opentext™

Open Enterprise Server 2023

OES Cluster Services for Linux Administration Guide

October 2022

Legal Notices

Copyright 2023 Open Text

The only warranties for products and services of Open Text and its affiliates and licensors (“Open Text”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Contents

About This Guide 13

1 Overview of OES Cluster Services 15

1.1 Why Should I Use Clusters? 15
1.2 Benefits of OES Cluster Services 16
1.3 Product Features 16
1.4 Clustering for High Availability 17
1.5 Shared Disk Scenarios 19
1.5.1 Using Fibre Channel Storage Systems 19
1.5.2 Using iSCSI Storage Systems 20
1.5.3 Using Shared SCSI Storage Systems 21
1.6 Terminology 21
1.6.1 The Cluster 21
1.6.2 Cluster Resources 22
1.6.3 Failover Planning 24

2 What's New or Changed for OES Cluster Services 27

2.1 What's New or Changed in OES Cluster Services (OES 2023) 27

3 Planning for a Cluster 29

3.1 Determining Your Design Criteria 29
3.2 Using Cluster Best Practices 29
3.3 Planning the LAN Connectivity 30
3.3.1 VLAN 30
3.3.2 Channel Bonding 31
3.3.3 Spanning Tree Protocol 31
3.3.4 IP Addresses 31
3.3.5 Name Resolution 31
3.4 Planning the Shared Storage Connectivity 32
3.5 Planning the Shared Storage Solution 32
3.6 Planning the eDirectory Deployment 32
3.6.1 Object Location 32
3.6.2 Cluster OU Context 33
3.6.3 Cluster OU Partitioning and Replication 33
3.7 Planning for Shared Storage as Cluster Resources 33
3.8 Planning for OES Services as Cluster Resources 33

4 Planning for OES Cluster Services 35

4.1 Cluster Administration Requirements 35
4.1.1 Cluster Installation Administrator 35
4.1.2 NCS Proxy User 36
4.1.3 Cluster Administrator or Administrator-Equivalent User 38
4.2 IP Address Requirements 38

4.3 Volume ID Requirements . . . 38
4.4 Hardware Requirements . . . 39
4.5 Virtualization Environments . . . 39
4.6 Software Requirements for Cluster Services . . . 39
4.6.1 Open Enterprise Server 2023 . . . 40
4.6.2 OES Cluster Services . . . 40
4.6.3 NetIQ eDirectory 9.2.6 . . . 40
4.6.4 SLP . . . 43
4.6.5 iManager 3.2.6 . . . 44
4.6.6 Clusters Plug-in for iManager . . . 44
4.6.7 Storage-Related Plug-Ins for iManager . . . 45
4.6.8 SFCB and CIMOM . . . 46
4.6.9 OES Credential Store . . . 47
4.6.10 Web Browser . . . 47
4.7 Software Requirements for Cluster Resources . . . 47
4.7.1 NCP Server for Linux . . . 48
4.7.2 Storage Services File System for Linux . . . 48
4.7.3 LVM Volume Groups and Linux POSIX File Systems . . . 49
4.7.4 NCP Volumes on Linux POSIX File Systems . . . 49
4.7.5 Dynamic Storage Technology Shadow Volume Pairs . . . 49
4.7.6 NCP File Access . . . 49
4.7.7 OES CIFS . . . 50
4.7.8 OES Domain Services for Windows . . . 50
4.8 Shared Disk Configuration Requirements . . . 50
4.8.1 Shared Devices . . . 50
4.8.2 SBD Partitions . . . 51
4.8.3 Shared iSCSI Devices . . . 53
4.8.4 Shared RAID Devices . . . 53
4.9 SAN Rules for LUN Masking . . . 53
4.10 Multipath I/O Configuration Requirements . . . 54
4.10.1 Path Failover Settings for Device Mapper Multipath . . . 54
4.10.2 Modifying the Port Down Retry Setting in the modprobe.conf.local File . . . 54
4.10.3 Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File . . . 55
4.10.4 Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS . . . 56

5 Installing, Configuring, and Repairing OES Cluster Services 57

5.1 OES Cluster Services Licensing . . . 58
5.2 Extending the eDirectory Schema to Add Cluster Objects . . . 58
5.2.1 Prerequisites for Extending the Schema . . . 58
5.2.2 Extending the Schema . . . 59
5.3 Assigning Install Rights for Container Administrators or Non-Administrator Users . . . 60
5.4 Installing OES Cluster Services . . . 61
5.4.1 Before You Install OES Cluster Services . . . 61
5.4.2 Installing OES Cluster Services during an OES Installation . . . 62
5.4.3 Installing OES Cluster Services on an Existing OES Server . . . 63
5.5 Configuring OES Cluster Services . . . 64
5.5.1 Initializing and Sharing a Device to Use for the SBD Partition . . . 64
5.5.2 Opening the Open Enterprise Server Configuration Page . . . 65
5.5.3 Using Different LDAP Credentials for the Cluster Configuration . . . 66
5.5.4 Accessing the OES Cluster Services Configuration Page in YaST . . . 66
5.5.5 Configuring a New Cluster . . . 67

5.5.6 Adding a Node to an Existing Cluster . . . 69
5.6 Configuring Additional Administrators . . . 71
5.7 Installing or Updating the Clusters Plug-in for iManager . . . 72
5.7.1 Prerequisites for Installing or Updating the Clusters Plug-In . . . 72
5.7.2 Installing the Clusters Plug-In . . . 73
5.7.3 Uninstalling and Reinstalling the Clusters Plug-In . . . 73
5.7.4 Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later . . . 74
5.8 Patching OES Cluster Services . . . 74
5.9 Removing the NCS Configuration Information from a Node . . . 75
5.10 Removing a Node from a Cluster . . . 77
5.11 Adding a Single Node That Was Previously in a Cluster . . . 78
5.12 Adding Multiple Nodes That Were Previously in a Cluster . . . 79
5.13 Adding a Node After Another Node is Permanently Removed from the Cluster . . . 80
5.14 What's Next . . . 80

6 Joining an OES 2023 Cluster Node and Cluster Resource to an Active Directory Domain 81

6.1 Joining the Cluster Node to an Active Directory Domain . . . 81
6.1.1 Prerequisites . . . 81
6.1.2 NSS-AD Support Pattern . . . 81
6.1.3 Configuring NSS for Active Directory . . . 81
6.2 Joining the Cluster Resource to an Active Directory Domain . . . 82
6.2.1 Prerequisites . . . 82
6.2.2 Steps for Joining the Cluster Resource . . . 82
6.3 Verifying the Service Principals and Computer Objects . . . 84
6.4 Removing the Cluster Resource from an Active Directory Domain . . . 86

7 Upgrading OES Clusters 89

7.1 Requirements and Guidelines for Upgrading Clusters . . . 89
7.1.1 Behavior of an NSS Pool Resource with Media Version 44.03 and Above in Mixed Node Cluster . . . 89
7.2 Upgrading Nodes in the Cluster (Rolling Cluster Upgrade) . . . 91

8 Configuring Cluster Policies, Protocols, and Properties 95

8.1 Understanding Cluster Settings . . . 95
8.1.1 Cluster Policies . . . 96
8.1.2 Cluster Priorities . . . 97
8.1.3 Cluster Protocols . . . 97
8.1.4 RME Groups . . . 98
8.1.5 BCC . . . 98
8.1.6 Cascade Failover Prevention . . . 98
8.1.7 Monitoring the eDirectory Daemon . . . 98
8.1.8 Cluster Node Reboot Behavior . . . 99
8.1.9 STONITH . . . 99
8.1.10 Cluster Information . . . 99
8.2 Setting Up a Personalized List of Clusters to Manage . . . 99
8.2.1 Adding Clusters to Your My Clusters List . . . 100
8.2.2 Viewing Information about Clusters in Your My Clusters List . . . 100
8.2.3 Personalizing the My Clusters Display . . . 101

8.2.4 Removing Clusters from a My Clusters List . . . 102
8.3 Selecting a Cluster to Manage . . . 102
8.4 Configuring Quorum Membership and Timeout Policies . . . 104
8.4.1 Quorum Membership (Number of Nodes) . . . 104
8.4.2 Quorum Timeout . . . 104
8.5 Configuring Cluster Protocols . . . 105
8.5.1 Heartbeat . . . 105
8.5.2 Tolerance . . . 106
8.5.3 Master Watchdog . . . 106
8.5.4 Slave Watchdog . . . 106
8.5.5 Maximum Retransmits . . . 106
8.6 Configuring Cluster Event Email Notification . . . 106
8.6.1 Configuring an Email Port . . . 107
8.6.2 Configuring the Notification Policies . . . 107
8.7 Configuring Cascade Failover Prevention . . . 108
8.7.1 Understanding the Cascade Failover Prevention Quarantine . . . 108
8.7.2 Releasing a Resource from Quarantine . . . 109
8.7.3 Enabling or Disabling Cascade Failover Prevention . . . 110
8.8 Configuring NCS to Monitor the eDirectory Daemon (ndsd) . . . 111
8.8.1 Understanding eDirectory Monitoring . . . 112
8.8.2 Displaying the Current Settings for NDSD Monitoring . . . 113
8.8.3 Configuring or Modifying NDSD Monitoring . . . 113
8.8.4 Disabling NDSD Monitoring . . . 114
8.9 Configuring the Cluster Node Reboot Behavior . . . 114
8.10 Configuring STONITH . . . 115
8.10.1 Sample Script for HP iLO . . . 115
8.10.2 Sample Script for Web-Based Power Switches . . . 116
8.11 Viewing or Modifying the NCS Proxy User Assignment in the NCS Management Group . . . 117
8.12 Viewing or Modifying the Cluster Master IP Address or Port . . . 118
8.13 Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster . . . 120
8.13.1 Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster . . . 120
8.13.2 Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources . . . 122
8.14 What's Next . . . 125

9 Managing Clusters 127

9.1 Starting and Stopping OES Cluster Services . . . 127
9.1.1 Starting OES Cluster Services . . . 128
9.1.2 Stopping OES Cluster Services . . . 128
9.1.3 Enabling and Disabling the Automatic Start of OES Cluster Services . . . 129
9.2 Taking the Cluster Down . . . 130
9.3 Selecting a Cluster to Manage . . . 130
9.4 Monitoring Cluster and Resource States . . . 131
9.5 Responding to a Start, Failover, or Failback Alert . . . 134
9.6 Viewing the Cluster Event Log in iManager . . . 135
9.7 Viewing the Cluster Event Log at the Command Line . . . 136
9.8 Viewing Summaries of Failed or Incomplete Events from the *.out Log Files . . . 137
9.9 Generating a Cluster Configuration Report . . . 138
9.10 Repairing a Cluster . . . 139
9.11 Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node . . . 140
9.11.1 Guidelines for Onlining and Offlining Resources . . . 140

9.11.2 Using iManager to Online and Offline Resources 140
9.11.3 Using the Cluster Online Command 141
9.11.4 Using the Cluster Offline Command 141
9.11.5 Using the Cluster Migrate Command 141
9.12 Cluster Migrating Resources to Different Nodes 141
9.13 Removing (Leaving) a Node from the Cluster 143
9.13.1 Removing One Node at a Time 143
9.13.2 Removing All Nodes at the Same Time (Taking the Cluster Down) 144
9.14 Joining a Node to the Cluster (Rejoining a Cluster) 144
9.15 Shutting Down OES Cluster Services When Servicing Shared Storage 145
9.16 Enabling or Disabling Cluster Maintenance Mode 145
9.17 Viewing the Cluster Node Properties 145
9.18 Creating or Deleting Cluster SBD Partitions 146
9.18.1 Requirements and Guidelines for Creating an SBD Partition 147
9.18.2 Before You Create a Cluster SBD Partition 149
9.18.3 Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL 149
9.18.4 Mirroring an Existing SBD Partition with NSSMU 152
9.18.5 Creating a Mirrored Cluster SBD Partition with SBDUTIL 154
9.18.6 Removing a Segment from a Mirrored Cluster SBD Partition 157
9.18.7 Deleting a Non-Mirrored Cluster SBD Partition 158
9.18.8 Deleting a Mirrored Cluster SBD 159
9.18.9 Additional Information about SBD Partitions 160
9.19 Customizing Cluster Services Management 160

10 Configuring and Managing Cluster Resources 163

10.1 Planning for Cluster Resources 163
10.1.1 Naming Conventions for Cluster Resources 164
10.1.2 Using Parameter Values with Spaces in a Cluster Script 164
10.1.3 Using Double Quotation Marks in a Cluster Script 164
10.1.4 Script Length Limits 165
10.1.5 Planning Cluster Maintenance 165
10.1.6 Number of Resources 165
10.1.7 Linux POSIX File System Types 165
10.2 Setting Up a Personalized List of Resources to Manage 165
10.2.1 Adding Resources to a My Resources List 166
10.2.2 Viewing Information about Resources in a My Resources List 166
10.2.3 Personalizing the My Resources Display 167
10.2.4 Removing Resources from a My Resources List 168
10.2.5 Updating the My Resources List after Renaming or Deleting a Resource 168
10.3 Using Cluster Resource Templates 169
10.3.1 Default Resource Templates 169
10.3.2 Viewing or Modifying a Resource Template in iManager 170
10.3.3 Creating a Resource Template 171
10.3.4 Synchronizing Locally Modified Resource Templates with Templates in eDirectory 173
10.4 Creating Cluster Resources 173
10.5 Configuring a Load Script for a Cluster Resource 174
10.6 Configuring an Unload Script for a Cluster Resource 175
10.7 Enabling Monitoring and Configuring the Monitor Script 176
10.7.1 Understanding Resource Monitoring 176
10.7.2 Configuring Resource Monitoring 178
10.7.3 Monitoring Services that Are Critical to a Resource 179
10.7.4 Example Monitor Scripts 180

10.8 Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration 181
10.9 Configuring the Start, Failover, and Failback Modes for Cluster Resources 182
10.9.1 Understanding Cluster Resource Modes 182
10.9.2 Setting the Start, Failover, and Failback Modes for a Resource 183
10.10 Configuring Preferred Nodes and Node Failover Order for a Resource 184
10.11 Configuring Resource Priorities for Load Order 185
10.12 Configuring Resource Mutual Exclusion Groups 185
10.12.1 Understanding Resource Mutual Exclusion 186
10.12.2 Setting Up RME Groups 187
10.12.3 Viewing RME Group Settings 188
10.13 Controlling Resource Monitoring 189
10.14 Changing the IP Address of a Cluster Resource 189
10.14.1 Pool Cluster Resource 189
10.14.2 Non-Pool Cluster Resources 190
10.15 Renaming a Cluster Resource 191
10.16 Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service . . 192
10.17 Additional Information for Creating Cluster Resources 195
10.17.1 Creating Storage Cluster Resources 195
10.17.2 Creating Service Cluster Resources 195
10.17.3 Creating Virtual Machine Cluster Resources 195

11 Quick Reference for Clustering Services and Data 197

12 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes 201

12.1 Requirements for Creating Pool Cluster Resources 202
12.1.1 OES Cluster Services 202
12.1.2 Resource IP Address 202
12.1.3 eDirectory 202
12.1.4 OES Storage Services 202
12.1.5 Shared Storage 204
12.1.6 NCP Server for Linux 204
12.1.7 OES CIFS for Linux 205
12.1.8 Domain Services for Windows 206
12.2 Guidelines for Using Pool Cluster Resources 206
12.2.1 Guidelines for Shared Pools 206
12.2.2 Guidelines for Shared Volumes 206
12.3 Initializing and Sharing a SAN Device 207
12.3.1 Initializing and Sharing a Device with iManager 208
12.3.2 Initializing and Sharing a Device with NSSMU 208
12.3.3 Initializing and Sharing a Device with NLVM 209
12.4 Creating Cluster-Enabled Pools and Volumes 210
12.4.1 Creating a Cluster-Enabled Pool and Volume with iManager 210
12.4.2 Creating a Cluster-Enabled Pool and Volume with NSSMU 213
12.5 Cluster-Enabling an Existing NSS Pool and Its Volumes 217
12.6 Configuring a Load Script for the Shared NSS Pool 222
12.6.1 Sample NSS Load Script 223
12.6.2 Specifying a Custom Mount Point for an NSS Volume 223
12.6.3 Specifying Custom Mount Options 224
12.6.4 Loading OES CIFS as an Advertising Protocol 224
12.7 Configuring an Unload Script for the Shared NSS Pool 224
12.7.1 Sample NSS Unload Script 225

12.7.2 Using Volume Dismount Commands . . . 225
12.7.3 Unloading OES CIFS as an Advertising Protocol . . . 226
12.8 Configuring a Monitor Script for the Shared NSS Pool . . . 226
12.8.1 Sample NSS Monitor Script . . . 226
12.8.2 Monitoring OES CIFS as an Advertising Protocol . . . 227
12.9 Adding Advertising Protocols for NSS Pool Cluster Resources . . . 227
12.10 Adding NFS Export for a Clustered Pool Resource . . . 228
12.10.1 Sample NFS Load Script . . . 228
12.10.2 Sample NFS Unload Script . . . 229
12.11 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes . . . 229
12.11.1 Understanding NSS Mirroring . . . 229
12.11.2 Requirements for NSS Mirroring . . . 231
12.11.3 Creating and Mirroring NSS Pools on Shared Storage . . . 231
12.11.4 Verifying the NSS Mirror Status in the Cluster . . . 233
12.12 Renaming a Clustered NSS Pool . . . 233
12.12.1 Renaming a Shared Pool with iManager . . . 234
12.12.2 Renaming a Shared Pool with NSSMU . . . 235
12.12.3 Renaming a Shared Pool and Its Resource Objects . . . 237
12.12.4 Renaming a Shared Pool in a DST Cluster Resource . . . 240
12.13 Renaming a Clustered NSS Volume . . . 242
12.13.1 Renaming a Shared NSS Volume with iManager . . . 243
12.13.2 Renaming a Shared NSS Volume with NSSMU . . . 245
12.13.3 Renaming a Shared NSS Volume in a DST Cluster Resource . . . 247
12.14 Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume) . . . 250
12.15 Changing the Name Space for a Clustered NSS Volume . . . 255
12.16 Adding a Volume to a Clustered Pool . . . 256
12.17 Changing an Assigned Volume ID . . . 257
12.18 Expanding the Size of a Clustered Pool . . . 257
12.18.1 Planning to Expand a Clustered Pool . . . 258
12.18.2 Increasing the Size of a Clustered Pool by Extending the Size of Its LUN . . . 259
12.18.3 Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN . . . 262
12.18.4 Verifying the Expanded Pool Cluster Resource on All Nodes . . . 265
12.19 Deleting NSS Pool Cluster Resources . . . 266
12.20 Disabling Clustering for a Pool . . . 266
12.21 Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool) . . . 268
12.22 Deleting a Clustered Pool . . . 271
12.22.1 Deleting a Clustered Pool on the Master Node . . . 271
12.22.2 Deleting a Clustered Pool on a Non-Master Node . . . 272

13 Configuring and Managing Cluster Resources for Shared LVM Volume Groups 275

13.1 Requirements for Creating LVM Cluster Resources . . . 275
13.1.1 OES Cluster Services . . . 276
13.1.2 Linux Logical Volume Manager 2 (LVM2) . . . 276
13.1.3 Clustered Logical Volume Manager Daemon (CLVMD) . . . 276
13.1.4 Resource IP Address . . . 276
13.1.5 Shared Storage Devices . . . 276
13.1.6 All Nodes Must Be Present . . . 277
13.1.7 NCP File Access with OES NCP Server . . . 277

13.1.8 Linux File Access Protocols . . . 278
13.2 Initializing a SAN Device . . . 278
13.2.1 Using NSSMU to Initialize a Device . . . 279
13.2.2 Using NLVM Commands to Initialize a Device . . . 279
13.3 Configuring an LVM Volume Group Cluster Resource with NSS Management Tools . . . 280
13.3.1 Sample Values . . . 281
13.3.2 Creating an LVM Volume Group Cluster Resource with NSSMU . . . 282
13.3.3 Creating an LVM Volume Group Cluster Resource with NLVM Commands . . . 286
13.3.4 Configuring the LVM Cluster Resource Settings . . . 289
13.3.5 Viewing or Modifying the LVM Resource Scripts . . . 290
13.3.6 Sample LVM Resource Scripts . . . 291
13.4 Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template . . . 295
13.4.1 Creating a Shared LVM Volume with LVM Commands . . . 295
13.4.2 Creating a Generic File System Cluster Resource for an LVM Volume Group . . . 300
13.4.3 Sample Generic LVM Resource Scripts . . . 305
13.5 Creating a Virtual Server Object for an LVM Volume Group Cluster Resource . . . 308
13.5.1 Creating an NCP Virtual Server for the LVM Resource . . . 309
13.5.2 Adding NCP Virtual Server Commands to the Resource Scripts . . . 309
13.5.3 Sample LVM Resource Scripts with an NCP Virtual Server . . . 311
13.6 Enabling NCP File Access for a Clustered LVM Volume . . . 314
13.6.1 Adding NCP Volume Commands to the Resource Scripts . . . 315
13.6.2 Creating a Shared NCP Volume Object . . . 317
13.6.3 Sample LVM Resource Scripts Enabled for NCP File Access . . . 318
13.7 Modifying the LVM Resource Scripts for Other Services . . . 321
13.8 Renaming the Mount Point Path for a Clustered LVM Volume . . . 321
13.9 Renaming a Clustered LVM Logical Volume . . . 323
13.10 Disabling Clustering for an LVM Volume Group and Logical Volume . . . 325
13.11 Deleting a Clustered LVM Volume Group and Logical Volume . . . 329
13.12 Deleting a Clustered LVM Volume (Created in NSSMU or NLVM) . . . 330
13.12.1 Deleting a Cluster-Enabled LVM Volume on the Master Node . . . 331
13.12.2 Deleting a Cluster-Enabled LVM Volume on a Non-Master Node . . . 332
13.13 Linux LVM Management Tools . . . 333
13.13.1 Using the YaST Expert Partitioner . . . 333
13.13.2 Using LVM Commands . . . 333

14 Configuring OES Cluster Services in a Virtualization Environment 337

14.1 Prerequisites for Xen Host Server Environments . . . 338
14.2 Virtual Machines as Cluster Resources . . . 338
14.2.1 Creating a Xen Virtual Machine Cluster Resource . . . 339
14.2.2 Configuring Virtual Machine Load, Unload, and Monitor Scripts . . . 340
14.2.3 Setting Up Live Migration . . . 347
14.3 Virtual Machines as Cluster Nodes . . . 347
14.4 Virtual Cluster Nodes in Separate Clusters . . . 348
14.5 Mixed Physical and Virtual Node Clusters . . . 349
14.6 Additional Information . . . 351

15 Troubleshooting OES Cluster Services 353

15.1 Migration or Online of Resources Causes Cluster Crash . . . 354
15.2 Cluster Node Crashing Issue . . . 354

15.3 Unable to Manage NCS from iManager When NSS is not Installed on the Node 355
15.4 Cluster Service Does Not Come Up After Upgrading to OES 2018 or Later if Service Proxy is Configured 355
15.5 Supported SBD Signatures in Cluster Nodes with Different OES Versions 356
15.6 Cluster Resource Is Not Accessible After a Network Service Restart 356
15.7 File Location Information 356
15.8 Version Issues 357
15.9 Diagnosing a Poison Pill Condition 357
15.10 Diagnosing Cluster Problems 357
15.11 Master IP Address Problem 358
15.12 Why didn't my resource go to the node I specified? 358
15.13 iManager: Error 500 Occurs while Connecting to a Cluster 359
15.14 iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5 359
15.15 iManager: Problems with Clusters Plug-In After Update for iManager 2.7.5 360
15.16 A Device Name Is Required to Create a Cluster Partition 360
15.17 Cluster Resource Goes Comatose After Migration or Failover 360
15.18 Cannot Authenticate to Remote Servers during Cluster Configuration 360
15.19 Cannot Connect to an iSCSI Target 361
15.20 Is There a way to Uninstall OES Cluster Services from a Server? 361
15.21 Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems 361
15.22 Error NFS Stale Handle Reported on Host when Cluster Failover Occurs 361
15.23 Error 20897 - This node is not a cluster member 362
15.24 Error 499 - Could not delete this resource Data_Server 362
15.25 Error 603 - smdr.novell is not registered with SLP for a new cluster resource 362
15.26 Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server 363
15.27 Communications Errors 363
15.28 Connections are not Getting Established if a Cluster Resource is Migrated from OES 2015 or Later to OES 11 SP2 Node 364

16 Security Considerations 365

16.1 Cluster Administration Rights 365
16.2 Ports 365
16.3 Email Alerts 366
16.4 Log Files 366
16.5 Configuration Files 366

A Console Commands for OES Cluster Services 369

A.1 Cluster Management Commands 369
A.2 Business Continuity Clustering Commands 378
A.3 OES Installation extend_schema Command 379
A.3.1 Syntax 379
A.3.2 Options 379
A.3.3 Example 380
A.4 SBD Utility 380
A.4.1 Syntax 380
A.4.2 Options 381

A.4.3 Return Value . . . 383
A.4.4 Examples . . . 384
A.5 DotOutParser Utility . . . 385
A.5.1 Syntax . . . 385
A.5.2 Options . . . 385
A.5.3 Example . . . 386
A.6 CSMPORT Utility (Cluster Segment Manager Import/Export) . . . 386
A.6.1 Syntax . . . 387
A.6.2 Options . . . 387
A.6.3 Examples . . . 389
A.7 ncs-configd.py Script (Pushing Configuration Updates to All Nodes) . . . 394
A.8 ncs_install.py Script (Extending the Schema, Monitoring NDSD, or Removing NCS Configuration) . . . 394
A.8.1 Extending the Schema . . . 395
A.8.2 Monitoring NDSD . . . 396
A.8.3 Removing NCS Configuration. . . . 397
A.9 ncs_ncpserv.py Script (Creating an NCP Virtual Server Object for a Clustered LVM Volume Group). . . . 399
A.9.1 Syntax . . . 399
A.9.2 Options . . . 399
A.9.3 Examples . . . 400
A.10 ncs_resource_scripts.pl Script (Modifying Resource Scripts Outside of iManager). . . . 401
A.10.1 Syntax . . . 401
A.10.2 Options . . . 402
A.10.3 Examples . . . 402
A.11 NCPCON Mount Command for NCP and NSS Volumes. . . . 403
A.11.1 Syntax . . . 403
A.11.2 NCP Volume Examples . . . 404
A.11.3 NSS Volume Examples . . . 404
A.11.4 Other Examples. . . . 405
A.12 Querylog Utility . . . 405
A.12.1 Syntax . . . 405
A.12.2 Options . . . 405
A.12.3 Examples . . . 406

B Files for OES Cluster Services 407

C Electing a Master Node 411

C.1 Master-Election Algorithm for OES 11 and Earlier . . . 411
C.2 New Master-Election Algorithm for OES 11 SP1 and Later . . . 413

D Clusters Plug-In Changes for NetIQ iManager 3.2.6 415

About This Guide

This guide describes how to install, upgrade, configure, and manage OES Cluster Services for Open Enterprise Server (OES) 2023. It is divided into the following sections:

NOTE: Micro Focus is now OpenText.

- Chapter 1, “Overview of OES Cluster Services,” on page 15
- Chapter 2, “What’s New or Changed for OES Cluster Services,” on page 27
- Chapter 3, “Planning for a Cluster,” on page 29
- Chapter 4, “Planning for OES Cluster Services,” on page 35
- Chapter 5, “Installing, Configuring, and Repairing OES Cluster Services,” on page 57
- Chapter 6, “Joining an OES 2023 Cluster Node and Cluster Resource to an Active Directory Domain,” on page 81
- Chapter 7, “Upgrading OES Clusters,” on page 89
- Chapter 8, “Configuring Cluster Policies, Protocols, and Properties,” on page 95
- Chapter 9, “Managing Clusters,” on page 127
- Chapter 10, “Configuring and Managing Cluster Resources,” on page 163
- Chapter 11, “Quick Reference for Clustering Services and Data,” on page 197
- Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 201
- Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 275
- Chapter 14, “Configuring OES Cluster Services in a Virtualization Environment,” on page 337
- Chapter 15, “Troubleshooting OES Cluster Services,” on page 353
- Chapter 16, “Security Considerations,” on page 365
- Appendix A, “Console Commands for OES Cluster Services,” on page 369
- Appendix B, “Files for OES Cluster Services,” on page 407
- Appendix C, “Electing a Master Node,” on page 411
- Appendix D, “Clusters Plug-In Changes for NetIQ iManager 3.2.6,” on page 415

Audience

This guide is intended for cluster administrators, or anyone who is involved in installing, configuring, and managing OES Cluster Services.

Understanding of file systems and services that are used in the cluster is assumed.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation. For online documentation see (https://www.microfocus.com/documentation/open-enterprise-server/2023/)

Documentation Updates

The latest version of this *OES Cluster Services for Linux Administration Guide* is available on the OES 2023 documentation website (https://www.microfocus.com/documentation/open-enterprise-server/2023/).

Additional Documentation

For information about other OES products, see the OES 2023 documentation website (https://www.microfocus.com/documentation/open-enterprise-server/2023/).

For links to information about clustering OES services and Linux services with Cluster Services, see Chapter 11, “Quick Reference for Clustering Services and Data,” on page 197.

For information about using Novell Cluster Services in a VMware virtual environment, see the *OES Cluster Services Implementation Guide for VMware*.

For information about converting clusters from NetWare to Linux, see the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*.

For information about Novell Cluster Services for NetWare 6.5 SP8, see the “Clustering NetWare Services” list on the NetWare 6.5 SP8 Clustering (High Availability) documentation website (http://www.novell.com/documentation/nw65/cluster-services.html#clust-config-resources).

1 Overview of OES Cluster Services

OES Cluster Services for Open Enterprise Server (OES) is a multiple-node server clustering system that ensures high availability and manageability of critical network resources including data, applications, and services. It stores information in NetIQ eDirectory about the cluster and its resources. It supports failover, failback, and cluster migration of individually managed cluster resources. You can cluster migrate a resource to a different server in order to perform rolling cluster maintenance or to balance the resource load across member nodes.

- Section 1.1, “Why Should I Use Clusters?,” on page 15
- Section 1.2, “Benefits of OES Cluster Services,” on page 16
- Section 1.3, “Product Features,” on page 16
- Section 1.4, “Clustering for High Availability,” on page 17
- Section 1.5, “Shared Disk Scenarios,” on page 19
- Section 1.6, “Terminology,” on page 21

1.1 Why Should I Use Clusters?

A server cluster is a group of redundantly configured servers that work together to provide highly available access for clients to important applications, services, and data while reducing unscheduled outages. The applications, services, and data are configured as cluster resources that can be failed over or cluster migrated between servers in the cluster. For example, when a failure occurs on one node of the cluster, the clustering software gracefully relocates its resources and current sessions to another server in the cluster. Clients connect to the cluster instead of an individual server, so users are not aware of which server is actively providing the service or data. In most cases, users are able to continue their sessions without interruption.

Each server in the cluster runs the same operating system and applications that are needed to provide the application, service, or data resources to clients. Storage is shared among servers. In case of failures, data can remain available through a different path. Clustering software monitors the health of each of the member servers by listening for its heartbeat, a simple message that lets the others know it is alive.

The cluster’s virtual server provides a single point for accessing, configuring, and managing the cluster servers and resources. The virtual identity is bound to the cluster’s master node and remains with the master node regardless of which member server acts the master node. The master server also keeps information about each of the member servers and the resources they are running. If the master server fails, the control duties are passed to another server in the cluster.

1.2 Benefits of OES Cluster Services

OES Cluster Services provides high availability for data and services running on OES servers. You can configure up to 32 OES servers in a high-availability cluster, where resources can be dynamically relocated to any server in the cluster. Each cluster resource can be configured to automatically fail over to a preferred server if there is a failure on the server where it is currently running. In addition, costs are lowered through the consolidation of applications and operations onto a cluster.

Cluster Services allows you to manage a cluster from a single point of control and to adjust resources to meet changing workload requirements (thus, manually “load balance” the cluster). Resources can also be cluster migrated manually to allow you to troubleshoot hardware. For example, you can move applications, websites, and so on to other servers in your cluster without waiting for a server to fail. This helps you to reduce unplanned service outages and planned outages for software and hardware maintenance and upgrades.

OES Cluster Services clusters provide the following benefits over stand-alone servers:

- Increased availability of applications, services, and data
- Improved performance
- Lower cost of operation
- Scalability
- Disaster recovery
- Data protection
- Server consolidation
- Storage consolidation

1.3 Product Features

OES Cluster Services includes several important features to help you ensure and manage the availability of your network resources:

- Support for shared SCSI, iSCSI, or Fibre Channel storage subsystems. Shared disk fault tolerance can be obtained by implementing RAID on the shared disk subsystem.
- Multi-node all-active cluster (up to 32 nodes). Any server in the cluster can restart resources (applications, services, IP addresses, and file systems) from a failed server in the cluster.
- A single point of administration through the browser-based Micro Focus iManager, which allows you to remotely manage the cluster. You can use the Clusters plug-in to manage and monitor clusters and cluster resources, and to configure the settings and scripts for cluster resources.
- The ability to tailor a cluster to the specific applications and hardware infrastructure that fit your organization.
- Dynamic assignment and reassignment of server storage as needed.
- The ability to use email to automatically notify administrators of cluster events and cluster state changes.

1.4 Clustering for High Availability

A OES Cluster Services for Linux cluster consists of the following components:

- 2 to 32 OES servers, each containing at least one local disk device.
- Cluster Services software running on each Linux server in the cluster.
- A shared disk subsystem connected to all servers in the cluster (optional, but recommended for most configurations).
- Equipment to connect servers to the shared disk subsystem, such as one of the following:
  - High-speed Fibre Channel cards, cables, and switches for a Fibre Channel SAN
  - Ethernet cards, cables, and switches for an iSCSI SAN
  - SCSI cards and cables for external SCSI storage arrays

The benefits that OES Cluster Services provides can be better understood through the following scenario.

Suppose you have configured a three-server cluster, with a web server installed on each of the three servers in the cluster. Each of the servers in the cluster hosts two websites. All the data, graphics, and web page content for each website is stored on a shared disk system connected to each of the servers in the cluster. Figure 1-1 depicts how this setup might look.

Figure 1-1 *Three-Server Cluster*

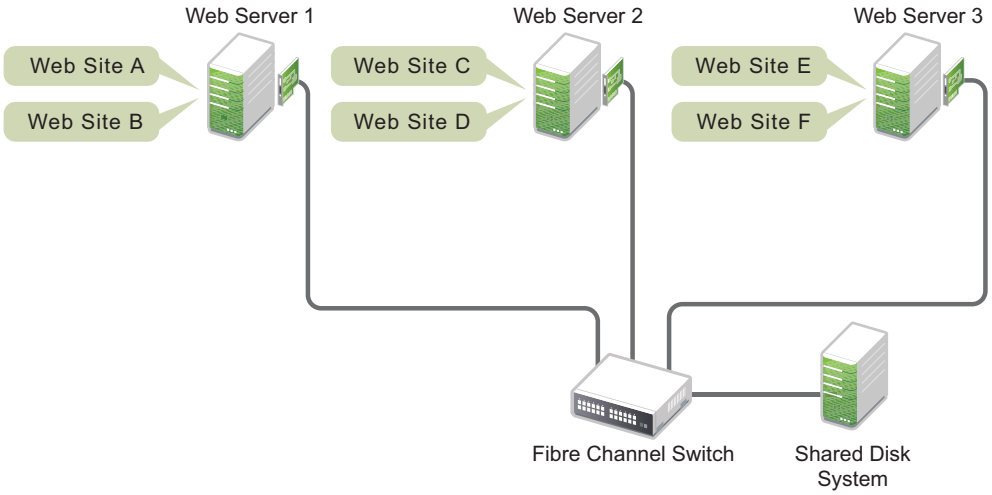

During normal cluster operation, each server is in constant communication with the other servers in the cluster and performs periodic polling of all registered resources to detect failure.

Suppose Web Server 1 experiences hardware or software problems and the users who depend on Web Server 1 for Internet access, email, and information lose their connections. Figure 1-2 shows how resources are moved when Web Server 1 fails.

Figure 1-2 *Three-Server Cluster after One Server Fails*

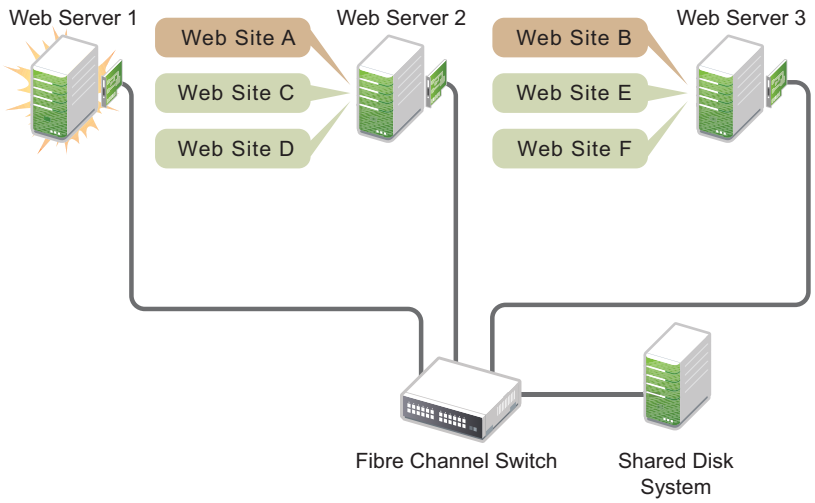

Web Site A moves to Web Server 2 and Web Site B moves to Web Server 3. IP addresses and certificates also move to Web Server 2 and Web Server 3.

When you configured the cluster, you decided where the websites hosted on each web server would go if a failure occurred. You configured Web Site A to move to Web Server 2 and Web Site B to move to Web Server 3. This way, the workload once handled by Web Server 1 is evenly distributed.

When Web Server 1 failed, Cluster Services software did the following:

- Detected a failure.
- Remounted the shared data directories (that were formerly mounted on Web Server 1) on Web Server 2 and Web Server 3 as specified.
- Restarted applications (that were running on Web Server 1) on Web Server 2 and Web Server 3 as specified.
- Transferred IP addresses to Web Server 2 and Web Server 3 as specified.

In this example, the failover process happened quickly and users regained access to website information within seconds, and in most cases, without logging in again.

Now suppose the problems with Web Server 1 are resolved, and Web Server 1 is returned to a normal operating state. Web Site A and Web Site B automatically fail back (that is, they are moved back to Web Server 1) if failback is configured for those resources, and Web Server operation returns to the way it was before Web Server 1 failed.

OES Cluster Services also provides resource migration capabilities. You can move applications, websites, and so on to other servers in your cluster without waiting for a server to fail.

For example, you could manually move Web Site A or Web Site B from Web Server 1 to either of the other servers in the cluster. You might want to do this to upgrade or perform scheduled maintenance on Web Server 1, or to increase performance or accessibility of the websites.

1.5 Shared Disk Scenarios

Typical cluster configurations normally include a shared disk subsystem connected to all servers in the cluster. The shared disk subsystem can be connected via high-speed Fibre Channel cards, cables, and switches, or it can be configured to use shared SCSI or iSCSI. If a server fails, another designated server in the cluster automatically mounts the shared disk directories previously mounted on the failed server. This gives network users continuous access to the directories on the shared disk subsystem.

- Section 1.5.1, “Using Fibre Channel Storage Systems,” on page 19
- Section 1.5.2, “Using iSCSI Storage Systems,” on page 20
- Section 1.5.3, “Using Shared SCSI Storage Systems,” on page 21

1.5.1 Using Fibre Channel Storage Systems

Fibre Channel provides the best performance for your storage area network (SAN). Figure 1-3 shows how a typical Fibre Channel cluster configuration might look.

Figure 1-3 *Typical Fibre Channel Cluster Configuration*

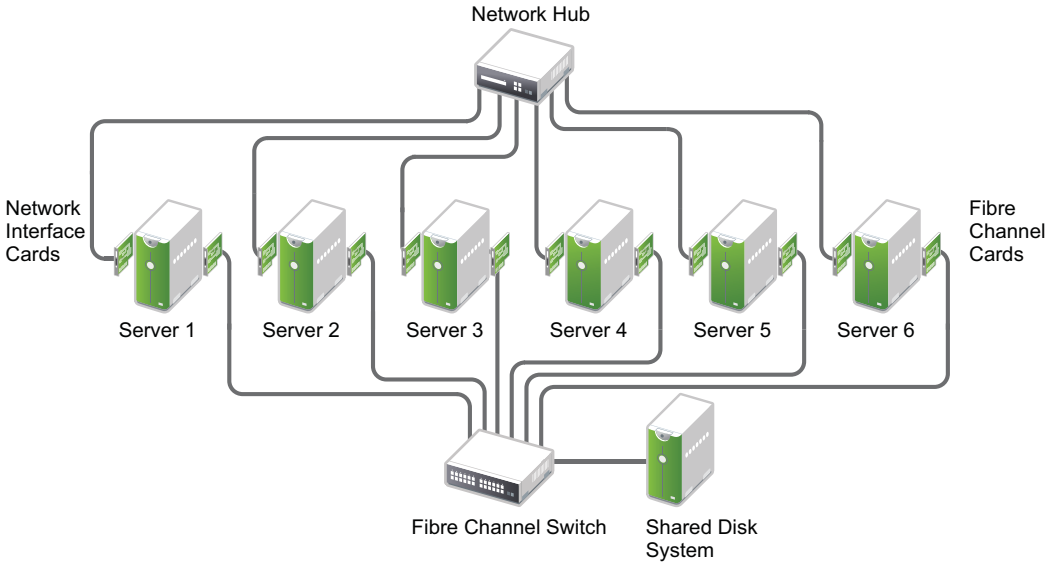

1.5.2 Using iSCSI Storage Systems

iSCSI is an alternative to Fibre Channel that can be used to create a lower-cost SAN with Ethernet equipment. Figure 1-4 shows how a typical iSCSI cluster configuration might look.

Figure 1-4 *Typical iSCSI Cluster Configuration*

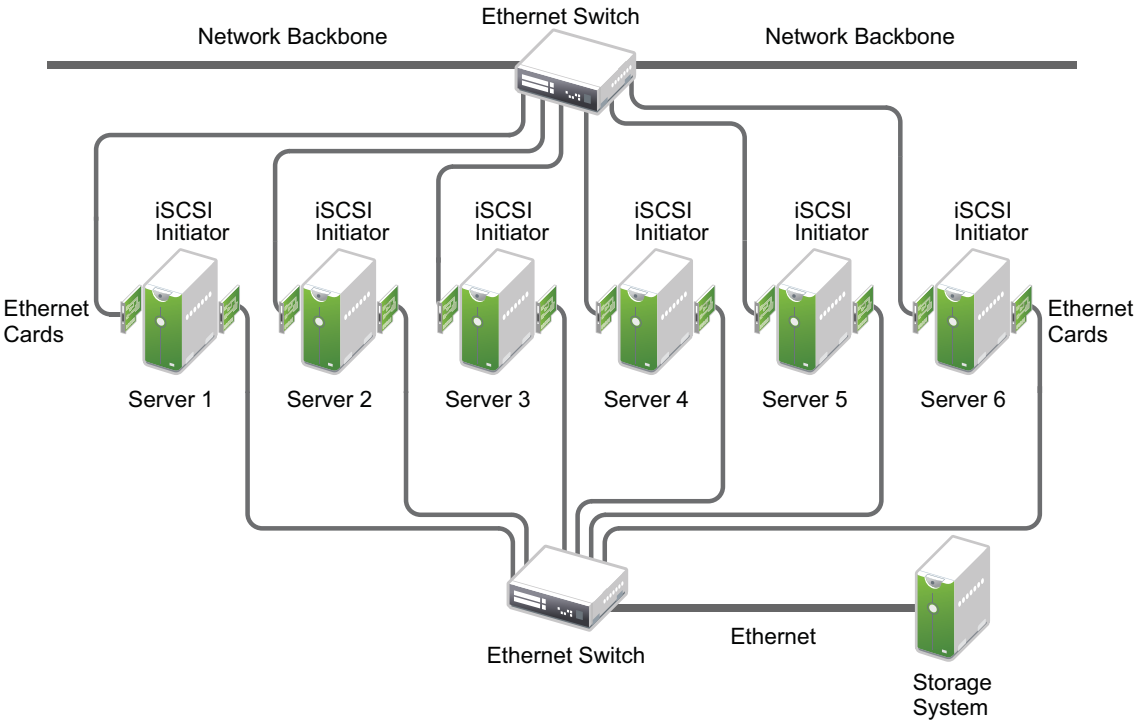

1.5.3 Using Shared SCSI Storage Systems

You can configure your cluster to use shared SCSI storage systems. This configuration is also a lower-cost alternative to using Fibre Channel storage systems. Figure 1-5 shows how a typical shared SCSI cluster configuration might look.

Figure 1-5 *Typical Shared SCSI Cluster Configuration*

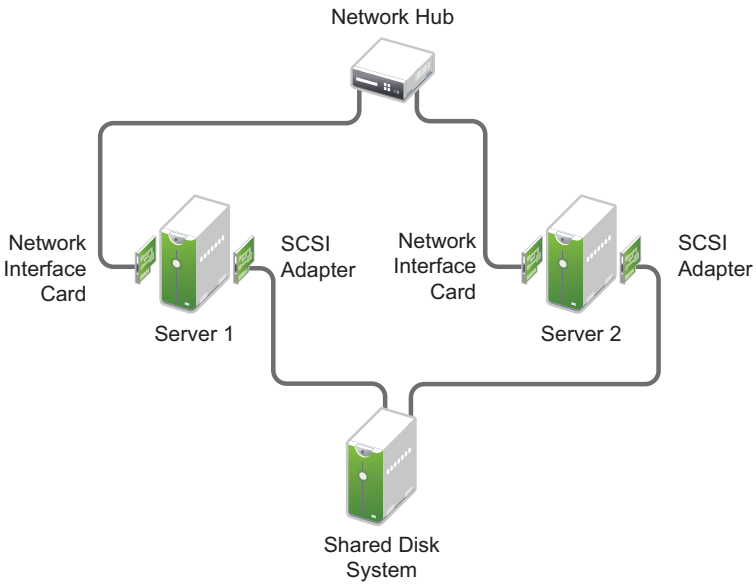

1.6 Terminology

Before you start working with a Cluster Services cluster, you should be familiar with the terms described in this section:

- Section 1.6.1, “The Cluster,” on page 21
- Section 1.6.2, “Cluster Resources,” on page 22
- Section 1.6.3, “Failover Planning,” on page 24

1.6.1 The Cluster

A cluster is a group 2 to 32 servers configured with OES Cluster Services so that data storage locations and applications can transfer from one server to another to provide high availability to users.

- “Cluster IP Address” on page 22
- “Server IP Address” on page 22
- “Master Node” on page 22
- “Slave Node” on page 22
- “Split-Brain Detector (SBD)” on page 22
- “Shared Storage” on page 22

Cluster IP Address

The unique static IP address for the cluster.

Server IP Address

Each server in the cluster has its own unique static IP address.

Master Node

The first server that comes up in an cluster is assigned the cluster IP address and becomes the master node. The master node monitors the health of the cluster nodes. It also synchronizes updates about the cluster to eDirectory. If the master node fails, Cluster Services migrates the cluster IP address to another server in the cluster, and that server becomes the master node. For information about how a new master is determined, see Appendix C, “Electing a Master Node,” on page 411.

Slave Node

Any member node in the cluster that is not currently acting as the master node.

Split-Brain Detector (SBD)

A small shared storage device where data is stored to help detect and prevent a split-brain situation from occurring in the cluster. If you use shared storage in the cluster, you must create an SBD for the cluster.

A split brain is a situation where the links between the nodes fail, but the nodes are still running. Without an SBD, each node thinks that the other nodes are dead, and that it should take over the resources in the cluster. Each node independently attempts to load the applications and access the data, because it does not know the other nodes are doing the same thing. Data corruption can occur. An SBD’s job is to detect the split-brain situation, and allow only one node to take over the cluster operations.

Shared Storage

Disks or LUNs attached to nodes in the cluster via SCSI, Fibre Channel, or iSCSI fabric. Only devices that are marked as shareable for clustering can be cluster-enabled.

1.6.2 Cluster Resources

A cluster resource is a single, logical unit of related storage, application, or service elements that can be failed over together between nodes in the cluster. The resource can be brought online or taken offline on one node at a time.

- “Resource IP Address” on page 23
- “NCS Virtual Server” on page 23
- “Resource Templates” on page 23
- “Service Cluster Resource” on page 23

- "Pool Cluster Resource" on page 24
- "Linux POSIX Volume Cluster Resource" on page 24
- "NCP Volume Cluster Resource" on page 24
- "DST Volume Cluster Resource" on page 24
- "Cluster Resource Scripts" on page 24

Resource IP Address

Each cluster resource in the cluster has its own unique static IP address.

NCS Virtual Server

An abstraction of a cluster resource that provides location independent access for users to the service or data. The user is not aware of which node is actually hosting the resource. Each cluster resource has a virtual server identity based on its resource IP address. A name for the virtual server can be bound to the resource IP address.

Resource Templates

A resource template contains the default load, unload, and monitor scripts and default settings for service or file system cluster resources. Resource templates are available for the following OES services and file systems:

- OES DHCP
- OES DNS
- Generic file system (for LVM-based Linux POSIX volumes)
- NSS file system (for NSS pool resources)
- Generic IP service
- OES iPrint
- MySQL
- Novell Samba
- CIS_Scale_Template
- CIS_Template

Personalized templates can also be created. See Section 10.3, "Using Cluster Resource Templates," on page 169.

Service Cluster Resource

An application or OES service that has been cluster-enabled. The application or service is installed on all nodes in the cluster where the resource can be failed over. The cluster resource includes scripts for loading, unloading, and monitoring. The resource can also contain the configuration information for the application or service.

Pool Cluster Resource

A cluster-enabled OES Storage Services pool. Typically, the shared pool contains only one NSS volume. The file system must be installed on all nodes in the cluster where the resource can be failed over. The NSS volume is bound to an NCS Virtual Server object (NCS:NCP Server) and to the resource IP address. This provides location independent access to data on the volume for NCP and OES CIFS clients.

Linux POSIX Volume Cluster Resource

A cluster-enabled Linux POSIX volume. The volume is bound to the resource IP address. This provides location-independent access to data on the volume via native Linux protocols such as FTP. You can optionally create an NCS Virtual Server object (NCS:NCP Server) for the resource as described in Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308.

NCP Volume Cluster Resource

An NCP volume (or share) that has been created on top of a cluster-enabled Linux POSIX volume. The NCP volume is re-created by a command in the resource load script whenever the resource is brought online. The NCP volume is bound to an NCS Virtual Server object (NCS:NCP Server) and to the resource IP address. This provides location-independent access to the data on the volume for NCP clients in addition to the native Linux protocols such as FTP. You must create an NCS Virtual Server object (NCS:NCP Server) for the resource as described in Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308.

DST Volume Cluster Resource

A cluster-enabled Dynamic Storage Technology volume made up of two shared NSS volumes. Both shared volumes are managed in the same cluster resource. The primary volume is bound to an NCS Virtual Server object (NCS:NCP Server) and to the resource IP address. This provides location independent access to data on the DST volume for NCP and OES CIFS clients.

Cluster Resource Scripts

Each cluster resource has a set of scripts that are run to load, unload, and monitor a cluster resource. The scripts can be personalized by using the Clusters plug-in for iManager.

1.6.3 Failover Planning

- “Heartbeat” on page 25
- “Quorum” on page 25
- “Preferred Nodes” on page 25
- “Resource Priority” on page 25
- “Resource Mutual Exclusion Groups” on page 25
- “Failover” on page 25
- “Fan-Out Failover” on page 25

- "Failback" on page 25
- "Cluster Migrate" on page 26
- "Leave a Cluster" on page 26
- "Join a Cluster" on page 26

Heartbeat

A signal sent between a slave node and the master node to indicate that the slave node is alive. This helps to detect a node failure.

Quorum

The administrator-specified number of nodes that must be up and running in the cluster before cluster resources can begin loading.

Preferred Nodes

One or more administrator-specified nodes in the cluster that can be used for a resource. The order of nodes in the Preferred Nodes list indicates the failover preference. Any applications that are required for a cluster resource must be installed and configured on the assigned nodes.

Resource Priority

The administrator-specified priority order that resources should be loaded on a node.

Resource Mutual Exclusion Groups

Administrator-specified groups of resources that should not be allowed to run on the same node at the same time. This Clusters plug-in feature is available only for clusters running OES 2 SP3 and later.

Failover

The process of automatically moving cluster resources from a failed node to an assigned functional node so that availability to users is minimally interrupted. Each resource can be failed over to the same or different nodes.

Fan-Out Failover

A configuration of the preferred nodes that are assigned for cluster resources so that each resource that is running on a node can fail over to different secondary nodes.

Failback

The process of returning cluster resources to their preferred primary node after the situation that caused the failover has been resolved.

Cluster Migrate

Manually triggering a move for a cluster resource from one node to another node for the purpose of performing maintenance on the old node, to temporarily lighten the load on the old node, and so on.

Leave a Cluster

A node leaves the cluster temporarily for maintenance. The resources on the node are cluster migrated to other nodes in their preferred nodes list.

Join a Cluster

A node that has previously left the cluster rejoins the cluster.

2 What's New or Changed for OES Cluster Services

This section describes the enhancements and changes in OES Cluster Services since the initial release Open Enterprise Server (OES) 2023.

2.1 What's New or Changed in OES Cluster Services (OES 2023)

Cluster Services in a VMware environment provides the following enhancements and changes in OES 2023.

VMware vMotion Support

Beginning with OES 2023, OES Cluster Services supports live migration of OES guest virtual machines using VMware vMotion. For more information, see VMware and vMotion Support in the OES Cluster Services Implementation Guide for VMware.

3 Planning for a Cluster

The success of your high-availability cluster solution depends on its stability and robustness. Use the guidelines in this section to design your OES Cluster Services cluster and cluster environment.

IMPORTANT: For information about the system requirements for installing and using Cluster Services, see Chapter 4, “Planning for OES Cluster Services,” on page 35.

- Section 3.1, “Determining Your Design Criteria,” on page 29
- Section 3.2, “Using Cluster Best Practices,” on page 29
- Section 3.3, “Planning the LAN Connectivity,” on page 30
- Section 3.4, “Planning the Shared Storage Connectivity,” on page 32
- Section 3.5, “Planning the Shared Storage Solution,” on page 32
- Section 3.6, “Planning the eDirectory Deployment,” on page 32
- Section 3.7, “Planning for Shared Storage as Cluster Resources,” on page 33
- Section 3.8, “Planning for OES Services as Cluster Resources,” on page 33

3.1 Determining Your Design Criteria

The purpose of designing a resilient cluster is to ensure that your essential data and services are highly available. Setting up data and services as cluster resources allows them to be moved between nodes in the same cluster. This helps eliminate or minimize the downtime caused by a server failure or maintenance.

You can determine what data and services to set up as cluster resources by asking and answering the following questions:

- ❒ What are the key services that drive your business?
- ❒ What services are essential for business continuance?
- ❒ What is the cost of downtime for the essential services?
- ❒ Based on their mission-critical nature and cost of downtime, what services are the highest priority for business continuance?
- ❒ What data is necessary to support the highest-priority services?
- ❒ How much data is involved, and how important is it?

3.2 Using Cluster Best Practices

Using the following cluster best practices can help you avoid potential problems with your cluster:

- Ensure that eDirectory is stable before implementing a cluster.

- Ensure that you have full Read/Write replicas of the entire eDirectory tree co-located in the data center where you are setting up the cluster.
- Ensure that IP addresses are unique.
- Consistently apply IP address assignments for each cluster and its cluster resources.
- Make IP address changes for the cluster and cluster resources only by using the procedure described in Section 8.13.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 122.

  IP address changes for cluster resources should always be made on the Protocols page of the iManager Clusters plug-in, not directly in load, unload, and monitor scripts. This is the only way to change the IP address on the virtual NCS:NCP Server object in eDirectory.
- Ensure that Volume IDs used for a cluster resources are unique across all nodes.

  Each cluster node automatically assigns volume ID 0 to volume SYS and volume ID 1 to volume _ADMIN. Cluster-enabled volumes use high volume IDs, starting from 254 in descending order. Volume IDs can be assigned in the cluster load script. You can view the volume IDs assigned on a node by using the `ncpcon volumes` command.

  The Client for Open Enterprise Server uses the volume ID to access a volume.
- Consider each node’s configuration requirements for each of the services it is intended to host.
- Create failover matrixes for each cluster resource so that you know what service is supported and which nodes are the preferred nodes for failover.

3.3 Planning the LAN Connectivity

The primary objective of LAN connectivity in a cluster is to provide uninterrupted heartbeat communications. Use the guidelines in this section to design the LAN connectivity for the cluster:

- Section 3.3.1, “VLAN,” on page 30
- Section 3.3.2, “Channel Bonding,” on page 31
- Section 3.3.3, “Spanning Tree Protocol,” on page 31
- Section 3.3.4, “IP Addresses,” on page 31
- Section 3.3.5, “Name Resolution,” on page 31

3.3.1 VLAN

Use a dedicated VLAN (virtual local area network) for each cluster.

The cluster protocol is non-routable, so you cannot direct communications to specific IP addresses. Using a VLAN for the cluster nodes provides a protected environment for the heartbeat process and ensures that heartbeat packets are exchanged only between the nodes of a given cluster.

When you use a VLAN, no foreign host can interfere with the heartbeat. For example, using a VLAN avoids broadcast storms that slow traffic and can result in false split-brain situations.

3.3.2 Channel Bonding

Servers should be redundantly cabled to the network in order to provide LAN fault tolerance, preferably at both the adapter level and the link level. Consider connecting cluster nodes to redundant access switches for fault tolerance.

Use channel bonding for the server adapters. Channel bonding combines Ethernet interfaces on a host computer for redundancy or increased throughput. Higher-level software uses a single virtual-network interface, and the channel bonding driver handles the complex choice of which physical-network interface to use. Channel bonding helps increase the availability of an individual cluster node, which helps avoid or reduce the occurrences of failover caused by slow LAN traffic. See the `/usr/src/linux/Documentation/bonding.txt` document.

3.3.3 Spanning Tree Protocol

Use the Spanning Tree Protocol (STP) to eliminate network topology loops. When you configure STP, ensure that the Portfast Bridge Protocol Data Unit (BPDU) guard feature is enabled, or consider using Rapid Spanning Tree Protocol (RSTP, IEEE 802.11w).

The default settings for STP inhibit the heartbeat for over 30 seconds whenever there is a change in link status. Test your STP configuration with Cluster Services running to ensure that a node is not cast out of the cluster when a broken link is restored.

3.3.4 IP Addresses

Plan your IP address assignment so that it is consistently applied across each cluster. For each cluster, provide a dedicated IP address range with sufficient addresses for the cluster. The addresses do not need to be contiguous.

You need a unique static IP address for each of the following components of a cluster:

- Cluster (master IP address)
- Cluster nodes
- Cluster resources (file system resources and service resources such as DHCP, DNS, SLP, FTP, and so on)

3.3.5 Name Resolution

Ensure that SLP is properly configured for name resolution. See Section 4.6.4, “SLP,” on page 43.

3.4 Planning the Shared Storage Connectivity

The primary objective of the shared storage connectivity in a cluster is to provide solid and stable connectivity between cluster nodes and the storage system. Before installing OES Cluster Services and setting up a cluster, ensure that the storage configuration is established and verified.

Use the guidelines in this section to design the storage connectivity for a cluster:

- Use host-based multipath I/O management. See the following resources:
  - Section 4.10, “Multipath I/O Configuration Requirements,” on page 54
  - “Managing Multipath I/O for Devices” (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-multipath.html) in the *SLES 12 Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html)
- Connect each node via two fabrics to the storage area network (SAN).
- Use redundant SAN connections to provide fault-tolerant connectivity between the cluster nodes and the shared storage devices.
- Use LUN masking to exclusively assign each LUN to one or more host connections. See Section 4.9, “SAN Rules for LUN Masking,” on page 53.

3.5 Planning the Shared Storage Solution

Use the guidelines in this section to design the shared storage solution for a cluster:

- For maximum flexibility, you should create only one cluster resource per LUN.

  A LUN cannot be concurrently accessed by servers belonging to different clusters. This means that all resources on a given LUN can be active only in a given cluster at any given time.
- You should use only one LUN per pool, and only one volume per pool. If you use multiple LUNs for a given shared NSS pool, all LUNs must fail over together.

  It is possible to create multiple pools per LUN or to use multiple LUNs per pool, but these alternatives are not recommended.

3.6 Planning the eDirectory Deployment

Your NetIQ eDirectory solution for each cluster must consider the following configuration elements. Your approach should be consistent across all clusters.

- Section 3.6.1, “Object Location,” on page 32
- Section 3.6.2, “Cluster OU Context,” on page 33
- Section 3.6.3, “Cluster OU Partitioning and Replication,” on page 33

3.6.1 Object Location

Cluster nodes and Cluster objects can exist in any container in the eDirectory tree. The Virtual Server object for the cluster and the objects for cluster resources are automatically created in the eDirectory context of the server where the cluster resource is created and cluster-enabled.

IMPORTANT: You should create cluster resources on the master node of the cluster.

3.6.2 Cluster OU Context

Before you create a new cluster, use iManager to create an OU container for the cluster, and use the OU container for the Cluster objects and Server objects.

3.6.3 Cluster OU Partitioning and Replication

Partition the Cluster OU, replicate it to dedicated eDirectory servers that are holding a replica of the parent partition, and replicate it to all cluster nodes. This helps prevent resources from being stuck in an NDS Sync state when a cluster resource's configuration is modified.

If you do not want to put a replica of eDirectory on the node, you must configure one or multiple LDAP servers for the node to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. For information about how to modify the LDAP server list that is used by a cluster, see Section 8.13.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 120.

3.7 Planning for Shared Storage as Cluster Resources

OES Cluster Services supports using cluster resources for the following file systems and storage solutions:

File System or Storage Solution	See
Storage Services (NSS) pools	Chapter 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 201
Linux LVM volume groups and logical volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 275
NCP (NetWare Control Protocol) volumes (NCP shares on cluster-enabled Linux POSIX volumes)	"Configuring NCP Volumes with OES Cluster Services" in the *OES 2023: NCP Server for Linux Administration Guide*
Dynamic Storage Technology (DST) volumes (NSS volumes configured in a shadow volume pair)	"Configuring DST Shadow Volume Pairs with OES Cluster Services" in the *OES 2023: Dynamic Storage Technology Administration Guide*

3.8 Planning for OES Services as Cluster Resources

OES Cluster Services supports using cluster resources for the following OES services:

Service	See
Apache Web Server	"Apache HTTP Server" in the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*
Certificate Server (NetIQ eDirectory Server Certificates)	"eDirectory Server Certificates" in the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*
CIFS (OES CIFS)	"Configuring CIFS with Cluster Services for an NSS File System" in the *OES 2023: OES CIFS for Linux Administration Guide*
CIS (Cloud Integrated Storage)	"Installing and Configuring CIS Server with Cluster Services" in the OES 2023: CIS Administration Guide
DFS VLDB (OES Distributed File Services Volume Location Database)	"Clustering Distributed File Services" in the *OES 2023: Distributed File Services Administration Guide for Linux*
DHCP Server on a Linux POSIX volume	"Configuring DHCP with OES Cluster Services for the Linux File System" in the *OES 2023: DNS/DHCP Services for Linux Administration Guide*
DHCP Server on an NSS volume	"Configuring DHCP with OES Cluster Services for the NSS File System" in the *OES 2023: DNS/DHCP Services for Linux Administration Guide*
DNS Server	"Configuring DNS with OES Cluster Services" in the *OES 2023: DNS/DHCP Services for Linux Administration Guide*
iPrint	"Configuring iPrint with OES Cluster Services" in the *OES 2023: iPrint Administration Guide*
MySQL	"High Availability and Scalability" in the *MySQL 5.x Reference Manual* (http://dev.mysql.com/doc/refman/5.5/en/ha-overview.html) "MySQL" in the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*
PureFTPd (OES FTP)	"Cluster Enabling Pure-FTPd in an OES Environment" in the *OES 2023: Planning and Implementation Guide*

4 Planning for OES Cluster Services

This section describes the requirements for installing and using OES Cluster Services on Open Enterprise Server (OES) 2023 servers.

IMPORTANT: For information about designing your cluster and cluster environment, see Chapter 3, “Planning for a Cluster,” on page 29.

- Section 4.1, “Cluster Administration Requirements,” on page 35
- Section 4.2, “IP Address Requirements,” on page 38
- Section 4.3, “Volume ID Requirements,” on page 38
- Section 4.4, “Hardware Requirements,” on page 39
- Section 4.5, “Virtualization Environments,” on page 39
- Section 4.6, “Software Requirements for Cluster Services,” on page 39
- Section 4.7, “Software Requirements for Cluster Resources,” on page 47
- Section 4.8, “Shared Disk Configuration Requirements,” on page 50
- Section 4.9, “SAN Rules for LUN Masking,” on page 53
- Section 4.10, “Multipath I/O Configuration Requirements,” on page 54

4.1 Cluster Administration Requirements

You use different credentials to install and set up the cluster and to manage the cluster. This section describes the tasks performed and rights needed for those roles.

- Section 4.1.1, “Cluster Installation Administrator,” on page 35
- Section 4.1.2, “NCS Proxy User,” on page 36
- Section 4.1.3, “Cluster Administrator or Administrator-Equivalent User,” on page 38

4.1.1 Cluster Installation Administrator

Typically, a tree administrator user installs and sets up the first cluster in a tree, which allows the schema to be extended. However, the tree administrator can extend the schema separately, and then set up the necessary permissions for a container administrator to install and configure the cluster.

NOTE: If the eDirectory administrator user name or password contains special characters (such as $, #, and so on), some interfaces in iManager and YaST might not handle the special characters. If you encounter problems, try escaping each special character by preceding it with a backslash (\) when you enter credentials.

- "eDirectory Schema Administrator" on page 36
- "Container Administrator" on page 36

eDirectory Schema Administrator

A tree administrator user with credentials to do so can extend the eDirectory schema before a cluster is installed anywhere in a tree. Extending the schema separately allows a container administrator to install a cluster in a container in that same tree without needing full administrator rights for the tree.

For instructions, see Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 58.

IMPORTANT: It is not necessary to extend the schema separately if the installer of the first cluster server in the tree has the eDirectory rights necessary to extend the schema.

Container Administrator

After the schema has been extended, the container administrator (or non-administrator user) needs the following eDirectory rights to install Cluster Services:

- Attribute Modify rights on the NCP Server object of each node in the cluster.
- Object Create rights on the container where the NCP Server objects are.
- Object Create rights where the cluster container will be.

For instructions, see Section 5.3, "Assigning Install Rights for Container Administrators or Non-Administrator Users," on page 60.

4.1.2 NCS Proxy User

During the cluster configuration, you must specify an NCS Proxy User. This is the user name and password that Cluster Services uses when the cluster management tools exchange information with eDirectory.

Cluster Services supports the OES Common Proxy User enablement feature of eDirectory. The proxy user is represented in eDirectory as a User object named `OESCommonProxy_<server_name>.<context>`. If the OES Common Proxy user is enabled in eDirectory when you configure a node for the cluster, the default NCS Proxy User is set to use the server's OES Common Proxy User. You can alternatively specify the LDAP Admin user or another administrator user.

The specified NCS Proxy User for the node is automatically assigned as a member in the `<cluster_name>_MGT_GRP.<context>` group that resides in the Cluster object container. The group accommodates the server-specific NCS Proxy Users that you assign when you configure each

node for the cluster. Members are added to the group as you configure each node for a cluster. Each member of the group has the necessary rights for configuring the cluster and cluster resources and for exchanging information with eDirectory.

IMPORTANT: You can modify this default administrator user name or password for the user name assigned as the NCS Proxy User after the install by following the procedure in Section 8.13, "Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster," on page 120.

Consider the following caveats for the three proxy user options:

- "OES Common Proxy User" on page 37
- "LDAP Admin User" on page 37
- "Another Administrator User" on page 37

OES Common Proxy User

If you specify the OES Common Proxy user for a cluster and later disable the Common Proxy user feature in eDirectory, the LDAP Admin user is automatically assigned to the <*cluster_name*>_MGT_GRP.<*context*> group, and the OES Common Proxy user is automatically removed from the group.

If an OES Common Proxy User is renamed, moved, or deleted in eDirectory, eDirectory takes care of the changes needed to modify the user information in the <*cluster_name*>_MGT_GRP.<*context*> group.

If a cluster node is removed from the tree, the OES Common Proxy User for that server is one of the cluster objects that needs to be deleted from eDirectory.

For information about enabling or disabling the OES Common Proxy User, see the *OES 2023: Installation Guide*. For caveats and troubleshooting information for the OES Common Proxy user, see the *OES 2023: Planning and Implementation Guide*.

LDAP Admin User

If you specify the LDAP Admin user as the NCS Proxy User, you typically continue using this identity while you set up the cluster and cluster resources. After the cluster configuration is completed, you create another user identity to use for NCS Proxy User, and grant that user sufficient administrator rights as specified in "Cluster Administrator or Administrator-Equivalent User" on page 38.

Another Administrator User

You can specify an existing user name and password to use for the NCS Proxy User. Cluster Services adds this user name to the <*cluster_name*>_MGT_GRP.<*context*> group.

4.1.3 Cluster Administrator or Administrator-Equivalent User

After the install, you can add other users (such as the tree administrator) as administrator equivalent accounts for the cluster by configuring the following for the user account:

- Give the user the Supervisor right to the Server object of each of the servers in the cluster.
- Linux-enable the user account with Linux User Management (LUM).
- Make the user a member of a LUM-enabled administrator group that is associated with the servers in the cluster.

4.2 IP Address Requirements

- ❒ Each server in the cluster must be configured with a unique static IP address.
- ❒ You need additional unique static IP addresses for the cluster and for each cluster resource and cluster-enabled pool.
- ❒ All IP addresses used by the master cluster IP address, its cluster servers, and its cluster resources must be on the same IP subnet. They do not need to be contiguous addresses.

4.3 Volume ID Requirements

A volume ID is a value assigned to represent the volume when it is mounted by NCP Server on an OES server. The Client for Open Enterprise Server accesses a volume by using its volume ID. Volume ID values range from 0 to 254. On a single server, volume IDs must be unique for each volume. In a cluster, volume IDs must be unique across all nodes in the cluster.

Unshared volumes are typically assigned low numbers, starting from 2 in ascending order. Volume IDs 0 and 1 are reserved. Volume ID 0 is assigned by default to volume SYS. Volume ID 1 is assigned by default to volume _ADMIN.

Cluster-enabled volumes use high volume IDs, starting from 254 in descending order. When you cluster-enable a volume, Cluster Services assigns a volume ID in the resource's load script that is unique across all nodes in a cluster. You can modify the resource's load script to change the assigned volume ID, but you must manually ensure that the new value is unique.

In a OES Business Continuity Clustering (BCC) cluster, the volume IDs of BCC-enabled clustered volumes must be unique across all nodes in every peer cluster. However, clustered volumes in different clusters might have the same volume IDs. Duplicate volume IDs can prevent resources from going online if the resource is BCC-migrated to a different cluster. When you BCC-enable a volume, you must manually edit its load script to ensure that its volume ID is unique across all nodes in every peer cluster. You can use the `ncpcon volumes` command on each node in every peer cluster to identify the volume IDs in use by all mounted volumes. Compare the results for each server to identify the clustered volumes that have duplicate volume IDs assigned. Modify the load scripts to manually assign unique volume IDs.

4.4 Hardware Requirements

The following hardware requirements for installing OES Cluster Services represent the minimum hardware configuration. Additional hardware might be necessary depending on how you intend to use Cluster Services.

- ❒ A minimum of two Linux servers, and not more than 32 servers in a cluster
- ❒ At least 512 MB of additional memory on each server in the cluster
- ❒ One non-shared device on each server to be used for the operating system
- ❒ At least one network card per server in the same IP subnet.

  In addition to Ethernet NICs, Cluster Services supports VLAN on NIC bonding in OES (with the latest patches applied) or later. No modifications to scripts are required. You can use eth*x* or vlan*x* interfaces in any combination for nodes in a cluster.

In addition, each server must meet the requirements for Open Enterprise Server 2023. See "Meeting All Server Software and Hardware Requirements" in the *OES 2023: Installation Guide*.

NOTE: Although identical hardware for each cluster server is not required, having servers with the same or similar processors and memory can reduce differences in performance between cluster nodes and make it easier to manage your cluster. There are fewer variables to consider when designing your cluster and failover rules if each cluster node has the same processor and amount of memory.

If you have a Fibre Channel SAN, the host bus adapters (HBAs) for each cluster node should be identical and be configured the same way on each node.

4.5 Virtualization Environments

Xen and KVM virtualization software is included with SUSE Linux Enterprise Server. OES Cluster Services supports using Xen or KVN virtual machine (VM) guest servers as nodes in a cluster. You can install Cluster Services on the guest server just as you would a physical server.

4.6 Software Requirements for Cluster Services

Ensure that your system meets the following software requirements for installing and managing OES Cluster Services:

- Section 4.6.1, "Open Enterprise Server 2023," on page 40
- Section 4.6.2, "OES Cluster Services," on page 40
- Section 4.6.3, "NetIQ eDirectory 9.2.6," on page 40
- Section 4.6.4, "SLP," on page 43
- Section 4.6.5, "iManager 3.2.6," on page 44
- Section 4.6.6, "Clusters Plug-in for iManager," on page 44
- Section 4.6.7, "Storage-Related Plug-Ins for iManager," on page 45
- Section 4.6.8, "SFCB and CIMOM," on page 46

- Section 4.6.9, “OES Credential Store,” on page 47
- Section 4.6.10, “Web Browser,” on page 47

4.6.1 Open Enterprise Server 2023

OES Cluster Services for Linux supports Open Enterprise Server 2023. OES Cluster Services is one of the OES Services patterns.

We recommend having uniform nodes in the cluster. The same release version of OES must be installed and running on each node in the cluster.

Mixed-mode clusters with different operating system platforms are supported during rolling cluster upgrades or conversions for the following scenarios:

Upgrading from	See:
OES 11 SP2 or later	Chapter 7, “Upgrading OES Clusters,” on page 89
NetWare 6.5 SP8	*OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*

4.6.2 OES Cluster Services

OES Cluster Services is required for creating and managing clusters and shared resources on your OES servers. Cluster Services is one of the OES Services patterns on OES.

4.6.3 NetIQ eDirectory 9.2.6

NetIQ eDirectory 9.1 is required for managing the Cluster object and Cluster Node objects for Cluster Services. eDirectory must be installed and running in the same tree where you create the cluster. eDirectory can be installed on any node in the cluster, on a separate server, or in a separate cluster. You can install an eDirectory master replica or replica in the cluster, but it is not required to do so for Cluster Services.

For information about using eDirectory, see *NetIQ eDirectory Administration Guide*.

IMPORTANT: Because the cluster objects and their settings are stored in eDirectory, eDirectory must be running and working properly whenever you modify the settings for the cluster or the cluster resources.

In addition, ensure that your eDirectory configuration meets the following requirements:

- “eDirectory Tree” on page 41
- “eDirectory Context” on page 41
- “Cluster Object Container” on page 41
- “Cluster Objects Stored in eDirectory” on page 41
- “LDAP Server List” on page 43

eDirectory Tree

All servers in the cluster must be in the same eDirectory tree.

eDirectory Context

If you are creating a new cluster, the eDirectory context where the new Cluster object will reside must be an existing context. Specifying a new context during the Cluster Services configuration does not create a new context.

Cluster Object Container

We recommend that the Cluster object and all of its member Server objects and Storage objects be located in the same OU context. Multiple Cluster objects can co-exist in the same eDirectory container. In iManager, use **Directory Administration > Create Object** to create a container for the cluster before you configure the cluster.

If the servers in the cluster are in separate eDirectory containers, the user that administers the cluster must have rights to the cluster server containers and to the containers where any cluster-enabled pool objects are stored. You can do this by adding trustee assignments for the cluster administrator to a parent container of the containers where the cluster server objects reside. See "eDirectory Rights" in the *NetIQ eDirectory Administration Guide* for more information.

Renaming a pool involves changing information in the Pool object in eDirectory. If Server objects for the cluster nodes are in different containers, you must ensure that the shared pool is active on a cluster node that has its NCP server object in the same context as the Pool object of the pool you are going to rename. For information about renaming a shared pool, see Section 12.12, "Renaming a Clustered NSS Pool," on page 233.

Cluster Objects Stored in eDirectory

Table 4-1 shows the cluster objects that are automatically created and stored in eDirectory under the Cluster object ( ) after you create a cluster:

Table 4-1 *Cluster Objects*

Icon	eDirectory Object
	Master_IP_Address_Resource
	Cluster Node object (*servername*)

Icon	eDirectory Object
	Resource Template objects. There are some default templates: AV_Template DHCP_Template DNS_Template Generic_FS_Template Generic_IP_Service iFolder_Template iPrint_Template MySQL_Template Samba_Template Xen_Template XenLive_Template CIS_Scale_Template CIS_Template

Table 4-2 shows the cluster objects that are added to eDirectory when you add nodes or create cluster resources:

Table 4-2 *Cluster Resource Objects*

Icon	eDirectory Object
	Cluster Node object (*servername*)
	NSS Pool Resource object (*poolname_SERVER*)
	Resource object

Table 4-3 shows the cluster objects that are added to eDirectory when you add nodes or create cluster resources in a OES Business Continuity Cluster, which is made up of OES Cluster Services clusters:

Table 4-3 *BCC Cluster Resource Objects*

Icon	eDirectory Object
	BCC NSS Pool Resource object
	BCC Resource Template object
	BCC Resource object

LDAP Server List

If eDirectory is not installed on a node, it looks to the LDAP server list for information about which LDAP server to use. As a best practice, you should list the LDAP servers in the following order:

- Local to the cluster
- Closest physical read/write replica

For information about configuring a list of LDAP servers for the cluster, see Section 8.13.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 120.

4.6.4 SLP

SLP (Service Location Protocol) is a required component for Novell Cluster Services on Linux when you are using NCP to access file systems on cluster resources. NCP requires SLP for the `ncpcon bind` and `ncpcon unbind` commands in the cluster load and unload scripts. For example, NCP is needed for NSS volumes and for NCP volumes on Linux POSIX file systems.

SLP is not automatically installed when you select Novell Cluster Services. SLP is installed as part of the eDirectory configuration during the OES installation. You can enable and configure SLP on the eDirectory Configuration - NTP & SLP page. For information, see “Specifying SLP Configuration Options” in the *OES 2023: Installation Guide*.

When the SLP daemon (`slpd`) is not installed and running on a cluster node, any cluster resource that contains the `ncpcon bind` command goes comatose when it is migrated or failed over to the node because the bind cannot be executed without SLP.

The SLP daemon (`slpd`) must also be installed and running on all nodes in the cluster when you manage the cluster or cluster resources.

NCP Server re-registers cluster resource virtual NCP servers with SLP based on the setting for the eDirectory advertise-life-time (`n4u.nds.advertise-life-time`) parameter. The parameter is set by default to 3600 seconds (1 hour) and has a valid range of 1 to 65535 seconds.

You can use the `ndsconfig set` command to set the `n4u.nds.advertise-life-time` parameter. To reset the parameter in a cluster, perform the following tasks on each node of the cluster:

1. Log in to the node as the `root` user, then open a terminal console.
2. Take offline all of the cluster resources on the node, or cluster migrate them to a different server. At a command prompt, enter

   ```
   cluster offline <resource_name>
   ```

   or

   ```
   cluster migrate <resource_name> <target_node_name>
   ```

3. Modify the eDirectory SLP advertising timer parameter (`n4u.nds.advertise-life-time`), then restart ndsd and slpd. At a command prompt, enter

```
ndsconfig set n4u.nds.advertise-life-time=<value_in_seconds>

rcndsd restart

rcslpd restart
```

4 Bring online all of the cluster resources on the node, or cluster migrate the previously migrated resources back to this node.

```
cluster online <resource_name>

or

cluster migrate <resource_name> <node_name>
```

5 Repeat the previous steps on the other nodes in the cluster.

OpenSLP stores the registration information in cache. You can configure the SLP Directory Agents to preserve a copy of the database when the SLP daemon (`slpd`) is stopped or restarted. This allows SLP to know about registrations immediately when it starts.

For more information about configuring and managing SLP, see "Configuring OpenSLP for eDirectory" in the NetIQ eDirectory Administration Guide.

4.6.5 iManager 3.2.6

iManager is required for configuring and managing clusters on OES.

iManager must be installed on at least one computer in the same tree as the cluster. It can be installed in the cluster or not in the cluster. For information about using iManager, see the iManager documentation website (https://www.netiq.com/documentation/imanager-32/).

For SFCB (Small Footprint CIM Broker) and CIMOM requirements, see Section 4.6.8, "SFCB and CIMOM," on page 46.

For browser configuration requirements, see "Web Browser" on page 47.

4.6.6 Clusters Plug-in for iManager

The Clusters plug-in for iManager provides the Clusters role where you can manage clusters and cluster resources with OES Cluster Services. The plug-in can be used on all operating systems supported by iManager and iManager Workstation.

The following components must be installed in iManager:

- Clusters (`ncsmgmt.rpm`)
- Common code for storage-related plug-ins (`storagemgmt.rpm`)

  See "Storage-Related Plug-Ins for iManager" on page 45.

If iManager is also installed on the server, these files are automatically installed in iManager when you install Cluster Services.

The Clusters plug-in also provides an integrated management interface for OES Business Continuity Clustering (BCC). The additional interface is present only if BCC is installed on the server. See the following table for information about the versions of BCC that are supported. BCC is sold separately from OES. For purchasing information, see the BCC product page.

BCC Release	OES Support	iManager and Clusters Plug-In
BCC 2.6	OES 2018 SP1 or later	Novell iManager 3.1 or later Requires the Clusters plug-in for OES 2018 SP1 or later with the latest patches applied. See the BCC administration guide.

4.6.7 Storage-Related Plug-Ins for iManager

In OES 11 and later, the following storage-related plug-ins for iManager share code is common in the `storagemgmt.rpm` file:

Product	Plug-In	NPM File
OES CIFS	File Protocols > CIFS	`cifsmgmt.rpm`
OES Cluster Services	Clusters	`ncsmgmt.rpm`
OES Distributed File Services	Distributed File Services	`dfsmgmt.rpm`
OES Storage Services	Storage	`nssmgmt.rpm`

These additional plug-ins are needed when working with the NSS file system. Ensure that you include the common `storagemgmt.rpm` plug-in module when installing any of these storage-related plug-ins.

IMPORTANT: If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to ensure that the common code works for all plug-ins.

Ensure that you uninstall the old version of the plug-ins before you attempt to install the new versions of the plug-in files.

The plug-in files are included on the installation disk. The latest storage-related plug-ins can be downloaded as a single zipped download file from the Downloads website (http://download.novell.com). For information about installing plug-ins in iManager, see Downloading and Installing Plug-in Modules in the NetIQ iManager Administration Guide.

For information about working with storage-related plug-ins for iManager, see "Understanding Storage-Related Plug-Ins" in the *OES 2023: NSS File System Administration Guide for Linux*.

4.6.8 SFCB and CIMOM

The Small Footprint CIM Broker (SFCB) replaces OpenWBEM for CIMOM activities in OES 11 and later. SFCB provides the default CIMOM and CIM clients. When you install any OES components that depend on WBEM, SFCB and all of its corresponding packages are installed with the components. See "Small Footprint CIM Broker (SFCB)" in the *OES 2023: Planning and Implementation Guide*.

IMPORTANT: SFCB must be running and working properly whenever you modify the settings for the cluster or the cluster resources.

Port 5989 is the default setting for Secure HTTP (HTTPS) communications. If you are using a firewall, the port must be opened for CIMOM communications. Ensure that the CIMOM broker daemon is listening on port 5989. Log in as the `root` user on the cluster master node, open a terminal console, then enter the following at the command prompt:

```
netstat -an |grep -i5989
```

The Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a user name and password) between iManager and the `_admin` volume on the OES server that 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.

IMPORTANT: SFCB is automatically PAM-enabled for Linux User Management (LUM) as part of the OES installation. Users not enabled for LUM cannot use the CIM providers to manage OES. The user name that you use to log in to iManager when you manage a cluster and the BCC cluster must be an eDirectory user name that has been LUM-enabled.

For more information about the permissions and rights needed by the administrator user, see Section 4.1, "Cluster Administration Requirements," on page 35.

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

You can use the following commands to start, stop, or restart SFCB:

To perform this task	At a command prompt, enter as the `root` user
To start SFCB	`rcsblim-sfcb start` or `systemctl start sblim-sfcb.service`
To stop SFCB	`rcsblim-sfcb stop` or `systemctl stop sblim-sfcb.service`
To check SFCB status	`rcsblim-sfcb status` or `systemctl status sblim-sfcb.service`
To restart SFCB	`rcsblim-sfcb restart` or `systemctl restart sblim-sfcb.service`

For more information, see "Web Based Enterprise Management using SFCB" (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-wbem.html) in the *SUSE Linux Enterprise Server Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-administration.html).

4.6.9 OES Credential Store

OES Cluster Services requires OES Credential Store to be installed and running on each node in the cluster.

To check that OES Credential Store is running correctly, at the command prompt enter `oescredstore -l` as the `root` user.

4.6.10 Web Browser

For information about supported web browsers for iManager, see "System Requirements for iManager Server" in the *NetIQ iManager Installation Guide*.

The Clusters plug-in for iManager might not operate properly if the highest priority Language setting for your web browser is set to a language other than one of the supported languages in iManager. To view a list of supported languages and codes in iManager, select the **Preferences** tab, click **Language**. The language codes are Unicode (UTF-8) compliant.

To avoid display problems, in your web browser, select **Tools > Options > Languages**, and then set the first language preference in the list to a supported language. You must also ensure the Character Encoding setting for the browser is set to Unicode (UTF-8) or ISO 8859-1 (Western, Western European, West European).

- In a Mozilla browser, select **View > Character Encoding**, then select the supported character encoding setting.
- In an Internet Explorer browser, select **View > Encoding**, then select the supported character encoding setting.

4.7 Software Requirements for Cluster Resources

Ensure that your system meets the following software requirements for creating and managing storage cluster resources:

- Section 4.7.1, "NCP Server for Linux," on page 48
- Section 4.7.2, "Storage Services File System for Linux," on page 48
- Section 4.7.3, "LVM Volume Groups and Linux POSIX File Systems," on page 49
- Section 4.7.4, "NCP Volumes on Linux POSIX File Systems," on page 49
- Section 4.7.5, "Dynamic Storage Technology Shadow Volume Pairs," on page 49
- Section 4.7.6, "NCP File Access," on page 49
- Section 4.7.7, "OES CIFS," on page 50
- Section 4.7.8, "OES Domain Services for Windows," on page 50

4.7.1 NCP Server for Linux

NCP Server for Linux is required in order to create virtual server names (NCS:NCP Server objects) for cluster resources. This includes storage and service cluster resources. To install NCP Server, select the **NCP Server and Dynamic Storage Technology** option during the install.

NCP Server for Linux also allows you to provide authenticated access to data by using the OES Trustee model. The NCP Server component must be installed and running before you can cluster-enable the following storage resources:

- NSS pools and volumes
- NCP volumes on Linux POSIX file systems
- Dynamic Storage Technology shadow volume composed of a pair of NSS volumes
- Linux Logical Volume Manager volume groups that use an NCS:NCP Server object, such as those created by using the Logical Volume Manger (NLVM) commands or the NSS Management Utility (NSSMU)

WARNING: Cross-protocol file locking is required when using multiple protocols for data access on the same volume. This helps prevent possible data corruption that might occur from cross-protocol access to files. The NCP Cross-Protocol File Lock parameter is enabled by default when you install NCP Server. If you modify the Cross-Protocol File Lock parameter, you must modify the setting on all nodes in the cluster.

NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource. If a file is opened with multiple protocols when the migration or failover begins, the file should be closed and reopened after the migration or failover to acquire cross-protocol locks on the new node.

See "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2023: NCP Server for Linux Administration Guide*.

For information about configuring and managing NCP Server for Linux, see the *OES 2023: NCP Server for Linux Administration Guide*.

For information about creating and cluster-enabling NCP volumes on Linux POSIX file systems, see "Configuring NCP Volumes with OES Cluster Services" in the *OES 2023: NCP Server for Linux Administration Guide*.

4.7.2 Storage Services File System for Linux

Storage Services (NSS) file system on Linux provides the following capabilities used by OES Cluster Services:

- Initializing and sharing devices used for the SBD (split-brain detector) and for shared pools. See Section 4.8.2, "SBD Partitions," on page 51.
- Creating and cluster-enabling a shared pool. See Chapter 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 201.
- Creating and cluster-enabling a shared Linux Logical Volume Manager (LVM) volume group. See Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 275.

The NSS pool configuration and NCS pool cluster resource configuration provide integrated configuration options for the following advertising protocols:

- NetWare Core Protocol (NCP), which is selected by default and is mandatory for NSS. See "NCP Server for Linux" on page 48.
- OES CIFS. See "OES CIFS" on page 50.

4.7.3 LVM Volume Groups and Linux POSIX File Systems

OES Cluster Services supports creating shared cluster resources on Linux Logical Volume Manager (LVM) volume groups. You can configure Linux POSIX file systems on the LVM volume group, such as Ext3, XFS, and ReiserFS. LVM and Linux POSIX file systems are automatically installed as part of the OES installation.

After the cluster is configured, you can create LVM volume group cluster resources as described in Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 275.

NCP Server is required if you want to create a virtual server name (NCS:NCP `Server` object) for the cluster resource. You can add an NCP volume (an NCP share) on the Linux POSIX file system to give users NCP access to the data. See Section 4.7.1, "NCP Server for Linux," on page 48.

4.7.4 NCP Volumes on Linux POSIX File Systems

After you cluster-enable an LVM volume group, OES Cluster Services supports creating NCP volumes on the volume group's Linux POSIX file systems. NCP Server is required. See Section 4.7.1, "NCP Server for Linux," on page 48.

For information about creating and cluster-enabling NCP volumes, see "Configuring NCP Volumes with OES Cluster Services" in the *OES 2023: NCP Server for Linux Administration Guide*.

4.7.5 Dynamic Storage Technology Shadow Volume Pairs

OES Cluster Services supports clustering for OES Dynamic Storage Technology (DST) shadow volume pairs on OES 11 and later. DST is installed automatically when you install NCP Server for Linux. To use cluster-enabled DST volume pairs, select the **NCP Server and Dynamic Storage Technology** option during the install.

For information about creating and cluster-enabling Dynamic Storage Technology volumes on Linux, see "Configuring DST Shadow Volume Pairs with OES Cluster Services" in the *OES 2023: Dynamic Storage Technology Administration Guide*.

4.7.6 NCP File Access

OES Cluster Services requires NCP file access to be enabled for cluster-enabled NSS volumes, NCP volumes, and DST volumes, even if users do not access files via NCP. This is required to support access control via the OES Trustee model. See Section 4.7.1, "NCP Server for Linux," on page 48.

4.7.7 OES CIFS

OES Cluster Services supports using OES CIFS as an advertising protocol for cluster-enabled NSS pools and volumes.

OES CIFS is not required to be installed when you install OES Cluster Services, but it must be installed and running in order for the **CIFS Virtual Server Name** field and the **CIFS** check box to be available. Select the check box to enable OES CIFS as an advertising protocol for the NSS pool cluster resource. A default **CIFS Virtual Server Name** is suggested, but you can modify it. The CIFS daemon should also be running before you bring resources online that have CIFS enabled.

To install OES CIFS, select the **OES CIFS** option from the OES Services list during the install. For information about configuring and managing the OES CIFS service, see the *OES 2023: OES CIFS for Linux Administration Guide*.

4.7.8 OES Domain Services for Windows

OES Cluster Services supports using clusters in Domain Services for Windows (DSfW) contexts. If Domain Services for Windows is installed in the eDirectory tree, the nodes in a given cluster can be in the same or different DSfW subdomains. Port 1636 is used for DSfW communications. This port must be opened in the firewall.

For information using Domain Services for Windows, see the *OES 2023: Domain Services for Windows Administration Guide*.

4.8 Shared Disk Configuration Requirements

A shared disk subsystem is required for a cluster in order to make data highly available. The OES Cluster Services software must be installed in order to be able to mark devices as shareable, such as the devices you use for clustered pools and the device you use for the SBD (split-brain detector) during the cluster configuration.

Ensure that your shared storage devices meet the following requirements:

- Section 4.8.1, “Shared Devices,” on page 50
- Section 4.8.2, “SBD Partitions,” on page 51
- Section 4.8.3, “Shared iSCSI Devices,” on page 53
- Section 4.8.4, “Shared RAID Devices,” on page 53

4.8.1 Shared Devices

OES Cluster Services supports the following shared disks:

- Fibre Channel LUN (logical unit number) devices in a storage array
- iSCSI LUN devices
- SCSI disks (shared external drive arrays)

Before configuring OES Cluster Services, the shared disk system must be properly set up and functional according to the manufacturer's instructions.

Prior to installation, verify that all the drives in your shared disk system are recognized by Linux by viewing a list of the devices on each server that you intend to add to your cluster. If any of the drives in the shared disk system do not show up in the list, consult the OES documentation or the shared disk system documentation for troubleshooting information.

Prepare the device for use in a cluster resource:

- **NSS Pool:** For new devices, you must initialize and share the device before creating the pool. For an existing pool that you want to cluster-enable, use NSSMU or iManager to share the device.

  All devices that contribute space to a clustered pool must be able to fail over with the pool cluster resource. You must use the device exclusively for the clustered pool; do not use space on it for other pools or for Linux volumes. A device must be marked as Shareable for Clustering before you can use it to create or expand a clustered pool.
- **Linux LVM volume group:** For new devices, use an unpartitioned device that has been initialized. Do not mark the device as shared because doing so creates a small partition on it. LVM uses the entire device for the volume group. For an existing volume group, no not mark the device as shared.

If this is a new cluster, connect the shared disk system to the first server so that the SBD cluster partition can be created during the Cluster Services install. See Section 4.8.2, “SBD Partitions,” on page 51.

4.8.2 SBD Partitions

If your cluster uses physically shared storage resources, you must create an SBD (split-brain detector) partition for the cluster. You can create an SBD partition in YaST as part of the first node setup, or by using the SBD Utility (sbdutil) before you add a second node to the cluster. Both the YaST new cluster setup and the SBD Utility (`sbdutil`) support mirroring the SBD partition.

An SBD must be created before you attempt to create storage objects like pools or volumes for file system cluster resources, and before you configure a second node in the cluster. NLVM and other NSS management tools need the SBD to detect if a node is a member of the cluster and to get exclusive locks on physically shared storage.

For information about how SBD partitions work and how to create an SBD partition for an existing cluster, see Section 9.18, “Creating or Deleting Cluster SBD Partitions,” on page 146.

- “Preparing the SAN Devices for the SBD” on page 52
- “Initializing and Sharing a Device for the SBD” on page 52
- “Determining the SBD Partition Size” on page 52

Preparing the SAN Devices for the SBD

Use the SAN storage array software to carve a LUN to use exclusively for the SBD partition. For 512 byte device, you should have at least 20 MB of free available space and the minimum partition size is 8 MB. For 4096 (4Kn) device, you should have at least 80 MB of free available space and the minimum partition size is 64 MB. Connect the LUN device to all nodes in the cluster.

For device fault tolerance, you can mirror the SBD partition by specifying two devices when you create the SBD. Use the SAN storage array software to carve a second LUN of the same size to use as the mirror. Connect the LUN device to all nodes in the cluster.

The device you use to create the SBD partition can be a software RAID device or a hardware RAID device.

Initializing and Sharing a Device for the SBD

Before you use YaST to set up the a cluster, you must initialize each SAN device that you created for the SBD, and mark each as Shareable for Clustering.

IMPORTANT: The OES Cluster Services software must already be installed in order to be able to mark the devices as shareable.

After you install OES Cluster Services, but before you configure the cluster, you can initialize a device and set it to a shared state by using NSSMU, the Storage plug-in for iManager, OES Linux Volume Manager (NLVM) commands, or an NSS utility called `ncsinit`.

If you configure a cluster before you create an SBD, NSS tools cannot detect if the node is a member of the cluster and cannot get exclusive locks to the physically shared storage. In this state, you must use the `-s` NLVM option with the `nlvm init` command to override the shared locking requirement and force NLVM to execute the command. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

When you mark the device as Shareable for Clustering, share information is added to the disk in a free-space partition that is about 4 MB in size. This space becomes part of the SBD partition.

Determining the SBD Partition Size

When you configure a new cluster, you can specify how much free space to use for the SBD, or you can specify the **Use Maximum Size** option to use the entire device. If you specify a second device to use as a mirror for the SBD, the same amount of space is used. If you specify to use the maximum size and the mirror device is bigger than the SBD device, you will not be able to use the excess free space on the mirror for other purposes.

Because an SBD partition ends on a cylinder boundary, the partition size might be slightly smaller than the size you specify. When you use an entire device for the SBD partition, you can use the **Use Maximum Size** option, and let the software determine the size of the partition.

4.8.3 Shared iSCSI Devices

If you are using iSCSI for shared disk system access, ensure that you have installed and configured the iSCSI initiators and targets (LUNs) and that they are working properly. The iSCSI target devices must be mounted on the server before the cluster resources are brought online.

4.8.4 Shared RAID Devices

We recommend that you use hardware RAID in the shared disk subsystem to add fault tolerance to the shared disk system.

Consider the following when using software RAIDs:

- NSS software RAID is supported for shared disks for NSS pools. Any RAID0/5 device that is used for a clustered pool must contribute space exclusively to that pool; it cannot be used for other pools. This allows the device to fail over between nodes with the pool cluster resource. Ensure that its component devices are marked as Shareable for Clustering before you use a RAID0/5 device to create or expand a clustered pool
- Linux software RAID can be used in shared disk configurations that do not require the RAID to be concurrently active on multiple nodes. Linux software RAID cannot be used underneath clustered file systems (such as OCFS2, GFS, and CXFS) because OES Cluster Services does not support concurrent activation.

  WARNING: Activating Linux software RAID devices concurrently on multiple nodes can result in data corruption or inconsistencies.

4.9 SAN Rules for LUN Masking

When you create a OES Cluster Services system that uses shared storage space, it is important to remember that all of the servers that you grant access to the shared device, whether in the cluster or not, have access to all of the volumes on the shared storage space unless you specifically prevent such access. OES Cluster Services arbitrates access to shared volumes for all cluster nodes, but cannot protect shared volumes from being corrupted by non-cluster servers.

LUN masking is the ability to exclusively assign each LUN to one or more host connections. With it you can assign appropriately sized pieces of storage from a common storage pool to various servers. See your storage system vendor documentation for more information on configuring LUN masking.

Software included with your storage system can be used to mask LUNs or to provide zoning configuration of the SAN fabric to prevent shared volumes from being corrupted by non-cluster servers.

IMPORTANT: We recommend that you implement LUN masking in your cluster for data protection. LUN masking is provided by your storage system vendor.

4.10 Multipath I/O Configuration Requirements

If you use shared devices with multipath I/O capability, ensure that your setup meets the requirements in this section.

- Section 4.10.1, “Path Failover Settings for Device Mapper Multipath,” on page 54
- Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 54
- Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 55
- Section 4.10.4, “Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS,” on page 56

4.10.1 Path Failover Settings for Device Mapper Multipath

When you use Device Mapper Multipath (DM-MP) with OES Cluster Services, ensure that you set the path failover settings so that the paths fail when path I/O errors occur.

The default setting in DM-MP is to queue I/O if one or more HBA paths is lost. OES Cluster Services does not migrate resources from a node set to the Queue mode because of data corruption issues that can be caused by double mounts if the HBA path is recovered before a reboot.

IMPORTANT: The HBAs must be set to Failed mode so that OES Cluster Services can automatically fail over storage resources if a disk paths go down.

Change the Retry setting in the `/etc/modprobe.d/99-local.conf` and `/etc/multipath.conf` files so that DM-MP works correctly with OES Cluster Services. See Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 54 and Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 55.

Also consider changes as needed for the retry settings in the HBA BIOS. See Section 4.10.4, “Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS,” on page 56.

4.10.2 Modifying the Port Down Retry Setting in the modprobe.conf.local File

The port_down_retry setting specifies the number of times to attempt to reconnect to a port if it is down when using multipath I/O in a cluster. Ensure that you have installed the latest HBA drivers from your HBA vendor. Refer to the HBA vendor’s documentation to understand the preferred settings for the device, then make any changes in the `/etc/modprobe.conf.local` file.

For example, for QLogic HBAs, ensure that you have installed the latest `qla-driver`. Ensure that you verify the vendor’s preferred settings before making the changes.

4.10.3 Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File

The goal of multipath I/O is to provide connectivity fault tolerance between the storage system and the server. When you configure multipath I/O for a stand-alone server, the retry setting protects the server operating system from receiving I/O errors as long as possible. It queues messages until a multipath failover occurs and provides a healthy connection. However, when connectivity errors occur for a cluster node, you want to report the I/O failure in order to trigger the resource failover instead of waiting for a multipath failover to be resolved. In cluster environments, you must modify the retry setting so that the cluster node receives an I/O error in relation to the cluster SBD verification process (recommended to be 50% of the heartbeat tolerance) if the connection is lost to the storage system. In addition, you want the multipath I/O fail back to be set to manual in order to avoid a ping-pong of resources because of path failures.

Use the guidance in the following sections to configure the polling interval, no path retry and failback settings in the `/etc/multipath.conf` file:

- "Polling Interval" on page 55
- "No Path Retry" on page 55
- "Failback" on page 56
- "Example of Multipath I/O Settings" on page 56

Polling Interval

The polling interval for multipath I/O defines the interval of time in seconds between the end of one path checking cycle and the beginning of the next path checking cycle. The default interval is 5 seconds. An SBD partition has I/O every 4 seconds by default. A multipath check for the SBD partition is more useful if the multipath polling interval value is 4 seconds or less.

IMPORTANT: Ensure that you verify the `polling_interval` setting with your storage system vendor. Different storage systems can require different settings.

No Path Retry

We recommend a retry setting of "fail" or "0" in the `/etc/multipath.conf` file when working in a cluster. This causes the resources to fail over when the connection is lost to storage. Otherwise, the messages queue and the resource failover cannot occur.

IMPORTANT: Ensure that you verify the retry settings with your storage system vendor. Different storage systems can require different settings.

```
features "0"
no_path_retry fail
```

The value `fail` is the same as a setting value of 0.

Failback

We recommend a `failback` setting of "`manual`" for multipath I/O in cluster environments in order to prevent multipath failover ping-pong.

```
failback "manual"
```

IMPORTANT: Ensure that you verify the failback setting with your storage system vendor. Different storage systems can require different settings.

Example of Multipath I/O Settings

For example, the following code shows the default `polling_interval`, `no_path_retry`, and `failback` commands as they appear in the `/etc/multipath.conf` file for EMC storage:

```
defaults
{
    polling_interval    5
#    no_path_retry        0
    user_friendly_names     yes
    features 0
}

devices {
  device {
        vendor "DGC"
        product ".*"
        product_blacklist "LUNZ"
        path_grouping_policy "group_by_prio"
        path_checker "emc_clariion"
        features "0"
        hardware_handler "1 emc"
        prio "emc"
        failback "manual"
        no_path_retry fail    #Set MP for failed I/O mode, any other non-zero values sets
the HBAs for Blocked I/O mode
  }
}
```

For information about configuring the `multipath.conf` file, see "Managing Multipath I/O for Devices" (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-multipath.html) in the *SLES 12 Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html).

4.10.4 Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS

In the HBA BIOS, the default settings for the **Port Down Retry** and **Link Down Retry** values are typically set too high for a cluster environment. For example, there might be a delay of more than 30 seconds after a fault occurs before I/O resumes on the remaining HBAs. Reduce the delay time for the HBA retry so that its timing is compatible with the other timeout settings in your cluster.

For example, you can change the **Port Down Retry** and **Link Down Retry** settings to 5 seconds in the QLogic HBA BIOS:

```
Port Down Retry=5
Link Down Retry=5
```

5 Installing, Configuring, and Repairing OES Cluster Services

This section describes how to install the OES Cluster Services software on Open Enterprise Server 2023 servers, how to configure the cluster on the first node, and how to configure other nodes for the cluster.

IMPORTANT: Before you install or configure OES Cluster Services, ensure that you understand the requirements for it and have configured the environment as described in Chapter 4, “Planning for OES Cluster Services,” on page 35.

See the following resources for information about rolling cluster upgrades or conversions to the latest version of OES Cluster Services:

Upgrading from	See:
OES 11 SP3 OES 2015 SP1 OES 2018 OES 2018 SP1	Chapter 7, “Upgrading OES Clusters,” on page 89
NetWare 6.5 SP8	*OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*

- Section 5.1, “OES Cluster Services Licensing,” on page 58
- Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,” on page 58
- Section 5.3, “Assigning Install Rights for Container Administrators or Non-Administrator Users,” on page 60
- Section 5.4, “Installing OES Cluster Services,” on page 61
- Section 5.5, “Configuring OES Cluster Services,” on page 64
- Section 5.6, “Configuring Additional Administrators,” on page 71
- Section 5.7, “Installing or Updating the Clusters Plug-in for iManager,” on page 72
- Section 5.8, “Patching OES Cluster Services,” on page 74
- Section 5.9, “Removing the NCS Configuration Information from a Node,” on page 75
- Section 5.10, “Removing a Node from a Cluster,” on page 77
- Section 5.11, “Adding a Single Node That Was Previously in a Cluster,” on page 78
- Section 5.12, “Adding Multiple Nodes That Were Previously in a Cluster,” on page 79
- Section 5.13, “Adding a Node After Another Node is Permanently Removed from the Cluster,” on page 80
- Section 5.14, “What’s Next,” on page 80

5.1 OES Cluster Services Licensing

OES Cluster Services supports up to 32 nodes in a single cluster. You receive a OES Cluster Services entitlement that covers an unlimited number of two-node clusters. To add nodes to a two-node cluster, you can purchase a paper license for them for an additional fee. For more information, see the OES Cluster Services for Open Enterprise Server How-to-Buy website (http://www.novell.com/products/openenterpriseserver/ncs/howtobuy.html).

5.2 Extending the eDirectory Schema to Add Cluster Objects

The first time that you install OES Cluster Services in an eDirectory tree, the directory schema for the tree is extended to include the Cluster object container and the following types of objects in it:

- Cluster Node objects
- Cluster Resource objects
- Cluster Template objects
- Volume Resource objects

A tree administrator user with the eDirectory credentials to do so can extend the eDirectory schema before a cluster is installed anywhere in a tree. This allows container administrators (or non-administrator users) to install a cluster in a container in that same tree without needing full administrator rights for the tree. After the schema has been extended, you must assign some eDirectory rights to the container administrators (or non-administrator users) who will install OES Cluster Services clusters.

If the schema is not extended separately, the installer of the first cluster server in the tree must be an administrator with credentials to extend the eDirectory schema. The schema is automatically extended during the install. Subsequent cluster servers can be installed by container administrators (or non-administrator users) with sufficient rights to install OES Cluster Services.

IMPORTANT: For information about the eDirectory rights needed to install OES Cluster Services in a tree after the schema has been extended, see Section 5.3, “Assigning Install Rights for Container Administrators or Non-Administrator Users,” on page 60.

See the following sections for information about extending the schema before you install OES Cluster Services in a tree.

- Section 5.2.1, “Prerequisites for Extending the Schema,” on page 58
- Section 5.2.2, “Extending the Schema,” on page 59

5.2.1 Prerequisites for Extending the Schema

This procedure assumes that no clusters currently exist in the tree, and the schema needs to be extended for cluster objects.

You need the tree administrator credentials for extending the eDirectory schema.

You need the following information about the tree where you want to install OES Cluster Services clusters:

Table 5-1 *Tree Information Needed for the Schema Expansion*

Parameter	Description	Example
port_num	The port number you assigned for eDirectory communications in the tree where you plan to install clusters. The default port is 636.	`636`
admin_username	The typeful fully distinguished user name of the administrator who has the eDirectory rights needed to extend the schema.	`cn=admin,o=example`
admin_password	The password of the administrator user.	`password`
server_ip_address	The IP address of the eDirectory server that contains the schema files.	`10.10.10.102`

5.2.2 Extending the Schema

You need to extend the schema only one time in the tree where you will be installing clusters.

IMPORTANT: It is not necessary to extend the schema separately from the OES Cluster Services installation if the installer of the first cluster server in the tree has the eDirectory rights necessary to change the schema, because the schema can be automatically extended during the install.

To extend the schema before OES Cluster Services is installed in the tree, the tree administrator user can expand the schema as follows:

1. As a user with rights to modify the eDirectory schema, log in as the `root` user, then open a terminal console.
2. In a text editor, create a text file, specify the configuration information for the OES Cluster Services cluster in it, then save the file.

   The following lines are an example of the content of the file, with sample values. The directives are self-explanatory.

   IMPORTANT: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

   ```
   CONFIG_NCS_LDAP_IP="10.1.1.102"
   CONFIG_NCS_LDAP_PORT="636"
   CONFIG_NCS_ADMIN_DN="cn=admin.o=context"
   CONFIG_NCS_ADMIN_PASSWORD="password"
   ```

3. As the `root` user, enter the following command at a command prompt:

   ```
   mkdir -p /var/opt/novell/install
   ```

4. As the `root` user, enter the following command at a command prompt:

   ```
   /opt/novell/ncs/install/ncs_install.py -e -f configuration_filename
   ```

Replace *configuration_filename* with the actual name of the file that you created in Step 2, such as

```
/opt/novell/ncs/install/ncs_install.py -e -f /root/Desktop/my_schema_info.txt
```

5 Delete the configuration file (*configuration_filename*) that you created.

This file contains a password in clear text. For security reasons, ensure that you delete the file when you are done.

6 Continue with Section 5.3, “Assigning Install Rights for Container Administrators or Non-Administrator Users,” on page 60.

5.3 Assigning Install Rights for Container Administrators or Non-Administrator Users

If the eDirectory schema has been extended in the tree where you want to create clusters, the container administrator or non-administrator user needs the following eDirectory rights to install and configure OES Cluster Services. These rights are also required if a different user configures OES Cluster Services after the install as described in Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 66.

- Attribute Modify rights on the NCP Server object of each node in the cluster.

  To set the Attribute Modify rights for the user on the nodes’ NCP Server objects:

  1. In iManager, select **Rights > Modify Trustees**.
  2. Select the NCP server object for the node, then click **Add Trustee**.
  3. For **Entry Rights**, set the **Browse** right.
  4. For **All Attributes Rights**, set the **Compare**, **Read**, and **Write** rights.
  5. Click **Apply** to save and apply your changes.
  6. Repeat Step 1 to Step 5 for the NCP Server object of each server that you plan to add to the cluster.

- Object Create rights on the container where the NCP Server objects are.

  To set the Object Create rights for the user on the container where the NCP Server objects are:

  1. In iManager, select **Rights > Modify Trustees**.
  2. Select the Container object, then click **Add Trustee**.
  3. For **Entry Rights**, set the **Browse**, **Create**, and **Rename** rights.
  4. For **All Attributes Rights**, set the **Compare**, **Read**, and **Write** rights.
  5. Click **Apply** to save and apply your changes.

- Object Create rights where the cluster container will be.

  This step is needed if the container for the Cluster object is different than the container for the NCP Server objects.

  To set the Object Create rights for the user on the container where the Cluster objects will be:

  1. In iManager, select **Rights > Modify Trustees**.
  2. Select the Container object, then click **Add Trustee**.
  3. For **Entry Rights**, set the **Browse**, **Create**, and **Rename** rights.

4. For **All Attributes Rights**, set the **Compare**, **Read**, and **Write** rights.
5. Click **Apply** to save and apply your changes.

For information about eDirectory rights, see "eDirectory Rights" in the *NetIQ eDirectory Administration Guide*.

5.4 Installing OES Cluster Services

OES Cluster Services for Linux is included on OES. It is necessary to install OES on every server that you want to add to a cluster. You can have up to 32 nodes in each cluster. See Section 5.1, "OES Cluster Services Licensing," on page 58.

Installing OES Cluster Services does the following:

- If the eDirectory schema has not already been extended for cluster objects in the tree, the schema is extended.
- OES Cluster Services software is installed on the server.

You can install OES Cluster Services when you install the OES platform, or you can install it later on an existing OES server.

- Section 5.4.1, "Before You Install OES Cluster Services," on page 61
- Section 5.4.2, "Installing OES Cluster Services during an OES Installation," on page 62
- Section 5.4.3, "Installing OES Cluster Services on an Existing OES Server," on page 63

5.4.1 Before You Install OES Cluster Services

Ensure that your system meets the requirements and guidelines in Chapter 4, "Planning for OES Cluster Services," on page 35. Before you install OES Cluster Services, verify that your setup meets these requirements:

- "Using a Local eDirectory Database" on page 61
- "Rights to Extend the Schema" on page 61
- "Shared Devices for the SBD" on page 62
- "SBD Required before You Install the Second Node in the Cluster" on page 62
- "Using Remote eDirectory Server" on page 62

Using a Local eDirectory Database

If you want OES Cluster Services to use a local eDirectory database on the same server, you must install and configure eDirectory before you configure OES Cluster Services.

Rights to Extend the Schema

If the eDirectory schema was not previously extended in the tree as described in Section A.3, "OES Installation extend_schema Command," on page 379, the user who installs OES Cluster Services must be an administrator user with the rights to extend the schema, such as the tree administrator.

Shared Devices for the SBD

If you have a shared disk system and the server where you are installing OES Cluster Services is the first node in a cluster, install the OES Cluster Services software, then initialize the device you need for the Split Brain Detector (SBD) partition and mark it as shareable. If you plan to mirror the SBD, prepare two devices of the same size. You can create the SBD when you configure Cluster Services on the first node, or you can use the `sbdutil` utility to create it before you install Cluster Services on a second node in the cluster.

SBD Required before You Install the Second Node in the Cluster

If you have a shared disk system and the server where you are installing Cluster Services is the second node in a cluster, verify that a cluster partition for the cluster's SBD exists on the first cluster node before you begin the install on the second node.

IMPORTANT: The cluster SBD partition is not required unless you have shared storage.

Typically, you create the SBD when you configure the cluster on the first node as described in Section 5.5.5, "Configuring a New Cluster," on page 67. A one-node cluster that has shared disk storage can be configured without an SBD, but the SBD must be created before you add another node. For information about manually creating an SBD, see:

- Section 9.18.3, "Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL," on page 149
- Section 9.18.5, "Creating a Mirrored Cluster SBD Partition with SBDUTIL," on page 154

Using Remote eDirectory Server

If eDirectory is not installed locally and you are using remote eDirectory server, then copy the eDirectory CA certificate from the remote eDirectory server `/etc/opt/novell/certs/SSCert.pem` to `/etc/pki/trust/anchors/` and run `/usr/sbin/update-ca-certificates`.

5.4.2 Installing OES Cluster Services during an OES Installation

This section describes only those steps in the install that are directly related to installing OES Cluster Services. For detailed instructions on installing OES, see the *OES 2023: Installation Guide*.

Repeat the following procedure for each server that you want to add to the cluster:

1. From the DVD boot menu, select **Installation**, click **Next**, then continue through the OES install until you get to the Installation Settings page.
2. On the Installation Settings page, click **Software** to open the Software Selection and System Tasks page.
3. Under **Open Enterprise Server**, select **OES Cluster Services** and any other OES components that you want to install, then click **Accept**.

   When you select OES Cluster Services, the following basic services for managing OES are automatically selected:

   - OES Backup / Storage Management

- OES Linux User Management
- OES Remote Manager

The following Open Enterprise Server are not automatically selected, but are required for managing and configuring OES Cluster Services:

- iManager must be installed on at least one server in the same tree.
- eDirectory must already be installed on at least one server in the tree where you are installing the cluster. You can install a replica on the cluster server.

Select other protocols and services according to your planned setup. See Section 4.7, “Software Requirements for Cluster Resources,” on page 47.

IMPORTANT: If you deselect a pattern after selecting it, you are instructing the installation program to not install that pattern and all of its dependent patterns. Rather than deselecting a pattern, click **Cancel** to cancel your software selections, then click the **Software** heading again to choose your selections again.

Selecting only the patterns that you want to install ensures that the patterns and their dependent patterns and packages are installed.

If you click **Accept**, then return to software pattern selection page, the selections that you made become your base selections and must be deselected if you want to remove them from the installation proposal.

---

4. Continue through the installation process, but do not configure OES Cluster Services when you reach the Open Enterprise Server Configuration page. You configure it later.
5. After the install, use the Software Updater (or other update methods) to download and apply any updates that are in the SLES and OES Updates channels.

   See “Updating (Patching) an OES Server” in the *OES 2023: Installation Guide*.
6. Ensure that your SAN storage is working properly for the server.
7. Continue with Section 5.5, “Configuring OES Cluster Services,” on page 64.

5.4.3 Installing OES Cluster Services on an Existing OES Server

If you did not install OES Cluster Services during the OES 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. On the Software Selection page under **Open Enterprise Server**, select **OES Cluster Services** and any other compatible OES components that you want to install.

   Services that you have already installed are indicated by a blue check mark in the status check box next to the service.

   For information about the options, see Step 3 in Section 5.4.2, “Installing OES Cluster Services during an OES Installation,” on page 62.
4. Click **Accept** to begin the install, then click **Continue** to accept changed packages.
5. Continue through the installation process, but do not configure OES Cluster Services when you reach the Open Enterprise Server Configuration page. You configure it later.

6 After the install, use the Software Updater (or other update methods) to download and apply any updates that are in the SLES and OES Updates channels.

   See "Updating (Patching) an OES Server" in the *OES 2023: Installation Guide*.

7 Ensure that your SAN storage is working properly for the server.

8 Continue with Section 5.5, "Configuring OES Cluster Services," on page 64.

5.5 Configuring OES Cluster Services

After you install OES Cluster Services, you use the **Open Enterprise Server > OES Install and Configuration** tool in YaST to set up the cluster or to add a node to the cluster.

IMPORTANT: The YaST-based configuration is not used to modify the settings for an existing cluster. For information about modifying the settings for an existing cluster, see Section 8.13, "Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster," on page 120.

If you are creating a new cluster, the OES Cluster Services configuration does the following:

- Creates a new Cluster object and a Cluster Node object in eDirectory.
- Creates a special cluster partition for the Split Brain Detector (SBD) if you have a shared disk system.

If you are adding a server to an existing cluster, the OES Cluster Services configuration does the following:

- Creates a new Cluster Node object in eDirectory.

Use the following procedures to configure OES Cluster Services on a node:

- Section 5.5.1, "Initializing and Sharing a Device to Use for the SBD Partition," on page 64
- Section 5.5.2, "Opening the Open Enterprise Server Configuration Page," on page 65
- Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 66
- Section 5.5.4, "Accessing the OES Cluster Services Configuration Page in YaST," on page 66
- Section 5.5.5, "Configuring a New Cluster," on page 67
- Section 5.5.6, "Adding a Node to an Existing Cluster," on page 69

5.5.1 Initializing and Sharing a Device to Use for the SBD Partition

If you plan to use shared storage in the cluster, the cluster needs a split-brain detector (SBD). You must create the SBD before you add a second node to the cluster. You can create the SBD when you set up the cluster on the first node. In preparation, you must initialize the partition, and mark its

device as shareable for clustering. If you plan to mirror the SBD, you must also initialize the partition that you want to use as the mirrored SBD, and mark its device as shareable for clustering. For a 512 byte device, you need at least 20 MB for the SBD. For a 4096 (4Kn) device, it is 80 MB.

1. Log in as the `root` user on the server that will be the first node of the cluster, then open a terminal console.
2. At the command prompt, enter

   ```
   nssmu
   ```

3. In the NSSMU, select **Devices** and press Enter.
4. In the **Devices** list, select the device that you want to use for the SBD.
5. If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3, then press Y (Yes) to confirm and continue.

   WARNING: Initializing a disk destroys all of the data on it.

6. Select the DOS or GPT partitioning scheme for the device, then press Enter.

   DOS supports devices up to 2 TB in size. GPT supports devices of any size.

   Wait for the page to refresh before continuing.
7. Press F6 to mark the device as shareable for clustering.

   The **Shareable for Clustering** value changes from **No** to **Yes**.
8. If you plan to mirror the SBD, repeat Step 4 through Step 7 for the second device.
9. Exit NSSMU.
10. Continue with Section 5.5.2, “Opening the Open Enterprise Server Configuration Page,” on page 65.

5.5.2 Opening the Open Enterprise Server Configuration Page

1. Log in to the server as the `root` user.
2. Verify that the OES 2023 Media is mounted on or available to the server.
3. In YaST, select **Open Enterprise Server > OES Install and Configuration**.
4. On the Software Selection page under **Open Enterprise Server**, verify that the **OES Cluster Services** option is already installed as indicated by a check mark.
5. Click **Accept** to proceed to the Open Enterprise Server Configuration page.
6. Do one of the following:
   - **Same Administrator:** To use the same administrator credentials that were used to install OES Cluster Services, continue with Section 5.5.4, “Accessing the OES Cluster Services Configuration Page in YaST,” on page 66.
   - **Different Administrator:** To use different administrator credentials than those used to install OES Cluster Services, continue with Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 66.

5.5.3 Using Different LDAP Credentials for the Cluster Configuration

You can use different user credentials to configure OES Cluster Services than were used during the installation of Open Enterprise Server on the server by reconfiguring the settings for the **LDAP Configuration of Open Enterprise Services** option.

For information about what rights are needed, see Section 5.3, “Assigning Install Rights for Container Administrators or Non-Administrator Users,” on page 60.

1. On the Open Enterprise Server Configuration page under **LDAP Configuration of Open Enterprise Services**, click the **disabled** link to enable re-configuration.

   The sentence changes to **Reconfiguration is enabled**.
2. Click the **LDAP Configuration of Open Enterprise Services** link to open the LDAP Configuration page.
3. Specify the following values:
   - **Admin name and context:** The user name and context (in LDAP form) of the container administrator user (or non-administrator user) who has the eDirectory rights needed to install OES Cluster Services.
   - **Admin password:** The password of the container administrator (or non-administrator user).
4. Click **Next**.

   The install returns to the Open Enterprise Server Configuration page.
5. Continue with Section 5.5.4, “Accessing the OES Cluster Services Configuration Page in YaST,” on page 66.

5.5.4 Accessing the OES Cluster Services Configuration Page in YaST

1. On the Open Enterprise Server Configuration page under **OES Cluster Services**, locate the OES Cluster Services link.

   The configuration is currently disabled.
2. Click the **disabled** link to enable configuration.

   The sentence changes to **Configure is enabled**.
3. Click the **OES Cluster Services** link to open the OES Cluster Services Configuration page.
4. If you are prompted for credentials, specify the password of the specified Administrator user, then click **OK**.

   This is either the LDAP administrator identity that was used when the Open Enterprise Server were installed, or the identity configured after the install in Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 66.
5. On the OES Cluster Services Configuration page, continue with one of the following:
   - Section 5.5.5, “Configuring a New Cluster,” on page 67
   - Section 5.5.6, “Adding a Node to an Existing Cluster,” on page 69

5.5.5 Configuring a New Cluster

Perform the following configuration for the first node that you configure for a cluster:

1. Go to the OES Cluster Services Configuration page as described in Section 5.5.4, “Accessing the OES Cluster Services Configuration Page in YaST,” on page 66.
2. On the first configuration page, specify the following settings for a new cluster, then click **Next**:

Parameter	Action
New or existing cluster	Select **New Cluster**.
Directory Server Address	Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default. The IP addresses shown are the LDAP servers available for this service to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.13.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 120.
Cluster FDN	Click **Browse** and navigate the tree to select the container where you want to create the cluster, then click **OK**. The FDN is automatically added to the field with a suggested cluster name. You can specify a different cluster name. You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots. For example: `cn=clus1,ou=ncs,o=mycompany` You must specify an existing context. Specifying a new context does not create a new context. Cluster names must be unique. You cannot create two clusters with the same name in the same eDirectory tree.
Cluster IP Address	The cluster IP address represents the cluster on the network. It is different than the server IP address. For example: `10.10.10.44` The cluster IP address provides a single point for cluster access, configuration, and management. A Master IP Address resource is created automatically during the Cluster Services installation. The cluster IP address is bound to the master node and remains with the master node regardless of which server is the master node. The cluster IP address is required to be on the same IP subnet as the other servers in the same cluster, which is necessary for certain external network management programs to get cluster status alerts.

Parameter	Action
Select the storage device with shared media	Select the initialized and shared device that you want to use for the SBD from the drop-down list. For a new cluster, you can create the SBD partition now, or you can create it manually at any time before you add a second node to the cluster. A split-brain detector (SBD) is required if you plan to use shared disks in the cluster. For information about SBD requirements, see Section 4.8.2, "SBD Partitions," on page 51. The drop-down menus show only devices that have been initialized and shared. See Section 5.5.1, "Initializing and Sharing a Device to Use for the SBD Partition," on page 64.
Select the device where the mirror partition will be created	If you want to mirror the SBD, select a different shared device from the drop-down list.
Desired partition size of shared media	Specify a size of 20 MB or greater for 512 byte device and 80 MB or greater for 4096 (4Kn) device to use for the SBD partition, or select **Use maximum size** to use the entire device. If you selected a device for the SBD mirror, the specified size is also used for the mirror segment.

3 On the Proxy User Configuration page, specify one of the following users as the NCS Proxy user, then click **Next**.

Parameter	Action
OES Common Proxy User	If the OES Common Proxy User is enabled in eDirectory, the **Use OES Common Proxy User** check box is automatically selected and the **NCS Proxy User Name** and **Specify NCS Proxy User Password** fields are populated with the credentials of the OES Common Proxy User.
LDAP Admin User	If the OES Common Proxy User is disabled in eDirectory, the **Use OES Common Proxy User** check box is automatically deselected and the **NCS Proxy User Name** and **Specify NCS Proxy User Password** fields are populated with the credentials of the LDAP Admin user. The fields are also automatically populated with the LDAP Admin credentials if you deselect the **Use OES Common Proxy User** check box.
Another Admin User	Deselect the **Use OES Common Proxy User** check box, then specify the credentials of an existing administrator user.

You can reset the default settings by clicking **Back** to return to the OES Cluster Services Configuration page, then clicking **Next** to continue again to the Proxy User Configuration page.

4 On the third Configuration page, configure the following settings, then click **Finish**.

Parameter	Action
IP address of this node	From the drop-down list, select the IP address that OES Cluster Services will use for this node. Some servers have multiple IP addresses. This step lets you choose which IP address OES Cluster Services uses.
Start Clustering Services now	Select the check box to start OES Cluster Services software on this node after configuring it.

5 On the OES Server Configuration page, scroll down to the OES Cluster Services entry to review the summary of the Cluster Services configuration, then click **Next**.

6 Continue through the setup process, then click **Finish** to exit the OES Configuration

7 Start OES Cluster Services using one of these methods:

Setup Condition	Instructions
Start Cluster Services now was enabled	OES Cluster Services starts automatically after the configuration completes.

Start Cluster Services now was disabled	Start OES Cluster Services manually by using one of these methods:
	◆ Reboot the cluster server.
	◆ At a command prompt, enter one of the following as the `root` user:
	`rcnovell-ncs start`
	or
	`systemctl start novell-ncs.service`

8 Continue with Section 5.6, “Configuring Additional Administrators,” on page 71.

5.5.6 Adding a Node to an Existing Cluster

Perform the following configuration for each node that you add to an existing cluster:

1 If you have not previously configured the SBD partition for the cluster, create an SBD partition for the cluster by using one of the following procedures:

- Section 9.18.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 149
- Section 9.18.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 154

If you have a shared disk system attached to your cluster servers, an SBD partition is required and must be created before you configure the second node in the cluster.

IMPORTANT: An SBD partition requires at least 20 MB of free space on a 512 byte device and at least 80 MB of free space on a 4096 (4Kn) device that has been previously initialized and marked as shareable for clustering.

2 Go to the OES Cluster Services Configuration page as described in Section 5.5.4, “Accessing the OES Cluster Services Configuration Page in YaST,” on page 66.

3 Specify the following settings for adding this node to an existing cluster, then click **Next**:

Parameter	Action
New or existing cluster	Select **Existing Cluster**.
Directory Server Address	Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default. The IP addresses shown are the LDAP servers available for this service to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.13.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 120.
Cluster FDN	Click **Browse** and navigate the tree to select the Cluster object of the existing cluster, then click **OK**. The FDN is automatically added to the field. You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots. For example: `cn=clus1,ou=ncs,o=mycompany`
Select the device where the mirror partition will be created	If you want to mirror the SBD, select a different shared device from the drop-down list.

4 On the Proxy User Configuration page, specify one of the following users as the NCS Proxy user for this node, then click **Next**.

Parameter	Action
OES Common Proxy User	If the OES Common Proxy User is enabled in eDirectory, the **Use OES Common Proxy User** check box is automatically selected and the **NCS Proxy User Name** and **Specify NCS Proxy User Password** fields are populated with the credentials of the OES Common Proxy User.
LDAP Admin User	If the OES Common Proxy User is disabled in eDirectory, the **Use OES Common Proxy User** check box is automatically deselected and the **NCS Proxy User Name** and **Specify NCS Proxy User Password** fields are populated with the credentials of the LDAP Admin user. The fields are also automatically populated with the LDAP Admin credentials if you deselect the **Use OES Common Proxy User** check box.
Another Admin User	Deselect the **Use OES Common Proxy User** check box, then specify the credentials of an existing administrator user.

You can reset the default settings by clicking **Back** to return to the OES Cluster Services Configuration page, then clicking **Next** to continue again to the Proxy User Configuration page.

5 On the third Configuration page, configure the following settings, then click **Finish**.

Parameter	Action
IP address of this node	From the drop-down list, select the IP address that OES Cluster Services will use for this node. Some servers have multiple IP addresses. This step lets you choose which IP address OES Cluster Services uses.
Start Clustering Services now	Select the check box to start OES Cluster Services software on this node after configuring it.

6 On the OES Server Configuration page, scroll down to the OES Cluster Services entry to review the summary of the Cluster Services configuration, then click **Next**.

7 Continue through the setup process, then click **Finish** to exit the OES Configuration

8 Start OES Cluster Services by using one of these methods:

Setup Condition	Instructions
Start Cluster Services now was enabled	OES Cluster Services starts automatically after the configuration completes.

Start Cluster Services now was disabled	Start OES Cluster Services manually by using one of these methods:
	◆ Reboot the cluster server.
	◆ At a command prompt, enter the following as the `root` user:
	`rcnovell-ncs start`
	or
	`systemctl start novell-ncs.service`

9 Continue with Section 5.6, “Configuring Additional Administrators,” on page 71.

5.6 Configuring Additional Administrators

The Administrator user that you specify as the Proxy User during OES Cluster Services install process is automatically configured as the administrator for the cluster with the following setup:

- The user is a trustee and has the Supervisor right to the Server object of each server node in the cluster.
- The user is enabled for Linux with Linux User Management. This gives the user a Linux UID in addition to the users eDirectory GUID.
- The user is a member of a LUM-enabled administrator group associated with the servers in the cluster.

IMPORTANT: To allow other administrators (such as the tree administrator) to manage the cluster, the users' user names must be similarly configured.

You can modify the default administrator user name or password after the install by following the procedure in Section 8.13.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 120.

5.7 Installing or Updating the Clusters Plug-in for iManager

The Clusters plug-in for iManager was reorganized in OES 11 SP1. On the Roles and Tasks page under **Clusters**, the **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, and **Cluster Options** menu options have been replaced with two options: **My Clusters** and **My Resources**.

- **My Clusters:** The logged-in cluster administrator can set up a personalized list of clusters to manage, which allows their status to be viewed at-a-glance. The list persists between the administrator's logins to iManager in the same tree. See Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

  My Resources: The logged-in cluster administrator can set up a personalized list of resources to manage, which allows their status to be viewed at-a-glance. The list persists between the administrator's logins to iManager in the same tree. See Section 10.2, "Setting Up a Personalized List of Resources to Manage," on page 165.

Links on the these pages take you to the familiar cluster options presented as tabs: **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, and **Cluster Options**.

Use the information in the following sections to install or update the Clusters plug-in.

- Section 5.7.1, "Prerequisites for Installing or Updating the Clusters Plug-In," on page 72
- Section 5.7.2, "Installing the Clusters Plug-In," on page 73
- Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 73
- Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later," on page 74

5.7.1 Prerequisites for Installing or Updating the Clusters Plug-In

The Clusters plug-in supports the management of OES and NetWare clusters and resources. It cannot be installed on a NetWare server.

The Clusters plug-in requires the following components to be installed in iManager:

- Clusters (`ncsmgmt.rpm`)
- Common code for storage-related plug-ins (`storagemgmt.rpm`)

  See "Storage-Related Plug-Ins for iManager" on page 45.

5.7.2 Installing the Clusters Plug-In

If you are installing a new instance of iManager 2.7.5 or later and the Clusters plug-in on a computer, follow the instructions in "Downloading and Installing Plug-in Modules" in the *NetIQ iManager Administration Guide*. Ensure that you install the Clusters plug-in (`ncsmgmt.rpm`) as well as the common code for storage-related plug-ins (`storagemgmt.rpm`).

Verify that the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory contains the class loader file `commons-lang-2.6.jar` (or later version), then remove all earlier versions of the file. The Clusters plug-in for iManager 2.7.5 or later requires `commons-lang-2.6.jar` (or later version) of the class loader file. iManager uses one class loader for all of its plug-ins. If multiple versions of the class loader file exist in the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory, iManager loads whichever one comes first alphabetically. You must remove the earlier versions (such as `commons-lang.jar` or `commons-lang-2.4.jar`) to ensure that iManager loads the 2.6 or later version of the class loader file. Restart Tomcat when you are done.

To configure role-based services for the Clusters plug-in, see "RBS Configuration" in the *NetIQ iManager Administration Guide*.

5.7.3 Uninstalling and Reinstalling the Clusters Plug-In

Use the procedure in this section to update an instance of the Clusters plug-in, or to add the Clusters plug-in on a computer where other storage-related plug-ins are already installed.

1. In iManager, uninstall the currently installed storage-related plug-ins.
2. Copy the new version of the plug-in files to the iManager plug-ins location, manually overwriting the older version of the plug-in in the `packages` folder with the newer version of the plug-in.
3. In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.
4. If you are updating the Clusters plug-in, delete the contents of the `/var/opt/novell/tomcat6/work/Catalina/localhost/nps` work folder.

   You must manually delete the old compiled JSPs that are cached in the `nps` work folder. This ensures that only the latest files are cached.
5. Verify that the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory contains the class loader file `commons-lang-2.6.jar` (or later version), then remove all earlier versions of the file from the directory.

   The Clusters plug-in for iManager 2.7.5 or later requires `commons-lang-2.6.jar` (or later version) of the class loader file. iManager uses one class loader for all of its plug-ins. If multiple versions of the class loader file exist in the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory, iManager loads whichever one comes first alphabetically. You must remove the earlier versions (such as `commons-lang.jar` or `commons-lang-2.4.jar`) to ensure that iManager loads the 2.6 or later version of the class loader file.
6. Restart Tomcat by entering the following command, or reboot the server if a reboot is required by any of the patches applied.

   ```
   rcnovell-tomcat restart
   ```

   or

```
systemctl restart novell-tomcat.service
```

7 If you use Role-Based Services, continue with Section 5.7.4, “Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later,” on page 74.

5.7.4 Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later

In OES 11 SP2 or later, the Clusters plug-in replaces the **Clusters**, the **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, and **Cluster Options** menu options with the **My Clusters** and **My Resources** options. The plug-in is updated automatically to show new tasks, and hide old ones. However, the Role-Based Services role for Clusters is not updated automatically.

If you use Role-Based Services, you must reinstall the Clusters plug-in on the Role-Based Services Configuration page in order to pick up the modified menu options:

1 In iManager, click **Configuration**, then select **Role Based Services > RBS Configuration.**
2 Select the **iManager 2.x Collections** tab.
3 On the iManager 2.x Collections page, view the entry for **Role Based Services 2.novell**.

   There are numbers in the columns for **Modules**, **Installed**, **Out-of-Date**, and **Not-Installed**.
4 Click the number link in the **Out-of-Date** column for **Role Based Services 2.novell**.
5 Select the check box next to **Clusters**, then click **Update**.

   This updates the RBS configuration information for the Clusters plug-in.
6 Click **OK** to confirm the reinstallation.
7 After the reinstallation is complete, click **OK** to confirm the success message.
8 Log out of iManager.
9 Restart Tomcat by entering the following command, or reboot the server if a reboot is required by any of the patches applied.

   ```
   rcnovell-tomcat restart
   ```

   or

   ```
   systemctl restart novell-tomcat.service
   ```
10 Log in to iManager to use the updated Clusters plug-in.

5.8 Patching OES Cluster Services

You should stop OES Cluster Services when applying maintenance patches for OES Cluster Services.

1 Log in to the node as the `root` user, then open a terminal console.
2 Use the `cluster leave` command to remove the node from the cluster.
3 Stop OES Cluster Services by entering

   ```
   rcnovell-ncs stop
   ```

   or

```
systemctl stop novell-ncs.service
```

4 Apply the latest patches for OES and OES Cluster Services.

5 Start OES Cluster Services by entering

```
rcnovell-ncs start
```

or

```
systemctl start novell-ncs.service
```

5.9 Removing the NCS Configuration Information from a Node

You can use the `/opt/novell/ncs/install/ncs_install.py -r` script to remove OES Cluster Services configuration information from a node to return it to a pre-configured state. This is useful for the following scenarios:

- If the OES Cluster Services configuration is unsuccessful, you might have a partial configuration that prevents you from properly configuring the node for a cluster.
- You want to move a node from one cluster to another cluster.

If this is a one-node cluster and you created an SBD, you must delete the SBD partition (and possibly re-initialize if it is a standalone disk). If the SBD is mirrored, delete both the mirror and the SBD partition before you re-configure the cluster for the single-node.

After you remove the cluster configuration information, you can start over with the configuration of OES Cluster Services on the node.

1 Log in as the `root` user to the node that you want to manage, then open a terminal console.

2 In a text editor, create a text file, specify the credentials for the OES Cluster Services cluster in it, then save the file.

This file contains a password in clear text. For security reasons, ensure that you delete the file when you are done.

The following lines are an example of the content of a file, with sample values. The directives are self-explanatory.

IMPORTANT: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

```
CONFIG_NCS_CLUSTER_DN="cn=clus1,ou=ncs,o=novell"
CONFIG_NCS_ADMIN_DN="cn=admin,o=novell"
CONFIG_NCS_ADMIN_PASSWORD="password"
```

3 If the node is configured and working in a cluster, leave the cluster. At a command prompt, enter

```
cluster leave
```

This cluster migrates any resources running on the node to other nodes in the cluster. When the node has successfully left the cluster, the following message is displayed:

```
Leaving...
No longer a member of cluster cluster_name
```

4 Stop OES Cluster Services if it is running. At the command prompt, enter

```
rcnovell-ncs stop
```

or

```
systemctl stop novell-ncs.service
```

Wait until the service has been unloaded.

5 At the command prompt, enter

```
/opt/novell/ncs/install/ncs_install.py -r [-f
<configuration_file_name>]
```

Without the `-f` option or applicable variables in the configuration file, `/opt/novell/ncs/install/ncs_install.py` prompts you to enter needed information.

6 You can check the return code (`echo $?`) of the command to see if the command is successful (`0`). Detailed log information is appended to `/var/opt/novell/install/ncslog`.

```
07/16/13 13:19:49: *** Starting NCS Post-Install ***
07/16/13 13:19:49: config type is: Existing Cluster.
07/16/13 13:19:49: ldap ip is: 10.10.10.37.
07/16/13 13:19:49: ldap port is: 636.
07/16/13 13:19:49: ldap server list is: ldaps://10.10.10.37:636.
07/16/13 13:19:49: Admin dn is: cn=admin,o=novell.
07/16/13 13:19:49: serverName is: myserver.
07/16/13 13:19:49: serverDn is: cn=myserver,O=novell.
07/16/13 13:19:49: server IP is: 10.10.10.37.
07/16/13 13:19:49: cluster dn is: cn=cl134,ou=ncs,o=novell.
07/16/13 13:19:49: Extend schema flag is: False.
07/16/13 13:19:49: Create cluster object flag is: False.
07/16/13 13:19:49: Add node flag is: False.
07/16/13 13:19:49: Delete node flag is: True.
07/16/13 13:19:49: Create node name file flag is: False.
07/16/13 13:19:49: Create cluster lib config file flag is: False.
07/16/13 13:19:49: Create SBD flag is: False.
07/16/13 13:19:49: Start requested flag is: False.
07/16/13 13:19:49: DSFW Enabled is: False.
07/16/13 13:19:49: Set Proxy User: False.
07/16/13 13:19:49: Binding to LDAP server ...
07/16/13 13:19:49: Deleting server 'cn=myserver,O=novell' from cluster
'cn=cl134,ou=ncs,o=novell' ...
07/16/13 13:19:49: Server 'cn=myserver,O=novell' deleted from cluster
'cn=cl134,ou=ncs,o=novell'.
07/16/13 13:19:49: *** Ending NCS Post-Install ***
```

7 Configure OES Cluster Services on the node as described in Section 5.5, “Configuring OES Cluster Services,” on page 64.

5.10 Removing a Node from a Cluster

You can use the procedure in this section to permanently remove a cluster node from the cluster. You must remove information about the node, even if you will replace the old server with a new server that has the same node name and IP address.

To delete a cluster node from a cluster:

1 Log in as the `root` user to the node in the cluster that you want to remove, then enter the following at a command prompt:

```
cluster leave
```

When the node has successfully left the cluster, the following message is displayed:

```
No longer a member of cluster cluster_name
```

2 In a web browser, open iManager, then log in to the eDirectory tree that contains the cluster you want to manage.

IMPORTANT: Log in as an administrator user who has sufficient rights in eDirectory to delete and modify eDirectory objects.

3 In iManager, delete the node's Cluster Node object from the cluster container:

IMPORTANT: If you plan to again add the node that is being deleted to the cluster later, it is recommended to take a note of the node number from the cluster configuration report before continuing with deletion. For information on how to generate the cluster configuration report, Section 9.9, "Generating a Cluster Configuration Report," on page 138.

3a Select **Directory Administration > Delete Objects**.

3b Browse to the Cluster container ( ) of the cluster, locate and select the Cluster Node object ( ) for the node in the container, then click **OK**.

3c On the Delete Objects page, click **OK**, then click **OK** again to confirm the deletion of the Cluster Node object.

4 Select **Directory Administration > Modify Object**, select the eDirectory Server object for the node that is being removed from the cluster, remove its NCS attributes, then click **OK** to save and apply your changes.

5 If the deleted cluster node persists in iManager, run the following command as the root user on the master node first and then on all the slave nodes that are a member of the cluster. It is safe to run the command on an active cluster:

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

6 (Optional) Build a replacement server with the same node name and IP address, and add it to the cluster. See Section 5.11, "Adding a Single Node That Was Previously in a Cluster," on page 78.

IMPORTANT: When a node is removed temporarily from the cluster for rebuilding it, you must not extend the cluster with an additional node to avoid the new node getting the node number of the temporarily removed node.

5.11 Adding a Single Node That Was Previously in a Cluster

Use the procedure in this section to completely rebuild a cluster node that has been destroyed or to reinstall a cluster node with a newer release of OES. The name and IP address of the node must be same as that of the node it replaces.

IMPORTANT: ◆While the node is being rebuild it must not have access to the shared storage. Reinstalling a server with access to shared storage can result in complete loss of shared resources.

- When a node is removed temporarily from the cluster for rebuilding it, you must not extend the cluster with an additional node to avoid the new node getting the node number of the temporarily removed node.

Prerequisites

- Before you begin replacing the cluster node, it is recommended to create a cluster configuration report and note the resources that are configured to use the node being replaced. For information on creating a cluster configuration report, see Section 9.9, “Generating a Cluster Configuration Report,” on page 138.
- You also need to note if there are any eDirectory replicas stored on this node.

  In iManager > **View Objects** > **Browse**, browse to the NCP server context representing the cluster node. Select the NCP Server object and perform the task Replica View. For information on viewing a partition’s replica, see Viewing a Partition’s Replicas in the NetIQ eDirectory Administration Guide.
- In preparation of the reinstallation, all eDirectory objects representing the node being replaced and its services must be removed from eDirectory.

  In iManager > **View Objects** > **Search**, enter *<HostName>* as the search pattern. This will list all eDirectory objects representing the node. First delete the NCP Server object only. This can take a couple of minutes.

  When the NCP Server object is deleted, click the **Multiple Select** link and repeat the search. Delete all the remaining objects related to the node being reinstalled including the object of class Unknown representing the SYS volume and the cluster node object.

  If the deleted node is still known to the active cluster nodes, force the cluster to recognize that the node has been permanently removed. See Step 5 of Section 5.10, “Removing a Node from a Cluster,” on page 77.
- Verify that the deletions have been processed properly. Verify that the deleted node is no longer known to any of the active cluster nodes by using the command `cluster view`.

Replacing the Cluster Node

1. Install OES with the same patterns that are installed on the other cluster nodes including OES Cluster Services and apply the latest patches. Ensure that you use the same node name and IP address of the node that was previously in the cluster.
2. Add any eDirectory replica that was stored on the server before it was removed from the tree.

3 Deploy the same `multipath.conf` and `bindings` file (if any) as used by the other cluster nodes. Restore access to the shared storage and reboot the node being rebuild. When the node is up, use `nlvm list` devices to verify that it has access to the same devices as the active cluster nodes.

4 Verify that the reinstalled server has access to the SBD of the cluster it will be added to by using `sbdutil -f -s -n <ClusterName>`.

5 Start YaST oes-install. On the **Micro Focus Open Enterprise Server Configuration** page under **OES Cluster Services**, enable the configuration, then click the **OES Cluster Services** link to open the OES Cluster Services Configuration wizard.

For information see, Section 5.5.4, “Accessing the OES Cluster Services Configuration Page in YaST,” on page 66.

6 Configure the replacement node for the existing cluster. Enter or browse to the cluster name. Leave the other settings at their default value and click **Next**. If your cluster node has multiple NICs ensure to select the right IP address on the third page of the cluster configuration.

See Section 5.5.6, “Adding a Node to an Existing Cluster,” on page 69.

After you restart OES Cluster Services, the replacement server joins the cluster with its old identity.

5.12 Adding Multiple Nodes That Were Previously in a Cluster

Use the information in this section to rebuild multiple cluster nodes that have been destroyed or to reinstall multiple cluster node with a newer release of OES. The names and IP addresses of the nodes must be same as that of the nodes they replace.

You can determine the node number from the node list under the **Number** column in the cluster configuration report taken before deleting the node object. For information on how to generate the cluster configuration report, see Section 9.9, “Generating a Cluster Configuration Report,” on page 138.

Figure 5-1 *Cluster Configuration Report*

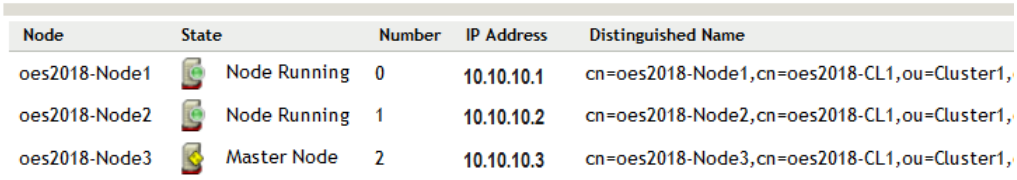

Node	State	Number	IP Address	Distinguished Name
oes2018-Node1	Node Running	0	10.10.10.1	cn=oes2018-Node1,cn=oes2018-CL1,ou=Cluster1,
oes2018-Node2	Node Running	1	10.10.10.2	cn=oes2018-Node2,cn=oes2018-CL1,ou=Cluster1,
oes2018-Node3	Master Node	2	10.10.10.3	cn=oes2018-Node3,cn=oes2018-CL1,ou=Cluster1,

Cluster Services internally identifies cluster nodes by the node number that is permanently assigned to a cluster node when NCS is configured through YaST. For NCS configuration in YaST, see Step 5 of Section 5.11, “Adding a Single Node That Was Previously in a Cluster,” on page 78.

WARNING: ◆When reinstalling multiple cluster nodes at the same time, the nodes must be added back to the cluster in a sequence with the lowest node number first.

- If you add nodes back to the cluster in any other sequence, the configuration of the reinstalled nodes does not match the configuration of the active cluster nodes. Joining such a mis-configured node back into the cluster has the potential to bring down the whole cluster.

To avoid this, it is strongly recommended to reinstall only one node at a time. If you absolutely need to reinstall multiple nodes at a time, you must perform the procedure in Section 5.11, “Adding a Single Node That Was Previously in a Cluster,” on page 78 for each of the nodes and it is critical to do the NCS configuration also in the sequence with the lowest node number first.

In case you have done the NCS configuration in an incorrect sequence, do not join the reconfigured nodes to the cluster. Immediately contact Technical Support for assistance.

IMPORTANT: When a node is removed temporarily from the cluster for rebuilding it, you must not extend the cluster with an additional node to avoid the new node getting the node number of the temporarily removed node.

5.13 Adding a Node After Another Node is Permanently Removed from the Cluster

When you add a new node to a cluster after another node is permanently removed, the new node is now assigned the node number of the node that is removed. Because of this, in the cluster view output, the new node name appears in the place where the removed node name used to appear before. As this can lead to confusion, it is strongly advised against such a configuration. You can replace the node by reinstalling it on a new hardware.

For example, in a three node cluster, you permanently remove Node2 and then add a new node Node4. The cluster view output is as follows:

```
# cluster view
Cluster CL1
This node Node3 [ epoch 7 master node Node3 ]
Cluster nodes [ Node1, Node4, Node3 ]
```

5.14 What’s Next

After installing OES Cluster Services on OES servers (physical servers or virtual guest servers), you can configure and manage the cluster and cluster resources. See the following:

- Chapter 8, “Configuring Cluster Policies, Protocols, and Properties,” on page 95
- Chapter 9, “Managing Clusters,” on page 127.
- Chapter 10, “Configuring and Managing Cluster Resources,” on page 163
- Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 201
- Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 275

If you install OES Cluster Services at the host level of an OES virtualized server, you can create cluster resources for the virtual machines. See Section 14.2, “Virtual Machines as Cluster Resources,” on page 338.

6 Joining an OES 2023 Cluster Node and Cluster Resource to an Active Directory Domain

This section describes how to join an OES 2023 cluster node and cluster resource to an Active Directory domain.

- Section 6.1, “Joining the Cluster Node to an Active Directory Domain,” on page 81
- Section 6.2, “Joining the Cluster Resource to an Active Directory Domain,” on page 82
- Section 6.3, “Verifying the Service Principals and Computer Objects,” on page 84
- Section 6.4, “Removing the Cluster Resource from an Active Directory Domain,” on page 86

6.1 Joining the Cluster Node to an Active Directory Domain

- Section 6.1.1, “Prerequisites,” on page 81
- Section 6.1.2, “NSS-AD Support Pattern,” on page 81
- Section 6.1.3, “Configuring NSS for Active Directory,” on page 81

6.1.1 Prerequisites

- We recommend that you upgrade all nodes to OES 2023 in order to provide users uninterrupted access to their data.
- The DNS names of the domain controller and the OES 2023 server must resolve each other.
- The clocks on the OES 2023 server and Active Directory server must be in sync in order for Kérberos to work properly.
- You join a computer to Active Directory with a domain administrator, or a domain administrator equivalent, or a standard domain user account with sufficient rights.

  If you plan to designate a standard user, refer to Microsoft documents for configuration details.

6.1.2 NSS-AD Support Pattern

The OES installation process remains same as in previous releases. Selecting the NSS AD Support pattern also selects and installs all the other dependencies.

NSS AD Support requires configuration in order for the installation to proceed further. If you want to configure it later, disable it, and then continue with the installation.

6.1.3 Configuring NSS for Active Directory

Accept the advisory notes and then click Next to proceed the installation further.

AD Domain Name: Displays the domain name based on this server's LAN settings. Specify the domain name if you want to join to a different domain.

AD User Name: Specify the user credentials of the Administrator, or an administrator equivalent, or a standard domain user with sufficient rights to create computer objects.

AD Supervisor Group: Specify the AD supervisor group name. The AD users belonging to this group will have supervisory rights for all the volumes associated with that OES server.

Container to create Computer Object: By default, the domain join process creates the computer object for this server to the default container CN=**Computers**. If you specify a different Organizational Unit, for example, *OU=OES2023Servers*, it adds the computer object to that OU. Create the OU before you start configuring NSS for Active Directory.

Use pre-created computer object: Selecting this check box maps this server to the Computer Object that exists in the container you specified. Before you start configuring NSS for Active Directory, create the Computer Object with the NetBIOS name of the node.

Novell Identity Translator: OES includes a service named Novell Identity Translator (NIT) that dynamically provides Linux user IDs on OES 2018 or later systems to eDirectory and Microsoft Active Directory users for NSS file access. NIT can be configured to work with either eDirectory only or both eDirectory and Active Directory.

For more information about command line options, see novell-ad-util in the OES 2023: NSS AD Administration Guide.

6.2 Joining the Cluster Resource to an Active Directory Domain

- Section 6.2.1, “Prerequisites,” on page 82
- Section 6.2.2, “Steps for Joining the Cluster Resource,” on page 82

6.2.1 Prerequisites

- The cluster pools that you plan to join must be in the Active state.
- Each cluster pool must have at least one volume and the volume must be in an active and mounted state.
- Upgrade all nodes to OES 2023 and join them to an Active Directory domain in order to provide users with uninterrupted access to their data.

6.2.2 Steps for Joining the Cluster Resource

This process includes four main tasks:

1. Upgrading the pool media format.
2. Enabling the Active Directory identities flag for a volume that you plan to provision for the Active Directory users.
3. Joining the cluster resources to an Active Directory Domain.
4. Verifying the service principals for the node and the resource.

Upgrading the Pool Media Format and Active Directory-enabling the Volumes

First, upgrade the existing NSS32 cluster pools with the new metadata structure to provision NSS resources to Active Directory users. This is a one-time activity. Once completed, you cannot return the media to its previous state.

If you create NSS32 pools, you have to explicitly upgrade them with the new media format. However, the NSS 64 pools in OES 2015 or later are shipped with the new media format.

After the media upgrade, you have to enable the Active Directory identities flag for the volumes in NSS32 and NSS64 pools.

Issue the commands at the NSS Console as the root user.

1. Run the `/pools` command and verify that the cluster pool is in the active state.
2. Run the `/ZLSSUpgradeCurrentPoolMediaFormatToAD=<POOL_NAME>` command.

   A successful upgrade is indicated by an appropriate message.
3. Run the `/volumes` command and verify that the volume is in the active state.
4. Run the `/ADIdentities=<VOLUME_NAME>` command.

   The successful addition of Active Directory identities is indicated by an appropriate message.
5. Next, run the `/volumes` command again and verify that the **AD Enabled** attribute appears in the Attributes column.
6. Create a 64-bit pool and a volume. Although the 64-bit pool is created with the upgraded media, the volume is not AD-enabled.
7. Run the `/ADIdentities=<VOLUME_NAME>` command to add Active Directory identities.

For more information, see Upgrading the NSS Media Format in the OES 2023: NSS File System Administration Guide for Linux.

Joining the Cluster Resources to an Active Directory Domain

Next, you will join the resources to an Active Directory domain.

To join the resources, use the user credentials of the Administrator, or an administrator equivalent, or a standard domain user with sufficient rights to create computer objects.

1. Run `kinit` and then open the NSSMU console. You can also open the NSSMU console without running `kinit`.
2. Select **Pools** from the menu, and then select the pool that you want to join.
3. Press `J` on your keyboard.
4. Specify valid **Active Directory Information**, then **Proceed**.

   If `kinit` has already been performed and you want to retain the credentials, press **Y** to skip. Otherwise, press **N** to run `kinit`. This will destroy the credentials cache and you will need to provide new credentials.
5. Specify the Administrator credentials or the credentials of the user that has sufficient rights, and then press **ok**.

NOTE: The credentials that you provide are valid only for this session.

For more information, see Joining Cluster Pools to the AD Domain in the OES 2023: NSS File System Administration Guide for Linux.

For more information about commandline options, see novell-ad-util in the OES 2023: NSS AD Administration Guide.

6.3 Verifying the Service Principals and Computer Objects

This section describes how to verify the service principals created on the OES cluster node and the computer objects created on the Active Directory domain.

Verify that the service principals and computer objects are created for the cluster node and cluster resource.

1 From the terminal console, run the `klist -k` command.

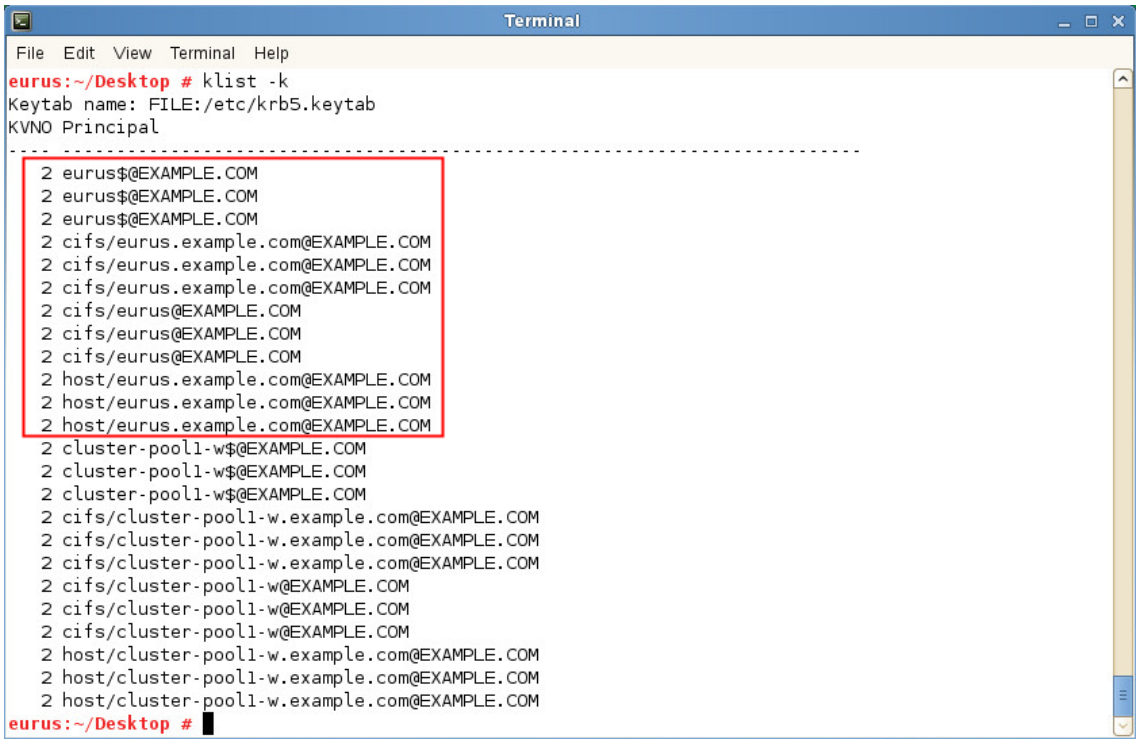

These are the Service Principals of the cluster node.

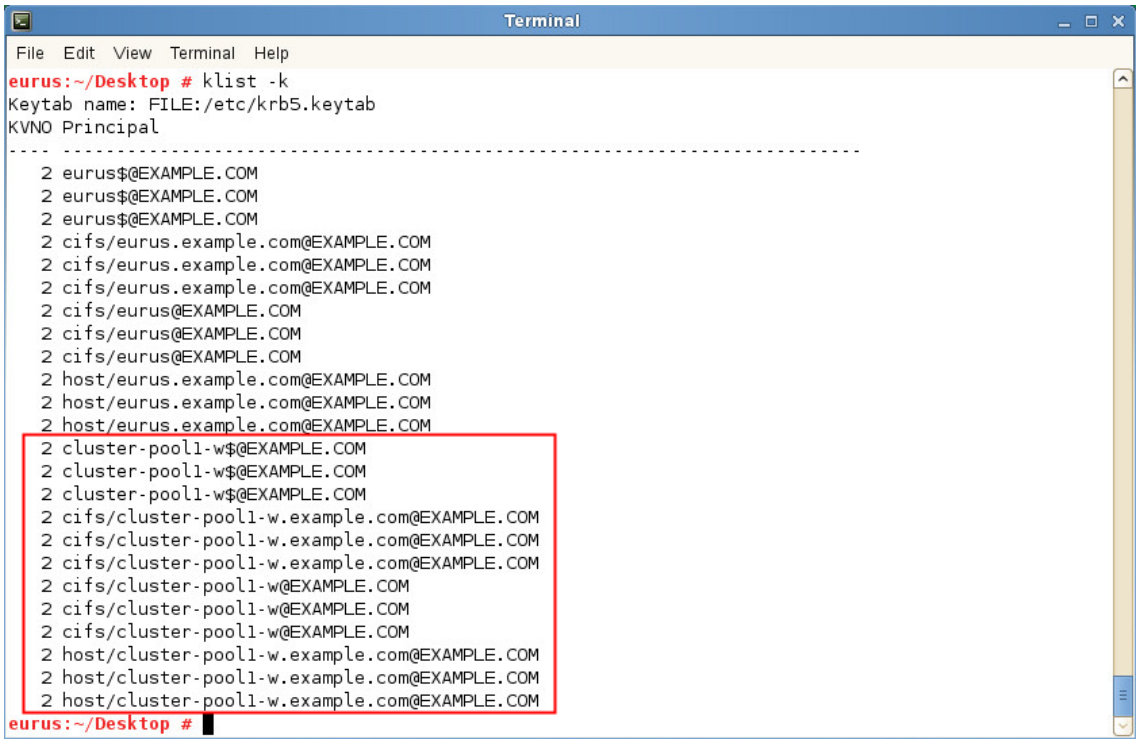

These are the Service Principals of the cluster resource.

2 Open the Active Directory users and computers console.

You will notice that computer objects have been created or mapped to the pre created computer objects for the cluster node and the cluster resource.

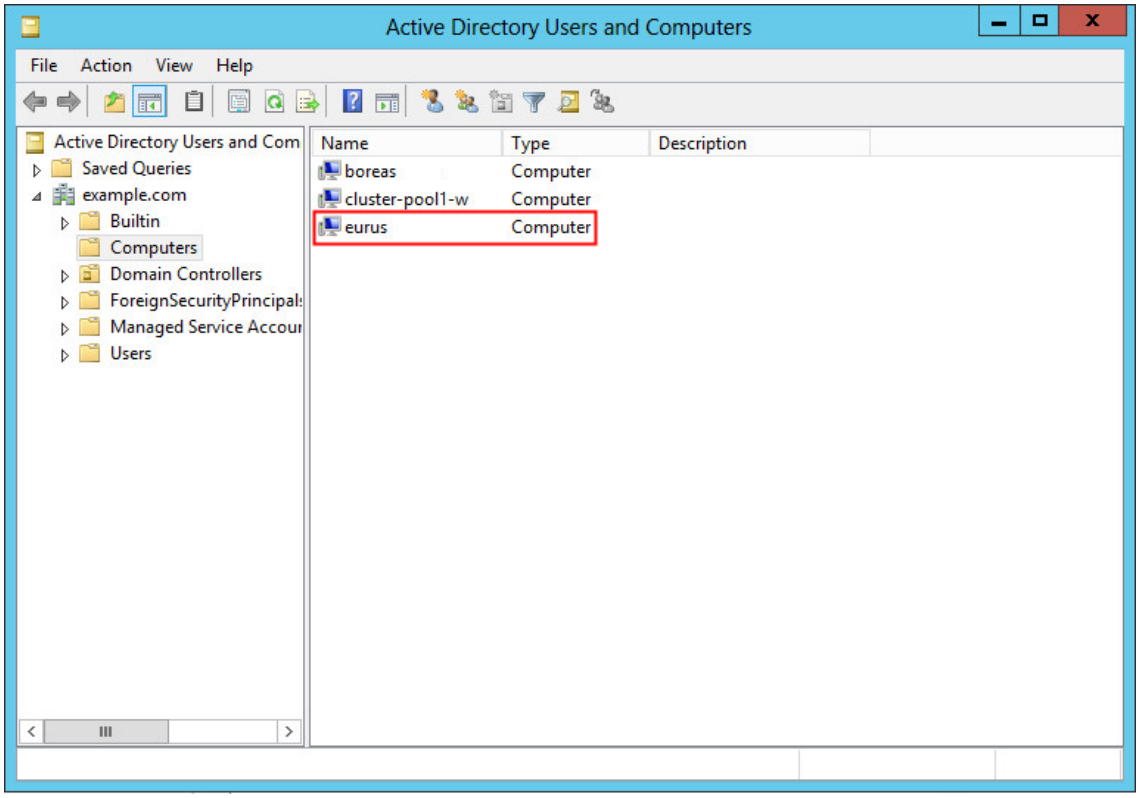

This is the Cluster node that you joined to the Active Directory domain.

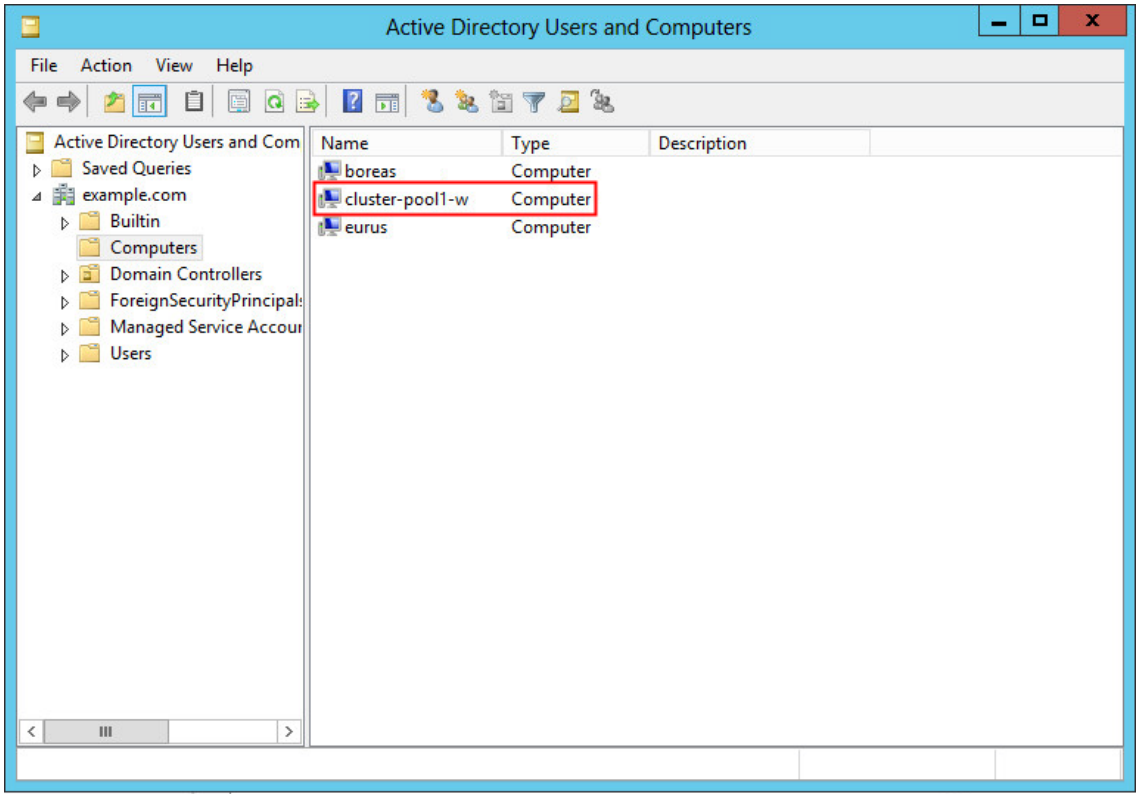

This is the Cluster resource that you joined to the Active Directory domain.

6.4 Removing the Cluster Resource from an Active Directory Domain

Executing the `./novell-ad-util --leave-domain --cluster-resource <resource_FDN_eDir_format> --domain-name <DOMAIN_NAME>` command will disjoin the cluster resource from the Active Directory domain.

Example:

`./novell-ad-util --leave-domain --cluster-resource .cn=CLUSTER-OES2023-POOL-SERVER.o=novell.t=NSSAD_CLUSTER. --domain-name EXAMPLE.COM`

How do I remove stale entries of keytab for unjoined cluster resources on all cluster nodes in the cluster?

When you disjoin a cluster resource from an Active Directory domain, novell-ad-util removes the keytab entries of that resource from the default keytab file, `/etc/krb5.keytab`, and deletes the volume keytab file. For example, `/media/nss/vol1/._NETWARE/vol.keytab` on the node where the resource is running.

Before disjoining the resource, if you have migrated it to other cluster nodes, all the cluster nodes where the resource is migrated will have the default keytab entries.

When you disjoin the cluster resource, the default keytab entries for that specific cluster node and the volume keytab entries will be removed. However, the default keytab entries will still be seen on those nodes where the resource was migrated.

To remove the stale entries, execute the following command respectively all nodes other than the node that you used for the resource disjoin:

```
./novell-ad-util --purge 0 --cluster-resource <cluster dn> --domain-name
<domain name>
```

This command removes the keytab entries of the cluster resource <cluster dn> specified; it will not remove the volume keytab file.

How do I migrate resources in a Mixed-Mode (nodes joined to the AD domain and nodes not joined to the AD domain) environment?

In a mixed-mode state, if you migrate a resource from a node joined to an AD domain to a node that is not, AD users provisioned to that resource will lose access to their data. However, eDirectory users will continue to have access without any hindrance.

At a later point in time, if you join the other node (not joined to AD domain) to the AD domain, AD users will still not have access to their data, even though the resource is up and running. To enable AD users access to their data, offline the resource and bring it back online.

7 Upgrading OES Clusters

This section describes how to upgrade an OES Cluster Services and the file system resources. See the individual service guides for information about upgrading clustered services.

- Section 7.1, “Requirements and Guidelines for Upgrading Clusters,” on page 89
- Section 7.2, “Upgrading Nodes in the Cluster (Rolling Cluster Upgrade),” on page 91

7.1 Requirements and Guidelines for Upgrading Clusters

As you upgrade a cluster to the next release version for the OES software, you have a mixed-mode cluster made up of nodes running the old platform and the next SP platform. In addition to the Chapter 4, “Planning for OES Cluster Services,” on page 35, consider the following rules and recommendations for mixed clusters during an upgrade to the next release:

- Mixed-mode clusters should be considered a temporary configuration that exists only during an upgrade.
- Adding a new cluster node with the old version to a mixed-mode cluster is not supported.
- NSS pool cluster resources can fail over to any node in the mixed-mode cluster.
- LVM volume group cluster resources can fail over to any node in the mixed-mode cluster.
- CSM cluster resources can fail over to any node in the mixed-mode cluster.
- No storage management functions should be executed while a cluster is in mixed-mode unless you must do so as part of the documented upgrade process. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

7.1.1 Behavior of an NSS Pool Resource with Media Version 44.03 and Above in Mixed Node Cluster

A mixed node cluster environment can contain OES 11 SP2, OES 2015, OES 2015 SP1, OES 2018, OES 2018 SP1, OES 2018 SP2 and OES 2018 SP3 nodes. The AD media upgraded NSS32 and NSS64 pools cannot be loaded in OES 11 SP2 nodes.

When a resource comes online, it tries to load on a node based on the preferred node assignment (OES 2015 or later nodes), and the older OES cluster nodes are skipped from the preferred node list. If OES 2015 or later nodes are not available in the cluster, the resources are moved to “unassigned” state.

However, if the upgraded (OES 2015 or later) node comes up or any new OES (OES 2015 or later) node is added to the cluster, the resource will automatically loads on the OES 2015 or later node.

IMPORTANT: ◆All nodes in the cluster must be patched with the latest OES patches. Otherwise, the AD media resource goes to comatose state on nodes earlier than OES 2015.

- All nodes in the cluster must belong to OES 11 SP2 or later. If any OES node in the cluster is older than OES 11 SP2, this feature does not work.

In NSS32 and NSS64 pool type, the pool status can be determined by the media version. Table 7-1 provides information about the NSS32 and NSS64 pool type media version.

Table 7-1 *NSS Pool Media Version Details*

Pool Media Version	Description
44.03	◆ **44:** An NSS32 pool that supports AD. ◆ **03:** All volumes are upgraded to support hard links.
51.00	◆ **51:** An NSS64 pool that supports AD by default. ◆ **00:** All volumes are upgraded to support hard links.
45.00	◆ **45:** An NSS32 pool that supports Trustee Index. ◆ **00:** Some volumes are not upgraded to support Trustee Index.
52.00	◆ **52:** An NSS64 pool that supports Trustee Index. ◆ **00:** Some volumes are not upgraded to support Trustee Index.
45.01	◆ **45:** An NSS32 pool that supports Trustee Index. ◆ **01:** All volumes are upgraded to support Trustee Index.
52.01	◆ **52:** An NSS64 pool that supports Trustee Index. ◆ **01:** All volumes are upgraded to support Trustee Index.

Similarly, the volume status can also be determined by the media version. Table 7-2 provides information about the NSS volume media version.

Table 7-2 *NSS Volume Media Version Details*

Volume Media Version	Description
38.05	◆ **38:** An NSS volume that supports Hard Link. ◆ **05:** An NSS volume is upgraded to support hard links.
41.00	◆ **41:** An NSS volume that supports Trustee Index. ◆ **00:** An NSS volume is not upgraded to support Trustee Index.

Volume Media Version	Description
41.01	• **41:** An NSS volume that supports Trustee Index. • **01:** An NSS volume is upgraded to support Trustee Index.

7.2 Upgrading Nodes in the Cluster (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade lets you keep your cluster up and running and lets your users continue to access cluster resources while the upgrade is being performed.

During a rolling cluster upgrade, one OES server is upgraded at a time to the next SP release. You can also add new nodes running the next release. You should complete the upgrade as soon as possible. Don't leave the cluster in a mixed-mode state for an extended period.

To perform a rolling cluster upgrade:

1. Make a note of the OES components that are installed on the existing cluster nodes.

   You will probably want to install the same OES components on any new nodes that you add to the cluster.

2. If you are adding new nodes, prepare the nodes for the cluster, but do not allow them to join the cluster.

   2a Install the OES Services from the OES Install Media, including OES Cluster Services, but do not configure the node for the cluster at this time.

   2b Verify that the server is functioning properly.

   2c In YaST, go to the OES Cluster Services Configuration page as described in Section 5.5.4, "Accessing the OES Cluster Services Configuration Page in YaST," on page 66.

   2d Configure the node for the existing OES cluster as described in Section 5.5.6, "Adding a Node to an Existing Cluster," on page 69. Complete the following fields:

Parameter	Action
LDAP administrator user password	Specify the password of the user name shown, then click **OK**. Typically, this is the LDAP administrator identity that was used when the OES Services were installed, or the identity configured after the install in Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 66.
New or existing cluster	Select **Existing Cluster**.

