

Novell Cluster Services™ 1.8.7 For Linux Administration Guide

Novell® Open Enterprise Server

2 SP2

August 9, 2010

www.novell.com



Legal Notices

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

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

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

Copyright © 2007–2010 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide	13
1 Overview of Novell Cluster Services	15
1.1 Why Should I Use Clusters?	15
1.2 Benefits of Novell Cluster Services	15
1.3 Product Features	16
1.4 Clustering for High-Availability	16
1.5 Shared Disk Scenarios	18
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
2 What's New	23
2.1 What's New (July 2010 Patches)	23
2.2 What's New (June 2010 Patches)	23
2.3 What's New (January 2010 Patches)	23
2.4 What's New (OES 2 SP2)	24
2.4.1 Improved Error Reporting	24
2.4.2 Improved Time Calculations	24
2.4.3 Specifying the Size of the SBD Partition	24
2.4.4 Customizing Translation Syntax	24
2.4.5 Migration Tool Support for Cluster Configurations	25
2.4.6 New iFolder Resource Template	25
2.4.7 Removed MPK Calls from the Code	25
2.4.8 Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade	25
2.5 What's New (OES 2 SP1)	25
2.5.1 Schema Extension	26
2.5.2 Installation by Container Administrator	26
2.5.3 Behavior Change for Adding a Node	26
2.5.4 Attribute NCS: GIPC Config Is No Longer Maintained	27
2.5.5 Support for Novell AFP for Linux	27
2.5.6 Support for Novell CIFS for Linux	27
2.5.7 Support for Domain Services for Windows	27
2.6 What's New (OES 2)	27
3 Installing and Configuring Novell Cluster Services on OES 2 Linux	29
3.1 System Requirements for Novell Cluster Services	29
3.1.1 Hardware Requirements	29
3.1.2 Software Requirements	30
3.1.3 Configuration Requirements	36
3.1.4 Shared Disk System Requirements	38
3.1.5 SAN Rules for LUN Masking	40
3.1.6 Using Device Mapper Multipath with Novell Cluster Services	40
3.2 Novell Cluster Services Licensing	42
3.3 Extending the eDirectory Schema to Add Cluster Objects	42
3.3.1 Prerequisites for Extending the Schema	42
3.3.2 Extending the Schema	43

3.4	Assigning Install Rights for Container Administrators	44
3.5	Installing Novell Cluster Services.	44
3.5.1	Installing Novell Cluster Services during a OES 2 Linux Installation	45
3.5.2	Installing Novell Cluster Services on an Existing OES 2 Linux Server	47
3.6	Configuring Novell Cluster Services.	48
3.6.1	Opening the Novell Open Enterprise Server Configuration Page	49
3.6.2	Using Different LDAP Credentials for the Cluster Configuration	50
3.6.3	Enabling the Novell Cluster Services Configuration Page in YaST	51
3.6.4	Configuring a New Cluster.	51
3.6.5	Adding a Node to an Existing Cluster	54
3.7	Configuring Additional Administrators	55
3.8	What's Next	56
4	Upgrading OES 2 Linux Clusters	57
4.1	Requirements for Upgrading Clusters	57
4.2	Upgrading OES 2 Clusters (Rolling Cluster Upgrade)	57
4.3	Upgrade Issues for OES 2 SP2	58
4.3.1	Updating the iFolder Resource Template	58
5	Upgrading OES 1 Linux Clusters to OES 2 Linux	59
5.1	Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux	59
5.2	Upgrading Existing OES 1 Linux Cluster Nodes to OES 2 (Rolling Cluster Upgrade)	60
5.3	Adding New OES 2 Linux Cluster Nodes to Your OES 1 Linux Cluster.	61
5.4	Modifying Cluster Resource Scripts for Mixed OES 1 Linux and OES 2 Linux Clusters	61
5.5	Finalizing the Cluster Upgrade.	61
6	Converting NetWare 6.5 Clusters to OES 2 Linux	63
6.1	Guidelines for Converting Clusters from NetWare to OES 2 Linux	63
6.1.1	Supported Mixed-Node Clusters	64
6.1.2	SBD Devices Must Be Marked as Shareable for Clustering	64
6.1.3	Syntax Translation Issues for Load and Unload Scripts	64
6.1.4	Case Sensitivity Issues	65
6.1.5	Adding New NetWare Nodes to a Mixed-Node Cluster	65
6.1.6	Converting Multiple NetWare Cluster Nodes to OES 2 Linux	65
6.1.7	Converting Nodes that Contain the eDirectory Master Replica	66
6.1.8	Failing Over Cluster Resources on Mixed-Node Clusters	66
6.1.9	Managing File Systems in Mixed-Node Clusters	66
6.1.10	Using Novell iManager in Mixed-Node Clusters	66
6.1.11	Using Novell Remote Manager Is Not Supported in Mixed-Node Clusters	67
6.1.12	Using ConsoleOne Is Not Supported for Mixed-Node Clusters.	67
6.1.13	Using the Monitor Function in Mixed-Node Clusters Is Not Supported	67
6.2	Guidelines for Converting NSS Pool Resources from NetWare to Linux.	67
6.2.1	NSS Pool Cluster Migration	67
6.2.2	NSS File System Migration to NCP Volumes or Linux POSIX File Systems	68
6.2.3	Estimated Time Taken to Build the Trustee File on Linux	68
6.3	Guidelines for Converting Service Cluster Resources from NetWare to Linux	68
6.3.1	Overview of All NetWare 6.5 SP8 Services	69
6.3.2	Apache Web Server.	71
6.3.3	Apple Filing Protocol (AFP)	71
6.3.4	Archive and Version Services	71
6.3.5	CIFS.	72
6.3.6	DFS VLDB	72
6.3.7	DHCP Server	73

6.3.8	DNS Server	75
6.3.9	eDirectory Server Certificates	76
6.3.10	iPrint	78
6.3.11	QuickFinder Server	79
6.4	Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion)	82
6.5	Adding New OES 2 Linux Nodes to Your NetWare Cluster	85
6.6	Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters	86
6.6.1	Comparing Script Commands for NetWare and Linux	86
6.6.2	Comparing Master IP Address Scripts	87
6.6.3	Comparing NSS Pool Resource Scripts	88
6.6.4	Comparing File Access Protocol Resource Script Commands	89
6.7	Customizing the Translation Syntax for Converting Load and Unload Scripts	91
6.8	Finalizing the Cluster Conversion	92
7	Configuring Cluster Policies and Priorities	95
7.1	Understanding Cluster Settings	95
7.1.1	Cluster Policies	95
7.1.2	Cluster Protocols Properties	96
7.2	Configuring Quorum Membership and Timeout Properties	96
7.2.1	Quorum Triggers (Number of Nodes)	97
7.2.2	Quorum Triggers (Timeout)	97
7.3	Configuring Cluster Protocol Properties	97
7.3.1	Heartbeat	98
7.3.2	Tolerance	98
7.3.3	Master Watchdog	98
7.3.4	Slave Watchdog	98
7.3.5	Maximum Retransmits	98
7.4	Configuring Cluster Event E-Mail Notification	99
7.5	Viewing the Cluster Node Properties	100
7.6	Modifying the Cluster IP Address and Port Properties	100
7.7	What's Next	100
8	Managing Clusters	101
8.1	Starting and Stopping Novell Cluster Services	101
8.1.1	Starting Novell Cluster Services	102
8.1.2	Stopping Novell Cluster Services	102
8.1.3	Enabling and Disabling the Automatic Start of Novell Cluster Services	102
8.2	Monitoring Cluster and Resource States	103
8.3	Generating a Cluster Configuration Report	105
8.4	Cluster Migrating Resources to Different Nodes	106
8.5	Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node	106
8.6	Removing (Leaving) a Node from the Cluster	107
8.7	Joining a Node to the Cluster	107
8.8	Configuring the EVMS Remote Request Timeout	107
8.9	Shutting Down Linux Cluster Servers When Servicing Shared Storage	108
8.10	Enabling or Disabling Cluster Maintenance Mode	108
8.11	Preventing a Cluster Node Reboot after a Node Shutdown	109
8.11.1	OES 2 SP2 with Patches and Later	109
8.11.2	OES 2 SP2 Release Version and Earlier	109
8.12	Renaming a Pool for a Pool Cluster Resource	110
8.13	Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster	110

8.13.1	Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster . . .	110
8.13.2	Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources . . .	111
8.14	Adding a Node That Was Previously in the Cluster	114
8.15	Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node	114
8.16	Creating or Deleting Cluster SBD Partitions	115
8.16.1	Prerequisites for Creating an SBD Partition	115
8.16.2	Before Creating an Cluster SBD Partition	116
8.16.3	Creating a Non-Mirrored Cluster SBD Partition	116
8.16.4	Creating a Mirrored Cluster SBD Partition	117
8.16.5	Deleting a Non-Mirrored Cluster SBD Partition	118
8.16.6	Deleting a Mirrored Cluster SBD Partition	119
8.16.7	Removing a Segment from a Mirrored Cluster SBD Partition	119
8.17	Customizing Cluster Services Management	120

9 Configuring and Managing Cluster Resources 123

9.1	Planning Cluster Maintenance	123
9.2	Naming Conventions for Cluster Resources	123
9.3	Creating Cluster Resource Templates	124
9.3.1	Default Resource Templates	124
9.3.2	Creating a Resource Template	125
9.4	Creating Cluster Resources	127
9.5	Configuring a Load Script for a Cluster Resource	127
9.6	Configuring an Unload Script for a Cluster Resource	128
9.7	Enabling Monitoring and Configuring the Monitor Script	129
9.7.1	Configuring Resource Monitoring	130
9.7.2	Example Monitoring Scripts	131
9.7.3	Monitoring Services Critical to Clustering	131
9.8	Setting Start, Failover, and Failback Modes for Cluster Resources	131
9.8.1	Understanding Cluster Resource Modes	132
9.8.2	Viewing or Modifying the Start, Failover, and Failback Modes for a Resource	132
9.9	Assigning Nodes to a Resource	133
9.10	Configuring Resource Priorities for Load Order	134
9.11	Changing the IP Address of a Cluster Resource	134
9.12	Deleting Cluster Resources	134
9.12.1	Deleting a Cluster Resource on a Master Node	135
9.12.2	Deleting a Cluster Resource on a Non-Master Node	135
9.13	Additional Information for Creating Cluster Resources	136
9.13.1	Creating Storage Cluster Resources	137
9.13.2	Creating Service Cluster Resources	137
9.13.3	Creating Virtual Machine Cluster Resources	137

10 Configuring Cluster Resources for Shared NSS Pools and Volumes 139

10.1	Planning for Shared NSS Pools and Volumes	139
10.1.1	Shared Storage	140
10.1.2	Novell Cluster Services	140
10.1.3	Novell Storage Services	140
10.1.4	IP Address	141
10.1.5	NCP Server for Linux	141
10.1.6	Novell CIFS for Linux	141
10.1.7	Novell AFP for Linux	142
10.2	Considerations for Working with Shared NSS Pools and Volumes in the Cluster	142
10.3	Creating NSS Shared Disk Partitions and Pools	143

10.3.1	Initializing Shared Devices	143
10.3.2	Enabling Sharing on a Device	143
10.3.3	Creating Shared NSS Pools	144
10.4	Creating NSS Volumes on a Shared Pool	148
10.4.1	Using iManager to Create NSS Volumes	148
10.4.2	Using NSSMU to Create NSS Volumes	149
10.5	Cluster-Enabling an Existing NSS Pool and Its Volumes	150
10.6	Adding Advertising Protocols	153
10.7	Configuring a Load Script for the Shared NSS Pool	155
10.8	Configuring an Unload Script for the Shared NSS Pool	156
10.9	Configuring a Monitor Script for the Shared NSS Pool	156
10.10	Mirroring and Cluster-Enabling Shared NSS Pools and Volumes	156
10.10.1	Understanding NSS Mirroring	157
10.10.2	Requirements for NSS Mirroring	158
10.10.3	Creating and Mirroring NSS Pools on Shared Storage	159
10.10.4	Verifying the NSS Mirror Status in the Cluster	160
10.11	Deleting NSS Pool Cluster Resources	161
10.12	Changing the Volume ID	161
10.13	What's Next	162

11 Configuring Cluster Resources for Shared Linux POSIX Volumes 163

11.1	Requirements for Shared Linux POSIX Volumes	163
11.2	Creating Linux POSIX Volumes on Shared Disks	164
11.2.1	Removing Existing Formatting and Segment Managers	164
11.2.2	Creating a Cluster Segment Manager Container	165
11.2.3	Adding a Non-CSM Segment Manager Container	166
11.2.4	Creating an EVMS Volume	167
11.2.5	Making a File System on the EVMS Volume	167
11.3	Cluster-Enabling a Linux POSIX Volume on a Shared Disk	168
11.3.1	Logging in to iManager	169
11.3.2	Creating a Cluster Resource for a Linux POSIX Volume	169
11.3.3	Configuring a Load Script for a Linux POSIX Volume Cluster Resource	170
11.3.4	Configuring an Unload Script for a Linux POSIX Volume Cluster Resource	171
11.3.5	Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource	172
11.3.6	Configuring Policies for a Linux POSIX Volume Cluster Resource	173
11.4	Creating a Virtual Server Name for the Cluster Resource	175
11.4.1	Using ncs_ncpserv.py to Create an NCS:NCP Server Object	176
11.4.2	Using iManager to Create an NCS:NCP Server Object	177
11.5	Sample Scripts for a Linux POSIX Volume Cluster Resource	179
11.5.1	Sample Load Script for the Linux POSIX Volume Cluster Resource	180
11.5.2	Sample Unload Script for the Linux POSIX Volume Cluster Resource	180
11.5.3	Sample Monitor Script for a Linux POSIX Volume Cluster Resource	181
11.6	Expanding EVMS Volumes on Shared Disks	181
11.6.1	Expanding a Volume to a Separate Disk	181
11.6.2	Moving a Volume to a Larger Disk	182
11.7	Deleting Shared Storage	183
11.8	Known Issues for Working with Cluster Resources for Linux POSIX Volumes	183
11.8.1	Dismount Volumes before Onlining a Comatose Resource	183
11.8.2	Cluster Services Must Be Running When Using EVMS	183
11.8.3	Close EVMS Utilities When They Are Not In Use	183
11.8.4	Do Not Migrate Resources When EVMS Tools Are Running	184
11.9	What's Next	184

12	Configuring Novell Cluster Services in a Xen Virtualization Environment	185
12.1	Prerequisites for Xen Host Server Environments	185
12.2	Virtual Machines as Cluster Resources	186
12.2.1	Creating a Xen Virtual Machine Cluster Resource	187
12.2.2	Configuring Virtual Machine Load, Unload, and Monitor Scripts	188
12.2.3	Setting Up Live Migration.	193
12.3	Virtual Machines as Cluster Nodes	193
12.4	Virtual Cluster Nodes in Separate Clusters	194
12.5	Mixed Physical and Virtual Node Clusters	195
12.6	Additional Information	197
13	Troubleshooting Novell Cluster Services	199
13.1	File Location Information	199
13.2	Version Issues	200
13.3	Diagnosing a Poison Pill Condition	200
13.4	Diagnosing Cluster Problems.	200
13.5	Cluster Search Times Out (Bad XML Error).	201
13.6	A Device Name Is Required to Create a Cluster Partition	202
13.7	Cluster Resource Goes Comatose Immediately After Migration or Failover	202
13.8	smdr.novell Is Not Registered with SLP for a New Cluster Resource	202
13.9	Cluster View Displays the Wrong Cluster Node Name	202
13.10	NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)	203
13.11	Could Not Delete This Resource Data_Server (Error 499)	203
13.12	Problem Authenticating to Remote Servers during Cluster Configuration	203
13.13	Problem Connecting to an iSCSI Target	203
13.14	Problem Deleting a Cluster Resource or Clustered Pool	204
13.15	Can't Find the Prevent Cascading Failover Option	204
13.16	Is there a way to uninstall Novell Cluster Services from a server?	204
13.17	Cluster Resource Is Stuck in "NDS Sync" or "eDirectory Sync" State	204
13.18	Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems	205
13.19	Nodes Reboot Randomly after January or March 2010 Updates.	205
14	Security Considerations	207
14.1	Cluster Administration Rights.	207
14.2	Ports	207
14.3	E-Mail Alerts.	207
14.4	Log Files.	208
A	Console Commands for Novell Cluster Services	209
A.1	Cluster Management Commands	209
A.2	Business Continuity Clustering Commands	212
A.3	extend_schema Command	212
A.4	SBD Utility	213

B	Files for Novell Cluster Services	215
C	Comparing Novell Cluster Services for Linux and NetWare	219
D	Comparing Clustering Support for OES 2 Services on Linux and NetWare	225
E	Documentation Updates	231
E.1	August 2, 2010 (Maintenance Patch)	231
E.1.1	Configuring Cluster Resources for Shared Linux POSIX Volumes	232
E.1.2	Configuring Cluster Resources for Shared NSS Pools and Volumes	232
E.1.3	Converting NetWare 6.5 Clusters to OES 2 Linux	232
E.1.4	Managing Clusters	233
E.1.5	What's New	233
E.2	June 14, 2010 (Maintenance Patch)	233
E.2.1	Configuring and Managing Cluster Resources	233
E.2.2	Configuring Cluster Policies and Priorities	234
E.2.3	Console Commands for Novell Cluster Services	234
E.2.4	Converting NetWare 6.5 Clusters to OES 2 Linux	234
E.2.5	Installing and Configuring Novell Cluster Services on OES 2 Linux	234
E.2.6	Managing Clusters	235
E.2.7	Troubleshooting Novell Cluster Services	235
E.2.8	Upgrading OES 2 Linux Clusters	236
E.2.9	What's New	236
E.3	March 15, 2010	236
E.3.1	Configuring and Managing Cluster Resources	236
E.3.2	Configuring Cluster Resources for Shared Linux POSIX Volumes	237
E.3.3	Configuring Novell Cluster Services in a Xen Virtualization Environment	237
E.3.4	Troubleshooting Novell Cluster Services	238
E.4	February 19, 2010	238
E.4.1	Installing and Configuring Novell Cluster Services on OES 2 Linux	238
E.4.2	Managing Clusters	238
E.5	February 10, 2010	238
E.5.1	Configuring Cluster Resources for Linux POSIX Volumes	239
E.6	January 29, 2010 (Maintenance Patch)	239
E.6.1	Converting NetWare 6.5 Clusters to OES 2 Linux	239
E.6.2	Installing and Configuring Novell Cluster Services on OES 2 Linux	239
E.6.3	Managing Clusters	240
E.6.4	Troubleshooting Novell Cluster Services	240
E.6.5	What's New	240
E.7	January 20, 2010	240
E.7.1	Converting NetWare 6.5 Clusters to OES 2 Linux	241
E.7.2	What's New	241
E.8	January 4, 2010	241
E.8.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	241
E.8.2	Converting NetWare 6.5 Clusters to OES 2 Linux	242
E.9	December 15, 2009	242
E.9.1	Installing and Configuring Novell Cluster Services on OES 2 Linux	242
E.10	December 10, 2009	242
E.10.1	Converting NetWare 6.5 Clusters to OES 2 Linux	242
E.10.2	Installing and Configuring Novell Cluster Services on OES 2 Linux	243
E.10.3	Managing Clusters	243
E.10.4	Troubleshooting Novell Cluster Services	243
E.11	November 2009 (OES 2 SP2)	243
E.11.1	Configuring Novell Cluster Services in a Xen Host Environment	244

E.11.2	Converting NetWare 6.5 Clusters to OES 2 Linux	244
E.11.3	Console Commands for Novell Cluster Services	244
E.11.4	Installing Novell Cluster Services on OES 2 Linux	244
E.11.5	Managing Clusters	245
E.11.6	Troubleshooting Novell Cluster Services	245
E.11.7	Upgrading OES 2 Linux Clusters	245
E.11.8	What's New	245
E.12	July 30, 2009	245
E.12.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	246
E.12.2	Console Commands for Novell Cluster Services	246
E.12.3	Converting NetWare 6.5 Clusters to OES 2 Linux	246
E.12.4	Installing Novell Cluster Services on OES 2 Linux	247
E.12.5	Managing Clusters	247
E.12.6	Security Considerations	247
E.12.7	Troubleshooting Novell Cluster Services	247
E.12.8	What's New	247
E.13	June 22, 2009	248
E.13.1	Converting NetWare 6.5 Clusters to OES 2 Linux	248
E.13.2	Managing Clusters	248
E.14	June 5, 2009	248
E.14.1	Configuring Cluster Policies and Priorities	248
E.14.2	Configuring Cluster Resources for Shared NSS Pools and Volumes	249
E.14.3	Console Commands for Novell Cluster Services	249
E.14.4	Converting NetWare 6.5 Clusters to OES 2 Linux	249
E.14.5	Installing Novell Cluster Services on OES 2 Linux	250
E.14.6	Managing Clusters	250
E.15	May 6, 2009	250
E.15.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	251
E.15.2	Configuring and Managing Cluster Resources	251
E.15.3	Converting NetWare 6.5 Clusters to OES 2 SP1 Linux	251
E.15.4	Installing Novell Cluster Services on OES 2 Linux	252
E.15.5	Troubleshooting Novell Cluster Services	252
E.15.6	Upgrading OES 2 Linux Servers	252
E.16	March 3, 2009	252
E.16.1	Configuring and Managing Cluster Resources	253
E.16.2	Configuring Cluster Resources for Shared NSS Pools and Volumes	253
E.16.3	Converting NetWare 6.5 Clusters to OES 2 Linux	253
E.16.4	Managing Clusters	253
E.16.5	What's Next	253
E.17	February 13, 2009	254
E.17.1	Converting NetWare 6.5 Clusters to OES 2 Linux	254
E.18	February 3, 2009	254
E.18.1	Configuring and Managing Cluster Resources	255
E.18.2	Configuring Cluster Resources for Shared NSS Pools and Volumes	255
E.18.3	Converting NetWare 6.5 Clusters to OES 2 Linux	255
E.18.4	Installing Novell Cluster Services on OES 2 Linux	255
E.19	January 13, 2009	255
E.20	December 2008 (OES 2 SP1)	256
E.20.1	Comparison of Clustering OES 2 Services for Linux and NetWare	256
E.20.2	Comparison of Novell Cluster Services for Linux and NetWare	256
E.20.3	Configuring and Managing Cluster Resources	256
E.20.4	Configuring Cluster Resources for Shared Linux POSIX Volumes	257
E.20.5	Configuring Cluster Resources for Shared NSS Pools and Volumes	257
E.20.6	Configuring Novell Cluster Services in a Virtualization Environment	258
E.20.7	Console Commands for Novell Cluster Services	258
E.20.8	Converting NetWare 6.5 Cluster to OES 2 Linux	258
E.20.9	Installing Novell Cluster Services on OES 2 Linux	259

E.20.10	Managing Clusters	259
E.20.11	Overview of Novell Cluster Services	259
E.20.12	Troubleshooting Novell Cluster Services	260
E.20.13	Upgrading OES 2 Clusters	260
E.20.14	What's New	260
E.21	June 4, 2008	260
E.21.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	261
E.21.2	Configuring Cluster Resources for Shared Linux POSIX Volumes	261
E.21.3	Installation and Setup	262
E.21.4	Managing Novell Cluster Services	262
E.22	May 2, 2008	263
E.22.1	Installation and Setup	263
E.22.2	Managing Novell Cluster Services	264

About This Guide

This guide describes how to install, upgrade, configure, and manage Novell Cluster Services. It is intended for cluster administrators and is divided into the following sections:

- ♦ Chapter 1, “Overview of Novell Cluster Services,” on page 15
- ♦ Chapter 2, “What’s New,” on page 23
- ♦ Chapter 3, “Installing and Configuring Novell Cluster Services on OES 2 Linux,” on page 29
- ♦ Chapter 4, “Upgrading OES 2 Linux Clusters,” on page 57
- ♦ Chapter 5, “Upgrading OES 1 Linux Clusters to OES 2 Linux,” on page 59
- ♦ Chapter 6, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 63
- ♦ Chapter 7, “Configuring Cluster Policies and Priorities,” on page 95
- ♦ Chapter 8, “Managing Clusters,” on page 101
- ♦ Chapter 9, “Configuring and Managing Cluster Resources,” on page 123
- ♦ Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139
- ♦ Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163
- ♦ Chapter 12, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 185
- ♦ Chapter 13, “Troubleshooting Novell Cluster Services,” on page 199
- ♦ Appendix A, “Console Commands for Novell Cluster Services,” on page 209
- ♦ Appendix B, “Files for Novell Cluster Services,” on page 215
- ♦ Appendix C, “Comparing Novell Cluster Services for Linux and NetWare,” on page 219
- ♦ Appendix D, “Comparing Clustering Support for OES 2 Services on Linux and NetWare,” on page 225

Audience

This guide is intended for anyone involved in installing, configuring, and managing Novell Cluster Services.

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, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

The latest version of this *Novell Cluster Services for Linux Administration Guide* is available on the OES 2 documentation Web site (<http://www.novell.com/documentation/oes2/cluster-services.html>).

Additional Documentation

For information about creating cluster resources for various Linux services on your OES 2 Linux server, refer to the clustering sections in the individual guides. See the “Clustering Linux Services” list on the Clustering (High Availability) Documentation Web site (<http://www.novell.com/documentation/oes2/cluster-services.html#clust-config-resources>).

For information about Novell Cluster Services 1.8.5 for NetWare, see the “Clustering NetWare Services” list on the NetWare 6.5 SP8 Clustering (High Availability) Documentation Web site (<http://www.novell.com/documentation/nw65/cluster-services.html#clust-config-resources>).

Overview of Novell Cluster Services

1

Novell Cluster Services is a server clustering system that ensures high availability and manageability of critical network resources including data, applications, and services. It is a multi-node clustering product for Linux that is enabled for Novell eDirectory and supports failover, failback, and migration (load balancing) of individually managed cluster resources.

- ◆ [Section 1.1, “Why Should I Use Clusters?,” on page 15](#)
- ◆ [Section 1.2, “Benefits of Novell Cluster Services,” on page 15](#)
- ◆ [Section 1.3, “Product Features,” on page 16](#)
- ◆ [Section 1.4, “Clustering for High-Availability,” on page 16](#)
- ◆ [Section 1.5, “Shared Disk Scenarios,” on page 18](#)

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. Shared devices are connected to and mounted on only one server at a time. 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 Novell Cluster Services

Novell Cluster Services provides high availability from commodity components. You can configure up to 32 OES 2 Linux servers in a high-availability cluster, where resources can be dynamically relocated to any server in the cluster. Resources can be configured to automatically fail over to one or multiple different preferred servers in the event of a server failure. In addition, costs are lowered through the consolidation of applications and operations onto a cluster.

Novell 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, Web sites, 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.

Novell Cluster Services clusters provide the following benefits over standalone 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

Novell 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 Novell iManager cluster configuration and monitoring GUI. iManager also lets you remotely manage your cluster.
- ◆ 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 e-mail to automatically notify administrators of cluster events and cluster state changes.

1.4 Clustering for High-Availability

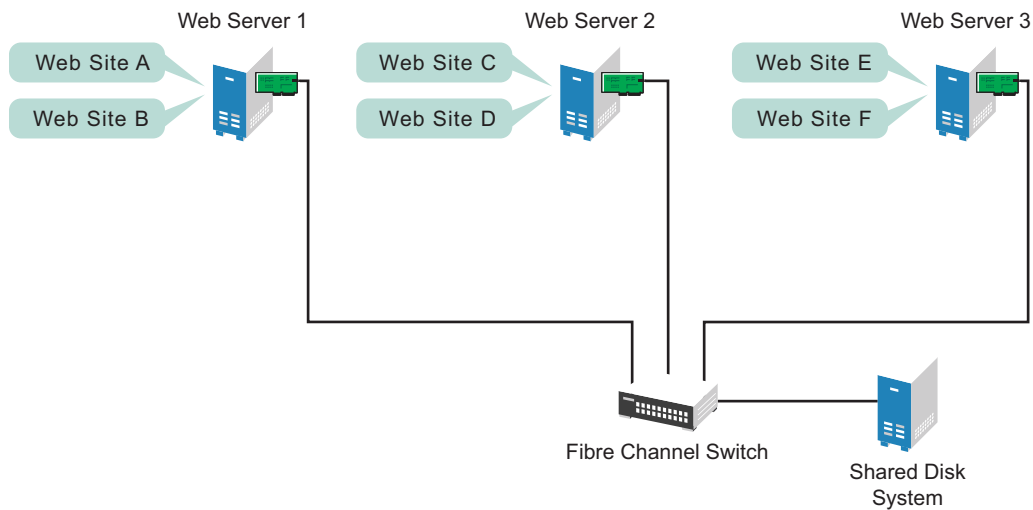
A Novell Cluster Services for Linux cluster consists of the following components:

- ◆ 2 to 32 OES 2 Linux servers, each containing at least one local disk device.
- ◆ Novell 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 Novell 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 Web sites. All the data, graphics, and Web page content for each Web site 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, e-mail, 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 Web sites 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, Novell 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 Web site 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 will automatically fail back, or be moved back to Web Server 1, and Web Server operation will return to the way it was before Web Server 1 failed.

Novell Cluster Services also provides resource migration capabilities. You can move applications, Web sites, etc. to other servers in your cluster without waiting for a server to fail.

For example, you could have manually moved 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 just to increase performance or accessibility of the Web sites.

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.

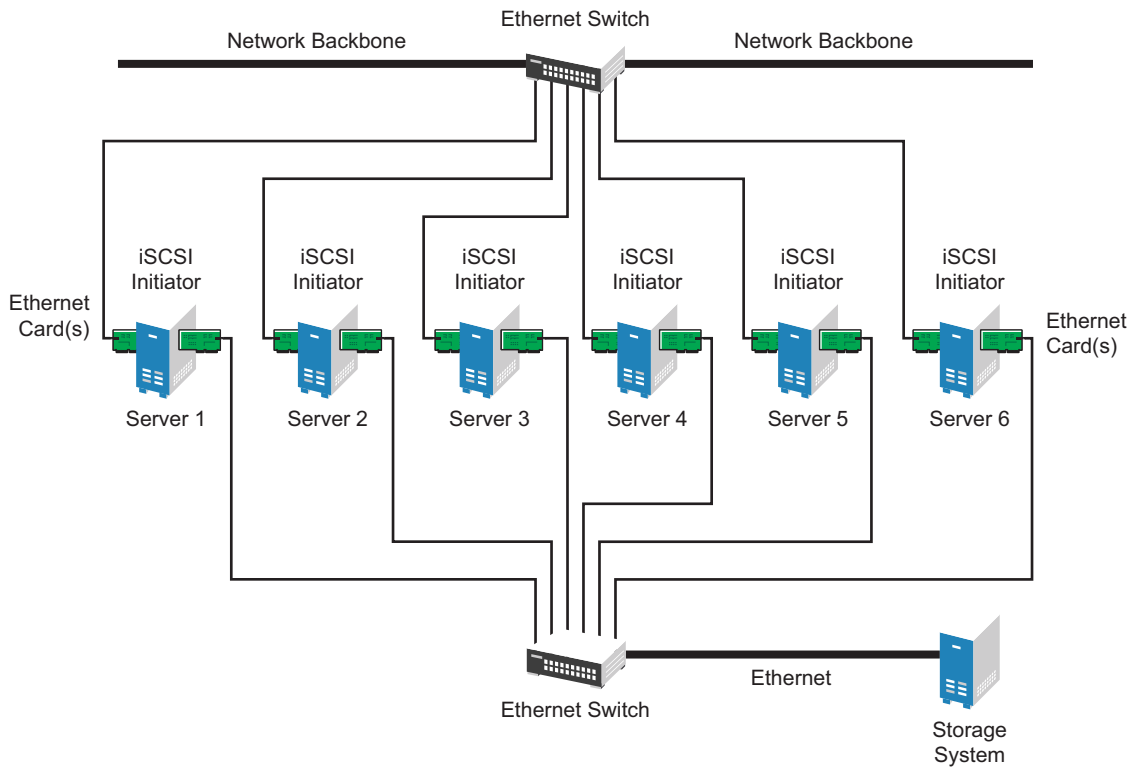
Figure 1-3 Typical Fibre Channel Cluster Configuration



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*



What's New

2

This section describes changes and enhancements that were made to Novell Cluster Services for Linux since the initial release of Novell Open Enterprise Server (OES) 2 Linux.

- ♦ [Section 2.1, “What’s New \(July 2010 Patches\),” on page 23](#)
- ♦ [Section 2.2, “What’s New \(June 2010 Patches\),” on page 23](#)
- ♦ [Section 2.3, “What’s New \(January 2010 Patches\),” on page 23](#)
- ♦ [Section 2.4, “What’s New \(OES 2 SP2\),” on page 24](#)
- ♦ [Section 2.5, “What’s New \(OES 2 SP1\),” on page 25](#)
- ♦ [Section 2.6, “What’s New \(OES 2\),” on page 27](#)

2.1 What’s New (July 2010 Patches)

The July 2010 Maintenance patches for OES 2 SP1 and OES 2 SP2 contain bug fixes for Novell Cluster Services. See the patch information for specific bug numbers.

2.2 What’s New (June 2010 Patches)

In addition to bug fixes, the June 2010 Novell Cluster Services Maintenance patches for OES 2 SP1 Linux and OES 2 SP2 Linux provide the following changes for Novell Cluster Services:

- ♦ **E-Mail:** The subject line of e-mail messages sent by Novell Cluster Services now provides information about the cluster name, resource name, action taken, and node name. For example:

```
CL1: POOL1_SERVER online on NODE1
```

Previously, the subject line was `Cluster Resource Change`.

- ♦ **Cluster Report:** The *Clusters > Cluster Manager > Run Report* output now includes the Monitoring script for each cluster resource.
- ♦ **Panic:** Novell Cluster Services has been modified to handle panic situations for a node according to the Linux kernel panic setting on the node. For information, see [Section 8.11, “Preventing a Cluster Node Reboot after a Node Shutdown,” on page 109](#).

2.3 What’s New (January 2010 Patches)

The January 2010 patch for OES 2 SP1 Linux contains bug fixes for Novell Cluster Services and Novell Business Continuity Clustering 1.2 for OES 2 SP1 Linux.

The January 2010 patch for OES 2 SP2 Linux contains bug fixes for Novell Cluster Services and adds support for Novell Business Continuity Clustering 1.2.1 for OES 2 SP2 Linux.

2.4 What's New (OES 2 SP2)

In addition to bug fixes, the following changes and enhancements were made for Novell Cluster Services for Linux in OES 2 SP2.

- ♦ [Section 2.4.1, “Improved Error Reporting,” on page 24](#)
- ♦ [Section 2.4.2, “Improved Time Calculations,” on page 24](#)
- ♦ [Section 2.4.3, “Specifying the Size of the SBD Partition,” on page 24](#)
- ♦ [Section 2.4.4, “Customizing Translation Syntax,” on page 24](#)
- ♦ [Section 2.4.5, “Migration Tool Support for Cluster Configurations,” on page 25](#)
- ♦ [Section 2.4.6, “New iFolder Resource Template,” on page 25](#)
- ♦ [Section 2.4.7, “Removed MPK Calls from the Code,” on page 25](#)
- ♦ [Section 2.4.8, “Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade,” on page 25](#)

2.4.1 Improved Error Reporting

This release provides improved error reporting for file protocol errors.

2.4.2 Improved Time Calculations

This release improves the way time is calculated so that the inter-packet gap between two heartbeat packets is reduced.

This means that you will observe an increase in the instances of packets incrementing the 0x (less than one second) counter, and a decrease in the instances of packets incrementing the 1x (between one second and two seconds) counter.

For example:

```
cluster stats display

Report taken: startTime= Thu Jul 23 13:16:33 2009,
              endTime= Mon Jul 27 08:44:36 2009
              node=5, name=Cluster_06, heartbeat=1, tolerance=8
              0x=645550, 1x=6, 2x=1, 3x=2, 5x=2
```

2.4.3 Specifying the Size of the SBD Partition

When you configure the SBD partition during the cluster configuration of a new cluster (as described in [“Configuring a New Cluster” on page 51](#)), you can now specify the size of the partition.

2.4.4 Customizing Translation Syntax

Beginning in OES 2 SP2, Novell Cluster Services allows you to customize the translation syntax that is used for load and unload scripts in mixed-platform situations by defining new syntax translations in the `/var/opt/novell/ncs/customized_translation_syntax` file that you create. The `clstrlib.py` script reads the additional translation syntax from the syntax file, and

processes them in addition to the normal translations in the Cluster Translation Library. For information, see [Section 6.7, “Customizing the Translation Syntax for Converting Load and Unload Scripts,”](#) on page 91.

2.4.5 Migration Tool Support for Cluster Configurations

Support was added for migrating services and data in cluster configurations by using the OES 2 SP2 Migration Tool. For instructions on using the Migration Tool to migrate services and data, see the *OES 2 SP2: Migration Tool Administration Guide*.

2.4.6 New iFolder Resource Template

The Novell iFolder 3.x resource template has been modified for OES 2 SP2. For information about using the new template, see [Section 4.3.1, “Updating the iFolder Resource Template,”](#) on page 58.

2.4.7 Removed MPK Calls from the Code

MPK calls were removed from the Novell Cluster Services code. The MPK calls were replaced with POSIX and Linux functions. These changes were made in support of the MPK calls being removed from the Novell Storage Services (NSS) file system software to achieve performance enhancements for the NSS file system.

2.4.8 Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade

In the OES 2 SP1 release, a cluster restart was required at the end of a rolling cluster upgrade in order to properly update the names of the nodes, as described in [Section 2.5.3, “Behavior Change for Adding a Node,”](#) on page 26. This issue was resolved in OES 2 SP2. The rolling cluster upgrade process no longer requires a cluster restart.

2.5 What’s New (OES 2 SP1)

In addition to bug fixes, the following changes and enhancements were made for Novell Cluster Services for Linux in OES 2 SP1.

- ◆ [Section 2.5.1, “Schema Extension,”](#) on page 26
- ◆ [Section 2.5.2, “Installation by Container Administrator,”](#) on page 26
- ◆ [Section 2.5.3, “Behavior Change for Adding a Node,”](#) on page 26
- ◆ [Section 2.5.4, “Attribute NCS: GIPC Config Is No Longer Maintained,”](#) on page 27
- ◆ [Section 2.5.5, “Support for Novell AFP for Linux,”](#) on page 27
- ◆ [Section 2.5.6, “Support for Novell CIFS for Linux,”](#) on page 27
- ◆ [Section 2.5.7, “Support for Domain Services for Windows,”](#) on page 27

2.5.1 Schema Extension

The administrator of a Novell eDirectory tree can now extend the schema for cluster objects before clusters are installed in the tree. This allows container administrators to install Novell Cluster Services without needing tree-level administrator rights. See [Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 42.

2.5.2 Installation by Container Administrator

Container administrators can install Novell Cluster Services without needing tree-level administrator rights. Make sure you have the rights needed for the install. See [Section 3.4, “Assigning Install Rights for Container Administrators,”](#) on page 44.

2.5.3 Behavior Change for Adding a Node

In this release, a behavior change was made to address a deadlock defect. After adding a new node to the cluster, the new node cannot be displayed in the Clusters plug-in to iManager until the `ncs-configd.py -init` script is run, or until the cluster is restarted.

IMPORTANT: A Novell Cluster Services patch is available in the patch channel and on the [Novell Downloads Web site \(http://www.novell.com/downloads\)](http://www.novell.com/downloads) that allows you to add a new node seamlessly again on a Linux server. For cluster conversions, a cluster restart is still necessary after all NetWare nodes have been removed from the cluster.

Run one of the following commands in order to make `cluster view` display the new node’s name correctly. It is okay to run `ncs-configd.py` on an active cluster.

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

or

```
rcnovell-ncs restart
```

If you are converting a cluster from NetWare to Linux, you must restart the cluster instead so that `clstrlib.ko` is reloaded.

For example, if you install a server named `sales_03` in an existing cluster named `oes2_sales_cluster` with two existing member nodes named `sales_01` and `sales_02`, the new node is generically displayed as `Node_03` when you enter the `cluster view` command:

```
Cluster OES2_SALES_CLUSTER
This node SALES_02 [ epoch 4 master node SALES_02 ]
Cluster nodes [ SALES_01, SALES_02, Node_03 ]
```

After running the cluster configuration daemon or restarting Novell Cluster Services, the new node’s name is properly displayed as `SALES_03`, and the node is visible in iManager.

```
Cluster OES2_SALES_CLUSTER
This node SALES_02 [ epoch 4 master node SALES_02 ]
Cluster nodes [ SALES_01, SALES_02, SALES_03 ]
```

2.5.4 Attribute NCS: GIPC Config Is No Longer Maintained

Beginning in OES 2 SP1 Linux, the attribute *NCS:GIPC Config* in the Cluster object is no longer maintained. This applies to Linux clusters and mixed NetWare and Linux clusters.

2.5.5 Support for Novell AFP for Linux

This release supports Novell AFP (Apple Filing Protocol) for Linux in combination with Novell Storage Services (NSS) volumes on OES 2 SP1 Linux. See [“Novell AFP for Linux” on page 34](#).

2.5.6 Support for Novell CIFS for Linux

This release supports Novell CIFS for Linux in combination with NSS volumes on OES 2 SP1 Linux. See [“Novell CIFS for Linux” on page 34](#).

2.5.7 Support for Domain Services for Windows

This release supports using clusters in Domain Services for Windows contexts for OES 2 SP1 Linux. See [“Novell Domain Services for Windows” on page 34](#).

2.6 What’s New (OES 2)

The following changes and enhancements were added to Novell Cluster Services for Linux in OES 2.

- ❑ **Resource Monitoring:** See [Section 9.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 129](#).
- ❑ **Support for Xen Virtualization:** See [Chapter 12, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 185](#).

Installing and Configuring Novell Cluster Services on OES 2 Linux

3

Novell Cluster Services can be installed during the Novell Open Enterprise Server (OES) 2 Linux installation or afterwards on an existing OES 2 Linux server.

For information about upgrading a cluster server from OES 1 SP2 Linux to OES 2 Linux, see [Chapter 5, “Upgrading OES 1 Linux Clusters to OES 2 Linux,”](#) on page 59.

For information about converting a cluster server from NetWare 6.5 SP7 or later to OES 2 Linux, see [Chapter 6, “Converting NetWare 6.5 Clusters to OES 2 Linux,”](#) on page 63.

- ♦ [Section 3.1, “System Requirements for Novell Cluster Services,”](#) on page 29
- ♦ [Section 3.2, “Novell Cluster Services Licensing,”](#) on page 42
- ♦ [Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 42
- ♦ [Section 3.4, “Assigning Install Rights for Container Administrators,”](#) on page 44
- ♦ [Section 3.5, “Installing Novell Cluster Services,”](#) on page 44
- ♦ [Section 3.6, “Configuring Novell Cluster Services,”](#) on page 48
- ♦ [Section 3.7, “Configuring Additional Administrators,”](#) on page 55
- ♦ [Section 3.8, “What’s Next,”](#) on page 56

3.1 System Requirements for Novell Cluster Services

- ♦ [Section 3.1.1, “Hardware Requirements,”](#) on page 29
- ♦ [Section 3.1.2, “Software Requirements,”](#) on page 30
- ♦ [Section 3.1.3, “Configuration Requirements,”](#) on page 36
- ♦ [Section 3.1.4, “Shared Disk System Requirements,”](#) on page 38
- ♦ [Section 3.1.5, “SAN Rules for LUN Masking,”](#) on page 40
- ♦ [Section 3.1.6, “Using Device Mapper Multipath with Novell Cluster Services,”](#) on page 40

3.1.1 Hardware Requirements

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

- A minimum of two Linux servers, and not more than 32 servers in a cluster
- At least 512 MB of 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, each server must meet the requirements for Novell Open Enterprise Server 2 Linux. For information, see [“Meeting All Server Software and Hardware Requirements”](#) in the *OES 2 SP2: 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.

3.1.2 Software Requirements

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

- ♦ [“Novell Open Enterprise Server 2 Linux”](#) on page 30
- ♦ [“Novell eDirectory 8.8”](#) on page 31
- ♦ [“Novell iManager 2.7.3”](#) on page 31
- ♦ [“EVMS”](#) on page 32
- ♦ [“Linux POSIX File Systems”](#) on page 32
- ♦ [“NSS File System on Linux”](#) on page 33
- ♦ [“Dynamic Storage Technology Shadow Volume Pairs”](#) on page 33
- ♦ [“NCP Server for Linux”](#) on page 33
- ♦ [“Novell AFP for Linux”](#) on page 34
- ♦ [“Novell CIFS for Linux”](#) on page 34
- ♦ [“Novell Domain Services for Windows”](#) on page 34
- ♦ [“OpenWBEM”](#) on page 34
- ♦ [“SLP”](#) on page 35
- ♦ [“Xen Virtualization Environments”](#) on page 35

Novell Open Enterprise Server 2 Linux

Novell Cluster Services 1.8.5 (or later) for Linux supports Novell Open Enterprise Server 2 Linux or later. OES 2 Linux must be installed and running on each cluster server in the cluster. Novell Cluster Services is a component of the OES 2 services for OES 2 Linux.

You cannot mix two versions of OES 2 Linux in a single cluster except to support a rolling cluster upgrade. For OES 2 Linux cluster upgrade information, see [Chapter 4, “Upgrading OES 2 Linux Clusters,”](#) on page 57.

You cannot mix OES 2 Linux and OES 1 Linux in a single cluster except to support a rolling cluster upgrade. For OES 1 Linux to OES 2 Linux upgrades, see [Chapter 5, “Upgrading OES 1 Linux Clusters to OES 2 Linux,”](#) on page 59.

You cannot mix OES 2 Linux and NetWare 6.5 SP7 (or later) in a single cluster except to support a rolling cluster conversion. For cluster conversion information, see [Chapter 6, “Converting NetWare 6.5 Clusters to OES 2 Linux,”](#) on page 63.

Novell eDirectory 8.8

Novell eDirectory 8.8 or later is required for managing the Cluster object and Cluster Node objects for Novell Cluster Services. eDirectory must be installed somewhere in the same tree as 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 Novell Cluster Services.

For eDirectory configuration requirements, see [“eDirectory Configuration”](#) on page 36.

Novell iManager 2.7.3

Novell iManager 2.7.3 or later is required for configuring and managing clusters for OES 2 SP2. iManager must be installed on at least one server in your tree.

NOTE: A February 3, 2009 update to the Clusters plug-in for OES 2 SP1 Linux is available that can be used to manage the Novell Business Continuity Clustering (BCC) 1.2 for OES 2 SP1 Linux. You can download the update from the [Novell Downloads Web site \(http://download.novell.com/\)](http://download.novell.com/). For information about BCC 1.2, see [BCC 1.2: Administration Guide for OES 2 SP1 Linux \(http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html\)](http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html)

The OES 2 SP2 Linux release contains a Clusters plug-in that is required when using Novell Business Continuity Clustering 1.2.1 for OES 2 SP2 Linux. For information about BCC 1.2.1, see [BCC 1.2.1: Administration Guide for OES 2 SP2 Linux \(http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html\)](http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html).

To use the Clusters role in iManager, the following storage-related plug-ins must be installed:

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

The `storagemgmt.npm` contains code in common with other storage-related plug-ins for OES 2 SP1:

Product	Plug-In	NPM File
Novell Archive and Version Services	Archive Versioning	<code>arkmgmt.npm</code>
Novell Apple Filing Protocol (AFP) for OES 2 SP1 Linux and NetWare	File Protocols > AFP	<code>afpmgmt.npm</code>
Novell CIFS for OES 2 SP1 Linux and NetWare	File Protocols > CIFS	<code>cifsmgmt.npm</code>
Novell Distributed File Services	Distributed File Services	<code>dfsmgmt.npm</code>
Novell Storage Services	Storage	<code>nssmgmt.npm</code>

These additional plug-ins are needed when working with the NSS file system. Make sure that you include the common `storagemgmt.npm` 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 make sure the common code works for all plug-ins.

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

These iManager plug-ins support iManager 2.7.2 and later on all operating systems supported by iManager and iManager Workstation.

The latest Novell storage-related plug-ins can be downloaded as a single zipped download file from the [Novell Downloads Web site \(http://download.novell.com\)](http://download.novell.com). For information about installing plug-ins in iManager, see “[Downloading and Installing Plug-in Modules](#)” in the *Novell iManager 2.7.3 Administration Guide*.

To update storage-related plug-ins:

- 1 In iManager, uninstall the currently installed storage-related plug-ins.
- 2 Copy the new .npm files into 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.

For information about working with storage-related plug-ins for iManager, see “[Understanding Storage-Related Plug-Ins](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

For browser configuration requirements, see “[Web Browser](#)” on page 38.

EVMS

EVMS (Enterprise Volume Management System) 2.5.5-24.54.5 or later is automatically installed on the server when you install Novell Cluster Services. It provides the Cluster Segment Manager (CSM) for shared cluster resources.

