume and re-added to visibility lists. Although this number is non-zero, it doesn't mean problems were fixed or resolved. Instead, the number indicates the number of trustees found and included in the visibility list (whether they were there before the rebuild started or not).

After you run a visibility rebuild, make sure you run the visibility check again. If there are still errors, you must examine your trustees for problem before running a visibility rebuild again.

A.41.2 Syntax

nss /VisibilityRebuild=*volumename*

Rebuild the authorization visibility lists for an NSS volume.

A.41.3 Additional Information

For information about checking, repair, and troubleshooting the visibility list, see *NetWare 6 Trustee Rights: How They Work and What to Do When All Goes Wrong* (http://support.novell.com/techcenter/articles/ana20030202.html) in *Novell AppNotes* (2003, February 1).

A.42 Volume Management Commands

- Section A.42.1, "System Volume Commands (NetWare)," on page 494
- Section A.42.2, "Volumes Command," on page 495
- Section A.42.3, "Volume Activity Commands," on page 496
- Section A.42.4, "Encrypted Volume Activity Commands," on page 496
- Section A.42.5, "VolumeAuto Commands for Load Time," on page 497

A.42.1 System Volume Commands (NetWare)

NSS provides the following special commands for managing the system volume on NetWare.

nss /RenameToSYS

Lets you rename the `sys:` volume back to `sys:` if you have previously renamed it. At the prompt, enter the name of the volume you want to rename to `sys:`.

nss /SalvageSys

Lets you restore the `sys:` volume if you have deleted it (only if it has not yet been purged).

nss /ExpandSYS

Lets you expand the pool that contains your `sys:` volume. When prompted, enter the partition ID of the free space you plan to use to expand the pool. You can obtain the partition ID of all free partitions by using the `/ListFreeSpace` command.

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 `/ExpandSYS` command.

nss /SYSHotFixSize=*size*

Use this switch in conjunction with the /ExpandSYS command to set the Hotfix size (in KB) of the new partition.

Default: 100 KB

Range: 0 to the size, where 0=No HotFix

A.42.2 Volumes Command

NSS provides the `volumes` command and `/volumes` option for viewing a list of the currently mounted volumes, their status, and the attributes for NSS volumes. To view which attributes are currently set for a volume, enter

```
nss /volumes
```

You can also use enter the `volumes` command to get the same output.

The volumes utility for NCP provides additional information about the mounted volumes on a Linux server, such as its Linux path. For information, see Section B.24, “VOLUMES (Linux, NCP Utility),” on page 545.

For example, the NSS `volumes` command outputs state and attributes information in a tabular format:

```
     Volume Name            State                Attributes
----------------------  -----------  ----------------------------------
_ADMIN                  ACTIVE       Hardlinks
DATA1                   ACTIVE       Salvage
                                     Compression
                                     Data Shredding
DATA2                   ACTIVE       Salvage
VOL1                    ACTIVE       Salvage
                                     User Space Restrictions
                                     Directory Quotas
VOL2                    ACTIVE        Salvage
```

The following volume attributes are displayed on both NetWare and Linux:

Attribute	For Information
Compression	Section A.7, “Compression Commands,” on page 460
Data shredding	Section A.8, “Data Shredding Commands,” on page 463
Directory quotas	Section A.33.2, “Directory Quotas Commands,” on page 485
Encryption	Section A.42.4, “Encrypted Volume Activity Commands,” on page 496
File-level snapshots	Section A.13, “File-Level Snapshot Commands,” on page 467
Hard links	Section A.14, “Hard Links Commands,” on page 468
Migration (for near-line storage support)	Section A.21, “Migration Support Commands for Near-Line Storage,” on page 474
Salvage	Section A.36, “Salvage and Purge Commands,” on page 489
User quotas	Section A.33.3, “User Quotas Commands,” on page 486

The following volume attributes are displayed only on NetWare:

Attribute	For Information
Flush files immediately	Section A.17, “I/O Write Commands (NetWare),” on page 471
Modified File List (MFL)	Section A.22, “Modified File List (MFL) Commands (NetWare),” on page 474
Transaction Tracking System (TTS)	Section A.40, “Transaction Tracking System (TTS) Commands (NetWare),” on page 492

A.42.3 Volume Activity Commands

nss /ExtendMac=*volumename*

Enable extended Macintosh name space on the specified NSS volume on NetWare.

nss /ForceActivate=*volumename*

Forces an NSS volume to become active. For encrypted NSS volumes, this command cannot force an activation unless the volume has been previously activated with a password on the first activation after a reboot.

nss /ForceDeactivate=*volumename*

Forces an NSS volume to the deactive state. Does not prompt for open files.

nss /VolumeActivate=*volumename*

Activates the specified NSS volume.

nss /VolumeDeactivate=*volumename*

Deactivates the specified NSS volume.

nss /VolumeMaintenance=*volumename*

Places a specified volume into maintenance mode. Volumes can be put in maintenance mode, but maintenance occurs only at the storage pool level.

mount *volume_name*

Mount the specified unencrypted NSS volume or an encrypted NSS volume that has been previously activated with its password. If it has not been previously activated, it returns an error message, requesting more information.

mount all

Mount all unencrypted NSS volumes and all encrypted NSS volumes that have been previously activated with their passwords. Encrypted NSS volumes that were not previously activated return error messages, requesting more information.

A.42.4 Encrypted Volume Activity Commands

Use the commands in this section to display volume status and to activate, mount, deactivate, or dismount encrypted NSS volumes.

You must enter a password the first time the volume is activated or mounted following a system reboot. Thereafter, other environmental security and authentication measures control access to user data. On NetWare, the encryption password can be entered when mounting from NSSMU or activating the volume from the command line.

IMPORTANT: Use NSSMU to mount encrypted volumes on Linux the first time after a server reboot. Thereafter, you can use the Linux `mount` command.

You cannot use wildcard commands, such as `nss /VolumeAutoActivate`, to activate encrypted NSS volumes.

You cannot use the wildcard option of All as the volume name for volumes where the password has not previously been provided. 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.

nss /activate=*volume_name*

On NetWare, activate the specified unencrypted or encrypted NSS volume. If you are prompted for it, enter the encryption password. The password is required only on the first activation following a system reboot.

This command cannot be used in a script to activate encrypted NSS volumes.

nss /activate=all

On NetWare, activate all unencrypted NSS volumes and all encrypted NSS volumes that have been previously activated with their passwords.

mount *volume_name*

Mount the specified unencrypted NSS volume or an encrypted NSS volume that has been previously activated with its password. If it has not been previously activated, it returns an error message, requesting more information.

mount all

Mount all unencrypted NSS volumes and all encrypted NSS volumes that have been previously activated with their passwords. Encrypted NSS volumes that were not previously activated return error messages, requesting more information.

nss /volumeactivate=*volume_name*

On NetWare, activate the specified unencrypted or encrypted NSS volume. If you are prompted for it, enter the encryption password. The password is required only on the first activation following a system reboot.

nss /volumes

View the status of an encrypted and unencrypted NSS volumes. The encrypted volume returns a status of Encrypted.

A.42.5 VolumeAuto Commands for Load Time

Use the following command to view the volume's current load-time policies.

VolumeAutoDisplay=*volumename*

Displays the volume's current load-time policies.

Use the following commands in the `nssstart.cfg` file to control which volumes are active at load time. You cannot use these commands for encrypted NSS volumes. You must use NSSMU to activate the volume on the first time after restart so you can provide the password.

nss /VolumeAutoActivate=*volumename*

Activates the specified volume at load time.

nss /VolumeAutoDeactivate=*volumename*

Deactivates the specified volume at load time.

A.43 Upgrade Commands for NetWare 5.*x* or 6.*x* to NetWare 6.5

Use the commands in this section to upgrade legacy NSS volumes from NetWare 5.*x* or 6.*x* to the format of NetWare 6.5 initial release.

nss /UpgradeObjectsOnVolume=*volumename*

Upgrades the Storage objects for NetWare 5.*x* or 6.*x* NSS volumes mounted on NetWare 6.5 or later to the media format for the NetWare 6.5 initial release.

nss /ZLSSVolumeUpgrade

Allows the NetWare 5.x ZLSS VolumeUpgrade to the NetWare 6.5 format to occur during the upgrade process.

A.44 ZID Commands

Use the commands in this section to manage the file numbering, or ZIDs, of files on an NSS volume.

nss /ZIDNameSpace=*namespace*

Specify the name space (DOS, Long, Macintosh, or UNIX) the command `/ZIDToFileName` should use.

nss /ZIDtoFilename=*ZIDnumber*

For a specified ZID, reports the file's full path and filename for a given volume and name space. Use only with `/ZIDVolumeName` and `/ZIDNameSpace` to provide context for the command.

nss /ZIDVolumeName=*volumename*

The volume name the command `/ZIDToFileName` should use.

/(No)ReZID

Use this option only with the pool rebuild options (the `ravsui` utility on Linux or `nss /poolrebuild` on NetWare) to enable or disable the rebuild to reZID the volume(s) in the pool that is being rebuilt.

Default: Off

Range: On or Off

Examples

For guidelines and instructions for how to use the ReZID option with pool rebuild commands, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.

NSS Utilities

B

This section details the syntax and options for the following Novell® Storage Services™ utilities for Novell Open Enterprise Server 2 Linux and NetWare and for NetWare® 6.5.

- Section B.1, “ABORT REMIRROR (NetWare),” on page 499
- Section B.2, “ATTRIB (Linux),” on page 500
- Section B.3, “CDDVD (NetWare),” on page 502
- Section B.4, “COMPFIX,” on page 502
- Section B.5, “DISMOUNT (NetWare),” on page 508
- Section B.6, “FILER (NetWare),” on page 508
- Section B.7, “FLAG (NetWare),” on page 509
- Section B.8, “METAMIG (Linux),” on page 514
- Section B.9, “MIRROR STATUS (Netware),” on page 516
- Section B.10, “MOUNT (NetWare),” on page 517
- Section B.11, “NCOPY (NetWare),” on page 518
- Section B.12, “NDIR,” on page 518
- Section B.13, “NSSCON (Linux),” on page 519
- Section B.14, “NSSMU,” on page 521
- Section B.15, “NSSUPDATE (Linux),” on page 521
- Section B.16, “RAVSUI (Linux),” on page 522
- Section B.17, “RAVVIEW (Linux),” on page 525
- Section B.18, “REBUILD (NetWare),” on page 528
- Section B.19, “REFRESHIDS (Linux),” on page 529
- Section B.20, “REMIRROR PARTITION (NetWare),” on page 529
- Section B.21, “RIGHTS,” on page 530
- Section B.22, “SET (NetWare),” on page 536
- Section B.23, “VERIFY (NetWare),” on page 544
- Section B.24, “VOLUMES (Linux, NCP Utility),” on page 545

B.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.

B.1.1 Syntax

```
ABORT REMIRROR logical_partition_number
```

B.1.2 Option

Option	Description
logical_partition_number	Specify the number of the logical partition you want to stop remirroring.

B.1.3 Example

To stop remirroring logical partition 0x2, enter

```
ABORT REMIRROR 2
```

B.2 ATTRIB (Linux)

Use the Attribute (ATTRIB) utility to set NSS file system directory and file attributes on OES Linux.

IMPORTANT: This utility works only with directories and files in the NSS file system on Linux.

B.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.

B.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.
`-d, --dos`	Use DOS compatible attributes (that is, ro=ro,di,ri)
`-v, --version`	Displays the program version information.
`-h, --help`	Displays the ATTRIB help screen.

B.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. This attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare.	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.	Yes	No
sd	Subdirectory (status display only) indicates that the entry is a directory, not a file.	No	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
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

B.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.

B.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 function does not apply to Linux servers, because Linux mounts CDs and DVDs as Linux POSIX volumes.

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.

B.3.1 Syntax

```
[LOAD] CDDVD
```

B.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

B.4 COMPFIX

- Section B.4.1, “Purpose,” on page 503

- Section B.4.2, “COMPFIX for Linux,” on page 503
- Section B.4.3, “COMPFIX for NetWare,” on page 506

B.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.

B.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` user's 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
VOLUMENAME	Specifies the volume name of the compressed volume you want to repair, such as `VOL1:`.
FILENAME	Specifies 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: *VOL1:/dirpath/filename.xxx*

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=`*`path`*	Specifies the location of the log file. The default location is at the root of the compressed volume you are analyzing or fixing. **Default:** `./compfix.log`

Volume-Level Options (VOPTION)

Only one volume-level option can be selected.

Option	Description
`-D, --delete-all`	Deletes all non-fixable compressed files on the specified volume.
`-F, --fix-all`	Fixesall repairable compressed files on the specified volume.
`-L, --list-all`	Lists all problematic compressed files on the specified volume.
`-S, --fix-stats`	Fixes volume compression-related statistics. For accurate results, make sure the volume is in maintenance mode before issuing this command.
`-C, --clear-all`	Clears 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`	Checks if the specified file is eligible for the next background compression process.
`-d, --delete`	Deletes the specified compressed file if it is non-fixable.
`-f, --fix`	Tries to fix the specified file's compression-related problem.

Option	Description
-l, --list	Lists the specified file's compression-related information.
-c, --clear	Clears 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:	Fixes 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	Lists 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	Fixes an individual compressed file, myfile.xxx. Log the results in the default location of VOL1:/compfix.log.
compfix -F VOL1: --logpath=LOGS:/compfix.log	Fixes 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	Checks 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:. Logs the results in the specified location of LOGS:/compfix.log.

B.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 command at the server console. `Compfix.nlm` runs on processor 0 only.

```
compfix -[attribute] [-noheader] path [log_volume]
```

Parameters

Parameter	Description
path	Specifies 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. For example: *VOL1:* *VOL1:\dirpath\filename.xxx*
log_volume	If the *log_volume* is specified, the COMPFIX log data is written to the *log_volume*`:\compfix.txt` file. For example, if `LOGS:` is the *log_volume* , then the file is written to `LOGS:\compfix.txt`. If the *log_volume* is not specified, the COMPFIX log data is written to the root of the compressed volume such as *volumename*`:\compfix.txt` file. For example, *VOL1*`:\compfix.txt`.

Option

Option	Description
`-noheader`	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. If `-noheader` is not specified, COMPFIX validates compression headers.

Attributes

Multiple attributes can optionally be used for a single command.

Attribute	Description
-fs	Fixes compression statistics. Make sure the volume is in maintenance mode before you begin in order for statistics to be properly computed.
-lc	Lists all corrupted compressed files on the volume.
-df	Diagnoses an individual compressed file.
-ff	Fixes an individual compressed file.
-fc	Fixes all fixable corrupted compressed files on the volume.
-bg	Diagnoses 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 Cannot Compress (Cc) 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
compfix -fs VOL1:	Fixes compression statistics for the specified volume, VOL1:. The volume should be in maintenance mode to get accurate results. Log the results in the default location of VOL1:\compfix.txt.
compfix -lc VOL1: LOGS:	Lists all corrupted compressed files on the specified volume, VOL1:. Log the results in the specified location of LOGS:\compfix.txt.
compfix -ff VOL1:\dir1\dir2\myfile.xxx	Fixes an individual compressed file, myfile.xxx. Log the results in the default location of VOL1:\compfix.txt.
compfix -fc VOL1: LOGS:	Fixes all fixable corrupted compressed files on the specified volume, VOL1:. Log the results in the specified location of LOGS:\compfix.txt.
compfix -bg VOL1:\dir1\dir2\myfile.xxx	Diagnoses the background compression effect for an individual file, myfile.xxx. Logs the results in the default location of VOL1:\compfix.txt.

Command	Description
`compfix -cc VOL1: LOGS:`	Clears the Cannot Compress (Cc) attribute for files in the specified volume, `VOL1:`. Logs the results in the specified location of `LOGS:\compfix.txt`.

B.5 DISMOUNT (NetWare)

Use at the server console to make a volume unavailable to users.

B.5.1 Syntax

`dismount` *`volumename`*

B.5.2 Option

Parameter	Description
`volumename`	Specifies the name of the volume you want to dismount.

B.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.

B.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 Novell 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, change directory (`cd`) to the mapped drive, then issue the `filer` command from the command line.

B.6.1 Options

The following figure shows the functions of FILER.