Parameter	Action
Directory Server Address	Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default. You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.13.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 120.
Cluster FDN	Click **Browse**, navigate the tree to select the existing Cluster object for the OES 11 cluster, click **OK**, then click **Next**.
NCS Proxy User Name	Specify one of the following users as the NCS Proxy user for this server, then click **Next**: ◆ OES Common Proxy User (default if it was configured for the server) ◆ LDAP Admin User ◆ Another Administrator User
IP address of this node	From the drop-down list, select the IP address that OES Cluster Services will use for this node.
Start Cluster Services now	Deselect the check box. Do not start cluster services on this node or join the cluster until, after all of the Linux POSIX cluster resources have been taken offline.

2e Click **Finish**, review the summary, then click **Next** to complete the configuration.

2f Click **Finish** to close the OES Configuration dialog box.

3 For each of the new nodes in turn, start OES Cluster services manually by using one of the following methods, and allow the node to join the existing cluster.

- Reboot the cluster server.
- At a command prompt, enter the following as the `root` user:

  ```
  rcnovell-ncs start
  ```

  or

  ```
  systemctl start novell-ncs.service
  ```

4 For each new node, ensure that you modify the **Preferred Nodes** list accordingly for each of the cluster resources.

5 For each of the existing cluster nodes in turn:

5a Log in to the node as the `root` user, then open a terminal console.

5b Use the `cluster leave` command to remove the node from the cluster.

Any cluster resources that were running on the server should fail over to another server in its preferred nodes list.

You can also manually cluster migrate the resources to another server in the cluster prior to bringing down the server.

5c Stop OES Cluster Services by entering

```
rcnovell-ncs stop
```

or

```
systemctl stop novell-ncs.service
```

5d If the autostart of OES Cluster Services is enabled, you must disable autostart in order to prevent it from automatically loading OES Cluster Services during the upgrade. Enter

```
systemctl disable novell-ncs.service
```

5e Apply the next SP release as an upgrade or upgrade to next release, and fully patch the system.

NOTE: Post upgrade if the shared storage or SBD partition is unavailable on the machine, ensure they are available and then restart NCS. SBD partition might not be available because iscsi service is disabled post upgrade.You can enable the iscsi service using `systemctl enable iscsi.service`.

5f (Optional) Enable the autostart of OES Cluster Services for the node, enter

```
systemctl enable novell-ncs.service
```

5g Start OES Cluster Services by entering

```
rcnovell-ncs start
```

or

```
systemctl start novell-ncs.service
```

5h Use the `cluster join` command to rejoin the node to the cluster.

5i Modify the **Preferred Nodes** list accordingly for each of the cluster resources.

5j Repeat these steps for each existing node in the cluster.

6 After all of the existing nodes have been upgraded and any new nodes added, the conversion is complete.

8 Configuring Cluster Policies, Protocols, and Properties

After installing OES Cluster Services on one or more nodes in a cluster, you can configure the settings for the cluster to meet your needs and help you manage the cluster effectively. This additional configuration might consist of changing the values on some of the properties for the Cluster object.

- Section 8.1, “Understanding Cluster Settings,” on page 95
- Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99
- Section 8.3, “Selecting a Cluster to Manage,” on page 102
- Section 8.4, “Configuring Quorum Membership and Timeout Policies,” on page 104
- Section 8.5, “Configuring Cluster Protocols,” on page 105
- Section 8.6, “Configuring Cluster Event Email Notification,” on page 106
- Section 8.7, “Configuring Cascade Failover Prevention,” on page 108
- Section 8.8, “Configuring NCS to Monitor the eDirectory Daemon (ndsd),” on page 111
- Section 8.9, “Configuring the Cluster Node Reboot Behavior,” on page 114
- Section 8.10, “Configuring STONITH,” on page 115
- Section 8.11, “Viewing or Modifying the NCS Proxy User Assignment in the NCS Management Group,” on page 117
- Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118
- Section 8.13, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 120
- Section 8.14, “What’s Next,” on page 125

8.1 Understanding Cluster Settings

IMPORTANT: You must perform all Cluster Services configuration operations on the master node in the cluster. In iManager, select the Cluster object, not the Cluster Node objects.

- Section 8.1.1, “Cluster Policies,” on page 96
- Section 8.1.2, “Cluster Priorities,” on page 97
- Section 8.1.3, “Cluster Protocols,” on page 97
- Section 8.1.4, “RME Groups,” on page 98
- Section 8.1.5, “BCC,” on page 98
- Section 8.1.6, “Cascade Failover Prevention,” on page 98
- Section 8.1.7, “Monitoring the eDirectory Daemon,” on page 98
- Section 8.1.8, “Cluster Node Reboot Behavior,” on page 99

- Section 8.1.9, “STONITH,” on page 99
- Section 8.1.10, “Cluster Information,” on page 99

8.1.1 Cluster Policies

Cluster policies are configuration settings that determine how a cluster is accessed, how many nodes are needed to form a quorum on startup, and if email notifications are sent. You can manage cluster policies in iManager by going to **Clusters > My Clusters**, selecting a cluster, then selecting **Action > Edit Properties > Policies**. Table 8-1 describes the configurable cluster policies:

Table 8-1 *Cluster Policies*

Property	Description
Cluster IP address	Specifies the IP address for the cluster. You specify the IP address when you install Cluster Services on the first node of the cluster. The IP address is bound to the master node and remains with the master node regardless of which server is the master node. Rarely, you might need to modify this value. See Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118. **IMPORTANT:** If you modify the master IP address, you must restart the cluster, exit iManager, then re-launch iManager before iManager can continue to manage the cluster.
Port	Specifies the port used for cluster communication. The default cluster port number is 7023, and is automatically assigned when the cluster is created. You might need to modify this value if there is a port conflict. You can change the port number to any other value that does not cause a conflict. See Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118.
Quorum membership - Number of nodes	Specifies number of nodes that must be up and running in the cluster before resources begin to load. Specify a value between 1 and the number of nodes. Set this value to a number greater than 1 so that all resources are not automatically loaded on the first server that is brought up in the cluster. For example, if you set the membership value to 4, there must be four serves up in the cluster before any resource will load and start. For instructions, see Section 8.4, “Configuring Quorum Membership and Timeout Policies,” on page 104.
Quorum timeout	Specifies the maximum amount of time to wait for the specified quorum to be met. If the timeout period elapses before a quorum is achieved, resources automatically begin loading on whatever number of nodes are actually up and running in the cluster. You can specify the time in seconds or minutes. For instructions, see Section 8.4, “Configuring Quorum Membership and Timeout Policies,” on page 104.

Property	Description
Email notification	Enables or disables email notification for the cluster. If it is enabled, you can specify up to eight administrator email addresses for cluster events notification. Specifies the type of cluster events for notifications. You can receive notifications for only critical events such as node failure or a resource going comatose, or you can receive notifications for all cluster state changes and resource state changes. Specifies whether to receive messages in XML format. XML format messages can be interpreted with a parser and formatted to customize the message information for your specific needs. For instructions, see Section 8.6, "Configuring Cluster Event Email Notification," on page 106.

8.1.2 Cluster Priorities

Cluster priorities determine the load priority of individual cluster resources. You can manage cluster resource priorities in iManager by going to the **Clusters > My Clusters**, selecting a cluster, then selecting **Action > Edit Properties > Priorities**. See Section 10.11, "Configuring Resource Priorities for Load Order," on page 185.

8.1.3 Cluster Protocols

Table 8-2 describes the configurable cluster protocols properties that govern inter-node communication transmission and tolerances. You can manage cluster protocols policies in iManager by going to the **Clusters > My Clusters**, selecting a cluster, then selecting **Action > Edit Properties > Protocols**. See Section 8.5, "Configuring Cluster Protocols," on page 105.

Table 8-2 *Cluster Protocols*

Property	Description
Heartbeat	Specifies the interval of time in seconds between signals sent by each of the non-master nodes in the cluster to the master node to indicate that it is alive. The default is 1 second.
Tolerance	Specifies the maximum amount of time in seconds that a master node waits to get an alive signal from a non-master node before considering that node to have failed and removing it from the cluster. The default is 8 seconds.
Master watchdog	Specifies the interval of time in seconds between alive signals sent from the master node to non-master nodes to indicate that it is alive. The default is 1 second. Modify this parameter setting only when supervised by Micro Focus Technical Support (http://www.novell.com/support).

Property	Description
Slave watchdog	Specifies the maximum amount of time in seconds that the non-master nodes wait to get an alive signal from the master node before considering that the master node has failed, assigning another node to become the master node, and removing the old master node from the cluster. The default is 8 seconds. Modify this parameter setting only when supervised by Micro Focus Technical Support (http://www.novell.com/support).
Maximum retransmits	This value is set by default and should not be changed. The default is 30.

8.1.4 RME Groups

The Resource Mutual Exclusion (RME) Groups page allows you to view or define sets of resources that must not run on the same node at the same time. You can manage RME Groups in iManager by going to the **Clusters > My Clusters**, selecting a cluster, then selecting **Action > Edit Properties > RME Groups**. See Section 10.12, “Configuring Resource Mutual Exclusion Groups,” on page 185.

8.1.5 BCC

If OES Business Continuity Clustering (BCC) is installed in the cluster, the BCC properties page allows you to view or change policies for the selected cluster. You can manage BCC policies in iManager by going to **Clusters > My Clusters**, selecting a cluster, then selecting **Action > Edit Properties > BCC**. See the *Novell Business Continuity Clustering website* (http://www.novell.com/documentation/bcc/).

8.1.6 Cascade Failover Prevention

OES Cluster Services added the cascade failover prevention function that detects if a node has failed because of a bad cluster resource and prevents that bad resource from failing over to other servers in the cluster. It is enabled by default. See Section 8.7, “Configuring Cascade Failover Prevention,” on page 108.

8.1.7 Monitoring the eDirectory Daemon

Some clustered services rely on the eDirectory daemon (ndsd) to be running and available in order to function properly. NCS provides the ability to monitor the status of the eDirectory daemon (ndsd) at the NCS level. It is disabled by default. The monitor can be set independently on each node. On a node, if the eDirectory daemon does not respond to a status request within a specified timeout period, NCS can take one of three configurable actions: an ndsd restart, a graceful node restart, or a hard node restart. See Section 8.8, “Configuring NCS to Monitor the eDirectory Daemon (ndsd),” on page 111.

8.1.8 Cluster Node Reboot Behavior

The OES Cluster Services reboot behavior conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown. On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems. To control the cluster node reboot behavior, you can use the directive

`kernel.panic` in the Linux `/etc/sysctl.conf` file to prevent a reboot, or to allow an automatic reboot and to specify the number of seconds to delay the reboot. See Section 8.9, “Configuring the Cluster Node Reboot Behavior,” on page 114.

8.1.9 STONITH

The STONITH (shoot-the-other-node-in-the-head) capability allows OES Cluster Services to kill a suspect node by using remote power control. Unlike a poison pill, it does not require a response from the suspect node. STONITH is used after a poison pill is issued; it does not replace the poison pill. See Section 8.10, “Configuring STONITH,” on page 115.

8.1.10 Cluster Information

After you install OES Cluster Services, you use the **Open Enterprise Server > OES Install and Configuration** tool in YaST to set up the cluster or to add a node to an existing cluster. The YaST-based configuration is not used to modify the settings for an existing cluster. For information about modifying the settings for an existing cluster, see:

- Section 8.11, “Viewing or Modifying the NCS Proxy User Assignment in the NCS Management Group,” on page 117
- Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118
- Section 8.13, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 120.

8.2 Setting Up a Personalized List of Clusters to Manage

The My Clusters page in the Clusters plug-in for iManager allows each cluster administrator to set up a personalized list of clusters to manage. This allows an administrator to view at a glance the status of the specified clusters. The list persists between the administrator's logins to the same iManager server.

Your personalized list of clusters and display preferences are saved on the iManager server in the following location:

`/var/opt/novell/iManager/nps/WEB-INF/config/NDS<tree_name>-<user_name_and_context_without_dots>/ncs.xml`

For example, for tree AVALONTREE and user `admin.novell`, the file path is:

`/var/opt/novell/iManager/nps/WEB-INF/config/NDSAVALONTREE-adminnovell/ncs.xml`

The following sections describe how manage the list of clusters and to personalize the display.

- Section 8.2.1, “Adding Clusters to Your My Clusters List,” on page 100
- Section 8.2.2, “Viewing Information about Clusters in Your My Clusters List,” on page 100
- Section 8.2.3, “Personalizing the My Clusters Display,” on page 101
- Section 8.2.4, “Removing Clusters from a My Clusters List,” on page 102

8.2.1 Adding Clusters to Your My Clusters List

1. Log in to iManager as a cluster administrator.
2. In **Roles and Tasks**, select **Clusters > My Clusters**.

   The list is initially empty.
3. Click **Add** to open the eDirectory browser pop-up window.
4. Browse the tree where you are currently logged in to locate and select one or more Cluster objects, then click **OK**.

   Newly selected clusters are added to your personalized list.
5. (Optional) Personalize the display as described in Section 8.2.3, “Personalizing the My Clusters Display,” on page 101.

8.2.2 Viewing Information about Clusters in Your My Clusters List

1. Log in to iManager as a cluster administrator.
2. In **Roles and Tasks**, select **Clusters > My Clusters**.

   The My Clusters page displays your personalized list of clusters in the tree.
3. View the Status icon next to a Cluster object to understand the current overall state of the cluster.

State	Icon	Description
Normal		The cluster is up and running.
Stopped		The cluster is stopped. Administrator intervention is needed.
Critical		One or more nodes in the cluster have failed. Administrator intervention is needed.
Warning		An alert condition has occurred. A node or resource needs administrator attention.
Unknown		The state of the cluster cannot be determined.

4 The My Clusters list displays the following information about each cluster:

Parameter	Description
Name	The name assigned to the cluster.
Distinguished Name	(Not displayed by default) The dot-delimited fully distinguished name of the cluster, such as `clustername.context.domain`.
Master Node	The name of the currently assigned master node for the cluster.
IP Address	The unique IP address that represents the cluster.
Nodes	The number of active nodes in the cluster.
Epoch	The number of times the cluster state has changed. The cluster state changes every time a server joins or leaves the cluster.
BCC	A check mark in this column indicates that the cluster is enabled for a Business Continuity Cluster (BCC).

5 (Optional) Personalize the display as described in Section 8.2.3, “Personalizing the My Clusters Display,” on page 101.

8.2.3 Personalizing the My Clusters Display

The My Clusters page provides menus in the column headings that allow you to personalize the display. You can sort the entries, modify the columns, or filter the entries.

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Clusters**.

3 If the list is empty, you must add at least one cluster in order to be able to personalize the display. Select **Add**, browse to select the Cluster object, then click **OK**.

4 Select **Refresh**, then choose a refresh rate that is long enough to allow you to view all items in the list.

5 Click or mouse-over a column heading to activate its options, then click the arrow to access the menu for that column.

6 Perform one of the available actions when you access the column's menu. Some actions are not available for all fields.

- **Sort:** Select **Sort Ascending** (A to Z) or **Sort Descending** (Z to A) to specify the preferred sort order. Cluster objects are sorted based on values in the selected column. Numbers are treated as text values and are sorted alphabetically, not numerically.
- **Columns:** Select **Columns** to display the parameters, select or deselect a parameter's check box to add or remove the column, then press Enter to apply the changes. The parameters are described in Step 4 in Section 9.17, "Viewing the Cluster Node Properties," on page 145.
- **Filter:** Specify a filter to display only the Cluster objects with values that match.

  To add or modify a column's filter, select **Filter**, place the cursor in the **Filter** field, type the desired filter text, then press Enter to apply the filter. If a cluster's parameter value matches the specified filter, the Cluster object and its information are displayed.

  To remove a column's filter, select Filter, place the cursor in the **Filter** field, delete the filter text, then press Enter to refresh the display. The filter no longer applies for that column.

8.2.4 Removing Clusters from a My Clusters List

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Clusters**.

The My Clusters page displays your personalized list of clusters in the tree that you want to manage.

3 Select the check box next to the cluster that you want to remove from your personalized list, then click **Remove**.

The cluster is removed from your personalized list. This action does not delete the Cluster object or related objects.

8.3 Selecting a Cluster to Manage

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Clusters**.

The My Clusters page displays your personalized list of clusters in the tree that you want to manage. To add or remove clusters from your list, see Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

The Clusters plug-in sends a status query to the master IP address of each cluster in the list.

3 If the master IP address of a cluster has connection problems, the Clusters plug-in allows you to use the server IP address of a node (preferably the master node) to manage the cluster and fix the problem.

On the My Clusters page, if the cluster status query fails:

3a Read the message about the problem connecting to the master IP address.

The message appears each time that you revisit the My Clusters page until you resolve the IP address conflict on the network or by modifying the cluster's master IP address.

3b Do one of the following:

- **Use a server node (preferably the master):** Click **Yes** to agree to use the server IP address of a node to temporarily manage the cluster.
- **View error details:** Click **No** to view the **Cluster Communication Error** details. It provides some tips for diagnosing the root cause of the IP address conflict. Click **OK** to dismiss the message. Try to resolve the conflict on the network, then select **Clusters > My Clusters** to see if the connection issue is resolved.

3c Browse to choose the node (preferably the master node) to use to temporarily manage the cluster, then click **OK**.

3d If you need to modify the master IP address to fix the connection problem, see Section 8.12, "Viewing or Modifying the Cluster Master IP Address or Port," on page 118.

4 After a successful connection to the cluster master IP address, do any of the following:

- **Cluster Properties:** To view or manage the cluster policies, protocols, or properties, select the check box next to the cluster, then select **Action > Edit Properties**.

  The Cluster Properties dialog box opens. The browser's pop-up blocker must be disabled. See the following:

  - Section 8.4, "Configuring Quorum Membership and Timeout Policies," on page 104
  - Section 8.5, "Configuring Cluster Protocols," on page 105
  - Section 8.6, "Configuring Cluster Event Email Notification," on page 106

- **Resources:** To manage cluster resources, click the cluster name link.

  The browser opens on the Cluster Manager page in the same browser window. See Chapter 10, "Configuring and Managing Cluster Resources," on page 163.

- **Cluster Report:** To generate a report for the cluster, select the check box next to the cluster, then select **Action > Run Report**.

  The Cluster Report dialog box opens. The browser's pop-up blocker must be disabled. See Section 9.9, "Generating a Cluster Configuration Report," on page 138.

- **Repair an OES Cluster:** To repair an OES cluster, select the check box next to the cluster, then click **Action** > **Repair**.

  The Repair option is useful when you see a mismatch in the cluster resource list or the resource priority list. The cluster repair will make the resource list consistent with the resource objects.

  For example, if the cluster status command output is blank or the deleted resources are still listed, the repair option will make the resource list match the resource objects. Resources will then be displayed properly in CLI and iManager.

  NCS logs the corresponding repair option messages in the `/var/opt/novell/log/ncs/repair.log` file.

8.4 Configuring Quorum Membership and Timeout Policies

The quorum membership and timeout properties govern when cluster resources begin loading on cluster startup, failback, or failover.

1 In iManager, select **Clusters > My Clusters**.

2 Select the check box of the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Select **Action > Edit Properties** to open the Properties dialog box.

4 Click the **Policies** tab.

5 Under **Quorum Triggers**, specify the number of nodes that are required to form a quorum for the specified cluster.

See Section 8.4.1, “Quorum Membership (Number of Nodes),” on page 104.

6 Under **Quorum Triggers**, specify the amount of time in seconds to wait for the quorum to form before beginning to load the cluster resources without a quorum being formed.

See Section 8.4.2, “Quorum Timeout,” on page 104.

7 Click **Apply** or **OK** to save your changes.

8 Restart Cluster Services. Open a terminal console as the `root` user, then enter the following at the command prompt:

```
rcnovell-ncs restart
```

or

```
systemctl restart novell-ncs.service
```

8.4.1 Quorum Membership (Number of Nodes)

A cluster quorum is the number of nodes that must be running in the cluster before resources begin to load. When you first bring up servers in your cluster, OES Cluster Services reads the number specified in this field and waits until that number of servers is up and running in the cluster before it starts loading resources.

Set this value to a number greater than 1 so that all resources do not automatically load on the first server that is brought up in the cluster. For example, if you set the **Number of Nodes** value to 4, there must be four servers up in the cluster before any resource loads and starts.

8.4.2 Quorum Timeout

Timeout specifies the amount of time to wait for the number of servers defined in the **Number of Nodes** field to be up and running. If the timeout period elapses before the quorum membership reaches its specified number, resources automatically start loading on the servers that are currently up and running in the cluster. For example, if you specify a **Number of Nodes** value of 4 and a timeout value equal to 30 seconds, and after 30 seconds only two servers are up and running in the cluster, resources begin to load on the two servers that are up and running in the cluster.

The timeout is an attribute of the cluster. It gets initialized when the NCS modules are being loaded. If you change the value, restart Cluster Services (`rcnovell-ncs restart` or `systemctl restart novell-ncs.service`) to get the updated settings.

8.5 Configuring Cluster Protocols

The cluster protocols determine the transmit frequency and tolerance settings for all nodes in the cluster, including the master node. The master node is generally the first node brought online in the cluster, but if that node fails, any of the other nodes in the cluster can become the master.

IMPORTANT: If you change any protocol properties, you should restart all servers in the cluster to ensure that the changes take effect.

---

1. In iManager, select **Clusters > My Clusters**.
2. In your personalized **My Clusters** list, select the check box next to the cluster you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Click **Action > Edit Properties** to open the Properties dialog box.
4. Click the **Protocols** tab.

   The Protocols page allows you to view or change cluster protocol settings such as heartbeat, tolerance, master/slave watchdogs, and maximum retransmits. The Protocols page also lets you view the script used to configure the cluster protocol settings, but not to change it. Changes made to the protocols setting automatically update the scripts.
5. Specify values for the cluster protocols properties. See the following:
   - Heartbeat
   - Tolerance
   - Master Watchdog
   - Slave Watchdog
   - Maximum Retransmits
6. Click **Apply** or **OK** to save changes.
7. Restart all nodes in the cluster to make the changes take effect.

8.5.1 Heartbeat

Heartbeat specifies the amount of time between transmits for all nodes in the cluster except the master. For example, if you set this value to 1, non-master nodes in the cluster send a signal that they are alive to the master node every second.

8.5.2 Tolerance

Tolerance specifies the amount of time the master node gives all other nodes in the cluster to signal that they are alive. For example, setting this value to 4 means that if the master node does not receive an “I’m alive” signal from a node in the cluster within four seconds, that node is removed from the cluster.

8.5.3 Master Watchdog

Master Watchdog specifies the amount of time between transmits for the master node in the cluster. For example, if you set this value to 1, the master node in the cluster transmits an “I’m alive” signal to all the other nodes in the cluster every second.

Modify this parameter setting only when supervised by Micro Focus Technical Support (http://www.novell.com/support).

8.5.4 Slave Watchdog

Slave Watchdog specifies the amount of time the master node has to signal that it is alive. For example, setting this value to 5 means that if the non-master nodes in the cluster do not receive an “I’m alive” signal from the master within five seconds, the master node is removed from the cluster and one of the other nodes becomes the master node.

Modify this parameter setting only when supervised by Micro Focus Technical Support (http://www.novell.com/support).

8.5.5 Maximum Retransmits

This value is set by default, and should not be changed.

8.6 Configuring Cluster Event Email Notification

OES Cluster Services can automatically send out email messages for certain cluster events like cluster and resource state changes or nodes joining or leaving the cluster. The subject line provides information about the cluster name, resource name, action taken, and node name. For example:

```
CL1: POOL1_SERVER online on NODE1
```

You can enable or disable email notification for the cluster and specify up to eight administrator email addresses for cluster notification.

- Section 8.6.1, “Configuring an Email Port,” on page 107
- Section 8.6.2, “Configuring the Notification Policies,” on page 107

8.6.1 Configuring an Email Port

OES Cluster Services uses Postfix to send cluster email alerts, using the primary IP address. Postfix uses port 25 by default. If you have a cluster resource that uses SMTP (which also defaults to port 25), the port conflict between Postfix and SMTP can prevent that resource from working in the cluster. To avoid a port conflict, you can change the Postfix port configuration, or you can bind the clustered service to its resource’s secondary IP address.

To configure Postfix to use a different port, you can edit the `/etc/postfix/main.cf` file and change the values for the `inet_interfaces`, `mydestination`, and `mynetworks_style` directives. You can change the listen port for the `smtpd` process in the `/etc/postfix/master.cf` file. For more information about configuring Postfix, see the Postfix website (http://www.postfix.org/documentation.html).

To bind a service to its resource's secondary IP address, refer to the cluster documentation for the service. For example, the GroupWise Internet Agent (GWIA) uses SMTP. Because both Postfix and the GWIA default to using port 25, you must configure the GWIA to bind exclusively to its resource's secondary IP address in order to avoid a port conflict between Postfix and the GWIA. For information about how to set up the secondary IP address and to bind the GWIA to it, see the following sections in the *GroupWise 2012 Interoperability Guide* (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/a20gkue.html):

- "Selecting the GWIA Partition and Secondary IP Address" (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/bwe3bwa.html#bwe3bwk)
- "Forcing Use of the GWIA Secondary IP Address" (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/bwe3bx1.html#bwe3byz)

8.6.2 Configuring the Notification Policies

1. In iManager, select **Clusters > My Clusters**.
2. In your personalized **My Clusters** list, select the check box next to the cluster you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.
3. Click **Action > Edit Properties** to open the Properties dialog box.
4. Click the **Policies** tab.
5. Under **Notification**, select or deselect the **Enable Cluster Notification Events** check box to enable or disable email notification.
6. If you enable email notification, configure the notification settings:

   6a Add up to eight email addresses in the field provided.

   You can click the buttons next to the field to add, delete, or edit email addresses. Repeat this process for each email address you want on the notification list.

   6b Specify the type of cluster events you want administrators to receive messages for.

   - **Only Critical Events:** (Default) Select this option to receive notification only of critical events like a node failure or a resource going comatose.
   - **Verbose Messages:** Select this option to receive notification of all cluster state changes including critical events, resource state changes, and nodes joining and leaving the cluster.

   6c Specify whether you want to receive notification of all cluster state changes in XML format by selecting the **XML Messages** option.

   XML is selected by default. XML format messages can be interpreted and formatted with a parser that lets you customize the message information for your specific needs.
7. Click **Apply** or **OK** to save changes.

8.7 Configuring Cascade Failover Prevention

A cascading failover occurs when a bad cluster resource causes a server to fail, then fails over to another server causing it to fail, and then continues failing over to and bringing down additional cluster servers until possibly all servers in the cluster have failed.

OES Cluster Services provides a Cascade Failover Prevention function. It detects if a node has failed because of a bad cluster resource, and prevents that bad resource from failing over to other servers in the cluster.

- Section 8.7.1, “Understanding the Cascade Failover Prevention Quarantine,” on page 108
- Section 8.7.2, “Releasing a Resource from Quarantine,” on page 109
- Section 8.7.3, “Enabling or Disabling Cascade Failover Prevention,” on page 110

8.7.1 Understanding the Cascade Failover Prevention Quarantine

The cascade failover prevention quarantine puts a resource in a comatose state rather than letting it load on (and potentially cause to fail) other cluster nodes. A resource can be quarantined if the systematic analysis of the logged node failures determines the following conditions to be true:

- If the resource is likely responsible for several consecutive node failures, unrelated to interference from failures of other resources.

  The consecutive failures might have occurred on different nodes in the cluster. If the resource loads successfully on any node, the failure count for the resource starts over.
- If loading the resource will put the cluster in grave danger.

OES Cluster Services does the following to determine if a resource should be put into quarantine:

1. Traces the history of node failures for the suspected bad resource. This includes:
   - What node the resource was running on or loading on
   - If the node failed
   - The state the resource was in when the node failed
   - If there were other resources trying to load when the node failed
2. Repeats the above process until one of the following happens:
   - The end of the cluster log file is reached
   - Enough consecutive node failures are found
   - Found that the node did not fail
   - Found that the whole cluster was down
   - The entries in the log file are more than 365 days old

If the resource attempts to load on a node where it was previously loaded and there are additional nodes still available in the cluster, it will not be quarantined and will be allowed to load. Also, a resource is not quarantined when it is initially brought online.

For a particular resource, there is no fixed number of node failures that triggers a quarantine. Generally, three consecutive node failures triggers a resource quarantine. However, the actual number of failures considered can vary based on other factors like how many nodes are in the

cluster and what other resources are doing at the time of the node failures. For example, if no other resources are in a running or loading state when a resource loads but never reaches a running state, then two consecutive node failures may trigger resource quarantine.

Factors that might contribute to a resource being quarantined include:

- A large number of consecutive node failures (generally, three or more)
- No other resources are causing node failures
- The resource never reaches a running state

Factors that might help prevent a resource from being quarantined include:

- A small number of consecutive node failures (generally, one or two)
- The resource has failed on this node previously
- Other resources are causing node failures
- The resource reaches a running state
- There is one node left up and running in the cluster

The resource quarantine is disabled if:

- Cascade Failover Prevention is turned off.

  See "Releasing a Resource from Quarantine" on page 109.
- There is no shared storage (SAN) or SBD partition.
- There are enough nodes in the cluster to form a quorum.

8.7.2 Releasing a Resource from Quarantine

While a resource is in quarantine, you can still manually take the resource from the comatose state to an offline state, and then bring it online or cluster migrate it to other cluster nodes.

To get the resource out of quarantine so that is once again able to automatically fail over:

1. Log in to the node as the `root` user.
2. Disable Cascade Failover Prevention. See "Disabling Cascade Failover Prevention" on page 110.
3. Re-enable Cascade Failover Prevention. See "Manually Enabling Cascade Failover Prevention" on page 111.

8.7.3 Enabling or Disabling Cascade Failover Prevention

The Cascade Failover Prevention function is enabled by default when you install OES Cluster Services. You can control Cascade Failover Prevention by creating a configuration file in the `/etc/modprobe.d` folder, and using it to disable and enable the `ncs_cascade_failover_detection_flag`. Two detection modes are supported:

- A value of 0 disables the function.
- A value of 2 enables the function. This value is assumed of the configuration file does not exist, or if the flag line is not present in the file. The file does not exist by default.

The setting in the file is specific to a node, and is set separately on each node. After you modify the setting on a node, you must manually unload and reload OES Cluster Services software on the server to apply change. The setting persists through patches and upgrades.

- "Disabling Cascade Failover Prevention" on page 110
- "Manually Enabling Cascade Failover Prevention" on page 111

Disabling Cascade Failover Prevention

You might need to disable the Cascade Failover Prevention function for the following reasons:

- To get a resource out of quarantine and allow it to once again automatically fail over, you can temporarily disable the Cascade Failover Prevention function on the node where the resource was taken offline.
- To stop using the Cascade Failover Prevention function in the cluster, you can disable the function on every node in the cluster.

To disable Cascade Failover Prevention:

1. Log in to the node as the `root` user.
2. Navigate to the `/etc/modprobe.d` folder.
3. In a text editor, create a configuration file (such as `novell-ncs.conf`) under the `/etc/modprobe.d` folder.

   If the file already exists, open it.
4. Add the following content to the file. If the line already exists, change the flag value to 0.

   ```
   options crm ncs_cascade_failover_detection_flag=0
   ```

   The flag value 0 disables Cascade Failover Prevention for the node.
5. Save the file, then close the text editor.
6. Open a terminal console, then restart OES Cluster Services software to apply the modified Cascade Failover Prevention setting on this node:

   ```
   rcnovell-ncs restart
   ```

   or

   ```
   systemctl restart novell-ncs.service
   ```
7. Repeat this procedure on other nodes where you want to disable Cascade Failover Prevention.
8. (Optional) Re-enable Cascade Failover Prevention as described in "Manually Enabling Cascade Failover Prevention" on page 111.

Manually Enabling Cascade Failover Prevention

You might need to manually re-enable Cascade Failover Prevention:

- If you disabled the function to release resources from the quarantine
- If you disabled the function for other reasons

To enable Cascade Failover Prevention:

1. Log in to the node as the `root` user.
2. Navigate to the `/etc/modprobe.d` folder.
3. Use one of the following methods to enable Cascade Failover Prevention:
   - **Delete the file:** Delete the `/etc/modprobe.d/novell-ncs.conf` file that you created in "Disabling Cascade Failover Prevention" on page 110.
   - **Change the value to 2:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in "Disabling Cascade Failover Prevention" on page 110, change the flag value from 0 to 2, then save the file.

     ```
     options crm ncs_cascade_failover_detection_flag=2
     ```

     The flag value 2 enables Cascade Failover Prevention for the node.
   - **Delete the line:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in "Disabling Cascade Failover Prevention" on page 110, remove the `ncs_cascade_failover_detection_flag` line, then save the file.
4. Open a terminal console, then restart OES Cluster Services software to apply the modified Cascade Failover Prevention setting on this node:

   ```
   rcnovell-ncs restart
   ```

   or

   ```
   systemctl restart novell-ncs.service
   ```
5. Repeat this procedure on other nodes where you want to enable Cascade Failover Prevention.

8.8 Configuring NCS to Monitor the eDirectory Daemon (ndsd)

OES Cluster Services requires eDirectory to be running when you manage the cluster and resources. At other times, NCS uses a file cache of information retrieved from eDirectory to cover temporary eDirectory or LDAP outages. This allows normal cluster operations to be less dependent on eDirectory than some clustered services that require the eDirectory daemon (`ndsd`) to be running in order to function properly. An eDirectory outage might not impact NCS, but it can render the services provided by cluster resources to be non-functional.

NCS provides the ability to monitor the status of the eDirectory daemon (ndsd) at the NCS level. It is disabled by default. The monitoring can be set independently on each node. It runs whenever NCS is running on the node. The settings persist across system restarts, patches, or upgrades.

IMPORTANT: If you enable NDSD monitoring at the NCS level, you should remove or comment out the eDirectory status checks in individual monitor scripts.

- Section 8.8.1, "Understanding eDirectory Monitoring," on page 112
- Section 8.8.2, "Displaying the Current Settings for NDSD Monitoring," on page 113
- Section 8.8.3, "Configuring or Modifying NDSD Monitoring," on page 113
- Section 8.8.4, "Disabling NDSD Monitoring," on page 114

8.8.1 Understanding eDirectory Monitoring

OES Cluster Services can monitor the eDirectory daemon for nodes where eDirectory is installed. You might want to monitor the status of ndsd if you cluster services that depend on the eDirectory daemon running in order to function properly, such as for logins of CIFS or NCP users. These service resources might run on only a subset of the nodes in the cluster or on all of the nodes. Each node's monitoring is configured separately.

Each monitoring iteration includes an interval of time to wait until NCS sends a request for status to ndsd, and a timeout period to listen for a response. You must specify the interval and timeout periods in seconds. The minimum interval is 16 seconds. The minimum timeout is 8 seconds. The following actions occur during the iteration:

1. The iteration begins.
2. NCS waits for the specified interval.
3. The interval ends, and NCS sends a status query to ndsd.
4. The timeout period begins, and NCS listens for a response.
5. Three actions are possible:
   - If the ndsd status returns within the timeout period and it is good, the iteration ends and monitoring continues.
   - If the ndsd status returns within the timeout period and it is not good, the specified remedy action is taken.
   - If the timeout period elapses without a response from ndsd, the specified remedy action is taken.
6. The monitoring continues with a new iteration.

If the ndsd status is bad or no status is returned within the timeout period, NCS can take one of three configurable remedy actions: an `ndsd` restart, a graceful node restart, or a hard node restart.

- **restart-ndsd:** Restarting the eDirectory daemon (`ndsd`) is the least intrusive action. However, it has the least chance of solving the problem. Even if the restart succeeds, the affected services might not be able to take advantage of the success.
- **reboot-node:** A graceful node restart attempts to synchronize the file systems and terminate all running processes. This option might not resolve a service's problem if a process is stuck, and it can delay a resource failing-over to a another node.
- **hard-reboot:** A hard node restart has a better chance of forcing resources to promptly fail over to other nodes. However, it does not synchronize the file systems, and clients can lose data that has not been committed to disks.

You use the `/opt/novell/ncs/install/ncs_install.py` script to enable, disable, or modify monitoring of the eDirectory daemon for a node. Issue the command from a console prompt as the `root` user.

```
/opt/novell/ncs/install/ncs_install.py
  -m <display|disable|restart-ndsd|reboot-node|hard-reboot|help>
  [-i <interval_value>]
  [-t <timeout_value>]]
```

For information about the command syntax and options, see Section A.8.2, "Monitoring NDSD," on page 396.

8.8.2 Displaying the Current Settings for NDSD Monitoring

1 Open a terminal console, then enter as the root user:

```
/opt/novell/ncs/install/ncs_install.py -m display
```

Example 1: Display current settings for ndsd monitoring, and monitoring is enabled.

```
# /opt/novell/ncs/install/ncs_install.py -m display
NDSD monitoring is enabled.
interval = 90 seconds, timeout = 20 seconds and remedy action is rebooting node hard.
```

Example 2: Display current settings for ndsd monitoring, and monitoring is disabled.

```
# /opt/novell/ncs/install/ncs_install.py -m display
NDSD monitoring is not enabled.
```

8.8.3 Configuring or Modifying NDSD Monitoring

1 Open a terminal console, then enter as the root user:

```
/opt/novell/ncs/install/ncs_install.py -m <hard-reboot|reboot-node|restart-ndsd> -i
<interval_value> -t <timeout_value>
```

Example 1: Enable ndsd monitoring with an interval of 60 seconds, timeout of 15 seconds, and remedy action of rebooting gracefully.

```
# /opt/novell/ncs/install/ncs_install.py -m reboot-node -i 60 -t 15

   NDSD monitoring is enabled.
   interval = 60 seconds, timeout = 15 seconds and remedy action is rebooting node.
```

Example 2: Change the remedy action to restarting ndsd.

```
# /opt/novell/ncs/install/ncs_install.py -m restart-ndsd -i 60 -t 15

   NDSD monitoring is enabled.
   interval = 60 seconds, timeout = 15 seconds and remedy action is restarting ndsd.
```

8.8.4 Disabling NDSD Monitoring

Monitoring of the eDirectory daemon is disabled by default. If it is enabled on a node, you can use the following procedure to disable it.

1 Open a terminal console, then enter as the root user:

```
/opt/novell/ncs/install/ncs_install.py -m disable
```

Example: Disable ndsd monitoring.

```
# /opt/novell/ncs/install/ncs_install.py -m disable
ndsd monitoring is not enabled (3).
```

8.9 Configuring the Cluster Node Reboot Behavior

If LAN connectivity is lost between a cluster node and the other nodes in the cluster, it is possible that the lost node will be automatically shut down by the other cluster nodes. This is normal cluster operating behavior, and it prevents the lost node from trying to load cluster resources because it cannot detect the other cluster nodes. The OES Cluster Services reboot behavior conforms to the

kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown. On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems.

To control the cluster node reboot behavior, you can use the directive `kernel.panic` in the Linux `/etc/sysctl.conf` file to prevent a reboot, or to allow an automatic reboot and to specify the number of seconds to delay the reboot. Set the `kernel.panic` value to 0 for no reboot after a node shutdown. Set the `kernel.panic` value to a positive integer value to indicate that the server should be rebooted after waiting the specified number of seconds. For information about using the Linux `sysctl`, see the Linux man pages on `sysctl(8)` and `sysctl.conf(5)`.

1. As the `root` user, open the `/etc/sysctl.conf` file in a text editor.
2. If the `kernel.panic` token is not present, add it.

   ```
   kernel.panic = 0
   ```

3. Set the `kernel.panic` value to 0 or to a positive integer value, depending on the desired behavior.
   - **No Reboot:** To prevent an automatic cluster reboot after a node shutdown, set the `kernel.panic` token to value to `0`. This allows the administrator to determine what caused the kernel panic condition before manually rebooting the server. This is the recommended setting.

     ```
     kernel.panic = 0
     ```

   - **Reboot:** To allow a cluster node to reboot automatically after a node shutdown, set the `kernel.panic` token to a positive integer value that represents the seconds to delay the reboot.

     ```
     kernel.panic = <seconds>
     ```

     For example, to wait 1 minute (60 seconds) before rebooting the server, specify the following:

     ```
     kernel.panic = 60
     ```

4. Save your changes.

8.10 Configuring STONITH

The STONITH (shoot-the-other-node-in-the-head) capability allows OES Cluster Services to remotely kill a suspect node by using remote power control instead of using a poison pill. STONITH does not require any action from the node being killed, unlike poison pills, which allows it to kill non-responsive nodes.

Using STONITH requires that you have server power management technology for all nodes in the cluster. STONITH supports remote accessible cards integrated in a cluster node's hardware, such as Integrated Lights Out (iLO) from Hewlett-Packard (HP) and Dell Remote Access Card (DRAC) from Dell, and stand-alone web-based power switches. Refer to the vendor documentation to determine how your power management system works.

To use STONITH in OES Cluster Services, you must create an executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` script file to authenticate to your power controller and turn off the power, cycle the power, or reset the power for the node. The script should take the node number as the

only parameter. Node numbers are assigned as 0 to 31 for nodes 1 to 32. These are the same node numbers that appear in the `/var/opt/novell/ncs/gipc.conf` file. Creating the script file automatically enables STONITH for the node; you do not need to restart anything.

IMPORTANT: STONITH does not replace poison pills in OES Cluster Services. OES Cluster Services issues poison pills before running the STONITH script.

Use the following sample scripts as a guide for how to create a script that works with your power management technology:

- Section 8.10.1, “Sample Script for HP iLO,” on page 115
- Section 8.10.2, “Sample Script for Web-Based Power Switches,” on page 116

8.10.1 Sample Script for HP iLO

This section provides a sample script for HP iLO power management cards. The sample code assumes the following setup for the cluster and iLO cards:

- The iLO card on each node has been pre-configured to trust any instructions sent from each of the nodes in the same cluster.
- Each cluster node’s iLO card is assigned a sequential static IP address, beginning with 192.168.188.201 on node 0, 192.168.188.202 on node 2, and so on up to 192.168.188.232 on node 32.

  If you alternatively use DNS names in your script, the translation must be performed by the script.
- The iLO card’s command to reset power is `power reset`.

  Refer to the HP documentation to determine the commands available for your iLO cards.

For each node in the cluster, create the executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` file and add the script below. The presence of the file automatically enables STONITH for the cluster.

IMPORTANT: Ensure that you replace the sample information with the settings for your system.

```
#! /bin/bash

if [ -n "$1" ]; then
    echo "Recycling power of node number $1 ... "
    iloIP=$(printf "192.168.188.2%02d" $(($1+1)))
    nodeIP=`grep "^nodeid .* ${1}\>" /var/opt/novell/ncs/gipc.conf | cut -d' ' -f2`
    while [ 1 ]; do
        ssh ${iloIP} power reset
        echo
        sleep 2
        if ! ping -c 1 ${nodeIP} | grep "1 received" &> /dev/null
        then
            break
        fi
        sleep 5
    done
fi
```

8.10.2 Sample Script for Web-Based Power Switches

This section provides a sample script for web-based power switches. The sample code assumes the following setup for the cluster and power switch:

- Each cluster node is assigned a sequential static IP address, beginning with 10.10.189.100.
- Each cluster node is plugged in to a sequential outlet in the power switch, beginning with node 0 in the first outlet.
- The authentication information for the power switch management interface is:
  - **User name:** admin
  - **Password:** novell
  - **IP address:** 10.10.189.149
- The power management switch's command to cycle power is CCL.

  Each vendor can have different command options available. Refer to your vendor documentation to determine the commands used by your power switch.

For each node in the cluster, create the executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` file and add the script below. The presence of the file automatically enables STONITH for the cluster.

IMPORTANT: Ensure that you replace the sample information with the settings for your system.

```
#! /bin/bash

until [ -z "$1" ]
do
    nodeNum=`expr $1 + 1`
    echo "Recycling power of node number $nodeNum ... "
    while [ 1 ]; do
        curl -u admin:novell http://10.10.189.149/outlet?${nodeNum}=CCL
        echo
        sleep 2
        if ! ping -c 1 10.10.189.10${nodeNum} | grep "1 received" &> /dev/null
        then
            break
        fi
        sleep 5
    done

    shift
done
```

8.11 Viewing or Modifying the NCS Proxy User Assignment in the NCS Management Group

When you configure a cluster, the NCS Management Group (<*cluster_name*>_MGT_GRP) is created in the Cluster object container. The group contains the user that you specify for the NCS Proxy User on the Proxy User Configuration page as described in Step 3 in Section 5.5.5, "Configuring a New Cluster," on page 67:

- OES Common Proxy User

- LDAP Admin User
- Another administrator user

The specified user is automatically added as a member of the group. If you specify the OES Common Proxy User, each nodes' OES Common Proxy User is added to the group.

If the OES Common Proxy User is later disabled in eDirectory for a node, the LDAP Admin user automatically replaces that user in the group.

You can view a list of members in the <*cluster_name*>_MGT_GRP in iManager. Select the **Directory Administration > Modify Object** function in iManager, then search for the group.

IMPORTANT: Do not attempt use the Modify Object page to add or remove members of the <*cluster_name*>_MGT_GRP group. You must use the `/opt/novell/ncs/install/ncs_install.py` script to add or remove members.

You can modify the users assigned as the NCS Proxy User by using the `/opt/novell/ncs/install/ncs_install.py` script. You are prompted for the distinguished user name and password for the user that you want to assign. The specified user is added to the group, and the old user is removed from the group.

1. Log in to the server as the `root` user, then open a terminal console.
2. At the command prompt, enter

   ```
   /opt/novell/ncs/install/ncs_install.py -l -u
   ```

   In addition, you can change other configurations by providing a configuration file (such as `/root/ncs.conf`) to the above command. You must create the file and its content before issuing the command. For example:

   ```
   /opt/novell/ncs/install/ncs_install.py -l -f /root/ncs.conf -u
   ```
3. When you are prompted, provide the comma-delimited distinguished name of the user to be assigned to the NCS_Management group, such as

   ```
   cn=ncs_admin,o=novell
   ```
4. When you are prompted, provide the password for the user you specified.

8.12 Viewing or Modifying the Cluster Master IP Address or Port

The cluster IP address is assigned to the Master IP Address Resource of the cluster when you configure the first node in a OES Cluster Services cluster. The cluster IP address is separate from the server IP address and is required for some external programs to get cluster status alerts. The IP address is bound to the master node and remains with the master node regardless of which server is the master node.

The cluster IP address normally does not need to be changed, but it can be changed for any reason. For example:

- A duplicate IP address conflicts with the master IP address.

- You want to move the cluster to a different IP subnet.
- You want to reallocate IP addresses in the same subnet.

You normally connect to the master IP address to manage the cluster. If the master IP address has connection problems, the **Clusters > My Clusters** page in iManager allows you to use the server IP address of a node (preferably the master node) to manage the cluster and fix the problem. For example, the following connection problems can prevent you from connecting to the master IP address to manage the cluster:

- Another device on the network is using the master IP address, and the Master IP Resource cannot come online.
- There are underlying communications errors for the SFCB protocol, and you get an error in iManager when you try to connect to the Master IP Resource of the cluster. For example errors, see "iManager errors while trying to manage Cluster, Storage, User Quotas, CIFS or AFP" (Technical Information Document 7010295) (http://www.novell.com/support/kb/doc.php?id=7010295).

Connecting to a node's IP address is a temporary measure to allow you to resolve the master IP address problem. It is not intended as an alternative way to manage the cluster. The option is not available from the My Resources page.

The default cluster port number is 7023, and is automatically assigned when the cluster is created. The cluster port number does not need to be changed unless a conflict is created by another resource using the same port number. If there is a port number conflict, change the Port number to any other value that doesn't cause a conflict.

IMPORTANT: After you change the cluster IP address or port number, you must restart Cluster Services, exit iManager, and relaunch iManager before iManager can manage the cluster with the new IP address.

---

1. In iManager, select **Clusters > My Clusters**.
2. In your personalized **My Clusters** list, view the status of the cluster you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.
3. If the master IP address does not respond to the status query, an error message prompts you to temporarily manage the server with the server IP address of the master node (preferred) or of any active node.

   The message appears each time that you revisit the My Clusters page until you resolve the IP address conflict on the network or by modifying the cluster's master IP address.

   3a Read the error message, then do one of the following:

   - **Use a server node (preferably the master):** Click **Yes** to agree to use the server IP address of a node to temporarily manage the cluster. Continue to Step 3b.
   - **View error details:** Click **No** to view the **Cluster Communication Error** details. It provides some tips for diagnosing the root cause of the IP address conflict. Click **OK** to dismiss the message. Try to resolve the conflict on the network, then select **Clusters > My Clusters** (Step 1) to see if the connection issue is resolved.

   3b Browse to choose the node (preferably the master node) to use to temporarily manage the cluster, then click **OK**.

4 After the clusters connection is made, select the check box next to the cluster, then click **Action > Edit Properties** to open the Cluster Properties dialog box.

5 In the Cluster Properties dialog box, click the **Policies** tab.

Normally, the cluster name is displayed if you access the cluster properties via the Master_IP_Resource. The server name (such as `avalon.example`, shown below) is displayed if you access the properties via a server node in Step 3.

6 Under **Cluster Information**, specify the new value for the IP address or port.

7 Click **OK** to save and apply your changes. Wait for the success notification before continuing.

A cluster restart is required before iManager can continue to manage the cluster with its new master IP address.

8 Exit iManager and close the web browser.

9 Restart Cluster Services. Open a terminal console as the `root` user, then enter

```
rcnovell-ncs restart
```

or

```
systemctl restart novell-ncs.service
```

```
avalon:~/Desktop # rcnovell-ncs restart
Stopping Novell Cluster Services
Leaving...
No longer a member of cluster

#---------------------------------------------------------------+
#                                                               |
#   Please wait for the Novell Cluster Services KOs to unload   |
#                                                               |
#---------------------------------------------------------------+

Novell Cluster Services Unloaded

                                                                done

Starting Novell Cluster Services
Mounting adminfs at /admin ... already mounted                  done
                                                                done

avalon:~/Desktop # Joining...
Now a member of

        Cluster  clus1
        This node avalon [ epoch 0 master node avalon ]
        Cluster nodes [ avalon ]
```

10 Launch a web browser and log in to iManager to continue managing the cluster.

8.13 Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster

Use the instructions in this section to change the IP address of the cluster, information about the LDAP servers that the cluster uses, or the credentials used to administer the cluster.

- Section 8.13.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 120
- Section 8.13.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 122

8.13.1 Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster

You can modify the administrator credentials or LDAP server settings that you assigned when you created the cluster. If eDirectory is not installed on a node, it looks to the LDAP server list for information about which LDAP server to use.

You must modify this cluster information in the following cases:

- Changing the Administrator user name and password for the cluster
- Changing the password for the existing Administrator user name
- Changing the IP address information about the existing LDAP servers
- Assigning a different LDAP server for the cluster to use
- Changing the order of the servers in the LDAP server list

  As a best practice, you should list the LDAP servers in the following order:

  - Local to the cluster
  - Closest physical read/write replica
- Adding LDAP servers to the list of ones that the cluster can use.

You can modify these settings at any time. OES Cluster Services can be running or not running.

To modify the LDAP server IP address or administrator credentials in the OES Cluster Services configuration settings:

1. Ensure that the administrator user name that you plan to use meets the requirements specified in Section 4.1, “Cluster Administration Requirements,” on page 35.
2. Ensure that the LDAP servers that you want to specify for the cluster are in the same eDirectory tree, and that the IP addresses meet the requirements in Section 4.2, “IP Address Requirements,” on page 38.
3. Log in as the `root` user to the master node of the cluster.
4. In a text editor, create a text file, specify the configuration information for the OES Cluster Services cluster in it, then save the file.

   An example is shown below of the content of the file with sample values. The directives are self-explanatory.

   IMPORTANT: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

   The following lines are the content of a sample configuration file for a OES Cluster Services cluster when you have multiple LDAP servers.

   ```
   CONFIG_NCS_CLUSTER_DN="cn=svr1_oes2_cluster.o=context"
   CONFIG_NCS_LDAP_INFO="ldaps://10.1.1.102:636,ldaps://10.1.1.101:636"
   CONFIG_NCS_ADMIN_DN="cn=admin.o=context"
   CONFIG_NCS_ADMIN_PASSWORD="password"
   ```

5. As the `root` user, enter the following command at a command prompt:

   ```
   /opt/novell/ncs/install/ncs_install.py -l -f configuration_filename
   ```

Replace *configuration_filename* with the actual name of the file you created.

6 Delete the configuration file that you created.

7 For each of the other nodes in the cluster, log in as the `root` user for the node, then repeat Step 4 to Step 6.

Modifying the information on each node allows iManager to manage the cluster after a different node becomes the master. This step is necessary because credentials are stored on OES Credential Store, and OES Credential Store does not synchronize across cluster nodes.

8 Push this update to all nodes on the cluster by entering the following as the `root` user on one of the cluster nodes:

```
cluster exec "/opt/novell/ncs/bin/ncs-configd.py -init"
```

8.13.2 Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources

If you move a cluster to a different subnet, you must change the IP addresses of the cluster nodes and the cluster resources, information about the LDAP servers used by the cluster, and possibly the administrator credentials for the cluster.

When you move the cluster to a new IP subnet, you must replace the existing unique static IP addresses with ones that are valid in that subnet. You can make the IP address changes in the old location or the new location. If you start the servers in the different IP subnet with the old IP addresses, the cluster does not come up until you make the changes described in this section.

To modify the IP addresses of servers being used in a OES Cluster Services cluster, perform the following tasks in the order given:

- "Prerequisites" on page 122
- "Changing the IP Addresses of Cluster Resources" on page 122
- "Changing the IP Addresses of Servers in a Cluster" on page 124
- "Changing the IP Address of the Master IP Resource" on page 124
- "Modifying the Cluster Configuration Information" on page 124

Prerequisites

Before you begin, ensure that the IP addresses that you plan to use meet the requirements specified in Section 4.2, "IP Address Requirements," on page 38. Ensure that the administrator user name that you will use in the new location has sufficient rights as described in Section 4.1.3, "Cluster Administrator or Administrator-Equivalent User," on page 38.

Changing the IP Addresses of Cluster Resources

Before you modify the server IP address for a server in a cluster, you must change the IP addresses of all of the cluster resources that run on it:

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Offline the cluster resources whose IP addresses are changing. On the Cluster Manager page, select the check boxes next to the cluster resources, then click **Offline**.

4 For each NSS pool cluster resource, modify the IP address on the resource's Protocols page.

4a On the Cluster Manager or Cluster Options page, click the pool cluster resource to open its Properties page.

4b Click the **Protocols** tab.

4c Modify the IP address for the NSS pool cluster resource.

4d Click **OK** to save the changes.

This changes the IP address for the `NCS:NCP Server` object for the resource, and updates the instances of the IP address in the resource's load, unload, and monitor script.

Do not online the resource at this time.

5 For each non-NSS cluster resource, modify the IP address information for the RESOURCE_IP variable as needed in the resource's load, unload, and monitor scripts.

5a On the Cluster Manager page or Cluster Options page, click the non-NSS cluster resource to open its Properties page.

5b Click the **Scripts** tab, then click the **Load Script** link.

5c Edit the script by replacing the RESOURCE_IP variable's value with the new IP address. You might also need to edit the values used in the command lines if you did not use the variable there.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, volume group name, file system type, and mount device.

5d Click **Apply** to save the changed script.

5e Make similar changes to the non-NSS resource's **Unload Script** and **Monitor Script**.

5f Repeat this process for each non-NSS cluster resource.

Do not online the resources at this time.

6 For each Linux cluster resource that has an NCP virtual server object, delete and re-create the NCS:NCP Server Object.

6a In iManager, select the **Directory Administration > Delete Object**, select the NCS:NCP Server object for the Linux resource, then click **OK**.

6b Re-create the NCS:NCP Server object with the new IP address:

6b1 On the master cluster node, open a terminal console, then log in as the `root` user.

6b2 In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.

6b3 At the terminal console prompt, enter

```
./ncs_ncpserv.py -c lv_volumename -i resource_ip_address
```

Replace the *lv_volumename* and *resource_ip_address* with the information for your particular solution.

Do not use periods in cluster resource names. Client for Open Enterprise Servers interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

For example, to create the NCS:NCP Server object for the `lv44` cluster resource where the IP address is 10.10.10.44 and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

```
./ncs_ncpserv.py -c lv44 -i 10.10.10.44
```

The confirmation message is displayed:

```
NCP Server
'cn=cluster1_lv44_server,ou=clusters,ou=city,o=mycompany'
created.
```

7 Stop OES Cluster Services for every node in the cluster by entering the following at the command prompt as the `root` user:

```
rcnovell-ncs stop
```

or

```
systemctl stop novell-ncs.service
```

8 Continue with "Changing the IP Addresses of Servers in a Cluster" on page 124.

Changing the IP Addresses of Servers in a Cluster

After preparing the cluster resources for the IP address change and stopping OES Cluster Services (see "Changing the IP Addresses of Cluster Resources" on page 122), you are ready to change the IP addresses of the servers in the cluster.

1 For each server in the cluster, change the server's IP address by following the instructions "Changing an OES Server's IP Address" in the *OES 2023: Planning and Implementation Guide*.

2 The server IP address changes are not complete until you make those changes known to OES Cluster Services and eDirectory. Continue with "Changing the IP Address of the Master IP Resource" on page 124.

Changing the IP Address of the Master IP Resource

1 In iManager, click **Directory Administration**, then click **Modify Object**.

2 Browse to locate and select the Master IP Address Resource object of the cluster you want to manage.

3 In the **Valued Attributes** list, select the attribute **NCS:CRM Load Script**, click **Edit**, modify the IP address in the script, then click **OK**.

4 In the **Valued Attributes** list, select the attribute **NCS:CRM Unload Script**, click **Edit**, modify the IP address in the script, then click **OK**.

5 In the **Valued Attributes** list, select the attribute **NCS:CRM Monitor Script**, click **Edit**, modify the IP address in the script, then click **OK**.

6 At the bottom of the page, click **OK** to save the changes.

7 Continue with "Modifying the Cluster Configuration Information" on page 124.

Modifying the Cluster Configuration Information

Before restarting OES Cluster Services, you must update the cluster configuration information in OES Cluster Services and eDirectory with the new IP addresses. You might also need to update the IP address information for the LDAP server and administrator credentials that the cluster uses in the new subnet.

1 If the cluster is using a different LDAP server or administrator in the new IP subnet, change the LDAP server IP address and administrator credentials for the cluster in the OES Cluster Services configuration settings.

   Follow the procedure in Section 8.13.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 120.

2 For each node in the cluster, including that of the master IP address resource, modify the **NCS: Network Address** attribute of its Cluster Node object.

   2a In iManager, select **Directory Administration > Modify Object**.

   2b Browse to locate and select the Cluster Node object of the cluster node you want to manage.

   2c In the **Valued Attributes** list, select the attribute **NCS: Network Address**, click **Edit**, modify the IP address, then click **OK**.

   2d Repeat this process for each node in the cluster and the master IP resource.

3 For the cluster container of the cluster you want to manage, modify the **NCS: Network Address** and **Network Address** attributes of its Cluster object to specify the new IP address information. Both **TCP** and **UDP** addresses need to be replaced.

   3a In iManager, select **Directory Administration > Modify Object**.

   3b Browse to locate and select the Cluster object of the cluster you want to manage.

   3c In the **Valued Attributes** list, select the attribute **NCS: Network Address** (the attribute for the TCP address), click **Edit**, modify the IP address, then click **OK**.

   3d In the **Valued Attributes** list, select the attribute **Network Address** (the attribute for the UDP address), click **Edit**, modify the IP address, then click **OK**.

4 Ensure that the LDAP server is running before restarting OES Cluster Services.

   IMPORTANT: OES Cluster Services requires LDAP.

5 Ensure that NSS is running if there are NSS cluster resources that you will bring online.

6 Start OES Cluster Services by entering the following command at a command prompt as the root user:

```
rcnovell-ncs start
```

   or

```
systemctl start novell-ncs.service
```

7 Online the cluster resources:

7a In iManager, select **Clusters > My Clusters**.

7b Select the cluster that you want to manage.

7c On the Cluster Manager page, select the check boxes next to the cluster resources, then click **Online**.

8.14 What's Next

For information about managing the cluster, see Chapter 9, "Managing Clusters," on page 127.

After installing and configuring the cluster, you are ready to configure cluster resources for it. See the following:

- Chapter 10, "Configuring and Managing Cluster Resources," on page 163
- Chapter 11, "Quick Reference for Clustering Services and Data," on page 197
- Chapter 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 201
- Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 275

9 Managing Clusters

After you have installed, set up, and configured OES Cluster Services on Open Enterprise (OES) servers and configured cluster resources, use the information in this section to help you effectively manage your cluster. This section provides instructions for migrating resources, identifying cluster and resource states, and customizing cluster management.

IMPORTANT: For information about using console commands to manage a cluster, see Appendix A, “Console Commands for OES Cluster Services,” on page 369.

- Section 9.1, “Starting and Stopping OES Cluster Services,” on page 127
- Section 9.2, “Taking the Cluster Down,” on page 130
- Section 9.3, “Selecting a Cluster to Manage,” on page 130
- Section 9.4, “Monitoring Cluster and Resource States,” on page 131
- Section 9.5, “Responding to a Start, Failover, or Failback Alert,” on page 134
- Section 9.6, “Viewing the Cluster Event Log in iManager,” on page 135
- Section 9.7, “Viewing the Cluster Event Log at the Command Line,” on page 136
- Section 9.8, “Viewing Summaries of Failed or Incomplete Events from the *.out Log Files,” on page 137
- Section 9.9, “Generating a Cluster Configuration Report,” on page 138
- Section 9.10, “Repairing a Cluster,” on page 139
- Section 9.11, “Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node,” on page 140
- Section 9.12, “Cluster Migrating Resources to Different Nodes,” on page 141
- Section 9.13, “Removing (Leaving) a Node from the Cluster,” on page 143
- Section 9.14, “Joining a Node to the Cluster (Rejoining a Cluster),” on page 144
- Section 9.15, “Shutting Down OES Cluster Services When Servicing Shared Storage,” on page 145
- Section 9.16, “Enabling or Disabling Cluster Maintenance Mode,” on page 145
- Section 9.17, “Viewing the Cluster Node Properties,” on page 145
- Section 9.18, “Creating or Deleting Cluster SBD Partitions,” on page 146
- Section 9.19, “Customizing Cluster Services Management,” on page 160

9.1 Starting and Stopping OES Cluster Services

OES Cluster Services automatically starts after it is installed. It also automatically starts when you reboot your OES server.

IMPORTANT: If you are using iSCSI for shared disk system access, ensure that you have configured iSCSI initiators and targets to start prior to starting Cluster Services. You can do this by entering the following at the Linux terminal console:

```
systemctl enable iscsi.service
```

9.1.1 Starting OES Cluster Services

If you stop OES Cluster Services, you can restart it by doing the following:

1. Open a terminal console, then log in as the `root` user.
2. To start Cluster Services, at the command prompt, enter

   ```
   rcnovell-ncs start
   ```

   Or

   ```
   systemctl start novell-ncs.service
   ```

3. To verify the cluster service status, at the command prompt, enter

   ```
   rcnovell-ncs status
   ```

   Or

   ```
   systemctl status novell-ncs.service
   ```

4. Verify the cluster status:

   ```
   cluster status
   ```

9.1.2 Stopping OES Cluster Services

OES Cluster Services is a kernel-space application. As with any kernel-space application, unloading it is a best effort and is not guaranteed to stop the modules. If Cluster Services is busy providing services from the kernel when you attempt to stop it, the process might not stop, and a reboot might be required.

1. Open a terminal console, then log in as the `root` user.
2. Check whether the node is active in the cluster. Enter

   ```
   cluster view
   ```

   If the node is in the cluster (and not still in a process of joining the cluster), it is relatively safe to attempt to stop it.

3. To stop Cluster Services, at the command prompt, enter

   ```
   rcnovell-ncs stop
   ```

   or

   ```
   systemctl stop novell-ncs.service
   ```

4. To verify the cluster service status, at the command prompt, enter

```
rcnovell-ncs status
```

Or

```
systemctl status novell-ncs.service
```

5 If you get in a state where the node has not joined the cluster and Cluster Services cannot be stopped:

5a Disable autostart on reboot. Enter

```
systemctl disable novell-ncs.service
```

5b Reboot the server.

5c Start Cluster Services manually. Enter

```
rcnovell-ncs start
```

or

```
systemctl start novell-ncs.service
```

5d Diagnose why the node had trouble joining the cluster.

5e (Optional) Re-enable autostart for the node. Enter

```
systemctl enable novell-ncs.service
```

9.1.3 Enabling and Disabling the Automatic Start of OES Cluster Services

OES Cluster Services automatically starts by default after it is installed and on server reboot.

To cause OES Cluster Services to not start automatically after a server reboot:

1 Open a terminal console, then log in as the `root` user.

2 Enter the following at a Linux terminal console:

```
systemctl disable novell-ncs.service
```

3 Reboot the server.

4 After rebooting, you must manually start Cluster Services by entering

```
rcnovell-ncs start
```

or

```
systemctl start novell-ncs.service
```

To cause OES Cluster Services to resume starting automatically after a server reboot:

1 Open a terminal console, then log in as the `root` user.

2 Enter the following at a Linux terminal console:

```
systemctl enable novell-ncs.service
```

3 Reboot the server.

9.2 Taking the Cluster Down

The `cluster down` command concurrently removes all cluster nodes from the cluster. This command has the same effect as executing the `cluster leave` command on every server in the cluster. You are prompted to confirm the cluster down. For example, you might need to take the cluster down if an SBD partition fails, and you need to delete and recreate an SBD partition.

1 Log in as the `root` user on any node in the cluster, then enter the following at a command prompt:

```
cluster down
```

2 When you are prompted to confirm the `cluster down` command, specify **Yes**, then press Enter.

9.3 Selecting a Cluster to Manage

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Clusters**.

The My Clusters page displays your personalized list of clusters in the tree that you want to manage. To add or remove clusters from your list, see Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

The Clusters plug-in sends a status query to the master IP address of each cluster in the list.

3 If the master IP address of a cluster has connection problems, the Clusters plug-in allows you to use the server IP address of a node (preferably the master node) to manage the cluster and fix the problem.

On the My Clusters page, if the cluster status query fails:

3a Read the message about the problem connecting to the master IP address.

The message appears each time that you revisit the My Clusters page until you resolve the IP address conflict on the network or by modifying the cluster’s master IP address.

3b Do one of the following:

- **Use a server node (preferably the master):** Click **Yes** to agree to use the server IP address of a node to temporarily manage the cluster.
- **View error details:** Click **No** to view the **Cluster Communication Error** details. It provides some tips for diagnosing the root cause of the IP address conflict. Click **OK** to dismiss the message. Try to resolve the conflict on the network, then select **Clusters > My Clusters** to see if the connection issue is resolved.

3c Browse to choose the node (preferably the master node) to use to temporarily manage the cluster, then click **OK**.

3d If you need to modify the master IP address to fix the connection problem, see Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118.

4 After a successful connection to the cluster master IP address, do any of the following:

- **Cluster Properties:** To view or manage the cluster policies, protocols, or properties, select the check box next to the cluster, then select **Action > Edit Properties**.

The Cluster Properties dialog box opens. The browser's pop-up blocker must be disabled. See the following:

- Section 8.4, "Configuring Quorum Membership and Timeout Policies," on page 104
- Section 8.5, "Configuring Cluster Protocols," on page 105
- Section 8.6, "Configuring Cluster Event Email Notification," on page 106

- **Resources:** To manage cluster resources, click the cluster name link.

  The browser opens on the Cluster Manager page in the same browser window. See Chapter 10, "Configuring and Managing Cluster Resources," on page 163.

- **Cluster Report:** To generate a report for the cluster, select the check box next to the cluster, then select **Action > Run Report**.

  The Cluster Report dialog box opens. The browser's pop-up blocker must be disabled. See Section 9.9, "Generating a Cluster Configuration Report," on page 138.

- **Repair an OES Cluster:** To repair an OES cluster, select the check box next to the cluster, then click **Action** > **Repair**.

  The Repair option is useful when you see a mismatch in the cluster resource list or the resource priority list. The cluster repair will make the resource list consistent with the resource objects.

  For example, if the cluster status command output is blank or the deleted resources are still listed, the repair option will make the resource list match the resource objects. Resources will then be displayed properly in CLI and iManager.

  NCS logs the corresponding repair option messages in the `/var/opt/novell/log/ncs/repair.log` file.

9.4 Monitoring Cluster and Resource States

The **Cluster Manager** page in iManager gives you important information about the status of servers and resources in your cluster.

1. In iManager, select **Clusters > My Clusters**.

   View the overall status of the each of the clusters in your list.

2. In your personalized My Clusters list, select the name of the cluster you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

3. View the status of the cluster nodes and resources.

   If scrolling is needed to view the complete list, you can use the **Refresh** option to set a longer refresh rate for the page in order to allow enough time to view the status for all items.

   Master Server

   The master server in the cluster is identified by a yellow diamond in the middle of the server icon ( ). The master server is initially the first server in the cluster, but another server can become the master if the first server fails.

Operating State

An icon next to the cluster resource or cluster indicates its operating state. Table 9-1 describes each operating state and icon. When a resource is state is blank or has no colored icon, it is unassigned, offline, changing state, or in the process of loading or unloading.

Table 9-1 *Operating States for a Cluster and Its Resources*

State	Icon	Description
Normal		A green ball indicates that the node is running, or the resource is online.
Stopped		A red ball with a horizontal white line indicates that the node or resource is stopped and waiting for administrator intervention.
Offline		A white ball with a horizontal red line indicates that the node or resource is offline.
Critical		A white ball with a red X indicates that the node has failed, the resource is comatose, or a resource in a My Resources list cannot be not found (renamed or deleted).
Warning		A white ball with a yellow diamond indicates that an alert condition has occurred, and the resource needs administrator attention.
Unknown		A yellow ball with a question mark indicates that the state of the node or resource cannot be determined. Either the server is not currently a member of the cluster or its state is unknown.

Epoch

The Epoch number indicates the number of times the cluster state has changed. The cluster state changes every time a server joins or leaves the cluster.

Cluster Resource States

Table 9-2 identifies the different resource states and gives descriptions and possible actions for each state.

Table 9-2 *Cluster Resource States*

Resource State	Description	Possible Actions
Alert Alert types are: Failback Alert Failover Alert Start Alert	Alerts can result if the Start mode, Failover mode, or Failback mode for the resource has been set to Manual. The resource is waiting for the administrator to intervene to start, fail over, or fail back the resource on the specified server.	On the Cluster Manager page in iManager, select the check box next to the resource, click the **Respond to Alert** option, then resolve the alert. Depending on the resource state, you are prompted to start, fail over, or fail back the resource. See Section 9.5, "Responding to a Start, Failover, or Failback Alert," on page 134. After you clear the **Start Alert**, you can take the resource offline, then bring the resource online.
Comatose	The resource is not running properly and requires administrator intervention.	Click the **Comatose** status indicator and take the resource offline. After resource problems have been resolved, you can bring the resource online (return it to the running state). A comatose resource cannot be taken directly to an online state.
Loading	The resource is in the process of loading on a server.	None.
NDS_Sync	The properties of the resource have changed and the changes are still being synchronized in eDirectory.	None.
Offline	Offline status indicates the resource is shut down or is in a dormant or inactive state.	If desired, you can click the **Offline** status indicator and click the **Online** button to load the resource on the best node possible, given the current state of the cluster and the resource's preferred nodes list.
Quorum Wait	The resource is waiting for the quorum to be established so it can begin loading.	None.
Running	The resource is in a normal running state.	You can click the **Running** status indicator and choose to either migrate the resource to a different server in your cluster or unload (bring offline) the resource.

Resource State	Description	Possible Actions
Unassigned	There is not an assigned node available that the resource can be loaded on.	If desired, you can click the **Unassigned** status indicator and offline the resource. Offlining the resource prevents it from running on any of its preferred nodes if any of them join the cluster.
Unloading	The resource is in the process of unloading from the server it was running on.	None.

9.5 Responding to a Start, Failover, or Failback Alert

You might receive an alert ( ) for a resource if you have configured the resource's Start mode, Failover mode, or Failback mode settings to Manual. The alert types correspond to the mode setting:

- Start Alert
- Failover Alert
- Failback Alert

The resource is waiting for the administrator to intervene to start, fail over, or fail back the resource on the specified server.

If you attempt to offline a resource that is in the **Start Alert** state, nothing happens and the following message is displayed:

```
This operation cannot be completed. It is only available when the resource
is in an offline state.
```

You must clear the **Start Alert** before you can offline the resource. After the resource is offline, you can online the resource.

To respond to an alert:

1. In iManager, select **Clusters > My Clusters**.

   If the cluster you want to manage does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

2. In your personalized list of clusters, a critical ( ) status for a cluster indicates that a resource might be in a comatose state and require administrator intervention. Click the resource name to manage the resource.
3. On the Cluster Manager page, select the check box next to the resource that has the alert, then click **Respond to Alert**.
4. On the Respond to Alert page, click **Yes**, **No**, or **Cancel** to perform the desired action for the selected resource's alert state.
5. After clearing a **Start Alert**, go to the Cluster Manager page, select the check box next to the resource, then click **Offline**. After the resource is offline, select the check box again and click **Online** to bring the resource online.

9.6 Viewing the Cluster Event Log in iManager

The Cluster Event Log in Clusters-plug for iManager lets you view recent cluster events logged in the `/admin/Novell/Cluster/EventLog.xml` file. Events can be node specific (a node joins the cluster or leaves the cluster) or they can be resource specific (a resource comes online or goes offline).

By default, the Cluster Event Log displays 10 most recent events (newest to oldest) for all severity levels, all nodes, and all resources. You can use the Cluster Event Log Filter page to filter out events that are not of interest so that only the desired log entries are displayed in the report. You can filter out events according to these categories:

- Severity (Error, Normal, Warning)
- Node (by node name)
- Resource (by resource name)

Selecting a check box causes that entry type to appear in the log. Deselecting a check box filters out the associated entry type.

In the **Event Log Range** section of the filter, you can choose to report all events in the log, or you can get a report of logged events that occurred during a specified time range. The time range filter displays events that occurred during the interval of time before a specified date and time and after a specified date and time. For example, you might want to see events that occurred for a cluster resource before the date/time that a failover event was triggered and after the date/time of its previously known good state. By default, the **After** field displays the date and time of the most recent logged entry, and the **Before** field displays the date and time of the oldest available logged entry. The search includes the specified before and after date/time values. If you specify date/time values outside the range of entries that are available, the filter applies the default range.

In the **Details** section of the filter, you can control how many event log messages are displayed per page. The default is 10 events per page. If the number of entries for your search exceeds one page, you can click the right-arrow button and left-arrow button to page through the report. Click the double-left-arrow button to jump to the first page of the report. Click the double-right-arrow button to jump to the last page of the report.

In the **Details** section of the filter, you can also choose whether the log message portion of a log entry is displayed in the log.

To view the logged cluster events:

1. In iManager, select **Clusters > My Clusters**, then select the cluster.
2. Select the **Cluster Event Log** to view the latest cluster events.

   Ten most recent events are displayed by default in reverse chronological order.
3. (Optional) Generate a custom event report.
   - **3a** At the top of the Cluster Event Log page, click **Filter** to open the Event Log Filter page.
   - **3b** In the **Security** section, deselect the check box next to message types that are not of interest.
     - **Error**
     - **Normal**
     - **Warning**

For example, deselect **Normal** to omit normal events. Only error and warning messages will appear in the report.

3c In the **Node** section, deselect the check box next to nodes that are not of interest. Only events for selected nodes will appear in the report.

3d In the **Resource** section, deselect the check box next to cluster resources that are not of interest. Only events for selected resources will appear in the report.

3e In the **Event Log Range** section, select one of the following:

- Entire log
- Time range (inclusive)
  - Before this date/time
  - After this date/time

3f In the **Details** section, specify how many messages to display per page and whether the log message portion of a log entry is displayed in the report.

3g Click **OK** to apply the filter.

3h View the custom report.

4 Click **Clear Filter** to clear the current filter.

5 (Optional) Click **Clear Log** to remove all entries from the log.

All prior log entries are deleted and are no longer available.

9.7 Viewing the Cluster Event Log at the Command Line

You can view cluster events logged in the `/admin/Novell/Cluster/EventLog.xml` file by copying the `/admin/Novell/Cluster/EventLog.xml` file to a working location, and using the `cat` command to display the file.

1 Log in to the node as the `root` user, then open a terminal console.

2 Enter

```
cp /admin/Novell/Cluster/EventLog.xml /tmp/
```

```
cat /tmp/EventLog.xml
```

This returns the log content in an XML format. For example:

```
<logEntry>
  <timeStamp> 5/12/2013 23:28:38</timeStamp>
  <node>CG-03</node>
  <clusterEvent>Joined at epoch 2</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:27:16</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
  <clusterEvent>Running on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:27:15</timeStamp>
  <node>CG-02</node>
  <clusterEvent>Joined at epoch 1</clusterEvent>
</logEntry>

<logEntry>
```

```
  <timeStamp> 5/12/2013 23:14:32</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
  <clusterEvent>Running on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:14:19</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
  <clusterEvent>Loading on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:14:13</timeStamp>
  <node>CG-01</node>
  <clusterEvent>Joined at epoch 0</clusterEvent>
</logEntry>
```

9.8 Viewing Summaries of Failed or Incomplete Events from the *.out Log Files

In OES 11 SP2 and later, you can use the DotOutParser utility (`/opt/novell/ncs/bin/dotoutparser.pl`) print summaries of failed or incomplete events that have been recorded in a specified `/var/opt/novell/log/ncs/<resource_name>.<script>.out` log file. It also prints output (if any) from any commands that failed or got stuck. It optionally includes line numbers, so you can easily refer the summary output to the source lines in the `*.out` file. The higher the level of verbosity you specify, the more information you get in the summary output. By default, it prints the entries with plus (+) variable assignments. See Section A.5, “DotOutParser Utility,” on page 385.

1. Log in to the node as the `root` user.
2. Copy the `/var/opt/novell/log/ncs/<resource_name>.<script>.out` log files of interest to a working location on the node.
3. At a terminal console, enter

   ```
   /opt/novell/ncs/bin/dotoutparser.pl [-v|-vv|-vvv|-vvvv] /tmp/
   <resource_name>.<script>.out
   ```

   For example:

   ```
   /opt/novell/ncs/bin/dotoutparser.pl -v /tmp/POOL_09_SERVER.load.out

   Fri May  3 11:51:59 2013 (#1):
   2: + exit_on_error nss /poolact=POOL_09
   9: + return 0
   10: + exit_on_error ncpcon mount VOL_092=237
   25: + return 0
   26: + exit_on_error ncpcon mount VOL_091=238
   41: + return 0
   42: + exit_on_error add_secondary_ipaddress 192.168.1.9
   94: + return 0
   95: + exit_on_error ncpcon bind --ncpservername=CLUSTER1_POOL_09_SERVER --
   ipaddress=192.168.1.9
   97: ++ ncpcon bind --ncpservername=CLUSTER1_POOL_09_SERVER --ipaddress=192.168.1.9
   98: ... Executing " bind"
   99:
   100: Cannot bind server, failure reason = -632
   101:
   102: ... FAILED completion [elapsed time = 20 Seconds 2 msecs 671 usecs]
   107: + exit 1
   ```

9.9 Generating a Cluster Configuration Report

A cluster configuration report can help you diagnose problems with the cluster nodes and resources. The report contains the following information for the selected cluster:

Field	Description
Cluster name	The dot delimited distinguished name of the selected cluster (*clustername.context*).
Date	The date and time the report was generated.
Cluster status	Current Nodes Quorum Triggers List of resources (Name, State, Location, Lives, Up Since)
Cluster options	Quorum Triggers (Number of nodes, Timeout) Protocol Setting (Heartbeat, Tolerance, Master Watchdog, Slave Watchdog, Maximum retransmits) List of nodes (Name, State, Number, IP Address, and Distinguished Name) Resource Mutual Exclusion Groups
Resource details	Information listed (as it applies) by resource for the master IP address resource and each of the cluster resources: Policies (Follows masters, Ignore quorum, Start mode, Failover mode, Failback mode, Assigned nodes) Load script and timeout Unload script and timeout Monitor script Protocols Virtual server name CIFS server name IP address Advertising protocols

To generate a cluster configuration report:

1. In iManager, select **Clusters > My Clusters**.
2. In your personalized My Clusters list, select the check box next to the cluster you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Click **Action > Run Report**.

   When you select **Run Report**, the generated report opens in a new window, instead of in the child frame. This allows you to print and save the report from any web browser.
4. Use the scroll bar to view the report in the browser.
5. (Optional) Use the browser options to print the report.

6 (Optional) Save the report to a text file.

Because some of the output is scripted, you must copy and paste the report content instead of using the browser's **File > Save as** option.

6a Click the cursor in the browser window where the report results are displayed, press Ctrl+A to select all, press Ctrl+C to copy the report to the clipboard.

6b In a text editor, open a new document, then press Ctrl+V to paste the report in the document.

6c Save the report document.

9.10 Repairing a Cluster

Repair a NetWare Cluster: To repair a NetWare cluster,

1 In iManager, select **Clusters** > **My Clusters**.

2 In your personalized My Clusters list, select the check box next to the cluster you want to repair.

3 Select **Action** > **Repair**.

This option is useful if you are replacing an existing NetWare server. Install the NetWare operating system and OES Cluster services on the new server, and name server with the same name as the old server. Choose the Repair option to automatically configure OES Cluster Services on the replacement server. You must manually copy the `ldncs.ncf` and `uldncs.ncf` files from the `sys:\system` directory to this node from another node in the cluster.

Repair an OES Linux Cluster: To repair an OES cluster,

1 In iManager, select **Clusters** > **My Clusters**.

2 In your personalized My Clusters list, select the check box next to the cluster you want to repair.

3 Select **Action** > **Repair**.

The Repair option is useful when you see a mismatch in the cluster resource list or the resource priority list. The cluster repair will make the resource list consistent with the resource objects.

For example, if the cluster status command output is blank or the deleted resources are still listed, the repair option will make the resource list match the resource objects. Resources will then be displayed properly in CLI and iManager.

NCS logs the corresponding repair option messages in the `/var/opt/novell/log/ncs/repair.log` file.

9.11 Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node

After a cluster resource is enabled, the load and unload scripts take care of starting and stopping the services or mounting and dismounting the volumes that are configured in the resource. You start services and mount devices by onlining the resource. You stop services and unmount devices by offlining the resource.

- Section 9.11.1, “Guidelines for Onlining and Offlining Resources,” on page 140
- Section 9.11.2, “Using iManager to Online and Offline Resources,” on page 140
- Section 9.11.3, “Using the Cluster Online Command,” on page 141
- Section 9.11.4, “Using the Cluster Offline Command,” on page 141
- Section 9.11.5, “Using the Cluster Migrate Command,” on page 141

9.11.1 Guidelines for Onlining and Offlining Resources

Onlining a resource runs the load script, which loads the resource on its primary preferred node, or on an alternate preferred node if possible, according to the order in its **Preferred Nodes** list. No action is taken if the preferred nodes are not active in the cluster, or if there are Resource Mutual Exclusion conflicts on its active preferred nodes.

Offlining a resource runs the unload script and unloads the resource from the server. The resource cannot be loaded on any other servers in the cluster and remains unloaded until you load it again.

If monitoring is enabled for a resource and you start, stop, or restart services outside the cluster framework while the cluster resource is online, a monitoring failure occurs, which triggers a failover of the cluster resource.

If you edit the scripts for a resource that is online, the changes you made do not take effect until the resource is taken offline and brought online again.

If a resource is in a comatose state, you must take the resource offline before it can be brought online again.

There are maintenance situations where you might need to control services by using the standard interfaces. In these rare cases, you must first offline the cluster resource, then use non-cluster methods to start and stop a service or mount and dismount a volume outside the cluster framework. When you are done, online the cluster resource to resume normal cluster handling.

9.11.2 Using iManager to Online and Offline Resources

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. On the Cluster Manager page, select the check box next to the resource you want to manage, then click **Online** or click **Offline**.

9.11.3 Using the Cluster Online Command

To bring a resource online on its most preferred node that is currently active:

1 As the `root` user, enter the following at the command prompt of a cluster node:

```
cluster online resource_name
```

To bring a resource online on a specific active node:

1 As the `root` user, enter the following at the command prompt of a cluster node:

```
cluster online resource_name node_name
```

9.11.4 Using the Cluster Offline Command

To take a resource offline:

1 As the `root` user, enter the following at the command prompt of a cluster node:

```
cluster offline resource_name
```

9.11.5 Using the Cluster Migrate Command

The `cluster migrate` command can be used to move a specified resource from the node where it is currently running to a another node. If a node name is not provided, the resource is brought online according to its **Preferred Nodes** list if it is possible. You can alternatively specify a destination node from its **Preferred Nodes** list that is running in the cluster. It migrates the resource to the specified node if it is possible. No action is taken if the specified node is not in the resource's preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.

1 As the `root` user, enter the following at the command prompt of a cluster node:

```
cluster migrate resource_name [node_name]
```

9.12 Cluster Migrating Resources to Different Nodes

You can migrate cluster resources to different servers in your cluster without waiting for a failure to occur. You might want to migrate resources to lessen the load on a specific server, to free up a server so it can be brought down for scheduled maintenance, or to increase the performance of the resource or application by putting it on a faster machine. The node you choose must be in the resource's preferred nodes list and available to run the resource.

Migrating resources lets you balance the load and evenly distribute applications among the servers in your cluster.

Using iManager

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 On the Cluster Manager page, select the check box next to the resource you want to cluster migrate, then click **Migrate**.

A page appears, displaying a list of possible servers that you can migrate this resource to.

4 Select a server from the list to migrate the resource to, then click **OK** to migrate the resource to the selected server.

Using Console Commands

1 As the `root` user, enter the following at the command prompt of a cluster node:

```
cluster migrate <resource_name> {<node_name> | -list | -most | -next}
```

Migrates the specified resource from the node where it is currently running to the node you specify in the command. The node you migrate the resource to must be running in the cluster and also be in the resource’s **Preferred Nodes** list.

Options

<node_name>

Migrates the resource to the specified node if possible. No action is taken if the specified node is not in the resource’s preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.

`-l`, `-list`

Shows the preferred nodes list for the specified resource.

`-m`, `-most`

Migrates the specified resource to the most preferred node currently in the cluster. If the resource is running on such a node, no action is taken.

`-n`, `-next`

Migrates the resource to the node in the preferred node list that is next in order to the node where the resource is currently running. If there is no next node or all such nodes are not in the cluster currently, it searches from the beginning of the resource’s preferred nodes list. No action is taken if a new destination could not be determined after the search is exhausted.

Examples

Move `res1` from the current node to `node2` in the resource’s preferred nodes list. The resource is migrated only if there are no Resource Mutual Exclusion conflicts on the target node.

```
cluster migrate res1 node2
```

List the preferred nodes for resource AUTO_POOL_14_SERVER.

```
cluster migrate AUTO_POOL_14_SERVER -l

Preferred nodes for resource ‘AUTO_POOL_14_SERVER'
  CG-06
  CG-05
  CG-04
  CG-03
  CG-02
  CG-01
  CG-08
  CG-07

Status for Resource: AUTO_POOL_14_SERVER
Running on CG-05  Lives: 65
Revision: 5
```

9.13 Removing (Leaving) a Node from the Cluster

When a node leaves a cluster, the node is no longer visible to the other nodes in the cluster. You might want a node to leave the cluster to perform maintenance on the server, or to upgrade the server during a rolling cluster upgrade.