Updates to EVMS are received through the update channel for SUSE Linux Enterprise Server 10 SP 2 or later. Make sure that you install the latest patches for EVMS before you create any cluster resources for this server.

WARNING: EVMS administration utilities (`evms`, `evmsgui`, and `evmsn`) should not be running when they are not being used. EVMS utilities lock the EVMS engine, which prevents other EVMS-related actions from being performed. This affects both NSS and Linux POSIX volume actions.

NSS and Linux POSIX volume cluster resources should not be migrated while any of the EVMS administration utilities are running.

Linux POSIX File Systems

Novell Cluster Services supports creating shared cluster resources on Linux POSIX file systems, such as Ext3, XFS, and ReiserFS. Linux POSIX file systems are automatically installed as part of the OES 2 Linux installation.

NSS File System on Linux

Novell Cluster Services supports creating shared cluster resources by using Novell Storage Services file system on Linux. NSS is not a required component for Novell Cluster Services on Linux unless you want to create and cluster-enable NSS pools and volumes. You can also use NSS for the SBD partition. The *Novell Storage Services* option is not automatically selected when you select *Novell Cluster Services* option during the install.

NSS on Linux supports the following advertising protocols in concurrent use for a given cluster-enabled NSS pool:

- ♦ NetWare Core Protocol (NCP), which is selected by default and is mandatory for NSS. For information, see [“NCP Server for Linux” on page 33](#).
- ♦ Novell Apple Filing Protocol (AFP), which is available for NSS only in OES 2 SP1 Linux and later. For information, see [“Novell AFP for Linux” on page 34](#).
- ♦ Novell CIFS, which is available for NSS only in OES 2 SP1 Linux and later. For information, see [“Novell CIFS for Linux” on page 34](#).

Dynamic Storage Technology Shadow Volume Pairs

Novell Cluster Services supports clustering for Novell Dynamic Storage Technology (DST) shadow volume pairs on OES 2 Linux. DST is installed automatically when you install NCP Server for Linux. The *NCP Server and Dynamic Storage Technology* option is not automatically selected when you select *Novell Cluster Services* option during the install.

For information about creating and cluster-enabling Dynamic Storage Technology volumes on Linux, see [“Configuring DST Shadow Volumes with Novell Cluster Services for Linux”](#) in the *OES 2 SP2: Dynamic Storage Technology Administration Guide*.

NCP Server for Linux

NCP Server for Linux is required to be installed and running before you can cluster-enable NCP volumes on Linux POSIX file systems, NSS volumes, or Dynamic Storage Technology shadow volume pairs. This applies to physical servers and Xen-based virtual machine (VM) guest servers (DomU). NCP Server is required in order to create virtual cluster server objects for cluster resources for these volume types. The *NCP Server and Dynamic Storage Technology* option is not automatically selected when you select *Novell Cluster Services* option during the install.

NCP Server for Linux is required to be installed and running for NSS volumes, even if users access the volume only with other protocols such as Novell AFP for Linux, Novell CIFS for Linux, or Samba.

NCP Server for Linux is not required when you are cluster-enabling only Linux POSIX volumes where you plan to use only native Linux protocols for user access, such as Samba.

NCP Server for Linux is not required when running Novell Cluster Services on a Xen-based VM host server (Dom0) for the purpose of cluster-enabling the configuration files for Xen-based VMs. Users do not access these VM files.

For information about installing and configuring NCP Server for Linux, see the [OES 2 SP2: 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 Novell Cluster Services](#)” in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

Novell AFP for Linux

Novell Cluster Services supports using Novell AFP for Linux as an advertising protocol for cluster-enabled NSS pools and volumes on OES 2 SP1 Linux and later. Novell AFP is not required to be installed when you install Novell Cluster Services, but it must be installed and running when you create and cluster-enable the shared NSS pool in order for the *AFP* option to be available as an advertising protocol for the cluster resource.

For information about installing and configuring the Novell AFP service, see the *OES 2 SP2: Novell AFP For Linux Administration Guide*.

Novell CIFS for Linux

Novell Cluster Services supports using Novell CIFS for Linux as an advertising protocol for cluster-enabled NSS pools and volumes on OES 2 SP1 Linux and later. Novell CIFS is not required to be installed when you install Novell Cluster Services, but it must be installed and running when you cluster-enable the shared NSS pool in order for the *CIFS Virtual Server Name* and *CIFS* option to be available as an advertising protocol for the cluster resource.

For information about installing and configuring the Novell CIFS service, see the *OES 2 SP2: Novell CIFS for Linux Administration Guide*.

Novell Domain Services for Windows

Novell Cluster Services supports using clusters in Domain Services for Windows (DSfW) contexts for OES 2 SP1 Linux and later. 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 2 SP2: Domain Services for Windows Administration Guide*.

OpenWBEM

OpenWBEM must be configured to start with `chkconfig`, and be running when you manage the cluster with Novell iManager. For information on setup and configuration, see the *OES 2 SP2: OpenWBEM Services Administration Guide*.

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.

For OES 2 and later, 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 username and password) between iManager and the `_admin` volume on the OES 2 Linux that server you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

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

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

To check the status of WBEM:

- 1 As root in a console shell, enter
`rcowcimomd status`

To start WBEM:

- 1 As root in a console shell, enter
`rcowcimomd start`

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 Novell eDirectory configuration during the OES 2 Linux install. You can enable and configure SLP on the eDirectory Configuration - NTP & SLP page. For information, see “[Specifying SLP Configuration Options](#)” in the *OES 2 SP2: 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.

For more information, see “[Implementing the Service Location Protocol](http://www.novell.com/documentation/edir88/edir88/data/ba51b4b.html)” (<http://www.novell.com/documentation/edir88/edir88/data/ba51b4b.html>) in the *Novell eDirectory 8.8 Administration Guide*.

Xen Virtualization Environments

Xen virtualization software is included with SUSE Linux Enterprise Server. Novell Cluster Services supports using Xen virtual machine (VM) guest servers as nodes in a cluster. You can install Novell Cluster Services on the guest server just as you would a physical server. All templates except the Xen and XenLive templates can be used on a VM guest server. For examples, see [Chapter 12, “Configuring Novell Cluster Services in a Xen Virtualization Environment,”](#) on page 185.

Novell Cluster Services is supported to run on a Xen host server where it can be used to cluster the virtual machine configuration files on Linux POSIX file systems. Only the Xen and XenLive templates are supported for use in the XEN host environment. For information about setting up Xen and XenLive cluster resources, see [Section 12.2, “Virtual Machines as Cluster Resources,”](#) on page 186.

3.1.3 Configuration Requirements

Ensure that configuration requirements are met for these components:

- ♦ “IP Addresses” on page 36
- ♦ “eDirectory Configuration” on page 36
- ♦ “Cluster Services Installation Administrator” on page 36
- ♦ “Cluster Services Management Administrator” on page 37
- ♦ “Web Browser” on page 38

IP Addresses

- 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.
- 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.

eDirectory Configuration

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

If the servers in the cluster are in separate eDirectory containers, the user who 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” (<http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html>) in the *eDirectory 8.8 Administration Guide* for more information.

- 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 Novell Cluster Services install configuration does not create a new context.
- Multiple clusters can co-exist in the same eDirectory container.

Cluster Services Installation 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 (or non-administrator user) to install a cluster in a container in that same tree without needing full administrator rights for the tree. For instructions, see [Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 42.

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.

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

NOTE: If the eDirectory administrator username 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.

Cluster Services Management Administrator

The administrator credentials entered during the install are automatically configured to allow that user to manage the cluster. You can modify this default administrator username or password after the install by following the procedure in [Section 8.13, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110.

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.
- ◆ Make the user a member of a LUM-enabled administrator group that is associated with the servers in the cluster.

Web Browser

The browser that will be used to manage Novell Cluster Services must be set to a supported language. For information about supported languages, see the [Novell iManager documentation \(http://www.novell.com/documentation/imanager27/\)](http://www.novell.com/documentation/imanager27/).

The Cluster 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. To avoid problems, in your Web browser, click *Tools > Options > Languages*, and then set the first language preference in the list to a supported language.

Supported language codes are Unicode (UTF-8) compliant. To avoid display problems, make sure 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, click *View > Character Encoding*, then select the supported character encoding setting. In an Internet Explorer browser, click *View > Encoding*, then select the supported character encoding setting.

For a information about supported browsers, see “Using a Supported Web Browser” in the *Novell iManager 2.7.3 Administration Guide*.

3.1.4 Shared Disk System Requirements

A shared disk subsystem is required for each cluster in order to make data highly available. Make sure your shared storage devices meet the following requirements:

- ◆ “Shared Devices” on page 38
- ◆ “SBD Partitions” on page 39
- ◆ “Shared iSCSI Devices” on page 39
- ◆ “Shared RAID Devices” on page 39

Shared Devices

Novell 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 installing Novell 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 2 documentation or the shared disk system documentation for troubleshooting information.

Devices where you plan to create shared file systems (such as NSS pools and Linux POSIX file systems) must be unpartitioned devices that can be managed by EVMS. NSS automatically partitions the device and lays down a Cluster Segment Manager on the partition when you use NSS tools to create a pool. You use the EVMS GUI to partition and create shared Linux POSIX file systems.

If this is a new cluster, the shared disk system must be connected to the first server so that the cluster partition can be created during the Novell Cluster Services install.

SBD Partitions

If your cluster shares storage resources, you need to create a shared device that can be used to store a special cluster partition for the split-brain detector (SBD). It must have at least 20 MB of free disk space available for creating a special cluster partition for the SBD.

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

The Novell Cluster Services installation automatically allocates one cylinder on one drive of the shared disk system for the special cluster partition. Depending on the location of the cylinder, the actual amount of space used by the cluster partition might be less than 20 MB.

When you mark a device as *Shareable for Clustering*, the share information is added to the disk and consumes about 4 MB. This space is used in addition to the space needed for the SBD. For example, if you create a 1 GB (1024 MB) device for the SBD, 4 MB are used for the share information, leaving 1020 MB as available free space for the SBD partition.

If you want to mirror the SBD partition for greater fault tolerance, you need at least 20 MB of free disk space on a second shared disk where you want to create the mirror.

Before installing Novell Cluster Services or creating a new SBD partition, 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 you want to use as the mirrored SBD, and mark its device as *Shareable for Clustering*. This allows the installation software to recognize available partitions and present them for use during the install. You can initialize the device by using NSSMU or the Storage plug-in to iManager. Beginning in OES 2 SP2, the NSS utility called `ncsinit` is also available for initializing a device and setting it to a shared state.

For information about how SBD partitions work and how to create one after installing the first node, see [Section 8.16, “Creating or Deleting Cluster SBD Partitions,” on page 115](#).

Shared iSCSI Devices

If you are using iSCSI for shared disk system access, ensure that you have configured iSCSI initiators and targets (LUNs) prior to installing Novell Cluster Services.

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.
- ◆ 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 it does not support concurrent activation.

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

3.1.5 SAN Rules for LUN Masking

When you create a Novell 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. Novell 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.

3.1.6 Using Device Mapper Multipath with Novell Cluster Services

When you use Device Mapper Multipath (DM-MP) with Novell Cluster Services, make sure to 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. Novell 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 Novell Cluster Services can auto failover storage resources if a disk paths go down.

Change the setting in the `modprobe.conf.local` and `multipath.conf` files so that Novell Cluster Services works correctly with DM-MP. Also consider changes as needed for the retry settings in the HBA BIOS.

- ♦ [“Settings for the modprobe.conf.local File” on page 40](#)
- ♦ [“Settings for the multipath.conf File” on page 41](#)
- ♦ [“Settings for a QLogic HBA BIOS” on page 41](#)

Settings for the `modprobe.conf.local` File

Use the following setting in the `/etc/modprobe.conf.local` file:

```
options qla2xxx qlport_down_retry=1
```

There was a change in the latest kernel and qla-driver that influences the time-out. Without the latest patch, an extra five seconds is automatically added to the `port_down_retry` variable to determine the time-out value for option 1 (`dev_loss_tmo=port_down_retry_count+5`), and option 1 is the best choice. In the patch, the extra 5 seconds are no longer added to the `port_down_retry` variable to determine the time-out value for option 1 (`dev_loss_tmo=port_down_retry_count`). If you have installed the latest qla-driver, option 2 is the best choice.

For OES 2 SP2 and later, or if you have installed the latest kernel and qla-driver, use the following setting in the `/etc/modprobe.conf.local` file:

```
options qla2xxx qlport_down_retry=2
```

Settings for the `multipath.conf` File

Use the following settings for the device settings in the `/etc/multipath.conf` file:

```
failback "immediate"  
no_path_retry fail
```

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

For example:

```
defaults  
{  
    polling_interval    1  
#    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 "immediate"  
        no_path_retry fail    #Set MP for failed I/O mode, any other non-zero  
values sets the HBAs for Blocked I/O mode  
    }  
}
```

Settings for a QLogic HBA BIOS

For the QLogic 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. Change these settings to 5 seconds in the HBA BIOS. For example:

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

3.2 Novell Cluster Services Licensing

Novell Cluster Services supports up to 32 nodes in a single cluster. Novell Open Enterprise Server 2 customers receive a Novell Cluster Services entitlement that covers an unlimited number of two-node clusters. Customers who want to add nodes to any cluster can purchase a paper license for them for an additional fee. For information, see the [Novell Cluster Services for Open Enterprise Server How-to-Buy Web site \(http://www.novell.com/products/openenterpriseserver/ncs/howtobuy.html\)](http://www.novell.com/products/openenterpriseserver/ncs/howtobuy.html).

3.3 Extending the eDirectory Schema to Add Cluster Objects

When you install Novell Cluster Services in a tree, the eDirectory schema for the tree is extended to include the following types of objects:

- ◆ Cluster objects (containers)
 - ◆ Cluster Node objects
 - ◆ Cluster Resource objects
 - ◆ Cluster Template objects
 - ◆ Volume Resource objects

In OES 2 SP1 Linux and later, a tree administrator user with 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 Novell 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 Novell Cluster Services. For install rights information, see [“Cluster Services Installation Administrator” on page 36](#) and [Section 3.4, “Assigning Install Rights for Container Administrators,” on page 44](#).

- ◆ [Section 3.3.1, “Prerequisites for Extending the Schema,” on page 42](#)
- ◆ [Section 3.3.2, “Extending the Schema,” on page 43](#)

3.3.1 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 Administrator credentials for extending the eDirectory schema.

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

Table 3-1 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 username 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
server_ip_address	The IP address of the eDirectory server that contains the schema files.	10.10.10.1

3.3.2 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 Novell 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 separately from the first cluster installation in the tree, the tree administrator user modifies the schema files as follows:

- 1 On an OES 2 SP1 Linux (or later) server, open a terminal console, then log in as the `root` user to the tree.
- 2 In a text editor, create a text file, specify the configuration information for the Novell 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: Make sure to 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 terminal console prompt:

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

- 4 As the `root` user, enter the following command at a terminal console prompt:

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

Replace *configuration_filename* with the actual name of the file you created.

- 5 Delete the configuration file (*configuration_filename*) that you created.
- 6 Continue with [Section 3.4, “Assigning Install Rights for Container Administrators,”](#) on [page 44](#).

3.4 Assigning Install Rights for Container Administrators

After 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 rights to install Novell 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 information about assigning eDirectory rights, see “eDirectory Rights” (<http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html>) in the *eDirectory 8.8 Administration Guide*.

After assigning rights for these users, they can specify their credentials when they install Novell Cluster Services clusters in their containers. You can also specify a different administrator user with these rights to configure Novell Cluster Services as described in [Section 3.6.2, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 50.

3.5 Installing Novell Cluster Services

Novell Cluster Services for Linux is included on the add-on media for OES 2 Linux (or later). It is necessary to install OES 2 Linux on every server that you want to add to a cluster.

You can install up to 32 nodes in each cluster. For information, see [Section 3.2, “Novell Cluster Services Licensing,”](#) on page 42.

Installing Novell Cluster Services does the following:

- ♦ If the eDirectory schema has not already been extended for cluster objects, the schema is extended.
- ♦ Installs Novell Cluster Services software on the server.

You can create a new cluster or add a server to an existing cluster during the OES 2 Linux installation, or afterwards by using the *Open Enterprise Server > OES Install and Configuration* tool in YaST. For information, see [Section 3.6, “Configuring Novell Cluster Services,”](#) on page 48.

You can install Novell Cluster Services when you install OES 2 Linux, or afterwards on an existing OES 2 Linux server.

IMPORTANT: Before you begin, make sure your system meets the requirements and caveats in [Section 3.1, “System Requirements for Novell Cluster Services,”](#) on page 29.

If the eDirectory schema has not been extended in the tree where you want to create the cluster as explained in [Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 42, the administrator user who installs Novell Cluster Services must have the rights necessary to extend the schema.

- ♦ [Section 3.5.1, “Installing Novell Cluster Services during a OES 2 Linux Installation,”](#) on page 45
- ♦ [Section 3.5.2, “Installing Novell Cluster Services on an Existing OES 2 Linux Server,”](#) on page 47

3.5.1 Installing Novell Cluster Services during a OES 2 Linux Installation

This section describes only those steps in the install that are directly related to installing Novell Cluster Services. For detailed instructions on installing OES 2 Linux, see the [OES 2 SP2: Installation Guide](#).

IMPORTANT: If you want Novell Cluster Services to use a local eDirectory database on this server, we recommend that you install and configure eDirectory before installing Novell Cluster Services.

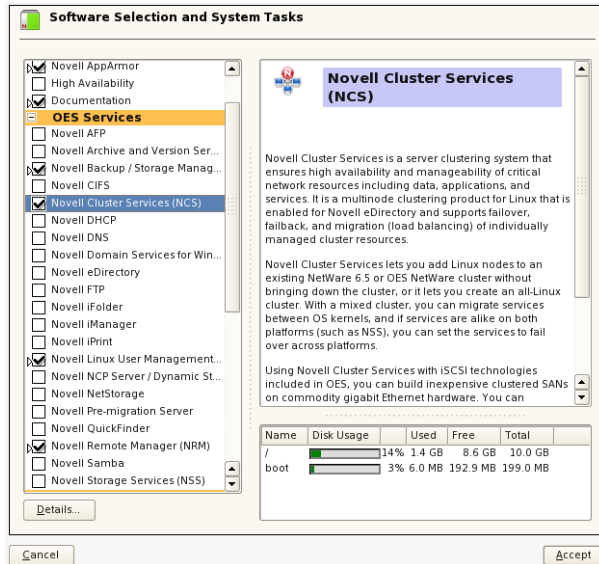
Repeat the following procedure for each server that you want to add to the cluster:

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

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

For information, see:

- ♦ [Section 8.16.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 116](#)
 - ♦ [Section 8.16.4, "Creating a Mirrored Cluster SBD Partition," on page 117](#)
- 2 Start the YaST install for SUSE Linux Enterprise Server 10 and continue to the Installation Mode page.
 - 3 Select *New Installation*, select *Include Add-On Products from Separate Media*, click *Next*, then continue through the OES 2 add-on part of the install until you get to the Installation Settings page.
 - 4 On the Installation Settings page, click *Software* to open the Software Selection and System Tasks page.
 - 5 Under *OES Services*, select *Novell Cluster Services* and any other OES components that you want to install, then click *Accept*.



When you select Novell Cluster Services, the following basic services for managing OES 2 are automatically selected:

- ◆ Novell Backup / Storage Management
- ◆ Novell Linux User Management
- ◆ Novell Remote Manager

The following OES services are not automatically selected, but are required for managing and configuring Novell Cluster Services:

- ◆ Novell iManager must be installed on at least one server in your data center.
- ◆ Novell 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 *Novell Storage Services* and *NCP Server and Dynamic Storage Technology* options if you want to cluster-enable NSS pools and volumes. This combination is also required to use cluster-enabled NSS volumes in DST shadow volume pairs.

Select *NCP Server and Dynamic Storage Technology* if you want to cluster-enable NCP volumes on Linux POSIX file systems.

Select other protocols and services according to your planned setup. For information, see [Section 3.1.2, “Software Requirements,” on page 30](#).

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.

- 6 Continue through the installation process until you reach the Novell Open Enterprise Server Configuration page, then do one of the following:
 - ♦ **Configure Novell Cluster Services now:** If eDirectory is already installed and running in your environment and you want to configure clustering for the server now, continue with [Section 3.6.3, “Enabling the Novell Cluster Services Configuration Page in YaST,” on page 51.](#)
 - ♦ **Configure Novell Cluster Services later:** Continue with the rest of the OES installation, but do not configure clustering. You can configure clustering later by using *YaST > Open Enterprise Server > OES Install and Configuration* to access the Novell Open Enterprise Server Configuration page. For information, see [Section 3.6, “Configuring Novell Cluster Services,” on page 48.](#)

3.5.2 Installing Novell Cluster Services on an Existing OES 2 Linux Server

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

IMPORTANT: If you want Novell Cluster Services to use a local eDirectory database on the existing server, we recommend that you install and configure eDirectory before installing Novell Cluster Services.

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

- 1 If you have a shared disk system and the server where you are installing Novell Cluster Services is the second node in a cluster, verify that a cluster partition for the cluster’s Split Brain Detector exists on the first cluster node before you begin the install on the second node.

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.

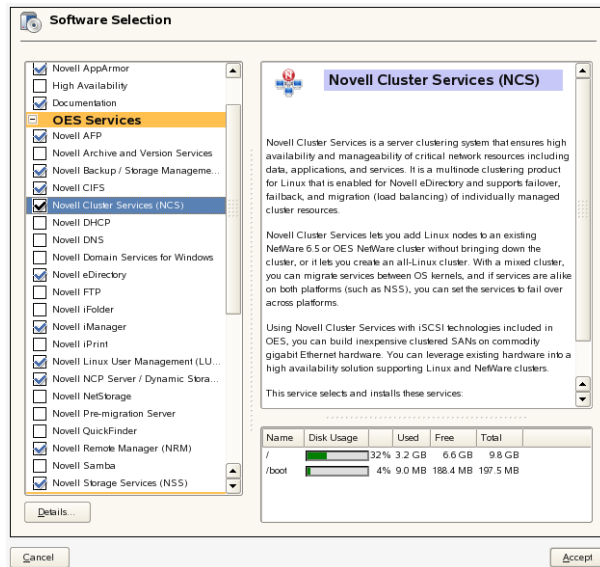
IMPORTANT: The cluster SBD partition is not required unless you have shared storage.

For information, see:

- ♦ [Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#)
 - ♦ [Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)
- 2 Log in to the server as the `root` user.
 - 3 In YaST, select *Open Enterprise Server > OES Install and Configuration*.



- 4 On the Software Selection page under *OES Services*, click *Novell 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 5 in Section 3.5.1, “Installing Novell Cluster Services during a OES 2 Linux Installation,”](#) on page 45.

- 5 Click *Accept* to begin the install, then click *Continue* to accept changed packages.
- 6 Continue through the installation process until you reach the Novell Open Enterprise Server Configuration page.
- 7 On the Novell Open Enterprise Server Configuration page, do one of the following:
- ◆ **Configure Novell Cluster Services now:** If eDirectory is already installed and running in your environment and you want to configure clustering for the server now, continue with [Section 3.6.3, “Enabling the Novell Cluster Services Configuration Page in YaST,”](#) on page 51.
 - ◆ **Configure Novell Cluster Services later:** Continue with the rest of the Novell Cluster Services installation, but do not configure clustering. You can configure clustering later by using *YaST > Open Enterprise Server > OES Install and Configuration* to access the Novell Open Enterprise Server Configuration page. For information, see [Section 3.6, “Configuring Novell Cluster Services,”](#) on page 48.

3.6 Configuring Novell Cluster Services

You can create a new cluster or add a server to an existing cluster during the OES 2 Linux installation, or afterwards by using the *Open Enterprise Server > OES Install and Configuration* tool in YaST.

IMPORTANT: If Novell Cluster Services is already configured on the server, see [Section 8.13, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110 for information about modifying an existing cluster configuration.

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

- ◆ Creates a new Cluster Node object in eDirectory.
- ◆ [Section 3.6.1, “Opening the Novell Open Enterprise Server Configuration Page,” on page 49](#)
- ◆ [Section 3.6.2, “Using Different LDAP Credentials for the Cluster Configuration,” on page 50](#)
- ◆ [Section 3.6.3, “Enabling the Novell Cluster Services Configuration Page in YaST,” on page 51](#)
- ◆ [Section 3.6.4, “Configuring a New Cluster,” on page 51](#)
- ◆ [Section 3.6.5, “Adding a Node to an Existing Cluster,” on page 54](#)

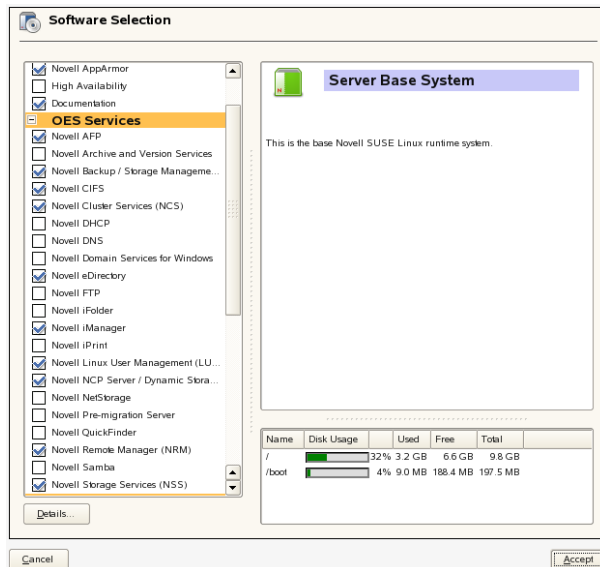
3.6.1 Opening the Novell Open Enterprise Server Configuration Page

- 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 *OES Services*, verify that the *Novell Cluster Services* option is already installed as indicated by a blue check mark, then click *Accept*.

It is okay to install other OES components at this time, but this setup focuses only on configuring Novell Cluster Services.



- 4 Click *Accept* to proceed to the Novell Open Enterprise Server Configuration page.

- 5 Do one of the following:

- ♦ **Same Administrator:** To use the same administrator credentials that were used to install Novell Cluster Services, continue with [Section 3.6.3, “Enabling the Novell Cluster Services Configuration Page in YaST,”](#) on page 51.
- ♦ **Different Administrator:** To use different administrator credentials than those used to install Novell Cluster Services, continue with [Section 3.6.2, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 50.

3.6.2 Using Different LDAP Credentials for the Cluster Configuration

You can use different user credentials to configure Novell Cluster Services than were used during the installation of OES Services 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 3.4, “Assigning Install Rights for Container Administrators,”](#) on page 44.

- 1 On the Novell 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 username and context (in LDAP form) of the container administrator user (or non-administrator user) who has the eDirectory rights needed to install Novell Cluster Services.
 - ♦ **Admin password:** The password of the container administrator (or non-administrator user).
- 4 Click *Next*.
The install returns to the Novell Open Enterprise Server Configuration page.
- 5 Continue with [Section 3.6.3, “Enabling the Novell Cluster Services Configuration Page in YaST,”](#) on page 51.


3.6.3 Enabling the Novell Cluster Services Configuration Page in YaST

- 1 On the Novell Open Enterprise Server Configuration page under *Novell Cluster Services*, locate the Novell 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 *Novell Cluster Services* link to open the Novell Cluster Services Configuration page.
- 4 If you are prompted for credentials, specify the password of the specified Administrator user, then click *OK*.
If you did not specify a different administrator user in [Section 3.6.2, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 50, this is the Administrator user whose credentials you specified for eDirectory when the OES Services were installed on the server.
- 5 On the Novell Cluster Services Configuration page, continue with one of the following:
 - ♦ [Section 3.6.4, “Configuring a New Cluster,”](#) on page 51
 - ♦ [Section 3.6.5, “Adding a Node to an Existing Cluster,”](#) on page 54

3.6.4 Configuring a New Cluster

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

- 1 [Go to the Novell Cluster Services Configuration page](#), then select *New Cluster*.

2 Specify the following settings for the new cluster, then click *Next*.

Fully Distinguished name (FDN) of the Cluster

IMPORTANT: Use the comma format illustrated in the example. Do not use dots.

This is the name you will give the new cluster and the eDirectory context where the new Cluster object will reside.

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 is separate from the server IP address, is required to be on the same IP subnet as the other servers in the same cluster, and is required for certain external network management programs to get cluster status alerts. The cluster IP address provides a single point for cluster access, configuration, and management. A Master IP Address resource that makes this possible 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.

SBD (Split Brain Detector) Partition

If you have a shared disk system attached to your cluster servers, Novell Cluster Services creates a small cluster partition for the Split Brain Detector (SBD) on that shared disk system. You can create the SBD partition as part of the first cluster node setup, or any time before you attempt to install the second node in the cluster by using one of the following procedures:

- ◆ [Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#)
- ◆ [Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

IMPORTANT: For each SBD partition, you must have at least 20 MB of free space on a device that has been initialized and marked as shareable for clustering.

If your cluster	Do this
Has no shared devices	Accept the default setting of <i>None</i> . The SBD is not needed.
Has shared devices, but a device that you want to use for the SBD partition has not yet been initialized and marked as shareable for clustering, or you want to set up the SBD later.	Accept the default setting of <i>None</i> . To create an SBD partition after the install, see Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116 . To create a mirrored SBD partitions after the install, see Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117 .
Has shared devices, and you want to create a single SBD partition. The partition you want to use has been initialized and marked as shareable for clustering.	Select the device where the SBD partition will be created and specify the size to use. For size requirements, see “SBD Partitions” on page 39 . For example, the device might be something similar to <code>sdc</code> .

If your cluster	Do this
Has shared devices, and you want to create a mirrored SBD partition for greater fault tolerance. The partitions you want to use have each been initialized and marked as shareable for clustering.	Select the device where the SBD partition will be created, specify the size to use, and select the device where you want to create the mirrored partition. For size requirements, see “SBD Partitions” on page 39 .

- 3 Select the IP address that Novell Cluster Services will use for this node.
Some servers have multiple IP addresses. This step lets you choose which IP address Novell Cluster Services uses.
- 4 If this is an existing server, choose whether to start Novell Cluster Services software on this node after configuring it, then click *Next*.
This option applies only to installing Novell Cluster Services on an existing server, because it starts automatically when the server reboots during the OES installation.
- 5 If you are configuring during the install, continue through the rest of the OES installation and setup process.
- 6 After the configuration is completed, start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
New install	Novell Cluster Services starts automatically when the server reboots during the OES installation.
Automatically start specified in Step 4	Novell Cluster Services should be running after the configuration completes.
Do not automatically start specified in Step 4	Start Novell Cluster Services manually by using one of these methods: <ul style="list-style-type: none"> ◆ Reboot the cluster server. ◆ At a terminal console prompt, go to the <code>/etc/init.d</code> directory, then enter the following as the <code>root</code> user: <pre>./novell-ncs start</pre> ◆ At a terminal console prompt, enter the following as the <code>root</code> user: <pre>rcnovell-ncs start</pre>

- 7 Use the Software Updater (or other update methods) to install any patches from the OES 2 Linux patch channel and any EVMS patches from the SUSE Linux Enterprise Server 10 SP3 patch channel.
- 8 Continue with [Section 3.7, “Configuring Additional Administrators,” on page 55](#).

3.6.5 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 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#)
- ♦ [Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

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 least 20 MB of free space on a device that has been previously initialized and marked as shareable for clustering.

- 2 [Go to the Novell Cluster Services Configuration page](#), then click *Existing Cluster*.
- 3 In the *Cluster FDN* field, specify the fully distinguished name (FDN) (the name and eDirectory context) of the cluster where you want to add this server.

IMPORTANT: Use the comma format illustrated in the example. Do not use dots.

- 4 Click *Next*.
- 5 Select the IP address that Novell Cluster Services will use for this node.
Some servers have multiple IP addresses. This step lets you choose which IP address Novell Cluster Services uses.
- 6 If this is an existing server, choose whether to start Novell Cluster Services software on this node after configuring it, then click *Finish*.
This option applies only to installing Novell Cluster Services on an existing server, because it starts automatically when the server reboots during the OES installation.
- 7 If you are configuring during the install, continue through the rest of the OES installation and setup process.
- 8 After the configuration is completed, start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
New install	Novell Cluster Services starts automatically when the server reboots during the OES installation.
Automatically start specified in Step 4	Novell Cluster Services should be running after the configuration completes.

Setup Condition	Instructions
Do not automatically start specified in Step 4	<p>Start Novell Cluster Services manually by using one of these methods:</p> <ul style="list-style-type: none"> ◆ Reboot the cluster server. ◆ At a terminal console prompt, go to the <code>/etc/init.d</code> directory, then enter the following as the <code>root</code> user: <pre>./novell-ncs start</pre> ◆ At a terminal console prompt, enter the following as the <code>root</code> user: <pre>rcnovell-ncs start</pre>

9 Use the Software Updater (or other update methods) to install any patches from the OES 2 Linux patch channel and any EVMS patches from the SUSE Linux Enterprise Server 10 SP3 patch channel.

10 Update the cluster configuration on the other nodes by doing one of the following as the `root` user on the master node in the cluster:

- ◆ Run the cluster configuration daemon by entering the following:

```
cluster exec "/opt/novell/ncs/bin/ncs-configd.py -init"
```

- ◆ Restart Novell Cluster Services by entering:

```
rcnovell-ncs restart
```

This step is necessary to make `cluster view` display the new node's name correctly, and allows the node to be displayed in iManager. For information, see [Section 2.5, "What's New \(OES 2 SP1\),"](#) on page 25.

11 Continue with [Section 3.7, "Configuring Additional Administrators,"](#) on page 55.

3.7 Configuring Additional Administrators

The Administrator user who you specify during Novell 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' usernames must be similarly configured.

You can modify the default administrator username or password after the install by following the procedure in [Section 8.13, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,"](#) on page 110.

3.8 What's Next

After installing Novell Cluster Services on OES 2 Linux servers (physical servers or virtual guest servers), you can configure and manage the cluster and cluster resources. For information, see the following:

- ◆ [Chapter 7, “Configuring Cluster Policies and Priorities,” on page 95](#)
- ◆ [Chapter 8, “Managing Clusters,” on page 101.](#)
- ◆ [Chapter 9, “Configuring and Managing Cluster Resources,” on page 123](#)
- ◆ [Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139](#)
- ◆ [Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163](#)

If you install Novell Cluster Services at the host level of an OES 2 Linux (Xen) server, you can create cluster resources for the virtual machines. For information, see [Section 12.2, “Virtual Machines as Cluster Resources,” on page 186.](#)

Upgrading OES 2 Linux Clusters

4

You can upgrade a Novell Open Enterprise (OES) 2 Linux (or later) cluster to OES 2 SP2 Linux. This section describes the upgrade process and how to manage the temporarily mixed cluster during the upgrade.

- ◆ [Section 4.1, “Requirements for Upgrading Clusters,” on page 57](#)
- ◆ [Section 4.2, “Upgrading OES 2 Clusters \(Rolling Cluster Upgrade\),” on page 57](#)
- ◆ [Section 4.3, “Upgrade Issues for OES 2 SP2,” on page 58](#)

4.1 Requirements for Upgrading Clusters

Make sure your environment meets the requirements for installing Novell Cluster Services that are described in [Section 3.1, “System Requirements for Novell Cluster Services,” on page 29](#).

4.2 Upgrading OES 2 Clusters (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade on your OES 2 Linux cluster lets you keep your cluster up and running. Your users can continue to access cluster resources while the upgrade is being performed.

During a rolling cluster upgrade, one server is upgraded to OES 2 SP2 Linux while the other servers in the cluster continue running older versions of OES 2 Linux. Then, if desired, another server can be upgraded, and then another, until all servers in the cluster have been upgraded.

You can also add OES 2 SP2 Linux servers to the existing OES 2 Linux (or later) cluster, and remove the old servers from the cluster.

You should complete the upgrade as soon as possible. Don't leave the cluster in a mixed version state for an extended period.

Make sure that any services that are available only in OES 2 SP1 or later (such as Novell CIFS or Novell AFP) are set up with preferred nodes for failover that are running OES 2 SP1 or later. Consider the following issues when working with a mixed-release OES 2 cluster:

- ◆ OES 2 SP1 features that are not available for the OES 2 Linux:
 - ◆ Novell AFP for Linux
 - ◆ Novell CIFS for Linux
 - ◆ 64-bit eDirectory 8.8.4 support
 - ◆ Domain Services for Windows support
 - ◆ NSS /`(No)atime` option
 - ◆ NSS /`PoixPermissionMask` option
 - ◆ NSS /`UnplugAlways` option
- ◆ The NSS default namespace on OES 2 SP1 Linux was changed to Long. This is not a problem for existing pools, but you should be aware of the difference when creating pools on OES 2 nodes versus OES 2 SP1 (or later) nodes.

To perform a rolling cluster upgrade of OES 2 Linux and Novell Cluster Services:

- 1 Bring down the OES 2 Linux cluster server you want to upgrade.
Any cluster resources that were running on the server should fail over to another server in the cluster.
You can also manually migrate the resources to another server in the cluster prior to bringing down the server. This prevents the resources from failing back to the node after you have completed the upgrade.
- 2 Upgrade the server by using the *Update* option on the OES 2 Linux installation.
See “[Upgrading to OES 2 SP2](#)” in the *OES 2 SP2: Installation Guide*.
- 3 If necessary, manually cluster migrate the resources that were previously loaded on this server server back to the upgraded server.
The resources can automatically fail back if both of the following apply:
 - ♦ The failback mode for the resources was set to Auto.
 - ♦ This Linux server is the preferred node for the resources.
- 4 Repeat [Step 2](#) through [Step 5](#) for each OES 2 Linux cluster server until your entire cluster has been upgraded.

4.3 Upgrade Issues for OES 2 SP2

This section contains known issues for upgrading clusters from OES 2 Linux or OES 2 SP1 Linux to OES 2 SP2.

- ♦ [Section 4.3.1, “Updating the iFolder Resource Template,” on page 58](#)

4.3.1 Updating the iFolder Resource Template

In OES 2 SP2 Linux, the Novell Cluster Services resource template for Novell iFolder 3.x has changed. The revised template is `/opt/novell/ncs/templates/iFolder_template.xml`. You must run a script on one of the nodes (master node is preferred) in the cluster after upgrading to SP2 in order to update the new resource template information in Novell Cluster Services and eDirectory.

After upgrading the cluster nodes from OES 2 Linux or OES 2 SP1 Linux to OES 2 SP2 Linux:

- 1 On the master node in the cluster, open a terminal console as the `root` user, then enter

```
/opt/novell/ncs/bin/ncstempl.py -U
```


The template changes take effect without rebooting the node or restarting the cluster.

Upgrading OES 1 Linux Clusters to OES 2 Linux

5

This section provides information to help you upgrade a Novell Open Enterprise (OES) 1 Linux cluster to OES 2 Linux, and how to manage the temporarily mixed cluster during the upgrade.

For information about managing an OES 1 Linux cluster, see the *OES: Novell Cluster Services 1.8.2 Administration Guide for Linux* (http://www.novell.com/documentation/oes/cluster_admin_lx/data/h4hgu4hs.html).

- ♦ Section 5.1, “Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux,” on page 59
- ♦ Section 5.2, “Upgrading Existing OES 1 Linux Cluster Nodes to OES 2 (Rolling Cluster Upgrade),” on page 60
- ♦ Section 5.3, “Adding New OES 2 Linux Cluster Nodes to Your OES 1 Linux Cluster,” on page 61
- ♦ Section 5.4, “Modifying Cluster Resource Scripts for Mixed OES 1 Linux and OES 2 Linux Clusters,” on page 61
- ♦ Section 5.5, “Finalizing the Cluster Upgrade,” on page 61

5.1 Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux

In addition to the [Section 3.1, “System Requirements for Novell Cluster Services,”](#) on page 29, consider the following rules and recommendations for mixed OES 1 Linux and OES 2 Linux clusters:

- ♦ Mixed OES 1 Linux and OES 2 Linux clusters should be considered a temporary configuration that exists only during an upgrade.
- ♦ Adding an OES 1 Linux cluster node to a mixed-node OES 1 Linux and OES 2 Linux cluster is not supported
- ♦ If you have configured resource monitoring for a resource running on an OES 2 Linux cluster node, resource monitoring does not function if the resource fails over or is migrated to an OES 1 Linux cluster node.
- ♦ The use of EVMS is recommended for upgrading file system resources.
- ♦ You should ensure that all resource policies are configured to the settings that existed before the upgrade.
- ♦ No storage management functions should be executed while a cluster is in a mixed-cluster mode. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

5.2 Upgrading Existing OES 1 Linux Cluster Nodes to OES 2 (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade from OES 1 Linux to OES 2 Linux 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 server is upgraded to OES 2 Linux while the other servers in the cluster continue running OES 1 Linux. Then, if desired, another server can be upgraded, and then another, until all servers in the cluster have been upgraded to OES 2 Linux. You should complete the upgrade as soon as possible. Don't leave the cluster in a mixed version state for an extended period.

To perform a rolling cluster upgrade from OES 1 Linux to OES 2 Linux:

- 1 Make a note of the OES components that are installed on the server you are upgrading.

You will probably want to install the same components on the node as you perform the upgrade.

NOTE: NSS, eDirectory, and NCP Server are not required components for Novell Cluster Services on OES 2, but are required components for Novell Cluster Services on OES 1 Linux. Ensure that you install these components on OES 2 Linux servers when you upgrade your OES 1 Linux cluster servers.

- 2 Bring down the OES 1 Linux cluster server you want to upgrade to OES 2.

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

You can also manually migrate the resources to another server in the cluster prior to bringing down the server. This prevents the resources from failing back to the node after you have completed the upgrade.

- 3 Upgrade the server by using the *Update* option on the OES 2 Linux installation, making sure to install the components that are currently installed on the server that you noted in [Step 1](#).

See “[Upgrading to OES 2 SP2](#)” in the *OES 2 SP2: Installation Guide*.

- 4 Repeat [Step 1](#) through [Step 3](#) for each OES 1 Linux cluster server until your entire cluster has been upgraded to OES 2.

- 5 (Conditional) If necessary, manually migrate the resources that were on the former OES 1 server to this Linux server.

The resources will automatically fail back if both of the following apply:

- ♦ The failback mode for the resources was set to Auto.
- ♦ This Linux server is the preferred node for the resources.

- 6 After the last OES 1 cluster node has been upgraded to OES 2, finalize the upgrade by following the instructions in [Section 5.5, “Finalizing the Cluster Upgrade,”](#) on page 61.

See also [Section 5.1, “Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux,”](#) on page 59 for more information on mixed OES 1 and OES 2 clusters and the rules that apply to them.

5.3 Adding New OES 2 Linux Cluster Nodes to Your OES 1 Linux Cluster

The process for adding a new OES 2 Linux cluster node to an existing OES 1 Linux cluster is the same as for adding an OES 1 cluster node to an OES 1 Linux cluster and adding a new OES 2 cluster node to an existing OES 2 cluster. See [Section 3.5, “Installing Novell Cluster Services,”](#) on page 44.

However, you should be aware of the rules that apply to mixed OES 1 and OES 2 clusters. See [Section 5.1, “Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux,”](#) on page 59.

5.4 Modifying Cluster Resource Scripts for Mixed OES 1 Linux and OES 2 Linux Clusters

OES 1 Linux and OES 2 Linux cluster resource load and unload scripts perform similar actions, but some template scripts differ in the functions used to perform those actions. OES 2 cluster template scripts have been upgraded and some of them now conform to the Open Cluster Framework (OCF).

Cluster resources created on an OES 1 Linux cluster server can run in a mixed version cluster on either OES 1 or OES 2 Linux cluster servers.

Cluster resources created on an OES 2 Linux cluster server that is part of an OES 1 and OES 2 Linux mixed cluster can also run either OES 1 and OES 2 cluster servers.

After completing the cluster upgrade to OES 2 you must finalize the cluster upgrade. See [Finalizing the Cluster Upgrade](#) below for more information. Any new cluster resources that are created using resource templates included with OES 2 Linux use upgraded templates after finalizing the cluster upgrade. New resources created after finalizing the upgrade cannot run on OES 1 Linux cluster nodes, so you should finalize the cluster upgrade only after all nodes have been upgraded to OES 2 Linux.

After you have finalized the cluster upgrade you, might want to change the scripts for existing cluster resources that were created by using templates so that they use the scripts in the upgraded templates. This is especially true if you want to use the new resource monitoring feature. There are three ways to do this:

- ◆ Copy and paste the scripts from the upgraded resource templates to existing resource scripts and then customize the scripts for the resource.
- ◆ Offline an existing resource and then create a new resource with a different name to replace the existing resource, using the upgraded templates.
- ◆ Offline and then delete an existing resource, then create a new resource with the same name to replace the existing resource, using the upgraded templates.

5.5 Finalizing the Cluster Upgrade

If you have upgraded all nodes in an OES 1 Linux cluster to OES 2 Linux, you must finalize the upgrade process by issuing the `cluster convert commit` command on one Linux cluster node. The `cluster convert commit` command upgrades load and unload scripts for cluster templates and adds monitor scripts.

Converting NetWare 6.5 Clusters to OES 2 Linux

6

You can use a rolling cluster conversion to convert a Novell Cluster Services cluster from NetWare 6.5 SP8 to Novell Open Enterprise Server (OES) 2 SP1 Linux. This section describes how to prepare for and perform the conversion, and how to manage the temporarily mixed cluster during the conversion.

- ◆ Section 6.1, “Guidelines for Converting Clusters from NetWare to OES 2 Linux,” on page 63
- ◆ Section 6.2, “Guidelines for Converting NSS Pool Resources from NetWare to Linux,” on page 67
- ◆ Section 6.3, “Guidelines for Converting Service Cluster Resources from NetWare to Linux,” on page 68
- ◆ Section 6.4, “Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion),” on page 82
- ◆ Section 6.5, “Adding New OES 2 Linux Nodes to Your NetWare Cluster,” on page 85
- ◆ Section 6.6, “Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters,” on page 86
- ◆ Section 6.7, “Customizing the Translation Syntax for Converting Load and Unload Scripts,” on page 91
- ◆ Section 6.8, “Finalizing the Cluster Conversion,” on page 92

For information about managing a NetWare cluster, see the “Clustering NetWare Services” list on the NetWare 6.5 SP8 Clustering (High Availability) Documentation Web site (<http://www.novell.com/documentation/nw65/cluster-services.html#clust-config-resources>).

6.1 Guidelines for Converting Clusters from NetWare to OES 2 Linux

In addition to Section 3.1, “System Requirements for Novell Cluster Services,” on page 29, consider the requirements and guidelines described in the following sections when converting clusters from NetWare to OES 2 Linux:

- ◆ Section 6.1.1, “Supported Mixed-Node Clusters,” on page 64
- ◆ Section 6.1.2, “SBD Devices Must Be Marked as Shareable for Clustering,” on page 64
- ◆ Section 6.1.3, “Syntax Translation Issues for Load and Unload Scripts,” on page 64
- ◆ Section 6.1.4, “Case Sensitivity Issues,” on page 65
- ◆ Section 6.1.5, “Adding New NetWare Nodes to a Mixed-Node Cluster,” on page 65
- ◆ Section 6.1.6, “Converting Multiple NetWare Cluster Nodes to OES 2 Linux,” on page 65
- ◆ Section 6.1.7, “Converting Nodes that Contain the eDirectory Master Replica,” on page 66
- ◆ Section 6.1.8, “Failing Over Cluster Resources on Mixed-Node Clusters,” on page 66
- ◆ Section 6.1.9, “Managing File Systems in Mixed-Node Clusters,” on page 66

- ◆ [Section 6.1.10, “Using Novell iManager in Mixed-Node Clusters,” on page 66](#)
- ◆ [Section 6.1.11, “Using Novell Remote Manager Is Not Supported in Mixed-Node Clusters,” on page 67](#)
- ◆ [Section 6.1.12, “Using ConsoleOne Is Not Supported for Mixed-Node Clusters,” on page 67](#)
- ◆ [Section 6.1.13, “Using the Monitor Function in Mixed-Node Clusters Is Not Supported,” on page 67](#)

6.1.1 Supported Mixed-Node Clusters

Mixed NetWare and OES 2 Linux nodes in the same cluster are supported as a temporary configuration while you are migrating a cluster from NetWare to Linux.

All NetWare servers must be either version 6.5 or 6.0 in order to exist in a mixed NetWare and OES 2 Linux cluster.

- ◆ Mixed NetWare 6.5 and OES 2 Linux clusters are supported so that you can convert a NetWare 6.5 cluster to OES 2 Linux.
- ◆ Mixed NetWare 6.0 and OES 2 Linux clusters are also supported so that you can convert a NetWare 6.0 cluster to OES 2 Linux.
- ◆ Mixed clusters consisting of NetWare 6.0 servers, NetWare 6.5 servers, and OES 2 Linux servers are not supported.