Figure B-1 *FILER Functions*

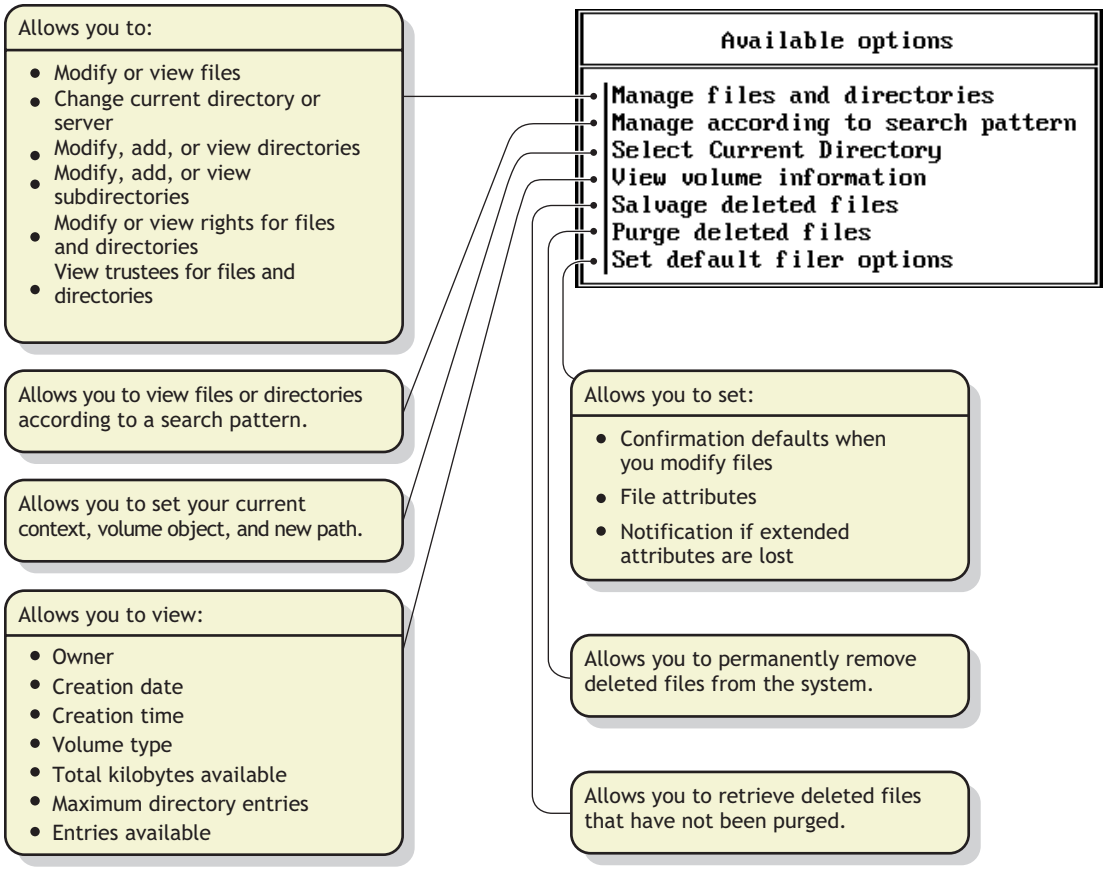

B.7 FLAG (NetWare)

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

B.7.1 Syntax

The `flag` command is a DOS command line 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]
```

B.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.
path	Specifies the path to the file or directory whose attributes or search modes you want to modify.
+, -	Sets 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 are 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.
attribute	Specifies one or more file or directory attributes.
	See Directory Attributes and File Attributes.
/*option*	Replaces *option* with any available option. See General Options.
/?	Views online help. All other parameters are ignored when /? is used.
/VER	Views the version number of the utility and the list of files it uses to execute. All other parameters are ignored when /VER is used.

General Options

Option	Description
/NAME=*name*	Changes the owner of a file or directory.
/GROUP=*name*	
/D	Displays details about a file or directory.
/DO	Displays or modifies only directories in the specified path.
/FO	Displays or modifies only files in the specified path.
/OWNER=*name*	Displays all files or directories owned by a user.
/M=*mode*	Modifies search modes of executable files. See Search Modes for Executable Files.

Option	Description
/S	Searches the subdirectory in the specified path and any subdirectories below that level.
/C	Scrolls continuously through output.

Search Modes for Executable Files

The syntax is

/M=*mode*

Replace *mode* with a mode number from the following table.

Mode	Description
0	Searches for instructions in the net.cfg file (the default mode).
1	Searches the path specified in the file. If there is no path, the file searches the default directory, and then all search drives.
2	Searches the path specified in the file. If there is no path, the file searches only the default directory.
3	Searchesthe 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 searches the path specified and then searches 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 searches 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.

B.7.3 Attributes

When you set or clear multiple attributes, include a space between attribute abbreviations.

Directory Attributes

Option	Description
ALL	Specifies 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)	Prevents 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 deletet 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)	Prevents the directory from being migrated to a secondary backup system (regardless of what the volume is set to).
H (Hidden)	Prevents the directory from being seen in the directory listing.
Ic (Immediate Compression)	Compresses 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)	Specifies no attributes.
P (Purge)	Purges the directory or files in a directory immediately when deleted.
Ri (Rename Inhibit)	Prevents the directory from being renamed.
Sy (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	Specifies 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)	Indicates that the file has been modified since the last backup.
Ci (Copy Inhibit)	Prevents files from being copied. This attribute overrides the trustee Read right and File Scan right. This attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare.
Dc (Don't Compress)	Prevents a file from being compressed.
Di (Delete Inhibit)	Prevents 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)	Prevents a file from being migrated to a secondary backup system (regardless of what the volume or directory is set to).
Ds (Don't Suballocate)	Prevents 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)	Prevents a filename from being displayed with the DOS DIR command. The file can't be copied or deleted.
Ic (Immediate Compression)	Compresses a file as soon as the OS can.
N (Normal)	Specifies the Rw attribute.
P (Purge)	Purges a file immediately if the file is deleted.
Ri (Rename Inhibit)	Prevents a file from being renamed.
Ro (Read Only)	Allows a file to only be read; it cannot be modified. Typically, it is used with Ri and Di. Ri and Di.
Rw (Read Write)	Allows a file to be read and written to.
Sh (Shareable)	Allows a file to be used by several users simultaneously. This attribute is usually used with Read Only.
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.
T (Transactional)	Protects 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 does 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 Section B.4, "COMPFIX," on page 502.
Co (Compressed)	The file has been compressed.
M (Migrated)	The file has been migrated.

B.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
```

B.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.

B.8.1 Syntax

METAMIG is located in the `/opt/novell/nss/sbin/metamig` directory.

```
metamig [OPTIONS]

metamig save volume [SOPTIONS]

metamig restore volume [ROPTIONS]
```

B.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`	Saves the indicated metadata to `stdout`.
`restore`	Restores 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.

B.8.3 Options

Several option types are available:

- “General Options” on page 515
- “SOptions (Save)” on page 515
- “ROptions (Restore)” on page 515
- “Mask” on page 515

General Options

The following options are general options available to actions related to save or restore.

Option	Description
`-v, --version`	Displays the program version information.
`-h, --help`	Displays the help screen.

SOptions (Save)

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 Mask.
`-n, --ncp`	Saves the trustee metadata from the NCP™ Trustee database instead of parsing the volume to generate the data. Use this option only if you are confident that the NCP Trustee database is available and current. For example, if the NCP Server has been turned off and its database is not yet resynchronized with the volume, do not use this command option until the database is again current.

ROptions (Restore)

The following options are options available to actions related to restore.

Option	Description
`-m, --meta=`*`mask`*	The types of metadata to be restored. If a value is not specified, it defaults to `all`. For information, see Mask.
`-f, --filter=`*`regexp`*	A filter that restores only files and directories that match the specified regular expression. Use a regular expression to specify one or more files and directories to be restored. For example, to set criteria to restore only files with names that start with the letter “a”, use this option:--filter=.*\\a.* Please refer to a programming textbook or search the Internet for information about how to construct regular expressions.
`-t, --tree=`*`treename`*	If this value is specified, this setting overrides the saved directory treename.
`-d, --details`	Displays 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 SOptions (Save) and ROptions (Restore).

Option	Description
t	Trustees
u	User quotas
d	Directory quotas
a	All metadata

B.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
```

B.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.

B.9.1 Syntax

```
MIRROR STATUS [logical_partition_number]
```

B.9.2 Option

Option	Description
(no option)	Displays a list of all logical disk partitions with their mirrored status.
logical_partition_number	Displays the partition's mirrored status and the devices that make up the members of the mirror group.

B.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.

B.10 MOUNT (NetWare)

On NetWare, use the `mount` command to make an NSS volume available to users. For information about using this command for encrypted volumes, see Section A.42.4, "Encrypted Volume Activity Commands," on page 496.

B.10.1 Syntax

`MOUNT` *`volume_name`* `|` `ALL`

Parameter	Description
`volume_name`	Specifies the name of the volume you want to mount.
`ALL`	Mounts all volumes not mounted.

B.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.

B.10.3 Examples

- To mount the `STATISTICS` volume, enter

  ```
  MOUNT STATISTICS
  ```

- To mount all unmounted volumes at once, enter

  ```
  MOUNT ALL
  ```

B.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 volume or directory you want to manage. Open a DOS Command Prompt window, go 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.

B.11.1 Syntax

Command	Description
`ncopy /R`	Retain compression on supported media.
`ncopy /R/`	Retain compression on unsupported media.

B.12 NDIR

The `ndir` command is a legacy MS-DOS utility that allows you to view file attributes and compression information for all files in a directory. NDIR requires the NetWare Core Protocol™ (NCP) protocol, so if you are using the command for NSS volumes on Linux, NCP Server must be configured and running.

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 volume or directory. Open a DOS Command Prompt window, go to the mapped drive, then issue the `ndir` command from the command line.

B.12.1 Syntax

At the MS-DOS command prompt on your Windows workstation, navigate to the mapped drive for the NSS volume, then enter

```
ndir path /option
```

The path is optional if you want the information for all directories in the mapped drive.

Help Options

Help Option	Description
`ndir /?`	Displays a list of help options.
`ndir /? for`	Displays help on display format.
`ndir /? sort`	Displays help on sorting features.

Help Option	Description
ndir /? res	Displays help on search filters (restrictions).
ndir /? at	Displays help on attribute filters.
ndir /? opt	Displays help on other options.
ndir /? syn	Displays help on syntax.
ndir /? all	Displays all help screens.

Compression and File Attributes Options

Command	Description
ndir path /R	Displays the trustee rights, inherited rights filter, and file attributes for all files in the specified volume or directory.
ndir path /COMP	Displays the compression information for all files in the specified volume or directory. It shows the amount of disk space saved by compressing data.
ndir path /COMP /VOL	Displays the compression information for the specified volume or directory. It also shows the disk space saved on the volume.

Example

```
C:\>ndir z: /comp /vol
Statistics for fixed volume <servername/vol>:
Space statistics are in KB (1024 bytes).
Total volume space:                          4,027,620  100.00%
Space used by 50,548 entries:                1,134,036   28.16%
Deleted space not yet purgeable:                     0    0.00%
                                           ------------ --------
Space remaining on volume:                   2,893,584   71.84%
Space available to <username>:               2,893,584   71.84%
Maximum directory entries:                 2,147,483,647
Available directory entries:               2,147,433,099  100.00%
Space used if files were not compressed:             0
Space used by compressed files:                      0
                                           ------------
Space saved by compressing files:                    0    0.00%
Uncompressed space used:                     2,069,352
Name spaces loaded: MAC, NFS, OS/2
```

B.13 NSSCON (Linux)

The NSS Console (NSSCON, `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.

Unlike NSS utilities, the NSS commands cannot be issued directly at the Linux terminal console. Therefore, you start the `nsscon` utility, then enter the usual NSS commands from the `nsscon` prompt. You can issue any NSS command that is valid for use on Linux at the nsscon command prompt. For information about NSS commands, see Appendix A, "NSS Commands," on page 451.

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 it 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.

- Section B.13.1, "Adding /opt/novell/nss/sbin to the PATH Environment Variable," on page 520
- Section B.13.2, "Starting nsscon," on page 520
- Section B.13.3, "Using nsscon in a Script," on page 520

B.13.1 Adding /opt/novell/nss/sbin to the PATH Environment Variable

You can 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.

B.13.2 Starting 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
     ```

B.13.3 Using nsscon in a Script

Only one instance of `nsscon` can be run at a time. If `nsscon` is already running when a script tries to run it, `nsscon` returns an error.

To work around this issue, you can send NSS commands directly to NSS via the `/dev/nsscmd` device.

For example, if you use a script to put pools into maintenance, use the following syntax:

```
echo "/PoolMaintenance=mypool" >/dev/nsscmd
```

This causes NSS to place MYPOOL into maintenance state even if `nsscon` is already running. Note that you do not echo an "`nss `" in front of commands to `/dev/nsscmd`.

If `nsscon` is running, then the output of your command is displayed by `nsscon` immediately. Otherwise, the output is held by NSS until `nsscon` is run. In OES 2 Linux, NSS holds 400K worth of output before starting to throw the oldest away. In OES 1 Linux, NSS holds only 32K of output before wrapping output.

B.14 NSSMU

The NSS Management Utility (NSSMU) is available for Linux and NetWare. For an NSSMU quick reference, see Section 9.2, "NSS Management Utility (NSSMU) Quick Reference," on page 112.

B.15 NSSUPDATE (Linux)

The NSS Update (`nssupdate`) utility is used to resize a pool to a larger size so that it can consume contiguous free space that follows the existing pool on a device.

You should not need to run this utility if you increase the size of a pool in the NSS-supported EVMS (Enterprise Volume Management System) environment. When working with NSSMU (NSS Management Utility) and the Storage plug-in for Novell iManager to manage NSS pools and volumes on EVMS-managed devices, NSS automatically makes any necessary changes to the pool's underlying structure.

You can use this tool to manually resize a pool if you are using the NSS file system on a device that is managed by a volume manager other than EVMS, such as LVM2 (Linux Volume Manager 2). After you have performed all steps to increase the size of the LVM2 device and partitions by using Linux tools, use the NSS Update utility to let NSS know to expand the pool size to the desired new size (specified in bytes).

B.15.1 Syntax

```
nssupdate -pool poolName -size sizeInBytes [-shared | -notshared]
```

Issue this command from a terminal console prompt as the `root` user.

B.15.2 Availability

Novell Open Enterprise Server 2 for Linux

B.15.3 Options

-pool *poolname*

Specifies the name of the pool you want to resize.

-size *sizeInBytes*

Specifies the new maximum size of the pool in bytes.

The value for sizeInBytes cannot exceed the actual size of the device. You can enter any value larger than the pool's current size and up to the size of the device. The space you designate must already be free for consumption by the pool. Make sure you understand exactly what space is free on the device so that you do not inadvertently overwrite any metadata stored at the end of the device. The pool begins at the same sector location as it currently does, and the extension is based on space that follows its current end location on the device.

[-shared | -notshared]

Sets the share state of the pool. The -shared option sets the share state to *Shareable in a Cluster* so that the pool can be shared in a cluster environment that is using Novell Cluster Services™ (NCS) for Linux. The -unshared option sets the share state to *Not Shareable*.

These options are intended to be used when you are using NSS pools with Novell Cluster Services clusters. NCS requires that the devices be managed by EVMS. You cluster-enable the shared device by placing the Cluster Segment Manager on the device before creating the file system on it. For the NSS file system, you create the clustered pool on the shared device and manage it with NSS management tools, such as NSSMU or the Storage plug-in to Novell iManager.

IMPORTANT: The shared state is not viable for NSS on non-EVMS-managed devices, because NCS currently requires the Cluster Segment Manager and EVMS.

B.15.4 Example

Open a terminal console, then log in as the `root` user to use this command.

`nssupdate -pool puserdata -size 2147483648`

Resizes the pool named `puserdata` to 2 GB (where 1 GB is 1024E3 bytes or 1,073,741,824 bytes).

B.16 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. You must log in as the `root` user to run this utility. You must place the pool in maintenance mode before starting the rebuild or verify process.

B.16.1 Syntax

```
/opt/novell/nss/sbin/ravsui

ravsui [OPTION]... [ROPTION]... rebuild poolname

ravsui [OPTION]... [VOPTION]... verify poolname
```

Replace *`poolname`* with the name of the pool you want to rebuild or verify, such as `POOL1`. Poolnames are case sensitive on Linux.

B.16.2 Arguments

The first mandatory argument specifies the action to be performed as rebuild or as verify.

- The rebuild action checks the integrity of the data in the pool and rebuilds the pool. A rebuild process can take up to several hours, depending on the size of the pool. The rebuild argument can be combined with one or more OPTION and ROPTION options.
- The verify action checks the integrity of the data in the pool; it does not perform the repair. A verify process can take several minutes, depending on the size of the pool. The verify argument can be combined with one or more OPTION and VOPTION options.

The second mandatory argument *poolname* is the target of the action. Replace *poolname* with the name of an NSS pool to be rebuilt or verified. Poolnames are case sensitive on Linux.

B.16.3 Options

This section describes the OPTION, ROPTION, and VOPTION options available for the RAVSUI (`ravsui`) utility. Mandatory arguments for long options are also mandatory for short options.

- "OPTION" on page 523
- "ROPTION" on page 523
- "VOPTION" on page 524

OPTION

General options can be used for both rebuild and verify actions.

Option	Description
`-a, --attach`	Attach to a rebuild or verify that it 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 Linux path where the log file is written that contains the results of the rebuild or verify action. **Default:** `/var/opt/novell/log/nss/rav/`
`-v, --version`	Displays version information and exits.

ROPTION

ROPTIONs can be used only for the rebuild action.

Option	Description
`-i, --iv-prune`	Prunes an internal volume.
`-l, --loss-file-limit=LIMIT`	Specifies the maximum number of files per volume to quietly prune. **Default:** 100
`-p, --purge-deleted-files`	Purges deleted files.
`-r, --rezid=ZID`	Specifies the threshold to cause a reZID of a volume. **IMPORTANT:** For NSS on Linux, a rebuild automatically causes a reZID of a volume if the 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 the rebuild, which increases the time to rebuild a volume by about 50%. For more information about reZID, see Section 17.3, “ReZIDing Volumes in an NSS Pool,” on page 223. **Default:** 0xeffffffff
`-u, --unknown-loss-prune`	Prunes if losses unknown.

VOPTION

VOPTIONs can be used only for the verify action.

Option	Description
`-q, --quick`	Skips cross-tree validations.

B.16.4 Files

/opt/novell/nss/sbin/ravsui

The Rebuild and Verify Simple User Interface (RAVSUI, `ravsui(8)`) utility file.

/var/opt/novell/log/nss/rav/

The default location of the log file for the RAVSUI utility’s rebuild or verify actions. The path is the directory where the file is stored and does not include the filename itself. You can specify a different path for the log file by using the -P=PATH option.

B.16.5 Note

Before you run the RAVSUI utility to perform a pool verify or pool rebuild, you must put the pool into maintenance mode.

Log in as the `root` user, then open a terminal console.

At a terminal prompt, enter

```
nsscon
```

In nsscon, enter

```
nss /PoolMaintenance=poolname
```

Replace *`poolname`* with the name of the pool you plan to rebuild or verify, such as `POOL1`.

If you do not place the pool in maintenance mode before starting the utility, you receive NSS Error 21726:

```
NSS error: PoolVerify results

   Status: 21726

     Name: zERR_RAV_STATE_MAINTENANCE_REQUIRED

   Source: nXML.cpp[1289]
```

B.16.6 Example

Log in as the `root` user, then open a terminal console.

At a terminal prompt, enter

```
nsscon
```

At the `nsscon` prompt, put the pool that you want to verify (such as `POOL1`) in maintenance mode by entering

```
nss /PoolMaintenance=POOL1
```

At the `nsscon` prompt, start the RAVSUI utility to verify the pool by entering

```
ravsui -q verify POOL1
```

This command checks the integrity of the data in `POOL1` in the current tree. It skips cross-tree validations.

B.17 RAVVIEW (Linux)

The Rebuild and Verify View (RAVVIEW) utility displays specified rebuild or verify log files in human-readable format. The log files are generated by the Rebuild and Verify Simple User Interface (RAVSUI, `ravsui(8)`) utility.

You must log in as the `root` user to run this utility. If the specified pool is active when you run the utility, the ZIDs of files that appear in the log are converted to their `/pathname/filename` format so that they are more easily understood by the reader.

B.17.1 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
```

Poolnames and filenames are case sensitive on Linux.