- Section 9.13.1, “Removing One Node at a Time,” on page 143
- Section 9.13.2, “Removing All Nodes at the Same Time (Taking the Cluster Down),” on page 144

9.13.1 Removing One Node at a Time

The `cluster leave` command allows you to remove a node from a cluster so that the node is not visible to other servers in the cluster.

1 Log in as the `root` user on the node in the cluster that you want to remove, then enter the following at a command prompt:

```
cluster leave
```

After the node has successfully left the cluster, the following message is displayed:

```
No longer a member of cluster cluster_name
```

9.13.2 Removing All Nodes at the Same Time (Taking the Cluster Down)

The `cluster down` command removes all cluster nodes from the cluster. This command has the same effect as executing the `cluster leave` command on every server in the cluster. You are prompted to confirm the cluster down.

1. Log in as the `root` user on any node in the cluster, then enter the following at a command prompt:

   ```
   cluster down
   ```

2. When you are prompted to confirm the `cluster down` command, specify **Yes**, then press Enter.

9.14 Joining a Node to the Cluster (Rejoining a Cluster)

The `cluster join` command allows you to add a node back to the cluster so that the node is again visible to other servers in the cluster.

While the node is not active in the cluster, any changes in storage objects are not updated automatically for the node. For example, the following storage object changes are automatically updated only for nodes that are currently active in the cluster:

- Creating, deleting, or expanding a clustered NSS pool
- Creating, deleting, or expanding a clustered LVM volume group
- Adding or removing segments in an NSS software RAID (that is used by a clustered NSS pool)
- Creating, mirroring, or re-creating an SBD partition
- Creating or deleting shared partitions

If OES Cluster Services is configured to autostart when a server is restarted, a node that has left the cluster automatically rejoins the cluster on reboot. A server restart automatically rescans devices, rescans storage objects, and remaps multipath devices. Thus, if you reboot a node that has left the cluster, a node automatically recognizes any storage object changes made while the node was not active in the cluster.

To rejoin a cluster without rebooting the server:

1. Log in as the `root` user on the server that you want to join the cluster, then enter the following at a command prompt:

   ```
   cluster join
   ```

2. (Storage object changes) If you modified shared storage objects while the node was not active in the cluster, you must rescan the storage objects in order for the node to recognize these changes:

   2a Log in to the node as the `root` user and open a terminal console.

   2b Scan for storage object changes. At the command prompt, enter

   ```
   nlvm rescan
   ```

9.15 Shutting Down OES Cluster Services When Servicing Shared Storage

If you need to power down or recycle your shared storage system (such as when you apply firmware changes to a SAN), you should shut down OES Cluster Services prior to doing so. One way to do this is to use the `cluster down` command on any of the nodes in the cluster. You are prompted to confirm the cluster down.

9.16 Enabling or Disabling Cluster Maintenance Mode

Cluster maintenance mode lets you temporarily suspend the cluster heartbeat while hardware maintenance is being performed. This is useful if you want to reset or power down the LAN switch without bringing down cluster servers.

Enabling the cluster maintenance mode from one cluster node puts the entire cluster in maintenance mode.

1 Log in as the `root` user to a node in the cluster, then enter the following at a command prompt:

```
cluster maintenance on
```

If the master server in the cluster is up, disabling the cluster maintenance mode from one cluster node brings the entire cluster out of maintenance mode. If the master server in the cluster goes down while the cluster is in cluster maintenance mode, you must disable cluster maintenance mode on all remaining cluster nodes in order to bring the cluster out of maintenance mode.

1 Log in as the `root` user to a node in the cluster, then enter the following at a command prompt:

```
cluster maintenance off
```

An expiry time can be set to automatically turn off the cluster maintenance mode. The specified time can be in any format, which is supported by the Linux 'date' utility. The ‘expiry_time’ option is applicable only when all the cluster nodes are on OES 2023.

1 Log in as the `root` user to a node in the cluster, then enter the following at a command prompt:

```
cluster maintenance on '2021-04-04 12:34:51'
```

All the cluster nodes will be in the maintenance mode till April 4,2021; time - 12:34:51. The cluster will be out of maintenance mode automatically at the specified time.

9.17 Viewing the Cluster Node Properties

You can view the cluster node number and IP address of the selected node as well as the distinguished name of the Linux Server object.

1 In iManager, select **Clusters > My Clusters**.

2 Click the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Select **Cluster Options**.

4 Select the check box next to the cluster node whose properties you want to view, then click the **Details** link.

5 View the node information:

- **Distinguished Name:** The distinguished name of the Server object.
- **Node Number:** The cluster node number.
- **IP Address:** The IP address of the server.

6 Click **OK** to return to the Cluster Options page.

9.18 Creating or Deleting Cluster SBD Partitions

If a single node (or group of nodes) somehow becomes isolated from other nodes, a condition called *split brain* results. Each side believes the other has failed, and forms its own cluster view that excludes the nodes it cannot see. Neither side is aware of the existence of the other. If the split brain is allowed to persist, each cluster will fail over the resources of the other. Since both clusters retain access to shared disks, corruption will occur when both clusters mount the same volumes.

OES Cluster Services provides a split-brain detector (SBD) function to detect a split-brain condition and resolve it, thus preventing resources from being loaded concurrently on multiple nodes. The SBD partition contains information about the cluster, nodes, and resources that helps to resolve the split brain condition.

OES Cluster Services requires an SBD partition for a cluster if its nodes use physically shared storage. Typically, you create the SBD when you configure the cluster on the first node. You can alternatively configure an SBD for the cluster after you configure the first node, but before you configure OES Cluster Services on the second node of the cluster. You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails.

An SBD must exist and the cluster must be enabled for shared disk access before you attempt to create shared storage objects in a cluster, such as pools and volumes. OES Linux Volume Manager (NLVM) and other NSS management tools need the SBD to detect if a node is a member of the cluster and to get exclusive locks on physically shared storage.

This section describes how to use the OES Cluster Services SBD Utility (`sbdutil`) to create and delete SBD partitions.

- Section 9.18.1, “Requirements and Guidelines for Creating an SBD Partition,” on page 147
- Section 9.18.2, “Before You Create a Cluster SBD Partition,” on page 149
- Section 9.18.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 149
- Section 9.18.4, “Mirroring an Existing SBD Partition with NSSMU,” on page 152
- Section 9.18.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 154
- Section 9.18.6, “Removing a Segment from a Mirrored Cluster SBD Partition,” on page 157
- Section 9.18.7, “Deleting a Non-Mirrored Cluster SBD Partition,” on page 158
- Section 9.18.8, “Deleting a Mirrored Cluster SBD,” on page 159
- Section 9.18.9, “Additional Information about SBD Partitions,” on page 160

9.18.1 Requirements and Guidelines for Creating an SBD Partition

Consider the requirements and guidelines in this section when you create a OES Cluster Services SBD (split-brain detector) partition for an existing cluster.

IMPORTANT: Check to see if a cluster SBD partition already exists before you create a new one. See Section 9.18.2, “Before You Create a Cluster SBD Partition,” on page 149.

- “Preparing OES Cluster Services” on page 147
- “Using a Shared Disk System” on page 147
- “Preparing a SAN Device for the SBD” on page 147
- “Initializing and Sharing a Device for the SBD” on page 148
- “Working with NLVM Commands in a Cluster” on page 148
- “Determining the SBD Partition Size” on page 148
- “Replacing an Existing SBD Partition” on page 148

Preparing OES Cluster Services

Before you create an SBD partition for an existing cluster, you must take the cluster down and stop OES Cluster Services software on all nodes. Do not restart OES Cluster Services and rejoin nodes to the cluster until after you create the new SBD and configure the Shared Disks flag attribute for the Cluster object.

You can mirror an existing SBD while the cluster is up and running.

Using a Shared Disk System

You must have a shared disk system (such as a Fibre Channel SAN or an iSCSI SAN) connected to your cluster nodes before you create a split-brain-detector (SBD) partition. See Section 4.8, “Shared Disk Configuration Requirements,” on page 50 for more information.

Preparing a SAN Device for the SBD

Use the SAN storage array software to carve a LUN to use exclusively for the SBD partition. The 512 byte device should have at least 20 MB of free available space and the 4096 (4Kn) byte device should have at least 80 MB of free available space. Connect the LUN device to all nodes in the cluster.

For device fault tolerance, you can mirror the SBD partition. Use the SAN storage array software to carve a second LUN of the same size to use as the mirror. Connect the LUN device to all nodes in the cluster.

The device you use to create the SBD partition must not be a software RAID device. A hardware RAID configured in a SAN array is seen as a regular device by the server.

If you attach new devices to the server while it is running, you should scan for new devices on each cluster node to ensure that the devices are recognized by all nodes. Log in as the `root` user, launch a terminal console, then enter

```
nlvm -s rescan
```

Initializing and Sharing a Device for the SBD

Before you use `sbdutil` to create an SBD, you must initialize each of the SAN devices that you created for the SBD and mark each device as Shareable for Clustering. When you mark the device as Shareable for Clustering, share information is added to the disk in a free-space partition that is about 4 MB in size. This space becomes part of the SBD partition.

When the cluster is down and OES Cluster Services is stopped, you can use NSSMU, the Storage plug-in for iManager, the `nlvm init` command, or an NSS utility called `ncsinit` to initialize a device and set it to a shared state. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

Working with NLVM Commands in a Cluster

If you have physically shared storage and the SBD does not exist, NSS management tools cannot detect if the node is a member of the cluster and cannot get exclusive locks to the physically shared storage. In this state, you can use the `-s` NLVM option with NLVM commands to override the shared locking requirement and force NLVM to execute the commands. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

Determining the SBD Partition Size

You use the `size` option with the SBD Utility to specify the desired size of the SBD partition. You can specify how much free space to use for the SBD, or you can specify the `-1` option to use the entire device (the maximum size). If you specify a device to use as a mirror, the same amount of space is used. If you specify to use the maximum size and the mirror device is bigger than the SBD device, you will not be able to use the excess free space on the mirror for other purposes.

Because an SBD partition ends on a cylinder boundary, the partition size might be slightly smaller than the size you specify. When you use an entire device for the SBD partition, you can use the `-1` option as the size, and let the software determine the size of the partition.

Replacing an Existing SBD Partition

To replace an existing SBD partition, you must first delete the old SBD partition, and then create the new one. To reuse the SBD partition's device, you must remove the SBD partition, then re-initialize and share the device.

You must take the cluster down and stop OES Cluster Services on all nodes before you delete the existing SBD partition. Do not restart OES Cluster Services and rejoin nodes to the cluster until after you create the new SBD.

9.18.2 Before You Create a Cluster SBD Partition

Before you create a OES Cluster Services SBD partition, you should ensure that an SBD does not already exist on your cluster. The new SBD partition is not recognized until the old partition is deleted.

1. Log in as the `root` user to any node in the cluster, and launch a terminal console.
2. At the command prompt, enter

   ```
   sbdutil -f
   ```

   The `-f` option tells you whether an SBD partition exists for the cluster, and identifies the SAN device where the SBD partition is located. It returns `NotFound` if the SBD partition does not exist.
3. If an SBD partition already exists, use one of the following methods to delete the existing SBD partition before attempting to create another one:
   - Section 9.18.7, “Deleting a Non-Mirrored Cluster SBD Partition,” on page 158
   - Section 9.18.8, “Deleting a Mirrored Cluster SBD,” on page 159

9.18.3 Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL

If you did not create a cluster SBD partition during the OES Cluster Services installation on the first node of the cluster, you can create it later by using the SBDUTIL utility (`/opt/novell/ncs/bin/sbdutil`). You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails. See the man page for `sbdutil` for more information on how to use it.

If a cluster partition does not exist, create one by using the SBDUTIL:

1. Ensure that nobody else is changing any storage on any nodes at this time.

   Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.
2. Take the cluster down:
   - **2a** Log in to any node in the cluster as the `root` user, then open a terminal console.
   - **2b** At the command prompt, enter

     ```
     cluster down
     ```
3. On each cluster node, stop OES Cluster Services:
   - **3a** Log in to the cluster node as the `root` user, then open a terminal console.
   - **3b** At the command prompt, enter

     ```
     rcnovell-ncs stop
     ```

     or

     ```
     systemctl stop ncs.service
     ```
   - **3c** After you have stopped OES Cluster Services on all nodes, continue with the next step.

4 Prepare a SAN device to use for the SBD partition:

4a Use the SAN storage array software to carve a device to use exclusively for the SBD partition.

4b Attach the device to all nodes in the cluster.

4c On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

5 Log in to any node in the cluster as the `root` user, then open a terminal console.

6 View a list of the devices and identify the leaf node name (such as `sdc`) of the SAN device that you want to use for the SBD partition. At the command prompt, enter

```
nlvm -s list devices
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.

For example, an uninitialized device reports a format of `None` and a shared state of `No`:

```
sdc size=102.00MB free=0KB format=None shared=No RAID=No
```

7 Initialize and share the device. At the command prompt, enter

```
nlvm -s init <device_name> format=msdos shared
```

WARNING: Initializing a device destroys all data on the device.

Replace *device_name* with the leaf node name (such as `sdc`) of the SAN device you want to use as the SBD partition.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the device as Shareable for Clustering.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

For example, the formatted device `sdc` reports a format of `MSDOS` and a shared state of `Yes`:

```
sdc size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
```

8 As the `root` user, enter the following at the command prompt:

```
sbdutil -c -n <cluster_name> -d <device_name> -s <size>
```

For the -n option, replace **cluster_name** with the name of the cluster, such as `cluster1`. This name must match the name of an existing cluster that has a Cluster object in eDirectory. The name is case sensitive.

For the -d option, replace *device_name* with the leaf node name of the device where you want to create the cluster partition, such as `sdc`.

For the -s option, use one of the following methods for specifying the size of the SBD partition. For information about size requirements, see "Preparing a SAN Device for the SBD" on page 147.

- **Specify a value:** Replace *size* with the size (in MB) to use for the SBD partition.

  For example, the following command creates the `/dev/nss/mycluster1.sbd` partition with a size of 200 MB:

  ```
  sbdutil -c -n mycluster1 -d sdc -s 200
  ```

- **Specify -1:** You can specify the size as -1 to use all free space on the device. This option allows OES Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

  For example, the following command creates the `/dev/nss/cl1.sbd` partition on the CX4-LUN000 device, and uses the entire device:

  ```
  sbdutil -c -n cl1 -d CX4-LUN000 -s -1
  ```

- **Use default size:** If the -s option is not used, the default size is 8 MB.

  For example, the following command creates the `/dev/nss/cl2.sbd` partition on the sdd device with the default size of 8 MB:

  ```
  sbdutil -c -n cl2 -d sdd
  ```

9 On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

10 (Optional) You can use the `sbdutil -f -s` command to view the path of the SBD partition.

11 Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the **NCS: Shared Disk Flag** attribute is enabled.

11a In a web browser, open iManager, then log in to the eDirectory tree that contains the cluster you want to manage.

IMPORTANT: Log in as an administrator user who has sufficient rights in eDirectory to delete and modify eDirectory objects.

11b Select **Directory Administration**, then select **Modify Object**.

11c Browse to locate and select the Cluster object of the cluster you want to manage, then click **OK**.

11d Under **Valued Attributes**, select the NCS: Shared Disk Flag, then click **Edit**.

11e Select (enable) the **NCS: Shared Disk Flag** check box, then click **OK**.

11f Click **Apply** to save changes.

12 On each cluster node, start OES Cluster Services:

12a Log in to the cluster node as the `root` user, then open a terminal console.

12b At the command prompt, enter

```
rcnovell-ncs start
```

or

```
systemctl start ncs.service
```

12c After you have restarted OES Cluster Services on all nodes, continue with the next step.

13 On each cluster node, join the cluster. At the command prompt, enter

```
cluster join
```

14 (Optional) Continue with Section 9.18.4, “Mirroring an Existing SBD Partition with NSSMU,” on page 152.

9.18.4 Mirroring an Existing SBD Partition with NSSMU

You can mirror an existing OES Cluster Services SBD partition to provide device fault tolerance. It is not necessary to take the cluster down or stop the cluster software.

This section describes how to use NSSMU to mirror the SBD partition. For information about using OES Linux Volume Manager (NLVM) commands to mirror an SBD, see “Mirroring an Existing SBD Partition with NLVM” in the *OES 2023: NLVM Reference*.

1 Prepare a SAN device to use as the mirror segment for the SBD partition:

1a Use the SAN storage array software to carve a device that is at least the size of the existing SBD partition’s device.

1b Attach the device to all nodes in the cluster.

1c On each node, log in as the `root` user and rescan for devices:

```
nlvm rescan
```

2 Initialize and mark the device as **Shareable for Clustering**:

2a Log in to any node of the cluster as the `root` user, then open a terminal console.

2b Launch NSSMU:

```
nssmu
```

2c In the NSSMU Main Menu, select **Devices** and press Enter.

2d In the **Devices** list, select the device that you want to use for the SBD mirror, press F3 to initialize the device, then press y (Yes) to confirm and continue.

WARNING: Initializing a device destroys all of the data on it.

2e Select the DOS partitioning scheme for the device, then press Enter.

DOS supports devices up to 2 TB in size. GPT supports devices of any size.

Wait for the page to refresh before continuing.

2f Press F6 to mark the device as shareable for clustering.

The **Shareable for Clustering** value changes from **No** to **Yes**.

2g Press Esc to return to the NSSMU Main Menu.

3 In NSSMU, mirror the SBD partition:

3a Select **Partitions** from the NSSMU Main Menu.

3b From the list of partitions, select the SBD partition that you want to mirror. The partition is named the same as the cluster with an `.sbd` extension, such as `clus1.sbd`.

3c Press F3 to open the dialog to create the RAID1 mirror.

3d Type the cluster name (such as `clus1`) for the RAID, then press Enter.

This is the same name as the SBD partition but without the `.sbd` extension.

3e From the list of available devices, select a device to use as the second segment of the mirror, then press the space bar to choose it.

When the device is selected, the asterisk next to the device stays there even if you move the cursor up and down in the list.

3f Press F3 again to accept your selection and create the mirror.

3g In the confirmation message, press y (Yes) to approve the RAID1 creation.

The SBD partitions are renamed with the partition name, and now have the extension of `.msbd0` (mirrored SBD partition 0) and `.msbd2` (mirrored SBD partition 1).

For example, for a cluster named `clus1`, the mirrored SBD partitions are `clus1.msbd0` and `clus1.msbd1`.

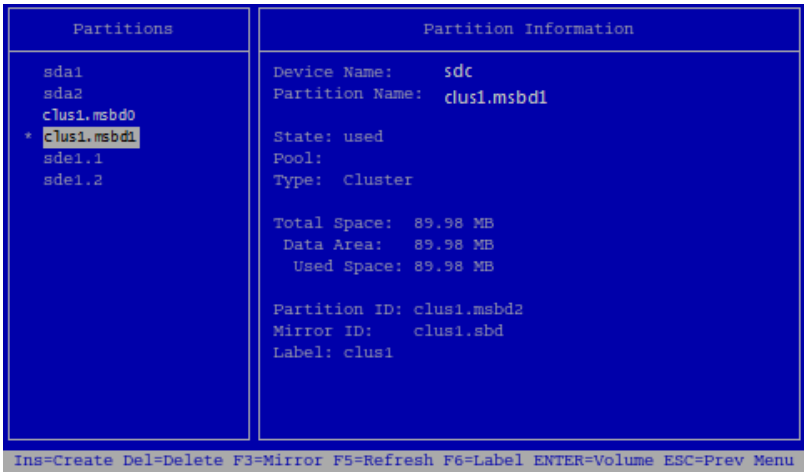

3h Press Esc to return to the NSSMU Main Menu.

4 In NSSMU, verify that the RAID1 device was created for the SBD:

4a Select **RAID Devices** from the NSSMU Main Menu.

4b Select the RAID1 device that you created to view details about the mirrored SBD device.

4c View the RAID Status to ensure that synchronization has begun.

Synchronization is complete when the status is **Synchronized**.

5 Press Esc twice to exit NSSMU.

9.18.5 Creating a Mirrored Cluster SBD Partition with SBDUTIL

You can create a mirrored OES Cluster Services SBD partition to provide device fault tolerance for the SBD. You must take the cluster down and stop the cluster software.

Use the procedure in this section to create a new mirrored SBD partition by using the `sbdutil` utility. For information about using OES Linux Volume Manager (NLVM) commands to create a new mirrored SBD partition, see "Creating a Mirrored SBD Partition with NLVM" in the *OES 2023: NLVM Reference*.

1. Ensure that nobody else is changing any storage on any nodes at this time.

   Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

2. Take the cluster down:

   2a Log in to any node in the cluster as the `root` user, then open a terminal console.

   2b At the command prompt, enter

   ```
   cluster down
   ```

3. On each cluster node, stop OES Cluster Services:

   3a Log in to the cluster node as the `root` user, then open a terminal console.

   3b At the command prompt, enter

   ```
   rcnovell-ncs stop
   ```

   or

   ```
   systemctl stop novell-ncs.service
   ```

   3c After you have stopped OES Cluster Services on all nodes, continue with the next step.

4. Prepare two SAN devices to use for the SBD partition:

   4a Use the SAN storage array software to carve two devices of equal size to use exclusively for the mirrored SBD partition.

   4b Attach the devices to all nodes in the cluster.

   4c On each node, log in as the `root` user and rescan for devices:

   ```
   nlvm -s rescan
   ```

   Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

5. Log in to any node in the cluster as the `root` user, then open a terminal console.

6. View a list of the devices and identify the leaf node name (such as `sdc`) of the two SAN devices that you want to use for the mirrored SBD partition. At the command prompt, enter

   ```
   nlvm -s list devices
   ```

   Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.

For example, uninitialized devices report a format of None and a shared state of No:

```
sdc size=102.00MB free=0KB format=None shared=No RAID=No
sdd size=102.00MB free=0KB format=None shared=No RAID=No
```

7 Initialize and share the two devices. At the command prompt, enter

```
nlvm -s init <device_name1>,<device_name2> format=msdos shared
```

WARNING: Initializing a device destroys all data on the device.

Replace *device_name1* and *device_name2* with the leaf node names (such as sdc and sdd) of the two SAN devices you want to use for the mirrored SBD partition.

Specify a partitioning format of msdos.

Specify the shared option to mark the devices as Shareable for Clustering.

Use the -s NLVM option to override the shared locking requirement and force the command to execute.

For example, to initialize devices sdc and sdd, enter

```
nlvm -s init sdc,sdd format=msdos shared
```

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

For example, the formatted devices sdc and sdd report a format of MSDOS and a shared state of Yes:

```
sdc size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
sdd size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
```

8 Create a mirrored SBD partition. At the command prompt, enter

```
sbdutil -c -n <cluster_name> -d <device_name1> -d <device_name2> -s
<size>
```

Replace **cluster_name** with the name of the cluster, such as cluster1. This name must match the name of an existing cluster that has a Cluster object in eDirectory. The name is case sensitive.

Replace *device_name1* and *device_name2* with the leaf node names (such as sdc and sdd) of the two SAN devices you want to use for the mirrored SBD partition. The cluster1.msbd0 mirrored SBD partition is created on the first device option instance in the command. The cluster1.msbd1 mirrored SBD partition is created on the second device option instance in the command.

For the -s option, replace *size* with the size (in MB) to use for each of the SBD RAID1 segments. Specify the size only once. Both devices should be the same size, but if they are not, the size of the RAID segments is determined by the size of the smaller device. Use one of the following methods for specifying the size of the SBD partition. For information about size requirements, see "Preparing a SAN Device for the SBD" on page 147.

- **Specify a value:** Replace *size* with the size (in MB) to use for the SBD partition.

  For example, the following command creates the `cl1.sbd` mirrored RAID device and the `cl1.msbd0` and `cl1.msbd1` partitions. Each segments is up to 1020 MB in size, depending on where the last cylinder boundary lies.

  ```
  sbdutil -c -n cl1 -d CX4-LUN000 -d CX4-LUN001 -s 1020
  ```

- **Specify -1:** You can specify the size as -1 to use all free space on the device. This option allows OES Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

  For example, the following command uses devices `sdc` and `sdd` to create the `cluster1.sbd` mirrored RAID device and the SBD mirrored partitions `cluster1.msbd0` and `cluster1.msbd1`. It uses all available free space on the smaller device as the segment size.

  ```
  sbdutil -c -n cluster1 -d sdc -d sdd -s -1
  ```

- **Use default size:** If the -s option is not used, the default size is 8 MB.

9 On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

10 (Optional) You can use the `sbdutil -f -s` command to verify the path and name of the SBD RAID device.

You can also use NSSMU, the Storage plug-in to iManager, or the `nlvm -s list partitions` command to view the partitions used by the RAID device.

11 Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the **NCS: Shared Disk Flag** attribute is enabled.

11a In a web browser, open iManager, then log in to the eDirectory tree that contains the cluster you want to manage.

IMPORTANT: Log in as an administrator user who has sufficient rights in eDirectory to delete and modify eDirectory objects.

11b Select **Directory Administration**, then select **Modify Object**.

11c Browse to locate and select the Cluster object of the cluster you want to manage, then click **OK**.

11d Under **Valued Attributes**, select the NCS: Shared Disk Flag, then click **Edit**.

11e Select (enable) the **NCS: Shared Disk Flag** check box, then click **OK**.

11f Click **Apply** to save changes.

12 On each cluster node, start OES Cluster Services:

12a Log in to the cluster node as the `root` user, then open a terminal console.

12b At the command prompt, enter

```
rcnovell-ncs start
```

or

```
systemctl start novell-ncs.service
```

12c After you have restarted OES Cluster Services on all nodes, continue with the next step.

13 On each cluster node, join the cluster. At the command prompt, enter

```
cluster join
```

9.18.6 Removing a Segment from a Mirrored Cluster SBD Partition

You can remove a segment from a mirrored cluster SBD partition and keep the remaining SBD partition. The software RAID definition remains, so if you delete the remaining partition later, you must delete the software RAID instead of simply deleting the partition as with a stand-alone SBD partition.

IMPORTANT: You can use NLVM commands to completely unmirror the SBD RAID1 device, and keep the SBD partition. See "Unmirroring a Mirrored SBD Partition with NLVM" in the *OES 2023: NLVM Reference*.

You can specify which segment to keep when you use NSSMU to remove a segment from a software RAID 1 (mirror) device.

1 Log in to any node as the `root` user, then launch a terminal console.

2 At the command prompt, enter

```
cluster maintenance on
```

This causes all cluster servers to enter maintenance mode.

3 Remove a segment from the SBD RAID1 device.

3a Launch NSSMU:

```
nssmu
```

3b In the NSSMU Main Menu, select **RAID Devices**.

3c Select the software RAID1 device for the cluster SBD that you want to manage.

3d Press Enter to show its member segments.

3e Select the member segment you want to delete, press Delete to remove the RAID segment, then press y (Yes) to confirm the removal.

3f Press Esc to return to the Software RAIDs page.

The RAID definition remains for the remaining segment of the mirrored SBD partition. The SBD RAID1 reports that it has one segment.

3g Press Esc twice to exit NSSMU.

4 At the terminal console of one cluster server, enter

```
cluster maintenance off
```

This causes all cluster servers to return to normal mode.

9.18.7 Deleting a Non-Mirrored Cluster SBD Partition

You might need to delete and re-create a OES Cluster Services SBD partition if the SBD becomes corrupted or its device fails. Use the procedure in this section to delete the SBD partition, then create a new SBD partition as described in Section 9.18.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 149 or Section 9.18.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 154.

IMPORTANT: You must take the cluster down and stop OES Cluster Services on all nodes before you delete the existing SBD partition. Do not restart OES Cluster Services and rejoin nodes to the cluster until after you create a new SBD.

---

1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

2 Take the cluster down:

2a Log in to any node in the cluster as the `root` user, then open a terminal console.

2b At the command prompt, enter

```
cluster down
```

3 On each cluster node, stop OES Cluster Services:

3a Log in to the cluster node as the `root` user, then open a terminal console.

3b At the command prompt, enter

```
rcnovell-ncs stop
```

or

```
systemctl stop novell-ncs.service
```

3c After you have stopped OES Cluster Services on all nodes, continue with the next step.

4 Delete the SBD partition.

4a Log in to any node in the cluster as the `root` user, then launch a terminal console.

4b Launch NSSMU:

```
nssmu
```

4c In the NSSMU Main Menu, select **Partitions**.

4d Select the SBD partition you want to delete, such as `clus1.sbd`.

4e Press Delete to delete the partition, then press `y (Yes)` to confirm the deletion.

4f Press Esc twice to exit NSSMU.

5 Use one of the following methods to create a new SBD partition:

- Section 9.18.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 149
- Section 9.18.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 154

Do not restart OES Cluster services and rejoin nodes to the cluster until after you create the new SBD.

9.18.8 Deleting a Mirrored Cluster SBD

You might need to delete and re-create a OES Cluster Services SBD partition if the SBD becomes corrupted or its device fails. Use the procedure in this section to delete a mirrored SBD partition, then create a new SBD partition as described in Section 9.18.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 149 or Section 9.18.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 154.

IMPORTANT: You must take the cluster down and stop OES Cluster Services on all nodes before you delete the existing mirrored SBD RAID1 device and its mirrored SBD partitions. Do not restart OES Cluster Services and rejoin nodes to the cluster until after you create a new SBD.

---

1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

2 Take the cluster down:

2a Log in to any node in the cluster as the `root` user, then open a terminal console.

2b At the command prompt, enter

```
cluster down
```

3 On each cluster node, stop OES Cluster Services:

3a Log in to the cluster node as the `root` user, then open a terminal console.

3b At the command prompt, enter

```
rcnovell-ncs stop
```

or

```
systemctl stop novell-ncs.service
```

3c After you have stopped OES Cluster Services on all nodes, continue with the next step.

4 Delete the mirrored software RAID that you used for the SBD partition:

4a Log in to any node in the cluster as the `root` user, then launch a terminal console.

4b Launch NSSMU:

```
nssmu
```

4c In the NSSMU Main Menu, select **RAID Devices**.

4d Select the software RAID1 for the mirrored SBD you want to delete.

4e Press Delete to delete the software RAID1 and its two mirrored SBD partitions, then press y (`Yes`) to confirm the deletion.

4f Press Esc twice to exit NSSMU.

5 Use one of the following methods to create a new SBD partition:

- Section 9.18.3, "Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL," on page 149
- Section 9.18.5, "Creating a Mirrored Cluster SBD Partition with SBDUTIL," on page 154

Do not restart OES Cluster services and rejoin nodes to the cluster until after you create the new SBD.

9.18.9 Additional Information about SBD Partitions

See the following resources for additional information about SBD partitions:

- To understand general requirements for SBD partitions, see Section 4.8.2, "SBD Partitions," on page 51.
- To configure the SBD when you configure the cluster on the first node, see Section 5.5.5, "Configuring a New Cluster," on page 67.
- To use OES Linux Volume Manager (NLVM) commands to create an SBD partition in a cluster, see "Creating or Mirroring an SBD Partition" in the *OES 2023: NLVM Reference*.
- To use NLVM commands to delete an SBD partition in a cluster, see "Deleting an SBD Partition with NLVM" in the *OES 2023: NLVM Reference*.
- For information about how the split brain detector works, see *NetWare Cluster Services: The Gory Details of Heartbeats, Split Brains, and Poison Pills* (TID 10053882) (http://support.novell.com/docs/Tids/Solutions/10053882.html).

9.19 Customizing Cluster Services Management

Some portions of OES Cluster Services management can be performed and custom by using virtual XML files that exist in the `/admin/Novell/Cluster` directory on Linux.

The cluster-related virtual XML files (management access points) are created on each server's `/admin/Novell/Cluster` directory. These files let you manage the cluster from any node in the cluster. This means that as long as the cluster is running, you can always access the cluster-related XML virtual files in the `/admin/Novell/Cluster` directory.

There are two types of virtual files in the `/admin/Novell/Cluster` directory: XML files and CMD files. The XML files are read-only and contain cluster configuration or cluster state information. The CMD files are write-then-read command files that are used to issue commands to the cluster and retrieve resulting status.

Table 9-3 lists the cluster-related virtual XML files and gives a brief description of each.

Table 9-3 *Cluster-Related Virtual XML Files*

Virtual XML File Name	Description
`Config.xml`	Provides the combined information from `ClusterConfig.xml`, `NodeConfig.xml`, `ResourceConfig.xml`, and `PoolConfig.xml`.
`ClusterConfig.xml`	Provides cluster configuration information.
`NodeConfig.xml`	Provides node configuration information for all nodes in the cluster that were active at the time the cluster was brought up.
`NodeState.xml`	Provides current information on the state of each node in the cluster (cluster membership).
`PoolConfig.xml`	Provides cluster-enabled pool and volume configuration information for each pool and volume.
`PoolState.xml`	Provides current information on the state of each cluster-enabled pool in the cluster.
`ResourceConfig.xml`	Provides resource configuration information for each resource in the cluster.
`ResourceState.xml`	Provides current information on the state of each resource in the cluster.
`State.xml`	Provides the combined information from `NodeState.xml`, `ResourceState.xml`, and `PoolState.xml`.

Table 9-4 lists the cluster-related CMD files and gives a brief description of each.

Table 9-4 *Cluster-Related CMD Files*

CMD File Name	Description
`Node.cmd`	Write-then-read command file used in conjunction with a Perl script to issue node-specific commands to the cluster and retrieve resulting node status and configuration information.
`Cluster.cmd`	Write-then-read command file used in conjunction with a Perl script to issue cluster-specific commands to the cluster and retrieve resulting cluster status and configuration information.
`Resource.cmd`	Write-then-read command file used in conjunction with a Perl script to issue resource-specific commands to the cluster and retrieve resulting resource status and configuration information.

10 Configuring and Managing Cluster Resources

After you create and configure a OES Cluster Services cluster, you are ready to create and configure cluster resources for the cluster. This section provides general instructions for creating cluster resources and configuring their behavior in the cluster.

For information about viewing and managing resource status on the cluster, see Chapter 9, “Managing Clusters,” on page 127.

- Section 10.1, “Planning for Cluster Resources,” on page 163
- Section 10.2, “Setting Up a Personalized List of Resources to Manage,” on page 165
- Section 10.3, “Using Cluster Resource Templates,” on page 169
- Section 10.4, “Creating Cluster Resources,” on page 173
- Section 10.5, “Configuring a Load Script for a Cluster Resource,” on page 174
- Section 10.6, “Configuring an Unload Script for a Cluster Resource,” on page 175
- Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176
- Section 10.8, “Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration,” on page 181
- Section 10.9, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 182
- Section 10.10, “Configuring Preferred Nodes and Node Failover Order for a Resource,” on page 184
- Section 10.11, “Configuring Resource Priorities for Load Order,” on page 185
- Section 10.12, “Configuring Resource Mutual Exclusion Groups,” on page 185
- Section 10.13, “Controlling Resource Monitoring,” on page 189
- Section 10.14, “Changing the IP Address of a Cluster Resource,” on page 189
- Section 10.15, “Renaming a Cluster Resource,” on page 191
- Section 10.16, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 192
- Section 10.17, “Additional Information for Creating Cluster Resources,” on page 195

10.1 Planning for Cluster Resources

Consider the guidelines in this section when planning for your cluster resources.

- Section 10.1.1, “Naming Conventions for Cluster Resources,” on page 164
- Section 10.1.2, “Using Parameter Values with Spaces in a Cluster Script,” on page 164
- Section 10.1.3, “Using Double Quotation Marks in a Cluster Script,” on page 164

- Section 10.1.4, “Script Length Limits,” on page 165
- Section 10.1.5, “Planning Cluster Maintenance,” on page 165
- Section 10.1.6, “Number of Resources,” on page 165
- Section 10.1.7, “Linux POSIX File System Types,” on page 165

10.1.1 Naming Conventions for Cluster Resources

Cluster resource names can be up to 63 characters. OES Cluster Services supports only alphanumeric characters and the underscore character in cluster resource names:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789_
```

Special characters, such as the following, are not supported in cluster resource names:

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

Because the NSS pool name and Linux volume names are automatically used in the cluster resource name, do not use special characters `!@#$%&( )` in names of NSS pools or Linux POSIX volumes that you plan to cluster enable.

10.1.2 Using Parameter Values with Spaces in a Cluster Script

In cluster scripts, if a parameter value contains spaces, you should enclose the value with both double-quotes (") and single-quotes ('). For example, `"'quoted text'"`. The preferred solution is to use both sets of quotation marks with the single-quotes on the inside. For OES 11 and later, the other order also works for spaces, but it reads a variable as text, not as a variable.

For example, either of the following techniques works for values with spaces:

```
exit_on_error echo "'Errors will be reported here.'"
```

```
exit_on_error echo '"Errors will be reported here."'
```

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

```
exit_on_error echo "\"Errors will be reported here.\""
```

Another alternative is to escape each of the spaces in the value. For example:

```
exit_on_error echo Errors\ will\ be\ reported\ here.
```

10.1.3 Using Double Quotation Marks in a Cluster Script

In cluster scripts, if a command requires double-quotation marks, you should enclose the quoted part of the command with both double-quotes (") and single-quotes ('), such as `"'quoted text'"`. The key is to use both sets of quotation marks with the single-quotes on the inside. This works for spaces and allows you to pass variables to the command. For example:

```
exit_on_error echo "'Errors will be reported here.'"
```

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

```
exit_on_error echo "\"Errors will be reported here.\""
```

10.1.4 Script Length Limits

Up to OES 2018 SP2, each cluster load, unload, or monitor script can be up to 3200 bytes in length. This limit includes commands, comments, and spaces. For example, non-special ASCII characters (including the space character) are 3.2 kilobytes per character, so you could have up to 3200 of these characters in a script. From OES 2018 SP3 onwards, a maximum of 1MB for the scripts is supported if the cluster is homogeneous with all the nodes being on OES 2018 SP3 and above. If it is a mixed node cluster with at least one node on a version lower than OES 2018 SP3, the script length is restricted to 3200 bytes/3.2KB.

IMPORTANT: Creating a script that exceeds the maximum script length can prevent a resource from loading. If your script commands and comments require more memory than 1 MB, you can spawn external scripts from a script.

10.1.5 Planning Cluster Maintenance

When performing cluster maintenance tasks, two cluster best practices should be observed:

- Perform maintenance tasks during non-peak hours so that users are minimally affected.
- Before performing maintenance on a node, cluster migrate its cluster resources to another node if you want the related users to be undisturbed.

10.1.6 Number of Resources

OES Cluster Services supports up to 254 resources in a cluster, regardless of the size of the cluster.

10.1.7 Linux POSIX File System Types

Ext3 is the default file system type used in the Generic File System template scripts. The Btrfs, Ext2, Ext3, ReiserFS, and XFS file systems have been tested and are fully supported.

10.2 Setting Up a Personalized List of Resources to Manage

The My Resources page in the Clusters plug-in for iManager allows each cluster administrator to set up a personalized list of cluster resources to manage. This allows an administrator to view at a glance the status of the specified resources. The list persists between the administrator's logins to the same iManager server.

Your personalized list of cluster resources and display preferences are saved on the iManager server in the following location:

`/var/opt/novell/iManager/nps/WEB-INF/config/NDS<`*`tree_name`*`>-<`*`user_name_and_context_without_dots`*`>/ncs.xml`

For example, for tree AVALONTREE and user `admin.novell`, the file path is:

```
/var/opt/novell/iManager/nps/WEB-INF/config/NDSAVALONTREE-adminnovell/
ncs.xml
```

The following sections describe how manage the list of resources and to personalize the display.

- Section 10.2.1, “Adding Resources to a My Resources List,” on page 166
- Section 10.2.2, “Viewing Information about Resources in a My Resources List,” on page 166
- Section 10.2.3, “Personalizing the My Resources Display,” on page 167
- Section 10.2.4, “Removing Resources from a My Resources List,” on page 168
- Section 10.2.5, “Updating the My Resources List after Renaming or Deleting a Resource,” on page 168

10.2.1 Adding Resources to a My Resources List

1. Log in to iManager as a cluster administrator.
2. In **Roles and Tasks**, select **Clusters > My Resources**.

   The list is initially empty.
3. Click **Add** to open the eDirectory browser pop-up window.
4. Browse the tree where you are currently logged in to locate and select one or more Cluster Resource objects, then click **OK**.

   Newly selected resources are added to your personalized list.
5. (Optional) Personalize the display as described in Section 10.2.3, “Personalizing the My Resources Display,” on page 167.

10.2.2 Viewing Information about Resources in a My Resources List

1. Log in to iManager as a cluster administrator.
2. In **Roles and Tasks**, select **Clusters > My Resources**.

   The My Resources page displays your personalized list of cluster resources in the tree.
3. View the Status icon next to a Cluster Resource object to understand the current state of the resource.

   If scrolling is needed to view the complete list, you can use the **Refresh** option to set a longer refresh rate for the page in order to allow enough time to view the status for all items.

State	Icon	Description
Normal		The resource is online.
Offline		The resource is offline. Administrator intervention is needed.

State	Icon	Description
Critical		The resource is comatose, or it cannot be found (renamed or deleted). Administrator intervention is needed. If you rename or delete a resource, its status changes to "Resource not found" when you next access or refresh the My Resources page. For a deleted resource, remove the resource from the list. For a renamed resource, remove the resource from the list, then add the resource using its new name.
Warning		An alert condition has occurred, and the resource needs administrator attention.
Unknown		The state of the cluster resource cannot be determined.

4 The My Resources list displays the following information about each resource:

Parameter	Description
Name	The name assigned to the resource.
Distinguished Name	(Not displayed by default) The dot-delimited fully distinguished name of the resource, such as `resourcename.clustername.context.domain`.
Cluster	The name of cluster for the resource.
Lives	The number of times the resource has been brought online.
Up Since	The date and time the resource was brought online.
BCC	A check mark in this column indicates that the resource is enabled for a Business Continuity Cluster (BCC).

5 (Optional) Personalize the display as described in Section 10.2.3, “Personalizing the My Resources Display,” on page 167.

10.2.3 Personalizing the My Resources Display

The My Resources page provides menus in the column headings that allow you to personalize the display. You can sort the entries, modify the columns, or filter the entries. Your preferences are saved between logins to iManager.

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Resources**.

3 Select **Refresh**, then choose a refresh rate that is long enough to allow you to view all items in the list.

4 Click or mouse-over a column heading to activate its options, then click the arrow to access the menu for that column.

5 Perform one of the available actions when you access the column's menu. Some actions are not available for all fields.

- **Sort:** Select **Sort Ascending** (A to Z) or **Sort Descending** (Z to A) to specify the preferred sort order. Resource objects are sorted based on values in the selected column. Numbers are treated as text values and are sorted alphabetically, not numerically.
- **Columns:** Select **Columns** to display the parameters, select or deselect a parameter's check box to add or remove the column, then press Enter to apply the changes. The parameters are described in Step 4 in Section 10.2.2, "Viewing Information about Resources in a My Resources List," on page 166.
- **Filter:** Specify a filter to display only the Cluster objects with values that match.

  To add or modify a column's filter, select **Filter**, place the cursor in the **Filter** field, type the desired filter text, then press Enter to apply the filter. If a cluster's parameter value matches the specified filter, the Cluster Resource object and its information are displayed.

  To remove a column's filter, select Filter, place the cursor in the **Filter** field, delete the filter text, then press Enter to refresh the display. The filter no longer applies for that column.

10.2.4 Removing Resources from a My Resources List

1 Log in to iManager as a cluster administrator.

2 In **Roles and Tasks**, select **Clusters > My Resources**.

3 Select the check box next to the resource that you want to remove from your personalized list, then click **Remove**.

  The resource is removed from your personalized list. This action does not delete the Cluster Resource object or its data.

10.2.5 Updating the My Resources List after Renaming or Deleting a Resource

In a **My Resources** list, if you rename or delete a resource, its status changes to `Critical: Resource not found`. You must manually update the resource list.

To update the list after deleting a resource:

1 In iManager, go to **Clusters > My Resources**.

2 In the **My Resources** list, select the check box next to the resource entry for the deleted resource, then click **Remove**.

To update the list after renaming a resource:

1 In iManager, go to **Clusters > My Resources**.

2 In the **My Resources** list, select the check box next to the resource entry, then click **Remove**.

3 Click **Add**, select the newly named resource, then click **OK**.

10.3 Using Cluster Resource Templates

Templates simplify the process of creating similar or identical cluster resources. For example, templates are helpful when you want to create multiple instances of the same resource on different servers. Several templates are provided for you. You can also create templates for any server application or resource you want to add to your cluster.

- Section 10.3.1, “Default Resource Templates,” on page 169
- Section 10.3.2, “Viewing or Modifying a Resource Template in iManager,” on page 170
- Section 10.3.3, “Creating a Resource Template,” on page 171
- Section 10.3.4, “Synchronizing Locally Modified Resource Templates with Templates in eDirectory,” on page 173

10.3.1 Default Resource Templates

OES Cluster Services provides several cluster resource templates that you can use on physical servers, virtualization host servers, and virtual machine (VM) guest servers (DomU). You can also create your own templates or personalize the default templates by using the Clusters plug-in for iManager. See Section 10.3.3, “Creating a Resource Template,” on page 171. Third-party templates might also available for third-party applications; see the vendor documentation.

Table 10-1 identifies the cluster resource templates that can be used for OES services and storage on physical servers and virtual machine guest servers:

Table 10-1 *Cluster Resource Templates for Physical Servers and VM Guest Servers*

Cluster Resource Template	OES 11 and Later Product
DHCP	Novell Dynamic Host Configuration Protocol using an NSS pool Novell Dynamic Host Configuration Protocol using a Linux POSIX File System
DNS	Novell Domain Name System
Generic File System (Generic_FS)	LVM volume groups
Generic IP Service	This template can be modified to create cluster resources for certain server applications that run on your cluster.
iPrint	Novell iPrint
MySQL	MySQL
Samba	
Cloud Integrated Storage (CIS_Template)	Cloud Integrated Storage
Third-party templates	See your vendor documentation.

Cluster Services provides the following templates for use by the Xen VM host server (Dom0). They are the only two templates supported for use in the virtualization host server environment. They are not supported for use by a VM guest server. The templates are generic and can be modified to start or stop different virtualization host server software.

Table 10-2 *Cluster Resource Templates for Virtualization Host Environments*

Cluster Resource Template	Use
Xen	Use to configure the cluster resource for a virtual machine. This template is available only on a virtualization host server. See Section 14.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 339.
XenLive	Use to configure the cluster resource for a virtual machine. It provides an additional function to allow a virtual machine resource migration (manual) without the need to boot or bring up the virtual machine on the cluster node where the virtual machine has been migrated. This template is available only on a virtualization host server. See Section 14.2.3, “Setting Up Live Migration,” on page 347.

10.3.2 Viewing or Modifying a Resource Template in iManager

You can view or manage the OES Cluster Services resource templates from the Cluster Options page in the Clusters plug-in for iManager. Ensure that you use the Clusters plug-in to modify the properties of the resource template. The modified settings apply for all subsequent cluster resources that you create from the template.

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Select the **Cluster Options** tab to access the templates.
4. Click the template name to view or modify the properties for the template.
5. Click the tabs to view or modify the default property settings for the resource template.

   The process for modifying a resource template is similar to modifying the settings for a cluster resource. If you modify a setting on the page, click **Apply** before you click the next tab in order to save the changes.

   - **Policies**

     Use the **Policies** tab to view or modify the default policies to use for the resource template. See Section 10.9, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 182.

- **Monitoring**

  Use the **Monitoring** tab to enable or disable monitoring. If monitoring is enabled, you can set the default behaviors for monitoring, including the interval to poll the resource's health, and the action to take if the resource fails to load on the maximum number of local restarts. See Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 176.

- **Preferred Nodes**

  Use the **Preferred Nodes** tab to view or modify the preferred nodes to use for cluster resources that are created with the selected template. See Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

- **Scripts**

  Use the **Scripts** tab to view or modify the default load, unload, and monitor scripts for the selected template. See the following:

  - Section 10.5, "Configuring a Load Script for a Cluster Resource," on page 174
  - Section 10.6, "Configuring an Unload Script for a Cluster Resource," on page 175
  - Section 10.7.4, "Example Monitor Scripts," on page 180

- **Business Continuity**

  The **Business Continuity** tab contains information only if the selected cluster is enabled for OES Business Continuity Clustering. If the page is enabled, you can view or set the default settings for resources created from the template. See the OES Business Continuity Clustering documentation website (http://www.novell.com/documentation/bcc).

6 Click OK

10.3.3 Creating a Resource Template

Templates help ensure that all of the necessary definition information, dependent services to be loaded, unloaded, or monitored, and the shared service or storage are entered correctly. You can use the default templates as a guide for what types of information to include in your personalized template, as well as the sequence of commands for loading, unloading, or monitoring a resource.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

  If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

3 Select the **Cluster Options** tab to access the templates.

4 Click the **New** link.

5 Specify **Template** as the resource type you want to create by clicking the **Template** radio button, then click **Next**.

6 In **Cluster Resource Name**, specify the name of the template you want to create.

7 If desired, in **Inherit from Template**, browse to the Cluster object and select the existing resource template in the Cluster container that you want to personalize for the new template.

8 Ensure that the **Define Additional Properties** check box is selected, then click **Next** to continue to the Load Script page.

9 On the Load Script page, configure the load script for the cluster resource template.

9a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

9b Edit or add any lines to the load script that are required to load dependent services such as web servers or file access protocols.

9c Edit or add the necessary commands to the script to load the resource on the server.

For example, this might include bind command for the NCP service and the mount commands for the shared disks and file systems.

9d Specify the default **Load Script Timeout** value, then click **Next** to continue to the Unload Script page.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

10 On the Unload Script page, configure the unload script for the cluster resource template.

10a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

10b Edit or add the necessary commands to the script to unload the resource from the server.

For example, this might include unbind command for the NCP service and the dismount commands for the shared disks and file systems.

10c Edit or add any lines to the unload script that are required to unload the dependent services that are loaded by this cluster resource.

10d Specify the default **Unload Script Timeout** value, then click **Next** to continue to the Monitor Script page.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose when migrating to another node. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

11 On the Monitor Script page, configure the monitor script for the cluster resource template.

11a Edit or add the variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

11b Edit or add the necessary commands to the script to monitor the resource on the server.

You can use the same commands that are used at the Linux terminal console.

The resource templates included with OES Cluster Services for Linux include resource monitor scripts that you can customize.

11c Specify the default **Monitor Script Timeout** value, then click **Next**.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action the administrator chooses for monitoring (comatose, migrate, or reboot) initiates. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

12 On the Resource Policies page, specify the default **Start**, **Failover**, and **Failback** modes, then click **Next.**

13 On the Resource Preferred Nodes page, specify the node assignments for the resource template, then click **Finish**.

The template you created is saved to the Cluster container of the cluster you selected. If you personalized an existing template, both the old template and the new template are in the container.

10.3.4 Synchronizing Locally Modified Resource Templates with Templates in eDirectory

You can also view the resource templates outside of iManager. They are cached as `/var/opt/novell/ncs/*_Template.*` files on the master node of the cluster. Typically, you modify the templates by using the Clusters plug-in to iManager as described in Section 10.3.2, "Viewing or Modifying a Resource Template in iManager," on page 170. However, if you add new template files or modify the files locally on the master cluster node, you must synchronize those changes with the resource templates and scripts that are held in the Cluster container in eDirectory.

To synchronize the locally modified resource templates with those held in the Cluster container in eDirectory,

1 Log in to the master node of the cluster as the `root` user, then open a terminal console.

2 At the command prompt, enter the following commands:

```
/opt/novell/ncs/bin/ncstempl.py
```

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

10.4 Creating Cluster Resources

Cluster resources must be created for every shared file system or server that you run on servers in your cluster. Cluster resources can include websites, email servers, databases, and any other server-based applications or services you want to make available to users at all times.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

3 Select the **Cluster Options** tab.

4 On the Cluster Options page, click the **New** link.

5 Specify Resource as the resource type you want to create by clicking the **Resource** radio button, then click **Next**.

6 Specify the name of the resource you want to create.

NOTE: Do not use periods in cluster resource names. Client for Open Enterprise Servers interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

---

7 In the **Inherit From Template** field, specify one of the available templates, such as the **Generic_FS_Template**.

For information about cluster resource templates, see Section 10.3, “Using Cluster Resource Templates,” on page 169.

8 Select the **Define Additional Properties** check box, then click **Next**.

9 If you are creating a new cluster resource, continue with “Configuring a Load Script for a Cluster Resource” on page 174.

10.5 Configuring a Load Script for a Cluster Resource

A load script is required for each resource, service, disk, or pool in your cluster. The load script specifies the commands to start the resource or service on a server.

Example load scripts are available in the following sections:

- Section 12.6, “Configuring a Load Script for the Shared NSS Pool,” on page 222
- “Sample Generic LVM Resource Load Script” on page 306

If you are creating a new cluster resource, the load script page should already be displayed. You can start with Step 6.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Select the **Cluster Options** tab.

4 Select the cluster resource to view its Properties page.

You can also select the check box next to the cluster resource, then click **Details**.

5 On the Properties page, click the **Scripts** tab, then click the **Load Script** link.

6 Edit or add the necessary commands to the script to load the resource on a server.

You can then add any lines to the load script that are required to load needed services like web servers, and so on.

You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, volume group name, file system type, and device.

---

7 Specify the **Load Script Timeout** value, then click **Apply** to save the script or, if you are creating a new cluster resource, click **Next**.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

8 Do one of the following:

- If you are configuring a new resource, click **Next**, then continue with Section 10.6, “Configuring an Unload Script for a Cluster Resource,” on page 175.
- Click **Apply** to save your changes.

  Changes for a resource’s properties are not applied while the resource is loaded or running on a server. Apply the updated script by taking the resource offline and then bringing it online on the same node. Alternatively, after the updated scripts have been synchronized from eDirectory to the source and destination nodes, the updated scripts are used automatically on system failover or cluster migration. For more information, see Section 10.8, “Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration,” on page 181.

10.6 Configuring an Unload Script for a Cluster Resource

Depending on your cluster application or resource, you can add an unload script to specify how the application or resource should terminate. An unload script is not required by all resources, but is required for cluster-enabled Linux partitions. Consult your application vendor or documentation to determine if you should add commands to unload the resource.

Programs should be unloaded in the reverse order of how they were loaded. This ensures that supporting programs are not unloaded before programs that rely on them in order to function properly.

Example unload scripts are available in the following sections:

- Section 12.7, “Configuring an Unload Script for the Shared NSS Pool,” on page 224
- “Sample Generic LVM Resource Unload Script” on page 307

If you are creating a new cluster resource, the unload script page should already be displayed. You can start with Step 6.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

  If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Select the **Cluster Options** tab.

4 Select the cluster resource to view its Properties page.

  You can also select the check box next to the resource, then click **Details**.

5 On the Properties page, click the **Scripts** tab, then click the **Unload Script** link.

6 Edit or add the necessary commands to the script to unload the resource on the server.

  You can add any lines to the unload script that are required to unload services that are loaded by this cluster resource.

  You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

7 Specify the **Unload Script Timeout** value, then click **Apply** to save the script or, if you are creating a new cluster resource, click **Next**.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose when migrating to another node. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

8 Do one of the following:

- If you are configuring a new resource, click **Next**, then continue with Section 10.7.2, “Configuring Resource Monitoring,” on page 178.
- Click **Apply** to save your changes.

  Changes for a resource’s properties are not applied while the resource is loaded or running on a server. Apply the updated script by taking the resource offline and then bringing it online on the same node. Alternatively, after the updated scripts have been synchronized from eDirectory to the source and destination nodes, the updated scripts are used automatically on system failover or cluster migration. For more information, see Section 10.8, “Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration,” on page 181.

10.7 Enabling Monitoring and Configuring the Monitor Script

Resource monitoring allows OES Cluster Services to detect a the resource failure independently of its ability to detect node failures. Monitoring is disabled by default. It is enabled separately for each cluster resource.

- Section 10.7.1, “Understanding Resource Monitoring,” on page 176
- Section 10.7.2, “Configuring Resource Monitoring,” on page 178
- Section 10.7.3, “Monitoring Services that Are Critical to a Resource,” on page 179
- Section 10.7.4, “Example Monitor Scripts,” on page 180

10.7.1 Understanding Resource Monitoring

When you enable resource monitoring, you must specify a polling interval, a failure rate, a failure action, and a timeout value. These settings control how error conditions are resolved for the resource.

- “Polling Interval” on page 177
- “Failure Rate” on page 177
- “Failure Action” on page 177
- “Timeout Value” on page 178
- “How Resource Monitoring Works” on page 178

Polling Interval

The monitor script runs at a frequency specified by the polling interval. By default, it runs every minute when the resource is online. You can specify the polling interval in minutes or seconds. The polling interval applies only to a given resource.

Failure Rate

The failure rate is the maximum number of failures (**Maximum Local Failures**) detected by the monitor script during a specified amount of time (**Time Interval**).

A failure action is initiated when the resource monitor detects that the resource fails more times than the maximum number of local failures allowed to occur during the specified time interval. For failures that occur before it exceeds the maximum, Cluster Services automatically attempts to unload and load the resource. The progress and output of executing a monitor script are appended to `/var/opt/novell/log/ncs/<resource_name>.monitor.out` file.

For example, if you set the failure rate to 3 failures in 10 minutes, the failure action is initiated if it fails 4 times in a 10 minute period. For the first 3 failures, Cluster Services automatically attempts to unload and load the resource.

Failure Action

The **Failover Action** indicates whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. With resource monitoring, the **Start**, **Failover**, and **Failback** Modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back (fail back) to the node it migrated from unless you manually migrate it back.

Set Resources as Comatose: (Default) If the failure action initiates, the resource is placed in a comatose state. Administrator action is required to take the resource offline, resolve the issue, and bring it online again on the same or different node.

Migrate the Resource Based on the Preferred Nodes List: If the failure action initiates and the resource is on its most preferred node, the resource migrates to the next available node in its **Preferred Nodes** list, which you previously ordered according to your failover order preferences. The resource is not automatically failed back to the original node. Administrator action is required to cluster migrate the resource to the node, as desired. Each time a failure action triggers a failover, the resource migrates to a different node, according to the order in its **Preferred Nodes** list and availability of the nodes.

Reboot the Hosting Node without Syncing or Unmounting Disks: If the failure action initiates, each of the resources on the hosting node will fail over to the next available node in its **Preferred Nodes** list because of the reboot. All resources on the node are failed over. This is a hard reboot, not a graceful one. The reboot option is normally used only for a mission-critical cluster resource that must remain available. The resources are not automatically failed back to the original node. Administrator action is required to cluster migrate them back to the node, as desired.

Timeout Value

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the configured failure action is initiated. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

How Resource Monitoring Works

1. The monitor script runs at the frequency you specify as the polling interval.
2. There are two conditions that trigger a response by OES Cluster Services:
   - An error is returned. Go to Step 3.
   - The script times out, and the process fails. Go to Step 4.
3. Cluster Services tallies the error occurrence, compares it to the configured failure rate, then does one of the following:
   - **Total errors in the interval are less than or equal to the Maximum Local Failures:** Cluster Services tries to resolve the error by offlining the resource, then onlining the resource.

     If this problem resolution effort fails, Cluster Services goes to Step 4 immediately regardless of the failure rate condition at that time.
   - **Total errors in the interval are more than the Maximum Local Failures:** Go to Step 4.
4. Cluster Services initiates the configured failure action. Possible actions are:
   - Puts the resource in a comatose state
   - Migrates the resource to another server
   - Reboots the hosting node (without synchronizing or unmounting the disks)

10.7.2 Configuring Resource Monitoring

The resource monitoring function allows you to monitor the health of a specified resource by using a script that you create or customize. If you want OES Cluster Services to check the health status of a resource, you must enable and configure resource monitoring for that resource. Enabling resource monitoring requires you to specify a polling interval, a failure rate, a failure action, and a timeout value.

If you are creating a new cluster resource, the Monitor Script page should already be displayed. You can start with Step 6.

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Select the **Cluster Options** tab.
4. Click the cluster resource to open its Properties page.

   You can also select the check box next to the resource, then click **Details**.

5 On the Properties page, click the **Monitoring** tab.

6 Select the **Enable Resource Monitoring** check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

7 For the polling interval, specify how often you want the resource monitor script for this resource to run.

You can specify the value in minutes or seconds.

8 Specify the number of failures (**Maximum Local Failures**) for the specified amount of time (**Time Interval**).

See "Failure Rate" on page 177.

9 Specify the **Failover Action** by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

See "Failure Action" on page 177.

10 Click the **Scripts** tab, then click the **Monitor Script** link.

11 Edit or add the necessary commands to the script to monitor the resource on the server.

The resource templates included with Cluster Services for Linux include resource monitor scripts that you can customize.

You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

You can use the same commands that would be used at the Linux terminal console. For example, see Section 10.7.3, "Monitoring Services that Are Critical to a Resource," on page 179.

12 Specify the **Monitor Script Timeout** value, then click **Apply** to save the script.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action you chose in Step 9 initiates.

13 Do one of the following:

- If you are configuring a new resource, click **Next**, then continue with Section 10.9.2, "Setting the Start, Failover, and Failback Modes for a Resource," on page 183.
- Click **Apply** to save your changes.

  Changes for a resource's properties are not applied while the resource is loaded or running on a server. Apply the updated script by taking the resource offline and then bringing it online on the same node. Alternatively, after the updated scripts have been synchronized from eDirectory to the source and destination nodes, the updated scripts are used automatically on system failover or cluster migration. For more information, see Section 10.8, "Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration," on page 181.

10.7.3 Monitoring Services that Are Critical to a Resource

In addition to monitoring the clustered service or storage objects, a resource monitor script can be used to monitor the status of services that are critical to the resource, such as Linux User Management (LUM), eDirectory, or other services.

IMPORTANT: NCS provides the ability to monitor the status of the eDirectory daemon (ndsd) at the NCS level. It is disabled by default. The monitoring can be set independently on each node. It runs whenever NCS is running on the node. See Section 8.8, “Configuring NCS to Monitor the eDirectory Daemon (ndsd),” on page 111.

If you enable NDSD monitoring at the NCS level, we recommend that you remove (or comment out) the eDirectory status check in individual monitor scripts to avoid excessive checking.

The resource monitor script runs only on the cluster server where the cluster resource is currently online. The script does not monitor the critical services on its assigned cluster server when the resource is offline. The monitor script does not monitor critical services for any other cluster node. Each monitor script acts independently, so be aware of the potential traffic you might generate if you include a status check in multiple monitoring scripts.

LUM Monitoring Example

You can use the `systemctl status namcd.service` command to monitor whether the Linux User Management service is running. Add the following command to the resource monitor script:

```
# (optional) status of the Linux User Management service
exit_on_error systemctl status namcd.service
```

Alternatively, you use the `namcd status` command to monitor whether the Linux User Management service is running and to automatically restart `namcd` if its daemon is not running. However, `namcd` creates messages in `/var/log/messages` with each check. Add the following command to the resource monitor script:

```
# (optional) status of the LUM service and restart if it is not loaded or
running
exit_on_error namcd status
```

eDirectory Monitoring Example

You can monitor the state of the eDirectory service in an individual monitoring script by adding the following command.

```
# (optional) status of the eDirectory service
exit_on_error systemctl status ndsd.service
```

For OES 11 SP2 and later, you can alternatively monitor NDSD at the NCS level on each node. See Section 8.8, “Configuring NCS to Monitor the eDirectory Daemon (ndsd),” on page 111.

10.7.4 Example Monitor Scripts

The resource templates included with OES Cluster Services for Linux include resource monitor scripts that you can customize.

Example monitor scripts are available in the following sections:

- **Clustered NSS pool:**
  - Section 12.8, “Configuring a Monitor Script for the Shared NSS Pool,” on page 226
- **Clustered LVM volume group:**
  - “Sample LVM Resource Monitor Scripts Created by NSS Management Tools” on page 294

- "Sample Generic LVM Resource Monitor Script" on page 308
- "Sample LVM Resource Monitor Script with an NCP Virtual Server" on page 314
- "Sample LVM Resource Monitor Script with an NCP Volume" on page 320

10.8 Applying Updated Resource Scripts by Offline/Offline, Failover, and Migration

For OES 11 SP2 (with the latest patches) and later, when you modify the load, unload, and monitor scripts for a resource and click **Apply** or **OK**, the updated resource scripts can take effect in any of the following ways:

- When an administrator takes the resource offline and then brings it online on the same node.
- When the resource fails over to another node, if both the source and destination nodes have received the updated scripts from eDirectory.
- When an administrator uses the `cluster migrate` command to load the resource on another node, if both the source and destination nodes have received the updated scripts from eDirectory.

In earlier releases, script changes do not take effect until the resource is taken offline and then brought online.

The ability of the scripts to be applied on fail-over or migration to a different node depends on the synchronization of the updated scripts from the master node to eDirectory, and from eDirectory to the nodes in the cluster. The destination node will not run the updated scripts unless the updated scripts are visible on both the source node (where the resource is currently running) and the destination node.

To verify that the updated scripts can take effect on the destination node:

1. In a text editor, verify that both the source node and destination node have the updated scripts (*`resource_name`*`.load`, *`resource_name`*`.unload` and *`resource_name`*`.monitor`) under the `/var/opt/novell/ncs/` directory.

   If the scripts are the same on both nodes, the updated scripts are applied on failover or migration between those nodes.

2. If the scripts are not the same on both nodes, pull the scripts down to the nodes from eDirectory:

   2a Ensure that eDirectory is fully synchronized.

   2b On the source node, pull down the scripts from eDirectory. Log on as the `root` user, open a terminal, then enter

   ```
   /opt/novell/ncs/bin/ncs-configd.py –init
   ```

   2c Repeat Step 2b for the destination node.

3. Repeat Step 1 to verify that both the source node and destination node have the updated scripts.

10.9 Configuring the Start, Failover, and Failback Modes for Cluster Resources

You can configure the start, failover, and failback of cluster resources to happen manually or automatically.

IMPORTANT: Cluster Services works with NCP user connections so that the user data sessions are resumed after failover. However, non-NCP users might experience service interruption and need to reconnect to the server after the failover. Applications using server based storage must be restarted on the client even with NCP unless they are NCP reconnect aware.

- Section 10.9.1, “Understanding Cluster Resource Modes,” on page 182
- Section 10.9.2, “Setting the Start, Failover, and Failback Modes for a Resource,” on page 183

10.9.1 Understanding Cluster Resource Modes

With the resource **Start** mode set to AUTO, the resource automatically starts on a server when the cluster is first brought up. If the resource **Start** mode is set to MANUAL, you can manually start the resource on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

With the resource **Failover** mode set to AUTO, the resource automatically starts on the next server in the Preferred Nodes list if there is a hardware or software failure. If the resource **Failover** mode is set to MANUAL, you can intervene after a failure occurs and before the resource is moved to another node.

With the resource **Failback** mode set to DISABLE, the resource does not fail back to its most preferred node when the most preferred node rejoins the cluster. If the resource **Failback** mode is set to AUTO, the resource automatically fails back to its most preferred node when the most preferred node rejoins the cluster. Set the resource **Failback** mode to MANUAL to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

The preferred node is the first server in the **Preferred Nodes** list for the resource.

IMPORTANT: Resources fail back only to the first node in their **Preferred Nodes** list. For example, if a resource has failed over to three servers since it originally ran on its preferred node, and the second server the resource was running on comes back up, the resource does not fail back to that second server.

Resources do not automatically move from node to node just because a node higher in the **Preferred Nodes** list rejoins the cluster, unless the **Failback** mode is set to AUTO and the first node in the **Preferred Nodes** list rejoins the cluster.

10.9.2 Setting the Start, Failover, and Failback Modes for a Resource

If you are creating a new cluster resource, the Resource Policies page should already be displayed. You can start with Step 6.

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Select the **Cluster Options** tab.
4. Select the cluster resource to open its Properties page.

   You can also select the check box next to the cluster resource, then click **Details**.
5. On the Properties page, click the **Policies** tab.
6. (Optional) Select the **Resource Follows Master** check box if you want to ensure that the resource runs only on the master node in the cluster.

   If the master node in the cluster fails, the resource fails over to whichever node becomes the master.
7. (Optional) Select the **Ignore Quorum** check box if you don’t want the cluster-wide timeout period and node number limit enforced.

   The quorum default values were set when you installed Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

   Selecting this box ensures that the resource is launched immediately on any server in the Preferred Nodes list as soon as any server in the list is brought online.
8. Specify the **Start**, **Failover**, and **Failback** modes for this resource.

   The default for both **Start** and **Failover** modes is AUTO, and the default for **Failback** mode is DISABLE.
9. Do one of the following:
   - If you are configuring a new resource, click **Next**, then continue with Section 10.10, “Configuring Preferred Nodes and Node Failover Order for a Resource,” on page 184.
   - Click **Apply** to save your changes.

     Changes for a resource’s properties are not applied while the resource is loaded or running on a server. You must offline, then online the resource to activate the changes for the resource.

10.10 Configuring Preferred Nodes and Node Failover Order for a Resource

The Preferred Nodes page allows you to assign nodes to use for the resource. You also sequence the list of nodes to specify the preferred order that the nodes will be tried when a resource is brought online after its current node fails or if you do not specify a node when you bring the resource online. Changes are not allowed for the Preferred Nodes list for the `Master_IP_Address_Resource`.

By default, a new resource is configured to run on all nodes, including future nodes. If its preferred node list is ever changed (directly or indirectly) by administrators, future nodes are not automatically added to the resource's preferred nodes list.

IMPORTANT: Ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

If you are creating a new cluster resource, the Preferred Nodes page should already be displayed. If you are assigning nodes for an existing resource, the Preferred Nodes page is displayed as part of the Resource Policies page. You can start with Step 6.

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.
3. Select the **Cluster Options** tab.
4. Select the cluster resource to open its Properties page.

   You can also select the box next to the cluster resource, then click the **Details** link.
5. On the Properties page, click the **Preferred Nodes** tab.
6. From the **Unassigned Nodes** area, select a node that the resource can use, then click the right-arrow button to move the selected node to the **Assigned Nodes** area.

   Repeat this step for all of the cluster nodes you want to assign to the resource.
7. From the **Assigned Nodes** area, select a node that you want to unassign from the resource, then click the left-arrow button to move the selected node to the **Unassigned Nodes** area.
8. Use one of the following methods to change the preferred failover order of the nodes assigned to a resource:
   - **Arrows:** Select one of the assigned nodes by clicking on it, then click the up-arrow and down-arrow buttons to move the node up or down in the list. The page refreshes between each click on an arrow.
   - **Edit:** Open the Preferred Nodes Edit function by clicking the **Edit** pen icon, list the assigned nodes in the preferred failover order with one node per line, then click **OK** to accept the revised order.

     Be careful to not alter the names of the nodes while editing the list. If you remove nodes from the list in the Edit view, the removed nodes are automatically added to the **Unassigned Nodes** area when you save.
9. At the bottom of the Preferred Nodes page, click **Apply** or **OK** to save the changes you made.

The new list and failover order take effect immediately for OES 2015 and later clusters, with the latest patches applied. It can take a few minutes for eDirectory to synchronize so the node can retrieve the updated list.

10 (OES 2023) Offline the resource, then online it again to apply the revised preferred nodes list and failover order.

11 (Optional) Use the following command on the node where the resource is running to list the preferred nodes for the resource, and verify that the resource has the updated list.

```
cluster migrate <resource_name> -list
```

10.11 Configuring Resource Priorities for Load Order

Cluster resource priorities control the load order of a resource relative to other cluster resources on the same cluster node when bringing up a cluster, or during a failover or failback. This is useful for ensuring that the most critical resources load first and are available to users before less critical resources.

The Resource Priority setting controls the order in which multiple resources start on a given node when the cluster is brought up or during a failover or failback. For example, if a node fails and two resources fail over to another node, the resource priority determines which resource loads first.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.

3 Select the **Cluster Options** tab.

4 On the Cluster Options page, click the **Properties** button above the **Cluster Objects** table to view the Cluster Properties dialog box.

5 In the Cluster Properties dialog box, click the **Priorities** tab.

6 Use either of the following methods to change the load order (from highest priority to lowest priority) of a resource relative to other cluster resources on the same node:

- **Arrows:** Select a resource in the list by clicking on it, then click the up-arrow or down-arrow buttons to move the resource up or down in the list. The page refreshes between each click on an arrow.
- **Edit:** Open the Resource Priority Edit function by clicking the **Edit** pen icon, list the resources in the preferred load order with one resource per line, then click **OK** to accept the revised order.

  Be careful to not alter the names of the resources while editing the list.

7 Click **Apply** or **OK** to save changes and close the Cluster Properties dialog box.

10.12 Configuring Resource Mutual Exclusion Groups

The Resource Mutual Exclusion (RME) Groups feature allows you to define sets of resources that must not run on the same node at the same time. You might need to use RME Groups for resources when an application does not want its resource to run concurrently on the same node as other

applications. You can also use RME Groups if you run multiple instances of a service in the same cluster and they should run on different nodes. This Clusters plug-in feature is available only for OES 2 SP3 and later servers.

For example, because of the designs and supported configurations of OES iPrint and GroupWise, the GroupWise and iPrint resources should not run on the same node at the same time. In addition, multiple iPrint server instances cannot be hosted on the same node. To avoid resource assignment conflicts, you could create an RME Group that consists of the GroupWise resource and the multiple iPrint resources.

- Section 10.12.1, “Understanding Resource Mutual Exclusion,” on page 186
- Section 10.12.2, “Setting Up RME Groups,” on page 187
- Section 10.12.3, “Viewing RME Group Settings,” on page 188

10.12.1 Understanding Resource Mutual Exclusion

Resource Mutual Exclusion is the practice of letting a node run only one of the resources in an RME Group at a time. The resources can run concurrently on different nodes.

- “How RME Groups Work” on page 186
- “Rules for Node Assignment” on page 186
- “Examples” on page 187
- “Planning for RME Groups” on page 187

How RME Groups Work

When RME Groups exist, OES Cluster Services does the following:

- It does not allow resources that are members of the same RME Group to run on the same node at the same time.
- It evaluates the exclusions together, but separately enforces each group of exclusions.

Rules for Node Assignment

When a resource is brought online, Cluster Services honors the resource settings in the following order:

1. Resource Follows Master
2. Resource Mutual Exclusion Group
3. Preferred Nodes list

When onlining, migrating, or failing over a resource onto a node, Cluster Services evaluates RME settings to check if the resource is a member of any RME Groups, then assesses whether any one of the other resources in the RME Groups that the resource belongs to is already running there. If another resource from the same RME Group is running on the node, the situation is handled as if the node was not on the resource’s Preferred Nodes list. The checking process repeats itself with another node as a candidate until a suitable node is found, or until the resource’s Preferred Nodes list is exhausted.

Resources that are not members of any of the RME Groups are not restricted. They are allowed to run at any time and in any combination on any node in their Preferred Nodes lists.

Examples

If you try to cluster migrate a resource to a node where another member of its RME group is running, the resource is not migrated and remains in the Running state on the node where it is currently loaded. Only one resource in the group is allowed to run on the node at a time.

If you try to cluster migrate multiple resources in an RME Group to the same node, the resources are brought online on different nodes. Only one of the resources (randomly selected) is brought online on the specified node. The other resources are migrated to different nodes in their Preferred Nodes lists.

Planning for RME Groups

An RME Group can contain any combination of the resources that are available to the cluster. Resources that are members of the same group are not allowed to run concurrently on a node. A resource can be a member of more than one group.

You can define up to four groups (Group A, Group B, Group C, and Group D). The group names are fixed; they cannot be custom. The selected check boxes in a column indicate which resources belong to that group.

RME Groups are supported when all cluster nodes are running a supported platform. If you are managing an older version of Cluster Services (without the RME code) from a new Clusters plug-in for iManager, the tables on the RME Groups page are shown as **No items**. You cannot set up RME Groups for clusters running on older platforms.

If you are upgrading from an OES 2 SP2 or NetWare cluster, the RMS Groups are not available. You should wait until the rolling cluster upgrade has been completed on all nodes before you define RME Groups. If you define RME Groups while the cluster is in a mixed-mode condition, the RME Groups are honored only if the resource fails over to a cluster node where the new OS is running.

10.12.2 Setting Up RME Groups

To define sets of resources that must not be assigned to the same node at the same time:

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. Select the **Cluster Options** tab.
4. On the Cluster Options page, click **Properties**, then click the **RME Groups** tab.

   The RME Groups page lists the resources in the cluster and columns for four possible RME Groups for the cluster. In each Group’s column, the check boxes correspond to the resources listed at the left.

   If you are managing an older version of Cluster Services (without the RME code) from a new Clusters plug-in for iManager, the tables are shown as **No items**. You cannot set up RME Groups for clusters running on older platforms.

5 To set up the members in a group, select the check boxes in the same column for two or more resources that must not be assigned on the same node at the same time.

The selected resources in a column are not allowed to run concurrently on a node. A resource can be a member of more than one group. You can define up to four groups (A, B, C, and D).

6 Click **OK** or **Apply** to save your changes.

10.12.3 Viewing RME Group Settings

The RME Groups settings are displayed in the following locations:

- "RME Groups in the Cluster Report" on page 188
- "RME Groups Configuration Page" on page 188

RME Groups in the Cluster Report

The Cluster Report has an RME Groups section that lists the member resources of each group.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster you want to manage.

3 Click **Run Report**, then scroll down to view the **RME Groups**.

The RME Groups section lists only the resources in the cluster that are members in each of the RME Groups.

For example, the following report shows the members for Groups A, B, and D. Group C has no defined members. The AUTO_POOL_12_SERVER resource belongs to two RME Groups. It has mutually exclusive relationships with two different groups of resources. These exclusions are managed separately.

Resource Mutual Exclusion (RME) Groups

Group A	Group B	Group C	Group D
AUTO_POOL_12_SERVER	AUTO_POOL_14_SERVER		AUTO_POOL_10_SERVER
AUTO_POOL_09_SERVER	AUTO_POOL_12_SERVER		AUTO_POOL_07_SERVER
AUTO_POOL_06_SERVER	AUTO_POOL_05_SERVER		Master_IP_Address_Resource
AUTO_POOL_02_SERVER			

RME Groups Configuration Page

The RME Groups configuration page displays all possible resources, and indicates membership by the check boxes that are selected in the same column under the RME Group Name.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

3 Select the **Cluster Options** tab.

4 Click **Properties**, then click the **RME Groups** tab.

The RME Groups configuration page lists the all of the resources in the cluster. The selected check boxes in the same group column indicate their memberships in the four possible RME Groups.

10.13 Controlling Resource Monitoring

You can use the `cluster monitor` command to start, stop, or view the status of monitoring for a specified cluster resource. Issue the command as the `root` user from the node where the resource is currently loaded:

```
cluster monitor <resourcename> {start | stop | status}
```

Monitoring must be enabled for the resource. For instructions, see Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

10.14 Changing the IP Address of a Cluster Resource

Typically, IP addresses for cluster resources are assigned for the life of the resource. You can modify the IP address of a cluster resource if desired. For example, you might need to assign a cluster and its resources to a different IP subnet.

Your setup must meet the following prerequisites:

- The resource must be in the **Offline** state in order to change its IP address.
- The new IP address must be unique in the network and available. It must satisfy the requirements described in Section 4.2, “IP Address Requirements,” on page 38.
- eDirectory must be running and working properly when you attempt to modify the IP address for the cluster resource.

Use the procedures in this section to modify the IP address for cluster resources:

- Section 10.14.1, “Pool Cluster Resource,” on page 189
- Section 10.14.2, “Non-Pool Cluster Resources,” on page 190

10.14.1 Pool Cluster Resource

You can change the IP address of a pool cluster resource by using the Protocols page in the resource’s Properties in the Clusters plug-in for iManager. Ensure that you take the pool cluster resource offline before attempting to change its IP address. When you save the change, the load, unload, and monitor scripts are automatically updated with the new IP address. It also automatically updates the resource’s IP address that is stored in its NCP Virtual Server object.

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster that you want to manage.

   If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 99.
3. On the Cluster Manager page, select the check box next to the pool cluster resource that you want to manage, then click **Offline**.

   Wait until the resource reports an Offline status before continuing.
4. Select the pool cluster resource to open its Properties page.
5. Select the **Protocols** tab.

6 In the IP Address field, type the new IP address for the cluster resource in IPv4 dotted decimal format (four octets of the address expressed individually in decimal and separated by periods).

For example, type:

```
10.10.10.43
```

7 Click **OK** to save and apply the changes.

The Cluster Options page opens. The new IP address is displayed for the pool cluster resource.

8 (Optional) Verify that the IP address has been changed in the pool cluster resource's load, unload, and monitor scripts.

8a On the Cluster Options page, select the cluster resource to open its Properties page, then select the **Scripts** tab.

8b Select the Load, Unload, and Monitor options in turn, and review the scripts to verify that the resource IP address has been changed in all of the relevant places.

9 Select the **Cluster Manager** tab.

10 On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Online**.

Wait until the resource reports an Online status before continuing to ensure that the resource is working properly.

10.14.2 Non-Pool Cluster Resources

The properties for non-pool cluster resources do not include a Protocols page. You can change the IP address of a non-pool cluster resource by changing the IP address value for the RESOURCE_IP variable in its load, unload, and monitor scripts.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.

3 On the Cluster Manager page, select the check box next to the non-pool cluster resource that you want to manage, then click **Offline**.

Wait until the resource reports an Offline status before continuing.

4 Select the cluster resource to open its Properties page.

5 Change the IP address value for the RESOURCE_IP variable in the load, unload, and monitor scripts:

5a Select the **Scripts** tab.

It automatically opens to the Load Script page.

5b On the Load Script page, type the new IP address in the RESOURCE_IP variable, then click **Apply**.

5c Click **Unload Script**, type the new IP address in the RESOURCE_IP variable, then click **Apply**.

5d Click **Monitor Script**, type the new IP address in the RESOURCE_IP variable, then click **Apply**.

5e Click **OK** to return to the Cluster Options page.

6 Select the **Cluster Manager** tab.

7 On the Cluster Manager page, select the check box next to the cluster resource, then click **Online**.

Wait until the resource reports an Online status before continuing to ensure that the resource is working properly.

10.15 Renaming a Cluster Resource

A `cluster rename` command is available that allows you to rename a cluster resource. Renaming the resource does not modify the virtual server name (the NCS:NCP Server object, such as `CL134-POOL1-SERVER`). Renaming the cluster resource modifies only the name of the Cluster Resource object; it does not rename the clustered pool or clustered LVM volume that is managed by the resource.

For an NSS pool cluster resource, renaming the cluster resource does not rename the pool and the resource scripts. For information about modifying the name of a cluster-enabled pool, see Section 12.12, “Renaming a Clustered NSS Pool,” on page 233.

For a Linux LVM cluster resource, renaming the cluster resource does not rename the LVM volume and update the resource scripts. For information about modifying the name of a cluster-enabled LVM volume, see Section 13.9, “Renaming a Clustered LVM Logical Volume,” on page 323.

Your setup must meet the following prerequisites:

- This command must be issued from the master node.
- The resource must be in the **Offline** state in order to be renamed.
- The new name must not exist prior to the renaming.
- eDirectory must be running when you attempt to rename cluster resource.

To rename the cluster resource:

1 Log in to the master node in the cluster as the `root` user, and open a terminal console.

2 Offline the cluster resource that you want to rename by entering

```
cluster offline <resource_name>
```

For example:

```
cluster offline POOL1_SERVER
```

You can use the `cluster status` command to verify that the resource reports an offline state before you continue.

3 Rename the cluster resource by entering

```
cluster rename <resource_name> <new_resource_name>
```

For example:

```
cluster rename POOL1_SERVER custom_name22
  Resource 'POOL1_SERVER' has been renamed to 'custom_name22'.
```

4 Online the cluster resource by entering

```
cluster online <new_resource_name>
```

For example:

```
cluster online custom_name22
```

10.16 Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service

You can delete a cluster resource for any of the following reasons:

- You want to re-create a Cluster Resource object. You want to remove the related objects in eDirectory and get the shared storage ready so that the resource can be re-created cleanly.

  To delete a resource and create a new one with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.
- You want to disable clustering for the shared storage managed by the resource. Afterwards, you will assign the SAN device to only a single node in the cluster. The storage will be available via the node's IP address instead of the cluster resource IP address.
- You want to disable clustering for the shared service. The service will be available via the node's IP address instead of the cluster resource IP address.
- You want to delete the shared storage managed by the resource. Afterwards, data is no longer available on any node in the cluster.

We strongly recommend that you delete cluster resource objects only from the master node in the cluster. Ensure that you take the cluster resource offline before you attempt to delete the cluster resource, or before you attempt to delete cluster-enabled shared storage managed by the resource.

WARNING: If you attempt to delete a storage cluster resource without first taking it offline, deletion errors occur, and the data associated with the clustered storage is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the **Storage > Pools** page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

Use the following procedure to delete a cluster resource:

1. If the resource is on a non-master node in the cluster, migrate it to the master node.

   As the *root* user, open a terminal console, then enter

   ```
   cluster migrate <resource_name> <master_node_name>
   ```

   The master node must be in the resource's preferred nodes list. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.
2. If the cluster resource is online, take it offline by using one of the following methods:
   - Enter the following at the command prompt as the `root` user:

     ```
     cluster offline <resource_name>
     ```

Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

- In iManager, go to **Clusters > My Clusters**, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click **Offline**.

  Refresh the Cluster Manager page to verify that the resource has a status of **Offline** before you continue.

3 In iManager, use the Clusters plug-in to delete the cluster resource.

3a Select **Clusters > My Clusters**, then select the cluster.

3b Select the **Cluster Options** tab.

3c Select the check box next to the resource, then click **Delete**.

This deletes the Cluster Resource object. It does not delete the storage or service represented by the object.

For a pool cluster resource, this also deletes the related NCP virtual server object, Pool object, and Volume object.

For an NCP-enabled LVM cluster resource, this also deletes the related NCP virtual server object and Volume object.

3d When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.

3e In the *Tree View* in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.

If necessary, you can manually delete the objects. In iManager, go to **Directory Administration > Delete Objects**, select the objects, then click **OK**.

4 If the deleted resource was a pool cluster resource, use the **Update eDirectory** function to re-create Storage objects for the pool and its volumes.

4a In iManager, select **Storage > Pools**, then select the master node if you plan to re-create the storage object, or select the node where you want the pool to reside as a locally available pool.

4b Select the pool, then click **Activate**.

4c Select the pool, then click **Update eDirectory**.

This creates a Pool object in eDirectory with a name format of *<server_name>_<pool_name>*_POOL.

4d Select **Storage > Volumes.** The server should still be selected.

4e Select the volume, then click **Mount**.

4f Select the volume, then click **Update eDirectory**.

This creates a Volume object in eDirectory with a name format of *<server_name>_<volume_name>*.

4g Repeat Step 4d through Step 4f for each volume in the pool.

4h In the Tree View, browse to verify that the Pool object and Volume object were created.

5 Do one of the following:

- **Re-create the cluster resource:** Use the Clusters plug-in to cluster-enable the storage area.
  - To cluster-enable a pool, see Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 217.
  - To cluster-enable a Linux LVM volume group and logical volume, see Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 300 and Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308.

  To re-create the cluster resource with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.
- **Use the non-clustered pool:** Unshare the devices that contribute space to the pool, then use a third-party SAN management tool to assign the devices to only the desired server.
  1. In iManager, go to **Storage** > **Devices**, then select the desired server (the one you specified in Step 4).
  2. Select the device.
  3. Deselect the **Shareable for Clustering** check box, then click **Apply**.

     Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the pool where you have disabled clustering.
  4. Repeat these steps for each device that contributes space to the pool.
  5. Use a third-party SAN management tool to assign the devices to only the desired server.
  6. Provide the node’s IP address to users.
- **Use the non-clustered Linux LVM volume:** The LVM volume group uses the entire device. Clustered LVM (cLVM) still recognizes the multiple-node assignments from the SAN.
  1. Use a third-party SAN management tool to assign the device to only the desired server.
  2. Provide the node’s IP address to users.
- **Use the non-clustered service:** Modify information for your users so that they access the local node’s IP address instead of a clustered service IP address.
- **Delete the pool or LVM volume group:** If you do not want to keep the data, delete the shared storage area.