Before converting NetWare 6.5 clusters or NetWare 6.0 clusters to OES 2 Linux, you must apply all of the latest service packs and patches for that version. For information, see [“Upgrading NetWare Clusters”](#) in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

If you have a NetWare 5.1 cluster, you must upgrade all nodes to a NetWare 6.5 cluster (with the latest service packs and patches) before adding new Linux cluster nodes to the cluster. For information, see [“Upgrading NetWare Clusters”](#) in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

6.1.2 SBD Devices Must Be Marked as Shareable for Clustering

Novell Cluster Services for Linux requires that the devices used for the SBD partition be explicitly marked as Shareable for Clustering. When converting a NetWare cluster, make sure that the SBD device, or both devices for a mirrored SBD, are marked as Shareable for Clustering before you add the first Linux node to the cluster.

6.1.3 Syntax Translation Issues for Load and Unload Scripts

When cluster migrating resources in a mixed-platform cluster from a NetWare cluster node to a OES 2 Linux cluster node, each cluster resource’s Load script and Unload script needs to be translated in-memory while the cluster contains mixed-platform nodes, and when the cluster is finally converted from NetWare to Linux. Executing a script that is valid for the NetWare platform is not necessarily recognized on the OES 2 Linux platform. This translation is done by the Cluster Translation Library script (`/opt/novell/ncs/bin/clstrlib.py`). If the commands in cluster resource’s load or unload scripts are not part of the normal translation library, the cluster resource can end up in a comatose state.

Beginning in OES 2 SP2, Novell Cluster Services allows you to customize the translation syntax that is used for load and unload scripts in mixed-platform situations by defining them in the `/var/opt/novell/ncs/customized_translation_syntax` file that you create. The `clstrlib.py` script reads the additional translation syntax from the syntax file. For information, see [Section 6.7, “Customizing the Translation Syntax for Converting Load and Unload Scripts,”](#) on page 91.

6.1.4 Case Sensitivity Issues

When adding a Linux node to the existing NetWare cluster, there are two areas where case sensitivity might be an issue:

- ♦ **Node name:** After you install the Linux node into the NetWare cluster, the Linux node is unable to join the cluster. To resolve this problem, edit the `/etc/opt/novell/ncs/nodename` file to modify the hostname of the node from lowercase (`clusnode1`) to all uppercase characters (`CLUSNODE1`), reboot the server, then run the `rcnovell-ncs start` command. This allows the cluster node to start and join the cluster.

NOTE: This case sensitivity issue has been resolved for OES 2 SP2 and later.

- ♦ **Cluster object name:** The Cluster object name (such as `cn=Clustername,ou=context,o=org`) is also present on the SBD partition. The SBD name (`Clustername.sbd`) matches the case of the object name in eDirectory. Running the `sbdutil -f` command displays the SBD name. If the case used when you enter the Cluster object name and SBD name during the Linux cluster install do not match the case used in eDirectory, the cluster install fails to detect the SBD partition.

6.1.5 Adding New NetWare Nodes to a Mixed-Node Cluster

You cannot add additional NetWare nodes to your cluster after adding a new Linux node or changing an existing NetWare cluster node to a Linux cluster node. If you want to add NetWare cluster nodes after converting part of your cluster to Linux, you must first remove the Linux nodes from the cluster.

6.1.6 Converting Multiple NetWare Cluster Nodes to OES 2 Linux

If you attempt to concurrently convert multiple NetWare cluster servers to OES 2 Linux, we strongly recommend that you use the old NetWare node IP addresses for your Linux cluster servers. You should record the NetWare node IP addresses before converting them to Linux.

If you must assign new node IP addresses, we recommend that you only convert one node at a time.

Another option if new cluster node IP addresses are required and new server hardware is being used is to shut down the NetWare nodes that are to be removed and then add the new Linux cluster nodes. After adding the new Linux cluster nodes, you can remove the NetWare cluster node-related objects as described in [Step 5 of Section 6.4, “Converting NetWare Cluster Nodes to OES 2 Linux \(Rolling Cluster Conversion\),”](#) on page 82.

IMPORTANT: Failure to follow these recommendations might result in NetWare server abends and Linux server restarts.

6.1.7 Converting Nodes that Contain the eDirectory Master Replica

When converting NetWare cluster servers to Linux, do not convert the server that has the master eDirectory replica first. If the server with the eDirectory master replica is a cluster node, convert it at the end of the rolling cluster conversion.

6.1.8 Failing Over Cluster Resources on Mixed-Node Clusters

Cluster resources that were created on NetWare cluster nodes and migrated or failed over to Linux cluster nodes can be migrated or failed back to NetWare cluster nodes.

Cluster resources that were originally created on Linux cluster nodes cannot be migrated or failed over to NetWare cluster nodes.

If you cluster migrate an NSS pool from a NetWare cluster server to a Linux cluster server, it could take several minutes for volume trustee assignments to synchronize after the migration. Users might have limited access to the migrated volumes until after the synchronization process is complete.

6.1.9 Managing File Systems in Mixed-Node Clusters

In a mixed cluster of NetWare and OES 2 Linux nodes, Linux POSIX file systems as cluster resources cannot be created until the entire cluster had been successfully converted to OES 2 Linux. Linux POSIX file systems as cluster resources cannot be migrated or failed over to NetWare cluster nodes.

Only NSS pool cluster resources that are created on a NetWare cluster node can be failed over between Linux and NetWare nodes of a mixed-node cluster.

NetWare-to-Linux (and vice versa) failover of NSS pool cluster resources requires that the Linux node be configured for NSS and that the version of NSS supports the NSS media format and features that are currently being used by the NSS pool cluster resource on NetWare.

No storage management functions should be executed while a cluster is in a mixed-cluster mode. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

WARNING: Attempting to reconfigure shared storage in a mixed cluster can cause data loss.

If you need to configure (or reconfigure) existing shared NSS pools and volumes in a mixed-node cluster, you must temporarily bring down all Linux cluster nodes prior to making changes, then make the configuration changes on a NetWare node. Make sure the resources are working properly on NetWare before having the Linux cluster nodes rejoin the cluster.

6.1.10 Using Novell iManager in Mixed-Node Clusters

Use Novell iManager 2.7.2 or later for all cluster administration in the mixed-node cluster. Using the Clusters plug-in to iManager is required to manage the cluster after the first OES 2 Linux node is added to the cluster.

The display of node IDs from the NetWare master node might be incomplete if you use other tools like ConsoleOne and Novell Remote Manager in a mixed-node cluster. However, you can use `cat /admin/Novell/Cluster/NodeConfig.xml` on any cluster node to get the node IDs.

6.1.11 Using Novell Remote Manager Is Not Supported in Mixed-Node Clusters

Do not use Novell Remote Manager when managing mixed-node clusters. Novell Remote Manager is not supported for cluster management on OES 2 Linux.

Because different time formats are used in the NCS Event log for NetWare and Linux, Novell Remote Manager might have difficulty displaying the time of logged events. To avoid this problem in a mixed-node cluster, use iManager to access the NCS Event log.

To reduce any confusion you might have when using Novell Remote Manager, you can unload module `pcluster.nlm` and delete its references in `ldnccs` and `uldncs`. This removes the *Cluster* tab in Novell Remote Manager.

6.1.12 Using ConsoleOne Is Not Supported for Mixed-Node Clusters

Do not use ConsoleOne when managing mixed-node clusters. ConsoleOne is not supported for cluster management on OES 2 Linux.

6.1.13 Using the Monitor Function in Mixed-Node Clusters Is Not Supported

In mixed-node clusters, the Monitor function in Novell Cluster Services for Linux is not available. You cannot enable the Monitor function or modify the Monitor script for cluster resources on the Linux nodes until the conversion is finalized and all nodes in the cluster are running OES 2 Linux.

6.2 Guidelines for Converting NSS Pool Resources from NetWare to Linux

- ♦ [Section 6.2.1, “NSS Pool Cluster Migration,” on page 67](#)
- ♦ [Section 6.2.2, “NSS File System Migration to NCP Volumes or Linux POSIX File Systems,” on page 68](#)
- ♦ [Section 6.2.3, “Estimated Time Taken to Build the Trustee File on Linux,” on page 68](#)

6.2.1 NSS Pool Cluster Migration

In the mixed-node cluster, NSS pool cluster resources created on NetWare can be failed over or cluster migrated to nodes that are running OES 2 Linux where NSS is installed and running. Some NSS features are not available or work differently on Linux. For information, see “[Cross-Platform Issues for NSS](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

Pool snapshots use different technologies on NetWare and Linux. Make sure to delete pool snapshots for all clustered pools before you begin the cluster conversion.

6.2.2 NSS File System Migration to NCP Volumes or Linux POSIX File Systems

To move data from NSS file systems on NetWare to NCP volumes or to Linux POSIX file systems on Linux, you must use the OES 2 SP1 Migration tool. For information, see “[Migrating File System from NetWare, OES 1 or OES 2 to OES 2 SP2 Linux](#)” in the *OES 2 SP2: Migration Tool Administration Guide*.

6.2.3 Estimated Time Taken to Build the Trustee File on Linux

When you migrate NSS volumes from NetWare to Linux, a `trustee.xml` file is built at the root of the volume.

Testing found that this process takes about one minute per 50,000 storage objects. Testing was done on the following configuration for the target server:

HP DL380 G5
2 Quadcore Intel Xeon CPU E5345 @ 2.33 GHz
12 GB RAM
1 Gigabit NIC
2 HBAs with 4 paths to the EMC DMX Symmetrix Storage with 4 gigabits per second (Gbps) bandwidth

6.3 Guidelines for Converting Service Cluster Resources from NetWare to Linux

Converting cluster resources for OES 2 services from NetWare to Linux might require more than a simple cluster migration from a NetWare node to a Linux node. For example, the service might require that you use migration tools to convert the service to Linux. Some services require post-conversion configuration to finalize the conversion. A few services on NetWare are not available on OES 2 Linux, so you must use the standard Linux service instead.

- ◆ [Section 6.3.1, “Overview of All NetWare 6.5 SP8 Services,” on page 69](#)
- ◆ [Section 6.3.2, “Apache Web Server,” on page 71](#)
- ◆ [Section 6.3.3, “Apple Filing Protocol \(AFP\),” on page 71](#)
- ◆ [Section 6.3.4, “Archive and Version Services,” on page 71](#)
- ◆ [Section 6.3.5, “CIFS,” on page 72](#)
- ◆ [Section 6.3.6, “DFS VLDB,” on page 72](#)
- ◆ [Section 6.3.7, “DHCP Server,” on page 73](#)
- ◆ [Section 6.3.8, “DNS Server,” on page 75](#)
- ◆ [Section 6.3.9, “eDirectory Server Certificates,” on page 76](#)
- ◆ [Section 6.3.10, “iPrint,” on page 78](#)
- ◆ [Section 6.3.11, “QuickFinder Server,” on page 79](#)

6.3.1 Overview of All NetWare 6.5 SP8 Services

See [Table 6-1](#) for information about converting cluster resources for NetWare 6.5 SP8 services:

Table 6-1 Guidelines for Converting Service Cluster Resources from NetWare to Linux

Service on NetWare 6.5 SP8	Cluster Migrate the Resource	Converting the Service to OES 2 Linux
Apache Web Server	Requires special handling	See Section 6.3.2, "Apache Web Server," on page 71.
Apple Filing Protocol (AFP)	Requires special handling	See Section 6.3.3, "Apple Filing Protocol (AFP)," on page 71.
Archive and Version Services	No	On Linux, you must configure a new cluster resource on a shared Linux POSIX file system. On NetWare, Archive and Versioning Services uses a MySQL database. On Linux, it uses a PostgreSQL database. The load script commands are also different. See Section 6.3.4, "Archive and Version Services," on page 71.
CIFS (Windows File Services)	Requires special handling	See Section 6.3.5, "CIFS," on page 72.
DFS VLDB (Distributed File Services volume location database)	Requires special handling	See Section 6.3.6, "DFS VLDB," on page 72.
DHCP Server	Requires special handling	See Section 6.3.7, "DHCP Server," on page 73.
DNS Server	Requires special handling	See Section 6.3.8, "DNS Server," on page 75.
eDirectory	Not clustered	See Section 6.1.7, "Converting Nodes that Contain the eDirectory Master Replica," on page 66.
eDirectory Certificate Server	Requires special handling	The Certificate Authority (CA) service is not cluster-enabled for NetWare or OES 2 Linux. There are no cluster-specific tasks for the CA itself. The Server Certificate service issues Server Certificate objects that might need to reside on each node in a cluster, depending on the service that is clustered. NetWare and Linux generate certificates differently, so the NetWare server's certificate is not reused for the OES 2 Linux server. See Section 6.3.9, "eDirectory Server Certificates," on page 76.
exteNd Application Server and MySQL	Not applicable	The exteNd Application Server was discontinued as an install option for NetWare 6.5 SP3. See MySQL .

Service on NetWare 6.5 SP8	Cluster Migrate the Resource	Converting the Service to OES 2 Linux
FTP	Not applicable	Use the standard FTP service for Linux.
iFolder	Requires special handling	<p>Novell iFolder 2.1x is not available on OES 2 Linux. You must upgrade to Novell iFolder 3.x.</p> <p>After you add a Novell iFolder 3.x server to the NetWare cluster and before you finalize the cluster conversion, use iFolder migration procedures to migrate the iFolder 2.1x server configuration and user data from the source NetWare node to the target Linux node. For information, see “Migrating iFolder Services” in the <i>Novell iFolder 3.8 Administration Guide</i>.</p>
iPrint	Requires special handling	See Section 6.3.10, “iPrint,” on page 78.
MySQL	Not applicable	<p>Use the MySQL 5.0.x software on OES2 Linux that is offered under the GPL.</p> <p>Configure the OES service to use MySQL 5.0.x on OES 2 Linux before setting up clustering for the related MySQL database.</p> <p>For Linux, use a procedure similar to the one on NetWare to set up a new cluster resource. Use the Linux commands for MySQL in the load and unload scripts. Use a Linux path on a shared Linux POSIX file system for the MySQL database. As a general reference, see “Configuring MySQL on Novell Clustering Services” in the <i>NW 6.5 SP8: Novell MySQL Administration Guide</i></p>
NetStorage	Yes	<p>Clustering the NetStorage service is supported for OES 2 SP1 Linux and later.</p> <p>For information, see “Configuring NetStorage with Novell Cluster Services” in the <i>OES 2 SP2: NetStorage for Linux Administration Guide</i>.</p>
NFS	Not applicable	Use standard NFS service for Linux.
QuickFinder (Server Synchronization Feature)	No	<p>You must create a new cluster resource. QuickFinder 5.0.x is supported only on OES 2 Linux. NetWare uses QuickFinder 4.2.0. QuickFinder does not support any automated procedure or scripts for a rolling upgrade from Netware to Linux.</p> <p>For information, see “Configuring QuickFinder Server for Novell Cluster Services” in the <i>NW 6.5 SP8 Novell QuickFinder Server 5.0 Administration Guide</i>.</p>
Tomcat	Not applicable	Use the standard Tomcat service for Linux.

6.3.2 Apache Web Server

- 1 On NetWare, offline the NSS pool cluster resource, then modify its load and unload scripts to remove the Apache start and stop commands.
- 2 On NetWare, online the cluster resource, then cluster migrate it to a Linux node.
- 3 After the cluster conversion is finished, use the standard Apache Web Server for Linux to set up the Apache service on the OES 2 Linux servers.
- 4 Use a procedure similar to the one on NetWare to set up the Apache configuration file, and copy it to every Linux node in the cluster. Point the service to the virtual IP address of the NSS pool cluster resource that contains the Web content.
- 5 On Linux, offline the cluster resource, then modify its load and unload scripts to add the Apache service start and stop commands for Linux.
- 6 Online the cluster resource.

6.3.3 Apple Filing Protocol (AFP)

Novell AFP for Linux is available beginning in OES 2 SP1 Linux. After you set up Novell AFP on the Linux node and before you finalize the NetWare-to-Linux conversion, use the AFP migration tool to convert the configuration. For information, see [“Migrating AFP from NetWare to OES 2 SP2 Linux”](#) in the *OES 2 SP2: Migration Tool Administration Guide*.

The commands in the scripts are also different. After the migration, modify the load and unload scripts on the Linux server. For information, see [Section 6.6.4, “Comparing File Access Protocol Resource Script Commands,”](#) on page 89.

AFP on Linux supports NCP cross-protocol file locking, which allows NCP, AFP, and CIFS users to access files on an NSS volume concurrently without data corruption by locking the files across protocols. Beginning in OES 2 SP2 Linux, the cross-protocol file locking parameter for NCP Server is enabled by default. It must be enabled on each node in the cluster if you plan to give both NCP users and AFP users access to NSS volume in the cluster. See [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

6.3.4 Archive and Version Services

Mixed-node cluster configurations are not supported by Novell Archive and Version Services.

Before you begin the conversion, make sure that Archive and Version Services is not running on the NetWare servers in the cluster.

- 1 Install Archive and Version Services on an OES 2 Linux node in the cluster.
- 2 Install Archive and Version Services on a second OES 2 Linux node in the cluster.
- 3 Using the database migration tools, migrate the data in the MySQL database on NetWare to the PostgreSQL database on one of the Linux nodes.
- 4 Cluster migrate the shared NSS pool resource that contains the volumes that were being archived from the NetWare server to a Linux node.
- 5 Remove the NetWare nodes from the cluster and finish the cluster conversion process.
- 6 On the Linux node where the primary NSS pool resources are active, use the *Clusters* plug-in in iManager to create an Archive Versioning cluster resource.

Other preparatory tasks are involved in the setup. Follow the procedure as described in “[Configuring Archive and Version Service for Novell Cluster Services](#)” in the *OES 2 SP2: Novell Archive and Version Services 2.1 for Linux Administration Guide*.

- 7 Copy the database files from the single-server location (`/var/opt/novell/arkmanager/data`) to the shared Linux POSIX volume that you created when you set up Archive and Version Services for clustering in [Step 6](#).

Use the `cp -a` command at a terminal console prompt to copy all files and retain the permissions.

- 8 Change the ownership of the new database location on the shared volume by entering the following at a terminal console prompt:

```
chown -R arkuser:arkuser_prggrp /shared/datapath
```

- 9 Edit the `/etc/opt/novell/arkmanager/conf/arkdatadir.conf` file to change the database location to new shared path.

- 10 Edit the `/opt/novell/arkmanager/bin/pg_restart.sh` file to change the line that starts the PostgreSQL database to the following:

```
su arkuser -c "postmaster -D /shared/datapath -h 127.0.0.1 -p 5432 -i"
```

- 11 Start Archive and Version Services by entering

```
rcnovell-ark start
```

You should see Archive and Version Services and the PostgreSQL database starting.

6.3.5 CIFS

Novell CIFS for Linux is available beginning in OES 2 SP1 Linux.

After you set up Novell CIFS on the Linux node and before you finalize the NetWare-to-Linux conversion, use the CIFS migration tool to convert the configuration. For information, see “[Migrating CIFS from NetWare to OES 2 SP2 Linux](#)” in the *OES 2 SP2: Migration Tool Administration Guide*.

The commands in the scripts are also different. After the migration, modify the load and unload scripts on the Linux server. For information, see [Section 6.6.4, “Comparing File Access Protocol Resource Script Commands,”](#) on page 89.

CIFS on OES 2 SP1 Linux does not support NCP cross-protocol file locking.

Beginning in OES 2 SP2 Linux, CIFS supports NCP cross-protocol file locking, which allows NCP, AFP, and CIFS users to access files on an NSS volume concurrently without data corruption by locking the files across protocols. Beginning in OES 2 SP2 Linux, the cross-protocol file locking parameter for NCP Server is enabled by default. It must be enabled on each node in the cluster if you plan to give both NCP users and CIFS users access to an NSS volume in the cluster. See “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

6.3.6 DFS VLDB

The Novell Distributed File Services volume location database (VLDB) `.dat` file format is the same on both NetWare and Linux. The shared NSS volume that contains the `.dat` file can be cluster migrated to the Linux server.

Use one of these two methods for migrating the VLDB from NetWare to Linux:

- ♦ [“Cluster Migrating the Shared NSS Volume for the VLDB” on page 73](#)
- ♦ [“Adding a Linux Server as a Replica Site” on page 73](#)

Cluster Migrating the Shared NSS Volume for the VLDB

Use this method if you want to use the same shared disk where the VLDB is currently stored.

- 1** Install Novell Storage Services and any dependent services on the Linux node, then add it to the mixed cluster that you are converting.
- 2** Cluster migrate the DFS cluster resource from NetWare to Linux.
- 3** On the Linux node where the VLDB is active, offline the DFS cluster resource.
- 4** Remove the NetWare clusters from the cluster by using the `cluster leave` command, then finish the cluster conversion.

This automatically updates the basic cluster commands in the cluster resource scripts.

- 5** Using the *Clusters* plug-in in iManager, modify the load script of the DFS cluster resource to change the `vldb` command to the Linux format. For example, change it from

```
vldb /dir=vldbpath
```

to

```
vldb -dir /vldbpath
```

- 6** Online the cluster resource.
- 7** Run a VLDB repair to ensure that the database is correct.

Adding a Linux Server as a Replica Site

Use this method if you want to use a different shared disk for the VLDB on Linux. You can do this by adding a DFS replica site on Linux.

- 1** Install OES 2 Linux on the server that you want to add to the cluster. Make sure Novell Storage Services and any dependent services are installed.
- 2** Create a shared NSS pool and volume on the OES 2 Linux server, or create a shared Linux POSIX volume.
- 3** In iManager, add the Linux server as the second VLDB replica site for the DFS management context, and point to the shared NSS volume as the VLDB location.
- 4** Allow the VLDB data to synchronize between the NetWare replica and the Linux replica.
- 5** In iManager, remove the NetWare instance of the replica site.
- 6** Add the Linux server to the mixed-node NetWare cluster.
- 7** Continue with the cluster conversion as described in [Section 6.4, “Converting NetWare Cluster Nodes to OES 2 Linux \(Rolling Cluster Conversion\),” on page 82.](#)

6.3.7 DHCP Server

The Novell DHCP Server for Linux is a standards compliant implementation that is based on the `bind` protocol. DHCP uses a different schema on Linux to store the configuration in eDirectory.

Novell DHCP Server for Linux supports using a shared Linux POSIX file system or a shared NSS file system for the cluster resource. For information, see [“Configuring DHCP with Novell Cluster Services for the NSS File System”](#) and [“Configuring DHCP with Novell Cluster Services for the Linux File System”](#) in the *OES 2 SP2: Novell DNS/DHCP Administration Guide for Linux*.

After you set up Novell DHCP Server on the Linux server, you can use the DHCP Migration utility to convert the configuration from NetWare to Linux. You cannot directly reuse the data. Use one of the following scenarios to migrate your DNS server data, then perform the post-migration tasks to set up clustering.

For information about prerequisites, see [“Migrating DHCP from NetWare to OES 2 SP2 Linux”](#) in the *OES 2 SP2: Migration Tool Administration Guide*.

- ◆ [“NetWare and Linux Clusters Are in the Same Tree”](#) on page 74
- ◆ [“NetWare and Linux Clusters Are in Different Trees”](#) on page 74
- ◆ [“Post-Migration Tasks”](#) on page 75

NetWare and Linux Clusters Are in the Same Tree

In this scenario, both the NetWare server and the OES 2 SP1 Linux server are in the same eDirectory tree. The NetWare source server must be running NetWare 5.1 or later versions. The Linux target server must be running OES 2 SP1 Linux on either 32-bit or 64-bit hardware.

Run the DHCP migration tool from one of the Linux nodes. Perform the Tree Level Migration with the same Source server (tree to which NetWare clustered nodes are attached) and Target server (tree to which the Linux clustered nodes are attached). This ensures that the entire NetWare DHCP configuration data is available for Linux DHCP. For information see [“NetWare and Linux Clusters Attached to the Same Tree”](#) in the *OES 2 SP2: Migration Tool Administration Guide*.

IMPORTANT: Before starting the DHCP server on the Linux cluster, stop the DHCP server on the Netware cluster.

NetWare and Linux Clusters Are in Different Trees

In this scenario, the NetWare server and the OES 2 SP1 Linux server are on different eDirectory trees. The NetWare source server must be running NetWare 5.1 or later versions. The Linux target server must be running OES 2 SP1 Linux on either 32-bit or 64-bit hardware.

Run the DHCP migration tool from one of the Linux nodes. Perform the Tree Level Migration with a different Source server (tree to which NetWare clustered nodes are attached) and Target server (tree to which the Linux clustered nodes are attached). This ensures that the entire NetWare DHCP configuration data is available for Linux DHCP. For information, see [“NetWare and Linux Clusters Attached to Different Trees”](#) in the *OES 2 SP2: Migration Tool Administration Guide*.

IMPORTANT: Before starting the DHCP server on the Linux cluster, stop the DHCP server on the Netware cluster.

Post-Migration Tasks

- 1 On the Linux node where you ran the migration, configure the DHCP service for clustering by using one of the following methods described in [OES 2 SP2: Novell DNS/DHCP Administration Guide for Linux](#):
 - ♦ “Configuring DHCP with Novell Cluster Services for the NSS File System”
 - ♦ “Configuring DHCP with Novell Cluster Services for the Linux File System”
- 2 Online the DHCP service cluster resource.
- 3 On the Linux node where you ran the migration:
 - 3a Open the `/mountpath/etc/dhcpd.conf` file in a text editor.

Replace `/mountpath` with the Linux path to the directory in the shared volume where DHCP-specific directories are created.
 - 3b In the `/mountpath/etc/dhcpd.conf` file, change the value for the `ldap-dhcp-server-cn` parameter to the cn needed for the migrated server, then save your changes.
 - 3c Copy the `migrated_server.leases` file from `/var/lib.dhcp/bd` directory to the `/mountpath/var/lib/dhcp/db/` directory, then rename it here to `dhcpd.leases`.
- 4 Stop the DHCP server on the NetWare cluster.
- 5 Start the DHCP server on the Linux cluster.

6.3.8 DNS Server

You can migrate the data from the Novell DNS Server on NetWare to a Novell DNS Server on Linux after you have installed and set up DNS services on an OES 2 SP1 Linux node in the cluster. You cannot directly reuse the data. Use one of the following scenarios to migrate your DNS server data, then perform the post-migration tasks.

For information about prerequisites, see “[Migrating DNS from NetWare to OES 2 SP2 Linux](#)” in the [OES 2 SP2: Migration Tool Administration Guide](#).

- ♦ “[NetWare and Linux Clusters Are in the Same Tree](#)” on page 75
- ♦ “[NetWare and Linux Clusters Are in Different Trees](#)” on page 76
- ♦ “[Post-Migration Tasks](#)” on page 76

NetWare and Linux Clusters Are in the Same Tree

In this scenario, both the NetWare server and the OES 2 SP1 Linux server are in the same eDirectory tree. The NetWare source server must be running NetWare 5.1 or later versions. The Linux target server must be running OES 2 SP1 Linux on either 32-bit or 64-bit hardware. Use iManager to move the DNS server from a NetWare NCP server to an OES 2 SP1 Linux NCP server.

Run the DNS migration tool from one of the Linux nodes. Perform the Tree Level migration with the same Source server (tree to which NetWare clustered nodes are attached) and Target server (tree to which the Linux clustered nodes are attached). This ensures that the entire NetWare DNS configuration data is available for Linux DNS. For information see “[Using iManager to Migrate Servers within the Same eDirectory Tree](#)” in the [OES 2 SP2: Migration Tool Administration Guide](#).

IMPORTANT: Before starting the DNS server on the Linux cluster, stop the DNS server on the Netware cluster.

NetWare and Linux Clusters Are in Different Trees

In this scenario, the NetWare server and the OES 2 SP1 Linux server are on different eDirectory trees. The NetWare source server must be running NetWare 5.1 or later versions. The Linux target server must be running OES 2 SP1 Linux on either 32-bit or 64-bit hardware.

Run the DNS migration tool from one of the Linux nodes. Perform the Tree Level Migration with a different Source server (tree to which NetWare clustered nodes are attached) and Target server (tree to which the Linux clustered nodes are attached). This ensures that the entire NetWare DNS configuration data is available for Linux DNS. For information see “[Using iManager to Migrate Servers across eDirectory Trees](#)” in the *OES 2 SP2: Migration Tool Administration Guide*.

IMPORTANT: Before starting the DNS server on the Linux cluster, stop the DNS server on the Netware cluster.

Post-Migration Tasks

See “[Post-Migration Procedure](#)” in the *OES 2 SP2: Migration Tool Administration Guide*.

6.3.9 eDirectory Server Certificates

Novell Certificate Server provides two categories of services: Certificate Authority (CA) and Server Certificates. The Certificate Authority services include the Enterprise CA and CRL (Certificate Revocation List). Only one server can host the CA, and normally that same server hosts the CRLs if they are enabled (although if you move the CA to a different server, the CRLs usually stay on the old server). The CA and CRL services are not cluster-enabled in either NetWare or OES 2 Linux, and therefore, there are no cluster-specific tasks for them.

Novell Certificate Server provides a Server Certificates service for NetWare and Linux. The service is not clustered. However, clustered applications that use the server certificates must be able to use the same server certificates on whichever cluster node they happen to be running. Use the instructions in the following sections to set up Server Certificate objects in a clustered environment to ensure that your cryptography-enabled applications that use Server Certificate objects always have access to them.

The eDirectory Server Certificate objects are created differently in OES 2 Linux and cannot be directly reused from the NetWare server. The differences and alternatives for setting up certificates on Linux are described in the following sections:

- ♦ “[Server Certificate Changes in OES 2 Linux](#)” on page 76
- ♦ “[Using Internal Certificates in a Cluster](#)” on page 77
- ♦ “[Using External Certificates in a Cluster](#)” on page 77

Server Certificate Changes in OES 2 Linux

When you install NetWare or OES 2 Linux in an eDirectory environment, the Server Certificate service can create certificates for eDirectory services to use. In addition, custom certificates can be created after the install by using iManager or command line commands.

For NetWare, all applications are integrated with eDirectory. This allows applications to automatically use the server certificates created by Novell Certificate Server directly from eDirectory. In a NetWare cluster, you might have copied the Server Certificate objects to all nodes in

the cluster using backup and restore functions as described in “[Server Certificate Objects and Clustering](http://www.novell.com/documentation/crt33/crtadmin/data/a2ebopb.html#acebe5n)” (<http://www.novell.com/documentation/crt33/crtadmin/data/a2ebopb.html#acebe5n>) in the *Novell Certificate Server 3.3.2 Administration Guide* (<http://www.novell.com/documentation/crt33/crtadmin/data/a2ebomw.html>).

For OES 2 Linux, many applications (such as Apache and Tomcat) are not integrated with eDirectory and therefore, cannot automatically use the certificates created by Novell Certificate Server directly from eDirectory. By default, these services use self-signed certificates, which are not in compliance with the X.509 requirements as specified in RFC 2459 and RFC 3280.

To address the difference, Novell Certificate Server offers an install option for OES 2 Linux called *Use eDirectory Certificates* that automatically exports the default eDirectory certificate *SSL Certificate DNS* and its key pair to the local file system in the following files:

```
/etc/ssl/servercerts/servercert.pem  
/etc/ssl/servercerts/serverkey.pem
```

Using Internal Certificates in a Cluster

Recent versions of Novell Certificate Server create default certificates that allow you to specify an alternative IP address or DNS address by adding it in the Subject Alternative Name extension. This requires that your DNS service be configured to reflect the cluster IP/DNS address as the default (or first) address. If the DNS service is set up correctly, the cluster applications can use the default certificates without needing any administration.

IMPORTANT: If the DNS service is not set up correctly, then you must use the process described for external certificates in “[Using External Certificates in a Cluster](#)” on page 77.

For OES 2 Linux clusters using the internal certificate method, make sure the DNS service is configured to use the cluster IP/DNS address. During the OES 2 Linux install, select the *Use eDirectory Certificates* option so that Novell Certificate Server automatically creates the *SSL Certificate DNS* certificate with the correct IP/DNS address. By selecting the *Use eDirectory Certificates* option during the install and using the cluster IP/DNS address, clustered applications should be able to access the certificates without needing further configuration for the Server Certificate object.

Using External Certificates in a Cluster

External (third-party) certificates create a Server Certificate object that includes the cluster's IP and/or DNS address. Create a backup of this certificate. For each server in the cluster, create a Server Certificate object with the same name by importing the previously created backup certificate and key pair to a location on that server. This allows all of the servers in the cluster to use and share the same certificate and key pair. After all cluster nodes have the certificate, configure the cluster applications to use the server certificate.

IMPORTANT: This cluster task can also be used for sharing internal certificates on the cluster nodes. In early versions of Novell Certificate Server, this was the only option available.

For information about exporting and using eDirectory Server Certificates for External Services, see “[Using eDirectory Certificates with External Applications](http://www.novell.com/documentation/crt33/crtadmin/data/bh9x78f.html)” (<http://www.novell.com/documentation/crt33/crtadmin/data/bh9x78f.html>) in the *Novell Certificate Server 3.3.2 Administration Guide* (<http://www.novell.com/documentation/crt33/crtadmin/data/a2ebomw.html>).

For OES 2 Linux clusters using the external certificate method, the solution is more complicated than for internal certificates. You must create the certificate for each server in the cluster just as you did for NetWare. You must also create a configuration on the SAS:Service object for each server so that the common certificate is automatically exported to the file system where the non-eDirectory enabled applications can use it.

6.3.10 iPrint

After adding the OES 2 Linux node to the NetWare cluster, you must use the following procedure to set up clustering for iPrint on Linux, then migrate the iPrint information from a NetWare shared NSS pool resource to a newly created Linux shared NSS pool resource on the Linux node.

- 1 On a Linux node, create a new shared NSS pool and volume.
- 2 Log in as the `root` user to the Linux node where the shared pool resource is active, go to the `/opt/novell/iprint/bin` directory, then run the `iprint_nss_relocate` script by entering

```
./iprint_nss_relocate -a admin_fdn -p password -n nss_volume_path -l cluster
```

Replace `admin_fdn` with the comma-delimited fully distinguished name of the iPrint administrator user (such as `cn=admin,o=mycompany`). Replace `password` with the actual password of the iPrint administrator user. Replace `nss_volume_path` with the Linux path (such as `/media/nss/NSSVOL1`) to the shared NSS volume where you want to relocate the iPrint configuration data.

- 3 Review the messages displayed on the screen to confirm the data migration from the local Linux path to the shared NSS path is completed.

For example, enter

```
./iprint_nss_relocate -a cn=admin,o=mycompany -p pass -n /media/nss/NSSVOL1 -l cluster
```

For information, see “[Executing the Script](#)” in the *OES 2 SP2: iPrint for Linux Administration Guide*.

- 4 For each Linux node in the cluster where iPrint is installed, set up clustering for iPrint.
 - 4a In iManager, select *Clusters > Cluster Manager*, then cluster migrate the shared NSS pool resource from the active Linux node to another Linux node.
 - 4b Log in to the Linux node as the `root` user, then run the `iprint_nss_relocate` script as described in [Step 2](#), using the same values.
- 5 In iManager, click *Clusters > Cluster Manager*, then select the Linux node where the shared NSS pool is currently active.
- 6 Select the Linux shared NSS pool, then go to the *Preferred Nodes* tab and move all of the remaining NetWare nodes from the *Assigned Nodes* to *Unassigned Nodes* column to prevent an inadvertent failback of the resource to a NetWare server.
- 7 In iManager, select *iPrint*, then create a Driver Store (*iPrint > Create Driver Store*) and a Print Manager (*iPrint > Create Print Manager*) on the Linux node with the IP or DNS name of the shared NSS pool resource.

IMPORTANT: Do not modify the load and unload scripts for the Linux shared NSS pool resource at this time.

- 8 Use the Migration tool to migrate data from the NetWare shared NSS pool to the Linux shared NSS pool.

For information about using the Migration tool for iPrint migration, see [“Migrating iPrint from NetWare to OES 2 Linux”](#) in the *OES 2 SP2: Migration Tool Administration Guide*.

- 8a** Start the migration tool from the target server, then authenticate by using the IP address or DNS name of the NetWare shared NSS pool resource.
 - 8b** For the target server, authenticate by using the IP address or DNS name of the Linux shared NSS pool resource.
 - 8c** Configure the Migration tool for migrating iPrint information, then proceed with the migration as described in [“Migrating iPrint from NetWare to OES 2 Linux”](#).
- 9** Edit the load and unload scripts for the Linux shared NSS pool resource. For information, see [“Prerequisites”](#) in the *OES 2 SP2: iPrint for Linux Administration Guide*.

6.3.11 QuickFinder Server

In a Novell Cluster Services cluster, you must install QuickFinder on each node in the cluster. This registers QuickFinder Server with each of the Web servers and application servers running on each server. On OES 2 Linux, QuickFinder is installed by default in the `/var/lib/qfsearch` directory. We recommend that you use the default path. After the installation, you must set up one or more virtual search servers to enable QuickFinder Server to work in a cluster.

When the Linux setup is completed, you are ready to manually migrate settings from the NetWare cluster to the Linux cluster. Set up QuickFinder on the OES 2 Linux cluster nodes, then manually migrate QuickFinder data from a NetWare node to an OES 2 Linux node.

For information about using the QuickFinder Server Manager and other procedures for QuickFinder, see the *OES 2: Novell QuickFinder Server 5.0 Administration Guide*.

- ♦ [“Prerequisites”](#) on page 79
- ♦ [“Setting Up QuickFinder Server on Linux Cluster Nodes”](#) on page 79
- ♦ [“Migrating QuickFinder Data from NetWare to Linux”](#) on page 80
- ♦ [“Post-Migration Considerations”](#) on page 81
- ♦ [“Searching the Cluster Volume”](#) on page 81

Prerequisites

Before you begin:

- 1** On one Linux node, create a Linux POSIX cluster resource where all of the indexes and virtual search server settings are to be stored.

For information, see [Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,”](#) on page 163.

Setting Up QuickFinder Server on Linux Cluster Nodes

On each OES 2 Linux node, do the following to set up QuickFinder for Linux:

- 1** Cluster migrate the Linux POSIX cluster resource to the OES 2 Linux node where you want to install QuickFinder
- 2** Install QuickFinder on the active cluster node.
- 3** Create a virtual search server to enable QuickFinder Server to work in a cluster.

Give each virtual search server the same name and location. After the first server is set up, any settings that you create on the shared volume are automatically displayed.

- 3a** On the active cluster node, open the QuickFinder Server Manager.
 - 3b** Click *Global Settings*, then click *Add New Virtual Server*.
 - 3c** In *Name*, specify the DNS name of the cluster.
 - 3d** In *Location*, specify the Linux path on the Linux POSIX cluster resource where all of the indexes and virtual search server settings will be located.
 - 3e** Click *Add*.
- 4** Repeat [Step 1](#) to [Step 3](#) for each of the nodes in the cluster.

Migrating QuickFinder Data from NetWare to Linux

Use the following steps to migrate QuickFinder Server data from a NetWare server to a corresponding Linux server. You must repeat the tasks for each NetWare server in the cluster. It assumes a one-to-one server replacement in the cluster.

WARNING: Migrating indexes and virtual search server settings from a QuickFinder Server running on NetWare to QuickFinder Server running on OES 2 Linux replaces the existing settings on the Linux server. If you want to merge your NetWare settings with the existing Linux settings, you must manually re-create the NetWare settings by using the QuickFinder Server Manager.

- 1** Open a Web browser, then access the QuickFinder Server Manager on the NetWare server.
`http://servername/qfsearch/admin`
- 2** Click *Global Settings* in the top toolbar.
- 3** Write down the paths for each virtual search server displayed in the *Location* column.
- 4** On the OES 2 Linux server where the shared volume is active, mount the NetWare server by using the `ncpmount` command.
- 5** Make a backup of the `/var/lib/qfsearch/SiteList.properties` file.
Make sure that you don't have a file with this name as a backup on the NetWare server.
- 6** Copy all `.properties` and `Cron.jobs` files from the root directory `sys:/qfsearch` on the NetWare server to `/var/lib/qfsearch` on the Linux server.
- 7** Copy `sys:/qfsearch/Sites` and all of its subdirectories to `/var/lib/qfsearch/Sites`.
- 8** Copy `sys:/qfsearch/Templates` and all of its subdirectories to `/var/lib/qfsearch/Templates`.
- 9** If any of the paths listed in [Step 3](#) are not under `sys:/qfsearch` (for example, if you installed a virtual search server somewhere other than the default location), you must also copy those paths to Linux.
For example, if you have the path `sys:/SearchSites/PartnerSite`, you must copy it to the Linux server. You could copy it to `/var/opt/SearchSites/PartnerSite` or `/var/lib/qfsearch/Sites/PartnerSite`.
- 10** Edit all NetWare paths in `/var/lib/qfsearch/SiteList.properties` to reflect the new Linux paths.
For example, change `sys:/qfsearch` to `/var/lib/qfsearch`.

Or, as in the example in [Step 9](#), change `sys:/SearchSites/PartnerSite` to `/var/opt/SearchSites/PartnerSite`.

Some paths might have one or two backslashes (\) that must be replaced with one forward slash (/). For example, `sys:\\qfsearch\\docs` needs to be changed to `/var/lib/qfsearch/docs`.

- 11 Update all NetWare paths in the properties and configuration files copied in the steps above to the Linux paths, and update any DNS names.

The following files must be updated:

```
AdminServlet.properties
Cron.jobs
Sites/Highlighter.properties
Sites/Print.properties
Sites/Search.properties
```

For each of the virtual search servers, modify the following:

- ◆ `qfind.cfg`
- ◆ Any of the above `.properties` files, if they exist.

IMPORTANT: Linux filenames are case sensitive.

The names of most properties files are mixed case, so make sure the files copied from NetWare are the correct case. You can compare them to the `.properties.sample` files on Linux.

You might also need to update paths in templates. If you have problems such as a template not being found or some properties not being set properly, check the case of the filename.

If you modified any “file” index paths to index directories on the Linux server, that index must be regenerated.

- 12 After all the files have been modified, run the following commands to set the access rights and the owner and groups so that the QuickFinder engine has rights to access the files.

As the `root` user, enter

```
chown -R root:www /var/lib/qfsearch
chmod -R 770 /var/lib/qfsearch
```

- 13 Repeat [Step 1](#) to [Step 12](#) for each NetWare and Linux pair of nodes.

Post-Migration Considerations

QuickFinder Server 5.0 indexes are not compatible with previous versions of QuickFinder Server. The indexes must be regenerated, and you cannot synchronize QuickFinder Server 5.0 indexes with indexes from a previous version of QuickFinder Server (and vice-versa).

Searching the Cluster Volume

To perform a search on the shared volume after the NetWare migration is complete:

- 1 Open a Web browser, then enter

```
http://DNS_CLUSTER/qfsearch/search
```

QuickFinder Server sees the DNS and sends the request to the appropriate virtual search server.

6.4 Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion)

Performing a rolling cluster conversion from NetWare 6.5 to OES 2 Linux lets you keep your cluster up and running and lets your users continue to access cluster resources while the conversion is being performed.

During a rolling cluster conversion, one server is converted to Linux while the other servers in the cluster continue running NetWare 6.5. Then, if desired, another server can be converted to OES 2 Linux, and then another, until all servers in the cluster have been converted to Linux. You can also leave the cluster as a mixed NetWare and Linux cluster.

The process for converting NetWare 6.0 cluster nodes to OES 2 Linux cluster nodes is the same as for converting NetWare 6.5 cluster nodes to OES 2 Linux cluster nodes.

IMPORTANT: Before you begin, make sure your system meets the following requirements and caveats:

- ♦ [Section 3.1, “System Requirements for Novell Cluster Services,”](#) on page 29
- ♦ [Section 6.1, “Guidelines for Converting Clusters from NetWare to OES 2 Linux,”](#) on page 63

If you are converting from NetWare on physical servers to OES 2 Linux on virtual servers (guest operating systems running on Xen virtual machines), you can use the same methods and processes as those used on a physical server. No additional changes or special configuration is required. For information, see [Section 12.5, “Mixed Physical and Virtual Node Clusters,”](#) on page 195.

To perform a rolling cluster conversion from NetWare 6.5 to OES 2 Linux:

- 1** Before you add the first Linux node to the NetWare cluster, if the NetWare cluster uses an SBD, make sure that the device (or devices) being used by the SBD are marked as *Shareable for Clustering*.

You can use NSSMU or iManager to mark the SBD devices as shareable. It is not necessary to bring the cluster down when changing the device attribute to *Shareable for Clustering*.

Using NSSMU:

- 1a** Log in to the master node of the NetWare cluster as the administrator user.
 - 1b** Enter `nssmu` at the server console prompt.
 - 1c** In the NSSMU main menu, select *Devices*.
 - 1d** In the *Devices* list, highlight the device that contains the SBD partition, then press F5 to select it.
 - 1e** Press F6 to mark the device as *Shareable for Clustering*.
 - 1f** If the SBD partition is mirrored, repeat [Step 1d](#) and [Step 1e](#) to also mark the mirror device as *Shareable for Clustering*.
 - 1g** Press Esc to exit NSSMU.
- 2** Make a note of the services that are installed on the server you are converting.
You might want to install the same components on the Linux node if they are available.
 - 3** On the NetWare server that you want to convert to Linux, remove eDirectory.

You can do this by running NWConfig, then selecting *Product Options > Directory Options <install NDS> > Remove Directory Services from this server.*

4 Bring down the NetWare server you want to convert to Linux.

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

You can also manually cluster migrate the resources to another server in the cluster prior to bringing down the server. This prevents the resources from failing back to the node after you have completed the upgrade.

5 In eDirectory, remove (delete) the Cluster Node object, the Server object, and all corresponding objects relating to the downed NetWare server.

Depending on your configuration, there could be 10 or more objects that relate to the downed NetWare server.

6 Run DSRepair from another server in the eDirectory tree to fix any directory problems.

If DSRepair finds errors or problems, run it multiple times until no errors are returned.

7 Install OES 2 Linux on the server, but do not install the Novell Cluster Services option in OES Services at this time.

You can use the same server name and IP address that were used on the NetWare server. This is suggested, but not required.

See the *OES 2 SP2: Installation Guide* for more information.

8 Set up and verify SAN connectivity for the Linux node.

Consult your SAN vendor documentation for SAN setup and connectivity instructions.

9 Install Novell Cluster Services and add the node to your existing NetWare 6.5 cluster.

9a Log in to the OES 2 Linux server as the `root` user.

9b In YaST, select *Open Enterprise Server > OES Install and Configuration.*

9c On the Software Selection page under *OES Services*, click *Novell Cluster Services.*

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 other install options, see [Step 5 in Section 3.5.1, “Installing Novell Cluster Services during a OES 2 Linux Installation,” on page 45.](#)

9d Click *Accept* to begin the install, then click *Continue* to accept changed packages.

9e Continue through the installation process until you reach the Novell Open Enterprise Server Configuration page.

9f Reconfigure *LDAP Configuration of Open Enterprise Services* to specify the credentials for the container administrator user (or non-administrator user) who has the eDirectory rights needed to install Novell Cluster Services.

For information about what rights are needed, see [Section 3.4, “Assigning Install Rights for Container Administrators,” on page 44.](#)

9f1 On the Novell 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.*

9f2 Click the *LDAP Configuration of Open Enterprise Services* link to open the LDAP Configuration page.

9f3 Specify the following values:

- ♦ **Admin name and context:** The username and context (in LDAP form) of the container administrator user (or non-administrator user) who has the eDirectory rights needed to install Novell Cluster Services.
- ♦ **Admin password:** The password of the container administrator (or non-administrator user).

9f4 Click *Next*.

The install returns to the Novell Open Enterprise Server Configuration page.

9g On the Novell Open Enterprise Server Configuration page under *Novell Cluster Services*, click the *disabled* link to enable configuration.

The sentence changes to *Configuration is enabled*.

9h Click the *Novell Cluster Services* link to open the Novell Cluster Services Configuration page.

9i Click *Existing Cluster*, specify the fully distinguished name (FDN) of the cluster, then click *Next*.

IMPORTANT: Use the comma format illustrated in the example. Do not use dots.

This is the name and eDirectory context of the cluster that you are adding this server to.

9j Select the IP address that Novell Cluster Services will use for this node.

Some servers have multiple IP addresses. This step lets you choose which IP address Novell Cluster Services uses.

9k Deselect *Start Services Now*.

9l Click *Next*, then continue through the rest of the OES installation.

9m After the install is complete, use the Software Updater (or other update methods) to install any patches from the OES 2 Linux patch channel and any EVMS patches from the SUSE Linux Enterprise Server 10 SP2 patch channel.

10 If you have a shared disk system on the cluster, enter `sbdutil -f` at the Linux terminal console to verify that the node can see the cluster (SBD) partition on the SAN.