B.17.2 Arguments

The first mandatory argument specifies the type of log file you want to view. This can be rtf, rtfn, vbf, or vbfn. The second argument specifies information about the log you want to view.

IMPORTANT: In the raw log files, files that were acted upon are identified by their ZIDs, not the path and filename. If the pool that was rebuilt or verified is active when you view its log, the utility converts each ZID to display the path and filename instead.

First Argument	Description	Second Argument
rtf	Specifies that you want to view a rebuild text file. You must specify the filename of the log you want to view. The utility displays any messages that occurred during the rebuild.	*filename*
rtfn	Specifies that you want to view the newest rebuild text file in a specified path for a specified poolname. Specify the path to the logs by using the -P=PATH option. The utility displays any messages that occurred during the rebuild.	*poolname*
vbf	Specifies that you want to view a verify binary file. You must specify the filename of the log you want to view. The utility converts the binary file into a human-readable format.	*filename*
vbfn	Specifies that you want to view the newest verify binary file in a specified path for a specified poolname. Specify the path to the logs by using the -P option. The utility converts the binary file into a human-readable format.	*poolname*

Replace *filename* with the path and name of the log file for the rebuild or verify process that you want to view.

Replace *poolname* with the name of the NSS pool that was rebuilt or verified.

B.17.3 Options

This section describes the OPTION, VOPTION, and NOPTION options available for the RAVVIEW utility. Mandatory arguments for long options are mandatory for short options too.

OPTION

General options can be used for viewing rebuild or verify logs. Use them in combination with any of the rtf, rtfn, vbf, or vbfn arguments.

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

VOPTIONs are available only for viewing the verify binary files. Use them in combination with the vbf or vbfn arguments.

Option	Description
`-a, --actions`	Displays action information.
`-H, --histograms`	Displays histograms.
`-o, --object-details`	Displays object details.

NOPTION

NOPTIONS are available for viewing the newest log for a specified pool. Use this option in combination with the rtfn and vbfn arguments.

Option	Description
`-P, --path=PATH`	Specifies the path to RAV files. **Default:** `/var/opt/novell/log/nss/rav/`

B.17.4 Files

/opt/novell/nss/sbin/ravview

The Rebuild and Verify View utility file.

/var/opt/novell/log/nss/rav/

The default location of the log file for the Rebuild and Verify Simple User Interface (RAVSUI, `ravsui(8)`) utility's rebuild or verify actions. The path is the directory where the file is stored and does not include the filename itself.

You can specify a different path for the log file by using the -P=PATH option. This should be the same path that you used when you ran the RAVSUI utility.

B.17.5 Note

In the raw log files, files that were acted upon are identified by their ZIDs, not the path and name of the file. If the pool that was rebuilt or verified is active when you view its log, the RAVVIEW utility converts each ZID to display the path and filename instead.

B.17.6 Example

Log in as the root user, then open a terminal console.

At the terminal prompt, start NSSCON by entering

```
nsscon
```

View the newest rebuilt text file for POOL1 by entering

```
ravview -V -P=/var/opt/novell/log/nss/rav/ rtfn POOL1
```

In this command, the RAVVIEW utility uses the newest rebuild log in the `/var/opt/novell/log/nss/rav/` directory for `POOL1`.

B.18 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 might be accustomed 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 might need to mount the pool and its volumes.

NOTE: This utility protects against system failures only, not hardware failures.

B.18.1 Syntax

To rebuild a single pool, specifying the name, enter

```
nss /poolrebuild=pool_name
```

B.18.2 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.

B.19 REFRESHIDS (Linux)

Use this command to force a reset of the NSS ID caches on your OES 2 Linux server.

B.19.1 Syntax

Issue the command as the `root` user in a terminal console on the OES 2 Linux server where the NSS volume exists.

```
refreshids
```

```
refreshids --help
```

B.20 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.

B.20.1 Syntax

```
REMIRROR PARTITION logical_partition_number
```

Parameter	Description
logical_partition_number	Specifies the number of the logical partition you want to start remirroring.

B.20.2 Example

To start remirroring for logical partition 0x4, type

```
REMIRROR PARTITION 4
```

B.21 RIGHTS

- Section B.21.1, “RIGHTS (Linux),” on page 530
- Section B.21.2, “RIGHTS (NetWare),” on page 534

B.21.1 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
`trustee`	Adds or modifies a trustee on a file or directory.
`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`	Specifies the rights to be given to this trustee. For more information, see "MASK" on page 533. If the No Rights (n) option is assigned, the trustee is removed. If rights are not specified, the default assignment is Read and File Scan rights.
`-f, --file=`*`filename`*	Specifies the name of file or directory to assign trustees to. *Filename* is the path for the file or directory. For example: `-f /users/`*`username`*`/`*`userfile.sxi`* `--file=/designs/topsecret` If a file or directory is not specified, the current directory is used.

DOPTIONS

Option	Description
`-f, --file=`*`filename`*	Specifies the name of file or directory to delete trustees from. *Filename* is the path for the file or directory. If a file or directory is not specified, the current directory is used.

IOPTIONS

Option	Description
`-r, --rights=MASK`	Specifies the rights to be passed through the filter. For more information, see "MASK" on page 533. If rights are not specified, the default assignment is All Rights.
`-f, --file=`*`filename`*	Specifies the name of the directory where the filter is to be applied. *Filename* is the path for the directory. If a directory is not specified, the current directory is used.

EROPTIONS

Option	Description
`-f, --file=`*`filename`*	Specifies the name of file or directory where effective rights are to be calculated. *Filename* is the path for the file or directory. If a file or directory is not specified, the current directory is used.

SOPTIONS

Option	Description
`-f, --file=`*`filename`*	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 Fully Distinguished Name of a Novell eDirectory object, including the tree name. Use the *`username.context.treename`* format, such as

joe.engineer.acme_tree

If you use special characters in a username, you must escape those special characters in the command line.

For example, the $ (dollar sign) is a special character reserved to the shell and must be escaped. For the bash shell, the command could be written in one of two ways on the command line:

```
rights -f /media/nss/DATA/stuff -r none \$j\$o\$e.engineer.acme_tree

rights -f /media/nss/DATA/stuff -r none '$j$o$e.engineer.acme_tree'
```

If you are using another shell, the special characters might need a different escape technique. In this case, please refer to the shell documentation for this information.

MASK

The mask is a string of characters, with each character representing certain 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	Has all rights to the file or directory. Also can grant or revoke the Access Control right.
Read	r	Grants the right to open and read files in the directory.
Write	w	Grants the right to open and write to files in the directory.
Create	c	Grants the right to create files and subdirectories. The user can also salvage (undelete) deleted files.
Erase	e	Grants the right to erase files and directories. The user can also purge deleted files.
Modify	m	Grants the right to modify the content of files and directories, and change file attributes.
File Scan	f	Grants the right to display and search on file and directory names in the file system structure.
Access Control	a	Grants the right to add and remove trustees, and change trustee rights to files and directories.
No Rights	none	Revokes all rights.
All Rights	all	Grants all rights except Supervisor (rwcemfa)

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 500.

B.21.2 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
`path`	Specifies the path to the file, directory, or volume you want to modify or view rights to (you must always specify a path).
`+ | -`	Adds or deletes the specified rights. See Using RIGHTS.
`rights`	Specifies one or more file or directory rights. See File and Directory Rights.
`/`*`option`*	Replaces *option* with any available option. See RIGHTS Options.
`/?`	Displays online help. All other parameters are ignored when /? is used.
`/VER`	Displays 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`	Scrolls continuously through output.
`/F`	Displays the Inherited Rights Filter (IRF).
`/I`	Displays the trustee and group rights that created the inherited rights, and view where the inherited rights came from.
`/NAME=`*`username`*	Displays or modifies rights for the user or group listed. Replace *username* with the name of the user or group whose rights you want to view or modify.
`/S`	Displays or modifies subdirectories below the current level.
`/T`	Displays 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	Has all rights to the file or directory. Also can grant or revoke the Access Control right.
Read	R	Grants the right to open and read files in the directory.
Write	W	Grants the right to open and write to files in the directory.
Create	C	Grants the right to create files and subdirectories. The user can also salvage (undelete) deleted files.
Erase	E	Grants the right to erase files and directories. The user can also purge deleted files.
Modify	M	Grants the right to modify the content of files and directories, and change file attributes.
File Scan	F	Grants the right to display and search on file and directory names in the file system structure.
Access Control	A	Grants the right to add and remove trustees, and change trustee rights to files and directories.
No Rights	N	Revokes all rights.
Remove	REM	Removes the specified user or group as a trustee of the specified file or directory.
All	ALL	Grants 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
  ```

B.22 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 for NetWare.

The default SET parameter values give maximum performance for most systems. Server parameter values seldom need to be modified.

B.22.1 Syntax

```
set [parameter]=[value]
```

Parameter Categories

This section describes only those parameters related to managing NSS file systems on NetWare. For information about other `SET` options, see " SET " in the *OES 2 SP1: Utilities Reference for NetWare*.

"Common File System Parameters" on page 537
"Disk Parameters" on page 541
"Novell Storage Services Parameters" on page 543

B.22.2 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 are not 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.

B.22.3 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 = number`	Specifies how many transactions can occur at the same time. **Supported Values:** 100 to 10000 (default) **Owning NLM:** `filesys.nlm`
`Maximum Concurrent Directory Cache Writes = number`	Specifies how many write requests from directory cache buffers are put in the elevator before the disk head begins a sweep across the disk. **Supported Values:** 5 to 500 **Default:** 75 A high number creates more efficient write requests. A low number creates more efficient read requests. **Owning NLM:** `filesys.nlm`
`Minimum File Delete Wait Time = time`	Specifies how long a deleted file remains salvageable on the volume. **Supported Values:** 0 seconds to 7 days **Default:** 1 minute 5.9 seconds Files deleted for less than this minimum are not automatically purged even if the volume is full and users can't create new files. **Owning NLM:** `filesys.nlm`
`Immediate Purge of Deleted Files = value`	**Supported Values:** On or Off (default) If this parameter is set to On, all files are purged immediately when they are deleted. **Owning NLM:** `filesys.nlm`

Parameter	Description
`Compression Daily Check Stop Hour = number`	Specifies the hour when you want the file compressor to stop scanning enabled volumes for files that need to be compressed. **Supported Values:** 0 to 23 **Default:** 6 Hours are specified by a 24-hour clock. For example, 0=midnight; 23=11 p.m. This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`
`Compression Daily Check Starting Hour = number`	Specifies the hour when you want the file compressor to start scanning enabled volumes for files that need to be compressed. **Supported Values:** 0 (default) to 23 Hours are specified by a 24-hour clock. For example, 0=midnight; 23=11 p.m. **Note:** 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. This parameter can be set in the `startup.ncf` file.
`Minimum Compression Percentage Gain = number`	Sets the minimum percentage a file must compress to remain in a compressed state. **Supported Values:** 0 to 50 **Default:** 20 This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`

Parameter	Description
`Enable File Compression = value`	Specifies whether file compression is suspended. **Supported Values:** On (default) or Off On allows file compression on compression-enabled volumes. Off suspends compression. Immediate Compression requests are queued until the value is reset to On, when the files meeting criteria are compressed. This parameter can be set in the `startup.ncf`. **Owning NLM:** `filesys.nlm`
`Maximum Concurrent Compressions = number`	Specifies the maximum concurrent or simultaneous compressions allowed. **Supported Values:** 1 to 8 **Default:** 2 Concurrent compressions can occur only if there are multiple volumes. This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`
`Convert Compressed to Uncompressed Option = value`	Specifies what the file system does with a decompressed version of a file after the server has decompressed it. **Supported Values:** 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 **Default:** 1 This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`

Parameter	Description
`Decompress Percent Disk Space Free to Allow Commit = number`	Specifies the percentage of free disk space required on a volume for file decompression to permanently change compressed files to decompressed. **Supported Values:** 0 to 75 **Default:** 10 This parameter prevents newly decompressed files from filling up the volume. This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`
`Decompress Free Space Warning Interval = number`	Specifies 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. **Supported Values:** 0 seconds to 29 days 15 hours 50 minutes 3.8 seconds **Default:** 31 minutes 18.5 seconds This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`
`Deleted Files Compression Option = number`	Specifies whether and when deleted files are compressed. **Supported Values:** 0 = Don't Compress deleted files 1 = Compress deleted files the next day 2 = Compress deleted files immediately **Default:** 1 This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`
`Days Untouched Before Compression = number`	Specifies the number of days the system waits after a file was last accessed before it is compressed. **Supported Values:** 0 to 100000 **Default:** 14 This parameter can be set in the `startup.ncf` file. **Owning NLM:** `filesys.nlm`

B.22.4 Disk Parameters

The following Disk parameters control aspects of disk reads and writes.

Parameter	Description
`Auto scan for devices = value`	Automatically scans for devices when a HAM is loaded. If it is turned off, the operator must issue a Scan for New Devices on the command line before the devices are seen. **Supported Values:** On, Off; Default: On **Owning NLM:** `nwpa.nlm`
`Auto load of CDM modules = value`	Automatically scans and loads defined generic Custom Device Modules (CDMs). If it is turned off, the desired CDMs must be loaded manually or added to the `startup.ncf` file. **Supported Values:** On, Off; Default: On **Owning NLM:** `nwpa.nlm`
`Sequential Elevator Depth = number`	Sets the maximum elevator depth for sequential requests. 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. **Supported Values:** 0 to 4294967295, Default: 8 You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`
`Enable IO Handicap Attribute = value`	Enables drivers and applications to inhibit read requests from one or more devices. **Supported Values:** On, Off; Default: Off. Setting this parameter to On enables the Inhibit attribute to function. Setting this parameter to Off prevents the attribute from functioning. **IMPORTANT:** Do not set this attribute to On unless instructed to do so by a device manufacturer. You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`
`Mirrored Devices Are Out of Sync Message Frequency = time`	Sets the frequency (in minutes) for checking out-of-sync devices. **Supported Values:** 5 to 9999, Default: 28 You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`

Parameter	Description
`Remirror Block Size = number`	Sets the remirror block size in 4 KB increments. (1=4 KB, 2=8 KB, 8=32 KB, etc.) **Supported Values:** 1 to 8, Default: 1 **Owning NLM:** `mm.nlm`
`Concurrent Remirror Requests = number`	Sets the number of remirror requests per Mirror object. **Supported Values:** 2 to 32 Default: 32 **Owning NLM:** `mm.nlm`
`Ignore Partition Ownership = value`	Specifies whether partitions owned by other servers in a cluster can be activated on this server. Setting this parameter to On allows partitions owned by other servers in the cluster to be activated on this server. **Supported Values:** On, Off; Default: Off **Owning NLM:** `mm.nlm`
`Ignore Disk Geometry = value`	Creates nonstandard and unsupported partitions. **Supported Values:** On, Off; Default: Off 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. **Caution:** Setting this parameter to On might harm other file systems contained on the disk. You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`
`Multi-path Support = value`	Enables or disables support for multipath devices or adapter failover. **Supported Values:** On, Off; Default: Off **Owning NLM:** `mm.nlm`
`Enable Hardware Write Back = value`	Enable hardware write back, if supported. 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. **Supported Values:** On, Off; Default: Off You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`

Parameter	Description
`Enable Disk Read After Write Verify = value`	Controls whether information written to disk is read back and compared with the original data. **Supported Values:** On, Off; Default: Off Setting this parameter to On might decrease performance significantly. You can set this parameter in the `startup.ncf` file. **Owning NLM:** `mm.nlm`
Auto LFVMount	Controls whether any local FAT volumes not recognized by DOS are automatically mounted and assigned a local drive letter when they are detected. **Supported Values:** On, Off; Default: Off On automatically mounts and assigns a local drive letter when the drive is detected. When set to Off, you must use the LFVMOUNT command to manually mount local FAT volumes. **Owning NLM:** `fatfs.nlm`

B.22.5 Novell Storage Services Parameters

The following are the default settings in NetWare 6.5, OES NetWare, and later 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`
Emulate Classic NetWare Directory Quotas	`Off`
NSS Work To Do Count	`50`
This parameter can be set only in the `nssstart.cfg` file.	
NSS Maximum Cache Balance Buffers Per Session	`1024`
NSS Cache Balance Timer	`30`
NSS Cache Balance Percent	`60`
NSS Cache Balance Enable	`On`

Parameters	Default Value
NSS Buffer Flush Timer	1
In OES 2 Linux, this timer is not used. It has been replaced by the Group Write timers for Journal, Metadata, and User Data.	
NSS Minimum OS Cache Buffers	256
NSS Minimum Cache Buffers	256
NSS File Flush Timer	10
NSS Closed File Cache Size	50000
This parameter can be set only in the nssstart.cfg file.	
NSS Open File Hash Shift	16
NSS Auth Cache Size	1024
This parameter can be set only in the nssstart.cfg file.	
NSS ASCII Name Cache Enable	On
This parameter can be set only in the nssstart.cfg file.	
NSS Name Cache Enable	On
This parameter can be set only in the nssstart.cfg file.	
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

B.23 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.

B.23.1 Syntax

The `poolverify` command is a read-only assessment of the pool. At the server console, enter