  WARNING: Deleting a pool or a Linux LVM volume group destroys all data on it.

  - For shared NSS pools and volumes, use NSSMU, the Storage plug-in to iManager, or the `nlvm delete pool <pool_name>` command. Deleting the pool automatically deletes the volumes on it.
  - For Linux LVM volume groups, use NSSMU or the `nlvm delete linux volume <volume_name>` command. These tools automatically delete the LVM logical volume and logical volume group. If the volume was NCP-enabled, it also deletes the related NCP volume.

10.17 Additional Information for Creating Cluster Resources

Cluster resources can be configured for storage, services, or virtual machines.

- Section 10.17.1, “Creating Storage Cluster Resources,” on page 195
- Section 10.17.2, “Creating Service Cluster Resources,” on page 195
- Section 10.17.3, “Creating Virtual Machine Cluster Resources,” on page 195

10.17.1 Creating Storage Cluster Resources

For information about creating cluster resources for shared storage on Linux, see the following:

Table 10-3 *Cluster-Enabling Shared Storage*

Shared Storage	Refer to
NSS pools and volumes	Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 201
NCP volumes	“Configuring NCP Volumes with OES Cluster Services” in the *OES 2023: NCP Server for Linux Administration Guide*
Dynamic Storage Technology shadow volume pairs (shared NSS pools and volumes configured as DST pairs)	“Configuring DST Shadow Volume Pairs with OES Cluster Services” in the *OES 2023: Dynamic Storage Technology Administration Guide*

10.17.2 Creating Service Cluster Resources

For information about creating cluster resources for various services, see Chapter 11, “Quick Reference for Clustering Services and Data,” on page 197.

10.17.3 Creating Virtual Machine Cluster Resources

If you install OES Cluster Services at the host level of an OES virtualized server, you can create cluster resources for the virtual machines. See Section 14.2, “Virtual Machines as Cluster Resources,” on page 338.

11 Quick Reference for Clustering Services and Data

You can use OES Cluster Services to provide high-availability access to OES services and some Linux services. See the tables in this section for links to information about how to configure cluster resources for services and data.

- Table 11-1, “Clustering OES Services with OES Cluster Services,” on page 197
- Table 11-2, “Clustering Storage with OES Cluster Services,” on page 198
- Table 11-3, “Clustering Linux Services with OES Cluster Services,” on page 199

Table 11-1 *Clustering OES Services with OES Cluster Services*

OES Service	For clustering information, see
Certificate Server	The NetIQ Certificate Server is not cluster-enabled. The Certificate Server service issues Server Certificate objects that might need to reside on each node in a cluster, depending on the service that is clustered. See “eDirectory Server Certificates” in the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*.
DFS VLDB (Distributed File Services volume location database)	“Clustering Distributed File Services” in the *OES 2023: Distributed File Services Administration Guide for Linux*.
DHCP Server	In the *OES 2023: DNS/DHCP Services for Linux Administration Guide*, see: ◆ “Configuring DHCP with OES Cluster Services for the NSS File System” ◆ “Configuring DHCP with OES Cluster Services for the Linux File System”
DNS Server	“Configuring DNS with OES Cluster Services” in the *OES 2023: DNS/DHCP Services for Linux Administration Guide*.
eDirectory	eDirectory is not clustered because it has its own replica system.
File, CIFS (Windows File Services)	“Configuring CIFS with Cluster Services for an NSS File System” in the *OES 2023: OES CIFS for Linux Administration Guide*.
iPrint	“Configuring iPrint with OES Cluster Services” in the *OES 2023: iPrint Administration Guide*.

Table 11-2 *Clustering Storage with OES Cluster Services*

Storage	For clustering information, see
Storage, Dynamic Storage Technology (DST) shadow volume pairs built with NSS volumes	"Configuring DST Shadow Volume Pairs with OES Cluster Services" in the *OES 2023: Dynamic Storage Technology Administration Guide*.
Storage, LVM volume groups and logical volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 275.
Storage, NCP volumes	◆ Section 13.6, "Enabling NCP File Access for a Clustered LVM Volume," on page 314. ◆ "Configuring NCP Volumes with OES Cluster Services" in the *OES 2023: NCP Server for Linux Administration Guide*. ◆ "Clustering LVM Volume Groups with Novell Cluster Services" in the *OES 2023: Linux POSIX Volume Administration Guide*. ◆ "Create Linux Volume" in the *OES 2023: NLVM Reference*.
Storage, NSS pools and volumes	Chapter 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 201.
Cloud Integrated Storage (CIS)	"Installing and Configuring CIS Server with Cluster Services" in the OES 2023: CIS Administration Guide.

Table 11-3 *Clustering Linux Services with OES Cluster Services*

Linux Service	For clustering information, see
Apache Web Server	"Apache HTTP Server" in the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*.
File, FTP	"Cluster Enabling Pure-FTPd in an OES Environment" in the *OES 2023: Planning and Implementation Guide*.
MySQL	"Configuring MySQL with Novell Cluster Services" in the *OES 2015 SP1: Web Services and Applications Guide*. A MySQL template is available that uses a shared LVM volume group and logical volume that you have already created.
Xen virtual machines	Section 14.2, "Virtual Machines as Cluster Resources," on page 338. You can use XEN and XENLive resource templates to configure virtual machines as cluster resources that can be failed over between host (Dom0) environments.

12 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

After you have installed and configured OES Cluster Services on a server, you can configure OES Storage Services (NSS) pool cluster resources. This section describes how to create and cluster-enable shared NSS pools volumes as pool cluster resources with OES Cluster Services.

- Section 12.1, “Requirements for Creating Pool Cluster Resources,” on page 202
- Section 12.2, “Guidelines for Using Pool Cluster Resources,” on page 206
- Section 12.3, “Initializing and Sharing a SAN Device,” on page 207
- Section 12.4, “Creating Cluster-Enabled Pools and Volumes,” on page 210
- Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 217
- Section 12.6, “Configuring a Load Script for the Shared NSS Pool,” on page 222
- Section 12.7, “Configuring an Unload Script for the Shared NSS Pool,” on page 224
- Section 12.8, “Configuring a Monitor Script for the Shared NSS Pool,” on page 226
- Section 12.9, “Adding Advertising Protocols for NSS Pool Cluster Resources,” on page 227
- Section 12.10, “Adding NFS Export for a Clustered Pool Resource,” on page 228
- Section 12.11, “Mirroring and Cluster-Enabling Shared NSS Pools and Volumes,” on page 229
- Section 12.12, “Renaming a Clustered NSS Pool,” on page 233
- Section 12.13, “Renaming a Clustered NSS Volume,” on page 242
- Section 12.14, “Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume),” on page 250
- Section 12.15, “Changing the Name Space for a Clustered NSS Volume,” on page 255
- Section 12.16, “Adding a Volume to a Clustered Pool,” on page 256
- Section 12.17, “Changing an Assigned Volume ID,” on page 257
- Section 12.18, “Expanding the Size of a Clustered Pool,” on page 257
- Section 12.19, “Deleting NSS Pool Cluster Resources,” on page 266
- Section 12.20, “Disabling Clustering for a Pool,” on page 266
- Section 12.21, “Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool),” on page 268
- Section 12.22, “Deleting a Clustered Pool,” on page 271

12.1 Requirements for Creating Pool Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in Chapter 4, “Planning for OES Cluster Services,” on page 35.

- Section 12.1.1, “OES Cluster Services,” on page 202
- Section 12.1.2, “Resource IP Address,” on page 202
- Section 12.1.3, “eDirectory,” on page 202
- Section 12.1.4, “OES Storage Services,” on page 202
- Section 12.1.5, “Shared Storage,” on page 204
- Section 12.1.6, “NCP Server for Linux,” on page 204
- Section 12.1.7, “OES CIFS for Linux,” on page 205
- Section 12.1.8, “Domain Services for Windows,” on page 206

12.1.1 OES Cluster Services

OES Cluster Services must be installed, configured, and running when you share devices and when you create manage the shared NSS pools and volumes. The cluster must be active.

12.1.2 Resource IP Address

Each NSS pool cluster resource requires a unique static IP address. The IP address is used to provide access to and failover capability for the pool cluster resource. Users access the pool by using the resource IP address instead of the server IP address where the pool is active. 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.

12.1.3 eDirectory

OES Cluster Services requires that eDirectory be running and working properly when you cluster-enable a pool and when you manage the pool.

12.1.4 OES Storage Services

NSS must be installed and running on each server in the cluster. For information about installing NSS and managing NSS pools and volumes, see the *OES 2023: NSS File System Administration Guide for Linux*. For information about using NLVM commands, see *OES 2023: NLVM Reference*.

In addition, the following requirements must be met:

- “Pool” on page 203
- “Volume” on page 203
- “Naming Conventions” on page 203

- "Pool Cluster Resource Name" on page 203
- "Location of Cluster, Server, and Storage Objects" on page 204

Pool

We recommend that your pool cluster resource be one device, one pool, one volume. You can create a pool up to 8 TB in size on a single device by formatting the device in GPT format.

You can create and cluster-enable a shared NSS pool by using the Storage plug-in for iManager, the server-based NSS Management Utility (`nssmu`), or the OES Linux Volume Manager (NLVM) command line interface. You can also use these tools to create NSS volumes on the shared pool. These tools interact with OES Cluster Services to create the pool cluster resource.

You can cluster-enable an existing NSS pool by using the Clusters plug-in for iManager.

Volume

You must create at least one volume in the shared pool. Typically, you create all volumes for a shared pool when you set up the pool cluster resource and before you need to cluster migrate or fail over the resource to a different node in the cluster.

For information about using antivirus software with NSS volumes, see "Antivirus Support for NSS" in the *OES 2023: NSS File System Administration Guide for Linux*.

Naming Conventions

Pool names and volume names must comply with the same naming conventions. Names must be 2 to 15 characters in length. OES Cluster Services supports pool names with characters A to Z, 0 to 9, and underscores (_). The name cannot begin or end with an underscore; it cannot contain multiple adjacent underscores (__). It does not allow the following special characters in a shared pool name:

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

IMPORTANT: Before you cluster-enable an existing NSS pool, you must rename the pool and its volumes to use names that do not contain special characters.

Pool Cluster Resource Name

The pool name is used to form the default name of the pool cluster resource and the virtual server (the NCS:NCP Server object). If you modify the resource name, ensure that the name conforms to the resource naming conventions as described in Section 10.1.1, "Naming Conventions for Cluster Resources," on page 164.

Storage Object	Default Name
Resource name	*<cluster_name>_<poolname>*
Clustered Volume object name	*<cluster_name>_<volume_name>* *<cluster_name>*-*<poolname>*-SERVER*<volume_name>*

Storage Object	Default Name
Virtual server name	*<cluster_name>*-*<poolname>*-SERVER

Location of Cluster, Server, and Storage Objects

The Server, Pool, Volume, Cluster Resource, and Cluster objects are recommended to be in the same context (such as ou=ncs,o=novell). It is supported for the objects to be in the same context or different contexts. If the objects are in different contexts, you might need to cluster migrate the pool cluster resource back to the context where the pool was created in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.

12.1.5 Shared Storage

You should carefully plan how you want to configure your shared storage prior to installing OES Cluster Services. Consider the guidelines and requirements in the following sections when planning your NSS storage solution.

- Section 4.8, “Shared Disk Configuration Requirements,” on page 50
- Section 4.9, “SAN Rules for LUN Masking,” on page 53
- Section 4.10, “Multipath I/O Configuration Requirements,” on page 54

12.1.6 NCP Server for Linux

NetWare Core Protocol (NCP) is the OES networking protocol used by the Client for Open Enterprise Server. NCP is automatically selected as an advertising protocol when you cluster-enable an NSS pool. This is necessary to provide authenticated access to data using the OES Trustee model.

OES Storage Services requires that the NCP Server for Linux service be installed and running on each node in the server. NCP Server must be running even if users access volumes on the shared NSS pool only via other protocols.

WARNING: Cross-protocol file locking is required when using multiple protocols for data access on the same volume. This helps prevent possible data corruption that might occur from cross-protocol access to files.

The NCP Cross-Protocol File Lock parameter is enabled by default when you install NCP Server. If you modify the Cross-Protocol File Lock parameter, you must modify the setting on all nodes in the cluster.

NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource. If a file is opened with multiple protocols when the migration or failover begins, the file should be closed and reopened after the migration or failover to acquire cross-protocol locks on the new node.

See “Configuring Cross-Protocol File Locks for NCP Server” in the *OES 2023: NCP Server for Linux Administration Guide*.

NCP Server for Linux is installed by selecting **NCP Server and Dynamic Storage Technology** from the OES Services menu in the YaST install interface. For information about NCP Server for Linux, see the *OES 2023: NCP Server for Linux Administration Guide*.

You must specify an NCP virtual server name for the cluster-enabled pool. By default, the suggested name is in the form of <*clusterName*>-<*poolName*>-SERVER. You can accept the suggested name, or specify a custom name for the NCP virtual server name. If you are also enabling CIFS as an advertising protocol, you might want to choose an NCP virtual server name that can also be used as the CIFS virtual server name. For information about using the NCP virtual server name as the CIFS virtual server name, see Section 12.1.7, "OES CIFS for Linux," on page 205.

12.1.7 OES CIFS for Linux

Common Internet File System (CIFS) is the Windows networking protocol. OES CIFS allows you to give clients access via CIFS to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. See "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2023: NCP Server for Linux Administration Guide*.

OES CIFS must be installed, configured, and working properly before you can specify CIFS as an advertising protocol when you cluster-enable an NSS pool. Otherwise, the CIFS option is not available when you configure the resource. OES CIFS for Linux is installed by selecting **OES CIFS** from the OES Services menu in the YaST install interface. For information about OES CIFS for Linux, see the *OES 2023: OES CIFS for Linux Administration Guide*.

You must specify a CIFS virtual server name for the cluster enabled pool. You can accept the suggested name, or specify a custom name for the CIFS virtual server name. The name can be up to 15 characters, which is a restriction of the CIFS protocol.

For users to collaborate effectively, all paths for user access should be identical, independent of the access protocol used. This is possible only if the same name is used for the NCP virtual server name and the CIFS virtual server name, and the name can be only up to 15 characters.

By default, the NCP virtual server name is suggested as the CIFS virtual server name. If the name is more than 15 characters, the CIFS virtual server name uses the rightmost 13 characters and adds `-W`. For example, an NCP virtual server name of `CLUSTER1-P_USERS` is modified to `STER1-P_USERS-W` for the CIFS virtual server name. If a default NCP virtual server name was used in the form of <*clusterName*>-<*poolName*>-SERVER and the name exceeds 15 characters, the CIFS virtual server name uses the rightmost 13 characters of the <*clusterName*>-<*poolName*> part of the name, and adds `-W`. For example, an NCP virtual server name of `CLUS1-P123-SERVER` is modified to `CLUS1-P123-W` for the CIFS virtual server name.

If an administrator user later changes the NCP virtual server name for a clustered pool, NSSMU automatically modifies the CIFS virtual server name accordingly. To use a custom name for the CIFS virtual server name, you can modify the CIFS virtual server name by using the CIFS management tools. The custom CIFS name will not affect the NCP virtual server name.

12.1.8 Domain Services for Windows

Cluster-enabled NSS volumes can be used in a Domain Services for Windows environment.

12.2 Guidelines for Using Pool Cluster Resources

Consider the guidelines in this section when working with shared NSS pools and volumes in the cluster:

- Section 12.2.1, “Guidelines for Shared Pools,” on page 206
- Section 12.2.2, “Guidelines for Shared Volumes,” on page 206

12.2.1 Guidelines for Shared Pools

- When the pool cluster resource is brought online, the pool is automatically activated by the resource load script. You typically do not activate the pool at the terminal console.
- The unit of failover is the device (disk or LUN) that contains the pool. If you use multiple devices in the pool, all of those devices must be failed over together.

  IMPORTANT: As a best practice, we recommend a single LUN (or device) per shared pool, and one volume per pool. (one LUN=>one pool=>one volume)

- If you delete a pool cluster resource, OES Cluster Services automatically removes the Pool Resource object and the virtual server object from eDirectory. After you delete the resource, you can unshare the device and use the pool locally, or you can delete the pool and its volumes.

  Ensure that you offline the cluster resource before you attempt to delete the cluster resource or the pool. See Section 10.16, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 192.
- Before you rename a cluster-enabled pool, ensure that you offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, deactivate the pool, then online the pool resource.

  OES Cluster Services automatically updates the pool resource load and unload scripts to reflect the name change. Also, NSS automatically changes the Pool Resource object name in eDirectory.

12.2.2 Guidelines for Shared Volumes

You must create at least one volume on a shared pool before you migrate it. As a best practice, we recommend that you create only one volume per shared pool.

When you create a volume, commands are added to the pool resource load and unload scripts to automatically mount and dismount the volume when the scripts run. You can modify the load script to comment out the mount command so that you can manually mount the volume on a node in the cluster where the pool resource has been activated.

Volumes on cluster-enabled pools do not appear as independent cluster resources. A cluster-enabled volume does not have an associated load and unload script or an assigned IP address. If you want each volume to be in a separate cluster resource, each volume must have its own pool.

When a server fails, the cluster resources fail over to other servers in the cluster. Because the cluster-enabled pool fails over, all volumes in the pool also fail over, but only the volumes that have been added to the load script are automatically mounted. Any volumes in the pool that are not in the load script must be mounted manually. For this reason, volumes that you do not want to fail over should be in separate pools that are not cluster-enabled.

When you create an encrypted NSS volume in a shared pool, you must mount the volume manually by using NSSMU and enter the password. NSS uses the password to create a key. Instead of storing it in the server memory as it does for non-shared volumes, NSS asks OES Cluster Services to store the key and to pass it to the other nodes. After all servers hold the key, the volume is available for access as long as any one of the servers is still participating actively in the cluster. If all of the servers in the cluster fail, you must repeat this manual mounting procedure when you recover the cluster and restart services.

If you delete a volume from a cluster-enabled pool, OES Cluster Services automatically removes the volume mount command from the pool cluster resource load script.

12.3 Initializing and Sharing a SAN Device

NSS pools can be up to 8 TB in size and use space from one or more devices. For clustered pool resources, we recommend one device per pool. Each of the devices that you want to use for shared pool must be initialized and marked as shareable for clustering.

You must initialize a device (disk or LUN) to set up its device format before you can create an NSS pool on it. You can also initialize a device to wipe its current structure and reconfigure it. Only initialized devices are visible to the NSS management tools when you create a pool.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that have never been initialized have a format of None. Devices that are being used for a OES Cluster Services SBD (split brain detector) partition also have a format of None; however, you should not use initialization to remove an SBD partition. For information about removing an SBD partition, see Section 9.18, “Creating or Deleting Cluster SBD Partitions,” on page 146. After the SBD partition is removed, you can initialize the device.

You can initialize and share a device by using the Storage plug-in for iManager, the OES Storage Services (NSS) Management Utility (NSSMU), or the OES Linux Volume Manager (NLVM) command line interface.

- Section 12.3.1, “Initializing and Sharing a Device with iManager,” on page 208
- Section 12.3.2, “Initializing and Sharing a Device with NSSMU,” on page 208
- Section 12.3.3, “Initializing and Sharing a Device with NLVM,” on page 209

12.3.1 Initializing and Sharing a Device with iManager

1. Ensure that the SAN device is attached to all nodes in the cluster.
2. Log in to iManager as an administrator user, then select Storage > Devices.
3. Browse to select the Cluster object ( ) of the cluster.

   Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.
4. From the **Devices** list, select the device where you want to create the pool.
5. If the device has not been previously initialized or if you want to delete the current partitioning structures on the device, click **Initialize Disk**, then click **OK** to confirm and continue.

   WARNING: Initializing a disk destroys all of the data on it.
6. If you are prompted, specify whether to use the GPT or DOS partitioning scheme, then click **OK** to complete the initialization.

   GPT can be used for any size device. DOS supports devices up to 2 TB in size.

   Wait for the page to refresh before continuing.
7. In the **Details** area, select the **Shareable for Clustering** check box, then click **OK** to apply the change.
8. Repeat Step 4 through Step 7 for each device that you want to use in the pool.
9. Exit iManager.
10. Continue with Section 12.4.1, “Creating a Cluster-Enabled Pool and Volume with iManager,” on page 210.

12.3.2 Initializing and Sharing a Device with NSSMU

1. Ensure that the SAN device is attached to all nodes in the cluster.
2. Log in as the `root` user to the master node in the cluster, then open a terminal console.
3. At the command prompt, enter:

   ```
   nssmu
   ```
4. In the NSSMU, select **Devices** and press Enter.
5. In the **Devices** list, select the device where you want to create the pool.
6. If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3 to initialize the selected device, then press Y (Yes) to confirm and continue.

   WARNING: Initializing a disk destroys all of the data on it.
7. Specify whether to use the GPT or DOS partitioning scheme, then click **OK** to complete the initialization.

   GPT can be used for any size device. DOS supports devices up to 2 TB in size.

   Wait for the page to refresh before continuing.

8 Press F6 to mark the device as shareable for clustering.

   The **Shareable for Clustering** value changes from **No** to **Yes**.

9 Repeat Step 5 through Step 8 for each device that you want to use in the pool.

10 Exit NSSMU.

11 Continue with Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,” on page 213.

12.3.3 Initializing and Sharing a Device with NLVM

1 Ensure that the SAN device is attached to all nodes in the cluster.

2 Log in as the `root` user to the master node in the cluster, then open a terminal console.

3 View a list of active devices by entering

```
nlvm list devices
```

4 Identify the device that you want to initialize. If a device contains data, be prepared to lose all data on the device when you initialize it.

   In the following example, the `sdd` and `sde` devices are uninitialized and display a format of `None`. The `clus1.sbd` device is a mirrored RAID device that is used for the SBD partition of a cluster named `clus1`. Do not select an SBD device.

```
avalon:~/Desktop # nlvm list devices
 Name            Size      Used      Free  Format Shared RAID Enabled
 sda          11.00GB    9.29GB    1.70GB  MSDOS   No    No
 sdb         102.00MB  102.00MB       0KB  MSDOS  Yes    No
 sdc         102.00MB  102.00MB       0KB  MSDOS  Yes    No
 sdd         512.00MB       0KB       0KB   None   No    No
 sde         512.00MB       0KB       0KB   None   No    No
 clus134.sbd  99.57MB   99.57MB       0KB   None  Yes     1    Yes
```

5 Initialize the device by entering

```
nlvm [--force] [--no-prompt] init <device_name> [format=<gpt|msdos>]
shared
```

   For command usage information, see “Init Device” in the *OES 2023: NLVM Reference*.

   You are automatically prompted to confirm the initialize action. Respond Y (Yes) or N (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

   Replace *device_name* with the node name of the device to be initialized, such as `sde`. The device name must be the first option after `init`.

   Specify `gpt` or `msdos` as the partitioning scheme to use when formatting the device.

   The `shared` option marks the device as shareable for clustering after the initialization is complete. A small partition is created on the device to store the shared setting.

   For devices that contain data, you can specify the `--force` option to force the initialization if the `init` command cannot delete any pools on the disk.

   For example, to initialize a device with the MSDOS partitioning scheme and mark it as shareable for clustering, enter

```
nlvm init sde format=msdos shared
```

6 List details about the device to verify that the device is formatted, and the amount of free space has increased.

```
nlvm list device <device_name>
```

For example, enter

```
nlvm list device sde
```

```
avalon:~/Desktop # nlvm list device sde
 Name= sde
  Size=512.00MB(1048576) Used=32KB(64) Free=511.95MB(1048479)
  Format=GPT Shared=No RAID=No M:M=8:80 H:S=255:32
```

12.4 Creating Cluster-Enabled Pools and Volumes

After you have installed and configured OES Cluster Services, you can create a pool cluster resource by creating a cluster-enabled NSS pool and volume.

- Section 12.4.1, “Creating a Cluster-Enabled Pool and Volume with iManager,” on page 210
- Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,” on page 213

12.4.1 Creating a Cluster-Enabled Pool and Volume with iManager

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Log in to iManager as an administrator user.

3 If you have not already done so, initialize and share the device you want to use for the clustered pool.

See Section 12.3, “Initializing and Sharing a SAN Device,” on page 207. We recommend one device per pool. If you use multiple devices, each device must be marked as **Shareable for Clustering**.

4 Create a clustered pool:

4a In **Roles and Tasks**, select **Storage > Pools**.

4b Browse to select the Cluster object ( ) of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster. The shared device will be assigned to this server and the pool and volume will be created there. After you create the clustered pool and volume, you can modify the resource’s preferred node order, then cluster migrate the volume to its most preferred node.

4c Click the **New** link to open the New Pool wizard.

4d Specify the new pool name, pool type, select the Upgrade Media to Support AD Users if you want to provision NSS resources to the Active Directory users, then click **Next**.

NOTE: In OES 2015 or later, if you want an NSS32-bit pool to support AD users, select Upgrade Media to Support AD Users. All NSS64-bit pools are by default AD media upgraded.

For more information, see Creating a Pool in the OES 2023: NSS File System Administration Guide for Linux.

4e Select the check box next to the shared device where you want to create the pool, then specify the size of the pool. We recommend one that you use one device per pool.

4f Select or deselect the **Mount On Creation** option.

The **Mount On Creation** option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is enabled by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you deselect the option, you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

4g Select the **Cluster Enable on Creation** check box.

This option is selected by default if the device is shared.

If you deselect this option, you are not shown the Cluster Pool Information page. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 217.

4h Select the **Force Create** check box, when you are creating a shared cluster pool that is media upgraded to support AD users, all cluster nodes from where the pool may be accessed must be upgraded to OES 2015 or later. If you have a mixed cluster node environment, where all your cluster nodes are not on OES 2015 or later, select the Force Create check box to force the pool creation. This shared cluster pool cannot be loaded on nodes earlier than OES 2015.

4i On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	(Optional) The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is `cluster1` and the pool cluster resource name is `POOL1_SERVER`, then the default virtual server name is `CLUSTER1-POOL1-SERVER`. You can use the suggested name, or specify a different Virtual Server name for the cluster resource. You can modify the cluster resource name after the resource has been created by using the `cluster rename` command. See Section 10.15, “Renaming a Cluster Resource,” on page 191. Changing the resource name does not modify the pool name or the virtual server name.

Parameter	Action
CIFS Server Name	(Optional) If OES CIFS is installed and running, you can use this field to specify the name of the CIFS virtual server that CIFS clients see when they browse the network. If OES CIFS is installed and running, but CIFS is disabled as an advertising protocol, this field is not available (dimmed). If OES CIFS is not installed and running, this field value is NOT_SUPPORTED. CIFS is disabled by default as an advertising protocol. You can select the **CIFS** check box to enable it. By default, the NCP virtual server name is suggested as the CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, "OES CIFS for Linux," on page 205. You can use the suggested name or specify a custom name for the CIFS virtual server name. If desired, specify a new name for the CIFS virtual server. The name can be up to 15 characters, which is a restriction of the CIFS protocol.
IP Address	Specify an IP address for the pool cluster resource. Tab between the address fields. The address is IPv4 format, such as 10.10.10.243. Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster hosts the pool.
Advertising Protocols	Select the check boxes of the advertising protocols (CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the OES Trustee model. Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol. If the OES CIFS protocols are not installed and running, selecting the corresponding CIFS check box has no effect.
Online Resource after Create	The check box is deselected by default and dimmed so that you cannot change the setting. The pool is currently active on the server. You must deactivate the pool from the server before attempting to bring the resource online. You should also configure the resource load, unload, and monitor scripts before you bring the resource online.
Define Additional Properties	Select the **Define Additional Properties** check box. This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.

4j Click **Finish**.

Typically, the pool creation takes less than a minute. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

The pool cluster resource should be online and active on the master node if the **Mount on Creation** option was enabled in Step 4f.

5 Create a volume on the clustered pool.

Repeat the following procedure for each cluster volume that you want to create on the shared pool. We recommend using only one volume per shared pool.

5a In iManager, select **Storage**, then select the **Volumes**.

5b Browse to select the Cluster object ( ) of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.

5c Click **New**.

5d Specify the new volume name, then click **Next**.

Each shared volume in the cluster must have a unique name across all nodes.

5e Select the check box next to the cluster pool where you want to create the volume, select **Allow the volume to grow to the size of the pool**, then click **Next**.

5f Review and change volume attributes by selecting or deselecting the check boxes next to the attributes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see "Volume Attributes" in the *OES 2023: NSS File System Administration Guide for Linux*.

5g Choose whether you want the volume activated and mounted when it is created, then click **Finish**.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

6 Verify that the pool cluster resource was created and is online.

6a In iManager, select **Clusters > My Clusters**.

6b Select the cluster.

If the cluster is not in the list, click **Add**, browse to select its Cluster object, then click **OK**.

6c On the Cluster Manager page in the Cluster Objects list, view the pool cluster resource, such as `POOL1_SERVER`.

If the resource is online, the state is **Running**.

7 Continue with Section 12.6, "Configuring a Load Script for the Shared NSS Pool," on page 222.

12.4.2 Creating a Cluster-Enabled Pool and Volume with NSSMU

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 On the master node, log in as the `root` user, then open a terminal console.

3 If you have not already done so, initialize and share the device you want to use for the clustered pool.

See Section 12.3, “Initializing and Sharing a SAN Device,” on page 207. We recommend one device per pool. If you use multiple devices, each device must be marked as **Shareable for Clustering**.

4 Start NSSMU by entering `nssmu` at the command prompt.

5 Create a clustered pool:

5a From the NSSMU **Main Menu**, select **Pools**.

5b On the Pools page, press Insert, select a pool type, type a name for the new pool you want to create, then press Enter.

5c From the list of available devices, select the shared device where you want the pool created (such as `sdc`), then press Enter.

Select the device you marked as shared for this purpose in Section 12.3, “Initializing and Sharing a SAN Device,” on page 207.

5d Specify the amount of space (in MB) to use, then press Enter.

5e Press F3 to accept the device and size settings.

The Cluster Pool Information page opens automatically because the selected device is shared.

5f For the **Activate On Creation** option, specify whether you want the pool to be activated when it is created, then continue to the next field by pressing Enter.

The **Activate On Creation** option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is set to **Yes** (enabled) by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you set the value to **No** (disabled), you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

5g Specify **Yes** for the **Cluster Enable on Creation** option.

This option is selected by default if the device is shared.

If you set the value to **No**, the values you set on the Cluster Pool Information page are ignored, and a pool cluster resource is not created. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 217.

5h On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	(Optional) You can use the suggested name or specify a custom name for the NCP Virtual Server name. The default NCP virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is `cluster1` and the pool cluster resource name is POOL1_SERVER, then the default virtual server name is CLUSTER1-POOL1-SERVER. To modify the resource's virtual server name, select the default name, press Enter, specify a new NCP Virtual Server name for the cluster resource, then press Enter. You can modify the cluster resource name after the resource has been created by using the `cluster rename` command. See Section 10.15, "Renaming a Cluster Resource," on page 191. Changing the resource name does not modify the pool name or the virtual server name.
CIFS Server Name	(Optional) If OES CIFS is installed and running, you can use this field to specify the name of the CIFS virtual server that CIFS clients see when they browse the network. This field is blank by default if OES CIFS is installed and running, but CIFS is disabled (set to **No**) as an advertising protocol. If OES CIFS is not installed and running, this field value is NOT_SUPPORTED. CIFS is disabled by default as an advertising protocol. To enable CIFS as an advertising protocol, press the down-arrow to go to the CIFS field, then press `y` (**Yes**) to enable CIFS. By default, the NCP virtual server name is suggested as the basis for CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, "OES CIFS for Linux," on page 205. You can use the suggested name or specify a custom name for the CIFS virtual server name. To modify the name, press the up-arrow to go the **CIFS Server Name** field, press Enter to select the suggested name, specify a new name for the CIFS virtual server, then press Enter. The name can be up to 15 characters, which is a restriction of the CIFS protocol.
IP Address	Press enter to select the field, specify an IP address for the pool cluster resource, then press Enter. Specify the address in IPv4 format, such as 10.10.10.243 Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

Parameter	Action
Advertising Protocols	Press y (**Yes**) to enable an advertising protocol (CIFS, NCP) that you want to enable for data requests to this shared pool. Do not disable NCP; NCP is required to support authenticated access to data via the OES Trustee model. Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol. If the OES CIFS protocols are not installed and running, enabling the corresponding CIFS option has no effect.

5i Select **Apply** to create and cluster-enable the pool.

Typically, the pool creation takes less than a minute. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

6 Press Esc to return to the NSSMU **Main Menu**.

7 Create a volume on the clustered pool.

Repeat the following procedure for each cluster volume that you want to create on the shared NSS pool. We recommend using only one volume per shared pool.

7a On the NSSMU **Main Menu**, select **Volumes**.

7b On the Volumes page, press Insert, type a name for the new volume you want to create, then press Enter.

Each shared volume in the cluster must have a unique name across all nodes.

7c Specify Y(es) to encrypt volume or N(o) to create a regular volume.

7d From the list of available pools, select the clustered pool where you want the volume to reside, then press Enter.

The new volume appears in the list of volumes.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

7e (Optional) From the list of volumes, select the newly created volume, press F8 to view more options, press Enter to open the **Properties** page to review and change volume attributes, then select **Apply** and press Enter to save any changes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see "Volume Attributes" in the *OES 2023: NSS File System Administration Guide for Linux*.

8 Exit NSSMU.

9 Verify that the pool cluster resource was created and is online.

9a In iManager, select **Clusters > My Clusters**.

9b Select the cluster.

9c On the Cluster Manager page in the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.

If the resource is online, the state is **Running**.

10 Continue with Section 12.6, “Configuring a Load Script for the Shared NSS Pool,” on page 222.

12.5 Cluster-Enabling an Existing NSS Pool and Its Volumes

Cluster-enabling an existing pool allows it (and its volumes) to be moved or mounted on different servers in the cluster in a manner that supports transparent client reconnect. This makes the data highly available to users.

You might need to cluster enable an existing pool in at least the following situations:

- You are converting stand-alone servers with SAN attached devices to a cluster.
- You are migrating NSS pools between two separate clusters.

  In this scenario, you must create a Pool object and Volume objects in the new cluster context by using the Update eDirectory option in NSSMU or the Storage plug-in for iManager.
- You are creating a cluster resource from a storage-based snapshot.

  In this scenario, you must create a Pool object and Volume objects for the activated pool snapshot and its volumes by using the Update eDirectory option in NSSMU or the Storage plug-in for iManager.

When you cluster-enable the exiting pool, its pool and volume information will be added automatically to the load, unload, and monitor scripts for the pool cluster resource. In addition, the volume entry is automatically removed from the `/etc/fstab` so that the pool cluster resource load script controls when the volume is mounted.

Before you cluster-enable the pool, we recommend that you deactivate it. The **Online Resource after Creation** option and pool state on the node when you begin determines the resource state and pool state after the resource is created, as described in the following table:

Pool State on the Node	Online Resource After Creation Setting	Resource State after Creation	Pool State after the Resource Is Created
Deactive	Selected	Online	Active as a resource on a preferred node.
Deactive	Deselected	Offline	Deactive. Modify the resource settings and scripts as needed, then bring the resource online on a preferred node.
Active	Deselected	Offline	Active locally on the same node, even if the node is a preferred node. Modify the resource settings and scripts as needed, deactivate the pool locally, then bring the resource online on a preferred node.

Pool State on the Node	Online Resource After Creation Setting	Resource State after Creation	Pool State after the Resource Is Created
Active	Selected	Comatose	Active locally on the same node, even if the node is a preferred node. Take the comatose resource offline. Modify the resource settings and scripts as needed, deactivate the pool locally, then bring the resource online on a preferred node.

The default name of a pool cluster resource is the pool name plus the word "SERVER", such as POOL1_SERVER. You can modify the resource name after the resource has been created by using the `cluster rename` command. See Section 10.15, "Renaming a Cluster Resource," on page 191. Changing the resource name does not modify the pool name or the virtual server name.

After the pool cluster resource is brought online, users can access the cluster-enabled volume by using the IP address or virtual server name of its pool cluster resource. The default virtual server name of the resource is the cluster name plus the default cluster resource name, such as CLUSTER1-POOL1-SERVER. You can specify the name of the resource's virtual server name when you cluster-enable the pool.

The procedure in this section describes how to enable clustering for an existing pool and its volumes. It assumes the following:

- The pool contains at least one volume.
- The pool name meets the naming conventions supported by OES Cluster Services. See "Naming Conventions" on page 203.

  If the pool name contains special characters, rename the pool by using NSS management tools. See "Renaming a Pool" in the *OES 2023: NSS File System Administration Guide for Linux*.
- The pool and its volumes have Storage objects in the eDirectory tree where you are setting up the cluster.

  If the Pool and Volume objects are missing, you can create (or re-create) them in eDirectory by using the **Update eDirectory** option in the Storage plug-in for NSS or in NSSMU. You can do this from any node in the cluster. See "Updating eDirectory Pool Objects" and "Updating eDirectory Volume Objects" in the *OES 2023: NSS File System Administration Guide for Linux*.
- The pool resides on a disk that has been enabled as **Sharable for Clustering**. If it is not shared, mark the device as shareable. You can use NSSMU, the `nlvm share` command, or the **Storage > Devices** task in OES iManager to change the device's share state.
- Each pool resource requires its own unique static IP address that is in the same IP subnet as the cluster IP address.
- To add CIFS as an advertising protocol when you cluster-enable an existing pool, you must install and configure OES CIFS on the server, and ensure that the CIFS daemon is running and functioning properly before you begin. The CIFS daemon should also be running before trying to online the resource.
- If you use special settings for mounting the volume, you must manually add those settings to the load script before you bring the resource online.

To cluster-enable an existing NSS pool:

1. Ensure that the SAN device is attached to all of the nodes in the cluster.
2. Log in to iManager as a cluster administrator user.
3. Deactivate the shared pool on the current server.
   - **3a** In **Roles and Tasks**, select **Storage > Pools**.
   - **3b** Browse to select the cluster where the pool is currently assigned.
   - **3c** In the Pools list, select the pool, then click **Deactivate**.

     Wait until the page refreshes and confirms that the pool is deactive.
4. Verify that the device is marked as shareable for clustering. If the pool uses multiple devices, you must verify the shared setting of each device.
   - **4a** In **Roles and Tasks**, select **Storage > Devices**.
   - **4b** Browse to select the cluster where the pool is currently assigned.
   - **4c** Select the device to view its Details page.
   - **4d** If the **Shareable for Clustering** check box is deselected, select the check box, click **Apply**, read the warning, then click **Continue** to enable sharing.

     Marking the device as shareable automatically modifies the Shared attribute on the eDirectory objects for the pools and volumes on the device.

     Wait a couple of minutes before continuing to allow the attribute change to synchronize in eDirectory. You can verify the change in the shared state of the pool and its volumes by returning to the **Pools** page and **Volumes** page.
   - **4e** Repeat this process for each device that is used by the pool.
5. In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click **Cluster Options**.
6. On the Cluster Options page, click the **New** link in the Cluster Objects toolbar.
7. On the Resource Type page, specify Pool ( ) as the resource type you want to create by clicking the **Pool** radio button.
8. At the bottom of the Resource Type page, click **Next**.
9. On the Cluster Pool Information page, specify the following information:

Parameter	Action
Pool Name	Specify the **Pool Name** by browsing to select the Pool object ( ) of the pool you want to cluster-enable, such as *servername*_POOLC1_POOL. After the pool is cluster enabled, the Pool object name is changed to *clustername*_POOLC1_POOL, such as clus1_POOLC1_POOL.

Parameter	Action
Virtual Server Name	(Optional) Change the default name of the Virtual Server for the cluster resource. The default virtual server name for the resource is the cluster name plus the cluster resource name. The default name of a pool cluster resource is the pool name plus the word "SERVER", such as POOL1_SERVER. For example, if the cluster name is clus1 and the pool cluster resource name is POOL1_SERVER, then the default virtual server name is CLUS1-POOL1-SERVER.
CIFS Server Name	If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters, which is a restriction of the CIFS protocol. By default, the NCP virtual server name is suggested as the CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, "OES CIFS for Linux," on page 205. You can use the suggested name or specify a custom name for the CIFS virtual server name. If OES CIFS is not installed and running, this field value is NOT_SUPPORTED.
IP Address	Specify an IP address for the pool cluster resource. Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.
Advertising Protocols	Select the check boxes of the advertising protocols (CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the OES Trustee model. Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol. If the OES CIFS protocol is not installed and running, selecting the corresponding CIFS check box has no effect.
Online Resource after Creation	The check box is deselected by default. Select this option only if the pool is known to be currently in a deactive state; otherwise the resource goes comatose on creation because it will not automatically deactivate the pool from the local node.
Define Additional Properties	Select the **Define Additional Properties** check box. This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.

10 At the bottom of the Cluster Pool Information page, click **Next**.

11 On the Resource Policies page, configure the resource policies for the start, failover, and failback modes.

See "Configuring the Start, Failover, and Failback Modes for Cluster Resources" on page 182.

12 At the bottom of the Resource Polices page, click **Next**.

13 On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource.

By default, a new resource is configured to run on all nodes, including future nodes. Add the nodes you do not want to use to the **Unassigned** list.

IMPORTANT: For **Assigned** nodes, ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

See "Configuring Preferred Nodes and Node Failover Order for a Resource" on page 184.

14 At the bottom of the Resource Preferred Nodes page, click **Finish**.

The pool cluster resource appears in the **Cluster Objects** list on the Cluster Options page, such as POOL1_SERVER.

The load, unload, and monitor scripts are automatically created with information about the pools and volumes. The volume entries are removed from the `/etc/fstab` file so that the resource load script can be used to control when the pools and volumes are mounted. The resource is offline.

15 Verify that the pool cluster resource was created and that it is offline.

15a In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click **Cluster Options**.

15b View the pool cluster resource in the Cluster Objects list, such as POOLC1_SERVER.

15c Select **Cluster Manager**, then view the pool cluster resource state as **Offline**.

16 Verify the cluster resource scripts for the pool cluster resource, and modify them if desired.

16a In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click **Cluster Options**.

16b Click the pool cluster resource link to open its Properties page.

16c Configure the following resource settings:

Property	Action
Monitoring	(Optional) Select the **Monitoring** tab, select **Enable Resource Monitoring**, and configure monitoring as described in Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 176.
Load script	Select the Scripts tab, select Load Script, then modify the script as described in Section 12.6, "Configuring a Load Script for the Shared NSS Pool," on page 222.
Unload Script	Select the Scripts tab, select Unload Script, then modify the unload script as described in Section 12.7, "Configuring an Unload Script for the Shared NSS Pool," on page 224.
Monitor Script	Select the Scripts tab, select Monitor Script, then modify the unload script as described in Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 176.

16d At the bottom of the page, click **Apply** or **Ok** to save your changes.

The changes are not applied until the next time the resource is brought online.

17 If you did not deactivate the pool before you began, manually deactivate the pool on its currently assigned node.

You must deactivate the pool from the server before attempting to bring the pool cluster resource online.

17a In **Roles and Tasks**, select **Storage > Pools**.

17b In the Server field, browse to select the Cluster object ( ) of the cluster.

The Cluster object is automatically associated with the node that is currently the master node in the cluster.

17c In the Pools list, select the pool, then click **Deactivate**.

Wait until the page refreshes and confirms that the pool is deactive.

18 Bring the resource online:

18a In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click **Cluster Manager**.

18b Select the check box next to the pool cluster resource, then click **Online** to bring the resource online.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

If the pool is comatose, take the resource offline, then bring it online. If it goes comatose again, check the changes you made to the cluster scripts, then try again.

12.6 Configuring a Load Script for the Shared NSS Pool

A cluster resource load script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access scripts for the cluster resource, see Section 10.5, “Configuring a Load Script for a Cluster Resource,” on page 174.

Newly created pool cluster resources use the following sequence for commands in a load script. The order is reversed in an unload script.

NSS (pool activation and volume mount)
[NFS] (If NFS is used, you manually add the command here.)
Secondary IP address
NCP
[CIFS]

If you modify the load script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the IP address, pool name, and volume name.

If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in Section 8.13, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 120.

- Section 12.6.1, “Sample NSS Load Script,” on page 223
- Section 12.6.2, “Specifying a Custom Mount Point for an NSS Volume,” on page 223
- Section 12.6.3, “Specifying Custom Mount Options,” on page 224
- Section 12.6.4, “Loading OES CIFS as an Advertising Protocol,” on page 224

12.6.1 Sample NSS Load Script

The following sample values are used in the NSS load script below:

Variable	Sample Value
Cluster resource’s virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

exit_on_error nss /poolact=SHPOOL43
exit_on_error ncpcon mount SHVOL43=252

exit_on_error add_secondary_ipaddress 10.10.10.43

exit_on_error ncpcon bind --ncpservername=NCS1-SHPOOL43-SERVER
   --ipaddress=10.10.10.43

exit 0
```

12.6.2 Specifying a Custom Mount Point for an NSS Volume

For NSS pool cluster resources, the default mount point of `/media/nss/volume_name` is assumed. In order to use a custom mount point, you must add some Linux commands to mount and dismount the volume from Linux at that location. For information about using a custom mount point path, see Section 12.14, “Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume),” on page 250.

12.6.3 Specifying Custom Mount Options

The load script for an NSS volume assumes that the default name space of LONG is used when mounting the clustered NSS volume. Supported name spaces include LONG, UNIX, DOS, and MAC. For information about name spaces for NSS volumes, see "Lookup Namespace" in the *OES 2023: NSS File System Administration Guide for Linux*.

You can specify the desired name space by adding the following switch to the `ncpcon mount` command in the load script:

```
/opt=ns=<long|unix|dos|mac>
```

The options are case-sensitive. In the `ncpcon mount` command, place the `/opt` switch before the volume information.

If you use a name space other than the default for the shared volume, or if you change the name space setting for shared volume by using NSSMU or the Storage plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume.

For example, to specify the LONG name space, add the `/opt=ns=long` switch as follows:

```
exit_on_error ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID>
```

For example, to specify the Unix name space, add the `/opt=ns=unix` switch as follows:

```
exit_on_error ncpcon mount /opt=ns=unix <VOLUMENAME>=<VOLUMEID>
```

After you modify the load script, you must take the pool cluster resource offline, then bring it online to apply the new name space setting for the volume.

12.6.4 Loading OES CIFS as an Advertising Protocol

If you enable OES CIFS as an advertising protocol for the pool, the following command is automatically added below the `bind` command:

```
exit_on_error novcifs --add
   '--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
   --ip-addr=10.10.10.43
```

The `--vserver` option uses the fully distinguished name of the NCP virtual server. The `--ip-addr` option uses the IP address of the pool cluster resource.

12.7 Configuring an Unload Script for the Shared NSS Pool

A cluster resource unload script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access the scripts for a cluster resource, see Section 10.6, "Configuring an Unload Script for a Cluster Resource," on page 175.

Newly created pool cluster resources use the following sequence for commands in an unload script. The order is reversed in a load script.

[CIFS]

NCP
Secondary IP address
[NFS] (If NFS is used, you manually add the command here.)
NSS (pool deactivation, which automatically dismounts the volumes)

If you modify the unload script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

- Section 12.7.1, “Sample NSS Unload Script,” on page 225
- Section 12.7.2, “Using Volume Dismount Commands,” on page 225
- Section 12.7.3, “Unloading OES CIFS as an Advertising Protocol,” on page 226

12.7.1 Sample NSS Unload Script

The following sample values are used in the NSS unload script below:

Variable	Sample Value
Cluster resource’s virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

ignore_error ncpcon unbind --ncpservername=NCS1-SHPOOL43-SERVER --
ipaddress=10.10.10.43

ignore_error del_secondary_ipaddress 10.10.10.43

ignore_error nss /pooldeact=SHPOOL43

exit 0
```

12.7.2 Using Volume Dismount Commands

By default, the pool deactivation automatically dismounts the volumes on the pool. The `ncpcon dismount` commands are not automatically added to the unload script. However, you can add `ncpcon dismount` commands for each volume before the pool deactivation line. For example:

```
ignore_error ncpcon dismount SHVOL43
ignore_error nss /pooldeact=SHPOOL43
```

Adding the `ncpcon dismount` command provides predictable and clean dismount behavior. It allows the volume dismount to occur and be logged independently of the pool deactivation, which can help with troubleshooting dismount issues. However, the ncpcon dismount commands are not automatically maintained and updated if you change a volume name or delete a volume. You must modify the unload script yourself.

12.7.3 Unloading OES CIFS as an Advertising Protocol

If OES CIFS is enabled as an advertising protocol for the pool, the following line is automatically added above the unbind command:

```
ignore_error novcifs
   --remove '--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
   --ip-addr=10.10.10.43
```

12.8 Configuring a Monitor Script for the Shared NSS Pool

A cluster resource monitor script is automatically generated for the pool when you cluster-enable it. It is disabled by default. To enable or disable monitoring, see Section 10.7.2, “Configuring Resource Monitoring,” on page 178.

If you modify the monitor script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

- Section 12.8.1, “Sample NSS Monitor Script,” on page 226
- Section 12.8.2, “Monitoring OES CIFS as an Advertising Protocol,” on page 227

12.8.1 Sample NSS Monitor Script

The following sample values are used in the NSS monitor script below:

Variable	Sample Value
Cluster resource’s virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

The following is a sample monitor script for an NSS pool cluster resource:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1
nsspool

exit_on_error ncpcon volume VOL1

exit_on_error status_secondary_ipaddress 10.10.10.41

exit 0
```

You can add lines to the monitor script that check the status of processes used by clustering. See Section 10.7.3, “Monitoring Services that Are Critical to a Resource,” on page 179.

12.8.2 Monitoring OES CIFS as an Advertising Protocol

If OES CIFS is enabled as an advertising protocol for the pool, the OES CIFS monitoring command is automatically added to help keep CIFS up and running.

```
exit_on_error rcnovell-cifs monitor
```

The command checks to see if the `cifsd` daemon is running and takes action if it is not running. If CIFS is running, it returns its status. If CIFS is dead or not running, it starts the `cifsd` daemon and returns its status. If the restart fails, the resource’s failover settings are invoked, such as failing over the resource to its preferred node, or going comatose.

Pool cluster resources that were created on OES 2 SP3 and earlier clusters might use the CIFS `status` command in the monitor script. It checks the status and immediately invokes the resource’s failover settings if the `cifsd` daemon is not running. To take advantage of the restart capability that is offered by the `monitor` command, you can modify the line to use “`monitor`” instead of “`status`”.

12.9 Adding Advertising Protocols for NSS Pool Cluster Resources

Shared NSS pools support three advertising protocols: NCP (default, required), and OES CIFS. You can add OES CIFS as advertising protocol for an existing pool cluster resource.

You must install OES CIFS on all the servers where you want to fail over the pool resource, but you set up the advertising protocol only once on the pool cluster resource. Before you begin, ensure that the pool is cluster migrated back to the server where it was originally created.

1. Log in to iManager as a cluster administrator user.
2. Offline the cluster resource that you want to modify.

   See Section 9.11, “Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node,” on page 140.
3. In **Roles and Tasks**, select **Storage**, then go the Volumes and Pools pages to verify that the volumes are not mounted and the pool is deactive.
4. In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then select **Cluster Options**.

5 On the Cluster Options page, use one of the following methods to open the Cluster Pool Properties page for the pool cluster resource that you want to manage:

- Select the check box next to the pool cluster resource, then click **Details**.
- Click the **Name** link for the pool cluster resource.

6 On the Cluster Pool Properties page, click the **Protocols** tab.

7 If you are enabling the resource for OES CIFS, select the **CIFS** check box and view the **CIFS Virtual Server Name**, then click **Apply**.

A default name (up to 15 characters) is suggested for the CIFS Virtual Server Name when you first add the advertising protocol.

If you later want to modify the CIFS Virtual Server Name, you must use the CIFS plug-in for iManager to manage the CIFS share name of an existing share. See "Setting CIFS General Server Parameters" in the *OES 2023: OES CIFS for Linux Administration Guide*.

8 Online the cluster resource. Click **Cluster Manager,** select the check box next to the resource, then click **Online**.

12.10 Adding NFS Export for a Clustered Pool Resource

You can add NFS support for a clustered NSS volume by adding an `exportfs(8)` command to the load script and unload script for a clustered pool resource. Basically, the order in load script should be NSS, NFS, IP, and NCP. The order is reversed in the unload script.

Before you use the `exportfs(8)` command in resource scripts, ensure that you install the `nfs-kernel-server` package on all of the cluster nodes in the resource's preferred nodes list. This contains the support utilities for the kernel `nfsd` daemon. It is not installed by default. You can use the **YaST > Software > Software Management** tool to find and install the package.

- Section 12.10.1, "Sample NFS Load Script," on page 228
- Section 12.10.2, "Sample NFS Unload Script," on page 229

12.10.1 Sample NFS Load Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# Activate the NSS pool and mount its volume
exit_on_error nss /poolact=POOL_16
exit_on_error ncpcon mount VOL_161=224

# Export the volume for NFS
exit_on_error exportfs -o rw,sync,no_root_squash,fsid=216 *:/media/nss/VOL_161

exit_on_error add_secondary_ipaddress 192.168.188.16

exit_on_error ncpcon bind --ncpservername=CLUS1_POOL_16_SERVER --ipaddress=192.168.188.16

exit 0
```

12.10.2 Sample NFS Unload Script

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

ignore_error ncpcon unbind --ncpservername=CLUS1_POOL_16_SERVER --
ipaddress=192.168.188.16

ignore_error del_secondary_ipaddress 192.168.188.16

#Unexport a volume for NFS
ignore_error exportfs -u *:/media/nss/VOL_161

# Deactivate the pool and volume
ignore_error nss /pooldeact=POOL_16

exit 0
```

12.11 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes

- Section 12.11.1, “Understanding NSS Mirroring,” on page 229
- Section 12.11.2, “Requirements for NSS Mirroring,” on page 231
- Section 12.11.3, “Creating and Mirroring NSS Pools on Shared Storage,” on page 231
- Section 12.11.4, “Verifying the NSS Mirror Status in the Cluster,” on page 233

12.11.1 Understanding NSS Mirroring

NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. It is an alternative to synchronous replication options that are provided by a SAN storage array.

Hardware configuration and placement for NSS mirroring can vary depending on geography and fault tolerance requirements. For example, Figure 12-1 depicts a simple hardware configuration in which one side of the mirrored NSS volume is located in a separate building from the rest of the cluster hardware. If a disaster occurs in one building, data is still safe on the mirrored NSS volume in the other building.

Figure 12-1 *Single Cluster with Mirrored NSS Volumes in Separate Buildings*

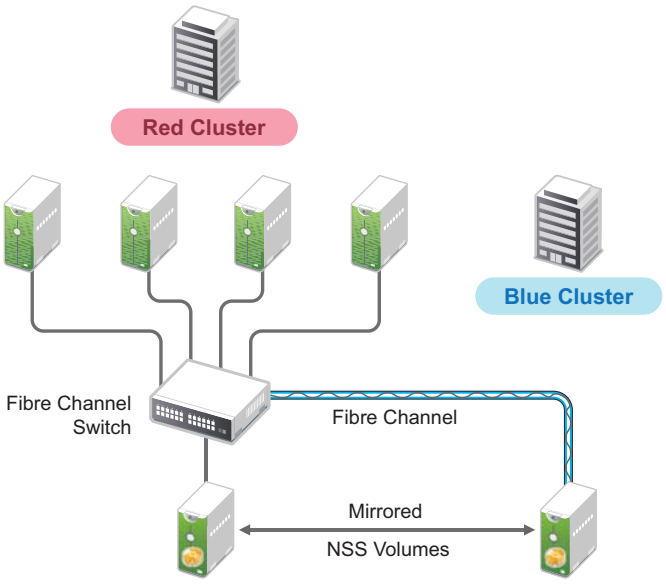

Figure 12-2 depicts a more complex hardware configuration in which two clusters are placed in separate buildings, but each has one side of a mirrored NSS volume in the building with the other cluster. If a disaster occur in one building, data is still safe and immediately accessible on the mirrored NSS volume in the other building.

Figure 12-2 *Two Clusters with Mirrored NSS Volumes in Separate Buildings*

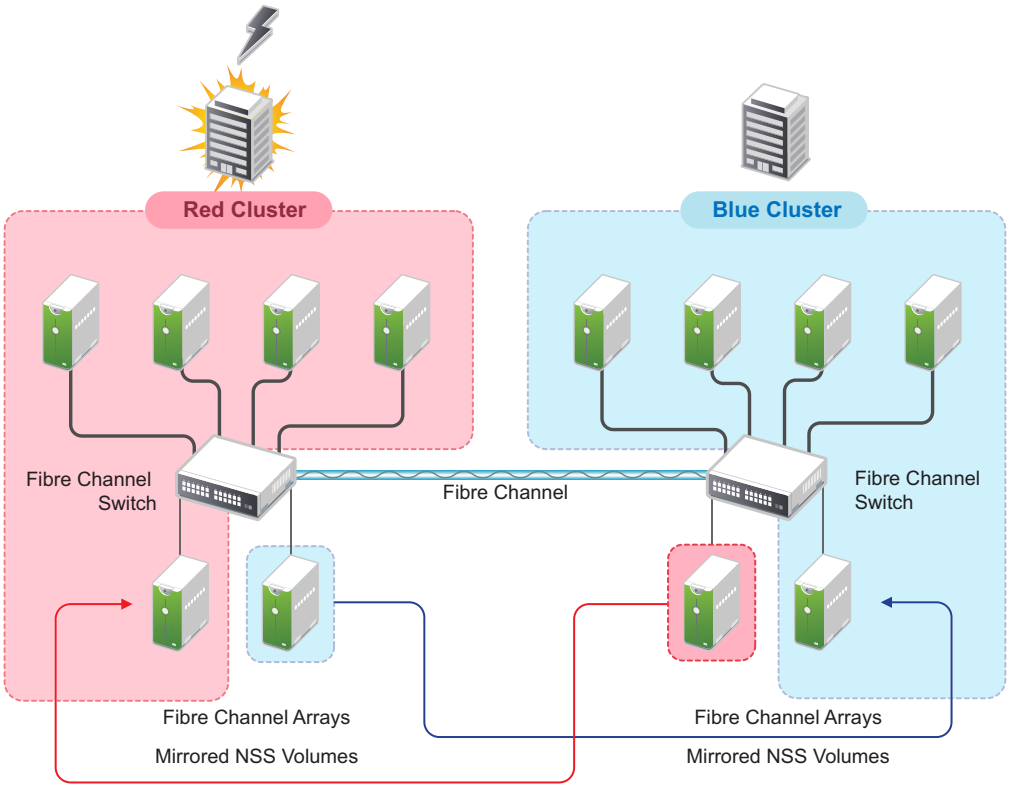

The maximum distance for a single stretch of fibre without optical repeaters to boost distance is about 10 kilometers. There are various devices available that convert Fibre Channel to IP in order to extend SANs to WAN scale.

NSS mirroring is synchronous, and so mirror I/Os must complete before the next I/O can be sent. Performance is a function of time to do an I/O over the longest link.

12.11.2 Requirements for NSS Mirroring

NSS partitions must be mirrored after they are created. If you have an existing partition that you want to mirror, you can either create another partition of equal size on another device to mirror the first partition to, or let the mirroring software automatically create another partition of equal size on another device.

OES Cluster Services should be installed and running prior to creating and mirroring partitions on shared storage.

When you create a OES Cluster Services system that utilizes shared storage space (a Storage Area Network or SAN), it is important to remember that all servers attached to the shared device, whether in the cluster or not, have access to all of the volumes on the shared storage space unless you specifically prevent such access. OES Cluster Services arbitrates access to shared volumes for all cluster nodes, but cannot protect shared volumes from being corrupted by non-cluster servers.

12.11.3 Creating and Mirroring NSS Pools on Shared Storage

Prior to creating and mirroring NSS partitions on shared storage, ensure that you have the following:

- OES 2023 is installed on all servers that will be part of the cluster.
- All servers in the cluster are connected to a shared storage system.
- One or more drive arrays are configured on the shared storage system.
- For 512 byte device, at least 20 MB of free space and for 4096 (4Kn) byte device at least 80 MB of free space on the shared storage system for a special cluster partition.
- To ensure disaster recovery, the device you select to mirror should be in a different storage array.

Use one of the following methods to mirror clustered pool resources:

- "Creating a Pool on a Shared Mirrored Device" on page 231
- "Mirroring the Partition of an Existing NSS Pool" on page 232

Creating a Pool on a Shared Mirrored Device

To create an NSS pool resource on a shared mirrored device:

1. Start NSSMU by entering `nssmu` at the command prompt of a cluster server.
2. Share the devices.
   - **2a** Select **Devices** from the NSSMU **Main Menu**.
   - **2b** Select the two devices that you want to use, then press F6 to mark the devices as **Sharable for Clustering**.

3 Create a mirrored RAID device using the shared devices.

3a Select **RAID Devices** from the NSSMU **Main Menu**.

3b Press Insert to create a software RAID, select RAID 1, then press Enter.

3c Use the arrow keys to select the two shared devices that you want to contribute space to the RAID, specify the amount of space to use, then press Enter.

The software RAID device automatically inherits the Sharable for Clustering setting.

4 Create a clustered NSS pool on the shared RAID device.

4a Select **Pools** from the NSSMU **Main Menu**.

4b Press Insert to create a new pool, then create the clustered pool as usual on the shared software RAID device as described in Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,” on page 213.

Mirroring the Partition of an Existing NSS Pool

You can use NSSMU to mirror a clustered pool’s partition. The second device must have as much space available as you know the pool’s real allocated size to be, and be marked as Sharable for clustering. To mirror the partition for an existing clustered NSS pool resource:

1 Start NSSMU by entering `nssmu` at the command prompt of a cluster server.

2 Share the second device.

2a Select **Devices** from the NSSMU **Main Menu**.

2b Select the device that you want to use, then press F6 to mark it as **Sharable for Clustering**.

To ensure disaster recovery, the device you select to mirror should be in a different storage array.

3 Select **Partitions** from the NSSMU **Main Menu**.

4 Choose the partition that contains the pool and press F3, enter the RAID name and select the second disk, then press F3 to accept.

This mirrors the existing pool and creates the RAID 1 device where each segment is the size of the pool.

5 Use one of the following methods to initiate mirroring for the newly created mirror:

- At the terminal 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 terminal 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.

6 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 that is greater than 0. It is fully synchronized at 100%.

12.11.4 Verifying the NSS Mirror Status in the Cluster

After you have configured NSS mirroring with OES Cluster Services, you should check to ensure that it is working properly in a cluster environment.

1 Ensure that the volumes on the cluster-enabled pool are mounted on an assigned server by entering `nss /volumes` at the terminal console.

2 Check the mirror status of the mirrored partition by entering `nlvm raid status` at the terminal console of the server where the NSS pool on the mirrored partition is active.

After entering `nlvm raid status`, you should see a message indicating that mirror status is 100 percent or a message indicating that the mirrored object is fully synchronized.

3 Migrate the pool to another server in the cluster and again check to ensure that the volumes on the pool are mounted by entering `volumes` at the terminal console.

4 Check the mirror status of the partition again by entering `nlvm raid status` at the terminal console.

12.12 Renaming a Clustered NSS Pool

When an NSS pool is cluster-enabled with OES Cluster Services, special steps must be followed to rename the pool. The pool cluster resource must be taken offline during the renaming process. You can use NSSMU or the Storage plug-in to iManager to rename the shared pool.

When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory, updates the pool name for its Volume objects, and updates the pool name that appears in the load, unload, and monitor scripts for the pool's cluster resource. The rename function does not change the names for the following resource objects:

- Cluster resource name (such as POOL1_SERVER)

  You can use the `cluster rename` command to rename the resource. See Section 10.15, "Renaming a Cluster Resource," on page 191.

- Virtual server name (the NCS:NCP Server object, such as CLUS1-POOL1-SERVER)
- CIFS virtual server name (such as CLUS1-POOL1-W), if CIFS is enabled as an advertising protocol for the resource

To rename a pool name in a way that also modifies names of resource objects, you can disable clustering for the pool, rename the pool, and then enable clustering for the renamed pool. This method is also described in Section 12.12.3, "Renaming a Shared Pool and Its Resource Objects," on page 237.

Because renaming requires that information is modified in eDirectory, you should rename the pool from the master node. NSS management tools require that the shared pool is active locally on a cluster node that has its NCP Server object in the same context as the Pool object of the pool you

want to rename. The server must also be in the resource's **Preferred Nodes** list to migrate to it. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

You can use any of the methods in this section to rename a shared pool.

- Section 12.12.1, "Renaming a Shared Pool with iManager," on page 234
- Section 12.12.2, "Renaming a Shared Pool with NSSMU," on page 235
- Section 12.12.3, "Renaming a Shared Pool and Its Resource Objects," on page 237
- Section 12.12.4, "Renaming a Shared Pool in a DST Cluster Resource," on page 240

12.12.1 Renaming a Shared Pool with iManager

1. Ensure that the pool cluster resource is running on the master node.

   If you need to migrate the resource:

   1a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

   1b On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Migrate**.

   1c Select the master node, then click **OK**.

   Wait for the resource to report a **Running** state before you continue.

2. Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Offline**.

   Wait for the resource to report an **Offline** state before you continue.

3. Activate the pool locally on the master node:

   3a In iManager, select **Storage > Pools**.

   3b Select the node where you took the resource offline.

   3c Select the pool, then click **Activate**.

4. Rename the shared pool:

   4a In the **Pools** list, select the pool, then click **Rename**.

   4b Read the warning message that the pool's volumes will be dismounted, then click **OK** to confirm that you want to rename the pool.

   4c In the **Rename Pool** dialog box, type the new name, then click **Finish**.

   4d Read the error message, then click **OK** to dismiss it.

   Error 23391 is `EDIR OBJECT NOT FOUND`. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

   4e View the renamed pool in the **Pools** list.

   After the rename, the pool is in a deactive state.

5. Verify that the pool name was modified in the pool resource scripts:

   5a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

   5b On the Cluster Manager page, click the name link of the pool resource.

   The pool resource name has not been modified by the pool rename. Its default name is *<old_pool_name>*_SERVER.

5c On the Resource Properties page, click the **Scripts** tab.

5d Verify that the new pool name is used in the load, unload, and monitor scripts.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

5d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click **Apply**.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

5d2 Click **Unload Script**, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

5d3 Click **Monitor Script**, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/
nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/
.pools/PUSERS
```

5e Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

6 Bring the pool cluster resource online:

6a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

6b On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Online**.

The resource is brought online on a node in its **Preferred Nodes** list, according to availability. Wait for the resource to report a **Running** state before you continue.

7 (Optional) Rename the pool cluster resource. For example, rename POOL1_SERVER as PUSERS_SERVER.

For instructions, see Section 10.15, “Renaming a Cluster Resource,” on page 191.

12.12.2 Renaming a Shared Pool with NSSMU

1 Log in to the master node as the `root` user, then launch a terminal console.

2 Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

3 Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

4 Activate the pool locally on the master node:

4a At the command prompt, enter

```
nssmu
```

4b Select **Pools** to view a list of pools. The pool state is **deactive**.

4c Select the pool, then press F7 **Activate**.

5 Rename the pool:

5a On the NSSMU Pools page, select the pool.

5b Press F6 **Rename**.

5c Read the warning message that the pool's volumes will be dismounted, then click **OK** to confirm that you want to rename the pool.

5d Specify the new name, then press Enter.

5e Read the error message, then press o (OK) or press Enter to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

5f After the rename, the pool is in a **deactive** state.

6 Press Esc twice to exit NSSMU.

7 Verify that the pool name was modified in the pool resource scripts:

7a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

7b On the Cluster Manager page, click the name link of the pool resource.

The pool resource name has not been modified by the pool rename.

7c On the Resource Properties page, click the **Scripts** tab.

7d Verify that the new pool name is used in the load, unload, and monitor scripts.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

7d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click **Apply**.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

7d2 Click **Unload Script**, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

7d3 Click **Monitor Script**, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/
nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/
.pools/PUSERS
```

7e Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

8 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

9 (Optional) Rename the pool cluster resource. For example, rename POOL1_SERVER as PUSERS_SERVER.

For instructions, see Section 10.15, “Renaming a Cluster Resource,” on page 191.

12.12.3 Renaming a Shared Pool and Its Resource Objects

To modify the pool name and its related resource objects, you can disable clustering for the pool, rename the pool, and then enable clustering for the renamed pool. Ensure that the pool is on the master node when you disable clustering. The Clusters plug-in in iManager must be used to disable and enable clustering for the pool. The rename procedure can be performed in iManager or NSSMU.

1 Ensure that the pool cluster resource is running on the master node.

If you need to migrate the resource:

1a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

1b On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Migrate**.

1c Select the master node, then click **OK**.

Wait for the resource to report a **Running** state before you continue.

2 Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Offline**.

Wait for the resource to report an **Offline** state before you continue.

3 In iManager, use the Clusters plug-in to delete the cluster resource for the shared pool:

3a Select **Clusters > My Clusters**, then select the cluster.

3b Select the **Cluster Options** tab.

3c Select the check box next to the resource, then click **Delete**.

3d When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.

3e In the **Tree View** in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.

If necessary, you can manually delete the objects. In iManager, go to **Directory Administration > Delete Objects**, select the objects, then click **OK**.

4 In iManager, re-create Storage objects for the pool and its volumes:

These objects are needed by NSS when you rename the pool. They are also used by OES Cluster Services when you re-create the resource.

4a In iManager, select **Storage > Pools**, then select the cluster node where you want to rename the pool.

4b In the **Pools** list, select the pool, then click **Activate**.

4c Select the pool, then click **Update eDirectory**.

This creates a Pool object in eDirectory with a name format of *<server_name>_<pool_name>*_POOL.

4d Select **Storage > Volumes.** The server should still be selected.

4e In the **Volumes** list, select the volume, then click **Mount**.

4f Select the volume, then click **Update eDirectory**.

This creates a Volume object in eDirectory with a name format of *<server_name>_<volume_name>*.

4g Repeat Step 4e through Step 4f for each volume in the pool.

5 Rename the pool.

The following rename procedure uses iManager. You can alternatively use NSSMU to rename the pool.

5a In iManager, select **Storage > Pools**. The server should still be selected.

5b In the **Pools** list, select the pool, then click **Rename**.

5c Read the warning message that the pool's volumes will be dismounted, then click **OK** to confirm that you want to rename the pool.

5d In the **Rename Pool** dialog box, type the new name, then click **Finish**.

5e Read the error message, then click **OK** to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

5f View the renamed pool in the **Pools** list.

After the rename, the pool is in a deactive state. Do not activate it on any nodes.

If you activate it to verify that it is working as desired with the new name, deactivate it again before you continue.

6 Use the Clusters plug-in to cluster-enable the pool and create new resource objects with the new pool name.

For detailed instructions, see Step 5 thru Step 18 in Section 12.5, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 217.

6a In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster.

6b Select the **Cluster Options** tab.

6c On the Cluster Options page, click the **New** link in the Cluster Objects toolbar.

6d On the Resource Type page, select the **Pool** radio button, then click **Next**.

6e On the Cluster Pool Information page, specify the following resource setup information, then click **Next**:

- Pool name

  Browse to select the Pool Object of the pool.

- Virtual server name

  A default name is suggested based on the pool name. You can specify a different name for the virtual server.

- IP address

  This is the IP address to use for the resource's virtual server.

- Advertising protocols
  - NCP (default, mandatory)
  - CIFS

    If you enable CIFS, a default name is suggested based on the NCP virtual server name. You can specify a different CIFS Server name.
- Deselect the **Online Resource after Creation** check box.
- Select the **Define Additional Properties** check box.

6f On the Resource Policies page, configure the resource policies for the start, failover, and failback modes, then click **Next**.

See "Configuring the Start, Failover, and Failback Modes for Cluster Resources" on page 182.

6g On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource, then click **Finish**.

See "Configuring Preferred Nodes and Node Failover Order for a Resource" on page 184.

The pool cluster resource appears in the **Cluster Objects** list on the Cluster Options page, such as POOL1_SERVER.

6h (Optional) View and modify the resource properties.

6i (Optional) Enable monitoring and configure the monitoring script.

6j Bring the resource online. Select the Cluster Manager tab, select the check box next to the resource, then click **Online**.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

If the pool goes comatose, take it offline, check that the pool is deactivated on the local server, then try again.

12.12.4 Renaming a Shared Pool in a DST Cluster Resource

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory; updates the pool name for its Volume objects; and updates the pool name that appears in the load, unload, and monitor scripts for the pool's cluster resource. However, renaming a secondary pool cannot also update its name in the scripts for the primary pool's resource, because the secondary pool's default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary pool.

Renaming the shared primary pool or secondary pool in the DST cluster resource does not affect the settings that define the DST shadow relationship.

The resource objects for a DST cluster resource are typically named when you cluster-enable the primary pool. Renaming a primary pool does not modify the names of the resource objects. To rename a primary pool name in a way that also modifies names of resource objects, you can disable clustering for the primary pool, rename the pool, and then enable clustering for the renamed pool. This process is described in Renaming a Shared Pool and Its Resource Objects.

Afterwards, reconfigure the primary pool cluster resource to include the necessary commands for the secondary pool and the DST shadow volume pair, just as you did when you first set up the DST cluster resource. Because the primary volume name and secondary volume name are not affected by the pool rename, you do not need to modify the related commands in the `/etc/opt/novell/ncp2nss.conf` file and `/etc/opt/novell/ncpserv.conf` file.

To rename a shared pool in a DST cluster resource:

1. Log in to the master node as the `root` user, then launch a terminal console.
2. Ensure that the pool cluster resource is running on the master node. Enter

   ```
   cluster status
   ```

   If you need to migrate the resource, enter the following at the command prompt:

   ```
   cluster migrate <resource_name> <master_node_name>
   ```

3. Take the pool cluster resource offline. At the command prompt, enter

   ```
   cluster offline <resource_name>
   ```

   You can verify the offline state by entering

   ```
   cluster status
   ```

4. Activate the pool to be renamed locally on the master node:
   - **4a** At the command prompt, enter

     ```
     nssmu
     ```

   - **4b** Select **Pools** to view a list of pools. The pool state is **Deactive**.
   - **4c** Select the pool, then press F7 **Activate**.

5 Rename the pool:

5a On the NSSMU Pools page, select the pool.

5b Press F6 **Rename**.

5c Read the warning message that the pool's volumes will be dismounted, then click **OK** to confirm that you want to rename the pool.

5d Specify the new name, then press Enter.

5e Read the error message, then press o (OK) or press Enter to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

5f After the rename, the pool is in a deactive state.

6 Press Esc twice to exit NSSMU.

7 Verify that the pool name was modified in the pool resource scripts:

7a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

7b On the Cluster Manager page, click the name link of the pool resource.

Remember that the pool resource name has not been modified by the pool rename.

7c On the Resource Properties page, click the **Scripts** tab.

7d Verify that the new pool name is used in the load, unload, and monitor scripts.

If you modified the secondary pool name, you must manually modify that line in the script.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

7d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click **Apply**.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

7d2 Click **Unload Script**, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

7d3 Click **Monitor Script**, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/
nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/
.pools/PUSERS
```

7e Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

8 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

12.13 Renaming a Clustered NSS Volume

When an NSS volume is in a cluster-enabled pool, special steps must be followed to rename it. Renaming the volume automatically changes the Volume object name, updates the volume information for the pool, and updates the volume name in the scripts for the pool cluster resource. Custom mount options are preserved in the ncpcon mount command. After the volume is renamed, the default mount point changes from `/media/nss/<old_volume_name>` to `/media/nss/<new_volume_name>`.

Because renaming requires that information is modified in eDirectory, you should rename the volume from the master node. NSS management tools require that the shared pool and volume are active locally on a cluster node that has its NCP Server object in the same context as the Pool object and Volume object of the volume you want to rename. The server must also be in the resource's **Preferred Nodes** list to migrate to it. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

Before you rename the volume, check the volume properties to ensure that the attribute **Allow mount point to be renamed** has been enabled. If the attribute is disabled, the volume name will be modified, and the mount point path is unchanged on the initial remount, making it appear that you have retained the old mount point path. However, for the shared volume, the default mount path of `/media/nss/<volume_name>` is enforced with the new name on subsequent mounts. If your intent is to use a custom mount point path for a shared NSS volume, see Section 12.14, "Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume)," on page 250.

If the pool cluster resource manages multiple volumes, we recommend that you rename only one volume at a time. Verify that the renamed volume and pool cluster resource are working as expected before you attempt to rename another volume in the shared pool.

The NCP Server and OES CIFS services are integrated with NSS and Cluster Services. The NCP and CIFS file access continues to work after the volume is renamed. However, you must ensure that your users have the proper path information for the volume after it is renamed. If you have other protocols and services configured for the volume, you might need to re-configure them to work with the new volume name.

After you rename the NSS volume, you must manually edit some lines in the resource scripts if you use a custom mount point path for the NSS volume. For more information, see Section 12.14, “Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume),” on page 250.

If the volume is used as the primary volume or secondary volume in a Dynamic Storage Technology resource script, special steps are required to remove references to the old names as a shadow volume pair. Some information must be removed manually on each node before you bring the modified resource online.

You can use the procedures in the following sections to rename shared volumes:

- Section 12.13.1, “Renaming a Shared NSS Volume with iManager,” on page 243
- Section 12.13.2, “Renaming a Shared NSS Volume with NSSMU,” on page 245
- Section 12.13.3, “Renaming a Shared NSS Volume in a DST Cluster Resource,” on page 247

12.13.1 Renaming a Shared NSS Volume with iManager

1. Ensure that the pool cluster resource is running on the master node.

   If you need to migrate the resource:

   1a. In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

   1b. On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Migrate**.

   1c. Select the master node, then click **OK**.

   Wait for the resource to report a **Running** state before you continue.

2. Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Offline**.

   Wait for the resource to report an **Offline** state before you continue.

3. Activate the pool locally on the master node:

   3a. In iManager, select **Storage > Pools**.

   3b. Browse to select the server where you want to activate the pool.

   3c. In the **Pools** list, select the shared pool that contains the volume, then click **Activate**.

4. Mount the volume locally on the master node:

   4a. In iManager, select **Storage > Volumes**.

   The server field is automatically populated with the server you chose in the previous step.

   4b. In the **Volumes** list, select the volume, then click **Mount**.

5. View the volume properties to ensure that the attribute **Allow Mount Point to be Renamed** has been enabled. If it is disabled, select it, then click **OK**.

6. Rename the volume:

   6a. In the **Volumes** list, select the volume, then click **Rename**.

   6b. In the **Rename Volume** dialog box, specify the new name, then click **OK**.

6c Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.

If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an Error 601 eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.

7 Dismount the volume:

7a In the **Volumes** list, select the volume, then click **Dismount**.

7b Read the warning message that open files will be closed by the action, then click **OK** to confirm the dismount.

7c Wait for the volume state to change to **Deactive, Not Mounted**, then continue.

8 Deactivate the pool:

8a In iManager, select **Storage > Pools**.

8b In the **Pools** list, select the pool, then click **Deactivate**.

8c Read the warning message that the pool's volumes will be dismounted, then click **OK** to continue.

8d When the page refreshes, verify that the pool state is **Deactive**.

9 Verify that the volume name was updated in the scripts for its pool cluster resource:

9a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

9b On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.

9c Click the **Scripts** tab.

9d On the Load Script page, verify that the volume name was changed in the `ncpcon mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click **Apply**.

```
exit_on_error ncpcon mount <new_volume_name>=<volume_id>
```

For example:

```
exit_on_error ncpcon mount USERS=252
```

9e Click **Unload Script**, then check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

The unload script does not contain any default volume commands.

9f Click **Monitor Script**, then verify that the volume name was changed in the `ncpcon volume` command. Also check any custom commands where the volume name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error ncpcon volume <new_volume_name>
```

For example:

```
exit_on_error ncpcon volume USERS
```

9g Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

10 Bring the pool cluster resource online:

10a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

10b On the Cluster Manager page, select the check box next to the pool cluster resource, then click **Online**.

The resource is brought online on a node in its **Preferred Nodes** list, according to availability. Wait for the resource to report a Running state before you continue.

12.13.2 Renaming a Shared NSS Volume with NSSMU

1 Log in to the master node as the `root` user, then launch a terminal console.

2 Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

3 Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

4 Activate the pool locally on the master node:

4a At the command prompt, enter

```
nssmu
```

4b On the NSSMU **Main Menu**, select **Pools**, then press Enter.

4c On the Pools page, select the pool, then press F7 **Activate**.

4d Press Esc to return to the NSSMU **Main Menu**.

5 Mount the volume locally on the master node:

5a On the NSSMU **Main Menu**, select **Volumes**, then press Enter.

5b On the Volumes page, select the volume, then press F7 **Mount**.

6 On the Volumes page, view the volume properties to ensure that the attribute **Allow renaming mount point** has been enabled. If it is disabled, select the attribute value, press y (Yes) to enable it, then select **Apply** and press Enter.

7 Rename the volume:

7a On the Volumes page, select the mounted NSS volume that you want to rename.

7b Press F6 **Rename**, specify the new name, then press Enter.

7c Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.

If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.

8 Dismount the volume:

8a On the NSSMU Volumes page, select the volume, press F7 **Dismount**.

8b Read the warning message that open files will be closed by the action, then press y (Yes) to confirm the dismount.

8c Press Esc to return to the NSSMU **Main Menu**.

9 Deactivate the pool on the cluster node:

9a On the NSSMU **Main Menu**, select **Pools**, then press Enter.

9b On the Pools page, select the pool, then press F7 **Deactivate**.

9c Read the warning message that the pool's volumes will be dismounted, then press y (Yes) to confirm the deactivation.

9d When the page refreshes, verify that the pool state is **deactive**.

9e Press Esc twice to exit NSSMU.

10 Verify that the volume name was updated in the scripts for its pool cluster resource:

10a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

10b On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.

10c Click the **Scripts** tab.

10d On the Load Script page, verify that the volume name was changed in the `ncpcon mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click **Apply**.

```
exit_on_error ncpcon mount <new_volume_name>=<volume_id>
```

For example:

```
exit_on_error ncpcon mount USERS=252
```

10e Click **Unload Script**, then check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

The unload script does not contain any default volume commands.

10f Click **Monitor Script**, then verify that the volume name was changed in the `ncpcon volume` command. Also check any custom commands where the volume name appears. If you modify the monitor script, click **Apply**.

```
exit_on_error ncpcon volume <new_volume_name>
```

For example:

```
exit_on_error ncpcon volume USERS
```

10g Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

11 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

12.13.3 Renaming a Shared NSS Volume in a DST Cluster Resource

A Dynamic Storage Technology (DST) shadow volume pair consists of two NSS volumes that are managed together. Users see a single merged view of the data stored on the two volumes. This section assumes some understanding of how clustered DST volumes are configured. For more information, see "Configuring DST Shadow Volume Pairs with OES Cluster Services" in the *OES 2023: Dynamic Storage Technology Administration Guide*.

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. Renaming a volume changes the Volume object name, updates the volume information for the pool, and updates the volume name in its pool's resource scripts. However, a volume rename for a secondary volume cannot also update the scripts for the primary pool's resource because its default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary volume.

When you rename a shared volume in a DST cluster resource, you also affect the settings that define the DST shadow relationship. The DST settings are not automatically updated when you rename the primary volume or secondary volume.

The DST relationship between the primary volume and secondary volume is defined by the `ncpcon mount` command in the DST cluster resource. DST settings that use the volume names are found in the following files that exist on each cluster node. You must manually remove the information on each node.

- `/etc/opt/novell/ncp2nss.conf`

  This file contains the `EXCLUDE_VOLUME <secondary_volume_name>` entry that prevents the secondary volume from being mounted in NCP. For example:

  ```
  EXCLUDE_VOLUME VOLD
  ```

  The rename function does not remove this entry from the file. Cluster Services also does not remove the entry from the file. After the volume is renamed, when you bring the resource online on a node, a new entry is made for the secondary volume with the new name. You must manually remove the entry for the old volume name from the file on each node as part of the rename process.

- `/etc/opt/novell/ncpserv.conf`

  This file contains the shadow volume entry. For example:

  ```
  SHADOW_VOLUME HOME /media/nss/VOLD
  ```

  After you rename the primary volume or secondary volume, a new entry is made in the file with the new primary volume name or secondary volume path. You must remove the entry for the DST volume from the file as part of the rename process.

After you rename the volume, new entries are added to the `ncp2nss.conf` file and `ncpserv.conf` file on the node where you bring the modified resource online.

To rename a volume in a shared DST shadow volume pair:

1. Log in to the master node as the `root` user, then launch a terminal console.
2. Ensure that the DST cluster resource is running on the master node. Enter

   ```
   cluster status
   ```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

3 Take the DST cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

4 Open the `/etc/opt/novell/ncpserv.conf` file in a text editor, remove the entry for the DST shadow volume pair, then save the file. Repeat this task for every node in the cluster.

```
SHADOW_VOLUME <primary_volume_name> /media/nss/<secondary_volume_name>
```

For example:

```
SHADOW_VOLUME HOME /media/nss/VOLD
```

5 If you are renaming the secondary volume, open the `/etc/opt/novell/ncp2nss.conf` file in a text editor, remove the EXCLUDE_VOLUME entry for the volume, then save the file. Repeat this task for every node in the cluster.

```
EXCLUDE_VOLUME <secondary_volume_name>
```

For example:

```
EXCLUDE_VOLUME VOLD
```

6 Activate the pool locally on the master node:

6a At the command prompt, enter

```
nssmu
```

6b On the NSSMU **Main Menu**, select **Pools**, then press Enter.

6c On the Pools page, select the pool that contains the volume you want to rename, then press F7 **Activate**.

6d Press Esc to return to the NSSMU **Main Menu**.

7 Mount the volume to be renamed locally on the master node:

7a On the NSSMU **Main Menu**, select **Volumes**, then press Enter.

7b On the Volumes page, select the volume, then press F7 **Mount**.

8 On the Volumes page, view the volume properties to ensure that the attribute **Allow renaming the mount point** has been enabled. If it is disabled, select the attribute value, press y (Yes) to enable it, then select **Apply** and press Enter.

9 Rename the volume:

9a On the Volumes page, select the mounted NSS volume that you want to rename.

9b Press F6 **Rename**, specify the new name, then press Enter.

9c Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.

If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.

10 Dismount the volume:

10a On the NSSMU Volumes page, select the volume, then press F7 **Dismount**.

10b Read the warning message that open files will be closed by the action, then press y (Yes) to confirm the dismount.

10c Press Esc to return to the NSSMU **Main Menu**.

11 Deactivate the pool that you activated in Step 6:

11a On the NSSMU **Main Menu**, select **Pools**, then press Enter.

11b On the Pools page, select the pool, then press F7 **Deactivate**.

11c Read the warning message that the pool's volumes will be dismounted, then press y (Yes) to confirm the deactivation.

11d When the page refreshes, verify that the pool state is **Deactive**.

11e Press Esc twice to exit NSSMU.

12 Modify the volume name in the scripts for its DST cluster resource:

12a In iManager, select **Clusters > My Clusters**, then select the cluster that you want to manage.

12b On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.

12c Click the **Scripts** tab.

12d On the Load Script page, modify the volume name in the `ncpcon mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click **Apply**.