`sbdutil -f` also tells you the device on the SAN where the SBD partition is located.

11 Reboot the server.

12 (Optional) Manually migrate the resources that were on the old server nodes to this Linux server.

Some cluster resources for services on NetWare cannot be used on Linux. For information, see [Section 6.3, “Guidelines for Converting Service Cluster Resources from NetWare to Linux,”](#) on page 68.

The resources can automatically fail back if all of the following apply:

- ♦ The failback mode for the resources was set to Auto.
- ♦ You used the same node number for this Linux server that was used for the former NetWare server.
This only applies if this Linux server is the next server added to the cluster.
- ♦ This Linux server is the preferred node for the resources.

6.5 Adding New OES 2 Linux Nodes to Your NetWare Cluster

You can add new OES 2 Linux cluster nodes to your existing NetWare 6.5 cluster without bringing down the cluster.

- 1** Before you add the first Linux node to the NetWare cluster, if the NetWare cluster uses an SBD, make sure that the device (or devices) being used by the SBD are marked as *Shareable for Clustering*.

You can use NSSMU or iManager to mark the SBD devices as shareable. It is not necessary to bring the cluster down when changing the device attribute to *Shareable for Clustering*.

Using NSSMU:

- 1a** Log in to the master node of the NetWare cluster as the administrator user.
 - 1b** Enter `nssmu` at the server console prompt.
 - 1c** In the NSSMU main menu, select *Devices*.
 - 1d** In the *Devices* list, highlight the device that contains the SBD partition, then press F5 to select it.
 - 1e** Press F6 to mark the device as *Shareable for Clustering*.
 - 1f** If the SBD partition is mirrored, repeat [Step 1d](#) and [Step 1e](#) to also mark the mirror device as *Shareable for Clustering*.
 - 1g** Press Esc to exit NSSMU.
- 2** Install OES 2 Linux on the new node, but do not install the Novell Cluster Services option from OES Services at this time.
See the “[OES 2 SP2: Installation Guide](#)” for more information.
 - 3** Set up and verify SAN connectivity for the new OES 2 Linux node.
Consult your SAN vendor documentation for SAN setup and connectivity instructions.
 - 4** Install Cluster Services and add the new node to your existing NetWare 6.5 cluster.
See [Section 3.5.2, “Installing Novell Cluster Services on an Existing OES 2 Linux Server,”](#) on [page 47](#) for more information.
 - 5** If you have a shared disk system on the cluster, enter `sbdutil -f` at the Linux terminal console to verify that the node can see the cluster (SBD) partition on the SAN.
`sbdutil -f` will also tell you the device on the SAN where the SBD partition is located.
 - 6** Start cluster software by going to the `/etc/init.d` directory and running `novell-ncs start`.
You must be logged in as root to run `novell-ncs start`.
 - 7** Add and assign cluster resources to the new Linux cluster node.
For information, see [Section 9.9, “Assigning Nodes to a Resource,”](#) on [page 133](#).

6.6 Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters

Novell Cluster Services includes specialized functionality to help NetWare and Linux servers coexist in the same cluster. This functionality is also beneficial as you migrate NetWare cluster servers to Linux. The translation between NetWare and Linux versions of the load and unload scripts is performed by the Cluster Translation Library script (`/opt/novell/ncs/bin/clstrlib.py`). It automates the conversion of the Master IP Address resource and cluster-enabled NSS pool resource load and unload scripts from NetWare to Linux.

The NetWare load and unload scripts are read from eDirectory, converted, and written into Linux load and unload script files. Those Linux load and unload script files are then searched for NetWare-specific command strings, and the command strings are then either deleted or replaced with Linux-specific command strings. Separate Linux-specific commands are also added, and the order of certain lines in the scripts is also changed to function with Linux.

This section compares NetWare commands in cluster scripts to their corresponding Linux commands that are used in Cluster Translation Library. If the commands in cluster resource's load or unload scripts are not part of the translation library, the cluster resource can end up in a comatose state.

IMPORTANT: Beginning in OES 2 SP2, Novell Cluster Services allows you to customize the translation syntax that is used for load and unload scripts in mixed-platform situations by defining new syntax translations to be used in addition to the normal translations. For information, see [Section 6.7, “Customizing the Translation Syntax for Converting Load and Unload Scripts,”](#) on page 91.

Unlike NetWare cluster load and unload scripts that are stored in eDirectory, the Linux cluster load and unload scripts are stored in files on Linux cluster servers. The cluster resource name is used in the load and unload script filenames. The path to the files is `/var/opt/novell/ncs/`.

IMPORTANT: Use the *Properties > Scripts* page in the Clusters plug-in in iManager whenever you make manual changes to the load and unload scripts. The changes are automatically saved to the files.

The normal translations performed by the Cluster Translation Library are described in the following sections:

- ◆ [Section 6.6.1, “Comparing Script Commands for NetWare and Linux,”](#) on page 86
- ◆ [Section 6.6.2, “Comparing Master IP Address Scripts,”](#) on page 87
- ◆ [Section 6.6.3, “Comparing NSS Pool Resource Scripts,”](#) on page 88
- ◆ [Section 6.6.4, “Comparing File Access Protocol Resource Script Commands,”](#) on page 89

6.6.1 Comparing Script Commands for NetWare and Linux

[Table 6-2](#) identifies some of the NetWare cluster load and unload script commands that are searched for and the Linux commands that they are replaced with (unless the commands are deleted).

Table 6-2 Cluster Script Command Comparison

Action	NetWare Cluster Command	Linux Cluster Command
Replace	IGNORE_ERROR add secondary ipaddress	ignore_error add_secondary_ipaddress
Replace	IGNORE_ERROR del secondary ipaddress	ignore_error del_secondary_ipaddress
Replace	del secondary ipaddress	ignore_error del_secondary_ipaddress
Replace	add secondary ipaddress	exit_on_error add_secondary_ipaddress
Delete	IGNORE_ERROR NUDP	(deletes the entire line)
Delete	IGNORE_ERROR HTTP	(deletes the entire line)
Replace	nss /poolactivate=	nss /poolact=
Replace	nss /pooldeactivate=	nss /pooldeact=
Replace	mount volume_name VOLID=number	exit_on_error ncpcon mount volume_name=number
Replace	NUDP ADD clusterservername ipaddress	exit_on_error ncpcon bind --ncpsservername=ncpsservername --ipaddress=ipaddress
Replace	NUDP DEL clusterservername ipaddress	ignore_error ncpcon unbind --ncpsservername=ncpsservername --ipaddress=ipaddress
Delete	CLUSTER CVSBIND	(deletes the entire line)
Delete	CIFS	(deletes the entire line)

6.6.2 Comparing Master IP Address Scripts

- ♦ [“Master IP Address Resource Load Script” on page 87](#)
- ♦ [“Master IP Address Resource Unload Script” on page 88](#)

Master IP Address Resource Load Script

This section provides examples of the master IP address resource load scripts on NetWare and Linux.

NetWare

```
IGNORE_ERROR set allow ip address duplicates = on
IGNORE_ERROR CLUSTER CVSBIND ADD BCCP_Cluster 10.1.1.175
IGNORE_ERROR NUDP ADD BCCP_Cluster 10.1.1.175
IGNORE_ERROR add secondary ipaddress 10.1.1.175
IGNORE_ERROR HTTPBIND 10.1.1.175 /KEYFILE:"SSL CertificateIP"
IGNORE_ERROR set allow ip address duplicates = off
```

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

ignore_error add_secondary_ipaddress 10.1.1.175 -np

exit 0
```

Master IP Address Resource Unload Script

This section provides examples of the master IP address resource unload scripts on NetWare and Linux.

NetWare

```
IGNORE_ERROR HTTPUNBIND 10.1.1.175
IGNORE_ERROR del secondary_ipaddress 10.1.1.175
IGNORE_ERROR NUDP DEL BCCP_Cluster 10.1.1.175
IGNORE_ERROR CLUSTER CVSBIND DEL BCCP_Cluster 10.1.1.175
```

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
ignore_error del_secondary_ipaddress 10.1.1.175
exit 0
```

6.6.3 Comparing NSS Pool Resource Scripts

- ♦ [“NSS Pool Resource Load Script” on page 88](#)
- ♦ [“NSS Pool Resource Unload Script” on page 89](#)

NSS Pool Resource Load Script

This section provides examples of the NSS pool resource load scripts on NetWare and Linux.

NetWare

```
nss /poolactivate=HOMES_POOL
mount HOMES VOLID=254
CLUSTER CVSBIND ADD BCC_CLUSTER_HOMES_SERVER 10.1.1.180
NUDP ADD BCC_CLUSTER_HOMES_SERVER 10.1.1.180
add secondary_ipaddress 10.1.1.180
CIFS ADD .CN=BCC_CLUSTER_HOMES_SERVER.OU=servers.O=lab.T=TEST_TREE.
```

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
exit_on_error nss /poolact=HOMES_POOL
exit_on_error ncpcon mount HOMES=254
exit_on_error add_secondary_ipaddress 10.1.1.180

exit_on_error ncpcon bind --ncpservername=BCC_CLUSTER_HOMES_SERVER
--ipaddress=10.1.1.180

exit 0
```

NSS Pool Resource Unload Script

This section provides examples of the NSS pool resource unload scripts on NetWare and Linux.

NetWare

```
del secondary ipaddress 10.1.1.180
CLUSTER CVSBIND DEL BCC_CLUSTER_HOMES_SERVER 10.1.1.180
NUDP DEL BCC_CLUSTER_HOMES_SERVER 10.1.1.180
nss /pooldeactivate=HOMES_POOL /overridetype=question
CIFS DEL .CN=BCC_CLUSTER_HOMES_SERVER.OU=servers.O=lab.T=TEST_TREE.
```

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

ignore_error ncpcon unbind --ncpservername=BCC_CLUSTER_HOMES_SERVER
--ipaddress=10.1.1.180

ignore_error del_secondary_ipaddress 10.1.1.180
ignore_error nss /pooldeact=HOMES_POOL
exit 0
```

6.6.4 Comparing File Access Protocol Resource Script Commands

- ♦ [“File Access Protocol Resource Load Scripts” on page 89](#)
- ♦ [“File Access Protocol Resource Unload Scripts” on page 90](#)

File Access Protocol Resource Load Scripts

This section provides examples of the file access protocol commands for load scripts on NetWare and Linux.

NetWare

Protocol	Script Command for Load Scripts
NCP	NUDP ADD NCS1_P1_SERVER 10.10.10.194
Novell AFP	AFPBIND ADD NCS1_P1_SERVER 10.10.10.204

Protocol	Script Command for Load Scripts
Novell CIFS	CIFS ADD .CN=NCS1_P1_SERVER.O=novell.T=CLUSTER.

Linux

Protocol	Script Command for Load Scripts
NCP	# mount the NCP volume exit_on_error ncpcon mount \$NCP_VOLUME=VOL_ID,PATH=\$MOUNT_POINT exit_on_error ncpcon bind --ncpservername=NCS1_P1_SERVER --ipaddress=10.10.10.194
Novell AFP	exit_on_error cluster_afp.sh add NCS1_P1_SERVER 10.10.10.204
Novell CIFS	exit_on_error novcifs --add --vserver=.CN=NCS1_P1_SERVER.O=novell.T=TREE-188. --ip- addr=\$CIFS_IP

File Access Protocol Resource Unload Scripts

This section provides examples of the Novell AFP commands for unload scripts on NetWare and Linux.

NetWare

Protocol	Script Command for Unload Scripts
NCP	NUDP DEL NCS1_P1_SERVER 10.10.10.194
Novell AFP	AFPBIND DEL NCS1_P1_SERVER 10.10.10.204
Novell CIFS	CIFS DEL .CN=NCS1_P1_SERVER.O=novell.T=CLUSTER.

Linux

Protocol	Script Command for Unload Scripts
NCP	ignore_error ncpcon unbind --ncpservername=NCS1_P1_SERVER --ipaddress=10.10.10.194 # dismount the NCP volume ignore_error ncpcon dismount \$NCP_VOLUME
Novell AFP	ignore_error cluster_afp.sh del NCS1_P1_SERVER 10.10.10.204
Novell CIFS	ignore_error novcifs --remove --vserver=.CN=NCS1_P1_SERVER.O=novell.T=TREE-188. --ip-addr=\$CIFS_IP ignore_error nss /pooldeact=OESPOOL /overridetype=question

6.7 Customizing the Translation Syntax for Converting Load and Unload Scripts

The syntax for load and unload scripts differs for NetWare and Linux platforms. A script that is valid for the NetWare platform is not necessarily recognized on the OES 2 Linux platform. In a mixed-platform cluster, a cluster resource's load script and unload script must be translated to use the proper syntax when running on the NetWare or Linux nodes. Translation occurs in-memory while the cluster contains mixed-platform nodes, and during the final cluster conversion of the cluster from NetWare to Linux.

The translation between NetWare and Linux versions of the load and unload scripts is performed by the Cluster Translation Library script (`/opt/novell/ncs/bin/clstrlib.py`). The normal translations in the library are described in [Section 6.6, "Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters,"](#) on page 86. If the commands in cluster resource's load or unload scripts are not part of the translation library, the cluster resource can end up in a comatose state.

Beginning in OES 2 SP2, Novell Cluster Services allows you to customize the translation syntax that is used for load and unload scripts in mixed-platform situations by defining new syntax translations in the `/var/opt/novell/ncs/customized_translation_syntax` file that you create. The `clstrlib.py` script reads the additional translation syntax from the syntax file, and processes them in addition to the normal translations in the Cluster Translation Library.

The customized translation supports using Python regular expressions to search for strings (`(\S+)`), digits (`(\d+)`), and other data types. The search is case insensitive.

NOTE: Refer to information about Python regular expressions to learn how to create searches for other data types.

In a text editor, create the `customized_translation_syntax` file with the additional translation syntax that you need, then copy the file to the `/var/opt/novell/ncs/` directory on each Linux node in the mixed-platform cluster.

The syntax file should contain a four-line command for each type of translation you want to add:

```
<R|D>
search_string
[replacement_data]
[preceding_data]
```

You can have any number of the four-line commands in the file. Use the following guidelines for creating the syntax translation commands:

Line	Description
<R D>	Specify whether to replace (R) all matches or to delete (D) all matches of the data type you are looking for in the load or unload script.
<code>search_string</code>	Specify the search string that is used to locate a line in the scripts.
<code>[replacement_data]</code>	Specify the replacement data used to replace a line matched by the search performed.
	Leave this line empty if there is no replacement.

Line	Description
[<i>preceding_data</i>]	Specify a line to be inserted before the first line that is matched by the search performed.
	Leave this line empty if there is no line to be inserted before the first matching line.

The following four lines are sample code for a search command in the `customized_translation_syntax` file. The fourth line is intentionally left empty.

```
R
^\s*bind\s+IP\s+(\S+)\s+(\S+)\s+address=(\d+\.\d+\.\d+\.\d+)
ignore_error bind IP \1 \2 address=\3\nexit_on_error ip addr add \3/32 dev \1
```

You can use the `cluster convert preview` command to verify that the `customized_translation_syntax` file is working as intended for a particular resource.

- 1 On the master node, open a terminal console as the `root` user, then enter

```
cluster convert preview resource_name
```

6.8 Finalizing the Cluster Conversion

If you have converted all nodes in a former NetWare cluster to OES 2 Linux, you must finalize the conversion process by issuing the `cluster convert` command on one Linux cluster node. The `cluster convert` command moves cluster resource load and unload scripts from the files where they were stored on Linux cluster nodes into eDirectory. This enables a Linux cluster that has been converted from NetWare to utilize eDirectory like the former NetWare cluster.

WARNING: After you finalize the cluster conversion, rollback to NetWare is not supported.

To finalize the cluster conversion:

- 1 Run `cluster convert preview resource_name` at the terminal console of one Linux cluster node.

Replace *resource_name* with the name of a resource that you want to preview.

The preview switch lets you view the resource load and unload script changes that will be made when the conversion is finalized. You can preview all cluster resources.

- 2 Run `cluster convert commit` at the terminal console of one Linux cluster node to finalize the conversion.

The `cluster convert commit` command generates or regenerates the cluster resource templates that are included with Novell Cluster Services for Linux. In addition to generating Linux cluster resource templates, this command deletes all NetWare cluster resource templates that have the same name as Linux cluster resource templates.

The cluster resource templates are automatically created when you create a new Linux cluster, but are not created when you convert an existing NetWare cluster to Linux.

- 3 Update the cluster configuration on all nodes by running the cluster configuration daemon. Enter the following command as the `root` user on every node in the cluster:

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

This removes the NetWare nodes from the list of nodes in the cluster so they are not displayed in iManager.

Configuring Cluster Policies and Priorities

7

After installing Novell 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 7.1, “Understanding Cluster Settings,” on page 95](#)
- ◆ [Section 7.2, “Configuring Quorum Membership and Timeout Properties,” on page 96](#)
- ◆ [Section 7.3, “Configuring Cluster Protocol Properties,” on page 97](#)
- ◆ [Section 7.4, “Configuring Cluster Event E-Mail Notification,” on page 99](#)
- ◆ [Section 7.5, “Viewing the Cluster Node Properties,” on page 100](#)
- ◆ [Section 7.6, “Modifying the Cluster IP Address and Port Properties,” on page 100](#)
- ◆ [Section 7.7, “What’s Next,” on page 100](#)

7.1 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 7.1.1, “Cluster Policies,” on page 95](#)
- ◆ [Section 7.1.2, “Cluster Protocols Properties,” on page 96](#)

7.1.1 Cluster Policies

[Table 7-1](#) describes the configurable cluster policies. You can manage cluster policies in iManager by going to the *Clusters > Cluster Options > Policies* page. For instructions, see [Section 7.2, “Configuring Quorum Membership and Timeout Properties,” on page 96](#).

Table 7-1 Cluster Policies

Property	Description
Cluster IP address	Specifies the IP address for the cluster. You specify the IP address when you install Novell Cluster Services on the first node of the cluster. Rarely, you might need to modify this value.
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.

Property	Description
Quorum membership	Specifies number of nodes that must be up and running in the cluster in order for cluster resources to begin loading. Specify a value between 1 and the number of nodes.
Quorum timeout	Specifies the maximum amount of time in seconds to wait for the specified quorum to be met before cluster resources begin loading on whatever number of nodes are actually up and running.
E-mail notification	Enables or disables e-mail notification for the cluster. If it is enabled, you can specify up to eight administrator e-mail addresses for cluster events notification.

7.1.2 Cluster Protocols Properties

Table 7-2 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 > Cluster Options > Protocols* page. For instructions, see [Section 7.3, “Configuring Cluster Protocol Properties,”](#) on page 97.

Table 7-2 Cluster Policies

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.
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.
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.
Maximum retransmits	This value is set by default and should not be changed.

7.2 Configuring Quorum Membership and Timeout Properties

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

- 1 In iManager, select *Clusters*, then select *Cluster Options*.
- 2 Specify the cluster name, or browse and select the Cluster object.
- 3 Click the *Properties* button under the cluster name.
- 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.
For information, see [Section 7.2.1, “Quorum Triggers \(Number of Nodes\),” on page 97](#).
- 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.
For information, see [Section 7.2.2, “Quorum Triggers \(Timeout\),” on page 97](#).
- 7 Click *Apply* or *OK* to save your changes.

7.2.1 Quorum Triggers (Number of Nodes)

The number of nodes required to form a cluster quorum is the number of nodes that must be running in the cluster before resources start to load. When you first bring up servers in your cluster, Novell 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 don't 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.

7.2.2 Quorum Triggers (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.

7.3 Configuring Cluster Protocol Properties

You can use the Cluster Protocol property pages to view or edit 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*, then select *Cluster Options*.
- 2 Specify the cluster name, or browse and select the Cluster object.
- 3 Click the *Properties* button under the cluster name.
- 4 Click the *Protocols* tab.

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.

For information, 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.

7.3.1 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.

7.3.2 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.

7.3.3 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.

If you are using multipath I/O to manage multiple paths between the server and the shared drive, make sure that you allow sufficient time in the watchdog setting for a path failover to avoid unnecessary cluster resource failovers between nodes. Test the failover time of the MPIO solution you are using, then adjust the watchdog setting upward accordingly.

7.3.4 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.

If you are using multipath I/O to manage multiple paths between the server and the shared drive, make sure that you allow sufficient time in the watchdog setting for a path failover to avoid unnecessary cluster resource failovers between nodes. Test the failover time of the MPIO solution you are using, then adjust the watchdog setting upward accordingly.

7.3.5 Maximum Retransmits

This value is set by default, and should not be changed.

7.4 Configuring Cluster Event E-Mail Notification

Novell Cluster Services can automatically send out e-mail 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
```

IMPORTANT: Novell Cluster Services uses Postfix to send e-mail alerts. If you have a cluster resource that uses SMTP, that resource might not work in the cluster unless you change the Postfix configuration.

For example, GroupWise uses SMTP and will not function as a cluster resource if Postfix uses the same port, which it does by default. In this case, Postfix must be configured to use a different port. You can do this by editing the `/etc/postfix/main.cf` file and changing the values for the `inet_interfaces`, `mydestination`, and `mynetworks_style` lines.

You also need to change the listen port for the `smtpd` process in the `/etc/postfix/master.cf` file. See [the Postfix Web site \(http://www.postfix.org\)](http://www.postfix.org) for more information on configuring Postfix.

You can enable or disable e-mail notification for the cluster and specify up to eight administrator e-mail addresses for cluster notification.

- 1 In iManager, select *Clusters*, then select *Cluster Options*.
- 2 Specify the cluster name, or browse and select the Cluster object.
- 3 Click the *Properties* button under the cluster name.
- 4 Click the *Policies* tab.
- 5 Under *Notification*, select or deselect the *Enable Cluster Notification Events* check box to enable or disable e-mail notification.
- 6 If you enable e-mail notification, add up to eight e-mail addresses in the field provided.
You can click the buttons next to the field to add, delete, or edit e-mail addresses. Repeat this process for each e-mail address you want on the notification list.
- 7 If you enable e-mail notification, 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.
- 8 If you enable e-mail notification, 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.
- 9 Click *Apply* or *OK* to save changes.

7.5 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*, then select *Cluster Options*.
- 2 Specify the cluster name, or browse and select the Cluster object.
- 3 Select the check box next to the cluster node whose properties you want to view, then click the *Details* link.
- 4 View the desired information, then click *OK*.

7.6 Modifying the Cluster IP Address and Port Properties

The cluster IP address is assigned when you install Novell Cluster Services. The cluster IP address normally does not need to be changed, but it can be if needed.

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.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Specify the cluster name, or browse and select the Cluster object.
- 3 Click the *Properties* button under the cluster name.
- 4 Click the *Policies* tab.
- 5 Specify the new value for the IP address or port.
- 6 Click *Apply* or *OK* to save your changes.

7.7 What's Next

After installing and configuring the cluster, you should configure cluster resources for it. For information, see the following:

- ♦ [Chapter 9, “Configuring and Managing Cluster Resources,” on page 123](#)
- ♦ [Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139](#)
- ♦ [Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163](#)

For information about managing the cluster, see [Chapter 8, “Managing Clusters,” on page 101](#).

Managing Clusters

8

After you have installed, set up, and configured Novell Cluster Services for your specific needs 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 Novell Cluster Services,”](#) on page 209.

- ◆ [Section 8.1, “Starting and Stopping Novell Cluster Services,”](#) on page 101
- ◆ [Section 8.2, “Monitoring Cluster and Resource States,”](#) on page 103
- ◆ [Section 8.3, “Generating a Cluster Configuration Report,”](#) on page 105
- ◆ [Section 8.4, “Cluster Migrating Resources to Different Nodes,”](#) on page 106
- ◆ [Section 8.5, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 106
- ◆ [Section 8.6, “Removing \(Leaving\) a Node from the Cluster,”](#) on page 107
- ◆ [Section 8.7, “Joining a Node to the Cluster,”](#) on page 107
- ◆ [Section 8.8, “Configuring the EVMS Remote Request Timeout,”](#) on page 107
- ◆ [Section 8.9, “Shutting Down Linux Cluster Servers When Servicing Shared Storage,”](#) on page 108
- ◆ [Section 8.10, “Enabling or Disabling Cluster Maintenance Mode,”](#) on page 108
- ◆ [Section 8.11, “Preventing a Cluster Node Reboot after a Node Shutdown,”](#) on page 109
- ◆ [Section 8.12, “Renaming a Pool for a Pool Cluster Resource,”](#) on page 110
- ◆ [Section 8.13, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110
- ◆ [Section 8.14, “Adding a Node That Was Previously in the Cluster,”](#) on page 114
- ◆ [Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,”](#) on page 114
- ◆ [Section 8.16, “Creating or Deleting Cluster SBD Partitions,”](#) on page 115
- ◆ [Section 8.17, “Customizing Cluster Services Management,”](#) on page 120

8.1 Starting and Stopping Novell Cluster Services

Novell Cluster Services automatically starts after it is installed. Novell Cluster Services also automatically starts when you reboot your Novell Open Enterprise Server (OES) 2 Linux 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 Novell Cluster Services. You can do this by entering the following at the Linux terminal console:

8.1.1 Starting Novell Cluster Services

If you stop Novell Cluster Services, you can restart it by doing the following:

- 1 Open a terminal console, then log in as the `root` user.
- 2 Use one of the following methods to start Novell Cluster Services:
 - ♦ At the terminal console prompt, go to the `/etc/init.d` directory and enter
`./novell-ncs start`
 - ♦ At the terminal console prompt, enter
`rcnovell-ncs start`

8.1.2 Stopping Novell Cluster Services

- 1 Open a terminal console, then log in as the `root` user.
- 2 Use one of the following methods to start Novell Cluster Services:
 - ♦ Go to the `/etc/init.d` directory and enter
`./novell-ncs stop`
 - ♦ At the terminal prompt, enter
`rcnovell-ncs stop`

8.1.3 Enabling and Disabling the Automatic Start of Novell Cluster Services

Novell Cluster Services automatically starts by default after it is installed and on server reboot.

To cause Novell 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:
`chkconfig novell-ncs off`
- 3 Reboot the server.
- 4 After rebooting, you must manually start Novell Cluster Services by entering
`rcnovell-ncs start`


To cause Novell 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:
`chkconfig novell-ncs on`
- 3 Reboot the server.

8.2 Monitoring Cluster and Resource States







The *Cluster Manager* link in iManager gives you important information about the status of servers and resources in your cluster.

- 1 In iManager, click *Cluster*, then click *Cluster Manager*.
- 2 Type the name of the desired cluster, or browse to locate and select the Cluster object.
A list of resources and resource states displays.

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.

Cluster servers and resources display the following icons for the different operating states:

Table 8-1 Cluster Operating States

State	Icon	Description
Normal		A green ball indicates that the server or resource is online or running.
Stopped		A red ball with a horizontal white line indicates that the node is stopped.
Offline		A white ball with a horizontal red line indicates that the node is offline.
Critical		A white ball with a red X indicates that the node has failed or is comatose.
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 is unknown.

When a resource is red, it is waiting for administrator intervention. When a resource is gray with no break in the icon, either that server is not currently a member of the cluster or its state is unknown. When a resource is blank or has no colored icon, it is unassigned, offline, changing state, or in the process of loading or unloading.

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.

[Table 8-2](#) identifies the different resource states and gives descriptions and possible actions for each state.

Table 8-2 Cluster Resource States

Resource State	Description	Possible Actions
Alert	Either the Start, Failover, or Failback mode for the resource has been set to Manual. The resource is waiting to start, fail over, or fail back on the specified server.	<p>Click the <i>Alert</i> status indicator. Depending on the resource state, you are prompted to start, fail over, or fail back the resource.</p> <p>If you attempt to offline a resource that is in the <i>Start Alert</i> state, nothing happens. You must clear the <i>Start Alert</i> before you can offline the resource.</p> <p>If you attempt to online a resource that is in the <i>Start Alert</i> state, you get the following warning:</p> <p>This operation cannot be completed. It is only available when the resource is in an offline state.</p> <p>When the resource is in state <i>Start Alert</i>, you must clear the alert before you can offline the resource. After the resource is offline, you can online the resource.</p>
Comatose	The resource is not running properly and requires administrator intervention.	Click the <i>Comatose</i> status indicator and bring the resource offline. After resource problems have been resolved, the resource can be brought back online (returned to the running 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 Novell eDirectory.	None.
Offline	Offline status indicates the resource is shut down or is in a dormant or inactive state.	Click the <i>Offline</i> status indicator and, if desired, click the <i>Online</i> 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.	Click the <i>Running</i> status indicator and choose to either migrate the resource to a different server in your cluster or unload (bring offline) the resource.
Unassigned	There isn't an assigned node available that the resource can be loaded on.	Click the <i>Unassigned</i> status indicator and, if desired, 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.

8.3 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 (<i>clustername.context</i>).
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 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 > Cluster Manager*.
- 2** Browse to select the cluster you want to manage.
The *Run Report* button is dimmed until the page is refreshed with the cluster information.
- 3** Click *Run Report*.
- 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 file.
 - 6a** Click the cursor in the browser frame 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.

8.4 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.

Migrating resources lets you balance the load and evenly distribute applications among the servers in your cluster.

Using iManager

- 1 In iManager, click *Clusters*, then click *Cluster Manager*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Select the check box next to the resource you want to migrate and 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 console command prompt of a cluster node:

```
cluster migrate resource_name node_name
```

8.5 Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node

You can start or stop a resource from running on a node by onlining or offlining the resource.

Offlining the resource unloads from the server. It cannot be loaded on any other servers in the cluster and remains unloaded until you load it again. If you edit a resource that is online, the changes you made do not take effect until the resource is taken offline and brought online again.

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. If monitoring is enabled 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.

There are maintenance situations where you might need to control services 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.

Using iManager

- 1 In iManager, select *Clusters*, then select *Cluster Manager*.

- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
A list of resources and resource states displays.
- 3 Select the check box next to the resource you want to manage, then click *Online* or click *Offline*.

Using Console Commands

To online a specified resource on the most preferred node that is currently active:

- 1 As the `root` user, enter the following at the console command prompt of a cluster node:
`cluster online resource_name`

To online a specified resource to a specific active node:

- 1 As the `root` user, enter the following at the console command prompt of a cluster node:
`cluster online resource_name node_name`

To unload a specified resource:

- 1 As the `root` user, enter the following at the console command prompt of a cluster node:
`cluster offline resource_name`

8.6 Removing (Leaving) a Node from the Cluster

You can remove a node from a cluster so that the node is not visible to other servers in the cluster by leaving the 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 terminal console prompt:

```
cluster leave
```

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

```
No longer a member of cluster clustername
```

8.7 Joining a Node to the Cluster

You can add the node back to the cluster so that the node is again visible to other servers in the cluster by joining the cluster.

- 1 Log in as the `root` user to the server that you want to join the cluster, then enter the following at a terminal console prompt:

```
cluster join
```

8.8 Configuring the EVMS Remote Request Timeout

Novell Cluster Services for Linux uses the EVMS Cluster Segment Manager for shared devices. When load and unload scripts run, EVMS mounts or dismounts shared devices. The default timeout for remote requests for EVMS is 12 seconds. If you have many shared devices to be mounted or dismounted and the timeout is exceeded, you might see EVMS locking errors and some cluster resources can go comatose.

You must increase the EVMS engine's `remote_request_timeout` value if you have many shared devices and it is taking longer than 12 seconds to process them as the load or unload scripts are run for a cluster migration. The time needed varies with the number of shared devices that are being processed. For example, for about 60 shared devices, set the timeout to 30 seconds.

To modify the EVMS timeout:

- 1 Open the `/etc/evms.conf` file in a text editor.
- 2 Edit the following section to increase the `remote_request_timeout` value:

```
engine {  
  
    # Timeout in seconds when waiting for a response from a remote node  
    remote_request_timeout = 12  
  
}
```

- 3 Save the file.
- 4 At the terminal console prompt, enter

```
evms_activate
```

8.9 Shutting Down Linux Cluster Servers When Servicing Shared Storage

If you need to power down or recycle your shared storage system, you should shut down Linux Cluster Servers prior to doing so.

8.10 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 terminal console 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 terminal console prompt:

```
cluster maintenance off
```

8.11 Preventing a Cluster Node Reboot after a Node Shutdown

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. By default, cluster nodes are configured to reboot after an automatic shutdown.

On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems.

- ♦ [Section 8.11.1, “OES 2 SP2 with Patches and Later,” on page 109](#)
- ♦ [Section 8.11.2, “OES 2 SP2 Release Version and Earlier,” on page 109](#)

8.11.1 OES 2 SP2 with Patches and Later

Beginning in the OES 2 SP2 Maintenance Patch for May 2010, the Novell Cluster Services reboot behavior conforms to the kernel panic setting for the Linux operating system.

- 1 As the `root` user, open the `/etc/sysctl.conf` file in a text editor.
- 2 Set the `kernel.panic` token to value to 0 or 1, 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 in the `/etc/sysctl.conf` file. This allows the administrator to determine what caused the kernel panic condition before manually rebooting the server. This is the recommended setting.
 - ♦ **Reboot:** To allow a cluster reboot after a node shutdown, set the `kernel.panic` token to value to 1 in the `/etc/sysctl.conf` file.

For information, see the Linux man pages on `sysctl` and `sysctl.conf`.

- 3 Save your changes.

8.11.2 OES 2 SP2 Release Version and Earlier

In OES 2 SP release version and earlier, you can modify the `opt/novell/ncs/bin/ldncs` file for the cluster to trigger the server to not automatically reboot after a shutdown.

- 1 Open the `opt/novell/ncs/bin/ldncs` file in a text editor.
- 2 Find the following line:

```
echo -n $TOLERANCE > /proc/sys/kernel/panic
```
- 3 Replace `$TOLERANCE` with a value of 0 to cause the server to not automatically reboot after a shutdown.
- 4 After editing the `ldncs` file, you must reboot the server to cause the change to take effect.

8.12 Renaming a Pool for a Pool Cluster Resource

When a cluster-enabled NSS pool is used as a Novell Cluster Services cluster resource, special steps must be followed if you decide to rename the pool.

- 1 Cluster migrate the pool resource to the node in the cluster where the pool was initially created.
- 2 Offline the pool cluster resource.
- 3 Activate the pool by using iManager or NSSMU.
- 4 Rename the pool.

After the rename, the pool is in a deactive state.

- 5 Online the pool cluster resource.

Novell 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.

8.13 Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster

Use the instructions in this section to change the IP addresses of the cluster, information about the LDAP server 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 110](#)
- ♦ [Section 8.13.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 111](#)

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. You must modify this cluster configuration information in the following cases:

- ♦ Changing the Administrator username and password for the cluster
- ♦ Changing the password for the existing Administrator username
- ♦ Changing the IP address information about the existing LDAP server
- ♦ Assigning a different LDAP server for the cluster to use
- ♦ Adding LDAP servers to the list of ones that the cluster can use

You can modify these settings at any time. Novell Cluster Services can be running or not running.

To modify the LDAP server IP address or administrator credentials in the Novell Cluster Services configuration settings:

- 1 Make sure that the IP addresses and administrator username that you plan to use meet the requirements specified in [Section 3.1.3, “Configuration Requirements,” on page 36](#).

- 2 Log in as the `root` user to the master node of the cluster.
- 3 In a text editor, create a text file, specify the configuration information for the Novell Cluster Services cluster in it, then save the file.

Two examples are shown below of the content of the file with sample values. The directives are self-explanatory.

IMPORTANT: Make sure to 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 Novell Cluster Services cluster when you have a single LDAP server.

```
CONFIG_NCS_CLUSTER_DN="cn=svr1_oes2_cluster.o=context "  
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"
```

If you have multiple LDAP servers, the syntax is slightly different. The following lines are the content of a sample configuration file for a Novell 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"
```

- 4 As the `root` user, enter the following command at a terminal console prompt:

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

Replace *configuration_filename* with the actual name of the file you created.

- 5 Delete the configuration file that you created.
- 6 For each of the other nodes in the cluster, log in as the `root` user for the node, then repeat [Step 3](#) to [Step 5](#).

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 CASA, and CASA does not synchronize across clustered nodes.

- 7 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 Novell Cluster Services cluster, perform the following tasks in the order given:

- ♦ [“Prerequisites” on page 112](#)
- ♦ [“Changing the IP Addresses of Cluster Resources” on page 112](#)
- ♦ [“Changing the IP Addresses of Servers in a Cluster” on page 113](#)
- ♦ [“Modifying the Cluster Configuration Information” on page 113](#)

Prerequisites

Before you begin, make sure that the IP addresses and administrator username that you plan to use meet the requirements specified in [Section 3.1.3, “Configuration Requirements,” on page 36](#).

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** Offline the cluster resources whose IP addresses are changing.
 - 1a** In iManager, click *Clusters*, then click *Cluster Manager*.
 - 1b** Browse to locate and select the Cluster object of the cluster you want to manage.
 - 1c** Select the check boxes next to the resources you want to take offline, then click *Offline*.
- 2** For each cluster resource, including the master IP address resource, modify information as needed in the resource load, unload, and monitor scripts.
 - 2a** In iManager, click *Clusters*, then click *Cluster Options*.
 - 2b** Browse to locate and select the Cluster object of the cluster you want to manage.
 - 2c** Select the check box next to the resource whose load script you want to edit, then click the *Details* link.
 - 2d** Click the *Scripts* tab, then click the *Load Script* link.
 - 2e** Edit the script by replacing variables with actual values for your new configuration, such as the resource IP address and container name where you will be moving the cluster. You might also need to edit the values used in the command lines.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, container name, file system type, and device.

- 2f** Click *Apply* to save the changed script.
 - 2g** Repeat [Step 4](#) through [Step 6](#) by making similar changes to the *Unload Script* and *Monitor Script*.
- 3** Stop Novell Cluster Services for every node in the cluster by entering the following at the terminal console prompt as the `root` user:

```
rcnovell-ncs stop
```
 - 4** Continue with [“Changing the IP Addresses of Servers in a Cluster” on page 113](#).

Changing the IP Addresses of Servers in a Cluster

After preparing the cluster resources for the IP address change and stopping Novell Cluster Services (see “[Changing the IP Addresses of Cluster Resources](#)” on page 112), 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 2 Server’s IP Address](#)” in the *OES 2 SP2: Planning and Implementation Guide*.
- 2 The server IP address changes are not complete until you make those changes known to Novell Cluster Services and eDirectory. Continue with “[Modifying the Cluster Configuration Information](#)” on page 113.

Modifying the Cluster Configuration Information

Before restarting Novell Cluster Services, you must update the cluster configuration information in Novell 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 Novell 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 110.

- 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, click *Directory Administration*, then click *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, click *Directory Administration*, then click *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 Make sure that LDAP server is running before restarting Novell Cluster Services.

IMPORTANT: Novell Cluster Services requires LDAP.

- 5 Make sure that NSS is running if there are NSS cluster resources that you will be onlining.

- 6 Start Novell Cluster Services by entering the following command at a terminal console prompt as the root user:

```
rcnovell-ncs start
```

- 7 Online the cluster resources:

- 7a In iManager, click *Clusters*, then click *Cluster Manager*.

- 7b Browse to locate and select the Cluster object of the cluster you want to manage.

- 7c Select the check boxes next to the resources you want to bring online, then click *Online*.

8.14 Adding a Node That Was Previously in the Cluster

- 1 If your storage array and devices are not configured, configure them before you install Novell Cluster Services.
- 2 If necessary, install OES 2 Linux and Novell Cluster Services, including the latest Service Pack on the server using the same node name and IP address of the node that was previously in the cluster.
- 3 If the Cluster object for the server is still present, delete the object.
For information, see [Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,”](#) on page 114.
- 4 Run the Novell Cluster Services installation by following the procedure outlined in [Section 3.5.2, “Installing Novell Cluster Services on an Existing OES 2 Linux Server,”](#) on page 47.

The node assumes its former identity.

8.15 Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node

In order to permanently remove a cluster node from one cluster and add it to another cluster, you must delete the node from the original cluster before you can reconfigure it for the new 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 terminal console prompt:

```
cluster leave
```

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

```
No longer a member of cluster clustername
```

- 2 In a Web browser, open iManager, then log in to the Novell 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, select *eDirectory Administration*, then browse to select and connect to the eDirectory server in the same tree that contains the master replica.
- 4 In iManager, delete the node’s Cluster Node object from the cluster container:
 - 4a Select *eDirectory Administration > Delete Objects*.

IMPORTANT: If a cluster SBD partition already exists, do not create another one. Delete the existing one first as described in [Section 8.16.2, “Before Creating an Cluster SBD Partition,”](#) on page 116.

8.16.2 Before Creating an Cluster SBD Partition

Before creating an SBD partition, you should make sure that an SBD does not already exist on your cluster.

- 1 As the `root` user, enter the following at the terminal console of a Linux cluster server:

```
sbdutil -f
```

This tells you whether an SBD partition exists and identifies the device on the SAN where the SBD partition is located.
- 2 If the SBD partition already exists, use one of the following methods to delete the existing partition before attempting to create another one:
 - ♦ [Section 8.16.5, “Deleting a Non-Mirrored Cluster SBD Partition,”](#) on page 118
 - ♦ [Section 8.16.7, “Removing a Segment from a Mirrored Cluster SBD Partition,”](#) on page 119

8.16.3 Creating a Non-Mirrored Cluster SBD Partition

If you did not create a cluster partition during the Novell Cluster Services installation on the first node of the cluster, you can create one on 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 doing the following:

- 1 As the `root` user, enter the following at the terminal console of a Linux cluster server:

```
sbdutil -c -n clustername -d device_name -s size
```

For the `-d` option, replace *device_name* with the name of the device where you want to create the cluster partition. Use the EVMSGUI (or EVMSN or EVMS) tool to check the names of the devices if needed, and only use the base (leaf) names (such as `sdb` or `mpathd`) with the `-d` option.

For the `-s` option, replace *size* with the size (in MB) to use for the SBD partition. The default size is 8 MB. The `-s` option is available in OES 2 SP1 and later. For information about size requirements, see [“SBD Partitions”](#) on page 39.

For example, you might enter something similar to the following:

```
sbdutil -c -n mycluster1 -d sdb -s 200
```
- 2 If the cluster has never had an SBD partition, modify the Cluster object in eDirectory to enable its *NCS: Shared Disk Flag* attribute.

This step is required only if Cluster Services was initially installed without an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the *NCS: Shared Disk Flag* attribute is enabled.

 - 2a In a Web browser, open iManager, then log in to the Novell 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.

- 2b** Click *eDirectory Administration*, then click *Modify Object*.
 - 2c** Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.
 - 2d** Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.
 - 2e** Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.
 - 2f** Click *Apply* to save changes.
- 3** Reboot all cluster nodes.

8.16.4 Creating a Mirrored Cluster SBD Partition

If the cluster has a shared disk system, you can achieve a greater level of fault tolerance for the SBD partition by mirroring it. Novell Cluster Services uses the NSS software RAID capability to mirror the SBD partition.

You can mirror the SBD partition when you install Novell Cluster Services on the first node, or you can create it afterwards by using the `sbdutil` utility or the `evmsgui` utility to create a mirrored cluster partition.

- ♦ [“Using SBDUTIL to Create a Mirrored Cluster SBD Partition” on page 117](#)
- ♦ [“Using EVMSGUI to Create a Mirrored Cluster SBD Partition” on page 118](#)

Using SBDUTIL to Create a Mirrored Cluster SBD Partition

- 1** As the `root` user, enter the following at the terminal console of a Linux cluster server:

```
sbdutil -c -n clustername -d device_name -d device_name -s size
```

Replace *device_name* with the name of the devices where you want to create the cluster partition and its mirror. Use the EVMSGUI (or EVMSN or EVMS) tool to check the names of the devices if needed, and only use the base (leaf) names (such as `sdb` or `mpathd`) with the `-d` option.

For the `-s` option, replace *size* with the size (in MB) to use for the SBD partition. The default size is 8 MB. The `-s` option is available in OES 2 SP1 and later. For information about size requirements, see [“SBD Partitions” on page 39](#).

For example, you might enter something similar to either of the following commands:

```
sbdutil -c -n mycluster1 -d sdb -d sdc -s 200
```

```
sbdutil -c -n cl1 -d CX4-LUN000 -d CX4-LUN001 -s 1020
```

- 2** If the cluster has never had an SBD partition, modify the Cluster object in eDirectory to enable its *NCS: Shared Disk Flag* attribute.

This step is required only if Cluster Services was initially installed without an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the *NCS: Shared Disk Flag* attribute is enabled.

- 2a** In a Web browser, open iManager, then log in to the Novell 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.

- 2b** Click *eDirectory Administration*, then click *Modify Object*.
 - 2c** Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.
 - 2d** Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.
 - 2e** Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.
 - 2f** Click *Apply* to save changes.
- 3** Reboot all cluster nodes.

Using EVMSGUI to Create a Mirrored Cluster SBD Partition

- 1** At the Linux terminal console of a cluster server, log in as the `root` user, then enter `evmsgui` to start the EVMS GUI utility.
- 2** In `evmsgui`, create an SBD partition:
 - 2a** Click *Action*, then click *Create*.
 - 2b** Click *Segment*, choose the *NetWare Segment Manager*, then click *Next*.
 - 2c** Select *Free Space Storage Object*, then click *Next*.
 - 2d** Specify 20 MB or larger as the size of the cluster partition, then choose *SBD* as the partition type.

For information about size requirements, see [“SBD Partitions” on page 39](#).
 - 2e** Specify the name of your cluster as the *Label*, then click *Create*.
 - 2f** Click *Save* to save your changes.
- 3** In `evmsgui`, mirror the newly created SBD partition:
 - 3a** Click *Segments*.
 - 3b** Locate the SBD partition and right-click it.
 - 3c** Select *Mirror Segment*, then click *OK*.

This creates a RAID 1 device with the name `/dev/evms/nodes/cluster.sbd`.
 - 3d** Click *Save* to save your changes.
- 4** Exit `evmsgui`.
- 5** Reboot all cluster nodes.

8.16.5 Deleting a Non-Mirrored Cluster SBD Partition

You must delete an existing SBD partition for the cluster before you attempt to create (or re-create) an SBD partition. The existing SBD partition might have been created during the Novell Cluster Services installation, or later by using the `sbdutil`.

- 1** At a Linux terminal console of a cluster server, log in as the `root` user.
- 2** Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.
- 3** Enter `nssmu` to open the NSS Utility, then select *Partitions*.

- 4 Select the SBD partition you want to delete.
- 5 Click *Delete* to delete the partition and its contents, then click *OK* to confirm the deletion.
- 6 If you have more than one node in the cluster, use one of the following methods to create a new SBD partition:
 - ♦ [Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#)
 - ♦ [Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)
- 7 Join the nodes to the cluster.

8.16.6 Deleting a Mirrored Cluster SBD Partition

You must delete an existing SBD partition for the cluster before you attempt to create (or re-create) an SBD partition. Because it is mirrored, you delete the software RAID device instead of deleting the two partitions separately. The existing mirrored SBD partition might have been created during the Novell Cluster Services installation, or later by using the SBDUTIL or the EVMS GUI tool.

- 1 At a Linux terminal console of a cluster server, log in as the `root` user.
- 2 Enter `cluster down` at the terminal console of one cluster server.
This causes all cluster servers to leave the cluster.
- 3 Enter `nssmu` to open the NSS Utility, then select *Software RAIDS*.
- 4 Select the software RAID 1 for the SBD partition you want to delete.
- 5 Click *Delete* to delete the software RAID and its member segments, then click *OK* to confirm the deletion.
- 6 If you have more than one node in the cluster, use one of the following methods to create a new SBD partition:
 - ♦ [Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#)
 - ♦ [Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)
- 7 Join the nodes to the cluster.

8.16.7 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 standalone SBD partition.

IMPORTANT: To get rid of the software RAID definition, you must delete the mirrored SBD partition as described in [Section 8.16.6, “Deleting a Mirrored Cluster SBD Partition,” on page 119](#), then re-create a non-mirrored SBD partition, as described in [Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116](#).

Beginning in OES 2 SP2, you can specify which segment to keep when you use NSSMU to remove a segment from a software RAID 1 (mirror) device.

- 1 At a Linux terminal console of a cluster server, log in as the `root` user.
- 2 Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.

- 3 Enter `nssmu` to open the NSS Utility, then select *Software RAIDS*.
- 4 Select the software RAID1 device for the cluster SBD partition that you want to manage.
- 5 Press Enter to show its member segments.
- 6 Select the member segment you want to delete, then select *Delete* to remove the segment.
The RAID definition will still exist for the remaining segment of the mirrored SBD partition, but it reports that it is not mirrored.
- 7 Join the nodes to the cluster.

8.17 Customizing Cluster Services Management

Some portions of Novell Cluster Services management can be performed and customized 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 8-3](#) lists the cluster-related virtual XML files and gives a brief description of each.