```
nss /poolverify=[poolname]
```

If you do not enter the NSS poolname parameter, you are 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.

B.23.2 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=[poolname]`

   If you do not enter the NSS pool name parameter, you are 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.

B.24 VOLUMES (Linux, NCP Utility)

Use this utility at the ncpcon (NCP Console utility) prompt (Linux) to list mounted volumes or information about a specified volume. For information about the NSS `volumes` command, see Section A.42.2, “Volumes Command,” on page 495.

B.24.1 Syntax

Command	Description
`volume name` `nvpcon volume name`	Displays details about the specified volume. Linux is case-senstive, so make sure that you enter the volume name in all caps, such as `volume VOL1`
`volumes` `ncpcon volumes`	Displays general information about all mounted volumes.

B.24.2 Using VOLUMES

When you execute VOLUMES at the ncpcon prompt, a list of the mounted volumes is displayed.

For example, on Linux, a simple list of volumes is displayed:

```
Mounted Volumes
        SYS
        USERS
        TEST
        _ADMIN
        VOL1
5 volumes mounted.
```

B.24.3 Using VOLUME Name

When you execute `volume` *`name`*, the screen displays detailed information about the specific volume.

The following is an example of the output on Linux in response to entering `volume USERS` at the ncpcon prompt, or entering `ncpcon volume USERS` at the terminal console prompt:

```
           Volume: USERS
           Status: online mounted NSS "user quotas" "directory quotas" salvageable
      Mount point: /media/nss/USERS
Shadow Mount point: (null)
         Capacity: 8.83 GB
               ID: 4
             GUID: 9a894a30-70a3-01dd-80-00-32b3b21ae612
        Pool Name: POOL1
```

Upgrading Legacy NSS and NetWare Traditional Volumes

C

This section describes how to convert legacy NetWare® 5 NSS volumes and NetWare Traditional volumes to NSS volumes for use with Novell® Open Enterprise Server NetWare.

- Section C.1, "Migrating Legacy NSS Volumes and NetWare Traditional Volumes to OES Linux," on page 547
- Section C.2, "Copying and Upgrading Legacy Volumes with Volume Copy Upgrade," on page 547
- Section C.3, "Upgrading Legacy NSS Volumes," on page 554
- Section C.4, "Upgrading Legacy NSS Volumes in Clusters," on page 555

C.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 to Linux. One way to do this is by using the Volume Copy Utility (VCU). For information about VCU, see Section C.2, "Copying and Upgrading Legacy Volumes with Volume Copy Upgrade," on page 547.

For information about other data migration options, see "Migrating and Consolidating Existing Servers and Data" in the *OES 2 SP1: Planning and Implementation Guide*.

C.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 to NSS volumes. Volume copy upgrade moves the volume to an NSS pool and all of the volume's file and directory attributes remain the same.

IMPORTANT: For Traditional NetWare volumes, VCU does not migrate the `deleted.sav` directory and the deleted files it contains to the new NSS volume. The deleted files are essentially purged. If you need to salvage deleted files, do it before using VCU to migrate the data.

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 C.2.1, "Understanding Volume Copy Upgrade," on page 548
- Section C.2.2, "Using VCU from the Command Line," on page 549
- Section C.2.3, "Using VCU from a Java Application," on page 551
- Section C.2.4, "After the Volume Copy or Volume Upgrade," on page 554

C.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 NetWare Traditional volumes to NSS volumes.
- You can run `vcu.nlm` from a Windows client.

VCU uses NetWare Traditional volume APIs that have 32 bit file sizes, so it works only with file sizes of up to 4 GB, which is the Traditional volume maximum. The APIs use local code-page ASCII for all filename representations. VCU cannot correctly copy filenames on mixed-language systems, where the ASCII code-page might be English (437) but where Unicode characters from other languages may exist in the filenames.

IMPORTANT: VCU uses traditional APIs only, and is not Unicode enabled.

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 24.10, "Viewing Compression Statistics with the NDIR Command," on page 343. 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	`vcu [/p|/l|/d|/m] origVolName nssPool`
Restore the Original Volume	`vcu [/r|/m] origVolName renamedVolName`
Help	`vcu /h`

The VCU parameters are listed in the following table:

VCU Parameters and Help	Description
`origVolName`	The name of the original NetWare Traditional volume that you want to copy to a new NSS volume of the same name.
`nssPool`	The pool where the new NSS volume resides.
`renamedVol`	The name assigned to the original NetWare Traditional volume that you successfully copied to the NSS volume.
`/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 (`Dst_Vol:error.out`).
`/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 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
```

C.2.2 Using VCU from the Command Line

- "Copying Volumes" on page 550
- "Restoring Volumes" on page 550

Copying Volumes

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

WARNING: Before using this utility to copy Traditional 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 548.

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 volume, the new NSS volume is named `vcu` *xxxx*`_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 554 and perform any required follow-up actions.

Restoring Volumes

1 At the server console, enter

```
vcu /r VCUTargetVolName VCUTargetVOlName_old
```

For example, if you want to restore `Vol2`, enter the following:

```
vcu /r Vol2_new Vol2_old
```

VCU deletes `Vol2_new`, then renames `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.

C.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 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 551.

The following graphic is an example of the VCU Copy screen.

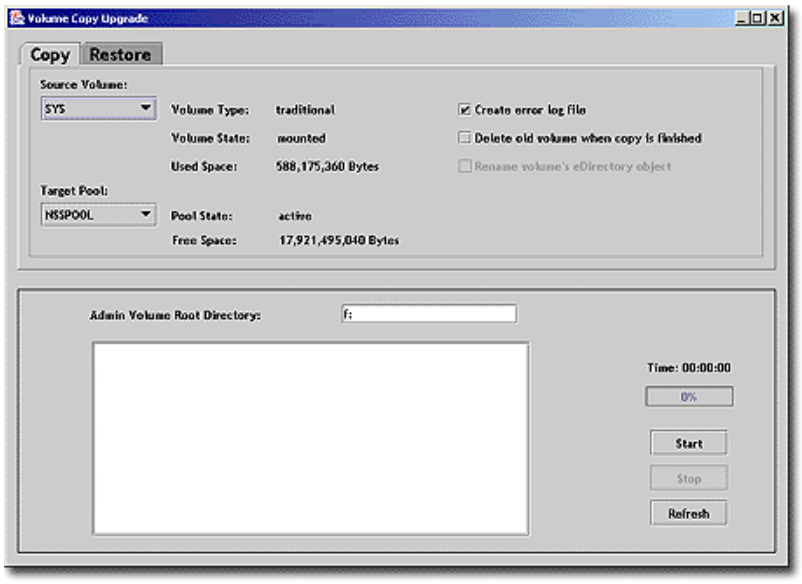

3 Click the *Copy* tab.

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 Section C.2, "Copying and Upgrading Legacy Volumes with Volume Copy Upgrade," on page 547.
`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 "Copying Volumes" on page 550.

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 554 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, 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 551.

The following graphic is an example of the VCU Restore dialog box.

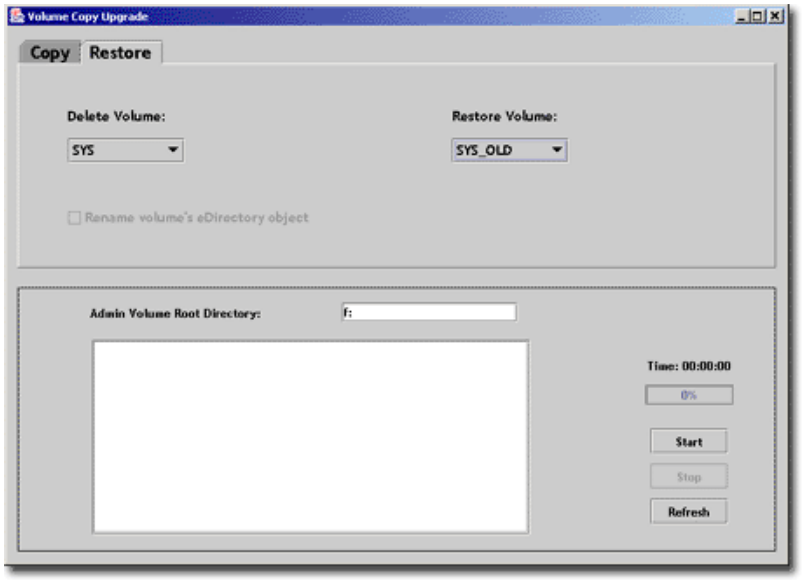

3 Click the *Restore* tab.

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.

C.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.

C.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 number of files 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 `nss /zlssvolumeupgrade` 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.

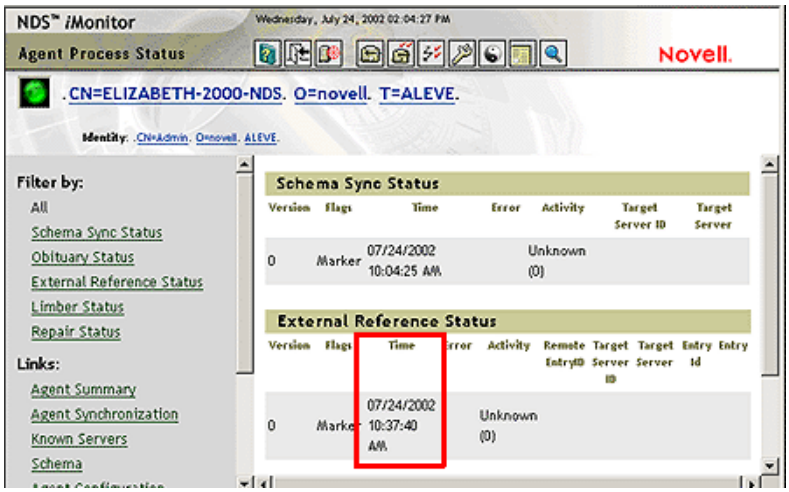

2 At the server console, enter one of the following:

- `nss /zlssvolumeupgrade=all`

  Specify `all` to upgrade all the NetWare 5 volumes on the selected server.

- `nss /zlssvolumeupgrade=volumename`

  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.

C.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 NetWare Clusters" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for NetWare Administration Guide*.

Using EVMS to Manage Devices with NSS Volumes (Linux)

D

In order to get the full features for Novell® Storage Services™ on Novell Open Enterprise Server 2 Linux servers, the devices you use for NSS volumes must be managed by the Enterprise Volume Management System (EVMS). This section describes how to configure EVMS services on Linux devices to allow you to use NSS.

- Section D.1, "Installing Linux with EVMS as the Volume Manager on the System Device," on page 557
- Section D.2, "Using Linux POSIX File Systems and NSS Pools and Volumes on the Same Device," on page 558
- Section D.3, "Viewing Free Space Available to EVMS-Managed Devices," on page 560
- Section D.4, "Making Devices Available to EVMS," on page 561
- Section D.5, "Additional EVMS Documentation," on page 567

D.1 Installing Linux with EVMS as the Volume Manager on the System Device

To use NSS on the same drive that contains the `/boot`, `swap`, and root (`/`) partitions, that drive must be managed by EVMS. The default volume manager in OES Linux is the Linux Volume Manager 2 (LVM2), not EVMS. Unless you modify the partitioning scheme during the operating system install, the system device is not managed by EVMS, and you cannot create NSS data volumes on that drive.

The simplest solution is to put your Linux system volumes on one device that is managed by LVM2 (the default configuration), then add devices that can be configured to use EVMS for your NSS pools and volumes.

Using EVMS to manage the system device allows you to later add NSS pools and volumes on any unparttioned free space on it. However, you must modify the partitioning scheme to use EVMS during the install. It is not possible to change the volume manager for the system device after the install.

From the YaST Installations Settings page, you go to the Partitioner page, then do the following to set up the EVMS partitioning scheme. For detailed instructions, see "Installing Linux with EVMS as the Volume Manager of the System Device" in the *OES2 SP1: Linux Installation Guide*.

- ❑ Delete the proposed LVM2-based partitioning solution.
- ❑ Create a boot partition of about 300 MB on the system disk with a `/boot` mount point.
- ❑ Create a system partition of 5 GB to 11 GB on the system disk where you will create the EVMS container for the `swap` and root (`/`) volumes.
- ❑ Do not partition the remaining space on the system device. This unpartitioned free space can be used after the install for NSS pools and volumes or for Linux POSIX file systems.
- ❑ Select EVMS Configuration, then select the system partition. This step ensures the device is managed by EVMS.

- ❑ Create an EVMS container for the system partition of type *Linux LVM2*.
- ❑ Create the `swap` volume of about 1 GB in size in the EVMS container.
- ❑ Create the / (root) volume of 4 GB to 10 GB in the EVMS container.

For example, Figure D-1 illustrates how a 30-GB device can be configured during the install to enable EVMS to be used as the volume manager. The `/boot` volume is placed at the beginning of the device. An EVMS container is created, then the `swap` and / (root) volumes are created in it. The remainder of the device is left as unpartitioned free space.

Figure D-1 *Example of a System Device Configured for EVMS and NSS*

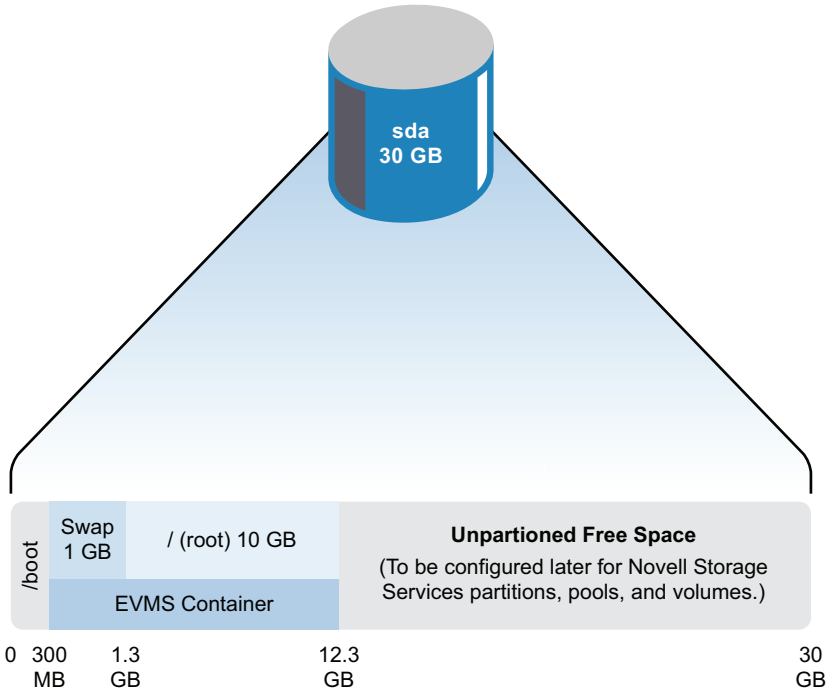

IMPORTANT: Only the boot and system partitions are configured. The remainder of the device remains as unpartitioned free space where you can create NSS pools and volumes after the install.