```
exit_on_error ncpcon mount
<primary_volume_name>=<volume_id>,shadowvolume=<secondary_volume_name>
```

For example, if you rename the primary volume from HOME to USERS, and the secondary volume is VOLD, change this line

```
exit_on_error ncpcon mount HOME=252,shadowvolume=VOLD
```

to this:

```
exit_on_error ncpcon mount USERS=252,shadowvolume=VOLD
```

If you rename the secondary volume name from VOLD to HOME_SH, and the primary volume is named HOME, change this line

```
exit_on_error ncpcon mount HOME=252,shadowvolume=VOLD
```

to this:

```
exit_on_error ncpcon mount HOME=252,shadowvolume=HOME_SH
```

12e Click **Unload Script**, then check any custom commands where the volume name appears. If you modify the unload script, click **Apply**.

The unload script does not contain any default volume commands.

12f If you renamed the primary volume, click **Monitor Script**, then modify the primary volume name in the `ncpcon volume` command. The secondary volume is not monitored because it is hidden from NCP. If you modify the monitor script, click **Apply**.

```
exit_on_error ncpcon volume <new_primary_volume_name>
```

For example, if you rename the primary volume from HOME to USERS, change this line

```
exit_on_error ncpcon volume HOME
```

to this:

```
exit_on_error ncpcon volume USERS
```

12g Click **Cancel** to return to the Cluster Manager page, or click **OK** to apply any changes you made in the scripts.

13 Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

The first time that the modified DST cluster resource is brought online on a node, the new entries are added to the `ncp2nss.conf` file and `ncpserv.conf` file.

12.14 Renaming the Mount Point Path for a Shared NSS Volume (Using a Custom Mount Point for a Shared NSS Volume)

This section describes how to modify the load and unload scripts to specify a custom mount point for a clustered NSS volume. The NSS volume attribute **Allow renaming mount point** must be enabled.

For local NSS volumes, the Linux mount point for the device that contains the volume is stored in the `/etc/fstab` file, and the NCP mount point for the volume is stored in the `/etc/opt/novell/ncpserv.conf` file. By design, clustered NSS volumes do not have entries in these files. The mounting and dismounting of clustered volumes is controlled by the load and unload scripts. The mount location is assumed to be the default `/media/nss/<volume_name>`.

In an NSS pool cluster resource's load script, the pool activation command automatically mounts the NSS volume in Linux at the default mount point `/media/nss/<volume_name>`. The `ncpcon mount` command creates an NCP volume instance for the NSS volume and also assumes the default mount point.

If you enable the NSS volume attribute **Allow renaming mount point**, you can modify the load and unload scripts to mount and dismount the NSS volume and NCP instance of the volume in a custom location. However, you must use a Linux command to dismount the volume from the default location, then mount it in the new location. Use the `path` option in the `ncpcon mount` command to specify the preferred mount point path. Before you bring the resource online by using the modified

load script, the custom mount point path must be created manually on all nodes where the resource is allowed to load. You can alternatively add the Linux `mkdir -p` command before the Linux `mount` command to create the path.

After you modify the load and unload scripts to use a custom mount location, you should use the `cluster online` command to mount the volume, and use the `cluster offline` command to dismount the volume. If you need to mount the volume locally, do not use NSSMU or the Storage plug-in in iManager to mount or dismount the volume. These tools rely on information in `/etc/fstab` file to provide the preferred mount point location, but the clustered volume has no entry there.

- To mount the volume locally, take the resource offline, activate the pool, then use the Linux `mount` command to mount it in Linux and the `ncpcon mount` command to mount it for NCP. Ensure that you specify the custom mount point location in the commands as you do in the load script.
- To dismount the volume locally, use the `ncpcon dismount` command to dismount the volume from NCP, then use the Linux `umount` command to dismount the volume from Linux, then deactivate the pool. Bring the resource online to mount the volume for the cluster.

The following sample values are used in the NSS pool cluster resource scripts in the following procedure. Replace the values with the actual values for your own clustered pool and volume.

Variable	Sample Value
Cluster resource's virtual server name	`CLUS1-POOL-D-SERVER`
Resource IP address	10.10.10.44
Device	`/dev/sdd`
Pool name	`POOL_D`
Volume name	`VOL_D`
Volume ID	252 (valid values are 0 to 254)
Advertising protocols enabled for the resource	NCP (You can add others.)
Default mount point location	`/media/nss/VOL_D`
New mount point location	`/usr/novell/nss/VOL_D`

To rename the mount point path for a clustered NSS volume to use a custom mount point location:

1 Verify that the NSS volume attribute **Allow renaming mount point** is enabled to allow the mount point to be renamed. The attribute is enabled by default.

1a Log in as the `root` user to the server where the NSS pool cluster resource is online.

1b Open a terminal console, then launch NSSMU:

```
nssmu
```

1c In the NSSMU **Main Menu**, select **Volumes**, then press Enter.

1d Select the volume in the Volumes list, then press Enter to open its Volume Properties page.

1e Use the Down arrow to navigate to and select the value for the **Allow renaming mount point** attribute.

1f If the value is set to no, press y (Yes) to enable it.

1g Use the arrow keys to navigate to and select **Apply**, then press Enter.

1h Press Esc twice to exit NSSMU.

2 Open iManager in a web browser, then log in as a cluster administrator user.

3 In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click the **Cluster Manager** tab.

4 On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the **Scripts** tab.

The Scripts tab automatically displays the load script.

5 On the Load Scripts page, view the default load script for the NSS pool cluster resource.

The following sample load script is enabled for NCP as an advertising protocol. The pool name is POOL_D and the volume name is VOL_D.

```
#!/bin/bash

. /opt/novell/ncs/lib/ncsfuncs

exit_on_error nss /poolact=POOL_D

exit_on_error ncpcon mount VOL_D=252

exit_on_error add_secondary_ipaddress 10.10.10.44

exit_on_error ncpcon bind --ncpservername=CLUS1-POOL-D-SERVER --ipaddress=10.10.10.44

exit 0
```

6 Modify the load script to specify the new mount point path:

6a After the pool activation command, add a Linux `umount` command to dismount the NSS volume from Linux.

```
# unmount the NSS volume from Linux (currently mounted at /media/nss/
volname)
ignore_error umount -t nssvol VOL_D
```

The pool activation automatically mounts the volume on Linux at `/media/nss/<volume_name>`. You must dismount the volume from Linux before you can remount it for Linux at the new location.

6b After the `umount` command, add a Linux `mkdir` command to create the new mount point path on nodes if it does not exist.

```
# create the custom mount point path if it does not exist on the node
ignore_error mkdir -p /usr/novell/nss/VOL_D
```

You can alternatively create the new path by using the `mkdir` command in a terminal console on each node where the resource is allowed to load. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

6c After the `mkdir` command, add a Linux `mount` command to mount the NSS volume at the new mount point path.

```
# mount the NSS volume using the Linux mount command
exit_on_error mount -t nssvol VOL_D /usr/novell/nss/VOL_D -o name=VOL_D
```

6d Add the `path` option to the `ncpcon mount` command in order to specify the custom mount point path.

```
# add the path command to create the NCP instance at the new location
exit_on_error ncpcon mount VOL_D=252,path=/usr/novell/nss/VOL_D
```

6e Review the modified load script to ensure the correct order of commands:

```
#!/bin/bash

. /opt/novell/ncs/lib/ncsfuncs

exit_on_error nss /poolact=POOL_D

# unmount the NSS volume from Linux (currently mounted at /media/nss/volname)
ignore_error umount -t nssvol VOL_D

# create the custom mount point path if it does not exist on the node
ignore_error mkdir -p /usr/novell/nss/VOL_D

# mount the NSS volume using the Linux mount command
exit_on_error mount -t nssvol VOL_D /usr/novell/nss/VOL_D -o name=VOL_D

# add the path to create the NCP instance at the new location
exit_on_error ncpcon mount VOL_D=252,path=/usr/novell/nss/VOL_D

exit_on_error add_secondary_ipaddress 10.10.10.44

exit_on_error ncpcon bind --ncpservername=CLUS1-POOL-D-SERVER --
ipaddress=10.10.10.44

exit 0
```

IMPORTANT: If you have multiple volumes with a non-default mount point, every `mount -t nssvol` command should be followed by the `ncpcon mount` command.

6f Click **Apply**.

Wait for the page to refresh before you continue.

7 On the Scripts tab, click **Unload Script**.

8 On the Unload Script page, view the default unload script for the NSS pool cluster resource.

The following sample unload script is enabled for NCP as an advertising protocol. The pool name is POOL_D and the volume name is VOL_D.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

ignore_error ncpcon unbind --ncpservername=CLUS1-POOL-D-SERVER --
ipaddress=10.10.10.44

ignore_error del_secondary_ipaddress 10.10.10.44

ignore_error nss /pooldeact=POOL_D

exit 0
```

9 Modify the unload script to dismount the volume and NCP instance from the custom mount point location:

9a After the `del_secondary_ipaddress` command, add an `ncpcon dismount` command to dismount the NCP instance from the new mount point location.

```
# dismount the NSS volume from NCP
ignore_error ncpcon dismount VOL_D
```

9b After the `ncpcon dismount` command, add a Linux `umount` command to unmount the NSS volume from Linux.

```
# unmount the NSS volume from Linux
ignore_error umount -t nssvol VOL_D
```

9c Review the modified unload script to ensure the correct order of commands:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

ignore_error ncpcon unbind --ncpservername=CLUS1-POOL-D-SERVER --
ipaddress=10.10.10.44

ignore_error del_secondary_ipaddress 10.10.10.44

# dismount the NSS volume from NCP
ignore_error ncpcon dismount VOL_D

# unmount the NSS volume from Linux
ignore_error umount -t nssvol VOL_D

ignore_error nss /pooldeact=POOL_D

exit 0
```

9d Click **Apply**.

Wait for the page to refresh before you continue.

10 (Optional) On the Scripts page, click **Monitor Script**, then view the monitor script.

No changes are required for the monitor script.

11 At the bottom of the **Scripts** page, click **OK** to save your changes and close the Properties page.

The changes do not take effect until the resource is taken offline and brought online.

12 Take the resource offline, then bring the resource online in order for the script changes to take effect.

12a On the Cluster Manager page, select the check box next to the resource, then click **Offline**. Wait for the status to report that it is offline, then continue.

12b Select the check box next to the resource, then click **Online**.

12c Verify that the resource comes online and reports a **Running** state.

If the resource goes into a Comatose state, there might be a mistake in the lines you added or modified in the script. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

13 In NSSMU, verify that the new mount point is used when the pool cluster resource is brought online:

13a Log in as the `root` user to the node that is hosting the resource, then start NSSMU by entering:

```
nssmu
```

13b From the NSSMU **Main Menu**, select **Volumes**, then press Enter.

13c In the **Volumes** list, select the clustered NSS volume.

13d View the volume details to verify that the mount point has changed.

For example, the mount point in this example should display `/usr/novell/nss/VOL_D`.

13e Press Escape twice to exit NSSMU.

14 In OES Remote Manager, verify that the new mount point is used for the NSS volume and its NCP instance when the pool cluster resource is brought online:

14a In a web browser, launch OES Remote Manager for the node that is hosting the resource, then log in as the `root` user.

14b On the File System Management page, view the mount information for the NSS volume under **File Systems**, and view its NCP instance under **NCP Volumes**. Note that both volumes are mounted at the new mount point location `/usr/novell/nss/VOL_D`.

14c Exit OES Remote Manager.

15 Open a Linux terminal console as the `root` user, then list the files in the new mount point location.

```
ll /usr/novell/nss/VOL_D
```

For example, the following is the list for a new NSS volume. You should see the default `~DFSINFO.8-P` file and the hidden `._NETWARE` directory.

12.15 Changing the Name Space for a Clustered NSS Volume

The load script for an NSS volume assumes that the default name space of LONG is used when mounting the clustered NSS volume. If you change the name space setting for shared volume by using NSSMU or the Storage plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume. See Section 12.6.3, “Specifying Custom Mount Options,” on page 224.

1 Open iManager in a web browser, then log in as a cluster administrator user.

2 In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster, then click the **Cluster Manager** tab.

3 Modify the mount point path value in the load scripts for the pool cluster resource:

3a On the Cluster Manager page, select the resource’s name link to open its Cluster Properties page, then click the **Scripts** tab.

The Scripts tab automatically displays the load script.

3b Modify the load script:

3b1 In the load script, add the `/opt=ns=<long|unix|dos|mac>` option for the `ncpcon mount` command and specify the custom name space to use for the volume. For example:

```
exit_on_error ncpcon mount /opt=ns=unix USERS=254
```

3b2 Click **Apply**, then click **OK** to acknowledge the confirmation message.

3c At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online.

4 Take the resource offline, and bring the resource online in order for the script changes to take effect.

4a On the Cluster Manager page, select the check box next to the resource, then click **Offline**. Wait for the status to report that it is offline, then continue.

4b Select the check box next to the resource, then click **Online**.

4c Verify that the resource comes online and reports a **Running** state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the script. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

12.16 Adding a Volume to a Clustered Pool

When you use NSSMU, the Storage plug-in for iManager, or the NLVM command line interface to add an NSS volume to a clustered pool, a `mount` command is automatically added to the load script for the pool cluster resource. Offline and online the primary pool cluster resource to apply the modified load script.

If you add a volume to the primary pool for a clustered DST shadow volume, the `mount` command is added twice in the primary pool's cluster load script, once after the primary pool's activation command and once after the secondary pool's activation command. You must manually delete the instance that occurs after the secondary pool's activation, then offline and online the primary pool cluster resource to apply the modified load script.

1. In iManager under **Roles and Tasks**, select **Storage > Volumes**.
2. Browse to select the Cluster object of the cluster where the pool is currently online.

   You can also select the server where the pool cluster resource is currently mounted.
3. Click **New** to open the **New Volume** wizard, then create the new volume on the shared pool.

   See "Creating Unencrypted NSS Volumes" in the *OES 2023: NSS File System Administration Guide for Linux*.
4. In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster you want to manage.

   If the cluster you want to manage is not in the list, you can add it. Click **Add**, browse to locate and select the Cluster object of the cluster, then click **OK**.
5. If you added a volume to the primary pool for a DST shadow volume, modify the load script to remove the instance of the mount command that immediately follows the pool activation line for the secondary pool.
   - **5a** Select **Cluster Options**.
   - **5b** Click the name link of the pool cluster resource to open its Properties page, then click the **Scripts** tab.
   - **5c** In the load script, delete the `mount` command for the new volume that follows the pool activation command for the secondary pool.
   - **5d** Click OK to save your changes.
   - **5e** Click **Cluster Manager** to return to the Cluster Manager page.
6. On the Cluster Manager page in the **Cluster Resource** list, select the check box next to the pool cluster resource, then click **Offline** to take the resource offline.

   Wait for the resource to report itself as Offline before you continue.
7. In the **Cluster Resource** list, select the check box next to the pool cluster resource, then click **Online** to bring the resource online.

   This applies the modified load script, which mounts the newly created volume.

12.17 Changing an Assigned Volume ID

OES Cluster Services supports NCP client access to cluster-enabled NSS volumes by using a unique volume ID to mount the volume in the cluster. The volume ID is used by an NCP client only for automatic reconnects to the volume after failover or migration of a cluster resource.

Valid volume ID values are 0 to 254 (up to 255 mounted volumes per server). 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. Values start at 254 for the first volume in the cluster and decrease for each new volume. You can view the volume IDs assigned on a node by using the `ncpcon volumes` command.

In older operating systems, there was a mounted volume limit of 64 volumes (values 0 to 63). Some older applications might have hardcoded the old maximum limit of 64 mounted volumes, and might not be able to handle volume IDs greater than 63. You can use the Clusters plug-in to iManager to modify the volume ID in the scripts for a given cluster resource in order to specify a value that works for the application.

Changing the volume ID does not affect the ability to log in to, back up, or access the data. However, there is a brief disruption of service as the cluster resource is offlined and onlined to apply the script changes. If you modify the volume ID for a volume in the cluster resource scripts, ensure that you do the following:

- Volume IDs that you manually assign must be unique across every volume on all servers in cluster.
- After the value is changed, you must offline and online the cluster resource for the volume in order to mount the volume with its new volume ID.
- After the volume is mounted with its new ID, the clients must log out and log in to the volume in order to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

  Some clients might cache the volume IDs. To reset the cached value, the client must be rebooted and reconnected to the volume.
- After the volume is mounted with its new ID, if the backup software is running on a client connection, you might need to restart the backup to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

12.18 Expanding the Size of a Clustered Pool

NSS pools are made up of segments of space from one or more devices up to a maximum total size of 8 TB (terabytes). Each device can be up to 2 TB in size if a DOS partitioning scheme is used, or up to 8 TB in size if a GPT partitioning scheme is used. You can expand a pool by adding more space from the same device or from a different shared device. If you extend the size of a LUN device that is currently being used by the pool, you must make the extended size visible to all nodes in the cluster

before you add the newly available space to the clustered pool. If you add a new shared LUN device, you must make the device visible to all nodes in the cluster before you add space from the device to the clustered pool.

- Section 12.18.1, “Planning to Expand a Clustered Pool,” on page 258
- Section 12.18.2, “Increasing the Size of a Clustered Pool by Extending the Size of Its LUN,” on page 259
- Section 12.18.3, “Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN,” on page 262
- Section 12.18.4, “Verifying the Expanded Pool Cluster Resource on All Nodes,” on page 265

12.18.1 Planning to Expand a Clustered Pool

Consider the guidelines in this section when you plan to expand the size of a clustered NSS pool.

- “Free Space on a Shared Device” on page 258
- “Extending a LUN Might Require a Service Outage” on page 258
- “Adding a LUN Might Require a Service Outage” on page 259
- “Multipath I/O Devices” on page 259

Free Space on a Shared Device

To expand the size of a pool, you need free unpartitioned space on the same shared device or on other shared devices that can be failed over to the same node in a cluster. To extend the size of an existing LUN, the free unpartitioned space must immediately follow the LUN on the same shared device.

Extending a LUN Might Require a Service Outage

OES Cluster Services and NSS allow you to perform an online extension of a LUN that is used by a clustered pool without taking the resource offline or stopping Cluster Services. However, not all SAN storage arrays and vendor-specific storage drivers fully support transparent online extension of LUNs. Refer to the third-party vendor documentation for information about how to extend a LUN. We recommend that you test the specific hardware environment and hardware configuration to confirm that it behaves correctly before you perform an online LUN extension.

IMPORTANT: To prevent any chance of data loss or corruption, we recommend that you back up the volumes before you expand a LUN.

After you extend the LUN size on the SAN storage array, you must scan for devices on each node to update the node’s storage map with its new size information before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node’s multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the extended LUN size.

Adding a LUN Might Require a Service Outage

OES Cluster Services and NSS allow you to add a LUN to a clustered pool without taking the resource offline or stopping Cluster Services. The LUN must be able to be failed over to the same node in a cluster as the pool cluster resource. Refer to the third-party vendor documentation for information about how to add a LUN.

After you create a new LUN on the SAN storage array and assign it to all nodes, you must scan for devices on each node to update the node's storage map with the new device before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node's multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the new LUN.

Multipath I/O Devices

Consider the following requirements and guidelines as you work with devices that have multiple I/O paths:

- The procedures in this section assume that you use the Linux Device Mapper - Multipath I/O (DM-MPIO) software to manage the multiple I/O paths for devices. You must modify the instructions accordingly if you are using a third-party multipath solution.
- Before you expand the pool size on the active node or allow the pool cluster resource to fail over to other nodes, you must update the multipath map on each node in the cluster so that DM-MPIO recognizes the LUN changes:
  - **Extend the LUN size:** On each node, you can restart the `multipathd` daemon, or you can use the `multipathd` command to resize the multipath map for the LUN.
  - **New LUN:** On each node, use the `multipath -v2` command to rebuild the node's multipath map.
- For a device with multiple I/O paths, use the device's multipath device name (such as `mpathb`) or device ID.
- Steps in this section that pertain only to multipath I/O devices are preceded by the keyword "MPIO".

12.18.2 Increasing the Size of a Clustered Pool by Extending the Size of Its LUN

One way to increase the size of a clustered pool is to add space to its LUN device on the SAN storage array, and then increase the size of the pool to use the newly available space. Before you extend the size of a LUN, ensure that you understand the cautions in "Extending a LUN Might Require a Service Outage" on page 258.

- "Extending the Size of a LUN" on page 260
- "Increasing the Pool Size by Using Free Space on the Same LUN" on page 262

Extending the Size of a LUN

You can extend the size of the LUN if there is free unpartitioned space on the device that immediately follows the LUN. If the LUN has multiple I/O paths, special handling is required to update the device size reported in the multipath map on each node. These steps are marked with the keyword “MPIO”.

1 Identify the device that is used by the pool cluster resource:

   1a In iManager, go to **Storage > Pools**, then select the Cluster object of the cluster.

   1b In the Pools list, select the pool, then view the device name of the LUN under the pool **Details** area.

   For example, a device with a single path has a device node name such as `sdb`. A device with multiple I/O paths has a multipath device name such as `mpathb`.

2 (Offline LUN extension only) If the SAN storage array or vendor-specific storage driver does not support online LUN extension, take the pool cluster resource offline. Open a terminal as the `root` user, then at the command prompt, enter:

```
cluster offline <resource_name>
```

3 On the SAN storage array, use the third-party vendor tools to extend the size of the existing LUN.

   Devices can be up to 2 TB in size with the DOS partitioning scheme. Devices can be up to 8 TB in size with the GPT partitioning scheme.

   For example, increase the size of the existing LUN from 100 GB to 150 GB.

4 Beginning with the active node, perform the following steps on each node in turn in order to update the node’s storage map with the new LUN size:

   4a Log in to the node as the `root` user, then open a terminal console.

   4b Scan for devices to recognize the new size for the LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan [--luns=XX]
```

   WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

   Use the `--luns=`*XX* option if you want to scan only for the extended LUN. You can use the `lsscsi(8)` command to locate the LUN number for the device. Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device.

   For information about other `rescan-scsi-bus.sh` options, see “Scanning for New Devices without Rebooting” (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-multipath.html#sec-multipath-best-practice-scandev) in the *SLES 12 Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html).

   4c Scan for storage objects. At the command prompt, enter:

```
nlvm rescan
```

4d (MPIO) If the device has multiple I/O paths, do either of the following to update the multipath map with the new size information for the extended LUN:

- **Resize the multipath map for the LUN:** At the command prompt, enter:

  ```
  multipathd -k'resize map <mpio_map_name>'
  ```

  There is no space between `-k` and the `r'esize map <mpio_map_name>'`option. You can use the `multipath -ll` command to locate the multipath map name for the device.

  For example, the following command resizes the map entry for a LUN device with the multipath map ID of `36001c230ce31da000eb0fa8d1ccb0c02`:

  ```
  multipathd -k'resize map 36001c230ce31da000eb0fa8d1ccb0c02'
  ```

  You can alternatively use the `multipathd -k` command to enter interactive mode and issue the `resize map <mpio_map_name>` command. From this mode, the available commands can be viewed by entering `help`. When you are finished entering commands, press Ctrl+D to quit.

- **Restart the multipathd daemon:** At the command prompt, enter one of the following commands:

  ```
  rcmultipathd restart
  ```

  ```
  /etc/multipathd restart
  ```

4e (MPIO) Verify that the new size of the LUN is correctly reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (sysfs, the device mapper, path checkers, and so on).

For example, the `mpathb` device reports a new size of 150 GB.

```
mpathb (36001438005de9b8d0000800007520000) dm-1 HP,HSV450
[size=150G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=100][active]
 \_ 4:0:1:7 sdae       65:224 [active][ready]
 \_ 2:0:1:7 sdu        65:64  [active][ready]
\_ round-robin 0 [prio=20][enabled]
 \_ 4:0:0:7 sdv        65:80  [active][ready]
 \_ 2:0:0:7 sdg        8:96   [active][ready]
```

4f Repeat Step 4a through Step 4e on each node in the cluster.

Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the extended LUN size.

5 In iManager, verify that the new device size is reported correctly to NSS:

5a In iManager, go to **Storage > Devices**, then select the Cluster object of the cluster.

5b Select the device (such as `mpathb`), then view the device capacity under the device's **Details** area.

5c If the new LUN size is reported correctly, continue to Step 6. If the size is not reported correctly, repeat Step 4, then check again.

6 (Offline LUN extension only) If you took the pool cluster resource offline in Step 2, you can now bring pool cluster resource online. Open a terminal as the `root` user, then at the command prompt, enter:

```
cluster online <resource_name>
```

7 After the new LUN size has been recognized by all nodes, continue with "Increasing the Pool Size by Using Free Space on the Same LUN" on page 262.

Increasing the Pool Size by Using Free Space on the Same LUN

The pool size is not automatically expanded to use the new free space on the LUN. After you have expanded the size of the existing LUN and the new size is recognized by all nodes in the cluster, you can increase the size of the pool by allocating space from the newly available free space.

1 In iManager, go to **Storage > Pools**, then select the Cluster object of the cluster.

2 On the Pool page, select the pool, then view the pool's current size in the **Pool Details** area.

The **Total Space** for the pool reports only the amount of space that has been allocated to the pool. The pool size is not automatically expanded to use the new free space on the LUN.

3 Increase the size of the pool by adding space from the same LUN:

3a On the Pool page, select the pool, then click **Increase Size**.

The Expand a Pool wizard opens and presents all devices with available free space.

3b In the list of devices, select the check box for the newly expanded LUN, such as `mpathb`, then specify the amount of space to add in MB.

Typically, this is the reported amount of free space available.

3c Click **Finish** to save and apply the change.

4 On the Pool page, select the pool, then verify that the **Total Space** field reports the new size.

5 Go to **Storage > Partitions**, then verify that a new partition exists for the LUN and is assigned to the pool.

6 For each volume on the pool, if the volume quota is set to **No Quota**, the volume can grow to the size of the pool. Go to **Storage > Volumes**, select the volume, then verify that its **Available Space** field has increased by the amount of space that was added to the pool.

7 Continue with "Verifying the Expanded Pool Cluster Resource on All Nodes" on page 265.

12.18.3 Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN

You can increase the size of a clustered pool by adding space to the pool from a different shared LUN device that can be failed over with the resource. Before you add a LUN, ensure that you understand the cautions in "Adding a LUN Might Require a Service Outage" on page 259.

- "Creating and Sharing a New LUN Device" on page 263
- "Increasing the Pool Size by Adding Free Space from a Different LUN" on page 265

Creating and Sharing a New LUN Device

If you do not have a shared device that can be failed over with the resource, you can create a new LUN and share it with all nodes in the cluster. If the LUN has multiple I/O paths, special handling is required to add the device to the multipath map on each node. These steps are marked with the keyword “MPIO”.

1 Use third-party tools for the SAN storage array to create a new shared LUN device, and assign it to each cluster node.

2 Beginning with the active node, perform the following steps on each node in turn in order to update the node’s storage map and multipath map with the new device:

2a Log in to the node as the `root` user, then open a terminal console.

2b Scan for devices to help the OS recognize the new LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan
```

WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

For information about other `rescan-scsi-bus.sh` options, see “Scanning for New Devices without Rebooting” (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-multipath.html#sec-multipath-best-practice-scandev) in the *SLES 12 Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html).

2c Use the `lsscsi` command to verify that the LUN is seen by the OS. At the command prompt, enter:

```
lsscsi
```

Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device. In the sample output below, the `/dev/sdc` and `/dev/sde` devices are two individual paths for LUN 2.

```
# lsscsi
[4:0:0:0]    disk    DGC      RAID 5           0216  /dev/sdb
[4:0:0:1]    disk    DGC      RAID 5           0216  /dev/sdf
[4:0:0:2]    disk    DGC      RAID 5           0216  /dev/sdc
[4:0:1:0]    disk    DGC      RAID 5           0216  /dev/sdd
[4:0:1:1]    disk    DGC      RAID 5           0216  /dev/sdg
[4:0:1:2]    disk    DGC      RAID 5           0216  /dev/sde
```

If the LUN is not seen, repeat this step using different `lsscsi` command line parameters to scan for all devices (such as `-w -c -l`). For information about command options, see the `lsscsi(8)` man page.

If the LUN is still not seen by the OS, a server restart might be required.

2d Scan for storage objects. At the command prompt, enter:

```
nlvm rescan
```

2e (MPIO) If the device has multiple I/O paths, rebuild the multipath map. At the command prompt, enter:

```
multipath -v2
```

2f (MPIO) Verify that the new LUN is reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (`sysfs`, the device mapper, path checkers, and so on).

2g Verify that the LUN is now visible to NSS on the node. In NSSMU, go to the Devices page, then verify that the new LUN is listed.

For example, the new device is `mpathj` on the active node. Remember that a multipath device might have a different multipath device name on each node.

2h Repeat Step 2a through Step 2g on each node in the cluster.

Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the new LUN.

3 In iManager, initialize the new LUN, then mark it as shareable for clustering:

3a In iManager, go to **Storage > Devices.**

3b On the Devices page, select the Cluster object of the cluster.

Selecting the Cluster object connects you to the current master node of the cluster.

3c In the Devices list, select the newly added LUN device.

WARNING: Ensure that you select the correct device. Initializing a device destroys all data on it.

3d Click **Initialize**, then click **OK** to confirm that you want to initialize the device.

If you are prompted for the partitioning scheme, specify your preferred partitioning scheme as DOS (up to 2 TB in size) or GPT (up to 8 TB in size).

3e In the Devices list, select the new LUN device.

3f In the **Details** area, select the **Shareable for Clustering** check box, then click **Apply**.

4 After the new LUN is recognized by all nodes, and has been initialized and shared, continue with "Increasing the Pool Size by Adding Free Space from a Different LUN" on page 265.

Increasing the Pool Size by Adding Free Space from a Different LUN

After you have set up a different shared device that can be failed over with the pool cluster resource, you can increase the size of the pool by allocating space from it.

1. In iManager, go to **Storage > Pools**, then select the Cluster object of the cluster.

   Selecting the Cluster object automatically connects you to the current master node of a cluster.

2. Select the pool, then view the pool's current size.

   The **Total Space** for the pool reports only the amount of space that has been allocated to the pool.

3. Increase the size of the pool by adding space from a different LUN:

   3a On the Pool page, select the pool, then click **Increase Size**.

   The Expand a Pool wizard opens and presents all devices with available free space.

   3b In the list of devices, select the check box for the new LUN, such as `mpathj`, then specify the amount of space to add in MB.

   Typically, this is the reported amount of free space available.

   3c Click **Finish** to save and apply the change.

4. On the Pool page, select the pool, then verify that the **Total Space** field reports the new size.

5. For each volume on the pool, if the volume quota is set to **No Quota**, the volume can grow to the size of the pool. Go to **Storage > Volumes**, select the volume, then verify that a volume's **Available Space** field has increased by the amount of space that was added to the pool.

6. After the pool size is increased, continue with Section 12.18.4, "Verifying the Expanded Pool Cluster Resource on All Nodes," on page 265.

12.18.4 Verifying the Expanded Pool Cluster Resource on All Nodes

Verify that the expanded pool size is reported correctly on each node in the cluster.

1. Log in to a node as the `root` user, then open a terminal console.

2. Cluster migrate the pool cluster resource to each node in turn, and use NSSMU to verify that the devices, partitions, pool, and volumes show the correct information on that server:

   2a Cluster migrate the pool cluster resource.

   ```
   cluster migrate resource_name node_name
   ```

   2b Launch NSSMU by entering

   ```
   nssmu
   ```

   2c On the Devices page, select the device and verify that the new size is reported.

   2d On the Partitions page, select each partition for the device, then verify that the partitions are shown for the expanded pool.

   2e On the Pools page, select the pool, then verify the total space available.

2f On the Volumes page, select each volume that has no quota set, then verify that the available space has increased by the amount of space added to its pool.

2g If the information is incorrect, wait a few minutes to allow time for the system to recognize the new partitions, then try again.

3 Repeat the previous validation for each node in the cluster.

12.19 Deleting NSS Pool Cluster Resources

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering, then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the **Storage > Pools** page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

See Section 10.16, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 192.

12.20 Disabling Clustering for a Pool

You can disable clustering for a pool by deleting the resource-related objects, and then re-creating the Pool and Volume objects on a desired node and unsharing the device. We strongly recommend that you delete the resource for a pool from the master node in the cluster.

The Cluster Options page of the Clusters plug-in for iManager provides a **Delete** option that automatically deletes the resource-related objects in eDirectory and updates the cluster information:

- Cluster-named Pool object
- Cluster-named Volume object for each of the pool’s volumes
- Cluster Resource object for the pool
- Virtual server for the cluster resource (NCS:NCP Server object)

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, you must offline the cluster resource for the pool before you mark the device as Not Shareable for Clustering. Then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the **Storage > Pools** page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

This deletes the resource-related objects, but not the storage area they represent.

Use the following procedure to disable clustering for a pool:

1. If the resource is on a non-master node in the cluster, migrate it to the master node.

   As the *root* user, open a terminal console, then enter

   `cluster migrate <resource_name> <master_node_name>`

   The master node must be in the resource's preferred nodes list. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

2. Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

   - Enter the following at the command prompt as the `root` user:

     `cluster offline <resource_name>`

     Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

   - In iManager, go to **Clusters > My Clusters**, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click **Offline**.

     Refresh the Cluster Manager page to verify that the resource has a status of **Offline** before you continue.

3. In iManager, use the Clusters plug-in to delete the cluster resource.

   3a Select **Clusters > My Clusters**, then select the cluster.

   3b Select the **Cluster Options** tab.

   3c Select the check box next to the resource, then click **Delete**.

   3d When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.

   3e In the **Tree View** in iManager, browse to verify that the Cluster Resource object and related objects were removed from eDirectory.

   If necessary, you can manually delete the objects. In iManager, go to **Directory Administration > Delete Objects**, select the objects, then click **OK**.

4. Use the **Update eDirectory** function to re-create Storage objects for the pool and its volumes.

   4a In iManager, select **Storage > Pools**, then select the node where you want the pool to reside as a locally available pool.

   4b Select the pool, then click **Activate**.

   4c Select the pool, then click **Update eDirectory**.

   This creates a Pool object in eDirectory with a name format of *<server_name>_<pool_name>*_POOL.

   4d Select **Storage > Volumes.** The server should still be selected.

4e Select the volume, then click **Mount**.

4f Select the volume, then click **Update eDirectory**.

This creates a Volume object in eDirectory with a name format of *<server_name>_<volume_name>*.

4g Repeat Step 4e through Step 4f for each volume in the pool.

4h In the **Tree View** in iManager, browse to verify that the server-named Pool object and Volume object were created.

5 Unshare the devices that contribute space to the pool, then use a third-party SAN management tool to assign the devices to only the desired server.

5a In iManager, go to **Storage** > **Devices**, then select the desired server (the one you specified in Step 4).

5b Select the device.

5c Deselect the **Shareable for Clustering** check box, then click Apply.

Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the pool where you have disabled clustering.

5d Repeat these steps for each device that contributes space to the pool.

5e Use a third-party SAN management tool to assign the devices to only the desired server.

6 Provide the node's IP address to users who access volumes on the non-clustered pool.

12.21 Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool)

You can re-create the resource objects for a pool by deleting the resource-related objects, creating new Pool and Volume objects, and then cluster-enabling the existing pool. This is the same as disabling and re-enabling clustering for a pool. We strongly recommend that you delete the resource for a pool from the master node in the cluster.

The Cluster Options page of the Clusters plug-in for iManager provides a **Delete** option that automatically deletes the resource-related objects in eDirectory and updates the cluster information:

- Cluster-named Pool object
- Cluster-named Volume object for each of the pool's volumes
- Cluster Resource object for the pool
- Virtual server for the cluster resource (NCS:NCP Server object)

This deletes the resource-related objects, but not the storage area they represent. To delete a resource and create a new one with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the device as Not Shareable for Clustering. Then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the **Storage > Pools** page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

Use the following procedure to delete and re-create a cluster resource for a pool:

1. If the resource is on a non-master node in the cluster, migrate it to the master node.

   As the *root* user, open a terminal console, then enter

   ```
   cluster migrate <resource_name> <master_node_name>
   ```

   The master node must be in the resource's preferred nodes list. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

2. Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

   - Enter the following at the command prompt as the `root` user:

     ```
     cluster offline <resource_name>
     ```

     Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

   - In iManager, go to **Clusters > My Clusters**, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click **Offline**.

     Refresh the Cluster Manager page to verify that the resource has a status of **Offline** before you continue.

3. In iManager, use the Clusters plug-in to delete the cluster resource.

   3a Select **Clusters > My Clusters**, then select the cluster.

   3b Select the **Cluster Options** tab.

   3c Select the check box next to the resource, then click **Delete**.

   3d When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.

   3e In the **Tree View** in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.

   If necessary, you can manually delete the objects. In iManager, go to **Directory Administration > Delete Objects**, select the objects, then click **OK**.

4. Use the **Update eDirectory** function to re-create Storage objects for the pool and its volumes.

   These objects are needed by OES Cluster Services when you re-create the resource.

   4a In iManager, select **Storage > Pools**, then select the master node if you plan to re-create the storage object, or select the node where you want the pool to reside as a locally available pool.

   4b Select the pool, then click **Activate**.

   4c Select the pool, then click **Update eDirectory**.

This creates a Pool object in eDirectory with a name format of *<server_name>_<pool_name>*_POOL.

4d Select **Storage > Volumes.** The server should still be selected.

4e Select the volume, then click **Mount**.

4f Select the volume, then click **Update eDirectory**.

This creates a Volume object in eDirectory with a name format of *<server_name>_<volume_name>*.

4g Repeat Step 4e through Step 4f for each volume in the pool.

4h In the **Tree View** in iManager, browse to verify that the Pool object and Volume object were created.

5 Use the Clusters plug-in to cluster-enable the pool:

For detailed instructions, see Step 5 thru Step 18 in Section 12.5, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 217.

To re-create the cluster resource with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

5a In **Roles and Tasks**, select **Clusters > My Clusters**, select the cluster.

5b Select the **Cluster Options** tab.

5c On the Cluster Options page, click the **New** link in the Cluster Objects toolbar.

5d On the Resource Type page, select the **Pool** radio button, then click **Next**.

5e On the Cluster Pool Information page, specify the following information, then click **Next**:

- Pool name
- Virtual server name
- IP address
- Advertising protocols (NCP, CIFS)
- If you enable CIFS, specify the CIFS Server name.
- Deselect the **Online Resource after Creation** check box.
- Select the **Define Additional Properties** check box.

5f On the Resource Policies page, configure the resource policies for the start, failover, and failback modes, then click **Next**.

See "Configuring the Start, Failover, and Failback Modes for Cluster Resources" on page 182.

5g On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource, then click **Finish**.

See "Configuring Preferred Nodes and Node Failover Order for a Resource" on page 184.

The pool cluster resource appears in the **Cluster Objects** list on the Cluster Options page, such as POOL1_SERVER.

5h (Optional) View the resource properties, and enable monitoring and configure the monitoring script.

5i Bring the resource online. Select the Cluster Manager tab, select the check box next to the resource, then click **Online**.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

If the pool goes comatose, take it offline, check that the pool is deactivated on the local server, then try again.

12.22 Deleting a Clustered Pool

We strongly recommend that you delete a cluster-enabled pool only from the master node in the cluster. This allows the cluster information to be automatically updated.

WARNING: Deleting a pool destroys all data on it.

NSS management tools delete the cluster-enabled pool and its related objects in eDirectory:

- Pool and its volumes from the file system and from NCP
- Cluster-named Pool object
- Cluster-named Volume object for each of the pool's volumes
- Cluster Resource object for the pool
- Virtual server for the cluster resource (NCS:NCP Server object)

When the pool resides on the master node, the cluster information is automatically updated.

When the pool resides on a non-master node, additional steps are required to update the cluster information. A cluster restart might be needed to force the information to be updated.

Use the following procedures to delete a cluster-enabled pool:

- Section 12.22.1, "Deleting a Clustered Pool on the Master Node," on page 271
- Section 12.22.2, "Deleting a Clustered Pool on a Non-Master Node," on page 272

12.22.1 Deleting a Clustered Pool on the Master Node

1. If the pool cluster resource is on a non-master node in the cluster, migrate it to the master node. As the `root` user, open a terminal console, then enter

   ```
   cluster migrate <resource_name> <master_node_name>
   ```

   To migrate the resource, the master node must be in the resource's preferred nodes list. To view or modify the list, see Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.
2. Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:
   - As the `root` user, enter

     ```
     cluster offline <resource_name>
     ```

     Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

- In iManager, go to **Clusters > My Clusters**, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click **Offline**.

  Refresh the Cluster Manager page to verify that the resource has a status of **Offline** before you continue.

3 Delete the pool on the master node by using NSSMU.

  You can alternatively use the Storage plug-in in iManager or the `nlvm delete pool <pool_name>` command.

  3a In NSSMU, select **Pools**, then press Enter.

  3b Select the deactive pool, then press Delete.

  3c Select OK to confirm, then press Enter.

4 In the **Tree View** in iManager, browse the objects to verify that the following objects were deleted:

- Pool object
- Volume object (for each volume in the pool)
- Pool cluster resource object (from the Cluster container)
- Virtual server for the resource (NCS:NCP Server object)

5 (Optional) Unshare the device:

  5a In iManager, go to **Storage** > **Devices**, then select the node where you want the unshared device to reside.

  5b Select the device.

  5c Deselect the **Shareable for Clustering** check box, then click **Apply**.

  Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the cluster-enabled pool you deleted.

  5d Repeat Step 5b to Step 5c for each device that contributes space to the pool.

  5e Use a third-party SAN management tool to assign the devices to only the desired server.

12.22.2 Deleting a Clustered Pool on a Non-Master Node

1 Log in as the `root` user to the non-master node where the cluster resource currently resides, then open a terminal console.

2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

```
cluster offline <resource_name>
```

  Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

3 At the command prompt on the non-master node, enter

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

4 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.

5 Delete the pool on the non-master node by using NSSMU.

You can alternatively use the Storage plug-in in iManager or the `nlvm delete pool <pool_name>` command.

5a In NSSMU, select **Pools**, then press Enter.

5b Select the pool, then press Delete.

5c Select OK to confirm, then press Enter.

6 In the **Tree View** in iManager, browse the objects to verify that the following objects were deleted:

- Pool object
- Volume object (for each volume in the pool)
- Pool cluster resource object (from the Cluster container)
- Virtual server for the resource (NCS:NCP Server object)

7 On the master node, log in as the `root` user, open a terminal console, then enter

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

8 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.

9 In iManager, select **Clusters > My Clusters**, select the cluster, then select the **Cluster Options** tab.

10 Click **Properties**, select the **Priorities** tab, then click **Apply** on the Priorities page.

11 At a command prompt, enter

```
cluster view
```

The cluster view should be consistent.

12 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

If the revision number increased, skip Step 13 and continue with Step 14.

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

13 If the revision number did not automatically update in the previous steps, restart OES Cluster Services by entering the following on one node in the cluster:

```
cluster restart [seconds]
```

For *seconds*, specify a value of 60 seconds or more.

For example:

```
cluster restart 120
```

14 (Optional) Unshare the device:

14a In iManager, go to **Storage** > **Devices**, then select the node where you want the unshared device to reside.

14b Select the device.

14c Deselect the **Shareable for Clustering** check box, then click **Apply**.

Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the cluster-enabled pool you deleted.

14d Repeat Step 14b to Step 14c for each device that contributes space to the pool.

14e Use a third-party SAN management tool to assign the devices to only the desired server.

13 Configuring and Managing Cluster Resources for Shared LVM Volume Groups

After you have installed and configured OES Cluster Services, you can create shared cluster resources for Linux Logical Volume Manager (LVM) volume groups. You create an LVM logical volume on the volume group, and add a Linux POSIX file system such as Btrfs, Ext2, Ext3, ReiserFS, and XFS. This section describes how to configure and manage an LVM volume group cluster resource.

- Section 13.1, “Requirements for Creating LVM Cluster Resources,” on page 275
- Section 13.2, “Initializing a SAN Device,” on page 278
- Section 13.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 280
- Section 13.4, “Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template,” on page 295
- Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308
- Section 13.6, “Enabling NCP File Access for a Clustered LVM Volume,” on page 314
- Section 13.7, “Modifying the LVM Resource Scripts for Other Services,” on page 321
- Section 13.8, “Renaming the Mount Point Path for a Clustered LVM Volume,” on page 321
- Section 13.9, “Renaming a Clustered LVM Logical Volume,” on page 323
- Section 13.10, “Disabling Clustering for an LVM Volume Group and Logical Volume,” on page 325
- Section 13.11, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 329
- Section 13.12, “Deleting a Clustered LVM Volume (Created in NSSMU or NLVM),” on page 330
- Section 13.13, “Linux LVM Management Tools,” on page 333

13.1 Requirements for Creating LVM Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in Chapter 4, “Planning for OES Cluster Services,” on page 35.

- Section 13.1.1, “OES Cluster Services,” on page 276
- Section 13.1.2, “Linux Logical Volume Manager 2 (LVM2),” on page 276
- Section 13.1.3, “Clustered Logical Volume Manager Daemon (CLVMD),” on page 276
- Section 13.1.4, “Resource IP Address,” on page 276
- Section 13.1.5, “Shared Storage Devices,” on page 276
- Section 13.1.6, “All Nodes Must Be Present,” on page 277
- Section 13.1.7, “NCP File Access with OES NCP Server,” on page 277
- Section 13.1.8, “Linux File Access Protocols,” on page 278

13.1.1 OES Cluster Services

OES Cluster Services must be installed, configured, and running when you create and manage the shared LVM volume group and logical volume. The cluster must be active.

13.1.2 Linux Logical Volume Manager 2 (LVM2)

The Linux Logical Volume Manager (LVM) 2 software supports LVM volume groups and logical volumes. LVM2 must be installed and running on each node in the cluster. LVM2 runs automatically on OES 11 and later servers; no separate installation or setup is required.

13.1.3 Clustered Logical Volume Manager Daemon (CLVMD)

The Linux Clustered Volume Manager Daemon (CLVMD, `clvmd`) software allows you to exclusively mount a shared volume group on one node at a time in a cluster. It distributes the LVM metadata updates around a cluster. CLVM must be installed and running on each node in the cluster. CLVMD runs automatically on OES 11 and later servers; no separate installation or setup is required.

IMPORTANT: Ensure that you have installed the latest patches for SUSE Linux Enterprise Server 12 SP2. Clustered LVM volume groups require Linux kernel version 2.6.32.45-0.3 or later.

13.1.4 Resource IP Address

Each cluster resource requires a unique static IP address that is in the same subnet as the IP addresses that are used for the cluster and cluster nodes. The IP address is used to provide access and failover capability for the cluster-enabled volume.

13.1.5 Shared Storage Devices

The shared SAN storage device that you use for an LVM volume group cluster resource must be initialized and have no partitions on it. When the device is used in a cluster resource, LVM uses the entire device for the volume group. Ensure that you size your LUNs accordingly. Use the SAN management tools to assign the LUN to all nodes in the cluster.

IMPORTANT: If you use NSS management tools to manage devices, do not enable the Shareable for Clustering option. Doing so adds a 4 KB partition to the device, which makes it unavailable to LVM.

13.1.6 All Nodes Must Be Present

LVM requires the presence of all the nodes in the cluster to modify the metadata on shared storage. This allows LVM to get the exclusive locks it needs to perform actions on shared storage.

Before you attempt to create or modify LVM volume group cluster resources:

- All of the nodes must be joined in the cluster and running properly.
- The `clvmd` daemon must be running on all nodes.

13.1.7 NCP File Access with OES NCP Server

OES NCP Server can be used to provide NCP file access to Linux POSIX file systems on OES servers. Its NCP volumes feature can be used to provide NCP access to files on an LVM volume group cluster resource. NCP Server must be installed, configured, and running on each node in the cluster.

- "Naming Conventions for NCP Volumes" on page 277
- "Creating an NCP Volume for a New Clustered LVM Volume" on page 277
- "Creating an NCP Volume on an Existing Clustered LVM Volume" on page 277
- "Using Antivirus Software with NCP Volumes" on page 278

Naming Conventions for NCP Volumes

NCP volume names can be up to 14 alphanumeric characters, using uppercase letters A through Z and numbers 0 through 9. Underscores (_) are allowed.

If you NCP enable a Linux volume as you create it with NSSMU or the `nlvm create linux volume` command, the NCP volume name uses the specified Linux volume name, but all letters are uppercase. NCP treats the Linux volume name as case insensitive. Ensure that the specified Linux volume name does not exceed 14 characters, does not use special characters, and is unique across all nodes in the cluster for both Linux and NCP.

Creating an NCP Volume for a New Clustered LVM Volume

You can configure NCP file access for an LVM volume group cluster resource when you create the resource by using NSSMU or the `nlvm create linux volume` command. With the NCP option enabled, these tools automatically add commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the volume group cluster resource.

Creating an NCP Volume on an Existing Clustered LVM Volume

You can create an NCP Virtual Server object for the LVM cluster resource to make the resource visible in the eDirectory tree. The virtual server alone does not provide NCP file access.

You can add NCP file access support to an existing LVM cluster resource by creating an NCP Volume object and adding commands to the scripts to manage the volume. For details about setting up NCP volumes on an existing clustered Linux volume, see "Configuring NCP Volumes with OES Cluster Services" in the *OES 2023: NCP Server for Linux Administration Guide*.

Using Antivirus Software with NCP Volumes

For information about using antivirus software with NCP volumes, see "Security" in the *OES 2023: Planning and Implementation Guide*.

13.1.8 Linux File Access Protocols

You can provide native Linux file access to files on an LVM volume group cluster resource for eDirectory users who are enabled for Linux User Management (LUM). The Linux file access protocols must also be LUM enabled on each node of the cluster. For information about LUM-enabling your eDirectory users and native Linux file access protocols, see the *OES 2023: Linux User Management Administration Guide*.

13.2 Initializing a SAN Device

Before you begin, you must initialize the SAN device to set up its device format. You can also initialize a device to wipe its current structure and reconfigure it.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that you want to use for a clustered Linux LVM volume group should contain no partitions and be in an unshared state. When you initialize the SAN device by using NSS management tools, ensure that the **Shareable for Clustering** option is disabled. CLVM handles the sharing for clustered LVM volumes.

IMPORTANT: NLVM and NSSMU do not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

Devices that have never been initialized have a format of `None` (NLVM commands) or `Uninitialized` (in NSSMU). Devices that are being used for a OES Cluster Services SBD (split brain detector) partition also have a format of `None`. You should not use the `nlvm init` command to remove an SBD partition. For information about removing an SBD partition, see Section 9.18, “Creating or Deleting Cluster SBD Partitions,” on page 146.

Use either procedure in this section to initialize the SAN device that you want to use for the LVM volume group. Do not mark it as shareable for clustering.

- Section 13.2.1, “Using NSSMU to Initialize a Device,” on page 279
- Section 13.2.2, “Using NLVM Commands to Initialize a Device,” on page 279

13.2.1 Using NSSMU to Initialize a Device

1. Ensure that the SAN device is attached to all of the nodes in the cluster.
2. Log in as the `root` user to the master node of the cluster, then open a terminal console.
3. Launch NSSMU:

   ```
   nssmu
   ```

4 In the NSSMU main menu, select **Devices**, then press Enter.

5 In the **Devices** list, select the SAN device (such as `sdf`), then view information about it.

A device that has never been initialized reports a partition type of **Uninitialized**. If the device contains partitions or data, be prepared to lose all data on the device when it is initialized. The clustered volume group requires the entire device.

6 Press F3 to initialize the selected device.

You are prompted to confirm the initialization. If you continue, any existing partitions on the device are deleted, and a new partition table is created.

7 Read the advisory message, then do one of the following:

- Press y (Yes) to confirm that you want to initialize the device, and continue to Step 8.
- Press n (No) (or press Esc) to cancel. Return to Step 5 and choose a different device.

8 Specify the partitioning scheme to use as DOS or GPT, then press Enter.

DOS supports devices up to 2 TB in size. GPT supports devices of any size.

9 Select the device and view its details to verify that the device is initialized and unshared (that is, **Shareable for Clustering** is set to **No**.) If **Shareable for Clustering** is set to **Yes**, press F6 to unshare the device.

This ensures there are no partitions on the device, and the entire device is available to Clustered LVM (CLVM). In OES 11 SP2 and later, if the device is shared but it contains no other partitions, NSSMU and NLVM will automatically clear the shared setting and use the entire device.

10 Press Esc twice to exit NSSMU.

13.2.2 Using NLVM Commands to Initialize a Device

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Log in as the `root` user to the master node of the cluster, then open a terminal console.

3 At the command prompt, enter

```
nlvm [--force] [--no-prompt] init <device_name> format=<gpt|msdos>
unshared
```

Replace *device_name* with the device node name of the device to be initialized, such as `sdf`. This must be the first command option to follow `init`.

Specify the partitioning scheme as `gpt` or `msdos`. The default is `msdos`.

Specify the `unshared` option to ensure that the device is not marked as shareable for clustering.

You can specify the `--force` NLVM option to force the initialization.

You are automatically prompted to confirm the initialize action. Respond by pressing y (Yes) or n (No). You can use the `--no-prompt` NLVM option to suppress the confirmation.

For example, enter

```
nlvm init sdf format=gpt unshared
```

4 Confirm that the device was initialized and is unshared. At the command prompt, enter

```
nlvm list device <device_name>
```

Replace *device_name* with the device node name of the device that you initialized, such as `sdf`.

The command returns information about the device, such as its major:minor number, device size, partitioning format, and whether the device is shared.

For example, enter

```
nlvm list device sdf
```

```
avalon:~/Desktop # nlvm list device sdf
 Name=sdf
  Size=512.00MB(1048576) Used=32KB(64) Free=511.95MB(1048479)
  Format=GPT Shared=No RAID=No M:M=8:80 H:S=255:32
```

13.3 Configuring an LVM Volume Group Cluster Resource with NSS Management Tools

This section describes how to use the NSSMU utility and NLVM commands to create a clustered LVM volume group and logical volume on a single SAN device that is assigned to all nodes in the cluster. The volume is exclusively mounted on only one node at a time. Clustered LVM manages the locks for the exclusive mount.

After you create the resource, you should view and modify the cluster resource settings and scripts before you bring the resource online. You can also add lines to its load script, unload script, and monitor script to customize the resource for other uses.

If you enable NCP when you create the LVM volume group in NSSMU or with `nlvm` commands, commands are automatically added to the resource scripts to mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. An NCP Virtual Server object is created for the resource and commands are added to the scripts to bind or unbind it with the resource IP address.

- Section 13.3.1, “Sample Values,” on page 281
- Section 13.3.2, “Creating an LVM Volume Group Cluster Resource with NSSMU,” on page 282
- Section 13.3.3, “Creating an LVM Volume Group Cluster Resource with NLVM Commands,” on page 286
- Section 13.3.4, “Configuring the LVM Cluster Resource Settings,” on page 289
- Section 13.3.5, “Viewing or Modifying the LVM Resource Scripts,” on page 290
- Section 13.3.6, “Sample LVM Resource Scripts,” on page 291

13.3.1 Sample Values

The configuration scenarios use following sample values. Ensure that you replace the sample values with information for your configuration.

Parameter	Sample Value
Device name for the shared SAN device	`/dev/sdf` The device is initialized and contains no partitions. It is not enabled as shareable for clustering.

Parameter	Sample Value
Volume group name	`vol44` By default, NSSMU uses the logical volume name as the LVM volume group name. If you use the NLVM `create linux volume` command to create the LVM volume group cluster resource, you can specify a different name for the volume group, such as `vg44`.
Volume name	`vol44` If you enable NCP for the volume, this name must comply with the "Naming Conventions for NCP Volumes" on page 277. Lowercase letters are automatically changed to uppercase for the NCP volume name.
NCP volume name	`VOL44` The NCP volume is created only if you enable NCP as you create the clustered LVM volume by using NSSMU or the NLVM `create linux volume` command. The mount point path for the NCP volume is the mount point for the LVM volume. This means that the NCP share is created at the root of the mounted LVM volume when the LVM resource is online.
Linux POSIX file system type	ext3 Valid values are btrfs (requires the `btrfsprogs` package), ext2, ext3, reiserfs, and xfs.
Make options for the file system	None (do not specify a value). Press Enter to continue. For a list of the supported file system options for the file system type you are making, see the `mkfs(8)` man page and the man page for the specific file system: `mkfs.btrfs(8)`, `mkfs.ext2(8)`, `mkfs.ext3(8)`, `mkfs.reiserfs(8)`, or `mkfs.xfs(8)`. **IMPORTANT:** The file system creation fails if you specify a make option that is not supported by the file system type.
Mount options	rw The Read/Write (rw) option is specified by default. For a list of available options that work with the file system type you are using, see the `mount(8)` man page.
Volume size	100 GB A 100 GB LUN is prepared in the shared storage subsystem. It is attached to the nodes in the cluster. The device must be initialized and contain no partitions. It should not be marked as Shareable for clustering. You are not prompted to enter a volume size. The clustered LVM volume group and logical volume use the entire device. When you select the device, all of the device's free available space is displayed in the **Free Size** field.

Parameter	Sample Value
Resource IP address	10.10.10.44 This is the IP address of the virtual server for the cluster resource. The address must be unique and in the same subnet as the cluster's IP address. Specify the IP address in IPv4 format.

Mount device	`/dev/vol44/vol44`
	
	The mount device path format is
	
	`/dev/<volume_group_name>/<logical_volume_name>`
	
	If you specify a different name for the volume group, such as `vg44`, by using the NLVM `create linux volume` command, the mount device path is `/dev/vg44/vol44`.

Mount point path	`/mnt/vol44`
	
	You can use any valid Linux path as the mount point. The default mount point location for LVM logical volumes created by NSSMU and NLVM is `/usr/novell/<lx_volume_name>`.
	
	NSSMU automatically creates the mount point path if it does not exist on this node. However, you must manually create the mount point on each of the other nodes in the cluster. If the path does not exist on a node when you fail over the resource to it, the resource goes comatose You can alternatively add the following line to the load script before the Linux file system mount command to create the path on a node if it does not exist:
	
	`ignore_error mkdir -p $MOUNT_PATH`

13.3.2 Creating an LVM Volume Group Cluster Resource with NSSMU

This section describes how to use NSSMU to create and cluster-enable an LVM volume group. NSSMU automatically uses the same script format as the Generic File System template (`Generic_FS_Template`) to create an LVM volume group cluster resource.

1. Ensure that the SAN device is attached to all of the nodes in the cluster.
2. Log in as the `root` user to the master node of the cluster, then open a terminal console.
3. If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.

   LVM allows only one device per volume group. See Section 13.2, "Initializing a SAN Device," on page 278.
4. Launch NSSMU:

   ```
   nssmu
   ```

5 In the NSSMU main menu, select **Linux Volumes**, then press Enter.

6 On the Linux Volumes page, press **Insert** to launch the volume creation wizard, then enter the following information as you are prompted for it:

Parameter	Action
Select LVM type	Select **Cluster Enabled LVM2 Volume**, then press Enter.
NCP Enable volume?	Specify whether to enable the LVM logical volume for NCP file access. Press y (yes) to allow NCP file access, or press n (No) to disallow NCP file access at this time. If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the resource. If you do not enable NCP at this time, you can manually create an NCP virtual server, create an NCP Volume object, and add NCP commands to the scripts. See Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308.
Enter volume name	Type the name for the LVM logical volume , then press Enter. The specified name is also used for the LVM volume group, which is used in the name of the LVM volume group cluster resource. If you enable NCP for the volume, the specified name is also applied to the NCP volume. An NCP volume name can be up to 14 alphanumeric characters. Underscores are allowed. All letters in the LVM volume name are uppercase for the NCP volume name. For example, if the LVM volume name is `vol44`, the NCP volume name is VOL44.
Enter volume IP address	Type the IP address to use for the LVM cluster resource in IPv4 format (such as 10.10.10.44), then press Enter.
Select Volume type	Select one of the following the Linux POSIX file systems, then press Enter: ◆ btrfs This option is displayed only if the `btrfsprogs` package is installed. ◆ ext2 ◆ ext3 ◆ resiserfs ◆ xfs
Enter full mount point path	Type the full mount point path for the LVM logical volume (such as `/mnt/vol44`), then press Enter. The default path is `/usr/novell/<lx_volume_name>`, such as `/usr/novell/vol44`. If NCP is enabled, the specified path is also used as the mount point path for the NCP volume.

Parameter	Action
Enter any make options	Press Enter to continue without specifying options, or specify the desired make options for the file system type you are making, then press Enter. For a list of the supported file system options for the file system type you are making, see the `mkfs(8)` man page and the man page for the specific file system: `mkfs.btrfs(8)`, `mkfs.ext2(8)`, `mkfs.ext3(8)`, `mkfs.reiserfs(8)`, or `mkfs.xfs(8)`. **IMPORTANT:** The file system creation fails if you specify a make option that is not supported by the specified file system type.
Enter any mount options	Press Enter to accept the default read and write options (rw). The Read/Write (rw) option is specified by default. You can add other mount options, or modify this option. For a list of available options that work with the specified file system type, see the `mount(8)` man page.
Device	From the list of available devices, select the device that you initialized in Section 13.2, “Initializing a SAN Device,” on page 278, then press Insert or Enter to select the device. The LVM volume group requires the entire device. You cannot specify the amount of space to use. The device’s free available space is displayed in the Free Size field.

7 Press F3 to accept the setup you entered for the volume group cluster resource.

The resource is created and brought online on the node where you created it. The resource is named <`volume_group_name`>`_resource`. In the example, the name of the volume group is the same as the logical volume, so the resource name is `vol44_resource`.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

8 In the Linux Volumes list, select the newly created volume and view information about it.

Parameter	Description
Device	Specifies the full device node path for the LVM logical volume. **Example:** `/dev/vol44/vol44`
Mount Point	When the resource is online, this specifies the path on the root file system where this volume is mounted. **Example:** `/mnt/vol44`
Mount options	When the resource is online, this specifies the mount options that are applied whenever this volume is automatically mounted after a reboot. **Example:** rw
Type	When the resource is online, this specifies the file system type. **Examples:** btrfs, ext2, ext3, reiserfs, xfs

Parameter	Description
Size	Specifies the amount of space reserved for this volume. **Example:** 99.58 GB
Mounted	Specifies whether the volume is mounted or unmounted. When the resource is brought online, the load script mounts the logical volume. **Value:** Yes or No
State	Specifies the availability for the file system. **Example:** Read/Write
LVM	Specifies whether the specified volume is an LVM logical volume. **Value:** Yes
NCP	Specifies whether the LVM logical volume is enabled for NCP. The NCP state cannot be determined for a clustered LVM volume. You can use the Clusters plug-in for iManager to determine if there are NCP commands in the load script. On the Cluster Options page, select the resource to view its properties, then click the **Scripts** tab.
Share State	Specifies whether the LVM logical volume is cluster enabled for a OES Cluster Services cluster. **Value:** Shareable for Clustering When an LVM logical volume group is clustered, CLVM manages the share state, not the device. NSSMU reports the device as Not Shareable for Clustering. NSSMU reports the clustered LVM volume is Shareable for Clustering.

9 Press Escape twice to exit NSSMU.

10 Continue with Section 13.3.4, “Configuring the LVM Cluster Resource Settings,” on page 289.

13.3.3 Creating an LVM Volume Group Cluster Resource with NLVM Commands

This section describes how to use NLVM commands to create and cluster-enable an LVM volume group. NLVM automatically uses the same script format as the Generic File System template (`Generic_FS_Template`) to create an LVM volume group cluster resource. The NLVM command allows you to specify a group name that is different than the volume name.

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Log in to the master node of the cluster as the `root` user, then open a terminal console.

3 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.

See Section 13.2, “Initializing a SAN Device,” on page 278.

4 Create a clustered LVM volume group and logical volume. At the command prompt, enter the following (all on the same line, of course):

```
nlvm [nlvm_options] create linux volume
     type=<btrfs|ext2|ext3|reiserfs|xfs>
     device=<device_name>
     [mp=</mount_path>]
     [mkopt=<option1[,option2,...]>]
     [mntopt=<option1[option2[]...]>]
     lvm
     name=<lvm_volume_name>
     [group=<lvm_group_name>]
     shared
     ip=<IP_address_for_LVM_volgroup_cluster_resource>
     [ncp]
```

For details about using this command, see "Create Linux Volume" in the *OES 2023: NLVM Reference*.

Parameters and Options

type=<btrfs|ext2|ext3|reiserfs|xfs>

You must specify the type of file system to use for the volume. Btrfs requires that the `btrfsprogs` package is installed on all nodes. Supported file systems are btrfs, ext2, ext3, reiserfs, and xfs.

device=<*device_name*>

You must specify the device to use. Replace *device_name* with the device node name of the device that you want to use for the volume, such as `sdf`. The entire device is used for the LVM volume group. You cannot specify a size to use. The device must already be initialized, contain no partitions, and must not be marked as shareable for clustering.

[mp=</*mount_path*>]

If a mount point path is not specified, the utility assigns a default mount path of `/usr/novell/<volume_name>`. Use the `mp` option to specify a custom mount point. Replace *mount_path* with the full Linux path of the mount point where the volume is to be mounted. The final directory name in the path can be the same or different than the specified volume name. If the path does not currently exist, it is automatically created on that node. You must manually create the path on other nodes in the cluster.

[mkopt=<option1[,option2,...]>]

You can use the `mkopt` option to specify the options to use when running `mkfs`. For a list of available options, see the `mkfs(8)` man page. No default option is specified. For a list of the supported file system options for the file system type you are making, see the `mkfs(8)` man page and the man page for the specific file system: `mkfs.btrfs(8)`, `mkfs.ext2(8)`, `mkfs.ext3(8)`, `mkfs.reiserfs(8)`, or `mkfs.xfs(8)`.

IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.

[mntopt=<option1[option2[]...]>]

You can use the `mntopt` option to specify the options to use when mounting the volume. For a list of available options, see the `mount(8)` man page. The default `mntopt` value is `rw` (Read/Write).

lvm

You must specify the `lvm` option to create an LVM volume group and logical volume. Use this option with the `name` option.

name=<*lvm_volume_name*>

Replace *lvm_volume_name* with a name for the LVM volume. If you do not specify the `group` option, this name is also used as the LVM volume group name, which is used in the cluster resource name. For LVM logical volume naming conventions, see "Linux LVM Volume Group and Logical Volume Names" in the *OES 2023: NLVM Reference*.

If you use the NCP option, the NCP volume name uses the same name as the LVM logical volume name, but all letters are uppercase. NCP volume names can be up to 14 alphanumeric characters, and underscores are allowed.

[group=<*lvm_group_name*>]

Replace *lvm_volume_group_name* with a name for the LVM volume group. The group name is also used in the cluster resource name. If you do not specify a volume group name, the group is automatically named the same as the LVM volume.

shared

You must specify the `shared` option to create a clustered LVM volume group and logical volume.

ip=<*IP_address_for_LVM_volgroup_cluster_resource*>

Replace *IP_address_for_LVM_volgroup_cluster_resource* with a static unique IP address to use for the LVM cluster resource. Specify the address in IPv4 format.

ncp

Specify the `ncp` option to NCP enable the LVM logical volume for NCP file access.

- If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the resource.
- If you do not enable NCP at this time, you can manually create a virtual server and add NCP file access later. See Section 13.5, "Creating a Virtual Server Object for an LVM Volume Group Cluster Resource," on page 308.

Example For example, at the command prompt, enter the following (all on the same line):

```
nlvm create linux volume
     type=ext3
     device=sdf
     mp=/mnt/vol44
     mntopt=rw
     lvm
     name=vol44
     group=vg44
     shared
     ip=10.10.10.44
     ncp
```

If the command is successful, the response is

```
Linux clustered volume vol44 created.
```

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

5 Verify that the cluster resource was created and brought online by entering

```
cluster status
```

The resource is named <*lv_name*>_resource. In the following example, `vol44_resource` is in the **Running** state.

```
avalon:~/Desktop # cluster status
Master_IP_Address_Resource                Running    avalon                1
lvmvol134_resource                        Running    avalon                1
vol44 resource                            Running    avalon                1
```

6 View the device information to see that the device shows the LVM format. At the command prompt, enter

```
nlvm list device <device_name>
```

Replace *device_name* with the name of the device that you used, such as `sdf`.

The command shows that the LVM container uses the entire device; the size and used space values are the same and the free space shows no space. The partitioning format is LVM.

For example, enter

```
nlvm list device sdf
```

```
avalon:~/Desktop # nlvm list device sdf
 Name=sdf
  Size=512.00MB(1048576) Used=512.00MB(1048576) Free=0KB(0)
  Format=LVM Shared=No RAID=No M:M=8:64 H:S=255:32
```

7 Continue with Section 13.3.4, “Configuring the LVM Cluster Resource Settings,” on page 289.

13.3.4 Configuring the LVM Cluster Resource Settings

After you create a clustered LVM volume group and logical volume, use the procedure in this section to verify that the LVM volume group cluster resource was created and is online. You can customize the resource policies, monitoring, and preferred nodes settings.

1 Open iManager in a web browser, then log in as a cluster administrator user.

2 In **Roles and Tasks**, select **Clusters > My Clusters**.

3 Select the cluster where you created the volume group cluster resource.

If the cluster you want to manage is not in the list, you can add it. Click **Add**, browse to locate and select the Cluster object of the cluster, then click **OK**.

4 In the list of resources, locate the new resource and notice the state of the resource. It should be online and running.

5 Click the resource’s name link to open its Properties page.

You can alternatively go to the Cluster Options page, then click the resource’s name link, or select the check box next to the resource and click **Details**.

6 On the Resource Policies page, view and modify the resource's Policy settings if needed. Click **Apply** if you make changes.

6a (Optional) Select the **Resource Follows Master** check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to the node that becomes the master.

6b (Optional) Select the **Ignore Quorum** check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed OES Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

Selecting this box ensures that the resource is launched immediately on any server in the Preferred Nodes list as soon as any server in the list is brought online.

6c By default, the Generic File System resource template sets the Start mode and Failover mode to **Auto** and the Failback Mode to **Disable**. You can change the default settings as needed.

- **Start Mode:** If the Start mode is set to **Auto**, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to **Manual**, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.
- **Failover Mode:** If the Failover mode is set to **Auto**, the resource automatically moves to the next server in the Preferred Nodes list if there is a hardware or software failure. If the Failover mode is set to **Manual**, you can intervene after a failure occurs and before the resource is started on another node.
- **Failback Mode:** If the Failback mode is set to **Disable**, the resource continues running on the node it has failed to. If the Failback mode is set to **Auto**, the resource automatically moves back to its preferred node when the preferred node is brought back online. Set the Failback mode to **Manual** to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

7 Enable and configure monitoring for the resource, then click **Apply**.

See Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 176.

7a In the Properties page, select the **Monitoring** tab.

7b Select the **Enable Resource Monitoring** check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

7c Specify the **Polling Interval** to control how often you want the resource monitor script for this resource to run.

You can specify the value in minutes or seconds. See "Polling Interval" on page 177.

7d Specify the number of failures (**Maximum Local Failures**) for the specified amount of time (**Time Interval**).

See "Failure Rate" on page 177.

7e Specify the **Failover Action** by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

See "Failure Action" on page 177.

8 Click the **Preferred Nodes** tab, assign preferred nodes for the resource by moving them from the **Unassigned** list to the **Assigned** list, then click **Apply**.

When you configure a volume group cluster resource with NSSMU or with NLVM commands, the node where you create it is automatically assigned as a preferred node for the resource.

IMPORTANT: For **Assigned** nodes, ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list in a text editor. In the editor, click **OK** to close the editor, then click **Apply** to save your changes.

9 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online again.

10 Continue with Section 13.3.5, "Viewing or Modifying the LVM Resource Scripts," on page 290.

13.3.5 Viewing or Modifying the LVM Resource Scripts

You can customize the scripts by adding lines for other products that use a shared LVM volume group resource. Compare the generic script with the templates for those products to identify what lines need to be added or modified.

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster where you created the volume group cluster resource.

3 Select the cluster resource's name link to open the Properties page, then click the **Scripts** tab.

The **Scripts** tab automatically opens to the load script.

4 On the **Load Script** page, view or modify the load script. Click **Apply** if you make changes.

Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the "Sample LVM Resource Load Scripts Created by NSS Management Tools" on page 291.

5 Click the **Unload Script** link to view or modify the unload script. Click **Apply** if you make changes.

Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the "Sample LVM Resource Unload Scripts Created by NSS Management Tools" on page 293.

6 Click the **Monitor Script** link to view or modify the monitor script. Click **Apply** if you make changes.

Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the "Sample LVM Resource Monitor Scripts Created by NSS Management Tools" on page 294.

7 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online again.

8 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect.

8a In Roles and Tasks, select **Clusters > My Clusters**.

8b On the Cluster Manager page, select the check box next to the resource, then click **Offline**.

Wait for the status to report that it is offline, then continue.

8c Select the check box next to the resource, then click **Online**.

8d Verify that the resource comes online and reports a **Running** state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

13.3.6 Sample LVM Resource Scripts

The sample LVM resource scripts in this section are automatically generated for an LVM volume group cluster resource when you use NLVM or NSSMU to create a clustered LVM2 volume, as described in Section 13.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 280. See Section 13.3.1, “Sample Values,” on page 281 for information about the sample values used in these scripts. Ensure that you replace sample values with those for your own system.

- “Sample LVM Resource Load Scripts Created by NSS Management Tools” on page 291
- “Sample LVM Resource Unload Scripts Created by NSS Management Tools” on page 293
- “Sample LVM Resource Monitor Scripts Created by NSS Management Tools” on page 294

Sample LVM Resource Load Scripts Created by NSS Management Tools

Compare the load scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- “Without NCP File Access” on page 292
- “With NCP File Access” on page 292

Without NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

With NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP virtual server to the resource IP address
exit_on_error ncpcon bind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Scripts Created by NSS Management Tools

Compare the unload scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- “Without NCP File Access” on page 293
- “With NCP File Access” on page 293

Without NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

With NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# unbind the NCP virtual server from the resource IP address
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

```
# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample LVM Resource Monitor Scripts Created by NSS Management Tools

Compare the monitor scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access. To use the script, you must also enable monitoring for the resource. See Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

- “Without NCP File Access” on page 294
- “With NCP File Access” on page 295

Without NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

With NCP File Access

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
```

```
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# check the LVM logical volume
exit_on_error status_lv $MOUNT_DEV

# check the NCP volume
exit_on_error ncpcon volume $NCP_VOLUME

# check the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

13.4 Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template

This section describes how to use Linux Logical Volume Manager (LVM) commands to create a shared LVM volume group. You use the Generic File System template (`Generic_FS_Template`) to create a volume group cluster resource that cluster-enables the exiting LVM volume group.

After you create the resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the `Generic_FS_Template` to the resource template for your product to determine which lines need to be added or modified.

- Section 13.4.1, “Creating a Shared LVM Volume with LVM Commands,” on page 295
- Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 300
- Section 13.4.3, “Sample Generic LVM Resource Scripts,” on page 305

13.4.1 Creating a Shared LVM Volume with LVM Commands

You can create a shared LVM volume group and logical volume by using native Linux LVM commands.

- “Sample Values for the LVM Volume Group and Logical Volume” on page 296
- “Creating a Shared LVM Volume Group (Quick Reference)” on page 296
- “Creating a Shared LVM Volume Group (Detailed Instructions)” on page 297

Sample Values for the LVM Volume Group and Logical Volume

The procedures in this section uses the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
LVM physical volume	`/dev/sdd` For clustering, we recommend only one device in the LVM volume group.
LVM volume group name	`clustervg01`
LVM logical volume	`clustervol01`
File system type	`ext3` This is the file system type that you make on the LVM logical volume, such as `btrfs`, `ext2`, `ext3`, `reiserfs`, or `xfs`. For information about `mkfs` command options for each file system type, see the `mkfs(8)` man page.
Logical volume path	`/dev/clustervg01/clustervol01`
Mount point for the logical volume	`/mnt/clustervol01`

Creating a Shared LVM Volume Group (Quick Reference)

You can create the volume group and logical volume by issuing the following LVM commands as the `root` user on the cluster node.

Before you begin, initialize the SAN device as described in Section 13.2, “Initializing a SAN Device,” on page 278.

When you are done, continue with Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 300.

Action	Command
1. Create the LVM physical volume.	`pvcreate <device>`
2. Create the clustered LVM volume group.	`vgcreate -c y <vg_name> <device>`
3. Activate the volume group exclusively on the node.	`vgchange -a ey <vg_name>`
4. Create the LVM logical volume.	`lvcreate -n <lv_name> -L size <vg_name>`
5. Add a file system to the LVM logical volume.	`mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]`
6. If it does not exist, create the mount point path.	`mkdir -p <full_mount_point_path>`

Action	Command
7. Mount the LVM logical volume.	`mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point>`

8. Dismount the LVM logical volume and deactivate the volume group.
 	
After you create the Generic LVM cluster resource, use the cluster resource to control when the volume is online or offline.
	`umount /dev/<vg_name>/<lv_name>`
	
	`vgchange -a n <vg_name>`

Creating a Shared LVM Volume Group (Detailed Instructions)

For detailed instructions, use the following procedure to create the LVM volume group and logical volume:

1. Log in as the Linux `root` user to the first node of the cluster, then open a terminal console.
2. Initialize the SAN device as described in Section 13.2, “Initializing a SAN Device,” on page 278.

   LVM allows only one device per volume group.
3. Create an LVM physical volume on the device (such as `/dev/sdd`) by entering:

   `pvcreate <device>`

   For example:

```
pvcreate /dev/sdd
  No physical volume label read from /dev/sdd
  Physical volume "/dev/sdd" successfully created
```

4. Display information about the physical volume by entering:

   `pvdisplay [device]`

   For example, to view information about `/dev/sdd`:

```
pvdisplay /dev/sdd
"/dev/sdd" is a new physical volume of "512 MB"
  --- NEW Physical volume ---
  PV Name               /dev/sdd
  VG Name
  PV Size               512 MB
  Allocatable           NO
  PE Size (KByte)       0
  Total PE              0
  Free PE               0
  Allocated PE          0
  PV UUID               dg5WG9-IJvq-MkEM-iRrC-hMKv-Zmwq-YY0N3m
```

5. Create an LVM volume group (such as `clustervg01`) on the physical volume by entering:

   `vgcreate -c y <vg_name> <device>`

   For example:

```
vgcreate -c y "clustervg01" /dev/sdd
  Clustered volume group "clustervg01" successfully created
```

The volume group is automatically activated.

6 Activate the volume group exclusively on the current server by entering:

`vgchange -a ey <vg_name>`

The -a option activates the volume. The `ey` parameter specifies the values `exclusively` and `yes`.

For example:

```
vgchange -a ey clustervg01
  logical volume(s) in volume group "clustervg01" now active
```

7 View information about the volume group by using the `vgdisplay` command.

`vgdisplay <vg_name>`

Notice that 4 MB of the device is used for the volume group's Physical Extent (PE) table. You must consider this reduction in available space on the volume group when you specify the size of the LVM logical volume in the next step (Step 8).

For example:

```
vgdisplay clustervg01
  --- Volume group ---
  VG Name               clustervg01
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  1
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                0
  Open LV               0
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               508.00 MB
  PE Size               4.00 MB
  Total PE              127
  Alloc PE / Size       0 / 0
  Free  PE / Size       127 / 508.00 MB
  VG UUID               rqyAd3-U2dg-HYLw-0SyN-loO7-jBH3-qHvySe
```

8 Create an LVM logical volume (such as `clustervol01`) on the volume group by entering:

`lvcreate -n <lv_name> -L size <vg_name>`

Specify the logical volume name, size, and the name of the volume group where you want to create it. The size is specified in megabytes by default.

The logical volume full path name is `/dev/<vg_name>/<lv_name>`.

For example:

```
lvcreate -n "clustervol01" -L 500 "clustervg01"
  Logical volume "clustervol01" created
```

This volume's full path name is `/dev/clustervg01/clustervol01`.

9 View information about the logical volume by entering:

`lvdisplay -v /dev/<vg_name>/<lv_name>`

For example:

```
lvdisplay -v /dev/clustervg01/clustervol01
    Using logical volume(s) on command line
  --- Logical volume ---
  LV Name                /dev/clustervg01/clustervol01
  VG Name                clustervg01
  LV UUID                nIfsMp-alRR-i4Lw-Wwdt-v5io-2hDN-qrWTLH
  LV Write Access        read/write
  LV Status              available
  # open                 0
  LV Size                500.00 MB
  Current LE             125
  Segments               1
  Allocation             inherit
  Read ahead sectors     auto
  - currently set to     1024
  Block device           253:1
```

10 Create a file system (such as Btrfs, Ext2, Ext3, ReiserFS, or XFS) on the LVM logical volume by entering:

`mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]`

For example:

```
mkfs -t ext3 /dev/clustervg01/clustervol01
mke2fs 1.41.9 (22-Aug-2009)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
128016 inodes, 512000 blocks
25600 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=67633152
63 block groups
8192 blocks per group, 8192 fragments per group
2032 inodes per group
Superblock backups stored on blocks:
    8193, 24577, 40961, 57345, 73729, 204801, 221185, 401409

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 29 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

11 If the mount point does not exist, create the full directory path for the mount point.

`mkdir -p <full_mount_point_path>`

For example:

```
mkdir -p /mnt/clustervol01
```

12 Mount the logical volume on the desired mount point by entering:

`mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point>`

For example:

```
mount -t ext3 /dev/clustervg01/clustervol01 /mnt/clustervol01
```

13 Dismount the volume and deactivate the LVM volume group by entering:

```
umount /dev/<vg_name>/<lv_name>
vgchange -a n "vg_name"
```

This allows you to use the load script and unload script to control when the volume group is activated or deactivated in the cluster.

For example, to unmount the volume and deactivate the `clustervg01` volume group, enter

```
umount /dev/clustervg01/clustervol01
vgchange -a n "clustervg01"
  0 logical volume(s) in volume group "clustervg01" now active
```

14 Continue with Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 300.

13.4.2 Creating a Generic File System Cluster Resource for an LVM Volume Group

You can use the OES Cluster Services Generic File System template (`Generic_FS_Template`) to create a cluster resource for the LVM volume group.

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Ensure that the existing LVM volume you want to cluster-enable is deactive on all nodes in the cluster.

See Step 13 in Section 13.4.1, “Creating a Shared LVM Volume with LVM Commands,” on page 295.

3 In iManager, select **Clusters > My Clusters.**

4 Select the cluster you want to manage, then click **Cluster Options**.

5 Click the **New** link.

6 Click the **Resource** radio button, then click **Next**.

7 Specify information to define your new cluster resource:

7a Specify the name of the resource you want to create, such as `clustervg01`.

Do not use periods in cluster resource names. Client for Open Enterprise Servers interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

7b In the **Inherit From Template** field, browse to select the **Generic_FS_Template**.

7c Deselect or select **Online Resource after Create**.

This option is deselected by default. Deselect this option to allow an opportunity to verify the scripts and resource policies before you bring the resource online for the first time.

Select this option to automatically online the resource on its most preferred node as soon as it is created and configured. The most preferred node is the first server in the assigned nodes list for the resource. You should also select the **Define Additional Properties** check box to configure the resource now.

7d Select the **Define Additional Properties** check box.

When this option is enabled, a wizard walks you through the load, unload, and monitor scripts, then allows you to set up the policies and preferred nodes for the cluster resource.

7e Click **Next**.

If you selected **Define Additional Properties**, continue with the setup. Otherwise, select the resource in the Cluster Objects list to open its properties page, and configure the resource.

8 Configure the load script for the resource by replacing the variables with your own values, specify the **Load Script Timeout** value, then click **Next**.

The following is the default Generic_FS template load script. The variable values must be specified before you can successfully bring the resource online. Be mindful that Linux path names are case sensitive. The mount path must already exist. You can add a `mkdir` command in the script for the mount point, or create the directory manually on each node before you allow the resource to fail over to other nodes. A sample script is available in "Sample Generic LVM Resource Load Script" on page 306.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
#define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

#activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

9 Configure the unload script for the resource by replacing the variables with your own values, specify the **Unload Script Timeout** value, then click **Next**.

The following is the default Generic_FS template unload script. The variable values must be specified before you bring the resource online. Be mindful that Linux path names are case sensitive. A sample script is available in "Sample Generic LVM Resource Unload Script" on page 307.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
#define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status

exit 0
```

10 Configure the monitor script for the resource by replacing the variables with your own values, specify the **Monitor Script Timeout** value, then click **Next**.

The following is the default Generic_FS template monitor script. Be mindful that Linux path names are case sensitive. The variable values must be specified before you bring the resource online. A sample script is available in "Sample Generic LVM Resource Monitor Script" on page 308.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
#define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

#check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

11 On the Resource Policies page, view and modify the resource's Policy settings:

11a (Optional) Select the **Resource Follows Master** check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to the node that becomes the master.

11b (Optional) Select the **Ignore Quorum** check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed OES Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

Selecting this box ensures that the resource is launched immediately on any server in the Preferred Nodes list as soon as any server in the list is brought online.

11c (Optional) By default, the Generic File System resource template sets the Start mode and Failover mode to **Auto** and the Failback Mode to **Disable**. You can change the default settings as needed.

- **Start Mode:** If the Start mode is set to **Auto**, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to **Manual**, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

- **Failover Mode:** If the Failover mode is set to **Auto**, the resource automatically moves to the next server in the Preferred Nodes list if there is a hardware or software failure. If the Failover mode is set to **Manual**, you can intervene after a failure occurs and before the resource is started on another node.
- **Failback Mode:** If the Failback mode is set to **Disable**, the resource continues running on the node it has failed to. If the Failback mode is set to **Auto**, the resource automatically moves back to its preferred node when the preferred node is brought back online. Set the Failback mode to **Manual** to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

11d Click **Next**.

12 On the Preferred Nodes page, assign preferred nodes for the resource by moving them from the **Unassigned** list to the **Assigned** list.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list a text editor. In the editor, click **OK** to close the editor.

13 Click **Finish**.

Typically, the resource creation takes less than 10 seconds. However, if you have a large tree or if the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

14 Verify that the resource was created by viewing its entry in the **Cluster Objects** list on the Cluster Options page.

15 (Optional) Enable and configure resource monitoring:

Resource monitoring is disabled by default. Enable the monitoring function and modify the settings for the resource. For detailed information, see Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

15a In iManager on the Cluster Manager page or Cluster Options page, select the resource name link to open its Properties dialog box, then select the **Monitoring** tab.

15b Select the **Enable Resource Monitoring** check box to enable resource monitoring for the selected resource.

15c Specify the **Polling Interval** to control how often you want the resource monitor script for this resource to run.

You can specify the value in minutes or seconds. See “Polling Interval” on page 177.

15d Specify the number of failures (**Maximum Local Failures**) for the specified amount of time (**Time Interval**).

See “Failure Rate” on page 177.

15e Specify the **Failover Action** by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

See “Failure Action” on page 177.

16 Bring the resource online:

16a Select **Clusters > My Clusters**.

16b On the Cluster Manager page, select the check box next to the new CSM cluster resource, then click **Online**.

Ensure that the resource successfully enters the **Running** state. If the resource goes comatose instead, it is probably because you made an error when typing values in the script definitions. Typical issues are that the target LVM volume is still active locally on a node, the path names were not entered as case sensitive, the IP address is not unique in the network, or the mount path does not exist.

Take the resource offline, go to the resource's **Properties > Scripts** page to review and modify its scripts as needed to correct errors, then try again to online the resource.

13.4.3 Sample Generic LVM Resource Scripts

This section contains sample scripts for the Generic LVM resource.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	`10.10.10.136`
MOUNT_FS	`ext3`
VOLGROUP_NAME	`clustervg01`
MOUNT_DEV	`/dev/$VOLGROUP_NAME/clustervol01`
MOUNT_POINT	`/mnt/clustervol01`

- "Sample Generic LVM Resource Load Script" on page 306
- "Sample Generic LVM Resource Unload Script" on page 307
- "Sample Generic LVM Resource Monitor Script" on page 308

Sample Generic LVM Resource Load Script

Use the following sample load script to complete the fields for your LVM volume group cluster resource:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

Sample Generic LVM Resource Unload Script

Use the following sample unload script to complete the fields for your LVM volume group cluster resource:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample Generic LVM Resource Monitor Script

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource. To use the script, you must also enable monitoring for the resource. See Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

13.5 Creating a Virtual Server Object for an LVM Volume Group Cluster Resource

You can create an NCP virtual server for an existing clustered LVM volume group cluster resource by using the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create an `NCS:NCP Server` object. Having a virtual server object allows the resource to be viewed in the **Browse** panel in iManager. You can also bind the virtual server name to the IP address to allow users to access the resource via an assigned name in addition to its IP address.