Table 8-3 Cluster-Related Virtual XML Files

Virtual XML Filename	Description
<code>Config.xml</code>	Provides the combined information from <code>ClusterConfig.xml</code> , <code>NodeConfig.xml</code> , <code>ResourceConfig.xml</code> , and <code>PoolConfig.xml</code> .
<code>ClusterConfig.xml</code>	Provides cluster configuration information.
<code>NodeConfig.xml</code>	Provides node configuration information for all nodes in the cluster that were active at the time the cluster was brought up.
<code>NodeState.xml</code>	Provides current information on the state of each node in the cluster (cluster membership).
<code>PoolConfig.xml</code>	Provides cluster-enabled pool and volume configuration information for each pool and volume.
<code>PoolState.xml</code>	Provides current information on the state of each cluster-enabled pool in the cluster.
<code>ResourceConfig.xml</code>	Provides resource configuration information for each resource in the cluster.
<code>ResourceState.xml</code>	Provides current information on the state of each resource in the cluster.
<code>State.xml</code>	Provides the combined information from <code>NodeState.xml</code> , <code>ResourceState.xml</code> , and <code>PoolState.xml</code> .

Table 8-4 lists the cluster-related CMD files and gives a brief description of each.

Table 8-4 *Cluster-Related CMD Files*

CMD Filename	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.

Configuring and Managing Cluster Resources

9

After you create and configure a Novell 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 8, “Managing Clusters,”](#) on page 101.

- ◆ [Section 9.1, “Planning Cluster Maintenance,”](#) on page 123
- ◆ [Section 9.2, “Naming Conventions for Cluster Resources,”](#) on page 123
- ◆ [Section 9.3, “Creating Cluster Resource Templates,”](#) on page 124
- ◆ [Section 9.4, “Creating Cluster Resources,”](#) on page 127
- ◆ [Section 9.5, “Configuring a Load Script for a Cluster Resource,”](#) on page 127
- ◆ [Section 9.6, “Configuring an Unload Script for a Cluster Resource,”](#) on page 128
- ◆ [Section 9.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 129
- ◆ [Section 9.8, “Setting Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 131
- ◆ [Section 9.9, “Assigning Nodes to a Resource,”](#) on page 133
- ◆ [Section 9.10, “Configuring Resource Priorities for Load Order,”](#) on page 134
- ◆ [Section 9.11, “Changing the IP Address of a Cluster Resource,”](#) on page 134
- ◆ [Section 9.12, “Deleting Cluster Resources,”](#) on page 134
- ◆ [Section 9.13, “Additional Information for Creating Cluster Resources,”](#) on page 136

9.1 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.
- ◆ When performing maintenance on a node, cluster migrate existing pools to another node if you want the related users to be undisturbed.

9.2 Naming Conventions for Cluster Resources

Cluster resource names can be up to 63 characters. Novell 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 is automatically used in the pool cluster resource name, do not use special characters !@#\$\$%&() in pool names for NSS pools that you plan to cluster enable.

9.3 Creating 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 9.3.1, “Default Resource Templates,” on page 124](#)
- ♦ [Section 9.3.2, “Creating a Resource Template,” on page 125](#)

9.3.1 Default Resource Templates

[Table 9-1](#) identifies the cluster resource templates that Novell Cluster Services provides for use on physical servers and Xen virtual machine (VM) guest servers (DomU). You can also create your own templates or personalize the default templates by using iManager. For information, see [Section 9.3.2, “Creating a Resource Template,” on page 125](#). Third-party templates might also be available for third-party applications; see the vendor documentation.

Table 9-1 Cluster Resource Templates for Physical Servers and Xen VM Guest Servers

Cluster Resource Template	OES 2 Linux Product
AV	Novell Archive and Version Services
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)	Linux POSIX file systems
Generic IP Service	This template can be modified to create cluster resources for certain server applications that run on your cluster.
iFolder	Novell iFolder
iPrint	Novell iPrint
MySQL	Novell MySQL
Samba	Novell Samba
Third-party templates	See your vendor documentation.

Novell Cluster Services provides the following templates for use by the Xen VM host server (Dom0). They are the only two templates supported for the host, and they are not supported for use by a VM guest server.

Table 9-2 Cluster Resource Templates for Xen-Based Virtualization Host Environments

Cluster Resource Template	Use
Xen	Automatically configure the cluster resource for the virtual machine.
XenLive	Automatically configure the cluster resource for the virtual machine. 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.

To view the default templates in iManager:

- 1 Start 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 that has iManager installed.
- 2 Enter your Administrator username and password.
- 3 In *Roles and Tasks*, click *Clusters*, then click *Cluster Options*.
- 4 Browse to locate and select the Cluster object of the cluster you want to manage.
A list of available resources and resource templates are displayed.

You can also view the templates outside of iManager. They are cached as `/var/opt/novell/ncs/*_Template.*` files on the master node of the cluster.

9.3.2 Creating a Resource Template

Templates help ensure that all of the necessary definition information, dependent services to be loaded and unloaded, and the shared service or storage are entered correctly when you are configuring multiple servers and clusters. You can use the default templates as a guide for what types of information to include in your personalized template.

- 1 Start 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 that has iManager installed.
- 2 Enter your Administrator username and password.
- 3 In *Roles and Tasks*, click *Clusters*, then click *Cluster Options*.
- 4 Browse to locate and select the Cluster object of the cluster you want to manage.
- 5 Click the *New* link.
- 6 Specify *Template* as the resource type you want to create by clicking the *Template* radio button, then click *Next*.
- 7 In *Cluster Resource Name*, specify the name of the template you want to create.
- 8 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.
- 9 Ensure that the *Define Additional Properties* check box is selected, then click *Next* to continue to the Load Script page.

- 10** On the Load Script page, configure the load 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, container name, file system type, and device.
 - 10b** Edit or add any lines to the load script that are required to load dependent services such as Web servers or file access protocols.
 - 10c** 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.
 - 10d** 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.
- 11** On the Unload Script page, configure the unload script for the cluster resource template.
 - 11a** Edit or add variables with example values for your template configuration, such as the mount point, IP address, container name, file system type, and device.
 - 11b** 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.
 - 11c** Edit or add any lines to the unload script that are required to unload the dependent services that are loaded by this cluster resource.
 - 11d** 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.
- 12** On the Monitor Script page, configure the monitor script for the cluster resource template.
 - 12a** Edit or add the variables with example values for your template configuration, such as the mount point, IP address, container name, file system type, and device.
 - 12b** 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 Novell Cluster Services for Linux include resource monitoring scripts that you can customize.
 - 12c** 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 or migrate) 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.
- 13** On the Resource Policies page, specify the default *Start*, *Failover*, and *Failback* modes, then click *Next*.

- 14 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.

9.4 Creating Cluster Resources

Cluster resources must be created for every resource or application you run on servers in your cluster. Cluster resources can include Web sites, e-mail servers, databases, and any other server-based applications or services you want to make available to users at all times.

- 1 Start your Internet browser and enter the URL for iManager.

The URL is `http://server_ip_address/nps/manager.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 Enter your username and password.

- 3 In the left column, click *Clusters*, then click the *Cluster Options* link.

iManager displays four links under Clusters that you can use to configure and manage your cluster.

- 4 Browse to locate and select the Cluster object of the cluster you want to manage, then 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. Novell clients 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 9.3, “Creating Cluster Resource Templates,” on page 124](#).

- 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 127](#).

9.5 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 10.7, “Configuring a Load Script for the Shared NSS Pool,” on page 155](#)
- ♦ [Section 11.5.1, “Sample Load Script for the Linux POSIX Volume Cluster Resource,” on page 180](#)

If you are creating a new cluster resource, the load script page should already be displayed. You can start with [Step 5](#).

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Select the check box next to the resource whose load script you want to edit, then click the *Details* link.
- 4 Click the *Scripts* tab, then click the *Load Script* link.
- 5 Edit or add the necessary commands to the script to load the resource on the server.

You can then add any lines to the load script that are required to load needed services like Web servers, etc.

You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, container name, file system type, and 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, container name, file system type, and device. 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 IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,” on page 110](#).

- 6 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.

- 7 If you are creating a new cluster resource, continue with [Section 9.6, “Configuring an Unload Script for a Cluster Resource,” on page 128](#).

9.6 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.

Example unload scripts are available in the following sections:

- ♦ [Section 10.8, “Configuring an Unload Script for the Shared NSS Pool,” on page 156](#)
- ♦ [Section 11.5.2, “Sample Unload Script for the Linux POSIX Volume Cluster Resource,” on page 180](#)

If you are creating a new cluster resource, the unload script page should already be displayed. You can start with [Step 5](#).

- 1** In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2** Browse to locate and select the Cluster object of the cluster you want to manage.
- 3** Select the check box next to the resource whose unload script you want to edit, then click the *Details* link.
- 4** Click the *Scripts* tab, then click the *Unload Script* link.
- 5** 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, container name, file system type, and device.
- 6** 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.
- 7** If you are creating a new cluster resource, continue with [Section 9.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 129](#).

9.7 Enabling Monitoring and Configuring the Monitor Script

Resource monitoring monitors the health of specified resources using scripts that you create or customize. Resource monitoring is functional only on Linux clusters. Monitoring is disabled by default.

- ♦ [Section 9.7.1, “Configuring Resource Monitoring,” on page 130](#)
- ♦ [Section 9.7.2, “Example Monitoring Scripts,” on page 131](#)
- ♦ [Section 9.7.3, “Monitoring Services Critical to Clustering,” on page 131](#)

9.7.1 Configuring Resource Monitoring

Although Novell Cluster Services can detect a node failure, it also has the ability to detect when an individual resource on a node has failed. If you want Novell 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, and a failure action.

If you are creating a new cluster resource, the Monitor Script page should already be displayed. You can start with [Step 5](#).

To configure resource monitoring:

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Select the check box next to the resource that you want to configure monitoring for, then click the *Details* link.
- 4 Click the *Monitoring* tab.
- 5 Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

- 6 For the polling interval, specify how often you want the resource monitoring script for this resource to run.

You can choose to specify the number in minutes or seconds.

- 7 Specify the number of failures (*Maximum Local Failures*) for the 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/run/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.

- 8 Specify whether you want the resource to be set to a comatose state, or to migrate to another server if the failure action initiates.

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 *Assigned Nodes* list, which you previously ordered according to your preferences. 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 *Assigned Nodes* list.

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.

- 9 Click the *Scripts* tab, then click the *Monitor Script* link.
- 10 Edit or add the necessary commands to the script to monitor the resource on the server.

You can use the same commands that would be used at the Linux terminal console.

The resource templates included with Novell Cluster Services for Linux include resource monitoring scripts that you can customize.

- 11 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 8](#) 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.

9.7.2 Example Monitoring Scripts

The resource templates included with Novell Cluster Services for Linux include resource monitoring scripts that you can customize.

Example monitor scripts are available in the following sections:

- ♦ [Section 10.9, “Configuring a Monitor Script for the Shared NSS Pool,” on page 156](#)
- ♦ [Section 11.5.3, “Sample Monitor Script for a Linux POSIX Volume Cluster Resource,” on page 181](#)

9.7.3 Monitoring Services Critical to Clustering

Monitoring scripts can also be used for monitoring critical services needed by the resources, such as Linux User Management (namcd) and Novell eDirectory (ndsd). However, the monitoring is in effect only where the cluster resource is running.

IMPORTANT: The 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.

For example, to monitor whether the namcd and ndsd services are running, add the following commands to the Monitor script:

```
# (optional) status of the eDirectory service
exit_on_error rcndsd status

# (optional) status of the Linux User Management service
exit_on_error namcd status
```

9.8 Setting 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 9.8.1, “Understanding Cluster Resource Modes,” on page 132](#)
- ♦ [Section 9.8.2, “Viewing or Modifying the Start, Failover, and Failback Modes for a Resource,” on page 132](#)

9.8.1 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 Assigned Nodes list in the event of 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 *Assigned Nodes* list for the resource.

IMPORTANT: Resources fail back only to the first node in their *Assigned 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 *Assigned Nodes* list rejoins the cluster, unless the *Failback* mode is set to AUTO and the first node in the *Assigned Nodes* list rejoins the cluster.

9.8.2 Viewing or Modifying the Start, Failover, and Failback Modes for a Resource

IMPORTANT: A cluster resource cannot be edited while it is loaded or running on a server. Make sure to offline the resource before attempting to modify its resource properties, policies, or scripts.

If you are creating a new cluster resource, the Resource Policies page should already be displayed. You can start with [Step 5](#).

- 1 In iManager, click *Clusters* and then click *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.

- 3 Select the box next to the resource whose *Start*, *Failover*, or *Failback* modes you want to view or edit, then click the *Details* link.
- 4 Click the *Policies* tab.
- 5 (Conditional) 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.
- 6 (Conditional) 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 Novell 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 Assigned Nodes list as soon as any server in the list is brought online.
- 7 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.
- 8 Continue with [“Assigning Nodes to a Resource” on page 133](#), or if you are creating a new cluster resource, click *Next*, then continue with [“Assigning Nodes to a Resource” on page 133](#).

9.9 Assigning Nodes to a Resource

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 5](#).

IMPORTANT: A cluster resource cannot be edited while it is loaded or running on a server. Make sure to offline the resource before attempting to modify its resource properties, policies, or scripts.

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Select the box next to the resource whose preferred node list you want to view or edit, then click the *Details* link.
- 4 Click the *Preferred Nodes* tab.
- 5 From the *Unassigned Nodes* list, select the server you want the resource assigned to, then click the right-arrow button to move the selected server to the *Assigned Nodes* list.

Repeat this step for all servers you want assigned to the resource.
- 6 From the *Assigned Nodes* list, select the servers you want to unassign from the resource, then click the left-arrow button to move the selected servers to the *Unassigned Nodes* list.
- 7 Click the up-arrow and down-arrow buttons to change the preferred failover order of the servers assigned to the resource or volume.
- 8 Click *Apply* to save node assignment changes.

9.10 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*, then select *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click the *Properties* button under the cluster name.
- 4 Click the *Priorities* tab.
- 5 To change the priority for a resource, select the resource in the list by clicking it, then click the up-arrow or down-arrow to move the resource up or down in the list.

This lets you change the load order of the resource relative to other cluster resources on the same node.

- 6 Click *Apply* or *OK* to save changes.

9.11 Changing the IP Address of a Cluster Resource

You can change the IP address of a cluster resource by using the Cluster Properties page of the Clusters plug-in in iManager. Make sure to offline the cluster resource before attempting to change its IP address. When you online the cluster resource, the load script is updated with the new IP address.

9.12 Deleting Cluster Resources

Make sure to offline the cluster resource before attempting to delete either the cluster resource or the clustered pool.

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.

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.

- ♦ [Section 9.12.1, “Deleting a Cluster Resource on a Master Node,” on page 135](#)
- ♦ [Section 9.12.2, “Deleting a Cluster Resource on a Non-Master Node,” on page 135](#)

9.12.1 Deleting a Cluster Resource on a 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. If the resource cannot be migrated to the master node, follow the procedure in [Section 9.12.2, “Deleting a Cluster Resource on a Non-Master Node,”](#) on page 135.

You might want to delete the shared storage area if you no longer need the data.

WARNING: Deleting a pool or Linux POSIX volume destroys all data on it.

- 1 If the resource is on a non-master node in the cluster, migrate it to the master node.
- 2 If the cluster resource is online, offline it before continuing by using one of the following methods:
 - ♦ Enter the following at the terminal console prompt as the `root` user:

```
cluster offline resource
```
 - ♦ In iManager, go to *Clusters > Cluster Manager*, specify the cluster you want to manage, select the cluster resource, then click *Offline*.
- 3 Delete the resource on the master node by using the appropriate storage management tool:
 - ♦ For shared NSS pools and volumes, use NSSMU or the Storage plug-in to iManager. Deleting the pool automatically deletes the volumes on it.
 - ♦ For shared Linux POSIX volumes, use `evmsgui`.
- 4 In eDirectory, look at the Cluster Resource objects in the Cluster container to verify that the resource has been deleted from the Cluster container.

If necessary, you can delete the Cluster Resource object and its virtual server object manually. In iManager, go to *Directory Administration > Delete Objects*, select the objects, then click *OK*.

9.12.2 Deleting a Cluster Resource on a Non-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. If the resource can be migrated, migrate it to the master node and follow the procedure in [Section 9.12.1, “Deleting a Cluster Resource on a Master Node,”](#) on page 135.

You might want to delete a cluster resource on a non-master node when deleting NSS pool and volume resources (by using NSSMU).

If you must delete a cluster resource while it resides a non-master node, use the following procedure:

- 1 Log in as the `root` user to the non-master node where the cluster resource currently resides, then open a terminal console.
- 2 If the cluster resource is online, offline it by entering

```
cluster offline resource
```
- 3 At the terminal console 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 resource on the non-master node by using the appropriate storage management tool:
 - ♦ For shared NSS pools and volumes, use NSSMU or the Storage plug-in to iManager.
 - ♦ For shared Linux POSIX volumes, use evmsgui.
- 6 In eDirectory, look at the objects in the Cluster container to verify that the resource has been deleted from the Cluster container.
- 7 On the master node, log in as the `root` user, then open a terminal console.
- 8 At the terminal console prompt on the master node, 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 > Cluster Options*, then browse to select the Cluster object.
- 11 Click *Properties*, select the *Priorities* tab, then click *Apply* on the Priorities page.
- 12 At the terminal console, 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, you are done. Do not 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.
- 14 If the revision number did not automatically update in the previous steps, restart Novell Cluster Services by entering the following on one node in the cluster:

```
 cluster restart
```

9.13 Additional Information for Creating Cluster Resources

Cluster resources can be configured for storage, services, or virtual machines.

- ♦ [Section 9.13.1, “Creating Storage Cluster Resources,” on page 137](#)
- ♦ [Section 9.13.2, “Creating Service Cluster Resources,” on page 137](#)
- ♦ [Section 9.13.3, “Creating Virtual Machine Cluster Resources,” on page 137](#)

9.13.1 Creating Storage Cluster Resources

For information about creating cluster resources for shared storage on Linux, see the following:

Table 9-3 *Cluster-Enabling Shared Storage*

Shared Storage	Refer to
NSS pools and volumes	Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139
Linux POSIX volumes	Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163
NCP volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 2 SP2: NCP Server for Linux Administration Guide</i>
Dynamic Storage Technology shadow volume pairs (shared NSS pools and volumes configured as DST pairs)	“Configuring DST Shadow Volumes with Novell Cluster Services for Linux” in the <i>OES 2 SP2: Dynamic Storage Technology Administration Guide</i>

9.13.2 Creating Service Cluster Resources

For information about creating cluster resources for various Linux services on your OES 2 Linux server, refer to the clustering sections in the individual guides. A list of the services is provided in [“Clustering Linux Services” on the Clustering \(High Availability\) Documentation Web site \(http://www.novell.com/documentation/oes2/cluster-services.html#clust-config-resources\)](#).

9.13.3 Creating Virtual Machine Cluster Resources

If you install Novell Cluster Services at the host level of an OES 2 Linux (Xen) server, you can create cluster resources for the virtual machines. For information, see [Section 12.2, “Virtual Machines as Cluster Resources,” on page 186](#).

Configuring Cluster Resources for Shared NSS Pools and Volumes

10

This section describes how to create shared disk partitions and Novell Storage Services (NSS) pools if they do not already exist and to configure the shared pools to work with Novell Cluster Services by cluster-enabling the pools.

- ◆ [Section 10.1, “Planning for Shared NSS Pools and Volumes,” on page 139](#)
- ◆ [Section 10.2, “Considerations for Working with Shared NSS Pools and Volumes in the Cluster,” on page 142](#)
- ◆ [Section 10.3, “Creating NSS Shared Disk Partitions and Pools,” on page 143](#)
- ◆ [Section 10.4, “Creating NSS Volumes on a Shared Pool,” on page 148](#)
- ◆ [Section 10.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 150](#)
- ◆ [Section 10.6, “Adding Advertising Protocols,” on page 153](#)
- ◆ [Section 10.7, “Configuring a Load Script for the Shared NSS Pool,” on page 155](#)
- ◆ [Section 10.8, “Configuring an Unload Script for the Shared NSS Pool,” on page 156](#)
- ◆ [Section 10.9, “Configuring a Monitor Script for the Shared NSS Pool,” on page 156](#)
- ◆ [Section 10.10, “Mirroring and Cluster-Enabling Shared NSS Pools and Volumes,” on page 156](#)
- ◆ [Section 10.11, “Deleting NSS Pool Cluster Resources,” on page 161](#)
- ◆ [Section 10.12, “Changing the Volume ID,” on page 161](#)
- ◆ [Section 10.13, “What’s Next,” on page 162](#)

10.1 Planning for Shared NSS Pools and Volumes

- ◆ [Section 10.1.1, “Shared Storage,” on page 140](#)
- ◆ [Section 10.1.2, “Novell Cluster Services,” on page 140](#)
- ◆ [Section 10.1.3, “Novell Storage Services,” on page 140](#)
- ◆ [Section 10.1.4, “IP Address,” on page 141](#)
- ◆ [Section 10.1.5, “NCP Server for Linux,” on page 141](#)
- ◆ [Section 10.1.6, “Novell CIFS for Linux,” on page 141](#)
- ◆ [Section 10.1.7, “Novell AFP for Linux,” on page 142](#)

10.1.1 Shared Storage

You should carefully plan how you want to configure your shared storage prior to installing Novell Cluster Services. Consider the guidelines and requirements in the following sections when planning your NSS storage solution.

- ♦ [Section 3.1.4, “Shared Disk System Requirements,” on page 38](#)
- ♦ [Section 3.1.5, “SAN Rules for LUN Masking,” on page 40](#)

10.1.2 Novell Cluster Services

You must install Novell Cluster Services before you can create shared disk partitions and pools on shared storage (storage area network, or SAN). Novell Cluster Services must be loaded and running when you create and use the shared NSS pools and volumes.

In a mixed-cluster of NetWare nodes and OES 1 or OES 2 Linux nodes, only NSS pool cluster resources that were created on a NetWare cluster node can be failed over between the Linux and NetWare nodes.

No storage management functions should be executed while a cluster is in a mixed-cluster mode. Do not attempt to create, delete, expand, or modify the properties for devices, partitions, pools, or volumes for any shared resources in the cluster.

WARNING: Attempting to reconfigure shared storage in a mixed cluster can cause data loss.

10.1.3 Novell Storage Services

In order to use NSS pools and volumes in cluster resources, NSS must be installed and running on each server in the cluster. You can use the Storage plug-in for Novell iManager or the server-based NSSMU (NSS Management Utility) tool to create and cluster-enable a shared NSS pool, and to create NSS volumes on those pools.

For information on NSS system and network requirements, see [“Requirements for Installing NSS”](#) in the *OES 2 SP2: NSS File System Administration Guide*.

For information on installing NSS when you install Novell Cluster Services, see [Step 5 on page 45](#).

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)

You must create at least one volume in the shared pool. Typically, all volumes for a shared pool are created when you initially 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.

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 to have the objects in the same context or in different context. If the objects are in different contexts, you might receive an eDirectory error when you attempt to modify the pool, create or modify the volumes, home directories, Distributed File Services junctions, or any other elements that are managed using eDirectory objects. To resolve the problem, you must cluster migrate the pool cluster resource back to the node where the pool was created in order to perform those management tasks.

10.1.4 IP Address

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

IMPORTANT: The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

10.1.5 NCP Server for Linux

NetWare Core Protocol (NCP) is the Novell networking protocol used by the Novell Client. NCP is automatically selected as an advertising protocol when you cluster-enable an NSS pool.

Novell Storage Services requires that the NCP Server for Linux service be installed, configured, 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 Novell CIFS or Novell AFP for data access to prevent possible data corruption that might occur from cross-protocol access to files. Beginning in OES 2 SP2, 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.

For information, see “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 2 SP2: 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 2 SP2: NCP Server for Linux Administration Guide*.

10.1.6 Novell CIFS for Linux

Common Internet File System (CIFS) is the Windows networking protocol. The Novell CIFS for Linux service is available beginning in the OES 2 SP1 Linux release. Novell CIFS is required when you want to give clients access via CIFS to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, make sure that the NCP Cross-Protocol File Locks parameter is enabled for NCP Server on all nodes in the cluster before you allow access to the data via Novell CIFS. For information, see “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

Novell CIFS must be installed, configured, and running before you can specify CIFS as an advertising protocol when cluster-enabling an NSS pool. Novell CIFS for Linux is installed by selecting *Novell CIFS* from the OES Services menu in the YaST install interface. For information about Novell CIFS for Linux, see the *OES 2 SP2: Novell CIFS for Linux Administration Guide*.

10.1.7 Novell AFP for Linux

Apple Filing Protocol (AFP) is the Macintosh networking protocol. The Novell AFP for Linux service is available beginning in the OES 2 SP1 Linux release. Novell AFP is required when you want to give Macintosh clients access via AFP to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, make sure that the NCP Cross-Protocol File Locks parameter is enabled for NCP Server on all nodes in the cluster before you allow access to the data via Novell AFP. For information, see “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

Novell AFP must be installed, configured, and running before you can specify AFP as an advertising protocol when cluster-enabling an NSS pool. Novell AFP for Linux is installed by selecting *Novell AFP* from the OES Services menu in the YaST install interface. For information about Novell AFP for Linux, see the *OES 2 SP2: Novell AFP For Linux Administration Guide*.

10.2 Considerations for Working with Shared NSS Pools and Volumes in the Cluster

Consider the following when working with shared NSS pools and volumes in the cluster:

- ◆ When the pool cluster resource is brought online, the pool is automatically activated by the resource load script. You don’t need to activate the pool at the terminal console.
- ◆ If you delete a volume from a cluster-enabled pool, Novell Cluster Services automatically removes the volume mount command from the resource load script.
- ◆ If you delete a cluster-enabled pool, Novell Cluster Services automatically removes the Pool Resource object and the virtual server object from Novell eDirectory. Make sure you offline the cluster resource before you attempt to delete the clustered pool or its cluster resource. For information, see [Section 9.12, “Deleting Cluster Resources,”](#) on page 134.
- ◆ Before you rename a cluster-enabled pool, make sure to offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, then online the pool resource.

Novell 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. However, does not modify the cluster resource name to use the new pool name.

- ◆ 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 receive an eDirectory error when you attempt to modify the pool, create or modify the volumes, home directories, Distributed File Services junctions, or any other elements that are managed using eDirectory objects. To resolve the problem, you must cluster migrate the pool cluster resource back to the node where the pool was created in order to perform those management tasks.

10.3 Creating NSS Shared Disk Partitions and Pools

NSS pools can be cluster-enabled at the same time they are created or they can be cluster-enabled at a later time after they are created. For information about managing NSS pools, see “[Managing NSS Pools](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

- ♦ [Section 10.3.1, “Initializing Shared Devices,”](#) on page 143
- ♦ [Section 10.3.2, “Enabling Sharing on a Device,”](#) on page 143
- ♦ [Section 10.3.3, “Creating Shared NSS Pools,”](#) on page 144

10.3.1 Initializing Shared Devices

All uninitialized devices that you want to use in a cluster-enabled pool must be initialized by using NSSMU or iManager before you can create cluster-enabled NSS pools on them.

Initialized devices with available free space are listed when you create an NSS pool. If a device that you want to use does not appear in the list of available devices, it might be because it has not yet been initialized. It might also mean that there is no free space on the device.

WARNING: The initialization process completely removes all the partitions on the device, and all of the data stored on the device is lost. Do not initialize devices that already contain data unless you intend to remove all of the data on the device.

To initialize a device:

- 1 Start NSSMU by entering `nssmu` at the terminal console of the master node of the cluster.
- 2 From the NSSMU main menu, click *Devices*.
A list of devices appears in the *Devices* list.
- 3 In the *Devices* list, select a device (such as `sdc`).
- 4 Press F3, *Initialize Disk*.
- 5 Exit NSSMU, or continue with [Section 10.3.2, “Enabling Sharing on a Device,”](#) on page 143.

10.3.2 Enabling Sharing on a Device

All devices that you want to use in a cluster-enabled pool must be enabled for sharing by using NSSMU or iManager before you can create cluster-enabled NSS pools on them.

Storage devices must be available via a shared storage area network (SAN) in order to be failed over between servers in the cluster. NSS supports Fibre Channel SAN and iSCSI SAN environments. External SCSI device arrays can also be configured for multiple server access.

Before enabling sharing on shared storage, Novell Cluster Services must be installed and running on your servers.

To enable sharing for a device:

- 1 Start NSSMU by entering `nssmu` at the terminal console of the master node of the cluster.
- 2 From the NSSMU main menu, click *Devices*.

A list of devices appears in the *Devices* list. Device names are not changeable and might be labeled with names such as *sdc*.

- 3 In the *Devices* list, select a device.
- 4 Press F6, *Share*.
- 5 Wait for the device details to refresh, then verify that the *Shareable for Clustering* parameter is now marked Yes.
- 6 Exit NSSMU.

10.3.3 Creating Shared NSS Pools

You can create a pool on the shared device by using iManager or NSSMU. The shared partition is automatically created for you when you create the pool.

- ♦ [“Using iManager to Create Shared NSS Pools” on page 144](#)
- ♦ [“Using NSSMU to Create Shared NSS Pools” on page 146](#)

Using iManager to Create Shared NSS Pools

- 1 Start 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 that has iManager installed or with the IP address for Apache-based services.
- 2 Specify your username and password, specify the tree information, then click *Login*.
- 3 In the left column, locate *Storage*, then click the *Pools* link.
- 4 Enter a cluster server name or browse and select the Cluster object.
- 5 Click the *New* link to open the New Pool wizard.
- 6 Specify the new pool name, then click *Next*.
- 7 Select the check box next to the shared device where you want to create the pool, then specify the size of the pool.
- 8 For the *Mount On Creation* option, specify whether you want the pool to be activated when it is created.

The *Mount On Creation* option determines if the pool you are creating is to be activated as soon as it is created. The *Activate On Creation* check box is selected (enabled) by default. If you deselect the check box, you must manually activate the pool later before it can be used. If the pool is cluster-enabled, activation is done automatically by when you bring the pool cluster resource online.

- 9 For the *Cluster Enable on Creation* option, specify whether you want the pool to be cluster-enabled when it is created by doing one of the following:
 - ♦ **Selected (default):** If you want to cluster-enable the pool at the same time it is created, leave the *Cluster Enable on Creation* check box selected, click *Next*, then continue with [Step 10 on page 145](#).

- ◆ **Deselected:** If you want to cluster-enable the pool at a later date, deselect the check box, then click *Finish* to complete the pool creation process.

To cluster-enable the existing pool, continue by creating a volume on the shared pool as described in “[Creating NSS Volumes on a Shared Pool](#)” on page 148, then use the *Clusters* plug-in to iManager to cluster-enable the existing pool as described in [Section 10.5](#), “[Cluster-Enabling an Existing NSS Pool and Its Volumes](#),” on page 150.

- 10** If you selected *Cluster Enable on Creation*, an additional Cluster Information page appears that allows you to specify the cluster information.

Specify the following shared pool clustering parameters:

- ◆ **Virtual Server Name:** The name assigned to the virtual server that represents the shared pool in the cluster.

When you cluster-enable a pool, a virtual Server object is automatically created in Novell eDirectory and given the name of the Cluster object plus the name of the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

- ◆ **CIFS Virtual Server Name:** The name (up to 15 characters) assigned to the virtual server for handling CIFS (Common Internet File System) requests. This is the name of the server as it appears in a Windows system.

By default, the suggested CIFS virtual server name follows the format of `clustername_poolname_w`. If the name does not comply with the 15-character maximum length, the default name is shortened by removing characters from the left.

- ◆ **IP Address:** The IP address that you want to assign the virtual server.

To specify an IP address, tab between the different entries; no dot is required in the fields. For example, if the IP address is `192.168.1.1`, type the following:

```
192 168 1 1
```

- ◆ **Advertising Protocols:** Protocols that give users native file access to data.

Select the check boxes of the protocols you want to enable for data requests to this shared pool.

If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.

NOTE: For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available. CIFS and AFP check boxes can be selected, but CIFS and AFP functionality does not apply to Linux. Selecting the check boxes has no effect.

- ◆ **AFP:** AFP is the Macintosh networking protocol. Selecting AFP causes commands to be added to the pool-resource load and unload scripts to activate the AFP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to AFP clients.
- ◆ **CIFS:** CIFS is the Windows networking protocol. Selecting CIFS causes commands to be added to the pool-resource load and unload scripts to activate the CIFS protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to CIFS/Samba clients.
- ◆ **NCP:** NCP is the Novell networking protocol used by the Novell Client. It is provided by default and is mandatory. Selecting NCP causes commands to be added to the pool-resource load and unload scripts to activate the NCP protocol on the

cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to Novell clients and that authentication and access to the file system is enforced according to the Novell Trustee Model.

11 Click *Finish*.

For NSS on Linux, the create time might take longer than expected. Typically, the pool creation takes less than a minute, and the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

12 Continue with [“Creating NSS Volumes on a Shared Pool” on page 148](#).

Using NSSMU to Create Shared NSS Pools

- 1** Start NSSMU by entering `nssmu` at the terminal console of the master node of the cluster.
- 2** From the NSSMU main menu, select *Pools*.
- 3** On the Pools page, press Insert, type a name for the new pool you want to create, then press Enter.
- 4** From the list of available devices, select the shared device where you want the pool created (such as `sdC`), then press Enter.
- 5** Specify the amount of space (in MB) to use, then press Enter.
- 6** 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 as soon as it is created. The *Activate On Creation* option is Yes (enabled) by default.

If you set the value to No, you must manually activate the pool later before it can be used. If the pool is cluster-enabled, activation is done automatically by when you bring the pool cluster resource online.

- 7** The *Cluster Enable on Creation* feature is enabled by default for shared devices. Specify whether you want the pool to be cluster-enabled when it is created by doing one of the following:
 - ♦ **Yes (default):** If you want to cluster-enable the pool at the same time it is created, accept the default entry (*Yes*), select *Create*, then continue with [Step 8](#).
 - ♦ **No:** If you want to cluster-enable the pool at a later date, change the default entry from *Yes* to *No*, then select *Apply* and press Enter to create the pool.

To cluster-enable an existing pool, continue by creating a volume on the shared pool as described in [“Creating NSS Volumes on a Shared Pool” on page 148](#), then use the *Clusters* plug-in to iManager to cluster-enable the existing pool as described in [Section 10.5](#), [“Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 150](#).
- 8** If you enabled *Cluster Enable on Creation*, specify the following fields in the area under Shared Pool Clustering Parameters:
 - ♦ **Virtual Server Name:** The name assigned to the virtual server that represents the shared pool in the cluster. Press Enter to modify the suggested name, type the name, then press Enter to continue to the next field.

When you cluster-enable a pool, a virtual Server object is automatically created in Novell eDirectory and given the name of the Cluster object plus the name of the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

- ◆ **CIFS Virtual Server Name:** The name (up to 15 characters) assigned to the virtual server for handling CIFS requests. This is the name of the server as it appears in a Windows system. Press Enter to modify the suggested name, type the name, then press Enter to continue to the next field.

By default, the suggested CIFS virtual server name follows the format of `clustername_poolname_w`. If the name does not comply with the 15-character maximum length, the default name is shortened by removing characters from the left.

- ◆ **IP Address:** Press Enter, specify the IP address that you want to assign the virtual server, then press Enter to continue to the next field.

Make sure to type the dots between the address fields. For example, if the IP address is `192.168.1.1`, type the following:

```
192.168.1.1
```

- ◆ **Advertising Protocols:** Protocols that give users native file access to data.

Specify Yes next to the protocols that you want to enable for data requests to this shared pool.

If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.

NOTE: On Linux, Novell CIFS and Novell AFP are available beginning in OES 2 SP1. For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available. CIFS and AFP check boxes can be selected, but CIFS and AFP functionality does not apply to Linux. Enabling the options has no effect.

- ◆ **NCP:** NCP is the Novell networking protocol used by the Novell Client. It is provided by default and is mandatory. Selecting NCP causes commands to be added to the pool-resource load and unload scripts to activate the NCP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to Novell clients and that authentication and access to the file system is enforced according to the Novell Trustee Model.
- ◆ **CIFS:** CIFS is the Windows networking protocol. Selecting CIFS causes commands to be added to the pool-resource load and unload scripts to activate the CIFS protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to CIFS/Samba clients.
- ◆ **AFP:** AFP is the Macintosh networking protocol. Selecting AFP causes commands to be added to the pool-resource load and unload scripts to activate the AFP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to AFP clients.

9 Select *Apply*, then press Enter to create and cluster-enable the pool.

10 Continue with [“Creating NSS Volumes on a Shared Pool” on page 148](#).

10.4 Creating NSS Volumes on a Shared Pool

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.

IMPORTANT: Before adding volumes on a shared pool, make sure that the pool is mounted on the server where the pool was created. This is necessary because the Pool object information is relative to the Server object of the server where the pool was created.

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.

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 Novell 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.

You can create NSS volumes by using the Storage plug-in to Novell iManager or the server-based NSSMU tool. For planning and management information, see the [OES 2 SP2: NSS File System Administration Guide](#).

- ♦ [Section 10.4.1, “Using iManager to Create NSS Volumes,” on page 148](#)
- ♦ [Section 10.4.2, “Using NSSMU to Create NSS Volumes,” on page 149](#)

10.4.1 Using iManager to Create NSS Volumes

Repeat the following procedure for each cluster volume that you want to create on the shared pool.

- 1 Make sure the pool resource is mounted on the same server in the cluster where the pool was created.
- 2 Start 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 that has iManager installed or with the IP address for Apache-based services.
- 3 Specify your username and password, specify the tree information, then click *Login*.
- 4 In the left column, select *Storage*, then select the *Volumes*.
- 5 Specify a cluster server name or browse and select one, then click *New*.
Select the server where the shared pool was originally created and is currently mounted.
- 6 Specify the new volume name, then click *Next*.
- 7 Select the check box next to the cluster pool where you want to create the volume and either specify the size of the volume (*Volume Quota*) or select *Allow the volume to grow to the size of the pool*, then click *Next*.

The volume quota is the maximum possible size of the volume. If you have more than one volume per pool, you can specify a quota for each volume rather than allowing multiple volumes to grow to the pool size.

- 8 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 2 SP2: NSS File System Administration Guide*.

- 9 Choose whether you want the volume activated and mounted when it is created, then click *Finish*.

10.4.2 Using NSSMU to Create NSS Volumes

Repeat the following procedure for each cluster volume that you want to create on the shared NSS pool.

- 1 Make sure the pool resource is mounted on the same server in the cluster where the pool was created.
- 2 Start NSSMU by entering `nssmu` at the terminal console of the cluster server where the pool was originally created and is currently mounted.
- 3 From the NSSMU main menu, select *Volumes*.
- 4 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.
- 5 Specify Y(es) to encrypt volume or N(o) to create a regular volume.
- 6 From the list of available pools, select the pool where you want the volume to reside, then press Enter.
The new volume appears in the list of volumes.
- 7 (Optional) From the list of volumes, select the newly created volume, then Press F5, *Properties*, to review and change volume attributes, then select *Apply* and press Enter to save any changes.

The default settings for new volumes are described in the following table:

Volume Property	Default Value
Backup	Yes
Compression	No
Data Shredding	0 (off), Times to shred data: 1
Directory Quotas	No
Modified File List (MFL)	No
Salvage Files	Yes
Snapshot—File Level	No
User Space Restrictions	No

Volume Property	Default Value
Flush Files Immediately	No
Migration	No
Allow Renaming Mount Point	Yes
Read Ahead Count in Blocks	16, Block Size=4.00 KB
Data Area for <i>poolname</i> is	<value> MB (This field is not modifiable.)
Volume Quota in MB	0 MB, 0=Grow to pool size
	The quota is the maximum possible size of the volume. If you have more than one volume per pool, you can specify a quota for each volume rather than allowing multiple volumes to grow to the pool size.

For information about volume attributes, see “[Volume Attributes](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

8 Exit NSSMU.

10.5 Cluster-Enabling an Existing NSS Pool and Its Volumes

Cluster-enabling a 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. If you have a shared disk system that is part of your cluster and you want the pools and volumes on the shared disk system to be highly available to NCP clients, you must cluster-enable those pools and volumes.

IMPORTANT: You can cluster-enable pools when you create them by following the procedures in [Section 10.3, “Creating NSS Shared Disk Partitions and Pools,”](#) on page 143 and [Section 10.4, “Creating NSS Volumes on a Shared Pool,”](#) on page 148.

This procedure describes how to enable clustering for an existing pool and its volumes. It assumes the following:

- ♦ The device and pool are enabled as Shareable for Clustering.
- ♦ The pool contains at least one volume.
- ♦ The pool resides on the server where it and its volumes were created.
- ♦ The pool and its volumes have Storage objects in the eDirectory tree where you are setting up the cluster.

If the objects are missing, you can use the Update eDirectory options in the Storage plug-in for NSS to create them in eDirectory. For information, see “[Updating eDirectory Pool Objects](#)” and “[Updating eDirectory Volume Objects](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

- ♦ The pool resides on a disk that has been enabled as *Sharable for Clustering*. For information, see [Section 10.3.2, “Enabling Sharing on a Device,”](#) on page 143.
- ♦ Each pool resource requires its own static IP address that is in the same IP subnet as the cluster IP address.

Volumes on cluster-enabled pools do not appear as independent cluster resources. NSS pools are resources, and the resource’s load and unload scripts are automatically generated for them. Each cluster-enabled NSS pool requires its own IP address. This means that each cluster-enabled volume does not have an associated load and unload script or an assigned IP address.

When a server fails, any cluster-enabled pools being accessed by that server 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.

If you want each volume to be its own cluster resource, each volume must have its own pool.

Some server applications don't require NCP client access to NSS volumes, so cluster-enabling pools and volumes might not be necessary.

IMPORTANT: Before you attempt to cluster-enable an existing pool and its volumes, the pool should be deactivated and its volumes should be dismantled.

Comment out (or remove) the volume’s entry in the `/etc/fstab` file. The load and unload scripts that are created when you cluster-enable the pool will be responsible for mounting and dismantling the volume after the pool is cluster enabled.

- 1 Start 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 that has iManager installed or with the IP address for Apache-based services.
- 2 Specify your username and password, specify the tree information, then click *Login*.
- 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*, click *Cluster Options*, then browse to locate and select the Cluster object of the cluster you want to manage.
- 5 On the Cluster Options page, click the *New* link.
- 6 Specify Pool as the resource type you want to create by clicking the *Pool* radio button, then click *Next*.
- 7 Type the name of the pool you want to cluster-enable, or browse and select it.
- 8 (Optional) Change the default name of the virtual Server object.
When you cluster-enable a pool, a Virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name is `cluster1_pool1_server`.
- 9 If Novell CIFS is running on the cluster server and you want to enable it for this pool, specify the *CIFS Virtual Server Name* (up to 15 characters), which is the name assigned to the virtual server for handling CIFS requests.

This is the name of the server as it appears in a Windows system.

If Novell CIFS is not installed and running, this field value is `NOT_SUPPORTED`.

10 Specify an IP address for the pool.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

11 Select the check boxes of the advertising protocols you want to enable for data requests to this shared pool.

If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.

NOTE: For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available. CIFS and AFP check boxes can be selected, but CIFS and AFP functionality does not apply to Linux. Selecting the check boxes has no effect.

- ♦ **AFP:** AFP is the Macintosh networking protocol. Selecting AFP causes commands to be added to the pool-resource load and unload scripts to activate the AFP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to AFP clients.
- ♦ **CIFS:** CIFS is the Windows networking protocol. Selecting CIFS causes commands to be added to the pool-resource load and unload scripts to activate the CIFS protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to CIFS/Samba clients.

IMPORTANT: In OES 2 SP1 Linux, the script is not modified properly. See the [OES 2 SP2: Readme](#) for the workaround.

- ♦ **NCP:** NCP is the Novell networking protocol used by the Novell Client. It is provided by default and is mandatory. When you cluster-enable the pool, NCP adds commands to the pool-resource load and unload scripts to activate the NCP protocol on the cluster. This lets you ensure that the cluster-enabled pool you are creating is highly available to Novell clients and that authentication and access to the file system is enforced according to the Novell Trustee Model.

12 (Optional) Select the *Online Resource after Create* check box.

This causes the NSS volume to automatically mount when the resource is created.

13 Ensure that the *Define Additional Properties* check box is selected, then click *Next*.

14 Complete the following tasks to configure the behavior of the NSS pool cluster resource you just created:

1. “[Setting Start, Failover, and Failback Modes for Cluster Resources](#)” on page 131
2. “[Assigning Nodes to a Resource](#)” on page 133
3. “[Configuring Resource Priorities for Load Order](#)” on page 134

15 If desired, modify the cluster load, unload, and monitor scripts.

Cluster resource load, unload, and monitor scripts are automatically generated for a pool when the cluster resource for the pool is created. You can modify the scripts and their timeout values to suit your needs. For examples of the scripts for NSS, see the following:

- ♦ [Section 10.7, “Configuring a Load Script for the Shared NSS Pool,”](#) on page 155
- ♦ [Section 10.8, “Configuring an Unload Script for the Shared NSS Pool,”](#) on page 156

10.6 Adding Advertising Protocols

Shared NSS pools support three advertising protocols: NCP (default, mandatory), Novell AFP, and Novell CIFS. You can add Novell AFP or Novell CIFS as advertising protocols for an existing pool cluster resource.

You must install Novell CIFS or Novell AFP 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, make sure the pool is migrated back to the server where it was originally created.

- 1 Start your Internet browser and enter the URL for iManager.

The URL is `http://server_ip_address/nps/manager.html`. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster that has iManager installed or with the IP address for Apache-based services.

- 2 Specify your username and password, specify the tree information, then click *Login*.

- 3 Offline the cluster resource that you want to modify.


For information, see [Section 8.5, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 106.




- 4 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.

- 5 In *Roles and Tasks*, select *Clusters*, click *Cluster Options*, then browse to locate and select the Cluster object of the cluster you want to manage.

- 6 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.



Cluster Objects		
New	Delete	Details
<input type="checkbox"/>	Type	Name
<input type="checkbox"/>		Master IP Address Resource
<input type="checkbox"/>		avalon
<input type="checkbox"/>		POOL1 SERVER

- 7 On the Cluster Pool Properties page, click the *Protocols* tab.

- 8 If you are enabling the resource for Novell AFP, select the *AFP* check box, then click *Apply*.

9 If you are enabling the resource for Novell CIFS, do the following:

9a On the Protocols page, specify the following parameters, then click *Apply*.

- ◆ *CIFS Virtual Server Name* (up to 15 characters)
- ◆ *CIFS* check box

A default name (up to 15 characters) is suggested for the CIFS Virtual Server Name when you first add the advertising protocol. You can modify the default name only at create time.

After you have applied this setting, the CIFS share is created with the name you specified. Thereafter, you cannot use this interface to change the CIFS Virtual Server Name. Use the CIFS plug-in for iManager to manage the CIFS share name of an existing share. For information, see “[Setting CIFS General Server Parameters](http://www.novell.com/documentation/oes2/file_cifs_lx/data/bdkfk5j.html#af3ozzn)” (http://www.novell.com/documentation/oes2/file_cifs_lx/data/bdkfk5j.html#af3ozzn) in the *OES 2 SP2: Novell CIFS for Linux Administration Guide*.

9b Run the `cifsPool.py` script to modify the load and unload scripts.

IMPORTANT: In OES 2 SP1 Linux, the Novell CIFS commands are not automatically added to the load and unload scripts for the NSS pool resource when you add Novell CIFS as an advertising protocol on an existing cluster-enabled pool. You must run the `cifsPool.py` script to add commands for Novell CIFS.

You can download the `cifsPool.py` script file as a patch [Novell Downloads Web site](http://download.novell.com/Download?buildid=sA5Tdb_U2CE~) (http://download.novell.com/Download?buildid=sA5Tdb_U2CE~).

At a terminal console prompt, log in as the `root` user, then enter

```
python cifsPool.py pool_resource_DN cifs_servername ldaps://  
ldap_server_ip:636 admin_DN admin_password
```

For example, enter

```
python cifsPool.py cn=CLUS1_POOL1_SERVER,cn=clus1,o=company CIFS_POOL  
ldaps://10.10.10.1:636 cn=admin,o=company pa$$w0rD
```

where the values are:

Parameter	Sample Value
pool_resource_DN	cn=CLUS1_POOL1_SERVER,cn=clus1,o=company
	NOTE: The cluster container is stored in eDirectory as a CN and not an OU.
cifs_servername (up to 15 characters)	CIFS_POOL1
ldap_server_ip	10.10.10.1
admin_DN	cn=admin,o=company
admin_password	pa\$\$w0rD

10 Online the cluster resource.

For information, see [Section 8.5, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 106.

10.7 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 9.5, “Configuring a Load Script for a Cluster Resource,”](#) on page 127.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, container name, file system type, and device. 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 IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110.

If you specified the following values for the variables in the template, your load script would appear like the script below.