D.2 Using Linux POSIX File Systems and NSS Pools and Volumes on the Same Device

You can use the free space on devices for NSS or Linux POSIX file systems by using EVMS to manage devices. Consider the guidelines below when working with EVMS-managed devices.

When you use NSS tools to create a pool on a new EVMS-managed device, the tools automatically carve out a partition with the DOS Segment Manager so that the device can be used later for either NSS or Linux POSIX file systems. Then the NetWare Segment Manager creates the NSS partition and pool.

IMPORTANT: Always use NSS tools (NSSMU and the Storage plug-in to iManager) to create NSS pools.

For EVMS-managed data devices, Linux POSIX file systems require that the Linux partitions be managed by the DOS Segment Manager (DOSSegMgr). Consider the following guidelines when working with evmsgui to create a Linux partition:

Scenario	To create the Linux partition in evmsgui:
Disk is not initialized.	1. Select *No* at the prompt to initialize. On a pure SLES system (no NSS), ignore this step. 2. Delete the disk object from the *Volumes* tab. 3. Add the DOS Segment Manager to the device. 4. Create the Linux partition as usual.
Free space is controlled by the DOS Segment Manager.	Create the Linux partition as usual.
Free space is controlled by the NetWare Segment Manager.	Do one of the following: ◆ If no partitions are on the device, remove the NetWare Segment Manager from the device, add the DOS Segment Manager, then create the partition as usual. **WARNING:** Changing the segment manager initializes the disk again (destroys existing data), so you only want to do this with disks that have no partitions on it, or if you do not want any of the partitions that are currently on the disk. ◆ If partitions exist, reboot the server to automatically give control of the device back to the DOS Segment Manager, then create the partition as usual.

If no partitions are on the device, do the following to add a DOS Segment Manager to a non-clustered device:

1 Log in as the `root` user, open a terminal console, then enter

WARNING: EVMS administration utilities (evms, evmsgui, and evmsn) should not be running when they are not being used. EVMS utilities lock the EVMS engine, which prevents other EVMS-related actions from being performed. This affects both NSS and Linux POSIX volume actions.

NSS and Linux POSIX volume cluster resources should not be migrated while any of the EVMS administration utilities are running.

2 If necessary, remove the NetWare Segment Manager:

2a Click the *Disks* tab, then locate and select the device, such as device `sdb`.

2b Right-click, then select *Remove segment manager from Object*.

This option appears only if there is an existing segment manager for the selected disk.

2c Select the listed segment manager, click *Remove*, then click *OK*.

WARNING: All data on the selected disk space is destroyed.

2d Click *Save*, then click *Save* again to save your changes.

3 Add the DOS Segment Manager.

3a From the evmsgui menu, click *Actions* > *Add* > *Segment Manager to Storage Object*.

3b On the Add Segment Manager to Storage Object page, choose *DOS Segment Manager*, then click *Next*.

3c On the Select Plugin Acceptable Objects page, choose the device where you want to add the segment manager, then click *Next*.

3d On the Configurable Options page, select the disk type (Linux is the default), click *Add*, then click *OK*.

3e Click *Save*, then click *Save* again to save your changes.

4 Create a segment for the DOS Segment Manager.

The DOS Segment Manager requires you to create a segment before creating an EVMS volume. Without a segment, the additional segment manager does not appear when you attempt to create an EVMS volume.

4a From the evmsgui menu, click *Actions* > *Create* > *Segment*.

4b On the Create Disk Segment page, select *DOS Segment Manager*, then click *Next*.

4c On the Select Plugin Acceptable Objects page, choose device where you want to add the segment, then click *Next*.

4d Specify the size of the segment, the partition type (such as Linux LVM), click *Create*, then click *OK*.

4e Click *Save*, then click *Save* again to save your changes.

For information about adding or changing segment managers when you are clustering a shared device with Novell Cluster Services, see "Creating Linux POSIX Volumes on Shared Disks" in the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.

D.3 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 already managed by EVMS, or they are free unpartitioned devices. After you select free space from an unpartitioned device, NSS automatically sets up the device to be managed by EVMS.

The OES 2 Linux install uses LVM2 by default as the volume manager of the system device, but NSS requires EVMS. The Linux 2.6 kernel prevents multiple volume managers from managing the same device. If you use LVM2 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 2 Linux install, the `sda` device contains the boot device and the root volume. Unless you have previously made EVMS the volume manager for that device, it is managed by LVM2 and cannot be used to create NSS pools. To prevent the display of space from `sda`, edit the `/etc/evms.conf` file as follows:

```
sysfs_devices {

...

exclude = [sda]

}
```

To exclude multiple devices (such as `/dev/sda` and `/dev/sdb`), use the following format:

```
exclude = [sda sdb]
```

D.4 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 ways to make the devices available to EVMS:

- Section D.4.1, "Mounting Your System Device with EVMS," on page 561
- Section D.4.2, "Converting LVM2 Volumes to EVMS Volumes," on page 566
- Section D.4.3, "Patching the 2.6 Kernel (Unsupported)," on page 567

D.4.1 Mounting Your System Device with EVMS

If you did not configure EVMS as the root partition manager during installation, the following post-install work-around 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.

- "Disabling boot.lvm and boot.md" on page 562
- "Enabling the boot.evms Service" on page 562
- "Editing the /etc/init.d/boot.evms Script" on page 562
- "Editing the /etc/fstab File" on page 563
- "Editing the Boot Loader File" on page 563
- "Forcing the RAM Disk to Recognize the Root Partition" on page 565

- "Rebooting the Server" on page 566
- "Verifying that EVMS Manages the Boot, Swap, and Root Partitions" on page 566

Disabling 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* > *System Services (Runlevel)* > *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 "Enabling the boot.evms Service" on page 562.

Enabling the boot.evms Service

Enable the `boot.evms` service so that it boots devices when you restart the server.

1. In YaST, click *System* > *System Services (Runlevel)* > *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 "Editing the /etc/init.d/boot.evms Script" on page 562.

Editing 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 affect until the server is restarted. Do not reboot at this time.

4 Continue with "Editing the /etc/fstab File" on page 563.

Editing 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*.

Although the `/boot` partition is not managed by EVMS, the `boot.evms` script forces EVMS to scan all the partitions at boot time, including the `/boot` partition, and it activates `/boot` under the `/dev/evms` directory. Therefore, this makes `/boot` a partition that is discovered by EVMS at startup, and requires that the device's path be listed under `/dev/evms` in the `fstab` file so it can be found when booting with `boot.evms`.

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.

   The changes do not take affect until the server is restarted. Do not reboot at this time.

6 Continue with "Editing the Boot Loader File" on page 563.

Editing the Boot Loader File

When you boot the system, the kernel reads the boot loader file for information about your system. For Grub, this is the `/boot/grub/menu.1st` file. For LILO, this is the `/etc/lilo.conf` file.

You must edit the boot loader file to modify the mount location of partitions so they are mounted under the `/dev/evms` directory. For example, change /dev/*sda1* to /dev/evms/*sda1*. Make sure to replace the path for all lines that contain device paths in the files. You can modify the boot loader file by editing fields in YaST, or use a text editor to modify the file directly.

IMPORTANT: When working in the boot loader 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.

To modify the boot loader file in the YaST Control Center:

1. Log in as the `root` user.
2. In Yast, select *System* > *Boot Loader*.
3. Modify the boot loader image so that the root file system is mounted as `/dev/evms/` instead of `/dev/`.
   - **3a** Select the boot loader image file, then click *Edit*.
   - **3b** Edit the device path in the *Root Device* field.

     For example, change the *Root Device* value from

     `/dev/sda2`

     to

     `/dev/evms/sda2`

     Replace *`sda2`* with the actual device on your machine.
   - **3c** Edit any device paths in the *Optional Kernel Command Line Parameter* field.
   - **3d** Click *OK* to save the changes and return to the Boot Loader page.
4. Modify the failsafe image so that the failsafe root file system is mounted as `/dev/evms/` instead of `/dev/`.
   - **4a** Select the failsafe image file, then click *Edit*.
   - **4b** Edit the device path in the *Root Device* field.
   - **4c** Check the *Other Kernel Parameters* field and make changes if needed.
   - **4d** Click *OK* to save the change and return to the Boot Loader page.
5. Click *Finish*.
6. Continue with "Forcing the RAM Disk to Recognize the Root Partition" on page 565.

To edit the boot loader file in a text editor:

1. Log in as the `root` user.
2. Open the boot loader file in a text editor.

   For Grub, this is the `/boot/grub/menu.1st` file. For LILO, this is the `/etc/lilo.conf` file.
3. Locate the line that contains the `root=` parameter.

   For example, if your root file system uses device *sda1*, look for a line similar to this:

   ```
   kernel (hd0,0)/vmlinuz root=/dev/sda1 vga=0x31a  splash=silent showopts
   ```
4. Modify the mount location from `/dev` to `/dev/evms` so it can be managed by EVMS.

   For example, after the change, the line looks like this:

```
kernel (hd0,0)/vmlinuz root=/dev/evms/sda1 vga=0x31a  splash=silent showopts
```

5 Repeat Step 3 and Step 4 to locate other lines in the file that need to be similarly modified.

6 Save the file.

The changes do not take affect until the server is restarted. Do not reboot at this time.

7 Continue with "Forcing the RAM Disk to Recognize the Root Partition" on page 565.

Forcing 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 because the `/initrd` file was configured during the install without using EVMS. 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 the `root_evms` parameter 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 the `root_evms` parameter from

```
root_evms=1
```

to

```
root_evms=
```

then save the file.

The changes do not take affect until the server is restarted. You reboot in the next task.

6 Continue with "Rebooting the Server" on page 566.

Rebooting the Server

1. Reboot the server to apply the post-install configuration settings.
2. Continue with "Verifying that EVMS Manages the Boot, Swap, and Root Partitions" on page 566.

Verifying 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.

D.4.2 Converting LVM2 Volumes to EVMS Volumes

The following instructions were modified from Converting to EVMS (http://evms.sourceforge.net/convert.html) at the EVMS Project Web site.

1. Open a terminal console, then log in as the `root` user.
2. Enter the following command to run the EVMS GUI:

   ```
   evmsgui
   ```

3. In the Volume panel, review the names that EVMS reports for your existing volumes to find the volume that represents your `/boot`, `swap`, and root (`/`) volumes, then write down the names for future reference.

   All storage objects that EVMS manages or activates (such as the `/boot` partition when using the `boot.evms` script) are mounted under `/dev/evms`. For example:

Storage Object	Mount Location
A disk segment of disk `sda` called `/dev/sda5`	`/dev/evms/sda5`
A software RAID device called `/dev/md1`	`/dev/evms/md/md1`
An LVM2 volume called `/dev/lvm_group/lvm_volume`	`/dev/evms/lvm/lvm_group/lvm_volume`

4. In a text editor, edit the `/etc/fstab` file to use the EVMS volume names.

   For example, change the following entry for an LVM2 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 file systems are mounted using the EVMS volume manager.

5 Update the boot scripts by doing the following:

- The command `evms_activate` must be run from your boot scripts in order to automatically activate your volumes so their file systems can be mounted.

  The `evms_activate` command discovers all EVMS volumes and activates kernel devices for all volumes that are not yet active. If EVMS volumes are listed in the `/etc/fstab`, `evms_activate` should be run before `/etc/fstab` is processed.

- If you run software-RAID (`md`) or LVM2 tools during your boot scripts, remove or disable those commands.

6 Update the boot loader.

To mount your root file system using EVMS, you must install the EVMS `init-ramdisk` image and update your boot-loader configuration. For information about updating the boot loader, see "Editing the Boot Loader File" on page 563.

7 Reboot your system.

When your system boots, the kernel loads the `init-ramdisk`, which runs the EVMS tools to activate your volumes and mount your root fleshiest. 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.

D.4.3 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 allows you to share your disks between EVMS and the kernel's built-in partitioning code, if that's the choice you want to make for your system.

WARNING: Adding this patch makes your server an unsupported configuration with respect to Novell Support.

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

D.5 Additional EVMS Documentation

For information about using EVMS on SUSE® Linux Enterprise Server 10 and OES 2 Linux servers, see the *SLES 10 SP2: Storage Administration Guide*.

For more information about the Enterprise Volume Management System, see the following open source documentation:

- EVMS Project at SourceForge.net (http://sourceforge.net/projects/evms/)
- *EVMS GUI User Guide* (http://evms.sourceforge.net/user_guide/#assignseggui)
- EVMS GUI Screen Shots (http://evms.sourceforge.net/gui_screen/)

Using NSS on Devices Managed by Non-EVMS Volume Managers (Linux)

E

The Novell® Storage Services™ file system can exist on devices managed by any volume manager. However, when using non-EVMS volume managers, some of the NSS features are not available because the NSS management tools require EVMS to carry out related commands.

WARNING: The LVM-based NSS configurations described in this appendix are untested and are not supported by Novell Support. Also, some key features, such as clustering with Novell Cluster Services™, are not available on non-EVMS managed volumes. For more information see Section E.1.2, "NSS Limitations for Non-EVMS Volume Managers," on page 570.

These instructions are only for those who are willing to assume the risks and limitations described. Everyone else should follow the instructions located in Chapter 3, "Installing and Configuring Novell Storage Services," on page 45, and use EVMS as their NSS volume manager.

This section describes the following:

- Section E.1, "FAQs About Using EVMS with NSS," on page 569
- Section E.2, "Guidelines for Using NSS with Non-EVMS Volume Managers," on page 570
- Section E.3, "Creating and Mounting NSS Pools and Volumes Using Linux Commands," on page 571
- Section E.4, "Configuring Default Mount Settings for NSS Pools and Volumes," on page 575
- Section E.5, "Expanding NSS Pools on Non-EVMS-Managed Devices," on page 577

E.1 FAQs About Using EVMS with NSS

- Section E.1.1, "Why Should I Use EVMS with NSS?," on page 569
- Section E.1.2, "NSS Limitations for Non-EVMS Volume Managers," on page 570

E.1.1 Why Should I Use EVMS with NSS?

Using EVMS with NSS allows you to take advantage of the full range of capabilities of NSS and the NSS management tools, including the Storage plug-in to Novell iManager 2.7, the NSS Management Utility (NSSMU), the NSS Console Utility (NSSCON), and other NSS utilities and commands. NSS tools require EVMS for management functions related to devices, software RAIDS, partitions, pool snapshots, and creating pools.

For information about limitations for NSS on devices managed by non-EVMS volume managers, see "NSS Limitations for Non-EVMS Volume Managers" on page 570.

E.1.2 NSS Limitations for Non-EVMS Volume Managers

The basic NSS file system works on devices that are managed by any volume manager. However, the following NSS capabilities are unavailable or have limited availability for devices managed by a non-EVMS volume manager:

- **NSS Management Tools:** The following NSS management tools are affected:
  - **iManager:** The storage-related plug-ins for Novell iManager require that EVMS be used as the volume manager of any devices where you want to create and manage NSS file systems. The tool does not work with devices managed by non-EVMS volume managers.

    For information about storage-related plug-ins, see Section 9.1, "Novell iManager and Storage-Related Plug-Ins," on page 97.
  - **NSSMU:** NSSMU requires EVMS to manage partitions, manage devices, create and manage software RAIDs, create and manage pool snapshots, and create and mirror pools.

    If you create the pool with `mkfs` at the command line on a non-EVMS device, you can only create a Pool object using the *Pools > Update NDS* option, and create and manage both non-encrypted and encrypted volumes for the pool using the *Volumes* page.

    For information about NSSMU on Linux, see Section 9.2.1, "NSSMU for Linux Quick Reference," on page 113.
  - **Command Line:** EVMS is required for NSS command line commands and utilities.
- **NSS Software RAIDS:** NSS requires EVMS to create and manage software RAID devices. However, you can optionally create software RAIDs for your device using the Linux Device Mapper utility.
- **NSS Pool Snapshots:** NSS requires EVMS to create and manage pool snapshots.
- **Cross-Platform Compatibility:** NSS pools that are 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 with NSS volumes.

To use NSS with a non-EVMS volume manager, create partitions and pools with the Linux `mkfs` command, then use NSSMU to manage the volumes. For information, see Section E.3, "Creating and Mounting NSS Pools and Volumes Using Linux Commands," on page 571.

E.2 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 by using the guidelines in this section.

- Section E.2.1, "Guidelines for the Server Install," on page 571
- Section E.2.2, "NSS Management Considerations," on page 571

E.2.1 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 by 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 by using the *Software Updates* feature in YaST.

E.2.2 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, partitions, and pool snapshots require 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 E.3.1, "Using the Linux mkfs Command to Create NSS Pools," on page 572.
- Manage existing NSS pools and create and manage NSS volumes with NSSMU. For information, see Section 9.2.1, "NSSMU for Linux Quick Reference," on page 113.
- File system management tools are unaffected; they should work regardless of the volume manager that is managing the device.

  For information, see the following:

  - *OES 2 SP1: File Systems Management Guide*
    - Section 9.4, "Novell NetStorage," on page 118
    - Section 9.6.2, "Novell Client for Windows XP/2003 and Vista," on page 124

E.3 Creating and Mounting NSS Pools and Volumes Using Linux Commands

Use the following procedure to create and mount an NSS pool and volume.