The `NCS:NCP Server` object does not give users NCP file access to the data on the Linux LVM logical volume. You can add commands to the load script to create an NCP volume on the Linux file system for NCP users. See “NCP Enabling a clustered LVM volume”.

The virtual server name is stored in eDirectory in an `NCS:NCP Server` object under the Cluster object where you created the resource. You must add a line to the load and unload scripts that identifies the name of this virtual server and a line that binds or unbinds the name to the IP address of the Linux POSIX cluster resource.

- Section 13.5.1, “Creating an NCP Virtual Server for the LVM Resource,” on page 309
- Section 13.5.2, “Adding NCP Virtual Server Commands to the Resource Scripts,” on page 309
- Section 13.5.3, “Sample LVM Resource Scripts with an NCP Virtual Server,” on page 311

13.5.1 Creating an NCP Virtual Server for the LVM Resource

You use the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create a virtual server object (`NCS:NCP Server`) in eDirectory for the LVM volume group cluster resource. If the resource does not have NCP volumes on it, you do not use the -v option. For information about the `ncs_ncpser.py` script, see Section A.9, "ncs_ncpserv.py Script (Creating an NCP Virtual Server Object for a Clustered LVM Volume Group)," on page 399.

1. On the master cluster node, open a terminal console, then log in as the `root` user.
2. In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.
3. At the command prompt, enter

   ```
   ./ncs_ncpserv.py -c vg_name -i resource_ip_address
   ```

   Replace the *vg_name* and *resource_ip_address* with the information for your particular solution.

   Do not use periods in cluster resource names. Client for Open Enterprise Servers interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

   For example, to create the NCS:NCP Server object for the `clustervg01` cluster resource where the IP address is 10.10.10.44 and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

   ```
   ./ncs_ncpserv.py -c clustervg01 -i 10.10.10.44
   ```

   The confirmation message is displayed:

   ```
   NCP Server 'cn=cluster1-clustervg01-
   server,ou=clusters,ou=city,o=mycompany' created.
   ```

4. Continue with Section 13.5.2, "Adding NCP Virtual Server Commands to the Resource Scripts," on page 309.

13.5.2 Adding NCP Virtual Server Commands to the Resource Scripts

After you have created an NCS:NCP Server object, you must modify the resource's load, unload, and monitor scripts to bind and unbind the NCS:NCP Server object with the resource IP address.

1. In iManager, select **Clusters > My Clusters**, then select the cluster.
2. Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page, then click the **Scripts** tab.
3. Modify the LVM resource's load script:
   - **3a** On the **Scripts** tab, click **Load Script**.
   - **3b** At the end of the definition area, add the following lines to define the virtual NCP server name:

     ```
     # define NCP server name
     NCP_SERVER=cluster1-clustervg01-server
     ```

     Replace the NCP server name with the name for your virtual NCP server.

3c Under the `add_secondary_ipaddress` command line, add a line to bind the NCP server name to the resource IP address:

```
# bind the NCP server name
exit_on_error ncpcon bind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

3d Verify the order of the commands against the "Sample LVM Resource Load Script with an NCP Virtual Server" on page 312.

3e Click **Apply** to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

4 Modify the LVM resource's unload script:

4a On the **Scripts** tab, click **Unload Script**.

4b At the end of the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1-clustervg01-server
```

Specify the same NCP virtual server name that you used in the load script.

4c Under the definition, add a line to unbind the NCP server name from the resource IP address:

```
# unbind the NCP server name
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

4d Verify the order of the commands against the "Sample LVM Resource Unload Script with an NCP Virtual Server" on page 313.

4e Click **Apply** to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

5 Modify the LVM resource's monitor script:

5a On the **Scripts** tab, click **Monitor Script**.

5b At the end of the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1-clustervg01-server
```

Specify the same NCP virtual server name that you used in the load script.

5c Verify the order of the commands against the "Sample LVM Resource Monitor Script with an NCP Virtual Server" on page 314.

5d Click **Apply** to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

6 At the bottom of the Properties dialog box, click **OK** to save your changes.

7 Activate the modified scripts by taking the resource offline and bringing it online.

7a Click the **Cluster Manager** tab.

7b Select the check box next to the LVM volume group cluster resource, then click **Offline**.

Wait until the resource is reports an Offline status before continuing.

7c Select the check box next to the LVM volume group cluster resource, then click **Online**.

Wait until the resource is reports an Online status before continuing.

8 Verify that an NCS:NCP `Server` object appears in the **Browse** panel in iManager.

8a In the iManager toolbar, click the **View Objects** icon.

8b In the left panel, click **Browse**.

8c Browse to the Cluster container to see the virtual server object for the cluster resource, such as `cluster1-clustervg01-server`.

9 (Optional) To add NCP file access for the Linux volume, continue with Section 13.6, "Enabling NCP File Access for a Clustered LVM Volume," on page 314

13.5.3 Sample LVM Resource Scripts with an NCP Virtual Server

This section contains sample scripts for the Generic LVM resource where commands have been added to bind and unbind an NCP Virtual Server object with the resource IP address.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	`10.10.10.136`
MOUNT_FS	`ext3`
VOLGROUP_NAME	`clustervg01`
MOUNT_DEV	`/dev/$VOLGROUP_NAME/clustervol01`
MOUNT_POINT	`/mnt/clustervol01`
NCP_SERVER	`clus1-clustervg01-server`

- "Sample LVM Resource Load Script with an NCP Virtual Server" on page 312
- "Sample LVM Resource Unload Script with an NCP Virtual Server" on page 313
- "Sample LVM Resource Monitor Script with an NCP Virtual Server" on page 314

Sample LVM Resource Load Script with an NCP Virtual Server

Use the following sample load script to complete the fields for your LVM volume group cluster resource:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# define NCP server name
NCP_SERVER=cluster1-clustervg01-server

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP server name
exit_on_error ncpcon bind --ncpservername=$NCP_SERVER --
ipaddress=$RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Script with an NCP Virtual Server

Use the following sample unload script to complete the fields for your LVM volume group cluster resource:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# define NCP server name
NCP_SERVER=cluster1-clustervg01-server

# unbind the NCP server name
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --
ipaddress=$RESOURCE_IP

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample LVM Resource Monitor Script with an NCP Virtual Server

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource. To use the script, you must also enable monitoring for the resource. See Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# define NCP server name
NCP_SERVER=cluster1-clustervg01-server

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# check the status of NCP services
rcndsd status
  if test $? != 0; then
     exit_on_error rcndsd restart
  fi
sleep 5

exit 0
```

13.6 Enabling NCP File Access for a Clustered LVM Volume

After you create an NCP virtual server object for the clustered LVM volume, you can enable NCP file access to a clustered LVM volume by adding NCP commands to the load, unload, and monitor script.

- Section 13.6.1, “Adding NCP Volume Commands to the Resource Scripts,” on page 315
- Section 13.6.2, “Creating a Shared NCP Volume Object,” on page 317
- Section 13.6.3, “Sample LVM Resource Scripts Enabled for NCP File Access,” on page 318

13.6.1 Adding NCP Volume Commands to the Resource Scripts

After you have created an NCS:NCP Server object and the Volume object for the NCP volume, you can add NCP file access for users by adding NCP volume commands to the resource's load, unload, and monitor scripts.

1. In iManager, select **Clusters > My Clusters**, then select the cluster.
2. Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page, then click the **Scripts** tab.
3. Modify the LVM resource's load script:
   - **3a** On the **Scripts** tab, click **Load Script**.
   - **3b** At the end of the definition area, add the following lines to define the NCP volume name:

     ```
     # define NCP volume name
     NCP_VOLUME=CLUSTERVOL01
     ```

     Replace the NCP server name with the name for your virtual NCP server. NCP volume names can be up to 14 alphanumeric characters. Supported characters are A to Z, 0 to 9, and underscores (_).
   - **3c** After the `add_secondary_ipaddress` command line, add a line to mount the NCP volume:

     ```
     # mount the NCP volume
     exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT
     ```
   - **3d** Verify the order of the commands against the "Sample LVM Resource Load Script with an NCP Volume" on page 318.
   - **3e** Click **Apply** to save your changes.

     The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.
4. Modify the LVM resource's unload script:
   - **4a** On the **Scripts** tab, click **Unload Script**.
   - **4b** At the end of the definition area, add the following lines to define the NCP volume name:

     ```
     # define NCP volume name
     NCP_VOLUME=CLUSTERVOL01
     ```

     Specify the same volume name that you used in the load script.
   - **4c** Under the `ncpcon unbind` command, add a line to dismount the NCP volume:

     ```
     # dismount the NCP volume
     ignore_error ncpcon dismount $NCP_VOLUME
     ```
   - **4d** Verify the order of the commands against the "Sample LVM Resource Unload Script with an NCP Volume" on page 319.
   - **4e** Click **Apply** to save your changes.

     The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

5 Modify the LVM resource's monitor script:

5a On the **Scripts** tab, click **Monitor Script**.

5b At the end of the definition area, add the following lines to define the NCP volume name:

```
# define NCP volume name
NCP_VOLUME=CLUSTERVOL01
```

Specify the same volume name that you used in the load script.

5c Verify the order of the commands against the "Sample LVM Resource Monitor Script with an NCP Volume" on page 320.

5d Click **Apply** to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

6 At the bottom of the Properties dialog box, click **OK** to save your changes.

7 Activate the modified scripts by taking the resource offline and bringing it online.

7a Click the **Cluster Manager** tab.

7b Select the check box next to the LVM volume group cluster resource, then click **Offline**.

Wait until the resource is reports an Offline status before continuing.

7c Select the check box next to the LVM volume group cluster resource, then click **Online**.

Wait until the resource is reports an Online status before continuing.

8 Verify that a clustered NCP Volume object appears in the **Browse** panel in iManager.

8a In the iManager toolbar, click the **View Objects** icon.

8b In the left panel, click **Browse**.

8c Browse to the cluster container to see the NCP volume, such as `clus1_CLUSTERVG01`.

9 Verify the NCP volume by using the Manage NCP Services plug-in for OES Remote Manager:

9a Log in as the root user to OES Remote Manager.

9b Select **Manage NCP Services**, then select **Manage Shares**.

9c Locate the NCP volume in the **Active Shares** list.

The clustered NCP volume is defined in the LVM volume group cluster resource load script. The NCP volume is mounted when you bring the resource online, and dismounted when you take the resource offline. You do not use the **Mount** and **Unmount** buttons in this interface.

9d Click the **Information** (i) icon to the left of the NCP volume to view details about the NCP volume and mount point.

9e Exit OES Remote Manager.

13.6.2 Creating a Shared NCP Volume Object

Use the following procedure to create an NCP volume object on the NCP virtual server you created in Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 308

1. Log in to iManager as a cluster administrator.
2. Select **Directory Administration > Create Object**.
3. On the Create Object page, select Volume, then click OK.
4. On the Create Volume page, complete the following information, then click **OK**.

Parameter	Action
Volume name	Specify the name that you defined in the load script in the format <**cluster_name**>_<**ncp_volumen_name**>, such as `cluster1_CLUSTERVOL01`.
Host server	Browse to select the NCP virtual server object for the LVM cluster resource, such as `cluster1-clustervg01-server`.
Context	Browse to select the context where you want to create the Volume object for the NCP volume, such as `novell`.
Host resource name	From the drop-down list of available NCP volumes on the server, select the NCP volume that you defined in the load script, such as CLUSTERVOL01.

5. After the Volume object is successfully created, click **OK** to dismiss the confirmation message.
6. Verify that the object is accessible.
   - **6a** In the iManager toolbar, select the **View Objects** icon.
   - **6b** Browse in the Tree view to locate and view information about the NCP volume’s Volume object.

13.6.3 Sample LVM Resource Scripts Enabled for NCP File Access

This section contains sample scripts for the Generic LVM resource where commands have been added to bind and unbind an NCP Virtual Server object with the resource IP address and to mount and dismount an NCP volume for the Linux file system.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	`10.10.10.136`
MOUNT_FS	`ext3`
VOLGROUP_NAME	`clustervg01`
MOUNT_DEV	`/dev/$VOLGROUP_NAME/clustervol01`
MOUNT_POINT	`/mnt/clustervol01`
NCP_SERVER	`clus1-clustervg01-server`
NCP_VOLUME	`CLUSTERVOL01`
Volume ID	`251` Ensure that the volume ID that you assign to the NCP volume is unique across all nodes in the cluster. Clustered NCP volume IDs are assigned from 254 to 0, in descending order.

- "Sample LVM Resource Load Script with an NCP Volume" on page 318
- "Sample LVM Resource Unload Script with an NCP Volume" on page 319
- "Sample LVM Resource Monitor Script with an NCP Volume" on page 320

Sample LVM Resource Load Script with an NCP Volume

Use the following sample load script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01
```

```
# define NCP server name
NCP_SERVER=cluster1-clustervg01-server
# define NCP volume name
NCP_VOLUME=CLUSTERVOL01

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP server name
exit_on_error ncpcon bind --ncpservername=$NCP_SERVER --
ipaddress=$RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Script with an NCP Volume

Use the following sample unload script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# define NCP server name
NCP_SERVER=cluster1-clustervg01-server
# define NCP volume name
NCP_VOLUME=CLUSTERVOL01

# unbind the NCP server name
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

```
# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME

# unmount the Linux volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample LVM Resource Monitor Script with an NCP Volume

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume. To use the script, you must also enable monitoring for the resource. See Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 176.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=clustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/clustervol01
# define the mount point
MOUNT_POINT=/mnt/clustervol01

# define NCP server name
NCP_SERVER=cluster1-clustervg01-server
# define NCP volume name
NCP_VOLUME=CLUSTERVOL01

# check the logical volume
exit_on_error status_lv $MOUNT_DEV
```

```
# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# check the status of NCP services
rcndsd status
  if test $? != 0; then
     exit_on_error rcndsd restart
  fi
sleep 5

exit 0
```

13.7 Modifying the LVM Resource Scripts for Other Services

After you create an LVM volume group cluster resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the Generic File System template to the resource template for your product to determine which lines need to be added or modified.

The OES Cluster Services resource templates are available in the Clusters plug-in for iManager. From the Cluster Options page, select the resource template to view its properties, then click the **Scripts** tab to access its default load, unload, and monitor scripts. You can also select the check box next to the resource template, then click **Details** to access its properties. See Section 10.3, “Using Cluster Resource Templates,” on page 169.

13.8 Renaming the Mount Point Path for a Clustered LVM Volume

For a clustered LVM volume, you can rename the mount point path by modifying the mount point variable in the cluster load, unload, and monitor scripts.

1. Open iManager in a web browser, then log in as an administrator user.
2. In **Roles and Tasks**, select **Clusters > My Clusters**.
3. Select the cluster where you created the volume group cluster resource.
4. On the Cluster Manager page, select the check box next to the resource, then click **Offline**.

   Wait for the status to report that it is offline, then continue.
5. Modify the mount point path value in the load, unload, and monitor scripts for the LVM volume group cluster resource:
   - **5a** On the Cluster Manager page, select the resource’s name link to open its Cluster Properties page, then click the **Scripts** tab.

     The Scripts tab automatically displays the load script.
   - **5b** Modify the load script:
     - **5b1** In the load script, type the new value for the mount point in the MOUNT_POINT variable:

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

5b2 Ensure that a `mkdir` command is added above the Linux `mount` command line in the load script to create the path on nodes if it does not exist.

```
# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT
```

You can alternatively make the new path by using the `mkdir` command in a terminal console on each node. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

5b3 Click **Apply**.

5c Click the **Unload Script** link, type the new value for the mount point in the MOUNT_POINT variable, then click **Apply**.

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

5d Click the **Monitor Script** link, type the new value for the mount point in the MOUNT_POINT variable, then click **Apply**.

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

5e At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is brought online.

6 Bring the resource online to allow the script changes to take effect.

6a In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster.

6b On the Cluster Manager page, select the check box next to the resource, then click **Online**.

6c Verify that the resource comes online and reports a **Running** state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

7 In NSSMU, verify that the new mount point is used when the clustered LVM volume resource is brought online:

7a Log in as the `root` user to the node that is hosting the resource, then start NSSMU by entering:

```
nssmu
```

7b From the NSSMU main menu, select **Linux Volumes**, then press Enter.

7c In the `Linux Volumes` list, select the clustered LVM volume.

7d View the volume details to verify that the mount point has changed.

7e Click Escape twice to exit NSSMU.

13.9 Renaming a Clustered LVM Logical Volume

For a clustered LVM logical volume, there is no easy tool for renaming the volume. You can rename the LVM logical volume by using the `lvrename` command, then modify the volume name in the cluster load, unload, and monitor scripts. The command does not change the volume group name. The LVM cluster resource must be offline while you rename the logical volume and modify the scripts.

Using the `lvrename` command to rename the LVM logical volume does not modify other related settings and objects that are used by the LVM cluster resource. You can optionally also modify the following:

- Mount point path
- Resource name
- If NCP is enabled for the LVM logical volume:
  - NCP virtual server name
  - NCP volume name
  - Volume object for the NCP volume on the virtual server

In the following procedure, you take the resource offline, modify the related parameters, modify the resource scripts with the new values, then bring the resource online.

1. Log in to the server as the `root` user, then open a terminal console.
2. Bring the LVM cluster resource offline.

   `cluster offline <lvm_resource_name>`

   For example, to offline the `vol44_resource`, enter

   `cluster offline vol44_resource`

   Wait until the resource is offline before continuing.
3. At the command prompt, enter

   `lvrename </dev/vg_name/old_lv_name> </dev/vg_name/new_lv_name>`

   Replace *vg_name* with the name of the volume group. If you created the LVM logical volume in NSSMU, the volume group name is the same as the logical volume name. If you created the LVM logical volume with NLVM or with LVM commands, you might have used a different name for the volume group.

   Replace *old_lv_name* with the current name of the LVM logical volume.

   Replace *new_lv_name* with the new name of the LVM logical volume.

   For example, to change the name of the logical volume on volume group **vghome** from `lv_users` to `lv_home`, enter

   `lvrename /dev/vghome/lv_users /dev/vghome/lv_home`

4 Modify the load, unload, and monitor scripts for the LVM cluster resource to use the new LVM logical volume name in the value for the MOUNT_DEV parameter.

4a Open iManager in a web browser, then log in as a cluster administrator user.

4b In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster where you created the LVM cluster resource.

4c On the Cluster Manager page, select the check box next to the resource, then click **Offline**.

Wait for the status to report that it is offline, then continue.

4d Modify the MOUNT_DEV value in the load, unload, and monitor scripts for the LVM volume group cluster resource:

4d1 On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the **Scripts** tab.

The Scripts tab automatically displays the load script.

4d2 Modify the load script:

In the load script, type the new value for the logical volume in the MOUNT_DEV variable, then click **Apply**.

```
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```

4d3 Click the **Unload Script** link, type the new value for the logical volume in the MOUNT_DEV variable, then click **Apply**.

```
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```

4d4 Click the **Monitor Script** link, type the new value for the logical volume in the MOUNT_DEV variable, then click **Apply**.

```
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```

4d5 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is brought online.

5 (Optional) Rename the mount point path that you use for the logical volume. The mount point should still be unmounted when you perform this task.

For information, see Section 13.8, "Renaming the Mount Point Path for a Clustered LVM Volume," on page 321.

6 (Optional) Rename the resource.

Ensure that the resource is offline, then enter

```
cluster rename <old_resource_name> <new_resource_name>
```

7 (Optional) If you created an NCP virtual server for the resource, rename the NCP virtual server name.

Use the **Directory Administration > Delete Object** task in iManager to delete the old NCP virtual server (the NCS:NCP Server object) for the resource, then use `/opt/novell/ncs/bin/ncs_ncpserv.py` to create a new virtual server for the resource. For information, see Section 13.5.1, “Creating an NCP Virtual Server for the LVM Resource,” on page 309. Ensure that you update the load, unload, and monitor scripts.

8 (Optional) If NCP file access is enabled for the LVM logical volume, rename the NCP volume and its Volume object.

8a Modify the NCP volume name for the NCP_VOLUME parameter in the load, unload, and monitor scripts for the LVM cluster resource. For information, see Section 13.5.2, “Adding NCP Virtual Server Commands to the Resource Scripts,” on page 309. Bring the resource online, but do not allow users to access the volume until the Volume object is re-created.

8b Use the **Directory Administration > Delete Object** task in iManager to delete the old Volume object, then create a new Volume object that uses the newly created NCP virtual server name and new NCP volume name. For information, see Section 13.6.2, “Creating a Shared NCP Volume Object,” on page 317.

9 Bring the resource online. At the command prompt, enter

cluster online <*resource_name*>

13.10 Disabling Clustering for an LVM Volume Group and Logical Volume

Before you attempt to disable clustering for an LVM volume group and logical volume, you must take the resource offline.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered volume group might not be recoverable.

Deleting a volume group cluster resource disables only the OES Cluster Services cluster settings. There are additional tasks to perform to prepare the LVM volume group and logical volume for non-clustered access on one of the cluster nodes. To delete a resource and create a new one with the same name, you must wait to create the new one until after eDirectory synchronizes all of the objects in the tree related to the deleted resource.

If the volume was NCP enabled, the definition of the NCP volume in the load script is automatically deleted, but you must manually remove its NCP Volume object in eDirectory. To re-create the NCP volume on the non-clustered Linux LVM logical volume, can use the Manage NCP Services plug-in for OES Remote Manager or use the `ncpcon create volume` command.

All resource configuration should happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node.

We strongly recommend that when you need to delete a cluster resource, that you do so only from the master node in the cluster. You might want to delete a cluster resource on a non-master node if the master node is not a preferred node for the resource, or if there is a resource mutual exclusion rule that prohibits the resource from being moved there. If the resource cannot be migrated to the master node, follow the additional steps that are annotated as “Non-Master”.

Use the procedure in this section to disable clustering for a Linux LVM volume group and logical volume. Afterwards, you can mount and dismount the volume on only one node in the cluster.

1. Ensure that the resource is online on the node where you want mount the volume after clustering is disabled.

   If the resource is not on the desired node, open a terminal console as the `root` user, then enter

   ```
   cluster migrate <resource_name> <from_node> <to_node>
   ```

   For example, if node2 is the to node, enter

   ```
   cluster migrate vg44_resource node1 node2
   ```

2. Log in to iManager as a cluster administrator.
3. In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster.
4. Take the volume group cluster resource offline, by using one of the following methods:
   - In iManager on the Cluster Manager page, select the check box next to the LVM volume group cluster resource (such as `vg44_resource`), then click **Offline**. Wait until the resource is offline to continue.
   - Open a terminal console as the `root` user, then enter

     ```
     cluster offline <resource_name>
     ```

   The unload script unbinds the NCP virtual server from the resource IP address, dismounts the NCP volume, dismounts the Linux LVM logical volume, and deactivates the Linux LVM volume group.
5. Delete the volume group cluster resource:
   - **5a** (Non-Master) If the volume group was taken offline on a non-master node, verify that the revision and resource information is the same on the non-master node as on the master node before you delete the resource.
     - **5a1** On the non-master node, log in as the `root` user, then open a terminal console. At a command prompt on the non-master node, enter

       ```
       /opt/novell/ncs/bin/ncs-configd.py -init
       ```

     - **5a2** Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.
   - **5b** Click the **Cluster Options** tab.
   - **5c** Select the check box next to the volume group cluster resource, then click **Delete**.
   - **5d** When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.
   - **5e** (Non-Master) If the volume group was taken offline on a non-master node, verify that the revision and resource information is the same on the master node as on the non-master node after you delete the resource.
     - **5e1** On the master node, log in as the `root` user, then open a terminal console. At a command prompt on the master node, enter

       ```
       /opt/novell/ncs/bin/ncs-configd.py -init
       ```

5e2 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.

5e3 In iManager, select **Clusters > Cluster Options**, then browse to select the Cluster object.

5e4 Click **Properties**, select the **Priorities** tab, then click **Apply** on the Priorities page.

5e5 At a command prompt, enter

```
cluster view
```

The cluster view should be consistent.

5e6 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

If the revision number increased, you are done. Do not continue with Step 13.

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

5e7 If the revision number did not automatically update in the previous steps, restart OES Cluster Services by entering the following on one node in the cluster:

```
cluster restart [seconds]
```

For *seconds*, specify a value of 60 seconds or more.

For example:

```
cluster restart 120
```

6 If the Linux volume was NCP enabled, delete the NCP volume object.

6a In iManager, select **Directory Administration > Delete Object.**

6b On the Delete Object page, browse to locate and select the resource's NCP Volume object, such as `clus1_VOL44.novell`, then click **OK**.

6c Click **OK** to delete the NCP volume object, then click **OK** again to close the success message.

7 If the Linux volume was NCP enabled, or if you manually created an NCP virtual server for the resource, delete the NCP virtual server object:

7a In iManager, select **Directory Administration > Delete Object.**

7b On the Delete Object page, browse to locate and select the resource's NCS:NCP Server object, such as `clus1-vg44-SERVER.ncs.novell`, then click **OK**.

7c Click **OK** to delete the NCP virtual server, then click **OK** again to close the success message.

7d Visually confirm that the resource does not appear in the list of Cluster objects, then exit iManager.

8 Log in as the `root` user on the server where the resource was online, then open a terminal console.

9 Deactive the volume group from Clustered LVM:

```
vgchange -c n <vg_name>
```

For example:

```
vgchange -c n vg44
```

10 Activate the volume group locally on the current node:

```
vgchange -a ey <vg_name>
```

For example:

```
vgchange -a ey vg44
```

11 Mount the LVM logical volume locally on the current node:

```
mount -t <fstype> <device_name> <full_mount_point_path>
```

Replace *fstype* with the file system type of the volume.

Replace *device_name* with the full device path of the logical volume, such as /dev/<*vg_name*>/<*lv_name*>.

Replace *full_mount_point_path* with the volume's mount point.

For example:

```
mount -t ext3 /dev/vg44/vol44 /mnt/vol44
```

12 Add an entry for the LVM volume in the `/etc/fstab` file to allow the volume to be mounted automatically on reboot. It also provides the automatic mount information for NSSMU to use to mount and dismount the volume on the node.

In a text editor, modify the `/etc/fstab` file to specify the mount point information and file system type.

For example, complete the line for the volume's device path, such as:

```
/dev/vg44/vol44   /mnt/vol44     ext3     rw          0 0
```

13 View details about the non-clustered LVM logical volume:

13a Open an terminal console as the `root` user, launch NSSMU:

```
nssmu
```

13b In the NSSMU main menu, select **Linux Volumes**, then press Enter.

13c On the Linux Volumes page, then select the volume to view its details.

The volume is no longer cluster enabled. You can mount and dismount the volume only on the current node.

14 (Optional) Create an NCP volume for the Linux volume by using OES Remote Manager. For information, see "Creating NCP Volumes on Linux File Systems" in the *OES 2023: NCP Server for Linux Administration Guide*.

To use the same volume name that was used in the cluster resource load script, wait to re-create the NCP volume until after eDirectory synchronizes all of the objects in the tree related to the deleted resource.

13.11 Deleting a Clustered LVM Volume Group and Logical Volume

Before you delete a clustered LVM volume group, you must take the volume group cluster resource offline, and delete the cluster resource object and related objects in eDirectory.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node.

1. Log in to iManager as a cluster administrator.
2. In **Roles and Tasks**, select **Clusters > My Clusters**, then select the cluster.
3. Take the volume group cluster resource offline:
   - **3a** On the Cluster Manager page, select the check box next to the volume group cluster resource, then click **Offline**. Wait until the resource is offline to continue.

     The unload script unbinds the NCP virtual server from the resource IP address, dismounts the NCP volume, dismounts the Linux LVM logical volume, and deactivates the Linux LVM volume group.
4. Delete the volume group cluster resource:
   - **4a** Click the **Cluster Options** tab.
   - **4b** Select the check box next to the volume group cluster resource, then click **Delete**.
   - **4c** When you are prompted to confirm the deletion, click **OK** to continue, or click **Cancel** to abort the deletion.
5. If the Linux volume was enabled for NCP file access, delete the NCP volume object.
   - **5a** In iManager, select **Directory Administration > Delete Object.**
   - **5b** On the Delete Object page, browse to locate and select the resource's NCP Volume object, then click **OK**.
   - **5c** Click **OK** to delete the NCP volume object, then click **OK** again to close the success message.
6. If the Linux volume was NCP enabled, or if you manually created an NCP virtual server for the resource, delete the NCP virtual server object.
   - **6a** In iManager, select **Directory Administration > Delete Object.**
   - **6b** On the Delete Object page, browse to locate and select the resource's NCS:NCP Server object, then click **OK**.
   - **6c** Click **OK** to delete the NCP virtual server, then click **OK** again to close the success message.
7. Delete the LVM volume group and logical volume:
   - **7a** Log in as the `root` user to the cluster node where the resource was online, then open a terminal console.
   - **7b** At a command prompt, launch NSSMU:

     ```
     nssmu
     ```
   - **7c** In the NSSMU main menu, select **Linux Volumes**, then press Enter.
   - **7d** In the **Linux Volumes** list, select the clustered linux volume, then press Delete.

7e When you are prompted to confirm the delete action, press y (Yes) to continue, or press n (No) to cancel the delete action.

Deleting the volume and volume group puts the device in an uninitialized state.

7f Press Esc to return to the main menu.

8 Re-initialize the device:

8a In the NSSMU main menu, select **Devices**, then press Enter.

8b In the **Devices** list, select the device.

8c Press F3 to initialize the device.

8d When you are prompted to confirm the initialization action, read the warning message, then press y (Yes) to continue, or press n (No) to cancel the action.

8e Select **DOS** or **GPT** as the partitioning type, then press Enter.

8f Visually verify that the device was initialized and the specified partition type is displayed.

8g Press Esc twice to exit NSSMU.

9 Verify that the LVM volume group is deleted by entering the following the native LVM command:

```
vgdisplay
```

13.12 Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)

We strongly recommend that you delete a cluster-enabled LVM volume only from the master node in the cluster. This allows the cluster information to be automatically updated.

The procedures in this section assume that you created the clustered LVM volume in NSSMU or NLVM. There are default naming conventions applied by these tools that might not apply to LVM volume groups that you created and cluster-enabled by using native LVM tools and the Generic File System template.

WARNING: Deleting an LVM volume destroys all data on it.

NSSMU and the `nlvm delete linux volume <volume_name>` command delete the cluster-enabled LVM volume as well as the related objects in eDirectory:

- Linux LVM volume group and logical volume from the file system
- Cluster Resource object for the LVM resource
- If the LVM volume is NCP-enabled:
  - Volume object for the LVM volume
  - Virtual server for the LVM resource (NCS:NCP Server object)

When the LVM volume resides on the master node, the cluster information is automatically updated.

When the LVM volume resides on a non-master node, additional steps are required to update the cluster information. A cluster restart might be needed to force the information to be updated.

Use the following procedures to delete a cluster-enabled LVM volume:

- Section 13.12.1, “Deleting a Cluster-Enabled LVM Volume on the Master Node,” on page 331
- Section 13.12.2, “Deleting a Cluster-Enabled LVM Volume on a Non-Master Node,” on page 332

13.12.1 Deleting a Cluster-Enabled LVM Volume on the Master Node

1. If the LVM resource is on a non-master node in the cluster, migrate it to the master node. As the `root` user, open a terminal console, then enter

   ```
   cluster migrate <resource_name> <master_node_name>
   ```

   To migrate the resource, the master node must be in the resource’s preferred nodes list.

2. Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline.

   As the `root` user, enter

   ```
   cluster offline <resource_name>
   ```

   Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.

3. Delete the LVM volume on the master node by using NSSMU.

   You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.

   3a In NSSMU, select **Linux Volumes**, then press Enter.

   3b Select the unmounted LVM volume, then press Delete.

   3c Select OK to confirm, then press Enter.

4. In the **Tree View** in iManager, browse the objects to verify that the following objects were deleted:

   - LVM resource object (from the Cluster container)
   - If the LVM volume was NCP-enabled:
     - Volume object for the LVM volume
     - Virtual server for the LVM resource (NCS:NCP Server object)

5. Re-initialize the device that contained the LVM volume.

   When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.

   5a In NSSMU, select **Devices**, then press Enter.

   5b Select the device, then press F3 (Initialize).

   5c Press y (Yes) to confirm.

   5d Select the partitioning scheme as DOS or GPT, then press Enter.

6. (Optional) Use a third-party SAN management tool to assign the device to only one desired server.

13.12.2 Deleting a Cluster-Enabled LVM Volume on a Non-Master Node

1. Log in as the `root` user to the non-master node where the cluster resource currently resides, then open a terminal console.
2. Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

   `cluster offline <resource_name>`

   Use the `cluster status` command to verify that the resource has a status of **Offline** before you continue.
3. At the command prompt on the non-master node, enter

   `/opt/novell/ncs/bin/ncs-configd.py -init`
4. Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.
5. Delete the LVM volume on the master node by using NSSMU.

   You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.

   5a In NSSMU, select **Linux Volumes**, then press Enter.

   5b Select the unmounted LVM volume, then press Delete.

   5c Select OK to confirm, then press Enter.
6. In the **Tree View** in iManager, browse the objects to verify that the following objects were deleted:
   - LVM resource object (from the Cluster container)
   - If the LVM volume was NCP-enabled:
     - Volume object for the LVM volume
     - Virtual server for the LVM resource (NCS:NCP Server object)
7. Re-initialize the device that contained the LVM volume.

   When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.

   7a In NSSMU, select **Devices**, then press Enter.

   7b Select the device, then press F3 (Initialize).

   7c Press y (Yes) to confirm.

   7d Select the partitioning scheme as DOS or GPT, then press Enter.
8. On the master node, log in as the `root` user, open a terminal console, then enter

   `/opt/novell/ncs/bin/ncs-configd.py -init`
9. Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.
10. In iManager, select **Clusters > My Clusters**, select the cluster, then select the **Cluster Options** tab.
11. Click **Properties**, select the **Priorities** tab, then click **Apply** on the Priorities page.
12. At a command prompt, enter

```
cluster view
```

The cluster view should be consistent.

13 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

If the revision number increased, skip Step 14.

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

14 If the revision number did not automatically update in the previous steps, restart OES Cluster Services by entering the following on one node in the cluster:

```
cluster restart [seconds]
```

For *seconds*, specify a value of 60 seconds or more.

For example:

```
cluster restart 120
```

15 (Optional) Use a third-party SAN management tool to assign the devices to only the desired server.

13.13 Linux LVM Management Tools

LVM tools are available in the YaST Expert Partitioner and in native Linux command line commands.

- Section 13.13.1, "Using the YaST Expert Partitioner," on page 333
- Section 13.13.2, "Using LVM Commands," on page 333

13.13.1 Using the YaST Expert Partitioner

You can access the LVM tools in the YaST Expert Partitioner by using the desktop menus, or enter the following in a terminal console:

```
yast2 disk
```

For information about using the Partitioner, see "LVM Configuration" (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-lvm.html) in the *SLES 12: Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html).

13.13.2 Using LVM Commands

For information about using LVM commands, see the man pages for the commands described in Table 13-1. Perform the commands as the `root` user.

Table 13-1 *LVM Commands*

Command	Description
`pvcreate <device>`	Initializes a device (such as `/dev/sdb`) for use by LVLM as a physical volume.
`pvdisplay <device>`	Displays information about the LVM physical volume, such as whether it is currently being used in a logical volume.
`vgcreate -c y <vg_name> <dev1> [dev2...]`	Creates a clustered volume group with one or more specified devices.
`vgchange -a [ey | n] <vg_name>`	Activates (`-a ey`) or deactivates (`-a n`) a volume group and its logical volumes for input/output. **IMPORTANT:** Ensure that you use the `ey` option to exclusively activate a volume group on a cluster node. This option is used by default in the load script.
`vgremove <vg_name>`	Removes a volume group. Before using this command, remove the logical volumes, then deactivate the volume group.
`vgdisplay <vg_name>`	Displays information about a specified volume group. To find the total physical extent of a volume group, enter `vgdisplay vg_name | grep "Total PE"`
`lvcreate -L size -n <lv_name> <vg_name>`	Creates a logical volume of the specified size.
`lvremove</dev/vg_name/lv_name>`	Removes a logical volume, such as `/dev/vg_name/lv_name`. Before using this command, close the logical volume by dismounting it with the `umount` command.
`vgextend <vg_name><device>`	Adds a specified physical volume to an existing volume group
`vgreduce <vg_name> <device>`	Removes a specified physical volume from an existing volume group. **IMPORTANT:** Ensure that the physical volume is not currently being used by a logical volume. If it is, you must move the data to another physical volume by using the `pvmove` command.
`lvextend -L size</dev/vg_name/lv_name>`	Extends the size of a specified logical volume. Afterwards, you must also expand the file system to take advantage of the newly available space.
`lvreduce -L size </dev/vg_name/lv_name>`	Reduces the size of a specified logical volume. **IMPORTANT:** Ensure that you reduce the size of the file system first before shrinking the volume, otherwise you risk losing data.

Command	Description
`lvrename </dev/vg_name/old_lv_name> </dev/vg_name/new_lv_name>`	Renames an existing LVM logical volume in a volume group from the old volume name to the new volume name. It does not change the volume group name.

14 Configuring OES Cluster Services in a Virtualization Environment

OES Cluster Services is installed and configured in a virtualization environment by using the same methods and processes as those used on a physical Open Enterprise Server (OES) 11 or later server. It can be used in the host and guest virtual environments.

For a virtual machine, you can use any virtualization hypervisor (such as Xen, KVM, VMware, and Hyper-V) that supports the 64-bit SUSE Linux Enterprise Server 15 SP4 operating system, which is the platform used for OES 2023.

For information about using OES with Xen and KVM, see the following resources:

- For information on setting up virtualized OES servers, see "Installing OES on a VM" in the *OES 2023: Installation Guide*.
- To get started with Xen virtualization, see the SLES 12 SP5 Virtualization Guide (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html).
- To get started with KVM virtualization, see the SLES 12 SP5 Virtualization Guide (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html).
- To get started with third-party virtualization platforms, such as Hyper-V from Microsoft and the different VMware product offerings, refer to the documentation for the product you are using. For information about setting up a cluster with nodes hosted on VMware, see *OES Cluster Services Implementation Guide for VMware*.

The following sections provide information to help you deploy and run Cluster Services in a virtual environment: Although many different cluster virtualization scenarios are possible, only those outlined in the following sections have been tested.

- Section 14.1, "Prerequisites for Xen Host Server Environments," on page 338
- Section 14.2, "Virtual Machines as Cluster Resources," on page 338
- Section 14.3, "Virtual Machines as Cluster Nodes," on page 347
- Section 14.4, "Virtual Cluster Nodes in Separate Clusters," on page 348
- Section 14.5, "Mixed Physical and Virtual Node Clusters," on page 349
- Section 14.6, "Additional Information," on page 351

14.1 Prerequisites for Xen Host Server Environments

You can install OES Cluster Services on a Xen virtualization host server (Dom0) running OES. You can create cluster resources that contain the virtual machine information by using the Xen and XenLive resource templates. You can fail over or cluster migrate these virtual machine cluster resources to different physical nodes in your cluster. Only the Xen and XenLive templates can be used in the Xen host server environment. You can modify these templates to work with other types of hypervisors.

All templates except Xen and XenLive are valid on guest servers (DomU). When Cluster Services is installed on an OES guest server environment, no additional changes or special configuration are required. You can create clusters consisting of all virtual nodes, or use a combination of virtual and physical nodes.

When using Cluster Services to create cluster resources in a Xen host server environment, one network adapter on each cluster node must be available without Xen bridging to allow OES Cluster Services to monitor for network state changes in the cluster. You need at least one other network adapter on each cluster node to use as a bridge for the virtual servers that will be hosted on the Xen host.

For each cluster node, determine which network adapter you want to use for OES Cluster Services communications and network state monitoring, then do the following:

- Disable bridging on the network adapter.
- Assign the IP address of the host server to the network adapter.

For information about why this configuration is necessary, see "TID 7004595: Novell Cluster Services is unable to detect a network link down when installed on a Xen domain 0 (dom0)" (http://www.novell.com/support/viewContent.do?externalId=7004595&sliceId=1) on Micro Focus Support (http://www.novell.com/support).

14.2 Virtual Machines as Cluster Resources

In this scenario, you have an OES cluster configured on physical machines. OES and Xen are installed and configured on each node along with OES Cluster Services. This part of the OES Cluster Services configuration does not differ from that of an OES cluster without virtualization.

You can create OES virtual machines on each cluster node and configure those virtual machines to be cluster resources. You can then fail over or migrate virtual machine cluster resources (entire virtual machines) to different physical nodes in your cluster.

Figure 14-1 depicts how this setup might look. OES Cluster Services (NCS) is installed and running on the virtual machine (VM) host server.

Figure 14-1 *Virtual Machines as Cluster Resources*

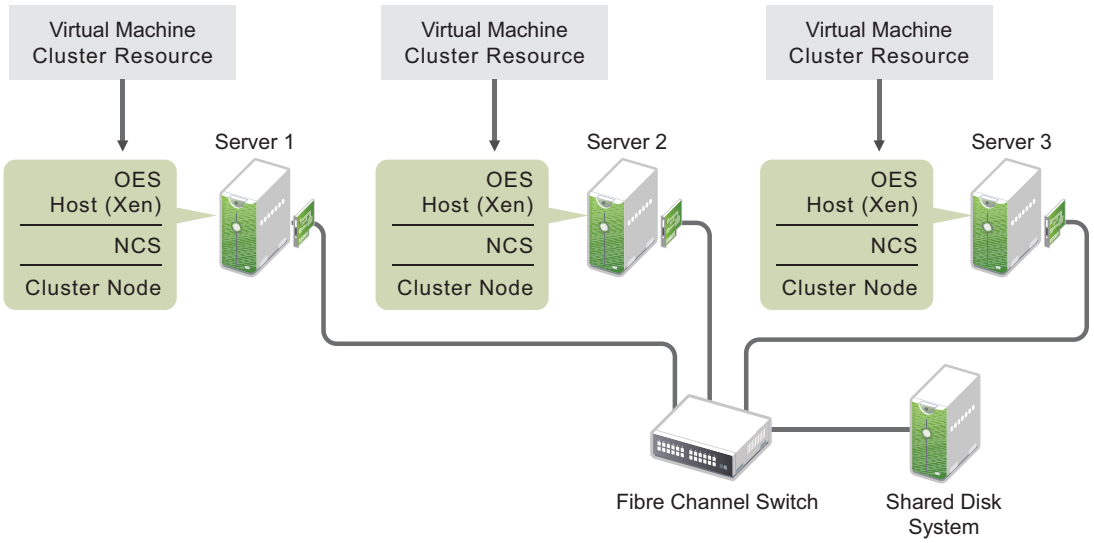

The following sections describe how to create a cluster resource and its cluster scripts for each virtual machine:

- Section 14.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 339
- Section 14.2.2, “Configuring Virtual Machine Load, Unload, and Monitor Scripts,” on page 340
- Section 14.2.3, “Setting Up Live Migration,” on page 347

14.2.1 Creating a Xen Virtual Machine Cluster Resource

OES Cluster Services includes two Xen (virtual machine) resource templates, which greatly simplify the process for creating a virtual machine cluster resource. Much of the virtual machine cluster resource configuration is performed automatically by the Xen resource templates. The two templates are named Xen_Template and XenLive_Template. Both templates perform similar functions to automatically configure the cluster resource for the virtual machine.

The XenLive template provides an additional function to allow a manual virtual machine resource migration without the need to boot or bring up the virtual machine on the cluster node where the virtual machine has been migrated. This lets clients continue to access a virtual machine that has been migrated without reconnecting or waiting for the virtual machine to boot or load on the target node.

IMPORTANT: The live migrate function is only useful for a manual virtual machine resource migration, and does not work for a virtual machine resource failover or failback.

Ensure that your Xen setup is working properly before you attempt to set up the OES Cluster Services clustering for your virtual machines in the Xen host environment. Refer to the *Introduction to Xen Virtualization* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-xen-basics.html) to find out how to set up XEN and XEN virtual machines.

To configure a virtual machine as a cluster resource:

1 Open your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/imanager.html. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.

2 Specify your user name and password, specify the tree where you are installing the cluster, then click **Login**.

3 In iManager, select **Clusters > My Clusters**, select the cluster, then click **Cluster Options**.

4 On the Cluster Options page, click **New**.

5 Click the **Resource** radio button to specify **Resource** as the resource type you want to create, then click **Next**.

6 Specify a name for the virtual machine resource.

7 In the **Inherit From Template** field, browse to the Cluster object container, then select the desired Xen template name from the list of templates in the container.

The Xen templates are named **Xen_Template** and **XenLive_Template**.

8 Select the **Define Additional Properties** check box, click **Next**, then continue with "Configuring Virtual Machine Load, Unload, and Monitor Scripts" on page 340.

14.2.2 Configuring Virtual Machine Load, Unload, and Monitor Scripts

The Xen resource templates configure the virtual machine resource by automatically creating load, unload, and monitor scripts, setting failover and failback modes, and assigning the virtual machine as a resource to all nodes in the cluster.

The load, unload, and monitor scripts for virtual machine cluster resources do not need to be modified if all the following are true:

- The resource name is the same as the virtual machine name.
- The configuration file name is the same as the virtual machine name.
- The mount point directory name is the same as the virtual machine name.
- You are using the Reiser file system.

If you are not modifying the scripts, continue the setup by configuring the resource policies and the resource server assignments. See Section 10.9, "Configuring the Start, Failover, and Failback Modes for Cluster Resources," on page 182 and Section 10.10, "Configuring Preferred Nodes and Node Failover Order for a Resource," on page 184.

If you are modifying the scripts, continue with the following sections:

- "Configuring the Load Script" on page 341
- "Configuring the Unload Script" on page 342
- "Configuring the Monitor Script" on page 344

Configuring the Load Script

The virtual machine resource load script page should already be displayed. The load script contains commands to start the virtual machine. You can customize some commands for your specific configuration.

1 View and, if necessary, edit the lines in the script for your specific directory structure, mount point, configuration file, and file system type (in the Xen_Template).

See the following examples of the default Xen_Template and XenLive_Template load scripts:

- "Sample Xen_Template Load Script" on page 341
- "Sample XenLive_Template Load Script" on page 342

2 Click **Next** and continue with "Configuring the Unload Script" on page 342.

Sample Xen_Template Load Script

The Xen_Template load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=$OCF_RESOURCE_INSTANCE

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/$CONTAINER_NAME/$VOLUME_NAME

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

#activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# mount the file system
exit_on_error ocf_start Filesystem

# start the service
exit_on_error ocf_start Xen

# return status
exit 0
```

Sample XenLive_Template Load Script

The XenLive_Template load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# start the service
if [ -n "$NCS_TOFROM" ]
then
    exit_on_error ocf_migrate_from Xen
else
    exit_on_error ocf_start Xen
fi

# return status
exit 0
```

Configuring the Unload Script

The virtual machine resource unload script page should now be displayed. The unload script contains commands to stop the virtual machine. You can customize some commands for your specific configuration.

1. View and, if necessary, edit the lines in the script for your specific directory structure, mount point, configuration files, and file system type (in the Xen_Template).

   Use the same values that you specified in the load script.

   See the following examples of the default Xen_Template and XenLive_Template unload scripts:

   - “Sample Xen_Template Unload Script” on page 342
   - “Sample XenLive_Template Unload Script” on page 344

2. Click **Next**, then continue the setup by configuring the resource policies and the resource server assignments.

   See Section 10.9, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 182 and Section 10.10, “Configuring Preferred Nodes and Node Failover Order for a Resource,” on page 184.

3. If you want to enable monitoring for the resource, continue with “Configuring the Monitor Script” on page 344.

Sample Xen_Template Unload Script

The Xen_Template unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=$OCF_RESOURCE_INSTANCE

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/$CONTAINER_NAME/$VOLUME_NAME

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# stop the service
exit_on_error ocf_stop Xen

# umount the file system
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error ocf_stop Filesystem

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status
exit 0
```

Sample XenLive_Template Unload Script

The XenLive_Template unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_CRM_meta_migrate_target=$NCS_TOFROM

RC=0
# stop the service
if [ -n "$NCS_TOFROM" ]
then
    RC=`ocf_migrate_to Xen`
    if [ $RC -ne 0 ]
    then
        ignore_error ocf_stop Xen
    fi
else
    ignore_error ocf_stop Xen
fi

# return status
exit $RC
```

Configuring the Monitor Script

The Xen_Template and XenLive Template each include a resource monitor script that you can customize. You use the script to monitor the health of a virtual machine cluster resource.

Resource monitoring is disabled by default. If you want to enable resource monitoring for a virtual machine cluster resource, you must enable it prior to customizing the resource monitor script.

- "Enabling Resource Monitoring" on page 344
- "Viewing or Modifying the Monitor Script" on page 345
- "Sample Xen_Template Monitor Script" on page 346
- "Sample XenLive_Template Monitor Script" on page 346

Enabling Resource Monitoring

To enable resource monitoring for a virtual machine cluster resource:

1. In iManager, select **Clusters > My Clusters**.
2. Select the cluster.
3. On the Cluster Manager page, select the check box next to the virtual machine resource, then click the **Details** link.

4 Click the **Monitoring** tab, then select the **Enable Resource Monitoring** check box to enable resource monitoring for the resource.

   Resource monitoring is disabled by default.

5 For the polling interval, specify how often you want the resource monitor script for this resource to run.

   You can choose to specify the number in minutes or seconds.

6 Specify the number of failures (**Maximum Local Failures**) for the specified amount of time (**Time Interval**).

   If the resource monitor detects that the resource fails the number of times specified in the amount of time specified, a failure action initiates.

7 Specify whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

   If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its **Preferred Nodes** list. The resource remains on the server it has migrated to unless you migrate it to another server or the failure action initiates again, in which case it again migrates to the next server in its **Preferred Nodes** list.

   If the failure action initiates and you chose the option to reboot the hosting node without synchronizing or unmounting the disks, each of the resources on the hosting node will fail over to the next server in its **Preferred Nodes** list because of the reboot. This is a hard reboot, not a graceful one.

   With resource monitoring, the **Failover**, **Failback**, and **Start** modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back to the node it migrated from unless you manually migrate it back.

Viewing or Modifying the Monitor Script

To view or customize the monitor script for the virtual machine's cluster resource:

1 In iManager, select **Clusters > My Clusters**.

2 Select the cluster.

3 On the Cluster Manager page, select the check box next to the virtual machine resource that you created, then click the **Details** link.

4 Click the **Scripts** tab, then click the **Monitor Script** link.

5 View or edit the commands in the script that monitor the resource on the server.

   You can use the same commands that would be used at the Linux terminal console.

   See the following examples of the default Xen_Template and XenLive_Template monitor scripts:

   - "Sample Xen_Template Monitor Script" on page 346
   - "Sample XenLive_Template Monitor Script" on page 346

6 Specify the **Monitor Script Timeout** value, then click **Apply** to save the script.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action initiates based on your settings in Step 7 of "Enabling Resource Monitoring" on page 344. Cluster Services marks the monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

Sample Xen_Template Monitor Script

The Xen_Template monitor script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=$OCF_RESOURCE_INSTANCE

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/$CONTAINER_NAME/$VOLUME_NAME

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

#check the logical volume
exit_on_error status_lv $MOUNT_DEV

# status of the file system
exit_on_error ocf_status Filesystem

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

Sample XenLive_Template Monitor Script

The XenLive_Template monitor script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

14.2.3 Setting Up Live Migration

Live migrations use the XenLive template. You can manually copy the virtual machine configuration files to the same path on each node of the cluster, or you can set up an OCFS2 file system for the configuration files. Do one of the following:

- Manually copy the configuration file for the virtual machine to the same directory (the path must be the same) on each cluster node where the virtual machine will run.
- Configure the OCFS2 file system on a shared disk system and copy the virtual machine configuration file to a directory on the file system. You also must ensure that all cluster nodes where the virtual machine will run have access to the OCFS2 file system on the shared disk system.

  Ensure that your OCFS2 file system is working properly before you attempt to use it with OES Cluster Services.

  For detailed information about using OCFS2, see the *OCFS2 Project* (http://oss.oracle.com/projects/ocfs2/) on the Oracle website.

14.3 Virtual Machines as Cluster Nodes

In this scenario, you have OES and Xen installed and configured on each node (physical machine). You create the OES virtual machine on each physical machine and install and configure OES Cluster Services on each virtual machine. The combined virtual machines (cluster nodes) comprise one cluster.

You can then create and configure cluster resources that run on the virtual cluster nodes. The process for creating and configuring cluster resources on a virtual cluster node is the same as on a physical cluster node.

Cluster resources can be failed over or migrated between virtual cluster nodes that are on the same physical node or on separate physical nodes.

Figure 14-2 depicts using virtual machines as cluster nodes.

Figure 14-2 *Virtual Machines as Cluster Nodes*

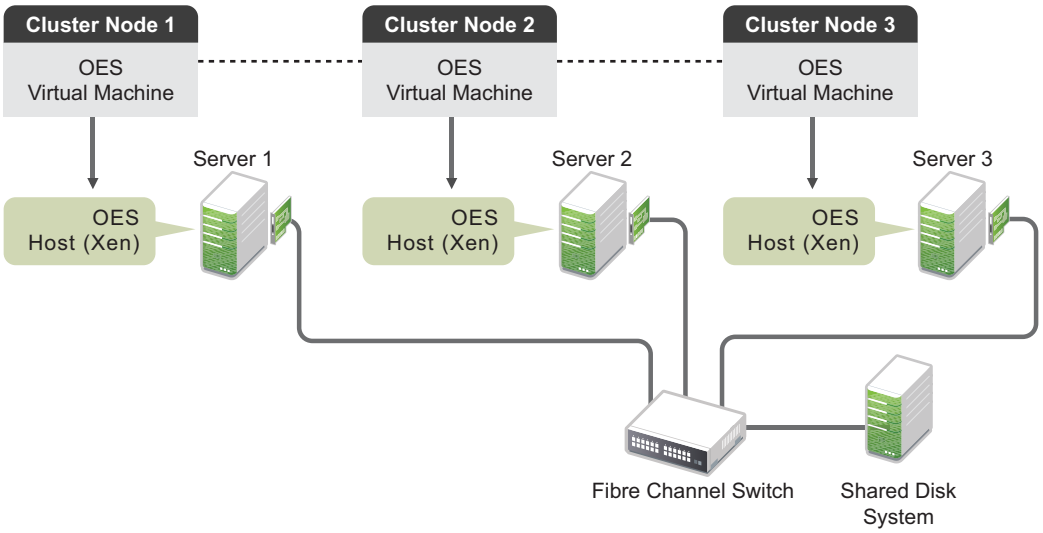

14.4 Virtual Cluster Nodes in Separate Clusters

In this scenario, you have OES and Xen installed and configured on each node (physical machine). You create multiple OES virtual machines on each physical machine and install and configure OES Cluster Services on each virtual machine. During the Cluster Services installation, you create separate clusters of virtual cluster nodes, with each virtual cluster node residing on a separate physical machine. This way you have multiple clusters of virtual cluster nodes on fewer physical machines.

You can then create and configure cluster resources on each virtual cluster node and cluster. The process for creating and configuring cluster resources on a virtual cluster node is the same as on a physical cluster node.

Figure 14-3 depicts using virtual cluster nodes in separate clusters.

Figure 14-3 *Virtual Cluster Nodes in Separate Clusters*

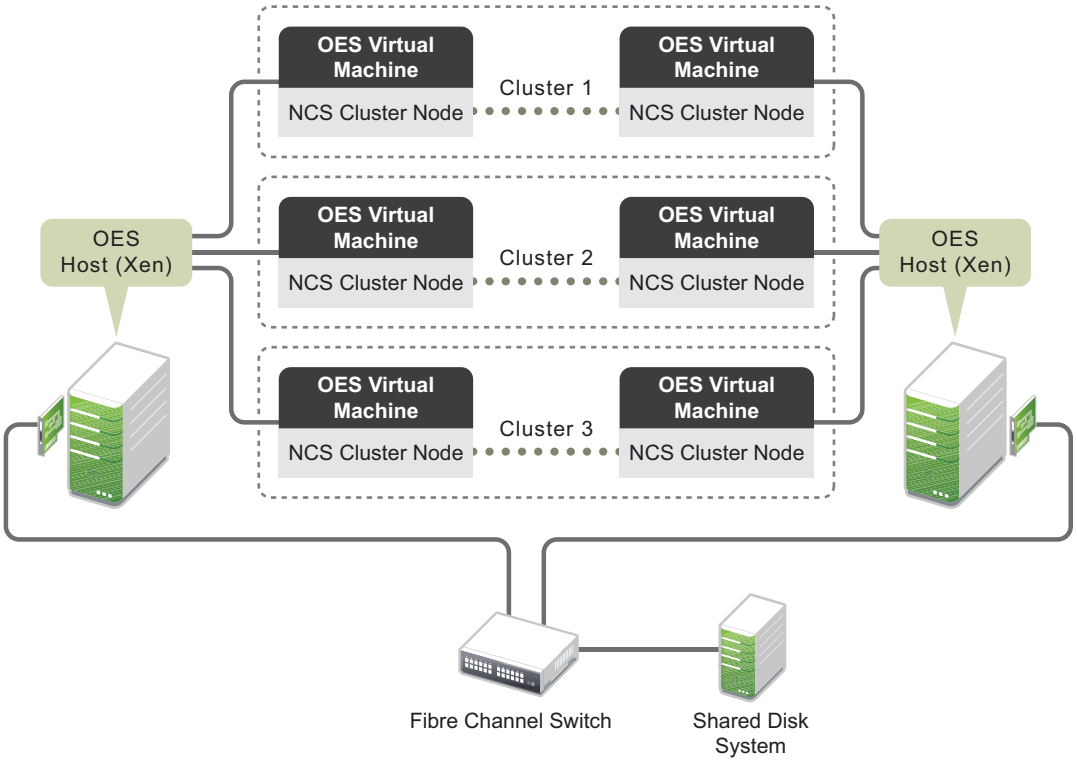

14.5 Mixed Physical and Virtual Node Clusters

This is a temporary scenario that is used for upgrading cluster nodes or converting clusters from physical to virtual cluster nodes.

This can be done through several different methods. One method is to add a virtual OES cluster node to an existing physical OES cluster. To do this, you install an OES and Xen server (physical machine). You then create a OES virtual machine on the physical machine and install and configure OES Cluster Services on the OES virtual machine. During the OES Cluster Services installation, you add the OES virtual cluster node to your existing OES cluster (physical nodes).

You can then migrate the desired resources from the physical cluster nodes to the virtual cluster node. This lets you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical cluster nodes with virtual cluster nodes.

Figure 14-4 depicts how this setup might look.

Figure 14-4 *Mixed Physical and Virtual Node Cluster*

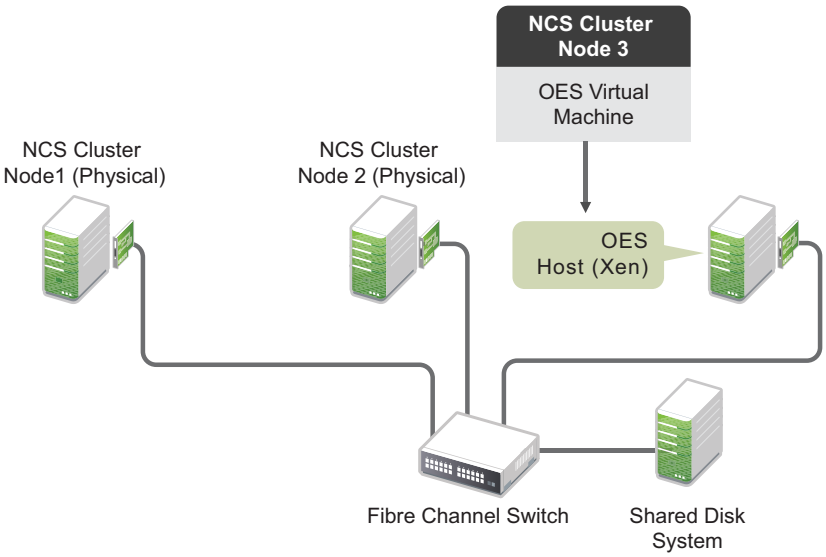

Another method is to install OES Cluster Services on physical nodes and create a separate cluster for each node. You then install an OES with Xen server (physical machine) and create OES virtual machines and install OES Cluster Services on each virtual machine. You can then add one virtual OES cluster node to each cluster to create multiple two-node clusters, each containing one physical and one virtual cluster node.

This allows you to migrate the desired resources from each physical cluster node to the virtual cluster node in the same cluster. Using this setup, you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical cluster nodes in each cluster with virtual cluster nodes.

Figure 14-5 depicts how this setup might look.

Figure 14-5 *Separate Mixed Physical and Virtual Node Clusters*

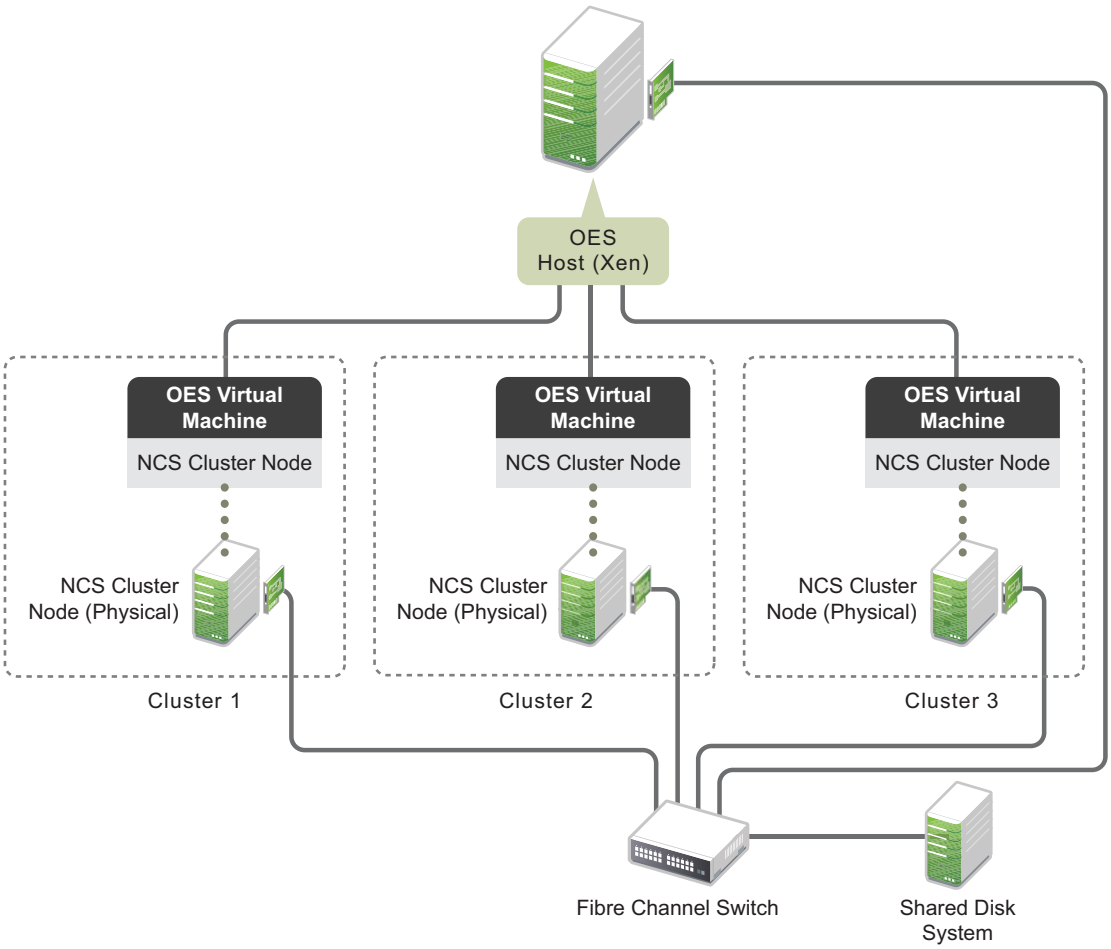

14.6 Additional Information

To get started with virtualization, see "Introduction to Xen Virtualization" (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-xen-basics.html) in the *SLES Virtualization Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html) guide.

For information on setting up virtualized OES servers, see "Installing OES on a VM" in the *OES 2023: Installation Guide*.

15 Troubleshooting OES Cluster Services

This section describes known issues and frequently asked questions for managing OES Cluster Services.

- Section 15.1, “Migration or Online of Resources Causes Cluster Crash,” on page 354
- Section 15.2, “Cluster Node Crashing Issue,” on page 354
- Section 15.3, “Unable to Manage NCS from iManager When NSS is not Installed on the Node,” on page 355
- Section 15.4, “Cluster Service Does Not Come Up After Upgrading to OES 2018 or Later if Service Proxy is Configured,” on page 355
- Section 15.5, “Supported SBD Signatures in Cluster Nodes with Different OES Versions,” on page 356
- Section 15.6, “Cluster Resource Is Not Accessible After a Network Service Restart,” on page 356
- Section 15.7, “File Location Information,” on page 356
- Section 15.8, “Version Issues,” on page 357
- Section 15.9, “Diagnosing a Poison Pill Condition,” on page 357
- Section 15.10, “Diagnosing Cluster Problems,” on page 357
- Section 15.11, “Master IP Address Problem,” on page 358
- Section 15.12, “Why didn’t my resource go to the node I specified?,” on page 358
- Section 15.13, “iManager: Error 500 Occurs while Connecting to a Cluster,” on page 359
- Section 15.14, “iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5,” on page 359
- Section 15.15, “iManager: Problems with Clusters Plug-In After Update for iManager 2.7.5,” on page 360
- Section 15.16, “A Device Name Is Required to Create a Cluster Partition,” on page 360
- Section 15.17, “Cluster Resource Goes Comatose After Migration or Failover,” on page 360
- Section 15.18, “Cannot Authenticate to Remote Servers during Cluster Configuration,” on page 360
- Section 15.19, “Cannot Connect to an iSCSI Target,” on page 361
- Section 15.20, “Is There a way to Uninstall OES Cluster Services from a Server?,” on page 361
- Section 15.21, “Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems,” on page 361
- Section 15.22, “Error NFS Stale Handle Reported on Host when Cluster Failover Occurs,” on page 361
- Section 15.23, “Error 20897 - This node is not a cluster member,” on page 362
- Section 15.24, “Error 499 - Could not delete this resource Data_Server,” on page 362
- Section 15.25, “Error 603 - smdr.novell is not registered with SLP for a new cluster resource,” on page 362

- Section 15.26, “Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server,” on page 363
- Section 15.27, “Communications Errors,” on page 363
- Section 15.28, “Connections are not Getting Established if a Cluster Resource is Migrated from OES 2015 or Later to OES 11 SP2 Node,” on page 364

15.1 Migration or Online of Resources Causes Cluster Crash

If the split-lock detection is enabled by default on the cluster node when the resource is migrated or bought online, the cluster crashes.

By permanently disabling the `split_lock_detect`, this problem can be fixed:

1. Log in to the node as the root user, then open a terminal console.
2. Verify if `split_lock_detect` is enabled.

   ```
   cat /proc/cpuinfo | grep split_lock_detect
   ```

   The command will display `split lock detection: enabled`.
3. Open the file `/etc/default/grub`.
4. Locate `"GRUB_CMDLINE_LINUX_DEFAULT="` and add `"split_lock_detect=off"` at the end of the line.

   ```
   "GRUB_CMDLINE_LINUX_DEFAULT=<…> split_lock_detect=off"
   ```
5. Update the bootloader.

   ```
   update-bootloader -refresh
   ```
6. Reboot the cluster node.
7. Verify if `split_lock_detect` is disabled by running any one of the following commands:
   - `journalctl -b |grep split`
   - `dmesg -T`

   The command will display `split lock detection: disabled`.

15.2 Cluster Node Crashing Issue

If you stop the cluster service on a cluster node while running **Update 2 OES 2023 Patch** or the latest kernel version, the cluster node will crash. To avoid the cluster node from crashing, follow the below steps:

1. Log in to the node as the root user, then open a terminal console.
2. Disable the Novell Cluster Service.

   ```
   systemctl disable novell-ncs.service
   ```
3. Remove the node from the cluster.

   ```
   cluster leave
   ```
4. Ensure that all cluster resources are moved to other cluster node.
5. Stop the Novell Cluster Service.

```
systemctl stop novell-ncs.service
```

NOTE: Stopping the Novell Cluster Service leads to cluster node crash.

6 Reboot the server.

7 Enable the Novell Cluster Service on the cluster node.

```
systemctl enable novell-ncs.service
```

8 Apply the **Update 3 OES 2023 Patch** on the cluster node.

```
zypper patch
```

9 Reboot the server.

15.3 Unable to Manage NCS from iManager When NSS is not Installed on the Node

If NSS or iManager is not installed on the master node and you try to manage NCS using iManager, iManager fails to retrieve the cluster status with an error. This is because the `novell-nss-admin-session-sfcb-provider` rpm is not installed on the node where NCS is running. To resolve this issue, install `novell-nss-admin-session-sfcb-provider` rpm on all the nodes.

NOTE: The `novell-nss-admin-session-sfcb-provider` rpm is automatically installed when NSS or iManager is installed.

15.4 Cluster Service Does Not Come Up After Upgrading to OES 2018 or Later if Service Proxy is Configured

Cluster service configured with service proxy fails to come up after upgrading to OES 2018 or later. This is because the service proxy users are not migrated to OES Credential Store (OCS). To resolve this issue, perform the following:

1 Login as root user.

2 Run the following command:

```
/opt/novell/ncs/bin/ncs_update_proxy_cred
<Proxy_user_with_full_context> <Proxy_user_password>
```

For example,

```
/opt/novell/ncs/bin/ncs_update_proxy_cred cn=proxyuser1,o=novell
proxypass1
```

3 Verify the service entry is present in OES Credential Store by using the following command:

```
oescredstore -l
```

4 Start NCS by using the following command:

```
systemctl start novell-ncs.service
```

15.5 Supported SBD Signatures in Cluster Nodes with Different OES Versions

Based on the OES versions running in cluster nodes, the following are the three types of SBD signatures available:

- If the node is running on OES 11, "SBD*" is used as SBD signature to support Hardlink media.
- If the node is running on OES 2015, "SBD)" is used as SBD signature to support AD media.
- If the node is running on OES 2015 SP1, "SBD(" is used as SBD signature to support Trustee Index media.
- If the node is running on OES 2018, "SBD’ " is used as SBD signature.

15.6 Cluster Resource Is Not Accessible After a Network Service Restart

If you restart the network services, cluster resource IP addresses will get removed from the network interfaces. As a result, the cluster resources will not be accessible through those IP addresses. Network service restart can happen as part of LAN settings reconfiguration through YaST or through the `rcnetwork restart` command. In order to recover from this state, follow the steps mentioned below:

1. Run the `cluster leave` command to remove the node from the cluster. This will cause the cluster resources running on this node migrate to another node in the cluster.
2. Run the `cluster join` command to rejoin the node to the cluster.
3. (Optional) You can also manually migrate the resources back to this node by running the cluster migrate command.

15.7 File Location Information

Information about OES Cluster Services scripts and configuration can be found in the following locations. Ensure that you use the Clusters plug-in for iManager to configure cluster and resource settings and to edit the load, unload, and monitor scripts. You can view cluster events by using the Clusters plug-in for iManager. Go to **Clusters > My Clusters**, select the cluster, then click **Event Log**.

Information	File Location
Load scripts Unload scripts Monitor scripts Cluster Services configuration files	`/var/opt/novell/ncs`
Load and unload script output files (.out) Most resource load and unload issues can be solved with the information in the .out files.	`/var/opt/novell/log/ncs`
Cluster Services scripts, such as ldncs and unldncs	`/opt/novell/ncs/bin`

Information	File Location
Resource templates	`/opt/novell/ncs/templates`

15.8 Version Issues

Knowing the location and purpose of the files that make up OES Cluster Services can be useful in helping you troubleshoot problems and resolve version issues. See Appendix B, “Files for OES Cluster Services,” on page 407.

15.9 Diagnosing a Poison Pill Condition

A poison pill is given to a node if it stops sending out the OES Cluster Services heartbeat packages, including the case when the node hangs or reboots.

To evaluate the poison pill condition on the node, look at the `/var/log/messages` log from the node that rebooted (was given the poison pill). Check the messages right before the restart. Normally, you can spot the process that caused the server to hang or reboot.

You can run the `cluster stats display` command on the surviving master node to show when the OES Cluster Services heartbeat packages stopped coming from the rebooted node. To be sure that is the case, you can also run a LAN trace to record the packages among the nodes.

Other events that might cause a pause in sending out heartbeat packages include the following:

- Antivirus software scanning the `/admin` file system
- Network traffic control (including the firewall)
- `check_oesnss_vol.pl` running under Nagios
- Packages that have been recently updated through the YaST Online Update or patch channels

15.10 Diagnosing Cluster Problems

To help OES Technical Support diagnose your cluster problems, some or all of the following information might be requested:

- Cluster statistics. A report created by the cluster statistics gathering tool.
- Cluster configuration. An HTML report created by the **Action > Run Report** option on the **Clusters > My Clusters** page of the Clusters plug-in to iManager.
- SAN configuration:
  - Host bus adapter, hub or switch type
  - Device driver name and revision
  - Storage controller type
  - Storage controller firmware revision
- SBD configuration - Single or mirrored partition?

- LAN configuration:
  - Network interface card, hub, or switch type
  - Device driver name and revision
  - Dedicated heartbeat or shared public LAN?
- Server configuration:
  - Type, memory, and number of CPUs
  - Software revisions and patches
  - List of RPMs
  - `/opt/novell/ncs/bin/ldncs` file
- LAN packet trace
- Console log files
- Abend log files
- Server coredump files
- Observable Coincidences? For example:
  - Are poison pills experienced at the same time of day as scheduled maintenance activities?
  - Are server inventory or health checks, virus scans, backups, or periodic eDirectory activity?

15.11 Master IP Address Problem

If the master IP address has a problem, such as a duplicate IP address on the network, the Clusters plug-in might not be able to connect to the master IP address to allow you to manage the cluster. The My Clusters page uses the IP address found in the Cluster object by default. If the master IP address does not respond, you cannot connect to the master IP address to manage the cluster in iManager.

In OES 11 SP2 and later, you are prompted to use the server IP address of a specified node to manage the cluster while you fix the problem. See Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118.

15.12 Why didn’t my resource go to the node I specified?

In a failover situation, the cluster looks for the best solution for the node. It goes to the most preferred node if it can. If that node is not available, it attempts to go to the next node in the resource’s preferred nodes list. It considers if the node is available and if it has an Resource Mutual Exclusion rules that must be obeyed for resources that are already mounted on the target node. If all of the nodes in the resource’s preferred nodes list are not available or have RME conflicts, the resource goes comatose instead of failing over to one of the nodes in its preferred nodes list.

In a cluster migration situation, the resource is currently online and you specify a target node. The resource cannot go to the specified node if any of the following conditions exist, and it stays online on the original node. You can view the target node’s log to understand why the resource was not available.

- The specified node is not in the resource’s preferred nodes list.

- The specified node is in the resource's preferred nodes list, but the node is currently running another resource that is a Resource Mutual Exclusion conflict for the resource you are attempting to migrate.
- The specified node is in the resource's preferred nodes list, but the node is not available.
- The specified node is in the resource's preferred nodes list, but the node does not support the NSS media on the pool resource.

NOTE: The cluster resource migration fails with the following error in `/Var/log/messages`.

```
POOL_REVISION_CHECK (sfcbd:0:13): Cannot online resource '<Resource-
Name>' on node '<Node_Name>', because NSS on the node may not understand
the newer NSS media associated with the resource
```

15.13 iManager: Error 500 Occurs while Connecting to a Cluster

iManager might hang or time out while trying to connect to a server with the Cluster plug-in or the Storage plug-in. Error 500 is reported. This error can occur if your login user name for iManager is not a user that is enabled with Linux User Management (LUM) on the target server.

To avoid this problem, ensure that the user name is LUM enabled on the target server. The Cluster plug-in connects to the master node in the cluster (that is, the node hosting the `Master IP Resource`).

For more information, see *Technical Information Document: iManager error 500 while trying to managing cluster, storage, user quotas, CIFS, or AFP (TID 7010295)* (http://www.novell.com/support/kb/doc.php?id=7010295) in the Micro Focus Knowledgebase (http://www.novell.com/support/).

15.14 iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5

After installing or updating the Clusters plug-in to version 3.3.1 for iManager 2.75, you get a Status Code 500 error whenever you attempt to manage the cluster from iManager.

This problem can occur if iManager loads a version of the class loader file earlier than `/var/opt/novell/iManager/nps/WEB-INF/lib/commons-lang-2.6.jar`. iManager uses one class loader for all of its plug-ins. If multiple versions of the class loader file exist in the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory, iManager loads whichever one comes first alphabetically.

To resolve this problem, you must remove the earlier versions (such as `commons-lang.jar` or `commons-lang-2.4.jar`) to ensure that iManager loads the 2.6 or later version of the class loader file, then restart Tomcat.

For more information, see Section 5.7, "Installing or Updating the Clusters Plug-in for iManager," on page 72.

This problem has been resolved in the Cluster plug-in for iManager 2.7.6 and later.

15.15 iManager: Problems with Clusters Plug-In After Update for iManager 2.7.5

The Clusters plug-in for iManager 2.7.5 in OES 11 SP1 was reorganized to display two tasks in the left panel: **My Clusters** and **My Resources**. If you have problems displaying or using the plug-in after updating to this version, ensure that you performed the following tasks after the update:

- Step 5 in Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 73
- Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later," on page 74

15.16 A Device Name Is Required to Create a Cluster Partition

If you are planning to work with shared-disk NSS pools and volumes, you must create a Split Brain Detector (SBD) partition when you configure the first node on the cluster by selecting the name of a device (such as `sdc`) that is already initialized and marked shareable for clustering. If you don't enter a device name, you must manually create an SBD before you configure clustering on a second node and before you attempt to create storage objects or cluster-enable NSS pools.

15.17 Cluster Resource Goes Comatose After Migration or Failover

When the SLP daemon (`slpd`) is not installed and running on a cluster node, any cluster resource that contains the `ncpcon bind` command goes comatose when it is migrated or failed over to the node because the bind cannot be executed without SLP.

See "SLP" on page 43.

When the backup of cluster NCP POSIX Volumes is in progress, the cluster resource goes comatose when it is migrated or failed over to another node.

15.18 Cannot Authenticate to Remote Servers during Cluster Configuration

During the cluster installation and configuration, if you choose **Remote System** on the OES Cluster Services LDAP Configuration page and you have LDAP configured to point to a NetWare 6.0 or earlier NetWare server, the cluster configuration fails.

To work around this problem, you must edit the `/etc/openldap/ldap.conf` file. Either disable certificates (`TLS_REQCERT <level>` line) or change the file that contains the certificates (`TLS_CACERT <filename>` line). See the `ldap.conf` man page for more information.

15.19 Cannot Connect to an iSCSI Target

If you are connecting to an iSCSI target that already has NSS partitions and pools created on it, you might not be able to access those NSS partitions and pools until you reboot the Linux initiator server. This is required for each Linux initiator server that will access the iSCSI target.

For instructions on configuring an OES 11 or later server as an iSCSI initiator and connecting to an iSCSI target, go to "Mass Storage over IP Networks--iSCSI" (https://documentation.suse.com/sles/12-SP5/html/SLES-all/cha-iscsi.html) in the *SLES 12 Storage Administration Guide.* (https://documentation.suse.com/sles/12-SP5/html/SLES-all/stor-admin.html)

15.20 Is There a way to Uninstall OES Cluster Services from a Server?

There is no utility to remove OES Cluster Services and its related eDirectory objects. See Section 5.10, "Removing a Node from a Cluster," on page 77.

15.21 Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems

Running the `supportconfig` command concurrently on multiple nodes can cause the multiple nodes to lose communications with the others and go down.

Micro Focus Support provides the `supportconfig` package to customers in order to gather configuration, log, and error information from both the installation of the OS and from varying applications. The `supportconfig` command was designed without clustering considerations. The command is normally run after a server experiences problems. It is not a good idea to run the command concurrently on multiple nodes in a healthy cluster. If you need to run the `supportconfig` command on multiple nodes in a cluster, run it on one node at a time.

15.22 Error NFS Stale Handle Reported on Host when Cluster Failover Occurs

The host I/O reports an `NFS Stale Handle` error immediately after failover is performed on the cluster.

This problem can occur if a user attempts to access the IP address of the resource before the volume is mounted or after the volume is dismounted. You can modify the order of commands in a load script to: NSS, NFS, IP, NCP, and CIFS. Reverse the order in an unload script. This allows storage to be mounted before the IP address is bound, and to be dismounted after the IP address is unbound. Adjust your scripts to this order, then take the resource offline and bring it online to apply the changes. For sample scripts, see Section 12.10, "Adding NFS Export for a Clustered Pool Resource," on page 228.

15.23 Error 20897 - This node is not a cluster member

If OES Cluster Services is installed on a node, but an SBD does not exist, NSS tools (including OES Linux Volume Manager (NLVM) commands, NSSMU, NSS console commands, and the Storage plug-in for iManager) return the following error:

```
Error 20897 - This node is not a cluster member.
```

In a OES Cluster Services cluster, NSS tools use the cluster's SBD to detect if a node is a cluster member and to lock against concurrent changes to physically shared storage. Without an SBD, NLVM cannot detect whether a node is a member of the cluster and cannot acquire the locks it needs to execute tasks. In this state, you can use the -s option with NLVM commands to prepare a device and create an SBD partition. To minimize the risk of corruption, you must ensure that nobody else is changing any storage on any nodes at the same time.

For information about creating an SBD partition, see Section 9.18, "Creating or Deleting Cluster SBD Partitions," on page 146.

15.24 Error 499 - Could not delete this resource Data_Server

When you try to delete a cluster resource, you get a message "Could Not Delete This Resource Data_Server (Error 499). The error means that the object does not exist in eDirectory. This error can occur if you try to delete an orphaned cluster resource. An orphaned cluster resource does not have a pool or volume associated with it, does not have an IP address, and the scripts are empty.

To delete an orphaned cluster resource:

1 On the master node of the cluster, open a terminal as the `root` user. At the command prompt, enter the following commands:

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

and if the resource is still there,

```
echo -n "exec rm -f /var/opt/novell/ncs/data_server.*" > /proc/ncs/
cluster
```

15.25 Error 603 - smdr.novell is not registered with SLP for a new cluster resource

You might get an error after you create a cluster resource, indicating that `smdr.novell` is not registered with SLP for cluster resources, but the `smdr.novell` service for the node is registered.

```
Error: "cluster--<212>: Read ResVol error -603"
```

The first time a cluster resource is created, smdrd cannot figure it out. The smdrd service is designed to work with resources that contain a data volume.

Ensure that the resource contains a volume and is configured properly, then restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

1 Log in to the server as the `root` user, open a terminal console, then enter

```
rcnovell-smdrd restart
```

15.26 Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server

You might get a Migration Authentication error when you attempt to connect to a newly created cluster resource on the target server.

```
2013-05-03 12:34:20,095 INFO  - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_IP_address>,
trying with DNS name <resource_DNS_name>
2013-05-03 12:34:20,096 INFO  - Migration Framework:Authentication:
2013-05-03 12:34:20,096 INFO  - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_DNS_name>
```

There are two possible causes:

- The Storage Management Service (SMS) `smdr.novell` service running on the target server is not aware of the newly created resource.

  The Migration tool uses SMS to transfer data between nodes. When a cluster resource is first created, the `smdrd` service is not automatically aware of the new resource. See Section 15.25, “Error 603 - smdr.novell is not registered with SLP for a new cluster resource,” on page 362.

  Restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

  1. Log in to the target server as the `root` user, open a terminal console, then enter

     ```
     rcnovell-smdrd restart
     ```

- The cluster resource does not contain a volume.

  A properly created pool cluster resource or LVM cluster resource contains at least one volume. Even if you have restarted the `smdrd` service, the service by design does not recognize a resource unless it contains a data volume.

  Ensure that the resource contains a volume and is configured properly, then restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

  1. Create a volume on the pool or LVM volume group, update the resource scripts as needed, then take the resource offline and bring it online to apply the changes.
  2. Log in to the target server as the `root` user, open a terminal console, then enter

     ```
     rcnovell-smdrd restart
     ```

15.27 Communications Errors

When you manage a cluster or cluster resources with the Clusters plug-in for iManager, CIMOM errors can occur for the following reasons:

- The CIMOM broker daemon must be running and listening on port 5989. See the Section 4.6.8, “SFCB and CIMOM,” on page 46.
- You must be able to connect to the cluster master node. Verify the ability to connect to the cluster master node by pinging the cluster IP address from a terminal console on a different computer.

If there is a conflict for the master IP address, you can modify the master IP address only by using the Clusters plug-in for iManager as described in Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118. After you fixed the conflicting IP address problem, you must close the browser and restart the cluster, and then open a browser and log in to iManager again in order to continue to manage the cluster.

- The cluster IP address must be the same in the **NCS:Network Address** attribute and **Network Address** attribute for the Cluster object. Addresses might not match if you modified the master IP address for the cluster and did not restart the cluster.

Other possible issues are described in the Cluster Communications Error message (see Figure 15-1).

Figure 15-1 *Cluster Communications Error*

15.28 Connections are not Getting Established if a Cluster Resource is Migrated from OES 2015 or Later to OES 11 SP2 Node

Description: In a mixed-node cluster environment, that is, OES 2015 or later nodes and the nodes running earlier versions of OES (OES 2 SP3, OES 11, OES 11 SP2), if you migrate a cluster resource from OES 2015 or later to earlier versions of OES, the existing connections accessing that resource might not establish automatically.

Cause: The client maintains a runtime cache for each server with which it communicates, whose cached entry indicates the protocol supported by the server.

When a cluster resource is on an OES 2015 or later node, the client maps the resource using SMB 2 protocol. If the same resource is migrated to earlier versions of OES, the client assumes that this server is also SMB 2 protocol capable and without negotiating with the server attempts to reconnect using SMB 2, which the node rejects since it's not SMB 2 protocol capable.

Action: To resolve this problem, the user must login to the server running earlier versions of OES afresh to access the resource. When doing a fresh mapping, the client first enumerates and negotiates the server's capability and then connects using the highest protocol that the server is capable of.

16 Security Considerations

This section describes security issues and recommendations for OES Cluster Services for Open Enterprise Server (OES) 2023. It is intended for security administrators or anyone who is responsible for the security of the system. It requires a basic understanding of OES Cluster Services. It also requires the organizational authorization and the administrative rights to effect the configuration recommendations.

- Section 16.1, “Cluster Administration Rights,” on page 365
- Section 16.2, “Ports,” on page 365
- Section 16.3, “Email Alerts,” on page 366
- Section 16.4, “Log Files,” on page 366
- Section 16.5, “Configuration Files,” on page 366

16.1 Cluster Administration Rights

For information about the rights needed to install and manage OES Cluster Services, see the following sections:

- “NetIQ eDirectory 9.2.6” on page 40
- “OES Domain Services for Windows” on page 50
- “SFCB and CIMOM” on page 46
- “SLP” on page 43
- “eDirectory Tree” on page 41
- “Cluster Installation Administrator” on page 35
- “Cluster Administrator or Administrator-Equivalent User” on page 38

16.2 Ports

For each cluster, you can specify the port used for cluster communication. The default cluster port number is 7023, and is automatically assigned when the cluster is created. You might need to modify this value if there is a port conflict. See Section 8.12, “Viewing or Modifying the Cluster Master IP Address or Port,” on page 118.