Variable	Your 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/ncsfncs

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 --ncpsservername=NCS1_SHPOOL43_SERVER --
ipaddress=10.10.10.43

exit 0
```

If you change the name space for an existing shared volume by using NSSMU or the NSS 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. Otherwise, the cluster assumes the default name space for mounting the volume. You can do this by using the `/opt=ns=<long|unix|dos|mac>` switch in the command. For information about default name spaces for NSS volumes, see [“Lookup Namespace”](#) in the *OES 2 SP2: NSS File System Administration Guide*.

For example, to specify the Long name space, add the `/opt=ns=long` switch as follows:

```
ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID>
```

For example, to specify the Unix name space, add the `/opt=ns=unix` switch as follows:

```
ncpcon mount /opt=ns=unix <VOLUMENAME>=<VOLUMEID>
```

10.8 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 9.6, “Configuring an Unload Script for a Cluster Resource,”](#) on page 128.

If you specified the following values for the variables in the template, your unload script would appear like the script below.

Variable	Your 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 --ncpservname=NCS1_SHPOOL43_SERVER --
ipaddress=10.10.10.43

ignore_error del_secondary_ipaddress 10.10.10.43
ignore_error nss /pooldeact=SHPOOL43

exit 0
```

10.9 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 9.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 129.

After you enable monitoring, you can modify the script as needed to suit your needs. For information, see [Section 9.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 129.

10.10 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes


- ◆ [Section 10.10.1, “Understanding NSS Mirroring,”](#) on page 157
- ◆ [Section 10.10.2, “Requirements for NSS Mirroring,”](#) on page 158
- ◆ [Section 10.10.3, “Creating and Mirroring NSS Pools on Shared Storage,”](#) on page 159
- ◆ [Section 10.10.4, “Verifying the NSS Mirror Status in the Cluster,”](#) on page 160

10.10.1 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 array.

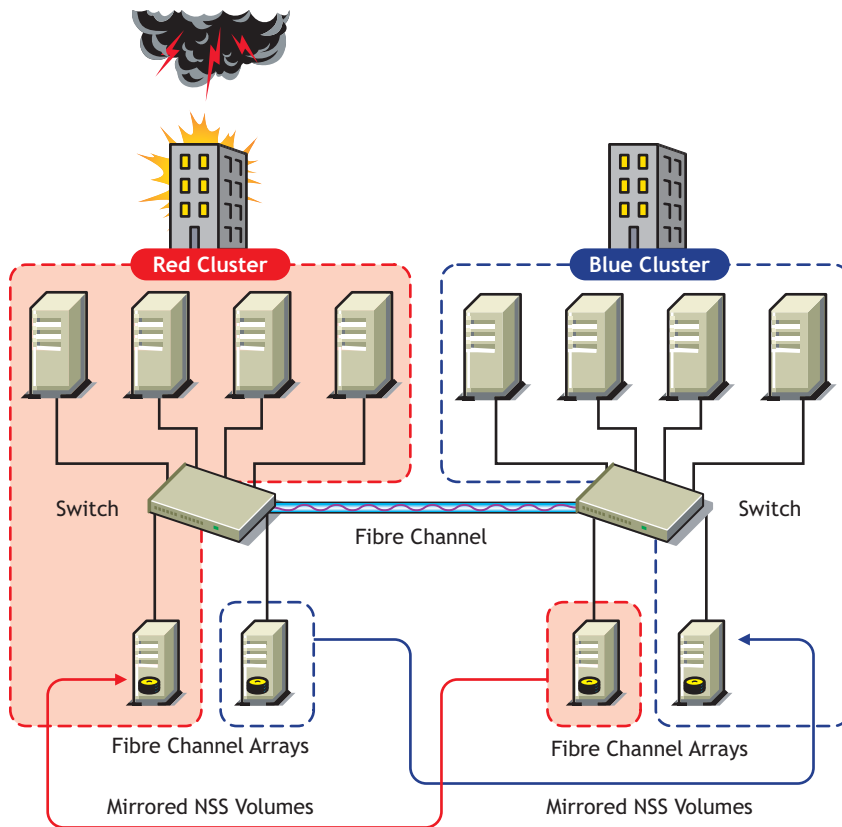
Hardware configuration and placement for NSS mirroring can vary depending on geography and fault tolerance requirements. For example, [Figure 10-1](#) depicts a simple hardware configuration in which one side of the mirrored NSS volume is located in a separate building from the rest of the cluster hardware. If a disaster occurs in one building, data is still safe on the mirrored NSS volume in the other building.

Figure 10-1 Single Cluster with Mirrored NSS Volumes in Separate Buildings



[Figure 10-2](#) 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 10-2 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.

10.10.2 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.

Novell Cluster Services should be installed and running prior to creating and mirroring partitions on shared storage.

When you create a Novell 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. Novell Cluster Services arbitrates access to shared volumes for all cluster nodes, but cannot protect shared volumes from being corrupted by non-cluster servers.

10.10.3 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 2 Linux 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
- ♦ At least 20 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 159](#)
- ♦ [“Mirroring the Partition of an Existing NSS Pool” on page 159](#)

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 terminal console 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.

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 terminal console 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** Press F3, choose the partition that contains the pool (even though it might not look like it has sufficient free space) and select the second shareable partition, then select *Yes*.
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%.

10.10.4 Verifying the NSS Mirror Status in the Cluster

After you have configured NSS mirroring with Novell 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 `volumes` at the terminal console.
- 2** Check the mirror status of the mirrored partition by entering `mirror status` at the terminal console of the server where the NSS pool on the mirrored partition is active.
After entering `mirror 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 `mirror status` at the terminal console.

10.11 Deleting NSS Pool Cluster Resources

Make sure to 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, connect to the Cluster object, not to Cluster Node objects. On the *Storage > Pools* page for iManager, connect to the master node. Run NSSMU only on the master node.

For information, see [Section 9.12, “Deleting Cluster Resources,”](#) on page 134.

10.12 Changing the Volume ID

Novell 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.

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, make sure to 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.

10.13 What's Next

For information on managing the cluster, see [Chapter 8, “Managing Clusters,”](#) on page 101.

Configuring Cluster Resources for Shared Linux POSIX Volumes

11

This section describes how to create shared disk partitions and Linux POSIX volumes as cluster resources with Novell Cluster Services.

- ◆ [Section 11.1, “Requirements for Shared Linux POSIX Volumes,” on page 163](#)
- ◆ [Section 11.2, “Creating Linux POSIX Volumes on Shared Disks,” on page 164](#)
- ◆ [Section 11.3, “Cluster-Enabling a Linux POSIX Volume on a Shared Disk,” on page 168](#)
- ◆ [Section 11.4, “Creating a Virtual Server Name for the Cluster Resource,” on page 175](#)
- ◆ [Section 11.5, “Sample Scripts for a Linux POSIX Volume Cluster Resource,” on page 179](#)
- ◆ [Section 11.6, “Expanding EVMS Volumes on Shared Disks,” on page 181](#)
- ◆ [Section 11.7, “Deleting Shared Storage,” on page 183](#)
- ◆ [Section 11.8, “Known Issues for Working with Cluster Resources for Linux POSIX Volumes,” on page 183](#)
- ◆ [Section 11.9, “What’s Next,” on page 184](#)

11.1 Requirements for Shared Linux POSIX Volumes

- ❑ EVMS must be installed and running on each server in the cluster.
- ❑ You must install Novell Cluster Services before you create shared disk partitions on shared storage (storage area network, or SAN).
- ❑ Cluster Segment Manager (CSM) containers require Novell Cluster Services to be running on all nodes that access the CSM container. Do not create or make to modifications to EVMS objects unless Novell Cluster Services is running.
- ❑ You should carefully plan how you want to configure your shared storage prior to installing Novell Cluster Services. For planning information, see the following:
 - ◆ [Section 3.1.4, “Shared Disk System Requirements,” on page 38](#)
 - ◆ [Section 3.1.5, “SAN Rules for LUN Masking,” on page 40](#)
- ❑ Each cluster-enabled Linux volume needs its own static IP address. The IP address is used to provide access and failover capability to the cluster-enabled volume. The IP address assigned to the volume remains assigned to it regardless of which server in the cluster is accessing the volume.
- ❑ In a mixed cluster of NetWare and OES 1 or 2 Linux nodes, Linux POSIX file systems as cluster resources cannot be created until the entire cluster is migrated to OES 2 Linux. Linux POSIX file systems as cluster resources cannot be migrated or failed over to NetWare cluster nodes. Only NSS pool cluster resources that are created on a NetWare cluster node can be failed over between Linux and NetWare nodes of a mixed-node cluster.

NetWare-to-Linux failover requires that the Linux node be configured for NSS and that the version of NSS supports the NSS media format and features being used by the NSS pool cluster resource.

No storage management functions should be executed while a cluster is in a mixed-cluster mode. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

11.2 Creating Linux POSIX Volumes on Shared Disks

The EVMS (Enterprise Volume Management System) provides a Cluster Segment Manager (CSM) that helps prevent data corruption caused by multiple nodes accessing the same data. For this reason, we recommend using EVMS to create Linux partitions, volumes, and file systems on shared storage for use with Novell Cluster Services. EVMS virtual volumes can more easily be expanded and failed over to different cluster servers than can physical devices. You can create partitions and volumes for any of the journaled Linux file systems (such as Ext3 and ReiserFS). For information and examples for using EVMS, enter `man evms` at the Linux terminal console to reference the `evms` man page.

WARNING: EVMS administration utilities (`evms`, `evmsgui`, and `evmsn`) should not be running when they are not being used. EVMS utilities lock the EVMS engine, which prevents other EVMS-related actions from being performed. This affects both NSS and Linux POSIX volume actions.

NSS and Linux POSIX volume cluster resources should not be migrated while any of the EVMS administration utilities are running.

Complete the tasks in this section to create a Linux POSIX volume and file system on a shared disk:

- ◆ [Section 11.2.1, “Removing Existing Formatting and Segment Managers,” on page 164](#)
- ◆ [Section 11.2.2, “Creating a Cluster Segment Manager Container,” on page 165](#)
- ◆ [Section 11.2.3, “Adding a Non-CSM Segment Manager Container,” on page 166](#)
- ◆ [Section 11.2.4, “Creating an EVMS Volume,” on page 167](#)
- ◆ [Section 11.2.5, “Making a File System on the EVMS Volume,” on page 167](#)

For information about how to increase the size of an EVMS volume on a shared disk, see [Section 11.6, “Expanding EVMS Volumes on Shared Disks,” on page 181](#).

11.2.1 Removing Existing Formatting and Segment Managers

The EVMS Cluster Segment Manager must be the first segment manager laid down on the space you want to use for the shared Linux volume. Before you can assign the CSM to a partition, you must make sure that the space is free of other formatting and segment managers.

WARNING: Clearing the existing formatting for a device destroys all data on it.

EVMS by default shows unformatted free space on disks as compatibility volumes. If any of the space you know to be free on the shared disk (that you plan to use in your cluster) is showing up as a compatibility volume, you must delete the compatibility volume in order to see the space as being free disk space.

If the shared disk that you plan to use with your cluster is configured with segment managers other than the Cluster Segment Manager, you must also delete the other segment managers so that the CSM is the first one on the partition.

- 1** At a Linux terminal console, log in as the `root` user, then enter
`evmsgui`
- 2** Locate and delete any compatibility volumes on the area you want to use for the shared volume:
 - 2a** Click the *Volumes* tab, then locate any compatibility volumes that are on the space you want to use.
 - 2b** Right-click the volume of interest, select *Display details*, click the *Page 2* tab, verify from the *Status* field that the volume is a compatibility volume, then click *OK* to dismiss the dialog box.
 - 2c** If the volume is a compatibility volume, continue with [Step 2d](#). Otherwise, go to [Step 3](#).
 - 2d** Click the *Volumes* tab, right-click the volume, then select *Delete*.

WARNING: All data on the selected volume is destroyed.

 - 2e** Select the volume, then click *Recursive Delete*.
 - 2f** If a *Response Required* pop-up appears, click the *Write zeros* button.
 - 2g** If another pop-up appears, click *Continue* to write 1024 bytes to the end of the volume.
 - 2h** Click *Save*, then click *Save* again to save your changes.
- 3** Locate and delete any existing non-CSM segment managers on the area you want to use for the shared volume, such as NetWare Segment Manager or DOS Segment Manager:
 - 3a** Click the *Disks* tab, then locate and select the device that you want to use for the shared volume, such as device `sdb`.
 - 3b** Right-click, then select *Remove segment manager from Object*.
This option appears only if there is an existing segment manager for the selected disk.
 - 3c** Select the listed non-CSM segment manager, click *Remove*, then click *OK*.

WARNING: All data on the selected disk space is destroyed.

 - 3d** Click *Save*, then click *Save* again to save your changes.
 - 3e** Repeat [Step 3a](#) to [Step 3d](#) until the *Remove segment manager from Object* option is no longer available when you right-click the device.
- 4** Continue with [Section 11.2.2, “Creating a Cluster Segment Manager Container,”](#) on page 165.

11.2.2 Creating a Cluster Segment Manager Container

A Cluster Segment Manager container can provide exclusive access to shared storage. In order to use a Linux POSIX volume as a cluster resource, you must assign an EVMS CSM to the shared device where you want to create the shared volume. The EVMS Cluster Segment Manager must be the first segment manager laid down on the space you want to use for the shared Linux volume. You use the Cluster Segment Manager plug-in for EVMS to create a CSM container.

IMPORTANT: CSM containers require Novell Cluster Services to be running on all nodes that access the CSM container. Make sure Novell Cluster Services is running before you continue.

- 1 In *evmsgui*, click *Actions > Create > Container*.
- 2 On the Create Storage Container page, select *Cluster Segment Manager*, then click *Next*.
- 3 On the Select Plugin Acceptable Objects page, select the disks (storage objects) you want to place in the container (such as *sdb*), then click *Next*.
- 4 On the Configuration Options page, select the cluster server node where you are creating the container, specify *Private* as the type, then specify a name for the container.

Specify a name that is easy to associate with the cluster resource you plan to create for the container. For example, if the IP address you plan to assign to the resource is 10.10.10.44, you might name the container `csm44`.

The name must be one word, must consist of standard alphanumeric characters, and must not be any of the following reserved words:

Container
Disk
EVMS
Plugin
Region
Segment
Volume

- 5 Click *Create*, then click *OK*.
- 6 Click *Save*, then click *Save* again to save your changes.
- 7 Continue with [Section 11.2.3, “Adding a Non-CSM Segment Manager Container,”](#) on page 166.

11.2.3 Adding a Non-CSM Segment Manager Container

After creating a CSM container, you can optionally add a non-CSM segment manager container on top of the CSM container you just created. The benefit of this is that other non-CSM segment manager containers allow you to create multiple smaller EVMS volumes on your shared disk. You can then add additional EVMS volumes or expand or shrink existing EVMS volumes to utilize or create additional free space on your shared disk. In addition, this means that you can also have different file system types on your shared disk.

A CSM container uses the entire EVMS disk, which means that creating additional volumes or expanding or shrinking volumes is not possible. And, because only one EVMS volume is possible in the container, only one file system type is allowed in that container.

- 1 In *evmsgui*, click *Actions > Add > Segment Manager to Storage Object*.
- 2 On the Add Segment Manager to Storage Object page, choose the desired segment manager (such as *DOS Segment Manager*), then click *Next*.
Most of the segment managers will work. The DOS segment manager is added by default for some EVMS operations.
- 3 On the Select Plugin Acceptable Objects page, choose the CSM container storage object where you want to add the segment manager, then click *Next*.

- 4 On the Configurable Options page, select the disk type (Linux is the default), click *Add*, then click *OK*.
- 5 Click *Save*, then click *Save* again to save your changes.
- 6 If you added a DOS segment manager, create a segment for it:

IMPORTANT: Some segment managers such as the DOS segment manager require you to create a segment before creating an EVMS volume. Without a segment, the additional segment manager does not appear when you attempt to create an EVMS volume.

- 6a In *evmsgui*, click *Actions > Create > Segment*.
 - 6b On the Create Disk Segment page, select *DOS Segment Manager*, then click *Next*.
 - 6c On the Select Plugin Acceptable Objects page, choose the CSM container storage object (such as `csm44/sdb_freespace1`) where you want to add the segment, then click *Next*.
 - 6d Specify the size of the segment, the partition type (such as Linux LVM), click *Create*, then click *OK*.
 - 6e Click *Save*, then click *Save* again to save your changes.
- 7 Continue with [Section 11.2.4, “Creating an EVMS Volume,”](#) on page 167.

11.2.4 Creating an EVMS Volume

- 1 In *evmsgui*, click *Actions > Create > EVMS Volume*.
- 2 On the Create EVMS Volume page, select the container you just created (either the CSM container or the additional segment manager container), then specify a name for the volume (such as `shvol144`).
- 3 Click *Create*, then click *OK*.
- 4 Click *Save*, then click *Save* again to save your changes.
- 5 Click the *Volumes* tab to verify that the EVMS volume was created.
For example, a volume named `shvol144` would be listed as `/dev/evms/csm44/shvol144`.
- 6 Continue with [Section 11.2.5, “Making a File System on the EVMS Volume,”](#) on page 167.

11.2.5 Making a File System on the EVMS Volume

- 1 In *evmsgui*, click the *Disks* tab, then activate the CSM container:
 - 1a On the Disks page, right-click the CSM container, then select *Activate*.
 - 1b On the Activate page, select the CSM container, click *Activate*, then click *OK*.
 - 1c Click *Save*, then click *Save* again to save your changes.
- 2 Click the *Volumes* tab, then activate the EVMS volume:
 - 2a On the Volumes page, right-click the EVMS volume, then select *Activate*.
 - 2b On the Activate page, select the volume, click *Activate*, then click *OK*.
 - 2c Click *Save*, then click *Save* again to save your changes.
- 3 Make the file system on the EVMS volume:
 - 3a On the Volumes page, right-click the volume, then select *Make File System*.

- ♦ [Section 11.3.5, “Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource,” on page 172](#)
- ♦ [Section 11.3.6, “Configuring Policies for a Linux POSIX Volume Cluster Resource,” on page 173](#)

11.3.1 Logging in to iManager

When you create a new cluster resource, a Cluster Resource object is added in eDirectory. Log in to iManager as a user with the credentials necessary to extend the schema for this object.

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/manager.html`. Replace *server_ip_address* with the IP address or DNS name of an OES server in the cluster that has iManager installed or with the IP address for Apache-based services.
- 2 Specify your username and password.
- 3 Specify the tree where the cluster is located, then click *Login*.
- 4 Continue with [“Creating a Cluster Resource for a Linux POSIX Volume” on page 169](#).

11.3.2 Creating a Cluster Resource for a Linux POSIX Volume

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse and select the Cluster object of the cluster where you want to add a cluster resource.
When the page refreshes, a list of existing cluster resources is displayed.
- 3 Click *New*.
- 4 Specify *Resource* as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.
- 5 Specify the name of the resource you want to create.
This is the name of the resource for the cluster-enabled Linux POSIX volume.
You can use any name. You might apply a similar naming convention for Linux volumes as is used for NSS pools (*clustname_poolname_server*) by using the EVMS volume name. For example, `CLUS1_SHVOL44_SERVER`.
Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.
- 6 In the *Inherit From Template* field, browse to the Cluster container object, then locate and select the *Generic_FS_Template* in that container.
- 7 Deselect *Online on creation*.
You want to configure the cluster resource before you online it for the first time. Whenever you modify cluster resource settings, you must offline, then online the resource in order for the changes to take effect.
- 8 Select the *Define Additional Properties* check box, then click *Next*.
This takes you through several pages where you can define the load, unload, and monitoring scripts, then configure the cluster resource properties.
- 9 Continue with [“Configuring a Load Script for a Linux POSIX Volume Cluster Resource” on page 170](#).

11.3.3 Configuring a Load Script for a Linux POSIX Volume Cluster Resource

The resource load script specifies the commands to start the resource (including mounting the file system) on a server in the cluster. A load script is required for each Linux volume that you cluster-enable.

EVMS containers are the unit of failover for Linux POSIX volumes. Because the EVMS container is the unit of failover, all volumes in a container also fail over, but only the volumes that are mounted through the cluster resource load script are automatically mounted on failover. Any volumes in the container that are not mounted through the resource load script must be mounted manually.

The generic file system template you specified in [Step 6](#) when you created the cluster resource contains scripts that you must customize for your specific configuration. You need the following values for each of the scripts:

- ♦ a static IP address that you want to assign to the cluster resource
- ♦ the name of the Cluster Segment Manager container that you created in [Section 11.2.2, “Creating a Cluster Segment Manager Container,”](#) on page 165
- ♦ the EVMS volume name that you created in [Section 11.2.4, “Creating an EVMS Volume,”](#) on page 167
- ♦ the file system type you made on the EVMS volume in [Section 11.2.5, “Making a File System on the EVMS Volume,”](#) on page 167
- ♦ the Linux path of the mount point where you want to mount the EVMS volume

To continue configuring the new cluster resource:

- 1 On the Load Script page, edit or add the necessary commands to the script to load the resource on the server.

The load script from the generic file system template should appear similar to the following script. Modify the variable values in the definition lines. For an example of the customized load script, see [Section 11.5.1, “Sample Load Script for the Linux POSIX Volume Cluster Resource,”](#) on page 180.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, container name, file system type, and device.

If you later 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 IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=reiserfs
#define the container name
container_name=name
# define the device
```

```

MOUNT_DEV=/dev/evms/$container_name/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# activate the container
exit_on_error activate_evms_container $container_name $MOUNT_DEV
$NCS_TIMEOUT

# 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

```

- 2 (Optional) Modify the mount command to add mount options that are available for the file system.

Check the function `mount_fs` in the `/opt/novell/ncs/lib/ncsfuncs` file for additional mount options that are available for the file system you are using.

For example, the Ext3 file system supports mount options for access control lists (ACL) and extended attributes (USER_XATTR). You could add the following to the load script for a shared Ext3 volume to take advantage of those options:

```

#define mount options
MOUNT_OPTIONS=acl,user_xattr

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS $MOUNT_OPTIONS

```

- 3 If the path to the mount point does not exist on other nodes, you can add the following lines in the script before the line to mount the file system:

```

# create the mount point path when loading on a new node
ignore_error mkdir -p $MOUNT_POINT

```

- 4 Below the script editing area, specify the *Load 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 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.

- 5 Continue with [Section 11.3.4, “Configuring an Unload Script for a Linux POSIX Volume Cluster Resource,”](#) on page 171.

11.3.4 Configuring an Unload Script for a Linux POSIX Volume Cluster Resource

The cluster resource unload script specifies the commands to stop the resource (including unmounting the file system) on a server in the cluster. The unload script is required for each Linux POSIX volume that you cluster-enable.

The generic file system template you specified in [Step 6](#) when you created the cluster resource contains an unload script that you must edit to supply information specific to your file system resource. Use the same values for the defined variables that you used in the load script in [Section 11.3.3, “Configuring a Load Script for a Linux POSIX Volume Cluster Resource,”](#) on page 170.

To continue configuring the new cluster resource:

- 1 On the Unload Script page, edit or add the necessary commands to the script to unload or stop the resource on the server.

The unload script from the generic file system template should appear similar to the following script. Modify the variable values in the definition lines. For an example of the customized unload script, see [Section 11.5.2, “Sample Unload Script for the Linux POSIX Volume Cluster Resource,”](#) on page 180.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=reiserfs
# define the container name
container_name=name
# define the device
MOUNT_DEV=/dev/evms/$container_name/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# 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

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# deactivate the container
exit_on_error deactivate_evms_container $container_name $NCS_TIMEOUT

# return status

exit 0
```

- 2 Below the script editing area, specify the *Unload 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 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.

- 3 Continue with [Section 11.3.5, “Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource,”](#) on page 172.

11.3.5 Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource

The cluster resource monitor script specifies the commands to monitor the cluster resource.

The generic file system template you specified in [Step 6](#) when you created the cluster resource contains a monitor script that you must edit to supply information specific to your file system resource. Use the same values for the defined variables that you used in the load script in [Section 11.3.3, “Configuring a Load Script for a Linux POSIX Volume Cluster Resource,”](#) on page 170.

To continue configuring the new cluster resource:

- 1 On the Monitor Scripts page, select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

- 2 On the Monitor Scripts page, edit or add the necessary commands to the script to monitor resource on the server.

The monitor script from the generic file system template should appear similar to the following script. Modify the variable values in the definition lines. For an example of the customized monitor script, see [Section 11.5.3, “Sample Monitor Script for a Linux POSIX Volume Cluster Resource,”](#) on page 181.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=reiserfs
#define the container name
container_name=name
# define the device
MOUNT_DEV=/dev/evms/$container_name/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# 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
```

- 3 Below the script editing area, specify the *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 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..

- 4 Continue with [“Setting Start, Failover, and Failback Modes for Cluster Resources”](#) on page 131 to configure the cluster resource behavior for the shared storage you have cluster-enabled.

11.3.6 Configuring Policies for a Linux POSIX Volume Cluster Resource

You can configure the start, failover, and failback of cluster resources to happen manually or automatically.

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 *Assigned Nodes* list in the event of 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 *Assigned Nodes* list for the resource.

IMPORTANT: Resources only fail back to the first node in their *Assigned Nodes* list. So 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 will not failback to that second server.

Resources do not automatically move from node to node just because a node higher in the *Assigned Nodes* list rejoins the cluster, unless the *Failback* mode is set to AUTO and the first node in the *Assigned Nodes* list rejoins the cluster.

To continue configuring a new cluster resource:

- 1** (Conditional) 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.
- 2** (Conditional) 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 Novell 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 *Assigned Nodes* list as soon as any server in the list is brought online.
- 3** 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.
- 4** Click *Next*.
- 5** Assign nodes to the resource as described in [Section 9.9, "Assigning Nodes to a Resource,"](#) on page 133.

- 6 Configure resource priorities for load order as described in [Section 9.10, “Configuring Resource Priorities for Load Order,”](#) on page 134.
- 7 (Optional) Continue with [Section 11.4, “Creating a Virtual Server Name for the Cluster Resource,”](#) on page 175.

11.4 Creating a Virtual Server Name for the Cluster Resource

NOTE: The procedures described in this section are the same as those found in [“Creating a Virtual NCP Server Object for NCP Volumes”](#) in the *OES 2 SP2: NCP Server for Linux Administration Guide*. If you create a virtual server object now, you can use the virtual server name if you set up NCP volumes later. You don’t need to re-create the virtual server name.

Creating a virtual server name for the cluster resource allows users to access the resource via an assigned name instead of by its IP address. The virtual server name is stored in Novell eDirectory in an `NCS:NCP Server` object under the Cluster object where you created the resource. The `NCS:NCP Server` object is not created automatically when you create a Linux POSIX volume as it is with clustered NSS pools. After you cluster-enable the Linux POSIX volume, you can create a new `NCS:NCP Server` object where you specify a virtual server name for the resource.

After you create the `NCS:NCP Server` object, you must modify the load script to identify the name of this virtual server and bind it to the IP address of the cluster resource. You must modify the unload script to identify the name of this virtual server and unbind it from the IP address of the cluster resource.

You can use any name as the virtual server name. You might apply a similar naming convention for Linux volumes as is used for NSS pools (`clustername_poolname_server`) by using the EVMS volume name or the NCP volume name. For example, `CLUS1_SHVOL44_SERVER` or `CLUS1_USERS_SERVER`.

IMPORTANT: Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

For example, when you cluster-enable an NSS pool, a virtual NCP Server object (`NCS:NCP Server`) is automatically created in Novell eDirectory and given the name of the Cluster object plus the name of the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`.

IMPORTANT: The following procedure uses a similar naming convention as that used for NSS; however, you can use a different naming convention.

Use one of the two methods in this section to create the `NCS:NCP Server` object for the volume to give it a virtual server name and modify the load and unload scripts.

- ♦ [Section 11.4.1, “Using `ncs_ncpserv.py` to Create an NCS:NCP Server Object,”](#) on page 176
- ♦ [Section 11.4.2, “Using iManager to Create an NCS:NCP Server Object,”](#) on page 177

11.4.1 Using `ncs_ncpserv.py` to Create an NCS:NCP Server Object

The script `/opt/novell/ncs/bin/ncs_ncpserv.py` is available to help you create the NCS:NCP Server object for your Linux POSIX volume. The `ncs_ncpserv.py` script creates a virtual NCP Server object (NCS:NCP Server) in eDirectory under the Cluster object, and associates it with none, one, or multiple volumes that you specify. This allows clients to access files on the Linux POSIX volume via the virtual server name. Commands are automatically added to the load script to bind the specified volumes to the IP address of the cluster resource. Commands are automatically added to the unload script to unbind the volumes from the IP address of the cluster resource.

- ♦ [“Syntax” on page 176](#)
- ♦ [“Examples” on page 176](#)
- ♦ [“Procedure” on page 176](#)

Syntax

At the terminal console prompt on the master node of the cluster, enter the following command as the root user:

```
./ncs_ncpserv.py -c ncs_ncp_server_name -i ip_address [-v <volumename |  
"volumename:volumename:volumename..."> ]
```

Replace the *ncs_ncp_server_name*, *ip_address*, and *volumename* with the information for your particular solution. 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.

If the `-v` option is not specified, all of the volumes that currently exist on the EVMS volume 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.

Examples

For example, to include all of the volumes on the SHVOL44 cluster resource where the IP address is 10.10.10.44, enter

```
./ncs_ncpserv.py -c CLUS1_SHVOL44_SERVER -i 10.10.10.44
```

For example, to specify a single volume named USERS on the SHVOL44 cluster resource where the resource IP address is 10.10.10.44, enter

```
./ncs_ncpserv.py -c CLUS1_SHVOL44_SERVER -i 10.10.10.44 -v USERS
```

For example, to specify multiple volumes on the SHVOL44 cluster resource where the resource IP address is 10.10.10.44, enter

```
./ncs_ncpserv.py -c CLUS1_SHVOL44_SERVER -i 10.10.10.44 -v  
"USERS:MY_NNCP_VOL06:cn=CLUS_02_MY_NNCP_VOL07,o=novell"
```

Procedure

- 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 terminal console prompt, enter

```
./ncs_ncpserv.py -c ncs_ncp_server_name -i ncp_ip_address -v  
ncp_volumename
```

Replace the `ncp_server_name`, `ip_address`, and `ncp_volumename` with the information for your particular solution.

For example, to include the NCP volume `USERS` on the `SHVOL44` cluster resource where the IP address is `10.10.10.44`, enter

```
./ncs_ncpserv.py -c CLUS1_SHVOL44_SERVER -i 10.10.10.44 -v USERS
```

4 In iManager, view the load script for the Linux POSIX volume cluster resource to verify that the load script was automatically updated with lines that define and bind the virtual NCP server name and its IP address.

For example, if you named the virtual server's Server object as `CLUS1_SHVOL44_SERVER`, and the IP address of the cluster resource is `10.10.10.44`, you should see the following new lines in the script:

```
# define NCP server name  
NCP_SERVER=CLUS1_SHVOL44_SERVER  
# define resource IP address  
RESOURCE_IP=10.10.10.44  
  
# bind the NCP volume  
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --  
ipaddress=$RESOURCE_IP
```

5 In iManager, view the unload script for the Linux POSIX volume cluster resource to verify that the unload script was automatically updated with lines that define and unbind the virtual NCP server name and its IP address.

For example, if you named the virtual server's Server object as `CLUS1_SHVOL44_SERVER`, and the IP address of the cluster resource is `10.10.10.44`, you should see the following new lines in the script:

```
# define NCP server name  
NCP_SERVER=CLUS1_SHVOL44_SERVER  
# define resource IP address  
RESOURCE_IP=10.10.10.44  
  
# unbind the NCP volume  
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --  
ipaddress=$RESOURCE_IP
```

11.4.2 Using iManager to Create an NCS:NCP Server Object

1 Create an NCS:NCP Server object in eDirectory with the name you want to use for the cluster resource.

1a In iManager, use one of the following methods to open the Create Object page:

- ◆ In the *Roles and Tasks* view, select *Directory Administration > Create Object*.
- ◆ In the *View Objects > Tree* view, expand the context to locate the Cluster object in the tree, click the Cluster object name, then select *New > Create Object*.

The Create Object page opens in a dialog box.

- ◆ In the *View Objects > Browse* view, browse to locate the Cluster object, click the Cluster object name, then select *Create Object* from the pop-up menu.

- 1b** Select *Show all object classes* so that the virtual server object (NCS:NCP Server) is available.
- 1c** From the list of *Available Object Classes*, select *NCS:NCP Server*, then click *OK*.
- 1d** Specify the *NCP Server Name*, such as `CLUS1_SHVOL44_SERVER`.
- 1e** Under the *Context to use for the server*, browse to select the Cluster object context such as `cluster.novell`.
- 1f** Click *OK*.

The object is created in the eDirectory tree where you are logged in, and a confirmation message acknowledges the successful creation.

For example, if you are logged in to MYTREE, the distinguished name of the new NCP server is `CLUS1_SHVOL44_SERVER.cluster.novell.MYTREE`.

- 2** Verify that the NCS:NCP Server object was created. In the *View Objects > Tree* view, expand the context to locate the Cluster object in the tree, click the Cluster object name, then locate the NCP Server object by the name you gave it (such as `CLUS1_SHVOL44_SERVER`) in the list of Cluster objects on the right.

- 3** In iManager, modify the load script for the Linux POSIX volume cluster resource by adding a line that binds the virtual NCP server name to the IP address.

For example, if you named the virtual server's Server object as `CLUS1_SHVOL44_SERVER`, and the IP address of the cluster resource is 10.10.10.44, add the following lines:

```
# define NCP server name
NCP_SERVER=CLUS1_SHVOL44_SERVER
# define resource IP address
RESOURCE_IP=10.10.10.44

# bind the NCP volume
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

- 4** In iManager, modify the unload script for the Linux POSIX volume cluster resource by adding a line that unbinds the virtual NCP server name from the IP address.

For example, if you named the virtual server's Server object as `CLUS1_SHVOL44_SERVER`, and the IP address of the cluster resource is 10.10.10.44, add the following lines:

```
# define NCP server name
NCP_SERVER=CLUS1_SHVOL44_SERVER
# define resource IP address
RESOURCE_IP=10.10.10.44

# unbind the NCP volume
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

11.5 Sample Scripts for a Linux POSIX Volume Cluster Resource

The scripts in this section are based on the sample values in the following table. Make sure to replace the sample values with the ones you used in your solution.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, container name, file system type, and device.

If you later 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 IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110.

Variable	Template Value	Sample Value	Description
RESOURCE_IP	a.b.c.d	10.10.10.44	IP address of the virtual cluster server for this cluster resource
container_name	name	csn44	The name you gave to the cluster segment manager.
MOUNT_DEV	/dev/evms/ \$container_name/ volume_name	/dev/evms/ \$container_name/ shvol44	The Linux path for the EVMS volume you created, such as shvol44.
MOUNT_FS	reiserfs	ext3	The file system type you made on the EVMS volume.
MOUNT_POINT	/mnt/mount_point	/mnt/shvol44	The mount location for the EVMS volume you created. This example shows a mount location with a directory named the same as the EVMS volume name (shvol44). You can mount the EVMS volume anywhere.

- ♦ [Section 11.5.1, “Sample Load Script for the Linux POSIX Volume Cluster Resource,”](#) on page 180
- ♦ [Section 11.5.2, “Sample Unload Script for the Linux POSIX Volume Cluster Resource,”](#) on page 180
- ♦ [Section 11.5.3, “Sample Monitor Script for a Linux POSIX Volume Cluster Resource,”](#) on page 181

11.5.1 Sample Load Script for the Linux POSIX Volume Cluster Resource

The following is an example load script for the cluster resource for a Linux POSIX volume:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/shvol44
# define the mount point
MOUNT_POINT=/mnt/shvol44

# activate the container
exit_on_error activate_evms_container $container_name $MOUNT_DEV $NCS_TIMEOUT

# mount the file system
ignore_error mkdir -p $MOUNT_POINT
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
```

11.5.2 Sample Unload Script for the Linux POSIX Volume Cluster Resource

The following is an example unload script for the cluster resource for a Linux POSIX volume:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/shvol44
# define the mount point
MOUNT_POINT=/mnt/shvol44

# 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

# del the IP address
```

```

ignore_error del_secondary_ipaddress $RESOURCE_IP

# deactivate the container
ignore_error deactivate_evms_container $container_name $NCS_TIMEOUT

# return status
exit 0

```

11.5.3 Sample Monitor Script for a Linux POSIX Volume Cluster Resource

The following is an example monitor script for the cluster resource for a Linux POSIX volume:

```

#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/shvol44
# define the mount point
MOUNT_POINT=/mnt/shvol44

# 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

# (optional) status of the eDirectory service
#exit_on_error rcnstd status

# (optional) status of the Linux User Management service
# exit_on_error namcd status

exit 0

```

11.6 Expanding EVMS Volumes on Shared Disks

As your storage needs increase, it might become necessary to add more disk space or drives to your shared storage system. EVMS provides features that allow you to expand or move existing volumes.

The two supported methods for creating additional space for an existing volume are:

- ◆ [Section 11.6.1, “Expanding a Volume to a Separate Disk,” on page 181](#)
- ◆ [Section 11.6.2, “Moving a Volume to a Larger Disk,” on page 182](#)

11.6.1 Expanding a Volume to a Separate Disk

- 1 Enter `evmsgui` at a terminal console prompt as the `root` user.
- 2 In `evmsgui`, unmount the file system for the shared Linux POSIX volume you want to expand.

- 3 In *evmsgui*, click the *Volumes* tab, right-click the volume you want to expand then select *Add Feature*.
- 4 Select *Drive Linking Feature*, then click *Next*.
- 5 Provide a name for the drive link, click *Add*, then save your changes.
- 6 Click *Actions*, select *Create*, and then select *Container*.
- 7 Select the *Cluster Segment Manager*, click *Next*, then select the disk you want to expand the volume to.
The entire disk is used for the expansion, so you must select a disk that does not have other volumes on it.
- 8 Provide the same settings information (same name, type (Private), owning node, and so on) as the existing container for the volume, then save your changes.
- 9 Click the *Volumes* tab, right-click the volume, then click *Expand*.
- 10 Select the volume that you are expanding, then click *Next*.
- 11 Verify the current volume size and the size of the volume after it is expanded, then click *Next*.
The expanded volume size should include the size of the disk the volume is being expanded to.
- 12 Select the storage device the volume is being expanded to, select *Expand*, and save your changes.
- 13 Click *Save* twice, then exit *evmsgui*.

11.6.2 Moving a Volume to a Larger Disk

- 1 Enter *evmsgui* at a terminal console prompt as the *root* user.
- 2 In *evmsgui*, unmount the file system for the volume you want to move by right-clicking the volume and selecting *Unmount*.
- 3 Add a larger disk to the CSM container:
 - 3a In *evmsgui*, click *Actions*, select *Create*, then click *Container*.
 - 3b Select the *Cluster Segment Manager*, then click *Next*.
 - 3c Select the larger disk you want to move the volume to.
The entire disk is used for the expansion, so you must select a disk that does not have other volumes on it.
 - 3d Provide the same settings information (name, type (Private), owning node, and so on) as the existing container for the volume, then save your changes.
 - 3e Click *Save* and exit *evmsgui*.
- 4 Restart *evmsgui*, click the *Containers* tab, then expand the container so that the objects under the container appear.
The new disk should appear as part of the container.
- 5 Right-click the object for the disk where the volume resides and select *Replace*.
- 6 Select the object for the disk where the volume will be moved, then click *Next*.
- 7 Save your changes.
Saving your changes could take a while, depending on volume size and other factors.
- 8 Click *Save*, exit *evmsgui*, then restart *evmsgui*.

- 9 Click the *Volumes* tab, right-click the volume, then select *Check/Repair filesystem*.
This runs the repair process and ensures that no problems exist on the moved volume.
- 10 Click the *Disks* tab, right-click the disk the volume was moved from, then select *Remove from container*.
- 11 Click *Save* twice, then exit *evmsgui*.

11.7 Deleting Shared Storage

Make sure to offline the cluster resource before attempting to delete either the cluster resource or the shared storage resource.

WARNING: If you attempt to delete a cluster resource or the shared storage without first offlining it, deletion errors occur, and the data associated with the shared storage is not recoverable.

11.8 Known Issues for Working with Cluster Resources for Linux POSIX Volumes

- ♦ [Section 11.8.1, “Dismount Volumes before Onlining a Comatose Resource,” on page 183](#)
- ♦ [Section 11.8.2, “Cluster Services Must Be Running When Using EVMS,” on page 183](#)
- ♦ [Section 11.8.3, “Close EVMS Utilities When They Are Not In Use,” on page 183](#)
- ♦ [Section 11.8.4, “Do Not Migrate Resources When EVMS Tools Are Running,” on page 184](#)

11.8.1 Dismount Volumes before Onlining a Comatose Resource

If a Linux POSIX volume resource goes comatose, you must verify that the volume is not mounted on any node in the cluster before bringing the resource online again. Failure to do so might result in corruption.

WARNING: To avoid corruption, make sure the Linux POSIX volume in a comatose resource is dismounted from any node before attempting to online it.

11.8.2 Cluster Services Must Be Running When Using EVMS

Novell Cluster Services must be running on all nodes in the cluster whenever you make any modifications to the cluster or cluster resources by using the EVMS management tools or the Clusters plug-in in iManager.

11.8.3 Close EVMS Utilities When They Are Not In Use

The EVMS management utilities (*evms*, *evmsn*, *evmsgui*) lock the EVMS engine while running, potentially blocking other EVMS-related actions from taking place. This affects NSS pool and volume actions and Linux POSIX volume actions.

Do not run the EVMS management utilities when they are not being actively used. Do not concurrently open multiple instances of the utilities. Do not open the utilities when using NSSMU or the Clusters or Storage plug-ins for iManager.

11.8.4 Do Not Migrate Resources When EVMS Tools Are Running

Cluster resources for NSS pools and Linux POSIX volumes should not be migrated while the EVMS management utilities (evms, evmsn, and evmsgui) are running.

11.9 What's Next

Complete the following tasks to configure the behavior of the cluster resource you just created:

- ❑ [“Assigning Nodes to a Resource” on page 133](#)
- ❑ [“Configuring Resource Priorities for Load Order” on page 134](#)

For information on managing the cluster, see [Chapter 8, “Managing Clusters,” on page 101](#).

IMPORTANT: You can create an NCP volume on the shared Linux POSIX volume and cluster-enable the shared NCP volume. For information, see [“Configuring NCP Volumes with Novell Cluster Services”](#) in the *OES 2 SP2: NCP Server for Linux Administration Guide*.

Configuring Novell Cluster Services in a Xen Virtualization Environment

12

Novell Cluster Services is installed and configured in a virtualization environment (on guest servers) by using the same methods and processes as those used on a physical server. No additional changes or special configuration are required.

You can install Novell Cluster Services on the virtualization host server and create cluster resources that contain the virtual machines. You can fail over or migrate virtual machine cluster resources (entire virtual machines) to different physical nodes in your cluster. Resource templates are provided for Xen virtualization environments.

IMPORTANT: All templates except Xen and XenLive are valid in guest servers (DomU) in the virtualization environment. Only the Xen and XenLive templates can be used in the OES 2 Linux (Xen) host environment (that is, in Dom0, but not in DomU).

Although many different cluster virtualization scenarios are possible, only those outlined in the sections below have been tested:

- ♦ [Section 12.1, “Prerequisites for Xen Host Server Environments,” on page 185](#)
- ♦ [Section 12.2, “Virtual Machines as Cluster Resources,” on page 186](#)
- ♦ [Section 12.3, “Virtual Machines as Cluster Nodes,” on page 193](#)
- ♦ [Section 12.4, “Virtual Cluster Nodes in Separate Clusters,” on page 194](#)
- ♦ [Section 12.5, “Mixed Physical and Virtual Node Clusters,” on page 195](#)
- ♦ [Section 12.6, “Additional Information,” on page 197](#)

12.1 Prerequisites for Xen Host Server Environments

When using Novell 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 Novell 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 Novell 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)” (<http://www.novell.com/support/viewContent.do?externalId=7004595&sliceId=1>) on Novell Support (<http://www.novell.com/support>).

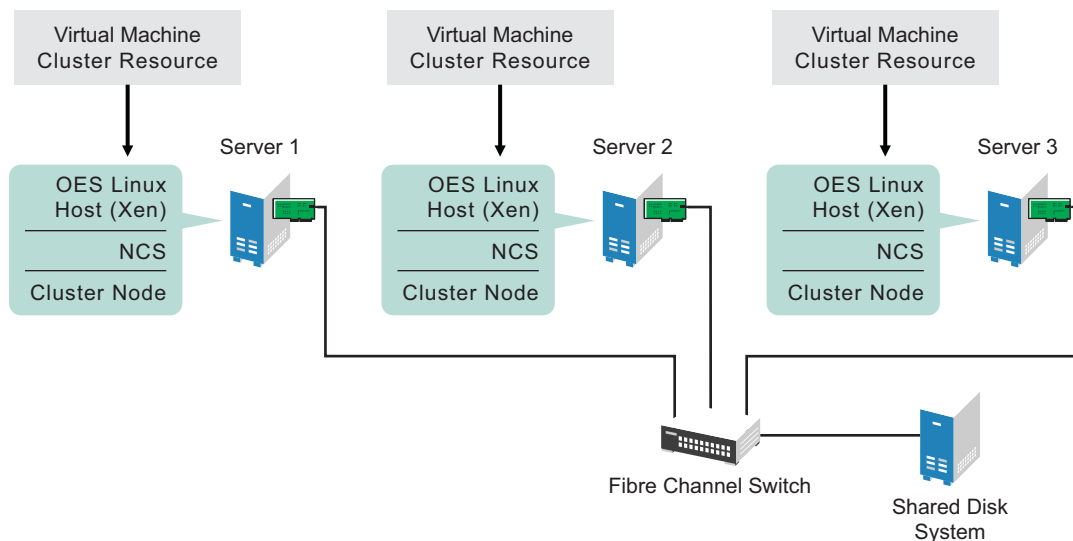
12.2 Virtual Machines as Cluster Resources

In this scenario, you have an OES 2 Linux cluster configured on physical machines. OES 2 Linux and Xen are installed and configured on each node along with Novell Cluster Services. This part of the Novell Cluster Services configuration does not differ from that of an OES 2 Linux cluster without virtualization.

You then create either Linux or NetWare 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 12-1 depicts how this setup might look. Novell Cluster Services (NCS) is installed and running on the virtual machine (VM) host server.

Figure 12-1 Virtual Machines as Cluster Resources



The following sections describe how to create a cluster resource and its cluster scripts for each virtual machine:

- ◆ [Section 12.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 187](#)
- ◆ [Section 12.2.2, “Configuring Virtual Machine Load, Unload, and Monitor Scripts,” on page 188](#)
- ◆ [Section 12.2.3, “Setting Up Live Migration,” on page 193](#)

12.2.1 Creating a Xen Virtual Machine Cluster Resource

Novell 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.

Make sure that your Xen setup is working properly before you attempt to set up the Novell Cluster Services clustering for your virtual machines in the Xen host environment. Refer to the *Virtualization with Xen* (http://www.novell.com/documentation/sles10/book_virtualization_xen/data/book_virtualization_xen.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 username and password, specify the tree where you are installing the cluster, then click *Login*.
- 3 In iManager, select *Clusters*, then click *Cluster Options*.
Under *Clusters*, iManager displays four links that you can use to configure and manage your cluster.
- 4 Browse and select the cluster name.
- 5 On the Cluster Options page, click *New*.
- 6 Click the *Resource* radio button to specify *Resource* as the resource type you want to create, then click *Next*.
- 7 Specify a name for the virtual machine resource.
- 8 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`.
- 9 Select the *Define Additional Properties* check box, click *Next*, then continue with “[Configuring Virtual Machine Load, Unload, and Monitor Scripts](#)” on page 188.

12.2.2 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 filename 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. For information, see [Section 9.8, “Setting Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 131 and [Section 9.9, “Assigning Nodes to a Resource,”](#) on page 133.

If you are modifying the scripts, continue with the following sections:

- ◆ [“Configuring the Load Script”](#) on page 188
- ◆ [“Configuring the Unload Script”](#) on page 189
- ◆ [“Configuring the Monitor Script”](#) on page 191

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 188
- ◆ [“Sample XenLive_Template Load Script”](#) on page 189

- 2 Click *Next* and continue with [“Configuring the Unload Script”](#) on page 189.

Sample Xen_Template Load Script

The Xen_Template load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns

# filesystem settings
export OCF_RESKEY_device=/dev/evms/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=
```

```

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

# 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

# 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 190](#)
 - ◆ [“Sample XenLive_Template Unload Script” on page 190](#)
- 2 Click *Next*, then continue the setup by configuring the resource policies and the resource server assignments.

For information, see [Section 9.8, “Setting Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 131 and [Section 9.9, “Assigning Nodes to a Resource,”](#) on page 133.

- 3 If you want to enable monitoring for the resource, continue with [“Configuring the Monitor Script”](#) on page 191.

Sample Xen_Template Unload Script

The Xen_Template unload script appears similar to the following example:

```
#!/bin/bash
./opt/novell/ncs/lib/ncsfuncs

# filesystem settings
export OCF_RESKEY_device=/dev/evms/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# stop the service
ignore_error ocf_stop Xen

# umount the file system
ignore_error ocf_stop Filesystem

# 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

# 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
fi
```

```

    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 monitoring 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 monitoring script.