- Section E.3.1, "Using the Linux mkfs Command to Create NSS Pools," on page 572
- Section E.3.2, "Creating a Partition," on page 572
- Section E.3.3, "Creating and Mounting an NSS Pool," on page 573
- Section E.3.4, "Creating and Mounting an NSS Volume," on page 574

E.3.1 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 fs-type ] [ fs-options ] filesys`

Option or Parameter	Description

`-t fs-type`	Specifies the type of file system to be built, such as nsspool or nssvol. For example:
	`-t nsspool`
	`-t nssvol`

`fs-options`	File-system-specific options to be passed to the real file system builder.
	When creating an NSS pool, use this fs-option:
	`-n poolname`
	Replace *poolname* with the actual name of the pool you want to create.
	When creating an NSS volume, use this fs-option:
	`-n volname`
	Replace *volname* with the actual name of the volume you want to create.

`filesys`	When using nsspool as the fs-type, replace *filesys* with the device name (*devname* such as `/dev/hda1` or `/dev/sdb2`). When using nssvol as the fs-type, replace *filesys* with the pool name (*poolname* such as `POOL2`).

Examples

`mkfs -t nsspool -n poolname devname`

`mkfs -t nssvol -n volname poolname`

E.3.2 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 *Mount Point* field blank.

8 Click *OK* to create the partition.

9 Continue with Section E.3.3, "Creating and Mounting an NSS Pool," on page 573.

E.3.3 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 E.3.2, "Creating a Partition," on page 572.

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 Select the pool from the *Pools* list, then press *F4* (NDS Update).

6 Continue with Section E.3.4, “Creating and Mounting an NSS Volume,” on page 574.

E.3.4 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 E.3.3, “Creating and Mounting an NSS Pool,” on page 573.

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 E.3.3, “Creating and Mounting an NSS Pool,” on page 573.