If you enable Email Notification, Cluster Services uses Postfix to send messages. The default port number is 25. For information about how to address port conflicts, see Section 8.6.1, “Configuring an Email Port,” on page 107.

You must specify the eDirectory port when you install clusters. Port 636 is the default port number for eDirectory communications in the tree.

If you are using a firewall, the port must be opened for CIMOM communications between the cluster and iManager. Port 5989 is the default setting for secure HTTP (HTTPS) communications.

16.3 Email Alerts

You can enable or disable email notification for the cluster and specify up to eight administrator email addresses for cluster notification.

OES Cluster Services uses Postfix to send email alerts. If you have a cluster resource that uses SMTP, that resource might not work in the cluster unless you change the Postfix configuration. See Section 8.6, “Configuring Cluster Event Email Notification,” on page 106.

16.4 Log Files

NCS normally writes log messages to `/var/log/messages`.

NCS installation writes log messages to `/var/opt/novell/install/ncslog`.

NCS configuration writes log messages to `/var/log/YaST2/y2log`.

You can view cluster events logged in the `/admin/Novell/Cluster/EventLog.xml` file by using the Clusters plug-in for iManager. Go to **Clusters > My Clusters**, select the cluster, then click **Event Log**. See Section 9.6, “Viewing the Cluster Event Log in iManager,” on page 135.

You can also copy the `/admin/Novell/Cluster/EventLog.xml` file to a working location, and then use the `cat` command to display the file. See Section 9.7, “Viewing the Cluster Event Log at the Command Line,” on page 136.

NCS writes cluster resource runtime messages to files under `/var/opt/novell/log/ncs/`. Output log files are named with the resource name, the type of script (load, unload, or monitor), and the `.out` extension (`<resource_name>.<load|unload|monitor>.out`). For example: `/var/opt/novell/log/ncs/myresource.unload.out` reports the log for the `myresource` unload script messages.

At `/etc/logrotate.d/novell-ncs-monitor-log`, you can configure settings for log rotation. By default, the `monitor.out` files are rotated every mid-night. The `monitor.out` files are compressed and stored at `/var/opt/novell/log/ncs/archives`. A file once rotated is saved in the system for 365 days. The Administrator has to manage disk usage.

In OES 11 SP2 and later, you can print summaries of failed or incomplete events in the `*.out` log files by copying the file to a working location and using the DotOutParser utility. See Section A.5, “DotOutParser Utility,” on page 385.

16.5 Configuration Files

You should never modify configuration files directly, unless you are directed to do so by a documented procedure or Micro Focus Support.

The cluster configuration file is `/etc/opt/novell/ncs/clstrlib.conf`.

The cluster configuration for a specific node is `/etc/opt/novell/ncs/<nodename>`.

The cluster node numbers are found in the `/var/opt/novell/ncs/gipc.conf` file.

The OES Cluster Services reboot behavior after shutting down a node conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown. You can use directive `kernel.panic` in the `/etc/sysctl.conf` file to allow an automatic reboot and to specify the number of seconds to delay the reboot. See Section 8.9, “Configuring the Cluster Node Reboot Behavior,” on page 114.

If you use Device Mapper Multipath for devices, ensure that you modify the settings in the `/etc/modprobe.conf.local` file and the `/etc/multipath.conf` file, so that multipath I/O works correctly for OES Cluster Services cluster resource fail over and fail back. See Section 4.10, “Multipath I/O Configuration Requirements,” on page 54.

A Console Commands for OES Cluster Services

OES Cluster Services provides several console commands to help you perform certain cluster-related tasks.

- Section A.1, “Cluster Management Commands,” on page 369
- Section A.2, “Business Continuity Clustering Commands,” on page 378
- Section A.3, “OES Installation extend_schema Command,” on page 379
- Section A.4, “SBD Utility,” on page 380
- Section A.5, “DotOutParser Utility,” on page 385
- Section A.6, “CSMPORT Utility (Cluster Segment Manager Import/Export),” on page 386
- Section A.7, “ncs-configd.py Script (Pushing Configuration Updates to All Nodes),” on page 394
- Section A.8, “ncs_install.py Script (Extending the Schema, Monitoring NDSD, or Removing NCS Configuration),” on page 394
- Section A.9, “ncs_ncpserv.py Script (Creating an NCP Virtual Server Object for a Clustered LVM Volume Group),” on page 399
- Section A.10, “ncs_resource_scripts.pl Script (Modifying Resource Scripts Outside of iManager),” on page 401
- Section A.11, “NCPCON Mount Command for NCP and NSS Volumes,” on page 403
- Section A.12, “Querylog Utility,” on page 405

A.1 Cluster Management Commands

To execute a cluster console command, enter `cluster` followed by the command. For example, if you want to display cluster statistics, enter `cluster stats display` at the terminal console. You can also enter `cluster help` at the command prompt to get information on the commands and their functions.

The functions of many of the commands can also be performed using iManager. See the other sections of this document for additional information.

Table A-1 lists the cluster-related terminal console commands and gives a brief description of each command.

Table A-1 *Cluster Console Commands*

Cluster Console Command	Description
ALERT {resource}{YES|NO}	The resource start, failover, or failback mode is set to manual and the resource is waiting for administrator intervention to start on a node, or to fail over or fail back to another node. Alert types are Start Alert, Failback Alert, and Failover Alert. For the Start Alert, you must address the alert, then offline and online the resource. Specify the resource name in the command and use the YES or NO switch to accept or reject an alert on a resource to control fail over, fail back, or manual start. **Example** `cluster alert res1 yes`

CONVERT preview *[resource]*
 	
CONVERT preview
 	
CONVERT commit
	Finalizes the cluster conversion from NetWare to Linux after all nodes in a mixed cluster have been converted to Linux. The CLUSTER CONVERT command can be executed only on Linux cluster nodes.
	
	Specify a resource name with the Preview option to view the resource load and unload script changes prior to finalizing the conversion. If you do not provide a resource name, it displays the information for all resources.
	
	Use the Commit switch without specifying a resource to finalize the conversion for all cluster resources.
	
	Examples
	
	`cluster convert preview res1`
	`cluster convert preview`
	`cluster convert commit`

DOWN	Removes all cluster nodes from the cluster. This command has the same effect as executing the CLUSTER LEAVE command on every server in the cluster. You are prompted to confirm the cluster down. **Example** `cluster down`
EXEC "path_to_script"	Executes the specified script on all nodes in the cluster. **Example** `cluster exec "command"`

Cluster Console Command	Description
`INFO {All, Basic, Notification, Priority, Protocol, Summary}`	Displays information on cluster configuration. **Options** **`all`** Displays a combination of Basic, Notification, Priority, and Protocol information. **`basic`** Displays IP address, port, and cluster quorum settings. **`notification`** Displays cluster e-mail notification settings. **`priority`** Displays the resource priority list. **`protocol`** Displays the cluster protocol settings. **`summary`** Displays the cluster protocol summary. **Example** `cluster info protocol`
`JOIN`	Adds the node where the command is executed to the cluster and makes the node visible to other servers in the cluster. OES Cluster Services software must already be installed and running on a node for it to join the cluster. **Example** `cluster join`
LEAVE	Removes the node where the command is executed from the cluster. The node will not be visible to other servers in the cluster. **Example** `cluster leave`

Cluster Console Command	Description

`MAINTENANCE {ON|OFF}{EXPIRY_TIME}`	Turning this switch on lets you temporarily suspend the cluster heartbeat while hardware maintenance is being performed. This is useful if you want to reset or power down the LAN switch without bringing the cluster servers down.
	
	Turning this switch on from one cluster server puts the entire cluster in maintenance mode.
	
	Running the command without the {ON|OFF} parameter reports the maintenance status of the cluster (that is, whether maintenance mode is on or off).
	
	An expiry time can be set to automatically turn off the cluster maintenance mode. The specified time can be in any format, which is supported by the Linux 'date' utility. The ‘expiry_time’ option is applicable only when all the cluster nodes are on OES 2023.
	
	Examples
	
	`cluster maintenance on`
	`cluster maintenance`
	
	`cluster maintenance on '2021-04-04 12:34:51'`
	`All the cluster nodes will be in the maintenance mode till April 4,2021; time - 12:34:51. The cluster will be out of maintenance mode automatically at the specified time.`

Cluster Console Command	Description

`MIGRATE {resource} { <node_name> | -list | -most | -next }`	Migrates the specified resource from the node where it is currently running to the node that you specify in the command. The node that you specify must be running in the cluster and also be in the resource's **Preferred Nodes** list.
	
	Options
	
	<node_name>
	
	Migrates the resource to the specified node if possible. No action is taken if the specified node is not in the resource's preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.
	
	-l, -list
	
	Shows the preferred nodes list for the specified resource.
	
	-m, -most
	
	Migrates the specified resource to the most preferred node currently in the cluster. If the resource is running on such a node, no action is taken.
	
	-n,-next
	
	Migrates the resource to the node in the preferred node list that is next in order to the node where the resource is currently running. If there is no next node or all such nodes are not in the cluster currently, it searches from the beginning of the resource's preferred nodes list. No action is taken if a new destination could not be determined after the search is exhausted.
	
	Examples
	
	Move res1 from the current node to node2 in the resource's preferred nodes list. The resource is migrated only if there are no Resource Mutual Exclusion conflicts on the target node.
	
	`cluster migrate res1 node2`
	
	List the preferred nodes for resource POOL_14_SERVER.
	
	`cluster migrate POOL_14_SERVER -l`
	
	`Preferred nodes for resource 'POOL_14_SERVER'`
	`  CG-06`
	`  CG-05`
	`  CG-04`
	`  CG-03`
	`  CG-02`
	`  CG-01`
	`  CG-08`
	`  CG-07`
	
	`Status for Resource: POOL_14_SERVER`
	`Running on CG-05  Lives: 65`
	`Revision: 5`

Cluster Console Command	Description
`MONITOR {resource} <start | stop | status>`	Manually starts, stops, or checks the status of resource monitoring. The resource must be running on the node where you issue the command. Resource monitoring must be enabled for the resource in order to use this command. **Example** `cluster monitor sh_pool_res_01 status`
`OFFLINE {resource}`	Unloads the specified resource from the node where it is currently running. **Example** `cluster offline res1`

`ONLINE {resource}{node name}`	Starts the specified resource on the most preferred node that is currently active. You can start the resource on a different node by specifying that node in the command. The node that you specify must be running in the cluster and also be in the resource's **Preferred Nodes** list.
	
	Example
	
	`cluster online res1`
	`cluster online res1 node1`

`POOLS`	Lists the NSS pools on the shared disk system that are accessible by OES Cluster Services. **Example** `cluster pools`
`RENAME <old_resource_name> <new_resource_name>`	Renames a pool cluster resource. This command must be issued from the master node. The resource must be in offline state to be renamed. The new name must not exist prior to the renaming. Renaming the resource does not modify the resource's virtual server name (NCS:NCP Server object). You can delete and re-create the NCS:NCS Server object with a new name by using iManager. **Example** `cluster rename POOL1_SERVER custom_name22`

Cluster Console Command	Description
RESOURCES <-i | -v | -c | -p | -u | -a>	Lists the information about the resources that currently exist in the cluster. The resources do not need to be online or running. **Options** **-i** Displays the IP address of all resources in the cluster. **-v** Displays the virtual server name of all resources in the cluster. **-c** Displays the CIFS server name of all resources in the cluster. **-p** Displays the preferred nodes of all resources in the cluster. **-u** Displays the unassigned nodes of all resources in the cluster. **-a** Displays the IP address, virtual server name, CIFS server name, preferred nodes, and unassigned nodes of all resources in the cluster. **Example** `cluster resources -i`
RESOURCE {resource}	Displays the IP address, virtual server name, CIFS server name, preferred nodes, and unassigned nodes of the specified resource. **Example** `cluster resource res1`
PREFERRED_NODES {resource}	Displays a list of preferred nodes for the specified resource. **Example** `cluster preferred_nodes res1`
UNASSIGNED_NODES {resource}	Displays a list of nodes in the cluster that are not assigned to the specified resource. **Example** `cluster unassigned_nodes res1`

Cluster Console Command	Description
`RESTART [seconds]`	Restarts OES Cluster Services software on all servers in the cluster. The `cluster leave` process begins immediately. Specify a value of 60 or more seconds as the time to wait before the `cluster join` begins. The default setting is 60 seconds. The default setting is used if a value is not provided, if the value provided is smaller than 60 seconds, or if the input is invalid. You are prompted to confirm the cluster restart. **Example** `cluster restart 60`
`SCAN FOR NEW DEVICES`	**IMPORTANT:** This command is deprecated on Linux. You can use the `nlvm rescan` command on OES 11 and later. You can alternatively use the `rescan-scsi-bus.sh` script to scan for the new devices on Linux without rebooting. See "Scanning for New Devices without Rebooting" (https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-multipath.html#sec-multipath-best-practice-scandev) in the *SLES 12 Storage Administration Guide* (https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-storage.html). **WARNING:** In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

Cluster Console Command	Description
`SET {Parameter} {Value}`	Sets cluster parameters individually for the cluster. See Chapter 8, "Configuring Cluster Policies, Protocols, and Properties," on page 95 for more information on cluster parameters. Specify one of the following parameters and a value for that parameter: `IPADDRESS` sets the cluster master IP address to the specified value. If you change the cluster IP address, you must restart cluster software on all cluster nodes. `PORT` sets, or lets you change, the IP port number of the cluster master. `QUORUMWAIT` sets the amount of time in seconds that the cluster waits before resources start to load. `QUORUM` sets the number of nodes that must be running in the cluster before resources will start to load. `HEARTBEAT` sets the amount of time in seconds between transmits for all nodes in the cluster except the master. `TOLERANCE` sets the amount of time in seconds that the master node gives all other nodes in the cluster to signal that they are alive. `MASTERWATCHDOG` sets the amount of time in seconds between transmits for the master node in the cluster. `SLAVEWATCHDOG` sets the amount of time in seconds that the slave nodes give the master node in the cluster to signal that it is alive. `MAXRETRANSMITS` sets the maximum number of times transmits will be attempted between the master node and slave nodes. `ENABLEEMAIL` enables and disables email notification. You can set the value to OFF to disable email notification, or either CRITICAL or VERBOSE to enable email notification. `EMAILADDRESSES` lets you specify the email addresses used for email notification. The addresses should be separated by spaces. Using this parameter without specifying any addresses clears existing addresses that have been set previously. `EMAILOPTIONS` sets the email notification options. Specify XML as the value to receive email notification in XML format. Not specifying any value with this parameter turns notification in XML format off. `RESOURCEPRIORITY` sets the resource priority for the cluster. **Example** `cluster set ipaddress 10.1.1.1`
`STATS {Display, Clear}`	The Display parameter reports the node number, node name, and heartbeat information to the console screen. For more information, see Section 8.5.1, "Heartbeat," on page 105. The Clear parameter resets the various stats counters. This command is useful for troubleshooting network problems related to the cluster heartbeat protocol. **Example** `cluster stats display`

Cluster Console Command	Description
`STATUS {resource}`	Reports the status of the specified resource. This includes the number of times the resource has been migrated or failed over to another server, the resource state, and the node where the resource is currently running. **Example** `cluster status res1`
`VIEW`	Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster. **Example** `cluster view`
`SCRIPT {load-script | unload-script | monitor-script | all} {resource}`	Displays the scripts for the specified resource. **Options** **`load-script`** Displays the load script for the specified resource. **`unload-script`** Displays the unload script for the specified resource. **`monitor-script`** Displays the monitor script for the specified resource. **`all`** Displays load, unload, and monitor scripts for the specified resource **Example** `cluster script load-script res1`
`RESOURCE-PROTOCOL {resource}`	Displays the set value for the VirtualServerName, CifsServerName, and IPAddress of the specified resource. **Example** `cluster resource-protocol res1`
`RESOURCE-POLICY {resource}`	Displays the set value for the start mode, failover mode, and failback mode on the specified resource. **Example** `cluster resource-policy res1`

A.2 Business Continuity Clustering Commands

If OES Business Continuity Clustering (BCC) is installed and running on the cluster, you can use the following additional `cluster` commands to manage the BCC cluster environment.

```
connections
credentials
disable
enable
nsmi
resetresources
```

For information about using these commands, see Console Commands for BCC in the OES 2023: BCC Administration Guide.

A.3 OES Installation extend_schema Command

A tree administrator user with credentials to do so can use the OES Installation `extend_schema` command to extend the eDirectory schema before a cluster is installed anywhere in a tree. This allows container administrators (or non-administrator users) to install a cluster in a container in that same tree without needing full administrator rights for the tree. You need to extend the schema only one time in the tree where you will be installing clusters.

- Section A.3.1, “Syntax,” on page 379
- Section A.3.2, “Options,” on page 379
- Section A.3.3, “Example,” on page 380

A.3.1 Syntax

```
/opt/novell/oes-install/util/extend_schema --port <port_num> <admin_username>
<admin_password> <ldap_server_ip_address> <schema_file>
```

To extend the schema, the tree administrator user modifies the following schema files in the given order:

`/opt/novell/ncs/schema/ncs.ldif`

`/opt/novell/ncs/schema/ncpserver.preldif`

A.3.2 Options

Replace the parameters with the credentials to access and location of the eDirectory schema files.

Parameter	Description	Example
port_num	The port number you assigned for eDirectory communications in the tree where you plan to install clusters. The default port is 636.	636
admin_username	The typeful fully distinguished user name of the administrator who has the eDirectory rights needed to extend the schema.	`cn=admin,o=example`
admin_password	The password of the administrator user.	`pas5W0rd`

Parameter	Description	Example
ldap_server_ip_address	The IP address of the eDirectory server that contains the schema files.	`10.10.10.1`

A.3.3 Example

For example, enter the following commands in the order shown, using the values for your particular solution:

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example
pas5W0rd 10.1.1.1 /opt/novell/ncs/schema/ncs.ldif
```

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example
pas5W0rd 10.1.1.1 /opt/novell/ncs/schema/ncpserver.preldif
```

A.4 SBD Utility

The SBD utility (sbdutil) allows you to create, find, or view a OES Cluster Services SBD partition.

IMPORTANT: The cluster SBD partition is not required unless you have shared storage in the cluster.

We recommend that you carve out a LUN/disk of 20 MB in size for 512 byte device and 80 MB for 4096 (4Kn) byte device to use for the SBD. If you mirror the SBD, you need to carve out a separate LUN/disk of equal size to use. Before you begin, each of these small disks must be initialized and marked as shareable for clustering. You can initialize the device by using the OES Storage Services (NSS) Management Utility (`nssmu(8)`) or the Storage plug-in to iManager. The NSS utility called `ncsinit(8)` is available for initializing a device and setting it to a shared state.

- Section A.4.1, “Syntax,” on page 380
- Section A.4.2, “Options,” on page 381
- Section A.4.3, “Return Value,” on page 383
- Section A.4.4, “Examples,” on page 384

A.4.1 Syntax

```
sbdutil [-c|-f|-i|-I|-v|-Z|-E|-r][-d device] [-p partition] [-n
cluster_name] [-s size][-k filename][-m filename]

sbdutil -c -d device [-s [size][-Z] [-n cluster_name]
sbdutil -f [-s] [-n cluster_name]
sbdutil -i -p partition [-s] [-n cluster_name]
sbdutil -I
sbdutil -v [all<slotnumber>] [-p partition] [-s] [-n cluster_name]
sbdutil -r -weg [-R resource_name] [-N node_name] [-A time] [-B time][-x
filename] [-j filename] [-u filename]
sbdutil -E -k filename
sbdutil -m filename
```

Log in to a node in the cluster where you want to create the SBD partition, then enter the command at a terminal console as the `root` user or any other user in `admin` or `ncsgroup`. If the command succeeds, the partition name is printed. See Section A.4.3, “Return Value,” on page 383 for more information.

A.4.2 Options

`-c`

Create an SBD partition. This option requires at least one device to be specified. You can create a mirrored SBD by supplying more than one device with multiple instances of the -d option.

IMPORTANT: Do not create an SBD partition for a cluster that already has an SBD partition. If you need to re-create the SBD for a cluster, delete its existing SBD first.

To delete an SBD:

1. Enter `cluster down` at the server console of one cluster server.

   This causes all cluster servers to leave the cluster.

2. Delete the SBD partition.

   You can use `nssmu` or other utilities to delete the SBD partition.

`-f`

Find the SBD partition.

`-i`

Initializes the SBD partition except the cluster event logs. If there is information in the `sbdutil -v` view from old nodes that have been removed from the cluster, you can run `sbdutil -i` to initialize the view. This can be done without any issues while the cluster is running.

`-I`

Initializes the entire SBD partition.

`-v`

View the used slots on the SBD partition.

Optional Paremeters:

1. `all`: To view all the slots on the SBD partition.
2. `slotnumber`: To view the specified number of slots of the SBD partition.

`-d <device>`

The device where you want to create an SBD partition. You can create a mirrored SBD by supplying more than one device with multiple instances of the -d option. Specify only the base (leaf) names (such as `sdb` or `mpathd`) with the -d option.

`-p <partition>`

Use this partition instead of searching for one.

-n <*cluster_name*>

Use the specified cluster name instead of getting the name from `cluster.xml`. If this option is not specified, the SBD partition is named by default with the cluster name that is found in the `cluster.xml` file on the node.

-s

Assume the device is a shared disk system instead of checking `cluster.xml`. An optional partition size (in MB) can also be specified when creating a partition (-c). The default size is 8 MB. Some of the allocated space is used for storing metadata such as the partition table and MBR (master boot record) for the partition.

Specify the size as -1 to use all free space on the device. This option allows OES Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

-Z

Use the entire device when creating a partition by using (\-c) option.

-E

To backup the cluster event logs in the specified format.

-k *file_name*

or.

--raw

To export the cluster event logs that are backed up to a file in a raw format. To be used along with the -E option.

-m *file_name*

or

--restore

To restore the backed up cluster event logs to a newly created sbd partition. It is highly recommended to perform the following sequence when restoring logs:

1. Cluster is down.
2. Restore cluster event logs. Any existing cluster event logs in the partition are over-ridden.

-r

Prints all log entries from the oldest to the newest.

Options Used with `-r`

-R *resource_name*

Prints the log entires of the specified resource.

-N *node_name*

Prints the log entries of the specified node.

-w -e -g

Prints the warning, error, and normal log entries respectively.

-A *time*

Prints the log entries after the specified time.

-B *time*

Prints the log entries before the specified time.

-x

or

--xml

Exports the log in XML format.

-j

or

--json

Exports the log in json format.

-u

or

--csv

Exports the log in CSV format.

A.4.3 Return Value

If the command succeeds, it prints the partition name to the screen and returns "0". Otherwise, no partition name is printed out, and a non-zero error code is returned.

You can use "`echo $?`" at the end of the command to check the return code:

```
/opt/novell/ncs/bin/sbdutil -c -n <cluster_name> -d <device> -d <device> -
s <size>; echo $?
```

Examples

The following example demonstrates a successful command where the `mycluster1.sbd` partition is created:

```
/opt/novell/ncs/bin/sbdutil -c -n mycluster1 -d MD_EMC_001G_A -d
MD_EMC_001G_B -s -1; echo $?
/dev/nss/mycluster1.sbd
0
```

The following example demonstrates a failed command where no partition is created, and a non-zero error code is returned:

```
/opt/novell/ncs/bin/sbdutil -c -n mycluster1 -d MD_EMC_001G_A -d
MD_EMC_001G_B -s 2048; echo $?

2
```

A.4.4 Examples

This program normally runs as `root` on a node in the cluster that will be using the SBD.

sbdutil -f

This tells you whether an SBD partition exists and identifies the device on the SAN where the SBD partition is located.

sbdutil -c -n mycluster1 -d sdb -s -1

Creates an SBD partition for the cluster named `mycluster1` on the LUN `/dev/sdb`. The size is set to use all of the free space on the device.

sbdutil -c -n mycluster1 -d sdb -d sdc

Creates a mirrored SBD partition for the cluster named `mycluster1` on the LUN `/dev/sdb` with the mirror on `/dev/sdc`.

sbdutil -r

Lets you view recent cluster events. Events can be node specific (a node joined or left the cluster) or they can be resource specific (a resource changed state from loading to running).

Sample output:

```
Mon Sep 15 15:41:36 2014 Node 'CG-04' Joined at epoch 5Mon Sep 15
15:41:37 2014 Resource 'Master_IP_Address_Resource' Running on CG-01
```

sbdutil -r -w -R resource1 -R resource2 -A '10 days ago' -B '12/31/2019 14:23'

Prints the warning log entries of resource1 and resource2 from 10 days ago till the time 14:23 on 31st December, 2019.

sbdutil -r -R POOL1_SERVER -B '2020-01-20 15:37:44'

Sample output:

```
Mon Jan 20 15:33:12 2020 Resource 'POOL1_SERVER' NDS Sync
Mon Jan 20 15:36:18 2020 Resource 'POOL1_SERVER' Offline
Mon Jan 20 15:36:18 2020 Resource 'POOL1_SERVER' Offline
```

sbdutil -r -R POOL1_SERVER -A '2020-01-20 15:33:12' -B '2020-01-20 15:37:44'

Sample output:

```
Mon Jan 20 15:36:18 2020 Resource 'POOL1_SERVER' Offline
Mon Jan 20 15:36:18 2020 Resource 'POOL1_SERVER' Offline
```

sbdutil -r -R CP5_SERVER --xml logs.xml --json logs.json

This command exports the event logs of the resource `'CP5_SERVER'` in xml and json file formats.

sbdutil -v7 or sbdutil -v 7

This command allows you to view the first 7 slots of the SBD partition.

sbdutil -E --raw backup_logs / sbdutil -E -k backup_logs

This command takes a backup of the event logs into `backup_logs` which can be later used for restoration purposes.

sbdutil -m backup_logs

This command restores the event logs to SBD Partition.

A.5 DotOutParser Utility

The DotOutParser utility (`/opt/novell/ncs/bin/dotoutparser.pl`) prints summaries of failed or incomplete events that have been recorded in a specified `/var/opt/novell/log/ncs/<resource_name>.<script>.out` log file. It also prints output (if any) from any commands that failed or got stuck. You can include ignored errors if desired. Line numbers allow you to easily refer the summary output to the source lines in the `*.out` file. You can omit the line numbers if desired.

The verbosity level controls the amount of information printed to the summary. The higher the level of verbosity you specify, the more information you get in the summary output.

- Section A.5.1, “Syntax,” on page 385
- Section A.5.2, “Options,” on page 385
- Section A.5.3, “Example,” on page 386

A.5.1 Syntax

At the command prompt on the cluster node, enter the following command as the `root` user:

```
/opt/novell/ncs/bin/dotoutparser.pl [-a] [-n] [-i] [-v|-vv|-vvv|-vvvv] <file_name|->
```

You can copy the `/var/opt/novell/log/ncs/ <resource_name>.<script>.out` log file of interest to a working location, then execute the command on the copy.

A.5.2 Options

-?, -h, --help

Print usage and exit.

***file_name*|-**

(Mandatory) Specify the `*.out` log file to be parsed, or specify `-` (hyphen) for standard input.

`/var/opt/novell/log/ncs/<resource_name>.<script>.out`

-a

(Optional) Print all records in the specified file.

-n

(Optional) Do not add line numbers to the summary output.

-i

(Optional) Print ignored errors in addition to the failed or incomplete events.

-v, -vv, -vvv, -vvvv

(Optional) Specify the verbose level (1 to 4). The higher the level of verbosity you specify, the more information you get in the summary output.

-v

(Default) Print the entries with plus (+) variable assignments.

-vv

Print the `-v` level entries and entries starting with ++ (sub-command).

-vvv

Print the `-vv` level entries and entries starting with +++ (more sub-command).

-vvvv

Print all entries in the specified file.

A.5.3 Example

```
# /opt/novell/ncs/bin/dotoutparser.pl /usr/src/temp/POOL_09_SERVER.load.out

Fri May  3 11:51:59 2013 (#1):
2: + exit_on_error nss /poolact=POOL_09
9: + return 0
10: + exit_on_error ncpcon mount VOL_092=237
25: + return 0
26: + exit_on_error ncpcon mount VOL_091=238
41: + return 0
42: + exit_on_error add_secondary_ipaddress 192.168.1.9
94: + return 0
95: + exit_on_error ncpcon bind --ncpservername=CLUSTER1_POOL_09_SERVER --
ipaddress=192.168.1.9
97: ++ ncpcon bind --ncpservername=CLUSTER1_POOL_09_SERVER --ipaddress=192.168.1.9
98: ... Executing " bind"
99:
100: Cannot bind server, failure reason = -632
101:
102: ... FAILED completion [elapsed time = 20 Seconds 2 msecs 671 usecs]
107: + exit 1
```

A.6 CSMPORT Utility (Cluster Segment Manager Import/Export)

The OES Cluster Services Cluster Segment Manager (CSM) Import/Export utility allows you to use a Linux POSIX volume that was cluster-enabled with a CSM container. This conversion is required because the Enterprise Volume Segment Manager (EVMS) was deprecated on SUSE Linux Enterprise Server 11, and the CSM is no longer available on OES 11 and later servers.

The utility can be used to scan for CSM containers and check if they can be imported for use in the cluster. This is the default mode when you run it without any options, or with only the `-v` option. By default, the utility searches block devices (`/dev/sd*`) for containers. Information about each device is displayed. Details about the CSM container are present if the container exists and can be imported for use in the cluster.

The utility also allows you to do the following:

- Search for a specified CSM container and check if it can be imported for use in the cluster.
- Import a specified CSM container for exclusive use in the cluster, and generate a corresponding Device Mapper object for it as `/dev/mapper/csm_name` .
- Export a specified CSM container by removing its Device Mapper object from the `/dev/mapper` directory. This allows the container to be imported by other nodes.

You can specify devices to search for any of the options. Wildcards such as the question mark (?) and asterisk (*) can be used when specifying devices. A question mark replaces a character in that position. An asterisk replaces any number of characters in that position.

- Section A.6.1, “Syntax,” on page 387
- Section A.6.2, “Options,” on page 387
- Section A.6.3, “Examples,” on page 389

A.6.1 Syntax

Open a console on a cluster node, then issue the command as the `root` user.

```
/opt/novell/ncs/bin/csmport [-v|-s] [-h] [-e devmapper_object | -i
container_name | -c container_name] [device...]
```

A.6.2 Options

No options

Use the command with no options to search the block devices (`/dev/sd*`) for containers.

```
/opt/novell/ncs/bin/csmport
```

[devices]

Use the command with no options and specify the devices to search in order to specify which devices to search for CSM containers. By default, the utility searches block devices (`/dev/sd*`) for containers.

```
/opt/novell/ncs/bin/csmport [device...]
```

When specifying devices, wildcards (? and *) can be used.

Examples

The following command searches the multipath device `/dev/mapper/MD_EMC_005G` for CSM containers.

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_005G
```

The following command uses three question mark wildcards to search multipath devices with any three characters in those positions in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_???G
```

The following command uses an asterisk wildcard to search multipath devices with any number of characters in that position in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_*
```

-h

Prints the help page for this command.

-c <*container_name*> [device...]

Searches the `/dev/sd*` devices or specified devices for the named CSM container.

```
/opt/novell/ncs/bin/csmport -c container_name
```

-i <*container_name*> [device...]

Use this option to import a CSM container on a node for exclusive use in the cluster. The utility achieves this by searching for the container and if found, extracting the volume inside the container and exposing it exclusively at `/dev/mapper/container_name`.

The utility uses the container name instead of the Linux POSIX volume name because some CSM containers might contain compatibility volumes, which do not have names. The `/dev/mapper/container_name` object can be mounted at the old mount point that was used for the shared Linux POSIX volume in that container. The mount point can have the same name or different name than the Device Mapper object.

```
/opt/novell/ncs/bin/csmport -i container_name
```

Example

Suppose that you created a Linux POSIX cluster resource with the following settings on an OES 2 server or an OES 1 server:

```
RESOURCE_IP:            10.10.10.44
CONTAINER_NAME:         csm44
EVMS_VOLUME_NAME:       lxvol44
MOUNT_DEV:              /dev/evms/$CONTAINER_NAME/$EVMS_VOLUME_NAME
MOUNT_POINT:            /mnt/users
```

On an OES 11 or later server, you can use the following command to generate the Device Mapper object at `/dev/mapper/csm44`:

```
/opt/novell/ncs/bin/csmport -i csm44
```

The Device Mapper object can be mounted at the old mount point.

The resulting settings are:

```
RESOURCE_IP:            10.10.10.44
CONTAINER_NAME:         csm44
MOUNT_DEV:              /dev/mapper/$CONTAINER_NAME
MOUNT_POINT:            /mnt/users
```

-e <*devmapper_object*> [device...]

Use this option to export the Device Mapper object from the node. It removes the object from the `/dev/mapper` directory. Basically, the option undoes what the `-i` option did, so that the container can be imported by other nodes.

```
/opt/novell/ncs/bin/csmport -e devmapper_object
```

Example

The following command exports the Device Mapper object `/dev/mapper/csm44`.

```
/opt/novell/ncs/bin/csmport -e csm44
```

-v

Verbose mode can be used with any of the options.

-s

Silent mode can be used with any of the options.

A.6.3 Examples

This section provides examples of typical tasks you perform when using the `csmport` command.

- "Scanning for CSM Containers" on page 389
- "Importing a CSM Container for Exclusive Use on a Node in a Cluster" on page 393
- "Exporting a CSM Container" on page 393
- "Searching for a Container on a Specified Device" on page 394
- "Searching for Containers on Multiple Devices with Wildcards" on page 394
- "Importing a CSM Container on a Specified Device" on page 394

Scanning for CSM Containers

Scan for CSM containers and check if they can be imported for use in the cluster.

```
/opt/novell/ncs/bin/csmport -v
```

In the following sample output, only device `/dev/sdi` has a CSM container `CSM5GNoEVMSVol` and can be imported to `/dev/mapper/CSM5GNoEVMSVol` for its exclusive use on the cluster node.

```
Scan for CSM containers

  Search device '/dev/sdn' for CSM containers.
    Retrieve device information.
    Device '/dev/sdn': major=8, minor=208.
    Refresh device '/dev/sdn'.
    Open device '/dev/sdn' for reading.
    Check CSM metadata on device '/dev/sdn'.
      /dev/sdn is not a CSM container.

  Search device '/dev/sdm' for CSM containers.
    Retrieve device information.
    Device '/dev/sdm': major=8, minor=192.
    Refresh device '/dev/sdm'.
    Open device '/dev/sdm' for reading.
    Check CSM metadata on device '/dev/sdm'.
      /dev/sdm is not a CSM container.

  Search device '/dev/sdh' for CSM containers.
    Retrieve device information.
    Device '/dev/sdh': major=8, minor=112.
    Refresh device '/dev/sdh'.
    Open device '/dev/sdh' for reading.
    Check CSM metadata on device '/dev/sdh'.
      /dev/sdh has a CSM container.
      Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
      Container name ='AJTestContainer01' on disk.
      Container is accessible (1).

  Search device '/dev/sdc' for CSM containers.
    Retrieve device information.
    Device '/dev/sdc': major=8, minor=32.
    Refresh device '/dev/sdc'.
```

```
    Open device '/dev/sdc' for reading.
    Check CSM metadata on device '/dev/sdc'.
      /dev/sdc is not a CSM container.

  Search device '/dev/sdl' for CSM containers.
    Retrieve device information.
    Device '/dev/sdl': major=8, minor=176.
    Refresh device '/dev/sdl'.
    Open device '/dev/sdl' for reading.
    Check CSM metadata on device '/dev/sdl'.
      /dev/sdl is not a CSM container.

  Search device '/dev/sdi' for CSM containers.
    Retrieve device information.
    Device '/dev/sdi': major=8, minor=128.
    Refresh device '/dev/sdi'.
    Open device '/dev/sdi' for reading.
    Check CSM metadata on device '/dev/sdi'.
      /dev/sdi has a CSM container.
      Container name ='CSM5GNoEVMSVol' on disk.
      Container is accessible (4).
      No EVMS volume.
      New name: /dev/mapper/CSM5GNoEVMSVol
      Size: 10485757 sectors.
    Device '/dev/sdi' has a CSM container 'CSM5GNoEVMSVol' and can be
converted to /dev/mapper/CSM5GNoEVMSVol.

  Search device '/dev/sdk' for CSM containers.
    Retrieve device information.
    Device '/dev/sdk': major=8, minor=160.
    Refresh device '/dev/sdk'.
    Open device '/dev/sdk' for reading.
    Check CSM metadata on device '/dev/sdk'.
      /dev/sdk is not a CSM container.

  Search device '/dev/sdg' for CSM containers.
    Retrieve device information.
    Device '/dev/sdg': major=8, minor=96.
    Refresh device '/dev/sdg'.
    Open device '/dev/sdg' for reading.
    Check CSM metadata on device '/dev/sdg'.
      /dev/sdg is not a CSM container.

  Search device '/dev/sdb' for CSM containers.
    Retrieve device information.
    Device '/dev/sdb': major=8, minor=16.
    Refresh device '/dev/sdb'.
    Open device '/dev/sdb' for reading.
    Check CSM metadata on device '/dev/sdb'.
      /dev/sdb is not a CSM container.

  Search device '/dev/sdt' for CSM containers.
    Retrieve device information.
    Device '/dev/sdt': major=65, minor=48.
    Refresh device '/dev/sdt'.
```

```
    Open device '/dev/sdt' for reading.
    Failed to read the 1st sector of device '/dev/sdt': Input/output error
(5).

  Search device '/dev/sdx' for CSM containers.
    Retrieve device information.
    Device '/dev/sdx': major=65, minor=112.
    Refresh device '/dev/sdx'.
    Open device '/dev/sdx' for reading.
    Failed to read the 1st sector of device '/dev/sdx': Input/output error
(5).

  Search device '/dev/sdj' for CSM containers.
    Retrieve device information.
    Device '/dev/sdj': major=8, minor=144.
    Refresh device '/dev/sdj'.
    Open device '/dev/sdj' for reading.
    Check CSM metadata on device '/dev/sdj'.
      /dev/sdj has a CSM container.
      Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
      Container name ='CMS4GCon' on disk.
      Container is accessible (4).

  Search device '/dev/sdy' for CSM containers.
    Retrieve device information.
    Device '/dev/sdy': major=65, minor=128.
    Refresh device '/dev/sdy'.
    Open device '/dev/sdy' for reading.
    Failed to read the 1st sector of device '/dev/sdy': Input/output error
(5).

  Search device '/dev/sdu' for CSM containers.
    Retrieve device information.
    Device '/dev/sdu': major=65, minor=64.
    Refresh device '/dev/sdu'.
    Open device '/dev/sdu' for reading.
    Failed to read the 1st sector of device '/dev/sdu': Input/output error
(5).

  Search device '/dev/sdd' for CSM containers.
    Retrieve device information.
    Device '/dev/sdd': major=8, minor=48.
    Refresh device '/dev/sdd'.
    Open device '/dev/sdd' for reading.
    Failed to read the 1st sector of device '/dev/sdd': Input/output error
(5).

  Search device '/dev/sdq' for CSM containers.
    Retrieve device information.
    Device '/dev/sdq': major=65, minor=0.
    Refresh device '/dev/sdq'.
    Open device '/dev/sdq' for reading.
    Check CSM metadata on device '/dev/sdq'.
      /dev/sdq is not a CSM container.
```

```
  Search device '/dev/sds' for CSM containers.
    Retrieve device information.
    Device '/dev/sds': major=65, minor=32.
    Refresh device '/dev/sds'.
    Open device '/dev/sds' for reading.
    Failed to read the 1st sector of device '/dev/sds': Input/output error
(5).

  Search device '/dev/sdv' for CSM containers.
    Retrieve device information.
    Device '/dev/sdv': major=65, minor=80.
    Refresh device '/dev/sdv'.
    Open device '/dev/sdv' for reading.
    Failed to read the 1st sector of device '/dev/sdv': Input/output error
(5).

  Search device '/dev/sdaa' for CSM containers.
    Retrieve device information.
    Device '/dev/sdaa': major=65, minor=160.
    Refresh device '/dev/sdaa'.
    Open device '/dev/sdaa' for reading.
   Failed to read the 1st sector of device '/dev/sdaa': Input/output error
(5).

  Search device '/dev/sde' for CSM containers.
    Retrieve device information.
    Device '/dev/sde': major=8, minor=64.
    Refresh device '/dev/sde'.
    Open device '/dev/sde' for reading.
    Check CSM metadata on device '/dev/sde'.
      /dev/sde is not a CSM container.

  Search device '/dev/sdz' for CSM containers.
    Retrieve device information.
    Device '/dev/sdz': major=65, minor=144.
    Refresh device '/dev/sdz'.
    Open device '/dev/sdz' for reading.
    Failed to read the 1st sector of device '/dev/sdz': Input/output error
(5).

  Search device '/dev/sdw' for CSM containers.
    Retrieve device information.
    Device '/dev/sdw': major=65, minor=96.
    Refresh device '/dev/sdw'.
    Open device '/dev/sdw' for reading.
    Failed to read the 1st sector of device '/dev/sdw': Input/output error
(5).

  Search device '/dev/sdr' for CSM containers.
    Retrieve device information.
    Device '/dev/sdr': major=65, minor=16.
    Refresh device '/dev/sdr'.
    Open device '/dev/sdr' for reading.
    Failed to read the 1st sector of device '/dev/sdr': Input/output error
(5).
```

```
  Search device '/dev/sdp' for CSM containers.
    Retrieve device information.
    Device '/dev/sdp': major=8, minor=240.
    Refresh device '/dev/sdp'.
    Open device '/dev/sdp' for reading.
    Failed to read the 1st sector of device '/dev/sdp': Input/output error
(5).

  Search device '/dev/sdo' for CSM containers.
    Retrieve device information.
    Device '/dev/sdo': major=8, minor=224.
    Refresh device '/dev/sdo'.
    Open device '/dev/sdo' for reading.
    Failed to read the 1st sector of device '/dev/sdo': Input/output error
(5).

  Search device '/dev/sdf' for CSM containers.
    Retrieve device information.
    Device '/dev/sdf': major=8, minor=80.
    Refresh device '/dev/sdf'.
    Open device '/dev/sdf' for reading.
    Check CSM metadata on device '/dev/sdf'.
      /dev/sdf is not a CSM container.

  Search device '/dev/sda' for CSM containers.
    Retrieve device information.
    Device '/dev/sda': major=8, minor=0.
    Refresh device '/dev/sda'.
    Open device '/dev/sda' for reading.
    Check CSM metadata on device '/dev/sda'.
      /dev/sda is not a CSM container.
```

Importing a CSM Container for Exclusive Use on a Node in a Cluster

Use the `-i` option to import the CSM container named `CSM5GNoEVMSVol` for exclusive use on a node in a cluster.

```
/opt/novell/ncs/bin/csmport –i CSM5GNoEVMSVol
```

This command generates the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`. You can map this object to the desired mount point.

Exporting a CSM Container

Use the -e option to export the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`.

```
/opt/novell/ncs/bin/csmport –e CSM5GNoEVMSVol
```

This removes the object and allows the container to be imported by other nodes.

Searching for a Container on a Specified Device

Specify a device to be searched in order to speed up the search for a container or containers. For example, to search for CSM containers on a multpath device `/dev/mapper/MD_EMC_005G`, enter

```
/opt/novell/ncs/bin/csmport /dev/mapper/MD_EMC_005G
```

Searching for Containers on Multiple Devices with Wildcards

Specify multiple devices to be searched and use wildcards to include multiple devices of that type.

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_???G /dev/disk/by-id/dm-uuid-mpath-*
```

Importing a CSM Container on a Specified Device

When importing the `CSM5GNoEVMSVol` container, specify the multipath device `/dev/mapper/MD_EMC_005G` in order to speed up the search for the container.

```
/opt/novell/ncs/bin/csmport -i CSM5GNoEVMSVol /dev/mapper/MD_EMC_005G
```

A.7 ncs-configd.py Script (Pushing Configuration Updates to All Nodes)

You can use the `/opt/novell/ncs/bin/ncs-configd.py` script to push configuration updates to all nodes in a cluster.

Syntax

As the `root` user, enter

```
/opt/novell/ncs/bin/ncs-configd.py -init
```

OES Cluster Services can be running or not running when you issue the command.

A.8 ncs_install.py Script (Extending the Schema, Monitoring NDSD, or Removing NCS Configuration)

You can use the `/opt/novell/ncs/install/ncs_install.py` script to extend the eDirectory schema in a tree for cluster objects or to monitor the eDirectory daemon on a node.

- Section A.8.1, “Extending the Schema,” on page 395
- Section A.8.2, “Monitoring NDSD,” on page 396
- Section A.8.3, “Removing NCS Configuration,” on page 397

A.8.1 Extending the Schema

You can extend the eDirectory schema for cluster objects before OES Cluster Services is installed for the first time in the tree. This allows an eDirectory tree administrator to expand the schema, and a container administrator to install the software. The schema needs to be expanded only one time in the tree. See Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,” on page 58.

After you run the script, the schema includes the Cluster object container and the following types of objects in it:

- Cluster Node objects
- Cluster Resource objects
- Cluster Template objects
- Volume Resource objects

Syntax

Run the following command as the `root` user:

```
/opt/novell/ncs/install/ncs_install.py -e -f <configuration_filename>
```

Options

-e

Expand a specified eDirectory schema to add cluster objects. Use with the `-f` option to specify a file that contains information about the LDAP server where the schema is stored.

-f <*configuration_file_name*>

Use with the `-e` option to specify information about the LDAP server where the schema to be expanded is stored.

In a text editor, create a text file, specify the configuration information for the OES Cluster Services cluster in it, then save the file. This file contains a password in clear text. For security reasons, ensure that you delete the file when you are done.

The following lines are an example of the content of a file, with sample values. The directives are self-explanatory.

IMPORTANT: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

```
CONFIG_NCS_LDAP_IP="10.1.1.102"
CONFIG_NCS_LDAP_PORT="636"
CONFIG_NCS_ADMIN_DN="cn=admin.o=context"
CONFIG_NCS_ADMIN_PASSWORD="password"
```

Example

Create a configuration file on the server desktop and name it `mytree_ncs_schema.txt`. Open a terminal console on the server where you created the configuration file, then enter the following commands as the `root` user:

```
mkdir -p /var/opt/novell/install

/opt/novell/ncs/install/ncs_install.py -e -f /root/Desktop/mytree_ncs_schema.txt
```

A.8.2 Monitoring NDSD

You can monitor the status of the eDirectory daemon (ndsd) at the NCS level. It is disabled by default. The monitoring can be set independently on each node. On a node, if the eDirectory daemon does not respond to a status request within a specified timeout period, NCS can take one of three configurable remedy actions: an ndsd restart, a graceful node restart, or a hard node restart. See Section 8.8, “Configuring NCS to Monitor the eDirectory Daemon (ndsd),” on page 111.

Syntax

```
/opt/novell/ncs/install/ncs_install.py
  -m <display|disable|restart-ndsd|reboot-node|hard-reboot|help>
  [-i <interval_value>]
  [-t <timeout_value>]]
```

The settings apply to an individual node. Different nodes can have different settings.

The settings take effect immediately. You do not need to start or restart any services.

Options

-m <display|disable|restart-ndsd|reboot-node|hard-reboot|help>

Monitor the eDirectory daemon (ndsd) status on this node. Use this option with the -i and -t options.

display

Display the current settings for monitoring of the eDirectory ndsd daemon on this node.

disable

Turn off the monitoring of the eDirectory ndsd daemon on this node.

restart-ndsd

Enable monitoring of the eDirectory ndsd daemon on this node and set a remedy action of a restart for the eDirectory ndsd daemon.

reboot-node

Enable monitoring of the eDirectory ndsd daemon on this node and set a remedy action of a graceful restart for the node.

hard-reboot

Enable monitoring of the eDirectory ndsd daemon on this node and set a remedy action of a hard restart for the node.

help

Display the usage information for the -m option.

-i <*interval_value*>

Specify the interval of time in seconds to check the status of ndsd on this node. The minimum interval is 16 seconds. Use this option with the -m and -t options.

-t <*timeout_value*>

Specify a timeout value in seconds to wait for a response from ndsd before enforcing the configured remedy action. The minimum timeout is 8 seconds. Use this option with the `-m` and `-i` options.

Examples

Example 1: Display current settings for ndsd monitoring.

```
# /opt/novell/ncs/install/ncs_install.py -m display
NDSD monitoring is enabled.
interval = 90 seconds, timeout = 20 seconds and remedy action is reboot node hard.
```

Example 2: Enable ndsd monitoring with an interval of 60 seconds, timeout of 15 seconds, and remedy action of rebooting gracefully.

```
# /opt/novell/ncs/install/ncs_install.py -m reboot-node -i 60 -t 15
NDSD monitoring is enabled.
interval = 60 seconds, timeout = 15 seconds and remedy action is reboot node.
```

Example 3: Change the remedy action to restart ndsd.

```
# /opt/novell/ncs/install/ncs_install.py -m restart-ndsd -i 60 -t 15
NDSD monitoring is enabled.
interval = 60 seconds, timeout = 15 seconds and remedy action is restart ndsd.
```

Example 4: Disable ndsd monitoring.

```
# /opt/novell/ncs/install/ncs_install.py -m disable
NDSD monitoring is not enabled (3).
```

Example 5: View usage for ndsd monitoring.

```
# /opt/novell/ncs/install/ncs_install.py -m help
Usage: /opt/novell/ncs/install/ncs_install.py -m display|disable|restart-ndsd|reboot-
node|hard-reboot|help [-i interval_value] [-t timeout_value]
  -m display: Display the current NDSD monitoring settings
  -m disable: Disable NDSD monitoring
  -m restart-ndsd: Enable NDSD monitoring and set remedy action to restarting NDSD
  -m reboot-node: Enable NDSD monitoring and set remedy action to rebooting node
  -m hard-reboot: Enable NDSD monitoring and set remedy action to rebooting node hard
  -m help: Print this help page
  -i interval_value: Specify in seconds how frequently the monitor process checks NDSD
  -t timeout_value: Specify in seconds how long the monitor process waits for the checking
operation
NDSD monitoring is not enabled (1).
```

A.8.3 Removing NCS Configuration

You can use the `/opt/novell/ncs/install/ncs_install.py -r` script to remove cluster configuration information from a node to return it to a pre-configured state. This is useful for the following scenarios:

- If the OES Cluster Services configuration is unsuccessful, you might have a partial configuration that prevents you from properly configuring the node for a cluster.
- You want to move a node from one cluster to another cluster.

After you remove the cluster configuration information, you can start over with the configuration of OES Cluster Services on the node as described in Section 5.5, “Configuring OES Cluster Services,” on page 64.

Syntax

You must stop NCS (`rcnovell-ncs stop` or `systemctl stop novell-ncs.service`) before you run the command. Run the command as the `root` user.

`/opt/novell/ncs/install/ncs_install.py -r [-f <configuration_file_name>]`

Without the `-f` option or applicable variables in the configuration file, `/opt/novell/ncs/install/ncs_install.py` prompts you to enter needed information.

You can check the return code ($?) of the command to see if the command is successful. Detailed log information is appended to `/var/opt/novell/install/ncslog`.

Options

-r

Clean up an unsuccessful OES Cluster Services configuration on a node, or remove the configuration information so that you can move a node from one cluster to another.

-f <*configuration_file_name*>

Use with the `-r` option to specify credentials for the node with the unsuccessful OES Cluster Services configuration.

In a text editor, create a text file, specify the credentials for the OES Cluster Services cluster in it, and then save the file. This file contains a password in clear text. For security reasons, ensure that you delete the file when you are done.

The following lines are an example of the content of a file, with sample values. The directives are self-explanatory.

IMPORTANT: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

```
CONFIG_NCS_CLUSTER_DN="cn=clus1,ou=ncs,o=novell"
CONFIG_NCS_ADMIN_DN="cn=admin.o=novell"
CONFIG_NCS_ADMIN_PASSWORD="password"
```

Example

A cluster configuration process failed. You must remove the unsuccessful OES Cluster Services configuration information from the node before you can re-try the configuration.

1. Open a terminal console as the `root` user.
2. Stop OES Cluster Services if it is running.

   ```
   rcnovell-ncs stop
   ```

   or

   ```
   systemctl stop novell-ncs.service
   ```

3. Do either of the following:
   - In a text editor, create a configuration file on the server desktop, add the cluster credentials, and save the file as `my_cluster_credentials.conf`.

At the command prompt, enter

```
/opt/novell/ncs/install/ncs_install.py -r -f /root/Desktop/
my_cluster_credentials.conf
```

- At the command prompt, enter the command, then reply with the credentials when you are prompted:

```
# /opt/novell/ncs/install/ncs_install.py -r
Please enter the DN for the admin (cn=admin,o=novell):
please enter the password for cn=admin,o=novell:
```

4 Check the return code to see if the command is successful:

```
$?
```

5 After the successful removal of the cluster configuration information, you can configure OES Cluster Services on the node as described in Section 5.5, “Configuring OES Cluster Services,” on page 64.

A.9 ncs_ncpserv.py Script (Creating an NCP Virtual Server Object for a Clustered LVM Volume Group)

The `/opt/novell/ncs/bin/ncs_ncpserv.py` script creates a virtual NCP Server object (`NCS:NCP Server`) in eDirectory, and associates it with none, one, or multiple NCP volumes that you specify. It automatically renames the NCP Volume objects to use the cluster name instead of the server name where the NCP volume was created. NCP clients access files on the Linux POSIX volume via the virtual server name.

- Section A.9.1, “Syntax,” on page 399
- Section A.9.2, “Options,” on page 399
- Section A.9.3, “Examples,” on page 400

A.9.1 Syntax

At the command prompt on the master node of the cluster, enter the following command as the `root` user:

```
/opt/novell/ncs/bin/ncs_ncpserv.py -c <lx_volume_name> -i <resource_ip_address>
     [-v <ncp_volume_name| "ncp_volume_name1:ncp_volume_name2:..."> ]
```

A.9.2 Options

-c <*lx_volume_name*>

Create a virtual NCP Server object for the specified volume.

-i <*resource_ip_address*>

Specify a unique static IP address to use for the cluster resource.

-v <*ncp_volume_name* | "*ncp_volume_name1*:*ncp_volume_name2*:...">

Use this option to specify one or multiple volumes on the Linux cluster resource that need a virtual NCP Server object to be created and bound to the specified IP address. If the -v option is not specified, all of the NCP volumes that currently exist on the Linux POSIX cluster resource are bound to the IP address.

If you enter multiple volume names, use colons to delimit the names and put quotation marks around the list of names. Volume names can be listed by the volume name (`MY_NNCP_VOL06`) or by the volume distinguished name (`cn=CLUS_02_MY_NNCP_VOL06,o=novell`), or any combination of the two methods.

A.9.3 Examples

In the following examples, the resource IP address is 10.10.10.44, the cluster name is `cluster1` and the cluster context is `ou=clusters,o=mycompany`.

Example 1

To specify a single NCP volume named USERS on the `lxvol44` cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v USERS
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.
```

Example 2

To specify multiple NCP volumes on the `lxvol44` cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v
"USERS:MY_NCP_VOL06:cn=servername_MY_NCP_VOL07,ou=clusters,o=novell"
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.

Object 'cn=servername_MY_NCP_VOL06,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL06,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL06
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster_MY_NCP_VOL06,ou=clusters,o=mycompany' are linked with each
other.

Object 'cn=servername_MY_NCP_VOL07,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL07,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL07
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster_MY_NCP_VOL07,ou=clusters,o=mycompany' are linked with each
other.
```

A.10 ncs_resource_scripts.pl Script (Modifying Resource Scripts Outside of iManager)

You can use the `/opt/novell/ncs/bin/ncs_resource_scripts.pl` script to modify resource scripts without using iManager. You can add, remove, or modify the commands in the script, or retrieve its scripts to search for information from the script. You can issue the command from the command line or in a script. This capability is very useful when you need to make the same script change to multiple resources in a cluster.

Given a resource name, the `ncs_resource_sripts.pl` tool does one of two things:

- **Read:** Retrieves the specified resource's scripts and prints them to standard output.
- **Write:** Gathers the script information from standard input and updates the specified resource's scripts.

- Section A.10.1, "Syntax," on page 401
- Section A.10.2, "Options," on page 402
- Section A.10.3, "Examples," on page 402

A.10.1 Syntax

At the command prompt on the master node of the cluster, enter the following command as the `root` user:

```
/opt/novell/ncs/bin/ncs_resource_scripts.pl <--help|-h|-?>
/opt/novell/ncs/bin/ncs_resource_scripts.pl [--read|--write] resource_name
```

You can redirect standard I/O to a file, and change the textual file any way you want.

A return code 0 indicates that the operation is successful. Otherwise, errors will be printed in standard error.

A.10.2 Options

--help, -h, -?

Prints usage help and exits.

--read

Retrieves the specified resource's scripts and prints them to standard output.

You can redirect standard output to a file, and modify the scripts in the file.

--write

Gathers the script information from standard input and updates the specified resource's scripts.

You can redirect standard input to a file that contains the modified scripts.

A.10.3 Examples

In the following examples, enter the `/opt/novell/ncs/bin/ncs_resource_scripts.pl` command at the command prompt on the master node of the cluster as the `root` user.

- "Example 1: Read the Scripts in Standard I/O" on page 402
- "Example 2: Redirect the Scripts from Standard I/O to a File" on page 402
- "Example 3: Delete the CIFS Workaround Line from Scripts" on page 403
- "Example 4: Get the Volume Name from the Resource Load Script" on page 403

Example 1: Read the Scripts in Standard I/O

At the command prompt, enter the following commands to read and write the scripts to and from the same file:

```
/opt/novell/ncs/bin/ncs_resource_scripts.pl --read P_119_93_SERVER
```

Example 2: Redirect the Scripts from Standard I/O to a File

At the command prompt, enter the following commands to read and write the scripts to and from the same file:

```
# /opt/novell/ncs/bin/ncs_resource_scripts.pl --read POOL_01_SERVER > tmp.scripts
```

In a text editor, modify the scripts as needed in the output file `tmp.scripts`, then redirect the content of the modified file to be used with the --write option.

```
# /opt/novell/ncs/bin/ncs_resource_scripts.pl --write POOL_01_SERVER < tmp.scripts
```

Example 3: Delete the CIFS Workaround Line from Scripts

For scripts where an extra CIFS line was added as a workaround for a timing issue, remove the line. At the command prompt, enter the following command:

```
# /opt/novell/ncs/bin/ncs_resource_scripts.pl --read POOL_01_SERVER | sed '/novcifs -sln/
d' | /opt/novell/ncs/bin/ncs_resource_scripts.pl --write POOL_01_SERVER
```

Example 4: Get the Volume Name from the Resource Load Script

Use `grep` with the `ncs_resource_scripts.pl` tool to parse the standard output and get the volume name from the `ncpcon mount` command in the resource load script. At the command prompt, enter the following command:

```
# /opt/novell/ncs/bin/ncs_resource_scripts.pl POOL_01_SERVER | grep -m1 -o
"^exit_on_error[[:space:]]\+ncpcon[[:space:]]\+mount[[:space:]]\+\w\+" | cut -d' ' -f4
```

A.11 NCPCON Mount Command for NCP and NSS Volumes

Use the `ncpcon mount` command to mount an NCP volume or NSS volume on your Linux server. This command makes the volume accessible to NCP clients.

- Section A.11.1, “Syntax,” on page 403
- Section A.11.2, “NCP Volume Examples,” on page 404
- Section A.11.3, “NSS Volume Examples,” on page 404
- Section A.11.4, “Other Examples,” on page 405

A.11.1 Syntax

`mount < all | volume_name | volume_name=volume_id,path=/volume_mntpoint >`

Replace **volume_name** with the name of the NSS volume or NCP volume, such as `VOL1`. To mount all local NCP and NSS volumes, replace **volume_name** with `all`. A clustered NCP volume or NSS volume is mounted when its cluster resource is brought online.

Replace *volume_id* with a value from 0 to 254 as the server volume ID. It is not necessary to manually specify a volume ID for a locally mounted volume. NCP automatically assigns unique volume IDs to locally mounted NCP volumes and NSS volumes in increasing order, from 0 to 254. IDs 0 and 1 are reserved for the `sys` and `_admin` volumes. When the command is used in a cluster resource script, the volume ID must be specified to ensure that the volume ID is unique across all cluster nodes where the volume will be mounted. By convention in clusters, the volume IDs are assigned in decreasing order, from 254 to 0.

A.11.2 NCP Volume Examples

For an NCP volume created on a Linux POSIX file system, replace `/volume_mntpoint` with the Linux path of the mount point for the NCP share. Typically, this path is mount point of the Linux volume; that is, the NCP share is created at the root of the Linux volume. The path must exist on all nodes in the cluster.

If an NCP volume is mounted locally, the mount path is stored in the `/etc/fstab` file, so it is not necessary to specify a mount path. For example:

```
ncpcon mount VOL1
```

The mount path for the NCP volume is required when you use the command in a cluster resource load script. For example, if you use NSSMU to create and NCP enable a clustered Linux LVM volume `myvol` with a mount point of `/usr/novell/myvol`, the NCP volume name (share name) is MYVOL, and the resource load script includes the following definitions for variables used in the command:

```
# define the mount point path of the Linux volume
MOUNT_PATH=/usr/novell/myvol

# define the NCP volume name
NCP_VOLUME=MYVOL

exit_on_error ncpcon mount $NCP_VOL=253,path=$MOUNT_PATH
```

A.11.3 NSS Volume Examples

For an NSS volume, the default mount path is `/media/nss/<nss_volume_name>`. For example, if you create an NSS volume named USERS, the default mount path is `/media/nss/USERS`. If you use the default path for NSS volumes, it is not necessary to include the path option for local mounts:

```
ncpcon mount USERS
```

The default path is also not required in a cluster script, but a volume ID that is unique across all cluster nodes must be specified:

```
exit_on_error ncpcon mount USERS=252
```

In a cluster, you cannot rename the mount point path. The default mount point path of `/media/nss/<volume_name>` applies.

```
exit_on_error ncpcon mount <volumename>=<volume_ID>
```

For example:

```
exit_on_error ncpcon mount USERS=252
```

The volume is mounted at `/media/nss/USERS`.

You can specify mount options for the NSS volume by using the `/opt` switch in the command. For example, you can specify the name space to use by using the `/opt=ns=<long|unix|dos|mac>` command.

```
# specify the name space to use when mounting the volume
exit_on_error ncpcon mount USERS=252 /opt=ns=unix
```

A.11.4 Other Examples

```
ncpcon mount sys
ncpcon mount all
```

A.12 Querylog Utility

The querylog lets you view recent cluster events. Events can be node specific (a node joined or left the cluster) or they can be resource specific (a resource changed state from loading to running)

You can specify a relative date range (a period of time that is relative to the current date) to find log records. Specifying the relative date range enables you save your queries and execute them as and when required without needing to alter the date range. Examples of relative date range: 10 days ago, Yesterday, Last Friday, and so forth.

- Section A.12.1, “Syntax,” on page 405
- Section A.12.2, “Options,” on page 405
- Section A.12.3, “Examples,” on page 406

A.12.1 Syntax

```
/opt/novell/ncs/bin/querylog [-ehiow] [-aAbB date][-n name] [-l number]
```

Enter the command at a terminal console prompt as the root user or any other user in `admin` or `ncsgroup`. If the command succeeds, the log entries are displayed.

A.12.2 Options

`-a` *`date`*

Prints all the log entries made on the specified date and continues until the date and time the command was executed.

`-b` *`date`*

Prints all the log entries made on and before the specified date. The output includes the entries since NCS started logging events until the specified date.

`-A` *`date`*

Prints all the log entries made after the specified date. The entries made before and on the specified date are excluded.

`-B` *`date`*

Prints all the log entries made before the specified date. The entries made on and after the specified date are excluded.

`-h`

Prints the help page.

`-l` *`number`*

Limits the output to the number of entries or lines.

`-o`

Prints the oldest log entry first.

`-n` *`name`*

Specify the name of the resource or node for which you want to view logs.

-i

Filters and prints only the informative log entries.

-w

Filters and prints only the warning log entries.

-e

Filters and prints only the error log entries.

A.12.3 Examples

This program normally runs as `root`.

/opt/novell/ncs/bin/querylog

Prints all the log entries since NCS started logging events until the date and time the command was executed.

The output will be in a reverse chronological order, which means the most recent event is displayed first.

/opt/novell/ncs/bin/querylog -oe -B 'yesterday 16:44'

Prints all the error log entries made before 16:44 the previous day from the current date.

The output will be in a chronological order, which means the oldest entry is displayed first.

/opt/novell/ncs/bin/querylog -w -e -a '10 days ago' -b '10/21/2014 14:23' -n name_of_resource_1 -n name_of_resource_2 -n name_of_node1 -o -l66

Filters and prints up to 166 error and warning log entries for the specified node(s) and resource(s).

The output will be in a chronological order, which means the oldest entry is displayed first.

The output includes all the log entries made in the past 10 days, that is, 14:23 October 11, 2014 until 14:23 on October 21, 2014. The log entries made at 14:23 on October 21, 2014 and after are excluded.

B Files for OES Cluster Services

Knowing the location and purpose of the files that make up OES Cluster Services can be useful in helping you troubleshoot problems and resolve version issues. Table B-1 lists the path and purpose for some of the files that are part of OES Cluster Services.

Table B-1 *Location and Purpose of OES Cluster Services Files*

NCS File Name and Path	Purpose
`/usr/lib/systemd/system/OES-ncs.service`	Systemd unit file
`/etc/opt/novell/ncs/<node_name>`	Configuration file for a specific cluster node
`/etc/opt/novell/ncs/clstrlib.conf`	Cluster configuration file
`/lib/modules/kernel_dir/updates/ncs/clstrlib.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/cma.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/cmsg.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/crm.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/css.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/cvb.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/gipc.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/sbd.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
`/lib/modules/kernel_dir/updates/ncs/sbdlib.ko`	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.

NCS File Name and Path	Purpose
/lib/modules/*kernel_dir*/updates/ncs/vipx.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/updates/ncs/vll.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/opt/novell/ncs/bin/adminfs	Cluster management (iManager and CLI)
/opt/novell/ncs/bin/ClusterCli.pl	Cluster CLI engine
/opt/novell/ncs/bin/ClusterCliSnapinInterface.pm	Cluster CLI engine
/opt/novell/ncs/bin/ClusterCliUtils.pm	Cluster CLI engine
/opt/novell/ncs/bin/csmport	CSM Import/Export utility
/opt/novell/ncs/bin/dotoutparser.pl	DotOutParser utility
/opt/novell/ncs/bin/ldncs	Loads NCS; used by the `Cluster Start` command.
/opt/novell/ncs/install/ncs_install.py	Script to extend the eDirectory schema or to configure eDirectory monitoring
/opt/novell/ncs/bin/ncs_ncpserv.py	Virtual NCS NCP Server Object script
/opt/novell/ncs/bin/NCS_STONITH_SCRIPT file	An executable file that you create to authenticate to your power controller and turn off the power, cycle the power, or reset the power for the node. The presence of the file enables the STONITH feature for the node.
/opt/novell/ncs/bin/ncs-configd.py	Cluster configuration daemon
/opt/novell/ncs/bin/ncs-emaild	Cluster email daemon
/opt/novell/ncs/bin/ncs-resourced.py	Daemon used to run load and unload scripts.
/opt/novell/ncs/bin/ncstempl.py	Used to install or update cluster resource templates in the Cluster container
/opt/novell/ncs/bin/sbdutil	SBD partition utility
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Alert.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Convert.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Down.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Exec.pm	Cluster CLI command

NCS File Name and Path	Purpose
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Info.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Join.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Leave.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Maintenance.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Migrate.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Monitor.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Offline.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Online.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Pools.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Rename.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Resources.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Restart.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Set.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Stats.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Status.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_View.pm	Cluster CLI command
/opt/novell/ncs/bin/uldncs	Unloads NCS; used by the `Cluster Stop` command.
/opt/novell/ncs/lib/ncs-1.0.0.so	NCS snap-in
/opt/novell/ncs/lib/ncsfuncs	Shared library commands for load/unload scripts
/opt/novell/ncs/schema/ncpserver.ldif	NCS schema file
/opt/novell/ncs/schema/ncpserver.preldif	NCS schema file

NCS File Name and Path	Purpose
/opt/novell/ncs/schema/ncs.ldif	NCS schema file
/opt/novell/ncs/schema/ncs.sch	NCS schema file
/opt/novell/oes-install/util/	Path to the extend_schema command
/usr/include/ncssdk.h	NCS SDK
/usr/lib/libncssdk.so	NCS SDK
/usr/lib/libncssdk.so.1.0.0	NCS SDK
/usr/sbin/rcnovell-ncs	Link to /usr/sbin/service
/usr/share/man/man7/sbdutil.7.gz	SBDUTIL man page
/usr/share/man/man8/cluster.8.gz	CLUSTER man page
/usr/share/man/man8/csmport.8.gz	CSMPORT man page

C Electing a Master Node

In a OES Cluster Services cluster, the master node monitors the health of the cluster nodes. It also synchronizes updates about the cluster to eDirectory. The first server that comes up in a cluster is automatically assigned the cluster IP address and becomes the master node. If the master node fails, Cluster Services migrates the cluster IP address to another server in the cluster, and that server becomes the master node.

If a slave node fails to receive heartbeat packages from a master node within a predefined tolerance, the slave assumes that the master has left the cluster or is experiencing problems. The slave node, together with other slave nodes, tries to elect a new master node. The goal is to identify a node with a higher IP address that can be seen by the most member nodes. Each slave node follows the same process to identify a master candidate. If multiple masters are elected, the SBD (split-brain detector) guarantees that only one master survives and fences the other masters and their members.

OES Cluster Services for Open Enterprise Server (OES) 11 SP1 introduces some intelligence in the master election process when the master leaves a cluster (voluntarily or involuntarily). The same master is elected under both the old and new algorithms, but the conclusion is reached sooner, especially for larger clusters. The new algorithm substantially reduces the time needed for master election in some cases. The algorithm is available for OES 11 SP1 and later as well as in the September 2012 Scheduled Maintenance for OES 11 and OES 2 SP3.

This section briefly describes the old and new master-election algorithms.

- Section C.1, “Master-Election Algorithm for OES 11 and Earlier,” on page 411
- Section C.2, “New Master-Election Algorithm for OES 11 SP1 and Later,” on page 413

C.1 Master-Election Algorithm for OES 11 and Earlier

In Cluster Services for OES 11 and earlier, after losing contact with the master node, a slave node attempts to find a node with a higher IP address that will act as master. The logic flow for the master-election algorithm is represented in Figure C-1, “Master-Election Algorithm for OES 11 Initial Release and Earlier Versions,” on page 412.

Before polling other nodes, a node checks its own network adapter status. If the adapter is down, the node immediately promotes itself as the master. Rather than polling and waiting for replies that cannot come back, the node can jump directly to the ultimate conclusion for that node that it is the only node remaining in the cluster. This shortcut can substantially shorten the election process if the network outage is caused by a network adapter failure. However, if the cause for a network outage occurs further upstream by a cable or switch failure, the node follows the election process, unaware that it is awaiting a response that cannot arrive.

If a node’s adapter is working, it begins a process to find a live node with a higher IP address than its own:

1. A node determines if there are other nodes (excluding the old master) with an IP address higher than its own. If it has the highest IP address, it promotes itself as the master.

2. If there are nodes with higher IP addresses, a node asks a member node with the highest IP address (excluding the old master) to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. If a heartbeat package arrives, the node becomes a slave of its elected master.
3. If no heartbeat package arrives, the node then picks the member node with the second highest IP address and repeats the same procedure.
4. If there are no responses from nodes with a higher IP address than its own, a node promotes itself as the master.

If more than one node is elected as master at the end of the process, the SBD guarantees that only one master survives and fences the other masters and their members. Describing how the SBD addresses all possible situations to determine the master is beyond the scope of this document. Let's consider a simple scenario with multiple master candidates. The SBD kills any master candidate node whose adapter is down. If a master candidate was the old master, it becomes the new master. Otherwise, the master candidate with the highest IP address wins.

The master election process can take time because a node sequentially asks each potential master node to become the master and waits for a response before trying the next higher IP address. If a large cluster (8+ nodes) loses LAN connections among all the nodes, it can take up to 3 minutes to elect all of the new masters. The node with lowest IP address tries almost all of the other nodes.

Figure C-1 *Master-Election Algorithm for OES 11 Initial Release and Earlier Versions*

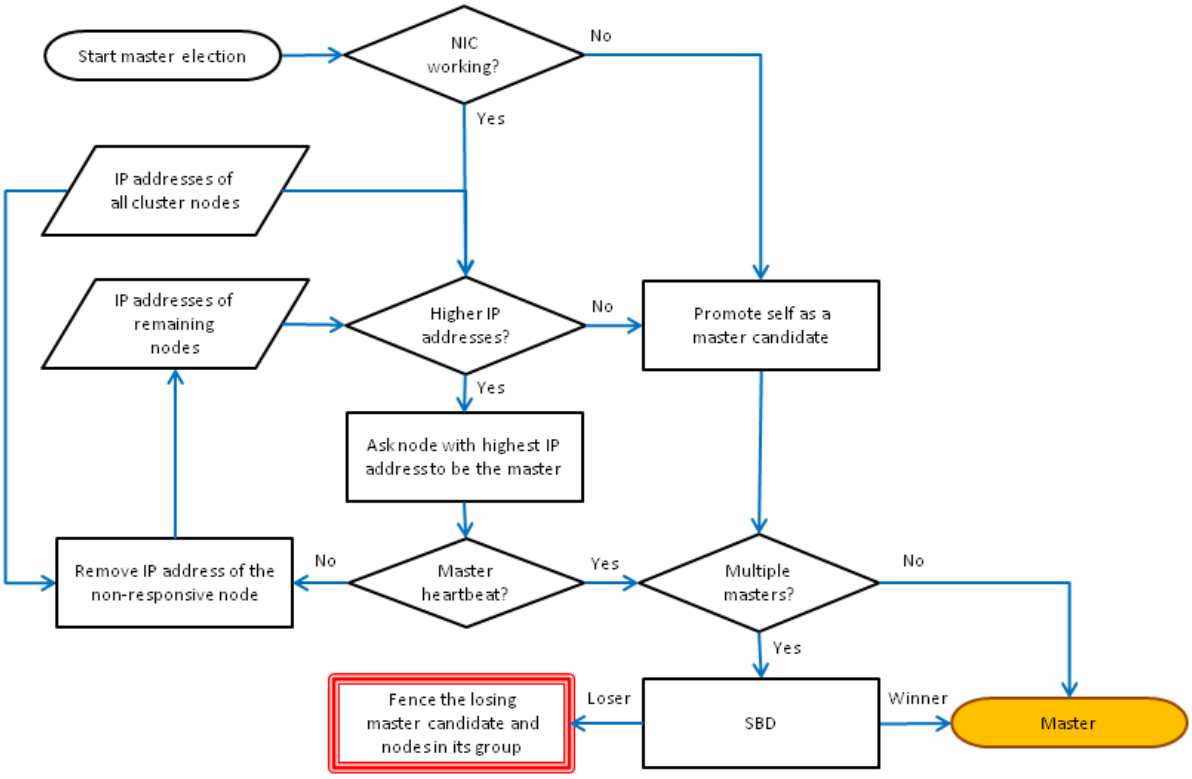

C.2 New Master-Election Algorithm for OES 11 SP1 and Later

In Cluster Services for OES 11 SP1 and later, after losing contact with the master node, a slave node attempts to find a node with a higher IP address that will act as master, but it also uses intelligence to determine which other nodes are capable of responding. The new algorithm elects the same master as the old algorithm, but the time to make the decision is deterministic and the process is more efficient. The logical flow for the master-election algorithm is represented in Figure C-2, “Master-Election Algorithm for OES 11 SP1 and Later,” on page 414.

In the new master-election algorithm, after a slave determines that its adapter is working, it broadcasts a ping to all of the cluster nodes. The purpose of pinging is to identify whether the node can communicate with other nodes and which ones. This helps address a situation where LAN communications are down between some nodes, such as a cable or switch failure.

If there are nodes with higher IP addresses than its own, a node asks a member node with the highest IP address (excluding the old master) to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. If a heartbeat package arrives, the node becomes a slave of its elected master.

Sending the first master request in the old manner serves two purposes:

- It maintains some compatibility with older versions of Cluster Services.
- It provides some action while waiting for the ping replies to come in.

If the member with the highest IP address starts acting as the new master (in time), the election is over for this node.

If no heartbeat packages arrive, the node then focuses on member nodes with higher IP addresses that have either pinged or replied to its ping. This helps the node by-pass waiting periods for requests to nodes that are unable to communicate. The node picks the member node with the highest IP address, and asks it to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. The request to this node is almost certain to succeed since it is known to be able to communicate.

If no heartbeat packages arrive, the node picks the member node with the next highest IP address that has either pinged or replied to the ping, and the election goes on.

If there are no responses from nodes with a higher IP address than its own, a node promotes itself as the master.

If more than one node is elected as master at the end of the process, the SBD guarantees that only one master survives and fences the other masters and their members.

The benefits of the new master-election algorithm are:

- The new algorithm elects the same master that would be determined with the old algorithm.
- The time to elect a master node is deterministic, regardless of the size of the cluster or the nature of the problem.
- The traditional master-election logic is mostly preserved.
- The first master request goes out while waiting for nodes to respond to pings, or to time out.

- The second master request goes to a node that is known to be able to communicate, and is almost certain to succeed.
- Election time is substantially decreased in some cases.

Figure C-2 *Master-Election Algorithm for OES 11 SP1 and Later*

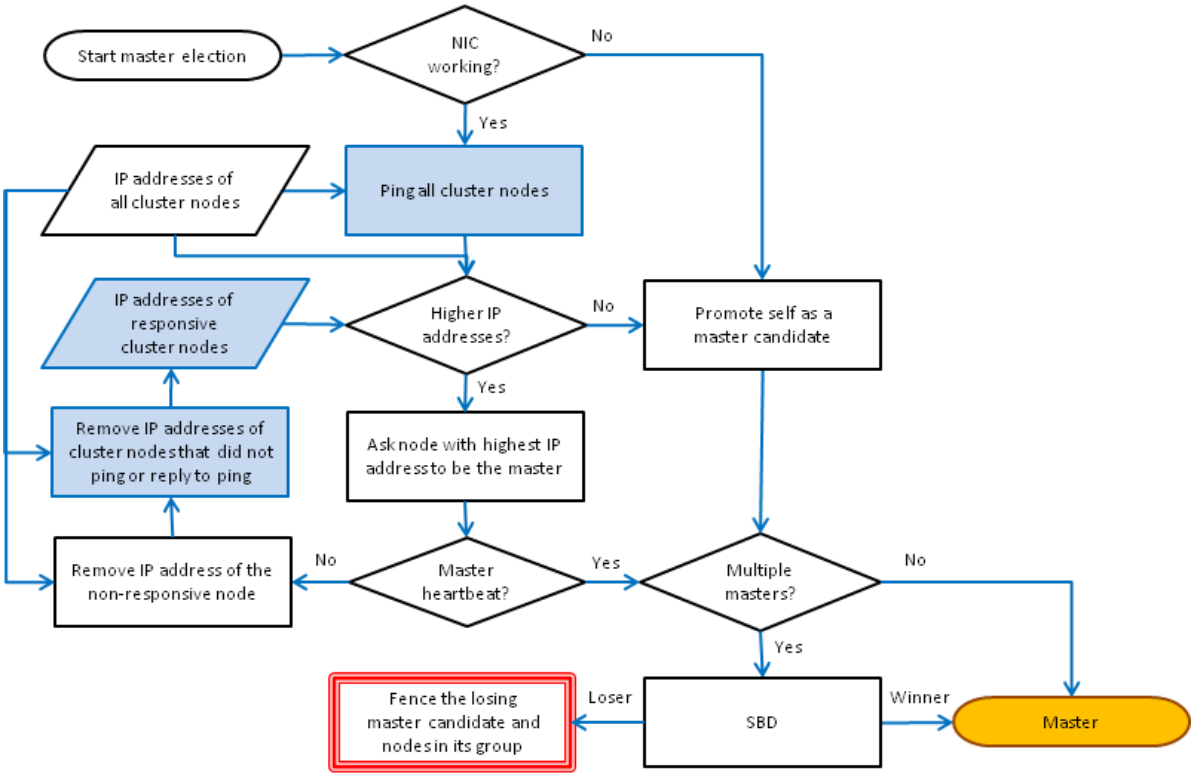

D Clusters Plug-In Changes for NetIQ iManager 3.2.6

The Clusters plug-in for iManager supports the management of OES clusters and resources. The availability of OES Cluster Services features depends on the version of OES Cluster Services and the server platform that are installed on the cluster being managed. You can use either the old plug-in or the new plug-in to manage a cluster and its resources.

Understanding the Cluster Plug-In's New Interface

The Clusters plug-in for iManager has been reorganized. In **Roles and Tasks** under **Clusters**, the **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, and **Cluster Options** menu items have been replaced with two menu options:

- **My Clusters:** The logged-in cluster administrator can set up a personalized list of clusters to manage. This allows an administrator to view at a glance the status of multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of clusters and display preferences persist between the administrator's logins to iManager on the same server. See Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 99.
- **My Resources:** The logged-in cluster administrator can set up a personalized list of cluster resources to manage. This allows an administrator to view at a glance the status of multiple cluster resources for multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of cluster resources and display preferences persist between the administrator's logins to iManager on the same server. See Section 10.2, "Setting Up a Personalized List of Resources to Manage," on page 165.

From a custom My Clusters page or a My Resources page, you can click the name link of a cluster to manage the cluster. The old plug-in's **Clusters** menu options are available here as tabs: **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, and **Cluster Options**.

Upgrading the Clusters Plug-In

If you use Role-Based Services (RBS), upgrading the Clusters plug-in does not automatically update the RBS settings. The RBS Configuration page reports that the Clusters plug-in is out-of-date. The plug-in must be reinstalled on the RBS Configuration page in order to pick up the **My Clusters** and **My Resources** menu options. See Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In after Upgrading to OES 11 SP2 and Later," on page 74.

Comparing Tasks in the Old and New Clusters Plug-Ins

Use the tables in this section to understand how to perform tasks in the old Clusters plug-in for iManager 2.7.4 and the new Clusters plug-in for iManager 2.7.5 and later.

- Table D-1, "Selecting a Cluster to Manage," on page 416
- Table D-2, "Managing a Cluster," on page 417

- Table D-3, “Accessing Cluster Properties,” on page 418
- Table D-4, “Managing a Cluster Resource or Resource Properties,” on page 419

Table D-1 describes how to select the cluster you want to manage.

Table D-1 *Selecting a Cluster to Manage*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
1. In **Roles and Tasks**, select **Clusters**, then select **Cluster Manager**, **BCC Manager**, **Cluster Event Log**, or **Cluster Options**. 2. Specify the cluster name, or browse and select the Cluster object. You must repeat the selection process to manage a different cluster. There is no capability to concurrently view the status of multiple clusters.	1. In **Roles and Tasks**, select **Clusters > My Clusters**. The list is initially empty. 2. On the My Clusters page, click **Add** to open the eDirectory browser pop-up window. 3. Browse the tree where you are currently logged in to locate and select one or more clusters that you want to manage, then click **OK**. The selected clusters are added to the list. Your personalized list is displayed each time you log in to this iManager server and return to the My Clusters page. Use the **Add** and **Remove** options to modify your list as needed. 4. Do any of the following from this page: ◆ Concurrently view the status of multiple clusters. ◆ Click the name link of a cluster to manage it. The Cluster Manager page opens. You can also select the other tabs from this page: **BCC Manager**, **Cluster Event Log**, or **Cluster Options**. ◆ Select the check box next to the cluster, then select **Actions > Edit Properties** to manage the cluster properties. ◆ Select the check box next to the cluster, then select **Actions > Run Report**.

Table D-2 describes how to access the cluster management pages: Cluster Manager, BCC Manager, Cluster Event Log, and Cluster Options. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in Table D-1.

Table D-2 *Managing a Cluster*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
Cluster Manager 1. In **Roles and Tasks**, select **Clusters > Cluster Manager**. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Manager** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Manager** tab.
BCC Manager 1. In **Roles and Tasks**, select **Clusters > BCC Manager**. 2. Specify the cluster name, or browse and select the Cluster object.	**BCC Manager** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **BCC Manager** tab.
Cluster Event Log 1. In **Roles and Tasks**, select **Clusters > Cluster Event Log**. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Event Log** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Event Log** tab.
Cluster Options 1. In **Roles and Tasks**, select **Clusters > Cluster Options**. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Options** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Options** tab.

Table D-3 describes how to access the cluster properties pages: Policies, Priorities, Protocols, Resource Mutual Exclusion, and Business Continuity. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in Table D-1.

Table D-3 *Accessing Cluster Properties*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
	From My Resources 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Select the check box next to the cluster, then click **Actions > Edit Properties.**
From Cluster Options 1. In **Roles and Tasks**, select **Clusters > Cluster Options**. 2. Specify the cluster name, or browse and select the Cluster object. 3. Click **Properties**.	**From Cluster Options** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Options** tab. 4. Click **Properties**.

Table D-4 describes how to view or manage cluster resources, including their properties: Policies, Monitoring, Preferred Nodes, Scripts (Load, Unload, and Monitor), and Business Continuity.

Table D-4 *Managing a Cluster Resource or Resource Properties*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
There is no capability to concurrently view the status of cluster resources across multiple clusters.	**From My Resources:** 1. In **Roles and Tasks**, select **Clusters > My Resources**. The initial list is empty. 2. On the My Resources page, click **Add** to open the eDirectory browser pop-up window. 3. Browse the tree where you are currently logged in to locate and select the cluster resources that you want to manage for one or more clusters, then click **OK**. The selected cluster resources are added to the list. Your personalized list is displayed each time you log in to this iManager server and return to the My Resources page. Use the **Add** and **Remove** options to modify your list as needed. 4. Do any of the following from this page: ◆ Concurrently view the status of multiple resources across multiple clusters. ◆ Click the name link of a resource to open its **Resource Properties**. ◆ Click the name link of the resource's cluster to manage the cluster.

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
From Cluster Manager: 1. In **Roles and Tasks**, select **Clusters > Cluster Manager**. 2. Specify the cluster name, or browse and select the Cluster object. 3. Do any of the following: ◆ View the status of resources on this cluster. ◆ Click the name link of a resource to open its **Resource Properties**. ◆ Select the check box next to a resource, then click one of the following options: Online Offline Migrate Respond to Alert	**From Cluster Manager:** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Manager** tab. 4. Do any of the following: ◆ View the status of resources on this cluster. ◆ Click the name link of a resource to open its **Resource Properties**. ◆ Select the check box next to a resource, then click one of the following options: Online Offline Migrate Respond to Alert
From Cluster Options: 1. In **Roles and Tasks**, select **Clusters > Cluster Options**. 2. Specify the cluster name, or browse and select the Cluster object. 3. Click the name link of a resource to open its **Resource Properties**. You can also select the check box next to the resource, then click the **Details** link.	**From Cluster Options:** 1. In **Roles and Tasks**, select **Clusters > My Clusters**. 2. Click the name link of the cluster. 3. Select the **Cluster Options** tab. 4. Click the name link of a cluster resource to open its **Resource Properties**. You can also select the check box next to the resource, then click the **Details** link.