- ♦ [“Enabling Resource Monitoring” on page 191](#)
- ♦ [“Viewing or Modifying the Monitor Script” on page 192](#)
- ♦ [“Sample Xen_Template Monitor Script” on page 192](#)
- ♦ [“Sample XenLive_Template Monitor Script” on page 192](#)

Enabling Resource Monitoring

To enable resource monitoring for a virtual machine cluster resource:

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse and select the Cluster object.
- 3 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 monitoring 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, or to migrate to another server if the failure action initiates.

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 *Assigned 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 *Assigned Nodes* list.

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, click *Clusters*, then click *Cluster Options*.
- 2 Browse and select the Cluster object.
- 3 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 192](#)
 - ♦ [“Sample XenLive_Template Monitor Script” on page 192](#)
- 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 191](#). 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/ncsfuns

# filesystem settings
export OCF_RESKEY_device=/dev/evms/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# 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

# 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
```

12.2.3 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.

Make sure that your OCFS2 file system is working properly before you attempt to use it with Novell Cluster Services.

An overview of OCFS2 is available in “Oracle Cluster File System 2” (http://www.novell.com/documentation/sles10/sles_admin/data/cha_ocfs2.html) in the *SUSE Linux Enterprise Server 10 SP2 Administration Guide*. For detailed information about using OCFS2, see the *OCFS2 Project* (<http://oss.oracle.com/projects/ocfs2/>) on the Oracle Web site.

For information about setting up live migration, see *Configuring a Xen VM for Live Migration with a Cluster* (<http://www.novell.com/coolsolutions/feature/19676.html>) in *Novell Cool Solutions* (<http://www.novell.com/communities/coolsolutions>).

12.3 Virtual Machines as Cluster Nodes

In this scenario, you have OES 2 Linux (Xen) installed and configured on each node (physical machine). You then create either a NetWare or a Linux virtual machine on each physical machine and install and configure Novell Cluster Services on each virtual machine. The combined virtual machines (cluster nodes) comprise one cluster.

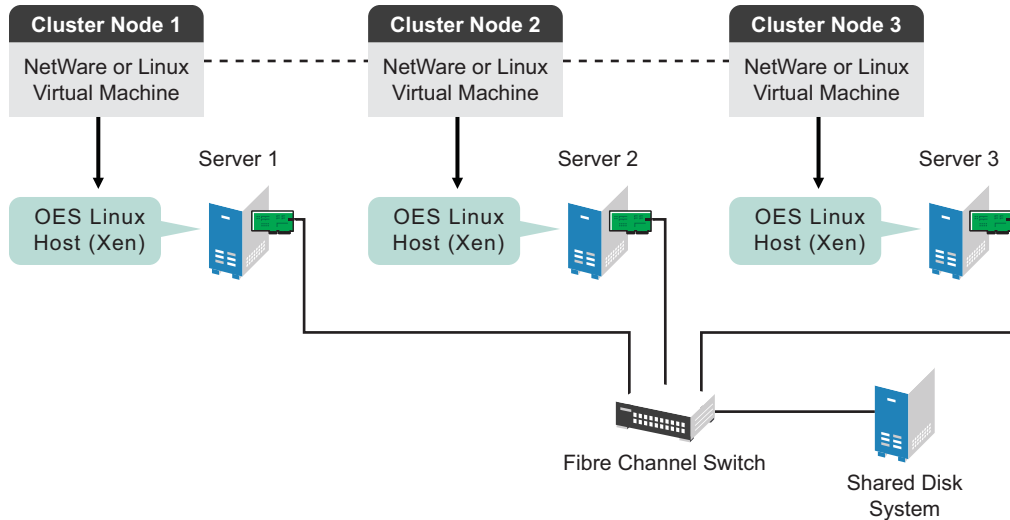
IMPORTANT: All virtual cluster nodes in the same cluster should be either Linux or NetWare. Do not mix Linux and NetWare cluster nodes in the same cluster.

You can then create and configure cluster resources on each virtual cluster node. 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 12-2 depicts using virtual machines as cluster nodes.

Figure 12-2 *Virtual Machines as Cluster Nodes*



12.4 Virtual Cluster Nodes in Separate Clusters

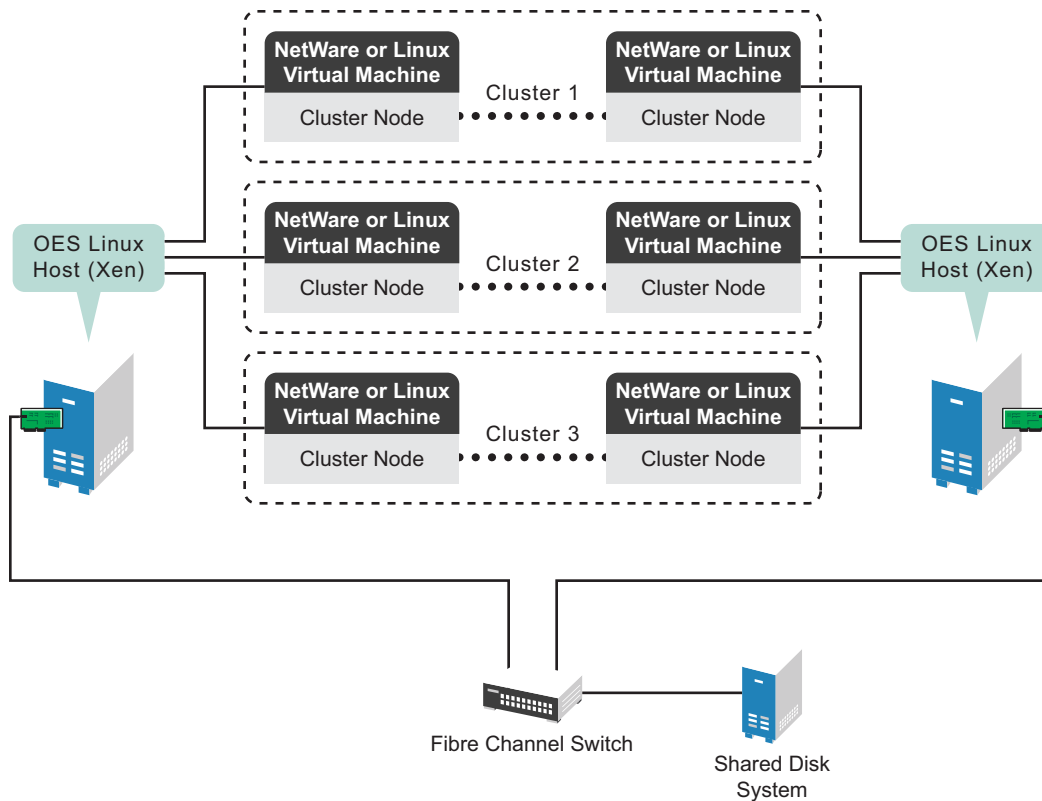
In this scenario, you have OES 2 Linux (Xen) installed and configured on each node (physical machine). You then create multiple NetWare or Linux virtual machines on each physical machine and install and configure Novell Cluster Services on each virtual machine. During the Novell Cluster Services installation, 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.

IMPORTANT: All virtual cluster nodes in the same cluster should be either Linux or NetWare. Do not mix Linux and NetWare cluster nodes in the same cluster.

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 12-3 depicts using virtual cluster nodes in separate clusters.

Figure 12-3 Virtual Cluster Nodes in Separate Clusters



12.5 Mixed Physical and Virtual Node Clusters

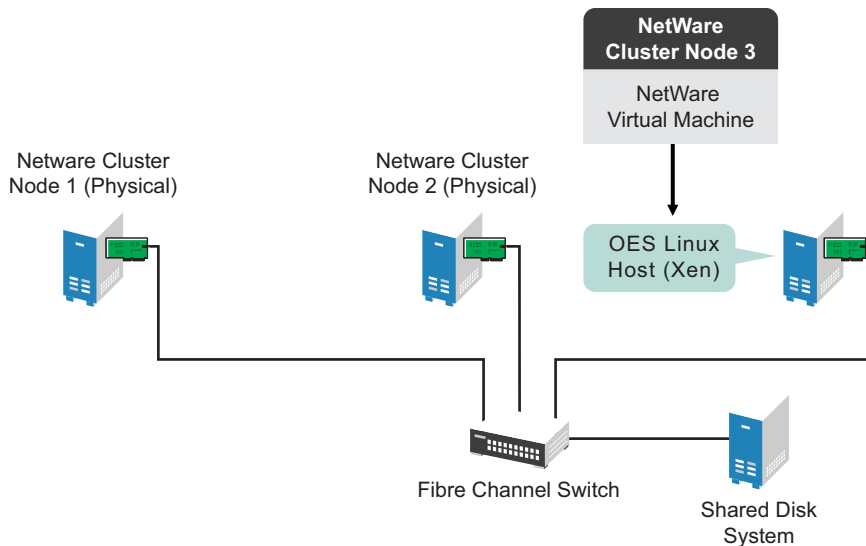
This is a temporary scenario that is used for upgrading cluster nodes or converting clusters from physical to virtual cluster nodes.

This can be done through several different methods. One method is to add a virtual NetWare cluster node to an existing physical NetWare cluster. To do this, you install an OES 2 Linux (Xen) server (physical machine). You then create a NetWare virtual machine on the physical machine and install and configure Novell Cluster Services on the NetWare virtual machine. During the Novell Cluster Services installation, you add the NetWare virtual cluster node to your existing NetWare cluster (NetWare physical nodes).

You can then migrate the desired resources from NetWare physical cluster nodes to the NetWare virtual cluster node. This lets you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical NetWare cluster nodes with virtual NetWare cluster nodes.

Figure 12-4 depicts how this setup might look.

Figure 12-4 *Mixed Physical and Virtual Node Cluster*

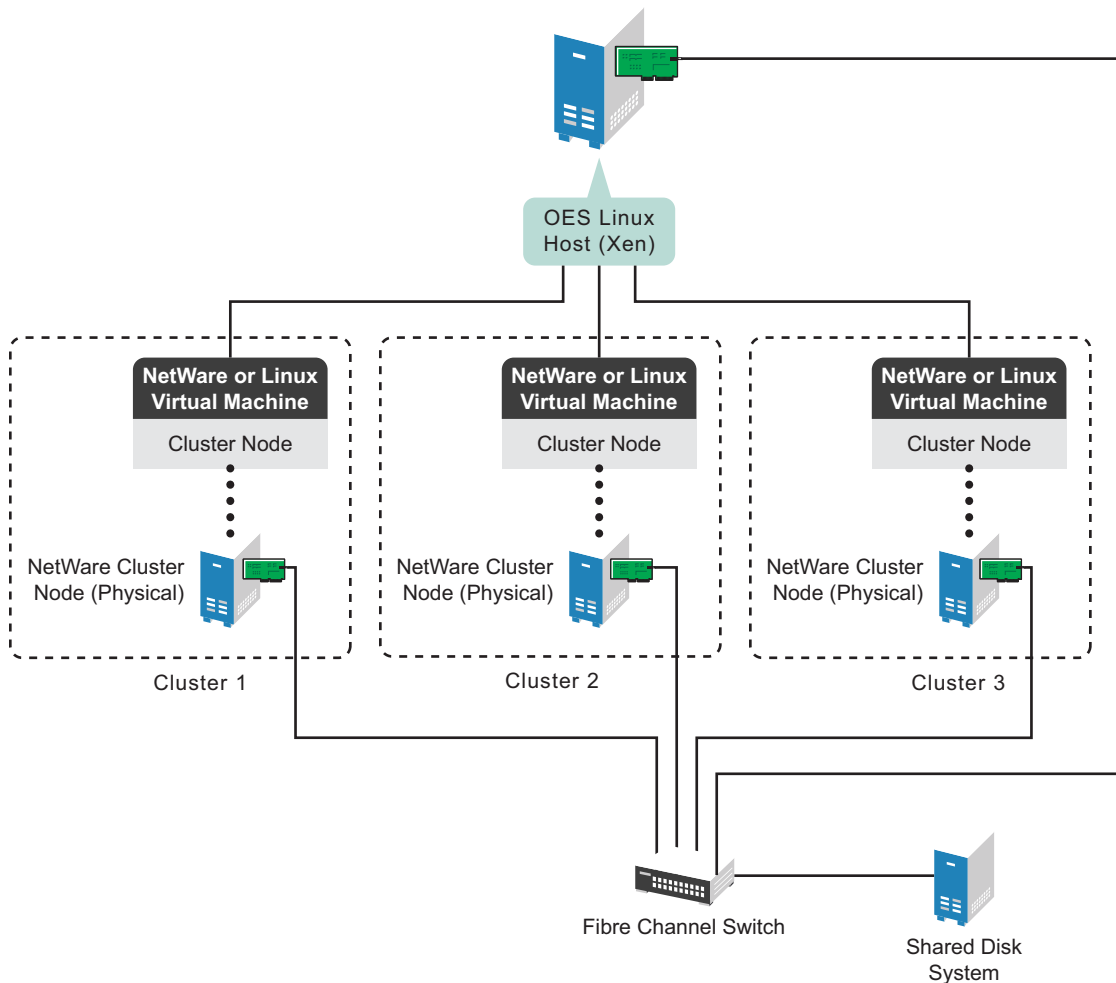


Another method is to install Novell Cluster Services on physical NetWare nodes and create a separate cluster for each node. You then install an OES 2 Linux (Xen) server (physical machine) and create NetWare virtual machines and install Novell Cluster Services on each virtual machine. You can then add one virtual NetWare 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 NetWare cluster node to the NetWare 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 NetWare cluster nodes in each cluster with virtual NetWare cluster nodes.

Figure 12-5 depicts how this setup might look.

Figure 12-5 *Separate Mixed Physical and Virtual Node Clusters*



12.6 Additional Information

To get started with virtualization, see “[Introduction to Xen Virtualization](http://www.novell.com/documentation/sles10/book_virtualization_xen/data/sec_xen_basics.html)” (http://www.novell.com/documentation/sles10/book_virtualization_xen/data/sec_xen_basics.html) in the *Virtualization with Xen* (http://www.novell.com/documentation/sles10/book_virtualization_xen/data/book_virtualization_xen.html) guide.

For information on setting up virtualized NetWare, see “[Installing, Upgrading, or Updating OES on a Xen-based VM](#)” in the *OES 2 SP2: Installation Guide* guide.

For information on setting up virtualized OES 2 Linux, see “[Installing, Upgrading, or Updating OES on a Xen-based VM](#)” in the *OES 2 SP2: Installation Guide* guide.

Troubleshooting Novell Cluster Services

13

This section describes known issues and frequently asked questions for managing Novell Cluster Services.

- ◆ [Section 13.1, “File Location Information,” on page 199](#)
- ◆ [Section 13.2, “Version Issues,” on page 200](#)
- ◆ [Section 13.3, “Diagnosing a Poison Pill Condition,” on page 200](#)
- ◆ [Section 13.4, “Diagnosing Cluster Problems,” on page 200](#)
- ◆ [Section 13.5, “Cluster Search Times Out \(Bad XML Error\),” on page 201](#)
- ◆ [Section 13.6, “A Device Name Is Required to Create a Cluster Partition,” on page 202](#)
- ◆ [Section 13.7, “Cluster Resource Goes Comatose Immediately After Migration or Failover,” on page 202](#)
- ◆ [Section 13.8, “smdr.novell Is Not Registered with SLP for a New Cluster Resource,” on page 202](#)
- ◆ [Section 13.9, “Cluster View Displays the Wrong Cluster Node Name,” on page 202](#)
- ◆ [Section 13.10, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted \(Linux\),” on page 203](#)
- ◆ [Section 13.11, “Could Not Delete This Resource Data_Server \(Error 499\),” on page 203](#)
- ◆ [Section 13.12, “Problem Authenticating to Remote Servers during Cluster Configuration,” on page 203](#)
- ◆ [Section 13.13, “Problem Connecting to an iSCSI Target,” on page 203](#)
- ◆ [Section 13.14, “Problem Deleting a Cluster Resource or Clustered Pool,” on page 204](#)
- ◆ [Section 13.15, “Can’t Find the Prevent Cascading Failover Option,” on page 204](#)
- ◆ [Section 13.16, “Is there a way to uninstall Novell Cluster Services from a server?,” on page 204](#)
- ◆ [Section 13.17, “Cluster Resource Is Stuck in “NDS Sync” or “eDirectory Sync” State,” on page 204](#)
- ◆ [Section 13.18, “Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems,” on page 205](#)
- ◆ [Section 13.19, “Nodes Reboot Randomly after January or March 2010 Updates,” on page 205](#)

13.1 File Location Information

Information about Novell Cluster Services scripts and configuration can be found in the following locations:

Information	File Location
Load scripts	<code>/var/opt/novell/ncs</code>
Unload scripts	
Cluster Services configuration files	
Load and unload script output files (.out)	<code>/var/opt/novell/log/ncs</code>
Most resource load and unload issues can be solved with the information in the .out files.	
Cluster Services scripts, such as <code>ldncls</code> and <code>unldncls</code>	<code>/opt/novell/ncs/bin</code>
Resource templates	<code>/opt/novell/ncs/templates</code>

13.2 Version Issues

Knowing the location and purpose of the files that make up Novell Cluster Services can be useful in helping you troubleshoot problems and resolve version issues. For information, see [Appendix B, “Files for Novell Cluster Services,”](#) on page 215.

13.3 Diagnosing a Poison Pill Condition

A poison pill is given to a node if it stops sending out the Novell 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 Novell 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

13.4 Diagnosing Cluster Problems

To help Novell 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 *Clusters > Cluster Manager > Run Report* option in 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?

13.5 Cluster Search Times Out (Bad XML Error)

If you are using in Novell eDirectory 8.7.3x, time-outs are possible when you search from iManager for eDirectory objects (such as NCP Server objects, Volume objects, and Cluster objects) because the Object Class attribute is not indexed by default. The LDAP sub-tree search can take over 30 seconds, which causes the query to time out. For example, a Cluster object search from the Cluster Options page returns the error:

```
Bad XML found during parsing when accessing cluster options
```

We recommend that you create a value index on the objects' Object Class attribute. (Object Class is considered an attribute for indexing purposes.) This helps to reduce the time needed for the sub-tree search from over 30 seconds to 10 to 50 milliseconds. For instructions, see [“Creating an Index”](#) in the *Novell eDirectory 8.8 Administration Guide*.

Building indexes speeds up the sub-tree search, even if some partitions being searched do not contain these types of objects. For example, searching for a Cluster object in a context that contains only users is not expected to return results; however, the Object Class search is still performed, and benefits from having an index present.

The sub-tree search performance issue is resolved in the eDirectory 8.8.x release with the addition of the AncestorID feature.

13.6 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 install a shared-disk cluster by entering a device name for the cluster Split Brain Detector (SBD) partition at cluster creation (new cluster) time. If you don't enter a device name, you won't be able to cluster-enable NSS pools. On Linux, names are case sensitive.

13.7 Cluster Resource Goes Comatose Immediately After Migration or Failover

When the SLP daemon (`sldap`) 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.

For information, see [“SLP” on page 35](#).

13.8 `smdr.novell` Is Not Registered with SLP for a New Cluster Resource

You might get an error after creating 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. Restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource. and advertise it correctly.

- 1 Log in to the server as the `root` user, open a terminal console, then enter
`novell-smdrd restart`

13.9 Cluster View Displays the Wrong Cluster Node Name

In OES 2 SP1, a behavior change was made to address a deadlock defect. After adding a new node to the cluster, you must run the `/opt/novell/ncs/bin/ncs-configd.py -init` script or the `rcnovell-ncs restart` command in order to make `cluster view` and the Clusters plug-in for iManager display the new node's name correctly.

- 1 Open a terminal console on the master node, then log in as the root user.
- 2 At the console prompt, enter
`cluster exec "/opt/novell/ncs/bin/ncs-configd.py -init"`

13.10 NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)

In some environments, a timing problem prevents the NDP user space application (`ndpapp`) from loading, and in turn, NSS cannot be loaded until the problem resolves itself. You can increase the UDEV event handling limits to 1024 to circumvent this problem. For instructions, see “[NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted \(Linux\)](#)” in the *OES 2 SP2: NSS File System Administration Guide*.

13.11 Could Not Delete This Resource Data_Server (Error 499)

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 console prompt as the `root` user, then enter the following commands:

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

echo -n "exec rm -f /var/opt/novell/ncs/data_server.*" > /proc/ncs/cluster
```

13.12 Problem Authenticating to Remote Servers during Cluster Configuration

During the Novell Open Enterprise Server (OES) 2 Linux cluster installation and configuration, if you choose *Remote System* on the Novell 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.

13.13 Problem Connecting 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 either reboot the Linux initiator server or run the `evms_activate` command at the Linux terminal console. This is required for each Linux initiator server that will access the iSCSI target.

For instructions on configuring an OES 2 Linux server as an iSCSI initiator and connecting to an iSCSI target, go to “[Mass Storage over IP Networks--iSCSI](http://www.novell.com/documentation/sles10/sles_admin/data/cha_inst_system_iscsi.html)” (http://www.novell.com/documentation/sles10/sles_admin/data/cha_inst_system_iscsi.html) in the *SUSE Linux Enterprise Server 10 Administration Guide*.

13.14 Problem Deleting a Cluster Resource or Clustered Pool

If you attempt to delete a cluster resource or clustered pool without first offlining the cluster resource, deletion errors occur, and the data associated with the clustered pool is not recoverable.

To avoid this problem, offline the cluster resource before attempting to delete either the cluster resource or the clustered pool.

13.15 Can't Find the Prevent Cascading Failover Option

Novell Cluster Services for NetWare provides a Preventing Cascading Failovers option. Its behavior was controlled using the `clstrlib /hmo=off` option in `ldncs.ncf`. Is there an equivalent option for Linux?

Cascading failover prevention was not implemented for Linux because user space applications (cluster resources) typically cannot cause a kernel failure and cascade failure of other kernels due to failing over user space code, as compared to ring0 NLMs on NetWare.

13.16 Is there a way to uninstall Novell Cluster Services from a server?

There is no utility to remove Novell Cluster Services and its related eDirectory objects. For information, see *How to Remove Clustering Services and the Clustering DS Objects (TID 3317095)* (<http://www.novell.com/support/search.do?usemicrosite=true&searchString=3317095>) on the Novell Support (<http://www.novell.com/support>) Web site.

13.17 Cluster Resource Is Stuck in "NDS Sync" or "eDirectory Sync" State

In OES 2, a cluster resource can become stuck in an `NDS Sync` or `eDirectory Sync` state if the character case for cluster information stored in the `/etc/opt/novell/ncs/clstrlib.conf` file does not match the case used for information stored in eDirectory. The `clstrlib.conf` file is automatically populated with information that you enter as you manage a cluster or cluster resource with tools such as iManager, YaST, and commands. If you use the wrong case when typing cluster information, that information is written to the file.

For example, you might enter the wrong case when performing the following cluster management tasks:

- ♦ In the YaST Software Installer, you type the cluster name and context when adding a node to an existing cluster instead of browsing to choose the information.

IMPORTANT: In OES 2 SP1 and later releases, the installer has been modified to make sure that the information that is added to the `/etc/opt/novell/ncs/clstrlib.conf` file matches the information in eDirectory.

- ◆ You modify the cluster configuration for an existing cluster as described in [Section 8.13, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 110.
- ◆ You modify the case for a container name or object name for an eDirectory object.

For information about how to make corrections in the `/etc/opt/novell/ncs/clstrlib.conf` file to resolve the `NDS Sync` or `eDirectory Sync` state, see [Cluster Resource Stuck in "NDS SYNC" or "edirectory Sync" State \(TID 7001394\)](#) (<http://www.novell.com/support/viewContent.do?externalId=7001394&sliceId=2>) on the [Novell Support](#) (<http://www.novell.com/support>) Web site.

A patch that helps prevent the occurrence of mismatched cluster information is available in the OES 2 SP1 patch channel and OES 2 SP2 patch channel in their respective March 2010 patch releases.

IMPORTANT: After the patch is applied, the `/etc/opt/novell/ncs/clstrlib.conf` file is case-insensitive.

13.18 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.

Novell 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.

13.19 Nodes Reboot Randomly after January or March 2010 Updates

After applying the January or March 2010 maintenance patches for OES 2 SP1 or OES 2 SP2, some Novell Cluster Services nodes reboot randomly.

This problem was caused by a code error in the NCS timer interruption handler. A fix is available. For information, see [Novell Cluster Services Nodes Rebooting Randomly after March 2010 Updates \(TID 7005916\)](#) (<http://www.novell.com/support/search.do?usemicrosite=true&searchString=7005916>) on the [Novell Support](#) (<http://www.novell.com/support>) Web site.

This section describes security issues and recommendations for Novell Cluster Services for Novell Open Enterprise Server 2 Linux. It is intended for security administrators or anyone who is responsible for the security of the system. It requires a basic understanding of Novell Cluster Services. It also requires the organizational authorization and the administrative rights to effect the configuration recommendations.

- ♦ [Section 14.1, “Cluster Administration Rights,” on page 207](#)
- ♦ [Section 14.2, “Ports,” on page 207](#)
- ♦ [Section 14.3, “E-Mail Alerts,” on page 207](#)
- ♦ [Section 14.4, “Log Files,” on page 208](#)

14.1 Cluster Administration Rights

For information about the rights needed to install and manage Novell Cluster Services, see the following sections:

- ♦ [“Novell eDirectory 8.8” on page 31](#)
- ♦ [“Novell Domain Services for Windows” on page 34](#)
- ♦ [“OpenWBEM” on page 34](#)
- ♦ [“SLP” on page 35](#)
- ♦ [“eDirectory Configuration” on page 36](#)
- ♦ [“Cluster Services Installation Administrator” on page 36](#)
- ♦ [“Cluster Services Management Administrator” on page 37](#)

14.2 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. For information, see [Section 7.6, “Modifying the Cluster IP Address and Port Properties,” on page 100](#).

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 with OpenWBEM.

14.3 E-Mail Alerts

You can enable or disable e-mail notification for the cluster and specify up to eight administrator e-mail addresses for cluster notification.

Novell Cluster Services uses Postfix to send e-mail alerts. If you have a cluster resource that uses SMTP, that resource might not work in the cluster unless you change the Postfix configuration.

For information, see [Section 7.4, “Configuring Cluster Event E-Mail Notification,”](#) on page 99.

14.4 Log Files

Look for events in the NCS Event Log in the Clusters plug-in to Novell iManager 2.7.2 or later.

Log files for cluster resources are written to the `/var/run/ncs` directory. Output log files are named with the resource name, load or unload to indicate which script, and the `.out` extension. For example: `/var/run/ncs/myresource.unload.out` reports the log for the `myresource` unload script.

Console Commands for Novell Cluster Services

A

Novell Cluster Services provides several console commands to help you perform certain cluster-related tasks.

- ◆ [Section A.1, “Cluster Management Commands,” on page 209](#)
- ◆ [Section A.2, “Business Continuity Clustering Commands,” on page 212](#)
- ◆ [Section A.3, “extend_schema Command,” on page 212](#)
- ◆ [Section A.4, “SBD Utility,” on page 213](#)

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 console 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
<code>ALERT {resource}{YES NO}</code>	The resource start, failover, or failback mode is set to manual and the resource is waiting to start on a node, or to fail over or fail back to another node. Specify the resource name in the command and use the YES or NO switch to specify whether you want the resource to fail over, fail back, or start.
<code>CONVERT {Preview, Commit}{Resource}</code>	Finalizes the cluster conversion from NetWare to Linux after all nodes in a mixed cluster have been converted to Linux. Specify a resource name with the Preview switch to view the resource load and unload script changes prior to finalizing the conversion. Use the Commit switch without specifying a resource to finalize the conversion for all cluster resources. The <code>CLUSTER CONVERT</code> command can only be executed on Linux cluster nodes.
<code>DOWN</code>	Removes all cluster nodes from the cluster. Has the same effect as executing the <code>CLUSTER LEAVE</code> command on every server in the cluster.
<code>EXEC "path_to_script"</code>	Executes the specified script on all nodes in the cluster.

Cluster Console Command	Description
INFO {All, Basic, Notification, Priority, Protocol, Summary}	<p>Displays information on cluster configuration.</p> <p>All displays a combination of Basic, Notification, Priority, and Protocol information.</p> <p>Basic displays IP address, port, and cluster quorum settings.</p> <p>Notification displays cluster e-mail notification settings.</p> <p>Priority displays the resource priority list.</p> <p>Protocol displays the cluster protocol settings.</p> <p>Summary displays the cluster protocol summary.</p>
JOIN	<p>Adds the node where the command is executed to the cluster and makes the node visible to other servers in the cluster. Novell Cluster Services software must already be installed on a node for it to join the cluster.</p>
LEAVE	<p>Removes the node where the command is executed from the cluster. The node will not be visible to other servers in the cluster.</p>
MAINTENANCE {ON OFF}	<p>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.</p> <p>Turning this switch on from one cluster server puts the entire cluster in maintenance mode.</p>
MIGRATE {resource}{node name}	<p>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 <i>Assigned Nodes</i> list.</p>
OFFLINE {resource}	<p>Unloads the specified resource from the node where it is currently running.</p>
ONLINE {resource}{node name}	<p>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.</p>
POOLS	<p>Lists the NSS pools on the shared disk system that are accessible by Novell Cluster Services.</p>
RESOURCES	<p>Lists all resources that currently exist in the cluster. The resources do not need to be online or running.</p>
RESTART {seconds}	<p>Restarts Novell Cluster Services software on all servers in the cluster.</p>

Cluster Console Command	Description
SCAN FOR NEW DEVICES	<p>IMPORTANT: This command is obsolete on Linux.</p> <p>You can use the <code>rescan-scsi-bus.sh</code> script to scan for the new devices on Linux without rebooting. For information, see “Scanning for New Devices without Rebooting” (http://www.novell.com/documentation/sles10/stor_admin/data/scandev.html) in the <i>SLES 10 SP3: Storage Administration Guide</i> (http://www.novell.com/documentation/sles10/stor_admin/data/bookinfo.html).</p>
SET {Parameter} {Value}	<p>Sets cluster parameters individually for the cluster. See Chapter 7, “Configuring Cluster Policies and Priorities,” on page 95 for more information on cluster parameters.</p> <p>Specify one of the following parameters and a value for that parameter:</p> <p>IPADDRESS sets the cluster IP address to the specified value. If you change the cluster IP address, you must restart cluster software on all cluster nodes.</p> <p>PORT sets, or lets you change, the cluster port number.</p> <p>QUORUMWAIT is the amount of time in seconds that the cluster waits before resources start to load.</p> <p>QUORUM is the number of nodes that must be running in the cluster before resources will start to load.</p> <p>HEARTBEAT is the amount of time in seconds between transmits for all nodes in the cluster except the master.</p> <p>TOLERANCE is the amount of time in seconds that the master node gives all other nodes in the cluster to signal that they are alive.</p> <p>MASTERWATCHDOG is the amount of time in seconds between transmits for the master node in the cluster.</p> <p>SLAVEWATCHDOG is the amount of time in seconds that the slave nodes give the master node in the cluster to signal that it is alive.</p> <p>MAXRETRANSMITS is the maximum number of times transmits will be attempted between the master node and slave nodes.</p> <p>ENABLEEMAIL enables and disables e-mail notification. You can set the value to OFF to disable e-mail notification, or either CRITICAL or VERBOSE to enable e-mail notification.</p> <p>EMAILADDRESSES lets you specify the e-mail addresses used for e-mail notification. The addresses should be separated by spaces. Using this parameter without specifying any addresses clears existing addresses that have been set previously.</p> <p>EMAILOPTIONS sets the e-mail notification options. Specify XML as the value to receive e-mail notification in XML format. Not specifying any value with this parameter turns notification in XML format off.</p>

Cluster Console Command	Description
STATS {Display, Clear}	Reports the node number, node name, and heartbeat information. You must switch to the log console screen to see cluster statistics.
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.
VIEW	Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster.

A.2 Business Continuity Clustering Commands

If Novell Business Continuity Clustering (BCC) 1.2 for OES 2 SP1 Linux 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” (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bcccommands.html) in the *BCC 1.2: Administration Guide for Linux* (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html).

A.3 extend_schema Command

```
/opt/novell/oes-install/util/extend_schema --port port_num admin_username admin_password  
server_ip_address schema_file
```

A tree administrator user with credentials to do so must use the `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.

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
```

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 username 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
server_ip_address	The IP address of the eDirectory server that contains the schema files.	10.10.10.1

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 Novell 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 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 NSSMU or the Storage plug-in to iManager. In OES 2 SP2, the NSS utility called `ncsinit` is available for initializing a device and setting it to a shared state.

Syntax