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
```

E.4 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 E.4.1, "Understanding Entries in the /etc/fstab Configuration File," on page 575
- Section E.4.2, "Adding NSS Pool and Volume Mount Information to /etc/fstab," on page 576

E.4.1 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.

Table E-1 *Options in the /etc/fstab File*

Column in /etc/fstab	Description	Examples
Device name	The location of the device or partition you want to mount.	`/dev/hda3` `/dev/sdb1`
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.	`/mnt/.pool2` `/media/nss/NSSV1`
File system type	The file system type of the device or partition.	`nsspool` `nssvol`

Column in /etc/fstab	Description	Examples

Mount options	Lists the comma-delimited mount options for the device or partition. Use a comma without spaces between options.
	
	auto or noauto: Use auto if you want the volume to mount automatically after a system reboot. Use noauto if you want the device to be mounted only when you explicitly issue the mount command. Do not use the auto command for encrypted NSS volumes; they require a password to be entered on the first mount after a system reboot.
	
	rw: Use rw to mount the device as read-write.
	
	name=<poolname | volname>: Specify the *poolname* or *volname* of the partition to be mounted.
	
	noatime: Use noatime for volumes when you want a file's access time (atime) to not be updated for reads.
	
	nodiratime: Use nodiratime for volumes when you want a directory's access time to not be updated for enumerations.
		`noauto,rw,name=POOL2`
		
		`noauto,rw,name=NSSV1`
		
		`name=POOL2,NSSV1,noatime, nodiratime`
		
		`name=NSSV1,noatime,nodira time`

Dump option number	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. Set this value to 0 for NSS pools and volumes.	`0`
File system check number	Designates if the file system should be checked with the Linux fsck utility. A value of 0 indicates that fsck should ignore this file system. Set this value to 0 for NSS pools and volumes.	`0`

E.4.2 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 E.4.1, "Understanding Entries in the /etc/fstab Configuration File," on page 575.

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 `/etc/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.

E.5 Expanding NSS Pools on Non-EVMS-Managed Devices

After using Linux utilities to move the pool to a larger device and to increase the partition size to the size of the device, use the NSS Update utility to increase the pool size up to the size of the new device. For instructions see, Section B.15, “NSSUPDATE (Linux),” on page 521.

Comparison of NSS on NetWare and NSS on Linux

F

This section compares features and capabilities of Novell® Storage Services™ on the NetWare® and Linux platforms for Novell Open Enterprise Server 2 servers.

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 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 terminal console (NSSMU, RIGHTS, NSSCON, ATTRIB, RAVSUi, RAVVIEW, `nbackup(1)` (http://www.novell.com/documentation/oes/smsadmin/data/nbackup.1.html))
	NSS commands in the server console	NSS commands in the NSS Console (NSSCON)
	Novell Remote Manager for NetWare	Novell Remote Manager for Linux (for Dynamic Storage Technology shadow volumes and for managing NCP Server connections to NSS volumes)
	Novell NetStorage for NetWare	Novell NetStorage for Linux
File system trustees, trustee rights, and inherited rights filter to control access to directories and files	Files and Folders plug-in to iManager	Files and Folders plug-in to iManager
	Novell Remote Manager for NetWare	
	Novell NetStorage for NetWare (via Web browser only, not WebDAV)	Novell NetStorage for Linux (via Web browser only, not WebDAV)
	Novell Client™	Novell Client
	RIGHTS utility for NetWare	RIGHTS utility for Linux
File system directory and file attributes to control functions available for directories and files	Files and Folders plug-in to iManager	Files and Folders plug-in to iManager
	Novell NetStorage for NetWare	Novell NetStorage for Linux
	Novell Client	Novell Client
	Novell Remote Manager for NetWare	

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 Linux
Directory quotas management (requires the Directory Quotas attribute for the volume)	Files and Folders plug-in to iManager Novell NetStorage for NetWare Novell Remote Manager for NetWare	Files and Folders plug-in to iManager Novell NetStorage for Linux
User space quota management (requires the User Space Quotas attribute for the volume)	Storage plug-in to Novell iManager	Storage plug-in to Novell iManager (OES SP1 and later)
Default mount location for NSS pools	Not applicable	`/opt/novell/nss/mnt/.pools/poolname`
Default mount location for NSS volumes	Server root	`/media/nss/volumename`
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 access protocols	NCP Native File Access Protocols (AFP, CIFS, and NFS)	NCP Novell AFP for Linux (beginning in OES 2 SP1) Novell CIFS for Linux (beginning in OES 2 SP1) CIFS/Samba using Novell Samba Linux NFS (version 3) Linux NFS and CIFS/Samba require users to be Linux-enabled with Linux User Management.
Interface	64-bit	64-bit
Character format	Unicode	Unicode
Maximum device size recognized (physical or logical)	2 TB	2 TB
Maximum software RAID device size (combined total for all member segments)	2 TB	2 TB
Minimum software RAID segment size	12 MB per segment	12 MB per segment
Maximum partition size	2 TB Valid Range: 10 MB to 2 TB	2 TB Valid Range: 10 MB to 2 TB
Maximum number of partitions per pool	No practical limit	No practical limit

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 Linux
Maximum pool size (created by 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.	Up to 8 TB, depending on the pool size and available space in the pool.
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. **NOTE:** NSS can support this. You are limited by the browser and application ability, of course.	Up to 8 trillion, regardless of how many name spaces are loaded. **NOTE:** NSS can support this. You are limited by the browser and application ability, of course.
Maximum number of files open concurrently	1 million	1 million
Maximum number of volumes per server	255 plus the `sys:` 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, thanks to journaling.	Requires only a few seconds, thanks to journaling.
Volume name space	Accommodates all name spaces (DOS, Macintosh, Long, and UNIX). Long is the default name space. In OES SP1 and later, the Lookup Name Space attribute allows you to set the default name space used when mounting volumes. Directory names and filenames are case insensitive.	Accommodates all name spaces (DOS, Macintosh, Long, and UNIX). Long is the default name space. The Lookup Name Space attribute allows you to set the default name space used when mounting volumes. Directory names and filenames are case insensitive with the Long name space.
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.

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 Linux
Error correction and data recovery time on system failure	Journaling file system logs changes. 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.	Journaling file system logs changes. 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.
Repair of corrupted pools and volume	Ongoing journaling of the file system; if the pool metadata structure is corrupted, use the NSS verify and rebuild functions.	Ongoing journaling of the file system; if the pool metadata structure is 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, Media Manager multipath I/O	Use a native Linux multiple path I/O solution.
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 5 (striping with parity)
Volume encryption	Yes	Yes, 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
Directory quotas	Yes	Yes
User space quotas (user space restrictions)	Yes	Yes, for OES SP1 and later
Salvage or purge deleted files, directories, or volumes	Yes	Yes

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 Linux
Transaction Tracking System™ (TTS™)	Yes	Not supported. If you need content tracking and trustee support, use NCP volumes on Linux reiser or ext3 file systems, then set the file system's journaling mode to the Journaling option.
Read ahead blocks	Yes	Yes
File save time	Provides 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.	Provides the group write options and timers. For information, see Section 31.3, "Configuring or Tuning Group I/O," on page 415.
File-level snapshot (make a temporary snapshot copy of an open file for backup)	Yes	Yes
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.	Yes, for OES 2 Linux and later. The stored-on pool must be on a separate partition.
Backup systems support	*OES 2 SP1: Storage Management Services Administration Guide* *OES 2 SP1: SBCON Administration Guide*	*OES 2 SP1: Storage Management Services Administration Guide* For an overview of backup resources on Linux, see Chapter 29, "Managing Backup and Restore for Data and Trustee Information," on page 399.
Distributed File Services for moving and splitting NSS volumes	Yes	Yes, for OES 2 Linux and later
Novell Archive and Version Services	Yes	Yes, for OES 2 Linux and later
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 POSIX file system options instead.

Feature Description	NSS for OES 2 SP1 NetWare and NetWare 6.5 SP8	NSS for OES 2 SP1 Linux
CD and DVD image files	Activate as read-only NSS volumes.	No; use Linux POSIX file system options instead.
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 file system options instead.
Operating system version detection	Default process	Default process
Cache balancing for NSS cache buffers	Yes; for information, see "Tuning NSS Performance on NetWare" on page 405.	Yes; for information, see "Tuning NSS Performance on Linux" on page 413.
Tuning I/O write behavior	Set I/O tuning parameters for NSS on NetWare. For information, see "Setting the File and Buffer Flush Timers" on page 410.	Set group I/O write parameters for NSS on Linux. For information, see "Configuring or Tuning Group I/O" on page 415.
Dynamic Storage Technology (DST) For information about NSS volume attributes and features that are supported in DST shadow volumes, see "Guidelines for Using NSS Volumes in DST Shadow Volumes" in the *OES 2 SP1: Dynamic Storage Technology Administration Guide*.	DST supports using NetWare iSCSI target devices to store NSS volumes in the shadow volume pair. The target devices are attached to the OES 2 Linux server by using the Linux iSCSI initiator software. For information, see "iSCSI Block Storage Devices" in the *OES 2 SP1: Dynamic Storage Technology Administration Guide*.	DST supports NSS volumes on OES 2 Linux servers as the primary or secondary volume in the shadow volume. For information, see the *OES 2 SP1: Dynamic Storage Technology Administration Guide*.

Comparison of NSS on NetWare and the NetWare Traditional File System

G

This section compares features and capabilities of the Novell® Storage Services™ (NSS) File System on NetWare® and the NetWare Traditional File System.

- Section G.1, "Features of NSS and Traditional File Systems on NetWare," on page 585
- Section G.2, "SET Parameters Shared by NSS and Traditional File Systems," on page 589
- Section G.3, "SET Parameters for the NSS File System," on page 590
- Section G.4, "SET Parameters for the Traditional File System," on page 591
- Section G.5, "Comparison of Directory and File Attributes," on page 593

G.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™ *OES 2: NetStorage for NetWare Administration Guide* (via Web browser only, not WebDAV) Novell Remote Manager for NetWare Rights utility for NetWare	Novell Client *OES 2: NetStorage for NetWare Administration Guide* (via Web browser only, not WebDAV)
Interface	64-bit	32-bit
File access protocols	NCP™ Native File Access Protocols (AFP, CIFS, AFP)	NCP
Character format	Unicode	ASCII double-byte

Feature Description	NSS File System	NetWare Traditional File System
Maximum device size recognized (physical or logical)	2 TB	2 TB
Minimum device size (physical or logical) (needed to create a pool or Traditional volume)	10 MB	10 MB
Maximum software RAID device size (combined total of all member segments)	2 TB	Not applicable
Minimum software RAID segment size	12 MB per segment	Not applicable
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.	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 (where each is 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 (The practical limit is the number of files that can be handled by the file browser or application.)	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 `sys:` volume. You can mount NSS volumes beyond 256, but they are not visible or accessible through the normal Netware APIs.	255 plus the `sys:` volume.

Feature Description	NSS File System	NetWare Traditional File System
Maximum size of a volume	Up to 8 TB, depending on the pool size and available space in the pool. If a pool contains multiple volumes, the cumulative administrative maximum sizes of all volumes can exceed the pool size by 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.	Up to 1 TB. 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.
Time to mount a volume	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).	Up to several minutes, depending on volume size. 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. 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.
Volume name space	Accommodates all name spaces and requires less memory to mount than Traditional volumes.	Each name space you add to a Traditional volume requires additional server memory. If you add name space support to a volume and do not have enough memory, that volume cannot be mounted.
Minimum server memory required to activate a volume	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.	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.
File access time	Same for each file, regardless of its location on the volume.	Depends on the file's location on the volume.

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 The smallest file size that is eligible for compression is 8192 + 1 bytes. NSS does not have suballocation so each file takes up a multiple of 4K instead of 512 bytes.	Yes The smallest file size that is eligible for compression is 512 + 1 bytes.
Volume encryption	Yes	No
Error correction and data recovery time on system failure	Journaling file system logs changes. 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. If the volume is corrupted, repair can take up to several hours, depending on the volume size.	Must verify the entire volume contents with vrepair.nlm, which can take several hours, depending on the volume size.
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)
Data shredding	Yes, up to 7 times	No
Transaction Tracking System™ (TTS™)	Yes Requires exclusive use per server. To use TTS, first disable TTS for all Traditional volumes, and then enable TTS for your NSS volumes.	Yes Requires exclusive use per server. Use for Traditional or for NSS, but not for both on the same server. TTS is enabled by default.
Directory quotas and User space quotas	Yes	Yes, enabled by default

Feature Description	NSS File System	NetWare Traditional File System
Salvage subsystem	Yes, for deleted directories and files	Yes, for deleted files
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
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 list of modified files. 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
filename locks	Unnecessary; superseded by technology.	Yes
Data migration	Yes	Yes

G.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.

Common File System SET Parameters	Default Value
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

G.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.

Common File System SET Parameters for NSS File System	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
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

Common File System SET Parameters for NSS File System	Default Value
NSS Buffer Flush Timer	1
In OES 2 Linux, this timer is not used. It is replaced by the Group Write timers for Journal, Metadata, and User Data.	
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

G.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.

SET Parameters for the Traditional File System	Default Value
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

SET Parameters for the Traditional File System	Default Value
Minimum Directory Cache Buffers	500
Maximum Number Of Internal Directory Handles	100
Maximum Number Of Directory Handles	20
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

G.5 Comparison of Directory and File Attributes

Attribute Code	Description	Applies to Files	Applies to Directories	NSS	Traditional
A	Archive Needed identifies files that have been modified since the last backup. This attribute is assigned automatically.	Yes	No	Yes	Yes
Ci	Copy Inhibit prevents Macintosh users from copying a file. This attribute overrides Read and File Scan trustee rights. This attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare.	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

Attribute Code	Description	Applies to Files	Applies to Directories	NSS	Traditional
Ri	Rename Inhibit prevents the directory or filename from being modified.	Yes	Yes	Yes	Yes
Ro	Read Only prevents a file from being modified.	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
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

Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems

H

This section compares features and capabilities of the Novell® Storage Services™ file system on Novell Open Enterprise Server 2 Linux to those of NCP volumes on Linux POSIX file systems such as Ext3 and Reiser. For information, see "Managing NCP Volumes" in the *OES 2 SP1: NCP Server for Linux Administration Guide*.

For information to help you choose from among the numerous Linux file system offerings, see the following:

- *File System Primer* (http://wiki.novell.com/index.php/File_System_Primer)
- *SUSE Linux Enterprise Server 10 File Systems Features Overview* (http://www.novell.com/linux/filesystems/features.html)

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
Management interfaces	Novell iManager Storage plug-in NSSMU for Linux Novell iManager with various storage-related plug-ins. For information, see Section 9.1, "Novell iManager and Storage-Related Plug-Ins," on page 97. EVMS (`evmsgui`) NSS utilities in a terminal console NSS commands in the NSS Console (NSSCON) Novell Remote Manager for Linux (browse only)	YaST > Partitioner for managing devices Novell Remote Manager for Linux (Managing Shares) NCP commands in the NCP Console (NCPCON) YaST EVMS (`evmsgui`) for managing Linux devices and volumes. Various Linux commands and utilities in a terminal console.
File system trustees and trustee rights to control access to directories and files	Yes, works with or without concurrent running of NCP Server.	Yes, requires NCP Server to enforce the rights and access on the extended attributes.

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
File access protocols	NCP Novell AFP for Linux (beginning in OES 2 SP1) Novell CIFS for Linux (beginning in OES 2 SP1) CIFS/Samba using Novell Samba Linux NFS (version 3) Linux NFS and CIFS/Samba require users to be Linux-enabled with Linux User Management.	NCP CIFS/Samba using Novell Samba
File system directory and file attributes to control functions available for directories and files	Files and Folders plug-in to iManager Novell NetStorage Novell Client™ Novell Remote Manager for Linux. See "Displaying Key NSS Directory and File Attributes as Linux POSIX Permissions" in the *OES 2 SP1: File Systems Management Guide*.	Not applicable. Use POSIX file and directory attributes.
Directory quotas	Yes, requires the Directory Quotas attribute to be enabled. Novell NetStorage Novell Client	No
User space quotas (user space restrictions)	Yes, for OES Linux SP1 and later	Yes, if the Linux file system being used under the NCP share supports user quotas and the Linux file system resides on a local, iSCSI, or Fibre Channel drive. All users of the NCP volume must be LUM enabled. Manage the user quotas using the Linux file system tools.
Default mount location for NSS pools	`/opt/novell/nss/mnt/.pools/poolname`	Not applicable.

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
Volume name space	Long is the default name space, which is case insensitive. You can specify the UNIX name spaces on mounting the NSS volume to make its directory names and filenames case sensitive. Using UNIX name space slows performance compared to using Long. For example: `mount ns=namespace` Valid name space values are dos, mac, long, or unix.	UNIX; no support for case insensitive names.
Salvage for deleted volumes, directories, and files	Yes	No
Volume encryption	Yes, for OES Linux SP1 and later	Yes, for Reiser
File compression	Yes	No
Data shredding (secure deletion)	Yes, up to 7 times	No
Online resizing of volumes and pools	Yes	Yes, depending on the file system
Multiple I/O paths to storage media For information, see "Managing Multipath I/O for Devices" in the *SLES 10 SP2: Storage Administration Guide*.	No; NSS-specific multipath I/O tools as are not available on Linux. Use the Linux Device Mapper driver support for mutlipath I/O on devices where you plan to create NSS file systems.	Use the Linux Device Mapper driver support for mutlipath I/O on devices. (NCP is not required to make this work.)
Software RAID support	RAID 0, 1, and 5.	RAID 0, 1, 4, 5 and 6. RAID 10 can be created using the Linux `mdadm(8)` command as a complex RAID using the RAID10 option, or as a nested RAID. For information, see the *SLES 10 SP2: Storage Administration Guide*.
Pool snapshot (retain point-in-time version of a pool using block-level copy on write)	Yes, using iManager or NSSMU.	Depends on the file system. EVMS supports device snapshots for the devices it manages. (NCP is not required to make this work.)

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
Hard links	Yes; enhanced hard links support is available in OES 2 and later. For information, see Chapter 27, “Managing Hard Links,” on page 375.	Yes
Backup support	Yes, using Novell Storage Management Services™ for Linux. For information, See Chapter 29, “Managing Backup and Restore for Data and Trustee Information,” on page 399.	No. Use third-party solutions.
Data migration from NSS volumes on NetWare	Yes	Yes
Novell Archive and Version Services	Yes. For information, see the *OES 2 SP1: Novell Archive and Version Services 2.1 for Linux Administration Guide*.	No
Novell Distributed File Services For information, see the *OES 2: Novell Distributed File Services Administration Guide*.	Yes, for OES 2 Linux and later. NSS volumes on OES 2 Linux can contain junctions or be a junction target. NSS volumes on OES 1 Linux can be a junction target, but junctions are not supported in the volume.	Only as targets of junctions in OES 2 and later. DFS does not support junctions on NCP volumes.
Dynamic Storage Technology For information, see the *OES 2 SP1: Dynamic Storage Technology Administration Guide*.	Yes	Not available in the initial OES 2 release.
Novell Cluster Services for Linux For information, see the *OES 2 SP1: Novell Cluster Services 1.8.5 for Linux Administration Guide*.	Yes For information, see “Creating NSS Shared Disk Partitions and Pools”.	Yes; cluster the Linux POSIX file system, then create the NCP volume on it. For information, see “Creating Linux POSIX Volumes on Shared Disks”.
Novell Transaction Tracking System™ (TTS™)	No	Use the Journal mode for Linux POSIX file systems that support journaling.
Operating system version detection	Default process	Default process
Device maintenance support	Activate and deactivate devices by pool.	Activate and deactivate devices using Linux tools.

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
Cache balancing for NSS cache buffers	You can specify a minimum cache buffer size. For information, see "Tuning NSS Performance on Linux" on page 413.	Integrated with the Linux file system cache.
CD and DVD device recognition	No; not managed by NSS. Use Linux services to mount CDs and DVDs as Linux volumes.	Yes, default
Ability to access DOS partitions as on a NetWare server	No; not managed by NSS. Use Linux services instead.	Yes, using Linux services.
Default mount location for NSS volumes	`/media/nss/volumename`	Not applicable.
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	For a 32-bit OS: ◆ 4 TB for Ext2 or Ext3 ◆ 16 TB (minus 1 Byte) for Reiser For a 64-bit OS: ◆ 2 to 32 TB for Ext2 or Ext3, depending on the block size ◆ 16 TB (minus 1 Byte) for Reiser
Maximum software RAID size (combined total of all member segments)	2 TB	See Maximum device size recognized.
Minimum software RAID segment size	12 MB per segment	Depends on the file system.
Maximum partition size	2 TB Valid Range: 10 MB to 2 TB	Up to 16 TB, depending on the file system and block size as noted above.
Maximum number of partitions (logical or physical devices) per pool	No practical limit	Not applicable.
Maximum pool size	8 TB (using 4 or more partitions of up to 2 TB each)	Up to the partition size, depending on the file system.
Minimum pool size	10 MB	Not applicable.

Feature Description	NSS on OES 2 SP1 Linux	NCP Volumes on Linux POSIX File Systems
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 Section 19.2, “Guidelines for NSS Volumes,” on page 258.	Up to the partition size, depending on the file system.
Maximum file size	Up to 8 TB, depending on the volume size and available space in the volume.	2 GB to 2 TB for Ext2 or Ext3, depending on the block size. Up to 8 TB for Reiser.
Maximum number of files per volume (In practice, how many files be managed is limited only by the file browser’s and application’s ability to list and access the files.)	Up to 8 trillion (10E12), regardless of how many name spaces are loaded. Up to 4 billion (10E9) files in a single directory.	Up to 8 trillion (10E12), regardless of how many name spaces are loaded.
Maximum number of files open concurrently	1 million (10E6)	Millions (10E6), depending on the file system
Maximum number of volumes per server	Unlimited NSS data volumes, but only 255 can be mounted at a time	Unlimited
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.
Time to repair corrupted volume	Up to several hours, depending on the volume size.	Up to several hours, depending on the volume size

NSS Nomenclature

This section describes the nomenclature used for key Novell® Storage Services™ media objects in Novell Open Enterprise Server 2. 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® or OES where the media object first appeared.

NOTE: All ZLSS (NSS Journaled Storage System) file blocks are 4 KB in size.

Media Object	Definition	Version Where First Used
Area Seed	To improve performance for OES 2 Linux, metadata blocks use an area seed logic to make sure that related metadata blocks are physically stored near each other. For information about configuring area size, see Section 31.3.1, "Viewing the Metadata Area Size," on page 416.	OES 2 Linux (not available on NetWare)
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.	NetWare 5.0
Checkpoint	Four blocks (one in each Superblock) that track where to start playing the journal if the server crashes. The checkpoint contains the metadata of the Journal.	NetWare 5.0
Directory B-Tree	Used to implement an NSS volume's Directory Quota feature.	NetWare 5.0 Support Packs
Entry ID (EID)	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. 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. 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.	NetWare 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.	NetWare 6.5
Globally Unique ID (GUID)	A globally unique identifier within eDirectory. The scope of this uniqueness is within one tree, although no actual checking is done to ensure this. GUIDs are 128 bits and are unique for each object. GUIDs allow an object to be referenced no matter which server you are accessing.	NetWare 5.0

Media Object	Definition	Version Where First Used
Journal	The file used to quickly make the file system consistent after a server crash. The journal is sometimes referred to as a zlog.	NetWare 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.	NetWare 6.0
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.	NetWare 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.	NetWare 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 Superblock.	NetWare 6.0 Support Packs
Name Tree B-Tree	The B-Tree that tracks the directory structure of a volume.	NetWare 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.	NetWare 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.	NetWare 5.0
Purge Tree B-Tree	Used to store information about all salvageable files. This tree is used when the file system needs to autopurge files to create free blocks.	NetWare 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 Superblock. All other blocks are allocated from a file on the internal volume.	NetWare 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 Checkpoints, Superblock Header, Snapshot, and Multiple Server Access Prevention.	NetWare 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.	NetWare 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.	NetWare 5.0 Support Packs

Media Object	Definition	Version Where First Used
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.	NetWare 5.0
ZID	A numeric ID within the NSS file system, used to reference an object with the object store (also known as the "beast tree"). ZIDs are 64 bits and are unique for each volume. The number of ZIDs is limited to 4 billion (4xE9) because of 32-bit limitations from the client. ZIDs are not reused, so if your processes create and delete lots of files in a volume, the numbering of ZIDs might approach this limit. and cause files not to be allowed to be created. If this occurs, you should rebuild the pool to renumber the files, referred to as "reZIDing". For more information, see Section 17.3, "ReZIDing Volumes in an NSS Pool," on page 223.	NetWare 6.0

Documentation Updates

J

This section contains information about documentation content changes made to the *OES 2: Novell Storage Services File System Administration Guide* since the initial release of Novell® Open Enterprise Server 2. 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 title page, to determine the release date of this guide. For the most recent version of the *OES 2: Novell Storage Services File System Administration Guide*, see the OES 2 documentation Web site (http://www.novell.com/documentation/oes2/storage.html#nss).

In this section, content changes appear in reverse chronological order, according to the publication date. Within a dated entry, changes are grouped by chapter and sequenced alphabetically. 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 J.1, “December 2008 (OES 2 SP1, NetWare 6.5 SP8),” on page 605
- Section J.2, “June 27, 2008,” on page 614
- Section J.3, “May 14, 2008,” on page 616
- Section J.4, “February 22, 2008,” on page 617
- Section J.5, “January 14, 2008,” on page 618
- Section J.6, “December 14, 2007,” on page 620
- Section J.7, “November 14, 2007,” on page 621

J.1 December 2008 (OES 2 SP1, NetWare 6.5 SP8)

Updates were made to the following sections. The changes are explained below.

- Section J.1.1, “Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems,” on page 606
- Section J.1.2, “Comparison of NSS on NetWare and NSS on Linux,” on page 606
- Section J.1.3, “Cross-Platform Issues for NSS,” on page 607
- Section J.1.4, “Installing and Configuring Novell Storage Services,” on page 607
- Section J.1.5, “Management Tools for NSS,” on page 607
- Section J.1.6, “Managing Backup and Restore for Data and Trustee Information,” on page 607
- Section J.1.7, “Managing Files and Folders on NSS Volumes,” on page 608
- Section J.1.8, “Managing NSS Pool Snapshots,” on page 608
- Section J.1.9, “Managing NSS Pools,” on page 608
- Section J.1.10, “Managing NSS Software RAID Devices,” on page 609
- Section J.1.11, “Managing NSS Volumes,” on page 610
- Section J.1.12, “Managing Space Quotas for Volumes, Directories, and Users,” on page 610

- Section J.1.13, “Migrating NSS Devices from NetWare to OES 2 Linux,” on page 610
- Section J.1.14, “NSS Commands,” on page 611
- Section J.1.15, “NSS Utilities,” on page 611
- Section J.1.16, “Planning for NSS Storage Solutions,” on page 612
- Section J.1.17, “Salvaging and Purging Deleted Volumes, Directories, and Files,” on page 612
- Section J.1.18, “Securing Access to NSS Volumes, Directories, and Files,” on page 612
- Section J.1.19, “Troubleshooting the NSS File System,” on page 613
- Section J.1.20, “Tuning NSS Performance on NetWare,” on page 613
- Section J.1.21, “Using EVMS to Manage Devices with NSS Volumes (Linux),” on page 613
- Section J.1.22, “Using NSS in a Virtualization Environment,” on page 614
- Section J.1.23, “What’s New,” on page 614

J.1.1 Comparison of NSS on Linux and NCP Volumes on Linux POSIX File Systems

Location	Change
“File access protocols” on page 596	This entry is new.
“User space quotas (user space restrictions)” on page 596	For NCP volumes on Linux: Yes, if the Linux file system being used under the NCP share supports user quotas and the Linux file system resides on a local, iSCSI, or Fibre Channel drive. All users of the NCP volume must be LUM enabled. Manage the user quotas using the Linux file system tools.
“Volume name space” on page 597	For NSS on OES 2 SP1 Linux, the default name space used when mounting NSS volumes is changed from UNIX to Long. This improves performance over using UNIX.
“Software RAID support” on page 597	NSS management tools on Linux do not support creating nested RAID 10 and RAID 15.

J.1.2 Comparison of NSS on NetWare and NSS on Linux

Location	Change
“File access protocols” on page 580	Novell AFP for Linux (beginning in OES 2 SP1) Novell CIFS for Linux (beginning in OES 2 SP1)
“Volume name space” on page 581	For NSS on OES 2 SP1 Linux, the default name space used when mounting NSS volumes is changed from UNIX to Long. This improves performance as compared to using UNIX.
“Software RAID support” on page 582	NSS management tools on Linux do not support creating nested RAID 10 and RAID 15.

J.1.3 Cross-Platform Issues for NSS

Location	Change
Section 7.5.1, “Storage-Related Plug-Ins for Novell iManager 2.7,” on page 90	The support matrix was updated to include Novell CIFS for Linux and Novell AFP for Linux in OES 2 SP1.
Section 7.5.5, “Management Capabilities for Software RAIDs,” on page 92	NSSMU for Linux does not support creating nested RAID 10 and 15 devices.

J.1.4 Installing and Configuring Novell Storage Services

Location	Change
“EVMS Requirements” on page 46	**IMPORTANT:** If you move devices that contain NSS pools cross-platform (such as reassigning SAN-based devices from a NetWare server to an OES 2 Linux server), NSS recognizes the pools and automatically uses EVMS to manage the devices.
Section 3.6, “Updating NSS on OES 2 Linux,” on page 52	This section is new.

J.1.5 Management Tools for NSS

Location	Change
Section 9.1.1, “Understanding Storage-Related Plug-Ins,” on page 98	The DFS, CIFS, AFP, and NFS features have been separated from the NSS `.npm` file, and now have `.npm` files of their own.

J.1.6 Managing Backup and Restore for Data and Trustee Information

Location	Change
Section 29.2, “Using the Event File List to Refine the Backup,” on page 400	This section is new.
Section 29.5.2, “NetWare,” on page 402	This section is new.

J.1.7 Managing Files and Folders on NSS Volumes

Location	Change
Section 28.2.1, “Prerequisites,” on page 388	A folder must be empty before you can delete it.
Section 28.5, “Viewing or Modifying File or Folder Properties,” on page 391	The Copy Inhibit attribute works only for clients using Macintosh operating systems to access NSS volumes on NetWare. If the directory quota exceeds the volume quota, the volume quota is enforced.

J.1.8 Managing NSS Pool Snapshots

Location	Change
Section 18.2.1, “Differences Between Snapshots on Linux and NetWare,” on page 232	This section is new.
Section 18.2.9, “Guidelines for Viewing a List of Snapshots that are Stored in a Pool (NetWare),” on page 237	This section is new.

J.1.9 Managing NSS Pools

Location	Change
Section 16.2, “Creating a Pool,” on page 198	Added information about using Novell CIFS for Linux and Novell AFP for Linux when cluster-enabling pools on OES 2 SP1 Linux and later.

J.1.10 Managing NSS Software RAID Devices

Location	Change
Section 13.5, “Creating Software RAID Devices with iManager,” on page 174 and Section 13.6, “Creating Software RAID Devices with NSSMU,” on page 176	**IMPORTANT:** Unallocated partitions (that is, partitions that are not mirrored and do not contain pools or other file systems), are deleted in order to present the unused space as free space for use by the RAID. No data loss occurs by this action.

Location	Change
Section 13.9, “Creating a Software RAID 10 with NSSMU (NetWare),” on page 178	NSS management tools on Linux do not support creating nested software RAID 10 device.
Section 13.10, “Creating a Software RAID 15 with NSSMU (NetWare),” on page 179	NSS management tools on Linux do not support creating nested software RAID 15 device.
Section 13.12, “Increasing the Size of a Software RAID Device,” on page 180	**IMPORTANT:** If the software RAID device is shared in a cluster, connect to the node where the RAID is currently active to manage the RAID and increase the size of the RAID. A device must be the same size or larger than the segment size being used in the RAID. You might need to add or initialize a new device, then try again. While restriping, the new device is considered a failed device until it is completely resynchronized. While expanding a RAID 5 device, if one of the drives goes down (either one of the existing segments or the newly added segment), the pool deactivates. If you remove any device from a RAID 5 other than the one that was just added for restripe, it considers that as a two-disk error, and deactivates the RAID and the pool.
Section 13.13, “Restriping a Software RAID 0 or 5 Device,” on page 182	When expanding a RAID 5 on Linux, if the newly added drive goes down during the restripe, the restriping continues without the new partition and puts the RAID in a degraded state with one partition missing. If the same partition comes back online, it finishes the restripe. If the partition has completely failed, after the degraded restriping is complete, you can add a new replacement partition, and the RAID restripes to fix it. When expanding a RAID 5 on NetWare, if the newly added drive fails during the restripe, this is considered a fault and the device and its pools are automatically deactivated. If the same partition comes back online, it finishes the restripe. If the partition is failed, this is not handled. You must recover from backup. (Handling this type of failure with a degraded restripe recovery (same as for RAID 5 on Linux) is planned for a future release.)
Section 13.14, “Replacing a Failed Segment in a Software RAID,” on page 183	**IMPORTANT:** If a second segment fails before the restriping is completed for the first drive replacement, this is considered a two-drive failure. You must recover data from a backup copy.

J.1.11 Managing NSS Volumes

Location	Change
Section 19.2.1, “Guidelines for Sizing Volumes,” on page 258	Added information about how NSS reports space usage to management tools.

Location	Change
“Lookup Namespace” on page 257	For OES 2 SP1 Linux and later, Long is the default name space instead of UNIX. For Linux, NCP tools can be used only when Long or UNIX is set as the primary name space. The UNIX name space supports some special characters that are not allowed in the Long name space, such as characters 0x01 through 0x07 and 0x10 through 0x1f. If you need to use these special characters in filenames, choose UNIX as the primary name space.

J.1.12 Managing Space Quotas for Volumes, Directories, and Users

Location	Change
Section 25.1, “Understanding Space Quotas,” on page 345	For OES 2 SP1, it is no longer necessary to restart eDirectory after enabling or disabling the Directory Quotas attribute.
Section 25.4.6, “Configuring User Space Quotas on Volumes After Upgrading or Migrating from OES 1 Linux (Linux),” on page 358	You cannot use the Linux `chown` command to change the creator field for the NSS file system. It changes the root user’s view of who is reported as the owner user in the Linux path, but the change has no effect on the NSS metadata.

J.1.13 Migrating NSS Devices from NetWare to OES 2 Linux

Location	Change
Section 11.1.1, “Media Format,” on page 141	This section was modified for clarity.

J.1.14 NSS Commands

Location	Change
Section A.5.5, “UnplugAlways Command for the Read Queue,” on page 457	The `/(no)unplug` option is new for OES 2 SP1.
Section A.25, “noatime and atime Commands,” on page 478	`/atime` and `/noatime` options are new for NetWare 6.5 SP8 (same as OES 2 SP1 NetWare). Added clarification about the volume ID value in the cluster examples.

Location	Change
Section A.27, “Opportunistic Locking (Oplock) Commands,” on page 481	Added a definition of opportunistic locking.
Section A.29.1, “Pool Status,” on page 482	Added a definition for `/ZLSSPoolIOErrors`.
Section A.32, “POSIX Permission Mask Command (Linux),” on page 485	The `/PosixPermissionMask` option is new for OES 2 SP1 Linux.
Section A.40, “Transaction Tracking System (TTS) Commands (NetWare),” on page 492	To disable TTS for volume VOL1, enter the following at the server console: `nss /noTransaction=VOL1`
Section A.42.2, “Volumes Command,” on page 495	Added an example of the output from the `volumes` command.

J.1.15 NSS Utilities

Location	Change
Section B.16, “RAVSUI (Linux),” on page 522	This section was updated for clarity.
Section B.17, “RAVVIEW (Linux),” on page 525	This section was updated for clarity.
Section B.24, “VOLUMES (Linux, NCP Utility),” on page 545	This section was updated for clarity.

J.1.16 Planning for NSS Storage Solutions

Location	Change
Section 5.1.7, “Storage Features,” on page 66	This section is new.
Section 5.2, “Compatibility and Interoperability Issues for NSS,” on page 66	This section is new.

Location	Change
Section 5.4.5, “UTF-8 Naming Considerations in Mixed-Language Environments (NetWare),” on page 70	This section is new.
“Linux-Enabled eDirectory Users” on page 73	In OES 2 SP1 and later, the modifier is reported as the user instead of as the `root` user or nobody user.
“File Access for Users” on page 74	This section is new.
Section 5.7, “Antivirus Support for NSS,” on page 77	This section is new.
Section 5.9, “NSS Support for Memory Mapped Files,” on page 78	This section is new.

J.1.17 Salvaging and Purging Deleted Volumes, Directories, and Files

This chapter was revised for clarity.

J.1.18 Securing Access to NSS Volumes, Directories, and Files

Location	Change
“eDirectory Users” on page 72	Moved to the planning section.
“Linux-Enabled eDirectory Users” on page 73	Users who access NSS via Samba, NFS, or third-party AFP protocols must be Linux-enabled in order for user quotas to work for them. Moved to the planning section.
Section 5.5.5, “How NSS Uses Novell Linux User Management,” on page 74	Moved to the planning section.
Section 5.6, “File Access for Users,” on page 74	In OES 2 SP1 Linux and later, users can also access NSS volumes using Novell CIFS, Novell AFP, and Domain Services for Windows. Moved to the planning section.

J.1.19 Troubleshooting the NSS File System

Location	Change
Section 33.9, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux),” on page 434	This section is new.
Section 33.13, “Server Hangs When Using an NSS Volume as a Netatalk Share,” on page 436	This section is new.

J.1.20 Tuning NSS Performance on NetWare

Location	Change
Section 30.8, “Tuning NSS Write Performance on NetWare,” on page 412	This section is new.

J.1.21 Using EVMS to Manage Devices with NSS Volumes (Linux)

Location	Change
“Using Linux POSIX File Systems and NSS Pools and Volumes on the Same Device” on page 558	This section is new.

J.1.22 Using NSS in a Virtualization Environment

Location	Change
“Using RAIDs” on page 80	Updated with support status for using software RAIDs in a Xen-based virtualization environment.

J.1.23 What's New

Location	Change
Section 2.1, "What's New for NSS (OES 2 SP1 and NetWare 6.5 SP8)," on page 37	This section is new.

J.2 June 27, 2008

Updates were made to the following sections. The changes are explained below.

- Section J.2.1, "Managing Devices," on page 614
- Section J.2.2, "Managing Multipath I/O to Devices (NetWare)," on page 614
- Section J.2.3, "Managing Space Quotas for Volumes, Directories, and Users," on page 615
- Section J.2.4, "NSS Commands," on page 615
- Section J.2.5, "Salvaging and Purging Deleted Volumes, Directories, and Files," on page 615
- Section J.2.6, "Using NSS in a Virtualization Guest Server Environment," on page 615

J.2.1 Managing Devices

Location	Change
"Enabling Write-Through Cache Management on SCSI Devices and RAID Controllers (Linux)" on page 138	This section was moved here from the Troubleshooting section. It was revised for clarity.

J.2.2 Managing Multipath I/O to Devices (NetWare)

Location	Change
"Enabling Multipath" on page 190	This section is new.

J.2.3 Managing Space Quotas for Volumes, Directories, and Users

Location	Change
"Configuring User Space Quotas on Volumes After Upgrading or Migrating from OES 1 Linux (Linux)" on page 358	Updated to reflect the OES 1 Linux to OES 2 Linux migration scenario where users were not Linux-enabled on OES 1.

J.2.4 NSS Commands

The command section titled "POSIX Rights for Files Command (Linux)" was removed. The /InitialUNIXRights option was not released in the FCS version.

J.2.5 Salvaging and Purging Deleted Volumes, Directories, and Files

Location	Change
"Understanding the NSS Salvage System" on page 361	After the deleted data enters a Purge state by manually starting a purge or by autopurging, deleted files can longer be salvaged (do not return to a Salvageable state).
Section 26.4, "Viewing, Salvaging, or Purging Deleted NSS Volumes in a Pool," on page 368	**Pause Purge/Restart Purge:** *Pause Purge* disables autopurging for the selected volumes so that purging does not begin automatically for a deleted volume when its Purge Delay time elapses. During the purge delay or while autopurging is disabled, the deleted volume is salvageable. *Restart Purge* enables autopurging for the selected volumes. The deleted volume is purged when the purge delay time elapses. This option does not make it possible to salvage a deleted volume that has already entered a Purge state.

J.2.6 Using NSS in a Virtualization Guest Server Environment

Location	Change
"Running NSS on the Host Server Is Not Supported" on page 79	Updated with information about using clustered NSS pools and volumes in a virtualization environment.

J.3 May 14, 2008

Updates were made to the following sections. The changes are explained below.

- Section J.3.1, "Installing and Configuring Novell Storage Services," on page 616

- Section J.3.2, “Managing Devices,” on page 616
- Section J.3.3, “NSS Commands,” on page 616
- Section J.3.4, “NSS Utilities,” on page 617

J.3.1 Installing and Configuring Novell Storage Services

Location	Change
“EVMS Requirements” on page 46	NSS requires EVMS version 2.5.5-24.54.5 or later. This mandatory update is available in the SUSE Linux Enterprise Server Support Pack 1 patch channel as of March 21, 2008. It is also available in the SLES 10 SP2 release. **WARNING:** The 54.5 or later version is required to prevent and avoid data corruption from occurring when you update EVMS while NSS volumes are mounted.

J.3.2 Managing Devices

Location	Change
“Device Size Limit” on page 127	Changed “disk partitioner” to “the storage vendor’s or third-party disk carving tool”.

J.3.3 NSS Commands

Location	Change
Section A.26, “noatime and nodiratime Support for Linux open, mount, nfsmount, and /etc/fstab (Linux),” on page 479	The noatime option disables the updating of access time for both files and directories so that reading a file does not update their access times (atime). Typically, you need to use only the `noatime` option so that atime is not updated for the accessed file and its directory when the file is accessed.
Section A.25.2, “Using noatime in a Cluster Load Script,” on page 478	This section is new.

J.3.4 NSS Utilities

Location	Change
Section B.13, “NSSCON (Linux),” on page 519	This section was reorganized for clarity.

Location	Change
Section B.13.3, “Using nsscon in a Script,” on page 520	This section is new.

J.4 February 22, 2008

Updates were made to the following sections. The changes are explained below.

- Section J.4.1, “Installing and Configuring Novell Storage Services,” on page 617
- Section J.4.2, “Management Tools for NSS,” on page 618
- Section J.4.3, “NSS Utilities,” on page 618
- Section J.4.4, “Troubleshooting the NSS File System,” on page 618

J.4.1 Installing and Configuring Novell Storage Services

Location	Change
“EVMS Requirements” on page 46	**NOTE:** The following data corruption problem that arose after the OES 2 Linux release can be easily resolved by upgrading to OES 2 SP1 Linux. **WARNING:** Data corruption can occur if NSS pools and volumes are mounted when you install EVMS updates. For details and workarounds, see the following resources in the *OES 2 Readme (http://www.novell.com/documentation/oes2/pdfdoc/readme-20/readme-20.pdf)*: ◆ Possible Data Corruption Can Occur If NSS Pools Are Mounted During an EVMS Update ◆ EVMS Updates Are Mandatory to Prevent Data Corruption for NSS Pools and Volumes

J.4.2 Management Tools for NSS

Location	Change
"EVMS (Linux)" on page 100	**NOTE:** The following data corruption problem that arose after the OES 2 Linux release can be easily resolved by upgrading to OES 2 SP1 Linux. **WARNING:** Data corruption can occur if NSS pools and volumes are mounted when you install EVMS updates. For details and workarounds, see the following resources in the *OES 2 Readme (http://www.novell.com/documentation/oes2/pdfdoc/readme-20/readme-20.pdf)*: ◆ Possible Data Corruption Can Occur If NSS Pools Are Mounted During an EVMS Update ◆ EVMS Updates Are Mandatory to Prevent Data Corruption for NSS Pools and Volumes
"Understanding Storage-Related Plug-Ins" on page 98	This section was edited for clarity.

J.4.3 NSS Utilities

Location	Change
"NSSCON (Linux)" on page 519	NSSCON can be used for all NSS commands listed in Appendix A, "NSS Commands," on page 451, except those commands that are marked explicitly for NetWare.

J.4.4 Troubleshooting the NSS File System

Location	Change
"eDirectory Error 672 When Creating an NSS Pool on Linux" on page 433	Run `nssAdminInstall` at a Linux terminal console as the `root` user.

J.5 January 14, 2008

Updates were made to the following sections. The changes are explained below.

- Section J.5.1, "Managing Devices," on page 619
- Section J.5.2, "Managing Files and Folders," on page 619
- Section J.5.3, "Management Tools for NSS," on page 619
- Section J.5.4, "Preface," on page 619

- Section J.5.5, “Securing Access to NSS Volumes, Directories, and Files,” on page 620
- Section J.5.6, “What’s New for NSS,” on page 620

J.5.1 Managing Devices

Location	Change
“Planning for Device Sharing” on page 136	Devices that contain NetWare Traditional partitions and volumes can be marked as Shareable for Clustering, but Novell Cluster Services supports clustering only for NSS volumes on those devices.

J.5.2 Managing Files and Folders

Location	Change
“Viewing or Modifying File or Folder Properties” on page 391	This information was moved from the Securing Access section.
“Viewing or Modifying File Ownership” on page 394	This information was moved from the Securing Access to NSS Volumes, Directories, and Files section.

J.5.3 Management Tools for NSS

Location	Change
“Deleted Files” on page 111	Salvage and purge of deleted files and directories is available only for NSS volumes where the volume’s Salvage attribute is enabled.
“Properties” on page 111	Directory quotas management is available only for NSS volumes where the volume’s Directory Quotas attribute is enabled.

J.5.4 Preface

Location	Change
“Additional Documentation” on page 26	The “Storage and File Systems” section in the *OES 2 SP1: Planning and Implementation Guide* describes considerations for choosing a storage solution and system-wide caveats for implementing the different storage solutions. *OES 2 SP1: Storage and File Services Overview* describes typical requirements for system storage, and identifies the various storage products and services in Novell Open Enterprise Server 2 that address those requirements.

J.5.5 Securing Access to NSS Volumes, Directories, and Files

Location	Change
"Configuring File System Trustees, Trustee Rights, Inherited Rights Filters, and Attributes" on page 297	General information about files and folders was moved to "Viewing or Modifying File or Folder Properties" on page 391. File ownership information was moved to "Viewing or Modifying File Ownership" on page 394.

J.5.6 What's New for NSS

Location	Change
Section 2.2, "What's New for NSS (OES 2 and NetWare 6.5 SP7)," on page 39	Directory quotas management is available only for NSS volumes where the volume's Directory Quotas attribute is enabled. Salvage and purge of deleted files and directories is available only for NSS volumes where the volume's Salvage attribute is enabled.

J.6 December 14, 2007

Updates were made to the following sections. The changes are explained below.

- Section J.6.1, "Installing and Configuring Novell Storage Services," on page 620
- Section J.6.2, "Managing Backup and Restore for Data and Trustee Information," on page 621
- Section J.6.3, "Management Tools for NSS," on page 621
- Section J.6.4, "NSS Commands," on page 621
- Section J.6.5, "Planning Your Multipath I/O Solution for NSS on NetWare," on page 621

J.6.1 Installing and Configuring Novell Storage Services

Location	Change
"Upgrading from OES 1 Linux to OES 2 Linux" on page 53	This section is new. It addresses read performance problems you might experience for NSS volumes after upgrading from OES 1 Linux to OES 2 Linux.

J.6.2 Managing Backup and Restore for Data and Trustee Information

Location	Change
"Using Extended Attributes (xAttr) Commands (Linux)" on page 400	Procedures were added for enabling and disabling support for the Linux `xattr` API.

J.6.3 Management Tools for NSS

Location	Change
"Understanding Storage-Related Plug-Ins" on page 98	Added the paths and filenames of all storage-related plug-in files.

J.6.4 NSS Commands

Location	Change
"Read Ahead Blocks and Allocate Ahead Blocks Commands" on page 486	A note was added to address read performance problems you might experience for NSS volumes after upgrading from OES 1 Linux to OES 2 Linux.

J.6.5 Planning Your Multipath I/O Solution for NSS on NetWare

Location	Change
"Planning Your Multipath I/O Solution for NSS on NetWare" on page 189	You should set a non-zero priority on all of the paths for the device. If you set a non-zero priority on any of the paths, zero is considered the highest priority until all paths are set to a non-zero value.

J.7 November 14, 2007

Updates were made to the following sections. The changes are explained below.

- Section J.7.1, "Comparison of NSS on NetWare and NSS on Linux," on page 622
- Section J.7.2, "Extended Attributes (XAttr) Commands (Linux)," on page 622

J.7.1 Comparison of NSS on NetWare and NSS on Linux

Location	Change
Appendix F, “Comparison of NSS on NetWare and NSS on Linux,” on page 579	Added an entry for Dynamic Storage Technology.

J.7.2 Extended Attributes (XAttr) Commands (Linux)

Location	Change
“Using the Linux cp Command to Copy Files with Extended Attributes” on page 466	When you enable listxattr on OES 2, the Linux cp command requires that you use the --preserve=all option when you also want to copy the extended attributes. You must issue the command as the root user. `cp --preserve=all /path/file1 /newpath/file1`
Section A.11.3, “Additional Information,” on page 467	Added a link to the listxattr tool in Novell Cool Solutions.