```
sbdutil [-c|-f|-i|-v] [-s] [-d device] [-d device] [-p partition] [-n
cluster_name]
```

Enter the command at a terminal console prompt as the `root` user or any other user in `admin` or `ncsgroup`.

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.

-f

Find the SBD partition.

-v

View 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. Use the `EVMSGUI` (or `EVMSN` or `EVMS`) tool to find the names of the devices you want to use, and only use 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 this cluster name instead of getting it from `cluster.xml`.

-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.

Files for Novell Cluster Services

B

Knowing the location and purpose of the files that make up Novell 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 Novell Cluster Services.

Table B-1 Location and Purpose of Novell Cluster Services Files

NCS Filename and Path	Purpose
/etc/init.d/novell-ncs	LSB Compliant Service
/etc/opt/novell/ncs/nodename	This node's name
/lib/evms/2.3.3/ncs-1.0.0.so	EVMS snap-in
/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/Snapins/ ClusterCliSnapin_Alert.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ ClusterCliSnapin_Down.pm	Cluster CLI command
/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_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_Resources.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ ClusterCliSnapin_Restart.pm	Cluster CLI command

NCS Filename and Path	Purpose
/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/adminfs	Cluster management (iManager and CLI)
/opt/novell/ncs/bin/ldncs	Loads NCS; used by the Cluster Start command.
/opt/novell/ncs/bin/ncs-configd.py	Cluster configuration daemon
/opt/novell/ncs/bin/ncs-emaild	Cluster e-mail 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 cluster resource templates.
/opt/novell/ncs/bin/sbdutil	SBD partition utility
/opt/novell/ncs/bin/uldncs (not yet implemented)	Unloads NCS; used by the Cluster Stop command.
/opt/novell/ncs/lib/ncs-1.0.0.so	EVMS snap-in
/opt/novell/ncs/lib/ncsfuncs	Shared library commands for load/unload scripts
/opt/novell/ncs/schema/ncpserver.preldif	NCS schema file
/opt/novell/ncs/schema/ncpserver.ldif	NCS schema file
/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 /etc/init.d/novell-ncs
/usr/share/man/man7/sbdutil.7.gz	SBDUTIL man page
/var/opt/novell/ncs/hangcheck-timer.conf	hang check option (comparable to CPU hog)
/lib/modules/kernel_dir/ncs/clstrlib.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.

NCS Filename and Path	Purpose
/lib/modules/ <i>kernel_dir</i> /ncs/cma.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/cmsg.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/crm.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/css.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/cvb.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/gipc.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/sbd.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/sbdlb.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/vipx.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
/lib/modules/ <i>kernel_dir</i> /ncs/vll.ko	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.

Comparing Novell Cluster Services for Linux and NetWare



Table C-1 compares the features and capabilities of Novell Cluster Services 1.8.5 for Linux on Novell Open Enterprise Server (OES) 2 Linux to Novell Cluster Services 1.8.5 for NetWare on NetWare 6.5 Support Pack 7 or later.

Table C-1 Comparison of Novell Cluster Services for OES 2 Linux and NetWare 6.5 SP7 or Later

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Operating system	OES 2 Linux (OES 2 built on SUSE Linux Enterprise Server 10 SP1) or later	NetWare 6.5 SP7 or later
Two-node cluster with OES 2 license	Yes	Yes
Up to 32 nodes in a single cluster	Yes	Yes
Guest servers on Xen VMs as cluster nodes	Yes	Yes
Business Continuity Clustering support	BCC 1.2 for OES 2 SP1 Linux and later	BCC 1.1 SP2 for NetWare 6.5 SP8 and later
Administrator users	The administrator user whose credentials you provide during the install is the cluster administrator. The tree admin is not automatically given rights. Rights must be granted manually. For information, see Section 3.7, "Configuring Additional Administrators," on page 55.	The administrator user whose credentials you provide during the install is the cluster administrator. For NetWare, rights are automatically extended to the tree administrator user.
Directory-based cluster configuration	Common schema for NetWare and Linux	Common schema for NetWare and Linux

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Schema extension during the Novell Cluster Services install	<p>In OES 2 Linux, the user who installs Novell Cluster Services must have schema extension rights.</p> <p>In OES 2 SP1 Linux and later, the schema extension is performed separately from the Novell Cluster Services install. A user with schema extension rights extends the schema. For information, see Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,” on page 42.</p> <p>Afterwards, any administrator with sufficient rights can install Novell Cluster Services. For information, see Section 3.4, “Assigning Install Rights for Container Administrators,” on page 44.</p>	Yes; the user who installs Novell Cluster Services must have schema extension rights.
Forward migration for Novell Cluster Services	OES 1 SP2 Linux to OES 2 Linux or later. Down cluster and rolling cluster upgrade are supported.	<p>NetWare 6.5 SP6 to NetWare 6.5 SP7 or later. Down cluster and rolling cluster upgrade are supported.</p> <p>NetWare 6.0 to NetWare 6.5 SP7 or later. Down cluster and rolling cluster upgrade are supported.</p> <p>NetWare 5.1 to NetWare 6.5 SP7 or later. Only the down cluster upgrade is supported.</p>
Cluster conversion from NetWare to Linux	<p>NetWare 6.5 SP7 or later to OES 2 Linux or later</p> <p>NetWare 6.5 SP6 to OES 2 Linux or later</p> <p>NetWare 6.0 with latest service packs and patches to OES 2 Linux. Uses the same process as NetWare 6.5 to OES 2 Linux.</p>	Not applicable

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Mixed-node Linux and NetWare clusters	<p>Supported only for rolling cluster conversions from NetWare to Linux.</p> <p>NSS pools created on Linux cannot fail over to NetWare nodes.</p> <p>In a mixed-cluster environment, you cannot add storage to the cluster or modify the existing storage pools.</p> <p>After adding a Linux node to a NetWare cluster, it is no longer possible to add new NetWare nodes. Only Linux nodes can be added.</p>	<p>Supported only for rolling cluster conversions from NetWare to Linux.</p> <p>NSS pools created on NetWare can fail over to Linux nodes.</p>
SBD (split-brain detector)	Yes; during the install on the first server in the cluster, or by using the sbdutil after the install and before adding a second node to the cluster.	Yes; during the install on the first server in the cluster.
Mirrored SBD	Yes; by using evmsgui after the install and before adding a second node to the cluster.	Yes; during the install on the first server in the cluster.
Shared disks	<p>Fibre Channel SAN LUNs</p> <p>iSCSI SAN LUNs</p> <p>SCSI disks (shared external drive arrays)</p>	<p>Fibre Channel SAN LUNs</p> <p>iSCSI SAN LUNs</p> <p>SCSI disks (shared external drive arrays)</p>
Cluster-aware shared devices	Yes; requires using the EVMS Cluster Segment Manager to provide cluster-awareness similar to NetWare Media Manager.	Yes; provided by NetWare Media Manager.
Requires Novell Storage Services (NSS)	Required only if you create NSS pools as cluster resources.	<p>Yes; NSS is the default file system on NetWare.</p> <p>Novell Cluster Services is not supported on NetWare traditional volumes.</p>
Disk format when using shared NSS pools	NetWare Segment Manager	NetWare Segment Manager
Requires NCP (NetWare Core Protocol)	Requires NCP Server to be installed if you create cluster resources that use NCP, such as NSS pools, NCP volumes, and Dynamic Storage Technology shadow volume pairs.	Yes; NCP is the default file access protocol on NetWare.

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
NSS pools as cluster resources	<p>Yes</p> <p>For information, see Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139.</p> <p>Shareable for Clustering</p> <p>Multiple-Server Activation Prevention (MSAP)</p> <p>Cluster volume broker; Linux kernel module handles NSS pool events.</p>	<p>Yes</p> <p>For information, see “Setting Up Cluster Resources for Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide</i>.</p> <p>Shareable for Clustering</p> <p>Multiple-Server Activation Prevention (MSAP)</p> <p>Cluster volume broker</p>
Linux POSIX file systems as cluster resources	<p>Yes</p> <p>For information, see Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163.</p>	Not applicable
NCP volumes on Linux POSIX file systems as cluster resources	<p>Yes</p> <p>For information, see “Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 2 SP2: NCP Server for Linux Administration Guide</i>.</p>	Not applicable
Dynamic Storage Technology shadow volume pairs as cluster resources	<p>Yes; by combining the load and unload scripts for shared NSS pools and managing the pair as a single cluster resource.</p> <p>For information, see “Configuring DST Shadow Volumes with Novell Cluster Services for Linux” in the <i>OES 2 SP2: Dynamic Storage Technology Administration Guide</i>.</p>	Not supported by Dynamic Storage Technology.
Xen virtual machines as cluster resources	<p>Yes</p> <p>For information, see Section 12.2, “Virtual Machines as Cluster Resources,” on page 186.</p>	Not applicable
iManager 2.7	Yes	Yes
Clusters plug-in and Storage Management plug-in for iManager	Yes	Yes
Cluster-enabling NSS pools by using the NSS plug-in for iManager	Yes	Yes

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Cluster-enabling NSS pools by using the NSS Management Utility (NSSMU)	Yes	Yes
Command line interface	Yes; using the terminal console	Yes; using the terminal console
XML-based API	Yes; same as NetWare except that it uses the <code>/_adminfs</code> path on Linux.	Yes; same as for Linux except that it uses the <code>_admin:</code> volume on NetWare.
Load, unload, and monitor scripts	Yes Script commands differ. Scripts are automatically translated from NetWare commands to Linux commands during the cluster conversion from NetWare to Linux. For a comparison of script commands, see Section 6.6, "Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters," on page 86.	Yes
NCP support for accessing files on shared NSS pools	Yes	Yes
NCP support for accessing files on shared NCP volumes on Linux POSIX file systems	Yes	Not applicable
Novell AFP support for accessing files on shared NSS pools	Yes; in OES 2 SP1 Linux and later	Yes
Novell CIFS support for accessing files on shared NSS pools	Yes; in OES 2 SP1 Linux and later. Cross-protocol locking is not supported in OES 2 SP1.	Yes; with cross-protocol locking
Linux Samba/CIFS support for accessing files on shared NSS pools on Linux	Yes; with support for cross-protocol locking that allows both NCP users and Linux Samba/CIFS users to concurrently access data. Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Linux Samba/CIFS support for accessing files on shared NCP volumes on Linux POSIX file systems	Yes; with support for cross-protocol locking that allows both NCP users and Linux Samba/CIFS users to concurrently access data. Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Linux Samba/CIFS support for accessing files on shared Linux POSIX file systems	Yes Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Leverage Heartbeat 2 resource agents	Yes	Not applicable
LAN fault tolerance	Channel bonding For information, see <code>/usr/src/linux/Documentation/bonding.txt</code>	NIC teaming For information, see “ NIC Teaming ” in the <i>NW 6.5 SP8: TCP/IP Administration Guide</i>
Multipath I/O	Device Mapper - Multipath I/O, or third-party MPIO solutions For information, see “ Managing Multipath I/O for Devices ” in the <i>SLES 10 SP3: Storage Administration Guide</i> .	Media Manager Multipath I/O, or third-party MPIO solutions For information, see “ Managing Multipath I/O to Devices (NetWare) ” in the <i>OES 2 SP2: NSS File System Administration Guide</i> .

Comparing Clustering Support for OES 2 Services on Linux and NetWare

[Table D-1](#) compares clustering support for Novell Open Enterprise Server (OES) 2 services when using Novell Cluster Services 1.8.5 on OES 2 Linux and NetWare 6.5 SP7 or later.

NSS pool cluster resources can be cluster migrated from NetWare to Linux as part of a cluster conversion. If the resource contains data only, no additional steps are required. However, clustered services can require special handling. For information, see [Section 6.3, “Guidelines for Converting Service Cluster Resources from NetWare to Linux,”](#) on page 68.

Table D-1 Comparison of Clustering Support for OES 2 Services on Linux and NetWare

Service	OES 2 NetWare	OES 2 Linux	Comments
AFP (Apple Filing Protocol)	Yes See “Setting Up for Macintosh” in the <i>NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide</i> .	Yes; for OES 2 SP1 Linux and later See “Configuring AFP with Novell Cluster Services for an NSS File System” in the <i>OES 2 SP2: Novell AFP For Linux Administration Guide</i> .	For information about cluster migrating an AFP service cluster resource from NetWare to Linux, see Section 6.3.3, “Apple Filing Protocol (AFP),” on page 71.
Apache Web Server	Yes See “Apache with Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide</i> .	Yes; use the standard Apache Web Server for Linux.	For information about cluster migrating an Apache service cluster resource from NetWare to Linux, see Section 6.3.2, “Apache Web Server,” on page 71.

Service	OES 2 NetWare	OES 2 Linux	Comments
Archive and Version Services (Novell)	<p>Yes</p> <p>See “Installing and Configuring New Archive Servers” in the <i>NW 6.5 SP8: Novell Archive and Version Services 2.1 Administration Guide</i>.</p>	<p>Yes</p> <p>See “Configuring Archive and Version Service for Novell Cluster Services” in the <i>OES 2 SP2: Novell Archive and Version Services 2.1 for Linux Administration Guide</i>.</p>	<p>On NetWare, it uses the MySQL database.</p> <p>On Linux, it uses PostgreSQL instead of MySQL. The archive database must reside on a Linux POSIX file system.</p> <p>For information about cluster migrating an Apache service cluster resource from NetWare to Linux, see Section 6.3.4, “Archive and Version Services,” on page 71.</p>
CIFS (Windows File Services)	<p>Yes; Novell CIFS</p> <p>See “Setting Up for Windows” in the <i>NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide</i>.</p>	<p>Yes; Novell CIFS</p> <p>See “Configuring CIFS with Novell Cluster Services for an NSS File System” in the <i>OES 2 SP2: Novell CIFS for Linux Administration Guide</i>.</p>	<p>For information about converting a cluster resource from NetWare to Linux, see Section 6.3.5, “CIFS,” on page 72.</p>
DFS (Novell Distributed File Services) Volume location database (VLDB)	<p>Yes</p> <p>See “Clustering Novell Distributed File Services” in the <i>NW 6.5 SP8: Novell Distributed File Services Administration Guide</i>.</p>	<p>Yes</p> <p>See “Clustering Novell Distributed File Services” in the <i>OES 2 SP2: Novell Distributed File Services Administration Guide</i>.</p>	<p>For information about converting a cluster resource from NetWare to Linux, see Section 6.3.6, “DFS VLDB,” on page 72.</p>
DHCP	<p>Yes</p> <p>See “Clustering in NetWare 6.5” in the <i>NW 6.5 SP8: Novell DNS/DHCP Services Administration Guide</i>.</p>	<p>Yes</p> <p>See “Configuring DHCP with Novell Cluster Services for the Linux File System” in the <i>OES 2 SP2: Novell DNS/DHCP Administration Guide for Linux</i>.</p>	<p>DHCP for Linux supports using a shared Linux POSIX file system or a shared NSS (Novell Storage Services) pool (supported as of OES 2 SP1) for the cluster resource.</p> <p>For information about converting a cluster resource from NetWare to Linux, see Section 6.3.7, “DHCP Server,” on page 73.</p>


Service	OES 2 NetWare	OES 2 Linux	Comments
DNS	Yes See “Creating a Cluster-Enabled DNS Server” in the <i>NW 6.5 SP8: Novell DNS/DHCP Services Administration Guide</i> .	Yes See “Configuring DNS with Novell Cluster Services” in the <i>OES 2 SP2: Novell DNS/DHCP Administration Guide for Linux</i> .	The association of the DNS Server object with the NCP Server must be changed manually after the cluster conversion. It is not done as part of the simple cluster migration. For information about converting a cluster resource from NetWare to Linux, see Section 6.3.8, “DNS Server,” on page 75.
Dynamic Storage Technology service	Not available	Can be used in a cluster, but the service is not clustered. See also Storage, DST shadow volume pairs .	DST runs on each OES 2 Linux node and you set the global server-level parameters to be the same on each one.
eDirectory 8.8	No	No	eDirectory has its own redundancy built in (multiple replicas) and would not benefit from being clustered.
eDirectory Certificate Server	Yes See “Server Certificate Objects and Clustering” in the <i>Novell Certificate Server 3.3.1 Administration Guide</i> (http://www.novell.com/documentation/crt33/).	Yes See “Server Certificate Objects and Clustering” in the <i>Novell Certificate Server 3.3.2 Administration Guide</i> .	For information about converting a cluster resource from NetWare to Linux, see Section 6.3.9, “eDirectory Server Certificates,” on page 76.
exteNd Application Server and MySQL	Yes; NetWare 6.5 SP2 or earlier. See “Configuring Novell exteNd Application Server and MySQL with Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide</i> .	Not available on Linux.	This install option was discontinued beginning in NetWare 6.5 SP3. See MySQL .
FTP Server	Yes See “Cluster-Enabling NetWare FTP Server” in the <i>NW 6.5 SP8: Novell FTP Administration Guide</i> .	No; standard Linux solution	

Service	OES 2 NetWare	OES 2 Linux	Comments
Identity Manager 3.6 Bundle Edition	Can be used in a cluster, but is not clustered. Requires Identity Manager 3.5.	Can be used in a cluster, but is not clustered.	In OES 2 SP1 Linux, it is 32-bit only and requires a 32-bit operating system.
iPrint	Yes See the NW 6.5 SP8: iPrint Administration Guide .	Yes See the OES 2 SP2: iPrint for Linux Administration Guide .	For information about converting a cluster resource from NetWare to Linux, see Section 6.3.10, "iPrint," on page 78 .
MySQL	Yes See "Configuring MySQL on Novell Clustering Services" in the NW 6.5 SP8: Novell MySQL Administration Guide .	Yes; use the standard MySQL service for Linux. MySQL 5.0.x on OES2 Linux is offered under the GPL.	For Linux, use a procedure similar to the one on NetWare. Use the Linux commands for MySQL in the load and unload scripts. Use a Linux path on a shared Linux POSIX file system for the MySQL database.
NCP Server	Can be used in a cluster, but is not clustered.	Can be used in a cluster, but is not clustered. See also Storage, NCP volumes on Linux POSIX file systems .	NCP Server runs on each server node in the cluster. It should have the same configuration on each node of the cluster.
NetStorage	Yes See "Configuring NetStorage with Novell Cluster Services" in the NW 6.5 SP8: NetStorage Administration Guide .	Yes See "Configuring NetStorage with Novell Cluster Services" in the OES 2 SP2: NetStorage for Linux Administration Guide .	No known issues.
NFS	Yes See "Cluster-Enabling Native File Access for UNIX" in the NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide .	No; native to Linux	
Novell iFolder 2.1.x	Yes See "Configuring iFolder on Novell Cluster Services" in the OES 2 SP1: Novell iFolder 2.1 Installation and Administration Guide .	Not applicable	iFolder 2.1x is replaced by iFolder 3.6 and later in OES 2 Linux.

Service	OES 2 NetWare	OES 2 Linux	Comments
Novell iFolder 3.6	Not applicable	Yes	See “Clustering iFolder 3.8 Servers with Novell Cluster Services for Linux” in the <i>Novell iFolder 3.8 Administration Guide</i> .
Printing	Yes	Yes	See iPrint .
QuickFinder	Yes	Yes	See Search .
Search (QuickFinder)	Yes; requires QuickFinder 4.2.0. Install QuickFinder on each server in your cluster, and use virtual search servers. See “Configuring QuickFinder Server for Novell Cluster Services” in the <i>NW 6.5 SP8 Novell QuickFinder Server 5.0 Administration Guide</i> .	Yes; requires QuickFinder 5.0.1. Install QuickFinder on each server in your cluster, and use virtual search servers. See “Configuring QuickFinder Server for Novell Cluster Services” in the <i>OES 2: Novell QuickFinder Server 5.0 Administration Guide</i> .	For information about converting a cluster resource from NetWare to Linux, see Section 6.3.11, “QuickFinder Server,” on page 79.
Storage, DST shadow volume pairs	Not applicable	Yes See “Configuring DST Shadow Volumes with Novell Cluster Services for Linux” in the <i>OES 2 SP2: Dynamic Storage Technology Administration Guide</i> .	DST shadow volumes are on shared NSS pools that are created separately, then managed in the same load/unload scripts.
Storage, NCP volumes on Linux POSIX file systems	Not applicable	Yes For information, see “Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 2 SP2: NCP Server for Linux Administration Guide</i> .	The NCP Server service is not clustered; its volumes can be clustered.
Storage, NetWare Traditional volumes	No	Not applicable	

Service	OES 2 NetWare	OES 2 Linux	Comments
Storage, NSS pools and volumes	<p>Yes</p> <p>See “Setting Up Cluster Resources for Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide</i>.</p>	<p>Yes</p> <p>See Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139.</p>	<p>NSS itself is not clustered. The NSS service must be installed and running on each server node where you want to load the shared pools. Server-level NSS parameter changes must be made manually on each node.</p> <p>Shared NSS pools and volumes that were created on NetWare can be migrated or failed over from a NetWare server to a Linux server during a cluster conversion. For information, see Chapter 6, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 63.</p> <p>Most NSS features are available on both platforms. For information see “Cross-Platform Issues for NSS” in the <i>OES 2 SP2: NSS File System Administration Guide</i>.</p>
Tomcat	<p>Yes</p> <p>See “Configuring Tomcat and Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide</i>.</p>	<p>Yes; native to Linux</p> <p>Use a similar procedure to the one outlined for Tomcat on NetWare, but use the Linux locations and files.</p>	<p>You cannot convert the NetWare Tomcat configuration for a Linux server.</p>
Xen guest servers as nodes in a cluster	<p>Virtualized NetWare nodes can be used in NetWare clusters. Nodes can be any combination of virtual and physical servers.</p>	<p>Virtualized OES 2 Linux nodes can be used in OES 2 Linux clusters. Nodes can be any combination of virtual and physical servers.</p>	<p>See Chapter 12, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 185.</p>
Xen virtual machines on the host server	<p>Not applicable</p>	<p>Yes; use the Xen and XenLive templates.</p>	<p>See Section 12.2, “Virtual Machines as Cluster Resources,” on page 186.</p>

Documentation Updates



This section contains information about documentation content changes made to the *Novell Cluster Services 1.8.5 for Linux Administration Guide* since the initial 1.8.5 release for Novell Open Enterprise Server 2. If you are an existing user, review the change entries to readily identify modified content. If you are a new user, simply read the guide in its current state.

This document was updated on the following dates:

- ♦ [Section E.1, “August 2, 2010 \(Maintenance Patch\),” on page 231](#)
- ♦ [Section E.2, “June 14, 2010 \(Maintenance Patch\),” on page 233](#)
- ♦ [Section E.3, “March 15, 2010,” on page 236](#)
- ♦ [Section E.4, “February 19, 2010,” on page 238](#)
- ♦ [Section E.5, “February 10, 2010,” on page 238](#)
- ♦ [Section E.6, “January 29, 2010 \(Maintenance Patch\),” on page 239](#)
- ♦ [Section E.7, “January 20, 2010,” on page 240](#)
- ♦ [Section E.8, “January 4, 2010,” on page 241](#)
- ♦ [Section E.9, “December 15, 2009,” on page 242](#)
- ♦ [Section E.10, “December 10, 2009,” on page 242](#)
- ♦ [Section E.11, “November 2009 \(OES 2 SP2\),” on page 243](#)
- ♦ [Section E.12, “July 30, 2009,” on page 245](#)
- ♦ [Section E.13, “June 22, 2009,” on page 248](#)
- ♦ [Section E.14, “June 5, 2009,” on page 248](#)
- ♦ [Section E.15, “May 6, 2009,” on page 250](#)
- ♦ [Section E.16, “March 3, 2009,” on page 252](#)
- ♦ [Section E.17, “February 13, 2009,” on page 254](#)
- ♦ [Section E.18, “February 3, 2009,” on page 254](#)
- ♦ [Section E.19, “January 13, 2009,” on page 255](#)
- ♦ [Section E.20, “December 2008 \(OES 2 SP1\),” on page 256](#)
- ♦ [Section E.21, “June 4, 2008,” on page 260](#)
- ♦ [Section E.22, “May 2, 2008,” on page 263](#)

E.1 August 2, 2010 (Maintenance Patch)

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.1.1, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 232](#)
- ♦ [Section E.1.2, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 232](#)
- ♦ [Section E.1.3, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 232](#)

- ◆ [Section E.1.4, “Managing Clusters,” on page 233](#)
- ◆ [Section E.1.5, “What’s New,” on page 233](#)

E.1.1 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
“Configuring a Load Script for a Linux POSIX Volume Cluster Resource” on page 170	Step 2 on page 171 describes how to use mount options in the load script.
“Creating a Virtual Server Name for the Cluster Resource” on page 175	This section was modified to clarify the the <code>NCS:NCP Server</code> object is created under the Cluster object for the cluster where you created the shared Linux POSIX volume.

E.1.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Novell Storage Services” on page 140	The Server, Pool, Volume, Cluster Resource, and Cluster objects are recommended to be in the same context (such as <code>ou=ncs,o=novell</code>). If they are in different contexts, you might need to cluster migrate the pool cluster resource back to the node where the pool was created.
“Considerations for Working with Shared NSS Pools and Volumes in the Cluster” on page 142	

E.1.3 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
“Managing File Systems in Mixed-Node Clusters” on page 66	If you need to configure (or reconfigure) existing shared NSS pools and volumes in a mixed-node cluster, you must temporarily bring down all Linux cluster nodes prior to making changes, then make the configuration changes on a NetWare node. Make sure the resources are working properly on NetWare before having the Linux cluster nodes rejoin the cluster.

E.1.4 Managing Clusters

Location	Change
"OES 2 SP2 with Patches and Later" on page 109	To allow an automatic reboot, set the <code>kernel.panic</code> token to value to 1 in the <code>/etc/sysctl.conf</code> file.

E.1.5 What's New

Location	Change
"What's New (July 2010 Patches)" on page 23	This section is new.

E.2 June 14, 2010 (Maintenance Patch)

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.2.1, "Configuring and Managing Cluster Resources," on page 233](#)
- ◆ [Section E.2.2, "Configuring Cluster Policies and Priorities," on page 234](#)
- ◆ [Section E.2.3, "Console Commands for Novell Cluster Services," on page 234](#)
- ◆ [Section E.2.4, "Converting NetWare 6.5 Clusters to OES 2 Linux," on page 234](#)
- ◆ [Section E.2.5, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 234](#)
- ◆ [Section E.2.6, "Managing Clusters," on page 235](#)
- ◆ [Section E.2.7, "Troubleshooting Novell Cluster Services," on page 235](#)
- ◆ [Section E.2.8, "Upgrading OES 2 Linux Clusters," on page 236](#)
- ◆ [Section E.2.9, "What's New," on page 236](#)

E.2.1 Configuring and Managing Cluster Resources

Location	Change
"Naming Conventions for Cluster Resources" on page 123	This section is new.

E.2.2 Configuring Cluster Policies and Priorities

Location	Change
“Configuring Cluster Event E-Mail Notification” on page 99	The subject line provides information about the cluster name, resource name, action taken, and node name.

E.2.3 Console Commands for Novell Cluster Services

Location	Change
“SCAN FOR NEW DEVICES” on page 211	You can use the <code>rescan-scsi-bus.sh</code> script to scan for the new devices on Linux without rebooting. For information, see “ Scanning for New Devices without Rebooting ” (http://www.novell.com/documentation/sles10/stor_admin/data/scandev.html) in the <i>SLES 10 SP3: Storage Administration Guide</i> (http://www.novell.com/documentation/sles10/stor_admin/data/bookinfo.html).

E.2.4 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
“Apple Filing Protocol (AFP)” on page 71	Beginning in OES 2 SP2 Linux, the cross-protocol file locking parameter for NCP Server is enabled by default.
“CIFS” on page 72	Beginning in OES 2 SP2 Linux, the cross-protocol file locking parameter for NCP Server is enabled by default.

E.2.5 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
“SBD Partitions” on page 39	When you mark a device as <i>Shareable for Clustering</i> , the share information is added to the disk and consumes about 4 MB. This space is used in addition to the space needed for the SBD. For example, if you create a 1 GB (1024 MB) device for the SBD, 4 MB are used for the share information, leaving 1020 MB as available free space for the SBD partition.

E.2.6 Managing Clusters

Location	Change
“Generating a Cluster Configuration Report” on page 105	The Monitoring script is included in the report for each resource.
“Generating a Cluster Configuration Report” on page 105	Step 6 provides instructions for how to save the report to a file.
“Preventing a Cluster Node Reboot after a Node Shutdown” on page 109	Novell Cluster Services has been modified to handle panic situations for a node according to the Linux kernel panic setting on the node.
“Creating a Non-Mirrored Cluster SBD Partition” on page 116	In Step 1 on page 116 , added the -s option to the command and example. Step 2 on page 116 is required only if Cluster Services was initially installed without an an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the <i>NCS: Shared Disk Flag</i> attribute is enabled.
“Using SBDUTIL to Create a Mirrored Cluster SBD Partition” on page 117	In Step 1 on page 117 , added the -s option to the command and example. Step 2 on page 117 is required only if Cluster Services was initially installed without an an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the <i>NCS: Shared Disk Flag</i> attribute is enabled.

E.2.7 Troubleshooting Novell Cluster Services

Location	Change
“Diagnosing a Poison Pill Condition” on page 200	This section is new.
“Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems” on page 205	This section is new.
“Nodes Reboot Randomly after January or March 2010 Updates” on page 205	This section is new.

E.2.8 Upgrading OES 2 Linux Clusters

Location	Change
"Updating the iFolder Resource Template" on page 58	You must run a script on one of the nodes (master node is preferred) in the cluster after upgrading to SP2. The template changes take effect without rebooting the node or restarting the cluster.

E.2.9 What's New

Location	Change
"What's New (June 2010 Patches)" on page 23	This section is new.

E.3 March 15, 2010

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.3.1, "Configuring and Managing Cluster Resources," on page 236](#)
- ◆ [Section E.3.2, "Configuring Cluster Resources for Shared Linux POSIX Volumes," on page 237](#)
- ◆ [Section E.3.3, "Configuring Novell Cluster Services in a Xen Virtualization Environment," on page 237](#)
- ◆ [Section E.3.4, "Troubleshooting Novell Cluster Services," on page 238](#)

E.3.1 Configuring and Managing Cluster Resources

Location	Change
Section 9.1, "Planning Cluster Maintenance," on page 123	This section is new.

Location	Change
Section 9.3.2, “Creating a Resource Template,” on page 125	Cluster Services marks the load, unload, or 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.
Section 9.5, “Configuring a Load Script for a Cluster Resource,” on page 127	
Section 9.6, “Configuring an Unload Script for a Cluster Resource,” on page 128	
Section 9.7.1, “Configuring Resource Monitoring,” on page 130	

E.3.2 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 11.3.3, “Configuring a Load Script for a Linux POSIX Volume Cluster Resource,” on page 170	Cluster Services marks the load, unload, or 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.
Section 11.3.4, “Configuring an Unload Script for a Linux POSIX Volume Cluster Resource,” on page 171	
Section 11.3.5, “Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource,” on page 172	

E.3.3 Configuring Novell Cluster Services in a Xen Virtualization Environment

Location	Change
“Viewing or Modifying the Monitor Script” on page 192	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.

E.3.4 Troubleshooting Novell Cluster Services

Location	Change
"File Location Information" on page 199	This section is new.

E.4 February 19, 2010

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.4.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 238](#)
- ♦ [Section E.4.2, "Managing Clusters," on page 238](#)

E.4.1 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
"Novell iManager 2.7.3" on page 31	The OES 2 SP2 Linux release contains a Clusters plug-in that is required when using Novell Business Continuity Clustering 1.2.1 for OES 2 SP2 Linux. For information about BCC 1.2.1, see BCC 1.2.1: Administration Guide for OES 2 SP2 Linux (http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html).

E.4.2 Managing Clusters

Location	Change
Section 8.17, "Customizing Cluster Services Management," on page 120	The references to <code>/_admin</code> should be <code>/admin</code> on Linux.

E.5 February 10, 2010

Updates were made to the following section. The changes are explained below.

- ♦ [Section E.5.1, "Configuring Cluster Resources for Linux POSIX Volumes," on page 239](#)

E.5.1 Configuring Cluster Resources for Linux POSIX Volumes

Location	Change
Section 11.4, “Creating a Virtual Server Name for the Cluster Resource,” on page 175	This section is new.

E.6 January 29, 2010 (Maintenance Patch)

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.6.1, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 239](#)
- ◆ [Section E.6.2, “Installing and Configuring Novell Cluster Services on OES 2 Linux,” on page 239](#)
- ◆ [Section E.6.3, “Managing Clusters,” on page 240](#)
- ◆ [Section E.6.4, “Troubleshooting Novell Cluster Services,” on page 240](#)
- ◆ [Section E.6.5, “What’s New,” on page 240](#)

E.6.1 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
“SBD Devices Must Be Marked as Shareable for Clustering” on page 64	This section is new.
“Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion)” on page 82	Step 1 on page 82 is new. Before you add the first Linux node to the NetWare cluster, mark the NetWare cluster SBD devices as Shareable for Clustering.
“Adding New OES 2 Linux Nodes to Your NetWare Cluster” on page 85	Step 1 on page 85 is new. Before you add the first Linux node to the NetWare cluster, mark the NetWare cluster SBD devices as Shareable for Clustering

E.6.2 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
“Novell eDirectory 8.8” on page 31	You can install an eDirectory master replica or replica in the cluster, but it is not required to do so for Novell Cluster Services. (eDirectory must be installed somewhere in the same tree.)

Location	Change
“SBD Partitions” on page 39	If you plan to mirror the SBD, you must also initialize the partition you want to use as the mirrored SBD, and mark its device as shareable for clustering.
“Configuring a New Cluster” on page 51	For each SBD partition, you must have at least 20 MB of free space on a device that has been initialized and marked as shareable for clustering.
“Adding a Node to an Existing Cluster” on page 54	Step 1 on page 54 is new. 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.

E.6.3 Managing Clusters

Location	Change
“Generating a Cluster Configuration Report” on page 105	.This section is new.

E.6.4 Troubleshooting Novell Cluster Services

Location	Change
Section 13.4, “Diagnosing Cluster Problems,” on page 200	Cluster configuration. An HTML report created by the <i>Clusters > Cluster Manager > Run Report</i> option in the Clusters plug-in to iManager.
“Cluster Resource Is Stuck in “NDS Sync” or “eDirectory Sync” State” on page 204	.This section is new.

E.6.5 What’s New

Location	Change
“What’s New (January 2010 Patches)” on page 23	.This section is new.

E.7 January 20, 2010

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.7.1, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 241](#)
- ◆ [Section E.7.2, “What’s New,” on page 241](#)

E.7.1 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
“Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion)” on page 82	.Step 11 was removed. A cluster restart is no longer required. For information, see Section 2.4.8, “Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade,” on page 25.
“Finalizing the Cluster Conversion” on page 92	.Step 3 was revised. A cluster restart is no longer required. For information, see Section 2.4.8, “Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade,” on page 25.

E.7.2 What’s New

Location	Change
“Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade” on page 25	.This section is new.

E.8 January 4, 2010

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.8.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,”](#) on page 241
- ♦ [Section E.8.2, “Converting NetWare 6.5 Clusters to OES 2 Linux,”](#) on page 242

E.8.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Using iManager to Create Shared NSS Pools” on page 144	In Step 9 on page 144 , if you choose to cluster-enable the pool later, you must create a volume on the shared pool before you attempt to cluster-enable the pool.
“Using NSSMU to Create Shared NSS Pools” on page 146	In Step 7 on page 146 , if you choose to cluster-enable the pool later, you must create a volume on the shared pool before you attempt to cluster-enable the pool.
“Cluster-Enabling an Existing NSS Pool and Its Volumes” on page 150	The pool must contain at least one volume before you attempt to cluster-enable the existing pool.
“Using NSSMU to Create NSS Volumes” on page 149	Information was modified about how to set volume attributes in NSSMU.

E.8.2 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
"Using Internal Certificates in a Cluster" on page 77	. This section was modified for clarity.

E.9 December 15, 2009

Updates were made to the following section. The changes are explained below.

- ♦ [Section E.9.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 242](#)

E.9.1 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
Section 3.3.2, "Extending the Schema," on page 43	Added Step 3 to create the <code>/var/opt/novell/install</code> directory.

E.10 December 10, 2009

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.10.1, "Converting NetWare 6.5 Clusters to OES 2 Linux," on page 242](#)
- ♦ [Section E.10.2, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 243](#)
- ♦ [Section E.10.3, "Managing Clusters," on page 243](#)
- ♦ [Section E.10.4, "Troubleshooting Novell Cluster Services," on page 243](#)

E.10.1 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.4, "Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion)," on page 82	Added the following instruction in Step 3 on page 82 : You can do this by running NWConfig, then selecting <i>Product Options > Directory Options <install NDS> > Remove Directory Services from this server.</i>

E.10.2 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
Section 3.3.2, “Extending the Schema,” on page 43	The procedure has been modified to create a configuration file and use it with the <code>/opt/novell/ncs/install/ncs_install.py</code> script to extend the schema.

E.10.3 Managing Clusters

Location	Change
Section 8.16, “Creating or Deleting Cluster SBD Partitions,” on page 115	The procedures in this section were reorganized and modified for clarity.
Section 8.16.7, “Removing a Segment from a Mirrored Cluster SBD Partition,” on page 119	This section is new.

E.10.4 Troubleshooting Novell Cluster Services

Location	Change
Section 13.16, “Is there a way to uninstall Novell Cluster Services from a server?,” on page 204	This section is new.

E.11 November 2009 (OES 2 SP2)

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.11.1, “Configuring Novell Cluster Services in a Xen Host Environment,” on page 244](#)
- ◆ [Section E.11.2, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 244](#)
- ◆ [Section E.11.3, “Console Commands for Novell Cluster Services,” on page 244](#)
- ◆ [Section E.11.4, “Installing Novell Cluster Services on OES 2 Linux,” on page 244](#)
- ◆ [Section E.11.5, “Managing Clusters,” on page 245](#)
- ◆ [Section E.11.6, “Troubleshooting Novell Cluster Services,” on page 245](#)
- ◆ [Section E.11.7, “Upgrading OES 2 Linux Clusters,” on page 245](#)
- ◆ [Section E.11.8, “What’s New,” on page 245](#)

E.11.1 Configuring Novell Cluster Services in a Xen Host Environment

Location	Change
Section 12.1, "Prerequisites for Xen Host Server Environments," on page 185	This section is new.

E.11.2 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.7, "Customizing the Translation Syntax for Converting Load and Unload Scripts," on page 91	This section is new.
Section 6.8, "Finalizing the Cluster Conversion," on page 92	Modified Step 3 to state that you must run <code>/opt/novell/ncs/bin/ncs-configd.py -init</code> on each of the nodes in the cluster, or restart Cluster Services on each of the nodes in the cluster.

E.11.3 Console Commands for Novell Cluster Services

Location	Change
Section A.4, "SBD Utility," on page 213	This section is new.

E.11.4 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Settings for the modprobe.conf.local File" on page 40	For OES 2 SP2 and later, or if you have installed the latest kernel and qla-driver, use the following setting in the <code>/etc/modprobe.conf.local</code> file: <code>options qla2xxx qlport_down_retry=2</code>
"Settings for the multipath.conf File" on page 41	For <code>no_path_retry</code> , a value of <code>fail</code> has the same meaning as a value of 0, and it is more easily understood. <code>failback "immediate"</code> <code>no_path_retry fail</code>

E.11.5 Managing Clusters

Location	Change
Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116	Use the EVMSGUI (or EVMSN or EVMS) tool to check the names of the devices if needed, and only use the base (leaf) names with the -d option.
“Using SBDUTIL to Create a Mirrored Cluster SBD Partition” on page 117	This section is new.

E.11.6 Troubleshooting Novell Cluster Services

Location	Change
Section 13.8, “smdr.novell Is Not Registered with SLP for a New Cluster Resource,” on page 202	This section is new.

E.11.7 Upgrading OES 2 Linux Clusters

Location	Change
Section 4.3, “Upgrade Issues for OES 2 SP2,” on page 58	This section is new.

E.11.8 What’s New

Location	Change
Section 2.4, “What’s New (OES 2 SP2),” on page 24	This section is new.

E.12 July 30, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.12.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 246](#)
- ◆ [Section E.12.2, “Console Commands for Novell Cluster Services,” on page 246](#)
- ◆ [Section E.12.3, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 246](#)
- ◆ [Section E.12.4, “Installing Novell Cluster Services on OES 2 Linux,” on page 247](#)
- ◆ [Section E.12.5, “Managing Clusters,” on page 247](#)

- ◆ [Section E.12.6, “Security Considerations,” on page 247](#)
- ◆ [Section E.12.7, “Troubleshooting Novell Cluster Services,” on page 247](#)
- ◆ [Section E.12.8, “What’s New,” on page 247](#)

E.12.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.2, “Considerations for Working with Shared NSS Pools and Volumes in the Cluster,” on page 142	Before you rename a cluster-enabled pool, make sure to offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, then online the pool resource.
“Mirroring and Cluster-Enabling Shared NSS Pools and Volumes” on page 156	This section was updated to correct errors in the procedures.

E.12.2 Console Commands for Novell Cluster Services

Location	Change
“Business Continuity Clustering Commands” on page 212	This section is new.

E.12.3 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.1.10, “Using Novell iManager in Mixed-Node Clusters,” on page 66	This section is new. Use Novell iManager 2.7.2 or later for all cluster administration in the mixed-node cluster.
Section 6.1.11, “Using Novell Remote Manager Is Not Supported in Mixed-Node Clusters,” on page 67	Do not use Novell Remote Manager when managing mixed-node clusters. Novell Remote Manager is not supported for cluster management on OES 2 Linux.
Section 6.1.12, “Using ConsoleOne Is Not Supported for Mixed-Node Clusters,” on page 67	This section is new. Do not use ConsoleOne when managing mixed-node clusters. ConsoleOne is not supported for cluster management on OES 2 Linux.
Section 6.8, “Finalizing the Cluster Conversion,” on page 92	WARNING: After you finalize the cluster conversion, rollback to NetWare is not supported.

E.12.4 Installing Novell Cluster Services on OES 2 Linux

Location	Change
“Cluster Services Installation Administrator” on page 36	Added information about how to set the rights that the user needs for installing Novell Cluster Services.
“Web Browser” on page 38	Added information about character encoding settings using Unicode (UTF-8).
“Using Device Mapper Multipath with Novell Cluster Services” on page 40	This section is new.

E.12.5 Managing Clusters

Location	Change
Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,” on page 114	Corrected a typographical error. Step 5 on page 115 should read “remove its NCS attributes”, not NCP attributes.

E.12.6 Security Considerations

This section is new.

E.12.7 Troubleshooting Novell Cluster Services

Location	Change
Section 13.4, “Diagnosing Cluster Problems,” on page 200	This section is new.
Section 13.11, “Could Not Delete This Resource Data_Server (Error 499),” on page 203	This section is new.

E.12.8 What’s New

Location	Change
Section 2.5.4, “Attribute NCS: GIPC Config Is No Longer Maintained,” on page 27	This section is new.

E.13 June 22, 2009

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.13.1, “Converting NetWare 6.5 Clusters to OES 2 Linux,”](#) on page 248
- ♦ [Section E.13.2, “Managing Clusters,”](#) on page 248

E.13.1 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.1.11, “Using Novell Remote Manager Is Not Supported in Mixed-Node Clusters,” on page 67	This section is new.

E.13.2 Managing Clusters

Location	Change
Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,” on page 114	Step 5 on page 115 is new.

E.14 June 5, 2009

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.14.1, “Configuring Cluster Policies and Priorities,”](#) on page 248
- ♦ [Section E.14.2, “Configuring Cluster Resources for Shared NSS Pools and Volumes,”](#) on page 249
- ♦ [Section E.14.3, “Console Commands for Novell Cluster Services,”](#) on page 249
- ♦ [Section E.14.4, “Converting NetWare 6.5 Clusters to OES 2 Linux,”](#) on page 249
- ♦ [Section E.14.5, “Installing Novell Cluster Services on OES 2 Linux,”](#) on page 250
- ♦ [Section E.14.6, “Managing Clusters,”](#) on page 250

E.14.1 Configuring Cluster Policies and Priorities

Location	Change
Section 7.3.2, “Tolerance,” on page 98	Added examples for the Tolerance setting.

E.14.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 150	Comment out (or remove) the volume’s entry in the <code>/etc/fstab</code> file. The load and unload scripts that are created when you cluster-enable the pool will be responsible for mounting and dismounting the volume after the pool is cluster enabled.
Section 10.7, “Configuring a Load Script for the Shared NSS Pool,” on page 155	Added information about using the <code>/opt=ns=<long unix mac dos></code> switch to specify the name space.

E.14.3 Console Commands for Novell Cluster Services

Location	Change
Section A.1, “Cluster Management Commands,” on page 209	The <code>cluster scan for new devices</code> command is obsolete on Linux. Use Linux commands to scan for new devices.

E.14.4 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.2.3, “Estimated Time Taken to Build the Trustee File on Linux,” on page 68	Added information about how long it might take to build the <code>trustee.xml</code> file during a migration from NetWare to Linux.
Section 6.4, “Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion),” on page 82	In Step 10 on page 84 , added clarification that checking for an SBD is necessary only if the cluster has a shared disk system.
Section 6.5, “Adding New OES 2 Linux Nodes to Your NetWare Cluster,” on page 85	In Step 5 on page 85 , added clarification that checking for an SBD is necessary only if the cluster has a shared disk system.

E.14.5 Installing Novell Cluster Services on OES 2 Linux

Location	Change
Section 3.1.4, “Shared Disk System Requirements,” on page 38	The following statement was added for clarification: <hr/> IMPORTANT: The cluster SBD partition is not required unless you have shared storage. <hr/>
Section 3.5.1, “Installing Novell Cluster Services during a OES 2 Linux Installation,” on page 45	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.
Section 3.5.2, “Installing Novell Cluster Services on an Existing OES 2 Linux Server,” on page 47	In Step 1 on page 47 , added clarification that creating an SBD is necessary only if the cluster has a shared disk system.
Section 3.6.1, “Opening the Novell Open Enterprise Server Configuration Page,” on page 49	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.

E.14.6 Managing Clusters

Location	Change
Section 8.16.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 116	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system. Added an example for the sbdutil utility.
Section 8.16.4, “Creating a Mirrored Cluster SBD Partition,” on page 117	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.

E.15 May 6, 2009

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.15.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 251](#)
- ♦ [Section E.15.2, “Configuring and Managing Cluster Resources,” on page 251](#)
- ♦ [Section E.15.3, “Converting NetWare 6.5 Clusters to OES 2 SP1 Linux,” on page 251](#)
- ♦ [Section E.15.4, “Installing Novell Cluster Services on OES 2 Linux,” on page 252](#)
- ♦ [Section E.15.5, “Troubleshooting Novell Cluster Services,” on page 252](#)
- ♦ [Section E.15.6, “Upgrading OES 2 Linux Servers,” on page 252](#)

E.15.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.6, “Adding Advertising Protocols,” on page 153	Step 9a was modified to clarify that you can specify a different name for the Novell CIFS Server only when you create the share. Use the CIFS iManager plug-in to manage the name of existing CIFS shares.

E.15.2 Configuring and Managing Cluster Resources

Location	Change
Section 9.3.2, “Creating a Resource Template,” on page 125	The following sentence in Step 5 was modified for clarification: If the script does not complete within the specified time, the resource becomes comatose when migrating to another node.
Section 9.6, “Configuring an Unload Script for a Cluster Resource,” on page 128	The following sentence in Step 6 was modified for clarification: If the script does not complete within the specified time, the resource becomes comatose when migrating to another node.
Section 9.7.3, “Monitoring Services Critical to Clustering,” on page 131	This section is new.
Section 9.8, “Setting Start, Failover, and Failback Modes for Cluster Resources,” on page 131	The following statement was added for clarification. There is no behavior change. 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. Client applications using server based storage must be restarted even with NCP unless the applications are NCP reconnect aware.

E.15.3 Converting NetWare 6.5 Clusters to OES 2 SP1 Linux

Location	Change
Section 6.1.9, “Managing File Systems in Mixed-Node Clusters,” on page 66	If you need to configure (or reconfigure) existing shared NSS pools and volumes or trustee assignments in a mixed-node cluster, you must temporarily remove all Linux cluster nodes by issuing the <code>cluster leave</code> command prior to making any NSS changes.

E.15.4 Installing Novell Cluster Services on OES 2 Linux

Location	Change
Section 3.3, “Extending the eDirectory Schema to Add Cluster Objects,” on page 42	Removed references to <code>ncpsrvr.ldif</code> . This file is a by-product of the extension and is not part of the setup.

E.15.5 Troubleshooting Novell Cluster Services

Location	Change
Section 13.15, “Can’t Find the Prevent Cascading Failover Option,” on page 204	This section is new.

E.15.6 Upgrading OES 2 Linux Servers

Location	Change
Section 4.2, “Upgrading OES 2 Clusters (Rolling Cluster Upgrade),” on page 57	<p>You can also add OES 2 SP1 Linux servers (or later versions) to the OES 2 Linux cluster, and remove the old OES 2 Linux servers from the cluster.</p> <p>Make sure that any services that are available only in OES 2 SP1 (such as Novell CIFS or Novell AFP) are set up with preferred nodes for failover that are running OES 2 SP1.</p> <p>Several guidelines are provided about the mixed-version cluster.</p>

E.16 March 3, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.16.1, “Configuring and Managing Cluster Resources,” on page 253](#)
- ◆ [Section E.16.2, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 253](#)
- ◆ [Section E.16.3, “Converting NetWare 6.5 Clusters to OES 2 Linux,” on page 253](#)
- ◆ [Section E.16.4, “Managing Clusters,” on page 253](#)
- ◆ [Section E.16.5, “What’s Next,” on page 253](#)

E.16.1 Configuring and Managing Cluster Resources

Location	Change
Section 9.12, "Deleting Cluster Resources," on page 134	The procedure was modified for clarification and accuracy.

E.16.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.5, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 150	Before enabling clustering for a pool and its volumes, make sure they have Storage objects in the eDirectory tree.

E.16.3 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Step 3 in Section 6.8, "Finalizing the Cluster Conversion," on page 92	You must restart the cluster so that clstrlib.ko is reloaded: <code>rcnovell-ncs restart</code>

E.16.4 Managing Clusters

Location	Change
Step 6 in Section 8.15, "Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node," on page 114	If you are converting a cluster from NetWare to Linux, you must restart the cluster instead so that clstrlib.ko is reloaded: <code>rcnovell-ncs restart</code>

E.16.5 What's Next

Location	Change
Section 2.5.3, "Behavior Change for Adding a Node," on page 26	IMPORTANT: A Novell Cluster Services patch is available in the patch channel and on the Novell Downloads Web site (http://www.novell.com/downloads) that allows you to add a new node seamlessly again on a Linux server. For cluster conversions, a cluster restart is still necessary after all NetWare nodes have been removed from the cluster.

Location	Change
Section 2.5.3, "Behavior Change for Adding a Node," on page 26	If you are converting a cluster from NetWare to Linux, you must restart the cluster instead so that clstrlib.ko is reloaded: <code>rcnovell-ncs restart</code>

E.17 February 13, 2009

Errata were corrected. Updates were made to the following section. The changes are explained below.

- ◆ [Section E.17.1, "Converting NetWare 6.5 Clusters to OES 2 Linux," on page 254](#)

E.17.1 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.2.2, "NSS File System Migration to NCP Volumes or Linux POSIX File Systems," on page 68	This section is new.
Table 6-1, "Guidelines for Converting Service Cluster Resources from NetWare to Linux," on page 69	Added a link to the eDirectory caveat.
Section 6.6, "Translation of Cluster Resource Scripts for Mixed NetWare and Linux Clusters," on page 86	The path to the files is <code>/var/opt/novell/ncs/</code> .

E.18 February 3, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.18.1, "Configuring and Managing Cluster Resources," on page 255](#)
- ◆ [Section E.18.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 255](#)
- ◆ [Section E.18.3, "Converting NetWare 6.5 Clusters to OES 2 Linux," on page 255](#)
- ◆ [Section E.18.4, "Installing Novell Cluster Services on OES 2 Linux," on page 255](#)

E.18.1 Configuring and Managing Cluster Resources

Location	Change
Step 7 in Section 9.7, "Enabling Monitoring and Configuring the Monitor Script," on page 129	Clarified the definition of <i>Maximum Local Failures</i> .

E.18.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Step 9b in Section 10.6, "Adding Advertising Protocols," on page 153	You can download the <code>cifsPool.py</code> script file as a patch Novell Downloads Web site (http://download.novell.com/Download?buildid=sA5Tdb_U2CE~) .

E.18.3 Converting NetWare 6.5 Clusters to OES 2 Linux

Location	Change
Section 6.3.3, "Apple Filing Protocol (AFP)," on page 71	Added information about using the NCP cross-protocol file locking feature on OES 2 SP1 Linux.

E.18.4 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Novell iManager 2.7.3" on page 31	NOTE: A February 3, 2009 update to the Clusters plug-in for OES 2 SP1 Linux is available that can be used to manage the Novell Business Continuity Clustering (BCC) 1.1 SP2 for NetWare 6.5 SP8 and Novell Business Continuity Clustering 1.2 for OES 2 SP1 Linux. You can download the update from the Novell Downloads Web site (http://download.novell.com/) .

E.19 January 13, 2009

Updates were made to correct errata.

E.20 December 2008 (OES 2 SP1)

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.20.1, “Comparison of Clustering OES 2 Services for Linux and NetWare,” on page 256](#)
- ◆ [Section E.20.2, “Comparison of Novell Cluster Services for Linux and NetWare,” on page 256](#)
- ◆ [Section E.20.3, “Configuring and Managing Cluster Resources,” on page 256](#)
- ◆ [Section E.20.4, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 257](#)
- ◆ [Section E.20.5, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 257](#)
- ◆ [Section E.20.6, “Configuring Novell Cluster Services in a Virtualization Environment,” on page 258](#)
- ◆ [Section E.20.7, “Console Commands for Novell Cluster Services,” on page 258](#)
- ◆ [Section E.20.8, “Converting NetWare 6.5 Cluster to OES 2 Linux,” on page 258](#)
- ◆ [Section E.20.9, “Installing Novell Cluster Services on OES 2 Linux,” on page 259](#)
- ◆ [Section E.20.10, “Managing Clusters,” on page 259](#)
- ◆ [Section E.20.11, “Overview of Novell Cluster Services,” on page 259](#)
- ◆ [Section E.20.12, “Troubleshooting Novell Cluster Services,” on page 260](#)
- ◆ [Section E.20.13, “Upgrading OES 2 Clusters,” on page 260](#)
- ◆ [Section E.20.14, “What’s New,” on page 260](#)

E.20.1 Comparison of Clustering OES 2 Services for Linux and NetWare

This section is new.

E.20.2 Comparison of Novell Cluster Services for Linux and NetWare

This section is new.

E.20.3 Configuring and Managing Cluster Resources

Location	Change
Section 9.3, “Creating Cluster Resource Templates,” on page 124	This section was modified for clarity.
Section 9.12, “Deleting Cluster Resources,” on page 134	Added procedures for deleting cluster resources on the master node and the non-master node.

E.20.4 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 11.8, “Known Issues for Working with Cluster Resources for Linux POSIX Volumes,” on page 183	This section is new.

E.20.5 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.3, “Creating NSS Shared Disk Partitions and Pools,” on page 143	Added information about using Novell CIFS for Linux and Novell AFP for Linux when cluster-enabling pools on OES 2 SP1 Linux and later.
Section 10.4, “Creating NSS Volumes on a Shared Pool,” on page 148	IMPORTANT: When adding volumes on a shared pool, make sure that the pool is mounted on the server where the pool was created. This is necessary because the Pool object information is relative to the Server object of the server where the pool was created. Added information about using encrypted NSS volumes on a shared pool.
Section 10.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 150	Added information about using Novell CIFS for Linux and Novell AFP for Linux when cluster-enabling pools on OES 2 SP1 Linux and later.
Section 10.6, “Adding Advertising Protocols,” on page 153	This section is new.
“Mirroring the Partition of an Existing NSS Pool” on page 159	This section is new.
Section 10.12, “Changing the Volume ID,” on page 161	This section is new.

E.20.6 Configuring Novell Cluster Services in a Virtualization Environment

Location	Change
Chapter 12, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 185	IMPORTANT: All templates except Xen and XenLive are valid in guest servers (DomU) in the virtualization environment. Only the Xen and XenLive templates can be used in the OES 2 Linux (Xen) host environment (that is, in Dom0, but not in DomU).
Section 12.2, “Virtual Machines as Cluster Resources,” on page 186	An overview of OCFS2 is available in “ Oracle Cluster File System 2 ” (http://www.novell.com/documentation/sles10/sles_admin/data/cha_ocfs2.html) in the <i>SUSE Linux Enterprise Server 10 SP2 Administration Guide</i> . For detailed information about using OCFS2, see the <i>OCFS2 Project</i> (http://oss.oracle.com/projects/ocfs2/) on the Oracle Web site.

E.20.7 Console Commands for Novell Cluster Services

Location	Change
Section A.3, “extend_schema Command,” on page 212	This section is new.

E.20.8 Converting NetWare 6.5 Cluster to OES 2 Linux

Location	Change
Section 6.1.13, “Using the Monitor Function in Mixed-Node Clusters Is Not Supported,” on page 67	This section is new.
Section 6.3, “Guidelines for Converting Service Cluster Resources from NetWare to Linux,” on page 68	This section is new.
Section 6.6.4, “Comparing File Access Protocol Resource Script Commands,” on page 89	This section is new.
Section 6.4, “Converting NetWare Cluster Nodes to OES 2 Linux (Rolling Cluster Conversion),” on page 82	This procedure has been updated.

E.20.9 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Novell eDirectory 8.8" on page 31	If the eDirectory administrator username or password contains special characters (such as \$, #, and so on), make sure to escape each special character by preceding it with a backslash (\) when you enter credentials.
"NCP Server for Linux" on page 33	This section is new.
"Novell AFP for Linux" on page 34	This section is new.
"Novell CIFS for Linux" on page 34	This section is new.
"SLP" on page 35	This section is new.
"Novell Domain Services for Windows" on page 34	This section is new.
Section 3.1.3, "Configuration Requirements," on page 36	Added information about the credentials needed by container administrators for installing Novell Cluster Services in trees where the schema has already been extended.
Section 3.3, "Extending the eDirectory Schema to Add Cluster Objects," on page 42	This section is new. It explains a new feature for extending the schema in a Novell eDirectory tree before clusters are installed.
Section 3.5.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 45	This section has been revised for OES 2 SP1.

E.20.10 Managing Clusters

Location	Change
Section 8.13, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 110	The procedure can also be used to modify the administrator user, or to modify the password of the existing administrator user. The procedure has been updated for clarity.

E.20.11 Overview of Novell Cluster Services

Location	Change
Section 1.1, "Why Should I Use Clusters?," on page 15	This section is new.
Section 1.4, "Clustering for High-Availability," on page 16	This scenario was relocated from Section 1.2, "Benefits of Novell Cluster Services," on page 15 .

Location	Change
Section 1.5, “Shared Disk Scenarios,” on page 18	This section was reorganized for clarity.

E.20.12 Troubleshooting Novell Cluster Services

Location	Change
Section 13.7, “Cluster Resource Goes Comatose Immediately After Migration or Failover,” on page 202	This section is new.
Section 13.10, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux),” on page 203	This section is new.

E.20.13 Upgrading OES 2 Clusters

This section is new.

E.20.14 What’s New

Location	Change
Section 2.5, “What’s New (OES 2 SP1),” on page 25	This section is new.

E.21 June 4, 2008

Updates were made to the following sections. The changes are explained below.

- ◆ [Section E.21.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 261](#)
- ◆ [Section E.21.2, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 261](#)
- ◆ [Section E.21.3, “Installation and Setup,” on page 262](#)
- ◆ [Section E.21.4, “Managing Novell Cluster Services,” on page 262](#)

E.21.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 10.1, “Planning for Shared NSS Pools and Volumes,” on page 139	This section is new.
Section 10.2, “Considerations for Working with Shared NSS Pools and Volumes in the Cluster,” on page 142	This section is new.
Section 10.3.1, “Initializing Shared Devices,” on page 143	This section is new.
Section 10.3.2, “Enabling Sharing on a Device,” on page 143	This section is new.
Section 10.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 150	Make sure you offline the cluster resource before you attempt to delete the clustered pool or its cluster resource.
Section 10.10, “Mirroring and Cluster-Enabling Shared NSS Pools and Volumes,” on page 156	This section is new.
Section 10.11, “Deleting NSS Pool Cluster Resources,” on page 161	This section is new.

E.21.2 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 11.1, “Requirements for Shared Linux POSIX Volumes,” on page 163	This section is new.
Section 11.2, “Creating Linux POSIX Volumes on Shared Disks,” on page 164	Procedures were corrected.
Section 11.5, “Sample Scripts for a Linux POSIX Volume Cluster Resource,” on page 179	Script commands were corrected.

Location	Change
Section 11.7, “Deleting Shared Storage,” on page 183	This section is new.

E.21.3 Installation and Setup

The following topics were moved to separate sections:

- ♦ [Chapter 7, “Configuring Cluster Policies and Priorities,” on page 95](#)
- ♦ [Chapter 9, “Configuring and Managing Cluster Resources,” on page 123](#)
- ♦ [Chapter 10, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 139](#)
- ♦ [Chapter 11, “Configuring Cluster Resources for Shared Linux POSIX Volumes,” on page 163](#)

The following additional changes were made:

Location	Change
Section 3.1, “System Requirements for Novell Cluster Services,” on page 29	This section is new.
Section 3.2, “Novell Cluster Services Licensing,” on page 42	This section was updated.
Section 5.1, “Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux,” on page 59	This section is new.
Section 6.1, “Guidelines for Converting Clusters from NetWare to OES 2 Linux,” on page 63	This section is new.

E.21.4 Managing Novell Cluster Services

The following topics were moved to separate sections:

- ♦ [Chapter 9, “Configuring and Managing Cluster Resources,” on page 123](#)
- ♦ [Chapter 13, “Troubleshooting Novell Cluster Services,” on page 199](#)
- ♦ [Appendix A, “Console Commands for Novell Cluster Services,” on page 209](#)
- ♦ [Appendix B, “Files for Novell Cluster Services,” on page 215](#)

The following additional changes were made:

Location	Change
Chapter 8, “Managing Clusters,” on page 101	This section was edited for clarity.
Section 8.2, “Monitoring Cluster and Resource States,” on page 103	This section was edited for clarity.
Section 8.5, “Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node,” on page 106	This section is new.
Section 8.6, “Removing (Leaving) a Node from the Cluster,” on page 107	This section is new.
Section 8.7, “Joining a Node to the Cluster,” on page 107	This section is new.
Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,” on page 114	The procedure was edited for clarity.

E.22 May 2, 2008

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.22.1, “Installation and Setup,” on page 263](#)
- ♦ [Section E.22.2, “Managing Novell Cluster Services,” on page 264](#)

E.22.1 Installation and Setup

Location	Change
Section 3.5.1, “Installing Novell Cluster Services during a OES 2 Linux Installation,” on page 45	A one-node cluster can be configured without an SBD (split-brain detector). However, for adding a second node in the existing cluster (without SBD), you need additional pre-configuration of creating the SBD partition on the shared disk by using <code>sbdutil</code> on the first node.
Section 3.5.2, “Installing Novell Cluster Services on an Existing OES 2 Linux Server,” on page 47	Corrected the procedure to use the YaST Control Center. You cannot install Novell Cluster Services by going directly to the <code>yast2 ncs</code> option at the command line. A one-node cluster can be configured without an SBD (split-brain detector). However, for adding a second node in the existing cluster (without SBD), you need additional pre-configuration of creating the SBD partition on the shared disk by using <code>sbdutil</code> on the first node.

Location	Change
“Configuring an Unload Script for a Linux POSIX Volume Cluster Resource” on page 171	In OES 2, use deactivate_evms_container instead of deport_evms_container.

E.22.2 Managing Novell Cluster Services

Location	Change
Section 8.2, “Monitoring Cluster and Resource States,” on page 103	When the resource is in Start Alert state, you must clear the alert before you can offline the resource. After the resource is offline, you can online the resource.
Section 8.15, “Deleting a Cluster Node from a Cluster, or Reconfiguring a Cluster Node,” on page 114	Details have been added for clarification.
Section 8.13, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,” on page 110	This section is new.